

SmartStack with Cisco UCS and Nimble CS700, with Citrix XenDesktop 7.7 VDI 2500 Seat Deployment with Graphics Support

Last Updated: June 20, 2016



About Cisco Validated Designs (CVD)

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, 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, CCIP, CCNA, CCNP, CCSP, 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. (0809R)

© 2016 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	7
Solution Overview	8
Introduction	8
Audience	8
SmartStack Program Benefits	8
Technology Overview	9
Cisco Unified Computing System.....	9
Cisco UCS Differentiators	11
Cisco UCS 5108 Blade Server Chassis	12
Cisco UCS 6200 Series Fabric Interconnects.....	13
Cisco UCS Fabric Extenders	14
Cisco UCS Manager	14
Cisco UCS B-Series M4 Servers.....	15
Cisco Nexus 9000 Series Platform Switches.....	16
Cisco Nexus 1000v.....	17
Cisco MDS 9100 Series Fabric Switches	18
Nimble Storage – CS700 Hybrid Array	20
Nimble Storage Predictive Flash Platform	20
Citrix XenApp and XenDesktop 7.7	23
Citrix Provisioning Services 7.7	25
Citrix Desktop Studio for XenApp 7.7	26
Benefits for Desktop Administrators.....	26
Citrix Provisioning Services Solution	27
Citrix Provisioning Services Infrastructure	27
Solution Architecture	29
Solution Design.....	35
Compute.....	35
Network.....	36
Storage.....	37
Design Considerations.....	39
Management Connectivity	39
QoS and Jumbo Frames	39
Cisco UCS Server – vSphere Configuration.....	39
Cisco UCS Server – Virtual Switching using Nexus 1000V	40
Cisco Nexus 9000 Series – vPC Best Practices.....	42
High Availability	45
Scalability	47

Solution Validation	49
Configuration Topology for Scalable XenDesktop Mixed Workload Desktop Virtualization Solution on Cisco Unified Computing System and Nimble Storage	49
Cisco Unified Computing System Configuration	49
Base Cisco UCS System Configuration	49
Enable Server Uplink Ports.....	51
Create Resource Pools	54
Create VLANs	62
Create VSANs.....	63
Create Host Firmware Package	65
Set Jumbo Frames in Cisco UCS Fabric.....	66
Create Network Control Policy for Cisco Discovery Protocol.....	66
Create Power Control Policy	67
Cisco UCS System Configuration for Cisco UCS B-Series	68
Create Server BIOS Policy	68
Configure Update Default Maintenance Policy	71
Create vNIC Templates for Cisco UCS B-Series	72
Create vHBA Templates for Cisco UCS B-Series	74
Create Service Profile Templates for Cisco UCS B-Series	75
Nimble CS700 Configuration.....	86
Nimble CS700 Adaptive Array System Configuration.....	86
Nimble Management Tools: InfoSight.....	89
Register and Login to InfoSight.....	89
Configure Arrays to Send Data to InfoSight.....	89
Nimble Management Tools: vCenter Plugin.....	90
Register vCenter Plugins.....	90
Configure Arrays to Monitor your Virtual Environment	91
Configure Setup Email Notifications for Alerts.....	92
Data Storage Layout	92
Create Initiator Groups.....	93
Create Volumes	94
Configure MDS 9100 Series	100
Boot from SAN Benefits.....	105
SAN Configuration on the Cisco MDS 9148 Switches	105
Install and Configure ESXi 6 U1a	107
VMware ESXi 6.0	107
Download Cisco Custom Image for ESXi 6 Update1	107
Install ESXi.....	107
Set Up Management Networking for ESXi Hosts	108
Download VMware vSphere Client.....	110

Download VMware vSphere CLI 6	110
Log in to VMware ESXi Hosts by Using VMware vSphere Client.....	111
ESXi Host VM-Host-01	111
Install and Configure vCenter 6.....	111
Install the Nimble Connection Manager.....	123
Building the Virtual Machines and Environment	124
Install and Configure VSUM and Nexus 1000v	126
Install Cisco Virtual Switch Update Manager	126
Install Cisco Virtual Switch Update Manager	126
About the Cisco VSUM GUI.....	130
Install Cisco Nexus 1000V using Cisco VSUM.....	131
Perform Base Configuration of the Primary VSM	133
Add VMware ESXi Hosts to Cisco Nexus 1000V	136
Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V	137
Cisco Nexus 1000V vTracker.....	139
Nexus 1000V Configuration	140
Installing and Configuring Citrix License Server	152
Installing and Configuring Citrix Provisioning Server 7.7	156
Preparing the Master Targets	180
Create the Provisioning Services Master Image	181
Creating PVS vDisks	181
Installing and Configuring Citrix XenDesktop/XenApp 7.7	187
Configure the XenDesktop Site Hosts and Storage	197
Configure XenDesktop HDX Policies	201
Installing and Configuring Citrix Storefront.....	206
Test Setup and Configurations.....	212
Cisco UCS Test Configuration for Single Blade Scalability	212
Cisco UCS Test Configuration for Blade Server VDI Scalability	213
Testing Methodology and Success Criteria	216
Testing Procedure	216
Success Criteria	217
VSImax 4.1.x Description.....	218
Server-Side Response Time Measurements	218
Calculating VSImax v4.1.x	219
Test Results.....	222
Single-Server Recommended Maximum Workload For B-Series	222
Single Server on B200-M4 with 195 Users Test Results	223
VDI Scaling with 2500 Users Test Results.....	226
Solution Design Considerations	227

2500 XenDesktop Users VDI Scale Testing with LoginVSI with and without PVS RAM Cache	229
2500 XenDesktop VM's+ 16 XenApp Servers BootStorm from vCenter	230
Single Server Testing Utilizing the NVIDIA M6 Card and vGPU	235
Cisco UCS B200 M4 Blade Server with NVIDIA M6 GRID Card	235
Install and Configure NVIDIA M6 Card.....	236
Physical Installation of the NVIDIA M6 Card into the Cisco UCS B200 M4 Server	236
Before You Begin.....	237
Install the NVIDIA VMware VIB Driver	238
Configure a VM with a vGPU	241
Install the GPU Drivers into Windows VM	243
Install and configure NVIDIA Grid License Server	244
Testing Methodology and Results for the NVIDIA M6 Cards	248
Internet Explorer 11 Configuration	249
Test Configurations	250
GPU Performance Metrics.....	250
Validated Hardware and Software.....	252
Bill of Materials (BOM)	254
Summary	256
About Authors	257
Acknowledgements	257
Appendix A – Cisco Nexus 9372 Switch Configuration	258
Switch A Configuration	258
Switch B Configuration	270
Appendix B – Cisco MDS 9148 Switch Configuration	282
MDS- A Switch Configuration	282
MDS- B Switch Configuration	296

Executive Summary

Cisco Validated Designs (CVD) are systems and solutions that have been designed, tested and documented to facilitate and accelerate customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of a customer. CVDs deliver a validated design, documentation and support information to guide customers from design to deployment.

Cisco, Nimble and Citrix **have partnered to deliver a series of SmartStack™ solutions that combine** Cisco Unified Computing System servers, Cisco Nexus family of switches, and Nimble Storage arrays into an enterprise VDI solution with graphics support.

Customers looking to implement a VDI solution using shared data center infrastructure face a number of challenges. A constant challenge is achieving the levels of IT agility and efficiency necessary to meet business objectives. Addressing these challenges requires having an optimal solution with the following characteristics:

- **Availability:** Helps ensure applications and services are accessible at all times with no single point of failure
- **Flexibility:** Ability to support new services without requiring infrastructure modifications.
- **Efficiency:** Facilitate efficient operation of the infrastructure through re-usable policies and API management.
- **Ease of deployment and management** to minimize operating costs.
- **Scalability:** Ability to expand and grow with some degree of investment protection
- **Minimal risk** by ensuring optimal design and compatibility of integrated components

SmartStack enables a data center platform with the above characteristics by delivering an integrated architecture that incorporates compute, storage and network design best practices. SmartStack minimizes IT risk by testing the integrated architecture to ensure compatibility between integrated components. SmartStack also addresses IT pain points by providing documented design guidance, deployment guidance and support that can be used in all stages (planning, design and implementation) of a deployment.

The SmartStack solution outlined in this document delivers a converged infrastructure platform designed for Enterprise and Cloud datacenters. SmartStack incorporates compute, network and storage best practices to deliver a resilient, scalable and flexible datacenter architecture. The design uses Cisco UCS servers for compute, VMware vSphere 6.0U1 hypervisor, Citrix XenDesktop 7.7, Citrix Provisioning Server 7.7, NVidia M6 Graphics Cards, Intel Iris Processors for graphics support, Cisco Nexus and MDS switches for network and Fibre Channel-attached Nimble CS700 hybrid array for storage.

Solution Overview

Introduction

This document outlines the deployment procedures for implementing a VDI and SmartStack infrastructure solution using Citrix VDI technologies. This guide is based on the SmartStack solution validation that was done using VMware vSphere 6.0 U1a, Cisco UCS B-Series, Nimble CS700 Adaptive (or Hybrid) Array, Citrix VDI Xen technologies with graphics support, and Cisco Nexus switches.

SmartStack is a pre-designed, validated integrated infrastructure architecture for the data center. SmartStack solution portfolio combines Nimble® Storage arrays, Cisco® UCS servers, Cisco MDS fabric switches and Cisco Nexus switches into a single, flexible architecture. SmartStack solutions are designed and validated to minimize deployment time, project risk, and overall IT costs.

SmartStack is designed for high availability, with no single points of failure while maintaining cost-effectiveness and flexibility in design to support a variety of workloads in Enterprise and cloud datacenters. SmartStack design can support different hypervisor options, and also be sized and optimized to support different use cases and requirements.

This document describes the SmartStack® integrated VDI solution for Enterprise and cloud deployments built using Cisco UCS, VMware vSphere 6.0U1, Cisco Nexus and MDS switches, Fibre Channel-attached Nimble CS700 array, Citrix XenDesktop 7.7, Citrix Provisioning Server 7.7, NVidia M6 Graphics Cards for graphics support.

Audience

The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

SmartStack Program Benefits

The SmartStack program is the result of a joint partnership between Cisco and Nimble Storage to deliver a series of infrastructure and application solutions optimized and validated on Cisco UCS, Nimble Storage and Cisco Nexus switches. Customers must use Cisco UCS, Nimble Storage and one of the approved application stacks to be a valid SmartStack solution and they must also have valid support contracts with Cisco and Nimble Storage.

Cisco and Nimble Storage have a solid, joint support program focused on SmartStack solution, from customer account and technical sales representatives to professional services and technical support engineers. The support alliance provided by Cisco and Nimble Storage provides customers and channel partners with direct access to technical expert who can collaborate with cross vendors and have access to shared lab resources to resolve potential issues.

Technology Overview

SmartStack is a data center architecture for Enterprise or Cloud deployments and uses the following infrastructure components for compute, network and storage:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus and MDS Switches
- Nimble Storage arrays

The validated SmartStack design covered in this document uses the following models of the above infrastructure components.

- Cisco UCS 5100 Series Blade Server Chassis with 2200 Series Fabric Extenders (FEX)
- Cisco UCS B-Series Blade Servers
- Cisco UCS 6200 Series Fabric Interconnects (FI)
- Cisco Nexus 9300 Series Platform switches
- Cisco MDS 9100 Series Fabric switches
- Nimble CS700 Hybrid Array
- NVidia M6 Graphics Cards for Blade Servers

The above components are integrated using design and component best practices to deliver a converged infrastructure for Enterprise and cloud data centers.

The next section provides a technical overview of the compute, network, storage and management components of the SmartStack solution.

Cisco Unified Computing System

The Cisco Unified Computing System™ (**Cisco UCS**) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources 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 where all resources are managed through a unified management domain.

The main components of the Cisco UCS are:

Compute - The system is based on an entirely new class of computing system that incorporates rack mount server, blade servers and modular servers based on Intel processors.

Network - The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.

Virtualization - The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.

Storage access - Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. Also, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity. SmartStack solutions can support either iSCSI or Fibre Channel based access. This design covers only Fibre Channel connectivity.

Management: The system uniquely integrates all the system components, enabling the entire solution to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager provides an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application-programming interface (API) to manage all system configuration and operations. Cisco UCS Manager helps increase IT staff productivity, enabling storage, network, and server administrators to collaborate on defining service profiles for applications. Service profiles are logical representations of desired physical configurations and infrastructure policies. They help automate provisioning and increase business agility, allowing data center managers to provision resources in minutes instead of days.

The Cisco Unified Computing System in the SmartStack architecture consists of the following components:

- [Cisco UCS Manager](#) provides unified management of all software and hardware components in the Cisco Unified Computing System and manages servers, networking, and storage from a single interface.
- [Cisco UCS 6200 Series Fabric Interconnects](#) is a family of line-rate, low-latency, lossless, 10-Gbps Ethernet and Fibre Channel over Ethernet interconnect switches providing the management and communication backbone for the Cisco Unified Computing System.
- [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](#) increase performance, efficiency, versatility and productivity with these Intel based blade servers.
- [Cisco UCS Adapters](#) wire-once architecture offers a range of options to converge the fabric, optimize virtualization and simplify management.
- [Cisco Nexus 1000V Series Switches](#) are virtual machine access switches for VMware vSphere environments that provide full switching capabilities and Layer 4 through Layer 7 services to virtual machines.
- [Cisco UCS Blade Server M6 GPU - GRID 2.0 SW Required for VDI](#) you can now expand your virtualization footprint without compromising performance or user experience while also increasing security. This means, you can empower your workforce to create anything around the world, from any location with the ease and flexibility.

The optional Cisco UCS components of the SmartStack solution are:

- [Cisco UCS Central](#) provides a scalable management platform for managing multiple, globally distributed Cisco UCS domains with consistency by integrating with Cisco UCS Manager to provide global configuration capabilities for pools, policies, and firmware.
- [Cisco UCS Performance Manager](#) is purpose-built data center management solution that provides a single pane-of-glass visibility of a converged heterogeneous infrastructure datacenter for performance monitoring and capacity planning.

Cisco Unified Computing System has revolutionizing the way servers are managed in data-center. The next section takes a detailed look at the unique differentiators of Cisco UCS and Cisco UCS Manager®.

Cisco UCS Differentiators

- Embedded Management – Servers in the system are managed by embedded software in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers.
- Unified Fabric – There is a single Ethernet cable to the FI from the server chassis (blade or modular or rack) for LAN, SAN and management traffic. This converged I/O results in reduced cables, SFPs and adapters – reducing capital and operational expenses of overall solution.
- Auto Discovery – By simply inserting a blade server in the chassis or connecting a rack server to the FI, discovery and inventory of compute resource occurs automatically without any intervention. Auto-discovery combined with unified fabric enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily without additional connections to the external LAN, SAN and management networks.
- Policy Based Resource Classification – When a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD showcases the policy based resource classification of Cisco UCS Manager.
- Combined Rack, Blade and Modular Server Management – Cisco UCS Manager can manage B-series blade servers, C-series rack servers and M-series modular servers under the same Cisco UCS domain. Along with stateless computing, this feature makes compute resources truly agnostic to the hardware form factor.
- Model based Management Architecture – Cisco UCS Manager architecture and management database is model based and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.
- Policies, Pools, Templates – The management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- Loose Referential Integrity – In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts from different domains, such as network, storage, security, server and virtualization the flexibility to work independently to accomplish a complex task.

- **Policy Resolution** – In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to another policy by name is resolved in the organization hierarchy with closest policy match. If no policy with **specific name is found in the hierarchy of the root organization, then special policy named “default”** is searched. This policy resolution logic enables automation friendly management APIs and provides great flexibility to owners of different organizations.
- **Service Profiles and Stateless Computing** – A service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- **Built-in Multi-Tenancy Support** – The combination of policies, pools and templates, loose referential integrity, policy resolution in organization hierarchy and a service profiles based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.
- **Extended Memory** – The enterprise-class Cisco UCS B200 M4 blade server extends the capabilities **of Cisco’s Unified Computing System portfolio in a half-width blade form factor**. The Cisco UCS B200 M4 harnesses the power of the latest Intel® Xeon® E5-2600 v3 Series processor family CPUs with up to 1536 GB of RAM (using 64 GB DIMMs) – allowing huge VM to physical server ratio required in many deployments, or allowing large memory operations required by certain architectures like big data.
- **Virtualization Aware Network** – Cisco VM-FEX technology makes the access network layer aware about host virtualization. This prevents domain pollution of compute and network domains with virtualization when virtual network is managed by port-profiles defined by the network **administrators’ team**. VM-FEX also off-loads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.
- **Simplified QoS** – Even though Fibre Channel and Ethernet are converged in Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a fundamental building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server architecture. A Cisco UCS 5108 Blade Server chassis is six rack units (6RU) high and can house up to eight half-width or four full-width Cisco UCS B-series blade servers.

For a complete list of blade servers supported, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>.


There are four hot-swappable power supplies that are accessible from the front of the chassis. These power supplies are 94 percent efficient and can be configured to support non-redundant, N+1, and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per

power supply), and two I/O bays that can support Cisco UCS 2000 Series Fabric Extenders. The two fabric extenders can be used for both redundancy and bandwidth aggregation. A passive mid-plane provides up to 80 Gbps of I/O bandwidth per server slot and up to 160 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 40 Gigabit Ethernet standards.

Cisco UCS 5108 blade server chassis uses a unified fabric and fabric-extender technology to simplify and reduce cabling by eliminating the need for dedicated chassis management and blade switches. The unified fabric also reduces TCO by reducing the number of network interface cards (NICs), host bus adapters (HBAs), switches, and cables that need to be managed, cooled, and powered. This architecture enables a single Cisco UCS domain to scale up to 20 chassis with minimal complexity.

For more information, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-5100-series-blade-server-chassis/index.html>.

Figure 1 Cisco UCS 5108 Blade Server Chassis



Cisco UCS 6200 Series Fabric Interconnects

Cisco UCS Fabric Interconnects are a family of line-rate, low-latency, lossless 1/10 Gigabit Ethernet and Fiber Channel over Ethernet (FCoE), and 4/2/1 and 8/4/2 native Fibre Channel switches. Cisco UCS Fabric Interconnects are the management and connectivity backbone of the Cisco Unified Computing System. Each chassis or rack server connects to the FI using a single Ethernet cable for carrying all network, storage and management traffic. Cisco UCS Fabric Interconnects provide uniform network and storage access to servers and are typically deployed in redundant pairs.

Cisco UCS Manager provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS) across multiple chassis with blade servers, rack servers and thousands of virtual machines. The Cisco UCS Management software (Cisco UCS Manager) runs as an embedded device manager in a clustered pair fabric interconnects and manages the resources connected to it. An instance of Cisco UCS Manager with all Cisco UCS components managed by it forms a Cisco UCS domain, which can include up to 160 servers.

The Cisco UCS Fabric Interconnect family is currently comprised of the Cisco 6100 Series, Cisco 6200 Series and Cisco 6300 Series of Fabric Interconnects.



Cisco UCS 6248UP Fabric Interconnects were used for this CVD.

For more information, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6200-series-fabric-interconnects/index.html>

Figure 2 Cisco UCS 6248UP Fabric Interconnect



Cisco UCS Fabric Extenders

The Cisco UCS Fabric extenders multiplexes and forwards all traffic from servers in a chassis to a parent Cisco UCS Fabric Interconnect over from 10-Gbps unified fabric links. All traffic, including traffic between servers on the same chassis, or between virtual machines on the same server, is forwarded to the parent fabric interconnect, where network profiles and policies are maintained and managed by the Cisco UCS Manager. The Fabric extender technology was developed by Cisco. Up to two fabric extenders can be deployed in a Cisco UCS chassis.

The Cisco UCS Fabric Extender family currently comprises of Cisco UCS 2200 and Cisco Nexus 2000 Series of Fabric Extenders. The Cisco UCS 2200 Series Fabric Extenders come in two flavors as outlined below.

- The Cisco UCS 2204XP Fabric Extender has four 10 Gigabit Ethernet, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2204XP has sixteen 10 Gigabit Ethernet ports connected through the mid-plane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 80 Gbps of I/O to the chassis.
- The Cisco UCS 2208XP Fabric Extender has eight 10 Gigabit Ethernet, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the mid-plane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 160 Gbps of I/O to the chassis.



Cisco UCS 2208 Fabric Extenders were used for this CVD.

For more information, see: http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6200-series-fabric-interconnects/data_sheet_c78-675243.html

Figure 3 Cisco UCS 2208 Series Fabric Extenders



Cisco UCS Manager

Cisco Unified Computing System (UCS) Manager provides unified, embedded management for all software and hardware components in the Cisco UCS. Using [Cisco Single Connect](#) technology, it manages, controls, and administers multiple chassis for thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive GUI, a

command-line interface (CLI), or an XML API. The Cisco UCS Manager resides on a pair of Cisco UCS 6200 Series Fabric Interconnects using a clustered, active-standby configuration for high availability.

Cisco UCS Manager offers unified embedded management interface that integrates server, network, and storage. Cisco UCS Manager performs auto-discovery to detect inventory, manage, and provision system components that are added or changed. It offers comprehensive set of XML API for third part integration, exposes 9000 points of integration and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

Service profiles benefit both virtualized and non-virtualized environments and increase the mobility of non-virtualized servers, such as when moving workloads from server to server or taking a server offline for service or upgrade. Profiles can also be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing virtual machine mobility.

For more information on Cisco UCS Manager, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

Cisco UCS B-Series M4 Servers



Cisco UCS B200 M4 blade servers with Cisco Virtual Interface Card 1340 were used for this CVD.

The enterprise-class **Cisco UCS B200 M4 Blade Server extends the capabilities of Cisco's Unified Computing System** portfolio in a half-width blade form factor. The Cisco UCS B200 M4 uses the power of the latest Intel® Xeon® E5-2600 v3 Series processor family CPUs with up to 768 GB of RAM (using 32 GB DIMMs), two solid-state drives (SSDs) or hard disk drives (HDDs), and up to 80 Gbps throughput connectivity. The Cisco UCS B200 M4 Blade Server mounts in a Cisco UCS 5100 Series blade server chassis or Cisco UCS Mini blade server chassis. It has 24 total slots for registered ECC DIMMs (RDIMMs) or load-reduced DIMMs (LR DIMMs) for up to 768 GB total memory capacity (Cisco UCS B200 M4 configured with **two CPUs using 32 GB DIMMs**). It supports one connector for Cisco's VIC 1340 or 1240 adapter, which provides Ethernet and FCoE. There is also a second mezzanine card slot that can be used for the NVidia M6 Graphics cards.

For more information, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b200-m4-blade-server/index.html>

Figure 4 Cisco UCS B200 M4 Blade Server



Cisco VIC 1340

The Cisco UCS Virtual Interface Card (VIC) 1340 is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities is enabled for two ports of 40-Gbps Ethernet. The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces

to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1340 supports Cisco® Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS Fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.

For more information, see: <http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1340/index.html>

Figure 5 Cisco VIC 1340



Cisco Nexus 9000 Series Platform Switches

The Cisco Nexus 9000 family of switches offer both modular (9500 switches) and fixed (9300 switches) 1/10/40/100 Gigabit Ethernet switch configurations designed to operate in one of two modes:

- Application Centric Infrastructure (ACI) mode that uses an application centric policy model with simplified automation and centralized management
- Cisco NX-OS mode for traditional architectures – the SmartStack design in this document uses this mode

Architectural Flexibility

- Delivers high performance and density, and energy-efficient traditional 3-tier or leaf-spine architectures
- Provides a foundation for Cisco ACI, automating application deployment and delivering simplicity, agility, and flexibility

Scalability

- Up to 60-Tbps of non-blocking performance with less than 5-microsecond latency
- Up to 2304 10-Gbps or 576 40-Gbps non-blocking layer 2 and layer 3 Ethernet ports
- Wire-speed virtual extensible LAN (VXLAN) gateway, bridging, and routing

High Availability

- Full Cisco In-Service Software Upgrade (ISSU) and patching without any interruption in operation
- Fully redundant and hot-swappable components
- A mix of third-party and Cisco ASICs provide for improved reliability and performance

Energy Efficiency

- The chassis is designed without a mid-plane to optimize airflow and reduce energy requirements
- The optimized design runs with fewer ASICs, resulting in lower energy use
- Efficient power supplies included in the switches are rated at 80 Plus Platinum

Investment Protection

- Cisco 40-Gb bidirectional transceiver allows for reuse of an existing 10 Gigabit Ethernet cabling plant for 40 Gigabit Ethernet
- Designed to support future ASIC generations
- Easy migration from NX-OS mode to ACI mode



A pair of Cisco Nexus 9372PX Platform switches were used in this CVD.

For more information, refer to: <http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>

Cisco Nexus 1000v

Cisco Nexus 1000V Series Switches provide a comprehensive and extensible architectural platform for virtual machine (VM) and cloud networking. Integrated into the VMware vSphere hypervisor and fully compatible with VMware vCloud® Director, the Cisco Nexus 1000V Series provides:

- Advanced virtual machine networking using Cisco NX-OS operating system. Capabilities include PortChannels (LACP), IEEE 802.1Q VLAN trunking, Jumbo Frame support and Virtual Extensible Local Area Network (VXLAN) for cloud deployments
- Cisco vPath technology for efficient and optimized integration of Layer 4-7 virtual networking services (e.g. Firewall)
- Mobile virtual machine security and network policy. Advanced security capabilities include Storm Control, BPDU Guard, Dynamic Host Configuration Protocol (DHCP) snooping, IP source guard, Dynamic Address Resolution Protocol (ARP) Inspection, and Cisco TrustSec® security group access (SGA), Security Group Tagging (SGT) and Security Group ACL (SGACL) support.
- Non-disruptive operational model for your server virtualization and networking teams
- Policy-based virtual machine connectivity
- Quality of service (QoS)
- Monitoring: NetFlow, Switch Port Analyzer (SPAN), and Encapsulated Remote SPAN (ERSPAN)
- Easy deployment using Cisco Virtual Switch Update Manager (VSUM) which allows you to install, upgrade, monitor and also migrate hosts to Cisco Nexus 1000V using the VMware vSphere web client.

- Starting with Cisco Nexus 1000V Release 4.2(1)SV2(1.1), a plug-in for the VMware vCenter Server, known as vCenter plug-in (VC plug-in) is supported on the Cisco Nexus 1000V virtual switch. It provides the server administrators a view of the virtual network and a visibility into the networking aspects of the Cisco Nexus 1000V virtual switch. The server administrator can thus monitor and manage networking resources effectively with the details provided by the vCenter plug-in. The vCenter plug-in is supported only on VMware vSphere Web Clients where you connect to VMware vCenter through a browser. The vCenter plug-in is installed as a new tab in the Cisco Nexus 1000V as part of the user interface in vSphere Web Client.

For more information, refer to:

- <http://www.cisco.com/en/US/products/ps9902/index.html>
- <http://www.cisco.com/en/US/products/ps10785/index.html>

Cisco MDS 9100 Series Fabric Switches

The Cisco® MDS 9148S 16G Multilayer Fabric Switch is the next generation of the highly reliable, flexible, and low-cost Cisco MDS 9100 Series switches. It combines high performance with exceptional flexibility and cost effectiveness. This powerful, compact one rack-unit (1RU) switch scales from 12 to 48 line-rate 16 Gbps Fibre Channel ports. Cisco MDS 9148S is powered by Cisco NX-OS and delivers advanced storage networking features and functions with ease of management and compatibility with the entire Cisco MDS 9000 Family portfolio for reliable end-to-end connectivity.

Cisco MDS 9148S is well suited as the following:

- Top-of-rack switch in medium-sized deployments
- Edge switch in a two-tiered (core-edge) data center topology
- Standalone SAN in smaller departmental deployments

The main features and benefits of Cisco MDS 9148S are summarized in the table below.

Table 1 Cisco MDS 9148S Features and Benefits

Feature	Benefit
Up to 48 autosensing Fibre Channel ports are capable of speeds of 2, 4, 8, and 16 Gbps, with 16 Gbps of dedicated bandwidth for each port. Cisco MDS 9148S scales from 12 to 48 high-performance Fibre Channel ports in a single 1RU form factor.	High Performance and Flexibility at Low Cost
Supports dual redundant hot-swappable power supplies and fan trays, PortChannels for Inter-Switch Link (ISL) resiliency, and F-port channeling for resiliency on uplinks from a Cisco MDS 9148S operating in NPV mode.	High-Availability Platform for Mission-Critical Deployments
Intelligent diagnostics/Hardware based slow port detection and Cisco Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) and Cisco Fabric Analyzer to capture and analyze network traffic. Fibre Channel ping and traceroute to identify exact path and timing of flows. Cisco Call Home for added	Enhanced performance and monitoring capability. Increase reliability, faster problem resolution, and reduce service costs

Feature	Benefit
reliability.	
In-Service Software Upgrades	Reduce downtime for planned maintenance and software upgrades
Aggregate up to 16 physical ISLs into a single logical PortChannel bundle with multipath load balancing.	High performance ISLs and optimized bandwidth utilization
Virtual output queuing on each port by eliminating head-of-line blocking	Helps ensure line-rate performance
PowerOn Auto Provisioning to automate deployment and upgrade of software images.	Reduces administrative costs
Smart zoning for creating and managing zones	Reduces consumption of hardware resources and administrative time
SAN management through a command-line interface (CLI) or Cisco Prime DCNM for SAN Essentials Edition, a centralized management tool. Cisco DCNM task-based wizards simplify management of single or multiple switches and fabrics including management of virtual resources end-to-end, from the virtual machine and switch to the physical storage.	Simplified Storage Management with built-in storage network management and SAN plug-and-play capabilities. Sophisticated Diagnostics
Fabric-wide per-VSAN role-based authentication, authorization, and accounting (AAA) services using RADIUS, Lightweight Directory Access Protocol (LDAP), Microsoft Active Directory (AD), and TACACS+. Also provides VSAN fabric isolation, intelligent, port-level packet inspection, Fibre Channel Security Protocol (FC-SP) host-to-switch and switch-to-switch authentication, Secure File Transfer Protocol (SFTP), Secure Shell Version 2 (SSHv2), and Simple Network Management Protocol Version 3 (SNMPv3) implementing Advanced Encryption Standard (AES). Other security features include control-plane security, hardware-enforced zoning, broadcast zones, and management access.	Comprehensive Network Security Framework
Virtual SAN (VSAN) technology for hardware-enforced, isolated environments within a physical fabric. Access control lists (ACLs) for hardware-based, intelligent frame processing. Advanced traffic management features, such as fabric-wide quality of service (QoS) and Inter-VSAN Routing (IVR) for resource sharing across vSANs. Zone-based QoS simplifies configuration and administration by using the familiar zoning concept.	Intelligent Network Services and Advanced Traffic Management for better and predictable network service without compromising scalability, reliability, availability, and network security
Common software across all platforms by using Cisco NX-OS and Cisco Prime DCNM across the fabric.	Reduce total cost of ownership (TCO) through consistent provisioning, management, and diagnostic capa-


Feature	Benefit
	bilities across the fabric

Nimble Storage – CS700 Hybrid Array

The Nimble Storage CS700 is designed for consolidating multiple large-scale critical applications with aggressive performance demands. This array handles IO-intensive workloads like larger scale VDI, Oracle or SQL server databases, while provides a compelling performance and IOPs per dollar.

For more information, refer to: <https://www.nimblestorage.com/products-technology/>

Figure 6 Nimble CS700 Array



Nimble Storage Predictive Flash Platform

The Nimble Storage Predictive Flash platform enables enterprise IT organizations to implement a single architectural approach to dynamically cater to the needs of varying workloads, driven by the power of predictive analytics. Predictive Flash is the only storage platform that optimizes across performance, capacity, data protection, and reliability within a dramatically smaller footprint.

Predictive Flash is built upon Nimble’s CASL™ architecture, NimbleOS and InfoSight™, the company’s cloud-connected predictive analytics and management system. CASL scales performance and capacity seamlessly and independently. InfoSight leverages the power of deep data analytics to provide customers with precise guidance on the optimal approach to scaling flash, CPU, and capacity around changing application needs, while ensuring peak storage health.

NimbleOS Architecture

The Nimble Storage operating system, NimbleOS is based on its patented Cache Accelerated Sequential Layout (CASL™) architecture. CASL leverages the unique properties of flash and disk to deliver high performance and capacity – all within a dramatically small footprint. CASL and InfoSight™ form the foundation of the Predictive Flash platform, which allows for the dynamic and intelligent deployment of storage resources to meet the growing demands of business-critical applications.

Nimble Storage InfoSight

Using systems modeling, predictive algorithms, and statistical analysis, InfoSight™ solves storage administrators’ most difficult problems. InfoSight also ensures storage resources are dynamically and intelligently deployed to satisfy the changing needs of business-critical applications, a key facet of Nimble Storage’s Predictive Flash platform. At the heart of InfoSight is a powerful engine comprised of deep data

analytics applied to telemetry data gathered from Nimble arrays deployed across the globe. More than 30 million sensor values are collected per day per Nimble Storage array. The InfoSight Engine transforms the millions of gathered data points into actionable information that allows customers to realize significant operational efficiency through:

- Maintaining optimal storage performance
- Projecting storage capacity needs
- Proactively monitoring storage health and getting granular alerts
- Proactively diagnoses and automatically resolves complex issues, freeing up IT resources for value-creating projects
- Ensures a reliable data protection strategy with detailed alerts and monitoring
- Expertly guides storage resource planning, determining the optimal approach to scaling cache, IOPS to meet changing SLAs
- Identifies latency and performance bottlenecks through the entire virtualization stack

For more information, refer to: <https://www.nimblestorage.com/infosight/architecture/>

In-Line Compression

CASL uses fast, in-line compression for variable application block sizes to decrease the footprint of inbound write data by as much as 75 percent. Once there are enough variable-sized blocks to form a full write stripe, CASL writes the data to disk. If the data being written is active data, it is also copied to SSD cache for faster reads. Written data is protected with triple-parity RAID.

Thin-Provisioning and Efficient Capacity Utilization

Capacity is only consumed as data is written. CASL efficiently reclaims free space on an ongoing basis, preserving write performance with higher levels of capacity utilization. This avoids fragmentation issues that hamper other architectures.

Accelerated Write Performance

By sequencing random write data, CASL's writes to disk are orders of magnitude faster than other storage systems' random writes. The CS700, Nimble's top-of-the-line array, delivers double the write IOPS of a single MLC flash drive with a 7,200-RPM hard disk.

Read Performance

CASL accelerates read performance by dynamically caching hot data in flash, delivering sub-millisecond read latency and high throughput across a wide variety of demanding enterprise applications.

Adaptive Flash

CASL leverages flash as a true read cache, as opposed to a bolt-on tier. This enables Nimble arrays to easily adapt to changing workloads. As the architectural foundation of Adaptive Flash, CASL allows flash to flexibly scale for higher performance, especially benefitting those applications that work best when their entire working sets reside in flash.

Intelligent Caching

CASL leverages flash as a true read cache, as opposed to a bolt-on tier. This enables Nimble arrays to easily adapt to changing workloads. As the architectural foundation of Adaptive Flash, CASL allows flash to flexibly scale for higher performance, especially benefitting those applications that work best when their entire working sets reside in flash.

Adaptive Flash Service Levels

Flash can be allocated to individual workloads on a per-volume basis according to one of three user-assignable service levels:

- All Flash: The entire workload is pinned in cache for deterministic low latency. Ideal for latency-sensitive workloads or single applications with large working sets or high cache churn.
- Auto Flash: Default service level where workload active data is dynamically cached. Ideal for applications requiring high performance, or a balance of performance and capacity.
- No Flash: No active data is cached in flash. Recommended for capacity-optimized workloads without high performance demands.

Efficient, Fully Integrated Data Protection

All-inclusive snapshot-based data protection is built into the Adaptive Flash platform. Snapshots and production data reside on the same array, eliminating the inefficiencies inherent to running primary and backup storage silos. And, InfoSight ensures that customers' data protection strategies work as expected through intuitive dashboards and proactive notifications in case of potential issues.

SmartSnap: Thin, Redirect-on Write Snapshots

Nimble snapshots are point-in-time copies capturing just changed data, allowing three months of frequent snapshots to be easily stored on a single array. Data can be instantly restored, as snapshots reside on the same array as primary data.

SmartReplicate: Efficient Replication

Only compressed, changed data blocks are sent over the network for simple and WAN-efficient disaster recovery.

Zero-Copy Clones

Nimble's snapshots allow fully functioning copies, or clones of volumes, to be quickly created. Instant clones deliver the same performance and functionality as the source volume, an advantage for virtualization, VDI, and test/development workloads.

Application-Consistent Snapshots

Nimble enables instant application/VM-consistent backups using VSS framework and VMware integration, using application templates with pre-tuned storage parameters.

SmartSecure: Flexible Data Encryption

NimbleOS enables encryption of individual volumes with little to no performance impact. Encrypted volumes can be replicated to another Nimble target, and data can be securely shredded.

REST API

NimbleOS has a RESTful API that allows for powerful, secure and scriptable management of storage objects. Using this API, an administrator can use the scripting or orchestration to interact with infrastructure components in a well-defined, repeatable manner.

For more information, see: <https://www.nimblestorage.com/products-technology/casl-architecture/>

Citrix XenApp and XenDesktop 7.7

Citrix XenApp and XenDesktop are application and desktop virtualization solutions built on a unified architecture so they're simple to manage and flexible enough to meet the needs of all your organization's users. XenApp and XenDesktop have a common set of management tools that simplify and automate IT tasks. You use the same architecture and management tools to manage public, private, and hybrid cloud deployments as you do for on premises deployments.

Citrix XenApp delivers the following:

- XenApp published apps, also known as server-based hosted applications: These are applications hosted from Microsoft Windows servers to any type of device, including Windows PCs, Macs, smartphones, and tablets. Some XenApp editions include technologies that further optimize the experience of using Windows applications on a mobile device by automatically translating native mobile-device display, navigation, and controls to Windows applications; enhancing performance over mobile networks; and enabling developers to optimize any custom Windows application for any mobile environment.
- XenApp published desktops, also known as server-hosted desktops: These are inexpensive, locked-down Windows virtual desktops hosted from Windows server operating systems. They are well suited for users, such as call center employees, who perform a standard set of tasks.
- Virtual machine-hosted apps: These are applications hosted from machines running Windows desktop **operating systems for applications that can't be hosted in a server environment.**
- Windows applications delivered with Microsoft App-V: These applications use the same management tools that you use for the rest of your XenApp deployment.

Citrix XenDesktop 7.7 includes significant enhancements to help customers deliver Windows apps and desktops as mobile services while addressing management complexity and associated costs. Enhancements in this release include:

- Unified product architecture for XenApp and XenDesktop—the FlexCast Management Architecture (FMA). This release supplies a single set of administrative interfaces to deliver both hosted-shared applications (RDS) and complete virtual desktops (VDI). Unlike earlier releases that separately provisioned Citrix XenApp and XenDesktop farms, the XenDesktop 7.7 release allows administrators to deploy a single infrastructure and use a consistent set of tools to manage mixed application and desktop workloads.
- Support for extending deployments to the cloud. This release provides the ability for hybrid cloud provisioning from Amazon Web Services (AWS) or any Cloud Platform-powered public or private cloud. Cloud deployments are configured, managed, and monitored through the same administrative consoles as deployments on traditional on-premises infrastructure.

- Enhanced HDX technologies. Since mobile technologies and devices are increasingly prevalent, Citrix has engineered new and improved HDX technologies to improve the user experience for hosted Windows apps and desktops.
- A new version of StoreFront. The StoreFront 2.5 release provides a single, simple, and consistent aggregation point for all user services. Administrators can publish apps, desktops, and data services to StoreFront, from which users can search and subscribe to services.
- **Remote power control for physical PCs. Remote PC access supports “Wake on LAN” that adds the ability to power on physical PCs remotely. This allows users to keep PCs powered off when not in use to conserve energy and reduce costs.**
- Full AppDNA support. AppDNA provides automated analysis of applications for Windows platforms and suitability for application virtualization through App-V, XenApp, or XenDesktop. Full AppDNA functionality is available in some editions.
- Additional virtualization resource support. As in this Cisco Validated Design, administrators can configure connections to VMware vSphere 5.5 hypervisors.

Citrix XenDesktop delivers:

- VDI desktops: These virtual desktops each run a Microsoft Windows desktop operating system rather than running in a shared, server-based environment. They can provide users with their own desktops that they can fully personalize.
- Hosted physical desktops: This solution is well suited for providing secure access powerful physical machines, such as blade servers, from within your data center.
- Remote PC access: This solution allows users to log in to their physical Windows PC from anywhere over a secure XenDesktop connection.
- Server VDI: This solution is designed to provide hosted desktops in multitenant, cloud environments.
- Capabilities that allow users to continue to use their virtual desktops: These capabilities let users continue to work while not connected to your network.



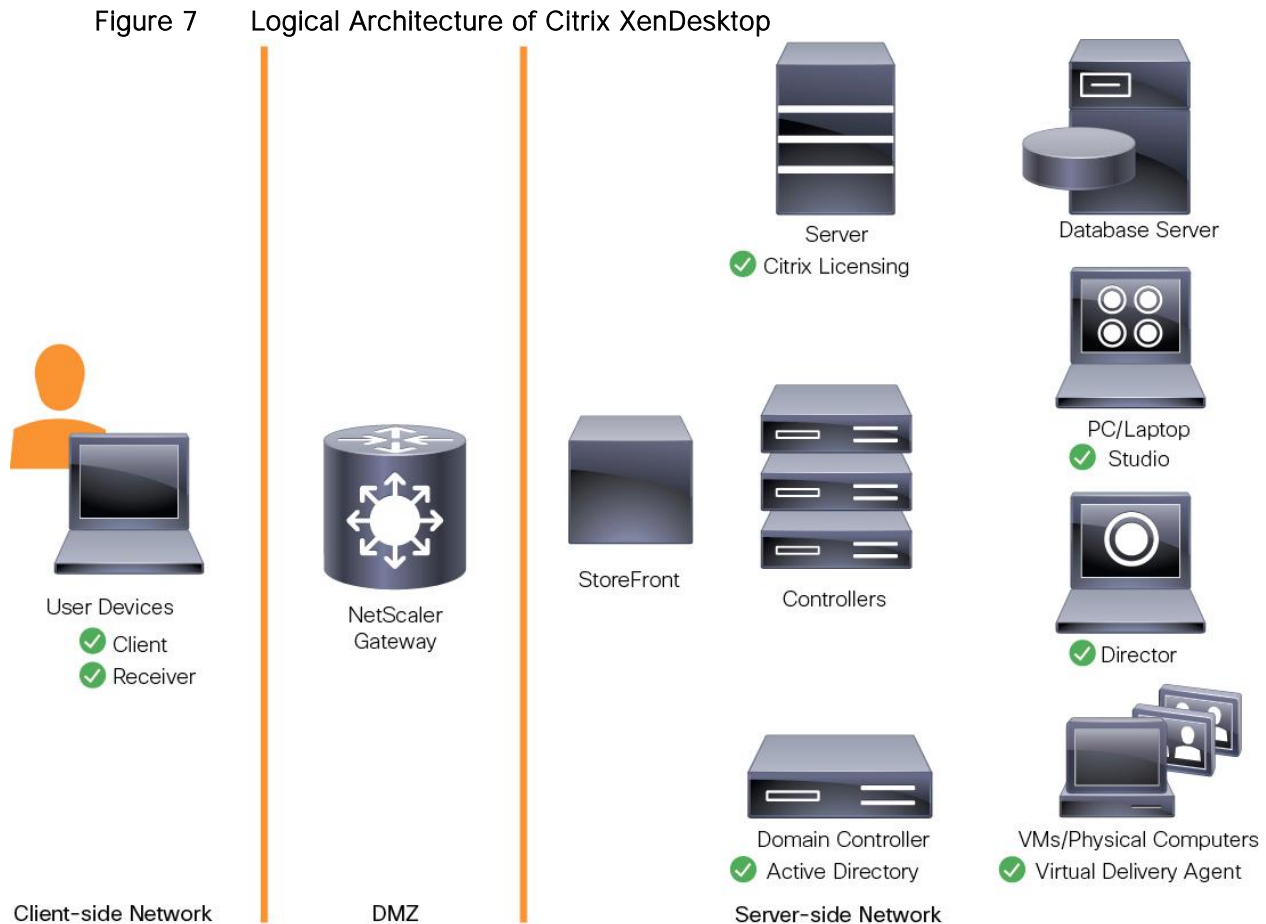
Some XenDesktop editions include the features available in XenApp.

Release 7.7 of XenDesktop includes new features that make it easier for users to access applications and desktops and for Citrix administrator to manage applications:

- The session prelaunch and session linger features help users quickly access server-based hosted applications by starting sessions before they are requested (session prelaunch) and keeping application sessions active after a user closes all applications (session linger).
- Support for unauthenticated (anonymous) users means that users can access server-based hosted applications and server-hosted desktops without presenting credentials to Citrix StoreFront or Receiver.
- Connection leasing makes recently used applications and desktops available even when the site database is unavailable.
- Application folders in Citrix Studio make it easier to administer large numbers of applications.

Other new features in this release allow you to improve performance by specifying the number of actions that can occur on a site's host connection; display enhanced data when you manage and monitor your site; and anonymously and automatically contribute data that Citrix can use to improve product quality, reliability, and performance.

For more information about the features new in this release, see [Citrix XenDesktop Release 7.7](#).



Citrix Provisioning Services 7.7

Most enterprises struggle to keep up with the proliferation and management of computers in their environments. Each computer, whether it is a desktop PC, a server in a data center, or a kiosk-type device, must be managed as an individual entity. The benefits of distributed processing come at the cost of distributed management. It costs time and money to set up, update, support, and ultimately decommission each computer. The initial cost of the machine is often dwarfed by operating costs.

Citrix PVS takes a very different approach from traditional imaging solutions by fundamentally changing the relationship between hardware and the software that runs on it. By streaming a single shared disk image (vDisk) rather than copying images to individual machines, PVS enables organizations to reduce the number of disk images that they manage, even as the number of machines continues to grow, simultaneously providing the efficiency of centralized management and the benefits of distributed processing.

In addition, because machines are streaming disk data dynamically and in real time from a single shared image, machine image consistency is essentially ensured. At the same time, the configuration, applications,

and even the OS of large pools of machines can be completely changed in the time it takes the machines to reboot.

Using PVS, any vDisk can be configured in standard-image mode. A vDisk in standard-image mode allows many computers to boot from it simultaneously, greatly reducing the number of images that must be maintained and the amount of storage that is required. The vDisk is in read-only format, and the image cannot be changed by target devices.

These same benefits apply to vDisks that are streamed to bare metal servers, which is the way we utilized PVS in this study.

Citrix Desktop Studio for XenApp 7.7

If you manage a pool of servers that work as a farm, such as Citrix XenApp servers or web servers, maintaining a uniform patch level on your servers can be difficult and time consuming. With traditional imaging solutions, you start with a clean golden master image, but as soon as a server is built with the master image, you must patch that individual server along with all the other individual servers. Rolling out patches to individual servers in your farm is not only inefficient, but the results can also be unreliable. Patches often fail on an individual server, and you may not realize you have a problem until users start complaining or the server has an outage. After that happens, getting the server resynchronized with the rest of the farm can be challenging, and sometimes a full reimaging of the machine is required.

With Citrix PVS, patch management for server farms is simple and reliable. You start by managing your golden image, and you continue to manage that single golden image. All patching is performed in one place and then streamed to your servers when they boot. Server build consistency is assured because all your servers use a single shared copy of the disk image. If a server becomes corrupted, simply reboot it, and it is instantly back to the known good state of your master image. Upgrades are extremely fast to implement. After you have your updated image ready for production, you simply assign the new image version to the servers and reboot them. You can deploy the new image to any number of servers in the time it takes them to reboot. Just as important, rollback can be performed in the same way, so problems with new images do not need to take your servers or your users out of commission for an extended period of time.

Benefits for Desktop Administrators

Because Citrix PVS is part of Citrix XenApp, **desktop administrators can use PVS's streaming technology to** simplify, consolidate, and reduce the costs of both physical and virtual desktop delivery. Many organizations are beginning to explore desktop virtualization. Although virtualization addresses **many of IT's needs for** consolidation and simplified management, deploying it also requires deployment of supporting infrastructure. Without PVS, storage costs can make desktop virtualization too costly for the IT budget. However, with PVS, IT can reduce the amount of storage required for VDI by as much as 90 percent. And with a single image to manage instead of hundreds or thousands of desktops, PVS significantly reduces the cost, effort, and complexity for desktop administration.

Different types of workers across the enterprise need different types of desktops. Some require simplicity and standardization, and others require high performance and personalization. XenApp can meet these requirements in a single solution using Citrix FlexCast delivery technology. With FlexCast, IT can deliver every type of virtual desktop, each specifically tailored to meet the performance, security, and flexibility requirements of each individual user.

Not all desktop applications can be supported by virtual desktops. For these scenarios, IT can still reap the benefits of consolidation and single-image management. Desktop images are stored and managed centrally in the data center and streamed to physical desktops on demand. This model works particularly well for standardized desktops such as those in lab and training environments and call centers and thin-client devices used to access virtual desktops.

Citrix Provisioning Services Solution

Citrix PVS streaming technology allows computers to be provisioned and re-provisioned in real time from a single shared disk image. With this approach, administrators can completely eliminate the need to manage and patch individual systems. Instead, all image management is performed on the master image. The local hard drive of each system can be used for runtime data caching or, in some scenarios, removed from the system entirely, which reduces power use, system failure rate, and security risk.

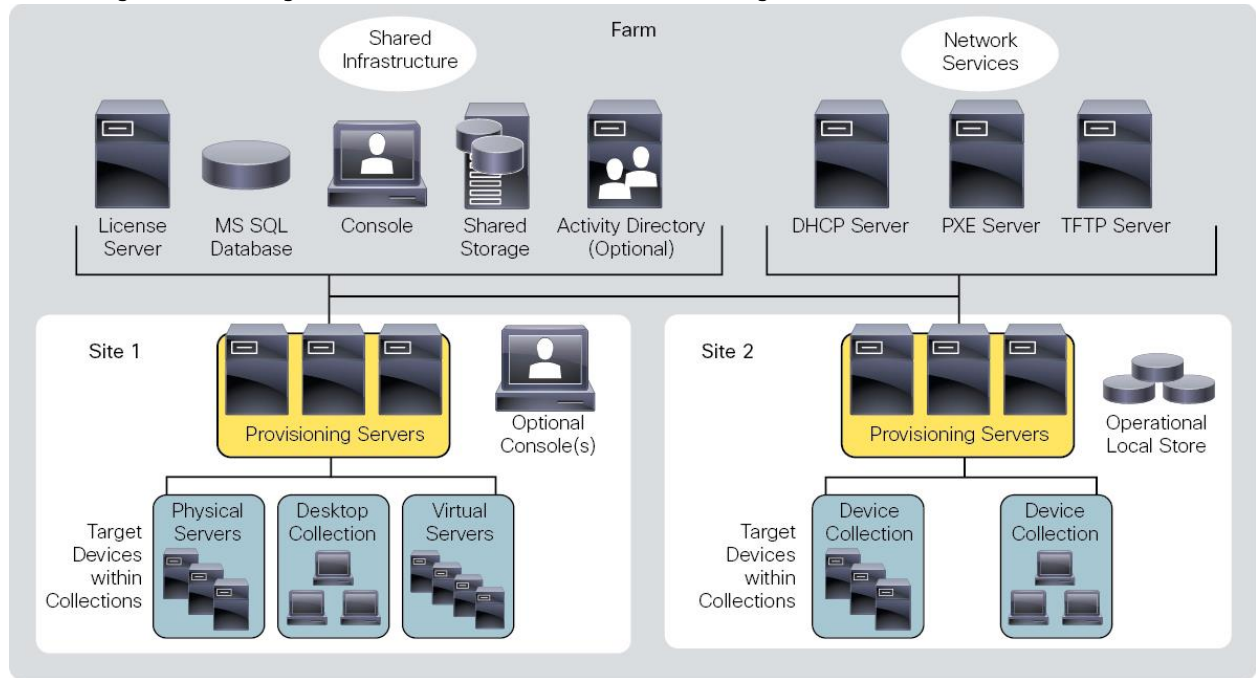
The PVS solution's infrastructure is based on software-streaming technology. After PVS components are installed and configured, a vDisk is created from a device's hard drive by taking a snapshot of the OS and application image and then storing that image as a vDisk file on the network. A device used for this process is referred to as a master target device. The devices that use the vDisks are called target devices. vDisks can exist on a PVS, file share, or in larger deployments, on a storage system with which PVS can communicate (iSCSI, SAN, network-attached storage [NAS], and Common Internet File System [CIFS]). vDisks can be assigned to a single target device in private-image mode, or to multiple target devices in standard-image mode.

Citrix Provisioning Services Infrastructure

The Citrix PVS infrastructure design directly relates to administrative roles within a PVS farm. The PVS administrator role determines which components that administrator can manage or view in the console.

A PVS farm contains several components. 0 provides a high-level view of a basic PVS infrastructure and shows how PVS components might appear within that implementation.

Figure 8 Logical Architecture of Citrix Provisioning Services



Solution Architecture

SmartStack solution delivers a converged infrastructure platform that incorporates compute, network and storage best practices from Cisco and Nimble to deliver a resilient, scalable and flexible datacenter architecture for Enterprise and cloud deployments.

The following platforms are integrated in this SmartStack architecture:

- Cisco Unified Computing System (UCS) - B-Series blade servers
- Cisco UCS 6200 Series Fabric Interconnects (FI) - unified access to storage and LAN networks
- Cisco Nexus 9300 series switches - connectivity to users, other LAN networks and Cisco UCS domains
- Cisco MDS 9100 fabric switches - SAN fabric providing Fibre Channel (FC) connectivity to storage
- Nimble CS700 array - SAN boot of hybrid storage array with SSDs and HDDs
- Cisco Nexus 1000V - access layer switch for virtual machines
- VMware vSphere 6.0 U1a - Hypervisor
- Cisco UCS Blade Server NVidia M6 GPU

This SmartStack architecture uses Cisco UCS B-series and M-series rack mount servers for compute. Cisco UCS B-series servers are housed in a Cisco UCS 5108 blade server chassis that can support up to eight half-width blades or four full-width blades. Each server supports a number of converged network adaptors (CNAs) that converge LAN and SAN traffic onto a single interface rather than requiring multiple network interface cards (NICs) and host bus adapters (HBAs) per server. **Cisco's Virtual Interface Cards (VICs) support 256 virtual interfaces and supports Cisco's VM-FEX technology** (see link below). Two CNAs are typically deployed on a server for redundant connections to the fabric. Cisco VIC is available as a Modular LAN on Motherboard (mLOM) card and as a Mezzanine Slot card. A half-height PCI Express (PCIe) card form-factor is also available but exclusively for Cisco UCS C-series rack mount servers. are available. For more information on the different models of Cisco UCS VIC adapters available, see:

<http://www.cisco.com/c/en/us/products/interfaces-modules/unified-computing-system-adapters/index.html>

All compute resources in the data center connect into a redundant pair of Cisco UCS fabric interconnects that provide unified access to storage and other parts of the network. Each blade server chassis requires 2 Fabric Extender (FEX) modules that extend the unified fabric to the blade server chassis and connect into the FIs using 2, 4 or 8 10GbE links. Each Cisco UCS 5100 Series chassis can support up to two fabric extenders that provide active-active data plane forwarding with failover for higher throughput and availability. FEX is a consolidation point for all blade server I/O traffic, which simplifies operations by providing a single point of management and policy enforcement. For detailed information about FEX see:

<http://www.cisco.com/c/en/us/solutions/data-center-virtualization/fabric-extender-technology-fex-technology/index.html>.

Two second generation models of FEX are currently available for the Cisco UCS blade server chassis. Cisco UCS 2204XP and 2208XP connect to the unified fabric using multiple 10GbE links. The number of 10GbE

links that are supported for connecting to an external fabric interconnect and to the blade servers within the chassis are shown in the table below. The maximum I/O throughput possible through each FEX is also shown.

Table 2 Blade Server Fabric Extender Models

Blade Server Models	Internal Facing Links to Blade Servers	External Facing Links to FI	Aggregate I/O Bandwidth
Cisco UCS 2204XP	16 x 10GbE	Up to 4 x 10GbE	40Gbps per FEX
Cisco UCS 2208XP	32 x 10GbE	Up to 8 x 10GbE	80Gbps per FEX

By deploying a pair of Cisco 2208XP FEX, a Cisco UCS 5100 series chassis with blade servers can get up to maximum of 160Gbps of I/O throughput for the servers on that chassis. For additional details on the FEX, see:

http://www.cisco.com/en/US/prod/collateral/ps10265/ps10276/data_sheet_c78-675243.html

The rack mount servers can also benefit from a FEX architecture to aggregate I/O traffic from several rack mount servers but unlike the blade servers where the FEX modules fit into the back of the chassis, rack mount servers require a standalone FEX chassis. Cisco 2200 Series FEX model is functionally equivalent to the above blade server FEX models. Other than the physical cabling required to connect ports on Cisco Nexus FEX 2300 to servers and FIs, the discovery and configuration is same as that of the blade server to FI connectivity. For data centers migrating their access-switching layer to 40GbE speeds, Cisco Nexus 2300 series FEX is the newest model. For additional details on the Cisco Nexus 2000 Series FEX, see:

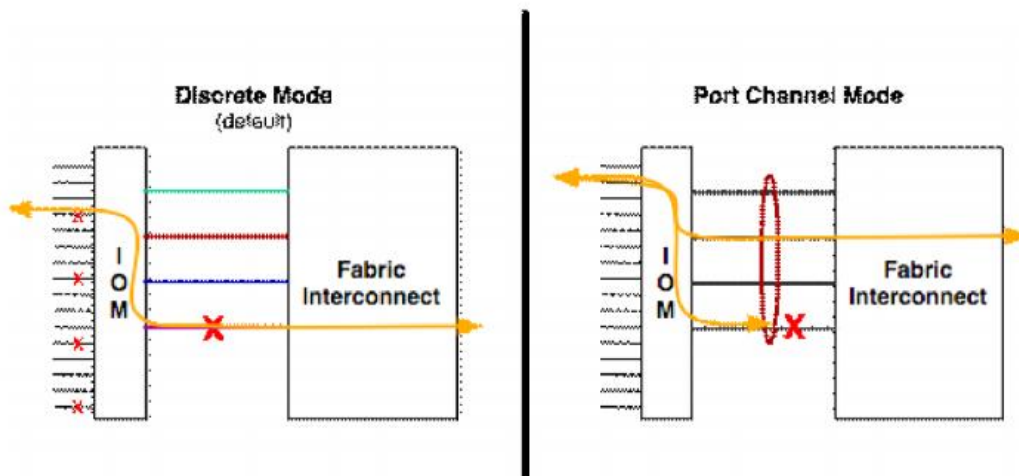
<http://www.cisco.com/c/en/us/products/switches/nexus-2000-series-fabric-extenders/index.html>

Rack mount servers can also be deployed by directly connecting into the FIs, without using Cisco UCS FEX chassis. However, this could mean additional operational overhead by having to manage server policies individually but nevertheless a valid option in smaller deployments.

The modular server chassis connect directly into FIs using break out cables that go from each 40G QSFP+ port on the server side to 4x10GbE ports on each FI. A second QSFP+ is used to connect to the secondary FI. For additional details on the Cisco Nexus 2000 Series FEX, see:

<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-m-series-modular-servers/index.html>

The fabric extenders in a blade server chassis connect externally to the unified fabric using multiple 10GbE links – the number of links depends on the aggregate I/O bandwidth required. Each FEX connect into a Cisco UCS 6200 series fabric interconnect using up to 4 or 8 10GbE links depending on the model of FEX. The links can be deployed as independent links (discrete Mode) or grouped together using link aggregation (port channel mode). In discrete mode, each server is pinned to a FEX link going to a port on the fabric interconnect and if the link goes down, the server's connection also goes down through the FEX link. In port channel mode, the flows from the server will be redistributed across the remaining port channel members. This is less disruptive overall and therefore port channel mode is preferable.



Cisco UCS system provides the flexibility to individually select components of the system that meet the performance and scale requirements of a customer deployment. There are several options for blade and rack servers, network adapters, FEX and Fabric Interconnects that make up the Cisco UCS system.

Compute resources are grouped into an infrastructure layer and application data layer. Servers in the infrastructure layer are dedicated for hosting virtualized infrastructure services that are necessary for deploying and managing the entire data center. Servers in the application data layer are for hosting business applications and supporting services that Enterprise users will access to fulfill their business function.

The architecture can support any hypervisor but VMware vSphere will be used to validate the architectures. High availability features available in the hypervisor will be leveraged to provide resiliency at the virtualization layer of the stack.

SmartStack architecture with Nimble Storage array can support block storage using either iSCSI or Fibre Channel (FC). The focus of this CVD will be fibre channel access to a Nimble CS700 array. For more details on an iSCSI based SmartStack design using Nimble CS 300 array, see the following links in the Cisco Design Zone for SmartStack:

- Design Guide:
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/smartstack_cs300_mini.html
- Deployment Guide:
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/smartstack_cs300_mini_deploy.html

The CS-Series arrays are based on **Nimble's** CPU-driven architecture that provide enterprises with the ability to scale performance and capacity independently. The Predictive Flash platform is based on the NimbleOS and its **patented Cache Accelerated Sequential Layout (CASL™) architecture**, and **InfoSight™**, a cloud-based management and support system. CASL maximizes the unique properties of flash and disk to deliver high performance and capacity with a dramatically small footprint. CASL and **InfoSight™** form the foundation of the Predictive Flash platform, which allows for the dynamic and intelligent deployment of storage resources to meet the growing demands of business-critical applications.

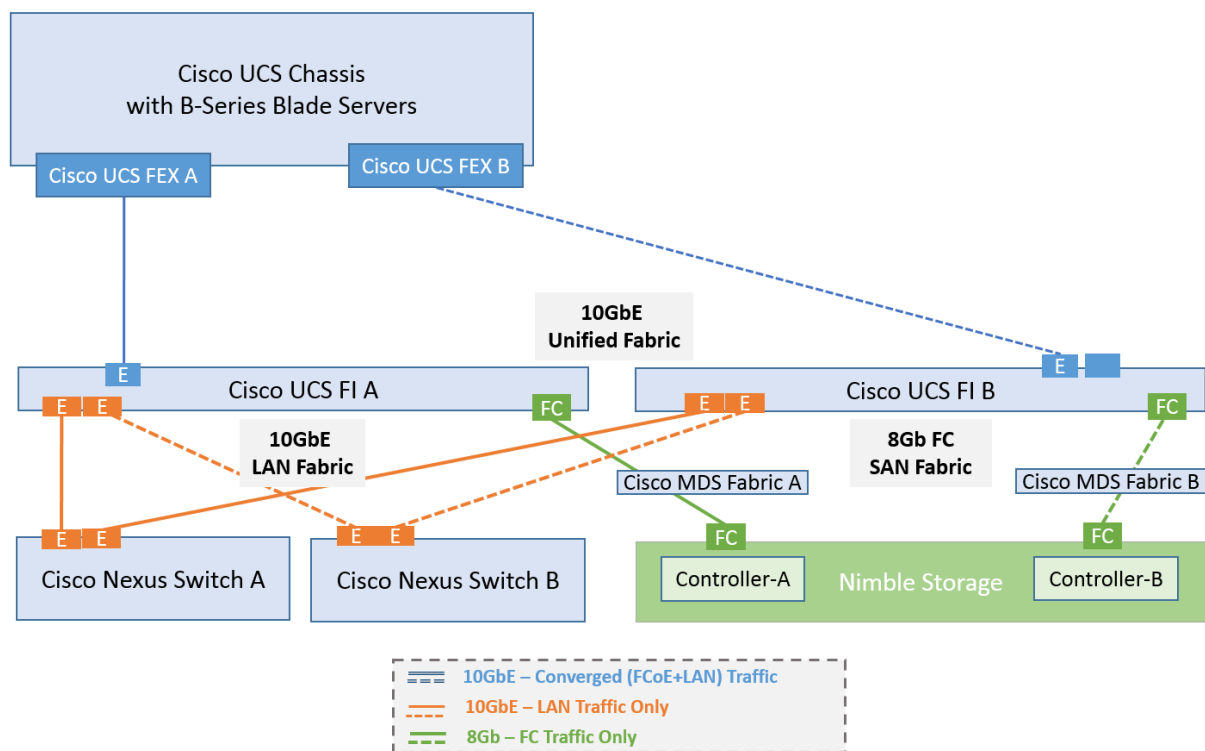
Nimble's CS700 is a hybrid array designed for large scale enterprise applications with high performance needs. The base array can support up to 12 hard disk drives (HDD) and 4 solid-state drives (SSD). The maximum raw capacity of the base array using 6TB HDD drives is 72TB, with an effective capacity of around 50-100 TB. The flash capacity of the base array using 1.6 TB SSD drives is approximately 3.2-7.6 TB. The capacity and performance can be extended using expansion shelves, with the best model providing up to 90TB of raw capacity and 66-132TB of effective capacity per shelf. The SSD expansion shelf can deliver an additional 30.72TB of flash capacity. CS700 supports up to 6 expansion shelves which can be combined to meet the performance and capacity needs of enterprise data centers. For additional details of CS700, see:

<http://www.nimblestorage.com/products-technology/products-specs/>

If a unified fabric from compute to storage is required or preferable, iSCSI access can be used but the focus of this SmartStack solution is on a FC based access which currently requires a dedicated SAN network. The traffic does traverse a unified fabric between Cisco UCS and Fabric Interconnects but then diverge onto separate LAN and SAN networks as shown in Figure 9.

This SmartStack architecture utilizes a dedicated SAN network for block storage traffic. Figure 9 shows a unified fabric (blue lines) between Cisco UCS servers and Fabric Interconnects which then splits into separate SAN (green lines) and LAN (orange lines) network. If a unified fabric is required or preferable end-to-end from compute to storage, iSCSI block storage access can be used. Alternatively, a Cisco Nexus 5000 series or MDS switch can be used for FCoE to FC connectivity since FCoE capability is not directly available to the Nimble Storage array. The focus of this SmartStack solution is on a dedicated fibre channel based storage access design using Cisco MDS fabric switches.

Figure 9 SmartStack Architecture - High Level



The SmartStack architecture is a highly resilient design with redundant Unified, LAN and SAN fabrics that includes component and link level redundancy when accessing these fabrics. The unified fabric and LAN

connections are 10Gb Ethernet links and the SAN connections are 8Gb FC links. Multiple 10GbE links and FC links are deployed in the SmartStack architecture with link aggregation using Link Aggregation Control Protocol (LACP) to provide higher aggregate bandwidth and resiliency across the different fabrics. Use of LACP is strongly recommended when available for improved failover convergence time and protection from misconfigurations. Static or manual bundling is an alternative but is less preferable and therefore not used in this architecture.

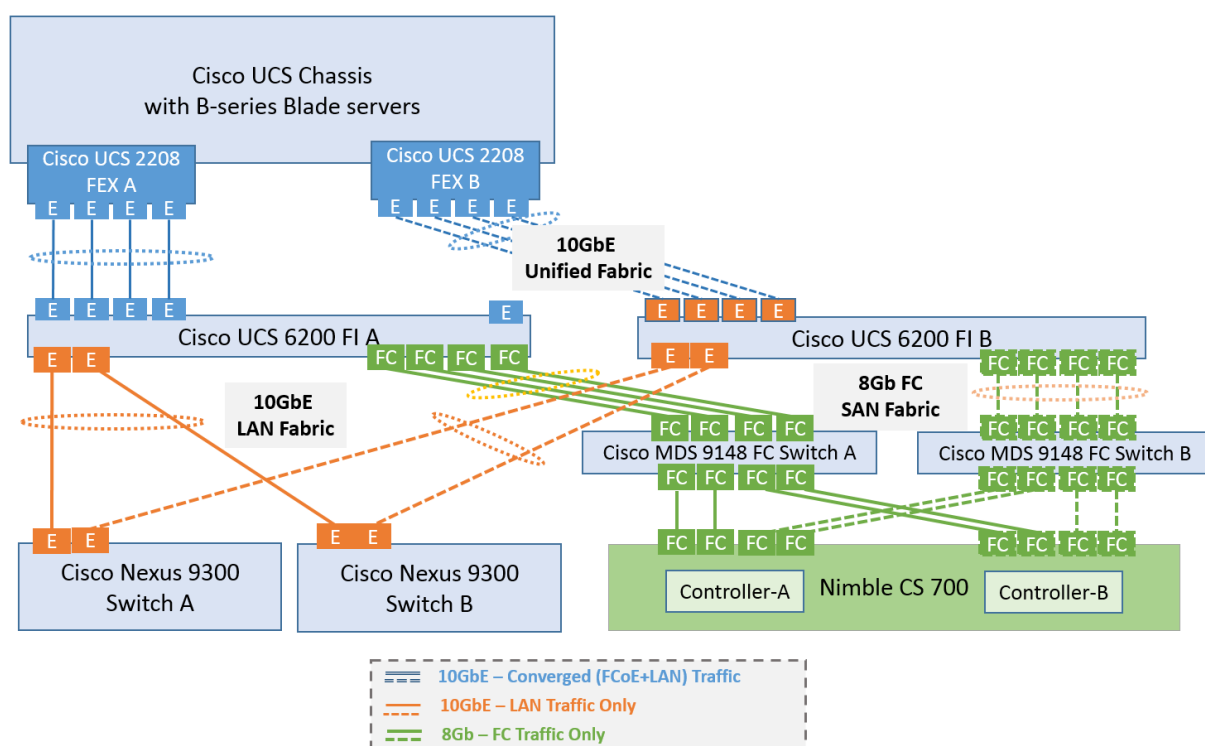
The data center LAN network in this SmartStack architecture uses a pair of Cisco Nexus 9300 switches that serve as the access/aggregation layer of the data center network. In this design, the Nexus 9300s provide reachability to users and other parts of the network. In larger deployments, an additional layer of hierarchy can be added using Cisco Nexus 9500 series switches as an aggregation/core layer in a classic data center tiered design or as a spine in a spine/leaf design. Cisco Nexus 9000 family of switches can operate in one of two modes: Application Centric Infrastructure (ACI) for newer cloud architectures or in Cisco NX-OS standalone mode for the more traditional data center architectures. In the SmartStack architecture, the Cisco 9300s are deployed in NX-OS standalone mode and provide investment protection by enabling a pathway to **Cisco's Application Centric Infrastructure (ACI)**. **Cisco ACI was designed to help Enterprises transition their** data centers to a cloud architecture that is dynamic, where applications can be quickly deployed and scaled, and adapt with the needs of the business. To enable this, Cisco ACI provides a flexible, scalable, application centric, policy-based framework based on open APIs that accelerate and simplify datacenter deployments while providing centralized orchestration, visibility and management.

Virtual Port Channel (vPC) or link-aggregation capabilities of the Cisco Nexus 9000 family of switches are used on the network links to Cisco UCS Fabric Interconnects. A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single PortChannel. vPC provides Layer 2 multipathing with load balancing by allowing multiple parallel paths between nodes that result in increased bandwidth and redundancy. A vPC-based architecture is therefore highly resilient and robust and scales the available Layer 2 bandwidth by using all available links. Other benefits of vPCs include:

- Provides a loop-free topology
- Eliminates Spanning Tree Protocol blocked ports
- Uses all available uplink bandwidth
- Provides fast convergence for both link and device failures
- Provides higher aggregate bandwidth by adding links – same as Port Channels
- Helps ensure high availability

The SAN network in this SmartStack architecture is 8Gb FC based and uses Cisco MDS 9100 switches.

Figure 10 SmartStack Architecture



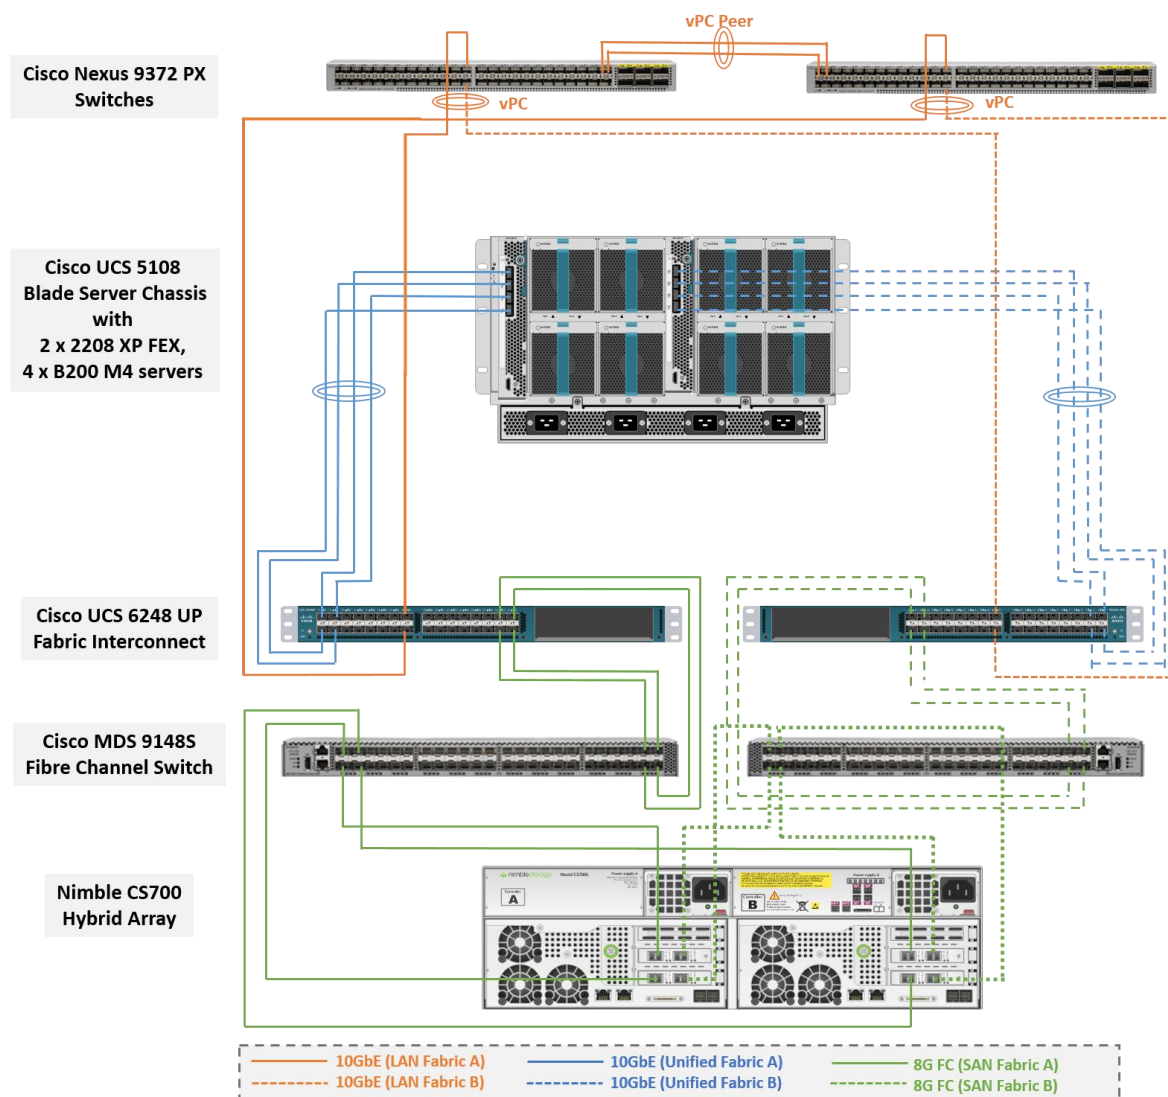
The aggregate throughput of the two fabrics depends on the number of 10GbE and 8Gb FC links deployed to meet the throughput needs of a given customer deployment. The figure above shows a 4 x 10GbE access to the unified fabric from Cisco UCS blade server chassis and 10GbE from the rack mount server. This can be scaled up to a max of 8x10 GbE on the UCS blade server chassis. The LAN and SAN fabric provides a 2x10 GbE and 4x8Gb FC access respectively from each FI. This can also be scaled higher by adding additional 10GbE and 8Gb FC links.

The SmartStack platform will be managed locally using Cisco UCS Manager, VMware vCenter and Nimble Management software. The storage array will also be remotely monitored from the cloud using Nimble **InfoSight™** to provide insights into I/O and capacity usage, trend analysis for capacity planning and expansion, and for pro-active ticketing and notification when issues occur.

Solution Design

The end-to-end SmartStack design is shown in Figure 11.

Figure 11 SmartStack Design



Compute

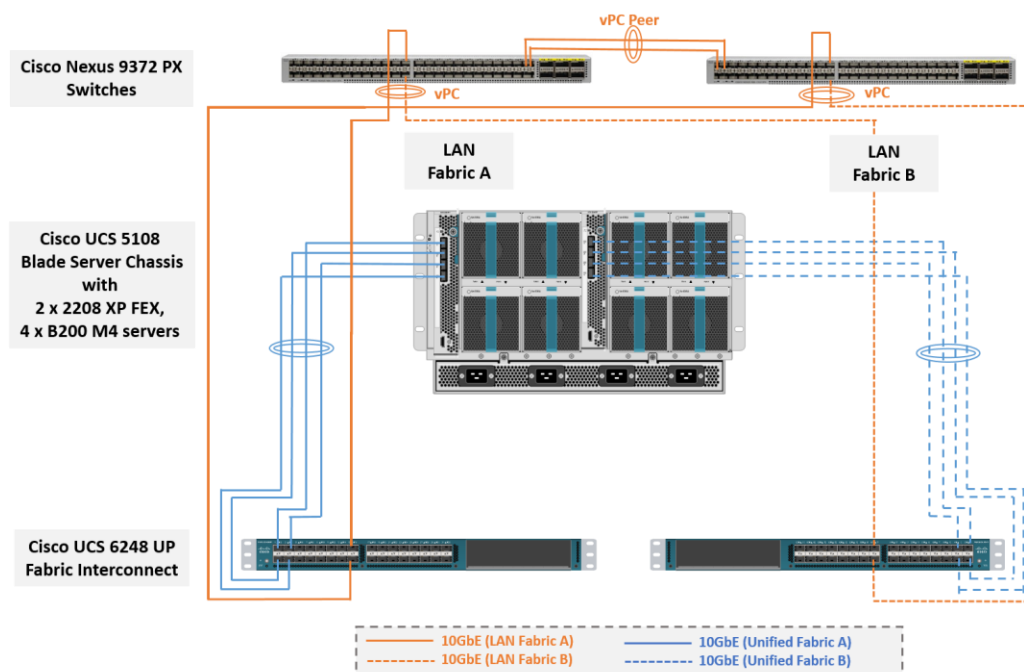
SmartStack design uses Cisco UCS with 16x Cisco B200M4 half-width blades to provide the compute resources. A Cisco M4308 M-Series Chassis is also included in the design to run XenApp sessions on. The hypervisor layer in the design is provided by VMware ESXi 6.0 U1a. Features available at the hypervisor layer (for example, VMware clustering, high availability) are leveraged where possible for a robust design. The blade server chassis is connected via FEX modules on the back of the chassis to a pair of Cisco UCS 6248 FIs. A pair of Cisco 2204 XP fabric extenders is used in this design. The FEX to FI connectivity uses 8x10GbE links, 4 from FEX-A to FI-A and 4 from FEX-B to FI-B to provide an aggregate access bandwidth of 80Gbps

to the unified fabric. The FIs are connected to LAN and SAN network using 10GbE and 8Gb FC links. The FI provides 40Gbps of aggregate bandwidth to the LAN network and 64Gbps to the SAN network. Link aggregation using port channels are used on the unified fabric FEX-FI connections and virtual port channels on the Nexus-FI LAN connections.

Network

The LAN network design is shown in Figure 12.

Figure 12 SmartStack LAN Fabric Design



SmartStack LAN Design

The LAN network provides network reachability to the applications hosted on Cisco UCS servers in the data center. The infrastructure consists of a pair of Nexus 9372 PX switches deployed in NX-OS standalone mode. Redundant 10Gbps links from each Nexus switch are connected to ports on each FI and provide 20Gbps of bandwidth through each Nexus. Virtual PortChannels (vPCs) are used on the Nexus links going to each FI. VLAN Trunking is enabled on these links as multiple application data vlans, management and vMotion traffic needs to traverse these links. Jumbo Frames are also enabled in the LAN network to support vMotion between multiple UCS domains. See Design Practices section for other Nexus 9000 best practices in the design.

The SAN network provides fibre channel connectivity to the Nimble storage array and consists of a pair of MDS switches. The MDS switches form completely separate fabrics (SAN fabric A, SAN fabric B) and use a dual vSAN (vSAN-A, vSAN-B) design to provide two redundant and completely diverse paths to the Nimble array.

Link aggregation using port channels are used to aggregate 4 x 8G FC links to provide 32G of FC bandwidth on each SAN Fabric between Cisco FI and MDS switches. Link aggregation is not used on the links to Nimble array but 2 links from each SAN fabric connects to both controllers to provide 32G of FC bandwidth to the active controller. Four links from the SAN fabric, 2 from SAN Fabric A and 2 from SAN Fabric B, also connect to the backup controller so that both controllers have 32B FC access to the SAN fabric.

Cisco MDS switches are deployed with N-Port ID Virtualization (NPIV) enabled to support the virtualized environment running on Cisco UCS blade and rack servers. NPIV is necessary to provide isolation in virtualized environments where multiple virtual machines are running on a single server but a LUN needs to be presented to only one VM and not all VMs running on the server. Without NPIV, LUNs would be presented to the host and as a result, all VMs running on that host. To support NPIV on the Cisco UCS servers, the Cisco UCS Fabric Interconnects that connect the servers to the SAN fabric, are enabled for N-Port Virtualization (NPV) mode by configuring to operate in end-host mode (as opposed to FC switching mode). NPV enables Cisco FIs to proxy fabric login, registration and other messages from the servers to the SAN Fabric without being a part of the SAN fabric. This is important for keeping the limited number of Domain IDs that Fabric switches require to a minimum.

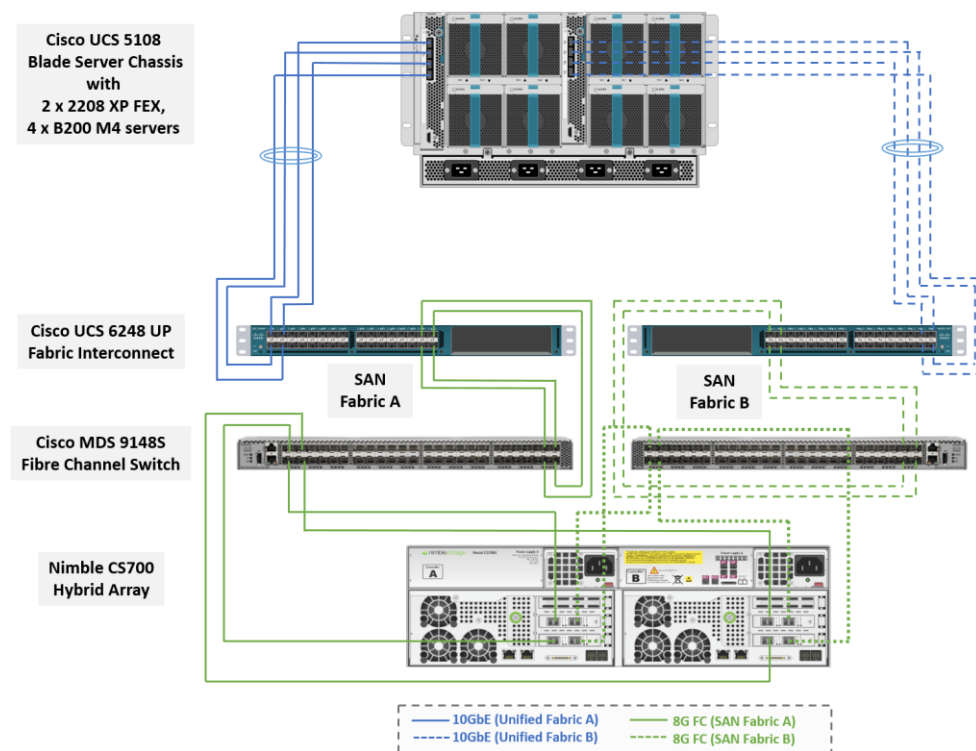
The design also uses the following best practices:

- Jumbo frames on unified fabric links between Cisco UCS and fabric interconnects
- QoS policy for traffic prioritization on the unified fabric
- Port-channels with multiple links are used in both the unified fabric and SAN network for higher aggregate bandwidth and redundancy.
- Zoning is single initiator (vHBA) to multiple targets.

Storage

The Nimble storage design is shown in Figure 13.

Figure 13 SmartStack Storage Design



SmartStack Block Storage Design

SmartStack design uses Nimble's CS700 hybrid array to provide block storage. A base configuration with 54.4TB of raw capacity (12 x 4TB HDDs, 4 x 1.6TB SSDs) was deployed in the array used for validation. Nimble's CS700 supports the addition of up to 6 expansion shelves or up to 4 CS-series (any model) arrays in a scale-out cluster to increase performance and capacity.

Each Nimble Storage controller supports up to 3 FC interface cards, each with dual 16G FC ports. This SmartStack design uses 8G fabric connectivity with two FC interface cards to provide 32G of FC bandwidth per controller. For additional FC bandwidth, a third FC card can be deployed on each controller but this interface is typically used for 10GbE connections to other arrays in a scale-out cluster for data replication traffic. The links between a pair of Cisco MDS and Fabric Interconnect switches are aggregated using 4x8G FC links to deliver 32G of bandwidth across the SAN fabric to each controller.

This SmartStack design uses FC SAN boot to boot the servers. The Service Profile used to configure and deploy Cisco UCS servers is configured to include a boot policy that points to the Nimble Storage array. The boot policy specifies a primary and secondary SAN path to the two controllers on the array where the boot volumes reside. A second boot policy is also configured but with the primary and secondary paths reversed from that of the first boot profile. The second boot policy is used to load balance SAN boot across different paths when multiple servers are booting. This is an optional aspect of the design that can be helpful in larger deployments for distributing load when multiple servers have to be simultaneously booted. Each server has a dedicated boot volume (40GB) on the Nimble storage array. Nimble Storage arrays provide an ACL at the initiator level to only allow connections from the appropriate Cisco UCS blade. During the initial SAN boot, the server attaches to all primary and secondary connections to both active and standby controller. This provides for normal boot operations even when a controller or primary path is offline. The hosts are

configured with the Nimble Connection Manager and Path Selection Policy which optimize MPIO settings. This will allow for proper FC path management and failover connectivity.

The following sections of this document provides more details on the connectivity and high availability aspects of this design.

Design Considerations

Management Connectivity

This SmartStack design uses a separate out-of-band management network to configure and manage compute, storage and network components in the solution. Management ports on each physical device (Cisco UCS FI, Nimble CS700 Array Controllers, Cisco Nexus and MDS switches) in the solution are connected to a separate, dedicated management switch.

Management access is also needed to ESXi hosts, vCenter VM and other management VMs but this is done in-band in this design. However, if out-of-band management to these components are required, the disjoint layer 2 feature available in Cisco UCS running end-host mode can be used. This would require additional uplink port(s) on the Cisco UCS FI to connect to the management switches. Additional out-of-band management vlans and vNICs will also be needed and associated with the uplink ports. The additional vNICs are necessary since a server vNIC can only be associated with a single uplink.

QoS and Jumbo Frames

Cisco UCS, Nexus, and MDS switches in the SmartStack solution provide QoS policies and features for handling congestion and traffic spikes that can occur in a SmartStack environment. SmartStack support different types of traffic (e.g. vMotion, FCOE) and the QoS capabilities in these components can alleviate and provide the priority that certain types of traffic require.

SmartStack design also uses jumbo frames with an MTU of 9000 Bytes across the LAN, SAN and Unified Fabric links. Jumbo frames increase the throughput between devices by enabling larger sized frames to be sent and received on the wire while reducing the CPU resources necessary to process them. Jumbo frames were enabled during validation on the LAN network links in the Nexus switching layer, the SAN MDS fabric and on the Unified Fabric links using Cisco UCS QoS system classes.

Fabric Extender Attached Design

For larger scale deployments, Cisco UCS C-series can be connected using standalone Fabric Extenders, namely the 1RU FEX 2232PP. Functionally, the standalone FEX is identical to the Cisco UCS 2204 and 2208 IOM modules that are deployed on the Cisco UCS 5108 blade server chassis for connecting to Cisco UCS Fabric Extenders. Similarly, the Cisco VIC on each Cisco UCS C-series server connect to both Fabric Interconnects using two Cisco FEX 2232PPs. The FEX and Fabric Interconnects form port channels automatically based on the chassis discovery policy providing a link resiliency to the Cisco UCS C-series server. This is identical to the behavior of the IOM to Fabric Interconnect connectivity. Logically, the virtual circuits formed within the Cisco UCS domain are consistent between B and C series deployment models and the virtual constructs formed at the vSphere are unaware of the platform in use.

Cisco UCS Server – vSphere Configuration

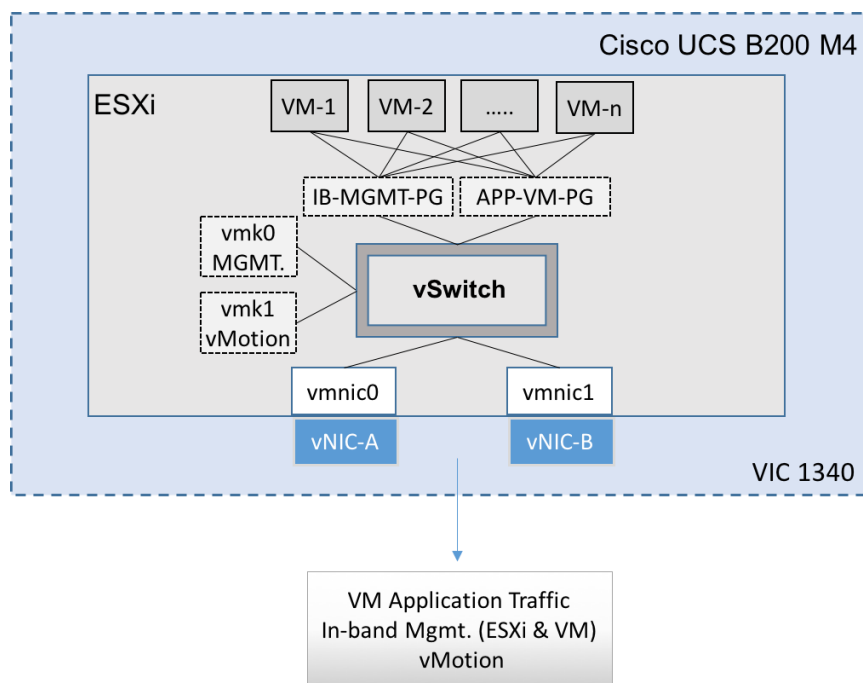
Cisco UCS B200M4 blade servers with Cisco 1340 VIC and Cisco UCS C220M4 rack servers with Cisco 1227 VIC running vSphere 6.0 U1 were validated in this SmartStack design. Cisco UCS servers were

assigned to a VMware High Availability(HA) cluster to mitigate against host failures. Two VMware HA clusters were used in validation – one for infrastructure management and services (e.g. VMware vCenter) and one for applications that users access. The Cisco VIC on each server presents multiple vPCIe devices to ESXi. vSphere identifies these virtual adapters as vmnics. In this SmartStack design, the following virtual adapters (vNICs) were used with –A connected to unified fabric A and –B to unified fabric B resulting in each ESXi node being dual homed to the external network.

- Two vNICs (vNIC-A, vNIC-B) for application VM, in-band management and vMotion traffic

The connectivity within each ESXi server and the vNIC presented to ESXi are shown in Figure 14.

Figure 14 Cisco UCS Server Networking - vSwitch



The SmartStack architecture uses two port groups (IB-MGMT-PG) for in-band management of the VMs and APP-VM-PG for application traffic. The design also used two VMkernel NICs (vmk), each with its own port group for host level functionality:

- vmk0 - ESXi management
- vmk1 - vMotion interface

The ESXi management interface is for host to vCenter connectivity, direct access to ESXi shell and VMware cluster communication. The vMotion interfaces are private subnets supporting data access and VM migration across the SmartStack infrastructure.

Cisco UCS Server – Virtual Switching using Nexus 1000V

A Cisco Nexus 1000V virtual distributed switch is used to provide connectivity between virtual machines and host connectivity to external networks. Cisco Nexus 1000v is an optional component of the SmartStack solution. Cisco Nexus 1000v is fully integrated into the VMware virtual infrastructure, including VMware vCenter and vCloud Director and extends the network edge to the hypervisor and virtual machines. Cisco

Nexus 1000V is compatible with any upstream physical access layer switch that is compliant with Ethernet standards including Cisco Nexus switches and switches from other network vendors.

The Cisco Nexus 1000v comprises of the following components and operationally emulates a modular switch where:

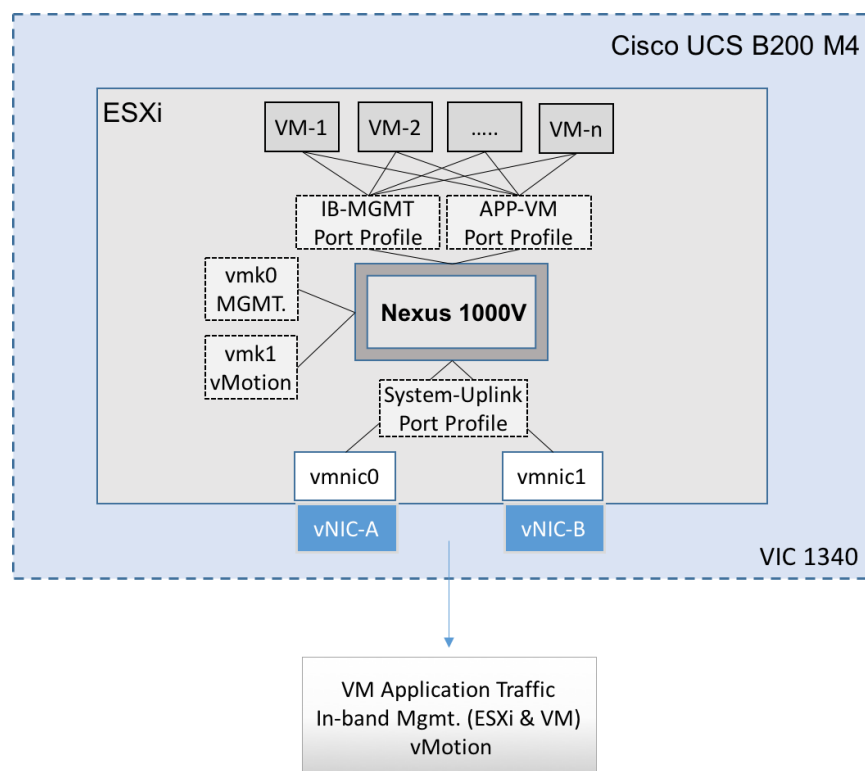
- Virtual Supervisor Module (VSM) – is the control and management plane of the virtual switch. VSM is deployed as an external virtual machine and runs NX-OS to manage multiple Virtual Ethernet Modules as one logical modular switch.
- Cisco Virtual Ethernet Module (VEM) – virtual line card or module within the virtual switch that VMs connect into. VEM is embedded in each VMware vSphere host and replaces the VMware Virtual Switch (vSwitch) functionality.
- Operating inside the VMware ESXi hypervisor, Cisco Nexus 1000V VEM uses the VMware vNetwork Distributed Switch (vDS) API, jointly developed by Cisco and VMware to provide policy-based VM connectivity, Layer 2 switching and advanced networking functions. The tight integration makes Cisco Nexus 1000V fully aware of key virtualization features such as VMware vMotion and Distributed Resource Scheduler (DRS). VEM provides switching functionality based on the configuration information it receives from the VSM. In the event of a communication loss with the VSM, VEM continues to forward traffic based on last known configuration or Nonstop Forwarding (NSF). VEM therefore provides reliability and advanced switching functionality.

Cisco Nexus 1000V VSM controls multiple VEMs as one logical modular switch with the VEM running as software on each server representing the line cards on a switch. VSM is integrated into VMware vCenter server so that the datacenter administrator can manage and monitor the network configuration on the Cisco Nexus 1000V switches. Configuration done through the VSM are automatically propagated to all VEMs managed by a given VSM. For high availability, VSMs can be redundantly deployed providing rapid, stateful failover as the VSMs. VSMs also provide port-profiles as a mechanism for grouping ports by category that enables the solution to scale to a high number of ports. VSM can also be accessed and managed through CLI, SNMP, XML API and CiscoWorks LAN management Solution.

Figure 15 shows the virtual networking within ESXi on a single Cisco UCS server. The Cisco Nexus 1000V VEM running on the ESXi node is registered to a VSM running on the infrastructure cluster and integrated into VMware vCenter. The Cisco VIC on each server presents multiple vPCIe devices to ESXi that are identified as vmnics. This SmartStack design uses two virtual adapters (vNIC-A, vNIC-B) with vNIC-A connected to unified fabric A and vNIC-B connected to unified fabric B. Host traffic (application VMs, in-band management, vMotion) are distributed across these vNICs. The ESXi vmnics are presented as Ethernet interfaces on Cisco Nexus 1000V. Cisco Nexus 1000V provides port profiles to address the dynamic nature **of server virtualization from the network's** perspective. Port profiles, defined on the VSM, serve as templates that define the network, security and service level policies for groups of virtual machines. Cisco Nexus 1000v aggregates the Ethernet uplinks into a single port channel named the "System-Uplink" port profile for fault tolerance and improved throughput. The port profiles can then be applied to individual virtual machine Ethernet interfaces through VMware vCenter.

Cisco Nexus 1000v provides link failure detection, disabling Cisco UCS Fabric Failover within the vNIC template is recommended.

Figure 15 Cisco UCS Server Networking – Nexus 1000V (Optional)



The SmartStack architecture uses two port profiles (IB-MGMT) for in-band management of the VMs and APP-VM for the application traffic used in validation. The design also uses two VMkernel NICs (vmk), each with its own port profile for host level functionality:

- vmk0 - ESXi management
- vmk1 - vMotion interface
- The ESXi management interface is for host to vCenter connectivity, direct access to ESXi shell and VMware cluster communication. The vMotion interfaces are private subnets supporting data access and VM migration across the SmartStack infrastructure.

The Cisco Nexus 1000v also supports Cisco's MQC to assist in uniform operation and enforcement of QoS policies across the infrastructure. The Cisco Nexus 1000v supports marking at the edge and policing traffic from VM-to-VM.

For more information about "Best Practices in Deploying Cisco Nexus 1000V Series Switches on Cisco UCS B and C Series Cisco UCS Manager Servers" refer to:

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c11-558242.html

Cisco Nexus 9000 Series – vPC Best Practices

SmartStack incorporates the following Cisco Nexus 9000 design best practices and recommendations.

vPC Peer Keepalive Link Considerations

- Recommendations for a vPC peer keepalive link is dedicated 1Gbps Layer3 link, followed by out-of-band management interface (mgmt0) and lastly, routing the keepalive link over an existing Layer3 infrastructure between the existing vPC peers. vPC peer keepalive link should not be routed over a vPC peer-link. The out-of-band management network was used as the vPC peer keepalive link in this SmartStack design.

vPC Peer Link Considerations

- Only vPC vlans are allowed on the vPC peer-links. For deployments that require non-VPC vlan traffic to be exchanged between vPC peer switches, deploy a separate Layer 2 link for this traffic.
- Only required vlans are allowed on the vPC peer links and member ports – prune all others to minimize internal resource consumption.
- Ports from different line cards should be used to provide redundancy for vPC peer links. It was not possible to do this on the fixed module Cisco Nexus 9372PX switches used in this SmartStack design.

vPC General Considerations

- vPC peer switches deployed using same bridge-id and spanning tree VLAN priority by configuring the peer-switch command on both vPC peer switches. This feature improves convergence and allows peer switches to appear as a single spanning-tree root in the Layer 2 topology.
- vPC role priority specified on both Cisco Nexus peer switches. vPC role priority determines which switch will be primary and which one will be secondary. The device with the lower value will become the primary. By default, this value is 32677. Cisco recommends that the default be changed on both switches. Primary vPC devices are responsible for BPDU and ARP processing. Secondary vPC devices are responsible for shutting down member ports and vlan interfaces when peer-links fail.
- vPC convergence time of 30s (default) was used to give routing protocol enough time to converge post-reboot. The default value can be changed using delay-restore <1-3600> and delay-restore interface-vlan <1-3600> commands. If used, this value should be changed globally on both peer switches to meet the needs of your deployment.
- vPC peer switches enabled as peer-gateways using peer-gateway command on both devices. Peer-gateway allows a vPC switch to act as the active gateway for packets that are addressed to the router MAC address of the vPC peer allowing vPC peers to forward traffic.
- vPC auto-recovery enabled to provide a backup mechanism in the event of a vPC peer-link failure due to vPC primary peer device failure or if both switches reload but only one comes back up. This feature allows the one peer to assume the other is not functional and restore the vPC after a default delay of 240s. This needs to be enabled on both switches. The time to wait before a peer restores the vPC can be changed using the command: auto-recovery reload-delay <240-3600>.
- Cisco NX-OS can synchronize ARP tables between vPC peers using the vPC peer links. This is done using a reliable transport mechanism that the Cisco Fabric Services over Ethernet (CFS over E) protocol provides. For faster convergence of address tables between vPC peers, ip arp synchronize command was enabled on both peer devices in this SmartStack design.

vPC Member Link Considerations

- LACP used for port channels in the vPC. LACP should be used when possible for graceful failover and protection from misconfigurations
- LACP mode active-active used on both sides of the port channels in the vPC. LACP active-active is recommended, followed by active-passive mode and manual or static bundling if the access device does not support LACP. Port-channel in mode active-active is preferred as it initiates more quickly than port-channel in mode active-passive.
- LACP graceful-convergence disabled on port-channels going to Cisco UCS FI. LACP graceful-convergence is ON by default and should be enabled when the downstream access switch is a Cisco Nexus device and disabled if it is not.
- Only required vlans are allowed on the vPC peer links and member ports – prune all others to minimize internal resource consumption.
- Source-destination IP, L4 port and VLAN are used load-balancing hashing algorithm for port-channels. This improves fair usage of all member ports forming the port-channel. The default hashing algorithm is source-destination IP and L4 port.

vPC Spanning Tree Considerations:

- Bridge Assurance enabled on vPC peer links by specifying spanning-tree port type network. Bridge Assurance should be disabled on vPC member ports.
- Spanning port type specified as edge or edge trunk on host facing interfaces connecting to Cisco UCS FI.
- BPDU Guard feature enabled on host-facing interfaces connecting to Cisco UCS FI. This was done by globally enabling it on all edge ports using the spanning-tree port type edge bpduguard default command.
- BPDU Filtering feature enabled on host-facing interfaces connecting to Cisco UCS FI. This was done by globally enabling it on all edge ports using the spanning-tree port type edge bpdufilter default command.
- Loop Guard was disabled (default setting) in this design. If necessary, they can be enabled globally using spanning-tree loopguard default or at the interface level using spanning-tree guard loop.
- Root Guard enabled on vpc member ports connected to access devices to make sure that vPC peer switches remain the spanning tree root – using interface level command spanning-tree guard root

Other Considerations

- Unidirectional Link Detection (UDLD) was enabled globally using feature udld and on vPC peer links and member ports to Cisco UCS FI.
- HSRP specific
 - Interface vlans should be defined as passive interfaces to avoid routing peer information
 - Disable IP redirection on HSRP interface vlans
 - Use default timer for HSRP/VRRP

- If the SmartStack design outlined in this CVD is connected to additional aggregation/core layer Cisco Nexus switches in a two-tiered design for scalability or other expansion purposes, the following guidelines should be followed.
 - In a two-tiered data center design using Cisco Nexus switches, vPCs can also be used between the Cisco Nexus switches in each tier using a double-sided vPC topology. In such a design, the vPC domain identifiers must be different as this information is exchanged through LACP protocol and using the same vPC domain identifiers will generate continuous flaps on vPC between the different Cisco Nexus network layers.
 - If modular Cisco Nexus switches are used, redundancy should be provided by using ports from different line cards.
 - Deploy dedicated Layer 3 link(s) for exchanging routing information between peer switches in a two-tiered design or other topologies where Layer 3 is used between the tiers. The vPC peer-link should not be used.

Last but not least, review the criteria for vPC Type-1 and Type-2 consistency checks in the link provided below to avoid issues in your deployment.

- vPC Design Guide:
http://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf
- Nexus 9000 NX-OS Release 6.x Configuration Guide:
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/interfaces/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Interfaces_Configuration_Guide.html

High Availability

SmartStack platform was designed for maximum availability of the complete infrastructure (compute, network, storage, virtualization) with no single points of failure.

Compute and Virtualization

- Cisco UCS system provides redundancy at the component and link level and end-to-end path redundancy to storage array and LAN network.
- Cisco UCS 5108 blade server platform is highly redundant with redundant power supplies, fans and fabric extender modules. C-series rack mount servers also have redundant power supplies and fans. The rack mount servers are directly connected to the upstream fabric interconnects in this SmartStack design.
- Each fabric extender on the Cisco UCS 5108 blade server is deployed with 4x10GbE links to the unified fabric. The links are part of a port channel to ensure rerouting of flows to other links in the event of a link failure.
- Each server is deployed using vNICs and vHBAs that provide redundant connectivity to the unified fabric. NIC failover is enabled between Cisco UCS Fabric Interconnects using Cisco UCS manager. This is done for all management and virtual machine vNICs.

Solution Design

- VMware vCenter is used to deploy VMware HA clusters to allow VMs to failover in the event of a server failure. VMware vMotion and VMware HA are enabled to auto restart VMs after a failure. Host Monitoring is enabled to monitor heartbeats of all ESXi hosts in the cluster for faster detection. Admission Control is also enabled on the blade servers to ensure the cluster has enough resources to accommodate a single host failure.
- VMware vSphere hosts use SAN multipathing to access LUNs on the Nimble array. If any component (NIC, HBA, FEX, FI, MDS, Nimble controller, cables) along a path fails, all storage traffic will reroute to an alternate path. When both paths are active, traffic is load balanced.

Network

- Link aggregation using port channels and virtual port channels are used throughout the SmartStack design for higher bandwidth and availability.
- Port channels are used on unified fabric links between fabric extender and fabric interconnects. Virtual port channels are used between FIs and Nexus switches. VPCs provide higher availability than port channels as it can continue to forward traffic even if one of the Nexus switches fail because VPCs distribute member links of port-channel across different Nexus switches.
- Pair of Cisco Nexus 9000 series switches are used in the datacenter LAN fabric to provide redundancy in the event of a switch failure.
- Pair of Cisco MDS switches are used in the SAN fabric to provide redundancy in the event of a switch failure.
- MDS and Nexus switches are highly redundant with redundant power supplies, fans and have out-of-band management access.
- The two MDS switches form two separate fabrics and provide two distinct physical paths to storage for redundancy. FI-A to MDS-A to Nimble array is SAN Fabric A and FI-B to MDS-B to Nimble array is SAN Fabric B. Dual VSANs are used across these fabrics with vSAN-A on Fabric A and vSAN-B on Fabric B. The dual vSANs represent two redundant paths to the storage array with traffic load balanced across both vSANs when there is no failure.

Storage

- The Nimble CS700 array has redundant storage controllers which allow for an active / standby configuration.
- The CS700 has redundant power supplies with diverse cabling and data paths to each controller.
- Each Nimble storage controller is redundantly connected to the SAN fabric. Each controller is connected using 4x 8Gb FC links to upstream MDS switches with 2x8Gb links going to MDS-A switch and 2x8Gb links going to MDS-B switch in the SAN fabric. This will allow 32GB network bandwidth for each controller.
- FC target connections are configured in a Dual fabric / dual vSAN switch fabric. This configuration is used across the SAN fabric and unified fabric for redundant connectivity to storage.
- Each Service Profile has a boot profile with redundant paths to primary and secondary FC targets on the Nimble Storage array.

- All VMware datastore volumes utilizes Nimble PSP_Directed for proper path failover and load distribution.

Scalability

For higher performance and capacity, SmartStack solutions can scale up by adding compute, network, storage or management subsystems individually or scale out with consistent and predictable performance by deploying additional units of the complete SmartStack solution.

Management

- Cisco UCS Manager residing on a clustered pair of Cisco UCS Fabric Interconnects that makes up a UCS domain can manage up to 160 servers (8 servers per chassis x 20 chassis) in a single UCS domain.
- Cisco UCS Central can be deployed to provide a single pane for managing multiple Cisco UCS domains – for up to 10,000 servers. Cisco UCS Central complements Cisco UCS Manager to provide centralized inventory, faults, and logs, global policies and pools, firmware management, global visibility and flexible administrative controls. Cisco UCS Central is a manager of managers that interfaces with Cisco UCS Manager in each domain to manage multiple, globally distributed Cisco UCS domains.

Storage

Scale-to-Fit

With Nimble Storage's CS700 [Predictive Flash platform](#), it is easy to accommodate application growth by scaling performance, capacity, or both—efficiently and non-disruptively. With Nimble Storage scale-to-fit, organizations can:

- Flexibly scale flash to accommodate a wide variety of application working sets. With the addition of an All Flash Shelf, the CS700 can support up to 38 TB in SSD.
- Scale capacity by adding additional HDDs or expansion shelves. The CS700 platform can support an effective capacity of up to 892 TB with expansion shelves.
- Scale up performance by upgrading compute for greater throughput and IOPS
- Scale capacity and performance together by clustering any combination of Nimble Storage arrays – **see next section for Nimble's scale-out capabilities**

Scale-Out

Nimble Storage's Predictive Flash platform features a scale-out architecture that makes it easy to scale capacity and performance beyond the physical limitations of a single array. With Nimble Storage scale-out, organizations can:

- Group up to 4 Nimble arrays (any model) for higher scalability and performance
- One management console to administer all storage hardware resources in the group as a single entity
- Dynamic load balancing across arrays in a group to eliminate performance hot spots

Solution Design

- Multi-array data striping, enabling any application to fully leverage the collective hardware resources of the scale-out group
- Flexible configuration of any combination of Nimble Storage arrays in the group, maximizing storage ROI
- Seamless reconfiguration and hardware refreshes, without downtime

Solution Validation

Configuration Topology for Scalable XenDesktop Mixed Workload Desktop Virtualization Solution on Cisco Unified Computing System and Nimble Storage

Cisco Unified Computing System Configuration

This section talks about the Cisco UCS configuration that was done as part of the infrastructure build-out. The racking, power and installation of the chassis are described in the install guide (see www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html) and it is beyond the scope of this document. For more information about each step, refer to the following document, Cisco UCS Manager Configuration Guides – GUI and Command Line Interface (CLI) [Cisco UCS Manager - Configuration Guides - Cisco](#)

Base Cisco UCS System Configuration

To configure the Cisco Unified Computing System, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS 6248 Fabric Interconnect cluster address.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.

Cisco UCS Manager Software to Version 3.1.1e

The Cisco UCS chassis comes with UCSM 3.1.1e release. This document assumes the use of Cisco UCS Manager Software version 3.1.1e. To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 Fabric Interconnect software to a higher version of the firmware,) refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Add a Block of IP Addresses for Out-of-Band KVM Access

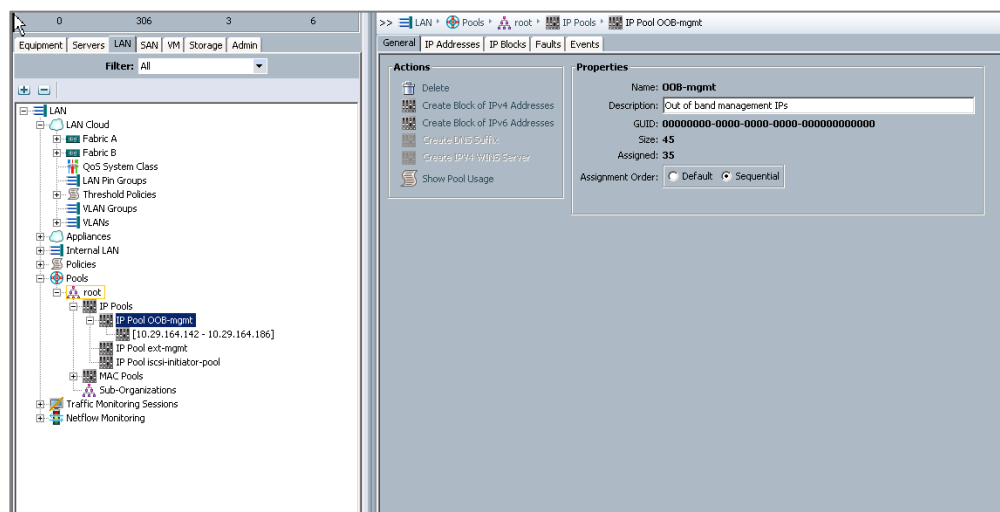
To create a block of IP addresses for server keyboard, video, mouse (KVM) access in the Cisco UCS environment, complete the following steps:



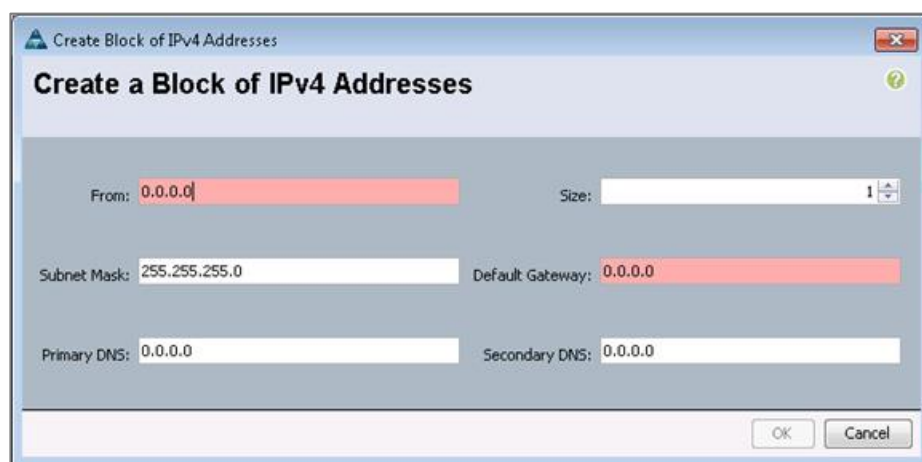
This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select Pools > root > IP Pools > IP Pool ext-mgmt.

Solution Validation



3. In the Actions pane, select Create Block of IP Addresses.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.



5. Click OK to create the IP block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Admin tab.
2. Select All > Timezone Management.
3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click Save Changes, and then click OK.
5. Click Add NTP Server.

Solution Validation

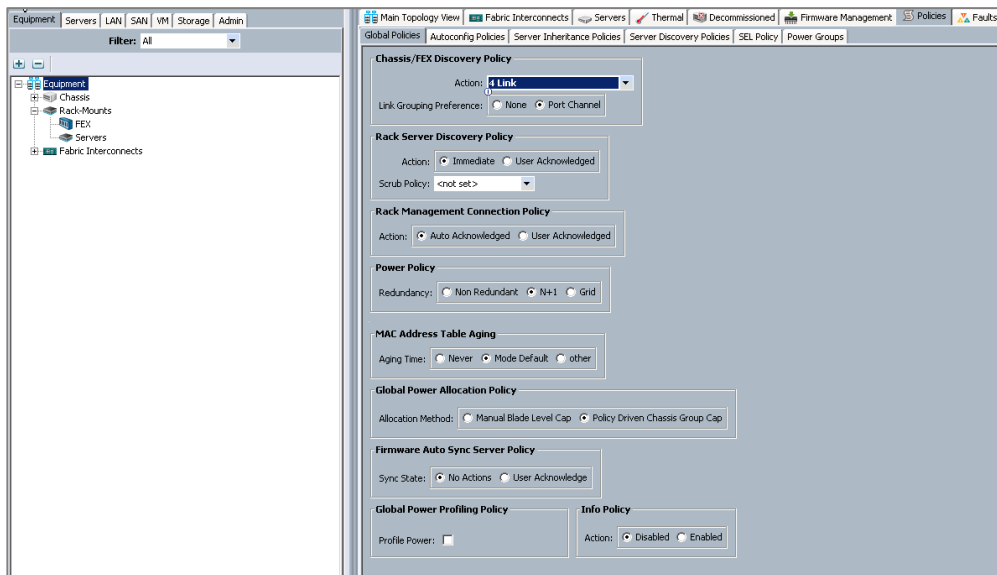
6. Enter <<var_global_ntp_server_ip>> and click OK.
7. Click OK.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and Cisco UCS M-Series chassis.

To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment node and select Equipment in the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to 4-link.
4. Set the Link Grouping Preference to Port Channel.



5. Click Save Changes.
6. Click OK.

Enable Server Uplink Ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
3. Expand Ethernet Ports.

4. Select ports 1 through 16 that are connected to the Cisco IO Modules of the two B-Series 5108 Chassis and the two M-Series 4308 Chassis, right-click them, and select Configure as Uplink Port.
5. Click Yes to confirm uplink ports and click OK.
6. In the left pane, navigate to Fabric Interconnect A. In the right pane, navigate to the Physical Ports tab > Ethernet Ports tab. Confirm that ports have been configured correctly in the If Role column.

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status
Port 1	1	1	00:2A:6A:03:DF:68	Server	Physical	Up
Port 2	1	2	00:2A:6A:03:DF:69	Server	Physical	Up
Port 3	1	3	00:2A:6A:03:DF:6A	Server	Physical	Up
Port 4	1	4	00:2A:6A:03:DF:6B	Server	Physical	Up
Port 5	1	5	00:2A:6A:03:DF:6C	Server	Physical	Up
Port 6	1	6	00:2A:6A:03:DF:6D	Server	Physical	Up
Port 7	1	7	00:2A:6A:03:DF:6E	Server	Physical	Up
Port 8	1	8	00:2A:6A:03:DF:6F	Server	Physical	Up
Port 9	1	9	00:2A:6A:03:DF:70	Server	Physical	Up
Port 10	1	10	00:2A:6A:03:DF:71	Server	Physical	Up
Port 11	1	11	00:2A:6A:03:DF:72	Server	Physical	Up
Port 12	1	12	00:2A:6A:03:DF:73	Server	Physical	Up
Port 13	1	13	00:2A:6A:03:DF:74	Server	Physical	Up
Port 14	1	14	00:2A:6A:03:DF:75	Server	Physical	Up
Port 15	1	15	00:2A:6A:03:DF:76	Server	Physical	Up
Port 16	1	16	00:2A:6A:03:DF:77	Server	Physical	Up
Port 17	1	17	00:2A:6A:03:DF:78	Network	Physical	Up
Port 18	1	18	00:2A:6A:03:DF:79	Network	Physical	Up
Port 19	1	19	00:2A:6A:03:DF:7A	Network	Physical	Up
Port 20	1	20	00:2A:6A:03:DF:7B	Network	Physical	Up

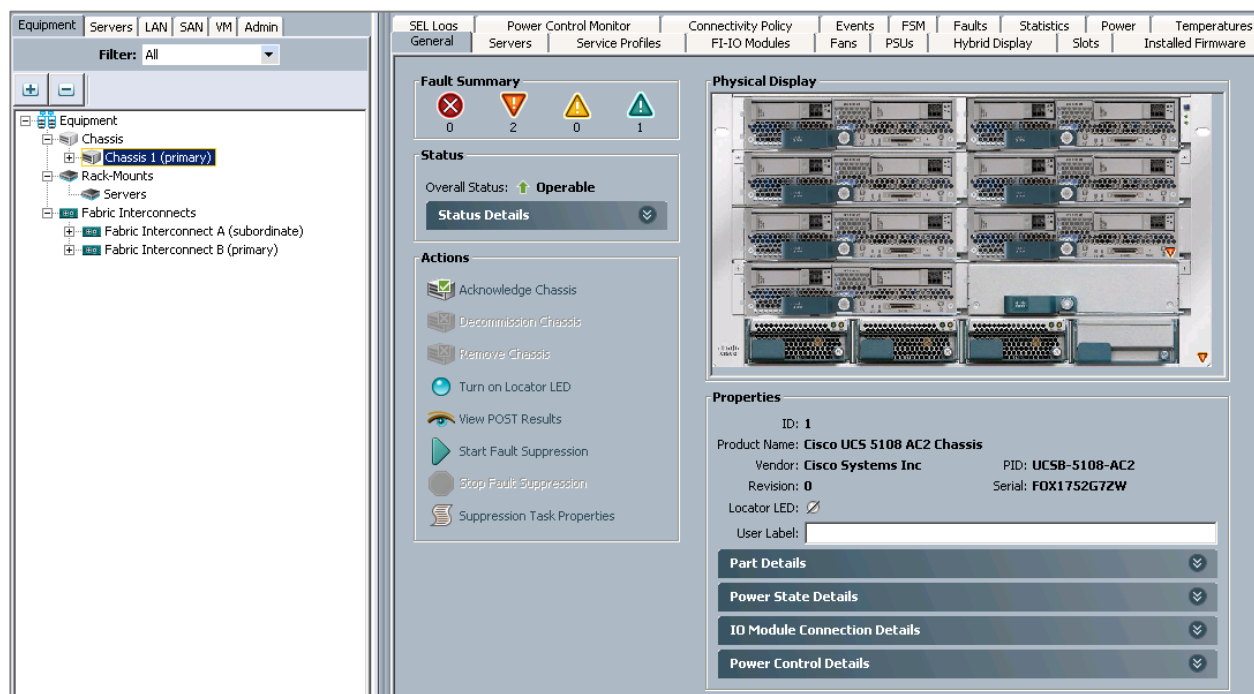
7. Repeat the above steps for Fabric Interconnect B.

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status
Port 1	1	1	00:2A:6A:09:84:A8	Server	Physical	Up
Port 2	1	2	00:2A:6A:09:84:A9	Server	Physical	Up
Port 3	1	3	00:2A:6A:09:84:AA	Server	Physical	Up
Port 4	1	4	00:2A:6A:09:84:AB	Server	Physical	Up
Port 5	1	5	00:2A:6A:09:84:AC	Server	Physical	Up
Port 6	1	6	00:2A:6A:09:84:AD	Server	Physical	Up
Port 7	1	7	00:2A:6A:09:84:AE	Server	Physical	Up
Port 8	1	8	00:2A:6A:09:84:AF	Server	Physical	Up
Port 9	1	9	00:2A:6A:09:84:B0	Server	Physical	Up
Port 10	1	10	00:2A:6A:09:84:B1	Server	Physical	Up
Port 11	1	11	00:2A:6A:09:84:B2	Server	Physical	Up
Port 12	1	12	00:2A:6A:09:84:B3	Server	Physical	Up
Port 13	1	13	00:2A:6A:09:84:B4	Server	Physical	Up
Port 14	1	14	00:2A:6A:09:84:B5	Server	Physical	Up
Port 15	1	15	00:2A:6A:09:84:B6	Server	Physical	Up
Port 16	1	16	00:2A:6A:09:84:B7	Server	Physical	Up
Port 17	1	17	00:2A:6A:09:84:B8	Network	Physical	Up
Port 18	1	18	00:2A:6A:09:84:B9	Network	Physical	Up
Port 19	1	19	00:2A:6A:09:84:BA	Network	Physical	Up
Port 20	1	20	00:2A:6A:09:84:BB	Network	Physical	Up

Acknowledge Cisco UCS Chassis

To acknowledge all Cisco UCS chassis, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the Equipment tab.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.

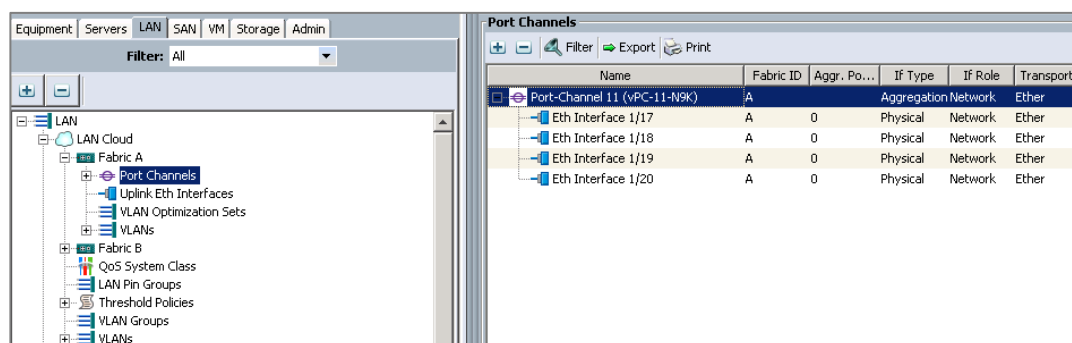


- Click Yes and then click OK to complete acknowledging the chassis.

Create Uplink Port Channels to Cisco Nexus 9372PX Switches

To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

- In Cisco UCS Manager, click the LAN tab in the navigation pane.
- In this procedure, two port channels are created: one from Fabric A to both Cisco Nexus 9372PX switches and one from Fabric B to both Cisco Nexus 9372PX switches.
- Under LAN > LAN Cloud, expand node Fabric A tree.

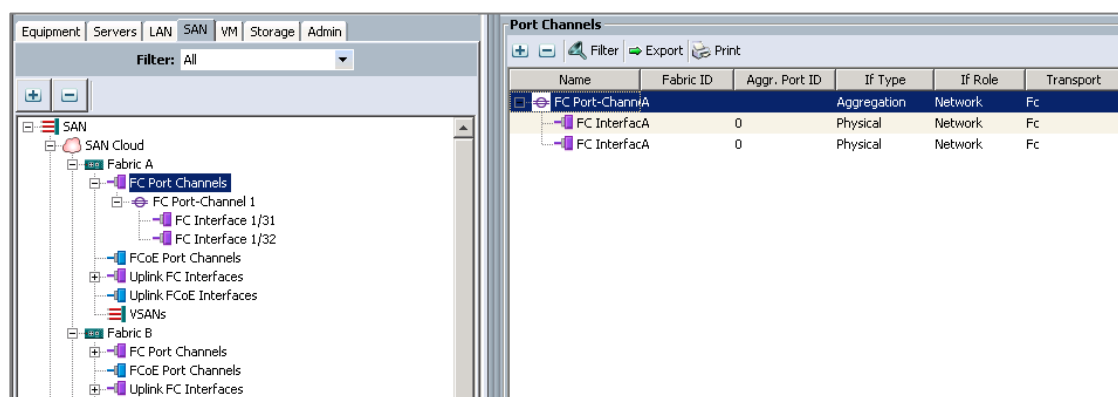


- Verify that four ports are configured as Ethernet network connectivity. If not, click the interface and click 'Configure as Uplink Port'.
- Repeat the above steps for Fabric B.

Create Uplink Port Channels to Cisco MDS 9148 Switches

To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. In this procedure, two port channels are created: One from Fabric A to Cisco MDS 9148 switch A and one from Fabric B to Cisco MDS 9148 switch B.
3. Under SAN > SAN Cloud, expand node Fabric A tree.



4. Repeat the above steps for Fabric B.

Create an Organization

Organizations are used to organize resources and restrict access to various groups within the IT organization, thereby enabling multi-tenancy of the compute resources.



Although this document does not assume the use of organizations this procedure provides instructions for creating one.

To configure an organization in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, from the New menu in the toolbar at the top of the window, select Create Organization.
2. Enter a name for the organization.
3. Optional: Enter a description for the organization.
4. Click OK.
5. Click OK in the confirmation message.

Create Resource Pools

This section details how to create the MAC address, iSCSI IQN, iSCSI IP, UUID suffix and server pools.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root.



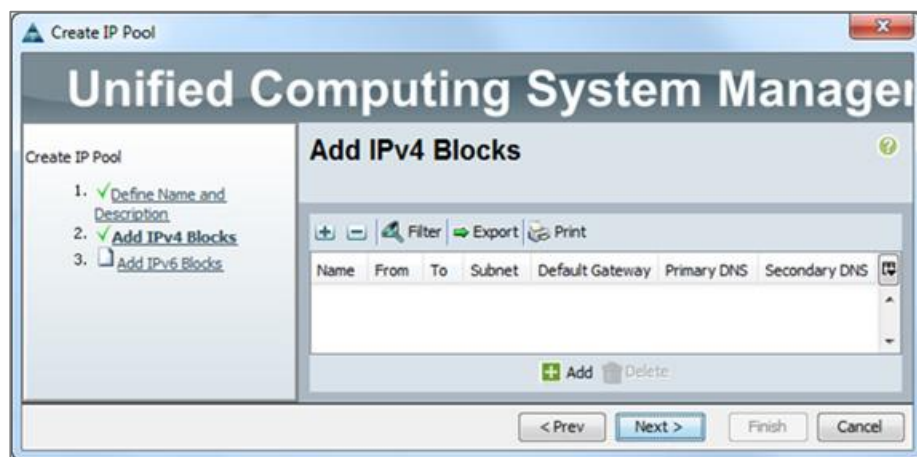
In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC_Pool_A as the name for MAC pool.
6. Optional: Enter a description for the MAC pool.

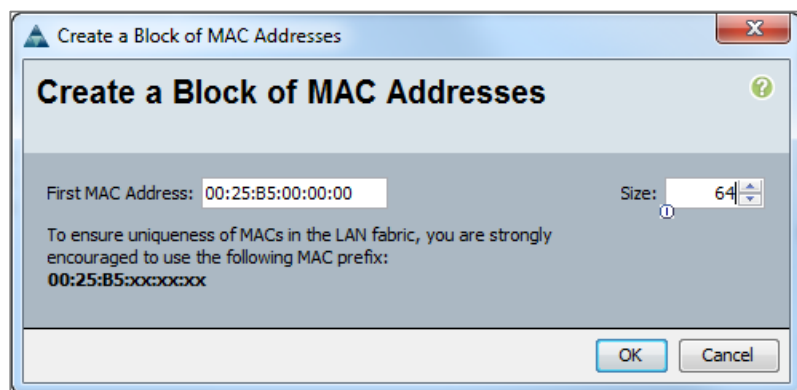


Keep the Assignment Order at Default.

7. Click Next.

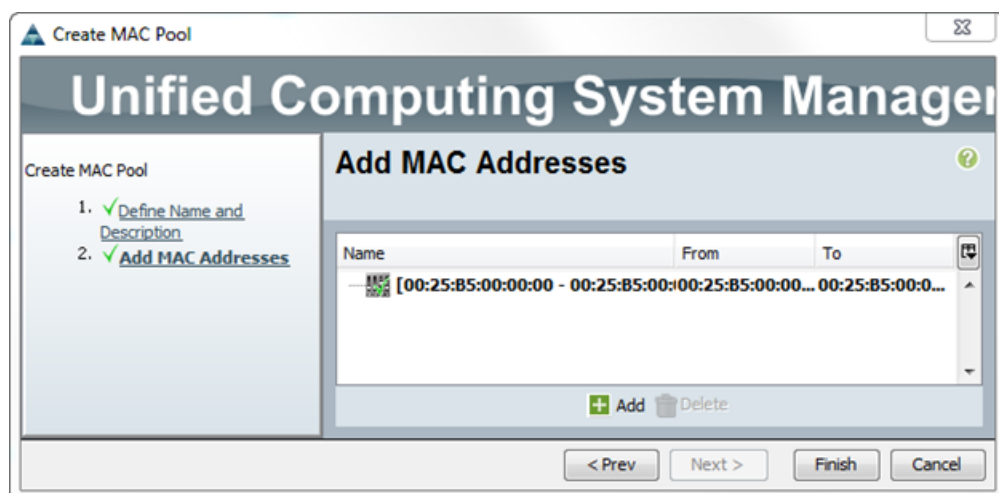


8. Click Add.
9. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.



10. Click OK.

11. Click Finish.



12. In the confirmation message, click OK.

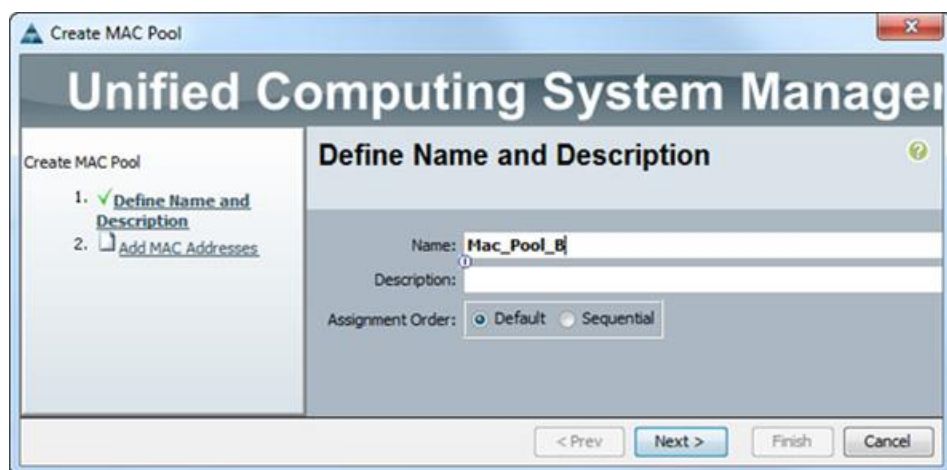
13. Right-click MAC Pools under the root organization.

14. Select Create MAC Pool to create the MAC address pool.

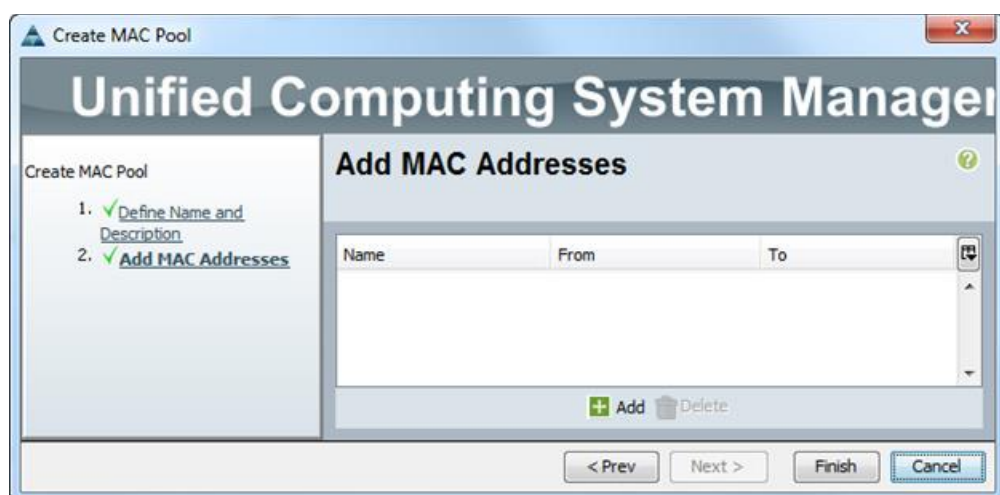
15. Enter MAC_Pool_B as the name for MAC pool.

16. Optional: Enter a description for the MAC pool.

17. Select Default for the Assignment Order.



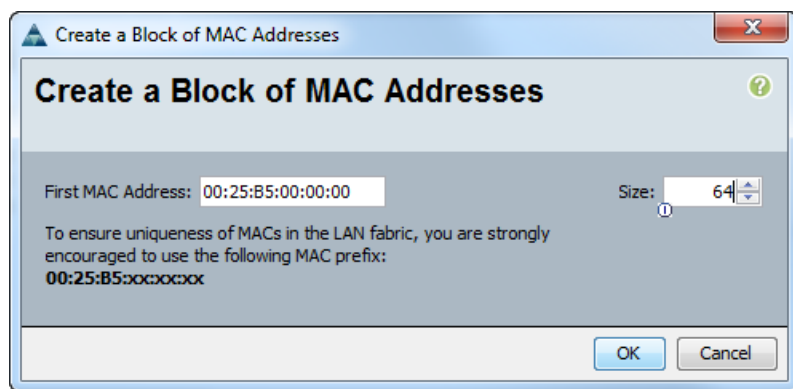
18. Click Next.



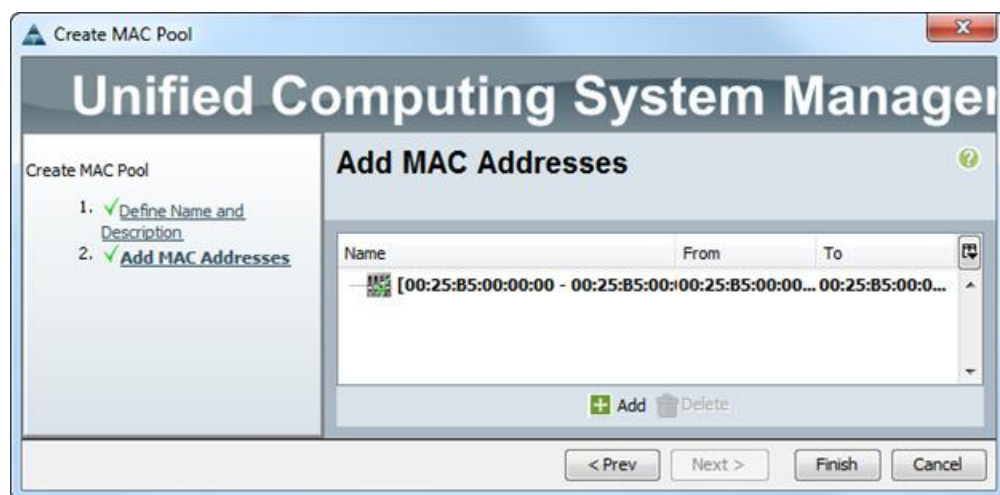
19. Click Add.

20. Specify a starting MAC address.

21. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.



22. Click OK.



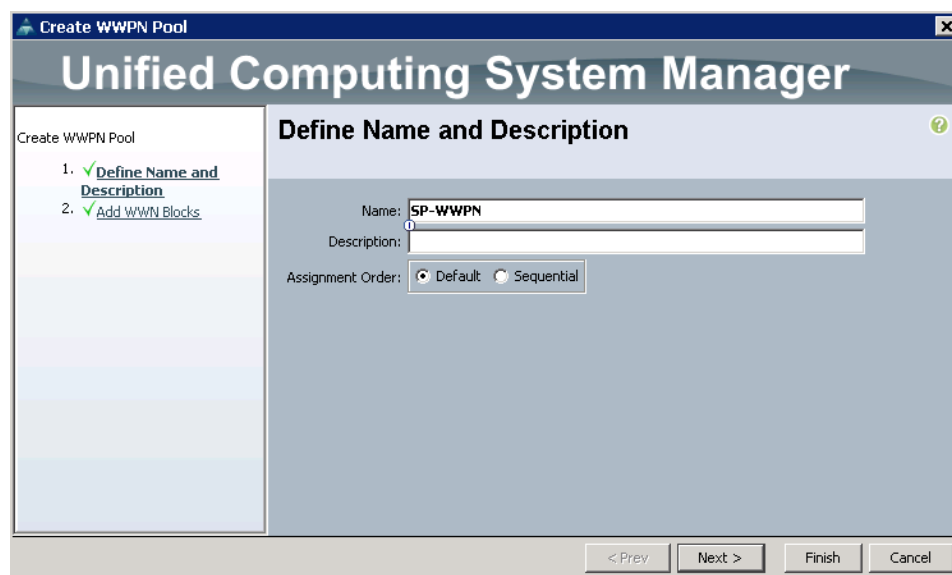
23. Click Finish.

24. In the confirmation message, click OK.

Create WWPN Pools

To configure the necessary WWPN pools for the Cisco UCS environment, complete the following steps:

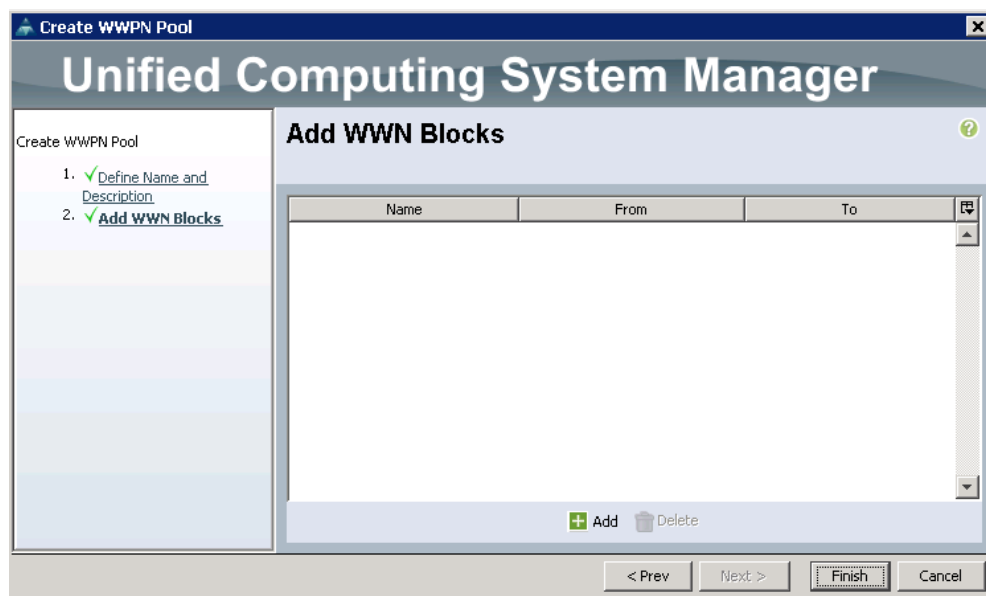
1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. Under WWPN Pools, right click WWPN Pools and select Create WWPN Pool.
4. Name it SP-WWPN.



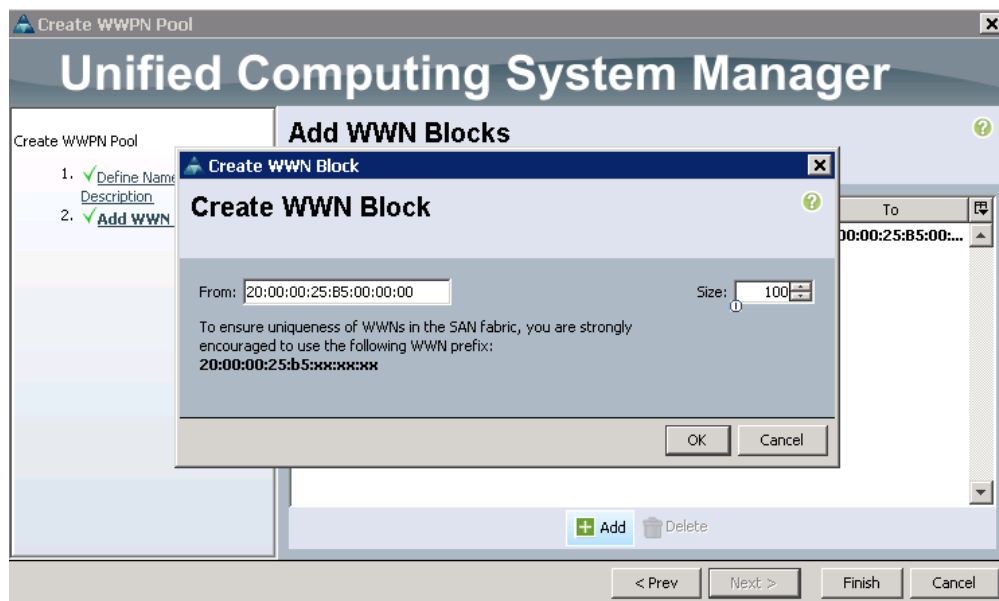
5. Assignment order can remain Default.

6. Click Next

7. Click Add to add block of Ports



8. Enter number of WWPNs. For this study we entered 100.



9. Click Finish.

Create WWNN Pools

To configure the necessary WWNN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > root.
3. Under WWNN Pools, right-click WWNN Pools and select Create WWNN Pool.

4. Name it SP-WWNN.

Create WWNN Pool

Unified Computing System Manager

Create WWNN Pool

1. **Define Name and Description**
2. Add WWN Blocks

Define Name and Description

Name:

Description:

Assignment Order: ☒ Default ☐ Sequential

< Prev Next > Finish Cancel

5. Assignment order can remain Default.
6. Click Next
7. Click Add to add block of Ports

Create WWNN Pool

Unified Computing System Manager

Create WWNN Pool

1. Define Name and Description
2. **Add WWN Blocks**

Create WWN Block

From: Size:

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:
20:00:00:25:b5:xx:xx:xx

OK Cancel

+ Add - Delete

< Prev Next > Finish Cancel

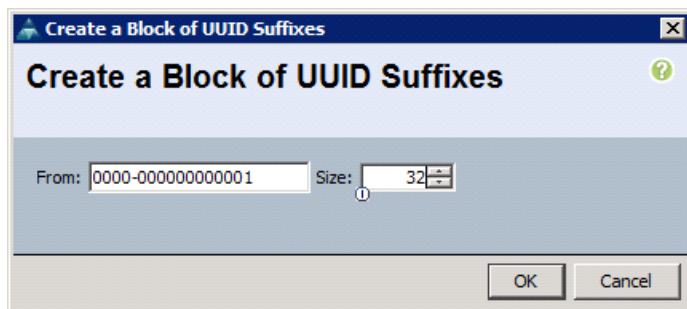
8. Enter number of WWNNs. For this study we did 100.
9. Click Finish.

Create UUID Suffix Pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

Solution Validation

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID_Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Click Next.
9. Click Add to add a block of UUIDs.
10. Keep the From field at the default setting.
11. Specify a size for the UUID block that is sufficient to support the available blade or server resources.



Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter Infra_Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.

Solution Validation

8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the Infra_Pool server pool.
9. Click Finish.
10. Click OK.

Create VLANs

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, complete the following steps:

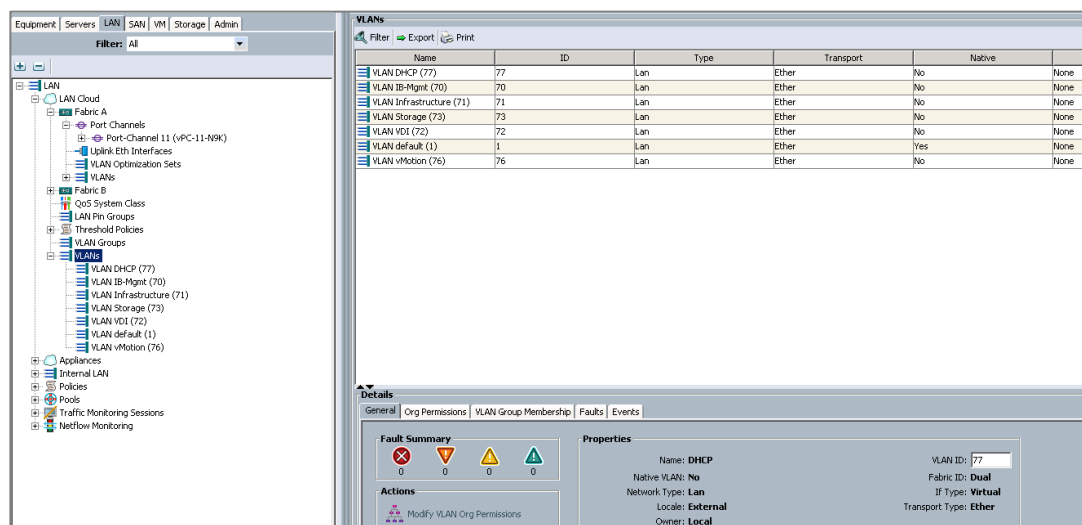
1. In Cisco UCS Manager, click the LAN tab in the navigation pane.



In this procedure, five VLANs are created.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter MGMT as the name of the VLAN to be used for in-band management traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter <<var_mgmt_id>> as the ID of the management VLAN.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.

10. Repeat the above steps to create all VLANs and configure the Default VLAN as native.



Name	ID	Type	Transport	Native	
VLAN DHCP (77)	77	Lan	Ether	No	None
VLAN IB-Mgmt (70)	70	Lan	Ether	No	None
VLAN Infrastructure (71)	71	Lan	Ether	No	None
VLAN Storage (73)	73	Lan	Ether	No	None
VLAN VDI (72)	72	Lan	Ether	No	None
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN vMotion (76)	76	Lan	Ether	No	None

Details	
General	Org Permissions VLAN Group Membership Faults Events
Fault Summary <div> 0 0 0 0 </div> Actions Modify VLAN Org Permissions	Properties Name: DHCP Native VLAN: No Network Type: Lan Locale: External Owner: Local VLAN ID: 77 Fabric ID: Dual IF Type: Virtual Transport Type: Ether

Create VSANs

To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.



In this procedure, two VSANs are created. When these VSANs are created, be sure to add them to the port-channel uplink created earlier.

2. Select SAN > SAN Cloud.
3. Under Fabric A, right-click VSANs.
4. Select Create VSANs.
5. Enter VSAN3 as the name of the VSAN to be used for in-band management traffic.
6. Select Fabric A for the scope of the VSAN.
7. Enter 3 as the ID of the VSAN.
8. Click OK, and then click OK again.

Create VSAN

Name:

FC Zoning Settings

FC Zoning: ☒ Disabled ☐ Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

☐ Common/Global ☒ Fabric A ☐ Fabric B ☐ Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.
Enter the VSAN ID that maps to this VSAN.

VSAN ID:

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.
Enter the VLAN ID that maps to this VSAN.

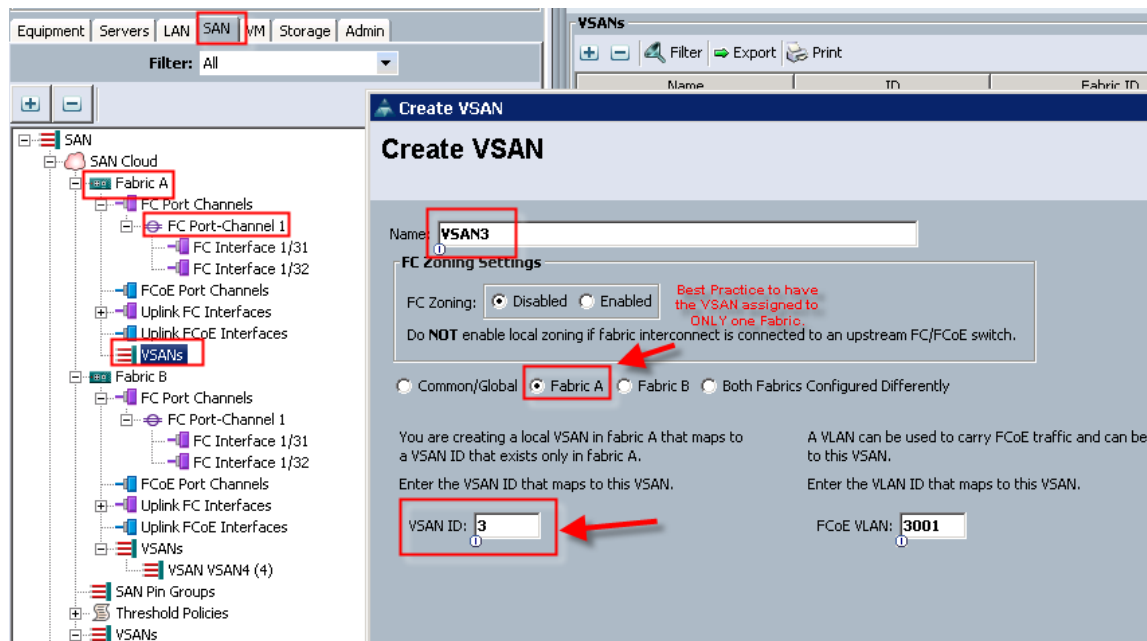
FCoE VLAN:

OK Cancel

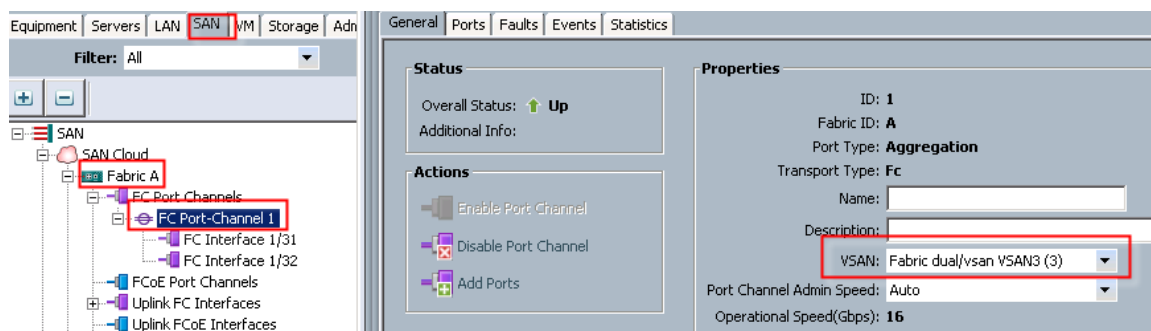
9. Repeat the above steps on Fabric B with VSAN 4 to create the VSANs necessary for this solution.
10. When done with both sides, go into the port-channel created earlier in the section 'Create uplinks for MDS 9148' and add the respective VSANs to their port channels. VSAN3 in this study is assigned to Fabric A and VSAN4 is assigned to Fabric B.



VSAN3 should only be on Fabric A and 4 on B.



11. Go to the Port-Channel for each Fabric and assign the VSAN appropriately.



Create Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter VM-Host as the name of the host firmware package.
6. Leave Simple selected.

7. Select the version 3.1.1e for the Blade Package.
8. Click OK to create the host firmware package.

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes in the bottom of the window.
6. Click OK.

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, complete the following steps:

Solution Validation

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable_CDP as the policy name.
6. For CDP, select the Enabled option.
7. Click OK to create the network control policy.

The screenshot shows the 'Create Network Control Policy' dialog box. The title bar reads 'Create Network Control Policy'. The main title is 'Create Network Control Policy'. The form contains the following fields and options:

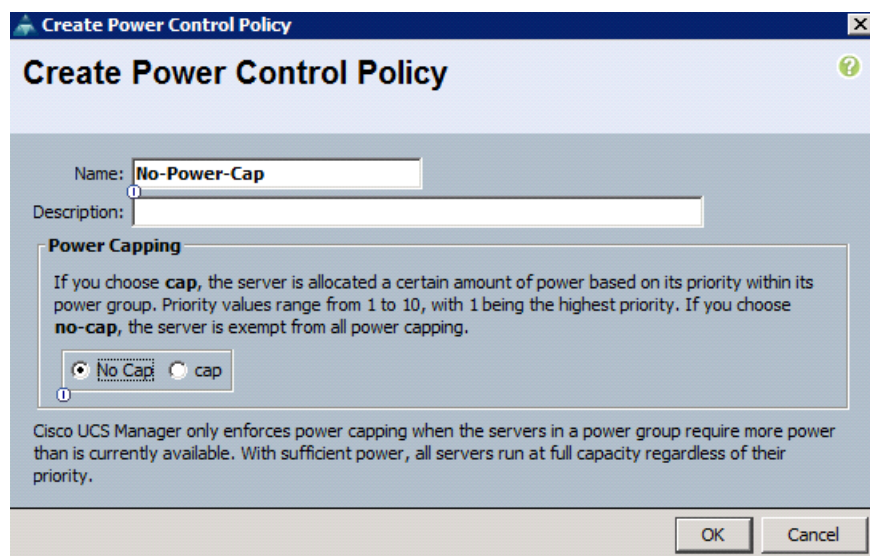
- Name:** A text box containing 'Enable_CDP'.
- CDP:** Two radio buttons: 'Disabled' and 'Enabled'. The 'Enabled' option is selected.
- MAC Register Mode:** Two radio buttons: 'Only Native Vlan' and 'All Host Vlan'. The 'Only Native Vlan' option is selected.
- Action on Uplink Fail:** Two radio buttons: 'Link Down' and 'Warning'. The 'Link Down' option is selected.
- MAC Security:** A section with a 'Forge:' label and two radio buttons: 'Allow' and 'Deny'. The 'Allow' option is selected.

At the bottom right, there are 'OK' and 'Cancel' buttons.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.




Cisco UCS System Configuration for Cisco UCS B-Series

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter B200-M4-BIOS as the BIOS policy name.
6. Repeat for M-Series changing the name to reflect M-Series.
7. In this paper we used GPUs, our VDI machines on the B200 servers. For the use of GPUs they need to be enabled in the BIOS policies. In this study we had two separate BIOS policies for B Series. Each one required a setting to enable GPU usage and they are as follows:
 - a. For B-Series use GPU BIOS Policy Settings:



8. Configure the remaining BIOS policies as follows and click Finish.

MainAdvancedBoot OptionsServer ManagementEvents

Actions

Delete

Show Policy Usage

Use Global

Properties

Name: B200-M4_BIOS

Description:

Owner: Local

Reboot on BIOS Settings Change: ☒

Quiet Boot: ☒ disabled ☐ enabled ☐ Platform Default

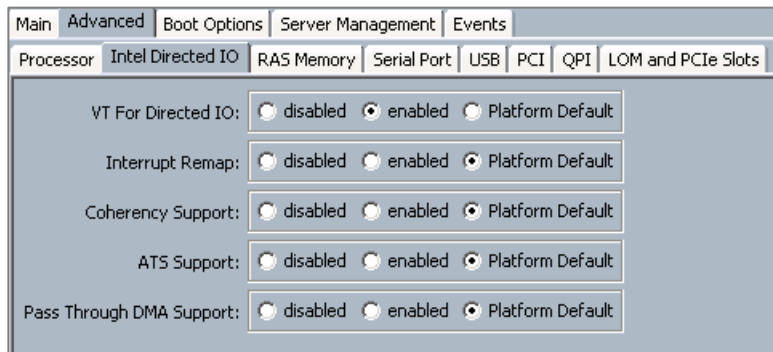
Post Error Pause: ☐ disabled ☐ enabled ☒ Platform Default

Resume Ac On Power Loss: ☐ stay-off ☐ last-state ☐ reset ☒ Platform Default

Front Panel Lockout: ☐ disabled ☐ enabled ☒ Platform Default

69

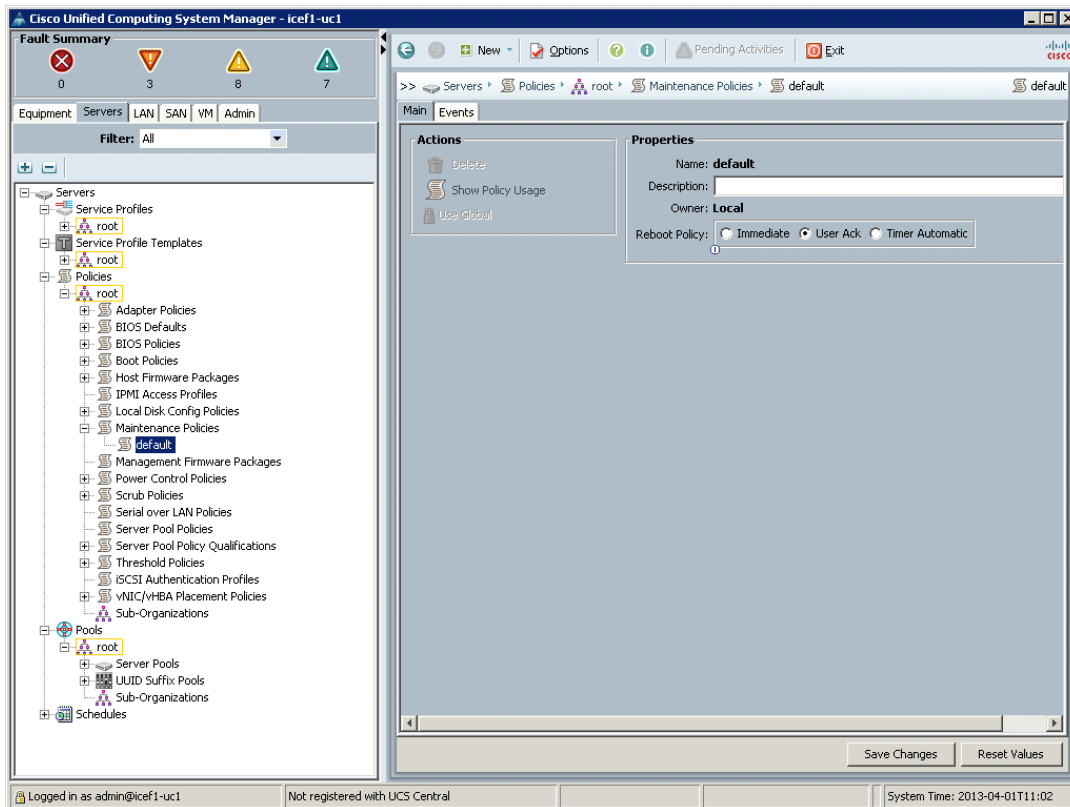
Main	Advanced	Boot Options	Server Management	Events			
Processor	Intel Directed IO	RAS Memory	Serial Port	USB	PCI	QPI	LOM and PCIe Slots
<p>Turbo Boost: <input type="radio"/> disabled <input checked="" type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Enhanced Intel Speedstep: <input type="radio"/> disabled <input checked="" type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Hyper Threading: <input type="radio"/> disabled <input checked="" type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Core Multi Processing: <input type="text" value="all"/></p> <p>Execute Disabled Bit: <input type="radio"/> disabled <input checked="" type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Virtualization Technology (VT): <input type="radio"/> disabled <input checked="" type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Hardware Pre-fetcher: <input type="radio"/> disabled <input type="radio"/> enabled <input checked="" type="radio"/> Platform Default</p> <p>Adjacent Cache Line Pre-fetcher: <input type="radio"/> disabled <input type="radio"/> enabled <input checked="" type="radio"/> Platform Default</p> <p>DCU Streamer Pre-fetch: <input type="radio"/> disabled <input type="radio"/> enabled <input checked="" type="radio"/> Platform Default</p> <p>DCU IP Pre-fetcher: <input type="radio"/> disabled <input type="radio"/> enabled <input checked="" type="radio"/> Platform Default</p> <p>Direct Cache Access: <input type="radio"/> disabled <input checked="" type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Processor C State: <input checked="" type="radio"/> disabled <input type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Processor C1E: <input checked="" type="radio"/> disabled <input type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Processor C3 Report: <input checked="" type="radio"/> disabled <input type="radio"/> acpi-c2 <input type="radio"/> acpi-c3 <input type="radio"/> Platform Default</p> <p>Processor C6 Report: <input type="radio"/> disabled <input checked="" type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>Processor C7 Report: <input checked="" type="radio"/> disabled <input type="radio"/> enabled <input type="radio"/> Platform Default</p> <p>CPU Performance: <input type="text" value="enterprise"/></p> <p>Max Variable MTRR Setting: <input type="radio"/> auto-max <input type="radio"/> 8 <input checked="" type="radio"/> Platform Default</p> <p>Local X2 APIC: <input type="radio"/> xapic <input type="radio"/> x2apic <input type="radio"/> auto <input checked="" type="radio"/> Platform Default</p> <p>Power Technology: <input type="text" value="performance"/></p> <p>Energy Performance: <input type="text" value="performance"/></p> <p>Frequency Floor Override: <input type="radio"/> disabled <input type="radio"/> enabled <input checked="" type="radio"/> Platform Default</p> <p>P-STATE Coordination: <input type="radio"/> hw-all <input type="radio"/> sw-all <input type="radio"/> sw-any <input checked="" type="radio"/> Platform Default</p> <p>DRAM Clock Throttling: <input type="text" value="Platform Default"/></p> <p>Channel Interleaving: <input type="text" value="Platform Default"/></p> <p>Rank Interleaving: <input type="text" value="Platform Default"/></p> <p>Demand Scrub: <input type="radio"/> disabled <input type="radio"/> enabled <input checked="" type="radio"/> Platform Default</p> <p>Patrol Scrub: <input type="radio"/> disabled <input type="radio"/> enabled <input checked="" type="radio"/> Platform Default</p> <p>Altitude: <input type="text" value="Platform Default"/></p>							



Configure Update Default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Click Save Changes.
6. Click OK to accept the change.



Create vNIC Templates for Cisco UCS B-Series

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Template_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure that the VM checkbox is not selected.
9. Select Updating Template as the Template Type.
10. Under VLANs, select the checkboxes for MGMT, Default, VDI, Infra, and vMotion.
11. Set Native-VLAN as the native VLAN.
12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_A.
14. In the Network Control Policy list, select CDP_Enabled.
15. Click OK to create the vNIC template.
16. Click OK.

Create vNIC Template

Name:

Description:

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

Target

☒ Adapter ☐ VM

Warning
If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: ☒ Initial Template ☐ Updating Template

VLANs

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	default	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	Infra_21	<input type="radio"/>
<input checked="" type="checkbox"/>	Mgmt_20	<input type="radio"/>
<input checked="" type="checkbox"/>	NFS-Storage	<input type="radio"/>
<input checked="" type="checkbox"/>	VDI_22	<input type="radio"/>
<input checked="" type="checkbox"/>	VLAN-OOB	<input type="radio"/>

Create VLAN

MTU:

Warning
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMQ

Dynamic vNIC Connection Policy:

Ok Cancel

17. In the navigation pane, select the LAN tab.
18. Select Policies > root.
19. Right-click vNIC Templates.
20. Select Create vNIC Template.

21. Enter vNIC_Template_B as the vNIC template name.
22. Select Fabric B.
23. Do not select the Enable Failover checkbox.
24. Under Target, make sure the VM checkbox is not selected.
25. Select Updating Template as the template type.
26. Under VLANs, select the checkboxes for MGMT, NFS-Storage, Default, VDI, Infra, and vMotion.
27. Set Native-VLAN as the native VLAN.
28. For MTU, enter 9000.
29. In the MAC Pool list, select MAC_Pool_B.
30. In the Network Control Policy list, select CDP_Enabled.
31. Click OK to create the vNIC template.
32. Click OK.

Create vHBA Templates for Cisco UCS B-Series

To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter vHBA_Template_A as the vHBA template name.
6. Keep Fabric A selected.
7. Select VSAN3 for Fabric A from the drop down.
8. Change to Updating Template.
9. For Max Data Field keep 2048.
10. Select SP-WWPN (created earlier) for our WWPN Pool.
11. Leave the remaining as is.

12. Click OK.

13. In the navigation pane, select the LAN tab.

14. Select Policies > root.

15. Right-click vHBA Templates.

16. Select Create vHBA Template.

17. Enter vHBA_Template_B as the vHBA template name.

18. Select Fabric B.

19. Select VSAN4 for Fabric B from the drop-down.

20. Change to Updating Template.

21. For Max Data Field keep 2048.

22. Select SP-WWPN (created earlier) for our WWPN Pool.

23. Leave the remaining as is.

24. Click OK.

Create Service Profile Templates for Cisco UCS B-Series

To create service profile templates for the Cisco UCS B-Series environment, complete the following steps:

1. Under the Servers tab in UCSM Select Service Profile Templates

Solution Validation

2. Right-click and select Create Service Profile Template
3. Name the template B-Series
4. Change to Updating Template
5. Select UUID pool created earlier

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. **Identify Service Profile Template**
2. Storage Provisioning
3. Networking
4. SAN Connectivity
5. Zoning
6. vNIC/vHBA Placement
7. vMedia Policy
8. Server Boot Order
9. Maintenance Policy
10. Server Assignment
11. Operational Policies

Identify Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name:

The template will be created in the following organization. Its name must be unique within this organization.

Where: **org-root**

The template will be created in the following organization. Its name must be unique within this organization.

Type: ☐ Initial Template ☒ Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.

UUID

UUID Assignment:

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev Next > Finish Cancel

6. Click Next
7. Click Next through Storage Provisioning
8. Under Networking, Select Expert
9. Click Add

The screenshot shows the 'Create Service Profile Template' wizard in the Unified Computing System Manager. The 'Networking' step is active, showing options for Dynamic vNIC Connection Policy and LAN connectivity configuration. A table for vNICs is present with columns for Name, MAC Address, Fabric ID, and Native VLAN. Below the table are 'Delete', 'Add', and 'Modify' buttons. A red arrow points to the 'Add' button. The left sidebar shows the progress of the wizard steps.

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ☒ Identify Service Profile Template
2. ☒ Storage Provisioning
3. ☒ **Networking**
4. ☐ SAN Connectivity
5. ☐ Zoning
6. ☐ vNIC/vHBA Placement
7. ☐ vMedia Policy
8. ☐ Server Boot Order
9. ☐ Maintenance Policy
10. ☐ Server Assignment
11. ☐ Operational Policies

Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) + Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity? ☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

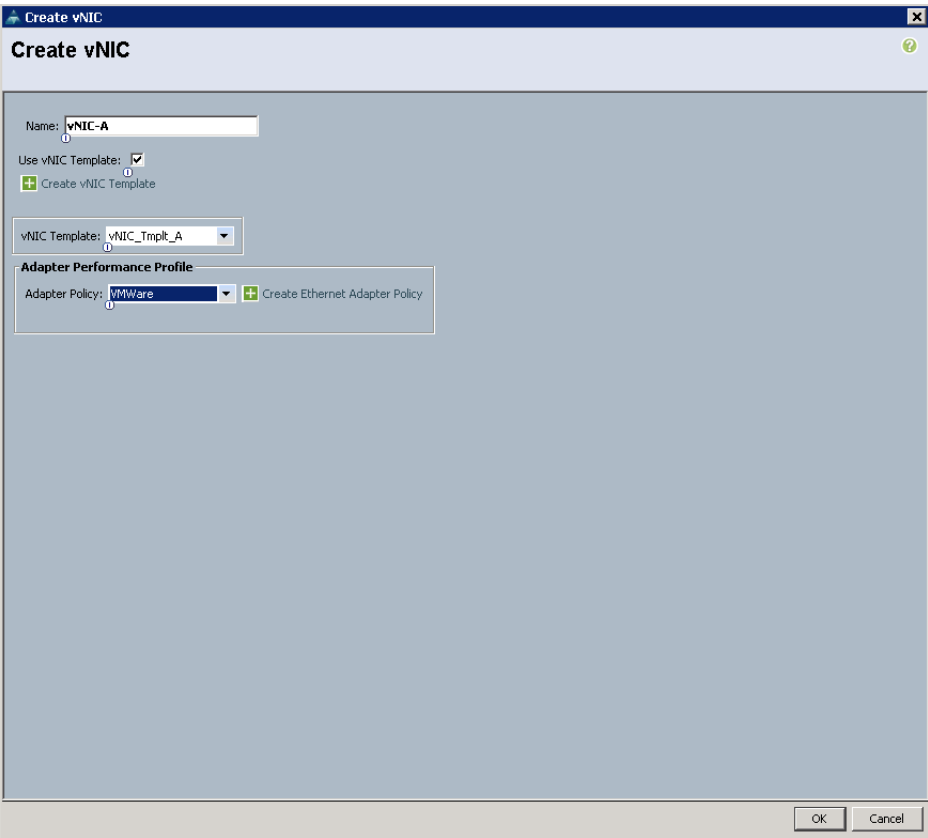
Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN

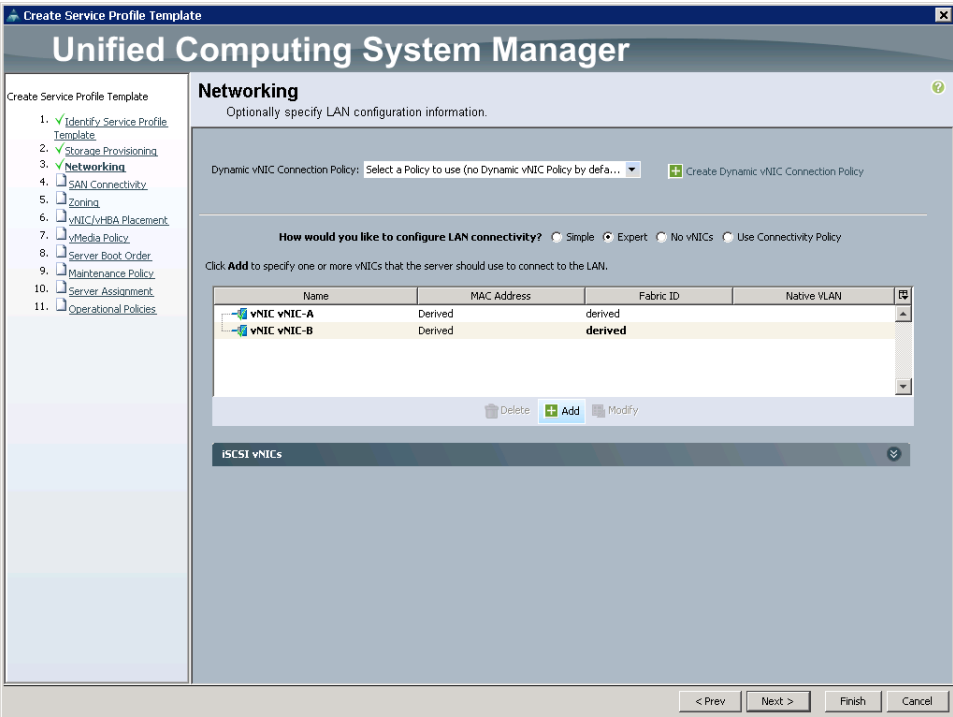
✖ Delete + Add ✎ Modify

ISCSI vNICs

10. Name it vNIC-A.
11. Select check box for Use vNIC Template.
12. Under vNIC template select the vNIC_Tmpl_A.
13. For Adapter Policy select VMware.

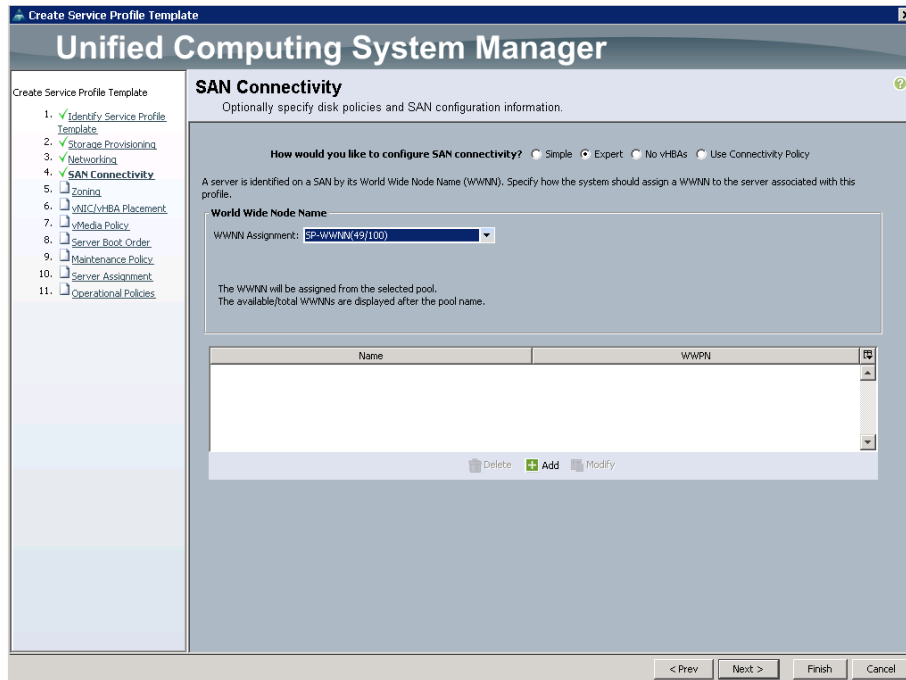


14. Repeat networking steps for vNIC_Tmpl_B




15. Click Next

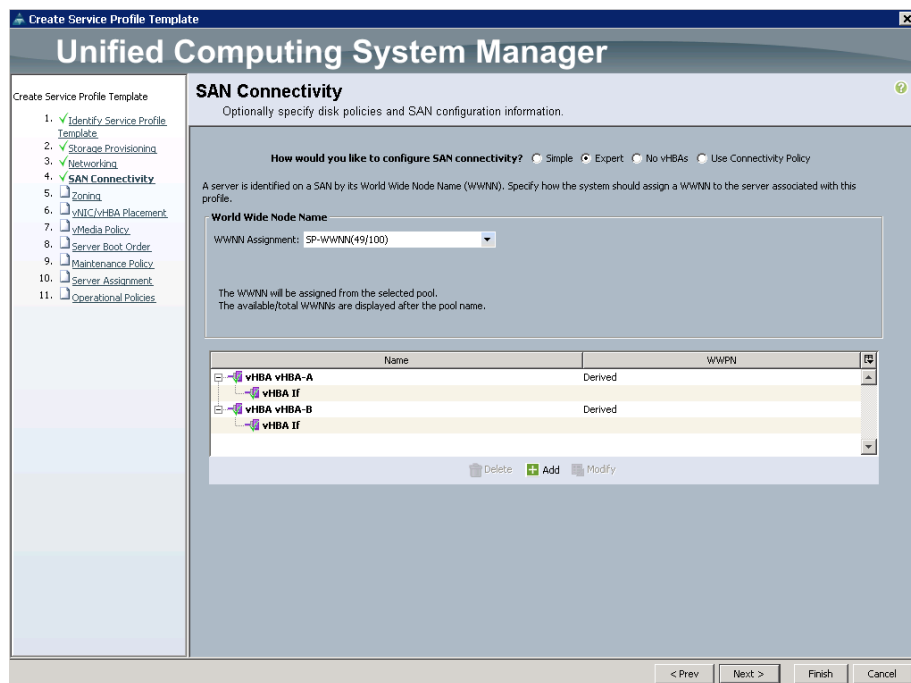
16. Under SAN Connectivity, select Expert
17. Select WWNN Assignment from the Pool created earlier
18. Click Add



19. Name the adapter vHBA-A
20. Select vHBA Template: vHBA-A
21. Select Adapter Policy : VMWare



22. Repeat steps for vHBA-B on Fabric B.

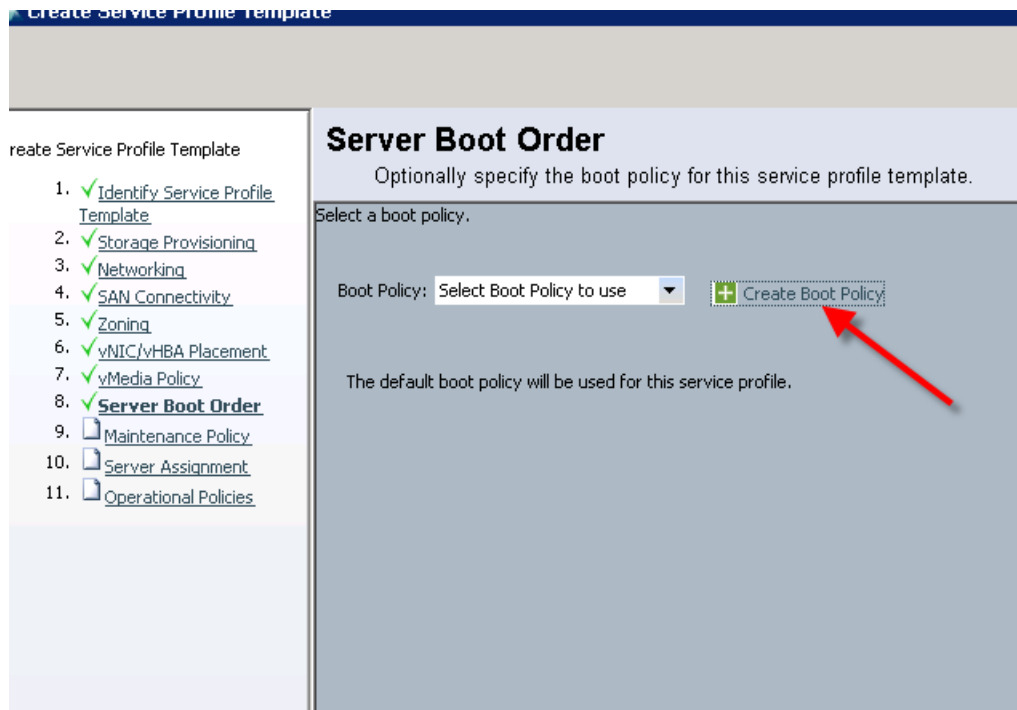


23. No Zoning will be used.

24. Click Next through vNIC/vHBA Placement policy.

25. Click Next through vMedia Policy.

26. Click Create Boot Policy to create a Boot From SAN policy.



27. Add Remote CD/DVD to install OS from ISO image.

Create Boot Policy

Name:

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☒

Boot Mode: ☒ Legacy ☐ Uefi

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Add External USB

Add Embedded Local LUN

Add Embedded Local Disk

Add CD/DVD

Add Local CD/DVD

Add Remote CD/DVD

Add Floppy

Add Local Floppy

Add Remote Floppy

Add Remote Virtual Drive

vNICs

vHBAs

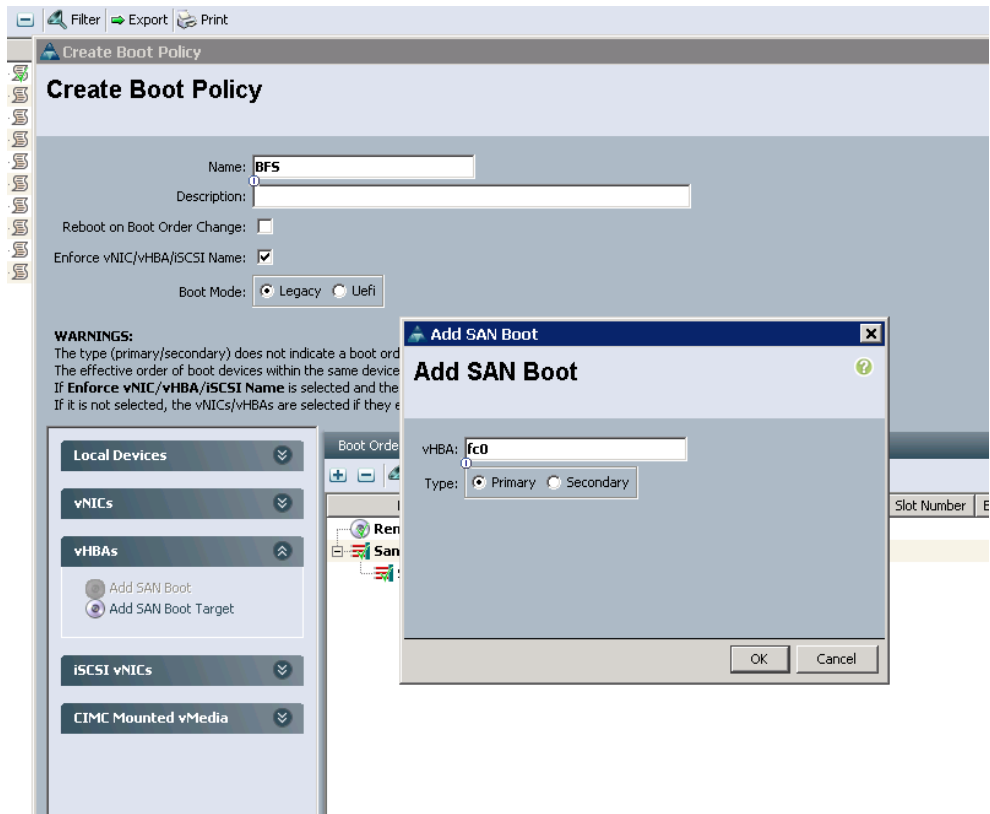
Add SAN Boot

Boot Order

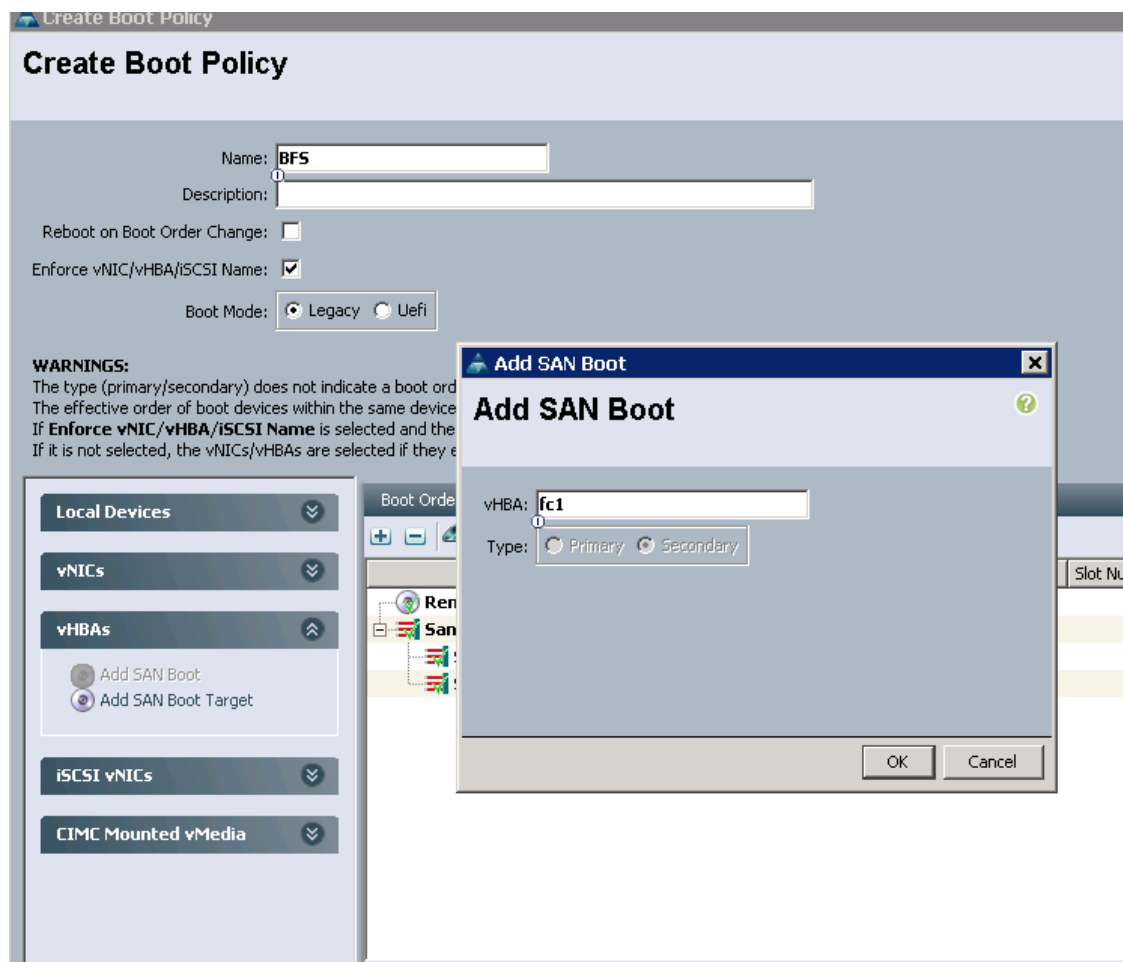
Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN	Slot
Remote CD/DVD	1					

Move Up Move Down Delete

28. Click Add SAN Boot to add an HBA for Boot From SAN. Label it 'fc0'



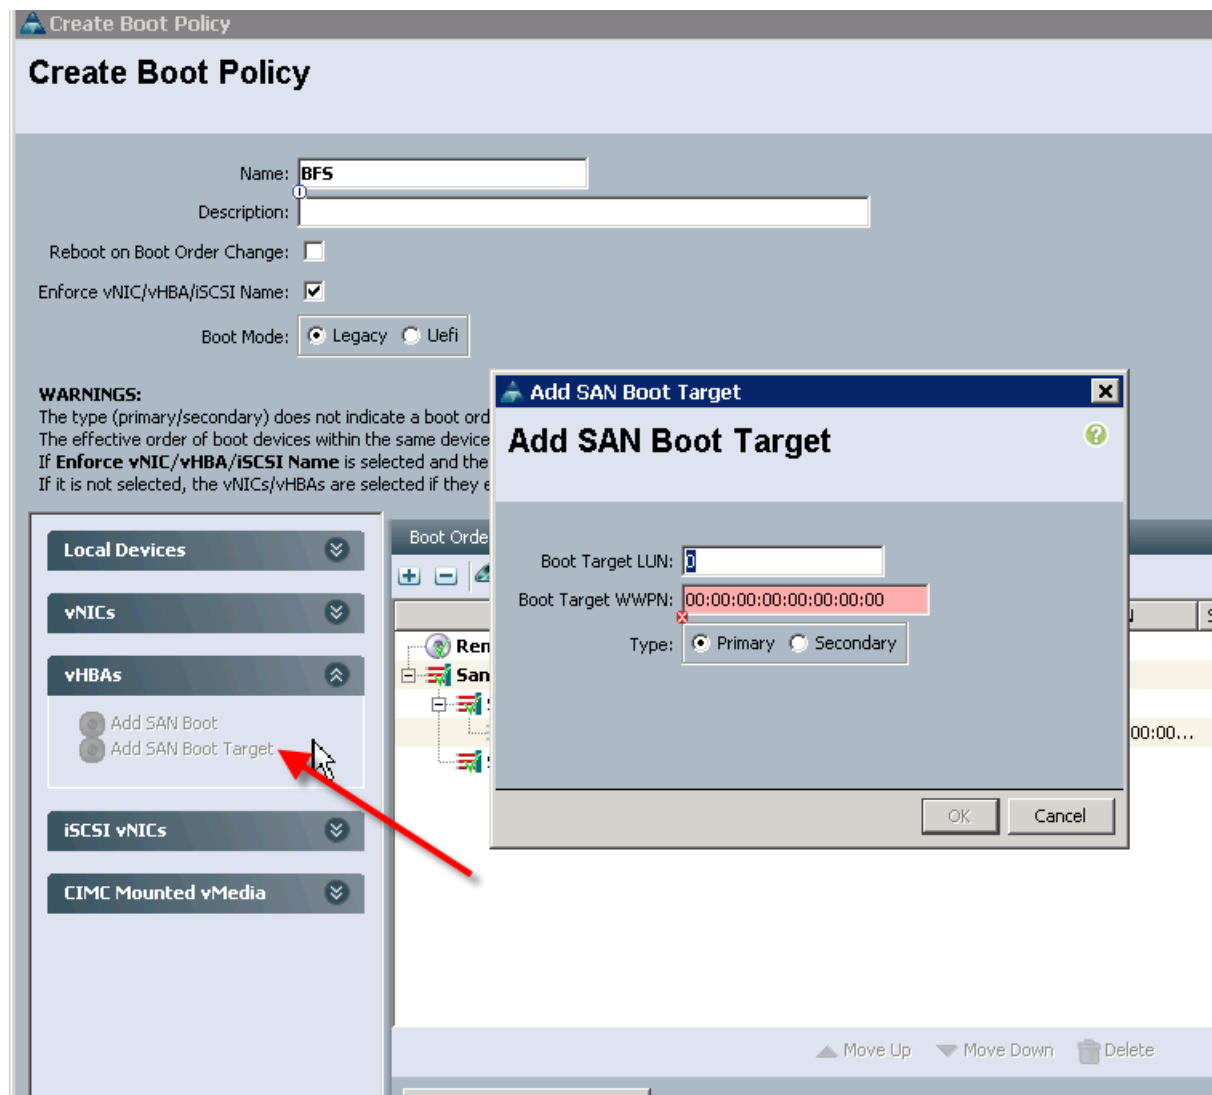
29. Add a second SAN Boot vHBA and label it 'fc1



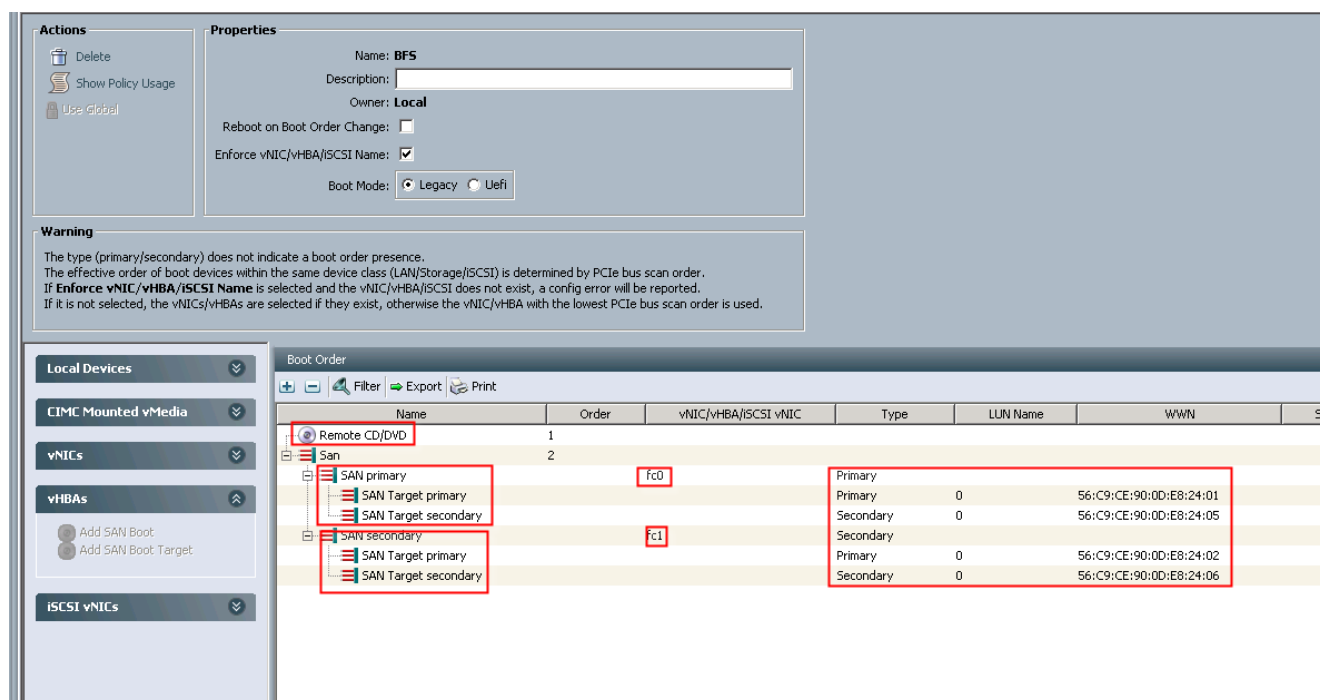
30. Add a SAN Boot Target to enter the WWPN for the Nimble Array. We will add 4 SAN Boot Targets. One Primary and one Secondary per vHBA.

The boot targets will be added in the following order:

- SAN Primary interface (fc0)
 - Target HBA-> Controller A: HBA FC5.1
 - Target HBA -> ControllerB: HBA FC5.1
- SAN Secondary interface (fc1)
 - Target HBA-> Controller A: HBA FC6.1
 - Target HBA -> ControllerB: HBA FC6.1



31. The final boot policy should look like this with the WWN filled representing the WWPN that belong to the FC ports on the Nimble CS700 Adaptive Array.



32. Click through the defaults for the remaining Service Profile Template Wizard. We can now create Service Profiles from this template and assign to our B-Series servers in our chassis. In this Solution, 16 Service Profiles were created and assigned to 2x fully populated 5108 Chassis with B200 M4 Servers. All Servers have ESXi installed and boot from the Nimble CS700 SAN.

Nimble CS700 Configuration

Nimble CS700 Adaptive Array System Configuration

This section provides the procedure for initializing a Nimble Storage array and setting up basic IP connectivity.

Nimble Setup Manager

The Nimble Setup manager can be downloaded from Infosight at this location: <http://infosightweb.nimblestorage.com/InfoSight/cgi-bin/downloadFile?ID=documents/Setup-NimbleNWT-x64.2.3.2.287.zip>

The Nimble Setup manager is part of the Nimble Storage Windows Toolkit. In this instance the Nimble Setup Manager is the only component that needs to be installed.

Initialize the Nimble Storage Array

To initialize the Nimble Storage Array, complete the following steps:

1. In the Windows Start menu, click Nimble Storage > Nimble Setup Manager.
2. Select one of the uninitialized arrays from the Nimble Setup Manager list and click Next.



If the array is not visible in Nimble Setup Manager, verify that the array's eth1 ports of both controllers are on the same subnet as the Windows host.

Configure the Nimble OS using the GUI

To configure the Nimble OS using the GUI, complete the following steps:

1. Choose the appropriate group option and click Next.
 - Set up the array but do not join a group. Continue to Step 5.
 - Add the array to an existing group.

If you chose to join an existing group, your browser automatically redirects to the login screen of the group leader array. See *Add Array to Group Using the GUI* to complete the configuration.

2. Provide or change the following initial management settings and click Finish:
 - Array name
 - Group name
 - Management IP address and subnet mask for the eth1 interface
 - Default gateway IP address
 - Optional. Administrator password
3. You may see a warning similar to There is a problem with this website's security certificate. It is safe to ignore this warning and click Continue.



If prompted, you can also download and accept the certificate. Alternatively, create your own. See the cert command in the *Nimble Command Line Reference Guide*. Also, if Internet Explorer v7 displays a blank page, clear the browser's cache. The page should be visible after refreshing the browser.

4. In the login screen, type the password you set and click Log In. From this point forward, you are in the Nimble OS GUI. The first time you access the Nimble OS GUI, the Nimble Storage License Agreement appears.
5. In the Nimble Storage License Agreement, read the agreement, scroll to the bottom, check the acknowledgment box, and then click Proceed.
6. Provide the Subnet Configuration information for the following sections and click Next:
 - a. Management IP: IP address, Network and Subnet Mask



The Management IP is used for the GUI, CLI, and replication. It resides on the management subnet and floats across all "Mgmt only" and "Mgmt + Data" interfaces on that subnet. Note: in this instance you only need to configure the Management network. No IP data network connectivity is required.

- b. Subnet: Subnet label, Network, Netmask, Traffic Type(Data only, Mgmt Only, Mgmt +Data), MTU
- 7. Maximum Transmission Unit (MTU) – Standard (1500) Provide Interface Assignment information for the following sections and click Next:
 - a. Interface Assignment: For each IP interface, assign it a subnet and a Data IP address within the specified network. For inactive interface, assign the "None" subnet.
 - b. Diagnostics:
 - i. Controller A diagnostics IP address will be on the same subnet as the management IP address.
 - ii. Controller B diagnostics IP address will be on the same subnet as the management IP address.
- 8. Provide the following Domain information and click Next:
 - a. Domain Name
 - b. DNS Servers: Type the hostname or IP address of your DNS server. You can list up to five servers.
- 9. Provide the following Time information and click Next:
 - a. Time Zone: Choose the time zone the array is located in.
 - b. Time (NTP) Server: Type the hostname or IP address of your NTP server.
- 10. Provide Support information for the following sections and click Finish.
- 11. Email Alerts:
 - a. From Address: This is the email address used by the array when sending email alerts. It does not need to be a real email address. Include the array name for easy identification.
 - b. Send to Address: Nimble recommends that you check the Send event data to Nimble Storage Support check box.
 - c. SMTP server hostname or IP address
 - d. AutoSupport:
 - i. Checking the Send AutoSupport data to Nimble Storage check box enables Nimble Storage Support to monitor your array, notify you of problems, and provide solutions.
 - e. HTTP Proxy: AutoSupport and software updates require an HTTPS connection to the Internet, either directly or through a proxy server. If a proxy server is required, check the Use HTTP Proxy check box and provide the following information to configure proxy server details:
 - i. HTTP proxy server hostname or IP address
 - ii. HTTP proxy server port
 - iii. Proxy server user name
 - iv. Proxy server password



The system does not test the validity of the SMTP server connection or the email addresses that you provided.

12. Click Finish. The Setup Complete screen appears. Click Continue.

13. The Nimble OS home screen appears. Nimble Storage array setup is complete.

Nimble Management Tools: InfoSight

Register and Login to InfoSight

Before You Begin

It can take up to 24 hours for the array to appear in InfoSight after the first data set is sent. Data sets are configured to be sent around midnight array local time. Changes made right after the data set is sent at midnight might not be reflected in InfoSight for up to 48 hours.

Procedure

1. Log in to the InfoSight portal at <https://infosight.nimblestorage.com>.
2. Click Enroll now to activate your account. If your email is not already registered, contact your InfoSight Administrator. If there is no existing, InfoSight Administrator (Super User) registered against your account or you are not sure, contact Nimble Storage Support for assistance.
3. Select the appropriate InfoSight role and enter the array serial number for your customer account. If this is the first account being created for your organization, you should select the Super User role. The number of super users is limited to the total number of arrays that are associated with an account.
4. Click Submit.
5. A temporary password is sent to the email address that you specified. You must change your password the first time you log in.

Configure Arrays to Send Data to InfoSight

To take full advantage of the InfoSight monitoring and analysis capabilities, configure your Nimble arrays to send data sets, statistics, and events to Nimble Storage Support. InfoSight recommendations and automatic fault detection are based on data that is sent from your arrays and processed by InfoSight. If you did not configure your Nimble arrays to send this data to Nimble Storage Support during the initial setup, you can change the configuration at any time from the Administration menu in the GUI or by running the group `--edit` command in the CLI.

Before You Begin

This procedure must be performed in the array GUI.

Procedure

1. From the Administration menu in the array GUI, select Alerts and Monitoring > AutoSupport / HTTP Proxy.
2. On the AutoSupport page, select Send AutoSupport data to Nimble Storage Support.
3. Click Test AutoSupport Settings to confirm that AutoSupport is set up correctly.
4. Click Save.

Nimble Management Tools: vCenter Plugin

Register vCenter Plugins

The vCenter plugin from Nimble Storage allows for single pane of glass administration directly from vCenter as well as integration with Nimble Infosight analytics. The first step is performed on the array GUI, not in Infosight. Nimble Storage has integration vCenter through plugin registration. This allows for datastore creation and management using vCenter. The vCenter plugin is supported on ESX 5.5 update 1 and later.



The plugin is not supported for:

- Multiple datastores located on one LUN
 - One datastore spanning multiple LUNs
 - LUNs located on a storage device not made by Nimble
-

Procedure

Use a vCenter account that has sufficient privileges to install a plugin (usually a user assigned to the Administrator role). You need to know the vCenter hostname or IP address. The plugin is part of the Nimble OS. To take advantage of it, you must first register the plugin with a vCenter Server. Multiple plugins can be registered on the Nimble array. In turn, each array that registers the plugin adds a tab to the vSphere client. The tab name for the datastore page is "datacenter page-Nimble-<groupname>". To register the vCenter plugins, complete the following steps:

1. From the Nimble OS GUI main menu, select Administration > vCenter Plugin.
2. If the fields are not already filled, enter the vCenter server host name or IP address, user name, and password.
3. Click View Status to see the current status of the plugin.
4. Click Register. If a Security Warning message appears, click Ignore.
5. If you are not sure of which subnet_label to select, the selection that appears when the dialog opens is probably the correct one.

The Nimble Storage plugin must be registered with your vCenter servers before it can be used by your vCenter clients.

Subnet	mgmt-data	Choose a subnet that vCenter clients can route to. The management subnet is often the best choice.
vCenter Host	vcenter.yourcompany.com	Port 443
Username	Administrator	
Password	

Register Unregister View Status

- Click View Status again to ensure that the plugin has been registered.
- Restart the vSphere client.

View a List of Installed Plugins

A list of all registered plugins on the array can be discovered by completing the following steps:

- Log into the Nimble OS CLI.
- At the command prompt, type:

```
vmwplugin --list --username <username> --password <password> --server
<server_hostname-address> --port port_number <port number>
```



If no port number is specified, port 443 is selected by default. A list of installed plugins displays.

Configure Arrays to Monitor your Virtual Environment

To configure arrays to monitor your virtual environment, complete the following steps:

- Log in to <https://infosight.nimblestorage.com>
- Go to Administration > Virtual Environment.
- In the Virtual Environment list, find the array group for which you want to monitor the virtual environment.
- Click Configure.
- The Configure Group dialog box opens.
- Verify that your software version is up to date and that your vCenter plugin is registered.
- Select Enable in the VM Streaming Data list.

8. Click Update.

Configure Setup Email Notifications for Alerts

To configure setup email notifications for alerts, complete the following steps:

1. On the Wellness page, click Daily Summary Emails.
2. Check Subscribe to daily summary email.
3. Enter an email address for delivery of the email alerts.
4. (Optional) You can click Test to send a test email to the email address that you indicated.
5. Click Submit to conclude the email alerts setup.

Data Storage Layout

For this solution, the layout for the Nimble CS700 Adaptive Array are listed in the following table. This is the recommended sizes and Performance Policies for best performance in this VDI solution.

Table 3 Layout for the Nimble CS700 Adaptive Array

Volumes	Name	Size	VMware / OS Connect / Boot	Performance Policy
Infrastructure	INFRA-DS	2TB	VMware	VMware ESX
Infrastructure Boot 1	Infra1Boot	10GB	Boot	VMware ESX
Infrastructure Boot 2	Infra2Boot	10GB	Boot	VMware ESX
vDisk	vDisks	500GB		Windows File Server
User Profiles	Profiles	3TB	File Share	Windows File Server
Boot LUN 1	SP-VDI-01	10GB	Boot	VMware ESX
Boot LUN 2	SP-VDI-02	10GB	Boot	VMware ESX
Boot LUN 3	SP-VDI-03	10GB	Boot	VMware ESX
Boot LUN 4	SP-VDI-04	10GB	Boot	VMware ESX
Boot LUN 5	SP-VDI-05	10GB	Boot	VMware ESX
Boot LUN 6	SP-VDI-06	10GB	Boot	VMware ESX
Boot LUN 7	SP-VDI-07	10GB	Boot	VMware ESX
Boot LUN 8	SP-VDI-08	10GB	Boot	VMware ESX
Boot LUN 9	SP-VDI-09	10GB	Boot	VMware ESX
Boot LUN 10	SP-VDI-10	10GB	Boot	VMware ESX
Boot LUN 11	SP-VDI-11	10GB	Boot	VMware ESX
Boot LUN 12	SP-VDI-12	10GB	Boot	VMware ESX
Boot LUN 13	SP-VDI-13	10GB	Boot	VMware ESX
Boot LUN 14	SP-VDI-14	10GB	Boot	VMware ESX
Write Cache 1-25	WC-01-25	1TB	VMware	VMware ESX

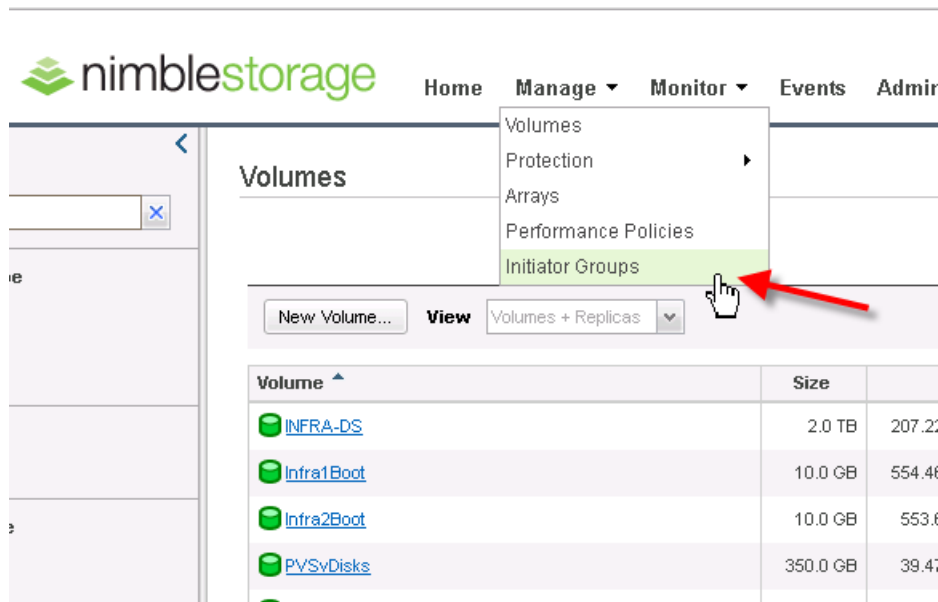
Create Initiator Groups

We must create our initiator groups to allow FC access to the array from the ESXi hosts. The following steps will show how to create an Initiator group in the Nimble GUI.

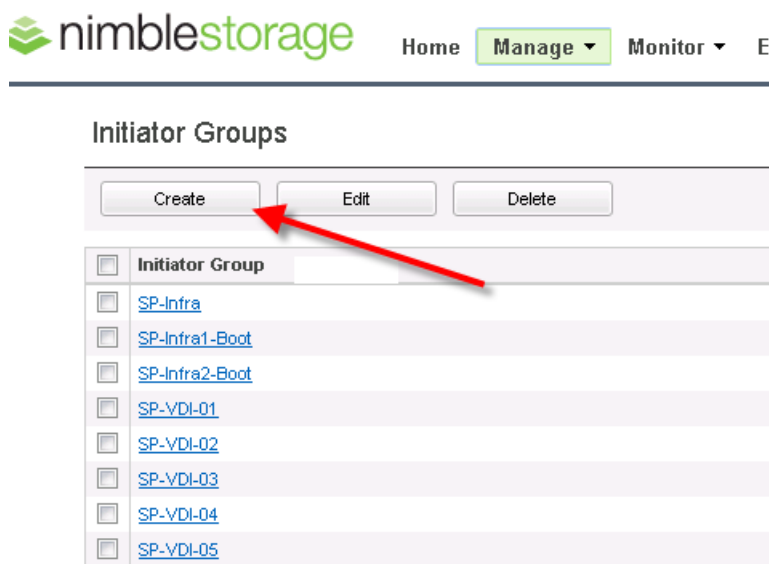


An initiator group will need to be created for every ESXi host.

1. On the Nimble Home screen, select Manage.



2. Select Initiator Groups and click Create.



3. Enter the name of your initiator group and the WWPN for each ESXi host vHBA.


Create an Initiator Group

Initiator Groups are a convenient way to limit volume access to only the specific initiators that are members of the group.

Name:

Initiators

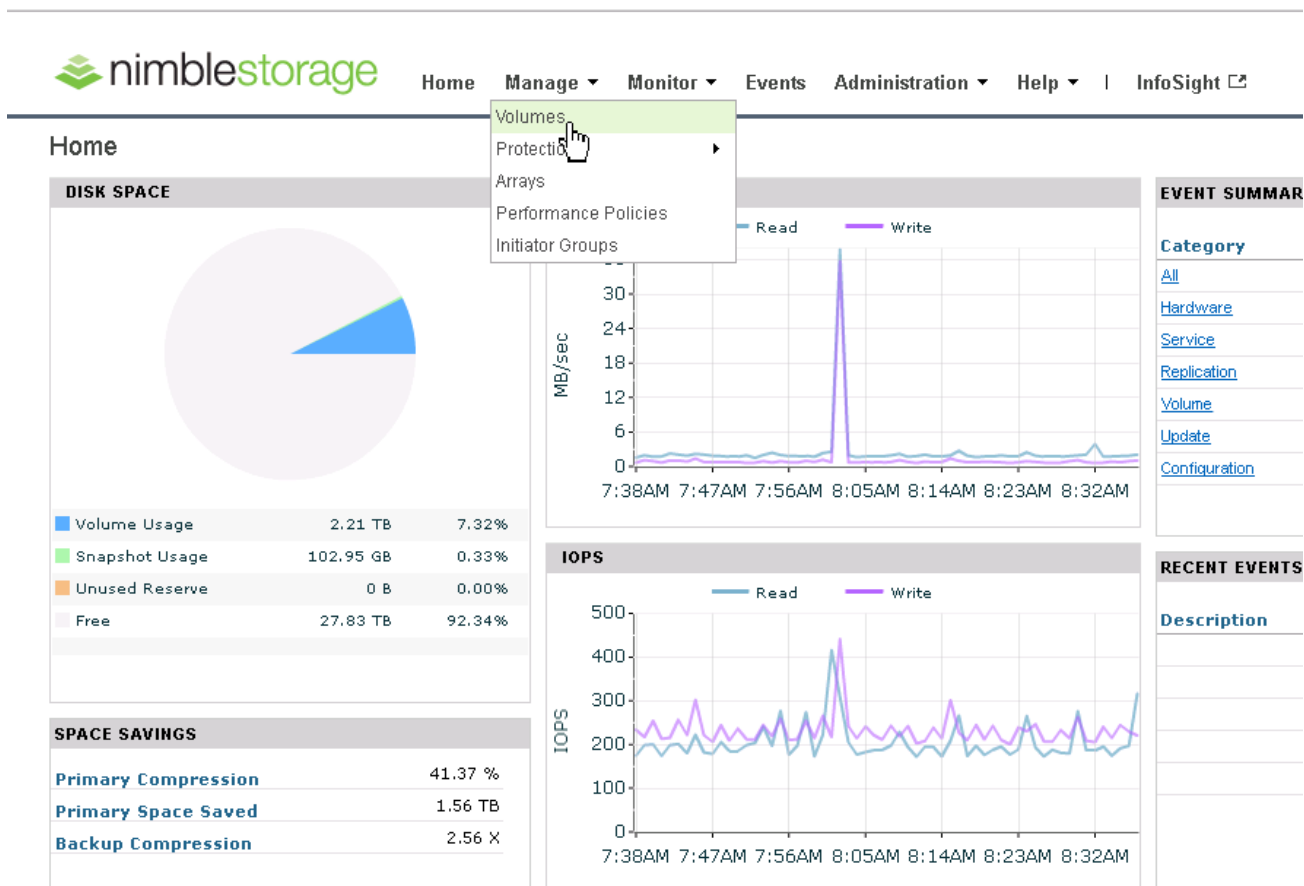
Specify an alias and WWPN for each initiator. To gain access an initiator must match both if provided.

Alias (Optional)	WWPN 	
<input type="text" value="vHBA0"/>	<input type="text" value="20:00:00:25:B5:00:00:2C"/>	<input type="button" value="X"/>
<input type="text" value="vHBA1"/>	<input type="text" value="20:00:00:25:b5:00:00:2f"/>	<input type="button" value="X"/>

Create Volumes

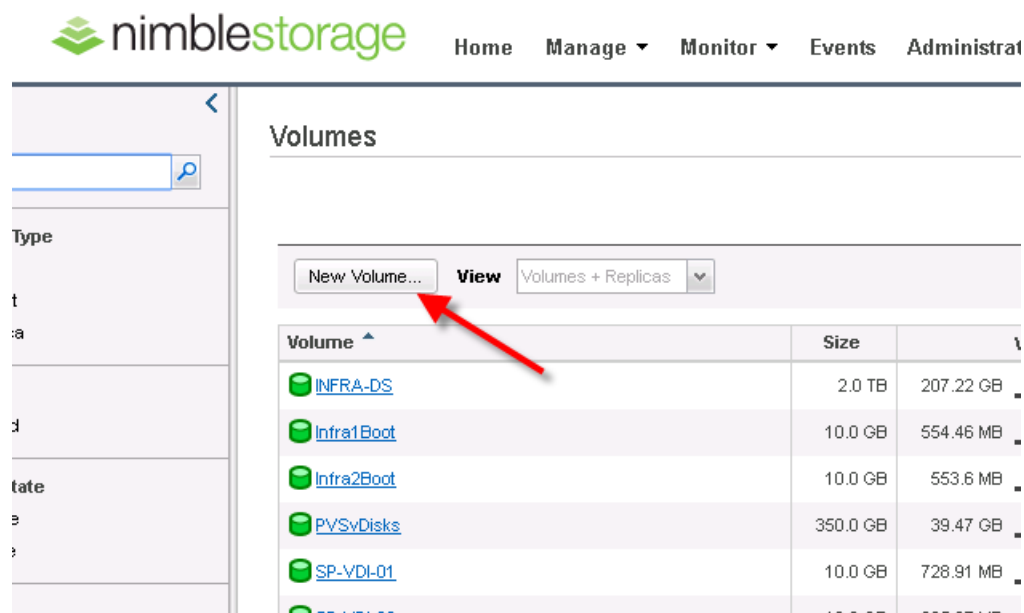
To create a volume on the Nimble CS700 Adaptive Array, complete the following steps:

1. On the Nimble Home page, select Manage > Volumes.



The screenshot shows the Nimble Storage management interface. The top navigation bar includes links for Home, Manage, Monitor, Events, Administration, Help, and InfoSight. The 'Manage' menu is open, showing options: Volumes, Protection, Arrays, Performance Policies, and Initiator Groups. The 'Volumes' option is highlighted. The main content area displays the 'Home' dashboard with a 'DISK SPACE' pie chart, a table of disk usage, and a 'SPACE SAVINGS' section. The 'DISK SPACE' table shows Volume Usage (2.21 TB, 7.32%), Snapshot Usage (102.95 GB, 0.33%), Unused Reserve (0 B, 0.00%), and Free space (27.83 TB, 92.34%). The 'SPACE SAVINGS' section shows Primary Compression (41.37%), Primary Space Saved (1.56 TB), and Backup Compression (2.56 X). On the right, there are two line graphs: 'MB/sec' (Read and Write) and 'IOPS' (Read and Write), both showing a significant spike around 8:05 AM. The 'EVENT SUMMAR' and 'RECENT EVENTS' sections are also visible on the right side of the dashboard.

2. Click New Volume.



The screenshot shows the Nimble Storage web interface. The top navigation bar includes the Nimble Storage logo and links for Home, Manage, Monitor, Events, and Administration. The left sidebar contains a search bar and a list of volume types. The main content area is titled 'Volumes' and features a 'New Volume...' button, a 'View' button, and a dropdown menu set to 'Volumes + Replicas'. Below this is a table listing existing volumes.

Volume	Size	Used
INFRA-DS	2.0 TB	207.22 GB
Infra1Boot	10.0 GB	554.46 MB
Infra2Boot	10.0 GB	553.6 MB
PVSvDisks	350.0 GB	39.47 GB
SP-VDI-01	10.0 GB	728.91 MB
SP-VDI-02	10.0 GB	995.87 MB

2. Input Volume Name, Description and Select your Performance Policy. In this example we used the Windows Files Server Performance policy because the volume being created was used for User Data. In our ESXi policies we utilized VMware ESX Performance Policy.

Assign to the proper initiator group (Initiator Groups for this solution include one per ESXi Host with both of their vHBA WWPN added to the initiator group).


Create a volume


Create a volume

General > Space > Protection > Performance

Volume Name

Description Optional

Performance Policy 

Data Encryption  Disabled

ACCESS CONTROL

This access control entry will be applied to both the volume and its associated snapshots. Access control can be modified and refined after the volume is created.

Grant access to initiator group

3. Enter the size of the volume desired and click Next.

Create a volume

Create a volume

General > **Space** > Protection > Performance

With built-in thin provisioning and in-line compression, you can create a volume with a size (as reported to the application) that exceeds the available free space.

Size ⓘ GB ▾

SPACE ⓘ

	% of Size	Bytes
Volume Reserve ⓘ <input type="text" value="0"/>	0	0.0 GB
Volume Quota <input type="text" value="100"/>	100	500.0 GB
Volume Warning <input type="text" value="80"/>	80	400.0 GB
Snapshot Reserve <input type="text" value="0"/>	0	0.0 GB
Snapshot Quota <input type="text"/>	Unlimited	<input checked="" type="checkbox"/> Unlimited snapshot quota
Snapshot Warning <input type="text" value="0"/>	0	0.0 GB

STORAGE

Capacity	30.14 TB
Used Space	2.31 TB
Free Space	27.83 TB

BackNextFinishCancel

- (Optional) Select your Protection Plan for this volume. For this project we created a snapshot schedule for our Infrastructure Volumes only. For WriteCache data and other volatile data we did not assign a protection plan.

Create a volume

Create a volume

General > Space > **Protection** > Performance

PROTECTION ⓘ

Volumes assigned to a volume collection are protected according to the volume collection's protection schedule. Standalone volumes can be protected using a protection template or by creating a custom protection schedule.

☐ No volume collection

☒ Join volume collection

☐ Create new volume collection

☐ Protect as standalone volume

PROTECTION SCHEDULES

Infrastructure	Schedule Trigger: Native
	Snapshot every: 1 day
	At: 00:00
	Retain up to: 30 Snapshots
	On the following days: Sun, Mon, Tue, Wed, Thu, Fri, Sat

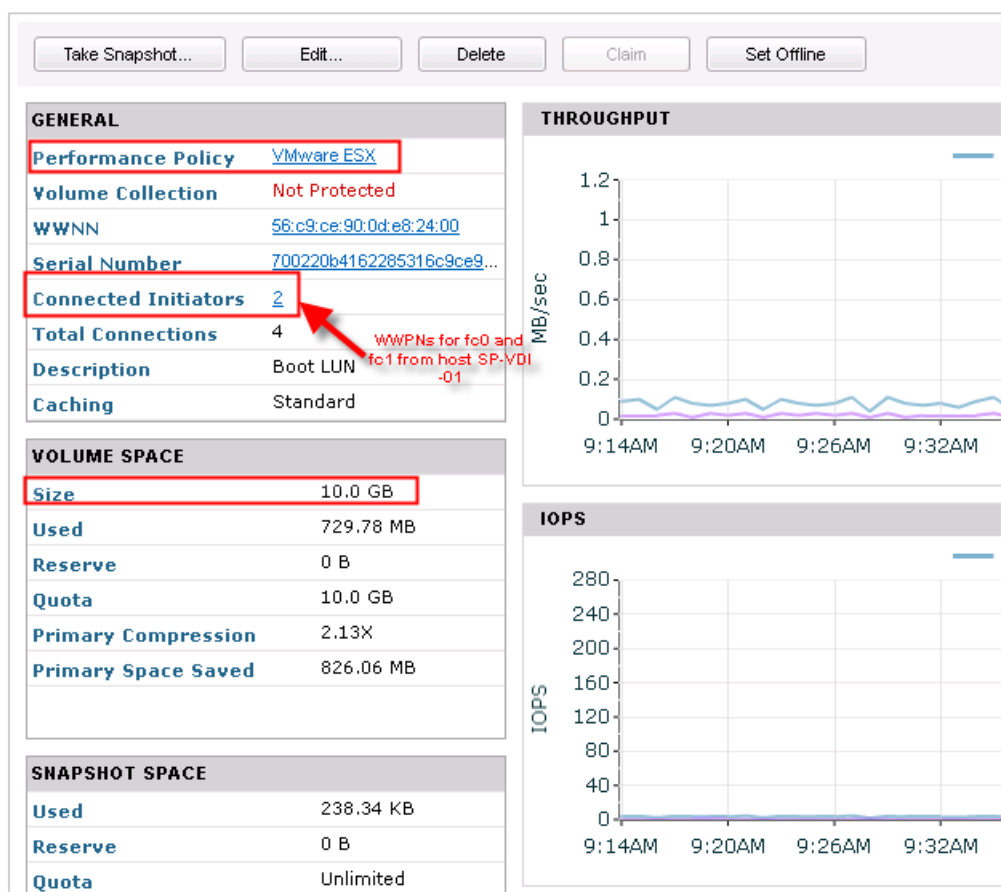
BackNextFinishCancel

5. Click Finish.

The screenshot shows a 'Create a volume' wizard window. The title bar says 'Create a volume'. Inside the window, the title 'Create a volume' is at the top. Below it is a breadcrumb navigation: 'General > Space > Protection > Performance'. The 'Performance' tab is selected. Below the breadcrumb is a section titled 'CACHING'. Under 'CACHING', there is a label 'Volume Caching' followed by two radio buttons: 'Normal (default)' (which is selected) and 'Pinned'. To the right of the 'Pinned' radio button is a blue information icon. At the bottom of the window are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

6. For a Boot LUN, we created a 10GB Volume with the ESX Performance Policy and assigned it to the initiator group for its host.

Volumes > SP-VDI-01 ● Online



Configure MDS 9100 Series

In this solution we utilized the Cisco MDS 9148 Switches for Fiber Channel Switching. For racking, cable and initial setup of the MDS switches, please refer to the Quick Start Guide:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/hw/9148/quick/guide/MDS_9148_QSG.pdf

When the MDS switch is racked and can be logged into it can now be configured to communicate with the Cisco UCS Fabric Interconnects.

In this study, we used two separate fabrics each with their own unique VSAN. Fabric A is configured for VSAN3 while Fabric B for VSAN4. In our initial UCS configuration you will see where we configured fiber cables on ports 31 and 32 and configured a FC port-channel. FI-A's FC port channel is configured for VSAN3 and FI-B's FC port-channel for VSAN4.

Figure 16 Fabric A

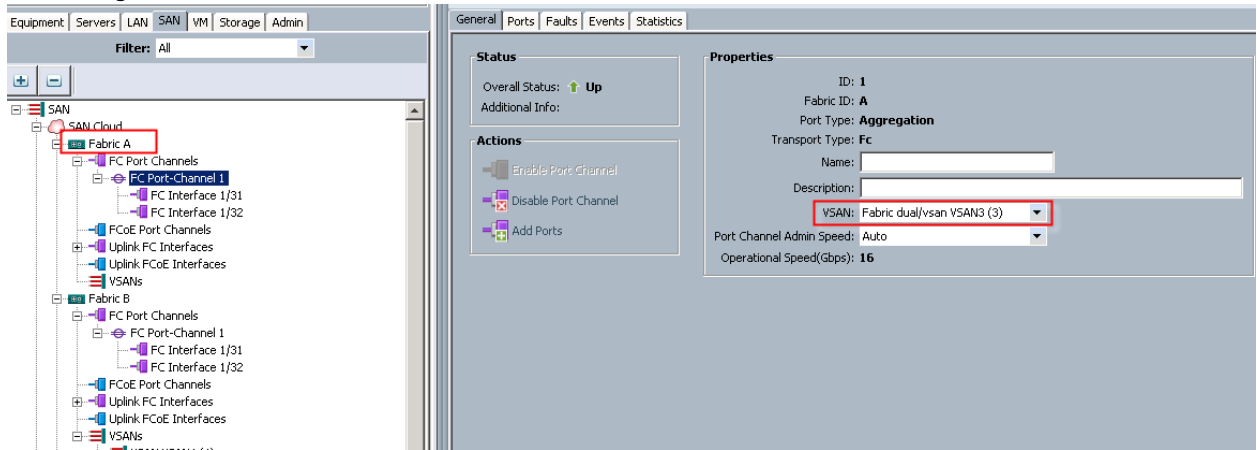
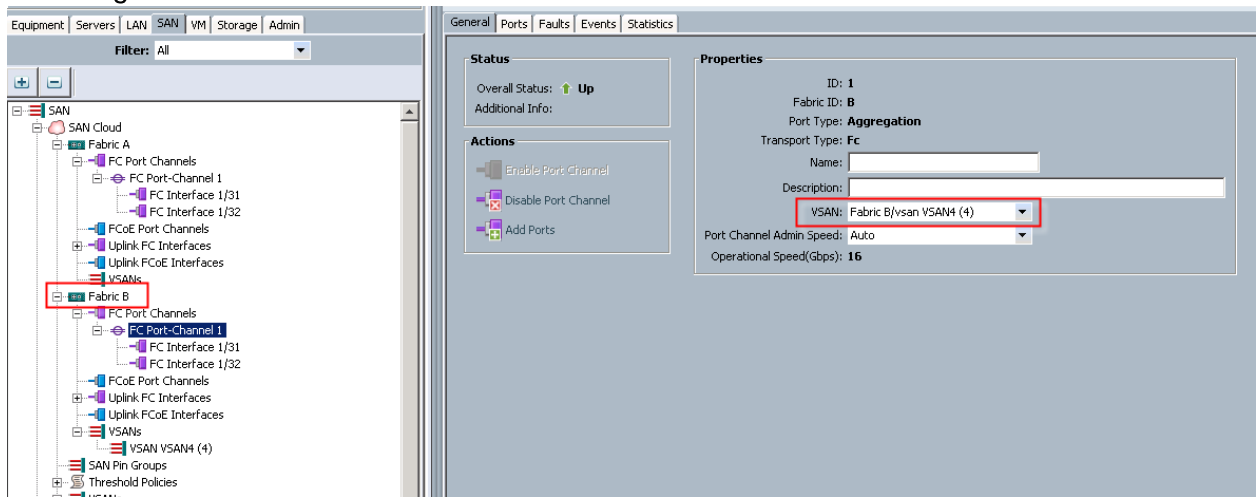
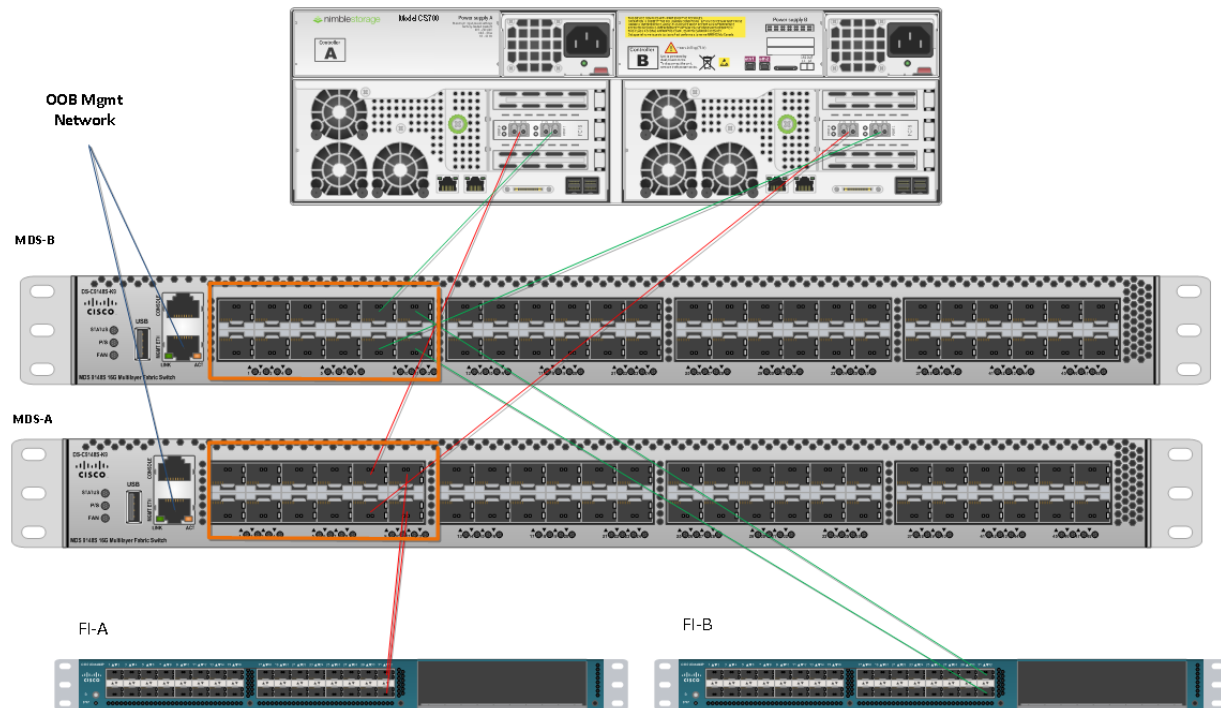


Figure 17 Fabric B



Physically, the Fabric Interconnects ports 31 and 32 run to the MDS switch ports 11 and 12.



Ports 11 and 12 on the MDS switches are configured in a port-channel as well while ports 9 and 10 are plugged into the Nimble CS700 Adaptive Array.

Device Manager 6.2(9a) - MDS-B [admin]

Device Physical Interface FC FICON IP Security Admin Logs Help

VSAN 4 Ports All

Device Summary

STATUS: Console, PS, FAN

6.2(9a)

Legend: Up (Green), Down (Yellow), Fail (Red), Minor (Orange), Unreachable (Grey), OutOfService (Blue)

MDS-B - Port Channels

General FC Interfaces Trunk Config Trunk Failures Auto Create FLOGI

Channel	Admin Mode	Oper Mode	Force	MemberList By Interface	MemberList LoadBalanced	LastAction Status	LastAction FailureCause	LastAction Time	CreationTime
channel1	active	active	<input type="checkbox"/>	fc1/11-fc1/12	fc1/11-fc1/12	successful		n/a	n/a

1 row(s)

Device Manager 6.2(9a) - MDS-A [admin]

Device Physical Interface FC FICON IP Security Admin Logs Help

VSAN 3 Ports All

Device Summary

STATUS: Console, PS, FAN

6.2(9a)

Legend: Up (Green), Down (Yellow), Fail (Red), Minor (Orange), Unreachable (Grey), OutOfService (Blue)

MDS-A - Port Channels

General FC Interfaces Trunk Config Trunk Failures Auto Create FLOGI

Channel	Admin Mode	Oper Mode	Force	MemberList By Interface	MemberList LoadBalanced	LastAction Status	LastAction FailureCause	LastAction Time	CreationTime
chann...	active	active	<input type="checkbox"/>	fc1/11-fc1/12	fc1/11-fc1/12	successful		2016/02/27-15:3...	2016/02/27-15:...

1 row(s)

After the ports and port channels are configured, the next steps are to configure the zones and Active Zoneset Database in the MDS switches. The below commands show how to add in a single host on both MDS A and B. You will need to configure all hosts that will access the Nimble Array in these commands. Then entire MDS switch configuration is included in this document in Error! Reference source not found..

MDS-A

Configure Terminal

Zoneset name SP-Infra-A vsan 3

Zone name {ESXi hostname-fc0} vsan 3

Solution Validation

```
Member pwwn {ESXi Host pwwn for fc0}
Member pwwn {Nimble pwwn Controller A, Port 1}
Member pwwn {Nimble pwwn Controller B, Port 1}
Zone commit vsan 3
Zoneset name SP-Infra-A vsan 3
Member {ESXi hostname-fc0}
Exit
Zoneset activate name SP-Infra-A vsan 3
Zone commit vsan 3
Exit
Copy running-config startup-config
```

MDS-B

```
Configure Terminal
Zoneset name SP-Infra-B vsan 4
Zone name {ESXi hostname-fc1} vsan 4
Member pwwn {ESXi Host pwwn for fc1}
Member pwwn {Nimble pwwn Controller A, Port 2}
Member pwwn {Nimble pwwn Controller B, Port 2}
Zone commit vsan 4
Zoneset name SP-Infra-B vsan 4
Member {ESXi hostname-fc1}
Exit
Zoneset activate name SP-Infra-A vsan 4
Zone commit vsan 4
Exit
Copy running-config startup-config
```

After these commands are saved you can add the vHBA WWPN to any volume initiator groups in the Nimble CS700 Array to grant access to the volumes.

Boot from SAN Benefits

Booting from SAN is another key feature which helps in moving towards stateless computing in which there is no static binding between a physical server and the OS/applications it is tasked to run. 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 were to be moved to another server, the pwwn of the HBAs and the Boot from SAN (BFS) policy also moves along with it. The new server now takes the same exact character of the old server, providing the true unique stateless nature of the Cisco UCS Blade Server.

The key benefits of booting from the network:

- **Reduce Server Footprints:** 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.
- **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 functionality of the servers at the primary site, the remote site can take over with minimal downtime.
- **Recovery from server failures 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 its image. As a result, boot from SAN can greatly reduce the time required for server recovery.
- **High Availability:** A typical data center is highly redundant in nature - redundant paths, redundant disks and redundant storage controllers. When operating system images are stored on disks in the SAN, it supports high availability and eliminates the potential for 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 only need to be in production for hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage 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.

With Boot from SAN, the image resides on a SAN LUN and the server communicates with the SAN through a host bus adapter (HBA). The HBAs BIOS contain the instructions that enable the server to find the boot disk. All FCoE-capable Converged Network Adapter (CNA) cards supported on Cisco UCS B-series blade servers support Boot from SAN.

After power on self-test (POST), the server hardware component fetches the boot device that is designated as the boot device in the hardware BIOS settings. When the hardware detects the boot device, it follows the regular boot process.



SAN Configuration on the Cisco MDS 9148 Switches

To configure the Cisco MDS 9148 switches for boot from SAN, we must identify the world wide port name for the Nimble CS700 Adaptive Array controllers. In this solution, the WWPN are as follows:




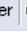

For Controller A:

Interface	Array Name	Controller	Link Status	Online	Speed	WWNN	WWPN
fc1.1	SP-NMBL	A		Yes	16 Gbps	56:c9:ce:90:0d:e8:24:00	56:c9:ce:90:0d:e8:24:01
fc2.1	SP-NMBL	A		Yes	16 Gbps	56:c9:ce:90:0d:e8:24:00	56:c9:ce:90:0d:e8:24:02

For Controller B:

fc1.1	SP-NMBL	B		Yes	16 Gbps	56:c9:ce:90:0d:e8:24:00	56:c9:ce:90:0d:e8:24:05
fc2.1	SP-NMBL	B		Yes	16 Gbps	56:c9:ce:90:0d:e8:24:00	56:c9:ce:90:0d:e8:24:06

When the WWPN are identified they can now be added to our Boot From SAN (BFS) Policy in Cisco UCS Manager.

Boot Order						
    						
Name	Order	vNIC/vHBA/iSCSI vNIC	Type	LUN Name	WWN	
Remote CD/DVD	1					
San	2					
SAN primary		fc0	Primary			
SAN Target primary			Primary	0	56:C9:CE:90:0D:E8:24:01	
SAN Target secondary			Secondary	0	56:C9:CE:90:0D:E8:24:02	
SAN secondary		fc1	Secondary			
SAN Target primary			Primary	0	56:C9:CE:90:0D:E8:24:05	
SAN Target secondary			Secondary	0	56:C9:CE:90:0D:E8:24:06	

When these steps are complete and the BFS policy is assigned to our Blade Servers, you will have an option to install ESXi onto a 10GB LUN presented from the Nimble Array.

Install and Configure ESXi 6 U1a

VMware ESXi 6.0

This section provides detailed instructions for installing VMware ESXi 6 Update1 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi 6 Update1

To download the Cisco Custom Image for ESXi 6 Update 1, complete the following steps:

1. Click the following link [vmware login page](#).
2. Type your email or customer number and the password and then click Log in.
3. Click on the following link [CiscoCustomImage6.0](#).
4. Click Download Now.
5. Save it to your destination folder.



This ESXi 6.0 Cisco custom image includes updates for the fnic and eNIC drivers. The versions that are part of this image are: eNIC: 2.1.2.59; fnic: 1.6.0.12

Install ESXi

To install VMware ESXi to the SAN-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the menu that is displayed.
2. After the installer is finished loading, press Enter to continue with the installation.
3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
4. Select the Nimble LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that existing partitions will be removed from the volume. Press F11 to continue with the installation.

8. After the installation is complete, clear the Mapped checkbox (located in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.



The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

9. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Because the media cannot be ejected and it is read-only, simply click Yes to unmap the image.
10. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host.

To configure the ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the `<<var_ib-mgmt_vlan_id>>` and press Enter.
6. From the Configure Management Network menu, select IP Configuration and press Enter.
7. Select the Set Static IP Address and Network Configuration option by using the space bar.
8. Enter the IP address for managing the first ESXi host: `<<var_vm_host_01_ip>>`.
9. Enter the subnet mask for the first ESXi host.
10. Enter the default gateway for the first ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Select the IPv6 Configuration option and press Enter.
13. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.
14. Select the DNS Configuration option and press Enter.



Because the IP address is assigned manually, the DNS information must also be entered manually.

15. Enter the IP address of the primary DNS server.

16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the fully qualified domain name (FQDN) for the first ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host reboots. After reboot, press F2 and log back in as root.
22. Select Test Management Network to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.
24. Press Enter to exit the window.
25. Press Esc to log out of the VMware console.

<p>Troubleshooting Mode Options</p> <p>Disable ESXi Shell Disable SSH Modify ESXi Shell and SSH timeouts Modify DCUI idle timeout Restart Management Agents</p>	<p>ESXi Shell</p> <p>ESXi Shell is Enabled</p> <p>Change current state of the ESXi Shell</p>
<p>Troubleshooting Mode Options</p> <p>Disable ESXi Shell Disable SSH Modify ESXi Shell and SSH timeouts Modify DCUI idle timeout Restart Management Agents</p>	<p>SSH Support</p> <p>SSH is Enabled</p> <p>Change current state of SSH</p>
<p>Configure Management Network</p> <p>Network Adapters VLAN (optional)</p> <p>IPv4 Configuration IPv6 Configuration DNS Configuration Custom DNS Suffixes</p>	<p>Network Adapters</p> <p>vmnic0 (MLOM Slot; relative bdf 03:00.0) vmnic1 (MLOM Slot; relative bdf 04:00.0)</p> <p>The adapters listed here provide the default network connection to and from this host. When two or more adapters are used, connections will be fault-tolerant and outgoing traffic will be load-balanced.</p>
<p>Configure Management Network</p> <p>Network Adapters VLAN (optional)</p> <p>IPv4 Configuration IPv6 Configuration DNS Configuration Custom DNS Suffixes</p>	<p>VLAN (optional)</p> <p>60</p> <p>A VLAN is a virtual network within a physical network. Because several VLANs can co-exist on the same physical network segment, VLAN configuration and partitioning is often more flexible, better isolated, and less expensive than flat networks based on traditional physical topology.</p> <p>If you are unsure how to configure or use a VLAN, it is safe to leave this option unset.</p>

<p>Configure Management Network</p> <hr/> <p>Network Adapters VLAN (optional)</p> <p>IPv4 Configuration</p> <p>IPv6 Configuration DNS Configuration Custom DNS Suffixes</p>	<p>IPv4 Configuration</p> <hr/> <p>Manual</p> <p>IPv4 Address: 10.10.60.100 Subnet Mask: 255.255.255.0 Default Gateway: 10.10.60.1</p> <p>This host can obtain an IPv4 address and other networking parameters automatically if your network includes a DHCP server. If not, ask your network administrator for the appropriate settings.</p>
<p>Configure Management Network</p> <hr/> <p>Network Adapters VLAN (optional)</p> <p>IPv4 Configuration IPv6 Configuration DNS Configuration Custom DNS Suffixes</p>	<p>IPv6 Configuration</p> <hr/> <p>IPv6 is disabled.</p> <p>This host can be configured to support IPv6. A restart of the host will be required to enable or disable IPv6.</p>
<p>Configure Management Network</p> <hr/> <p>Network Adapters VLAN (optional)</p> <p>IPv4 Configuration IPv6 Configuration DNS Configuration Custom DNS Suffixes</p>	<p>DNS Configuration</p> <hr/> <p>Manual</p> <p>Primary DNS Server: 10.10.61.30 Alternate DNS Server: 10.10.61.31</p> <p>Hostname C1-Blade1</p>
<p>Configure Management Network</p> <hr/> <p>Network Adapters VLAN (optional)</p> <p>IPv4 Configuration IPv6 Configuration DNS Configuration Custom DNS Suffixes</p>	<p>Custom DNS Suffixes</p> <hr/> <p>sp.local</p> <p>When using short, unqualified names, DNS queries will attempt to locate the specified host by appending the suffixes listed here in the order shown until a match is found or the list is exhausted.</p> <p>If no suffixes are specified here, a default suffix list is derived from the local domain name.</p>

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-01 management IP address.
2. Download and install the vSphere Client.



This application is downloaded from the VMware website and Internet access is required on the management workstation.

Download VMware vSphere CLI 6

To download VMware vSphere CLI 6, complete the following steps:

1. Click the following link [VMware vSphere CLI 6.0](#)

Install and Configure ESXi 6 U1a

2. Select your OS and Click Download.
3. Save it to your destination folder.
4. Run the VMware-vSphere-CLI.exe
5. Click Next.
6. Accept the terms for the license and click Next.
7. Click Next on the Destination Folder screen.
8. Click Instal.
9. Click Finish.



Install VMware vSphere CLI 6.0 on the management workstation.

Log in to VMware ESXi Hosts by Using VMware vSphere Client

ESXi Host VM-Host-01

To log in to the `VM-Host-01` ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-01` as the host you are trying to connect to: `<<var_vm_host_01_ip>>`.
2. Enter `root` for the user name.
3. Enter the root password.
4. Click Login to connect.


Install and Configure vCenter 6

Install and Configure VMware vCenter Appliance


To build the VMWare vCenter VM, complete the following steps:

1. From the vSphere 6 download page on the VMware Web site, download the vCenter ISO file for the vCenter Server appliance onto your system.
2. Open the vSphere ISO via Windows Explorer and double-click the `vcsa-setup.htm` file

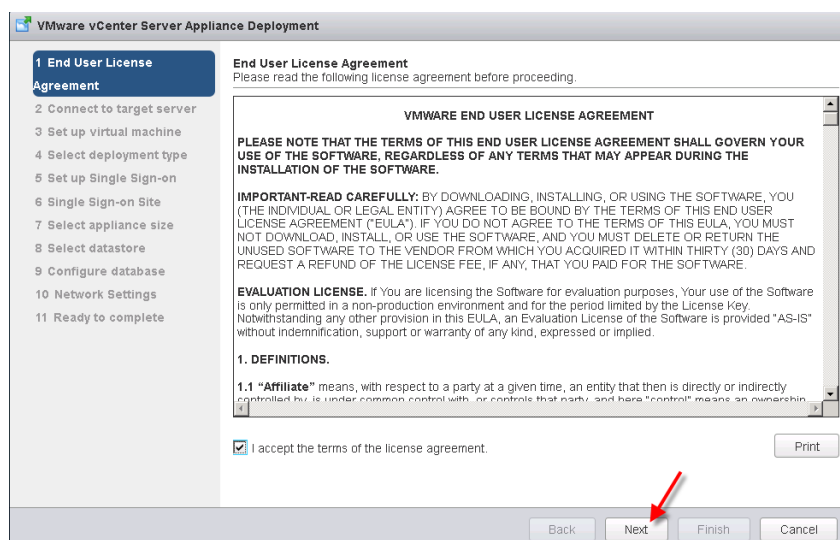
Install and Configure ESXi 6 U1a



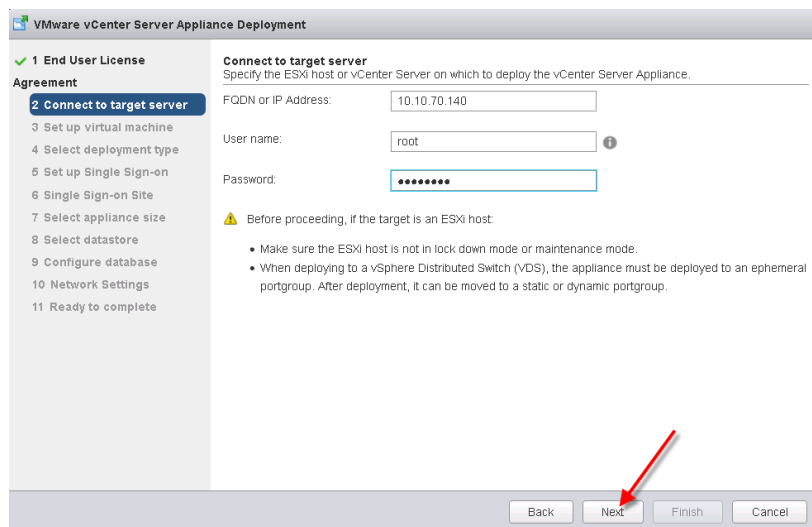
A browser will open with an option to Install.



3. Follow the onscreen prompts. Accept EULA.



4. Enter the IP of the ESXi host the vCenter Appliance will reside. Click Next.



The screenshot shows the 'VMware vCenter Server Appliance Deployment' wizard. On the left, a list of steps is shown, with '2 Connect to target server' selected. The main area is titled 'Connect to target server' and contains the following fields: 'FQDN or IP Address' with the value '10.10.70.140', 'User name' with the value 'root', and 'Password' with masked characters. Below these fields, a warning icon and text state: 'Before proceeding, if the target is an ESXi host:'. This is followed by two bullet points: 'Make sure the ESXi host is not in lock down mode or maintenance mode.' and 'When deploying to a vSphere Distributed Switch (vDS), the appliance must be deployed to an ephemeral portgroup. After deployment, it can be moved to a static or dynamic portgroup.' At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'. A red arrow points to the 'Next' button.

5. Click Yes to accept Certificate Warning.

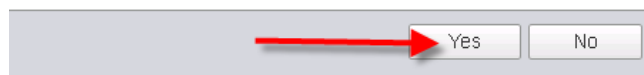
Certificate Warning

An untrusted SSL certificate is installed on 10.10.70.140 and secure communication cannot be guaranteed. Depending on your security policy, this issue might not represent a security concern.

The SHA1 thumbprint of the certificate is:

60:F5:06:FB:7F:82:A4:BD:DD:BD:63:6F:4E:6F:0F:FA:9E:2D:88:CC

To accept and continue, press Yes



The screenshot shows the bottom of the 'Certificate Warning' dialog box. It contains two buttons: 'Yes' and 'No'. A red arrow points to the 'Yes' button.

6. Click Yes.
7. Provide a name for the vCenter appliance, then click Next to continue.

VMware vCenter Server Appliance Deployment

✓ 1 End User License Agreement
 ✓ 2 Connect to target server
3 Set up virtual machine
 4 Select deployment type
 5 Set up Single Sign-on
 6 Single Sign-on Site
 7 Select appliance size
 8 Select datastore
 9 Configure database
 10 Network Settings
 11 Ready to complete

Set up virtual machine
 Specify virtual machine settings for the vCenter Server Appliance to be deployed.

Appliance name: ⓘ

OS user name: root

OS password: ⓘ

Confirm OS password:

Back Next Finish Cancel

8. Select Install vCenter Server with and Embedded Platform Services Controller (unless your environment already has a PSC)

VMware vCenter Server Appliance Deployment

✓ 1 End User License Agreement
 ✓ 2 Connect to target server
 ✓ 3 Set up virtual machine
4 Select deployment type
 5 Set up Single Sign-on
 6 Single Sign-on Site
 7 Select appliance size
 8 Select datastore
 9 Configure database
 10 Network Settings
 11 Ready to complete

Select deployment type
 Select the services to deploy onto this appliance.

vCenter Server 6.0 requires a Platform Services Controller, which contains shared services such as Single Sign-On, Licensing, and Certificate Management. An embedded Platform Services Controller is deployed on the same Appliance VM as vCenter Server. An external Platform Services Controller is deployed in a separate Appliance VM. For smaller installations, consider vCenter Server with an embedded Platform Services Controller. For larger installations with multiple vCenter Servers, consider one or more external Platform Services Controllers. Refer to the vCenter Server documentation for more information.

Note: Once you install vCenter Server, you can only change from an embedded to an external Platform Services Controller with a fresh install.

Embedded Platform Services Controller

☒ Install vCenter Server with an Embedded Platform Services Controller

External Platform Services Controller

☐ Install Platform Services Controller
☐ Install vCenter Server (Requires External Platform Services Controller)

Diagram illustrating the deployment options:

- Embedded Platform Services Controller:** A single box labeled "VM or Host" containing "Platform Services Controller" and "vCenter Server".
- External Platform Services Controller:** A box labeled "VM or Host" containing "Platform Services Controller" is connected to two separate boxes labeled "VM or Host", each containing "vCenter Server".

Back Next Finish Cancel

9. Create a new SSO domain (unless your environment already has and SSO domain. Multiple SSO domains can co-exist).

Install and Configure ESXi 6 U1a

VMware vCenter Server Appliance Deployment

✓ 1 End User License Agreement
✓ 2 Connect to target server
✓ 3 Set up virtual machine
✓ 4 Select deployment type
5 Set up Single Sign-on
6 Select appliance size
7 Select datastore
8 Configure database
9 Network Settings
10 Ready to complete

Set up Single Sign-on (SSO)
Create or join a SSO domain. An SSO configuration cannot be changed after deployment.

☒ Create a new SSO domain
☐ Join an SSO domain in an existing vCenter 6.0 platform services controller

vCenter SSO User name: administrator

vCenter SSO Password: •••••••• ⓘ

Confirm password: ••••••••

SSO Domain name: vsphere.local ⓘ

SSO Site name: Example: Default-First-Site ⓘ

⚠ Before proceeding, make sure that the vCenter Single Sign-On domain name used is different than your Active Directory domain name.

Back Next Finish Cancel

10. Select the proper appliance size for your deployment. In our study, Medium was sufficient.

VMware vCenter Server Appliance Deployment

✓ 1 End User License Agreement
✓ 2 Connect to target server
✓ 3 Set up virtual machine
✓ 4 Select deployment type
✓ 5 Set up Single Sign-on
6 Select appliance size
7 Select datastore
8 Configure database
9 Network Settings
10 Ready to complete

Select appliance size
Specify a deployment size for the new appliance

Appliance size: Medium (up to 400 hosts, 4,000 VMs)

Description:
This will deploy a Medium VM configured with 8 vCPUs and 24 GB of memory and requires 300 GB of disk space. This option contains vCenter Server with an embedded Platform Services Controller.

Back Next Finish Cancel

11. In our study we used the embedded PostgreSQL database.

The screenshot shows the 'VMware vCenter Server Appliance Deployment' wizard at step 8, 'Configure database'. The left sidebar lists steps 1 through 10, with step 8 highlighted. The main area is titled 'Configure database' and contains the instruction 'Configure the database for this deployment'. There are two radio button options: 'Use an embedded database (PostgreSQL)' (which is selected) and 'Use Oracle database'. At the bottom, there are four buttons: 'Back', 'Next' (highlighted with a blue border), 'Finish', and 'Cancel'.

12. Enter Network Settings for appliance.

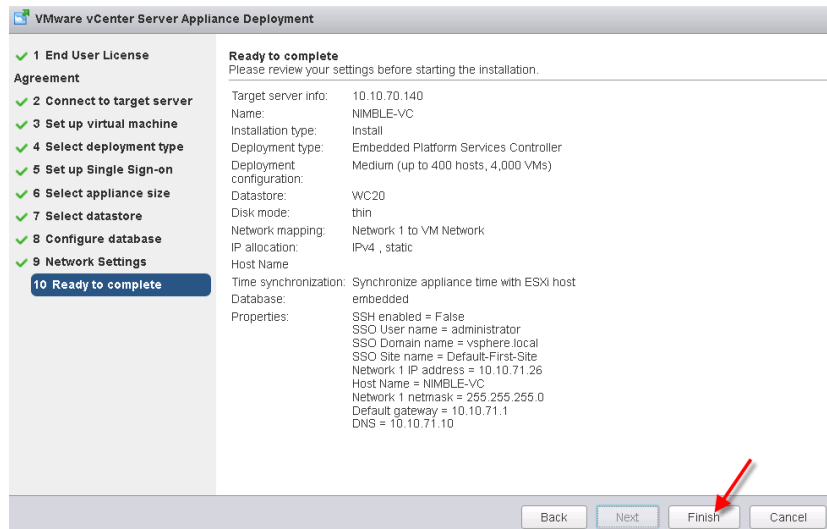


It is important to note at this step that you should create a DNS A record for your appliance prior to running the install. The services will fail to startup and your install will fail if it cannot resolve properly.

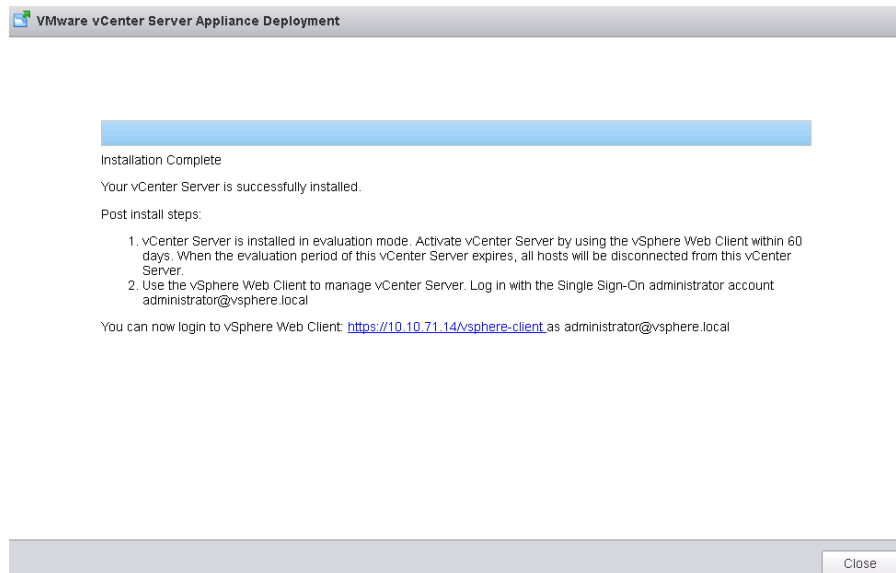
The screenshot shows the 'VMware vCenter Server Appliance Deployment' wizard at step 9, 'Network Settings'. The left sidebar lists steps 1 through 10, with step 9 highlighted. The main area is titled 'Network Settings' and contains the instruction 'Configure network settings for this deployment'. It includes several input fields: 'Choose a network' (dropdown menu showing 'VM Network'), 'IP address family' (dropdown menu showing 'IPv4'), 'Network type' (dropdown menu showing 'static'), 'Network address' (text field with '10.10.71.26'), 'System name [FQDN or IP address]' (text field with 'NIMBLE-VC'), 'Subnet mask' (text field with '255.255.255.0'), 'Network gateway' (text field with '10.10.71.1'), and 'Network DNS Servers (separated by commas)' (text field with '10.10.71.10'). At the bottom, there are two radio button options for 'Configure time sync': 'Synchronize appliance time with ESXi host' (unselected) and 'Use NTP servers (Separated by commas)' (selected). At the bottom right, there are four buttons: 'Back', 'Next' (highlighted with a blue border), 'Finish', and 'Cancel'.

13. Review the Install Settings and click Finish.

Install and Configure ESXi 6 U1a

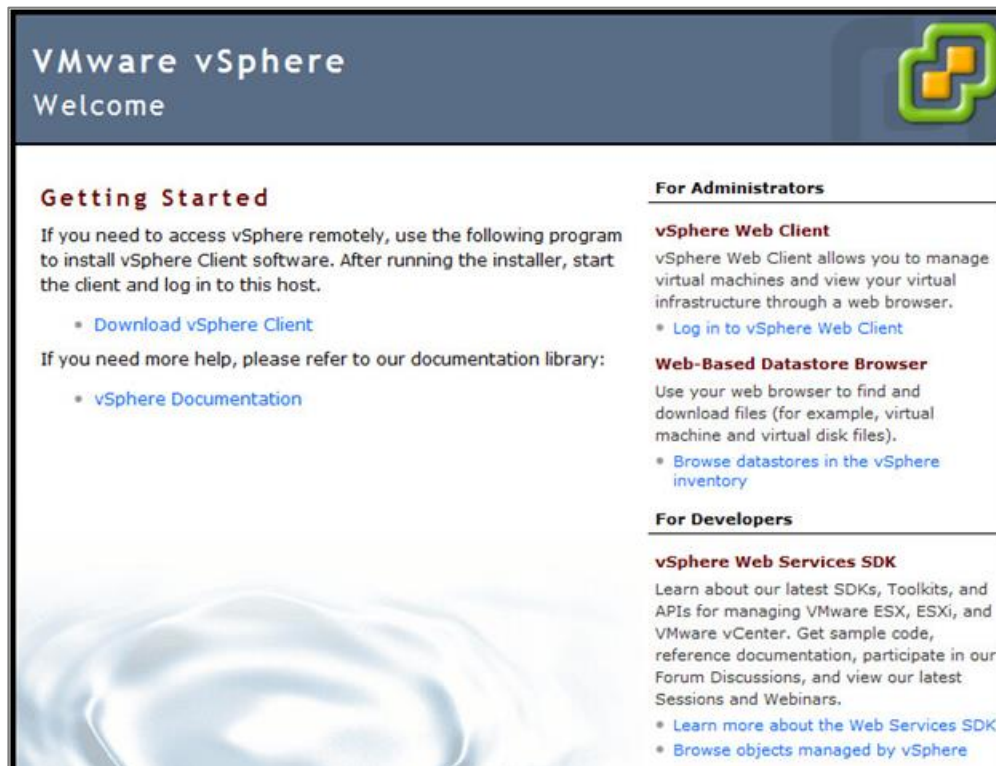


14. When your install completes successfully, you can now login to your Web Client and begin adding hosts and configuring clusters.

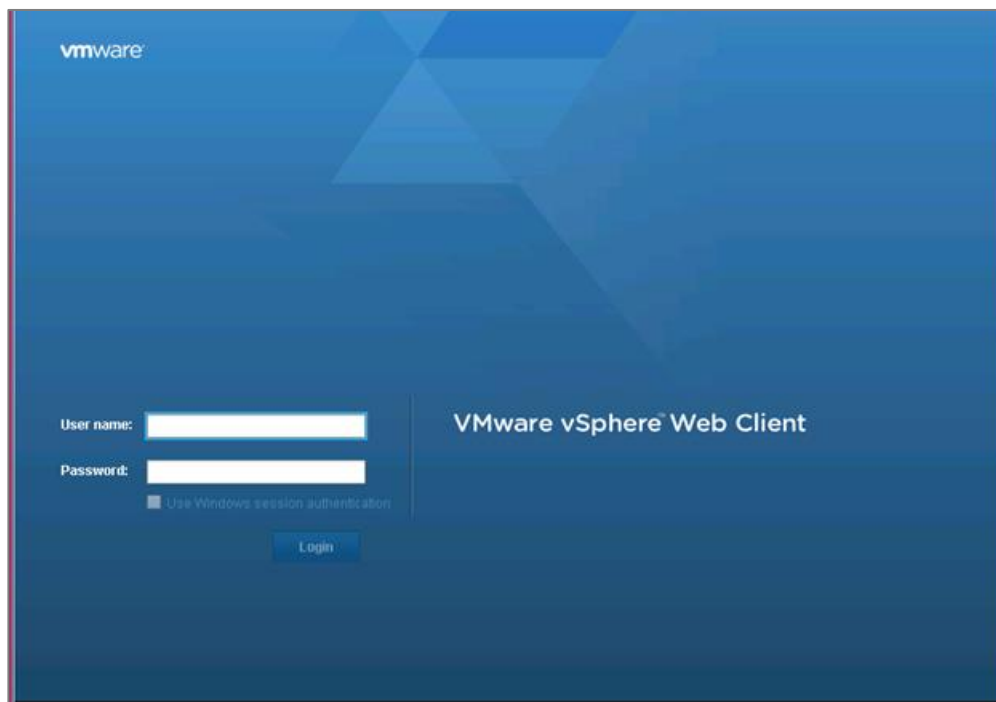


15. Log into the vSphere Web Client.

16. Using a web browser, navigate to `https://<var_vcenter_ip>`.



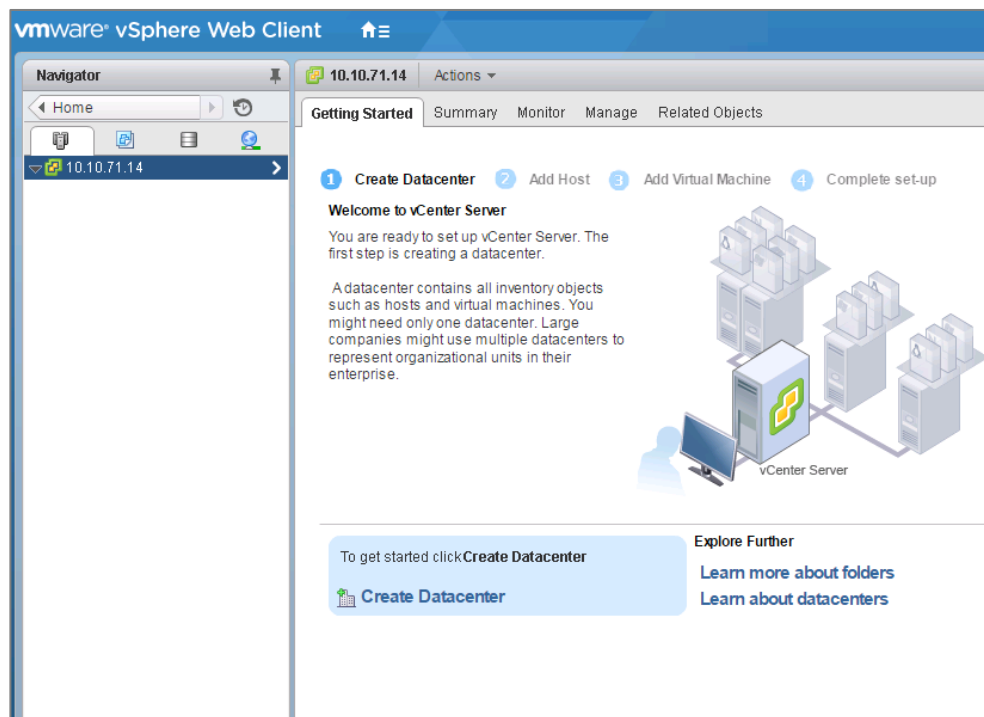
17. Click the link labeled Log in to vSphere Web Client.



18. If prompted, run the VMWare Remote Console Plug-in.

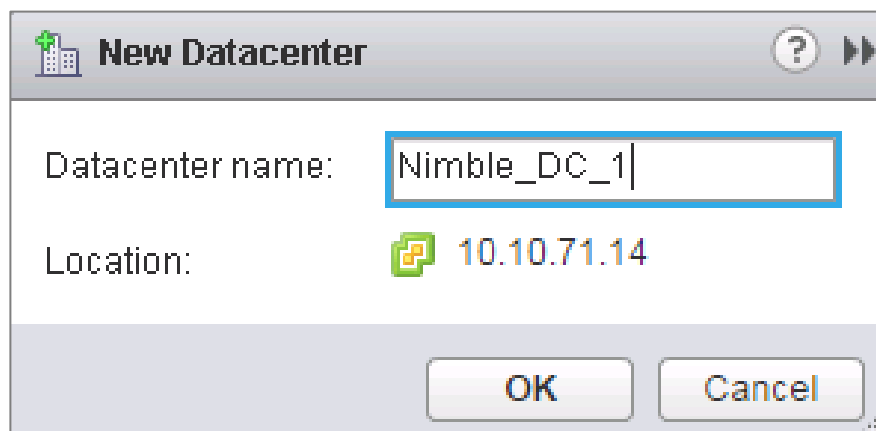
19. Log in using the root user name and password.

20. Click the vCenter link on the left panel.



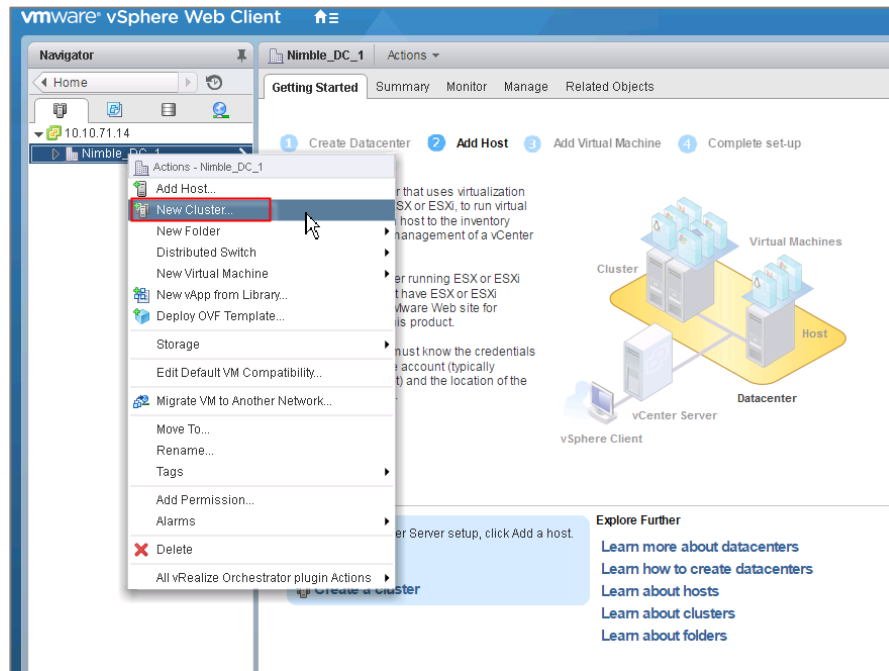
21. Click the Datacenters link on the left panel.

22. To create a Datacenter, click the icon in the center pane which has the green plus symbol above it.



23. Type Nimble_DC_1 as the Datacenter name.

24. Click the vCenter server available in the list. Click OK to continue.



25. Right-click Datacenters > Nimble_DC_1 in the list in the center pane, then click New Cluster.
26. Name the cluster Nimble_Infrastructure.
27. Select DRS. Retain the default values.
28. Select vSphere HA. Retain the default values.

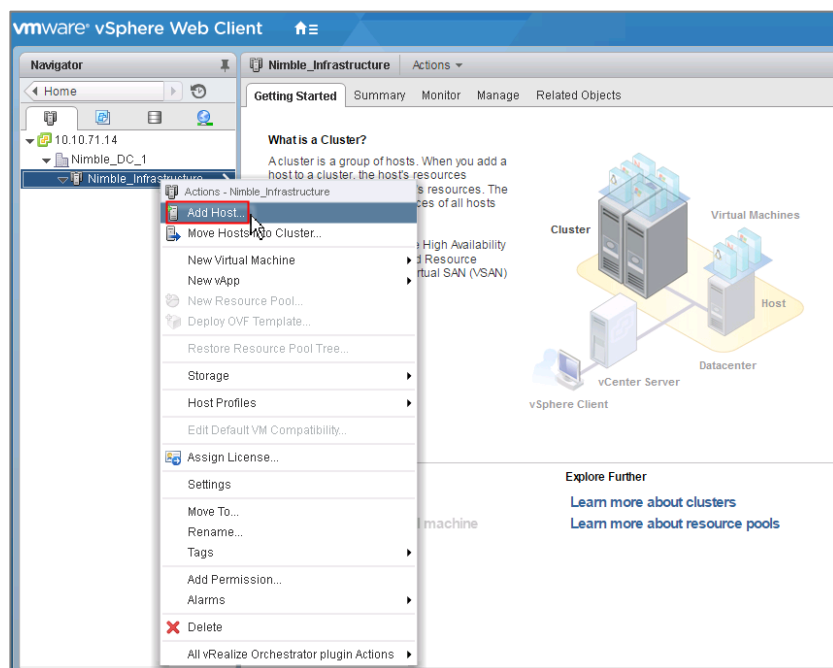
Name	Nimble_Infrastructure
Location	Nimble_DC_1
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ——— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	<input checked="" type="checkbox"/> Enable admission control Admission control will prevent powering on VMs that violate availability constraints
Policy	Specify the type of the policy that admission control should enforce. <input checked="" type="radio"/> Host failures cluster tolerates: 1 <input type="radio"/> Percentage of cluster resources reserved as failover spare capacity: Reserved failover CPU capacity: 25 % CPU Reserved failover Memory capacity: 25 % Memory
VM Monitoring	<input type="checkbox"/> Turn ON <input checked="" type="checkbox"/> Turn OFF
Monitoring Sensitivity	Low ——— High
EVC	Intel® "Ivy Bridge" Generation
Virtual SAN	<input type="checkbox"/> Turn ON



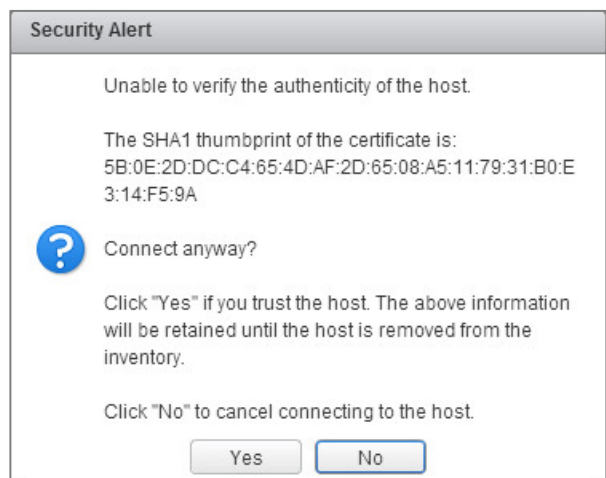
If mixing Cisco UCS B M3 and M4 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, refer to Enhanced vMotion Compatibility (EVC) Processor Support.

29. Click OK to create the new cluster.

30. Click Nimble_DC_1 in the left pane.



31. Right-click Nmble_Infrastructure in the center pane and click Add Host.
32. Type <<var_esx_host_1_ip>> and click Next.
33. Type root as the user name and <<var_esx_host_password>> as the password. Click Next to Continue.




34. Click Yes to accept the certificate.
35. Review the host details, and click Next to continue.
36. Assign a license, and click Next to continue.
37. Click Next to continue.
38. Click Next to continue.

39. Review the configuration parameters then click Finish to add the host.
40. Repeat this for the other hosts.

Install the Nimble Connection Manager

In order to utilize the Nimble provided PSP to manage the paths ESXi will use for the Nimble Volumes we need to install the Nimble Connection Manager VIB.



This can be installed using the **esxcli** command or through VMware's Update Manager. We used Update Manager to deploy this VIB to our hosts.

1. Put ESXi host into maintenance mode.
2. SSH into ESXi host.
3. run command "esxcli software vib install -d <http://update.nimblestorage.com/esx6/ncm>" which will automatically download and install the files for you.
4. Take ESXi host out of maintenance mode.

Update Manager Administration for sp-vdi-vc.sp.local			
Getting Started	Baselines and Groups	Configuration	Events
Notifications	Patch Repository	ESXi Images	VA Upgrades
Patch Name	Product	Vendor	Patch ID
Cisco Nexus 1000V 5.2(1)SV3(1.5a)	embeddedEsx 6.0.0	Cisco Systems, Inc.	VEM600-201508198119-BG
Cisco Nexus 1000V 5.2(1)SV3(1.5a)	embeddedEsx 5.0.0	Cisco Systems, Inc.	VEM500-201508198101-BG
Cisco Nexus 1000V 5.2(1)SV3(1.5a)	embeddedEsx 5.1.0	Cisco Systems, Inc.	VEM510-201508198107-BG
Cisco Nexus 1000V 5.2(1)SV3(1.5a)	embeddedEsx 5.5.0	Cisco Systems, Inc.	VEM550-201508198113-BG
Nimble Connection Manager	embeddedEsx 6.0.0, esx 6.0.0	Nimble Storage	nimble-ncm-2.3.1-600006
Version 346.68	embeddedEsx 6.0	NVIDIA	NVD.NVIDIA_bootbank_NVIDIA-vgx-VMware_346.68-10EM.600.0.0.2494585
Version 352.54	embeddedEsx 6.0	NVIDIA	NVD.NVIDIA_bootbank_NVIDIA-kepler-VMware_352.54-10EM.600.0.0.2494585
Version 346.69	embeddedEsx 6.0	NVIDIA	NVD.NVIDIA_bootbank_NVIDIA-VMware_346.69-10EM.600.0.0.2494585
Version 346.42	embeddedEsx 6.0	NVIDIA	NVD.NVIDIA_bootbank_NVIDIA-vgx-VMware_346.42-10EM.600.0.0.2159203
Version 352.83	embeddedEsx 6.0	NVIDIA	NVD.NVIDIA_bootbank_NVIDIA-vGPU-VMware_352.83-10EM.600.0.0.2494585

Building the Virtual Machines and Environment

Software Infrastructure Configuration

This section details the configuration for the software infrastructure components that comprise the solution.

Install and configure the infrastructure virtual machines following the guidance provided in Table 4 .

Table 4 Test Infrastructure Virtual Machine Configuration

Configuration	Citrix XenDesktop Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Operating system	Microsoft Windows Server 2012 R2	Microsoft Windows Server 2012 R2
Virtual CPU amount	4	4
Memory amount	8 GB	12 GB
Network	VMXNET3 Infrastructure vLAN	VMXNET3 Infrastructure vLAN
Disk-1 (OS) size and location	40 GB Infra-DS volume	40 GB Infra-DS volume
Disk-2 size and location	–	500 GB PVS-vDisk volume using CIFS
Configuration	Microsoft Active Directory DCs Virtual Machines	Citrix Licensing Virtual Machines
Operating system	Microsoft Windows Server 2012 R2	Microsoft Windows Server 2012 R2

Configuration	Citrix XenDesktop Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Virtual CPU amount	4	2
Memory amount	4 GB	4 GB
Network	VMXNET3 Infrastructure vLAN	VMXNET3 Infrastructure vLAN
Disk size and location	40 GB Infra-DS volume	40 GB Infra-DS volume
Configuration	Microsoft SQL Server Virtual Machine	Citrix Storefront Virtual Ma- chines
Operating system	Microsoft Windows Server 2012 R2 Microsoft SQL Server 2012 SP1	Microsoft Windows Server 2012 R2
Virtual CPU amount	4	2
Memory amount	12 GB	4 GB
Network	VMXNET3 Infrastructure vLAN	VMXNET3 Infrastructure vLAN
Disk-1 (OS) size and location	60 GB Infra-DS volume	40 GB Infra-DS volume
Disk-2 size and location	-	
Disk-3 size and location	-	

Install and Configure VSUM and Nexus 1000v

Install Cisco Virtual Switch Update Manager

Verifying the Authenticity of the Cisco-Signed Image (Optional)

Before you install the Nexus1000v-vsum.1.5.x-pkg.zip image, you have the option to validate its authenticity. In the zip file, there is a signature.txt file that contains an SHA-512 signature and an executable script that can be used to verify the authenticity of the Nexus1000v-vsum.1.5.x-pkg.zip image.

To set up the primary Cisco Nexus 1000V VSM on the Cisco Nexus 1110-X A, complete the following steps:



Verifying the authenticity of an image is optional. You can still install the image without validating its authenticity.

1. Copy the following files to a directory on the Linux machine:

- Nexus1000v-vsum.1.5.x-pkg.zip image
- signature.txt file
- cisco_n1k_image_validation_v_1_5_x script

2. Make sure the script is executable.

```
chmod 755 cisco_n1k_image_validation_v_1_5_x
```

3. Run the script.

```
./cisco_n1k_image_validation_v_1_5_x -s signature.txt Nexus1000v-vsum.1.5.x-pkg.zip
```

4. Run the script.

```
./cisco_n1k_image_validation_v_1_5_x -s signature.txt Nexus1000v-vsum.1.5.x-pkg.zip
```

5. Check the output. If the validation is successful, the following message displays:

```
Authenticity of Cisco-signed image Nexus1000v-vsum.1.5.x-pkg.zip has been successfully verified!
```

Install Cisco Virtual Switch Update Manager

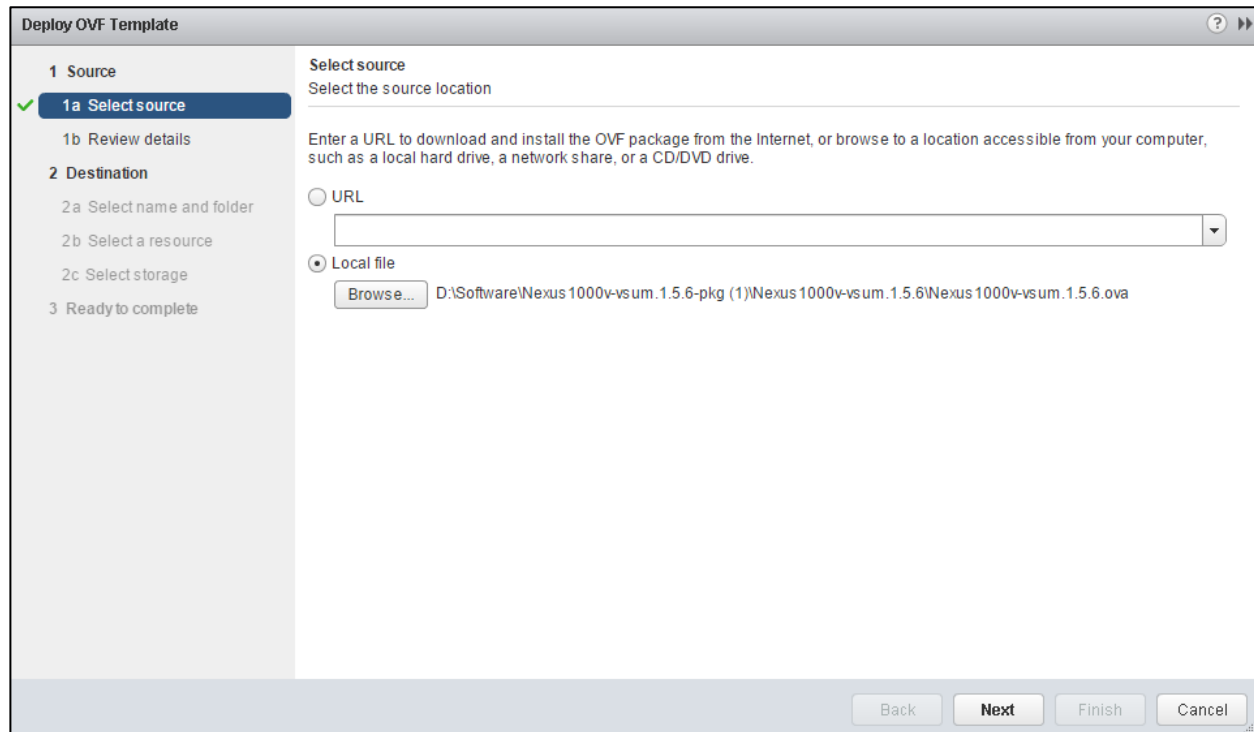
VMware vSphere Web Client

To install the Cisco Virtual Switch Upgrade Manager from OVA in the VMware virtual environment, complete the following steps:

1. Log into the VMware vSphere Web Client.
2. In the pane on the right, click VMs and Templates.
3. In the center pane, select Actions > Deploy OVF Template.
4. Select Browse and browse to and select the Nexus1000v-vsum.1.5.x.ova file.

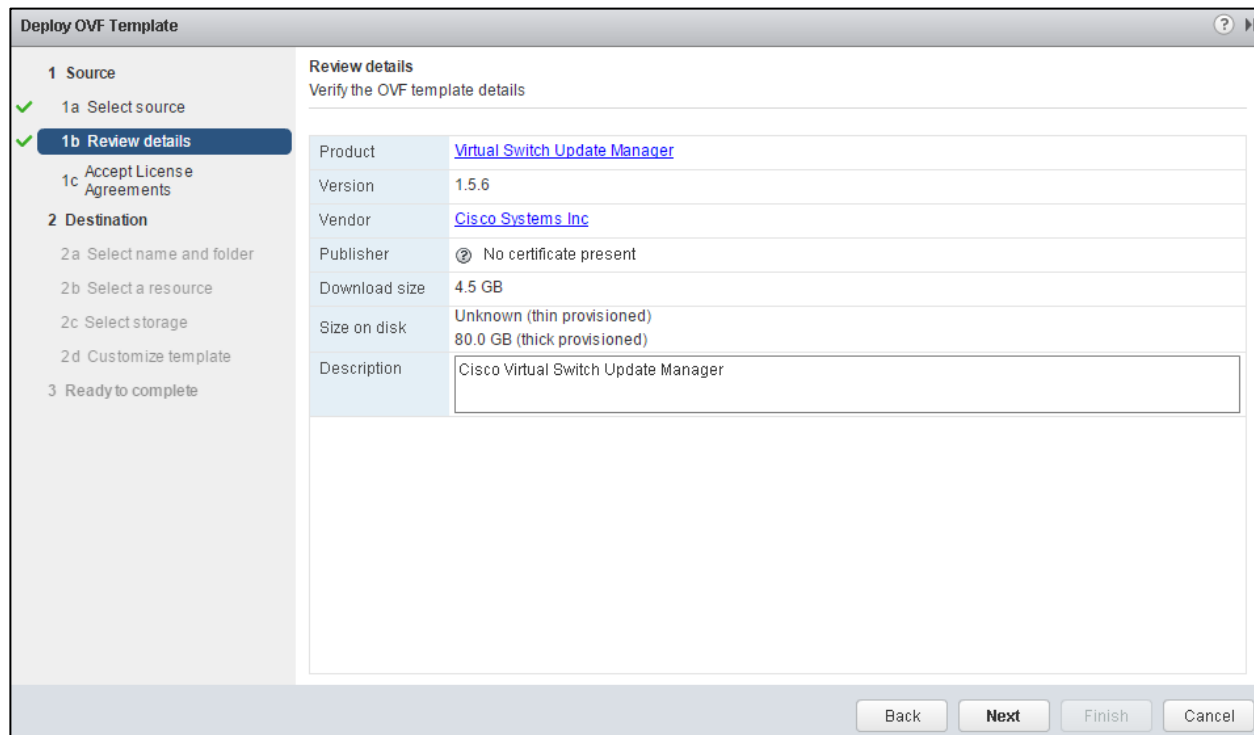
Install and Configure ESXi 6 U1a

5. Click Open.
6. Click Next.



The screenshot shows the 'Deploy OVF Template' wizard. The left sidebar has a tree view with '1 Source' expanded, showing '1a Select source' (checked with a green checkmark) and '1b Review details'. Below '1 Source' is '2 Destination' with sub-items '2a Select name and folder', '2b Select a resource', '2c Select storage', and '3 Ready to complete'. The main area is titled 'Select source' and 'Select the source location'. It contains instructions: 'Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.' There are two radio buttons: 'URL' (unselected) and 'Local file' (selected). Below 'Local file' is a 'Browse...' button and a text field containing the path 'D:\Software\Nexus 1000v-vsum.1.5.6-pkg (1)\Nexus 1000v-vsum.1.5.6\Nexus 1000v-vsum.1.5.6.ova'. At the bottom are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

7. Review the details and click Next.

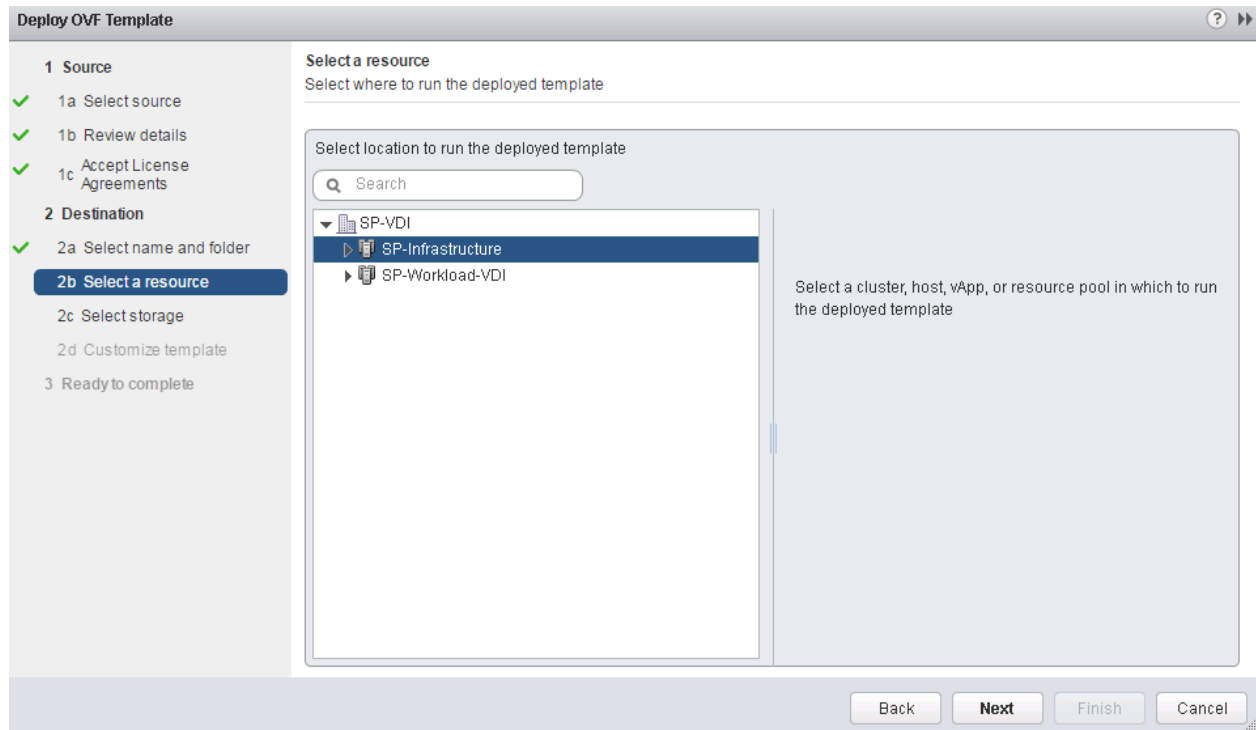


The screenshot shows the 'Deploy OVF Template' wizard at the 'Review details' step. The left sidebar shows '1a Select source' and '1b Review details' (checked with a green checkmark). The main area is titled 'Review details' and 'Verify the OVF template details'. It contains a table with the following information:

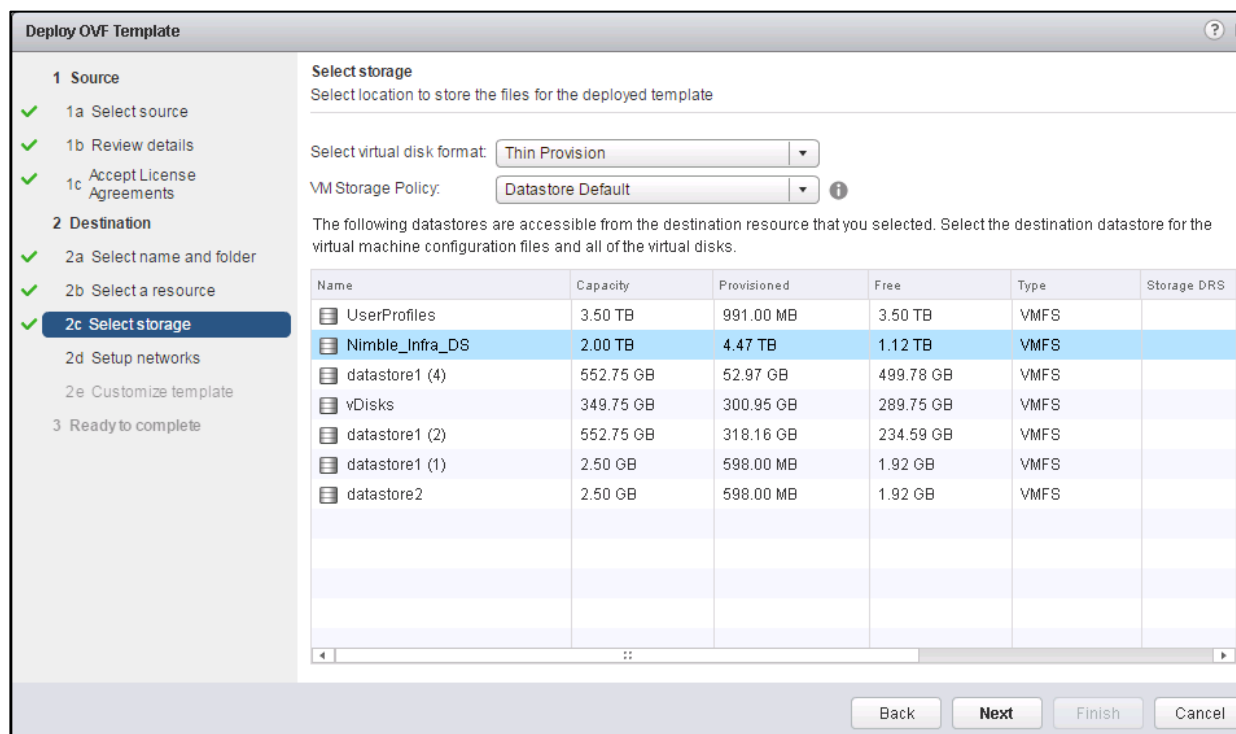
Product	Virtual Switch Update Manager
Version	1.5.6
Vendor	Cisco Systems Inc
Publisher	Ⓢ No certificate present
Download size	4.5 GB
Size on disk	Unknown (thin provisioned) 80.0 GB (thick provisioned)
Description	Cisco Virtual Switch Update Manager

At the bottom are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

8. Click Accept to accept the License Agreement and click Next.
9. Name the Virtual Machine, select the VSphere_DC datacenter and click Next.



10. Select the SP-Infrastructure cluster and click Next.
11. Select Nimble_Infra_DS and the Thin Provision virtual disk format and click Next.



12. Select the MGMT Network and click Next.
13. Fill in the Networking Properties.
14. Expand the vCenter Properties and fill those in.
15. Click Next.
16. Review all settings and click Finish.
17. Wait for the Deploy OVF template task to complete.
18. Select the Home button in VMware vSphere Web Client and select Hosts and Clusters.
19. Expand the Infrastructure cluster and select the Virtual Switch Update Manager VM.
20. In the center pane, select Launch Remote Console. If a security warning pops up, click Allow.
21. If a security certificate warning pops up, click Connect Anyway.
22. Power on the Virtual Switch Update Manager VM.
23. When the VM has completely booted up, log out and log back into the VMware vSphere Web Client.

Deploy OVF Template

1 Source

- 1a Select source
- 1b Review details
- 1c Accept License Agreements

2 Destination

- 2a Select name and folder
- 2b Select storage
- 2c Setup networks
- 2d Customize template**

3 Ready to complete

Customize template
Customize the deployment properties of this software solution

All properties have valid values [Show next...](#) [Collapse all...](#)

Default Gateway: Gateway IP for the management interface (e.g., 192.168.0.1)
10.10.80.1

DNS Server 1: The domain name server IP. Optional. Needed to resolve vCenter's FQDN if entered.
10.10.80.10

DNS Server 2: Secondary DNS Server IP (e.g., 10.10.10.10). Optional.

vCenter Properties: 5 settings

IP Address or FQDN (Fully Qualified Domain Name): The IP address or FQDN (e.g., foo.example.com) of the vCenter to register with.
10.10.80.26

Username: vCenter username. User must be able to manage extensions.
administrator@vsphere.local

Password: Password for the above username.
Enter password:
Confirm password:

Back Next Finish Cancel

About the Cisco VSUM GUI

The following lists the details of the Cisco VSUM GUI:

- Cisco VSUM is a virtual appliance that is registered as a plug-in to the VMware vCenter Server.
- The Cisco VSUM is the GUI that you use to install, migrate, monitor, and upgrade the VSMs in high availability (HA) or standalone mode and the VEMs on ESX/ESXi hosts.

Figure 18 VMware vSphere Web Client–Home Page

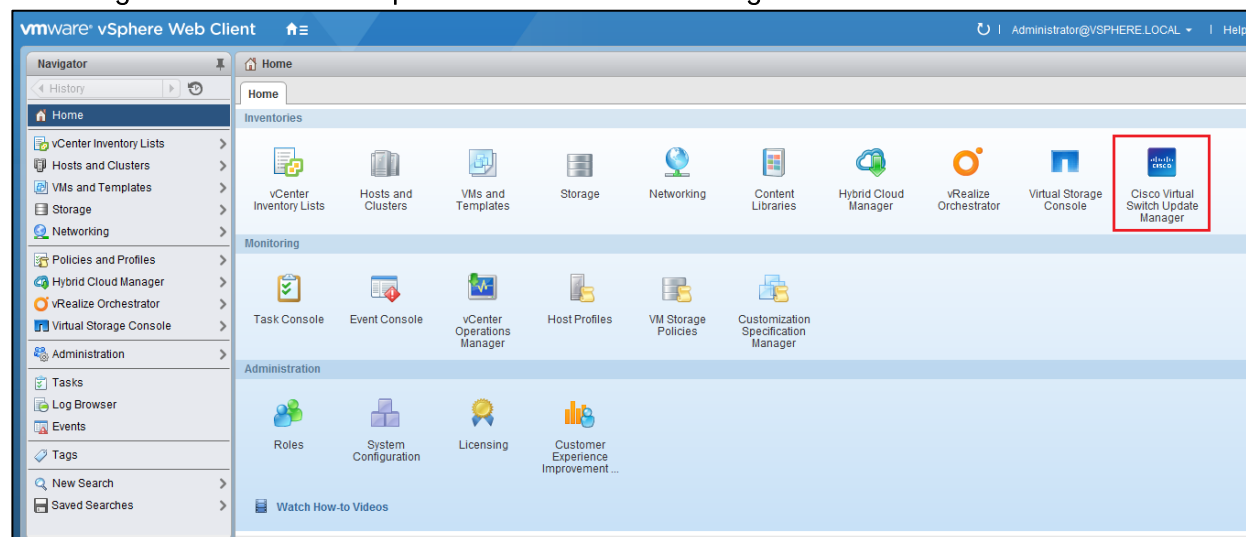
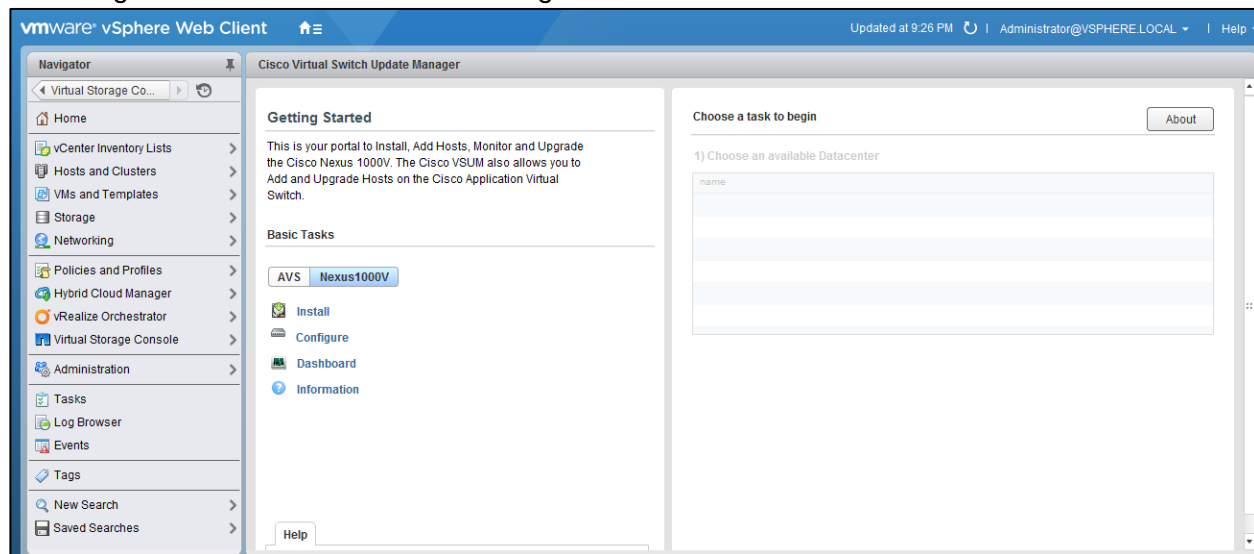


Figure 19 Cisco VSUM—Home Page



Install Cisco Nexus 1000V using Cisco VSUM

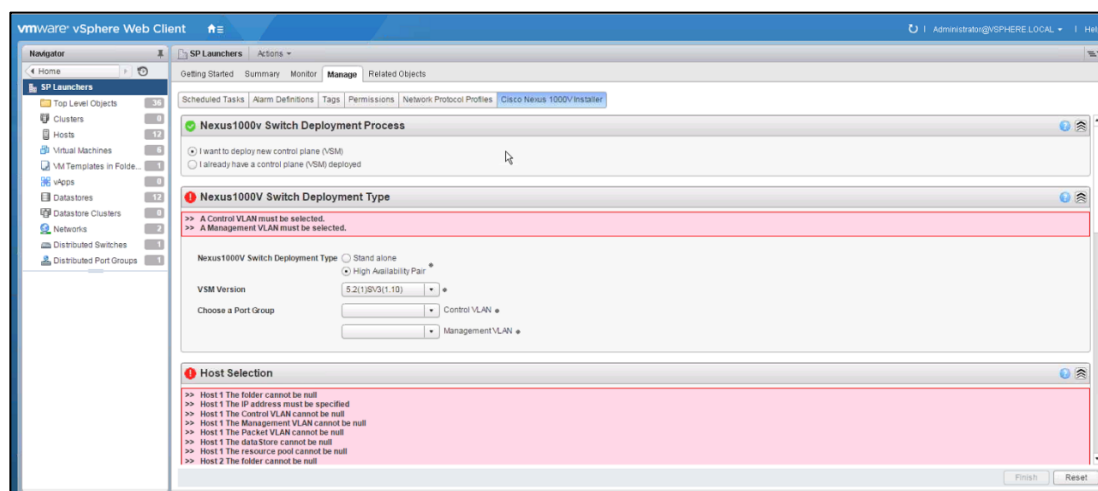
VMware vSphere Web Client

To install the Cisco Nexus 1000V switch by creating a new VSM, complete the following steps:



Optionally, an existing VSM can be used that is provided by a Cisco Nexus Cloud Services Platform (CSP).

1. Log in to VMware vSphere Web Client and choose Home > Cisco Virtual Switch Update Manager > Nexus 1000V > Install, and then choose the data center. The installation screen appears.



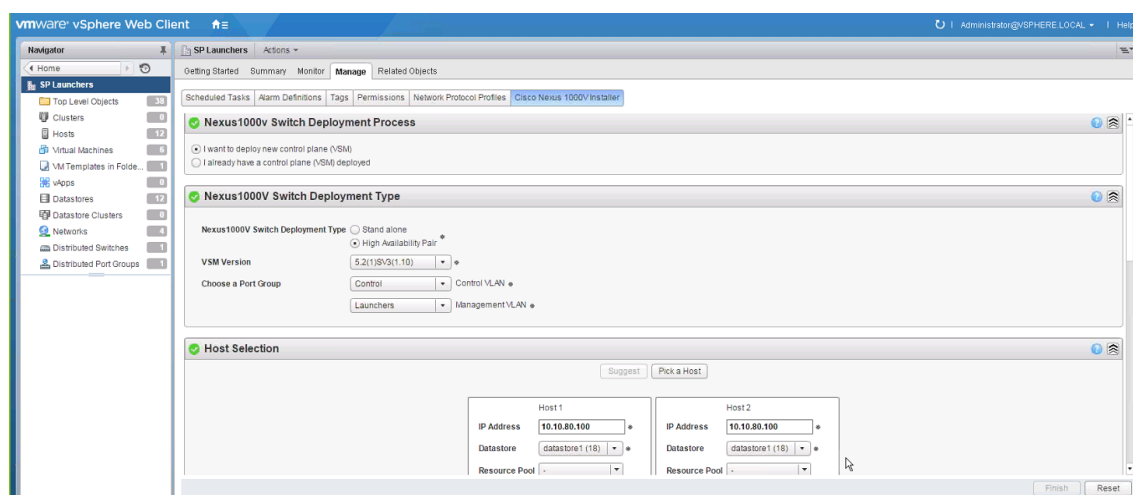
2. In the Nexus 1000v Switch Deployment area, choose I want to deploy new control plane (VSM).
3. In the Cisco Nexus 1000V Switch Deployment Type area, install the switches as an HA pair. By default, the High Availability Pair is selected.
4. Choose the control port group for the switch.

5. Choose the management port group for the switch.

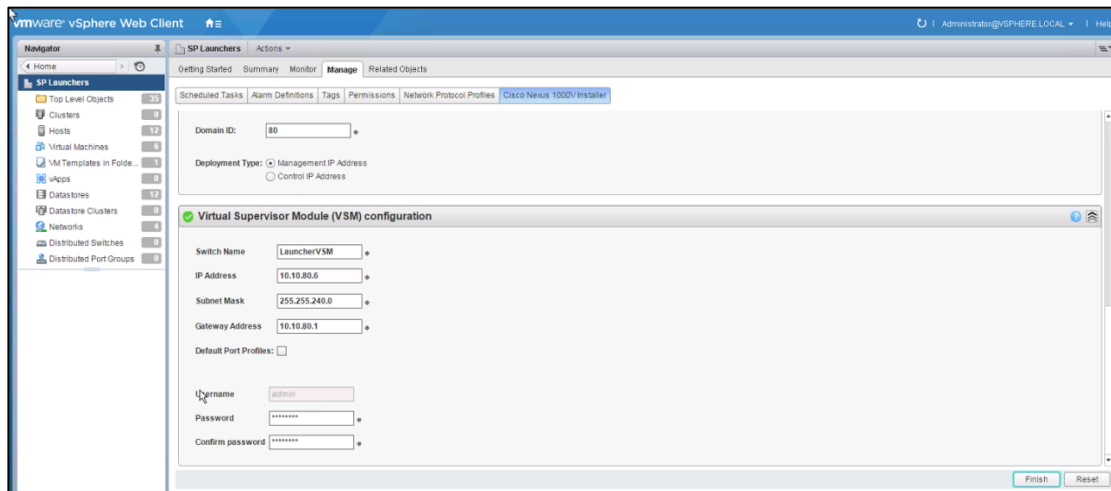


The Cisco Nexus 1000V VSM uses the management network to communicate with vCenter Server and ESXi. The management and control port group can use the same VLAN.

6. In the Host Selection area, click Suggest to choose two hosts based on the details provided in the Cisco Nexus 1000V Switch Deployment Type area. The IP address of the hosts on which the switch will be deployed.
7. The primary switch is deployed on Infrastructure Host 1 and the secondary switch is deployed on Infrastructure Host 2. Click Pick a Host to override the system choices.
8. Choose the system-selected datastore that you want to override. Choose Nimble_Infra_DS as the datastore for each host.



9. In the Switch Configuration area, enter 70 as the domain ID for the switch.
10. The domain ID is common for both the primary and secondary switches and it should be unique for every new switch. The range for the domain is from 1 to 1023.
11. In the Virtual Supervisor Module (VSM) configuration area, enter the Switch Name, IP Address, Subnet Mask, and Gateway Address.
12. Do not select Default Port Profiles.
13. Enter the Password and Confirm Password for Admin.



14. Click Finish to install the Cisco Nexus 1000V switch.



The Cisco Nexus 1000V installation is confirmed when the primary task Create Nexus 1000v Switch has the status Completed. A typical installation of the switch takes about 4 minutes.

Perform Base Configuration of the Primary VSM

SSH Connection to Primary VSM

To perform the base configuration of the primary VSM, complete the following steps:

1. Using an SSH client, log in to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands:



Any VLAN that has a VMkernel port should be assigned as a system VLAN on both the **uplink** and the **vEthernet** ports of the virtual switch.

```
config t
ntp server <<var_switch_a_ntp_ip>> use-vrf management
ntp server <<var_switch_b_ntp_ip>> use-vrf management
vlan <<var_ib-mgmt_vlan_id>> 70
name IB-MGMT-VLAN
vlan <<var_nfs_vlan_id>> 73
name NFS-VLAN
vlan <<var_vmotion_vlan_id>> 76
name vMotion-VLAN
```



The Nexus 1000V is currently limited to 1024 Max ports per profile. This solution is comprised of 2,512 virtual desktop machines for the user workload and requires three dedicated port-profiles.

```

vlan <<var_vdi_vlan_id>> 77
name DHCP
vlan <<var_vdi_vlan_id>> 77
name DHCP2
vlan <<var_vdi_vlan_id>> 77
name DHCP3
vlan <<var_vm-traffic_vlan_id>> 71
name Infrastructure
vlan <<var_native_vlan_id>> 1
name Native-VLAN
exit
port-profile type ethernet system-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>> 1
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-infra_vlan_id>> 70-79,164
channel-group auto mode on mac-pinning
no shutdown
system vlan <<var_ib-mgmt_vlan_id>>, <<var_nfs_vlan_id>>,
<<var_vmotion_vlan_id>>, <<var_vm-infra_vlan_id>> 70-79,164
system mtu 9000
state enabled
port-profile type vethernet IB-MGMT-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>> 70
no shutdown
system vlan <<var_ib-mgmt_vlan_id>> 70
state enabled
port-profile type vethernet NFS-VLAN
vmware port-group
switchport mode access

```



```
switchport access vlan <<var_nfs_vlan_id>> 73
no shutdown
system vlan <<var_nfs_vlan_id>> 73
state enabled
port-profile type vethernet vMotion-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vmotion_vlan_id>> 76
no shutdown
system vlan <<var_vmotion_vlan_id>> 76
state enabled

port-profile type vethernet VM-INFRA-VLAN
vmware port-group
switchport mode access
switchport access vlan <<var_vm-infra_vlan_id>> 71
no shutdown
system vlan <<var_vm-infra_vlan_id>> 71
state enabled
port-profile type vethernet n1kv-L3
capability l3control
vmware port-group
switchport mode access
switchport access vlan <<var_ib-mgmt_vlan_id>> 70
no shutdown
system vlan <<var_ib-mgmt_vlan_id>> 70
state enabled

port-profile type vethernet DHCP
vmware port-group
switchport mode access
switchport access vlan <<var_vdi_1_vlan_id>> 77
```

Install and Configure ESXi 6 U1a

```
no shutdown
max-ports 1024
system vlan <<var_vdi_1_vlan_id>> 77
state enabled

port-profile type vethernet DHCP2
vmware port-group
switchport mode access
switchport access vlan <<var_vdi_1_vlan_id>> 77
no shutdown
max-ports 1024
system vlan <<var_vdi_1_vlan_id>> 77
state enabled

port-profile type vethernet DHCP3
vmware port-group
switchport mode access
switchport access vlan <<var_vdi_1_vlan_id>> 77
no shutdown
max-ports 1024
system vlan <<var_vdi_1_vlan_id>> 77
state enabled

switchport access vlan <<var_ib-mgmt_vlan_id>> 70
no shutdown
system vlan <<var_ib-mgmt_vlan_id>> 70
state enabled
exit
copy run start
```

Add VMware ESXi Hosts to Cisco Nexus 1000V

VMware vSphere Web Client

To add VMware ESXi hosts, complete the following steps:

1. Back in the VMware vSphere Web Client, from the Home tab, select Cisco Virtual Switch Update Manager.
2. Under Basic Tasks, select Nexus 1000V.
3. Select Configure.
4. Select the SP-VDI datacenter on the right.
5. Select the VSM on the lower right.
6. Click Manage.
7. In the center pane, select the Add Host tab.
8. Expand the Infrastructure ESXi Cluster and select one of the Infrastructure Management Hosts.
9. Click Suggest.
10. Scroll down to Physical NIC Migration and expand each ESXi host.
11. On both hosts, unselect vmnic0, and select vmnic1. For vmnic1, select the system-uplink Profile.
12. Scroll down to VM Kernel NIC Setup and expand both ESXi hosts.
13. All VMkernel ports should already have the appropriate checkboxes selected.
14. Scroll down to VM Migration and expand both ESXi hosts.
15. Select the IB-MGMT-VLAN profile for the VSUM and vCenter Virtual Machines.
16. Click Finish.



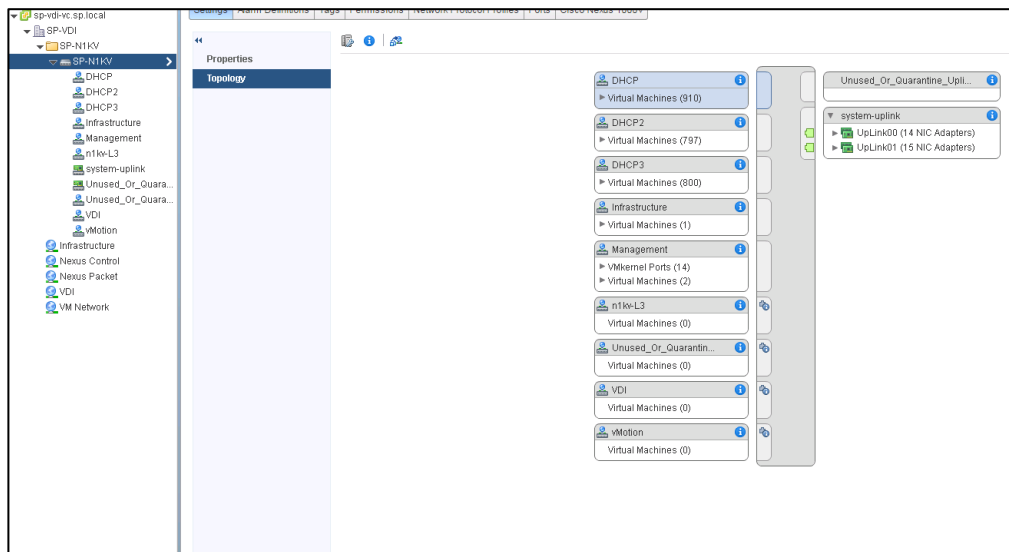
The progress of the virtual switch installation can be monitored from the c# interface.

Migrate ESXi Host Redundant Network Ports to Cisco Nexus 1000V

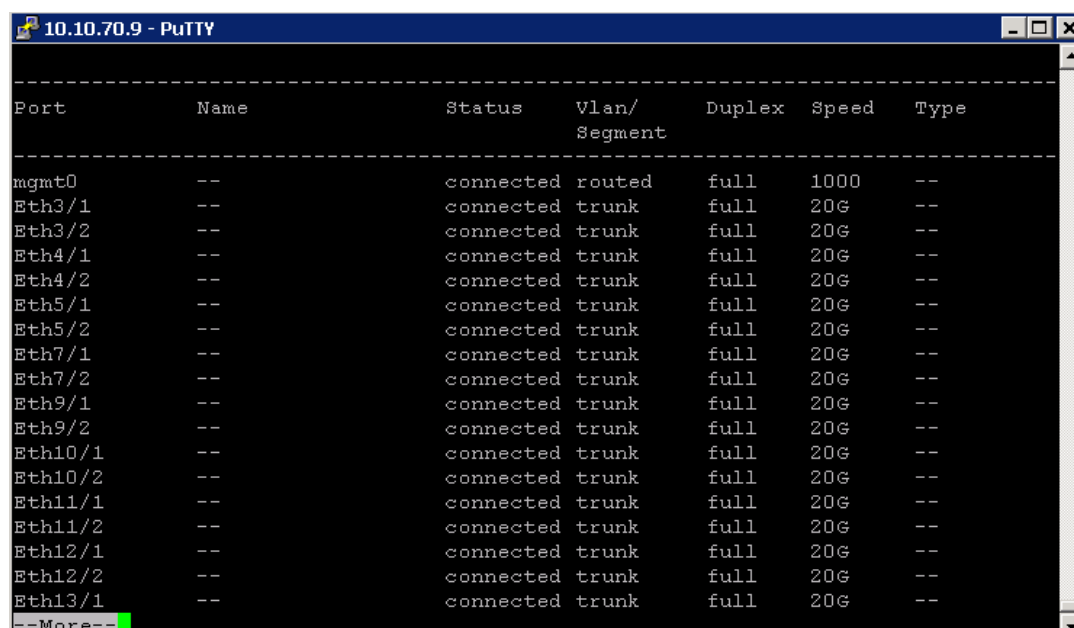
To migrate the ESXi host redundant network ports, complete the following steps:

1. In the VMware vSphere Web Client window, select Home > Hosts and Clusters.
2. On the left expand the Datacenter and cluster, and select the first VMware ESXi host.
3. In the center pane, select the Manage tab, then select Networking.
4. Select vSwitch0. All of the port groups on vSwitch0 should be empty. Click the red X under Virtual switches to delete vSwitch0.
5. Click Yes to remove vSwitch0. It may be necessary to refresh the Web Client to see the deletion.

6. The Nexus 1000V VSM should now be the only virtual switch. Select it and select the third icon above it under Virtual switches (Manage the physical network adapters connected to the selected switch).
7. Click the green plus sign to add an adapter.
8. For UpLink01, select the system-uplink port group and make sure vmnic0 is the Network adapter. Click OK.
9. Click OK to complete adding the Uplink. It may be necessary to refresh the Web Client to see the addition.

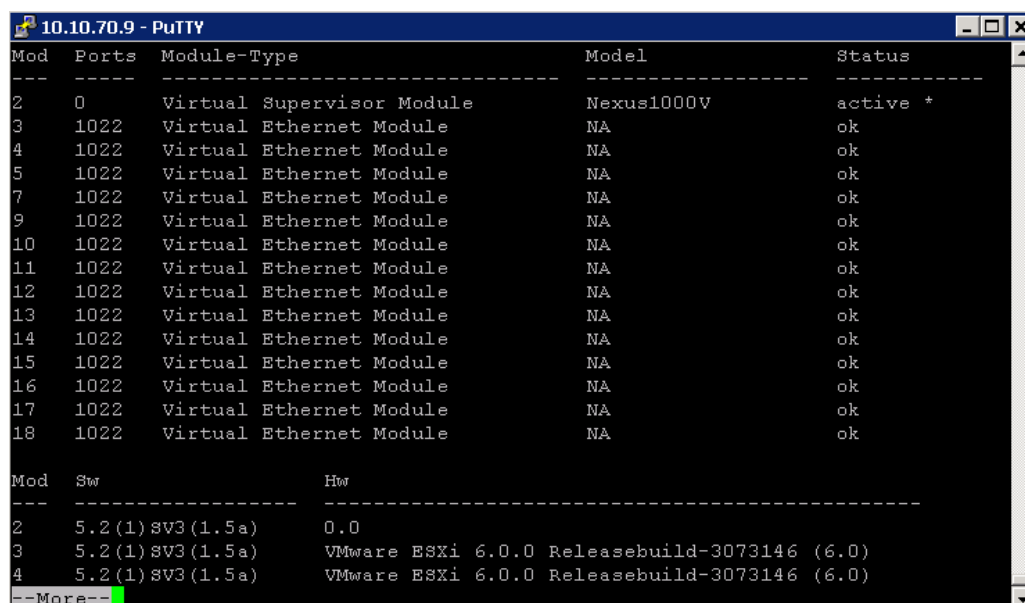


10. Repeat this procedure for the other ESXi host.
11. From the SSH client that is connected to the Cisco Nexus 1000V, run show interface status to verify that all interfaces and port channels have been correctly configured.



Port	Name	Status	Vlan/ Segment	Duplex	Speed	Type
mgmt0	--	connected	routed	full	1000	--
Eth3/1	--	connected	trunk	full	20G	--
Eth3/2	--	connected	trunk	full	20G	--
Eth4/1	--	connected	trunk	full	20G	--
Eth4/2	--	connected	trunk	full	20G	--
Eth5/1	--	connected	trunk	full	20G	--
Eth5/2	--	connected	trunk	full	20G	--
Eth7/1	--	connected	trunk	full	20G	--
Eth7/2	--	connected	trunk	full	20G	--
Eth9/1	--	connected	trunk	full	20G	--
Eth9/2	--	connected	trunk	full	20G	--
Eth10/1	--	connected	trunk	full	20G	--
Eth10/2	--	connected	trunk	full	20G	--
Eth11/1	--	connected	trunk	full	20G	--
Eth11/2	--	connected	trunk	full	20G	--
Eth12/1	--	connected	trunk	full	20G	--
Eth12/2	--	connected	trunk	full	20G	--
Eth13/1	--	connected	trunk	full	20G	--

12. Run show module and verify that the one ESXi host is present as a module.



Mod	Ports	Module-Type	Model	Status
2	0	Virtual Supervisor Module	Nexus1000V	active *
3	1022	Virtual Ethernet Module	NA	ok
4	1022	Virtual Ethernet Module	NA	ok
5	1022	Virtual Ethernet Module	NA	ok
7	1022	Virtual Ethernet Module	NA	ok
9	1022	Virtual Ethernet Module	NA	ok
10	1022	Virtual Ethernet Module	NA	ok
11	1022	Virtual Ethernet Module	NA	ok
12	1022	Virtual Ethernet Module	NA	ok
13	1022	Virtual Ethernet Module	NA	ok
14	1022	Virtual Ethernet Module	NA	ok
15	1022	Virtual Ethernet Module	NA	ok
16	1022	Virtual Ethernet Module	NA	ok
17	1022	Virtual Ethernet Module	NA	ok
18	1022	Virtual Ethernet Module	NA	ok

Mod	Sw	Hw
2	5.2 (1) SV3 (1.5a)	0.0
3	5.2 (1) SV3 (1.5a)	VMware ESXi 6.0.0 Releasebuild-3073146 (6.0)
4	5.2 (1) SV3 (1.5a)	VMware ESXi 6.0.0 Releasebuild-3073146 (6.0)

13. Repeat the above steps to migrate the remaining ESXi hosts to the Nexus 1000V.

14. Run: copy run start.

Cisco Nexus 1000V vTracker

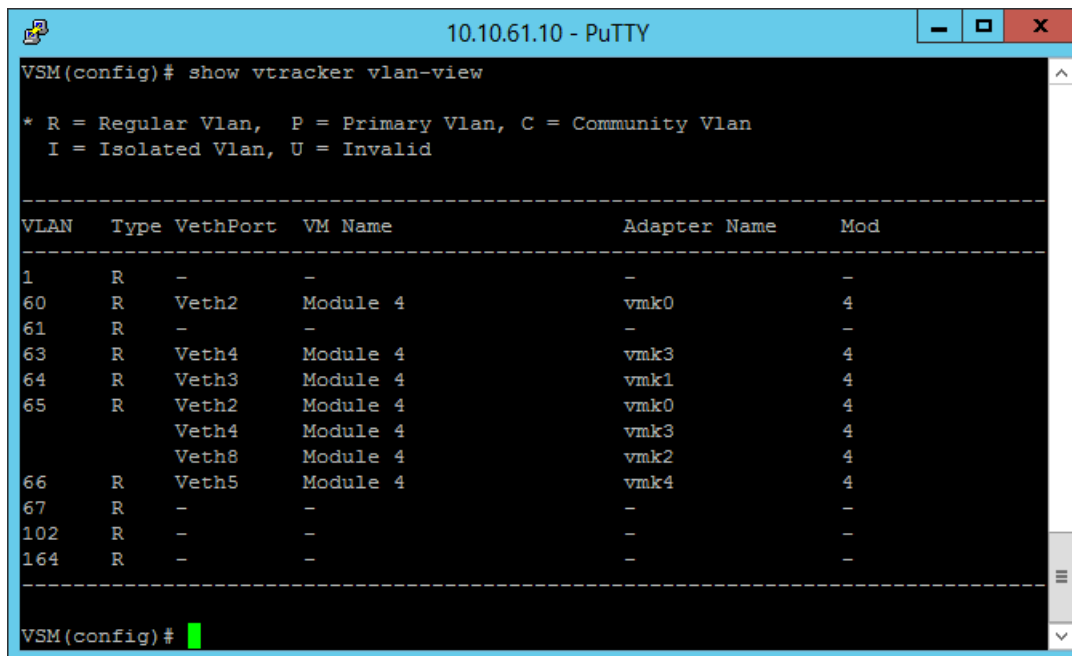
SSH Connection to Primary VSM

The vTracker feature on the Cisco Nexus 1000V switch provides information about the virtual network environment. To connect SSH to the primary VSM, complete the following steps:

1. From an SSH interface connected to the Cisco Nexus 1000V VSM, enter the following:

Install and Configure ESXi 6 U1a

```
config t
feature vtracker
copy run start
show vtracker upstream-view
show vtracker vm-view vnic
show vtracker vm-view info
show vtracker module-view pnuc
show vtracker vlan-view
copy run start
```



The screenshot shows a PuTTY terminal window titled "10.10.61.10 - PuTTY". The terminal displays the command "VSM(config)# show vtracker vlan-view" and its output. The output includes a legend for VLAN types (R = Regular, P = Primary, C = Community, I = Isolated, U = Invalid) and a table of VLAN configurations.

VLAN	Type	VethPort	VM Name	Adapter Name	Mod
1	R	-	-	-	-
60	R	Veth2	Module 4	vmk0	4
61	R	-	-	-	-
63	R	Veth4	Module 4	vmk3	4
64	R	Veth3	Module 4	vmk1	4
65	R	Veth2	Module 4	vmk0	4
		Veth4	Module 4	vmk3	4
		Veth8	Module 4	vmk2	4
66	R	Veth5	Module 4	vmk4	4
67	R	-	-	-	-
102	R	-	-	-	-
164	R	-	-	-	-

Nexus 1000V Configuration

```
SP-N1KV# sho ru
!Command: show running-config
!Time: Fri Mar 18 23:37:14 2016

version 5.2(1)SV3(1.5a)
hostname SP-N1KV

feature telnet

username admin password 5 $1$qxPYV8oS$OqsLtCVU/ZC8oatwK7fmU1 role network-admin
```

```
username admin keypair generate rsa

banner motd #Nexus 1000v Switch
#

ssh key rsa 2048
ip domain-lookup
ip host SP-N1KV 10.10.70.9
errdisable recovery cause failed-port-state
vem 3
    host id e811ca25-ef8f-e511-0000-00000000000e
vem 4
    host id e811ca25-ef8f-e511-0000-00000000003d

vem 5_[K
    host id e811ca25-ef8f-e511-0000-00000000000c
vem 6
    host id e811ca25-ef8f-e511-0000-00000000003f
vem 7
    host id e811ca25-ef8f-e511-0000-00000000003e
vem 8
    host id e811ca25-ef8f-e511-0000-00000000002f
vem 9
    host id e811ca25-ef8f-e511-0000-00000000001c
vem 10
    host id e811ca25-ef8f-e511-0000-00000000002c
vem 11
    host id e811ca25-ef8f-e511-0000-00000000001e
vem 12
    host id e811ca25-ef8f-e511-0000-00000000003c
vem 13
    host id e811ca25-ef8f-e511-0000-00000000001f
```

```
vem 14
    host id e811ca25-ef8f-e511-0000-00000000002e
vem 15
    host id e811ca25-ef8f-e511-0000-00000000002d
vem 16
    host id e811ca25-ef8f-e511-0000-00000000001d
vem 17
    host id e811ca25-ef8f-e511-0000-00000000000d
vem 18
    host id e811ca25-ef8f-e511-0000-00000000000f
snmp-server user admin network-admin auth md5 0x1abb14a596b559408ae5ed99ed9e84c0
priv 0x1abb14a596b559408ae5ed99ed9e84c0 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vrf context management
    ip route 0.0.0.0/0 10.10.70.1
vlan configuration 701-702
vlan 1,70-72,76-77,701-702
vlan 70
    name Management
vlan 71
    name Infrastructure
vlan 72
    name VDI

vlan 76_[K
    name vMotion
```



```
vlan 77
    name DHCP
vlan 701
    name Control
vlan 702
    name Packet

port-channel load-balance ethernet source-mac
port-profile default max-ports 32
port-profile type ethernet Unused_Or_Quarantine_Uplink
    shutdown
    description Port-group created for Nexus 1000V internal usage. Do not use.
    state enabled
    vmware port-group
port-profile type vethernet Unused_Or_Quarantine_Veth
    shutdown
    description Port-group created for Nexus 1000V internal usage. Do not use.
    state enabled
    vmware port-group
port-profile type ethernet system-uplink
    switchport mode trunk
    switchport trunk native vlan 1
    switchport trunk allowed vlan 70-79
    system mtu 9000
    channel-group auto mode on mac-pinning
    no shutdown
    system vlan 70-72,76-77
    state enabled
    vmware port-group
port-profile type vethernet Management
    switchport mode access
    switchport access vlan 70
```

```
no shutdown
capability l3control
system vlan 70
state enabled
vmware port-group
port-profile type vethernet Infrastructure
switchport mode access
switchport access vlan 71
no shutdown
system vlan 71
state enabled
vmware port-group
port-profile type vethernet VDI
switchport mode access
switchport access vlan 72
no shutdown
system vlan 72
state enabled
vmware port-group
port-profile type vethernet vMotion
switchport mode access
switchport access vlan 76
no shutdown
system vlan 76
state enabled
vmware port-group
port-profile type vethernet DHCP
switchport mode access
switchport access vlan 77
no shutdown
max-ports 1024
system vlan 77
```

```
state enabled
vmware port-group
port-profile type vethernet DHCP2
switchport mode access
switchport access vlan 77
no shutdown
max-ports 1024
system vlan 77
state enabled
vmware port-group
port-profile type vethernet DHCP3
switchport mode access
switchport access vlan 77
no shutdown
max-ports 1024
system vlan 77
state enabled
vmware port-group
port-profile type vethernet nlkv-L3
switchport mode access
switchport access vlan 70
state enabled
vmware port-group

interface port-channel1

inherit port-profile system-uplink
vem 4
mtu 9000

interface port-channel2
```

```
inherit port-profile system-uplink  
vem 5  
mtu 9000
```

```
interface port-channel3  
inherit port-profile system-uplink  
vem 7  
mtu 9000
```

```
interface port-channel4  
inherit port-profile system-uplink  
vem 9  
mtu 9000
```

```
interface port-channel5  
inherit port-profile system-uplink  
vem 10  
mtu 9000
```

```
interface port-channel6  
inherit port-profile system-uplink  
vem 11  
mtu 9000
```

```
interface port-channel7  
inherit port-profile system-uplink  
vem 12  
mtu 9000
```

```
interface port-channel8  
inherit port-profile system-uplink
```

```
vem 13  
mtu 9000
```

```
interface port-channel9  
  inherit port-profile system-uplink  
  vem 15  
  mtu 9000
```

```
interface port-channel10  
  inherit port-profile system-uplink  
  
  vem 16  
  mtu 9000
```

```
interface port-channel11  
  inherit port-profile system-uplink  
  vem 18  
  mtu 9000
```

```
interface port-channel12  
  inherit port-profile system-uplink  
  vem 3  
  mtu 9000
```

```
interface port-channel13  
  inherit port-profile system-uplink  
  vem 14  
  mtu 9000
```

```
interface port-channel14  
  inherit port-profile system-uplink  
  vem 17
```

```
mtu 9000
```

```
interface mgmt0  
  ip address 10.10.70.9/24
```

```
interface Ethernet3/1  
  inherit port-profile system-uplink
```

```
interface Ethernet3/2  
  inherit port-profile system-uplink
```

```
interface Ethernet4/1  
  inherit port-profile system-uplink
```

```
interface Ethernet4/2  
  
  inherit port-profile system-uplink
```

```
interface Ethernet5/1  
  inherit port-profile system-uplink
```

```
interface Ethernet5/2  
  inherit port-profile system-uplink
```

```
interface Ethernet7/1  
  inherit port-profile system-uplink
```

```
interface Ethernet7/2  
  inherit port-profile system-uplink
```

```
interface Ethernet9/1
    inherit port-profile system-uplink
```

```
interface Ethernet9/2
    inherit port-profile system-uplink
```

```
interface Ethernet10/1
    inherit port-profile system-uplink
```

```
interface Ethernet10/2
    inherit port-profile system-uplink
```

```
interface Ethernet11/1
    inherit port-profile system-uplink
```

```
interface Ethernet11/2
    inherit port-profile system-uplink
```

```
interface Ethernet12/1
    inherit port-profile system-uplink
```

```
interface Ethernet12/2
    inherit port-profile system-uplink
```

```
interface Ethernet13/1
    inherit port-profile system-uplink
```

```
interface Ethernet13/2
    inherit port-profile system-uplink
```

```
interface Ethernet14/1
```

```
inherit port-profile system-uplink

_K
interface Ethernet14/2
inherit port-profile system-uplink

interface Ethernet15/1
inherit port-profile system-uplink

interface Ethernet15/2
inherit port-profile system-uplink

interface Ethernet16/1
inherit port-profile system-uplink

interface Ethernet16/2
inherit port-profile system-uplink

interface Ethernet17/1
inherit port-profile system-uplink

interface Ethernet17/2
inherit port-profile system-uplink

interface Ethernet18/1

inherit port-profile system-uplink

interface Ethernet18/2
inherit port-profile system-uplink
```



```
interface control0
line console
line vty
boot kickstart bootflash:/n1000v-dk9-kickstart.5.2.1.SV3.1.5a.bin sup-1
boot system bootflash:/n1000v-dk9.5.2.1.SV3.1.5a.bin sup-1
boot kickstart bootflash:/n1000v-dk9-kickstart.5.2.1.SV3.1.5a.bin sup-2
boot system bootflash:/n1000v-dk9.5.2.1.SV3.1.5a.bin sup-2
svs-domain
    domain id 70
    control vlan 1
    packet vlan 1
    svcs mode L3 interface mgmt0
    switch-guid 3f4da4f4-cf89-4968-96ab-bdd61db67f1c
    enable l3sec
svs connection SP-VDI
    protocol vmware-vim
    remote ip address 10.10.71.26 port 80
    vmware dvs uuid "c6 35 09 50 15 cb 1e 2f-c8 42 86 a3 19 28 98 d0" datacenter-n
ame SP-VDI
    max-ports 9000
    connect
vservice global type vsg
    no tcp state-checks invalid-ack
    no tcp state-checks seq-past-window
    no tcp state-checks window-variation
    no bypass asa-traffic
    no l3-frag
vservice global
    idle-timeout
        tcp 30
        udp 4
```

Install and Configure ESXi 6 U1a

```
icmp 4
layer-3 4
layer-2 2
nsc-policy-agent
registration-ip 0.0.0.0
shared-secret *****
log-level
```

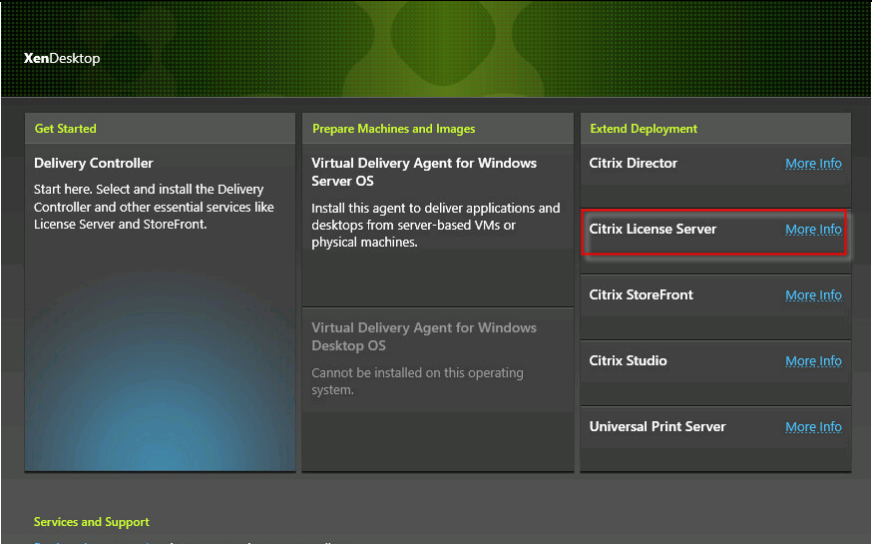
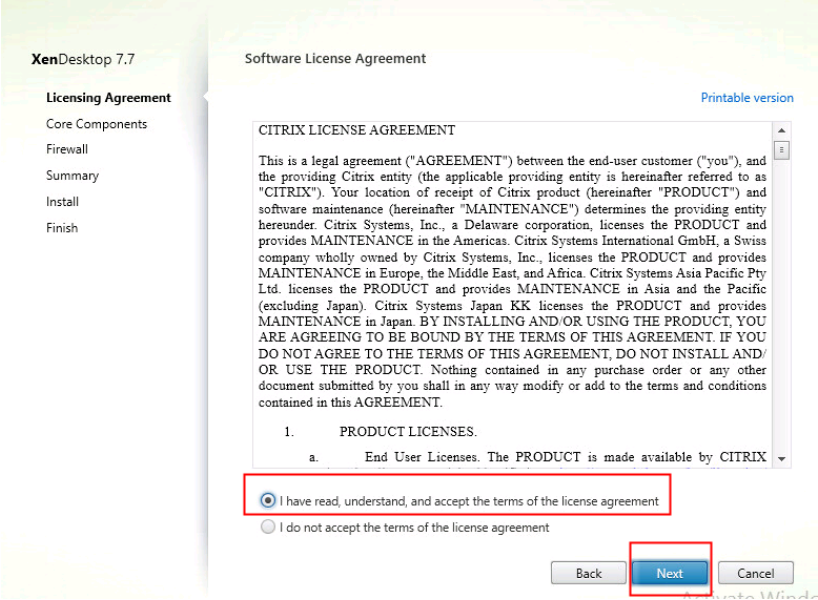
Installing and Configuring Citrix License Server

To install and configure the Citrix Licexe Server, complete the following steps:

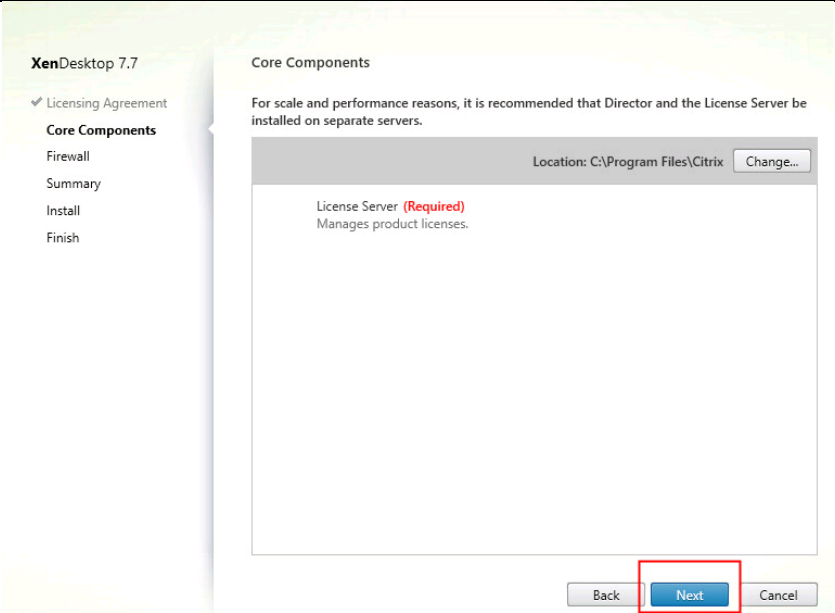
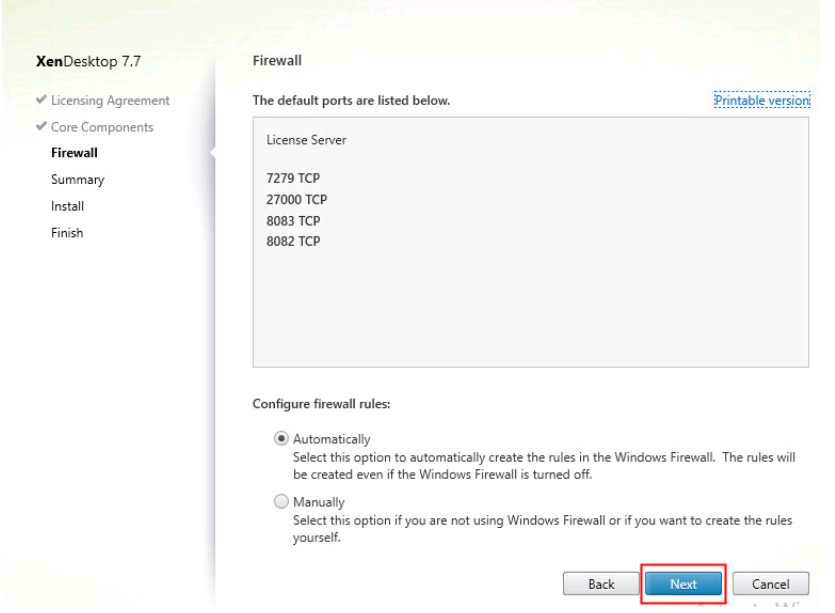


In this study we used a dedicated Citrix License Server (SP-LIC).

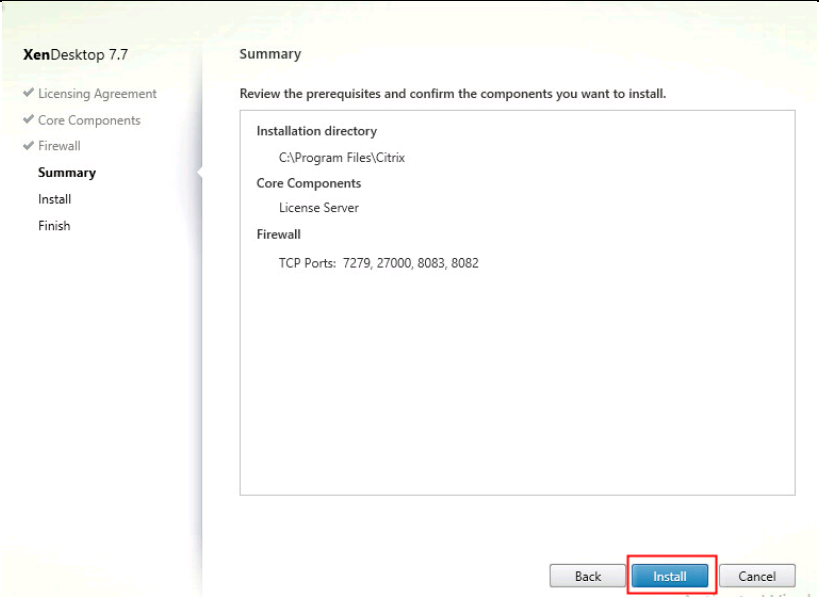
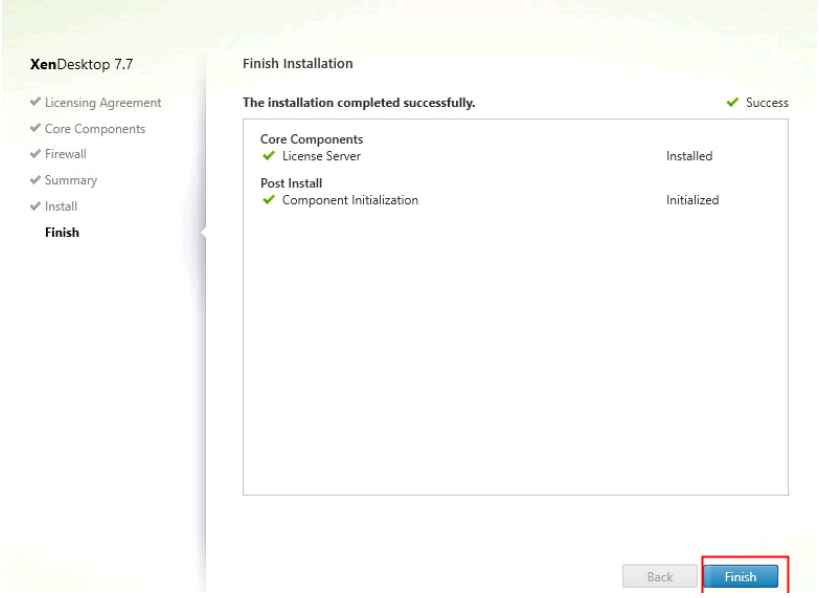
Instructions	Visual
<div>1. Insert the Citrix XenDesktop 7.7 ISO and let AutoRun launch the installer.</div> <div>2. Click the XD Start button.</div>	

Instructions	Visual
<p>3. Click Citrix License Server</p>	
<p>4. Read the Citrix License Agreement.</p> <p>5. If acceptable, select the radio button labeled "I accept the terms in the license agreement."</p> <p>6. Click Next.</p>	

Instructions	Visual
--------------	--------

Instructions	Visual
1. Click Next	 The screenshot shows the 'Core Components' step of the XenDesktop 7.7 installation. On the left, a navigation pane lists 'Licensing Agreement', 'Core Components', 'Firewall', 'Summary', 'Install', and 'Finish'. The 'Core Components' section is active. The main area contains a message: 'For scale and performance reasons, it is recommended that Director and the License Server be installed on separate servers.' Below this is a 'Location' field set to 'C:\Program Files\Citrix' with a 'Change...' button. A 'License Server (Required)' section states 'Manages product licenses.' At the bottom, there are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a red rectangular box.
2. Click Next	 The screenshot shows the 'Firewall' configuration step of the XenDesktop 7.7 installation. The navigation pane on the left now highlights 'Firewall'. The main area is titled 'Firewall' and includes a link for 'Printable version'. It lists default ports for the 'License Server': 7279 TCP, 27000 TCP, 8083 TCP, and 8082 TCP. Below this, under 'Configure firewall rules:', there are two radio button options: 'Automatically' (selected) and 'Manually'. The 'Automatically' option has a description: 'Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.' The 'Manually' option has a description: 'Select this option if you are not using Windows Firewall or if you want to create the rules yourself.' At the bottom, there are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a red rectangular box.

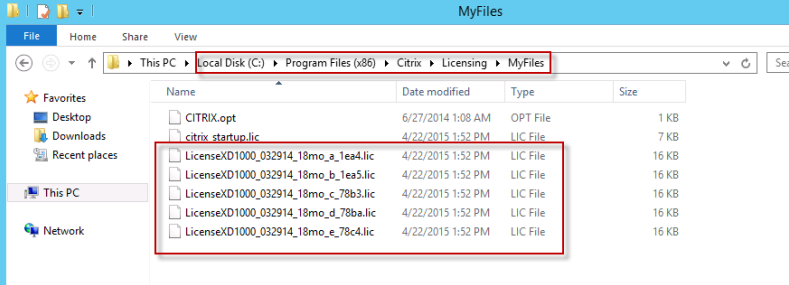

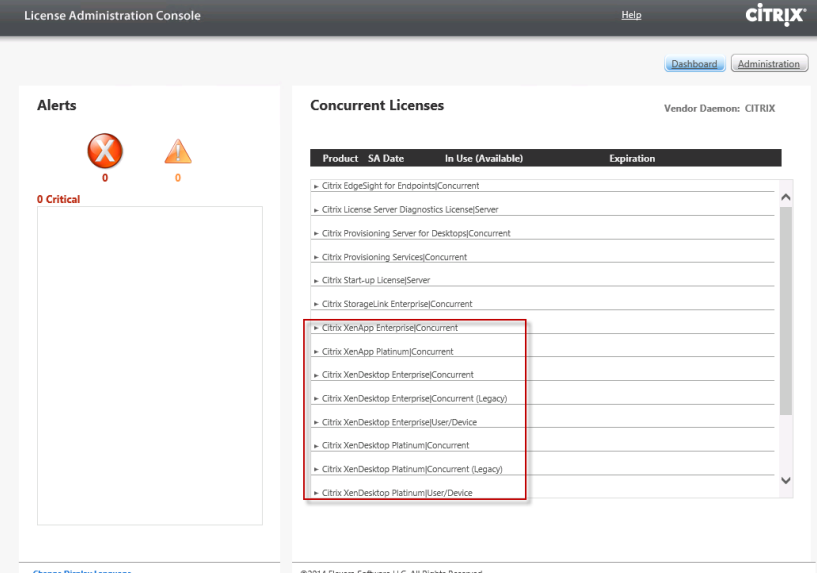
Instructions	Visual
--------------	--------

Instructions	Visual
3. Click Next	
4. Click Finish	



Before configuring XenDesktop or PVS, install the Citrix Licenses.

Instructions	Visual
--------------	--------

Instructions	Visual
<p>1. Copy the license files to the default location (C:\Program Files (x86)\Citrix\Licensing\MyFiles) on the license server</p> <p>2. Restart the server or licensing services so that the licenses are activated.</p>	
<p>3. Run the application Citrix License Administration Console</p>	
<p>4. Confirm that the license files have been read and enabled correctly.</p>	

Installing and Configuring Citrix Provisioning Server 7.7

In most implementations, there is a single vDisk providing the standard image for multiple target devices. Thousands of target devices can use a single vDisk shared across multiple Provisioning Services (PVS) servers in the same farm, simplifying virtual desktop management. This section describes the installation and configuration tasks required to create a PVS implementation.

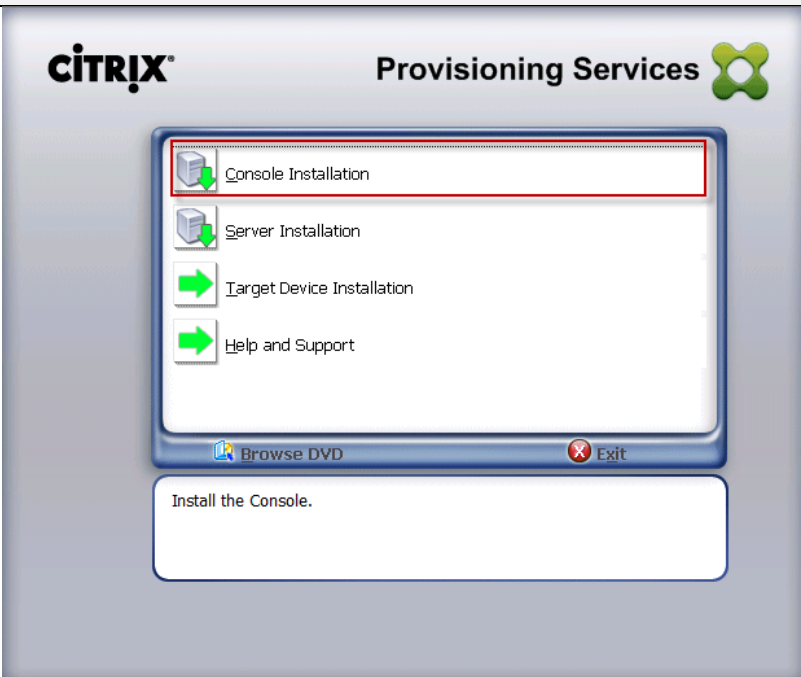
The PVS server can have many stored vDisks, and each vDisk can be several gigabytes in size. Your streaming performance and manageability can be improved using a RAID array, SAN, or NAS. PVS software and hardware requirements are available [at http://support.citrix.com/proddocs/topic/provisioning-7/pvs-install-task1-plan-6-0.html](http://support.citrix.com/proddocs/topic/provisioning-7/pvs-install-task1-plan-6-0.html).


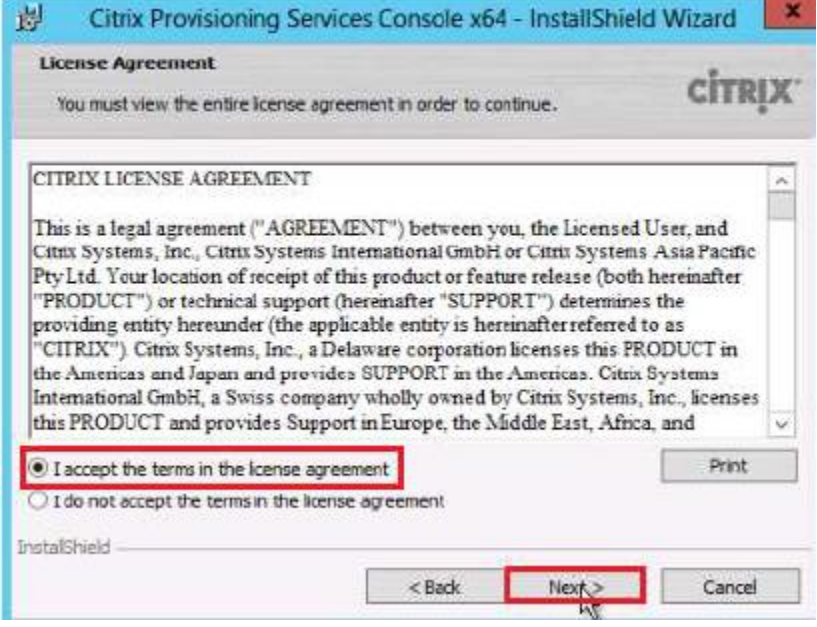
Prerequisites

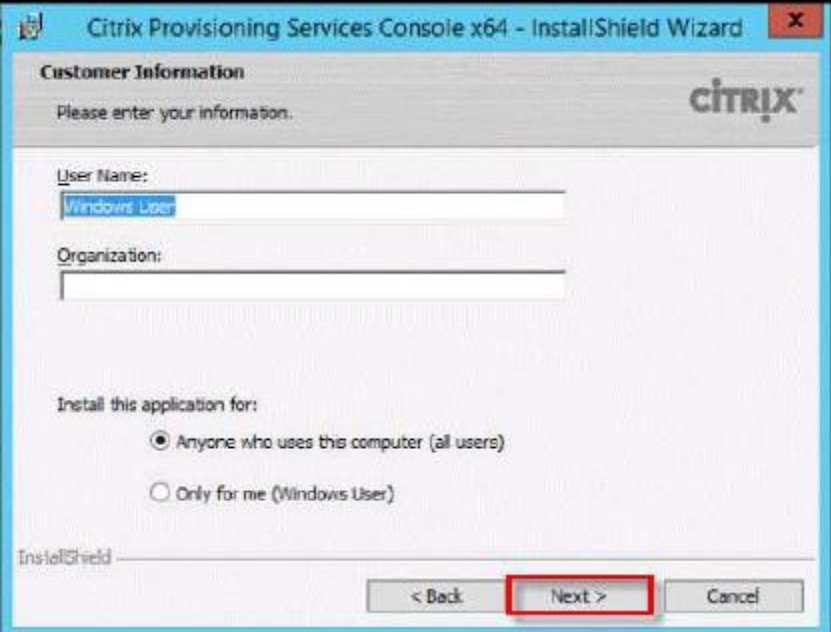
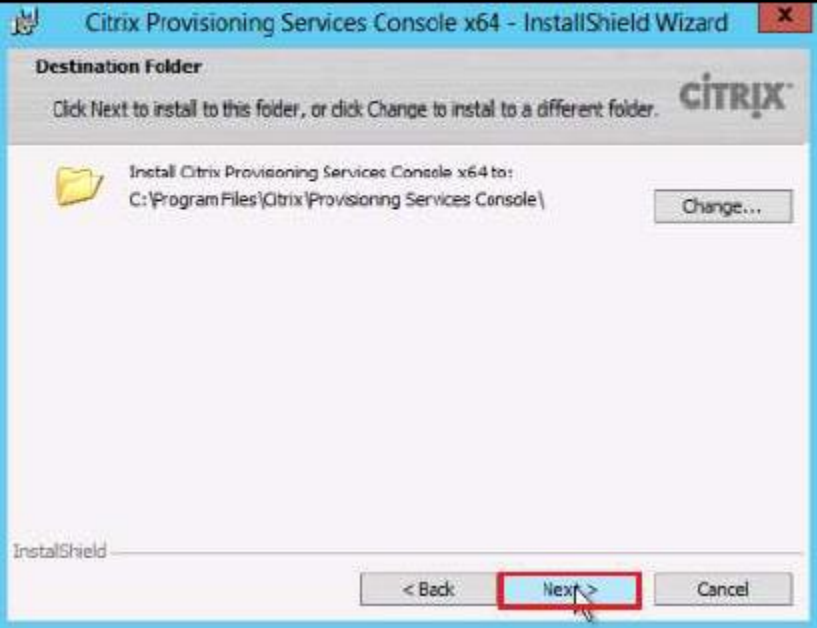
In this study, we used an existing SQL Always-On pair. For more information about configuring a SQL HA pair, please reference [https://msdn.microsoft.com/en-us/library/ff878265\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/ff878265(v=sql.110).aspx).

Only one MS SQL database is associated with a farm. You can choose to install the Provisioning Services database software on an existing SQL database, if that machine can communicate with all Provisioning Servers within the farm, or with a new SQL Express database machine, created using the SQL Express software that is free from Microsoft.

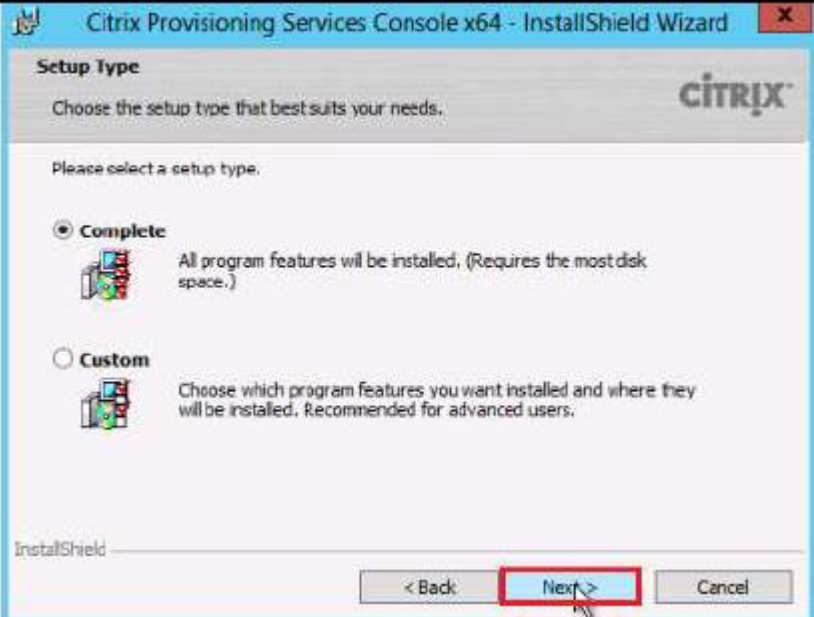
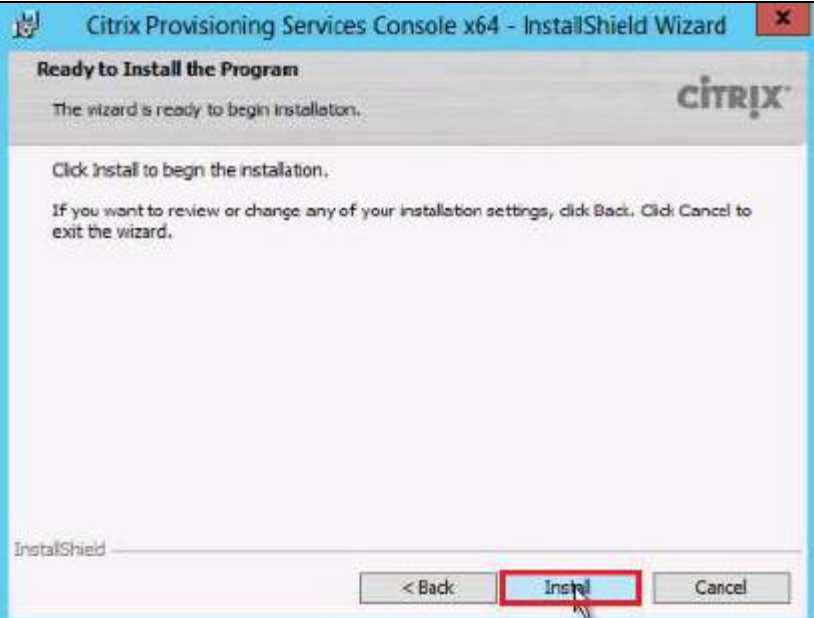
The following MS SQL 2008, MS SQL 2008 R2, MS SQL 2012, MS SQL 2012 R2 and MS SQL 2014 Server (32 or 64-bit editions) databases can be used for the Provisioning Services database: SQL Server Express Edition, SQL Server Workgroup Edition, SQL Server Standard Edition, SQL Server Enterprise Edition. Microsoft SQL were installed separately for this CVD.

Instructions	Visual
<ol style="list-style-type: none"> 1. Insert the Citrix Provisioning Services 7.7 ISO and let AutoRun launch the installer. 2. Click the Console Installation button. 	

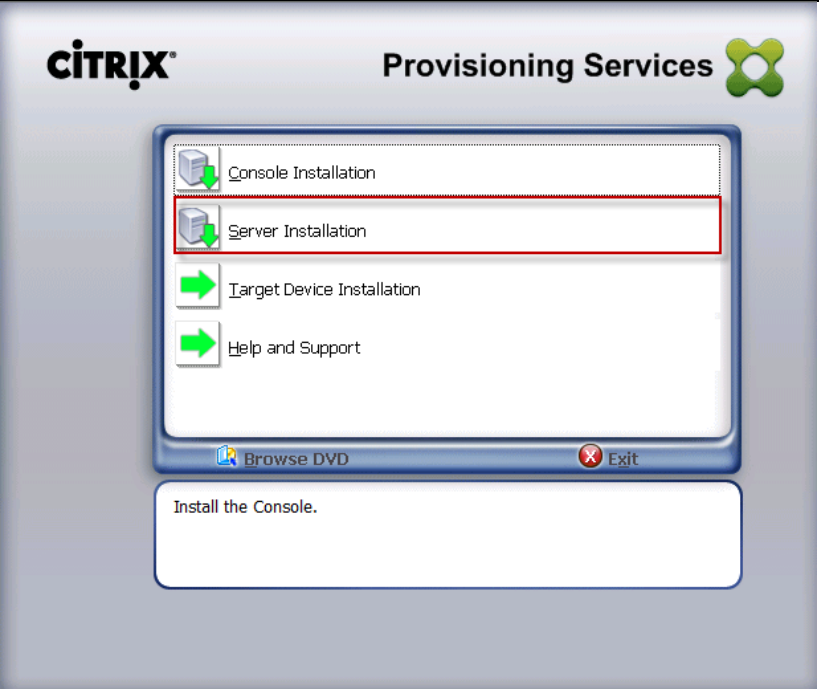
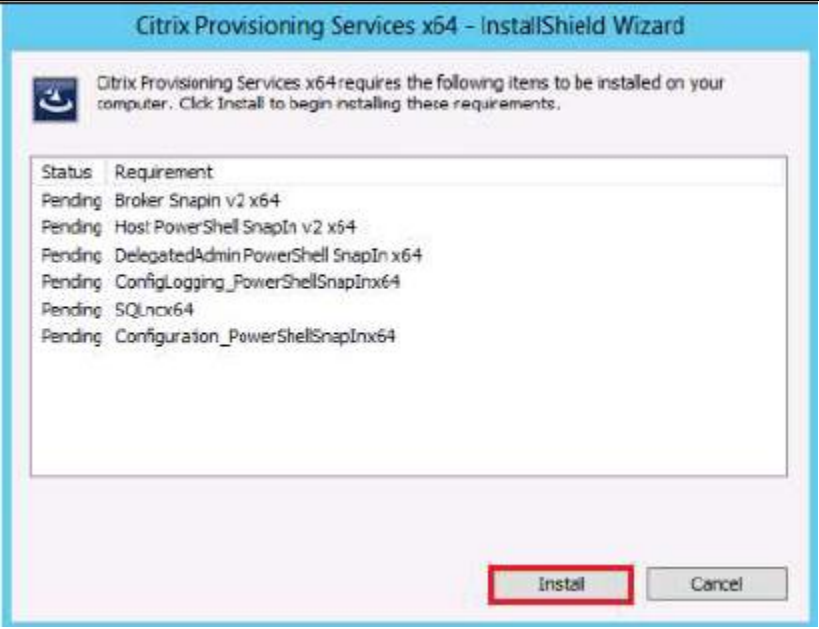
Instructions	Visual
3. Click Next	
<p>4. Read the Citrix License Agreement.</p> <p>5. If acceptable, select the radio button labeled “I accept the terms in the license agreement.”</p> <p>6. Click Next</p>	
Instructions	Visual

Instructions	Visual
<div>1. Optionally provide User Name and Organization.</div> <div>2. Click Next</div>	
<div>3. Accept the default path.</div> <div>4. Click Next</div>	

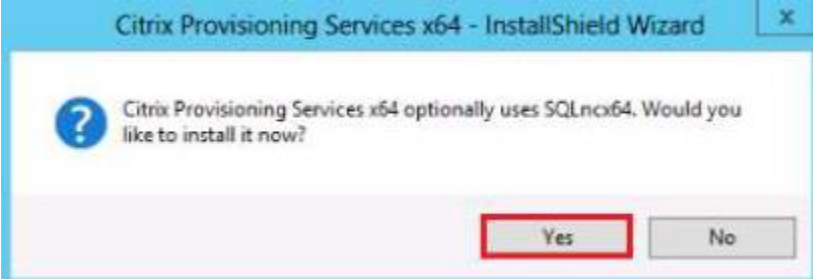

Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Leave the Complete radio button selected.</div> <div>2. Click Next</div>	 <p>The screenshot shows the 'Setup Type' window of the Citrix Provisioning Services Console x64 - InstallShield Wizard. It prompts the user to 'Choose the setup type that best suits your needs.' There are two options: 'Complete' (selected) and 'Custom'. The 'Complete' option states 'All program features will be installed, (Requires the most disk space.)'. The 'Custom' option states 'Choose which program features you want installed and where they will be installed. Recommended for advanced users.' At the bottom, there are three buttons: '< Back', 'Next >' (highlighted with a red rectangle), and 'Cancel'.</p>
<div>3. Click the Install button to start the console installation.</div>	 <p>The screenshot shows the 'Ready to Install the Program' window of the Citrix Provisioning Services Console x64 - InstallShield Wizard. It states 'The wizard is ready to begin installation.' and 'Click Install to begin the installation.' It also provides instructions: 'If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.' At the bottom, there are three buttons: '< Back', 'Install' (highlighted with a red rectangle), and 'Cancel'.</p>

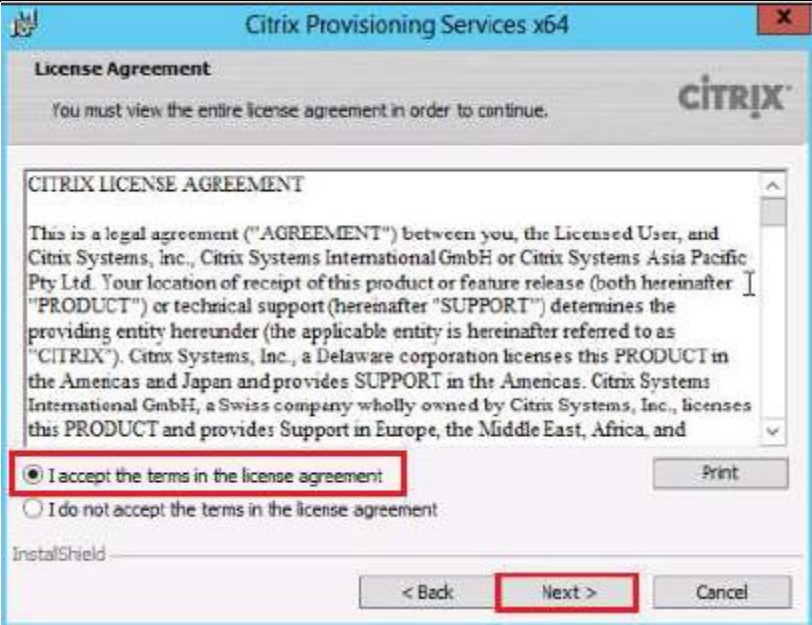

Instructions	Visual
--------------	--------

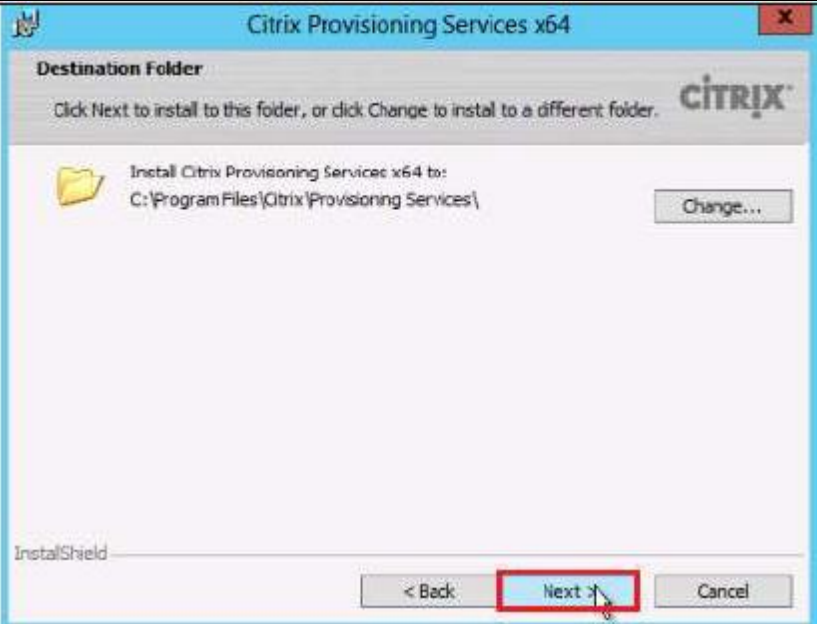
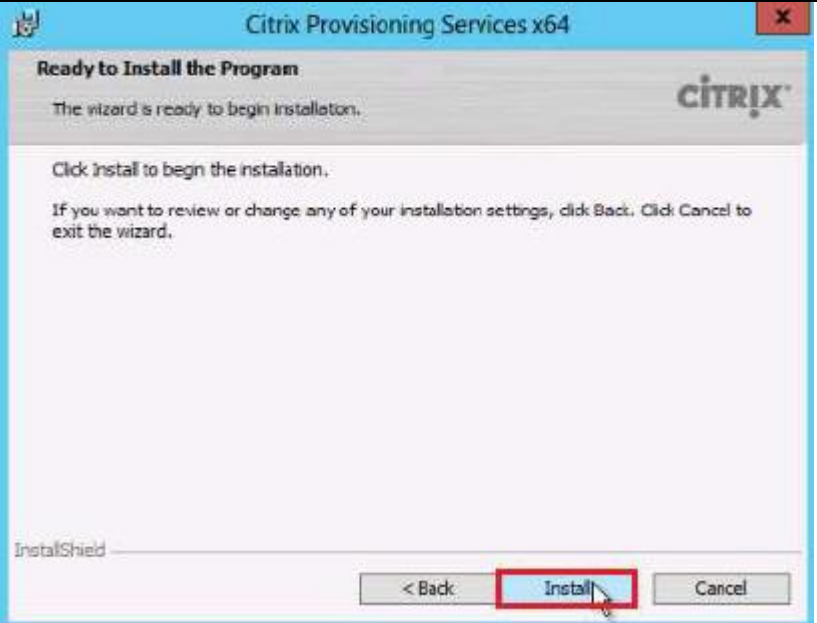
Instructions	Visual														
<div>1. From the main installation screen, select Server Installation.</div> <div>2. The installation wizard will check to resolve dependencies and then begin the PVS server installation process.</div>	 <p>The screenshot shows the Citrix Provisioning Services main installation window. It features the Citrix logo and the title 'Provisioning Services'. Inside a blue-bordered frame, there are four options: 'Console Installation', 'Server Installation' (highlighted with a red rectangle), 'Target Device Installation', and 'Help and Support'. At the bottom of the frame are 'Browse DVD' and 'Exit' buttons. Below the frame, a text box says 'Install the Console.'</p>														
<div>3. Click Install on the prerequisites dialog.</div>	 <p>The screenshot shows the 'Citrix Provisioning Services x64 - InstallShield Wizard' prerequisites dialog. It states that Citrix Provisioning Services x64 requires several items to be installed. A table lists these requirements, all with a 'Pending' status. The 'Install' button at the bottom right is highlighted with a red rectangle.</p> <table border="1"><thead><tr><th>Status</th><th>Requirement</th></tr></thead><tbody><tr><td>Pending</td><td>Broker Snapin v2 x64</td></tr><tr><td>Pending</td><td>Host PowerShell SnapIn v2 x64</td></tr><tr><td>Pending</td><td>DelegatedAdmin PowerShell SnapIn x64</td></tr><tr><td>Pending</td><td>ConfigLogging_PowerShellSnapInx64</td></tr><tr><td>Pending</td><td>SQLncx64</td></tr><tr><td>Pending</td><td>Configuration_PowerShellSnapInx64</td></tr></tbody></table>	Status	Requirement	Pending	Broker Snapin v2 x64	Pending	Host PowerShell SnapIn v2 x64	Pending	DelegatedAdmin PowerShell SnapIn x64	Pending	ConfigLogging_PowerShellSnapInx64	Pending	SQLncx64	Pending	Configuration_PowerShellSnapInx64
Status	Requirement														
Pending	Broker Snapin v2 x64														
Pending	Host PowerShell SnapIn v2 x64														
Pending	DelegatedAdmin PowerShell SnapIn x64														
Pending	ConfigLogging_PowerShellSnapInx64														
Pending	SQLncx64														
Pending	Configuration_PowerShellSnapInx64														

Instructions	Visual
--------------	--------

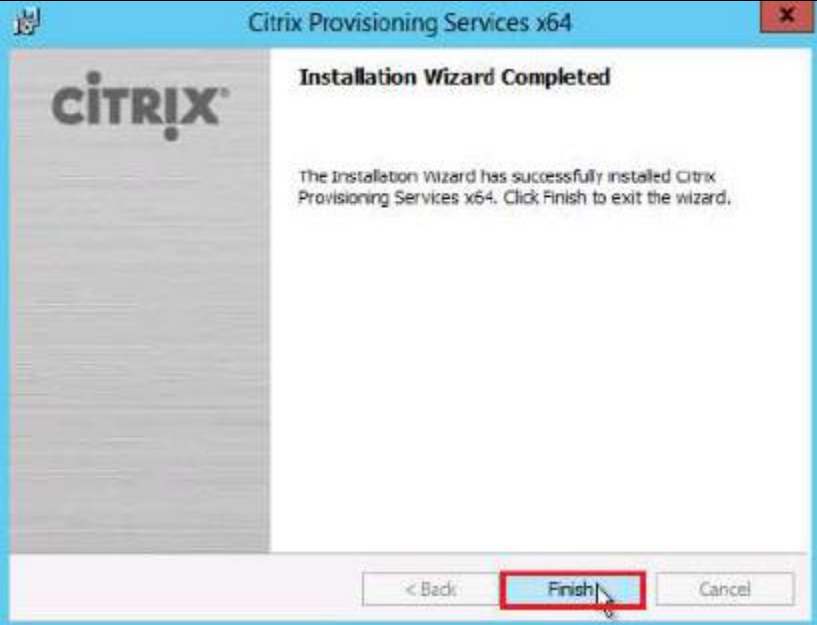

Instructions	Visual
<p>1. Click Yes when prompted to install the SQL Native Client.</p>	
<p>2. Click Next when the Installation wizard starts.</p>	

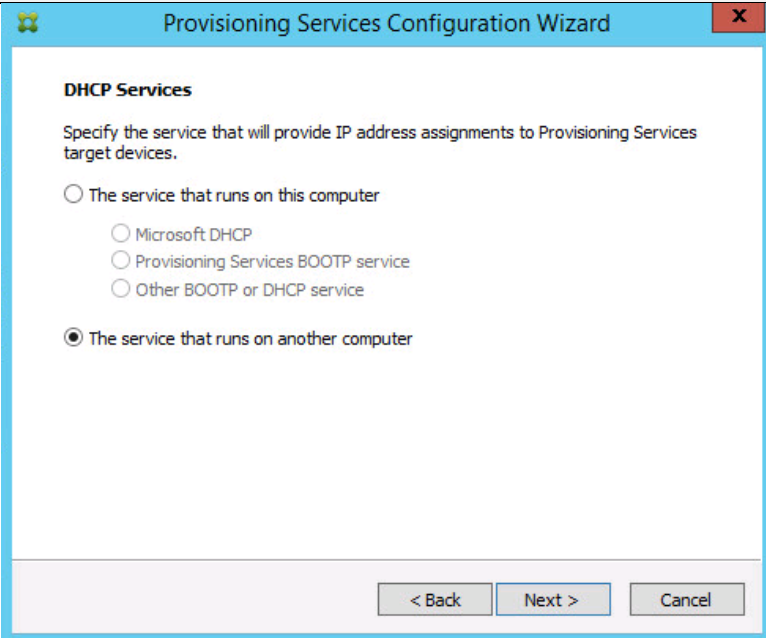
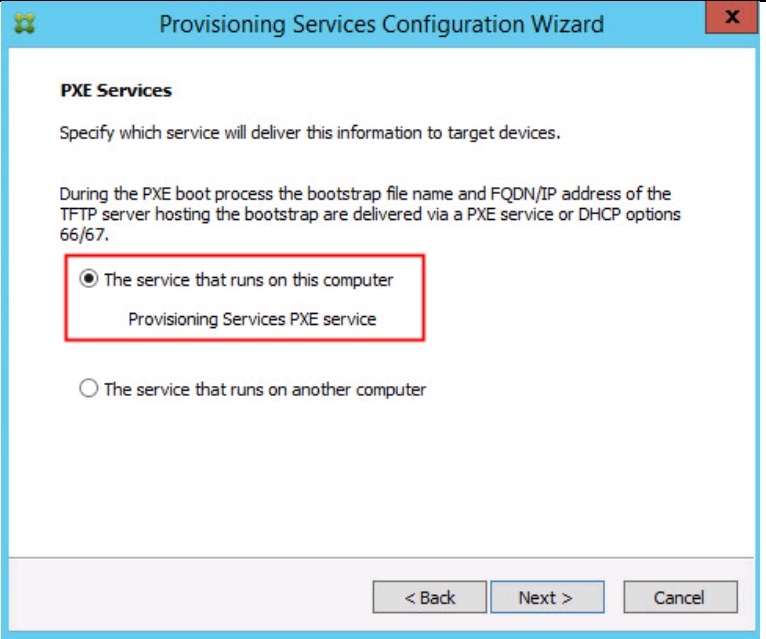
Instructions	Visual
--------------	--------

Instructions	Visual
<p>1. Review the license agreement terms.</p> <p>2. If acceptable, select the radio button labeled “I accept the terms in the license agreement.”</p> <p>3. Click Next</p>	 <p>The screenshot shows the 'Citrix Provisioning Services x64' window with the 'License Agreement' tab selected. The text of the license agreement is displayed in a scrollable area. Below the text, there are two radio buttons: 'I accept the terms in the license agreement' (which is selected and highlighted with a red box) and 'I do not accept the terms in the license agreement'. At the bottom right, there is a 'Next >' button highlighted with a red box, along with '< Back' and 'Cancel' buttons. A 'Print' button is also visible on the right side of the agreement text area.</p>
<p>4. Provide User Name, and Organization information.</p> <p>Select who will see the application.</p> <p>5. Click Next</p>	 <p>The screenshot shows the 'Citrix Provisioning Services x64' window with the 'Customer Information' tab selected. It prompts the user to 'Please enter your information.' There are two text input fields: 'User Name:' (containing 'Windows User') and 'Organization:'. Below these, there is a section 'Install this application for:' with two radio buttons: 'Anyone who uses this computer (all users)' (selected) and 'Only for me (Windows User)'. At the bottom right, the 'Next >' button is highlighted with a red box, along with '< Back' and 'Cancel' buttons. The 'InstallShield' logo is visible in the bottom left corner.</p>
Instructions	Visual

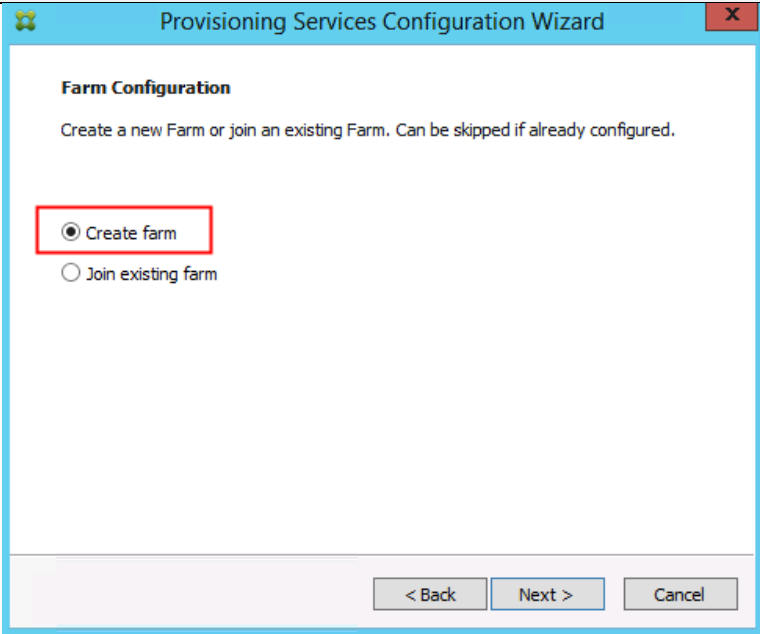
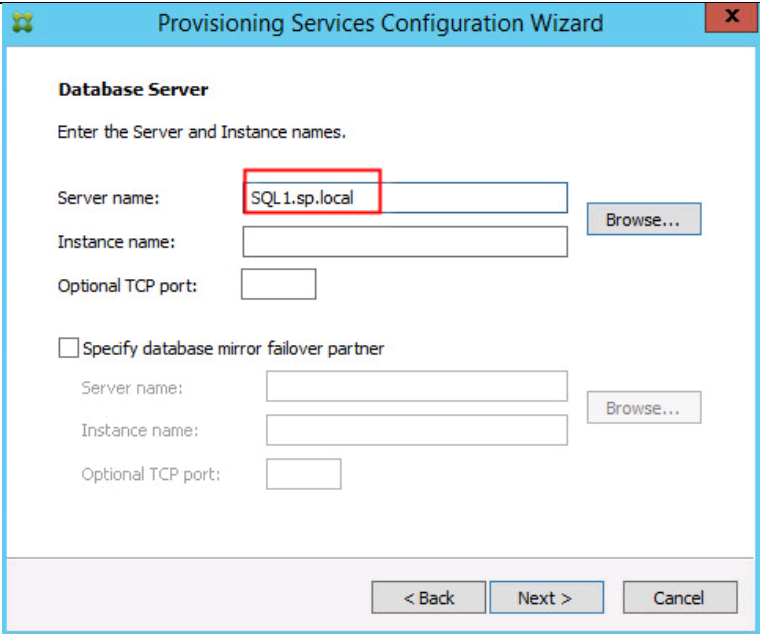
Instructions	Visual
<div>1. Accept the default installation location.</div> <div>2. Click Next</div>	 <p>The screenshot shows the 'Citrix Provisioning Services x64' installation window. The title bar is blue with the Citrix logo. The main area is titled 'Destination Folder' and contains the text: 'Click Next to install to this folder, or click Change to install to a different folder.' Below this, it says 'Install Citrix Provisioning Services x64 to: C:\Program Files\Citrix\Provisioning Services\'. There is a 'Change...' button to the right. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red rectangle and a mouse cursor is pointing at it. The 'InstallShield' logo is visible in the bottom left corner.</p>
<div>3. Click Install to begin the installation.</div>	 <p>The screenshot shows the 'Citrix Provisioning Services x64' installation window at the 'Ready to Install the Program' step. The title bar is blue with the Citrix logo. The main area is titled 'Ready to Install the Program' and contains the text: 'The wizard is ready to begin installation.' and 'Click Install to begin the installation.' Below this, it says: 'If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.' At the bottom, there are three buttons: '< Back', 'Install', and 'Cancel'. The 'Install' button is highlighted with a red rectangle and a mouse cursor is pointing at it. The 'InstallShield' logo is visible in the bottom left corner.</p>

Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Click Finish when the install is complete.</div>	<div>The screenshot shows the 'Citrix Provisioning Services x64' window with the title 'Installation Wizard Completed'. The Citrix logo is on the left. The text states: 'The Installation Wizard has successfully installed Citrix Provisioning Services x64. Click Finish to exit the wizard.' At the bottom, there are three buttons: '< Back', 'Finish', and 'Cancel'. The 'Finish' button is highlighted with a red rectangle.</div>
<div>2. The PVS Configuration Wizard starts automatically.</div> <div>3. Click Next</div>	<div>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The Citrix logo is on the left. The title is 'Welcome to the Configuration Wizard'. The text says: 'The Configuration Wizard provides an easy way to setup a "basic" Server configuration. For advanced configurations, see the Installation and Configuration Guide. You can always run the Configuration Wizard again later from the Start Menu.' At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red rectangle.</div>
Instructions	Visual

Instructions	Visual
<div>1. Since the PVS server is not the DHCP server for the environment, select the radio button labeled, “The service that runs on another computer.”</div> <div>2. Click Next</div>	<div></div>
<div>3. Since this server will be a PXE server, select the radio button labeled, “The service that runs on this computer.”</div> <div>4. Click Next</div>	<div></div>

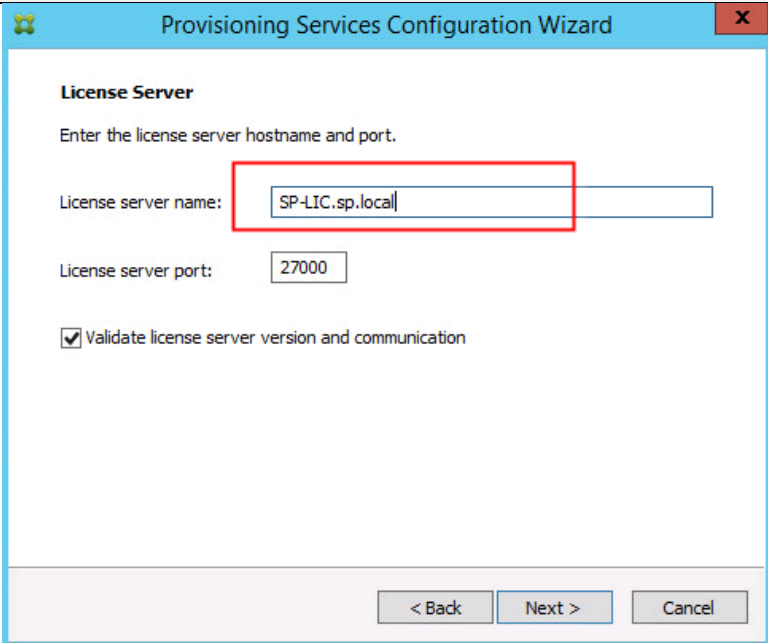
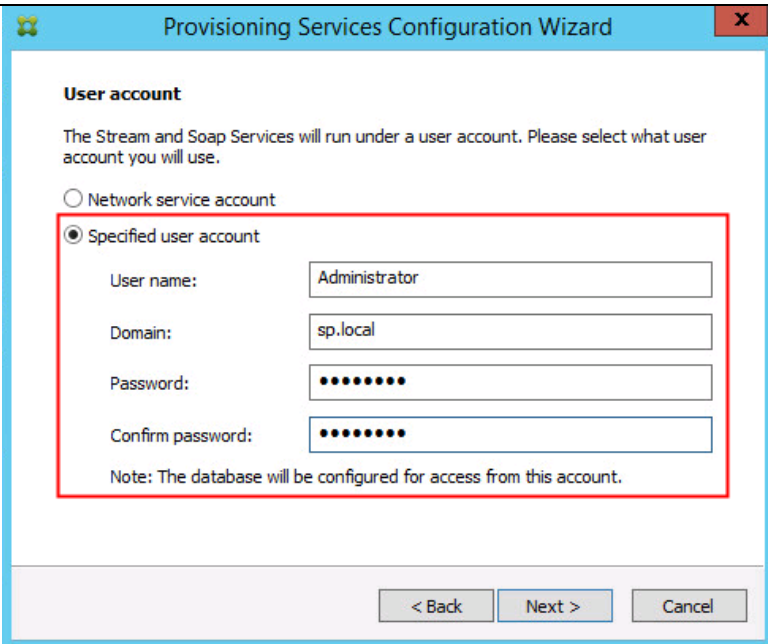
Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Since this is the first server in the farm, select the radio button labeled, “Create farm”.</div> <div>2. Click Next</div>	
<div>3. Enter the FQDN of the SQL server.</div> <div>4. Click Next</div>	


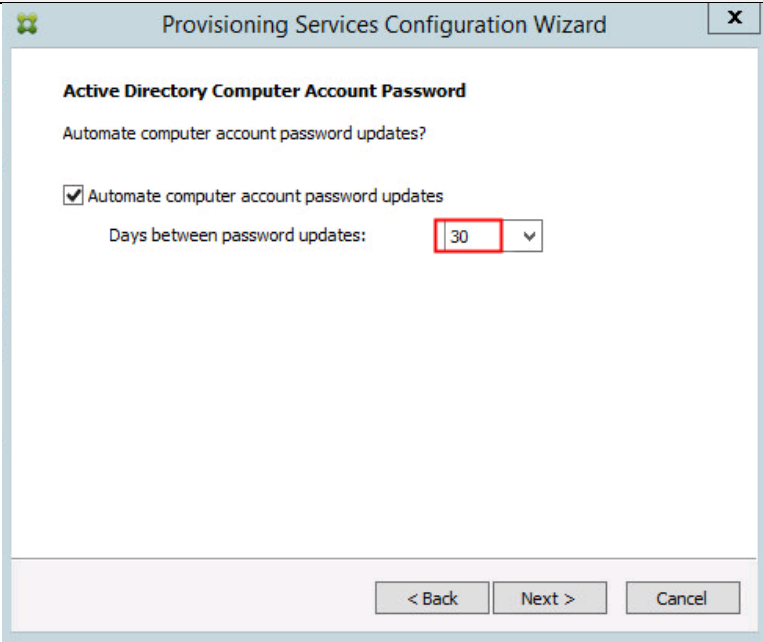
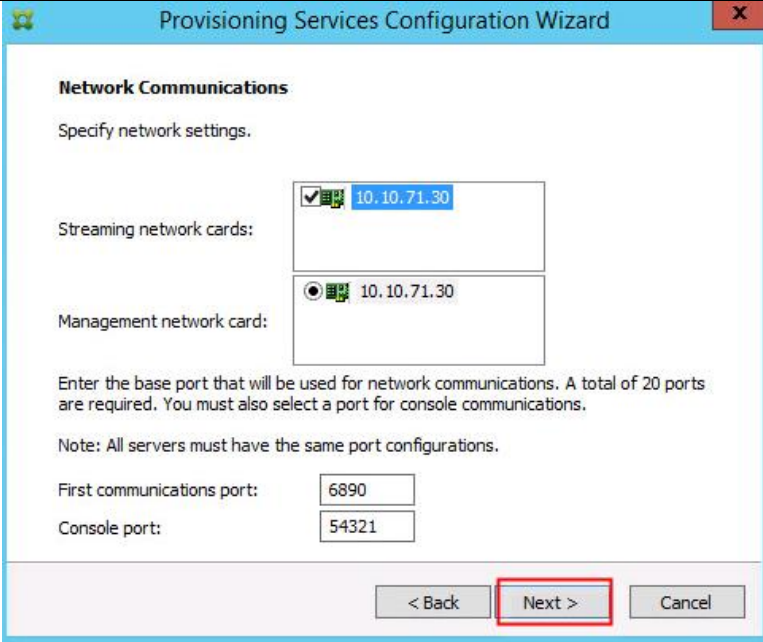
Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Provide a database name.</div> <div>2. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>New Farm</div><div>Enter the new Database and Farm names.</div><div><div>Database name:</div><div>PVS</div></div><div><div>Farm name:</div><div>Farm</div></div><div><div>Site name:</div><div>Site</div></div><div><div>Collection name:</div><div>Collection</div></div><div><div><input checked="" type="radio"/> Use Active Directory groups for security</div><div><input type="radio"/> Use Windows groups for security</div></div><div><div>Farm Administrator group:</div><div>sp.local/Builtin/Administrators</div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div>
<div>3. Provide a vDisk Store name and the storage path to the Nimble vDisk share.</div> <div>*Create the share On a Windows File Server that has a data disk on the Nimble CS700 Array using the Windows File Server Performance Policy</div> <div>4. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>New Store</div><div>Enter a new Store and default path.</div><div><div>Store name:</div><div>Store</div></div><div><div>Default path:</div><div>\\SP-FILE01\vdisk\Store</div><div>Browse...</div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div><div><div>File share data drive on Nimble CS700 Array</div></div></div>

Instructions	Visual
--------------	--------

Instructions	Visual
<ol style="list-style-type: none"> 1. Provide the FQDN of the license server. 2. Optionally, provide a port number if changed on the license server. 3. Click Next 	
<ol style="list-style-type: none"> 4. If an Active Directory service account is not already setup for the PVS servers, create that account prior to clicking Next on this dialog. 5. Select the Specified user account radio button. 6. Complete the User name, Domain, Password, and Confirm password fields, using the PVS account information created earlier. 7. Click Next 	

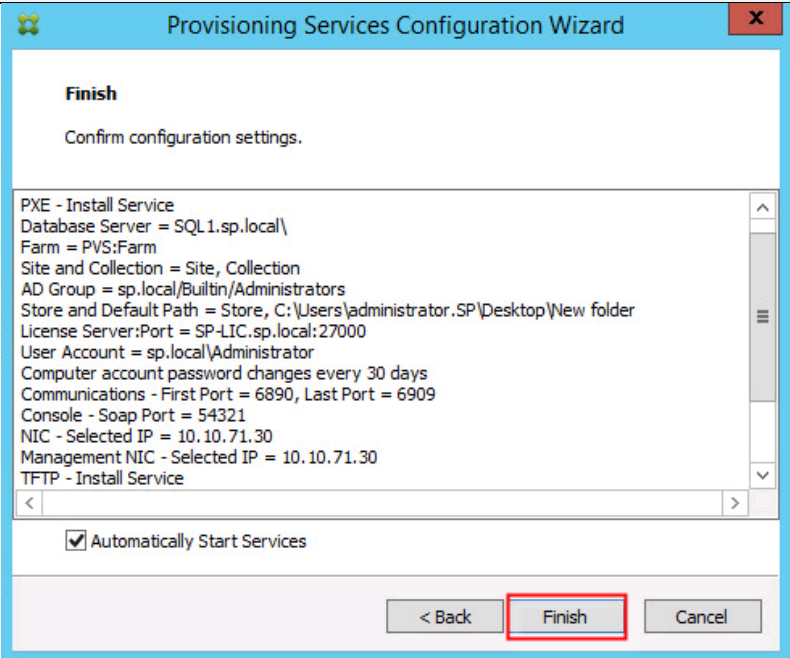
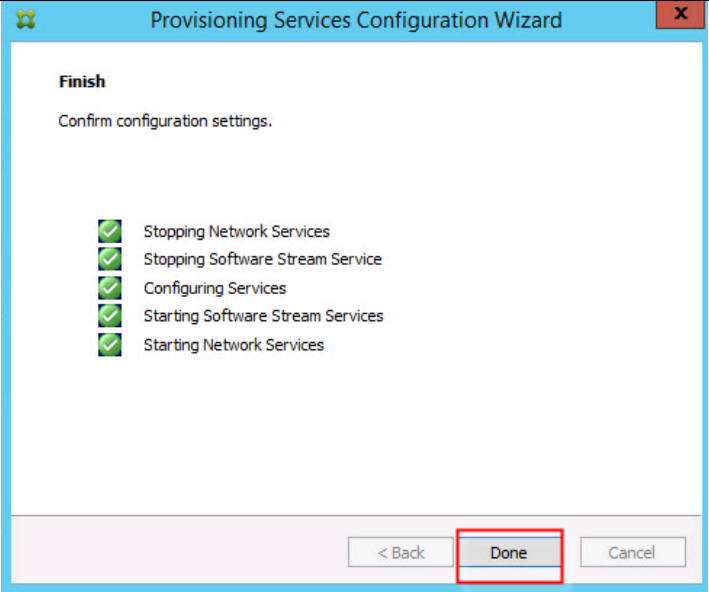
Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Set the Days between password updates to 30.</div> <div> This will vary per environment. "30 days" for the configuration was appropriate for testing purposes.</div> <div>2. Click Next</div>	
<div>3. Keep the defaults for the network cards.</div> <div>4. Click Next</div>	

Instructions	Visual
--------------	--------

Instructions	Visual												
<div>1. Select Use the Provisioning Services TFTP service checkbox.</div> <div>2. Click Next</div>	<div><div>Provisioning Services Configuration Wizard</div><div><div>TFTP Option and Bootstrap Location</div><div>Typically only one TFTP server is deployed as part of Provisioning Services.</div><div><div><input checked="" type="checkbox"/> Use the Provisioning Services TFTP service</div><div>C:\ProgramData\Citrix\Provisioning Services\Tftpboot\ARDBP32.BIN</div><div>Browse...</div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div>												
<div>3. Make sure that the IP Addresses for all PVS servers are listed in the Stream Servers Boot List</div> <div>4. Click Next</div>	<div><div>Provisioning Services Configuration Wizard</div><div><div>Stream Servers Boot List</div><div>Specify at least 1 and at most 4 boot servers.</div><div>The bootstrap file specifies what servers target devices may contact to complete the boot process.</div><div><table><tr><th>Server IP Address</th><th>Server Port</th><th>Device Subnet Mask</th><th>Device Gateway</th></tr><tr><td>10.10.71.30</td><td>6910</td><td></td><td></td></tr><tr><td>10.10.71.31</td><td>6910</td><td></td><td></td></tr></table><div><div>Add</div><div>Edit</div><div>Remove</div><div>Move up</div><div>Move down</div></div><div>Advanced...</div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div>	Server IP Address	Server Port	Device Subnet Mask	Device Gateway	10.10.71.30	6910			10.10.71.31	6910		
Server IP Address	Server Port	Device Subnet Mask	Device Gateway										
10.10.71.30	6910												
10.10.71.31	6910												

Instructions	Visual
--------------	--------

Instructions	Visual
1. Click Finish to start installation.	
2. When the installation is completed, click the Done button.	

Install Additional PVS Servers

Complete the same installation steps on the additional PVS servers up to the configuration step where it asks to Create or Join a farm. In this CVD, we repeated the procedure to add the second PVS servers.

Instructions	Visual
--------------	--------

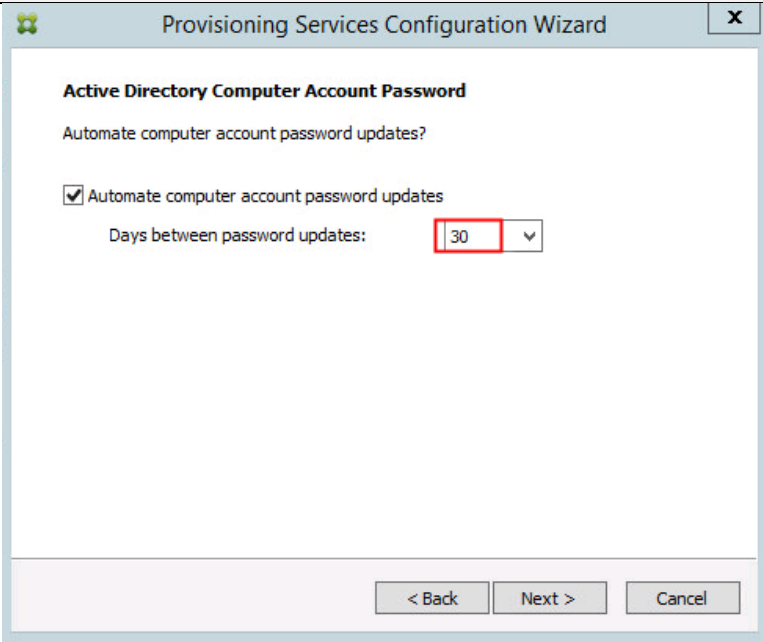
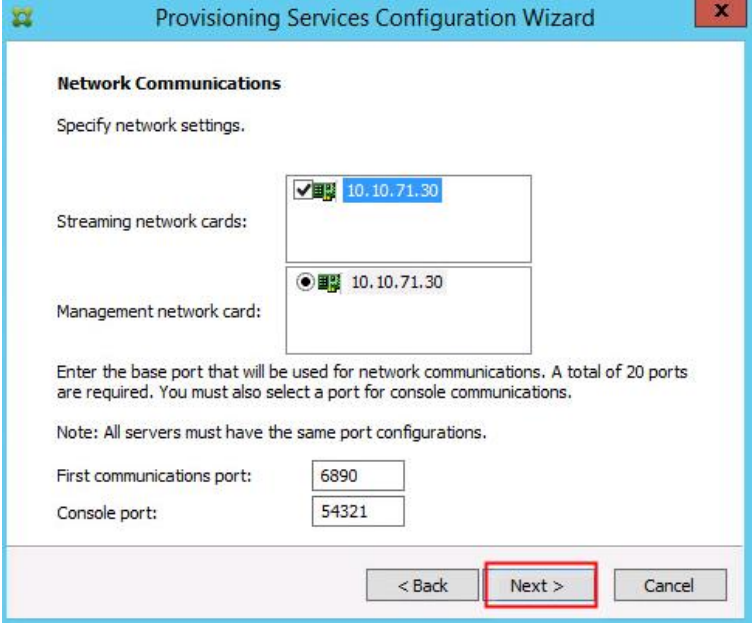
Instructions	Visual
<div>1. On the Farm Configuration dialog, select “Join existing farm.”</div> <div>2. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>Farm Configuration</div><div>Create a new Farm or join an existing Farm. Can be skipped if already configured.</div><div><div><div>Create farm</div><div>Join existing farm</div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div>
<div>3. Provide the FQDN of the SQL Server.</div> <div>4. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>Database Server</div><div>Enter the Server and Instance names.</div><div><div>Server name:</div><div>SQL1.sp.local</div><div>Browse...</div></div><div><div>Instance name:</div><div></div></div><div><div>Optional TCP port:</div><div></div></div><div><div><input type="checkbox"/> Specify database mirror failover partner</div><div><div>Server name:</div><div></div><div>Browse...</div></div><div><div>Instance name:</div><div></div></div><div><div>Optional TCP port:</div><div></div></div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div>
Instructions	Visual

Instructions	Visual
<div>1. Accept the Farm Name.</div> <div>2. Click Next.</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>Existing Farm</div><div>Select the Farm.</div><div>Farm name: PVS:Farm</div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div>
<div>3. Accept the Existing Site.</div> <div>4. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>Site</div><div>Select a Site or enter a new Site and Collection.</div><div><div><div>Existing site</div><div>Site name: Site</div></div><div><div>New site</div><div>Site name:</div><div>Collection name: Collection</div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div>

Instructions	Visual
--------------	--------

Instructions	Visual		
<div>1. Accept the existing vDisk store.</div> <div>2. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>Store</div><div>Select a Store or enter a new Store and default path.</div><div><div><div>Existing store</div><div>Store name: Store</div></div><div><div>New store</div><div>Store name:</div><div>Default path:</div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div> <tr><td><div>3. Provide the PVS service account information.</div><div>4. Click Next</div></td><td><div><div><div>Provisioning Services Configuration Wizard</div><div><div>User account</div><div>The Stream and Soap Services will run under a user account. Please select what user account you will use.</div><div><div>Network service account</div><div><div><div>Specified user account</div><div><div>User name: Administrator</div><div>Domain: sp.local</div><div>Password: </div><div>Confirm password: </div></div><div>Note: The database will be configured for access from this account.</div></div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div></td></tr>	<div>3. Provide the PVS service account information.</div> <div>4. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>User account</div><div>The Stream and Soap Services will run under a user account. Please select what user account you will use.</div><div><div>Network service account</div><div><div><div>Specified user account</div><div><div>User name: Administrator</div><div>Domain: sp.local</div><div>Password: </div><div>Confirm password: </div></div><div>Note: The database will be configured for access from this account.</div></div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div>
<div>3. Provide the PVS service account information.</div> <div>4. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>User account</div><div>The Stream and Soap Services will run under a user account. Please select what user account you will use.</div><div><div>Network service account</div><div><div><div>Specified user account</div><div><div>User name: Administrator</div><div>Domain: sp.local</div><div>Password: </div><div>Confirm password: </div></div><div>Note: The database will be configured for access from this account.</div></div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div>		

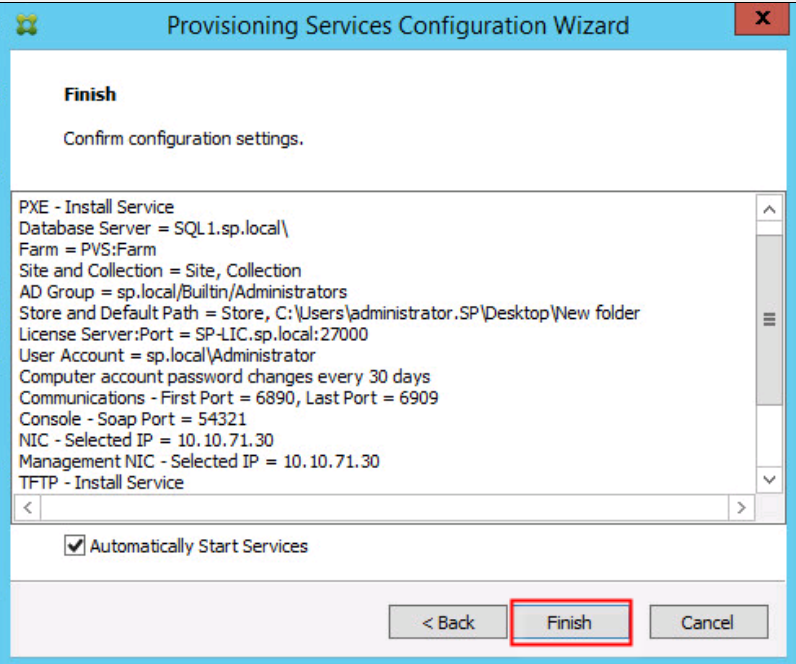
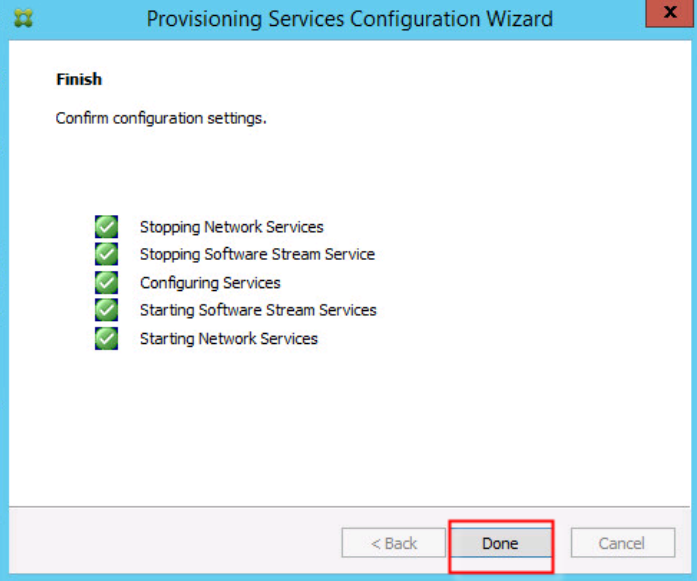
Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Set the Days between password updates to 30.</div> <div>2. Click Next</div>	<div></div>
<div>3. Accept the network card settings.</div> <div>4. Click Next</div>	<div></div>

Instructions	Visual
--------------	--------

Instructions	Visual												
<div>1. Select Use the Provisioning Services TFTP service checkbox.</div> <div>2. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>TFTP Option and Bootstrap Location</div><div>Typically only one TFTP server is deployed as part of Provisioning Services.</div><div><div><input checked="" type="checkbox"/> Use the Provisioning Services TFTP service</div><div>C:\ProgramData\Citrix\Provisioning Services\Tftpboot\ARDBP32.BIN</div><div>Browse...</div></div></div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div>												
<div>3. Make sure that the IP Addresses for all PVS servers are listed in the Stream Servers Boot List</div> <div>4. Click Next</div>	<div><div><div>Provisioning Services Configuration Wizard</div><div><div>Stream Servers Boot List</div><div>Specify at least 1 and at most 4 boot servers.</div><div>The bootstrap file specifies what servers target devices may contact to complete the boot process.</div><div><table><tr><th>Server IP Address</th><th>Server Port</th><th>Device Subnet Mask</th><th>Device Gateway</th></tr><tr><td>10.10.71.30</td><td>6910</td><td></td><td></td></tr><tr><td>10.10.71.31</td><td>6910</td><td></td><td></td></tr></table></div><div><div>Add</div><div>Edit</div><div>Remove</div><div>Move up</div><div>Move down</div></div><div>Advanced...</div><div><div>< Back</div><div>Next ></div><div>Cancel</div></div></div></div></div>	Server IP Address	Server Port	Device Subnet Mask	Device Gateway	10.10.71.30	6910			10.10.71.31	6910		
Server IP Address	Server Port	Device Subnet Mask	Device Gateway										
10.10.71.30	6910												
10.10.71.31	6910												

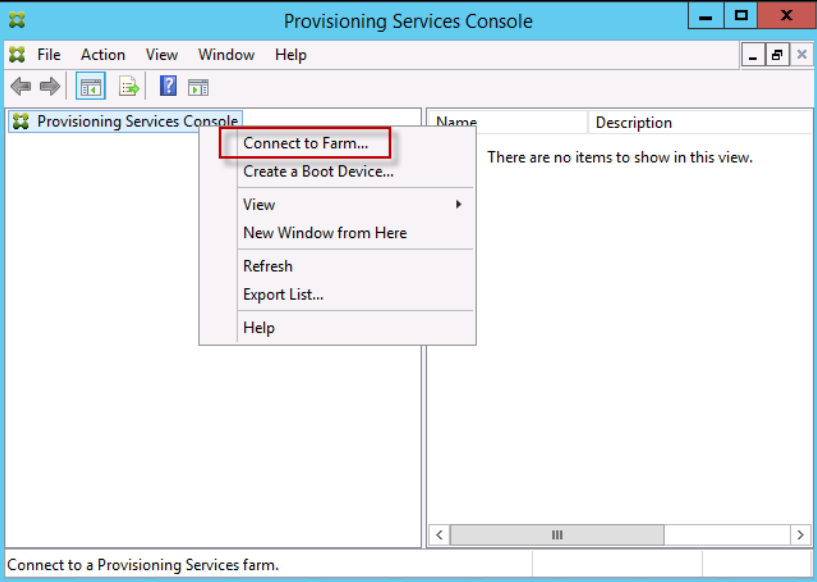
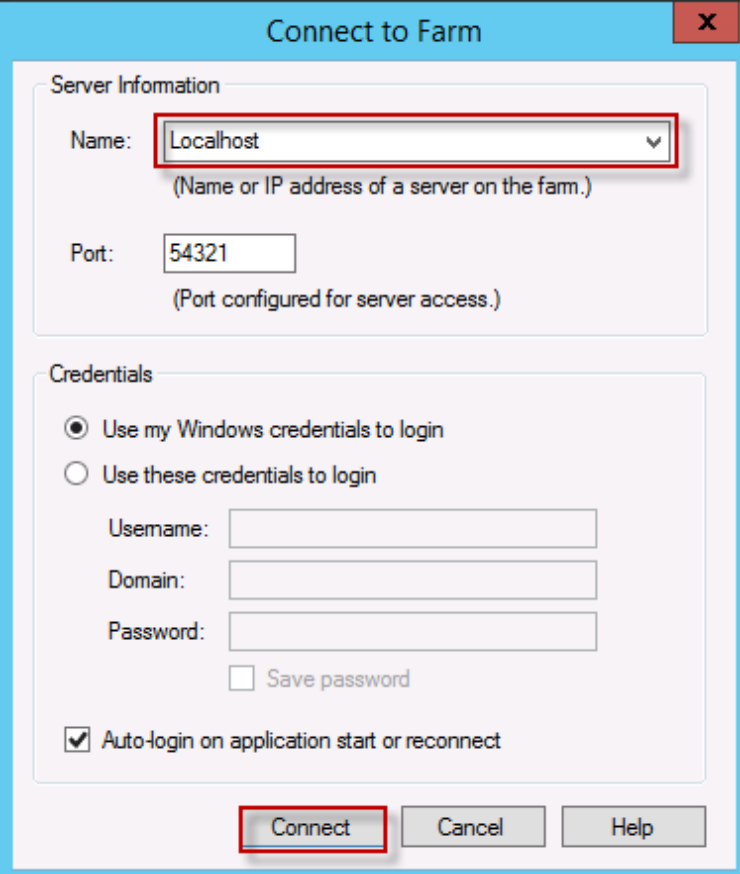
Instructions	Visual
--------------	--------

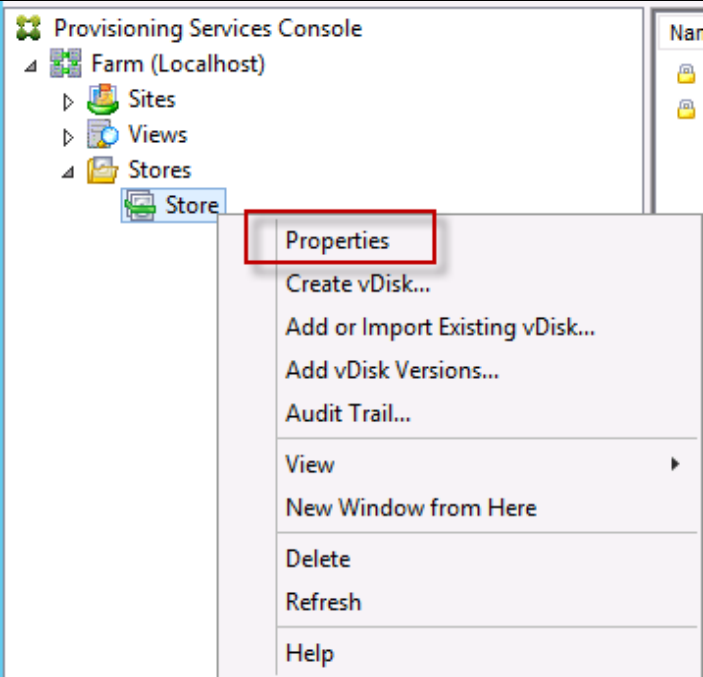
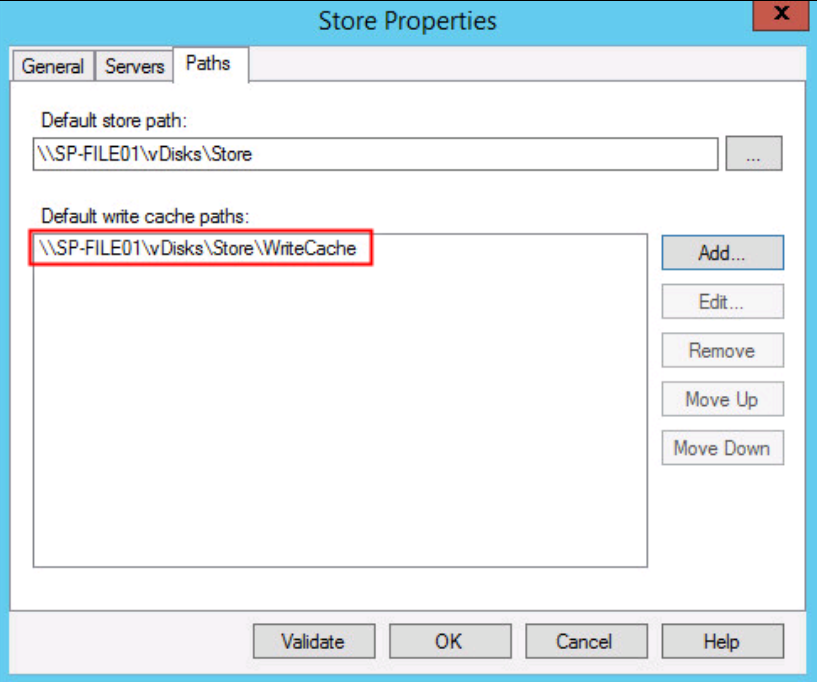
Instructions	Visual
1. Click Finish to start the installation process.	
2. Click Done when the installation finishes	

Optionally, you can install the Provisioning Services console on the second PVS server following the procedure in section Installing Provisioning Services.

After completing the steps to install the second PVS server, launch the Provisioning Services Console to verify that the PVS Servers and Stores are configured and that DHCP boot options are defined.

Instructions	Visual
--------------	--------

Instructions	Visual
1. Launch Provisioning Services Console and select Connect to Farm	 The screenshot shows the 'Provisioning Services Console' application window. The menu bar includes 'File', 'Action', 'View', 'Window', and 'Help'. The 'File' menu is open, and the 'Connect to Farm...' option is highlighted with a red rectangular box. Other menu items include 'Create a Boot Device...', 'View', 'New Window from Here', 'Refresh', 'Export List...', and 'Help'. The main pane shows a table with columns 'Name' and 'Description', and a message 'There are no items to show in this view.' The status bar at the bottom says 'Connect to a Provisioning Services farm.'
2. Enter localhost for the PVS1 server 3. Click Connect	 The screenshot shows the 'Connect to Farm' dialog box. It has two sections: 'Server Information' and 'Credentials'. In the 'Server Information' section, the 'Name' dropdown is set to 'Localhost' and is highlighted with a red box. Below it is the 'Port' field set to '54321'. In the 'Credentials' section, the radio button for 'Use my Windows credentials to login' is selected. There are input fields for 'Username', 'Domain', and 'Password'. A 'Save password' checkbox is unchecked. The 'Auto-login on application start or reconnect' checkbox is checked. At the bottom, the 'Connect' button is highlighted with a red box, along with 'Cancel' and 'Help' buttons.
Instructions	Visual

Instructions	Visual
1. Select Store Properties from the drop-down menu	
2. In the Store Properties dialog, add the Default store path to the list of Default write cache paths. 3. Click Validate. If the validation is successful, click OK to continue.	

Preparing the Master Targets

This section provides guidance around creating the golden (or master) images for the environment. VMs for the master targets must first be installed with the software components needed to build the golden images. For this CVD, the images contain the basics needed to run the Login VSI workload.

The master target HVD(VDI) VMs were configured as follows in Table 5


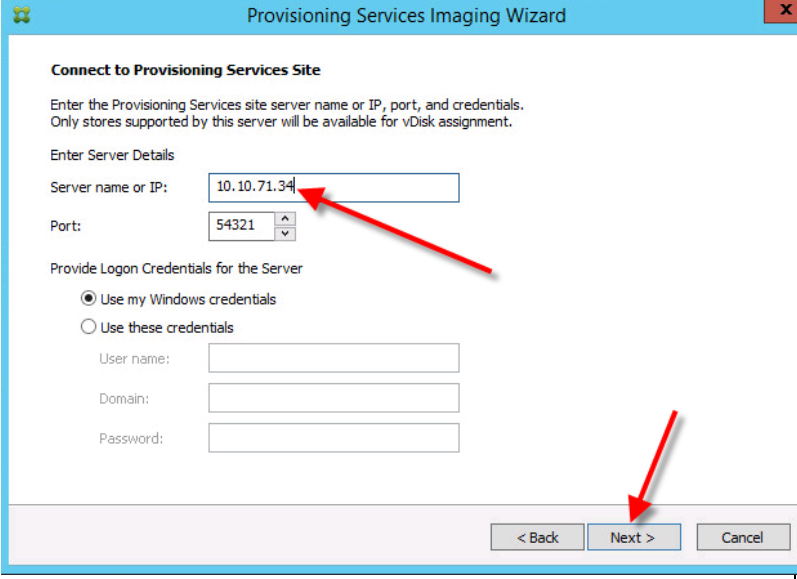
Table 5 VDI and RDS Configurations

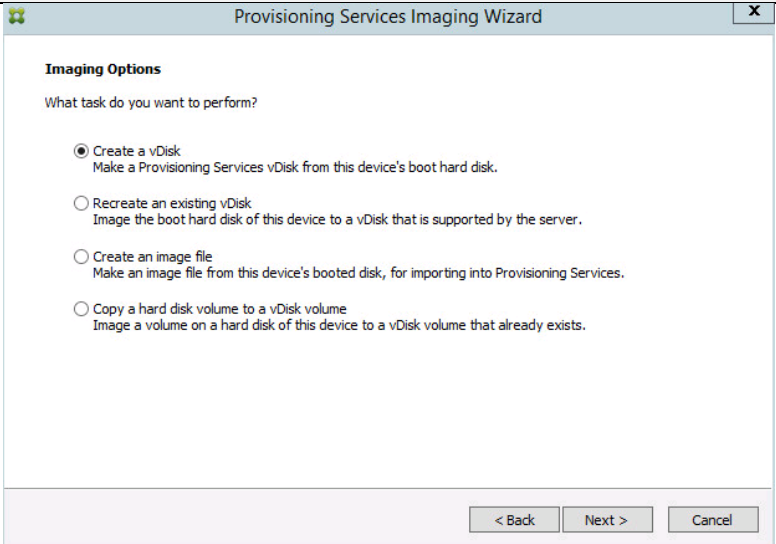
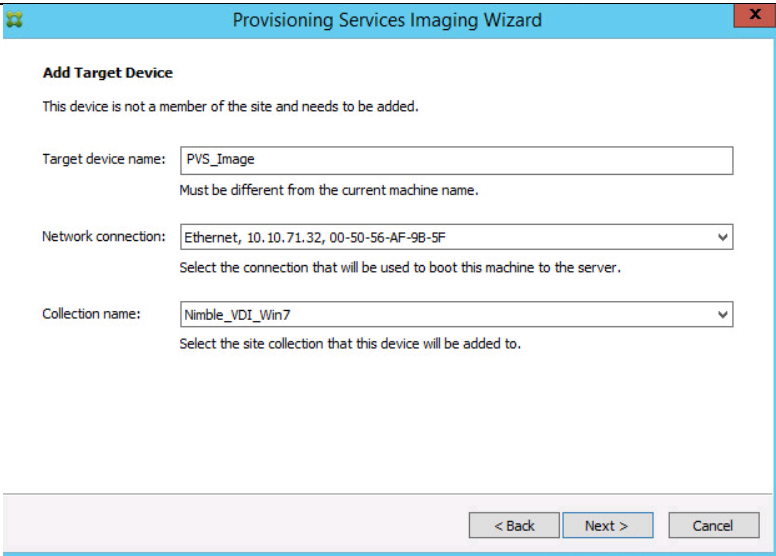
Configuration	VDI Virtual Machines	
Operating system	Microsoft Windows 7 SP1 32-bit	
Virtual CPU amount	2	
Memory amount	1.5-GB reserve for all guest memory	
Network	VMXNET3 VDI vLAN	
Citrix PVS vDisk size and location	24 GB (dynamic) PVS-vDisk volume	
Citrix PVS write cache Disk size	4 GB	
Citrix PVS write cache RAM cache size	64 MB	
Additional software used for testing	Microsoft Office 2010 Login VSI 4.1.4	

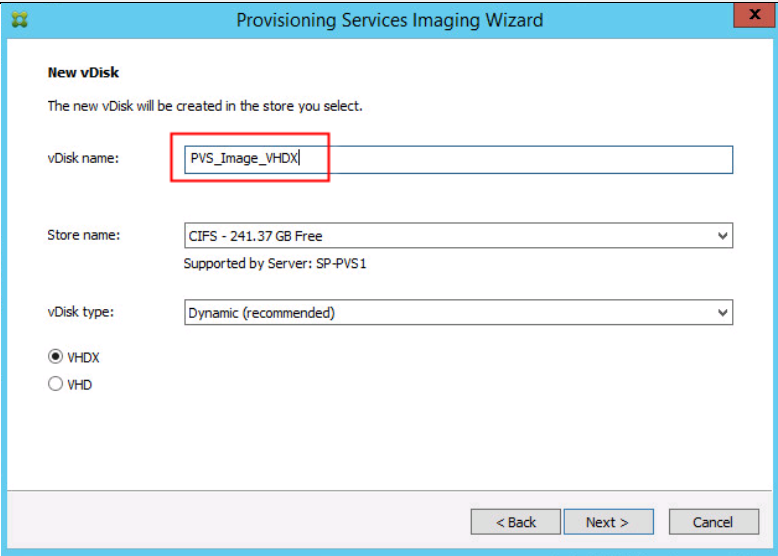
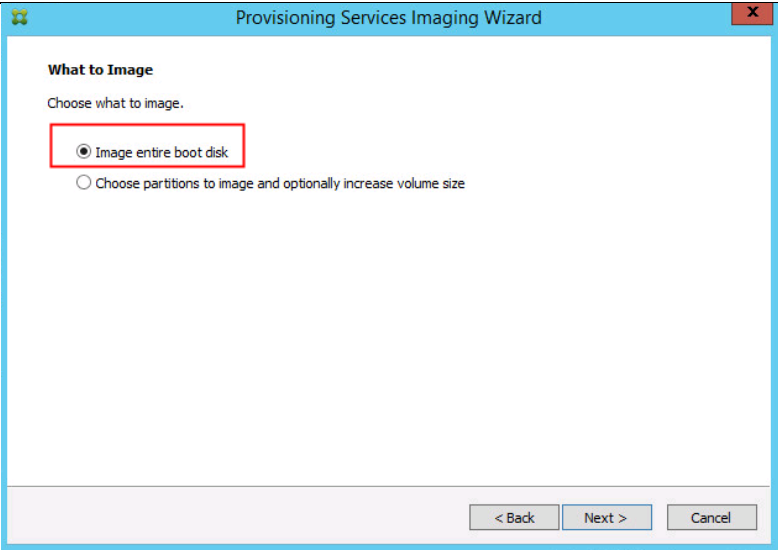
Create the Provisioning Services Master Image

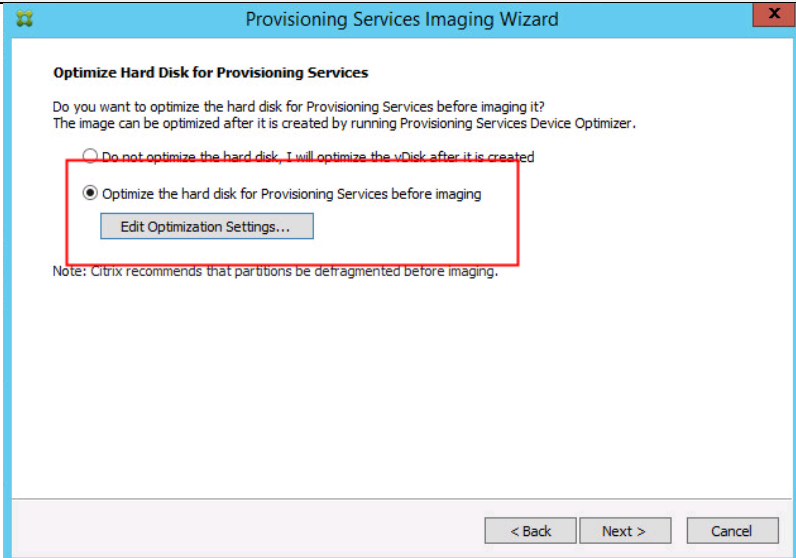
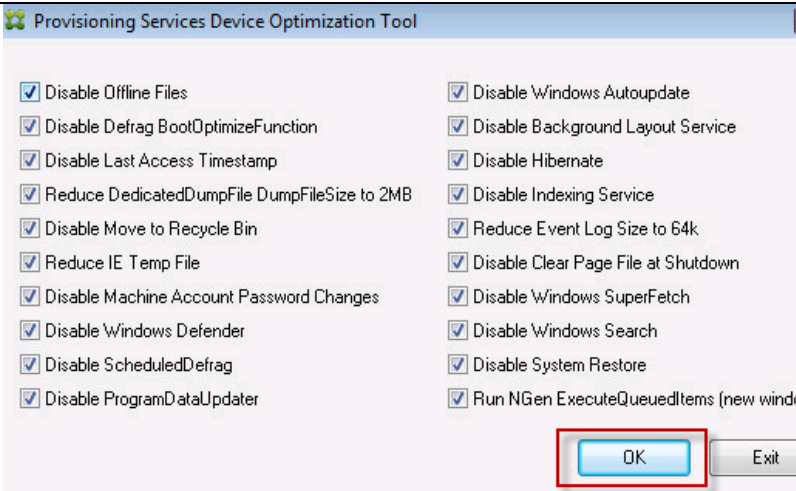
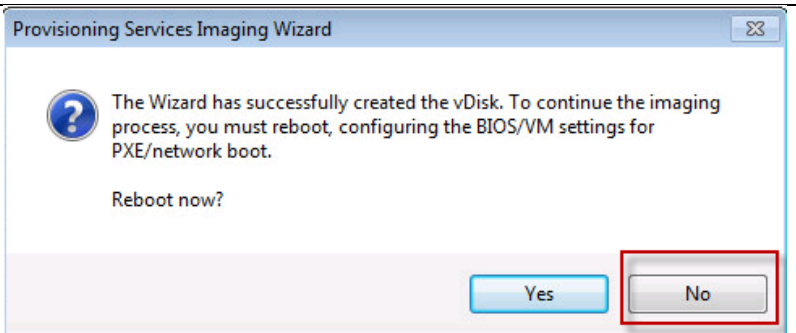
Creating PVS vDisks

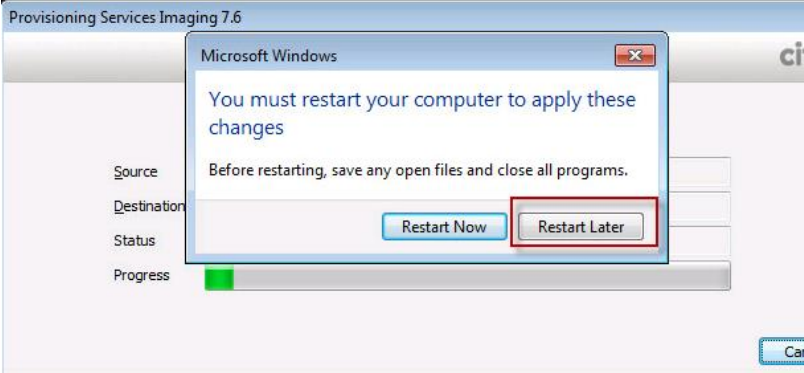
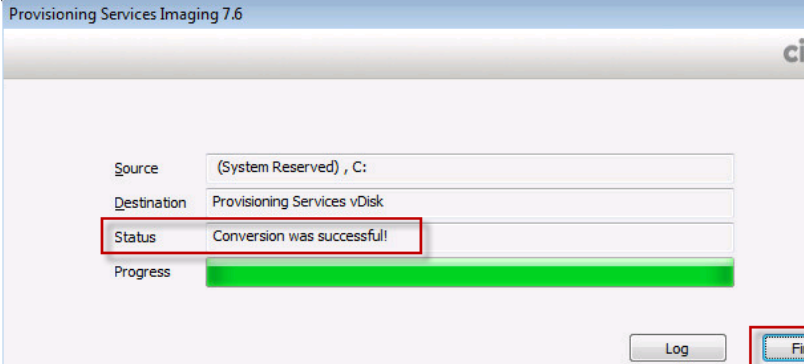
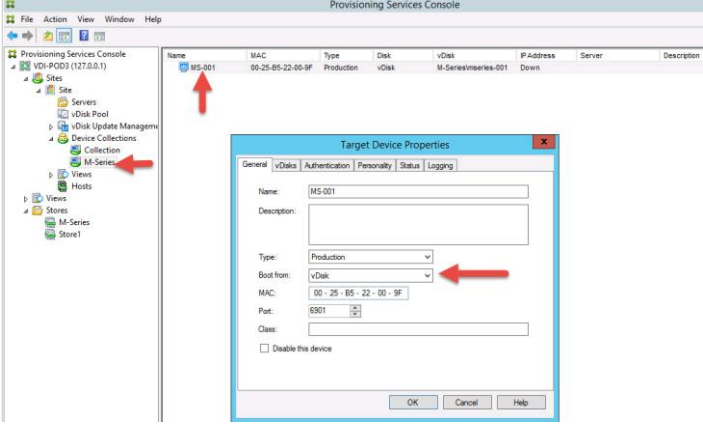
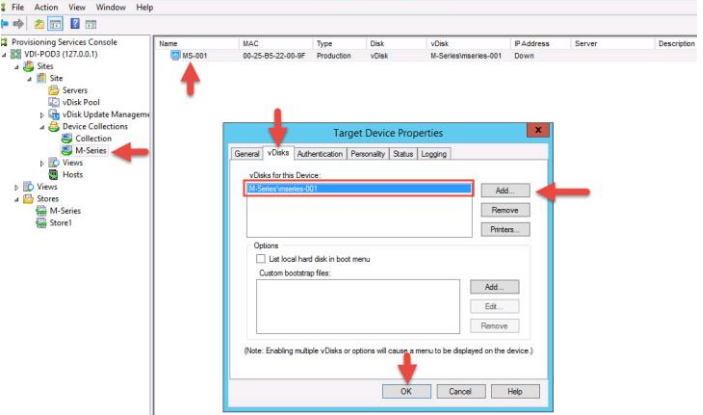
Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. The PVS Imaging Wizard's Welcome page appears.</div> <div>2. Click Next</div>	
<div>3. The Connect to Farm page appears. Enter the name or IP address of a Provisioning Server within the farm to connect to and the port to use to make that connection.</div> <div>4. Use the Windows credentials (default) or enter different credentials.</div> <div>5. Click Next.</div>	

Instructions	Visual
<ol style="list-style-type: none">1. Select Create new vDisk.2. Click Next.	 <p>The screenshot shows the 'Provisioning Services Imaging Wizard' window. The title bar is light blue with a close button (X) on the right. The main content area is white and titled 'Imaging Options'. Below the title, it asks 'What task do you want to perform?'. There are four radio button options: 'Create a vDisk' (selected), 'Recreate an existing vDisk', 'Create an image file', and 'Copy a hard disk volume to a vDisk volume'. Each option has a brief description. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.</p>
	 <p>The screenshot shows the 'Provisioning Services Imaging Wizard' window. The title bar is light blue with a close button (X) on the right. The main content area is white and titled 'Add Target Device'. Below the title, it says 'This device is not a member of the site and needs to be added.' There are three input fields: 'Target device name' (text box with 'PVS_Image'), 'Network connection' (dropdown menu with 'Ethernet, 10.10.71.32, 00-50-56-AF-9B-5F'), and 'Collection name' (dropdown menu with 'Nimble_VDI_Win7'). Each field has a brief description. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.</p>

Instructions	Visual
<div>1. Define volume sizes on the Configure Image Volumes page.</div> <div>2. Click Next.</div>	 <p>The screenshot shows the 'New vDisk' step of the Provisioning Services Imaging Wizard. The title bar says 'Provisioning Services Imaging Wizard'. Below the title, it says 'New vDisk' and 'The new vDisk will be created in the store you select.' There are three input fields: 'vDisk name:' with the value 'PVS_Image_VHDX' (highlighted by a red box), 'Store name:' with a dropdown menu showing 'CIFS - 241.37 GB Free', and 'vDisk type:' with a dropdown menu showing 'Dynamic (recommended)'. Below these fields are two radio buttons: 'VHDX' (selected) and 'VHD'. At the bottom right are three buttons: '< Back', 'Next >', and 'Cancel'.</p>
<div>3. The Add Target Device page appears.</div> <div>4. Select the Target Device Name, the MAC address associated with one of the NICs that was selected when the target device software was installed on the master target device, and the Collection to which you are adding the device.</div> <div>5. Click Next.</div>	 <p>The screenshot shows the 'What to Image' step of the Provisioning Services Imaging Wizard. The title bar says 'Provisioning Services Imaging Wizard'. Below the title, it says 'What to Image' and 'Choose what to image.' There are two radio buttons: 'Image entire boot disk' (selected and highlighted by a red box) and 'Choose partitions to image and optionally increase volume size'. At the bottom right are three buttons: '< Back', 'Next >', and 'Cancel'.</p>

Instructions	Visual
<p>A Summary of Farm Changes appears.</p> <p>1. Select Optimize for Provisioning Services</p>	
<p>2. The PVS Optimization Tool appears. Select the appropriate optimizations and click OK.</p> <p>3. Review the configuration and click Finish</p>	
<p>The vDisk creation process begins. A dialog appears when the creation process is complete.</p> <p>1. At the reboot now prompt select No and manually shutdown the machine.</p>	

Instructions	Visual
<p>After restarting, log into the VDI or RDS master target. The PVS Imaging conversion process begins, converting C: to the PVS vDisk.</p> <p>1. If prompted to Restart select Restart Later</p>	
<p>A message is displayed when the conversion is complete</p> <p>2. Click Finish</p> <p>Machine can be turned off.</p>	
<p>3. Return to the PVS console, in the properties of the VM, change the Boot from: field to vDisk on the General tab.</p>	
<p>4. Add the newly created vDisk on the vDisks tab and click OK.</p>	

Instructions	Visual
5. Boot the machine from the vDisk and verify there are no errors. Leave the vDisk in 'Private Mode' so changes will be saved. Next step is to install the XenDesktop 7.7 VDA.	

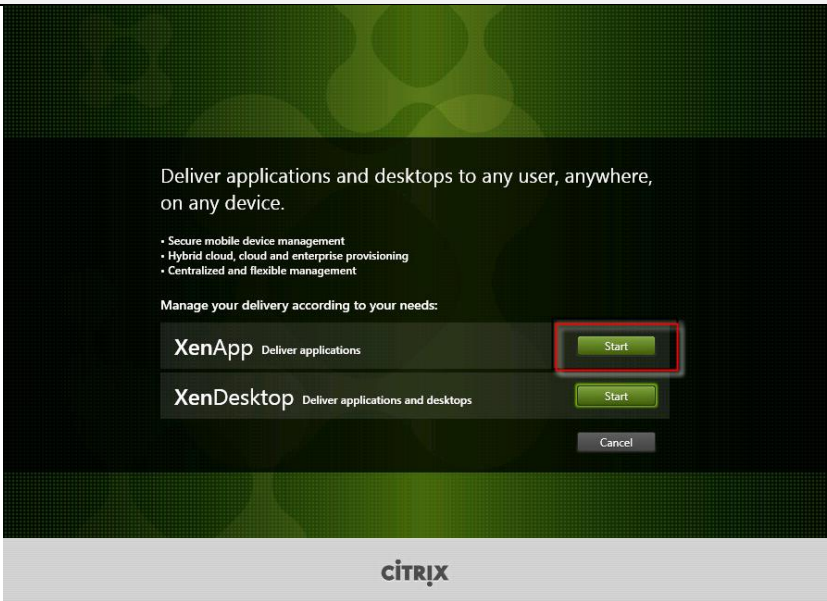
Installing and Configuring Citrix XenDesktop/XenApp 7.7

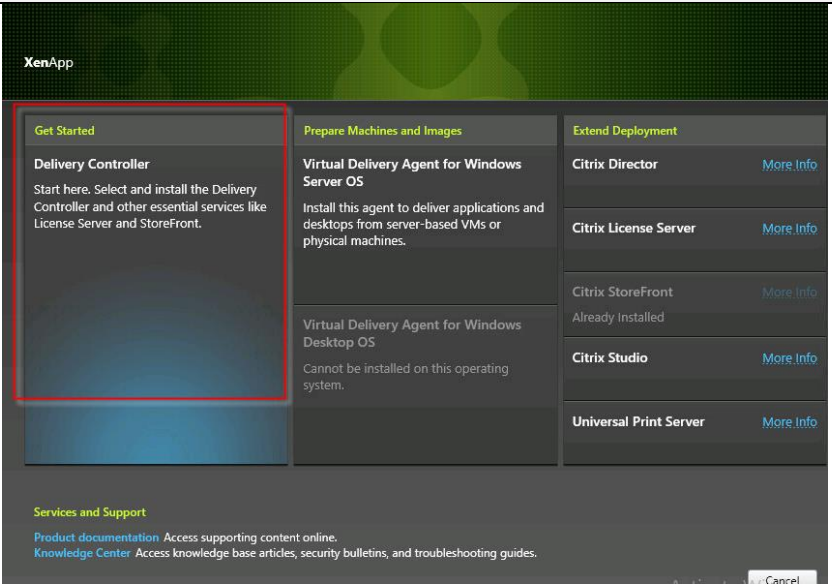
Each Citrix Infrastructure component runs on a VMware ESXi 6 virtual Windows 2012 R2 virtual machine on the Customer's existing infrastructure. The configuration for each virtual machine is as follows:

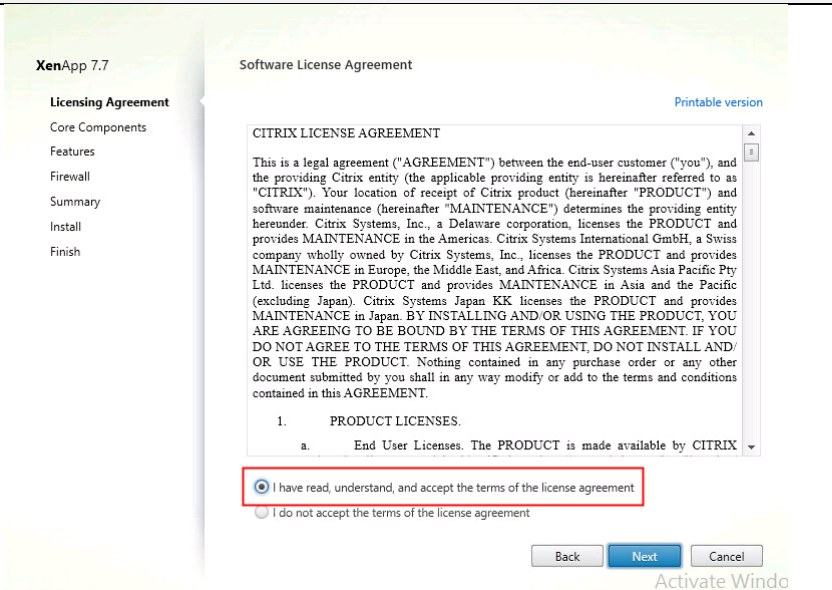
- Citrix XenApp Delivery Controllers (2) 2 vCPUs, 16GB RAM, 1 10Gb vNIC, 50GB Thick Provisioned vDisk
- Citrix Licensing Server (1) 1 vCPU, 4GB RAM, 1 10Gb vNIC, 30GB Thick Provisioned vDisk
- Citrix StoreFront Server (2) 2 vCPU, 4GB RAM, 1 10Gb vNIC, 30GB Thick Provisioned vDisk
- Citrix Provisioning Server (2) 2 vCPU 32GB RAM, 1 10Gb vNIC, 50GB Thick Provisioned vDisk

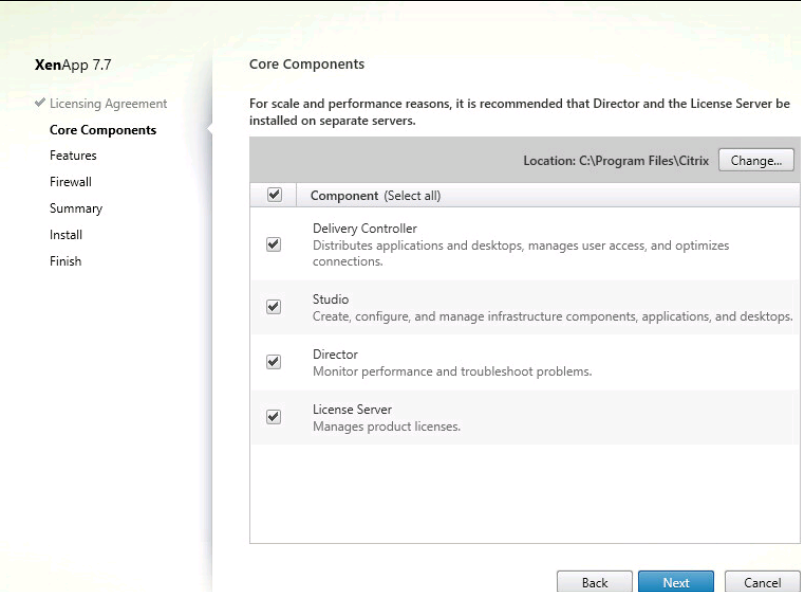
Install XenApp Delivery Controller, Citrix Licensing and StoreFront

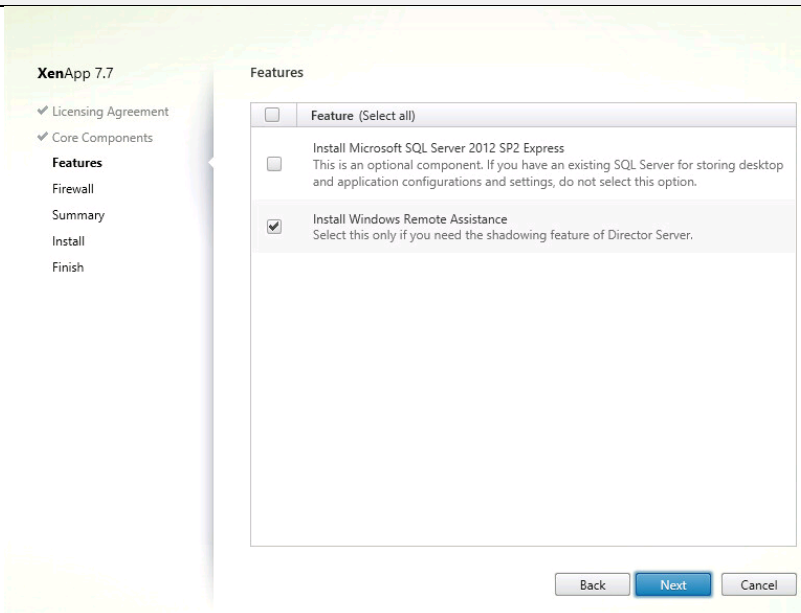
The process of installing the XenApp Delivery Controller also installs other key XenApp software components, including Studio, which is used to create and manage infrastructure components, and Director, which is used to monitor performance and troubleshoot problems.

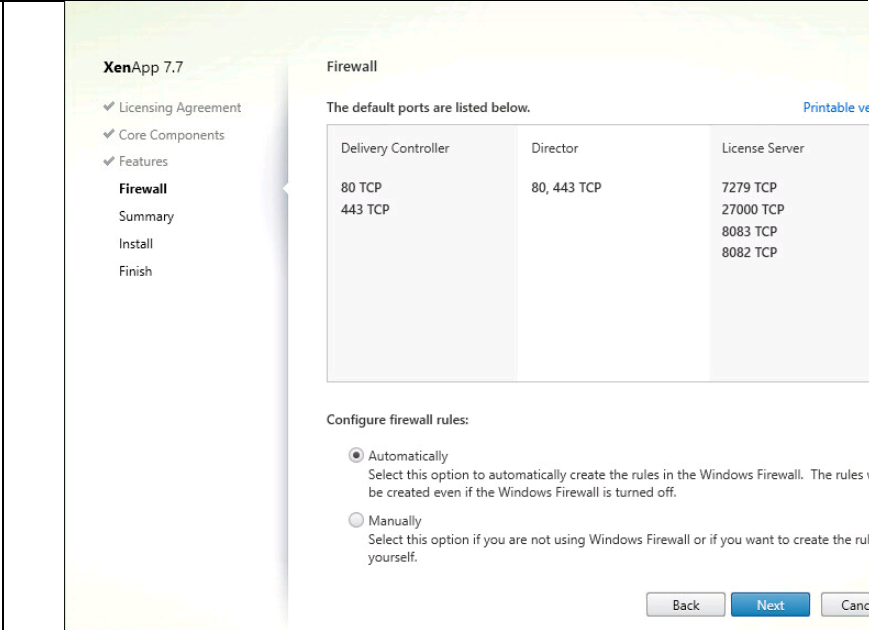
Instructions	Visual
<p>1. To begin the installation, connect to the first XenApp server and launch the installer from the Citrix XenApp 7.7 ISO.</p> <p>2. Click Start</p>	

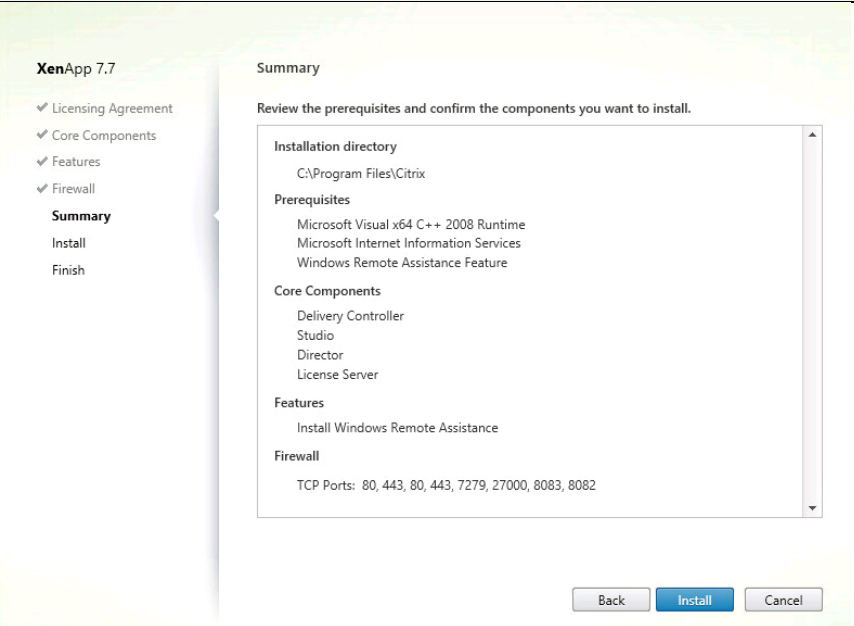
Instructions	Visual
<p>The installation wizard presents a menu with three subsections.</p> <p>3. Click “Get Started – Delivery Controller.”</p>	

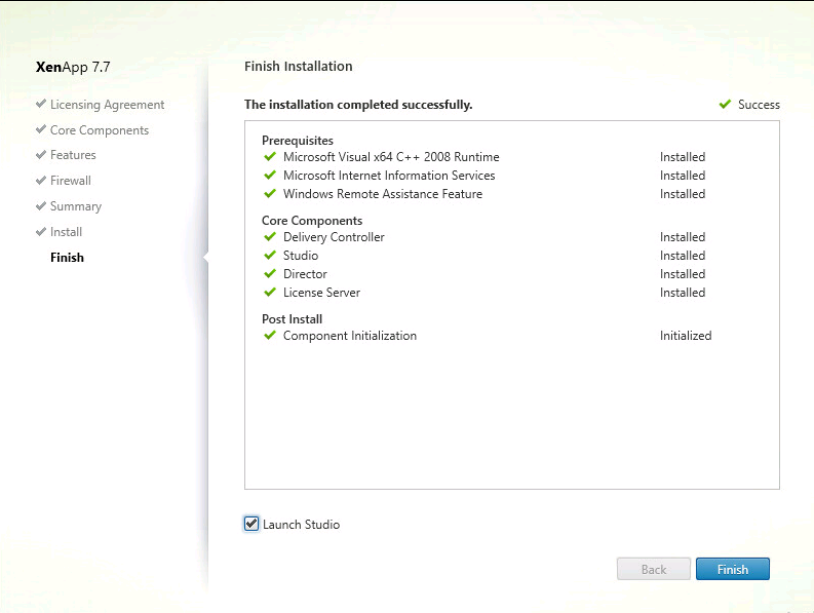
Instructions	Visual
<p>1. Read the Citrix License Agreement.</p> <p>2. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.</p> <p>3. Click Next</p>	

Instructions	Visual
<p>4. Select the components to be installed:</p> <p>Delivery Controller</p> <p>Studio</p> <p>License Server</p> <p>StoreFront</p> <p>5. Click Next</p>	

Instructions	Visual
<p>1. Since a SQL Server will be used to Store the Database, leave “Install Microsoft SQL Server 2012 SP1 Express” unchecked.</p> <p>2. Click Next</p>	

Instructions	Visual
<div>3. Select the default ports and automatically configured firewall rules.</div> <div>4. Click Next</div>	

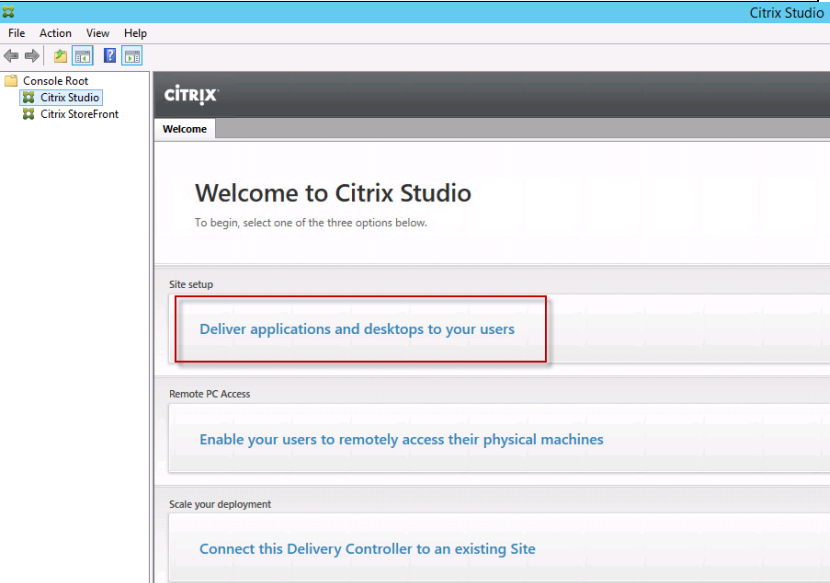
Instructions	Visual
<div>The Summary screen is shown.</div> <div>1. Click the Install button to begin the installation.</div>	

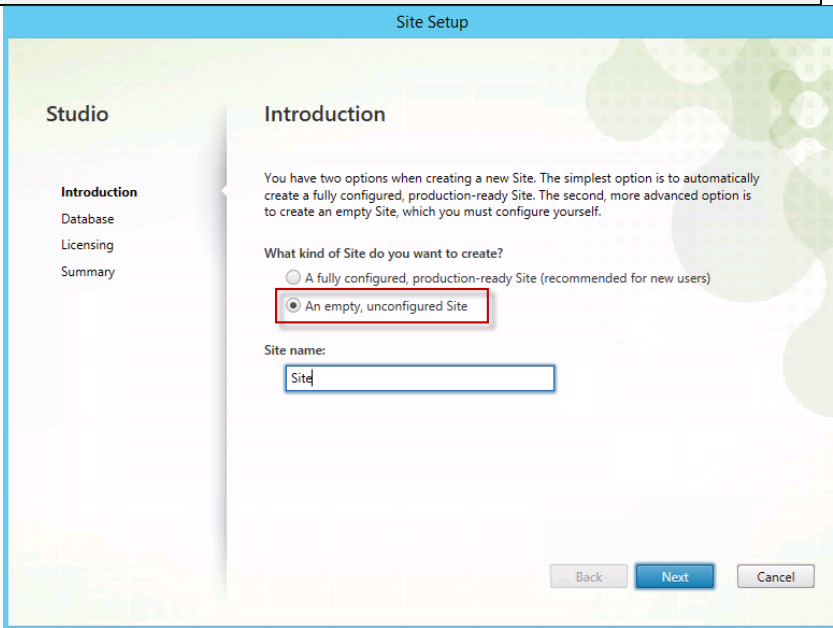
Instructions	Visual
<p>The installer displays a message when the installation is complete</p> <p>2. Click Finish</p> <p>3. (Optional) Check Launch Studio to launch Citrix Studio Console.</p>	

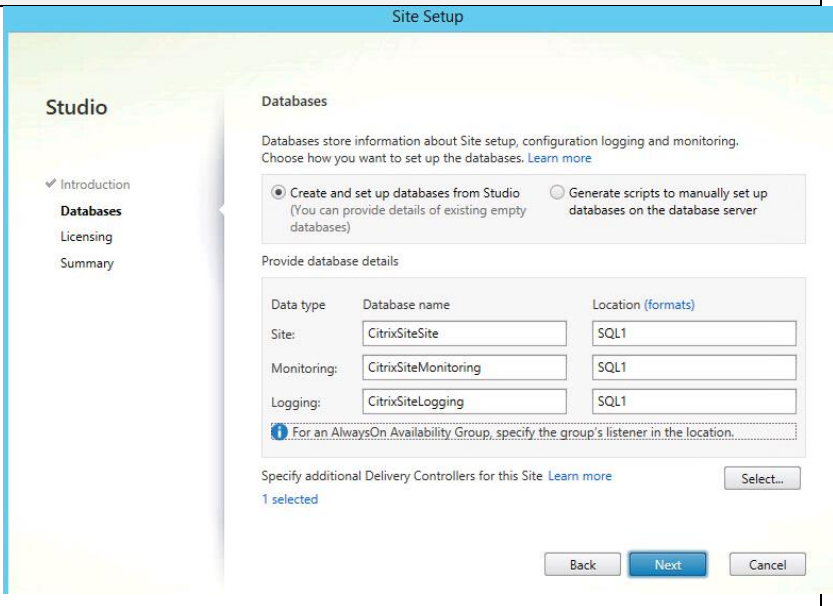
Configure the XenDesktop Site

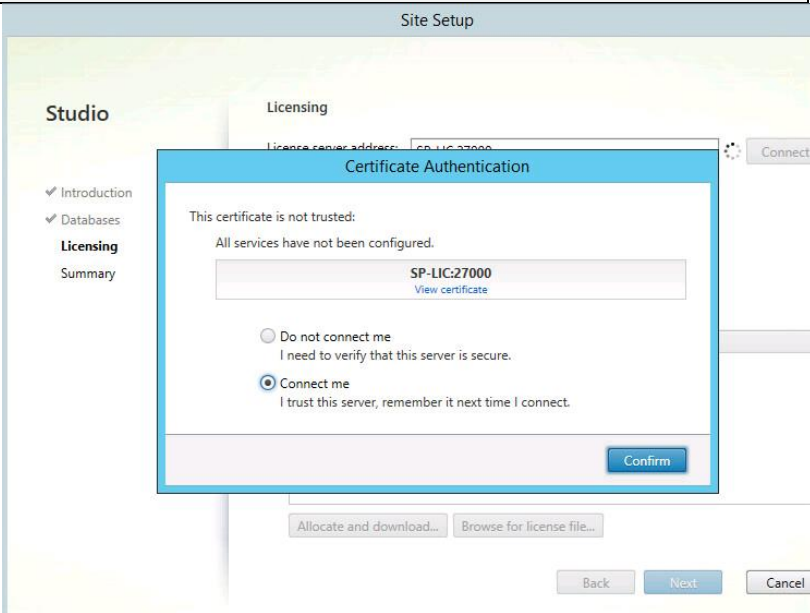
Citrix Studio is a management console that allows you to create and manage infrastructure and resources to deliver desktops and applications. Replacing Desktop Studio from earlier releases, it provides wizards to set up your environment, create workloads to host applications and desktops, and assign applications and desktops to users.

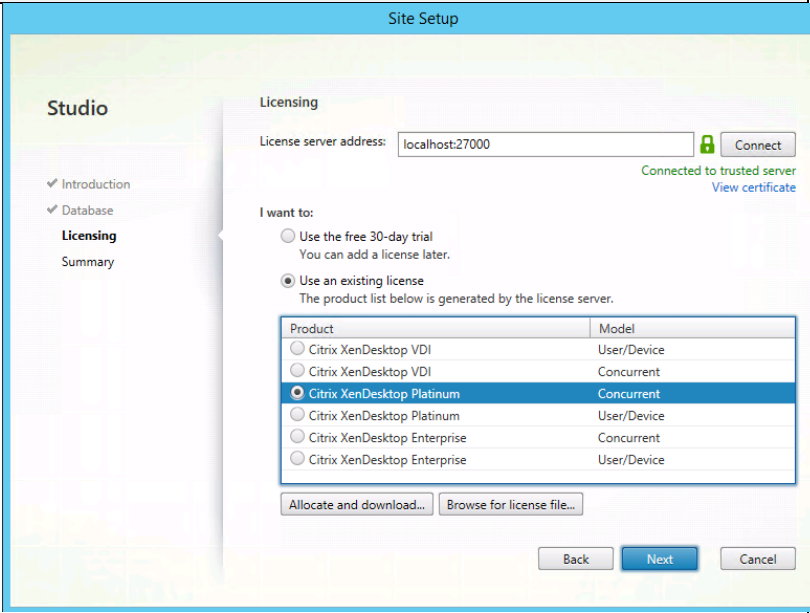
Citrix Studio launches automatically after the XenDesktop Delivery Controller installation, or if necessary, it can be launched manually. Studio is used to create a Site, which is the core XenApp 7.7 environment consisting of the Delivery Controller and the Database.

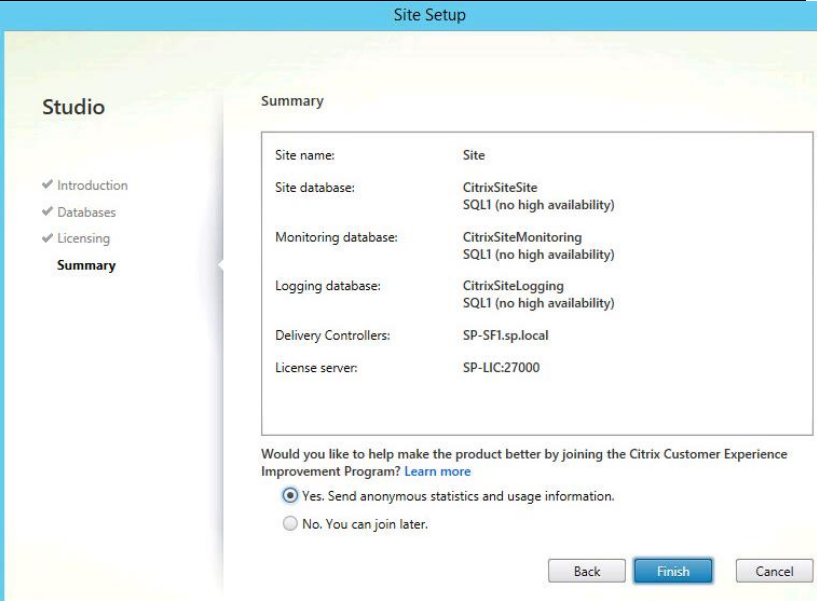
Instructions	Visual
<p>1. Click the Deliver applications and desktops to your users button.</p>	

Instructions	Visual
<p>2. Select the “An empty, unconfigured Site” radio button.</p> <p>3. Enter a site name.</p> <p>4. Click Next</p>	

Instructions	Visual
<p>1. Provide the Database Server location.</p> <p>2. Click the Next.</p>	

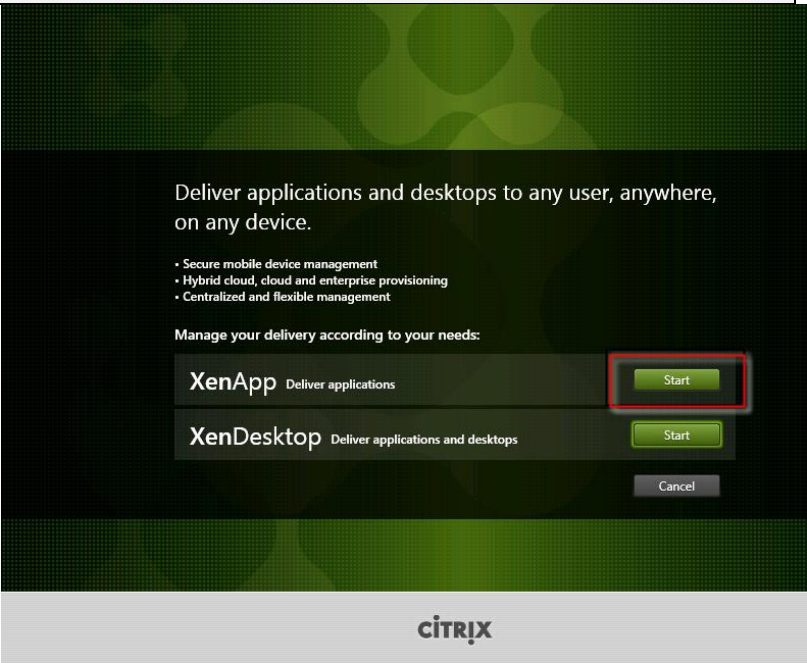
Instructions	Visual
<p>3. Enter the License server, accept Certificate warning if necessary and click Confirm.</p>	 <p>The screenshot shows the 'Site Setup' window in Citrix Studio. The 'Licensing' tab is active. A 'Certificate Authentication' dialog box is displayed in the foreground. The dialog states: 'This certificate is not trusted: All services have not been configured.' It shows a certificate identifier 'SP-LIC:27000' with a 'View certificate' link. There are two radio buttons: 'Do not connect me' (unselected) and 'Connect me' (selected). Below the radio buttons, it says 'I trust this server, remember it next time I connect.' A 'Confirm' button is at the bottom right of the dialog. In the background, the 'Licensing' section of the 'Site Setup' window is visible, showing the 'License server address' as 'localhost:27000' and a 'Connect' button. The left sidebar shows 'Studio' with 'Introduction', 'Databases', 'Licensing', and 'Summary' options.</p>

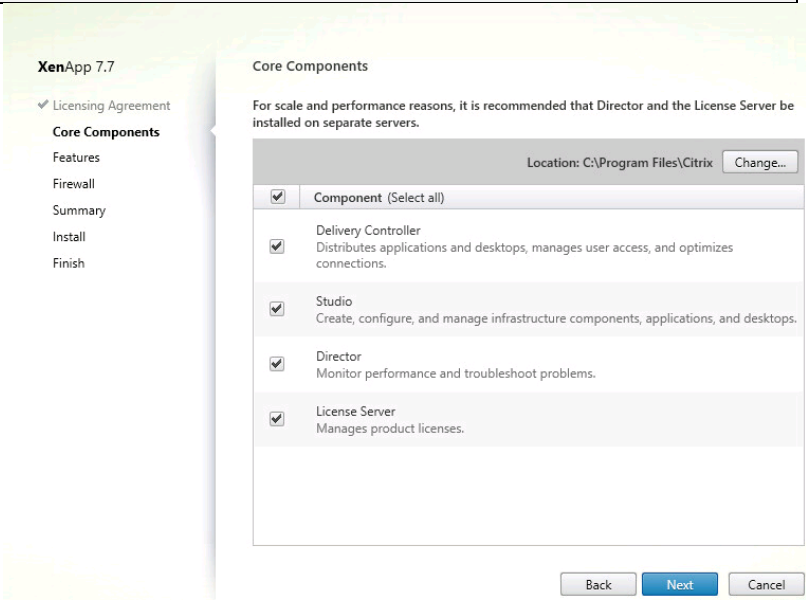
Instructions	Visual														
<p>1. Provide the FQDN of the license server.</p> <p>2. Click Connect to validate and retrieve any licenses from the server.</p> <p>If no licenses are available, you can use the 30-day free trial or activate a license file.</p> <p>3. Select the appropriate product edition using the license radio button</p> <p>4. Click Next</p>	 <p>The screenshot shows the 'Site Setup' window in Citrix Studio. The 'Licensing' tab is active. The 'License server address' is 'localhost:27000' and the 'Connect' button is highlighted with a green lock icon and the text 'Connected to trusted server'. Below this, the 'I want to:' section has two radio buttons: 'Use the free 30-day trial' (unselected) and 'Use an existing license' (selected). Below the radio buttons, it says 'The product list below is generated by the license server.' A table is displayed with the following data:</p> <table border="1"> <thead> <tr> <th>Product</th> <th>Model</th> </tr> </thead> <tbody> <tr> <td><input type="radio"/> Citrix XenDesktop VDI</td> <td>User/Device</td> </tr> <tr> <td><input type="radio"/> Citrix XenDesktop VDI</td> <td>Concurrent</td> </tr> <tr> <td><input checked="" type="radio"/> Citrix XenDesktop Platinum</td> <td>Concurrent</td> </tr> <tr> <td><input type="radio"/> Citrix XenDesktop Platinum</td> <td>User/Device</td> </tr> <tr> <td><input type="radio"/> Citrix XenDesktop Enterprise</td> <td>Concurrent</td> </tr> <tr> <td><input type="radio"/> Citrix XenDesktop Enterprise</td> <td>User/Device</td> </tr> </tbody> </table> <p>At the bottom of the table, there are 'Allocate and download...' and 'Browse for license file...' buttons. The 'Next' button is highlighted in blue. The left sidebar shows 'Studio' with 'Introduction', 'Database', 'Licensing', and 'Summary' options.</p>	Product	Model	<input type="radio"/> Citrix XenDesktop VDI	User/Device	<input type="radio"/> Citrix XenDesktop VDI	Concurrent	<input checked="" type="radio"/> Citrix XenDesktop Platinum	Concurrent	<input type="radio"/> Citrix XenDesktop Platinum	User/Device	<input type="radio"/> Citrix XenDesktop Enterprise	Concurrent	<input type="radio"/> Citrix XenDesktop Enterprise	User/Device
Product	Model														
<input type="radio"/> Citrix XenDesktop VDI	User/Device														
<input type="radio"/> Citrix XenDesktop VDI	Concurrent														
<input checked="" type="radio"/> Citrix XenDesktop Platinum	Concurrent														
<input type="radio"/> Citrix XenDesktop Platinum	User/Device														
<input type="radio"/> Citrix XenDesktop Enterprise	Concurrent														
<input type="radio"/> Citrix XenDesktop Enterprise	User/Device														

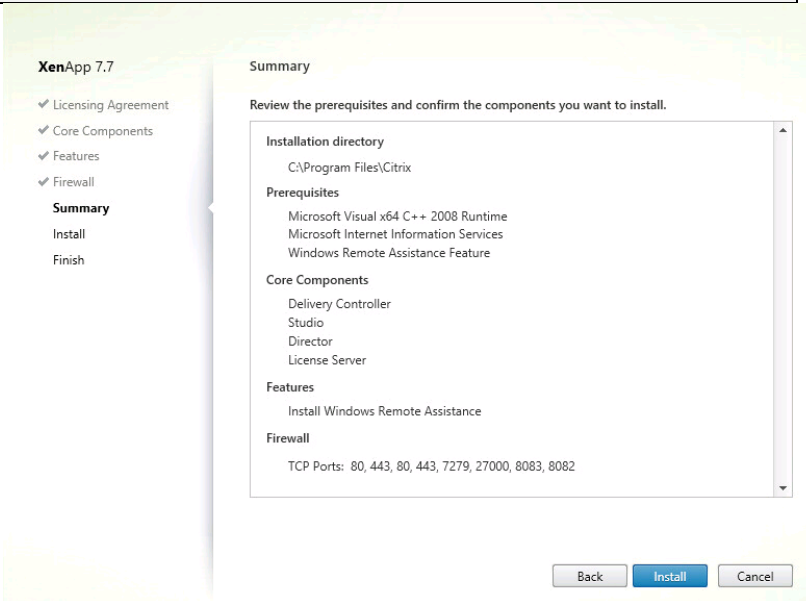
Instructions	Visual
5. Click Finish to complete initial setup.	

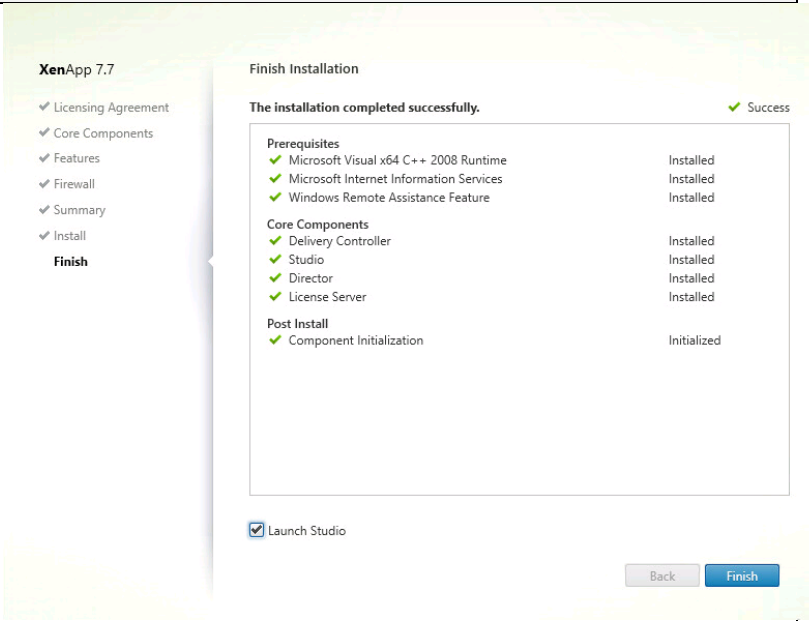
Additional XenDesktop Controller Configuration

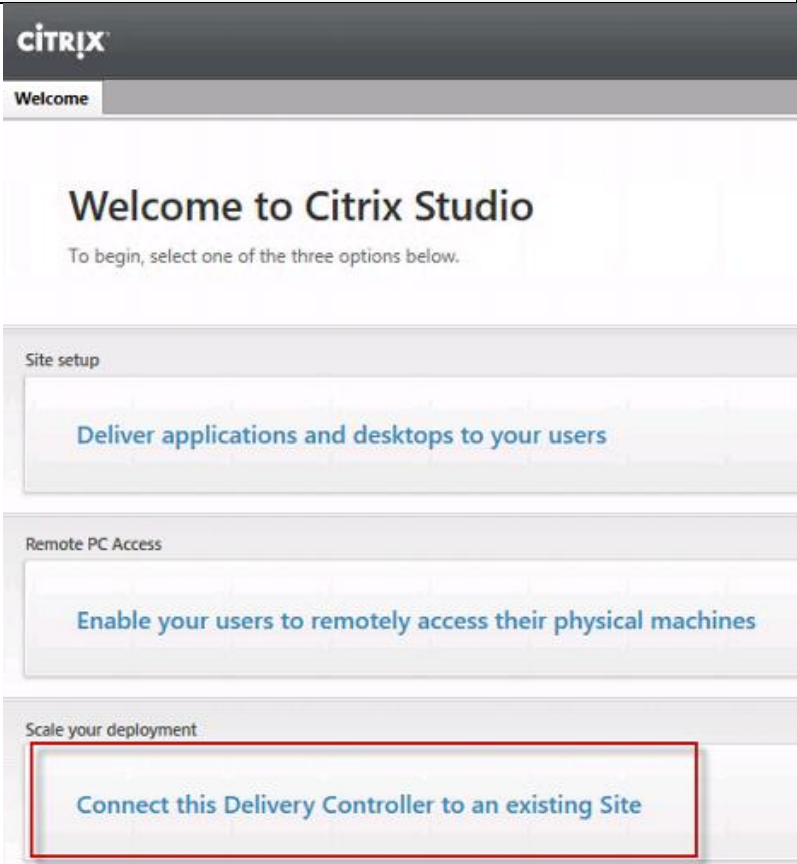
After the first controller is completely configured and the Site is operational, you can add additional controllers. In this CVD, we created two Delivery Controllers.

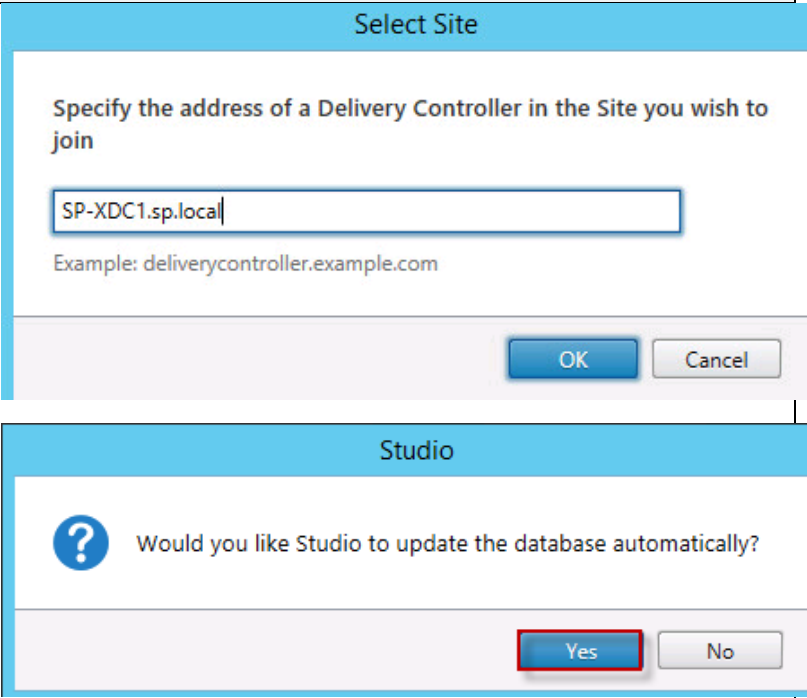
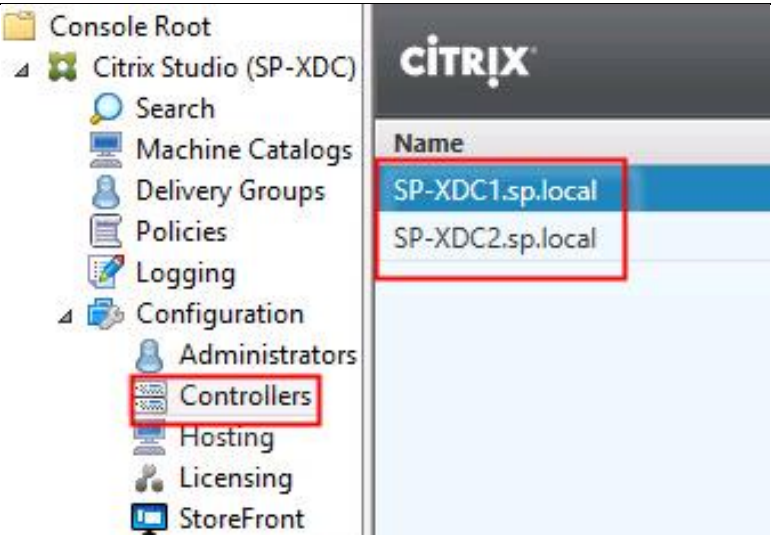
Instructions	Visual
<p>1. To begin the installation of the second Delivery Controller, connect to the second XenApp server and launch the installer from the Citrix XenApp 7.7 ISO.</p> <p>2. Click Start</p>	

Instructions	Visual
<p>3. Select the components to be installed:</p> <p>Delivery Controller</p> <p>Studio</p> <p>Director</p> <p>StoreFront</p> <p>4. Click Next</p>	

Instructions	Visual
<p>1. Repeat the same steps used to install the first Delivery Controller.</p> <p>2. Review the Summary configuration.</p> <p>3. Click Install</p>	

Instructions	Visual
<p>4. Confirm all selected components were successfully installed.</p> <p>5. Verify the Launch Studio checkbox is checked.</p> <p>6. Click Finish</p>	 <p>XenApp 7.7</p> <ul style="list-style-type: none"> ✓ Licensing Agreement ✓ Core Components ✓ Features ✓ Firewall ✓ Summary ✓ Install Finish <p>Finish Installation</p> <p>The installation completed successfully. ✓ Success</p> <p>Prerequisites</p> <ul style="list-style-type: none"> ✓ Microsoft Visual x64 C++ 2008 Runtime Installed ✓ Microsoft Internet Information Services Installed ✓ Windows Remote Assistance Feature Installed <p>Core Components</p> <ul style="list-style-type: none"> ✓ Delivery Controller Installed ✓ Studio Installed ✓ Director Installed ✓ License Server Installed <p>Post Install</p> <ul style="list-style-type: none"> ✓ Component Initialization Initialized <p><input checked="" type="checkbox"/> Launch Studio</p> <p>Back Finish</p>

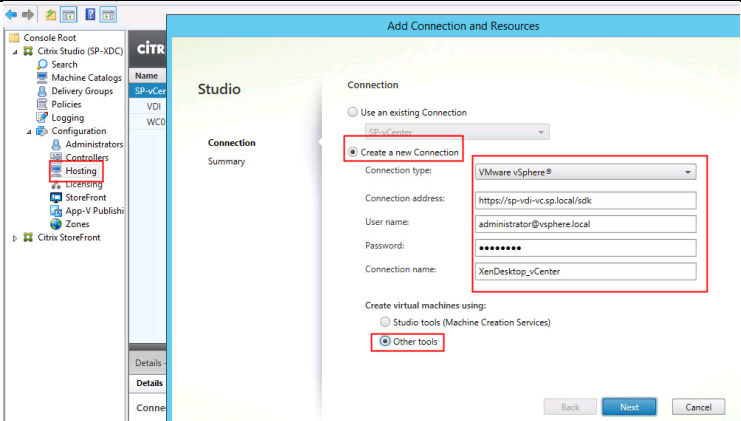
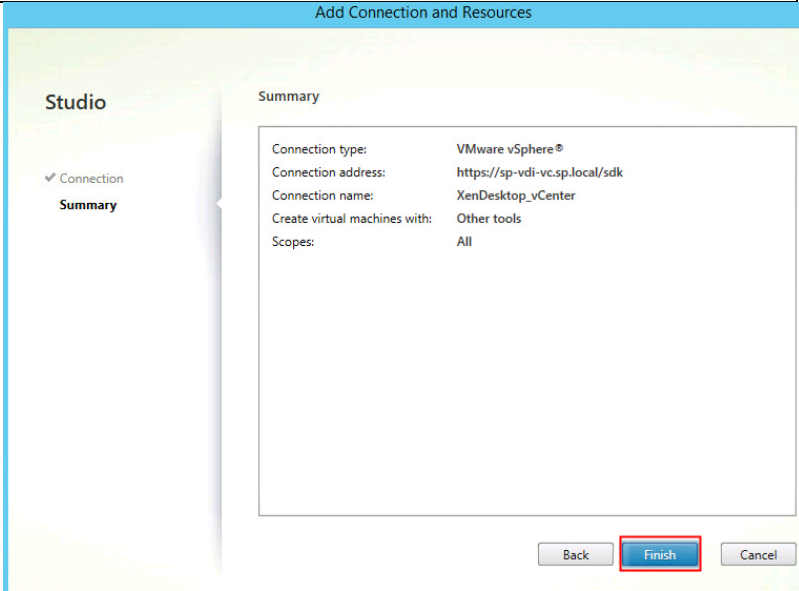
Instructions	Visual
<p>1. Click the Connect this Delivery Controller to an existing Site button.</p>	 <p>CITRIX</p> <p>Welcome</p> <p>Welcome to Citrix Studio</p> <p>To begin, select one of the three options below.</p> <p>Site setup</p> <p>Deliver applications and desktops to your users</p> <p>Remote PC Access</p> <p>Enable your users to remotely access their physical machines</p> <p>Scale your deployment</p> <p>Connect this Delivery Controller to an existing Site</p>

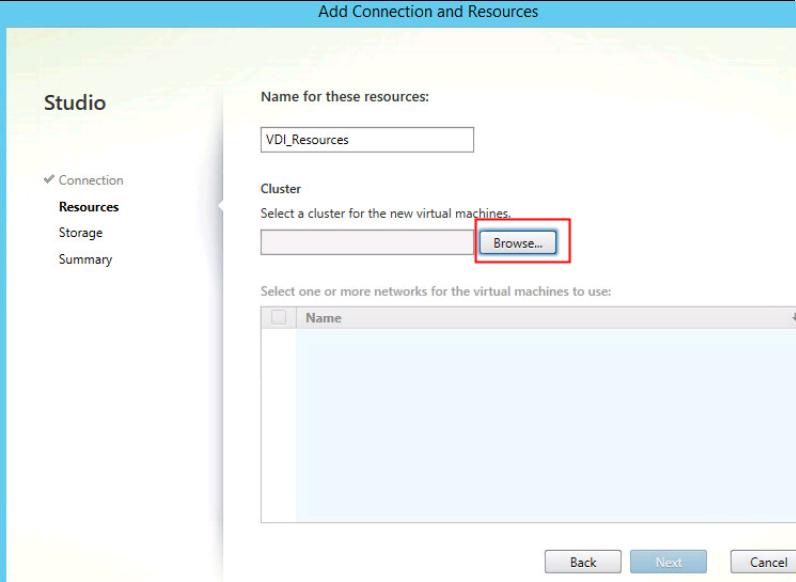
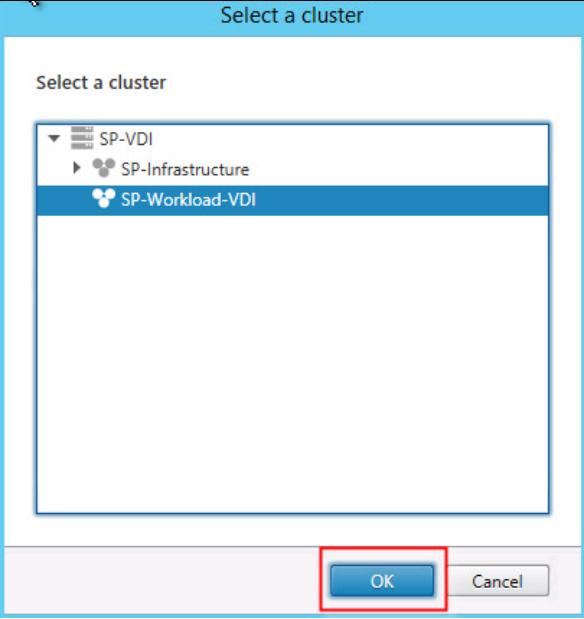
Instructions	Visual
<p>2. Enter the FQDN of the first delivery controller.</p> <p>3. Click OK</p> <p>4. Click Yes to allow the database to be updated with this controller's information automatically.</p>	
<p>5. When complete, verify the Delivery Controller has been added to the list of Controllers.</p>	

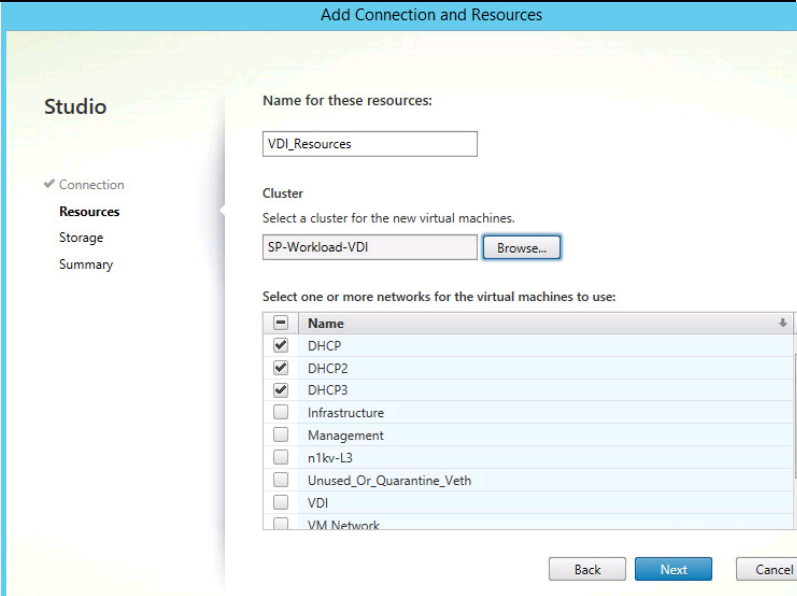
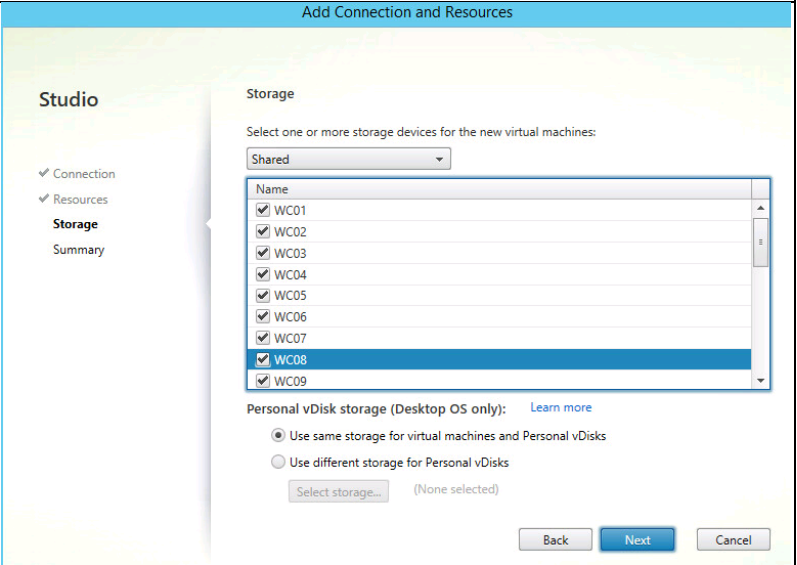
Configure the XenDesktop Site Hosts and Storage

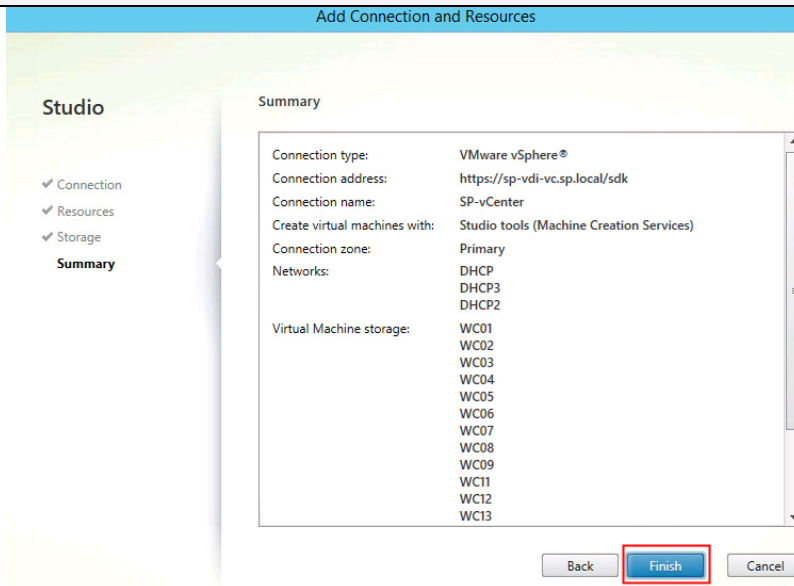
We added the VMware Infrastructure to the Delivery Controller to be able to provision machines.

Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Click Hosting and click New Connection.</div> <div>2. Click Next</div>	<div></div>
<div>3. Click Finish to create connection.</div>	<div></div>
Instructions	Visual

Instructions	Visual
<p>4. Under the newly created connection, begin to add resources. Name the resources 'VDI_Resources'</p> <p>5. Click Browse to select a VSphere cluster.</p>	
<p>6. Select the Cluster the VMs will reside on.</p>	
Instructions	Visual

Instructions	Visual
7. Select the Network Resources and click Next.	
8. Select the Storage datastores the VMs will reside on. 9. Click Next.	

Instructions	Visual
10. Click Finish.	

Configure XenDesktop HDX Policies

The following Citrix Policies were implemented for our testing:

1. For The Windows 7 VDI Delivery Group the following Policies for user profile management were used.

Policies

Policies

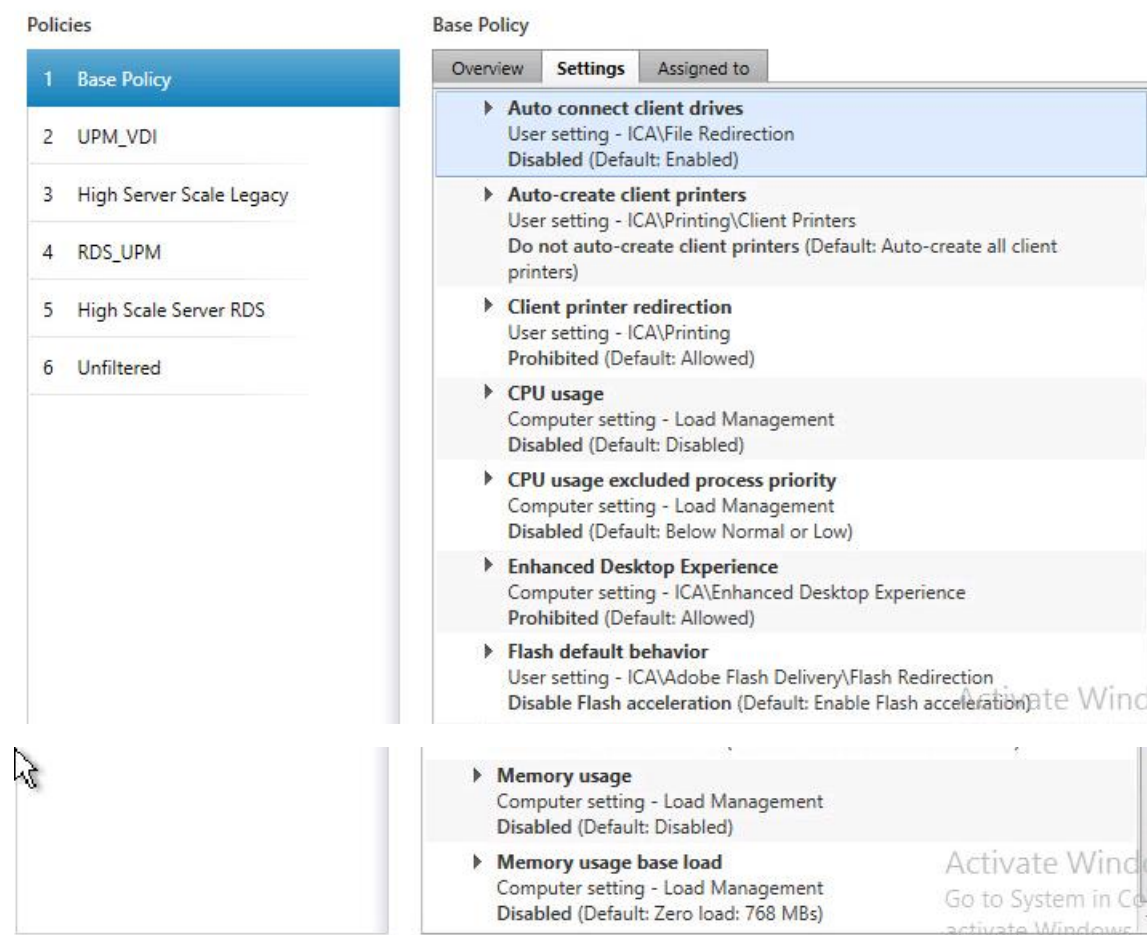
- 1 Base Policy
- 2 UPM_VDI
- 3 High Server Scale Legacy
- 4 RDS_UPM
- 5 High Scale Server RDS
- 6 Unfiltered

UPM_VDI

Overview Settings Assigned to


- ▶ **Active write back**
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)
- ▶ **Delete locally cached profiles on logoff**
Computer setting - Profile Management\Profile handling
Enabled (Default: Disabled)
- ▶ **Enable Profile management**
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)
- ▶ **Exclusion list - directories**
Computer setting - Profile Management\File system\Exclusions
AppData\Local;AppData\LocalLow;AppData\Roaming;\$Recycle.Bin
(Default:)
- ▶ **Path to user store**
Computer setting - Profile Management\Basic settings
\\sp-file01\UserProfiles\#SAMAccountName# (Default: Windows)
- ▶ **Process logons of local administrators**
Computer setting - Profile Management\Basic settings
Enabled (Default: Disabled)

2. The Following Base Policy was applied to all objects in the XenDesktop Site.

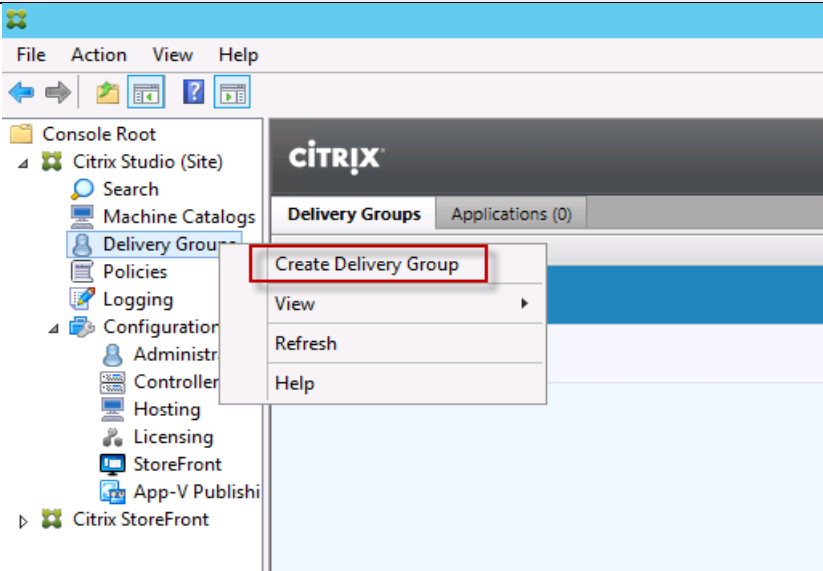
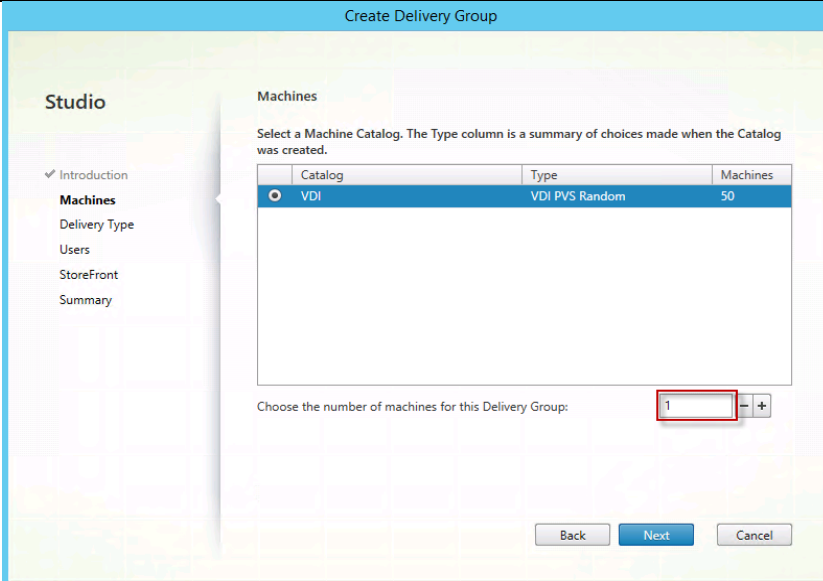


Creating Delivery Groups

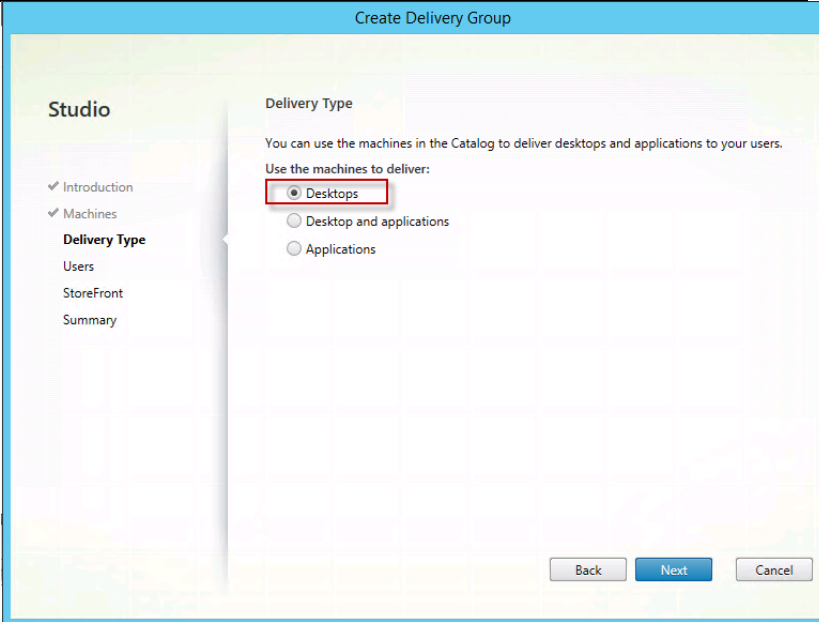
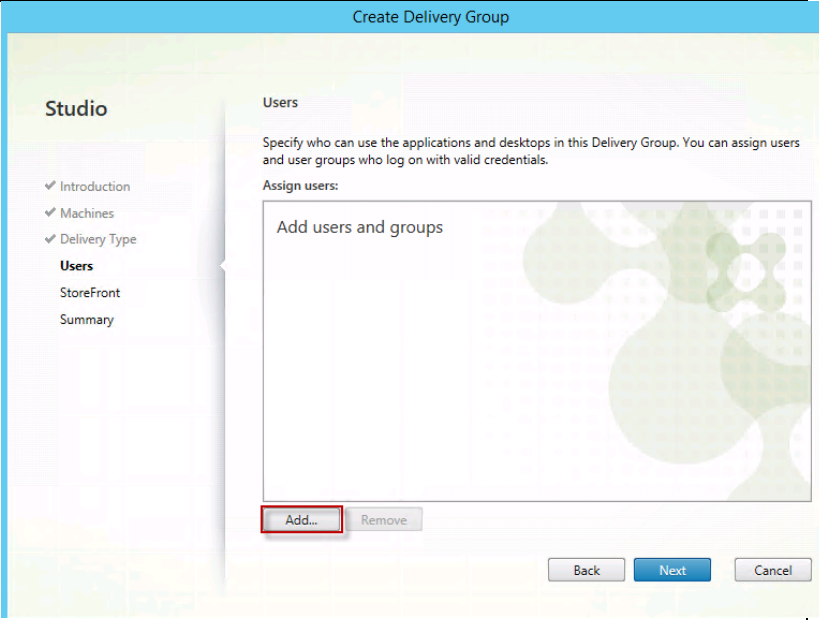
Delivery Groups are collections of machines that control access to desktops and applications. With Delivery Groups, you can specify which users and groups can access which desktops and applications.

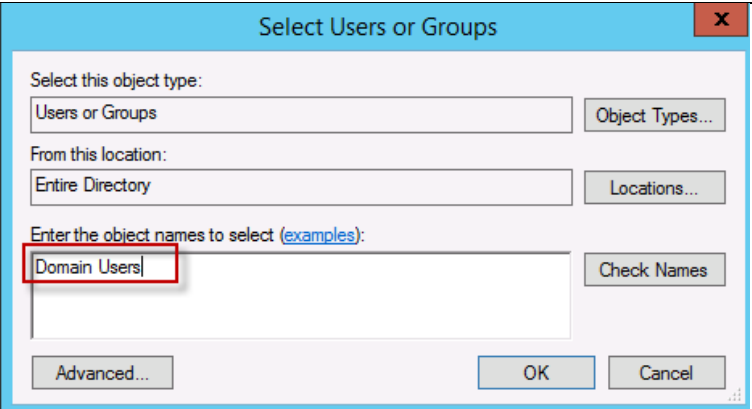
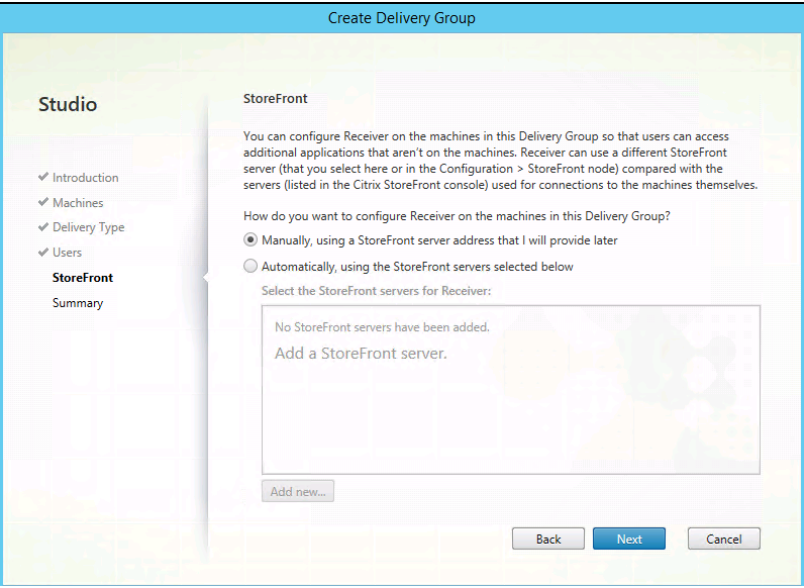
 The instructions below outline the procedure to create a Delivery Group for VDI desktops. When you have completed these steps, repeat the procedure to a Delivery Group for RDS desktops.

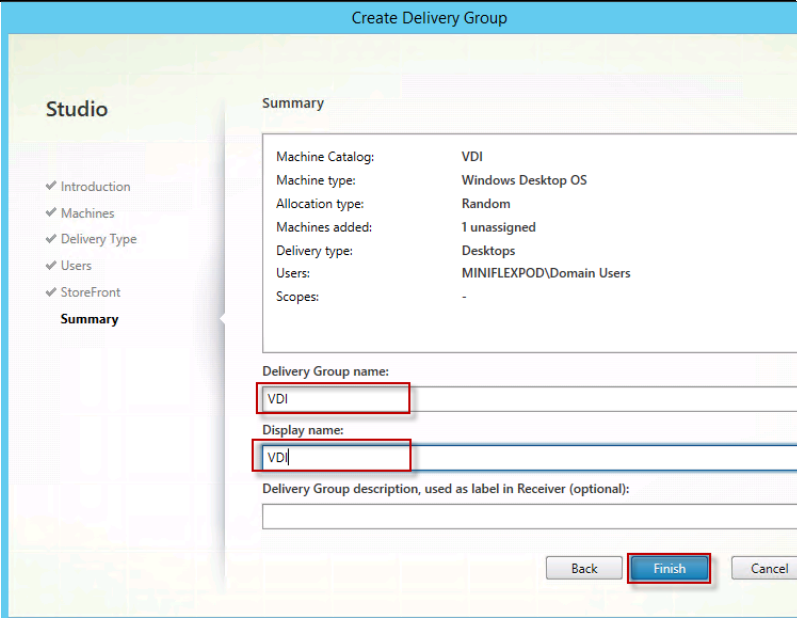
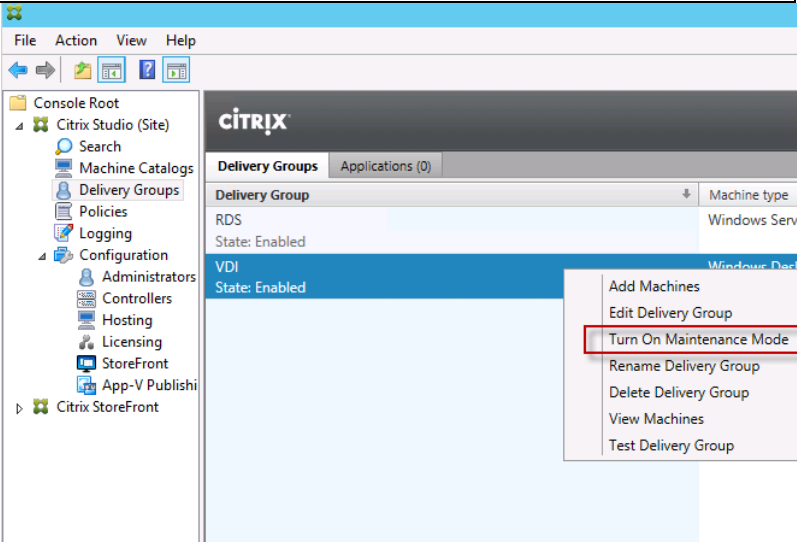
Instructions	Visual
--------------	--------

Instructions	Visual						
<div>1. Connect to a XenDesktop server and launch Citrix Studio.</div> <div>2. Choose Create Delivery Group from the drop-down menu.</div>	 <p>The screenshot shows the Citrix Studio console interface. On the left, the 'Console Root' tree is expanded to 'Citrix Studio (Site)', then 'Machine Catalogs', and finally 'Delivery Groups'. A right-click context menu is open over 'Delivery Groups', with 'Create Delivery Group' highlighted by a red rectangle. The main pane on the right shows the 'Delivery Groups' tab with 'Applications (0)' listed.</p>						
<div>3. Specify the Machine Catalog and increment the number of machines to add.</div> <div>4. Click Next</div>	 <p>The screenshot shows the 'Create Delivery Group' wizard in Citrix Studio. The 'Machines' step is active, displaying a table to select machine catalogs. A red rectangle highlights the '1' in the spinner control for 'Choose the number of machines for this Delivery Group:'. The 'Next' button is highlighted in blue.</p> <table><tr><th>Catalog</th><th>Type</th><th>Machines</th></tr><tr><td>VDI</td><td>VDI PVS Random</td><td>50</td></tr></table>	Catalog	Type	Machines	VDI	VDI PVS Random	50
Catalog	Type	Machines					
VDI	VDI PVS Random	50					

Instructions	Visual
--------------	--------

Instructions	Visual
<div>1. Specify what the machines in the catalog will deliver: Desktops, Desktops and Applications, or Applications.</div> <div>2. Select Desktops.</div> <div>3. Click Next</div>	 <p>The screenshot shows the 'Create Delivery Group' wizard in the 'Delivery Type' step. On the left, the 'Studio' sidebar has 'Introduction', 'Machines', 'Delivery Type', 'Users', 'StoreFront', and 'Summary' items. 'Delivery Type' is selected. The main area has the heading 'Delivery Type' and a sub-heading 'Use the machines to deliver:'. There are three radio buttons: 'Desktops' (which is selected and highlighted with a red box), 'Desktop and applications', and 'Applications'. At the bottom right are 'Back', 'Next', and 'Cancel' buttons.</p>
<div>To make the Delivery Group available, you must add users.</div> <div>4. Click Add...</div>	 <p>The screenshot shows the 'Create Delivery Group' wizard in the 'Users' step. The 'Studio' sidebar on the left now has 'Users' selected. The main area has the heading 'Users' and a sub-heading 'Assign users:'. Below this is a box titled 'Add users and groups'. At the bottom of this box are 'Add...' (highlighted with a red box) and 'Remove' buttons. At the bottom right of the wizard are 'Back', 'Next', and 'Cancel' buttons.</p>
Instructions	Visual

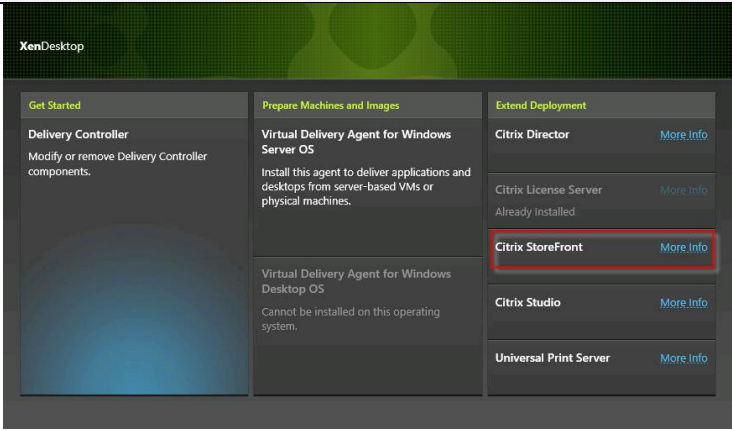
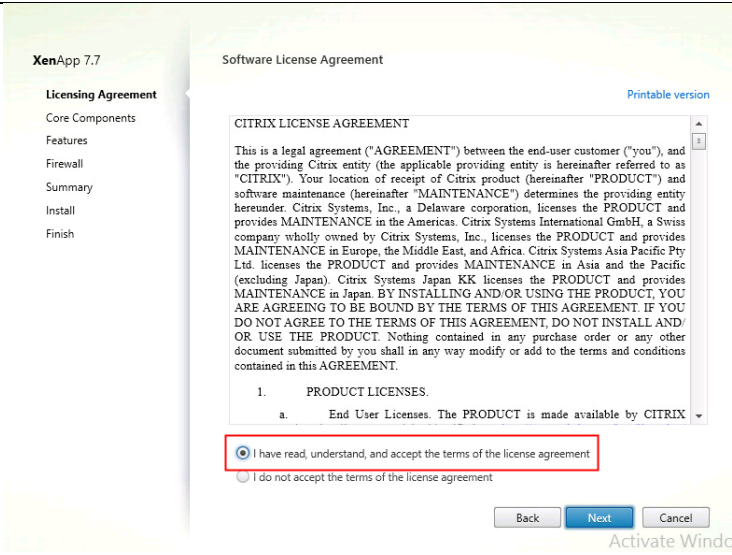
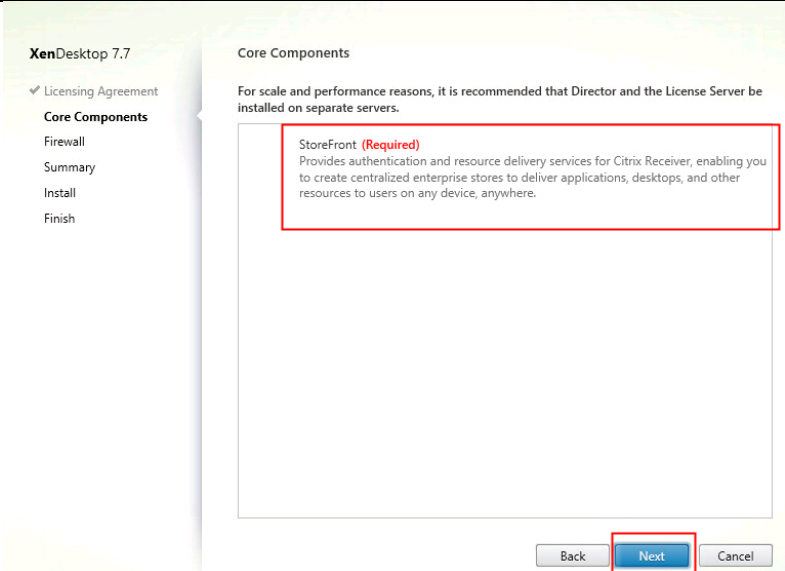
<div>1. In the Select Users or Groups dialog, add users or groups.</div> <div>2. Click OK. When users have been added, click Next on the Assign dialog (shown above).</div>	<div></div>
<div>3. Enter the StoreFront configuration for how Receiver will be installed on the machines in this Delivery Group. Click “Manually, using a StoreFront server address that I will provide later.”</div> <div>4. Click Next</div>	<div></div>
<div>Instructions</div>	<div>Visual</div>

Instructions	Visual
<div>1. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Display name (for example, VDI or RDS).</div> <div>2. Click Finish</div>	
<div>4. Citrix Studio lists the created Delivery Groups and the type, number of machines created, sessions, and applications for each group in the Delivery Groups tab.</div> <div>5. On the pull-down menu, select “Turn on Maintenance Mode.”</div>	

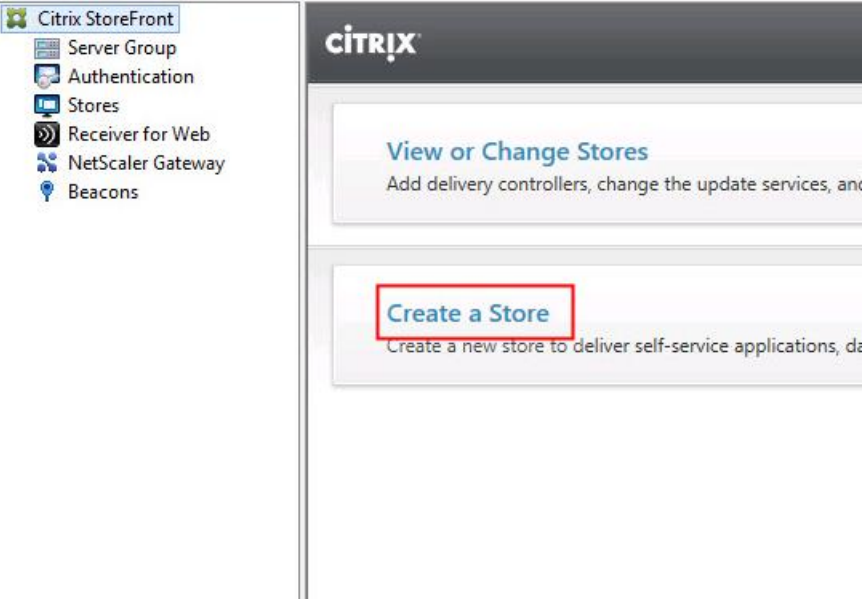
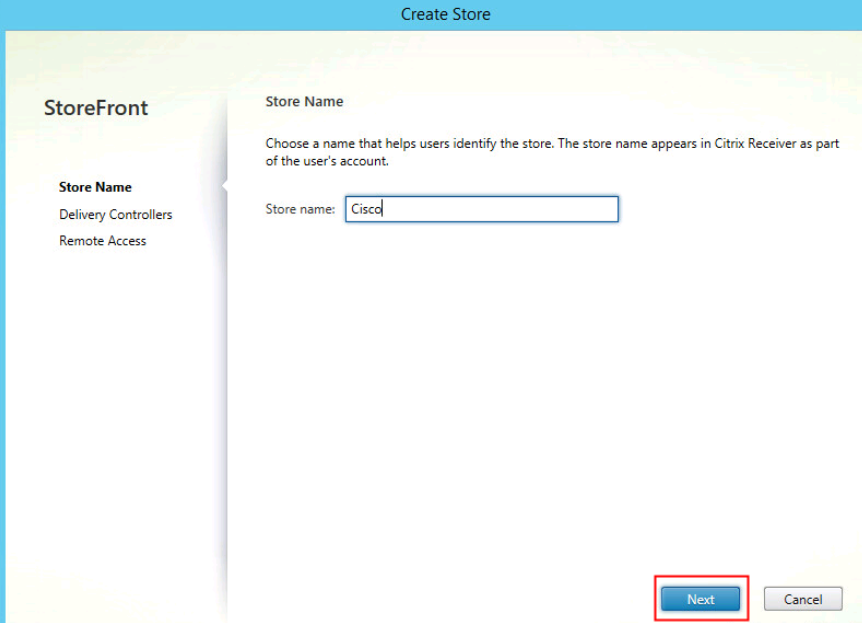
Installing and Configuring Citrix Storefront

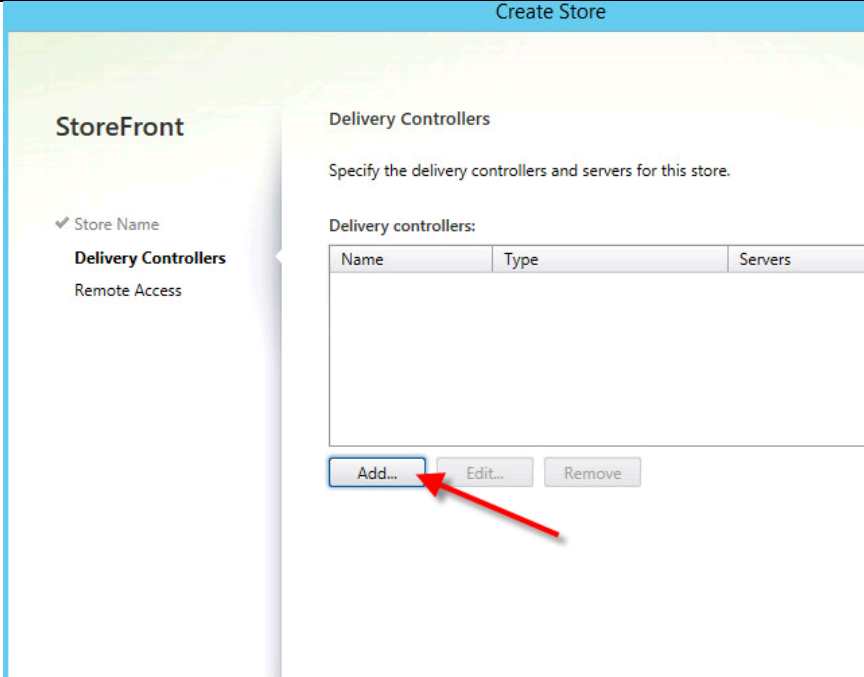
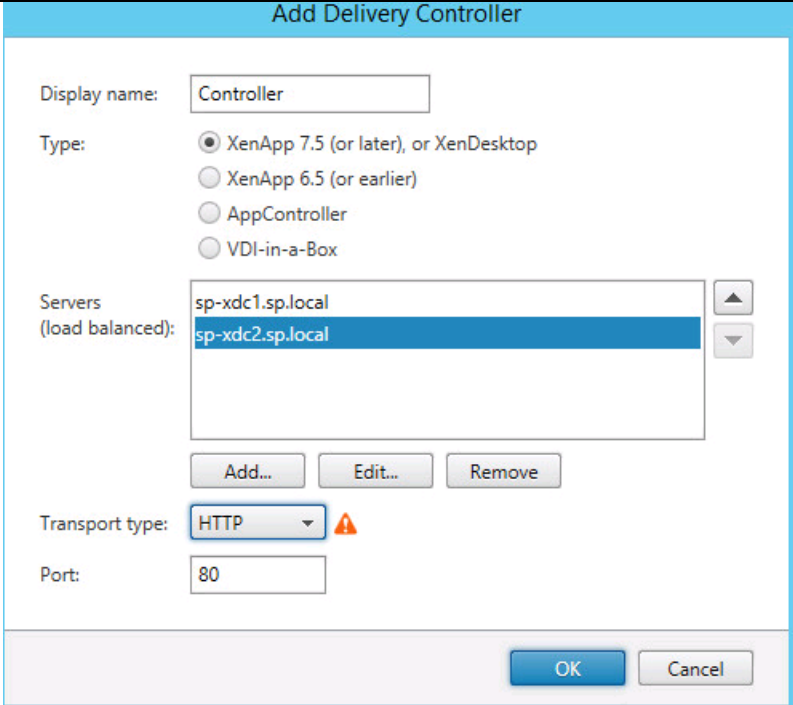
In this study we used a dedicated server for StoreFront functionality. To configure the StoreFront Site, complete the following steps:

Instructions	Visual
--------------	--------

Instructions	Visual
<p>1. Copy the license files to the default location (C:\Program Files (x86)\Citrix\Licensing\MyFiles) on the license server</p> <p>2. Restart the server or licensing services so that the licenses are activated.</p>	 <p>The screenshot shows the XenDesktop installation wizard. The 'Extend Deployment' tab is selected, showing a list of components. 'Citrix StoreFront' is highlighted with a red box, indicating it is the next step in the process. Other components like 'Citrix Director', 'Citrix License Server', 'Citrix Studio', and 'Universal Print Server' are also listed with 'More Info' links.</p>
<p>3. Run the application Citrix License Administration Console</p>	 <p>The screenshot shows the Citrix License Administration Console. The 'Licensing Agreement' tab is selected. The 'I have read, understand, and accept the terms of the license agreement' radio button is selected and highlighted with a red box. The 'Next' button is visible at the bottom right.</p>
<p>4. Confirm that the license files have been read and enabled correctly.</p>	 <p>The screenshot shows the XenDesktop Core Components configuration screen. The 'StoreFront (Required)' section is highlighted with a red box, indicating it is a required component. The 'Next' button is highlighted with a red box at the bottom right.</p>

Instructions	Visual
5. Click Next	<div><div><div>XenDesktop 7.7</div><div><div>✓ Licensing Agreement</div><div>✓ Core Components</div><div>Firewall</div><div>Summary</div><div>Install</div><div>Finish</div></div></div><div><div>Firewall</div><div>The default ports are listed below.Printable version</div><div><div>StoreFront</div><div>80, 443 TCP</div></div><div><div>Configure firewall rules:</div><div><div><input checked="" type="radio"/> Automatically</div><div>Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.</div></div><div><div><input type="radio"/> Manually</div><div>Select this option if you are not using Windows Firewall or if you want to create the rules yourself.</div></div></div><div><div>Back</div><div>Next</div><div>Cancel</div></div></div></div>
6. Click Install	<div><div><div>XenDesktop 7.7</div><div><div>✓ Licensing Agreement</div><div>✓ Core Components</div><div>✓ Firewall</div><div>Summary</div><div>Install</div><div>Finish</div></div></div><div><div>Summary</div><div>Review the prerequisites and confirm the components you want to install.</div><div><div>Installation directory</div><div>C:\Program Files\Citrix</div><div>Prerequisites</div><div>Microsoft Internet Information Services</div><div>Core Components</div><div>StoreFront</div><div>Firewall</div><div>TCP Ports: 80, 443</div></div><div><div>Back</div><div>Install</div><div>Cancel</div></div></div></div>

Instructions	Visual
7. When install completes, open the StoreFront Console and select Create a Store	 <p>The screenshot shows the Citrix StoreFront console interface. On the left is a navigation pane with icons and labels for 'Citrix StoreFront', 'Server Group', 'Authentication', 'Stores', 'Receiver for Web', 'NetScaler Gateway', and 'Beacons'. The main area on the right has a dark header with the 'CITRIX' logo. Below the header, there are two prominent buttons: 'View or Change Stores' and 'Create a Store'. The 'Create a Store' button is highlighted with a red rectangular box. Text below the buttons describes the actions for each.</p>
8. Name your Store	 <p>The screenshot shows the 'Create Store' wizard in the Citrix StoreFront console. The title bar at the top says 'Create Store'. The main content area is divided into two sections. On the left, under the 'StoreFront' heading, are links for 'Store Name', 'Delivery Controllers', and 'Remote Access'. The 'Store Name' section is active, showing instructions: 'Choose a name that helps users identify the store. The store name appears in Citrix Receiver as part of the user's account.' Below this is a text input field labeled 'Store name:' containing the text 'Cisco'. At the bottom right of the wizard, there are two buttons: 'Next' and 'Cancel'. The 'Next' button is highlighted with a red rectangular box.</p>

Instructions	Visual
<p>9. Click Add to add your deliver controllers.</p>	
<p>10. Add your XenDesktop Delivery Controllers using the FQDN.</p> <p>In this study we used HTTP as a Transport Type.</p>	

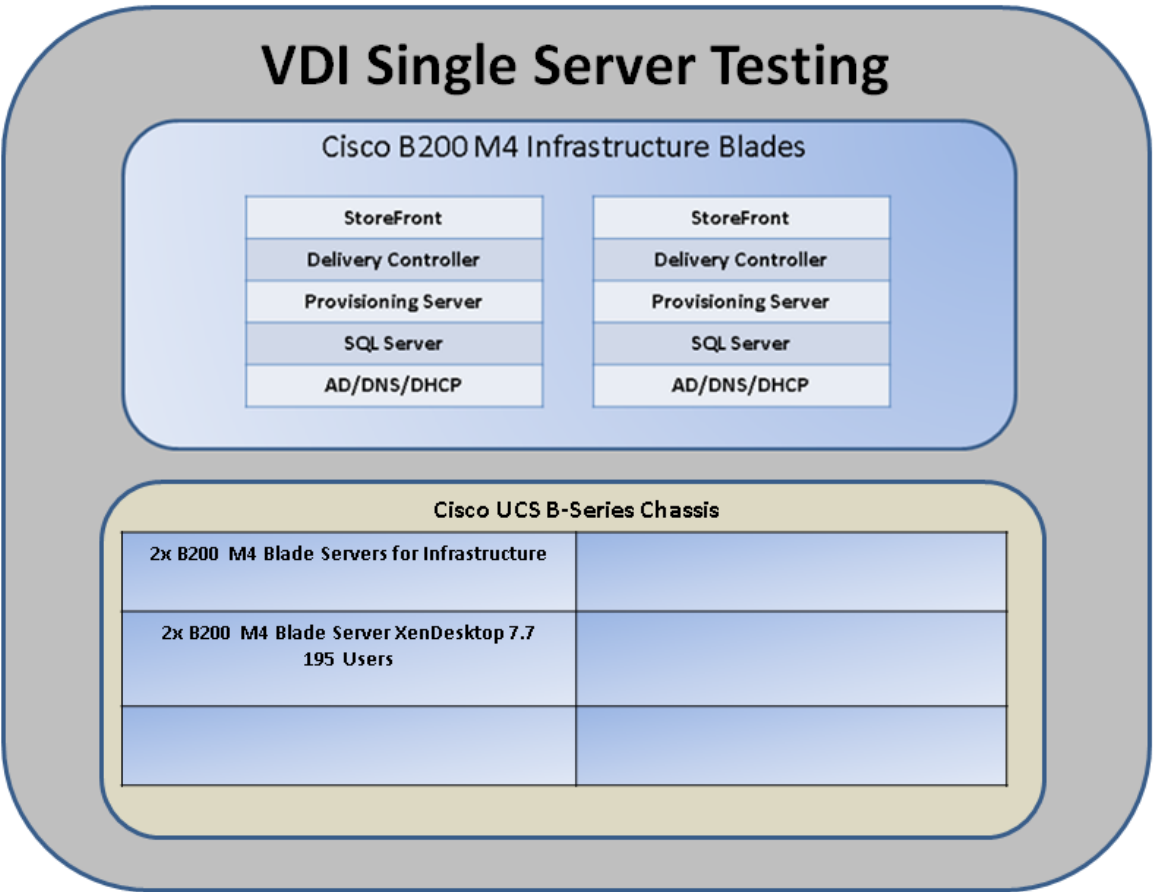
Instructions	Visual
<div>11. Select None for Remote Access Type.</div> <div>12. Click Create.</div> <div><p>*We recommend a NetScaler Appliance to Load Balance two individual StoreFront Servers as a Best Practice.</p></div>	<div><div><div>StoreFront</div><div><div>✓ Store Name</div><div>✓ Delivery Controllers</div><div>Remote Access</div></div></div><div><div>Remote Access</div><div>Add NetScaler Gateway appliances to provide user access from external networks.</div><div>Remote access:<div><div><input checked="" type="radio"/> None</div><div><input type="radio"/> No VPN tunnel ⓘ</div><div><input type="radio"/> Full VPN tunnel ⓘ</div></div></div><div>NetScaler Gateway appliances:<div></div><div>Add...</div></div><div>Default appliance:<div></div></div><div><div>Back</div><div>Create</div></div></div></div>

Test Setup and Configurations

In this project, we tested a single Cisco UCS B200 M4 blade in a single chassis and 14 Cisco UCS B200 M4 blades in two chassis to illustrate linear scalability for each workload studied.

Cisco UCS Test Configuration for Single Blade Scalability

Figure 20 Cisco UCS B200 M4 Server for Single Server Scalability XenApp 7. 7 VDI with PVS 7.7



Hardware components:

- Cisco UCS 5108 B-Series Server Chassis
- 2 Cisco UCS 6248UP Fabric Interconnects
- 1 Cisco B200-M4 B-Series Blade Server XenDesktop VDI workload.
- 2 Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v3 CPUs at 2.6 GHz, with 256 GB of memory per blade server [16 GB x 15 DIMMs at 2133 MHz]) for all Infrastructure blades
- Cisco vNIC CNA (1 per blade)

Test Setup and Configurations

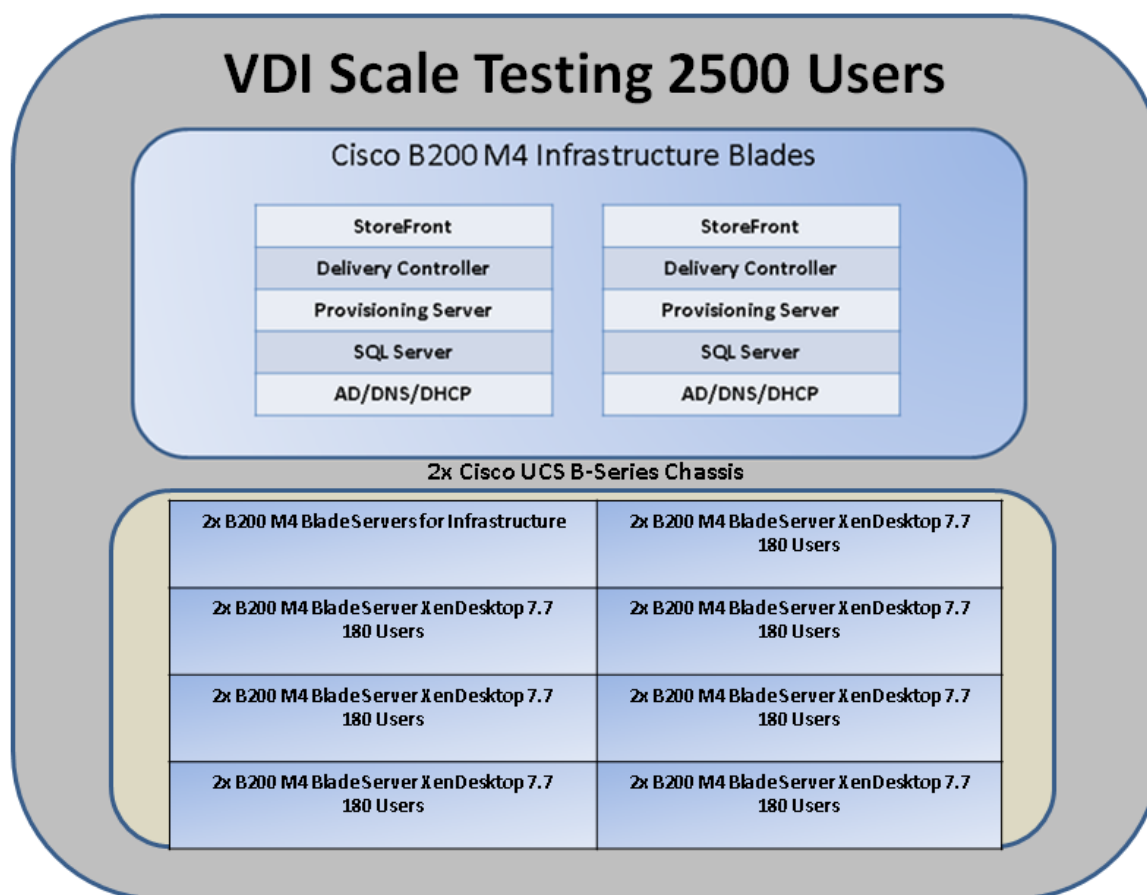
- 2 Cisco Nexus 9300 Access Switches
- Nimble CS700 Adaptive Array

Software components:

- Cisco UCS firmware 3.1.1e
- Citrix XenDesktop 7.7 VDI Hosted Virtual Desktops
- Citrix Provisioning Server 7.7
- Citrix User Profile Manager
- Microsoft SQL Server 2012
- Microsoft Windows 7, 2vCPU, 1.7GB RAM, 40 GB vdisk
- Microsoft Office 2010
- Login VSI 4.1.4

Cisco UCS Test Configuration for Blade Server VDI Scalability

Figure 21 14x Cisco UCS B200 M4 Server for Blade Server VDI Scalability XenDesktop 7.7 VDI with PVS 7.7



Hardware components:

- Cisco UCS 5108 B-Series Server Chassis
- 2 Cisco UCS 6248UP Fabric Interconnects
- 14x Cisco B200-M4 M-Series Blade Servers for XenDesktop VDI workload.
- In this study we tested the environment to maximize the resources. In an N+1 scenario the VDI scale workload should be a total of 2300 VMs to accommodate a single blade failure.
- 2x Cisco UCS B200 M4 Blade Servers (2 Intel Xeon processor E5-2660 v3 CPUs at 2.6 GHz, with 256 GB of memory per blade server [16 GB x 15 DIMMs at 2133 MHz]) for all Infrastructure blades
- Cisco vNIC CNA (1 per blade)
- 2 Cisco Nexus 9300 Access Switches
- Nimble CS700 Adaptive Array

Software components:

- Cisco UCS firmware 3.1.1e
- Citrix XenDesktop 7.7 VDI Hosted Virtual Desktops
- Citrix Provisioning Server 7.7
- Citrix User Profile Manager
- Microsoft SQL Server 2012
- Microsoft Windows 7, 2vCPU, 2GB RAM, 40 GB vdisk
- Microsoft Office 2010
- Login VSI 4.1.4

Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco labs in San Jose, California.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the Citrix XenApp RDS Hosted Shared models under test.

Test metrics were gathered from the virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

You can obtain additional information and a free test license from <http://www.loginvsi.com>.

Testing Procedure

The following protocol was used for each test cycle in this study to insure consistent results.

Pre-Test Setup for Single and Multi-Blade Testing

All machines were shut down utilizing the XenApp 7.7 Administrator.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the **required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.**

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. Additionally, we require all sessions started, whether 60 single server users or 900 full scale test users to become active within two minutes after the last session is launched.

In addition, Cisco requires that the Login VSI Benchmark method is used for all single server and scale testing. This assures that our tests represent real-world scenarios. For each of the three consecutive runs on single server tests, the same process was followed. Complete the following steps:

1. Time 0:00:00 Start PerfMon Logging on the following systems:
 - Infrastructure and VDI Host Blades used in test run
 - All Infrastructure VMs used in test run (AD, SQL, brokers, image mgmt., etc.)
2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
3. Time 0:05: Boot RDS Machines using XenDesktop Studio UCSM KVM.
4. Time 0:06 First machines boot.
5. Time 0:35 Single Server or Scale target number of RDS Servers registered on XD.



No more than 60 Minutes of rest time is allowed after the last desktop is registered on the XD Studio or available in View Connection Server dashboard. Typically a 30 minute rest period for Windows 7 desktops and 15 minutes for RDS VMs is sufficient.

6. Time 1:35 Start Login VSI 4.1.4 Office Worker Benchmark Mode Test, setting auto-logoff time at 900 seconds, with Single Server or Scale target number of desktop VMs utilizing sufficient number of Launchers (at 20-25 sessions/Launcher).
7. Time 2:23 Single Server or Scale target number of desktop VMs desktops launched (48 minute benchmark launch rate).
8. Time 2:25 All launched sessions must become active.



All sessions launched must become active for a valid test run within this window.

9. Time 2:40 Login VSI Test Ends (based on Auto Logoff 900 Second period designated above).
10. Time 2:55 All active sessions logged off.



All sessions launched and active must be logged off for a valid test run. The XD Studio or View Connection Dashboard must show that all desktops have been returned to the registered/available state as evidence of this condition being met.

11. Time 2:57 All logging terminated; Test complete.
12. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows 7 machines.
13. Time 3:30 Reboot all hypervisors.
14. Time 3:45 Ready for new test sequence.

Success Criteria

Our “pass” criteria for this testing follows is Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, memory utilization, storage utilization, and network utilization. We will use Login VSI to launch version 4.1 Office Worker workloads. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The Citrix Desktop Studio or VMware Horizon with View Connection Server Dashboard will be monitored throughout the steady state to make sure of the following:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered, unavailable or available state at any time during steady state

Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. Stuck sessions define a test failure condition.

Cisco requires three consecutive runs with results within +/- 1% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/- 1% variability are accepted. (All test data from partner run testing must be supplied along with proposed white paper.)

We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing. Cisco UCS M-Series Modular Servers with XenApp 7.6 Test Results

The purpose of this testing is to provide the data needed to validate Citrix XenApp 7.7 Hosted Shared Desktop with Citrix Provisioning Services 7.7 using Microsoft Windows Server 2012 R2 sessions on Cisco UCS M-Series M142-M4 Modular Compute Cartridges on Local Storage.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of Citrix products.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish single blade performance and multi-blade, linear scalability.

VSImax 4.1.x Description

The philosophy behind Login VSI is different to conventional benchmarks. In general, most system benchmarks are steady state benchmarks. These benchmarks execute one or multiple processes, and the measured execution time is the outcome of the test. Simply put: the faster the execution time or the bigger the throughput, the faster the system is according to the benchmark.

Login VSI is different in approach. Login VSI is not primarily designed to be a steady state benchmark (however, if needed, Login VSI can act like one). Login VSI was designed to perform benchmarks for SBC or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer and Adobe PDF reader. By gradually increasing the amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what its true maximum user capacity is.

After a test is performed, the response times can be analyzed to calculate the maximum active session/desktop capacity. Within Login VSI this is calculated as VSImax. When the system is coming closer to its saturation point, response times will rise. When reviewing the average response time it will be clear the response times escalate at saturation point.

This VSImax is the “Virtual Session Index (VSI)”. With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

Server-Side Response Time Measurements

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts on every target system, and are initiated **at logon within the simulated user’s desktop session context.**

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI the choice has been made to execute the scripts completely server side. This is the only practical and platform independent solution, for a benchmark like Login VSI.

Calculating VSImax v4.1.x

The simulated desktop workload is scripted in a 48 minute loop when a simulated Login VSI user is logged on, performing generic Office worker activities. After the loop is finished it will restart automatically. Within each loop the response times of sixteen specific operations are measured in a regular interval: sixteen times in within each loop. The response times of these five operations are used to determine VSImax.

The five operations from which the response times are measured are:

- Notepad File Open (NFO)

Loading and initiating VSINotepad.exe and opening the openfile dialog. This operation is handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

- Notepad Start Load (NSLD)

Loading and initiating VSINotepad.exe and opening a file. This operation is also handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

- Zip High Compression (ZHC)

This action copy's a random file and compresses it (with 7zip) with high compression enabled. The compression will very briefly spike CPU and disk IO.

- Zip Low Compression (ZLC)

This action copy's a random file and compresses it (with 7zip) with low compression enabled. The compression will very briefly disk IO and creates some load on the CPU.

- CPU

Calculates a large array of random data and spikes the CPU for a short period of time.

These measured operations within Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, etc. These operations are specifically short by nature. When such operations become consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. If such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

Figure 22 Sample of a VSI Max Response Time Graph, Representing a Normal Test





Figure 23 Sample of a VSI Test Response Time Graph with a Clear Performance Issue



When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

The response times are very different per measurement type, for instance Zip with compression can be around 2800 ms, while the Zip action without compression can only take 75ms. These response time of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

In comparison to previous VSImax models, this weighting much better represent system performance. All actions have very similar weight in the VSImax total. The following weighting of the response times are applied.

The following actions are part of the VSImax v4.1 calculation and are weighted as follows (US notation):

- Notepad File Open (NFO): 0.75
- Notepad Start Load (NSLD): 0.2
- Zip High Compression (ZHC): 0.125
- Zip Low Compression (ZLC): 0.2
- CPU: 0.75

This weighting is applied on the baseline and normal Login VSI response times.

With the introduction of Login VSI 4.1 we also created a new method to calculate the basephase of an environment. With the new workloads (Taskworker, Powerworker, etc.) enabling 'basephase' for a more reliable baseline has become obsolete. The calculation is explained below. In total 15 lowest VSI response time samples are taken from the entire test, the lowest 2 samples are removed and the 13 remaining samples are averaged. The result is the Baseline. The calculation is as follows:

- Take the lowest 15 samples of the complete test
- From those 15 samples remove the lowest 2
- Average the 13 results that are left is the baseline

The VSImax average response time in Login VSI 4.1.x is calculated on the amount of active users that are logged on the system.

Always a 5 Login VSI response time samples are averaged + 40% of the amount of “active” sessions. For example, if the active sessions is 60, then latest 5 + 24 (=40% of 60) = 31 response time measurement are used for the average calculation.

To remove noise (accidental spikes) from the calculation, the top 5% and bottom 5% of the VSI response time samples are removed from the average calculation, with a minimum of 1 top and 1 bottom sample. As a result, with 60 active users, the last 31 VSI response time sample are taken. From those 31 samples the top 2 samples are removed and lowest 2 results are removed (5% of 31 = 1.55, rounded to 2). At 60 users the average is then calculated over the 27 remaining results.

VSImax v4.1.x is reached when the VSImax + a 1000 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSImax response time can grow 2 - 3x the baseline average. In end-user computing, a 3x increase in response time in comparison to the baseline is typically regarded as the maximum performance degradation to be considered acceptable.

In VSImax v4.1.x this latency threshold is fixed to 1000ms, this allows better and fairer comparisons between two different systems, especially when they have different baseline results. Ultimately, in VSImax v4.1.x, the performance of the system is not decided by the total average response time, but by the latency is has under load. For all systems, this is now 1000ms (weighted).

The threshold for the total response time is: average weighted baseline response time + 1000ms.

When the system has a weighted baseline response time average of 1500ms, the maximum average response time may not be greater than 2500ms (1500+1000). If the average baseline is 3000 the maximum average response time may not be greater than 4000ms (3000+1000).

When the threshold is not exceeded by the average VSI response time during the test, VSImax is not hit and the amount of sessions ran successfully. This approach is fundamentally different in comparison to previous VSImax methods, as it was always required to saturate the system beyond VSImax threshold.

Lastly, VSImax v4.1.x is now always reported with the average baseline VSI response time result. For **example: “The VSImax v4.1 was 125 with a baseline of 1526ms”**. This helps considerably in the comparison of systems and gives a more complete understanding of the system. The baseline performance helps to understand the best performance the system can give to an individual user. VSImax indicates what the total user capacity is for the system. These two are not automatically connected and related:

When a server with a very fast dual core CPU, running at 3.6 GHZ, is compared to a 10 core CPU, running at 2,26 GHZ, the dual core machine will give an individual user better performance than the 10 core machine. This is indicated by the baseline VSI response time. The lower this score is, the better performance an individual user can expect.

However, the server with the slower 10 core CPU will easily have a larger capacity than the faster dual core system. This is indicated by VSImax v4.1.x, and the higher VSImax is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSImax method is introduced: VSImax v4.1. This methodology gives much better insight in system performance and scales to extremely large systems.

Test Results

Single-Server Recommended Maximum Workload For B-Series

For Citrix XenDesktop 7.7 VDI Desktop use cases, the recommended maximum workload was determined based on both Login VSI Medium workload with flash end user experience measures and blade server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the blade can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to insure that end user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95 percent.



Memory should never be oversubscribed for Desktop Virtualization workloads.

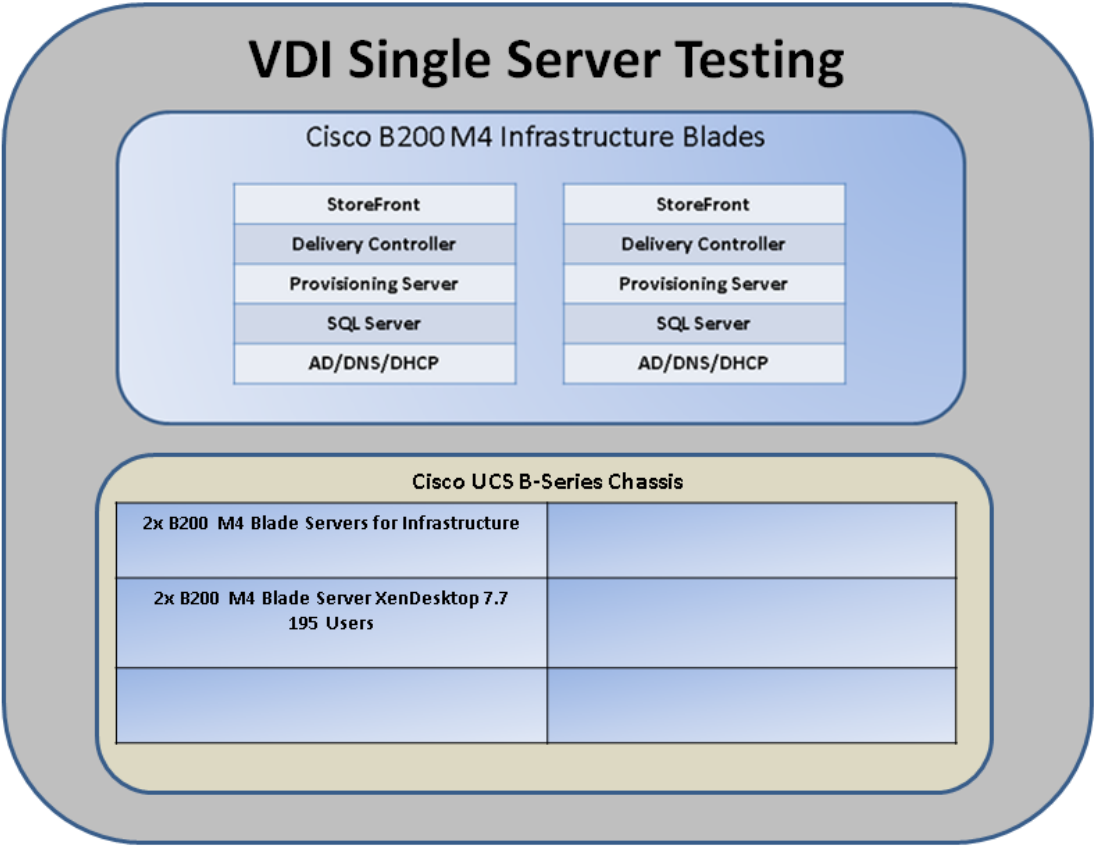


Callouts have been added throughout the data charts to indicate each phase of testing.

Test Phase	Description
Boot	Start all RDS and VDI virtual machines at the same time
Login	The Login VSI phase of test is where sessions are launched and start executing the workload over a 48 minutes duration
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files
Logoff	Sessions finish executing the Login VSI workload and logoff


Single Server on B200-M4 with 195 Users Test Results

Figure 24 Single-Server Recommended Maximum Workload for VDI with 195 Users



The recommended maximum workload for a Cisco UCS B200-M4 cartridge server with E5-2680 v3 processors and 384GB of RAM is 195 Windows 7 Desktops with 2vCPU and 1.7GB RAM.

Figure 25 B-200 M4 Single Server | XenDesktop 7.7 VDI | VSI Score



Performance data for the server running the workload follows:

Figure 26 B-200 M4 Single Server | XenDesktop 7.7 VDI | Host CPU Utilization
Single Server 195 Users Physical Cpu(_Total)\% Core Util Time




Figure 27 B-200 M4 Single Server | XenDesktop 7.7 VDI | Host Memory Utilization

B200-M4 Single Server Citrix XenApp 7.7

195 Users

Memory\NonKernel MBytes





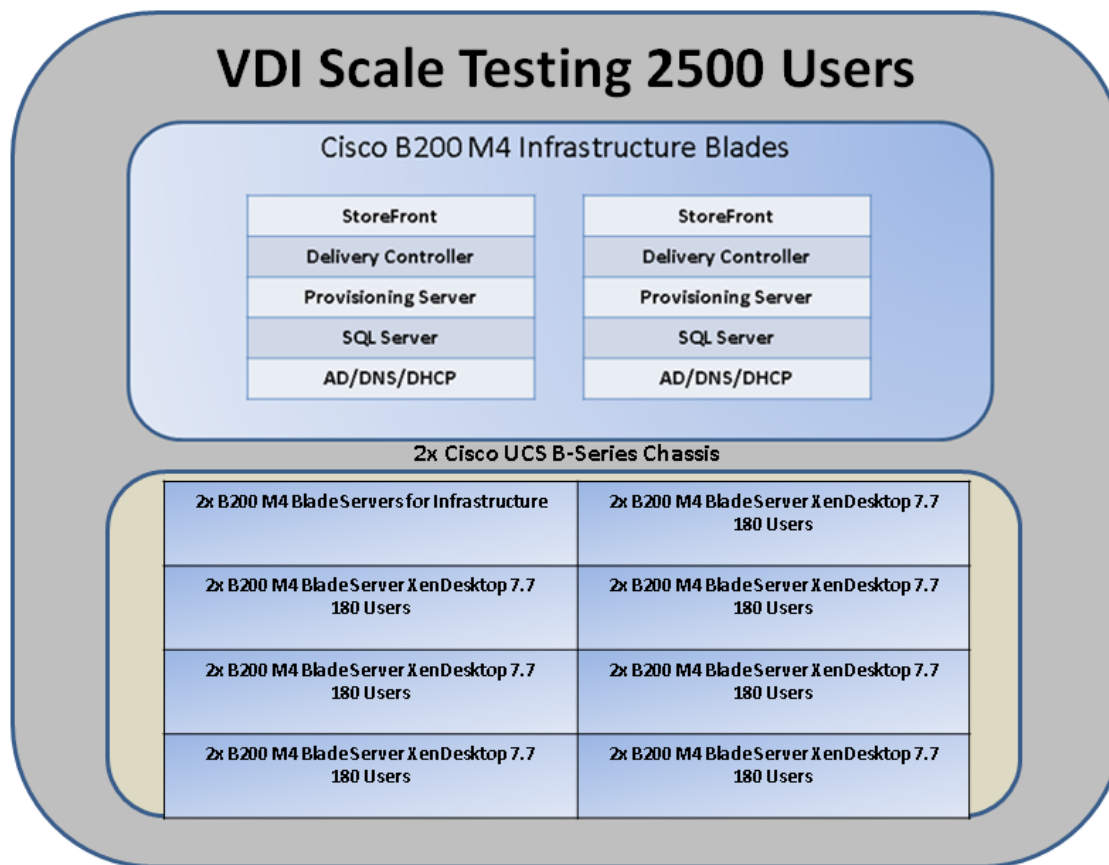
Figure 28 B-200 M4 Single Server | XenDesktop 7.7 VDI | Host Network Utilization

Single Server Bandwidth Usage; 195 Users




VDI Scaling with 2500 Users Test Results

Figure 29 VDI Server Scale Recommended Maximum Workload for VDI with 2500 Users



The recommended maximum workload for 14x Cisco UCS B200-M4 cartridge server with E5-2680 v3 processors and 384GB of RAM is 2500 Windows 7 Desktops with 2vCPU and 2GB RAM.

Figure 30 Cisco UCS B-200 M4 Server Scale w/ 2500 Users | XenDesktop 7.7 VDI | VSI Score



Solution Design Considerations

Each test detailed within this document utilizes the Citrix Provisioning Server's write cache on RAM overflow to disk feature. The server write cache on RAM overflow to disk feature allows virtual desktop administrators to utilize the free RAM available in each server blade to accelerate read operations targeted at the boot drives. Additionally, all VDI clients used in each test detailed in this document were non-persistent. The hybrid VDI CVD design easily scales to over 10,000 virtual desktops with no measurable impact to storage performance.



If the hybrid VDI CVD design requires the use of persistent VDI clients, the server write cache on RAM overflow to disk feature will continue to accelerate read operations, but will result in additional storage I/O activity. As a result, storage array sizing may be different from the specified design within this document.



Nimble Storage recommends using the storage array-sizing tool available in Nimble's InfoSight to help guide individual VDI storage sizing requirements: <https://infosight.nimblestorage.com/InfoSight/#sizing>

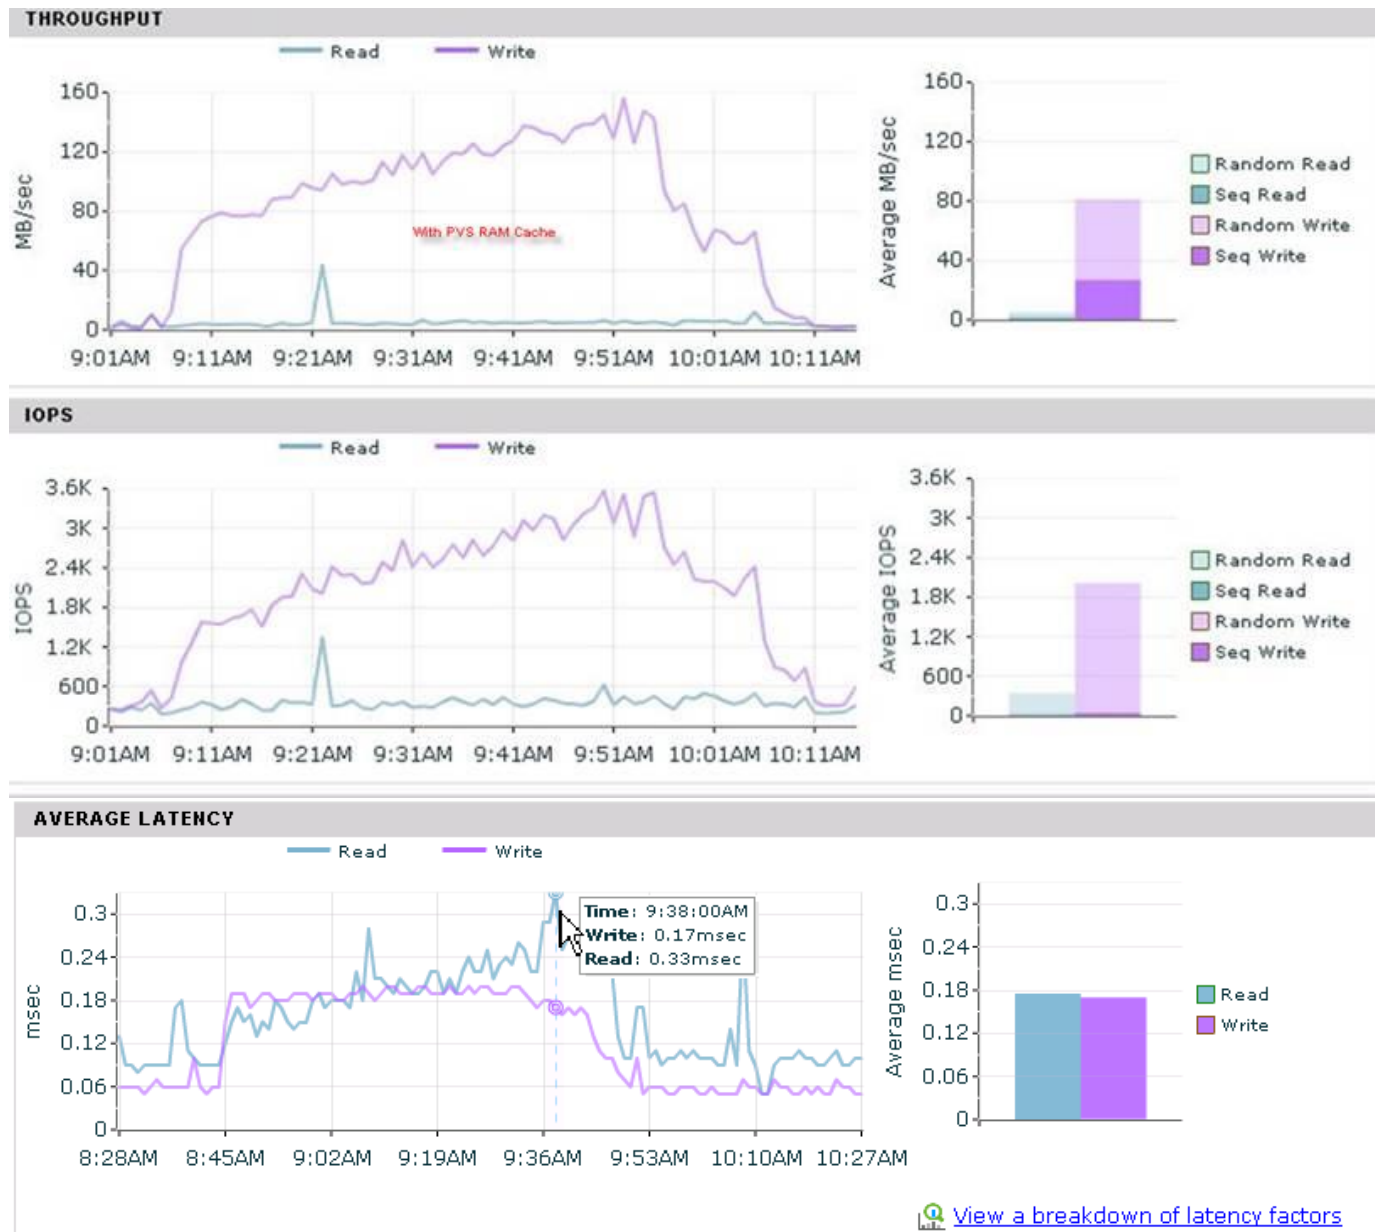
In the following charts, we demonstrated how the Nimble Array performed when all of the users PVS Write Cache was placed on the array. All 2500 VDI desktops utilized the cache to local hard disk feature in Citrix PVS while the M-Series servers were redirected to the PVS server to show the load for all 3000 users caching on the Nimble Array.

This is a steady state – LoginVSI test for 3000 users as described above.

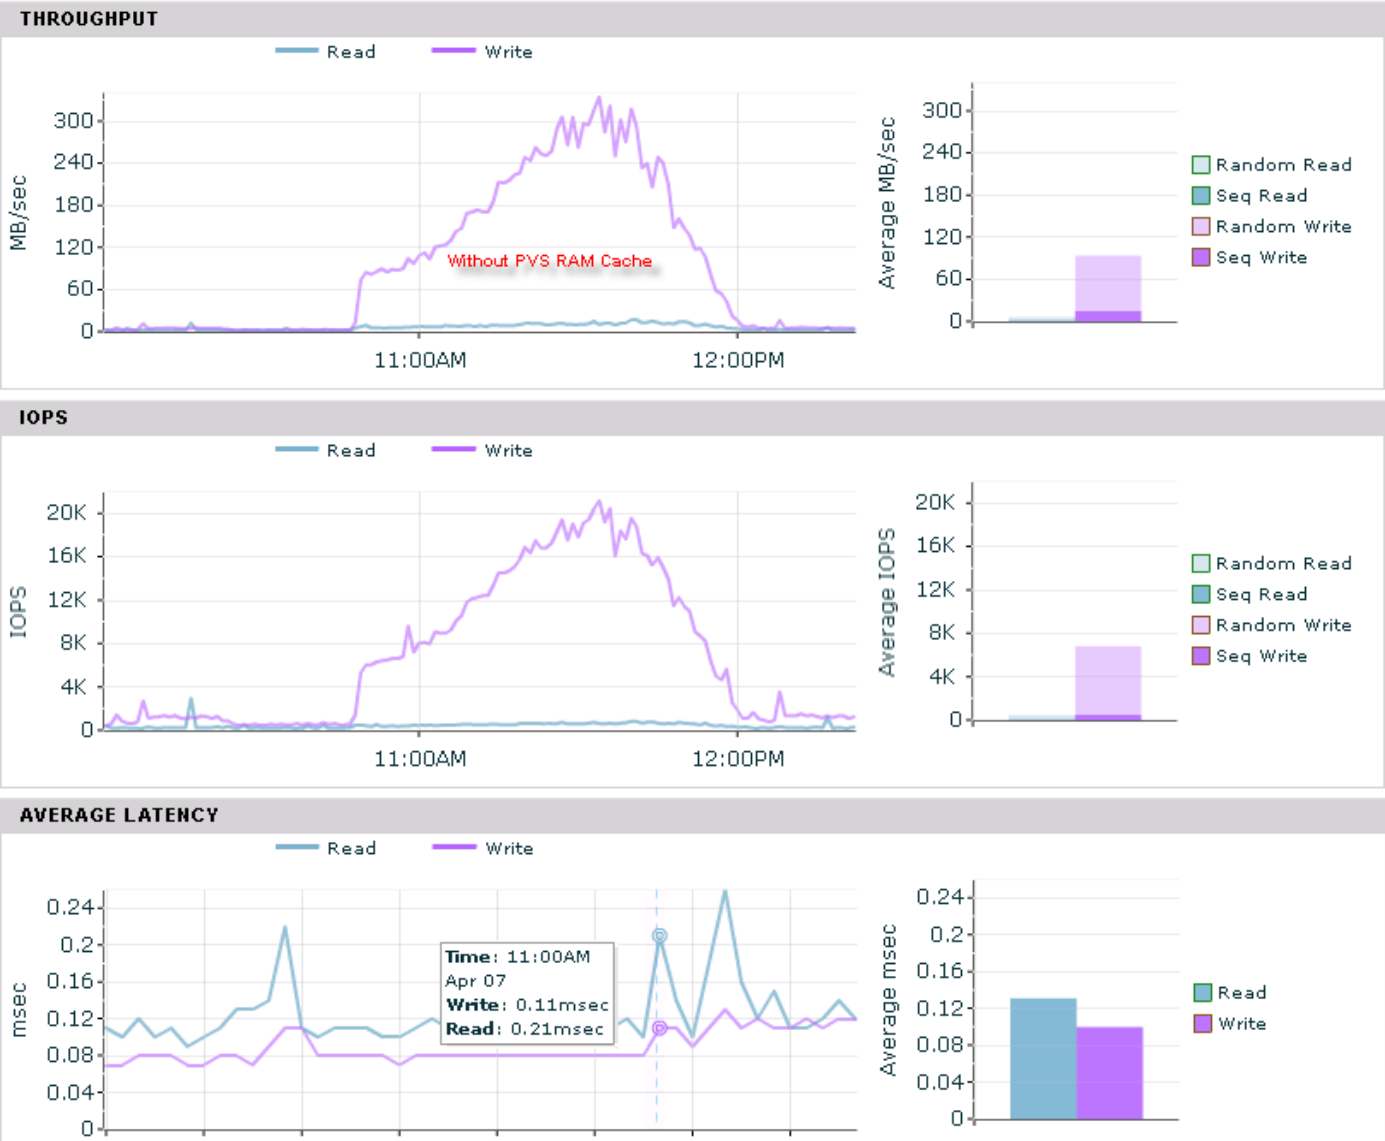


2500 XenDesktop Users VDI Scale Testing with LoginVSI with and without PVS RAM Cache

With PVS RAM Cache



Without PVS RAM Cache



2500 XenDesktop VM's+ 16 XenApp Servers BootStorm from vCenter






Figure 31 Volume performance for Infrastructure Data Stores

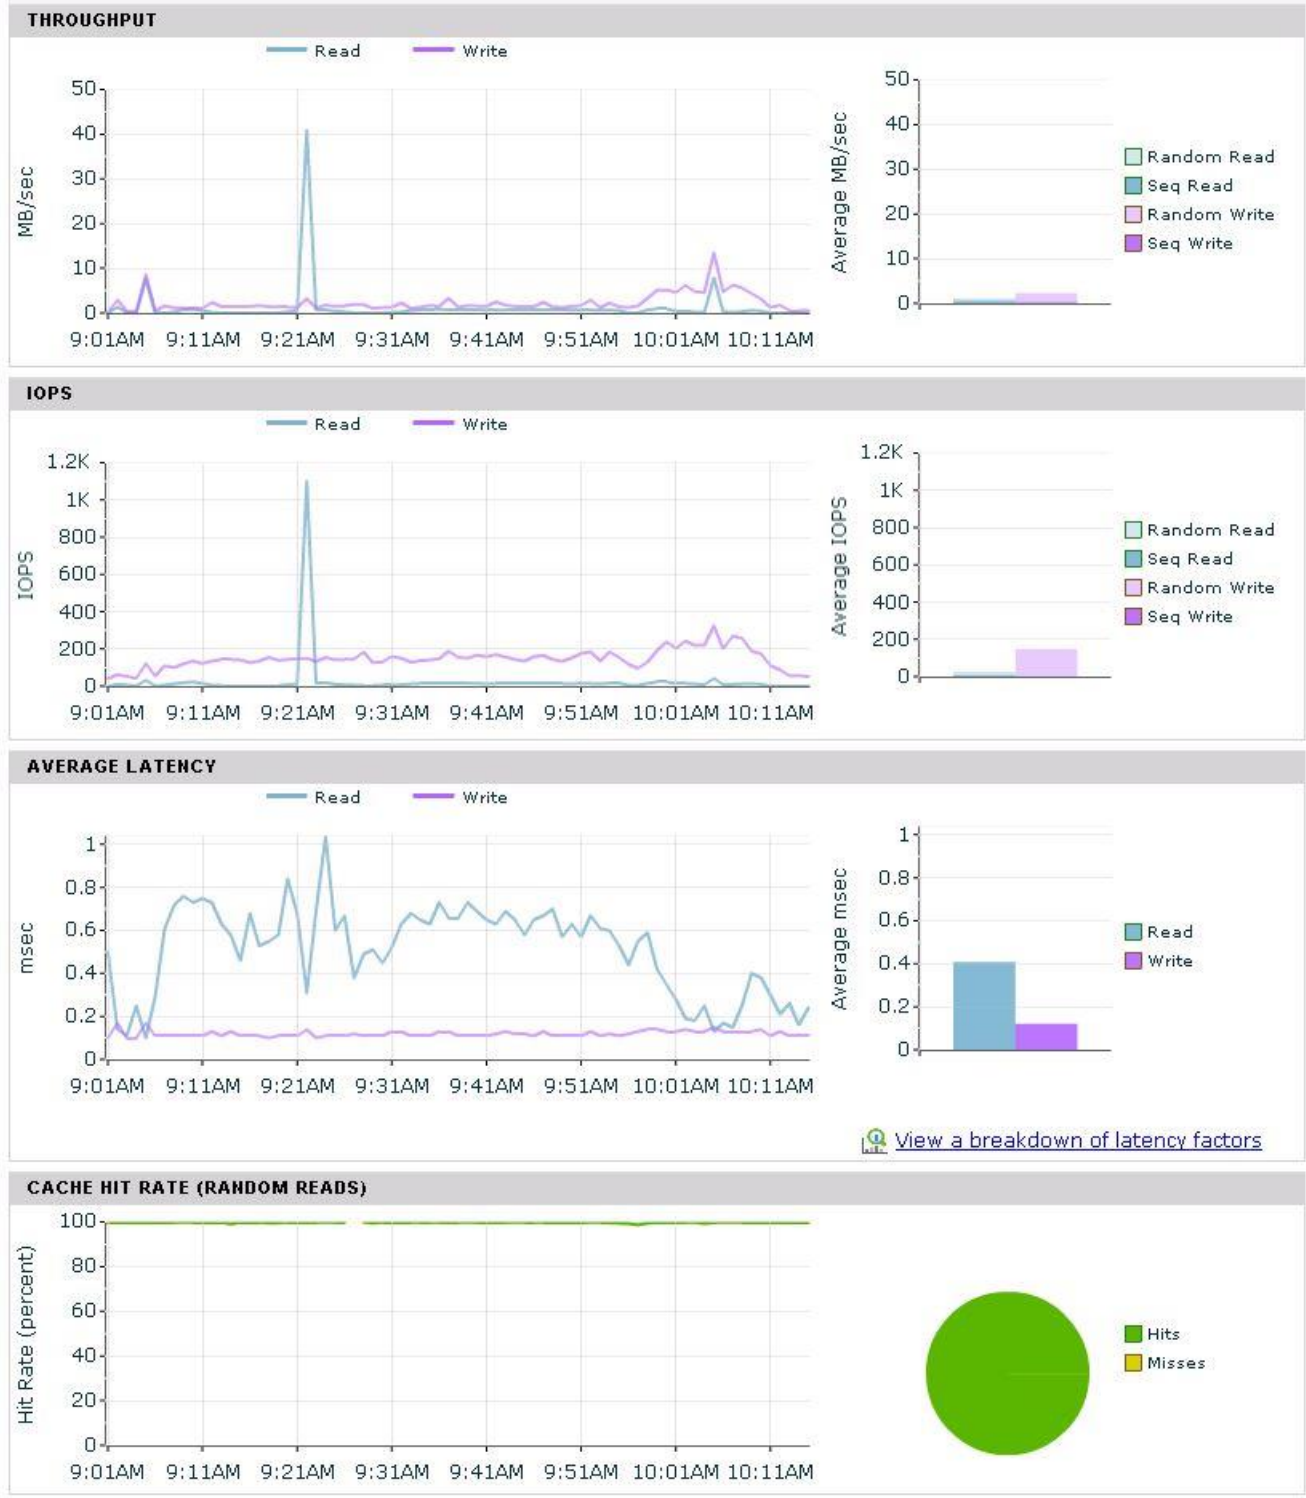


Figure 32 Performance for Provisioning Services Volume



Figure 33 Write Cache volume performance for Single Server

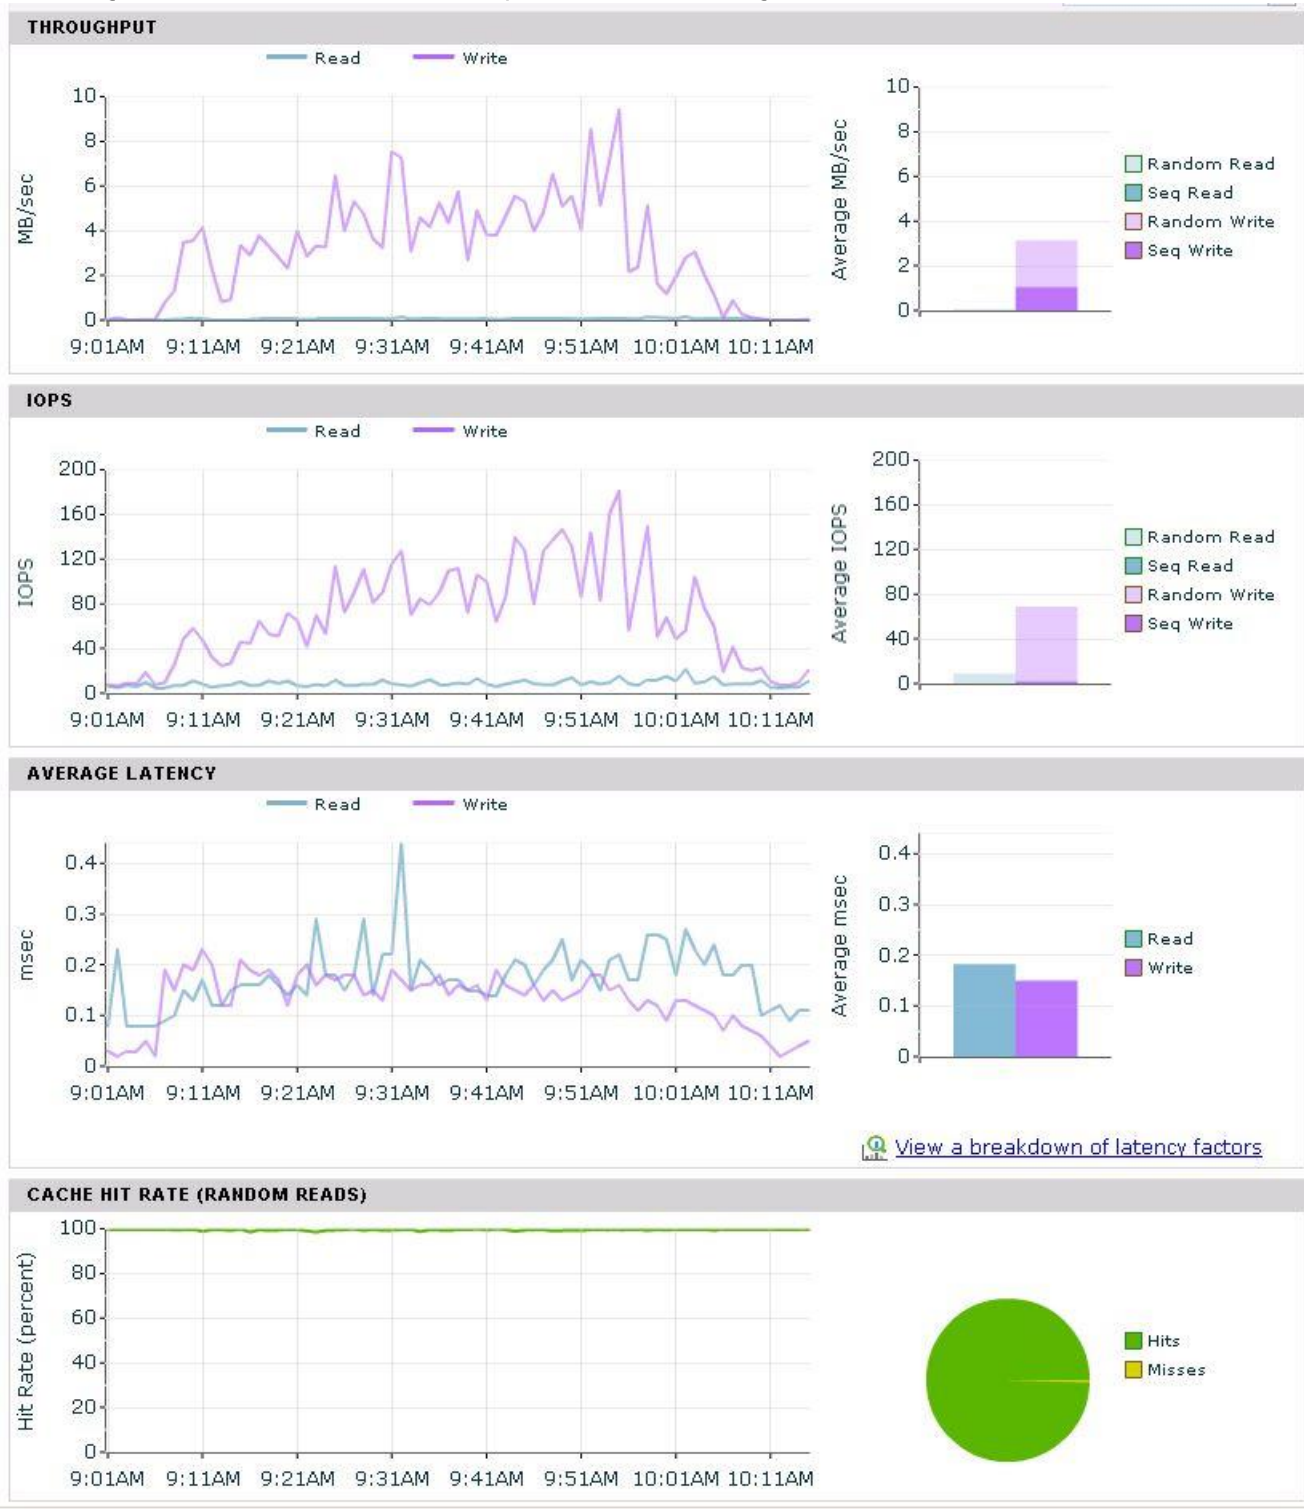



Figure 34 vHBA Physical disk Adapter Performance Throughput for a Single Host During VDI Testing



Single Server Testing Utilizing the NVIDIA M6 Card and vGPU

Cisco UCS B200 M4 Blade Server with NVIDIA M6 GRID Card

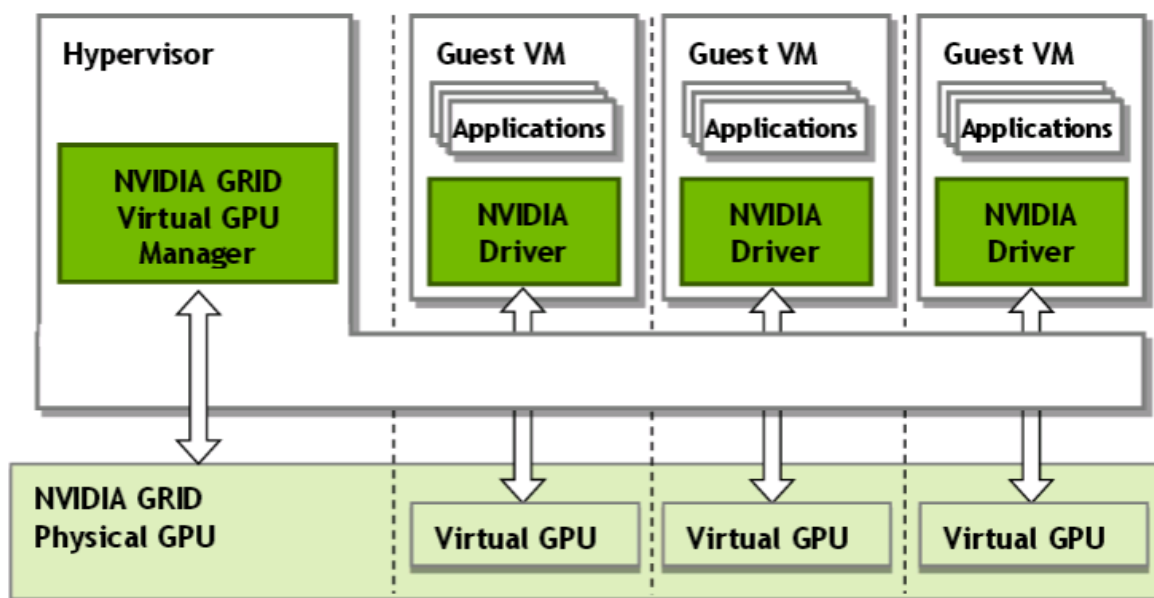
Earlier this year (2016) NVIDIA released the M6 GRID cards capable of fitting into a Cisco B200 M4 Server. This release opened up a new world of possibilities in our already robust VDI deployments and adding the capability of Graphics Desktops to our traditional VDI solutions. Our solution here is meant to run in parallel on the same B200 Hardware as a non-graphics user sessions. We have determined that we can run 180-195 individual users on a single blade. With the introduction of the M6 card, we are able to run 150-165 user sessions on the same blade that will be shared with vGPU enabled VMs.

The difference between a CPU and a GPU is that CPU only has a dozen or so cores (in our case the 2680v4 has 12) and runs tasks sequentially. Where a GPU has many more cores and runs tasks in parallel.

Under the control of NVIDIA's Virtual GPU Manager that runs on ESXi, the GRID physical GPUs are able to support multiple virtual GPU devices (vGPU) that can be assigned directly to our guest VMs.

Guest VMs can use the vGPUs the same as a physical GPU that has been passed through directly to the VM from the hardware. **With an NVIDIA driver loaded in the guest VM's OS the end user is able** to utilize GPU functionality. 0 illustrates the architecture of the NVIDIA vGPU.

Figure 35 NVIDIA vGPU Architecture




Install and Configure NVIDIA M6 Card

In this study we tested the NVIDIA M6 card deployed as a vGPU to Citrix XenDesktop 7.7 virtual desktops. We used the Power User vGPU Profile to enable a density of 16 desktops using 512MB of video memory.

Physical Installation of the NVIDIA M6 Card into the Cisco UCS B200 M4 Server

The NVIDIA M6 graphics processing unit (GPU) provides graphics and computing capabilities to the server. The GPU package consists of the three elements shown in Figure 36.

Figure 36 NVIDIA M6 GPU Package



1	NVIDIA M6 GPU (CPU and heat sink)	2	T-shaped wrench
3	Custom standoff		

Before You Begin

Before installing the NVIDIA M6 GPU, do the following:

- Remove any adapter card, such as a VIC 1380, VIC 1280, or PT extender card from slot 2. You cannot use any other card in slot 2 when the NVIDIA M6 GPU is installed.
- Upgrade your Cisco UCS system to a version of Cisco UCS Manager that supports this card. Refer to the latest version of the Release Notes for Cisco UCS Software at the following URL for information about supported hardware: <http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-release-notes-list.html>.

To install the NVIDIA M6 GPU, complete the following steps:

1. Use the T-shaped wrench that comes with the GPU to remove the existing standoff at the back end of the motherboard.
2. Install the custom standoff in the same location at the back end of the motherboard.
3. Position the GPU over the connector on the motherboard and align all captive screws to the standoff posts (callout 1).

- 4. Tighten the captive screws (callout 2).

Figure 37 Installing the NVIDIA MG GPU





Figure 38 illustrates the GPU installed in a Cisco UCS B200 M4 blade server.

Figure 38 Installed NVIDIA M6 GPU



1	Front of server	2	Custom standoff screw
---	-----------------	---	-----------------------

Install the NVIDIA VMware VIB Driver

To install the NVIDIA VMware VIB Driver, complete the following steps:

- 1. Download the latest drivers and software packages from NVidia’s Web Site.

2. Upload the VIB file to the /tmp directory of the ESXi host.

```
10.10.70.144 - PuTTY
login as: root
Using keyboard-interactive authentication.
Password:
The time and date of this login have been sent to the system logs.

VMware offers supported, powerful system administration tools. Please
see www.vmware.com/go/sysadmintools for details.

The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
[root@SP-VDI-14:~] cd /tmp
[root@SP-VDI-14:/tmp] ls
NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
apa_start.log
dpafifo
nfsgssd_krb5cc
probe.session
```

3. Install the latest driver: `esxcli software vib install -v /tmp/{Latest Driver Package Name}`



Host must be in Maintenance Mode to install.

```
The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
[root@SP-VDI-14:~] cd /tmp
[root@SP-VDI-14:/tmp] ls
NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
apa_start.log
dpafifo
nfsgssd_krb5cc
probe.session
vem-vmkbinding.log
vemdpd_cpu_mhz
vemdpd_mem_kb
vmware-root
[root@SP-VDI-14:/tmp] esxcli software vib install -v /tmp/NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
```








A message will validate that the VIB installed correctly.

```
[root@SP-VDI-14:/tmp] esxcli software vib install -v /tmp/NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  VIBs Installed: NVIDIA_bootbank_NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-1OEM.600.0.0.2494585
  VIBs Removed:
  VIBs Skipped:
[root@SP-VDI-14:/tmp]
```

4. Validate the driver was installed by running the command 'nvidia-smi' command.

```
[root@SP-VDI-14:/tmp] esxcli software vib install -v /tmp/NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-10EM.600.0.0.2494585.vib
Installation Result
  Message: Operation finished successfully.
  Reboot Required: false
  VBIs Installed: NVIDIA_bootbank_NVIDIA-vGPU-VMware_ESXi_6.0_Host_Driver_352.83-10EM.600.0.0.2494585
  VBIs Removed:
  VBIs Skipped:
[root@SP-VDI-14:/tmp] nvidia-smi
Wed Mar 23 23:40:35 2016
+-----+
| NVIDIA-SMI 352.83      Driver Version: 352.83      |
+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+
|  0 Tesla M6             On      | 0000:81:00.0  Off  |           Off       |
| N/A   44C    P8      16W / 100W |  14MiB /  8191MiB |      0%      Default |
+-----+-----+-----+
+-----+
| Processes:
| GPU      PID   Type   Process name                      GPU Memory
|-----|
| No running processes found
+-----+
[root@SP-VDI-14:/tmp] █
```

5. By Default the M6 cards come in Compute mode. We will utilize them in Graphics mode in this study. You will need to download the gpumodeswitch utility from NVidia’s web site. In this exercise, we used the boot ISO which loads a Linux environment with the gpumodeswitch utility already loaded.

Name ^	Type	Compressed size	Password
 gpumodeswitch	File	766 KB	No
 gpumodeswitch	Application	618 KB	No
 gpumodeswitch	Virtual CloneDrive	47,289 KB	No
 gpumodeswitch	Compressed (zipped) Folder	47,268 KB	No
 GRID gpumodeswitch User Guide	Firefox HTML Document	691 KB	No
 LICENSES	Text Document	19 KB	No
 nvflash64.sys	System file	8 KB	No

6. Mount the ISO file through the UCSM KVM and reboot the host.
7. When the Linux shell loads, enter the command: gpumodeswitch -gpumode graphics



Type ‘Y’ when prompted to switch all adapters to Graphics. When it completes, reboot back into ESXi.

```
# gpumodeswitch --gpumode graphics


NVIDIA GPU Mode Switch Utility Version 1.02
Copyright (C) 2015, NVIDIA Corporation. All Rights Reserved.

Update GPU Mode of all adapters to "graphics"?
Press 'y' to confirm or 'n' to choose adapters or any other key to abort:
```

Configure a VM with a vGPU

To configure a vGPU for a VM, complete the following steps:

1. Select Edit Settings in the VSphere Web client for the VM you want to add to the vGPU.
2. Select the Virtual Hardware tab.
3. In the New device section, select Shared PCI Device to add the NVIDIA GRID Card.



4. Select the GPU Profile you want to run. In this study, we wanted to achieve a density of 16 vGPU machines on this host so we chose Profile 'grid_m6-0b' which allocates 512Mb per VM for a total of 16 per blade with the M6 Card.

NMBL005 - Edit Settings

Virtual Hardware | VM Options | SDRS Rules | vApp Options

CPU: 2
 Memory: 2048 MB
 Hard disk 1: 6 GB
 SCSI controller 0: LSI Logic SAS
 Network adapter 1: DHCP (SP-N1KV) ☒ Connect...
 Floppy drive 1: Client Device ☐ Connect...
 Video card: Specify custom settings
 PCI device 0: **NVIDIA GRID vGPU**
 GPU Profile: **grid_m6-0b**
 VMCI device
 Other Devices

grid_m6-8q
 grid_m6-4q
 grid_m6-2q
 grid_m6-2b
 grid_m6-1q
 grid_m6-1b

New device: ----- Select ----- Add

Compatibility: ESXi 6.0 and later (VM version 11)

OK Cancel

*GPU Profiles for the M6 are as follows:

Card	Physical GPUs	GRID Virtual GPU	Intended Use Case	Frame Buffer (Mbytes)	Virtual Display Heads	Max Resolution per Display Head	Maximum vGPUs	
							Per GPU	Per Board
Tesla M6	1	M6-8Q	Designer	8192	4	3840x2160	1	1
		M6-4Q	Designer	4096	4	3840x2160	2	2
		M6-2Q	Designer	2048	4	2560x1600	4	4
		M6-1Q	Power User, Designer	1024	2	2560x1600	8	8
		M6-0Q	Power User, Designer	512	2	2560x1600	16	16
		M6-2B	Power User	2048	2	2560x1600	4	4
		M6-1B	Power User	1024	2	2560x1600	8	8
		M6-0B	Power User	512	2	2560x1600	16	16

Install the GPU Drivers into Windows VM

It is important to note that the drivers installed with the Windows VDI desktop must match the version that accompanies the driver for the ESXi host. So if you downgrade or upgrade the ESXi host vib, you must do the same with the NVIDIA driver in your Windows master image.

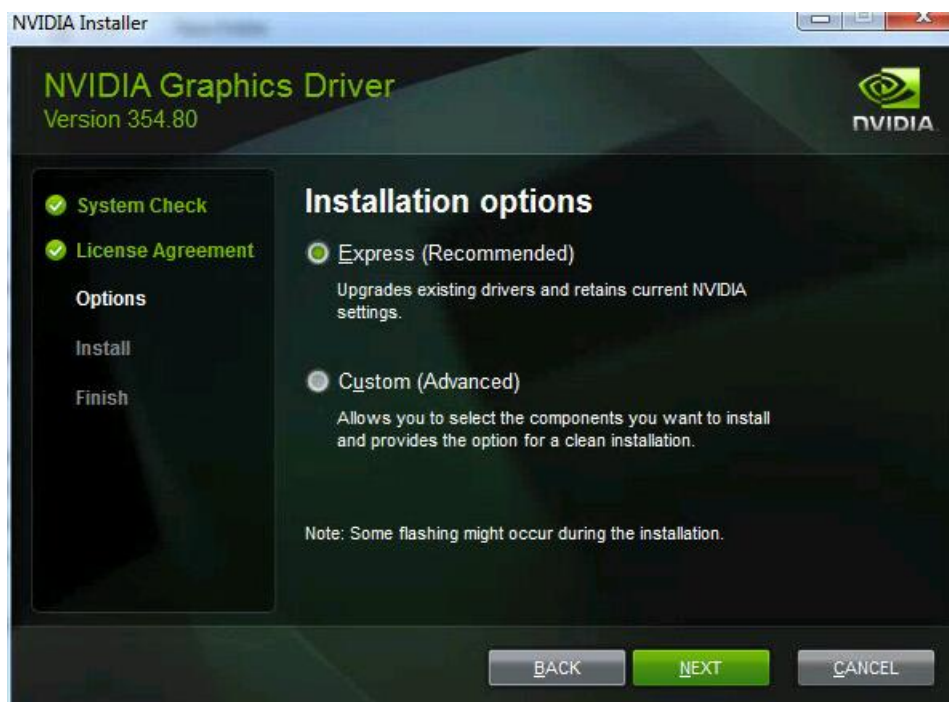
In this study we used ESXi Host Driver version 352.83 and 354.80 for the Windows VDI image. These drivers come in the same download package from NVIDIA.

To install the GPU drivers into your Windows VM, complete the following steps:

1. Since our image is deployed via Citrix PVS, first place the image in Private Mode.
2. Double-click file '354.80_grid_win8_win7_international'



3. Select Agree and Continue.



4. Click Next to use Express Installation.



5. The driver and software will be installed and click 'Finish' to complete install.

Install and configure NVIDIA Grid License Server

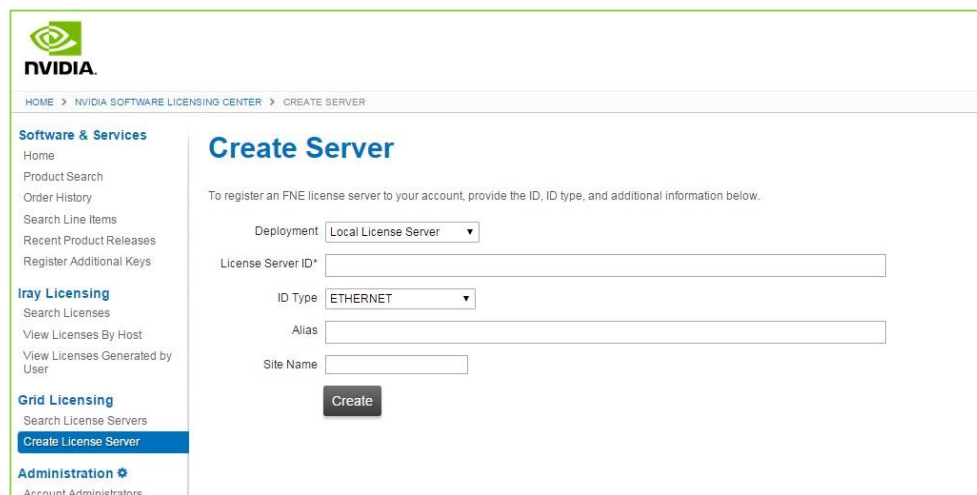
To use NVIDIA's vGPU features you must setup a Grid Licensing server. The detailed instructions for setting up a Grid License server can be found in the Grid Quick Start guide:

<http://images.nvidia.com/content/grid/pdf/grid-2.0-quick-start-guide.pdf>

The license server requires a fixed IP address. The IP address may be assigned through DHCP or can be statically configured. The server's Ethernet MAC address is used as a unique identifier when registering the server and generating licenses in NVIDIA's licensing portal. The server runs on either Windows or Linux.

To create a server interface, complete the following steps:

1. Select **Create License Server** from under **GRID Licensing** in the left pane of the NVIDIA Software Licensing Center page to display the Create Server page.



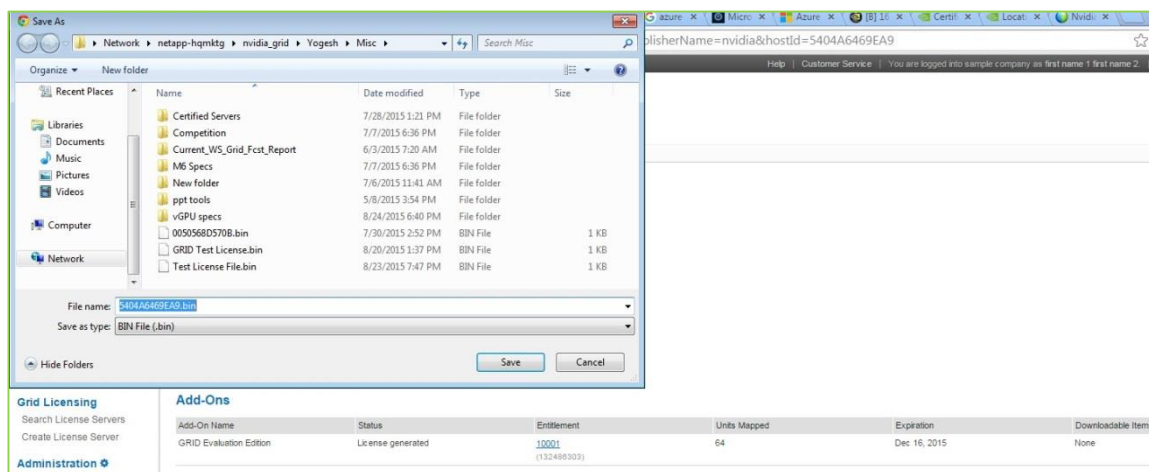
The screenshot shows the NVIDIA Software Licensing Center interface. The left sidebar contains navigation links under 'Software & Services', 'Iray Licensing', 'Grid Licensing', and 'Administration'. The 'Create License Server' link under 'Grid Licensing' is highlighted. The main content area is titled 'Create Server' and includes a sub-header: 'To register an FNE license server to your account, provide the ID, ID type, and additional information below.' The form contains the following fields: 'Deployment' (a dropdown menu set to 'Local License Server'), 'License Server ID*' (a text input field), 'ID Type' (a dropdown menu set to 'ETHERNET'), 'Alias' (a text input field), and 'Site Name' (a text input field). A 'Create' button is located below the 'Site Name' field.

2. Fill in your server details on the Create Server page.

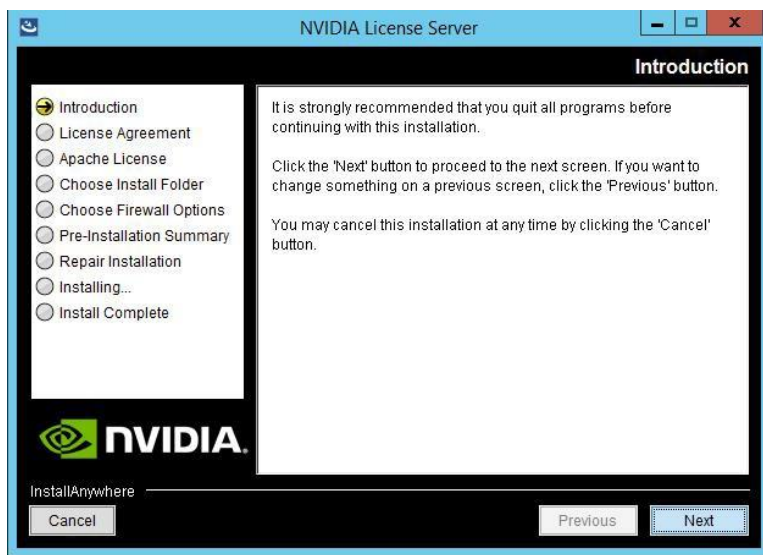


The License Server ID field is the MAC address of the VM of the License server.

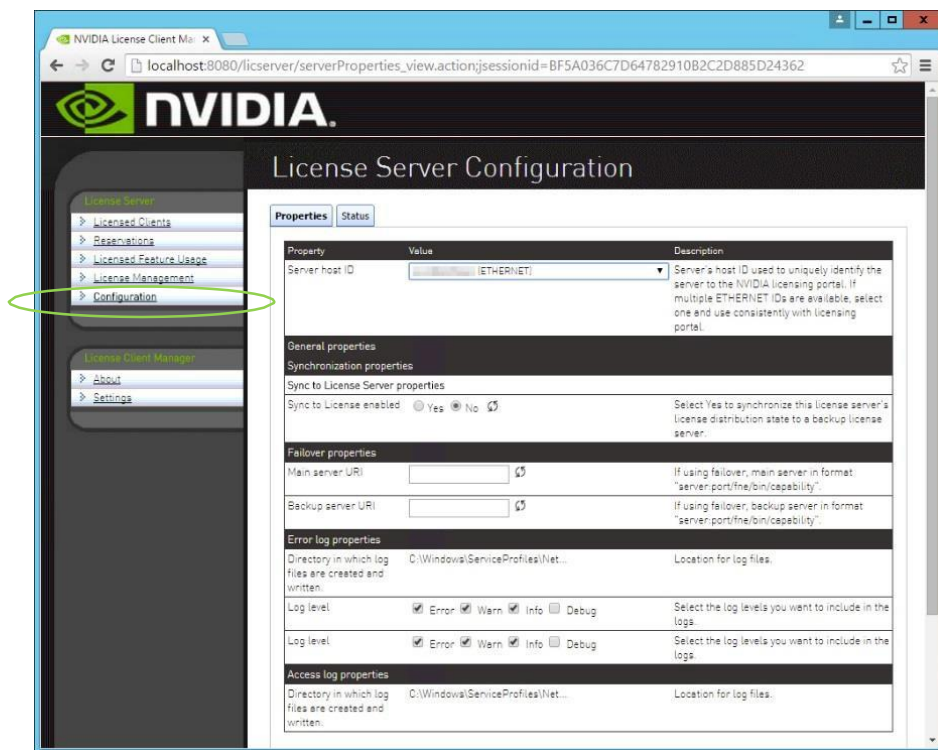
3. Save the .bin file onto your license server for installation. Java is required to install the NVIDIA GRID License Server. The package comes in a **.zip** file.



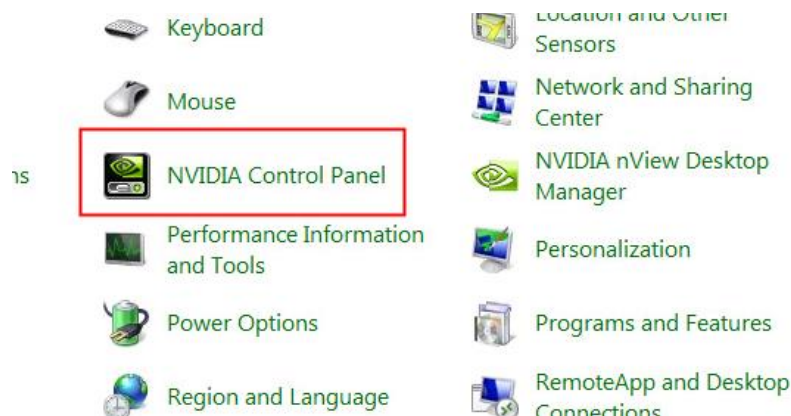
4. Unzip the license server installer.
5. Run setup.exe and follow the installation wizard.



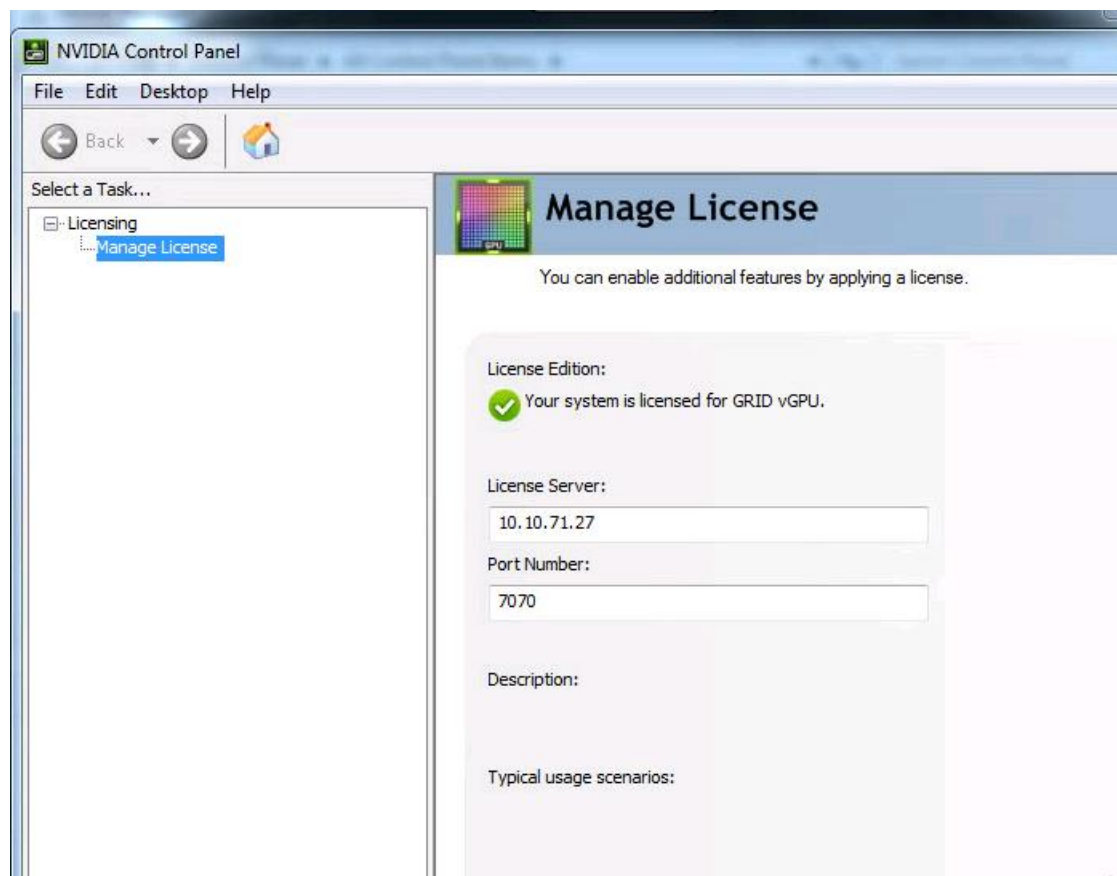
6. Go to <http://<FQDN of the license Server>:8080/licserver> to display the License Server Configuration page. You will need the License Server's MAC Address to generate a license .bin file on the portal.



7. Select Configuration from the menu in the left pane.
8. Use the License Server Configuration menu to install the .bin file:
 - a. Select Choose File.
 - b. Use the file browser to locate the .bin file downloaded from the licensing portal web site.
9. When the License server is properly installed, we must point our master image to the license server so the VMs with vGPUs can obtain a license.
 - a. In Windows – Control Panel, double click the NVidia Control Panel.



- b. In the Control Panel, enter the IP or FQDN of the Grid License Server. You should receive a result similar to the below image.



Testing Methodology and Results for the NVIDIA M6 Cards

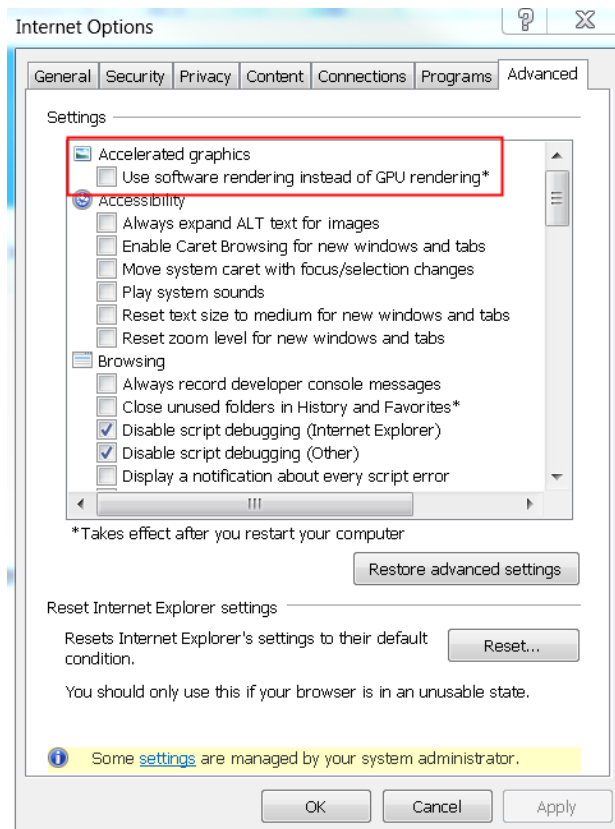


In this study we have shown how the Login VSI Knowledge Worker workload can successfully run on the Cisco UCS B200 M4 Servers along with the performance results and charts. To incorporate the NVIDIA Grid M6 cards using the vGPU we were able to demonstrate that 16 vGPU enabled VMs could run simultaneously on a B200 M4 with 100+ standard, non-vGPU VMs running the Knowledge Worker workload. What this showed was the ability for VDI administrators to add Power User or Graphics workers to the same hardware as task workers.

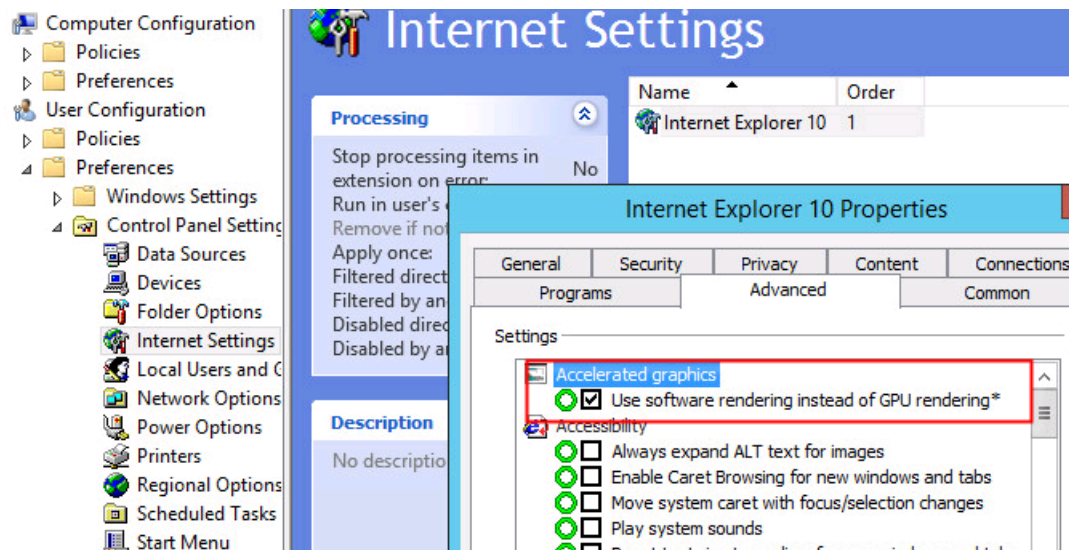
To illustrate the usage of the vGPU in the graphics VMs, we developed a test workload that incorporates software already used in the LoginVSI knowledge worker workload. We will illustrate how Internet Explorer 11 can utilize the vGPU while running online HTML5 videos.

Internet Explorer 11 Configuration

In the Advanced Settings tab of Internet Explorer 11, the setting 'Use software rendering instead of GPU rendering*' is disabled by default, thus allowing the system GPU to render content in IE.

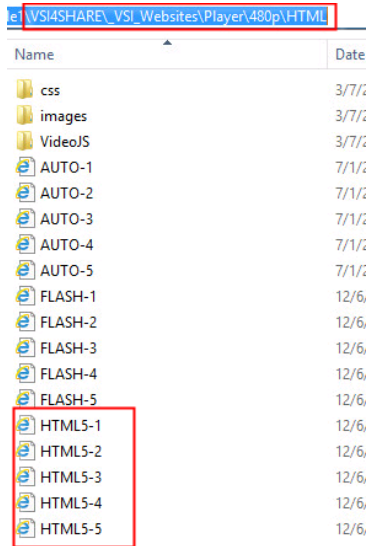


Using a Group Policy Preference, we were able to Enable and Disable quickly.



Test Configurations

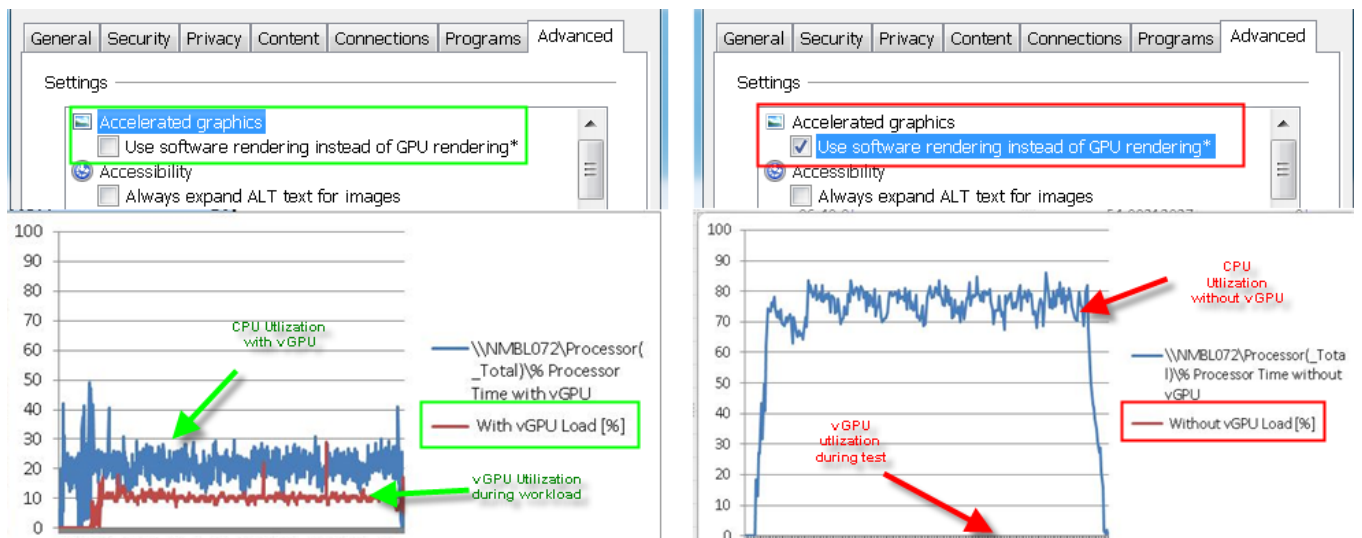
1. To heavily utilize the IE rendering feature we created a custom workload with LoginVSI 4.1.4 that launches Internet Explorer sessions running the 480p HTML5 videos included on the VSI File Share.



Name	Date
css	3/7/12
images	3/7/12
VideoJS	3/7/12
AUTO-1	7/1/12
AUTO-2	7/1/12
AUTO-3	7/1/12
AUTO-4	7/1/12
AUTO-5	7/1/12
FLASH-1	12/6/11
FLASH-2	12/6/11
FLASH-3	12/6/11
FLASH-4	12/6/11
FLASH-5	12/6/11
HTML5-1	12/6/11
HTML5-2	12/6/11
HTML5-3	12/6/11
HTML5-4	12/6/11
HTML5-5	12/6/11

2. The workload is programmed to launch each of the 5 HTML5 videos, 2 times for a total of 10, in a staggered fashion over a period of 30 minutes.

During our typical workloads, we enable this setting to allow the software rendering instead of GPU rendering, however with the introduction of VMs with vGPUs, we ran a workload with high graphics utilization and this setting disabled to observe IE offload its rendering to the vGPU. The following were our results.



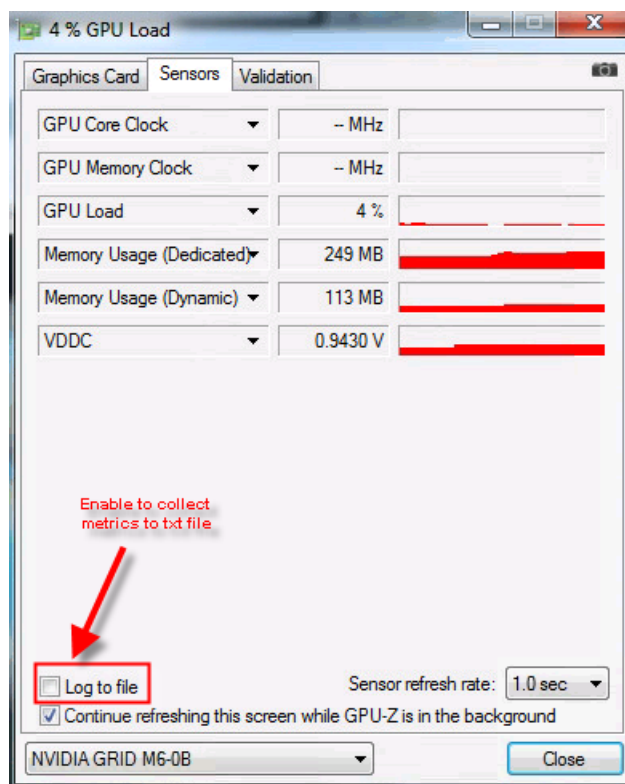
GPU Performance Metrics

GPU metrics and how to gather and present the performance data is not always a straight forward discussion. The `nvidia-smi` command that is built in when you install the GPU Manager VIB can show vGPU utilization for the 16 VMs deployed.

NVIDIA-SMI 352.83 Driver Version: 352.83									
GPU Name		Persistence-M		Bus-Id		Disp.A		Volatile Uncorr. ECC	
Fan	Temp	Perf	Pwr:Usage/Cap	Memory-Usage		GPU-Util		Compute M.	
0	Tesla M6		On	0000:81:00.0		Off		Off	
N/A	42C	P8	16W / 100W	6720MiB / 8191MiB		0%		Default	
Processes:									
GPU	PID	Type	Process name				GPU Memory Usage		
0	152218	C+G	NMBL008				416MiB		
0	152233	C+G	NMBL015				416MiB		
0	152234	C+G	NMBL005				416MiB		
0	152235	C+G	NMBL014				416MiB		
0	152236	C+G	NMBL003				416MiB		
0	152237	C+G	NMBL010				416MiB		
0	152245	C+G	NMBL002				416MiB		
0	152246	C+G	NMBL007				416MiB		
0	152324	C+G	NMBL004				416MiB		
0	152325	C+G	NMBL009				416MiB		
0	152326	C+G	NMBL072				416MiB		
0	152337	C+G	NMBL006				416MiB		
0	152338	C+G	NMBL012				416MiB		
0	152339	C+G	NMBL011				416MiB		
0	152340	C+G	NMBL013				416MiB		
0	173345	C+G	NMBL001				416MiB		

GPU-Z

The other tool used to gather metrics in this scenario was GPU-Z. GPU-Z is a popular tool that can measure the GPU Load usage and other metrics. We were able to collect GPU Load usage metrics to make our graphs by Logging to a text file. This in conjunction with PerfMon allowed us to measure the GPU vs. CPU load when our Internet Explorer HTML-5 videos were running.



Validated Hardware and Software

Table 6 lists all the components and software versions used in validating the SmartStack design. Cisco and Nimble Storage provides interoperability matrices that should also be consulted to ensure support for a specific SmartStack implementation.

- Cisco UCS Hardware and Software Interoperability Tool: <http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>
- Interoperability Matrix for Cisco Nexus and MDS 9000 Products: <http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-device-support-tables-list.html>
- Nimble Support Matrix: https://infosight.nimblestorage.com/InfoSight/cgi-bin/viewPDFFile?ID=array/pubs_support_matrix_2_3_rev_f.pdf (requires login)
- VMware Compatibility Guide: <http://www.vmware.com/resources/compatibility/search.php>

Table 6 Infrastructure Components and Software Revisions

	Components	Software Version	Comments
Network	Cisco Nexus 9372PX (N9k-9372PX)	6.1(2)I3(4a)	Cisco Platform Switch for ToR, MoR, EoR deployments; Provides connectivity to users and other networks and deployed in NX-OS Standalone mode

	Components	Software Version	Comments
	Nexus 1000V	5.2(1)SV3(1.5a)	(Optional) Distributed Virtual Switch
	Cisco UCS 6248UP FI	3.1.1e	Fabric Interconnect with embedded management
	Cisco MDS 9148S	6.2(13a)	16G Multilayer Fabric Switch
Compute	Cisco UCS 5108	3.1.1e	Blade Server Chassis
	Cisco UCS B200M4 servers	3.1.1e	Blade Servers
	Cisco ENIC Driver	2.3.0.6*	Cisco VIC Ethernet driver
	Cisco FNIC Driver	1.6.0.24**	Cisco VIC FCoE driver
	Adapter Firmware	4.1(1d)	Cisco VIC Adapter Firmware
Manage-ment	Cisco UCS Manager	3.1.1e	Embedded Management
	vCenter plugin for Nimble	TBD	
	vCenter plugin for UCS	TBD	
Storage	Nimble CS700	NimbleOS 2.3.12	
	Nimble NCM for ESXi	2.3.1	Build: 2.3.1-600006
Virtualiza-tion	VMware vSphere	6.0 U1a	Cisco ISO Available
	VMware vCenter Server	6.0 U1	Appliance
Tools	LoginVSI	4.1.4.2	
Other	Microsoft Active Directory/DNS	2012R2	

* During this study ENIC driver version 2.3.0.6 was tested but it is highly recommended to upgrade this driver to the latest version available.

** During this study FNIC driver version 1.6.0.24 was tested but it is highly recommended to upgrade this driver to the latest version available.

Bill of Materials (BOM)

The BOM below lists the major components validated but it is not a comprehensive list.

Table 7 SmartStack Bill of Materials

Line	SKU	Description	Quantity
1.0	UCSB-B200-M4	UCS B200 M4 w/o CPU, mem, drive bays, HDD, mezz	1
1.1	UCS-CPU-E52620D	2.40 GHz E5-2620 v3/85W 6C/15MB Cache/DDR4 1866MHz	2
1.2	UCS-MR-1X162RU-A	16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	8
1.3	UCSB-MLOM-40G-03	Cisco UCS VIC 1340 modular LOM for blade servers	1
1.4	UCSB-HS-EP-M4-F	CPU Heat Sink for UCS B200 M4/B420 M4 (Front)	1
1.5	UCSB-HS-EP-M4-R	CPU Heat Sink for UCS B200 M4/B420 M4 (Rear)	1
1.6	UCSB-LSTOR-BK	FlexStorage blanking panels w/o controller, w/o drive bays	2
1.7		NVidia M6 GPU card for B200 M4	1
1.8	C1UCS-OPT-OUT	Cisco ONE Data Center Compute Opt Out Option	1
2.0	UCSB-5108-AC2-UPG	UCS 5108 Blade Server AC2 Chassis, 0 PSU/8 fans/0 FEX	1
2.1	N01-UAC1	Single phase AC power module for UCS 5108	1
2.2	N20-FAN5	Fan module for UCS 5108	8
2.3	N20-CBLKB1	Blade slot blanking panel for UCS 5108/single slot	7
2.4	N20-CAK	Accessory kit for UCS 5108 Blade Server Chassis	1
2.5	N20-FW014	UCS Blade Server Chassis FW Package 3.1	1
2.6	UCSB-5108-PKG-HW	UCS 5108 Packaging for chassis with half width blades.	1
2.7	UCSB-PSU-2500ACDV	2500W Platinum AC Hot Plug Power Supply - DV	2
2.8	CAB-C19-CBN	Cabinet Jumper Power Cord, 250 VAC 16A, C20-C19 Connectors	2
2.9	UCS-IOM-2208XP	UCS 2208XP I/O Module (8 External, 32 Internal 10Gb ports)	2
3.0	UCS-FI-6248UP-UPG	UCS 5108 Blade Server AC2 Chassis, 0 PSU/8 fans/0 FEX	1
3.1	N10-MGT014	UCS Manager 3.1	2
3.2	UCS-ACC-6248UP	UCS 6248UP Chassis Accessory Kit	1
3.3	UCS-PSU-6248UP-AC	UCS 6248UP Power Supply/100-240VAC	2
3.4	UCS-BLKE-6200	UCS 6200 Series Expansion Module Blank	1
3.5	UCS-FAN-6248UP	UCS 6248UP Fan Module	2
3.6	UCS-FI-DL2	UCS 6248 Layer 2 Daughter Card	1
3.7	CAB-9k12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	2
3.8	UCS-L-6200-10G-C	2 nd Gen FI License to connect C-direct only	1
5.0	DS-C9148S-D12P8K9	MDS 9148S 16G FC switch, w/ 12 active ports + 8G SW SFPs	1
5.1	DS-SFP-FC8G-SW	8Gbps Fibre Channel SW SFP+, LC	12
5.2	DS-9148S-KIT-CSCO	MDS 9148S Accessory Kit for Cisco	1
5.3	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	2
5.0	CS700-4F-48T-6400FS	Nimble CS700 FC Connectivity with 12 x 4TB HDDs and 4 x 1.6TB	1

		SSDs	
--	--	------	--

Summary

SmartStack delivers an infrastructure platform for Enterprise VDI deployments and cloud datacenters using Cisco and Fibre Channel-attached Nimble Storage CS700 array. SmartStack is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives.

About Authors

Jeff Nichols, Technical Marketing Engineer, Cisco UCS Solutions Engineering, Cisco Systems, Inc.

Jeff Nichols is a Cisco Unified Computing System architect, focusing on Virtual Desktop and Application solutions with extensive experience with VMware ESX/ESXi, XenDesktop, XenApp and Microsoft Remote Desktop Services. He has expert product knowledge in application, desktop and server virtualization across all three major hypervisor platforms and supporting infrastructures including but not limited to Windows Active Directory and Group Policies, User Profiles, DNS, DHCP and major storage platforms

Bharath Ram, Sr. Technical Marketing Engineer, Nimble Storage Inc.

Bharath Ram is an SME for Virtual Desktop and Application infrastructure. He has extensive knowledge and experience in designing and implementing large scale Citrix and VMWare VDI solutions for Healthcare and Insurance domains. As a Technical Marketing Engineer, he works on creating whitepapers and reference architecture for VDI based solutions for Nimble Storage.

Steve Sexton, Technical Marketing Engineer, Nimble Storage Inc.

Steve Sexton has over 15 years of experience in both Network and Storage industries. For the last five years he has been focused on Cisco UCS technologies and integration efforts. He holds **Bachelor's and Master's** degrees in Industrial Technology from Appalachian State University.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Mike Brennan, Technical Marketing Manager, Cisco Systems Inc.
- Bill Heffelfinger, Sr. Director Technical Marketing, Nimble Storage Inc.
- Arun Garg, Director, Solutions Product Management, Nimble Storage Inc.
- Matt Miller, Senior Product Marketing Manager, Nimble Storage Inc.

Appendix A – Cisco Nexus 9372 Switch Configuration

Switch A Configuration

```
SP-N9K-A# sho ru
!Command: show running-config

version 6.1(2)I3(3a)
switchname SP-N9K-A
vdc SP-N9K-A id 1
  allocate interface Ethernet1/1-54
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 512
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature telnet
cfs ipv4 distribute
cfs eth distribute
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp

username admin password 5 $1$9oM5JvS/$QwdAYtzhL1yKttz24UHmT/  role network-admin
no password strength-check
```

```
ip domain-lookup
copp profile strict
snmp-server user admin network-admin auth md5 0xb33f045ecc70424cfca355d6a205e2f8
priv 0xb33f045ecc70424cfca355d6a205e2f8 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vlan 1,70-80
vlan 70
    name IB-MGMT-VLAN
vlan 71
    name SP-Infra
vlan 72
    name VDI
vlan 73
    name Storage-1
vlan 74
    name ISCSI-BOOT-1
vlan 75
    name ISCSI-BOOT-2
vlan 76
    name vMotion
vlan 77
    name VLAN77
vlan 78
    name VLAN78
vlan 79
    name VLAN79
vlan 80
```

```
name Launcher80

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
route-map exit permit 10
service dhcp
ip dhcp relay

ipv6 dhcp relay
vrf context management
    ip route 0.0.0.0/0 10.29.164.1
vpc domain 10
    peer-switch
    role priority 10
    peer-keepalive destination 10.29.164.132 source 10.29.164.131
    delay restore 150
    peer-gateway
    auto-recovery

interface Vlan1
    no ip redirects
    no ipv6 redirects
    no shutdown

interface Vlan70
    no ip redirects
    ip address 10.10.70.4/24
    hsrp version 2
    hsrp 70
        preempt delay minimum 240
        priority 130
```

```
timers 1 3
ip 10.10.70.1
description IB Mgmt
no shutdown

interface Vlan71
no ip redirects
ip address 10.10.71.4/24
no ipv6 redirects
hsrp version 2
hsrp 71
preempt
priority 130
ip 10.10.71.1
ip dhcp relay address 10.10.71.21
description Infrastructure
no shutdown

interface Vlan72
no ip redirects
ip address 10.10.72.5/24
no ipv6 redirects
hsrp version 2
hsrp 72
preempt
ip 10.10.72.1
ip dhcp relay address 10.10.71.21
description VDI
no shutdown

interface Vlan73
no ip redirects
```

```
ip address 10.10.73.5/25
no ipv6 redirects
hsrp version 2
hsrp 73
    preempt
    ip 10.10.73.1
description Storage 1
no shutdown
```

```
interface Vlan74
    no ip redirects
    ip address 10.10.74.5/25
    no ipv6 redirects
    hsrp version 2
    hsrp 74
        preempt
        ip 10.10.74.1
description ISCSI Boot 1
no shutdown
```

```
interface Vlan75
    no ip redirects
    ip address 10.10.75.5/25
    no ipv6 redirects
    hsrp version 2
    hsrp 75
        preempt
        ip 10.10.75.1
description ISCSI Boot 2
no shutdown
```

```
interface Vlan76
```



```
no ip redirects
ip address 10.10.76.5/25
no ipv6 redirects
hsrp version 2
hsrp 76
    preempt
    ip 10.10.76.1
description vMotion
no shutdown

interface Vlan77
    no ip redirects
    ip address 10.3.0.2/19
    no ipv6 redirects
    hsrp version 2
    hsrp 77
        preempt
        priority 130
        ip 10.3.0.1
    ip dhcp relay address 10.10.71.21
    no shutdown

interface Vlan80
    no ip redirects
    ip address 10.10.80.4/20
    no ipv6 redirects
    hsrp version 2
    hsrp 80
        preempt
        ip 10.10.80.1
    ip dhcp relay address 10.10.71.21
    no shutdown
```

```
interface port-channel10
  description vPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 1,70-80
  spanning-tree port type network
  vpc peer-link
```

```
interface port-channel11
  description SP-FI-A
  switchport mode trunk
  switchport trunk allowed vlan 1,70-80
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11
```

```
interface port-channel12
  description SP-FI-B
  switchport mode trunk
  switchport trunk allowed vlan 1,70-80
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12
```

```
interface port-channel13
  description Launcher-A
  switchport mode trunk
  switchport trunk allowed vlan 1,70-80
  spanning-tree port type edge trunk
  mtu 9216
  vpc 13
```

```
interface port-channel14
  description Launcher-B
  switchport mode trunk
  switchport trunk allowed vlan 1,70-80
  spanning-tree port type edge trunk
  mtu 9216
  vpc 14
```

```
interface Ethernet1/1

  switchport access vlan 70
  spanning-tree port type edge
  speed 1000
```

```
interface Ethernet1/2
  switchport access vlan 71
  speed 1000
```

```
interface Ethernet1/3
```

```
interface Ethernet1/4
```

```
interface Ethernet1/5
```

```
interface Ethernet1/6
```

```
interface Ethernet1/7
```

```
interface Ethernet1/8
```

```
interface Ethernet1/9
```

```
interface Ethernet1/10
```

```
interface Ethernet1/11
```

```
interface Ethernet1/12
```

```
interface Ethernet1/13
```

```
interface Ethernet1/14
```

```
interface Ethernet1/15
```

```
interface Ethernet1/16
```

```
interface Ethernet1/17
```

```
interface Ethernet1/18
```

```
interface Ethernet1/19
```

```
description Launcher-A:19
```

```
shutdown
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,70-80
```

```
mtu 9216
```

```
channel-group 13 mode active
```

```
interface Ethernet1/20
```

```
description Launcher-B:20
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,70-80
```

```
mtu 9216
```

```
channel-group 14 mode active
```

```
interface Ethernet1/21
```

```
interface Ethernet1/22
```

```
interface Ethernet1/23
```

```
interface Ethernet1/24
```

```
interface Ethernet1/25
```

```
    switchport mode trunk
```

```
    switchport trunk allowed vlan 1,70-80
```

```
    mtu 9216
```

```
    channel-group 11 mode active
```

```
interface Ethernet1/26
```

```
    switchport mode trunk
```

```
    switchport trunk allowed vlan 1,70-80
```

```
    mtu 9216
```

```
    channel-group 12 mode active
```

```
interface Ethernet1/27
```

```
    description SP-FI-A:1/27
```

```
    switchport mode trunk
```

```
    switchport trunk allowed vlan 1,70-80
```

```
    mtu 9216
```

```
    channel-group 11 mode active
```

```
interface Ethernet1/28
```

```
    description SP-FI-B:1/28
```

```
    switchport mode trunk
```

```
    switchport trunk allowed vlan 1,70-80
```

```
mtu 9216
channel-group 12 mode active

interface Ethernet1/29
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    mtu 9216
    channel-group 13 mode active

interface Ethernet1/30
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    mtu 9216
    channel-group 14 mode active

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39
```

```
interface Ethernet1/40
```

```
interface Ethernet1/41
```

```
interface Ethernet1/42
```

```
interface Ethernet1/43
```

```
interface Ethernet1/44
```

```
interface Ethernet1/45
```

```
interface Ethernet1/46
```

```
interface Ethernet1/47
```

```
    description VPC Peer SP-N9K-B:1/47
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    channel-group 10 mode active
```

```
interface Ethernet1/48
```

```
    description VPC Peer SP-N9K-B:1/48
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    channel-group 10 mode active
```

```
interface Ethernet1/49
```

```
interface Ethernet1/50
```

```
interface Ethernet1/51

interface Ethernet1/52

interface Ethernet1/53

interface Ethernet1/54

interface mgmt0
    vrf member management
    ip address 10.29.164.131/24
line console
line vty
boot nxos bootflash:/n9000-dk9.6.1.2.I3.3a.bin

SP-N9K-A# exit
```

Switch B Configuration

```
SP-N9K-B# sho ru
!Command: show running-config

version 6.1(2)I3(3a)
switchname SP-N9K-B
vdc SP-N9K-B id 1
    allocate interface Ethernet1/1-54
    limit-resource vlan minimum 16 maximum 4094
    limit-resource vrf minimum 2 maximum 4096
    limit-resource port-channel minimum 0 maximum 512
    limit-resource u4route-mem minimum 248 maximum 248
```



```
limit-resource u6route-mem minimum 96 maximum 96
limit-resource m4route-mem minimum 58 maximum 58
limit-resource m6route-mem minimum 8 maximum 8

feature telnet
cfs eth distribute
feature udld
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp

username admin password 5 $1$6vbjAAN7$/ciP/uU95xAu3lce49DsO/ role network-admin
ip domain-lookup
copp profile strict
snmp-server user admin network-admin auth md5 0xc9cdb92eef56e64d61657edda45c102a
priv 0xc9cdb92eef56e64d61657edda45c102a localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vlan 1,70-80
vlan 70
    name IB-MGMT
vlan 71
    name SP-Infra
vlan 72
    name VDI
```

```
vlan 73
    name Storage-1
vlan 74
    name ISCSI-BOOT-1
vlan 75
    name ISCSI-BOOT-2
vlan 76
    name vMotion
vlan 77
    name VLAN77
vlan 78
    name VLAN78
vlan 79
    name VLAN79
vlan 80
    name Launcher80

spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
service dhcp
ip dhcp relay
ipv6 dhcp relay
vrf context management
    ip route 0.0.0.0/0 10.29.164.1

vpc domain 10
    role priority 20
    peer-keepalive destination 10.29.164.131 source 10.29.164.132
    delay restore 150
    peer-gateway
    auto-recovery
```

```
interface Vlan1
  no ip redirects
  no ipv6 redirects

interface Vlan70
  no ip redirects
  ip address 10.10.70.5/24
  no ipv6 redirects
  hsrp version 2
  hsrp 70
    preempt
    ip 10.10.70.1
  no shutdown

interface Vlan71
  no ip redirects
  ip address 10.10.71.6/24
  no ipv6 redirects
  hsrp version 2
  hsrp 71
    preempt
    ip 10.10.71.1
  no shutdown

interface Vlan72
  no ip redirects
  ip address 10.10.72.6/24
  no ipv6 redirects
  hsrp version 2
```

```
hsrp 72
  preempt
  ip 10.10.72.1
ip dhcp relay address 10.10.71.21
description VDI
no shutdown

interface Vlan77
  no ip redirects

  ip address 10.3.0.3/19
  no ipv6 redirects
  hsrp version 2
  hsrp 77
    preempt
    priority 130
    ip 10.3.0.1
  ip dhcp relay address 10.10.71.21
  no shutdown

interface Vlan80
  no ip redirects
  ip address 10.10.80.5/20
  no ipv6 redirects
  hsrp version 2
  hsrp 80
    preempt
    ip 10.10.80.1
  ip dhcp relay address 10.10.71.21
  no shutdown

interface port-channel10
```

```
description vPC peer-link

switchport mode trunk
switchport trunk allowed vlan 1,70-80
spanning-tree port type network
vpc peer-link

interface port-channel11
description SP-FI-A
switchport mode trunk
switchport trunk allowed vlan 1,70-80
spanning-tree port type edge trunk
mtu 9216
vpc 11

interface port-channel12
description SP-FI-B
switchport mode trunk
switchport trunk allowed vlan 1,70-80
spanning-tree port type edge trunk
mtu 9216
vpc 12

interface port-channel13
switchport mode trunk
switchport trunk allowed vlan 1,70-80
spanning-tree port type edge trunk
mtu 9216
vpc 13

interface port-channel14
switchport mode trunk
```

```
switchport trunk allowed vlan 1,70-80
spanning-tree port type edge trunk
mtu 9216
vpc 14
```

```
interface Ethernet1/1
switchport access vlan 71
speed 1000
```

```
interface Ethernet1/2
```

```
interface Ethernet1/3
```

```
interface Ethernet1/4
```

```
interface Ethernet1/5
```

```
interface Ethernet1/6
```

```
interface Ethernet1/7
```

```
interface Ethernet1/8
```

```
interface Ethernet1/9
```

```
interface Ethernet1/10
```

```
interface Ethernet1/11
```

```
interface Ethernet1/12
```

```
interface Ethernet1/13
```

```
interface Ethernet1/14
```

```
interface Ethernet1/15
```

```
interface Ethernet1/16
```

```
interface Ethernet1/17
```

```
interface Ethernet1/18
```

```
interface Ethernet1/19
```

```
description Launcher-A:19-b
shutdown
switchport mode trunk
switchport trunk allowed vlan 1,70-80
mtu 9216
channel-group 13 mode active
```

```
interface Ethernet1/20
```

```
description Launcher-B:20-b
shutdown
switchport mode trunk
switchport trunk allowed vlan 1,70-80
mtu 9216
channel-group 14 mode active
```

```
interface Ethernet1/21
```

```
interface Ethernet1/22
```

```
interface Ethernet1/23
```

```
interface Ethernet1/24
```

```
interface Ethernet1/25
```

```
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    mtu 9216
    channel-group 11 mode active
```

```
interface Ethernet1/26
```

```
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    mtu 9216
    channel-group 12 mode active
```

```
interface Ethernet1/27
```

```
    description SP-FI-B:1/27
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    mtu 9216
    channel-group 11 mode active
```

```
interface Ethernet1/28
```

```
    description SP-FI-A:1/28
    switchport mode trunk
    switchport trunk allowed vlan 1,70-80
    mtu 9216
    channel-group 12 mode active
```

```
interface Ethernet1/29
```

```
    switchport mode trunk
```



```
switchport trunk allowed vlan 1,70-80
mtu 9216
channel-group 13 mode active

interface Ethernet1/30
switchport mode trunk
switchport trunk allowed vlan 1,70-80
mtu 9216
channel-group 14 mode active

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41
```

```
interface Ethernet1/42
```

```
interface Ethernet1/43
```

```
interface Ethernet1/44
```

```
interface Ethernet1/45
```

```
interface Ethernet1/46
```

```
interface Ethernet1/47
```

```
description VPC Peer SP-N9K-A:1/47
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,70-80
```

```
channel-group 10 mode active
```

```
interface Ethernet1/48
```

```
description VPC Peer SP-N9K-A:1/48
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,70-80
```

```
channel-group 10 mode active
```

```
interface Ethernet1/49
```

```
interface Ethernet1/50
```

```
interface Ethernet1/51
```

```
interface Ethernet1/52
```

```
interface Ethernet1/53
```

```
interface Ethernet1/54
```

```
interface mgmt0
```

```
    vrf member management
```

```
    ip address 10.29.164.132/24
```

```
line console
```

```
line vty
```

```
boot nxos bootflash:/n9000-dk9.6.1.2.I3.3a.bin
```

```
SP-N9K-B# exit
```

Appendix B – Cisco MDS 9148 Switch Configuration

MDS- A Switch Configuration

```

MDS-A# sho ru
!Command: show running-config
!Time: Mon Mar 21 01:32:00 2016

version 6.2(9a)
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
feature fcsp
role name default-role
    description This is a system defined role and applies to all users.
    rule 5 permit show feature environment
    rule 4 permit show feature hardware
    rule 3 permit show feature module
    rule 2 permit show feature snmp
    rule 1 permit show feature system
username admin password 5 $1$loX7vizP$00IbhSFcpX6WufBmOMKB.1 role network-admin
ip domain-lookup
ip host MDS-A 10.29.164.64
aaa group server radius radius
snmp-server user admin network-admin auth md5 0x6c81eb7167a2e69497a60698ca3957da
    priv 0x6c81eb7167a2e69497a60698ca3957da localizedkey
snmp-server host 10.155.160.192 traps version 2c public udp-port 1163
snmp-server host 10.29.164.130 traps version 2c public udp-port 1163
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR

```

```

rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
vsan database

```

```

    vsan 3 name "VSAN3-A"

```

```

device-alias database

```

```

    device-alias name SP-Infra1-fc0 pwwn 20:00:00:25:b5:00:00:2f
    device-alias name SP-Infra2-fc0 pwwn 20:00:00:25:b5:00:00:0f
    device-alias name SP-VDI-01-fc0 pwwn 20:00:00:25:b5:00:00:2c

```

```

device-alias commit

```

```

fcdomain fcid database

```

```

    vsan 3 wwn 20:1f:00:2a:6a:d3:df:80 fcid 0x300000 dynamic
    vsan 3 wwn 20:20:00:2a:6a:d3:df:80 fcid 0x300100 dynamic
    vsan 3 wwn 52:4a:93:72:0d:21:6b:01 fcid 0x300200 dynamic
    vsan 3 wwn 52:4a:93:72:0d:21:6b:11 fcid 0x300300 dynamic
    vsan 3 wwn 52:4a:93:72:0d:21:6b:10 fcid 0x300400 dynamic
    vsan 1 wwn 52:4a:93:72:0d:21:6b:11 fcid 0x290000 dynamic
    vsan 1 wwn 52:4a:93:72:0d:21:6b:10 fcid 0x290100 dynamic
    vsan 1 wwn 20:20:00:2a:6a:d3:df:80 fcid 0x290200 dynamic
    vsan 1 wwn 24:01:00:2a:6a:d3:df:80 fcid 0x290300 dynamic
    vsan 3 wwn 24:01:00:2a:6a:d3:df:80 fcid 0x300500 dynamic
    vsan 3 wwn 56:c9:ce:90:0d:e8:24:01 fcid 0x300600 dynamic
    vsan 3 wwn 56:c9:ce:90:0d:e8:24:05 fcid 0x300700 dynamic
    vsan 1 wwn 52:4a:93:72:0d:21:6b:00 fcid 0x290400 dynamic

```

```

interface port-channel1

```

```

    channel mode active
    switchport rate-mode dedicated

```

```

vsan database

```

```

    vsan 3 interface port-channel1

```

```
vsan 3 interface fc1/9
vsan 3 interface fc1/10
switchname MDS-A
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9a.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.9a.bin
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
```

```
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12

!Active Zone Database Section for vsan 3
zone name SP-Infra1-fc0 vsan 3
    member pwwn 20:00:00:25:b5:00:00:2f
!
    [SP-Infra1-fc0]
    member pwwn 56:c9:ce:90:0d:e8:24:01
```

```
member pwnn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-Infra2-fc0 vsan 3
```

```
member pwnn 20:00:00:25:b5:00:00:0f
```

```
! [SP-Infra2-fc0]
```

```
member pwnn 56:c9:ce:90:0d:e8:24:01
```

```
member pwnn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-01-fc0 vsan 3
```

```
member pwnn 20:00:00:25:b5:00:00:2c
```

```
! [SP-VDI-01-fc0]
```

```
member pwnn 56:c9:ce:90:0d:e8:24:01
```

```
member pwnn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-02-fc0 vsan 3
```

```
member pwnn 20:00:00:25:b5:00:00:0c
```

```
member pwnn 56:c9:ce:90:0d:e8:24:05
```

```
member pwnn 56:c9:ce:90:0d:e8:24:01
```

```
zone name SP-VDI-03-fc0 vsan 3
```

```
member pwnn 20:00:00:25:b5:00:00:4b
```

```
member pwnn 56:c9:ce:90:0d:e8:24:01
```

```
member pwnn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-04-fc0 vsan 3
```

```
member pwnn 20:00:00:25:b5:00:00:2b
```

```
member pwnn 56:c9:ce:90:0d:e8:24:01
```

```
member pwnn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-05-fc0 vsan 3
```

```
member pwnn 20:00:00:25:b5:00:00:0b
```

```
member pwnn 56:c9:ce:90:0d:e8:24:01
```



```
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-06-fc0 vsan 3
```

```
member pwwn 20:00:00:25:b5:00:00:4a
```

```
member pwwn 56:c9:ce:90:0d:e8:24:01
```

```
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-07-fc0 vsan 3
```

```
member pwwn 20:00:00:25:b5:00:00:2a
```

```
member pwwn 56:c9:ce:90:0d:e8:24:01
```

```
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-08-fc0 vsan 3
```

```
member pwwn 20:00:00:25:b5:00:00:0a
```

```
member pwwn 56:c9:ce:90:0d:e8:24:01
```

```
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-09-fc0 vsan 3
```

```
member pwwn 20:00:00:25:b5:00:00:59
```

```
member pwwn 56:c9:ce:90:0d:e8:24:01
```

```
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-10-fc0 vsan 3
```

```
member pwwn 20:00:00:25:b5:00:00:5c
```

```
member pwwn 56:c9:ce:90:0d:e8:24:01
```

```
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-11-fc0 vsan 3
```

```
member pwwn 20:00:00:25:b5:00:00:29
```

```
member pwwn 56:c9:ce:90:0d:e8:24:01
```

```
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-12-fc0 vsan 3
    member pwn 20:00:00:25:b5:00:00:09
    member pwn 56:c9:ce:90:0d:e8:24:01
    member pwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-13-fc0 vsan 3
    member pwn 20:00:00:25:b5:00:00:48
    member pwn 56:c9:ce:90:0d:e8:24:01
    member pwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-14-fc0 vsan 3
    member pwn 20:00:00:25:b5:00:00:18
    member pwn 56:c9:ce:90:0d:e8:24:01
    member pwn 56:c9:ce:90:0d:e8:24:05
```

```
zoneset name SP-Infra-A vsan 3
    member SP-Infra1-fc0
    member SP-Infra2-fc0
    member SP-VDI-01-fc0
    member SP-VDI-02-fc0
    member SP-VDI-03-fc0
    member SP-VDI-04-fc0
    member SP-VDI-05-fc0
    member SP-VDI-06-fc0
    member SP-VDI-07-fc0
    member SP-VDI-08-fc0
    member SP-VDI-09-fc0
    member SP-VDI-10-fc0
    member SP-VDI-11-fc0
    member SP-VDI-12-fc0
    member SP-VDI-13-fc0
    member SP-VDI-14-fc0
```

```
zoneset activate name SP-Infra-A vsan 3
do clear zone database vsan 3
!Full Zone Database Section for vsan 3
zone name SP-Infra1-fc0 vsan 3
    member pwwn 20:00:00:25:b5:00:00:2f
!
    [SP-Infra1-fc0]
    member pwwn 56:c9:ce:90:0d:e8:24:01
    member pwwn 56:c9:ce:90:0d:e8:24:05

zone name SP-Infra2-fc0 vsan 3
    member pwwn 20:00:00:25:b5:00:00:0f
!
    [SP-Infra2-fc0]
    member pwwn 56:c9:ce:90:0d:e8:24:01
    member pwwn 56:c9:ce:90:0d:e8:24:05

zone name SP-VDI-01-fc0 vsan 3
    member pwwn 20:00:00:25:b5:00:00:2c
!
    [SP-VDI-01-fc0]
    member pwwn 56:c9:ce:90:0d:e8:24:01
    member pwwn 56:c9:ce:90:0d:e8:24:05

zone name SP-VDI-02-fc0 vsan 3
    member pwwn 20:00:00:25:b5:00:00:0c
    member pwwn 56:c9:ce:90:0d:e8:24:05
    member pwwn 56:c9:ce:90:0d:e8:24:01

zone name SP-VDI-03-fc0 vsan 3
    member pwwn 20:00:00:25:b5:00:00:4b
    member pwwn 56:c9:ce:90:0d:e8:24:01
    member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-04-fc0 vsan 3
  member pwwn 20:00:00:25:b5:00:00:2b
  member pwwn 56:c9:ce:90:0d:e8:24:01
  member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-05-fc0 vsan 3
  member pwwn 20:00:00:25:b5:00:00:0b
  member pwwn 56:c9:ce:90:0d:e8:24:01
  member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-06-fc0 vsan 3
  member pwwn 20:00:00:25:b5:00:00:4a

  member pwwn 56:c9:ce:90:0d:e8:24:01
  member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-07-fc0 vsan 3
  member pwwn 20:00:00:25:b5:00:00:2a
  member pwwn 56:c9:ce:90:0d:e8:24:01
  member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-08-fc0 vsan 3
  member pwwn 20:00:00:25:b5:00:00:0a
  member pwwn 56:c9:ce:90:0d:e8:24:01
  member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-09-fc0 vsan 3
  member pwwn 20:00:00:25:b5:00:00:59
  member pwwn 56:c9:ce:90:0d:e8:24:01
  member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-10-fc0 vsan 3
```

```
member pwwn 20:00:00:25:b5:00:00:5c
member pwwn 56:c9:ce:90:0d:e8:24:01
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-11-fc0 vsan 3
member pwwn 20:00:00:25:b5:00:00:29
member pwwn 56:c9:ce:90:0d:e8:24:01
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-12-fc0 vsan 3
member pwwn 20:00:00:25:b5:00:00:09
member pwwn 56:c9:ce:90:0d:e8:24:01
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-13-fc0 vsan 3
member pwwn 20:00:00:25:b5:00:00:48
member pwwn 56:c9:ce:90:0d:e8:24:01
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zone name SP-VDI-14-fc0 vsan 3
member pwwn 20:00:00:25:b5:00:00:18
member pwwn 56:c9:ce:90:0d:e8:24:01
member pwwn 56:c9:ce:90:0d:e8:24:05
```

```
zoneset name SP-Nimble-A vsan 3
```

```
zoneset name SP-Infra-A vsan 3
```

```
member SP-Infra1-fc0
member SP-Infra2-fc0
member SP-VDI-01-fc0
member SP-VDI-02-fc0
member SP-VDI-03-fc0
```

```
member SP-VDI-04-fc0
member SP-VDI-05-fc0
member SP-VDI-06-fc0
member SP-VDI-07-fc0
member SP-VDI-08-fc0
member SP-VDI-09-fc0
member SP-VDI-10-fc0
member SP-VDI-11-fc0
member SP-VDI-12-fc0
member SP-VDI-13-fc0
member SP-VDI-14-fc0
```

```
interface fc1/1
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/2
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/3
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/4
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/5
  port-license acquire
```

```
interface fc1/6
  port-license acquire
```

```
interface fc1/7
  no port-license
  no shutdown
```

```
interface fc1/8
  no port-license
  no shutdown
```

```
interface fc1/9
  port-license acquire
  no shutdown
```

```
interface fc1/10
  port-license acquire
  no shutdown
```

```
interface fc1/11
  port-license acquire
  channel-group 1 force
  no shutdown
```

```
interface fc1/12
  port-license acquire
  channel-group 1 force
  no shutdown
```

```
interface fc1/13
```

```
interface fc1/14
```

```
interface fc1/15
```

```
interface fc1/16
```

```
interface fc1/17
```

```
interface fc1/18
```

```
interface fc1/19
```

```
interface fc1/20
```

```
interface fc1/21
```

```
interface fc1/22
```

```
interface fc1/23
```

```
interface fc1/24
```

```
interface fc1/25
```

```
interface fc1/26
```

```
interface fc1/27
```

```
interface fc1/28
```



```
interface fc1/29
```

```
interface fc1/30
```

```
interface fc1/31
```

```
interface fc1/32
```

```
interface fc1/33
```

```
interface fc1/34
```

```
interface fc1/35
```

```
interface fc1/36
```

```
interface fc1/37
```

```
interface fc1/38
```

```
interface fc1/39
```

```
interface fc1/40
```

```
interface fc1/41
```

```
interface fc1/42
```

```
interface fc1/43
```

```
interface fc1/44
```

```
interface fc1/45
```

```
interface fc1/46
```

```
interface fc1/47
```

```
interface fc1/48
```

```
interface mgmt0
```

```
ip address 10.29.164.64 255.255.255.0
```

```
ip default-gateway 10.29.164.1
```

```
MDS-A# exit
```

MDS- B Switch Configuration

```
MDS-B# sho ru
```

```
!Command: show running-config
```

```
version 6.2(9a)
```

```
power redundancy-mode redundant
```

```
feature npiv
```

```
feature fport-channel-trunk
```

```
role name default-role
```

```
description This is a system defined role and applies to all users.
```

```
rule 5 permit show feature environment
```

```
rule 4 permit show feature hardware
```

```
rule 3 permit show feature module
```

```
rule 2 permit show feature snmp
```

```

rule 1 permit show feature system
username admin password 5 $1$dcmFCg/p$0ZC5U6uhI65oOePpHfAzn0 role network-admin
no password strength-check
ip domain-lookup
ip host MDS-B 10.29.164.128
aaa group server radius radius
snmp-server user admin network-admin auth md5 0xc9e1af5dbb0bbac72253a1bef037bbbe
priv 0xc9e1af5dbb0bbac72253a1bef037bbbe localizedkey
snmp-server host 10.155.160.192 traps version 2c public udp-port 1164

```

```

snmp-server host 10.29.164.130 traps version 2c public udp-port 1164
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
vsan database
vsan 4 name "SP-FAB-B"

```

```

device-alias database
device-alias name SP-Infra1-fc1 pwn 20:00:00:25:b5:00:00:3f
device-alias name SP-Infra2-fc1 pwn 20:00:00:25:b5:00:00:1f

```

```

device-alias commit

```

```

fcdomain fcid database
vsan 4 wwn 24:01:00:2a:6a:d9:84:c0 fcid 0x5b0000 dynamic
vsan 4 wwn 56:c9:ce:90:0d:e8:24:02 fcid 0x5b0100 dynamic
vsan 4 wwn 56:c9:ce:90:0d:e8:24:06 fcid 0x5b0200 dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:5a fcid 0x5b0001 dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:1b fcid 0x5b0002 dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:19 fcid 0x5b0003 dynamic

```

```

vsan 4 wwn 20:00:00:25:b5:00:00:1a fcid 0x5b0004 dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:3a fcid 0x5b0005 dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:1f fcid 0x5b0006 dynamic
!
[SP-Infra2-fc1]
vsan 4 wwn 20:00:00:25:b5:00:00:58 fcid 0x5b0007 dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:3c fcid 0x5b0008 dynamic

vsan 4 wwn 20:00:00:25:b5:00:00:3f fcid 0x5b0009 dynamic
!
[SP-Infra1-fc1]
vsan 4 wwn 20:00:00:25:b5:00:00:5b fcid 0x5b000a dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:3b fcid 0x5b000b dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:38 fcid 0x5b000c dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:1c fcid 0x5b000d dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:49 fcid 0x5b000e dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:08 fcid 0x5b000f dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:39 fcid 0x5b0010 dynamic
vsan 4 wwn 20:00:00:25:b5:00:00:37 fcid 0x5b0011 dynamic

interface port-channel1

channel mode active

switchport rate-mode dedicated

vsan database
vsan 4 interface port-channel1
vsan 4 interface fc1/9
vsan 4 interface fc1/10

switchname MDS-B

line console

line vty

boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.9a.bin

```

```
boot system bootflash:/m9100-s5ek9-mz.6.2.9a.bin

interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
```

```
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface fc1/1
interface fc1/2
interface fc1/11
interface fc1/12
```

```
!Active Zone Database Section for vsan 4
zone name SP-VDI-01-fc1 vsan 4
```

```
member pwn 20:00:00:25:b5:00:00:3c
member pwn 56:c9:ce:90:0d:e8:24:02
member pwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-02-fc1 vsan 4
member pwn 20:00:00:25:b5:00:00:1c
member pwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-03-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:5b
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
zone name SP-VDI-04-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:3b
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-05-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:1b
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-06-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:5a
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-07-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:3a
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-08-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:1a
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-09-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:49
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-10-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:39

    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-11-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:19
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-12-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:58
    member pwwn 56:c9:ce:90:0d:e8:24:06
    member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
zone name SP-VDI-13-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:38
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-14-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:08
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-Infra1-fc1 vsan 4
```



```

        member pwnn 20:00:00:25:b5:00:00:3f
!
        [SP-Infra1-fc1]
        member pwnn 56:c9:ce:90:0d:e8:24:02
        member pwnn 56:c9:ce:90:0d:e8:24:06

```

```

zone name SP-Infra2-fc1 vsan 4
        member pwnn 20:00:00:25:b5:00:00:1f
!
        [SP-Infra2-fc1]
        member pwnn 56:c9:ce:90:0d:e8:24:02
        member pwnn 56:c9:ce:90:0d:e8:24:06

```

```

zoneset name SP-Infra-B vsan 4

```

```

        member SP-VDI-01-fc1
        member SP-VDI-02-fc1
        member SP-VDI-03-fc1
        member SP-VDI-04-fc1
        member SP-VDI-05-fc1
        member SP-VDI-06-fc1
        member SP-VDI-07-fc1
        member SP-VDI-08-fc1
        member SP-VDI-09-fc1
        member SP-VDI-10-fc1
        member SP-VDI-11-fc1
        member SP-VDI-12-fc1
        member SP-VDI-13-fc1
        member SP-VDI-14-fc1
        member SP-Infra1-fc1
        member SP-Infra2-fc1

```

```

zoneset activate name SP-Infra-B vsan 4

```

```

do clear zone database vsan 4

```

```

!Full Zone Database Section for vsan 4

```

```
zone name SP-VDI-01-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:3c
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-02-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:1c
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-03-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:5b
    member pwwn 56:c9:ce:90:0d:e8:24:06
    member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
zone name SP-VDI-04-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:3b
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-05-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:1b
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-06-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:5a
    member pwwn 56:c9:ce:90:0d:e8:24:02
    member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-07-fc1 vsan 4
    member pwwn 20:00:00:25:b5:00:00:3a
```

```
member pwnn 56:c9:ce:90:0d:e8:24:02
member pwnn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-08-fc1 vsan 4
```

```
member pwnn 20:00:00:25:b5:00:00:1a
member pwnn 56:c9:ce:90:0d:e8:24:02
member pwnn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-09-fc1 vsan 4
```

```
member pwnn 20:00:00:25:b5:00:00:49
member pwnn 56:c9:ce:90:0d:e8:24:02
member pwnn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-10-fc1 vsan 4
```

```
member pwnn 20:00:00:25:b5:00:00:39
member pwnn 56:c9:ce:90:0d:e8:24:02
member pwnn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-11-fc1 vsan 4
```

```
member pwnn 20:00:00:25:b5:00:00:19
member pwnn 56:c9:ce:90:0d:e8:24:02
member pwnn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-12-fc1 vsan 4
```

```
member pwnn 20:00:00:25:b5:00:00:58
member pwnn 56:c9:ce:90:0d:e8:24:06
member pwnn 56:c9:ce:90:0d:e8:24:02
```

```
zone name SP-VDI-13-fc1 vsan 4
```

```
member pwnn 20:00:00:25:b5:00:00:38
member pwnn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-VDI-14-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:08
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-Infra1-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:3f
```

```
! [SP-Infra1-fc1]
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zone name SP-Infra2-fc1 vsan 4
```

```
member pwwn 20:00:00:25:b5:00:00:1f
```

```
! [SP-Infra2-fc1]
```

```
member pwwn 56:c9:ce:90:0d:e8:24:02
```

```
member pwwn 56:c9:ce:90:0d:e8:24:06
```

```
zoneset name SP-Infra-B vsan 4
```

```
member SP-VDI-01-fc1
```

```
member SP-VDI-02-fc1
```

```
member SP-VDI-03-fc1
```

```
member SP-VDI-04-fc1
```

```
member SP-VDI-05-fc1
```

```
member SP-VDI-06-fc1
```

```
member SP-VDI-07-fc1
```

```
member SP-VDI-08-fc1
```

```
member SP-VDI-09-fc1
```

```
member SP-VDI-10-fc1
```

```
member SP-VDI-11-fc1
```

```
member SP-VDI-12-fc1
member SP-VDI-13-fc1
member SP-VDI-14-fc1
member SP-Infra1-fc1
member SP-Infra2-fc1
```

```
interface fc1/1
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/2
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/3
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/4
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/5
  port-license acquire
```

```
interface fc1/6
  port-license acquire
```

```
interface fc1/7
  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/8

  switchport trunk mode off
  port-license acquire
  no shutdown
```

```
interface fc1/9
  port-license acquire
  no shutdown
```

```
interface fc1/10
  port-license acquire
  no shutdown
```

```
interface fc1/11
  port-license acquire
  channel-group 1 force
  no shutdown
```

```
interface fc1/12
  port-license acquire
  channel-group 1 force
  no shutdown
```

```
interface fc1/13
```

```
interface fc1/14
```

```
interface fc1/15
```

```
interface fc1/16
```

```
interface fc1/17
```

```
interface fc1/18
```

```
interface fc1/19
```

```
interface fc1/20
```

```
interface fc1/21
```

```
interface fc1/22
```

```
interface fc1/23
```

```
interface fc1/24
```

```
interface fc1/25
```

```
interface fc1/26
```

```
interface fc1/27
```

```
interface fc1/28
```

```
interface fc1/29
```

```
interface fc1/30
```

```
interface fc1/31
```

```
interface fc1/32
```

```
interface fc1/33
```

```
interface fc1/34
```

```
interface fc1/35
```

```
interface fc1/36
```

```
interface fc1/37
```

```
interface fc1/38
```

```
interface fc1/39
```

```
interface fc1/40
```

```
interface fc1/41
```

```
interface fc1/42
```

```
interface fc1/43
```

```
interface fc1/44
```

```
interface fc1/45
```



```
interface fc1/46
```

```
interface fc1/47
```

```
interface fc1/48
```

```
interface mgmt0
```

```
    ip address 10.29.164.128 255.255.255.0
```

```
ip default-gateway 10.29.164.1
```

```
MDS-B#    exit
```