



REFERENCE ARCHITECTURE FOR 500-SEAT AND 1000-SEAT VIRTUAL DESKTOP INFRASTRUCTURE

CITRIX XENDESKTOP 7 BUILT ON CISCO UCS B200-M3 BLADE SERVERS WITH EMC VNXE AND MICROSOFT HYPER-V 2012

December 2013



Table of Contents

1. Overview	9
1.1. About this Document.....	9
2. Solution Component Benefits	9
2.1. Benefits of Cisco Unified Computing System	9
2.1.1. Benefits of Cisco Nexus 5548UP	10
2.2. Benefits of EMC VNX Family of Storage Controllers	11
2.2.1. The EMC VNX Family	11
2.2.2. VNXe Series – Simple, Efficient, Affordable	11
2.2.3. VNX Series - Simple, Efficient, Powerful.....	12
2.3. Benefits of Microsoft Windows Server 2012 with Hyper-V	12
2.4. Benefits of Citrix XenDesktop 7	13
2.5. Audience	13
3. Summary of Main Findings	14
4. Architecture.....	15
4.1. Hardware Deployed	15
4.2. Logical Architecture.....	16
4.3. Software Revisions	17
4.4. Configuration Guidelines	18
4.4.1. VLAN.....	18
4.4.2. Hyper-V Clustering	19
5. Infrastructure Components	20
5.1. Cisco Unified Computer System (UCS)	20
5.1.1. Cisco Unified Computing Components	20
5.1.2. Cisco Fabric Interconnects.....	22
5.1.3. Cisco IO Modules (Fabric Extenders)	22
5.1.4. Cisco UCS Chassis	23
5.1.5. Cisco UCS Manager	23
5.1.6. Cisco UCS B200 M3Blade Servers.....	23
5.1.7. Cisco Virtual Interface Card (VIC) Converged Network Adapter	24
5.2. EMC VNXe3300.....	25
5.2.1. Advantages and Value Proposition.....	25

5.2.2.	Software suites available.....	26
5.2.3.	Software packs available.....	26
5.3.	Microsoft Technologies.....	26
5.3.1.	Windows Server 2012.....	26
5.3.2.	Failover Clustering	27
5.3.3.	Clustered Shared Volumes.....	27
5.3.4.	Networking Support	27
5.3.5.	Hyper-V.....	28
5.3.6.	Hyper-V Server 2012	28
5.3.7.	SQL Server 2012.....	28
5.3.8.	System Center Virtual Machine Manager 2012 SP1	29
5.4.	Citrix XenDesktop 7	29
5.4.1.	Enhancements in This Release	29
5.4.2.	FlexCast Technology	30
5.4.3.	High-Definition User Experience (HDX) Technology	32
5.4.4.	Citrix XenDesktop 7 Desktop and Application Services	32
5.4.5.	Provisioning Services 7	33
6.	Solution Architecture	34
6.1.	Citrix Design Fundamentals	36
6.1.1.	Citrix Hosted Shared Desktop Design Fundamentals.....	37
6.2.	EMC Storage Architecture Design.....	42
6.2.1.	High Availability.....	42
6.2.2.	Data Protection	43
6.3.	Solution Validation	43
6.3.1.	Configuration Topology for Scalable Citrix XenDesktop 7 Hybrid Virtual Desktop Infrastructure on Cisco Unified Computing System and EMC Storage.....	43
6.4.	Configuration Topology for Citrix XenDesktop 7 on Cisco Unified Computing System with VNXe Storage.....	45
6.5.	Cisco UCS Configuration.....	46
6.6.	Base Cisco UCS Configuration.....	47
6.6.1.	Firmware Update.....	47
6.6.2.	Acknowledge the Chassis	48
6.6.3.	Server Port Configuration.....	49
6.6.4.	Uplink Port Configuration.....	50

6.6.5.	VNXe Appliance Port Configuration.....	53
6.6.6.	KVM IP Address Pool	54
6.6.7.	MAC Address Pool.....	56
6.6.8.	UUID Suffix Pool	56
6.6.9.	IQN Pool	57
6.6.10.	Server Pool and Related Policies (Optional)	57
6.6.11.	Local Disk Configuration.....	59
6.6.12.	BIOS Policy.....	60
6.6.13.	Power Control Policy	61
6.6.14.	Scrub Policy.....	62
6.6.15.	Host Firmware Package.....	63
6.6.16.	QOS Policies	64
6.6.17.	Steps to Enable QOS on the Cisco Unified Computing System	66
6.6.18.	iSCSI Adapter Policy.....	68
6.6.19.	VLANs and vNIC Templates.....	69
6.6.20.	iSCSI Boot Policy	73
6.6.21.	Service Profile Template	74
6.6.22.	Create Service Profiles	80
6.6.23.	Configure iSCSI Boot LUNS for Each Service Profile.....	81
6.7.	LAN Configuration.....	84
6.8.	SAN Configuration.....	84
6.8.1.	Boot from SAN benefits	84
6.9.	EMC VNXe Storage Configuration	85
6.9.1.	iSCSI Host Configuration.....	86
6.9.2.	iSCSI Server Configuration	89
6.9.3.	iSCSI Storage Creation.....	92
6.9.4.	CIFS Server Creation	96
6.9.5.	CIFS Share Storage Creation – Profile Share	98
6.9.6.	CIFS Share Storage Creation – vDisk Share	103
6.10.	Installing and Configuring Microsoft Server 2012.....	106
6.10.1.	Infrastructure Servers.....	106
6.10.2.	VDI Hosts.....	120
6.10.3.	Local Configuration Tasks After Install.....	123

6.11.	Installing and Configuring SQL Server 2012 SP1	123
6.11.1.	Pass-Through Storage Configuration	124
6.11.2.	SQL Server 2012 – Installation Pre-requisites.....	129
6.11.3.	SQL Server 2012 Installation	129
6.11.4.	AlwaysOn Application Group	139
6.11.5.	Log File Management	147
6.12.	Installing and Configuring System Center 2012 Virtual Machine Manager	148
6.13.	Installing and Configuring Cisco Nexus 1000V Virtual Switch Manager and Virtual Ethernet Modules	154
6.13.1.	Installing Cisco Nexus 1000V Virtual Switch.....	155
6.13.2.	Installing and Configuring the Virtual Supervisor Modules (VSM)	159
6.13.3.	Cisco Nexus 1000V for Hyper-V Configuration	173
6.14.	Installing and Configuring Citrix XenDesktop	187
6.14.1.	Installing Provisioning Services	191
6.14.2.	Installation of the Second PVS Server	212
6.14.3.	Installing the Delivery Controller	219
6.14.4.	XenDesktop Controller Configuration.....	228
6.14.5.	Additional XenDesktop Controller Configuration.....	233
6.14.6.	Configure VM Hosting, Connections, and Resources.....	235
6.14.7.	Installing and Configuring StoreFront	239
7.	Desktop Delivery Infrastructure - Citrix	245
7.1.	Overview of desktop delivery.....	245
7.1.1.	PVS vDisk Image Management	246
7.1.2.	Overview – Golden Image Creation.....	246
7.2.	Installing the XenDesktop 7 Virtual Desktop Agent.....	247
7.3.	Citrix User Profile Management Servers – CITRIX.....	251
7.3.1.	Install and Configuration of User Profile Manager.....	251
7.4.	Microsoft Windows 7 and Windows Server 2012 Golden Image Creation	254
7.4.1.	Microsoft Windows 7 and Windows Server 2012 OS Configurations	254
7.4.2.	Installing the PVS Target Device Software.....	255
7.4.3.	Running the PVS Imaging Wizard.....	257
7.4.4.	Installation of Login VSI Software	261
7.4.5.	Optimization for PVS and XenDesktop 7	262

7.4.6.	Conversion to PVS vDisk.....	263
7.4.7.	Write-Cache Drive Sizing and Placement.....	263
7.5.	Citrix Provisioning Services	264
7.5.1.	Creating the Virtual Machine Manager Templates	264
7.5.2.	Process to Create Virtual Desktops using PVS Wizard	271
8.	Test Setup and Configurations	281
8.1.	Cisco UCS Test Configuration for Single Blade Scalability of Hosted Shared Desktops	282
8.2.	Cisco UCS Test Configuration for Single Blade Scalability of Hosted Virtual Machines	283
8.3.	Cisco UCS Test Configuration for Single Blade Scalability for a Mixed Hosted Shared and Hosted Virtual Workload.....	285
8.4.	Cisco UCS Test Configuration for a Single-Chassis 500-User Configuration	286
8.5.	Cisco UCS Test Configuration for a Two-Chasses 1000-User Configuration	288
8.6.	Test Methodology and Success Criteria	289
8.6.1.	Load Generation.....	289
8.6.2.	User Workload Simulation – Login VSI	289
8.6.3.	Testing Procedure	291
8.6.4.	Success Criteria.....	293
9.	Login VSI Test Result	296
9.1.	Cisco UCS B200-M3 Single-Server Scalability Results for Hosted Shared Desktops	296
9.2.	Cisco UCS B200-M3 Single-Server Scalability Results for Hosted Virtual Desktops	300
9.3.	Cisco UCS B200-M3 Recommended Single-Server Mixed Desktop Workload	303
9.4.	Cisco UCS 500-User Scalability Results.....	305
9.4.1.	EMC VNXe Performance.....	306
9.4.2.	VDI Host	307
9.4.3.	XenDesktop Controllers	311
9.4.4.	Provisioning Services	312
9.4.5.	SQL Servers	313
9.4.6.	System Center Virtual Machine Manager and StoreFront.....	314
9.4.1.	Active Directory.....	315
9.5.	Cisco UCS 1000-User Scalability Results	317
9.5.1.	EMC VNXe Performance.....	317
9.5.2.	VDI Host	319
9.5.3.	XenDesktop Controllers	323

9.5.4.	Provisioning Services	324
9.5.5.	SQL Servers	325
9.5.6.	System Center Virtual Machine Manager and StoreFront	326
9.5.7.	Active Directory	327
10.	Scalability Considerations and Guidelines	329
10.1.	Cisco UCS Configuration	329
10.2.	EMC VNXe Storage Configuration	329
10.2.1.	Capacity planning	329
10.2.2.	Performance Planning	330
10.2.3.	Scalability Planning	330
10.3.	Microsoft Windows Server 2012 with Hyper-V 2012 Configuration	330
10.3.1.	Virtual Machines	330
10.3.2.	Server Running Hyper-V	331
10.3.4.	Failover Clusters and Hyper-V	332
10.4.	Citrix XenDesktop 7 Configuration – Citrix	333
11.	Other Considerations	333
11.1.	Power Outages and Boot Sequence	333
11.1.1.	Recommended Boot Sequence	334
11.2.	Microsoft Cluster Maintenance	334
11.3.	SQL Server AlwaysOn Groups	334
12.	References	335
12.1.	Cisco Reference Documents – Cisco	335
12.2.	EMC Reference Documents – EMC	335
12.3.	Microsoft Reference Documents	336
12.4.	Citrix Reference Documents	337
13.	Appendix	338
13.1.	Performance Charts for Scalability Tests	338
13.2.	Sample Cisco Nexus 6248-UP Configurations	352
13.2.1.	6248UP – A	377
13.3.	Sample Cisco Nexus 1000V Configuration	401
13.3.1.	N1000V-1	406
13.4.	Sample PowerShell Scripts	411
13.4.1.	Update Virtual Machines created by XenDesktop Wizard	411

13.4.2.	Enable Dynamic Memory	413
13.4.3.	Disable Dynamic Memory.....	414
13.4.4.	Query the XenDesktop Database Connection Strings	415
13.4.5.	Test the XenDesktop Database Connection String	416
13.4.6.	Change the XenDesktop Database Connection String	416
14.	Acknowledgements	417

1. Overview

1.1. About this Document

This document describes the reference architecture for a 500-seat and 1000-seat virtual desktop infrastructure using Citrix XenDesktop 7 built on Cisco UCS B200-M3 Blade Servers with EMC VNXe3300 and Microsoft Hyper-V 2012.

Industry trends indicate a vast data center transformation toward shared infrastructures. Enterprise customers are moving away from silos of information and toward shared infrastructures, to virtualized environments, and eventually to the cloud to increase agility and reduce costs.

This document provides the architecture and design of a virtual desktop infrastructure that can grow from 500 users to 1000 users. The infrastructure is 100% virtualized on Microsoft Hyper-V Server 2012 with third-generation Cisco UCS B-Series B200 M3 Blade Servers iSCSI booting from an EMC VNXe3300 storage array.

The virtual desktops are powered using Citrix Provisioning Server 7 and Citrix XenDesktop 7, with a mix of hosted shared desktops (70%) and pooled desktops (30%) to support the user population. Where applicable, the document provides best practice recommendations and sizing guidelines for customer deployments of XenDesktop 7 on the Cisco Unified Computing System.

2. Solution Component Benefits

Each of the components of the overall solution materially contributes to the value of functional design contained in this document.

2.1. Benefits of Cisco Unified Computing System

Cisco Unified Computing System™ is the first converged data center platform that combines industry-standard, x86-architecture servers with networking and storage access into a single converged system. The system is entirely programmable using unified, model-based management to simplify and speed deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud computing environments.

Benefits of the Unified Computing System include:

Architectural flexibility

- Cisco UCS B-Series blade servers for infrastructure and virtual workload hosting
- Cisco UCS C-Series rack-mount servers for infrastructure and virtual workload Hosting
- Cisco UCS 6200 Series second generation fabric interconnects provide unified blade, network and storage connectivity
- Cisco UCS 5108 Blade Chassis provide the perfect environment for multi-server type, multi-purpose workloads in a single containment

Infrastructure Simplicity

- Converged, simplified architecture drives increased IT productivity
- Cisco UCS management results in flexible, agile, high performance, self-integrating information technology with faster ROI
- Fabric Extender technology reduces the number of system components to purchase, configure and maintain
- Standards-based, high bandwidth, low latency virtualization-aware unified fabric delivers high density, excellent virtual desktop user-experience

Business Agility

- Model-based management means faster deployment of new capacity for rapid and accurate scalability
- Scale up to 16 chassis and up to 128 blades in a single Cisco UCS management domain
- Leverage Cisco UCS Management Packs for System Center 2012 for integrated management

2.1.1. Benefits of Cisco Nexus 5548UP

The Cisco Nexus 5548UP Switch delivers innovative architectural flexibility, infrastructure simplicity, and business agility, with support for networking standards. For traditional, virtualized, unified, and high-performance computing (HPC) environments, it offers a long list of IT and business advantages, including:

Architectural Flexibility

- Unified ports that support traditional Ethernet, Fiber Channel (FC), and Fiber Channel over Ethernet (FCoE)
- Synchronizes system clocks with accuracy of less than one microsecond, based on IEEE 1588
- Offers converged Fabric extensibility, based on emerging standard IEEE 802.1BR, with Fabric Extender (FEX) Technology portfolio, including the Nexus 1000V Virtual Distributed Switch

Infrastructure Simplicity

- Common high-density, high-performance, data-center-class, fixed-form-factor platform
- Consolidates LAN and storage
- Supports any transport over an Ethernet-based fabric, including Layer 2 and Layer 3 traffic
- Supports storage traffic, including iSCSI, NAS, FC, RoE, and IBoE
- Reduces management points with FEX Technology

Business Agility

- Meets diverse data center deployments on one platform
- Provides rapid migration and transition for traditional and evolving technologies
- Offers performance and scalability to meet growing business needs

Specifications At-a-Glance

- A 1 -rack-unit, 1/10 Gigabit Ethernet switch
- 32 fixed Unified Ports on base chassis and one expansion slot totaling 48 ports
- The slot can support any of the three modules: Unified Ports, 1/2/4/8 native Fiber Channel, and Ethernet or FCoE
- Throughput of up to 960 Gbps

2.2. Benefits of EMC VNX Family of Storage Controllers

2.2.1. The EMC VNX Family

The EMC VNX Family delivers industry leading innovation and enterprise capabilities for file, block, and object storage in a scalable, easy-to-use solution. This next-generation storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's enterprises.

All of this is available in a choice of systems ranging from affordable entry-level solutions to high performance, petabyte-capacity configurations servicing the most demanding application requirements. The VNX family includes the VNXe Series, purpose-built for the IT generalist in smaller environments , and the VNX Series , designed to meet the high-performance, high scalability, requirements of midsize and large enterprises.

Figure 1 VNX Family



2.2.2. VNXe Series – Simple, Efficient, Affordable

The VNXe Series was designed with the IT generalist in mind and provides an affordable, integrated storage system for small-to-medium businesses as well as remote offices, and departments in larger enterprise businesses. The VNXe series provides true storage consolidation with a unique application-driven approach that eliminates the boundaries between applications and their storage.

This simple application-driven approach to managing shared storage makes the VNXe series ideal for IT generalists/managers and application administrators who may have limited storage expertise. EMC

Unisphere for the VNXe series enables easy, wizard-based provisioning of storage for Microsoft, Exchange, file shares, iSCSI volumes, VMware, and Hyper-V. VNXe supports tight integration with VMware to further facilitate efficient management of virtualized environments. Complemented by Unisphere Remote, the VNXe is also ideal for remote office-branch office (ROBO) deployments. Built-in efficiency capabilities, such as file de-duplication with compression and thin provisioning result in streamlined operations and can save up to 50 percent in upfront storage costs. Software packs aimed at facilitating backup, remote data protection, and disaster recovery include features such as easy-to-configure application snapshots.

The VNXe series supports high availability by using redundant components – power supplies, fans, and storage processors – as well as dynamic failover and failback. Additionally, the VNXe series supports the ability to upgrade system software or hardware while the VNXe system is running. It also delivers single click access to a world of resources such as comprehensive online documentation, training, and how-to-videos to expand your knowledge and answer questions.

2.2.3. VNX Series - Simple, Efficient, Powerful

The EMC VNX flash-optimized unified storage platform delivers innovation and enterprise capabilities for file, block, and object storage in a single, scalable, and easy-to-use solution. Ideal for mixed workloads in physical or virtual environments, VNX combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's virtualized application environments.

VNX includes many features and enhancements designed and built upon the first generation's success. These features and enhancements include:

- More capacity with multicore optimization with Multicore Cache, Multicore RAID, and Multicore FAST Cache (MCx™)
- Greater efficiency with a flash-optimized hybrid array
- Better protection by increasing application availability with active/active
- Easier administration and deployment by increasing productivity with new Unisphere® Management Suite

Next-Generation VNX is built to deliver even greater efficiency, performance, and scale than ever before.

2.3. Benefits of Microsoft Windows Server 2012 with Hyper-V

Microsoft Windows Server 2012 with Hyper-V builds on the architecture and functionality of Microsoft Hyper-V 2008 R2 allowing you to run the largest workloads in your virtualized environment. Windows Server 2012 with Hyper-V offers support for up to 64 virtual processors, 1 terabyte of memory per guest VM, and 4,000 virtual machines on a 64-node cluster. With Hyper-V, you can support Offloaded Data Transfer and improved Quality of Service to enforce minimum bandwidth requirements (even for network storage). High-availability options include incremental backup support, enhancements in clustered environments to support virtual Fiber Channel adapters within the virtual machine, and inbox

NIC Teaming. Windows Server 2012 Hyper-V can also use server message block file shares for virtual storage. This new option is simple to provision and offers performance capabilities and features that rival those available with Fiber Channel storage area networks. The Hyper-V Extensible Switch within Windows Server 2012 with Hyper-V gives you an open, extensible switch to help support security and management needs. You can build your own extensions, or use partner extensions to support these needs. Hyper-V works with Microsoft System Center 2012 SP1 management tools to handle your multi-server virtualization environment. With new management support for Hyper-V, you can fully automate management tasks and help reduce the administrative overhead costs of your environment.

Hyper-V provides a dynamic, reliable, and scalable virtualization platform combined with a single set of integrated management tools to manage both physical and virtual resources, enabling creation of an agile and dynamic data center.

2.4. Benefits of Citrix XenDesktop 7

There are many reasons to consider a virtual desktop solution. An ever growing and diverse base of users, an expanding number of traditional desktops, an increase in security mandates and government regulations, and the introduction of Bring Your Own Device (BYOD) initiatives are factors that add to the cost and complexity of delivering and managing desktop and application services.

Citrix XenDesktop™ 7 transforms the delivery of Microsoft Windows apps and desktops into a secure, centrally managed service that users can access on any device, anywhere. The release focuses on delivering these benefits:

- Mobilizing Microsoft Windows application delivery, bringing thousands of corporate applications to mobile devices with a native-touch experience and high performance
- Reducing costs with simplified and centralized management and automated operations

Securing data by centralizing information and effectively controlling access

Citrix XenDesktop 7 promotes mobility, allowing users to search for and subscribe to published resources, enabling a service delivery model that is cloud-ready.

The release follows a new unified FlexCast 2.0 architecture for provisioning all Windows apps and desktops either on hosted-shared RDS servers or VDI-based virtual machines. The new architecture combines simplified and integrated provisioning with personalization tools. Whether a customer is creating a system to deliver just apps or complete desktops, Citrix XenDesktop 7 leverages common policies and cohesive tools to govern infrastructure resources and access.

2.5. Audience

This document describes the architecture and deployment procedures of an infrastructure comprised of Cisco, EMC, Microsoft and Citrix virtualization. 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 deploy the solution described in this document.

3. Summary of Main Findings

The combination of technologies from Cisco Systems, Inc, Citrix Systems, Inc., Microsoft and EMC produced a highly efficient, robust and affordable Virtual Desktop Infrastructure (VDI) for a hosted virtual desktop deployment. Key components of the solution included:

- This design is Cisco's Desktop Virtualization Simplified Design, with compute and storage converged at the Cisco UCS Fabric Interconnect. In this design, the Cisco UCS Fabric Interconnects are uplinked directly to the Layer 3 network, reducing the solution footprint and cost. This design is well suited for smaller deployments of virtual desktop infrastructure.
- Local storage in the form of two 400 GB Enterprise SSD's provides fast local storage for the Citrix Provisioning Services write-cache drives and significantly reduces the impact on the primary EMC VNXe3300 storage array.
- Cisco UCS B200 M3 half-width blade with dual 12-core 2.7 GHz Intel Ivy Bridge (E5-2697v2) processors and 384GB of memory supports 25% more virtual desktop workloads than the previously released Sandy Bridge processors on the same hardware.
- The 500-user design is based on using one Unified Computing System chassis with three Cisco UCS B200 M3 blades for virtualized desktop workloads and one Cisco UCS B200 M3 blade for virtualized infrastructure workloads.
- The 1000-user design is based on using two Cisco Unified Computing System chassis with five Cisco UCS B200 M3 blades for virtualized desktop workloads and one Cisco UCS B200 M3 blade for virtualized infrastructure workloads.
- All log in and start workloads up to steady state were completed in 30-minutes without pegging the processor, exhausting memory or storage subsystems.
- The rackspace required to support the 500 users was a single rack of approximately 22 rack units. The space required to support 1000 users in a fully redundant configuration was only 28 RUs, which translates to an additional Cisco UCS 5108 chassis.
- Pure Virtualization: This Cisco Validated Design presents a validated design that is 100% virtualized on Microsoft Hyper-V 2012. All of the Windows 7 SP1 virtual desktops and supporting infrastructure components, including Active Directory, Provisioning Servers, SQL Servers, and XenDesktop delivery controllers, were hosted as virtual servers.
- Cisco maintains our industry leadership with our new Cisco UCS Manager 2.1.3(a) software that simplifies scaling, guarantees consistency, and eases maintenance.
- Our 10G unified fabric story gets additional validation on second generation Cisco UCS 6200 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- EMC's VNXe3300 system provides storage consolidation and outstanding efficiency for up to 1000 users.
- Citrix XenDesktop™ 7 follows a new unified product architecture that supports both hosted-shared desktops and applications (RDS) and complete virtual desktops (VDI). This new XenDesktop release simplifies tasks associated with large-scale VDI management. This modular solution supports seamless delivery of Windows apps and desktops as the number of users increase. In addition, HDX enhancements help to optimize performance and improve the user

experience across a variety of endpoint device types, from workstations to mobile devices including laptops, tablets, and smartphones.

- For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the XenDesktop 7 RDS virtual machines did not exceed the number of hyper-threaded cores available on the server. In other words, maximum performance is obtained when not overcommitting the CPU resources for hosted shared desktops.

4. Architecture

4.1. Hardware Deployed

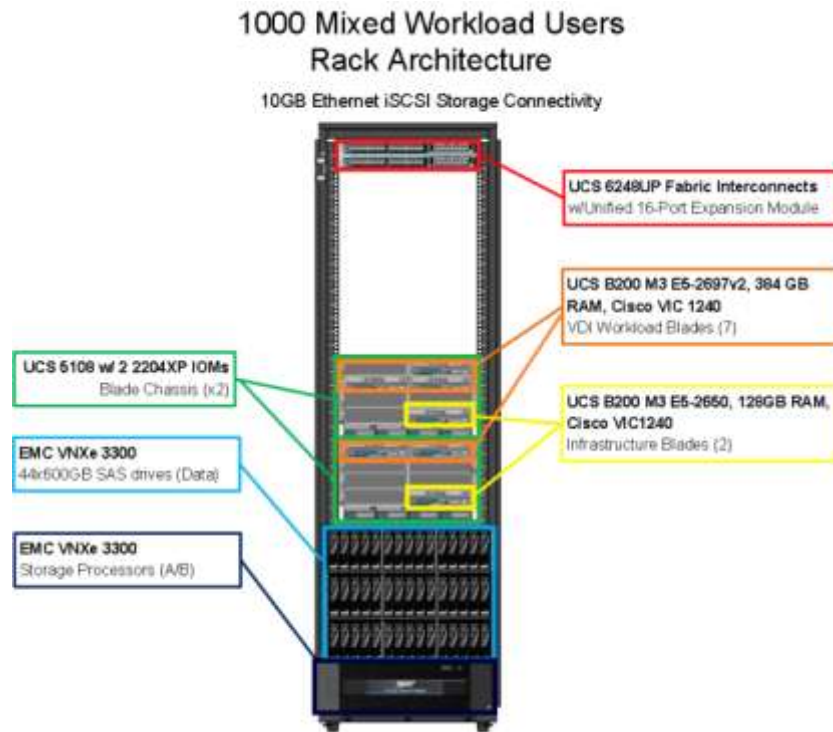
The architecture deployed is highly modular. While each customer's environment might vary in its exact configuration, when the reference architecture contained in this document is built, it can easily be scaled as requirements and demands change. This includes scaling both up (adding additional resources within a Cisco UCS Domain) and out (adding additional Cisco UCS Domains and VNX Storage arrays).

The 500- and 1000-user XenDesktop 7 solution includes Cisco networking, Cisco Unified Computing System, and EMC VNXe storage, which fits into a single data center rack, including the access layer network switches.

This Cisco Validated Design document details the deployment of the 500- and 1000-user configurations for a mixed XenDesktop workload featuring the following software:

- Citrix XenDesktop 7 Pooled Hosted Virtual Desktops with PVS write cache on Tier0 storage
- Citrix XenDesktop 7 Shared Hosted Virtual Desktops with PVS write cache on Tier0 storage
- Citrix Provisioning Server 7
- Citrix User Profile Manager
- Citrix StoreFront 2.0
- Cisco Nexus 1000V Distributed Virtual Switch
- Microsoft Windows Hyper-V 2012 Hypervisor
- Microsoft System Center 2012 Virtual Machine Manager SP1
- Microsoft SQL Server 2012 SP1

Figure 2: Workload Architecture



The workload contains the following hardware as shown in Figure 2: Workload Architecture:

- Two Cisco UCS 6248UP Series Fabric Interconnects
- Two Cisco UCS 5108 Blade Server Chassis (1 for each 500-users of capacity) with two 2204XP IO Modules per chassis
- Five Cisco UCS B200 M3 Blade Servers with Intel E5-2697v2 processors, 384GB RAM, and VIC1240 mezzanine cards for the mixed desktop virtualization workloads.
- Two Cisco UCS B200 M3 Blade Servers with Intel E5-2650 processors, 128 GB RAM, and VIC1240 mezzanine cards for the infrastructure virtualization workloads
- One EMC VNXe3300 dual controller storage system for HA, 44 SAS disks across 3 shelves, 10GE ports for network connectivity.

The EMC VNXe3300 disk shelf configurations are detailed in Section 5.4 Storage Architecture Design later in this document.

4.2. Logical Architecture

The logical architecture of the validated design is very similar between the 500-user and 1000-user configuration. The design was architected to support 500 users within a single chassis and four blades. The 1000-users configuration would require seven blades across two chassis, which would also add physical redundancy for the chassis. If full redundancy is required within the 500-user configuration, a second infrastructure host (INFRA-2) can be added or the virtual machines hosted on INFRA-2 could be placed on VDI 1-4. The table below outlines all the servers in the two configurations.

Table 1: Infrastructure Architecture

Server Name	Location	Purpose
INFRA-1	Physical – Chassis 1	Clustered Windows 2012 Datacenter server for infrastructure guests
VDI1-2	Physical – Chassis 1	Mixed workload Hyper-V 2012 server
VDI1-3	Physical – Chassis 1	Mixed workload Hyper-V 2012 server
VDI1-4	Physical – Chassis 1	Mixed workload Hyper-V 2012 server – (N+1) spare capacity
INFRA-2	Physical – Chassis 2	Clustered Windows 2012 Datacenter server for infrastructure guests (1000-user configuration)
VDI2-1	Physical – Chassis 2	Mixed workload Hyper-V 2012 server (1000 – user configuration)
VDI2-2	Physical – Chassis 2	Mixed workload Hyper-V 2012 server (1000-user configuration)
AD-DC1	Virtual – INFRA-1	Active Directory Domain Controller
EXC1	Virtual – INFRA-1	XenDesktop 7 controller
PVS1	Virtual – INFRA-1	Provisioning Services streaming server
SCVMM1/ SCVMM2	Virtual – INFRA-1	System Center 2012 Virtual Machine Manager Server
SFS1	Virtual – INFRA-1	StoreFront Services server
SQL1	Virtual – INFRA-1	SQL Server (primary) for AlwaysOn groups
Nexus1000V 1	Virtual – INFRA-1	Nexus 1000-V VSM HA node
Nexus1000V 3	Virtual – INFRA-1	Nexus 1000-V VSM HA node
HSDGold	Virtual – INFRA-1	Used to manage the PVS golden image for the Hosted Shared Desktop server image
XDGold	Virtual – INFRA-1	Used to managed the PVS golden image for the Windows 7 XenDesktop VDI image
AD-DC2	Virtual – INFRA-2	Active Directory Domain Controller
EXC2	Virtual – INFRA-2	XenDesktop 7 controller
PVS2	Virtual – INFRA-2	Provisioning Services streaming server
SQL2	Virtual – INFRA-2	SQL Server (secondary) for AlwaysOn groups
Nexus1000V 2	Virtual – INFRA-2	Nexus 1000-V VSM HA node
Nexus1000V 4	Virtual – INFRA-2	Nexus 1000-V VSM HA node

4.3. Software Revisions

This section includes the software versions of the primary products installed in the environment.

Table 2: Software Revisions

Vendor	Product	Version
Cisco	UCS Firmware	2.1(3a)
Cisco	UCS Manager	2.1(3a)
Cisco	Nexus 1000V for Hyper-V	5.2(1) SM1 (5.1)
Citrix	XenDesktop	7.0.0.3018
Citrix	Provisioning Services	7.0.0.46

Citrix	StoreFront Services	2.0.0.90
Microsoft	System Center 2012 Virtual Machine Manager SP1	3.1.6027.0
Microsoft	Windows Server 2012 DataCenter	6.2.9200 Build 9200
Microsoft	Hyper-V Server 2012	6.2.9200 Build 9200
Microsoft	SQL Server 2012 SP1	11.0.30000.0 (x64)

4.4. Configuration Guidelines

This section provides guidelines for situations where additional guidance may be necessary.

4.4.1. VLAN

The VLAN configuration recommended for the environment includes a total of eight VLANs as outlined in the table below.

Table 3: VLAN Configuration

VLAN Name	VLAN ID	Use
MGMT	60	Management. Used for the Hyper-V hosts and physical hardware. Should always be assigned to the first vNIC on any host and never connected to a Hyper-V virtual switch.
INFRA	61	Infrastructure. Used for all the virtualized infrastructure hosts, such as the XenDesktop Controllers, Provisioning Servers, SQL Servers, etc.
PVS-VDI	62	Provisioning Services and VDI. Used as the only network available for the provisioned hosted shared and virtual desktops.
CSV	63	Clustered Shared Volumes and Cluster heartbeat. Used only on the infrastructure cluster hosts for cluster communication and data volume traffic.
LMIGR	64	LiveMigration for Infrastructure Cluster. Used only on the infrastructure cluster hosts for live migration of guests between the two hosts.
STORAGE-A	65	iSCSI traffic on Fabric A. Used only for iSCSI traffic on channel A.
STORAGE-B	66	iSCSI traffic on Fabric B. Used only for iSCSI traffic on channel B.
STORAGE-NULL	99	Null Storage VLAN. Used temporarily during the Windows 2012 install to prevent the install from detecting multiple storage paths to the iSCSI volume.

As described in section 4.2 Logical Architecture section, the only clustering in the design is between INFRA-1 and INFRA-2 when 1000-users are involved or full redundancy is required for the infrastructure hosts. If INFRA-2 is not included in the design, then the clustering VLANs 63 and 64 can be omitted.

4.4.2. Hyper-V Clustering

This section describes the guidelines about configuring Microsoft Hyper-V Clustering.

4.4.2.1. *Network Configuration Guidelines*

Microsoft recommends having a minimum of two networks for your failover cluster: a public network that allows clients to connect to the cluster and a separate network that is used only for communication between the clustered servers. You can configure additional networks for specific storage options or for redundancy as needed.

When you use identical network adapters for a network, also use identical communication settings on those adapters (for example, Speed, Duplex Mode, Flow Control, and Media Type). Also, compare the settings between the network adapter and the switch it connects to and make sure that no settings are in conflict.

If you have private networks that are not routed to the rest of your network infrastructure, make sure that each of these private networks uses a unique subnet. This is necessary even if you give each network adapter a unique IP address. For example, if you have a cluster node in a central office that uses one physical network, and another node in a branch office that uses a separate physical network, do not specify 10.0.0.0/24 for both networks, even if you give each adapter a unique IP address.

4.4.2.2. *Prestage Cluster Computer Objects in Active Directory*

When you create a failover cluster by using the Create Cluster Wizard or by using Windows PowerShell, you must specify a name for the cluster. If you have sufficient permissions when you create the cluster, the cluster creation process automatically creates a computer object in AD DS that matches the cluster name. This object is called the cluster name object or CNO. Through the CNO, virtual computer objects (VCOs) are automatically created when you configure clustered roles that use client access points. For example, if you create a highly available file server with a client access point that is named FileServer1, the CNO will create a corresponding VCO in AD DS.

To create the CNO automatically, the user who creates the failover cluster must have the Create Computer objects permission to the organizational unit (OU) or the container where the servers that will form the cluster reside. To enable a user or group to create a cluster without having this permission, a user with appropriate permissions in AD DS (typically a domain administrator) can prestage the CNO in AD DS. This also provides the domain administrator more control over the naming convention that is used for the cluster, and control over which OU the cluster objects are created in.

Instructions about how to pre-stage the Cluster Name Object can be found here:

<http://technet.microsoft.com/en-us/library/dn466519.aspx>

4.4.2.3. *Quorum Configuration Guidelines*

The cluster software automatically configures the quorum for a new cluster, based on the number of nodes configured and the availability of shared storage. This is usually the most appropriate quorum configuration for that cluster. However, it is a good idea to review the quorum configuration after the cluster is created, before placing the cluster into production. To view the detailed cluster quorum

configuration, you can use the Validate a Configuration Wizard, or the Test-Cluster Windows PowerShell cmdlet, to run the Validate Quorum Configuration test. In Failover Cluster Manager, the basic quorum configuration is displayed in the summary information for the selected cluster, or you can review the information about quorum resources that returns when you run the Get-ClusterQuorum Windows PowerShell cmdlet.

At any time, you can run the Validate Quorum Configuration test to validate that the quorum configuration is optimal for your cluster. The test output indicates if a change to the quorum configuration is recommended and the settings that are optimal. If a change is recommended, you can use the Configure Cluster Quorum Wizard to apply the recommended settings.

After the cluster is in production, do not change the quorum configuration unless you have determined that the change is appropriate for your cluster.

4.4.2.4. Cluster Validation Tests

Before you create the failover cluster, we strongly recommend that you validate the configuration to make sure that the hardware and hardware settings are compatible with failover clustering. Microsoft supports a cluster solution only if the complete configuration passes all validation tests and if all hardware is certified for the version of Windows Server that the cluster nodes are running.

Note: *You must have at least two nodes to run all tests. If you have only one node, many of the critical storage tests do not run.*

The cluster validation tool can be launched through the Failover Cluster Manager, under Management, click Validate Configuration.

5. Infrastructure Components

This section describes the infrastructure components used in this Cisco Validated Design.

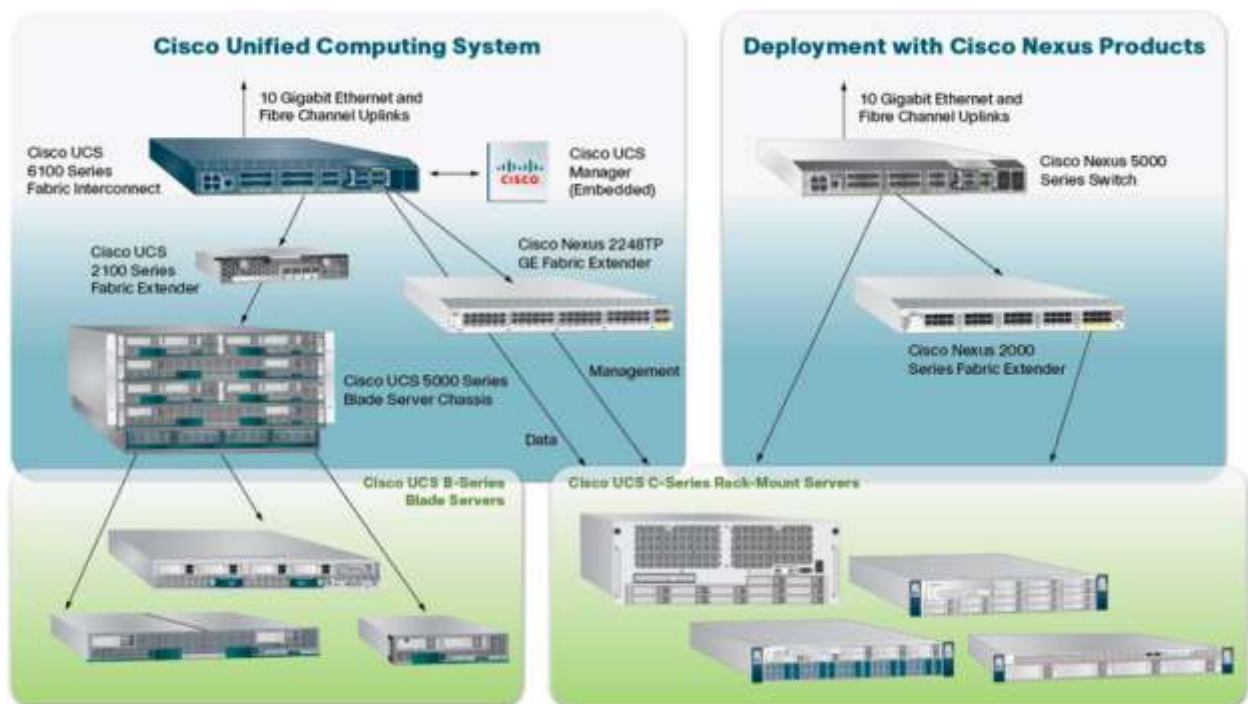
5.1. Cisco Unified Computer System (UCS)

The Cisco Unified Computing System™ (Cisco UCS™) is a next-generation data center platform that unites 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 in which all resources participate in a unified management domain.

5.1.1. Cisco Unified Computing Components

The Cisco UCS components are shown in the diagram below.

Figure 3: Cisco UCS Components



The Cisco Unified Computing System is designed from the ground up to be programmable and self-integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards, even the number and type of I/O interfaces is programmed dynamically, making every server ready to power any workload at any time.

With model-based management, administrators manipulate a model of a desired system configuration then associate a model's service profile with hardware resources. Once associated the system configures itself to match the model. This automation speeds provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco Fabric Extender technology reduces the number of system components to purchase, configure, manage, and maintain by condensing three network layers into one. It eliminates both blade server and hypervisor-based switches by connecting fabric interconnect ports directly to individual blade servers and virtual machines. Virtual networks are now managed exactly as physical networks are, but with massive scalability. This represents a radical simplification over traditional systems, reducing capital and operating costs while increasing business agility, simplifying and speeding deployment, and improving performance.

5.1.2. Cisco Fabric Interconnects

The Cisco UCS 6200 Series Fabric Interconnects are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system (Figure 2). The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet, Fiber Channel over Ethernet (FCoE), and Fiber Channel functions.

The Cisco UCS 6200 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the Cisco UCS 6200 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6200 Series provides both the LAN and SAN connectivity for all blades within its domain.

From a networking perspective, the Cisco UCS 6200 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, switching capacity of 2 terabits (Tb), and 320-Gbps bandwidth per chassis, independent of packet size and enabled services. The product family supports Cisco® low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the Interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

The Cisco UCS 6248UP is a 48-port Fabric Interconnect which provides low-latency throughput in excess of 1Tbps in a single rack unit (1 RU) form-factor. The Interconnect itself has 32 fixed ports of Fiber Channel, 10-Gigabit Ethernet, Cisco Data Center Ethernet, and FCoE SFP+ ports. One expansion module slot can provide an additional sixteen ports of Fiber Channel, 10-GE, Cisco Data Center Ethernet, and FCoE SFP+.

5.1.3. Cisco IO Modules (Fabric Extenders)

The Cisco UCS 2200 Series FEX is responsible for multiplexing and forwarding all traffic from blade servers in a chassis to a parent Cisco UCS Fabric Interconnect over the 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis, or VMs on the same blade, is forwarded to the parent interconnect, where network profiles are managed efficiently and effectively by the Fabric Interconnect. At the core of the Cisco UCS Fabric Extenders are ASIC processors developed by Cisco to multiplex all traffic.

Note: Up to two fabric extenders can be placed in a blade chassis.

Cisco UCS 2204 used in this architecture has eight 10GBASE-KR connections to the blade chassis mid-plane, with one connection per fabric extender for each of the chassis' eight half slots. This gives each half-slot blade server access to each of two 10-Gbps unified fabric-based networks through SFP+ sockets for both throughput and redundancy. It has 4 ports connecting up the fabric interconnect.

5.1.4. Cisco UCS Chassis

The Cisco UCS 5108 Series Blade Server Chassis is a 6 RU blade chassis that will accept up to eight half-width Cisco UCS B-Series Blade Servers or up to four full-width Cisco UCS B-Series Blade Servers, or a combination of the two. The Cisco UCS 5108 Series Blade Server Chassis can accept four redundant power supplies with automatic load-sharing and failover and two Cisco UCS (either 2100 or 2200 series) Fabric Extenders. The chassis is managed by Cisco UCS Chassis Management Controllers, which are mounted in the Cisco UCS Fabric Extenders and work in conjunction with the Cisco UCS Manager to control the chassis and its components.

A single Cisco UCS managed domain can theoretically scale to up to 40 individual chassis and 320 blade servers. At this time Cisco supports up to 20 individual chassis and 160 blade servers.

Basing the I/O infrastructure on a 10-Gbps unified network fabric allows the Cisco Unified Computing System to have a streamlined chassis with a simple yet comprehensive set of I/O options. The result is a chassis that has only five basic components:

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power supply bays with power entry in the rear, and hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two fabric extender slots accessible from the back panel
- Eight blade server slots accessible from the front panel

5.1.5. Cisco UCS Manager

The Cisco UCS 6200 Series Fabric Interconnect hosts and runs Cisco UCS Manager in a highly available configuration, enabling the fabric interconnects to fully manage all Cisco UCS elements. Connectivity to the Cisco UCS 5100 Series blade chassis is maintained through the Cisco UCS 2100 or 2200 Series Fabric Extenders in each blade chassis. The Cisco UCS 6200 Series interconnects support out-of-band management through a dedicated 10/100/1000-Mbps Ethernet management port as well as in-band management. Cisco UCS Manager typically is deployed in a clustered active-passive configuration on redundant fabric interconnects connected through dual 10/100/1000 Ethernet clustering ports.

5.1.6. Cisco UCS B200 M3 Blade Servers

Cisco UCS B200 M3 is a third generation half-slot, two-socket blade server. The Cisco UCS B200 M3 harnesses the power of the latest Intel® Xeon® processor E5-2600 product family, with up to 384 GB of RAM (using 16-GB DIMMs), two optional SAS/SATA/SSD disk drives, and up to dual 4x 10 Gigabit Ethernet throughput, utilizing our VIC 1240 LAN on motherboard (LOM) design. The Cisco UCS B200 M3 further extends the capabilities of Cisco Unified Computing System by delivering new levels of manageability, performance, energy efficiency, reliability, security, and I/O bandwidth for enterprise-class virtualization and other mainstream data center workloads.

5.1.7. Cisco Virtual Interface Card (VIC) Converged Network Adapter

A Cisco innovation, the Cisco UCS Virtual Interface Card (VIC) 1240 (Figure 4) is a 4-port 10 Gigabit Ethernet, Fiber Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the Cisco UCS M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional Port Expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet.

The Cisco UCS VIC 1240 enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs).

Figure 4: VIC 1240


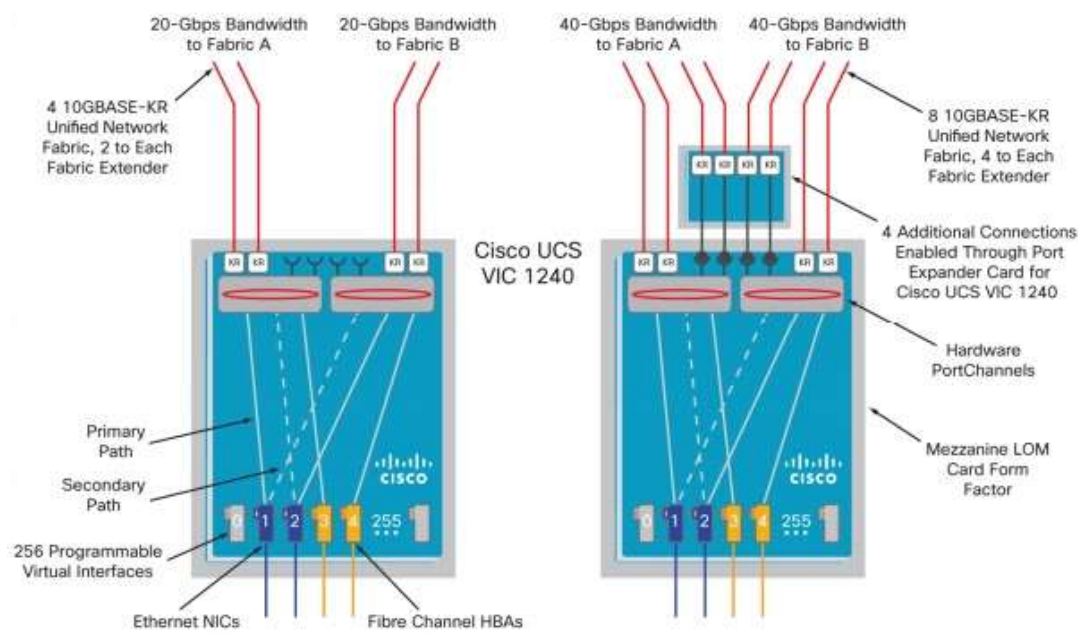


Figure 5: VIC 1240 Architecture



5.2. EMC VNXe3300

The EMC VNXe series redefines networked storage for the small business to small enterprise user, delivering an unequalled combination of features, simplicity, and efficiency. These unified storage systems provide true storage consolidation capability with seamless management and a unique application-driven approach that eliminates the boundaries between applications and their storage.

With scalability from six up to 150 disk drives and 450 terabytes of capacity, the VNXe series is ready to meet the needs of growing organizations with increasingly complex storage requirements. The VNXe3150™ is an ideal platform for businesses with physical server infrastructures, as well as those making the move to server virtualization to drive consolidation and greater efficiency. The VNXe3300™ includes all of the ease of use and application-driven management features of the VNXe3150, along with increased performance, scalability, and I/O expandability. Both systems share a comprehensive set of features including exceptional capacity utilization, data protection and availability solutions, and advanced support capabilities.

5.2.1. Advantages and Value Proposition

The EMC VNX™ family is optimized for virtual applications delivering industry-leading innovation and enterprise capabilities for file, block, and object storage in a scalable, easy-to-use solution. This next-generation storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today's enterprises.

The VNXe series is powered by Intel Xeon processor, for intelligent storage that automatically and efficiently scales in performance, while ensuring data integrity and security.

The VNXe series is purpose-built for the IT manager in smaller environments and the VNX series is designed to meet the high-performance, high-scalability requirements of midsize and large enterprises.

The table below lists the VNXe customer benefits.

Table 4: VNXe Benefits

Feature	
Next-generation unified storage, optimized for virtualized applications	✓
Capacity optimization features including compression, deduplication, thin provisioning, and application-centric copies	✓
High availability, designed to deliver five 9s availability	✓
Multiprotocol support for file and block	✓
Simplified management with EMC Unisphere™ for a single management interface for all NAS, SAN, and replication needs	✓

5.2.2. Software suites available

- Remote Protection Suite—Protects data against localized failures, outages, and disasters.
- Application Protection Suite—Automates application copies and proves compliance.
- Security and Compliance Suite—Keeps data safe from changes, deletions, and malicious activity.

5.2.3. Software packs available

- Total Value Pack—Includes all three protection software suites and the Security and Compliance Suite

5.3. Microsoft Technologies

5.3.1. Windows Server 2012

With Windows Server 2012, Microsoft delivers a server platform built on our experience of building and operating many of the world's largest cloud-based services and datacenter. Whether you are setting-up a single server for your small business or architecting a major new datacenter environment, Windows Server 2012 will help you cloud-optimize your IT so you can fully meet your organization's unique needs.

5.3.1.1. *Beyond Virtualization*

Offers a dynamic, multitenant infrastructure to help you scale and secure workloads and build a private cloud. Windows Server 2012 can help you provide:

- **Complete Virtualization Platform** - A fully-isolated, multitenant environment with tools that can help guarantee service level agreements, enable usage-based chargeback, and support self-service delivery.
- **Improved Scalability and Performance** - A high-density, scalable environment that you can modify to perform at an optimum level based on your needs.
- **Connecting to Cloud Services** - A common identity and management framework to enable highly secure and reliable cross-premises connectivity.

5.3.1.2. *The Power of Many Servers, the Simplicity of One*

Windows Server 2012 delivers a highly available and easy to manage cloud-optimized platform.

Windows Server 2012 can help you provide:

- **Flexible Storage** - Diverse storage choices that can help you achieve high performance, availability, and storage resource efficiency through virtualization and storage conservation.
- **Continuous Availability** - New and improved features that provide cost-effective, highly available services with protection against a wide range of failure scenarios.
- **Management Efficiency** - Automation of a broad set of management tasks and simplified deployment of workloads as you move toward full, lights-out automation.

5.3.1.3. *Every App, Any Cloud*

Microsoft Windows Server 2012 offers a cloud-optimized server platform that gives you the flexibility to build and deploy applications and websites on-premises, in the cloud, or across both. Windows Server 2012 can help you deliver:

- **Flexibility to Build On-Premises and in the Cloud** - A consistent set of tools and frameworks that enables developers to build symmetrical or hybrid applications across the datacenter and the cloud.
- **A Scalable and Elastic Infrastructure** - New features to help you increase website density and efficiency, plus frameworks, services, and tools to increase the scalability and elasticity of modern applications.
- **An Open Web and App Development Environment** - An open platform that enables mission-critical applications and provides enhanced support for open standards, open-source applications, and various development languages

5.3.2. Failover Clustering

Failover clusters provide high availability and scalability to many server workloads. These include server applications such as Microsoft Exchange Server, Hyper-V, Microsoft SQL Server, and file servers. The server applications can run on physical servers or virtual machines. In a failover cluster, if one or more of the clustered servers (nodes) fails, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they restart or move to another node. Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes.

5.3.3. Clustered Shared Volumes

Cluster Shared Volumes (CSVs) in a Windows Server 2012 failover cluster allow multiple nodes in the cluster to simultaneously have read-write access to the same LUN (disk) that is provisioned as an NTFS volume. With CSVs, clustered roles can fail over quickly from one node to another node without requiring a change in drive ownership, or dismounting and remounting a volume. CSVs also help simplify managing a potentially large number of LUNs in a failover cluster.

CSVs provide a general-purpose, clustered file system in Windows Server 2012, which is layered above NTFS. They are not restricted to specific clustered workloads. (In Windows Server 2008 R2, CSVs only supported the Hyper-V workload.) CSV applications include:

- Clustered virtual hard disk (VHD) files for clustered Hyper-V virtual machines
- Scale-out file shares to store application data for the Scale-Out File Server role

5.3.4. Networking Support

Windows Server 2012 makes it as straightforward to manage an entire network as a single server, giving you the reliability and scalability of multiple servers at a lower cost. Automatic rerouting around storage, server, and network failures enables file services to remain online with minimal noticeable downtime.

Plus Windows Server 2012 – together with System Center 2012 SP1 – provides an end-to-end Software Defined Networking solution across public, private, and hybrid cloud implementations.

Whatever your organization's needs, be it administering network assets to managing an extensive private and public cloud network infrastructure, Windows Server 2012 offers you solutions to today's changing business landscape.

5.3.5. Hyper-V

Windows Server 2012 with Hyper-V is a virtualization platform that has helped organizations of all sizes realize considerable cost savings and operational efficiencies. With industry leading size and scale, Hyper-V is the platform of choice for you to run your mission critical workloads.

Hyper-V in Windows Server 2012 greatly expands support for host processors and memory. It now includes support for up to 64 processors and 1 terabyte of memory for Hyper-V guests, a new VHDX virtual hard disk format with larger disk capacity of up to 64 terabytes, and additional resilience.

Using Windows Server 2012 with Hyper-V, you can take advantage of new hardware technology, while still utilizing the servers you already have. This way you can virtualize today, and be ready for the future.

Whether you are looking to help increase VM mobility, help increase VM availability, handle multi-tenant environments, gain bigger scale, or gain more flexibility, Windows Server 2012 with Hyper-V gives you the platform and tools you need to increase business flexibility with confidence. And you get the portability you need to virtualize on premises or extend your datacenter out to a hosting providing, helping you transform your datacenter into a cloud computing environment.

5.3.6. Hyper-V Server 2012

Microsoft Hyper-V Server 2012 is a hypervisor-based server virtualization product that enables you to consolidate workloads, helping organizations improve server utilization and reduce costs.

Hyper-V Server is a dedicated stand-alone product that contains the hypervisor, Windows Server driver model, virtualization capabilities, and supporting components such as failover clustering, but does not contain the robust set of features and roles as the Windows Server operating system. As a result Hyper-V Server produces a small footprint and requires minimal overhead. Organizations consolidating servers where no new Windows Server licenses are required or where the servers being consolidated are running an alternative OS may want to consider Hyper-V Server.

One of the most common uses for Hyper-V Server is in Virtual Desktop Infrastructure (VDI) environments. VDI allows a Windows client operating system to run on server-based virtual machines in the datacenter, which the user can access from a PC, thin client, or other client device. A full client environment is virtualized within a server-based hypervisor, centralizing users' desktops.

5.3.7. SQL Server 2012

Microsoft® SQL Server™ is a database management and analysis system for e-commerce, line-of-business, and data warehousing solutions. SQL Server 2012, the latest version, adds new high availability and disaster recovery solutions through AlwaysOn clusters and availability groups, xVelocity in-memory storage for extremely fast query performance, rapid data exploration and scalable business intelligence

through Power View and tabular modeling in Analysis Services, and new data management capability with Data Quality Services.

5.3.7.1. *AlwaysOn Application Groups*

The AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, that fail over together. An availability group supports a set of read-write primary databases and one to four sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations.

5.3.8. *System Center Virtual Machine Manager 2012 SP1*

Microsoft System Center 2012 provides a common management toolset to help you configure, provision, monitor, and operate your IT infrastructure. If your infrastructure is like that of most organizations, you have physical and virtual resources running heterogeneous operating systems. The integrated physical, virtual, private, and public cloud management capabilities in System Center 2012 can help ensure efficient IT management and optimized ROI of those resources.

Virtual Machine Manager (VMM) is a management solution for the virtualized datacenter, enabling you to configure and manage your virtualization host, networking, and storage resources in order to create and deploy virtual machines and services to private clouds that you have created.

Virtual Machine Manager uses a single pane of glass to manage multi-hypervisor virtualized environments such as Windows Server Hyper-V, Citrix XenServer, and VMware vSphere. This enables you to extend existing investments while you build your private cloud.

5.4. *Citrix XenDesktop 7*

5.4.1. *Enhancements in This Release*

Built on the Avalon™ architecture, Citrix XenDesktop™ 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:

- A new unified product architecture—the FlexCast 2.0 architecture—and administrative interfaces designed to deliver both hosted-shared applications (RDS) and complete virtual desktops (VDI). Unlike previous software releases that required separate Citrix XenApp farms and XenDesktop infrastructures, this new release allows administrators to deploy a single infrastructure and employ a consistent set of management tools for mixed desktop and app workloads.
- New and improved management interfaces. XenDesktop 7 includes two new purpose-built management consoles—one for automating workload provisioning and app publishing and the second for real-time monitoring of the infrastructure.

- Enhanced HDX technologies. Since mobile technologies and devices are increasingly pervasive, Citrix has engineered new and improved HDX technologies to improve the user experience for hosted Windows apps and desktops delivered on laptops, tablets, and smartphones.
- Unified App Store. The release includes a self-service Windows app store, implemented through Citrix StoreFront services, that provides a single, simple, and consistent aggregation point for all user services. IT can publish apps, desktops, and data services to the StoreFront, from which users can search and subscribe to services.

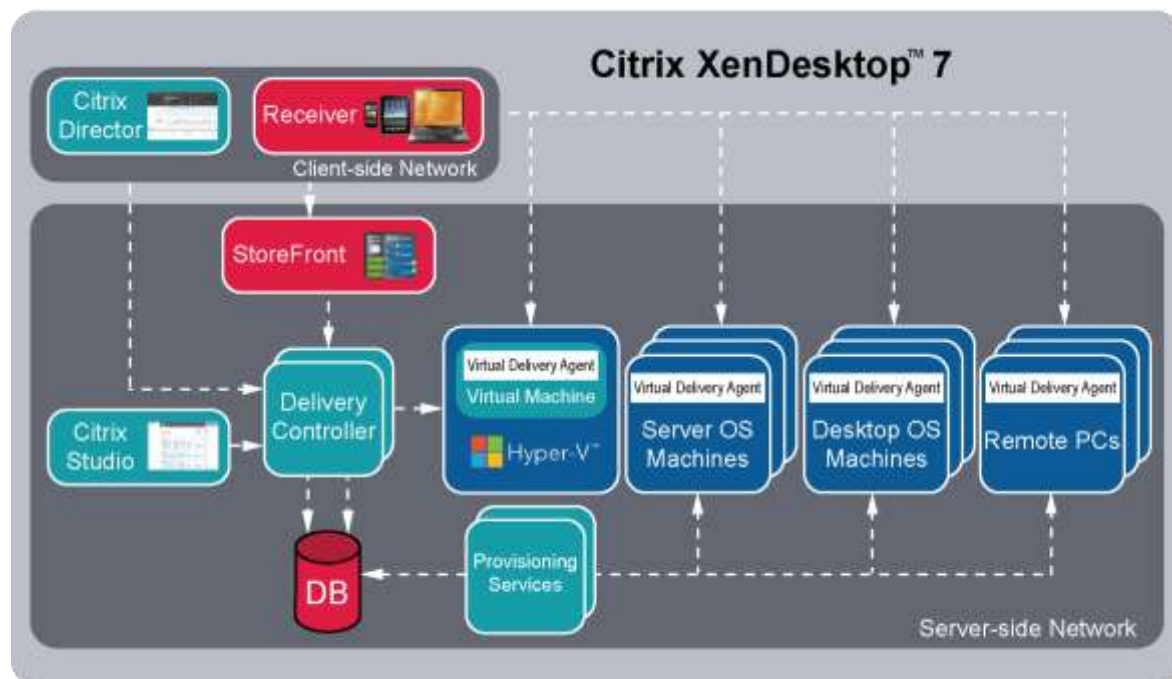
5.4.2. FlexCast Technology

In Citrix XenDesktop 7, FlexCast 2.0 technology is responsible for delivering and managing hosted-shared RDS apps and complete VDI desktops. By using Citrix Receiver with XenDesktop 7, users have a device-native experience on endpoints including Windows, Mac, Linux, iOS, Android, ChromeOS, and Blackberry.

The diagram below shows an overview of the unified FlexCast 2.0 architecture and underlying components, which are also described below:

- Citrix Receiver. Running on user endpoints, Receiver provides users with self-service access to resources published on XenDesktop servers. Receiver combines ease of deployment and use, supplying fast, secure access to hosted applications, desktops, and data. Receiver also provides on-demand access to Windows, Web, and Software-as-a-Service (SaaS) applications.
- Citrix StoreFront. StoreFront authenticates users and manages catalogs of desktops and applications. Users can search StoreFront catalogs and subscribe to published services through Citrix Receiver.

Figure 6: XenDesktop 7 Architecture



- Citrix Studio. Using the new and improved Studio interface, administrators can easily configure and manage XenDesktop deployments. Studio provides wizards to guide the process of setting up an environment, creating desktops, and assigning desktops to users, automating provisioning and application publishing. It also allows administration tasks to be customized and delegated to match site operational requirements.
- Delivery Controller. The Delivery Controller is responsible for distributing applications and desktops, managing user access, and optimizing connections to applications. Each site has one or more delivery controllers.
- Server OS Machines. These are virtual or physical machines (based on a Windows Server operating system) that deliver RDS applications or hosted shared desktops to users.
- Desktop OS Machines. These are virtual or physical machines (based on a Windows Desktop operating system) that deliver personalized VDI desktops or applications that run on a desktop operating system.
- Remote PC. XenDesktop with Remote PC allows IT to centrally deploy secure remote access to all Windows PCs on the corporate network. It is a comprehensive solution that delivers fast, secure remote access to all the corporate apps and data on an office PC from any device.
- Virtual Delivery Agent. A Virtual Delivery Agent is installed on each virtual or physical machine (within the server or desktop OS) and manages each user connection for application and desktop services. The agent allows OS machines to register with the Delivery Controllers and governs the HDX connection between these machines and Citrix Receiver.

- Citrix Director. Citrix Director is a powerful administrative tool that helps administrators quickly troubleshoot and resolve issues. It supports real-time assessment, site health and performance metrics, and end user experience monitoring. Citrix EdgeSight® reports are available from within the Director console and provide historical trending and correlation for capacity planning and service level assurance.
- Citrix Provisioning Services 7. This new release of Citrix Provisioning Services (PVS) technology is responsible for streaming a shared virtual disk (vDisk) image to the configured Server OS or Desktop OS machines. This streaming capability allows VMs to be provisioned and re-provisioned in real-time from a single image, eliminating the need to patch individual systems and conserving storage. All patching is done in one place and then streamed at boot-up. Citrix PVS 7 supports image management for both RDS and VDI-based machines, including support for image snapshots and rollbacks.

5.4.3. High-Definition User Experience (HDX) Technology


High-Definition User Experience (HDX) technology in this release is optimized to improve the user experience for hosted Windows apps on mobile devices. Specific enhancements include:

- HDX Mobile™ technology, designed to cope with the variability and packet loss inherent in today's mobile networks. HDX technology supports deep compression and redirection, taking advantage of advanced codec acceleration and an industry-leading H.264-based compression algorithm. The technology enables dramatic improvements in frame rates while requiring significantly less bandwidth. HDX technology offers users a rich multimedia experience and optimized performance for voice and video collaboration.
- HDX Touch technology enables mobile navigation capabilities similar to native apps, without rewrites or porting of existing Windows applications. Optimizations support native menu controls, multi-touch gestures, and intelligent sensing of text-entry fields, providing a native application look and feel.
- HDX 3D Pro uses advanced server-side GPU resources for compression and rendering of the latest OpenGL and DirectX professional graphics apps. GPU support includes both dedicated user and shared user workloads.

5.4.4. Citrix XenDesktop 7 Desktop and Application Services

IT departments strive to deliver application services to a broad range of enterprise users that have varying performance, personalization, and mobility requirements. Citrix XenDesktop 7 allows IT to configure and deliver any type of virtual desktop or app, hosted or local, and optimize delivery to meet individual user requirements, while simplifying operations, securing data, and reducing costs.

Figure 7: XenDesktop Controller



With previous product releases, administrators had to deploy separate XenApp farms and XenDesktop sites to support both hosted-shared RDS and VDI desktops. As shown above, the new XenDesktop 7 release allows administrators to create a single infrastructure that supports multiple modes of service delivery, including:

- Application Virtualization and Hosting (RDS). Applications are installed on or streamed to Windows servers in the data center and remotely displayed to users' desktops and devices.
- Hosted Shared Desktops (RDS). Multiple user sessions share a single, locked-down Windows Server environment running in the datacenter and accessing a core set of apps. This model of service delivery is ideal for task workers using low intensity applications, and enables more desktops per host compared to VDI.
- Pooled VDI Desktops. This approach leverages a single desktop OS image to create multiple thinly provisioned or streamed desktops. Optionally, desktops can be configured with a Personal vDisk to maintain user application, profile and data differences that are not part of the base image. This approach replaces the need for dedicated desktops, and is generally deployed to address the desktop needs of knowledge workers that run more intensive application workloads.
- VM Hosted Apps (16 bit, 32 bit, or 64 bit Windows apps). Applications are hosted on virtual desktops running Windows 7, XP, or Vista and then remotely displayed to users' physical or virtual desktops and devices.

This Cisco Validated Design focuses on delivering a mixed workload consisting of hosted shared desktops (RDS) and pooled VDI desktops.

5.4.5. Provisioning Services 7

One significant advantage to service delivery through RDS and VDI is how these technologies simplify desktop administration and management. Citrix Provisioning Services (PVS) takes the approach of streaming a single shared virtual disk (vDisk) image rather than provisioning and distributing multiple OS

image copies across multiple virtual machines. One advantage of this approach is that it constrains the number of disk images that must be managed, even as the number of desktops grows, providing image consistency. At the same time, using a single shared image (rather than hundreds or thousands of desktop images) significantly reduces the required storage footprint and dramatically simplifies image management.

Since there is a single master image, patch management is simple and reliable. All patching is done on the master image, which is then streamed as needed. When an updated image is ready for production, the administrator simply reboots to deploy the new image. Rolling back to a previous image is done in the same manner. Local hard disk drives in user systems can be used for runtime data caching or, in some scenarios, removed entirely, lowering power usage, system failure rates, and security risks.

After installing and configuring PVS components, 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. vDisks can exist on a Provisioning Server, file share, or in larger deployments (as in this Cisco Validated Design), on a storage system with which the Provisioning Server can communicate (through iSCSI, SAN, NAS, and CIFS). vDisks can be assigned to a single target device in Private Image Mode, or to multiple target devices in Standard Image Mode.

When a user device boots, the appropriate vDisk is located based on the boot configuration and mounted on the Provisioning Server. The software on that vDisk is then streamed to the target device and appears like a regular hard drive to the system. Instead of pulling all the vDisk contents down to the target device (as is done with some imaging deployment solutions), the data is brought across the network in real time, as needed. This greatly improves the overall user experience since it minimizes desktop startup time.

This release of PVS extends built-in administrator roles to support delegated administration based on groups that already exist within the network (Windows or Active Directory Groups). All group members share the same administrative privileges within a farm. An administrator may have multiple roles if they belong to more than one group.

6. Solution Architecture

An ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Computer (BYOC) to work programs are prime reasons for moving to a virtual desktop solution. The first step in designing a virtual desktop solution is to understand the user community and the type of tasks that are required to successfully execute their role. Users generally fall into one of the following classifications:

- **Knowledge Workers** today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even connect from coffee shops. These anywhere workers expect access to all of their applications and data wherever they are.
- **External Contractors** are increasingly part of everyday business. They need access to only certain portions of applications and data, yet administrators still have little control over the

devices they use and the locations from which they work. Consequently, the IT staff must choose between the cost of providing these workers a device or assuming the inherent security risk of allowing access to company data from unmanaged devices.

- **Task Workers** perform a set of well-defined, repetitive tasks. These workers run a limited set of applications that are less resource-intensive than applications run by knowledge workers. However, since task workers are interacting with customers, partners, and employees, they also have access to critical business data.
- **Mobile Workers** need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.
- **Shared Workstation** users are often found in state-of-the-art university and business computer labs, conference rooms or training centers. Shared workstation environments must constantly be updated with the latest operating systems and applications as the needs of the organization change.

After classifying the user population and evaluating the business requirements, the next step is to review and select the appropriate type of virtual desktop for each user classification. The five potential desktop environments are as follows:

- **Traditional PC:** A traditional PC is what traditionally constitutes a desktop environment: a physical device with a locally installed operating system.
- **Hosted Shared Desktop:** A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, multiple users simultaneously share a single installed instance of a server operating system, such as Microsoft Windows Server 2012. Each user receives a desktop "session" and works in an isolated memory space. Changes made by one user could impact the other users.
- **Hosted Virtual Desktop:** A hosted virtual desktop is a virtual desktop running either on a virtualization layer (such as Microsoft Hyper-V) or on bare metal hardware. The user does not work with and sit in front of the desktop, but instead the user interacts through a delivery protocol.
- **Streamed Applications:** Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly but the resources may only be available while they are connected to the network.
- **Local Virtual Desktop:** A local virtual desktop is a desktop running entirely on the user's local device and continues to operate when disconnected from the network. In this case, the user's local device is used as a type 1 hypervisor and is synced with the data center when the device is connected to the network.

For the purposes of the validation represented in this document, only hosted shared desktops and hosted virtual desktops were validated. Each of the sections provides some fundamental design decisions for this environment.

6.1. Citrix Design Fundamentals

With Citrix XenDesktop 7, the method you choose to provide applications or desktops to users depends on the types of applications and desktops you are hosting and available system resources, as well as the types of users and user experience you want to provide.

Table 5: Citrix Design Decisions

Server OS machines	<p>You want: Inexpensive server-based delivery to minimize the cost of delivering applications to a large number of users, while providing a secure, high-definition user experience.</p> <p>Your users: Perform well-defined tasks and do not require personalization or offline access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations.</p> <p>Application types: Any application.</p>
Desktop OS machines	<p>You want: A client-based application delivery solution that is secure, provides centralized management, and supports a large number of users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition.</p> <p>Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications.</p> <p>Application types: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines.</p> <p>Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.</p>
Remote PC Access	<p>You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYOD support without migrating desktop images into the datacenter.</p> <p>Your users: Employees or contractors that have the option to work from home, but need access to specific software or data on their corporate desktops to perform their jobs remotely.</p> <p>Host: The same as Desktop OS machines.</p> <p>Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.</p>

For the Cisco Validated Design described in this document, Hosted Shared (using Server OS machines) and Hosted Virtual Desktops (using Desktop OS machines) were configured and tested. The following sections discuss fundamental design decisions relative to this environment.

6.1.1. Citrix Hosted Shared Desktop Design Fundamentals

Citrix XenDesktop 7 integrates Hosted Shared and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, and manageable solution for delivering Windows applications and desktops as a service.

Users can select applications from an easy-to-use “store” that is accessible from tablets, smartphones, PCs, Macs, and thin clients. XenDesktop delivers a native touch-optimized experience with HDX high-definition performance, even over mobile networks.

6.1.1.1. Machine Catalogs

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Machine Catalog. In this Cisco Validated Design, VM provisioning relies on Citrix Provisioning Services to make sure that the machines in the catalog are consistent. In this Cisco Validated Design, machines in the Machine Catalog are configured to run either a Windows Server OS (for RDS hosted shared desktops) or a Windows Desktop OS (for hosted pooled VDI desktops).

6.1.1.2. Delivery Groups

To deliver desktops and applications to users, you create a Machine Catalog and then allocate machines from the catalog to users by creating Delivery Groups. Delivery Groups provide desktops, applications, or a combination of desktops and applications to users. Creating a Delivery Group is a flexible way of allocating machines and applications to users. In a Delivery Group, you can:

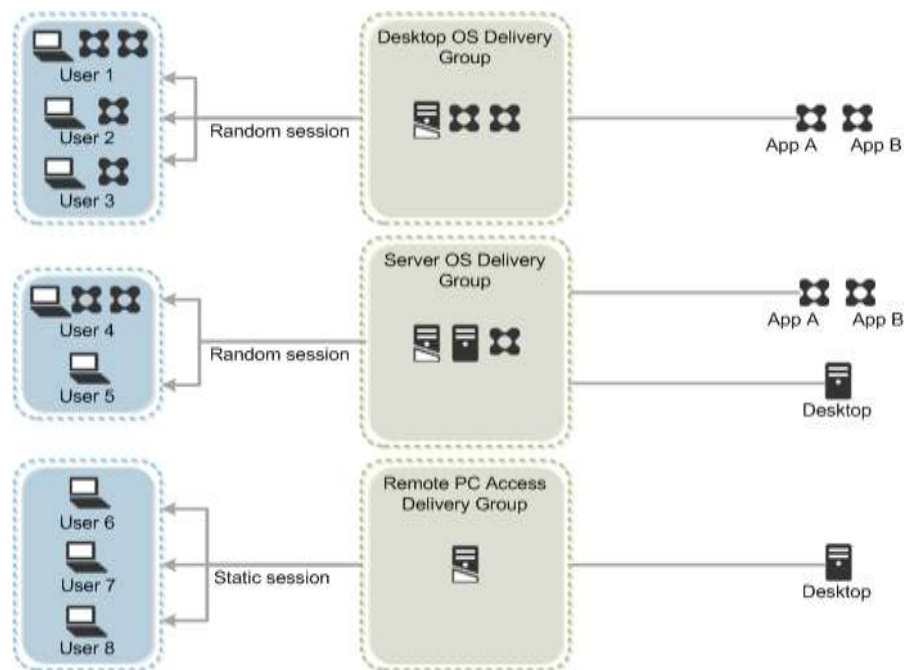
- Use machines from multiple catalogs
- Allocate a user to multiple machines
- Allocate multiple users to one machine

As part of the creation process, you specify the following Delivery Group properties:

- Users, groups, and applications allocated to Delivery Groups
- Desktop settings to match users' needs
- Desktop power management options

The graphic below shows how users access desktops and applications through machine catalogs and delivery groups. (Note that only Server OS and Desktop OS Machines are configured in this Cisco Validated Design solution to support hosted shared and pooled virtual desktops.)

Figure 8: User Access in XenDesktop 7



6.1.1.3. Hypervisor Selection

Citrix XenDesktop is hypervisor-agnostic, so any of the following three hypervisors can be used to host RDS- and VDI-based desktops:

- Hyper-V: Microsoft Windows Server 2012 with Hyper-V is available in a Standard, Server Core and free Hyper-V Server 2008 R2 versions. More information on Hyper-V can be obtained at the Microsoft web site: <http://www.microsoft.com/en-us/server-cloud/windows-server/default.aspx>.
- VMware vSphere: VMware vSphere comprises the management infrastructure or virtual center server software and the hypervisor software that virtualizes the hardware resources on the servers. It offers features like Distributed Resource Scheduler, vMotion, high availability, Storage vMotion, VMFS, and a multipathing storage layer. More information on vSphere can be obtained at the VMware web site: <http://www.vmware.com/products/datacenter-virtualization/vsphere/overview.html>.
- XenServer: Citrix® XenServer® is a complete, managed server virtualization platform built on the powerful Xen® hypervisor. Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. XenServer is designed for efficient management of Windows and Linux virtual servers and delivers cost-effective server consolidation and business continuity. More information on XenServer can be obtained at the web site: <http://www.citrix.com/products/xenserver/overview.html>

For this Cisco Validated Design, the hypervisor used was Microsoft Windows Server 2012 with Hyper-V.

6.1.1.4. Provisioning Services

Citrix XenDesktop 7 can be deployed with or without Citrix Provisioning Services (PVS). The advantage of using Citrix PVS is that it allows computers to be provisioned and re-provisioned in real-time from a single shared-disk image. In this way Citrix PVS greatly reduces the amount of storage required in comparison to other methods of provisioning virtual desktops.

Citrix PVS can create desktops as Pooled or Private:

- Private Desktop: A private desktop is a single desktop assigned to one distinct user.
- Pooled Desktop: A pooled virtual desktop uses Citrix PVS to stream a standard desktop image to multiple desktop instances upon boot.

When considering a PVS deployment, there are some design decisions that need to be made regarding the write cache for the virtual desktop devices that leverage provisioning services. The write cache is a cache of all data that the target device has written. If data is written to the PVS vDisk in a caching mode, the data is not written back to the base vDisk. Instead it is written to a write cache file in one of the following locations:

- Cache on device HD: Cache on local HD is stored in a file on a secondary local hard drive of the device. It gets created as an invisible file in the root folder of the local HD. The cache file size grows as needed, but it never gets larger than the original vDisk and frequently not larger than the free space on the original vDisk.
- Cache in device RAM: Cache is stored in client RAM (memory), The cache maximum size is fixed by a parameter in vDisk properties. All written data can be read from local RAM instead of going back to the server. RAM cache is faster than server cache and works in a high availability environment.
- Cache on server: Server cache is stored in a file on the server, or on a share, SAN, or other network storage resource. The file size grows as needed, but never gets larger than the original vDisk and frequently not larger than the free space on the original vDisk. It is slower than RAM cache because all reads/writes have to go to the server and be read from a file. Cache gets deleted when the device reboots; in other words, on every boot the device reverts to the base image. Changes remain only during a single boot session.
- Cache on device hard drive persisted: (Experimental Phase) This is the same as “Cache on device hard drive”, except that the cache persists. At this time, this write cache method is an experimental feature only, and is only supported for NT6.1 or later (Windows 7 and Windows 2008 R2 and later). This method also requires a different bootstrap.
- Cache on server persisted: This cache option allows for the saving of changes between reboots. Using this option, after rebooting, a target device is able to retrieve changes made from previous sessions that differ from the read only vDisk image. If a vDisk is set to Cache on server persistent, each target device that accesses the vDisk automatically has a device-specific, writable disk file created. Any changes made to the vDisk image are written to that file, which is not automatically deleted upon shutdown.

The alternative to Citrix Provisioning Services for pooled desktop deployments is Citrix Machine Creation Services (MCS), which is integrated directly with the XenDesktop Studio console.

For this study, we used PVS 7 for managing Pooled Desktops with cache on device HD of each virtual machine so that the design would scale up to many thousands of desktops. Provisioning Server 7 was used for Active Directory machine account creation and management as well as for streaming the shared disk to the hypervisor hosts.

6.1.1.5. System Center 2012 Virtual Machine Manager

Microsoft System Center Virtual Machine Manager (SCVMM) 2012 is a management solution for the virtualized datacenter, enabling you to configure and manage your virtualization host, networking, and storage resources in order to create and deploy virtual machines and services to private clouds that you have created. Microsoft System Center 2012 cloud and datacenter management solutions empower you with a common management toolset for your private and public cloud applications and services. SCVMM is an integral part of the System Center 2012 Application Management component.

6.1.1.6. Example Deployments

The following are two examples of typical XenDesktop deployments:

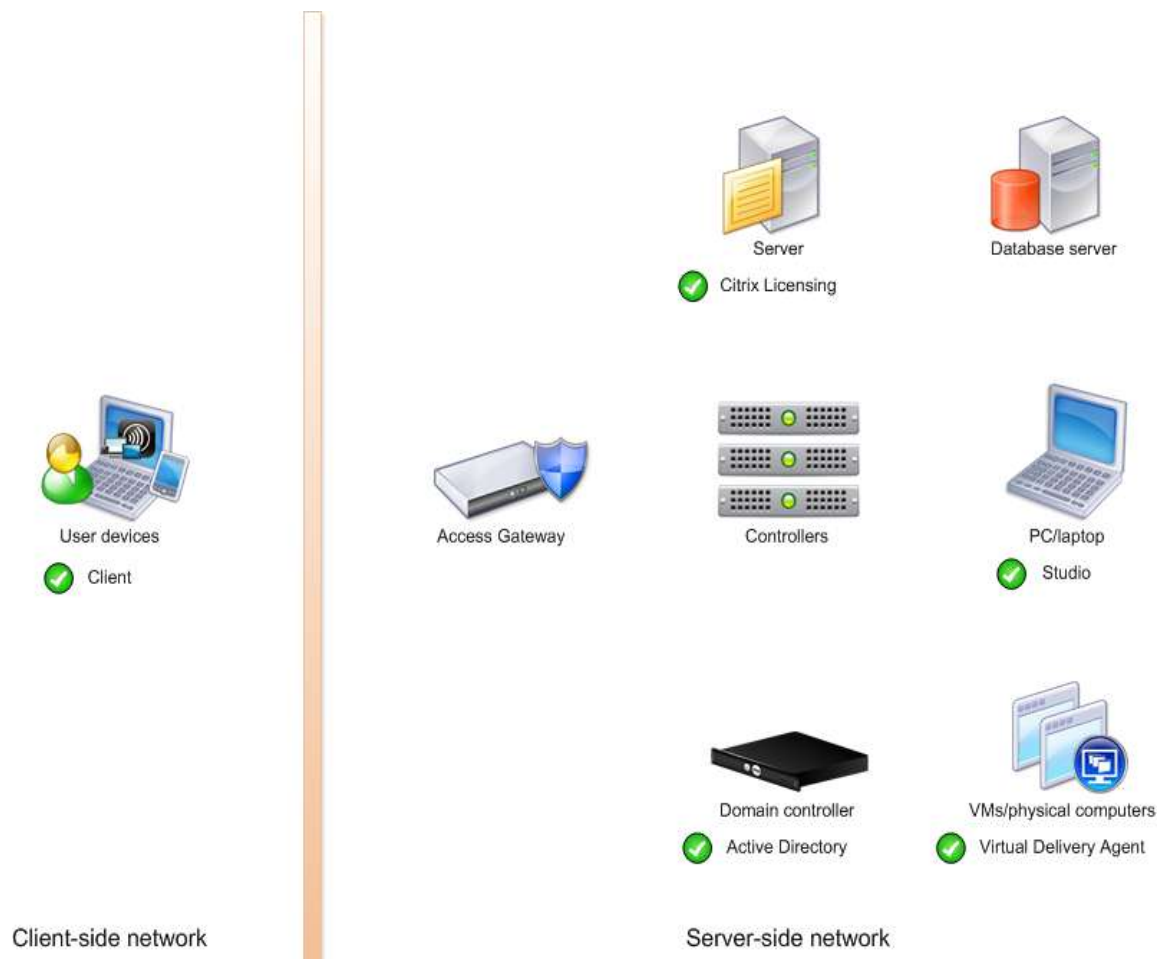
- A distributed components configuration
- A multiple site configuration

6.1.1.7. Distributed Components Configuration

You can distribute the components of your deployment among a greater number of servers, or provide greater scalability and failover by increasing the number of controllers in your site. You can install management consoles on separate computers to manage the deployment remotely. A distributed deployment is necessary for an infrastructure based on remote access through NetScaler Gateway (formerly called Access Gateway).

The diagram below shows an example of a distributed components configuration. A simplified version of this configuration is often deployed for an initial proof-of-concept (POC) deployment. The Cisco Validated Design described in this document deploys Citrix XenDesktop in a configuration that resembles this distributed components configuration shown.

Figure 9: Distributed components in XenDesktop 7



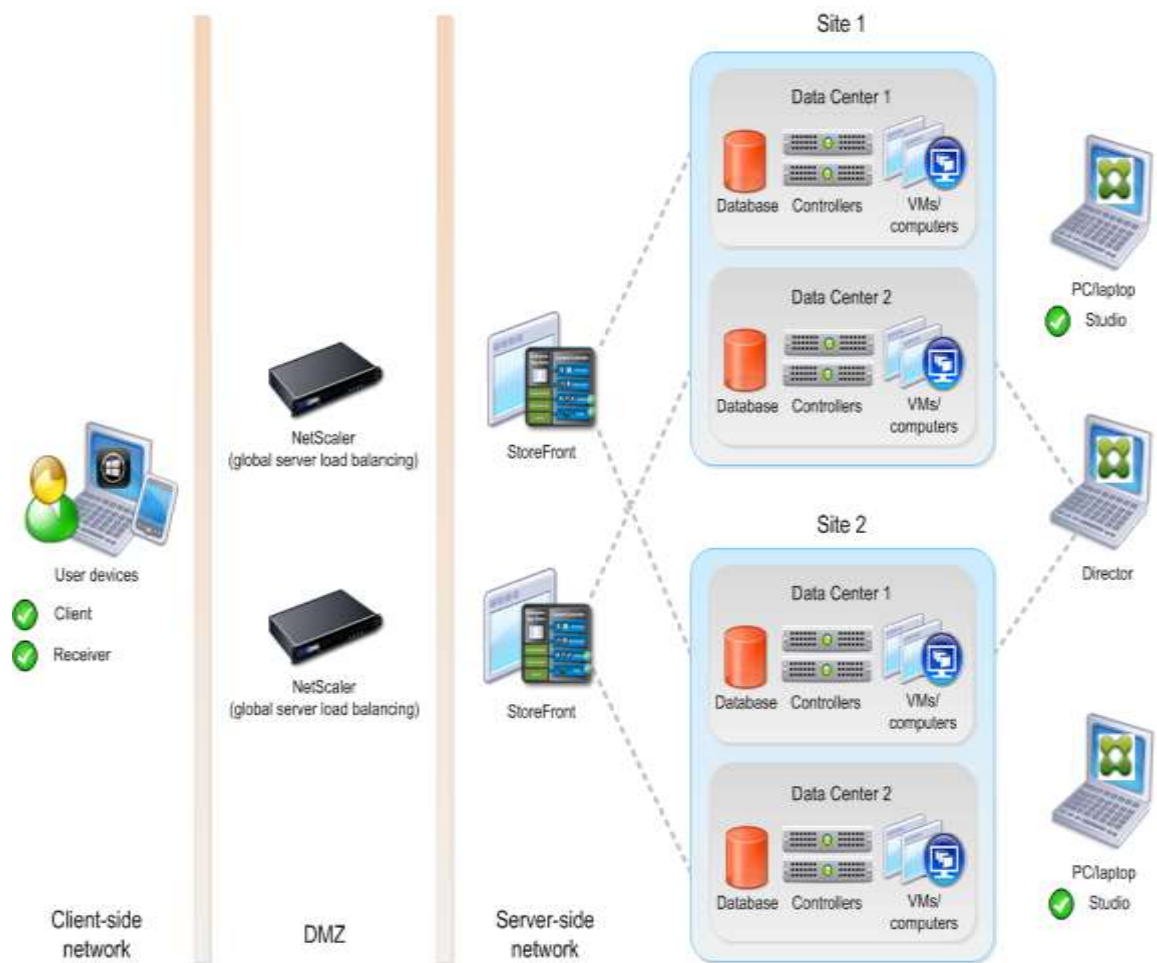
6.1.1.8. Multiple Site Configuration

If you have multiple regional sites, you can use Citrix NetScaler to direct user connections to the most appropriate site and StoreFront to deliver desktops and applications to users.

In the diagram below depicting multiple sites, each site is split into two data centers, with the database mirrored or clustered between the data centers to provide a high availability configuration. Having two sites globally, rather than just one, minimizes the amount of unnecessary WAN traffic. A separate Studio console is required to manage each site; sites cannot be managed as a single entity. You can use Director to support users across sites.

Citrix NetScaler accelerates application performance, load balances servers, increases security, and optimizes the user experience. In this example, two NetScalers are used to provide a high availability configuration. The NetScalers are configured for Global Server Load Balancing and positioned in the DMZ to provide a multi-site, fault-tolerant solution. Two Cisco blade servers host infrastructure services (AD, DNS, DHCP, Profile, SQL, Citrix XenDesktop management, and web servers).

Figure 10: XenDesktop 7 Multi-Site Architecture



6.2. EMC Storage Architecture Design

This section contains guidance on the EMC storage architecture design for high availability and data protection.

6.2.1. High Availability

The VNXe series of storage systems offer several built-in high-availability features. This high availability is provided through redundant components. If one component fails, the other one is available to back it up. The redundant components include Storage Processors (SPs), cooling fans, AC power cords, power supplies, I/O modules, and Link Controller Cards (LCCs). Network high availability is provided through link aggregation.

For network high-availability features to work, the cable on each SP needs to have the same connectivity. If Port 0 on SPA is plugged in to subnet X, Port 0 on SPB must also be plugged in to subnet X. This is necessary for both server and network failover. If a VNXe server is configured to use a port that is not connected on the peer SP, an alert is generated. Unisphere does not verify if they are plugged in

to the same subnet, but they should be, for proper failover. If you configure a server on a port that has no cable or connectivity, the traffic is routed over an SP interconnect path to the same port on the peer SP (just a single network connection for the entire system is not recommended).

For additional information about high availability in VNXe storage systems, refer to the *EMC VNXe High Availability* white paper on EMC online support (<https://support.emc.com>) VNXe Product Page.

6.2.2. Data Protection

A small to medium organization's data is one of its most valuable assets. Therefore, the company's highest priorities must include safeguarding the data. EMC VNXe series provides integrated features that meet customers' goals of business continuity and data protection. Data protection for VNXe systems is summarized in three categories: snapshots, replication, and backup. For additional information about these features, refer to the EMC VNXe Data Protection white paper on EMC online support (<https://support.emc.com>) VNXe Product Page.

6.3. Solution Validation

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

6.3.1. Configuration Topology for Scalable Citrix XenDesktop 7 Hybrid Virtual Desktop Infrastructure on Cisco Unified Computing System and EMC Storage

Figure 11: Architecture Block Diagram

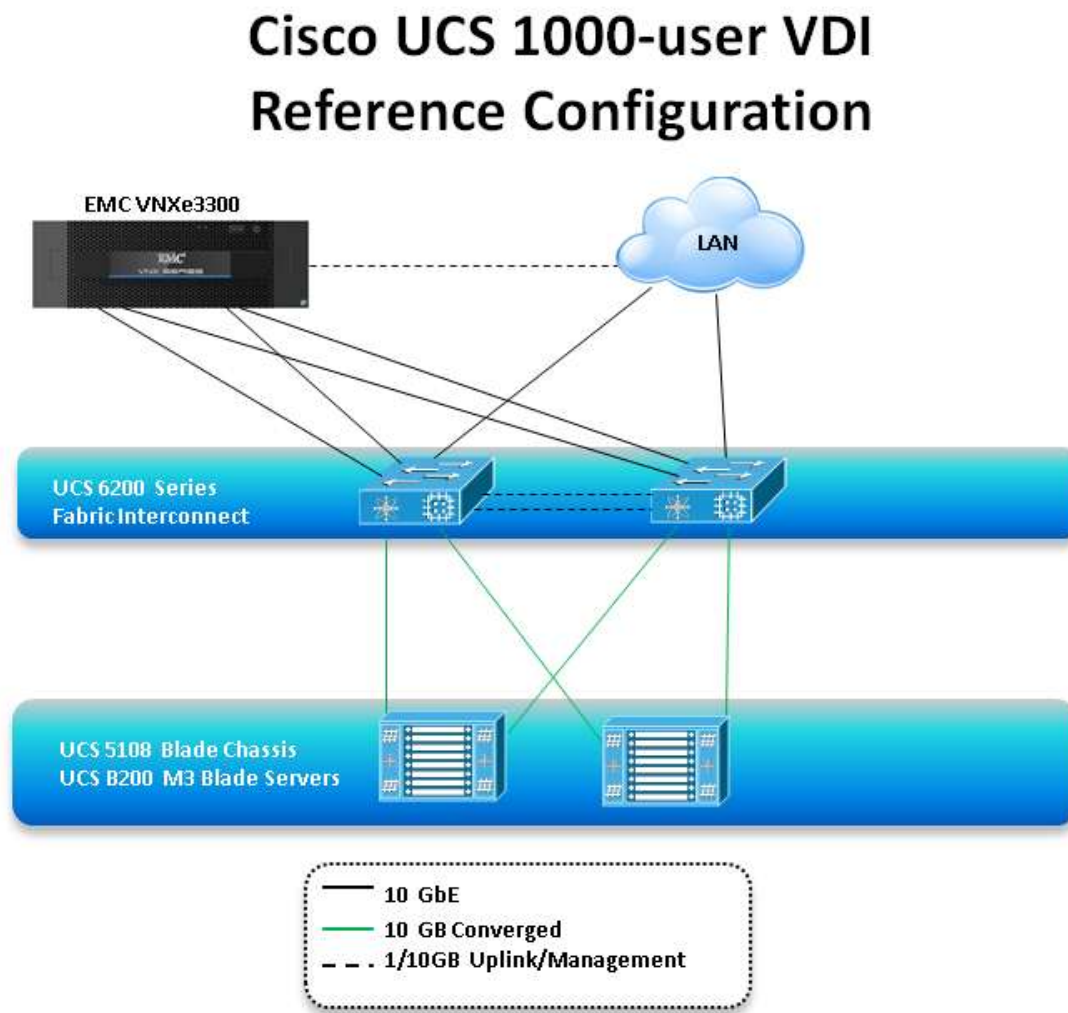
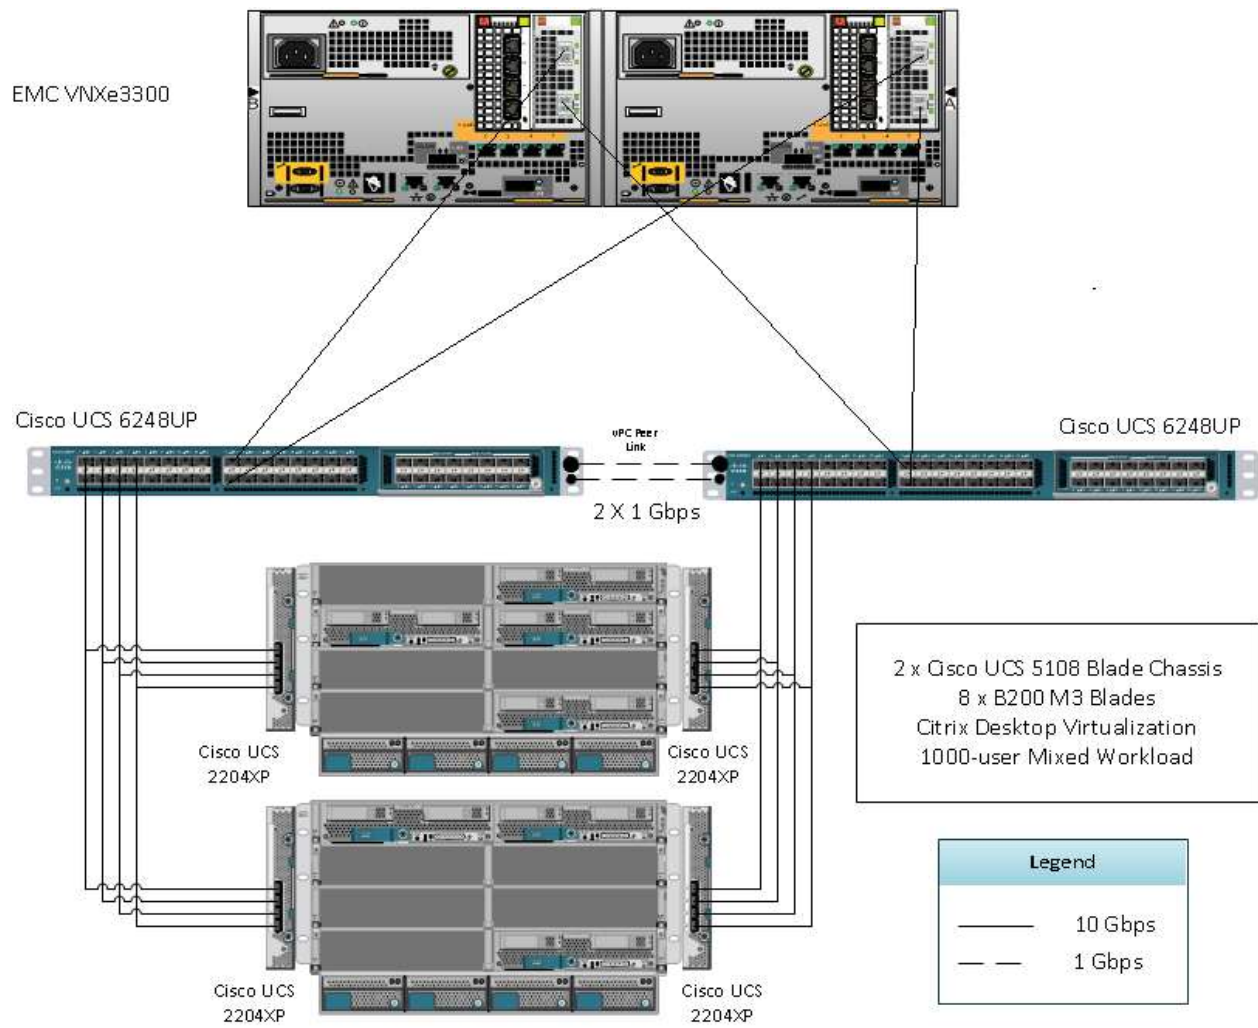


Figure 11 captures the architectural diagram for the purpose of this study. The architecture is divided into four distinct layers:

- Cisco UCS Compute Platform
- The Virtual Desktop Infrastructure that runs on Cisco UCS blade hypervisor hosts
- Network Access layer and LAN
- Storage Access Network (SAN) and EMC VNXe Storage array

Figure 12 details the physical configurations of the 500-seat and 1000-seat XenDesktop 7 environments built for this validation.

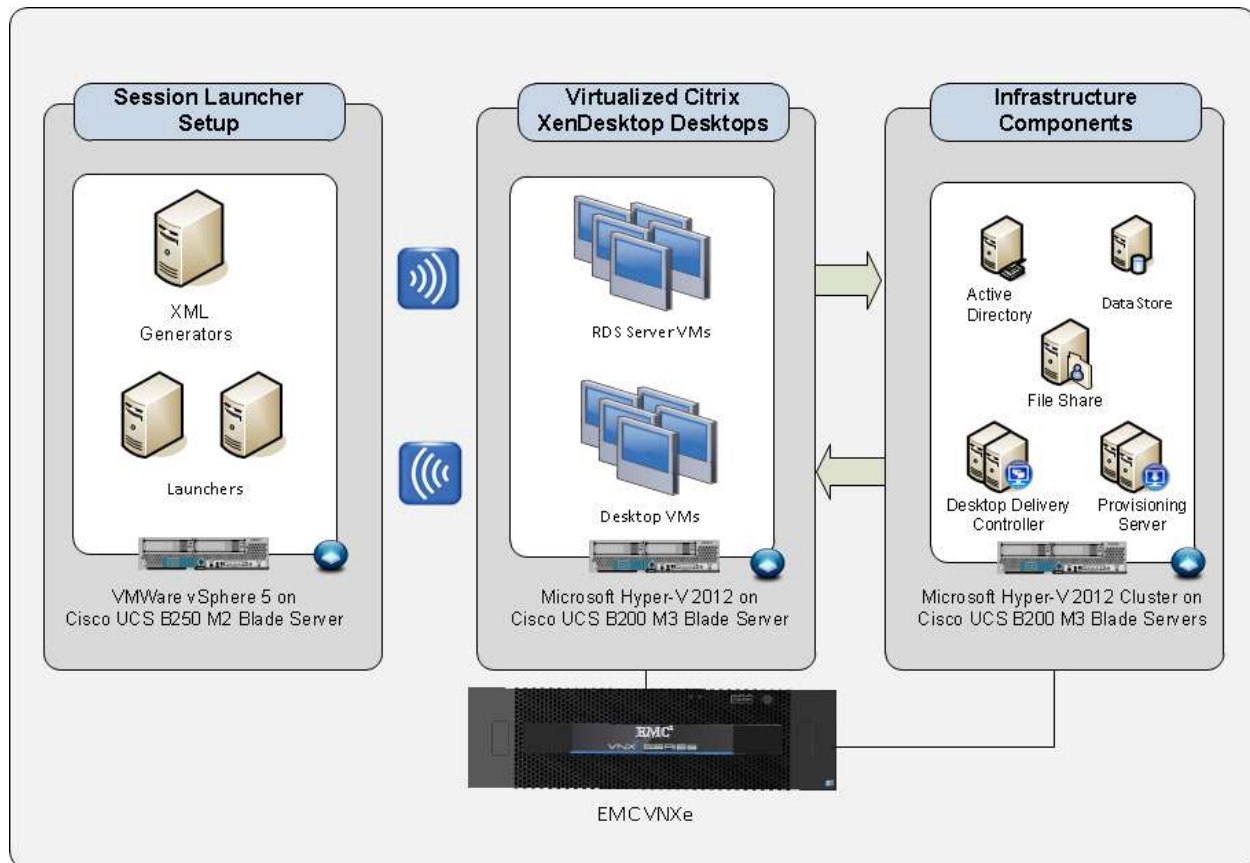
Figure 12: Detailed Architecture of Configurations



6.4. Configuration Topology for Citrix XenDesktop 7 on Cisco Unified Computing System with VNXe Storage

The diagram below provides an overview of the solution's topology.

Figure 13: Logical Architecture



6.5. Cisco UCS 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

http://www.cisco.com/en/US/docs/unified_computing/ucs/hw/chassis/install/ucs5108_install.html) and it is beyond the scope of this document. More details on each step can be found in the following documents:

- Cisco UCS CLI Configuration Guide
- http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/2.0/b_UCSM_CLI_Configuration_Guide_2_0.html
- Cisco UCS-M GUI Configuration Guide
http://www.cisco.com/en/US/partner/docs/unified_computing/ucs/sw/gui/config/guide/2.1/b_UCSM_GUI_Configuration_Guide_2_1.html

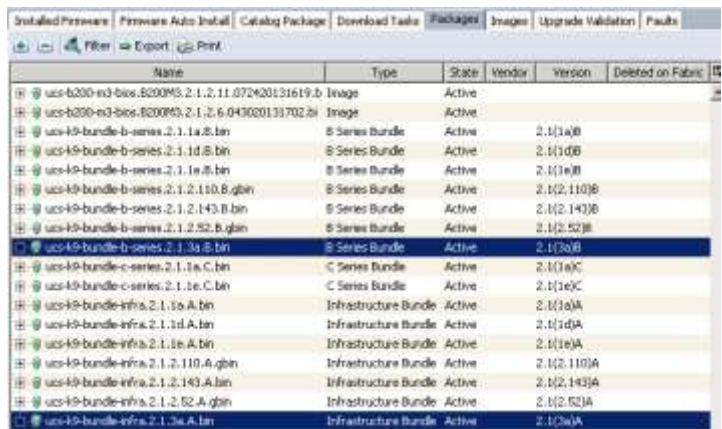
6.6. Base Cisco UCS Configuration

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

- Before beginning the Cisco UCS configuration the Fabric Interconnect must be installed and configured for remote access through a browser. To do this, bring up the Fabric interconnect and from a Serial Console connection set the IP address, gateway, and the hostname of the primary fabric interconnect. Now bring up the second fabric interconnect after connecting the dual cables between them. The second fabric interconnect automatically recognizes the primary and ask if you want to be part of the cluster, answer yes and set the IP address, gateway and the hostname.
- When this is completed, access to the FI can be done remotely. You will also configure the virtual IP address to connect to the FI, you need a total of three IP address to bring it online. You can also wire up the chassis to the FI, using either 1, 2 or 4 links per IO Module, depending on your application bandwidth requirement.
- Connect using your favorite browser to the Virtual IP and launch the UCS Manager (UCSM). The Java based UCSM will let you do everything that you could do from the CLI. We will highlight the GUI methodology here.
- Before continuing, you should use the Admin tab in the left pane, to configure logging, users and authentication, key management, communications, statistics, time zone and NTP services, and Licensing. Time zone Management (including NTP time source(s)) and uploading your license files are critical steps in this process.

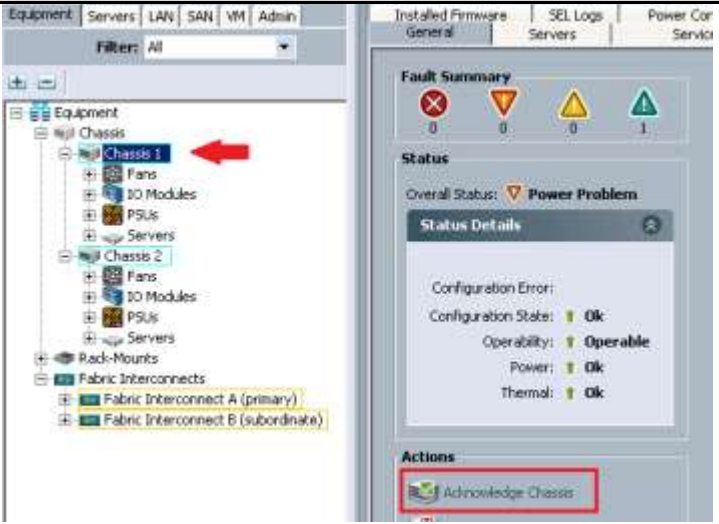
All steps outlined below in the configuration are completed from the Cisco UCS Manager user interface unless otherwise stated.

6.6.1. Firmware Update

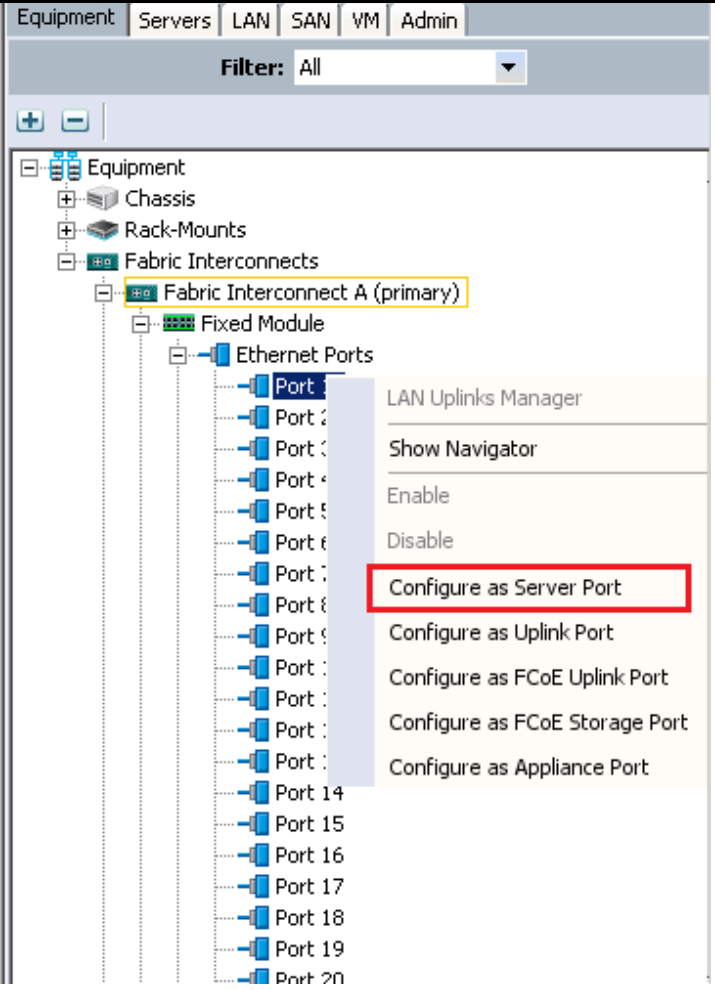
Instructions	Visual																																																																																																																																																			
<p>First check the firmware on the system and see if it is current. Visit the Cisco Software Download site to download the most current Cisco UCS Infrastructure and Cisco UCS Manager software.</p> <p>In the Equipment tab, select the Equipment node then the Firmware Management tab and the Packages sub-tab to view the packages on the system. Use the Download Tasks tab to download needed software to the FI. The firmware release used in this paper is 2.1(3a).</p> <p>If the firmware is not current, follow the installation and upgrade guide to upgrade the</p>	 <table><tr><th colspan="7">Installed Firmware Firmware Auto Install Catalog Package Download Tasks Packages Images Upgrade Validation Faults</th></tr><tr><th colspan="7">Filter Export Print</th></tr><tr><th>Name</th><th>Type</th><th>State</th><th>Vendor</th><th>Version</th><th>Deleted on Fabric</th><th></th></tr><tr><td>ucsb-6200-nd-bios-8200M3.2.1.2.11.072420131619.b</td><td>Image</td><td>Active</td><td></td><td></td><td></td><td></td></tr><tr><td>ucsb-6200-nd-bios-8200M3.2.1.2.6.043020131702.b</td><td>Image</td><td>Active</td><td></td><td></td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-b-series-2.1.1a.8.bin</td><td>B Series Bundle</td><td>Active</td><td></td><td>2.1(1a)B</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-b-series-2.1.1d.8.bin</td><td>B Series Bundle</td><td>Active</td><td></td><td>2.1(1d)B</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-b-series-2.1.1e.8.bin</td><td>B Series Bundle</td><td>Active</td><td></td><td>2.1(1e)B</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-b-series-2.1.2.110.B.gbin</td><td>B Series Bundle</td><td>Active</td><td></td><td>2.1(2.110)B</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-b-series-2.1.2.143.B.bin</td><td>B Series Bundle</td><td>Active</td><td></td><td>2.1(2.143)B</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-b-series-2.1.2.52.B.gbin</td><td>B Series Bundle</td><td>Active</td><td></td><td>2.1(2.52)B</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-b-series-2.1.3a.8.bin</td><td>B Series Bundle</td><td>Active</td><td></td><td>2.1(3a)B</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-c-series-2.1.1a.C.bin</td><td>C Series Bundle</td><td>Active</td><td></td><td>2.1(1a)C</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-c-series-2.1.1e.C.bin</td><td>C Series Bundle</td><td>Active</td><td></td><td>2.1(1e)C</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-infra-2.1.1a.A.bin</td><td>Infrastructure Bundle</td><td>Active</td><td></td><td>2.1(1a)A</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-infra-2.1.1d.A.bin</td><td>Infrastructure Bundle</td><td>Active</td><td></td><td>2.1(1d)A</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-infra-2.1.1e.A.bin</td><td>Infrastructure Bundle</td><td>Active</td><td></td><td>2.1(1e)A</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-infra-2.1.2.110.A.gbin</td><td>Infrastructure Bundle</td><td>Active</td><td></td><td>2.1(2.110)A</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-infra-2.1.2.143.A.bin</td><td>Infrastructure Bundle</td><td>Active</td><td></td><td>2.1(2.143)A</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-infra-2.1.2.52.A.gbin</td><td>Infrastructure Bundle</td><td>Active</td><td></td><td>2.1(2.52)A</td><td></td><td></td></tr><tr><td>ucsb-k9-bundle-infra-2.1.3a.A.bin</td><td>Infrastructure Bundle</td><td>Active</td><td></td><td>2.1(3a)A</td><td></td><td></td></tr></table>	Installed Firmware Firmware Auto Install Catalog Package Download Tasks Packages Images Upgrade Validation Faults							Filter Export Print							Name	Type	State	Vendor	Version	Deleted on Fabric		ucsb-6200-nd-bios-8200M3.2.1.2.11.072420131619.b	Image	Active					ucsb-6200-nd-bios-8200M3.2.1.2.6.043020131702.b	Image	Active					ucsb-k9-bundle-b-series-2.1.1a.8.bin	B Series Bundle	Active		2.1(1a)B			ucsb-k9-bundle-b-series-2.1.1d.8.bin	B Series Bundle	Active		2.1(1d)B			ucsb-k9-bundle-b-series-2.1.1e.8.bin	B Series Bundle	Active		2.1(1e)B			ucsb-k9-bundle-b-series-2.1.2.110.B.gbin	B Series Bundle	Active		2.1(2.110)B			ucsb-k9-bundle-b-series-2.1.2.143.B.bin	B Series Bundle	Active		2.1(2.143)B			ucsb-k9-bundle-b-series-2.1.2.52.B.gbin	B Series Bundle	Active		2.1(2.52)B			ucsb-k9-bundle-b-series-2.1.3a.8.bin	B Series Bundle	Active		2.1(3a)B			ucsb-k9-bundle-c-series-2.1.1a.C.bin	C Series Bundle	Active		2.1(1a)C			ucsb-k9-bundle-c-series-2.1.1e.C.bin	C Series Bundle	Active		2.1(1e)C			ucsb-k9-bundle-infra-2.1.1a.A.bin	Infrastructure Bundle	Active		2.1(1a)A			ucsb-k9-bundle-infra-2.1.1d.A.bin	Infrastructure Bundle	Active		2.1(1d)A			ucsb-k9-bundle-infra-2.1.1e.A.bin	Infrastructure Bundle	Active		2.1(1e)A			ucsb-k9-bundle-infra-2.1.2.110.A.gbin	Infrastructure Bundle	Active		2.1(2.110)A			ucsb-k9-bundle-infra-2.1.2.143.A.bin	Infrastructure Bundle	Active		2.1(2.143)A			ucsb-k9-bundle-infra-2.1.2.52.A.gbin	Infrastructure Bundle	Active		2.1(2.52)A			ucsb-k9-bundle-infra-2.1.3a.A.bin	Infrastructure Bundle	Active		2.1(3a)A		
Installed Firmware Firmware Auto Install Catalog Package Download Tasks Packages Images Upgrade Validation Faults																																																																																																																																																				
Filter Export Print																																																																																																																																																				
Name	Type	State	Vendor	Version	Deleted on Fabric																																																																																																																																															
ucsb-6200-nd-bios-8200M3.2.1.2.11.072420131619.b	Image	Active																																																																																																																																																		
ucsb-6200-nd-bios-8200M3.2.1.2.6.043020131702.b	Image	Active																																																																																																																																																		
ucsb-k9-bundle-b-series-2.1.1a.8.bin	B Series Bundle	Active		2.1(1a)B																																																																																																																																																
ucsb-k9-bundle-b-series-2.1.1d.8.bin	B Series Bundle	Active		2.1(1d)B																																																																																																																																																
ucsb-k9-bundle-b-series-2.1.1e.8.bin	B Series Bundle	Active		2.1(1e)B																																																																																																																																																
ucsb-k9-bundle-b-series-2.1.2.110.B.gbin	B Series Bundle	Active		2.1(2.110)B																																																																																																																																																
ucsb-k9-bundle-b-series-2.1.2.143.B.bin	B Series Bundle	Active		2.1(2.143)B																																																																																																																																																
ucsb-k9-bundle-b-series-2.1.2.52.B.gbin	B Series Bundle	Active		2.1(2.52)B																																																																																																																																																
ucsb-k9-bundle-b-series-2.1.3a.8.bin	B Series Bundle	Active		2.1(3a)B																																																																																																																																																
ucsb-k9-bundle-c-series-2.1.1a.C.bin	C Series Bundle	Active		2.1(1a)C																																																																																																																																																
ucsb-k9-bundle-c-series-2.1.1e.C.bin	C Series Bundle	Active		2.1(1e)C																																																																																																																																																
ucsb-k9-bundle-infra-2.1.1a.A.bin	Infrastructure Bundle	Active		2.1(1a)A																																																																																																																																																
ucsb-k9-bundle-infra-2.1.1d.A.bin	Infrastructure Bundle	Active		2.1(1d)A																																																																																																																																																
ucsb-k9-bundle-infra-2.1.1e.A.bin	Infrastructure Bundle	Active		2.1(1e)A																																																																																																																																																
ucsb-k9-bundle-infra-2.1.2.110.A.gbin	Infrastructure Bundle	Active		2.1(2.110)A																																																																																																																																																
ucsb-k9-bundle-infra-2.1.2.143.A.bin	Infrastructure Bundle	Active		2.1(2.143)A																																																																																																																																																
ucsb-k9-bundle-infra-2.1.2.52.A.gbin	Infrastructure Bundle	Active		2.1(2.52)A																																																																																																																																																
ucsb-k9-bundle-infra-2.1.3a.A.bin	Infrastructure Bundle	Active		2.1(3a)A																																																																																																																																																

<p>Cisco UCS Manager firmware. We will use UCS Policy in Service Profiles later in this document to update all Cisco UCS components in the solution.</p> <p>Note: The BIOS and Board Controller version numbers do not track the IO Module, Adapter, nor CIMC controller version numbers in the packages.</p>	
--	--

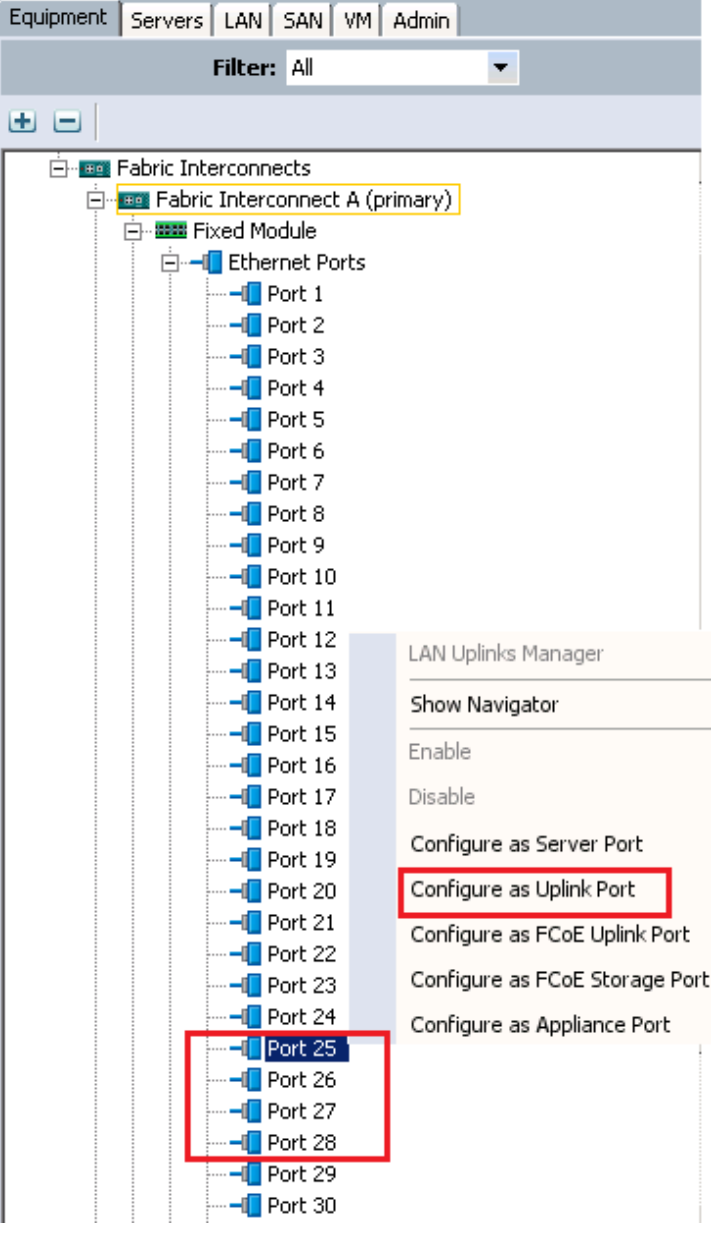
6.6.2. Acknowledge the Chassis

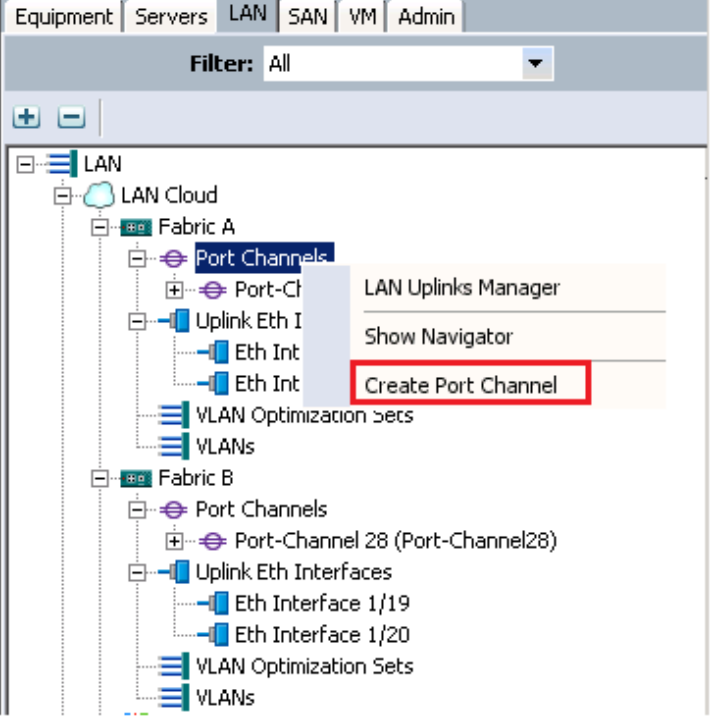
Instructions	Visual
<p>In the Equipment tab, expand the Chassis node. Then for both Chassis 1 and Chassis 2, select the General tab and click on the Acknowledge Chassis link to enable blade discovery and complete the firmware update.</p>	 <p>The screenshot shows the Cisco UCS Manager interface. On the left, the 'Equipment' tab is selected, and the 'Chassis' node is expanded. 'Chassis 1' and 'Chassis 2' are visible in the tree, with 'Chassis 1' highlighted by a red arrow. On the right, the 'General' tab for 'Chassis 1' is displayed. The 'Fault Summary' section shows four icons: a red 'X' (0), a yellow triangle (0), a green triangle (0), and a green circle (1). The 'Status' section shows 'Overall Status: Power Problem'. Below this, 'Status Details' are listed: 'Configuration Error: Configuration State: Ok', 'Operability: Operable', 'Power: Ok', and 'Thermal: Ok'. At the bottom, the 'Actions' section has a link 'Acknowledge Chassis' highlighted with a red box.</p>

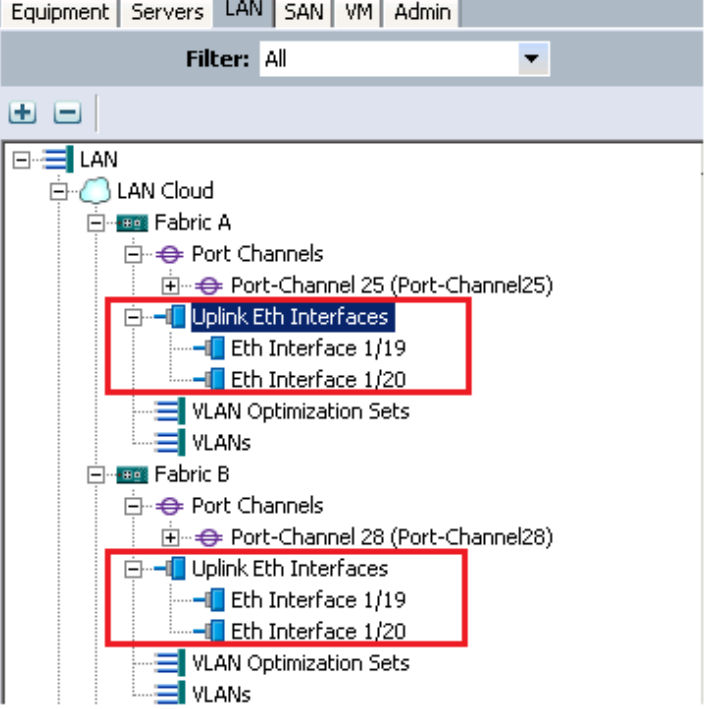
6.6.3. Server Port Configuration

Instructions	Visual
<p>In the Equipment tab, select Fabric Interconnects. For both Fabric Interconnect A and Fabric Interconnect B, right-click on the ports for connected to the IO Modules and configure them as Server Ports.</p> <p>For this validation Ports 1-8 on each Fabric Interconnect were connected to the four ports from each of the IO Modules in the two chassis.</p>	 <p>The screenshot displays the NCM interface with the 'Equipment' tab selected. The 'Filter' is set to 'All'. The tree view shows the hierarchy: Equipment > Chassis > Rack-Mounts > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module > Ethernet Ports. A right-click context menu is open over 'Port 1', with the option 'Configure as Server Port' highlighted by a red rectangle. Other menu options include 'LAN Uplinks Manager', 'Show Navigator', 'Enable', 'Disable', 'Configure as Uplink Port', 'Configure as FCoE Uplink Port', 'Configure as FCoE Storage Port', and 'Configure as Appliance Port'.</p>

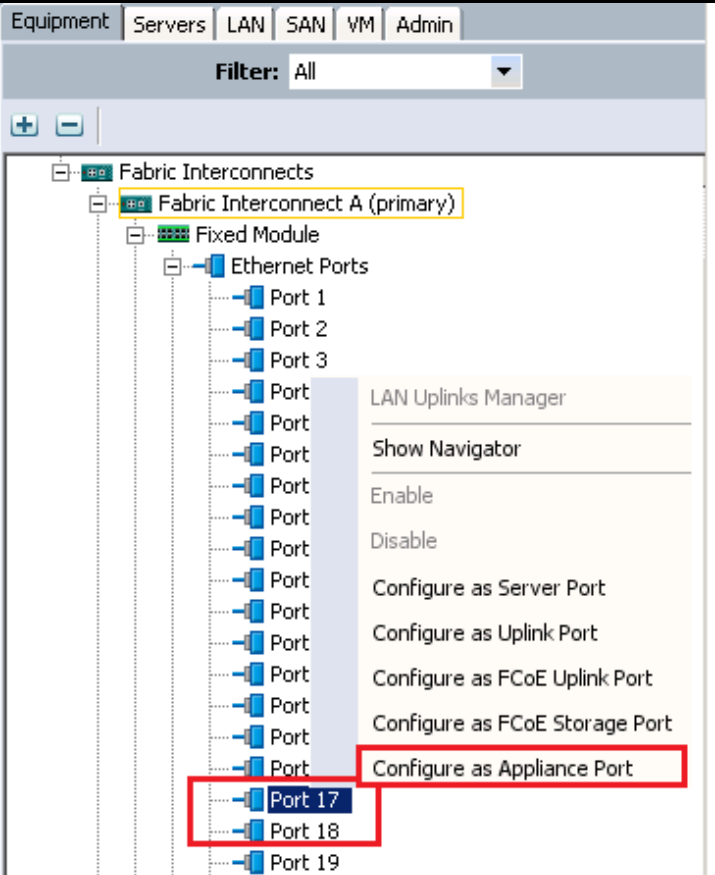
6.6.4. Uplink Port Configuration

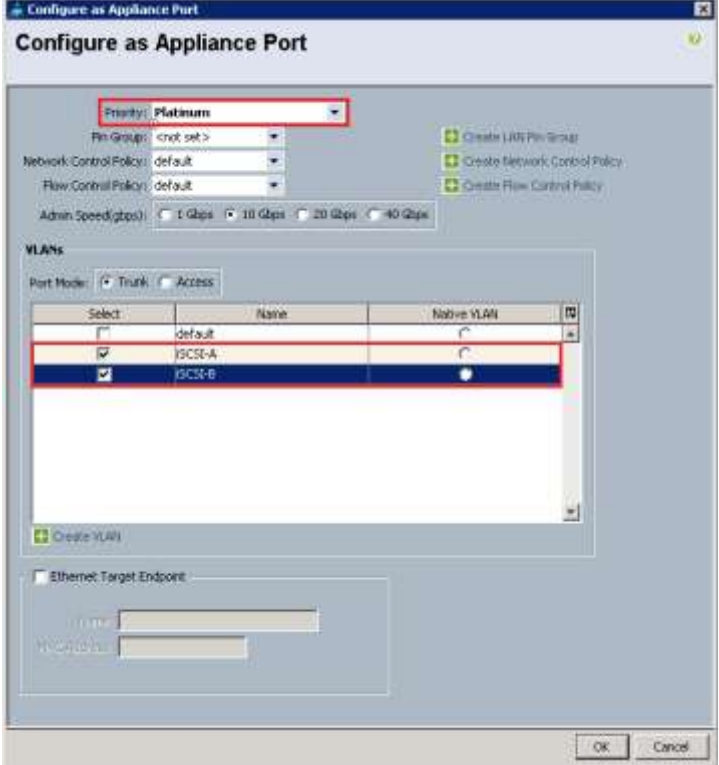
Instructions	Visual
<p>In the Equipment tab, select Fabric Interconnects. For both Fabric Interconnect A and Fabric Interconnect B, right-click on the ports connected to the core infrastructure switches and configure them as Uplink Ports.</p> <p>For this validation ports 25-26 on both Fabric Interconnects were the uplink ports to the core switch.</p>	

Instructions	Visual
<p>In the LAN tab, expand the LAN> LAN Cloud> and Fabric A (or Fabric B) nodes. Select the Port Channels node. Right-click and choose Create Port Channel to configure the port channel.</p> <p>For this validation, Port-Channel 25 (Fabric A) was configured for all Ethernet Uplink Ports 1/25 – 1/28.</p> <p>For this validation, Port-Channel 28 (Fabric B) was configured for all Ethernet Uplink Ports 1/25 – 1/28.</p>	 <p>The screenshot displays the LAN configuration interface. At the top, there are tabs for Equipment, Servers, LAN, SAN, VM, and Admin. Below the tabs is a filter dropdown set to 'All'. The main area shows a tree view of the LAN configuration. Under 'LAN', 'LAN Cloud' is expanded, showing 'Fabric A' and 'Fabric B'. 'Fabric A' is further expanded, showing 'Port Channels', 'Port-Channels', 'Uplink Eth I', 'Eth Int', and 'Eth Int'. A right-click context menu is open over the 'Port Channels' node, showing options: 'LAN Uplinks Manager', 'Show Navigator', and 'Create Port Channel' (which is highlighted with a red box). Below 'Fabric A', 'Fabric B' is also expanded, showing 'Port Channels', 'Port-Channel 28 (Port-Channel28)', 'Uplink Eth Interfaces', 'Eth Interface 1/19', 'Eth Interface 1/20', 'VLAN Optimization Sets', and 'VLANs'.</p>

Instructions	Visual
<p>In the LAN tab, expand the LAN> LAN Cloud> and Fabric A (or Fabric B) nodes. Select the Uplink Eth Interfaces node. Verify the uplink ports are available.</p>	 <p>The screenshot displays the LAN configuration interface. At the top, there are tabs for Equipment, Servers, LAN, SAN, VM, and Admin. Below the tabs is a filter dropdown set to 'All'. The main content area shows a tree view of the LAN configuration. The 'LAN' node is expanded, showing 'LAN Cloud' and 'Fabric A' (and 'Fabric B'). Under 'Fabric A', the 'Uplink Eth Interfaces' node is highlighted with a red box. This node contains two sub-items: 'Eth Interface 1/19' and 'Eth Interface 1/20'. Other nodes visible include 'Port Channels', 'Port-Channel 25 (Port-Channel25)', 'VLAN Optimization Sets', and 'VLANs'.</p>

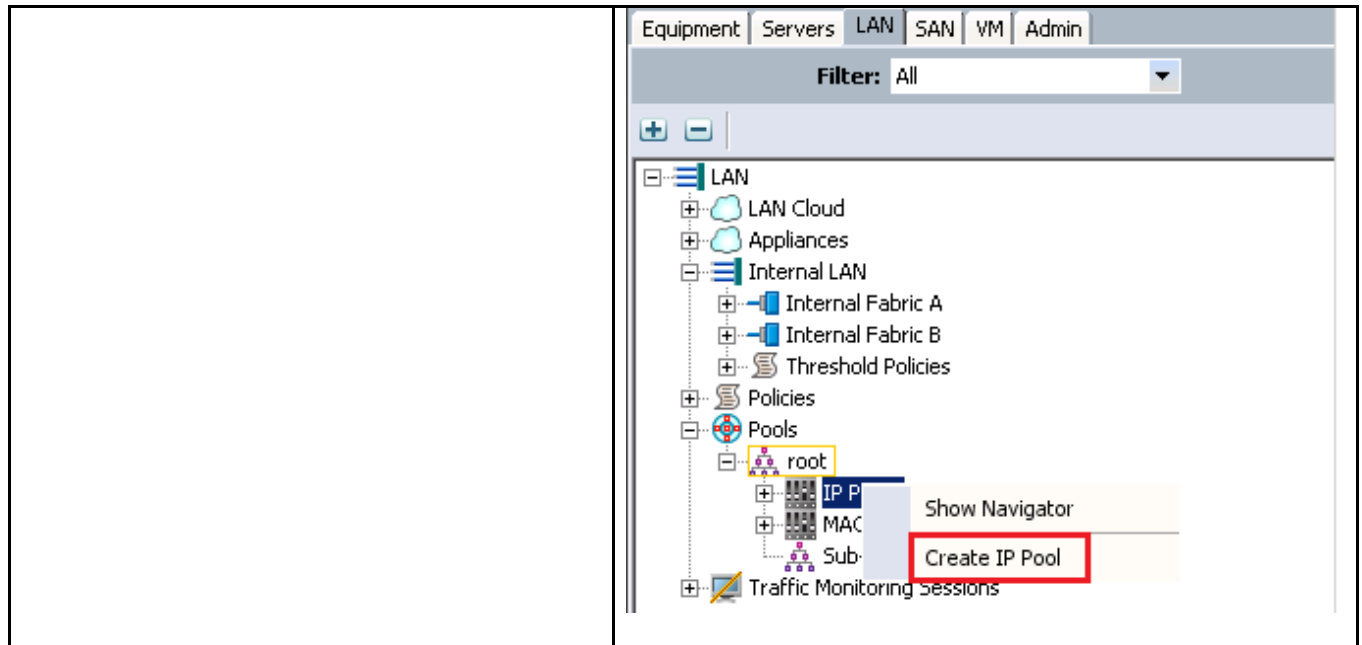
6.6.5. VNXe Appliance Port Configuration

Instructions	Visual
<p>In the Equipment tab, select Fabric Interconnects. For both Fabric Interconnect A and Fabric Interconnect B, right-click on the ports connected to the VNXe storage and configure them as Appliance Ports.</p> <p>For this validation ports 17 and 18 on both Fabric Interconnects were the uplink ports to the core switch.</p>	

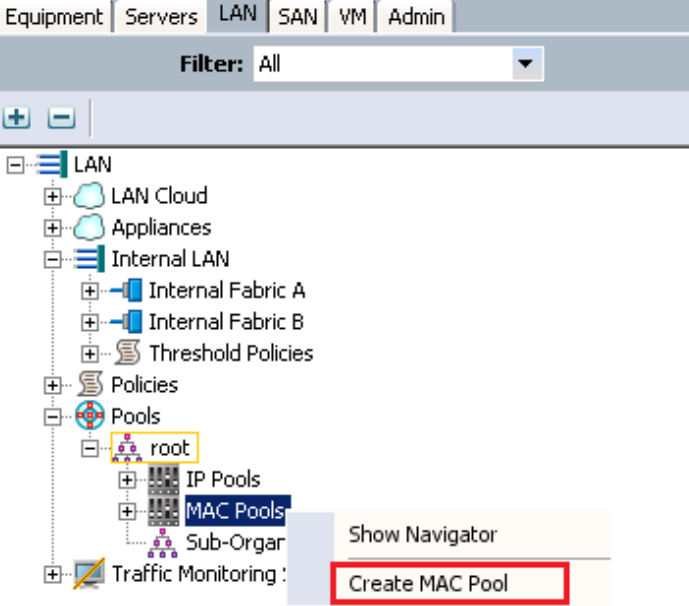
Instructions	Visual												
<p>Set Priority to Platinum</p> <p>Verify the Port Mode is set to Trunk</p> <p>Enable iSCSI-A and iSCSI-B vLANs</p>	 <p>The screenshot shows the 'Configure as Appliance Port' window. The 'Priority' is set to 'Platinum'. The 'Port Mode' is set to 'Trunk'. The 'VLANs' table has the following data:</p> <table><thead><tr><th>Select</th><th>Name</th><th>Native VLAN</th></tr></thead><tbody><tr><td><input type="checkbox"/></td><td>default</td><td>1</td></tr><tr><td><input checked="" type="checkbox"/></td><td>iSCSI-A</td><td></td></tr><tr><td><input checked="" type="checkbox"/></td><td>iSCSI-B</td><td>2</td></tr></tbody></table>	Select	Name	Native VLAN	<input type="checkbox"/>	default	1	<input checked="" type="checkbox"/>	iSCSI-A		<input checked="" type="checkbox"/>	iSCSI-B	2
Select	Name	Native VLAN											
<input type="checkbox"/>	default	1											
<input checked="" type="checkbox"/>	iSCSI-A												
<input checked="" type="checkbox"/>	iSCSI-B	2											

6.6.6. KVM IP Address Pool

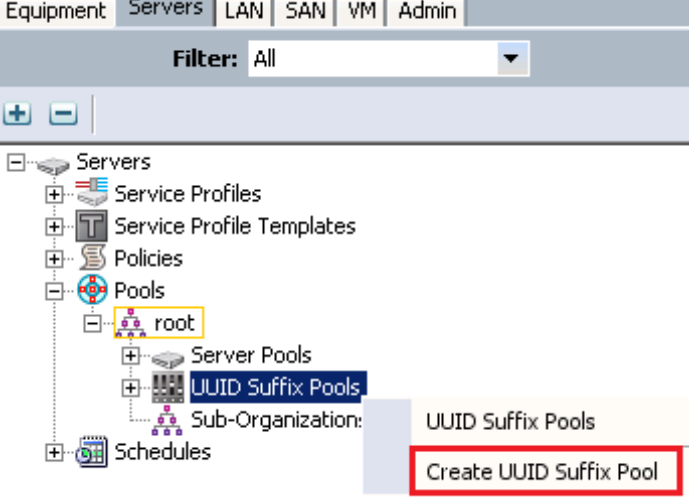
Instructions	Visual
<p>In the LAN tab, expand the Pools > root nodes. Right-click on the IP Pool node and select Create IP Pool from the context menu.</p> <p>Provide a block of IP addresses to be used for KVM access to the blades and CIMC hosts in the environment.</p> <p>This validation used the IP address range of 10.60.0.16 – 10.60.0.47 with a 255.255.255.0 subnet mask and a gateway of 10.60.0.1.</p>	



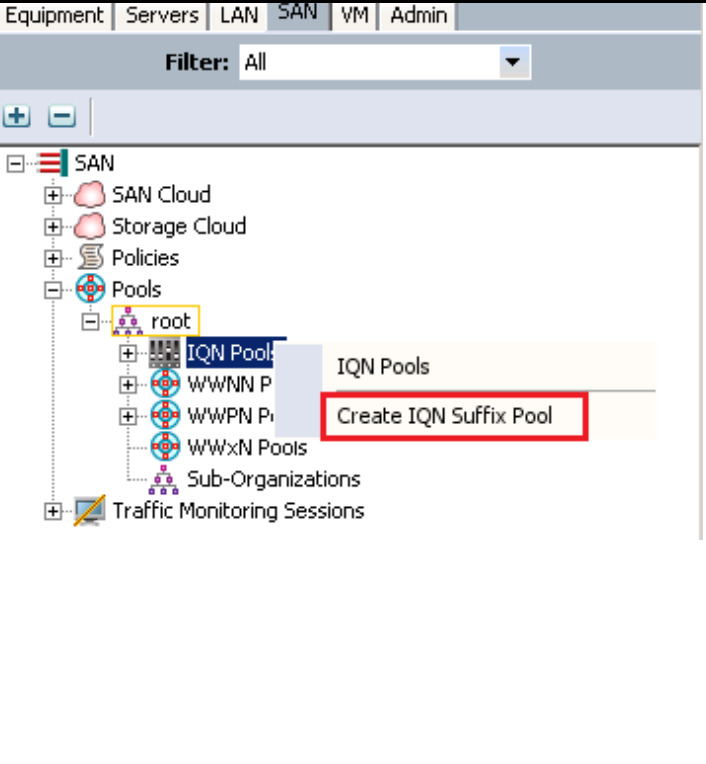
6.6.7. MAC Address Pool

Instructions	Visual
<p>In the LAN tab, expand the Pools > root nodes. Right-click on the MAC Pools node and select Create MAC Pool from the context menu.</p> <p>Provide a block of MAC addresses to be used by vNICs on the hosts. Keep the default prefix and create a MAC address pool with 256 addresses in it.</p> <p>This validation created a MAC address pool named HV-Mgmt which had the MAC address range of 00:25:b5:00:a0:00:00 – 00:25:b5:a0:00:ff.</p> <p>Note: Cisco recommends creating a new pool and not modifying the default pool.</p>	 <p>The screenshot shows the Cisco UCS Manager interface. At the top, there are tabs for Equipment, Servers, LAN, SAN, VM, and Admin. The 'LAN' tab is selected. Below the tabs is a 'Filter' dropdown set to 'All'. On the left, a navigation tree shows the hierarchy: LAN > Internal LAN > Pools > root > MAC Pools. The 'MAC Pools' node is highlighted. A context menu is open over 'MAC Pools', showing options like 'Show Navigator' and 'Create MAC Pool'. The 'Create MAC Pool' option is highlighted with a red rectangular box.</p>

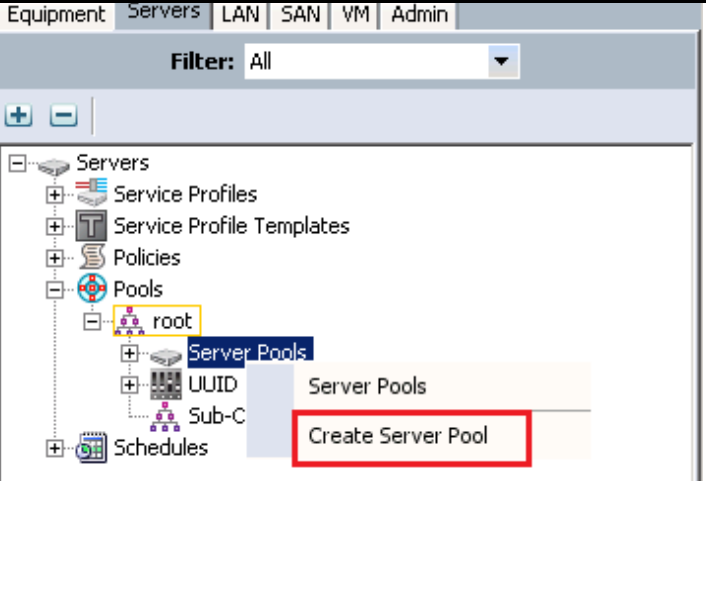
6.6.8. UUID Suffix Pool

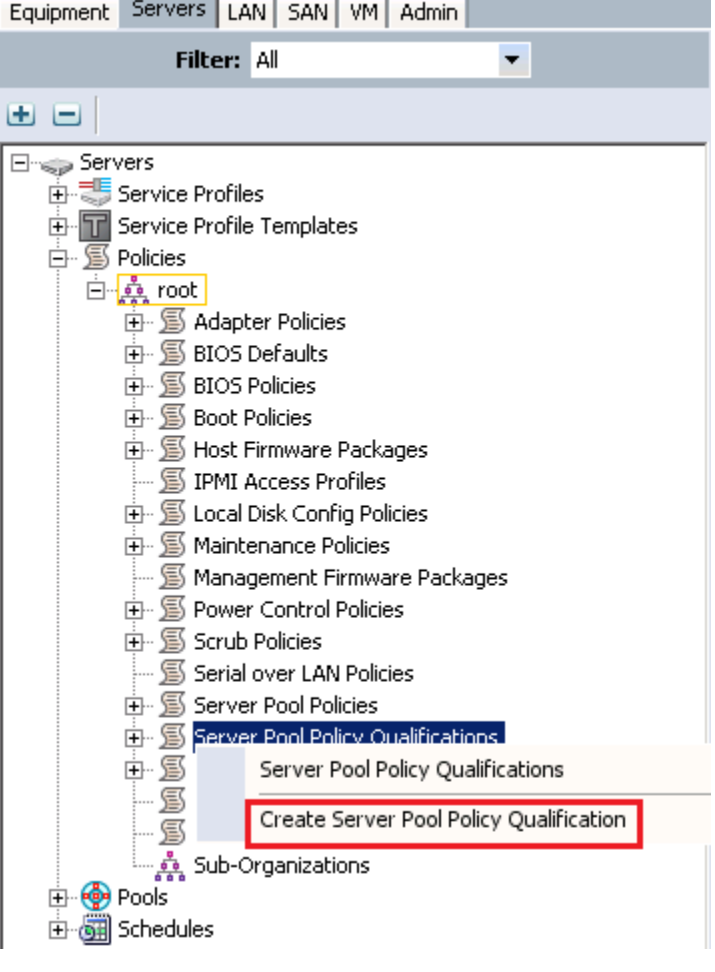
Instructions	Visual
<p>In the Servers tab, expand the Pools > root nodes. Select the UUID Suffix Pools node and right-click to choose the Create UUID Suffix Pool context menu item.</p> <p>Create a range of UUID suffixes to be used for the environment.</p> <p>For this validation a pool of 128 UUID's was generated using the default prefix.</p>	 <p>The screenshot shows the Cisco UCS Manager interface. At the top, there are tabs for Equipment, Servers, LAN, SAN, VM, and Admin. The 'Servers' tab is selected. Below the tabs is a 'Filter' dropdown set to 'All'. On the left, a navigation tree shows the hierarchy: Servers > Pools > root > UUID Suffix Pools. The 'UUID Suffix Pools' node is highlighted. A context menu is open over 'UUID Suffix Pools', showing options like 'UUID Suffix Pools' and 'Create UUID Suffix Pool'. The 'Create UUID Suffix Pool' option is highlighted with a red rectangular box.</p>

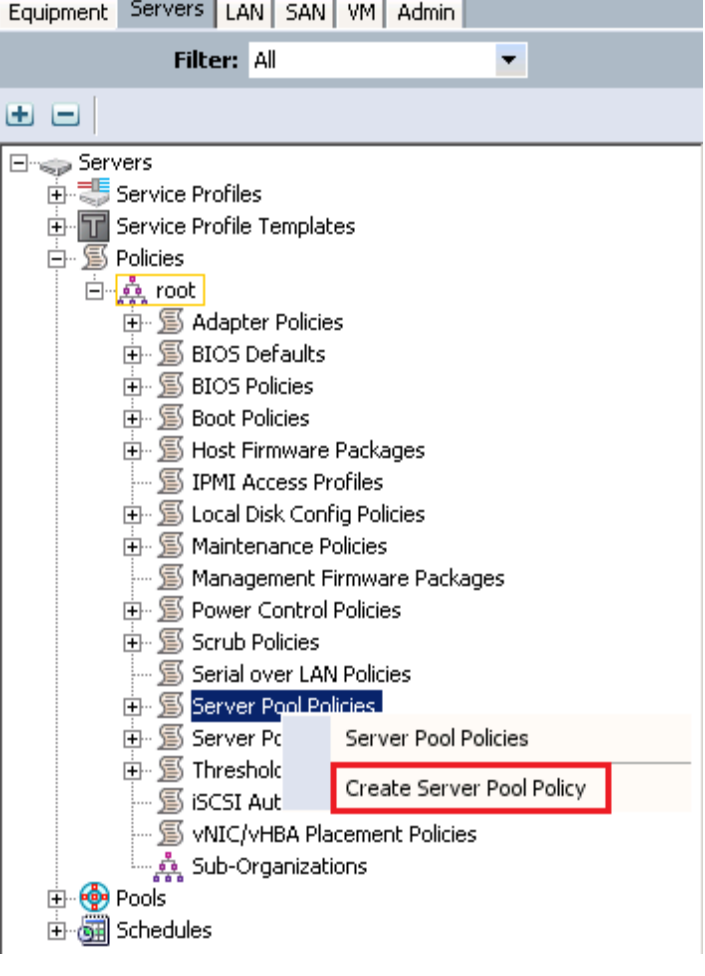
6.6.9. IQN Pool

Instructions	Visual
<p>In the SAN tab, expand the Pools > root nodes. Select the IQN Pools node and right-click to choose the Create IQN Suffix Pool context menu item.</p> <p>Create a range of 32 IQN suffixes to be used for the environment.</p> <p>For the validation, the prefix for this validation was iqn.1992-05.com.cisco and the suffix used was cvd: with a pool size was 32.</p> <p>Note: These IQN values will be used by the VNXe to setup target LUNs. They will also need to be statically assigned to the hosts later.</p>	

6.6.10. Server Pool and Related Policies (Optional)

Instructions	Visual
<p>In the Servers tab, expand Pools > root. Select Server Pools and right-click to choose Create Server Pool.</p> <p>If using a qualification policy do not add any servers to the pool at this time. If manually adding servers, skip the Qualification policy step next.</p> <p>For this validation, two server pools were created. One for the Infrastructure hosts and one for the Desktop hosts.</p>	

Instructions	Visual
<p>In the Servers tab, expand Policies > root. Select the Server Pool Policy Qualifications node and right-click to choose Create Server Pool Policy Qualification from the context menu.</p> <p>For this validation, two qualification pools were created. One for the Infrastructure hosts and one for the Desktop hosts. The qualification pools were based upon blade location, with the Infrastructure blades in slot 8 on both chassis and the remaining slots were considered Desktops hosts. However, any selection criteria may be used to distinguish between the blades for selection</p>	 <p>The screenshot shows the iDRAC web interface with the 'Servers' tab selected. The 'Policies' folder is expanded, and the 'root' folder is selected. The 'Server Pool Policy Qualifications' node is highlighted in blue. A context menu is open over this node, showing the option 'Create Server Pool Policy Qualification' in a red box. Other options in the menu include 'Server Pool Policy Qualifications' and 'Server Pool Policies'.</p>

Instructions	Visual
<p>In the Servers tab, expand Policies > root. Select the Server Pool Policies node and right-click to choose Create Server Pool Policy from the context menu.</p> <p>Use the Server Pool Policy to associate the Server Pool (Target Pool) created earlier with the Server Qualification Policy (Qualification) created earlier.</p> <p>In this validation, two Server Pool policies were created to associate the Infrastructure and Desktop Pools with their related qualification policies.</p>	 <p>The screenshot shows the 'Servers' tab with a filter set to 'All'. The 'Policies' folder is expanded, and the 'root' folder is selected. The 'Server Pool Policies' node is highlighted, and a context menu is open, showing the option 'Create Server Pool Policy' in red.</p>

6.6.11. Local Disk Configuration

Instructions	Visual
<p>In the Servers tab, expand the Policies > root nodes, then select Local Disk Config Policies. Right-click and choose Create Local Disk Configuration Policy from the context menu.</p> <p>Create a RAID0 policy for the SSD drives in the desktop hosts and a RAID1 policy for the drives in the infrastructure hosts.</p>	

EquipmentServersLANSANVMAdmin

Filter: All

Servers

Service Profiles

Service Profile Templates

Policies

root

Adapter Policies

BIOS Defaults

BIOS Policies

Boot Policies

Host Firmware Packages

IPMI Access Profiles

Local Disk Config Policies

M.

M.

Pc

Scrub Policies

Serial over LAN Policies

Server Pool Policies

Server Pool Policy Qualifications

Threshold Policies

ISCSI Authentication Profiles

vNIC/vHBA Placement Policies

Sub-Organizations

Pools

Schedules

Create Local Disk Configuration Policy

Properties

Name: RAID0

Description: RAID0 for SSD drives

Mode: RAID 0 Striped

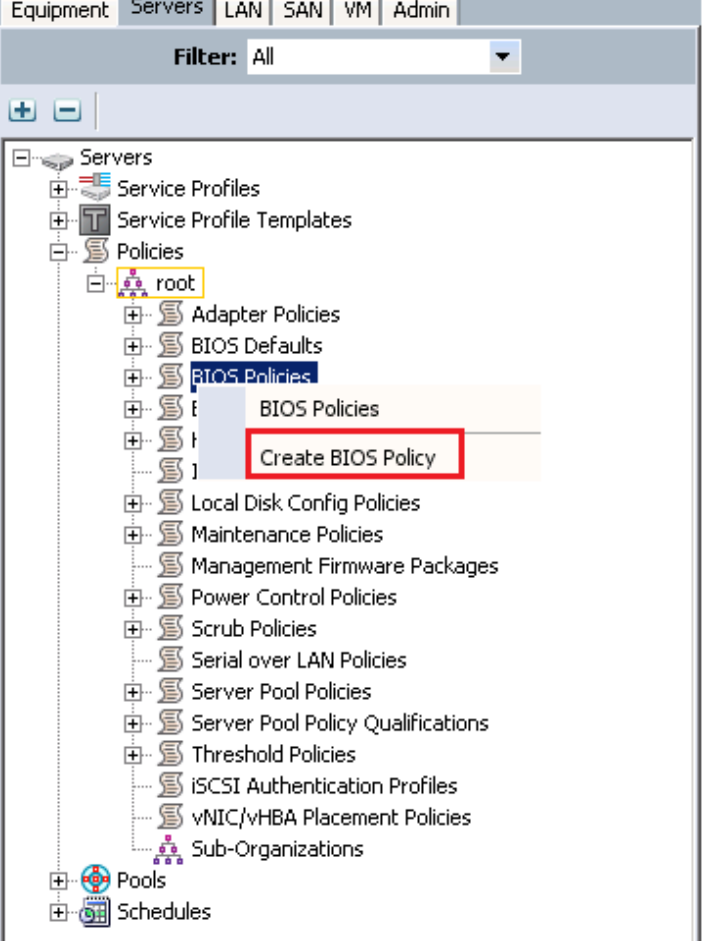
Protect Configuration: ☒

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server.

In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

6.6.12. BIOS Policy

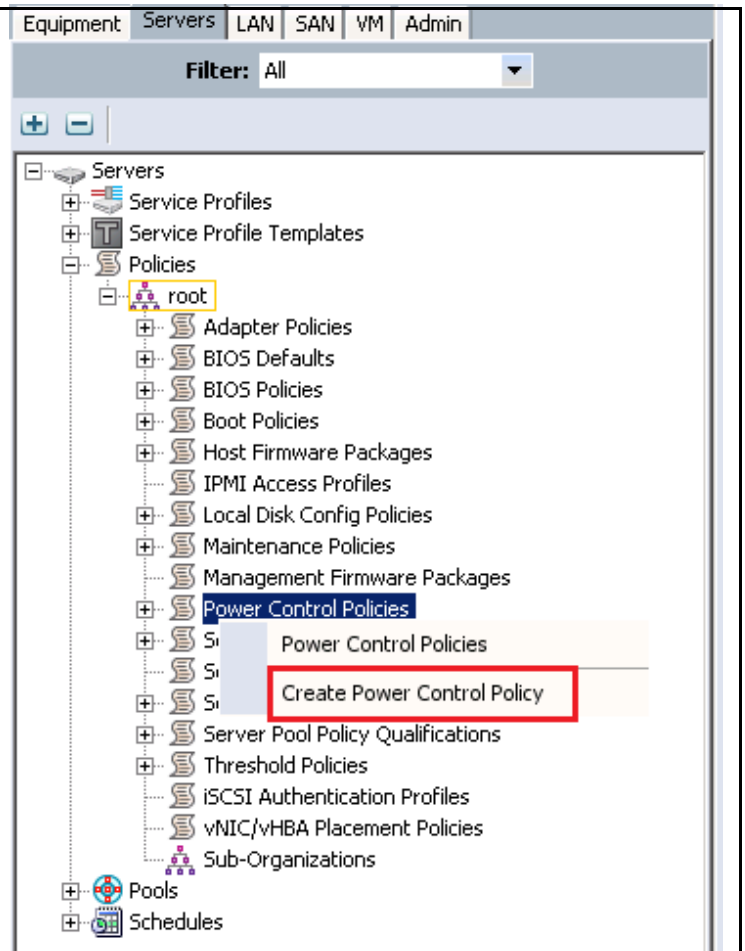
Instructions	Visual
In the Servers tab, expand the Policies > root nodes, then select BIOS Policies . Right-click and choose Create BIOS Policy from the context menu.	

<p>Verify the following are set in the BIOS policy or in the BIOS Defaults:</p> <p><u>Processor</u></p> <p>Turbo Boost: Enabled</p> <p>Enhanced Intel Speed Step: Enabled</p> <p>Hyper Threading: Enabled</p> <p>Execute Bit Disabled: Enabled</p> <p>Virtualization Technology (VT): Enabled</p> <p>Direct Cache Access: Enabled</p> <p>Processor C1E: Disabled</p> <p>CPU Performance: High-throughput</p> <p><u>Intel Directed I/O</u></p> <p>VT For Directed I/O: Enabled</p> <p><u>RAS Memory</u></p> <p>Memory RAS Config: Maximum-performance</p> <p>NUMA: Enabled</p> <p>LV DDR Mode: performance-mode</p>	 <p>The screenshot shows the 'Servers' tab in a management console. The 'Policies' folder is expanded, showing a 'root' node. A context menu is open for the 'root' node, with 'Create BIOS Policy' highlighted in red. The menu also includes options like 'Adapter Policies', 'BIOS Defaults', 'BIOS Policies', 'Local Disk Config Policies', 'Maintenance Policies', 'Management Firmware Packages', 'Power Control Policies', 'Scrub Policies', 'Serial over LAN Policies', 'Server Pool Policies', 'Server Pool Policy Qualifications', 'Threshold Policies', 'iSCSI Authentication Profiles', 'vNIC/vHBA Placement Policies', 'Sub-Organizations', 'Pools', and 'Schedules'.</p>
--	---

6.6.13. Power Control Policy

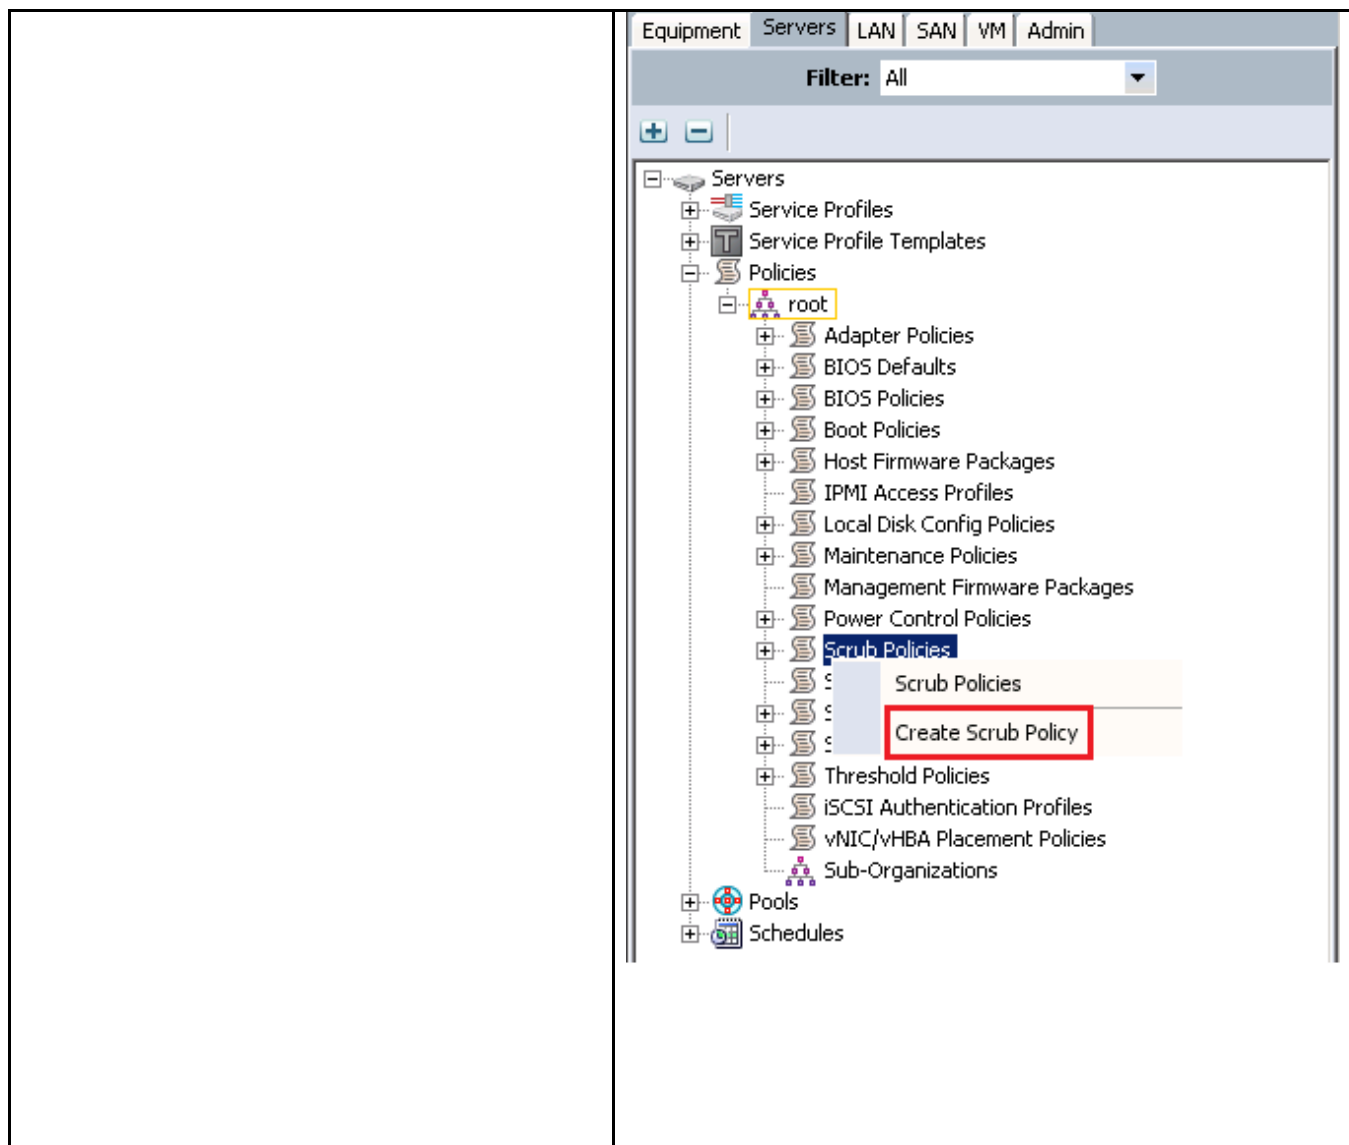
Instructions	Visual
<p>In the Servers tab, expand the Policies > root nodes, then select Power Control Policies. Right-click and choose Create Power Control Policy from the context menu.</p>	

Select the **No Cap** radio button for the policy to prevent power restrictions from limiting the performance.



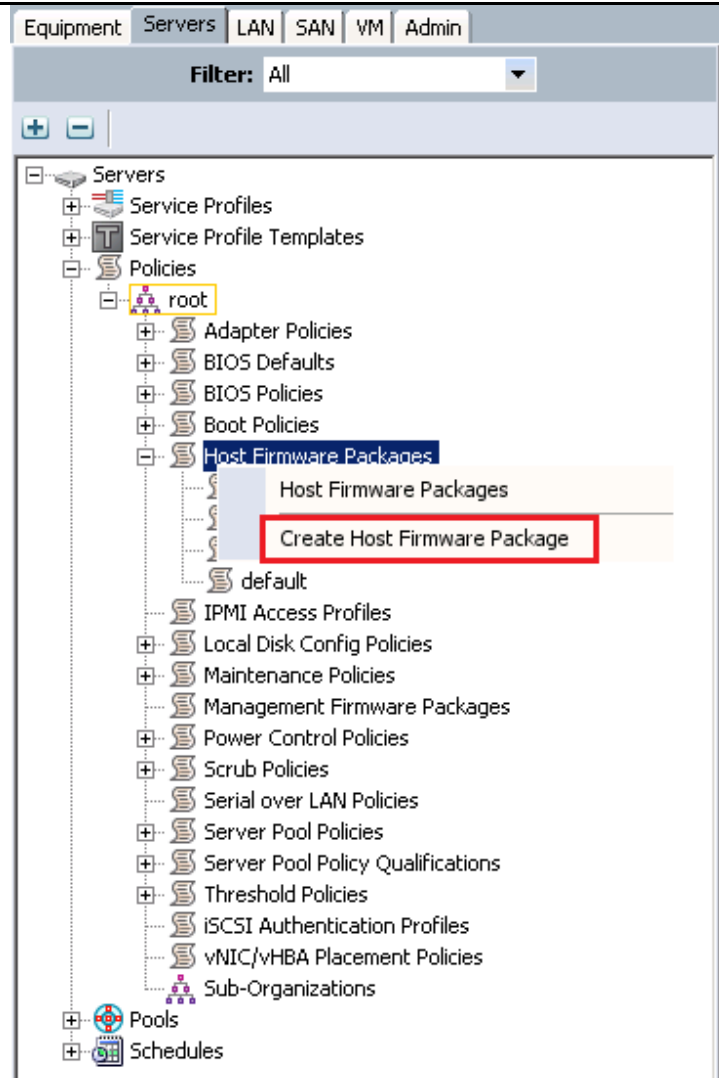
6.6.14. Scrub Policy

Instructions	Visual
<p>In the Servers tab, expand the Policies > root nodes, then select Scrub Policies. Right-click and choose Create Scrub Policy from the context menu.</p>	



6.6.15. Host Firmware Package

Instructions	Visual
<p>In the Servers tab, expand the Policies > root nodes, then select Host Firmware Packages. Right-click and choose Create Host Firmware Package from the context menu.</p> <p>Choose Advanced radio button</p>	

<p>For the Adapter tab</p> <ul style="list-style-type: none"> - select the Cisco UCS VIC 1240 card (UCSB-MLOM-40G-01) - set the firmware version to 2.1(3a) <p>For the CIMC tab</p> <ul style="list-style-type: none"> - select the UCSB-B200-M3 - set the firmware version to 2.1(3a) <p>For the BIOS tab</p> <ul style="list-style-type: none"> - select the UCSB-200-M3 - set the firmware version to B200M3.2.1.3a.0.082320131800 	 <p>The screenshot shows the Cisco UCS Manager interface. At the top, there are tabs for Equipment, Servers, LAN, SAN, VM, and Admin. Below these is a 'Filter' dropdown set to 'All'. The main area displays a tree structure under 'Servers'. The 'Policies' folder is expanded, showing a 'root' node. Under 'root', several policy folders are listed, including 'Host Firmware Packages'. This folder is expanded, showing a 'Create Host Firmware Package' button highlighted with a red rectangle. Other visible folders include 'Adapter Policies', 'BIOS Defaults', 'BIOS Policies', 'Boot Policies', 'default', 'IPMI Access Profiles', 'Local Disk Config Policies', 'Maintenance Policies', 'Management Firmware Packages', 'Power Control Policies', 'Scrub Policies', 'Serial over LAN Policies', 'Server Pool Policies', 'Server Pool Policy Qualifications', 'Threshold Policies', 'iSCSI Authentication Profiles', 'vNIC/vHBA Placement Policies', and 'Sub-Organizations'. At the bottom of the tree are 'Pools' and 'Schedules'.</p>
--	--

6.6.16. QoS Policies

Cisco Unified Computing System provides different system class of service to implement quality of service including:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames.

Applications like the Cisco Unified Computing System and other time sensitive applications have to adhere to a strict QoS for optimal performance.

6.6.16.1. System Class Configuration

Systems Class is the global operation where entire system interfaces are with defined QoS rules.

- By default system has Best Effort Class and FCoE Class.
- Best effort is equivalent in MQC terminology as “match any”
 - FCoE is special Class define for FCoE traffic. In MQC terminology “match cos 3”
- System class allowed with 4 more users define class with following configurable rules.
 - CoS to Class Map
 - Weight: Bandwidth
 - Per class MTU
 - Property of Class (Drop v/s no drop)
- Max MTU per Class allowed is 9216.
- Via Cisco UCS we can map one CoS value to particular class.
- Apart from FCoE class there can be only one more class can be configured as no-drop property.
- Weight can be configured based on 0 to 10 numbers. Internally system will calculate the bandwidth based on following equation (there will be rounding off the number).

$$\text{➤ \% b/w shared of given Class} = \frac{(\text{Weight of the given priority} * 100)}{\text{Sum of weights of all priority}}$$

6.6.16.2. Cisco UCS System Class Configuration

Cisco Unified Computing System defines user class names as follows.

- Platinum
- Gold
- Silver
- Bronze

Table 6: Name Table Map between Cisco Unified Computing System and the NXOS

Cisco UCS Names	NXOS Names
Best effort	Class-default
Platinum	Class-Platinum
Gold	Class-Gold
Silver	Class-Silver
Bronze	Class-Bronze

Table 7: Class to CoS Map by default in Cisco Unified Computing System

Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	Match any
Platinum	5

Gold	4
Silver	2
Bronze	1

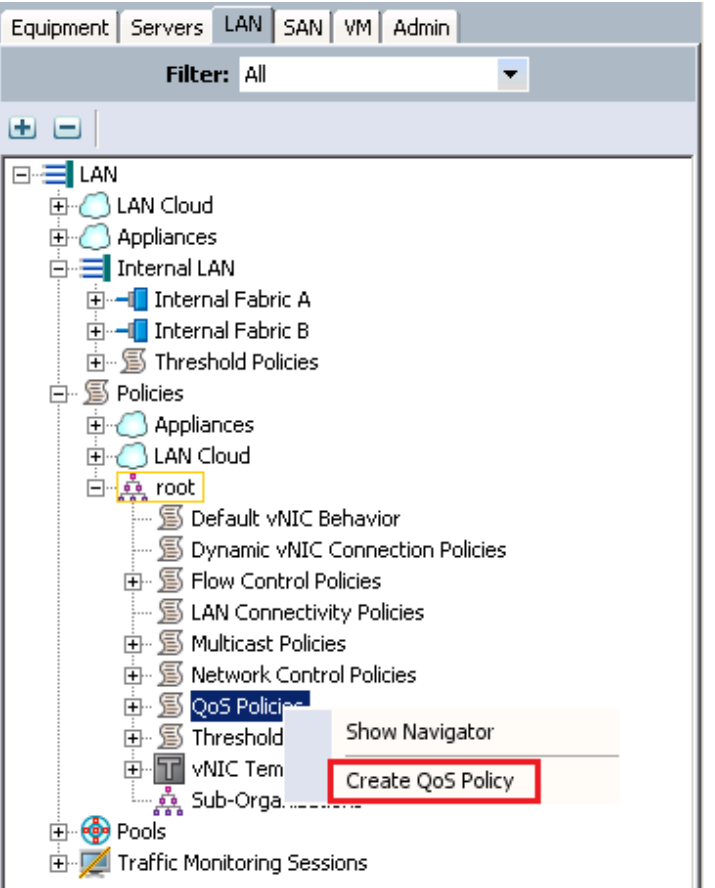
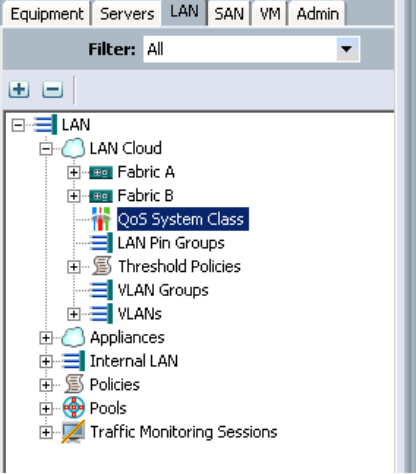
6.6.17. Steps to Enable QOS on the Cisco Unified Computing System

In this validation, we utilized four Cisco UCS QoS System Classes to priorities four types of traffic in the infrastructure:

Table 8: QoS Priority to vNIC and VLAN Mapping

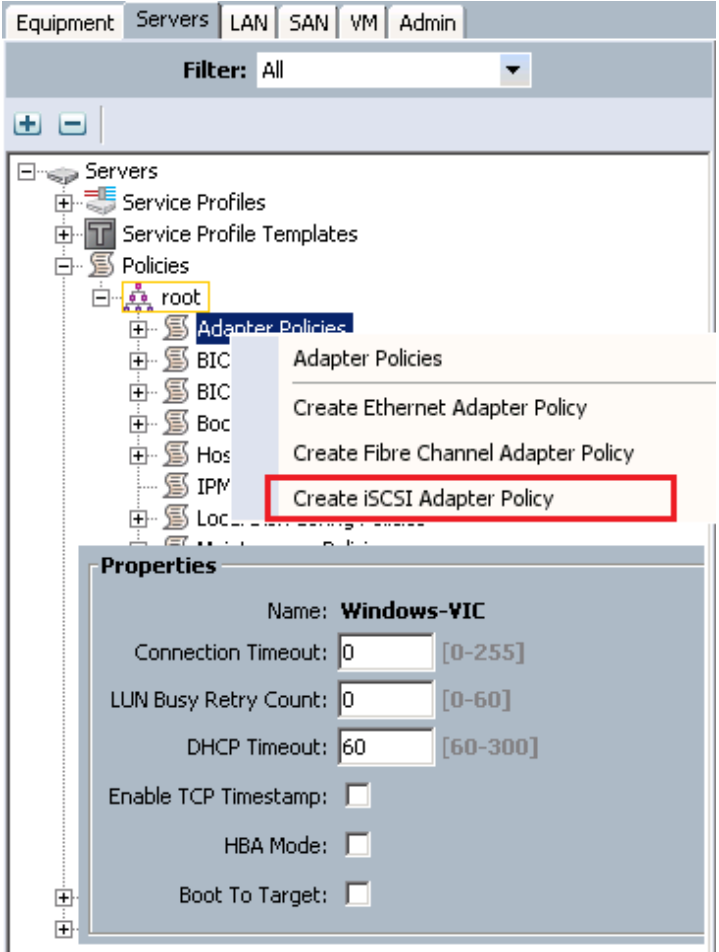
Cisco UCS Qos Priority	vNIC Assignment	MTU Size	VLAN Supported
Platinum	eth6, eth7	9000	65, 66 (iSCSI-A, iSCSI-B)
Gold	eth3	9000	62 (PVS - VDI)
Silver	eth0, eth1	Normal	60, 61 (Management, Infrastructure)
Bronze	eth4, eth5	9000	63, 64 (CSV, Live Migration)

Configure Platinum, Gold, Silver and Bronze policies by checking the enabled box. For the Platinum, Gold, and Bronze Policies, configure Jumbo Frames in the MTU column. To configure the QoS policies follow these steps:

Instructions	Visual
<p>In the LAN tab, expand the Policies > root nodes, then select QoS Policies. Right-click and choose Create QoS Policy from the context menu.</p> <p>Create four QoS policies with their corresponding priorities: Bronze, Silver, Gold, Platinum. The different QoS priorities will be used for the following vLANs:</p> <p>Bronze: CSV, LMIGR</p> <p>Silver: INFRA, MGMT</p> <p>Gold: PVS-VDI</p> <p>Platinum: iSCSI-A, iSCSI-B</p>	 <p>The screenshot shows the 'LAN' tab selected in the top navigation bar. The left sidebar displays a tree structure under 'LAN'. The 'Policies' node is expanded, and the 'root' node is selected. A right-click context menu is visible over the 'root' node, with the 'Create QoS Policy' option highlighted in a red box. Other options in the menu include 'Show Navigator'.</p>
<p>In the LAN tab, expand LAN Cloud</p> <p>Select the QoS System Class node.</p> <p>Select the Enabled check box for any QoS levels not already enabled.</p> <p>Set the MTU value for the Bronze, Gold, and Platinum QoS levels to 9000.</p> <p>Click Save Changes to save the changes.</p>	 <p>The screenshot shows the 'LAN' tab selected. The left sidebar displays a tree structure under 'LAN'. The 'LAN Cloud' node is expanded, and the 'QoS System Class' node is selected. The 'Enabled' checkbox for this node is checked. The 'MTU' value is set to 9000.</p>

Instructions	Visual																																																								
	<div><div>General</div><div>Events</div><div>QoS</div></div> <table><thead><tr><th>Priority</th><th>Enabled</th><th>CoS</th><th>Packet Drop</th><th>Weight</th><th>Weight (%)</th><th>MTU</th><th>Multicast Optimized</th></tr></thead><tbody><tr><td>Platinum</td><td><input checked="" type="checkbox"/></td><td>5</td><td><input checked="" type="checkbox"/></td><td>10</td><td>22</td><td>9000</td><td><input checked="" type="checkbox"/></td></tr><tr><td>Gold</td><td><input checked="" type="checkbox"/></td><td>4</td><td><input checked="" type="checkbox"/></td><td>9</td><td>20</td><td>9000</td><td><input checked="" type="checkbox"/></td></tr><tr><td>Silver</td><td><input checked="" type="checkbox"/></td><td>3</td><td><input checked="" type="checkbox"/></td><td>8</td><td>18</td><td>normal</td><td><input checked="" type="checkbox"/></td></tr><tr><td>Bronze</td><td><input checked="" type="checkbox"/></td><td>1</td><td><input checked="" type="checkbox"/></td><td>7</td><td>15</td><td>9000</td><td><input checked="" type="checkbox"/></td></tr><tr><td>Best Effort</td><td><input checked="" type="checkbox"/></td><td>Any</td><td><input checked="" type="checkbox"/></td><td>5</td><td>11</td><td>normal</td><td><input checked="" type="checkbox"/></td></tr><tr><td>Fibre Channel</td><td><input checked="" type="checkbox"/></td><td>3</td><td><input checked="" type="checkbox"/></td><td>5</td><td>14</td><td>N/A</td><td><input checked="" type="checkbox"/></td></tr></tbody></table>	Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized	Platinum	<input checked="" type="checkbox"/>	5	<input checked="" type="checkbox"/>	10	22	9000	<input checked="" type="checkbox"/>	Gold	<input checked="" type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	20	9000	<input checked="" type="checkbox"/>	Silver	<input checked="" type="checkbox"/>	3	<input checked="" type="checkbox"/>	8	18	normal	<input checked="" type="checkbox"/>	Bronze	<input checked="" type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	15	9000	<input checked="" type="checkbox"/>	Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	11	normal	<input checked="" type="checkbox"/>	Fibre Channel	<input checked="" type="checkbox"/>	3	<input checked="" type="checkbox"/>	5	14	N/A	<input checked="" type="checkbox"/>
Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized																																																		
Platinum	<input checked="" type="checkbox"/>	5	<input checked="" type="checkbox"/>	10	22	9000	<input checked="" type="checkbox"/>																																																		
Gold	<input checked="" type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	20	9000	<input checked="" type="checkbox"/>																																																		
Silver	<input checked="" type="checkbox"/>	3	<input checked="" type="checkbox"/>	8	18	normal	<input checked="" type="checkbox"/>																																																		
Bronze	<input checked="" type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	15	9000	<input checked="" type="checkbox"/>																																																		
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	11	normal	<input checked="" type="checkbox"/>																																																		
Fibre Channel	<input checked="" type="checkbox"/>	3	<input checked="" type="checkbox"/>	5	14	N/A	<input checked="" type="checkbox"/>																																																		

6.6.18. iSCSI Adapter Policy

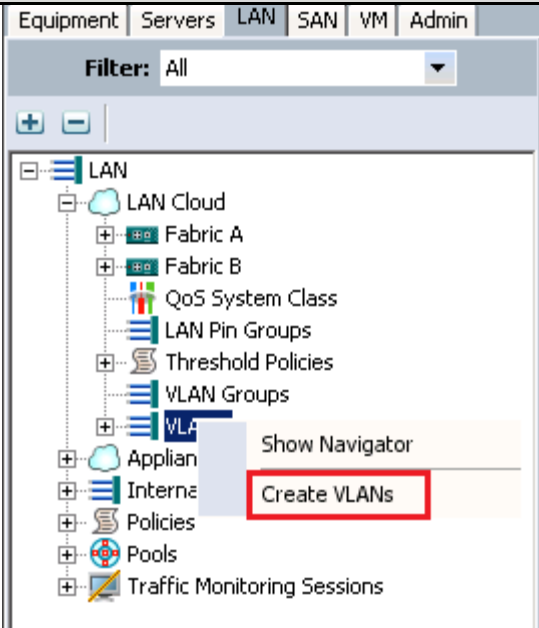

Instructions	Visual
<p>In the Servers tab, expand the Policies > root nodes, then select Adapter Policies. Right-click and choose Create iSCSI Adapter Policy from the context menu.</p>	 <p>Equipment Servers LAN SAN VM Admin</p> <p>Filter: All</p> <p>Servers</p> <ul style="list-style-type: none"> Service Profiles Service Profile Templates Policies <ul style="list-style-type: none"> root <ul style="list-style-type: none"> Adapter Policies <ul style="list-style-type: none"> Adapter Policies Create Ethernet Adapter Policy Create Fibre Channel Adapter Policy Create iSCSI Adapter Policy BIC BIC Boc Hos IPM Loc <p>Properties</p> <p>Name: Windows-VIC</p> <p>Connection Timeout: 0 [0-255]</p> <p>LUN Busy Retry Count: 0 [0-60]</p> <p>DHCP Timeout: 60 [60-300]</p> <p>Enable TCP Timestamp: <input type="checkbox"/></p> <p>HBA Mode: <input type="checkbox"/></p> <p>Boot To Target: <input type="checkbox"/></p>

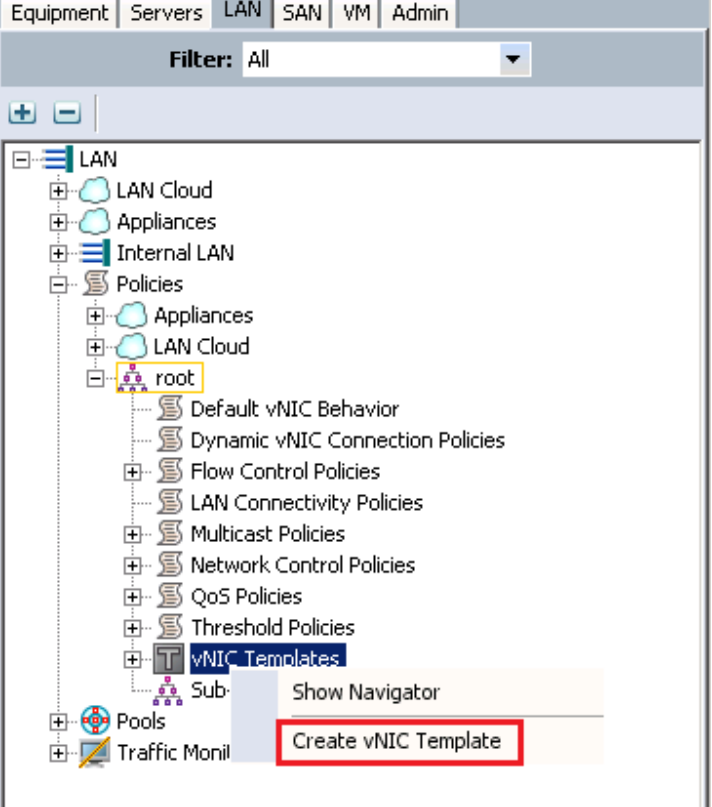
6.6.19. VLANs and vNIC Templates

In addition, to control network traffic in the infrastructure and assure priority to high value traffic, virtual LANs (VLANs) were created on the Nexus 5548s, on the Cisco UCS Manager (Fabric Interconnects,) and on the Nexus 1000V Virtual Switch Modules in the on the VDI Hosts. The virtual machines in the environment used the VLANs depending on their role in the system.

Note: A total of seven Virtual LANs were utilized for the project as defined above in section 4.4.1 VLAN.

VLANs are configured in Cisco UCS Manager on the LAN tab, LAN\VLANs node in the left pane of Cisco UCS Manager

Instructions	Visual
<p>In the LAN tab, expand LAN Cloud, and select the VLANs node. Right-click and choose Create VLANs.</p> <p>Repeat for each of the VLANs required for the environment.</p> <p>For this validation, the following VLANs were created:</p> <p>VLAN NAME: VLAN ID – Purpose</p> <p>MGMT: 60 – Management</p> <p>INFRA: 61 – Primary Infrastructure</p> <p>PVS-VDI: 62 – Virtual Desktops</p> <p>CSV: 63 – Cluster Shared Volumes</p> <p>LMIGR: 64 – Live Migration</p> <p>iSCSI-A:65 – iSCSI Fabric A</p> <p>iSCSI-B:66 – iSCSI Fabric B</p> <p>iSCSI-Null: 999 – Windows Installation*</p>	 

Instructions	Visual
<p>Fabric is set to Common/Global</p> <p>Sharing Type is set to None</p> <p>*iSCSI-Null is a special VLAN used temporarily to work around a known issue with Microsoft Windows Server 2012 install when multiple iSCSI paths exist but no multi-path software is running.</p>	
<p>In the LAN tab, expand the Policies > root nodes. Select the vNIC Templates node. Right-click and select the Create vNIC Template context menu item.</p>	 <p>The screenshot shows the Cisco UCS Manager interface with the 'LAN' tab selected. The navigation tree on the left shows the path: LAN > Policies > root > vNIC Templates. The 'vNIC Templates' node is selected, and a context menu is open with the 'Create vNIC Template' option highlighted in a red box. Other options in the menu include 'Show Navigator'.</p>

One of the unique value propositions for Cisco Unified Computing System with respect to end-to-end QoS is the ability to tag the traffic at the edge. For example, dedicate a VLAN for the EMC storage, configure Platinum policy with Jumbo frames and get an end-to-end QoS and performance guarantees

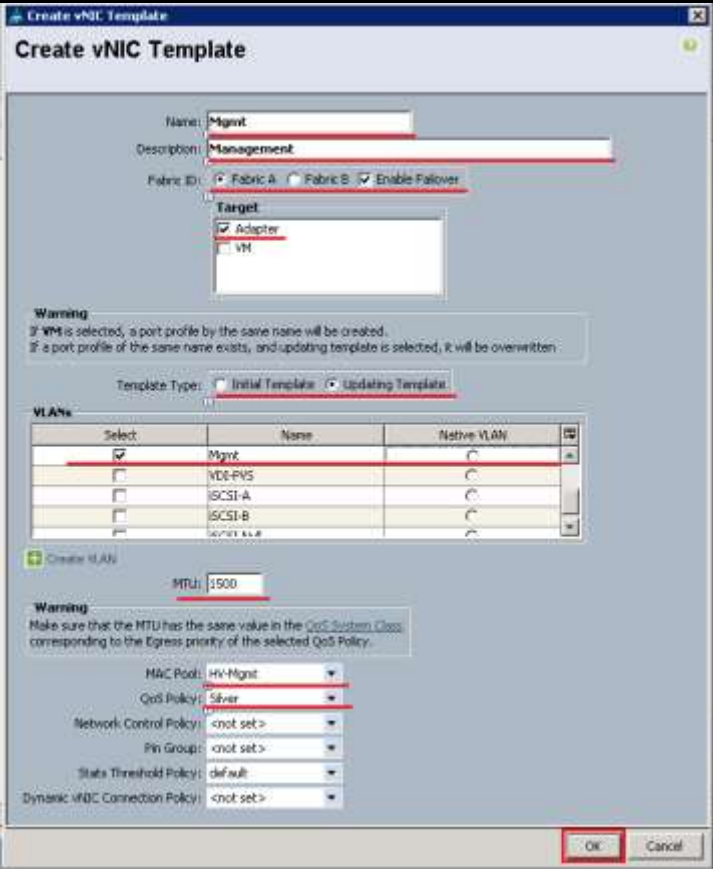
from the blade servers to the Cisco Nexus 1000V virtual distributed switches running on Hyper-V through the higher layer access switches.

The Cisco Nexus 1000V will only be used on the VDI hosts (and not the Infrastructure hosts) per Cisco's best practice guidelines for this release of Cisco Nexus 1000V for Hyper-V. The VDI hosts will then be able to offload the tagging to the Nexus 1000V and will be able to combine the Infrastructure, PVS-VDI, CSV, and Live Migration networks into a single pair of vNICs, referred to in this design as the Uplink vNICs.

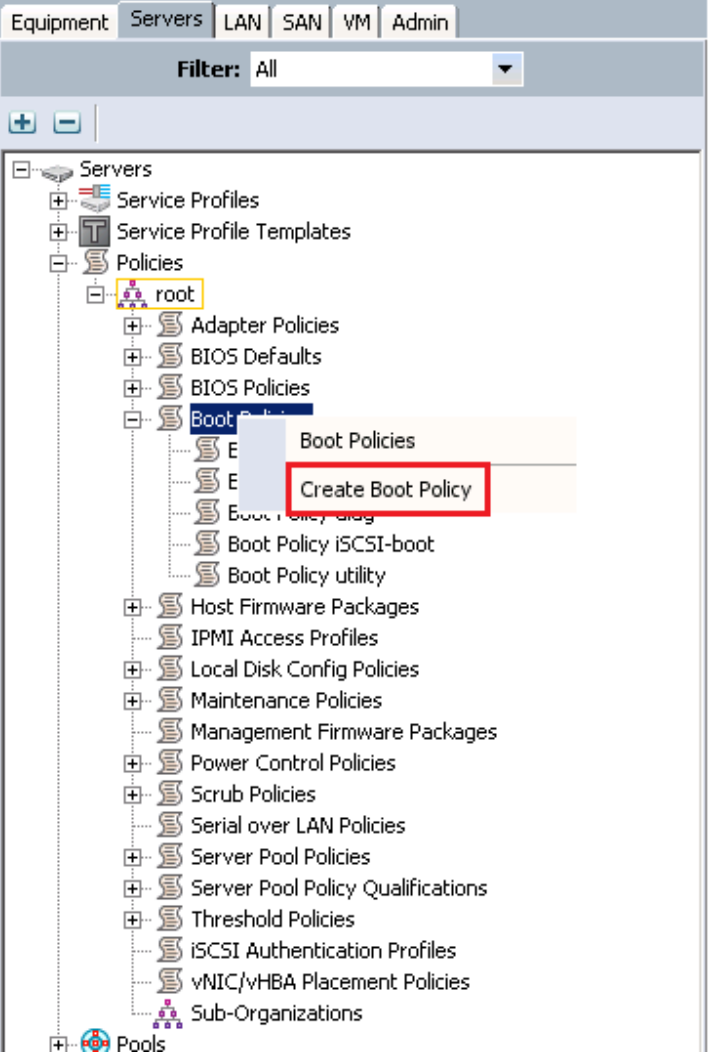
When creating the vNIC templates, you can link the QoS policy to the vNIC template along with the VLAN. Use the following table when creating the vNIC templates:

Table 9: vNIC Template Settings

VLAN	Fabric ID	Enable Failover	Native	MTU	Enabled VLANs	Qos Policy
MGMT	A	Yes	No	1500	Mgmt	Silver
INFRA	B	Yes	No	1500	Infra	Silver
PVS-VDI	A	Yes	No	9000	PVS-VDI	Gold
CSV	A	Yes	No	9000	CSV	Bronze
LMIGR	B	Yes	No	9000	LMigr	Bronze
UPLINK1	A	Yes	No	9000	Mgmt,Infra, PVS-VDI, LMigr	Silver
UPLINK2	B	Yes	No	9000	Mgmt,Infra, PVS-VDI, LMigr	Silver
iSCSI-A	A	No	No	9000	iSCSI-A	Platinum
iSCSI-B	B	No	No	9000	iSCSI-B	Platinum
iSCSI-Null	B	No	Yes	9000	iSCSI-Null	Platinum

Instructions	Visual																		
<p>Create vNICs for all the VLANs in the environment.</p> <p>Complete the vNIC Template, distributing the VLANs across the two fabrics and enabling failover for all but the iSCSI VLANs. Use an updating template so changes later propagate back to the server profiles.</p> <p>Set the MAC Pool to the pool created earlier.</p> <p>Set the QoS Policy based on the VLAN purpose.</p> <p>Click OK to save the changes.</p> <p>Repeat for all the vNICs as identified in Table 9: vNIC Template Settings above.</p> <p>For the Uplink ports, enable checkboxes for the applicable VLANs and do not enable any Native VLAN radio buttons.</p>	 <p>The screenshot shows the "Create vNIC Template" window. The "Name" field is "Mgmt" and the "Description" is "Management". Under "Fabric ID", "Fabric B" is selected and "Enable Failover" is checked. The "Target" section has "Adapter" and "VM" checked. A warning message states: "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." The "Template Type" is set to "Updating Template". Below is a table of VLANs:</p> <table><tr><th>Select</th><th>Name</th><th>Native VLAN</th></tr><tr><td><input checked="" type="checkbox"/></td><td>Mgmt</td><td><input type="radio"/></td></tr><tr><td><input type="checkbox"/></td><td>VDS-PVS</td><td><input type="radio"/></td></tr><tr><td><input type="checkbox"/></td><td>iSCSI-A</td><td><input type="radio"/></td></tr><tr><td><input type="checkbox"/></td><td>iSCSI-B</td><td><input type="radio"/></td></tr><tr><td><input type="checkbox"/></td><td>Server-LAN</td><td><input type="radio"/></td></tr></table> <p>Below the table, the "MTU" is set to 1500. Another warning states: "Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy." The "MAC Pool" is "HV-Mgmt", "QoS Policy" is "Silver", "Network Control Policy" is "<not set>", "Pin Group" is "<not set>", "State Threshold Policy" is "default", and "Dynamic vNIC Connection Policy" is "<not set>". The "OK" button is highlighted with a red box.</p>	Select	Name	Native VLAN	<input checked="" type="checkbox"/>	Mgmt	<input type="radio"/>	<input type="checkbox"/>	VDS-PVS	<input type="radio"/>	<input type="checkbox"/>	iSCSI-A	<input type="radio"/>	<input type="checkbox"/>	iSCSI-B	<input type="radio"/>	<input type="checkbox"/>	Server-LAN	<input type="radio"/>
Select	Name	Native VLAN																	
<input checked="" type="checkbox"/>	Mgmt	<input type="radio"/>																	
<input type="checkbox"/>	VDS-PVS	<input type="radio"/>																	
<input type="checkbox"/>	iSCSI-A	<input type="radio"/>																	
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>																	
<input type="checkbox"/>	Server-LAN	<input type="radio"/>																	

6.6.20. iSCSI Boot Policy

Instructions	Visual
<p>In the Servers tab, expand Policies > root nodes. Select the Boot Policies node. Right-click and choose Create Boot Policy from the context menu.</p>	 <p>The screenshot shows the configuration interface for iSCSI Boot Policy. The 'Servers' tab is selected, and the 'Policies' tree is expanded to 'root'. The 'Boot Policies' node is selected, and a context menu is open with 'Create Boot Policy' highlighted.</p>

In the Create Boot Policy dialog complete the following:

Expand Local Devices

Select **Add Local Disk**

Select **Add CD-ROM**

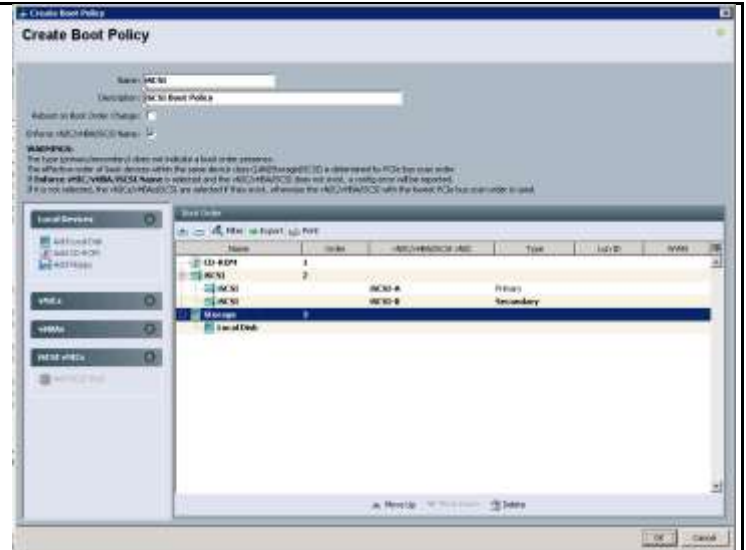
Expand iSCSI vNICs

Select **Add iSCSI Boot (iSCSI-A)**

Select **Add iSCSI Boot (iSCSI-B)**

Adjust boot order so it is CD-ROM, iSCSI-A, iSCSI-B, Local Disk.

Click **OK** to save changes.



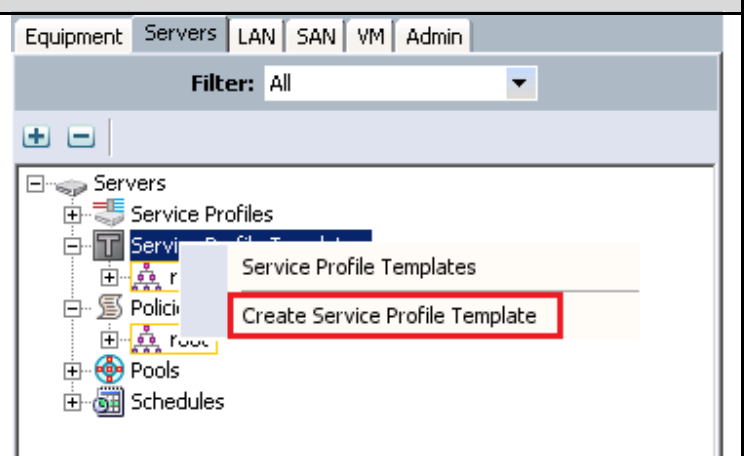
6.6.21. Service Profile Template

Instructions

In the **Servers** tab, select **Service Profile Templates**. Right-click and select **Create Service Profile Template**.

Start by creating a template for the Infrastructure hosts. Afterwards, repeat the process for the Desktops hosts.

Visual



Provide a **Name** and **Description** for the template.

Select **Initial Template** (because later it is unassociated for the iSCSI boot parameters)

Select the UUID pool created earlier. For this validation the pool is called HyperV-UUID.

Click **Next**.

Name: **Infrastructure**

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 Assignment: **HyperV-UUID(112/128)**

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.

Service Profile Template for Infrastructure hosts.

NETWORKING

Select the **Expert** radio button

Click the **Add** button

Add the following vNICs using the vNIC Templates as specified in the next steps.

vNICs

- Mgmt
- Infra
- PVS-VDI
- CSV
- LMigr
- iSCSI-A
- iSCSI-B

iSCSI vNICs

- iSCSI-A
- iSCSI-B

Click **Next**

Dynamic vNIC Connection Policy: Select a Policy to use (No Dynamic vNIC Policy by default) ☒ Create Dynamic vNIC Connection Policy

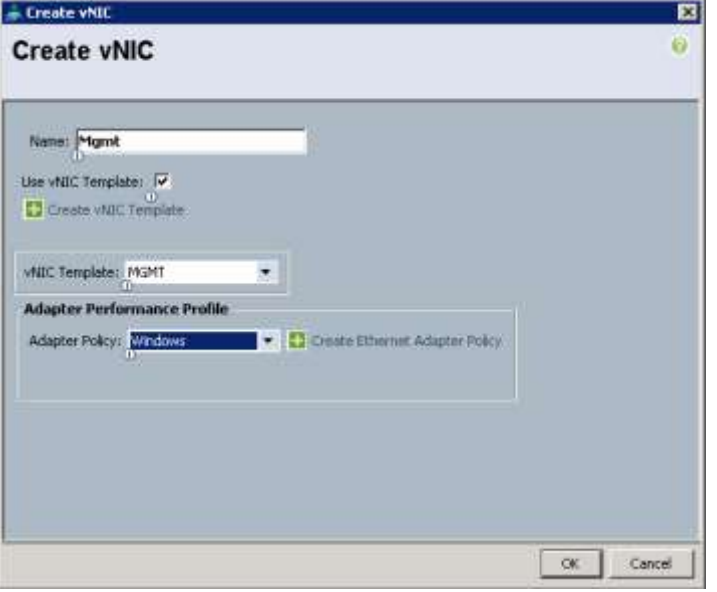
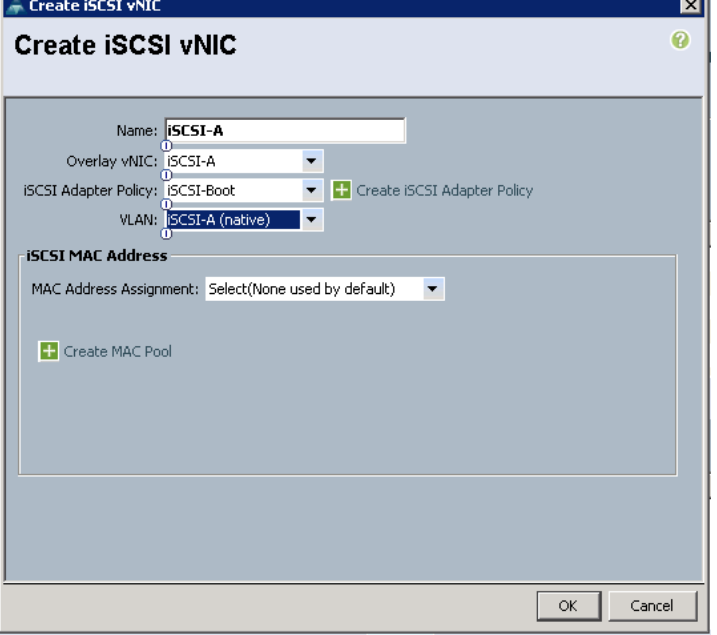

How would you like to configure LAN connectivity? ☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

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

Name	MAC Address	Port ID	Native VLAN
vNIC Mgmt	Derived	derived	
vNIC Infra	Derived	derived	
vNIC PVS-VDI	Derived	derived	
vNIC CSV	Derived	derived	
vNIC LMigr	Derived	derived	

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

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-B	iSCSI-B	iSCSI-Boot	Derived
iSCSI vNIC iSCSI-A	iSCSI-A	iSCSI-Boot	Derived

<p><u>Creating a vNIC</u></p> <p>Provide a Name</p> <p>Enable the vNIC Template</p> <p>Select the associated vNIC Template created earlier</p> <p>Select the Windows Adapter Performance Policy</p> <p>Click OK</p> <p>Repeat for the remaining vNICs.</p>	
<p><u>Creating iSCSI vNIC</u></p> <p>Provide a Name (iSCSI-A or iSCSI-B)</p> <p>Select the associated Overlay vNIC (iSCSI-A or iSCSI-B)</p> <p>Select the iSCSI Adapter Policy created earlier (iSCSI-Boot)</p> <p>Select the associated VLAN (iSCSI-A, iSCSI-B)</p> <p>Do Not select a MAC Address Assignment, leave it as None used by default.</p> <p>Click OK</p>	
<p>STORAGE</p> <p>Select the RAID1 Mirrored for the Local Storage.</p> <p>Select the No vHBAs radio button</p> <p>Click Next</p>	

Skip the Zoning page by clicking **Next**.

vNIC/vHBA PLACEMENT

The suggested order for the vNICs is shown.

Click **Next**.

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement:

System will perform automatic placement of vNICs and vHBAs based on PCI order.

Name	Address	Order	
vNIC Mgmt	Derived	1	
vNIC Infra	Derived	2	
vNIC PVS-VDI	Derived	3	
vNIC CSV	Derived	4	
vNIC LMIgr	Derived	5	
vNIC iSCSI-A	Derived	6	
vNIC iSCSI-B	Derived	7	

SERVER BOOT ORDER

Select iSCSI-Boot **Boot Policy** created earlier

Click **Next**.

Select a boot policy:

Boot Policy:

Name: iSCSI-Boot

Description:

Reboot on Boot Order Change: **No**

Enforce vNIC/vHBA/iSCSI Name: **Yes**

WARNINGS:

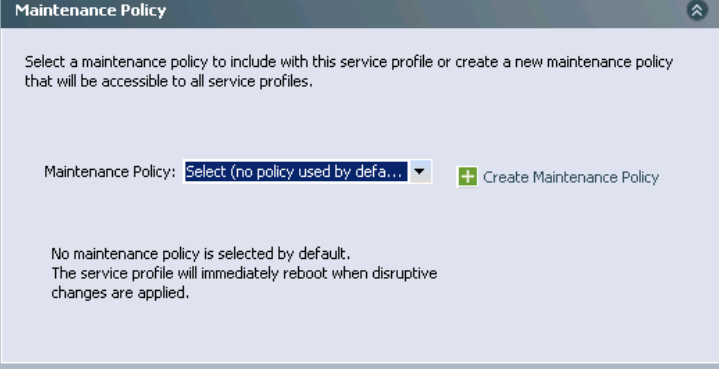


The type (primary/secondary) does not indicate a boot order preference.

The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCI 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.

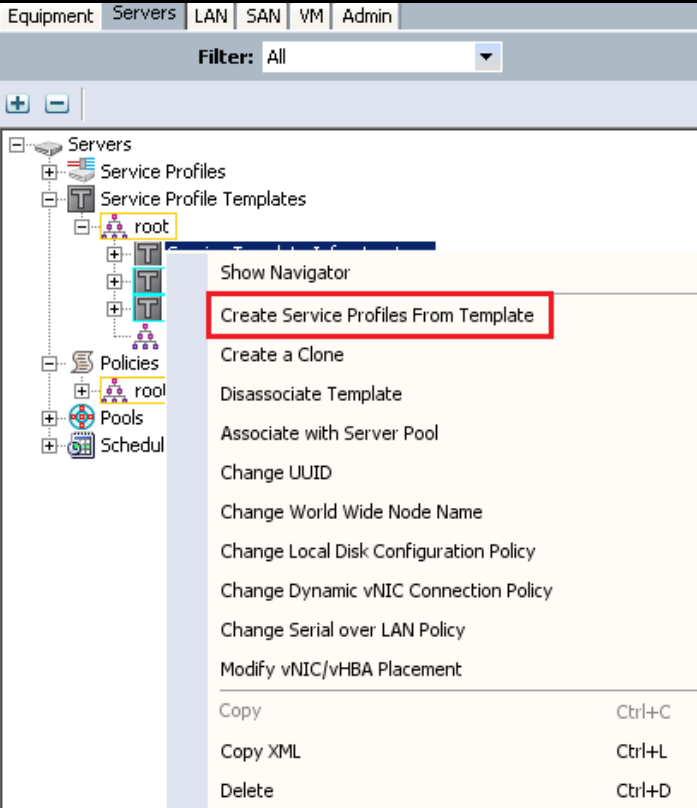
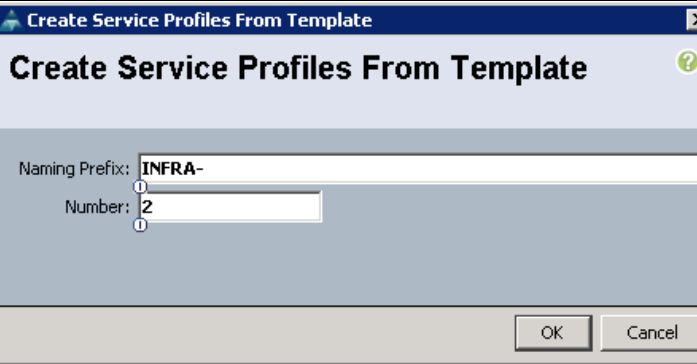
If it is not selected, the vNIC/vHBA/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCI bus scan order is used.

Boot Order						
Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN	
CD-ROM	1					
iSCSI	2					
iSCSI		iSCSI-A	Primary			
iSCSI		iSCSI-B	Secondary			
Storage	3					
Local Disk						

<p>MAINTENANCE POLICY</p> <p>Click Next</p>	
<p>SERVER ASSIGNMENT</p> <p>Set Pool Assignment to the one created earlier. (Infrastructure)</p> <p>Select the Down the radio button for desired power state. (Still need to set the iSCSI IQN information)</p> <p>Set Server Pool Qualification to the one created earlier (Infrastructure)</p> <p>Set Firmware Management to the one created earlier. (B200M3)</p> <p>Click Next</p>	
<p>OPERATIONAL POLICIES</p> <p>Set the Management IP Address Policy to the policy created earlier (KVMIPPool)</p> <p>Set the Scrub Policy to the policy created earlier (HyperVScrub)</p> <p>Click Finish</p>	

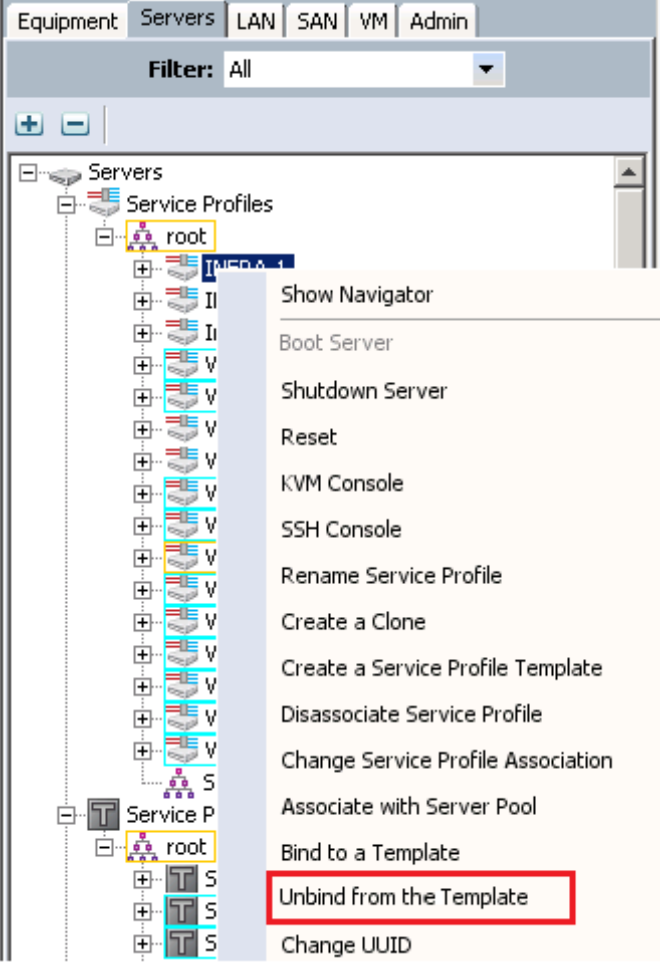
<p>Repeat for the Desktop Hosts using the values specified on the right in place of the earlier ones. Values not specified should be the same as the Infrastructure template.</p>	<p>Name: Desktop</p> <p>Server UUID Assignment: HyperV-UUID</p> <p>NETWORKING</p> <p><u>vNICs</u></p> <ul style="list-style-type: none"> - Mgmt - Uplink1 - Uplink2 - iSCSI-A - iSCSI-B <p><u>iSCSI vNICs</u></p> <ul style="list-style-type: none"> - iSCSI-A - iSCSI-B <p>STORAGE: RAID0 Striping</p> <p>SERVER ASSIGNMENT:</p> <p>Pool Assignment: Desktop</p> <p>Server Pool Qualification: Desktop</p>
---	--

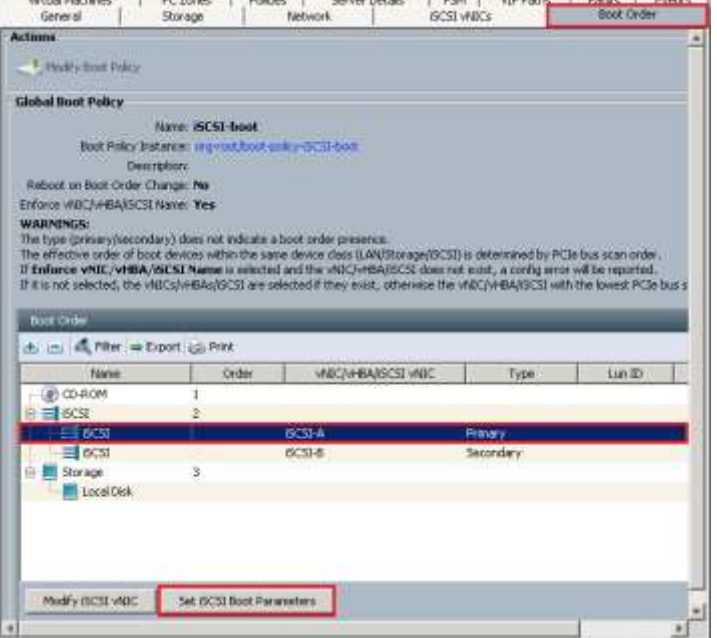
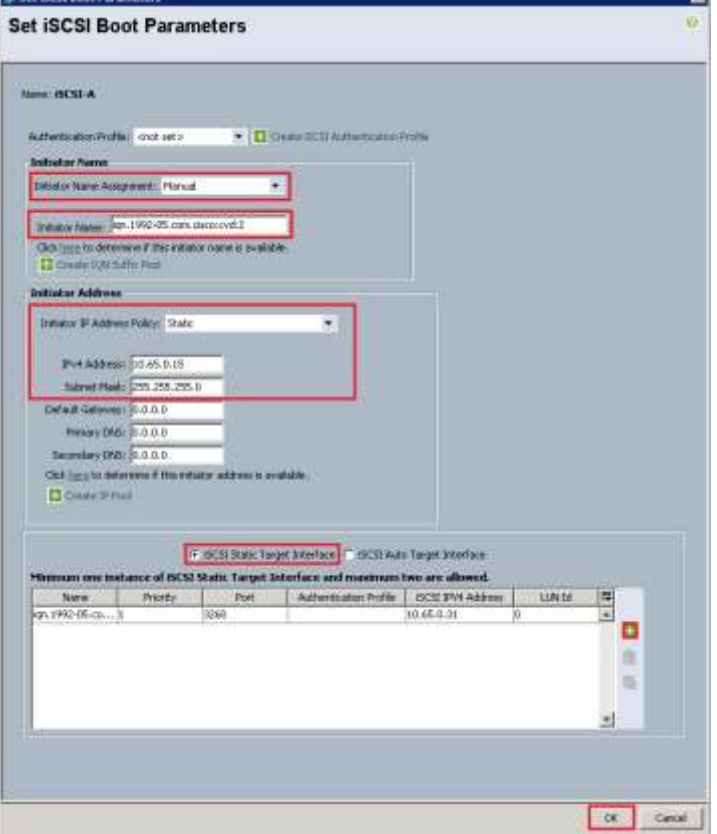
6.6.22. Create Service Profiles

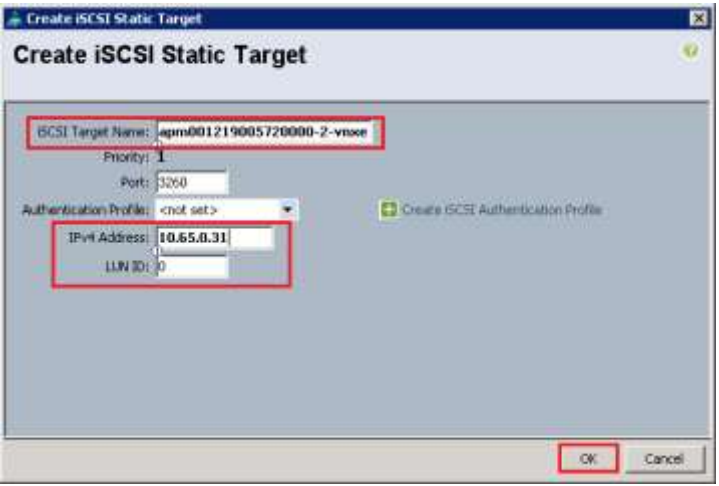
Instructions	Visual
<p>In the Servers tab, expand Service Profile Templates > root. Select the Infrastructure service profile created earlier. Right-click and choose Create Service Profiles From Template from the context menu.</p>	
<p>Provide a Naming Prefix for the Infrastructure servers: INFRA-.</p> <p>Provide the Number of server profiles to create for Infrastructure: 2</p> <p>Click OK to complete the creation of the profiles.</p> <p>Repeat for the Desktop servers in the solution with a Naming prefix of VDI-</p> <p>Number of Desktop server profiles to create: 500-user configuration: 3 1000-user configuration: 5</p>	

6.6.23. Configure iSCSI Boot LUNS for Each Service Profile

The new created service profiles will have to be unbound from the template and updated with iSCSI target information which is unique for each server.

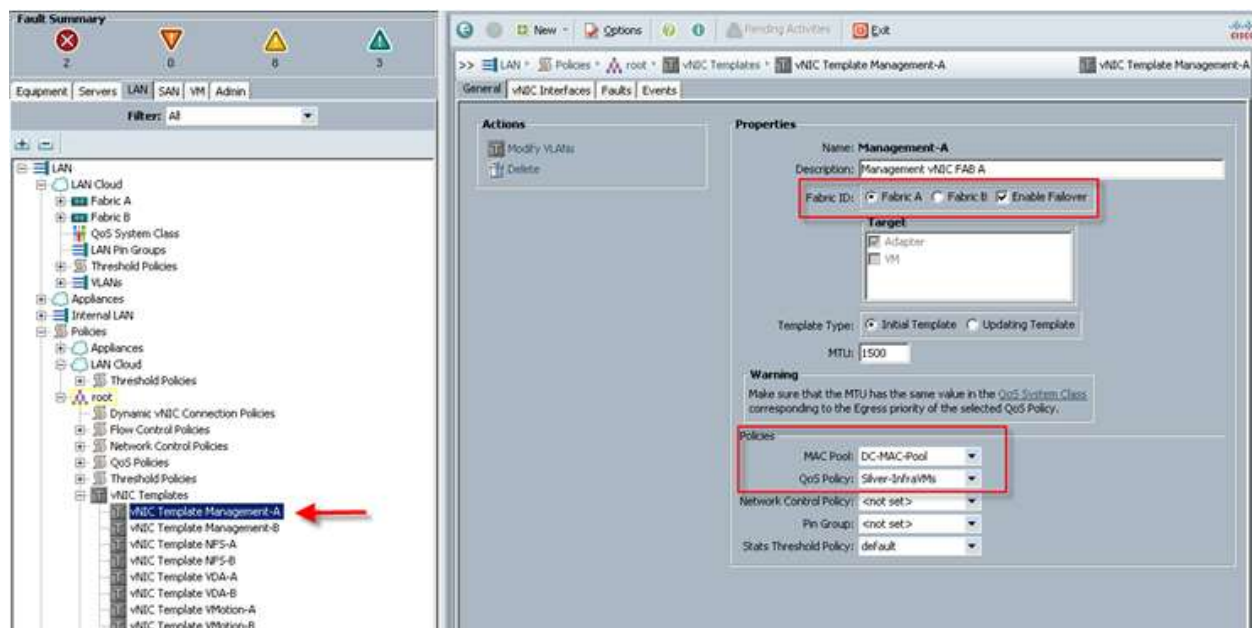
Instructions	Visual
<p>In the Servers tab, expand Service profiles > root. Select the first service profile. Right-click and choose Unbind from the Template from the context menu.</p>	 <p>The screenshot shows the 'Servers' tab in a management console. The 'Service Profiles' section is expanded, showing a 'root' folder. A right-click context menu is open over the first service profile, with the option 'Unbind from the Template' highlighted by a red rectangle. The menu options include: Show Navigator, Boot Server, Shutdown Server, Reset, KVM Console, SSH Console, Rename Service Profile, Create a Clone, Create a Service Profile Template, Disassociate Service Profile, Change Service Profile Association, Associate with Server Pool, Bind to a Template, Unbind from the Template, and Change UUID.</p>

Instructions	Visual
<p>Select a Service Profile</p> <p>Select the Boot Order tab.</p> <p>Select the first iSCSI boot NIC (iSCSI-A)</p> <p>Click the Set iSCSI Boot Parameters button</p>	
<p>The EMC VNXe uses a different IP address and iSCSI Target Name for each of the individual iSCSI boot LUNs. As such, the boot parameters must be configured separately for each service profile.</p> <p>Set Initiator Name Assignment to Manual</p> <p>Set Initiator Name to the unique name from the pool</p> <p>Set the Initiator IP Address Policy to Static</p> <p>Set the IPv4 Address to a unique IP address for this host</p> <p>Set the Subnet Mask to the correct subnet mask for the iSCSI VLAN</p> <p>Select the iSCSI Static Target Interface radio button</p>	

Instructions	Visual
<p>The unique name and IP address will need to be configured in the VNxe side.</p> <p>To add the VNxe Target information, click the + button.</p> <p>Provide the iSCSI Target Name of the VNxe</p> <p>Provide the IPv4 Address and LUN ID for this LUN on the VNxe</p> <p>Click OK to save the iSCSI Static Target information.</p> <p>Click OK to save the iSCSI Boot Parameters</p> <p><i>Repeat the process for the iSCSI-B boot NIC</i></p> <p><i>Repeat these steps for each service profile</i></p>	

Include the corresponding QoS Policy into each vNIC template using the QoS policy drop down, using the QoS Priority to vNIC and VLAN Mapping table above.

Figure 14: Utilize QoS Policy in vNIC Templates



6.7. LAN Configuration

The access layer LAN configuration consists of a pair of Cisco Nexus 5548s (N5Ks,) a family member of our low-latency, line-rate, 10 Gigabit Ethernet and FCoE switches for our VDI deployment.

Four 10 Gigabit Ethernet uplink ports are configured on each of the Cisco UCS 6248 fabric interconnects.

Note: The upstream configuration is beyond the scope of this document; there are some good reference documents in the [Appendix](#) that discuss about best practices of using the Cisco Nexus 5000 and 7000 Series Switches. For informational purposes, the upstream switch for this Cisco Validated Design was a Cisco Nexus 5500 series which included the Layer 3 module.

6.8. SAN Configuration

6.8.1. Boot from SAN benefits

Boot 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 iSCSI boot LUN 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 an iSCSI boot vNIC. The vNIC boot code contains the instructions that enable the server to find the boot disk. All Cisco VIC 1240 MLOM cards 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 BOIS settings. When the hardware detects the boot device, it follows the regular boot process.

In this study we deployed boot from SAN over iSCSI protocol.

6.9. EMC VNXe Storage Configuration

The figure below shows the physical storage layout of the disks in the reference architecture. This configuration accommodates PVS vDisks, user homes and profiles for up to 1000 virtual desktops, hypervisor boot LUNs, SQL databases, SCVMM library, cluster quorum disk, and cluster shared volume (CSV).

Figure 15: EMC VNXe Storage Configuration



The above storage layout is used for the following configurations. Note that VNXe provisioning wizards perform disk allocation and do not allow user selection.

- Seven SAS disks are allocated in a 6+1 RAID-5 pool that contains boot LUNs for the hypervisor OS.

- Seven SAS disks are allocated in a 6+1 RAID-5 pool that contains SQL databases, SCVMM library, cluster quorum disk, and CSV.
- Twenty eight SAS disks are allocated in a 6+1 RAID-5 pool that contains PVS vDisks, user home directories and profiles.
- Two SAS disks are used as hot spares and are contained in the VNXe hot spare pool.

If more capacity is required, larger drives may be substituted. To satisfy the load recommendations, the drives will all need to be 15k rpm and the same size. If differing sizes are utilized, storage layout algorithms may give sub-optimal results.

Table 10 shows the LUNs that need to be created, the size, the assigned host, the assigned storage pool, and the names used for this Cisco Validated Design.

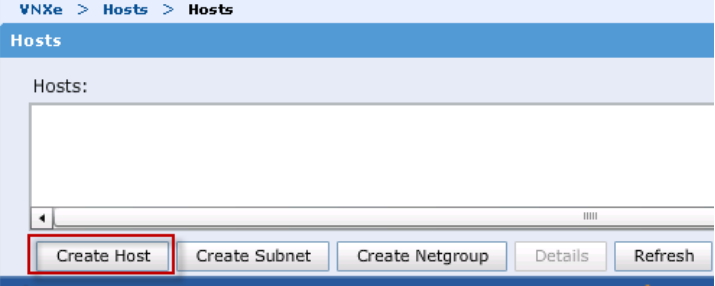

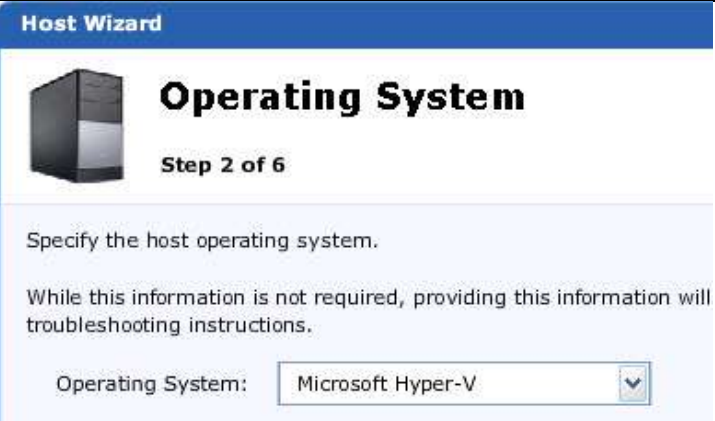
Table 10: Storage LUN Mapping



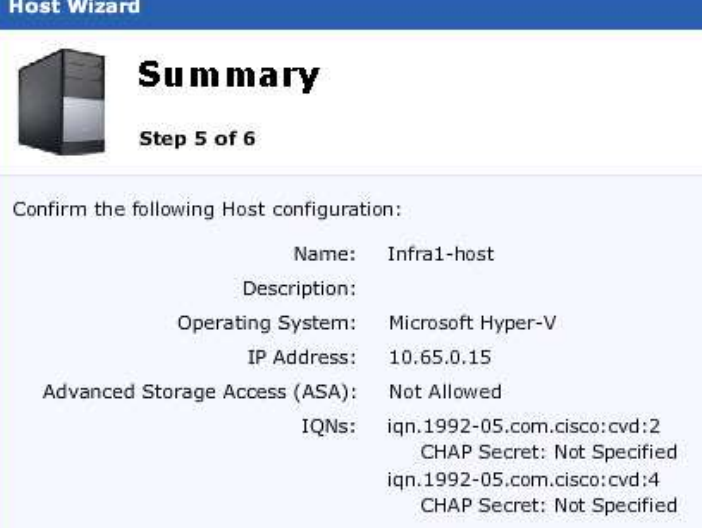
HostName	LUN	Size (GB)	Assigned Pool	Storage Type
Infra-1	Infra-1	50	Boot LUN	Generic iSCSI Storage
Infra-1	SQLDB1LUN	50	InfrastructurePool	Generic iSCSI Storage
Infra-1	SQLLog1LUN	50	InfrastructurePool	Generic iSCSI Storage
Infra-2	Infra-2	50	Boot LUN	Generic iSCSI Storage
Infra-2	SQLDB2LUN	50	InfrastructurePool	Generic iSCSI Storage
Infra-2	SQLLog2LUN	50	InfrastructurePool	Generic iSCSI Storage
Infra-1/Infra-2	InfraCSVLUN	750	InfrastructurePool	Generic iSCSI Storage
Infra-1/Infra-2	InfraQuorumLUN	2	InfrastructurePool	Generic iSCSI Storage
Infra-1/Infra-2	SCVMMLibraryLUN	200	InfrastructurePool	Generic iSCSI Storage
UserProfile2	CIFSServer1	200	CIFSPool	Shared Folders
UserProfile	CIFSServer2	200	CIFSPool	Shared Folders
vDisk	CIFSServer2	500	CIFSPool	Shared Folders
VDI1-2	VDI1-2	50	Boot LUN	Generic iSCSI Storage
VDI1-3	VDI1-3	50	Boot LUN	Generic iSCSI Storage
VDI1-4	VDI1-4	50	Boot LUN	Generic iSCSI Storage
VDI2-1	VDI2-1	50	Boot LUN	Generic iSCSI Storage
VDI2-2	VDI2-2	50	Boot LUN	Generic iSCSI Storage


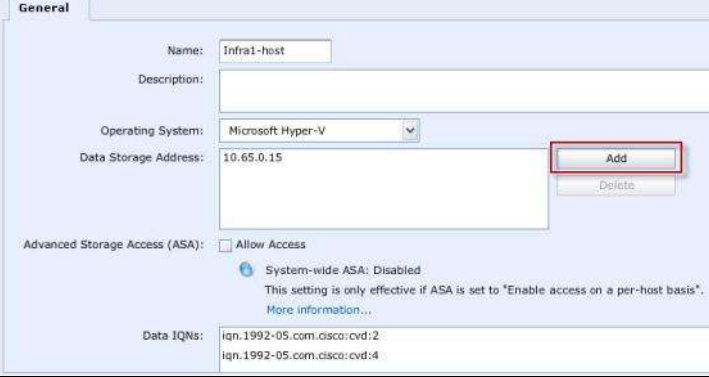
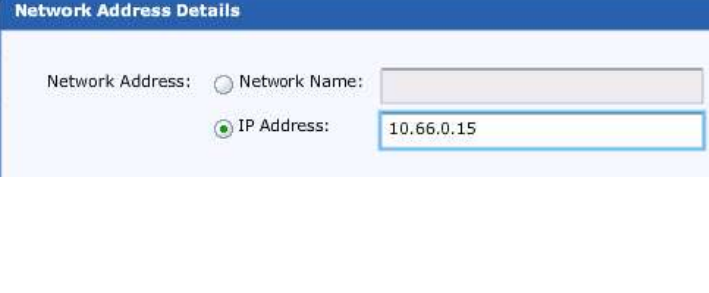
The rest of this section covers creating the hosts, storage pools, and LUNs required for the environment. Single examples are provided, but the process should be repeated for each of the hosts, pools, and LUNs required for the environment.

6.9.1. iSCSI Host Configuration

Complete the following steps in Unisphere to configure iSCSI hosts on VNXe3300 for each of the hosts in the environment.

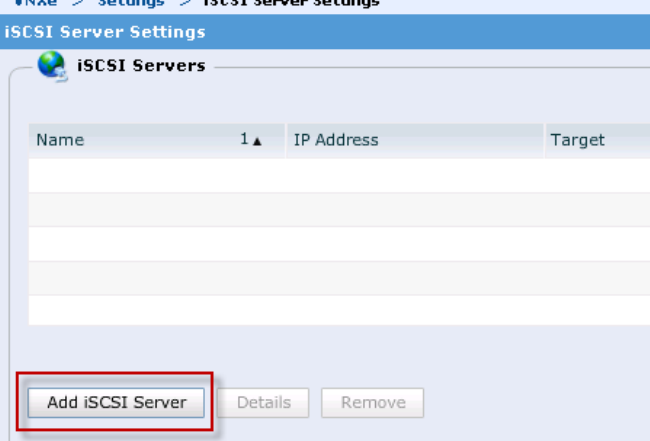
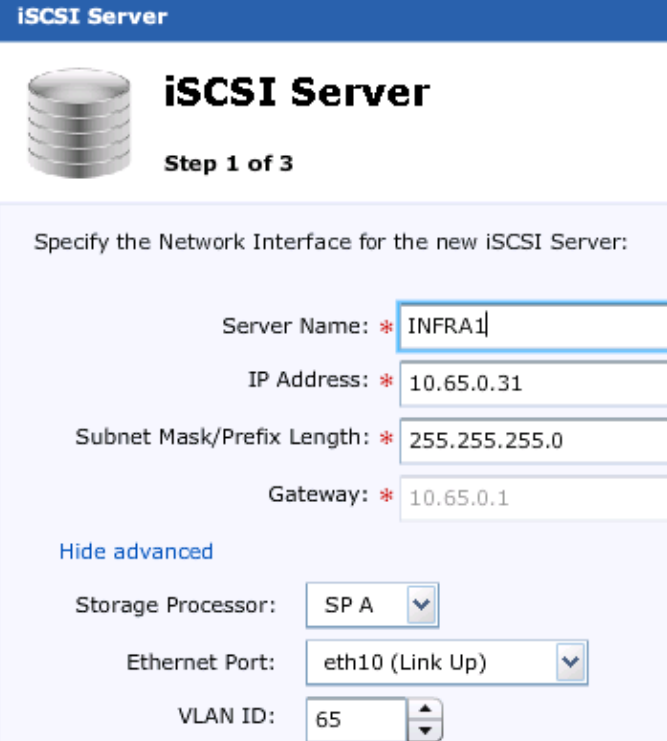
Instructions	Visual
<p>From the Hosts->Hosts area in Unisphere, select Create Host.</p>	
<p>Specify a name and description (optional) for the iSCSI host.</p>	
<p>Specify Microsoft Hyper-V as the host operating system.</p>	

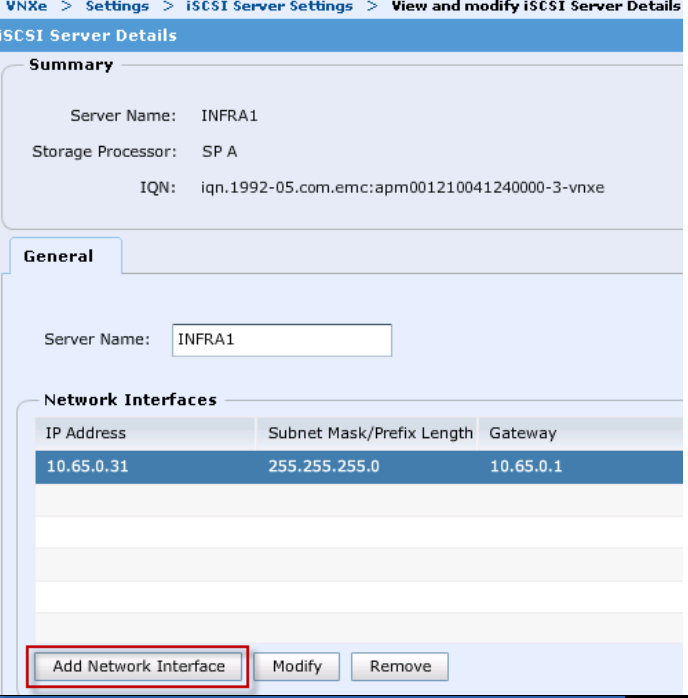
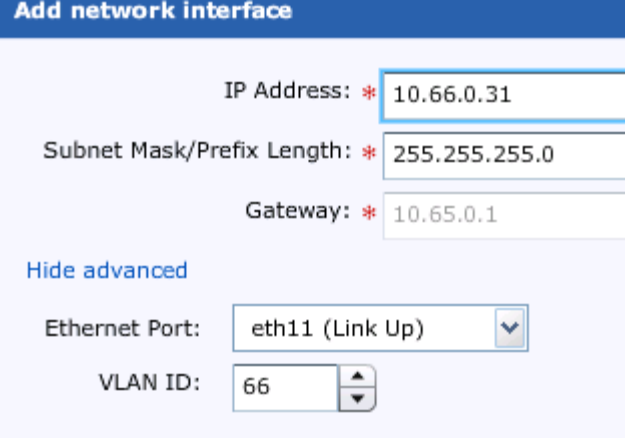
Instructions	Visual
<p>Specify an IP Address that will be used as an iSCSI initiator for the host.</p> <p>In this Cisco Validated Design, the IP address is the statically assigned IP address on the iSCSI-A VLAN.</p>	
<p>Specify all IQN addresses that will be used on the iSCSI host.</p> <p>The IQN addresses can be found in the corresponding service profile in Cisco UCS Manager. Generally, these IQNs are unique and created by the storage administrator. The typical format would be iqn.1992-05.<topdomain>.<domain>:<systemid>:<number></p> <p>Optionally, fill out the CHAP Secret fields if CHAP security needs to be enforced.</p>	
<p>Review the host configuration on the Summary page and select Finish to create the host entry.</p>	

Instructions	Visual
<p>Select the newly created host and Details to modify its properties.</p>	 <p>The screenshot shows the VNXe interface with the 'Hosts' tab selected. A list of hosts is shown, with 'Infra1-host' highlighted. Below the list, there are buttons for 'Create Host', 'Create Subnet', 'Create Netgroup', 'Details' (highlighted with a red box), 'Refresh', and 'Delete'.</p>
<p>To the right of the Data Storage Address field, select Add to add the secondary IP address of the host.</p>	 <p>The screenshot shows the 'General' tab for the 'Infra1-host' configuration. The 'Data Storage Address' field is set to '10.65.0.15'. To the right of this field, there is an 'Add' button (highlighted with a red box) and a 'Delete' button. Below this, there is a section for 'Advanced Storage Access (ASA)' with a checkbox for 'Allow Access' and a note about 'System-wide ASA: Disabled'.</p>
<p>Enter the secondary IP address in the IP Address field. Select OK and click the Apply Changes button to save the changes.</p> <p>The secondary IP address would be the statically assigned alternate IP address for the host on the iSCSI-B VLAN.</p>	 <p>The screenshot shows the 'Network Address Details' tab. The 'Network Address' section has two radio buttons: 'Network Name' and 'IP Address'. The 'IP Address' radio button is selected, and the 'IP Address' field is highlighted with a blue box and contains the value '10.66.0.15'.</p>

6.9.2. iSCSI Server Configuration

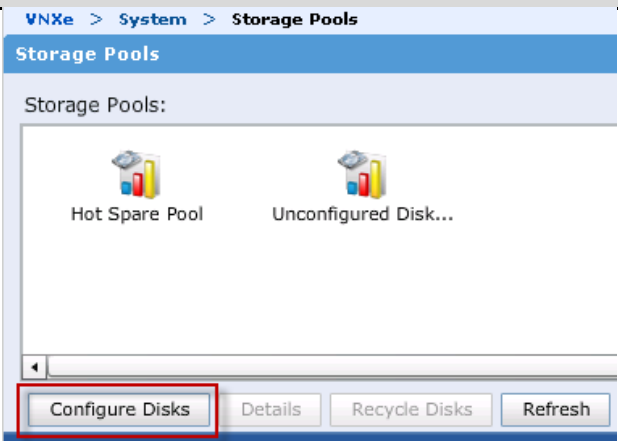
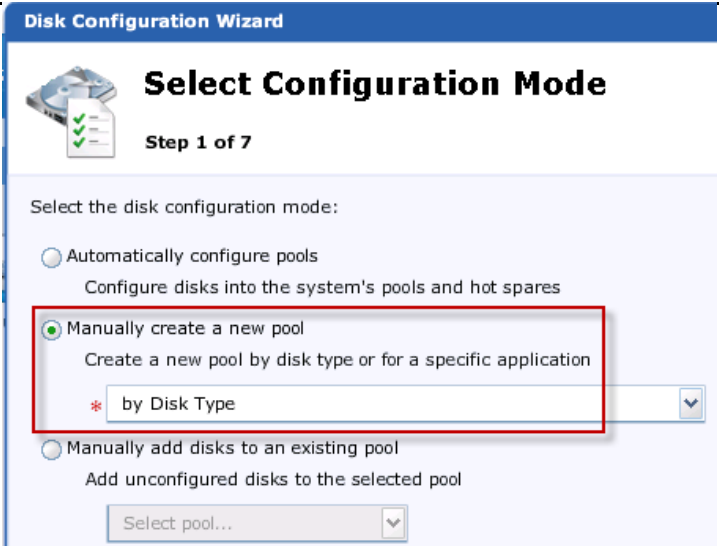

Complete the following steps in Unisphere to configure iSCSI servers on VNXe.





Instructions	Visual
<p>Prior to provisioning iSCSI storage, an iSCSI server must be created.</p> <p>From the Settings->iSCSI Server Settings area in Unisphere, select Add iSCSI Server.</p>	
<p>Fill out the iSCSI server network information as shown in the screenshot. Select Finish to create the iSCSI server.</p> <p>The Advanced settings will change for each of the iSCSI servers and should alternate Storage Processors (SPA/SPB) and Ethernet Ports (eth10/eth11) to load-balance across the resources.</p> <p>The VLAN ID will need to match the IP address specified for the iSCSI server and the Ethernet Port. For instance, in this Cisco Validated Design, eth10 maps to VLAN65 and eth11 maps to VLAN66.</p>	

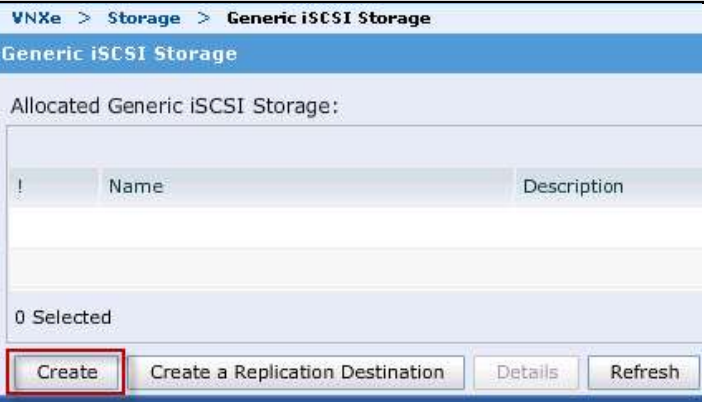


Instructions	Visual
<p>Select the newly created iSCSI server and select Details.</p>	
<p>From the iSCSI Server Details screen, select Add Network Interface to add a second interface to provide network redundancy for the iSCSI server.</p>	
<p>Fill out the network information of the second interface as shown in the screenshot. Select Add to create the second interface for the iSCSI server to provide HA across the resources.</p> <p>The VLAN ID will map to the Ethernet Port as mentioned earlier. For this Cisco Validated Design, eth10 maps to VLAN65 and eth11 to VLAN66.</p>	



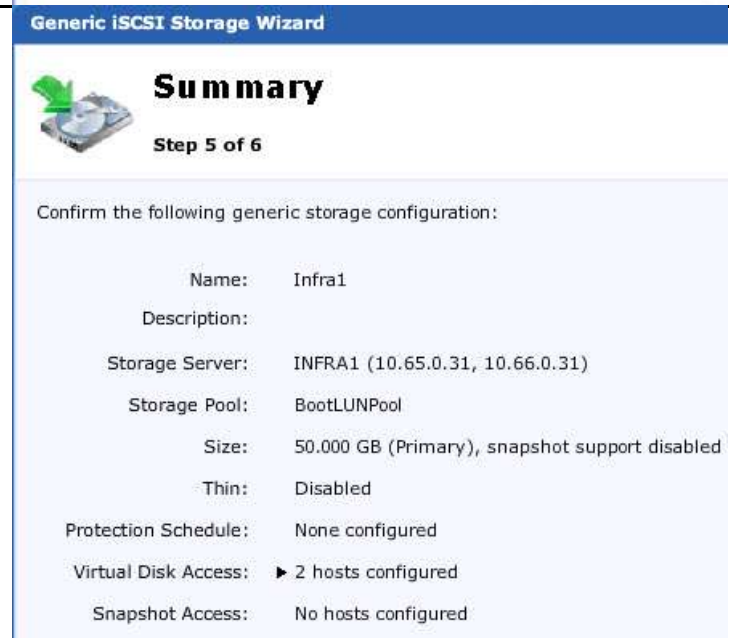
6.9.3. iSCSI Storage Creation

Complete the following steps in Unisphere to configure iSCSI LUNs on VNXe.



Instructions	Visual
<p>Prior to provisioning iSCSI storage, create a storage pool with the appropriate number of disks that will contain the iSCSI LUNs.</p> <p>From the System->Storage Pools area in Unisphere, select Configure Disks.</p>	
<p>Select Manually create a new pool by Disk Type.</p>	
<p>Specify a pool name of your choice.</p>	


Instructions	Visual															
Select SAS and Balanced Perf/Capacity when prompted for disk type and storage profile.	<div><div>Disk Configuration Wizard</div><div><div></div><div>Select Storage Type</div><div>Step 3 of 6</div><div></div></div><div>Please select the type of disks you want to use for this new pool.</div><div><table><tr><th>Disk Type</th><th>Max Capacity</th><th>Storage Profile</th></tr><tr><td>NL SAS</td><td>14.331 TB</td><td>High Capacity</td></tr><tr><td>SAS</td><td>3.145 TB</td><td>High Performance</td></tr><tr><td>SAS</td><td>6.291 TB</td><td>Balanced Perf/Capacity</td></tr><tr><td>EFD</td><td>0 GB (None Available)</td><td>Best Performance</td></tr></table></div><div>Show advanced</div><div>Uses SAS disks to provide a balanced level of storage performance and capacity. This pool type does not offer performance as high as High Performance pools, but it can be adequate for databases with low-to-average performance requirements.</div><div>General purpose SAS storage pool using RAID 5(6+1).</div></div>	Disk Type	Max Capacity	Storage Profile	NL SAS	14.331 TB	High Capacity	SAS	3.145 TB	High Performance	SAS	6.291 TB	Balanced Perf/Capacity	EFD	0 GB (None Available)	Best Performance
Disk Type	Max Capacity	Storage Profile														
NL SAS	14.331 TB	High Capacity														
SAS	3.145 TB	High Performance														
SAS	6.291 TB	Balanced Perf/Capacity														
EFD	0 GB (None Available)	Best Performance														
Specify the desired number of disks to be included in this storage pool.	<div><div>Disk Configuration Wizard</div><div><div></div><div>Select Amount of Storage</div><div>Step 4 of 6</div></div><div>Select the amount of storage to configure.</div><div><div>300GB SAS (15000 RPM) Disks:</div><div><div>Use 7 of 28 Disks</div><div>▼</div></div></div><div>Total Disks to Configure: 7</div></div>															
Review the pool configuration in the Summary page and select Finish to create the storage pool.	<div><div>Disk Configuration Wizard</div><div><div></div><div>Summary</div><div>Step 5 of 6</div></div><div>The disks will be configured into a new storage pool as indicated below.</div><div><div>Storage Pool Name:</div><div>BootLUNPool</div></div><div><div>Storage Pool Description:</div><div></div></div><div><div>300GB SAS (15000 RPM) Disks:</div><div>7</div></div></div>															

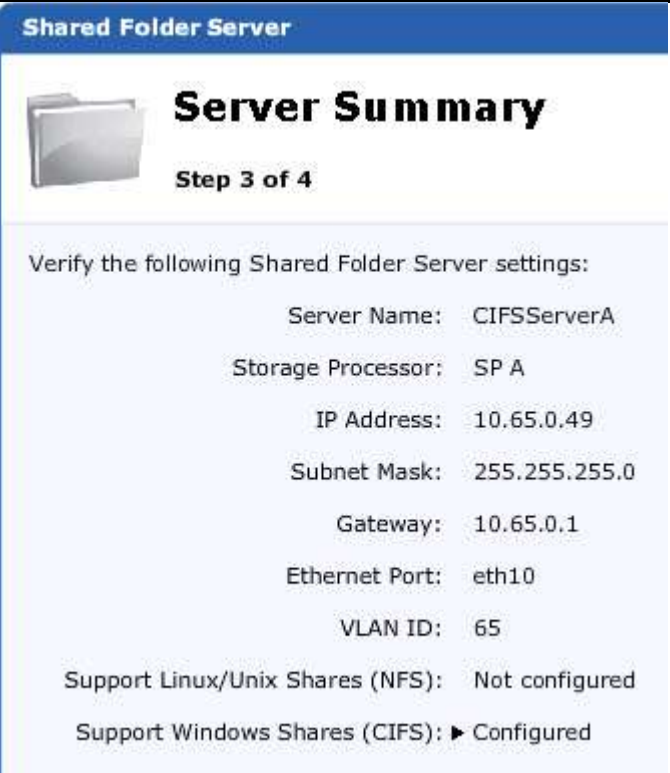
Instructions	Visual
<p>After the storage pool is created, navigate to Storage->Generic iSCSI Storage in Unisphere. Select Create to provision iSCSI storage.</p>	
<p>Specify a name and description (optional) for the storage resource.</p>	
<p>Specify a storage server (iSCSI server defined in the previous steps), a storage pool with enough space available, and the size of the iSCSI LUN to be created.</p> <p>The assigned storage processor (SPA/SPB) will be dependent on the iSCSI server selected. Load-balancing across the storage processors is done manually when creating the iSCSI servers as discussed above.</p>	

Instructions	Visual
<p>Choose the protection policy for replication and snapshots that fits your requirements.</p>	
<p>Select Virtual Disk in the Access column for the host to which iSCSI LUN access is to be granted.</p>	
<p>Review the iSCSI storage configuration in the Summary page and select Finish to provision the iSCSI LUN and grant access to the designated host(s).</p> <p>The Summary page shows that the host has both IP addresses (and subsequently the matching IQNs though not shown) assigned to the host on the iSCSI-A and iSCSI-B networks.</p>	

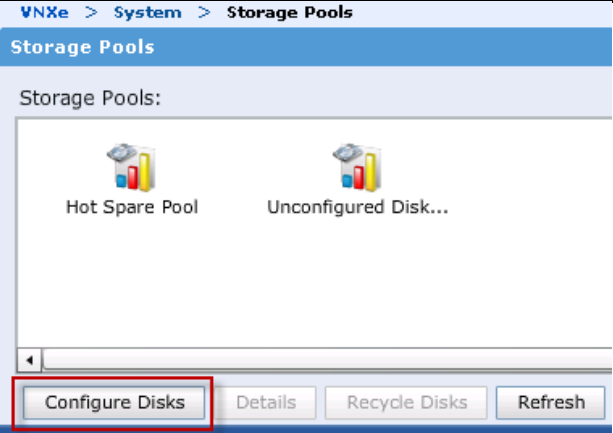
6.9.4. CIFS Server Creation

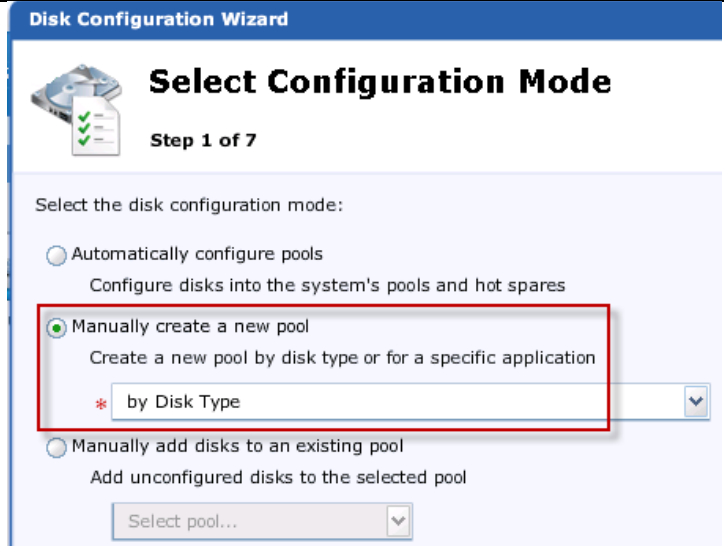

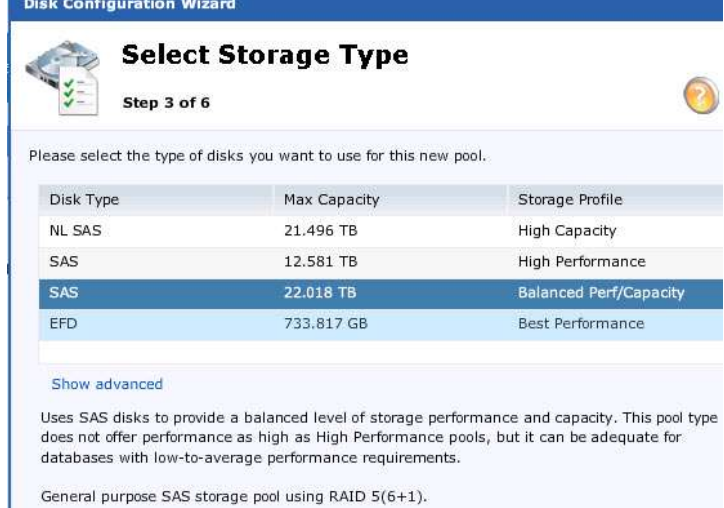
Instructions	Visual
<p>Prior to provisioning storage for CIFS share, a shared folder server must be created.</p> <p>From the Settings->Shared Folder Server Settings area in Unisphere, select Add Shared Folder Server.</p>	
<p>Fill out the CIFS server network information as shown in the screenshot.</p> <p>Repeat for the CIFS share on the other Storage Processor (SPB) if building the 1000-user configuration.</p>	

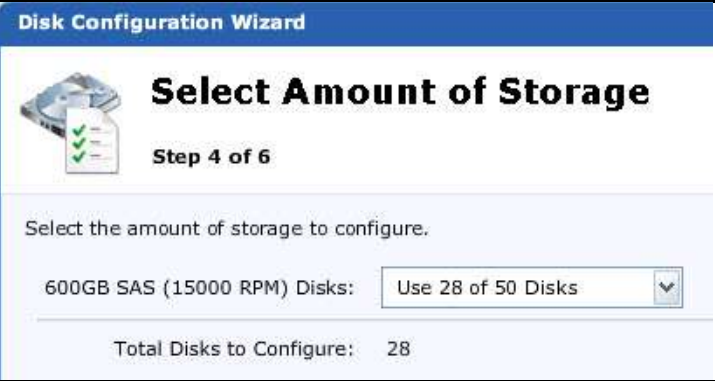

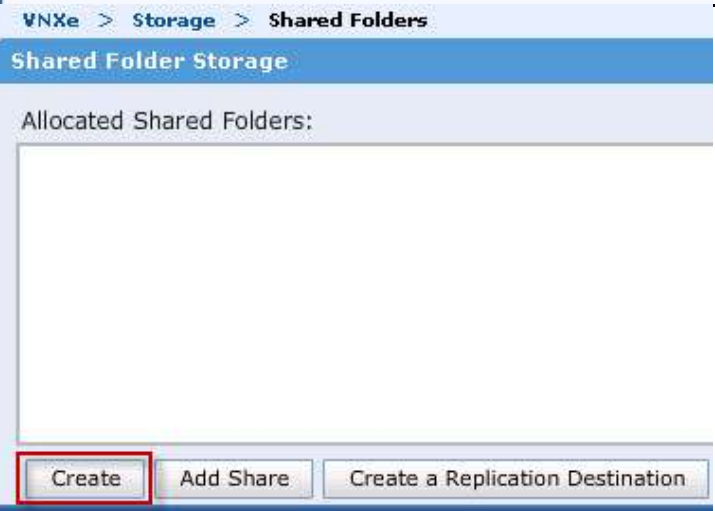
Instructions	Visual
<p>Deselect Linux/Unix shares (NFS) if NFS is not required.</p> <p>Select Windows shares (CIFS) and Join to the Active Directory. Specify a Windows domain, DNS server, and an administrator's credential used to join the CIFS server to the designated domain.</p>	<div data-bbox="797 226 1505 1150"> <div>Shared Folder Server</div> <div>  <h2>Shared Folder Types</h2> <p>Step 2 of 4</p> <p>Choose the type of shares the Shared Folder Server supports:</p> <div> <input type="checkbox"/> Linux/Unix shares (NFS) <input checked="" type="checkbox"/> Windows shares (CIFS) <input type="radio"/> Standalone <input checked="" type="radio"/> Join to the Active Directory </div> <div> <div>Windows Domain: *</div> <div>HV.POD.LOCAL</div> </div> <div> <div>DNS Servers: *</div> <div>10.82.0.10</div> <div>10.82.0.11</div> </div> <div> <div>IP Address:</div> <div></div> </div> <div> <div>User Name:</div> <div>administrator</div> </div> <div> <div>Password:</div> <div>*****</div> </div> <div> <div>Organizational Unit:</div> <div><Click to edit default></div> </div> </div> </div>




Instructions	Visual
<p>Review the shared folder server configuration in the Summary page and select Finish to create the shared folder server.</p>	 <p>Shared Folder Server</p> <p>Server Summary</p> <p>Step 3 of 4</p> <p>Verify the following Shared Folder Server settings:</p> <ul style="list-style-type: none"> Server Name: CIFSServerA Storage Processor: SP A IP Address: 10.65.0.49 Subnet Mask: 255.255.255.0 Gateway: 10.65.0.1 Ethernet Port: eth10 VLAN ID: 65 Support Linux/Unix Shares (NFS): Not configured Support Windows Shares (CIFS): ► Configured



6.9.5. CIFS Share Storage Creation – Profile Share


Instructions	Visual
<p>Prior to provisioning CIFS storage, create a storage pool with the appropriate number of disks that will contain the CIFS shares.</p> <p>From the System->Storage Pools area in Unisphere, select Configure Disks.</p>	 <p>VNXe > System > Storage Pools</p> <p>Storage Pools</p> <p>Storage Pools:</p> <ul style="list-style-type: none"> Hot Spare Pool Unconfigured Disk... <p>Configure Disks Details Recycle Disks Refresh</p>

Instructions	Visual
Select Manually create a new pool by Disk Type .	
Specify a pool name of your choice.	
Select SAS and Balanced Perf/Capacity when prompted for disk type and storage profile.	

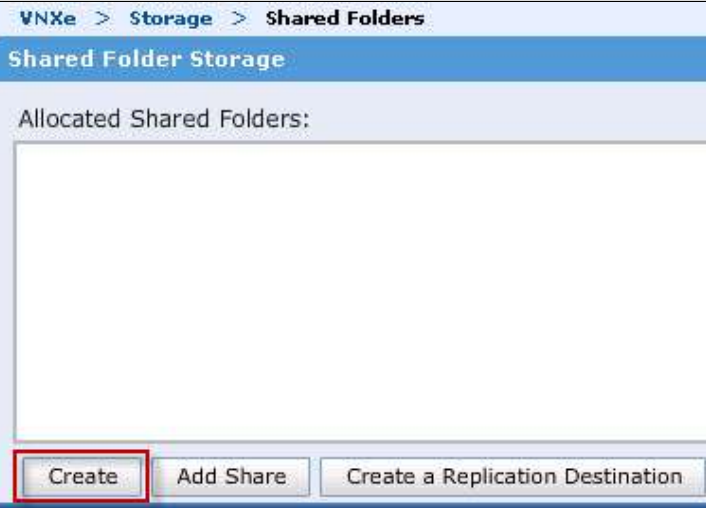
Instructions	Visual
<p>Specify the desired number of disks to be included in this storage pool.</p>	 <p>Disk Configuration Wizard</p> <p>Select Amount of Storage</p> <p>Step 4 of 6</p> <p>Select the amount of storage to configure.</p> <p>600GB SAS (15000 RPM) Disks: Use 28 of 50 Disks</p> <p>Total Disks to Configure: 28</p>
<p>Review the pool configuration in the Summary page and select Finish to create the storage pool.</p>	 <p>Disk Configuration Wizard</p> <p>Summary</p> <p>Step 5 of 6</p> <p>The disks will be configured into a new storage pool as indicated below.</p> <p>Storage Pool Name: CIFSSharePool</p> <p>Storage Pool Description:</p> <p>600GB SAS (15000 RPM) Disks: 28</p>
<p>After the storage pool is created, navigate to Storage->Shared Folders in Unisphere. Select Create to provision CIFS storage.</p>	 <p>VNXe > Storage > Shared Folders</p> <p>Shared Folder Storage</p> <p>Allocated Shared Folders:</p> <p>Create Add Share Create a Replication Destination</p>




Instructions	Visual																									
<p>Specify a name and description (optional) for the shared folder.</p>	<div><div>Shared Folder Wizard</div><div><div></div><div><h2>Specify Shared Folder Name</h2><p>Step 1 of 7</p><p>Specify a name and optional description for the shared folder:</p><div><div>Name: *</div><div><input type="text" value="UserProfile"/></div></div><div><div>Description:</div><div><input type="text"/></div></div></div></div></div>																									
<p>Specify a storage server (CIFS server defined in the previous steps), a storage pool with enough space available, and the size of the CIFS share to be created.</p>	<div><div>Shared Folder Wizard</div><div><div></div><div><h2>Configure Shared Folder Storage</h2><p>Step 2 of 7</p><p>Configure the storage for this shared folder:</p><p>Select a storage pool with available space on the selected shared folder server.</p><div><div>Storage Server:</div><div><div>CIFSServerA (SP A)</div><div>▼</div></div><div>More information...</div></div><table><thead><tr><th>Type</th><th>1 ▲</th><th>Pool</th><th>Available</th><th>Percent Used</th></tr></thead><tbody><tr><td>SAS</td><td></td><td>CIFSSharePool</td><td>5.974 TB</td><td><div><div></div></div>0%</td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr></tbody></table><div><div>Percent Available:</div><div><div></div></div><div>Percent Used:</div><div><div></div></div></div><div><div>Size: *</div><div><input type="text" value="200"/></div><div>GB</div><div>▼</div></div><div><input type="checkbox"/> Thin</div></div></div></div>	Type	1 ▲	Pool	Available	Percent Used	SAS		CIFSSharePool	5.974 TB	<div><div></div></div> 0%															
Type	1 ▲	Pool	Available	Percent Used																						
SAS		CIFSSharePool	5.974 TB	<div><div></div></div> 0%																						
<p>Choose the protection policy for replication and snapshots that fits your requirements.</p>	<div><div>Shared Folder Wizard</div><div><div></div><div><h2>Configure Protection</h2><p>Step 3 of 7</p><p>Configure protection storage for replication and snapshots:</p><div><div><input checked="" type="radio"/> Do not configure protection storage for this storage resource.</div><div>Replication and snapshots can be supported by allocating protection space at a later time.</div><div><input type="radio"/> Configure protection storage, do not configure a snapshot protection schedule.</div><div>An automated snapshot protection schedule may be configured at a later time.</div><div><input type="radio"/> Configure protection storage, protect data using snapshot schedule:</div><div><div>Default Protection</div></div></div><div><div>This schedule will create snapshots</div><div>Every day at 04:00, keep for 2 days</div></div><div>Note: Times are displayed in Local Time (UTC-0400) in 24-hour format</div></div></div></div>																									



Instructions	Visual
<p>Select Windows shares (CIFS) as the share type.</p>	 <p>Shared Folder Wizard</p> <p>Configure Shared Folder Attributes</p> <p>Step 4 of 7</p> <p>Configure the type of shares which will be exported from this shared folder:</p> <p>Share Type: <input checked="" type="radio"/> Windows shares (CIFS)</p> <p>CIFS shares use the CIFS/SMB protocol to share content in Windows environments.</p> <p><input type="radio"/> Linux/Unix shares (NFS)</p> <p>The selected storage server does not have NFS support enabled. You can enable this feature, which requires a license, from the Shared Folder Server Settings page.</p> <p>Show advanced</p>
<p>Select Create a Windows share and specify a share name of your choice.</p>	 <p>Shared Folder Wizard</p> <p>Configure Share</p> <p>Step 5 of 7</p> <p>Configure the share to be created for this shared folder:</p> <p>Local Path: /</p> <p><input checked="" type="checkbox"/> Create a Windows share</p> <p>Name: <input type="text" value="UserProfile"/></p> <p>Export Path: <input type="text" value="\\10.65.0.49\\UserProfile"/></p> <p>Description: <input type="text"/></p> <p>Show advanced</p>

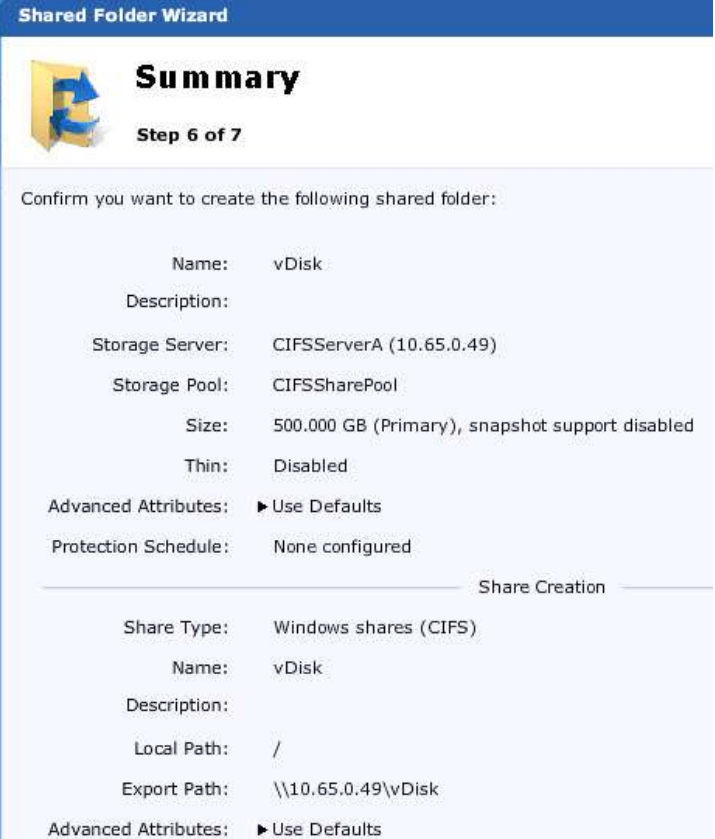
Instructions	Visual
<p>Review the shared folder configuration in the Summary page and select Finish to provision the CIFS share.</p>	 <p>Shared Folder Wizard</p> <p>Summary</p> <p>Step 6 of 7</p> <p>Confirm you want to create the following shared folder:</p> <p>Name: UserProfile</p> <p>Description:</p> <p>Storage Server: CIFSServerA (10.65.0.49)</p> <p>Storage Pool: CIFSSharePool</p> <p>Size: 200.000 GB (Primary), snapshot support disabled</p> <p>Thin: Disabled</p> <p>Advanced Attributes: ► Use Defaults</p> <p>Protection Schedule: None configured</p> <hr/> <p>Share Creation</p> <p>Share Type: Windows shares (CIFS)</p> <p>Name: UserProfile</p> <p>Description:</p> <p>Local Path: /</p> <p>Export Path: \\10.65.0.49\\UserProfile</p> <p>Advanced Attributes: ► Use Defaults</p>

6.9.6. CIFS Share Storage Creation – vDisk Share

Instructions	Visual
<p>Navigate to Storage->Shared Folders in Unisphere. Select Create to provision CIFS storage.</p>	 <p>VNXe > Storage > Shared Folders</p> <p>Shared Folder Storage</p> <p>Allocated Shared Folders:</p> <p>Create Add Share Create a Replication Destination</p>

Instructions	Visual																														
Specify a name and description (optional) for the shared folder.	<div><div>Shared Folder Wizard</div><div><div></div><div><h2>Specify Shared Folder Name</h2><p>Step 1 of 7</p><p>Specify a name and optional description for the shared folder:</p><div><div>Name: *</div><div><input type="text" value="vDisk"/></div></div><div><div>Description:</div><div><input type="text"/></div></div></div></div></div>																														
Specify a storage server (CIFS server defined in the previous steps), a storage pool with enough space available, and the size of the CIFS share to be created.	<div><div>Shared Folder Wizard</div><div><div></div><div><h2>Configure Shared Folder Storage</h2><p>Step 2 of 7</p><p>Configure the storage for this shared folder:</p><p>Select a storage pool with available space on the selected shared folder server.</p><div><div>Storage Server:</div><div><input type="text" value="CIFSServerA (SP A)"/></div><div>More information...</div></div><table><thead><tr><th>Type</th><th>1 ▲</th><th>Pool</th><th>Available</th><th>Percent Used</th></tr></thead><tbody><tr><td>SAS</td><td></td><td>CIFSSharePool</td><td>5.779 TB</td><td><div><div></div></div>3%</td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr></tbody></table><div><div>Percent Available:</div><div><div></div></div><div>Percent Used</div></div><div><div>Size: *</div><div><input type="text" value="500"/></div><div><div>GB</div><div></div></div></div><div><input type="checkbox"/> Thin</div></div></div></div>	Type	1 ▲	Pool	Available	Percent Used	SAS		CIFSSharePool	5.779 TB	<div><div></div></div> 3%																				
Type	1 ▲	Pool	Available	Percent Used																											
SAS		CIFSSharePool	5.779 TB	<div><div></div></div> 3%																											
Choose the protection policy for replication and snapshots that fits your requirements.	<div><div>Shared Folder Wizard</div><div><div></div><div><h2>Configure Protection</h2><p>Step 3 of 7</p><p>Configure protection storage for replication and snapshots:</p><div><div><input checked="" type="radio"/> Do not configure protection storage for this storage resource.</div><div>Replication and snapshots can be supported by allocating protection space at a later time.</div><div><input type="radio"/> Configure protection storage, do not configure a snapshot protection schedule.</div><div>An automated snapshot protection schedule may be configured at a later time.</div><div><input type="radio"/> Configure protection storage, protect data using snapshot schedule:</div><div><div>Default Protection</div></div></div><div><div>This schedule will create snapshots</div><div>Every day at 04:00, keep for 2 days</div></div><div>Note: Times are displayed in Local Time (UTC-0400) in 24-hour format</div></div></div></div>																														

Instructions	Visual
<p>Select Windows shares (CIFS) as the share type.</p>	
<p>Select Create a Windows share and specify a share name of your choice.</p>	

Instructions	Visual
<p>Review the shared folder configuration in the Summary page and select Finish to provision the CIFS share.</p>	 <p>Shared Folder Wizard</p> <p>Summary</p> <p>Step 6 of 7</p> <p>Confirm you want to create the following shared folder:</p> <p>Name: vDisk</p> <p>Description:</p> <p>Storage Server: CIFSServerA (10.65.0.49)</p> <p>Storage Pool: CIFSSharePool</p> <p>Size: 500.000 GB (Primary), snapshot support disabled</p> <p>Thin: Disabled</p> <p>Advanced Attributes: ► Use Defaults</p> <p>Protection Schedule: None configured</p> <hr/> <p>Share Creation</p> <p>Share Type: Windows shares (CIFS)</p> <p>Name: vDisk</p> <p>Description:</p> <p>Local Path: /</p> <p>Export Path: \\10.65.0.49\vDisk</p> <p>Advanced Attributes: ► Use Defaults</p>

6.10. Installing and Configuring Microsoft Server 2012

Two types of service profiles were required to support two different blade server types:

Table 11: Role/Server/OS Deployment

Role	Blade Server Used	Operating System Deployed
Infrastructure	UCS B200 M3 (E5-2650)	Microsoft Windows Server 2012 Datacenter
VDI Hosts	UCS B200 M3 (E5-2697v2)	Microsoft Windows Server Hyper-V 2012

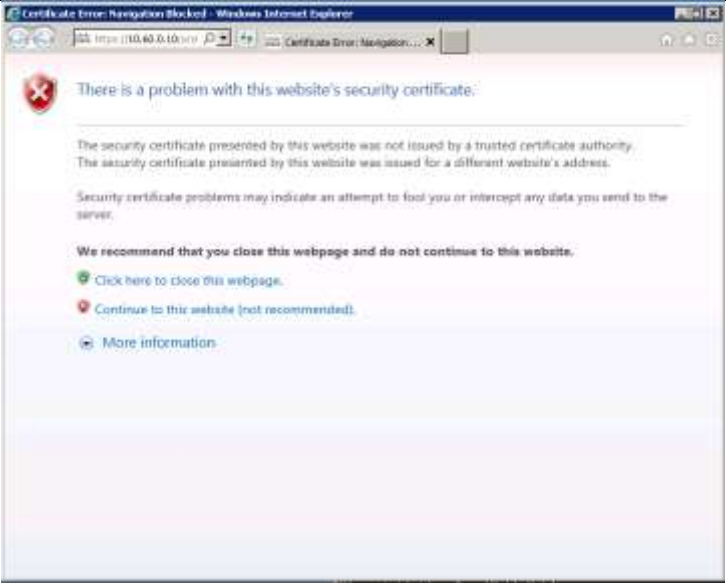

To support those different hardware platforms, service profile templates were created, utilizing various policies created earlier as documented.

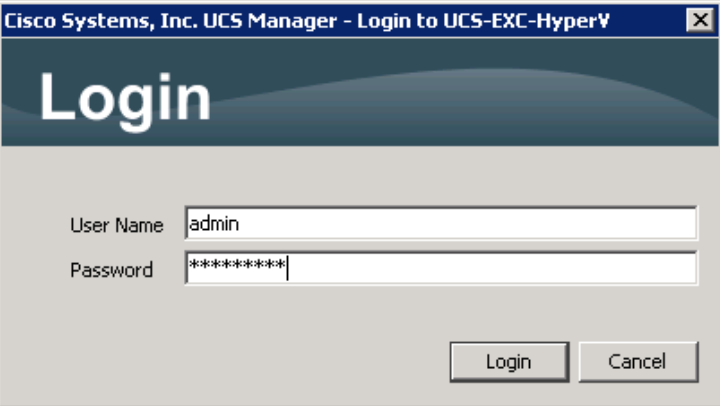
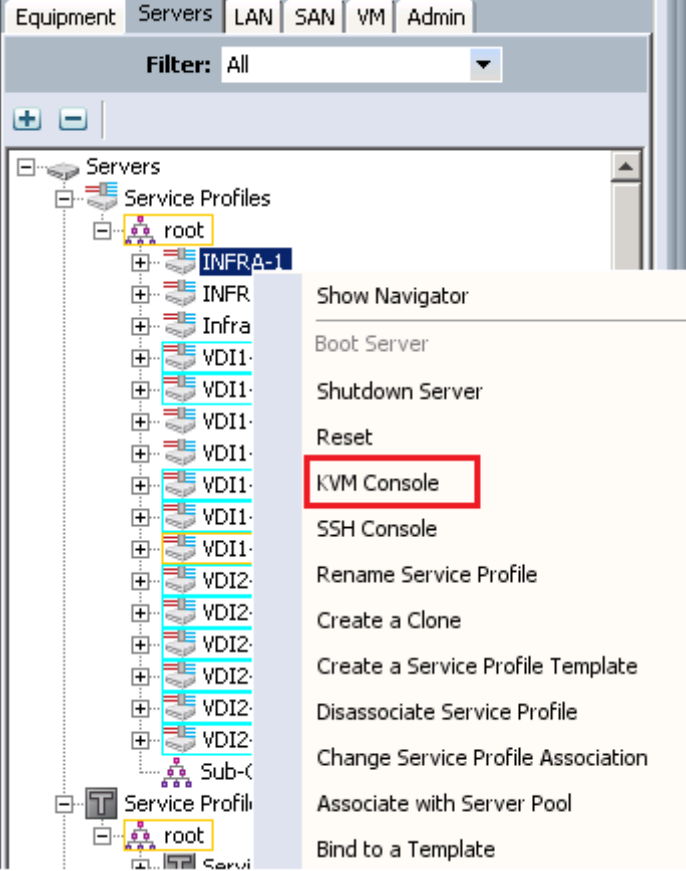
The service profile templates were then used to quickly deploy service profiles for each blade server in the Cisco Unified Computing System. When each blade server booted for the first time, the service profile was deployed automatically, providing the perfect configuration for Microsoft Windows installation.

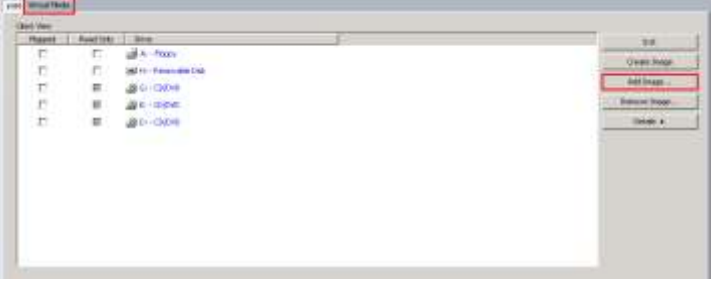
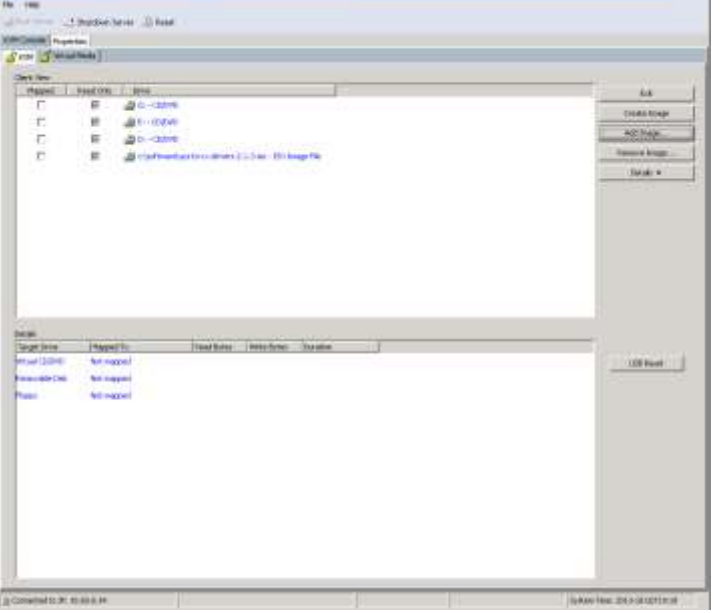
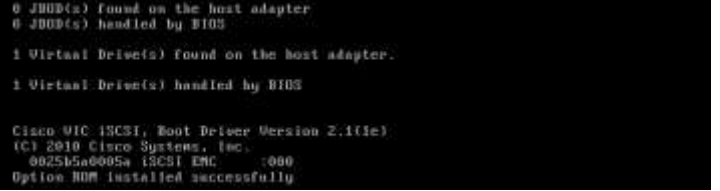
6.10.1. Infrastructure Servers



For this Cisco Validated Design, iSCSI storage was used to boot the hosts from LUNs on the VNXe3300 storage system. Prior to installing the operating system, storage groups were created, assigning to specific boot LUNs to individual hosts. (See Section 6.9 EMC VNXe Storage Configuration for details.)

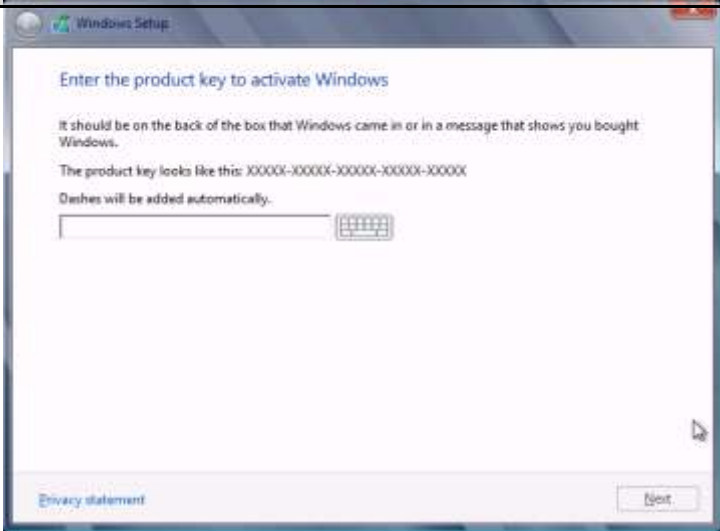
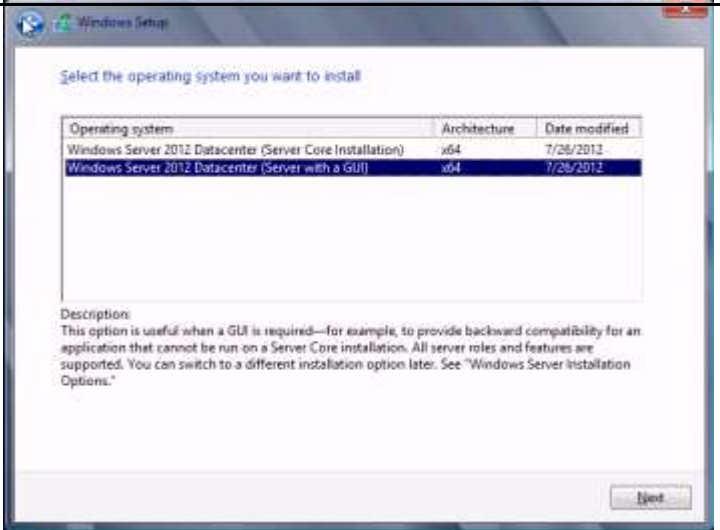
The table below provides the steps to install Windows Server 2012 on the Cisco UCS hosts.

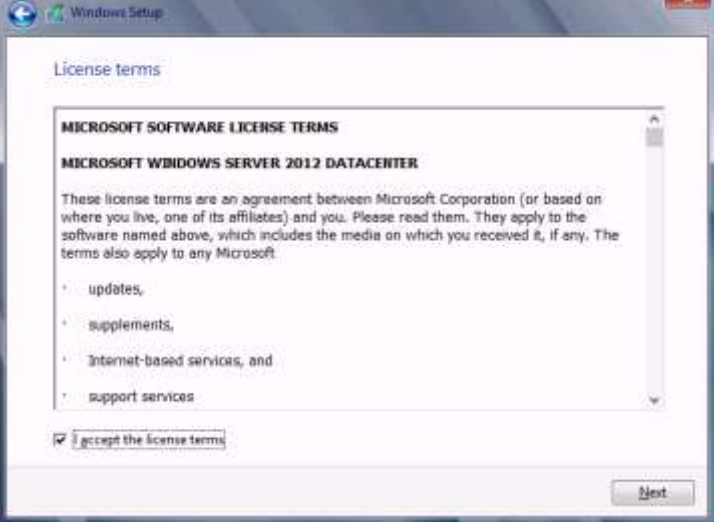
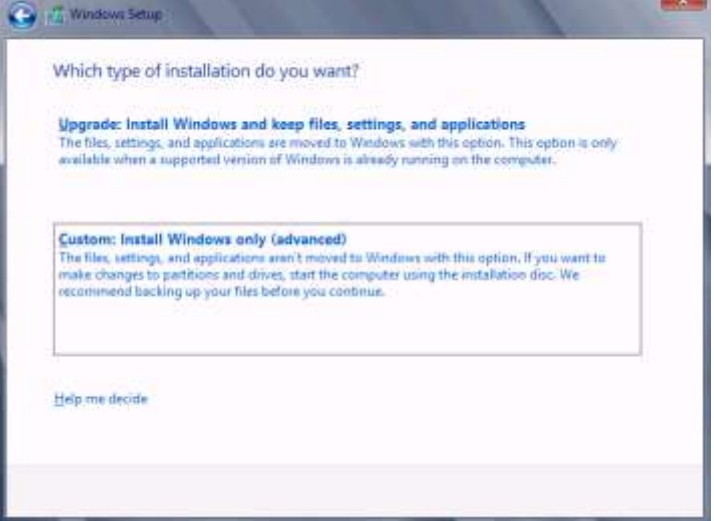
Instructions	Visual
<p>Open a browser and enter the address of the Cisco UCS Fabric Interconnect: https://xx.xx.xx.xx</p> <p>Click on Continue to this website (not recommended).</p>	
<p>Click Launch Cisco UCS Manager.</p>	

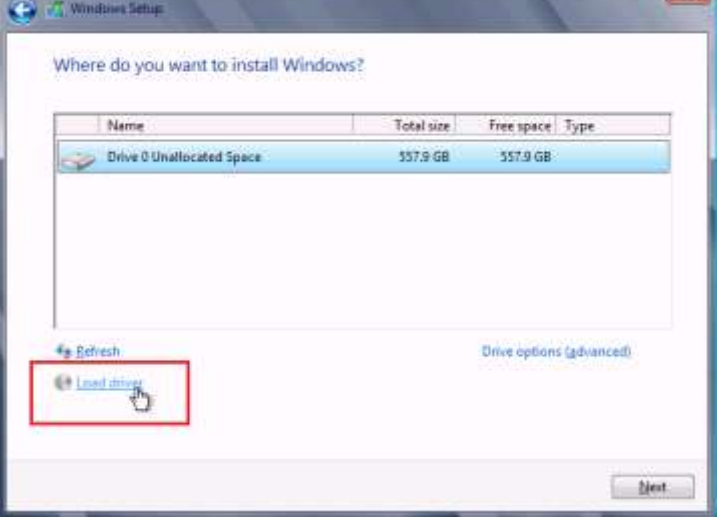
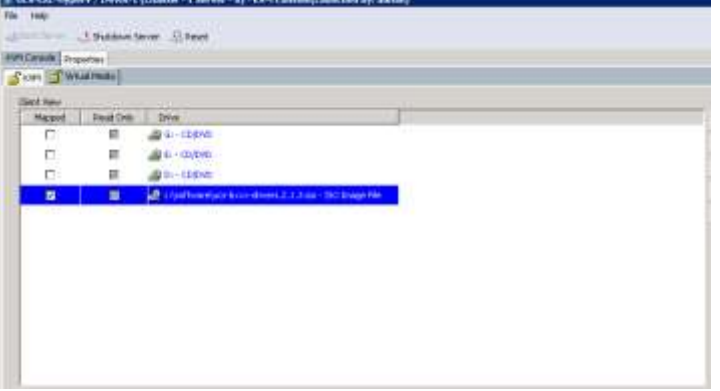
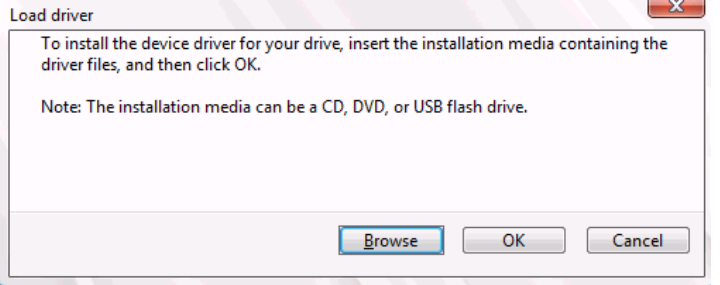
Instructions	Visual
<p>Enter the User Name: admin</p> <p>Enter the Password specified during the initial setup of the Cisco UCS System.</p> <p>Click Login.</p>	
<p>In the Servers tab, expand Service Profiles > root. Select the server you wish to install the operating system on. Right-click and choose KVM Console from the context menu.</p> <p>If you get a certificate warning, just click the “Always trust this certificate” checkbox and then click Run.</p>	
<p>Click the Virtual Media tab of the KVM Console.</p>	

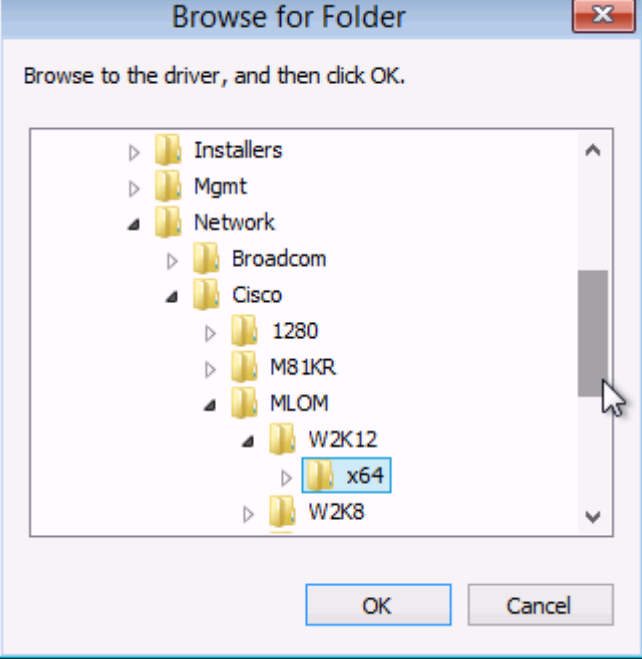
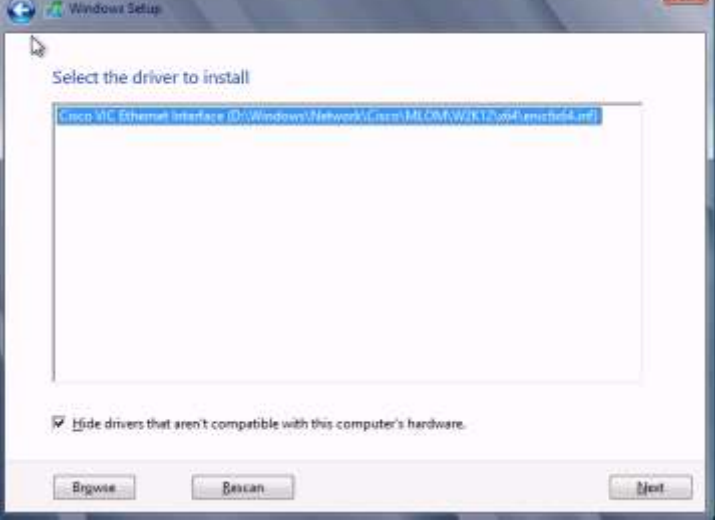
Instructions	Visual
<p>Click the Add Image... button.</p> <p>Add the image for the Windows Server 2012 ISO image by browsing to the ISO and clicking Open.</p>	
<p>Repeat the process for the ucs-bxxxx-drivers 2.1.3 ISO image.</p> <p>Click the Mapped checkbox for the Windows Server 2012 ISO image.</p>	
<p>NOTE: Due to a small bug in the Windows Installation when using iSCSI boot LUNs, you will need to ensure there is only one path to the iSCSI volume. If there are two paths the install will fail until you first write something to the disk.</p> <p>To disable one of the paths, just assign the eth7 vNIC temporarily to the iSCSI-Null VLAN created earlier.</p>	
<p>Reboot the Cisco UCS server and verify the iSCSI boot driver loads and connects to the EMC VNXe3300 LUN.</p>	 <pre> 0 JBOD(s) found on the host adapter 0 JBOD(s) handled by BIOS 1 Virtual Drive(s) found on the host adapter. 1 Virtual Drive(s) handled by BIOS Cisco VIC iSCSI, Boot Driver Version 2.1(1c) (C) 2010 Cisco Systems, Inc. 0025b5a6005a iSCSI EMC :000 Option ROM installed successfully </pre>


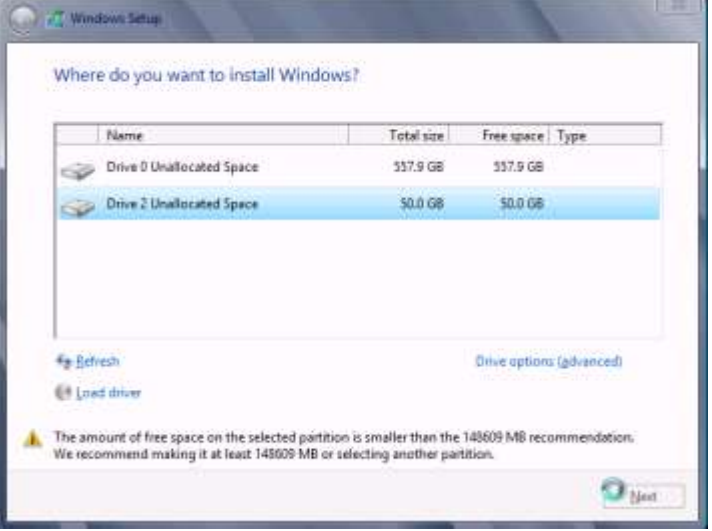
Instructions	Visual
<p>Select Language to install.</p> <p>Select Time and currency format.</p> <p>Select Keyboard or input method.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Windows Setup' window for Windows Server 2012. The window has a dark blue background with the Windows logo and 'Windows Server 2012' text. Below this, there are three dropdown menus: 'Language to install' (set to English (United States)), 'Time and currency format' (set to English (United States)), and 'Keyboard or input method' (set to US). At the bottom, there is a 'Next' button and a copyright notice: '© 2012 Microsoft Corporation. All rights reserved.'</p>
<p>Click Install now.</p>	 <p>The screenshot shows the 'Windows Setup' window for Windows Server 2012, but it is a smaller window within a larger application window. The background is dark blue with the Windows logo and 'Windows Server 2012' text. In the center, there is a large 'Install now' button. Below the button, there is a link that says 'Repair your computer' and a copyright notice: '© 2012 Microsoft Corporation. All rights reserved.'</p>


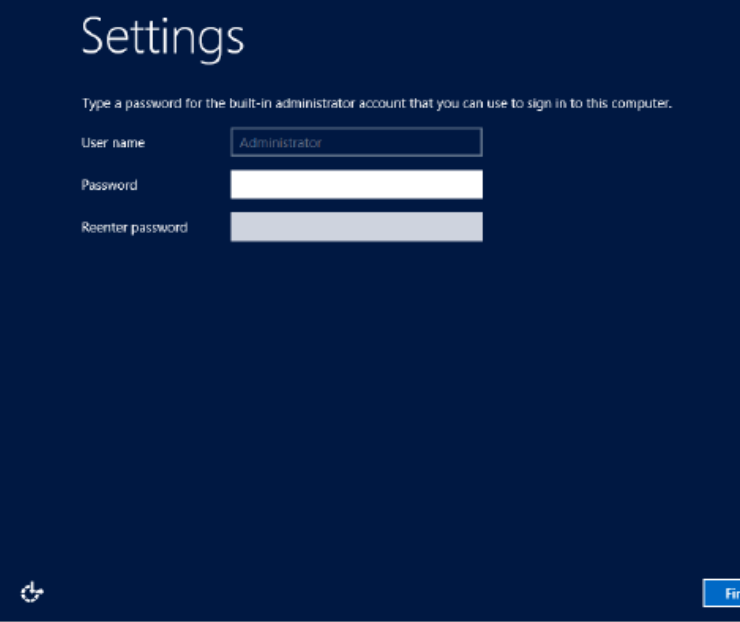
Instructions	Visual									
<p>Provide the Windows Server 2012 license key.</p> <p>Click Next.</p>	 <p>Windows Setup</p> <p>Enter the product key to activate Windows</p> <p>It should be on the back of the box that Windows came in or in a message that shows you bought Windows.</p> <p>The product key looks like this: XXXXX-XXXXX-XXXXX-XXXXX-XXXXX</p> <p>Dashes will be added automatically.</p> <p>Privacy statement</p> <p>Next</p>									
<p>Select Windows Server 2012 Datacenter (Server with a GUI).</p> <p>Click Next.</p>	 <p>Windows Setup</p> <p>Select the operating system you want to install</p> <table><thead><tr><th>Operating system</th><th>Architecture</th><th>Date modified</th></tr></thead><tbody><tr><td>Windows Server 2012 Datacenter (Server Core Installation)</td><td>x64</td><td>7/26/2012</td></tr><tr><td>Windows Server 2012 Datacenter (Server with a GUI)</td><td>x64</td><td>7/26/2012</td></tr></tbody></table> <p>Description: This option is useful when a GUI is required—for example, to provide backward compatibility for an application that cannot be run on a Server Core installation. All server roles and features are supported. You can switch to a different installation option later. See "Windows Server installation Options."</p> <p>Next</p>	Operating system	Architecture	Date modified	Windows Server 2012 Datacenter (Server Core Installation)	x64	7/26/2012	Windows Server 2012 Datacenter (Server with a GUI)	x64	7/26/2012
Operating system	Architecture	Date modified								
Windows Server 2012 Datacenter (Server Core Installation)	x64	7/26/2012								
Windows Server 2012 Datacenter (Server with a GUI)	x64	7/26/2012								

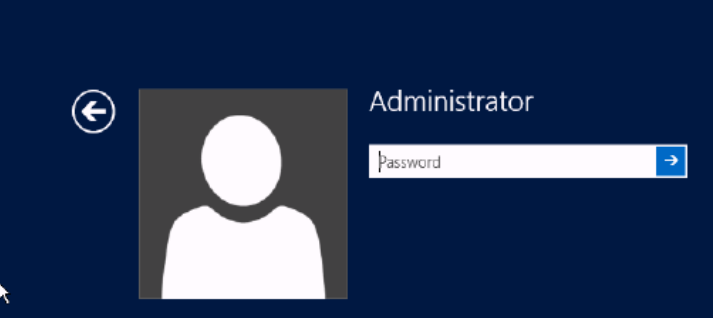
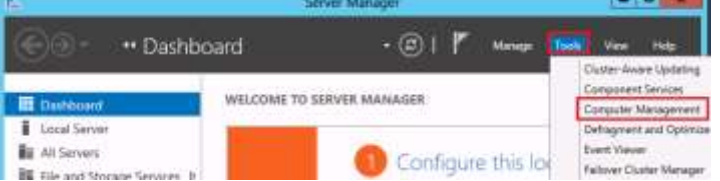
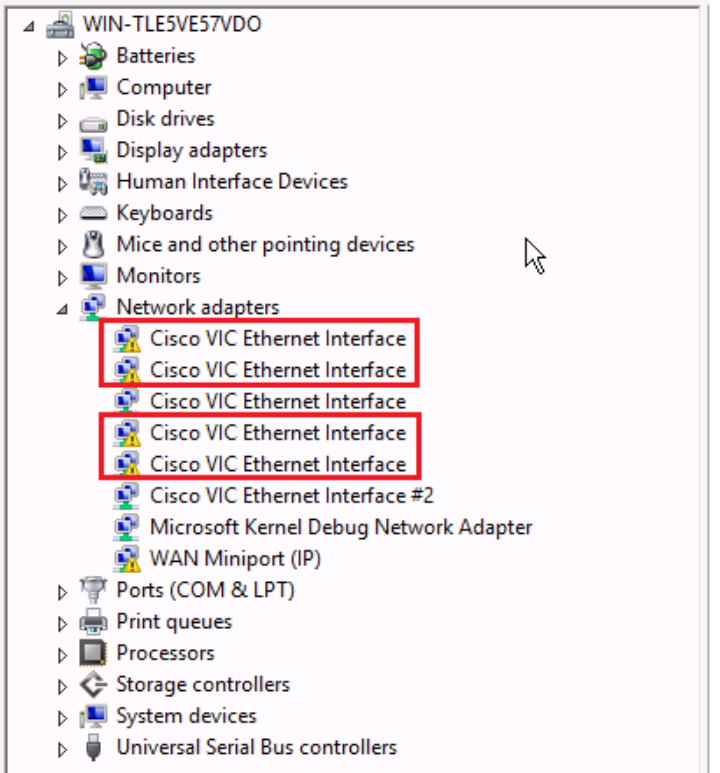
Instructions	Visual
<p>Review the License terms, and if agreeable, enable the checkbox labeled I accept the license terms.</p> <p>You must accept the license terms to continue.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Windows Setup' window with the 'License terms' tab selected. It displays the 'MICROSOFT SOFTWARE LICENSE TERMS' for 'MICROSOFT WINDOWS SERVER 2012 DATACENTER'. The terms are listed in a scrollable box, including updates, supplements, Internet-based services, and support services. At the bottom, the checkbox 'I accept the license terms' is checked, and a 'Next' button is visible in the bottom right corner.</p>
<p>Click Custom: Install Windows only (advanced).</p>	 <p>The screenshot shows the 'Windows Setup' window with the question 'Which type of installation do you want?'. Two options are presented: 'Upgrade: Install Windows and keep files, settings, and applications' and 'Custom: Install Windows only (advanced)'. The 'Custom' option is highlighted with a red border. Below the options is a 'Help me decide' link.</p>

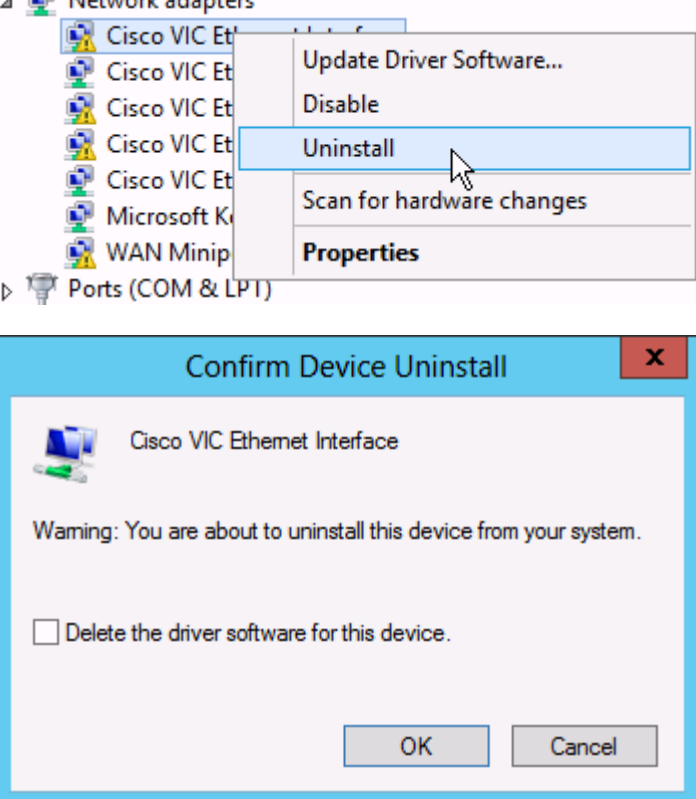
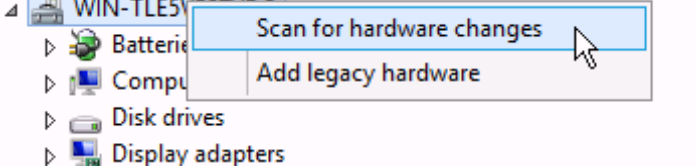
Instructions	Visual
<p>The local disk will be visible, but the iSCSI volume needs the Cisco UCS disk driver loaded.</p> <p>Click Load driver.</p>	
<p>Before loading the driver, you will need to insert the Cisco B-Series 2.1(3a) device driver CD.</p> <p>Select the Virtual Media tab.</p> <p>Uncheck the Mapped box for the Windows Server ISO (and confirm the unmap warning dialog).</p> <p>Check the Mapped box for the Cisco UCS driver CD.</p> <p>Return to the KVM tab.</p>	
<p>Click Browse to locate the driver install.</p>	

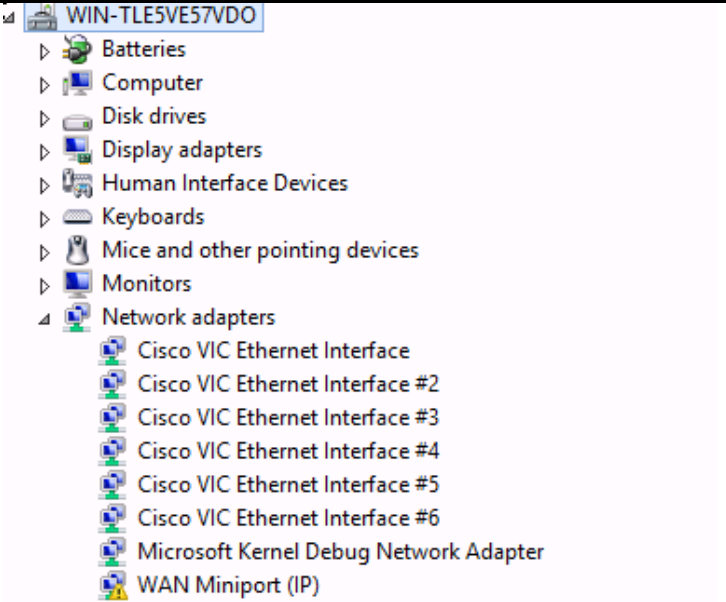
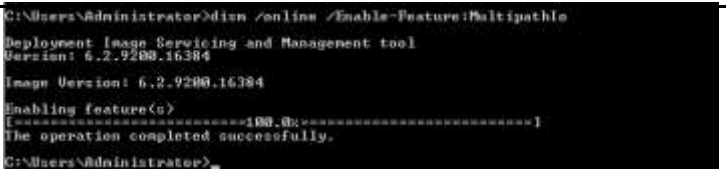
Instructions	Visual
<p>Navigate to the MLOM Windows Server 2012 x64 driver found at:</p> <p>.\\Windows\\Network\\Cisco\\MLOM\\W2K12\\x64</p> <p>Click OK.</p>	
<p>When the driver is located, click Next.</p>	

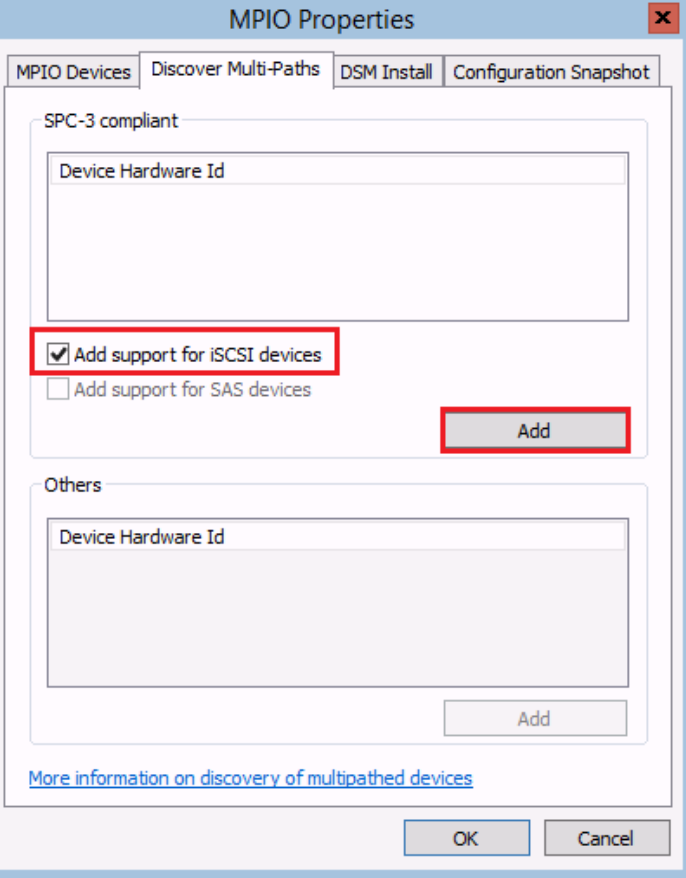
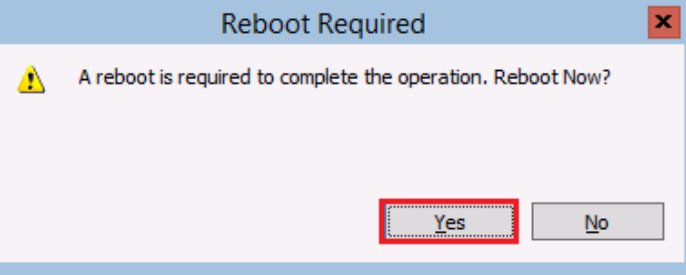
Instructions	Visual
<p>Before returning to the install, be sure to reinsert the Windows Server 2012 ISO.</p> <p>Select the Virtual Media tab.</p> <p>Uncheck the Mapped box for the Cisco UCS driver CD.</p> <p>Check the Mapped box for the Windows Server ISO</p> <p>Return to the KVM tab.</p>	
<p>Select the iSCSI LUN.</p> <p>Click Next.</p>	

Instructions	Visual
<p>Complete the Windows installation.</p>	 <p>The screenshot shows the 'Windows Setup' window. The title bar says 'Windows Setup'. The main heading is 'Installing Windows'. Below it, a message states: 'Your computer will restart several times. This might take a while.' A progress bar is shown with a green checkmark icon. The progress bar is divided into four segments: 'Copying Windows files' (completed), 'Getting files ready for installation (23%)' (current step), 'Installing features', 'Installing updates', and 'Finishing up'.</p>
<p>Enter the initial Administrator's password.</p> <p>Re-enter the Administrator's password.</p>	 <p>The screenshot shows the 'Settings' window for creating a password. The title bar says 'Settings'. The main heading is 'Settings'. Below it, a message states: 'Type a password for the built-in administrator account that you can use to sign in to this computer.' There are three input fields: 'User name' (containing 'Administrator'), 'Password', and 'Reenter password'. A blue 'Finish' button is visible in the bottom right corner.</p>

Instructions	Visual
<p>Login to the server.</p>	
<p>Launch the Computer Management tool from Server Manager.</p> <p>Server Manager >> Tools >> Computer Management.</p>	
<p>Select Device Manager from the left-hand pane.</p> <p>From the right-hand pane, locate the Cisco VIC Ethernet Interface cards that are showing a warning symbol.</p> <p>The Cisco VIC Ethernet Interfaces without the warning symbol are the two iSCSI interfaces. The remaining ones will need to be uninstalled.</p> <p>WARNING: <i>Uninstalling one of the Cisco VIC Ethernet Interfaces that is being used will blue screen the box and require a complete reinstall.</i></p> <p>Right-click the Cisco VIC Ethernet Interface.</p> <p>Select Uninstall from the context menu.</p>	

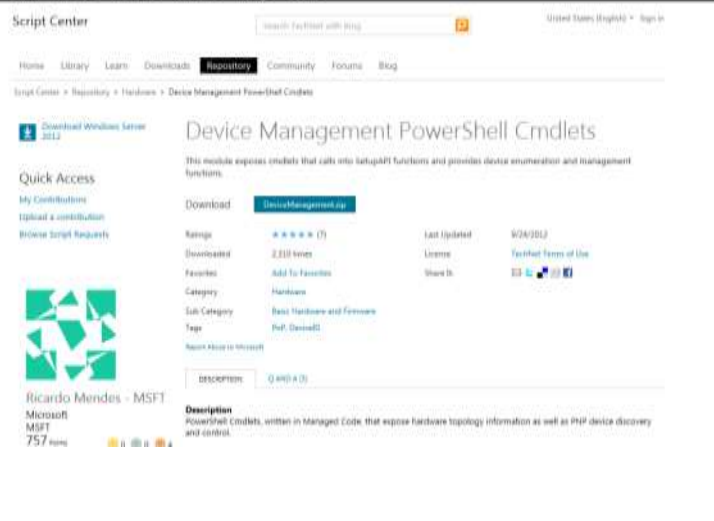

Instructions	Visual
<p>If prompted, click OK on the Confirm Device Uninstall dialog.</p> <p>Repeat for all Cisco VIC Ethernet Interfaces which are having problems.</p>	 <p>The visual shows two screenshots. The top screenshot is a Windows Device Manager window with 'Network adapters' expanded. A context menu is open for a 'Cisco VIC Ethernet Interface', showing options: 'Update Driver Software...', 'Disable', 'Uninstall' (highlighted), 'Scan for hardware changes', and 'Properties'. The bottom screenshot is a 'Confirm Device Uninstall' dialog box for the 'Cisco VIC Ethernet Interface'. It contains a warning: 'Warning: You are about to uninstall this device from your system.' and a checkbox labeled 'Delete the driver software for this device.' which is unchecked. 'OK' and 'Cancel' buttons are at the bottom right.</p>
<p>To reinstall the drivers correctly, complete the following:</p> <p>Select the computer name.</p> <p>Right-click and choose Scan for hardware changes from the context menu.</p> <p>When finished, all Cisco VIC Ethernet Interfaces should be functioning correctly.</p>	 <p>The visual shows a screenshot of the Windows Device Manager window. The 'WIN-TLE5V' category is selected, and a context menu is open over it. The menu options are 'Scan for hardware changes' (highlighted) and 'Add legacy hardware'. Other categories visible in the background include 'Batteries', 'Computers', 'Disk drives', and 'Display adapters'.</p>

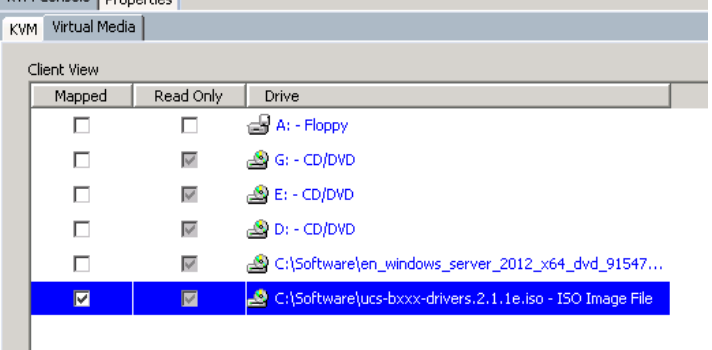
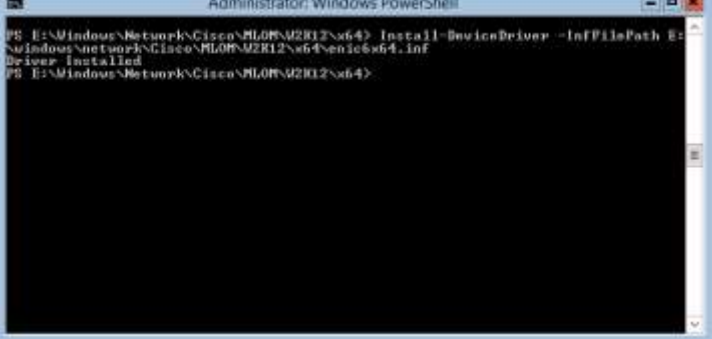
Instructions	Visual
	 <p>WIN-TLE5VE57VDO</p> <ul style="list-style-type: none"> Batteries Computer Disk drives Display adapters Human Interface Devices Keyboards Mice and other pointing devices Monitors Network adapters <ul style="list-style-type: none"> Cisco VIC Ethernet Interface Cisco VIC Ethernet Interface #2 Cisco VIC Ethernet Interface #3 Cisco VIC Ethernet Interface #4 Cisco VIC Ethernet Interface #5 Cisco VIC Ethernet Interface #6 Microsoft Kernel Debug Network Adapter WAN Miniport (IP)
<p>Note: To re-enable one of the paths, just assign the eth7 vNIC back to the iSCSI-B VLAN created earlier. Ideally, this step should be completed prior to enabling the MultipathIO software.</p>	
<p>Install Multipathing software by executing the following CaSe-SeNsItIvE command-line:</p> <p>Dism /online /Enable-Feature:Multipathio</p>	 <pre> C:\Users\Administrator>dism /online /Enable-Feature:Multipathio Deployment Image Servicing and Management tool Version: 6.2.9200.16384 Image Version: 6.2.9200.16384 Enabling feature(s) [Success] The operation completed successfully. C:\Users\Administrator> </pre>

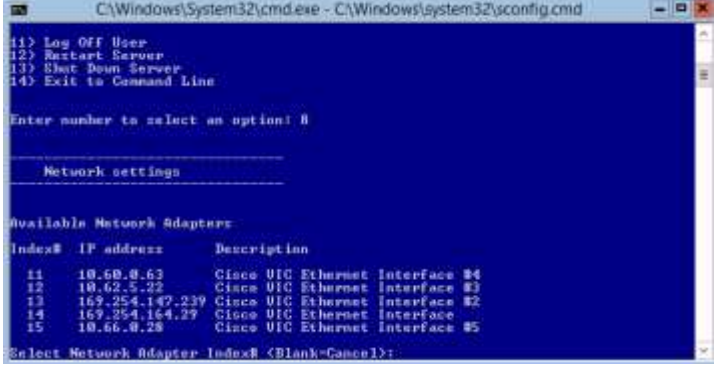
Instructions	Visual
<p>Enable Multipathing software from the MPIO control panel.</p> <p>From the command-line execute Mpiocpl.</p> <p>Enable the Add support for iSCSI devices checkbox.</p> <p>Click Add.</p>	
<p>Click “Yes” on the Reboot Required dialog.</p> <p>Repeat steps for the other Infrastructure server.</p>	

6.10.2. VDI Hosts

Windows Hyper-V 2012 does not include the full GUI, the steps previously identified to force the Cisco VIC Ethernet Interface drivers needs to be accomplished using PowerShell instead. Microsoft has a Device Management PowerShell cmdlet which is able to perform the VIC Ethernet Interface update properly.

Instructions	Visual
<p>Download the Device Management PowerShell Cmdlets from: http://gallery.technet.microsoft.com/scriptcenter/Device-Management-7fad2388</p> <p>When downloaded, they will need to be copied to the Hyper-V 2012 server. Since only the iSCSI interfaces are available, they can be copied through the iSCSI-A network or the can be burned to an ISO format and mounted through the Virtual Media tab.</p>	 <p>The screenshot shows the 'Script Center' page for 'Device Management PowerShell Cmdlets'. It includes a search bar, navigation tabs (Home, Library, Learn, Downloads, Repository, Community, Forums, Blog), and a download button for 'DeviceManagement.ps'. The page also displays a description of the module, its version (757), and a download count (2,119 times).</p>
<p>Copy the PowerShell Cmdlets to:</p> <p>C:\Windows\System32\WindowsPowerShell\v1.0\Modules\devmgmt</p> <p>Import the PowerShell Cmdlets from the devmgmt folder with this command:</p> <p>import-module .\DeviceManagement.psd1</p>	 <p>The screenshot shows a Windows PowerShell command prompt window. The user has navigated to the directory 'C:\Windows\System32\WindowsPowerShell\v1.0\Modules\devmgmt' and executed the command 'import-module .\DeviceManagement.psd1'. The output shows the files being copied and the module being imported successfully.</p>

Instructions	Visual
<p>Before loading the driver, you will need to insert the Cisco B-Series 2.1(3a) device driver CD.</p> <p>Select the Virtual Media tab.</p> <p>Check the Mapped box for the Cisco UCS driver CD.</p> <p>Return to the KVM tab.</p>	 <p>The screenshot shows the 'KVM Virtual Media' window with the 'Client View' tab selected. It displays a table with three columns: 'Mapped', 'Read Only', and 'Drive'. The 'Mapped' checkbox for the last row, 'C:\Software\ucs-bxxx-drivers.2.1.1e.iso - ISO Image File', is checked. The other rows represent standard drives (A:, G:, E:, D:) and a Windows server DVD.</p>
<p>Execute the Install-DeviceDriver cmdlet with this command:</p> <p>Install-DeviceDriver -InfFilePath E:\Windows\Network\Cisco\MLOM\W2K12\x64\enic6x64.inf</p> <p>Where E: is the drive letter for the mapped Cisco UCS ISO.</p>	 <p>The screenshot shows a Windows PowerShell window with the following command and output:</p> <pre>PS E:\Windows\Network\Cisco\MLOM\W2K12\x64> Install-DeviceDriver -InfFilePath E:\Windows\Network\Cisco\MLOM\W2K12\x64\enic6x64.inf Driver installed PS E:\Windows\Network\Cisco\MLOM\W2K12\x64></pre>

Instructions	Visual																		
<p>After the driver is installed correctly, all the Cisco VIC Ethernet Interface adapters will be visible from SCONFIG option 8.</p>	 <p>The screenshot shows a Windows command prompt window titled "C:\Windows\System32\cmd.exe - C:\Windows\system32\sconfig.cmd". The menu options are: 1) Log Off User, 2) Restart Server, 3) Shut Down Server, 4) Exit to Command Line. Below these is a prompt "Enter number to select an option: 8". The "Network settings" section is active, showing "Available Network Adapters". A table lists five adapters with their Index#, IP address, and Description.</p> <table><thead><tr><th>Index#</th><th>IP address</th><th>Description</th></tr></thead><tbody><tr><td>11</td><td>10.60.0.63</td><td>Cisco VIC Ethernet Interface #4</td></tr><tr><td>12</td><td>10.62.5.22</td><td>Cisco VIC Ethernet Interface #3</td></tr><tr><td>13</td><td>169.254.147.239</td><td>Cisco VIC Ethernet Interface #2</td></tr><tr><td>14</td><td>169.254.164.29</td><td>Cisco VIC Ethernet Interface #1</td></tr><tr><td>15</td><td>10.66.0.28</td><td>Cisco VIC Ethernet Interface #5</td></tr></tbody></table> <p>At the bottom, it says "Select Network Adapter Index# <Blank=Cancel>:".</p>	Index#	IP address	Description	11	10.60.0.63	Cisco VIC Ethernet Interface #4	12	10.62.5.22	Cisco VIC Ethernet Interface #3	13	169.254.147.239	Cisco VIC Ethernet Interface #2	14	169.254.164.29	Cisco VIC Ethernet Interface #1	15	10.66.0.28	Cisco VIC Ethernet Interface #5
Index#	IP address	Description																	
11	10.60.0.63	Cisco VIC Ethernet Interface #4																	
12	10.62.5.22	Cisco VIC Ethernet Interface #3																	
13	169.254.147.239	Cisco VIC Ethernet Interface #2																	
14	169.254.164.29	Cisco VIC Ethernet Interface #1																	
15	10.66.0.28	Cisco VIC Ethernet Interface #5																	

6.10.3. Local Configuration Tasks After Install

When completed, the standard configuration tasks can now be completed on each server, such as naming the server, joining the domain, etc. For this validation, the following configuration tasks were completed after the OS installation was finished:

1. Name the interfaces to match the Cisco UCS Virtual Ethernet Interfaces
 - a. **NOTE:** The order that the vNICs are loaded by Windows is non-deterministic. The best way to identify the interfaces is to match them based on the MAC addresses assigned in Cisco Unified Computing System.
2. Configure static IP address for management and infrastructure networks
3. Configure DNS servers
4. Name the server
5. Join the domain
6. Disable the Firewall
7. Enable Remote Desktop
8. Enable Remote Management
9. Install EMC PowerPath software
10. Windows Update
11. Share the local SSD drive out for XenDesktop write-cache drives, granting everyone access to the drive.

6.11. Installing and Configuring SQL Server 2012 SP1

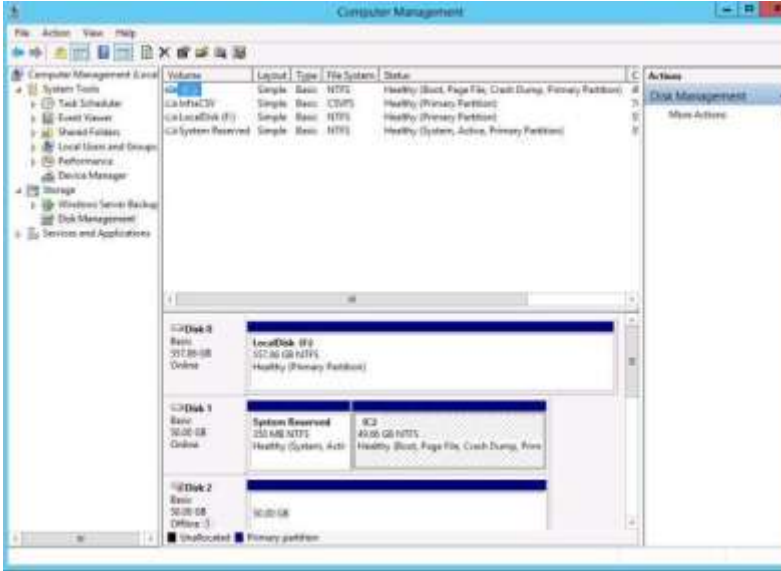
This section provides the instructions for installing and configuring Microsoft SQL Server 2012 SP1 for the System Center Virtual Machine Manager database.

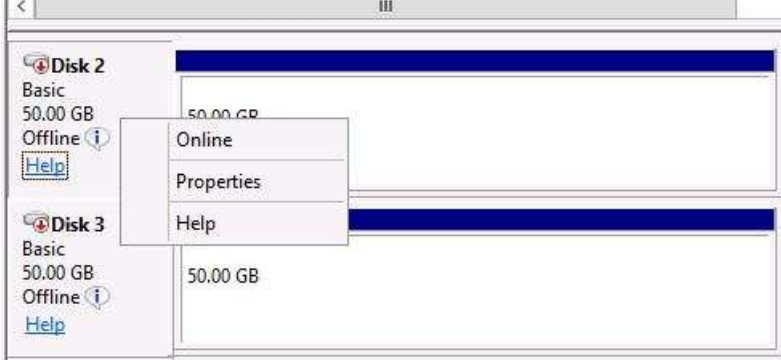


6.11.1. Pass-Through Storage Configuration


Hyper-V allows virtual machines to access storage mapped directly to the Hyper-V server without requiring the volume to be configured. When storage is mapped to the Hyper-V server, it will appear as a raw volume in an offline state in disk manager on the Hyper-V server. Bring the disk online, initialize the disk, and then place the disk in an offline state. To ensure that the Virtual Machine will have exclusive access to the storage, the disk should be left in an offline state. At this point, add the disk as pass-through storage to the SQL VMs using Windows Hyper-V Manager.

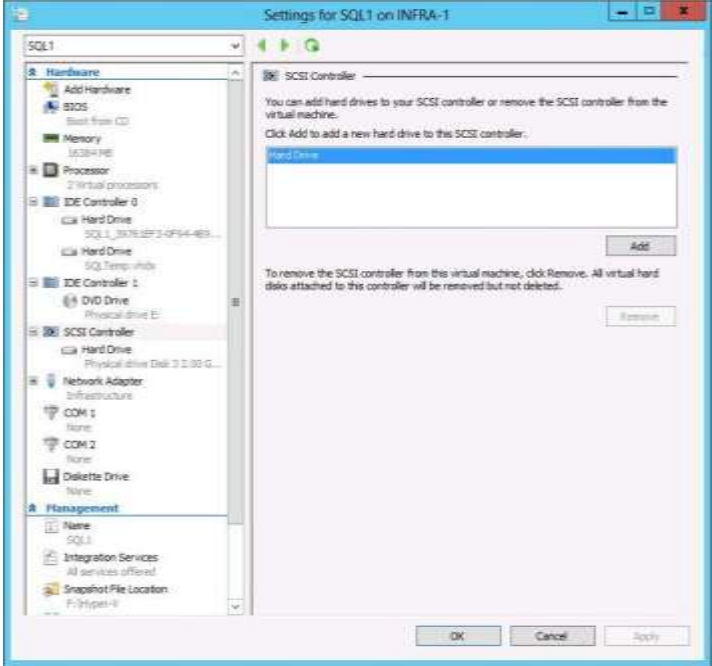
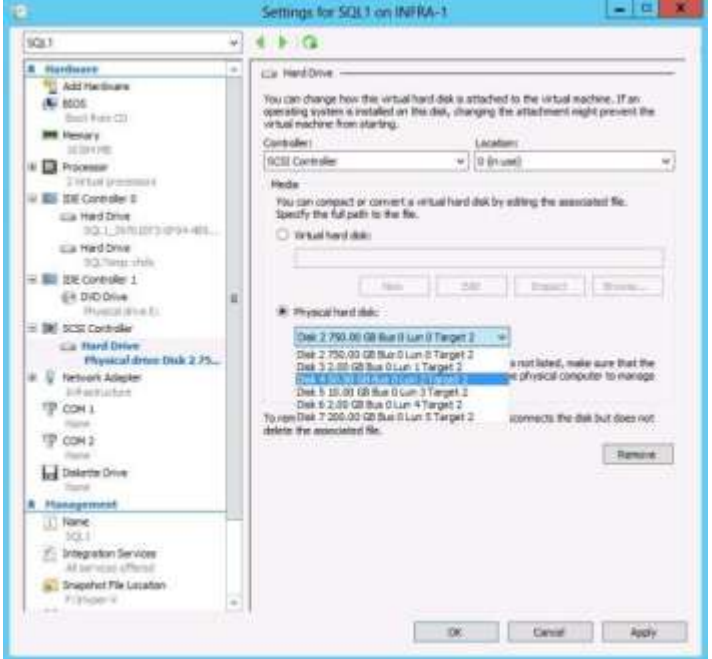
Prior to configuring the storage as pass-through LUNs to SQL VMs SQL 1 and SQL2, you will need to:

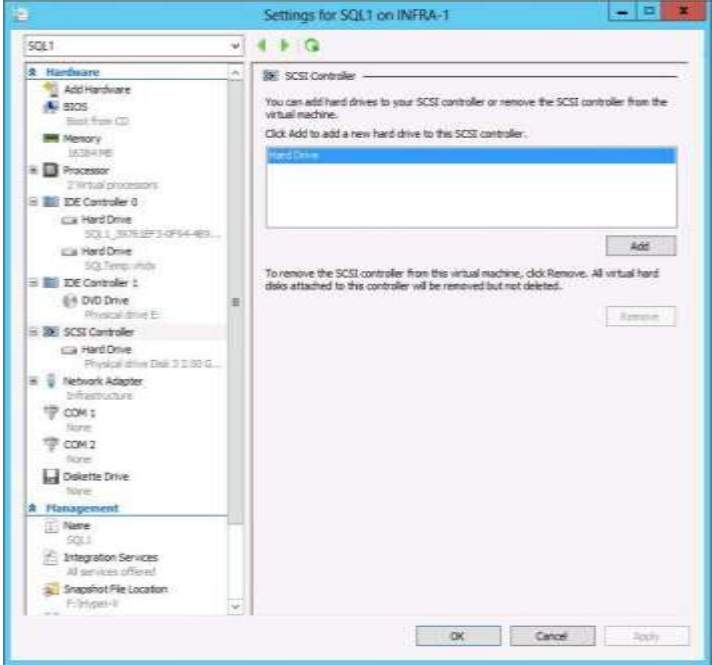
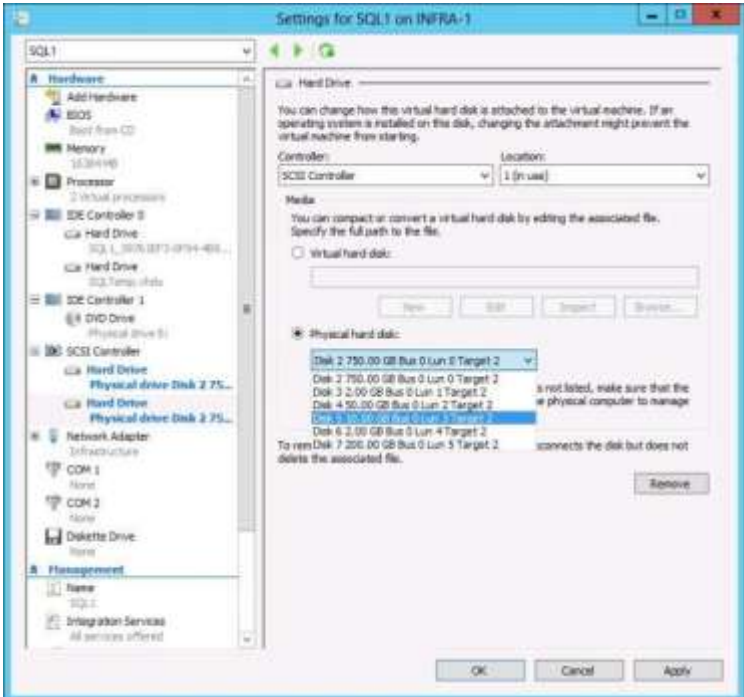
1. Make sure that the LUNs to be used for DB and Log files have been presented to the Hyper-V servers (INFRA-1 and INFRA-2).

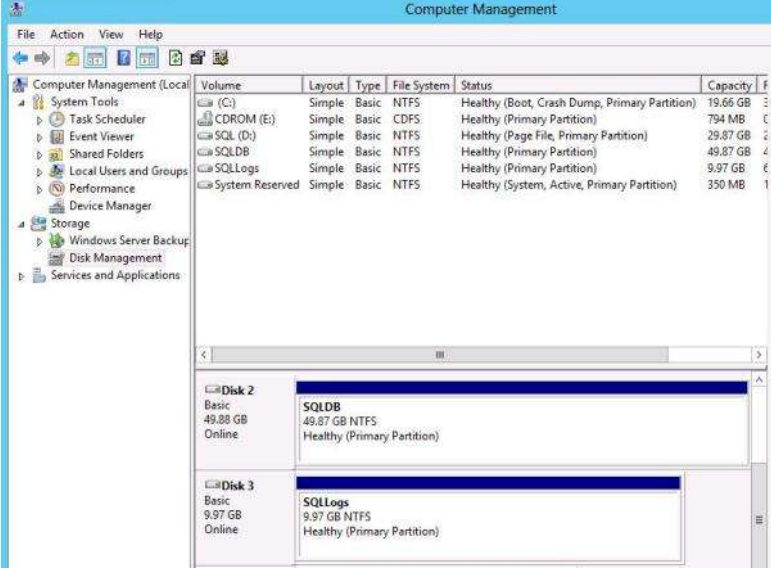
Instructions	Visual
Perform the following steps on both INFRA-1 and INFRA-2 .	
Open the Computer Management Tool . Select Disk Management in the left pane .	

Instructions	Visual
Right-click the disk and select Online .	
When Online , the disk will show as being Not Initialized .	
Right-click the disk and select Initialize Disk .	

Instructions	Visual
<p>Select either the MBR partition or GPT partition radio button type depending on the size of your volume. Volumes over 2TB need to be GPT. When finished, click OK.</p>	
<p>When a disk is initialized, it can once again be placed in an Offline state. If the disk is not in an Offline state, it will not be available for selection when configuring the Guest's storage.</p> <p>Right-click on the disk and select Offline.</p>	
<p>Add the pass-through disk to the virtual machines SQL1 and SQL2.</p> <p>Perform the following steps on both INFRA-1 and INFRA-2:</p>	

Instructions	Visual
<p>Open Windows Hyper-V Manager. Right click on the SQL1 or SQL2 VM and choose Settings.</p> <p>Click SCSI Controller in the left.</p> <p>Click Hard Drive.</p> <p>Click Add.</p>	
<p>Select Physical Hard Disk.</p> <p>In the drop down list, choose the disk to be used for the SQLDB (50GB).</p> <p>Click Apply.</p>	

Instructions	Visual
<p>Click SCSI Controller in the left pane.</p> <p>Click Hard Drive.</p> <p>Click Add.</p>	 <p>The screenshot shows the 'Settings for SQL1 on INFRA-1' window. The left pane is expanded to 'Hardware', and 'SCSI Controller' is selected. The right pane shows the 'SCSI Controller' settings, including a list of hard drives and an 'Add' button.</p>
<p>Select Physical Hard Disk.</p> <p>In the drop down list, choose the disk to be used for the SQLLOGs (10GB).</p> <p>Click Apply.</p>	 <p>The screenshot shows the 'Settings for SQL1 on INFRA-1' window. The left pane is expanded to 'Hardware', and 'SCSI Controller' is selected. The right pane shows the 'Hard Drive' settings, including a list of physical hard disks and a 'Physical hard disk' section with a dropdown menu.</p>

Instructions	Visual
<p>Verify that the pass-through disk has been successfully added to the virtual machines SQL1 and SQL2.</p> <p>Perform the following steps on both SQL-1 and SQL-2:</p>	
<p>Open the Computer Management tool.</p> <p>Click Disk Management in the left pane.</p> <p>Verify that the DB and LOG volumes are shown.</p> <p>Right click on the DB volume and select Online</p> <p>Right click on the LOG volume and select Online</p> <p>Create a new volume and call it SQLDB</p> <p>Create a new volume and call it SQLLogs</p>	

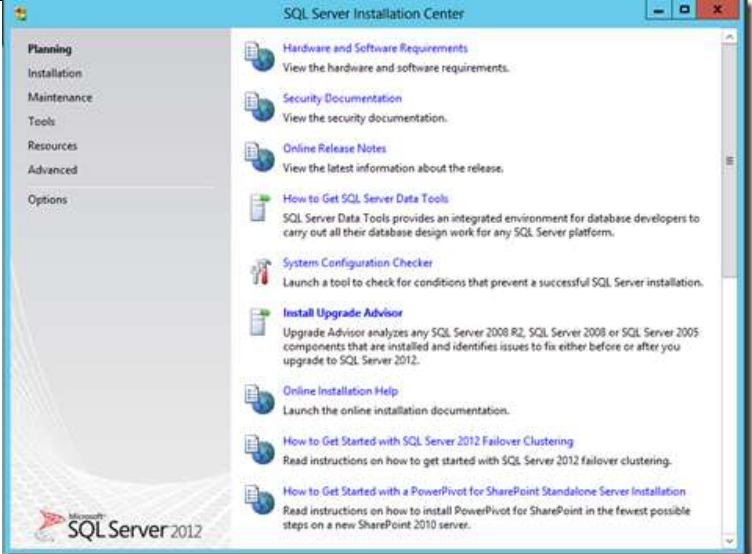
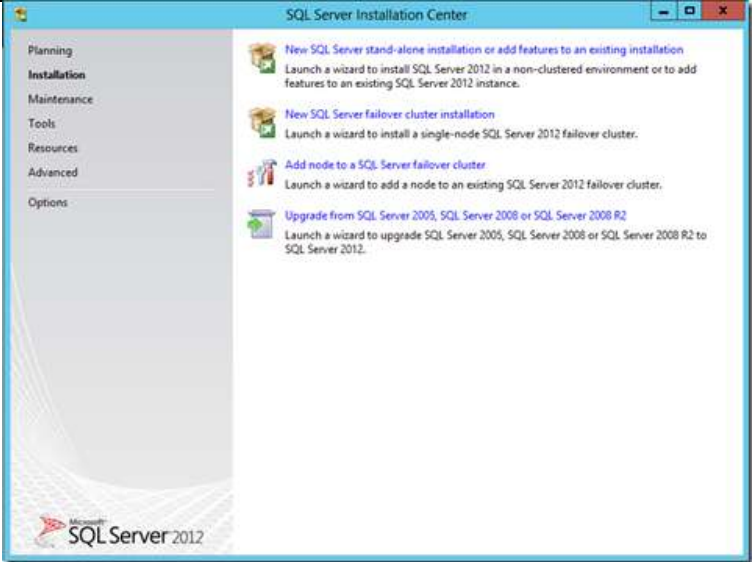
6.11.2. SQL Server 2012 – Installation Pre-requisites

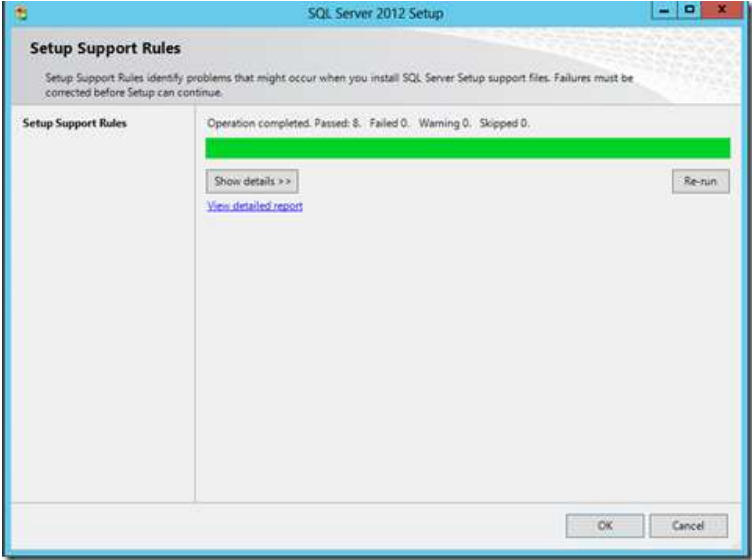
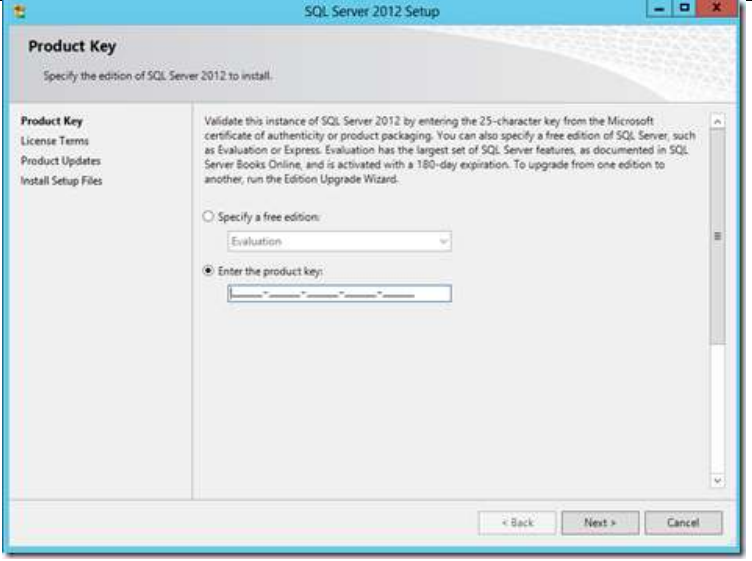
Prior to installing SQL Server 2012 SP1, you will need to do the following on SQL1 and SQL2:

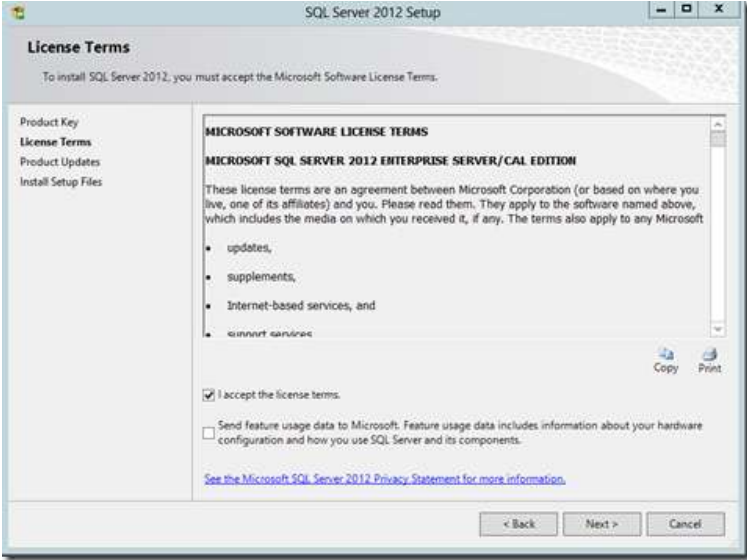
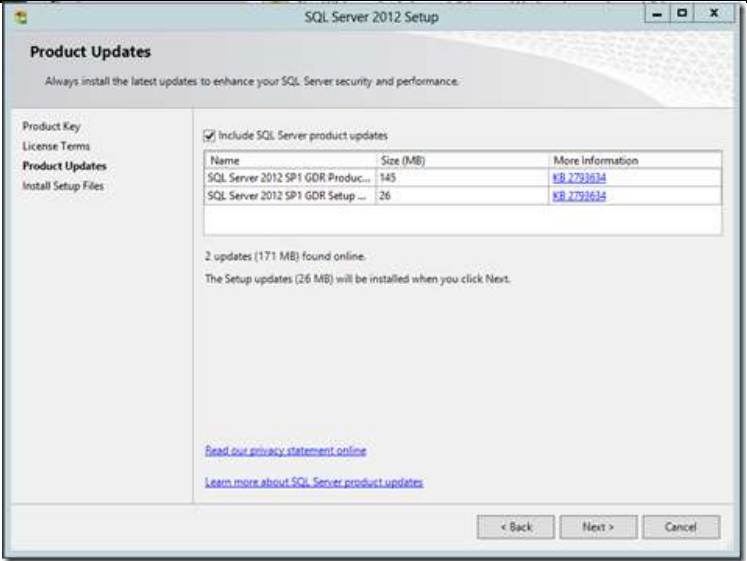
1. Make sure that .NET 3.5 SP1 has been installed on the guest operating system on both SQL1 and SQL2 (it is no longer installed by SQL server setup). For more information please see [Microsoft .NET Framework 3.5 Deployment Considerations](#).

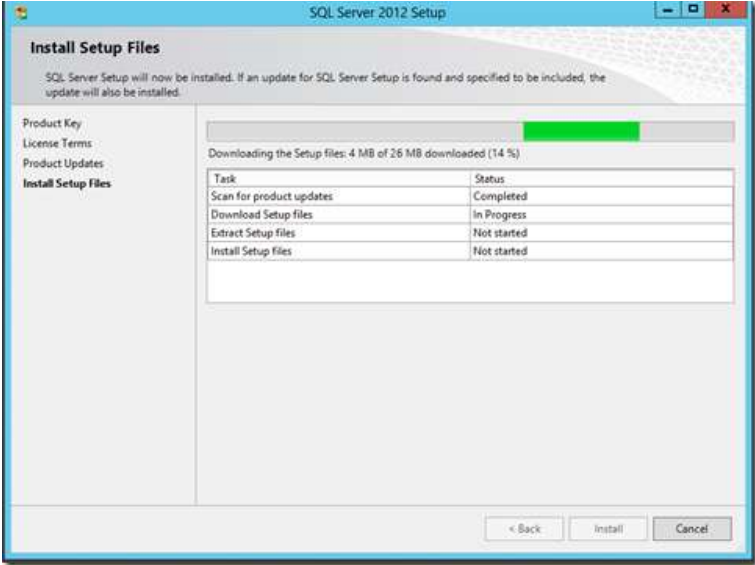
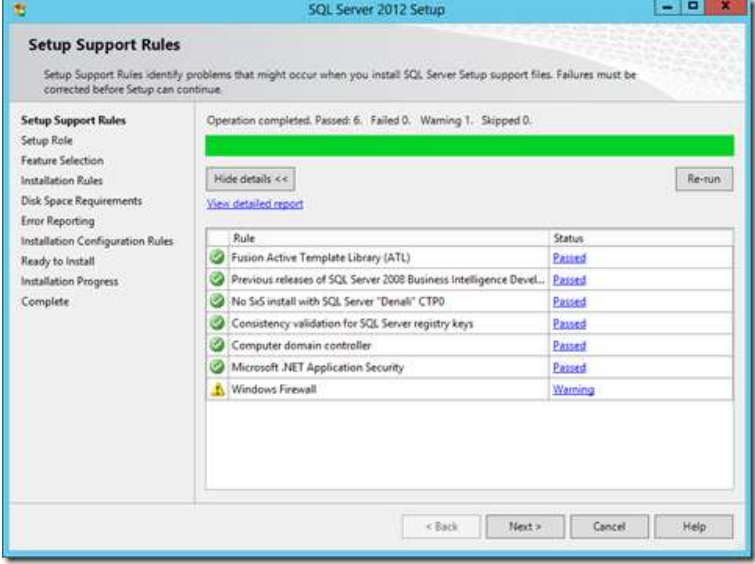
6.11.3. SQL Server 2012 Installation

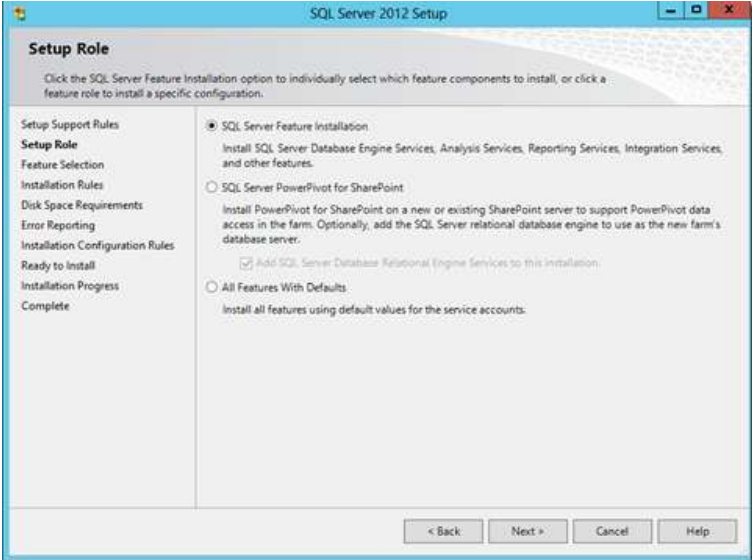
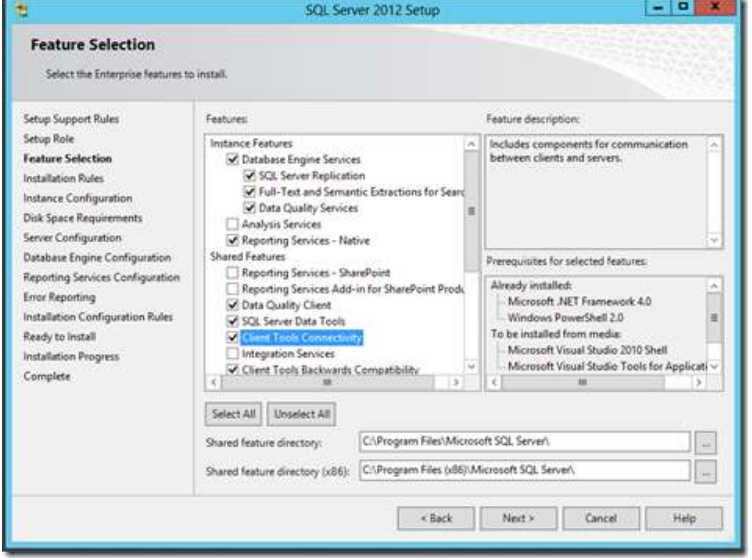
The table below provides the installation steps for SQL Server 2012.

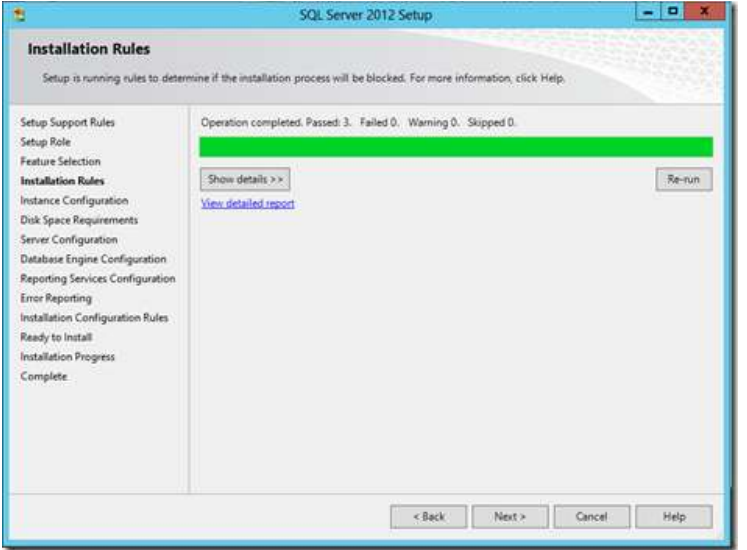
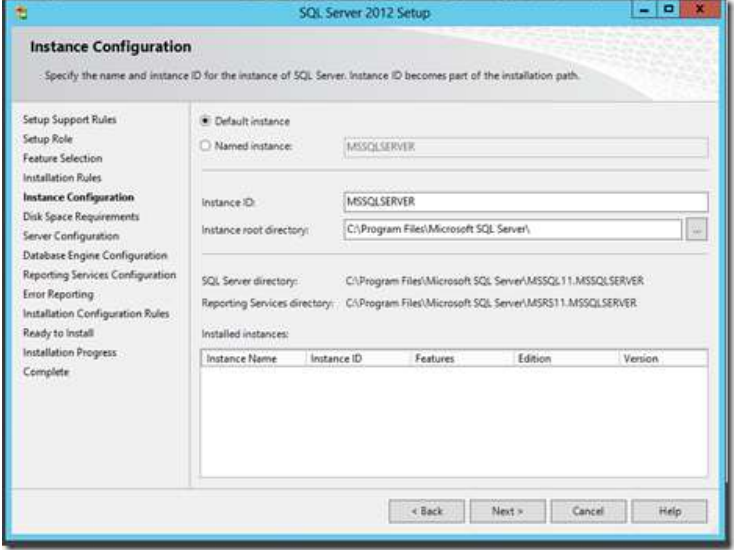
Instructions	Visual
<p>Perform the following steps on both SQL1 and SQL2.</p> <p>Launch the installation setup by double clicking on “setup.exe”.</p>	
<p>In the left side of the screen click Installation.</p>	
<p>In the pane on the right click on “New SQL Server stand—alone installation or add features to an existing installation”.</p>	

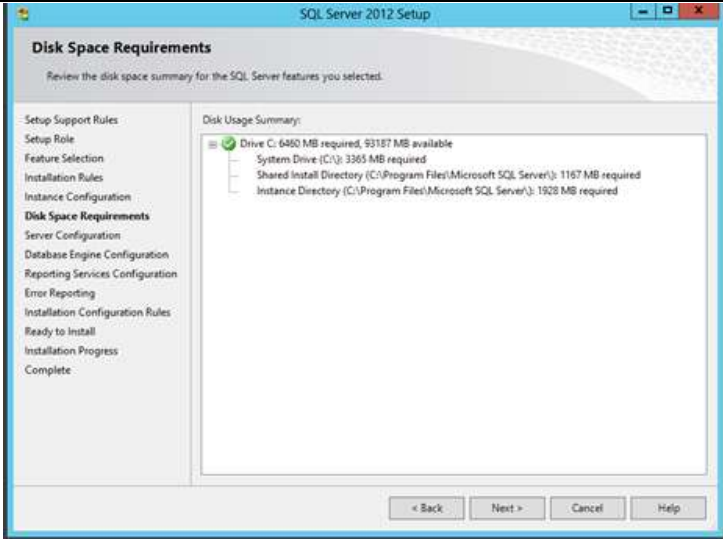
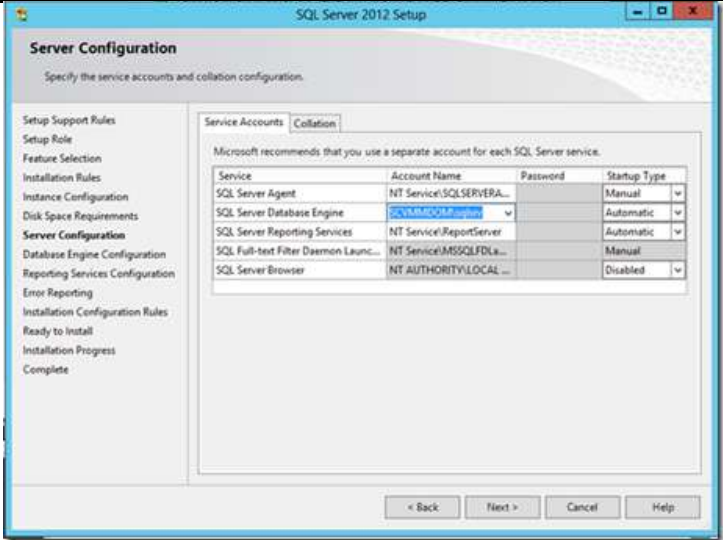
Instructions	Visual
<p>Review the installation advisor report.</p> <p>Click OK.</p>	
<p>Provide a valid Product Key.</p> <p>Click Next.</p>	

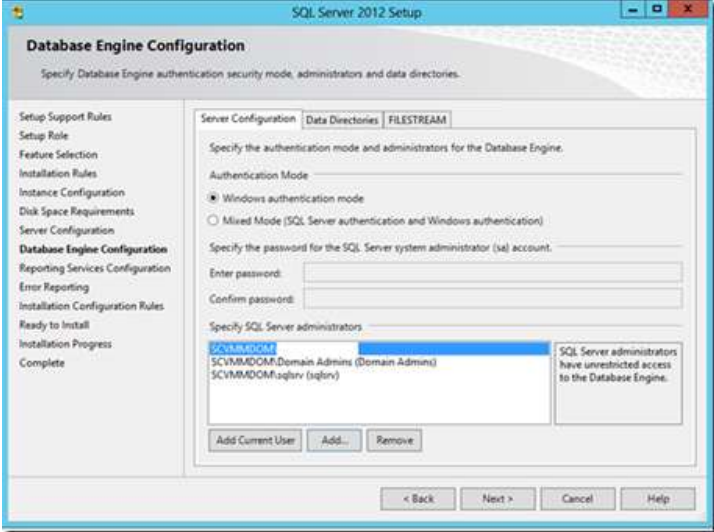
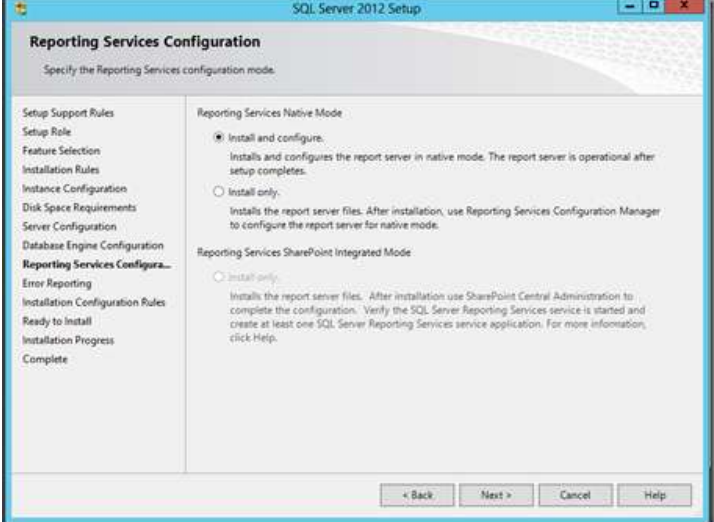
Instructions	Visual
<p>Mark the checkbox “I accept the license terms”</p> <p>Click Next.</p>	
<p>Click Next.</p>	

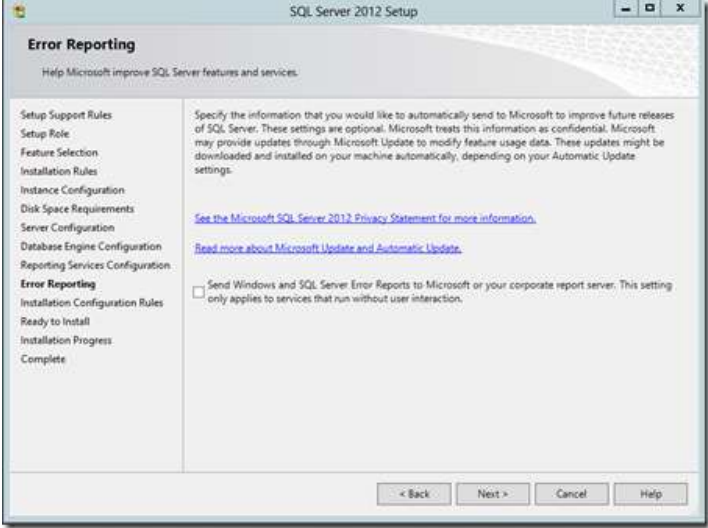
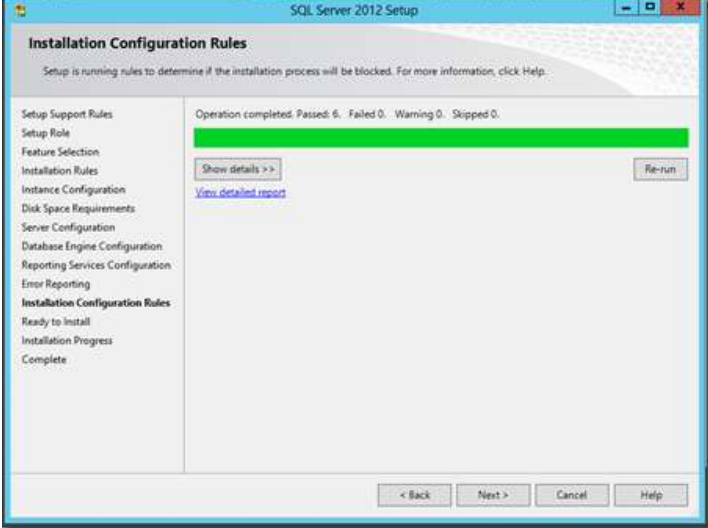
Instructions	Visual
<p>Note: To complete this step successfully, the server need to be connect to the internet.</p>	
<p>Click Next.</p>	

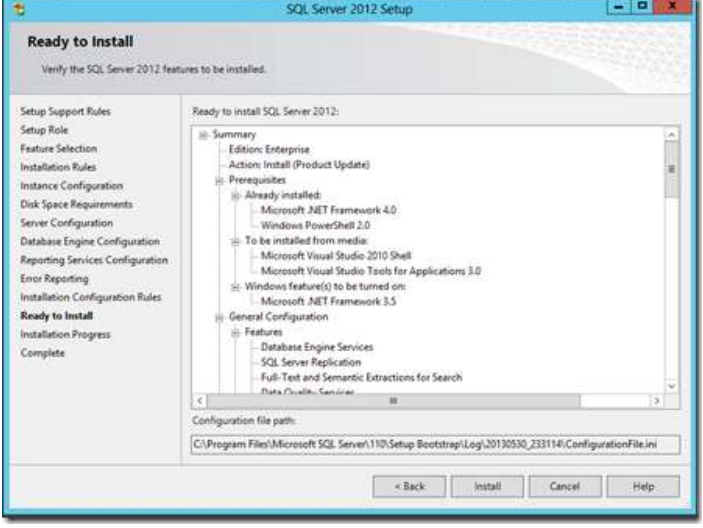
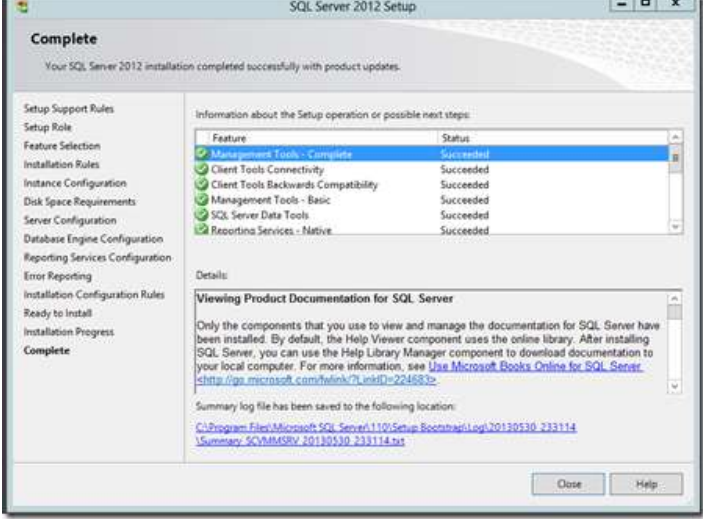
Instructions	Visual
<p>Click Next.</p>	
<p>Mark the relevant SQL roles.</p> <p>For our SCVMM install we installed the following:</p> <ul style="list-style-type: none"> -Database Engine Services (with all subs) -Reporting Services – Native -Data Quality Client -SQL Server Data Tools -Client Tools Backwards Compatibility -Management Tools - Complete <p>Click Next.</p>	

Instructions	Visual
<p>Click Next.</p>	
<p>Click Next.</p>	

Instructions	Visual																								
<p>Click Next.</p>																									
<p>Set “SQLSRV” as a service account for the SQL engine.</p> <p>Click Next.</p>	 <table><thead><tr><th>Service</th><th>Account Name</th><th>Password</th><th>Startup Type</th></tr></thead><tbody><tr><td>SQL Server Agent</td><td>NT Service\SQLSERVERA...</td><td></td><td>Manual</td></tr><tr><td>SQL Server Database Engine</td><td>SQLMOM\sqlservr</td><td></td><td>Automatic</td></tr><tr><td>SQL Server Reporting Services</td><td>NT Service\ReportServer</td><td></td><td>Automatic</td></tr><tr><td>SQL Full-text Filter Daemon Launc...</td><td>NT Service\MSSQLFDLa...</td><td></td><td>Manual</td></tr><tr><td>SQL Server Browser</td><td>NT AUTHORITY\LOCAL...</td><td></td><td>Disabled</td></tr></tbody></table>	Service	Account Name	Password	Startup Type	SQL Server Agent	NT Service\SQLSERVERA...		Manual	SQL Server Database Engine	SQLMOM\sqlservr		Automatic	SQL Server Reporting Services	NT Service\ReportServer		Automatic	SQL Full-text Filter Daemon Launc...	NT Service\MSSQLFDLa...		Manual	SQL Server Browser	NT AUTHORITY\LOCAL...		Disabled
Service	Account Name	Password	Startup Type																						
SQL Server Agent	NT Service\SQLSERVERA...		Manual																						
SQL Server Database Engine	SQLMOM\sqlservr		Automatic																						
SQL Server Reporting Services	NT Service\ReportServer		Automatic																						
SQL Full-text Filter Daemon Launc...	NT Service\MSSQLFDLa...		Manual																						
SQL Server Browser	NT AUTHORITY\LOCAL...		Disabled																						

Instructions	Visual
<p>Add the local Administrators group & the domain Administrator to the SQL Administrator group and “SQLSRV” service account.</p> <p>Click Next.</p>	
<p>Click Next.</p>	

Instructions	Visual
<p>Click Next.</p>	
<p>Click Next.</p>	

Instructions	Visual
<p>Click Next.</p>	
<p>Review the installation summary.</p> <p>Click Close.</p>	

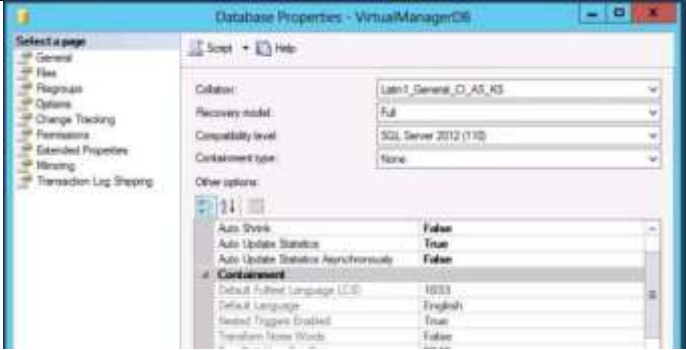
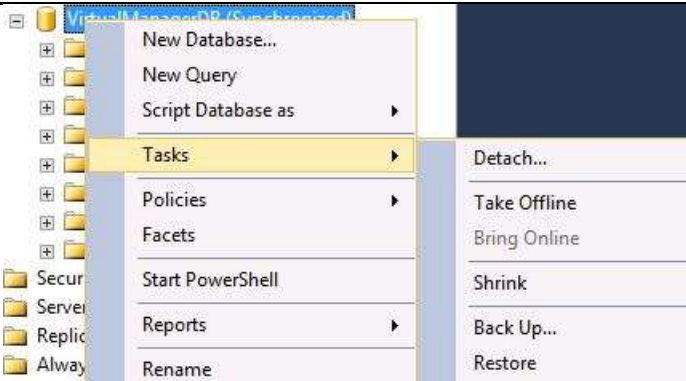
6.11.4. AlwaysOn Application Group

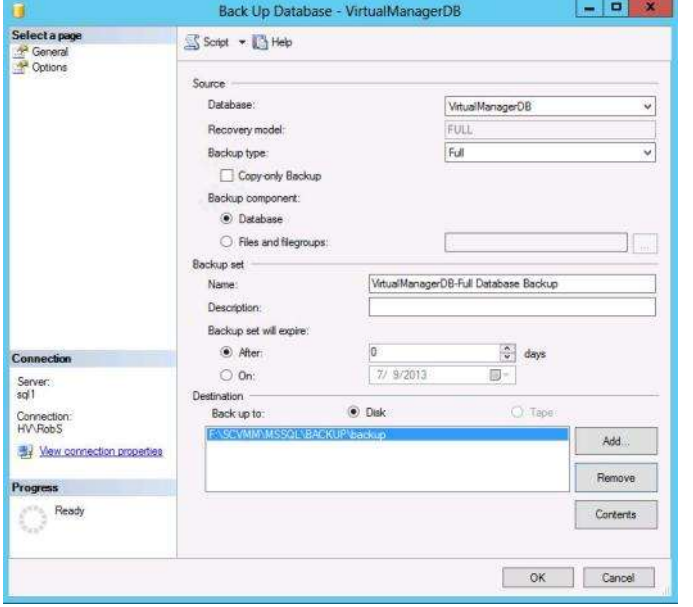
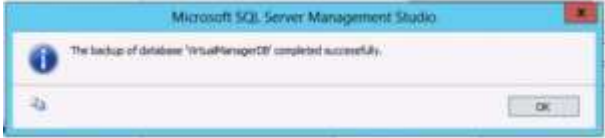
This section provides the instructions for configuring an Always On Availability Group for System Center Virtual Machine Manager database.

6.11.4.1. Installation Pre-requisites

Prior to creating the Always On Availability Group, you will need to do the following:

1. Make sure that the Recovery Model for the VirtualManagerDB database is set to Full
2. Perform a full backup of the VirtualManagerDB database
3. Create a shared network location accessible by both servers (SQL1 and SQL2)

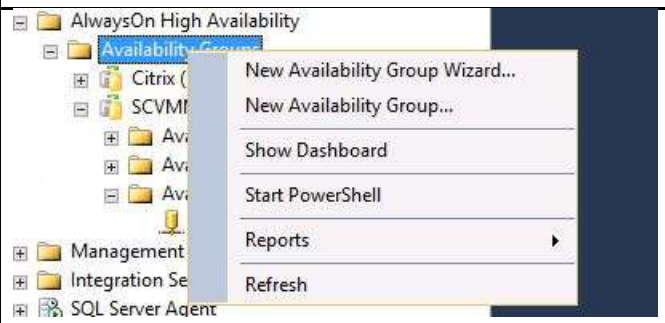
Instructions	Visual
<p>Perform the following steps on SQL1 only.</p> <p>Open the SQL Server Management Studio (with elevated privileges).</p> <p>Expand the databases in the left pane.</p>	
<p>Right-click the database called VirtualManagerDB and select Properties.</p> <p>In the Database Properties window, select Options from the left pane.</p> <p>Set the Recovery Model to FULL.</p> <p>Click OK.</p>	
<p>Right-click the database called VirtualManagerDB and select Tasks and then Back Up...</p>	

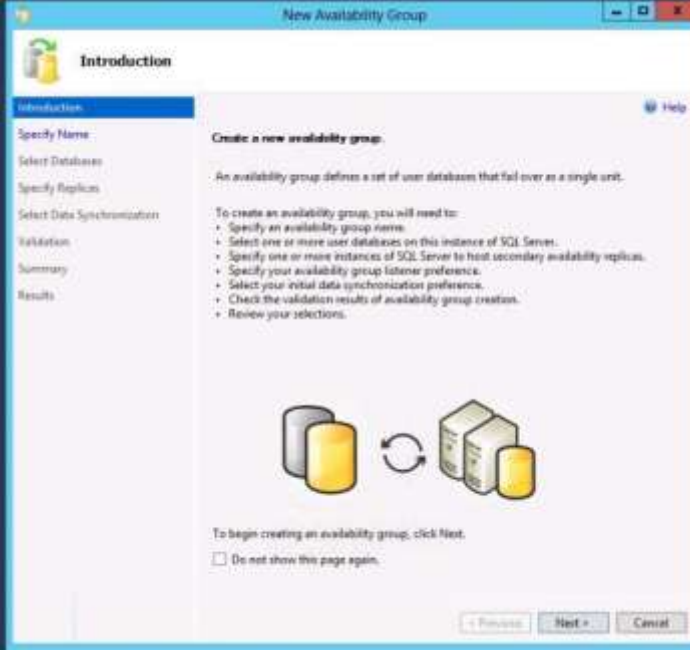
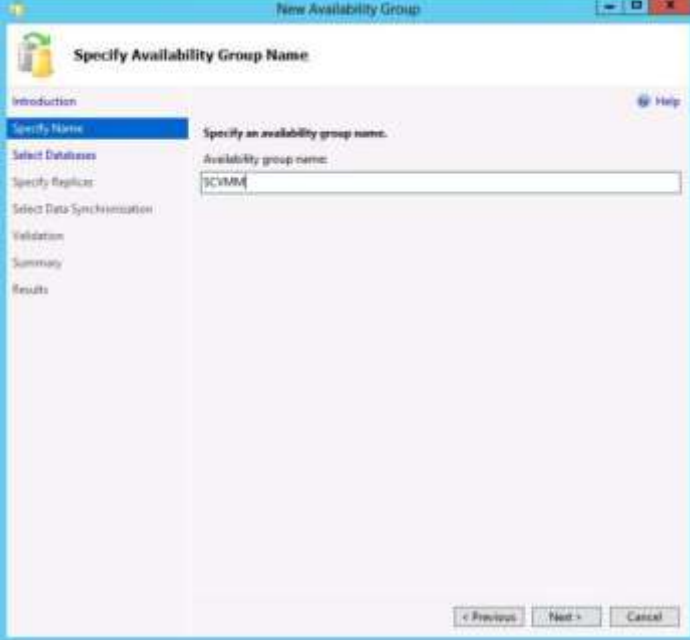
Instructions	Visual
<p>Set the Backup Type to FULL.</p> <p>Select the desired destination of the backup.</p> <p>Click OK.</p>	
<p>Click OK.</p>	
<p>Create a share that is accessible by both SQL1 and SQL2 that will be used in a later step when the Always On Availability Group is created.</p> <p>On INFRA-1 or INFRA-2, open File Explorer.</p> <p>Navigate to c:\clusterstorage\infracsv.</p> <p>Create a new folder called AlwaysOn.</p> <p>This folder will be available through a cluster shared volume.</p> <p>The full path will be <u>\\clusterfs\infracsv\AlwaysOn</u> .</p>	

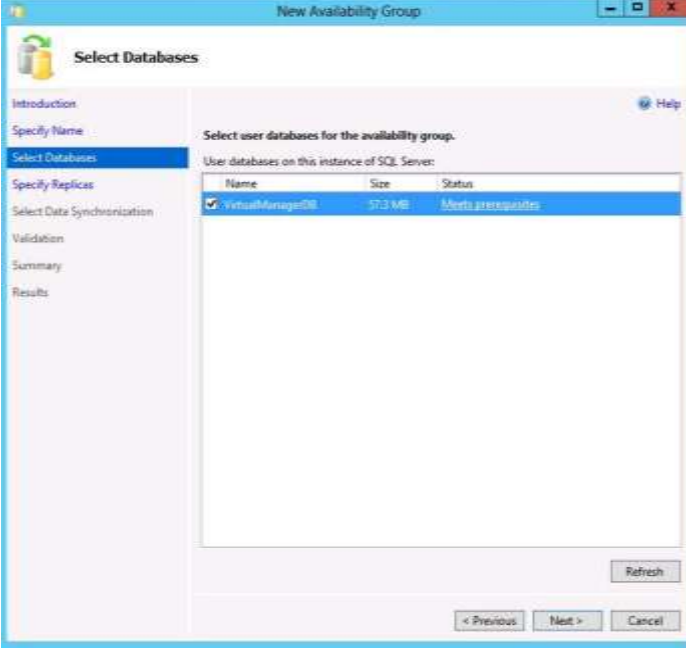
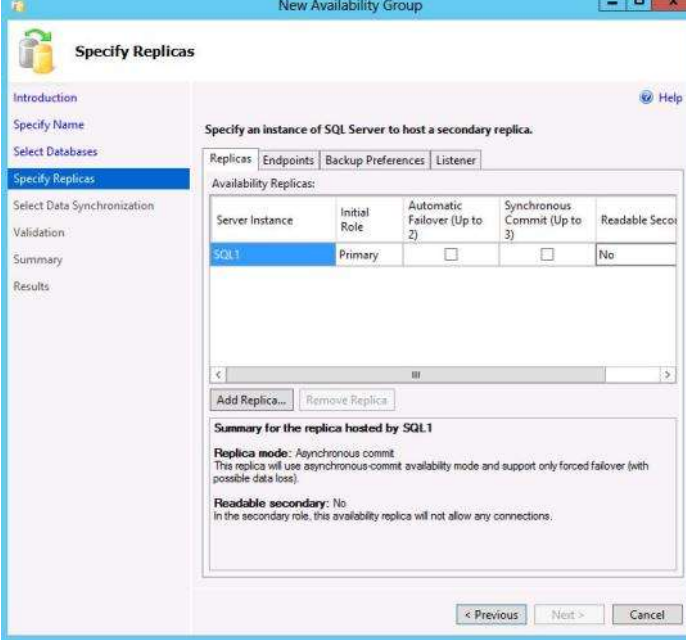
6.11.4.2. Creating the AlwaysOn Availability Group

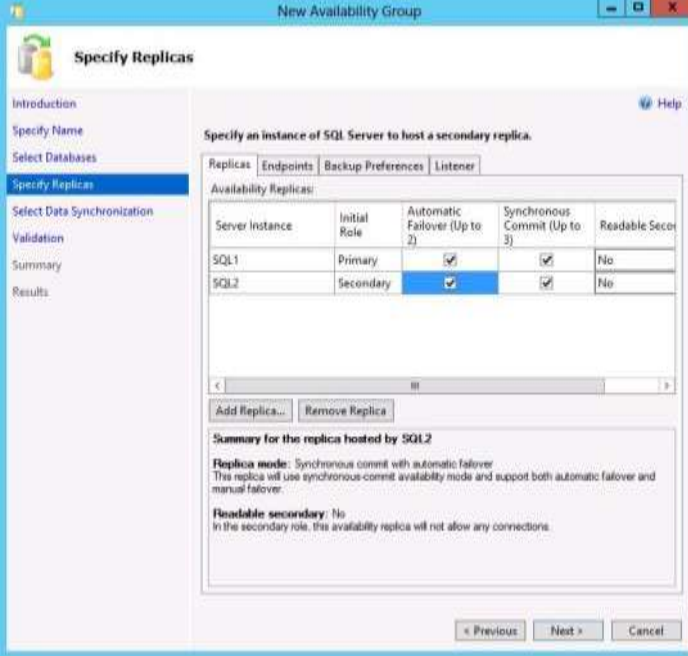
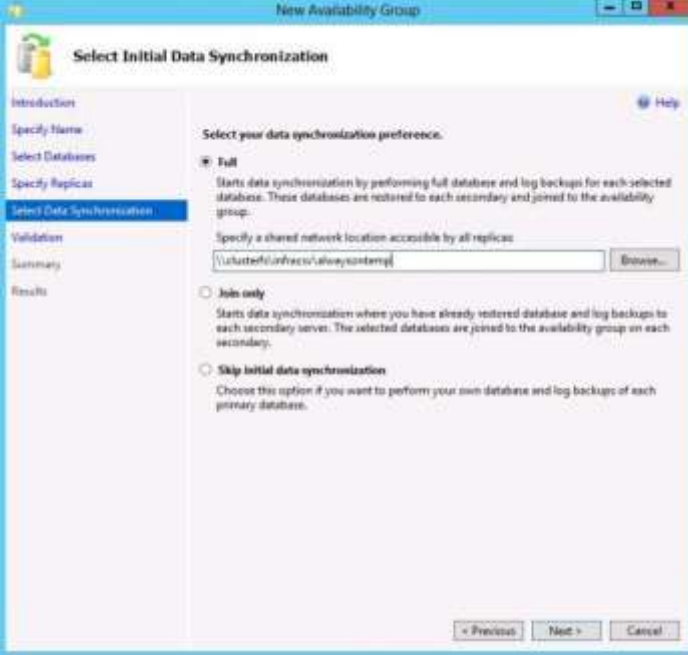
Before creating the Always On Availability Group the following prerequisite steps should be completed as per the instructions in the previous section.

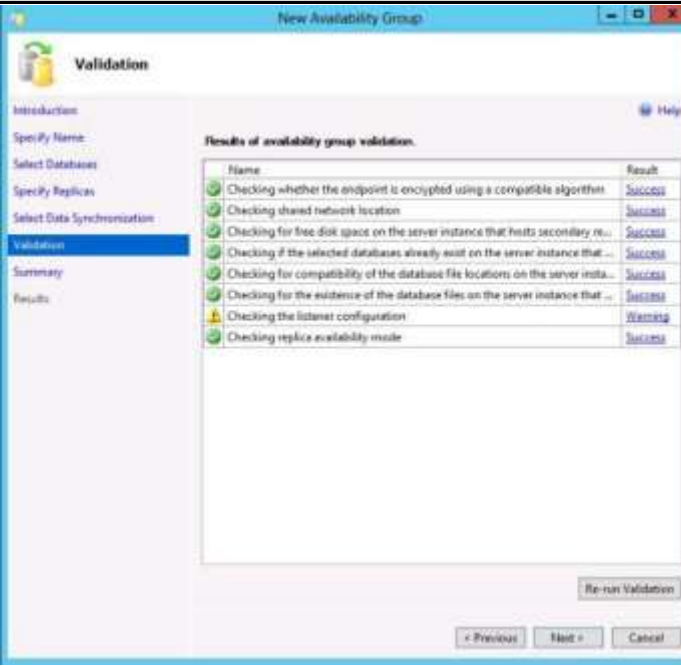
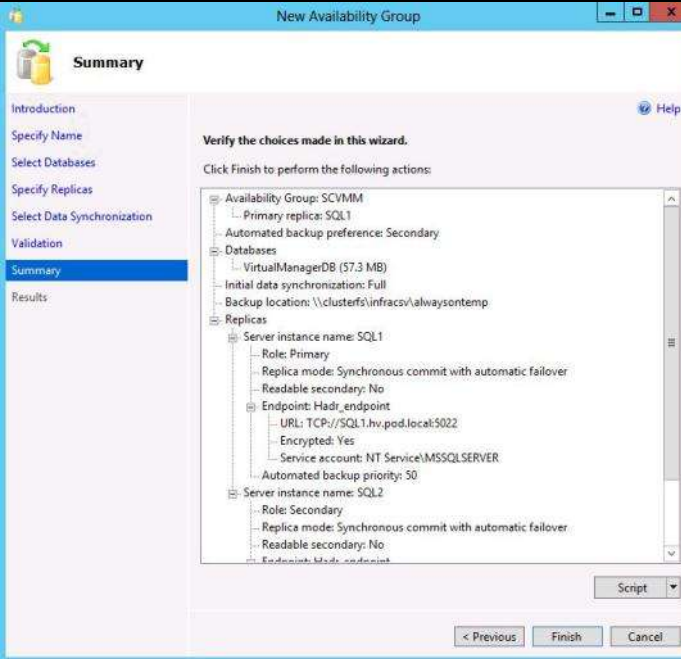
1. Make sure that the Recovery Model for the VirtualManagerDB database is set to Full
2. Perform a full backup of the VirtualManagerDB database
3. Create a shared network location accessible by both servers (SQL1 and SQL2)

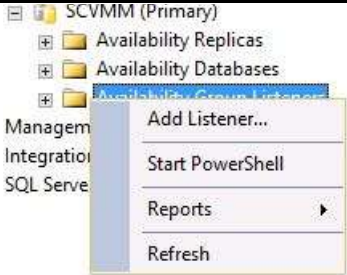
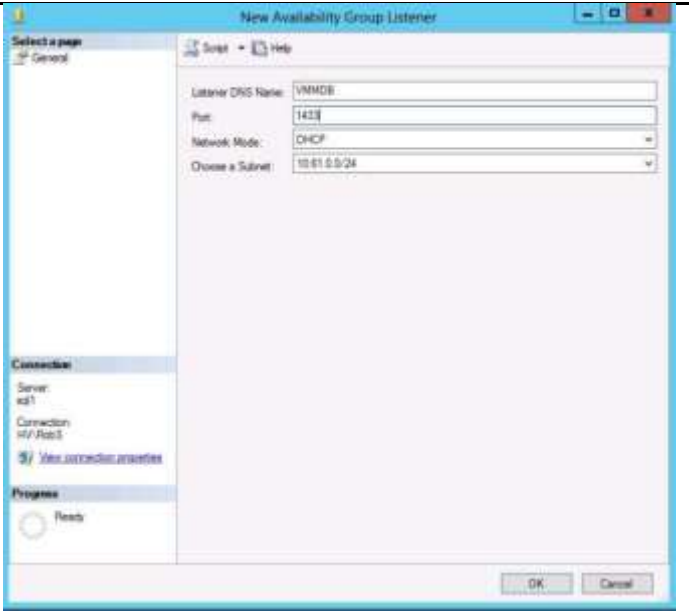
Instructions	Visual
<p>Perform the following steps on SQL1.</p> <p>Open the SQL Server Management Studio (with elevated privileges).</p> <p>Expand AlwaysOn High Availability in the left.</p>	
<p>Right-click Availability Groups and select New Availability Group Wizard...</p>	

Instructions	Visual
<p>Click Next.</p>	 <p>The screenshot shows the 'New Availability Group' wizard in SQL Server Enterprise Manager. The title bar reads 'New Availability Group'. The left pane shows a list of steps: Introduction (selected), Specify Name, Select Databases, Specify Replicas, Select Data Synchronization, Validation, Summary, and Results. The main pane is titled 'Introduction' and contains the following text:</p> <p>Create a new availability group.</p> <p>An availability group defines a set of user databases that fail over as a single unit.</p> <p>To create an availability group, you will need to:</p> <ul style="list-style-type: none"> • Specify an availability group name. • Select one or more user databases on this instance of SQL Server. • Specify one or more instances of SQL Server to host secondary availability replicas. • Specify your availability group listener preference. • Select your initial data synchronization preference. • Check the validation results of availability group creation. • Review your selections. <p>Below the text is a diagram showing two server icons connected by a circular arrow, representing high availability. At the bottom, there is a checkbox labeled 'Do not show this page again.' and three buttons: '< Previous', 'Next >', and 'Cancel'.</p>
<p>Specify the name you want to use for the Availability Group.</p> <p>In this example, we used SCVMM.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Specify Availability Group Name' page of the 'New Availability Group' wizard. The title bar reads 'New Availability Group'. The left pane shows the same list of steps as the previous screenshot, with 'Specify Name' now selected. The main pane is titled 'Specify an availability group name.' and contains a text box labeled 'Availability group name:' with the value 'SCVMM' entered. At the bottom, there are three buttons: '< Previous', 'Next >', and 'Cancel'.</p>

Instructions	Visual
<p>Select the databases that you want to be managed by the availability group.</p> <p>In this example we used VirtualManagerDB.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'New Availability Group' wizard at the 'Select Databases' step. The left pane lists the steps: Introduction, Specify Name, Select Databases (current), Specify Replicas, Select Data Synchronization, Validation, Summary, and Results. The main pane shows a table of user databases on the SQL server. The table has columns for Name, Size, and Status. One database, 'VirtualManagerDB', is selected and highlighted in blue. Its size is 57.3 MB and its status is 'Meets prerequisites'. At the bottom right, there are buttons for '< Previous', 'Next >', and 'Cancel'.</p>
<p>By default, the primary replica will be located on the SQL server where the database was originally created.</p>	 <p>The screenshot shows the 'New Availability Group' wizard at the 'Specify Replicas' step. The left pane lists the steps: Introduction, Specify Name, Select Databases, Specify Replicas (current), Select Data Synchronization, Validation, Summary, and Results. The main pane shows a table of availability replicas. The table has columns for Server Instance, Initial Role, Automatic Failover (Up to 2), Synchronous Commit (Up to 3), and Readable Secondary. One replica, 'SQL1', is listed with an initial role of 'Primary'. Below the table, there are buttons for 'Add Replica...' and 'Remove Replica'. At the bottom right, there are buttons for '< Previous', 'Next >', and 'Cancel'.</p>

Instructions	Visual
<p>Click Add Replica.</p> <p>Enter the name of the server instance where you would like the secondary replica to be created.</p> <p>In this example we used SQL2.</p> <p>Enable Automatic Failover and Synchronous Commit for both the primary and secondary replicas.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Specify Replicas' window in the 'New Availability Group' wizard. The 'Replicas' tab is active, displaying a table of replicas. The table has columns: Server Instance, Initial Role, Automatic Failover (Up to 2), Synchronous Commit (Up to 3), and Readable Secondary. Two replicas are listed: SQL1 with a Primary role and SQL2 with a Secondary role. Both have 'Automatic Failover' and 'Synchronous Commit' checked. Below the table, a summary for the replica hosted by SQL2 is shown, indicating 'Replica mode: Synchronous commit with automatic failover' and 'Readable secondary: No'.</p>
<p>Set your data synchronization preference to FULL</p> <p>Specify a network share accessible by all replicas. In this example we'll use the share that was created in the prerequisites section preceding this section.</p> <p>\\clusterfs\infracs\alwaysonshare</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Select Initial Data Synchronization' window in the 'New Availability Group' wizard. The 'Full' option is selected under 'Select your data synchronization preference.' The description for 'Full' states: 'Starts data synchronization by performing full database and log backups for each selected database. These databases are restored to each secondary and joined to the availability group.' A text box shows the network share path: \\clusterfs\infracs\alwaysonshare. Other options like 'Join only' and 'Skip initial data synchronization' are also visible but not selected.</p>

Instructions	Visual																		
<p>Setup will perform a validation check prior to creating the availability group.</p> <p>In some cases, the listener configuration will not pass the validation check. This is fine as we will create a new availability group listener in a later step.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Validation' step of the 'New Availability Group' wizard. The left sidebar lists the steps: Introduction, Specify Name, Select Databases, Specify Replicas, Select Data Synchronization, Validation (selected), Summary, and Results. The main area displays the 'Results of availability group validation' in a table:</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Result</th> </tr> </thead> <tbody> <tr> <td>Checking whether the endpoint is encrypted using a compatible algorithm</td> <td>Success</td> </tr> <tr> <td>Checking shared network location</td> <td>Success</td> </tr> <tr> <td>Checking for free disk space on the server instance that hosts secondary re...</td> <td>Success</td> </tr> <tr> <td>Checking if the selected databases already exist on the server instance that ...</td> <td>Success</td> </tr> <tr> <td>Checking for compatibility of the database file locations on the server insta...</td> <td>Success</td> </tr> <tr> <td>Checking for the existence of the database files on the server instance that ...</td> <td>Success</td> </tr> <tr> <td>Checking the listener configuration</td> <td>Warning</td> </tr> <tr> <td>Checking replica availability mode</td> <td>Success</td> </tr> </tbody> </table> <p>Buttons at the bottom: < Previous, Next >, Cancel, and a 'Re-run Validation' button.</p>	Name	Result	Checking whether the endpoint is encrypted using a compatible algorithm	Success	Checking shared network location	Success	Checking for free disk space on the server instance that hosts secondary re...	Success	Checking if the selected databases already exist on the server instance that ...	Success	Checking for compatibility of the database file locations on the server insta...	Success	Checking for the existence of the database files on the server instance that ...	Success	Checking the listener configuration	Warning	Checking replica availability mode	Success
Name	Result																		
Checking whether the endpoint is encrypted using a compatible algorithm	Success																		
Checking shared network location	Success																		
Checking for free disk space on the server instance that hosts secondary re...	Success																		
Checking if the selected databases already exist on the server instance that ...	Success																		
Checking for compatibility of the database file locations on the server insta...	Success																		
Checking for the existence of the database files on the server instance that ...	Success																		
Checking the listener configuration	Warning																		
Checking replica availability mode	Success																		
<p>Verify that the choices you made in previous steps are correct.</p> <p>Click Finish.</p>	 <p>The screenshot shows the 'Summary' step of the 'New Availability Group' wizard. The left sidebar lists the steps: Introduction, Specify Name, Select Databases, Specify Replicas, Select Data Synchronization, Validation, Summary (selected), and Results. The main area displays the 'Verify the choices made in this wizard' section with a tree view of the configuration:</p> <ul style="list-style-type: none"> Click Finish to perform the following actions: <ul style="list-style-type: none"> Availability Group: SCVMM <ul style="list-style-type: none"> Primary replica: SQL1 Automated backup preference: Secondary Databases <ul style="list-style-type: none"> VirtualManagerDB (57.3 MB) Initial data synchronization: Full Backup location: \\clusterfs\infracs\alwaysontemp Replicas <ul style="list-style-type: none"> Server instance name: SQL1 <ul style="list-style-type: none"> Role: Primary Replica mode: Synchronous commit with automatic failover Readable secondary: No Endpoint: Hadr_endpoint <ul style="list-style-type: none"> URL: TCP://SQL1.hv.pod.local:5022 Encrypted: Yes Service account: NT Service\MSSQLSERVER Automated backup priority: 50 Server instance name: SQL2 <ul style="list-style-type: none"> Role: Secondary Replica mode: Synchronous commit with automatic failover Readable secondary: No Endpoint: Hadr_endpoint <p>Buttons at the bottom: < Previous, Finish, Cancel, and a 'Script' button.</p>																		

Instructions	Visual
<p>In this step we configure the Availability Group Listener.</p> <p>Expand Availability Groups.</p> <p>Expand SCVMM.</p> <p>Right click on Availability Group Listener and select Add Listener...</p>	
<p>Enter VMMDb as the Listener DNS Name</p> <p>Enter 1433 as the Port</p> <p>Enter DHCP as Network Mode</p> <p>Enter 10.61.0.0/24 as the Subnet</p> <p>Click OK.</p>	

6.11.5. Log File Management

The log files for XenDesktop and System Center Virtual Machine Manager are prone to large amounts of activity. Since a Full recovery model is necessary for AlwaysOn groups, the following actions are recommended:

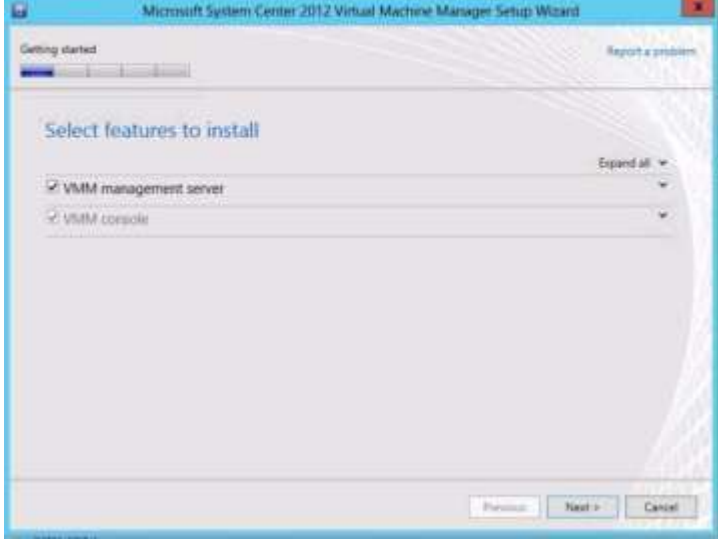
1. Verify that the Log file drives have sufficient space to hold the log files. For this environment, 50GB is the minimum recommended size.
2. Configure the maximum size of the log file for each database in SQL Management Studio so that the SQL LUN is not unexpectedly overrun. The PVS log file maximum is recommended to be 10GB, with 20GB used for the maximum size of VMM and XD.

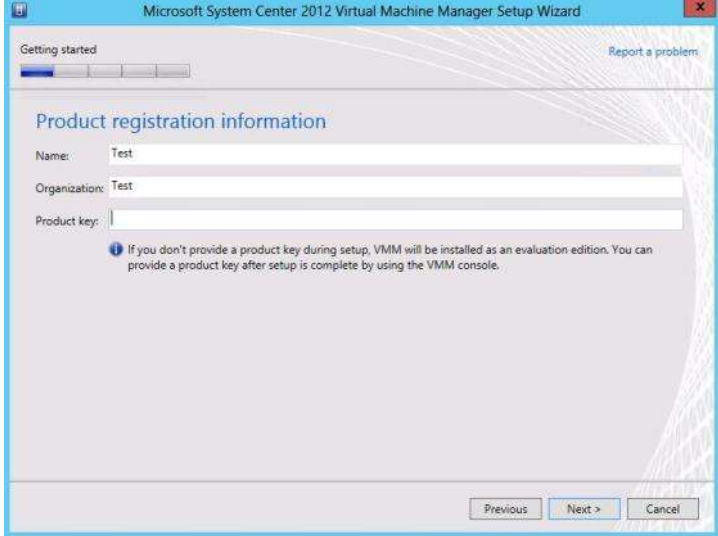
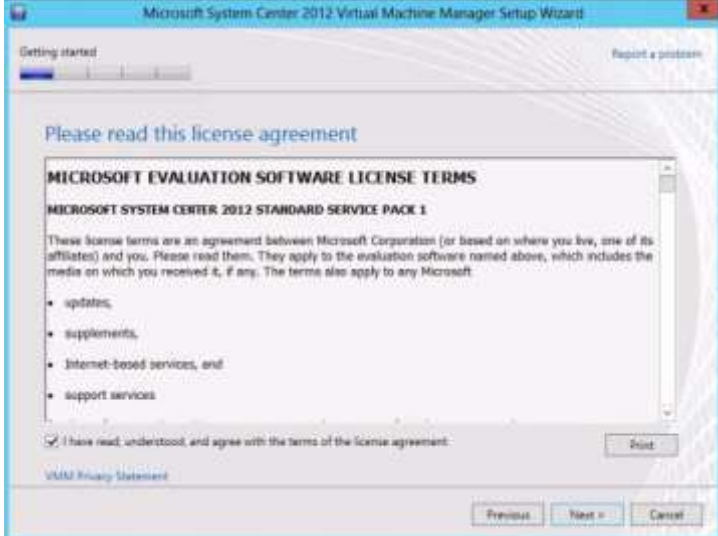
3. Configure SQL Agent to automatically backup and truncate the log files for the XenDesktop, Provisioning Services, and Virtual Machine Manager databases on a frequent basis, or set the backup job to occur when the log file size reaches 90% of capacity.


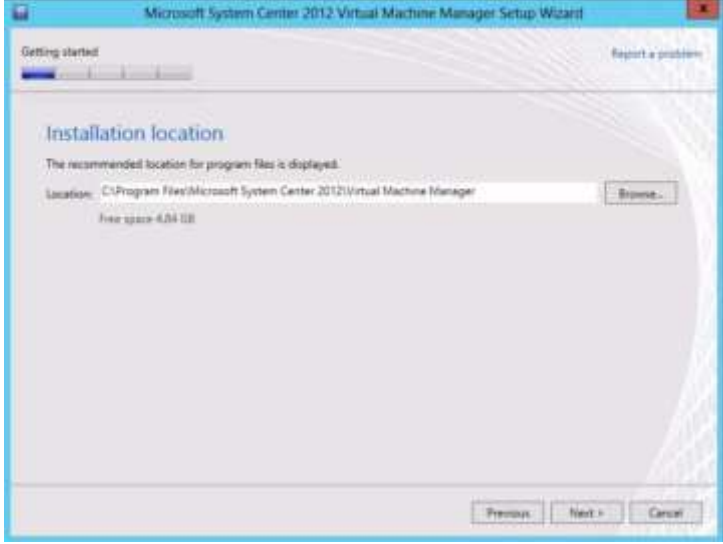
6.12. Installing and Configuring System Center 2012 Virtual Machine Manager

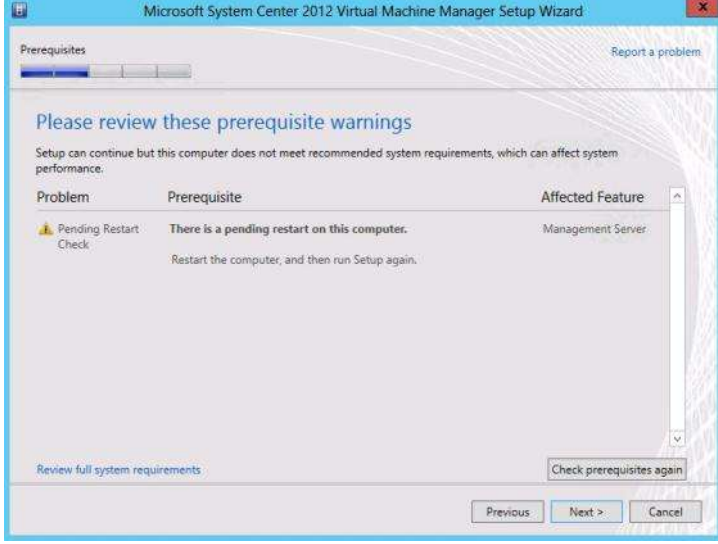
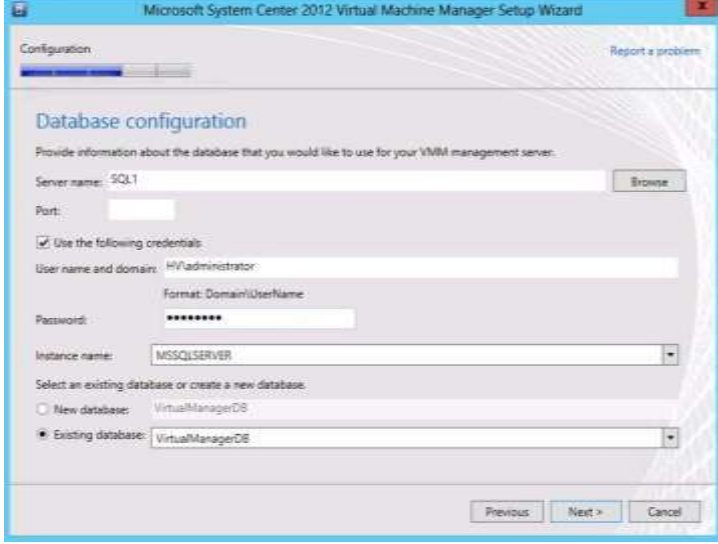
This section provides the instructions for installing and configuring System Center 2012 Virtual Machine Manager. Before starting the installation the following steps should be completed.

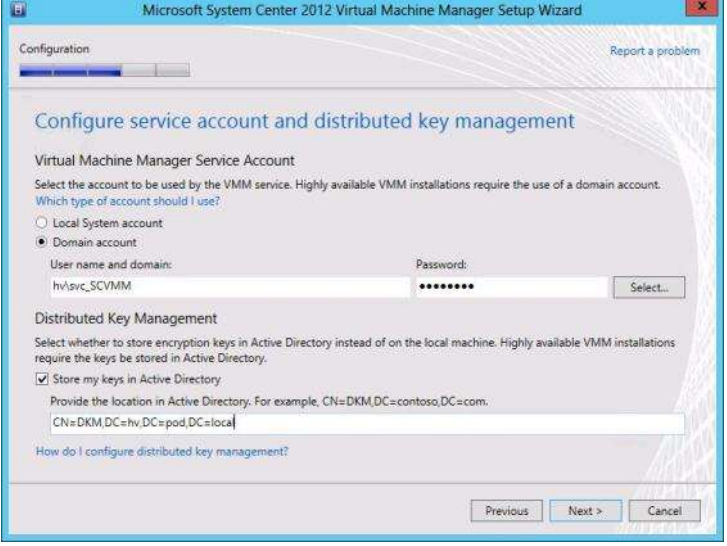
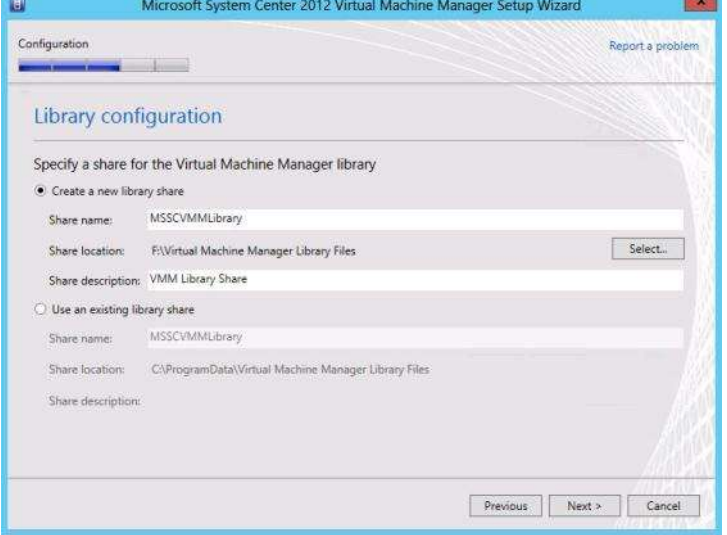
1. Install SQL server
2. Create a SQL instance called MSSQLSERVER
3. Create a SQL database called VirtualManagerDB

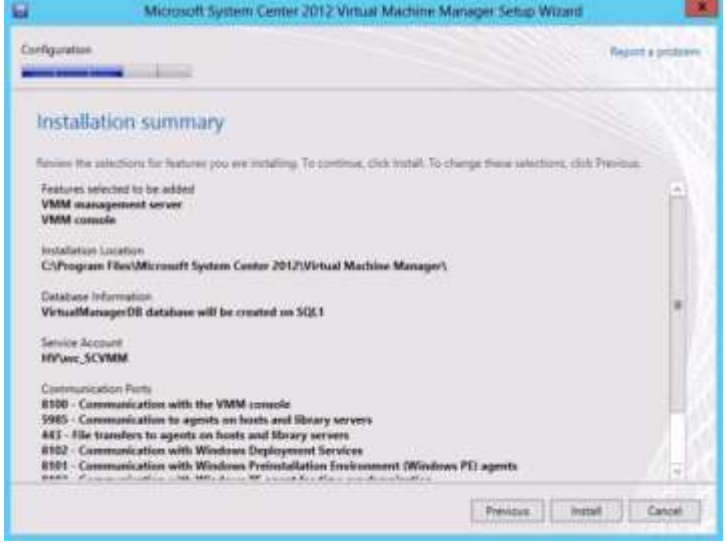
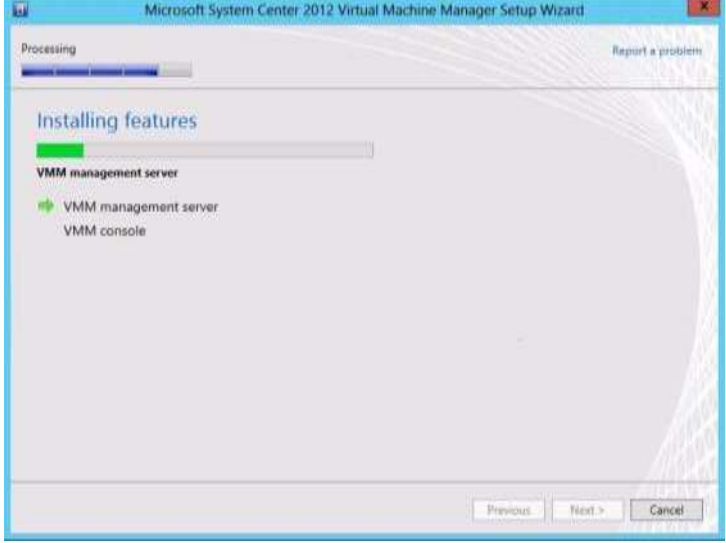
Instructions	Visual
<p>Perform the following steps on VMM1.</p> <p>Browse to the VMM 2012 SP1 installation files and run setup.exe.</p>	
<p>Select both the VMM Management Server and VMM Console.</p> <p>Click Next.</p>	

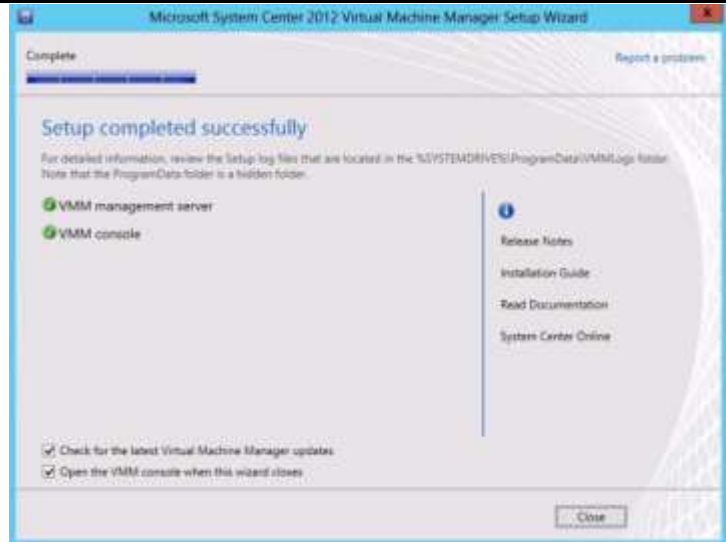
Instructions	Visual
<p>Enter the product registration information.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Product registration information' step of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The window has a blue title bar and a progress bar at the top. The main area contains three text input fields labeled 'Name:', 'Organization:', and 'Product key:'. The 'Name' and 'Organization' fields contain the text 'Test'. Below these fields is an information icon and a note: 'If you don't provide a product key during setup, VMM will be installed as an evaluation edition. You can provide a product key after setup is complete by using the VMM console.' At the bottom right, there are three buttons: 'Previous', 'Next >', and 'Cancel'.</p>
<p>Read the license agreement and if you agree with the terms of the agreement select the appropriate check box.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Please read this license agreement' step of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The window has a blue title bar and a progress bar at the top. The main area displays the 'MICROSOFT EVALUATION SOFTWARE LICENSE TERMS' and 'MICROSOFT SYSTEM CENTER 2012 STANDARD SERVICE PACK 1'. It includes a paragraph of text about the license terms and a bulleted list of services: updates, supplements, Internet-based services, and support services. At the bottom, there is a checkbox labeled 'I have read, understood, and agree with the terms of the license agreement:' which is checked. To the right of the checkbox is a 'Print' button. At the bottom right, there are three buttons: 'Previous', 'Next >', and 'Cancel'.</p>

Instructions	Visual
<p>Select whether you would like to join the Customer Experience Improvement Program.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Customer Experience Improvement Program (CEIP)' screen of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The window title is 'Microsoft System Center 2012 Virtual Machine Manager Setup Wizard'. At the top, there is a progress bar labeled 'Getting started' with four steps, the first of which is highlighted. A 'Report a problem' link is in the top right corner. The main heading is 'Customer Experience Improvement Program (CEIP)'. Below it, the text says 'If you choose to participate:'. Under 'Microsoft will:', there are two bullet points: 'Collect information about your software and hardware configurations.' and 'Collect information about how you use our software and services to identify trends and usage patterns.' Under 'Microsoft will not:', there are three bullet points: 'Collect your name or address.', 'Ask you to take surveys; nor will you be contacted by a sales representative.', and 'Prompt you with additional messages that might interrupt your work.' There are two radio buttons: 'Yes, I am willing to participate in the Customer Experience Improvement Program' (which is selected) and 'No, I am not willing to participate'. Below the radio buttons, it says 'You can stop participating at any time by changing a setting in Customer Experience Improvement Program Settings, found in Settings workspace of the VMM console.' There are three links: 'More about the Customer Experience Improvement Program', 'Privacy Statement for the Microsoft Customer Experience Improvement Program', and 'VMM Privacy Statement'. At the bottom right, there are three buttons: 'Previous', 'Next >', and 'Cancel'.</p>
<p>Select the installation location. The default is fine for most installations.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Installation location' screen of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The window title is 'Microsoft System Center 2012 Virtual Machine Manager Setup Wizard'. At the top, there is a progress bar labeled 'Getting started' with four steps, the second of which is highlighted. A 'Report a problem' link is in the top right corner. The main heading is 'Installation location'. Below it, the text says 'The recommended location for program files is displayed.' There is a text box labeled 'Location:' containing the path 'C:\Program Files\Microsoft System Center 2012\Virtual Machine Manager'. To the right of the text box is a 'Browse...' button. Below the text box, it says 'Free space: 4,04 GB'. At the bottom right, there are three buttons: 'Previous', 'Next >', and 'Cancel'.</p>

Instructions	Visual
<p>Click Next.</p>	
<p>Enter the name of the SQL server that contains the SQL instance that you configured in a previous step. In this example the server name is SQL1.</p> <p>Enter the port used to connect to the SQL instance.</p> <p>Enter the credentials used to connect to the SQL instance.</p> <p>Enter the name of the SQL instance. In this example we used MSSQLSERVER.</p> <p>Enter the name of the SQL database that was created in a previous step. In this example we used VirtualManagerDB.</p> <p>Click Next.</p>	

Instructions	Visual
<p>Select the account to be used by the VMM service. In this example we used svc_SCVMM.</p> <p>Select whether to store encryption keys in Active Directory instead of on the local machine and provide the desired path in AD.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Configure service account and distributed key management' step of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The 'Virtual Machine Manager Service Account' section has 'Domain account' selected. The 'User name and domain' field contains 'hvlsvc_SCVMM' and the 'Password' field is masked with dots. The 'Distributed Key Management' section has 'Store my keys in Active Directory' checked, and the 'Provide the location in Active Directory' field contains 'CN=DKM,DC=hv,DC=pod,DC=local'. Navigation buttons 'Previous', 'Next >', and 'Cancel' are at the bottom.</p>
<p>Choose to create a new library share, provide the share location and a description of the share. In this example we used the following:</p> <p>Share name: MSSCVMMLibrary</p> <p>Share location: "F:\Virtual Machine Manager Library Files"</p> <p>Share description: VMM Library Share.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Library configuration' step of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The 'Specify a share for the Virtual Machine Manager library' section has 'Create a new library share' selected. The 'Share name' field contains 'MSSCVMMLibrary', the 'Share location' field contains 'F:\Virtual Machine Manager Library Files', and the 'Share description' field contains 'VMM Library Share'. There is also an option to 'Use an existing library share' with its own fields. Navigation buttons 'Previous', 'Next >', and 'Cancel' are at the bottom.</p>

Instructions	Visual
<p>Review the installation summary to ensure that all of the desired configuration information is correct.</p> <p>Click Install.</p>	 <p>The screenshot shows the 'Installation summary' screen of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The window title is 'Microsoft System Center 2012 Virtual Machine Manager Setup Wizard'. The progress bar indicates 'Configuration' is complete. The 'Installation summary' section lists the following details:</p> <ul style="list-style-type: none"> Features selected to be added: VMM management server, VMM console Installation location: C:\Program Files\Microsoft System Center 2012\Virtual Machine Manager\ Database information: VirtualManagerDB database will be created on SQL I Service account: HV\svc_SCVMM Communication Ports: <ul style="list-style-type: none"> 8100 - Communication with the VMM console 5985 - Communication to agents on hosts and library servers 443 - File transfers to agents on hosts and library servers 8102 - Communication with Windows Deployment Services 8101 - Communication with Windows Preinstallation Environment (Windows PE) agents <p>At the bottom, there are 'Previous', 'Install', and 'Cancel' buttons.</p>
<p>The installation begins...</p>	 <p>The screenshot shows the 'Installing features' screen of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The window title is 'Microsoft System Center 2012 Virtual Machine Manager Setup Wizard'. The progress bar indicates 'Processing' is in progress. The 'Installing features' section shows a progress bar for 'VMM management server' and a list of features being installed:</p> <ul style="list-style-type: none"> VMM management server VMM console <p>At the bottom, there are 'Previous', 'Next >', and 'Cancel' buttons.</p>

Instructions	Visual
<p>If the setup completed successfully you will see the following screen.</p> <p>Click Close.</p>	

6.13. Installing and Configuring Cisco Nexus 1000V Virtual Switch Manager and Virtual Ethernet Modules

This section provides the instructions for installing and configuring the Cisco Nexus 1000V for Hyper-V virtual switch. Before starting the installation the following steps should be completed:

1. System Center Virtual Machine Manager installation
2. Download the latest Nexus 1000V installation [software](#) from Cisco
3. Copy the Nexus 100V for HyperV installation software package to the SCVMM server and extract it to a local directory.


The following table provides a list of parameters used during the installation process of the Nexus 1000V

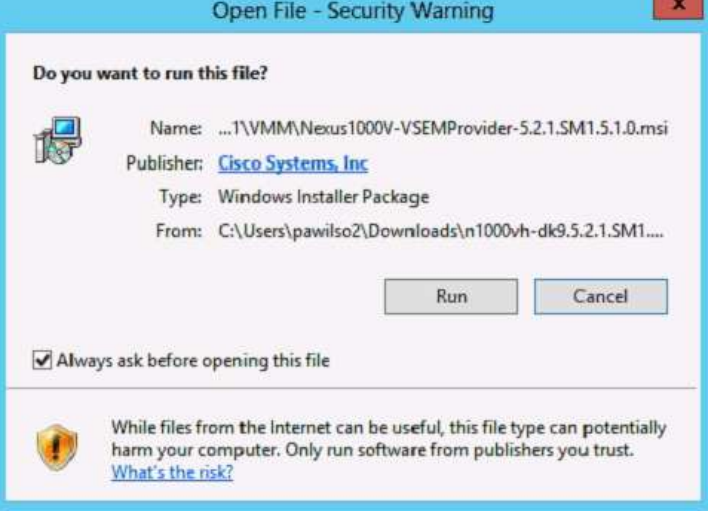

Table 12: Nexus 1000V Parameter Table

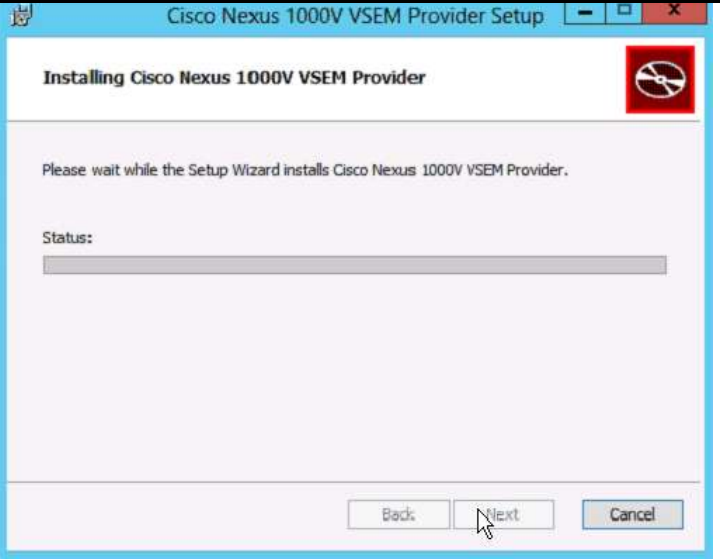
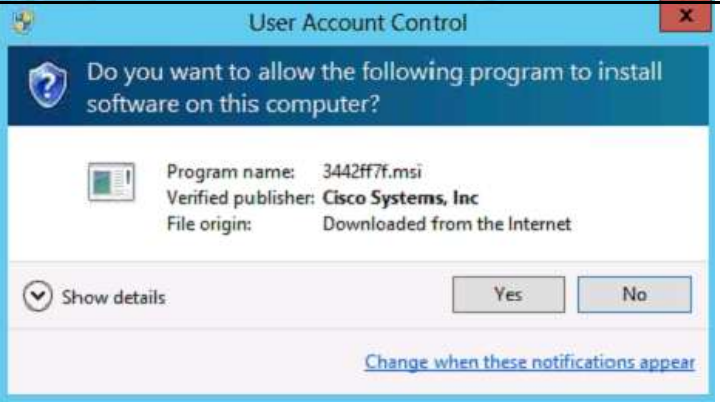
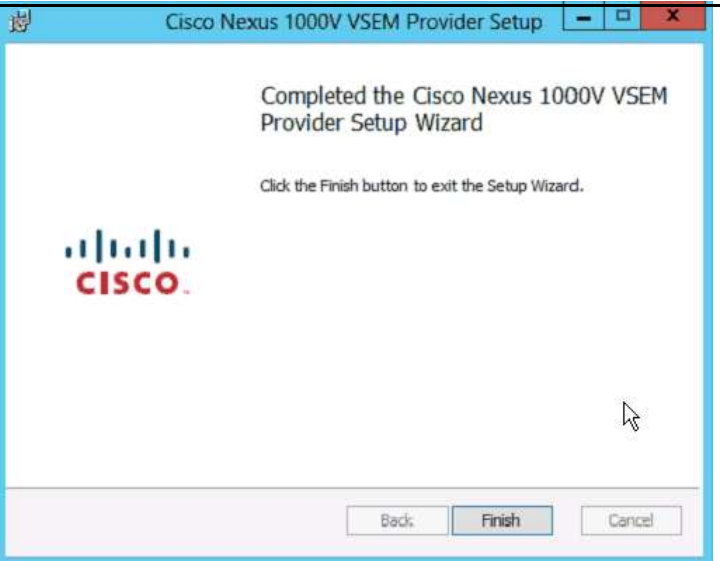
Parameter	500-user Value (Screen Shots)	1000-user Value (Not shown)
SCVMM Server Name	SCVMM1.hv.pod.local	SCVMM1.hv.pod.local
Virtual Machine Name	Nexus 1000V 1 / Nexus 1000V 2	Nexus 1000V 3 / Nexus 1000V 4
Admin Password	P@ssw0rd	P@ssw0rd
Domain ID	100	200
SwitchName	nexus1000v	nexus1000v3
Mgmt0 IP Address	10.61.0.10	10.61.0.13
Mgmt0 Subnet Mask	255.255.255.0	255.255.255.0
Mgmt0 Gateway IP	10.61.0.1	10.61.0.1
Network_Name	VDINetwork, ClusterNetwork	VDINetwork2, ClusterNetwork2
VDI_Pool_Name	VDI-Pool-1	VDI-Pool-2
Cluster_Pool_Name	Cluster-Pool-1	Cluster-Pool-2
Network_Name	VDINetwork, ClusterNetwork	VDINetwork2, ClusterNetwork2





VLAN_Pool_Name	VLAN-61-Pool, VLAN-62-Pool, VLAN-63-Pool, VLAN-64-Pool	VLAN-61-Pool-2, VLAN-62-Pool-2, VLAN-63-Pool-2, VLAN-64-Pool-2
Start_IP	10.61.0.11, 10.62.0.100, 10.63.0.100, 10.64.0.100	10.61.0.13, 10.62.0.151, 10.63.0.151, 10.64.0.151
End_IP	10.61.0.11, 10.62.0.150, 10.63.0.150, 10.64.0.150	10.61.0.13, 10.62.0.200, 10.63.0.200, 10.64.0.200
Network_ID	10.61.0.0, 10.62.0.0, 10.63.0.0, 10.64.0.0	10.61.0.0, 10.62.0.0, 10.63.0.0, 10.64.0.0
Network_Mask	255.255.255.0, 255.255.240.0, 255.255.255.0, 255.255.255.0	255.255.255.0, 255.255.240.0, 255.255.255.0, 255.255.255.0
Router_Address	10.61.0.1, 10.62.0.1, 10.63.0.1, 10.64.0.1	10.61.0.1, 10.62.0.1, 10.63.0.1, 10.64.0.1
VLAN_ID	VLAN-61, VLAN-62, VLAN-63, VLAN-64	VLAN-61-2, VLAN-62-2, VLAN-63-2, VLAN-64-2
VLAN_##	61, 62, 63, 64	61, 62, 63, 64
vEthernet_Port_profile	VDI-Port-profile	VDI-Port-profile2
Uplink_port_profile	Uplink-profile	Uplink-profile2
Uplink_profile_name	nexus1000v-uplink	nexus1000v-uplink2
AdminRunAsAccount	Nexus1000V Admin	Nexus1000V Admin
VDIPortClassName	VDI Port Class	VDI Port Class

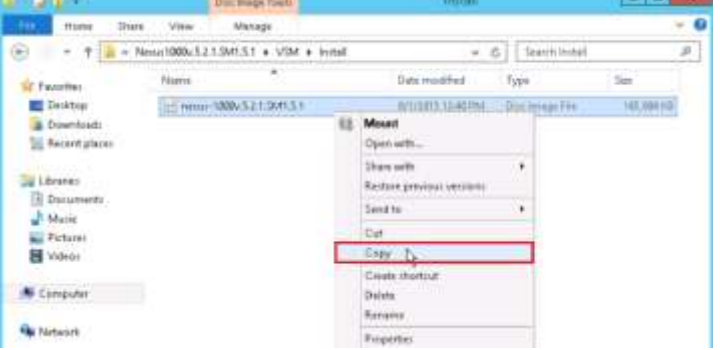
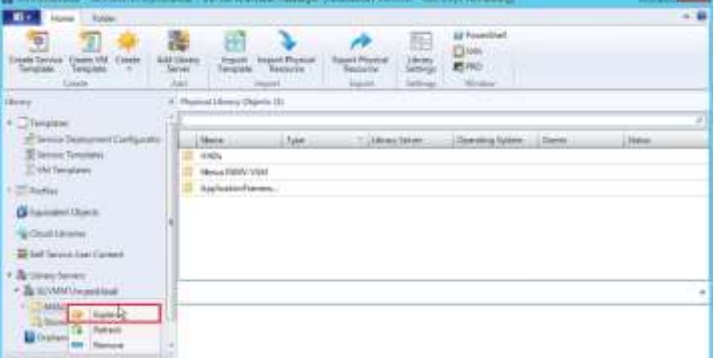
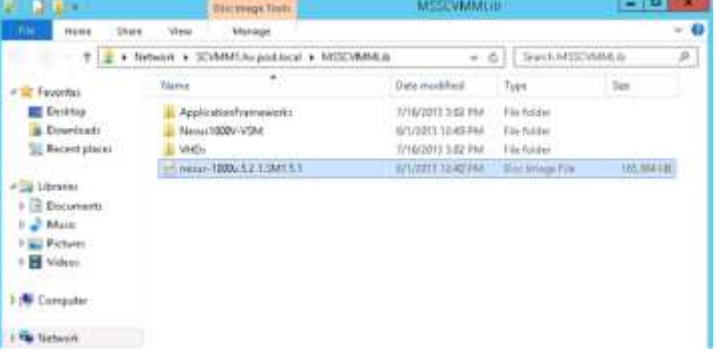
6.13.1. Installing Cisco Nexus 1000V Virtual Switch

Instructions	Visual
Login to the System Center Virtual Machine Manager server as an Administrator.	
From the SCVMM server, run the \VMMNexus-1000-VSEMPProvider-5.2.1.SM1.5.1.0.MSI package	

Instructions	Visual
<p>Click Run if prompted.</p>	
<p>Review the Cisco 1000V VSEM Provider License Agreement and if acceptable, enable the checkbox next to "I accept the terms in the License Agreement"</p> <p>Click Install to begin the installation.</p>	

Instructions	Visual
<p>The install begins immediately.</p>	
<p>If prompted by User Account Control, click Yes.</p>	
<p>When complete, the wizard will show the completed install.</p> <p>Click Finish.</p>	

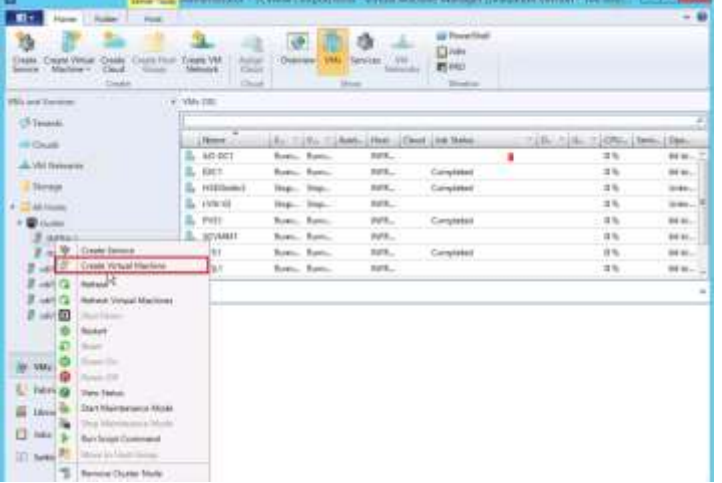
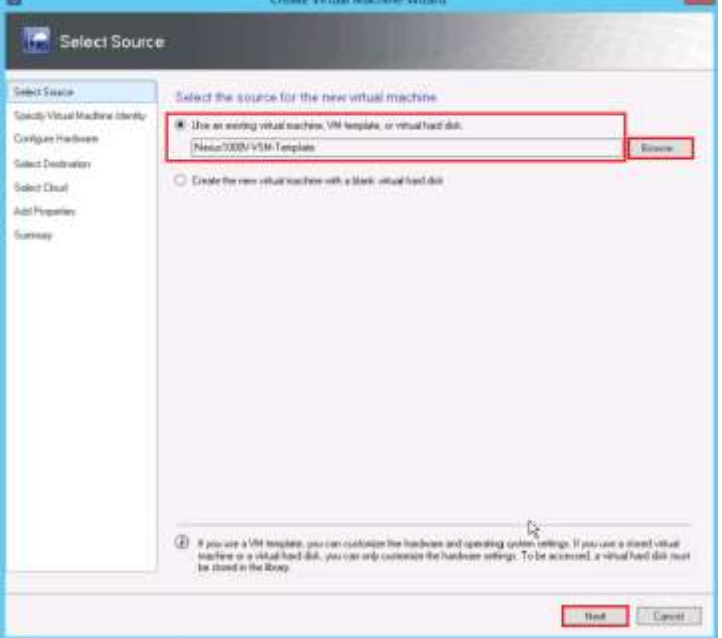
Instructions	Visual
<p>When complete, copy the \\VEM\\Nexus1000V-VEM-5.2.1.SM1.5.1.0.MSI MSI package to the c:\ProgramData\SwitchExtensionDrivers folder</p>	
<p>Launch System Center Virtual Machine Manager console as an administrator.</p>	
<p>Start Windows PowerShell as an Administrator.</p> <p>Run the \\VMM\\VSM_Template\Register-Nexus1000VSMTemplate.ps1 script.</p>	
<p>Navigate to the Library >> Templates section in the Virtual Machine Manager Console and verify the Nexus 1000V-VSM template was imported successfully and the status is OK.</p>	

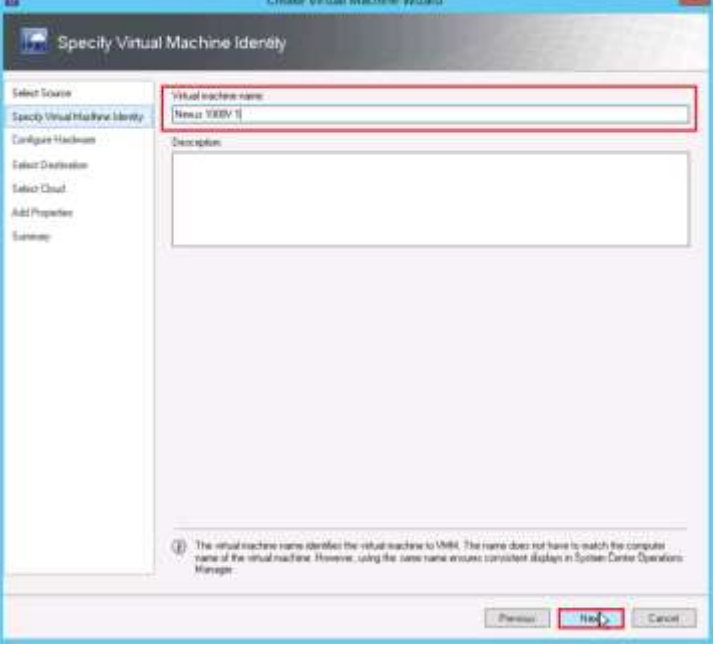
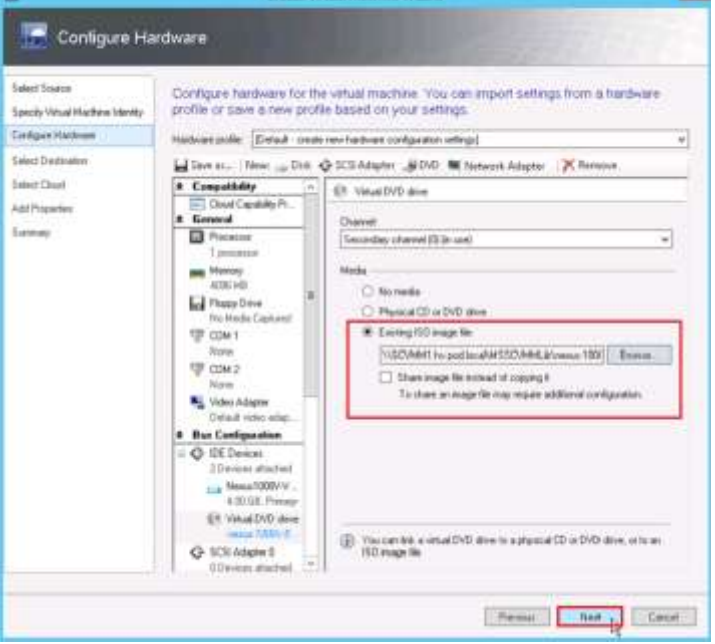
Instructions	Visual
Copy the ISO from the installation media at \\VSM\\Instal\\nexus-1000v.5.2.1.SM1.5.1.iso	
From the SCVMM console, select Library >> Library Servers >> [SCVMM Server Name] >> MSSCVMLib node and right-click. From the context menu select Explore .	
Paste the copied ISO into the MSSCVMM Library root folder (MSSCVMLib)	

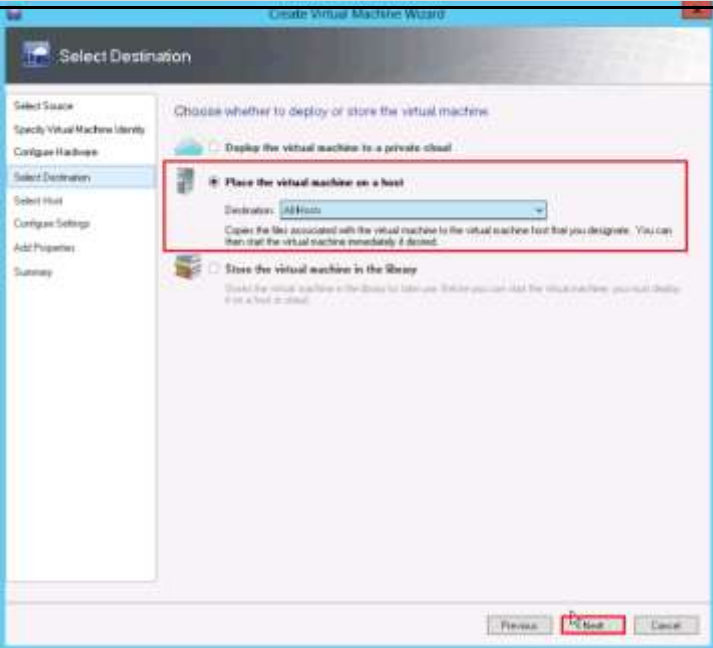
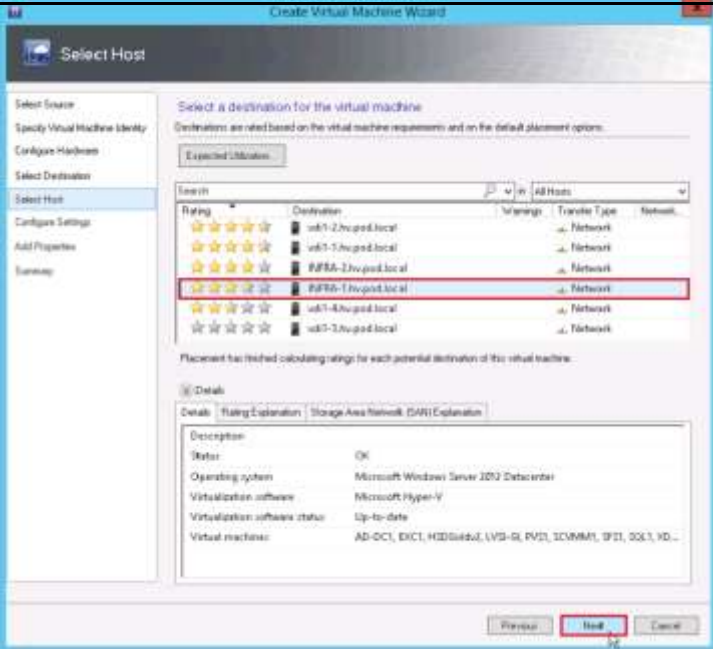
6.13.2. Installing and Configuring the Virtual Supervisor Modules (VSM)

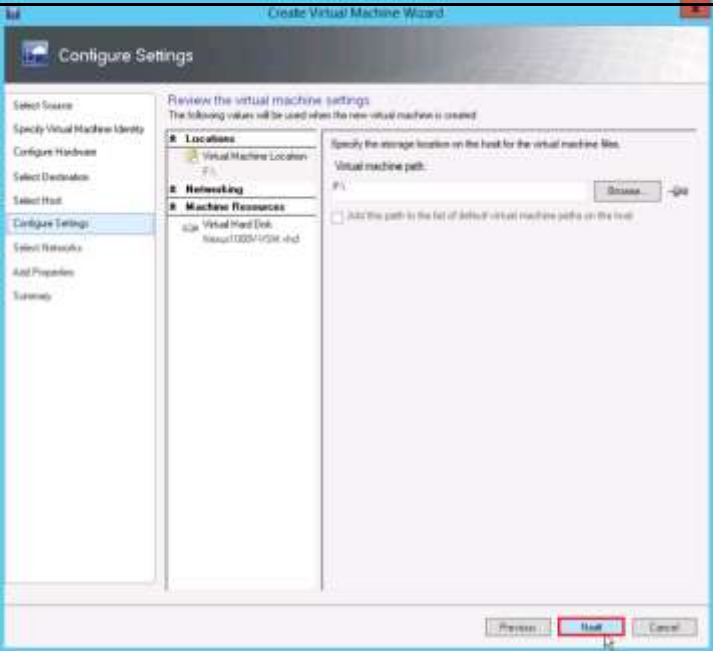
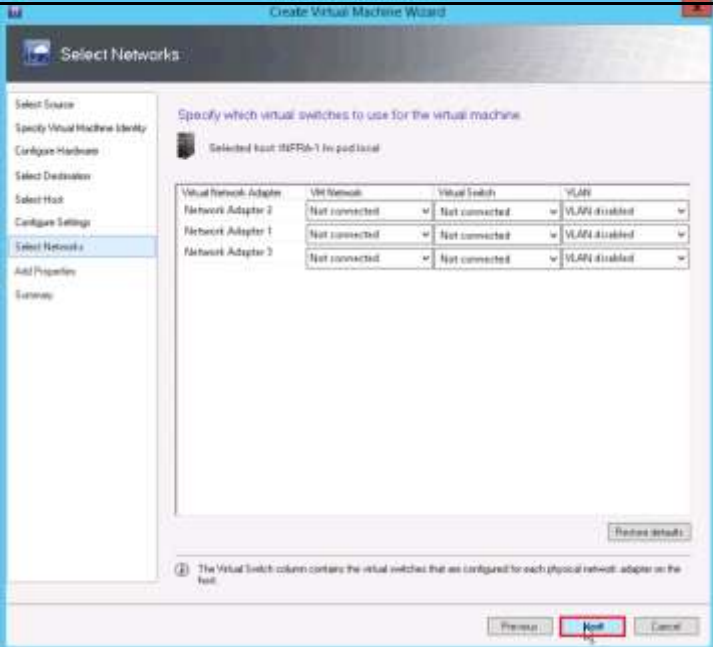
If building to support the 500-user configuration, you will only need on HA pair of VSMs for the environment. However, if building to support the 1000-user configuration, you will need two pairs of VSMs and the second pair will need a different **Domain ID** and **IP Address** assigned to them.

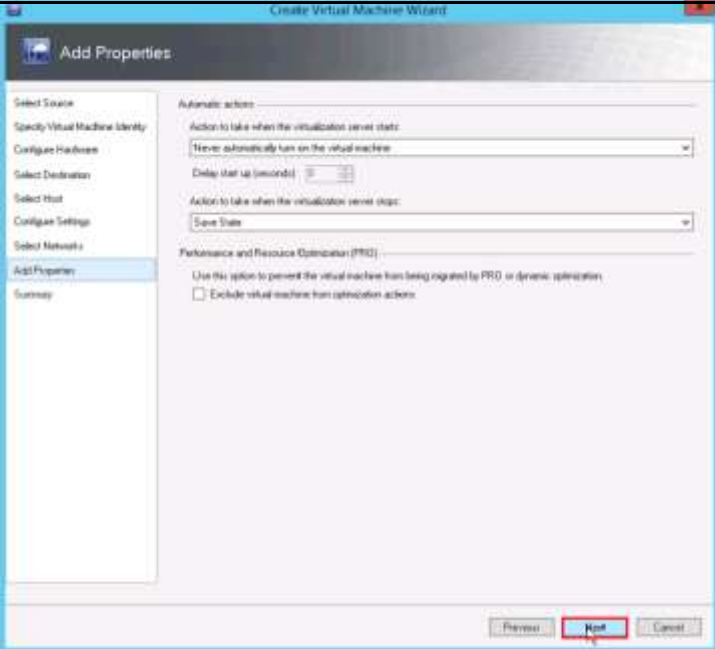
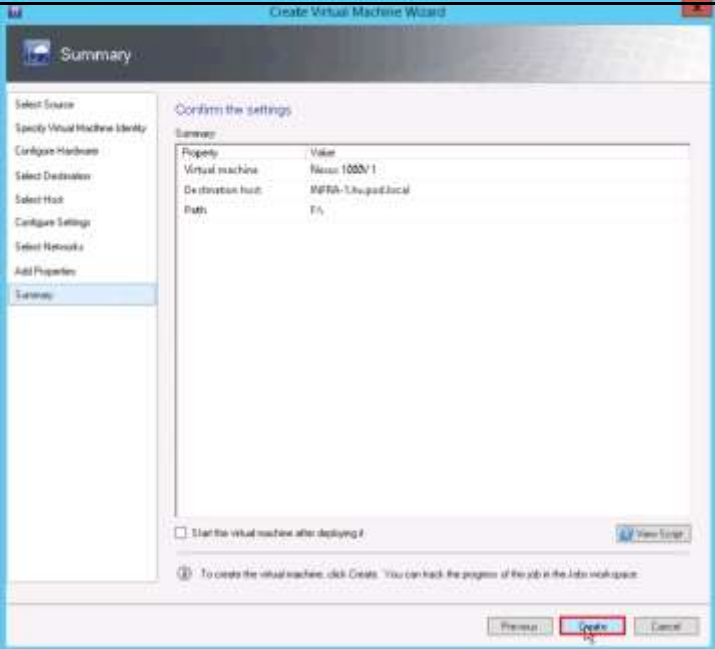
Instructions	Visual
--------------	--------

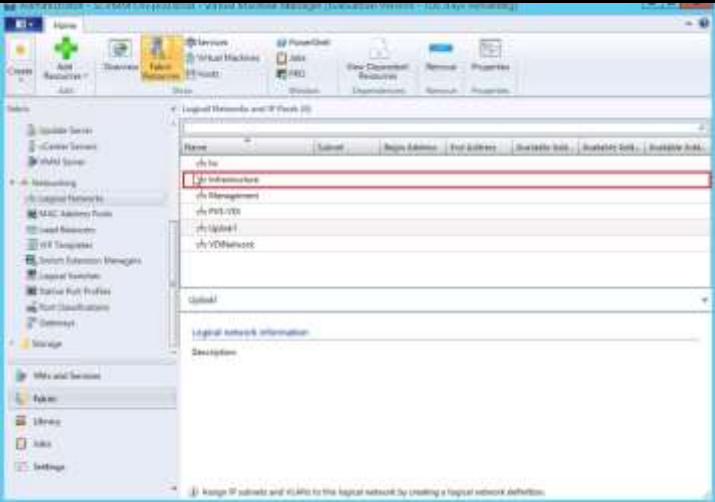
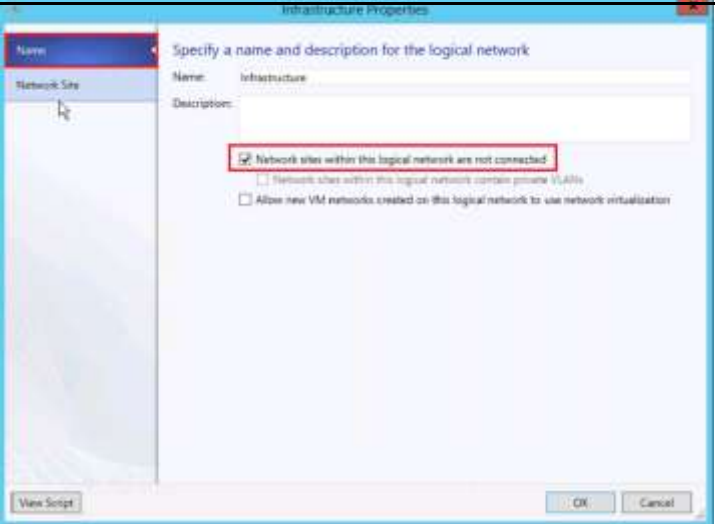
Instructions	Visual
<p>In the SCVMM Console, navigate to the VMs and Services tab.</p> <p>Right-click on the Infrastructure server and select Create Virtual Machine from the context menu.</p>	 <p>The screenshot shows the SCVMM console interface. On the left, the 'Infrastructure' server is selected in the tree view. A right-click context menu is open, and the 'Create Virtual Machine' option is highlighted. The main pane shows a list of existing VMs with columns for Name, CPU, Memory, Hard Disk, and Status.</p>
<p>The Create Virtual Machine Wizard starts.</p> <p>Click the Browse button and select the recently imported Nexus 1000V-VSM-Template.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Create Virtual Machine Wizard' with the 'Select Source' step. The 'Use an existing virtual machine, VM template, or virtual hard disk' radio button is selected. The text box below it contains 'Nexus1000V-VSM-Template', and the 'Browse' button is highlighted with a red box. The 'Next' button is also visible at the bottom right.</p>

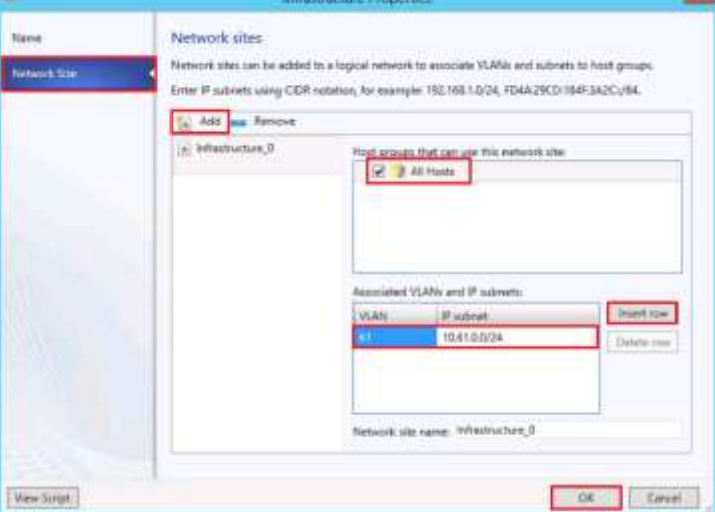
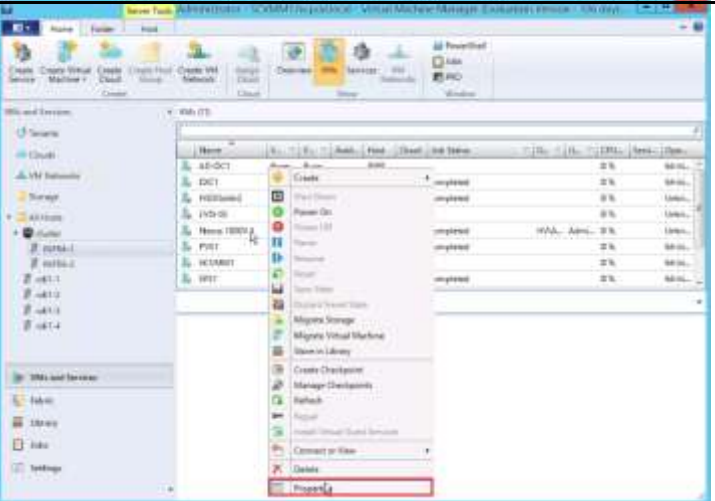
Instructions	Visual
<p>Provide a [Virtual Machine Name] for the Virtual Machine.</p> <p>Click Next.</p>	
<p>On the Configure Hardware tab, select the Virtual DVD drive and mount the nexus-1000v.5.2.1.SM1.5.1.iso added to the SCVMM library earlier.</p> <p>Click Next.</p>	

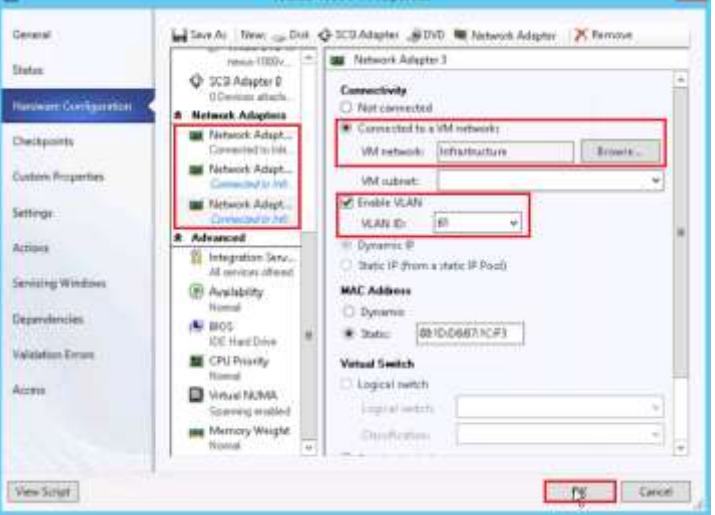
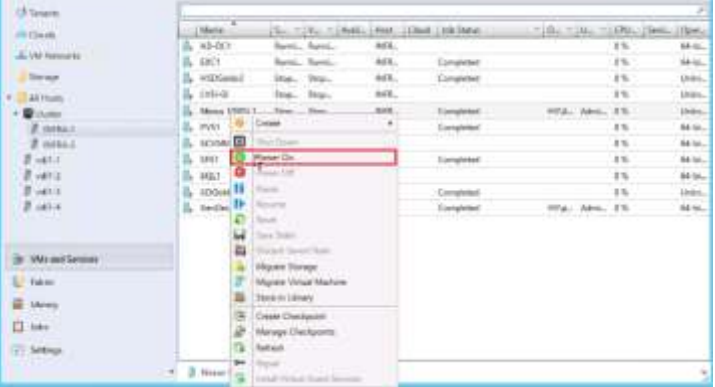
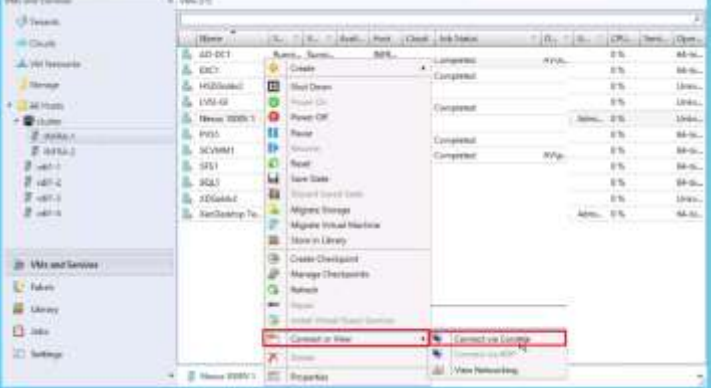
Instructions	Visual																																			
<p>On the Select Destination tab, select a Destination of All Hosts.</p> <p>Click Next.</p>																																				
<p>On the Select Host tab, select the infrastructure host.</p> <p>Click Next.</p>	 <table><tr><th>Rating</th><th>Destination</th><th>Warnings</th><th>Transfer Type</th><th>Network</th></tr><tr><td>★★★★★</td><td>win7-2.hyperv.local</td><td></td><td>Network</td><td>Network</td></tr><tr><td>★★★★★</td><td>win7-3.hyperv.local</td><td></td><td>Network</td><td>Network</td></tr><tr><td>★★★★★</td><td>P8R6-2.hyperv.local</td><td></td><td>Network</td><td>Network</td></tr><tr><td>★★★★★</td><td>P8R6-1.hyperv.local</td><td></td><td>Network</td><td>Network</td></tr><tr><td>★★★★★</td><td>win7-4.hyperv.local</td><td></td><td>Network</td><td>Network</td></tr><tr><td>★★★★★</td><td>win7-2.hyperv.local</td><td></td><td>Network</td><td>Network</td></tr></table>	Rating	Destination	Warnings	Transfer Type	Network	★★★★★	win7-2.hyperv.local		Network	Network	★★★★★	win7-3.hyperv.local		Network	Network	★★★★★	P8R6-2.hyperv.local		Network	Network	★★★★★	P8R6-1.hyperv.local		Network	Network	★★★★★	win7-4.hyperv.local		Network	Network	★★★★★	win7-2.hyperv.local		Network	Network
Rating	Destination	Warnings	Transfer Type	Network																																
★★★★★	win7-2.hyperv.local		Network	Network																																
★★★★★	win7-3.hyperv.local		Network	Network																																
★★★★★	P8R6-2.hyperv.local		Network	Network																																
★★★★★	P8R6-1.hyperv.local		Network	Network																																
★★★★★	win7-4.hyperv.local		Network	Network																																
★★★★★	win7-2.hyperv.local		Network	Network																																

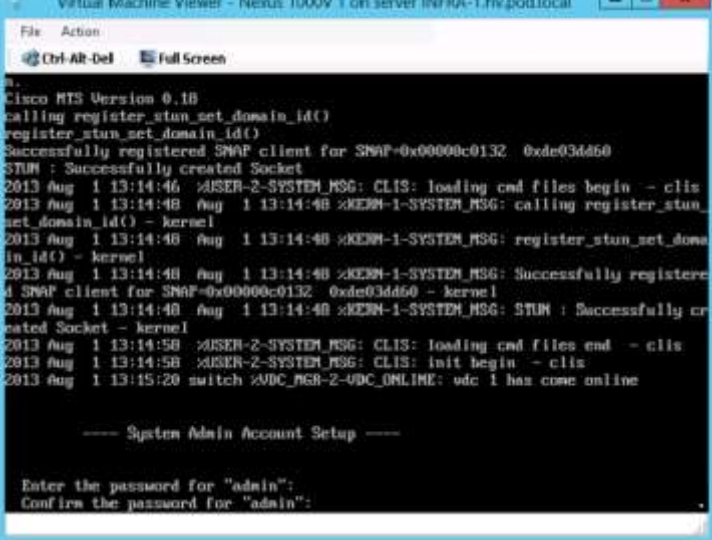
Instructions	Visual																
<p>On the Configure Settings tab, accept the defaults.</p> <p>Click Next</p>																	
<p>On the Select Networks tab, leave the defaults for now.</p> <p>Click Next.</p>	 <table data-bbox="961 1113 1477 1201"><thead><tr><th>Virtual Network Adapter</th><th>Virtual Network</th><th>Virtual Switch</th><th>VLAN</th></tr></thead><tbody><tr><td>Network Adapter 2</td><td>Not connected</td><td>Not connected</td><td>VLAN disabled</td></tr><tr><td>Network Adapter 1</td><td>Not connected</td><td>Not connected</td><td>VLAN disabled</td></tr><tr><td>Network Adapter 3</td><td>Not connected</td><td>Not connected</td><td>VLAN disabled</td></tr></tbody></table>	Virtual Network Adapter	Virtual Network	Virtual Switch	VLAN	Network Adapter 2	Not connected	Not connected	VLAN disabled	Network Adapter 1	Not connected	Not connected	VLAN disabled	Network Adapter 3	Not connected	Not connected	VLAN disabled
Virtual Network Adapter	Virtual Network	Virtual Switch	VLAN														
Network Adapter 2	Not connected	Not connected	VLAN disabled														
Network Adapter 1	Not connected	Not connected	VLAN disabled														
Network Adapter 3	Not connected	Not connected	VLAN disabled														

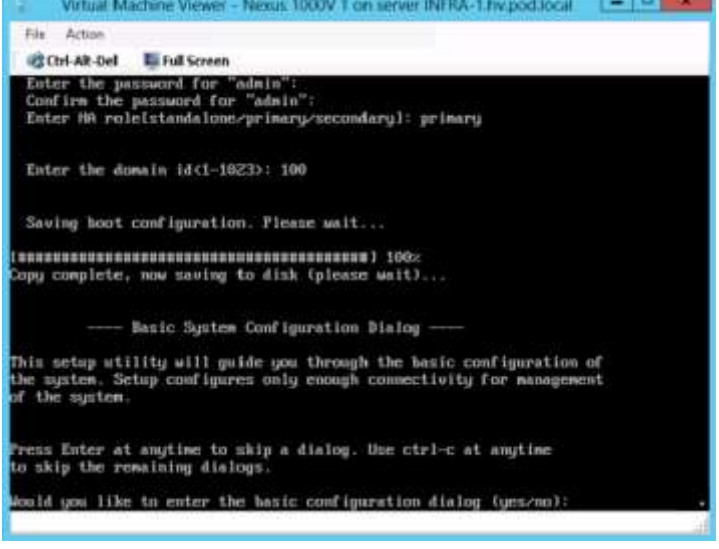
Instructions	Visual
<p>On the Add Properties tab, keep the defaults.</p> <p>Click Next.</p>	
<p>On the Summary tab, review the settings and if acceptable, click Create.</p> <p>Note: Do not auto-start the virtual machine because the networks still need to be setup.</p>	

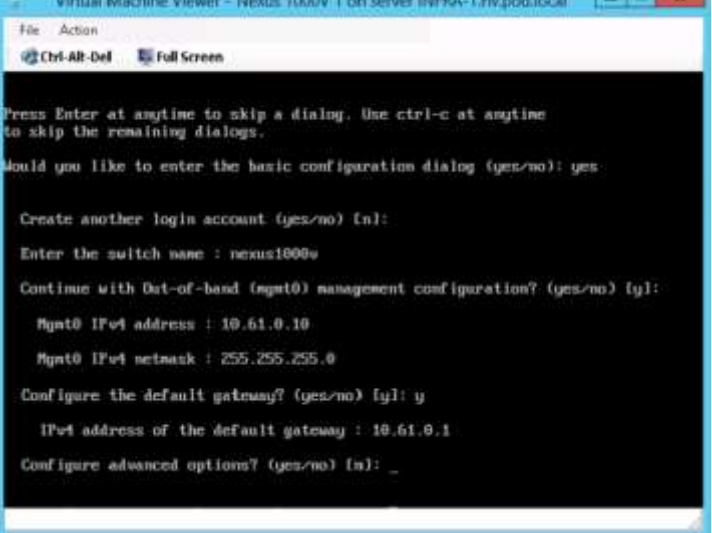
Instructions	Visual
<p>Select the Logical Networks node from the Fabric tab.</p> <p>Double-click the Infrastructure network to modify the properties.</p>	
<p>On the Name tab, enable the Network sites within this logical network are not connected checkbox.</p>	

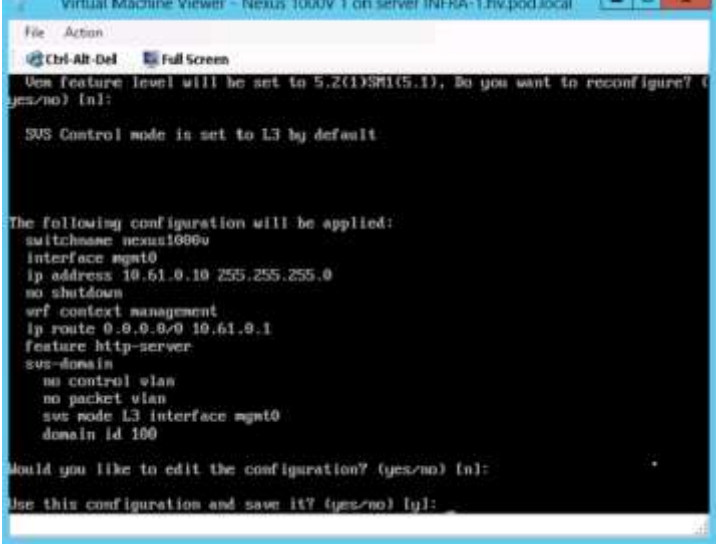
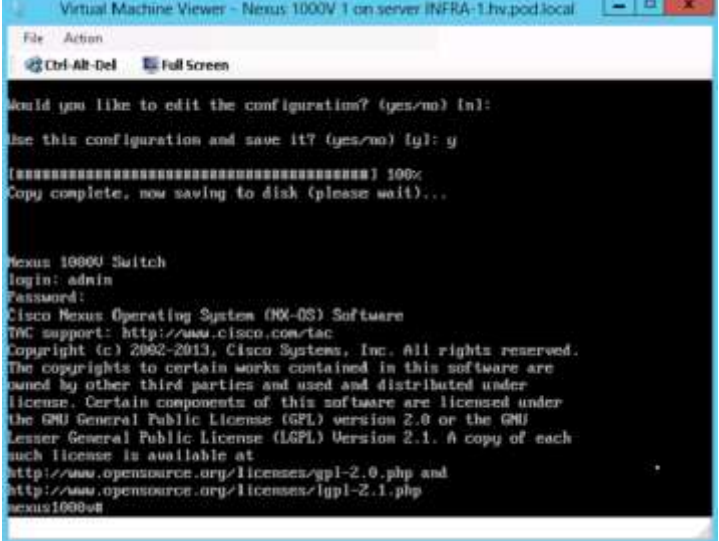
Instructions	Visual
<p>In the Properties dialog, select the Network Site tab.</p> <p>Click Add.</p> <p>Enable the All Hosts checkbox.</p> <p>Click Insert row.</p> <p>Set the VLAN to 61 and the IP subnet to 10.61.0.0/24</p> <p>Click OK.</p>	
<p>Locate the newly created VM on the VMs and Services tab.</p> <p>Right-click and select Properties from the context menu.</p>	

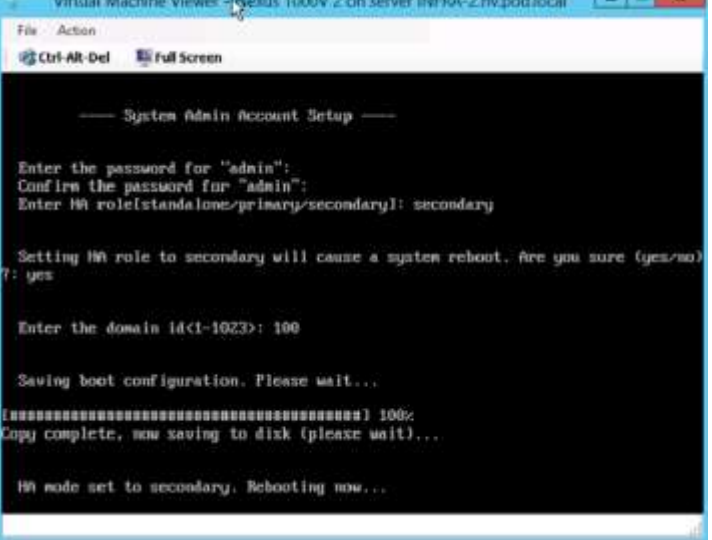
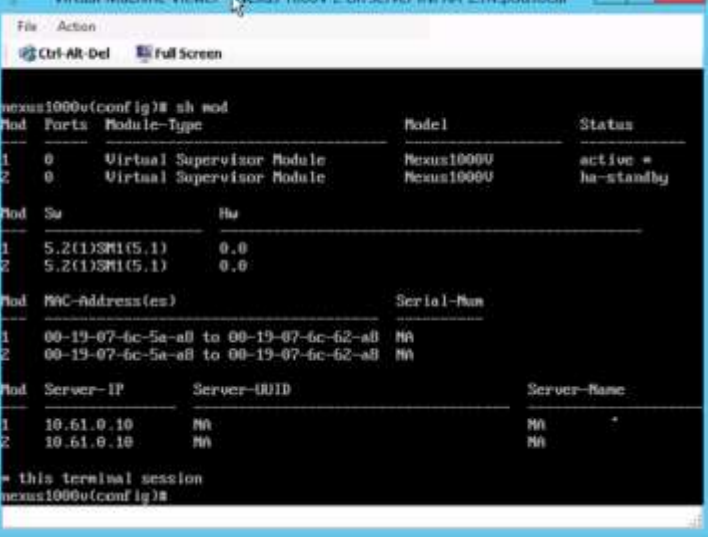
Instructions	Visual
<p>In the Hardware Configuration tab, select each of the three network adapters in turn and complete the following steps on each network adapter:</p> <ol style="list-style-type: none"> 1. Enable the Connected to a VM network radio button. 2. Select the Infrastructure VM network. 3. Enable the Enable VLAN checkbox. 4. Select the VLAN ID 61. <p>Click OK.</p>	
<p>After it finishes updating, right-click Cisco Nexus 1000V virtual machine.</p> <p>Select Power On from the context menu to start the virtual switch.</p>	
<p>Right-click Cisco Nexus 1000V virtual machine again.</p> <p>Select Connect or View, followed by Connect via Console to start a console session.</p>	

Instructions	Visual
<p>Provide a password for the admin account.</p> <p>Confirm the password for the admin account.</p> <p>Note: At this point the password must conform to composition checks, so select one that will meet those requirements, such as P@ssW0rd.</p> <p>When both Nexus 1000V virtual supervisor modules are setup, the admin password can be changed from the primary node. To change the admin password, run the following commands, where <newpass> is the new admin password.</p> <pre>#config terminal (config)#no password strength-check (config)#username admin password <newpass> (config)#copy run start (config)#exit</pre>	

Instructions	Visual
<p>Enter the HA Role: Primary (for first N1000V).</p> <p>Provide the Domain ID. [Domain ID]</p> <p>Note: The Domain ID <u>must</u> be the same on both the primary and secondary N1000V virtual supervisor modules. Furthermore, Domain IDs must be unique for each pair of N1000V virtual supervisor modules.</p>	

Instructions	Visual
<p>Would you like to enter the basic configuration dialog: Yes</p> <p>Create another login account: No</p> <p>Enter the switch name: [SwitchName]</p> <p>Continue with Out-of-band (mgmt0) management configuration: Yes</p> <p>Mgmt0 IPv4 address: [Mgmt0 IP Address]</p> <p>Mgmt0 IPv4 netmask: [Mgmt0 Subnet Mask]</p> <p>Configure the default gateway: Yes</p> <p>IPv4 address of the default gateway: [Mgmt0 Gateway IP]</p> <p>Configure advanced options: No</p>	

Instructions	Visual
<p>Do you want to reconfigure the VEM feature level: No</p> <p>Confirm the displayed configuration is correct.</p> <p>Would you like to edit the configuration: No</p> <p>Use this configuration and save it: Yes</p>	 <pre> Virtual Machine Viewer - Nexus 1000V 1 on server INFRA-1.hw.pod.local File Action Ctrl-Alt-Del Full Screen Use feature level will be set to 5.2(1)SM(5.1). Do you want to reconfigure? (yes/no) [n]: SWS Control mode is set to L3 by default. The following configuration will be applied: switchname nexus1000v interface mgmt0 ip address 10.61.0.10 255.255.255.0 no shutdown vrf context management ip route 0.0.0.0/0 10.61.0.1 feature http-server svs-domain no control vlan no packet vlan sws mode L3 interface mgmt0 domain id 100 Would you like to edit the configuration? (yes/no) [n]: Use this configuration and save it? (yes/no) [y]: </pre>
<p>The configuration is saved and the switch returns to a login prompt.</p> <p>Provide the username and password combination created earlier during setup.</p> <p>Verify network connectivity by pinging the gateway address.</p>	 <pre> Virtual Machine Viewer - Nexus 1000V 1 on server INFRA-1.hw.pod.local File Action Ctrl-Alt-Del Full Screen Would you like to edit the configuration? (yes/no) [n]: Use this configuration and save it? (yes/no) [y]: y [*****] 100% Copy complete, now saving to disk (please wait)... Nexus 1000V Switch login: admin Password: Cisco Nexus Operating System (NX-OS) Software TAC support: http://www.cisco.com/tac Copyright (c) 2002-2013, Cisco Systems, Inc. All rights reserved. The copyrights to certain works contained in this software are owned by other third parties and used and distributed under license. Certain components of this software are licensed under the GNU General Public License (GPL) version 2.0 or the GNU Lesser General Public License (LGPL) Version 2.1. A copy of each such license is available at http://www.opensource.org/licenses/gpl-2.0.php and http://www.opensource.org/licenses/lgpl-2.1.php nexus1000v# </pre>

Instructions	Visual
<p>Repeat the all the steps in this section for the second Cisco Nexus 1000V virtual supervisor module VM.</p> <p>Provide the admin password, which <u>must</u> be the same admin password used for the primary Cisco Nexus 1000V.</p> <p>Enter HA role: secondary</p> <p>System Reboot: yes</p> <p>Domain ID: [Domain ID]</p> <p>Note: The Domain ID <u>must</u> match the primary's Domain ID.</p>	 <p>The screenshot shows the 'System Admin Account Setup' window in a terminal. The user is configuring a secondary HA role. The steps shown are: entering the password for 'admin', confirming the password, entering the HA role as 'secondary', confirming the reboot, entering the domain ID as '100', and saving the boot configuration. The terminal output shows the HA mode set to secondary and the system rebooting.</p>
<p>Verify the two VSMs are functioning by running show module from the primary.</p> <p>Note: If building for 1000-users, repeat the section again to create another pair of VSMs with a different Domain ID.</p>	 <p>The screenshot shows the output of the 'show module' command in the primary VSM's configuration mode. The output displays two modules: 'Virtual Supervisor Module' and 'Virtual Supervisor Module'. The first module is 'active' and the second is 'ha-standby'. Below this, the output shows the 'show module' command output for the HA standby module, including the 'Mod', 'Sta', and 'Hw' columns. The output also shows the 'show module' command output for the HA standby module, including the 'Mod', 'Serial-Name', and 'Server-Name' columns.</p>

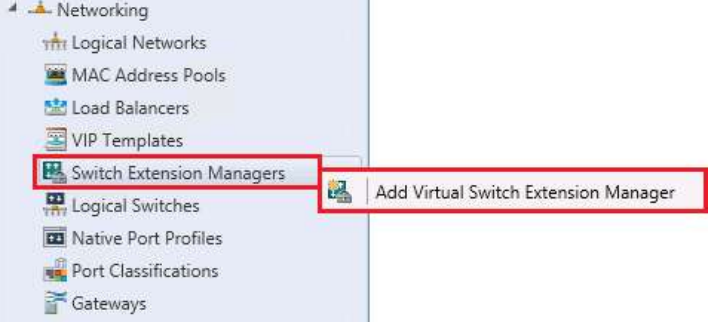
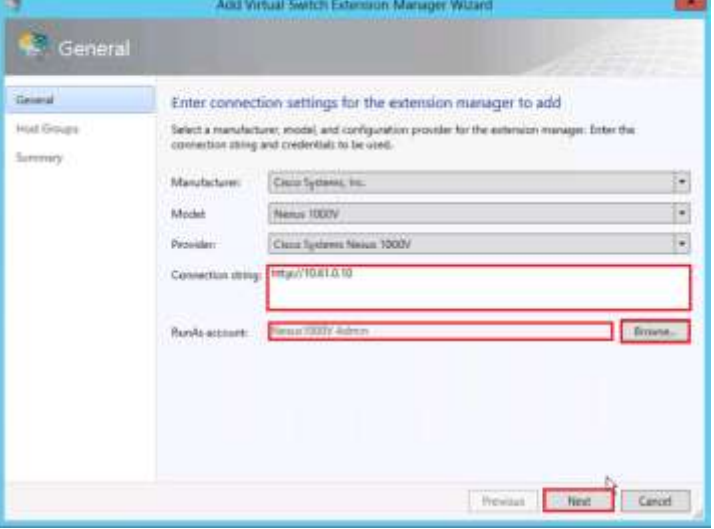
6.13.3. Cisco Nexus 1000V for Hyper-V Configuration

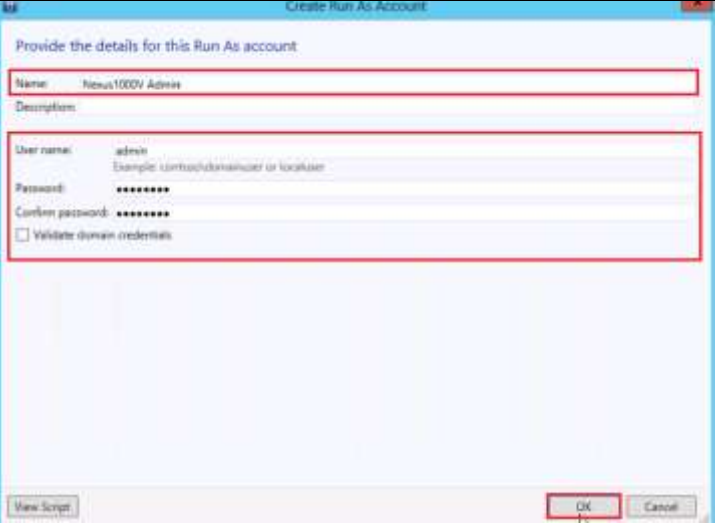
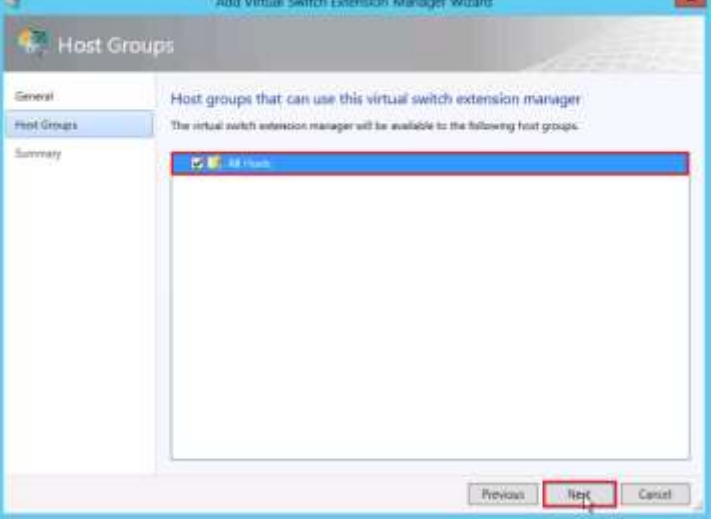
If building the 500-user environment, complete this section only once. However, if building for a 1000-user environment, two VSMs will be needed, and the section should be completed once for each pair of VSMs.

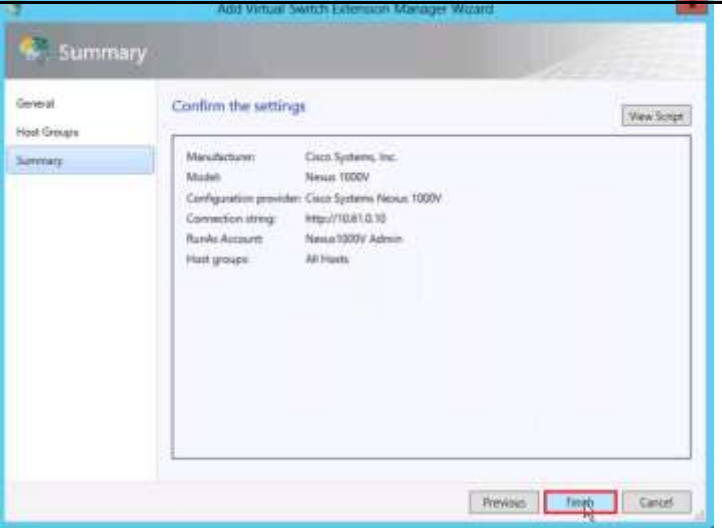
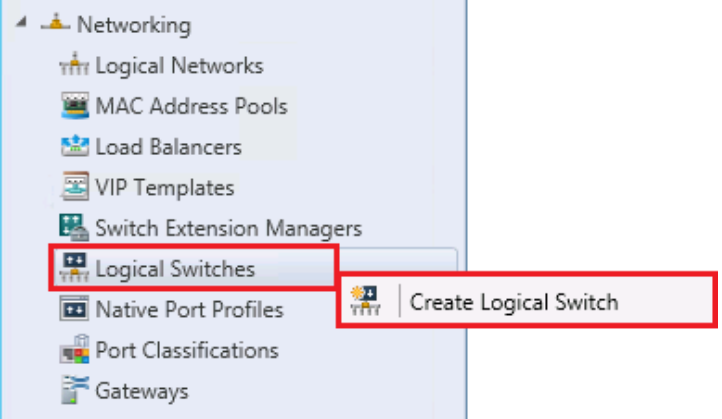
Instructions	Visual
<p>Use PuTTY or a similar SSH client to connect to the Mgmt0 interface on the primary VSM.</p>	
<p>Enter configuration mode:</p> <p>Config terminal</p>	<pre>nexus1000v# conf t Enter configuration commands, one per line. End with CNTL/Z.</pre>
<p>Create the logical networks for VDI and Cluster traffic:</p> <pre>(config)# nsm logical network [Network_Name] (config)# exit</pre>	<pre>nexus1000v(config)# nsm logical network VDINetwork nexus1000v(config-logical-net)# exit nexus1000v(config)# nsm logical network ClusterNetwork nexus1000v(config-logical-net)# exit</pre>
<p>Create the network segment pools for the VDI and the Cluster networks and bind them to the logical network:</p> <pre>(config)# nsm network segment pool [VDI_Pool_Name] (config-net-seg-pool)# member-of logical network [Network_Name] (config-net-seg-pool)# exit</pre>	<pre>nexus1000v(config)# nsm network segment pool VDI-Pool-1 nexus1000v(config-net-seg-pool)# member-of logical network VDINetwork nexus1000v(config-net-seg-pool)# exit nexus1000v(config)# nsm network segment pool Cluster-Pool-1 nexus1000v(config-net-seg-pool)# member-of logical network ClusterNetwork nexus1000v(config-net-seg-pool)# exit</pre>

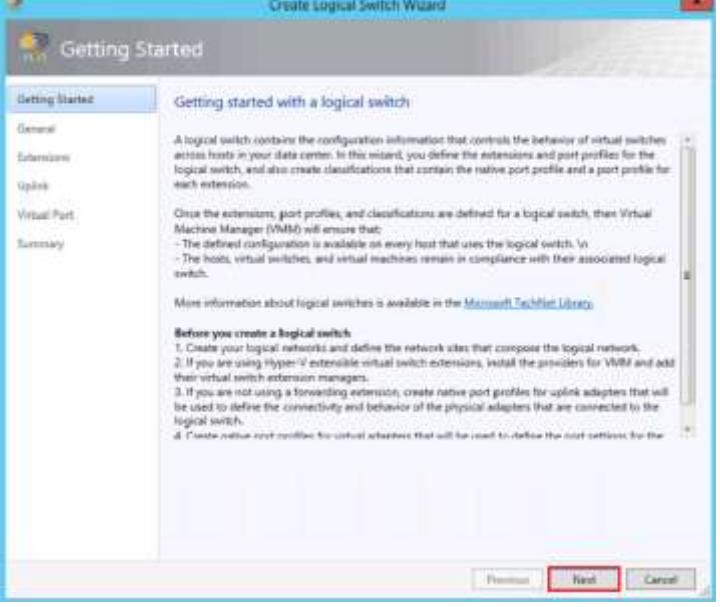
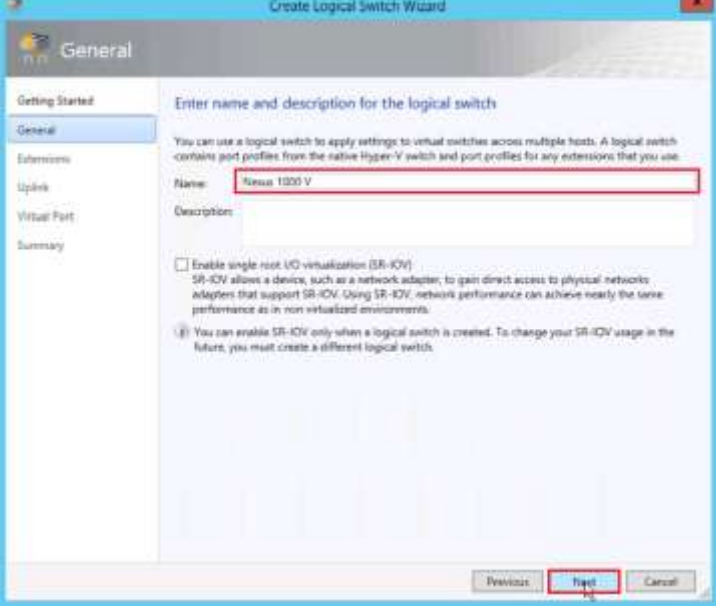
Instructions	Visual
<p>Create the IP Pool templates and set the DHCP range for Infrastructure (VLAN 61), PVS (VLAN 62), Cluster (VLAN 63), and LiveMigration (VLAN 64) networks</p> <pre>(config)# nsm ip pool template [VLAN_Pool_Name] (config-ip-pool-template)# ip address [Start_IP] [End_IP] (config-ip-pool-template)# network [Network_ID] [Network_Mask] (config-ip-pool-template)# default-router [Router_Address] (config-ip-pool-template)# exit</pre>	<pre>nexus1000v(config)# nsm ip pool template VLAN-61-Pool nexus1000v(config-ip-pool-template)# ip address 10.61.0.11 10.61.0.11 nexus1000v(config-ip-pool-template)# network 10.61.0.0 255.255.255.0 nexus1000v(config-ip-pool-template)# default-router 10.61.0.1 nexus1000v(config-ip-pool-template)# exit nexus1000v(config)# nsm ip pool template VLAN-62-Pool nexus1000v(config-ip-pool-template)# ip address 10.62.0.100 10.62.0.150 nexus1000v(config-ip-pool-template)# network 10.62.0.0 255.255.240.0 nexus1000v(config-ip-pool-template)# default-router 10.62.0.1 nexus1000v(config-ip-pool-template)# exit nexus1000v(config)# nsm ip pool template VLAN-63-Pool nexus1000v(config-ip-pool-template)# ip address 10.63.0.100 10.63.0.150 nexus1000v(config-ip-pool-template)# network 10.63.0.0 255.255.255.0 nexus1000v(config-ip-pool-template)# default-router 10.63.0.1 nexus1000v(config-ip-pool-template)# nsm ip pool template VLAN-64-Pool nexus1000v(config-ip-pool-template)# ip address 10.64.0.100 10.64.0.150 nexus1000v(config-ip-pool-template)# network 10.64.0.0 255.255.255.0 nexus1000v(config-ip-pool-template)# default-router 10.64.0.1 nexus1000v(config-ip-pool-template)# exit</pre>
<p>Create Network Segments and import the VLAN pool into the segment for the Infrastructure, PVS, Cluster, and LiveMigration VLANs. The VDI Pool will manage VLANs 61-62. The Cluster Pool will manage VLANs 63-64.</p> <pre>(config)# nsm network segment [VLAN_ID] (config-net-seg)# switchport mode access (config-net-seg)# switchport access vlan [VLAN_##] (config-net-seg)# member-of network segment pool [VDI_Pool_Name] (config-net-seg)# ip pool import template [VLAN_Pool_Name] (config-net-seg)# publish network segment (config-net-seg)# exit</pre>	<pre>nexus1000v(config-net-seg)# switchport mode access nexus1000v(config-net-seg)# switchport access vlan 61 nexus1000v(config-net-seg)# member-of network segment pool VDI-Pool-1 nexus1000v(config-net-seg)# ip pool import template VLAN-61-Pool nexus1000v(config-net-seg)# publish network segment nexus1000v(config-net-seg)# exit nexus1000v(config)# nsm network segment VLAN-62 nexus1000v(config-net-seg)# switchport mode access nexus1000v(config-net-seg)# switchport access vlan 62 nexus1000v(config-net-seg)# member-of network segment pool VDI-Pool-1 nexus1000v(config-net-seg)# ip pool import template VLAN-62-Pool nexus1000v(config-net-seg)# publish network segment nexus1000v(config-net-seg)# exit nexus1000v(config)# nsm network segment VLAN-63 nexus1000v(config-net-seg)# switchport mode access nexus1000v(config-net-seg)# switchport access vlan 63 nexus1000v(config-net-seg)# member-of network segment pool Cluster-Pool-1 nexus1000v(config-net-seg)# ip pool import template VLAN-63-Pool nexus1000v(config-net-seg)# publish network segment nexus1000v(config-net-seg)# exit nexus1000v(config)# nsm network segment VLAN-64 nexus1000v(config-net-seg)# switchport mode access nexus1000v(config-net-seg)# switchport access vlan 64 nexus1000v(config-net-seg)# member-of network segment pool Cluster-Pool-1 nexus1000v(config-net-seg)# ip pool import template VLAN-64-Pool nexus1000v(config-net-seg)# publish network segment nexus1000v(config-net-seg)# exit</pre>

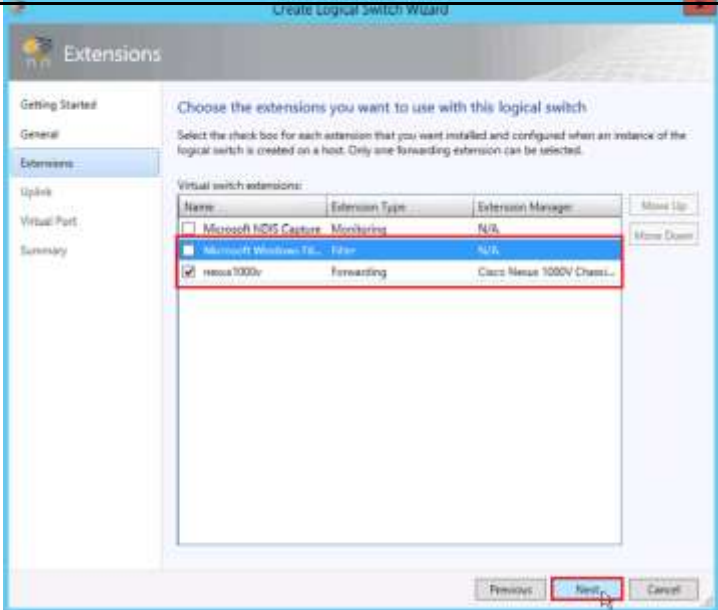
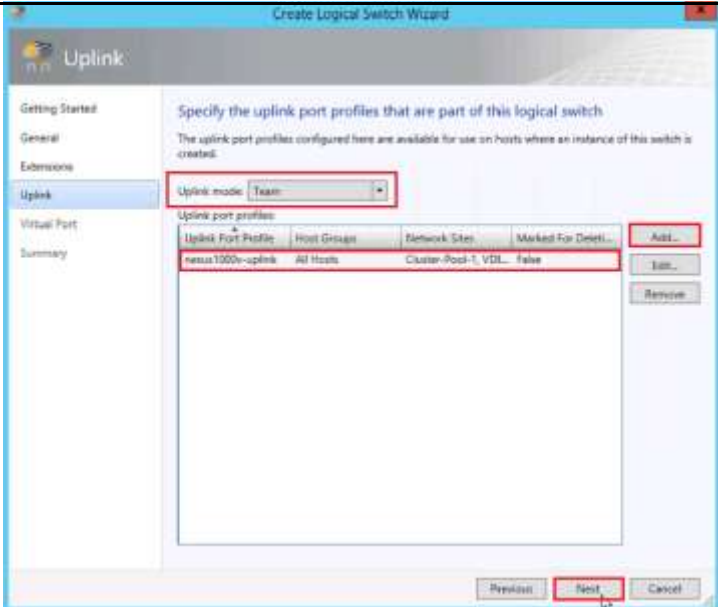
Instructions	Visual
<p>Create vEthernet and Ethernet Port Profiles</p> <pre> (config)# port-profile type vEthernet [vEthernet_Port_profile] (config-port-prof) # no shutdown (config-port-prof) # state enabled (config-port-prof) # max-ports 1024 (config-port-prof) # port-binding static auto expand (config-port-prof) # publish port-profile (config-port-prof) # exit (config)# port-profile type Ethernet Uplink-profile [Uplink_Port_profile] (config-port-prof) # no shutdown (config-port-prof) # state enabled (config-port-prof) # channel-group auto mode on (config-port-prof) # channel-group auto mode on mac-pinning (config-port-prof) # exit (config) # nsm network uplink [Uplink_profile_name] (config-uplink-net) # all network segment pool [VDI_Pool_Name] (config-uplink-net) # allow network segment pool[Cluster_Pool_Name] (config-uplink-net) # import port-profile [Uplink_Port_profile] (config-uplink-net) # publish network uplink (config-uplink-net) # exit (config)# copy run start </pre>	<pre> nexus1000v(config)# port-profile type vEthernet VDI-Port-profile nexus1000v(config-port-prof)# no shut nexus1000v(config-port-prof)# state enabled nexus1000v(config-port-prof)# max-ports 1024 nexus1000v(config-port-prof)# port-binding static auto expand nexus1000v(config-port-prof)# publish port-profile nexus1000v(config-port-prof)# exit nexus1000v(config)# port-profile type Ethernet Uplink-profile nexus1000v(config-port-prof)# no shut nexus1000v(config-port-prof)# state en nexus1000v(config-port-prof)# channel-group auto mode on nexus1000v(config-port-prof)# channel-group auto mode on mac-pinning nexus1000v(config-port-prof)# exit nexus1000v(config)# nsm network uplink nexus1000v-uplink nexus1000v(config-uplink-net)# allow network segment pool VDI-Pool-1 nexus1000v(config-uplink-net)# allow network segment pool Cluster-Pool-1 nexus1000v(config-uplink-net)# import port-profile Uplink-profile nexus1000v(config-uplink-net)# publish network uplink nexus1000v(config-uplink-net)# exit </pre>

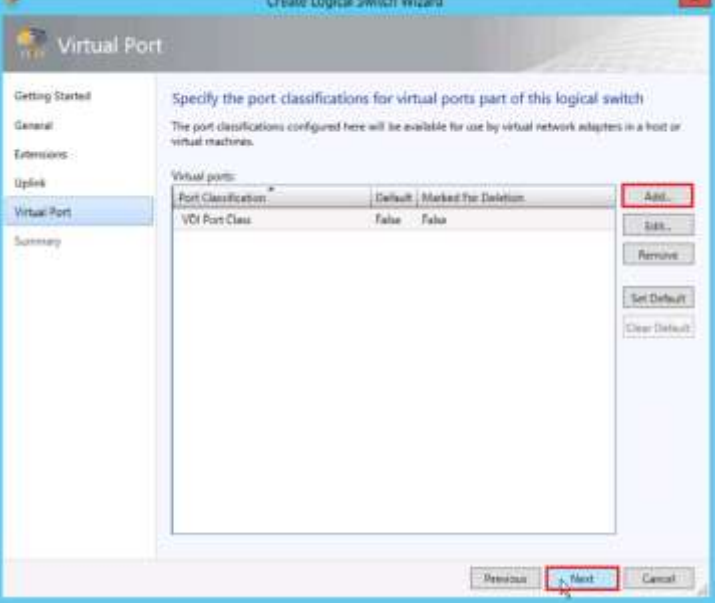
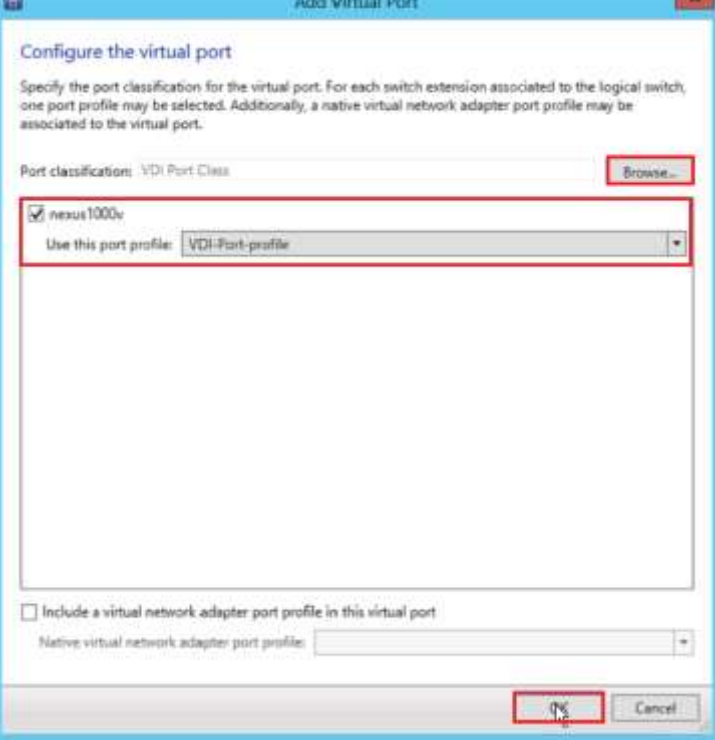
Instructions	Visual
<p>Return to the SCVMM Console and select the Fabric tab.</p> <p>Under Networking >> Switch Extension Manager.</p> <p>Right-click and choose Add Virtual Switch Extension Manager from the context menu.</p>	
<p>In the Add Virtual Switch Extension Manager Wizard, set the connection string to the http://<Mgmt0 IP Address></p> <p>Set the RunAs account to the Admin account on the VSM (see next step) by clicking the Browse button.</p> <p>Click Next when completed.</p>	

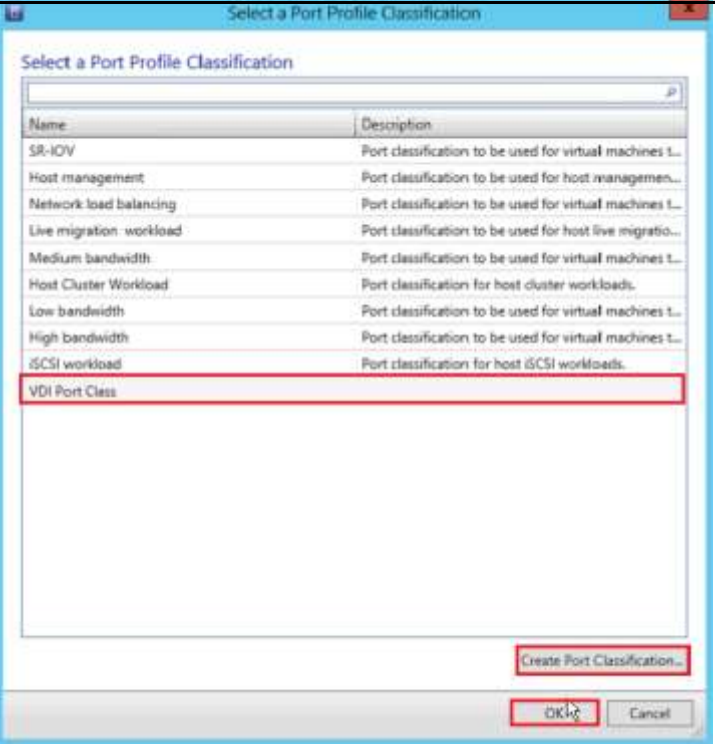
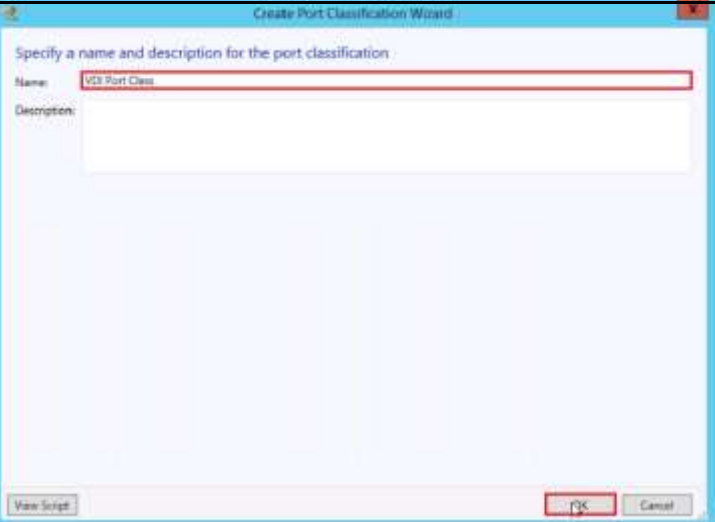
Instructions	Visual
<p>To add a RunAs Account for the VSM complete the following:</p> <p>Provide a Name: [AdminRunAsAccount]</p> <p>Username: admin</p> <p>Password: [Admin Password]</p> <p>Confirm Password: [Admin Password]</p> <p>Validate domain credentials: Disabled</p> <p>Click OK.</p>	
<p>Returning to the Add Virtual Switch Extension Manager Wizard.</p> <p>On the Hosts Groups tab, enable the All Hosts checkbox.</p> <p>Click Next.</p>	


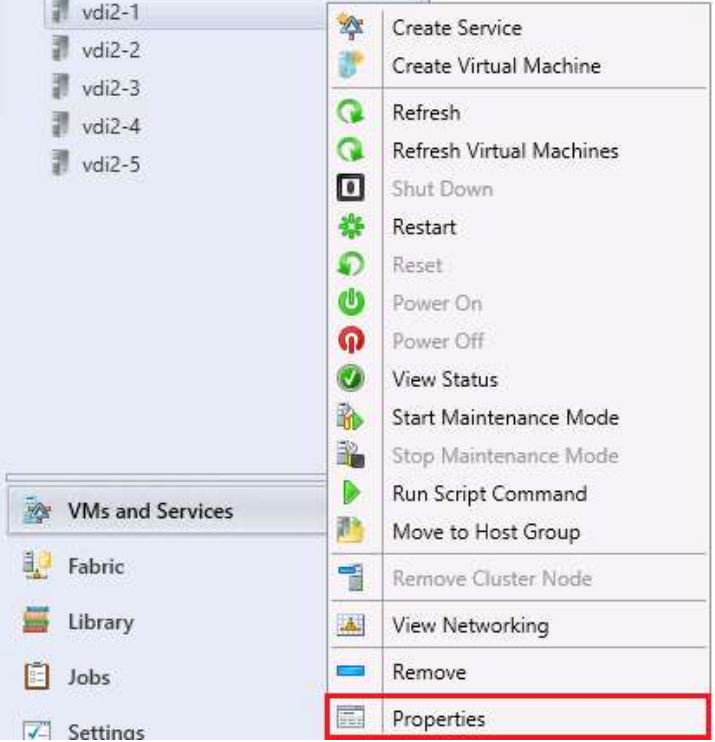
Instructions	Visual
<p>On the Summary tab, click Finish.</p> <p>Verify the switch was created successfully.</p>	
<p>On the Fabric tab, navigate to Networking >> Logical Switches</p> <p>Right-click Logical Switches and choose Create Logical Switch from the context menu.</p>	

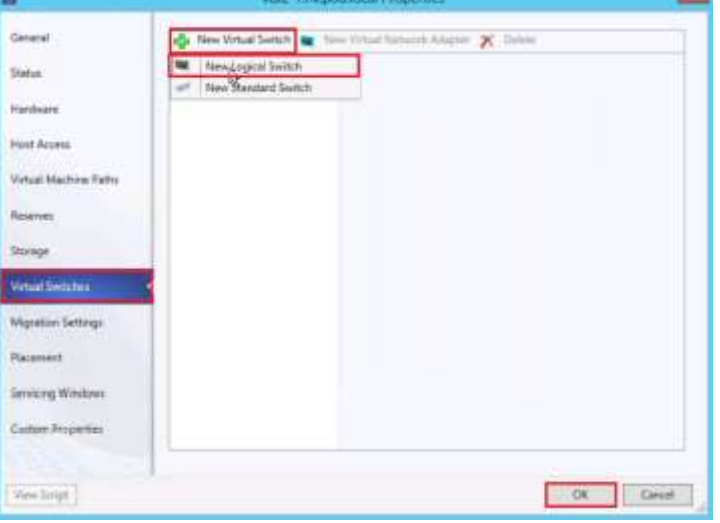
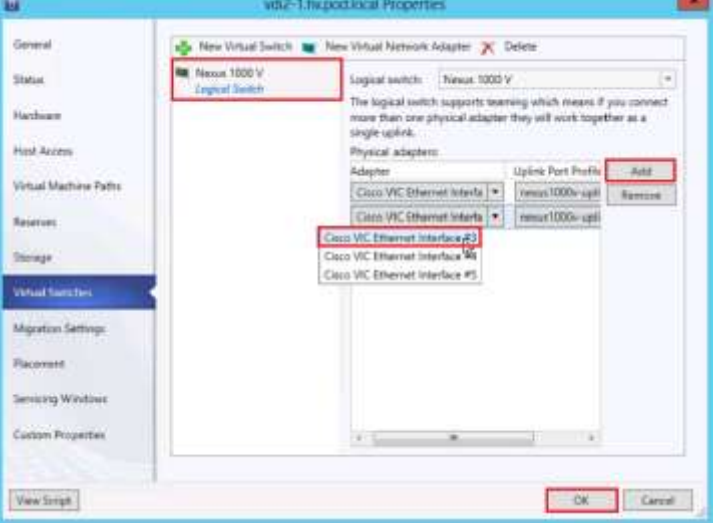
Instructions	Visual
<p>From the Getting Started tab, click Next.</p>	 <p>The screenshot shows the 'Create Logical Switch Wizard' with the 'Getting Started' tab selected. The left sidebar contains links for 'Getting Started', 'General', 'Extensions', 'Uplink', 'Virtual Port', and 'Summary'. The main content area is titled 'Getting started with a logical switch' and provides introductory text and a list of steps to follow before creating a logical switch. At the bottom right, the 'Next' button is highlighted with a red rectangular box.</p>
<p>From the General tab, provide a Name for the logical switch.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Create Logical Switch Wizard' with the 'General' tab selected. The left sidebar is the same as the previous screenshot. The main content area is titled 'Enter name and description for the logical switch'. It includes a text box for 'Name' containing 'Nexus 1000 V', which is highlighted with a red rectangular box. Below it is a 'Description' text box. There are also checkboxes for 'Enable single root I/O virtualization (SR-IOV)' and a note about SR-IOV usage. At the bottom right, the 'Next' button is highlighted with a red rectangular box.</p>


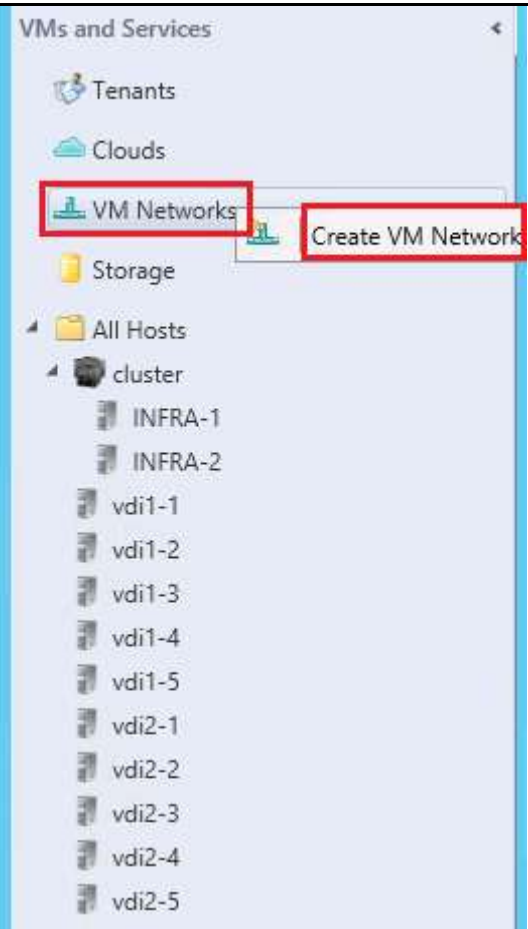
Instructions	Visual												
<p>From the Extensions tab, disable the Microsoft Windows Filtering checkbox and enable the Cisco Nexus1000v checkbox.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Create Logical Switch Wizard' with the 'Extensions' tab selected. The table lists the following extensions:</p> <table><tr><th>Name</th><th>Extension Type</th><th>Extension Manager</th></tr><tr><td>Microsoft HCVS Capture</td><td>Monitoring</td><td>N/A</td></tr><tr><td>Microsoft Windows Filtering</td><td>Filter</td><td>N/A</td></tr><tr><td>Cisco Nexus1000v</td><td>Forwarding</td><td>Cisco Nexus 1000V Control Plane</td></tr></table> <p>The 'Microsoft Windows Filtering' checkbox is unchecked, and the 'Cisco Nexus1000v' checkbox is checked. The 'Next' button is highlighted with a red box.</p>	Name	Extension Type	Extension Manager	Microsoft HCVS Capture	Monitoring	N/A	Microsoft Windows Filtering	Filter	N/A	Cisco Nexus1000v	Forwarding	Cisco Nexus 1000V Control Plane
Name	Extension Type	Extension Manager											
Microsoft HCVS Capture	Monitoring	N/A											
Microsoft Windows Filtering	Filter	N/A											
Cisco Nexus1000v	Forwarding	Cisco Nexus 1000V Control Plane											
<p>From the Uplink tab, set the Uplink mode to Team.</p> <p>Click Add.</p> <p>Select the Uplink port profile [Uplink_profile_name] created earlier</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Create Logical Switch Wizard' with the 'Uplink' tab selected. The 'Uplink mode' is set to 'Team'. The 'Add...' button is highlighted with a red box. Below the table, the 'Uplink port profiles' section shows a table with one entry:</p> <table><tr><th>Uplink Port Profile</th><th>Host Groups</th><th>Network Sites</th><th>Marked For Deletion</th></tr><tr><td>nexus1000v-uplink</td><td>All Hosts</td><td>Cluster-Pool-1, VDL...</td><td>false</td></tr></table> <p>The 'Add...' button is highlighted with a red box.</p>	Uplink Port Profile	Host Groups	Network Sites	Marked For Deletion	nexus1000v-uplink	All Hosts	Cluster-Pool-1, VDL...	false				
Uplink Port Profile	Host Groups	Network Sites	Marked For Deletion										
nexus1000v-uplink	All Hosts	Cluster-Pool-1, VDL...	false										

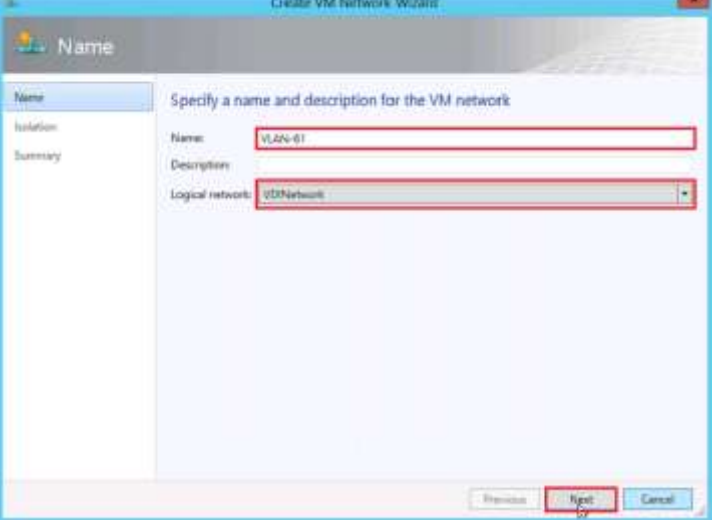
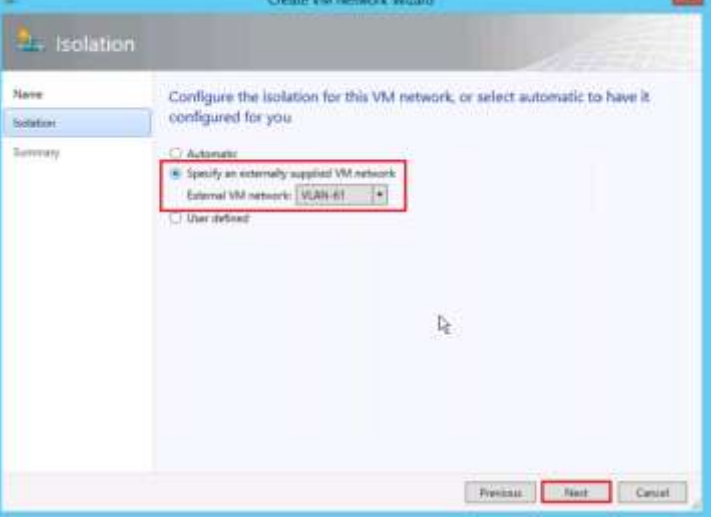
Instructions	Visual
<p>From the Virtual Port tab, click Add to add a port class which has not yet been defined. (see next three steps.)</p> <p>When the port classification has been created, click Next.</p>	 <p>The screenshot shows the 'Virtual Port' configuration window. On the left, the 'Virtual Port' tab is selected. In the main area, there is a table for 'Virtual ports' with columns 'Port Classification', 'Default', and 'Marked For Deletion'. The 'Add...' button is highlighted with a red box. At the bottom, the 'Next' button is also highlighted with a red box.</p>
<p>From the Add Virtual Port Dialog, click Browse to add a new Port classification profile (see next two steps)</p> <p>When the port classification has been added, enable the checkbox next to nexus1000v hostname and select the [vEthernet_Port_profile] from the drop-down box.</p> <p>Click OK</p>	 <p>The screenshot shows the 'Add Virtual Port' dialog. The 'Port classification' field is set to 'VDI Port Class'. The 'Browse...' button is highlighted with a red box. Below, the 'nexus1000v' checkbox is checked, and 'vEthernet-Port-profile' is selected in the 'Use this port profile' drop-down menu. At the bottom, the 'OK' button is highlighted with a red box.</p>


Instructions	Visual																						
<p>In the Select a Port-Profile Classification dialog, click the Create Port Classification button. (See the next step to complete the addition of a port profile classification.)</p> <p>When added, click OK.</p>	 <p>The screenshot shows the 'Select a Port Profile Classification' dialog box. It contains a table with the following entries:</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>SR-IOV</td> <td>Port classification to be used for virtual machines t...</td> </tr> <tr> <td>Host management</td> <td>Port classification to be used for host managemen...</td> </tr> <tr> <td>Network load balancing</td> <td>Port classification to be used for virtual machines t...</td> </tr> <tr> <td>Live migration workload</td> <td>Port classification to be used for host live migratio...</td> </tr> <tr> <td>Medium bandwidth</td> <td>Port classification to be used for virtual machines t...</td> </tr> <tr> <td>Host Cluster Workload</td> <td>Port classification for host cluster workloads.</td> </tr> <tr> <td>Low bandwidth</td> <td>Port classification to be used for virtual machines t...</td> </tr> <tr> <td>High bandwidth</td> <td>Port classification to be used for virtual machines t...</td> </tr> <tr> <td>iSCSI workload</td> <td>Port classification for host iSCSI workloads.</td> </tr> <tr> <td>VDI Port Class</td> <td></td> </tr> </tbody> </table> <p>Below the table is a 'Create Port Classification...' button. At the bottom right are 'OK' and 'Cancel' buttons.</p>	Name	Description	SR-IOV	Port classification to be used for virtual machines t...	Host management	Port classification to be used for host managemen...	Network load balancing	Port classification to be used for virtual machines t...	Live migration workload	Port classification to be used for host live migratio...	Medium bandwidth	Port classification to be used for virtual machines t...	Host Cluster Workload	Port classification for host cluster workloads.	Low bandwidth	Port classification to be used for virtual machines t...	High bandwidth	Port classification to be used for virtual machines t...	iSCSI workload	Port classification for host iSCSI workloads.	VDI Port Class	
Name	Description																						
SR-IOV	Port classification to be used for virtual machines t...																						
Host management	Port classification to be used for host managemen...																						
Network load balancing	Port classification to be used for virtual machines t...																						
Live migration workload	Port classification to be used for host live migratio...																						
Medium bandwidth	Port classification to be used for virtual machines t...																						
Host Cluster Workload	Port classification for host cluster workloads.																						
Low bandwidth	Port classification to be used for virtual machines t...																						
High bandwidth	Port classification to be used for virtual machines t...																						
iSCSI workload	Port classification for host iSCSI workloads.																						
VDI Port Class																							
<p>In the Create Port Classification Wizard, provide a [VDIPortClassName] for the port classification.</p> <p>Click OK.</p>	 <p>The screenshot shows the 'Create Port Classification Wizard' dialog box. It contains the following fields:</p> <ul style="list-style-type: none"> Name: VDI Port Class (highlighted with a red box) Description: (empty text area) <p>At the bottom right are 'OK' and 'Cancel' buttons, with 'OK' highlighted by a red box and a mouse cursor.</p>																						

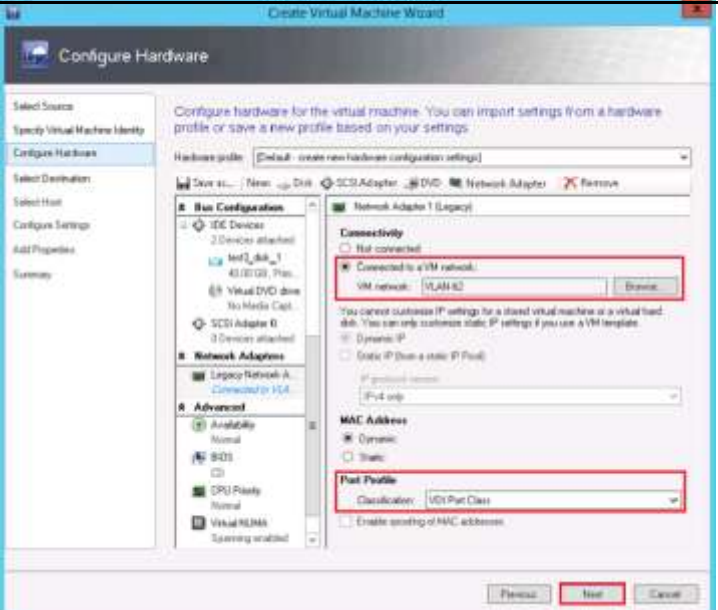
Instructions	Visual
<p>In the Summary tab, review the settings and click Finish to complete the wizard.</p>	
<p>Right-click Hyper-V host and select Properties from the context menu.</p>	

Instructions	Visual
<p>Select the Virtual Switches tab.</p> <p>Click New Virtual Switch.</p> <p>Select the New Logical Switch choice. (See next steps for adding a logical switch)</p> <p>Click OK when complete.</p>	
<p>Select the Cisco Nexus 1000 V logical switch if more than one exists.</p> <p>Under the Physical Adapters, select the Cisco VIC Ethernet adapter that is connected to the Uplink1 vNIC template.</p> <p>Note: The Cisco VIC Ethernet adapter number assigned to the Uplink1 vNIC template is random on each Hyper-V host, so verify the correct one by matching MAC address on the Hyper-V host with the MAC address assigned in the server profile.</p> <p>Click Add to add the Uplink2 vNIC adapter</p> <p>Click OK.</p>	

Instructions	Visual
<p>If prompted, click OK to continue with the changes.</p>	
<p>From the VMs and Services tab, right-click on VM Networks and then select Create VM Network from the context menu.</p>	

Instructions	Visual
<p>The Create VM Network Wizard starts.</p> <p>Provide a name.</p> <p>Select the Logical Network.</p>	 <p>The screenshot shows the 'Name' step of the 'Create VM Network Wizard'. The title bar says 'Create VM Network Wizard'. On the left, there are tabs for 'Name', 'Isolation', and 'Summary', with 'Name' being the active tab. The main area is titled 'Specify a name and description for the VM network'. It contains three input fields: 'Name' with the value 'VLAN-61', 'Description' which is empty, and 'Logical network' with a dropdown menu showing 'vDSNetwork1'. At the bottom right, there are three buttons: 'Previous', 'Next', and 'Cancel'. The 'Next' button is highlighted with a red box.</p>
<p>From the Isolation tab, select the “Specify an externally supplied VM network” radio button.</p> <p>Select the VLAN that matches with the VM network name.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Isolation' step of the 'Create VM Network Wizard'. The title bar says 'Create VM Network Wizard'. On the left, there are tabs for 'Name', 'Isolation', and 'Summary', with 'Isolation' being the active tab. The main area is titled 'Configure the isolation for this VM network, or select automatic to have it configured for you'. It contains three radio button options: 'Automatic', 'Specify an externally supplied VM network' (which is selected), and 'User defined'. Below the 'Specify an externally supplied VM network' option, there is a dropdown menu for 'External VM network' showing 'VLAN-61'. At the bottom right, there are three buttons: 'Previous', 'Next', and 'Cancel'. The 'Next' button is highlighted with a red box.</p>

Instructions	Visual										
<p>Review the Summary page before clicking Finish.</p> <p>Repeat for the remaining VLANs. When completed, all four should be configured.</p> <table border="1" data-bbox="186 506 764 789"> <thead> <tr> <th>Name</th><th>Logical Network</th></tr> </thead> <tbody> <tr> <td>VLAN-61</td><td>VDINetwork / VDINetwork2</td></tr> <tr> <td>VLAN-62</td><td>VDINetwork / VDINetwork2</td></tr> <tr> <td>VLAN-63</td><td>ClusterNetwork / ClusterNetwork2</td></tr> <tr> <td>VLAN-64</td><td>ClusterNetwork / ClusterNetwork2</td></tr> </tbody> </table>	Name	Logical Network	VLAN-61	VDINetwork / VDINetwork2	VLAN-62	VDINetwork / VDINetwork2	VLAN-63	ClusterNetwork / ClusterNetwork2	VLAN-64	ClusterNetwork / ClusterNetwork2	
Name	Logical Network										
VLAN-61	VDINetwork / VDINetwork2										
VLAN-62	VDINetwork / VDINetwork2										
VLAN-63	ClusterNetwork / ClusterNetwork2										
VLAN-64	ClusterNetwork / ClusterNetwork2										

<p>When creating a new virtual machine or template, to use the Nexus 1000V virtual switch, specify both the VM Network and the associated Port Profile in the VM properties.</p>	
--	---

6.14. Installing and Configuring Citrix XenDesktop

This section lists supported configurations for the Delivery Controller, Citrix Studio, Citrix Director, and Virtual Delivery Agents (VDAs) at the time of release. System requirements for other features and components (such as StoreFront, host systems, HDX, Receivers and plug-ins, Desktop Lock, and Provisioning Services) are described in their respective documentation.

Unless otherwise noted, the installer deploys component prerequisites automatically. This includes .NET 3.5 SP1, which is required when installing components on Windows 7 or Windows Server 2008 R2 systems. The installation media contains several third-party components. Before using the Citrix software, check for security updates from the third party, and install them.

If you install all the core components (Controller, Studio, Director, StoreFront, and Licensing) on a single server, you need a minimum of 3 GB of RAM; more is recommended. Performance will vary depending on your exact configuration, including the number of users, applications, desktops, and other factors.

Delivery Controller

Supported operating systems:

- Windows Server 2012.
- Windows Server 2008 R2, Standard and Enterprise Editions.

Requirements:

- Disk space: 100 MB.
- Microsoft .NET Framework 3.5 SP1 (required on Windows Server 2008 R2 only).
- Microsoft .NET 4.0.
- Windows PowerShell 2.0 (included with Windows Server 2008 R2) or 3.0 (included with Windows Server 2012).
- Visual C++ 2005, 2008 SP1, and 2010 Redistributable packages. The installer deploys these automatically. They are also available on the Citrix installation media in the Support folder.

Database

Supported Microsoft SQL Server versions for the Site Configuration Database (which initially includes the Configuration Logging Database and the Monitoring Database):

- SQL Server 2012 SP1 - Express, Standard, and Enterprise Editions. By default, the Express edition is installed when installing the Controller.
- SQL Server 2008 R2 SP2 - Express, Standard, Enterprise, and Datacenter Editions.

The following database features are supported (except on SQL Server Express Edition, which supports only standalone mode):

- SQL Server Clustered Instances
- SQL Server Mirroring
- SQL Server Replication
- SQL Server 2012 AlwaysOn Availability Groups

Note; Windows authentication is required for connections between the Controller and the SQL Server database.

Studio

Supported operating systems:

- Windows 8 Professional and Enterprise Editions.
- Windows 7 Professional and Enterprise Editions.
- Windows Server 2012 Essentials, Standard, and Enterprise Editions.
- Windows Server 2008 R2 SP1 Foundation, Standard, Enterprise, and Datacenter Editions.

Requirements:

- Disk space: 75 MB.
- Microsoft .NET Framework 3.5 SP1(required on Windows Server 2008 R2 only).
- Microsoft Management Console 3.0 (included with all supported operating systems).
- Windows PowerShell 2.0 (included with Windows 7 and Windows Server 2008 R2) or 3.0 (included with Windows 8 and Windows Server 2012).

Director

Supported operating systems: Windows Server 2012, Standard Edition or Windows Server 2008 R2 SP1, Standard and Enterprise Editions

Requirements:

- Disk space: 50 MB.
- Microsoft .NET Framework 4.0. The installer deploys this automatically. It is also available on the Citrix installation media in the Support\DotNet4 folder.
- Microsoft Internet Information Services (IIS) 7.0 and ASP.NET 2.0. If these are not already installed, you are prompted for the Windows Server installation media, then they are installed for you.
- Supported browsers for viewing Director: Internet Explorer 9 and 10, Firefox, Chrome

Virtual Delivery Agent (VDA) for Windows Desktop OS

Supported operating systems:

- Windows 8.
- Windows 7 SP1, Enterprise and Professional Editions.

The installer automatically deploys requirements such as the Microsoft .NET Framework and the Visual C++ Runtime Library. These components are also available on the Citrix installation media in the Support folders.

Remote PC Access uses this VDA, which you install on physical office PCs.

Several multimedia acceleration features (such as HDX MediaStream Windows Media Redirection) require that Microsoft Media Foundation be installed on the machine where you install the VDA. If the machine does not have Media Foundation installed, the multimedia acceleration features will not be installed and will not work. Do not remove Media Foundation from the machine after installing the Citrix software; otherwise, users will not be able to log on to the machine. On most Windows 7 and Windows 8 editions, the Media Foundation support is already installed and cannot be removed. However, N editions do not include certain media-related technologies; you can obtain that software from Microsoft or a third party.

You cannot install a VDA supplied with this release on a machine running Windows XP or Windows Vista; however, you can install an earlier Virtual Desktop Agent version on those operating systems, if needed. For more information, see [Install an earlier Virtual Desktop Agent on Windows XP or Windows Vista](#). Remote PC Access is not supported on Windows Vista operating systems.

Virtual Delivery Agent (VDA) for Windows Server OS

Supported operating systems:

- Windows Server 2012.
- Windows Server 2008 R2 SP1; Datacenter, Enterprise, and Standard Editions.

The installer automatically deploys requirements such as the Microsoft .NET Framework and the Visual C++ Runtime Library. These components are also available on the Citrix installation media in the Support folders.

The installer automatically enables the Remote Desktop Services role.

Several multimedia acceleration features (such as HDX MediaStream Windows Media Redirection) require that the Microsoft Media Foundation be installed on the machine where you install the VDA. If the machine does not have Media Foundation installed, the multimedia acceleration features will not be installed and will not work. Do not remove Media Foundation from the machine after installing the Citrix software; otherwise, users will not be able to log on to the machine. On most Windows Server 2008 R2 and Windows Server 2012 editions, the Media Foundation feature is installed through the Server Manager (for Windows Server 2012: ServerMediaFoundation, for Windows Server 2008 R2: DesktopExperience). However, N editions do not include certain media-related technologies; you can obtain that software from Microsoft or a third party.

The Print Spooler Service is enabled by default on the Windows server. If you disable this service, you cannot successfully install a VDA for Windows Server OS. Therefore, make sure that this service is enabled before installing a VDA.

Host

Supported hypervisors are XenServer, VMware vSphere, and Microsoft Hyper-V. For this Cisco Validated Design, we used Microsoft Windows Server 2012 SP1 with Hyper-V as the hypervisor.

Other

StoreFront requires 2 GB of disk space. See the StoreFront documentation for full system requirements.

Citrix License Server requires 40 MB of disk space. See the licensing documentation for full system requirements.

Universal Print Server - Supported operating systems (the Controller includes the Universal Print Server functionality; you need only install the Universal Print Server on your print servers):

- Windows Server 2008 R2 SP1
- Windows Server 2008 32-bit

The Microsoft Group Policy Management Console (GPMC) is required if you store Citrix policy information in Active Directory rather than the Site Configuration Database. For more information, see the Microsoft documentation.

You can install the Receiver for Windows when installing a VDA provided on the XenDesktop installation media. For system requirements information on other platforms, see the Receiver for Windows documentation.


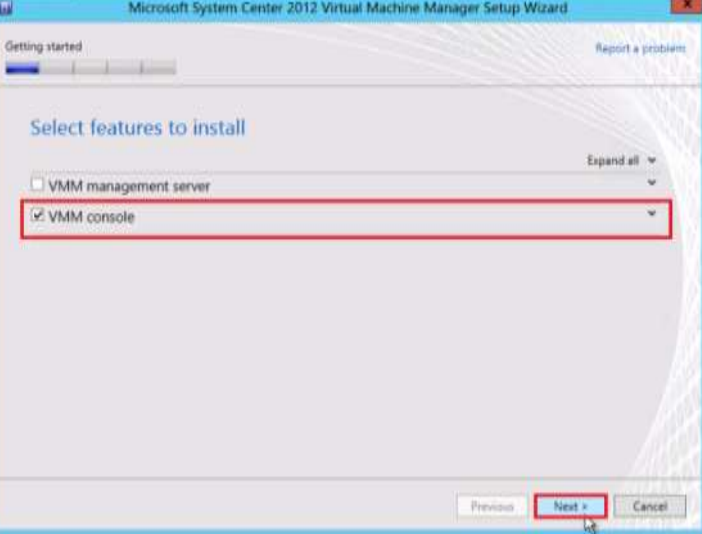
The Receiver for Linux and the Receiver for Mac are provided on the XenDesktop installation media. See their documentation for system requirements.

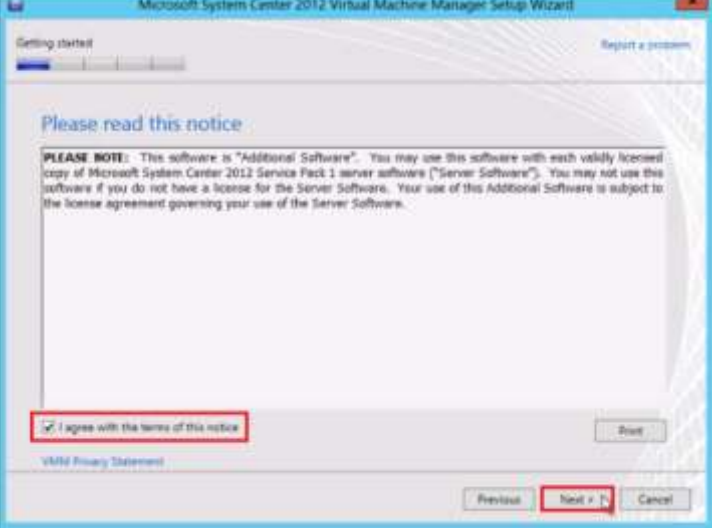
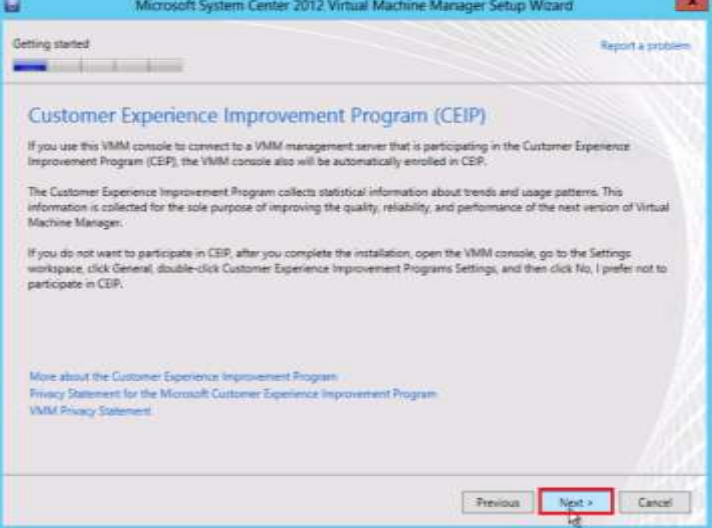
6.14.1. Installing Provisioning Services

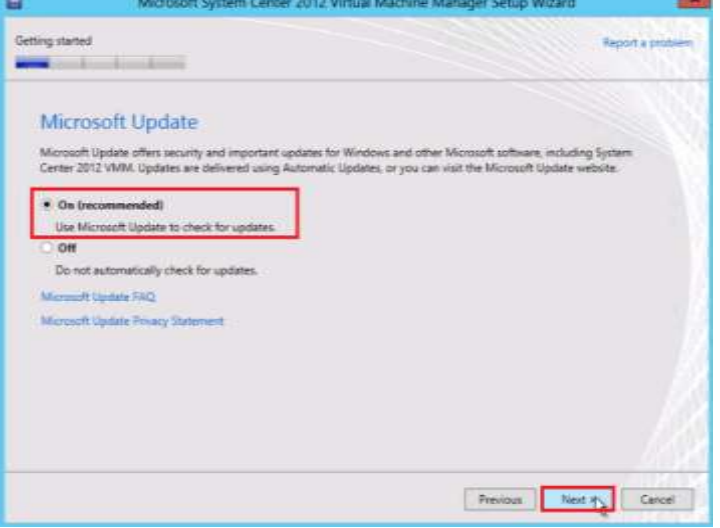
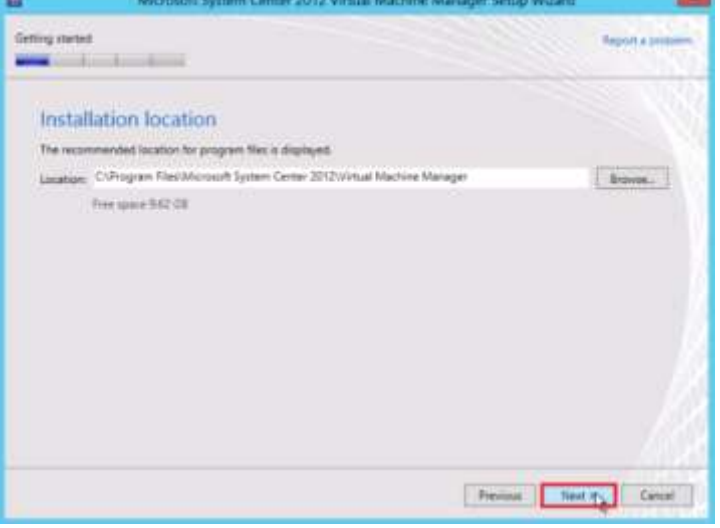
In most implementations, there is a single vDisk providing the standard image for multiple target devices. Thousands of target devices can use a single vDisk shared across multiple Provisioning Services (PVS) servers in the same farm, simplifying virtual desktop management. This section describes the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

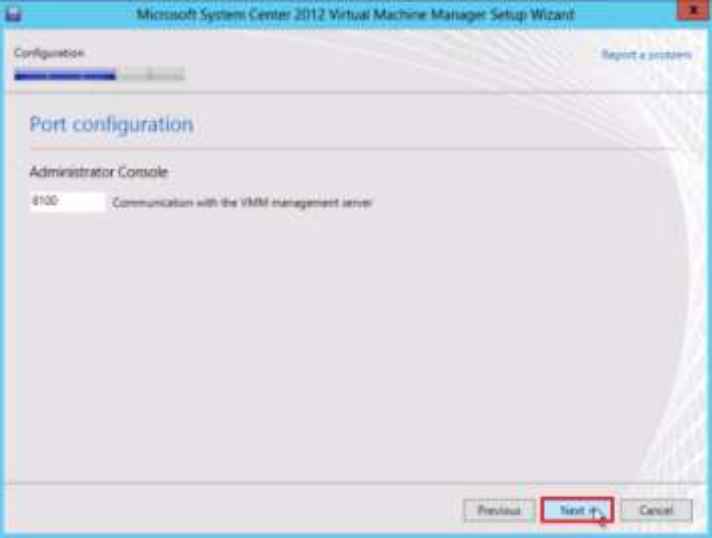
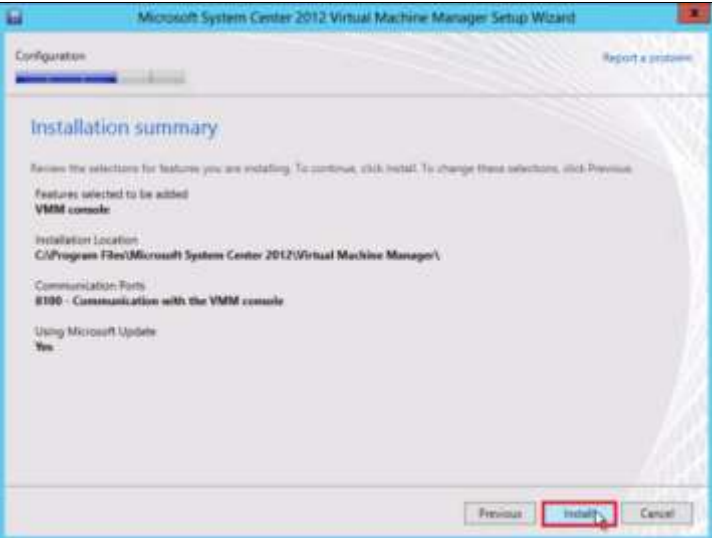
The PVS server can have many stored vDisks, and each vDisk can be several gigabytes in size. Your streaming performance can be improved using a RAID array, SAN, or NAS. PVS software and hardware requirements are available at <http://support.citrix.com>.

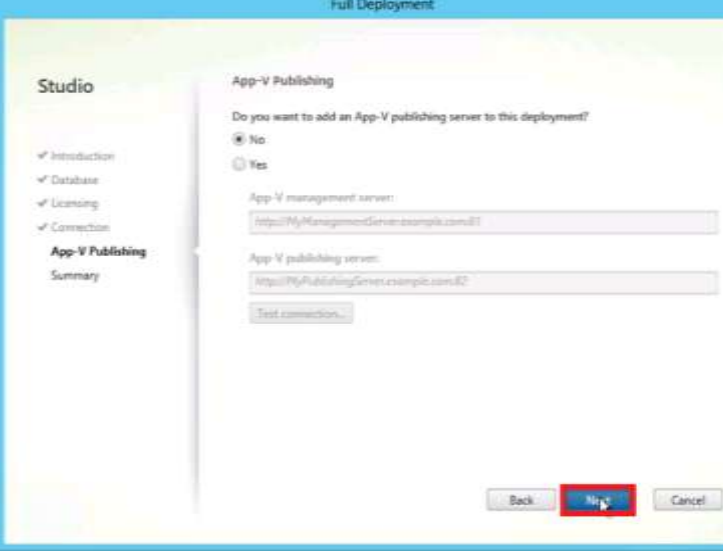
Microsoft System Center 2012 Virtual Machine Manager must be installed if the XenDesktop Setup Wizard will be used for creating the virtual machines. After inserting the System Center Virtual Machine Manager ISO, you should get a welcome screen.

Instructions	Visual
<p>Select the Install link.</p> <p>The installer will copy some files over to the host then present the feature screen.</p>	
<p>Select the VMM Console for install.</p> <p>Click Next.</p>	

Instructions	Visual
<p>Read the notice.</p> <p>I acceptable, enable the checkbox labeled “I agree with the terms of this notice.”</p> <p>Click Next.</p>	
<p>On the Customer Experience Improvement Program screen, click Next .</p>	

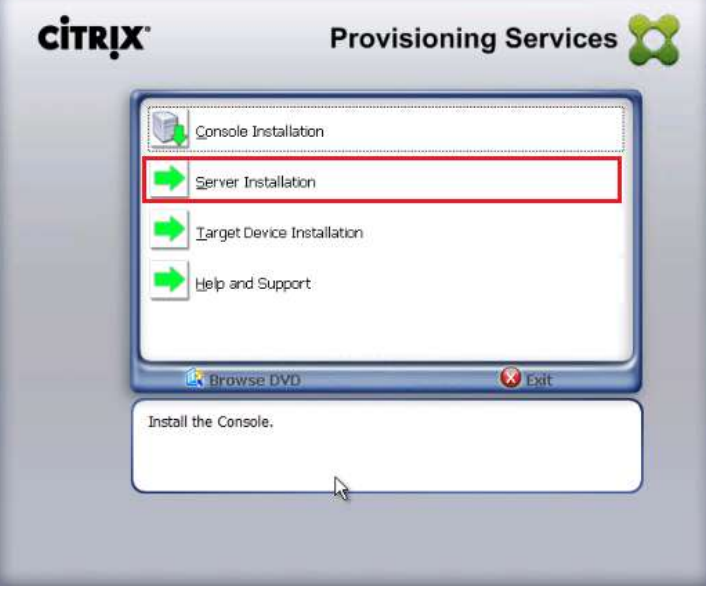
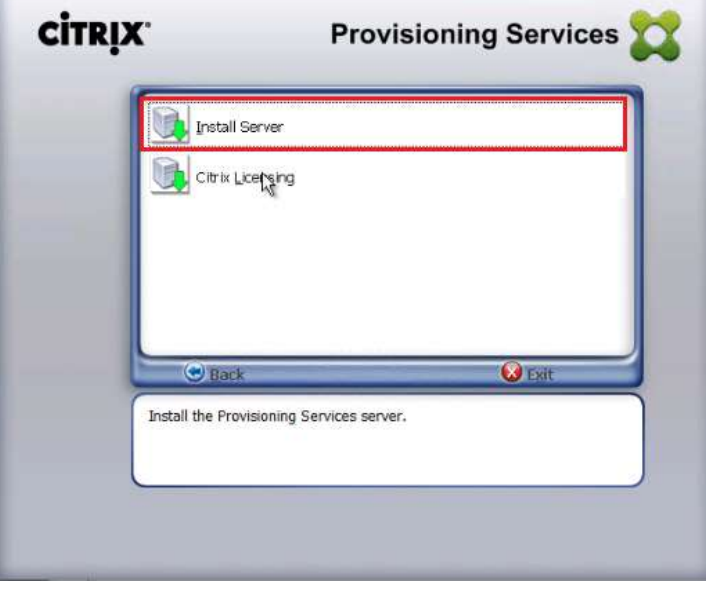
Instructions	Visual
<p>On the Microsoft Update screen enable the On (recommended) radio button.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Microsoft Update' screen of the setup wizard. It includes a progress bar at the top, a 'Report a problem' link, and a description of Microsoft Update. Two radio buttons are present: 'On (recommended)' which is selected and highlighted with a red box, and 'Off'. Below the buttons are links for 'Microsoft Update FAQ' and 'Microsoft Update Privacy Statement'. At the bottom, the 'Next' button is highlighted with a red box, along with 'Previous' and 'Cancel' buttons.</p>
<p>Accept the default installation location.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Installation location' screen of the setup wizard. It displays the recommended location for program files: 'C:\Program Files\Microsoft System Center 2012\Virtual Machine Manager'. There is a 'Browse...' button next to the location field. At the bottom, the 'Next' button is highlighted with a red box, along with 'Previous' and 'Cancel' buttons.</p>

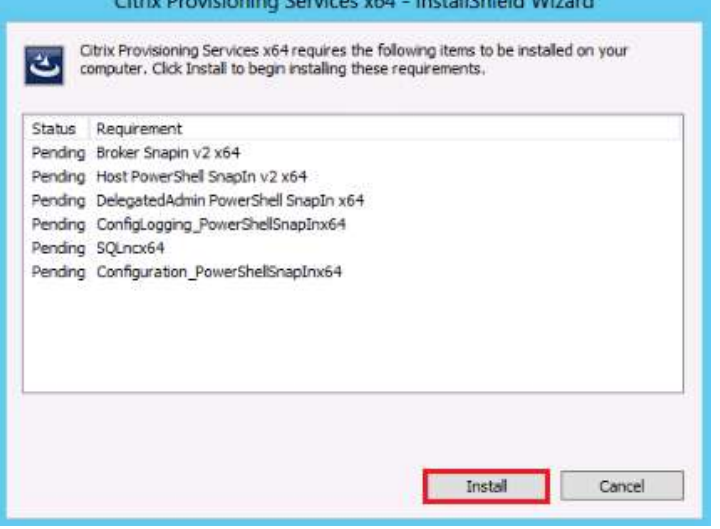
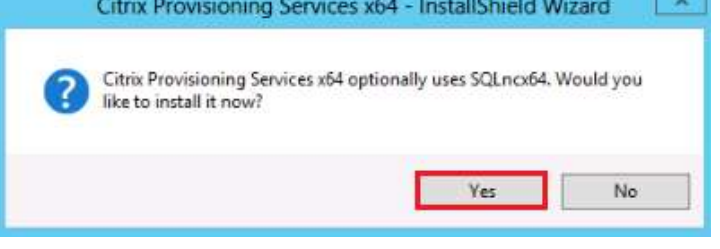
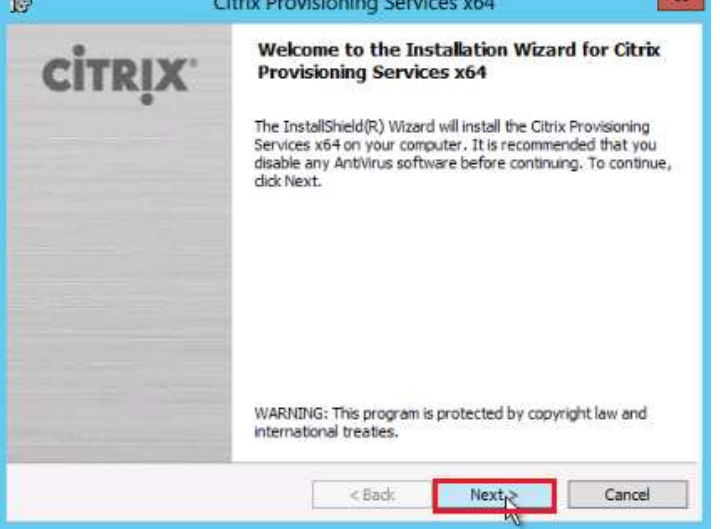
Instructions	Visual
<p>Accept the default communication port.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Port configuration' step of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The 'Administrator Console' section displays a text box with the value '8100' and the label 'Communication with the VMM management server'. The 'Next' button at the bottom right is highlighted with a red box.</p>
<p>Review the installation summary.</p> <p>Click Install.</p>	 <p>The screenshot shows the 'Installation summary' step of the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard. The summary lists the features to be installed, the installation location (C:\Program Files\Microsoft System Center 2012\Virtual Machine Manager\), the communication port (8100), and the option to use Microsoft Update (Yes). The 'Install' button at the bottom right is highlighted with a red box.</p>


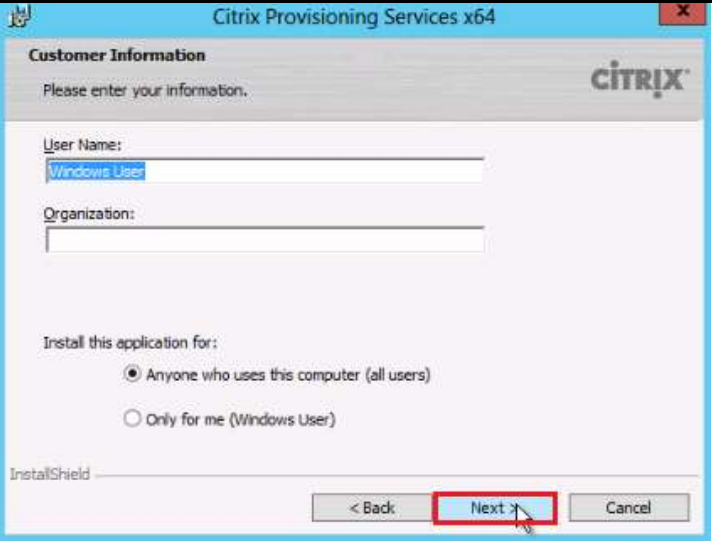
Instructions	Visual
<p>When installation is finished, disable the checkbox labeled Open the VMM console when this wizard closes.</p> <p>Click Close.</p> <p>Reboot the server before starting the XenDesktop configuration or the SCVMM configuration will fail.</p>	

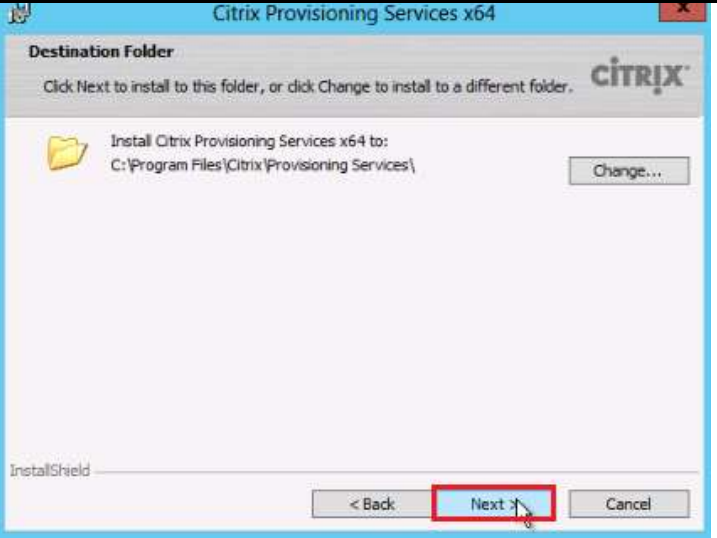
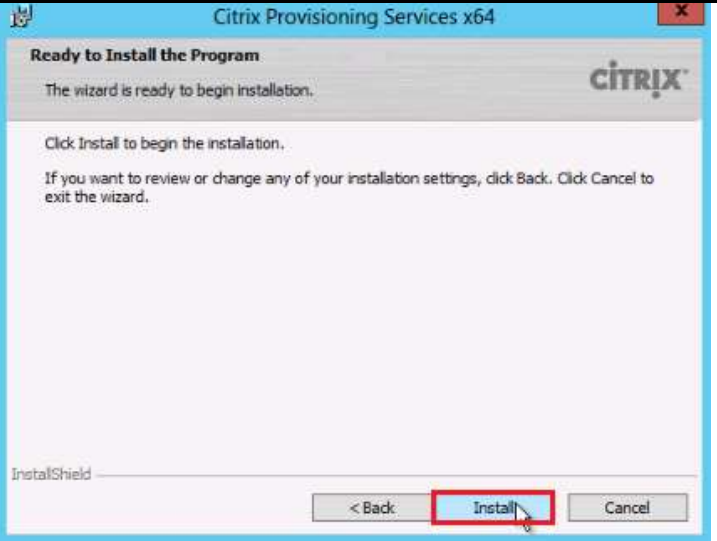
Only one MS SQL database is associated with a farm. You can choose to install the Provisioning Services database software on an existing SQL database, if that machine can communicate with all Provisioning Servers within the farm, or with a new SQL Express database machine, created using the SQL Express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

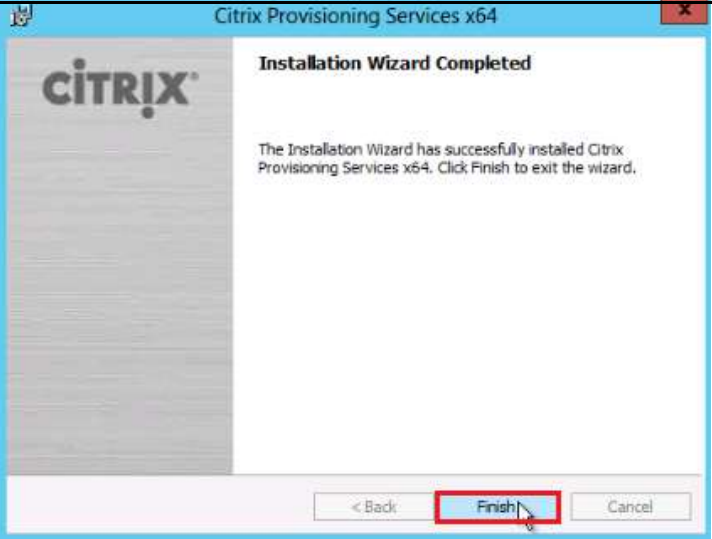

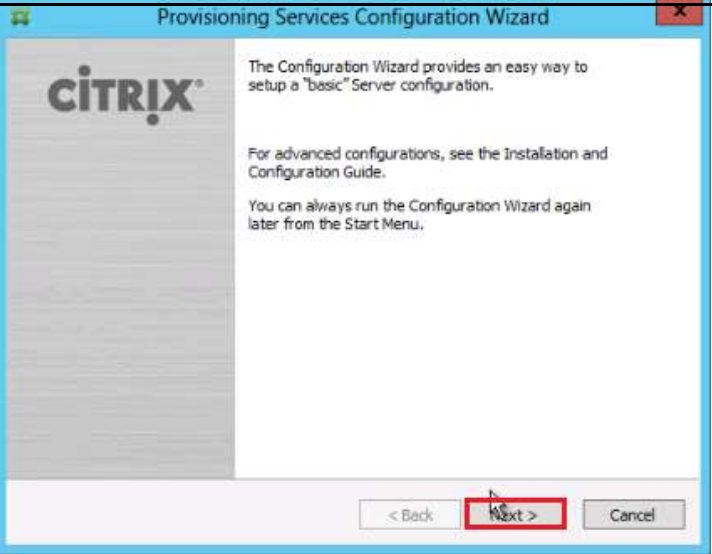
The following MS SQL 2008, MS SQL 2008 R2, and MS SQL 2012 Server (32 or 64-bit editions) databases can be used for the Provisioning Services database: SQL Server Express Edition, SQL Server Workgroup Edition, SQL Server Standard Edition, SQL Server Enterprise Edition. The SQL database for this Cisco Validated Design was created during the Site configuration process described previously.

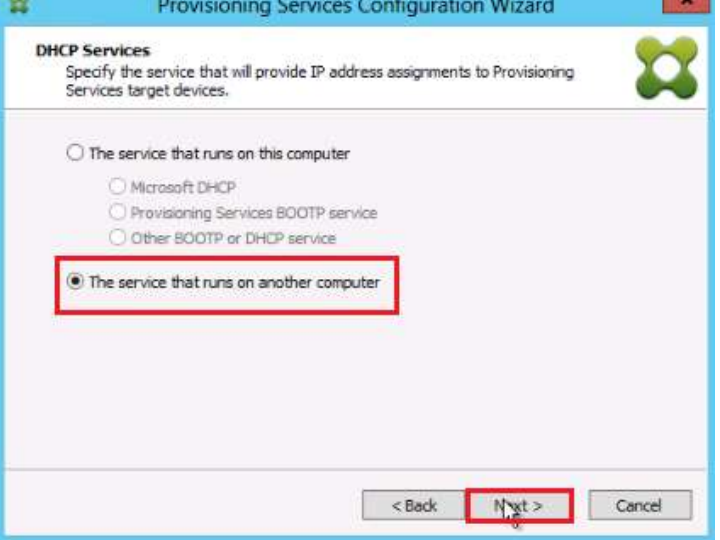
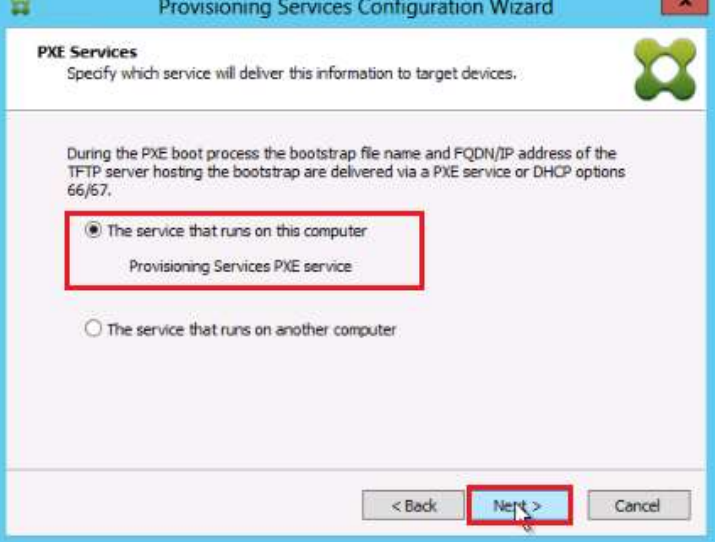
Instructions	Visual
<p>Insert the Provisioning Services ISO and let AutoRun launch the installer.</p> <p>Click Server Installation</p>	
<p>Click Install Server.</p> <p>The installation wizard will check to resolve dependencies and then begin the PVS server installation process. It is recommended that you disable anti-virus software prior to the installation.</p>	

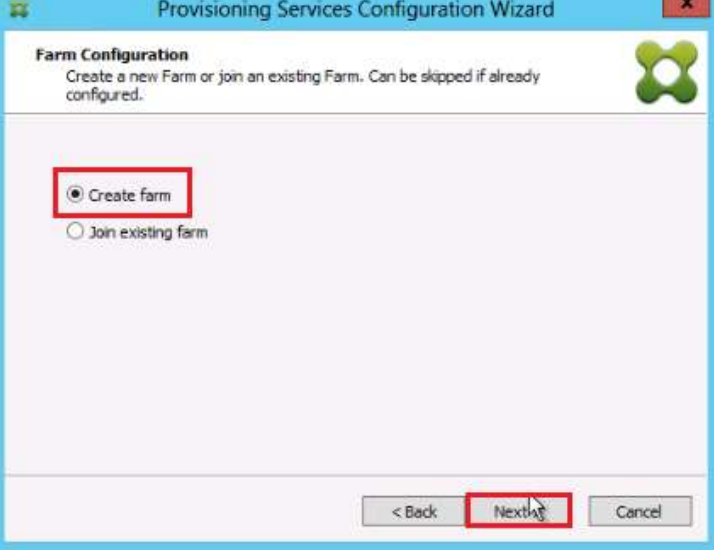
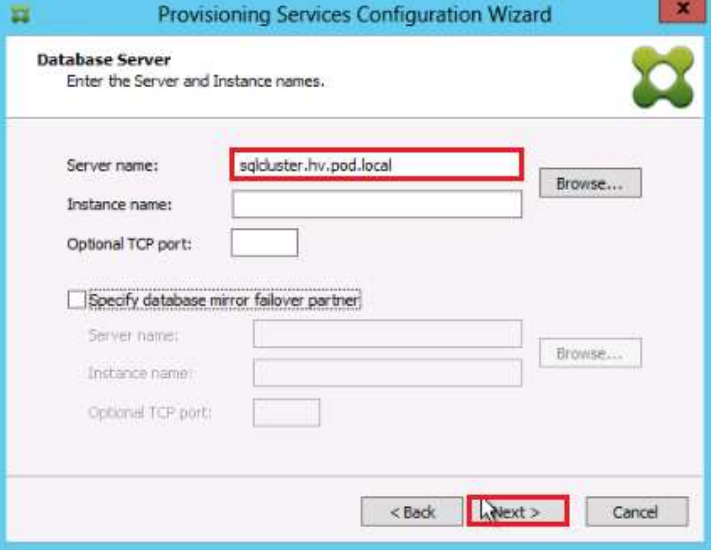
Instructions	Visual
Click Install on the PreRequisites dialog.	 <p>The screenshot shows the 'Citrix Provisioning Services x64 - InstallShield Wizard' window. It lists several requirements that are 'Pending': Broker Snapin v2 x64, Host PowerShell SnapIn v2 x64, DelegatedAdmin PowerShell SnapIn x64, ConfigLogging_PowerShellSnapInx64, SQLncx64, and Configuration_PowerShellSnapInx64. The 'Install' button at the bottom right is highlighted with a red rectangle.</p>
Click Yes when prompted to install the SQL Native Client.	 <p>The screenshot shows a dialog box asking: 'Citrix Provisioning Services x64 optionally uses SQLNcx64. Would you like to install it now?'. The 'Yes' button is highlighted with a red rectangle.</p>
Click Next when the Installation wizard starts.	 <p>The screenshot shows the 'Welcome to the Installation Wizard for Citrix Provisioning Services x64' window. It includes the Citrix logo and a warning about disabling antivirus software. The 'Next >' button at the bottom right is highlighted with a red rectangle.</p>

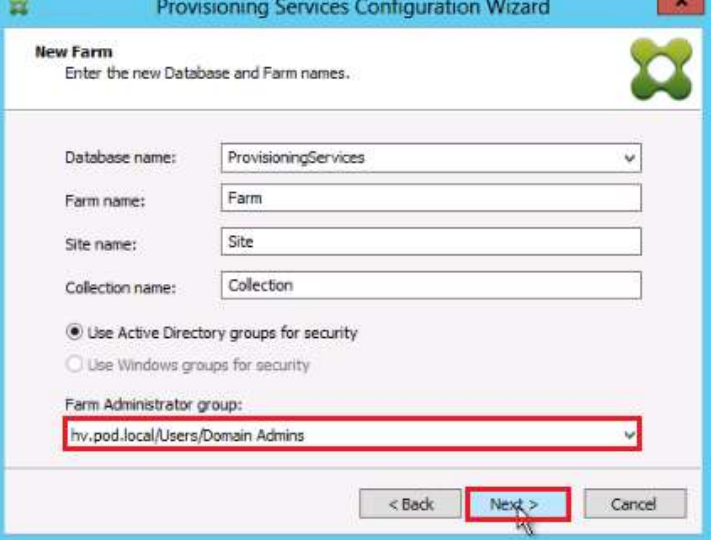
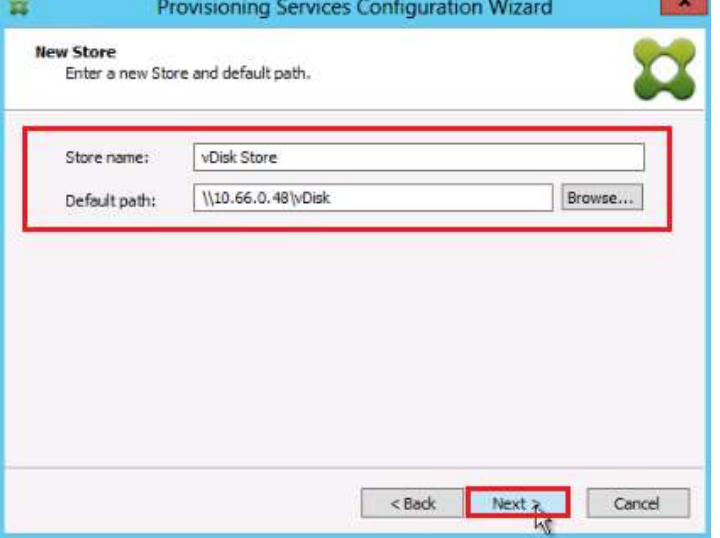
Instructions	Visual
<p>Review the license agreement terms.</p> <p>If acceptable, select the radio button labeled “I accept the terms in the license agreement.”</p> <p>Click Next.</p>	
<p>Provide User Name, and Organization information.</p> <p>Select who will see the application.</p> <p>Click Next.</p>	

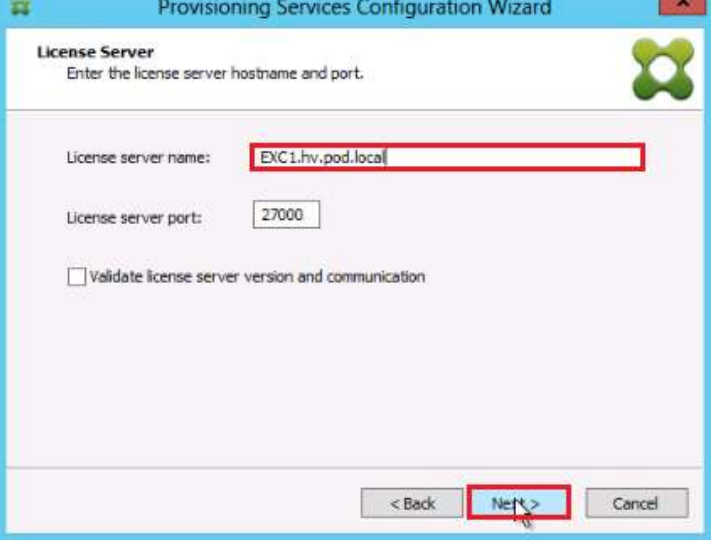
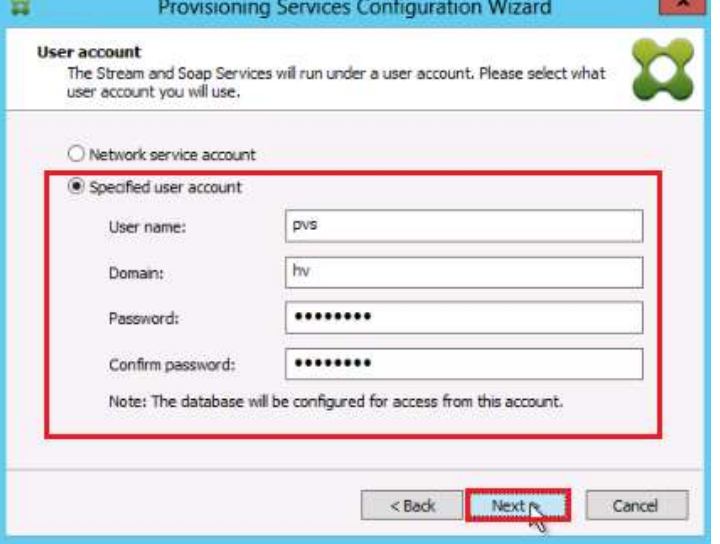
Instructions	Visual
<p>Accept the default installation location.</p> <p>Click Next.</p>	
<p>Click Install to begin the installation.</p>	

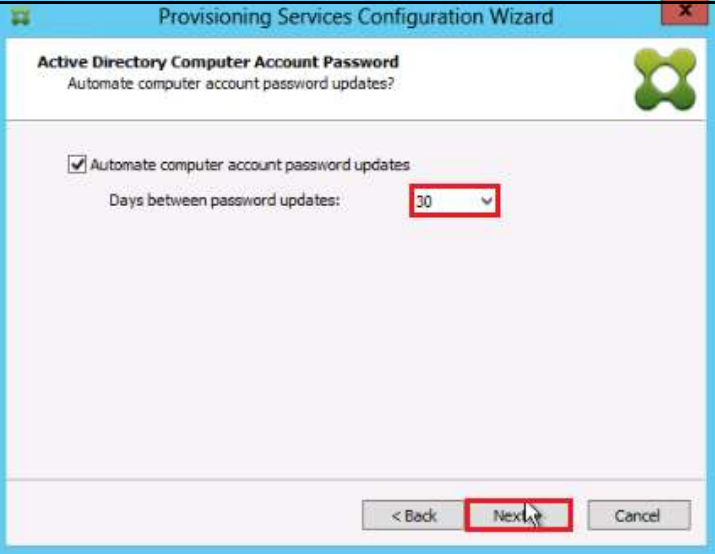
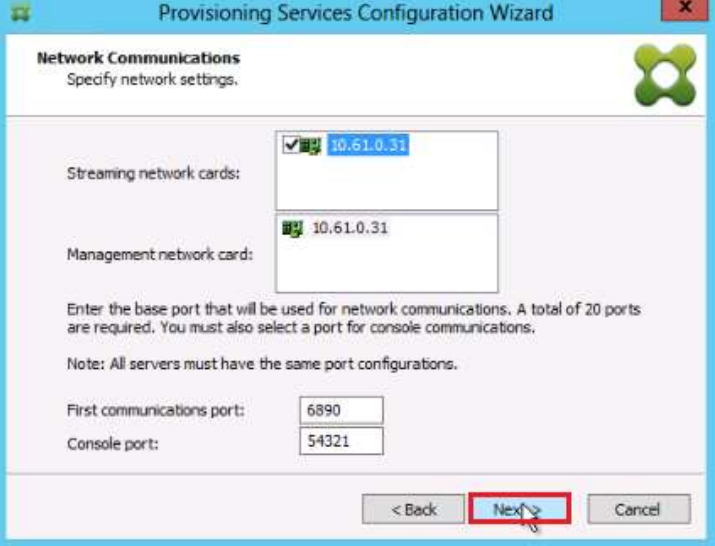
Instructions	Visual
Click Finish when the install is complete.	 <p>The screenshot shows the 'Citrix Provisioning Services x64' window with the title 'Installation Wizard Completed'. The text inside states: 'The Installation Wizard has successfully installed Citrix Provisioning Services x64. Click Finish to exit the wizard.' At the bottom, there are three buttons: '< Back', 'Finish', and 'Cancel'. The 'Finish' button is highlighted with a red rectangle, and a mouse cursor is pointing at it.</p>
Click OK to acknowledge the PVS console has not yet been installed.	 <p>The screenshot shows the 'Citrix Provisioning Services x64' window with a yellow warning icon. The text inside says: 'The PVS Console is not detected in your system. You will need the Console to log into your PVS Farm from this system. Please install it.' At the bottom, there is an 'OK' button highlighted with a red rectangle, and a mouse cursor is pointing at it.</p>
<p>The PVS Configuration Wizard starts automatically.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The text inside 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, and a mouse cursor is pointing at it.</p>

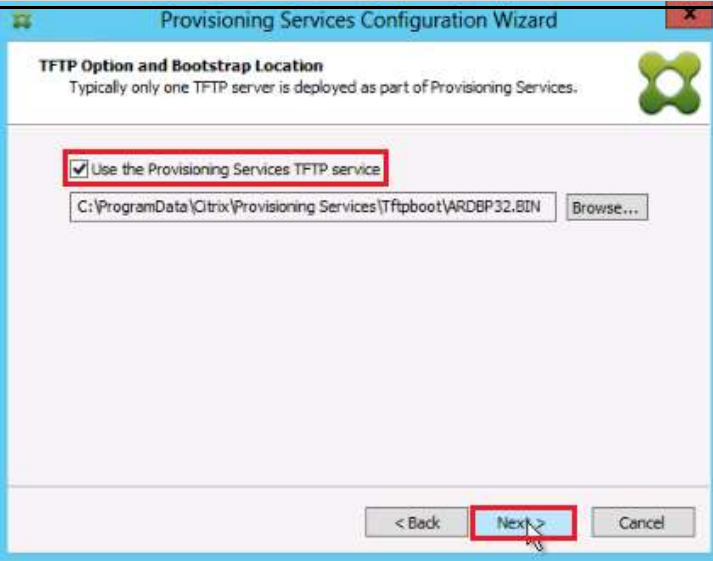
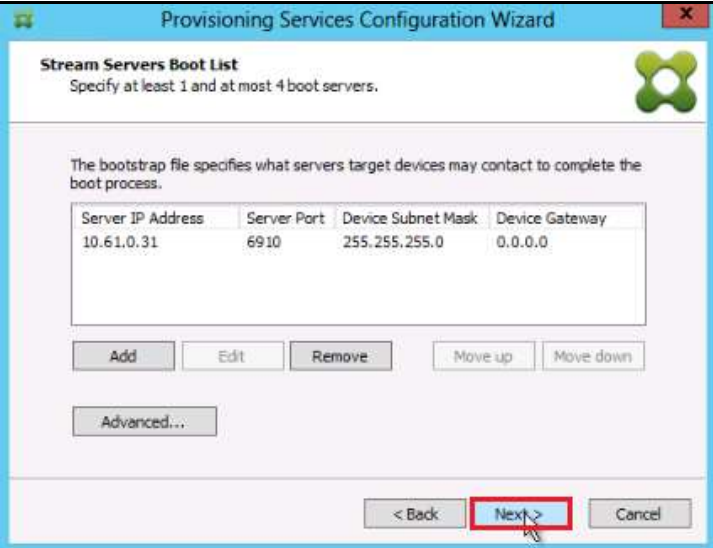
Instructions	Visual
<p>Since the PVS server is not the DHCP server for the environment, select the radio button labeled, “The service that runs on another computer.”</p> <p>Click Next.</p>	
<p>Since this server will be a PXE server, select the radio button labeled, “The service that runs on this computer.”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

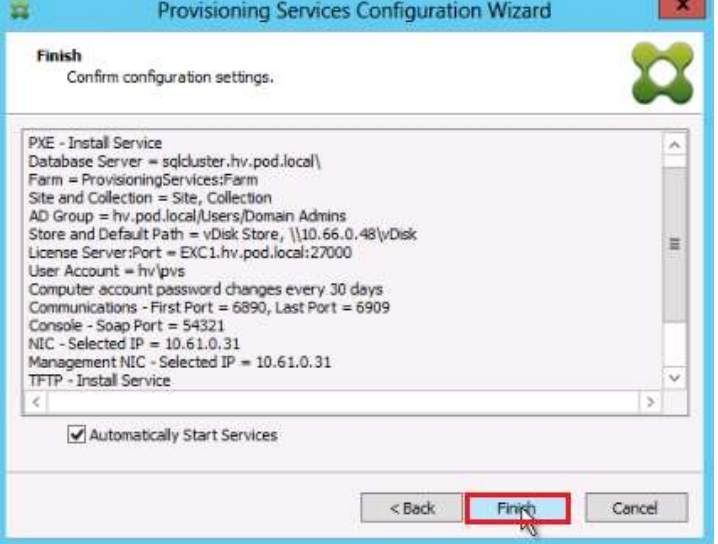
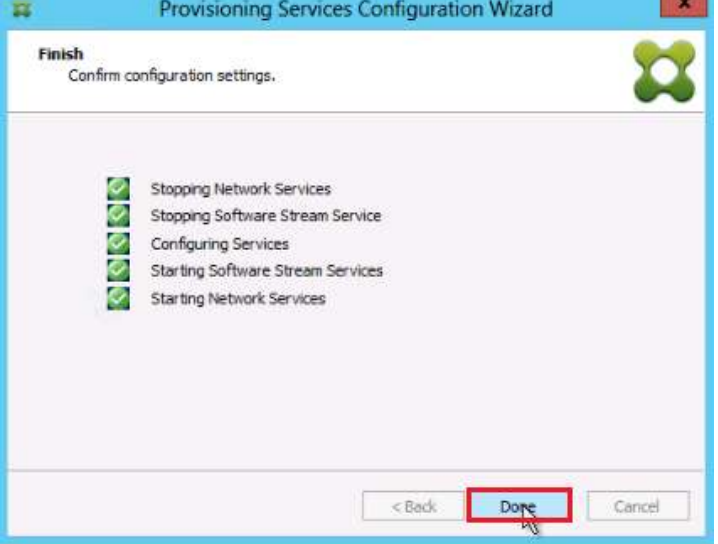
Instructions	Visual
<p>Since this is the first server in the farm, select the radio button labeled, “Create farm”.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The title bar says 'Provisioning Services Configuration Wizard'. The main heading is 'Farm Configuration'. Below it, it says 'Create a new Farm or join an existing Farm. Can be skipped if already configured.' There are two radio buttons: 'Create farm' (which is selected and highlighted with a red box) and 'Join existing farm'. At the bottom right, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>
<p>Enter the listener name of the SQL server AlwaysOn group that hosts the database as the Server name.</p> <p>Note: If using a cluster, instead of AlwaysOn groups, you will need to supply the instance name as well.</p> <p>Click Next</p>	 <p>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The title bar says 'Provisioning Services Configuration Wizard'. The main heading is 'Database Server'. Below it, it says 'Enter the Server and Instance names.' There are three text boxes: 'Server name:' (containing 'sqlcluster.hv.pod.local' and highlighted with a red box), 'Instance name:', and 'Optional TCP port:'. To the right of the 'Server name' box is a 'Browse...' button. Below these is a checkbox labeled 'Specify database mirror failover partner' which is unchecked. Under this checkbox are three more text boxes: 'Server name:', 'Instance name:', and 'Optional TCP port:'. To the right of the 'Server name' box under the checkbox is another 'Browse...' button. At the bottom right, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>

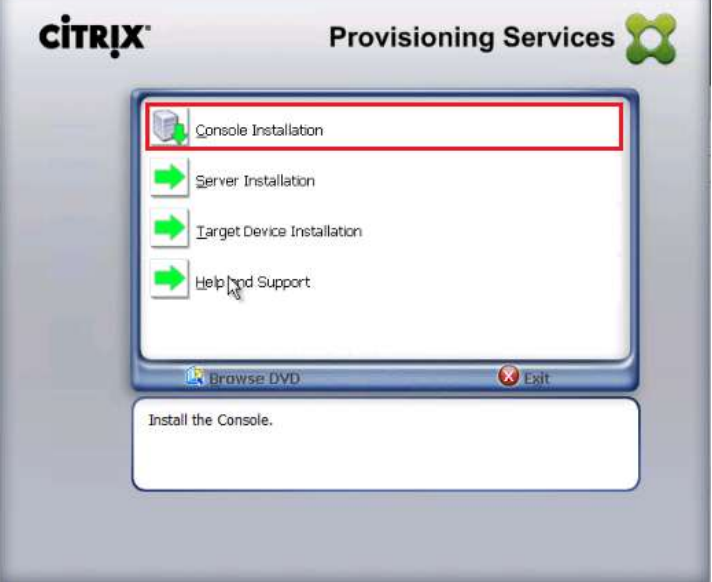

Instructions	Visual
<p>Optionally provide a Database name, Farm name, Site name, and Collection name for the PVS farm.</p> <p>Select the Domain Admins group for the Farm Administrator group.</p> <p>Click Next.</p>	
<p>Provide a vDisk Store name and the storage path the EMCVNxe vDisk share.</p> <p>Click Next.</p>	

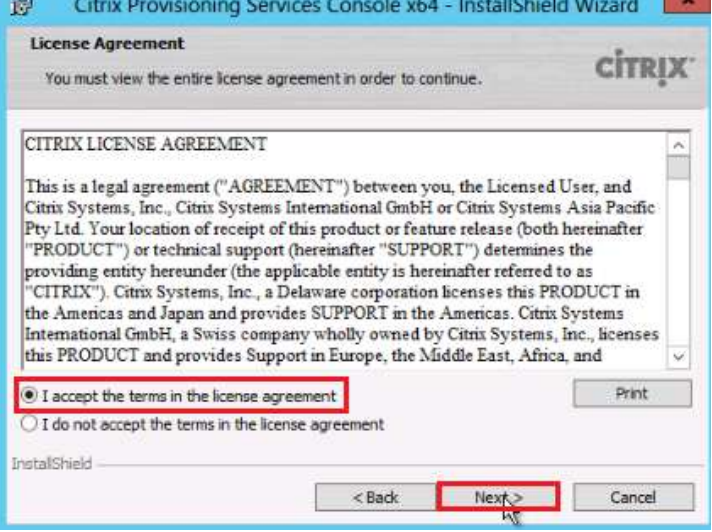
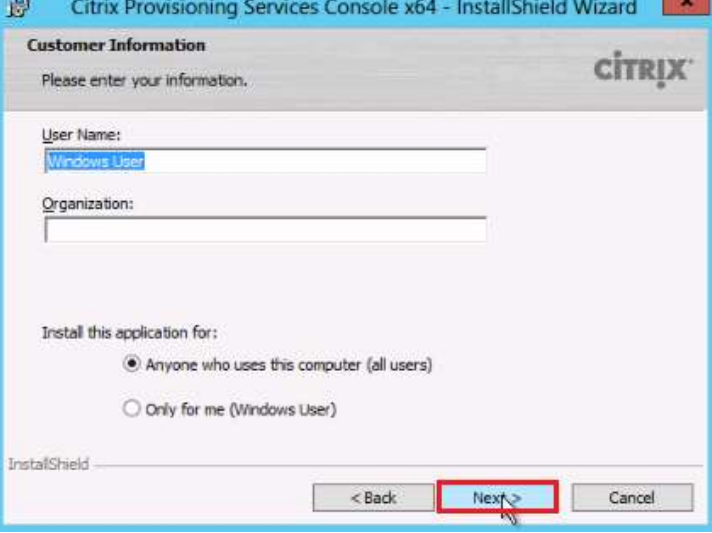
Instructions	Visual
<p>Provide the FQDN of the License Server, in this case it will be the first delivery controller.</p> <p>Optionally, provide a port number.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The title bar says 'Provisioning Services Configuration Wizard'. The main heading is 'License Server' with the instruction 'Enter the license server hostname and port.' Below this, there are two input fields: 'License server name' with the text 'EXC1.hv.pod.local' and 'License server port' with the text '27000'. There is a checkbox labeled 'Validate license server version and communication' which is unchecked. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red rectangular box.</p>
<p>If an active directory account is not already setup for the PVS servers, create that account (which only needs to be a domain member) prior to clicking Next on this dialog.</p> <p>Select the Specified user account radio button.</p> <p>Complete the User name, Domain, Password, and Confirm password fields, using the PVS account information created earlier.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The title bar says 'Provisioning Services Configuration Wizard'. The main heading is 'User account' with the instruction 'The Stream and Soap Services will run under a user account. Please select what user account you will use.' Below this, there are two radio buttons: 'Network service account' (unselected) and 'Specified user account' (selected). Under the 'Specified user account' section, there are four input fields: 'User name' with the text 'pvs', 'Domain' with the text 'hv', 'Password' with masked characters '.....', and 'Confirm password' with masked characters '.....'. Below these fields is a note: 'Note: The database will be configured for access from this account.' At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red rectangular box.</p>

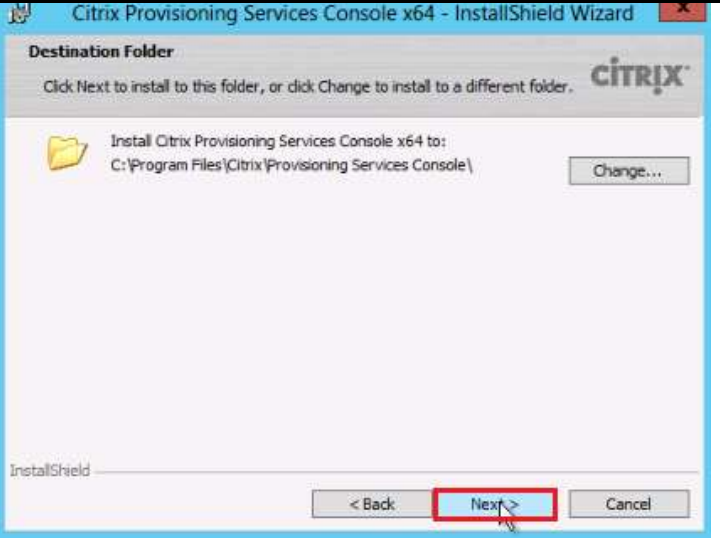
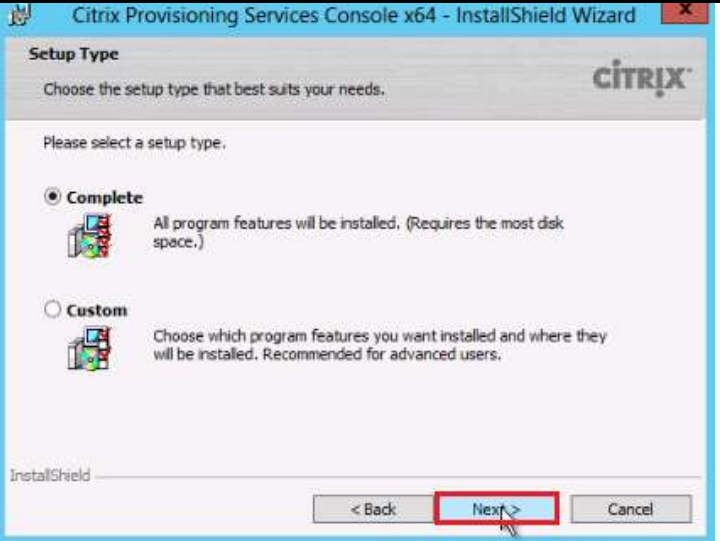
Instructions	Visual
<p>Set the Days between password updates to 30.</p> <p>Click Next.</p>	
<p>Keep the defaults for the network cards.</p> <p>Click Next.</p>	

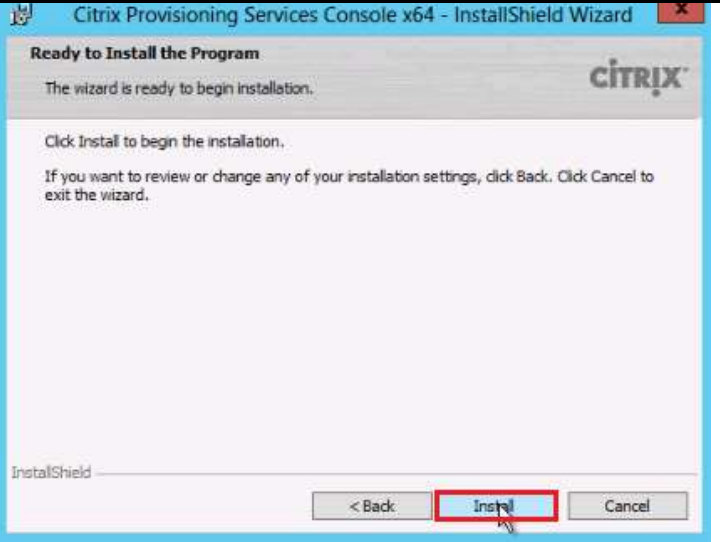

Instructions	Visual								
<p>Enable the Use the Provisioning Services TFTP service checkbox.</p> <p>Click Next.</p>	 <p>Provisioning Services Configuration Wizard</p> <p>TFTP Option and Bootstrap Location Typically only one TFTP server is deployed as part of Provisioning Services.</p> <p><input checked="" type="checkbox"/> Use the Provisioning Services TFTP service</p> <p>C:\ProgramData\Citrix\Provisioning Services\Tftpboot\ARDBP32.BIN Browse...</p> <p>< Back Next > Cancel</p>								
<p>Accept the default Stream Servers boot list.</p> <p>Click Next.</p>	 <p>Provisioning Services Configuration Wizard</p> <p>Stream Servers Boot List Specify at least 1 and at most 4 boot servers.</p> <p>The bootstrap file specifies what servers target devices may contact to complete the boot process.</p> <table><thead><tr><th>Server IP Address</th><th>Server Port</th><th>Device Subnet Mask</th><th>Device Gateway</th></tr></thead><tbody><tr><td>10.61.0.31</td><td>6910</td><td>255.255.255.0</td><td>0.0.0.0</td></tr></tbody></table> <p>Add Edit Remove Move up Move down</p> <p>Advanced...</p> <p>< Back Next > Cancel</p>	Server IP Address	Server Port	Device Subnet Mask	Device Gateway	10.61.0.31	6910	255.255.255.0	0.0.0.0
Server IP Address	Server Port	Device Subnet Mask	Device Gateway						
10.61.0.31	6910	255.255.255.0	0.0.0.0						

Instructions	Visual
<p>Click Finish to start the installation.</p>	
<p>When the installation is completed, click the Done button.</p>	

Instructions	Visual
<p>From the main installation screen.</p> <p>Click Console Installation .</p>	 <p>The screenshot shows the 'Citrix Provisioning Services' main window. It features a list of installation options: 'Console Installation' (highlighted with a red rectangle), 'Server Installation', 'Target Device Installation', and 'Help and Support'. Below the list are 'Browse DVD' and 'Exit' buttons. A status bar at the bottom says 'Install the Console.'</p>
<p>Click Next.</p>	 <p>The screenshot shows the 'Citrix Provisioning Services Console x64 - InstallShield Wizard' window. It includes the Citrix logo, a welcome message, and instructions to click 'Next'. A warning message is also present. At the bottom, the 'Next >' button is highlighted with a red rectangle, along with '< Back' and 'Cancel' buttons.</p>

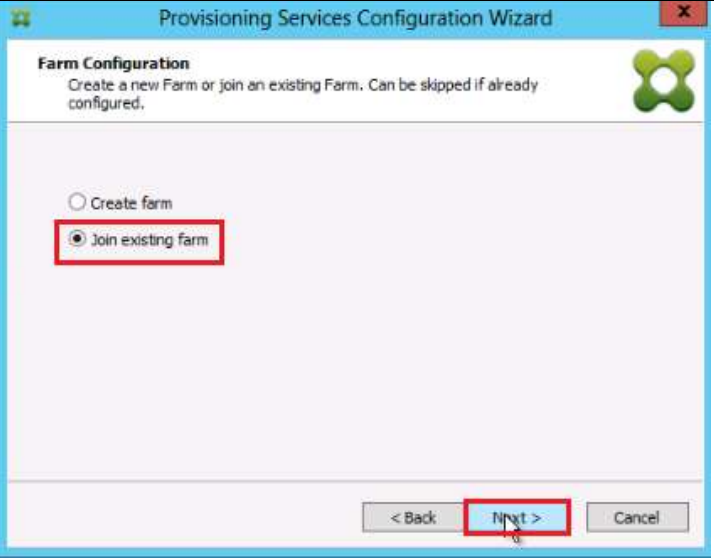
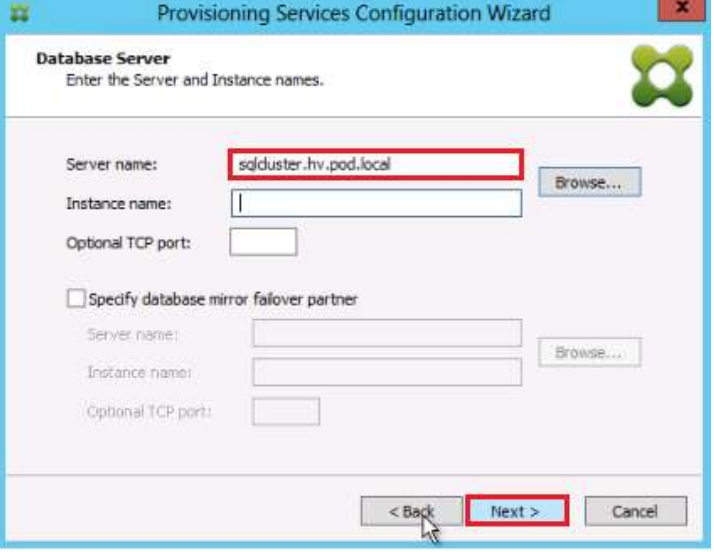
Instructions	Visual
<p>Read the Citrix License Agreement.</p> <p>If acceptable, select the radio button labeled “I accept the terms in the license agreement.”</p> <p>Click Next.</p>	
<p>Optionally provide User Name and Organization.</p> <p>Click Next.</p>	

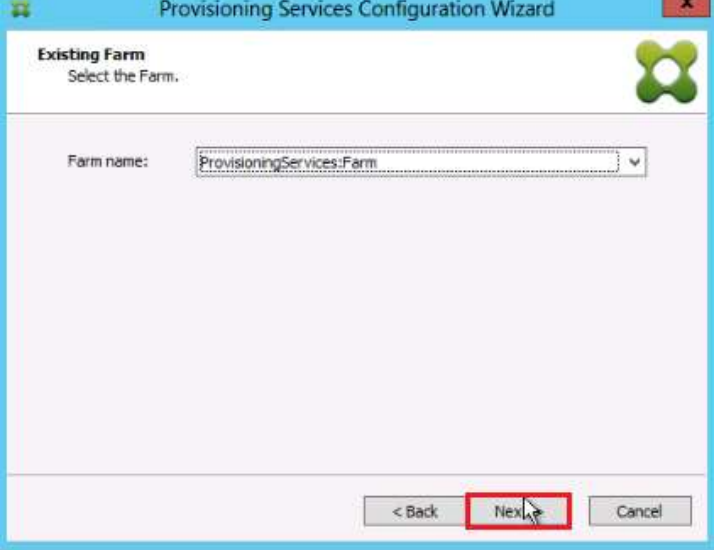
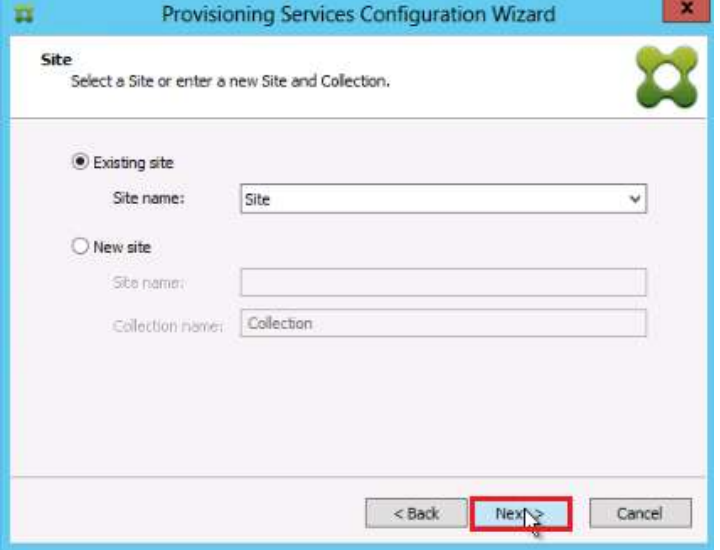
Instructions	Visual
<p>Accept the default path.</p> <p>Click Next.</p>	
<p>Leave the Complete radio button selected.</p> <p>Click Next.</p>	

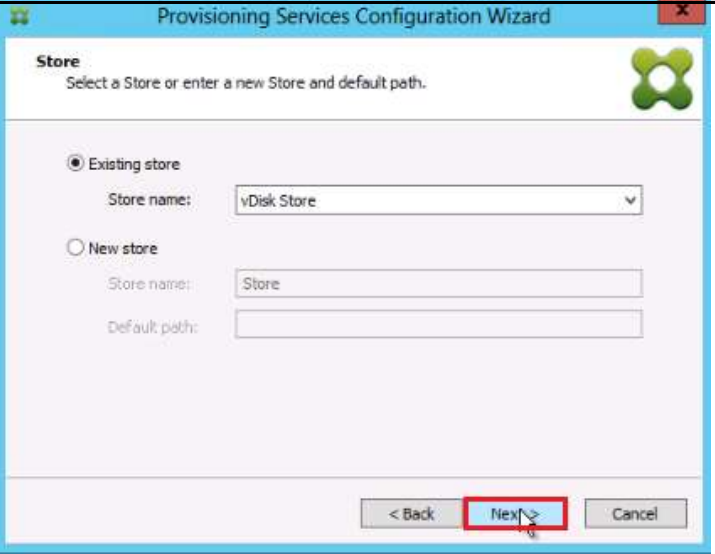
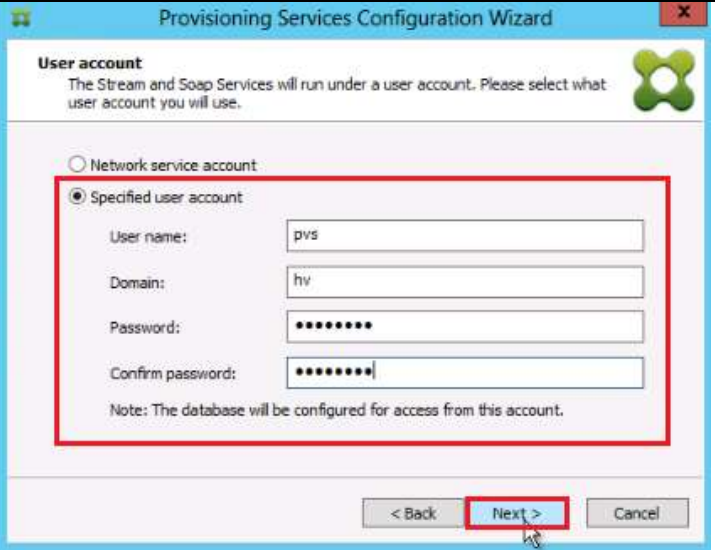
Instructions	Visual
Click Install to start the console installation.	
When the Installation completes, click Finish to close the dialog box.	

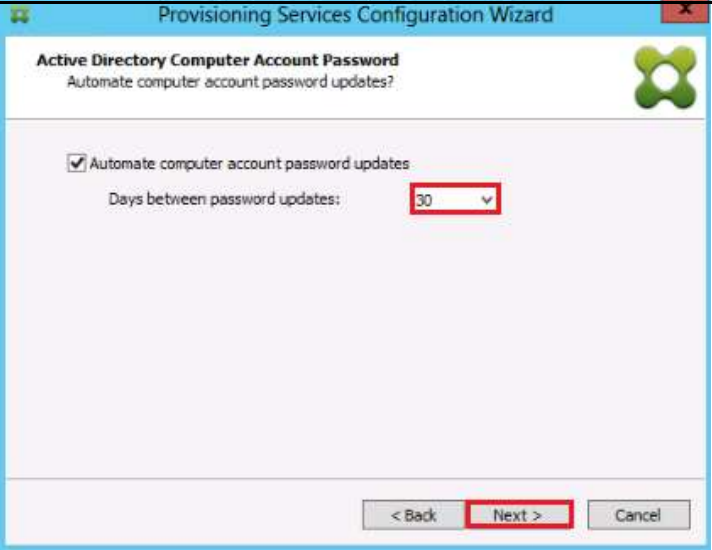
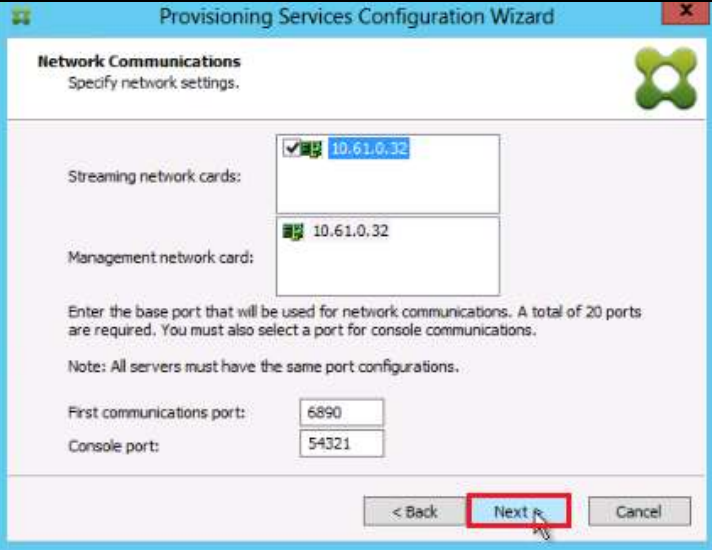
6.14.2. Installation of the Second PVS Server

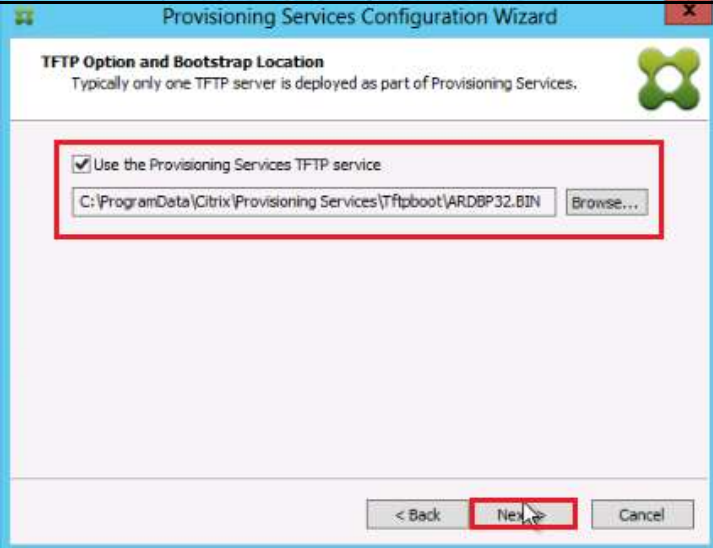
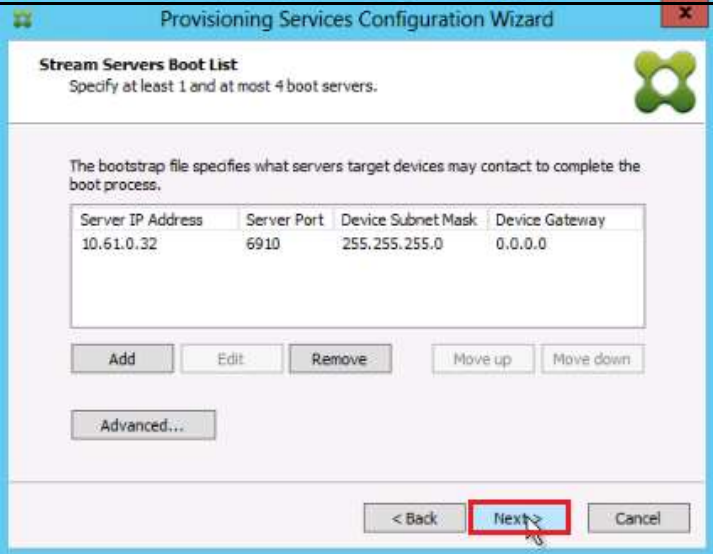
Complete the installation steps on the second PVS server as previously completed on the first PVS server to the configuration step where it asks to Create or Join a farm.

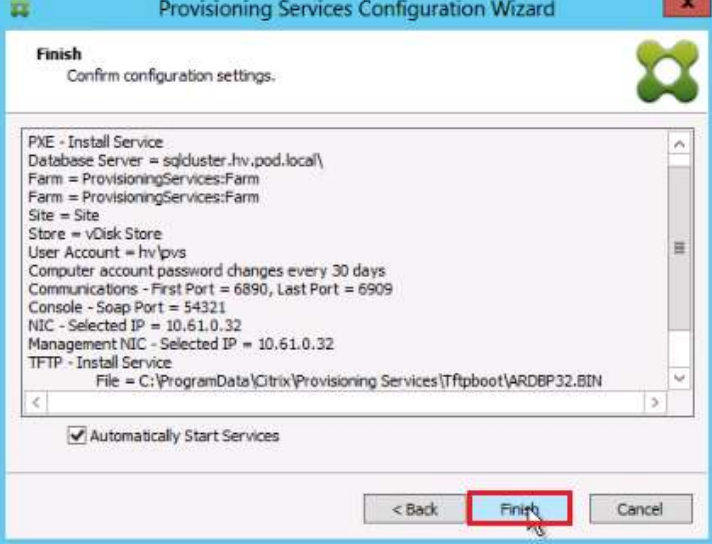
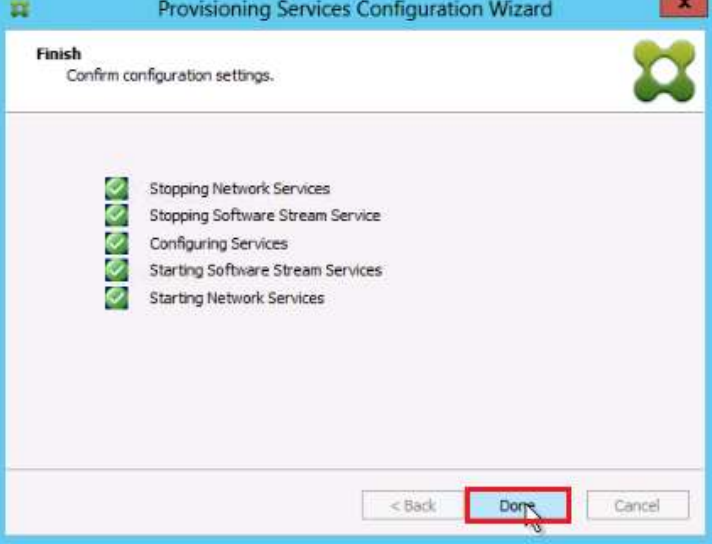
Instructions	Visual
<p>From the Farm Configuration dialog, select “Join existing farm.”</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The title bar says 'Provisioning Services Configuration Wizard'. The main heading is 'Farm Configuration'. Below it, it says 'Create a new Farm or join an existing Farm. Can be skipped if already configured.' There are two radio buttons: 'Create farm' and 'Join existing farm'. The 'Join existing farm' button is selected and highlighted with a red box. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red box.</p>
<p>Enter the listener name of the SQL server AlwaysOn group that hosts the database as the Server name.</p> <p>Note: If using a cluster, instead of AlwaysOn groups, you will need to supply the instance name as well.</p>	 <p>The screenshot shows the 'Provisioning Services Configuration Wizard' window. The title bar says 'Provisioning Services Configuration Wizard'. The main heading is 'Database Server'. Below it, it says 'Enter the Server and Instance names.' There are three text boxes: 'Server name:', 'Instance name:', and 'Optional TCP port:'. The 'Server name:' text box contains the text 'sqlcluster.hv.pod.local' and is highlighted with a red box. To the right of the 'Server name:' text box is a 'Browse...' button. Below these are four more text boxes: 'Server name:', 'Instance name:', 'Optional TCP port:', and 'Optional TCP port:'. The first 'Server name:' text box is highlighted with a red box. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red box.</p>

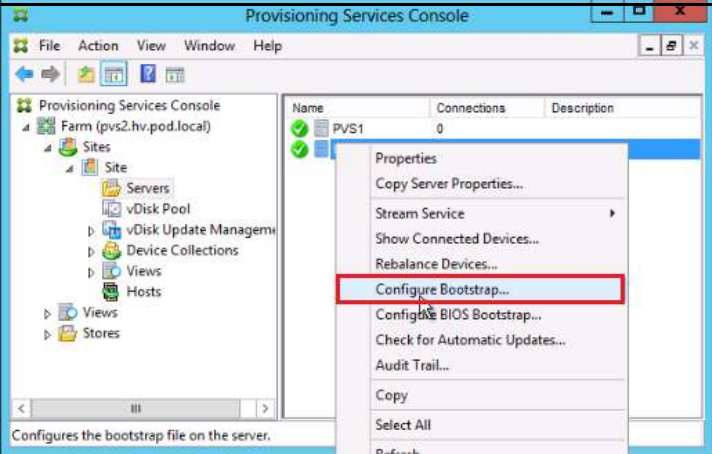
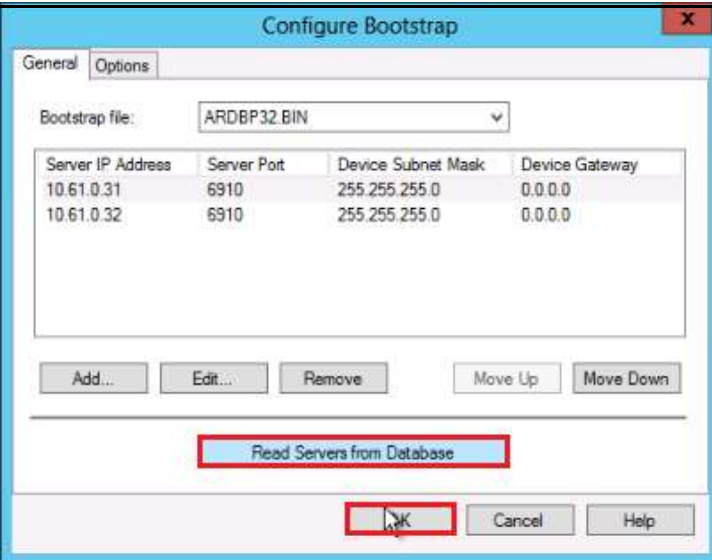
Instructions	Visual
<p>Accept the Farm Name.</p> <p>Click Next.</p>	
<p>Accept the Existing Site.</p> <p>Click Next.</p>	

Instructions	Visual
<p>Accept the Existing vDisk store.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Store' configuration screen of the Provisioning Services Configuration Wizard. The title bar reads 'Provisioning Services Configuration Wizard'. Below the title, it says 'Store' and 'Select a Store or enter a new Store and default path.' There are two radio buttons: 'Existing store' (selected) and 'New store'. Under 'Existing store', there is a 'Store name' dropdown menu currently showing 'vDisk Store'. Under 'New store', there are text boxes for 'Store name' (containing 'Store') and 'Default path' (empty). At the bottom, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>
<p>Provide the PVS user account information.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'User account' configuration screen of the Provisioning Services Configuration Wizard. The title bar reads 'Provisioning Services Configuration Wizard'. Below the title, it says 'User account' and 'The Stream and Soap Services will run under a user account. Please select what user account you will use.' There are two radio buttons: 'Network service account' and 'Specified user account' (selected). Under 'Specified user account', there are four text boxes: 'User name' (containing 'pvs'), 'Domain' (containing 'hv'), 'Password' (masked with dots), and 'Confirm password' (masked with dots). A note below these boxes states: 'Note: The database will be configured for access from this account.' At the bottom, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>

Instructions	Visual
<p>Set the Days between password updates to 30.</p> <p>Click Next.</p>	
<p>Accept the network card settings.</p> <p>Click Next.</p>	

Instructions	Visual								
<p>Enable the “Use the Provisioning Services TFTP Service” checkbox.</p> <p>Click Next.</p>	 <p>Provisioning Services Configuration Wizard</p> <p>TFTP Option and Bootstrap Location Typically only one TFTP server is deployed as part of Provisioning Services.</p> <p><input checked="" type="checkbox"/> Use the Provisioning Services TFTP service</p> <p>C:\ProgramData\Citrix\Provisioning Services\Tftpboot\ARDBP32.BIN Browse...</p> <p>< Back Next > Cancel</p>								
<p>Accept the Stream Servers list.</p> <p>Click Next.</p>	 <p>Provisioning Services Configuration Wizard</p> <p>Stream Servers Boot List Specify at least 1 and at most 4 boot servers.</p> <p>The bootstrap file specifies what servers target devices may contact to complete the boot process.</p> <table><thead><tr><th>Server IP Address</th><th>Server Port</th><th>Device Subnet Mask</th><th>Device Gateway</th></tr></thead><tbody><tr><td>10.61.0.32</td><td>6910</td><td>255.255.255.0</td><td>0.0.0.0</td></tr></tbody></table> <p>Add Edit Remove Move up Move down</p> <p>Advanced...</p> <p>< Back Next > Cancel</p>	Server IP Address	Server Port	Device Subnet Mask	Device Gateway	10.61.0.32	6910	255.255.255.0	0.0.0.0
Server IP Address	Server Port	Device Subnet Mask	Device Gateway						
10.61.0.32	6910	255.255.255.0	0.0.0.0						

Instructions	Visual
<p>Click Finish to start the installation process.</p>	
<p>Click Done when the installation finishes.</p>	


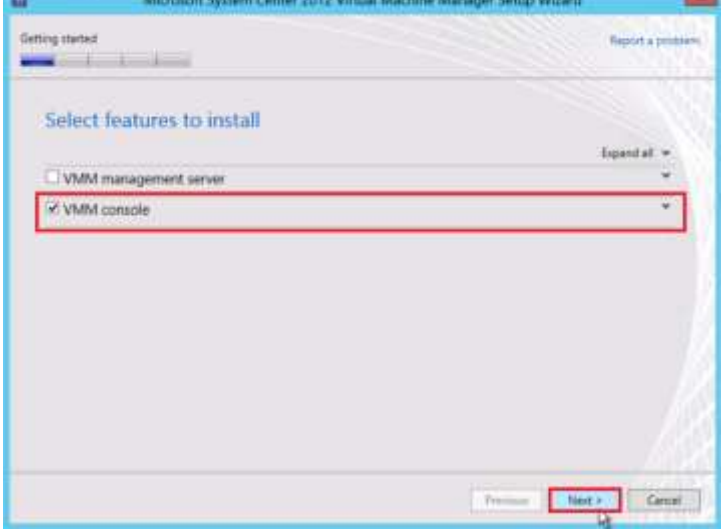
Instructions	Visual												
<p>With a second PVS server the boot strap should be updated on both servers. The boot strap server is responsible for providing vDisk assignments to target devices.</p> <p>Launch the PVS console.</p> <p>Select the server and right-click and select Configure Bootstrap... from the context menu.</p>	 <p>The screenshot shows the 'Provisioning Services Console' window. On the left, a tree view shows the hierarchy: Farm (pvs2.hv.pod.local) > Sites > Site > Servers. A server is selected, and a right-click context menu is open. The menu items include 'Properties', 'Copy Server Properties...', 'Stream Service', 'Show Connected Devices...', 'Rebalance Devices...', 'Configure Bootstrap...' (highlighted with a red rectangle), 'Configure BIOS Bootstrap...', 'Check for Automatic Updates...', 'Audit Trail...', 'Copy', 'Select All', and 'Refresh'. The status bar at the bottom says 'Configures the bootstrap file on the server.'</p>												
<p>Click the Read Servers from Database button to add the other PVS server.</p> <p>Click OK to save the changes</p> <p>Repeat for the other PVS server.</p>	 <p>The screenshot shows the 'Configure Bootstrap' dialog box. It has two tabs: 'General' and 'Options'. The 'General' tab is active, showing a 'Bootstrap file' dropdown set to 'ARDBP32.BIN'. Below is a table with server information:</p> <table><tr><th>Server IP Address</th><th>Server Port</th><th>Device Subnet Mask</th><th>Device Gateway</th></tr><tr><td>10.61.0.31</td><td>6910</td><td>255.255.255.0</td><td>0.0.0.0</td></tr><tr><td>10.61.0.32</td><td>6910</td><td>255.255.255.0</td><td>0.0.0.0</td></tr></table> <p>Below the table are buttons: 'Add...', 'Edit...', 'Remove', 'Move Up', and 'Move Down'. At the bottom, the 'Read Servers from Database' button is highlighted with a red rectangle. The 'OK' button is also highlighted with a red rectangle. Other buttons at the bottom are 'Cancel' and 'Help'.</p>	Server IP Address	Server Port	Device Subnet Mask	Device Gateway	10.61.0.31	6910	255.255.255.0	0.0.0.0	10.61.0.32	6910	255.255.255.0	0.0.0.0
Server IP Address	Server Port	Device Subnet Mask	Device Gateway										
10.61.0.31	6910	255.255.255.0	0.0.0.0										
10.61.0.32	6910	255.255.255.0	0.0.0.0										

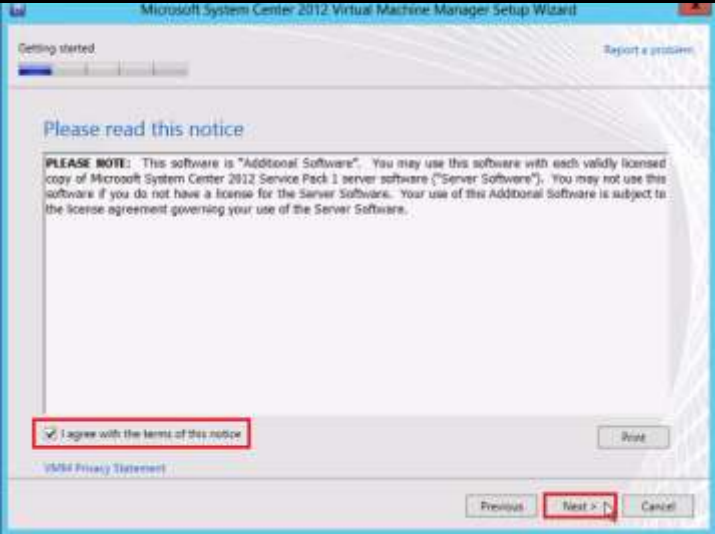
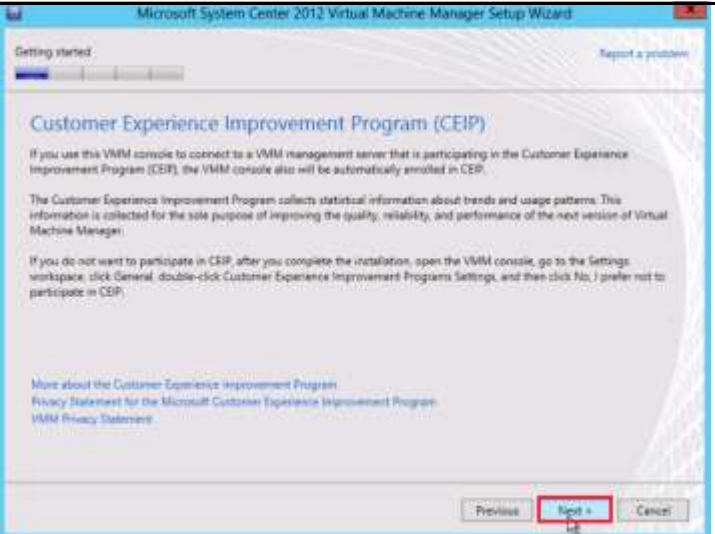
6.14.3. Installing the Delivery Controller

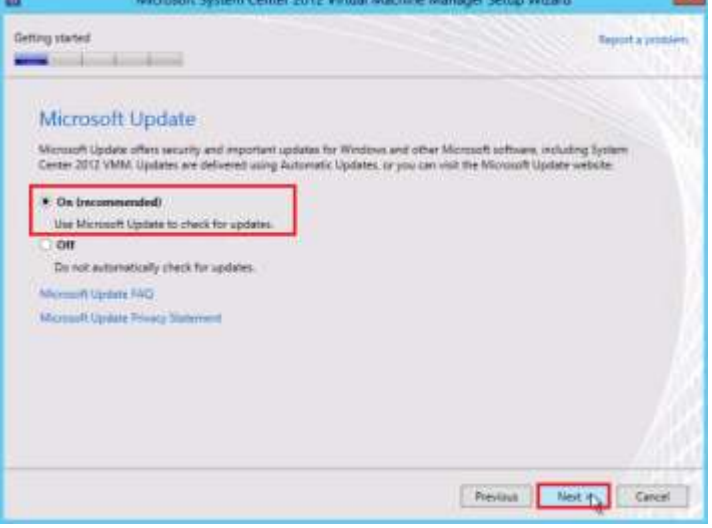
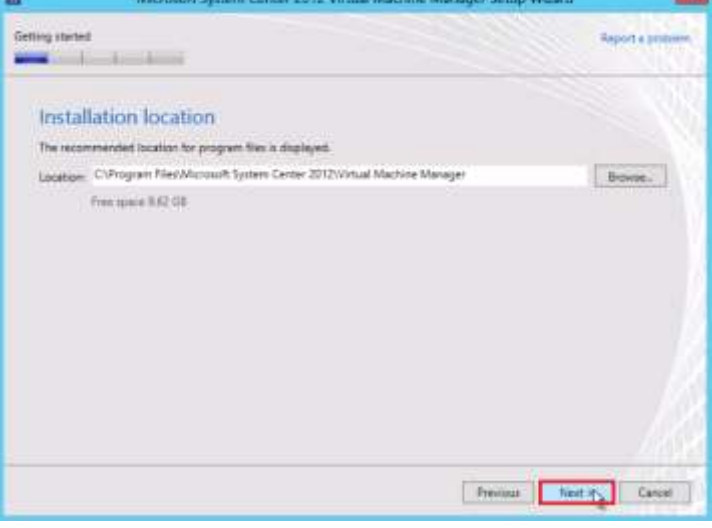
The steps in this section should be completed on all XenDesktop Delivery Controllers. There are two key tasks described in this section: (1) installing Microsoft System Center 2012 Virtual Machine Manager (SCVMM) and (2) installing the XenDesktop Delivery Controllers and other XenDesktop software components.


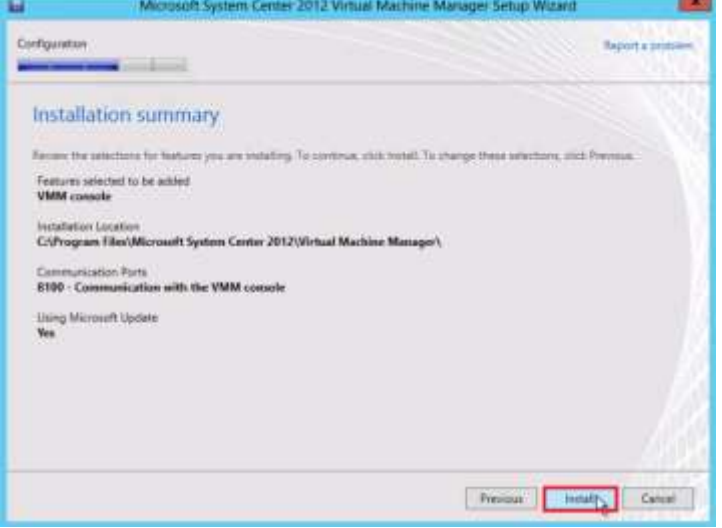
6.14.3.1. Installing Microsoft System Center 2012 Virtual Machine Manager (SCVMM)

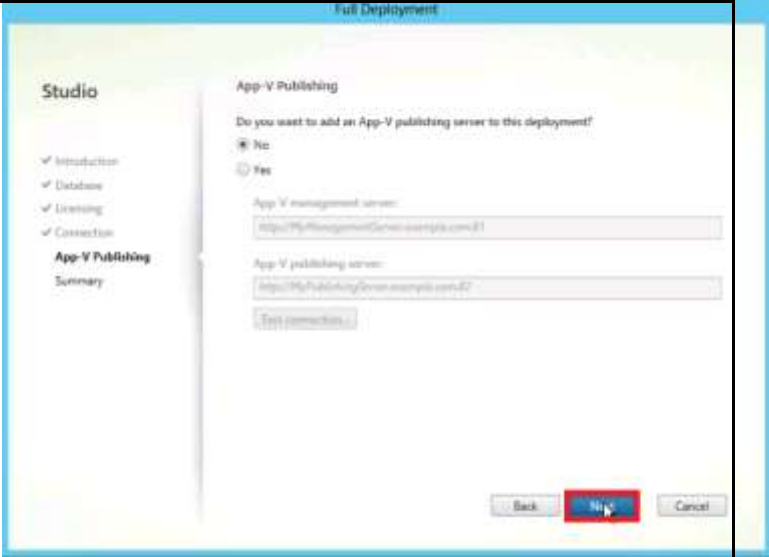
Microsoft System Center 2012 Virtual Machine Manager must first be installed before beginning XenDesktop software installation and configuration. After inserting the System Center Virtual Machine Manager ISO, you should see a Welcome screen.

Instructions	Visual
<p>Select the Install link.</p> <p>The installer will copy some files over to the host then present the feature screen.</p>	
<p>Select the VMM Console for install.</p> <p>Click Next.</p>	

Instructions	Visual
<p>Read the notice.</p> <p>I acceptable, enable the checkbox labeled “I agree with the terms of this notice.”</p> <p>Click Next.</p>	
<p>From the Customer Experience Improvement Program screen, click Next .</p>	


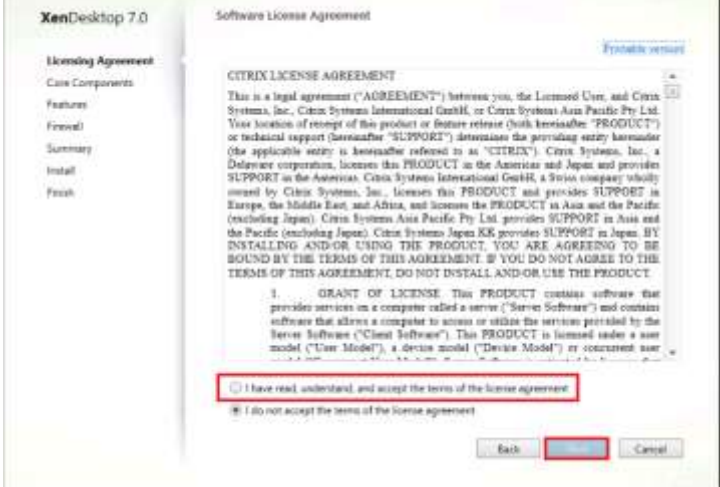
Instructions	Visual
<p>From the Microsoft Update screen enable the On (recommended) radio button.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Microsoft Update' screen of the setup wizard. It includes a progress bar at the top, a 'Report a problem' link, and instructions about updates. Two radio buttons are present: 'On (recommended)' (selected) and 'Off'. The 'Next' button at the bottom right is highlighted with a red box.</p>
<p>Accept the default installation location.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Installation location' screen. It states that the recommended location is displayed. The 'Location' field shows 'C:\Program Files\Microsoft System Center 2012\Virtual Machine Manager' with a 'Browse...' button next to it. Below the field, it indicates 'Free space 8.62 GB'. The 'Next' button at the bottom right is highlighted with a red box.</p>

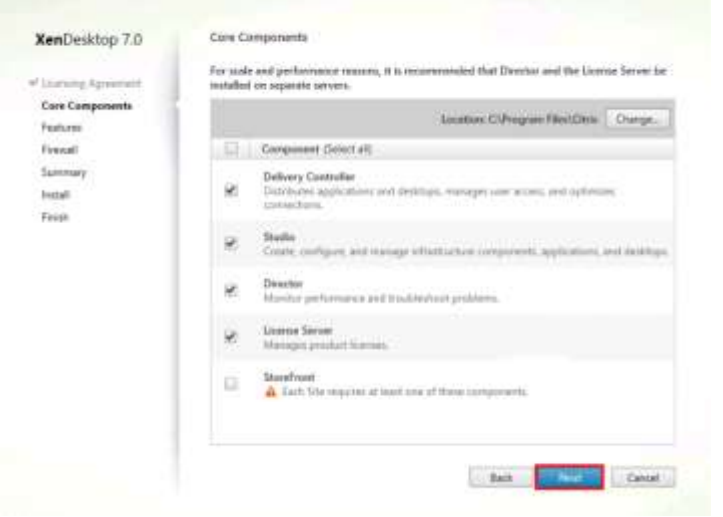
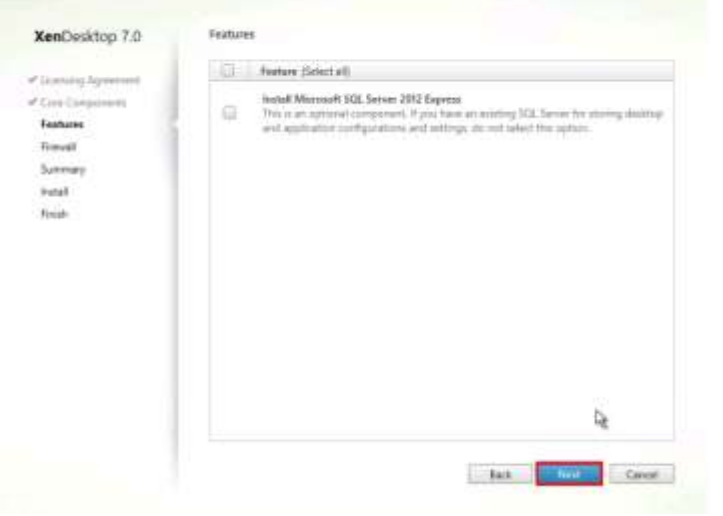
Instructions	Visual
<p>Accept the default communication port.</p> <p>Click Next.</p>	 <p>Microsoft System Center 2012 Virtual Machine Manager Setup Wizard</p> <p>Configuration</p> <p>Port configuration</p> <p>Administrator Console:</p> <p>8100 Communication with the VMM management server</p> <p>Previous Next Cancel</p>
<p>Review the installation summary.</p> <p>Click Install.</p>	 <p>Microsoft System Center 2012 Virtual Machine Manager Setup Wizard</p> <p>Configuration</p> <p>Installation summary</p> <p>Review the selections for features you are installing. To continue, click Install. To change these selections, click Previous.</p> <p>Features selected to be added</p> <p>VMM console</p> <p>Installation Location</p> <p>C:\Program Files\Microsoft System Center 2012\Virtual Machine Manager\</p> <p>Communication Ports</p> <p>8100 - Communication with the VMM console</p> <p>Using Microsoft Update</p> <p>Yes</p> <p>Previous Install Cancel</p>

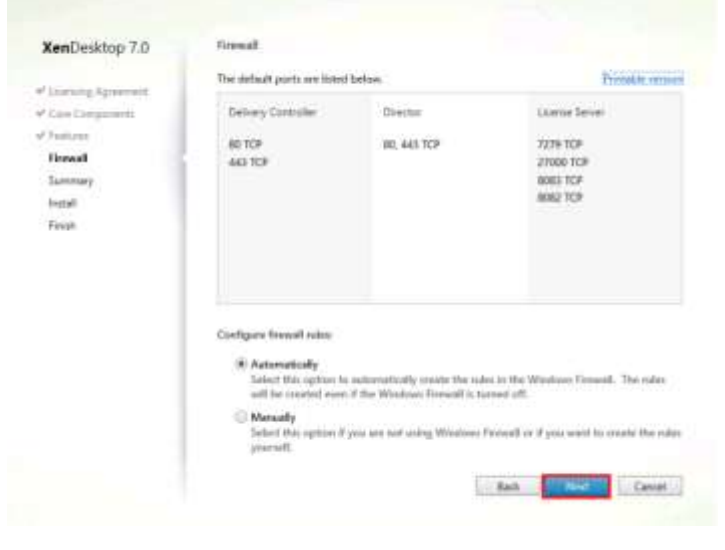
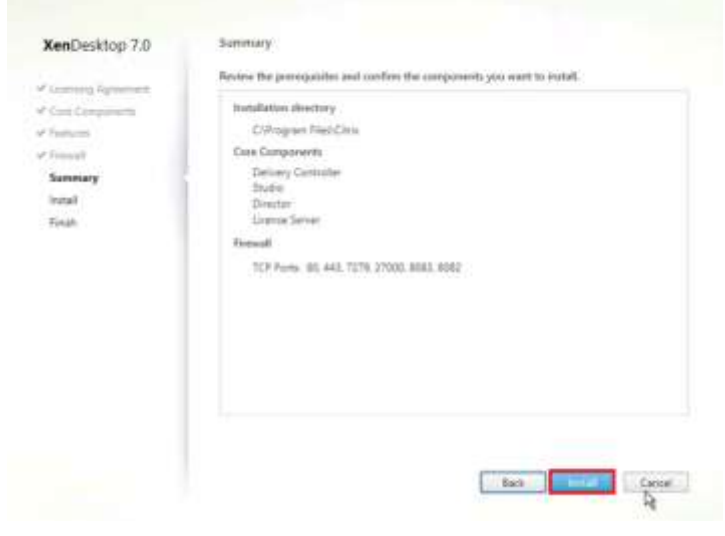
Instructions	Visual
<p>When installation is finished, disable the checkbox labeled Open the VMM console when this wizard closes.</p> <p>Click Close.</p> <p>Reboot the server before starting the XenDesktop configuration or the SCVMM configuration will fail.</p>	 <p>The screenshot shows the 'Full Deployment' wizard in Citrix Studio. The left pane shows the navigation tree with 'App-V Publishing' selected. The main pane is titled 'App-V Publishing' and contains the question 'Do you want to add an App-V publishing server to this deployment?'. The 'No' radio button is selected. Below this, there are two text boxes for 'App-V management server' and 'App-V publishing server', both containing placeholder URLs. A 'Test connections' button is located below the text boxes. At the bottom right, there are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a red box and a mouse cursor.</p>

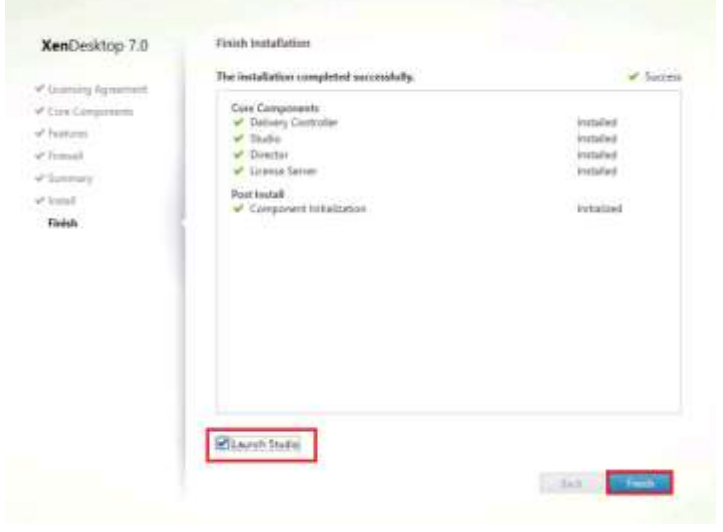
6.14.3.2. Installing XenDesktop Delivery Controllers and Other Components

After the installation of SCVMM on the VM, you must install the XenDesktop Delivery Controllers and other XenDesktop components. When you start the installation process from the Citrix XenDesktop 7 DVD, the installation wizard presents a menu with three topics.

Instructions	Visual
<p>To begin the installation of Delivery Controllers, click “Get Started - Delivery Controller.”</p>	 <p>The screenshot shows the XenDesktop 7.0 installation wizard. The 'Get Started' tab is selected and highlighted with a red box. It contains the 'Delivery Controller' section, which instructs the user to start here, select and install the Delivery Controller and other essential services like Licensing Server and Storefront. Other tabs visible include 'Prepare Machines and Images' and 'Extend Deployment'.</p>
<p>Read the Citrix License Agreement.</p> <p>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>Click Next.</p>	 <p>The screenshot displays the 'Software License Agreement' window for XenDesktop 7.0. The 'Citrix License Agreement' tab is active. The text of the agreement is visible, including the title 'CITRIX LICENSE AGREEMENT' and the introductory paragraph. At the bottom, there are two radio buttons: 'I have read, understand, and accept the terms of the license agreement' (which is selected and highlighted with a red box) and 'I do not accept the terms of the license agreement'. The 'Next' button is also highlighted with a red box.</p>

Instructions	Visual
<p>Select the components to be installed:</p> <ul style="list-style-type: none"> • Delivery Controller • Studio • Director • License Server (install only on one VM) <p>In this Cisco Validated Design, StoreFront is installed on a separate virtual machine, so uncheck StoreFront.</p> <p>Click Next.</p>	
<p>The Microsoft SQL Server is installed on a separate server, so uncheck the Install Microsoft SQL Server 2012 Express checkbox.</p> <p>Click Next.</p>	



Instructions	Visual															
<p>Select the default ports and automatically configured firewall rules.</p> <p>Click Next.</p>	 <p>XenDesktop 7.0</p> <ul style="list-style-type: none">✓ Learning Agreement✓ Core Components✓ FeaturesFirewallSummaryInstallFinish <p>Firewall</p> <p>The default ports are listed below.</p> <table><tr><th>Delivery Controller</th><th>Director</th><th>License Server</th></tr><tr><td>80 TCP</td><td>80, 443 TCP</td><td>7279 TCP</td></tr><tr><td>443 TCP</td><td></td><td>27000 TCP</td></tr><tr><td></td><td></td><td>8081 TCP</td></tr><tr><td></td><td></td><td>8082 TCP</td></tr></table> <p>Configure firewall rules:</p> <p><input checked="" type="radio"/> Automatically 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.</p> <p><input type="radio"/> Manually Select this option if you are not using Windows Firewall or if you want to create the rules yourself.</p> <p>Back Next Cancel</p>	Delivery Controller	Director	License Server	80 TCP	80, 443 TCP	7279 TCP	443 TCP		27000 TCP			8081 TCP			8082 TCP
Delivery Controller	Director	License Server														
80 TCP	80, 443 TCP	7279 TCP														
443 TCP		27000 TCP														
		8081 TCP														
		8082 TCP														
<p>The Summary screen is shown.</p> <p>Click Install to begin the installation.</p> <p>The installer displays a message when the installation is complete.</p>	 <p>XenDesktop 7.0</p> <ul style="list-style-type: none">✓ Learning Agreement✓ Core Components✓ Features✓ FirewallSummaryInstallFinish <p>Summary</p> <p>Review the prerequisites and confirm the components you want to install.</p> <p>Installation directory C:\Program Files\Citrix</p> <p>Core Components Delivery Controller Studio Director License Server</p> <p>Firewall TCP Ports: 80, 443, 7279, 27000, 8081, 8082</p> <p>Back Install Cancel</p>															

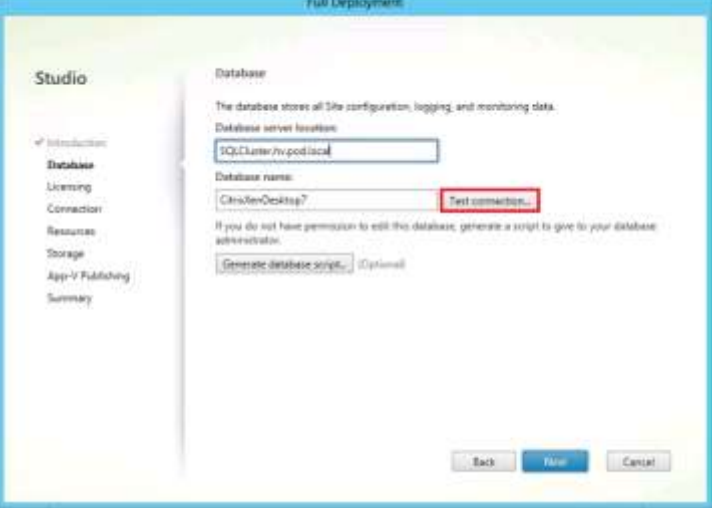
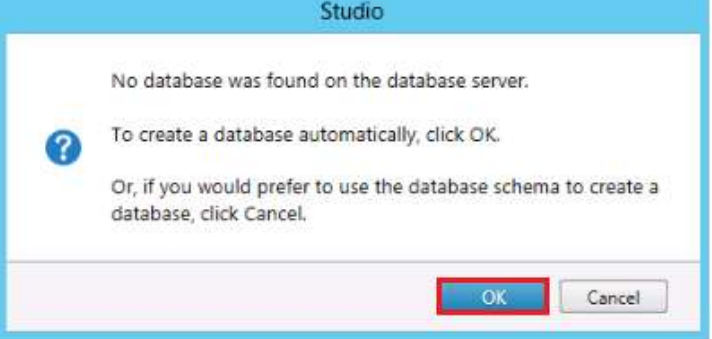
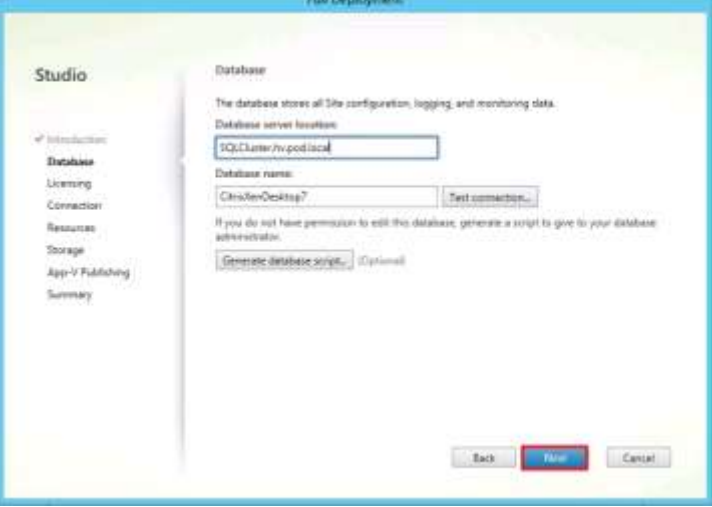
Instructions	Visual
<p>Verify the Launch Studio checkbox is enabled.</p> <p>Click Finish.</p>	

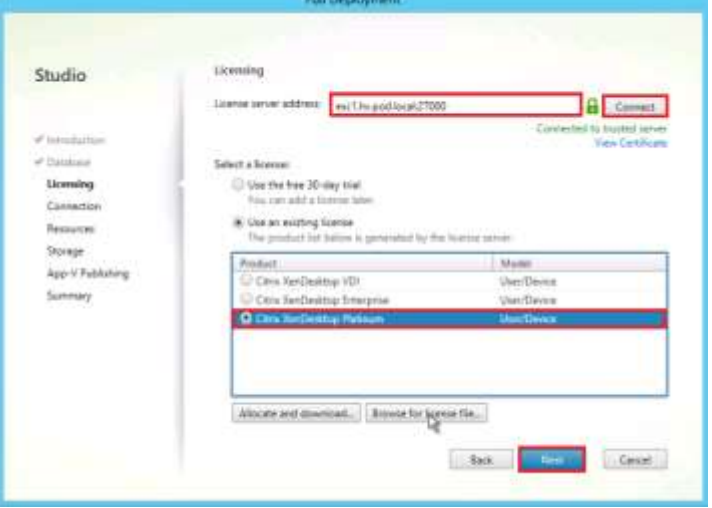
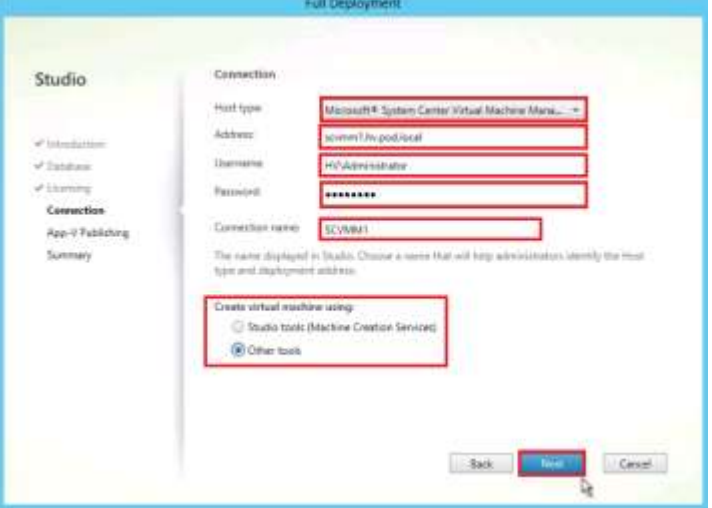
6.14.4. XenDesktop Controller Configuration

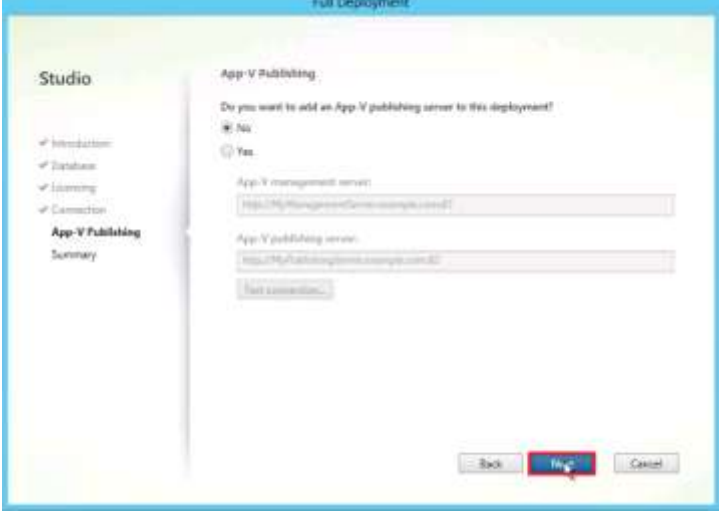
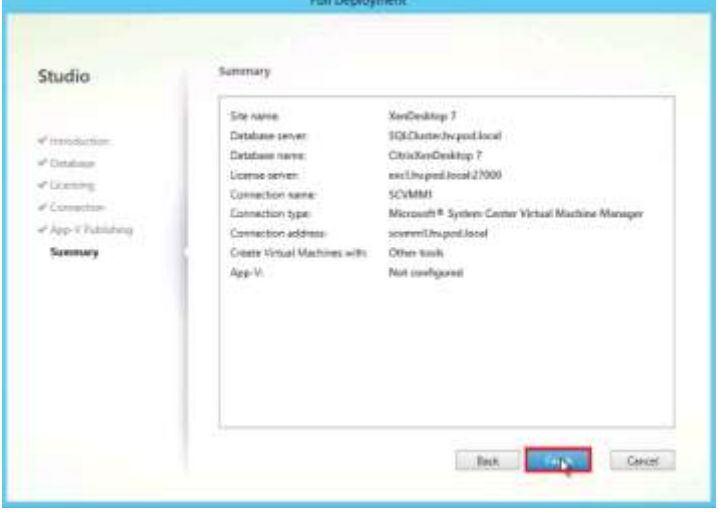
Citrix Studio is a management console that allows you to create and manage infrastructure and resources to deliver desktops and applications. Replacing Desktop Studio from earlier releases, it provides wizards to set up your environment, create workloads to host applications and desktops, and assign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

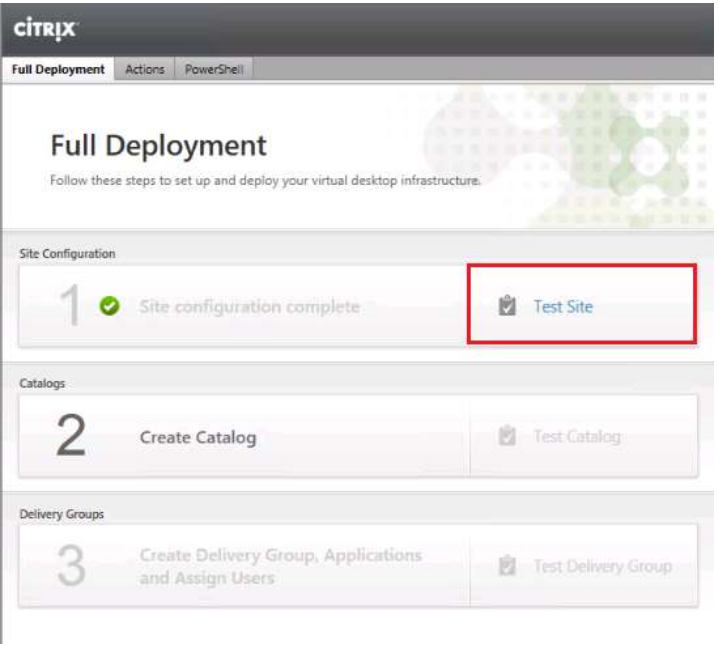
Citrix Studio launches automatically after component installation from DVD, or if necessary, it can be launched manually. You can use Studio to create a Site, which is the core XenDesktop 7 environment consisting of the Delivery Controller and the Database.

Instructions	Visual
<p>Click Get Started! Create a Site.</p>	 <p>The screenshot shows the Citrix Studio 'Welcome' page. At the top is the Citrix logo and a 'Welcome' tab. Below is a 'Welcome to Citrix Studio' heading with a subtext: 'Use this console to configure a fresh deployment, create a new Site, or extend your existing deployment.' There are three main sections: 'Full deployment' with a highlighted button 'Get started! Create a Site' (which includes the text 'Get started! Deploy applications and desktops for your organization. (Remote PC Access deployment can be added later)'), 'Remote PC Access deployment' with a button 'Provide secure remote access to physical PCs' (text: 'Build a deployment to allow users remote access to their physical PCs through a secure connection. (Full deployment can be added later)'), and 'Extend' with a button 'Scale out your deployment' (text: 'Add the Delivery Controller installed on this server to an existing Site.').</p>
<p>Select the “Configure a Site and start delivering applications and desktops to users” radio button.</p> <p>Enter a site name.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Full Deployment' configuration screen in Citrix Studio. On the left is a navigation pane with links: Introduction, Database, Licensing, Connection, Resources, Storage, App-V Publishing, and Summary. The main area is titled 'Build and customize a full production environment.' Under 'What would you like to create?', there are two radio buttons. The first, 'Configure the Site and start delivering applications and desktops to users (recommended for new users)', is selected. The second is 'Create an empty Site (recommended for advanced users who want to add the Site later)'. Below the radio buttons is a text field labeled 'Name your site:' containing the text 'XenDesktop1'. At the bottom right are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a red box.</p>

Instructions	Visual
<p>Provide the Database Server location using the AlwaysOn SQL Listener name.</p> <p>Click Test connection... to verify that the database is accessible.</p> <p>Note: If using a clustered database instead of the AlwaysOn configuration, then the SQL instance name must also be supplied. Ignore any errors and continue.</p>	
<p>Click OK to have the installer create the database.</p>	
<p>Click Next.</p>	

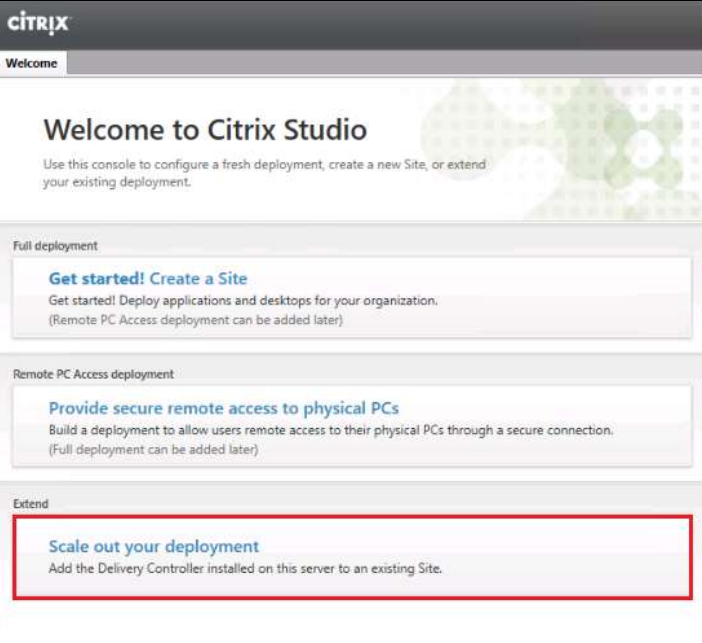
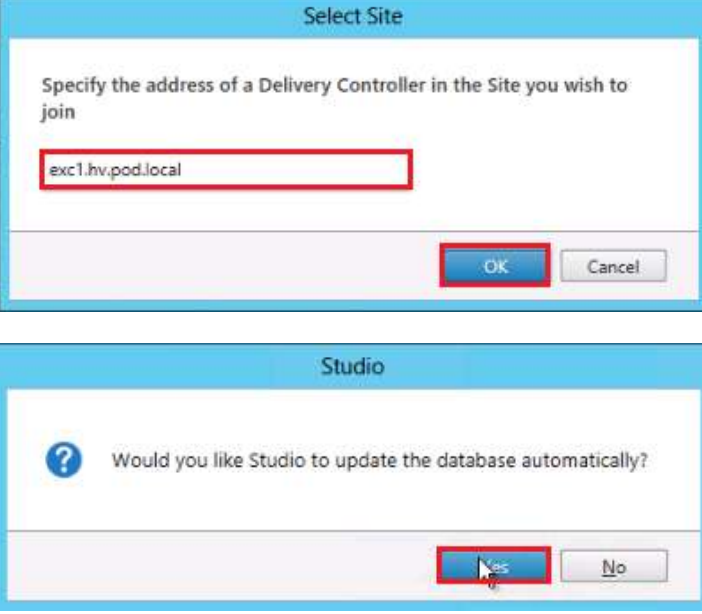
Instructions	Visual
<p>Provide the FQDN of the license server.</p> <p>Click Connect to validate and retrieve any licenses from the server.</p> <p>If no licenses are available, you can use the 30-day free trial or upload a license file.</p> <p>Select the Citrix XenDesktop Platinum license radio button.</p> <p>Click Next.</p>	
<p>Select the Host Type of Microsoft System Center Virtual Machine Manager</p> <p>Enter the FQDN of the SCVMM server</p> <p>Enter the username (in domain\username format) for the SCVMM account.</p> <p>Provide the password for the SCVMM account.</p> <p>Provide a connection name</p> <p>Select the Other tools radio button since Provisioning Services will be used.</p> <p>Click Next.</p>	

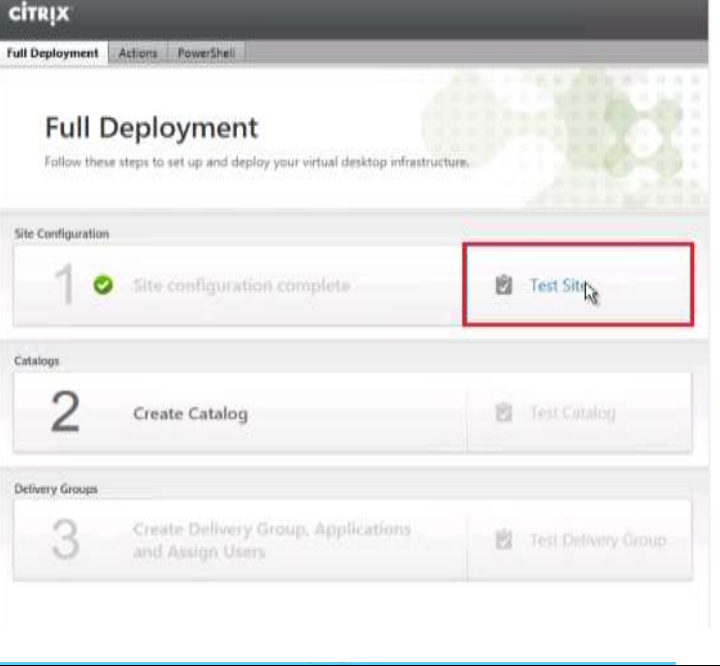
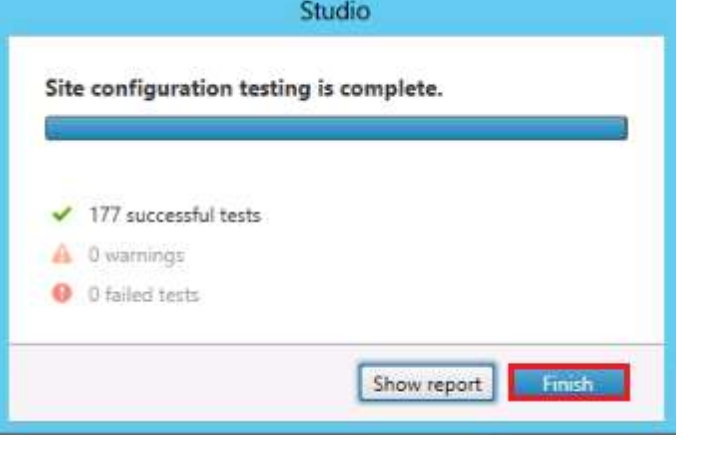
Instructions	Visual																				
<p>Click Next on the App-V Publishing dialog.</p>	 <p>The screenshot shows the 'Full Deployment' wizard in the 'App-V Publishing' step. On the left, the 'Studio' pane lists steps: Introduction, Database, Licensing, Connection, App-V Publishing, and Summary. The main area asks 'Do you want to add an App-V publishing server to this deployment?' with radio buttons for 'No' (selected) and 'Yes'. Below are text boxes for 'App-V management server' and 'App-V publishing server', both containing placeholder URLs. A 'Test connection...' button is below the publishing server field. At the bottom right, 'Back', 'Next' (highlighted with a red box), and 'Cancel' buttons are visible.</p>																				
<p>Click Finish to complete the deployment.</p>	 <p>The screenshot shows the 'Full Deployment' wizard in the 'Summary' step. The 'Studio' pane on the left highlights 'Summary'. The main area displays a table with deployment details. At the bottom right, 'Back', 'Finish' (highlighted with a red box), and 'Cancel' buttons are visible.</p> <table border="1" data-bbox="1003 909 1458 1098"> <thead> <tr> <th colspan="2">Summary</th> </tr> </thead> <tbody> <tr> <td>Site name</td> <td>XenDesktop 7</td> </tr> <tr> <td>Database server</td> <td>SQLCluster\sqlpool.local</td> </tr> <tr> <td>Database name</td> <td>CitrixXenDesktop 7</td> </tr> <tr> <td>License server</td> <td>sqlc1.us.pool.local:27000</td> </tr> <tr> <td>Connection name</td> <td>SCVMM1</td> </tr> <tr> <td>Connection type</td> <td>Microsoft® System Center Virtual Machine Manager</td> </tr> <tr> <td>Connection address</td> <td>sqlcvm1.us.pool.local</td> </tr> <tr> <td>Create Virtual Machines with</td> <td>Other tools</td> </tr> <tr> <td>App-V</td> <td>Not configured</td> </tr> </tbody> </table>	Summary		Site name	XenDesktop 7	Database server	SQLCluster\sqlpool.local	Database name	CitrixXenDesktop 7	License server	sqlc1.us.pool.local:27000	Connection name	SCVMM1	Connection type	Microsoft® System Center Virtual Machine Manager	Connection address	sqlcvm1.us.pool.local	Create Virtual Machines with	Other tools	App-V	Not configured
Summary																					
Site name	XenDesktop 7																				
Database server	SQLCluster\sqlpool.local																				
Database name	CitrixXenDesktop 7																				
License server	sqlc1.us.pool.local:27000																				
Connection name	SCVMM1																				
Connection type	Microsoft® System Center Virtual Machine Manager																				
Connection address	sqlcvm1.us.pool.local																				
Create Virtual Machines with	Other tools																				
App-V	Not configured																				

Instructions	Visual
<p>When the deployment is complete, click the Test Site button.</p> <p>All 177 tests should pass successfully.</p> <p>Click Finish.</p>	 <p>The screenshot shows the Citrix Studio 'Full Deployment' wizard. The 'Site Configuration' step is marked as complete with a green checkmark and the text 'Site configuration complete'. A red box highlights the 'Test Site' button. Below this, the 'Catalogs' and 'Delivery Groups' sections are visible, each with a 'Test' button. At the bottom, a 'Studio' dialog box displays 'Site configuration testing is complete.' with a progress bar and summary: '177 successful tests', '0 warnings', and '0 failed tests'. A red box highlights the 'Finish' button in this dialog.</p>

6.14.5. Additional XenDesktop Controller Configuration

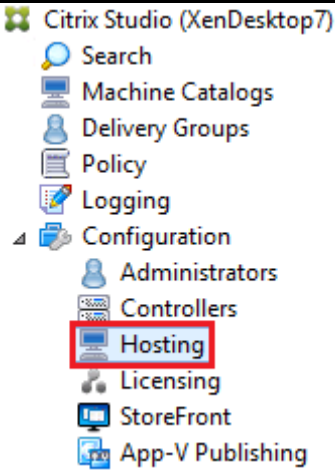
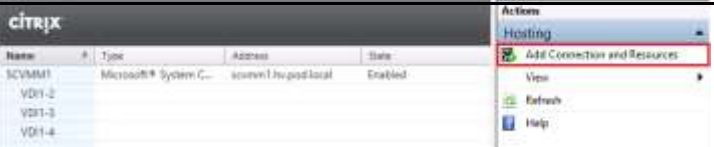
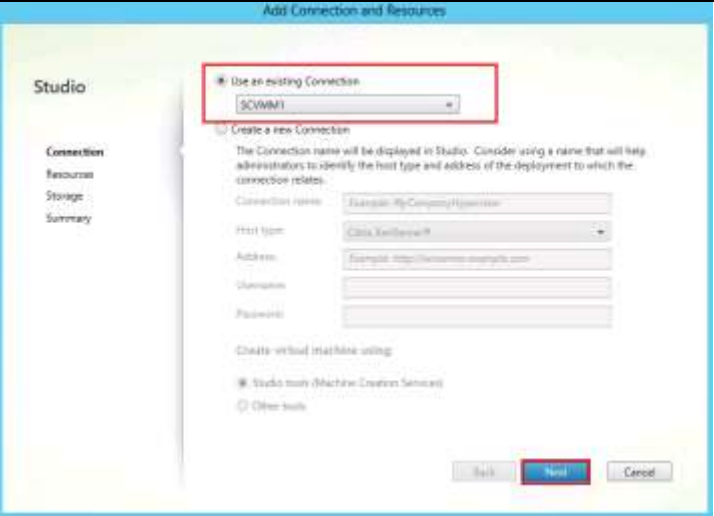
After the first controller is completely configured and the Site is operational, additional controllers can be added. Also, all the installation steps configured in XenDesktop Controller Configuration section must be completed prior to beginning this section.

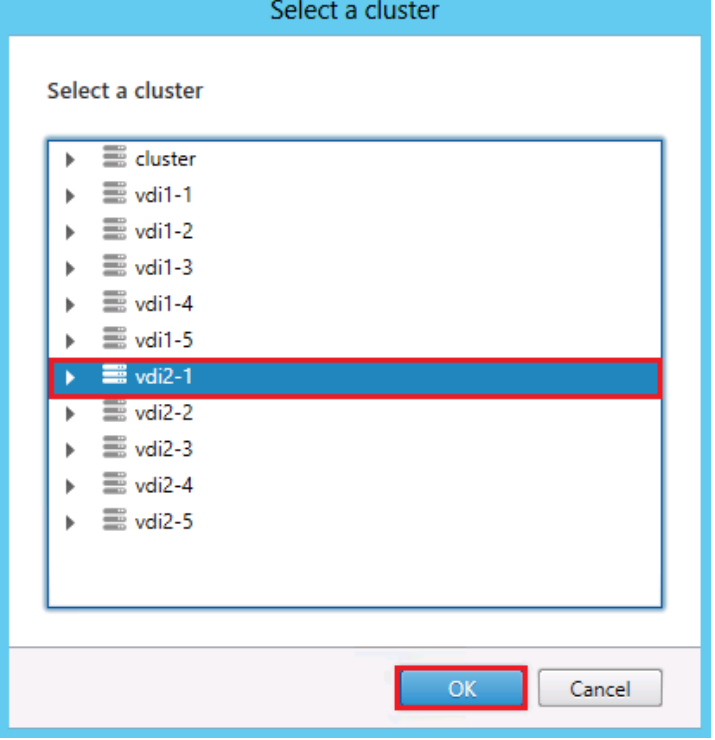
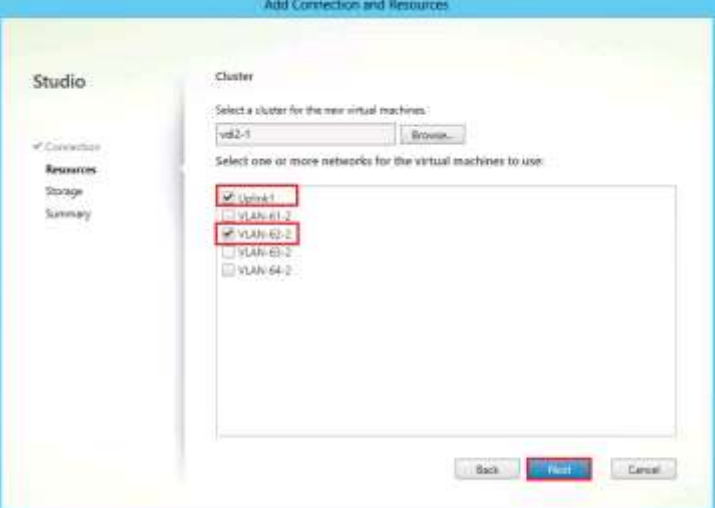
Instructions	Visual
<p>Click Scale out your deployment.</p>	 <p>The image shows the Citrix Studio 'Welcome' screen. At the top is the Citrix logo and a 'Welcome' tab. Below is the heading 'Welcome to Citrix Studio' with a subtext: 'Use this console to configure a fresh deployment, create a new Site, or extend your existing deployment.' There are three main sections: 'Full deployment' with a button 'Get started! Create a Site', 'Remote PC Access deployment' with a button 'Provide secure remote access to physical PCs', and 'Extend' with a button 'Scale out your deployment'. The 'Scale out your deployment' button is highlighted with a red rectangle.</p>
<p>Enter the FQDN of the first delivery controller configured earlier.</p> <p>Click OK.</p> <p>Click Yes to allow the database to be updated with this controller's information automatically.</p>	 <p>The image shows two dialog boxes. The first is 'Select Site' with the text 'Specify the address of a Delivery Controller in the Site you wish to join'. A text box contains 'exc1.hv.pod.local' and is highlighted with a red rectangle. Below the text box are 'OK' and 'Cancel' buttons, with the 'OK' button highlighted by a red rectangle. The second dialog box is 'Studio' with the question 'Would you like Studio to update the database automatically?'. It has 'Yes' and 'No' buttons, with the 'Yes' button highlighted by a red rectangle.</p>

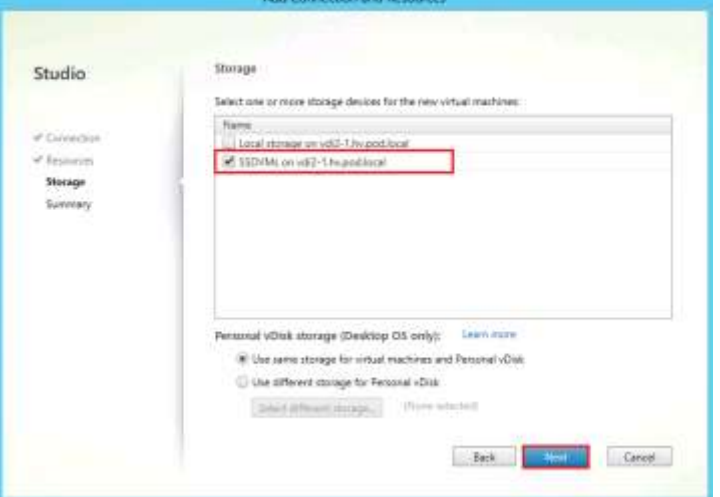
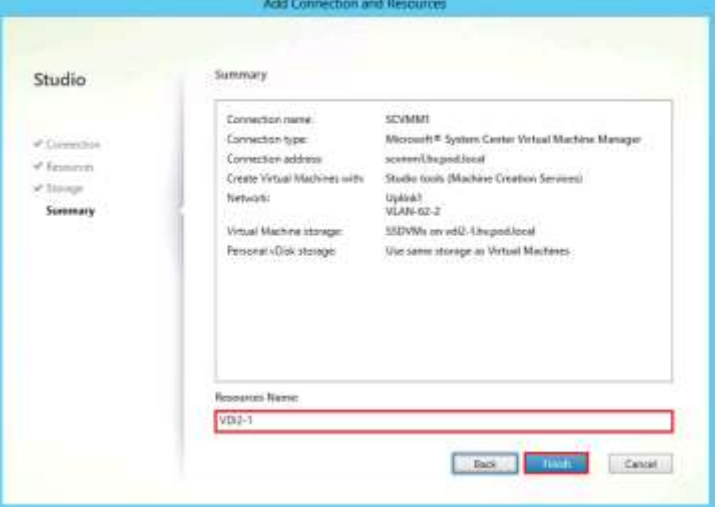
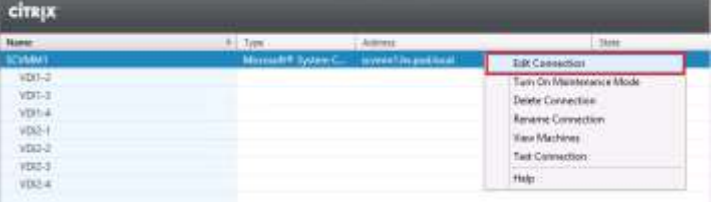
Instructions	Visual
<p>When complete, verify the site is functional by clicking Test Site.</p>	 <p>The screenshot shows the Citrix 'Full Deployment' wizard. The 'Site Configuration' step is highlighted with a large number '1' and a green checkmark, indicating it is complete. A red rectangular box highlights the 'Test Site' button in the top right corner of the 'Site Configuration' section. Below this, the 'Catalogs' section shows a large number '2' and a 'Create Catalog' button, with a 'Test Catalog' button to its right. The 'Delivery Groups' section shows a large number '3' and a 'Create Delivery Group, Applications and Assign Users' button, with a 'Test Delivery Group' button to its right.</p>
<p>Click Finish to close the test results dialog.</p>	 <p>The screenshot shows the 'Studio' test results dialog. It states 'Site configuration testing is complete.' with a progress bar. Below this, it shows '177 successful tests' with a green checkmark, '0 warnings' with a yellow warning icon, and '0 failed tests' with a red error icon. At the bottom, there are two buttons: 'Show report' and 'Finish'. The 'Finish' button is highlighted with a red rectangular box.</p>

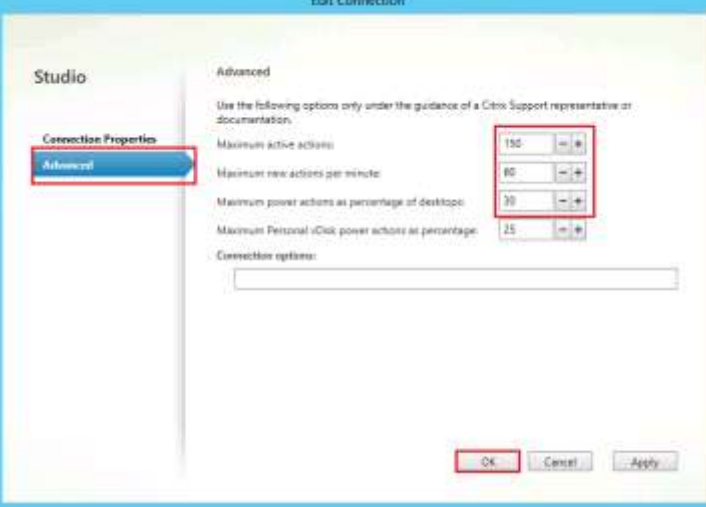
6.14.6. Configure VM Hosting, Connections, and Resources

Each of the individual Hyper-V hosts will need to be added to the XenDesktop configuration.

Instructions	Visual
<p>From the Citrix Studio expand the Configuration node.</p> <p>Select the Hosting node.</p>	
<p>From the right-side Action pane select the Add Connection and Resources link to start the Add Connection and Resources wizard.</p>	
<p>Select the SCVMM server that the host has been registered with.</p> <p>Click Next.</p>	


Instructions	Visual
<p>Select the host to add.</p> <p>Click OK.</p>	 <p>The screenshot shows a 'Select a cluster' dialog box. It contains a list of clusters: 'cluster', 'vdi1-1', 'vdi1-2', 'vdi1-3', 'vdi1-4', 'vdi1-5', 'vdi2-1', 'vdi2-2', 'vdi2-3', 'vdi2-4', and 'vdi2-5'. The 'vdi2-1' cluster is selected and highlighted with a red border. At the bottom of the dialog, there are 'OK' and 'Cancel' buttons. The 'OK' button is highlighted with a red box.</p>
<p>Select the networks that will be available to the virtual machines. Since the XenDesktop Setup Wizard does not support the Nexus 1000V logical switch, both the standard switch network (Uplink1) and the logical network (VLAN-62-2) will be selected.</p> <p>Click Next.</p> <p>Note: The Uplink1 network is used only for the Virtual Machine creation through the Wizard.</p>	 <p>The screenshot shows the 'Add Connection and Resources' dialog box. The 'Cluster' tab is active, showing 'vdi2-1' selected. Under 'Select one or more networks for the virtual machines to use:', 'Uplink1' and 'VLAN-62-2' are selected with checkboxes. At the bottom, there are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a red box.</p>

Instructions	Visual
<p>Select the location that reflects the share name created on the local SSD drives.</p> <p>Click Next.</p> <p>Note: If the share name is not visible, create the share and then use Refresh in the SCVMM console on the host. Then restart the Add Connections and Resources wizard.</p>	
<p>Provide a Name for the resource connection that will appear in the Citrix Studio console.</p> <p>Click Finish.</p>	
<p>Repeat the Add Connections and Resources for each of the remaining hosts.</p>	
<p>When all hosts are added to the environment, right-click on the SCVMM server and choose Edit Connection from the context menu.</p>	

Instructions	Visual
<p>Click Advanced .</p> <p>Set the Maximum Active Actions to 150.</p> <p>Set the Maximum new actions per minute to 60 (Recommended value: 10 * Number of Hosts).</p> <p>Set the Maximum power actions as percentage of desktops to 30.</p> <p>Click OK.</p>	

6.14.7. Installing and Configuring StoreFront

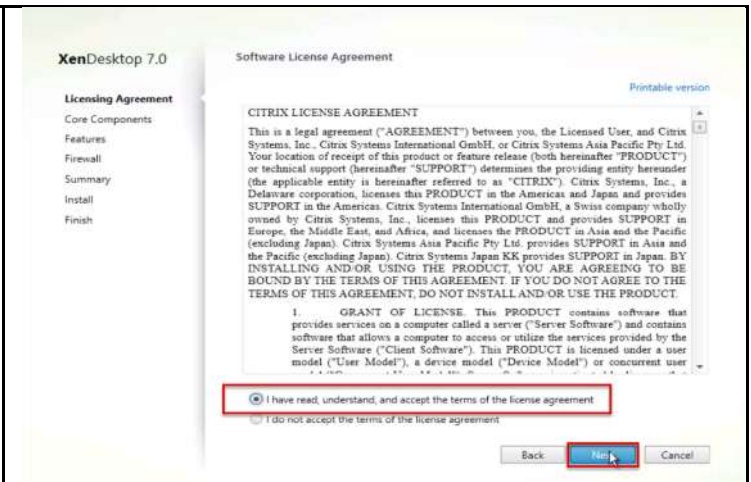
In this Cisco Validated Design, StoreFront is installed on a separate virtual machine from other XenDesktop components. Log into that virtual machine, and start the installation process from the Citrix XenDesktop 7 DVD. The installation wizard presents a menu with three topics.

<p>To begin the installation of StoreFront, click on “Get Started - Delivery Controller.”</p>	
--	--

Read the Citrix License Agreement.

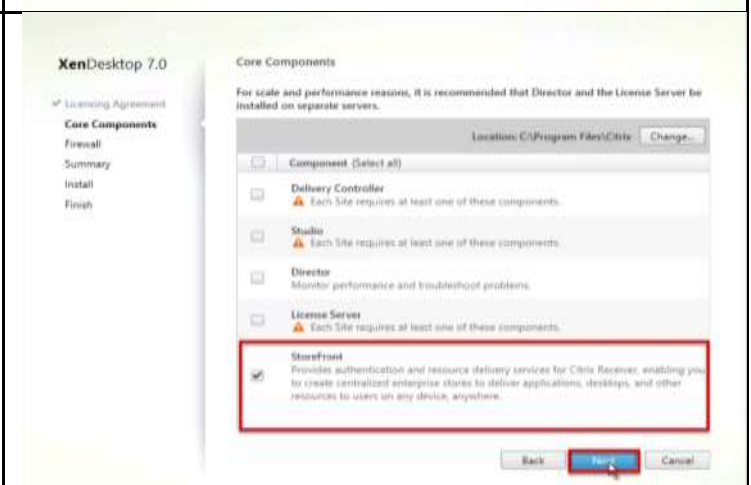
If acceptable, indicate your acceptance of the license by selecting the “**I have read, understand, and accept the terms of the license agreement**” radio button.

Click **Next**.



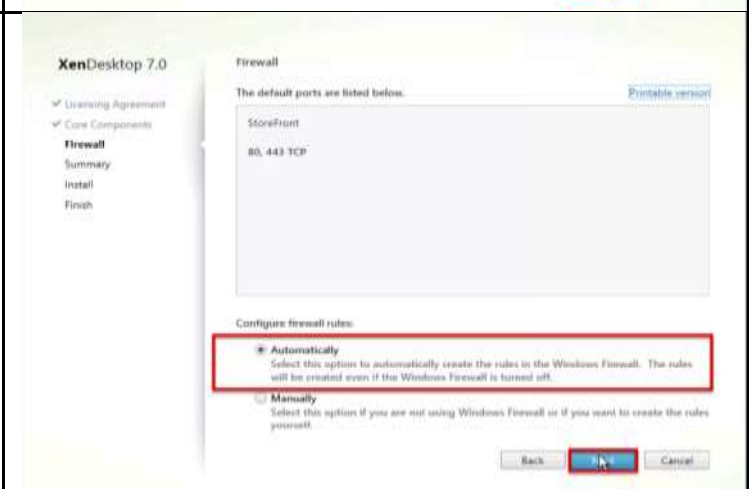
Select **StoreFront** as the component to be installed.

Click **Next**.



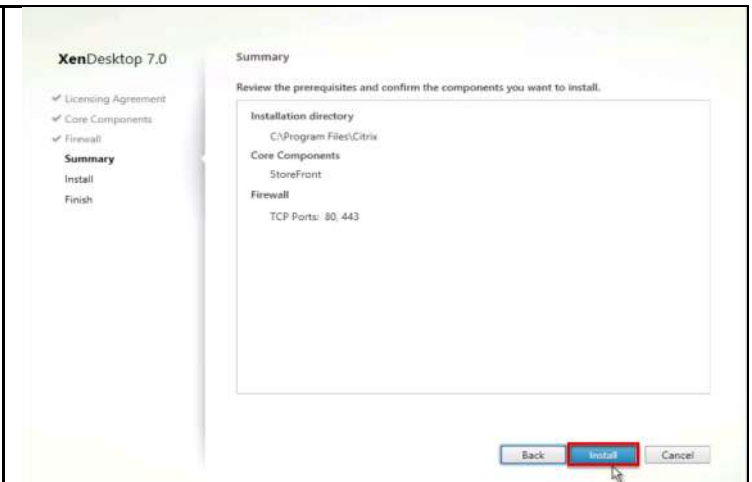
Select the default ports and automatically configured firewall rules.

Click **Next**.



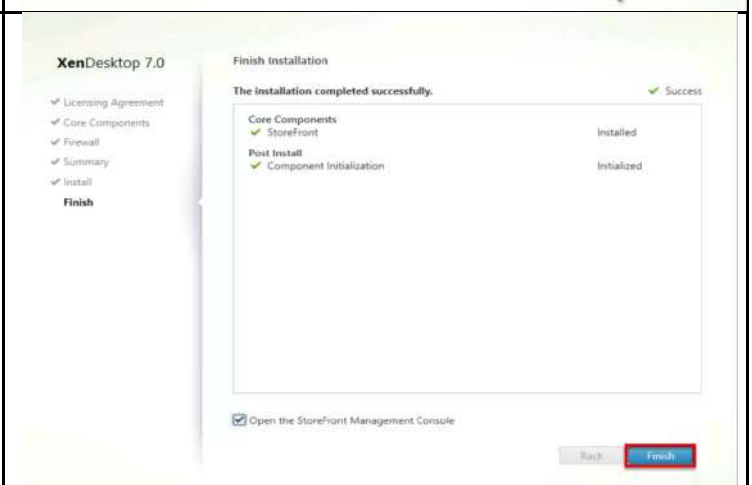
The Summary screen is shown.

Click the **Install** button to begin the installation.
The installer displays a message when the installation is complete.

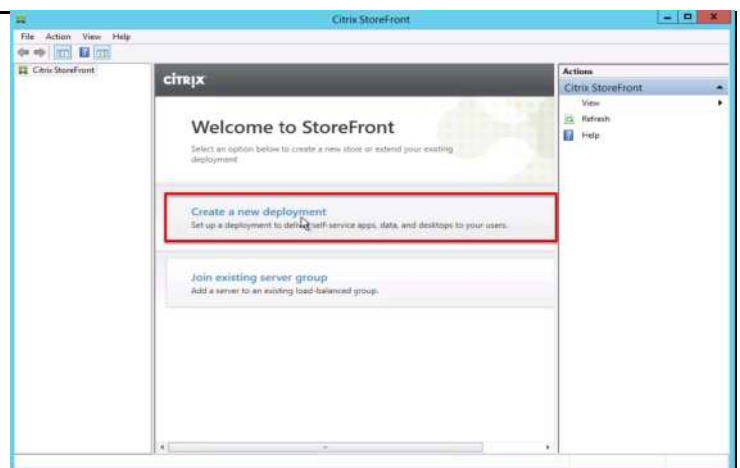


Verify that the checkbox to “**Open the StoreFront Management Console**” is enabled.

Click **Finish**.



Citrix StoreFront stores aggregate desktops and applications from XenDesktop sites, XenApp farms, and AppController, making these resources readily available to users. Citrix StoreFront launches automatically after installation from DVD, or if necessary, it can be launched manually. Click on the “**Create a new deployment**” button.



Enter the **Base URL** to be used to access StoreFront services.

Click **Next**.

The screenshot shows the 'Create New Deployment' dialog box. On the left, a sidebar lists 'Base URL', 'Store Name', 'Delivery Controllers', and 'Remote Access'. The 'Base URL' option is selected. The main area contains the text: 'Confirm the base URL for services hosted on this deployment. For multiple server deployments, specify the load-balanced URL for the server group.' Below this, the 'Base URL' field is populated with 'http://sfstfw.pod.local'. At the bottom right, the 'Next' button is highlighted with a red box, and the 'Cancel' button is visible next to it.

Enter a **Store Name**.

Click **Next**.

The screenshot shows the 'Create Store' dialog box. On the left, a sidebar lists 'Base URL', 'Store Name', 'Delivery Controllers', and 'Remote Access'. The 'Store Name' option is selected. The main area contains the text: '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, the 'Store name' field is populated with 'XenDesktop7'. At the bottom right, the 'Next' button is highlighted with a red box, and the 'Cancel' button is visible next to it.

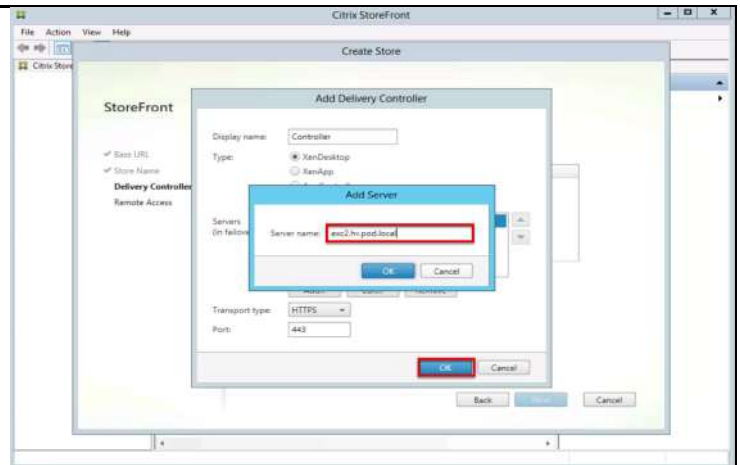
From the **Create Store** page, specify the XenDesktop Delivery Controller and servers that will provide resources to be made available in the store.

Click **Add**.

The screenshot shows the 'Create Store' dialog box at the 'Delivery Controllers' step. The sidebar on the left has 'Delivery Controllers' selected. The main area says 'Specify the delivery controllers and servers for this store.' Below this is a table with columns 'Name', 'Type', and 'Servers'. At the bottom of the table area, the 'Add' button is highlighted with a red box. Other buttons like 'Back', 'Next', and 'Cancel' are visible at the bottom of the dialog.

In the **Add Delivery Controller** dialog box, add servers for the XenDesktop Delivery Controller. List the servers in failover order.

Click **OK** to add each server to the list.



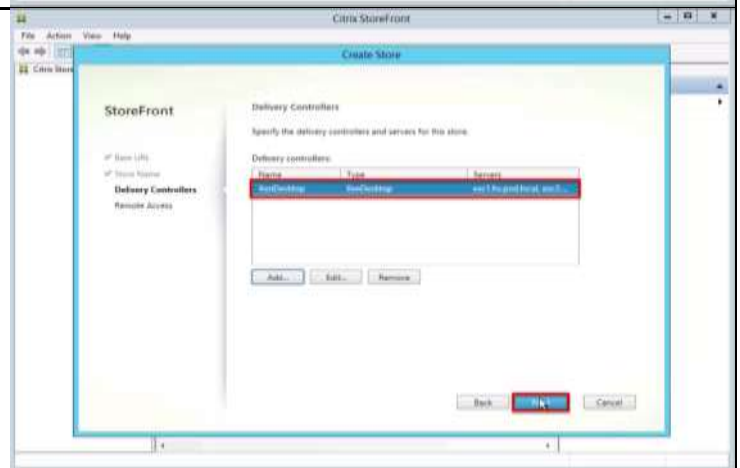
After adding the list of servers, specify a **Display name** for the DeliveryController. For testing purposes, set the default transport type and port to HTTP on port 80.

Click **OK** to add the Delivery Controller.




The **Delivery Controller** associated with the store now appears on the **Create Store** page.

Click **Next**.



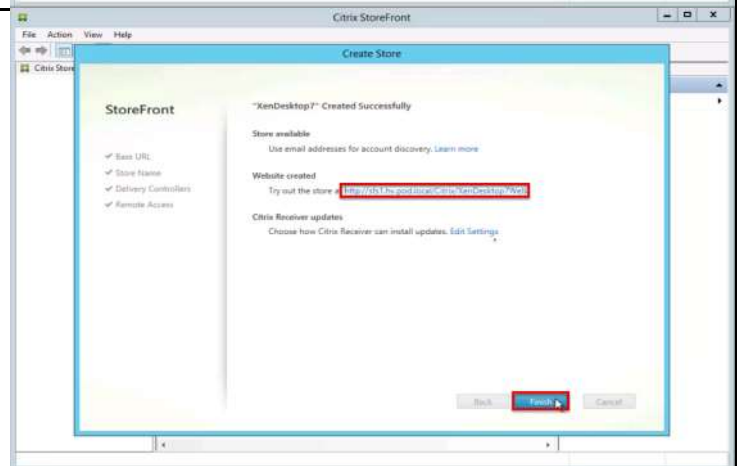
From the **Remote Access** page, accept **None** (the default).

Click **Create** to begin creating the store.




A message indicates when the store creation process is complete. The **Create Store** page lists the Website for the created store.


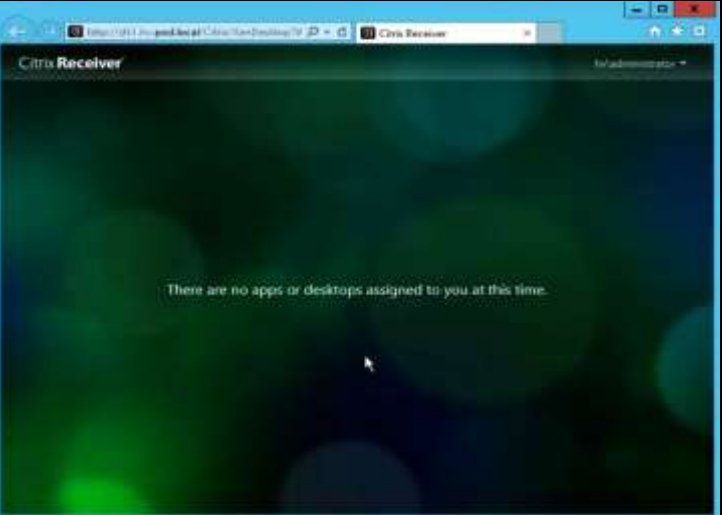
Click **Finish**.



To test the store, open it using the link in the message indicating that the store was successfully created.

Click **Log on**.



<p>Log into Citrix Receiver as the administrator and with the appropriate password.</p> <p>Click Log On.</p>	
<p>Since there are no apps or desktops configured in the store at this time, you will see a message to that effect.</p>	

7. Desktop Delivery Infrastructure - Citrix

7.1. Overview of desktop delivery

The advantage of using Citrix Provisioning Services (PVS) is that it allows VMs to be provisioned and re-provisioned in real-time from a single shared disk image called a virtual Disk (vDisk). By streaming a vDisk rather than copying images to individual machines, PVS allows organizations to manage a small number of disk images even when the number of VMs grows, providing the benefits of centralized management, distributed processing, and efficient use of storage capacity.

In most implementations, a single vDisk provides a standardized image to multiple target devices. Multiple PVS servers in the same farm can stream the same vDisk image to thousands of target devices. Virtual desktop environments can be customized through the use of write caches and by personalizing user settings through Citrix User Profile Management.

This section describes the installation and configuration tasks required to create standardized master vDisk images using PVS. This section also discusses write cache sizing and placement considerations, and how policies in Citrix User Profile Management can be configured to further personalize user desktops.

7.1.1. PVS vDisk Image Management

After installing and configuring PVS components, 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. vDisks can exist on a Provisioning Server, file share, or in larger deployments (as in this Cisco Validated Design), on a storage system with which the Provisioning Server can communicate (through iSCSI, SAN, NAS, and CIFS). A PVS server can access many stored vDisks, and each vDisk can be several gigabytes in size. For this solution, the vDisk was stored on a SMB3/CIFS share located on the EMC storage.

vDisks can be assigned to a single target device in Private Image Mode, or to multiple target devices in Standard Image Mode. In Standard Image mode, the vDisk is read-only, which means that multiple target devices can stream from a single vDisk image simultaneously. Standard Image mode reduces the complexity of vDisk management and the amount of storage required since images are shared. In contrast, when a vDisk is configured to use Private Image Mode, the vDisk is read/write and only one target device can access the vDisk at a time.

When a vDisk is configured in Standard Image mode, each time a target device boots, it always boots from a "clean" vDisk image. Each target device then maintains a write cache to store any writes that the operating system needs to make, such as the installation of user-specific data or applications.

7.1.2. Overview – Golden Image Creation

For this Cisco Validated Design, PVS supplies these master or golden vDisk images to the target devices:

Table 13: Golden Image Descriptions

vDisk Name	Server location	Description
HSDGold	Virtual– INFRA-1	The PVS golden image of Microsoft Windows Server 2012 for Hosted Shared Desktops.
XDGold	Virtual– INFRA-1	The PVS golden image of Microsoft Windows 7 SP1 for Hosted Virtual Desktops.

To build the vDisk images, OS images of Microsoft Windows 7 SP1 and Windows Server 2012 were initially installed, along with some additional software, and prepared as standard virtual machines on Microsoft Hyper-V 2012. These master OS images were then converted into a separate Citrix PVS vDisk files. During testing, Citrix PVS and the XenDesktop Delivery Controllers use the golden vDisk images to instantiate new desktop virtual machines on multiple Hyper-V targets.

In this Cisco Validated Design, virtual machines for the hosted shared and hosted virtual desktops were created using the XenDesktop Setup Wizard along with System Center Virtual Machine Manager 2012 PowerShell scripting. The XenDesktop Setup Wizard (XDSW) does the following:

1. Creates VMs on a XenDesktop hosted hypervisor server from an existing template.


2. Creates PVS target devices for each new VM within a new or existing collection matching the XenDesktop catalog name.
3. Assigns a Standard Image vDisk to VMs within the collection. Each virtual desktop is assigned a "Write Cache" (differencing disk) where any delta changes (writes) to the default image are recorded and is used by the virtual Windows operating system throughout its working life cycle. The Write Cache is written to a dedicated virtual hard disk created by thin provisioning and attached to each new virtual desktop using PowerShell scripts.
4. Adds virtual desktops to a XenDesktop Machine Catalog.

Virtual desktops were optimized according to best practices for performance. (In the steps later in this section, the "Optimize performance" checkbox was selected during the installation of the VDA, and the "Optimize for Provisioning Services" checkbox was selected during the PVS image creation process using the PVS Imaging Wizard.)

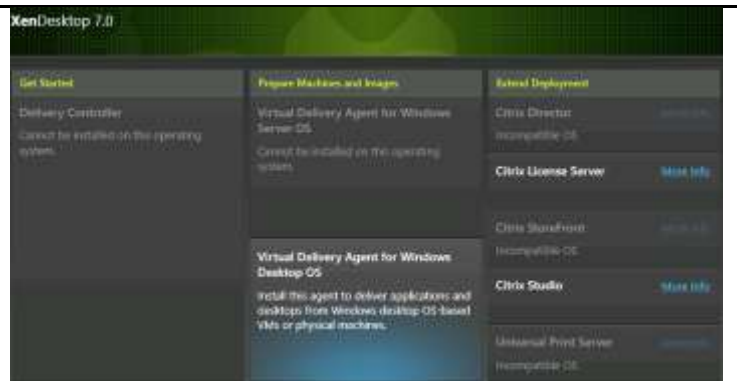
7.2. Installing the XenDesktop 7 Virtual Desktop Agent

Installed on the server and workstation operating systems, Virtual Delivery Agents (VDAs) enable connections for desktops and apps. The following procedure was used to install VDAs for both OS environments.

By default, when you install the Virtual Delivery Agent, Citrix User Profile Management 5.0 is installed silently on master images. (Using profile management as a profile solution is optional but was used for this Cisco Validated Design, and is described in a later section.)

Instructions	Visual
<p>Insert the XenDesktop ISO and let AutoRun launch the installer.</p> <p>Click Start on the Welcome Screen.</p>	 <p>The image shows the XenDesktop 7.0 'The Avalon Platform' welcome screen. It has a dark green background with white and yellow text. The title 'XenDesktop 7.0' is at the top left. Below it, 'The Avalon Platform' is written in a large, bold, white font. Underneath that, the tagline 'Deliver applications and desktops to any user, anywhere, on any device.' is in yellow. A list of three bullet points follows: 'Flexible application and desktop delivery', 'Centralized management and security', and 'Optimized for deployments of any size'. At the bottom right, there are two buttons: a green 'Start' button and a grey 'Cancel' button.</p>

To install the VDA for the Hosted VDI Desktops, select **Virtual Delivery Agent for Windows Desktop OS**. (After the VDA is installed for Hosted VDI Desktops, repeat the procedure to install the VDA for Hosted Shared Desktops. In this case, select **Virtual Delivery Agent for Windows Server OS** and follow the same basic steps.)



Select **“Users to be able to connect to desktop machines I create from this master image”**.

Click **Next**.



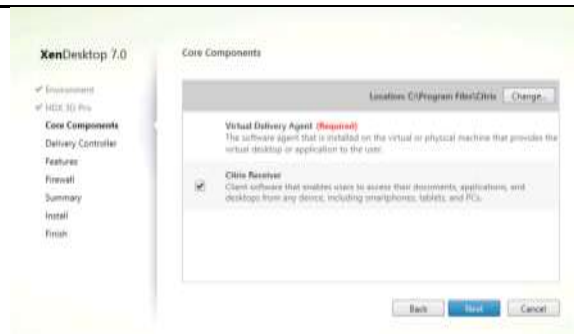
Select **“No, install the standard VDA”**.

Click **Next**.



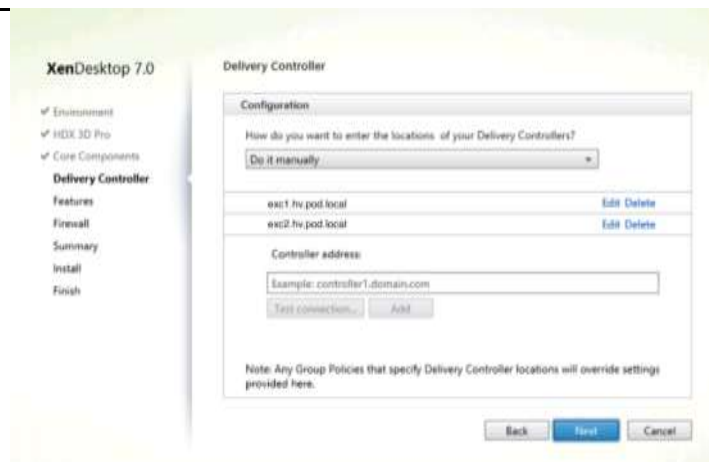
Select **Citrix Receiver**. Note: the Citrix Receiver was not installed in the virtual desktops for the Cisco Validated Design testing.

Click **Next**.



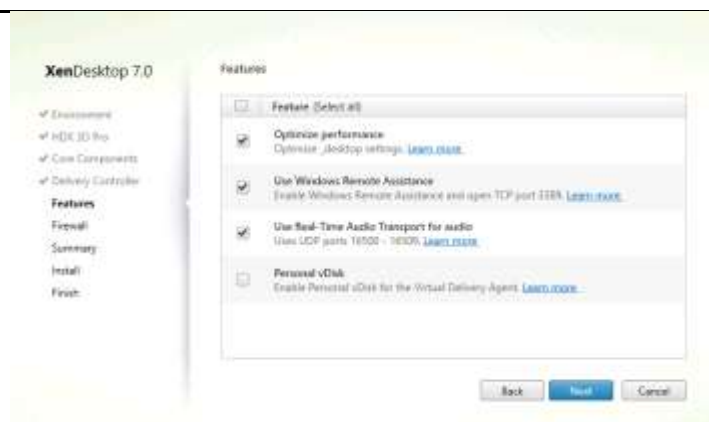
Select **“Do it manually”** and specify the location of your Delivery Controllers (in this case exc1.hv.pod.local and exc2.hv.pod.local).

Click **Next**.



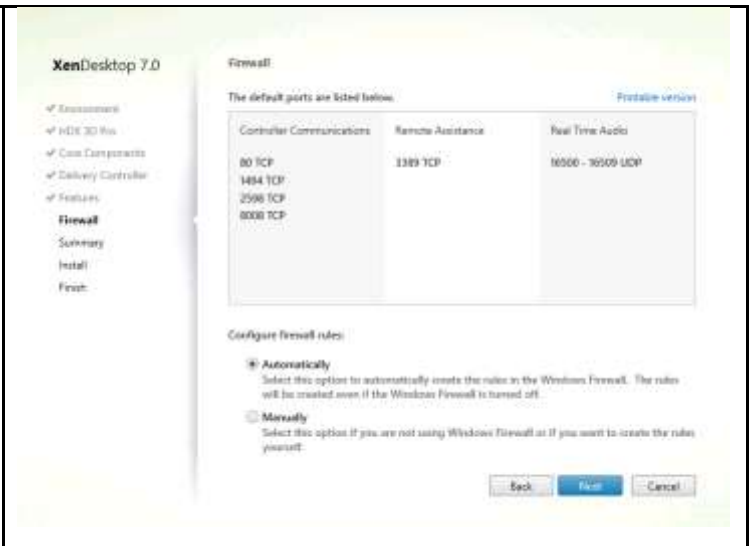
Select the default features.

Click **Next**.

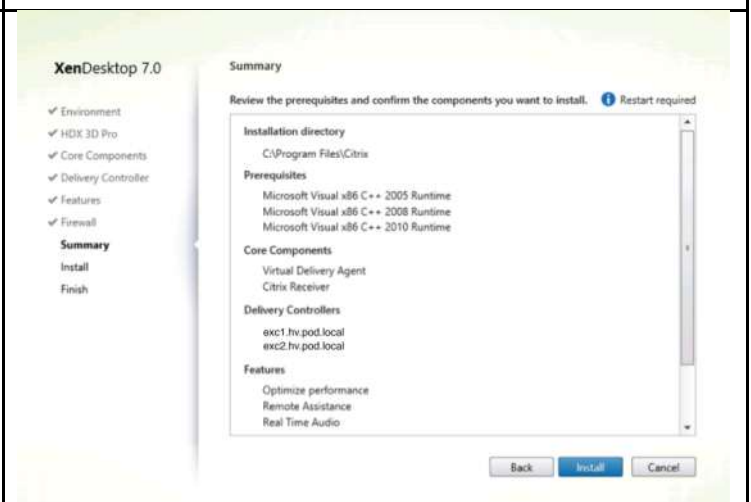


Allow the firewall rules to be configured **Automatically**.

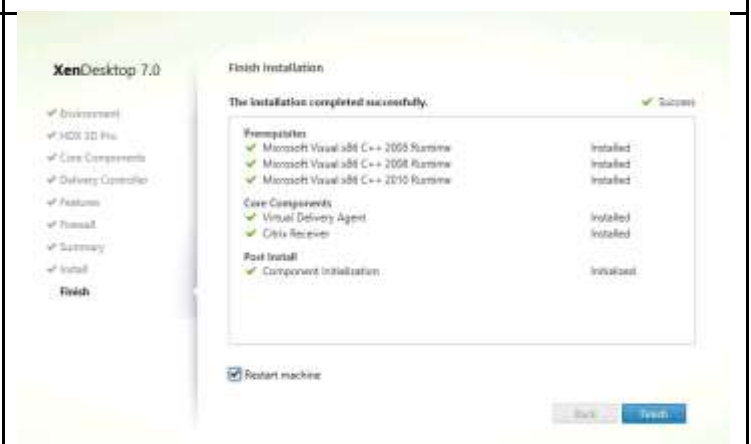
Click **Next**.



Verify the **Summary** and click **Install**.



Click **Finish** and the machine will restart.



Repeat the procedure so that VDAs are installed for both the Hosted VDI Desktops (using the Windows 7 OS image) and the Hosted Shared Desktops (using the Windows Server 2012 image).

7.3. Citrix User Profile Management Servers – CITRIX

Profile management from Citrix provides an easy, reliable, and high-performance way to manage user personalization settings in virtualized or physical Windows environments. It requires minimal infrastructure and administration, and provides users with fast logons and logoffs. A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Examples of settings that can be customized are:

- Desktop settings such as wallpaper and screen saver
- Shortcuts and Start menu setting
- Internet Explorer Favorites and Home Page
- Microsoft Outlook signature
- Printers

Some user settings and data can be redirected by means of folder redirection. However, if folder redirection is not used these settings are stored within the user profile.

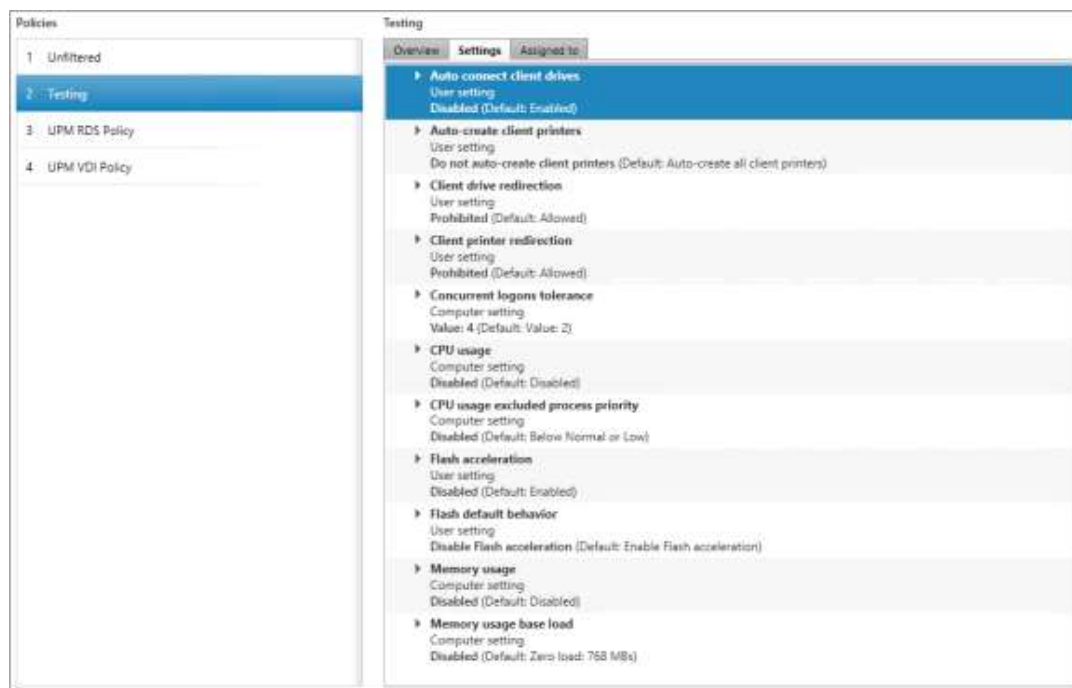
7.3.1. Install and Configuration of User Profile Manager

The first stage in planning a profile management deployment is to decide on a set of policy settings that together form a suitable configuration for your environment and users. The automatic configuration feature simplifies some of this decision-making for XenDesktop deployments. Screenshots of the User Profile Management interfaces that establish policies are shown below. Basic profile management policy settings are documented here: <http://support.citrix.com/proddocs/topic/xendesktop-7/cds-policies-rules-pm-basic-settings.html>.

Several XenDesktop Policy settings were required for the LoginVSI test. These policy settings were configured based on the table (and screenshot) below and were assigned to all domain users.

Setting Name	Setting Value
Auto connect Client Drives	Disabled
Auto-create Client printers	Do not auto-create client printers
Client drive redirection	Prohibited
Client printer redirection	Prohibited
Concurrent logons tolerance	4
CPU usage	Disabled
CPU usage excluded process priority	Disabled
Flash acceleration	Disabled
Flash default behavior	Disable Flash acceleration
Memory usage	Disabled
Memory usage base load	Disabled

Figure 16: XenDesktop Policy



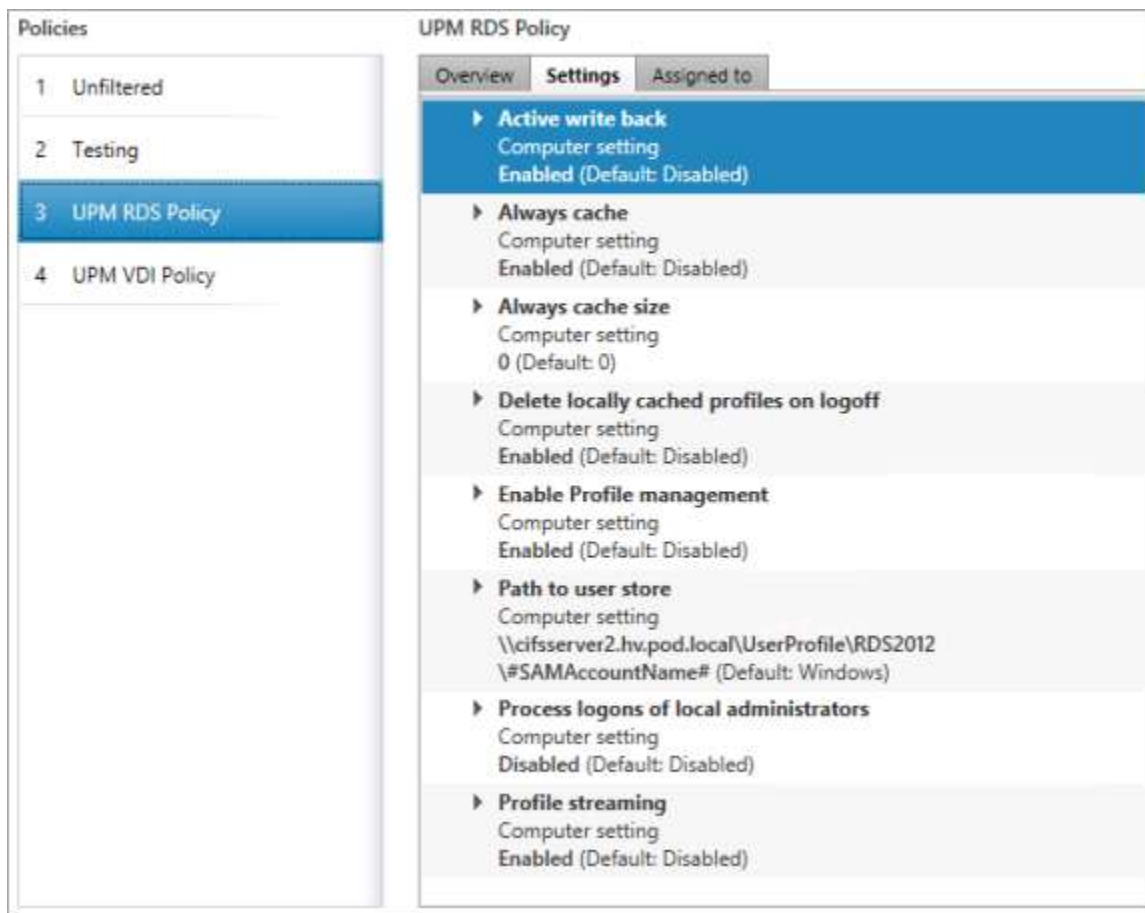
Separate XenDesktop policies were used for the Hosted Shared Desktops and the Hosted Virtual Desktops because the location of the user profiles was different. The policies are outlined below and were assigned to the appropriate delivery group.

Hosted Shared Desktop Users:

Setting Name	Setting Value
Active write back	Enabled
Always cache	Enabled
Always cache size	0
Delete locally cached profiles on logoff	Enabled
Enable Profile management	Enabled
Path to user store	\\cifsserver2.hv.pod.local\UserProfile\RDS2012\#SAMAccountName#
Process logons of local administrators	Disabled
Profile streaming	Enabled

Note: If building the 1000-user configuration, you will need to create a second policy with a different path to the user store on the other storage controller and assign it to another delivery group to split the traffic across the two CIFS shares.

Figure 17: RDS User Profile Manager Policy

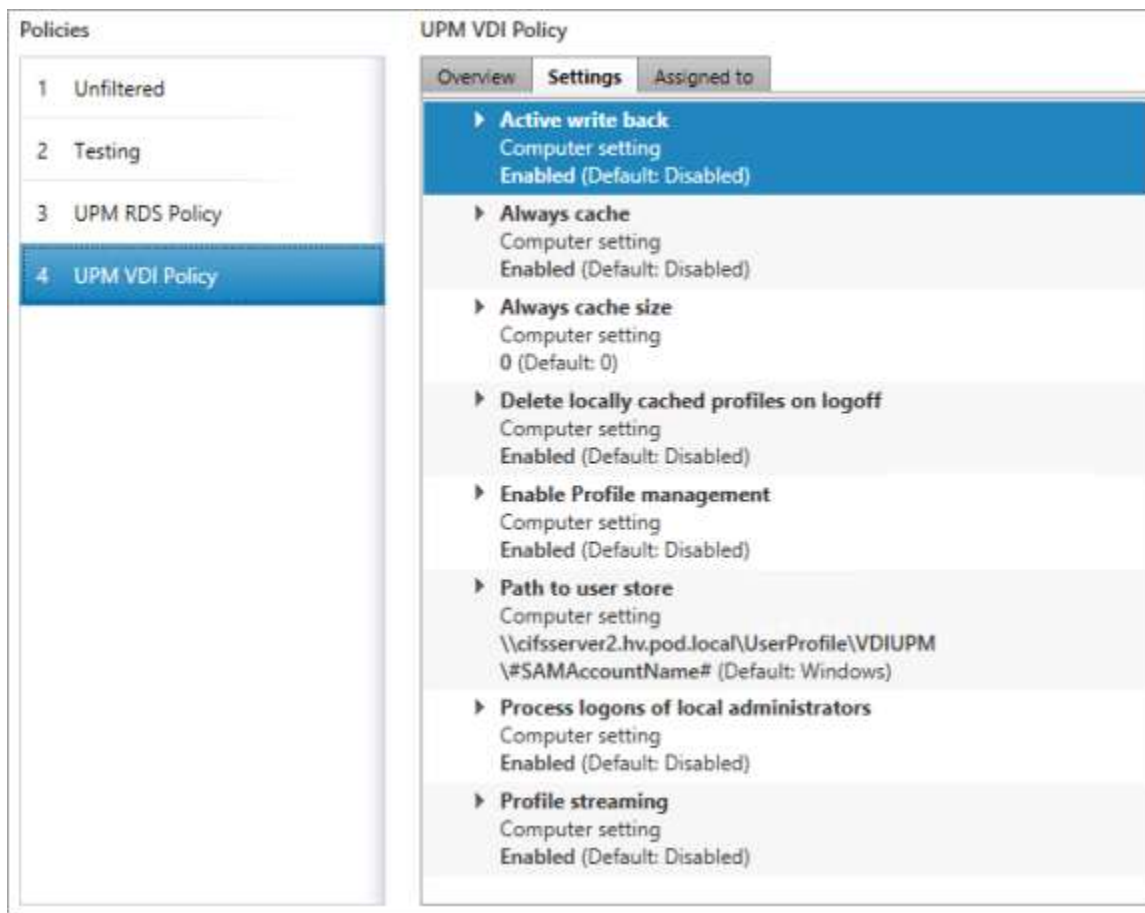


Hosted Virtual Desktops:

Setting Name	Setting Value
Active write back	Enabled
Always cache	Enabled
Always cache size	0
Delete locally cached profiles on logoff	Enabled
Enable Profile management	Enabled
Path to user store	\\cifsserver2.hv.pod.local\UserProfile\VDIUPM\#SMAccountName#
Process logons of local administrators	Enabled
Profile streaming	Enabled

Note: If building the 1000-user configuration, you will need to create a second policy with a different path to the user store on the other storage controller and assign it to another delivery group to split the traffic across the two CIFS shares.

Figure 18: VDI User Profile Manager Policy



7.4. Microsoft Windows 7 and Windows Server 2012 Golden Image Creation

This section provides the guidance around creating the golden, or master images, for the environment. In this case, the images only had the basics added as necessary to run the Login VSI medium workload.

7.4.1. Microsoft Windows 7 and Windows Server 2012 OS Configurations

The master VMs for the Hosted Virtual Desktops and Hosted Shared Desktops were configured as follows:

Table 14: OS Configurations

vDisk Feature	Hosted Virtual Desktops	Hosted Shared Desktops
Virtual CPUs	1 vCPU	6 vCPUs
Dynamic RAM	Startup/Minimum 1.5 GB; Maximum to 4 GB	Startup/Minimum 20 GB, Maximum to 32 GB
vDisk size	17 GB	40 GB
Virtual NICs	2 virtual NICs—one “legacy” NIC for PXE boot and one “synthetic” NIC for OS operation once the VM has booted	Same (2 virtual NICs)


vDisk OS	Microsoft Windows 7 Enterprise (x86)	Microsoft Windows Server 2012
Additional software	Microsoft Office 2010	Microsoft Office 2010
Test workload	Login VSI “medium” workload (knowledge worker)	Login VSI “medium” workload (knowledge worker)



The software installed on each image before cloning the vDisk included:


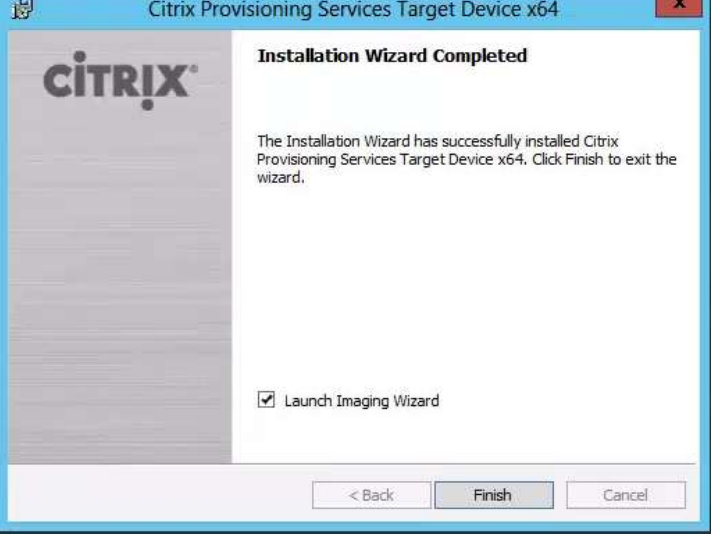
- Citrix Provisioning Server Target Device (32-bit used for HVD and 64-bit used for HSD)
- Microsoft Office Enterprise 2010
- Internet Explorer 8.0.7600.16385 (HVD only; Internet Explorer 10 is included with Windows Server 2012 by default)
- Adobe Reader 9.1.0
- Adobe Flash Player 10.0.22

7.4.2. Installing the PVS Target Device Software

A Master Target Device refers to a target device from which a hard disk image is built and stored on a vDisk. Provisioning Services then streams the contents of the vDisk created from the Master Target Device to other target devices. Follow this procedure to install the PVS Target Device x64 software.



Instructions	Visual
<p>Insert the Provisioning Services ISO and let AutoRun launch the installer.</p> <p>Click Target Device Installation .</p>	



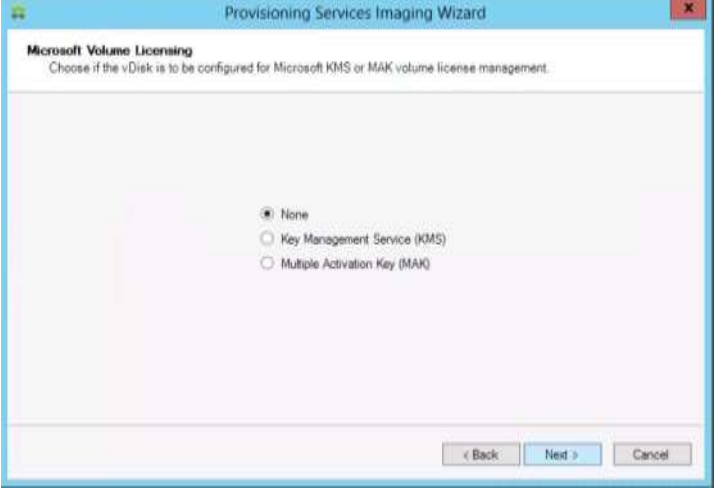
Instructions	Visual
<p>Click Target Device Installation.</p> <p>The installation wizard will check to resolve dependencies and then begin the PVS target device installation process.</p>	 <p>The screenshot shows the Citrix Provisioning Services console. At the top, the Citrix logo and 'Provisioning Services' are displayed. Below this, there is a list of options: 'Target Device Installation' and 'Install Upgrade Wizard'. The 'Target Device Installation' option is highlighted with a red rectangular box. At the bottom of the console, there is a status bar that says 'Install the Target Device.' with 'Back' and 'Exit' buttons.</p>
<p>The Welcome page appears.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Citrix Provisioning Services Target Device x64' setup window. The title bar reads 'Citrix Provisioning Services Target Device x64'. The window has a Citrix logo on the left and a main text area on the right. The text in the main area reads: 'Welcome to the Installation Wizard for Citrix Provisioning Services Target Device x64'. Below this, it says: 'Citrix Provisioning Services Target Device x64 Setup is preparing the Installation Wizard which will guide you through the program setup process. Please wait.' and 'Computing space requirements'. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.</p>

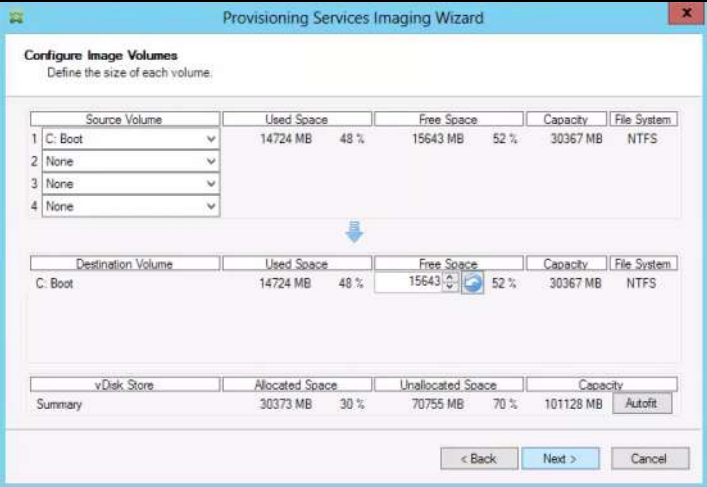
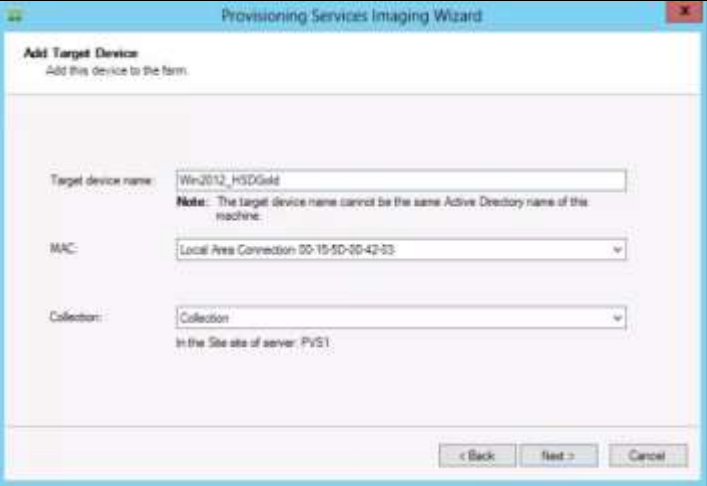
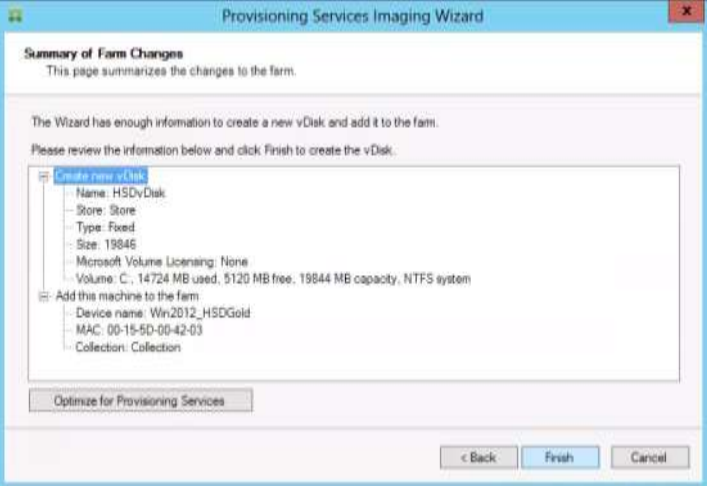
Instructions	Visual
<p>Select the Destination Folder for the PVS Target Device program and click Next. Confirm the Installation settings and Click Install.</p>	
<p>A confirmation screen appears indicating that the installation completed successfully.</p> <p>Click the checkbox to launch the Imaging Wizard and click Finish.</p>	

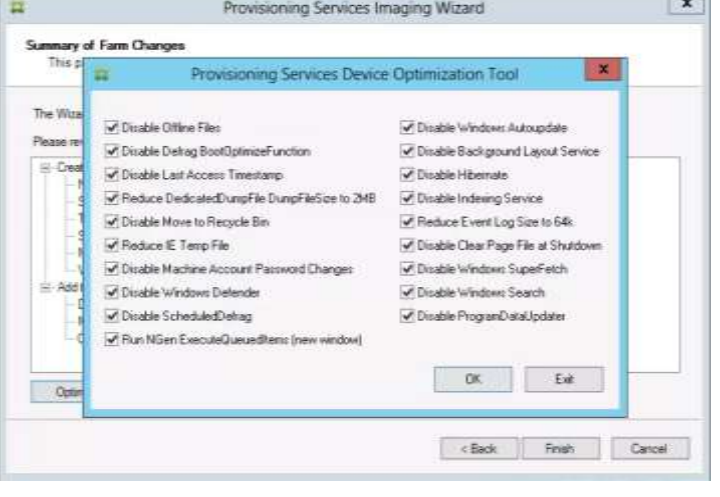
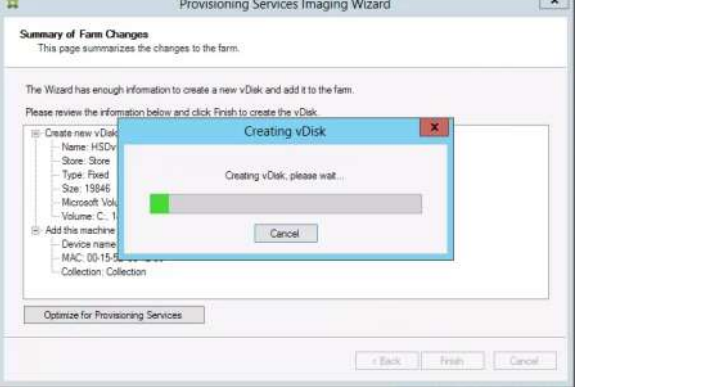
7.4.3. Running the PVS Imaging Wizard

The Imaging Wizard automatically creates the base vDisk image from a master target device.

Instructions	Visual
<p>The Imaging Wizard's Welcome page appears.</p> <p>Click Next.</p>	
<p>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. (For this Cisco Validated Design, the server name entered was PVS1.)</p> <p>Use the Windows credentials (default), or enter different credentials, then click Next. If using Active Directory, enter the appropriate password information.</p> <p>Click Next.</p>	

Instructions	Visual
<p>Select Create new vDisk.</p> <p>Click Next.</p>	
<p>The New vDisk dialog displays. Enter the name of the vDisk, such as XDGold for the Hosted VDI Desktop vDisk (Windows 7 OS image) and HSDGold for the Hosted Shared Desktop vDisk (Windows Server 2012 image). Select the Store where this vDisk will reside. Select the vDisk type, either Fixed or Dynamic, from the drop-down menu. (This Cisco Validated Design used a Fixed rather than Dynamic vDisk.)</p> <p>Click Next.</p>	
<p>From the Microsoft Volume Licensing page, select the volume license option to use for target devices. For this Cisco Validated Design, volume licensing is not used, so the None button is selected.</p> <p>Click Next.</p>	

Instructions	Visual
<p>Define volume sizes on the Configure Image Volumes page.</p> <p>Click Next.</p>	
<p>The Add Target Device page appears.</p> <p>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.</p> <p>Click Next.</p>	
<p>A Summary of Farm Changes appears.</p> <p>Select Optimize for Provisioning Services.</p>	

Instructions	Visual
<p>The PVS Optimization Tool appears. Select the appropriate optimizations and click OK.</p> <p>Review the configuration and click Finish.</p>	
<p>The vDisk creation process begins. A dialog appears when the creation process is complete. To continue the imaging process, you must reboot and configure the BIOS/VM settings for PXE/network boot.</p>	

Repeat the procedure to create vDisks for both the Hosted VDI Desktops (using the Windows 7 OS image) and the Hosted Shared Desktops (using the Windows Server 2012 image).

7.4.4. Installation of Login VSI Software

Tests were performed using Login VSI 3.7 (<http://www.loginvsi.com>), a load generation tool for VDI benchmarking that simulates production user workloads to generate desktop workloads and gather data about VDI performance. All tests were done using the default Medium workload to simulate the desktop activity of a typical knowledge worker. Login VSI generates an office productivity workload that includes Microsoft Office 2010 with Microsoft Outlook, Word, PowerPoint, and Excel, Internet Explorer with a Flash video applet, Java app, and Adobe® Acrobat® Reader. All applications are included for testing with the Login VSI software distribution except for Microsoft Office, which must be installed separately in the OS images.

We used standard Login VSI installation instructions (<http://www.loginvsi.com/documentation/v3/installation>) and best practices in preparing the master images.

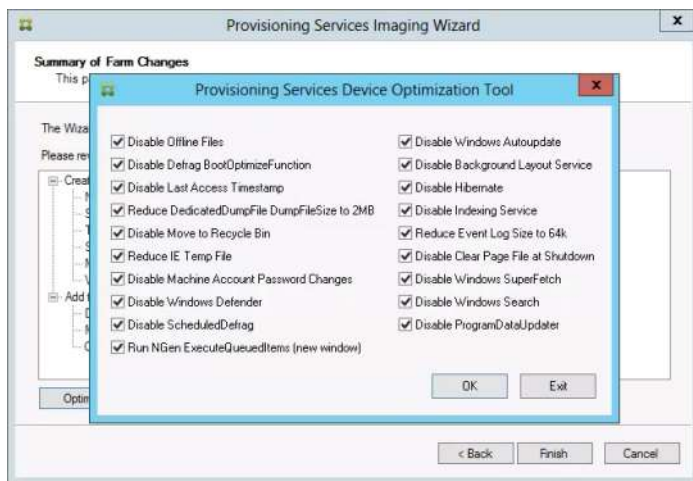
Follow these installation procedures to install the software:

1. [Download](#) the Login VSI archive file (VSI37.exe).
2. Unpack the Login VSI archive file. This file contains:
 - AD Setup
 - Target Setup
 - Launcher Setup
 - Analyzer Setup
3. [Configure the Active Directory with the AD Setup.](#)
4. [Configure a File Share \(VSIshare\) for logging.](#)
5. [Prepare the launcher workstations with the Launcher Setup.](#)
6. [Prepare the target platform with the Target Setup.](#)
7. [Install the analyzer with the Analyzer Setup.](#)

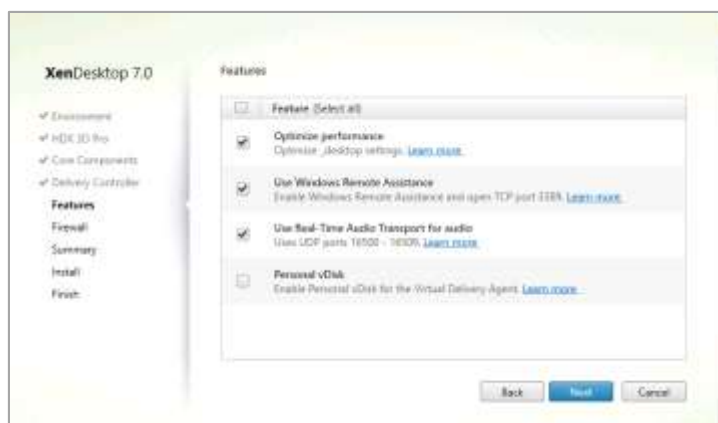
7.4.5. Optimization for PVS and XenDesktop 7

In the process of running the PVS Imaging Wizard, the PVS Optimization Tool appears. Be sure to select the default optimizations.

Figure 19: XenDesktop Optimizations



The Virtual Desktop Agent also includes default optimizations during installation. When this feature is enabled, the optimization tool is used for VDAs running in a VM on a hypervisor. VM optimization includes disabling offline files, disabling background defragmentation, and reducing event log size. For more information, see [CTX125874](#).



7.4.6. Conversion to PVS vDisk

After installing and configuring PVS components, 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. vDisks can exist on a Provisioning Server, file share, or in larger deployments (as in this Cisco Validated Design), on a storage system with which the Provisioning Server can communicate (through iSCSI, SAN, NAS, and CIFS). vDisks can be assigned to a single target device in Private Image Mode, or to multiple target devices in Standard Image Mode.

7.4.7. Write-Cache Drive Sizing and Placement

When considering a PVS deployment, there are some design decisions that need to be made regarding the write cache for the virtual desktop devices that leverage provisioning services. The write cache is a cache of all data that the target device has written. If data is written to the PVS vDisk in a caching mode, the data is not written back to the base vDisk. Instead it is written to a write cache file.

Note: It is important to consider Write Cache sizing and placement when scaling virtual desktops using PVS server.

There are several options as to where the Write Cache can be placed, such as on the PVS server, in hypervisor RAM, or on a device local disk (this is usually an additional vDisk for VDI instances). For this study, we used PVS 7 to manage desktops with write cache placed on the device local disk of each virtual machine, which allows the design to scale more effectively. For optimal performance, write cache files were stored on SSDs located each of the virtual desktop host servers.

For Citrix PVS pooled desktops, write cache size needs to be calculated based on how often the user reboots the desktop and type of applications used. We recommend using a write cache twice the size of RAM allocated to each individual VM. For example, if VM is allocated with 1.5GB RAM, use at least a 3GB write cache vDisk for each VM.

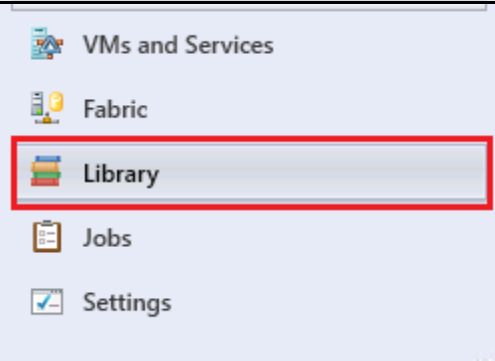
For this solution, 6GB virtual disks were assigned to the Windows 7-based virtual machines used in the desktop creation process. The PVS Target device agent installed in the Windows 7 gold image automatically places the Windows swap file on the same drive used by the PVS Write Cache when this mode is enabled. 50GB write cache virtual disks were used for the Server 2012 desktop machines.

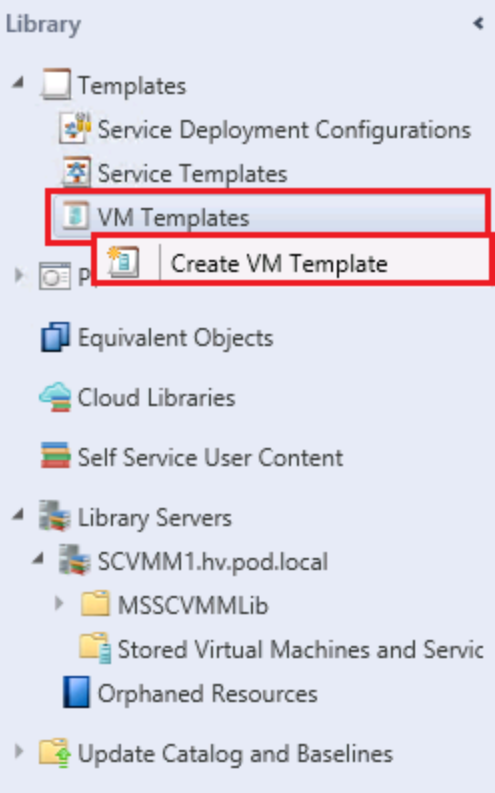
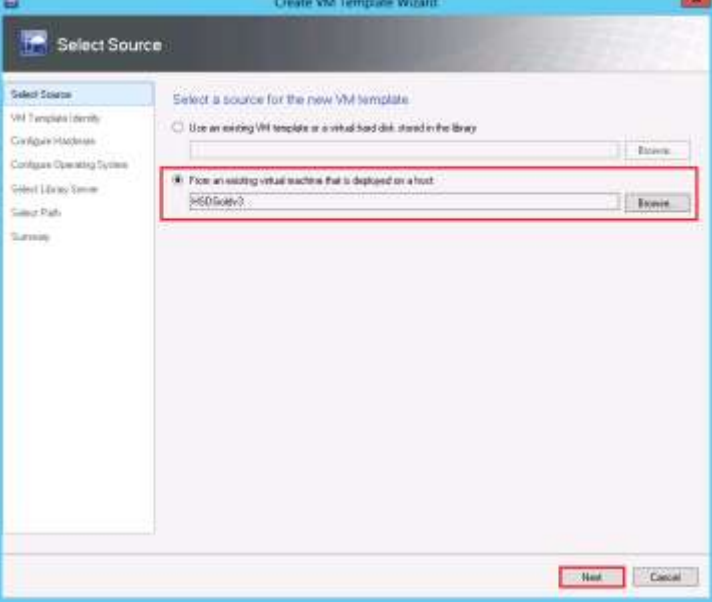
7.5. Citrix Provisioning Services

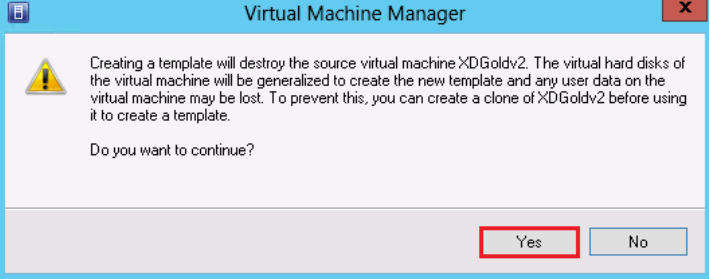
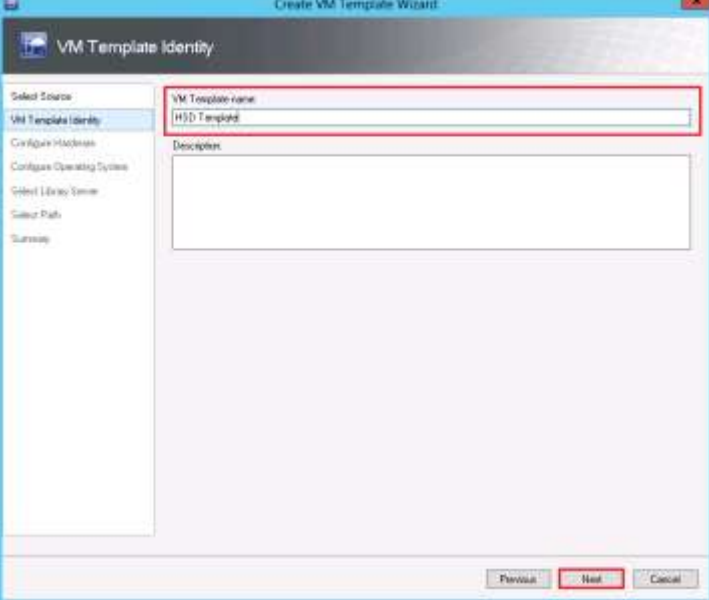
7.5.1. Creating the Virtual Machine Manager Templates

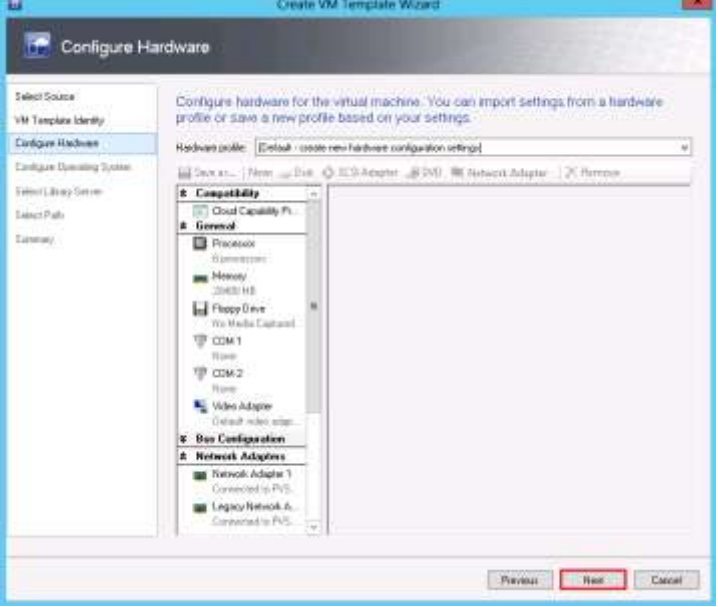
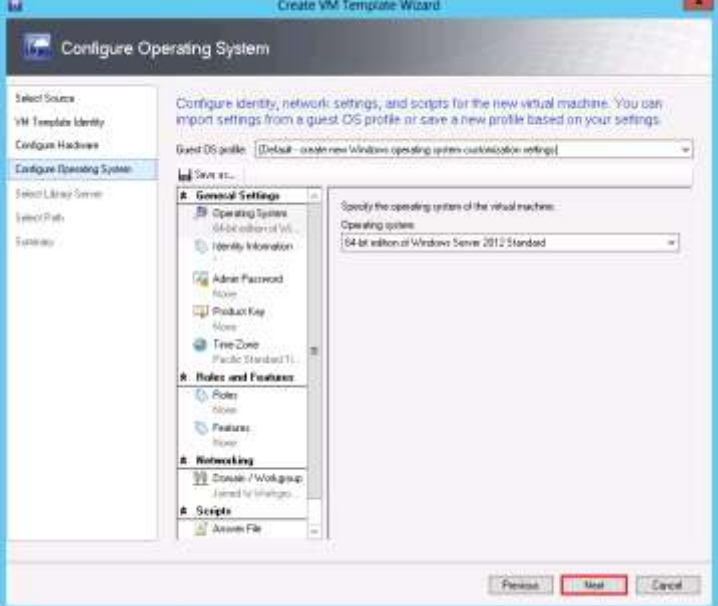
This section outlines the process for creating a template in System Center Virtual Machine Manager for both the hosted shared desktop and the hosted virtual desktop virtual machines. The template is required for the XenDesktop Setup Wizard.

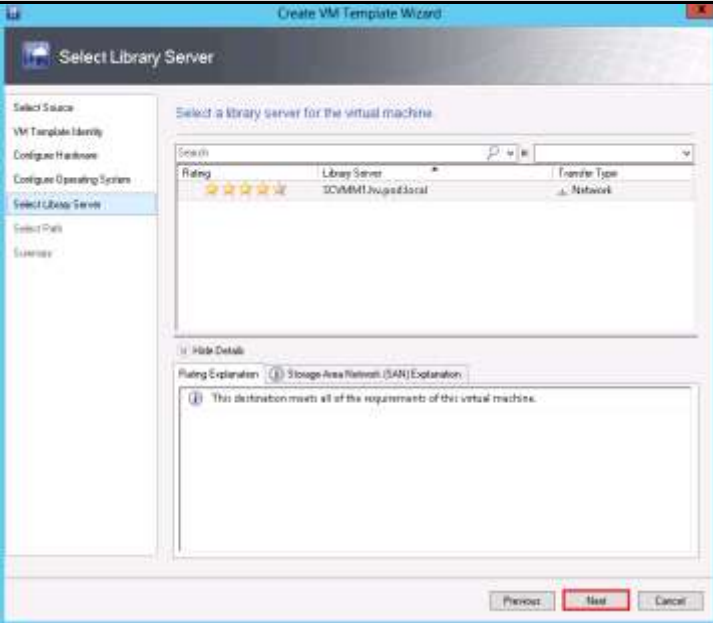
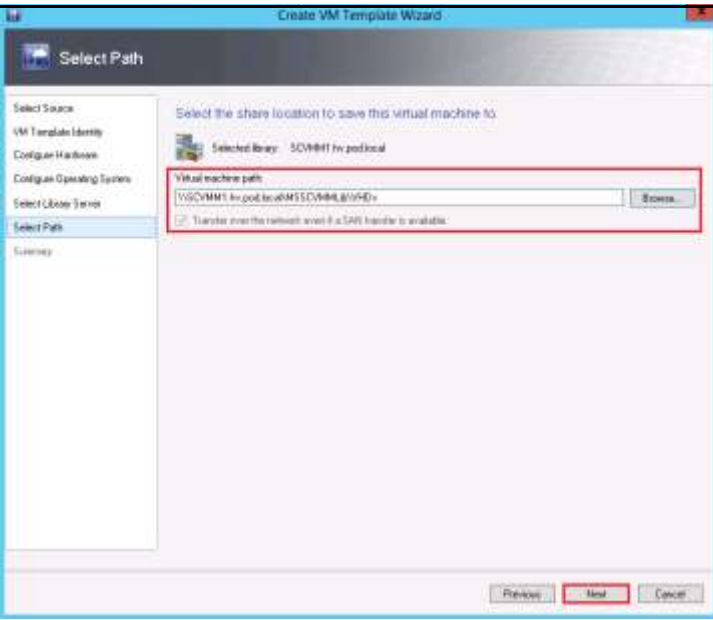
Before creating the Virtual Machine Manager template, first clone the existing Virtual Machines (Hosted Shared Desktop and Hosted Virtual Desktop) because the template creation process destroys the source virtual machine.

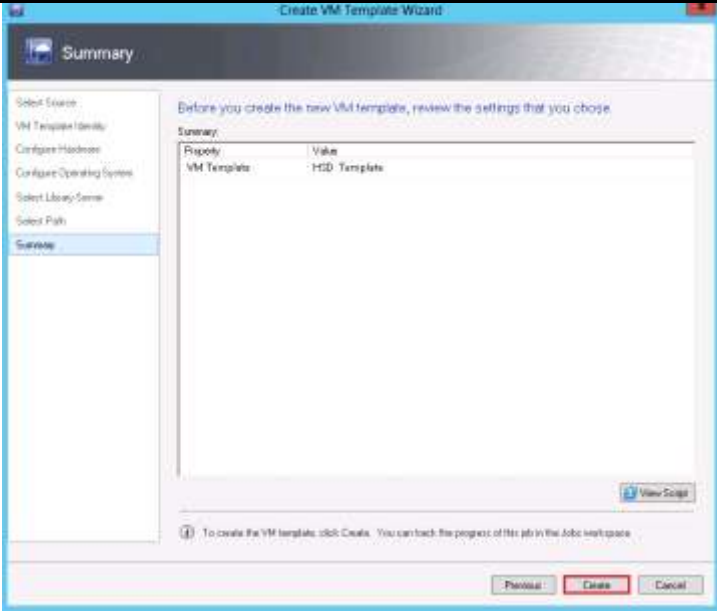
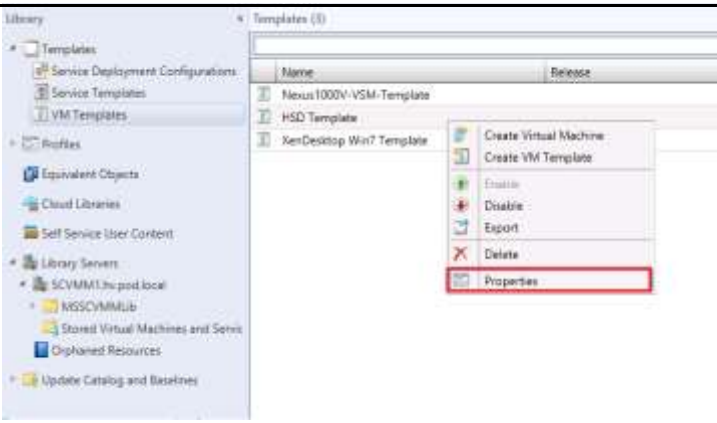
Instructions	Visual
Launch the System Center Virtual Machine Manager console.	
Select the Library tab.	 A screenshot of the System Center Virtual Machine Manager console. The left-hand navigation pane is visible, showing a list of tabs: 'VMs and Services', 'Fabric', 'Library', 'Jobs', and 'Settings'. The 'Library' tab is highlighted with a red rectangular border, indicating it is the selected option.

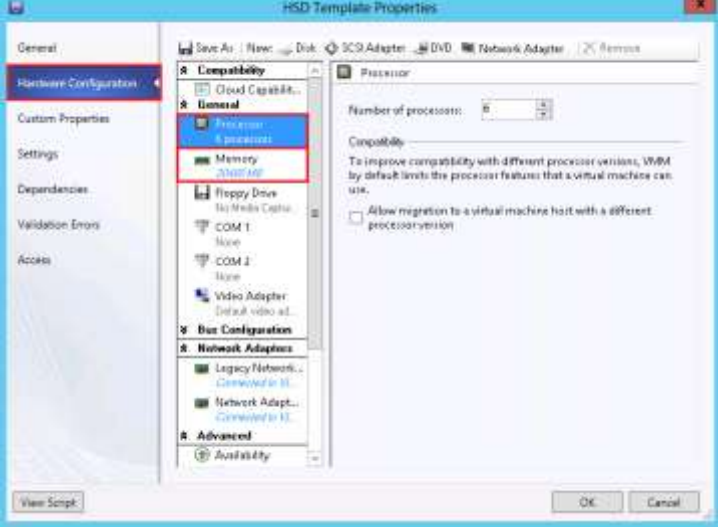
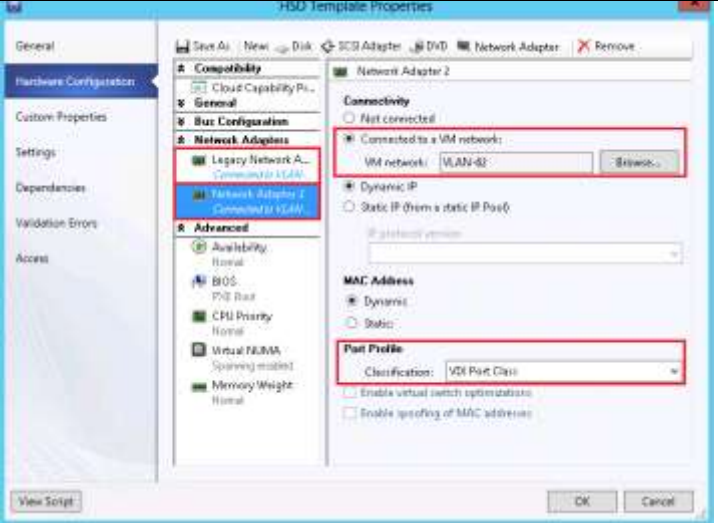
Instructions	Visual
<p>Select the VM Templates node and right-click.</p> <p>From the context menu, choose Create VM Template.</p>	 <p>The screenshot shows the 'Library' pane in vSphere. The 'VM Templates' node is selected and right-clicked, opening a context menu. The 'Create VM Template' option is highlighted with a red rectangle.</p>
<p>Select the previously created template Virtual machine.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Create VM Template Wizard' window, specifically the 'Select Source' step. The 'Select Source' list on the left has 'VM Template Identity' selected. The main area shows two radio button options: 'Use an existing VM template or a virtual hard disk stored in the library' (unselected) and 'Pick an existing virtual machine that is deployed on a host' (selected). A red rectangle highlights the 'Pick an existing virtual machine that is deployed on a host' option and the 'Browse...' button next to it. The 'Next' and 'Cancel' buttons are at the bottom right.</p>

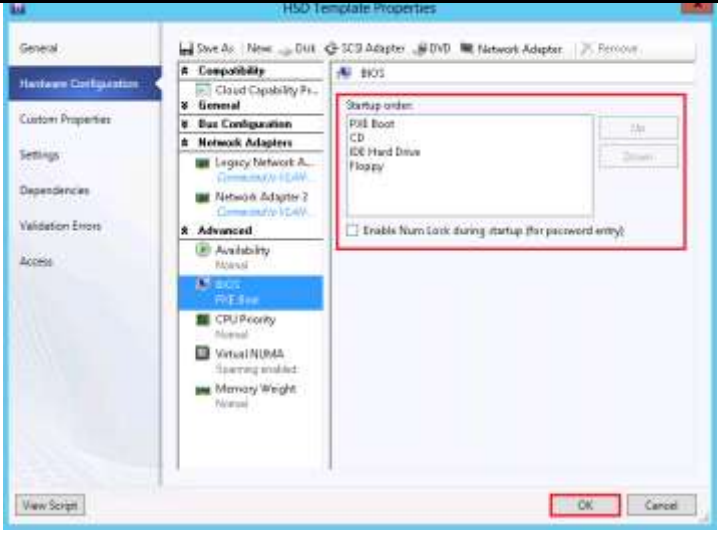
Instructions	Visual
<p>Click Yes to acknowledge your existing template will be generalized.</p> <p>Note: This is why a copy was made prior to creating the template. In the end, the disk itself will be replaced with a PVS write-cache VHDX.</p>	 <p>The screenshot shows a 'Virtual Machine Manager' dialog box with a yellow warning icon. The text inside reads: 'Creating a template will destroy the source virtual machine XDGoldv2. The virtual hard disks of the virtual machine will be generalized to create the new template and any user data on the virtual machine may be lost. To prevent this, you can create a clone of XDGoldv2 before using it to create a template. Do you want to continue?'. At the bottom right, there are two buttons: 'Yes' (highlighted with a red box) and 'No'.</p>
<p>Provide a template name.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Create VM Template Wizard' window, specifically the 'VM Template Identity' step. On the left is a navigation pane with steps: 'Select Source', 'VM Template Identity' (selected), 'Configure Hardware', 'Configure Operating System', 'Select Library/Server', 'Select Path', and 'Summary'. The main area has a 'VM Template name' text box containing 'H5D Template' (highlighted with a red box) and a larger 'Description' text box below it. At the bottom right, there are three buttons: 'Previous', 'Next' (highlighted with a red box), and 'Cancel'.</p>

Instructions	Visual
<p>From the Configure Hardware tab click Next.</p> <p>The actual configuration will be adjusted to match later.</p>	 <p>The screenshot shows the 'Configure Hardware' tab of the 'Create VM Template Wizard'. The left sidebar lists the steps: 'Select Source', 'VM Template Identity', 'Configure Hardware' (selected), 'Configure Operating System', 'Select Library Server', 'Select Path', and 'Summary'. The main area is titled 'Configure Hardware' and contains a 'Hardware profile' dropdown set to '[Default - create new hardware configuration settings]'. Below this are several expandable sections: 'Compatibility' (with 'Cloud Capability P...', 'Processors', 'Memory', 'Disk', 'Floppy Drive', 'CD/DVD', 'CD-ROM', and 'Video Adapter'), 'Base Configuration', 'Network Adapters', and 'Network Adapter 1'. At the bottom are 'Previous', 'Next' (highlighted in red), and 'Cancel' buttons.</p>
<p>From the Configure Operating System tab, click Next.</p>	 <p>The screenshot shows the 'Configure Operating System' tab of the 'Create VM Template Wizard'. The left sidebar is the same as the previous screenshot, with 'Configure Operating System' now selected. The main area is titled 'Configure Operating System' and contains a 'Guest OS profile' dropdown set to '[Default - create new Windows operating system customization settings]'. Below this are several expandable sections: 'General Settings' (with 'Operating System', 'Identity Information', 'Admin Password', 'Product Key', 'Time Zone', and 'Roles and Features'), 'Networking', and 'Script'. At the bottom are 'Previous', 'Next' (highlighted in red), and 'Cancel' buttons.</p>

Instructions	Visual
<p>From the Select Library Server, click Next</p>	
<p>Provide the Virtual Machine path to store in the VMM library.</p> <p>Click Next.</p>	

Instructions	Visual
<p>From the Summary tab, click Create to finish the creation process for the VMM template.</p>	
<p>From the VM Templates node, right-click on the newly created template. Choose Properties from the context menu.</p>	

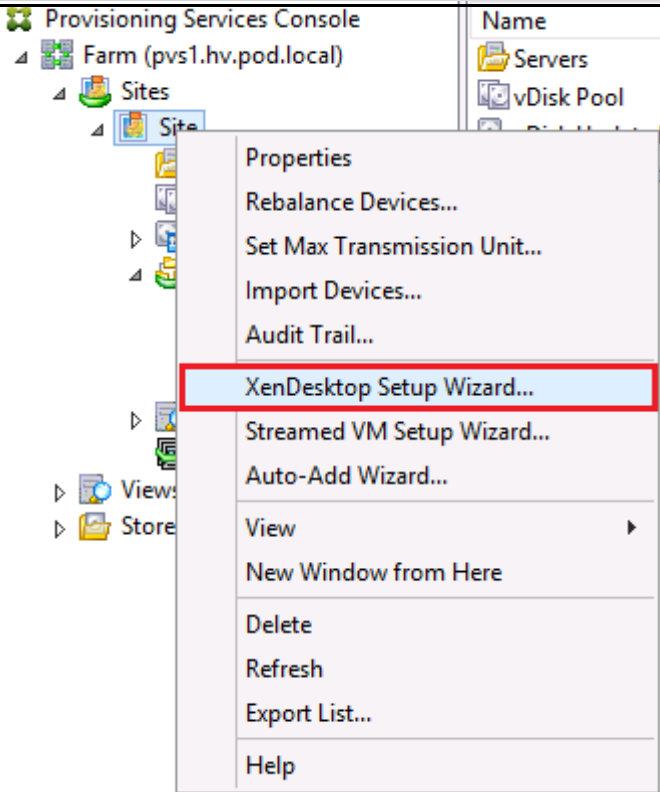
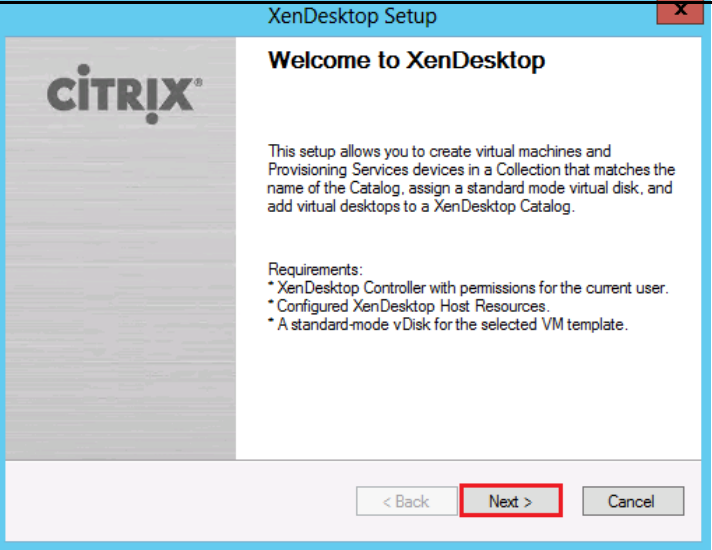
Instructions	Visual
<p>Select the Hardware Configuration tab.</p> <p>Verify the Processor and Memory configurations.</p> <p><u>Hosted Shared Desktops</u>: 6 vCPU with 20GB Static Memory or Dynamic Memory 20Gb minimum and startup with 32 GB maximum.</p> <p><u>Hosted Virtual Desktops</u>: 1 vCPU with Dynamic Memory 1.5GB minimum and startup with 4GB maximum.</p>	
<p>Set the Network Adapter and Legacy Network Adapter to use the Cisco Nexus 1000V VM Network and the Port Profile created earlier.</p>	

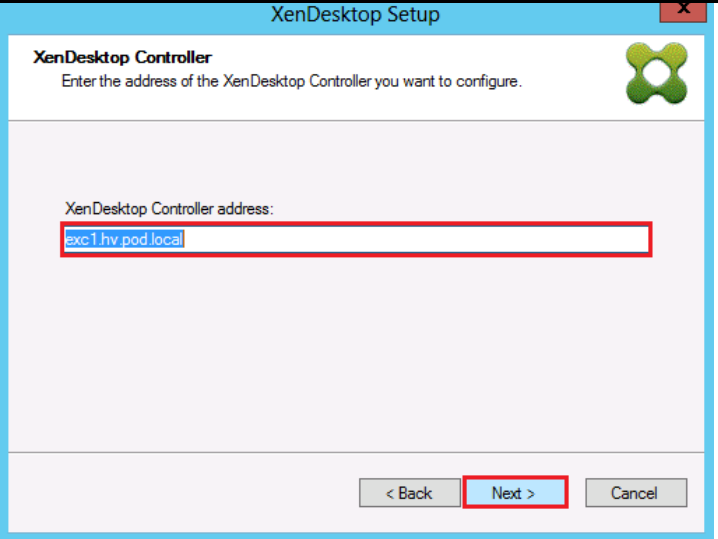
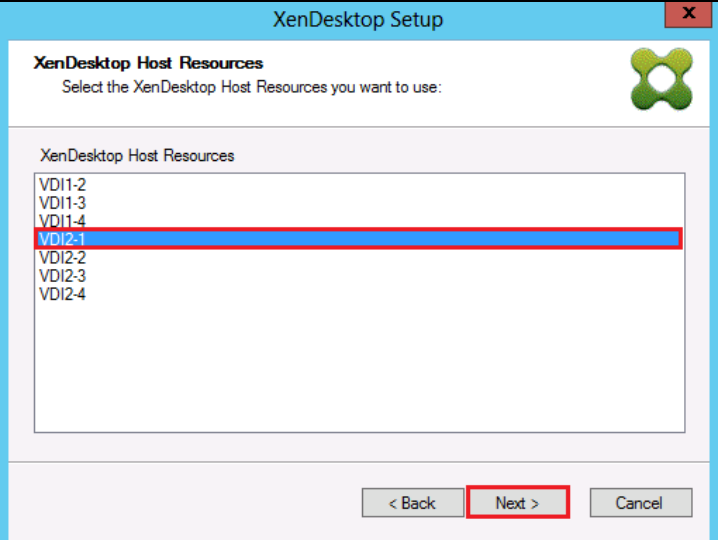
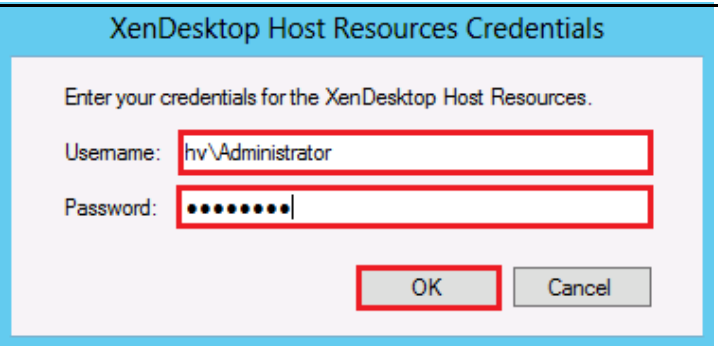
Instructions	Visual
<p>Select the BIOS setting and configure PXE Boot to be at the top of the list.</p> <p>Note: The XenDesktop wizard will actually move CD to the top, but this step ensures that PXE Boot is no longer at the bottom of the list.</p> <p>Click OK to save the template changes.</p>	
<p>Repeat the template process for the Hosted Virtual Desktop template.</p>	

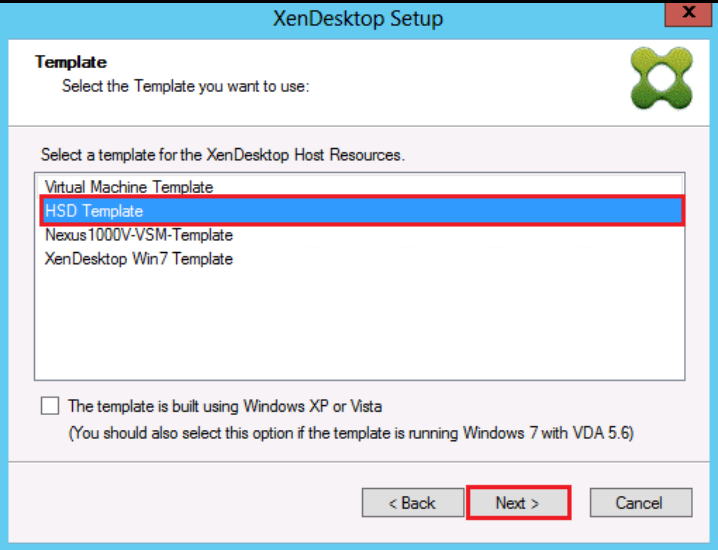
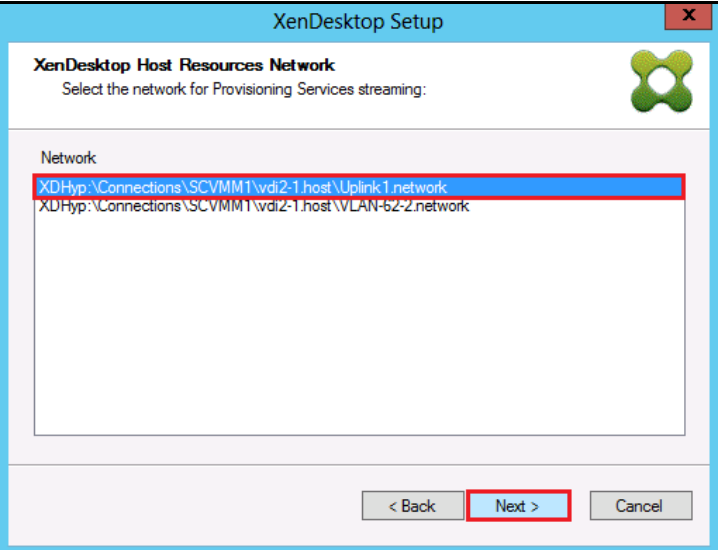
7.5.2. Process to Create Virtual Desktops using PVS Wizard

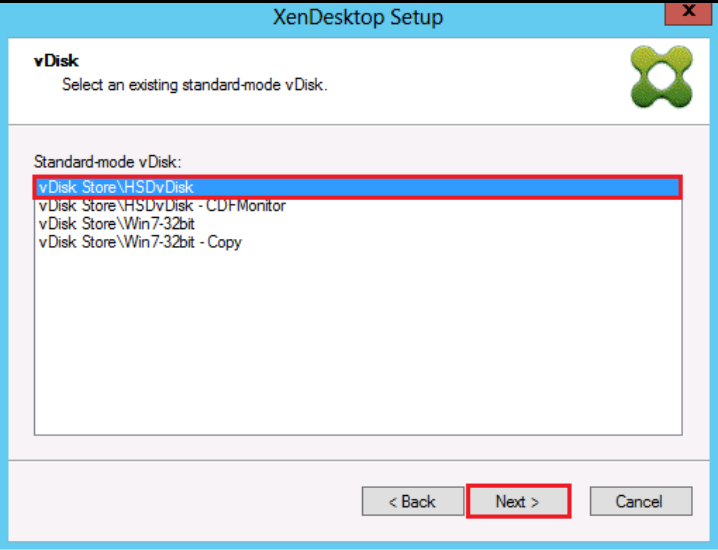
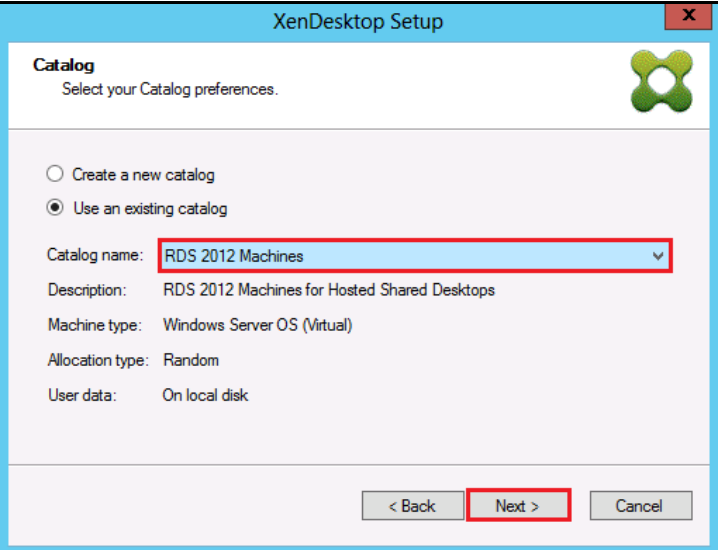
Using the Provisioning Services XenDesktop Setup Wizard provides an automated way to create the virtual machines, but it does not currently support the Nexus 1000V logical switch for networking or enable dynamic memory, so after creating the virtual machines a PowerShell script will need to be run to update the virtual machines to use the Nexus 1000V logical switch. The XenDesktop wizard can make the virtual machines as long as the network cards are on a Microsoft VirtualSwitch instead of the Nexus 1000V. Then the PowerShell script found in section 13.4.1, Update Virtual Machines created by XenDesktop Wizard of the Sample PowerShell Scripts section, can be used to move the VMs to the Cisco Nexus 1000V networks.

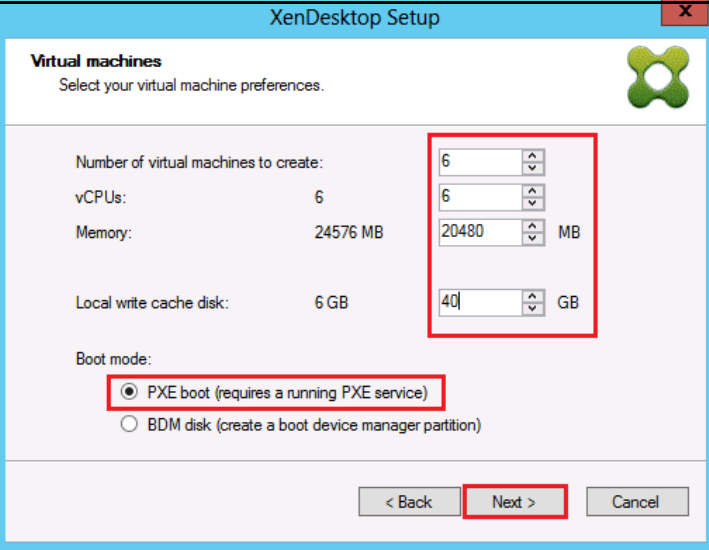
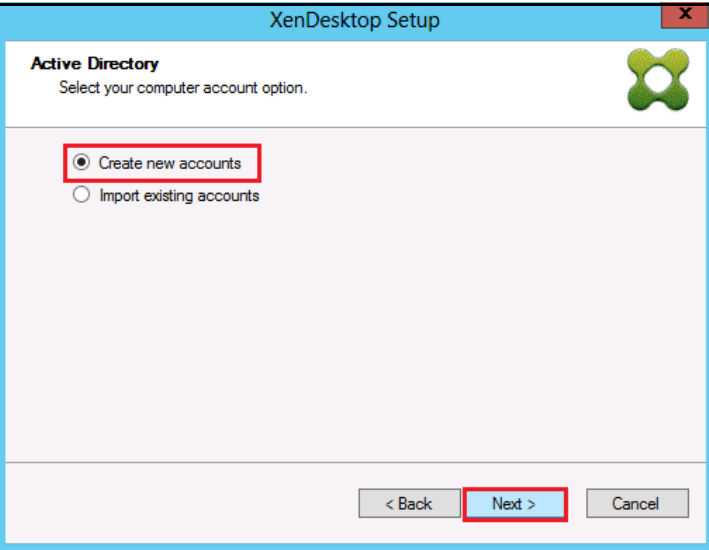
Note: A hotfix that supports the Nexus 1000V switch may be available for the PVS XenDesktop Wizard. If so, the use of the Microsoft Standard switch is not necessary and the Cisco Nexus 1000V networks will work.

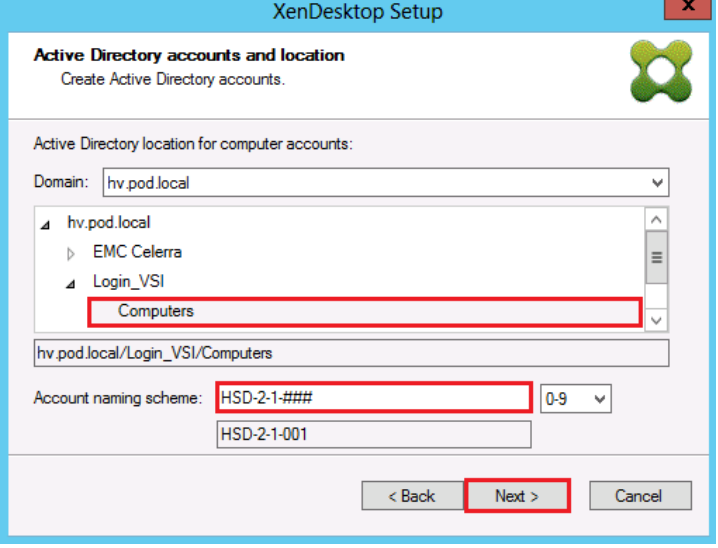
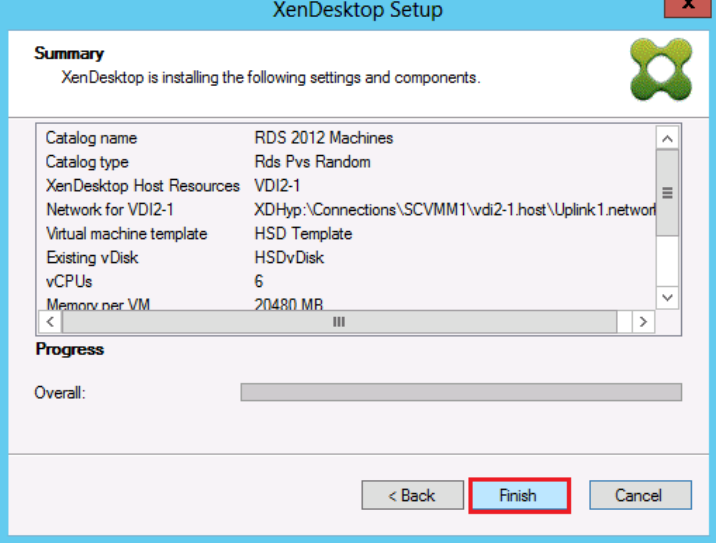
Instructions	Visual
<p>Start the XenDesktop Wizard from the Provisioning Services Console.</p> <p>Right-click on the Site.</p> <p>Choose XenDesktop Setup Wizard... from the context menu.</p>	 <p>The screenshot shows the Provisioning Services Console interface. In the left pane, a tree view shows 'Farm (pvs1.hv.pod.local)' expanded to 'Sites', with a 'Site' selected. A right-click context menu is open over the 'Site'. The menu items are: Properties, Rebalance Devices..., Set Max Transmission Unit..., Import Devices..., Audit Trail..., XenDesktop Setup Wizard... (highlighted with a red box), Streamed VM Setup Wizard..., Auto-Add Wizard..., View (with a sub-menu arrow), New Window from Here, Delete, Refresh, Export List..., and Help. The right pane shows a table with columns 'Name' and 'Servers', containing a row for 'vDisk Pool'.</p>
<p>From the opening dialog, click Next.</p>	 <p>The screenshot shows the 'XenDesktop Setup' dialog box. It has a blue title bar and a light blue border. On the left is the Citrix logo. The main area contains the text 'Welcome to XenDesktop' and a paragraph: 'This setup allows you to create virtual machines and Provisioning Services devices in a Collection that matches the name of the Catalog, assign a standard mode virtual disk, and add virtual desktops to a XenDesktop Catalog.' Below this, under 'Requirements:', there is a bulleted list: '• XenDesktop Controller with permissions for the current user.', '• Configured XenDesktop Host Resources.', and '• A standard-mode vDisk for the selected VM template.' At the bottom, there are three buttons: '< Back', Next > (highlighted with a red box), and 'Cancel'.</p>

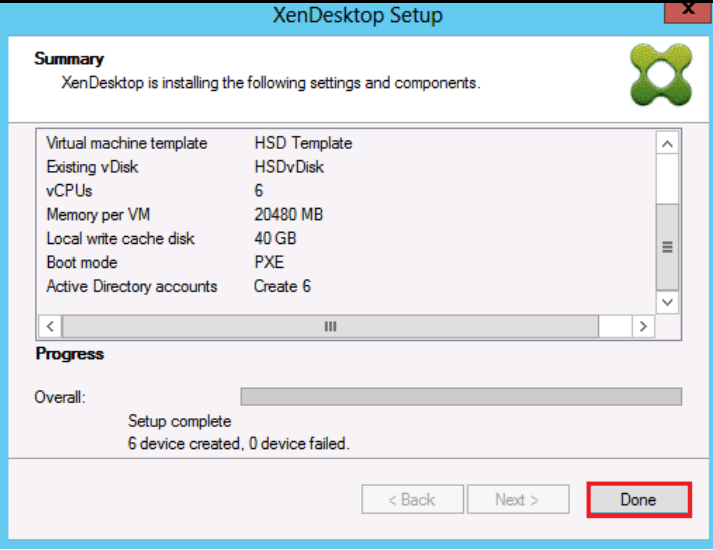
Instructions	Visual
<p>Enter the XenDesktop Controller address that will be used for the wizard operations.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'XenDesktop Setup' window with the 'XenDesktop Controller' tab selected. The instruction 'Enter the address of the XenDesktop Controller you want to configure.' is displayed. Below it, the text 'XenDesktop Controller address:' is followed by a text input field containing 'exc1.hv.pod.local'. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red rectangle.</p>
<p>Select the Hyper-V host where the virtual machines will be created.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'XenDesktop Setup' window with the 'XenDesktop Host Resources' tab selected. The instruction 'Select the XenDesktop Host Resources you want to use:' is displayed. Below it, a list of host resources is shown: VDI1-2, VDI1-3, VDI1-4, VDI2-1, VDI2-2, VDI2-3, and VDI2-4. The 'VDI2-1' entry is highlighted with a blue selection bar. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a red rectangle.</p>
<p>Provide the Authentication credentials (username and password) to the XenDesktop controller when prompted.</p> <p>Click OK.</p>	 <p>The screenshot shows the 'XenDesktop Host Resources Credentials' dialog box. The instruction 'Enter your credentials for the XenDesktop Host Resources.' is displayed. Below it, there are two input fields: 'Username:' containing 'hv\Administrator' and 'Password:' containing a masked password (represented by dots). At the bottom, there are two buttons: 'OK' and 'Cancel'. The 'OK' button is highlighted with a red rectangle.</p>

Instructions	Visual
<p>Select the Hosted Shared Desktop template (HSD Template) created earlier.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'XenDesktop Setup' window with the 'Template' tab selected. The instruction 'Select the Template you want to use:' is at the top. Below it, a list of templates is shown: 'Virtual Machine Template', 'HSD Template' (highlighted with a red border), 'Nexus1000V-VSM-Template', and 'XenDesktop Win7 Template'. At the bottom, there is a checkbox for 'The template is built using Windows XP or Vista' and navigation buttons: '< Back', 'Next >' (highlighted with a red border), and 'Cancel'.</p>
<p>Select the standard VM network switch, not the Nexus 1000V logical network*.</p> <p>Note: Selecting the Nexus 1000V network will result in a failed deployment. The correct logical switch will be updated later using PowerShell script.</p> <p>Click Next.</p> <p>*A private hotfix should be available from Citrix to fix this issue by the time this Cisco Validated Design is published.</p>	 <p>The screenshot shows the 'XenDesktop Setup' window with the 'XenDesktop Host Resources Network' tab selected. The instruction 'Select the network for Provisioning Services streaming:' is at the top. Below it, a list of networks is shown: 'XDHyp:\Connections\SCVMM1\vdz2-1.host\Uplink1.network' (highlighted with a red border) and 'XDHyp:\Connections\SCVMM1\vdz2-1.host\VLAN-b2-2.network'. At the bottom, there are navigation buttons: '< Back', 'Next >' (highlighted with a red border), and 'Cancel'.</p>

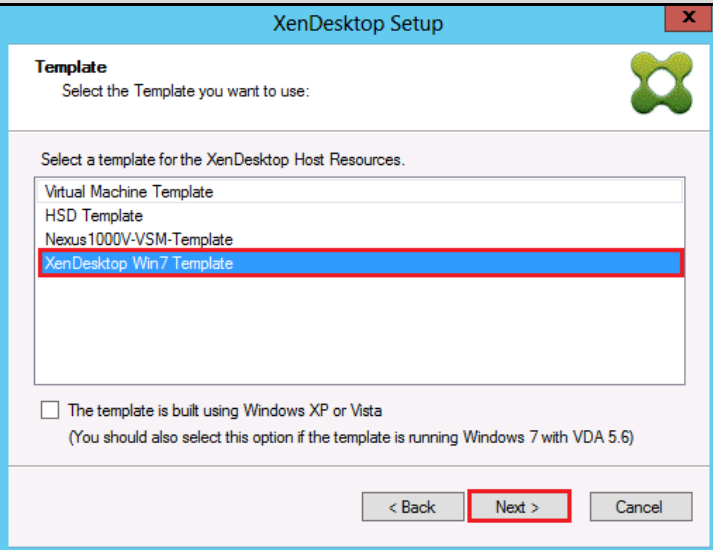
Instructions	Visual
<p>Select the vDisk that will be used to stream to the virtual machine.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'vDisk' selection window in XenDesktop Setup. The title bar says 'XenDesktop Setup'. Below the title bar, it says 'vDisk' and 'Select an existing standard-mode vDisk.' There is a list of vDisks: 'vDisk Store\HSDvDisk', 'vDisk Store\HSDvDisk - CDFMonitor', 'vDisk Store\Win7-32bit', and 'vDisk Store\Win7-32bit - Copy'. The first vDisk is highlighted with a red box. At the bottom, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>
<p>Select an existing catalog or choose to create a new catalog.</p> <p>Note: The catalog name is also used as the collection name in PVS site.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Catalog' selection window in XenDesktop Setup. The title bar says 'XenDesktop Setup'. Below the title bar, it says 'Catalog' and 'Select your Catalog preferences.' There are two radio buttons: 'Create a new catalog' and 'Use an existing catalog' (which is selected). Below the radio buttons, there is a 'Catalog name' dropdown menu with 'RDS 2012 Machines' selected (highlighted with a red box). Other fields include 'Description: RDS 2012 Machines for Hosted Shared Desktops', 'Machine type: Windows Server OS (Virtual)', 'Allocation type: Random', and 'User data: On local disk'. At the bottom, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>

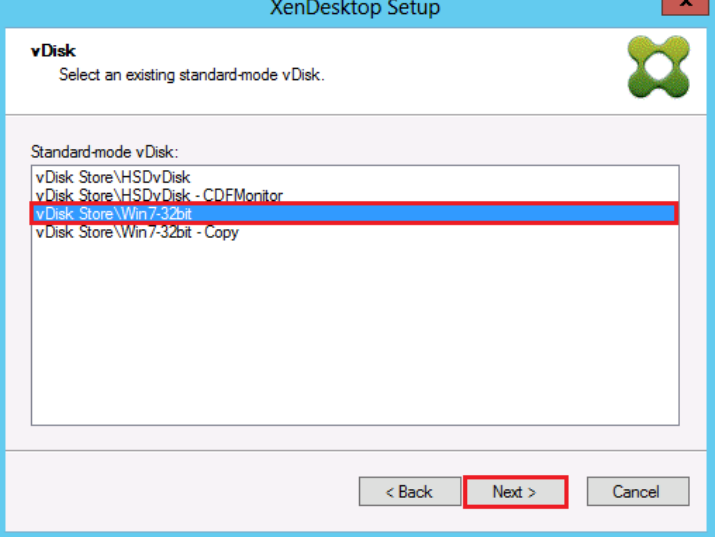
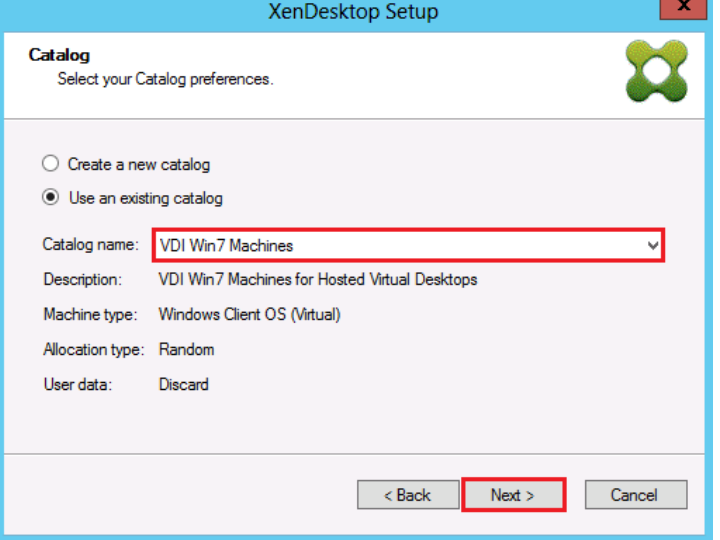
Instructions	Visual
<p>Provide the number of VMs to create Recommended to create 40 or less per run</p> <p>Provide number of vCPUs for the VM 6 for Hosted-Shared VMs</p> <p>Provide the amount of memory for the VM 20GB for Hosted-Shared VMs</p> <p>Provide the write-cache disk size 40GB for the Hosted-Shared VMs</p> <p>Select the PXE boot radio button.</p> <p>Click Next.</p>	
<p>Select the Create new accounts radio button.</p> <p>Click Next.</p>	

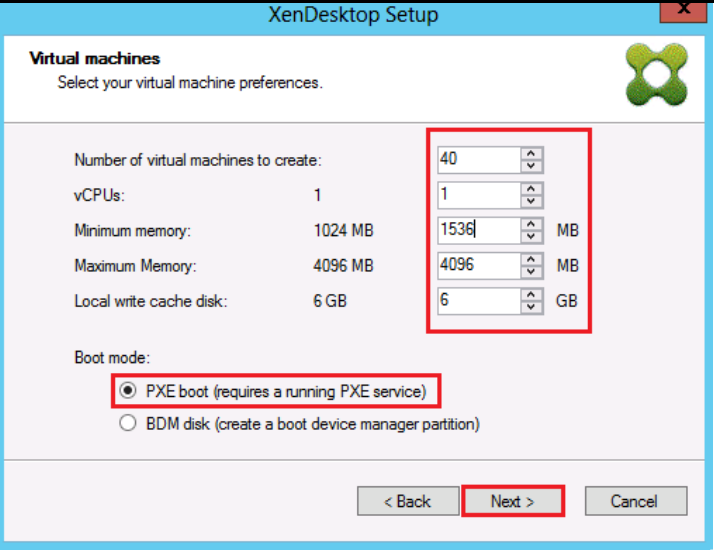
Instructions	Visual
<p>Select the Active Directory Location where the computer accounts should be created by the wizard.</p> <p>Provide the Account naming scheme. An example name is shown in the text box below the name scheme selection location.</p> <p>Click Next.</p>	
<p>Click Finish to begin the virtual machine creation.</p>	

Instructions	Visual
<p>Then the wizard is done creating the virtual machines, click Done.</p>	

When completing the Wizard for the virtual machines, the process will be the same except for the selections on the template, vDisk, catalog, memory, vCPUs, and write-cache size. The differences are shown in the table below.


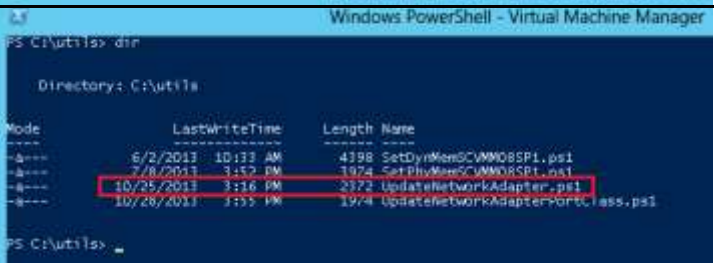
Instructions	Visual
<p>Select the Windows 7 template (XenDesktop Win7 Template) created earlier.</p> <p>Click Next.</p>	

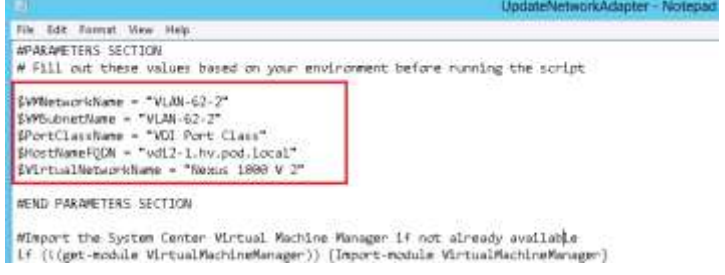
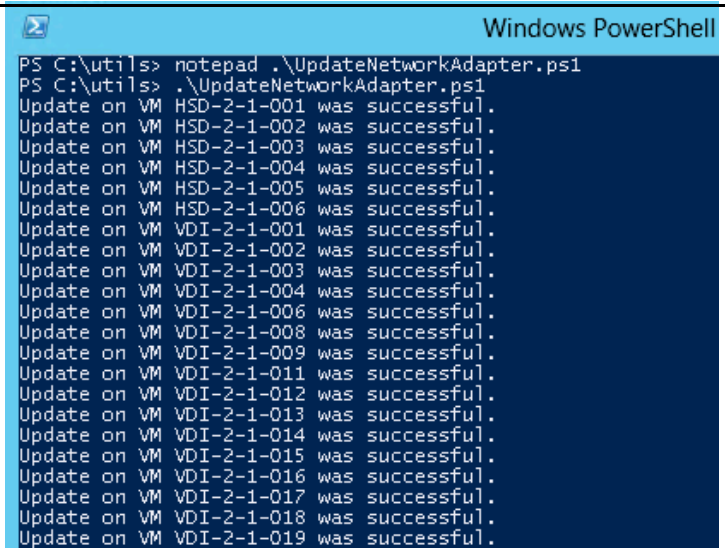
Instructions	Visual
<p>Select the vDisk that will be used to stream to the virtual machine.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'vDisk' selection window in XenDesktop Setup. The title bar says 'XenDesktop Setup'. Below the title bar, it says 'vDisk' and 'Select an existing standard-mode vDisk.' There is a list of vDisks: 'vDisk Store\HSDvDisk', 'vDisk Store\HSDyDisk - CDEMonitor', 'vDisk Store\Win7-32bit' (which is highlighted with a red box), and 'vDisk Store\Win7-32bit - Copy'. At the bottom, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>
<p>Select an existing catalog or choose to create a new catalog.</p> <p>Note: The catalog name is also used as the collection name in PVS site.</p> <p>Click Next.</p>	 <p>The screenshot shows the 'Catalog' selection window in XenDesktop Setup. The title bar says 'XenDesktop Setup'. Below the title bar, it says 'Catalog' and 'Select your Catalog preferences.' There are two radio buttons: 'Create a new catalog' and 'Use an existing catalog' (which is selected). Below the radio buttons, there is a 'Catalog name' dropdown menu with 'VDI Win7 Machines' selected (highlighted with a red box). Below the dropdown, there are labels for 'Description' (VDI Win7 Machines for Hosted Virtual Desktops), 'Machine type' (Windows Client OS (Virtual)), 'Allocation type' (Random), and 'User data' (Discard). At the bottom, there are three buttons: '< Back', 'Next >' (highlighted with a red box), and 'Cancel'.</p>

Instructions	Visual
<p>Provide the number of VMs to create Recommended to create 40 or less per run</p> <p>Provide number of vCPUs for the VM 1 for the Windows 7 VMs</p> <p>Provide dynamic memory settings for the VM 1.5GB Minimum 4 GB Maximum</p> <p>Provide the write-cache disk size 6GB for the Windows 7 VMs</p> <p>Select the PXE boot radio button.</p> <p>Click Next.</p>	

When all the VMs are built on a single host, the PowerShell script from section 13.4.1 Update Virtual Machines created by XenDesktop Wizard can be run from the SCVMM server to update the network adapter settings, the boot order, and the start and stop actions.

Note: If using a hotfixed version of the XenDesktop Setup Wizard the PowerShell script can be modified to not update the network adapters.

Instructions	Visual
<p>Login to the SCVMM server as an administrator and launch PowerShell from the QuickLaunch bar.</p>	
<p>Navigate to the location where the PowerShell script has been saved.</p>	 <pre> PS C:\utils> dir Directory: C:\utils Mode LastWriteTime Length Name ---- - d----- 6/2/2013 10:13 AM 4398 SetDynamSCVMMOSPI.ps1 d----- 7/8/2013 3:47 PM 1974 SetPhysSCVMMOSPI.ps1 d----- 10/25/2013 3:16 PM 2172 UpdateNetworkAdapter.ps1 d----- 10/28/2013 3:55 PM 1974 UpdateNetworkAdapterPortClass.ps1 PS C:\utils> </pre>

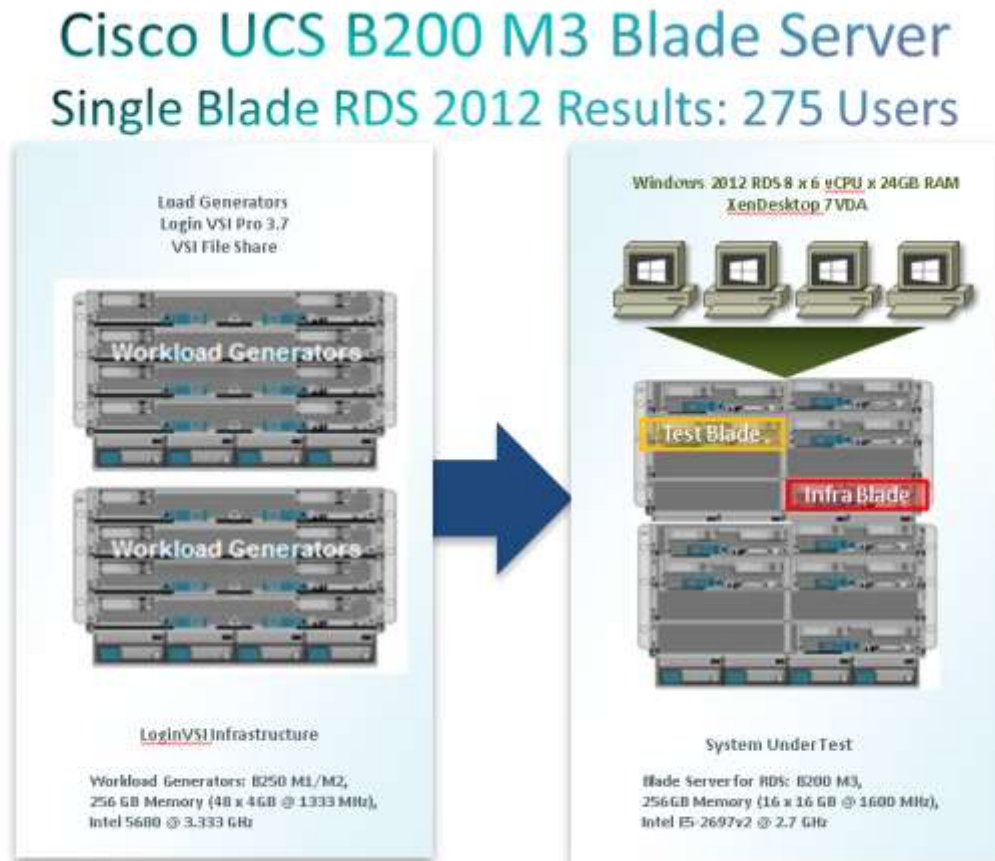
Instructions	Visual
<p>Use Notepad to edit the PowerShell script and to fill in the appropriate values in the PARAMETERS section for the environment.</p>	 <pre> File Edit Format View Help UpdateNetworkAdapter - Notepad #PARAMETERS SECTION # Fill out these values based on your environment before running the script \$VMNetworkName = "VLAN-62-2" \$VMSubnetName = "VLAN-62-2" \$PortClassName = "VMI Port Class" \$HostNameFQDN = "vdl2-1.hv.pod.local" \$VirtualNetworkName = "Nexus 1800 V 2" #END PARAMETERS SECTION #Import the System Center Virtual Machine Manager if not already available If (!(get-module VirtualMachineManager)) {Import-module VirtualMachineManager} </pre>
<p>Execute the PowerShell script by running it from the PowerShell window.</p> <p>PS C:\utils> .\UpdateNetworkAdapter.ps1</p> <p>The script should report each VM being successfully updated.</p>	 <pre> PS C:\utils> notepad .\UpdateNetworkAdapter.ps1 PS C:\utils> .\UpdateNetworkAdapter.ps1 Update on VM HSD-2-1-001 was successful. Update on VM HSD-2-1-002 was successful. Update on VM HSD-2-1-003 was successful. Update on VM HSD-2-1-004 was successful. Update on VM HSD-2-1-005 was successful. Update on VM HSD-2-1-006 was successful. Update on VM VDI-2-1-001 was successful. Update on VM VDI-2-1-002 was successful. Update on VM VDI-2-1-003 was successful. Update on VM VDI-2-1-004 was successful. Update on VM VDI-2-1-006 was successful. Update on VM VDI-2-1-008 was successful. Update on VM VDI-2-1-009 was successful. Update on VM VDI-2-1-011 was successful. Update on VM VDI-2-1-012 was successful. Update on VM VDI-2-1-013 was successful. Update on VM VDI-2-1-014 was successful. Update on VM VDI-2-1-015 was successful. Update on VM VDI-2-1-016 was successful. Update on VM VDI-2-1-017 was successful. Update on VM VDI-2-1-018 was successful. Update on VM VDI-2-1-019 was successful. </pre>

8. Test Setup and Configurations

This section provides an overview of the test configurations for this validated design along with a summary of the recommended results from the testing. More detailed information about the results can be found in Section 9 Login VSI Test Result.

8.1. Cisco UCS Test Configuration for Single Blade Scalability of Hosted Shared Desktops

Figure 20: RDS Single-Server Results



Hardware components

- Virtual Desktop Hosts: 1 X Cisco UCS B200-M3 (E5-2697v2 @ 2.7 GHz) blade server with 256GB of memory (16 GB X 16 DIMMS @ 1600 MHz), 2X 400GB Samsung SSD, and 1X VIC1240 Converged Network Adapter
- Infrastructure Servers: 2 X Cisco UCS B200-M3 (E5-2650) blade servers with 128 GB of memory (16 GB X 8 DIMMS @ 1600 MHz), 2X 600GB Seagate SAS 10K 6Gb, and 1X VIC1240 Converged Network Adapter
- Load Generators: 8 X Cisco UCS B250-M2 (5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) 1 X M81KR (Palo) Converged Network Adapter
- 2 X Cisco Fabric Interconnect 6248UP
- 1 X EMC VNXe System 3300, dual-controller storage system for HA, 4 X dual port 10 GbE cards, 44 X 600GB SAS drives for Infrastructure file shares and Boot LUNs

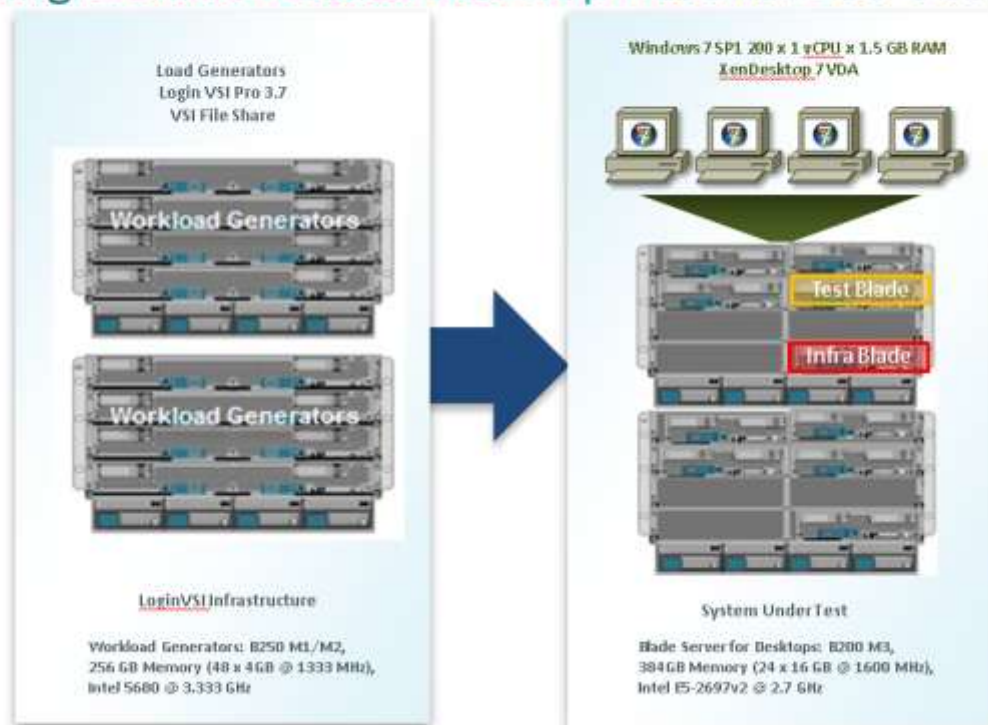
Software components

- Cisco UCS firmware 2.1(3a)
- Cisco Nexus 1000V virtual distributed switch
- XenDesktop 7
- Provisioning Server 7
- Citrix User Profile Manager
- Windows 2012 64-bit Remote Desktop Services, 8 X 6vCPU, 24 GB of static memory

8.2. Cisco UCS Test Configuration for Single Blade Scalability of Hosted Virtual Machines

Figure 21: VDI Single-Server Results

Cisco UCS B200 M3 Blade Server Single Blade Hosted Desktop Results: 200 Users



Hardware components

- Virtual Desktop Hosts: 1 X Cisco UCS B200-M3 (E5-2697v2 @ 2.7 GHz) blade server with 384GB of memory (16 GB X 24 DIMMS @ 1600 MHz), 2X 400GB Samsung SSD, and 1X VIC1240 Converged Network Adapter

- Infrastructure Servers: 2 X Cisco UCS B200-M3 (E5-2650) blade servers with 128 GB of memory (16 GB X 8 DIMMS @ 1600 MHz), 2X 600GB Seagate SAS 10K 6Gb, and 1X VIC1240 Converged Network Adapter
- Load Generators: 8 X Cisco UCS B250-M2 (5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) 1 X M81KR (Palo) Converged Network Adapter
- 2 X Cisco Fabric Interconnect 6248UP
- 1 X EMC VNXe System 3300, dual-controller storage system for HA, 4 X dual port 10 GbE cards, 44 X 600GB SAS drives for Infrastructure file shares and Boot LUNs

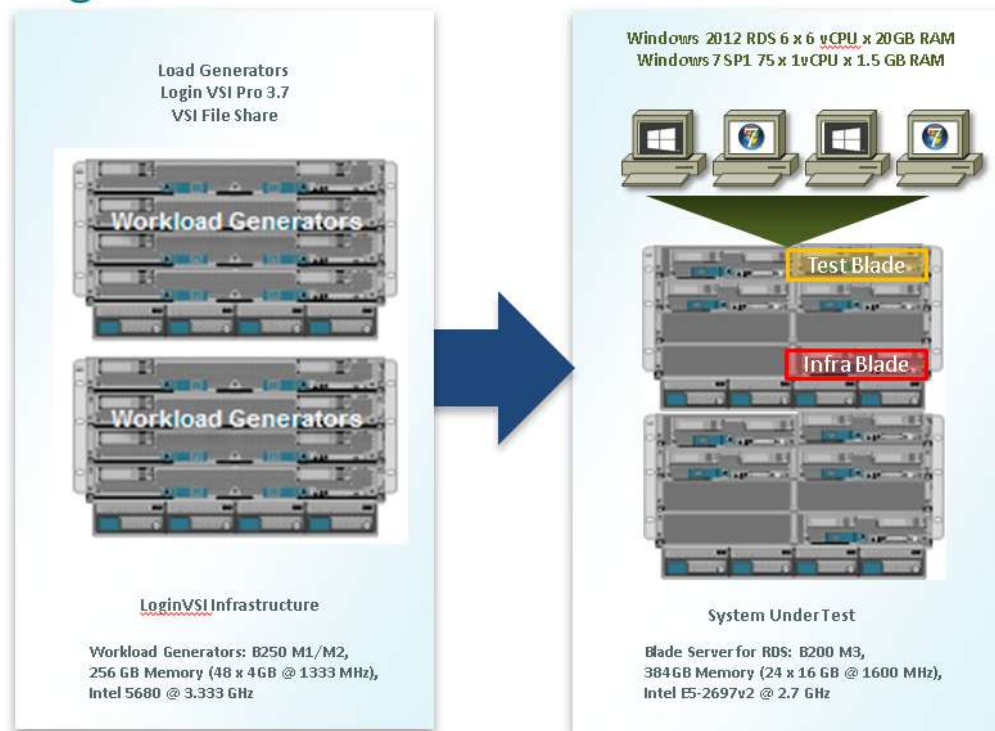
Software components

- Cisco UCS firmware 2.1(3a)
- Cisco Nexus 1000V virtual distributed switch
- XenDesktop 7
- Provisioning Server 7
- Citrix User Profile Manager
- Windows 7 SP1 32-bit, 1vCPU, 1.5 GB of static memory

8.3. Cisco UCS Test Configuration for Single Blade Scalability for a Mixed Hosted Shared and Hosted Virtual Workload

Figure 22: Mixed Workload Single-Server Results

Cisco UCS B200 M3 Blade Server Single Blade Mix Workload Results: 250 Users



Hardware components

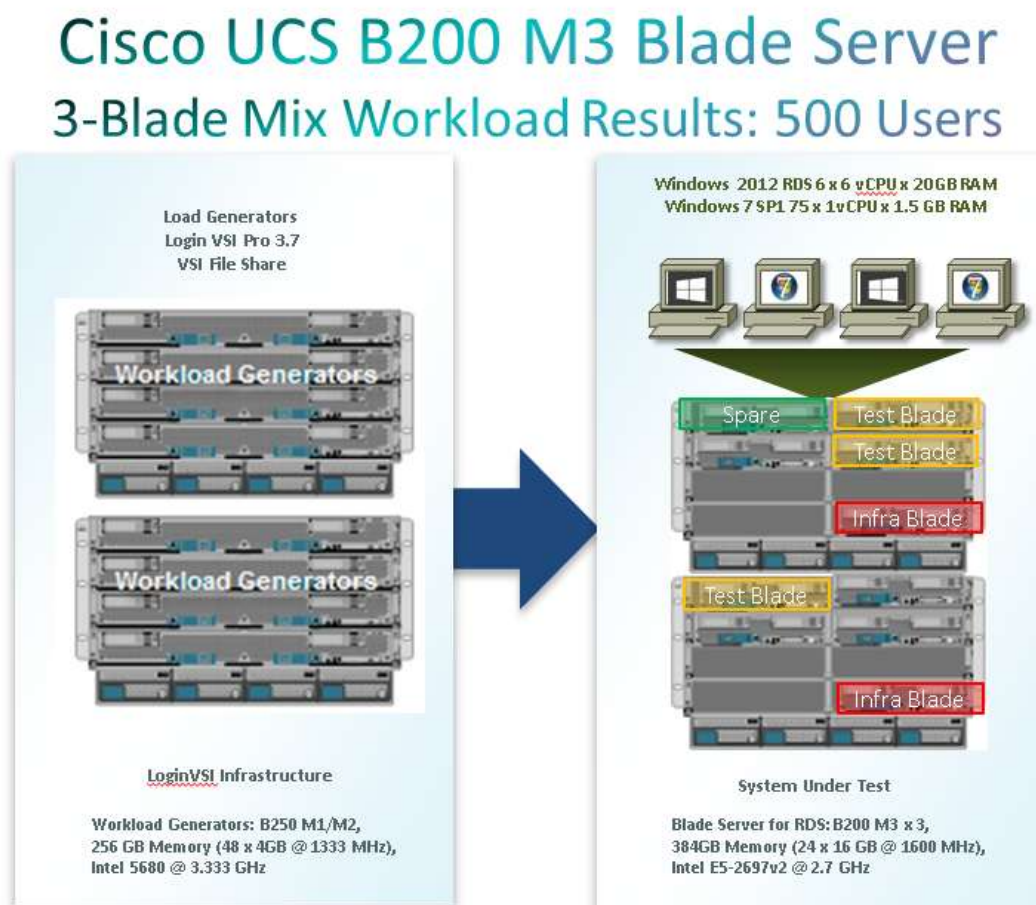
- Virtual Desktop Hosts: 1 X Cisco UCS B200-M3 (E5-2697v2 @ 2.7 GHz) blade server with 384GB of memory (16 GB X 24 DIMMS @ 1600 MHz), 2X 400GB Samsung SSD, and 1X VIC1240 Converged Network Adapter
- Infrastructure Servers: 2 X Cisco UCS B200-M3 (E5-2650) blade servers with 128 GB of memory (16 GB X 8 DIMMS @ 1600 MHz), 2X 600GB Seagate SAS 10K 6Gb, and 1X VIC1240 Converged Network Adapter
- Load Generators: 8 X Cisco UCS B250-M2 (5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) 1 X M81KR (Palo) Converged Network Adapter
- 2 X Cisco Fabric Interconnect 6248UP
- 1 X EMC VNXe System 3300, dual-controller storage system for HA, 4 X dual port 10 GbE cards, 44 X 600GB SAS drives for Infrastructure file shares and Boot LUNs

Software components

- Cisco UCS firmware 2.1(3a)
- Cisco Nexus 1000V virtual distributed switch
- XenDesktop 7
- Provisioning Server 7
- Citrix User Profile Manager
- 6 Windows 2012 64-bit Remote Desktop Services, 6vCPU, 20GB of dynamic memory
- 75 Windows 7 SP1 32-bit, 1vCPU, 1.5GB of dynamic memory

8.4. Cisco UCS Test Configuration for a Single-Chassis 500-User Configuration

Figure 23: Multi-Server 500-User Results



Hardware components

- Virtual Desktop Hosts: 3 X Cisco UCS B200-M3 (E5-2697v2 @ 2.7 GHz) blade server with 384GB of memory (16 GB X 24 DIMMS @ 1600 MHz), 2X 400GB Samsung SSD, and 1X VIC1240 Converged Network Adapter
- Infrastructure Servers: 2 X Cisco UCS B200-M3 (E5-2650) blade servers with 128 GB of memory (16 GB X 8 DIMMS @ 1600 MHz), 2X 600GB Seagate SAS 10K 6Gb, and 1X VIC1240 Converged Network Adapter
- Load Generators: 8 X Cisco UCS B250-M2 (5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) 1 X M81KR (Palo) Converged Network Adapter
- 2 X Cisco Fabric Interconnect 6248UP
- 1 X EMC VNXe System 3300, dual-controller storage system for HA, 4 X dual port 10 GbE cards, 44 X 600GB SAS drives for Infrastructure file shares and Boot LUNs

Software components

- Cisco UCS firmware 2.1(3a)
- Cisco Nexus 1000V virtual distributed switch
- XenDesktop 7
- Provisioning Server 7
- Citrix User Profile Manager
- 18 Windows 2012 64-bit Remote Desktop Services, 6 vCPU, 20GB of dynamic memory
- 225 Windows 7 SP1 32-bit, 1vCPU, 1.5 GB of dynamic memory

Figure 24: Multi-Server 1000-user Results

The diagram illustrates the test environment architecture, showing the flow from Load Generators to the System Under Test.

Load Generators: Two racks of Workload Generators are shown, each labeled "Workload Generators".

System Under Test: A rack of Blade Servers is shown, labeled "System Under Test". The rack contains several blades, including "Spare", "Test Blade", and "Infra Blade".

Configuration Details:

- Load Generators:**
 - Operating System: Windows 2012 RDS 6 x 6 vCPU x 20GB RAM
 - Operating System: Windows 7 SP1 75 x 1vCPU x 1.5 GB RAM
 - Software: Login VSI Pro 3.7
 - Feature: VSI File Share
- System Under Test:**
 - Blade Server for RDS: B200 M3 x 3, 384GB Memory (24 x 16 GB @ 1600 MHz), Intel E5-2697v2 @ 2.7 GHz

A large blue arrow points from the Load Generators to the System Under Test, indicating the direction of traffic.

- Virtual Desktop Hosts: 5 X Cisco UCS B200-M3 (E5-2697v2 @ 2.7 GHz) blade servers with 384GB of memory (16 GB X 24 DIMMS @ 1600 MHz), 2X 400GB Samsung SSD, and 1X VIC1240 Converged Network Adapter
- Infrastructure Servers: 2 X Cisco UCS B200-M3 (E5-2650) blade servers with 128 GB of memory (16 GB X 8 DIMMS @ 1600 MHz), 2X 600GB Seagate SAS 10K 6Gb, and 1X VIC1240 Converged Network Adapter
- Load Generators: 8 X Cisco UCS B250-M2 (5680 @ 3.333 GHz) blade servers with 192 GB of memory (4 GB X 48 DIMMS @ 1333 MHz) 1 X M81KR (Palo) Converged Network Adapter
- 2 X Cisco Fabric Interconnect 6248UP
- 1 X EMC VNXe System 3300, dual-controller storage system for HA, 4 X dual port 10 GbE cards, 44 X 600GB SAS drives for Infrastructure file shares and Boot LUNs

288

- Cisco UCS firmware 2.1(3a)
- Cisco Nexus 1000V virtual distributed switch
- XenDesktop 7
- Provisioning Server 7
- Citrix User Profile Manager
- 42 Windows 2012 64-bit Remote Desktop Services, 6 vCPU, 20GB of dynamic memory
- 525 Windows 7 SP1 32-bit, 1vCPU, 1.5 GB of dynamic memory

8.6. Test Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco Solution Labs with joint support from Citrix, Microsoft, and EMC resources.

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 Hosted VDI model under test.

Test metrics were gathered from the hypervisor, 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.

8.6.1. Load Generation

Within each test environment, load generators were utilized to put demand on the system to simulate multiple users accessing the XenDesktop 7 environment and executing a typical end-user workflow. To generate load within the environment, the Login VSI software application was used to generate the end user connection to the XenDesktop 7 environment, to provide unique user credentials to the Citrix StoreFront server, to initiate the workload, and to evaluate the end-user experience.

In the Hosted VDI test environment, sessions launchers were used simulate multiple users making a direct connection to XenDesktop 7 through a Citrix HDX protocol connection.

8.6.2. User Workload Simulation – Login VSI

One of the most critical factors of validating a XenDesktop deployment is identifying a real-world user workload that is easy for customers to replicate and standardized across platforms to allow customers to realistically test the impact of a variety of worker tasks. To accurately represent a real-world user workload, the Login VSI third-party tool from Login Consultants was used throughout the Hosted VDI testing.

The tool has the benefit of taking measurements of the in-session response time, providing an objective way to measure the expected user experience for individual desktop throughout large scale testing, including login storms.

The Virtual Session Indexer (Login Consultants' Login VSI 3.7) methodology, designed for benchmarking Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI) environments is completely platform and protocol independent and hence allows customers to easily replicate the testing results in their environment.

Login VSI calculates an index based on the amount of simultaneous sessions that can be run on a single machine.

Login VSI simulates a medium workload user (also known as knowledge worker) running generic applications such as: Microsoft Office 2007 or 2010, Internet Explorer 8 including a Flash video applet and Adobe Acrobat Reader (Note: For the purposes of this test, applications were installed locally, not streamed nor hosted on XenApp).

Like real users, the scripted Login VSI session will leave multiple applications open at the same time. The medium workload is the default workload in Login VSI and was used for this testing. This workload emulated a medium knowledge working using Office, IE, printing and PDF viewing.

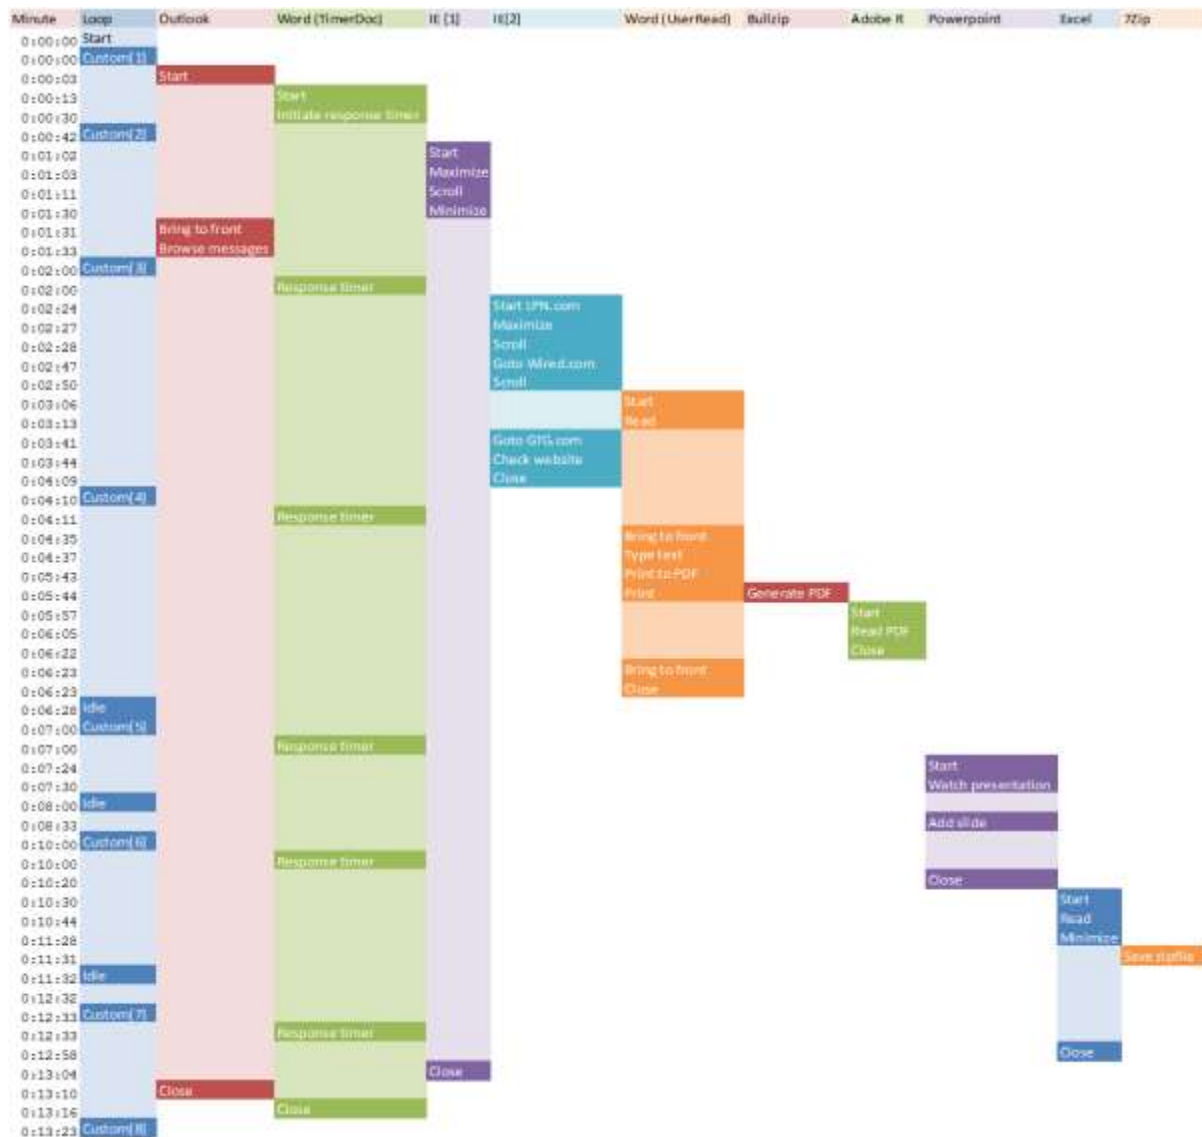
- When a session has been started the medium workload will repeat every 12 minutes.
- During each loop the response time is measured every 2 minutes.
- The medium workload opens up to 5 apps simultaneously.
- The type rate is 160ms for each character.
- Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007/2010, browse 10 messages.
- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com, Lonelyplanet.com and heavy
- 480 p Flash application gettheglass.com.
- Word 2007/2010, one instance to measure response time, one instance to review and edit document.
- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.
- Excel 2007/2010, a very large randomized sheet is opened.
- PowerPoint 2007/2010, a presentation is reviewed and edited.
- 7-zip: using the command line version the output of the session is zipped.

A graphical representation of the medium workload is shown below.

Figure 25: Login VSI Workload Timing Chart



You can obtain additional information on Login VSI from <http://www.LoginVSI.com>.

8.6.3. Testing Procedure

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

8.6.3.1. Test Run Setup

The test run setup for both the single-server and multi-server tests was the same. The following steps were completed:

1. Delivery Group(s) placed in maintenance mode
2. Reboot Citrix XenDesktop Controller(s)
3. Reboot Citrix StoreFront Server
4. Reboot Hyper-V Virtual Desktop Host(s)

5. Reboot Login VSI Launchers

Rebooting the Citrix infrastructure components and the VDI hosts insures that a clean environment was available for each test run.

8.6.3.2. Test Run Protocol

To simulate severe, real-world environments, Cisco requires that the log-on and start-work sequence, known as Ramp Up, must be completed in 30 minutes. Additionally, Cisco requires all test runs, whether performing a single-server or multi-server test, to be started and become active within two minutes after the session is launched.

For each of the three consecutive runs on single-server tests, this process was followed:

1. Started logging on Hyper-V Host Servers
2. Depending on whether testing Hosted Shared Desktop Groups or Hosted Virtual Desktops one of the following sequences was executed.
 - Take the Hosted Shared Desktop group out of maintenance mode and start all the Hosted Shared Desktop virtual machines then wait for them to register. (~5 minutes)
 - Take the Hosted Virtual Desktop group out of maintenance mode and wait for the virtual machines to register (~10 minutes)
3. Wait an additional 5 minutes for the hypervisor to settle down.
4. Start the Login VSI 3.7 test configured with an 1800-second parallel launching window and 840-second auto-logoff for all sessions. For the single-server tests 14 launchers were used.
5. Wait and verify all sessions have become active. (~30 minutes)
6. Wait for the auto-logoff timer to expire (~14 minutes) which is set to allow the sessions to complete at least one full loop, at which point the Login VSI places the logoff.txt file on the VSIShare to initiate logoff.
7. Wait until all the sessions have logged off. (~30 minutes)
8. Stop logging

For each of the three consecutive runs on multi-server tests, this process was followed:

1. Started logging on:
 - Hyper-V Host Servers
 - Citrix PVS Server(s)
 - Citrix Desktop Controller(s)
 - Citrix StoreFront Server
 - Microsoft SCVMM Server
 - Microsoft SQL Server(s)
 - Microsoft Domain Controller(s)
 - EMC VNXe Storage Processor(s)
2. Take the Hosted Shared Desktop group out of maintenance mode and start all the Hosted Shared Desktop virtual machines then wait for them to register. (~5 minutes)

3. Take the Hosted Virtual Desktop group out of maintenance mode and wait for all the Hosted Virtual Desktops to register. (~10 minutes)
4. Wait an additional 5 minutes for the hypervisor to settle down.
5. Start the Login VSI 3.7 test configured with an 1800-second parallel launching window and 840-second auto-logoff for all sessions. For the 500-user test 30 launchers were used. For the 1000-user test 61 launchers were used.
6. Wait and verify all sessions have become active. (~30 minutes)
7. Wait for the auto-logoff timer to expire (~14 minutes) which is set to allow the sessions to complete at least one full loop, at which point the Login VSI places the logoff.txt file on the VSIShare to initiate logoff.
8. Wait until all the sessions have logged off. (~30 minutes)
9. Stop logging

8.6.4. Success Criteria

Multiple metrics were captured during each test run, but the success criteria for considering a single test run as pass or fail was based on the key metric, VSI Max. The Login VSI Max evaluates the user response time during increasing user load and assesses the successful start-to-finish execution of all the initiated virtual desktop sessions. A successful test sequence consisted of **three consecutive test runs** which met the passing criteria within a 1% variance.

8.6.4.1. Login VSI Max

VSI Max represents the maximum number of users the environment can handle before serious performance degradation occurs. VSI Max is calculated based on the response times of individual users as indicated during the workload execution. The user response time has a threshold of 4000ms and all users response times are expected to be less than 4000ms in order to assume that the user interaction with the virtual desktop is at a functional level. VSI Max is reached when the response times reaches or exceeds 4000ms for 6 consecutive occurrences. If VSI Max is reached, that indicates the point at which the user experience has significantly degraded. The response time is generally an indicator of the host CPU resources, but this specific method of analyzing the user experience provides an objective method of comparison that can be aligned to host CPU performance.

8.6.4.2. Calculating VSImax

Typically the desktop workload is scripted in a 12-14 minute loop when a simulated Login VSI user is logged on. After the loop is finished it will restart automatically. Within each loop the response times of seven specific operations is measured in a regular interval: six times in within each loop. The response times if these seven operations are used to establish VSImax.

The seven operations from which the response times are measured are:

1. Copy new document from the document pool in the home drive
 - This operation will refresh a new document to be used for measuring the response time. This activity is mostly a file-system operation.
2. Starting Microsoft Word with a document

- This operation will measure the responsiveness of the Operating System and the file system. Microsoft Word is started and loaded into memory; also the new document is automatically loaded into Microsoft Word. When the disk I/O is extensive or even saturated, this will impact the file open dialogue considerably.
3. Starting the “File Open” dialogue
 - This operation is handled for small part by Microsoft Word and a large part by the operating system. The file open dialogue uses generic subsystems and interface components of the OS. The OS provides the contents of this dialogue.
 4. Starting “Notepad”
 - This operation is handled by the OS (loading and initiating notepad.exe) and by the Notepad.exe itself through execution. This operation seems instant from an end-user’s point of view.
 5. Starting the “Print” dialogue
 - This operation is handled for a large part by the OS subsystems, as the print dialogue is provided by the OS. This dialogue loads the print-subsystem and the drivers of the selected printer. As a result, this dialogue is also dependent on disk performance.
 6. Starting the “Search and Replace” dialogue
 - This operation is handled within the application completely; the presentation of the dialogue is almost instant. Serious bottlenecks on application level will impact the speed of this dialogue.
 7. Compress the document into a zip file with 7-zip command line
 - This operation is handled by the command line version of 7-zip. The compression will very briefly spike CPU and disk I/O.

These measured operations with 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 are 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. When such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

For these tests, Cisco utilized the VSImax Dynamic model exclusively.

8.6.4.3. VSImax Dynamic

VSImax Dynamic is calculated when the response times are consistently above a certain threshold. However, this threshold is now dynamically calculated on the baseline response time of the test.

Five individual measurements are weighted to better support this approach:

- Copy new doc from the document pool in the home drive: 100%
- Microsoft Word with a document: 33.3%
- Starting the “File Open” dialogue: 100%
- Starting “Notepad”: 300%

- Starting the “Print” dialogue: 200%
- Starting the “Search and Replace” dialogue: 400%
- Compress the document into a zip file with 7-zip command line 200%

A sample of the VSImax Dynamic response time calculation is displayed below:

Figure 26: VSImax Dynamic Results

Activity (RowName)	Result (ms)	Weight (%)	Weighted Result (ms)
Refresh document (RFS)	160	100%	160
Start Word with new doc (LOAD)	1400	33.3%	467
File Open Dialogue (OPEN)	350	100%	350
Start Notepad (NOTEPAD)	50	300%	150
Print Dialogue (PRINT)	220	200%	440
Replace Dialogue (FIND)	10	400%	40
Zip documents (ZIP)	130	200%	230
VSImax Dynamic Response Time			1837

The average VSImax response time is calculated based on the amount of active Login VSI users logged on to the system. For VSImax value to be reached the average VSImax response times need to consistently higher than a dynamically calculated threshold.

To determine this dynamic threshold, first the average baseline response time is calculated. This is done by averaging the baseline response time of the first 15 Login VSI users on the system.

The formula for the dynamic threshold is: Avg. Baseline Response Time x 125% + 3000. As a result, when the baseline response time is 1800, the VSImax threshold will now be $1800 \times 125\% + 3000 = 5250\text{ms}$.


When application virtualization is used, the baseline response time can vary widely per vendor and streaming strategy. Using the VSImax Dynamic model will provide a level playing field when comparing application virtualization or anti-virus applications. The resulting VSImax Dynamic scores are aligned again with saturation on a CPU, Memory or Disk level, also when the baseline response time are relatively high.

8.6.4.4. Determining VSImax

The Login VSI analyzer will automatically identify the “VSImax”. In the example below the VSImax is 98. The analyzer will automatically determine “stuck sessions” and correct the final VSImax score.

- Vertical axis: Response Time in milliseconds
- Horizontal axis: Total Active Sessions

Figure 27: Sample Login VSI Analyzer Graphic Output



- Red line: Maximum Response (worst response time of an individual measurement within a single session)
- Orange line: Average Response Time within for each level of active sessions
- Blue line: the VSImax average.
- Green line: Minimum Response (best response time of an individual measurement within a single session)

In our tests, the total number of users in the test run had to login, become active and run at least one test loop and log out without reaching the VSI Max to be considered a success.

9. Login VSI Test Result

This section provides the validation results of the Login VSI testing within the environment for the four configurations of single-server and multi-blade architectures, each in their own section. These sections provide data points for customers to reference when designing their own environment. The first two single-server scalability sections provided the information necessary to correctly identify the mix of hosted shared desktops versus hosted virtual desktops for the final two multi-blade tests.

9.1. Cisco UCS B200-M3 Single-Server Scalability Results for Hosted Shared Desktops


As noted earlier in section 8.6.4 Success Criteria, in order for a successful test sequence to be accepted, it had to pass in three consecutive test runs. The first test sequence was to determine the VSImax value for a single blade running only Hosted Shared Desktop sessions on Windows Server 2012 with XenDesktop 7. The maximum density was determined by testing different combinations of servers and vCPUs assigned to those servers.

The best performance was achieved when the number of vCPUs assigned to the virtual machines did not exceed the number of hyper-threaded cores available on the server. In other words, not overcommitting

the CPU resources provides the best user experience. For the E5-2697v2 processors this means that 24 cores with hyper-threading will enable 48 vCPUs. The highest density was observed at eight VMs each with six vCPUs assigned.

To achieve the VSI_{max} score of 299 for hosted shared desktops, 320 users were launched against the eight Windows 2012 virtual machines with six vCPUs using the VSI Medium Workload with flash. The Login VSI score was achieved on three consecutive runs reporting 298, 299, and 299 respectively. The graph from the first 299 run is shown below.


Figure 28: VSI_{max} for RDS Single-Server Results



The key performance charts for this representative test run are shown below along with any interesting observations.


One interesting observation around the CPU utilization is that with the Ivy Bridge CPUs, the processor is not pegged for an extended period of time during the test, which was previously observed with the Sandy Bridge processors. This effect is even more pronounced with the hosted virtual desktops.

Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time RDS 2012 | 320 Sessions | 8 VMs x 6vCPU | VSI Medium Workload




With the processor being less of a constraint within the environment, the disk performance becomes more prominent. The charts below depict the Disk Queue lengths and Disk IOs for the two SSD drives in a RAID0 array which were hosting the write-cache drives. The graphs show the write activity high at the beginning and during the steady-state phases, with the read activity spiking near the end of the test during logoff.

Disk Queue Lengths SSD Drive RDS 2012 | 8 VMs x 24 GB | VSI Medium Workload




Disk IO Operations SSD RAID0 RDS 2012 | 8 VMs x 24 GB | Medium VSI Workload




With memory the total memory assigned to the servers was 192GB (8x24GB out of a total 256GB). The memory was statically assigned. The graph below shows the memory utilization on the blade during the test was essentially static with 53GB available throughout the run.

Memory\Available MegaBytes RDS 2012 | 8 VMs x 24GB | VSI Medium Workload



Below is the networking performance during the test. As expected, based on the disk activity above, the early part of the test during ramp up shows the bulk of the receive traffic as the VMs accept files from the PVS server. The end of the test shows a spike in send traffic as the profiles are updated and data is read off of the disks. The valley in the middle provides a good understanding of when the steady-state phase was in place with very little network activity compared to the ramp up and logoff stages. For reference, 120000000 is approximately 915 Mbps and 80000000 is approximately 610 Mbps.




Based on the data gathered during testing, Cisco recommends for the VSI Medium workload (which includes Flash) that the number of users be kept around 275. This equates to approximately 35 users per XenDesktop 7 VM with 8 virtual machines on a blade.

9.2. Cisco UCS B200-M3 Single-Server Scalability Results for Hosted Virtual Desktops


The second phase of the testing involved determining the VSI_{max} value for hosted virtual desktops running Windows 7 SP1 32-bit operating system with XenDesktop 7.

A VSI_{max} of 205 was achieved with the three consecutive runs of 218-220 desktops on a single B200-M3 blade. The three VSI_{max} scores for the runs were 204, 207, and 205 respectively. The VSI_{max} graph for the 205 run is shown below.

Figure 29: VSImax for VDI Single-Server Results




The CPU resources were not pegged at any point during the test as seen from the graph below:




While CPU is probably still a factor in reaching higher numbers, it is now joined by the disks within this configuration. The charts below show the disk queue lengths and disk transfers per second for the two-disk RAID0 array of SSD drives that host the PVS write-cache drives. The disk queue length is the chart that shows the disks are starting to get busy with longer than ideal queue lengths.



Disk Queue Lengths SSD RAID0 Windows 7 | 219 Desktops | 1.5 GB x 1 vCPU | LoginVSI Medium Workload



Disk IO Operations SSD RAID0 Windows 7 | 219 Desktops | 1.5 GB x 1 vCPU | LoginVSI Medium Workload



The memory became the gating factor, so the memory on the blade was increased from 256GB to 384GB.



The recommended value is one where the CPU resources peak around 95%, since CPU resources are not the gating factor, Cisco recommends for the Medium VSI (with flash) workload to target 200 virtual desktops.

9.3. Cisco UCS B200-M3 Recommended Single-Server Mixed Desktop Workload

Using the information gained from single-server VSImax testing, the next step was to identify the best workload mix and blade use for the environment that met the project requirement of 70% hosted


shared desktops and 30% hosted virtual desktops. Based on the recommended loads from the single server testing, 3 blades would be required to support the fault-tolerant hosted shared desktop environment and 3 blades to support the fault-tolerant hosted virtual desktop environment, for a total of six blades.

Since all the workloads are virtualized, a better approach is to mix the workloads so that each blade consists of 70% hosted shared desktops and 30% hosted virtual desktops. This means only a single physical blade will be required to provide fault-tolerance for the environment. Using the recommended loads from the single-server testing results in a mixed blade of 6 hosted shared desktop VMs and 75 hosted virtual desktop VMs. The six hosted shared desktop VMs would support 175 users bringing the total users per physical blade to 250.

The mixed server workload was switched to dynamic memory instead of the static memory used in the single-server testing. This change was made after testing in the environment showed no significant difference in CPU Utilization between enabling dynamic memory and disabling it. This change more closely resembles “real world” environments and provides the power users with access to more memory as necessary. The memory in the Windows 2012 virtual machines was reduced from 24GB to 20GB, since testing showed only 14GB was being used by the individual Windows Server 2012 VMs. This extra memory, then became available to the dynamic memory pool for use if necessary by either the Windows 7 or Windows 2012 virtual machines, although it would not be needed with the VSI Medium workload.


A run was completed successfully without receiving a VSImax value as expected. The chart below shows the VSImax test for a single server at the recommended mix of virtual machines to support the 250 users.

Figure 30: VSImax Mixed Workload Single-Server Results



The performance charts from this test run are shown below and since the recommended workload is selected, the performance counters are within expected operating parameters. The CPU processor chart for the test period is shown below for a reference point.

Figure 31: CPU Processor Chart 250-user Mixed Workload




The mixed workload blade configuration allowed the 500-user to be contained within 2 blades and the 1000-user configuration to be contained within 4 blades. Adding one blade for fault-tolerance means the solution becomes 3 blades for 500-users and 5 blades for 1000 users. Of course, for full fault-tolerance, the blades should be housed in different chassis.

9.4. Cisco UCS 500-User Scalability Results

This section provides the testing results for the 500-user testing with three physical blades. As expected from the 250-user testing, the blades were able to manage the load without overburdening the processor.

Figure 32: VSImax 500-user Scalability Results




Below are the performance charts from the servers during the test.

9.4.1. EMC VNXe Performance

The following chart shows the IOPS of the VNXe array during the 500-user test. The VNXe array was used solely for the user profiles and home directories through the CIFS shares.

Figure 33: VNXe Performance 500-users Login VSI




9.4.2. VDI Host

This section contains charts for one of the VDI hosts which provides a representative sample of the host performance. The graphs for the remaining four hosts can be found in the Appendix Section **Error! eference source not found..**

9.4.2.1. Processor

The hosts were busy, but never pegged the CPU during the test.


Figure 34: Processor Performance VDI Host 500-user



9.4.2.2. Memory

The hosts were running with about 10GB of available RAM after all the VMs started up. The VMs were configured with dynamic memory. The six RDS hosts had 20GB minimum/startup configured and the 75 desktops had 1.5GB minimum/startup configured.


Figure 35: Memory Performance VDI Host 500-User



9.4.2.3. Disk


The only disk with interesting information is the E: drive which consisted of two 400GB Enterprise SSD drives in a RAID0 array. The C: drive, which was the boot iSCSI LUN on the VNxe, had almost no activity reported. The E: drive hosted the PVS write-cache drive, so majority of the traffic would be observed at the peak of the test as the last few desktops logged on. The SSD's are handling about 6000 IOPS per VDI host during the peak periods, which represents a significant reduction on the IOPS required on the backend VNxe storage array.

Figure 36: Disk IO Performance VDI Host 500-user



Disk performance can be quantified partially by outstanding disk queue lengths. In this case, with two disks on the RAID0 array, any queue lengths under 2 would be ideal. From the chart below, it appears that other than three intervals, all the sampled queue lengths are ideal.


Figure 37: Disk Queue Length Performance VDI Host 500-user



9.4.2.4. Network

For reference, 60,000,000 bytes per second equates to about 458Mbps. The network on the host, even at peak times was within normal expected loads.

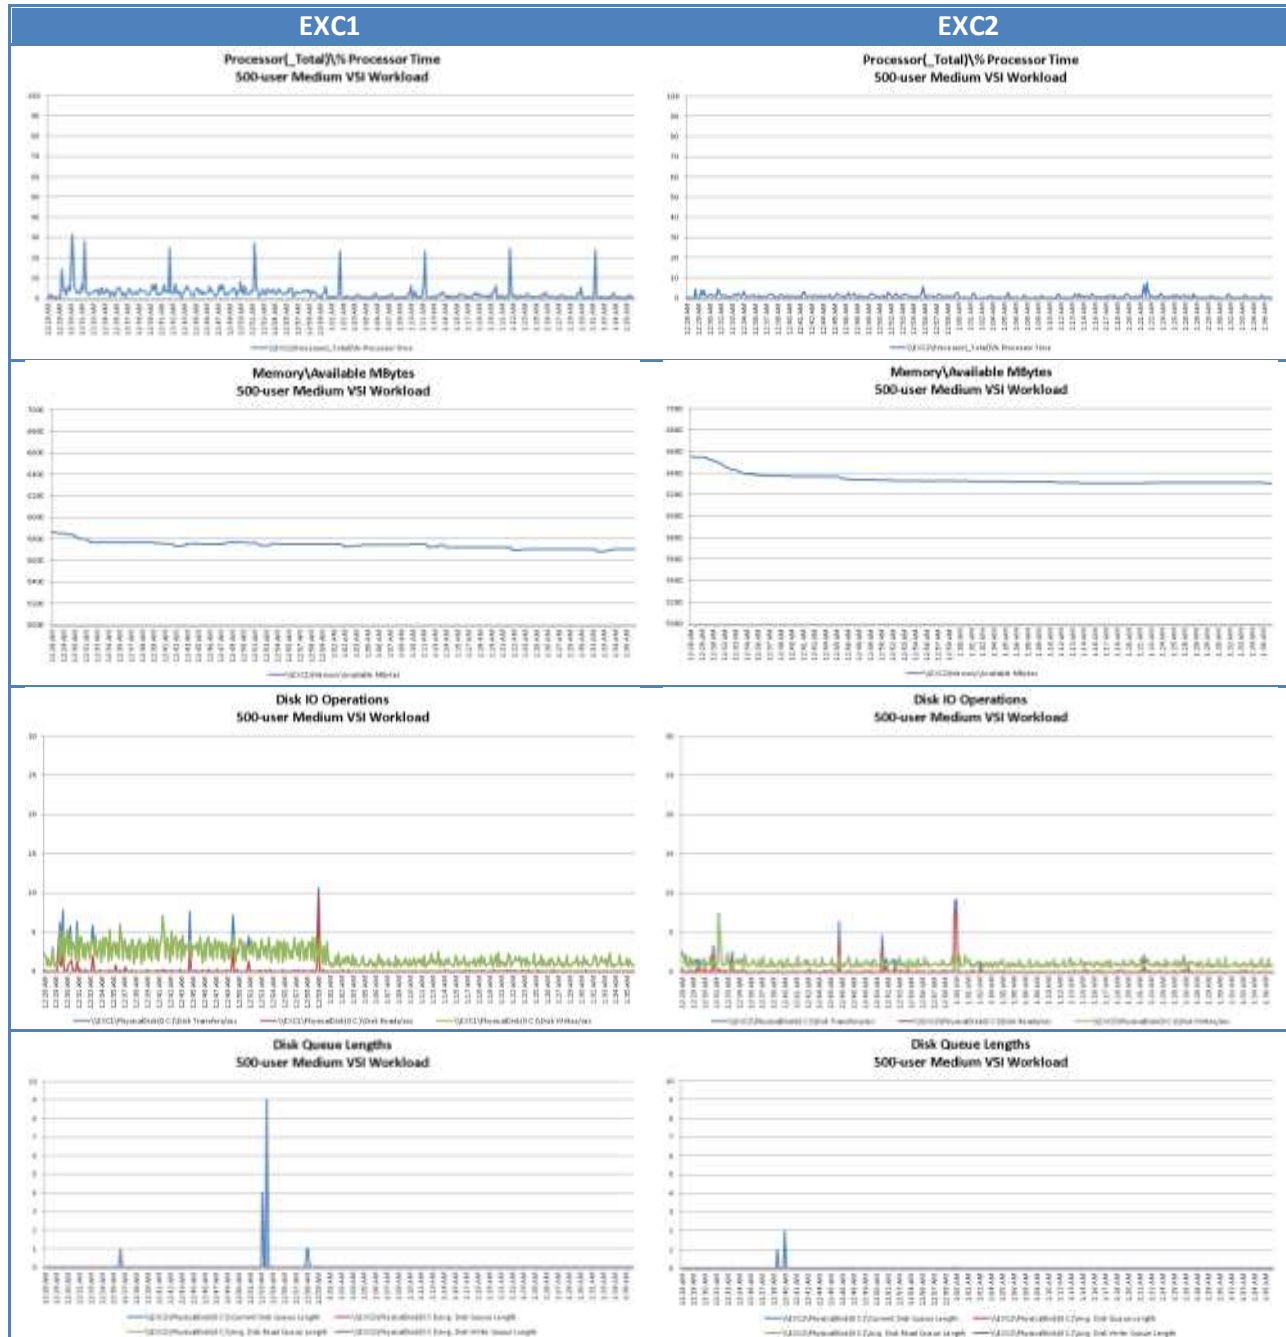
Figure 38: Network Performance VDI Host 500-user




The VDI hosts performed well within expected performance guidelines.

9.4.3. XenDesktop Controllers


Notice the largest hit for processor time occurred within the first 20 minutes of the test, when all the virtual machines were registering. The second XenDesktop controller did not handle as many registrations as the first one, which was also performing the communications with the SCVMM server.

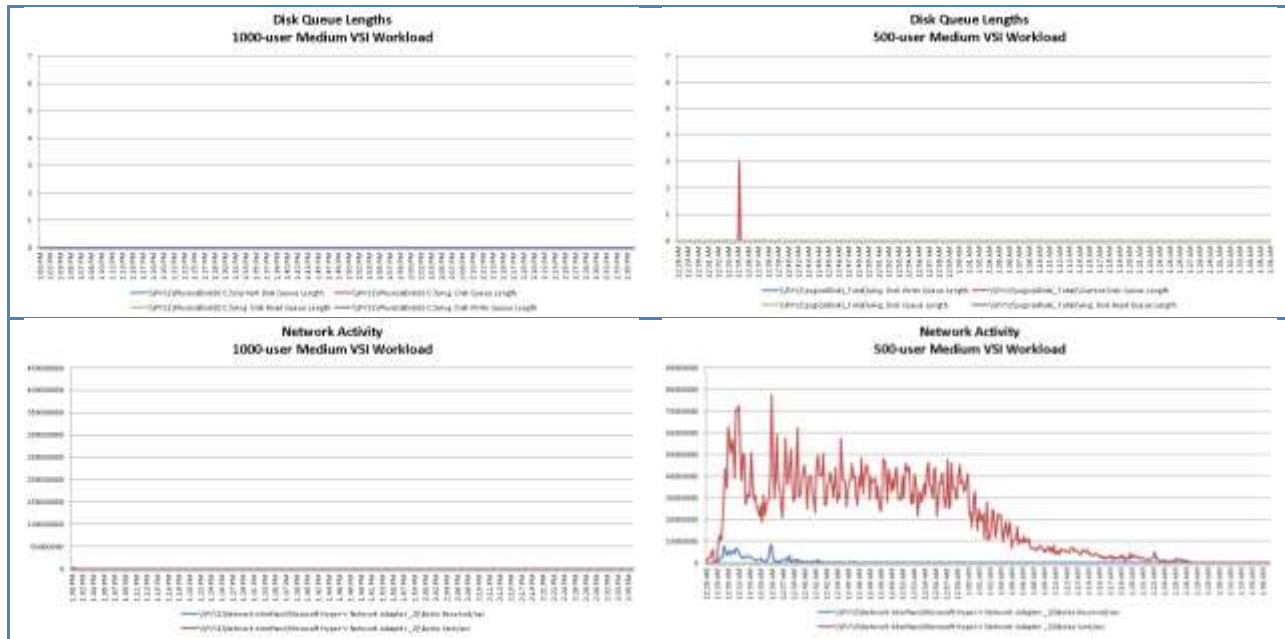




9.4.4. Provisioning Services


From the charts below, it appears that PVS2 was handling the majority of the streaming for the desktops during this test. Certainly, one PVS server is capable of supporting this environment in its entirety, with the second PVS server (PVS1) being available for fault-tolerance.





9.4.5. SQL Servers


From the charts below it is easy to tell that SQL1 was the primary database server in the cluster and that SQL2 was functioning as a standby server. On both servers the F: drive, which was the passthrough drive which holds the data files, was the busiest. The G: drive, which holds the log file, was fairly active on the Secondary server due primary to the log shipping configuration.

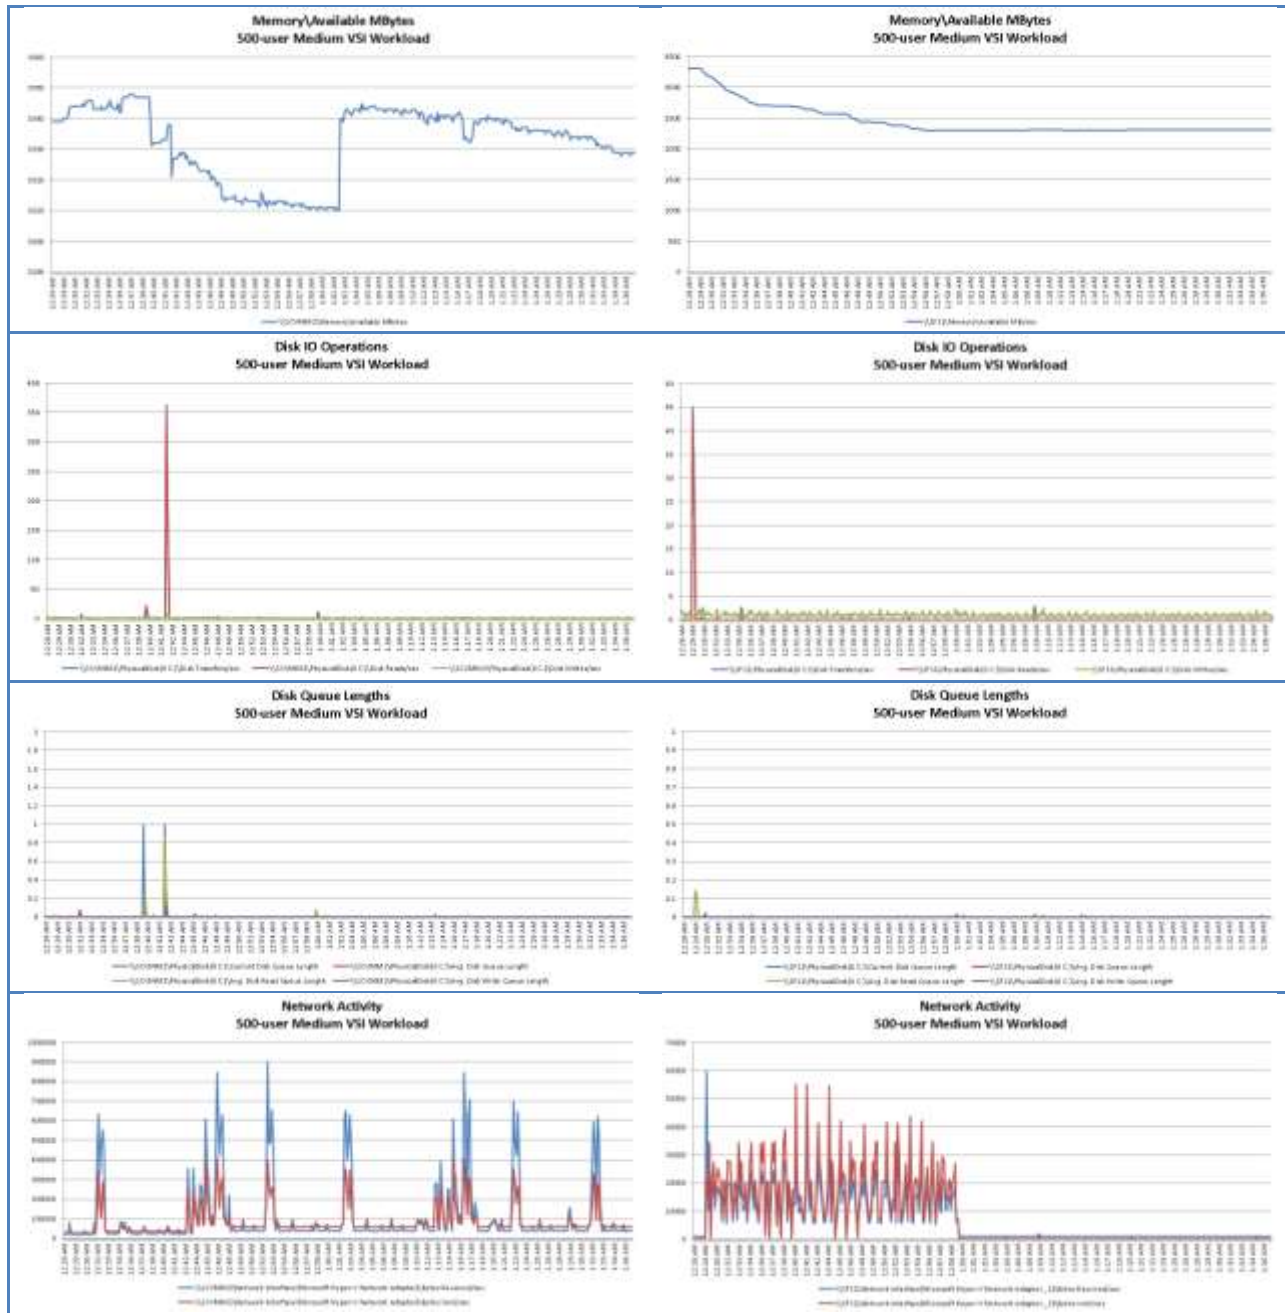




9.4.6. System Center Virtual Machine Manager and StoreFront

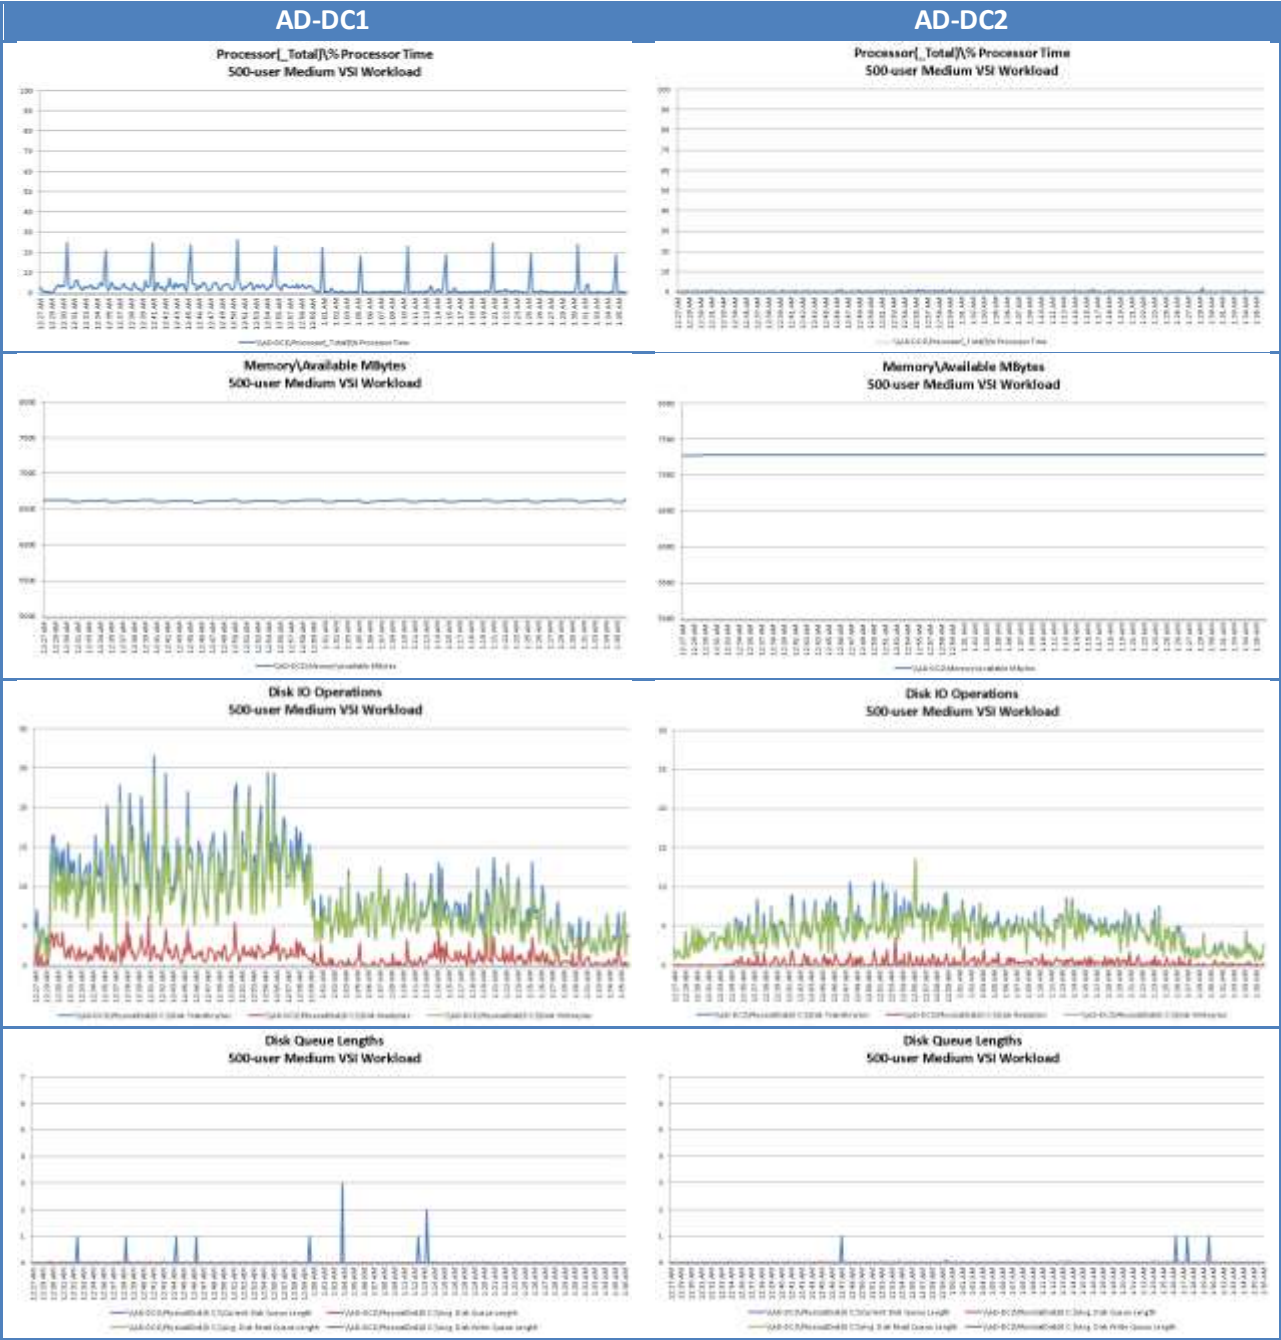
Since neither of these servers have a partner, they are presented together. The StoreFront server is just busy during the ramp-up portion of the test as it is responsible for handing out the tickets to allow the launcher sessions to connect. The SCVMM activity is likely related to the status polling intervals as it goes out and queries the virtual machines at a set interval.






9.4.1. Active Directory

The two Active Directory Controllers were not super busy during the process at all. The graphs are provided for completeness below.






9.5. Cisco UCS 1000-User Scalability Results

This section provides the testing results for the 1000-user testing with five physical blades, four to support the workload and one as a spare. As expected from the 500-user tests, the blades were able to manage the load without overburdening the processor.

Figure 39: VSImax 1000-user Scalability Results




Below are the performance charts for the servers during this test run. The first 20-minutes of the charts represents the time when all the XenDesktop servers and desktops were started and registered with the controllers, also known as the boot phase. The remaining hour is when the Login VSI test was running, with the first 30-minutes being the ramp-up phase and the next 20 minutes being the steady-state phase. The last few minutes would represent the logoff phase of the test.

9.5.1. EMC VNXe Performance


The chart below shows the IOPS on the VNXe array during the boot up of the virtual machines. During the boot up process the VNXe array was used the user profiles and home directories on the CIFS shares.

Figure 40: VNXe Bootup 1000-Users



The chart below shows the IOPS used during the Login VSI test. During the Login VSI test the VNXe was used solely for user home directories and profiles through the CIFS shares on the two storage processors.

Figure 41: VNXe Performance 1000-Users Login VSI




9.5.2. VDI Host

This section contains charts for one of the VDI hosts which provides a representative sample of the host performance. The graphs for the remaining four hosts can be found in the Appendix Section **Error! eference source not found..**

9.5.2.1. Processor

The hosts were quite busy, but never pegged the CPU during the test. The most intensive part was the boot up phase where all 405 virtual machines (81 per blade) were started within a 15-minute window.


Figure 42: Processor Performance VDI Host 1000-User



9.5.2.2. Memory

The hosts were running with about 10GB of available RAM after all the VMs started up. The VMs were configured with dynamic memory. The six RDS hosts had 20GB minimum/startup configured and the 75 desktops had 1.5GB minimum/startup configured.


Figure 43: Memory Performance VDI Host 1000-user



9.5.2.3. Disk


The only disk with interesting information is the E: drive which consisted of two 400GB Enterprise SSD drives in a RAID0 array. The C: drive, which was the boot iSCSI LUN on the VNXe, had almost no activity reported. The E: drive hosted the PVS write-cache drive, so some activity is expected during bootup and the majority would be observed at the peak of the test as the last few desktops logged on. The SSD's are handling just over 6000 IOPS per VDI host during the peak periods, which represents a significant reduction on the IOPS required on the backend VNXe storage array.

Figure 44: Disk IO Performance VDI Host 1000-User



Disk performance can be quantified partially by outstanding disk queue lengths. In this case, with two disks on the RAID0 array, any queue lengths under 2 would be ideal. From the chart below, it appears that the vast majority of the sampling intervals showed the queue lengths to be ideal and in times of heavy load the SSDs were able to catch up.


Figure 45: Disk Queue Lengths VDI Host 1000-User



9.5.2.4. Network

For reference, 80,000,000 bytes per second equates to about 610Mbps. The network on the host, even at peak times was within normal expected loads.

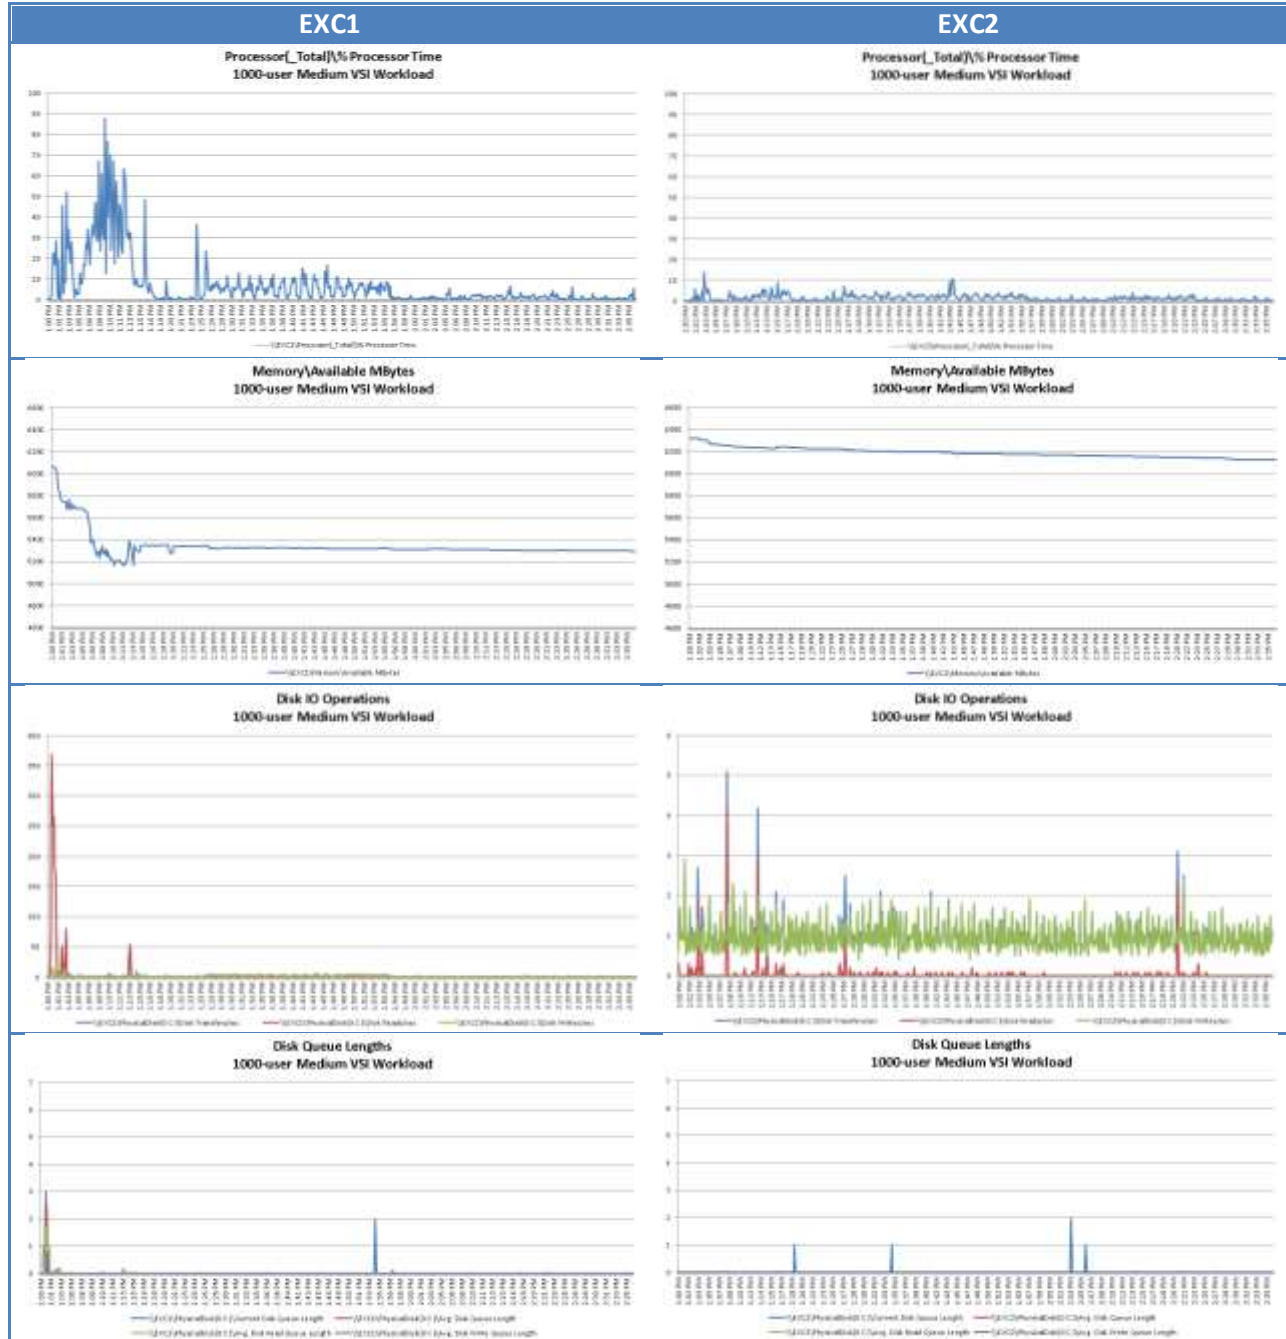
Figure 46: Network Performance VDI Host 1000-User




The VDI hosts performed well within expected performance guidelines.

9.5.3. XenDesktop Controllers


Notice the largest hit for processor time occurred within the first 20 minutes of the test, when all the virtual machines were registering. The second XenDesktop controller did not handle as many registrations as the first one, which was also performing the communications with the SCVMM server.






9.5.4. Provisioning Services


From the charts below, it appears that PVS2 was handling the majority of the streaming for the desktops during this test. Certainly, one PVS server is capable of supporting this environment in its entirety, with the second PVS server (PVS1) being available for fault-tolerance.

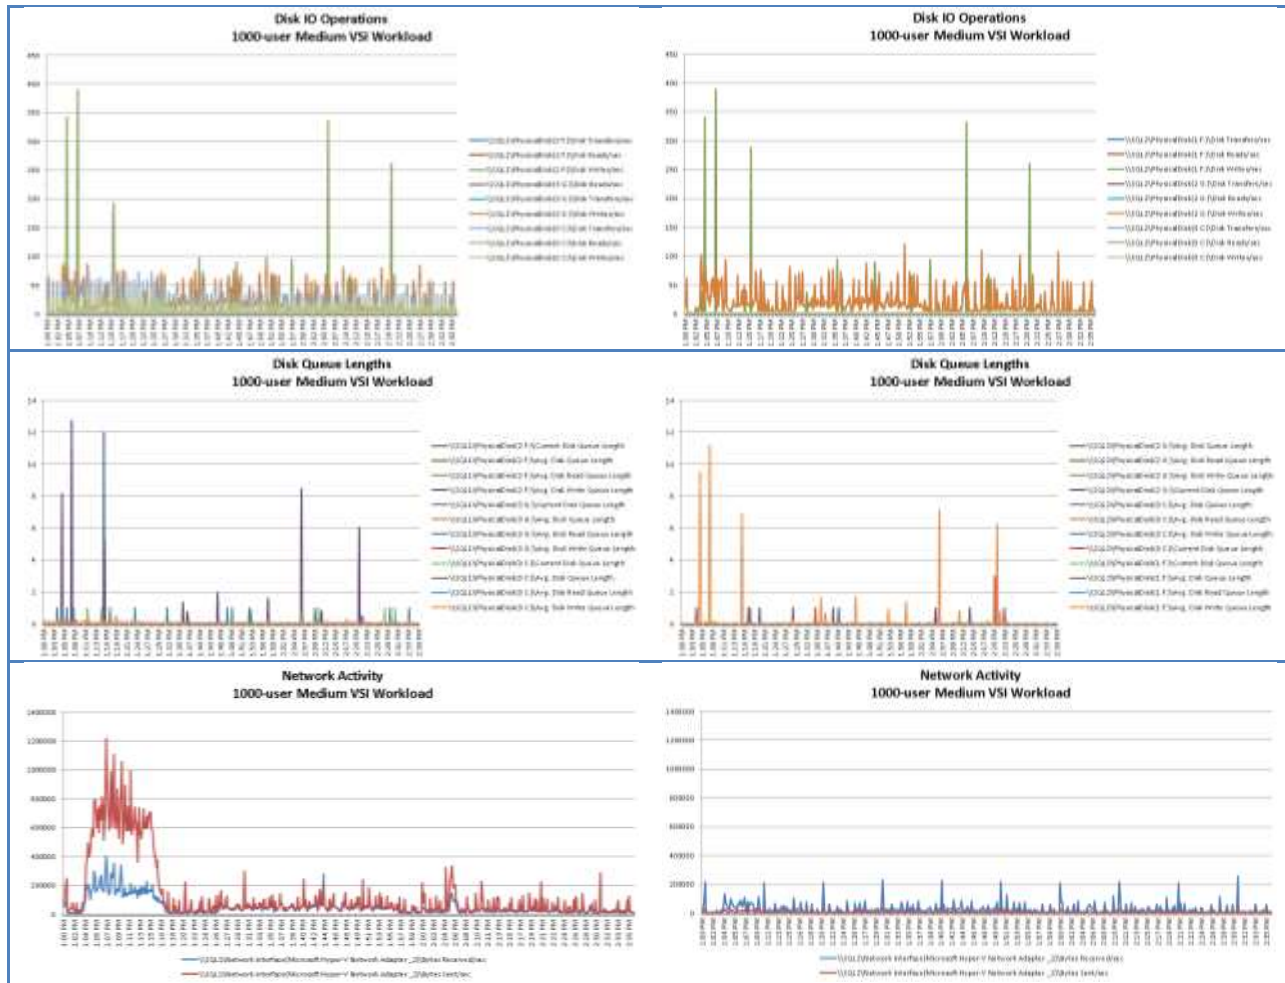




9.5.5. SQL Servers


From the charts below it is easy to tell that SQL1 was the primary database server in the cluster and that SQL2 was functioning as a standby server. On both servers the F: drive, which was the passthrough drive which holds the data files, was the busiest. The G: drive, which holds the log file, was fairly active on the Secondary server due primary to the log shipping configuration.






9.5.6. System Center Virtual Machine Manager and StoreFront

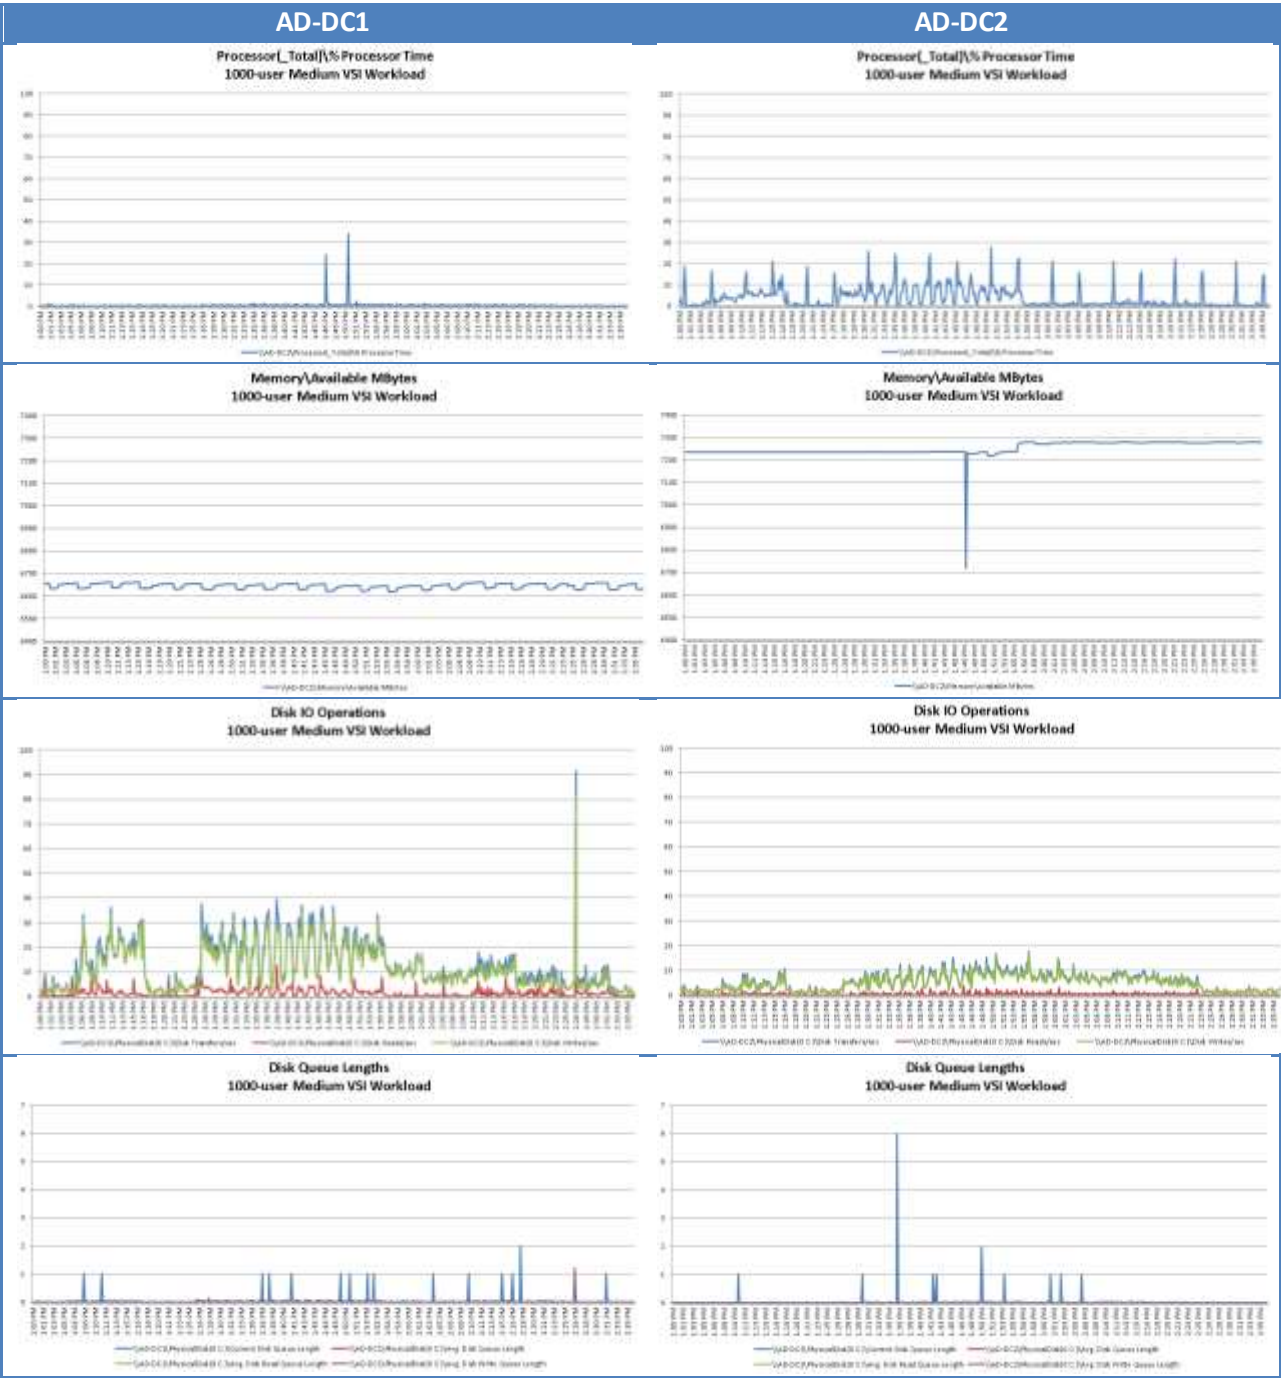
Since neither of these servers have a partner, they are presented together. The SCVMM server is quite busy during the 20-minute boot-up phase and marginally busy at times during the test. The presumption of those busy times during the test phase would be around the polling intervals to check virtual machine status. The StoreFront server is just busy during the ramp-up portion of the test as it is responsible for handing out the tickets to allow the launcher sessions to connect.






9.5.7. Active Directory

The two Active Directory Controllers were not super busy during the process at all. The graphs are provided for completeness below.





10. Scalability Considerations and Guidelines

Many factors should be considered when scaling beyond 1000-user configuration presented in this design. This section provides guidance around those factors to consider.

10.1. Cisco UCS Configuration

The results indicate a fairly linear scalability across the reference architecture. In other words, each additional blade will provide capacity to host an additional 250 users.

- Cisco UCS 2.0 management software supports up to 20 chassis within a single Cisco UCS domain on the second generation Cisco UCS Fabric Interconnect 6248 and 6296 models. A single Cisco UCS domain can grow to 160 blades.
- With Cisco UCS 2.1 and later management software, each Cisco UCS 2.1 Management domain is manageable by Cisco UCS Central, vastly increasing the reach and manageability of the Cisco Unified Computing System.
- As scale grows, the value of the combined Cisco UCS fabric, Cisco Nexus physical switches and Cisco Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100% of the time.
- Based on the number of uplinks from each chassis, we can calculate number of user sessions that can be hosted in a single Cisco UCS domain. Assuming eight links per chassis, four to each 6248, scaling beyond 10 chassis would require additional Cisco UCS fabric interconnects. A 30,000-seat infrastructure, with all support services can be built out of the reference architecture described in this Cisco Validated Design with eight links per chassis and 20 Cisco UCS chassis comprised of eight B200 M3 blades servers in each chassis.

10.2. EMC VNXe Storage Configuration

10.2.1. Capacity planning

You must calculate the required sizes of the iSCSI LUNs and/or CIFS shares in a Hyper-V environment. When calculating the capacity needed in the storage pools, ensure that extra space required for the protection storage for replication and snapshots is considered.

10.2.2. Performance Planning

VNXe has predefined storage pools that can be used based on the performance requirements and the drives in the platform. Solid State Drives (SSDs) provide the greatest performance followed by 15k rpm SAS drives. Do not use 7200 rpm NL-SAS drives for performance –sensitive applications. They are primarily used for capacity.

10.2.3. Scalability Planning

When scaling this solution, keep in mind that the VNXe3300 is rated for 1000 users. To scale above the 1000 users, a different EMC solution, such as the VNX series will be required.

Also, when using version 2.4.1.21171 of the VNXe Software, a hotfix is available that should be installed prior to putting a load on the EMC VNXe. The information for the hotfix is shown below.

ETA169041: VNXe: ESX hosts are unable to boot from iSCSI LUNs or hosts may lose connectivity to data stores

https://emc--c.na5.visual.force.com/apex/KB_ET?id=kA37000000000eF

Later versions of the EMC VNXe software will have this hotfix included.

10.3. Microsoft Windows Server 2012 with Hyper-V 2012 Configuration

Hyper-V in Windows Server® 2012 supports significantly larger configurations of virtual and physical components than in previous releases of Hyper-V. This increased capacity enables you to run Hyper-V on large physical computers and to virtualize high-performance, scale-up workloads. This topic lists the supported maximum configuration for the various components. As you plan your deployment of Hyper-V, consider the maximums that apply to each virtual machine as well as those that apply to the physical computer that runs the Hyper-V role.

10.3.1. Virtual Machines

The following table lists the maximums that apply to each virtual machine.

Table 15: Virtual Machine Maximums

Component	Maximum	Notes
Virtual processors	64	The number of virtual processors supported by a guest operating system might be lower. For more information, see the Hyper-V Overview .
Memory	1 TB	Review the requirements for the specific operating system to determine the minimum and recommended amounts.
Virtual hard disk capacity	64 TB supported by the VHDX format introduced in Windows Server 2012 and Windows® 8; 2040 GB supported by the VHD format.	Each virtual hard disk is stored on physical media as either a .vhdx or a .vhd file, depending on the format used by the virtual hard disk.

Virtual IDE disks	4	The startup disk (sometimes referred to as the boot disk) must be attached to one of the IDE devices. The startup disk can be either a virtual hard disk or a physical disk attached directly to a virtual machine.
Virtual SCSI controllers	4	Use of virtual SCSI devices requires integration services to be installed in the guest operating system. For a list of the guest operating systems for which integration services are available, see the Hyper-V Overview .
Virtual SCSI disks	256	Each SCSI controller supports up to 64 disks, which means that each virtual machine can be configured with as many as 256 virtual SCSI disks. (4 controllers x 64 disks per controller)
Virtual Fiber Channel adapters	4	As a best practice, we recommended that you connect each virtual Fiber Channel Adapter to a different virtual SAN.
Size of physical disks attached directly to a virtual machine	Varies	Maximum size is determined by the guest operating system.
Snapshots	50	The actual number may be lower, depending on the available storage. Each snapshot is stored as an .avhd file that consumes physical storage.
Virtual network adapters	12	<ul style="list-style-type: none"> • 8 can be the “network adapter” type. This type provides better performance and requires a virtual machine driver that is included in the integration services packages. • 4 can be the “legacy network adapter” type. This type emulates a specific physical network adapter and supports the Pre-execution Boot Environment (PXE) to perform network-based installation of an operating system.
Virtual floppy devices	1 virtual floppy drive	None.
Serial (COM) ports	2	None.

10.3.2. Server Running Hyper-V

The following table lists the requirements and maximums that apply to the server running Hyper-V.

Component	Maximum	Notes
Logical processors	320	Both of the following must be available and enabled in the BIOS: <ul style="list-style-type: none"> • Hardware-assisted virtualization • Hardware-enforced Data Execution

Prevention (DEP)		
Virtual processors per logical processor	No ratio imposed by Hyper-V.	None.
Running virtual machines per server	1024	None.
Virtual processors per server	2048	None.
Memory	4 TB	None.
Storage	Limited by what is supported by the management operating system. No limits imposed by Hyper-V.	Note Microsoft supports network-attached storage (NAS) for Hyper-V in Windows Server 2012 when using SMB 3.0. NFS-based storage is not supported.
Virtual storage area networks (SANs)	No limits imposed by Hyper-V	None.
Physical network adapters	No limits imposed by Hyper-V.	None.
Network adapter teams (NIC Teaming)	No limits imposed by Hyper-V.	For more information about NIC Teaming in Windows Server 2012, see NIC Teaming Overview .
Virtual switches	Varies; no limits imposed by Hyper-V.	The practical limit depends on the available computing resources.
Virtual network switch ports per server	Varies; no limits imposed by Hyper-V.	The practical limit depends on the available computing resources.

10.3.4. Failover Clusters and Hyper-V

The following table lists the maximums that apply to highly available servers running Hyper-V. It is important to do capacity planning to ensure that there will be enough hardware resources to run all the virtual machines in a clustered environment. For more information about requirements for failover clusters, see [Failover Clustering Hardware Requirements and Storage Options](#).

Table 16: Failover Cluster Maximums

Component	Maximum	Notes
Nodes per cluster	64	Consider the number of nodes you want to reserve for failover, as well as maintenance tasks such as applying updates. We recommend that you plan for enough resources to allow for 1 node to be reserved for failover, which means it remains idle until another node is failed over to it. (This is sometimes referred to as a passive node.) You can

		increase this number if you want to reserve additional nodes. There is no recommended ratio or multiplier of reserved nodes to active nodes; the only specific requirement is that the total number of nodes in a cluster cannot exceed the maximum of 64.
Running virtual machines per cluster and per node	8,000 per cluster	<p>Several factors can affect the real number of virtual machines that can be run at the same time on one node, such as:</p> <ul style="list-style-type: none"> • Amount of physical memory being used by each virtual machine. • Networking and storage bandwidth. • Number of disk spindles, which affects disk I/O performance.

10.4. Citrix XenDesktop 7 Configuration – Citrix

XenDesktop environments can scale to large numbers. When implementing Citrix XenDesktop hosted shared and hosted virtual desktops have the following considerations:

- Types of Storage in your environment
- Types of desktops that will be deployed
- Data protection requirements
- For Citrix Provisioning Server pooled desktops write cache size and placement

These and other various aspects of scalability considerations described in greater detail in “XenDesktop - Modular Reference Architecture” document and should be a part of any XenDesktop design.

When designing and deploying this Cisco Validated Design environment best practices were followed whenever possible.

The following practices are in particular worth mentioning here:

- Citrix always recommends using N+1 schema for virtualization host servers, to accommodate resiliency. In our RA environment, this recommendation is applied to all host servers.
- All Provisioning Server Network Adapters were configured to have a static IP and management.
- We used the XenDesktop Setup Wizard in PVS. Wizard does an excellent job of creating the desktops automatically and it's possible to run multiple instances of the wizard provided the deployed desktops are placed in different catalogs and have different naming conventions.
- To run wizard at a minimum you need to install the Provisioning Server, the XenDesktop Controller, and configure hosts, as well as create VM templates on all datastores where desktops will be deployed.

11. Other Considerations

This section provides some guidance around working within this reference architecture in an effort to help the reader avoid common pitfalls.

11.1. Power Outages and Boot Sequence

The best defense against power outages is to design the datacenter to have redundant power sources and battery backups such that power is always available. However, occasionally a power outage may be

schedule or last longer than the battery backups can support. In those instances, the boot sequence of the reference architecture components is critical in getting the environment back up in the least amount of time.

The Citrix XenDesktop, Provisioning Services, and Microsoft System Center Virtual Machine Manager components rely heavily on SQL for their databases, so the SQL servers should be online before attempting to boot any of the above mentioned components. Furthermore, the Citrix XenDesktop controllers require access to the System Center Virtual Machine Manager to bring desktops online. Finally, the Citrix XenDesktop controllers should be used to bring the virtual desktops online.

Note: The Hyper-V virtual machine settings should have all VMs (except perhaps the Active Directory domain controllers), set to do nothing on Start Action. This will allow an orderly process from bringing up the resources.

11.1.1. Recommended Boot Sequence

The recommended boot sequence for the environment is as follows:

1. Active Directory infrastructure
2. Storage infrastructure
3. Hyper-V infrastructure hosts
4. SQL Server infrastructure (bring up primary then secondary)
5. Microsoft System Center Virtual Machine Manager
6. Citrix Provisioning Services
7. Citrix XenDesktop Controllers (place delivery groups in maintenance mode)
8. Hyper-V desktop hosts
9. XenDesktop Hosted Shared Desktop servers (RDS)
10. XenDesktop Hosted Virtual Desktops

This sequence insures that all necessary resources are available prior to starting any resources that depend on those resources. Should the order not be followed, it is possible to recover usually by just rebooting the component with issues. Should that not work, the recommended approach is to perform an orderly power down of the environment and bring it up in the sequence provided.

11.2. Microsoft Cluster Maintenance

Enable cluster-aware updates so that both cluster servers are not rebooted simultaneously. Always bring up the primary server (whichever one was in the primary role at the time of shutdown) first so it can claim control of the cluster resources, such as disks. When shutting down Microsoft clusters, always shutdown the secondary nodes first, then finally shutdown the primary nodes. This order is especially important when working with SQL server and AlwaysOn groups.

11.3. SQL Server AlwaysOn Groups

When using SQL Server AlwaysOn groups, be sure you have completed the following:

- Enable database containment at the SQL Server level, before creating the XenDesktop database.

- Enable database containment (and full recovery model) on the XenDesktop database when creating it.
- Use the SQL Server AlwaysOn listener for the database location, rather than the SQL Server Cluster resource name or IP address, so that the failover will work correctly.

If unexpected connectivity issues occur with the AlwaysOn Group databases, sometimes the fastest recovery is to just remove the database from the group and then re-add it using the wizard.

In this environment, two AlwaysOn groups were created. The first AlwaysOn group was for the Microsoft System Center Virtual Machine Manager database and second group held the two Citrix databases: XenDesktop and Provisioning Services.

12. References

12.1. Cisco Reference Documents – Cisco

Cisco Unified Computing System Manager Home Page

<http://www.cisco.com/en/US/products/ps10281/index.html>

Cisco UCS B200 M3 Rack Server Resources

<http://buildprice.cisco.com/catalog/ucs/models/B200M3>

<http://www.cisco.com/en/US/products/ps10280/index.html>

Cisco UCS 6200 Series Fabric Interconnects

<http://www.cisco.com/en/US/products/ps11544/index.html>

Download Software for Cisco Nexus 1000V Switch for Microsoft Hyper-V – ACCOUNT REQUIRED

[http://software.cisco.com/download/release.html?mdfid=284786025&flowid=&softwareid=282088129&os=null&release=5.2\(1\)SM1\(5.1\)&relind=null&rellifecycle=null&reltype=null](http://software.cisco.com/download/release.html?mdfid=284786025&flowid=&softwareid=282088129&os=null&release=5.2(1)SM1(5.1)&relind=null&rellifecycle=null&reltype=null)

Download Cisco UCS Manager and Blade Software Version 2.1(3a) – ACCOUNT REQUIRED

[http://software.cisco.com/download/release.html?mdfid=283853163&flowid=&softwareid=283655681&os=null&release=2.1\(3a\)&relind=null&rellifecycle=null&reltype=null](http://software.cisco.com/download/release.html?mdfid=283853163&flowid=&softwareid=283655681&os=null&release=2.1(3a)&relind=null&rellifecycle=null&reltype=null)

Download Cisco UCS Central Software Version 2.1(3a) – ACCOUNT REQUIRED

[http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=2.1\(3a\)&flowid=22121](http://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=2.1(3a)&flowid=22121)

12.2. EMC Reference Documents – EMC

- EMC Infrastructure for Citrix XenDesktop 7, EMC VNX Series (NFS and FC), Citrix XenDesktop 7, VMware vSphere 5.1—Reference Architecture

- VSPEX Proven Infrastructure: EMC VSPEX End-User Computing, Citrix XenDesktop 5.6 with VMware vSphere 5.1 for up to 250 Virtual Desktops Enabled by EMC VNXe and EMC Next-Generation Backup
- VSPEX Proven Infrastructure: EMC VSPEX End-User Computing, Citrix XenDesktop 5.6 with VMware vSphere 5.1 for up to 2000 Virtual Desktops Enabled by EMC VNX and EMC Next-Generation Backup
- White Paper: Sizing EMC VNX Series for VDI Workload—An Architectural Guideline
- Reference Architecture: EMC Infrastructure for Citrix XenDesktop 5.6, EMC VNX Series (NFS), VMware vSphere 5.0, Citrix XenDesktop 5.6, and Citrix Profile Manager 4.1
- Proven Solutions Guide: EMC Infrastructure for Citrix XenDesktop 5.6 — EMC VNX Series (NFS), VMware vSphere 5.0, Citrix XenDesktop 5.6, and Citrix Profile Manager 4.1
- EMC VNXe Series Storage Systems—A Detailed Review
- EMC VNXe High Availability—Overview
- EMC VNXe Data Protection—Overview
- Deployment Guide: Microsoft Windows Server 2008 R2 Hyper-V on EMC VNXe Series

12.3. Microsoft Reference Documents

For more information on the products used in this Cisco Validated Design, please see:

- **Microsoft Windows Server 2012 with Hyper-V**

[Windows Server 2012](#)

[Windows Server 2012 capabilities](#)

[Failover Clustering feature in Windows Server 2012](#)

[What's New In Networking For Windows Server 2012](#)

[Networking Features in Windows Server 2012](#)

[Server Virtualization Features in Windows Server 2012](#)

[Windows Server Hyper-V Technical White Paper](#)

[Competitive Advantages of Windows Server 2012 with Hyper-V](#)

[Competitive Advantages of Hyper-V Server 2012 over the VMware vSphere Hypervisor.](#)

- **Microsoft SQL Server 2012**

[What's New in SQL Server 2012](#)

[SQL Server AlwaysOn Availability Groups](#)

[Failover Clustering and AlwaysOn Availability Groups](#)

- **Microsoft System Center**

[System Center 2012](#)


[What's New in System Center 2012 Virtual Machine Manager SP1](#)

12.4. Citrix Reference Documents


- Citrix Product Downloads
<http://www.citrix.com/downloads/xendesktop.html>
- Citrix Knowledge Center
<http://support.citrix.com>
- Definitive Guide to XenDesktop 7
<http://support.citrix.com/proddocs/topic/xendesktop/cds-xd-7landing-page.html>
- Citrix Provisioning Services
<http://support.citrix.com/proddocs/topic/technologies/pvs-provisioning.html>
- Citrix UPM 5.0
<http://support.citrix.com/proddocs/topic/user-profile-manager-5-x/upm-wrapper-kib.html>
- Login VSI
<http://www.loginvsi.com/documentation/v3>

13. Appendix


13.1. Performance Charts for Scalability Tests




Disk IO Operations - SSD RAID0 500-user Medium VSI Workload




Disk Queue Lengths - SSD RAID0 500-user Medium VSI Workload




Network Activity 500-user Medium VSI Workload




Hyper-V Hypervisor Logical Processor (_Total)\% Total Run Time 500-user Medium VSI Workload




Memory\Available MBytes 500-user Medium VSI Workload




Disk IO Operations - SSD RAID0 500-user Medium VSI Workload



Disk Queue Lengths - SSD RAID0 500-user Medium VSI Workload




Network Activity 500-user Medium VSI Workload




Hyper-V Hypervisor Logical Processor (_Total)\% Total Run Time
1000-user Medium VSI Workload




Memory\Available MBytes
1000-user Medium VSI Workload




Disk IO Operations - SSD RAID0 1000-user Medium VSI Workload




Disk Queue Lengths - SSD RAID0 1000-user Medium VSI Workload




Network Activity 1000-user Medium VSI Workload




Hyper-V Hypervisor Logical Processor (_Total)\% Total Run Time 1000-user Medium VSI Workload




Memory\Available MBytes 1000-user Medium VSI Workload




Disk IO Operations - SSD RAID0 1000-user Medium VSI Workload



Disk Queue Lengths - SSD RAID0 1000-user Medium VSI Workload




Network Activity 1000-user Medium VSI Workload




Hyper-V Hypervisor Logical Processor (_Total)\% Total Run Time
1000-user Medium VSI Workload




Memory\Available MBytes
1000-user Medium VSI Workload




Disk IO Operations - SSD RAID0 1000-user Medium VSI Workload




Disk Queue Lengths - SSD RAID0 1000-user Medium VSI Workload




Network Activity 1000-user Medium VSI Workload




Hyper-V Hypervisor Logical Processor (Total)\% Total Run Time 1000-user Medium VSI Workload




Memory\Available MBytes
1000-user Medium VSI Workload




Disk IO Operations - SSD RAID0
1000-user Medium VSI Workload



Disk Queue Lengths - SSD RAID0 1000-user Medium VSI Workload



Network Activity 1000-user Medium VSI Workload



13.2. Sample Cisco Nexus 6248-UP Configurations

!Command: show running-config
!Time: Thu Nov 21 10:54:46 2013

version 5.0(3)N2(2.11.3a)
feature fcoe
feature adapter-fex

no feature telnet
no telnet server enable
feature tacacs+
cfs ipv4 distribute
cfs eth distribute
feature private-vlan
feature port-security

feature lacp	rule 3 permit read-write feature
feature lldp	sam-ls-security-policy
feature fex	rule 2 permit read-write feature
logging level assoc_mgr 2	sam-ls-security
	rule 1 permit read-write feature
logging level aaa 2	sam-pn-security
logging level afm 2	role name server-compute
	rule 3 permit read-write feature
logging level cfs 2	sam-ls-compute
logging level enm 2	rule 2 permit read-write feature
logging level fex 2	sam-ls-server-oper
logging level fwm 2	rule 1 permit read-write feature
logging level msp 2	sam-ls-server-policy
logging level npv 2	role name server-profile
logging level pfm 2	rule 6 permit read-write feature
logging level vms 2	sam-ls-server-oper
logging level evmc 2	
logging level port 2	rule 5 permit read-write feature
logging level vshd 2	sam-ls-ext-access
logging level ethpm 2	rule 4 permit read-write feature
logging level track 2	sam-ls-server-policy
logging level xmlma 2	rule 3 permit read-write feature
logging level licmgr 2	sam-ls-config-policy
logging level radius 2	rule 2 permit read-write feature
logging level tacacs 2	sam-ls-server
logging level bootvar 2	rule 1 permit read-write feature
logging level monitor 2	sam-ls-config
logging level fcdomain 2	role name operations
logging level ascii-cfg 2	rule 2 permit read-write feature
logging level provision 2	sam-fault
	rule 1 permit read-write feature
logging level securityd 2	sam-operations
logging level pltfm_config 2	role name read-only
logging level port-channel 2	rule 1 permit read-write feature
logging level private-vlan 2	sam-read-only
logging level spanning-tree 2	role name KVM-Only
logging level port-resources 2	rule 1 permit read-write feature
role name server-equipment	sam-ls-ext-access
rule 3 permit read-write feature	role name network
sam-pn-maintenance	rule 12 permit read-write feature
rule 2 permit read-write feature	sam-ls-qos-policy
sam-pn-policy	rule 11 permit read-write feature
rule 1 permit read-write feature	sam-ls-network-policy
sam-pn-equipment	rule 10 permit read-write feature
role name facility-manager	sam-ls-qos
rule 1 permit read-write feature	rule 9 permit read-write feature
sam-power-mgmt	sam-ls-network
role name server-security	

rule 8 permit read-write feature
sam-ext-lan-qos
rule 7 permit read-write feature
sam-ext-lan-security
rule 6 permit read-write feature
sam-ext-lan-policy
rule 5 permit read-write feature
sam-ext-lan-config
rule 4 permit read-write feature
sam-pod-qos

rule 3 permit read-write feature
sam-pod-security
rule 2 permit read-write feature
sam-pod-policy
rule 1 permit read-write feature
sam-pod-config
role name storage
rule 6 permit read-write feature
sam-ls-storage-policy
rule 5 permit read-write feature
sam-ls-storage
rule 4 permit read-write feature
sam-ext-san-qos
rule 3 permit read-write feature
sam-ext-san-security
rule 2 permit read-write feature
sam-ext-san-policy
rule 1 permit read-write feature
sam-ext-san-config
role name admin
rule 1 permit read-write feature
sam-admin
role name aaa
rule 1 permit read-write feature
sam-aaa
no password strength-check

banner motd #Cisco UCS 6200
Series Fabric Interconnect
#

ip domain-lookup
ip name-server 171.70.168.183
171.68.226.120 175.25.205.7
aaa group server tacacs+ tacacs
switchname UCS-EXC-HyperV-A
logging event link-status default

errdisable recovery interval 60
errdisable recovery cause link-flap
errdisable recovery cause udd
errdisable recovery cause
bpduguard
errdisable recovery cause loopback
errdisable recovery cause pause-
rate-limit
class-map type qos class-fcoe

class-map type qos match-all class-
gold
match cos 4
class-map type qos match-all class-
bronze
match cos 1
class-map type qos match-all class-
silver
match cos 2
class-map type qos match-all class-
platinum
match cos 5
class-map type queuing class-fcoe
match qos-group 1
class-map type queuing class-gold
match qos-group 3
class-map type queuing class-
bronze
match qos-group 5
class-map type queuing class-silver
match qos-group 4
class-map type queuing class-
platinum
match qos-group 2
class-map type queuing class-all-
flood
match qos-group 2
class-map type queuing class-ip-
multicast
match qos-group 2

policy-map type qos
system_qos_policy
class class-platinum
set qos-group 2
class class-silver
set qos-group 4
class class-bronze

```

    set qos-group 5
class class-gold
    set qos-group 3
class class-fcoe
    set qos-group 1
policy-map type queuing
system_q_in_policy
class type queuing class-fcoe
    bandwidth percent 14
class type queuing class-platinum
    bandwidth percent 22
class type queuing class-gold
    bandwidth percent 20
class type queuing class-silver
    bandwidth percent 18
class type queuing class-bronze
    bandwidth percent 15

class type queuing class-default
    bandwidth percent 11
policy-map type queuing
system_q_out_policy
class type queuing class-fcoe
    bandwidth percent 14
class type queuing class-platinum
    bandwidth percent 22
class type queuing class-gold
    bandwidth percent 20
class type queuing class-silver
    bandwidth percent 18
class type queuing class-bronze
    bandwidth percent 15
class type queuing class-default
    bandwidth percent 11
policy-map type queuing org-
root/ep-qos-Gold
class type queuing class-default
    bandwidth percent 100
    shape 40000000 kbps 10240
policy-map type queuing org-
root/ep-qos-Bronze
class type queuing class-default
    bandwidth percent 100

    shape 40000000 kbps 10240
policy-map type queuing org-
root/ep-qos-Silver
class type queuing class-default

```

```

    bandwidth percent 100
    shape 40000000 kbps 10240
policy-map type queuing org-
root/ep-qos-Platinum
class type queuing class-default
    bandwidth percent 100
    shape 40000000 kbps 10240
class-map type network-qos class-
fcoe
    match qos-group 1
class-map type network-qos class-
gold
    match qos-group 3
class-map type network-qos class-
bronze
    match qos-group 5
class-map type network-qos class-
silver
    match qos-group 4
class-map type network-qos class-
platinum
    match qos-group 2
class-map type network-qos class-
all-flood
    match qos-group 2
class-map type network-qos class-
ip-multicast

    match qos-group 2
policy-map type network-qos
system_nq_policy
class type network-qos class-
platinum
    mtu 9000
class type network-qos class-silver
    mtu 9000
class type network-qos class-
bronze
    mtu 9000
class type network-qos class-gold
    mtu 9000
class type network-qos class-fcoe
    pause no-drop
    mtu 2158
class type network-qos class-
default
system qos

```

```

    service-policy type qos input
system_qos_policy
    service-policy type queuing input
system_q_in_policy
    service-policy type queuing output
system_q_out_policy
    service-policy type network-qos
system_nq_policy
fex 1
    pinning max-links 1
    description "FEX0001"

fex 2
    pinning max-links 1
    description "FEX0002"
fex management-instance
39619c74-90b7-11e2-8831-
547feef80284 fabric 1
no snmp-server enable traps entity
entity_mib_change
no snmp-server enable traps entity
entity_module_status_change
no snmp-server enable traps entity
entity_power_status_change
no snmp-server enable traps entity
entity_module_inserted
no snmp-server enable traps entity
entity_module_removed
no snmp-server enable traps entity
entity_unrecognised_module
no snmp-server enable traps entity
entity_fan_status_change
no snmp-server enable traps link
linkDown
no snmp-server enable traps link
linkUp
no snmp-server enable traps link
extended-linkDown
no snmp-server enable traps link
extended-linkUp
no snmp-server enable traps link
cieLinkDown
no snmp-server enable traps link
cieLinkUp
no snmp-server enable traps link
connUnitPortStatusChange
no snmp-server enable traps link
fcTrunkIfUpNotify

```

```

no snmp-server enable traps link
fcTrunkIfDownNotify
no snmp-server enable traps link
delayed-link-state-change
no snmp-server enable traps link
fcot-inserted

no snmp-server enable traps link
fcot-removed
no snmp-server enable traps rf
redundancy_framework
no snmp-server enable traps license
notify-license-expiry
no snmp-server enable traps license
notify-no-license-for-feature
no snmp-server enable traps license
notify-licensefile-missing
no snmp-server enable traps license
notify-license-expiry-warning
no snmp-server enable traps rmon
risingAlarm
no snmp-server enable traps rmon
fallingAlarm
no snmp-server enable traps rmon
hcRisingAlarm
no snmp-server enable traps rmon
hcFallingAlarm
snmp-server enable traps entity fru
ntp server 171.68.10.80
ntp server 171.68.10.150
no aaa user default-role

vrf context management
    ip name-server 171.70.168.183
171.68.226.120 175.25.205.7
vlan 1,60-66,999
vlan 4044
    name SAM-vlan-management
vlan 4047
    name SAM-vlan-boot

vlan 4048
    fcoe vsan 1
    name fcoe-vsan-4048
vlan 4049
port-channel load-balance ethernet
invalid invalid-hash

```



```

port-profile type vethernet
ucsm_internal_rackserver_portprofi
le
  switchport trunk allowed vlan
4044
  switchport mode trunk
  no shutdown
  vmware port-group
  max-ports 320
  state enabled

```

```

logging level sysmgr 2

```

```

interface port-channel25
  description U: Uplink
  switchport mode trunk
  pinning border
  switchport trunk allowed vlan
1,60-66,999
  speed 10000

```

```

interface port-channel1280
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1281
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1282
  switchport mode vntag
  no pinning server sticky
  speed 10000

```

```

interface port-channel1283
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky

  speed 10000

```

```

interface port-channel1284
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1304
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1310
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1317
  switchport mode vntag

  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1321
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1322
  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1329
  switchport mode vntag
  no pinning server sticky
  speed 10000

```

```

interface port-channel1333

  switchport mode vntag
  switchport vntag max-vifs 118
  no pinning server sticky
  speed 10000

```

```

interface port-channel1339
 switchport mode vntag
 switchport vntag max-vifs 118
 no pinning server sticky
 speed 10000

```

```

interface port-channel1341
 switchport mode vntag
 switchport vntag max-vifs 118
 no pinning server sticky
 speed 10000

```

```

interface port-channel1342
 switchport mode vntag
 switchport vntag max-vifs 118
 no pinning server sticky
 speed 10000

```

```

interface port-channel1344
 switchport mode vntag
 switchport vntag max-vifs 118
 no pinning server sticky
 speed 10000
 feature npv
 feature npiv

```

```

interface Ethernet1/1
 description S: Server
 no pinning server sticky
 switchport mode fex-fabric
 fex associate 1 chassis-serial
 FOX1642H32G module-serial
 FCH16487E86 module-sl
 ot left
 no shutdown

```

```

interface Ethernet1/2
 description S: Server
 no pinning server sticky
 switchport mode fex-fabric

```

```

 fex associate 1 chassis-serial
 FOX1642H32G module-serial
 FCH16487E86 module-sl

```

```

ot left
 no shutdown

```

```

interface Ethernet1/3
 description S: Server
 no pinning server sticky
 switchport mode fex-fabric
 fex associate 1 chassis-serial
 FOX1642H32G module-serial
 FCH16487E86 module-sl
 ot left
 no shutdown

```

```

interface Ethernet1/4
 description S: Server
 no pinning server sticky
 switchport mode fex-fabric
 fex associate 1 chassis-serial
 FOX1642H32G module-serial
 FCH16487E86 module-sl
 ot left
 no shutdown

```

```

interface Ethernet1/5
 description S: Server

 no pinning server sticky
 switchport mode fex-fabric
 fex associate 2 chassis-serial
 FOX1652G334 module-serial
 FCH1650JRUP module-sl
 ot left
 no shutdown

```

```

interface Ethernet1/6
 description S: Server
 no pinning server sticky
 switchport mode fex-fabric
 fex associate 2 chassis-serial
 FOX1652G334 module-serial
 FCH1650JRUP module-sl
 ot left
 no shutdown

```

```

interface Ethernet1/7
 description S: Server
 no pinning server sticky
 switchport mode fex-fabric

```

```

    fex associate 2 chassis-serial
    FOX1652G334 module-serial
    FCH1650JRUP module-sl
    ot left
    no shutdown

interface Ethernet1/8
    description S: Server
    no pinning server sticky
    switchport mode fex-fabric
    fex associate 2 chassis-serial
    FOX1652G334 module-serial
    FCH1650JRUP module-sl
    ot left
    no shutdown

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
    description A: Appliance
    untagged cos 5
    no pinning server sticky
    pinning server pinning-failure link-
    down
    no cdp enable
    switchport mode trunk
    switchport trunk allowed vlan 65-
    66
    no shutdown

interface Ethernet1/18
    description A: Appliance

```

```

    untagged cos 5
    no pinning server sticky
    pinning server pinning-failure link-
    down
    no cdp enable
    switchport mode trunk
    switchport trunk allowed vlan 65-
    66
    no shutdown

interface Ethernet1/19
    description U: Uplink
    pinning border
    switchport mode trunk
    switchport trunk allowed vlan
    1,60-66,999
    no shutdown

interface Ethernet1/20
    description U: Uplink
    pinning border
    switchport mode trunk
    switchport trunk allowed vlan
    1,60-66,999
    no shutdown

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25
    description U: Uplink
    pinning border
    switchport mode trunk
    switchport trunk allowed vlan
    1,60-66,999
    channel-group 25 mode active
    no shutdown

interface Ethernet1/26
    description U: Uplink
    pinning border

```

```
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
channel-group 25 mode active
no shutdown
```

```
interface Ethernet1/27
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
channel-group 25 mode active

no shutdown
```

```
interface Ethernet1/28
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
channel-group 25 mode active
no shutdown
```

```
interface Ethernet1/29
```

```
interface Ethernet1/30
```

```
interface Ethernet1/31
```

```
interface Ethernet1/32
```

```
interface Ethernet2/1
```

```
interface Ethernet2/2
```

```
interface Ethernet2/3
```

```
interface Ethernet2/4
```

```
interface Ethernet2/5
```

```
interface Ethernet2/6
```

```
interface Ethernet2/7
```

```
interface Ethernet2/8
```

```
interface Ethernet2/9
```

```
interface Ethernet2/10
```

```
interface Ethernet2/11
```

```
interface Ethernet2/12
```

```
interface Ethernet2/13
```

```
interface Ethernet2/14
```

```
interface Ethernet2/15
```

```
interface Ethernet2/16
```

```
interface mgmt0
shutdown force
ip address 10.60.0.11/24
```

```
interface Ethernet1/1/1
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/1
channel-group 1317
no shutdown
```

```
interface Ethernet1/1/2
no pinning server sticky
```

```
interface Ethernet1/1/3
```

```
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/1
channel-group 1317
no shutdown
```

```
interface Ethernet1/1/4
no pinning server sticky
```

```
interface Ethernet1/1/5
switchport vntag max-vifs 118
```

```

no pinning server sticky
switchport mode vntag
fabric-interface Eth1/2
channel-group 1283
no shutdown

interface Ethernet1/1/6
no pinning server sticky

interface Ethernet1/1/7

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/2
channel-group 1283
no shutdown

interface Ethernet1/1/8
no pinning server sticky

interface Ethernet1/1/9
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/3
channel-group 1280
no shutdown

interface Ethernet1/1/10
no pinning server sticky

interface Ethernet1/1/11

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/3
channel-group 1280
no shutdown

interface Ethernet1/1/12
no pinning server sticky

interface Ethernet1/1/13
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag

fabric-interface Eth1/4
channel-group 1281
no shutdown

interface Ethernet1/1/14
no pinning server sticky

interface Ethernet1/1/15

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/4
channel-group 1281
no shutdown

interface Ethernet1/1/16
no pinning server sticky

interface Ethernet1/1/17
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/1
channel-group 1341
no shutdown

interface Ethernet1/1/18
no pinning server sticky

interface Ethernet1/1/19

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/1
channel-group 1341
no shutdown

interface Ethernet1/1/20
no pinning server sticky

interface Ethernet1/1/21
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/2
channel-group 1282

```

```

no shutdown

interface Ethernet1/1/22
no pinning server sticky

interface Ethernet1/1/23

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/2
channel-group 1282
no shutdown

interface Ethernet1/1/24
no pinning server sticky

interface Ethernet1/1/25
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/3
channel-group 1344
no shutdown

interface Ethernet1/1/26
no pinning server sticky

interface Ethernet1/1/27

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/3
channel-group 1344
no shutdown

interface Ethernet1/1/28
no pinning server sticky

interface Ethernet1/1/29
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/4
channel-group 1284
no shutdown

interface Ethernet1/1/30
no pinning server sticky

interface Ethernet1/1/31

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/4
channel-group 1284
no shutdown

interface Ethernet1/1/32
no pinning server sticky

interface Ethernet1/1/33
no pinning server sticky
switchport mode trunk
switchport trunk native vlan 4044
switchport trunk allowed vlan
4044
no shutdown

interface Ethernet2/1/1
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/5

channel-group 1321
no shutdown

interface Ethernet2/1/2
no pinning server sticky

interface Ethernet2/1/3
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/5
channel-group 1321
no shutdown

interface Ethernet2/1/4
no pinning server sticky

interface Ethernet2/1/5
switchport vntag max-vifs 118

```

no pinning server sticky	fabric-interface Eth1/8
switchport mode vntag	channel-group 1339
fabric-interface Eth1/6	no shutdown
channel-group 1322	interface Ethernet2/1/14
no shutdown	no pinning server sticky
interface Ethernet2/1/6	interface Ethernet2/1/15
no pinning server sticky	switchport vntag max-vifs 118
interface Ethernet2/1/7	no pinning server sticky
switchport vntag max-vifs 118	switchport mode vntag
no pinning server sticky	fabric-interface Eth1/8
switchport mode vntag	channel-group 1339
fabric-interface Eth1/6	no shutdown
channel-group 1322	interface Ethernet2/1/16
no shutdown	no pinning server sticky
interface Ethernet2/1/8	interface Ethernet2/1/17
no pinning server sticky	switchport vntag max-vifs 118
interface Ethernet2/1/9	no pinning server sticky
switchport vntag max-vifs 118	switchport mode vntag
no pinning server sticky	fabric-interface Eth1/5
switchport mode vntag	channel-group 1329
fabric-interface Eth1/7	no shutdown
channel-group 1333	interface Ethernet2/1/18
no shutdown	no pinning server sticky
interface Ethernet2/1/10	interface Ethernet2/1/19
no pinning server sticky	switchport vntag max-vifs 118
interface Ethernet2/1/11	no pinning server sticky
switchport vntag max-vifs 118	switchport mode vntag
no pinning server sticky	fabric-interface Eth1/5
switchport mode vntag	channel-group 1329
fabric-interface Eth1/7	no shutdown
channel-group 1333	interface Ethernet2/1/20
no shutdown	no pinning server sticky
interface Ethernet2/1/12	interface Ethernet2/1/21
no pinning server sticky	switchport vntag max-vifs 118
interface Ethernet2/1/13	no pinning server sticky
switchport vntag max-vifs 118	switchport mode vntag
no pinning server sticky	fabric-interface Eth1/6
switchport mode vntag	

```

channel-group 1304
no shutdown

interface Ethernet2/1/22
no pinning server sticky

interface Ethernet2/1/23
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/6
channel-group 1304
no shutdown

interface Ethernet2/1/24
no pinning server sticky

interface Ethernet2/1/25
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/7

channel-group 1342
no shutdown

interface Ethernet2/1/26
no pinning server sticky

interface Ethernet2/1/27
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/7
channel-group 1342
no shutdown

interface Ethernet2/1/28
no pinning server sticky

interface Ethernet2/1/29
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/8

channel-group 1310

```

```

no shutdown

interface Ethernet2/1/30
no pinning server sticky

interface Ethernet2/1/31
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/8
channel-group 1310
no shutdown

interface Ethernet2/1/32
no pinning server sticky

interface Ethernet2/1/33
no pinning server sticky
switchport mode trunk
switchport trunk native vlan 4044
switchport trunk allowed vlan
4044

no shutdown

interface Vethernet1360
description server 1/8, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1284
channel 1360
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1363
description server 1/8, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down

```



```

no cdp enable

switchport trunk allowed vlan 61
bind interface port-channel1284
channel 1363
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1364
description server 1/8, VNICeth3
switchport mode trunk
untagged cos 4
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 62
bind interface port-channel1284
channel 1364
service-policy type queuing input
org-root/ep-qos-Gold
no shutdown

interface Vethernet1366
description server 1/8, VNICeth4
switchport mode trunk
untagged cos 1
no pinning server sticky

pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 63
switchport trunk allowed vlan 63
bind interface port-channel1284
channel 1366
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown

interface Vethernet1369
description server 1/8, VNICeth5
switchport mode trunk
untagged cos 1
no pinning server sticky
pinning server pinning-failure link-
down

```

```

no cdp enable
switchport trunk native vlan 64
switchport trunk allowed vlan 64
bind interface port-channel1284
channel 1369
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown

interface Vethernet1370
description server 1/8, VNICeth6
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1284
channel 1370
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet1373
description server 1/8, VNICeth7
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1284
channel 1373

service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet1374
description server 2/8, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky

```

pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1310
channel 1374
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1377
description server 2/8, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down

no cdp enable
switchport trunk allowed vlan 61
bind interface port-channel1310
channel 1377
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1378
description server 2/8, VNICeth3
switchport mode trunk
untagged cos 4
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk allowed vlan 62
bind interface port-channel1310
channel 1378
service-policy type queuing input
org-root/ep-qos-Gold
no shutdown

interface Vethernet1380
description server 2/8, VNICeth4
switchport mode trunk
untagged cos 1

no pinning server sticky

pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 63
switchport trunk allowed vlan 63
bind interface port-channel1310
channel 1380
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown

interface Vethernet1383
description server 2/8, VNICeth5
switchport mode trunk
untagged cos 1
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 64
switchport trunk allowed vlan 64
bind interface port-channel1310
channel 1383
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown

interface Vethernet1384
description server 2/8, VNICeth6
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1310
channel 1384
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet1387
description server 2/8, VNICeth7
switchport mode trunk
untagged cos 5

no pinning server sticky
pinning server pinning-failure link-down

no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66

bind interface port-channel1310
channel 1387
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet1501
description server 1/3, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1280
channel 1501
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1577
description server 1/3, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky

pinning server pinning-failure link-down
no cdp enable
switchport trunk allowed vlan 61-64

bind interface port-channel1280
channel 1577
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1580
description server 1/3, VNICeth2

switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down

no cdp enable
switchport trunk allowed vlan 61-64

bind interface port-channel1280
channel 1580
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1581
description server 1/3, VNICeth3
switchport mode trunk

untagged cos 5
no pinning server sticky
pinning server pinning-failure link-down

no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1280
channel 1581

service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet1584
description server 1/3, VNICeth4
switchport mode trunk
untagged cos 5

no pinning server sticky
pinning server pinning-failure link-down

no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1280
channel 1584

service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

interface Vethernet1803
  description server 2/2, VNICeth0
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 60
  switchport trunk allowed vlan 60
  bind interface port-channel1322
channel 1803
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

interface Vethernet1805
  description server 2/2, VNICeth1
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64

  bind interface port-channel1322
channel 1805
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

interface Vethernet1808
  description server 2/2, VNICeth2
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1322
channel 1808

```

```

  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet1809
  description server 2/2, VNICeth3
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down

  no cdp enable
  switchport trunk native vlan 65
  switchport trunk allowed vlan 65
  bind interface port-channel1322
channel 1809
  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

```

```

interface Vethernet1812
  description server 2/2, VNICeth4
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 66
  switchport trunk allowed vlan 66
  bind interface port-channel1322
channel 1812
  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

```

```

interface Vethernet1970
  description server 1/6, VNICeth0

  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 60
  switchport trunk allowed vlan 60

```

```
bind interface port-channel1282
channel 1970
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1972
description server 1/6, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1282
channel 1972
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1975
description server 1/6, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1282
channel 1975
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1987
description server 1/6, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
```

```
switchport trunk native vlan 66
switchport trunk allowed vlan 66
```

```
bind interface port-channel1282
channel 1987
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet1988
description server 2/5, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1329
channel 1988
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1990
description server 2/5, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
```

```
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
```

```
bind interface port-channel1329
channel 1990
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1993
description server 2/5, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
```

```

pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1329
channel 1993
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet1994
description server 2/5, VNICeth3
switchport mode trunk

```

```

untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1329
channel 1994
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet1997
description server 2/5, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1329
channel 1997
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet2000
description server 1/2, VNICeth0

```

```

switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1283
channel 2000
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2002
description server 1/2, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64

```

```

bind interface port-channel1283
channel 2002
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2005
description server 1/2, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1283
channel 2005
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2006

```

description server 1/2, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down

no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1283
channel 2006
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet2009
description server 1/2, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down

no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1283
channel 2009
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet2069
description server 2/6, VNICeth0

switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down

no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1304
channel 2069
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet2072
description server 2/6, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down

no cdp enable
switchport trunk allowed vlan 61
bind interface port-channel1304
channel 2072
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet2073
description server 2/6, VNICeth3
switchport mode trunk
untagged cos 4
no pinning server sticky
pinning server pinning-failure link-
down

no cdp enable
switchport trunk allowed vlan 62
bind interface port-channel1304
channel 2073
service-policy type queuing input
org-root/ep-qos-Gold
no shutdown

interface Vethernet2075
description server 2/6, VNICeth4
switchport mode trunk
untagged cos 1
no pinning server sticky
pinning server pinning-failure link-
down

no cdp enable
switchport trunk native vlan 63
switchport trunk allowed vlan 63

bind interface port-channel1304
channel 2075
service-policy type queuing input
org-root/ep-qos-Bronze

```

no shutdown

interface Vethernet2078
  description server 2/6, VNICeth5
  switchport mode trunk
  untagged cos 1
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 64
  switchport trunk allowed vlan 64
  bind interface port-channel1304
channel 2078
  service-policy type queuing input
org-root/ep-qos-Bronze
  no shutdown

interface Vethernet2079
  description server 2/6, VNICeth6
  switchport mode trunk
  untagged cos 5
  no pinning server sticky

  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 65
  switchport trunk allowed vlan 65
  bind interface port-channel1304
channel 2079
  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

interface Vethernet2083
  description server 2/6, VNICeth7
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 66
  switchport trunk allowed vlan 66
  bind interface port-channel1304
channel 2083

```

```

  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

interface Vethernet2094
  description server 1/4, VNICeth0
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 60
  switchport trunk allowed vlan 60
  bind interface port-channel1281
channel 2094
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

interface Vethernet2096
  description server 1/4, VNICeth1
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1281
channel 2096
  service-policy type queuing input
org-root/ep-qos-Silver

  no shutdown

interface Vethernet2099
  description server 1/4, VNICeth2
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64

```



```
bind interface port-channel1281
channel 2099
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2100
description server 1/4, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
```

```
switchport trunk allowed vlan 65
bind interface port-channel1281
channel 2100
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2103
description server 1/4, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1281
channel 2103
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2114
description server 2/1, VNICeth0
switchport mode trunk
untagged cos 2

no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
```

```
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1321
channel 2114
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2116
description server 2/1, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1321
channel 2116
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2119
description server 2/1, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1321
channel 2119
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2120
description server 2/1, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
```

```
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1321
channel 2120
service-policy type queuing input
org-root/ep-qos-Platinum

no shutdown
```

```
interface Vethernet2123
description server 2/1, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1321
channel 2123
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2124
description server 1/1, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
```

```
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1317
channel 2124
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2126
description server 1/1, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
```

```
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1317
channel 2126
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2129
description server 1/1, VNICeth2
switchport mode trunk
untagged cos 2

no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1317
channel 2129
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2130
description server 1/1, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1317
channel 2130
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2133

description server 1/1, VNICeth4
switchport mode trunk
```

```

untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1317
channel 2133
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet2158
description server 2/3, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1333
channel 2158

```

```

service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2160
description server 2/3, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1333
channel 2160
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2163
description server 2/3, VNICeth2

```

```

switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable

```

```

switchport trunk allowed vlan 61-
64
bind interface port-channel1333
channel 2163
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2164
description server 2/3, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1333
channel 2164
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet2167
description server 2/3, VNICeth4
switchport mode trunk
untagged cos 5

no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1333
channel 2167
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet2226
  description server 2/4, VNICeth0
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 60
  switchport trunk allowed vlan 60
  bind interface port-channel1339
channel 2226
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet2228
  description server 2/4, VNICeth1
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1339
channel 2228
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet2231
  description server 2/4, VNICeth2
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1339
channel 2231
  service-policy type queuing input
org-root/ep-qos-Silver

```

```

no shutdown

```

```

interface Vethernet2232
  description server 2/4, VNICeth3
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 65
  switchport trunk allowed vlan 65
  bind interface port-channel1339
channel 2232
  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

```

```

interface Vethernet2235
  description server 2/4, VNICeth4
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable

```

```

  switchport trunk native vlan 66
  switchport trunk allowed vlan 66
  bind interface port-channel1339
channel 2235
  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown
  clock timezone PST -8 0
  clock summer-time PDT 2 Sunday
March 03:00 1 Sunday November
02:00 60
  line console
  line vty
  system default switchport
shutdown
  ldap-server timeout 30
  ldap-server port 0
  aaa group server ldap ldap
svs veth auto-delete retention-
timer days 0 hours 0 mins 15

```

```

logging logfile messages 2
no logging monitor
logging level kernel 2
logging level user 2
logging level mail 2
logging level daemon 2
logging level auth 2
logging level syslog 2

logging level lpr 2
logging level news 2
logging level uucp 2
logging level cron 2
logging level authpri 2
logging level ftp 2
logging level local0 2
logging level local1 2
logging level local2 2
logging level local3 2
logging level local4 2
logging level local5 2
logging level local6 2
logging level local7 2
no logging console

```

13.2.1. 6248UP – A

```

!Command: show running-config
!Time: Mon Nov 18 10:35:07 2013

```

```

version 5.0(3)N2(2.11.3a)
feature fcoe
feature adapter-fex

no feature telnet
no telnet server enable
feature tacacs+
cfs ipv4 distribute
cfs eth distribute
feature private-vlan
feature port-security
feature lacp
feature lldp
feature fex
logging level assoc_mgr 2

logging level aaa 2
logging level afm 2

```

```

logging level cfs 2
logging level enm 2
logging level fex 2
logging level fwm 2
logging level msp 2
logging level npv 2
logging level pfm 2
logging level vms 2
logging level evmc 2
logging level port 2
logging level vshd 2
logging level ethpm 2
logging level track 2
logging level xmlma 2
logging level licmgr 2
logging level radius 2
logging level tacacs 2
logging level bootvar 2
logging level monitor 2
logging level fcdomain 2
logging level ascii-cfg 2
logging level provision 2

```

```

logging level securityd 2
logging level pltfm_config 2
logging level port-channel 2
logging level private-vlan 2
logging level spanning-tree 2
logging level port-resources 2
role name server-equipment
  rule 3 permit read-write feature
sam-pn-maintenance
  rule 2 permit read-write feature
sam-pn-policy
  rule 1 permit read-write feature
sam-pn-equipment
role name facility-manager
  rule 1 permit read-write feature
sam-power-mgmt
role name server-security
  rule 3 permit read-write feature
sam-ls-security-policy
  rule 2 permit read-write feature
sam-ls-security
  rule 1 permit read-write feature
sam-pn-security
role name server-compute

```

```

rule 3 permit read-write feature
sam-ls-compute
rule 2 permit read-write feature
sam-ls-server-oper
rule 1 permit read-write feature
sam-ls-server-policy
role name server-profile
rule 6 permit read-write feature
sam-ls-server-oper
rule 5 permit read-write feature
sam-ls-ext-access
rule 4 permit read-write feature
sam-ls-server-policy
rule 3 permit read-write feature
sam-ls-config-policy
rule 2 permit read-write feature
sam-ls-server
rule 1 permit read-write feature
sam-ls-config
role name operations
rule 2 permit read-write feature
sam-fault
rule 1 permit read-write feature
sam-operations
role name read-only
rule 1 permit read-write feature
sam-read-only
role name KVM-Only rule 1 permit
read-write feature sam-ls-ext-
access
role name network
rule 12 permit read-write feature
sam-ls-qos-policy
rule 11 permit read-write feature
sam-ls-network-policy
rule 10 permit read-write feature
sam-ls-qos
rule 9 permit read-write feature
sam-ls-network
rule 8 permit read-write feature
sam-ext-lan-qos
rule 7 permit read-write feature
sam-ext-lan-security
rule 6 permit read-write feature
sam-ext-lan-policy
rule 5 permit read-write feature
sam-ext-lan-config

```

```

rule 4 permit read-write feature
sam-pod-qos

```

```

rule 3 permit read-write feature
sam-pod-security
rule 2 permit read-write feature
sam-pod-policy
rule 1 permit read-write feature
sam-pod-config
role name storage
rule 6 permit read-write feature
sam-ls-storage-policy
rule 5 permit read-write feature
sam-ls-storage
rule 4 permit read-write feature
sam-ext-san-qos
rule 3 permit read-write feature
sam-ext-san-security
rule 2 permit read-write feature
sam-ext-san-policy
rule 1 permit read-write feature
sam-ext-san-config
role name admin
rule 1 permit read-write feature
sam-admin
role name aaa
rule 1 permit read-write feature
sam-aaa
no password strength-check

```

```

banner motd #Cisco UCS 6200
Series Fabric Interconnect
#

```

```

ip domain-lookup
ip name-server 171.70.168.183
171.68.226.120 175.25.205.7
aaa group server tacacs+ tacacs
switchname UCS-EXC-HyperV-B
logging event link-status default
errdisable recovery interval 60
errdisable recovery cause link-flap
errdisable recovery cause uddld
errdisable recovery cause
bpduguard
errdisable recovery cause loopback
errdisable recovery cause pause-
rate-limit

```

```

class-map type qos class-fcoe

class-map type qos match-all class-
gold
    match cos 4
class-map type qos match-all class-
bronze
    match cos 1
class-map type qos match-all class-
silver
    match cos 2
class-map type qos match-all class-
platinum
    match cos 5
class-map type queuing class-fcoe
    match qos-group 1
class-map type queuing class-gold
    match qos-group 3
class-map type queuing class-
bronze
    match qos-group 5
class-map type queuing class-silver
    match qos-group 4
class-map type queuing class-
platinum
    match qos-group 2
class-map type queuing class-all-
flood
    match qos-group 2
class-map type queuing class-ip-
multicast
    match qos-group 2

policy-map type qos
system_qos_policy
    class class-platinum
        set qos-group 2
    class class-silver
        set qos-group 4
    class class-bronze
        set qos-group 5
    class class-gold
        set qos-group 3
    class class-fcoe
        set qos-group 1
policy-map type queuing
system_q_in_policy
    class type queuing class-fcoe
        bandwidth percent 14
    class type queuing class-platinum
        bandwidth percent 22
    class type queuing class-gold
        bandwidth percent 20
    class type queuing class-silver
        bandwidth percent 18
    class type queuing class-bronze
        bandwidth percent 15

    class type queuing class-default
        bandwidth percent 11
policy-map type queuing
system_q_out_policy
    class type queuing class-fcoe
        bandwidth percent 14
    class type queuing class-platinum
        bandwidth percent 22
    class type queuing class-gold
        bandwidth percent 20
    class type queuing class-silver
        bandwidth percent 18
    class type queuing class-bronze
        bandwidth percent 15
    class type queuing class-default
        bandwidth percent 11
policy-map type queuing org-
root/ep-qos-Gold
    class type queuing class-default
        bandwidth percent 100
        shape 40000000 kbps 10240
policy-map type queuing org-
root/ep-qos-Bronze
    class type queuing class-default
        bandwidth percent 100

        shape 40000000 kbps 10240
policy-map type queuing org-
root/ep-qos-Silver
    class type queuing class-default
        bandwidth percent 100
        shape 40000000 kbps 10240
policy-map type queuing org-
root/ep-qos-Platinum
    class type queuing class-default
        bandwidth percent 100
        shape 40000000 kbps 10240

```

```

class-map type network-qos class-
fcoe
  match qos-group 1
class-map type network-qos class-
gold
  match qos-group 3
class-map type network-qos class-
bronze
  match qos-group 5
class-map type network-qos class-
silver
  match qos-group 4
class-map type network-qos class-
platinum
  match qos-group 2
class-map type network-qos class-
all-flood
  match qos-group 2
class-map type network-qos class-
ip-multicast

  match qos-group 2
policy-map type network-qos
system_nq_policy
  class type network-qos class-
platinum
    mtu 9000
  class type network-qos class-silver
    mtu 9000
  class type network-qos class-
bronze
    mtu 9000
  class type network-qos class-gold
    mtu 9000
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
  class type network-qos class-
default
system qos
  service-policy type qos input
system_qos_policy
  service-policy type queuing input
system_q_in_policy
  service-policy type queuing output
system_q_out_policy
  service-policy type network-qos
system_nq_policy

```

```

fex 1
  pinning max-links 1
  description "FEX0001"

fex 2
  pinning max-links 1
  description "FEX0002"
fex management-instance
39619c74-90b7-11e2-8831-
547feef80284 fabric 2
snmp-server enable traps entityfru
no snmp-server enable traps entity
entity_mib_change
no snmp-server enable traps entity
entity_module_status_change
no snmp-server enable traps entity
entity_power_status_change
no snmp-server enable traps entity
entity_module_inserted
no snmp-server enable traps entity
entity_module_removed
no snmp-server enable traps entity
entity_unrecognised_module
no snmp-server enable traps entity
entity_fan_status_change
no snmp-server enable traps link
linkDown
no snmp-server enable traps link
linkUp
no snmp-server enable traps link
extended-linkDown
no snmp-server enable traps link
extended-linkUp
no snmp-server enable traps link
cieLinkDown
no snmp-server enable traps link
cieLinkUp
no snmp-server enable traps link
connUnitPortStatusChange
no snmp-server enable traps link
fcTrunkIfUpNotify
no snmp-server enable traps link
fcTrunkIfDownNotify
no snmp-server enable traps link
delayed-link-state-change

no snmp-server enable traps link
fcot-inserted

```


no snmp-server enable traps link	state enabled
fcot-removed	
no snmp-server enable traps rf	logging level sysmgr 2
redundancy_framework	
no snmp-server enable traps license	
notify-license-expiry	interface port-channel28
no snmp-server enable traps license	description U: Uplink
notify-no-license-for-feature	switchport mode trunk
no snmp-server enable traps license	pinning border
notify-licensefile-missing	switchport trunk allowed vlan
no snmp-server enable traps license	1,60-66,999
notify-license-expiry-warning	speed 10000
no snmp-server enable traps rmon	
risingAlarm	
no snmp-server enable traps rmon	interface port-channel1290
fallingAlarm	switchport mode vntag
no snmp-server enable traps rmon	switchport vntag max-vifs 118
hcRisingAlarm	no pinning server sticky
no snmp-server enable traps rmon	speed 10000
hcFallingAlarm	
ntp server 171.68.10.80	interface port-channel1291
ntp server 171.68.10.150	switchport mode vntag
no aaa user default-role	switchport vntag max-vifs 118
	no pinning server sticky
	speed 10000
vrf context management	
ip name-server 171.70.168.183	
171.68.226.120 175.25.205.7	
vlan 1,60-66,999	interface port-channel1293
vlan 4044	switchport mode vntag
name SAM-vlan-management	switchport vntag max-vifs 118
vlan 4047	no pinning server sticky
name SAM-vlan-boot	speed 10000
vlan 4048	interface port-channel1294
fcoe vsan 1	switchport mode vntag
name fcoe-vsan-4048	switchport vntag max-vifs 118
vlan 4049	
port-channel load-balance ethernet	no pinning server sticky
invalid invalid-hash	speed 10000
port-profile type vethernet	
ucsm_internal_rackserver_portprofile	interface port-channel1295
le	switchport mode vntag
switchport trunk allowed vlan	no pinning server sticky
4044	speed 10000
switchport mode trunk	
no shutdown	interface port-channel1298
vmware port-group	switchport mode vntag
max-ports 320	switchport vntag max-vifs 118

no pinning server sticky speed 10000	switchport vntag max-vifs 118 no pinning server sticky speed 10000
interface port-channel1302 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000	interface port-channel1343 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000
interface port-channel1316 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000	interface port-channel1345 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000 feature npv feature npiv
interface port-channel1320 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000	interface Ethernet1/1 description S: Server no pinning server sticky switchport mode fex-fabric fex associate 1 chassis-serial FOX1642H32G module-serial FCH16457K1Z module-sl ot right no shutdown
interface port-channel1323 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000	interface Ethernet1/2 description S: Server no pinning server sticky switchport mode fex-fabric fex associate 1 chassis-serial FOX1642H32G module-serial FCH16457K1Z module-sl ot right no shutdown
interface port-channel1328 switchport mode vntag no pinning server sticky speed 10000	interface Ethernet1/3 description S: Server no pinning server sticky switchport mode fex-fabric
interface port-channel1332 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000	
interface port-channel1338 switchport mode vntag switchport vntag max-vifs 118 no pinning server sticky speed 10000	
interface port-channel1340 switchport mode vntag	

```
fex associate 1 chassis-serial
FOX1642H32G module-serial
FCH16457K1Z module-sl
ot right
no shutdown
```

```
interface Ethernet1/4
description S: Server
no pinning server sticky
switchport mode fex-fabric
fex associate 1 chassis-serial
FOX1642H32G module-serial
FCH16457K1Z module-sl
ot right
no shutdown
```

```
interface Ethernet1/5
description S: Server

no pinning server sticky
switchport mode fex-fabric
fex associate 2 chassis-serial
FOX1652G334 module-serial
FCH1650JRQK module-sl
ot right
no shutdown
```

```
interface Ethernet1/6
description S: Server
no pinning server sticky
switchport mode fex-fabric
fex associate 2 chassis-serial
FOX1652G334 module-serial
FCH1650JRQK module-sl
ot right
no shutdown
```

```
interface Ethernet1/7
description S: Server
no pinning server sticky
switchport mode fex-fabric
fex associate 2 chassis-serial
FOX1652G334 module-serial
FCH1650JRQK module-sl
ot right
no shutdown
```

```
interface Ethernet1/8
description S: Server
no pinning server sticky
switchport mode fex-fabric
fex associate 2 chassis-serial
FOX1652G334 module-serial
FCH1650JRQK module-sl
ot right
no shutdown
```

```
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
description A: Appliance
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport mode trunk
switchport trunk allowed vlan 65-
66
no shutdown
```

```
interface Ethernet1/18
description A: Appliance
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport mode trunk
```

```
switchport trunk allowed vlan 65-66
no shutdown
```

```
interface Ethernet1/19
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
no shutdown
```

```
interface Ethernet1/20
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
no shutdown
```

```
interface Ethernet1/21
```

```
interface Ethernet1/22
```

```
interface Ethernet1/23
```

```
interface Ethernet1/24
```

```
interface Ethernet1/25
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
channel-group 28 mode active
no shutdown
```

```
interface Ethernet1/26
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
channel-group 28 mode active
no shutdown
```

```
interface Ethernet1/27
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
channel-group 28 mode active

no shutdown
```

```
interface Ethernet1/28
description U: Uplink
pinning border
switchport mode trunk
switchport trunk allowed vlan
1,60-66,999
channel-group 28 mode active
no shutdown
```

```
interface Ethernet1/29
```

```
interface Ethernet1/30
```

```
interface Ethernet1/31
```

```
interface Ethernet1/32
```

```
interface Ethernet2/1
```

```
interface Ethernet2/2
```

```
interface Ethernet2/3
```

```
interface Ethernet2/4
```

```
interface Ethernet2/5
```

```
interface Ethernet2/6
```

```
interface Ethernet2/7
```

```
interface Ethernet2/8
```

```
interface Ethernet2/9
```

```
interface Ethernet2/10
```

```

interface Ethernet2/11

interface Ethernet2/12

interface Ethernet2/13

interface Ethernet2/14

interface Ethernet2/15

interface Ethernet2/16

interface mgmt0
shutdown force
ip address 10.60.0.12/24

interface Ethernet1/1/1
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/1
channel-group 1316
no shutdown

interface Ethernet1/1/2
no pinning server sticky

interface Ethernet1/1/3

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/1
channel-group 1316
no shutdown

interface Ethernet1/1/4
no pinning server sticky

interface Ethernet1/1/5
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/2
channel-group 1293
no shutdown

```

```

interface Ethernet1/1/6
no pinning server sticky

interface Ethernet1/1/7

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/2
channel-group 1293
no shutdown

interface Ethernet1/1/8
no pinning server sticky

interface Ethernet1/1/9
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/3
channel-group 1294
no shutdown

interface Ethernet1/1/10
no pinning server sticky

interface Ethernet1/1/11

switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/3
channel-group 1294
no shutdown

interface Ethernet1/1/12
no pinning server sticky

interface Ethernet1/1/13
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/4
channel-group 1290
no shutdown

interface Ethernet1/1/14
no pinning server sticky

```

```

interface Ethernet1/1/15
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/4
    channel-group 1290
    no shutdown

interface Ethernet1/1/16
    no pinning server sticky

interface Ethernet1/1/17
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/1
    channel-group 1340
    no shutdown

interface Ethernet1/1/18
    no pinning server sticky

interface Ethernet1/1/19
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/1
    channel-group 1340
    no shutdown

interface Ethernet1/1/20
    no pinning server sticky

interface Ethernet1/1/21
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/2
    channel-group 1295
    no shutdown

interface Ethernet1/1/22
    no pinning server sticky

interface Ethernet1/1/23
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/2
    channel-group 1295
    no shutdown

interface Ethernet1/1/24
    no pinning server sticky

interface Ethernet1/1/25
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/3
    channel-group 1345
    no shutdown

interface Ethernet1/1/26
    no pinning server sticky

interface Ethernet1/1/27
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/3
    channel-group 1345
    no shutdown

interface Ethernet1/1/28
    no pinning server sticky

interface Ethernet1/1/29
    switchport vntag max-vifs 118
    no pinning server sticky
    switchport mode vntag
    fabric-interface Eth1/4
    channel-group 1291
    no shutdown

interface Ethernet1/1/30
    no pinning server sticky

interface Ethernet1/1/31
    switchport vntag max-vifs 118

```

```

no pinning server sticky
switchport mode vntag
fabric-interface Eth1/4
channel-group 1291
no shutdown

interface Ethernet1/1/32
no pinning server sticky

interface Ethernet1/1/33
no pinning server sticky
switchport mode trunk
switchport trunk native vlan 4044
switchport trunk allowed vlan
4044
no shutdown

interface Ethernet2/1/1
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/5

channel-group 1320
no shutdown

interface Ethernet2/1/2
no pinning server sticky

interface Ethernet2/1/3
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/5
channel-group 1320
no shutdown

interface Ethernet2/1/4
no pinning server sticky

interface Ethernet2/1/5
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/6

channel-group 1323
no shutdown

```

```

interface Ethernet2/1/6
no pinning server sticky

interface Ethernet2/1/7
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/6
channel-group 1323
no shutdown

interface Ethernet2/1/8
no pinning server sticky

interface Ethernet2/1/9
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/7

channel-group 1332
no shutdown

interface Ethernet2/1/10
no pinning server sticky

interface Ethernet2/1/11
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/7
channel-group 1332
no shutdown

interface Ethernet2/1/12
no pinning server sticky

interface Ethernet2/1/13
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/8

channel-group 1338
no shutdown

interface Ethernet2/1/14

```

```

no pinning server sticky

interface Ethernet2/1/15
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/8
channel-group 1338
no shutdown

interface Ethernet2/1/16
no pinning server sticky

interface Ethernet2/1/17
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/5

channel-group 1328
no shutdown

interface Ethernet2/1/18
no pinning server sticky

interface Ethernet2/1/19
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/5
channel-group 1328
no shutdown

interface Ethernet2/1/20
no pinning server sticky

interface Ethernet2/1/21
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/6

channel-group 1302
no shutdown

interface Ethernet2/1/22
no pinning server sticky

interface Ethernet2/1/23
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/6
channel-group 1302
no shutdown

interface Ethernet2/1/24
no pinning server sticky

interface Ethernet2/1/25
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/7

channel-group 1343
no shutdown

interface Ethernet2/1/26
no pinning server sticky

interface Ethernet2/1/27
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/7
channel-group 1343
no shutdown

interface Ethernet2/1/28
no pinning server sticky

interface Ethernet2/1/29
switchport vntag max-vifs 118
no pinning server sticky
switchport mode vntag
fabric-interface Eth1/8

channel-group 1298
no shutdown

interface Ethernet2/1/30
no pinning server sticky

interface Ethernet2/1/31
switchport vntag max-vifs 118

```



```

no pinning server sticky
switchport mode vntag
fabric-interface Eth1/8
channel-group 1298
no shutdown

interface Ethernet2/1/32
no pinning server sticky

interface Ethernet2/1/33
no pinning server sticky
switchport mode trunk
switchport trunk native vlan 4044
switchport trunk allowed vlan
4044

no shutdown

interface Vethernet1361
description server 1/8, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1291
channel 1361
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1362
description server 1/8, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable

switchport trunk allowed vlan 61
bind interface port-channel1291
channel 1362
service-policy type queuing input
org-root/ep-qos-Silver

no shutdown

interface Vethernet1365
description server 1/8, VNICeth3
switchport mode trunk
untagged cos 4
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 62
bind interface port-channel1291
channel 1365
service-policy type queuing input
org-root/ep-qos-Gold
no shutdown

interface Vethernet1367
description server 1/8, VNICeth4
switchport mode trunk
untagged cos 1
no pinning server sticky

pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 63
switchport trunk allowed vlan 63
bind interface port-channel1291
channel 1367
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown

interface Vethernet1368
description server 1/8, VNICeth5
switchport mode trunk
untagged cos 1
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 64
switchport trunk allowed vlan 64
bind interface port-channel1291
channel 1368
service-policy type queuing input
org-root/ep-qos-Bronze

```

```

no shutdown

interface Vethernet1371

    description server 1/8, VNICeth6
    switchport mode trunk
    untagged cos 5
    no pinning server sticky
    pinning server pinning-failure link-
down
    no cdp enable
    switchport trunk native vlan 65
    switchport trunk allowed vlan 65
    bind interface port-channel1291
channel 1371
    service-policy type queuing input
org-root/ep-qos-Platinum
    no shutdown

interface Vethernet1372

    description server 1/8, VNICeth7
    switchport mode trunk
    untagged cos 5
    no pinning server sticky
    pinning server pinning-failure link-
down
    no cdp enable
    switchport trunk native vlan 66
    switchport trunk allowed vlan 66
    bind interface port-channel1291
channel 1372

    service-policy type queuing input
org-root/ep-qos-Platinum
    no shutdown

interface Vethernet1375

    description server 2/8, VNICeth0
    switchport mode trunk
    untagged cos 2
    no pinning server sticky
    pinning server pinning-failure link-
down
    no cdp enable
    switchport trunk native vlan 60
    switchport trunk allowed vlan 60
    bind interface port-channel1298
channel 1375

```

```

    service-policy type queuing input
org-root/ep-qos-Silver
    no shutdown

interface Vethernet1376

    description server 2/8, VNICeth1
    switchport mode trunk
    untagged cos 2
    no pinning server sticky
    pinning server pinning-failure link-
down

    no cdp enable
    switchport trunk allowed vlan 61
    bind interface port-channel1298
channel 1376
    service-policy type queuing input
org-root/ep-qos-Silver
    no shutdown

interface Vethernet1379

    description server 2/8, VNICeth3
    switchport mode trunk
    untagged cos 4
    no pinning server sticky
    pinning server pinning-failure link-
down
    no cdp enable
    switchport trunk allowed vlan 62
    bind interface port-channel1298
channel 1379
    service-policy type queuing input
org-root/ep-qos-Gold
    no shutdown

interface Vethernet1381

    description server 2/8, VNICeth4
    switchport mode trunk
    untagged cos 1

    no pinning server sticky
    pinning server pinning-failure link-
down
    no cdp enable
    switchport trunk native vlan 63
    switchport trunk allowed vlan 63
    bind interface port-channel1298
channel 1381

```

```
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown
```

```
interface Vethernet1382
description server 2/8, VNICeth5
switchport mode trunk
untagged cos 1
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 64
switchport trunk allowed vlan 64
bind interface port-channel1298
channel 1382
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown
```

```
interface Vethernet1385
description server 2/8, VNICeth6
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1298
channel 1385
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet1386
description server 2/8, VNICeth7
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
```

```
bind interface port-channel1298
channel 1386
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet1502
description server 1/3, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1294
channel 1502
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1578
description server 1/3, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
```

```
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1294
channel 1578
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1579
description server 1/3, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
```

```
switchport trunk allowed vlan 61-64
bind interface port-channel1294
channel 1579
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1582
description server 1/3, VNICeth3
switchport mode trunk

untagged cos 5
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1294
channel 1582
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet1583
description server 1/3, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1294
channel 1583
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet1804
description server 2/2, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down
```

```
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1323
channel 1804
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1806
description server 2/2, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk allowed vlan 61-64
```

```
bind interface port-channel1323
channel 1806
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1807
description server 2/2, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk allowed vlan 61-64
```

```
bind interface port-channel1323
channel 1807
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet1810
description server 2/2, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
```

pinning server pinning-failure link-down

no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1323
channel 1810
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet1811
description server 2/2, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1323
channel 1811
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet1971
description server 1/6, VNICeth0

switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1295
channel 1971
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1973
description server 1/6, VNICeth1
switchport mode trunk

untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk allowed vlan 61-64
bind interface port-channel1295
channel 1973
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1974
description server 1/6, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk allowed vlan 61-64
bind interface port-channel1295
channel 1974
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet1985
description server 1/6, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-down
no cdp enable
switchport trunk native vlan 999
switchport trunk allowed vlan 999

bind interface port-channel1295
channel 1985
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

interface Vethernet1986
  description server 1/6, VNICeth4
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 66
  switchport trunk allowed vlan 66
  bind interface port-channel1295
channel 1986
  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

```

```

interface Vethernet1989
  description server 2/5, VNICeth0
  switchport mode trunk
  untagged cos 2
  no pinning server sticky

  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 60
  switchport trunk allowed vlan 60
  bind interface port-channel1328
channel 1989
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet1991
  description server 2/5, VNICeth1
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1328
channel 1991
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet1992
  description server 2/5, VNICeth2

  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1328
channel 1992
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet1995
  description server 2/5, VNICeth3
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 65
  switchport trunk allowed vlan 65
  bind interface port-channel1328
channel 1995
  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

```

```

interface Vethernet1996
  description server 2/5, VNICeth4
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 66
  switchport trunk allowed vlan 66
  bind interface port-channel1328
channel 1996

```

```
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2001
description server 1/2, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
```

```
switchport trunk allowed vlan 60
bind interface port-channel1293
channel 2001
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2003
description server 1/2, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1293
channel 2003
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2004
description server 1/2, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky

pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
```

```
bind interface port-channel1293
channel 2004
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2007
description server 1/2, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1293
channel 2007
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2008
description server 1/2, VNICeth4

switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1293
channel 2008
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2070
description server 2/6, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
```

```

switchport trunk allowed vlan 60
bind interface port-channel1302
channel 2070
service-policy type queuing input
org-root/ep-qos-Silver

no shutdown

interface Vethernet2071
description server 2/6, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61
bind interface port-channel1302
channel 2071
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

interface Vethernet2074
description server 2/6, VNICeth3
switchport mode trunk
untagged cos 4
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 62

bind interface port-channel1302
channel 2074
service-policy type queuing input
org-root/ep-qos-Gold
no shutdown

interface Vethernet2076
description server 2/6, VNICeth4
switchport mode trunk
untagged cos 1
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 63

```

```

switchport trunk allowed vlan 63
bind interface port-channel1302
channel 2076
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown

interface Vethernet2077
description server 2/6, VNICeth5
switchport mode trunk
untagged cos 1
no pinning server sticky

pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 64
switchport trunk allowed vlan 64
bind interface port-channel1302
channel 2077
service-policy type queuing input
org-root/ep-qos-Bronze
no shutdown

interface Vethernet2080
description server 2/6, VNICeth6
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1302
channel 2080
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

interface Vethernet2082
description server 2/6, VNICeth7
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down

```



```
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1302
channel 2082
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2095
description server 1/4, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1290
channel 2095
```

```
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2097
description server 1/4, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1290
channel 2097
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2098
description server 1/4, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
```

```
pinning server pinning-failure link-
down
no cdp enable
```

```
switchport trunk allowed vlan 61-
64
bind interface port-channel1290
channel 2098
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2101
description server 1/4, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1290
channel 2101
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2102
description server 1/4, VNICeth4
switchport mode trunk
untagged cos 5

no pinning server sticky
pinning server pinning-failure link-
down
```

```
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1290
channel 2102
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2115
description server 2/1, VNICeth0
switchport mode trunk
```

```

untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1320
channel 2115
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2117
description server 2/1, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1320
channel 2117
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2118
description server 2/1, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1320
channel 2118
service-policy type queuing input
org-root/ep-qos-Silver

no shutdown

```

```

interface Vethernet2121

```

```

description server 2/1, VNICeth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1320
channel 2121
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet2122
description server 2/1, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable

```

```

switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1320
channel 2122
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet2125
description server 1/1, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1316
channel 2125
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2127
  description server 1/1, VNICeth1
  switchport mode trunk

  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1316
channel 2127
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet2128
  description server 1/1, VNICeth2
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1316
channel 2128
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet2131
  description server 1/1, VNICeth3
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 65
  switchport trunk allowed vlan 65
  bind interface port-channel1316
channel 2131
  service-policy type queuing input
org-root/ep-qos-Platinum

```

```

no shutdown

```

```

interface Vethernet2132
  description server 1/1, VNICeth4
  switchport mode trunk
  untagged cos 5
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 66
  switchport trunk allowed vlan 66
  bind interface port-channel1316
channel 2132

```

```

  service-policy type queuing input
org-root/ep-qos-Platinum
  no shutdown

```

```

interface Vethernet2159
  description server 2/3, VNICeth0
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down
  no cdp enable
  switchport trunk native vlan 60
  switchport trunk allowed vlan 60
  bind interface port-channel1332
channel 2159
  service-policy type queuing input
org-root/ep-qos-Silver
  no shutdown

```

```

interface Vethernet2161
  description server 2/3, VNICeth1
  switchport mode trunk
  untagged cos 2
  no pinning server sticky
  pinning server pinning-failure link-
down

  no cdp enable
  switchport trunk allowed vlan 61-
64
  bind interface port-channel1332
channel 2161

```

```
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2162
description server 2/3, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1332
channel 2162
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2165
description server 2/3, VNICeth3
switchport mode trunk
untagged cos 5

no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1332
channel 2165
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2166
description server 2/3, VNICeth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
```

```
bind interface port-channel1332
channel 2166
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown
```

```
interface Vethernet2227
description server 2/4, VNICeth0
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 60
switchport trunk allowed vlan 60
bind interface port-channel1338
channel 2227
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2229
description server 2/4, VNICeth1
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk allowed vlan 61-
64
bind interface port-channel1338
channel 2229
```

```
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown
```

```
interface Vethernet2230
description server 2/4, VNICeth2
switchport mode trunk
untagged cos 2
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
```

```

switchport trunk allowed vlan 61-
64
bind interface port-channel1338
channel 2230
service-policy type queuing input
org-root/ep-qos-Silver
no shutdown

```

```

interface Vethernet2233
description server 2/4, VNIC eth3
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable

```

```

switchport trunk native vlan 65
switchport trunk allowed vlan 65
bind interface port-channel1338
channel 2233
service-policy type queuing input
org-root/ep-qos-Platinum
no shutdown

```

```

interface Vethernet2234
description server 2/4, VNIC eth4
switchport mode trunk
untagged cos 5
no pinning server sticky
pinning server pinning-failure link-
down
no cdp enable
switchport trunk native vlan 66
switchport trunk allowed vlan 66
bind interface port-channel1338
channel 2234
service-policy type queuing input
org-root/ep-qos-Platinum

```

```

no shutdown
clock timezone PST -8 0
clock summer-time PDT 2 Sunday
March 03:00 1 Sunday November
02:00 60
line console
line vty

```

```

system default switchport
shutdown
ldap-server timeout 30
ldap-server port 0
aaa group server ldap ldap
svs veth auto-delete retention-
timer days 0 hours 0 mins 15
logging logfile messages 2
no logging monitor
logging level kernel 2
logging level user 2
logging level mail 2
logging level daemon 2
logging level auth 2
logging level syslog 2
logging level lpr 2
logging level news 2
logging level uucp 2
logging level cron 2
logging level authpri 2
logging level ftp 2
logging level local0 2
logging level local1 2
logging level local2 2

```

```

logging level local3 2
logging level local4 2
logging level local5 2
logging level local6 2
logging level local7 2
no logging console

```

13.3. Sample Cisco Nexus 1000V Configuration

```

!Command: show running-config
!Time: Tue Nov 19 19:57:29 2013

```

```

version 5.2(1)SM1(5.1)
hostname nexus1000v

```

```

no feature telnet

```

feature network-segmentation-manager

no password strength-check

banner motd #Nexus 1000V Switch

#

ip domain-lookup

errdisable recovery cause failed-port-state

vem 3

host id 749C6139-B790-E211-0000-00000000003E

vem 4

host id 749C6139-B790-E211-0000-00000000000F

vem 5

host id 749C6139-B790-E211-0000-000000000006E

vem 6

host id 749C6139-B790-E211-0000-000000000002F

vem 7

host id 749C6139-B790-E211-0000-000000000004E

vem 8

host id 749C6139-B790-E211-0000-000000000002E

vem 9

host id 749C6139-B790-E211-0000-000000000005E

vem 10

host id 749C6139-B790-E211-0000-000000000007E

vem 11

host id 749C6139-B790-E211-0000-000000000001F

vem 12

host id 749C6139-B790-E211-0000-000000000003F

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 10.61.0.1

vlan 1,61-64

port-channel load-balance ethernet source-mac

port-profile default max-ports 32

port-profile default port-binding static

port-profile type vethernet NSM_template_vlan

no shutdown

guid 2d26a674-2453-499f-8cd5-609d967ad7ac

```

    description NSM default port-profile for VLAN networks. Do not delete.
    state enabled
port-profile type vethernet NSM_template_segmentation
    no shutdown
    guid 4bc0ff48-39b2-414b-b90c-830e5ad65896
    description NSM default port-profile for VXLAN networks. Do not delete.
    state enabled
port-profile type vethernet VDI-Port-profile
    no shutdown
    guid 43f99166-d36a-4962-b4ef-8d2b30626201
    port-binding static auto expand
    max-ports 1024
    state enabled

publish port-profile
port-profile type ethernet Uplink-profile
    channel-group auto mode on mac-pinning
    no shutdown
    guid 62a0310b-e9b0-40d2-84da-6d5ae19d158c
    max-ports 512
    state enabled
port-profile type ethernet uplink_network_default_policy
    no shutdown
    guid af5708a0-978b-423d-9f2c-da3394c37f29
    max-ports 512
    description NSM created profile. Do not delete.
    state enabled
port-profile type ethernet nexus1000v-uplink
    inherit port-profile Uplink-profile
    switchport mode trunk
    switchport trunk allowed vlan 61-64
    guid 6d728913-d267-4ca3-ab74-e304a88edeff
    max-ports 512
    description NSM created profile. Do not delete.
    state enabled
port-profile type vethernet dynpp_43f99166-d36a-4962-b4ef-8d2b30626201_c5fcbaa8-
e3fb-4f41-ae33-6bcabfe9b742

inherit port-profile VDI-Port-profile
switchport mode access
switchport access vlan 62
guid 0d40e8de-9cdc-4424-ac18-5e9f42ca5784
description NSM created profile. Do not delete.
state enabled

interface port-channel1
    inherit port-profile nexus1000v-uplink

```

vem 3

```
interface port-channel2
  inherit port-profile nexus1000v-uplink
vem 4
```

```
interface port-channel3
  inherit port-profile nexus1000v-uplink
vem 5
```

```
interface port-channel4
  inherit port-profile nexus1000v-uplink
vem 6
```

```
interface port-channel5
  inherit port-profile nexus1000v-uplink
vem 7
```

```
interface port-channel6
  inherit port-profile nexus1000v-uplink
vem 8
```

```
interface port-channel7
  inherit port-profile nexus1000v-uplink
vem 9
```

```
interface port-channel8
  inherit port-profile nexus1000v-uplink
vem 10
```

```
interface port-channel9
  inherit port-profile nexus1000v-uplink
vem 11
```

```
interface port-channel10
  inherit port-profile nexus1000v-uplink
```

vem 12

```
interface mgmt0
  ip address 10.61.0.10/24
```

```
interface Ethernet4/1
  inherit port-profile nexus1000v-uplink
```

```
interface Ethernet5/1
```



```

inherit port-profile nexus1000v-uplink

interface Ethernet7/1
  inherit port-profile nexus1000v-uplink

interface control0
  no snmp trap link-status
  line console
  line vty
  boot kickstart bootflash:/n1000vh-dk9-kickstart.5.2.1.SM1.5.1.bin sup-1
  boot system bootflash:/n1000vh-dk9.5.2.1.SM1.5.1.bin sup-1
  boot kickstart bootflash:/n1000vh-dk9-kickstart.5.2.1.SM1.5.1.bin sup-2
  boot system bootflash:/n1000vh-dk9.5.2.1.SM1.5.1.bin sup-2
  svcs-domain

  domain id 100
  control vlan 1
  packet vlan 1
  svcs mode L3 interface mgmt0
  switch-guid cee527ce-ff22-413b-82d3-c15e2df44a86
  vservice global type vsg
  tcp state-checks
  vnm-policy-agent
  registration-ip 0.0.0.0
  shared-secret *****
  log-level info
  nsm ip pool template VLAN-61-Pool
  ip address 10.61.0.11 10.61.0.11
  network 10.61.0.0 255.255.255.0
  default-router 10.61.0.1
  nsm ip pool template VLAN-62-Pool
  ip address 10.62.0.100 10.62.0.150
  network 10.62.0.0 255.255.240.0
  default-router 10.62.0.1
  nsm ip pool template VLAN-63-Pool
  ip address 10.63.0.100 10.63.0.150
  network 10.63.0.0 255.255.255.0
  default-router 10.63.0.1

  nsm ip pool template VLAN-64-Pool
  ip address 10.64.0.100 10.64.0.150
  network 10.64.0.0 255.255.255.0
  default-router 10.64.0.1
  nsm logical network VDINetwork
  nsm logical network ClusterNetwork
  nsm network segment pool VDI-Pool-1
  guid 5ef2ba1d-212f-4067-b6b5-54210635541b
  member-of logical network VDINetwork

```

```

nsm network segment pool Cluster-Pool-1
  guid 20c1b5da-2a4b-4a53-9926-af7dbe1fbba0
  member-of logical network ClusterNetwork
nsm network segment VLAN-61
  guid 840e731b-0337-4f2e-9e0a-fe4fce2182d2
  member-of vmnetwork VLAN-61 guid 84a7093f-4c54-4840-a71c-cb8599359e48
  member-of network segment pool VDI-Pool-1
  switchport access vlan 61
  ip pool import template VLAN-61-Pool guid 424651dd-c4ae-4a30-baee-73686b0dd867
  publish network segment
  switchport mode access
nsm network segment VLAN-62
  guid c5fcbaa8-e3fb-4f41-ae33-6bcabfe9b742
  member-of vmnetwork VLAN-62 guid 2fcdca69-65e3-43e8-9d1a-21442c69a325

  member-of network segment pool VDI-Pool-1
  switchport access vlan 62
  ip pool import template VLAN-62-Pool guid 433450e9-0bc9-4cde-a84f-51f30fa8ccd9
  publish network segment
  switchport mode access
nsm network segment VLAN-63
  guid afee3cc3-8d51-4c41-9d8d-a910c8a6f74f
  member-of vmnetwork VLAN-63 guid 6c0c92c3-17a1-4067-bfd9-70d2e9ea58c4
  member-of network segment pool Cluster-Pool-1
  switchport access vlan 63
  ip pool import template VLAN-63-Pool guid ca9dab93-f4be-4fa4-b7af-cce56473858a
  publish network segment
  switchport mode access
nsm network segment VLAN-64
  guid 1352eb9b-770f-405a-b055-a5874d661175
  member-of vmnetwork VLAN-64 guid 7c4816f4-c0c6-4043-a063-fa199d397683
  member-of network segment pool Cluster-Pool-1
  switchport access vlan 64
  ip pool import template VLAN-64-Pool guid 20c3fe31-bf46-458b-9e87-d01f08e77032
  publish network segment
  switchport mode access
nsm network uplink nexus1000v-uplink
  import port-profile Uplink-profile

  allow network segment pool VDI-Pool-1
  allow network segment pool Cluster-Pool-1
  publish network uplink

```

13.3.1. N1000V-1

```

!Command: show running-config
!Time: Thu Nov 21 12:44:44 2013

```

```

version 5.2(1)SM1(5.1)

```

```

hostname nexus1000v-2

no feature telnet
feature network-segmentation-manager

no password strength-check

banner motd #Nexus 1000V Switch
#

ip domain-lookup
errdisable recovery cause failed-port-state
vem 3
    host id 749C6139-B790-E211-0000-00000000003E
vem 9
    host id 749C6139-B790-E211-0000-000000000005E
vem 10

    host id 749C6139-B790-E211-0000-000000000007E
vem 11
    host id 749C6139-B790-E211-0000-000000000001F
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 10.61.0.1
vlan 1,61-64

port-channel load-balance ethernet source-mac
port-profile default max-ports 32
port-profile default port-binding static
port-profile type vethernet NSM_template_vlan
    no shutdown
    guid e92a97d4-1ae1-4d43-a244-f0e5b2bf007d
    description NSM default port-profile for VLAN networks. Do not delete.
    state enabled

port-profile type vethernet NSM_template_segmentation
    no shutdown
    guid 1e06ea81-ccc1-4b83-8f14-87acb840cdd0
    description NSM default port-profile for VXLAN networks. Do not delete.
    state enabled
port-profile type vethernet VDI-Port-profile2
    no shutdown
    guid 12fdb4e1-b3ed-4dbd-81f2-7b4be0b64bbc

```

```

state enabled
publish port-profile
port-profile type ethernet Uplink-profile2
channel-group auto mode on mac-pinning
no shutdown
guid 28fbc2f1-d757-4235-bc60-e64ec2711df1
max-ports 512
state enabled
port-profile type ethernet uplink_network_default_policy
no shutdown
guid e3949f10-9713-4a1f-a267-b818490323c3
max-ports 512
description NSM created profile. Do not delete.
state enabled
port-profile type ethernet nexus1000v-uplink2

inherit port-profile Uplink-profile2
switchport mode trunk
switchport trunk allowed vlan 61-64
guid b8d35b70-81b4-4f5a-a15d-47080f85c92f
max-ports 512
description NSM created profile. Do not delete.
state enabled

interface port-channel1
inherit port-profile nexus1000v-uplink2
vem 3

interface port-channel2
inherit port-profile nexus1000v-uplink2
vem 10

interface port-channel3
inherit port-profile nexus1000v-uplink2
vem 9

interface port-channel4
inherit port-profile nexus1000v-uplink2

vem 11

interface mgmt0
ip address 10.61.0.12/24

interface Ethernet3/2
inherit port-profile nexus1000v-uplink2

```

```

interface Ethernet9/1
    inherit port-profile nexus1000v-uplink2

interface Ethernet10/1
    inherit port-profile nexus1000v-uplink2

interface Ethernet11/1
    inherit port-profile nexus1000v-uplink2

interface control0
    no snmp trap link-status
line console
line vty
boot kickstart bootflash:/n1000vh-dk9-kickstart.5.2.1.SM1.5.1.bin sup-1
boot system bootflash:/n1000vh-dk9.5.2.1.SM1.5.1.bin sup-1

boot kickstart bootflash:/n1000vh-dk9-kickstart.5.2.1.SM1.5.1.bin sup-2
boot system bootflash:/n1000vh-dk9.5.2.1.SM1.5.1.bin sup-2
svs-domain
    domain id 200
    control vlan 1
    packet vlan 1
    svcs mode L3 interface mgmt0
    switch-guid 862960f3-835d-48f1-be0c-dec4839d2de4
vservice global type vsg
    tcp state-checks
vnm-policy-agent
    registration-ip 0.0.0.0
    shared-secret *****
    log-level info
nsm ip pool template VLAN-61-Pool2
    ip address 10.61.0.13 10.61.0.13
    network 10.61.0.0 255.255.255.0
    default-router 10.61.0.1
nsm ip pool template VLAN-62-Pool2
    ip address 10.62.0.151 10.62.0.200
    network 10.62.0.0 255.255.240.0
    default-router 10.62.0.1
nsm ip pool template VLAN-63-Pool2

    ip address 10.63.0.151 10.63.0.200
    network 10.63.0.0 255.255.255.0
    default-router 10.63.0.1
nsm ip pool template VLAN-64-Pool2
    ip address 10.64.0.151 10.64.0.200
    network 10.64.0.0 255.255.255.0
    default-router 10.64.0.1
nsm logical network VDINetwork2

```

```

nsm logical network ClusterNetwork2
nsm network segment pool VDI-Pool-2
  guid d6beab2d-5a35-432f-8487-852d4bb9bb5c
  member-of logical network VDINetwork2
nsm network segment pool Cluster-Pool-2
  guid 5f7966c6-6fb1-449b-bcc6-5a5b1a0b34ab
  member-of logical network ClusterNetwork2
nsm network segment VLAN-61-2
  guid f211722b-f1b4-4c33-bab3-22fd7edec5ae
  member-of vmnetwork VLAN-61-2 guid 80a41c1c-c7aa-40eb-8696-ba4d0895885e
  member-of network segment pool VDI-Pool-2
  switchport access vlan 61
  ip pool import template VLAN-61-Pool2 guid d3f4d443-ac0d-4861-8d53-57103c70e76
9
  publish network segment

```

```

  switchport mode access
nsm network segment VLAN-62-2
  guid 39f86003-8e46-4521-88d5-f712eb4670c2
  member-of vmnetwork VLAN-62-2 guid 78a66b7d-e7c7-4dfc-b305-5cd891aa4a40
  member-of network segment pool VDI-Pool-2
  switchport access vlan 62
  ip pool import template VLAN-62-Pool2 guid 09b71666-d870-4619-8dc1-778f0ba84cb
a
  publish network segment
  switchport mode access

```

```

nsm network segment VLAN-63-2
  guid 5bf535bd-0ab0-4511-a0e5-cf0290fbbd0b
  member-of vmnetwork VLAN-63-2 guid 70a92b45-f60f-4886-8729-c4a656573b70
  member-of network segment pool Cluster-Pool-2
  switchport access vlan 63
  ip pool import template VLAN-63-Pool2 guid 5d3f9e00-8668-45e4-bef1-307629e629e
4
  publish network segment
  switchport mode access

```

```

nsm network segment VLAN-64-2
  guid a86dd83e-f96d-44d1-a79e-c63ae2f8bbd7
  member-of vmnetwork VLAN-64-2 guid 648affe3-4a92-4861-9609-0598b1ddc159
  member-of network segment pool Cluster-Pool-2

  switchport access vlan 64
  ip pool import template VLAN-64-Pool2 guid a44544ba-b10c-4fed-a57c-00c89790f04
3
  publish network segment
  switchport mode access

```

```

nsm network uplink nexus1000v-uplink2
  import port-profile Uplink-profile2
  allow network segment pool VDI-Pool-2

```

allow network segment pool Cluster-Pool-2
publish network uplink

13.4. Sample PowerShell Scripts

13.4.1. Update Virtual Machines created by XenDesktop Wizard

The following PowerShell script can be used to update the virtual machines after the XenDesktop Wizard creates them. The SCVMM cloning process used by the XenDesktop Wizard does not use all the properties from the template, such as the network adapter, VLAN, or boot order, so this PowerShell script updates the VMs to have the correct properties. This script must be run from the SCVMM server, since it uses SCVMM specific PowerShell cmdlets.

Before running, fill out the PARAMETERS section to match your environment. This script then performs the following actions:

1. Checks if the virtual machine is running, if so attempts to stop it.
2. Checks if the virtual machine is in a bad state and needs to be repaired, if so it repairs it.
3. Sets the Startup Action on the VM so it will not boot automatically upon a host restart. This prevents the logon storms and allows XenDesktop controller to pace the virtual machine starts
4. Sets the Stop Action on the VM to turn off VM, which prevents the creation of the BIN file (equal to the size of the RAM assigned for saving-state).
5. Sets the boot order to PXEBoot (boot from LegacyNic), CD, IDE, Floppy, so that the VMs can PXE boot off the PVS servers.
6. Sets each NetworkAdapter to use the Nexus 1000V logical switch and associated Port Profile.
7. Outputs status on each VM update for monitoring progress

This script is provided as a sample of how to complete these actions on System Center Virtual Machine Manager 2012 SP1.

```
# Purpose: Update VM settings on a single host after the XenDesktop
Setup Wizard runs:
#           Sets StartAction to not auto-start
#           Sets StopAction to Turn off
#           Moves PXE Boot to top of boot order
#           Sets all network adapters for every VM to the specified VM
Network as specified in the Parameters section.

# Author: Paul Wilson
# Date: 20 October 2013

#PARAMETERS SECTION
# Fill out these values based on your environment before running the
script.

$VMNetworkName = "VLAN-62"
$VMSubnetName = "VLAN-62"
$PortClassName = "VDI Port Class"
```

```

$HostNameFQDN = "vdi1-3.hv.pod.local"
$VirtualNetworkName = "Nexus 1000 V"

#END PARAMETERS SECTION

#Import the System Center Virtual Machine Manager if not already
available
if (!(get-module VirtualMachineManager)) {Import-module
VirtualMachineManager}

#Get the objects for the VMNetwork and Port Classification based upon
the names
#supplied in the Parameters section

$VMNetwork= Get-SCVMNetwork | where {$_.Name -eq $VMNetworkName}
$PortClass= Get-SCPortClassification | where {$_.Name -eq
$PortClassName}
$VirtualNetwork = Get-SCVirtualNetwork | where {($_.vmhost -eq
$HostNameFQDN) -and ($_.Name -like $VirtualNetworkName)}
$VMsubnet = get-scvmsubnet | where {$_.name -like $VMSubnetName}

#Get the VMs to update
$AllVMs = get-scvirtualmachine | where {$_.HostName -eq $HostNameFQDN}
| Sort-object Name

#Process each of the VMs on the specified host. Update all the network
adapters to use the
#supplied VLAN-ID, VM Network, and Port Classification.

foreach ($myvm in $allVms)
{
    Try
    {
        If ($myVM.status -eq "Running")
        {
            stop-vm -VM $myvm
        }
        if ($myvm.Status -eq "UpdateFailed")
        {
            write-output "Repairing $myvm..."
            $result=Repair-SCVirtualMachine -VM $myvm -Dismiss -Force
        }
        $result=set-scvirtualmachine -vm $myvm -startAction
"NeverAutoTurnOnVM" -stopAction "TurnOffVM" -BootOrder
PxeBoot,CD,IdeHardDrive,Floppy
        $myNetworkAdapters = get-SCVirtualnetworkadapter -VM $myvm
#This ForEach loop can be skipped if using the Citrix Hotfix.
        foreach ($na in $myNetworkAdapters)
        {
            $Error.clear()
            if ($na.VMNetwork -notlike $VMNetworkName)
            {

```



```

        $result=set-scvirtualnetworkadapter -VMNetwork $VMNetwork -
PortClassification $PortClass -VirtualNetworkAdapter $na -
VirtualNetwork

$VirtualNetwork -VMSubnet $VMSubnet
    }
    else
    {
        $result=set-scvirtualnetworkadapter -NoConnection -
VirtualNetworkAdapter $na
    }
}

if ($error.count -eq 0)
{
    $s = "Update on VM {0} was successful." -f $myvm.Name
}
else
{
    $s = "Update on VM {0} failed!!!" -f $myvm.Name
}

write-output $s
}
Catch
{
    $s2 = "Update on VM {0} failed!" -f $myvm.Name
    write-output $s2
}
}

```

13.4.2. Enable Dynamic Memory

By default the XenDesktop wizard creates virtual machines with static memory if your template has . The following script can be used to enable Dynamic Memory on a group of virtual machines. This script must be run from the Hyper-V host where the virtual machines reside. The \$NameMatch parameter can be used to filter which virtual machines are updated.

```

# Purpose: Enabled dynamic memory for a group of VMs based on
parameters supplied.

# Author: Paul Wilson
# Date: 15 August 2013

#PARAMETERS SECTION
# Fill out these values based on your environment before running the
script

$NameMatch = "VDI"
$minMem = 1536MB

```

```

$maxMem = 4GB

#END PARAMETERS SECTION

$AllVMs = get-vm | where {$_.Name -match $NameMatch}

foreach ($myvm in $allVms)
{
    Try
    {
        If ($myVM.state -eq "Running")
        {
            stop-vm -Name $myvm.Name -TurnOff
        }
        set-vmmemory -VM $myvm -DynamicMemoryEnabled $True -MinimumBytes
$minMem -MaximumBytes $maxMem -StartupBytes $minMem
        $s = "Update on VM {0} was successful." -f $myvm.Name

        write-output $s
    }
    Catch
    {
        $s2 = "Update on VM {0} failed!" -f $myvm.Name
        write-output $s2
    }
}

```

13.4.3. Disable Dynamic Memory

The following script can be used to disable dynamic memory. This script must be run from the Hyper-V host where the virtual machines reside. The \$NameMatch parameter can be used to filter which virtual machines are updated.

```

# Purpose: Disable dynamic memory for a group of VMs based on
parameters supplied.

# Author: Paul Wilson
# Date: 15 August 2013

#PARAMETERS SECTION
# Fill out these values based on your environment before running the
script

$NameMatch = "VDI"
$StaticMem = 1536MB

#END PARAMETERS SECTION

$AllVMs = get-vm | where {$_.Name -match $NameMatch}

foreach ($myvm in $allVms)

```

```

{
    Try
    {
        If ($myVM.state -eq "Running")
        {
            stop-vm -Name $myvm.Name -TurnOff
        }
        set-vmmemory -VM $myvm -DynamicMemoryEnabled $False -StartupBytes
$StaticMem
        $s = "Update on VM {0} was successful." -f $myvm.Name

        write-output $s
    }
    Catch
    {
        $s2 = "Update on VM {0} failed!" -f $myvm.Name
        write-output $s2
    }
}

```

13.4.4. Query the XenDesktop Database Connection Strings

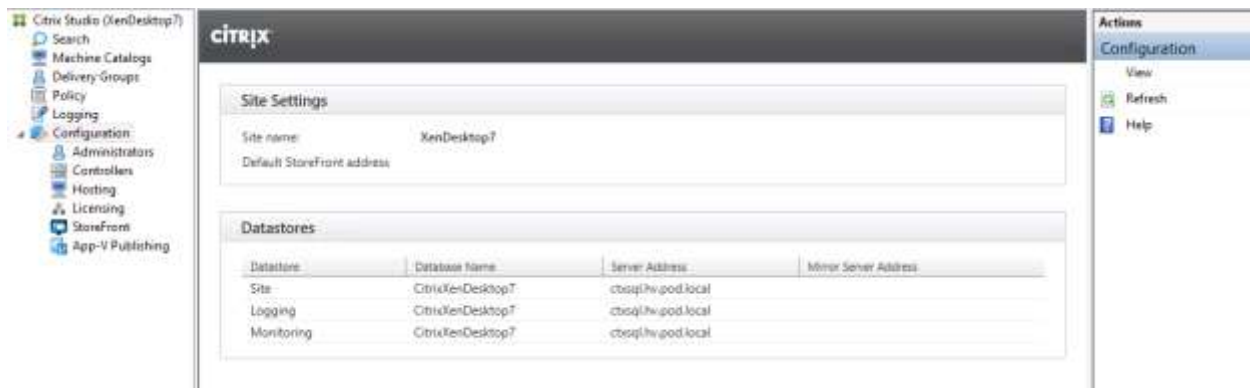
Execute these PowerShell commands from one of the XenDesktop controllers or a machine. If not executed from a XenDesktop controller add the -AdminAddress parameter to point to an existing controller. The response should result in the existing connection string for each of the database schemas and they should be the same. This step is recommended before changing the database strings in case you need to reverse the change.

```

Add-PSSnapin Citrix*
Get-AdminDBConnection
Get-ConfigDBConnection
Get-LogDBConnection
Get-AcctDBConnection
Get-HypDBConnection
Get-ProvDBConnection
Get-BrokerDBConnection
Get-MonitorDBConnection
Get-SfDBConnection
Get-EnvTestDBConnection

```

These commands are only for the Site datastore. To see the location for the Logging and Monitoring datastores, they can be viewed on the Configuration node of the Citrix Studio console.



13.4.5. Test the XenDesktop Database Connection String

Use the following PowerShell commands to test the connectivity to the new location of the XenDesktop database before actually switching it over. Testing before switching over is recommended since to switch over you need to delete the old connection string first.

```
Add-PSSnapin Citrix*
```

```
$cs = "Server=ctxsql.hv.pod.local;Initial  
Catalog=CitrixXenDesktop7;Integrated Security=True"
```

```
Test-AdminDBConnection -DBConnection $cs  
Test-ConfigDBConnection -DBConnection $cs  
Test-LogDBConnection -DBConnection $cs  
Test-AcctDBConnection -DBConnection $cs  
Test-HypDBConnection -DBConnection $cs  
Test-ProvDBConnection -DBConnection $cs  
Test-BrokerDBConnection -DBConnection $cs  
Test-MonitorDBConnection -DBConnection $cs  
Test-SfDBConnection -DBConnection $cs  
Test-EnvTestDBConnection -DBConnection $cs
```

Note: There should be no line breaks in the \$cs variable declaration, which is wrapped for readability.

13.4.6. Change the XenDesktop Database Connection String

Use the following PowerShell commands to change the location of the XenDesktop site datastore. These commands must be run on every XenDesktop controller in the environment.

```
Add-PSSnapin Citrix*
```

```
$cs = "Server=ctxsql.hv.pod.local;Initial  
Catalog=CitrixXenDesktop7;Integrated Security=True"
```

```
Set-AcctDBConnection -DBConnection $null  
Set-AcctDBConnection -DBConnection $cs  
Set-AdminDBConnection -DBConnection $null  
Set-AdminDBConnection -DBConnection $cs
```

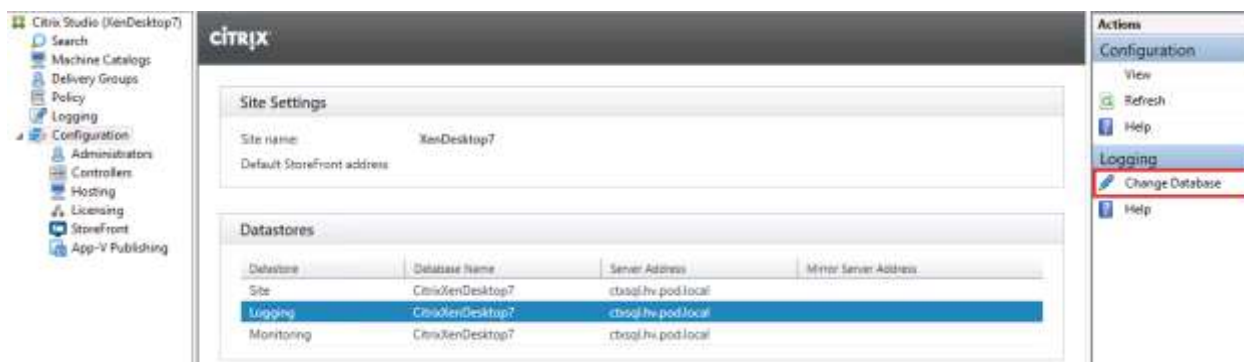
```

Set-BrokerDBConnection -DBConnection $null
Set-BrokerDBConnection -DBConnection $cs
Set-ConfigDBConnection -DBConnection $null
Set-ConfigDBConnection -DBConnection $cs
Set-EnvTestDBConnection -DBConnection $null
Set-EnvTestDBConnection -DBConnection $cs
Set-HypDBConnection -DBConnection $null
Set-HypDBConnection -DBConnection $cs
Set-LogDBConnection -DBConnection $null
Set-LogDBConnection -DBConnection $cs
Set-MonitorDBConnection -DBConnection $null
Set-MonitorDBConnection -DBConnection $cs
Set-ProvDBConnection -DBConnection $null
Set-ProvDBConnection -DBConnection $cs
Set-SfDBConnection -DBConnection $null
Set-SfDBConnection -DBConnection $cs

```

Note: There should be no line breaks in the \$cs variable declaration, which is wrapped for readability.

When finished use the Get-xxxDBConnection commands to verify that the database connection strings have changed and then reboot the XenDesktop controller. To change the datastore for the Logging and Monitoring databases, use Change Database command in Citrix Studio.



14. Acknowledgements

Cisco would like to thank the following individuals for their authoring of this document:

Frank Anderson, Principal Solutions Architect, Strategic Alliance at Citrix Systems: Frank Anderson is a Principal Solutions Architect at Citrix, focusing on Desktop and Application Virtualization. Responsibilities include solutions validation, strategic alliances, technical content creation, and testing/benchmarking.

Rob Simpson, Sr Program Manager WSSC Partner and Customer Ecosystem at Microsoft: Rob's been with Microsoft for over 8 years. He started out at the Microsoft Enterprise Engineering Center (EEC) as a customer engagement PM running Exchange and Active Directory customer proof of concepts. For the Exchange 2010 release, Rob moved over to the Exchange product group as a tester in the Scalability and Capacity Planning team. He then started the Exchange Tested Solutions Program showcasing partner

reference architectures for Exchange 2010. Rob then returned to the EEC and took over the EEC Partner Program. Recently Rob has focused on driving a number of VDI reference architecture projects.

Paul Wilson, Technical Marketing Engineer for Desktop Virtualization at Cisco : Paul Wilson is a technical marketing engineer specializing in desktop virtualization. Paul has been involved in performance testing virtualization solutions for over 10 years and specializes in Citrix and Microsoft offerings.

Ka-kit Wong, Solutions Engineer, Strategic Solutions Engineering at EMC: Ka-Kit Wong is a solutions engineer for desktop virtualization in EMC's Strategic Solutions Engineering group, where he focuses on developing End User Computing (EUC) validated solutions. He has been at EMC for more than 13 years, and his roles have included systems, performance and solutions testing in various positions. He holds a master of science degree in computer science from Vanderbilt University.

Cisco would also like to acknowledge the contributions of the following individuals to this document:

- Mike Brennan, Cisco Systems
- Sai Chaitanya, Cisco Systems
- Richard Dowell, Citrix Systems
- John Moran, EMC
- Hardik Patel, Cisco Systems
- Rajmohan Rajanayagam, EMC
- Derek Rice, Citrix Systems