Desktop-as-a-Service for Service Provider 2000-Seat Virtual Desktop Infrastructure

Citrix XenDesktop/XenApp 7.5 built on Cisco UCS B200-M3 Blades with EMC VNX5600 and VMware vSphere 5.5

Last Updated: April 8, 2015

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Jeff Nichols, Technical Marketing Engineer, VDI Performance and Solutions Team, Cisco Systems

Jeff Nichols is a Cisco Unified Computing System architect, focusing on Virtual Desktop and Application solutions with extensive experience with VMware ESX/ESXi, XenDesktop, XenApp and Microsoft Remote Desktop Services. He has expert product knowledge in application, desktop and server virtualization across all three major hypervisor platforms and supporting infrastructures including but not limited to Windows Active Directory and Group Policies, User Profiles, DNS, DHCP and major storage platforms.

Rob Briggs - Principal Solutions Architect, Citrix Systems, Inc.

Rob Briggs is an internationally experienced professional who has been involved in both the public and tech sectors for over 20 years. He is a highly accomplished IT systems leader, architect and analyst specializing in integration and design of advanced virtualized system solutions. His expertise helps teams bridge the gap between technology, solutions marketing and systems management.

Ka-kit Wong, Solutions Engineer, Strategic Solutions Engineering, 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.

Acknowledgements

For his support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to thank:

- Mike Brennan, Sr. Technical Marketing Engineer, VDI Performance and Solutions Team Lead, Cisco Systems, Inc.
- Phani Penmethsa, Technical Marketing Engineer, Cisco Systems, Inc.
- Kurtis Moody, Director, Product Marketing, Citrix Systems, Inc.
- Selma Wei, Principal Consultant, Citrix Systems, Inc.
- Scott Harris, Senior Director, Citrix Systems, Inc.
- Priyadarshan Ketkar, Senior Cloud Architect, Citrix Systems, Inc.
- Daniel L'Hommedieu, Product Manager, Citrix Systems, Inc.

About Cisco Validated Design (CVD) Program

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

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

Desktop-as-a-Service for Service Provider 2000-Seat Virtual Desktop Infrastructure

About this Document

This document provides a Reference Architecture for a 2000-Seat Virtual Desktop Infrastructure using Citrix XenDesktop 7.5 built on Cisco UCS B200-M3 blades with an EMC VNX5600 and the VMware vSphere ESXi 5.5 hypervisor platform.

Audience

This document is designed for use by the Desktop as a Service deployment and management technical team at the Cisco-Citrix-EMC Service Provider.

Business executives in the Service Provider organization can gain an overall understanding of the solutions components and the system architecture in the early sections of the Cisco Validated Design.

Overview

The landscape of desktop virtualization is changing constantly. New, high performance Cisco UCS Blade Servers and Cisco UCS unified fabric combined with the latest generation EMC VNX arrays result in a more compact, more powerful, more reliable and more efficient platform.

In addition, the advances in the Citrix XenDesktop 7.5 system, which now incorporates both traditional hosted virtual Windows 7 or Windows 8 desktops, hosted applications and hosted shared Server 2008 R2 or Server 2012 R2 server desktops (formerly delivered by Citrix XenApp,) provide unparalleled scale and management simplicity while extending the Citrix HDX FlexCast models to additional mobile devices

This document provides the architecture and design of a virtual desktop infrastructure for 2000 mixed use-case users. The infrastructure is 100% virtualized on VMware ESXi 5.5 with third-generation Cisco UCS B-Services B200 M3 blade servers booting via Fiber Channel from an EMC VNX5600 storage array. The virtual desktops are powered using Citrix Provisioning Server 7.1 and Citrix XenDesktop 7.5, with a mix of hosted shared desktops (90%) and pooled Server Virtual Desktops (ServerVDI) desktops

(10%) to support the user population. Where applicable, the document provides best practice recommendations and sizing guidelines for customer deployments of Citrix DaaS Solution on the Cisco Unified Computing System.

Multi-Tenant DaaS Architecture

As mobile technologies have become pervasive, businesses are looking to take advantage of Bring

Your Own Device (BYOD) and anywhere computing workstyles. User mobility is a compelling force behind the increase in Desktop-as-a-Service (DaaS) initiatives that are creating new and growing opportunities for Cisco Powered™ Cloud, Citrix® and EMC Service Provider Partners.

Employees that have ubiquitous access are more productive, responsive, and collaborative since they can work wherever and whenever it's convenient-at home, on the road, and in the office. Employees that have ubiquitous access are more productive, responsive, and collaborative since they can work wherever and whenever it's convenient-at home, on the road, and in the office.

To help providers rapidly deploy DaaS functionality and manage it at cloud scale, Cisco and Citrix engineers have validated this new reference architecture that is hosted on the Cisco Unified Data Center (UDC) architecture. On this scalable and cost-effective platform, the architecture delivers high user densities, outstanding End-User experience, proven Cisco network reliability and security, and easy deployment and manageability, while providing a rich mobile user experience.

Benefits Service Providers will gain by deploying this solution include:

- Simplified administration and management
- Intuitive infrastructure management and low total cost of ownership (TCO)
- Support for a variety of multi-tenancy isolation models
- Easy platform expansion
- Exceptional simplicity for provisioning and onboarding
- Optimized user experience.

Summary of Main Findings

The combination of technologies from Cisco Systems, Inc, Citrix Systems, Inc, VMware and EMC produced a highly efficient, robust and scalable 2000 seat mixed Desktop as a Service Infrastructure delivering outstanding end-user experience with the following concurrently running workloads:

- 200 Citrix XenDesktop 7.5 Server VDI Desktops using Windows Server 2012 R2
- 1800 Citrix XenDesktop 7.5 Hosted Shared Desktops.
- 10 Separate DaaS Tenants provisioned and running simultaneously.

The combined power of the Cisco Unified Computing System, Nexus switching and EMC storage hardware, VMware ESXi 5.5, Citrix XenDesktop 7.5, Citrix App Orchestrator 2.5, Citrix CloudPortal Services Manager and Cisco UCS Director produces a high-density per blade and per chassis mixed workload Virtual Desktop delivery system with the following:

- 14 Cisco UCS B200 M3 half-width blades with dual 10-core processors, 256GB of 1600 MHz memory that boot from SAN and house PVS write-cache were able to successfully run 2000 desktops spread across 10 different tenants ranging from a public shared model to private, isolated models.

- 2 Cisco UCS B200 M3 half-width blades with dual 10-core processors, 256GB of 1600 MHZ memory were able to handle all tenant infrastructure VMs (i.e. Active Directory, DNS, SQL, Citrix components)
- Pure Virtualization: We continue to present a validated design that is 100% virtualized on ESXi 5.5. All of the Windows Server 2012 R2 virtual desktops and supporting infrastructure components, including Active Directory, Profile Servers, Provisioning Servers, SQL Servers, and XenDesktop delivery controllers were hosted as virtual servers.
- We maintain our industry leadership with our new Cisco UCS Manager 2.2(2c) software that makes scaling simple, consistency guaranteed and maintenance simple. Combined with UCS Central, our Cisco UCS management scope extends to over 100,000 virtual desktops.
- Our 10G unified fabric story gets additional validation on second generation Cisco UCS 6200 Series Fabric Interconnects and second generation Nexus 5500 Series access switches as we run more challenging workload testing, maintaining unsurpassed user response times.
- EMC's VNX 5600 system provides storage consolidation and outstanding efficiency. Both block and NFS storage resources were provided by a single system.
- EMC's Fast Cache technology delivers predictable performance and continuous availability for end user computing environment ensuring a rich end user experience while enforcing compliance, data security, high- availability, and increasing IT productivity.
- EMC delivers simplified management to an End User Computing infrastructure through Unisphere for ease of configuration and management, plug-in technologies for simplified desktop provisioning, and integrations that deliver rich metrics for monitoring your VNX storage platform.
- Citrix HDX technology, extended in XenDesktop 7.5 software, provides excellent performance with host-rendered flash video and other demanding applications.
- Citrix XenDesktop 7.5 extends the flexibility of the solution design by adding hosted shared server desktops and Server VDI desktops.

Solution Component Benefits

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

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 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 20 Chassis and up to 160 blades in a single UCS management domain
- Scale to multiple UCS Domains with Cisco UCS Central within and across data centers globally
- Leverage UCS Management Packs for VMware vCenter 5.5 for integrated management

Benefits of Cisco Nexus Physical Switching

The Cisco Nexus product family includes lines of physical unified port layer 2, 10 GB switches, fabric extenders, and virtual distributed switching technologies. In our study, we utilized Cisco Nexus 5548UP physical switches and Cisco Nexus 1000V distributed virtual switches to deliver amazing end user experience

Cisco Nexus 5548UP Unified Port Layer 2 Switches

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 (ISCSI)
- 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 ISCSI
- Throughput of up to 960 Gbps.

Cisco Nexus 1000V Distributed Virtual Switch

Get highly secure, multitenant services by adding virtualization intelligence to your data center network with the Cisco Nexus 1000V Switch for VMware vSphere. This switch:

- Extends the network edge to the hypervisor and virtual machines
- Is built to scale for cloud networks
- Forms the foundation of virtual network overlays for the Cisco Open Network Environment and Software Defined Networking (SDN)

Important differentiators for the Cisco Nexus 1000V for VMware vSphere include:

- Extensive virtual network services built on Cisco advanced service insertion and routing technology
- Support for vCloud Director and vSphere hypervisor
- Feature and management consistency for easy integration with the physical infrastructure
- Exceptional policy and control features for comprehensive networking functionality
- Policy management and control by the networking team instead of the server virtualization team (separation of duties)

Use Virtual Networking Services

The Cisco Nexus 1000V Switch optimizes the use of Layer 4 - 7 virtual networking services in virtual machine and cloud environments through Cisco vPath architecture services.

Cisco vPath 2.0 supports service chaining so you can use multiple virtual network services as part of a single traffic flow. For example, you can simply specify the network policy, and vPath 2.0 can direct traffic:

- Through the Cisco ASA1000V Cloud Firewall for tenant edge security
- Through the Cisco Virtual Security Gateway for Nexus 1000V Switch for a zoning firewall

In addition, Cisco vPath works on VXLAN to support movement between servers in different Layer 2 domains. Together, these features promote highly secure policy, application, and service delivery in the cloud.

Benefits of EMC VNX5600 Storage Arrays

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 storage includes the following components:

- Host adapter ports (for block)-Provide host connectivity through fabric into the array.
- Data Movers (for file)-Front-end appliances that provide file services to hosts (optional if providing CIFS/SMB or NFS services).
- Storage processors (SPs)-The compute component of the storage array. SPs handle all aspects of data moving into, out of, and between arrays.
- Disk drives-Disk spindles and solid state drives (SSDs) that contain the host/application data and their enclosures.

Note The term Data Mover refers to a VNX hardware component, which has a CPU, memory, and input/output (I/O) ports. It enables the CIFS (SMB) and NFS protocols on the VNX array.

EMC Next-Generation VNX Series

Next-generation 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 with the new Unisphere® Management Suite

VSPEX is built with next-generation VNX to deliver even greater efficiency, performance, and scale than ever before.

Flash-Optimized Hybrid Array

VNX is a flash-optimized hybrid array that provides automated tiering to deliver the best performance to your critical data, while intelligently moving less frequently accessed data to lower-cost disks.

In this hybrid approach, a small percentage of flash drives in the overall system can provide a high percentage of the overall IOPS. Flash-optimized VNX takes full advantage of the low latency of flash to deliver cost-saving optimization and high performance scalability. EMC Fully Automated Storage Tiering Suite (FAST Cache and FAST VP) tiers both block and file data across heterogeneous drives and boosts the most active data to the flash drives, ensuring that customers never have to make concessions for cost or performance.

Data generally is accessed most frequently at the time it is created; therefore, new data is first stored on flash drives to provide the best performance. As the data ages and becomes less active over time, FAST VP tiers the data from high-performance to high-capacity drives automatically, based on

customer-defined policies. This functionality has been enhanced with four times better granularity and with new FAST VP solid-state disks (SSDs) based on enterprise multilevel cell (eMLC) technology to lower the cost per gigabyte.

FAST Cache uses flash drives as an expanded cache layer for the array to dynamically absorb unpredicted spikes in system workloads. Frequently accessed data is copied to the FAST Cache in 64 KB increments. Subsequent reads and/or writes to the data chunk are serviced by FAST Cache. This enables immediate promotion of very active data to flash drives. This dramatically improves the response times for the active data and reduces data hot spots that can occur within the LUN.

All VSPEX use cases benefit from the increased efficiency provided by the FAST Suite. Furthermore, VNX provides out-of-band, block-based deduplication that can dramatically lower the costs of the flash tier.

VNX Intel MCx Code Path Optimization

The advent of flash technology has been a catalyst in making significant changes in the requirements of midrange storage systems. EMC redesigned the midrange storage platform to efficiently optimize multicore CPUs to provide the highest performing storage system at the lowest cost in the market.

MCx distributes all VNX data services across all cores (up to 32), as shown in Figure 1. The VNX series with MCx has dramatically improved the file performance for transactional applications like databases or virtual machines over network-attached storage (NAS).

Figure 1 ***Next-Generation VNX with Multicore Optimization***

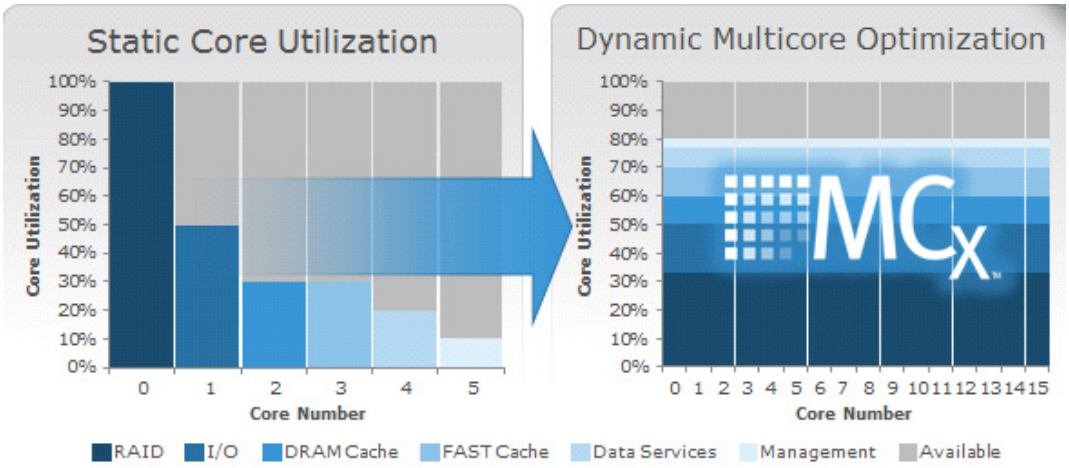

Multicore Cache

The cache is the most valuable asset in the storage subsystem; its efficient use is the key to the overall efficiency of the platform in handling variable and changing workloads. The cache engine has been modularized to take advantage of all the cores available in the system.

Multicore RAID

Another important improvement to the MCx design is how it handles I/O to the permanent back-end storage—hard disk drives (HDDs) and SSDs. The modularization of the back-end data management processing, which enables MCx to seamlessly scale across all processors, greatly increases the performance of the VNX system.

Performance Enhancements

VNX storage, enabled with the MCx architecture, is optimized for FLASH 1st and provides unprecedented overall performance; it optimizes transaction performance (cost per IOPS), bandwidth performance (cost per GB/s) with low latency, and capacity efficiency (cost per GB).

VNX provides the following performance improvements:

- Up to four times more file transactions when compared with dual controller arrays
- Increased file performance for transactional applications (for example, Microsoft Exchange on VMware over NFS) by up to three times, with a 60 percent better response time
- Up to four times more Oracle and Microsoft SQL Server OLTP transactions
- Up to six times more virtual machines

Active/Active Array Storage Processors

The new VNX architecture provides active/active array storage processors, as shown in Figure 2, which eliminate application timeouts during path failover because both paths are actively serving I/O.

Figure 2 ***Active/Active Processors Increase Performance, Resiliency, and Efficiency***

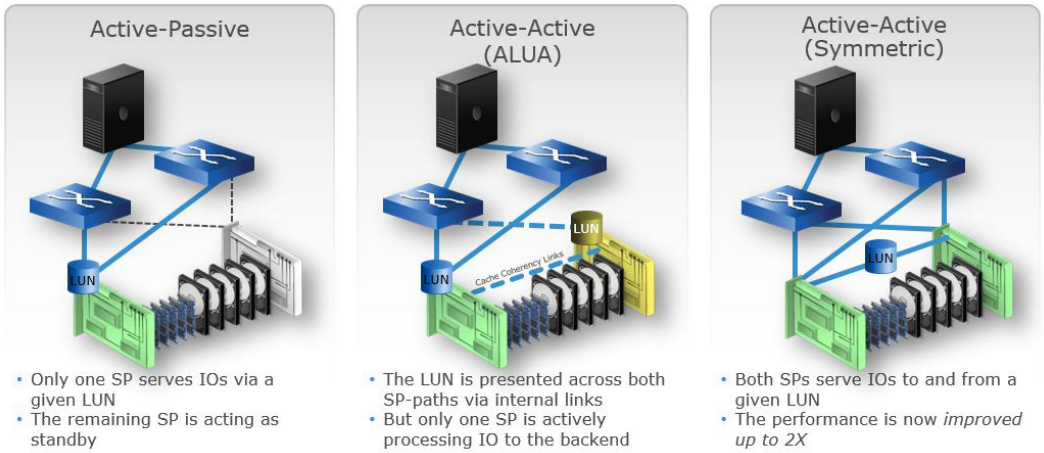

Load balancing is also improved, providing up to double the performance for applications. Active/active for block is ideal for applications that require the highest levels of availability and performance, but do not require tiering or efficiency services like compression, deduplication, or snapshot.

Note The active/active processors are available only for RAID LUNs, not for pool LUNs.

Benefits of VMware vSphere ESXi 5.5

VMware vSphere® 5.5 is the latest release of the flagship virtualization platform from VMware. VMware vSphere, known in many circles as "ESXi", for the name of the underlying hypervisor architecture, is a bare-metal hypervisor that installs directly on top of your physical server and partitions it into multiple virtual machines. Each virtual machine shares the same physical resources as the other virtual machines and they can all run at the same time. Unlike other hypervisors, all management functionality of vSphere is possible through remote management tools. There is no underlying operating system, reducing the install footprint to less than 150MB.

The following are some key features included with vSphere 5.5:

- Improved Security
- Extensive Logging and Auditing
- Enhanced vMotion
- New Virtual Hardware
- Active Directory Integration
- Centralized Management
- Stateless Firewall
- Centralized Management of Host Image and Configuration via Auto Deploy.

For more information on the vSphere ESXi hypervisor, go to:

http://www.vmware.com/products/esxi-and-esx/overview.html

Benefits of Citrix XenDesktop 7.5

Service Providers and other IT organizations are tasked with the challenge of provisioning Microsoft Windows apps and desktops while managing cost, centralizing control, and enforcing corporate security policy. Deploying Windows apps to users in any location, regardless of the device type and available network bandwidth, enables a mobile workforce that can improve productivity. With Citrix XenDesktop™ 7.5, IT can effectively control app and desktop provisioning while securing data assets and lowering capital and operating expenses.

In addition to providing a performance boost over the previous XenDesktop 7.1 version, the XenDesktop™ 7.5 release offers these benefits:

- **Comprehensive virtual desktop delivery for any use case.** The XenDesktop 7.5 release incorporates the full power of XenApp, delivering full desktops or just applications to users. Administrators can deploy both XenApp published applications and desktops (to maximize IT control at low cost) or personalized VDI desktops (with simplified image management) from the same management console. Citrix XenDesktop 7.5 leverages common policies and cohesive tools to govern both infrastructure resources and user access.
- **Simplified support and choice of BYO (Bring Your Own) devices**. XenDesktop 7.5 brings thousands of corporate Microsoft Windows-based applications to mobile devices with a native-touch experience and optimized performance. HDX technologies create a “high definition” user experience, even for graphics-intensive design and engineering applications.
- **Lower cost and complexity of application and desktop management.** XenDesktop 7.5 helps IT organizations take advantage of agile and cost-effective cloud offerings, allowing the virtualized infrastructure to flex and meet seasonal demands or the need for sudden capacity changes. IT organizations can deploy XenDesktop application and desktop workloads to private or public clouds, including Amazon AWS, Citrix Cloud Platform, and (in the near future) Microsoft Azure.

- **Protection of sensitive information through centralization.** XenDesktop decreases the risk of corporate data loss, enabling access while securing intellectual property and centralizing applications since assets reside in the datacenter.

Benefits of Citrix App Orchestration 2.5

Citrix App Orchestration 2.5 allows CSPs to orchestrate and automate the delivery of applications and desktops in multi-tenant environments and across multiple products, sites, and datacenters. With App Orchestration, service providers can:

- Manage XenApp and XenDesktop across multiple locations, including multiple datacenters in multiple versions, sites or farms, Active Directory domains, and datacenters, from a single unified interface.
- Provide consistent configuration across global deployments spanning multiple delivery sites, eliminating configuration drift and issues.
- Define tenant
- and user affinity to deliver offerings to primary and backup locations, for optimum continuity and fault tolerance.
- Provision desktops and applications on any supported hypervisor, including VMware ESXi as in this CVD. App Orchestration can incorporate externally provisioned VMs (e.g., provisioning through Cisco UCS Director as in this architecture).

App Orchestration 2.5 provides the following features to simplify cloud-scale administration for service providers:

- **Simplified management across datacenters.** App Orchestration simplifies how Citrix technologies can be provisioned and deployed on virtual servers across datacenters. Given pools of XenDesktop Session Machines and Delivery Controllers, App Orchestration automatically manages capacities across multiple sites and datacenters, even managing multiple product versions and farms/sites in physically different domains.
- **Multi-tenant configuration.** Support for different types of isolation models (e.g., Session-based, Server-based, and Site-based) on a per-application or per-desktop basis. There are two areas in App Orchestration 2.5 in which you can specify isolation levels: delivery isolation and tenant StoreFront isolation.
- **Simplified, secure network configuration with zero trust domains.** New in App Orchestration 2.5, zero trust domains eliminate the Active Directory requirement for private tenant domains. Instead, a domain agent resides in the private tenant domain, and App Orchestration uses SSL client certificates to validate agent identity. This approach eliminates the need for various ports to be opened on the firewall, resulting in a simplified and hardened network configuration for dedicated environments.
- **Quick application and desktop configuration.** App Orchestration enables the harmonious configuration and integration of XenDesktop, XenApp, NetScaler, and Active Directory. This helps automate the installation of farms/sites, Session Machines and StoreFront server groups. Automatic discovery of application information (including name, icon, command line, working directory, etc.) from a XenDesktop HSD host can save valuable administrative time.
- **App Orchestration web management console.** App Orchestration supplies a web-based management console to control App Orchestration activities. You can use the console to monitor workflows for deployment actions, such as creating Delivery Sites or adding Session Machines.

- **Easier patching of XenDesktop HSD hosts.** When you create a new version of a XenDesktop HSD host, App Orchestration can automate the tasks of gradually draining users from the older version servers to the newer ones, without any downtime or manual intervention.
- **Tenant Management.** The administrator defines tenants into the system, their desired level of isolation, and assigns resources to them directly. The console allows the administrator to easily view which resources (applications, desktops, XenDesktop HSD hosts, sites, etc.) are allocated to which tenants.
- **CloudPortal Services Manager integration.** Using App Orchestration with CloudPortal Services Manager, CSPs can enable multi-tenant self-servicing of application and desktop offerings that are configured through App Orchestration. This capability empowers a degree of self-support, delegating control to the tenant administrator to manage user subscription offerings.

For CSPs unfamiliar with Citrix App Orchestration, see the App Orchestration documentation at the Citrix eDocs site:

http://support.citrix.com/proddocs/topic/app-orchestration/cao-app-orchestration-25-landing.html

Benefits of Citrix CloudPortal Services Manager

Citrix CloudPortal Services Manager (CPSM) is a self-service portal that helps providers manage the delivery of cloud services and customer offerings. It drives App Orchestration operation by associating desired states with tenants and allowing services to be provisioned to users.

CPSM provides out-of-the-box support for Desktop-as-a-Service and Windows applications (powered by Citrix XenApp and Citrix XenDesktop), as well as popular business applications and services like Microsoft Exchange, Office, SharePoint, Lync, web and data hosting, and virtualization service management. Customers and sub-customers (for example, such as a reseller's customers) that lack IT expertise can add or change services, view reports, manage users, and perform day-to-day administration tasks through the self-service interface.

The integration of Citrix CloudPortal Services Manager and App Orchestration allows providers to scale their subscriber base while using the existing support staff, helping to increase provider profit margins.

DaaS Architectural Overview

This reference architecture enables Citrix Service Providers to deliver Windows applications and desktops as Desktop-as-a-Service (DaaS) through an integrated set of Citrix and partner technologies:

- Citrix XenDesktop unifies the delivery of hosted applications and desktops (XenApp) with virtual desktops (XenDesktop) using a single architecture and management experience.
- Citrix App Orchestration allows CSPs to automate and manage, at a high scale, the delivery of DaaS offerings in a multi-tenant environment.
- CloudPortal Services Manager supplies a portal to manage service delivery and subscriber offerings. It supports delegated management roles that enable down-channel partners and tenant administrators to self-provision and monitor provisioning requests.

Key Architectural Modules

Figure 3 illustrates the architecture for the DaaS software solution can be divided into four logical modules: (1) Infrastructure-as-a-Service (IaaS), (2) Multi-Tenant Citrx Farms and Sites, (3) Dashboards and Management, and (4) Endpoints and Offices.

Figure 3 Four Architectural Modules of the Daas Logical Architecture

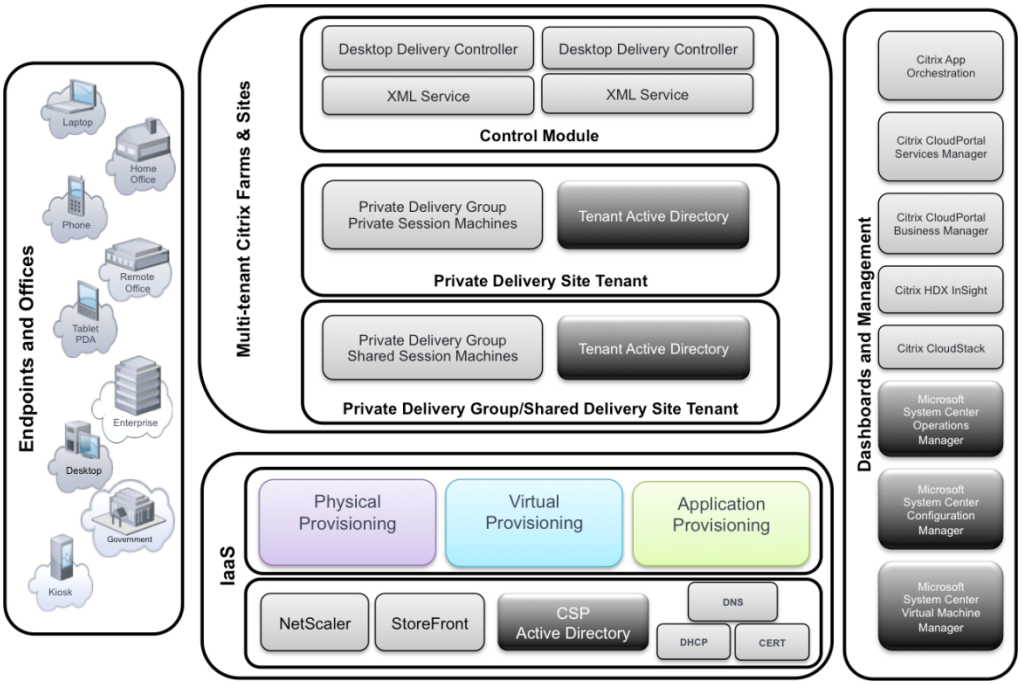

Infrastructure-as-a-Service (IaaS) Module

The foundation of the architecture is the Infrastructure-as-a-Service (IaaS) module, which is responsible for network, authentication, and provisioning functions (Figure 4). It has two sub-layers:

- Network infrastructure, including the implementation of Active Directory domains
- Provisioning infrastructure, including virtual, application, and physical provisioning

Figure 4 Infrastructure-as-a-Service (IaaS) Module

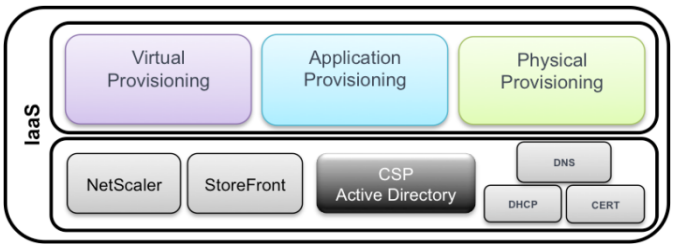

The IaaS module controls the system-wide network configuration, forest-level Active Directory management, remote access, and all layers of provisioning.

Multi-Tenant Citrix Farms and Sites Module

The Multi-Tenant Citrix Farms and Sites module (Figure 5) is the core component of the service provider datacenter — this logical block controls application and desktop delivery within the multi-tenant architecture.

Note In the unified XenDesktop 7.5 release, a "site" rather than a "farm" is the main XenDesktop environment consisting of Delivery Controllers and a database used to deliver both XenApp and XenDesktop services.

Within a multi-tenant datacenter, applications and desktops are virtualized and subscriber partitions and Active Directory boundaries are defined, while centralized XenDesktop Delivery Controllers govern application and desktop delivery across tenants.

Figure 5 Multi-Tenant Citrix Farms and Sites

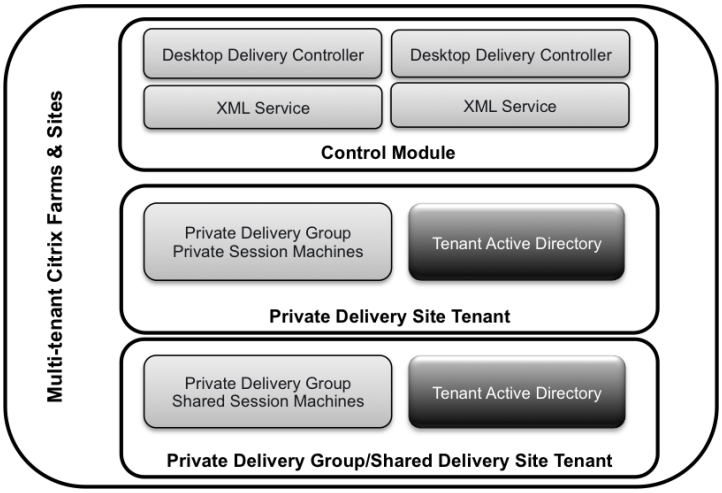

Dashboards and Cloud Platform Management Module

To successfully manage a service provider network, administrators need effective tools that are simple to use and scale efficiently as they add new tenants. Along with Cisco UCS Manager, Citrix App Orchestration and Citrix CloudPortal Services Manager enable a unified view across the entire infrastructure, across multiple datacenters, XenDesktop and XenApp sites, and servers. This end-to-end view gives providers the detailed information and wide spectrum of control necessary to provision applications and quickly and maintain service level agreements for subscribers. Additional tools such as HDX Insight and the Citrix Usage Collector also facilitate ease of management. The Citrix documentation http://www.citrix.com/edocs describes App Orchestration and CloudPortal Services Manager management capabilities in more detail.

Figure 6 Dashboards and Cloud Platform Management

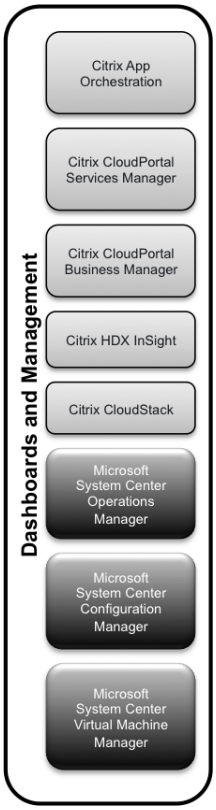

Endpoints and Offices Module

When applications, desktops and data are delivered as a service, the user is the ultimate judge of the endpoint experience. Service providers must deliver a consistent experience across a range of network bandwidths and devices to build and expand the subscriber base. Citrix Receiver and HDX technologies are the strategic components that make this possible.

With Citrix Receiver, CSPs have complete control over security, performance, and user experience with no need to own or manage the physical device or its location. Users simply install Citrix Receiver on their own device to gain access to their desktop and all of their business, web, Software-as-a-Service (SaaS), and native mobile applications.

With the introduction of Citrix's next-generation seamless application capabilities, applications that must execute on the endpoint device can now be presented within a user's cloud-hosted desktop. This capability enables 100% application compatibility within the CSP solution while also providing a smooth transition over time for application migrations from legacy endpoints and datacenters into the CSP hosted datacenter.

Figure 7 Endpoints and Offices Module

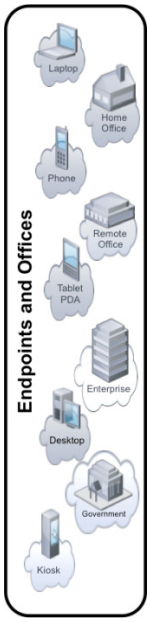

Key Concepts

The reference architecture uses the following key concepts to deliver DaaS in a multi-tenant environment.

Delivery Sites and Delivery Groups

App Orchestration uses Delivery Sites and Delivery Groups to isolate the provisioning of application and desktop services in this reference architecture:

- Delivery Sites. A Delivery Site is the core environment that contains the XenDesktop Delivery Controllers and the SQL Database used to deploy XenApp and XenDesktop services. Delivery Sites provision desktops and applications to users through App Orchestration.
- Delivery Group. A Delivery Group is a container for one or more virtual machines used to deliver applications and desktops to a specific group of users. A Delivery Group is associated with a shared or private Delivery Site. A Delivery Group can be shared among tenants or dedicated to a specific tenant, according to the isolation level of the subscriptions it is hosting.

When you create App Orchestration offerings, you choose a means of isolation for tenants who subscribe to the app or desktop. The isolation level refers to whether the Delivery Controllers and Session Machines used for the offering are shared with other tenants or private to the subscribing tenant.

Multi-Tenancy Isolation Models

Multi-tenancy capabilities provide economies of scale on a single infrastructure while providing the required isolation and data protection. Providers can make trade-offs regarding price and features to meet individual tenant requirements.

The three common multi-tenancy isolation models used in the market today — Shared Delivery Group/Shared Delivery Site, Private Delivery Group/Shared Delivery Site, and Private Delivery Group/Private Delivery Site — differ according to the type of isolation they employ (Figure 8). In a Shared Delivery Group/Shared Delivery Site, both the session machines and the delivery site are shared with other tenants. In a Private Delivery Group/Shared Delivery Site, session machines are private, but the delivery site is shared. In a Private Delivery Group/Private Delivery Site, both the session machines and the delivery site are private, and not available to other tenants. All three multi-tenant approaches can be delivered from the provider's datacenters.

Figure 8 Multi-Tenancy Isolation Models

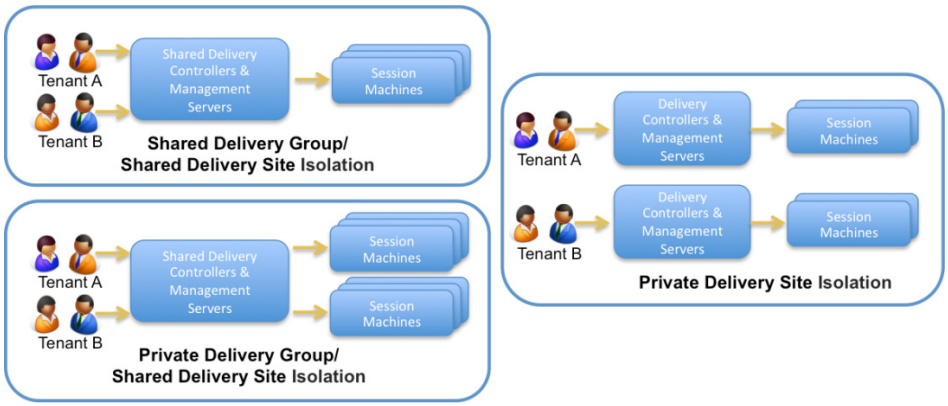

Tenant Domains

A domain contains machines hosting tenants' resources, users, or both. A domain can be either shared or private. A shared resource domain contains machines that that host resources for multiple tenants. In contrast, a private domain contains machines hosting resources for a single tenant.

Configuring Trust

In a previous version of the DaaS reference architecture (based on App Orchestration 2.1), one-way trust relationships were configured to permit access to CSP resources external to a Private Delivery Site tenant's Active Directory (AD) domain.

Figure 9 illustrates the CSP Management domain featured a 1-way, non-transitive trust with Private Delivery Site tenant domains to allow those tenants to access resources within the CSP domain.

Figure 9 Configuring Trust — Private Delivery Site-Isolated Tenants

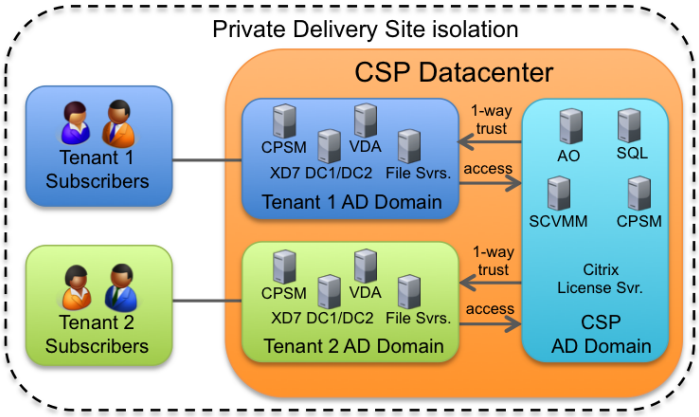

To simplify DaaS configurations with isolated private domains, App Orchestration 2.5 includes a Zero Trust Agent that provides the ability to manage all shared components from a single management domain (without requiring Active Directory one-way trusts to a private tenant domain). The Zero Trust Agent establishes a secure SSL encrypted communication channel to the App Orchestration configuration server and authenticates using SSL certificates. This feature provides a simpler and more secure model to orchestrate resources across tenant domains using SSL-client authentication. For detailed information, go to Configuring SSL for App Orchestration 2.5" and "Deploying the Zero Trust Agent in App Orchestration 2.5.

Virtual Provisioning

Virtualization is a fundamental enabler of an efficient cloud datacenter. As a best practice, virtualizing workloads enables dynamic scale and simplifies management. VMware ESXi 5.5 was used as the hypervisor technology in this CVD. Cisco UCS Director was used to provision virtual machines for the infrastructure, HSD, and VDI services deployed in this architecture.

Application Provisioning

Application provisioning provides a key element to the dynamic assembly capabilities within the system. Dynamic assembly is the process by which separate elements are combined in real-time to present a user with their specific, familiar, and personalized environment of operating system, desktop, application, and personalization settings.

Application virtualization is one of the key enablers of dynamic assembly, separating applications from the underlying OS. This also allows lifecycle management of the application as a discreet object. A further advantage to this separation of OS and application is the ability to deliver and manage a single application image across CSP tenants, personalized for each tenant's SLA through the policies associated with that tenant's Delivery Group partition.

App Orchestration provides the means of allocating applications and desktops to subscribers based on Delivery Groups and Delivery Sites that map to specific Active Directory OUs.

Citrix StoreFront provides users an enterprise app store that aggregates offerings in one place. Each StoreFront user is able to subscribe to their favorite application and desktop resources, which can then follow the user automatically between devices. In this reference architecture, App Orchestration configures private or shared StoreFront catalogs that manage desktop and application offerings for subscribers.

Architecture

Hardware Deployed

The architecture deployed is highly modular. While each Service Provider's environment might vary in its exact configuration, once 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 EMC VNX Storage arrays).

The 2000-user Citrix DaaS solution includes Cisco networking, Cisco UCS and EMC VNX storage, which fits into a single data center rack, including the access layer network switches.

This Cisco Validated Design details the deployment of the 2000-user configuration for a mixed XenDesktop workload featuring the following software:

- Citrix XenDesktop 7.5 Pooled Hosted Shared Desktops with PVS write cache on NFS storage
- Citrix XenDesktop 7.5 Pooled Server-VDI Virtual Desktops with PVS write cache on NFS storage
- Citrix Provisioning Server 7.1
- Citrix User Profile Manager
- Citrix StoreFront 2.1
- Cisco Nexus 1000V Distributed Virtual Switch
- VMware vSphere ESXi 5.5 Hypervisor
- Microsoft Windows Server 2012 R2 virtual machine Operating Systems
- Microsoft SQL Server 2012 SP1

Figure 10 Workload Architecture

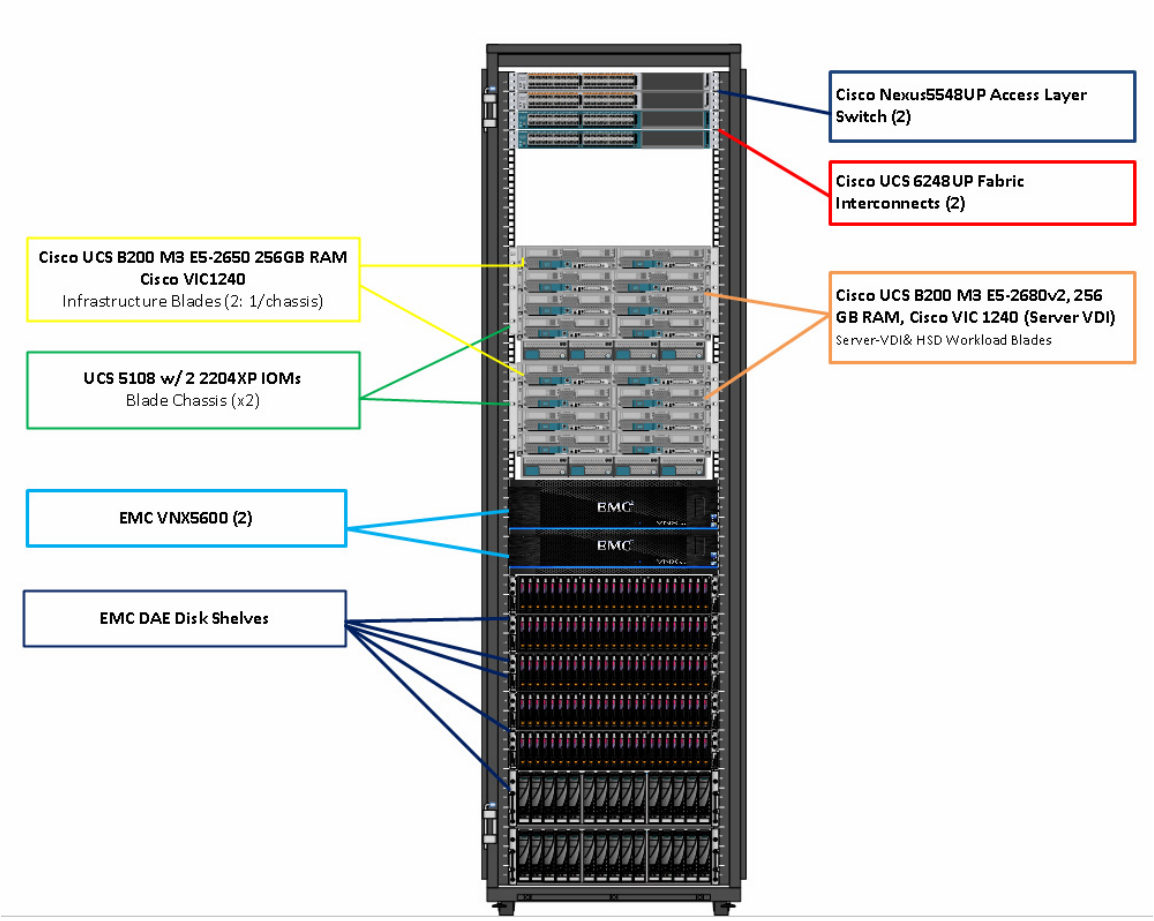

The workload contains the following hardware as shown in Figure 10:

- Two Cisco Nexus 5548UP Layer 2 Access Switches
- Two Cisco UCS 6248UP Series Fabric Interconnects
- Two Cisco UCS 5108 Blade Server Chassis with two 2204XP IO Modules per chassis
- Two Cisco UCS B200 M3 Blade servers with Intel E5-2650v2 processors, 256GB RAM, for Infrastructure of all DaaS components with N+1 server fault tolerance.
- Fourteen Cisco UCS B200 M3 Blade servers with Intel E5-2680v2 processors, 256 GB RAM, and VIC1240 mezzanine cards for the 2000 hosted shared and server VDI Windows Server 2012 server desktop workloads with N+1 server fault tolerance.
- EMC VNX5600 dual controller storage system, 10 disk shelves, 10GE ports for NFS andCIFS connectivity.
- (Not Shown) One Cisco UCS 5108 Blade Server Chassis with 3 Cisco UCS B200 M3 Blade servers with Intel E5-2650 processors, 128 GB RAM, and VIC1240 mezzanine cards for the Login VSI launcher infrastructure

Logical Architecture

The logical architecture of the validated is designed to support 2000 users within two chassis and fourteen blades, which provides physical redundancy for the chassis and blade servers for each workload.

Figure 11 outlines the logical architecture of the test environment.

Figure 11 Logical Architecture Overview

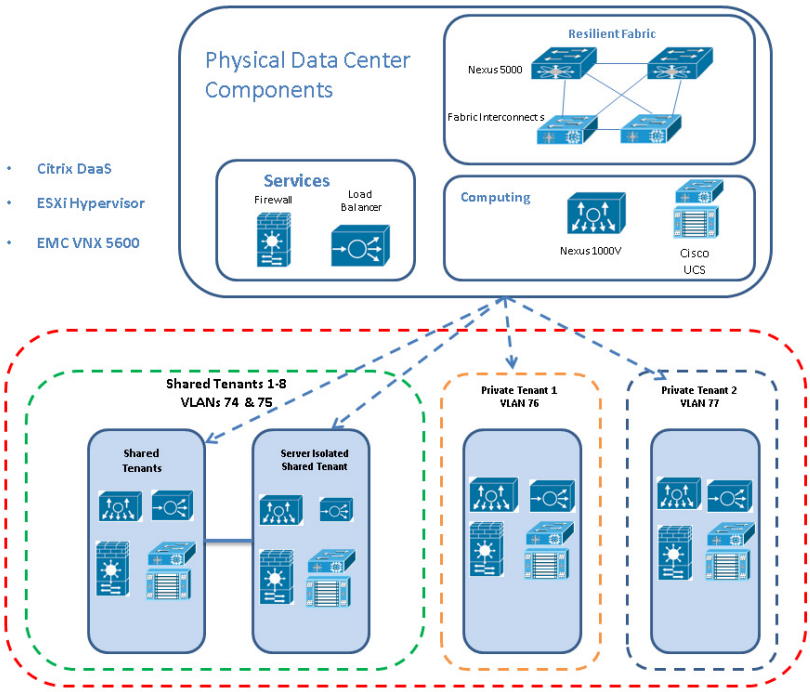

Table 1 outlines all the servers in the configurations

Table 1 Infrastructure Architecture

Server Name	Location	Purpose
CH1-BL1	Physical – Chassis 1	Windows 2012 Datacenter VMs ESXi 5.5 host (Infrastructure Guests)
CH1-BL2,3,4	Physical – Chassis 1	XenDesktop 7.5 RDS ESXi 5.5 Hosts
CH1-BL5,6,7,8	Physical – Chassis 1	XenDesktop 7.5 HVD ESXi 5.5 Host
CH2-BL1	Physical – Chassis 2	Windows 2012 Datacenter VMs ESXi 5.5 host (Infrastructure Guests)
CH2-BL2,3,4	Physical – Chassis 2	XenDesktop 7.5 RDS ESXi 5.5 Hosts
CH2-BL5,6,7,8	Physical – Chassis 2	XenDesktop 7.5 HVD ESXi 5.5 Hosts
Daas-DC01	Virtual – CH1-BL1	Active Directory Domain Controller
DaaS-XDC01	Virtual – CH1-BL1	XenDesktop 7.5 controller
PVS1	Virtual – CH1-BL1	Provisioning Services 7.1 streaming server
VCENTER	Virtual – CH1-BL1	vCenter 5.5 Server
DaaS-SF01	Virtual – CH1-BL1	StoreFront Services server
SQL01	Virtual – CH1-BL1	SQL Server (clustered)
Daas-n1kv	Virtual – CH1-BL1	Nexus 1000-V VSM HA node

CTXLIC	Virtual – CH1-BL1	XenDesktop 7.5 License server
N1KV-VSM-1	Virtual – CH1-BL1	Nexus 1000-V VSM HA primary node
DaaS-DC02	Virtual – CH2-BL1	Active Directory Domain Controller
DaaS-XDC02	Virtual – CH2-BL1	XenDesktop 7.5 controller
PVS2	Virtual – CH2-BL1	Provisioning Services 7.1 streaming server
DaaS-SF2	Virtual – CH2-BL1	StoreFront Services server
SQL02	Virtual – CH2-BL1	SQL Server (clustered)
N1KV-VSM-1	Virtual – CH2-BL1	Nexus 1000-V VSM HA backup node

Software Revisions

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

Table 2 Software Revision

Vendor	Product	Version
Cisco	UCS Component Firmware	2.2(2c)
Cisco	UCS Manager	2.2(2c)
Cisco	Nexus 1000V for VSphere	4.2(1)SV2(2.2)
Citrix	XenDesktop	7.5.0.4033
Citrix	Provisioning Services	7.1.0.4022
Citrix	StoreFront Services	2.1.0.17
Citrix	Netscaler Appliance	10.1 Build 128.8
Citrix	App Orchestrator	2.5
Citrix	Cloud Portal Services Manager	11.1
Citrix	Netscaler	10.1 Build 128.8
VMware	vCenter	5.5.0 Build 1476327
VMware	vSphere ESXi 5.5	5.5.0 Build 1746018
EMC	VAAI Plugin	1.0-11
EMC	Power Path for VMware	5.9 SP1 Build 011
EMC	VNX Block Operating System	05.33.000.5.051
EMC	VNX File Operating System	8.1.2-51

Configuration Guidelines

The 2000 User Citrix DaaS Solution Provider design described in this document provides details for configuring a fully redundant, highly-available configuration. Configuration guidelines are provided that refer to which redundant component is being configured with each step, whether that be A or B. For example Nexus A and Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are configured similarly.

Networking and Multi-tenancy

Cisco VMDC Architecture for Scalable, Secure and Resilient Infrastructure

The Cisco VMDC is a tested and validated reference architecture for the Cisco Unified Data Center. It provides a set of guidelines and best practices for the creation and deployment of a scalable, secure, and resilient infrastructure in the data center. The Cisco VMDC architecture demonstrates how to bring

together the latest Cisco routing and switching technologies, network services, data center and cloud security, automation, and integrated solutions with those of Cisco's ecosystem of partners to develop a trusted approach to data center transformation. Specific benefits include:

- Demonstrated solutions to critical technology-related problems in evolving IT infrastructure: Provides support for cloud computing, applications, desktop virtualization, consolidation and virtualization, and business continuance
- Reduced time to deployment: Provides best-practice recommendations based on a fully tested and validated architecture, helping enable technology adoption and rapid deployment
- Reduced risk: Enables enterprises and service providers to deploy new architectures and technologies with confidence
- Increased flexibility: Enables rapid, on-demand, workload deployment in a multitenant environment using a comprehensive automation framework with portal-based resource provisioning and management capabilities
- Improved operating efficiency: Integrates automation with a multitenant pool of computing, networking, and storage resources to improve asset use, reduce operation overhead, and mitigate operation configuration errors

The Cisco VMDC architecture, consisting of the Cisco Unified Data Center and Cisco Data Center Interconnect (DCI) together with other architectural components such as infrastructure abstraction, orchestration and automation, assurance, and integrated services and applications, as shown in Figure 12, provide comprehensive guidelines for deployment of cloud infrastructure and services at multiple levels.

Figure 12 Cisco VMDC Components

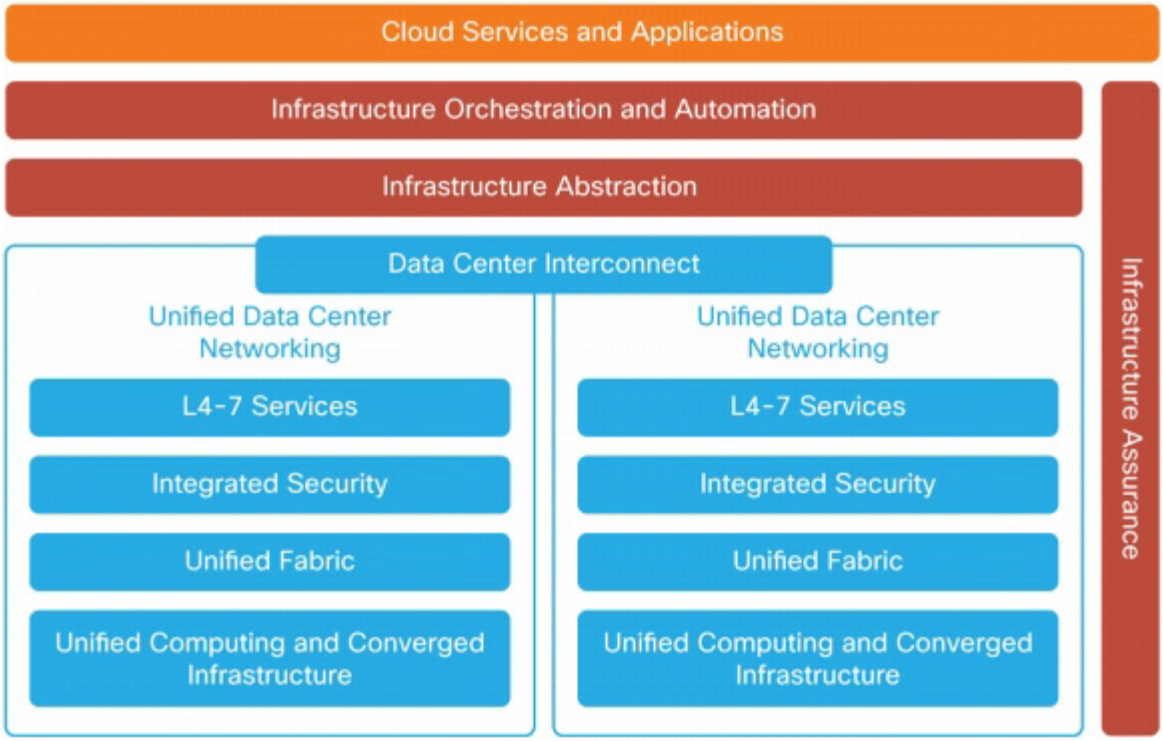

Cisco VMDC Architecture

This section describes the primary components of the Cisco VMDC architecture:

- Modular building blocks
- Resilient network fabric
- Multilayer end-to-end security
- Intelligent network-based services
- Efficient data center interconnection for business continuity
- Comprehensive cloud service management
- Integrated applications and services

Modular Building Blocks: Integrated Systems, Points of Delivery (PoDs), and Data Centers

The Cisco VMDC architecture builds on existing integrated system components such as VSPEX (using the Cisco UCS platform and EMC storage) and Vblock ™ Infrastructure Packages (using the Cisco UCS platform and EMC storage, integrated through the Virtual Computing Environment (VCE) coalition). In both cases, VMware software is used for virtualization; however, other hypervisors such as Microsoft Hyper-V, Red Hat KVM, and Citrix XenServer can also be used in designing basic integrated system building blocks. The Cisco VMDC architecture supports all variations.

Another modular building block of the Cisco VMDC architecture is the point of delivery (PoD), as shown in Figure 13, which contains standardized computing, storage, and networking components, as predefined integrated FlexPod or Vblock systems, and customized computing and storage systems, as needed by the deployment. The PoD concept and architecture is not limited to Cisco UCS and can be modified and extended to include other computing and storage stacks.

Figure 13 Point of Delivery

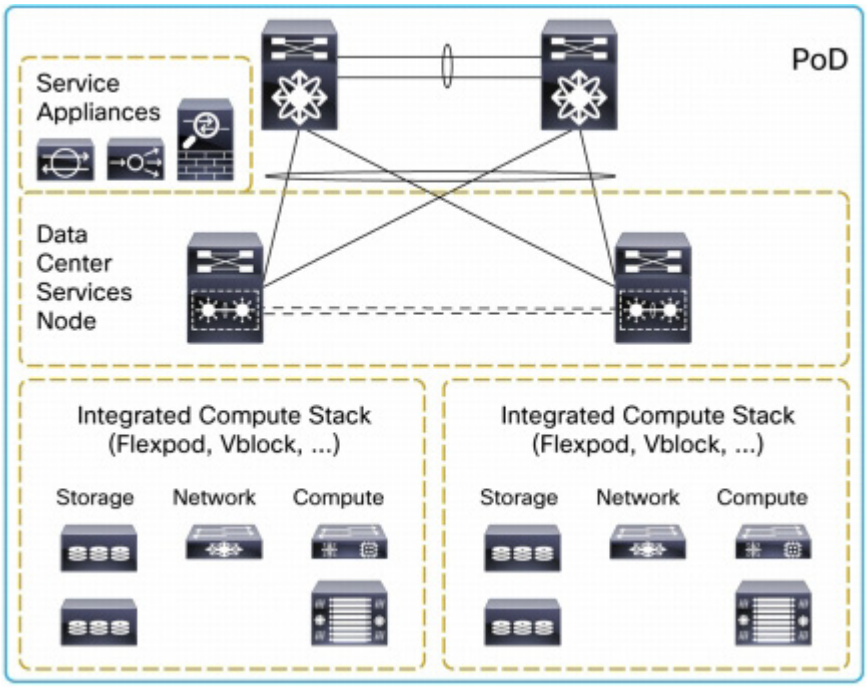

PoDs can contain localized network-based services such as firewalls and load balancers. PoDs can be created to house different application or service loads. Each PoD can be mapped to a class of applications, such as IT applications (print and file services, for example) or dedicated application servers (for SAP, for example). For orchestration and automation, a PoD represents a shared resource pool in a common administrative domain and can be autodiscovered by the operation software and configured for the specific defined service profile. Factors that go into PoD sizing and definition include:

- Storage capacity: Balance of computing needs to storage I/O operations per second (IOPS)
- Computing capacity: VMware vSphere cluster sizing and VMware vCenter domain management considerations
- Layer 2 scale considerations: MAC address and Address Resolution Protocol (ARP) capacity and VLAN scaling budgets
- Service insertion requirements: Scale and performance of network-based services such as load balancing and firewalling; a PoD may have its own dedicated set of load balancing and firewall service engines
- Application requirements: Requirements of specific applications; for example, some PoDs can be dedicated to VDI, and others can be dedicated to media applications
- Management requirements: Some PoDs can be dedicated to specific functions such as management

Multiple PoDs can be connected together to make up a data center, as shown in Figure 14. To design the optimal network for PoD connectivity, organizations need to consider the data center scale, network resiliency and tolerance to failure, security through traffic separation and protection, consumer and operator access control, and traffic characteristics and requirements of the applications and services that are hosted in the PoDs. The Cisco VMDC architecture provides all necessary details for these design considerations in a modular fashion, hence enabling creation of data centers that can grow and expand with ease.

Figure 14 Connecting Multiple PoDs

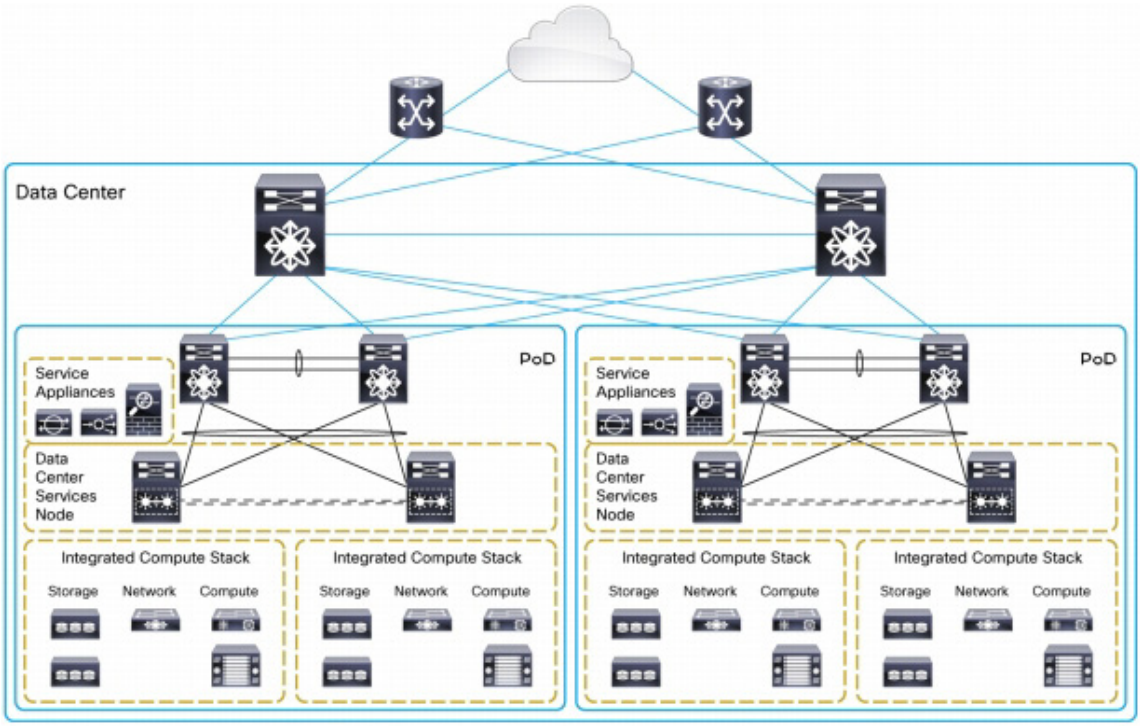

In addition to the basic connectivity within a data center, which may span multiple buildings in a campus or metropolitan network, the Cisco VMDC architecture specifies optimal connectivity between multiple sites that can be separated by farther distances. These specifications are part of the Cisco DCI module of Cisco VMDC. Use of resilient, secure, and efficient DCI methods enables application mobility at scale and the capacity to provide disaster recovery and business continuance.

Because the Cisco VMDC architecture is well documented, we did not include deployment steps for this flexible, modular system in this paper.

For more information on Cisco's VMDC architecture, go to:

http://www.cisco.com/c/en/us/solutions/collateral/enterprise/data-center-designs-cloud-computing/white_paper_c11-714729.html

VLAN

The VLAN configuration recommended for the environment includes a total of eight VLANs as outlined in the Table 3.

Table 3 VLAN Configuration

VLAN Name	VLAN ID	Use
Default	1	Default VLAN
MGMT-OB	70	Out of Band Management Network

MGMT-IB	71	In Band Management Network
Storage Traffic	72	IP Storage VLAN for NFS and CIFS
vMotion	73	vMotion
Shared Tenant Clients	74	VLAN for client machines in shared tenants
Server Isolated Clients	75	VLAN for clients in hardware isolated environment.
Private Tenant 1	76	VLAN for all of Private Tenant 1 (Infrastructure and Clients)
Private Tenant 2	77	VLAN for all of Private Tenant 1 (Infrastructure and Clients)
PXE for UCS-D	79	Native VLAN

VMware Clusters

We utilized four VMware Clusters in one data center to support the solution and testing environment:

- SP DaaS Service Provider DaaS
- XenDesktop RDS Clusters (Windows Server 2012 R2 hosted shared desktops)
- XenDesktop Server Virtual Desktop Cluster (Windows Server 2012 R2 ServerVDI)
- Infrastructure Cluster (vCenter, Active Directory, DNS, DHCP, SQL Clusters, XenDesktop Controllers, Provisioning Servers, and Nexus 1000V Virtual Switch Manager appliances, etc.)

Figure 15 vCenter Data Centers and Clusters Deployed

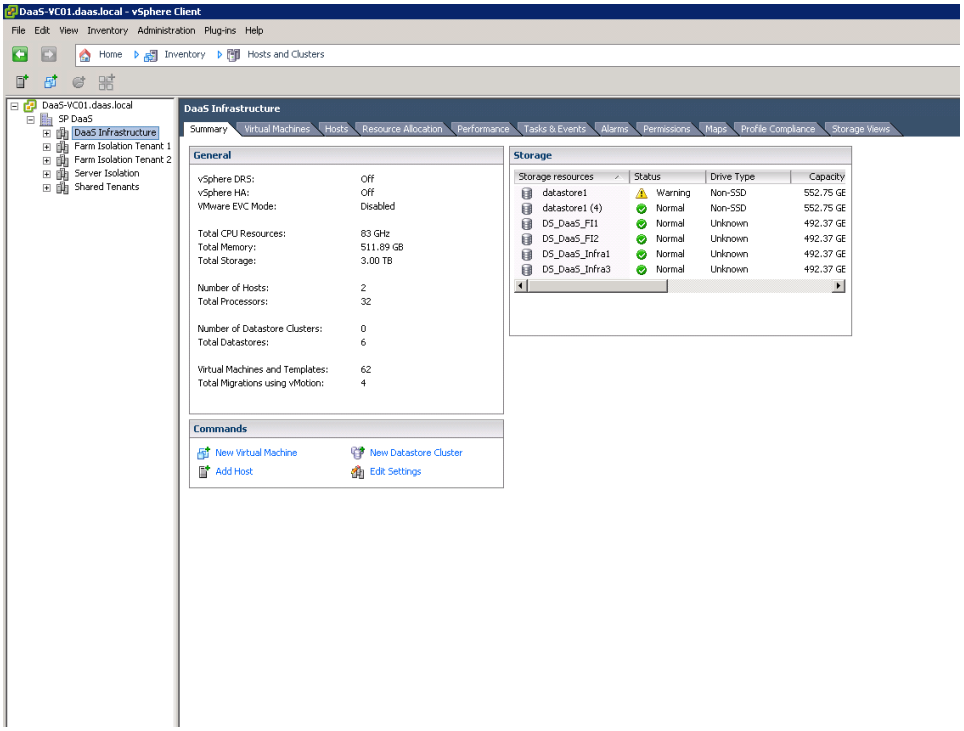

Infrastructure Components

This section describes the infrastructure components used in the solution outlined in this study.

Cisco Unified Computing System (UCS)

Cisco UCS is a set of pre-integrated data center components that comprises blade servers, adapters, fabric interconnects, and extenders that are integrated under a common embedded management system. This approach results in far fewer system components and much better manageability, operational efficiencies, and flexibility than comparable data center platforms.

Cisco Unified Computing System Components

Cisco UCS components are shown in Figure 16.

Figure 16 Cisco Unified Computing System Components

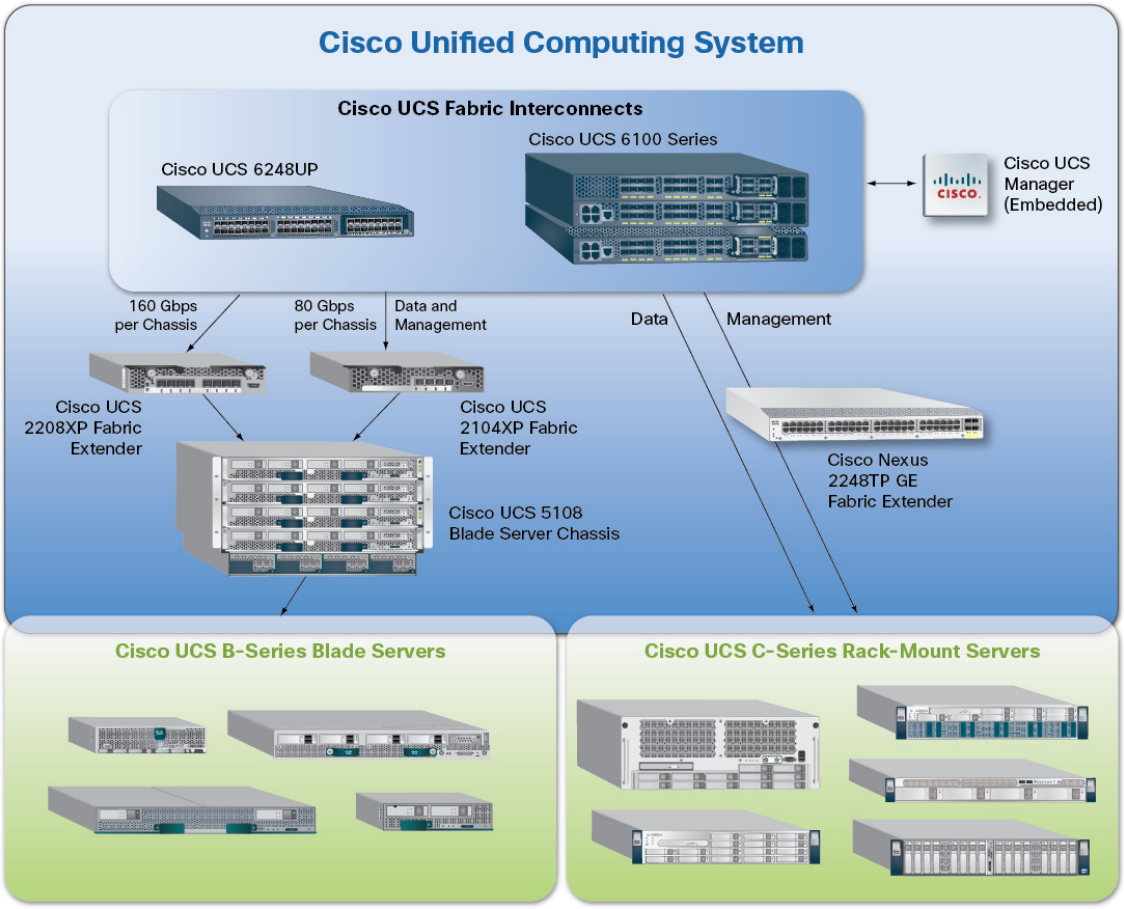

The Cisco UCS is designed from the ground up to be programmable and self integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards, even the number and type of I/O interfaces 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, associate a model's service profile with hardware resources and 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.

Fabric Interconnect

Cisco UCS Fabric Interconnects create a unified network fabric throughout the Cisco UCS. They provide uniform access to both networks and storage, eliminating the barriers to deploying a fully virtualized environment based on a flexible, programmable pool of resources.

Cisco Fabric Interconnects comprise a family of line-rate, low-latency, lossless 10-GE, Cisco Data Center Ethernet, and FCoE interconnect switches. Based on the same switching technology as the Cisco Nexus 5000 Series, Cisco UCS 6000 Series Fabric Interconnects provide the additional features and management capabilities that make them the central nervous system of Cisco UCS.

The Cisco UCS Manager software runs inside the Cisco UCS Fabric Interconnects. The Cisco UCS 6000 Series Fabric Interconnects expand the UCS networking portfolio and offer higher capacity, higher port density, and lower power consumption. These interconnects provide the management and communication backbone for theCisco UCS B-Series Blades and Cisco UCS Blade Server Chassis.

All chassis and all blades that are attached to the Fabric Interconnects are part of a single, highly available management domain. By supporting unified fabric, the Cisco UCS 6200 Series provides the flexibility to support LAN and SAN connectivity for all blades within its domain right at configuration time. Typically deployed in redundant pairs, the Cisco UCS Fabric Interconnect provides uniform access to both networks and storage, facilitating a fully virtualized environment.

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

Cisco UCS 6248UP 48-Port Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a 1 RU, 10-GE, Cisco Data Center Ethernet, FCoE interconnect providing more than 1Tbps throughput with low latency. It has 32 fixed ports of Fibre Channel, 10-GE, Cisco Data Center Ethernet, and FCoE SFP+ ports.

One expansion module slot can be up to sixteen additional ports of Fibre Channel, 10-GE, Cisco Data Center Ethernet, and FCoE SFP+.

Cisco UCS 6248UP 48-Port Fabric Interconnects were used in this study.

Cisco UCS 2200 Series IO Module

The Cisco UCS 2100/2200 Series FEX multiplexes and forwards all traffic from blade servers in a chassis to a parent Cisco UCS Fabric Interconnect over from 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 Extender are ASIC processors developed by Cisco that multiplex all traffic.

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

Cisco UCS 2104 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 via SFP+ sockets for both throughput and redundancy. It has 4 ports connecting up the fabric interconnect.

Cisco UCS 2208 has thirty-two 10GBASE-KR connections to the blade chassis midplane, 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 4x10-Gbps unified fabric-based networks via SFP+ sockets for both throughput and redundancy. It has 8 ports connecting up the fabric interconnect.

Cisco UCS 2208 fabric extenders were utilized in this study.

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 UCS 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

Cisco UCS B200 M3 Blade Server

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 v2 product family, with up to 768 GB of RAM (using 32GB 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 UCS 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.

In addition, customers who initially purchased Cisco UCS B200M3 blade servers with Intel E5-2600 series processors, can field upgrade their blades to the second generation E5-2600 processors, providing increased processor capacity and providing investment protection

Figure 17 Cisco UCS B200 M3 Server

Cisco UCS VIC1240 Converged Network Adapter

A Cisco® innovation, the Cisco UCS Virtual Interface Card (VIC) 1240 (Figure 1) is a 4-port 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the 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). In addition, the Cisco UCS VIC 1240 supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

Figure 18 Cisco UCS VIC 1240 Converged Network Adapter

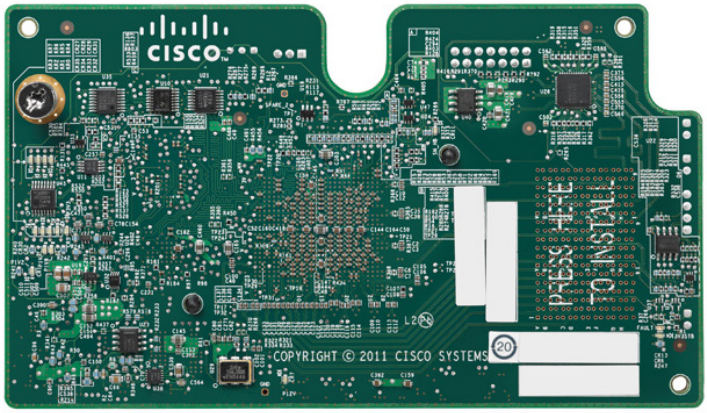

Figure 19 The Cisco UCS VIC1240 Virtual Interface Cards Are Deployed in the Cicso UCS B-Series B200 M3 Blade Servers

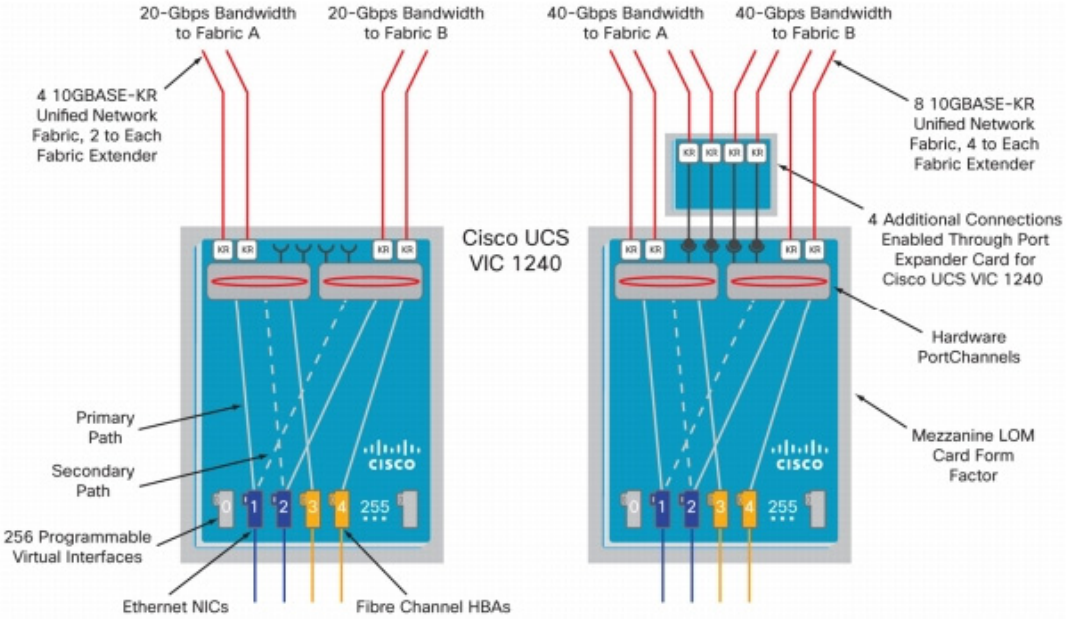

Cisco UCS Director

Cisco UCS Director improves consistency, efficiency, and speed within your organization. It accomplishes this by replacing time-consuming, manual provisioning and de-provisioning of data center resources with automated workflows. Cisco UCS Director reduces delivery time from weeks to minutes.

Superior IaaS Management Capabilities

Unified infrastructure provisioning and management delivers superior infrastructure as a service (IaaS) management capabilities with the following benefits:

- Out-of-box task library that spans Cisco and third-party hardware solutions to build your infrastructure in minutes
- Reduction of data center complexity by replacing manual provisioning and de-provisioning tasks with workflows that span computing, network, storage, and virtualization functions
- Lowering of capital expenses with real-time monitoring, dynamic load balancing, and optimum resource usage

Multi-Hypervisor Solution

With support for VMware, Microsoft Hyper-V, and Red Hat KVM hypervisors, Cisco UCS Director's task library supports creation, manipulation, and editing of virtual machines, hosts, and virtual networks. Broad OS guest support facilitates cloning of Windows and Linux virtual machines, and also monitors host and virtual machine memory and resource use. It offers:

- Support for NetApp Virtual Storage Console on Clustered Data ONTAP to deliver fast and efficient cloning of virtual machines with low storage use
- Support for Microsoft System Center Virtual Machine Manager networking models and Hyper-V infrastructure management
- Support for Cisco Nexus 1000V and Citrix network devices within Microsoft System Center Private Cloud

Multivendor Solution

Because your data center comprises diverse technologies, Cisco UCS Director delivers heterogeneous infrastructure management. The benefits include:

- Support for Hewlett Packard Onboard Administrator to install bare-metal blades and manage them using Cisco UCS Director's task library
- Extensive enhancements for VMware, VCE, and EMC solution components
- Integration of third-party solutions into the Cisco UCS Director management platform with a publicly available software development kit

EMC Storage Architecture Design

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.

EMC VNX5600 used in this solution provides comprehensive storage architecture for hosting all virtual desktop components listed below on a unified storage platform.

- ESXi OS is stored on an FC LUN from which each vSphere host is booted. The boot from SAN design allows UCS service profiles to be portable from one blade to another when the blades do not use local disks.
- PVS vDisk is hosted on a VNX CIFS share to provide central management of vDisk by eliminating duplicated copies of the same vDisk image.
- PVS write cache and infrastructure VMs hosted on VNX NFS datastores simplifies VM storage provisioning.
- User profiles defined by Citrix User Profile Management (UPM) and user home directories both reside on VNX CIFS shares that can leverage VNX deduplication, compression, and data protection.

Enhancements in XenDesktop 7.5

Citrix XenDesktop 7.5 includes significant enhancements to help customers deliver Windows apps and desktops as mobile services while addressing management complexity and associated costs. Enhancements in this release include:

- Unified product architecture for XenApp and XenDesktop—the FlexCast Management Architecture (FMA). This release supplies a single set of administrative interfaces to deliver both hosted-shared applications (RDS) and complete virtual desktops (VDI). Unlike earlier releases that separately provisioned Citrix XenApp and XenDesktop farms, the XenDesktop 7.5 release allows administrators to deploy a single infrastructure and use a consistent set of tools to manage mixed application and desktop workloads.

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

FlexCast Management Architecture (FMA) Technology

In Citrix XenDesktop 7.5, FlexCast Management Architecture (FMA) technology is responsible for delivering and managing hosted-shared RDS apps and complete VDI desktops. By using Citrix Receiver with XenDesktop 7.5, users have access to a device-native experience on a variety of endpoints, including Windows, Mac, Linux, iOS, Android, ChromeOS, and Blackberry devices.

Figure 20 Key components in a typical deployment using FlexCast technology

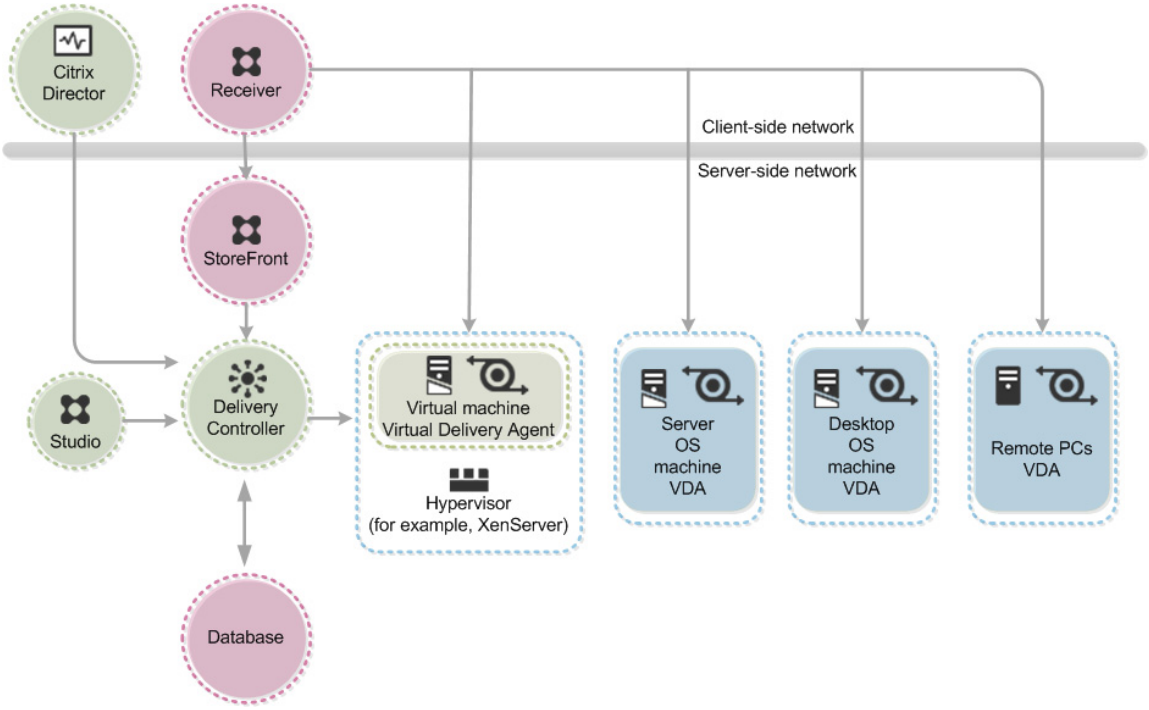

The diagram above shows the key components in a typical XenDesktop deployment:

- Director — Director is a web-based tool that enables IT support and help desk teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users.
- Receiver — Installed on user devices, Citrix Receiver provides users with quick, secure, self-service access to documents, applications, and desktops. Receiver provides on-demand access to Windows, Web, and Software as a Service (SaaS) applications.
- StoreFront — StoreFront authenticates users to sites hosting resources and manages stores of desktops and applications that users can access.
- Studio — Studio is the management console to set up the environment, create workloads to host applications and desktops, and assign applications and desktops to users.
- License server —At least one license server is needed to store and manage license files.
- Delivery Controller — Installed on servers in the data center, the Delivery Controller consists of services that communicate with the hypervisor to distribute applications and desktops, authenticate and manage user access, and broker connections between users and their virtual desktops and applications. The Controller manages the desktop state, starting and stopping them based on demand and administrative configuration. Each XenDesktop site has one or more Delivery Controllers.
- Hypervisor —Hypervisor technology is used to provide an enterprise-class virtual machine infrastructure that is the foundation for delivering virtual applications and desktops. Citrix XenDesktop is hypervisor-agnostic and can be deployed with Citrix XenServer, Microsoft Hyper-V, or VMware vSphere. For this CVD, the hypervisor used was VMware ESXi 5.5.
- Virtual Delivery Agent (VDA) — Installed on server or workstation operating systems, the VDA enables connections for desktops and apps. For Remote PC Access, install the VDA on the office PC.

- Machine Creation Services (MCS) — A collection of services that work together to create virtual servers and desktops from a master image on demand, optimizing storage utilization and providing a pristine virtual machine to users every time they log on. Machine Creation Services is fully integrated and administrated in Citrix Studio.
- Windows Server OS machines — These are VMs or physical machines based on Windows Server operating system used for delivering applications or hosted shared desktops to users.
- Desktop OS machines — These are VMs or physical machines based on Windows Desktop operating system used for delivering personalized desktops to users, or applications from desktop operating systems.
- Remote PC Access — User devices that are included on a whitelist, enabling users to access resources on their office PCs remotely, from any device running Citrix Receiver.

In addition, 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. PVS supports image management for both RDS and VDI-based machines, including support for image snapshots and rollbacks.

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. Real-time multimedia transcoding improves the delivery of Windows Media content (even in extreme network conditions). HDX technology offers a rich multimedia experience and optimized performance for voice and video collaborations.
- 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. In this release, HDX 3D Pro has been upgraded to support Windows 8.

HSD and VDI 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 7.5 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 21 XenDesktop 7.5

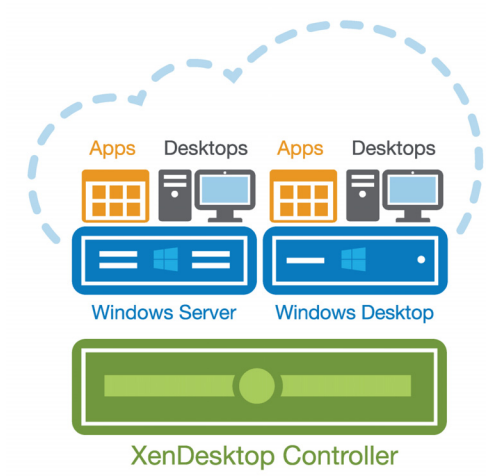

As depicted above, the XenDesktop 7.5 release allows administrators to create a single infrastructure that supports multiple modes of service delivery, including:

- Application Virtualization and Hosting (via XenApp). 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 CVD focuses on delivering a mixed workload consisting of hosted shared desktops (HSD based on RDS) and hosted virtual desktops (VDI).

Citrix Provisioning Services

Citrix XenDesktop 7.5 can be deployed with or without Citrix Provisioning Services (PVS). The advantage of using Citrix PVS is that it allows virtual machines to be provisioned and re-provisioned in real-time from a single shared-disk image. In this way administrators can completely eliminate the need to manage and patch individual systems and reduce the number of disk images that they manage, even as the number of machines continues to grow, simultaneously providing the efficiencies of a centralized management with the benefits of distributed processing.

The Provisioning Services solution's infrastructure is based on software-streaming technology. After installing and configuring Provisioning Services components, a single shared disk image (vDisk) is created from a device's hard drive by taking a snapshot of the OS and application image, and then storing that image as a vDisk file on the network. A device that is used during the vDisk creation process is the Master target device. Devices or virtual machines that use the created vDisks are called target devices.

When a target device is turned on, it is set to boot from the network and to communicate with a Provisioning Server. Unlike thin-client technology, processing takes place on the target device (Step 1).

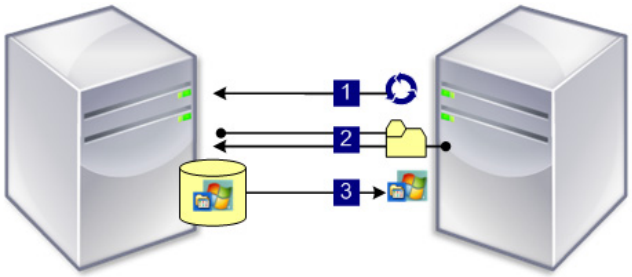

The target device downloads the boot file from a Provisioning Server (Step 2) and boots. Based on the boot configuration settings, the appropriate vDisk is mounted on the Provisioning Server (Step 3). The vDisk software is then streamed to the target device as needed, appearing as a regular hard drive to the system.

Instead of immediately pulling all the vDisk contents down to the target device (as with traditional imaging solutions), the data is brought across the network in real-time as needed. This approach allows a target device to get a completely new operating system and set of software in the time it takes to reboot. This approach dramatically decreases the amount of network bandwidth required and making it possible to support a larger number of target devices on a network without impacting performance

Citrix PVS can create desktops as Pooled or Private:

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

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

Locating the PVS Write Cache

When considering a PVS deployment, there are some design decisions that need to be made regarding the write cache for the target 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 hard drive.** Write cache exists as a file in NTFS format, located on the target-device's hard drive. This option frees up the Provisioning Server since it does not have to process write requests and does not have the finite limitation of RAM.
- **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 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 in device RAM.** Write cache can exist as a temporary file in the target device's RAM. This provides the fastest method of disk access since memory access is always faster than disk access.

- **Cache in device RAM with overflow on hard disk.** This method uses VHDX differencing format and is only available for Windows 7 and Server 2008 R2 and later. When RAM is zero, the target device write cache is only written to the local disk. When RAM is not zero, the target device write cache is written to RAM first. When RAM is full, the least recently used block of data is written to the local differencing disk to accommodate newer data on RAM. The amount of RAM specified is the non-paged kernel memory that the target device will consume.
- **Cache on a server.** Write cache can exist as a temporary file on a Provisioning Server. In this configuration, all writes are handled by the Provisioning Server, which can increase disk I/O and network traffic. For additional security, the Provisioning Server can be configured to encrypt write cache files. Since the write-cache file persists on the hard drive between reboots, encrypted data provides data protection in the event a hard drive is stolen.
- **Cache on server persisted.** This cache option allows for the saved changes between reboots. Using this option, a rebooted 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 this method of caching, 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.

In this CVD, PVS 7.1 was used to manage Pooled Desktops with cache on device storage for each virtual machine. This design enables good scalability to many thousands of desktops. Provisioning Server 7.1 was used for Active Directory machine account creation and management as well as for streaming the shared disk to the hypervisor hosts.

Citrix App Orchestration

Citrix App Orchestration allows CSPs to orchestrate and automate the delivery of applications and desktops in multi-tenant environments and across multiple products, sites, and datacenters. With App Orchestration, hosted service providers can:

- Manage XenApp and XenDesktop across multiple locations, including multiple datacenters in multiple versions, sites or farms, Active Directory domains, and datacenters
- Provide consistent configuration across global deployments spanning multiple delivery sites, eliminating configuration drift and issues
- Define tenant and user affinity to deliver offerings to primary and backup locations, for optimum continuity and fault tolerance
- Provision desktops and applications on any supported hypervisor. App Orchestration can incorporate externally provisioned VMs (e.g., provisioning via PVS as in this reference architecture).

App Orchestration 2.5 works with XenDesktop 7.5 to automate deployment of machine catalogs, Delivery Sites, and Delivery Groups for delivering applications and desktops (known as “offerings”) to users. Offerings are containers that help CSPs define a set of apps, desktops, and resources. They are designed so that tenant users can select them as needed from an application storefront.

App Orchestration provides a configuration system that enables a multi-tenant data model with flexible isolation concepts at its core (permitting separate isolation models for each service). It also features a simple user interface and automated workflows that control XenDesktop, Active Directory, and other components.

For detailed information about planning, installing, and configuring Citrix App Orchestration, see the App Orchestration documentation
http://support.citrix.com/proddocs/topic/app-orchestration/cao-app-orchestration-25-landing.html

Citrix CloudPortal Services Manager

Citrix CloudPortal Services Manager (CPSM) is a self-service portal that helps providers manage the delivery of cloud services and customer offerings. It drives App Orchestration operation by associating desired states with tenants and allowing services to be provisioned to users. CPSM provides out-of-the-box support for Desktop-as-a-Service and Windows applications (powered by Citrix XenApp and Citrix XenDesktop), as well as popular business applications and services like Microsoft Exchange, Office, SharePoint, Lync, web and data hosting, and virtualization service management. Customers and sub-customers (i.e., such as a reseller's customers) that lack IT expertise can add or change services, view reports, manage users, and perform day-to-day administration tasks through the self-service interface.

Figure 22 Citrix CloudPortal Services Manager Provisioning Capabilities

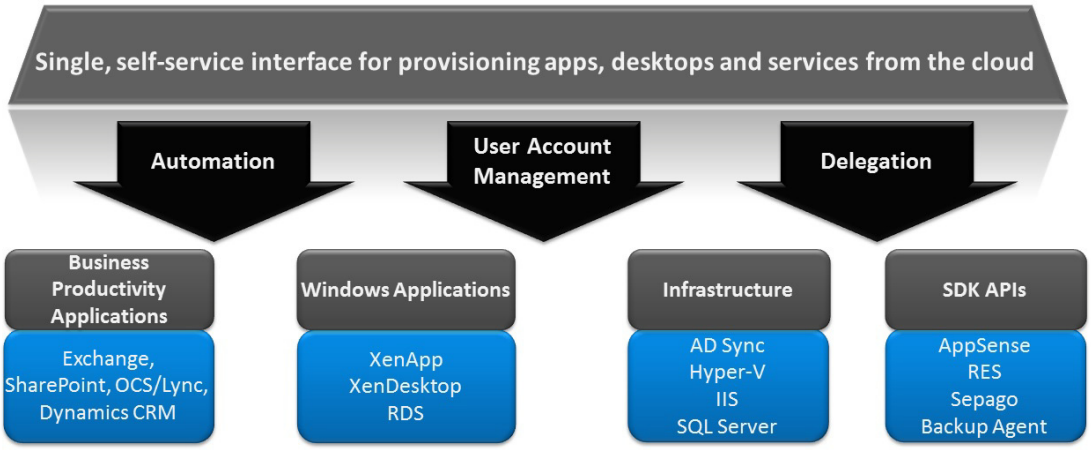

CloudPortal Services Manager provides the following features to simplify and streamline application service provisioning within the CSP reference architecture:

- **Secure delegation of administrative tasks.** Day-to-day administrative tasks, such as creating users, resetting passwords, and provisioning applications and services, are delegated to the customer (or sub-customer in reseller situations) for streamlined management, reduced support costs, and improved quality of service and response time.
- **Simplified user interface.** The easy-to-use self-service web portal interface enables helpdesk staff and customers to manage their application and service offerings without requiring intense training or expensive skillsets.
- **Consolidated management of multiple customer environments.** Multiple customers in a multi-tenant infrastructure can be managed from a single web-based interface, simplifying administration and enabling faster response times.
- **Wizard-driven interface for adding new customers and users.** Adding new users and customers is easy and simple — required information fields are displayed in a familiar, web-based form. The User Copy feature allows an existing user's profile and services to be replicated to a new user for even faster user creation. Multiple users can also be imported from a simple Microsoft Excel spreadsheet, allowing customers to get started quickly and efficiently.
- **Customer resource and limit configuration.** Establishing limit configurations prevents customers from overprovisioning services.

- **Reporting.** By tracking and monitoring the environment, CPSM supports the creation of customized reports for usage and billing.

Solution Architecture

Citrix DaaS Design Fundamentals

DaaS solutions can provide the "anytime, anywhere, any device" access to Windows desktops and applications that workers need for optimal efficiency. For subscribers, the Cisco and Citrix architecture brings a native user experience across a range of endpoints—smartphones, tablets, laptops, PCs, and Macs—even on low-bandwidth mobile networks. From the provider's perspective, implementing the Cisco and Citrix architecture on Cisco UDC platforms enables cost-effective services at cloud-scale across multiple tenants, infrastructure resources, and datacenters.

For customers with stringent security requirements, providers can deploy services using dedicated rather than shared resources. Sharing resources (and using session isolation) enables a lower cost model for tenants with less rigorous security needs.

Isolation Models

Three common multi-tenancy isolation models are used in the market today — Shared Delivery Group/Shared Delivery Site, Private Delivery Group/Shared Delivery Site, and Private Delivery Group/Private Delivery Site. The models differ according to the type of isolation they employ.

Shared Delivery Group/Shared Delivery Site Isolation Model

With the Shared Delivery Group/Shared Delivery Site isolation model, tenants share a single site infrastructure and session host, but each tenant's applications and desktops run within an isolated session on the same virtual machine. This approach is not recommended from a best practice or security perspective, but it is a common model in use for smaller providers today, particularly for those CSPs offering basic, standard desktop services where cost — not security — is the most significant business concern. This CVD does not demonstrate the deployment of this type of isolation model.

Figure 23 Multi-Tenancy: Shared Delivery Group isolation

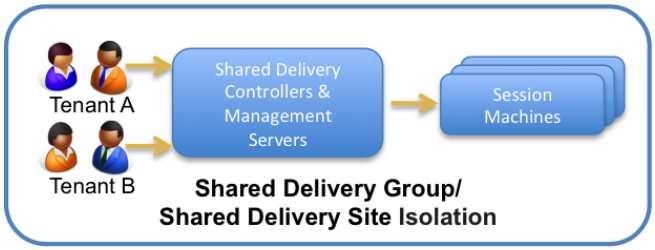

Key characteristics of this model include:

- Users from multiple tenants have isolated sessions on a shared virtual machine (called a Session Machine). This requires appropriate lockdown of Session Machines to minimize the possibility of a user on one tenant negatively affecting users of another tenant. However, there is still a chance that a user can compromise a server (thus affecting another tenant's users).

- User performance guarantees can be established by using the CPU Utilization Management feature.
- A separate web interface site can provide custom branding for each tenant. In addition, Microsoft Windows and Citrix policies in Active Directory can provide a highly customized experience to users (for example, wallpaper, theme, Citrix HDX settings, and so on).
- This method of multi-tenancy is extremely cost-effective because a CSP can spread infrastructure costs across multiple tenants.

Private Delivery Group/Shared Delivery Site Isolation Model

The Private Delivery Group/Shared Delivery Site isolation model provides isolation at the virtual machine layer. Tenants share a single XenDesktop Delivery Controller management network (including a shared XenDesktop site and infrastructure components). Session Machines, on the other hand, are connected to the tenant's private management network, supplying isolation through tenant-specific virtual machines. More and more CSPs are moving to this method. Although it does not provide the strict security of the Private Delivery Site isolation model (which is described next), for many tenants this model provides arguably the most optimal blend of isolation, performance, customization, self-service administration, and cost — a combination that translates into a very attractive offering.

Figure 24 Multi-Tenancy: Private Delivery Group/Shared Delivery Site Isolation

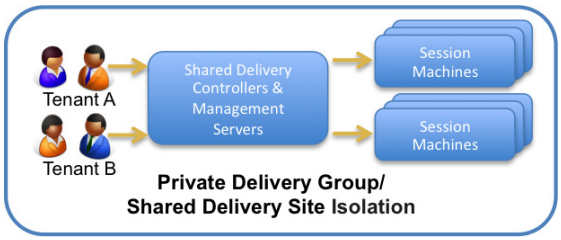

Key characteristics of this model include:

- Each tenant has a dedicated pool of session servers. Delivery Groups and Session Machine Catalogs in App Orchestration simplify deployment. As a best practice, administrators should still always lock down individual Session Machine hosts.
- Because users from one tenant can have sessions only on designated servers, a user cannot negatively impact the performance of another tenant's users. Administrators can further guarantee performance to users by using additional capabilities within XenDesktop.
- In addition to the customization capabilities mentioned in the shared deployment, each tenant can have customized machine images for RDS and VDI workloads.
- CSPs can allow tenants to perform some level of administration for their pool of session hosts or dedicated desktops (e.g., helpdesk activity for viewing which users are logged onto which servers, shadowing a session, or resetting a session).
- Though each tenant has dedicated session hosts or desktops, the costs might not be much higher than that of the shared model. This deployment method offers a blendof multi-tenancy capabilities at a very reasonable cost.

Private Delivery Site Isolation Model

In the Private Delivery Site isolation model, one XenDesktop site (or VDI-in-a-Box grid with dedicated infrastructure) is deployed per tenant. None of the infrastructure components are shared and Session Machines and Delivery Sites are connected to the tenant's private management network. This model is

best suited for tenants with stringent confidentiality and security requirements, such as federal agencies, healthcare, and so on, or those with heavy-duty performance or customization needs. These capabilities come at a cost, but most CSPs typically charge a premium for this type of service. It is understandably less common to see deployments of this nature, but important to understand that the option exists. This option is recommended for those environments where the tightest possible security, regardless of cost, is the primary requirement.

Figure 25 Multi-Tenancy: Private Delivery Site Isolation

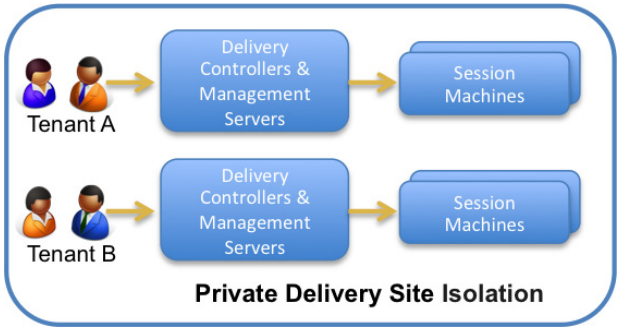

Key characteristics of this model include:

- Tenants are completely isolated including dedicated brokering operations.
- Performance guarantees are similar to the Private Delivery Group/Shared Delivery Site isolation model.
- The customized experience aspects remain the same as that of the Private Delivery Group/Shared Delivery Site isolation model.
- Service providers have the option to allow the tenant to perform a much higher level of self-service administration (for example, help desk activity, managing session hosts, managing applications, etc.).
- The costs are higher for this model because the infrastructure components are not shared between tenants.

Citrix Provisioning Services 7.1

A significant advantage to service delivery via 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, ensuring 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 CVD), on a storage system with which the Provisioning Server can communicate (via 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 XenDesktop site. An administrator may have multiple roles if they belong to more than one group.

Citrix App Orchestration

App Orchestration follows the principle of "Desired State". When a change to an orchestrated deployment occurs, such as creating a Delivery Site or adding a Session Machine to a catalog, the change is saved as a desired configuration in the database. The App Orchestration engine then issues all of the actions required to apply the change. These actions are called workflows, which the administrator can monitor from the App Orchestration management console. The configuration server applies changes asynchronously, allowing multiple operations to occur across different products in the correct sequence and over extended periods of time. If any failures result, they can be corrected and the system will complete the change.

Figure 26 Citrix App Orchestration

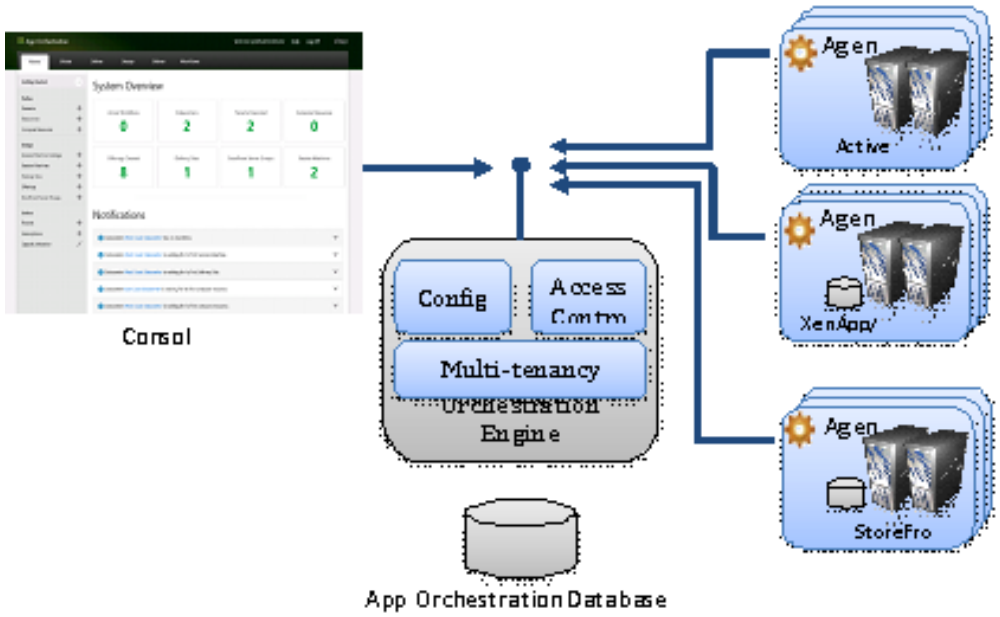

Orchestrating Multi-Tenant Isolation

Citrix App Orchestration simplifies the complex task of multi-tenant isolation by implementing the three primary isolation models discussed earlier in this document. All three multi-tenancy isolation models — Shared Delivery Group, Private Delivery Group/Shared Delivery Site, and Private Delivery Site — can be delivered from the same datacenter. The reference architecture implements these isolation models through App Orchestration Delivery Groups, Active Directory OUs, and Group Policy Objects (GPOs). When administrators create a subscriber offering, they must specify the level of delivery isolation. They must also specify the level of StoreFront isolation for each tenant imported into App Orchestration to define whether the StoreFront server group is shared and whether the tenant uses a private or shared store.

Citrix CloudPortal Services Manager

CloudPortal Services Manager is an easy-to-use web portal that helps service providers manage the delivery of cloud services and offerings to their customers. It provides detailed management for customers, users, services, and applications through a single interface. It supports:

- Channel and reseller enablement. Sub-customers can be nested within other customers to create a hierarchy for reseller and channel use cases. CPSM sites can be branded for a personalized reseller experience.
- Layered services. Services can be enabled at various layers — service provider, customer, and user — for easier and faster management.

Key CPSM Components

The CPSM cloud platform has four primary components, listed in the table below with the corresponding DNS alias and shown in Figure 27.

Component (DNS Alias)	Note
CSPM Web/User Interface (CortexWeb)	This is the Web server that is the user frontend. This server hosts the portal as well as the APIs that are used to access CPSM.
Provisioning Engine (CortexProvisioning)	This is the backend for the CPSM environment. The provisioning engine is workflow rules-based. The engine consists of MSMQ and the queue monitoring process. Rules and actions are stored in the SQL database.
Database (CortexSQL)	This is the SQL database used by the frontend and provisioning engine.
Reporting Services (CortexReports)	This server is used to generate reports regarding usage.

Figure 27 Citrix CloudPortal Services Manager Architecture

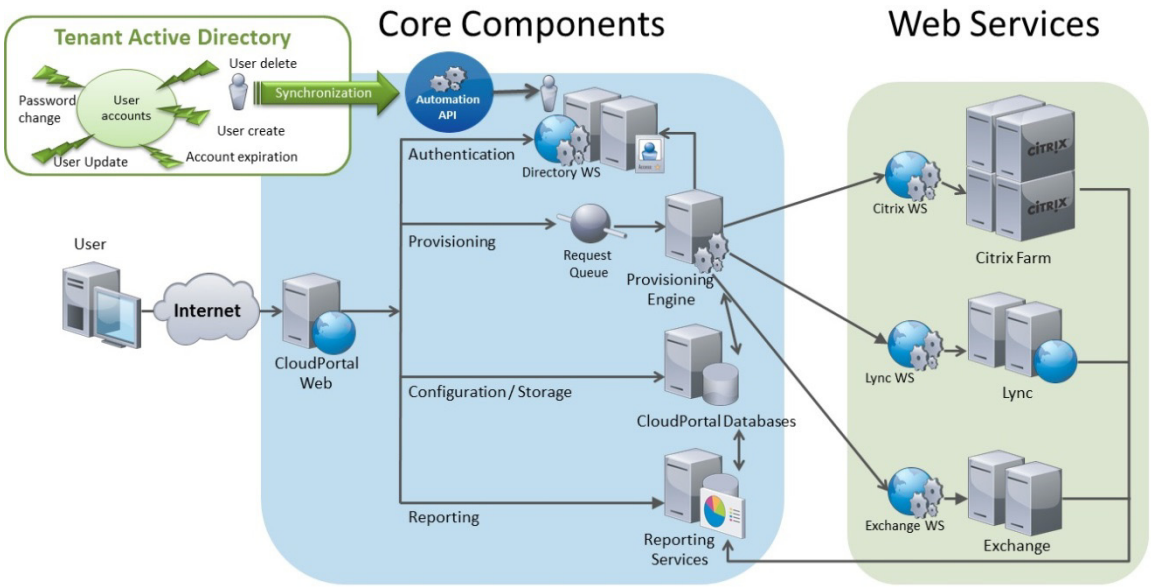

CloudPortal Services Manager User Interface

CloudPortal Services Manager (CPSM) provides a unified interface for CSP administration as well as delegated administration to resellers and end-customers. The CPSM Web UI (CortexWeb) is loosely coupled with the other CPSM components. This loose coupling provides several security benefits. The web server has no dependency on Active Directory so it can essentially operate outside of the managed domain. The website can be locked down and run with minimal administrative permissions while still allowing the CPSM system to complete administrative tasks.

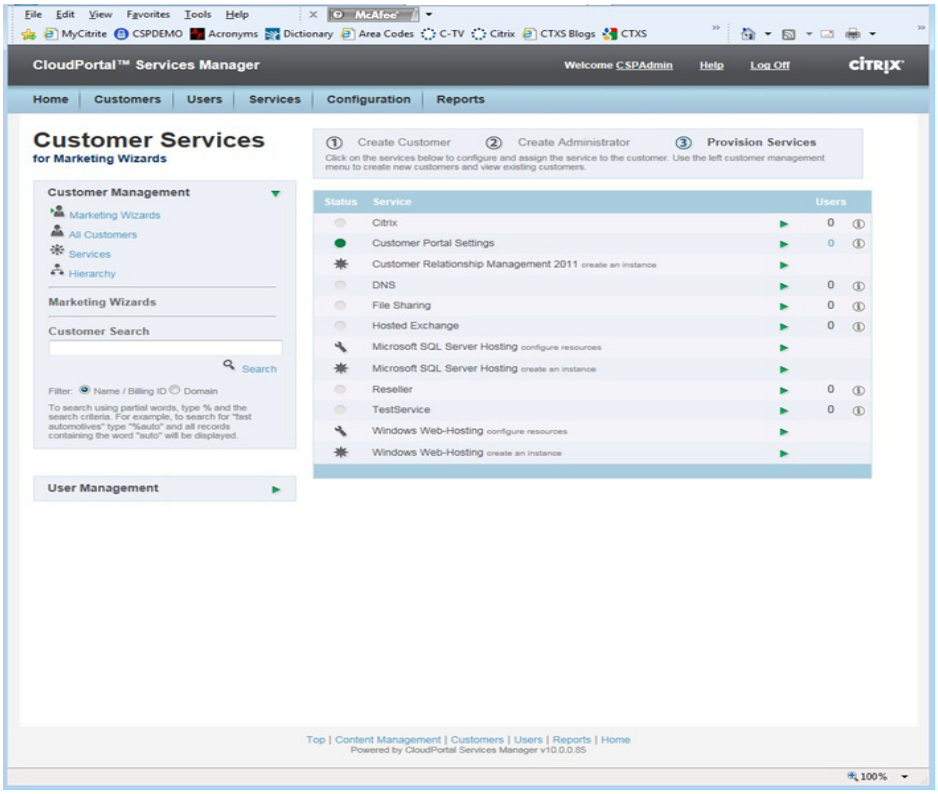

CPSM System Databases

A Microsoft SQL Server (CortexSQL) provides the backbone of the CPSM system. The CPSM databases store configuration information for all provisioned services, as well as all customer and user details. The databases also act as a cache mechanism for Active Directory, ensuring rapid user response without the need for slower AD queries. In addition, the databases store logging and auditing information for all provisioning transactions that pass through the system.

CPSM Provisioning Engine

The CPSM provisioning engine (CortexProvisioning) runs as a Windows Service. It monitors the provisioning queues for requests. When the provisioning engine receives a request, it follows provisioning rules to determine the actions required to complete the task.

The provisioning rules are easily customized using a simple Windows-based graphical interface that also provides a simple way to understand specific provisioning processes, which is helpful when troubleshooting problems. This interface can also be used to customize the provisioning process and to integrate new rules for custom services.

Each provisioning action performs a reusable piece of work, typically associated with provisioning applications. CPSM includes over 100 provisioning actions. Example actions include:

- Creating an Active Directory user
- Creating a security group in Active Directory
- Creating a folder in a file system
- Creating an address list in Microsoft Exchange

- Running a shell command or a Visual Basic script

All provisioning processes are built using provisioning actions, enabling quick setup with little coding, while giving the service providers visibility into the processes being executed in their environment.

Active Directory Web Service (ADWS)

The Active Directory web service provides a secure and simple interface to Active Directory. The CPSM website uses this service to perform real time tasks such as user authentication and password expiry status.

Reporting

CPSM uses Microsoft SQL Server Reporting Services to deliver usage reporting capability through the CloudPortal Services Manager user interface. CloudPortal Services Manager interacts directly with the reporting services web service interface and allows controlled publishing of reports to all users of the CloudPortal Services Manager system.

EMC VNX5600 Storage Configuration

Figure 28 shows the physical storage layout of the disks in the reference architecture to support up to 2000 desktops.

Figure 28 Storage Configuration

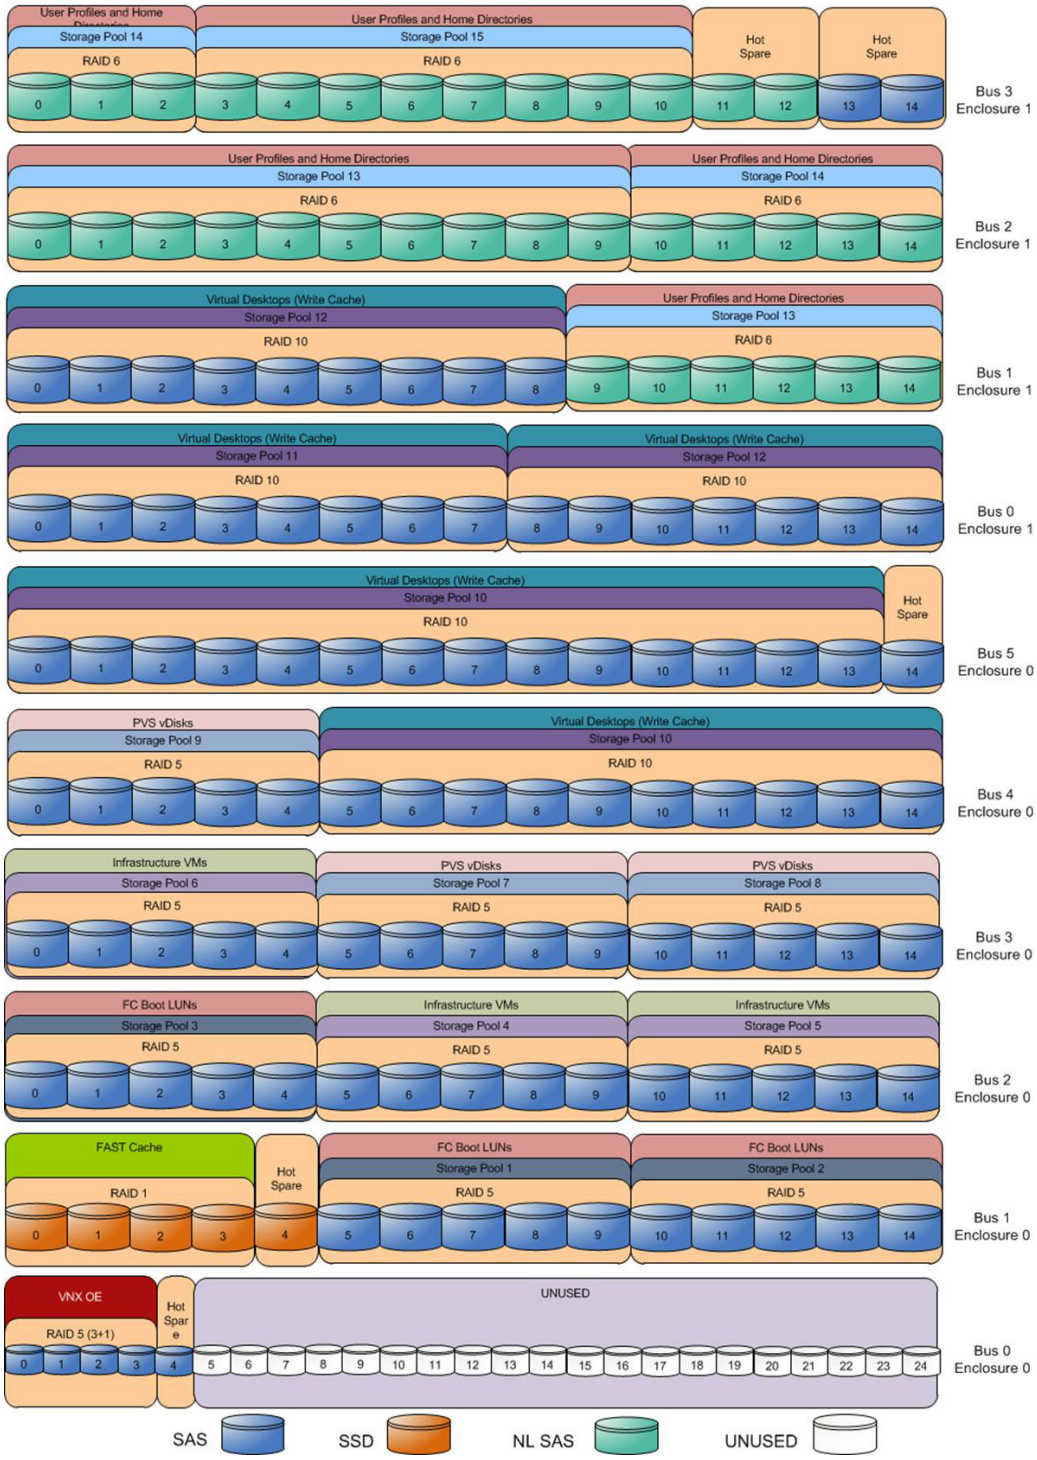

The above storage layout is used for the following configurations:

- Four SAS disks (0_0_0 to 0_0_3) are used for the VNX OE.

- The VNX series does not require a dedicated hot spare drive. Disks 0_0_4, 1_0_4, 5_0_14, and 3_1_11 to 3_1_14 are unbound disks that can be used as hot spares when needed. These disks are marked as hot spares in the diagram.
- Four 200GB Flash drives are used for EMC VNX FAST Cache. See the "EMC FAST Cache in Practice" section below to follow the FAST Cache configuration best practices.
- Five SAS disks (1_0_5 to 1_0_9) on the RAID 5 storage pool 1 are used to store the first pool of FC boot LUNs for vSphere hosts.
- Five SAS disks (1_0_10 to 1_0_14) on the RAID 5 storage pool 2 are used to store the second pool of FC boot LUNs for vSphere hosts.
- Five SAS disks (2_0_0 to 2_0_4) on the RAID 5 storage pool 3 are used to store the third pool of FC boot LUNs for vSphere hosts.
- Five SAS disks (2_0_5 to 2_0_9) on the RAID 5 storage pool 4 are used to store the first pool of infrastructure VMs.
- Five SAS disks (2_0_10 to 2_0_14) on the RAID 5 storage pool 5 are used to store the second pool of infrastructure VMs.
- Five SAS disks (3_0_0 to 3_0_4) on the RAID 5 storage pool 6 are used to store the third pool of infrastructure VMs.
- Five SAS disks (3_0_5 to 3_0_9) on the RAID 5 storage pool 7 are used to store the first pool of PVS vDisks.
- Five SAS disks (3_0_10 to 3_0_14) on the RAID 5 storage pool 8 are used to store the second pool of PVS vDisks.
- Five SAS disks (4_0_0 to 4_0_4) on the RAID 5 storage pool 9 are used to store the third pool of PVS vDisks.
- Twenty four SAS disks (4_0_5 to 4_0_14 and 5_0_0 to 5_0_13) on the RAID 10 storage pool 10 are used to store the first pool of PVS write cache allocated for the virtual desktops.
- Eight SAS disks (0_1_0 to 0_1_7) on the RAID 10 storage pool 11 are used to store the second pool of PVS write cache allocated for the virtual desktops.
- Sixteen SAS disks (0_1_8 to 0_1_14 and 1_1_0 to 1_1_8) on the RAID 10 storage pool 12 are used to store the third pool of PVS write cache allocated for the virtual desktops.
- Sixteen NL-SAS disks (1_1_9 to 1_1_14 and 2_1_0 to 2_1_9) on the RAID 6 storage pool 13 are used to store the first pool of user profiles and home directories.
- Eight NL-SAS disks (2_1_10 to 2_1_14 and 3_1_0 to 3_1_2) on the RAID 6 storage pool 14 are used to store the second pool of user profiles and home directories.
- Eight NL-SAS disks (3_1_3 to 3_1_10) on the RAID 6 storage pool 15 are used to store the third pool of user profiles and home directories.
- FAST Cache is enabled on all storage pools.
- Disks 0_0_5 to 0_0_24 are unbound. They are not used for testing this solution.
- All SAS disks used for this solution are 300GB.

The graphic below illustrates how the EMC DataPools were defined per tenant prior to UCS Director configured datastores.

# of users	# of tenants	Total users	Isolation Model	FlexCast Model (90/10 HSD/SVDI	User data (CIFS) 5GB per user	User profile (CIFS) 50MB per user	BootLUN(FC) 10GB per LUN	Infra(NFS) 500GB per DataStore	vDisk(CIFS)	Write Cache(NFS)	VLAN Isolation
50	4	200	Shared Site Private Group	Hosted Shared and Server VDI	DataPool1	DataPool1	BootLUNPool1	InfraPool1	vDiskPool1(1TB)	1.45 TB	Shared VLAN
100	1	100	Shared Site Private Group	Hosted Shared and Server VDI	DataPool2	DataPool2	BootLUNPool1	InfraPool1	vDiskPool1(1TB)	0.75 TB	Shared VLAN
150	2	300	Shared Site Private Group	Hosted Shared and Server VDI	DataPool1	DataPool1	BootLUNPool1	InfraPool1	vDiskPool1(1TB)	2.10TB	Shared VLAN
200	1	200	Private Delivery Site	Hosted Shared and Server VDI	DataPool3	DataPool3	BootLUNPool2	InfraPool2	vDiskPool2(500GB)	1.32 TB	Dedicated VLAN
500	1	500	Shared Site Private Group Server Isolated	Hosted Shared and Server VDI	DataPool2	DataPool2	BootLUNPool1	InfraPool1	vDiskPool1(1TB)	3.35 TB	Dedicated VLAN
700	1	700	Private Delivery Site	Hosted Shared and Server VDI	DataPool4	DataPool4	BootLUNPool3	InfraPool3	vDiskPool3(500GB)	4.63 TB	Dedicated VLAN
		2000			9.8 TB	98 GB	160GB	2.15 TB	2 TB	13.6 TB	

EMC FAST Cache in Practice

FAST Cache is best for small random I/O where data has skew; the higher the locality, the better the FAST Cache benefits.

General Considerations

EMC recommends first utilizing available flash drives for FAST Cache, which can globally benefit all LUNs in the storage system. Then supplement performance as needed with additional flash drives in storage pool tiers.

- Match the FAST Cache size to the size of active data set.
  - For existing EMC VNX/CX4 customers, EMC Pre-Sales have tools that can help determine active data set size.
- If active dataset size is unknown, size FAST Cache to be 5 percent of your capacity, or make up any shortfall with a flash tier within storage pools.
- Consider the ratio of FAST Cache drives to working drives. Although a small FAST Cache can satisfy a high IOPS requirement, large storage pool configurations will distribute I/O across all pool resources. A large pool of HDDs might be able to provide better performance than a few drives of FAST Cache.

Preferred application workloads for FAST Cache:

- Small-block random I/O applications with high locality
- High frequency of access to the same data
- Systems where current performance is limited by HDD capability, not SP capability

AVOID enabling FAST Cache for LUNs that are not expected to benefit, such as when:

- The primary workload is sequential.
- The primary workload is large-block I/O.

AVOID enabling FAST Cache for LUNs where the workload is small-block sequential, including:

- Database logs
- Circular logs
- VNX OE for File SavVol (snapshot storage)

Enabling FAST Cache on a Running System

When adding FAST Cache to a running system, it is recommended to enable FAST Cache on a few LUNs at a time, and then wait until those LUNs have equalized in FAST Cache before adding more LUNs.

FAST Cache can improve overall system performance if the current bottleneck is drive-related, but boosting the IOPS will result in greater CPU utilization on the SPs. Systems should be sized so that the maximum sustained utilization is 70 percent. On an existing system, check the SP CPU utilization of the system, and then proceed as follows:

- Less than 60 percent SP CPU utilization – enable groups of LUNs or one pool at a time; let them equalize in the cache, and ensure that SP CPU utilization is still acceptable before turning on FAST Cache for more LUNs/pools.
- 60-80 percent SP CPU utilization – scale in carefully; enable FAST Cache on one or two LUNs at a time, and verify that SP CPU utilization does not go above 80 percent.
- CPU greater than 80 percent – DON'T activate FAST Cache.

AVOID enabling FAST Cache for a group of LUNs where the aggregate LUN capacity exceeds 20 times the total FAST Cache capacity.

- Enable FAST Cache on a subset of the LUNs first and allow them to equalize before adding the other LUNs.

To enable FAST Cache as an array-wide feature in the system properties of the array in EMC Unisphere, complete the following steps:

1. Click the **FAST Cache** tab, then click **Create** and select the Flash drives to create the FAST Cache. RAID 1 is the only RAID type allowed. There are no user-configurable parameters for FAST Cache. In this solution, four 200GB SSD drives were used for FAST Cache.

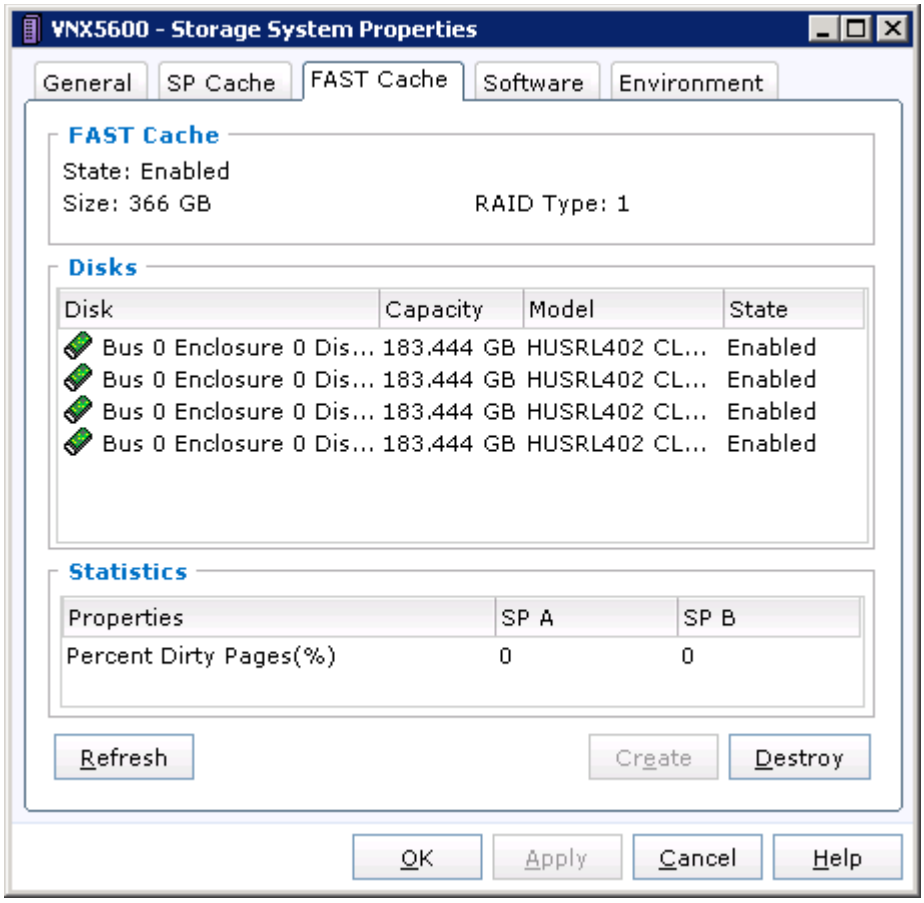

2. To enable FAST Cache for a particular pool, navigate to the **Storage Pool Properties** page in Unisphere, and then click the **Advanced** tab. Select **Enabled** to enable FAST Cache.

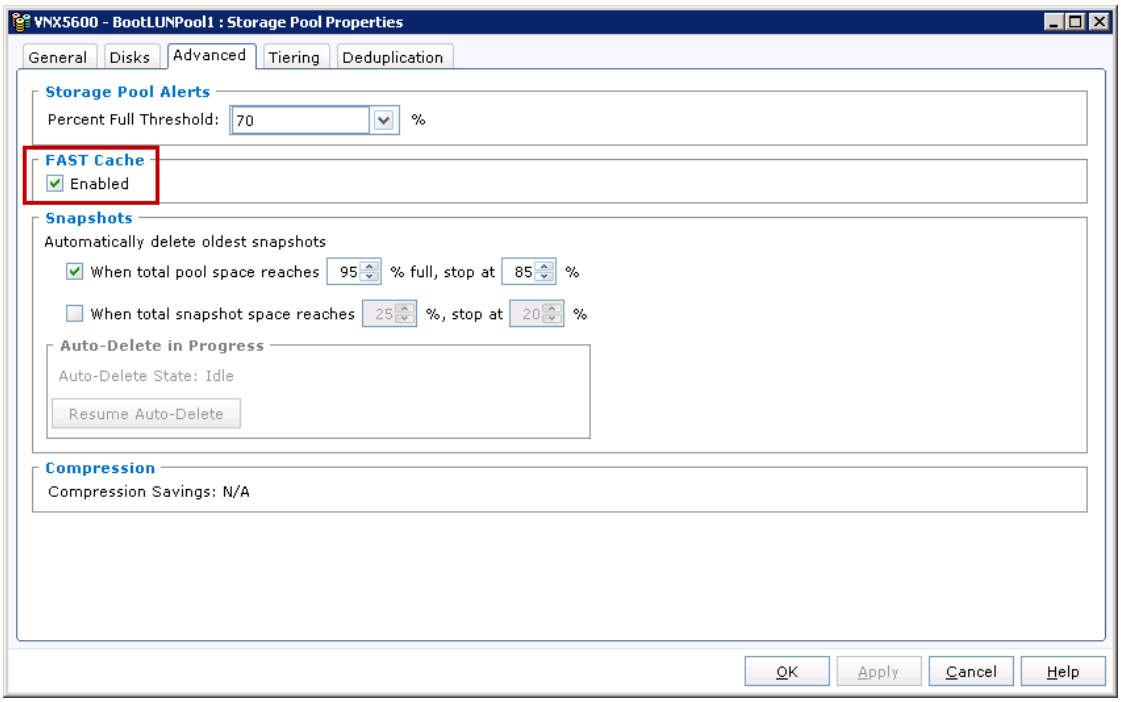

EMC Additional Configuration Information

The following tuning configurations optimize NFS/CIFS performance on the VNX5600 Data Movers:

NFS/CIFS Active Threads Per Data Mover

The default number of threads dedicated to serve NFS/CIFS requests is 512 per Data Mover on VNX5600. Some use cases such as the scanning of desktops might require more number of NFS active threads. It is recommended to increase the number of active NFS and CIFS threads to 2048 on the active Data Mover to support up to 2000 desktops.

1. The **nthreads** parameters can be set by using the following commands:

```
# server_param server_2 -facility nfs -modify nthreads -value 2048
# server_param server_2 -facility cifs -modify nthreads -value 2048
```

Reboot the Data Mover for the change to take effect.

2. Type the following command to confirm the value of the parameter:

```
# server_param server_2 -facility nfs -info nthreads
server_2 :
name                    = nthreads
facility_name           = nfs
default_value           = 384
current_value           = 2048
configured_value        = 2048
user_action             = none
change_effective        = immediate
range                   = (32,2048)
description            = Number of threads dedicated to serve nfs requests, and
memory dependent.
```

3. The values of NFS and CIFS active threads can also be configured by editing the properties of the **nthreads** Data Mover parameter in **Settings–Data Mover Parameters** menu in Unisphere. Type **nthreads** in the filter field, highlight the **nthreads** value you want to edit and select **Properties** to open the nthreads properties window.
4. Update the **Value** field with the new value and click **OK** . Perform this procedure for each of the **nthreads** Data Mover parameters listed menu. Reboot the Data Movers for the change to take effect.

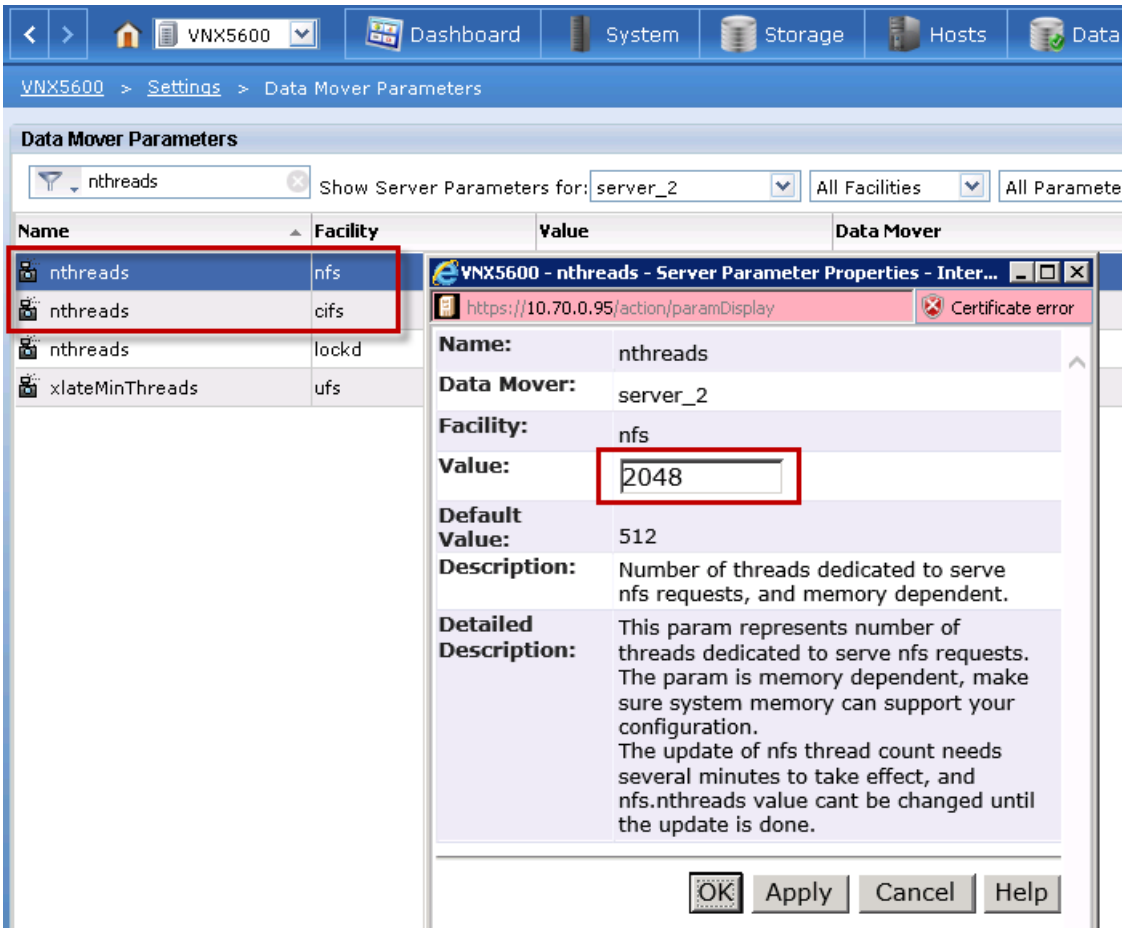

Solution Validation

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

Configuration Topology for a Cisco-Citrix DaaS Solution

Figure 29 illustrates the architectural diagram for the purpose of this study.

Figure 29 Cisco Solutions for EMC VSPEX XenDesktop 7.5 2000 Seat Architecture Block Diagram

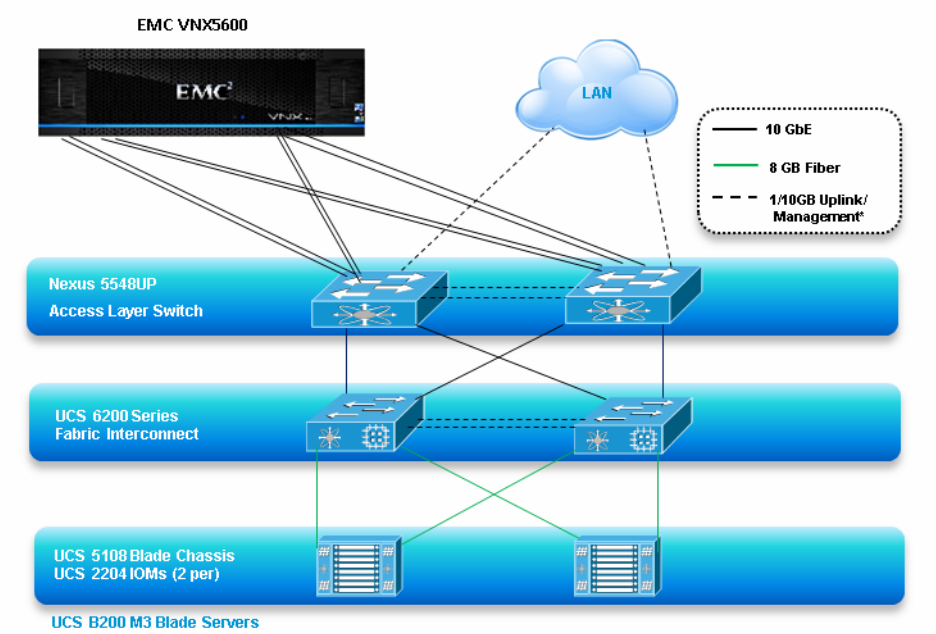

The architecture is divided into four distinct layers:

- Cisco UCS Compute Platform
- The Virtual Desktop Infrastructure and Virtual Desktops that run on Cisco UCS blade hypervisor hosts
- Network Access layer and LAN
- Storage Access through NFS on EMC VNX5600 deployment

Figure 30 details the physical configuration of the 2000 seat Citrix DaaS environment.

Figure 30 Detailed Architecture of the EMC VNX5600, 2000 Seat DaaS Solution

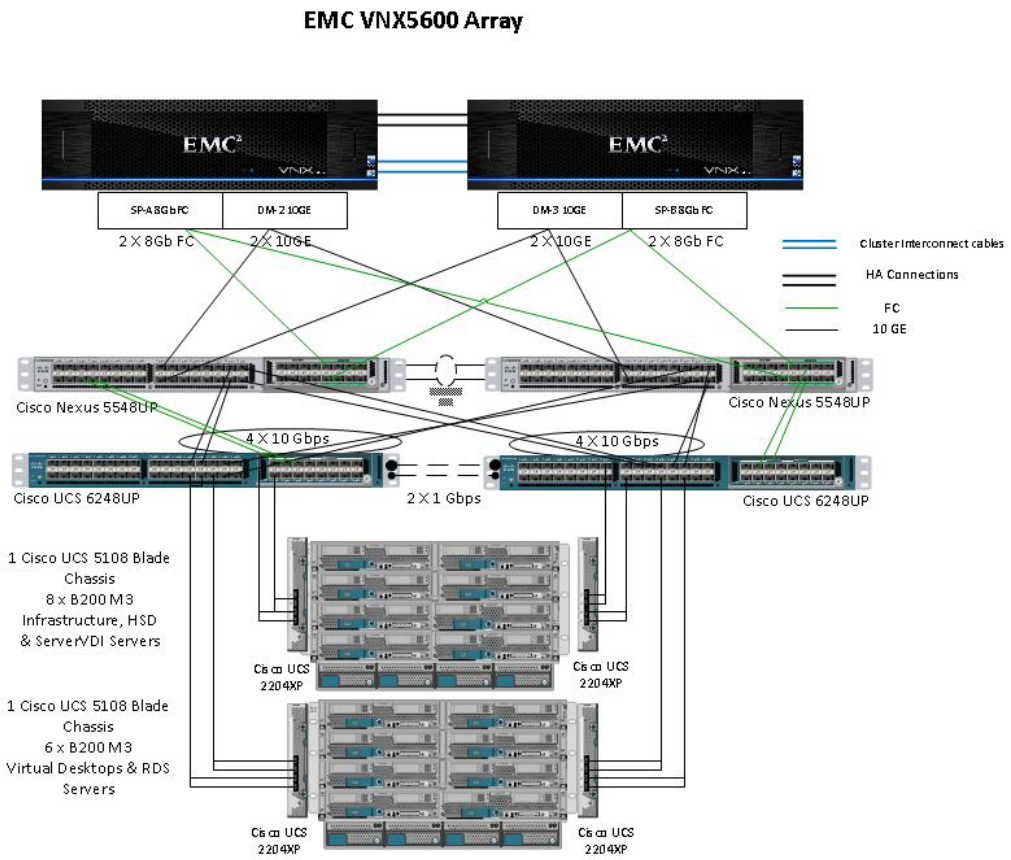

Cisco UCS Director Configuration

This section talks about the configuration that was done to utilize UCS Director to assist in the provisioning of Infrastructure and DaaS components. The installation and configuration of the UCS Director and UCS Baremetal agent are detailed here.

Cisco UCS Director Deployment

For this project we deployed the Cisco UCS Director and Baremetal Agent appliances to an already existing infrastructure environment that is independent of the infrastructure that the DaaS components will use. It is recommended that the Cisco UCS Director and Baremetal Agents run outside of the environments they will be deploying or managing.

For installation and deployment of Cisco UCS Director for DaaS, please refer to http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-director/vsphere-install-guide/4-1/b_Installing_UCSDirector_on_vSphere_41.html

Cisco UCS Director Bare Metal Agent Deployment

Refer to the “Cisco UCS Director Baremetal Agent Setup Guide” for information about installation and deployment of the “Cisco UCS Director Baremetal Agent”.:

http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-director/products-installation-guides-list.html

In order to properly utilize Cisco UCS Director a series of "Day-0 Tasks" must be completed as prerequisites. These tasks are as follows.

Storage: VNX (Day-0)

- Initialize Storage Controllers (SP-A, SP-B & Control Station)
- Configure Management IP Address for SP-A, SP-B & Control Station
- Create required Storage Pools for Block
- Create and present file Storage (Disk Volumes) to DataMovers
- Cabling: FC Cables from SP-A & SP-B to N5K pair, 10G Cables from Datamover to N5K pair, Create LACP (NIC Teamed Interface) for each Datamover
- Configure CIFS server with DNS and AD Integration

Cisco Nexus N5K

Initialize Nexus 5K Switches, configure Management IP Address (See configuration section)

- Cable from Compute and Storage
- Create and configure required vPC & port channels for UCSM & Storage (See N5K configuration)
- Create VSAN-A and B for SAN (See N5K configuration)
- Configure minimum VLANs for Management Network access (See N5K Configuration)

Cisco UCS Manager

Cisco UCS Director provides support for Cisco UCS (Unified Computing System) infrastructure. It provides auto- discovery, monitoring and complete visibility to manage all Cisco UCS components. Following section(s) explain adding Cisco UCS account into Cisco UCS Director to support VSPEX functionality.

- Rack, Stack and Cabled
- Configure Fis in cluster mode with Management IP Address
- Configure bare-minimum pools and policies and management VLAN for at least to run a single ESXi Server

VMware

1. Bring up at least one ESXi Server with Management Network

Note For assistance in installation and configuration of Day-0 ESXi and vCenter server please reference this Validated Design on pages 125-135)
http://www.cisco.com/c/dam/en/us/td/docs/unified_computing/ucs/UCS_CVDs/ucs_vspex_xd75.pdf

2. Install vCenter with required licenses

Note When Day-0 pre-requisite tasks are complete, you are ready to login to the Cisco UCS Director system and begin the configuration.

3. Log in to the system by entering the appliance IP Address into a web browser. (Default username and passowrd: ‘admin’)

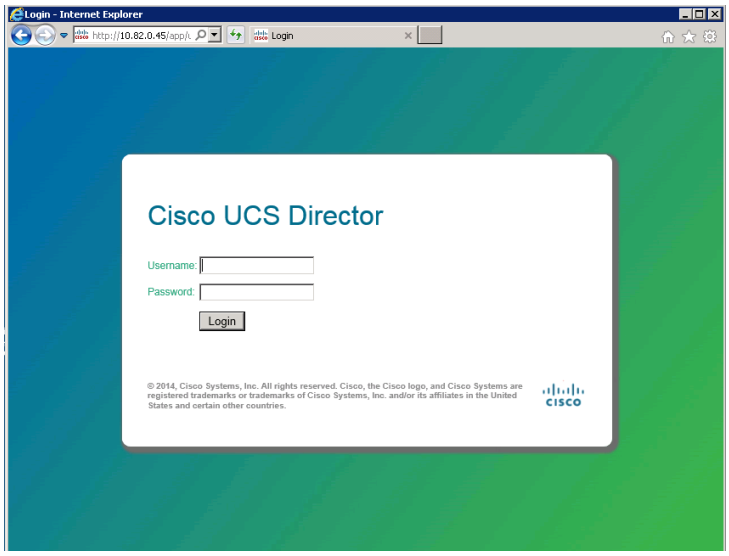

4. UCS Director Home Dashboard

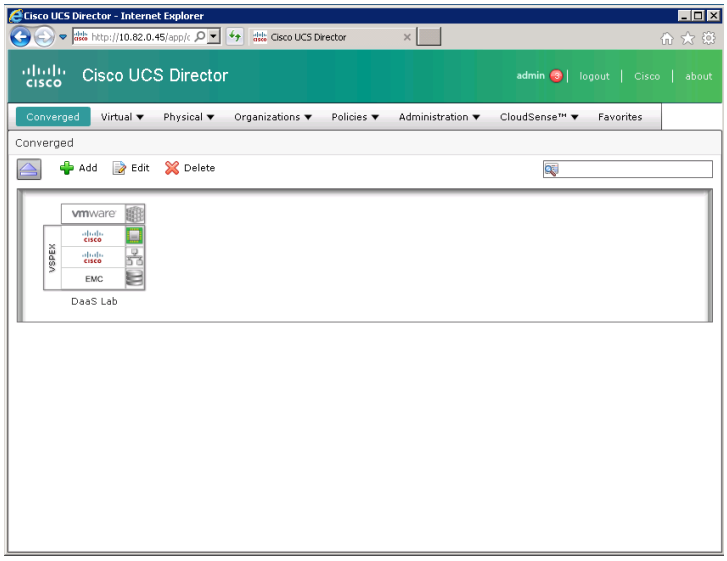

Add Data Center Pod

To add a Cisco UCS Manager (UCSM) account, a Data Center needs to be added first, complete the following steps:

1. Select Administrator →Physical Accounts Pod tab
2. Click 'Add' to add a Data Center.
3. Specify the Data Center Pod 'Name', select the 'Type' of Data Center and the location 'Address'. Then click on 'Add' to create the Data Center.

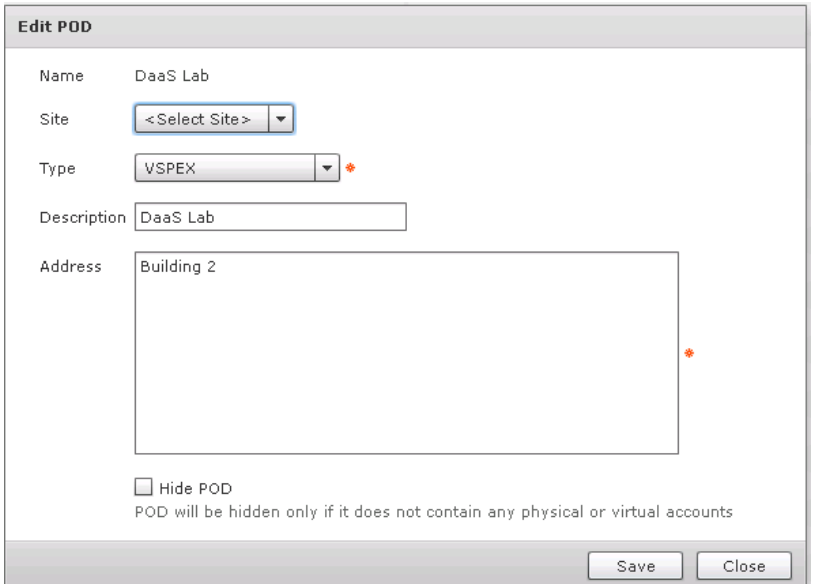

Add Cisco UCS Account

When a Data Center has been added, a Cisco UCS Manager account can be added; to do so, complete the following steps:

1. Select Administrator →Physical Accounts →Physical Accounts tab→ Add

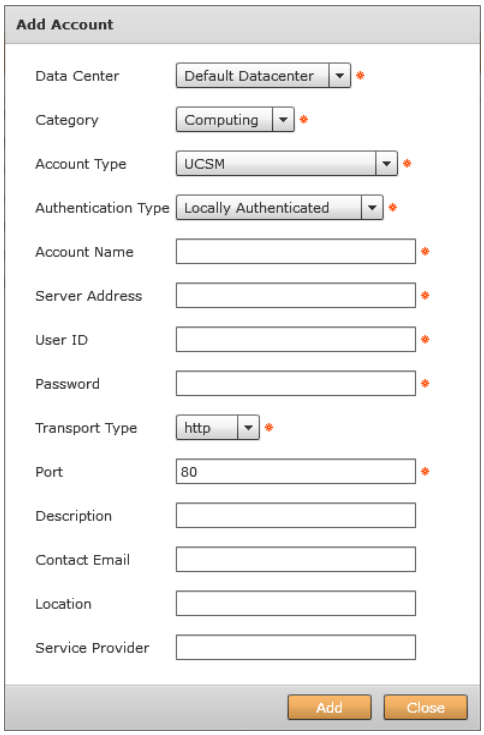

Field Name	Description
Data Center	Select the Data Center to which the UCSM account will be associated to.
Category Type	Sepcify the type of infrastructure. In this case 'Computing'.
Account Type	Select the account type. In this case UCSM.
Authentication Type	Specify the authentication Type , Locally Authenticated or Remotely Authenticated. Locally Authenticated User Accounts - A locally authenticated user account is authenticated directly through the fabric interconnect and can been enabled or disabled by anyone with admin or AAA privileges. Remotely Authenticated User Accounts - A remotely authenticated user account is any user account that is authenticated through LDAP, RADIUS, or TACACS+.
Account Name	Specify a name for the UCSM account.
Server Address	Specify the IP address of the UCSM.

User ID	Specify the user id UCSM.
Password	Specify the password for the UCSM.
Field Name	**Description**
Transport Type	Select the transport type, either http or https.
Port	Specify the port number of the UCSM.
Description	Specify the description if required.
Contact Email	Specify the email address if required.
Location	Specify the location of the UCSM if required.
Service Provide	Service provider name if any.

Cisco UCS Director will automatically discover all infrastructure elements in the Cisco UCSM account like Chassis, Servers, Fabric Interconnects, Service Profiles, Server Pools etc. in the newly added UCSM account. Typically the discovery process takes about 5 minutes.

Cisco UCS Director Configuration for EMC VNX 5600

Cisco UCS Director provides support for EMC VNX storage . It supports auto- discovery, monitoring and complete visibility to manage all the VNX (VNX 5600) components. Following section(s) explain adding EMC VNX account into Cisco UCS Director to support DaaS Storage functionality.

Add EMC VNX Account

1. Select Administration→Physical Accounts →Physical Accounts tab→Click on 'Add'

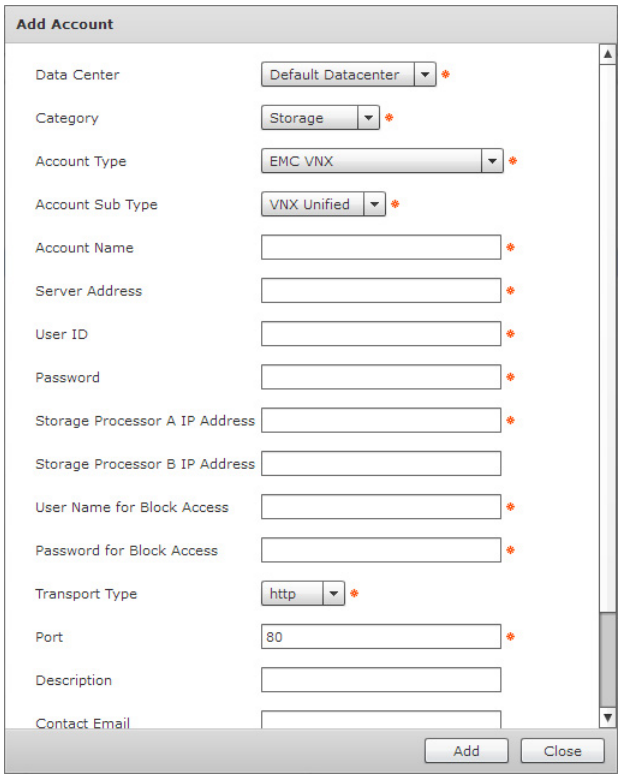

Table 4 Add Account Fields Explanation

Field Name	Description
Data Center	Select the Data Center to which the compute account is added.
Category Type	Sepcify the type of infrastructure. In this case 'Storage'.
Account Type	Select the account type. In this case EMC VNX.
Account Sub Type	Select the VNX account sub type – File, Block or Unified.
Account Name	Specify a name for the VNX account.
Server Address	Specify the IP address of the VNX.
User ID	Specify the user id VNX.
Password	Specify the password for the VNX.
Storage Processor A IP	Specify the IP address of Storage Processor A.

Storage Processor B IP	Specify the IP address of Storage Processor B.
User Name for Block Access	Specify the user name for block access.
Password for Block	Specify the password for block access.
Transport Type	Select the transport type, either http or https.
Port	Specify the port number of the VNX.
Description	Specify the description if required.
Contact Email	Specify the email address if required.
Location	Specify the location of the VNX if required.
Service Provide	Service provider name if any.

Cisco UCS Director Nexus 5K Configuration

Cisco UCS Director provides support for a multitude of Network devices. Users can add the devices to the Cisco UCS Director and monitor them. The device categories currently supported are:

- Cisco IOS devices
- Cisco Nexus OS devices
- Cisco UCS Fabric Interconnect

The following sections explain adding Cisco 5K device(s) account into Cisco UCS Director to support DaaS functionality.

Add Network Devices

To add Cisco Nexus Device(s) to the Cisco UCS Director required for this DaaS configuration's functionality:

1. Select Administrator→→Physical Accounts→→ Manage Network Elements tab →→ Click on 'Add Network Element'

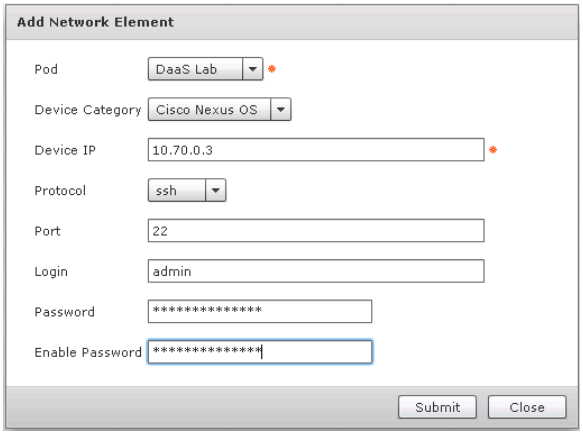

Field Name	Description
Data Center	Select the Data Center to which the other account compute and storage are added.
Device Category	Select the type of device category being added.
Device IP	Specify the IP address of the network device.
Protocol	Select the protocol used to communicate with the device. Either telnet or ssh can be used.
Port	The port number of the network device
Login	Specify the login id of the device
Password	Specify the device password
Enable Password	Certain devices require a separate password to enter in the command configuration mode. Specify any such password in this field.

Cisco UCS Director will discover the devices and collect inventory from the network devices and display them in the form of tabular reports. To view device details that is already added to Cisco UCS Director .

2. Select Physical →Network.
3. Select the Data Center name from the left column.
4. Select the 'Managed Network Element' tab.

To view details of a specific device:

1. Select a device from the list and click on 'View Details'

   This will display all information related to the device; Interfaces, Configurations, Port Profiles, Private VLANs, and Port Capabilities, etc.

Add VMware Data Center Cloud

To add a VMware Data Center to the Cisco UCS Director required for this DaaS configuration's functionality, complete the following steps:

1. Select Administrator→Virtual Accounts→→Click on 'Virtual Account'.

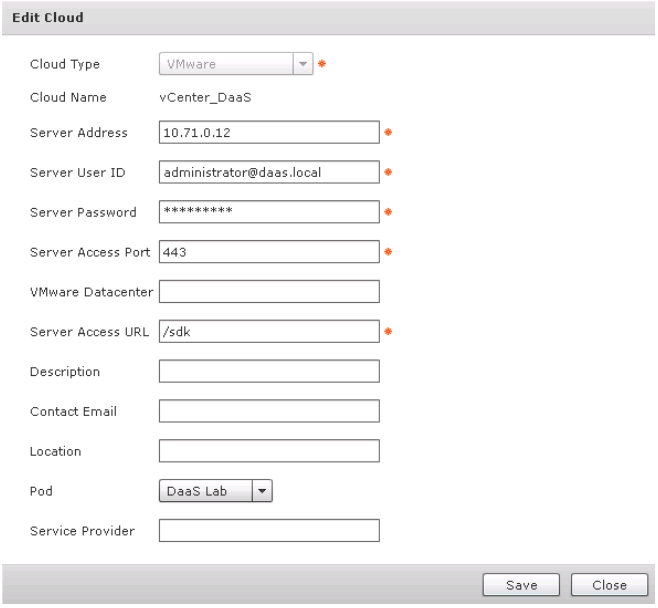

Field Name	Description
Cloud Type	Select Cloud Type to be used in UCS Director converged Data Center.
Cloud Name	Select Cloud name for display in UCS Director.
Server Address	Specify the IP address of the Vcenter machine.
Server User ID	Username for Vcenter administrator account.
Server	Vcenter administrator password.
Server Access	Port for Vcenter server communication.
VMware	VMware Datacenter name.
Server Access	URL of Vcenter server access.
Description	Description for Vcenter Datacenter
Contact E-mail	E-mail for Datacenter contacts.
Location	Location of Datacenter.
Pod	Pod created in initial setup
Service	Service Provider name if available.

Cisco UCS Director will discover the devices and collect inventory from the Vcenter Data Center cloud and display them in the form of tabular reports. To view device details that is already added to Cisco UCS Director .

Provisioning ESXi to Cisco UCS Blade Servers using Cisco UCS Director

In this project we provisioned a Cisco Custom ESXi 5.5 image to our Cisco B200 M3 servers using a custom workflow In UCS Director. We imported a custom workflow created by Cisco for deploying ESXi on EMC VNX systems. That custom workflow can be downloaded here:

https://communities.cisco.com/servlet/JiveServlet/download/55002-3-81893/Deploy_ESXi_Host_On_Vblock_3XX_with_VisionIO_v2_0.wfdx.zip

This workflow provisions ESXi Host booting from SAN and Integrates with vCenter Server and Nexus 1000v DVSwitch on VNX5600 Array.

To instuall, complete the following steps:

1. Download the attached .ZIP file below to your computer. *Remember the location of the saved file on your computer.
2. Unzip the file on your computer. Should end up with a .WFDX file.
3. Log in to UCS Director as a user that has "system-admin" privileges.
4. Navigate to "Policies-->Orchestration" and click on "Import".
5. Click "Browse" and navigate to the location on your computer where the .WFDX file resides. Choose the .WFDX file and click "Open".
6. Click "Upload" and then "OK" once the file upload is completed, Click "Next".
7. Click "Import".
8. A new folder called 'Vblock-3XX-Validated' should appear in "Policies-->Orchestration" that contains the imported workflow. You will now need to update the included tasks with information about the specific environment.

Prerequisites for Using this Workflow

- The DaaS Hardware Infrastructure components (VMware vCenter, Cisco UCSM, Cisco Nexus 5Ks, Cisco Nexus 1000v & EMC VNX5600 Storage Array) are added into UCS Director
- Cisco UCS Director BMA is integrated into Cisco UCS Director
- VMware ESXi Image (with N1Kv) from Cisco is configured in BMA. For more information please refer to Configure ESXi PXE Image.
- Modify 'ks.cfg' file under /opt/cnsaroot/templates/<ESXi_Image_Name> folder to enable ESXi Install on SAN Boot LUN
  - Comment the below line with '#"
    # install --firstdisk --overwritevmfs
  - Uncomment the below line by removing '#"

    install --firstdisk=remote --overwritevmfs
- Cisco UCS Service Profile Template is created with 2 vNICs & 2 vHBAs, Boot Order is configured with FC Boot First and LAN Second, FC Target Ports are added into SAN Boot Policy and Boot LUN ID is set to '0'. vHBA-0 desired order set to '0' and vHBA-1 desired order set to '1'. Service Profile's Power State set to 'Off'.

- Cisco UCS Service Profile Template vNICs are configured with BMA PXE VLAN and set it as Native VLAN.

Workflow Tasks

1. Create UCS Service Profile from Template
2. Select UCS Server
3. Setup PXE Boot with BMA Selection
4. Generic Configure SAN Zoning
5. Create VNX LUN
6. Create VNX Storage Group
7. Add VNX Host Initiator Entry - vHBA1
8. Add VNX Host Initiator Entry - vHBA2
9. Add VNX Hosts to VNX Storage Group
10. Add VNX LUN to Storage Group
11. Associate UCS Service Profile
12. Reset UCS Server
13. Monitor PXE Boot
14. Reset UCS Server
15. Wait for Specified Duration
16. Register Host with vCenter
17. Create VMware Port Group
18. Add Hosts to DVSwitch
19. Migrate vSwitch VMkernel Port to DVSwitch (Custom Task)
20. Migrate Default vSwitch to DVSwitch by Mapping Policy
21. Add Virtual Adapter (vMotion Interface)

User Inputs

The user inputs listed below should be provided by the user after executing the workflow.

User Input	User Input Description
ESXI_HOSTNAME	Enter ESXi Hostname
ESXI_CLUSTER	Select ESXi Cluster to which it need to be added
ESXI_MANAGEMENT_IP	Enter ESXi Management Network IP Address
ESXI_VMOTION_IP_OR_IPPOOL	Enter an IP Address or IP Pool range for VMotion interface
N1KV_L3_CONTROL_IP_OR_IPPOOL	Enter IP Address or IP Address Pool Range for N1Kv L3 Control Access Interface.
UCS_SERVER_BLADE	Select UCS Server Blade to install VMware ESXi Hypervisor

Admin Inputs

The Admin Input values are pre-defined in the workflow by the admin user. Edit the Workflow properties, go to user inputs section and modify the below inputs appropriately.

User Input	User Input Description
ESXI_ROOT_PASSWORD	Enter ESXi root user password
ESXI_MANAGEMENT_SUBNETMASK	Enter VMware ESXi Management Network's Subnet Mask
ESXI_MANAGEMENT_GATEWAY_IP	Enter VMware ESXi Management Network's Gateway IP
ESXI_MANAGEMENT_VLAN	Enter ESXi Management VLAN ID
ESXI_VMOTION_IP_POOL	Enter ESXi VMotion IP Pool (Ex: 192.168.100.100-192.168.100.200)
ESXI_VMOTION_SUBNET_MASK	Enter ESXi VMotion Subnet Mask
DOMAIN_NAME_SERVER_IP	Enter DNS Server IP Address
ESXI_HOST_LICENSE	Enter ESXi Host License Key
ESXI_VMOTION_DV_PORTGROUP	Select N1Kv DV Port Group for vMotion Interface
VCENTER_ACCOUNT	Select VMware vCenter Account
VCENTER_DATACENTER_NAME	Enter VMware vCenter Datacenter Name (The Datacenter in which the ESXi Host is going to be added)
ESXI_CLUSTER	Select ESXi Cluster to which host need to be added
N1KV_DVSWITCH	Select N1Kv DV Switch Name
N1KV_L3_CONTROL_PORTGROUP	Select N1Kv L3 Control Access Port Group
N1KV_MGMT_UPLINK_PORTGROUP	Select N1Kv Management Uplink Port Group
N1KV_L3_CONTROL_VLAN	Enter VLAN ID for N1Kv L3 Control Access Interface
N1KV_L3_CONTROL_IP_RANGE	Enter IP Range for N1Kv L3 Control Access Network (Ex: 192.168.99.100-192.168.99.200)
N1KV_L3_CONTROL_SUBNET_MASK	Enter Subnet Mask for N1Kv L3 Control Access Network
MGMT_PORTGROUP_MAPPING_POLICY	Select N1Kv Management Portgroup Mapping Policy
UCS_ORGANIZATION	Select UCS Organization
UCS_SP_TEMPLATE	Select UCS Service Profile Template
SAN_FABRIC_SWITCH_A1	Select SAN Fabric Switch A
SAN_FABRIC_SWITCH_B1	Select SAN Fabric Switch B
SANBOOT_TARGET_PORTS_FABRIC_A	Select VNX FC Target Port for SAN Zoning on Fabric A
SANBOOT_TARGET_PORTS_FABRIC_B	Select VNX FC Target Port for SAN Zoning on Fabric B
SAN_CONTROLLER_ACCOUNT	Select SAN Boot Storage Controller
SANBOOT_LUN_SIZE	Select SAN Boot LUN Size (GB)
SANBOOT_HOST_LUN_ID	Enter SAN Boot LUN Host ID (Should match UCS Service Profile Template Boot Policy)
VNX_FC_TARGET_PORT_FABRIC_A	Select VNX FC Target Port for SAN Boot on Fabric A
VNX_FC_TARGET_PORT_FABRIC_B	Select VNX FC Target Port for SAN Boot on Fabric B

Additional Workflow Configuration

After configuring the admin inputs, open the workflow using workflow designer and modify the "TASKS" by completing the following steps:

Note In upcoming versions, the below inputs should be able to configured as admin input values

1. Select UCS Server (Select UCS Server for ESXi Host): Open the Task > Click 'Next' > Click 'Next' > Click on 'Revalidate' button > Select appropriate 'UCSM' account, Click 'Next' > Click 'Submit'
2. Setup PXE Boot with BMA Selection (Configure PXE Server for ESXi Install): Open the Task > Click 'Next' > Click 'Next' > Select BMA Account > Select configured ESXi OS Image in BMA and Set appropriate 'Timezone', Click 'Next' > Click 'Submit'
3. Create VNX LUN (Create SANBoot LUN for ESXi Host):Open the Task > Click 'Next' > Click 'Next' > Select EMC VNX Account, Select Storage Pool Type, Select RAID Type, and Select RIAD Group Name for New LUN, Click 'Next' > Click 'Submit'
4. Add VNX Host Initiator Entry (Add ESXi vHBA-1 Host Initiator): Open the Task > Click 'Next' > Click 'Next' > Select EMC VNX Account, Click 'Next' > Click 'Submit'
5. Add VNX Host Initiator Entry (Add ESXi vHBA-2 Host Initiator): Open the Task > Click 'Next' > Click 'Next' > Select EMC VNX Account, Click 'Next' > Click 'Submit'
6. Register Host with vCenter (Add ESXi Host to Cluster on vCenter): Open the Task > Click 'Next' > Click 'Next' > Click on 'Revalidate' button, Click 'Next' > Click 'Submit'

User Inputs for Deploying ESXi to Blade Servers

Workflow User Inputs

Input Label	Admin Input Value	Input Description		
ESXI_HOSTNAME		Enter VMware ESXi Hostname	fal	ge
ESXI_ROOT_PASSWORD	**********	Enter ESXi root user password	fal	pa
SANBOOT_HOST_LUN_ID	0	Enter SAN Boot LUN Host ID	fal	ge
ESXI_MANAGEMENT_GATEWAY_IP	10.70.0.1	Enter VMware ESXi Management	fal	ip
DOMAIN_NAME_SERVER_IP	10.71.0.10	Enter DNS Server IP Address	fal	ip
ESXI_NFS_STORAGE_IP_RANGE	10.72.0.52-10.72.0.90		fal	St
ESXI_VMOTION_IP_POOL	10.73.0.109-10.73.0.129	Enter ESXi VMotion IP Pool	fal	St
MGMT_DV_PORTGROUP_MAPPING_POLI	2		fal	VM
SANBOOT_LUN_SIZE	20	Select SAN Boot LUN Size (GB)	fal	ge
ESXI_MANAGEMENT_SUBNETMASK	255.255.255.0	Enter VMware ESXi Management	fal	su
ESXI_NFS_STORAGE_SUBNET_MASK	255.255.255.0		fal	ge
ESXI_VMOTION_SUBNET_MASK	255.255.255.0	Enter ESXi VMotion Subnet Mask	fal	su
ESXI_MANAGEMENT_VLAN	70	Enter ESXi Management VLAN ID	fal	ge
SAN_FABRIC_SWITCH_A	DaaS Lab@10.70.0.2	Select SAN Fabric Switch A	fal	ne

SAN_FABRIC_SWITCH_B	DaaS Lab@10.70.0.3	Select SAN Fabric Switch B	fal	ne
SAN_CONTROLLER_ACCOUNT	DaaS Lab@VNX5600	Select SAN Boot Storage Controll	fal	EM
SANBOOT_TARGET_PORTS_FABRIC_B	DaaS Lab@VNX5600@@50		fal	GE
SANBOOT_TARGET_PORTS_FABRIC_A	DaaS Lab@VNX5600@@50		fal	GE
ESXI_HOST_LICENSE	R12C3-6825K-58341-0N88		fal	ge
ESXI_DV_NFS_STORAGE_PORTGROUP	Storage		fal	vn
UCS_ORGANIZATION	UCSM_SP;org-root	Select UCS Organization	fal	uc
UCS_SP_TEMPLATE	UCSM_SP;org-root;org-ro	Select UCS Service Profile Templ	fal	uc
SELECT_VMWARE_ACCOUNT	vCenter_DaaS		fal	vn
ESXI_VMOTION_DV_PORTGROUP	vMotion		fal	vn

Total 24 items

Submit Close

Figure 31 DaaS Workflow Diagram for Installing ESXi on Cisco UCS B200 M3 Servers

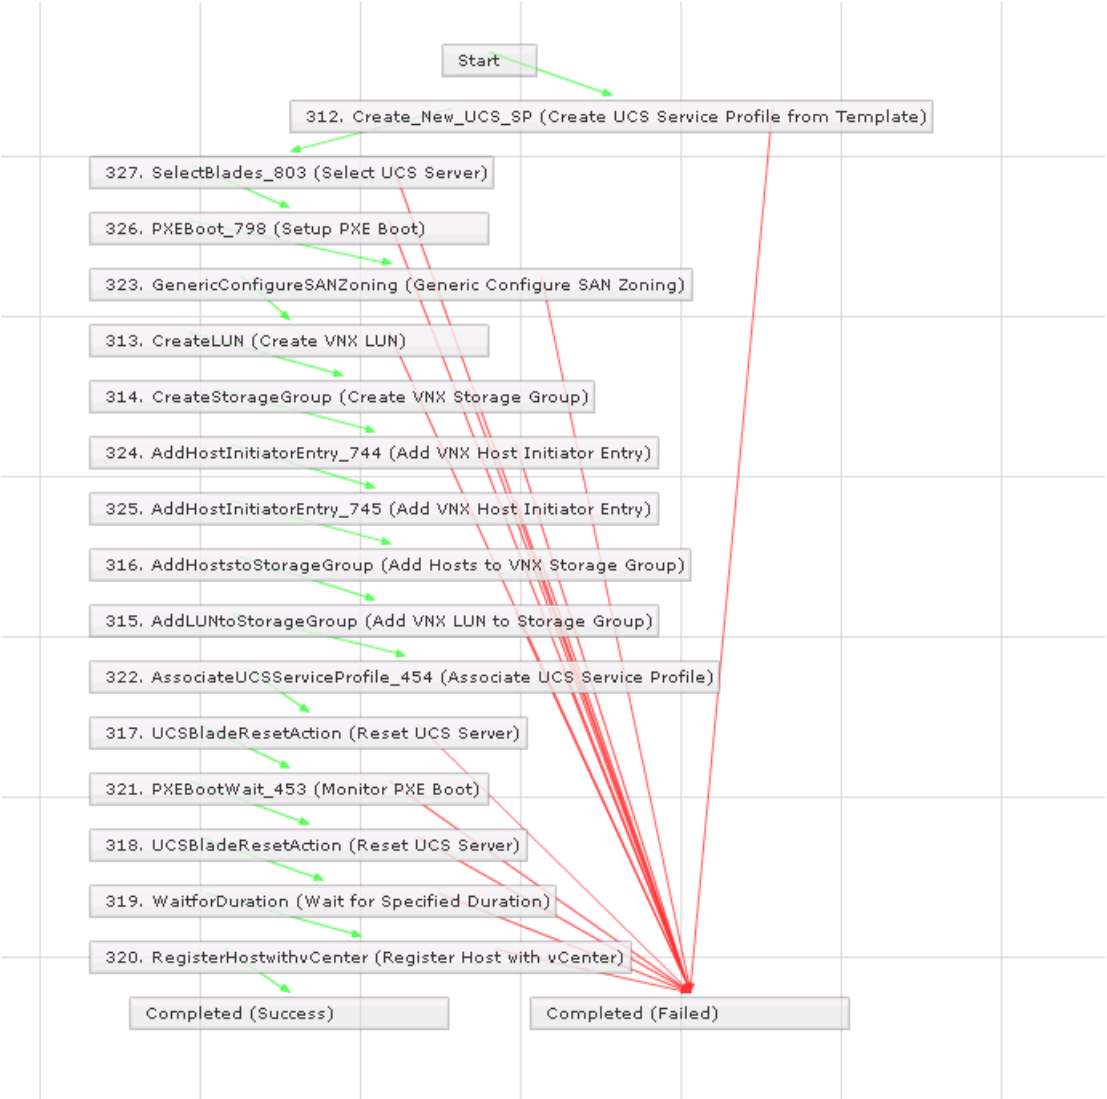

Deploying ESXi Install Workflow

To deploy ESXi to blades using the Cisco UCS Director workflow imported and customized earlier, complete the following steps.

1. From the main screen, Select ‘Policies’ → ‘Orchestration’.

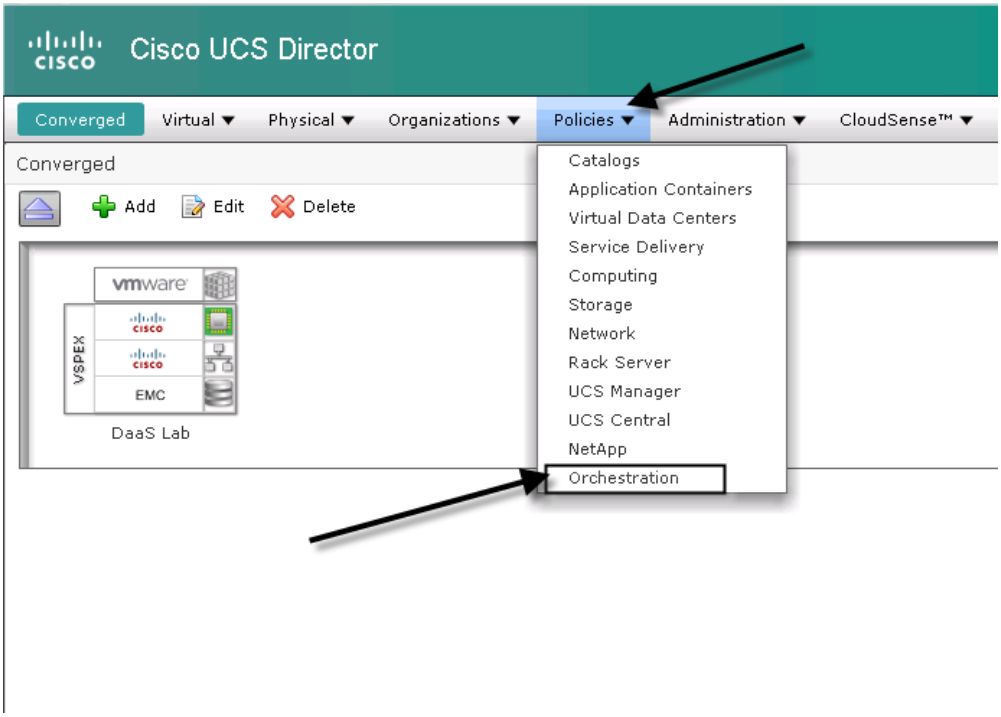

2. Double click the workflow created in the proper tenant.

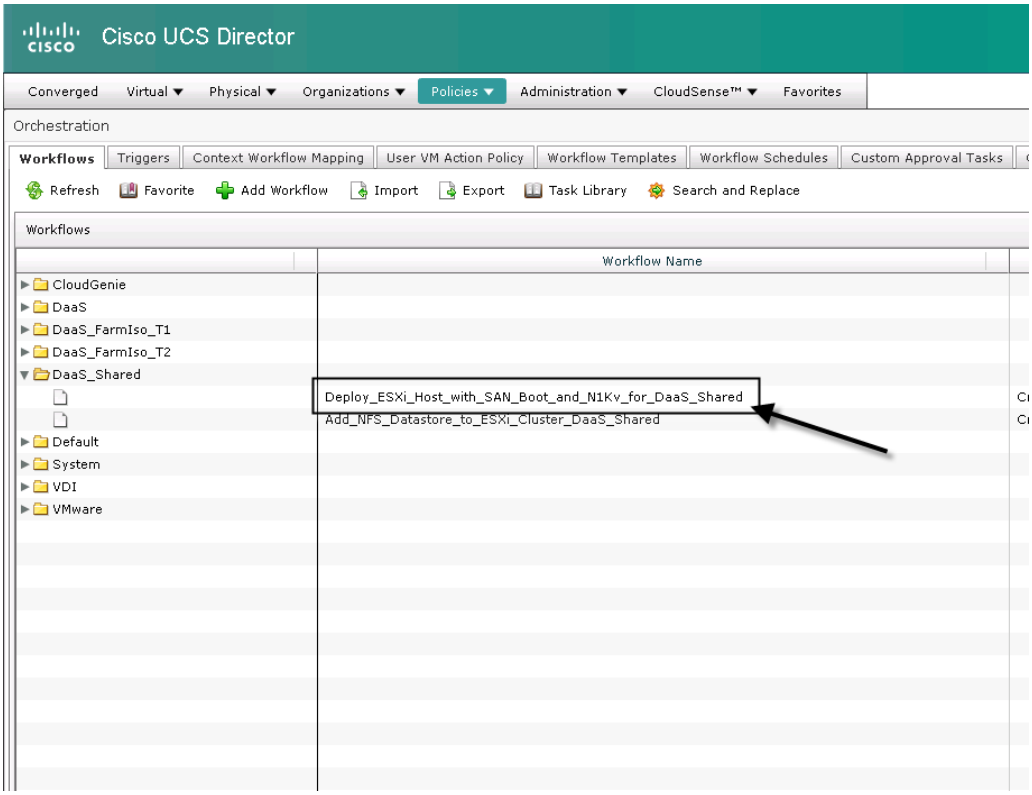

3. Double-click ‘Execute Now’

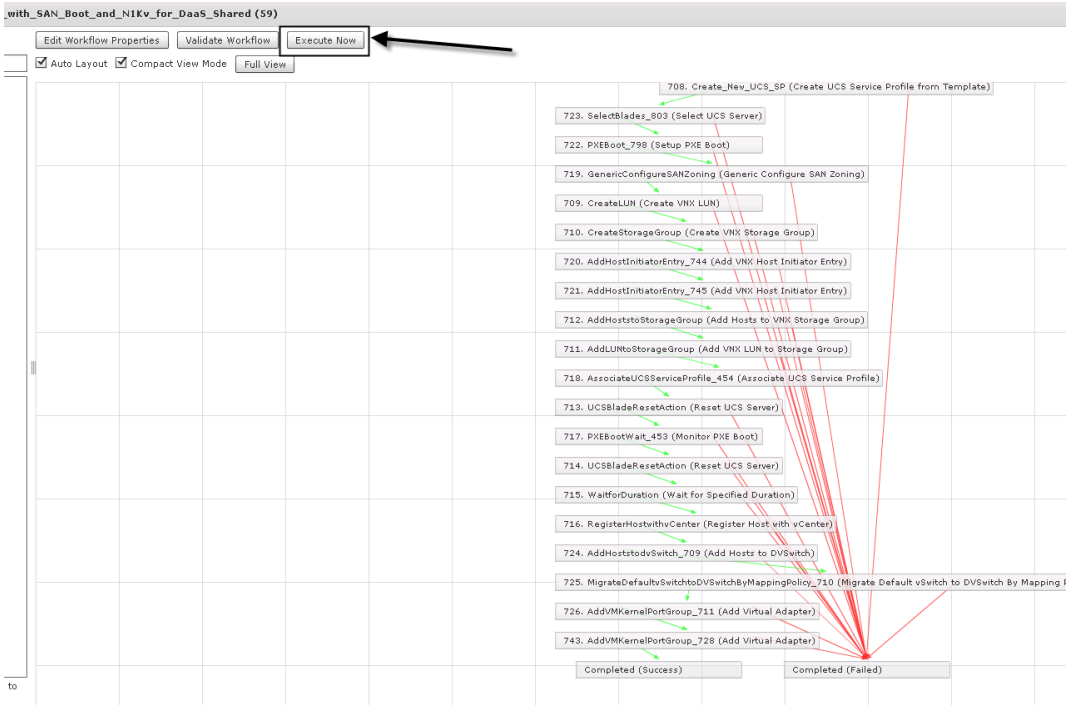

4. Enter the hostname of the blade you are deploying too and click ‘Submit’.

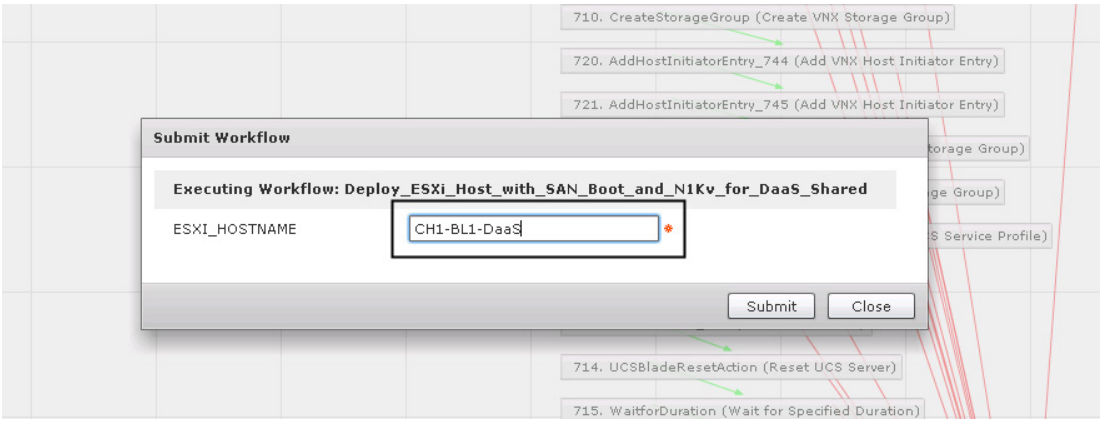

Cisco UCS Director Workflow for Adding NFS Datastores to the DaaS Environment

For this project we created a workflow to provision NFS datastores on demand for our Service Providers. Since we went through processes of onboarding 10 different tenants, some of which were private and some of which were shared infrastructures, we utilized different datastores for different tenants. We had to create new NFS datastores for user data (i.e. home drives, shared drives) and user profiles using Citrix UPM. We also used NFS datastores on the ESXi hosts to store the virtual machine files and write cache disks for the VMs provisioned by Citrix Provisioning Services. The following are steps and configuration we used to automate the NFS datastore provisioning using Cisco UCS Director.

Figure 32 Workflow for Adding NFS to Hosts

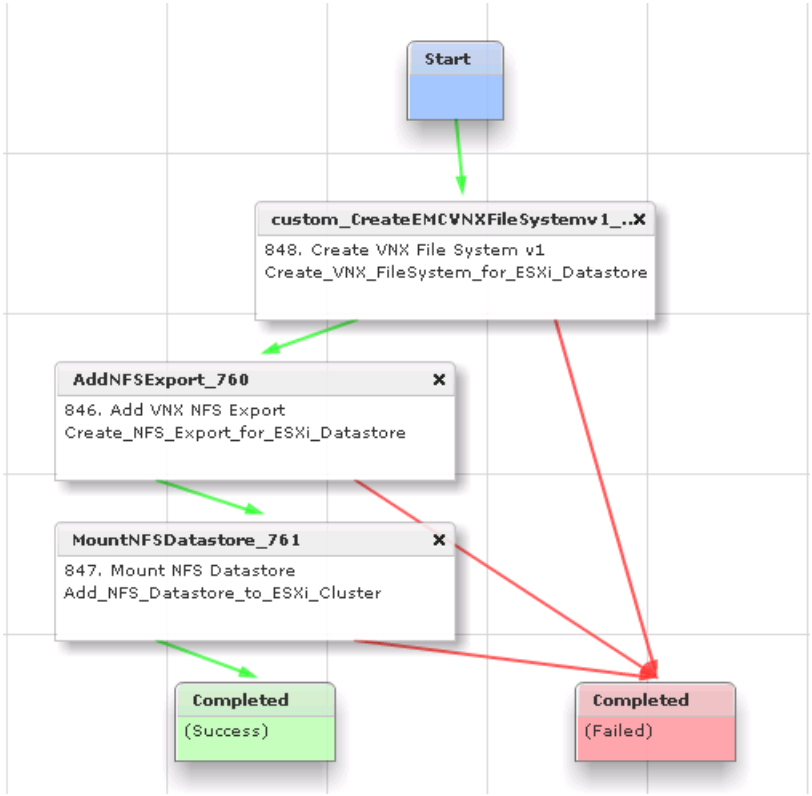

Figure 33 User Inputs for NFS Workflow

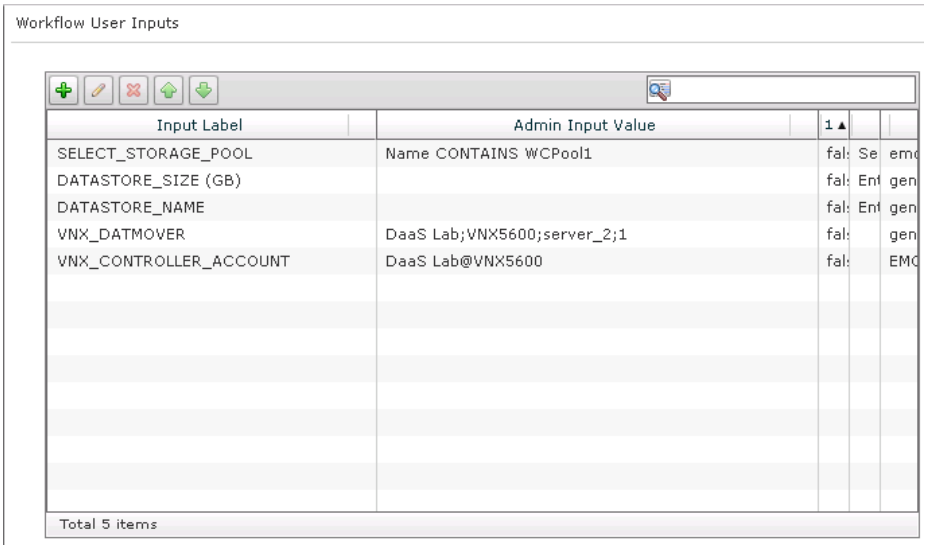

Workflow User Inputs

Input Label	Admin Input Value	1 ▲		
SELECT_STORAGE_POOL	Name CONTAINS WCPool1	fal:	Se	emc
DATASTORE_SIZE (GB)		fal:	Ent	gen
DATASTORE_NAME		fal:	Ent	gen
VNX_DATMOVER	DaaS Lab;VNX5600;server_2;1	fal:		gen
VNX_CONTROLLER_ACCOUNT	DaaS Lab@VNX5600	fal:		EMC

Total 5 items

Create a Catalog for Adding the NFS Datastore

Cisco UCS Director was used to provision storage for user data and Citrix PVS write cache data. To create a catalog, complete the following steps:

1. Under ‘Policies’ select ‘Catalogs

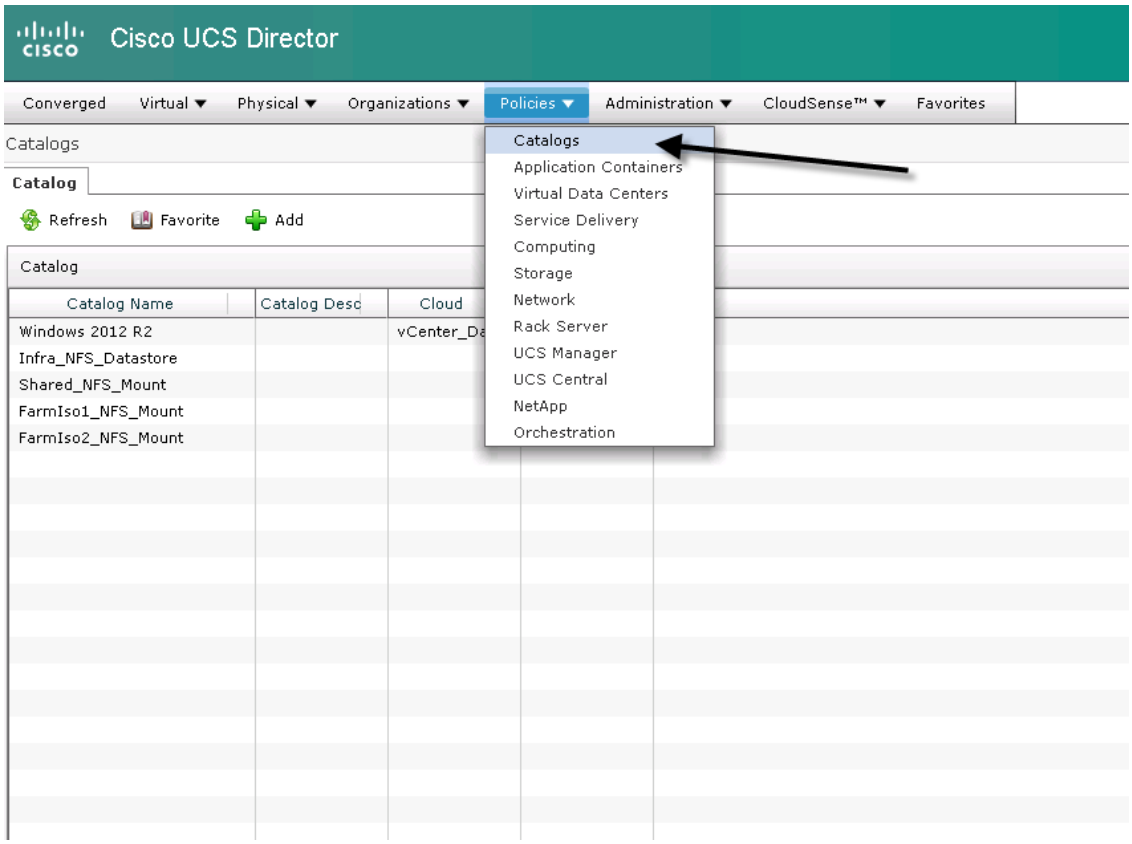

2. Select ‘Add' a new catalog.

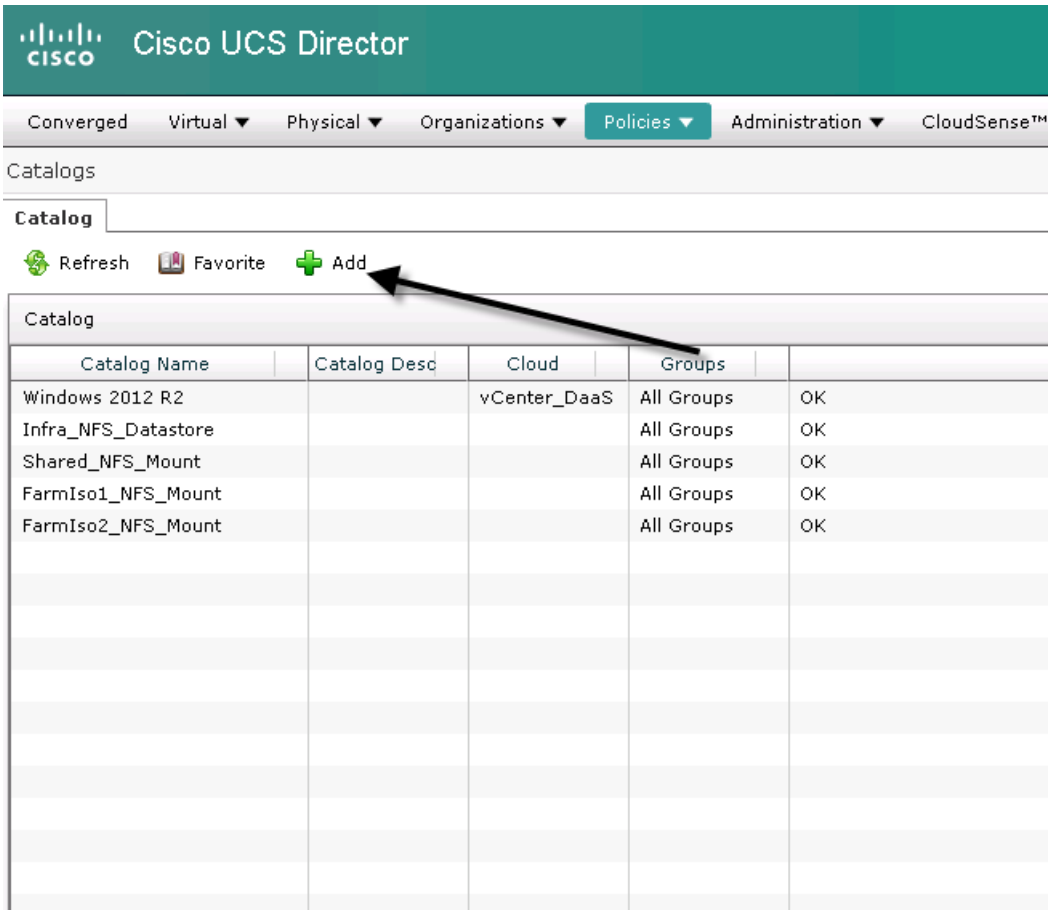

3. Create catalog steps
   a. Name the catalog
   b. Select ‘Advanced’ in catalog type
   c. Select ‘Workflow Icon’ for catalog Icon

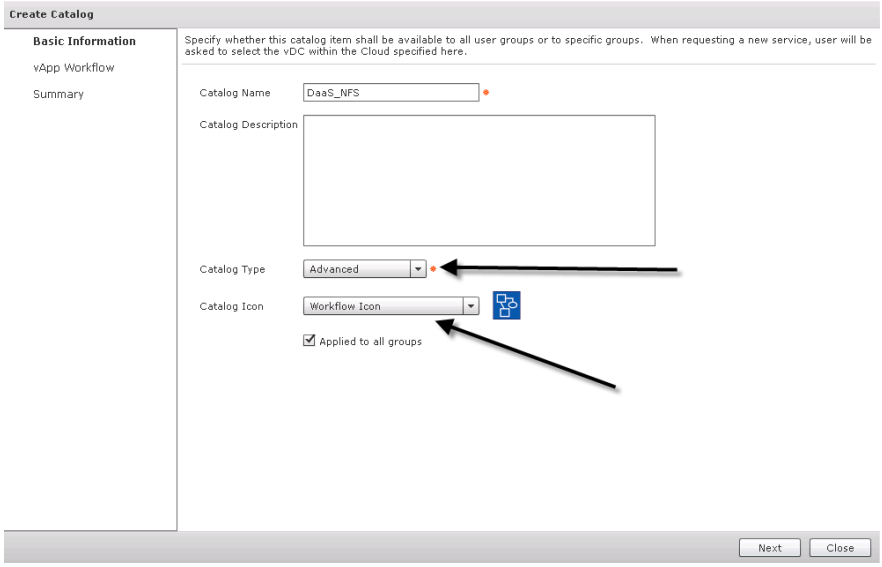

4. For Workflow, select the Workflow we created earlier.

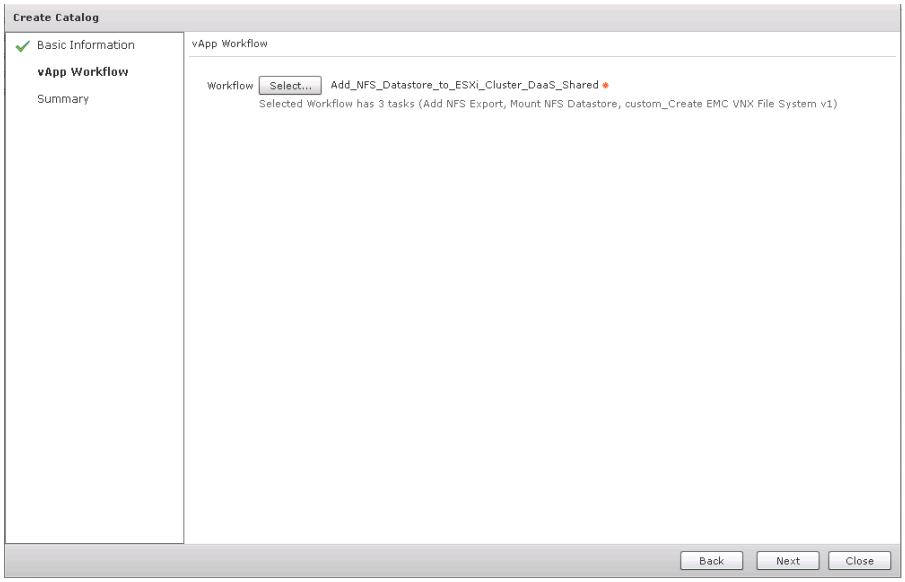

5. Select ‘Submit’ on the summary page

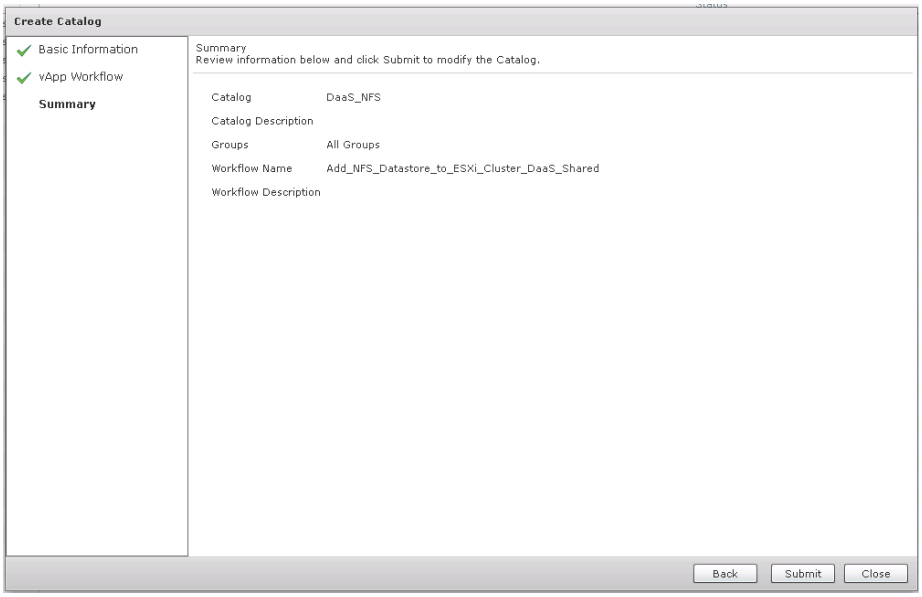

Create a Service Request for Creating and Mounting NFS Datastore to Hosts

To create the Service Request to mount an NFS datastore to the VMware hosts, complete the following steps:

1. Select ‘Organizations’ then ‘Service Requests’

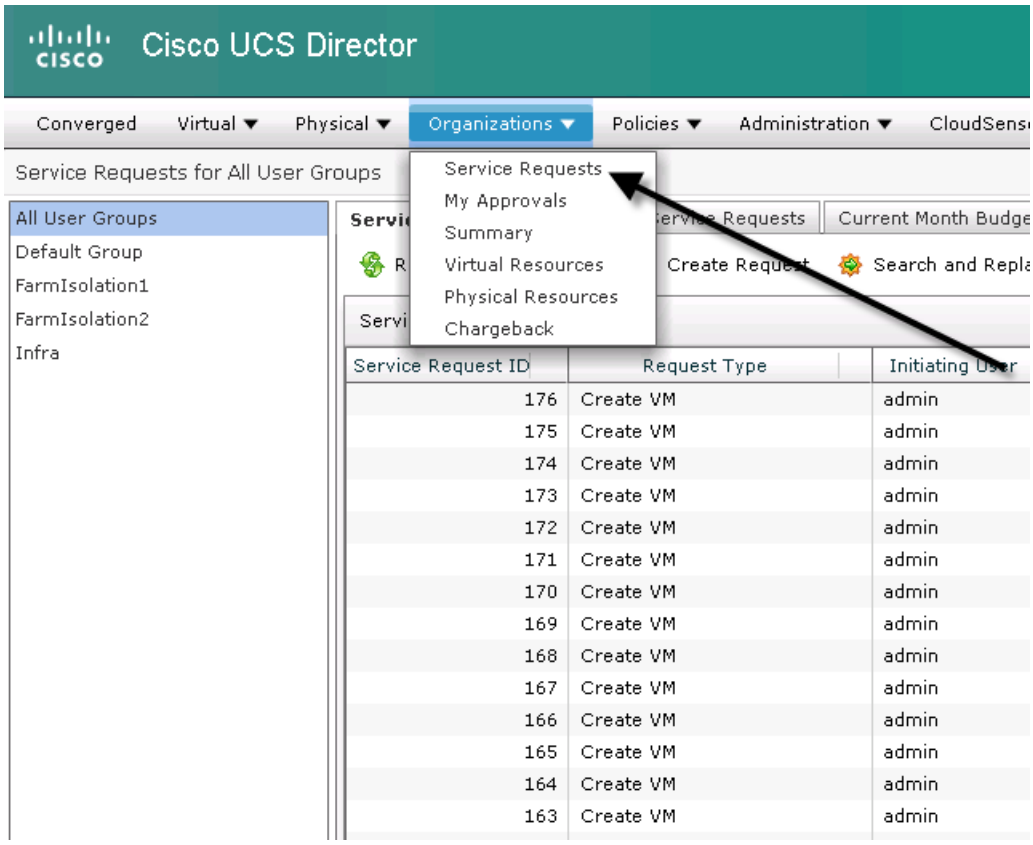

2. Click ‘Create Request’

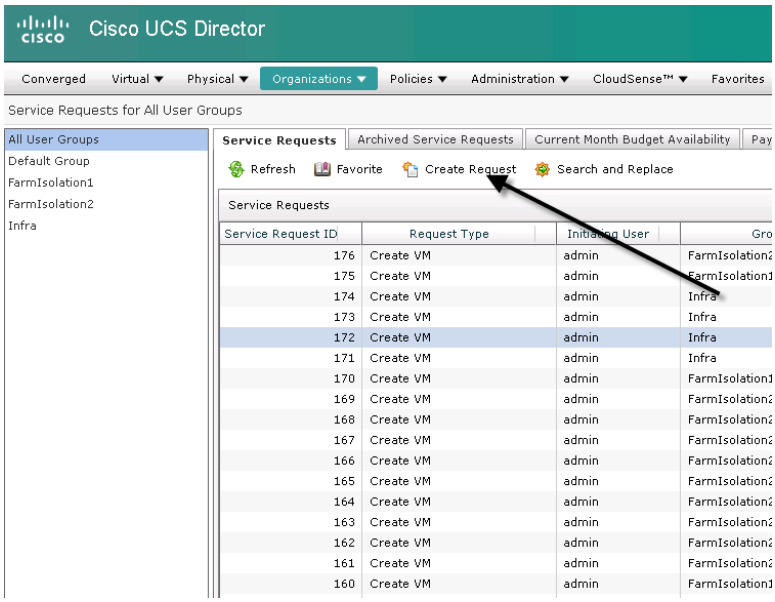

3. Enter inputs
   a. Select Group (pictured is Infrastructure, but change per tenant)

b. Catalog Type : Advanced

c. Select catalog created in prior section.

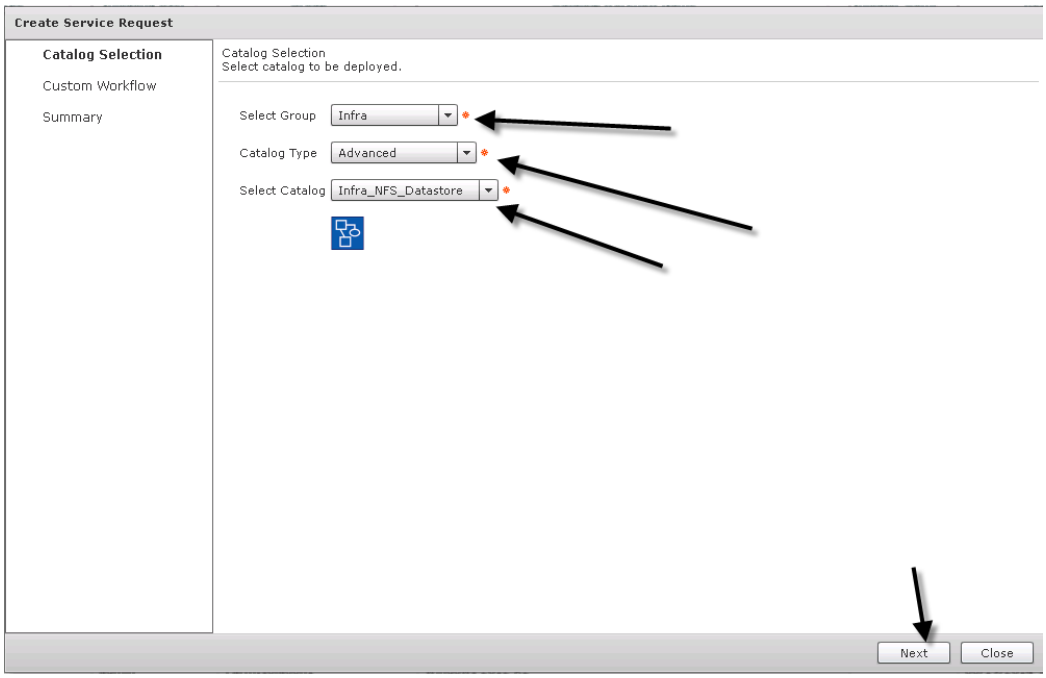

4. Enter DataStore name, Size and select from proper pool

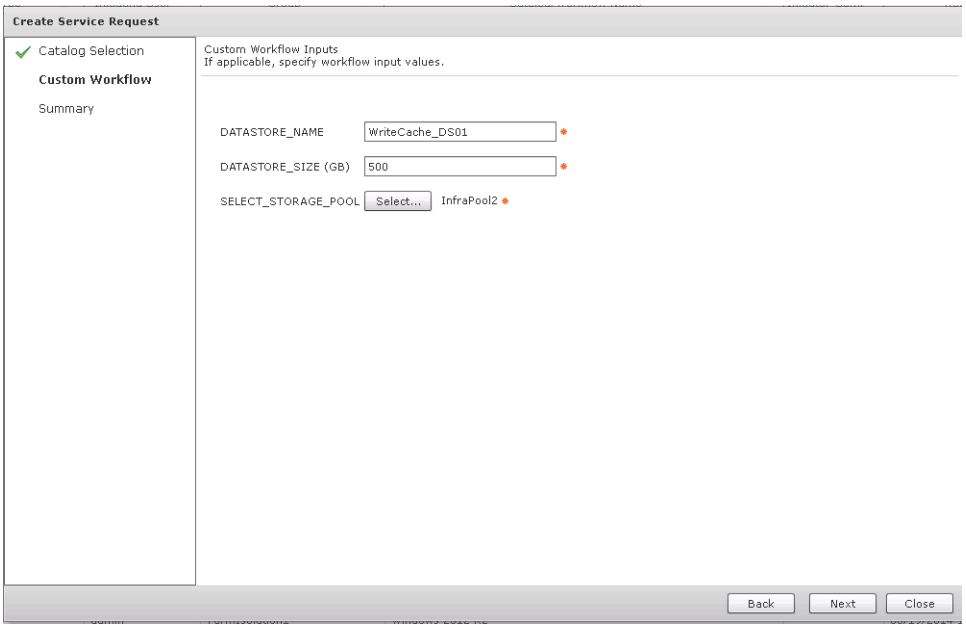

5. Check Summary and Submit your request

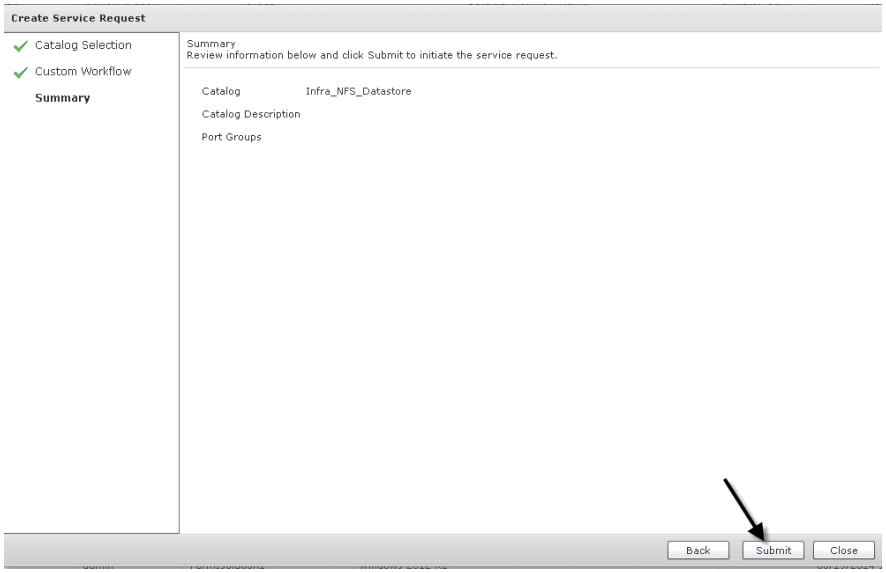

Create Virtual Machines in the VCenter_DaaS Cloud Using Cisco UCS Director

In our solution we created a VMware_DaaS Cloud that is part of our converged environment. This is detailed earlier in this document. With this Cloud, we were able to utilize Cisco UCS Director to provision virtual machines on demand. This process includes the creation of Virtual Data Centers and policies required to make them. Then creating a catalog to use when initiating the service request for VM provisioning.

Create Catalog for Virtual Machine Provisioning

1. Click 'Catalogs' under the Policies tab

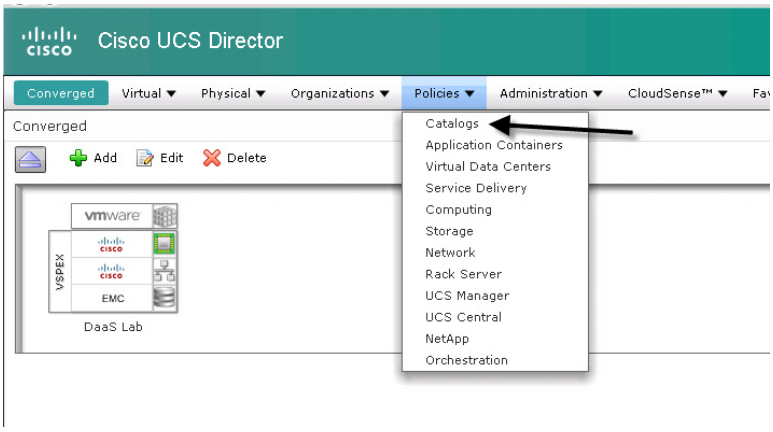

2. Click 'Add' to begin the process

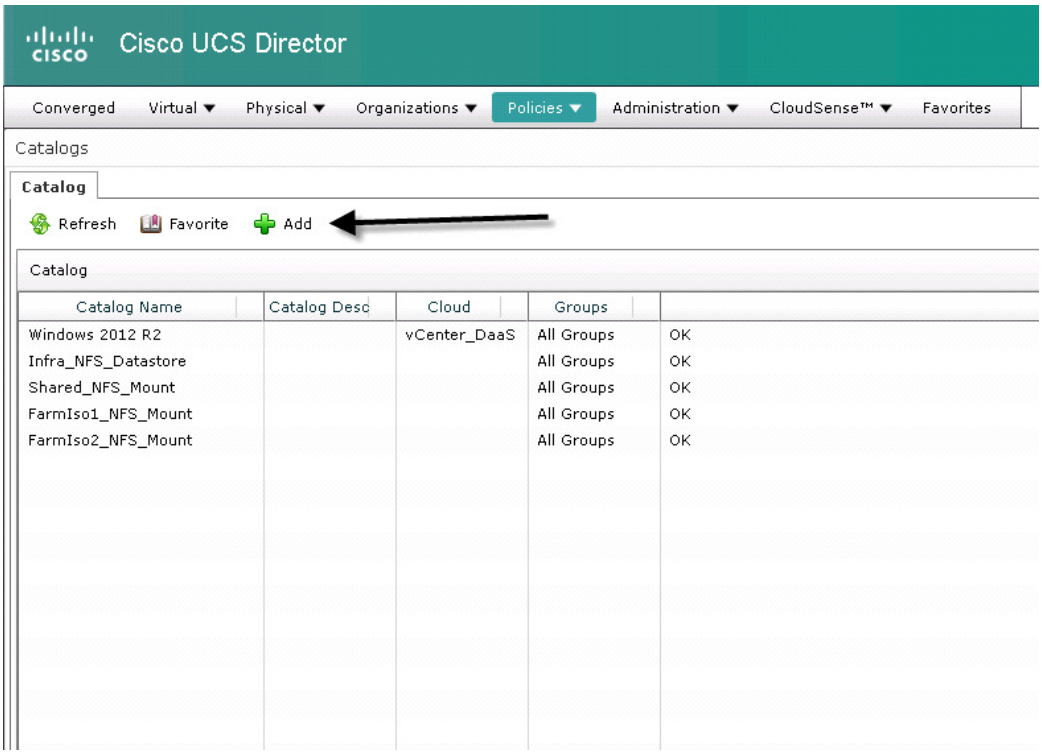

3. Enter a Catalog Name
   a. Catalog Type is ‘Standard’
   b. Select VM Operating System Type
   c. Click ‘All Groups’
   d. Select the Vcenter Cloud created earlier
   e. Select VM image Template from Vcenter

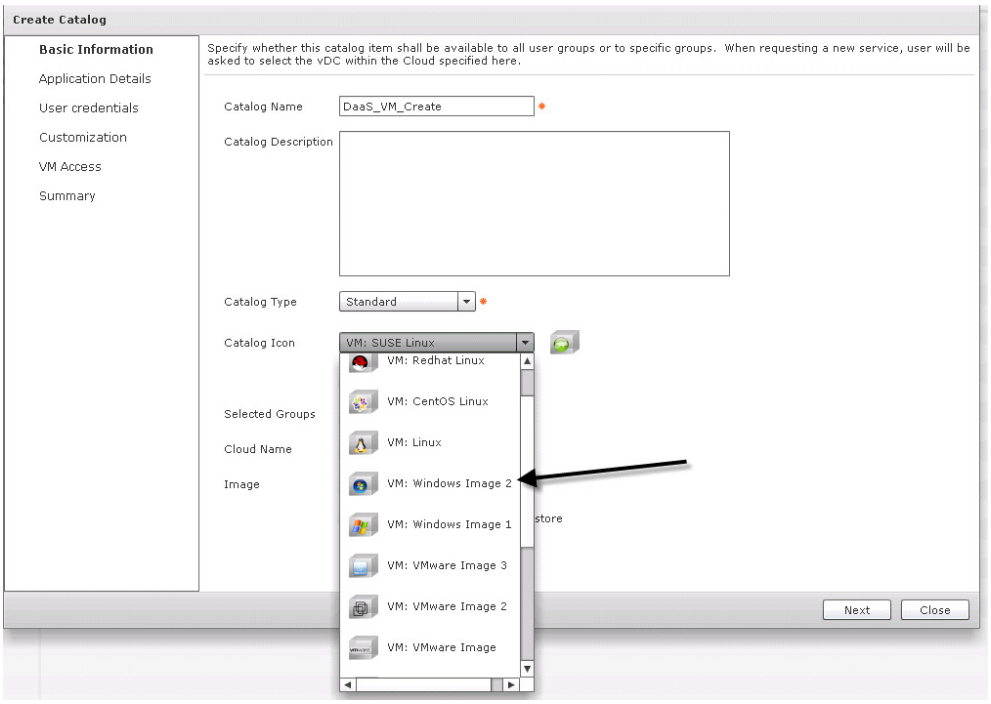

4. Click ‘Next’

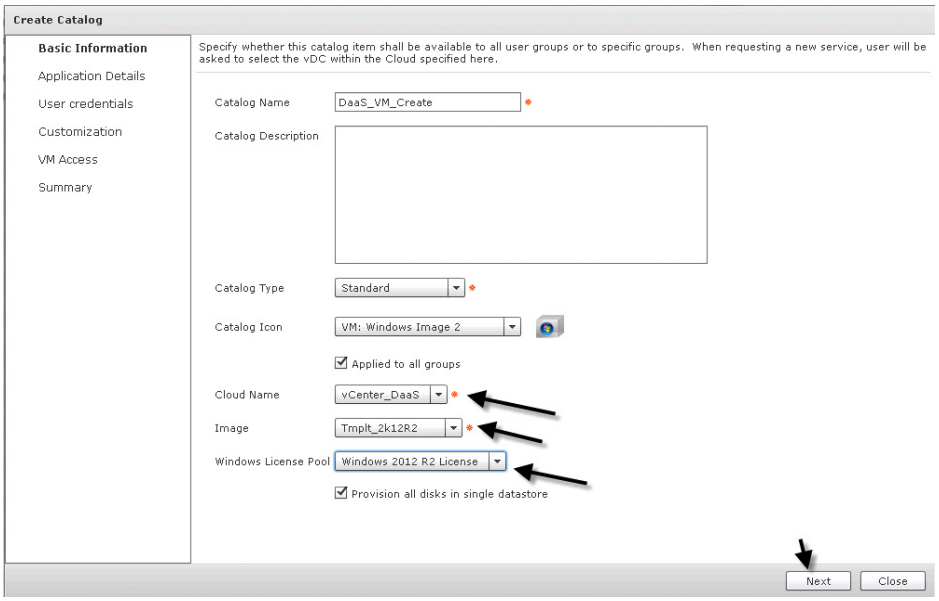

5. In this use case we used a Generic VM with no special settings for applications.

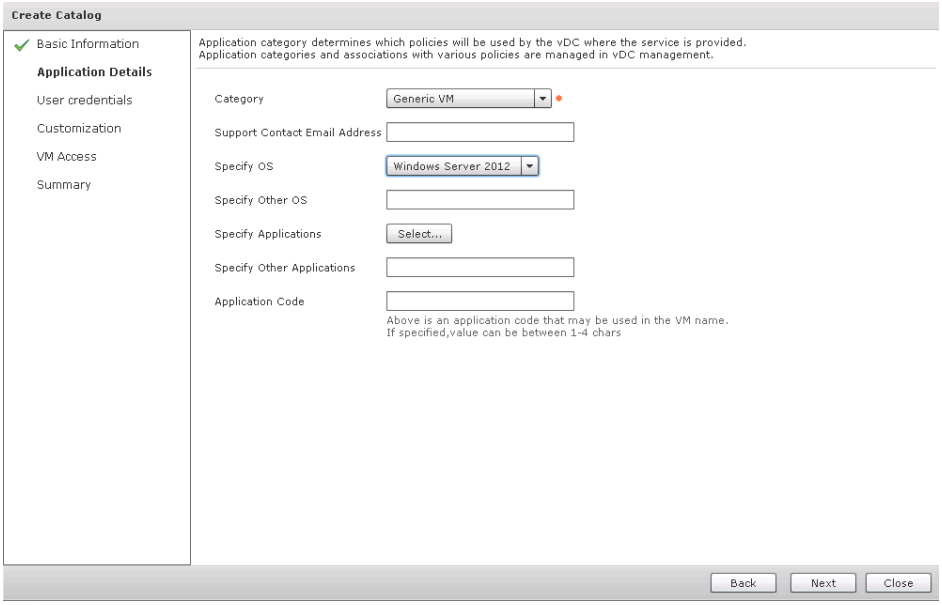

6. Click Next through credential options

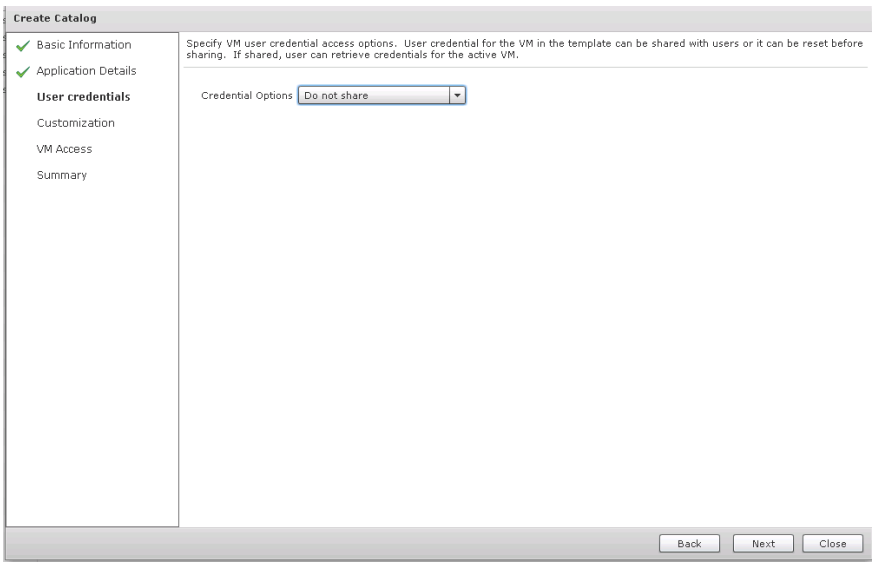

7. Click Next through the defaults.

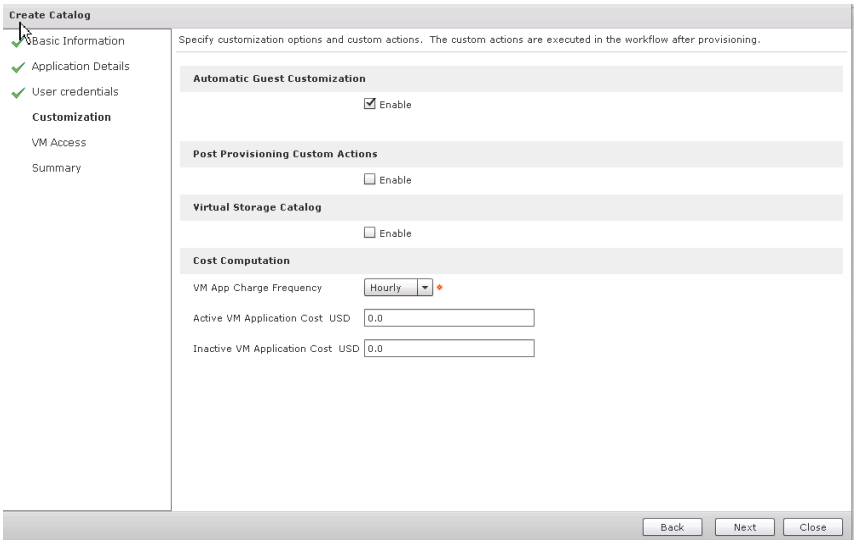

8. Click Next through the defaults

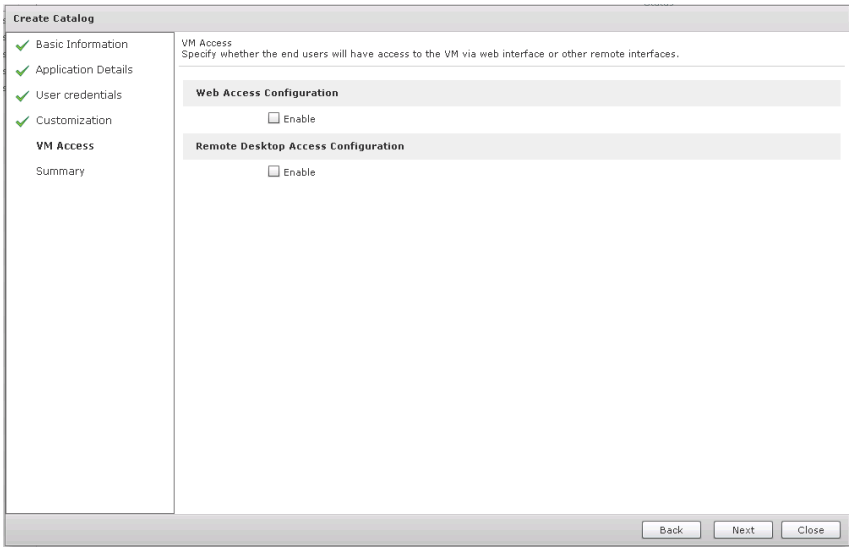

9. Click Submit to create Catalog

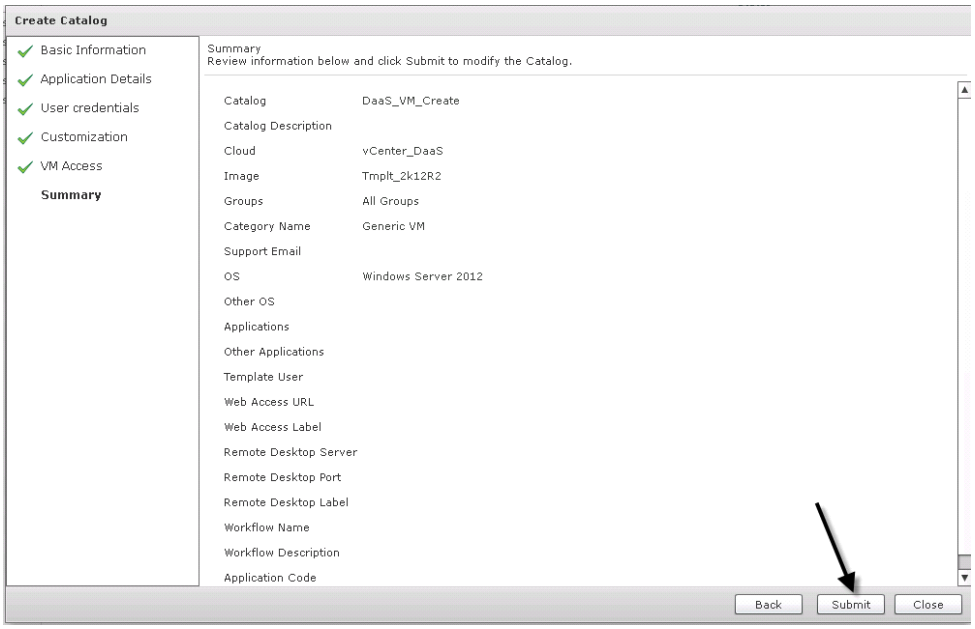

Create vDC (Virtual Data Center) Required for VM Provisioning Workflows

Before utilizing the catalog for provisioning virtual machines we created in the prior steps, we need to create vDCs or Virtual Data Centers for each tenant model used. Virtual Data Center is made up of a collection of policies that direct the virtual machines to be properly configured, named and placed when the provisioning workflows are executed. The policies we used for this project were System, Computing, Network and Storage. We did not utilize a Cost Model Policy in this project.

Create System Policy for vDC

1. Click ‘Service Delivery’ under the ‘Policies’ Tab

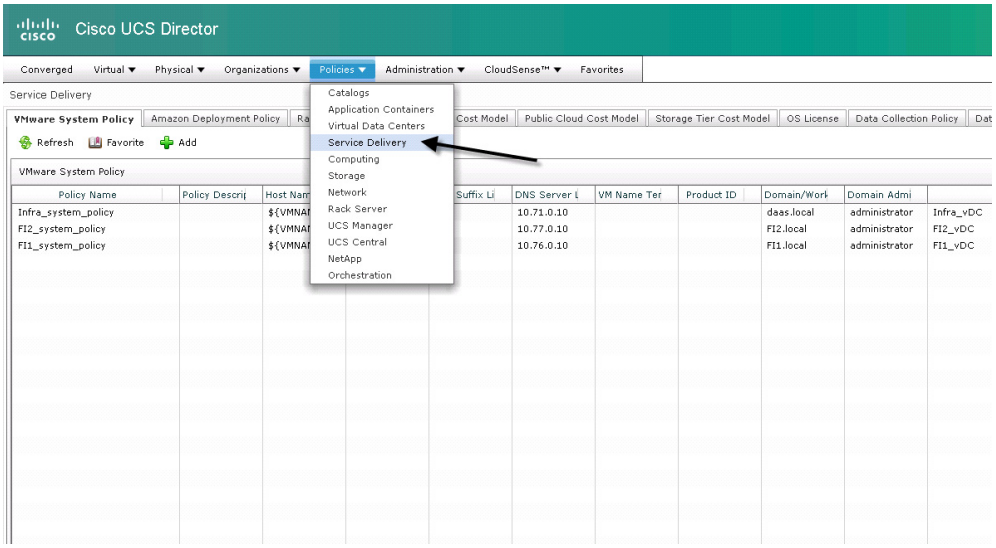

2. Click ‘Add’ to create policy.

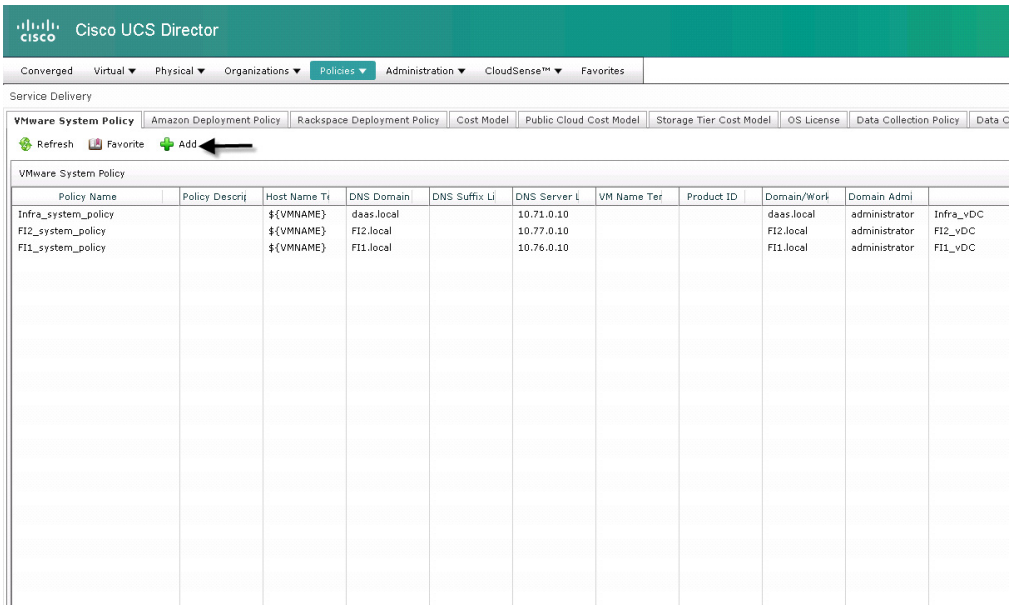

3. In the System Policy Information box we defined the following fields
   a. Policy Name
   b. Host Name Template
   c. DNS Domain
   d. Time Zone
   e. DNS server IP
   f. VM Image Type set to Windows and Linux

g. Define Windows License Model
h. Local Administrator password
i. Domain/workgroup set to ‘Domain’ so machines can auto join the domain
j. Domain name
k. Domain User account that can add machines to the domain
l. Above accounts Password

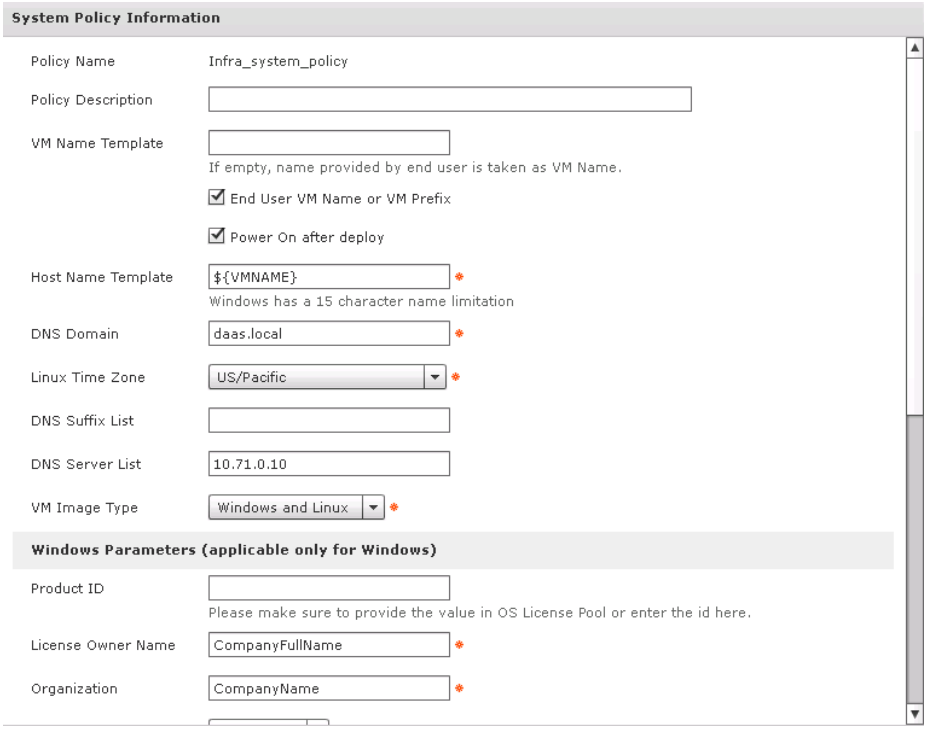

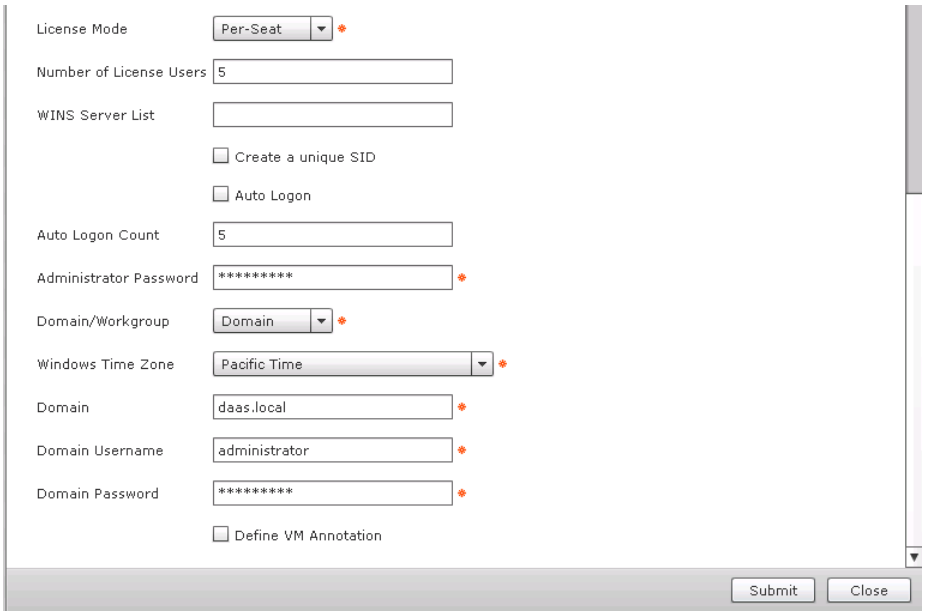

Create a Computing Policy for vDC

1. Under 'Policies' select 'Computing'

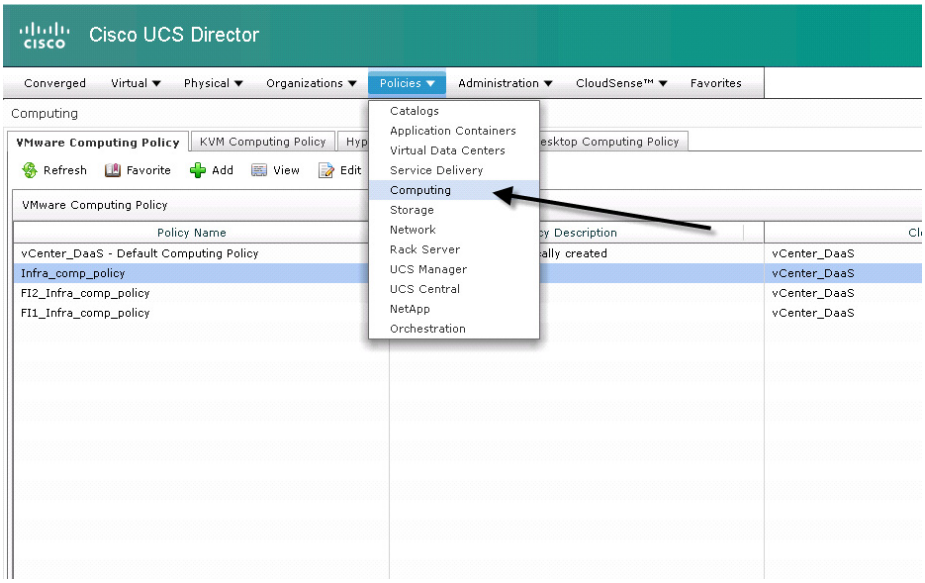

2. Click 'Add'

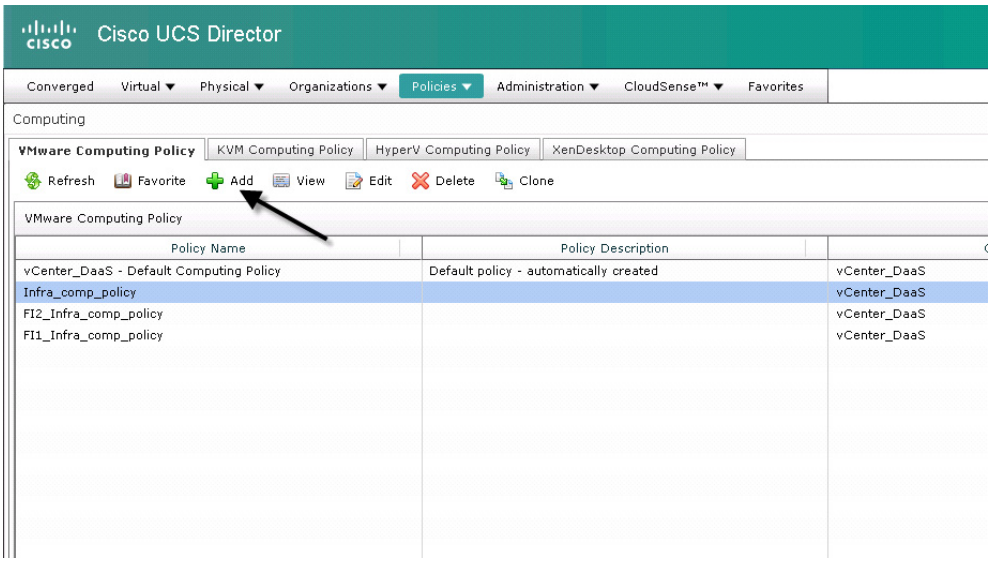

3. When creating the computing policy, we need to provide a name, specify our Vcenter Cloud created earlier and select the cluster that this VM will use. In our case, we defined a separate cluster per tenant model. Lastly we provided disk resizing values to allow us to resize VMs during provisioning. Once these parameters are defined, click 'Save' to create the System Policy.

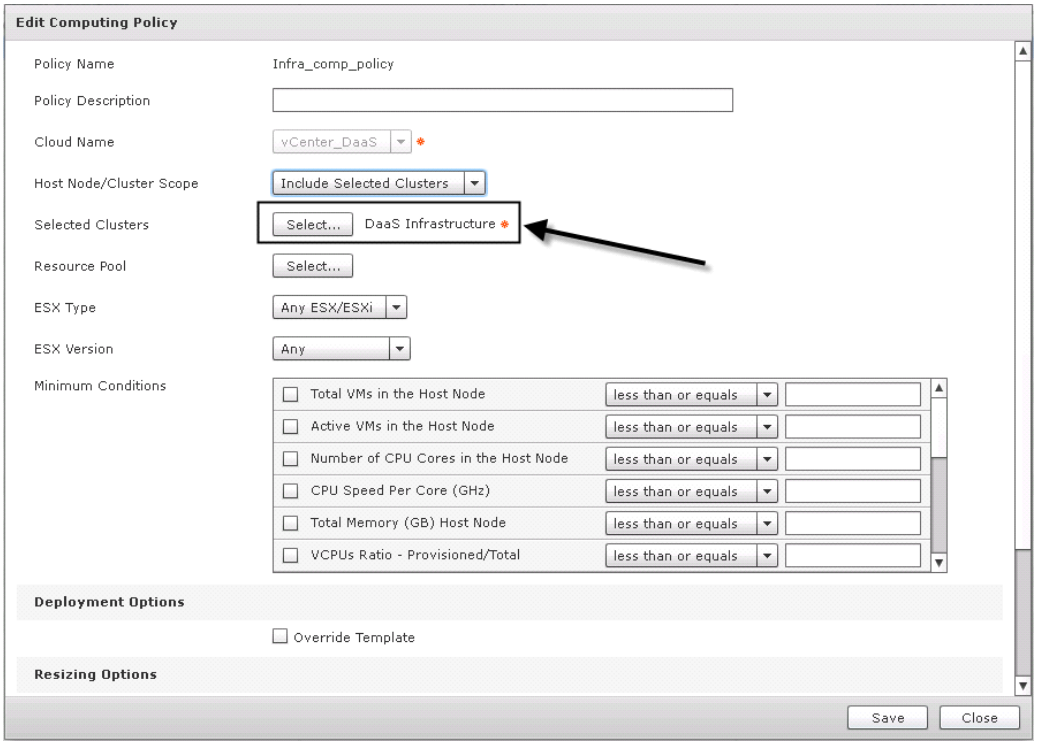

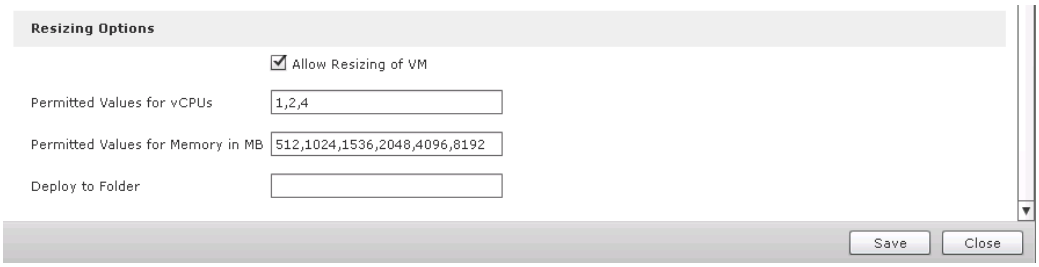

Create Network Policy for vDC

1. Select ‘Network’ under ‘Policies’ Tab

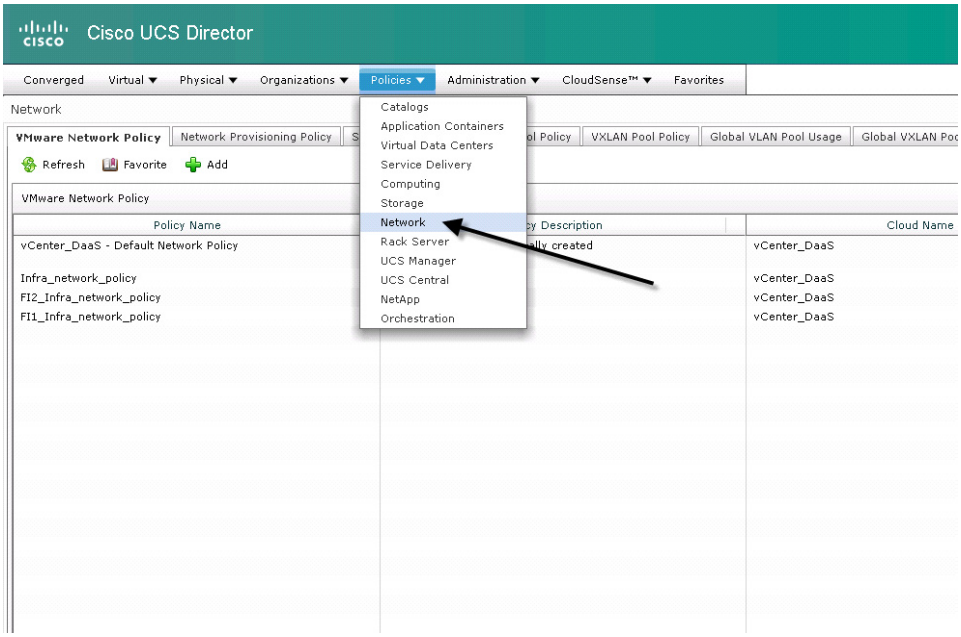

2. Click ‘Add’

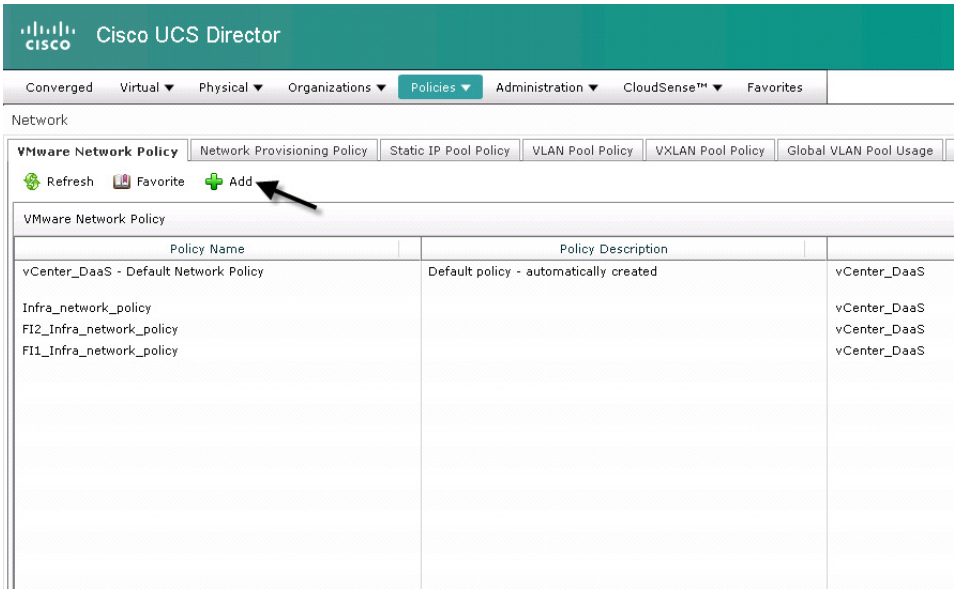

3. Select the Vcenter Cloud under Cloud Name, then click the '+' to add VM Networks

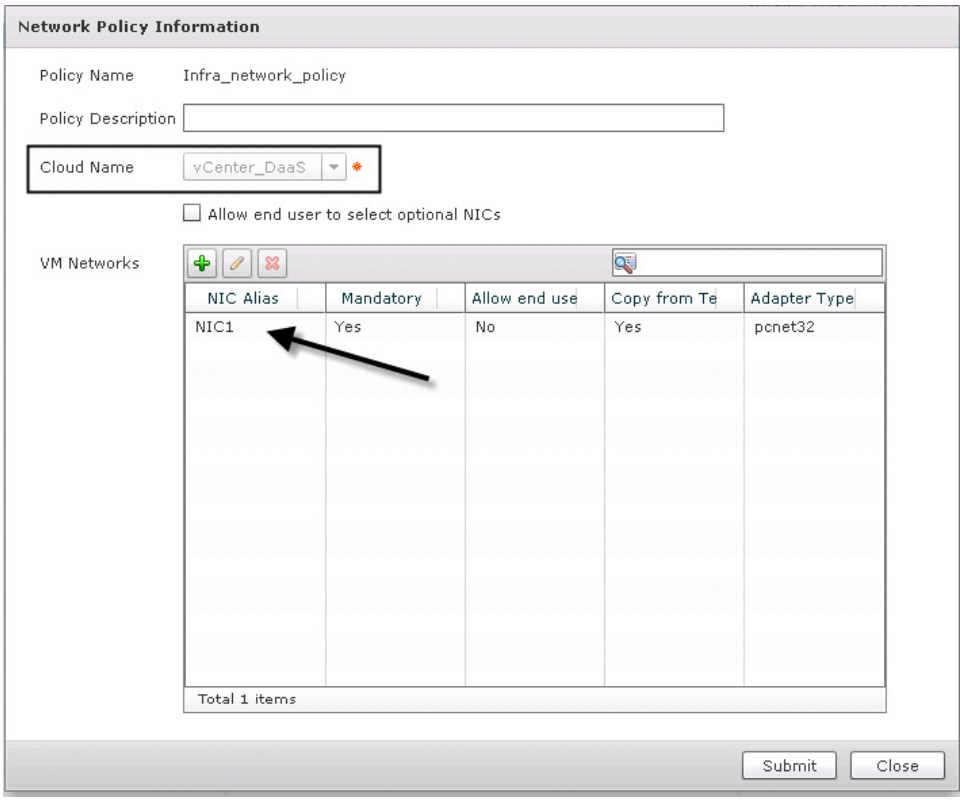

4. When adding VM Network, click the '+' button to add port groups found in the Vcenter Cloud.

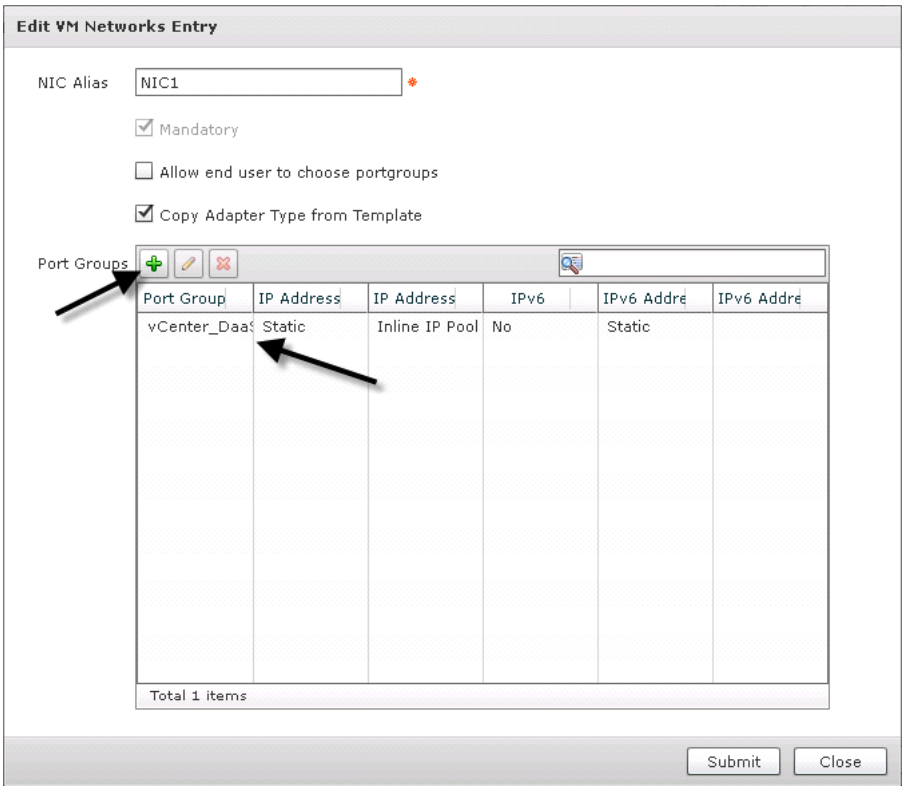

5. For Port Group Name, click ‘Select’ and find Vcenter Port group in the list.

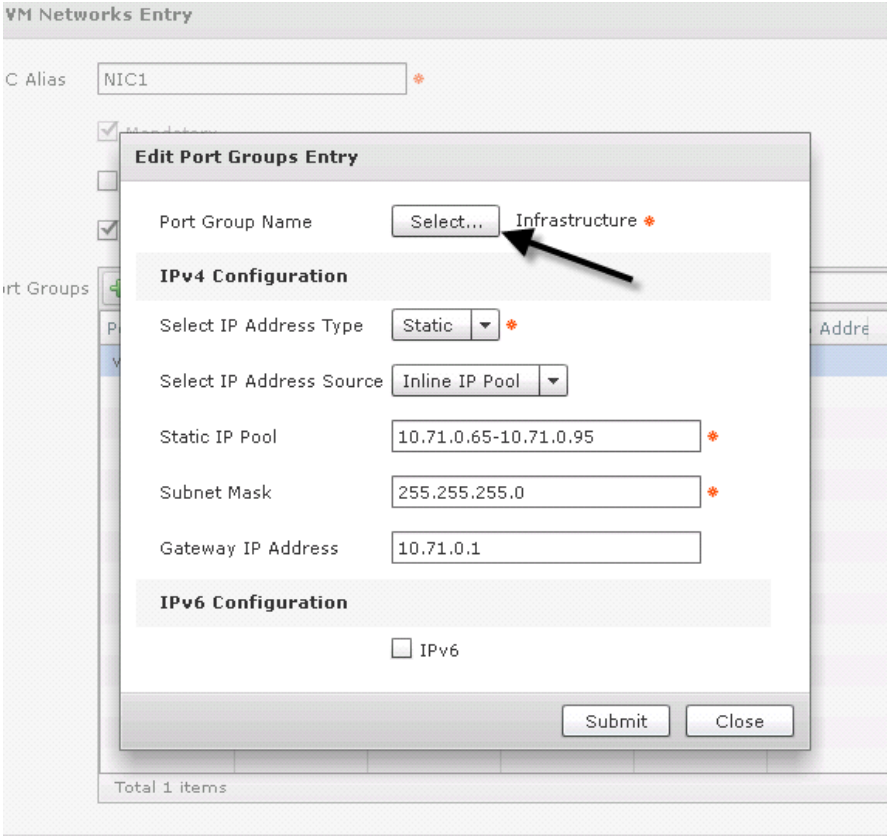

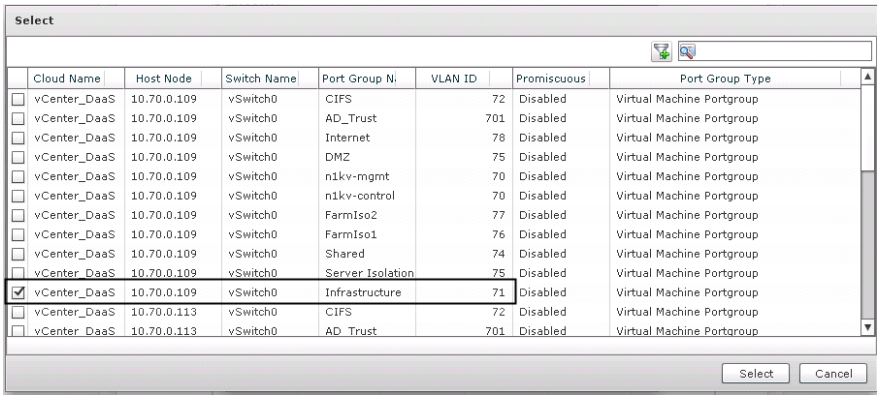

6. Define an Inline IP Pool to allocate a block of IPs for VMs to be configured with. We also defined the Subnet Mask and Default Gateway the VMs will use.

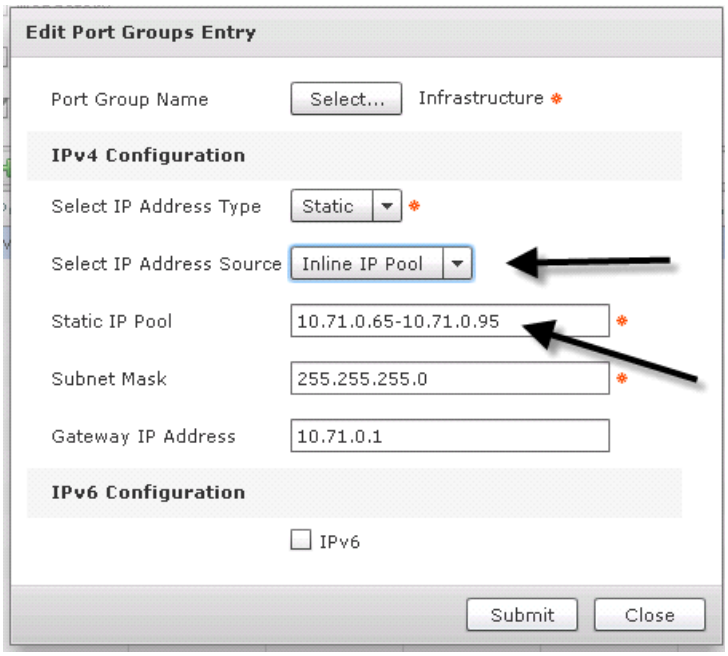

Create a Storage Policy for vDC

1. Select ‘Storage’ under ‘Policies’ Tab

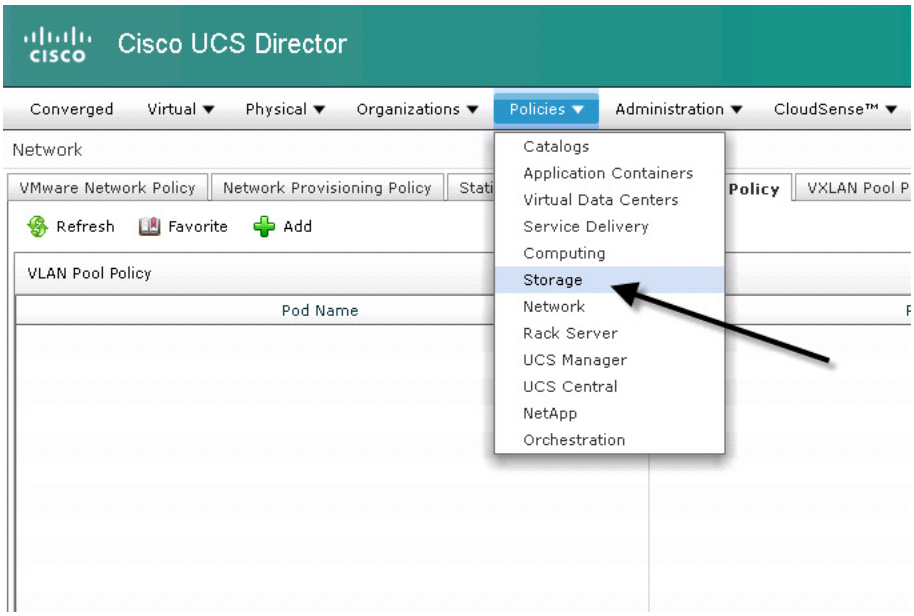

2. Select the ‘VMware Storage Policy’ Tab and then click ‘Add’

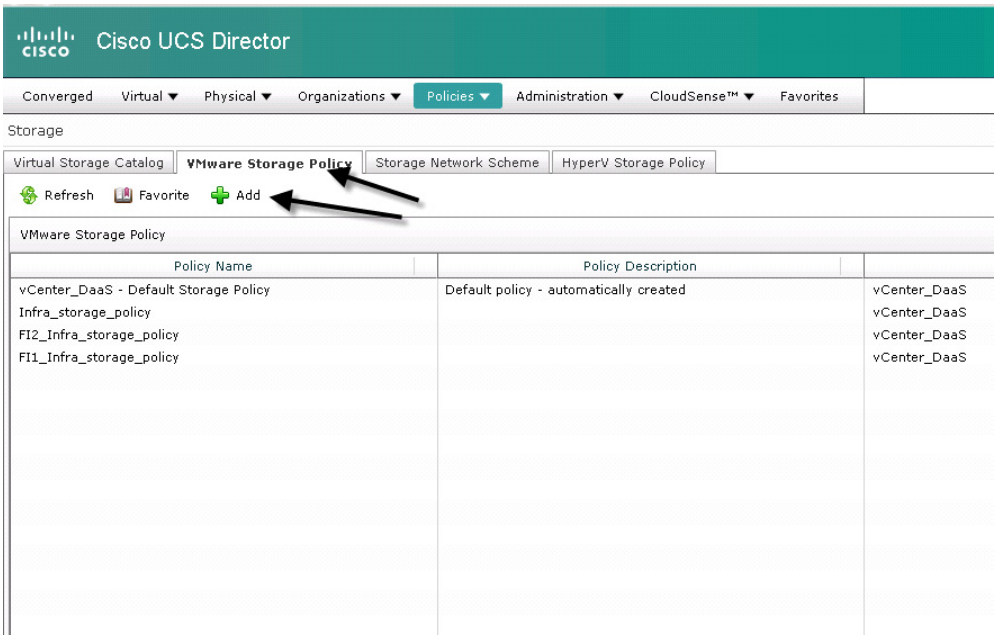

3. In Storage Resource Allocation Policy we define the Policy Name, Select the Vcenter Cloud we created earlier, and then select the Data Stores that we added to the Vcenter environment in our NFS Provisioning Workflows. We used NFS for this project. We also defined parameters to resize VM disk size upon provisioning.

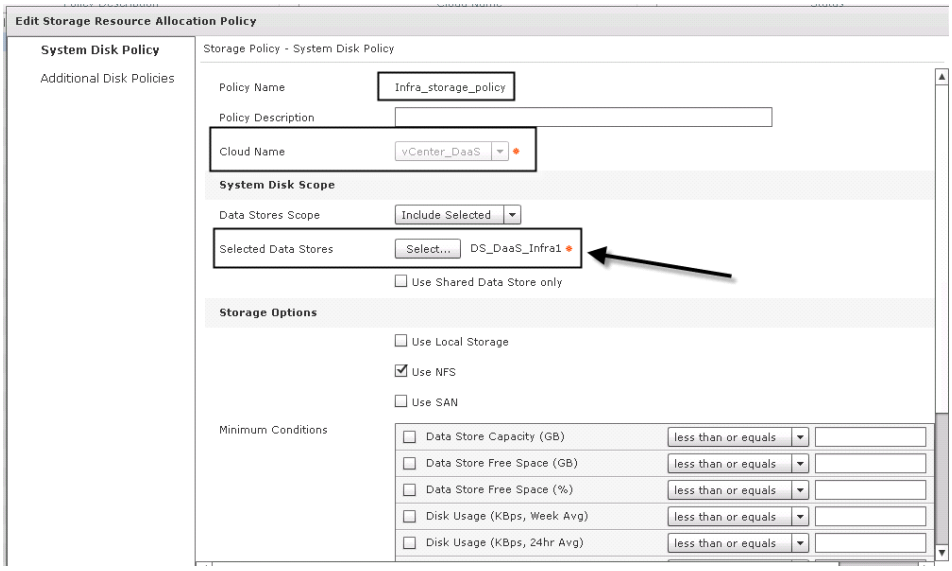

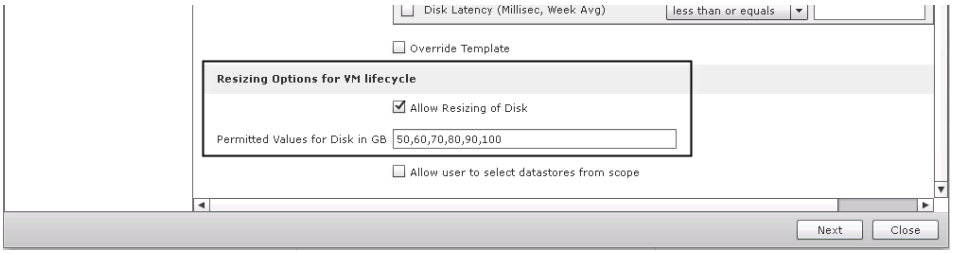

4. Image shows the Datastore selection screen.

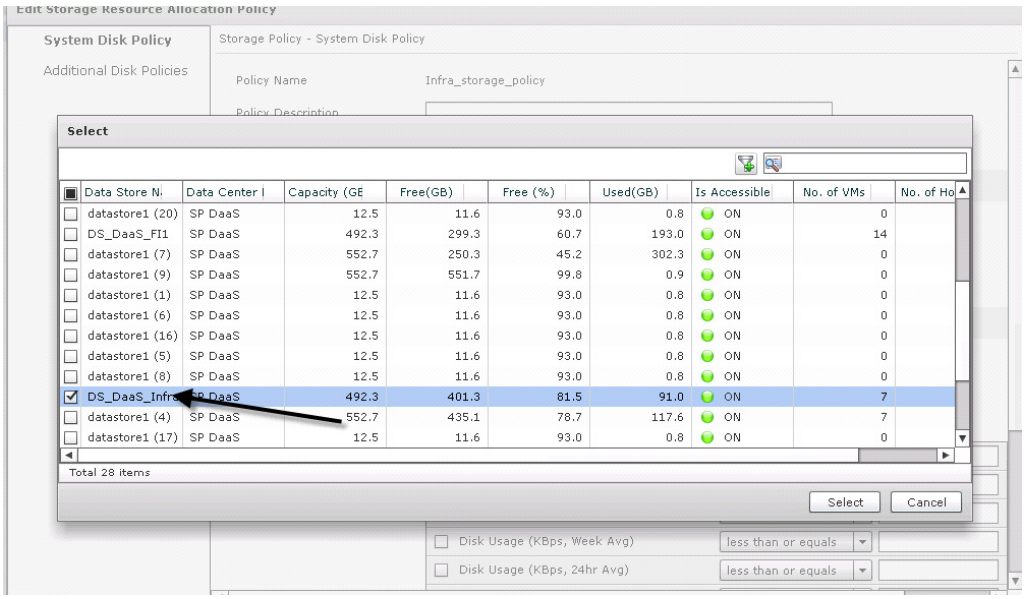

Cisco Unified Computing System Configuration

This section describes the Cisco UCS configuration that was done as part of the infrastructure build out. The racking, power and installation of the chassis are described in the install guide (see www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html) and it is beyond the scope of this document. More details on each step can be found in the following documents:

- Cisco UCS Manager Configuration Guides – GUI and Command Line Interface (CLI)
- Cisco UCS Manager - Configuration Guides - Cisco

Base Cisco UCS System Configuration

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

1. Bring up the Fabric Interconnect (FI) 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. Once this is done all 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, 4 or 8 links per IO Module, depending on your application bandwidth requirement. We connected four links to each module.

2. Connect using your favorite browser to the Virtual IP and launch the UCS-Manager. The Java based UCSM will let you do everything that you could do from the CLI. We will highlight the GUI methodology here.

3. Check the firmware on the system and see if it is current. Visit: Download Software for Cisco UCS Infrastructure and UCS Manager Software to download the most current UCS Infrastructure and UCS Manager software. Use the UCS Manager Equipment tab in the left pane, then the Firmware Management tab in the right pane and 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 2.2(2c).

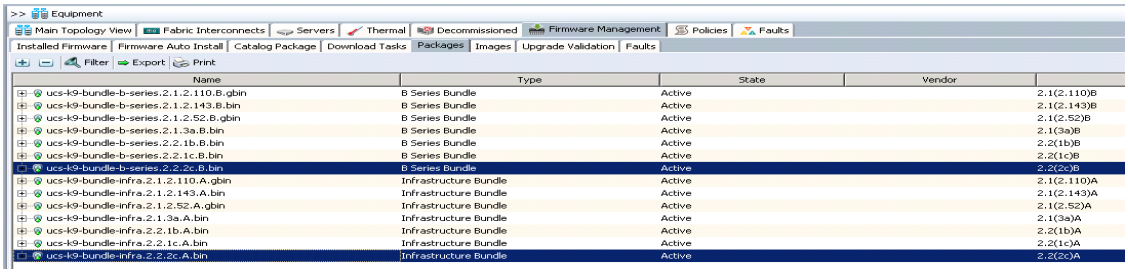

Name	Type	State	Vendor	
ucs-k9-bundle-b-series.2.1.2.110.B.gbin	B Series Bundle	Active		2.1(2.110)B
ucs-k9-bundle-b-series.2.1.2.143.B.bin	B Series Bundle	Active		2.1(2.143)B
ucs-k9-bundle-b-series.2.1.2.52.B.gbin	B Series Bundle	Active		2.1(2.52)B
ucs-k9-bundle-b-series.2.1.3a.B.bin	B Series Bundle	Active		2.1(3a)B
ucs-k9-bundle-b-series.2.2.1b.B.bin	B Series Bundle	Active		2.2(1b)B
ucs-k9-bundle-b-series.2.2.1c.B.bin	B Series Bundle	Active		2.2(1c)B
ucs-k9-bundle-b-series.2.2.2c.B.bin	B Series Bundle	Active		2.2(2c)B
ucs-k9-bundle-infra.2.1.2.110.A.gbin	Infrastructure Bundle	Active		2.1(2.110)A
ucs-k9-bundle-infra.2.1.2.143.A.bin	Infrastructure Bundle	Active		2.1(2.143)A
ucs-k9-bundle-infra.2.1.2.52.A.gbin	Infrastructure Bundle	Active		2.1(2.52)A
ucs-k9-bundle-infra.2.1.3a.A.bin	Infrastructure Bundle	Active		2.1(3a)A
ucs-k9-bundle-infra.2.2.1b.A.bin	Infrastructure Bundle	Active		2.2(1b)A
ucs-k9-bundle-infra.2.2.1c.A.bin	Infrastructure Bundle	Active		2.2(1c)A
ucs-k9-bundle-infra.2.2.2c.A.bin	Infrastructure Bundle	Active		2.2(2c)A

4. If the firmware is not current, follow the installation and upgrade guide to upgrade the UCS Manager firmware. We will use UCS Policy in Service Profiles later in this document to update all UCS components in the solution.

Note The Bios and Board Controller version numbers do not track the IO Module, Adapter, nor CIMC controller version numbers in the packages.

5. Configure and enable the server ports on the FI. These are the ports that will connect the chassis to the FIs.

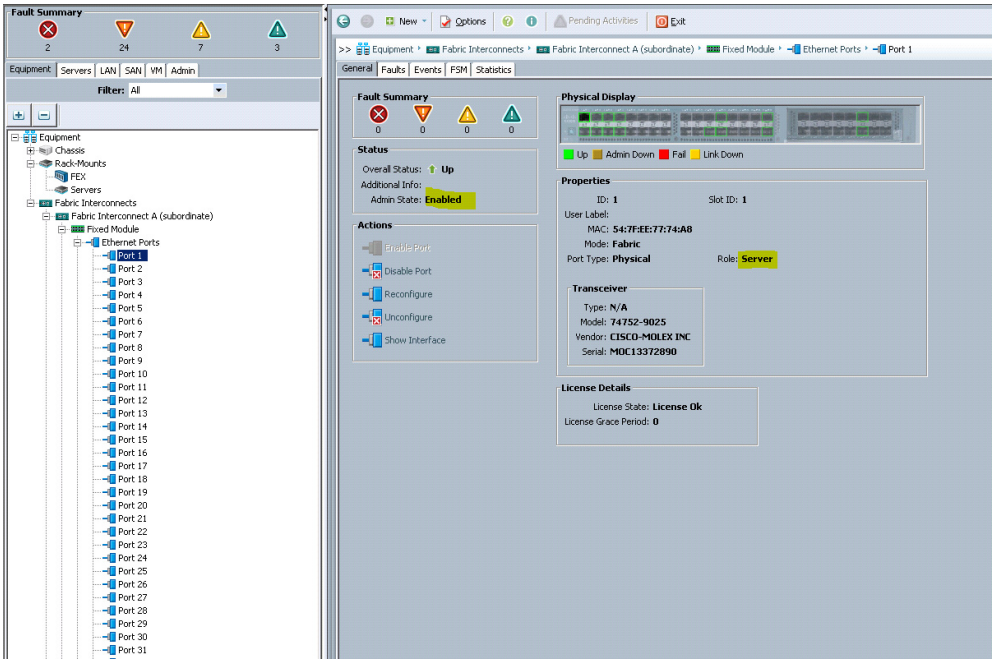

6. Configure and enable uplink Ethernet ports:

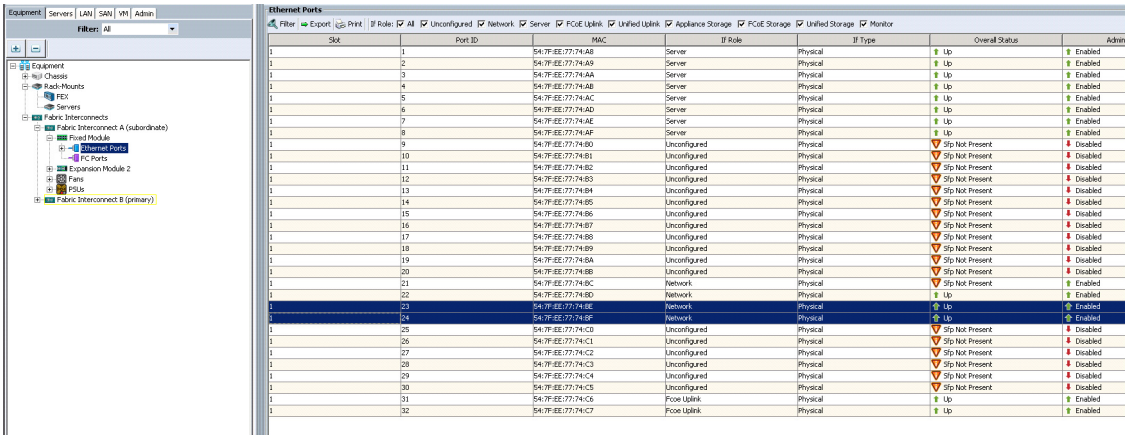

a. On the LAN tab in the Navigator pane, configure the required Port Channels and Uplink Interfaces on both Fabric Interconnects.

b. Configure the ethernet uplink port channels using the ethernet Network ports configured above

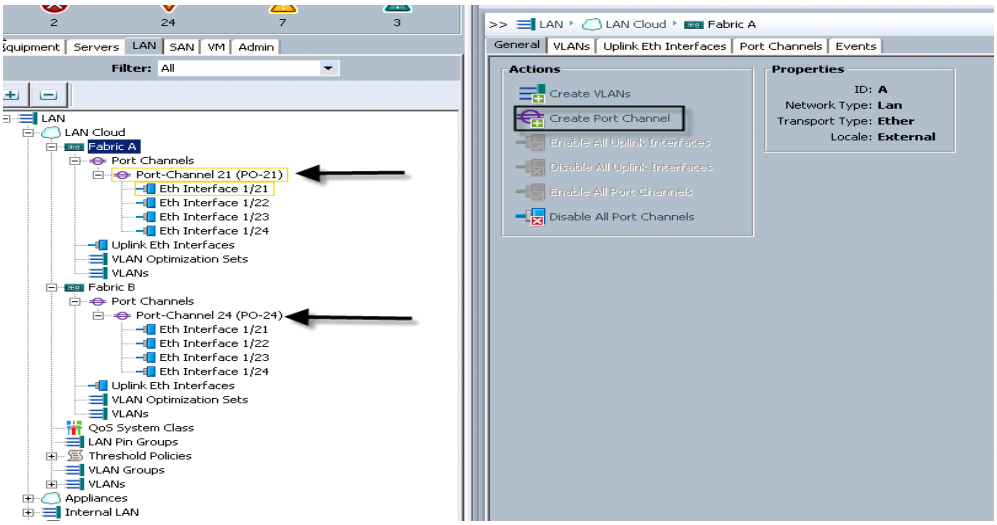

c. 5b Configure the FC uplink ports:

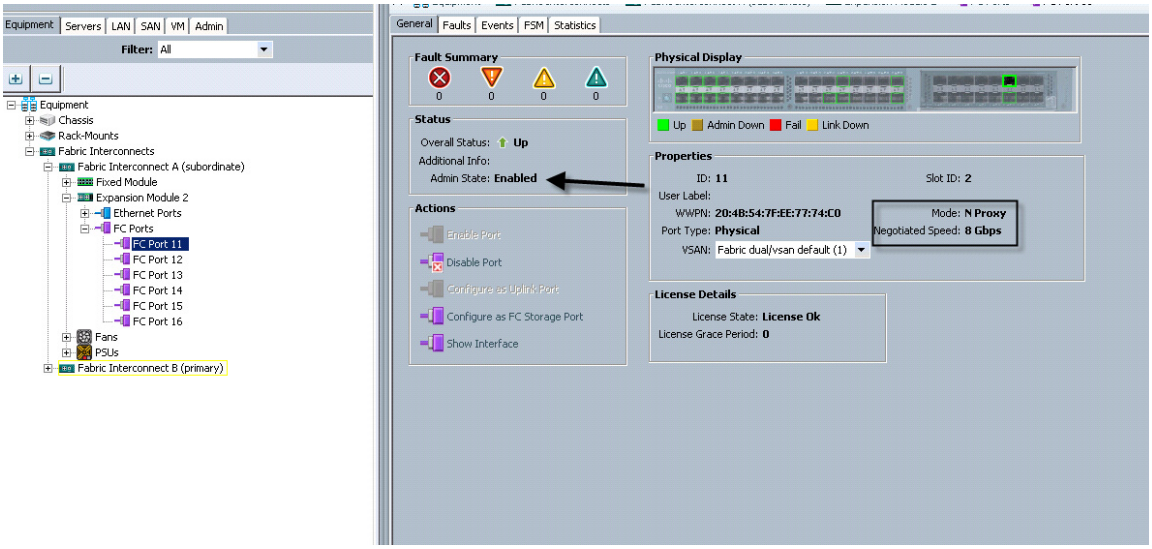

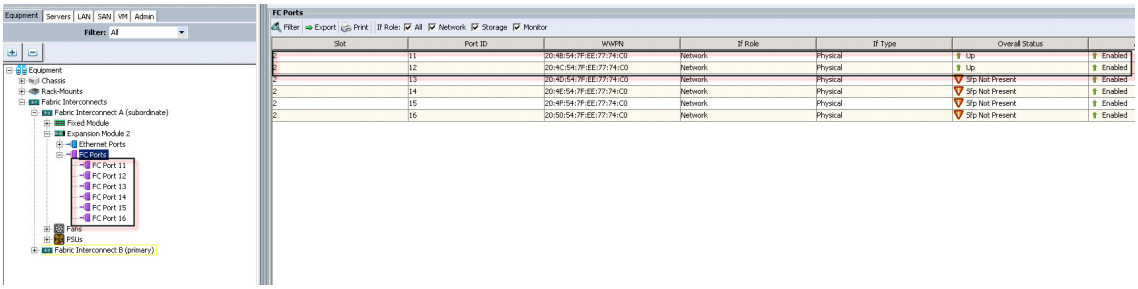

7. On the Equipment tab, expand the Chassis node in the left pane, the click on each chassis in the left pane, then click Acknowledge Chassis in the right pane to bring the chassis online and enable blade discovery.

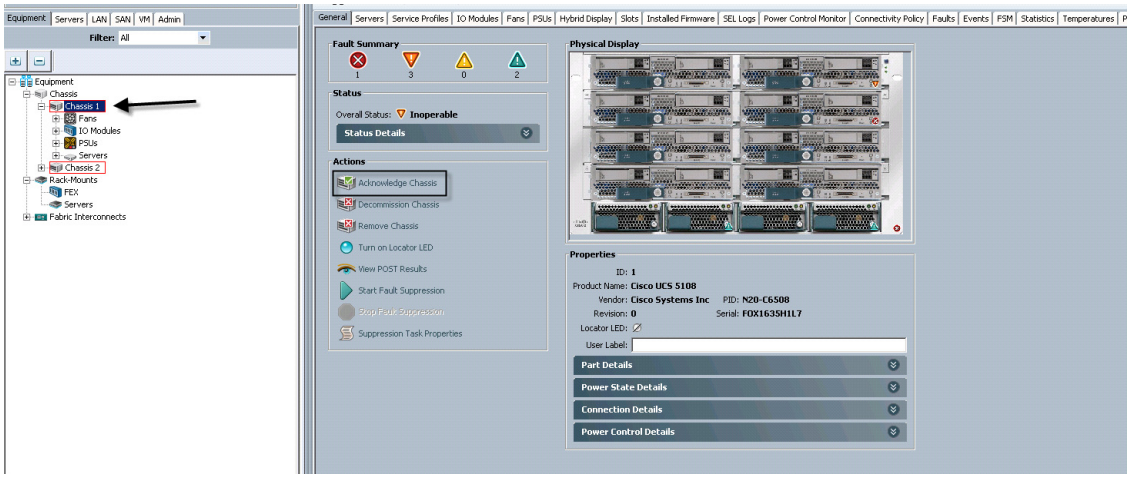

8. Use the Admin tab in the left pane, to configure logging, users and authentication, key management, communications, statistics, time zone and NTP services, and Licensing. Configuring your Management IP Pool (which provides IP based access to the KVM of each UCS Blade Server,) Time zone Management (including NTP time source(s)) and uploading your license files are critical steps in the process.

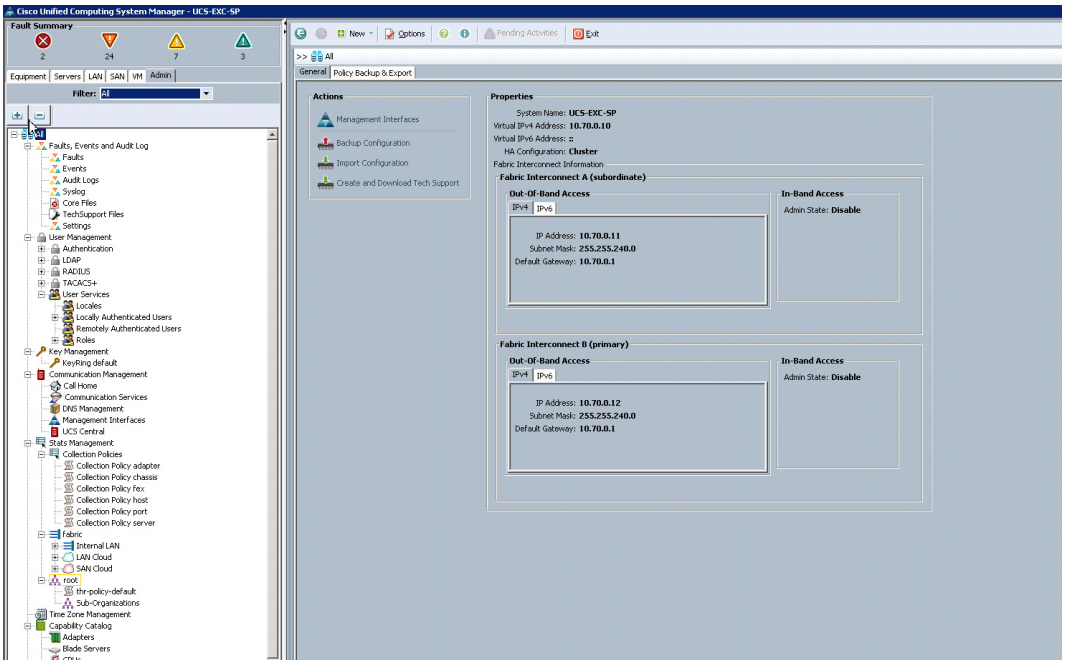

9. Create all the pools: MAC pool, UUID pool, WWNN pool, WWPN Pool, External Management IP Address Pool and Server pools

10. From the LAN tab in the navigator, under the Pools node, we created a MAC address pool of sufficient size for the environment. In this project, we created a single pool with two address ranges for expandability.

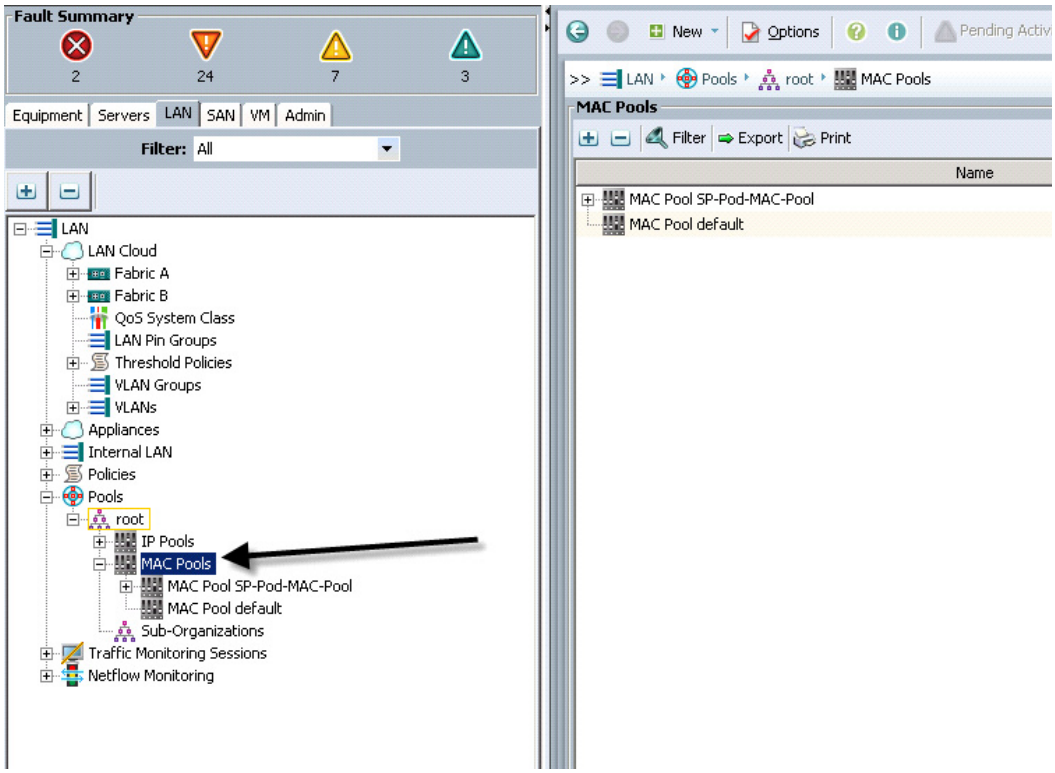

11. For Fiber Channel Connectivity, WWN and WWPN pools must be created from the SAN tab in the navigator pane, in the Pools node:

WWPN

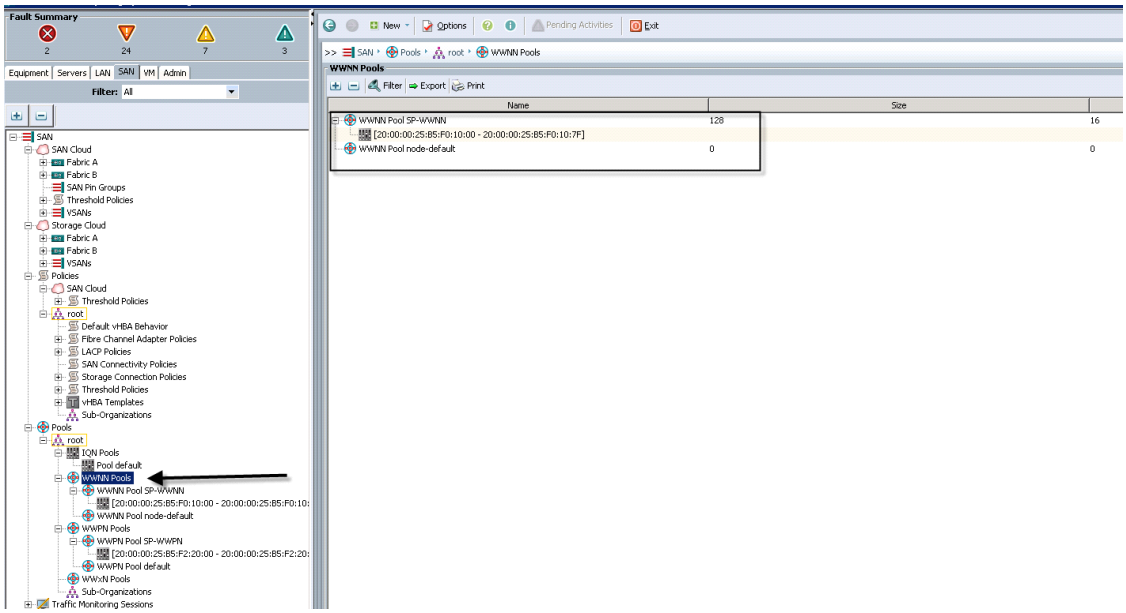

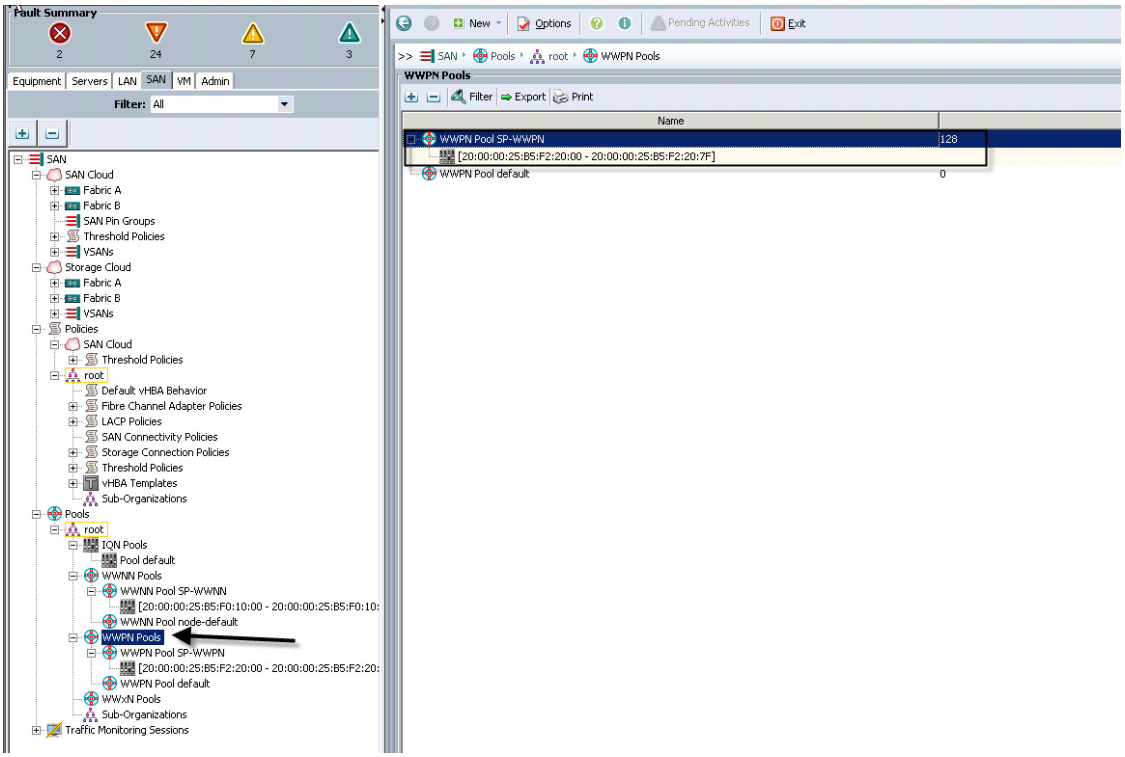

12. For this project, we used a single VSAN, the default VSAN with ID 1:

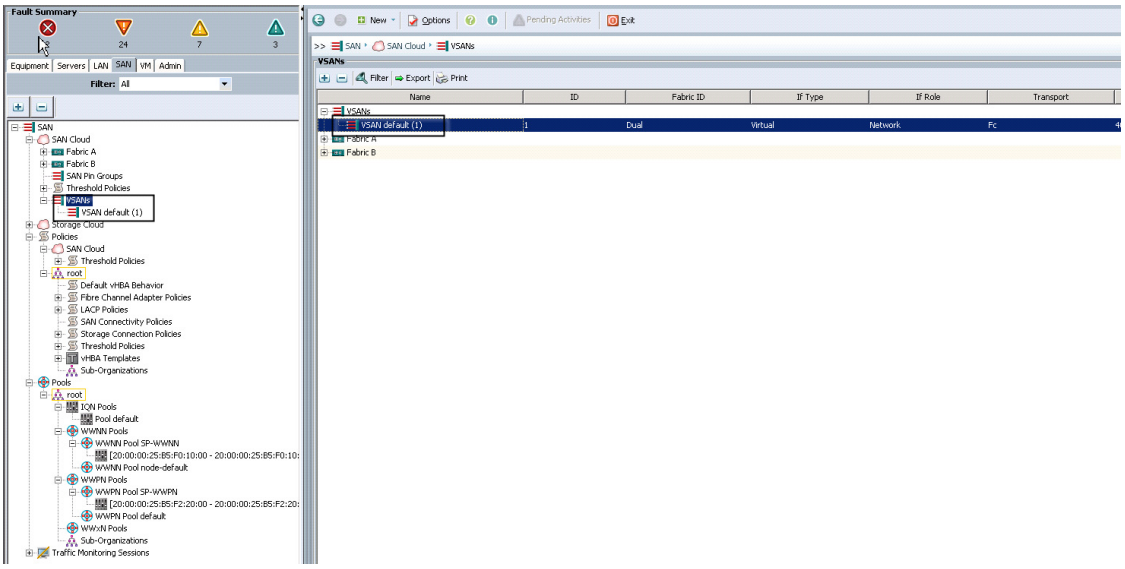

13. From the LAN under the Pools node, we created an External Management IP address pool for use by the Cisco UCS KVM connections to the blade servers in the study.

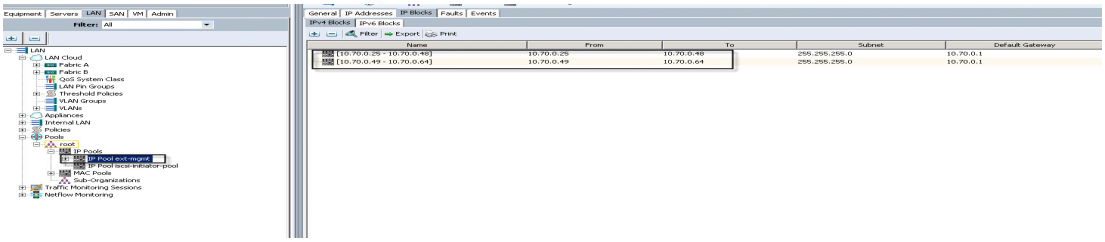

14. The next pool we created is the Server UUID pool. On the Servers tab in the Navigator page under the Pools node we created a single UUID Pool for the test environment. Each UCS Blade Server requires a unique UUID to be assigned by its Service profile.

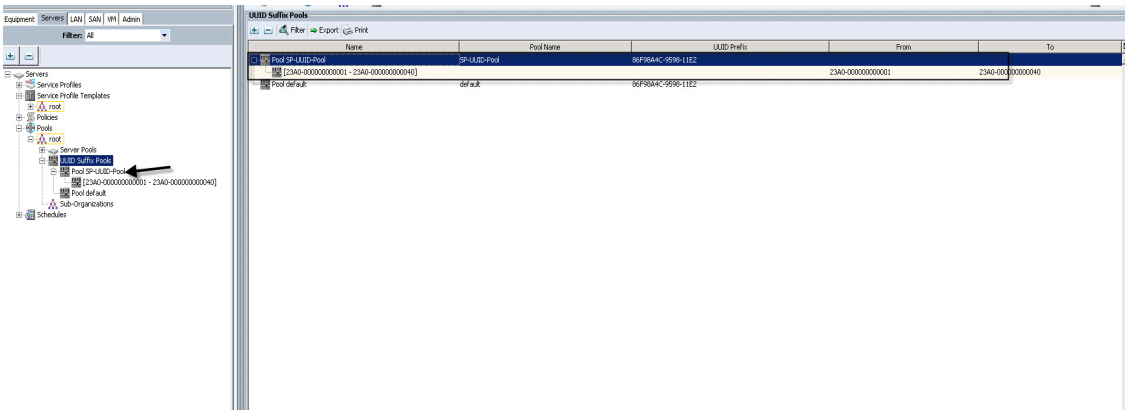

15. We created four Server Pools for use in our Service Profile Templates as selection criteria for automated profile association. Server Pools were created on the Servers tab in the navigation page under the Pools node. Only the pool name was created, no servers were added:

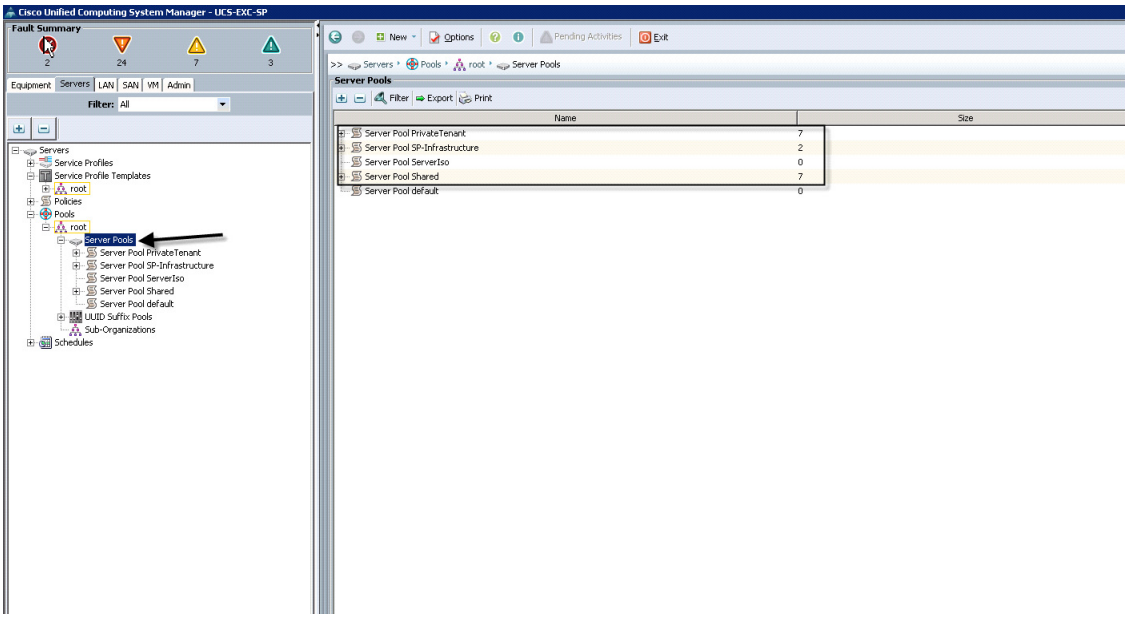

16. We created three Server Pool Policy Qualifications to identify the blade server model, its processor and the amount of RAM onboard for placement into the correct Server Pool using the Service Profile Template. In this case we used slot id to segregate tenant hosts into specific slots. (We could have used a combination of chassis, slot, server model, etc or any combination of those things to make the selection.)

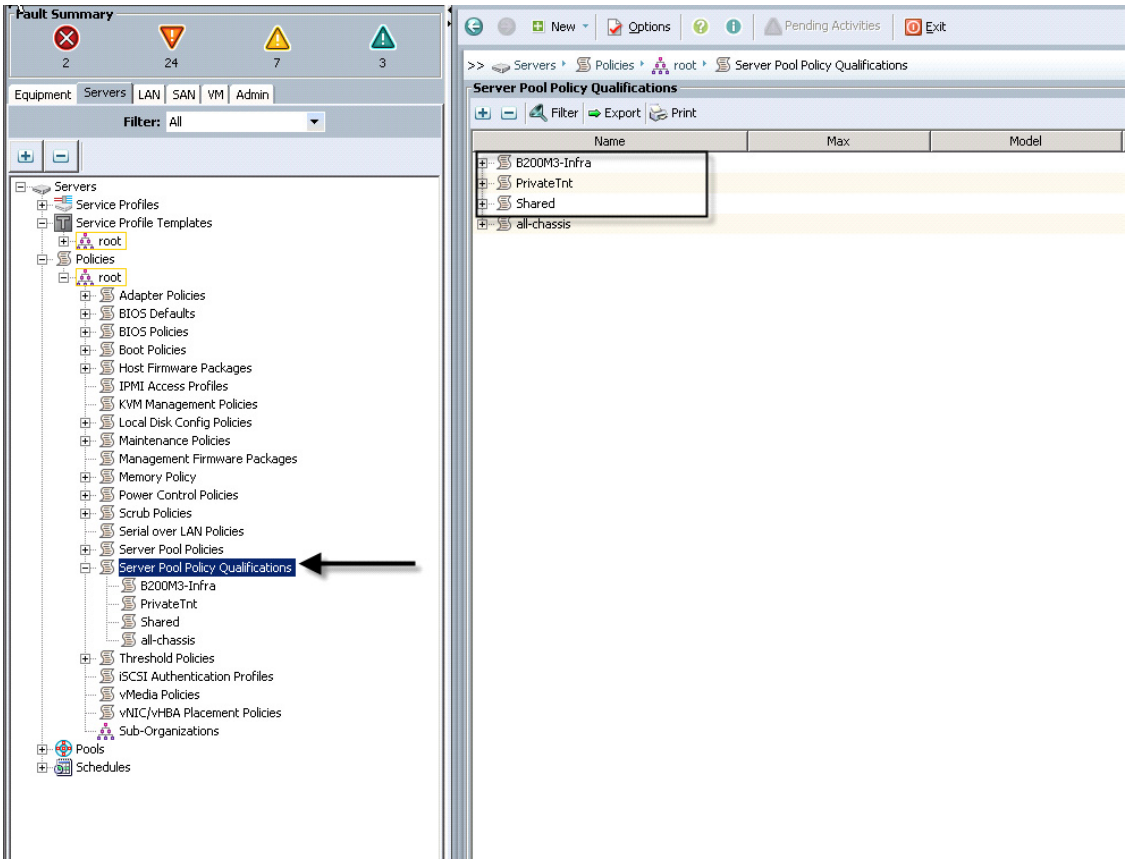

17. The next step in automating the server selection process is to create corresponding Server Pool Policies for each UCS Blade Server configuration, utilizing the Server Pool and Server Pool Policy Qualifications created earlier.

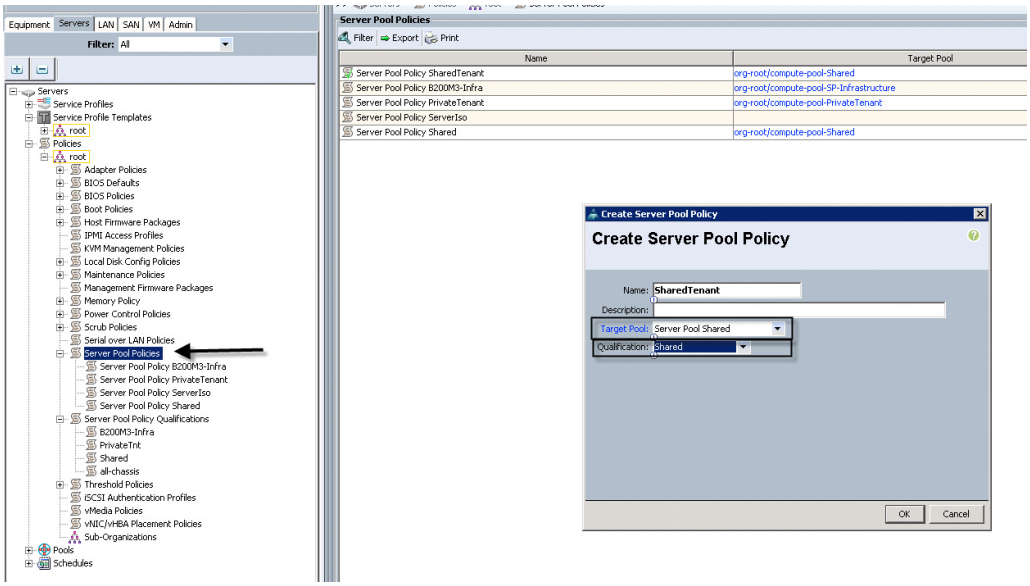

18. To create the policy, right-click the Server Pool Policy node, select Create Server Pool Policy, provide a name, description (optional,) select the Target Pool from the dropdown, the Qualification from the dropdown and click OK. Repeat for each policy to be created.
19. On the LAN tab in the navigator pane, configure the VLANs for the environment:

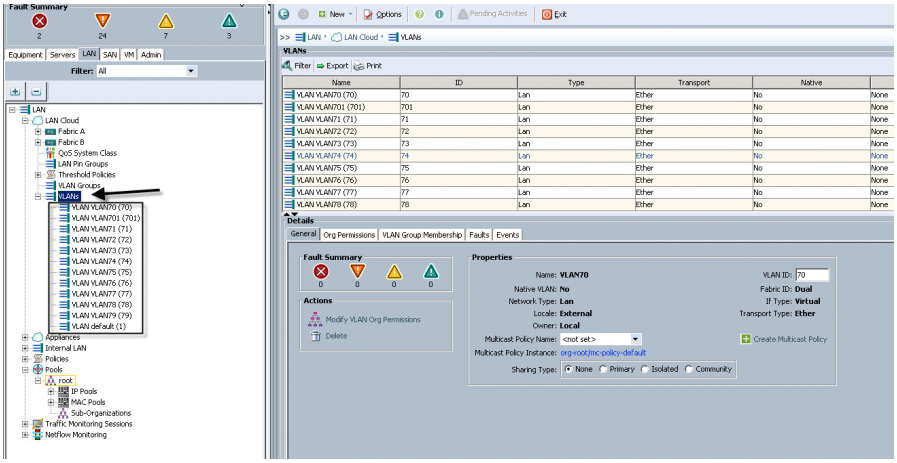

20. In this project we utilized nine VLANs to accommodate our four traffic types and a separate native VLAN for all traffic that was not tagged on the network. The Storage VLANs provided communications for NFS and CIFS storage traffic.
21. On the LAN tab in the navigator pane, under the policies node configure the vNIC templates that will be used in the Service Profiles. In this project, we utilize two virtual NICs per host, one to each Fabric Interconnect for resiliency. QoS is handled by Cisco Nexus 1000V for the VM-Network, so no QoS policy is set on the templates. Both in-band and out-of-band management VLANs are trunked to the eth0 and eth1 vNIC templates. The Default VLAN is 79 and used for PXE booting for UCS Director Baremetal Imaging.

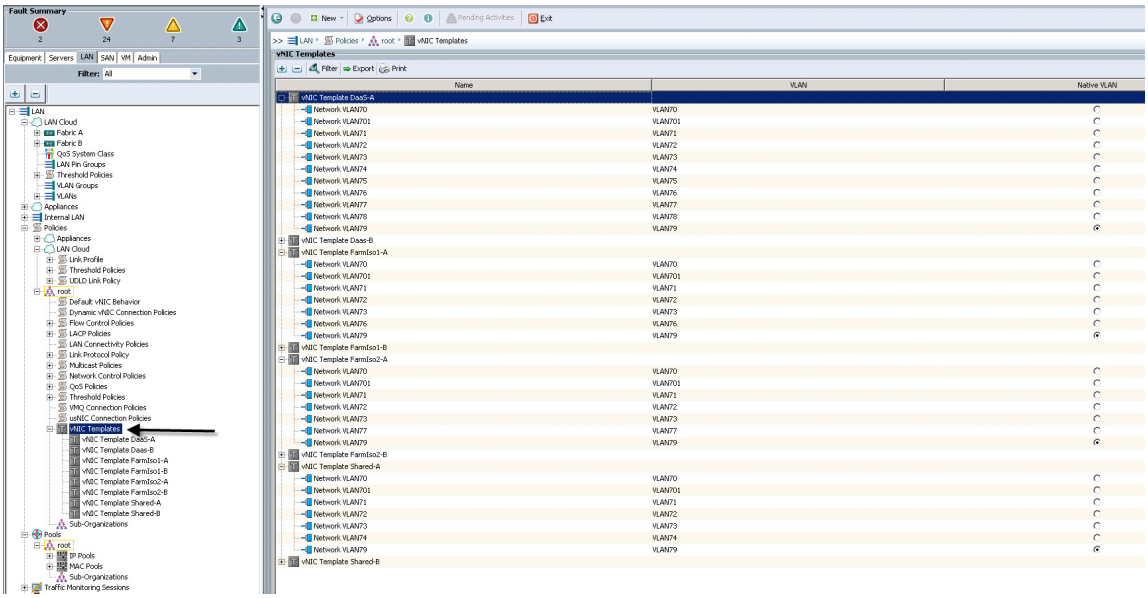

22. To create vNIC templates for eth0 and eth1 on both fabrics, select the Fabric ID, select all VLANs and set the MTU size to 9000, select the MAC Pool, then click OK.
23. Prepare the Perf-Cisco BIOS Policy. From the Server tab, Policies, Root node, right-click the BIOS Policies container and click Create BIOS Policy. Provide the policy name and step through the wizard making the choices indicated on the screen shots below:

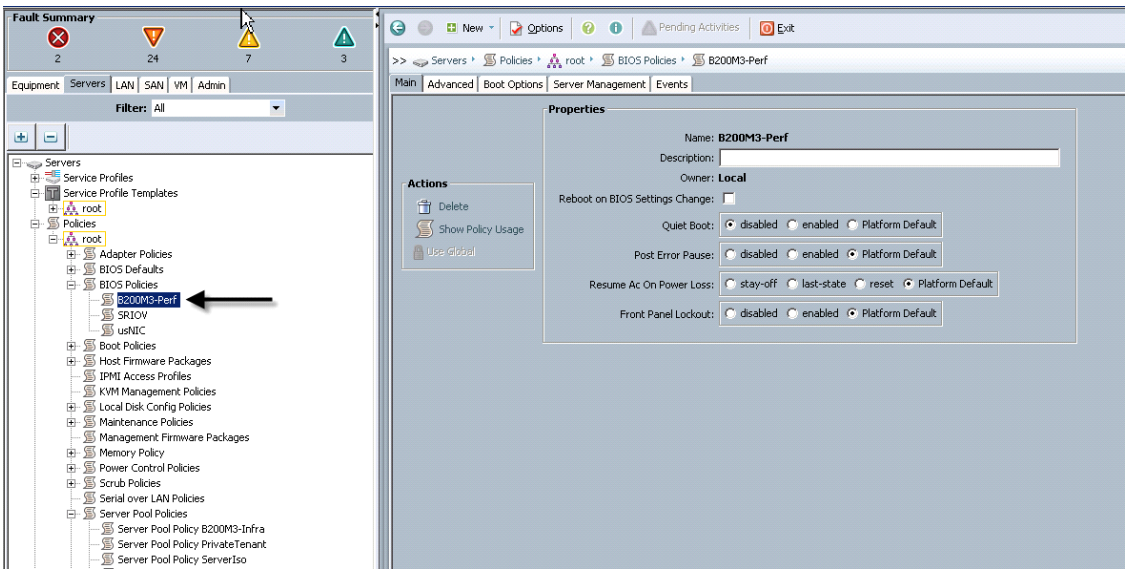

The Advanced Tab Settings

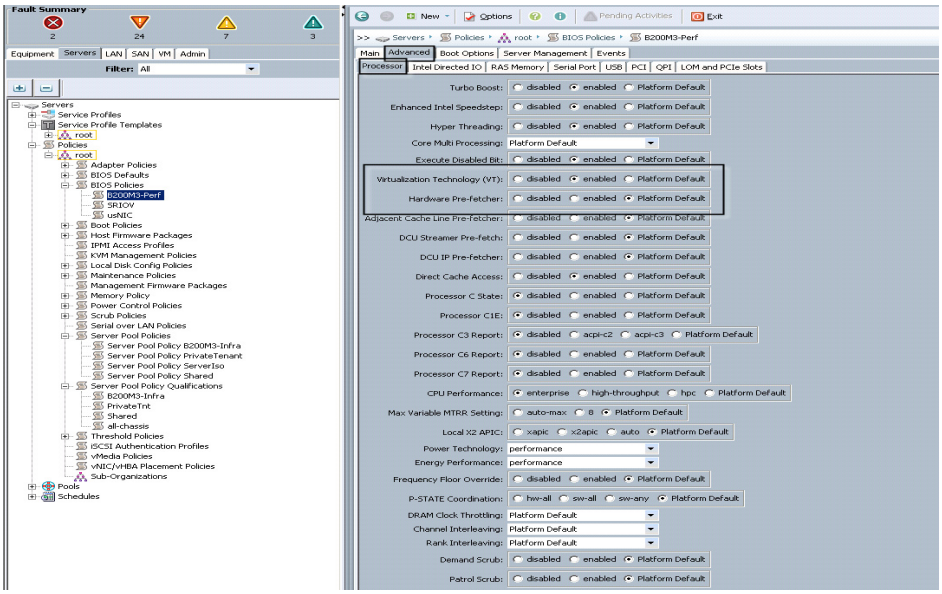

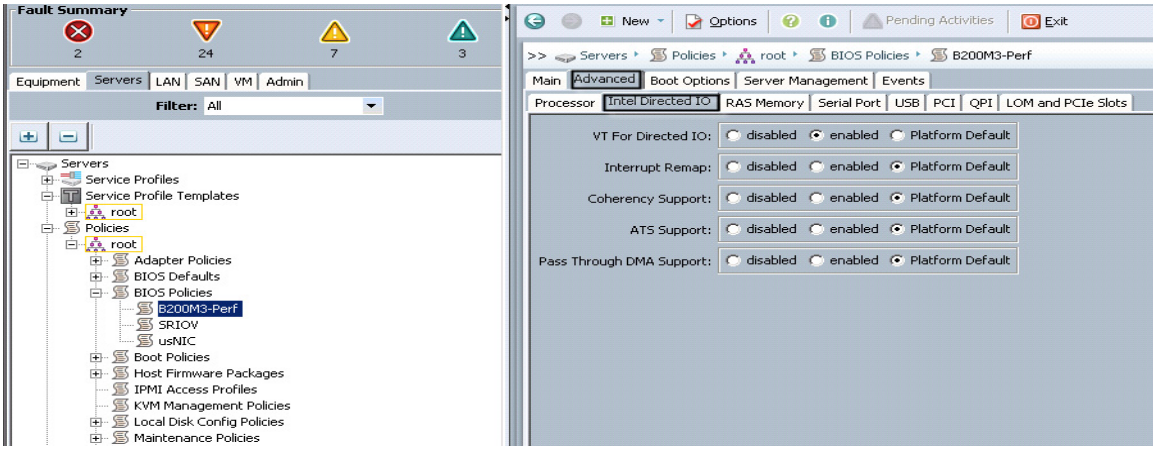

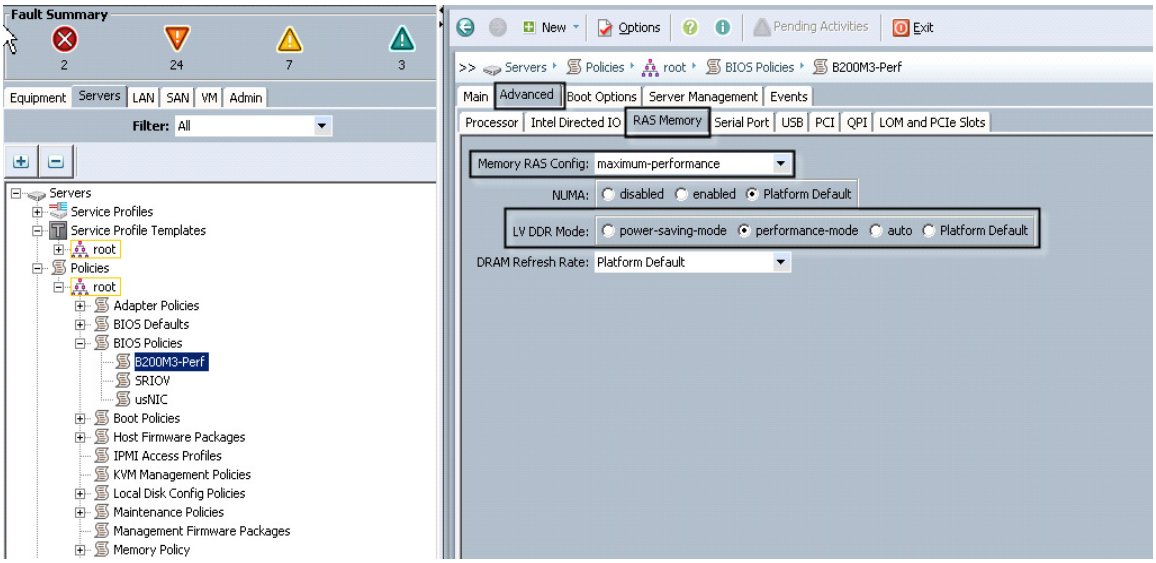

The remaining Advanced tab settings are at platform default or not configured. Similarly, the Boot Options and Server Management tabs'settings are at defaults. Many of the settings in this policy are the UCS B200 M3 BIOS default settings. We created this policy to illustrate the combined effect of the Platform Default and specific settings for this use case.

Note Be sure to Save Changes at the bottom of the page to preserve this setting. Be sure to add this policy to your blade service profile template.

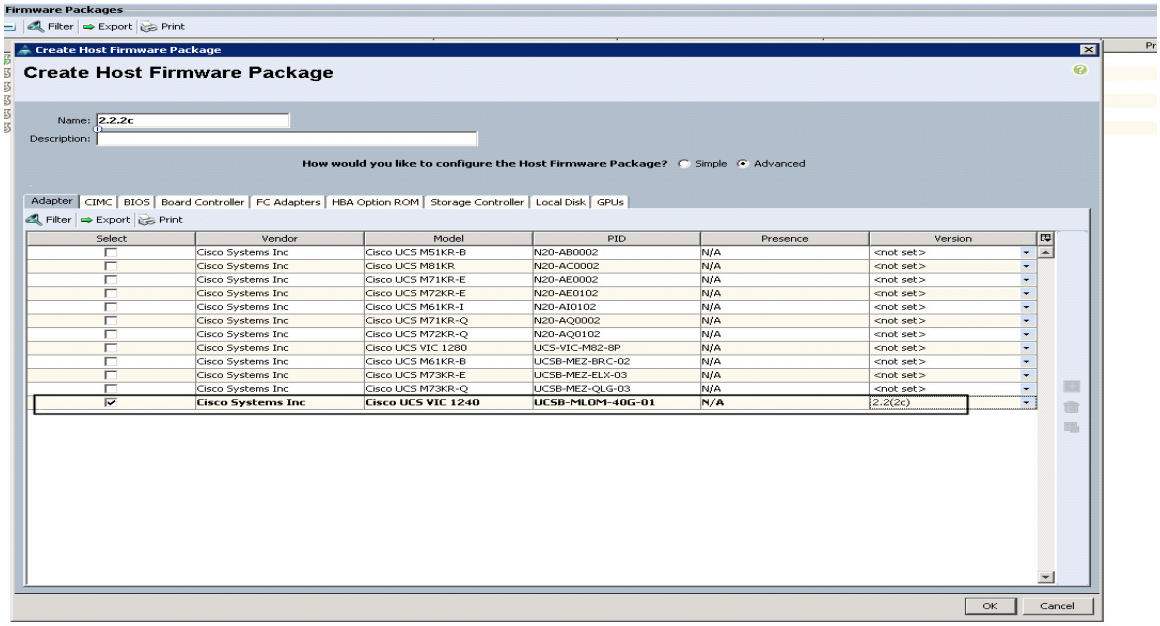

24. Continue through the CIMC, BIOS, Board Controller and Storage Controller tabs as follows:

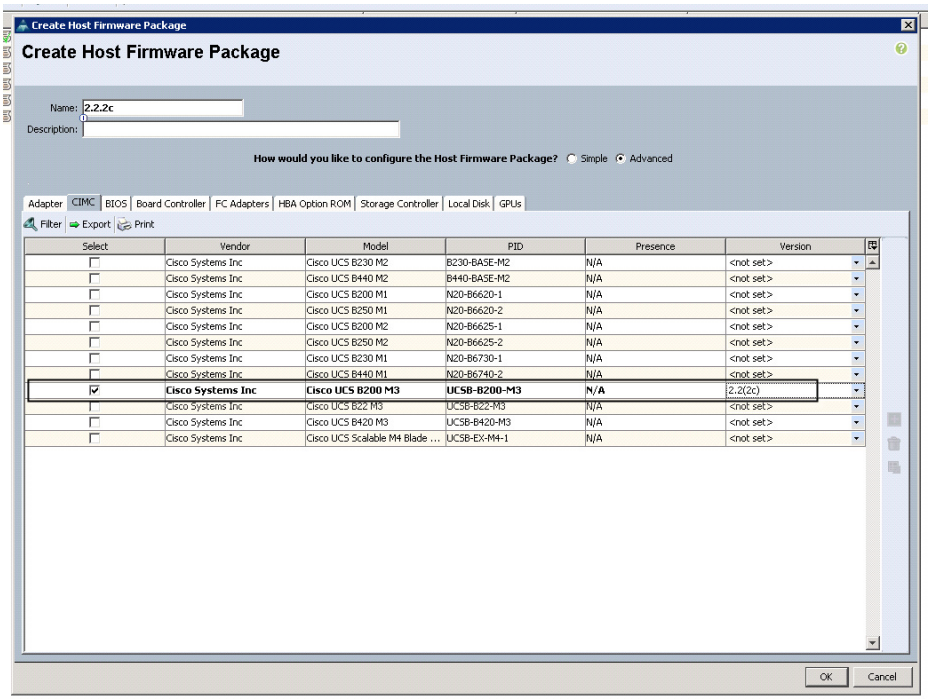

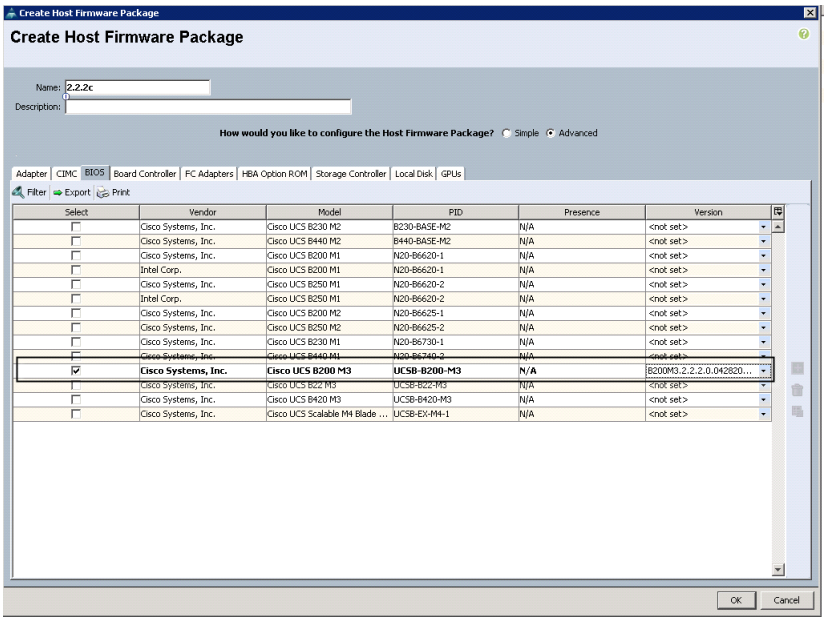

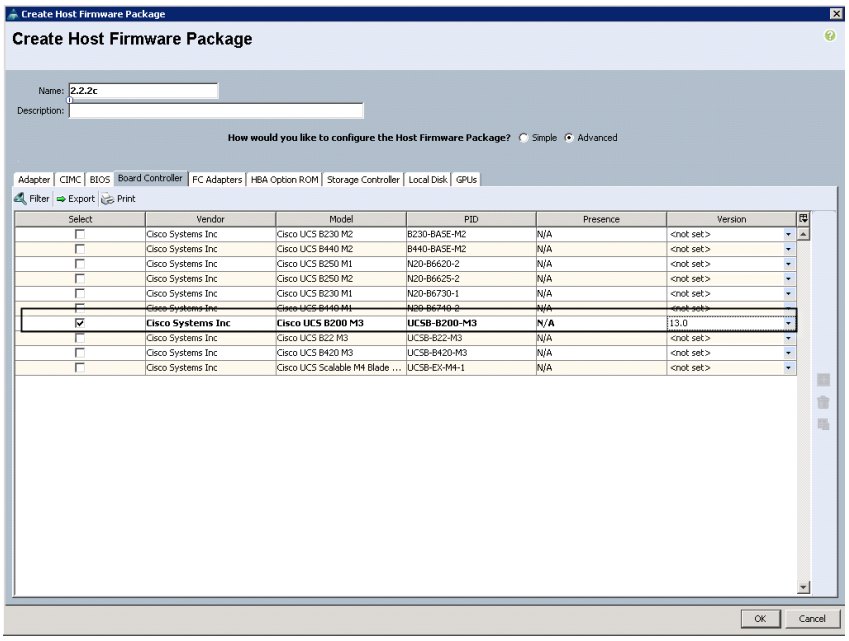

Note We did not use legacy nor third party FC Adapters nor HBA's so there was no configuration required on those tabs.

The result is a customized Host Firmware Package for the Cisco UCS B200 M3 blade servers.

Note For this project we utilized Fiber channel to boot from SAN.

25. In the Servers tab, expand Policies > root nodes. Select the Boot Policies node. Right-click and choose Create Boot Policy from the context menu.
26. In the Create Boot Policy dialog complete the following:
    a. Expand Local Devices
    b. Select Add CD-ROM
    c. Expand vHBAs
    d. Select Add SAN Boot (vHBA0) as Primary
    e. Select Add SAN Boot (vHBA1) as Secondary
    f. Adjust boot order so it is CD-ROM, vHBA0, vHBA1.
27. Click Save Changes.

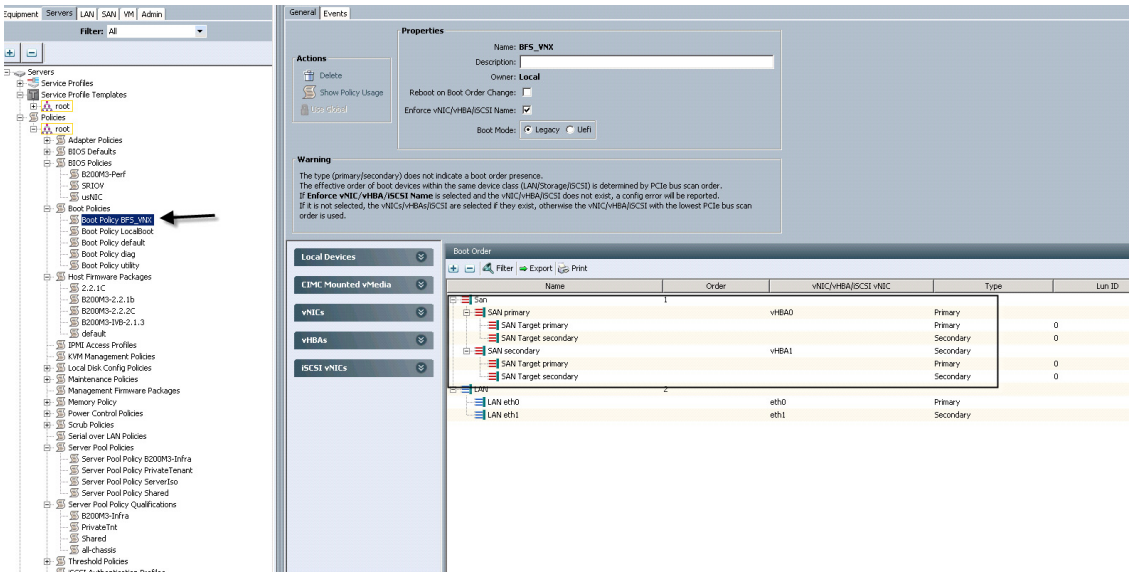

28. Create a service profile template using the pools, templates, and policies configured above. We created a total of five Service Profile Templates, one for each tenant model, Shared, Server Isolation, Private Tenant 1, Private Tenant 2 and Infrastructure Hosts (Infra)as follows:

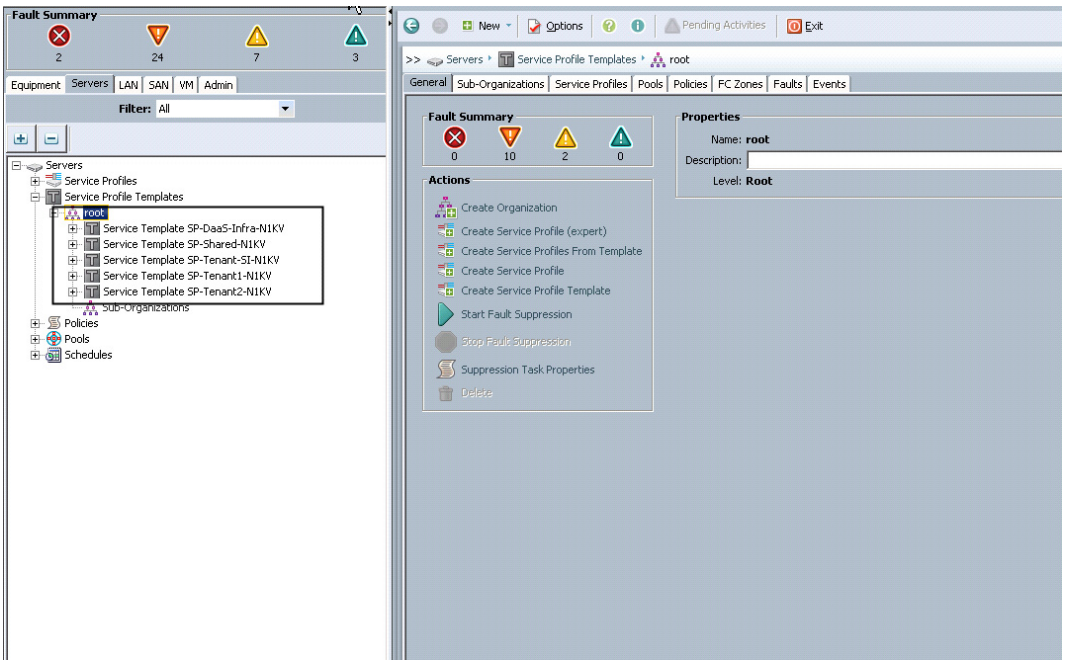

29. To create a Service Profile Template, right-click the Service Profile Templates node on the Servers tab and click Create Service Profile Template. The Create Service Profile template wizard will open.
30. Follow through each section, utilizing the policies and objects you created earlier, then click Finish.

Note On the Operational Policies screen, select the appropriate performance BIOS policy you created earlier to insure maximum LV DIMM performance.

Note For automatic deployment of service profiles from your template(s), you must associate a server pool that contains blades with the template.

31. On the Create Service Profile Template wizard, we entered a unique name, selected the type as updating, and selected the SP-UUIDs Suffix Pool created earlier, then clicked Next.

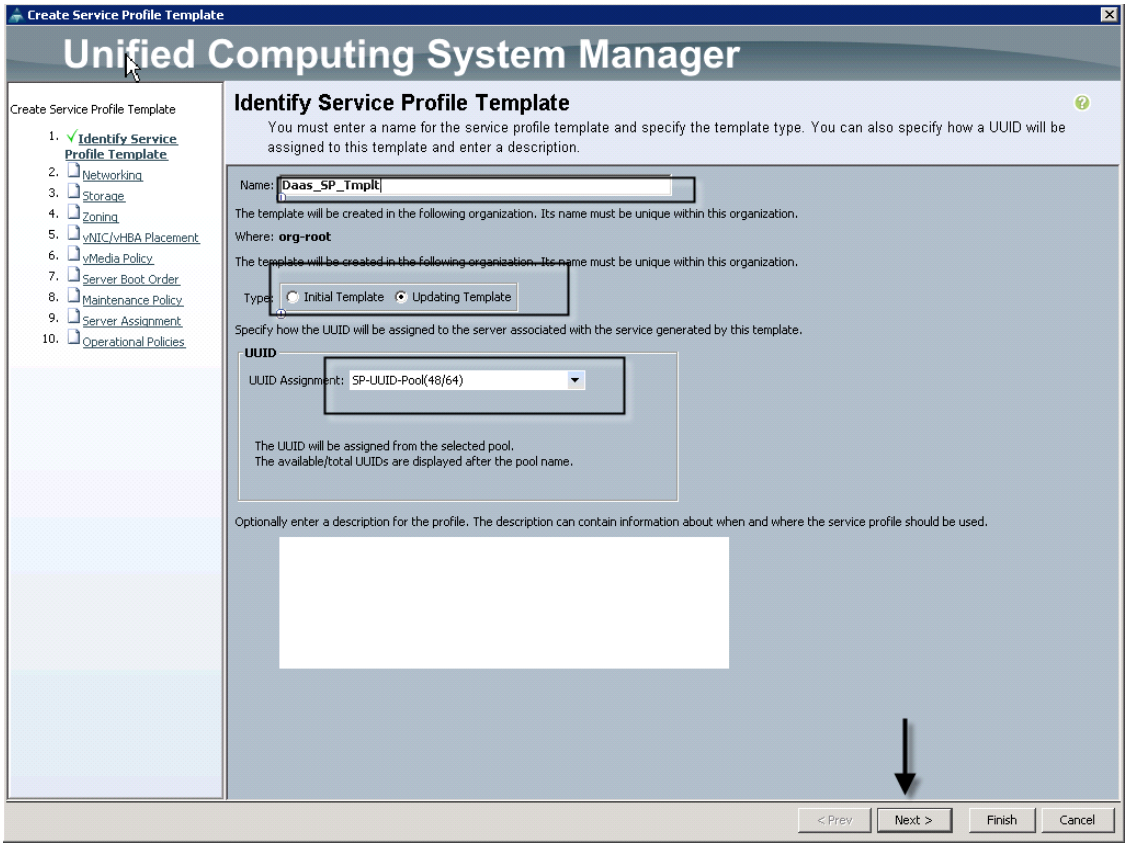

32. We selected the Expert configuration option on the Networking page and clicked Add in the adapters window:

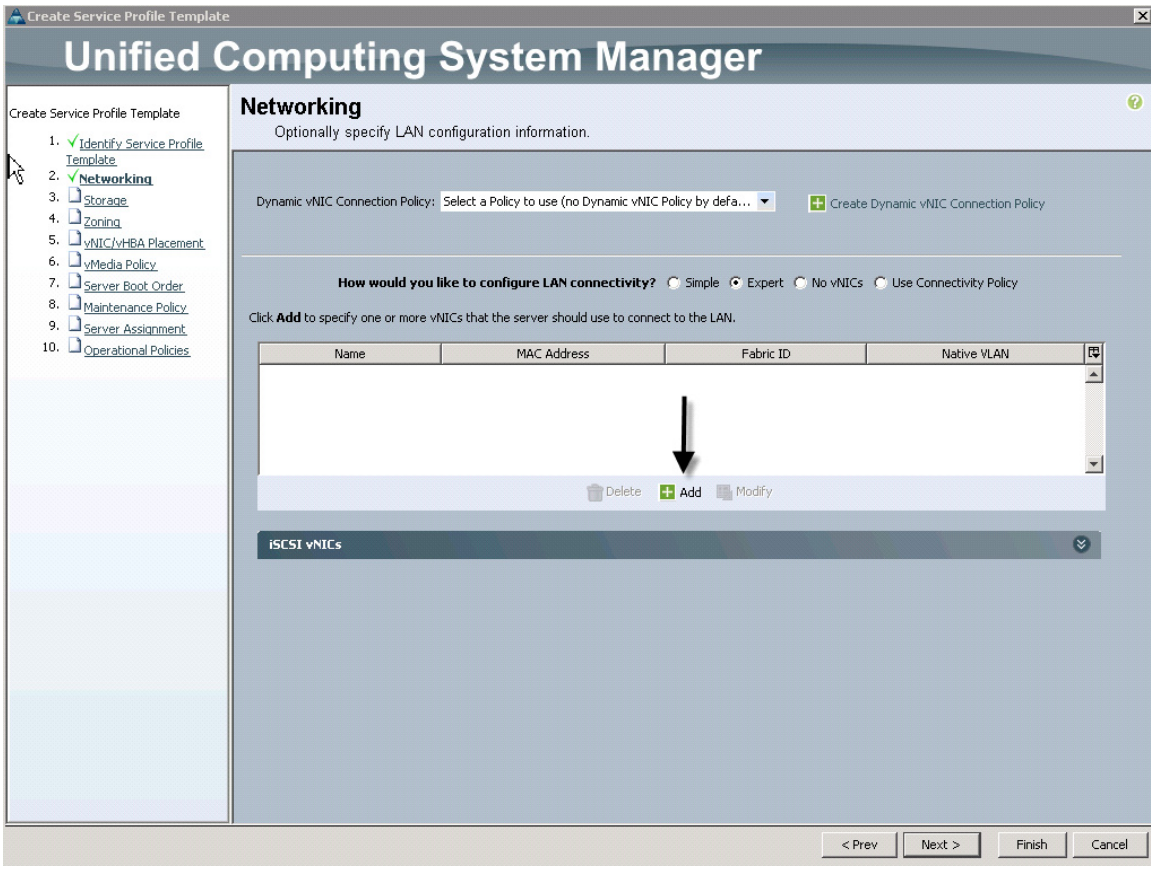

33. In the Create vNIC window, we entered a unique Name, checked the Use LAN Connectivity Template checkbox, selected the vNIC Template from the drop down, and the Adapter Policy the same way.

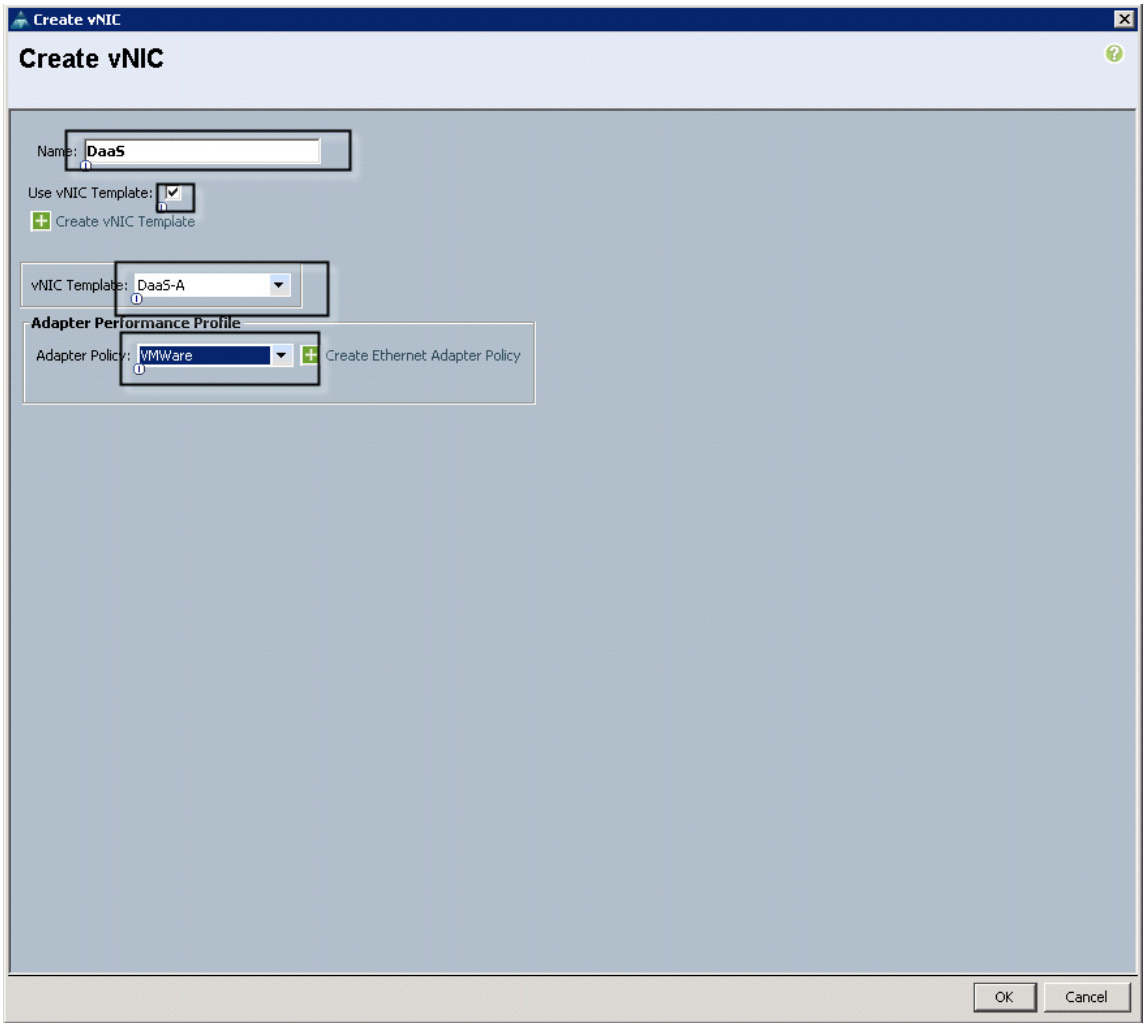

34. We repeated the process for the remaining vNIC , resulting in the following:

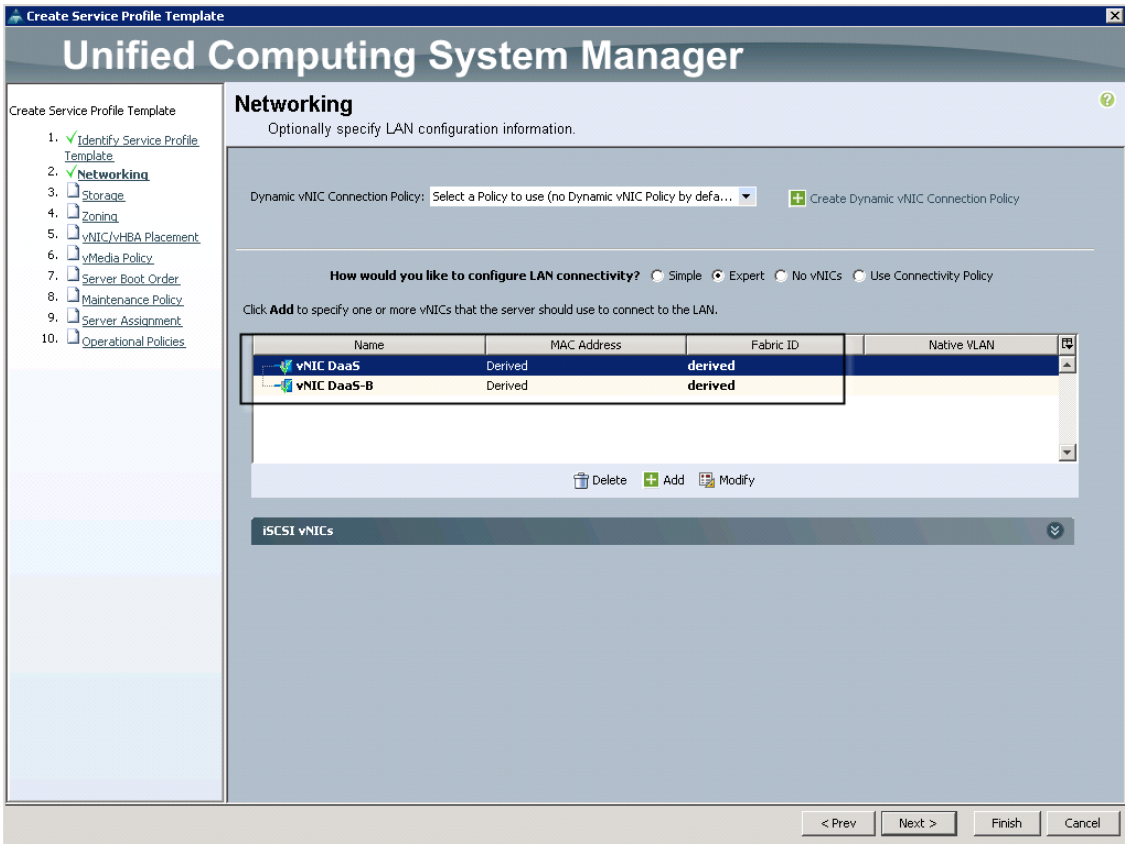

35. Click Next.
36. Click Next.

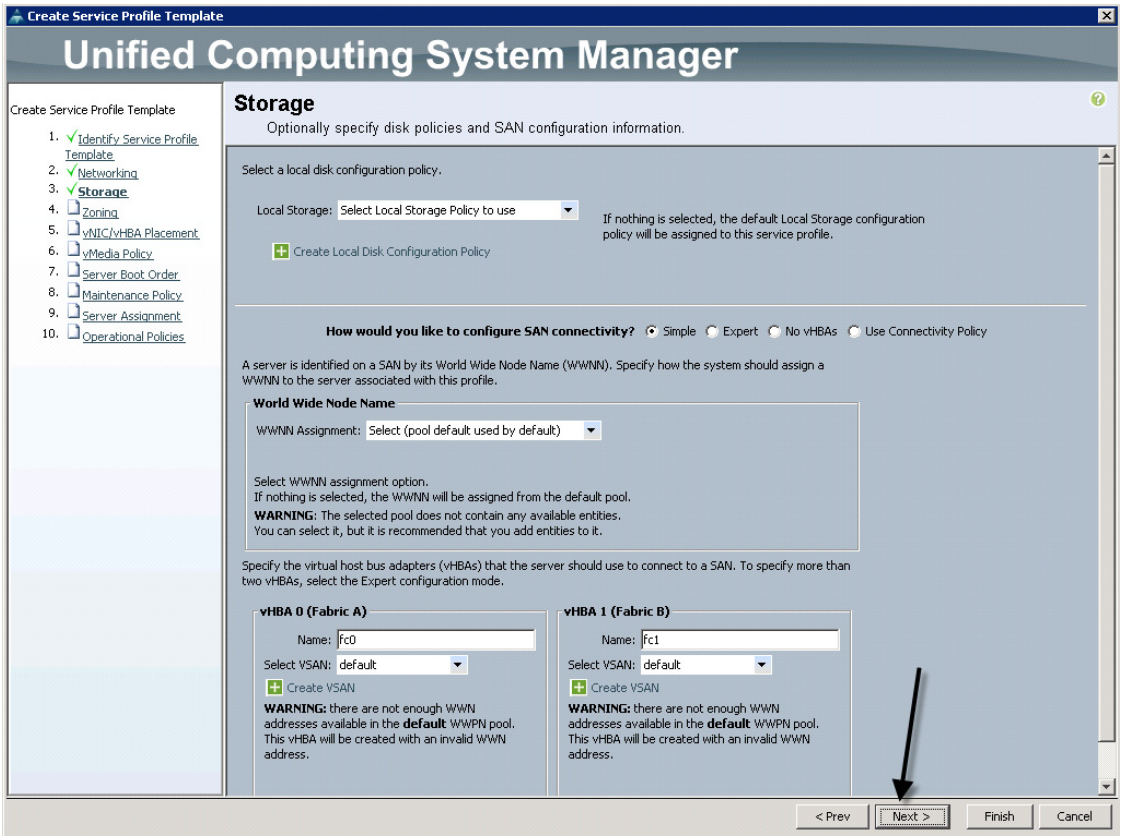

37. On the Zoning page, click Next.

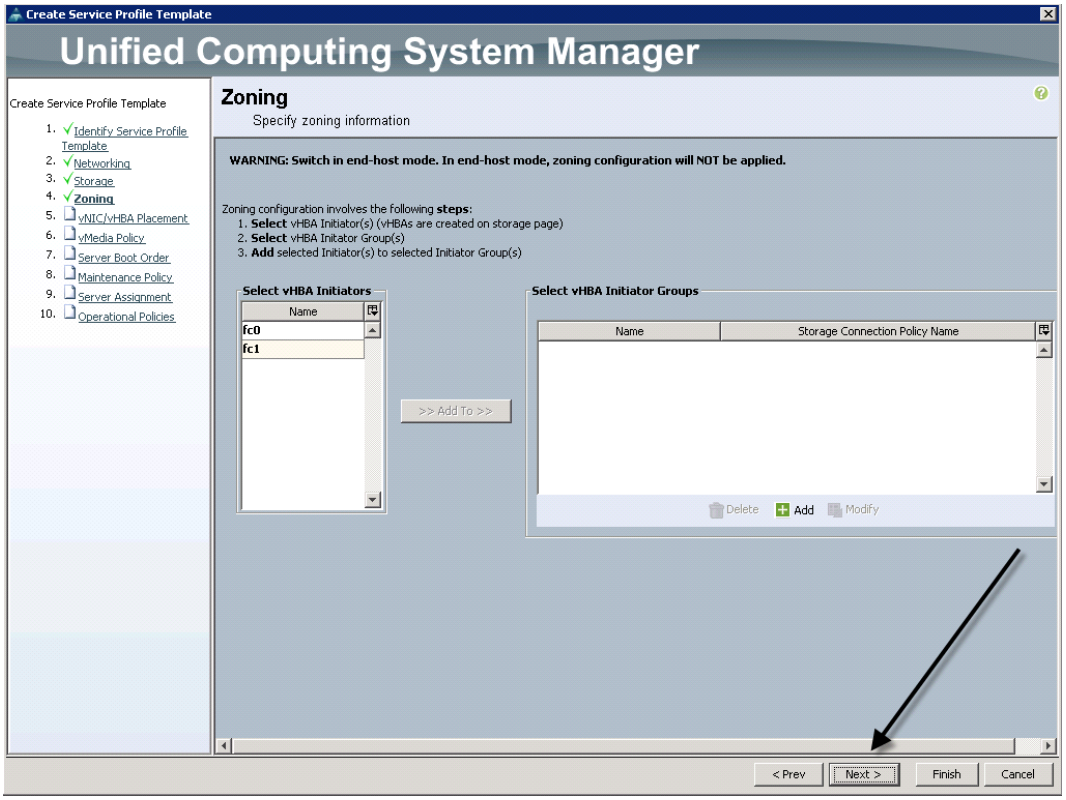

38. On the vNIC/vHBA Placement page, click Next.

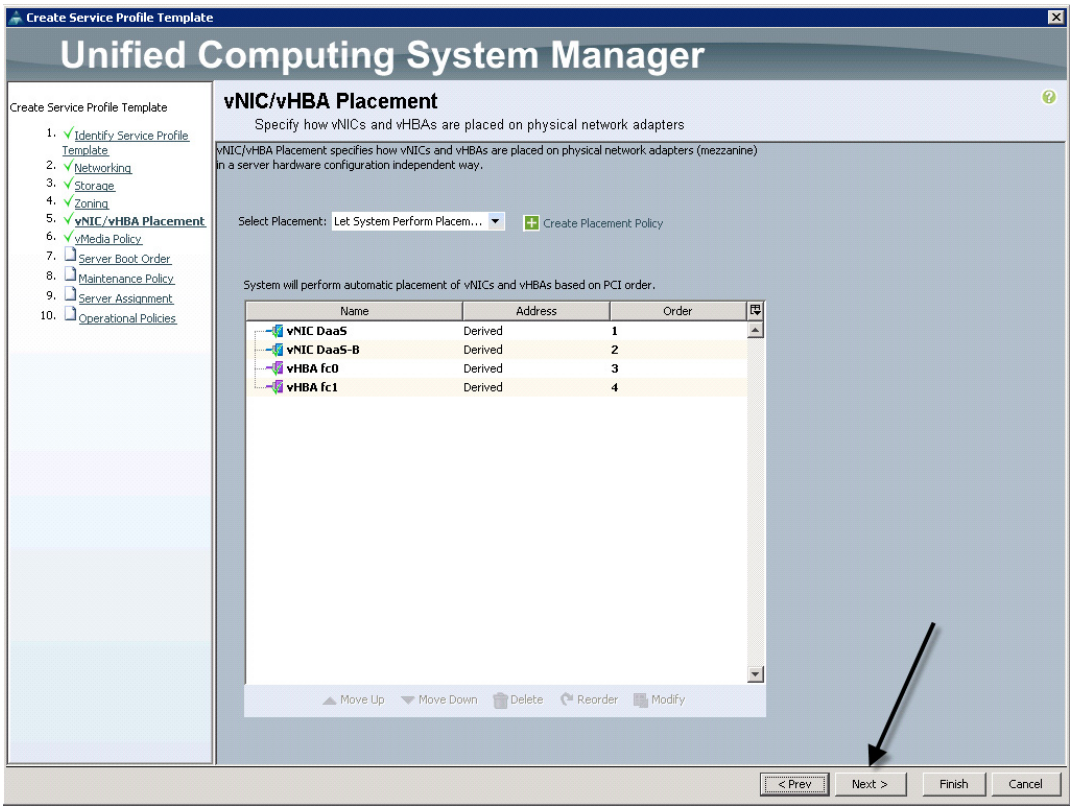

39. vMedia Policy click Next

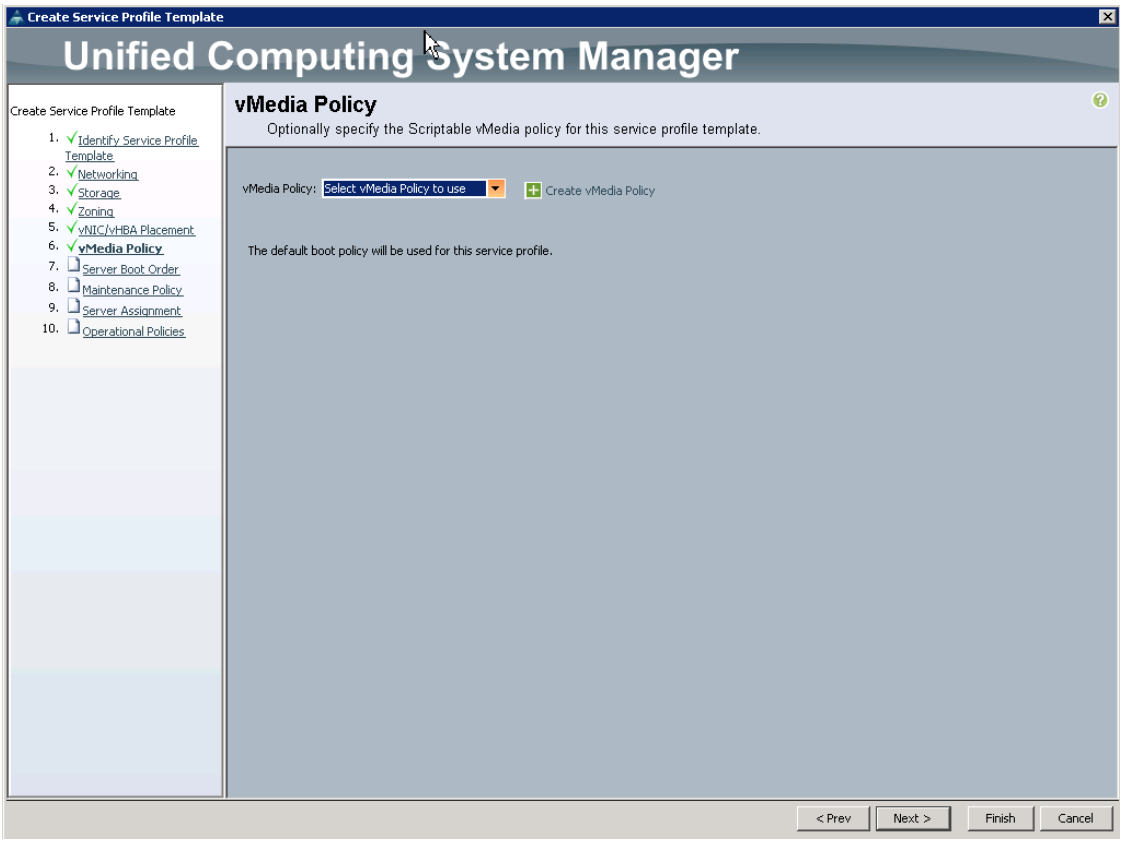

40. On the Server Boot Order page, select the Boot From SAN Boot policy BFS_VNX, created earlier from the drop-down, then click Next to proceed.

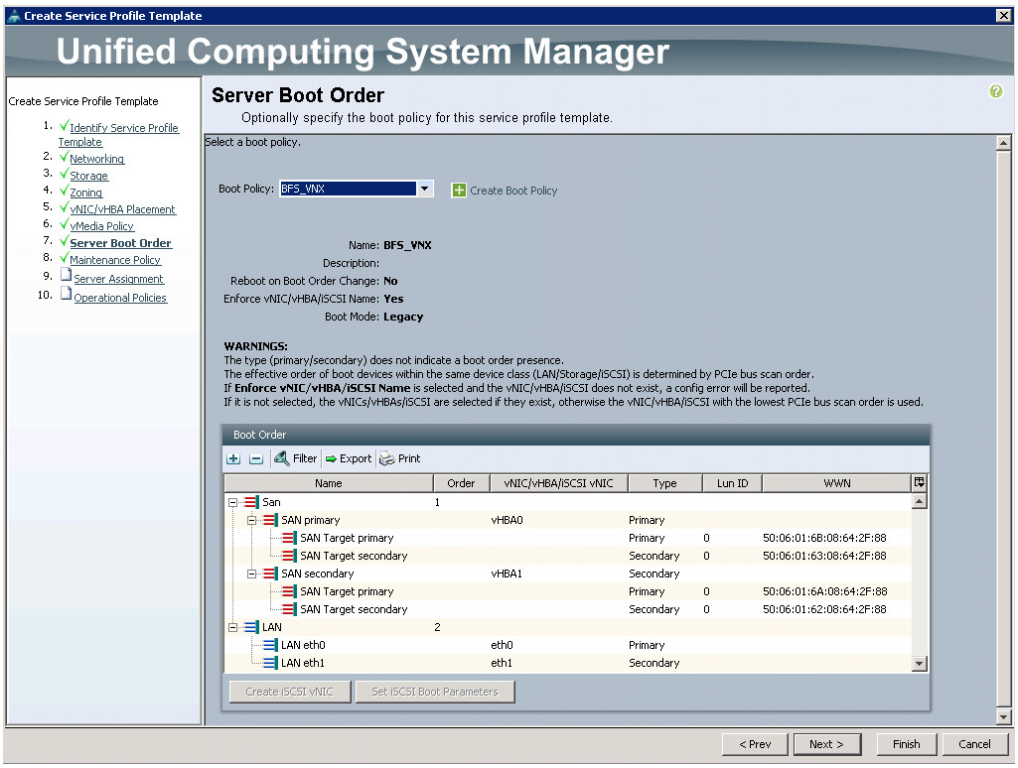

41. Do not create a Maintenance Policy for the project, click Next to continue.

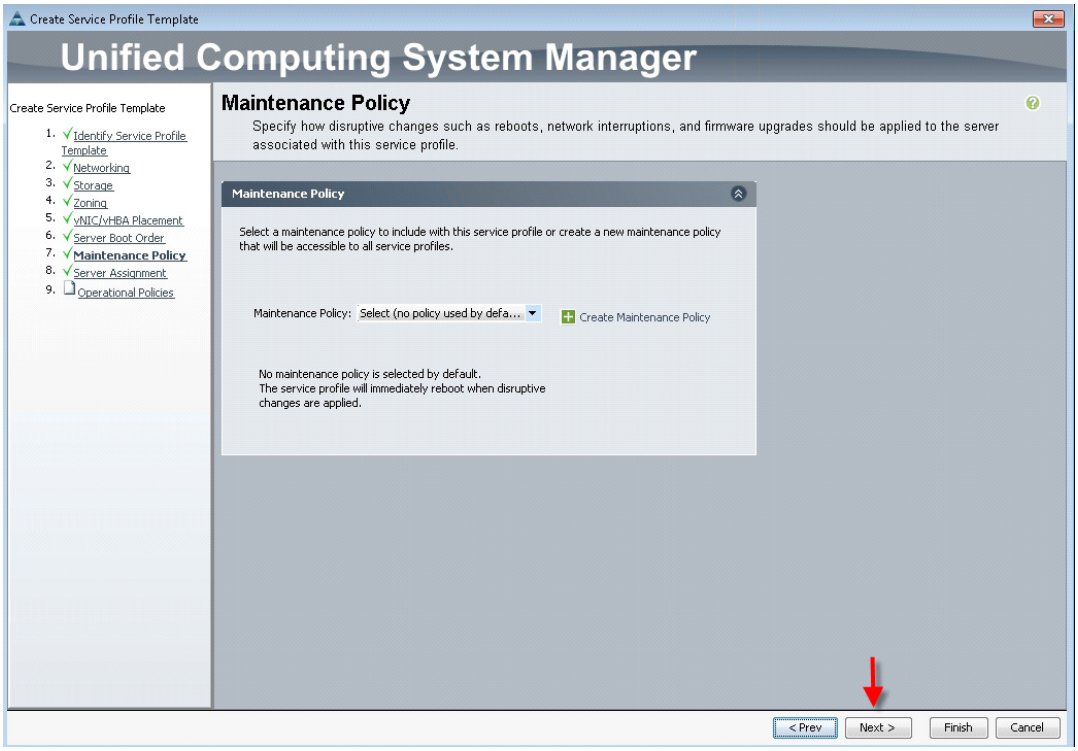

42. On the Server Assignment page, make the following selections from the drop-downs and click the expand arrow on the Firmware Management box as shown:

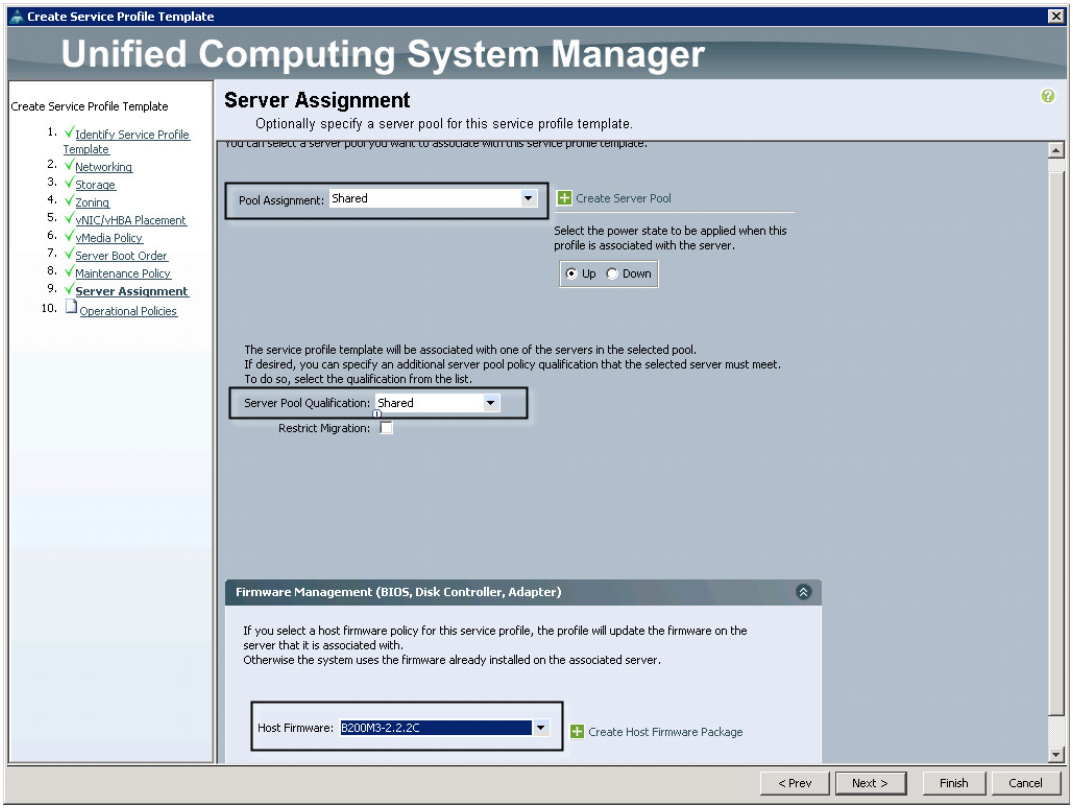

Note For the other four Service Profile Templates that were created for the project, we choose ServerIso, PrivateTenant or Infrastructure for Pool Assignments and PrivTnt or Infra for the Server Pool Qualification.

Note In all three cases, we utilized the Default Host Firmware Management policy for the Cisco UCS B200 M3 blades.

43. On the Operational Policies page, expand the BIOS Configuration drop-down and select the B200M3-Perf for User Workload Hosts, then click Finish to complete the Service Profile Template:

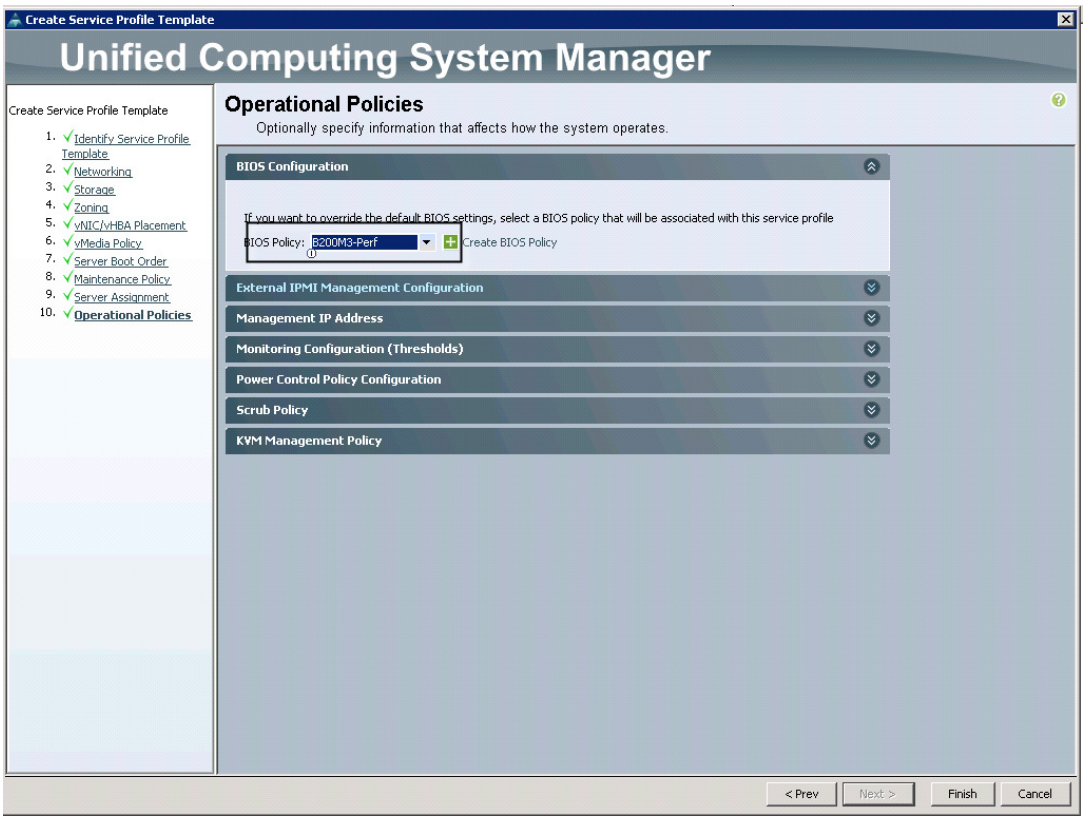

44. Repeat the Create Service Profile Template for the four remaining templates.

The result is a Service Profile Templates for each use case in the study and an Infrastructure template as shown below:

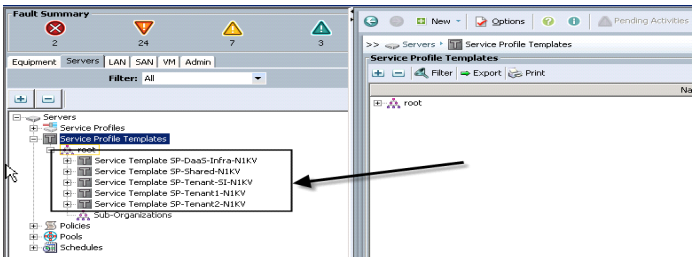

Now that the Service Profile Templates for each UCS Blade Server model used in the project have been created, utilize a UCS Director workflow to provision Service Profile Templates to blades during the Hypervisor installation.

QoS and CoS in Cisco Unified Computing System

Cisco Unified Computing System provides dfferent 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.

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 9217.
- Through Cisco Unified Computing System you 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).

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

Cisco UCS System Class Configuration

Cisco Unified Computing System defines user class names as follows.

- Platinum
- Gold
- Silver
- Bronze

Table 5 Name Table Map Between Cisco Unified Computing System and the NXOS

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

Table 6 Class to CoS Map by default in Cisco Unified Computing System

Cisco UCS Class Names	Cisco UCS Default Class Value
Best effort	Match any
Fc	3
Platinum	5
Gold	4
Silver	2
Bronze	1

Table 7 Default Weight in Cisco Unified Computing System

Cisco UCS Class Names	Weight
Best effort	5
Fc	5

Enable QOS on the Cisco Unified Computing System

For this study, we utilized four 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	VLAN Supported
Platinum	eth0, eth1	72 (Storage)
Gold	eth0, eth1	71,72,74,75,76,77 (VM Traffic)
Silver	eth0, eth1	70 (Management)
Bronze	eth0, eth1	73 (vMotion)

Note In this study, all VLANs were trunked to eth0 and eth1 and both use Best Effort QoS. Detailed QoS was handled by the Cisco Nexus 1000V and Nexus 5548 switches, but it is important that the UCS QoS System Classes match what the switches are using.

Configure Platinum, Gold, Silver and Bronze policies by checking the enabled box. For the Platinum Policy, used for NFS and CIFS storage, Bronze for vMotion and Best Effort were configured for Jumbo Frames in the MTU column. Notice the option to set no packet drop policy during this configuration. Click Save Changes at the bottom right corner prior to leaving this node.

Figure 34 Cisco UCS QoS System Class Configuration

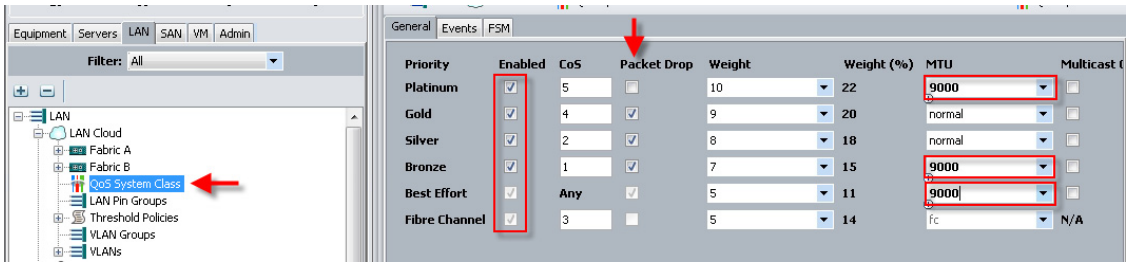

This is a unique value proposition for Cisco UCS with respect to end-to-end QOS. For example, we have a VLAN for the EMC VNX storage, configured Platinum policy with Jumbo frames and get an end-to-end QOS and performance guarantees from the Blade Servers running the Nexus 1000V virtual distributed switches through the Nexus 5548UP access layer switches.

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 FC switches for our DaaS deployment.

Cisco UCS and EMC VNX Ethernet Connectivity

Two 10 Gigabit Ethernet uplink ports and 2, 8 Gigabit FC ports are configured on each of the Cisco UCS 6248 fabric interconnects, and they are connected to the Cisco Nexus 5548 pair in a bow tie manner as shown below in a port channel.

The 6248 Fabric Interconnect is in End host mode, and switch mode for fiber channel as we are doing both fiber channel as well as Ethernet (NAS) data access as per the recommended best practice of the Cisco Unified Computing System. We built this out for scale and have provisioned 20 GB per Fabric Interconnect for ethernet (Figure 32) and 16GB per Fabric Interconnect for fiber.

The VNX5600s are also equipped with two dual-port 10GB adapters which are connected to the pair of N5Ks downstream. Both paths are active providing failover capability. This allows end-to-end 10G access for file-based storage traffic. We have implemented jumbo frames on the ports and have priority flow control on, with Platinum CoS and QoS assigned to the vNICs carrying the storage data access on the Fabric Interconnects.

Note The upstream configuration is beyond the scope of this document; there are some good reference document [4] that talks about best practices of using the Cisco Nexus 5000 and 7000 Series Switches. New with the Nexus 5500 series is an available Layer 3 module that was not used in these tests and that will not be covered in this document.

Figure 35 Ethernet Network Configuration with Upstream Cisco Nexus 5500 Series from the Cisco Unified Computing System 6200 Series Fabric Interconnects and EMC VNX5600

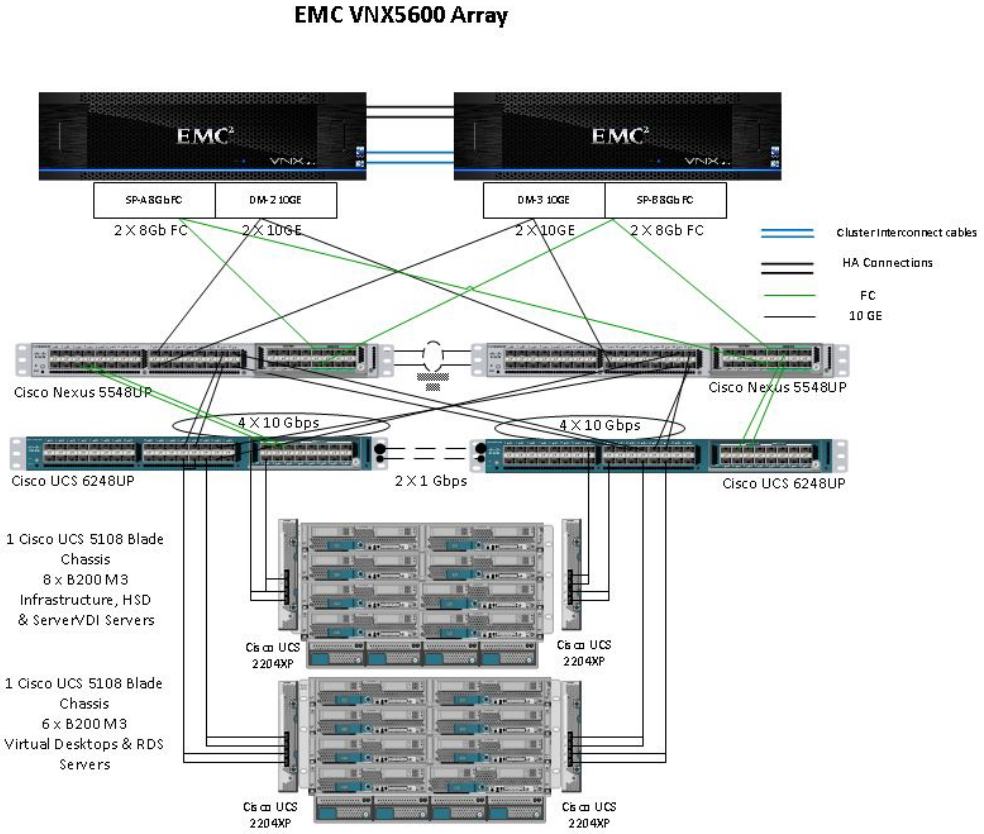

Cisco Nexus 1000V Configuration in L3 Mode

1. To download the Nexus1000 V 4.2(1) SV2 (2.2), click the link below.
   https://software.cisco.com/download/release.html?mdfid=282646785&softwareid=282088129&release=5.2(1)SV3(1.2)&flowid=42790
2. Extract the downloaded N1000V .zip file on the Windows host.
3. To start the N1000V installation, run the command below from the command prompt. (Make sure the Windows host has the latest Java version installed)

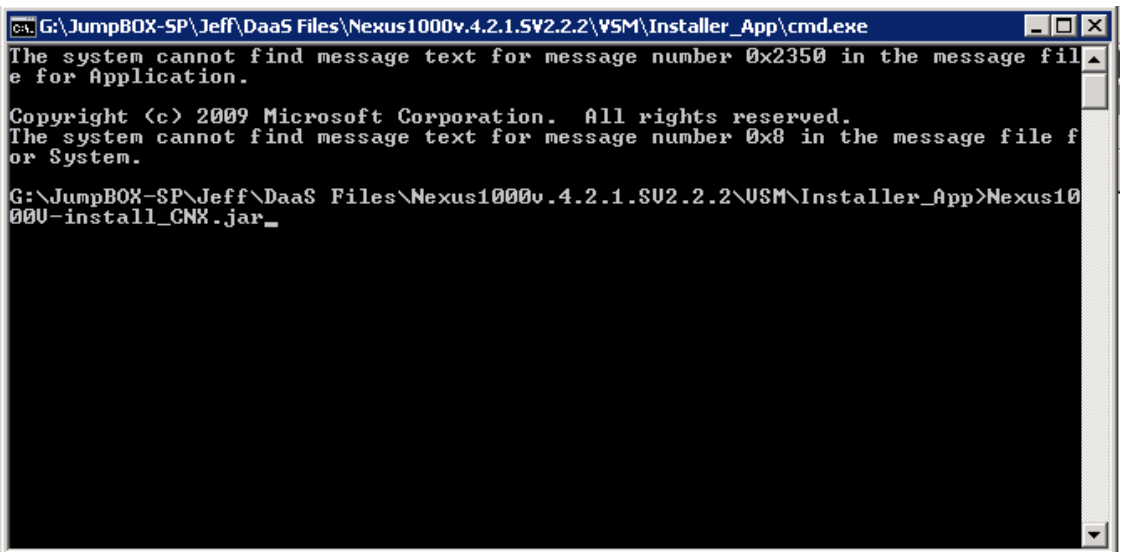

4. After running the installation command, you will see the “Nexus 1000V Installation Management Center”

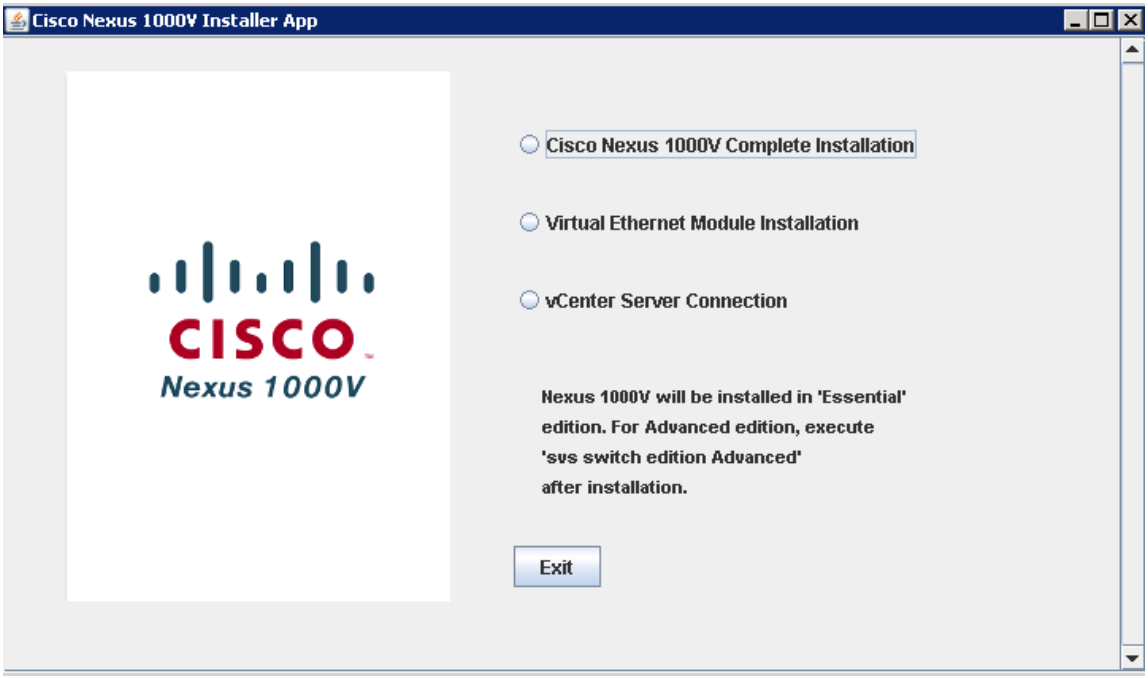

5. Type the vCenter IP and the logon credentials.

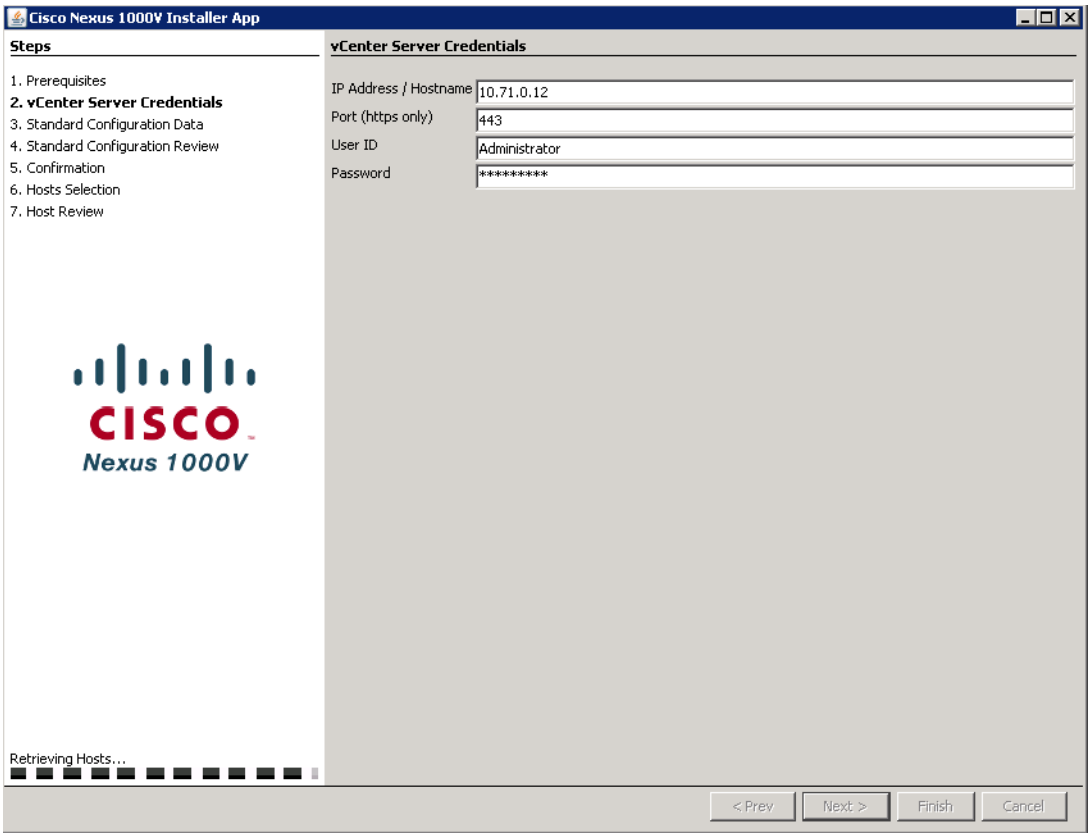

6. Select the ESX host on which to install N1KV Virtual Switch Manager.

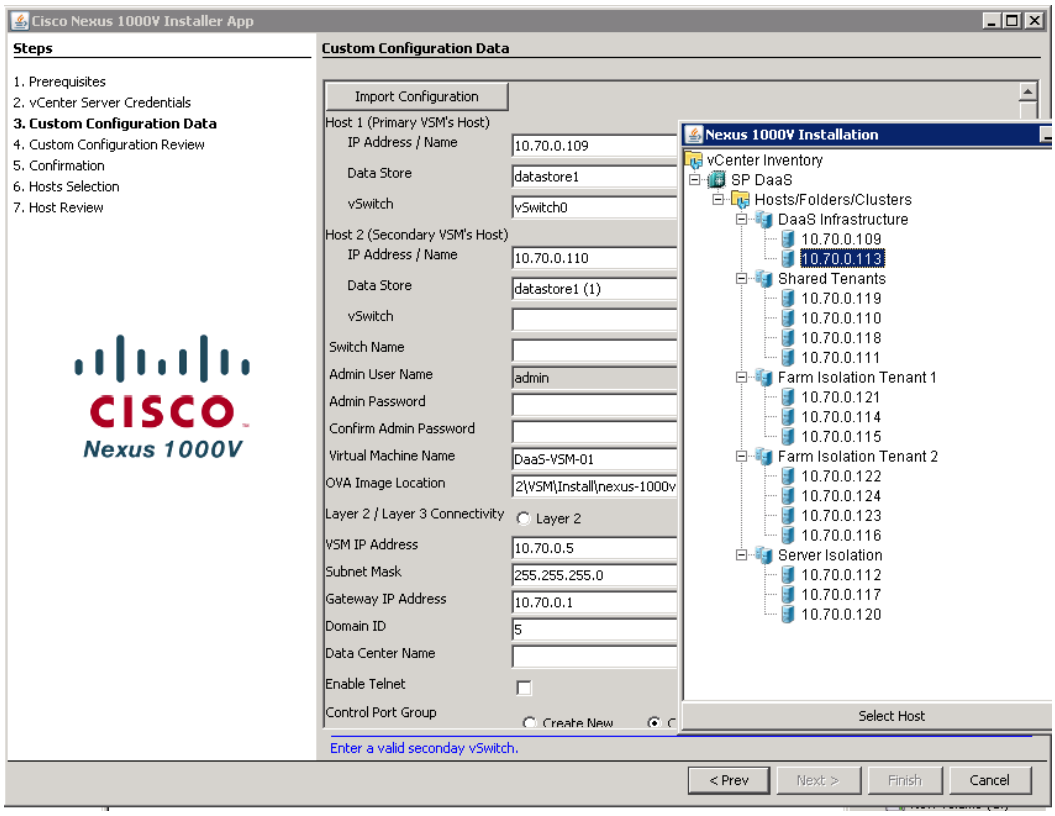

7. Configure all fields for the Custom Install.

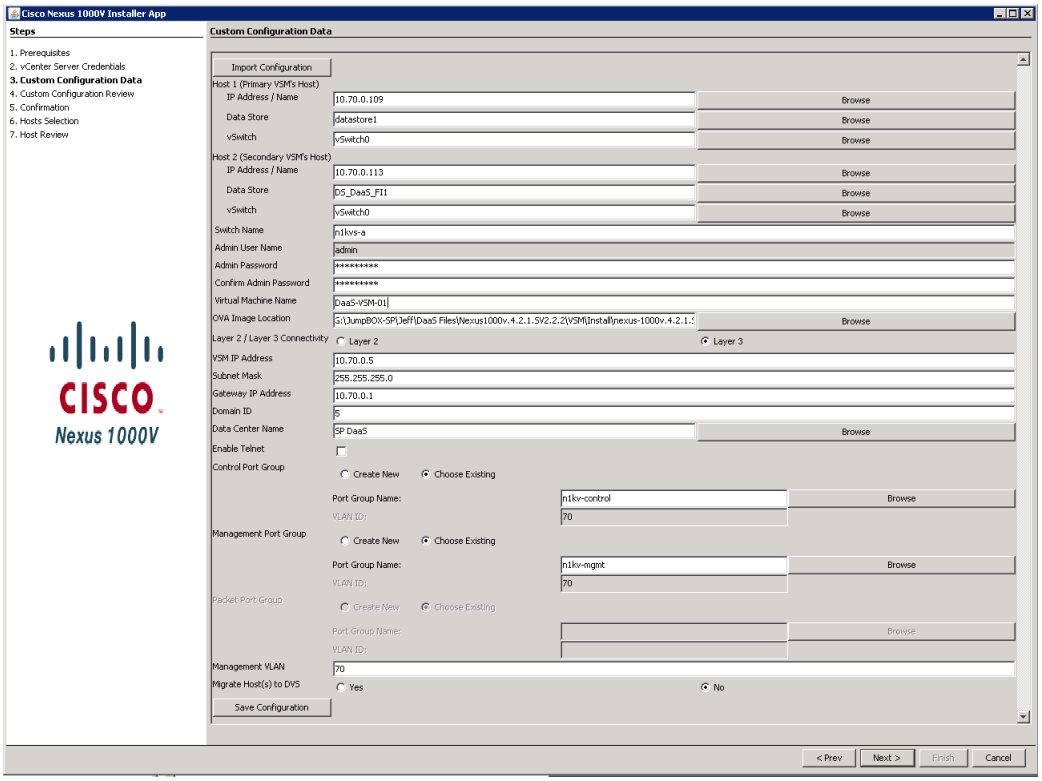

8. Confirm all settings and click next for the VSM Deployment.

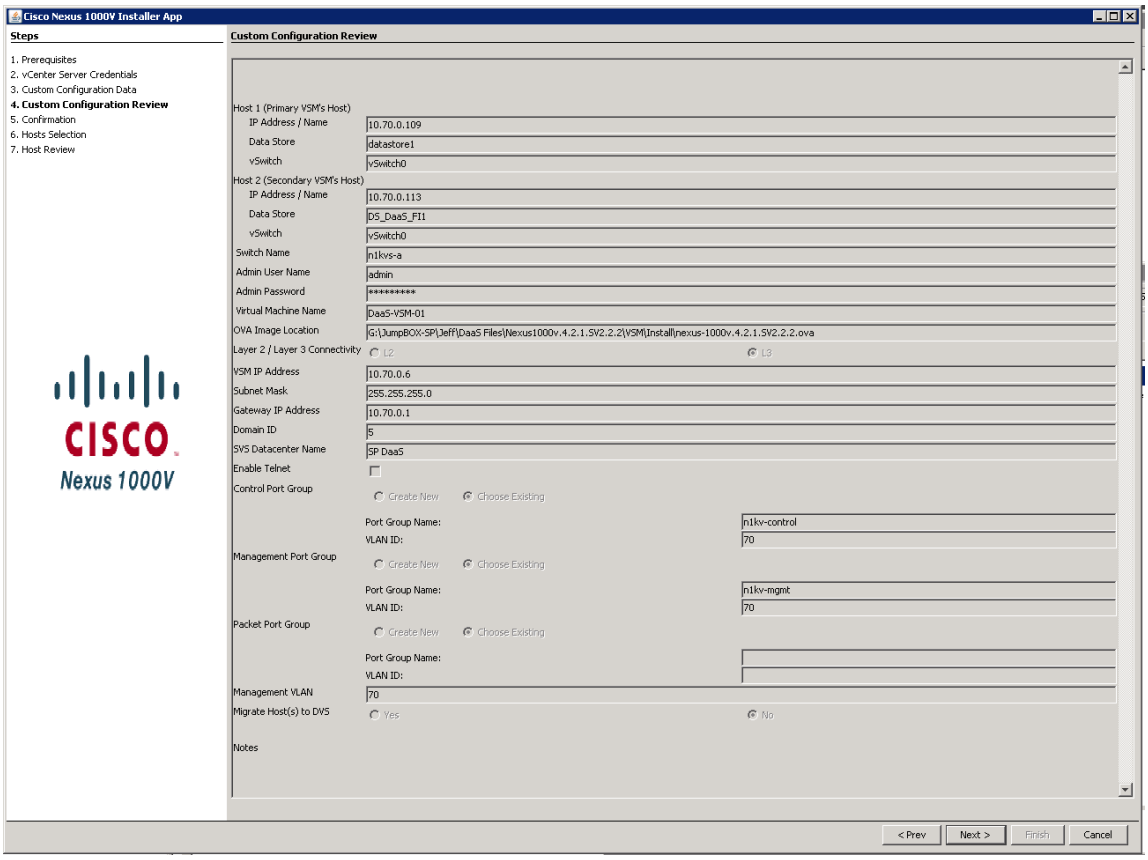

9. Install VEM to the other hosts

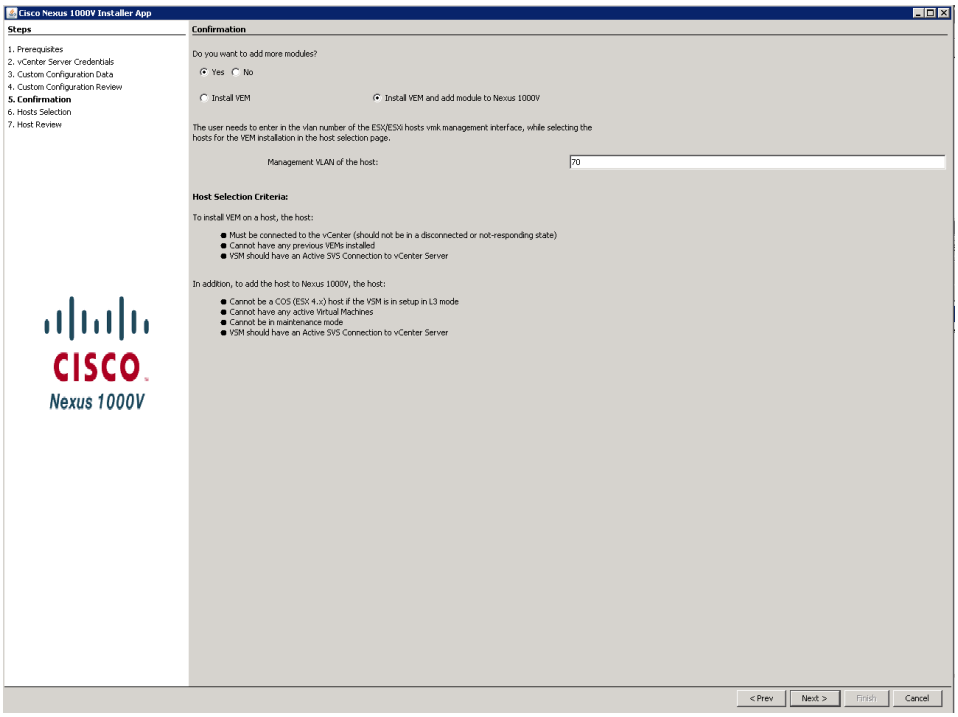

10. Select the Hosts that you will deploy the VEM.

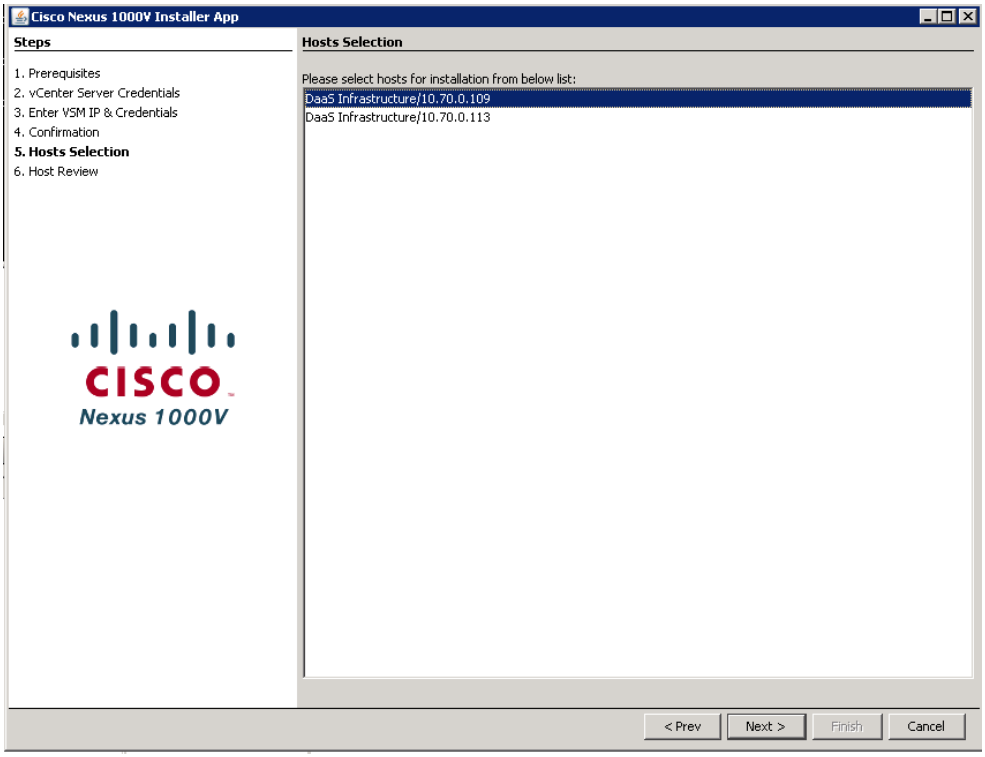

11. Host Review Page, click Finish

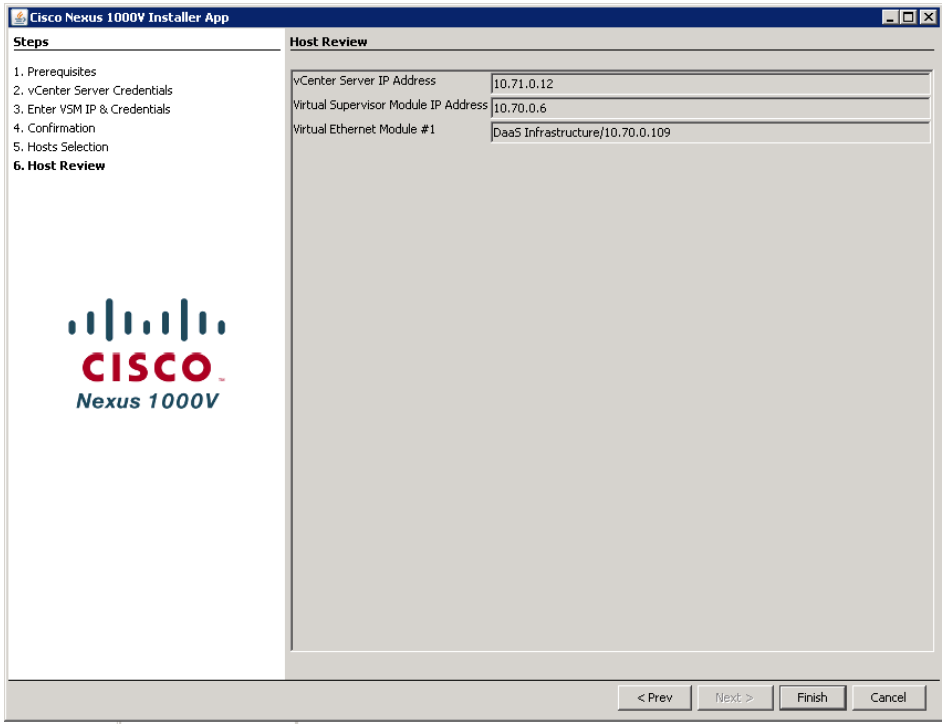

12. Log in (ssh or telnet) to the N1KV VSM with the IP address and configure VLAN for ESX Mgmt, Control, N1K Mgmt and also for Storage and vMotion purposes as mentioned below (VLAN ID differs based on your Network). First, create ip access lists for each QoS policy:

```
N1kvs# conf t
Enter the following configuration commands, one per line.  End with CNTL/Z.
ip access-list mark_Bronze
  10 permit ip any 10.73.0.0/24
  20 permit ip 10.73.0.0/24 any
ip access-list mark_Gold
10 permit ip any 10.74.0.0/24
                   11 permit ip any 10.76.0.0/24
                   12 permit ip any 10.77.0.0/24
                   20 permit ip 10.74.0.0/24 any
                   21 permit ip 10.76.0.0/24 any
                   22 permit ip 10.77.0.0/24 any
 ip access-list mark_Platinum
  10 permit ip any 10.72.0.0/24
  20 permit ip 10.72.0.0/24 any
ip access-list mark_Silver
     10 permit ip any 10.71.0.0/24
  20 permit ip 10.71.0.0/24 any
```

13. Create class maps for QoS policy

```
class-map type qos match-all Gold_Traffic
     match access-group name mark_Gold
class-map type qos match-all Bronze_Traffic
     match access-group name mark_Bronze
class-map type qos match-all Silver_Traffic
     match access-group name mark_Silver
class-map type qos match-all Platinum_Traffic
```

```
    match access-group name mark_Platinum
```

14. Create policy maps for QoS and set class of service

```
policy-map type qos DaaS
  class Platinum_Traffic
    set cos 5
  class Gold_Traffic
    set cos 4
  class Silver_Traffic
    set cos 2
  class Bronze_Traffic
    set cos 1
```

15. Set vlans for QoS

```
vlan 1,6,70-77,79
vlan 6
  name Native-VLAN
vlan 71
  name
Infrastructure
vlan
72
  name
NFS
vlan
73
  name
VMotion
vlan
74
  name
Shared
vlan 75
  name Server-Isolation
vlan 76
  name FarmIso1
vlan 77
  name FarmIso2
vlan 79
  name PXE
```

16. Create port profile for system uplinks and vethernet port groups. **Note:** There are existing port profiles created during the install. Do not modify or delete these port profiles.

```
port-profile type ethernet system-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 70-79,701
  system mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 70-74,76-77
  state enabled
port-profile type vethernet Management
  vmware port-group
  switchport mode access
  switchport access vlan 70
  service-policy type qos input DaaS
  no shutdown
```

```
  system vlan 70
  max-ports 254
  state enabled
port-profile type vethernet Storage
  vmware port-group
  switchport mode access
  switchport access vlan 72
  service-policy type qos input DaaS
  no shutdown
  system vlan 72
  state enabled
port-profile type vethernet vMotion
  vmware port-group
  switchport mode access
  switchport access vlan 73
  service-policy type qos input DaaS
  no shutdown
  system vlan 73
  state enabled
port-profile type vethernet SharedIso
  vmware port-group
  port-binding static auto expand
  switchport mode access
  switchport access vlan 74
  service-policy type qos input DaaS
  no shutdown
  system vlan 74
  state enabled
port-profile type vethernet FarmIso1
  vmware port-group
  port-binding static auto expand
  switchport mode access
  switchport access vlan 76
  service-policy type qos input DaaS
  no shutdown
  system vlan 76
  state enabled
port-profile type vethernet FarmIso2
  vmware port-group
  port-binding static auto expand
  switchport mode access
  switchport access vlan 77
  service-policy type qos input DaaS
  no shutdown
  system vlan 77
  state enabled
port-profile type vethernet Infrastructure
  vmware port-group
  port-binding static auto expand
  switchport mode access
  switchport access vlan 71
  service-policy type qos input DaaS
  no shutdown
  system vlan 71
  state enabled
port-profile type vethernet n1k-L3
  capability l3control
  vmware port-group
```

```
switchport mode access
switchport access vlan 70
service-policy type qos input DaaS
no shutdown
system vlan 70
```

17. After creating port profiles, make sure vCenter shows all the port profiles and port groups under the respective N1KV VSM. Then, Add the ESXi host to the VSM.
18. Go to Inventory →networking →select DVS for N1KV →click on tab for hosts.

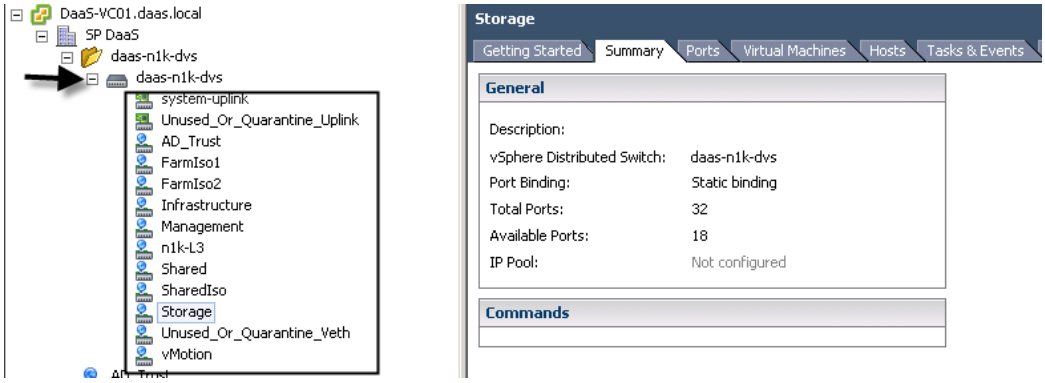

19. Right-click and select add host to vSphere Distributed Switch.

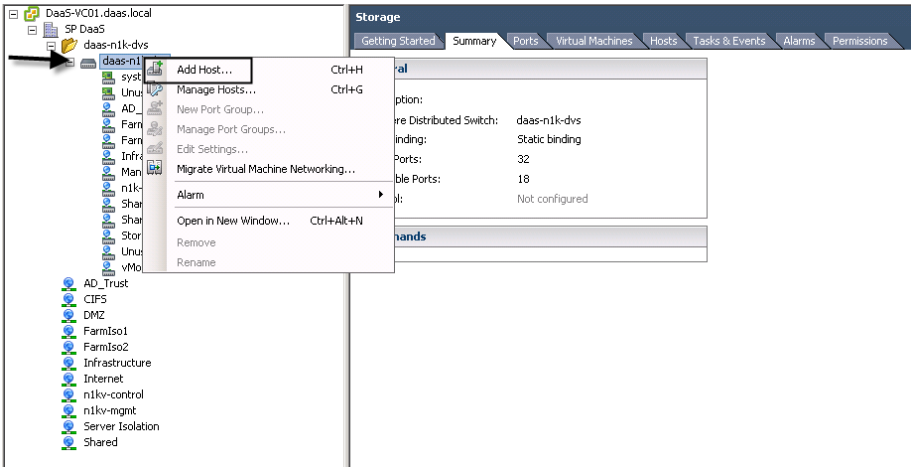

This will bring up ESXi hosts which are not part of existing configuration.

20. Select the ESX host to add, choose the vNICs to be assigned, click on select an uplink port-group drop down and select system-uplink for both vmnic0 and vmnic1. After selecting appropriate uplinks click Next.

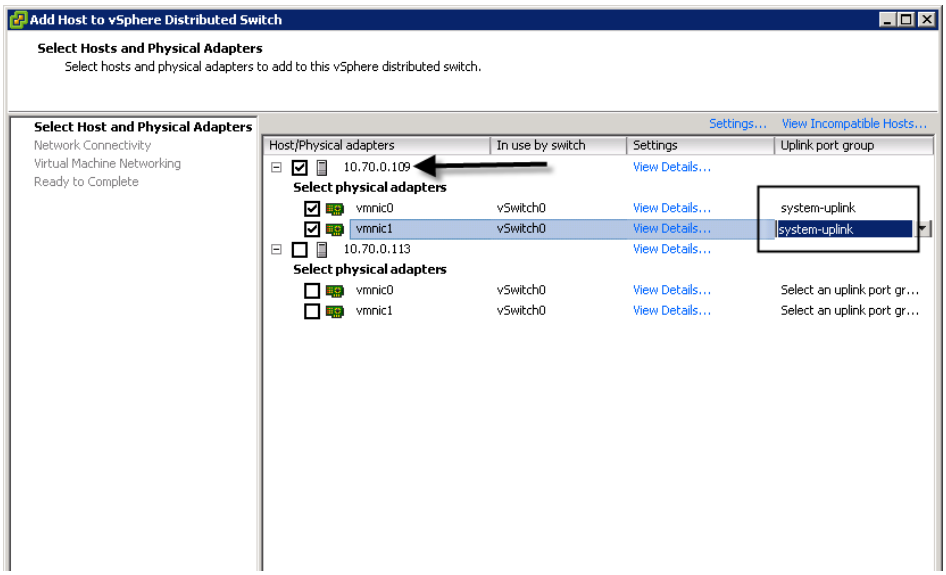

21. Network Connectivity tab select Destination port group for vmk0, then click Next

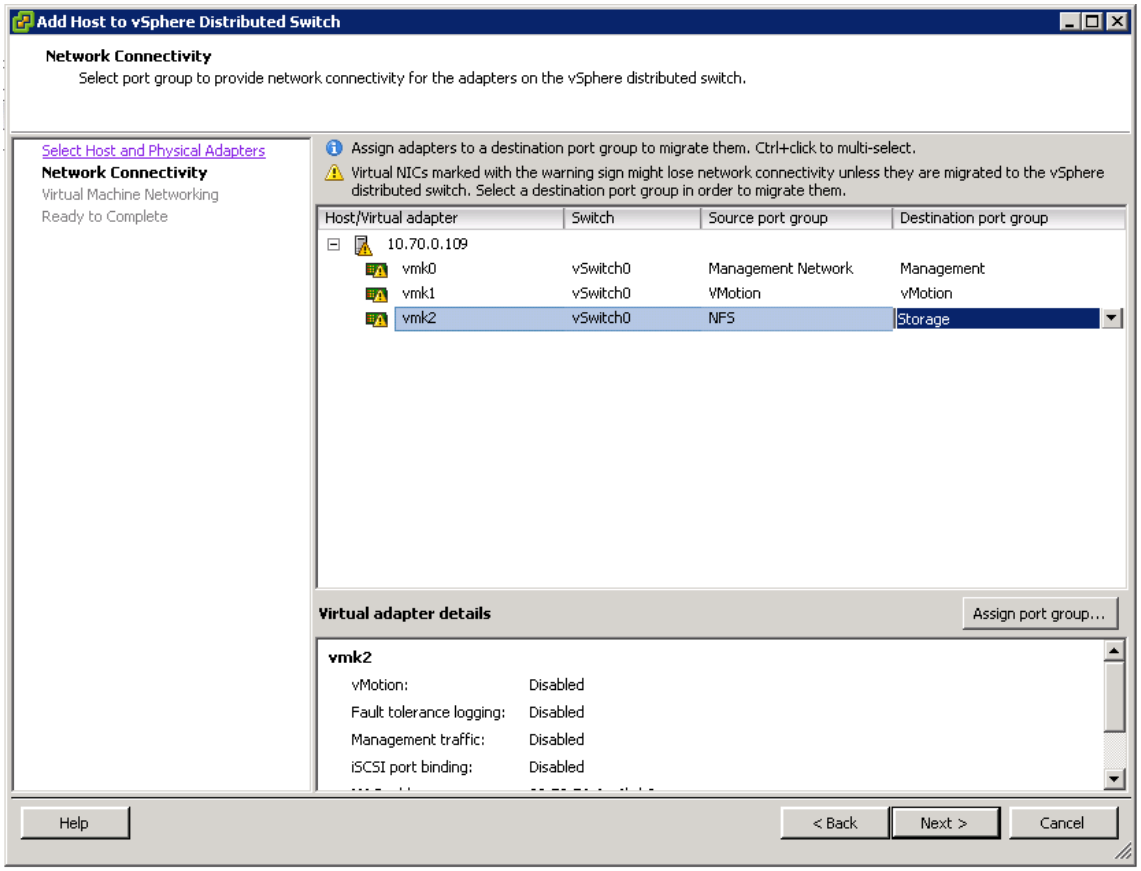

22. On the tab for virtual machine networking select VMs and assign them to a destination port-group if there is any. Otherwise click Next to Ready to complete.

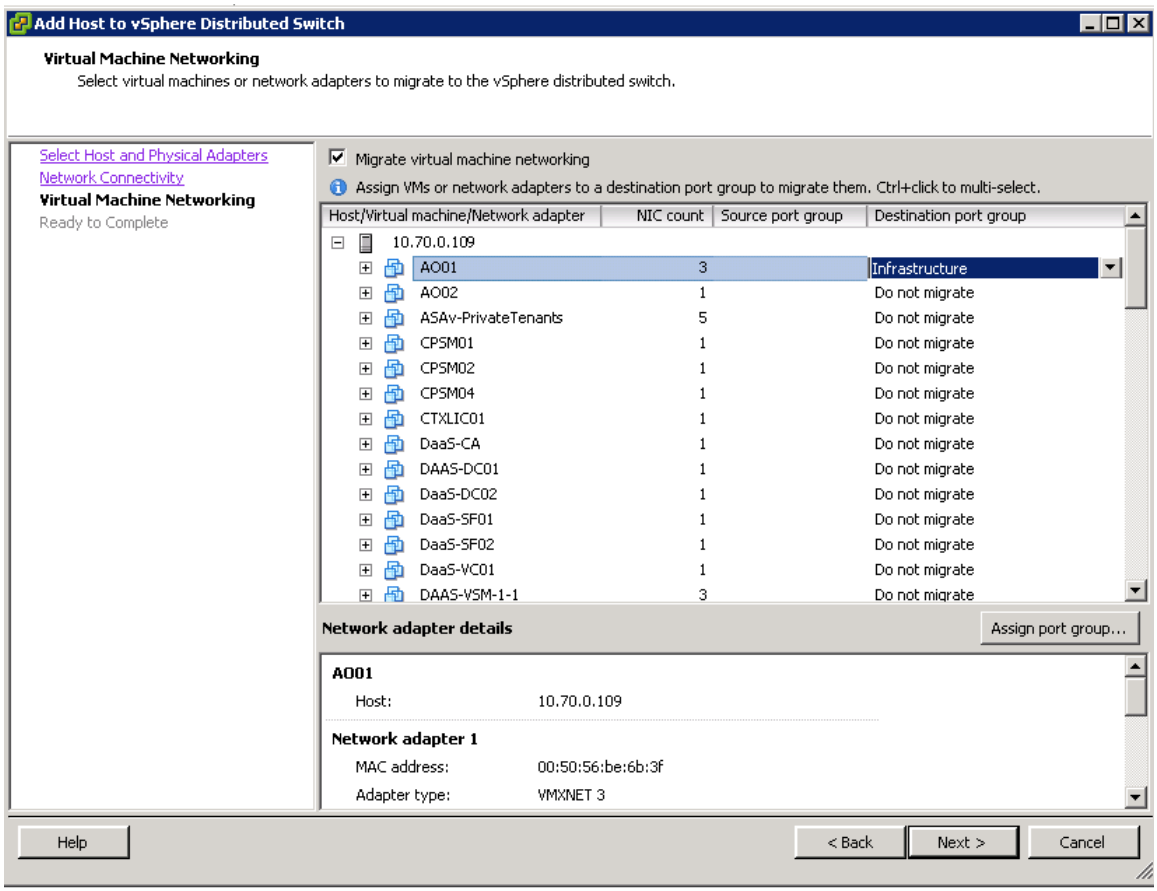

23. Verify the Settings and Click Finish to add the ESXi host part of N1KV DVS.

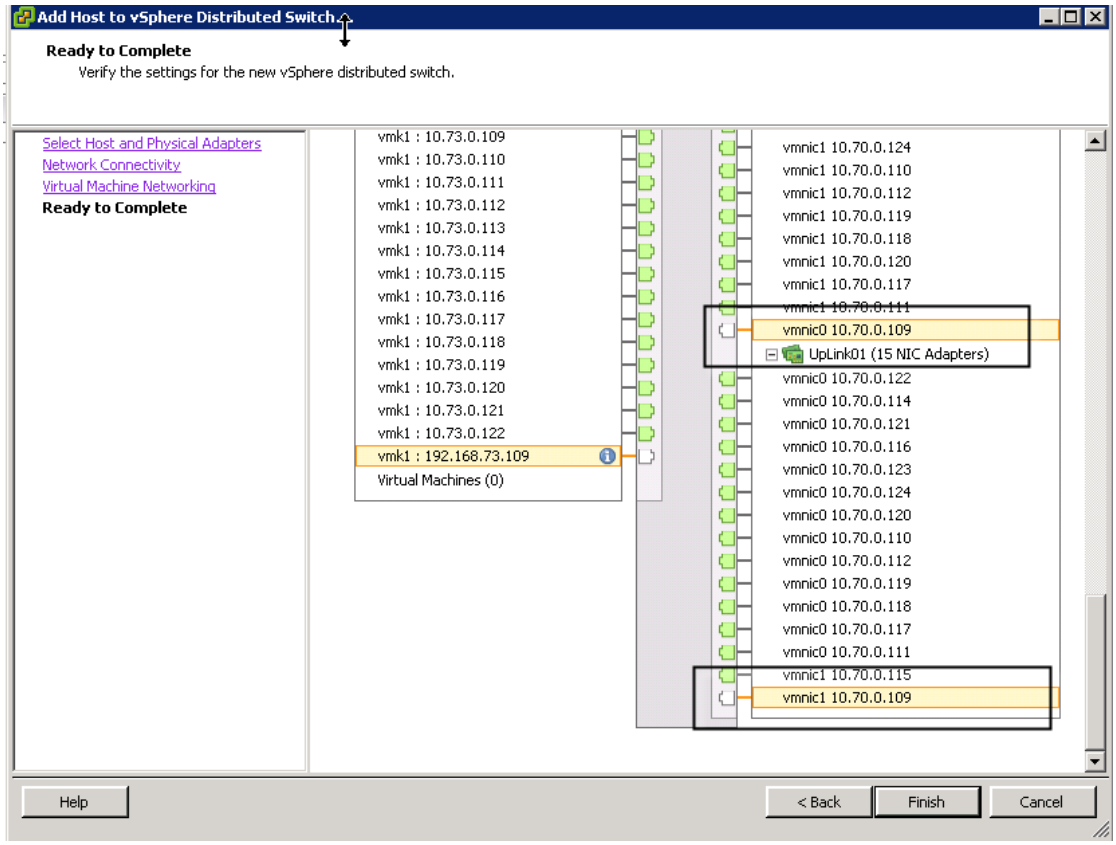

Note This will invoke VMware update manger (VUM) to automatically push the VEM installation for the selected ESXi hosts. After successful staging, install and remediation process, now the ESXi host will be added to N1KV VSM. From the vCenter task manager, quickly check the process of VEM installation.

In the absence of Update manager:

1. Upload vib file **cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib** for VEM installation to local or remote datastore which can be obtained by browsing to the management IP address for N1KV VSM.

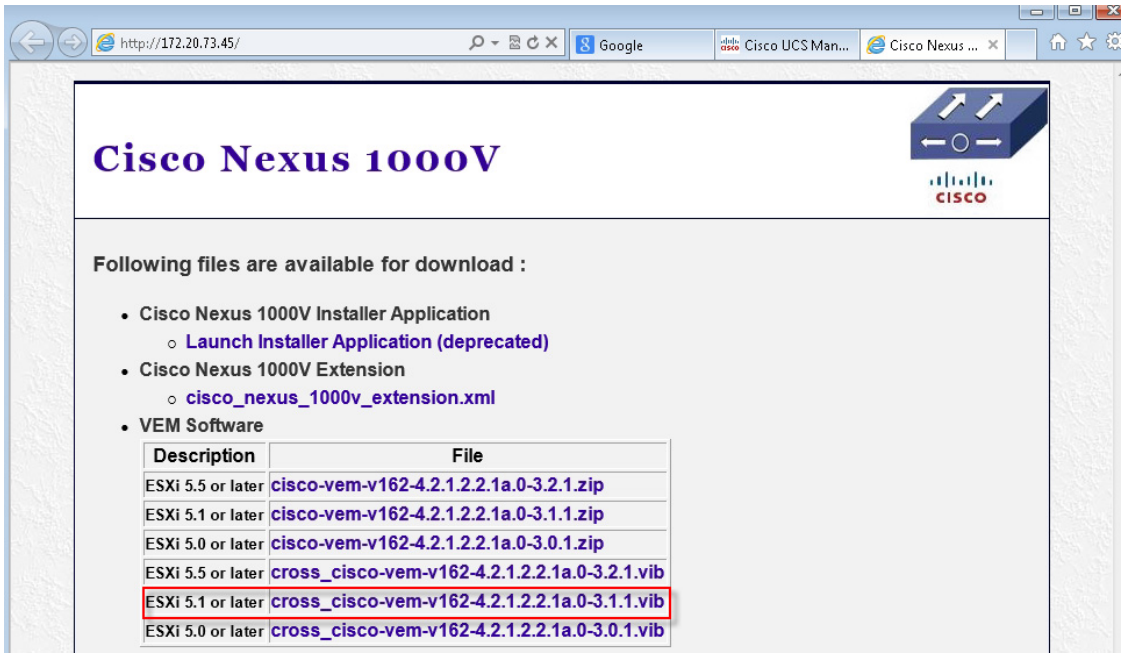

2. Login to each ESXi host using ESXi shell or SSH session.
3. Run following command:

```
esxcli software vib install -v /vmfs/volumes/ datastore/
cross_cisco-vem-v162-4.2.1.2.2.1a.0-3.1.1.vib
```

To verify the successful installation of ESXi VEM and the status of ESXi host:

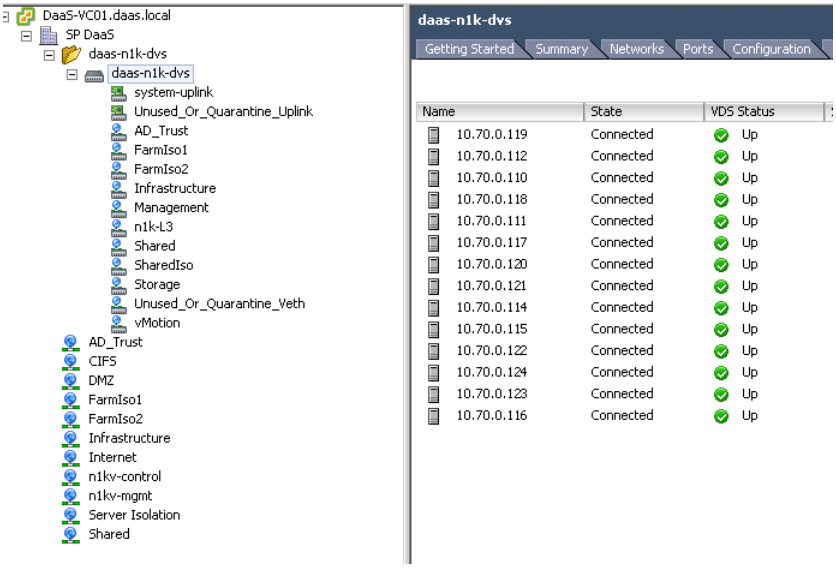

SAN Configuration

The pair of Nexus 5548UP switches was used in the configuration to connect between the 10 Gbps ethernet ports on the EMC VNX5600 and the 10 GE ports of the UCS 6248 Fabric Interconnects. We also used a pair of 8 Gbps Fiber Channel ports for Boot from SAN.

Boot from SAN Benefits

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

The key benefits of booting from the network:

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

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

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

Configuring Boot from SAN Overview

In this project we used a custom workflow in UCS Director to deploy our hypervisor to the blade hardware as well as configure the Nexus 5548Ups switches for zoning and the EMC VNX5600 for boot LUNs. Earlier in this solution we had configured our storage pools on the VNX5600 as part of a prerequisite for UCS Director workflows. The following are the prerequisite steps we completed as part of our "Day-0" tasks for UCS Director. These settings along with the custom workflow in UCS Director allowed us to configure a Boot from SAN solution and install the ESXi Hypervisor.

1. Create VSAN-A & B for SAN on Nexus 5548UP
2. Configuring Boot from SAN on EMC VNX
3. Cisco UCS configuration of Boot from SAN policy in the service profile template

In each of the following sections, each high-level phase will be discussed.

Create VSAN-A & B for SAN on Nexus 5548UP

For a detailed Cisco Nexus 5500 series switch configuration, refer to Cisco Nexus 5500 Series NX-OS SAN Switching Configuration Guide. (See the Reference Section of this document for a link.)

Configuring Boot from SAN on EMC VNX

To configure boot from SAN LUNs on EMC VNX, complete the following steps:

1. Create a storage pool from which LUNs will be provisioned. RAID type, drive number and type are specified in the dialogue box below. Five 300GB SAS drives are used in this example to create a RAID 5 pool. Uncheck "Schedule Auto-Tiering" to disable automatic tiering.

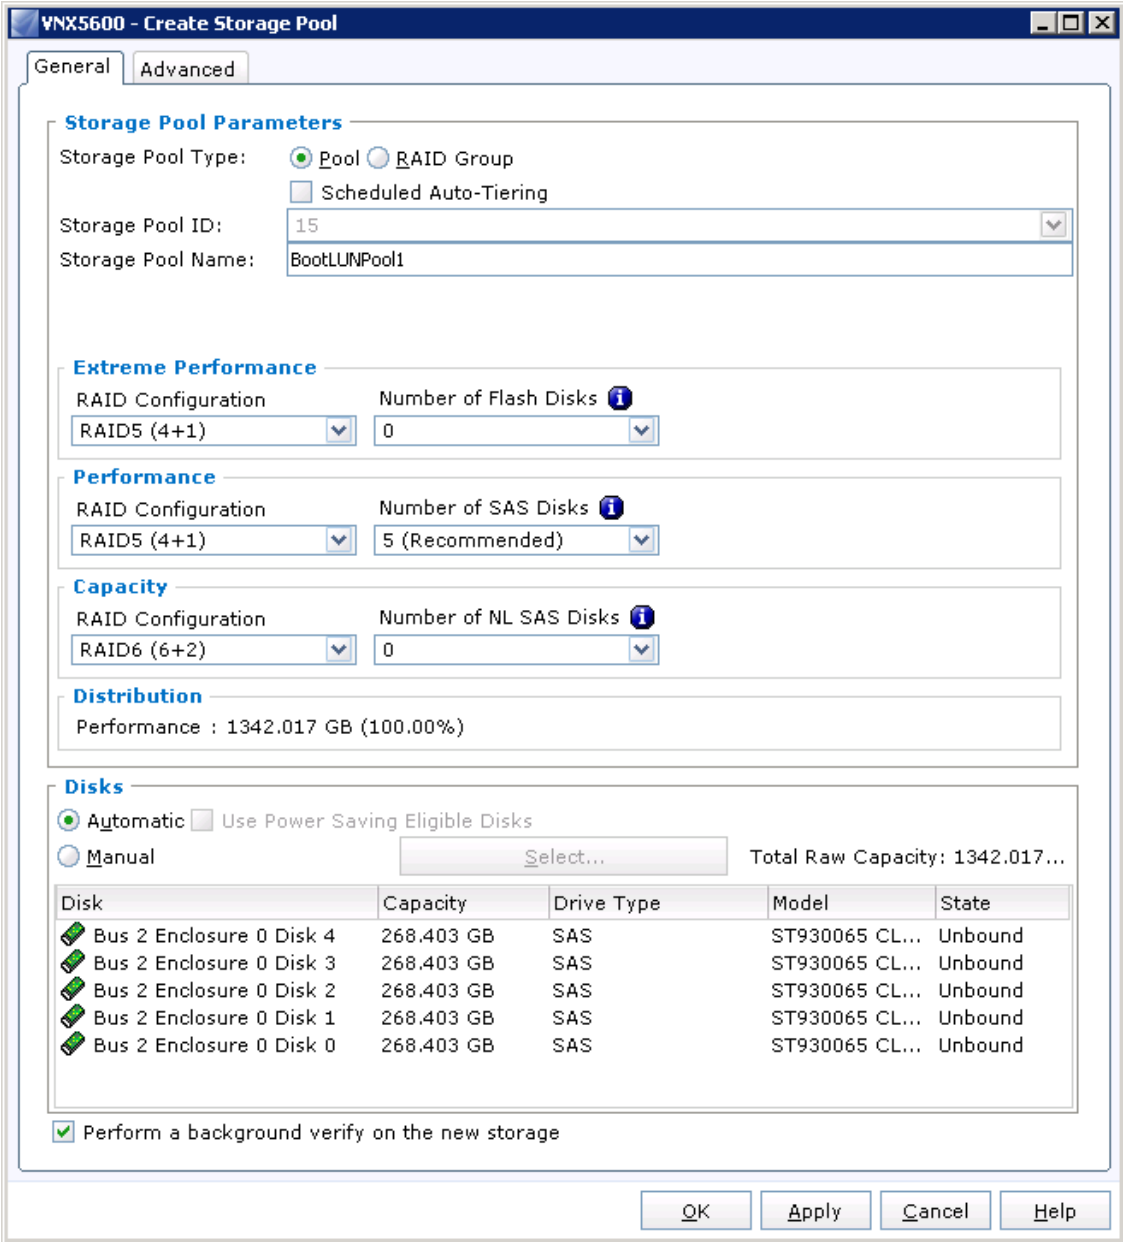

2. Provision LUNs from the storage pool and assign host-to-LUN mapping using UCS Director from here. After storage pools are created, our UCS Director workflow will handle creating the boot luns as well as VSAN configuration. In a typical Cisco CVD we would detail the steps we take to create Service Profiles and templates with a Boot from SAN Policy as well as how to connect and use them. This CVD utilized UCS Director for building out ESXi hosts. The Workflow for UCS Director handled the, usually manual tasks related to create a boot LUN and mount it to the blade host. See the UCS Director section for workflow details.

Solution Validation: Citrix DaaS Implementation

This CVD solution allows service providers to deliver Windows applications and desktops as Desktop-as-a-Service (DaaS) through an integrated set of Cisco, Citrix, and partner technologies. The procedures in this section build out the Citrix software components. These same procedures were followed to deploy the environment for scalability testing.

Topology Overview of the Validated Solution

The DaaS solution documented in this CVD constructs a provider environment that can support both shared and private delivery sites, as shown in the diagram below. The procedures create a shared delivery site that provisions hosted desktops and applications to shared tenants (using Private Delivery Group/Shared Delivery Site isolation). In addition, zero trust mechanisms are used to provision hosted desktops and applications to private site tenants (using Private Delivery Group/Private Delivery Site isolation).

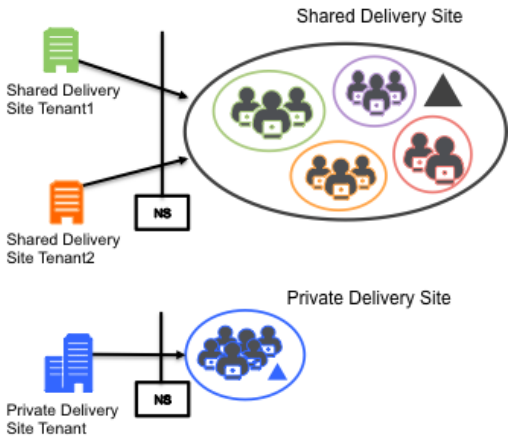

Both models are implemented using the Citrix software components:

- CloudPortal Services Manager simplifies the management of tenant on-boarding and user subscriptions. Delegated management roles allow tenant administrators to self-provision and monitor provisioning requests.
- Citrix App Orchestration allows providers to automate and manage the delivery of desktop and application offerings in a comprehensive multi-tenant environment that uses an array of isolation
- models. It enables a common management interface across all managed tenants.
- Citrix XenDesktop supports the delivery of hosted applications and desktops, supplying both Hosted Shared Desktop (HSD) and Server Virtual Desktop Infrastructure (VDI) services.
- Citrix NetScaler provides secure access to the Service Provider domains over SSL (TCP 443) across the public Internet.

Within this architecture, App Orchestration supplies the provider with a single unified interface to manage services for all tenants regardless of whether they are shared or private delivery sites. The provider can operate hosted desktops and application services and deliver capabilities to multiple tenants within a shared site as well as providing managed services to tenants in private sites.

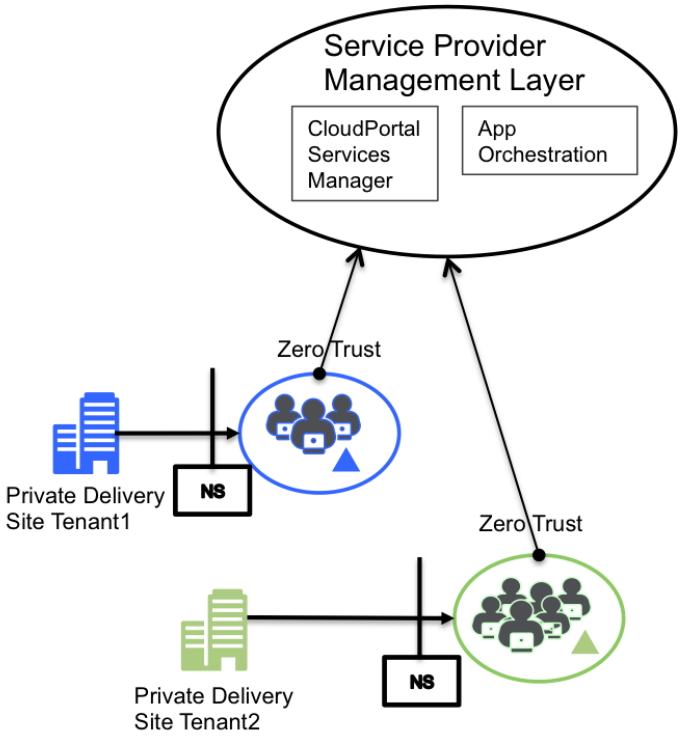

The diagram above emphasizes the relationship between a service provider and two managed private delivery site tenants. App Orchestration 2.5 provides a zero trust agent that simplifies connectivity between the App Orchestration configuration server and orchestrated delivery controllers. Domain trusts are no longer required between the target orchestrated domain and the App Orchestration domain.

The CVD creates an environment designed to support 10 tenants (2 private delivery site/8 shared delivery site) and 2000 users. It assumes a mixed user workload of 90% HSD and 10% Server VDI (SVDI). The diagram below illustrates the topology for the full implementation. It shows logical functions in the architecture and the required virtual servers, illustrating the overall datacenter implementation for the DaaS provider.

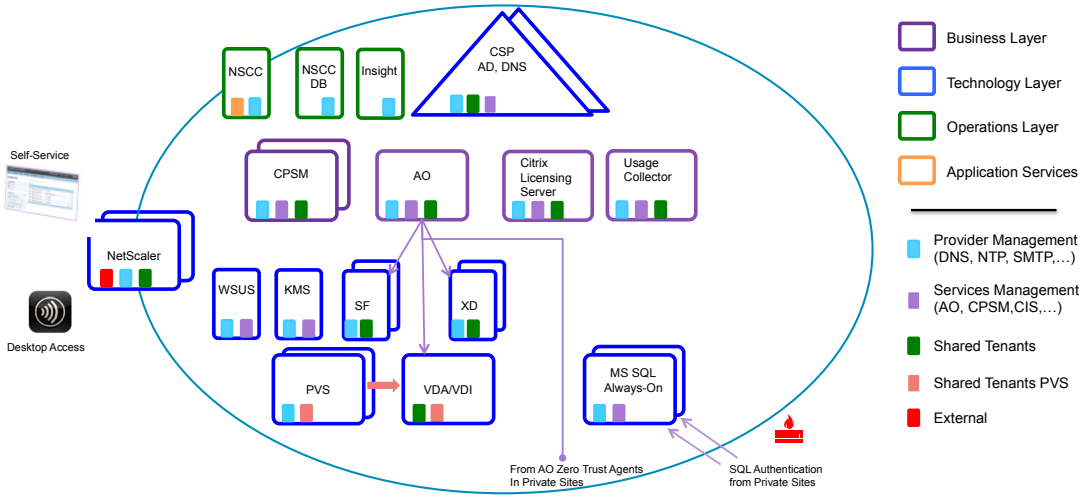

Critical Concepts

The following concepts and definitions are critical in understanding how to deploy this DaaS solution at provider sites:

- Zero Trust. By installing SSL certificates on all servers that communicate with App Orchestration, the service provider can create a solution that easily integrates with existing Active Directory domains. The Zero Trust Agent establishes a secure SSL-encrypted communication channel to the App Orchestration configuration server and authenticates using certificates. In the CVD implementation described here, zero trust communication is enabled between the shared delivery site infrastructure domain (daas.local) and the private delivery site tenant domains (e.g., fi1.local and fi2.local) .
- Unified tenant management with tiered and delegated tenant administration and self-service capabilities. App Orchestration supplies the provider with a common management interface for both shared and private tenants, allowing the provider to onboard tenant subscribers, provision session services, and manage offerings from a single interface for all tenants. In addition, a provider can choose to delegate subscriber management capabilities and enable self-service provisioning of hosted applications and desktops. CloudPortal Services Manager implements delegated management by defining tenant-specific administrators that can approve workflows for self-service hosted application and desktop provisioning.
- Network isolation. It is assumed that virtual networks are implemented to provide the necessary network security isolation. VLANs deployed for this CVD include:
  - Private VLANs for each Private Delivery Site tenant
  - Shared Tenant VLAN for Shared Delivery Site infrastructure and tenants
  - Provider Management VLAN (for DNS, NTP, SNMP, etc.)
  - Services Management vLAN
  - (for services such as App Orchestration, CPSM, and licensing)
  - Provisioning vLANs (one for PVS provisioning to the Shared Delivery Site tenants as well as one for each Private Delivery Site tenant)
  - Application vLAN (for back office applications)
  - Storage network (for resource separation)

The following documentation is useful when planning a deployment.

- Terminology in App Orchestration 2.5
- Getting Started with Citrix App Orchestration 2.5
- Known Issues for App Orchestration 2.5
- CloudPortal Services Manager 11.0 Documentation
- Known Issues for CloudPortal Services Manager
- Configuring SSL for App Orchestration 2.5

It is highly recommended that providers take advantage of experienced consultants in the Citrix Services organization to plan, perform, and assist with Citrix software installation and integration tasks. In doing so service providers can be confident in achieving an optimal deployment configuration that provides rigorous security, optimal scale, and ease of ongoing management and customer onboarding.

Prerequisites

The installation and integration procedures for Citrix App Orchestration and CloudPortal Services Manager in this environment depend on a number of prerequisites and assumptions. Refer to the documentation above for checklists and detailed explanations of the requirements and assumptions. Specifically:

- For CloudPortal Services Manager, see System Requirements for CloudPortal Services Manager, Firewall requirements for Services Manager, and Verify deployment readiness and create system databases.
- For App Orchestration, see the Setup Checklist in Getting Started with Citrix App Orchestration 2.5. The procedures here assume the checklist items have been completed in advance.

Prior to installing App Orchestration and CloudPortal Services Manager for this CVD, there are a few additional setup procedures that can be performed out-of-band:

- Installing and configuring Citrix Provisioning Services 7.1
- Installing Virtual Delivery Agents on Microsoft Windows server and workstation operating systems
- Setting up and configuring NetScaler functionality

These procedures are included in subsequent sections.

Shared Delivery Site VMs

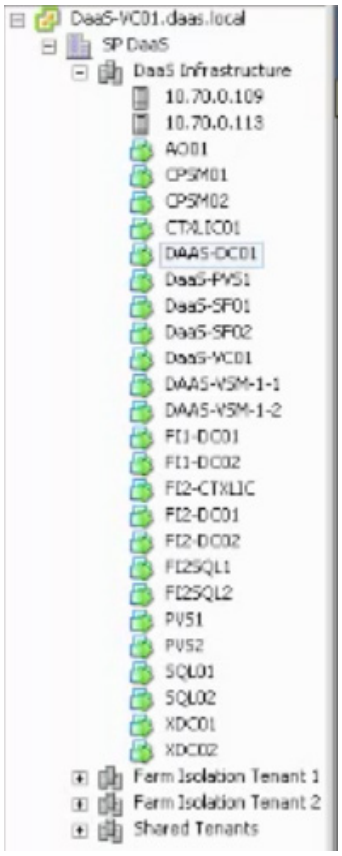

The procedures in this section assume that Cisco UCS Director has provisioned the following VMs (via VMware ESXi) with Windows Server 2012 R2 (minimum) to support Citrix software and infrastructure services:

- 1 domaincontroller with Windows Server 2012 R2
- 1 Citrix License Server
- 1 NetScaler Gateway
- 2 Servers
- for the CloudPortal Services Manager provisioning (one can also host CloudPortal Services Manager Web services)
- 1 server for the App Orchestration configuration server
- 1 server for PVS provisioning
- 1 server (minimum) for the Session Machines that will host applications and desktops for users
- 2 database servers for each delivery site running Microsoft SQL Server 2012 with all recommended updates installed
- 2 servers for the Delivery Controllers that make up each delivery site
- 2 servers for the StoreFront servers that make up one StoreFront server group

Private Delivery Site VMs

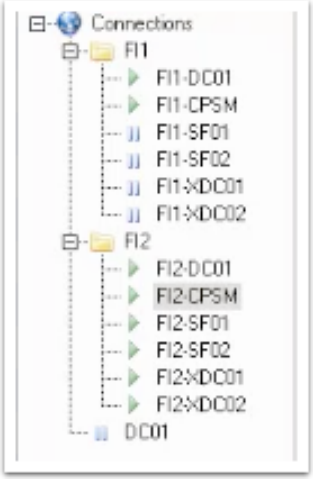

As shown in the screenshot, the CVD defines two private delivery site tenants: Farm Isolation Tenant 1 (FI1) and Farm Isolation Tenant 2 (FI2). Each private delivery site tenant has these dedicated infrastructure components:

- 1 domain controller with Windows Server 2012 R2
- 1 server to host CloudPortal Services Manager and the App Orchestration Zero Trust (Domain) Agent
- 2 servers for the StoreFront servers that make up one StoreFront server group
- 2 servers for the XenDesktop Delivery Controllers
- 1 server (minimum) for the Session Machines that will host applications and desktops for users

In addition, each tenant uses a dedicated Active Directory domain for user and group data. Tenant user data is isolated using Organizational Unit (OU) lockdown provided by CloudPortal Services Manager. Users and groups created under each tenant's OU are locked down to fully isolate one tenant from another.

Other Prerequisites and Setup Tasks

The following is a quick synopsis of additional prerequisites and assumptions:

- **Accessible media.** It is assumed that Citrix software installation packages are available on a partition accessible to the servers.
- **Active Directory.** Prepare the required Active Directory domain to be used as the shared resource domain (Windows Server 2012 R2 was used). Extend the Active Directory schema to include the standard Exchange attributes and prepare the environment for multi-tenancy.
- **Group Policy object.** Configure the App Orchestration Group Policy object that will be associated with all machines in the shared resource domain. Configure the PowerShell execution policy, PowerShell remoting, and remote administration with WMI.
- **DNS aliases.** Services Manager uses DNS aliases internally for the core components. Create CNAME records for the following roles and components:

Platform component	Alias
Database server	CORTEXSQL
Provisioning server	CORTEXPROVISIONING
Web server	CORTEXWEB
Reporting Services	CORTEXREPORTS

The XenDesktop zero trust agent in a private Delivery Site domain must have network access to the App Orchestration configuration servers and must be able to resolve the FQDN address of the configuration servers. To do this, create a DNS forward lookup zone on the DNS server for the isolated domain. For more information see
http://support.citrix.com/servlet/KbServlet/download/37594-102-711575/cao-zero-trust-agent.pdf

- **Other CloudPortal Services Manager requirements.** The Services Manager server role installer (Setup Tool) handles many prerequisites, such as installing .NET Framework 4.0, enabling Web Server roles, and enabling MSMQ features.
- **SSL certificates.** An SSL certificate is used to secure communication:
  - On each StoreFront server.
  - Between an App Orchestration agent on each XenApp or XenDesktop delivery controller and the domain agent installed on a dedicated machine residing behind a NAT-enabled device.
  - Between the CloudPortal Services Manager web console and the App Orchestration configuration server (to secure administrative tasks).

The SSL certificate is installed on the App Orchestration server, as described later in this section. See Configuring SSL for App Orchestration 2.5 for more information. A public SSL certificate is also installed on the NetScaler Active Gateway to secure access from user devices.

PVS Installation and Configuration

PVS installation can occur out–of-band from the installation of App Orchestration and CloudPortal Services Manager. This CVD assumes that UCS Director has allocated VMs for the PVS servers and the following procedures have been used to install PVS on those VMs.

Prerequisites

PVS software and hardware requirements are available at http://support.citrix.com/proddocs/topic/provisioning-7/pvs-install-task1-plan-6-0.html.

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

Instructions	Visual
From the Citrix Provisioning Services 7.1 ISO, let AutoRun launch the installer. Click the **Server Installation** button.	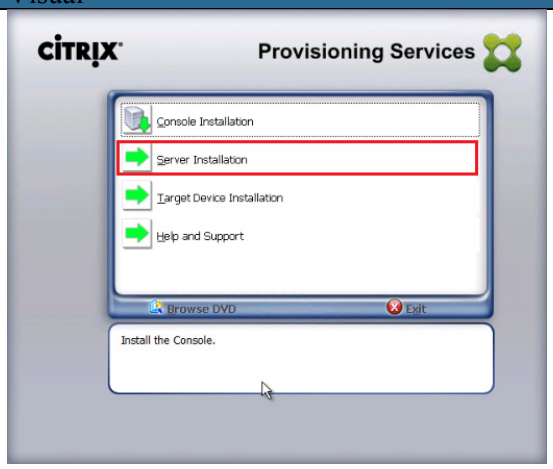
Click the **Install Server** button. The installation wizard will check to resolve dependencies and then begin the PVS server installation process. It is recommended that you temporarily disable anti-virus software prior to the installation.	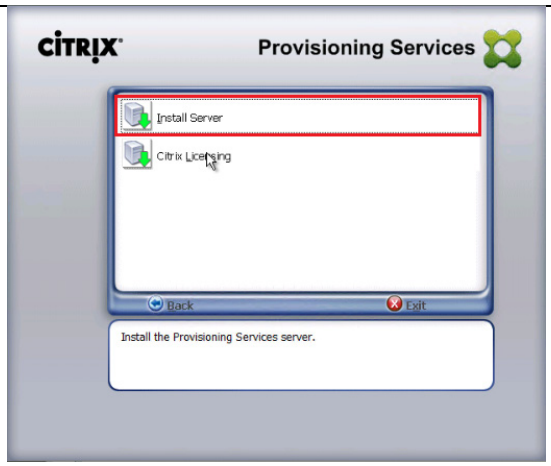

Click **Install** on the prerequisites dialog.	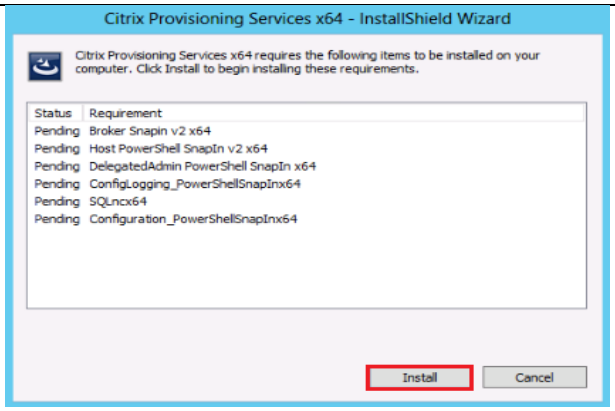
Click **Yes** when prompted to install the SQL Native Client.	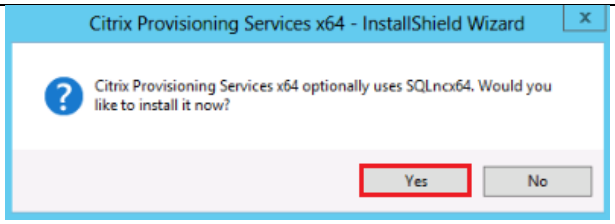
Click **Next** when the Installation wizard starts.	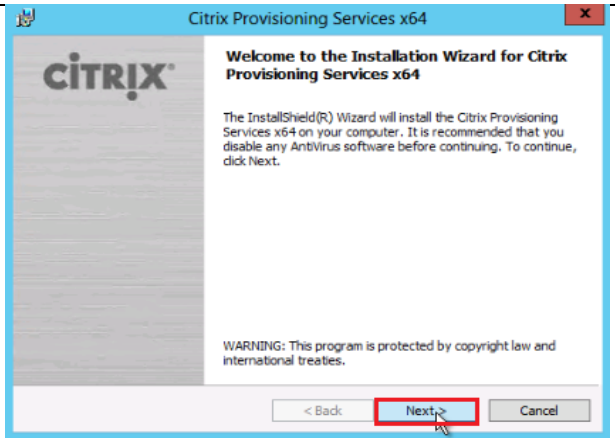
Review the license agreement terms. If acceptable, select the radio button labeled **"I accept the terms in the license agreement**." Click **Next**	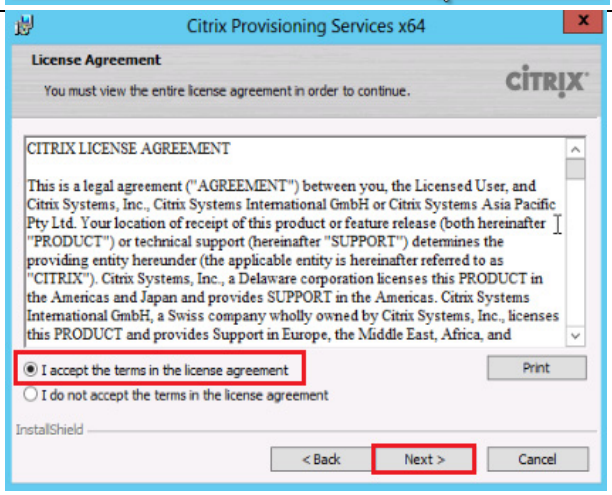

Provide User Name, and Organization information. Select who will see the application. Click **Next**	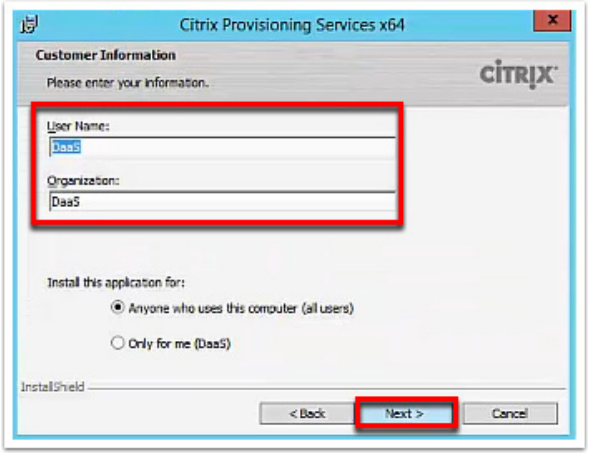
Accept the default installation location. Click **Next**	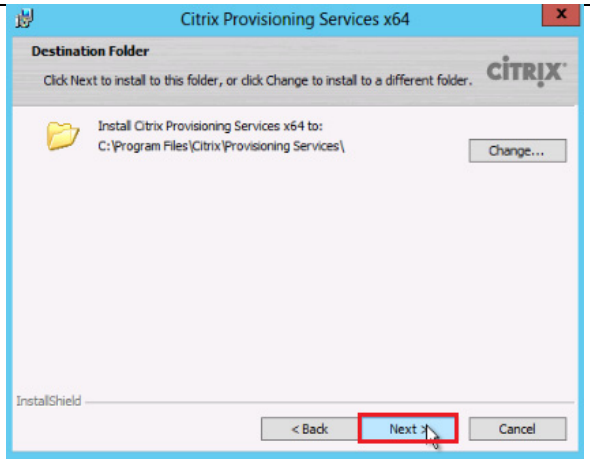
Click **Install** to begin the installation.	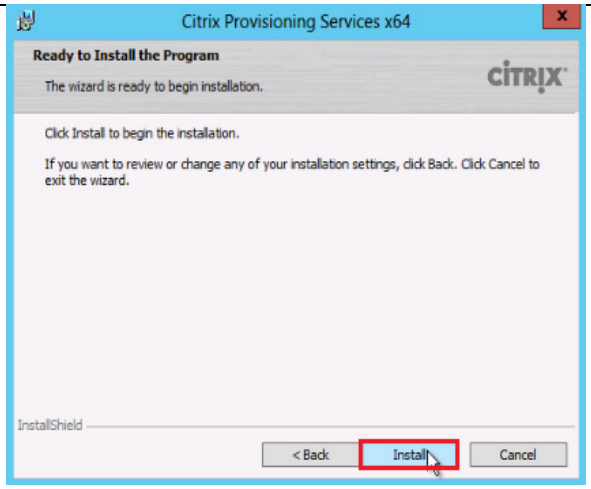

Click **Finish** when the install is complete.	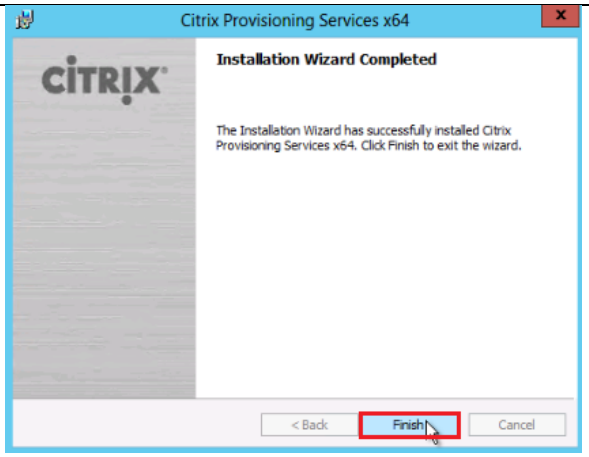
Click **OK** to acknowledge the PVS console has not yet been installed.	
The PVS Configuration Wizard starts automatically. Click **Next**	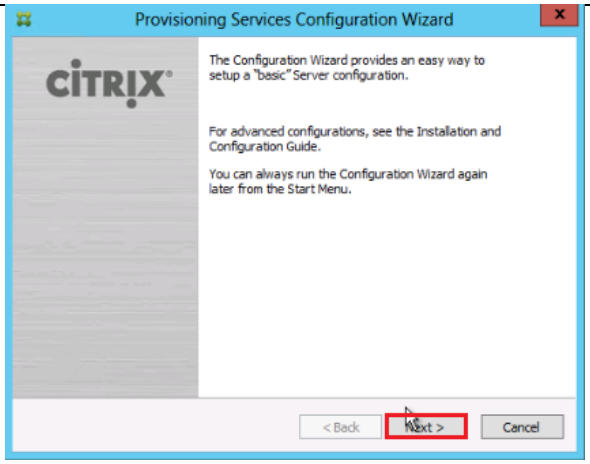
Since the PVS server is not the DHCP server for the environment, select the radio button labeled, "**The service that runs on another computer.**" Click **Next**	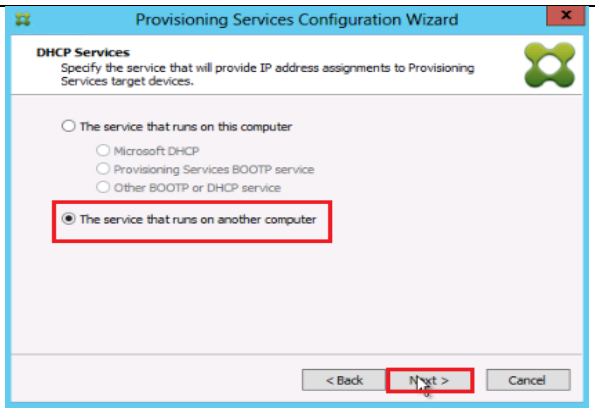

Since this server will be a PXE server, select the radio button labeled, "**The service that runs on this computer.**" Click **Next**	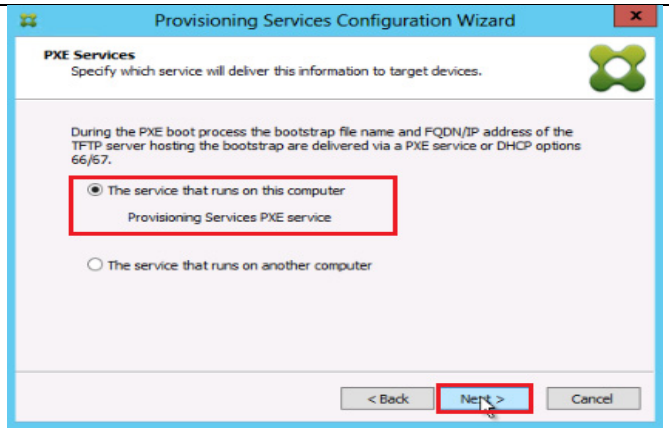
Since this is the first server in the farm, select the radio button labeled, "**Create farm**". Click **Next**	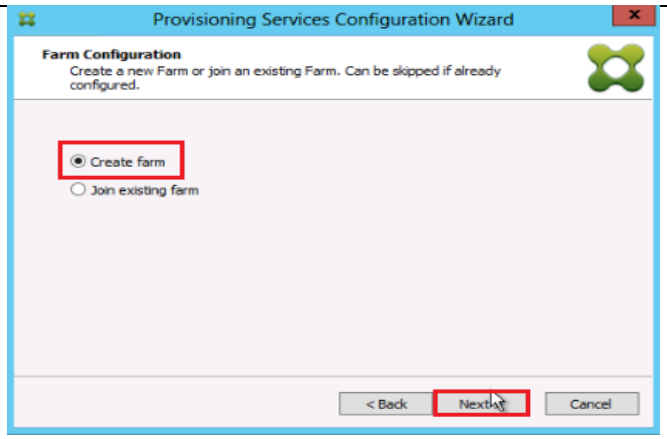
Enter the name of the SQL server. Note: If using a cluster, instead of AlwaysOn groups, you will need to supply the instance name as well. Click **Next**	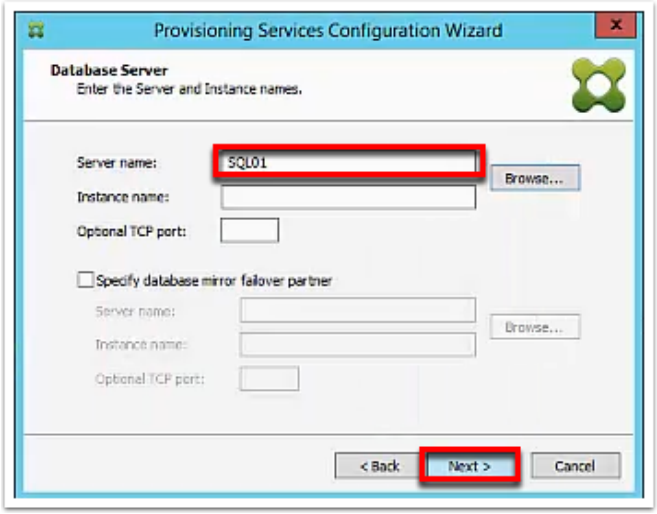

Optionally provide a Database name, Farm name, Site name, and Collection name for the PVS farm. Select the **Administrators** group for the Farm Administrator group. Click **Next**	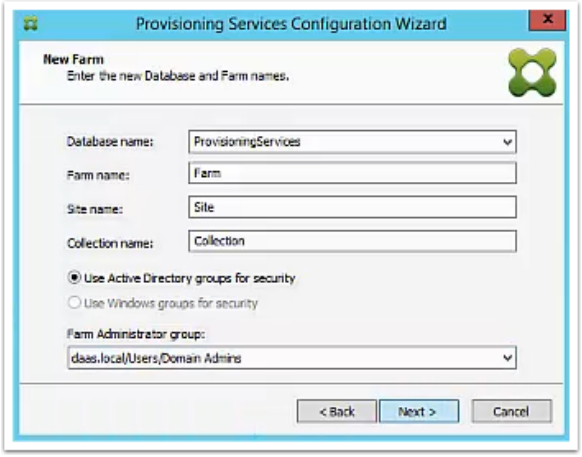
Provide a vDisk Store name and the storage path to the vDisk share. Click **Next**	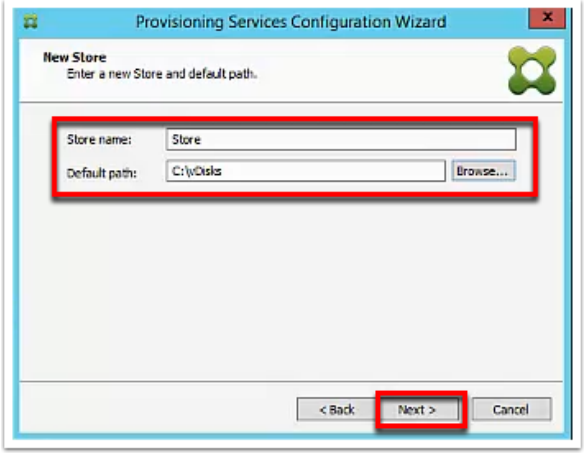
Provide the FQDN of the License Server. Optionally, provide a port number if changed on the license server. Click **Next**	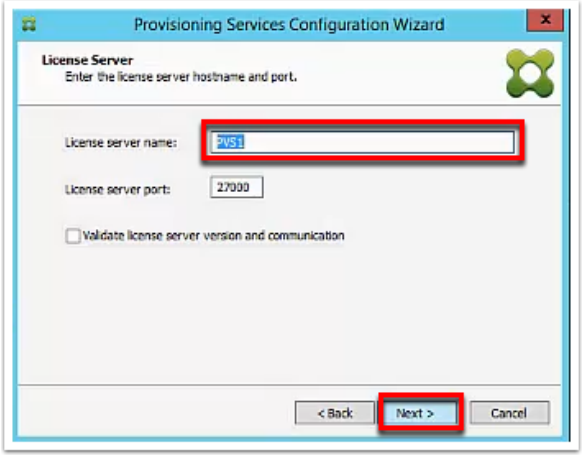

If an Active Directory service account is not already setup for the PVS servers, create that account prior to clicking Next on this dialog. Select the **Network service account** radio button. Click **Next**	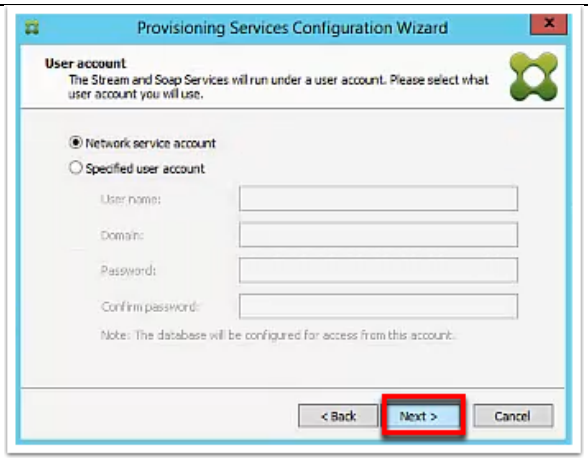
Set the Days between password updates. NOTE: This will vary per environment. “7 days” for the configuration was appropriate for testing purposes. Click **Next**	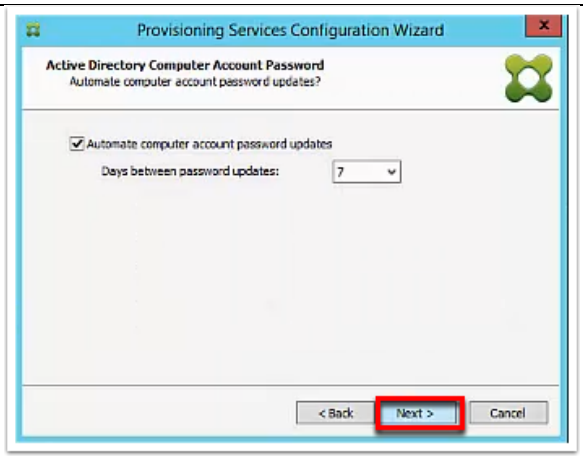
Keep the defaults for the network cards. Click **Next**	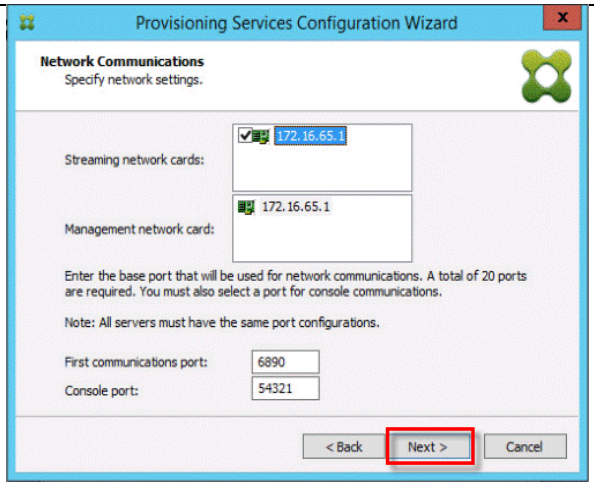

Leave the **Use the Provisioning Services TFTP service** checkbox enabled. Click **Next**	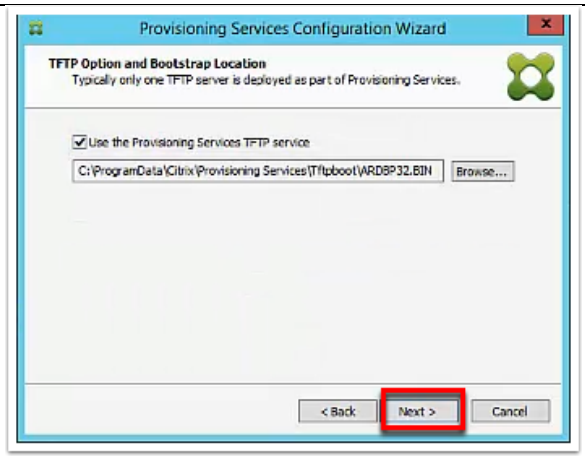
Confirm at least one boot server is listed. Add or remove boot servers as appropriate for your configuration. Click **Next**	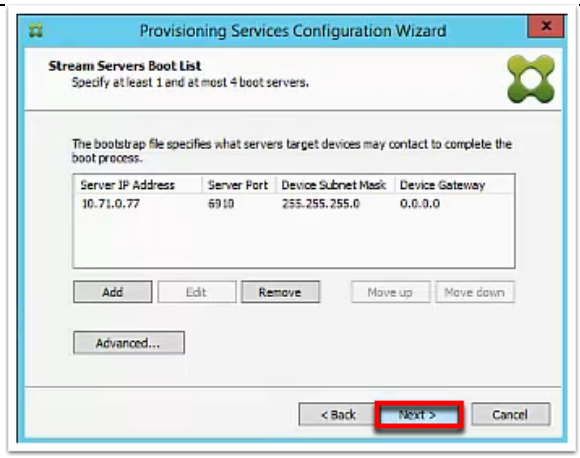
Click **Finish** to start the installation.	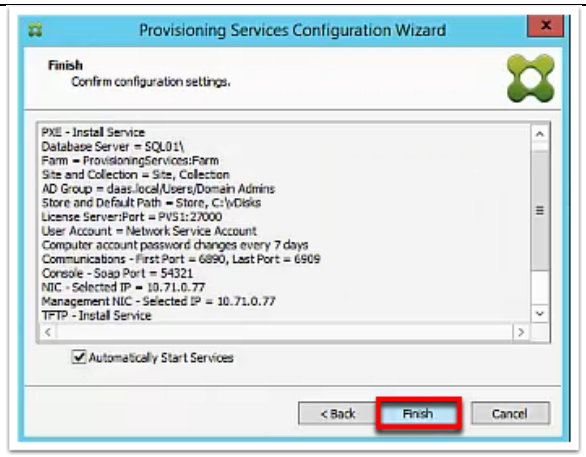

When the installation is completed, click the **Done** button.	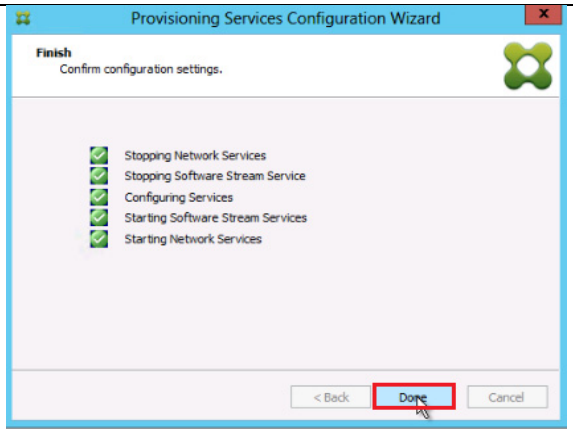
From the main installation screen, select **Console Installation**.	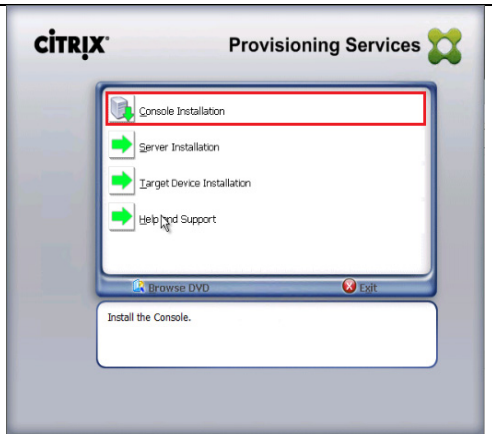
Click **Next**	

Read the Citrix License Agreement. If acceptable, select the radio button labeled "**I accept the terms in the license agreement**." Click **Next**	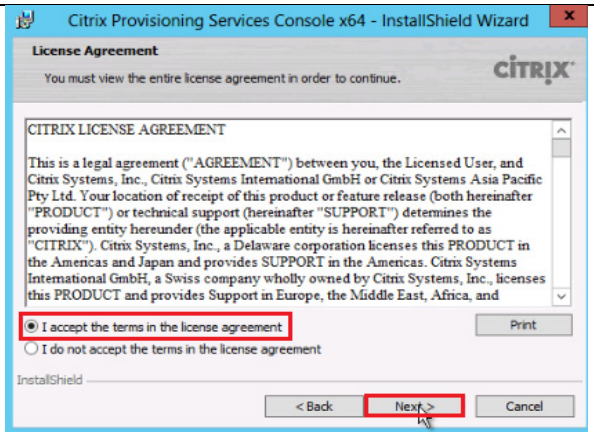
Optionally provide **User Name** and **Organization**. Click **Next**	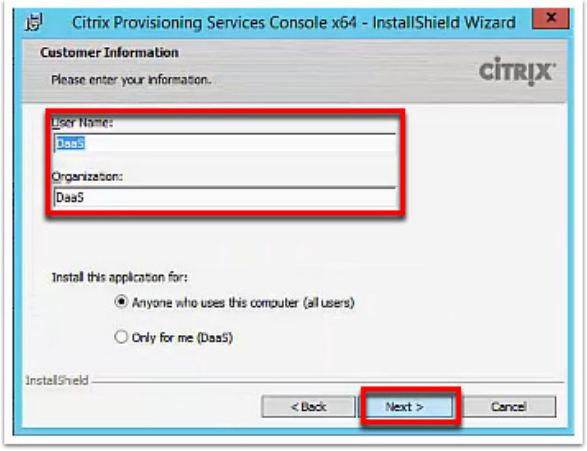
Accept the default path. Click **Next**	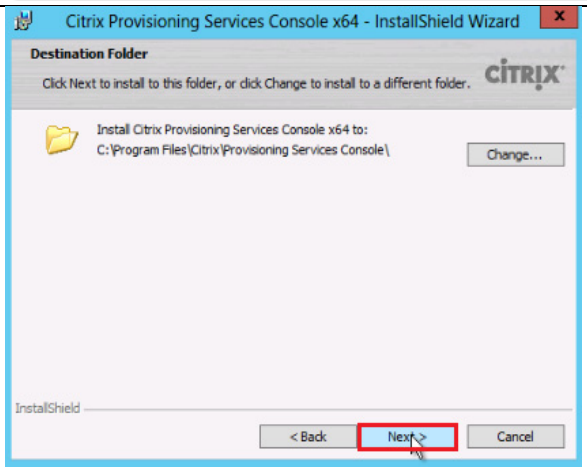

Leave the **Complete** radio button selected. Click **Next**	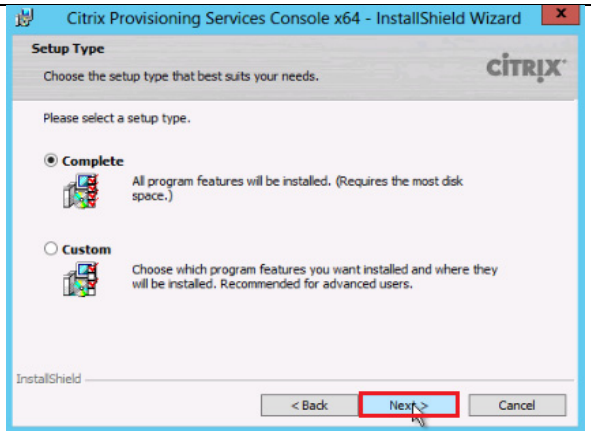
Click the **Install** button to start the console installation.	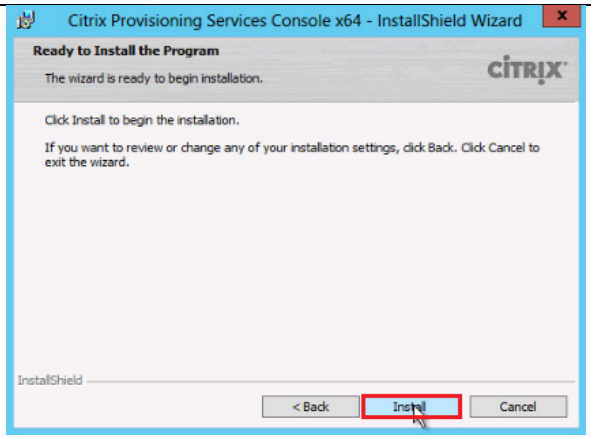
When the installation completes, click **Finish** to close the dialog box.	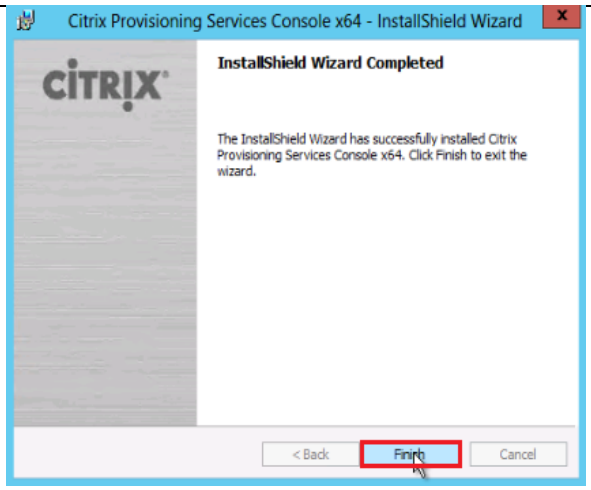

Configuring Store and Boot Properties for PVS1

Instructions	Visual
From the Windows Start screen for the Provisioning Server PVS1, launch the **Provisioning Services Console**.	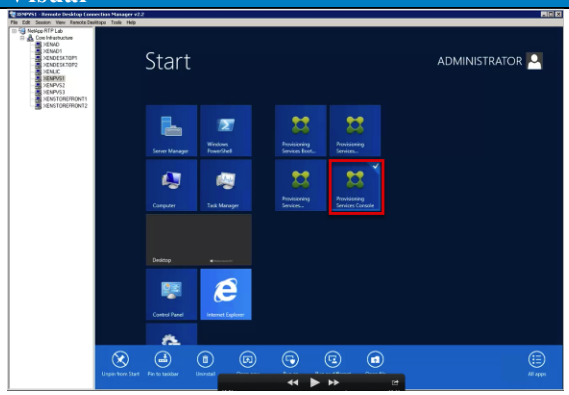
Select Connect to Farm.	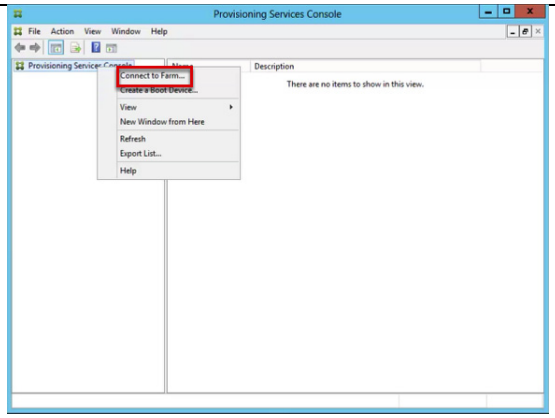
Enter **localhost** for the PVS1 server. Click **Connect**.	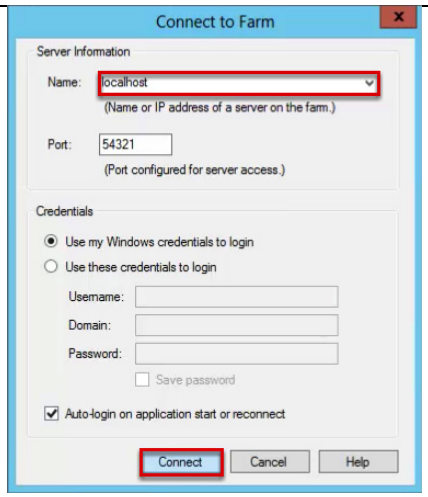

Select **Store Properties** from the pull-down menu.	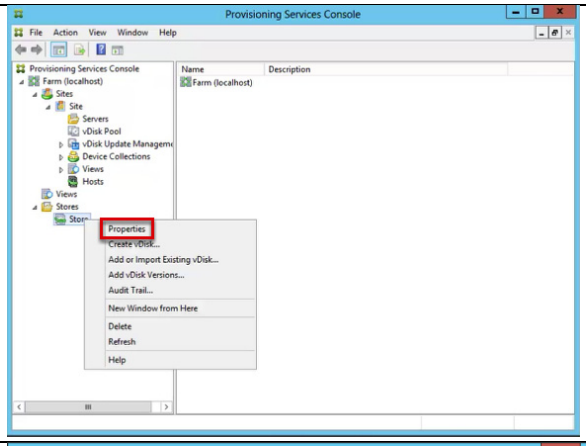
In the Store Properties dialog, add the Default store path to the list of Default write cache paths. Click **Validate**. If the validation is successful, click **OK** to continue.	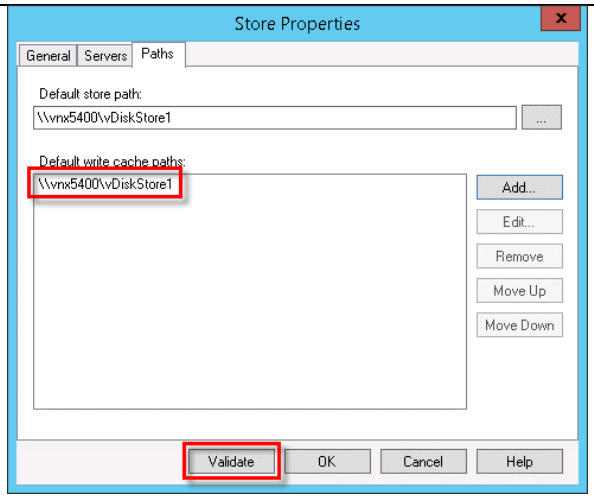

In this CVD, we repeated the procedure to add a second PVSserver (installing the Provisioning Services console on the second PVS virtual machines is optional).

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

VDA Installation and Configuration

Virtual Delivery Agents (VDAs) are installed on Microsoft Windows server and workstation operating systems, and enable connections for desktops and apps. VDA installation can occur out–of-band from the installation of App Orchestration and CloudPortal Services Manager. This CVD assumes that the following procedures have been used to install VDAs for both HSD and Server VDA environments.

By default, when you install the Virtual Delivery Agent, Citrix User Profile Management is installed silently on master images. (Using profile management as a profile solution is optional.)

Instructions	Visual
Launch the XenDesktop installer from the ISO image or DVD. Click **Start** on the Welcome Screen.	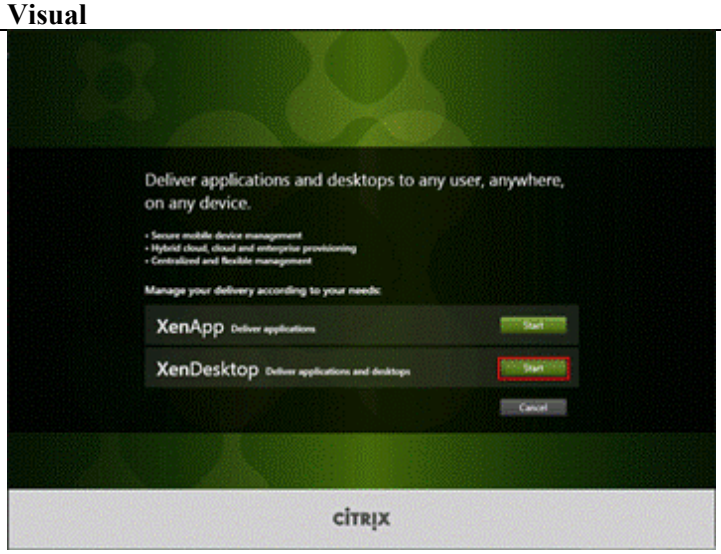
To install the VDA for the Hosted VDI Desktops, select **Virtual Delivery Agent for Windows Desktop OS**. (After the VDA is installed for Hosted VDI Desktops, repeat the procedure to install the VDA for Hosted Shared Desktops. In this case, select **Virtual Delivery Agent for Windows Server OS** and follow the same basic steps.)	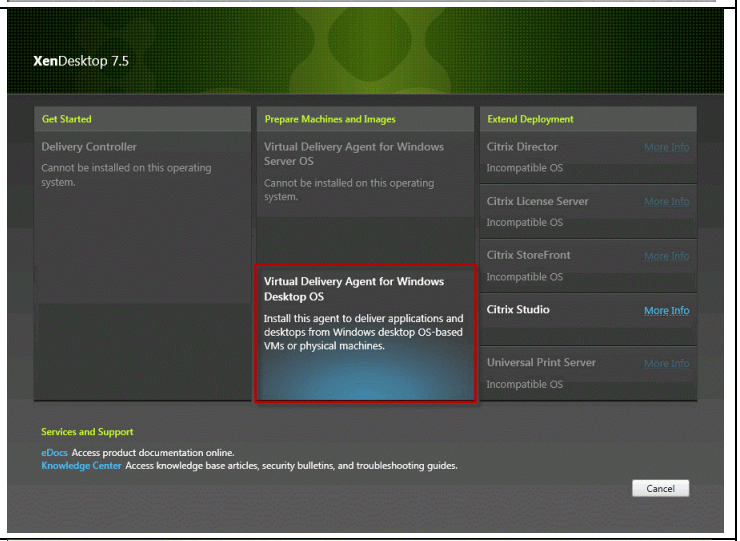
Select "**Create a Master Image".** Click **Next**	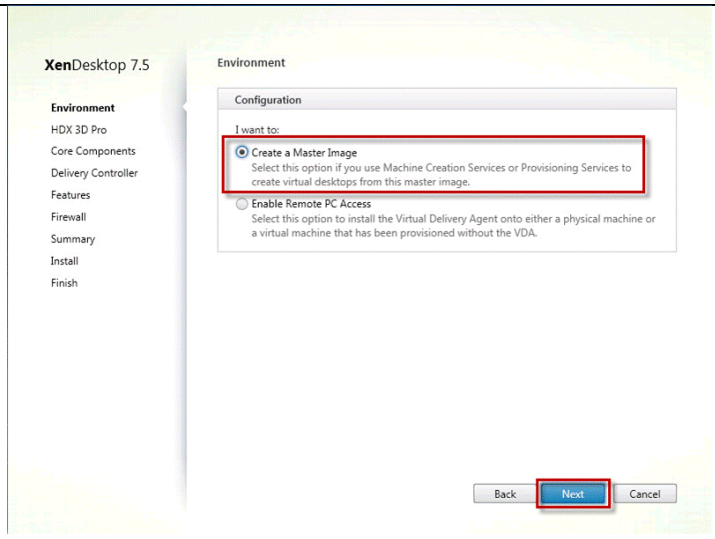

For the HVD vDisk, select "**No, install the standard VDA".** Click **Next**	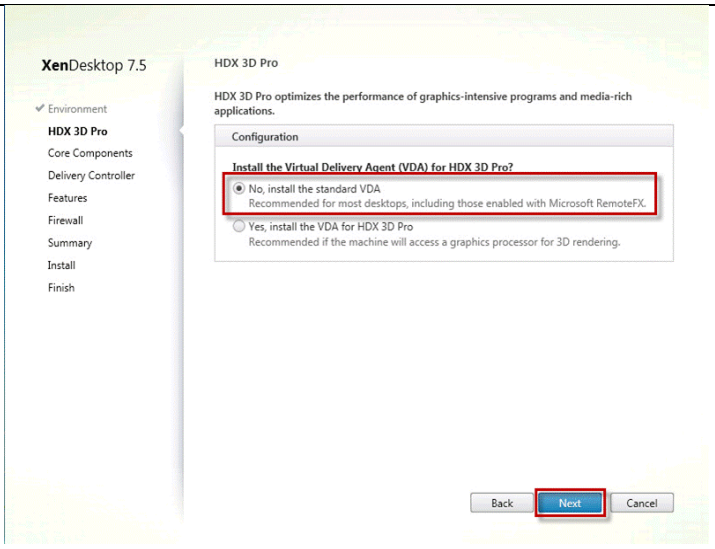
Select **Citrix Receiver.** Click **Next**	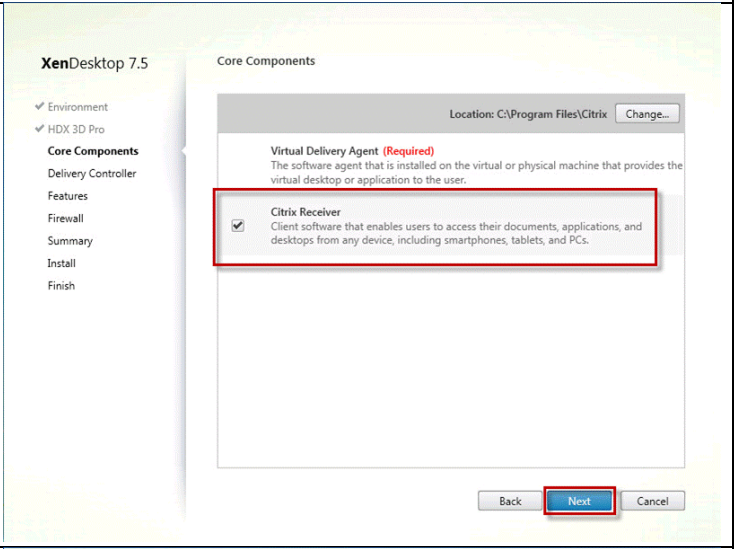
Select "**Do it manually"** and specify the FQDN of the Delivery Controllers. Click **Next**	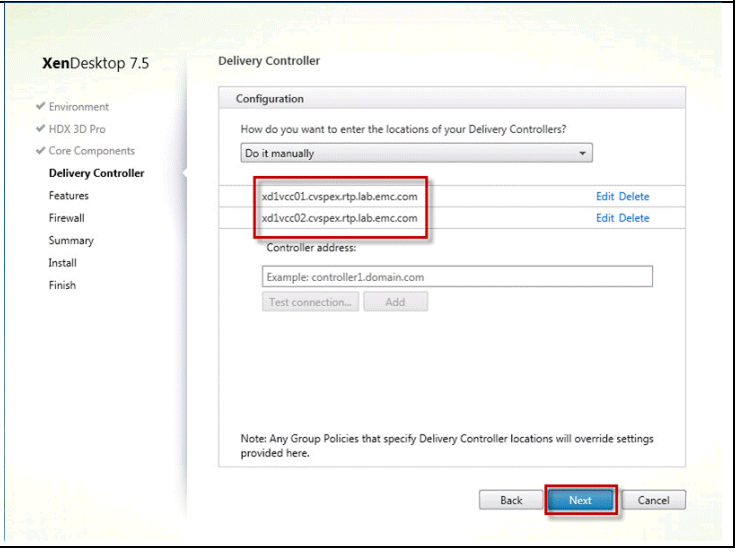

Accept the default features. Click **Next.**	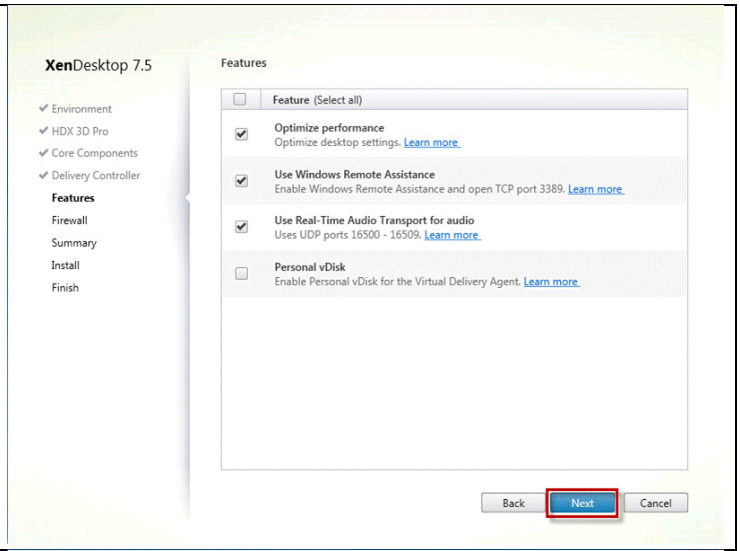
Allow the firewall rules to be configured **Automatically.** Click **Next**	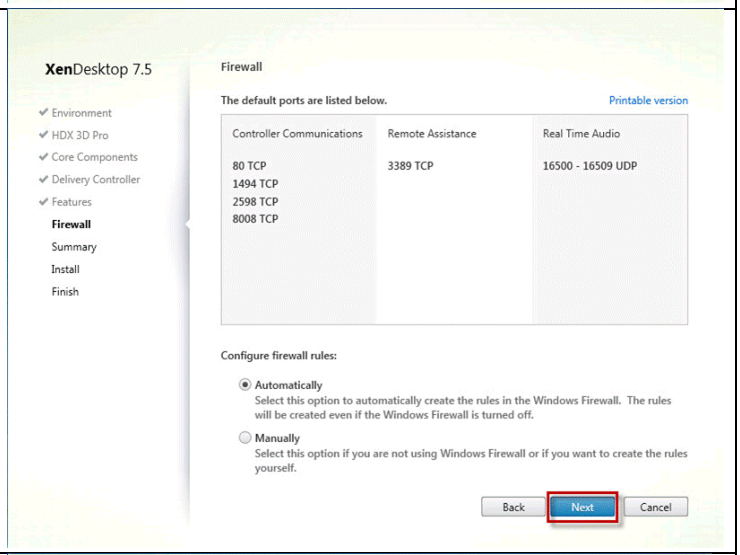
Verify the **Summary** and click **Install.**	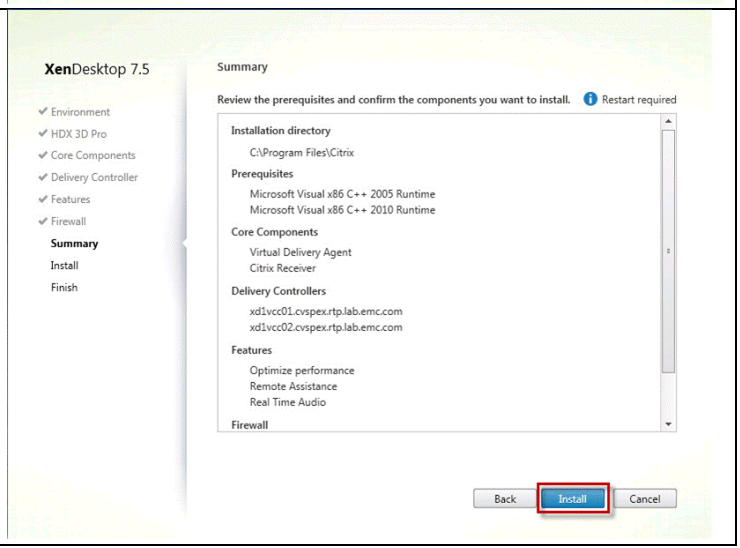

Check "**Restart Machine**". Click **Finish** and the machine will reboot automatically.	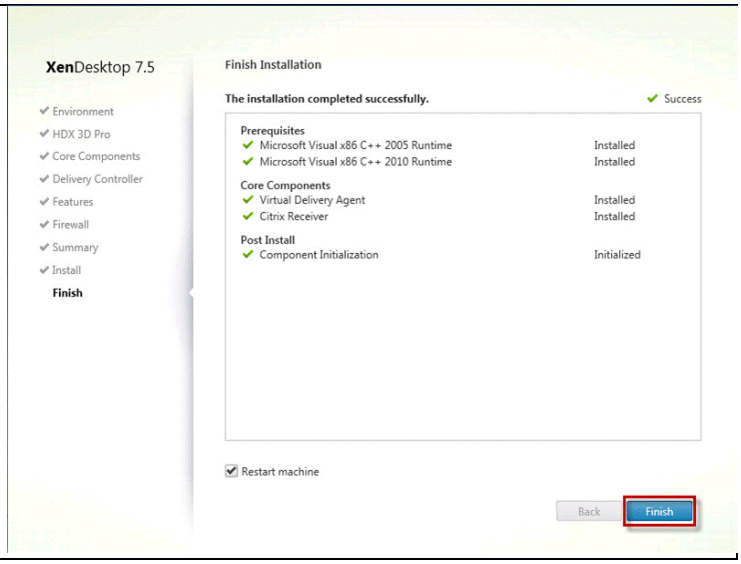

Repeat the procedure so that VDAs are installed for both HVD (using the Windows 7 OS image) and the HSD desktops (using the Windows Server 2012 image).

Citrix NetScaler Setup

This section describes the configuration of Citrix NetScaler components.

Citrix NetScaler VPX Configuration

Citrix NetScaler VPX provides the complete NetScaler all-in-one feature set in a simple, easy-to-install virtual appliance. NetScaler VPX is a web application delivery virtual appliance that accelerates internal and external web applications up to five times, optimizes application availability through advanced L4-L7 traffic management, increases security with an integrated application firewall, and substantially lowers costs by increasing web server efficiency.

There are several key configuration elements required for using NetScaler VPX as the Global Server Load Balancer (GSLB) in this DaaS architecture.

- NetScalers must be the authoritative DNS for the Access Gateway URLs.
- App Orchestration must be configured to have the tenant primary site in the first datacenter and the secondary site in a second datacenter. This is needed for each tenant requiring GSLB support.
- There must be back-end replication on the XenDesktop site for failover to succeed. This is typically implemented through storage replication of LUNS for the XenDesktop site and vDisk filesystems.

NetScaler IP Configuration

Tenant	IP Type	IP/Netmask	Purpose
Shared	NetScaler IP (NSIP)	10.71.0.20/24	NetScaler 1 Management
	NetScaler IP (NSIP)	10.71.0.21/24	NetScaler 2 Management
	Subnet IP (SNIP)	10.71.0.22/24	Source IP to Servers
	NetScaler Gateway IP	10.71.0.23/24	Client IP
Tenant 1	NetScaler IP (NSIP)	10.76.0.20/24	NetScaler 1 Management
	NetScaler IP (NSIP)	10.76.0.21/24	NetScaler 2 Management
	Subnet IP (SNIP)	10.76.0.22/24	Source IP to Servers
	NetScaler Gateway IP	10.76.0.23/24	Client IP
Tenant 2	NetScaler IP (NSIP)	10.77.0.20/24	NetScaler 1 Management
	NetScaler IP (NSIP)	10.77.0.21/24	NetScaler 2 Management
	Subnet IP (SNIP)	10.77.0.22/24	Source IP to Servers
	NetScaler Gateway IP	10.77.0.23/24	Client IP

Global Configurations

Tenant	Option	Value	Description
TCP Optimizations	Nagle's Algorithm	Enabled	Reduces number of TCP Packets
	Selective Acknowledgement (SACK)	Enabled	TCP Retransmission mechanism
	Windows Scaling	4	Increase TCP Receive Windows Size
	RNAT TCP Proxy	Enabled	Dynamic RNAT Proxy
HTTP Optimizations	Cookie Version	1	Uses UTP for time values
	Drop Invalid HTTP Requests	Enabled	Detects and drops invalid HTTP headers

Load Balancing

All services should be load balanced in the interest of increasing service availability. Load balancing also sources traffic to back-end servers from the NetScaler Subnet IP (SNIP). Load balanced services include LDAPS (TCP 636) to domain controllers, DNS (UDP/TCP 53) to DNS servers, and HTTP/HTTPS (TCP 80/443) to StoreFront servers, as shown in the table below.

Protocol	Port(s)	Target	Notes
LDAPS	TCP 636	Domain Controllers	NetScaler Authentication
DNS	TCP/UDP 53	Domain Controllers	DNS Requests
HTTPS	TCP 443	Storefront Servers	NetScaler ICA Proxy Target
Citrix XML	TCP 8080	Citrix XML Brokers	vIP to which XenApp/XenDesktop Server refer for XML Load Balancing

NetScaler Gateway

Citrix NetScaler Gateway VPX is a secure application, desktop and data access solution that provides administrators granular application- and device-level control while enabling user access from anywhere using SmartAccess and the XenMobile Micro VPN. It is a virtual appliance that runs on all industry standard hypervisors.

NetScaler Gateway offers a single point of management and tools to ensure the highest levels of information security across and outside the enterprise. At the same time, it empowers users with a single point of access-optimized for roles, devices, and networks-to the enterprise applications and

data they need. This unique combination of capabilities helps maximize the productivity of today's mobile workforce.

Authentication

NetScaler Gateway will proxy Active Directory authentication for clients.

Tenant	Setting	Detail
Shared	Name	LDAPS_authpol
	Priority	100
	Expression	Ns_true
	Profile	LDAPS_authsrv
	Profile Configuration	-serverIP 192.0.2.2 -serverPort 636 -ldapBase "DC=daas,DC=local" -ldapBindDn administrator@daas.local -ldapBindDnPassword d02d067f26957177a4 -encrypted -ldapLoginName samAccountName -searchFilter "memberOf=CN=DaasUsers,OU=Customers,DC=daas,DC=local" -groupAttrName memberOf -subAttributeName CN -secType SSL
Tenant 1	Name	LDAPS_authpol
	Priority	100
	Expression	Ns_true
	Profile	LDAPS_authsrv
	Profile Configuration	-serverIP 192.0.2.2 -serverPort 636 -ldapBase "DC=daas,DC=local" -ldapBindDn administrator@FL1.local -ldapBindDnPassword d02d067f26957177a4 -encrypted -ldapLoginName samAccountName -searchFilter "memberOf=CN=DaasUsers,OU=Customers,DC=FL1,DC=local" -groupAttrName memberOf -subAttributeName CN -secType SSL
Tenant 2	Name	LDAPS_authpol
	Priority	100
	Expression	Ns_true
	Profile	LDAPS_authsrv
	Profile Configuration	-serverIP 192.0.2.2 -serverPort 636 -ldapBase "DC=daas,DC=local" -ldapBindDn administrator@FL2.local -ldapBindDnPassword d02d067f26957177a4 -encrypted -ldapLoginName samAccountName -searchFilter "memberOf=CN=DaasUsers,OU=Customers,DC=FL2,DC=local" -groupAttrName memberOf -subAttributeName CN -secType SSL

Session

NetScaler VPX Session configuration provides different user experiences depending on the client used to access the DaaS as shown in the table below.

Tenant	Setting	Detail
Shared	Name	AG_receiver_pol
	Priority	50
	Expression	REQ.HTTP.HEADER X-Citrix-Gateway EXISTS && REQ.HTTP.HEADER User-Agent CONTAINS CitrixReceiver
	Profile	clientless_prof
	Profile Configuration	-dnsVserverName DNS_lbvsrv -transparentInterception OFF -defaultAuthorizationAction ALLOW -icaProxy ON -wihome "https://sf-lb.daas.local/Citrix/Tenant2Site" -wiPortalMode NORMAL -ntDomain daas -storefronturl "https://sf-lb.daas.local/Citrix/Tenant2Site"

Shared	Name	AG_browser_pol
	Priority	60
	Expression	REQ.HTTP.HEADER User-Agent NOTCONTAINS CitrixReceiver && REQ.HTTP.HEADER Referer EXISTS
	Profile	web_prof
	Profile Configuration	-dnsVserverName DNS_lbvsrv -transparentInterception OFF -defaultAuthorizationAction ALLOW -icaProxy ON -wihome "https://sf-lb.daas.local/Citrix/Tenant2Site" -wiPortalMode NORMAL -ntDomain daas
Tenant 1	Name	AG_receiver_pol
	Priority	50
	Expression	REQ.HTTP.HEADER X-Citrix-Gateway EXISTS && REQ.HTTP.HEADER User-Agent CONTAINS CitrixReceiver
	Profile	clientless_prof
	Profile Configuration	-dnsVserverName DNS_lbvsrv -transparentInterception OFF -defaultAuthorizationAction ALLOW -icaProxy ON -wihome "https://sf-lb.daas.local/Citrix/Tenant2Site" -wiPortalMode NORMAL -ntDomain daas -storefronturl "https://sf-lb.daas.local/Citrix/Tenant2Site"
Tenant 1	Name	AG_browser_pol
	Priority	60
	Expression	REQ.HTTP.HEADER User-Agent NOTCONTAINS CitrixReceiver && REQ.HTTP.HEADER Referer EXISTS
	Profile	web_prof
	Profile Configuration	-dnsVserverName DNS_lbvsrv -transparentInterception OFF -defaultAuthorizationAction ALLOW -icaProxy ON -wihome "https://sf-lb.daas.local/Citrix/Tenant2Site" -wiPortalMode NORMAL -ntDomain daas
Tenant 2	Name	AG_receiver_pol
	Priority	50
	Expression	REQ.HTTP.HEADER X-Citrix-Gateway EXISTS && REQ.HTTP.HEADER User-Agent CONTAINS CitrixReceiver
	Profile	clientless_prof
	Profile Configuration	-dnsVserverName DNS_lbvsrv -transparentInterception OFF -defaultAuthorizationAction ALLOW -icaProxy ON -wihome "https://sf-lb.daas.local/Citrix/Tenant2Site" -wiPortalMode NORMAL -ntDomain daas -storefronturl "https://sf-lb.daas.local/Citrix/Tenant2Site"
Tenant 2	Name	AG_browser_pol
	Priority	60
	Expression	REQ.HTTP.HEADER User-Agent NOTCONTAINS CitrixReceiver && REQ.HTTP.HEADER Referer EXISTS
	Profile	web_prof
	Profile Configuration	-dnsVserverName DNS_lbvsrv -transparentInterception OFF -defaultAuthorizationAction ALLOW -icaProxy ON -wihome "https://sf-lb.daas.local/Citrix/Tenant2Site" -wiPortalMode NORMAL -ntDomain daas

Virtual Server

Virtual server setup is summarized in the table below.

	Setting	Detail
Shared	Name	Tenant2_agvsrv
	IP/Port	10.77.0.23 TCP 443
	STA	http://10.77.0.70 http://10.0.77.0.69
Tenant 1	Name	Tenant2_agvsrv
	IP/Port	10.77.0.23 TCP 443
	STA	http://10.77.0.70 http://10.0.77.0.69
Tenant 2	Name	Tenant2_agvsrv
	IP/Port	10.77.0.23 TCP 443
	STA	http://10.77.0.70 http://10.0.77.0.69

NetScaler VPX Configuration Procedures on VMware ESXi

The following summary provides an overview of how to setup a NetScaler VPX appliance on VMware ESXi and how to configure the NetScaler appliance as described above. The procedures provided here follow steps in the online documentation (see http://edocs.citrix.com).

Downloading the NetScaler Virtual Appliance Setup Files

The NetScaler virtual appliance setup package for VMware ESX follows the Open Virtual Machine (OVF) format standard. You can download the files from MyCitrix.com. You need a My Citrix account to log on. If you do not have a My Citrix account, access the home page at http://www.mycitrix.com, click the New Users link, and follow the instructions to create a new My Citrix account.

When logged on, navigate the following path from the My Citrix home page: MyCitrix.com > Downloads > NetScaler > Virtual Appliances.

Copy the following files to a workstation on the same network as the ESX server. Copy all three files into the same folder.

- NSVPX-ESX-<release number>-<build number>-disk1.vmdk (for example, NSVPX-ESX-9.3-39.8-disk1.vmdk
- NSVPX-ESX-<release number>-<build number>.ovf (for example, NSVPX-ESX-9.3-39.8.ovf)
- NSVPX-ESX-<release number>-<build number>.mf (for example, NSVPX-ESX-9.3-39.8.mf)

For more information see
http://support.citrix.com/proddocs/topic/netscaler-vpx-10-1/ns-vpx-install-on-esx-wrapper-con.html

Labeling the physical network ports of VMware ESXi

Before installing a NetScaler virtual appliance, label of all the interfaces that you plan to assign to virtual appliances, in a unique format. Citrix recommends the following format: NS_NIC_1_1, NS_NIC_1_2, and so on. In large deployments, labeling in a unique format helps in quickly identifying the interfaces that are allocated to the NetScaler virtual appliance among other interfaces used by other virtual machines, such as Windows and Linux. Such labeling is especially important when different types of virtual machines share the same interfaces. To label the ports, complete the following steps:

1. Log on to the VMware ESXi server by using the vSphere client.
2. On the vSphere client, select the Configuration tab, and then click Networking.
3. At the top-right corner, click Add Networking.
4. In the Add Network Wizard, for Connection Type, select Virtual Machine, and then click Next.

5. Scroll through the list of vSwitch physical adapters, and choose the physical port that will map to interface 1/1 on the virtual appliances.
6. Enter NS_NIC_1_1 as the name of the vSwitch that will be associated with interface 1/1 of the virtual appliances.
7. Click Next to finish the vSwitch creation. Repeat the procedure, beginning with step 2, to add any additional interfaces to be used by your virtual appliances. Label the interfaces sequentially, in the correct format (for example, NS_NIC_1_2).

Setting Up the Initial Configuration by Using the NetScaler Virtual Appliance Console

The first task after installing a NetScaler virtual appliance on a virtualization host is to use the NetScaler virtual appliance console in the vSphere client to configure the following initial settings.

- NetScaler IP address (NSIP): This is the IP address at which you access a NetScaler or a NetScaler virtual appliance for management purposes. A physical NetScaler or virtual appliance can have only one NSIP. You must specify this IP address when you configure the virtual appliance for the first time. You cannot remove an NSIP address.
- Netmask: This is the subnet mask associated with the NSIP address.
- Default Gateway: A default gateway is needed on the virtual appliance to access it through SSH or the configuration utility from a remote administrative desktop on a different network.

To configure initial settings on the virtual appliance, use the management application on the virtual appliance Console:

1. Connect to the VMware ESXi server on which the virtual appliance is installed by using the vSphere client.
2. In the details pane on the Console tab, log on to the virtual appliance by using the administrator credentials. The default user name and password are “nsroot.”
3. At the prompts, enter the NSIP address, subnet mask, and default gateway and then save the configuration.

After you have set up an initial configuration through the NetScaler virtual appliance console in the management application, you can use either the NetScaler command line interface or the configuration utility to complete the configuration or to change the initial settings.

For more information see
http://support.citrix.com/proddocs/topic/netscaler-vpx-10-1/ns-vpx-config-basic-settings-set-initial-config-use-nsvpx-console-tsk.html

Load Balancing

The load balancing feature distributes user requests for web pages and other protected applications across multiple servers that all host (or mirror) the same content. You use load balancing primarily to manage user requests to heavily used applications, preventing poor performance and outages and ensuring that users can access protected applications. Load balancing also provides fault tolerance; when one server that hosts a protected application becomes unavailable, the feature distributes user requests to the other servers that host the same application.

Basic load balancing consists of the following 4 steps:

Step 1 Add a Server

From command line interface:

add server <name> (<ipAddress> | (<domain> [-ipv6Address (YES | NO)]) [-state (ENABLED | DISABLED)] [-comment <string>]

Example

add server DC01_srv 10.0.71.10 -state DISABLED -comment "Domain Controller"

To add a server by using the configuration utility, complete the following steps:

1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, click Add.
3. In the Create Server dialog box, specify values for the following parameters:
   - Server Name—name
   - IP Address—ipAddress (Select IP Address and type the address. Before typing an IPv6 address, select the IPv6 check box.)
   - Domain Name—domain (For a domain-name based server, select Domain Name and type the name of the server's domain.)
   - Enable after Creating—state
   - Comment—comment
4. If you specify the domain name of the server and you want the domain name to be resolved to an IPv6 address, select the IPv6 Domain check box.
5. Click Create, and then click Close. The server you named appears in the Servers pane.

Step 2 Add a Service

From the command line interface, at the command prompt, type:

add service <name> <serverName> <serviceType> <port>

To create a service by using the configuration utility, complete the following steps:

1. Navigate to Traffic Management > Load Balancing > Services.
2. In the details pane, click Add.
3. In the Create Service dialog box, specify values for the following parameters:
   - Service Name—name
   - Server—serverName
   - Protocol—serviceType
   - Port—port
4. Click Create, and then click Close. The service you created appears in the Services pane.

Step 3 Create a Virtual Server

From the command prompt:

add lb vserver <name> <serviceType> <ip> <port>

To create a virtual server by using the configuration utility, complete the following steps:

1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, click Add.
3. In the Create Virtual Server (Load Balancing) dialog box, specify values for the following parameters:
   - Name—name
   - IP Address—IPAddress

- Protocol—serviceType
- Port—port

4. Click Create, and then click Close. The virtual server you created appears in the Load Balancing Virtual Servers pane.

Step 4 Bind Services to Virtual Server

From the command prompt:

bind lb vserver <name> <serviceName>

From the configuration utility, complete the following steps:

1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, select the virtual server for which you want to bind the service.
3. Click Open.
4. In the Configure Virtual Server (Load Balancing) dialog box, on the Services tab, select the Active check box next to the service that you want to bind to the virtual server.
5. Click OK.

For more information see
http://support.citrix.com/proddocs/topic/netscaler-vpx-10-1/ns-lb-wrapper.html

NetScaler Gateway LDAPS Authentication

1. In the configuration utility, on the Configuration tab, expand NetScaler Gateway > Policies > Authentication/Authorization > Authentication.
2. Click LDAP.
3. In the details pane, on the Policies tab, click Add.
4. In Name, type a name for the policy.
5. Next to Server, click New.
6. In Name, type the name of the server.
7. Under Server, in IP Address and Port, type the IP address and port number of the LDAP server.
8. In Type, select either AD for Active Directory or NDS for Novell Directory Services.
9. Under Connection Settings, complete the following:
   a. In Base DN (location of users), type the base DN under which users are located.
   b. The base DN is usually derived from the Bind DN by removing the user name and specifying the group where users are located.
   c. In Administrator Bind DN, type the administrator bind DN for queries to the LDAP directory.
10. In Administrator Password and Confirm Administrator Password, type the administrator password for the LDAP server.
11. In Security Type, select the security type and then click Create.
12. To allow users to change their LDAP password, select Allow Password Change.

For more information see
http://support.citrix.com/proddocs/topic/netscaler-gateway-101/ng-ldap-authen-config-overview-con.html

NetScaler Gateway Session

To create a session profile for Receiver, WorxHome, or StoreFront, complete the following steps:

1. In the configuration utility, on the Configuration tab, in the navigation pane, expand NetScaler Gateway > Policies and then click Session.
2. In the details pane, on the Profiles tab, click Add.
3. In Name, type a name for the profile.
4. Click the Security tab and in Default Authorization Action, click Override Global, and then select ALLOW.
5. Click the Client Experience tab and then do the following:
   a. Next to Plug-in Type, click Override Global and then select Java.
   b. Next to Single Sign-on to Web Applications, click Override Global and then select the check box Single Sign-on to Web Applications.
   c. Next to Clientless Access, click Override Global and then select Off.
6. Click the Published Applications tab and then configure the following settings:
   a. Next to ICA Proxy, click Override Global, and then select ON.
   b. Next to Single Sign-on Domain, click Override Global, enter the domain name and then click Create. For example, enter mydomain.
   c. In Web Interface Address, click Override Global, and then type the web address for StoreFront. For example, enter https://storefront.t.com/Citrix/StoreWeb.

Note When you configure the StoreFront URL in NetScaler Gateway, such as https://<SFLite-FQDN>/Citrix/StoreWeb, the text StoreWeb is case sensitive.

7. Click Create.

To create the session profile for Receiver for Web, complete the following steps:

1. In the configuration utility, on the Configuration tab, in the navigation pane, expand NetScaler Gateway > Policies and then click Session.
2. In the details pane, on the Profiles tab, click Add.
3. In Name, type a name for the profile.
4. Click the Client Experience tab and then do the following:
   a. In Clientless Access, click Override Global and then select Allow.
   b. In Single Sign-on to Web Applications, click Override Global and then select the check box.
5. On the Published Applications tab, do the following:
   a. Next to ICA Proxy, click Override Global, and then select ON.
   b. Next to Web Interface Address, click Override Global and then enter the web address (URL) for StoreFront.

Note The StoreFront URL is case sensitive, such as https://<StoreFrontFQDN>/Citrix/<StoreWebName>/.

   c. In Single Sign-on Domain, type the domain name.
6. Click Create.

For more information see
http://support.citrix.com/proddocs/topic/netscaler-gateway-101/ng-connect-users-wrapper-con.html

Virtual Server

1. In the configuration utility, on the Configuration tab, in the navigation pane, expand NetScaler Gateway and click on Virtual Servers
2. In the details pane, click on Add
3. In Name type a name for the virtual server
4. In IP Address type the IP of the virtual server
5. In SSL, add the appropriate SSL certificate
6. Click on the Authentication tab
7. Click on Insert Policy to bind the authentication policy to the virtual server
8. Click on the Policies tab
9. Click on Insert Policy to bind the session policies to the virtual server
10. Click on the Published Applications tab
11. Click on the Add link in the Secure Ticket Authority section to add the STA servers URL (http://x.x.x.x or http://server.domain.local)

For more information see
http://support.citrix.com/proddocs/topic/netscaler-gateway-101/ng-install-wrapper-con.html

NetScaler Troubleshooting

The table below shows NetScaler Log Files, which can be useful in troubleshooting.

Log File	Purpose
/tmp/aaa.debug	Streaming file viewable with cat - Displays realtime NetScaler Gateway authentication events
/var/log/auth.log	logs ssh/sftp administrative logins to NetScaler
/var/log/cron	Logs cron job results
/var/log/httpaccess.log	Logs http/https administrative logins to NetScaler
/var/log/httperror.log	Logs NetScaler website errors
/var/log/license.log	Logs NetScaler licensing status
/var/log/messages	Logs history of commands in NetScaler shell
/var/log/ns.log	General NetScaler log
/var/log/nsvpn.log	Logs NetScaler Gateway access
/var/log/snmpd.log	Logs SNMP status
/var/nslog/aslearn.log	NetScaler Learning Server logs
/var/nslog/dmesg.boot	dmesg log file
/var/log/newnslog	Logfile readable with nsconmsg

The NetScaler stat command lists real time relevant statistics of most NetScaler objects. A few examples of the stat command are shown below.

Entity	Command
System Statistics	stat ns stat system stat cpu stat interface

NetScaler Object Statistics	stat serviceGroup stat lb vserver stat service
Protocol Statistics	stat ssl stat dns stat http

The NetScaler *nsconmsg* command lists useful metrics.

Command	Purpose
nsconmsg	Displays current real time NetScaler statistics
nsconmsg --help	Displays different options for nsconmsgconmsg
nsconmsg –d current –g cpu_use	Displays live CPU statistics
nscongmsg –K [newnslog-filename] –d event	Disaplys significant archived events
nsconmsg –d current –g pol_hits	Displays real time policy hits
nsconmsg –s ConSSL=2 –d oldconmsg	Displays real time SSL statistics

Tools for network tracing include the *nstcpdump.sh* script and the *nstrace.sh* script. The nstcpdump.sh script outputs network traffic details directly to stdout. It is not as detailed as nstrace.sh and more CPU-intensive, but useful for quickly viewing network traffic. The *nstrace.sh* script outputs information to a .cap or .pcap file, which is viewable with Wireshark and other tools.

Installing Citrix App Orchestration and CloudPortal Services Manager

The procedures in this section describe how to install the Citrix CloudPortal Services Manager and App Orchestration software components. It is assumed that the software prerequisites have been met in advance. The installation documentation includes checklists that should first be completed prior to software component installation.

Preparing and installing CloudPortal Services Manager software

Review the pre-installation checklist in the documentation:

- CloudPortal Services Manager 11.0 Documentation
- Known Issues for CloudPortal Services Manager

It is assumed that all prerequisites have been met prior to beginning software installation.

Instructions	Visual
A critical prerequisite is to establish DNS aliases for servers that run CloudPortal Services Manager roles: CORTEXSQL, CORTEXPROVISIONING, CORTEXWEB, and CORTEXREPORTS. If this prerequisite has not already been completed, set up the CPSM DNS aliases using `dnscommand` or the DNS Manager (see http://technet.microsoft.com/en-us/library/cc772053.aspx for instructions).	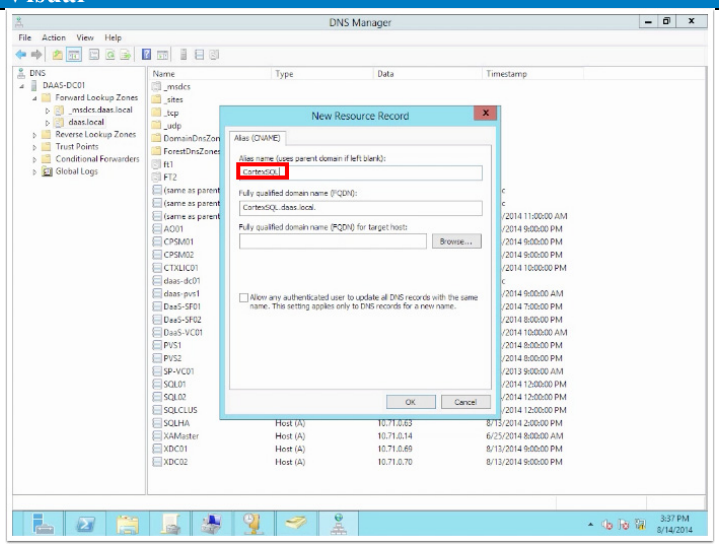
Remote Desktop Connect (RDC) to the first CloudPortal Services Manager server (CPSM01) and run the Roles & Features wizard in the Windows Server Manager to configure the required server roles. Configure the required prerequisites (e.g., .NET framework, MSMQ features, etc.). The CloudPortal Services Manager Setup Tool attempts to handle many unresolved prerequisites at installation (such as installing the .NET framework) if needed. Repeat as needed to prepare and verify the second CloudPortal Services Manager server (CPSM02).	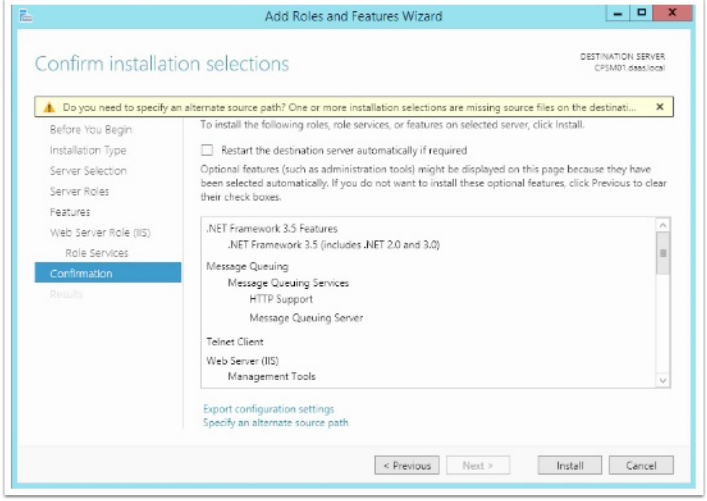

Using the Microsoft SQL Server Management Studio, check to make sure both Windows and SQL authentications are enabled. Verify that the installation domain account has sysadmin rights on the CloudPortal Services Manager SQL server instance, which is on the CORTEXSQL server (SQL01). Check the Data Warehouse/Reporting server (SQL02) to make sure that the SQL reporting service is configured, Reporting services URLs can be resolved, and the installation account has sufficient rights on the Data Warehouse SQL server and reporting services. Test the SQL connection to the CORTEXSQL server from the CORTEXPROVISIONING server and web server (CPSM01 and CPSM02).	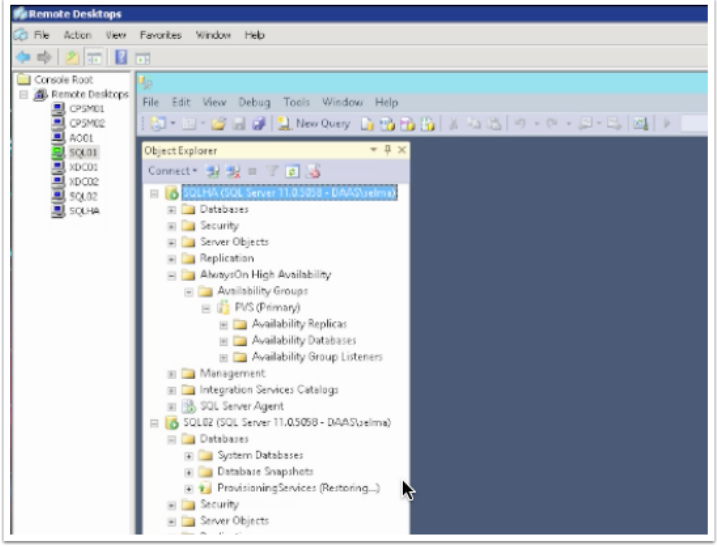
RDC into the first CloudPortal Services Manager server (CPSM01) and access the software distribution media. Launch **Setup.exe** to run the CloudPortal Services Manager control panel, and then click **Get Started**.	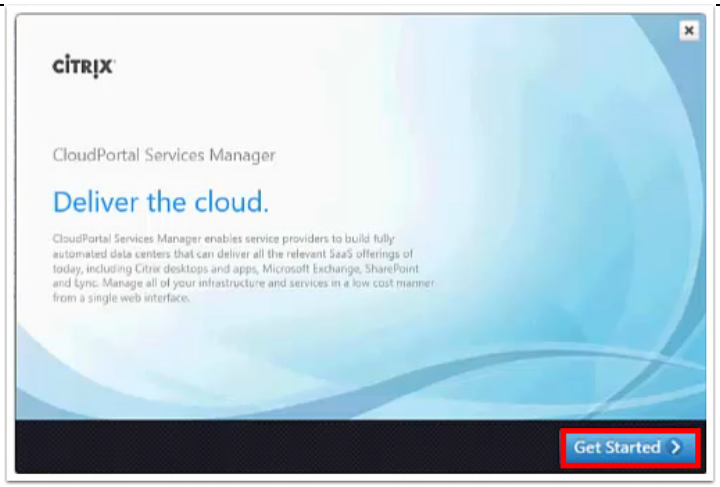
On the **Select Deployment Task** page, select **Install CloudPortal Services Manager**.	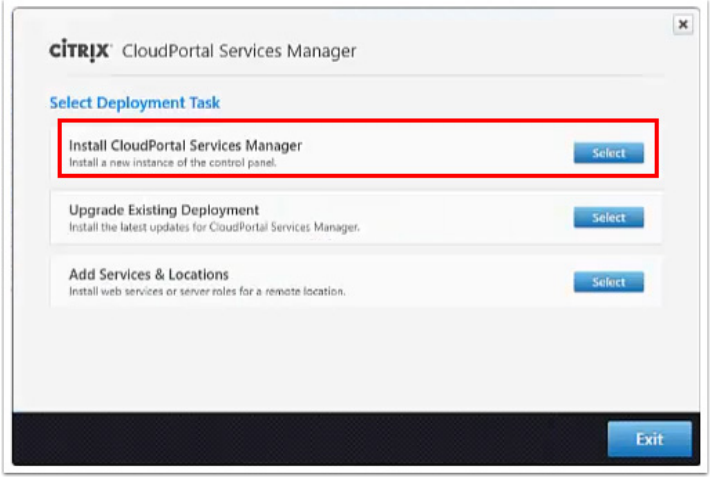

On the **Prepare Environment** page, click **Verify** for the **Extend Active Directory Scheme** task. Since the aliases have been defined already, the Setup tool indicated the completion of that step.	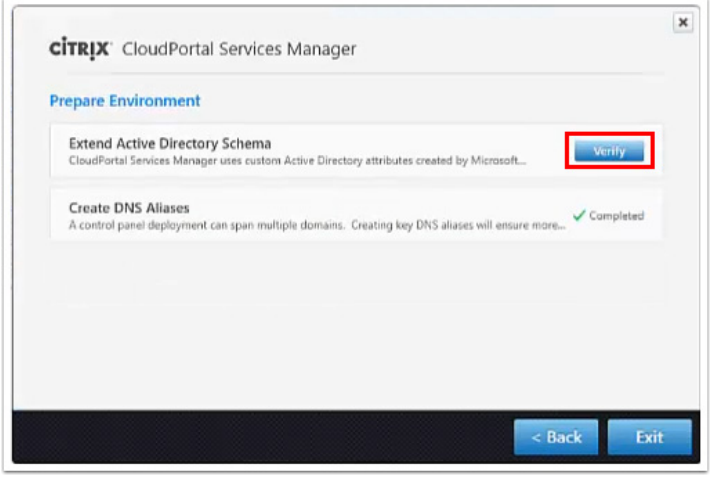

Preparing and installing Citrix App Orchestration software

Review the pre-installation checklist in the documentation:

- Getting Started with Citrix App Orchestration 2.5
- Known Issues for App Orchestration 2.5

Deploying App Orchestration typically occurs using the following phased approach:

Phase	Tasks
Prepare	• Download the software for App Orchestration and its components. • Prepare your environment and the machines you will use to deploy App Orchestration and design and deliver offerings.
Install	• Use the App Orchestration Install Center to install the required software on the configuration server, Delivery Controllers, Session Machines, and StoreFront servers.
Configure	• Configure App Orchestration's global settings.
Define	• Define additional domains. • Create additional datacenters. • Set up and configure compute resources. • Add instance configurations.
Design	• Create Delivery Sites. • Create a Session Machine Catalog for on-demand provisioning or external provisioning. • Create a StoreFront Server Group. • Create an offering.
Deliver	• Add a tenant and add users. • Adjust capacity. • Subscribe the tenant to the offering. • Enable tenant self-service provisioning with CloudPortal Services Manager.

Create Active Directory domains

App Orchestration manages multiple Active Directory domains: e.g., a shared resource domain and a default user domain.

- **Shared resource domain:** The domain where the App Orchestration configuration server resides. This domain contains components that are shared with multiple tenants. This is also where the App Orchestration root OU is created.
- **Default user domain:** The domain where App Orchestration user accounts reside. This domain also includes the tenant users and groups that will access offerings delivered from the shared resource domain.

Create Root OU and Admin Group

Delivery Groups in Citrix App Orchestration and CloudPortal Services Manager correspond to Active Directory Organizational Units (OUs). The implementation of App Orchestration relies on the proper configuration of the root App Orchestration OU.

Within the root App Orchestration OU, an AD group identifies the set of privileged CSP administrators that are permitted to perform management tasks across the CSP domain. Early implementation tasks for App Orchestration include creating the root OU, adding the administrative users to this OU, and adding them to the Domain administrative group.

Instructions	Visual
In the shared resource domain, create a new Active Directory OU that will be the root OU for the deployment. In this CVD, assign the name **CitrixAppOrchestration**.	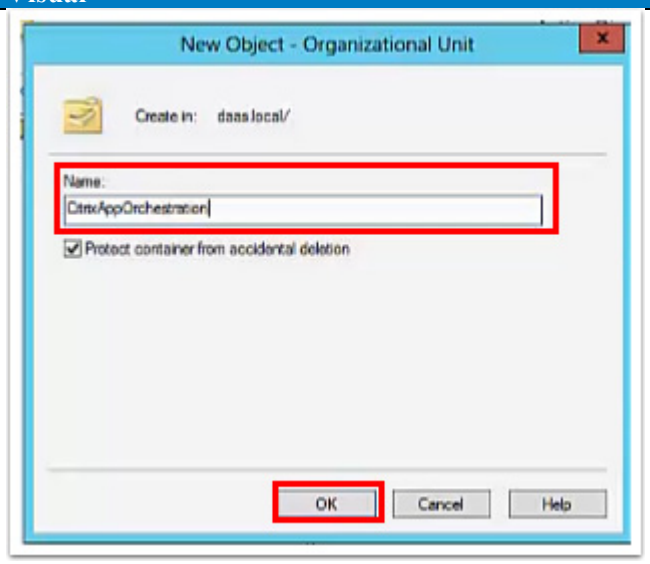

In the CitrixAppOrchestration OU, create a new Active Directory group. Assign the name **AOAdmins**.	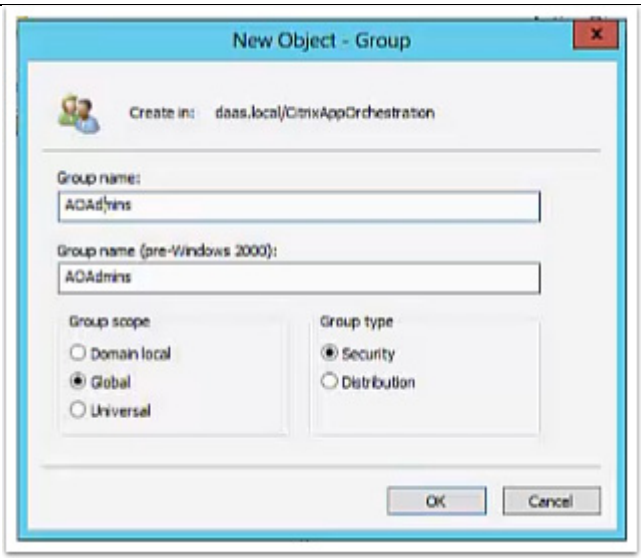
In the CitrixAppOrchestration OU, create a new user. Assign the logon name **AOAdmin**. Click **Next**.	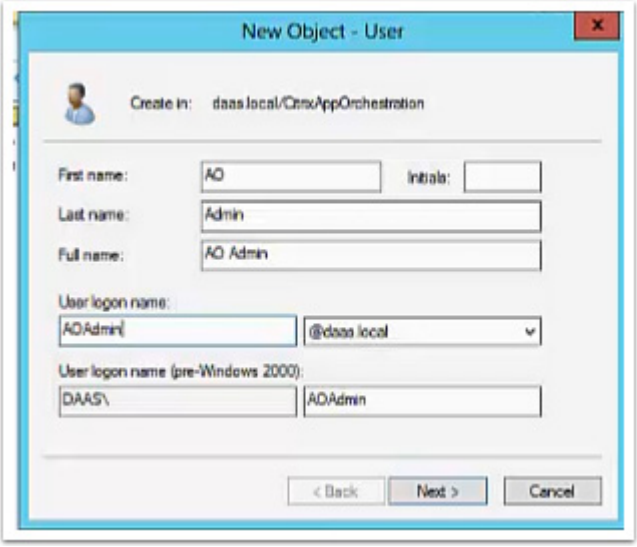
Add a password for this user, and disable all checkboxes on the **New Object – User** page. Click **Next**.	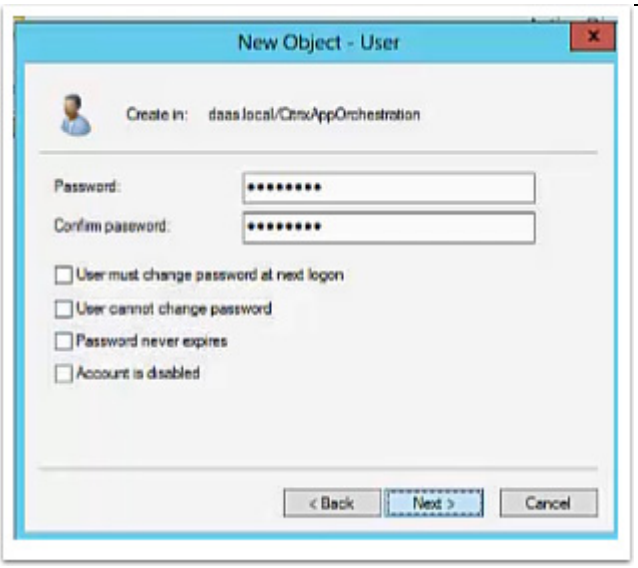

Confirm user information, and click **Finish**.	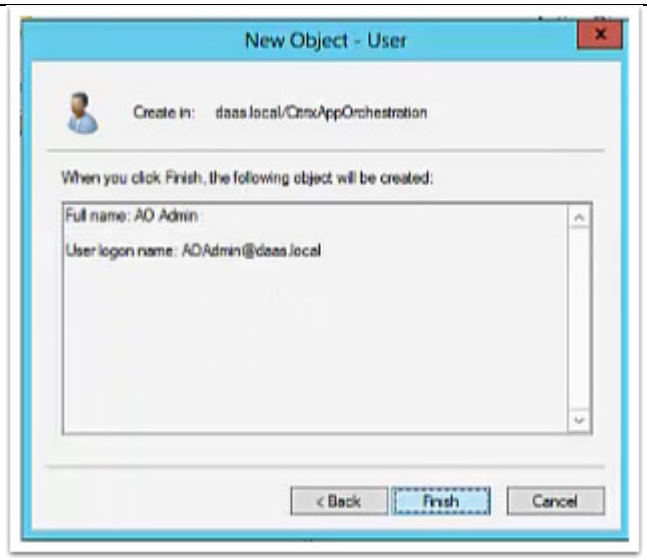
Next, add the new user to the group: Edit the properties for the new user (AOAdmin). Click **Member Of** tab. Then, click **Add**.	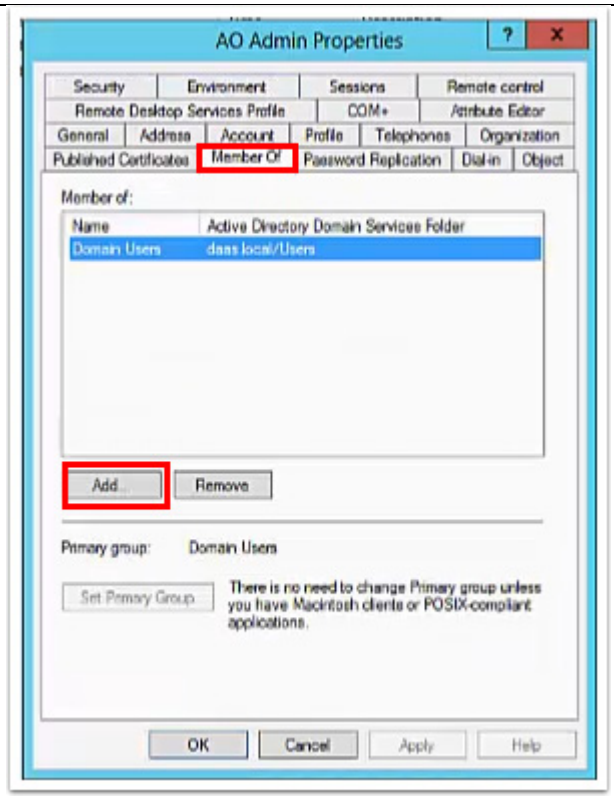

Enter the name of the group (**aoadmins**). Click **OK**.	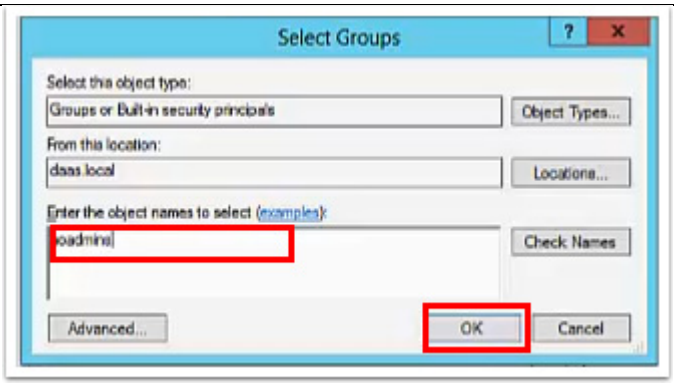
Next, edit the properties for the CitrixAppOrchestration OU. Select the **Security** tab. Click **Add.**	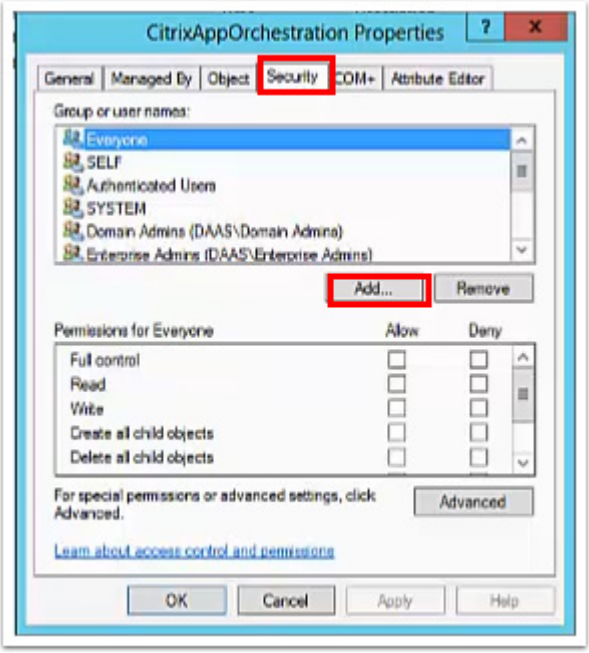
Enter the name of the new group (**aoadmins**). Click **OK**.	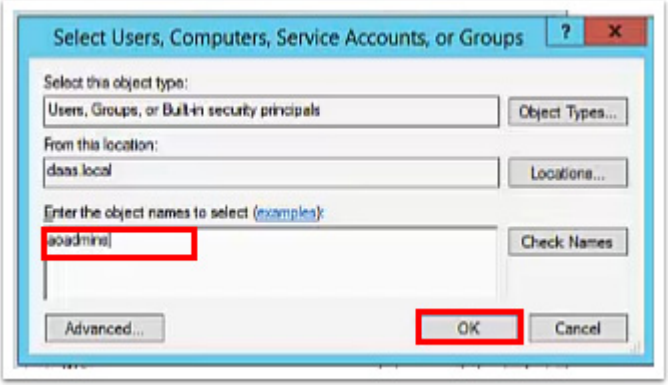

Edit the permissions for the AOAdmins group, and enable all **Allow** checkboxes**.**	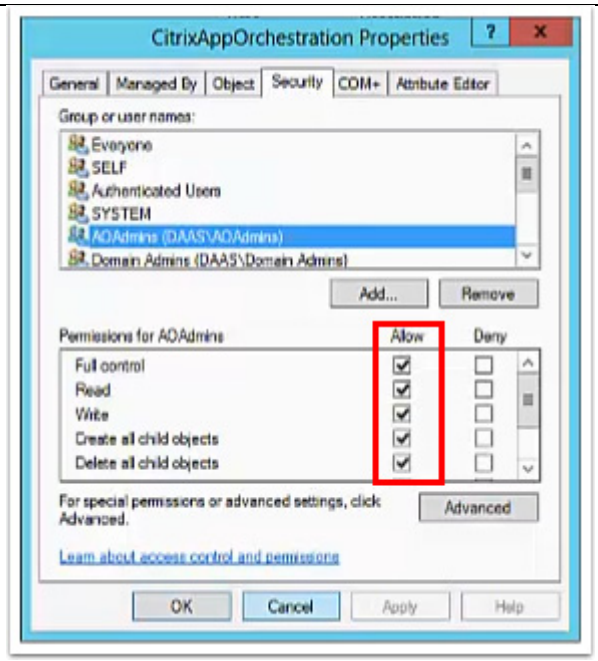

Install AO Software

The Citrix App Orchestration Install Center is used to install App Orchestration software and prepare machines for deployment.

Instructions	Visual
Launch the Citrix App Orchestration Install Center, and click on **App Orchestration Configuration Server.** When prompted on the welcome page, click **Start**. Then accept the end user license agreement and click **Install** to begin the installation.	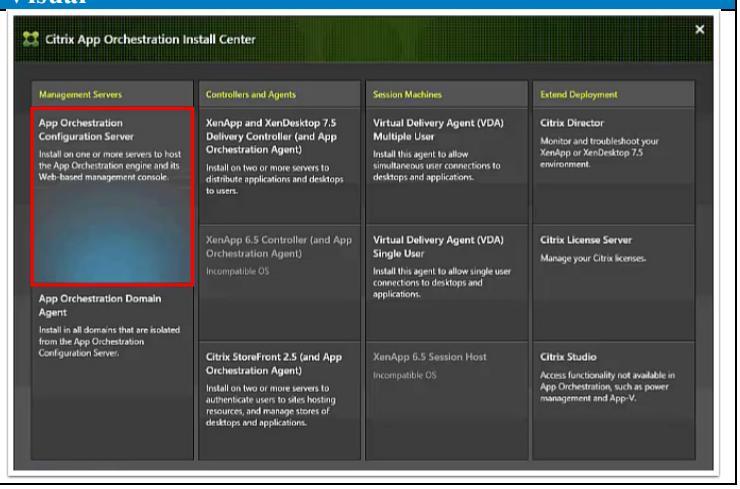

Installation progress is displayed.	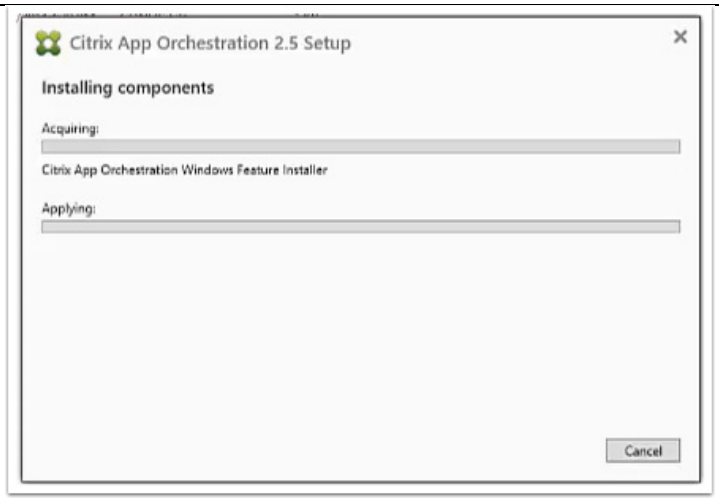
The following status message indicates the installation was successful.	

App Orchestration 2.5 High Availability Configuration

App Orchestration 2.5 provides failover capabilities from the application level. At the same time, Microsoft's SQL Server 2012/2014 Enterprise provides a new High Availability (HA) solution known as "AlwaysOn". With availability groups at the database level, SQL Server now allows you to create a group of databases which failover together as a unit from one replica/instance to another.

To support AlwaysOn availability groups, no special configuration is needed for App Orchestration. SQL Server provides the capability from a transaction level. This section will provide the step-by-step guide in SQL Server to enable this feature within App Orchestration 2.5 databases.

Process Overview

Citrix recommends the following installation sequence and key configuration steps to deploy an App Orchestration 2.5 farm that uses an AlwaysOn availability group

1. Select or create a Windows Server failover cluster.
2. Install SQL Server 2012 or 2014 on each cluster node.

3. Create and configure an availability group.
4. Install and configure App Orchestration multi-datacenter.
5. Add the App Orchestration databases to availability group.
6. Test failover for availability group.

Before You Begin

A SQL Server AlwaysOn Availability Group is not just a combination of database mirroring and database clustering. It is a completely new high availability and disaster recovery feature that co-exists with existing high availability and disaster recover options such as mirroring and log shipping. Before you begin deployment, review the following information about SQL Server AlwaysOn, the technologies that support AlwaysOn, and App Orchestration:

- Knowledge and skill requirements
- AlwaysOn Availability Group concepts
- Hardware and software requirements
- Permissions

Knowledge and Skill Requirements

To implement SQL Server AlwaysOn Availability Groups as a high availability and disaster recovery solution, several technologies interact and have to be installed and configured correctly. We recommend that the team responsible for setting up an AlwaysOn environment for App Orchestration products has a working knowledge of, and hands-on skills with the following technologies:

- Windows Server Failover Clustering (WSFC) services
- SQL Server 2012
- App Orchestration multi-datacenter deployment

High Availability Concepts for App Orchestration at Various Levels

- Multi-Controller for a delivery site – several controllers forms one delivery site, it can provide HA capability for the connection request from end users
- Multi-Sites for the subscription - by creating multiple delivery sites, the load can be distributed to different sites internally while creating the subscription
- Multi-Configuration Servers - you can configure multi-configuration servers to provide the high availability of configuration server, each server can talk to AO Agent
- Multi-Datacenter – To span different physical location, you can setup multiple datacenters in case one of the datacenter fail
- Multi-Databases – it leverage SQL servers HA capability where App Orchestration database resides, in this document, we only discuss “AlwaysOn” availability group

SQL Server AlwaysOn Availability Group concepts

A SQL Server Availability Group enables you to specify a set of databases that you want to fail over together as a single entity. When an availability group fails over to a target instance or target server, all the databases in the group fail over also. Because SQL Server 2012 can host multiple availability groups on a single server, you can configure AlwaysOn to fail over to SQL Server instances on different servers. This reduces the need to have idle high performance standby servers to handle the full load of the primary server, which is one of the many benefits of using availability groups.

An availability group consists of the following components:

- Replicas, which are a discrete set of user databases called availability databases that fail over together as a single unit. Every availability group supports one primary replica and up to four secondary replicas.
- A specific instance of SQL Server to host each replica and to maintain a local copy of each database that belongs to the availability group.

For the details about the benefit of AlwaysOn Availability Groups and overview of AlwaysOn Availability Groups terminology, see AlwaysOn Availability Groups(SQL Server).

Windows Server Failover Clustering

To create and use SQL Server 2012 AlwaysOn Availability Groups, you have to install SQL Server 2012 on a Windows Server Failover Clustering (WSFC) cluster. For more information, see Windows Server Failover Clustering (WSFC) with SQL Server.

Although configuring a WSFC cluster is out of the scope for this article, you should be aware of the following requirements before you install and configure a cluster:

- All the cluster nodes must be in the same Active Directory Domain Services (AD DS) domain.
- Each availability replica in an availability group must reside on a different node of the same Windows Server Failover Clustering (WSFC) cluster.
- The cluster creator must have the following accounts and permissions:
  - Have a domain account in the domain where the cluster will exist.
  - Have local administrator permissions on each cluster node.
  - HaveCreate Computer objects and Read All Properties permissions in AD DS. For more information, see Failover Cluster Step-by-Step Guide: Configuring Accounts in Active Directory.

A very important aspect of configuring failover clustering and AlwaysOn is determining the quorum votes that are needed for the cluster nodes.

Failover clustering is based on a voting algorithm where more than one half of the voters, or quorum, must be online and able to communicate with each other. Because a given cluster has a specific number of nodes and a specific quorum configuration, the cluster service is able to determine what constitutes a quorum. The cluster service will stop on all the nodes if the number of voters drops below the required majority.

For more information, seeWSFC Quorum Modes and Voting Configuration (SQL Server) and Configure Cluster Quorum NodeWeight Settings.

Configure an AlwaysOn Group for App Orchestration

You should have 2 SQL Server instances to form one group and one availability group listener should be created for the database connection

Prepare the Windows Server Cluster environment

Obtain access to or create a three node Windows Server Failover Clustering (WSFC) cluster that you can use to install SQL Server 2012 on each cluster node. The following reference material provides guidance and detailed steps to configure a Windows Server failover cluster:

- Failover Clusters in Windows Server 2008 R2

  This page provides links to Getting Started, Deployment, Operations, and Troubleshooting articles for Windows Server 2008 R2.

- Failover Clustering Overview

This page provides links to Getting Started, Deployment, Operations, and Troubleshooting articles for Windows Server 2012.

Prepare the SQL Server Environment

Before you can create an Availability Group for App Orchestration, you must prepare the SQL Server 2012 environment. To prepare the environment, complete the following tasks:

- Install the SQL Server prerequisites.
- Install SQL Server.
- Enable AlwaysOn.

Install SQL Server 2012

To install SQL Server 2012, complete the following steps:

1. Install SQL Server 2012 prerequisites on each cluster node.

   For more information, see Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups (SQL Server).

2. Install SQL Server on each cluster node.

   Choose the 2nd option “New SQL Server failover cluster installation”, For more information, see Installation for SQL Server 2012.

Enable AlwaysOn

Note You must enable AlwaysOn for each database server in the cluster.

To enable AlwaysOn, complete the following steps:

1. Your logon account must have the permission levels to create an availability group. The account must have membership in the db_owner fixed database role and either CREATE AVAILABILITY GROUP server permission, CONTROL AVAILABILITY GROUP permission, ALTER ANY AVAILABILITY GROUP permission, or CONTROL SERVER permission, it recommended to change the logon account to the same domain user account for each instance
2. Log on to the server that will host the primary replica and start SQL Server Configuration Manager.
3. In Object Explorer, select SQL Server Services, right-click SQL Server (<instance name>), where <instance name> is the name of a local server instance for which you want to enable AlwaysOn Availability Groups, and then click Properties.
4. Select the AlwaysOn High Availability tab.
5. Select the Enable AlwaysOn Availability Groups check box, and then click OK.
6. Although the change is saved you must manually restart the SQL Server service (MSSQLSERVER) to commit the change. The manual restart enables you to choose a restart time that is best for your business requirements.
7. Repeat the previous steps to enable AlwaysOn for SQL Server on the other cluster nodes.

For more information, see Enable and Disable AlwaysOn Availability Groups (SQL Server) (http://go.microsoft.com/fwlink/p/?LinkId=267140).

Create and Configure the Availability Group

Use the following procedure to create an availability group on the primary replica, which is AO-SRV1 in our example.

If there is no user databases are on the instance of connected server, which is true in our case, since we have not created App Orchestration database yet in this step, you need to create empty AO database before creating the availability group

To create the empty AO databases, complete the following steps:

1. Make sure that your logon account has the correct permissions for this task. You require one of the following permissions in the master database to create the new database:
   - CREATE DATABASE
   - CREATE ANY DATABASE
   - ALTER ANY DATABASE
2. Log on to the server that will host the primary replica, which is AO-SRV1 in our example.
3. Start Management Studio.
4. In Object Explorer, right-click
5. Databases and then clickNew Database.
6. In the New Database dialog box, type the Database name: which is "AppOrchestration" for this example, Options->Collation, in the drop box, select “Latin1_General_100_CI_AS_KS”, Click OK.
7. Because the New Availability Group Wizard will not create an availability group unless the user database was backed up, you have to back up the database and set the Recovery model to “Full”,From Options->Recovery model.
8. In Object Explorer expand Databases and right-click the database that you just created. Pick Tasks and then choose Back Up.

9. In the Back Up Database dialog box, click OK to accept all the default settings and create the back up.
10. Repeat above steps to create another database named “AppOrchestrationLogging”

Create the Availability Group

1. Make sure that your logon account has the required permissions to create an availability group. This requires membership in the db_owner fixed database role and either CREATE AVAILABILITY GROUP server permission, CONTROL AVAILABILITY GROUP permission, ALTER ANY AVAILABILITY GROUP permission, or CONTROL SERVER permission.
2. Log on to the server that will host the primary replica and start SQL Server Management Studio.
3. To start the New Availability Group Wizard, right-click AlwaysOn High Availability and then click New Availability Group Wizard.
4. Click Next to advance to the Specify Name page. Enter AO-AG1 as the name of the new availability group in the Availability group name: box.

   This name must be a valid SQL Server identifier, unique on the Windows Server Failover Clustering cluster and unique on the domain.
5. On the Select Databases page, all user databases that are eligible to become the primary database for the new availability group are listed on the User databases on this instance of SQL Server grid. Select the database which is “AppOrchestration” and “AppOrchestrationLogging”, and then click Next.
6. On theSpecify Replicas page, use the following tabs to configure the replicas for AO-AG1: Replicas,Endpoints, and Backup Preferences.
7. An availability group listener is a virtual network name that provides client connectivity to the database a given availability group. Availability group listeners direct incoming connections to the primary replica or to a read-only secondary replica. The listener provides fast application failover after an availability group fails over. For more information, see Availability Group Listeners, Client Connectivity, and Application Failover (SQL Server).
8. On the Listener tab, configure an availability group listener for our example, use the name AOListener, if your deployment span multi-subnet, you can add multiple IP address to the AOListener’s DNS records to provide fast connection
9. Select the desired configuration for each instance in the Selected instances grid, and then click Next.
10. Click Finish to create the availability group.
11. The Select Initial Data Synchronization page lets you select a synchronization preference and specify the shared network location that all replicas can access. For our environment accept the default, Full, which performs full database and log backups. ClickNext.
12. The Validation page of the wizard displays the results of six checks before it lets you continue with availability group creation. If all checks pass, click Next to continue. If any tests fail, you cannot continue until you correct the error and then click Re-run Validation to run the validation tests again. When all the tests pass, click Next to continue.
13. On theSummary page, verify the configuration of the replica that you are adding and then click Finish to save it. To change the configuration, click Previous to return to previous wizard pages.

Install App Orchestration and Configure Multi-datacenter

We recommend to read the Getting Started Guide and deploy a Multi-Datacenter Environment in App Orchestration 2.5, use the group listener DNS name as all database servers connection during the configuration, in this case, it's "AOListener" which is already created in the previous steps

1. After the installation of configuration server, run "Citrix App Orchestration Server Configuration"->Create a new deployment in the Database name field, input "AppOrchestration" which is the empty database created in the previous step ,in the Database server field, input availability group listener DNS Name
2. Install and join the 2nd configuration server to the existing deployment right after the new deployment, launch "Citrix App Orchestration Server Configuration"->Join an existing deployment, input the first configuration server's address and finish the wizard

Note If you have already configured the Delivery Sites, Storefront server group in App Orchestration before joining the 2nd configuration server, you should append the second configuration server address on each App Orchestration Agent machines

1. Open Registry editor, modify "ConfigurationServiceAddress" under "HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\CloudAppManagement\Agent", append 2nd configuration server information to it, like https://FQDN/cam/api
2. Restart the Agent services

Login Replication

App Orchestration logins are not replicated in the availability group, This occurs because login information is stored in the master database, the workaround is to manually copy the App Orchestration database's logins from the primary replica to the secondary replicas

1. Login to primary SQL instance by Management Studio
2. Go to each App Orchestration databases->Security->Users, you will find the machine account information like "domain name\host name$", record all App Orchestration servers information
3. Execute the SQL query "CREATE LOGIN [domain name\host name$] FROM WINDOWS WITH DEFAULT_DATABASE = [master]" on each secondary SQL instance's SQL query's window, replace "domain name\host name$" by the machine account information from previous step
4. Repeat these steps if you add any new database to the availability group later, keep the login account synced manually

Use Failover Test to Validate the AlwaysOn Installation

After you synchronize the App Orchestration data with the secondary replicas, the final step is to test failover.

You must run extensive failover tests to make sure that the behavior of the AlwaysOn environment is as expected and that you completely understand the configuration requirements and procedures related to SQL Server 2012 Availability Groups.

Test availability group failover by using either the planned manual failover described in Perform a Planned Manual Failover of an Availability Group (SQL Server) or the forced manual failover described in Perform a Forced Manual Failover of an Availability Group (SQL Server).

You can perform either of the previous failovers by using the Failover Wizard in SQL Server Management Studio, Transact-SQL, or Windows PowerShell in SQL Server 2012.

Then test the database connectivity and App Orchestration functionality, these tests include and are not limited to the following:

- Verify that published App/Desktop from different datacenter are completely functional.
- Can login to configuration server web console, verify that App Orchestration data is preserved and not corrupted.

Migrate Existing Database Deployment to AlwaysOn Availability Group

If you have already deployed the App Orchestration databases, connecting to the SQL instance directly and want to add HA capability by enabling “AlwaysOn” availability group, you need to:

- Backup the databases (2 databases need to backed up, AppOrchestration and AppOrchestrationLogging)
- Restore the databases to the primary replica
- Add the App Orchestration databases to the availability group
- Modify the databases connection string pointing the connection to group listener

Restore the Databases on the Primary Replica

After restore, you need create the login account manually on the primary replica and each secondary replicas in the group, detail steps, refer to Login Replication section

To modify the databases connection string, complete the following steps:

1. On the configuration servers, open the registry editor, find the key “DatabaseConnectionString”
2. Under “HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\CloudAppManagement\Configuration”, replace current SQL server FQDN with group listener’s VNN, other connection string should be modified also in the follows:

   HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\XDservices\ADIdentitySchema\DataStore\Connections\ConnectionString

   HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\XDservices\ConfigLoggingSiteSchema\DataStore\Connections\ConnectionString

   HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\XDservices\ConfigurationSchema\DataStore\Connections\ConnectionString

   HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\XDservices\DAS\DataStore\Connections\ConnectionString

   HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\XDservices\DesktopUpdateManagerSchema\DataStore\Connections\ConnectionString

   HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\XDservices\HostingUnitServiceSchema\DataStore\Connections\ConnectionString

   HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\XDservices\TrustServiceSchema\DataStore\Connections\ConnectionString

CPSM High Availability Configuration

CloudPortal Services Manager consists of multiple components:

Core Components:

- Databases
- Provisioning Engine
- Directory Web Service
- Web Portal & API
- Reporting

Optional Components:

- Service-specific web services or tools (e.g. XenDesktop Web Service)

This document provides a basic guidance for deploying these components to support highly availability based on Windows Server 2012, 2012R2, and SQL Server 2012.

Databases

CloudPortal Services Manager uses DNS Alias (CNAME) that points to the SQL server instance name in connection strings, this simplifies the potential future database move from one SQL server cluster/group to another.

Prerequisite

SQL Server 2012 AlwaysOn Database Availability Group is configured correctly and operational following the Microsoft guidance.

Installation

For CPSM to support the SQL server AlwaysOn in a single subnet, first make the CORTEXSQL DNS alias point to the Listener name as part of the preparation for the initial deployment.

Follow the normal process to start the CloudPortal Services Manager system database installation: http://support.citrix.com/proddocs/topic/ccps-115/ccps-install-database-create.html. If the listener configured for the AlwaysOn Availability group is on a port other than 1433, the correct port number must be specified during the installation step of “Create System Databases”:

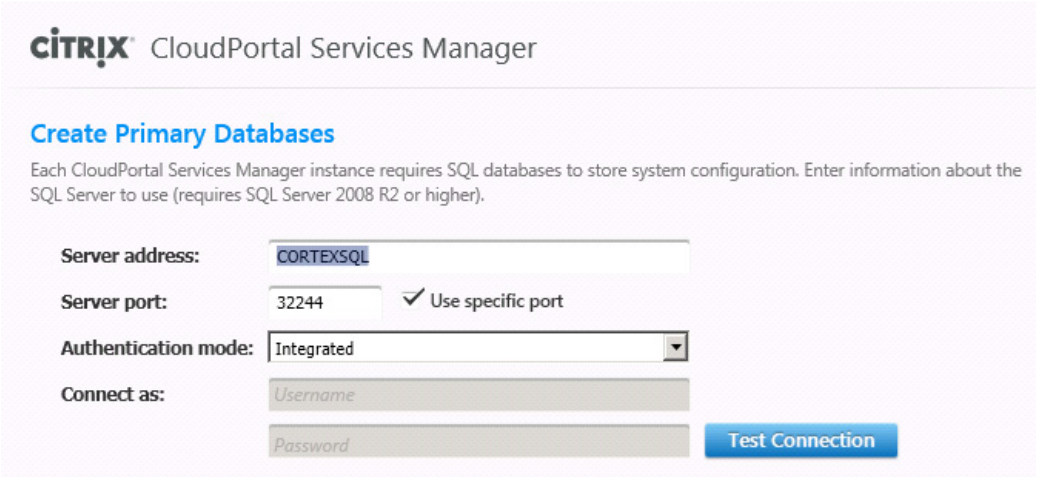

When the installer completes “Create System Databases”, verify that the two CPSM system databases OLM and OLMReports are restored on the primary replica.

Supporting multi-subnet failover (MultiSubnetFailover=True) requires the format of all CPSM connection strings to be re-evaluated and tested.

Adding CPSM Databases to the Availability Group

To add the CPSM databases OLM, OLMReports, and OLMReporting (Reporting database) to the database availability group, complete the following steps:

1. Logon to the SQL server that hosts the primary replica and start SQL Server Management Studio.
2. For each of the 3 databases, change the recovery model from “Simple” to “Full”, and do a full backup.

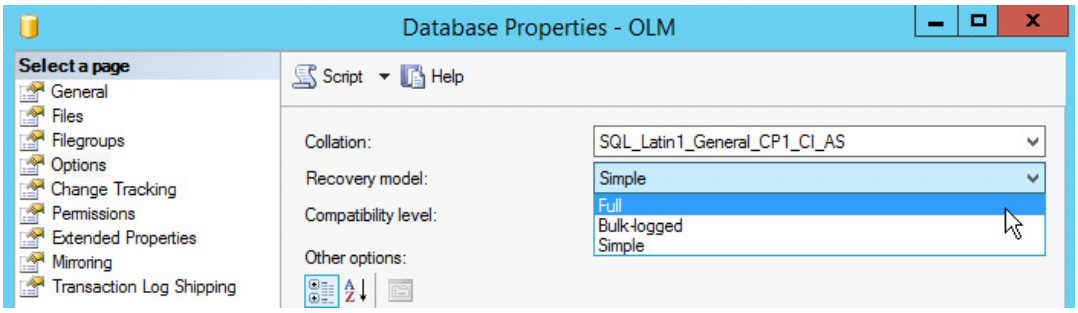

3. In Object Explorer, browse and expand the Availability Groups.
4. Right-click the relevant group, and then click Add Database.
5. On the Select Databases page, all databases that are eligible to become the primary database for the new availability group are listed in the table. The CPSM databases should be shown as “Meets requirements”. Use the checkboxes to select the 3 CPSM databases and click Next. Alternatively these databases can be selected and added individually.
6. On Select Initial Data Synchronization page, accept the default Full option, and Next.

7. If the Validation page displays the results of six checks as successful, click “Next” to continue. If any test fails, action must be taken to correct the error items and re-run the validation.
8. On the Summary page, verify the configuration of the replica, and then click Finish.

Login Replication

CloudPortal Services Manager SQL logins are not automatically replicated in the availability group so that need to be manually created on the secondary replicas.

From Management Studio, connect to the secondary replicas, replace the placeholders {…} with the actual account password in the below SQL statement and run:

```
USE Master
EXEC sp_addlogin 'CortexProp', '{CortexPropPassword}'
EXEC sp_addlogin 'OLMUser', '{OLMUserPassword}'
EXEC sp_addlogin 'OLMReportsUser', '{OLMReportsUserPassword}'
EXEC sp_addlogin 'OLMReportingUser', '{OLMReportingUserPassword}'
Go
```

Provisioning Engine

CloudPortal Services Manager provisioning engine is dependent on Microsoft Message Queuing, for high availability requirement, MSMQ needs to be clustered, and so as the CPSM provisioning engine.

Prerequisites

The provisioning server cluster (Windows Server 2012 Failover Cluster) is created, and in addition, all servers must be able to see a shared storage device (i.e. a SAN drive) and be able to take ownership of it. Shared storage is not a requirement for Windows Server 2012 clusters but is a requirement for some Microsoft services, in this case Microsoft Message Queuing.

Installation

Install and configure CPSM Provisioning role on each of the cluster nodes using the CPSM v11.x installer: http://support.citrix.com/proddocs/topic/ccps-115/ccps-install-config-roles-gui.html. If you configure the Provisioning role on the secondary nodes with the same service accounts, make sure the password matches the ones specified for the same accounts when configuring on the primary node.

To configure the cluster, complete the following steps:

1. On the cluster node, open Failover Cluster Manager.
2. Expand Cluster and right click on Roles and select Configure Roles.
3. Click Next and select Message Queuing and click on Next.

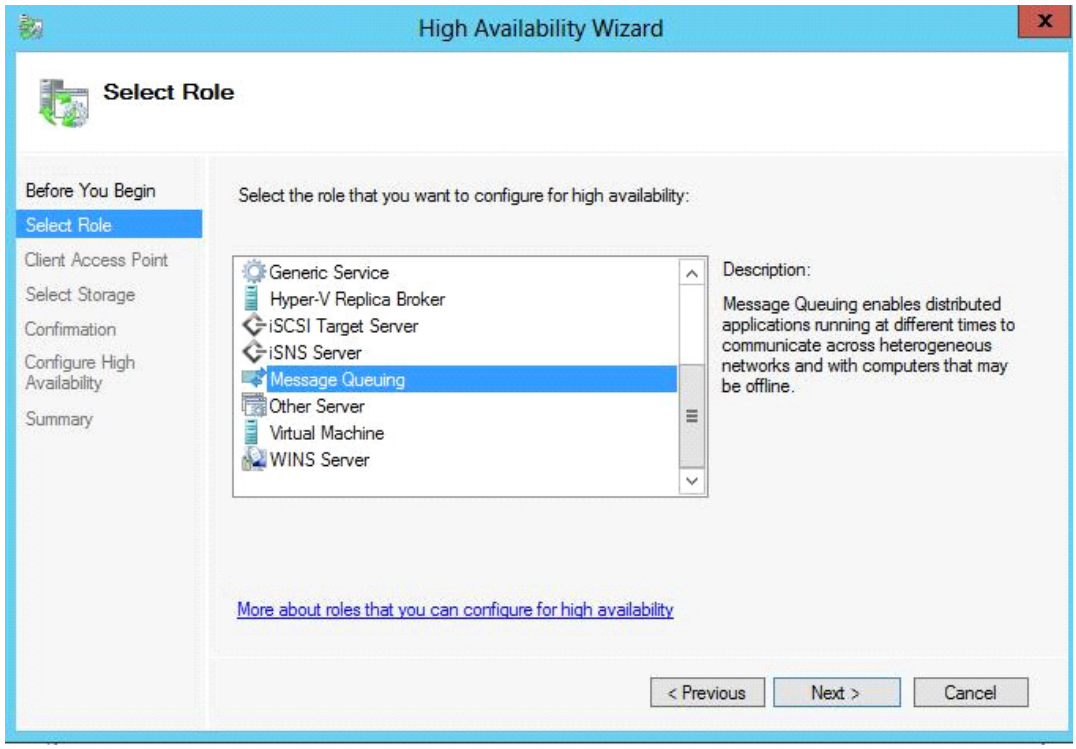

4. Enter the name that the clients will use to access this cluster role. Click Next.
5. Select the shared disk drive name to assign to the cluster role, and Next.

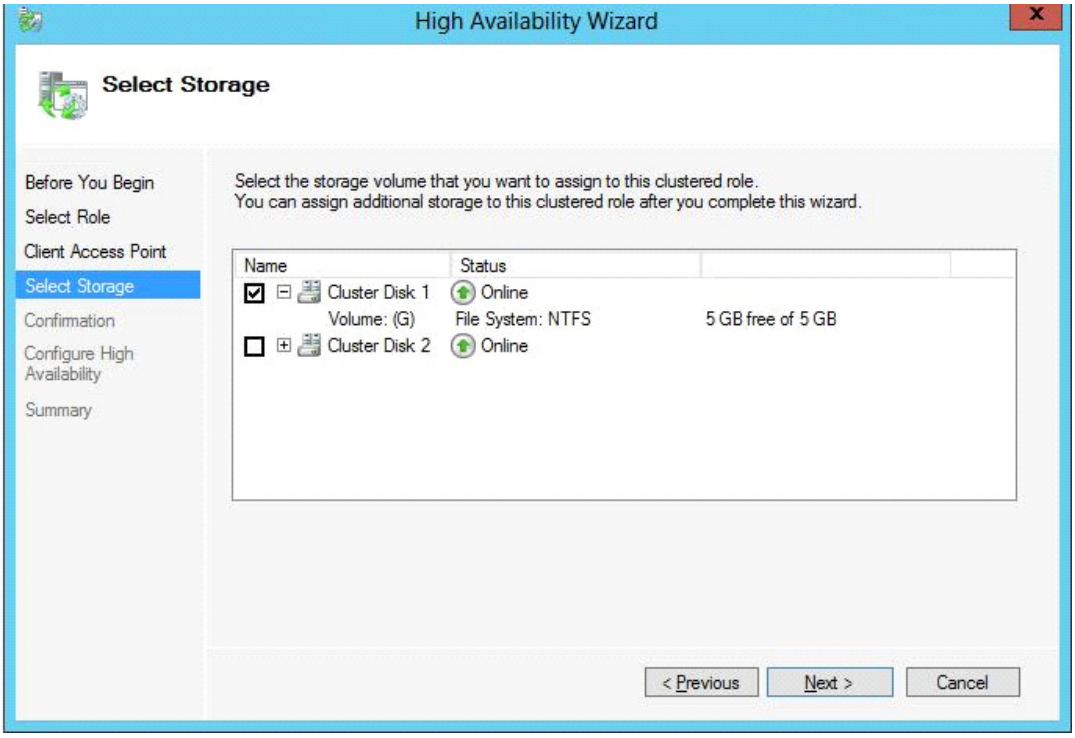

6. Click Finish when “High availability was successfully configured for the role” message and summary are displayed.
7. Right-click on the newly added role above, and select Add Resource > Generic Service.

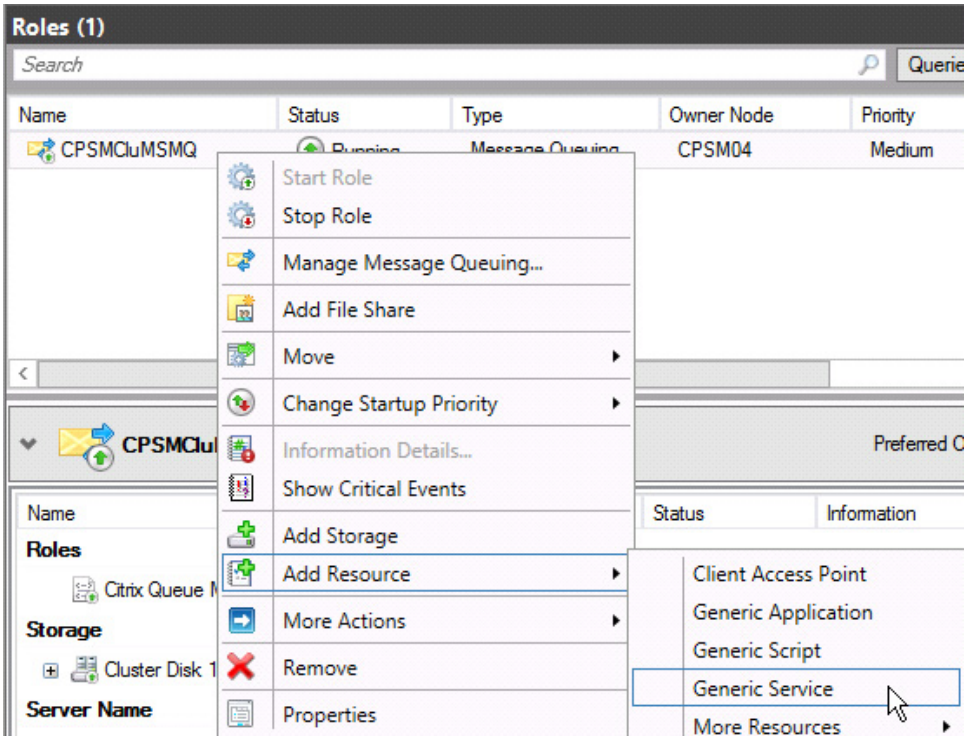

8. Select “Citrix Queue Monitor Service” from the list, click Next, follow the wizard and finish.
9. Right-click the new resource “Citrix Queue Monitor Service”, and select Properties. On Dependencies tab, add MSMQ-{your new cluster role} as the dependency, and click Apply.

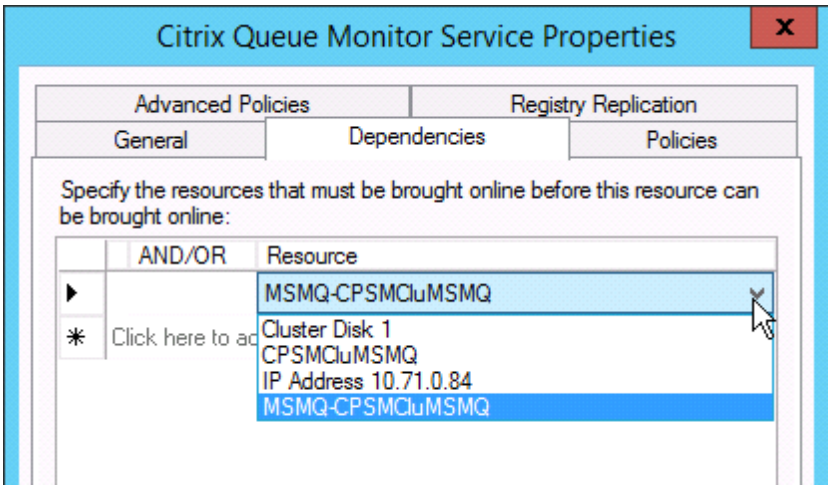

10. Select General tab, check “Use Network Name for computer name”, and OK. This step must be performed after step 9, otherwise an error would occur.

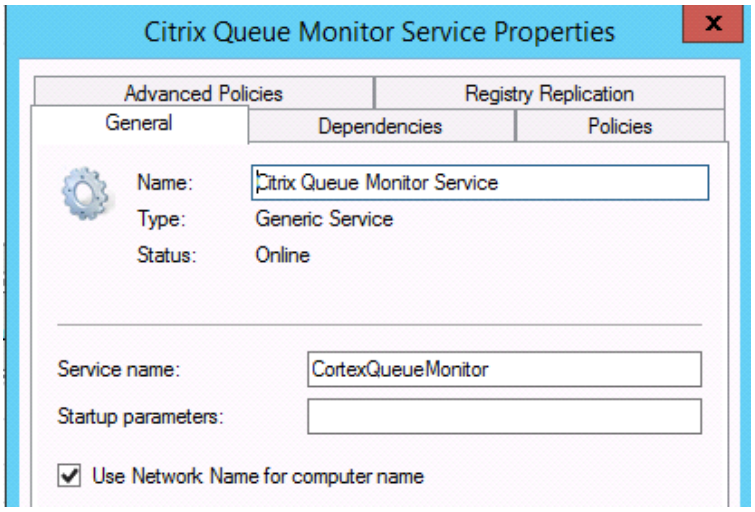

11. On each server node, stop “Citrix Queue Monitor Service”, and open Registry.
12. Navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CortexQueueMonitor, replace the value off the key “DependOnService” with MSMQ$MSMQ-{YourClusterRoleName}.

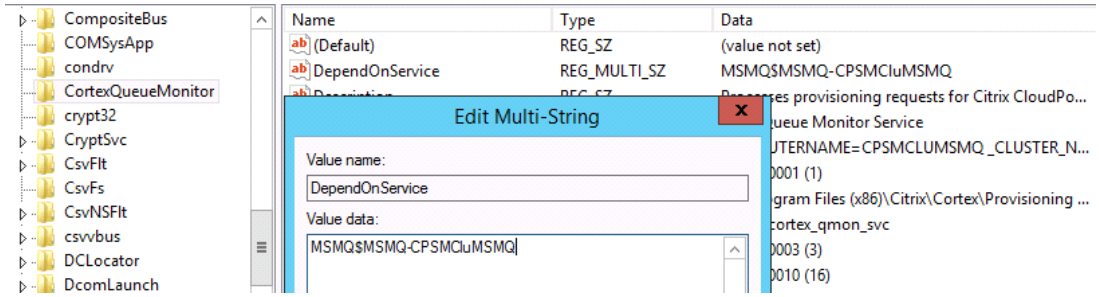

13. From Failover Cluster Manager, Bring Online “Citrix Queue Monitor Service”.
14. Logon to CloudPortal Services Manager portal as a Service Provider Administrator, go to Configuration > System Manager > Locations, and expand the relevant location.
15. Change the server name or IP address of the following queue paths to the Message Queuing cluster role name or IP (not the Windows server cluster name), and Save.

- Primary Queue Path
- Bulk Queue Path
- Usage Data Queue Path

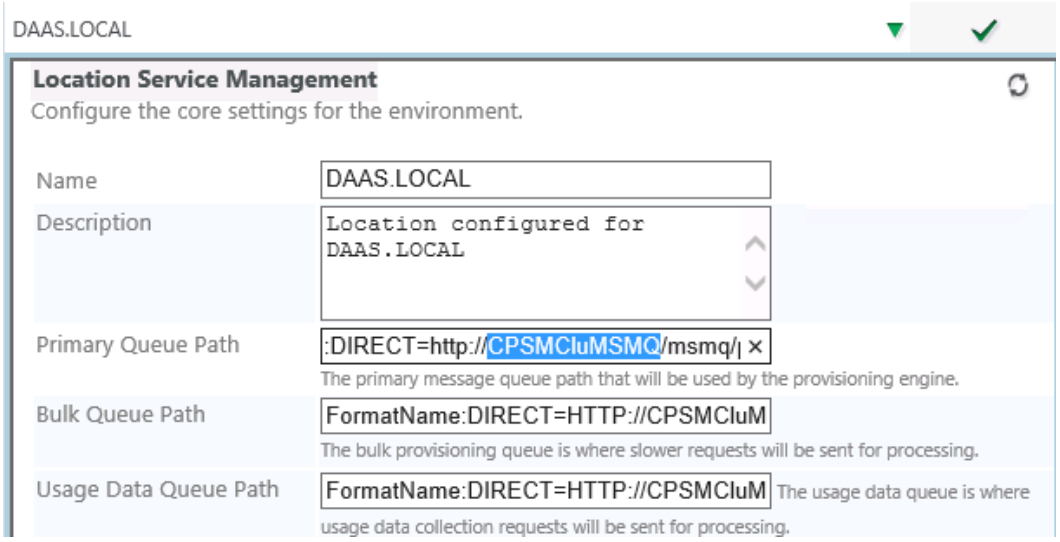

16. To view the message queue status, requests, and journals for monitoring and troubleshooting purposes, you can no longer use the local Message Queueing on each of the nodes, instead, from Failover Cluster Manager, select Message Queuing cluster role, and open Manage Message Queuing…

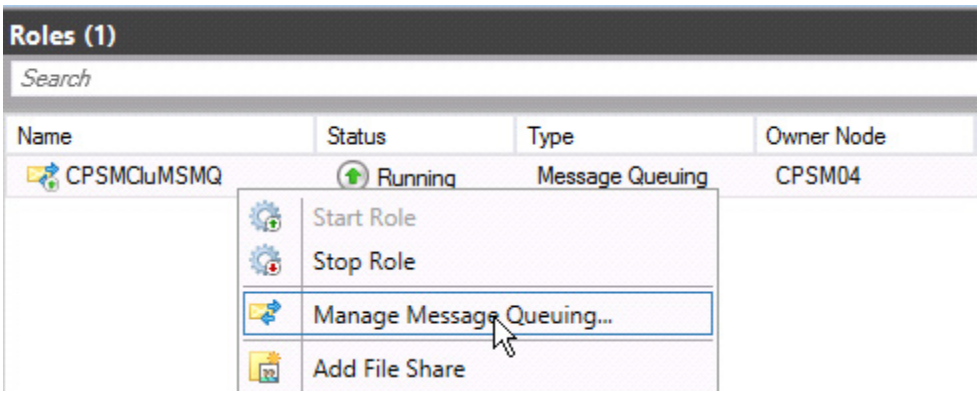

Directory Web Service

The directory web service is typically located on the same server as the Provisioning Engine and listens on port 8095. When the CPSM provisioning server is clustered, the Directory Web Service is also installed on all cluster nodes (refer to Provisioning Engine section).

To configure CPSM to achieve Directory Web Service high availability, complete the following steps.

1. Logon to CPSM portal as a service provider administrator.
2. Navigate to Configuration > System Manager > Servers. Click “Refresh Server List”, the provisioning server cluster name should appear in servers list.

3. If you are configuring for a remote location domain and the DNS of the server names may not be resolvable in the primary location domain, expand the server cluster, and enter IP address in Alias field, and Save.

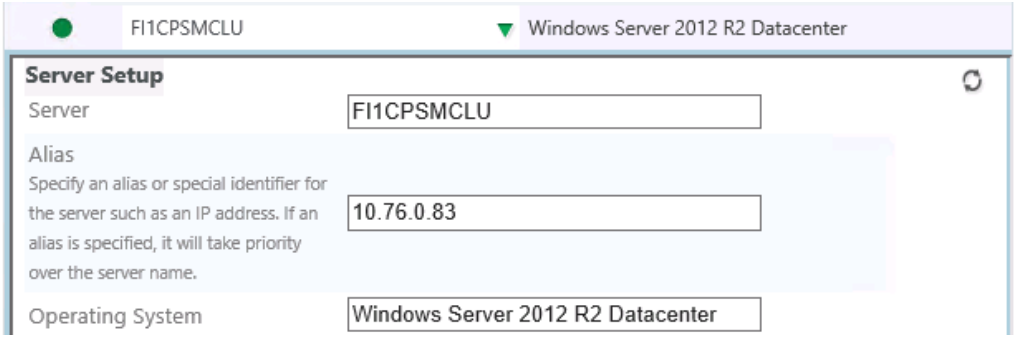

4. Click Server Roles link on the left or got to Configuration > System Manager > Server Roles, expand the provisioning cluster, tick Directory under Server Connection Components and Save.
5. Click Server Connections link or go to Configuration > System Manager > Server Connections, expand the existing entry for the Directory role, select the provisioning cluster name from the Server dropdown list, and click Save.
6. Test the connection via the icon on the right. It should go green if valid and there no firewall blocking issues.

Directory Web Service can also be load balanced, in this case the configuration should use the load balanced VIP address instead of the cluster name/address.

Web Portal and API

The web portal (interface) and API are under the same standard .NET framework 4.0 website in IIS which listens on ports 80 and/or 443. To enable high availability, it is recommended to run two (or more) nodes in a standard load-balancing setup. Sticky sessions are required.

To install the web portal, complete the following steps:

1. Install and configure Web server role on all load balanced servers: http://support.citrix.com/proddocs/topic/ccps-115/ccps-install-config-roles-gui.html. If you are adding additional web servers to an existing deployment with functional services, it is recommended to skip the Service Package Import (deselect the services and properties) or select Ignore to the properties of all enabled services.
2. The following items of CPSM web on the first/primary web server must be replicated to all load balanced web servers. It is recommended to replicate all the files in CortexDotNet and CortexAPI sub sites (except the web.config files specific to the local sites) from the first/primary web server to all the rest of the web servers.

- Images for branding
- Stylesheets for branding
- Any custom downloads
- Any custom DLLs or pages
- Web.config configuration changes

- IS Security and authentication changes

3. For AD Sync to use the same key for API on all web servers, the key must be exported from the first web server and imported to the rest of the web servers:

- Logon to the first/primary web server, open Command Prompt or Windows PowerShell (if UAC is on, use “Run as Administrator”), navigate to the directory: C:\Windows\Microsoft.NET\Framework\v4.0.30319>
- Run the following command (the export location can be anywhere, in this case we used C:\temp) aspnet_regiis -px "CortexAPI" "C:\temp\cortexapi.xml" –pri
- Copy the exported XML file to each of the load balanced web server.
- On the secondary load balanced web server open Command Prompt or Windows PowerShell (if UAC is on, use “Run as Administrator”), and navigate to the same directory: C:\Windows\Microsoft.NET\Framework\v4.0.30319>
- Run the following command to import the key: aspnet_regiis -pi "CortexAPI" "C:\temp\cortexapi.xml"

4. Add the same URL (internal and/or external) host headers to Cortex Management site on all web servers in IIS Manager.
5. Update the DNS for all the above host headers to point to the load balanced VIP address.
6. Recycle CortexMgmt application pool via IIS Manager on all web servers.

Reporting

The high availability of CloudPortal Services Manager Reporting role is dependent on the SQL Reporting Services HA configurations. The MS SQL Reporting Services achieves HA via a scale-out deployment so that they share the same report server database: https://msdn.microsoft.com/en-us/library/bb522745.aspx.

The data source DB (OLMReporting) of CPSM Reports can be added to AlwaysOn availability group (refer to Adding CPSM Databases to the Availability Group section), in this case when installing Reporting role via CPSM v11.x installer, the listener name and port should be specified for the Reporting SQL server.

Other Web Services

Similar to the Directory Web Service, other service integration related CPSM web services like Exchange, Lync, XenDesktop, XenApp, and IIS web services, etc. can be deployed to multiple servers for high availability:

1. Install and configure the web service on all HA servers using the installer: http://support.citrix.com/proddocs/topic/ccps-115/ccps-services-deploy.html.
2. Logon to CPSM portal as a service provider administrator to update the web service call configurations.
3. Navigate to Configuration > System Manager > Servers. Click “Refresh Server List”, if the cluster or VIP name does not exist on the list, click Add a Server link, enter the VIP name and click Add Server.

4. If you are configuring for a remote location domain and the DNS of the server names may not be resolvable in the primary location domain, expand the “server” name you have just added, enter the IP address in Alias field, and Save.
5. Click Server Roles link on the left or got to Configuration > System Manager > Server Roles, expand newly created VIP Server placeholder, tick the appropriate role under Server Connection Components and click Save.
6. Click Server Connections link or go to Configuration > System Manager > Server Connections, expand the existing entry for the web server or connection role to be updated, select the newly created VIP Server name from the Server dropdown list, and click Save.
7. Test the connection via the icon on the right. It should go green if valid and there no firewall blocking issues.

Initial Configuration of App Orchestration and CloudPortal Services Manager

This section describes initial software configuration procedures for the Citrix App Orchestration and CloudPortal Services Manager software, including:

- Running the PowerShell script to configure App Orchestration Group Policy and define policy settings
- Configuring SSL on the App Orchestration configuration server and creating a new App Orchestration deployment
- Configuring prerequisites for CloudPortal Services Manager, setting up server roles, and defining the primary location

Using the checklists in the documentation is recommended to avoid installation and configuration errors when performing these procedures.

Setting Up the App Orchestration Configuration Server

This section illustrates configuration steps that are required to set up the App Orchestration configuration server. It shows steps to prepare for setting up the configuration server (including how to create a Group Policy Object (GPO), configure the required policy settings, and install the SSL certificate). The procedure shows how to define settings and a delivery site for a new App Orchestration deployment. Generally these steps are performed by a top-level CSP administrator to initially set up and configure the environment.

Instructions	Visual
RDC to the App Orchestration server and launch **Citrix App Orchestration Infrastructure Tools** from the Apps menu.	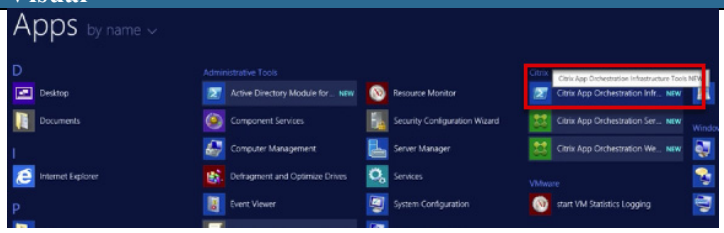
Run the Citrix-provided PowerShell script **New-CamGPO.ps1.** The script creates a Group Policy Object called CtxCloudAppManagement and configures the required policy settings.	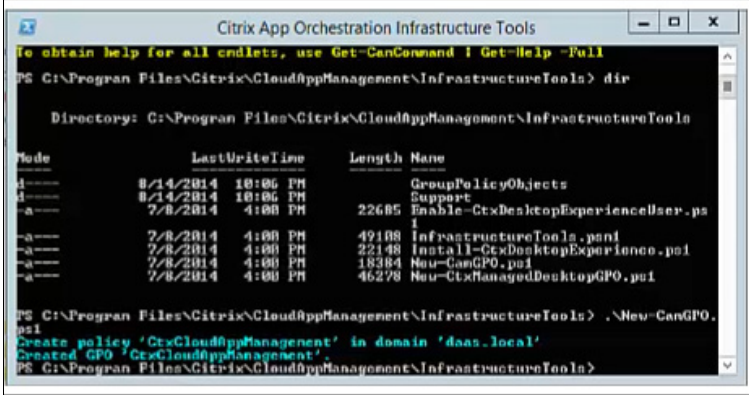
After you create this policy, run **gpupdate** to link the GPO to the following objects: · App Orchestration root OU in the shared resource domain. · App Orchestration root OU in each additional private tenant resource domain that you create.	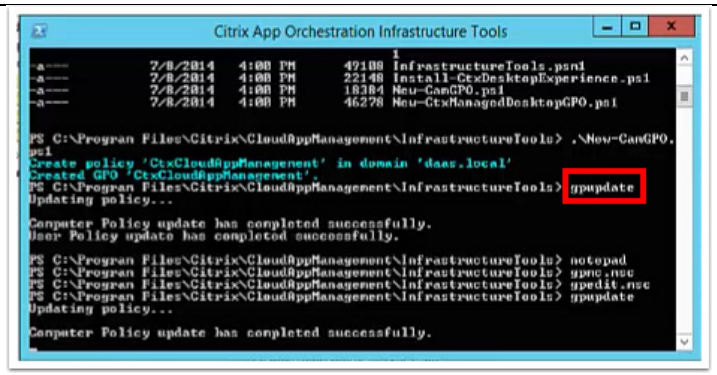

At the end of the App Orchestration Setup procedure, enabling the checkbox "Launch Server Config on exit" causes the Citrix App Orchestration Server Configuration tool to start automatically.	
Click **Create a new deployment.**	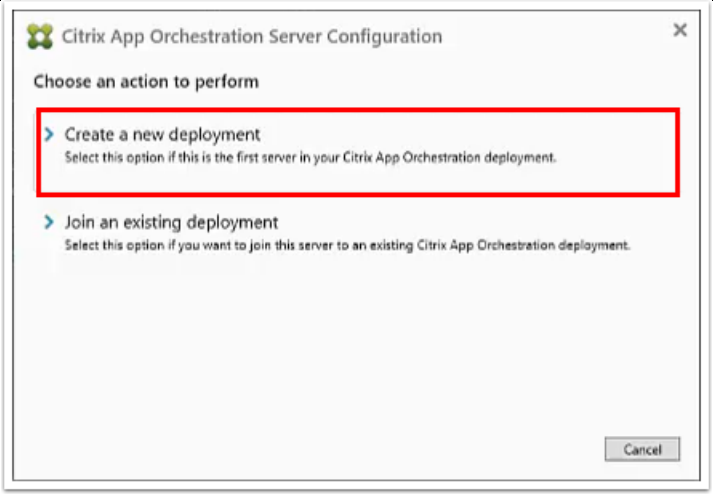
Enter the information for the new App Orchestration deployment, including the database name, the database server, and the administrative group for the DaaS Shared Infrastructure (e.g., DaaS\AOAdmins). Click **Yes** when prompted: "The specified database does not exist. Do you want to create it?"	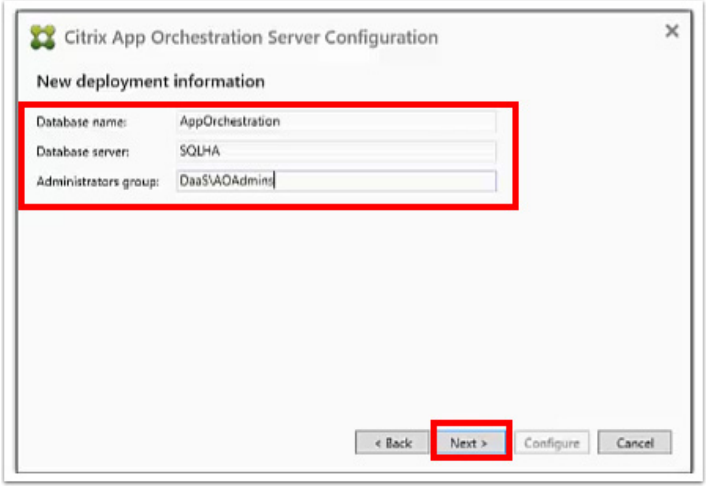

The SSL information screen advises that App Orchestration requires an SSL certificate to secure machine-to-machine communications. (It is assumed that an SSL certificate is established and available to meet this prerequisite.) Click **Browse certificates installed on the local machine.**	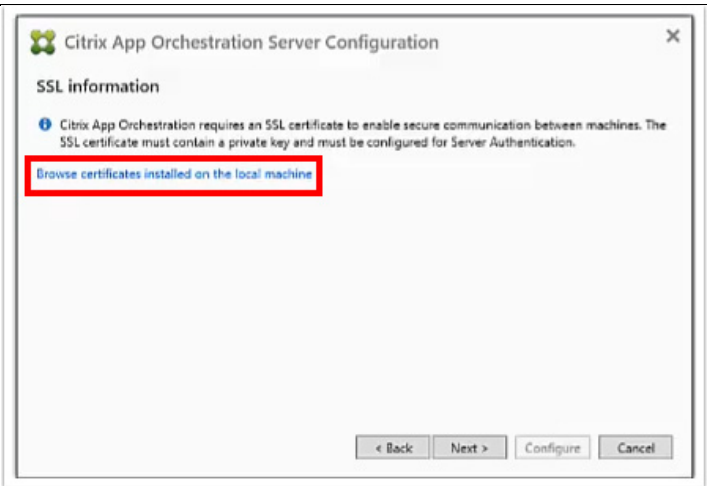
Select a certificate and click OK.	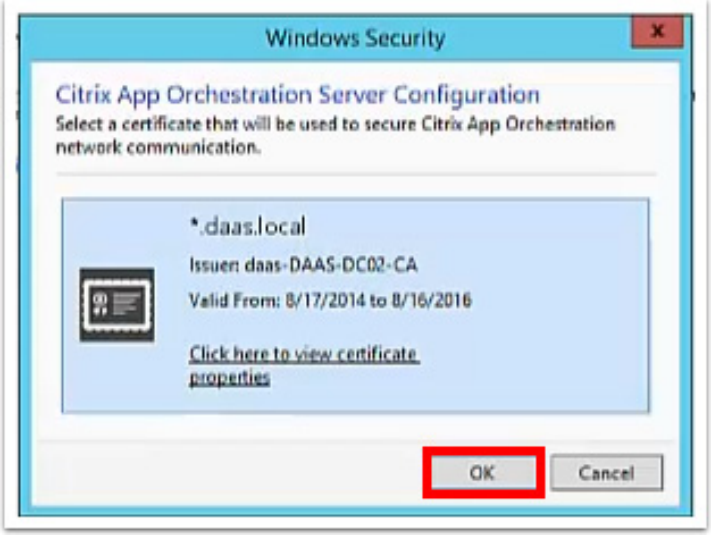
After you select the certificate, the configuration wizard is ready to bind it during the installation. Click **Next**.	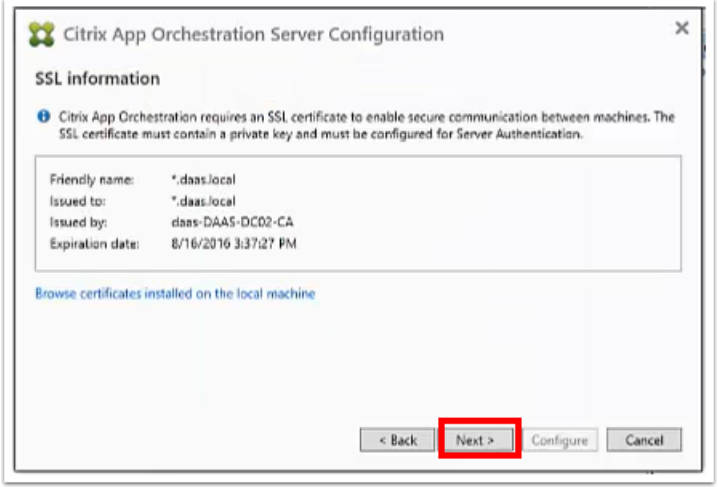

Click **Configure**.	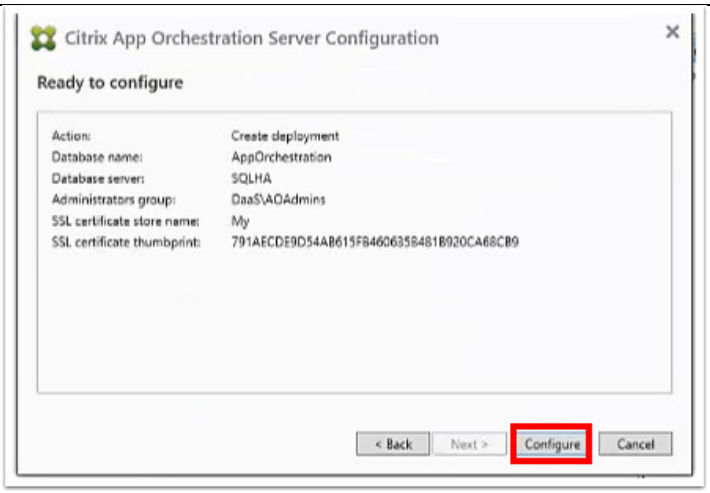
A progress screen shows the steps for App Orchestration Server Configuration as each one completes.	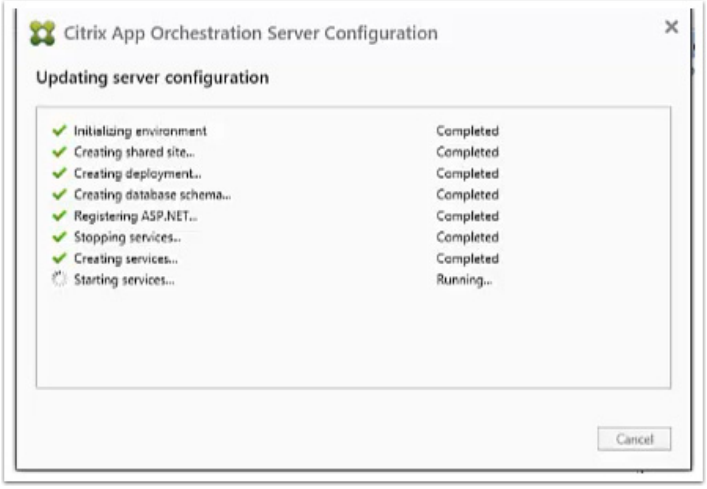
The “Configuration was successful” screen appears when the Server Configuration process is complete. Click **Close**.	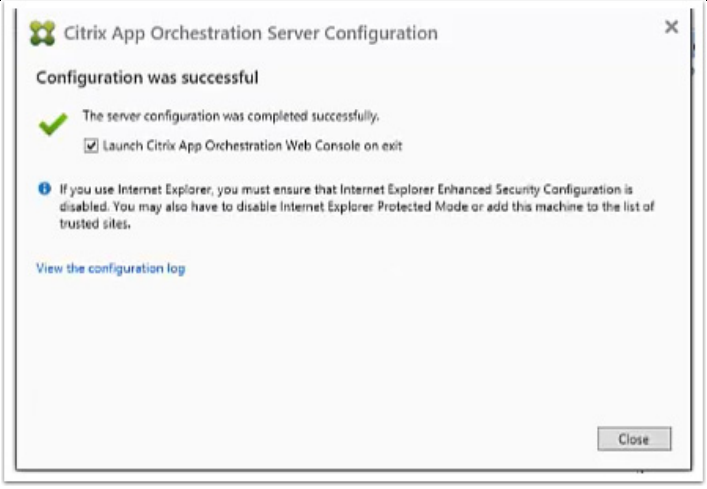

Defining Settings for the App Orchestration Deployment

This procedure shows how to define settings and a delivery site for a new App Orchestration deployment. As shown below, App Orchestration uses a phased deployment. This section focuses on the “Define” phase.

Review the following assumptions:

- Make sure each machine configured and deployed by App Orchestration has all of the minimum system requirements installed, including the Microsoft .NET Framework.
- Each machine under App Orchestration control requires PowerShell remoting.
- If there are multiple datacenters, make sure the App Orchestration configuration server can ping IP addresses in each datacenter.

Instructions	Visual
Launch App Orchestration web console on the App Orchestration server (AO01). Login as the administrator AOAdmin.	
Enter the App Orchestration Root OU and the user name (AOAdmin) and password for the App Orchestration administrator. Click **Next**.	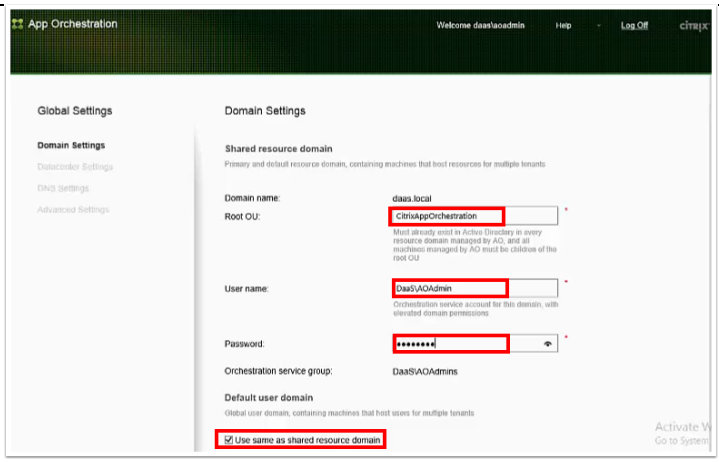
Enter the datacenter name (CiscoCVD) and the license server address. Click **Next**.	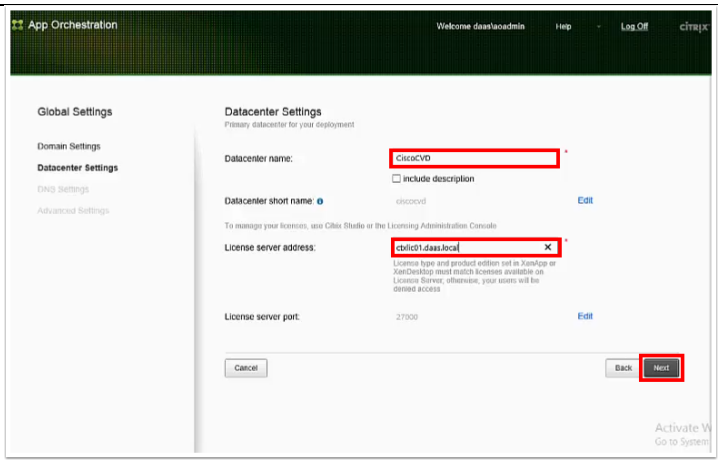

Enter the external DNS suffix (in this case, go.ciscocvd.com).	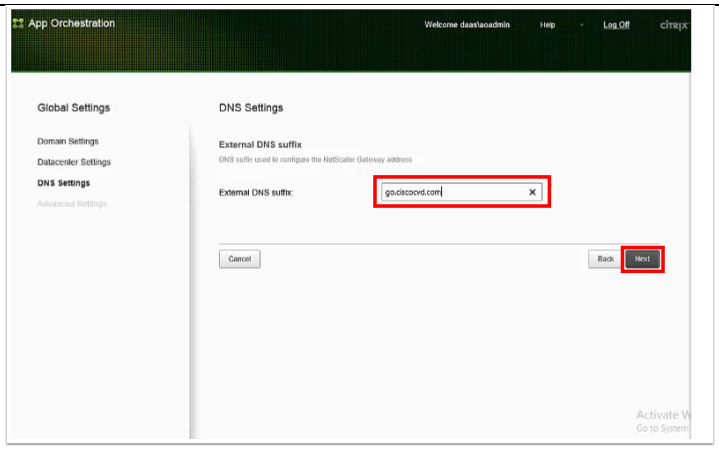
On the Advanced Settings screen, change the shared controller management network and shared Delivery Group management network to **Infrastructure**. Set the NetScaler Gateway address to http://go.ciscocvd.com . Note: the external NetScaler Gateway address can be changed later.	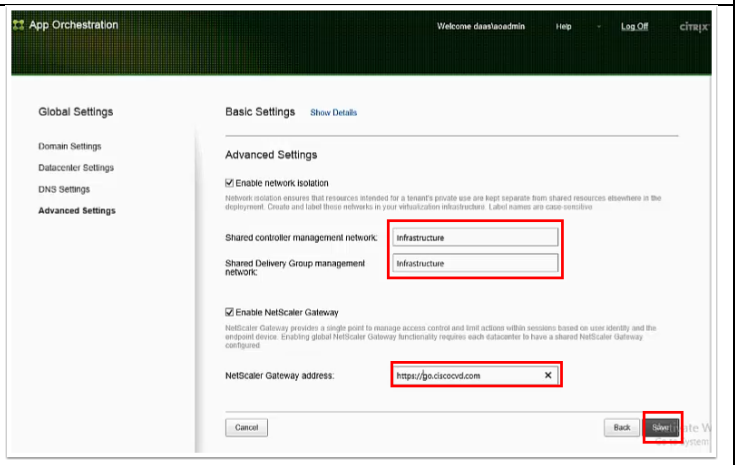
Configuration progress is visible in the App Orchestration Getting Started screen. As shown, the Define phase is complete and tasks in the Design phase (which is the focus of this section) still must be completed.	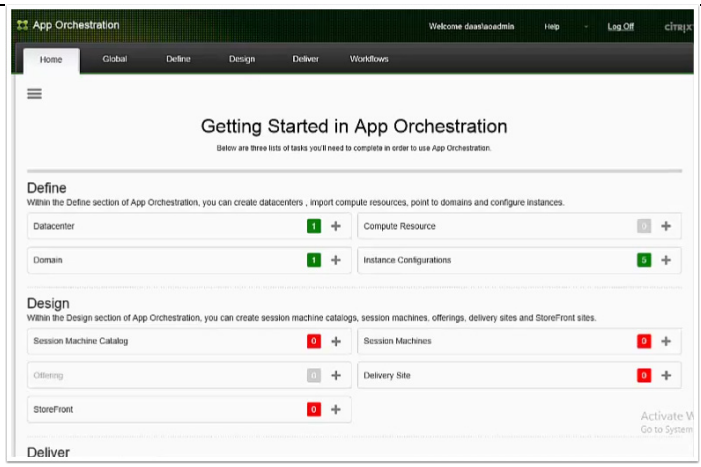

On the Delivery Controller (DC01), run the App Orchestration Install Center. Select the option to install the XenApp and XenDesktop 7.5 Delivery Controller and App Orchestration Agent. When prompted, click Yes to continue. Repeat the same process above to install on the second Delivery Controller (DC02).	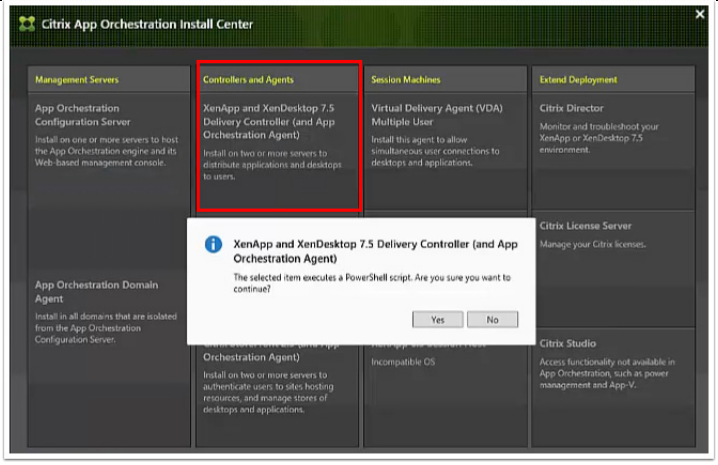

Setting Up the CloudPortal Services Manager Primary Location

Using the CloudPortal Services Manager console, server roles (provisioning, web, etc.) must be set up. Then, the primary CloudPortal Services Manager location is defined

Instructions	Visual
On the CPSM provisioning server (CPSM02), run **setup.exe** to launch the CloudPortal Services Manager console.	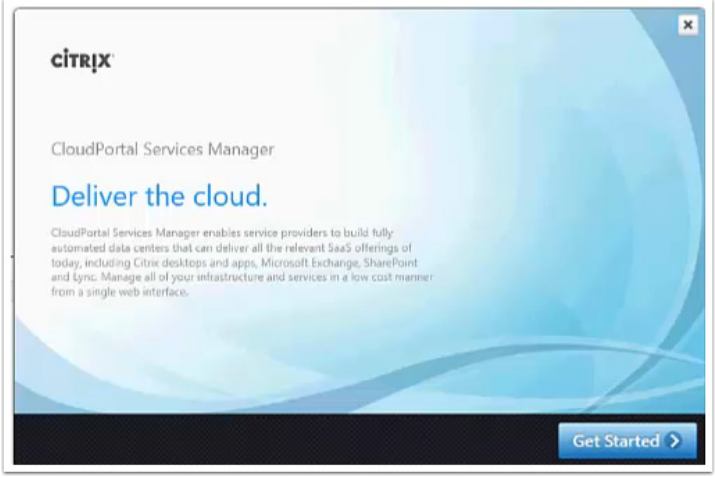
Select **Install CloudPortal Services Manager.**	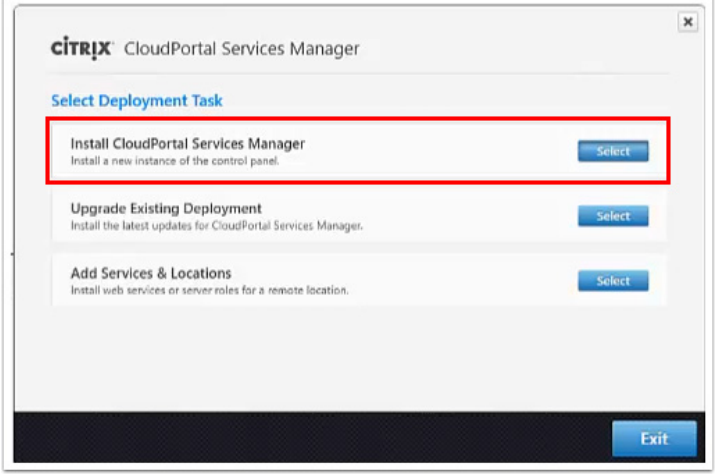

Select **Check Environment Prerequisites.**	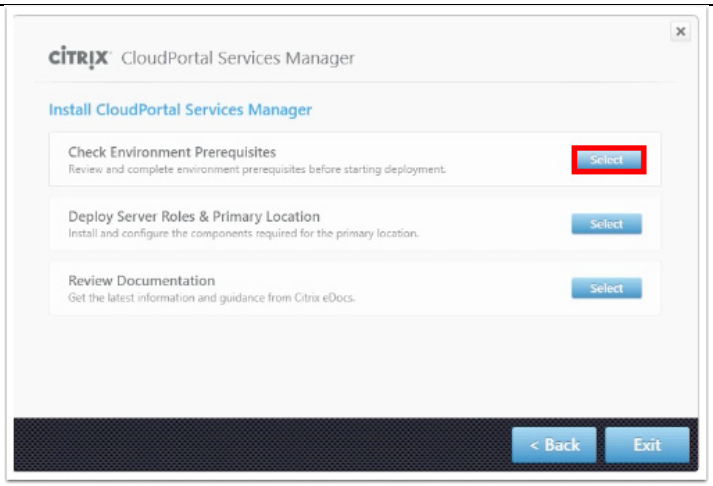
The prerequisites have been completed, so the environment is prepared. Click **Exit**.	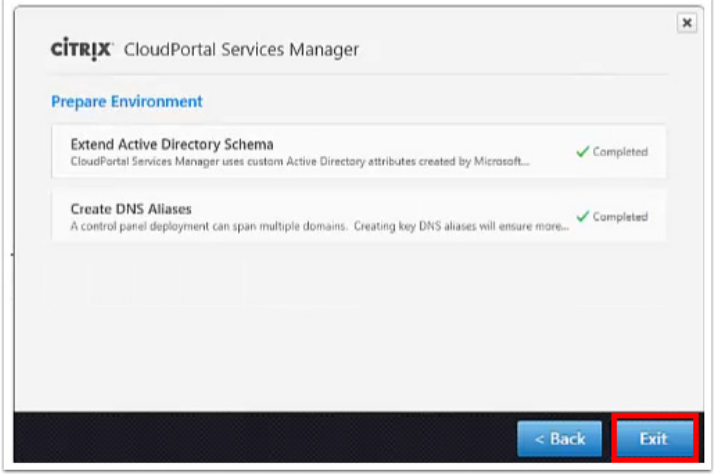
Select Deploy Server Roles & Primary Location	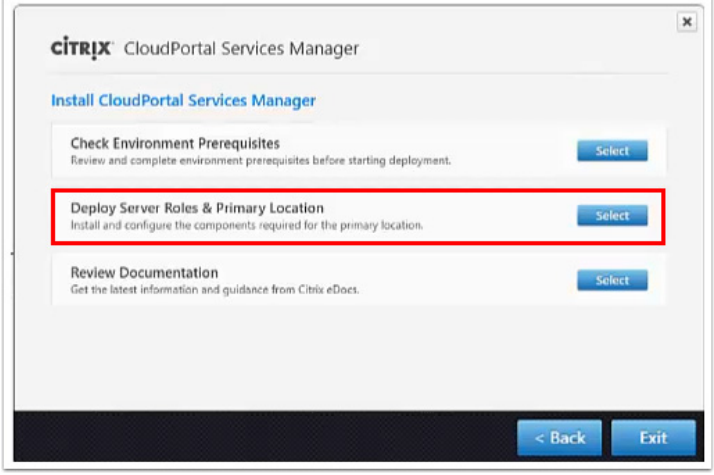

Click **Create System Databases**. This brings up **Install Configuration Tool**. When prompted, click **Install** to continue. Accept the License Agreement, and click **Next** to continue.	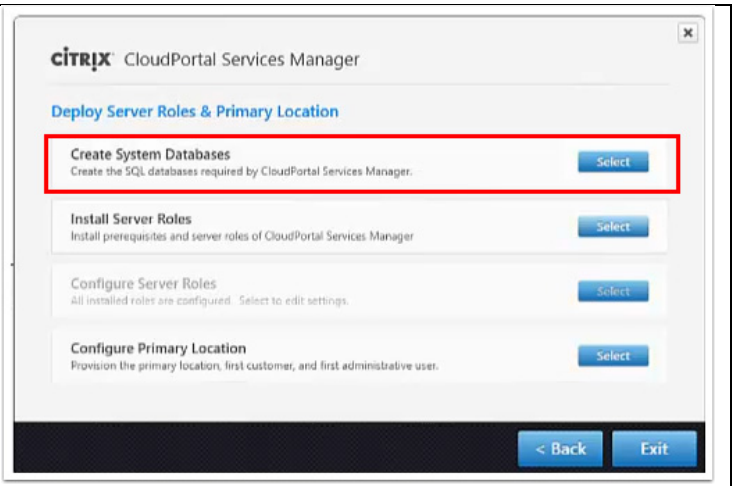
Click **Install** to confirm installation and continue.	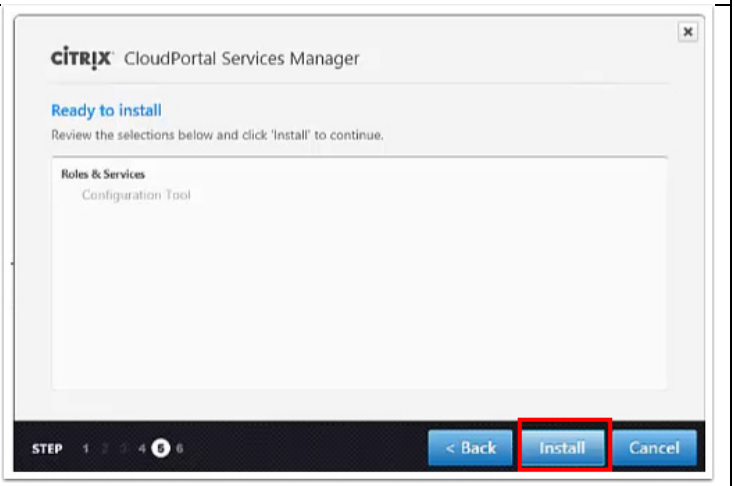
The installation completes. Click **Finish** to continue.	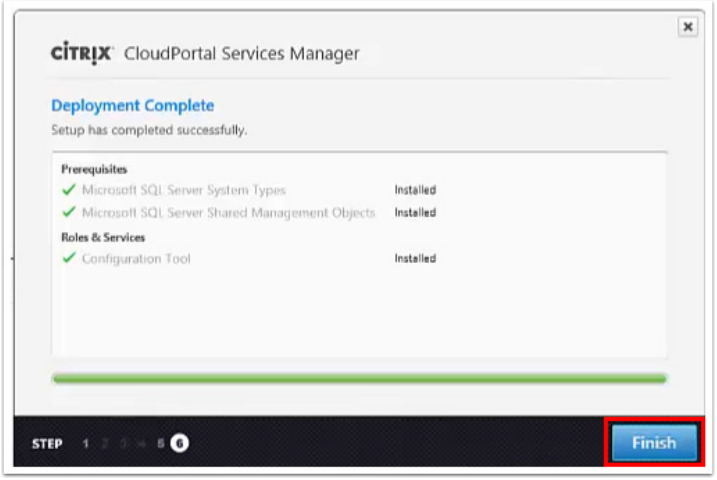

Create a deployment configuration file. This is an XML file that will be accessed from each machine that is configured. Click Next.	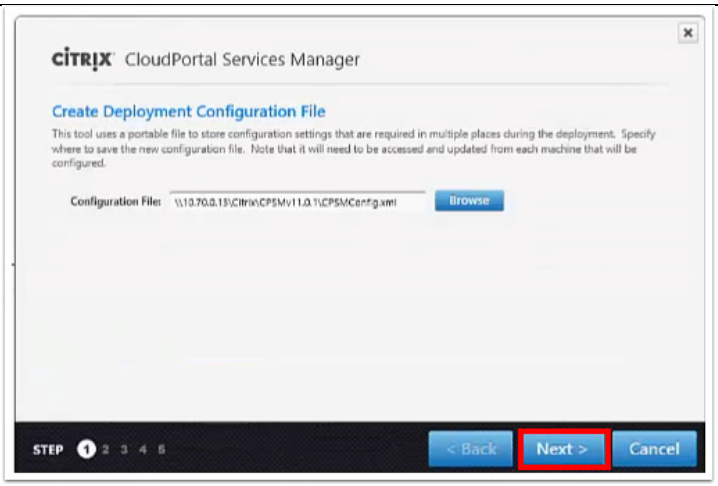
Enter the information for the primary database. Click **Test Connection** to confirm that the database can be accessed. Click Next.	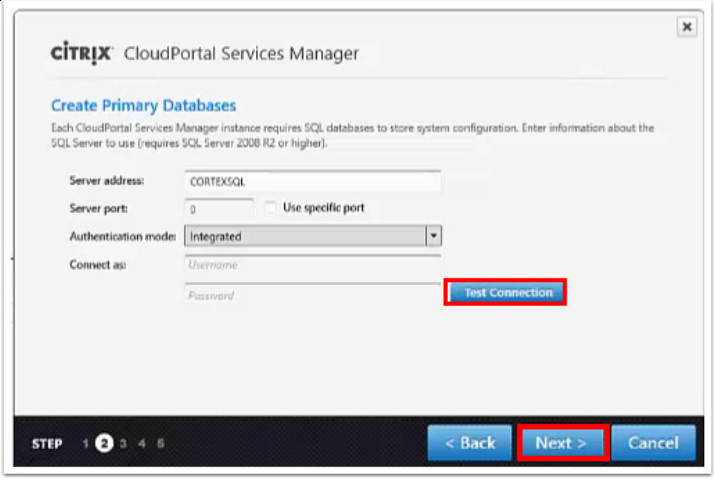
Configure database logins. Click **Next**.	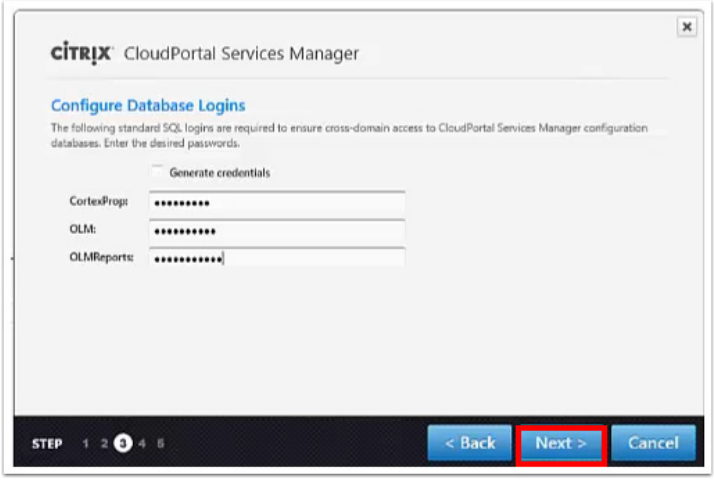

Review the summary and Click **Commit** to configure the CloudPortal Services Manager database.	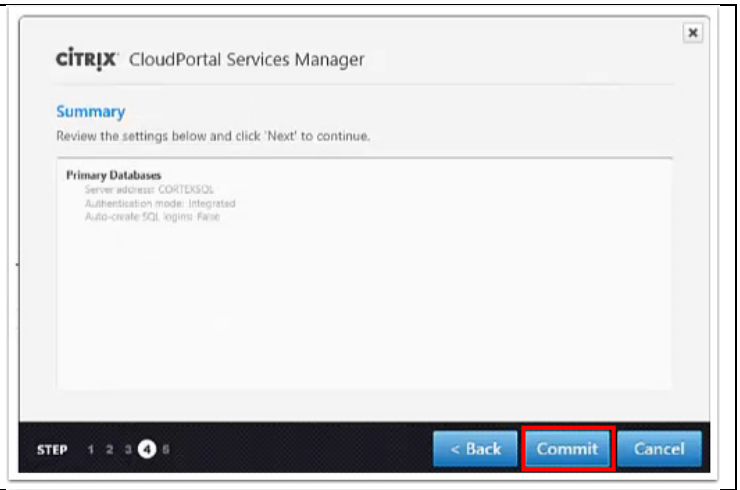
The configuration is complete. Click **Finish** to continue.	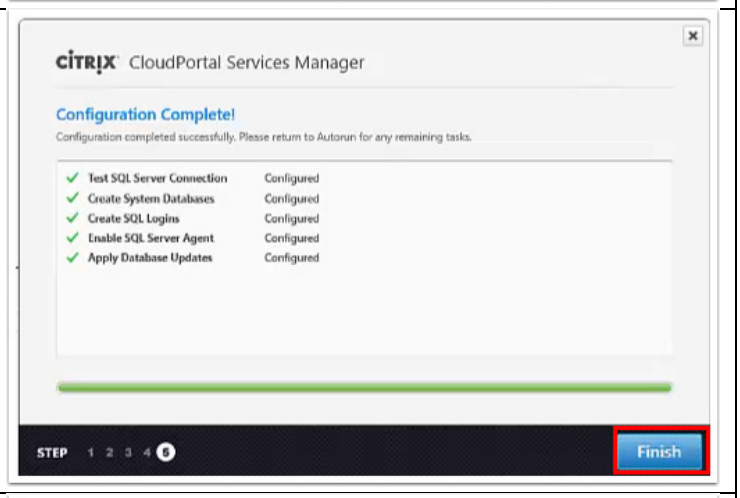
Click on **Install Server Roles**. When prompted, read and accept the License Agreement. Click **Next** to continue.	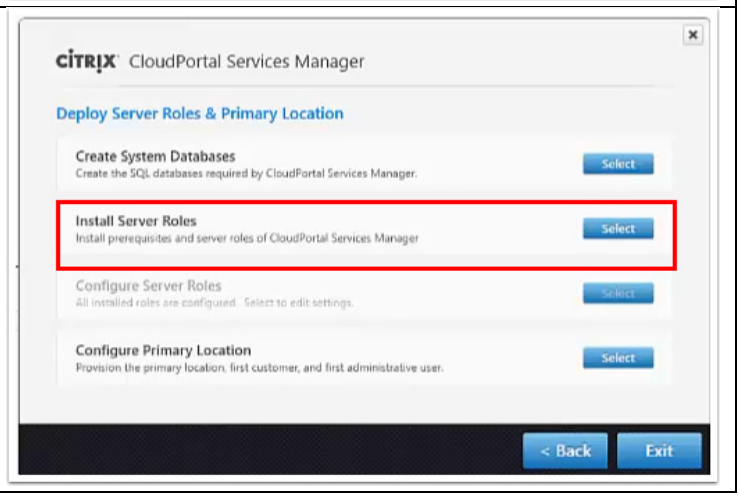

Select **Provisioning** and **Directory Web Service** as the components to install. Click **Next** to continue.	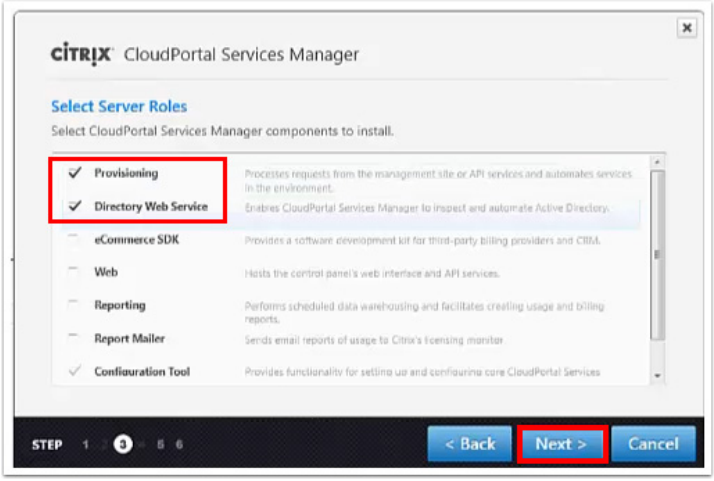
Review prerequisites. Click **Next** to continue.	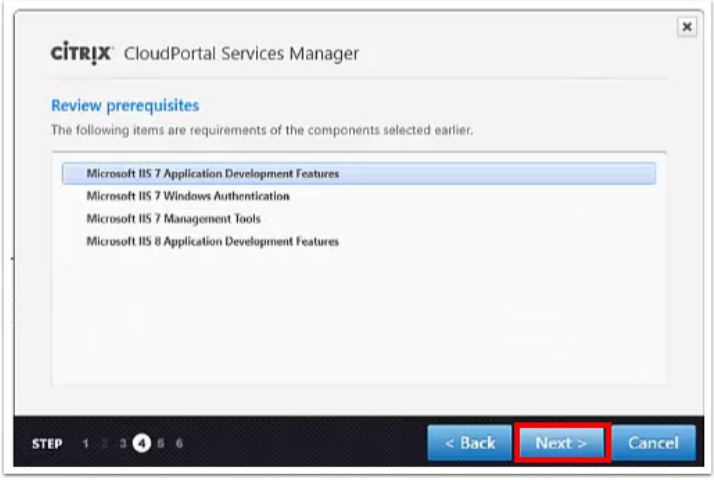
Review the selections and click **Install**.	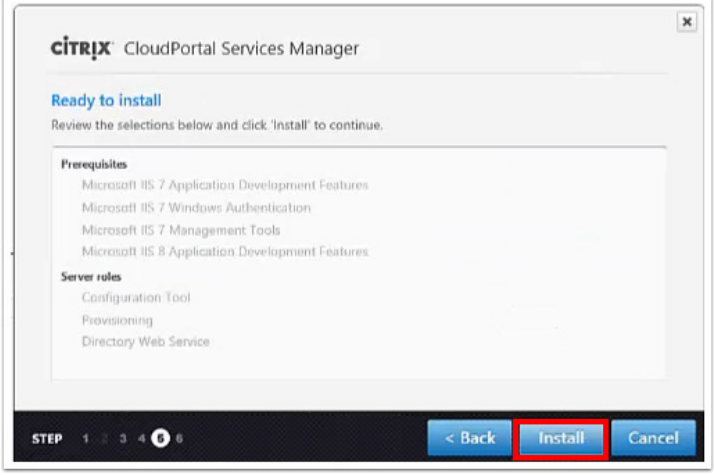

The CloudPortal Services Manager set up process is complete. Click **Finish**.	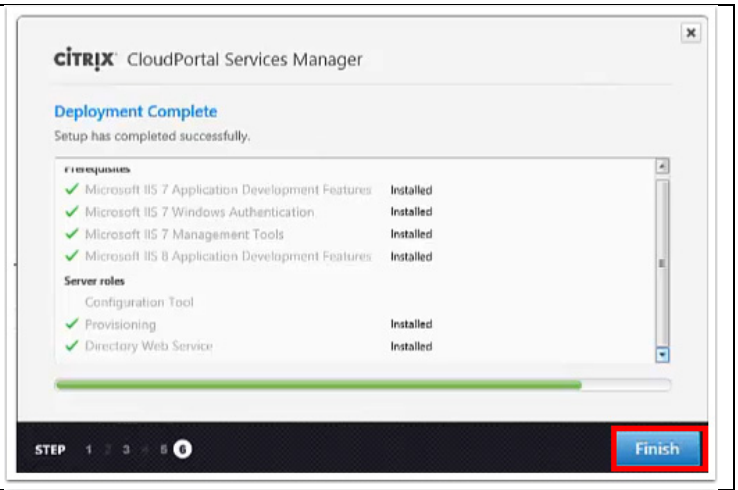
Next, select **Configure Server Roles.**	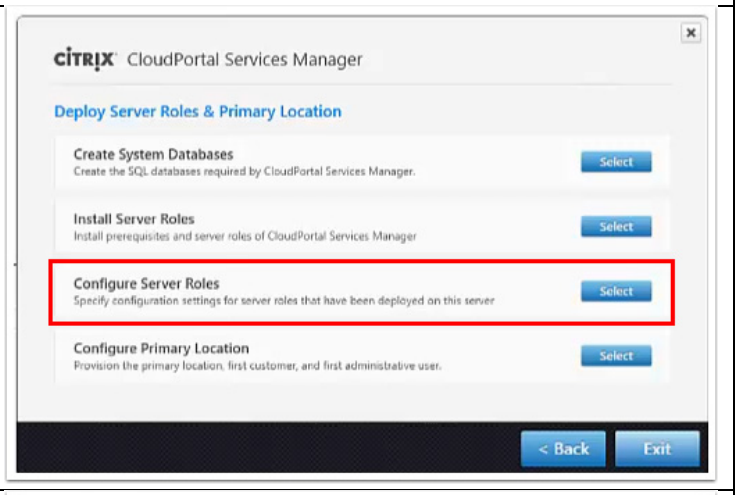
Specify the XML deployment configuration file created previously. Click **Next** to continue.	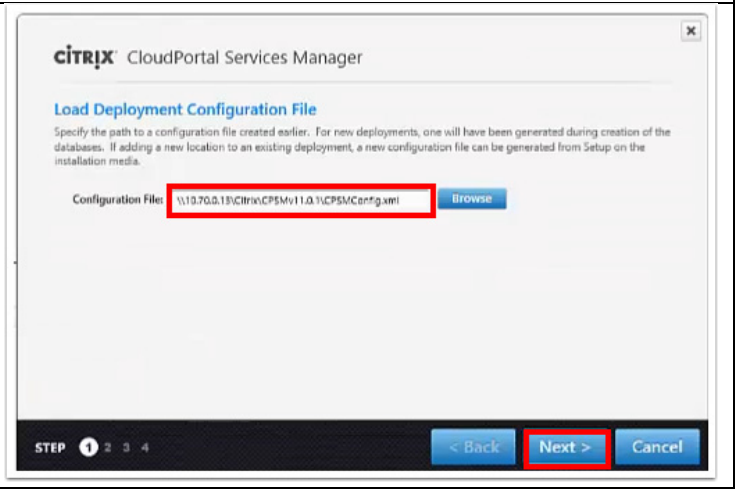

Select configuration tasks for **Provisioning** and **Directory Web Service**. Follow the prompts to configure the Mail Server, the Directory Web Service, the Queue Monitor, and Directory Monitor service. Click **Next** to continue.	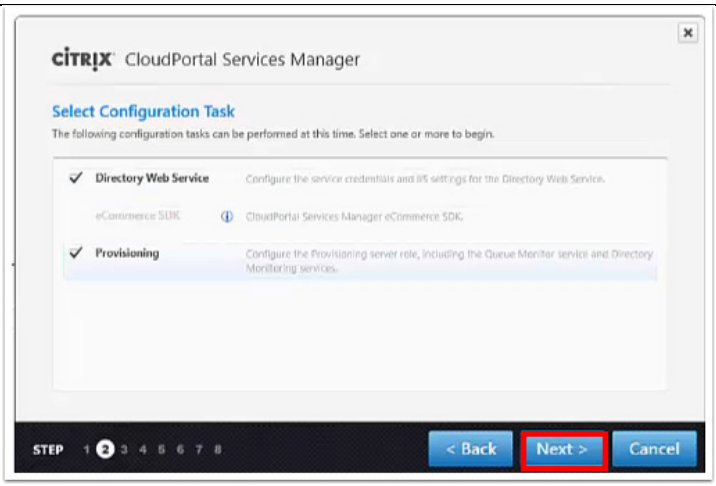
Review the summary and click **Commit**.	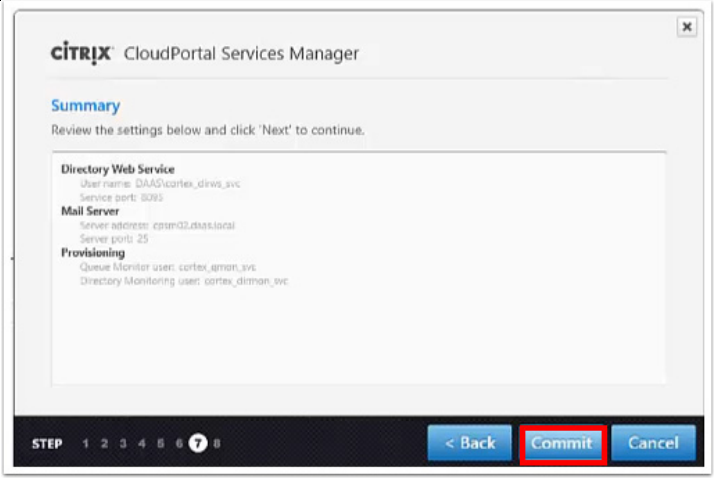
The task configuration is complete.	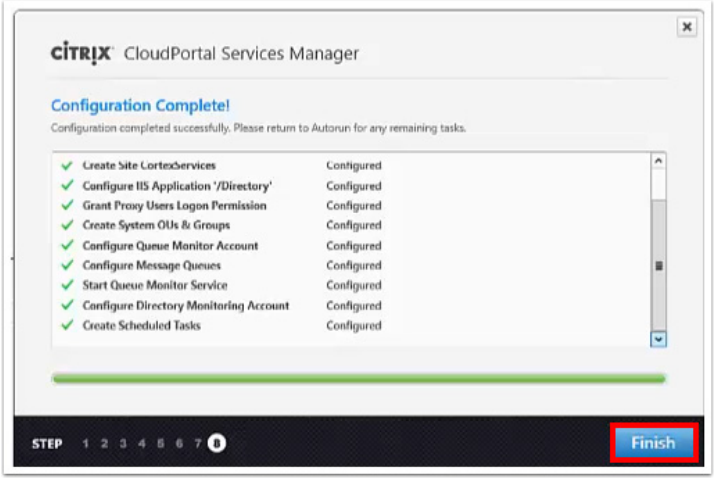

Repeat the **Install Server Roles** setup on the CloudPortal Services Manager web server (this time, the database is already created, so that step won't be necessary). Specify the same deployment configuration file that was created earlier.	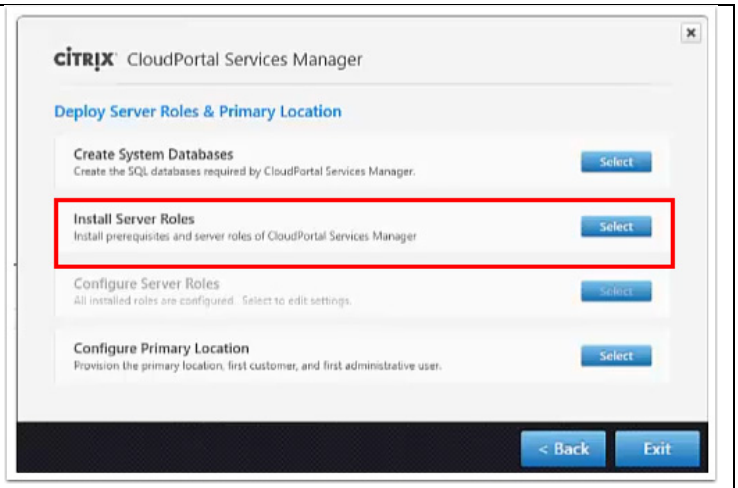
Select **Configuration Tool** and **Web** as the Configuration Tasks to be performed on this server. Click **Next** to continue.	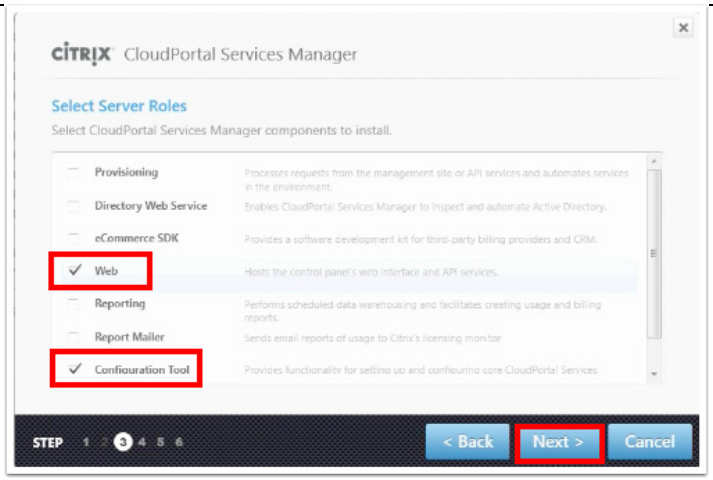
Click **Configure Server Roles** on the second CPSM server.	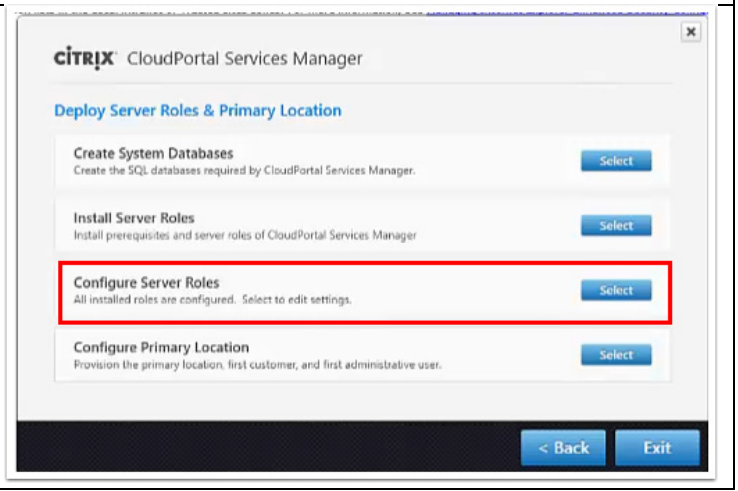

Specify the deployment configuration file created previously. Preview the **Service Package Import** page and click **Next** to continue.	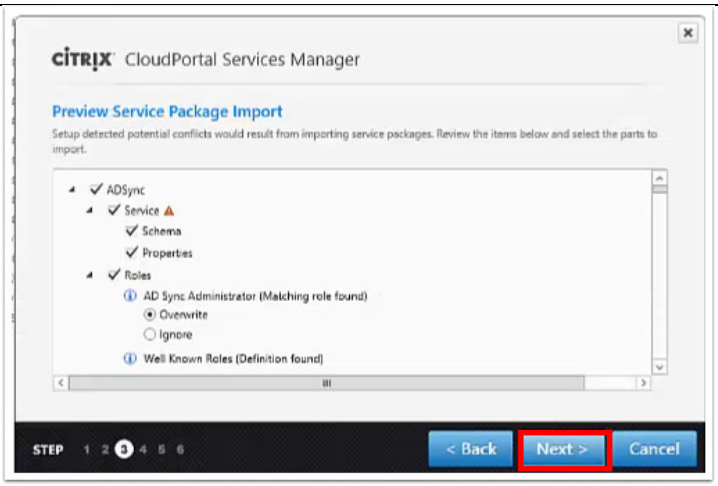
Follow the prompts to configure the Web Server. Click **Next** to continue.	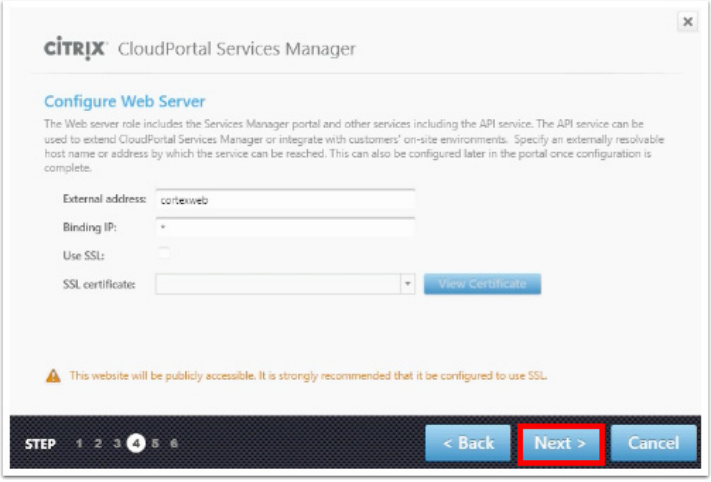
When the configuration is complete, click Finish. Restart the Queue Monitor Service on the CPSM provisioning server.	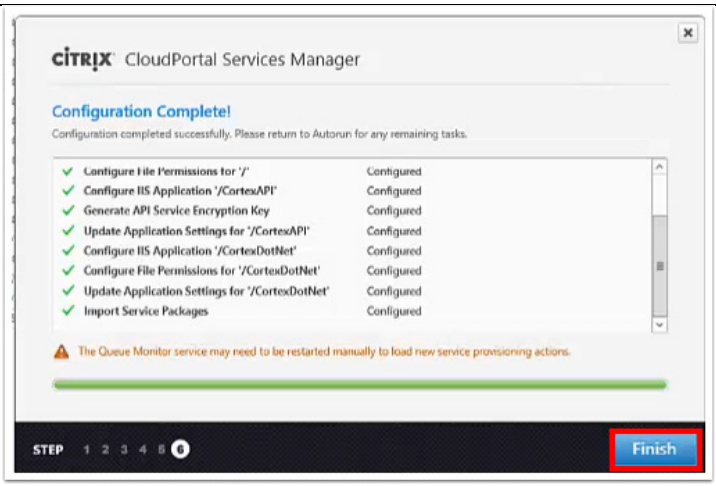

Click **Configure Primary Location.**	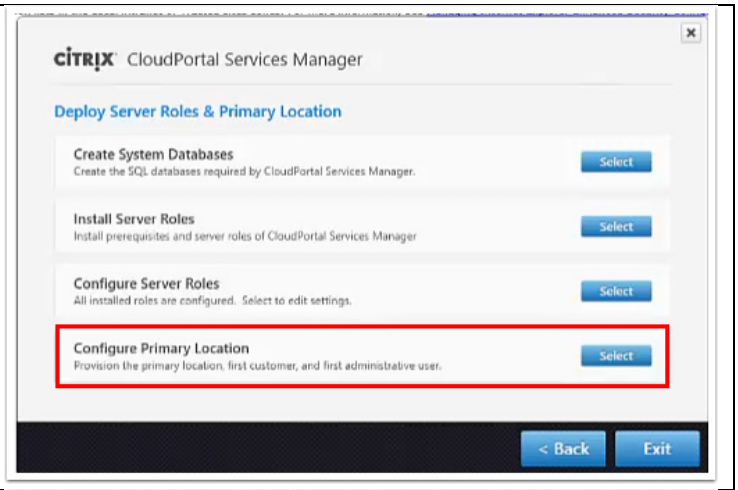
Specify the deployment configuration file created earlier. Click **Next** to continue.	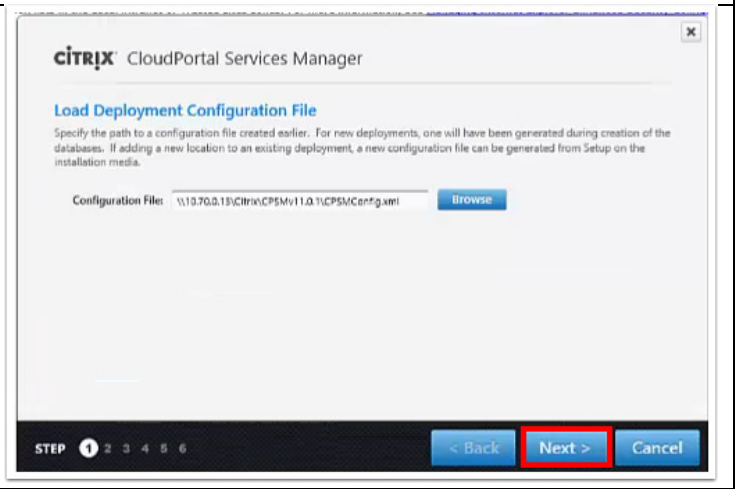
Configure the CloudPortal Services Manager primary location, which corresponds to an Active Directory domain or forest.. Click **Next** to continue.	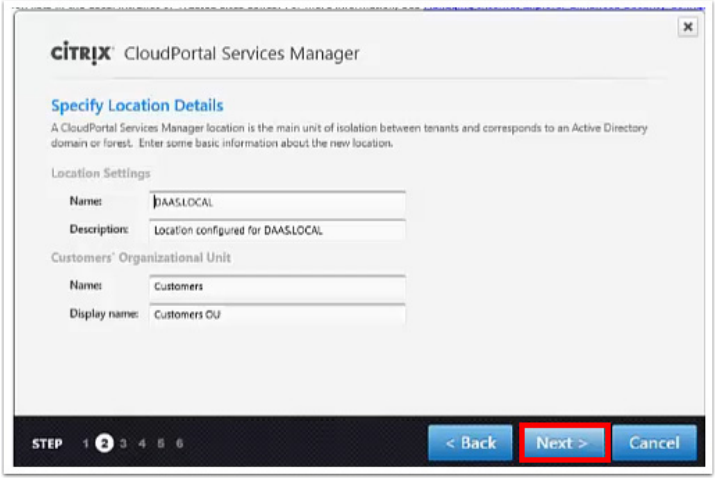

Enter details for the Service Provider. Click **Next** to continue.	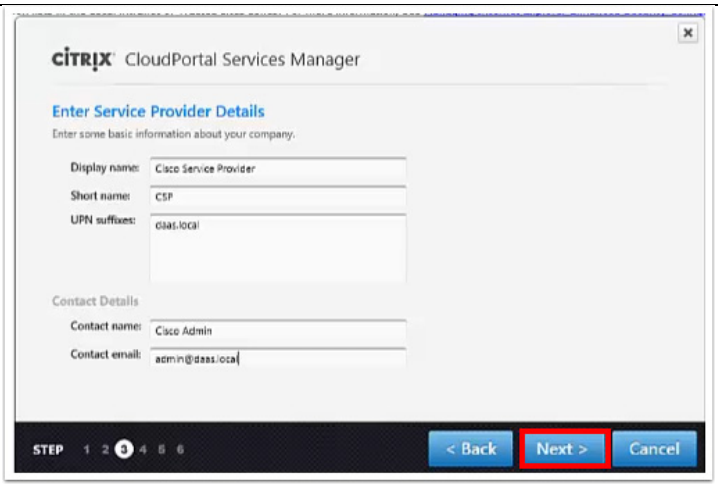
Create the First Administrator (cspadmin). Click **Next** to continue.	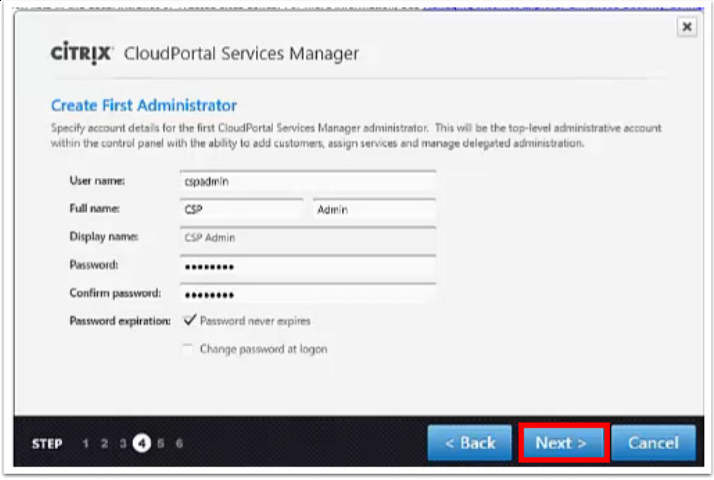
Review the summary and click **Commit**.	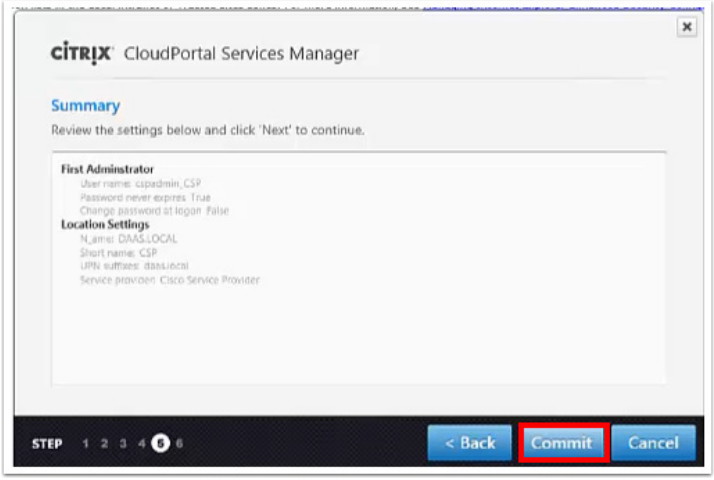

Allow Internet Explorer content from http://cortexweb .	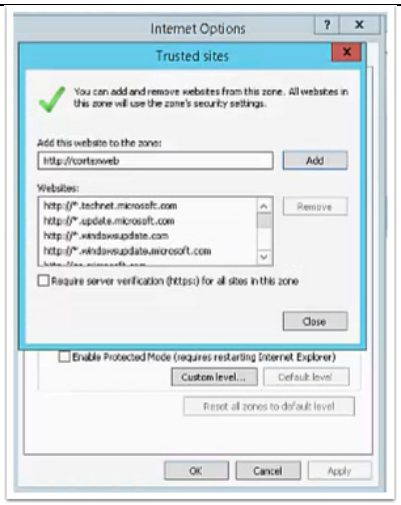
The configuration of the server roles and location for CloudPortal Services Manager is complete. Click **Finish**.	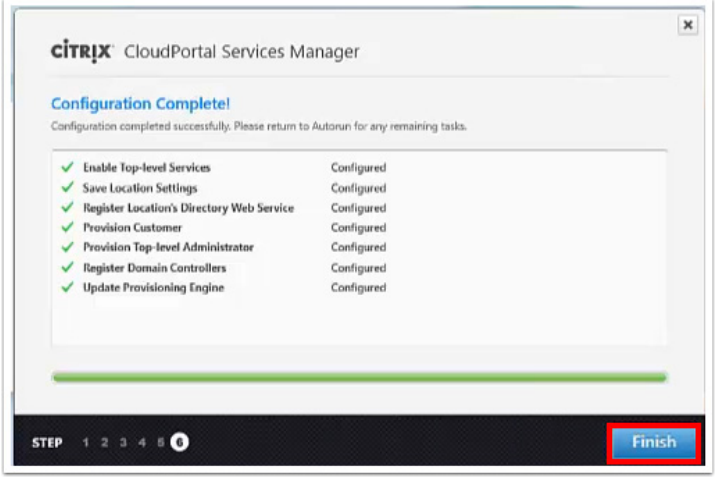

Setting Up the Shared Tenant Infrastructure

This section sets up the shared tenant infrastructure in the CVD. First, CloudPortal Services Manager is used to create and configure a new customer (named “Install Test Customer”) and import the new version of the Hosted Apps and Desktops service schema. Then, Citrix App Orchestration is used to create a new shared Delivery Site (named “SharedSite”) and a new StoreFront Server Group (named “SharedSFGroup”).

Create and Manage New Customer

CloudPortal Services Manager is used to create and manage a new customer (named “Install Test Customer”). After creating the customer, an administrator for this customer is created and the security roles and service deployment settings are configured.

Instructions	Visual

Log into CloudPortal Services Manager as the configured SP administrator (cspadmin_CSP). From the CloudPortal Services Manager **Customer** menu, select **New Customer**.	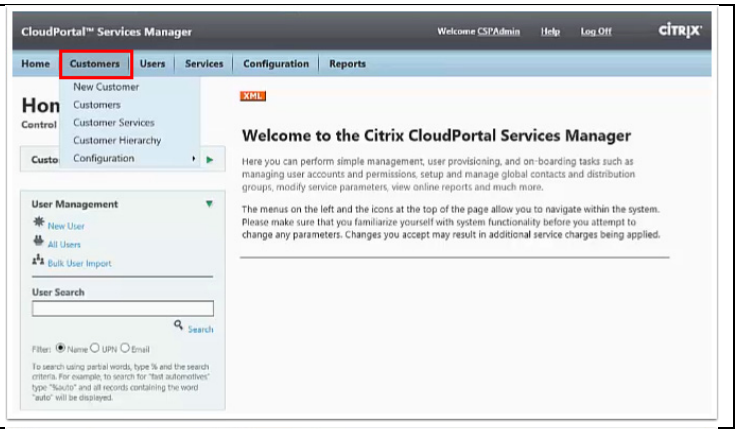
On the **Create Customer** page, complete the fields to add “Install Test Customer”**.** Click **Provision**.	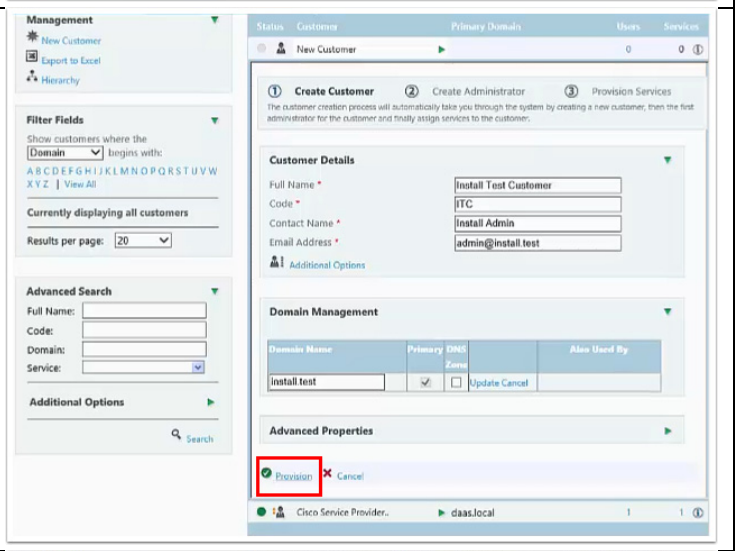
On the **Create Administrator** page, complete the fields and click **Provision**.	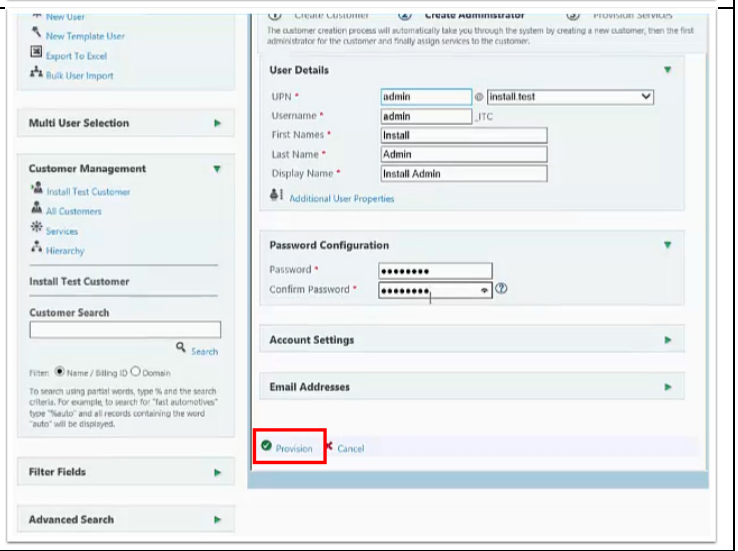

The new customer is now visible in the Active Directory hierarchy.	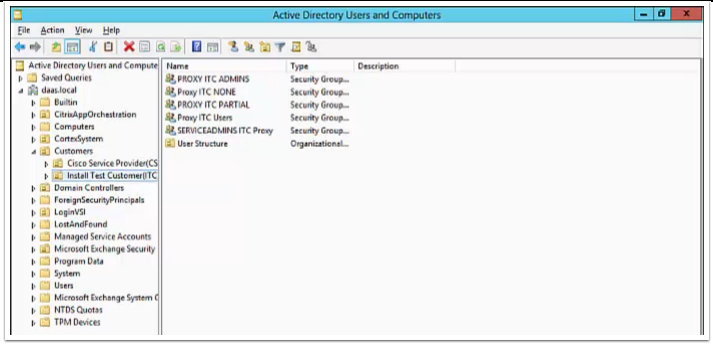
Next, configure the service deployment settings. In CloudPortal Service Manager, select **Configuration->System Manager->Service Deployment**. On the **Service Deployment** page, filter by **Active Directory Location Services**, and select **Customer Portal Settings**.	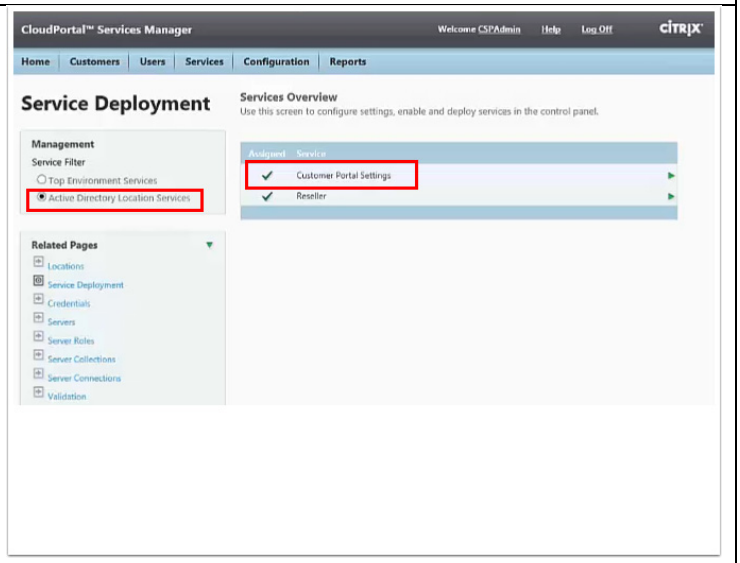

Select **Service Settings**. On the **Configure Service Settings** page, select the checkbox next to the **DNS Templates** property to enable editing, clear the default "tick" of the property, and then scroll down and click **Apply Changes**.	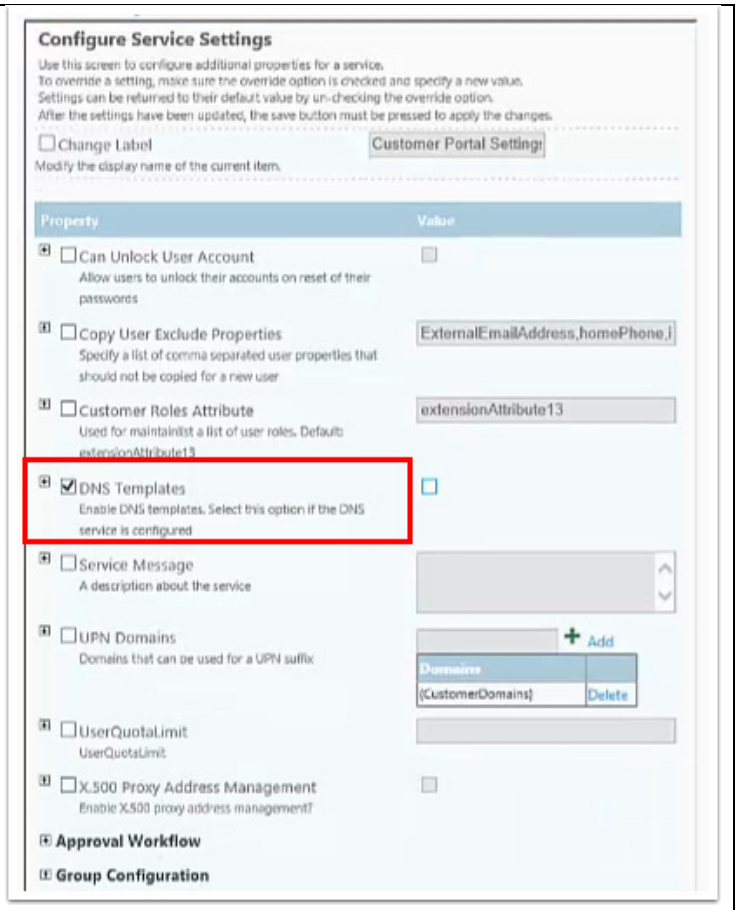
Click **Save**.	

Next, configure security roles. In CloudPortal Service Manager, select **Configuration->Security->Security Roles**. On the **Security Roles** page, select **Service Provider Administrator.**	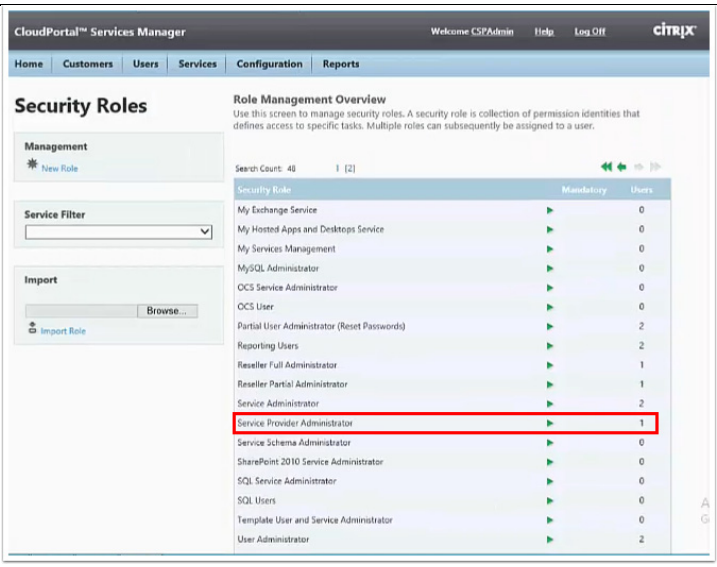
On the **Role Permissions**, select the **Menus** tab.	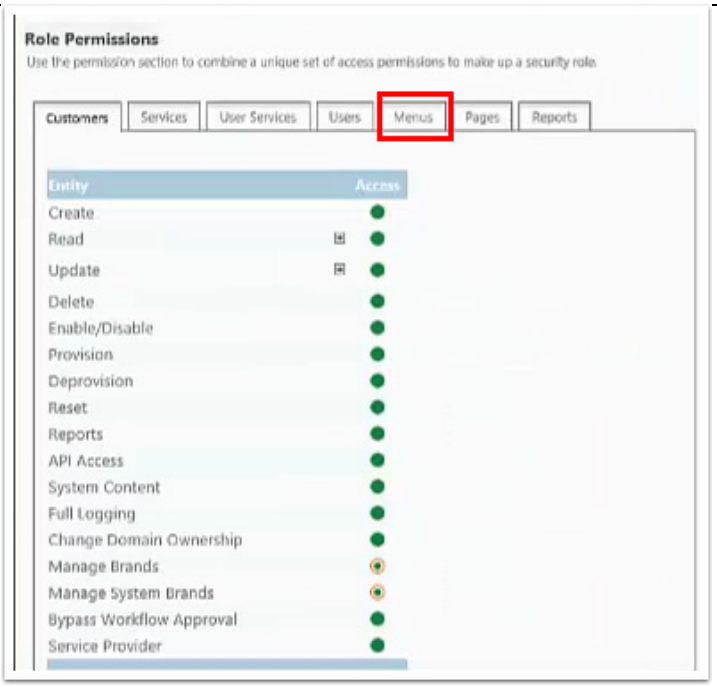

Expand the Services entry, and disable the checkboxes for the following entries: · Citrix · CRM 2011 · DNS · File Sharing Manager · MySQL · SharePoint 2010 · Virtual Machine · Windows Web Hosting [Only the checkbox for **Hosted Apps and Desktops** is enabled.] Click **Save**.	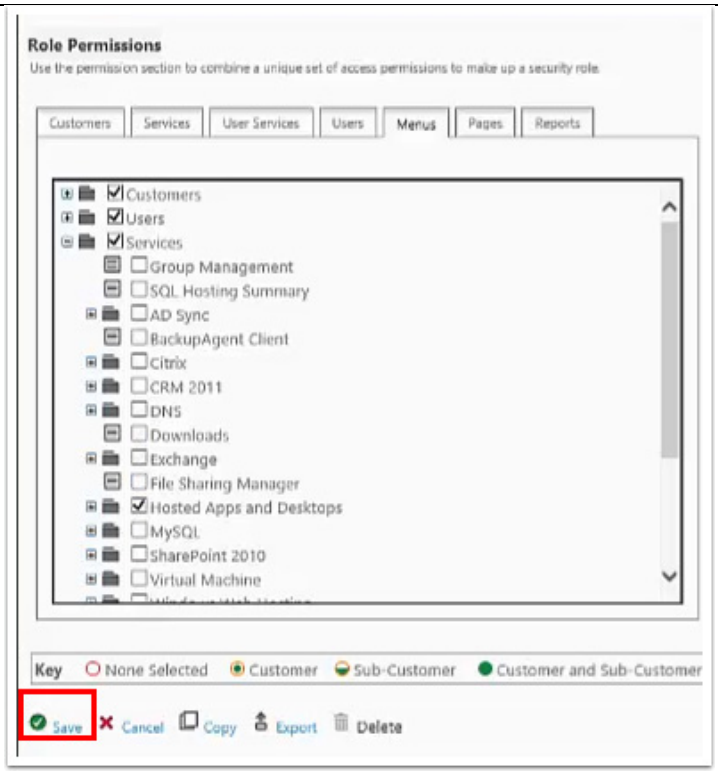
From CloudPortal Services Manager, select the **Users** tab. Expand the default administrator entry (cspadmin@daas.local) in the UPN column, and select **Edit User**.	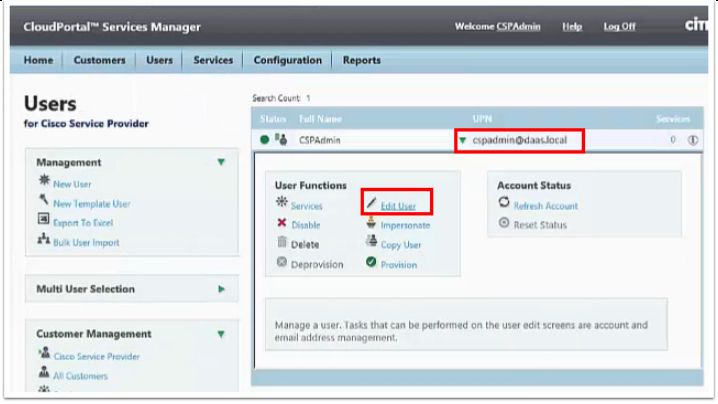

Expand **Account Settings**, and click **Advanced Options**.	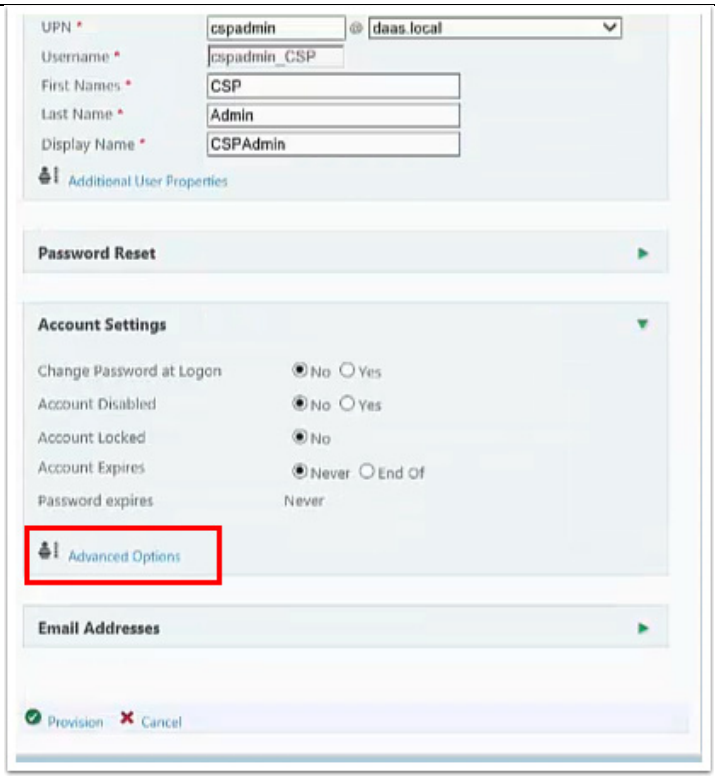
Check boxes for **All Services Schema Administrator** and **Service Schema Administrator**. Click **Provision**.	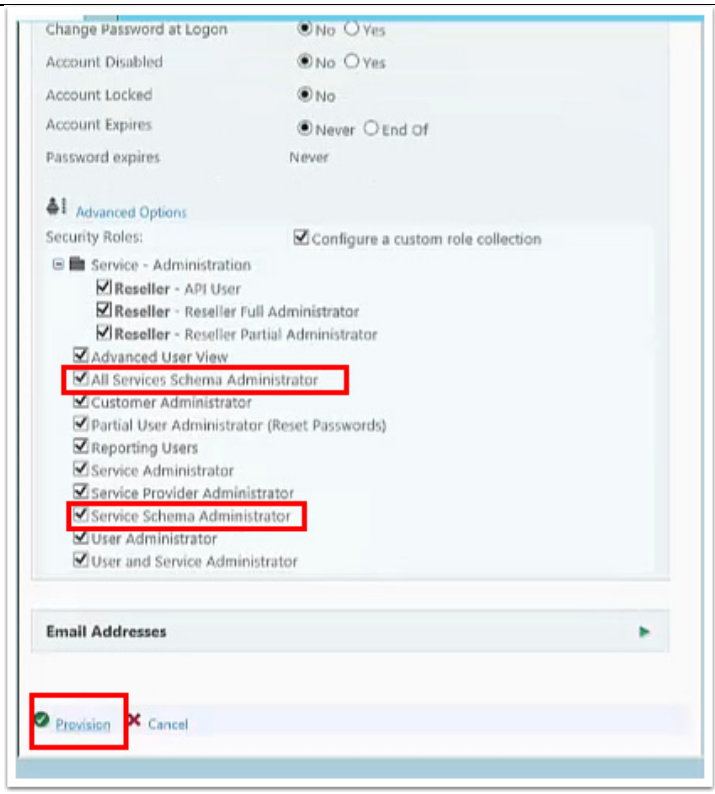

Import the Hosted Apps and Desktops (HaaD) Service

The system has an existing (v1) Hosted Apps and Desktops service by default. Delete this default service, and then import the new version (v2) of the Hosted Apps and Desktops service schema.

Instructions	Visual
First, delete the existing default Hosted Apps and Desktops service. Login to CloudPortal Services Manager as the configured SP administrator (cspadmin_CSP). Select **Configuration->System Manager ->Service Schema**. On the **Service Schema** page, click **Hosted Apps and Desktops**.	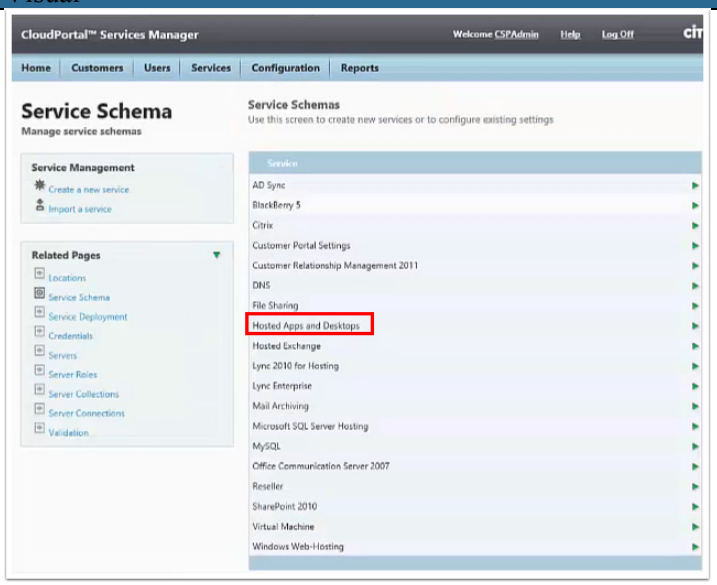
The Service Settings are displayed. Scroll down, and select **Delete**. When prompted, click **OK** to confirm the deletion.	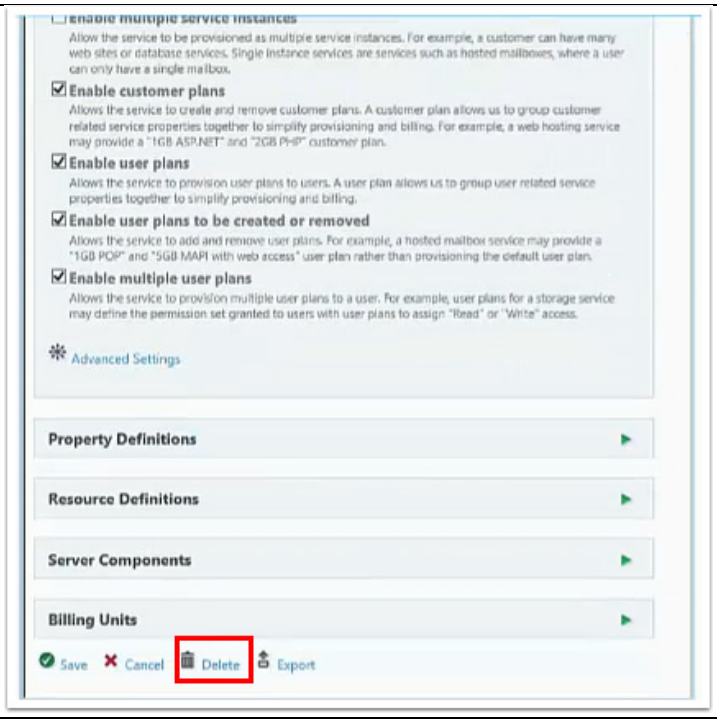

Next, import the new version of the Hosted Apps and Desktops service schema. From CloudPortal Services Manager, select **Configuration->System Manager->Service Schema**. On the **Service Schema** page (left), click **Import a service**.	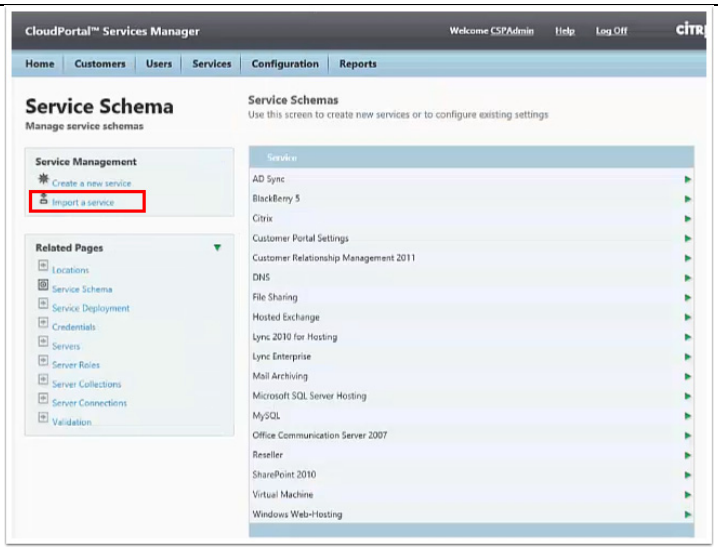
Browse, and select the import file (Hosted Apps and Desktops.package). Note: For CPSM v11.0.1, the package comes from the HaaD service release install media downloaded from Citrix.com.	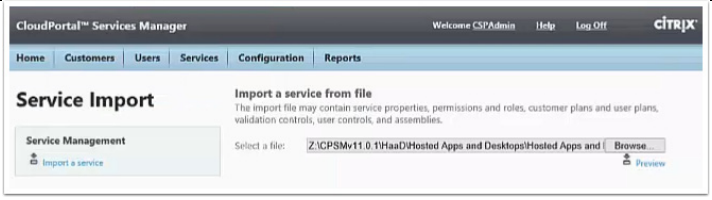
Review the information, and then click **Import**.	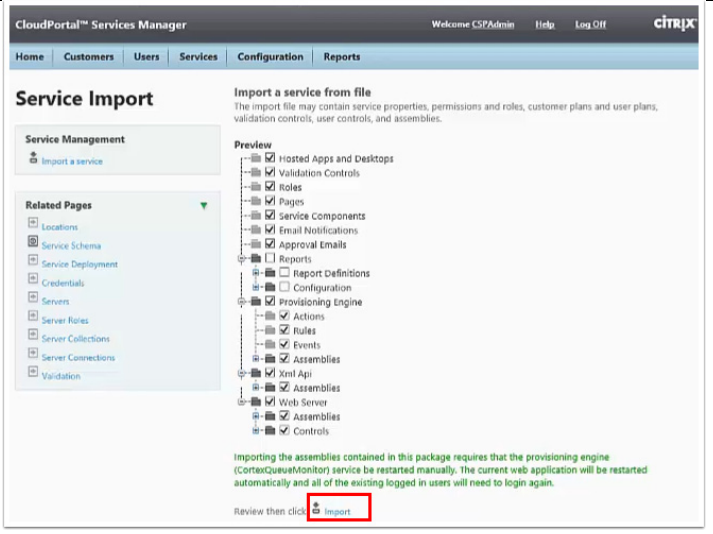

When the import operation completes, restart the Citrix Queue Monitor Service.	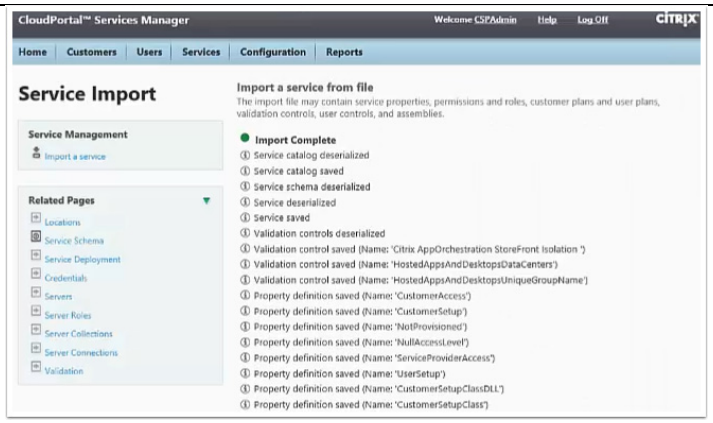

Create a new App Orchestration Delivery Site

This section illustrates adding a shared StoreFront server group. Before using Citrix App Orchestration to create a new StoreFront server group, the following prerequisites much be met:

- The .NET 3.5 Framework must be installed on Delivery Controllers (XDC01, XDC02)
- Internet Information Services Manager – import certificate for *.daas.local
- Active Directory user/groups – predefine user (SiteAdmin) and group (SiteAdmins) on Delivery Controller & SQL servers
- Microsoft SQL Server – add login for DAAS\SiteAdmins

Instructions	Visual
Run the Citrix App Orchestration Install Center on the Delivery Controller (e.g., DC01). Click on **XenApp and XenDesktop 7.6 Delivery Controller (and App Orchestration Agent)**. When prompted, click **Yes** to allow PowerShell script execution**.**	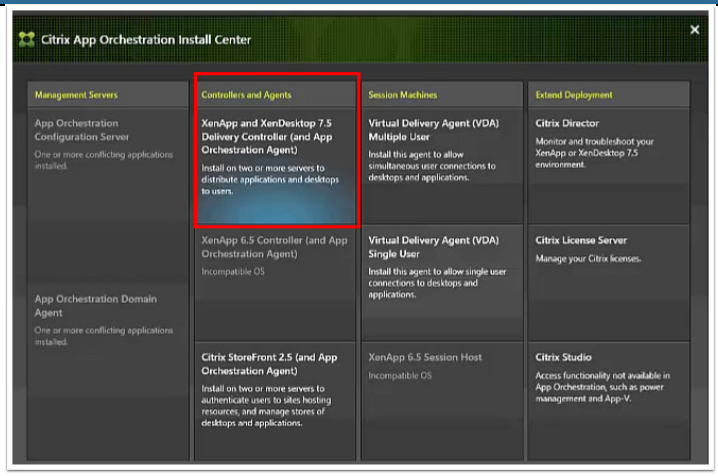

Installation messages are displayed. [Note: although the validation of the Citrix certificate failed in this example, the installation was successful.] Click **Close** to continue. Restart computer if instructed to do so. After restart, rerun App Orchestration Install Center and select 'XenApp and XenDesktop 7.5 Delivery Controller (and App Orchestration Agent)' to finish the operation.	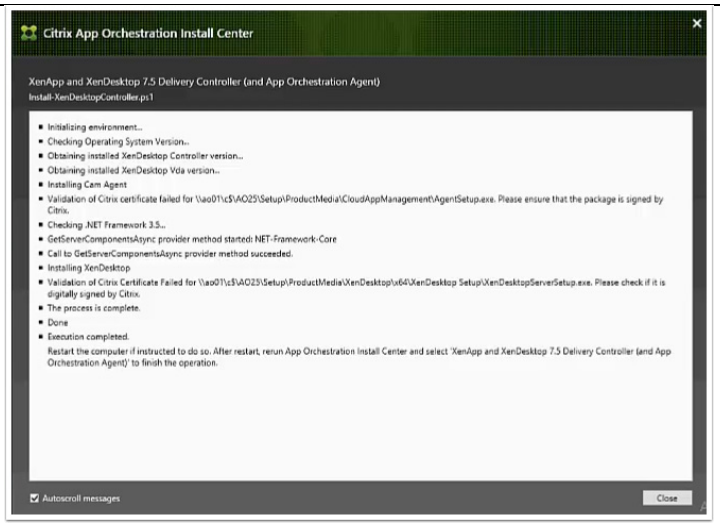
Next, create a new shared Delivery Site. From the Citrix App Orchestration web console, select **Design** tab. Then select **Create new** from the **Manage** menu.	
On the **Basic Settings** page, fill in the name of the new Delivery Site. In this CVD, the name "SharedSite" is used. Click **Next**.	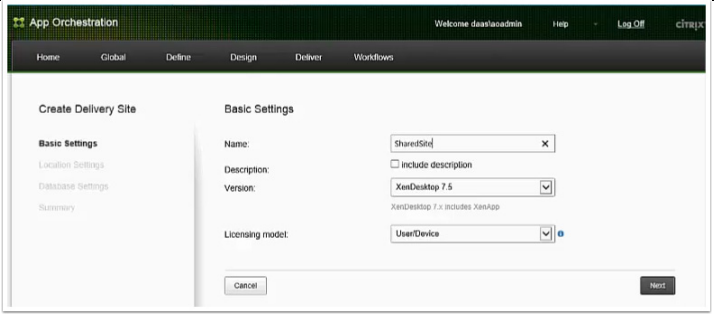
On the **Location Settings** page, fill in the fields. In this CVD, the following information was used: · Machine name: xcd01.daas.local, xcd02.daas.local · Domain: daas.local · Delivery site admin group: SiteAdmins · User name: SiteAdmin Click **Next**.	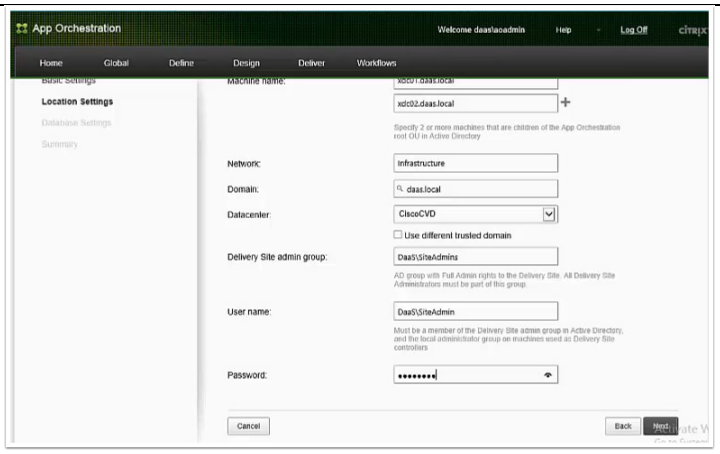

On the **Database Settings** page, fill in the fields. In this CVD, the following information was used: · Database server: SQLHA · Database name: SharedSite · User name: SiteAdmin Click **Next**.	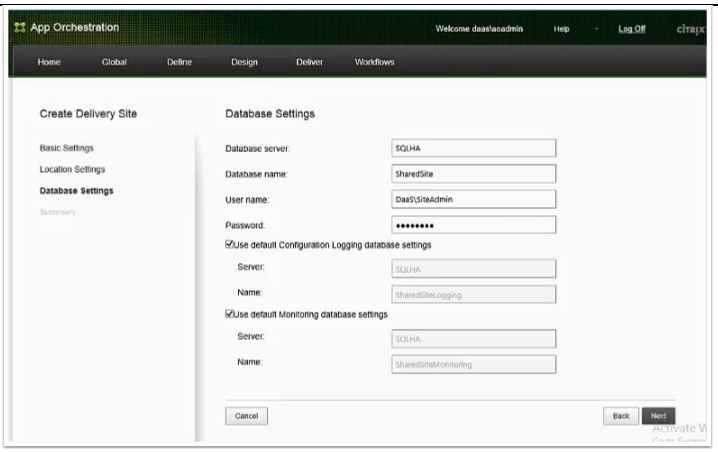
A summary is displayed. Click **Save**.	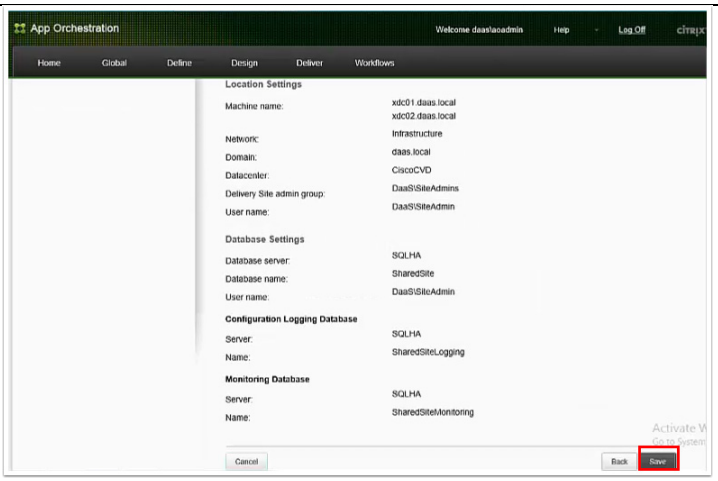
When the workflows complete, the new shared Delivery Site is displayed.	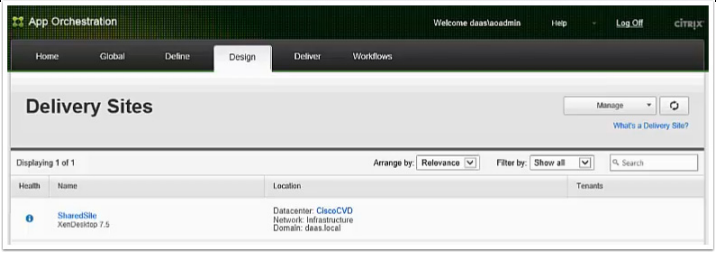

Create a StoreFront Server Group

This section illustrates adding a shared StoreFront server group. Citrix App Orchestration Install Center is used to install Citrix StoreFront 2.5, and then the Citrix App Orchestration web console is used to create the shared StoreFront server group (named “SharedSFGroup” in this CVD).

Before starting, a URL should be configured in DNS for the load balancer. The actual load balancer virtual IP address can be configured later in NetScaler. If at the time of installation the NetScaler IP address is not available, the DNS entry for the load balancer URL can temporarily point to one of the StoreFront servers.

Instructions	Visual
These procedures should be performed on both StoreFront servers (DaaS-SF01 and DaaS-SF02). Run the Citrix App Orchestration Install Center on the StoreFront server. Select **Citrix StoreFront 2.5 (and App Orchestration Agent)**. When prompted, click **Yes** to allow PowerShell script execution**.**	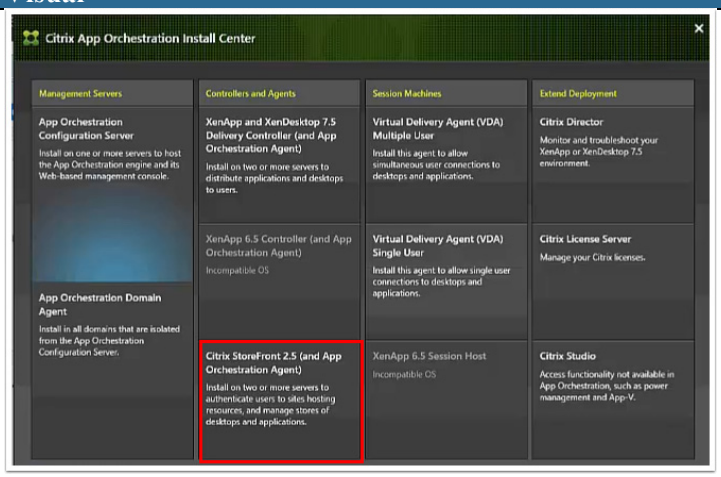
Installation messages are displayed. [Note: although the validation of the Citrix certificate failed in this example, the installation was successful.] Click **Close** to continue. Restart computer if instructed to do so. After restart, rerun App Orchestration Install Center and select 'XenApp and XenDesktop 7.5 Delivery Controller (and App Orchestration Agent)' to finish the operation.	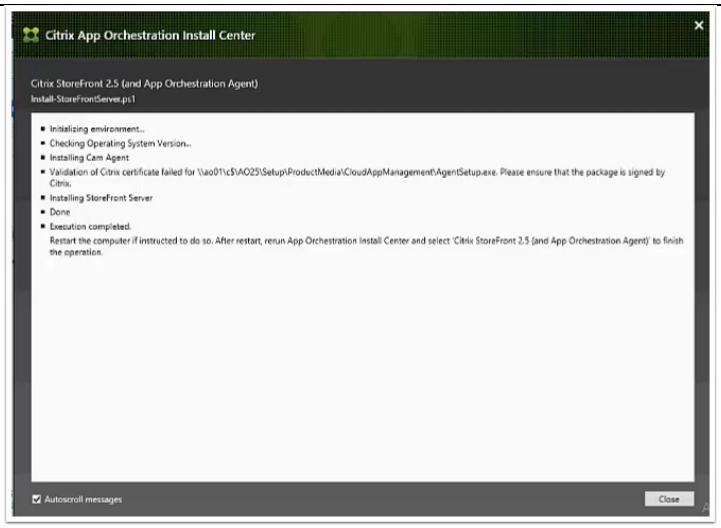
In Citrix App Orchestration web console, click the **Design** tab. Then click **Add StoreFront Server Group**.	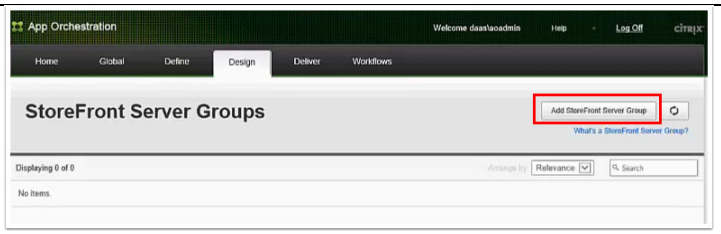
On the **Basic Settings** page, fill in the fields. In this CVD, the following information was used: · Name: SharedSFGroup · SSL cert: *.daas.local · Load balancer: https://sf-lb.daas.local Click **Next**.	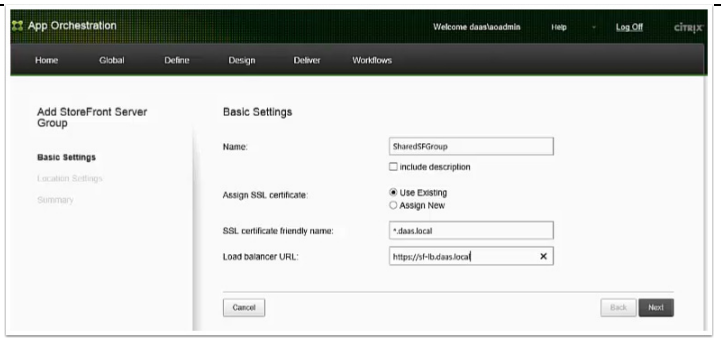

On the **Location Settings** page, fill in the fields. In this CVD, the following information was used: · Machine names: daas-sf01.daas.local, daas-sf02.daas.local · Network: Infrastructure · Resource domain: daas.local · Datacenter: CiscoCVD Click **Next**.	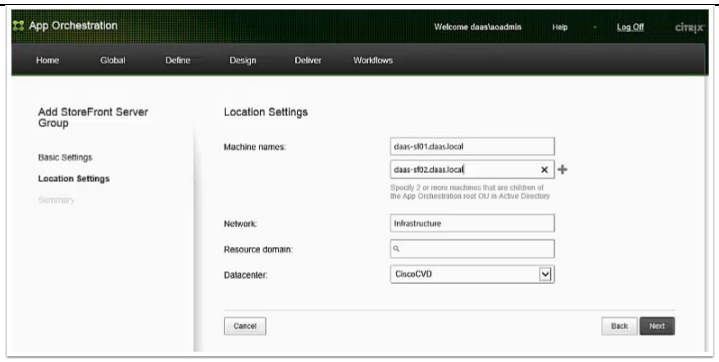
A summary is displayed. Click **Save**.	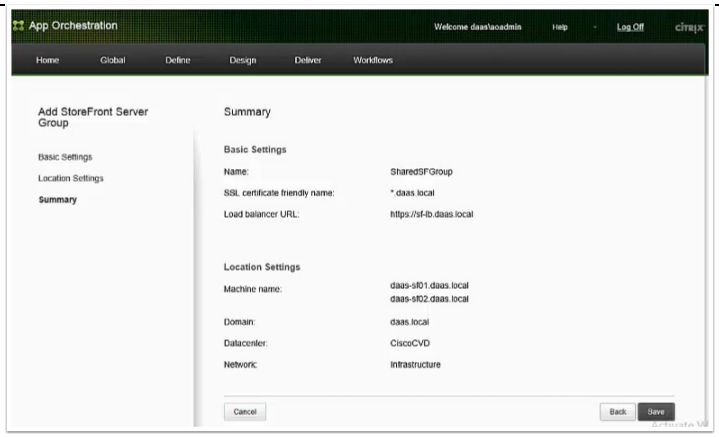
When the workflows complete, the new shared StoreFront Server Group is displayed.	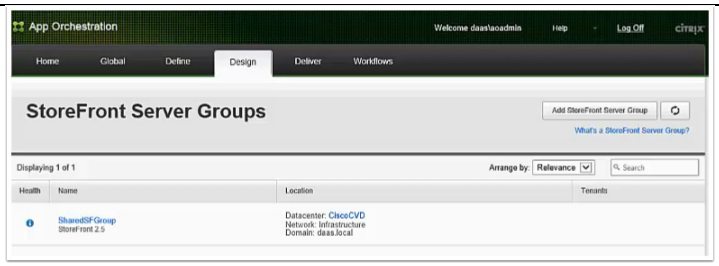
The Getting Started in App Orchestration status page now shows the StoreFront and Delivery Site design tasks now complete.	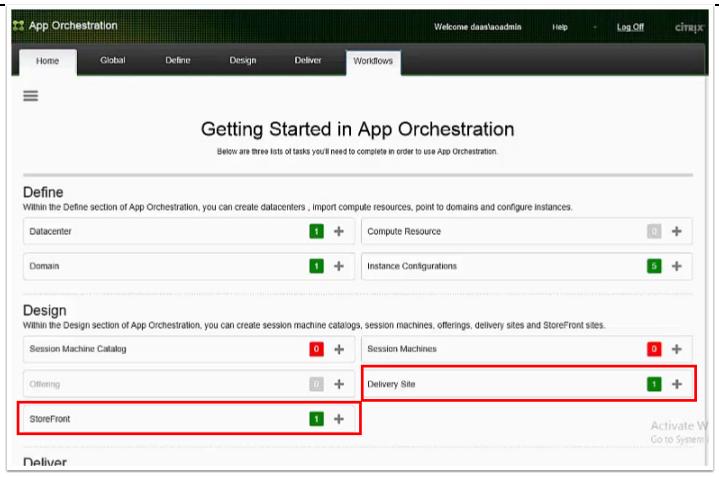

Configuring AO/CPSM Integration and Shared Delivery Site Offerings

CloudPortal Services Manager can extend an App Orchestration deployment to simplify customer and subscriber management, on-boarding, and self-service.

This section describes the software integration procedures for App Orchestration and CloudPortal Services Manager software and procedures to set up offerings for Shared Delivery Site tenants.

Configuring Hosted Apps and Desktops Service Roles and Connections

The Hosted Apps and Desktops web service for Services Manager allows service providers to manage and delegate end-user administration of applications, desktops, and resources. Service configuration typically involves these tasks, which are covered in this section:

- Create self-service account
- Add the servers associated with the service
- Enable the service at the top environment and location levels
- and specify service deployment with App Orchestration
- Assign service roles to the servers
- Add credentials for accessing the servers and management tools
- Add service connections for the Hosted Apps and Desktops service to establish communication between the servers and CloudPortal Services Manager
- Configure service settings

Check the box to install the **App Orchestration Configuration Tool**. Click **Next**.	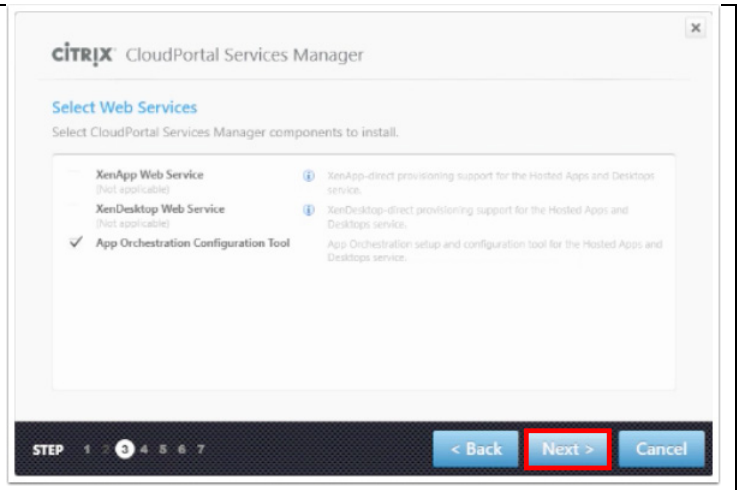

Review installation prerequisites and click Next. Review the **Ready to install** page. Click **Install**.	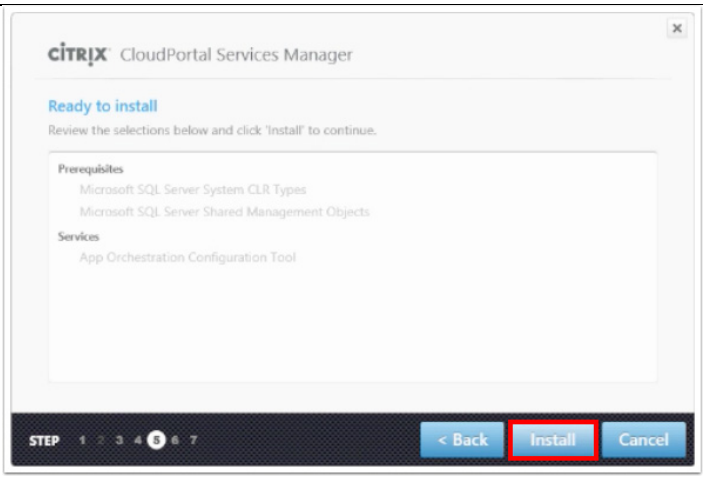
The **Deployment Complete** screen appears when the installation is complete. Click **Next**.	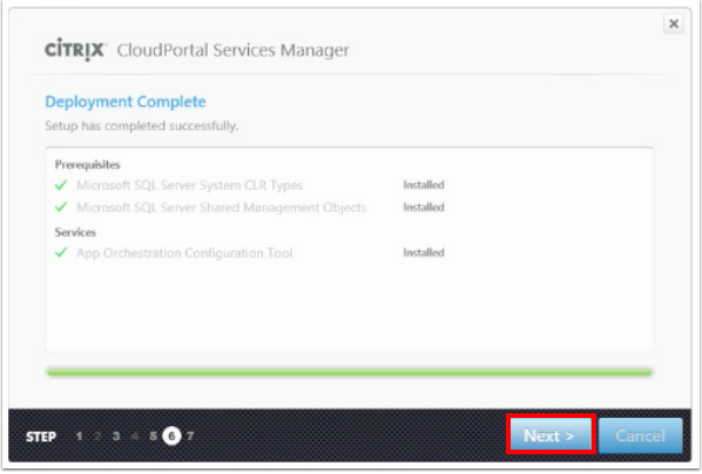
On the Configure Installed Components dialog, click **Configure** for the **App Orchestration Configuration Tool**.	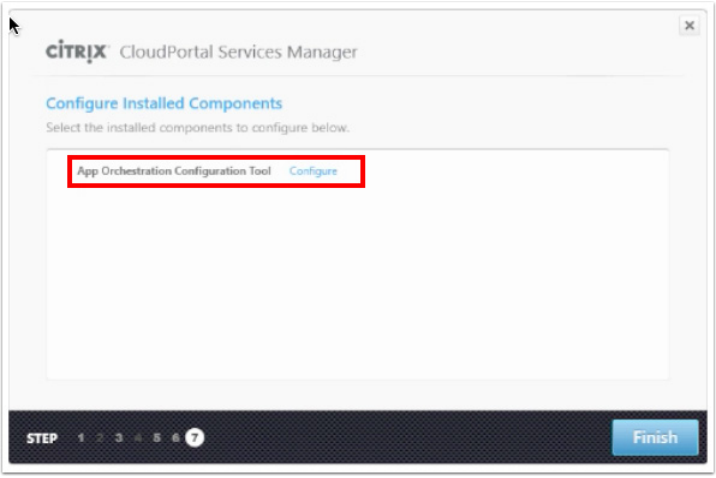

Enter the address of the App Orchestration Configuration Server (AO01.daas.local) and specify the administrative account with which to connect (DAAS\Administrator). Click **Test Connection** and if successful click **Next**.	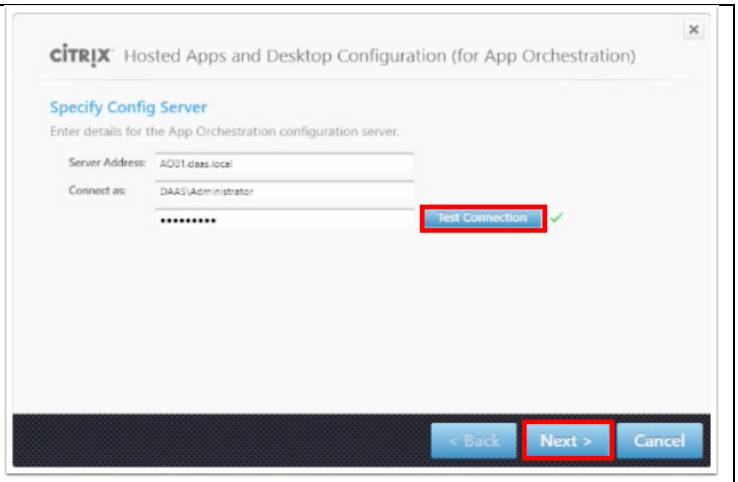
Create a service account (csm_haad_selfsvc) that will be used to enable self-service administration in App Orchestration. Click **Next**.	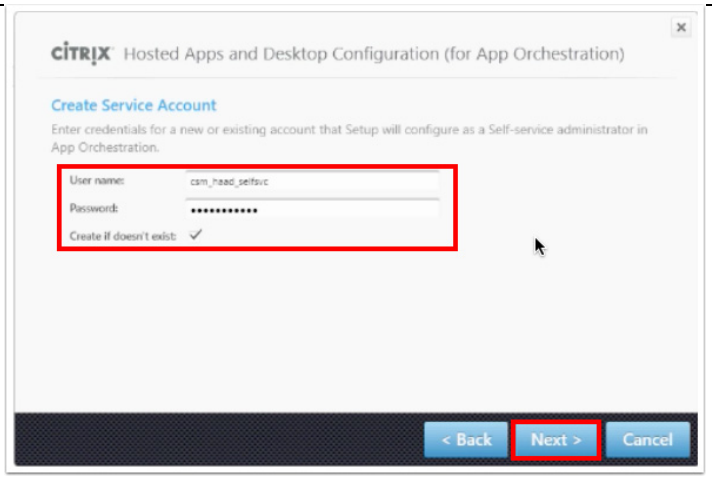
Specify the CloudPortal Services Manager database server (CORTEXSQL). Click **Next**. Review the configuration summary, **click Next**, and the configuration is applied.	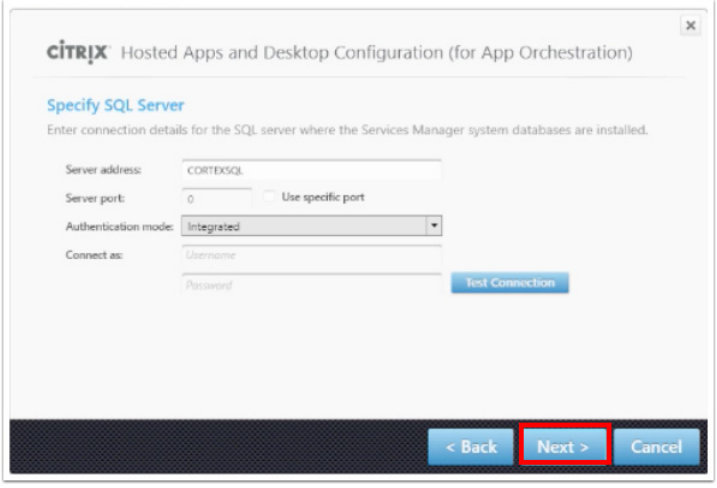

A notification screen displays when the configuration operations complete. Click **Finish**.	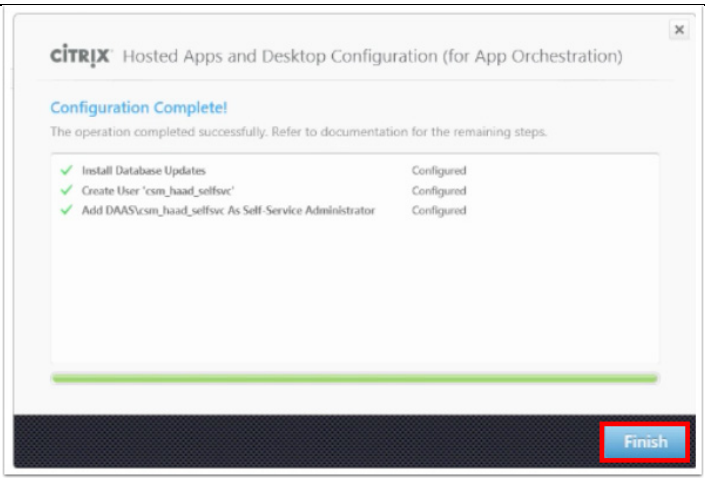
The next task is to add the servers associated with the Hosted Desktops and Applications service. On the CloudPortal Services Manager server (CPSM02), **login as the CSP Administrator** (cspadmin_csp) using the console.	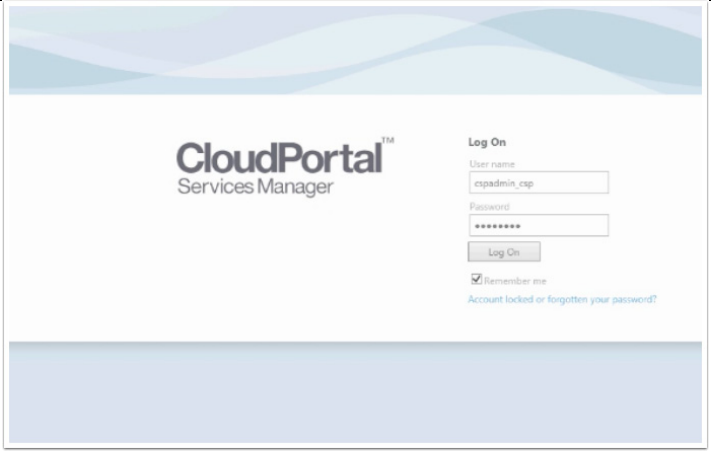
Select Servers from the Configuration -> System Manager menu. Select and expand the entry for the App Orchestration server. Enter a **server alias** with the FQDN (in this CVD, AO01.daas.local). Verify other App Orchestration server details and **Save**.	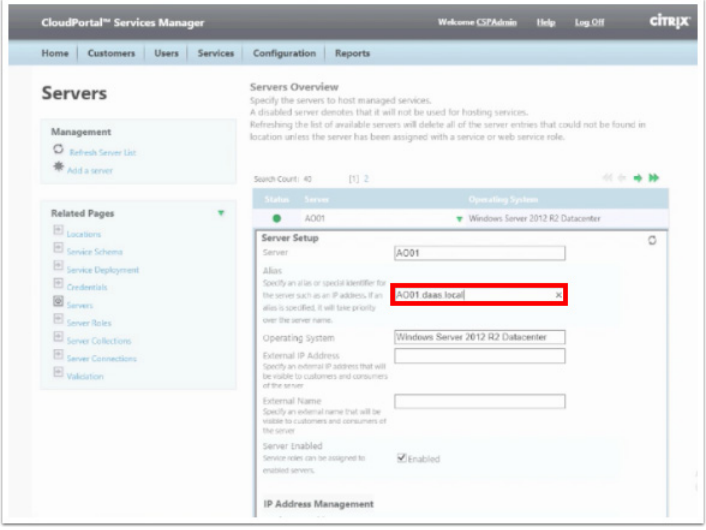

Verify the **Top Environment Service Configuration**. Review the Service Settings and Customer Plans for the Hosted Apps and Desktops service. Save any changes, if needed.	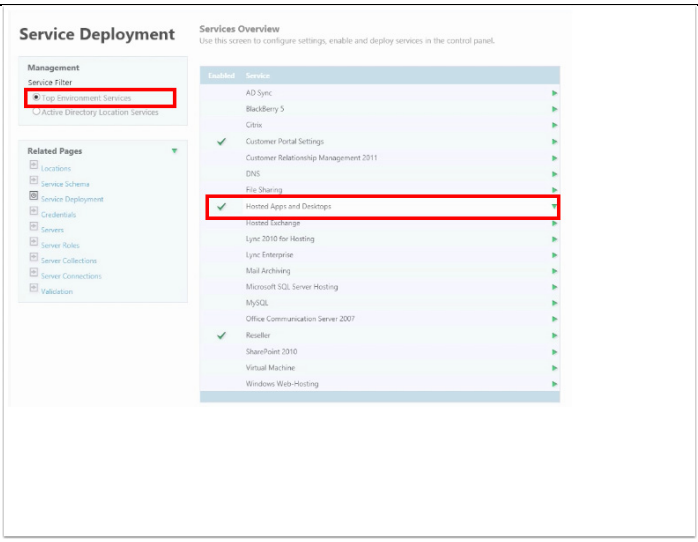
To configure the service at the location level, set the **Active Directory Location Services** as the service filter on the Service Deployment menu. Expand **Customer Plans** for the Hosted Apps and Desktops service. Unselect XenApp and XenDesktop so that **App Orchestration** is only selection. Click **Save**.	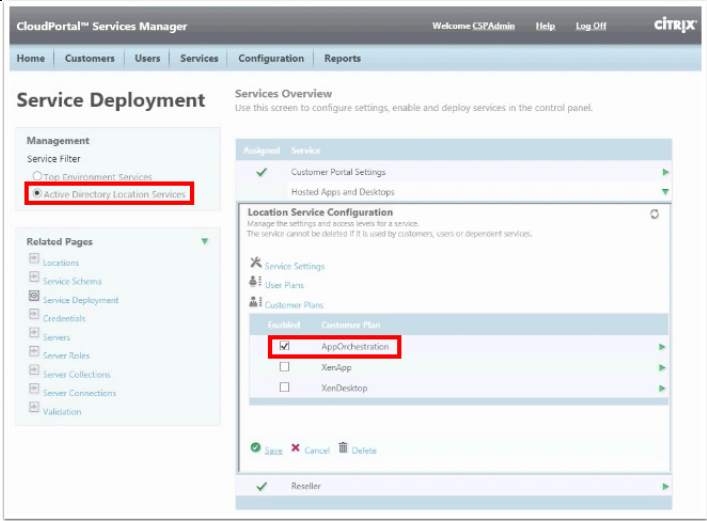
Using the **Server Roles** menu, assign the role of Hosted Apps and Desktops to the **App Orchestration** server (AO01). Click **Save**.	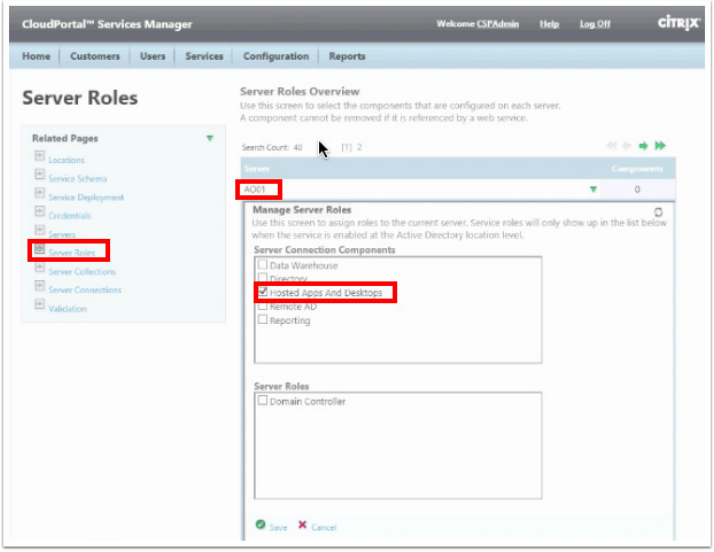

Using the **Credentials** menu, add DAAS\csm_haad_selfsvc and update.	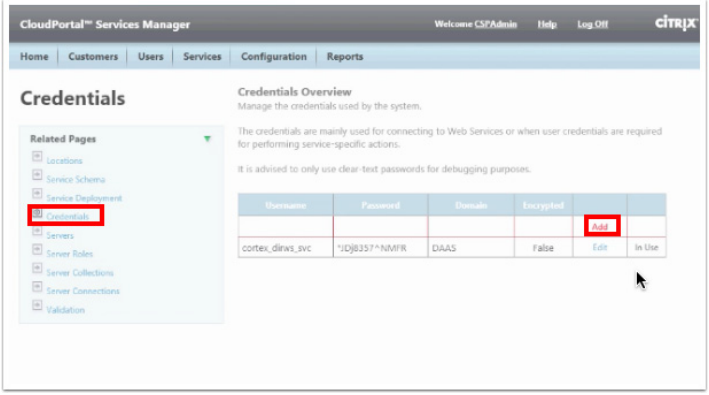
On the **Server Connections** page, enter the details under **Manage Server Connections**. Specify Hosted Apps and Desktops as the server role, the App Orchestration server (AO01), credentials, HTTPS as the protocol, and App Orchestration as the version. Click **Save**. Test the defined server connection by clicking on test connection icon.	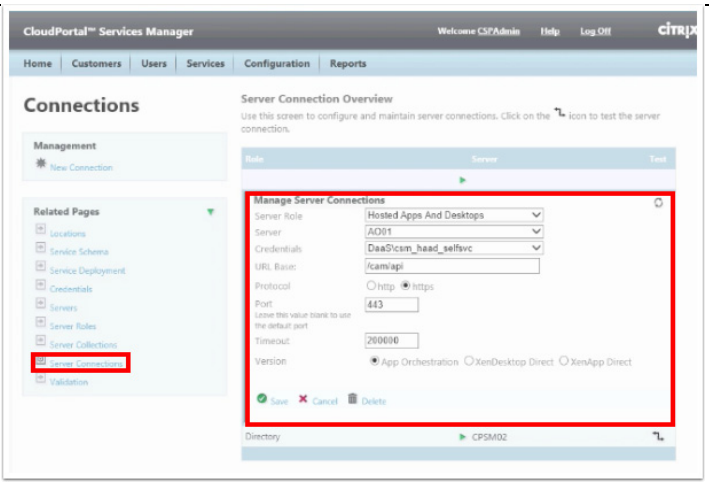
On the **Server Deployment** page, click on Active Directory Location Services as the service filter. Under the App Orchestration section, click on **Reload** for the App Orchestration Datacenter. Mark the **checkbox** for App Orchestration Datacenter and specify the **datacenter name** (CiscoCVD). Click **Apply changes**. Click **Save**.	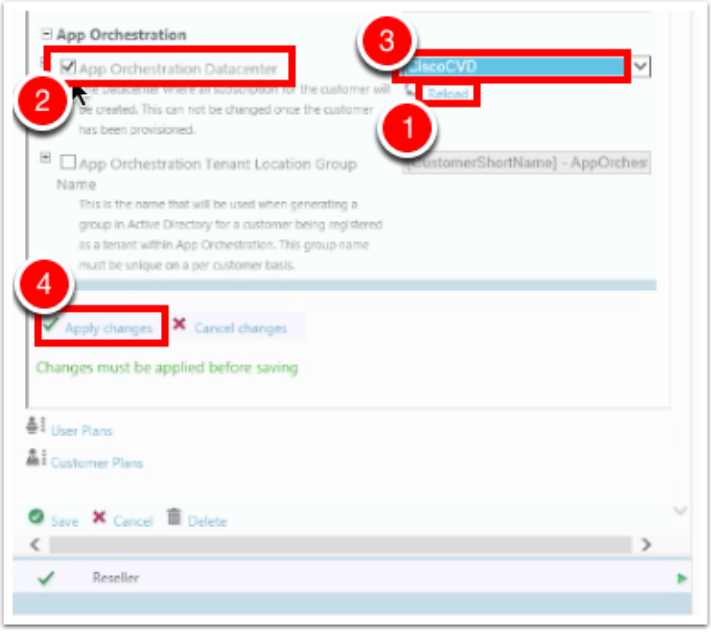

On the **Server Deployment** page under Customer Plans and Configure Service Settings, set up App Orchestration properties. Check the box to specify the **App Orchestration Private Network Name.** Enter the name (in this case, Infrastructure). Click **Apply changes**. Click **Save**.	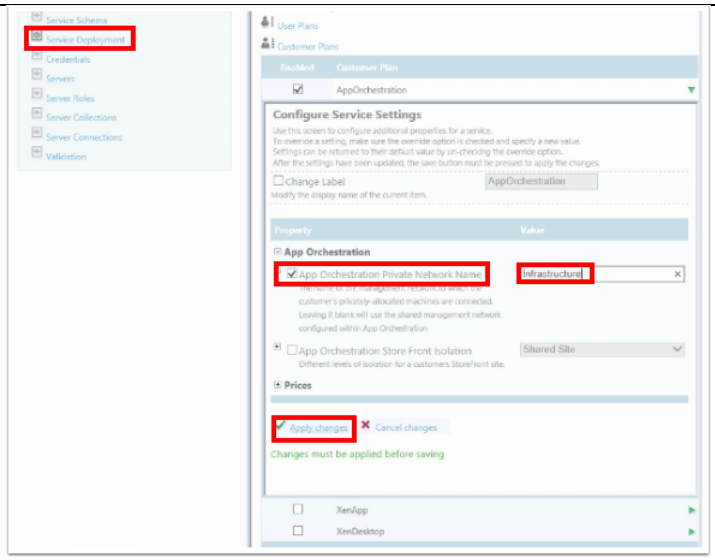

Configuring a Session Machine Catalog for Shared Delivery Site HSD Users

Session Machines host the applications and desktops that tenant users can access through Citrix Receiver. Session Machines are collected and organized in Session Machine Catalogs.

Configuring the Catalog

This section creates a Session Machine Catalog for the shared delivery site tenants that require Hosted Shared Desktop (HSD) services.

Instructions	Visual
Launch the App Orchestration Web Console and select **Session Machine Catalogs** from the Design tab. Select "**Created externally**" from the Use Machines pull-down. Specifying externally provisioned machines allows other means, such as Citrix PVS or PowerShell scripts, to provision servers for the catalog. When additional capacity is needed, App Orchestration can notify the administrator (they are not deployed automatically).	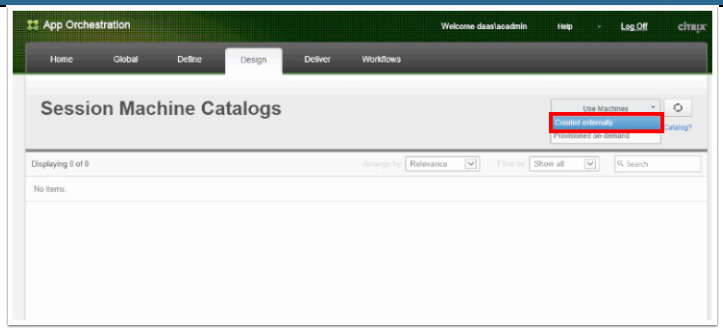

Configure Basic Settings for the Machine Catalog. To define a catalog for machines that support Hosted Shared Desktops, name the catalog **HSD** and check the radio button for a **Multi-User** OS. Click **Next**.	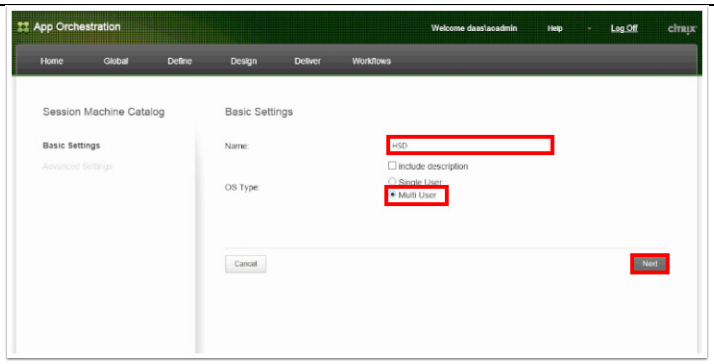
Review and modify the Advanced Settings, if necessary. Click **Save**.	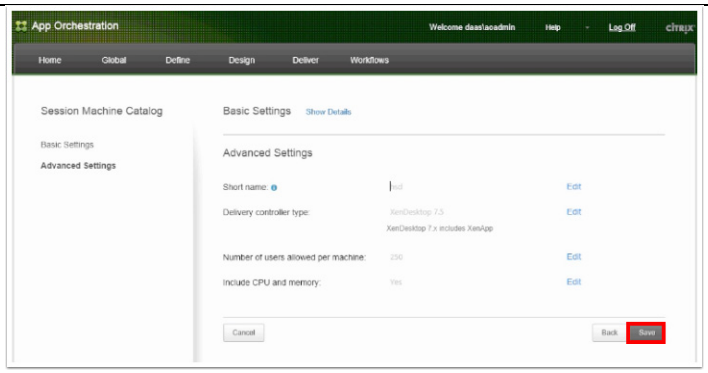

Adding Session Machines to the Catalog

This section adds Session Machines to a catalog. For this CVD, Cisco UCS Director initially created virtual machines that were imported into PVS prior to the configuration of App Orchestration and CloudPortal Services Manager.

Instructions	Visual
To add Session Machines to a Session Machine Catalog, first power on the VMs that will support hosted applications and desktops.	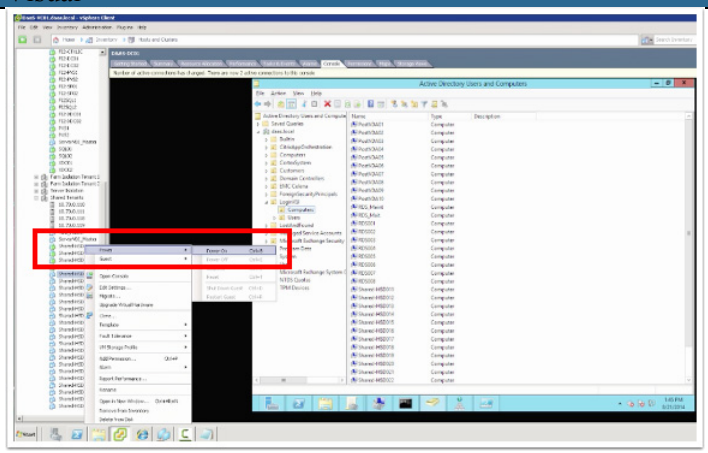

Launch the App Orchestration Web Console. On the Session Machine Catalogs page, select a **Session Machine Catalog** (the screenshot shows a catalog named CSP-SharedHSD.)	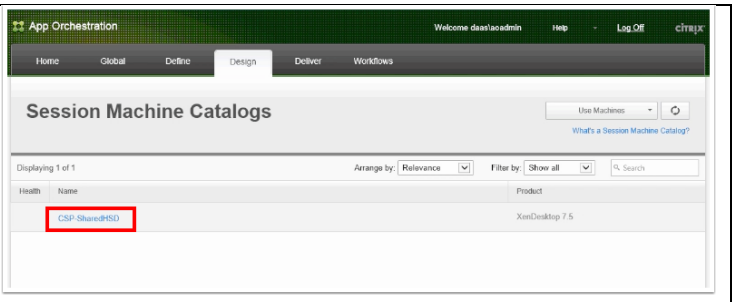
Click **Add Machines**.	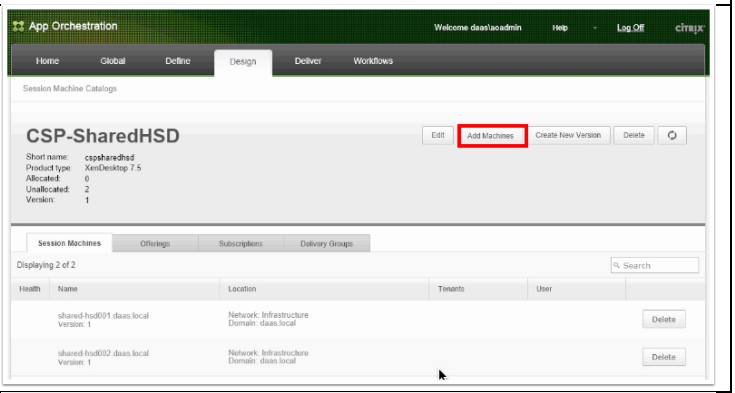
Specify **Basic Settings** (e.g., VM names, network, domain, and datacenter). Click **Save** to add the listed machines. Checking the Workflows page will show background tasks as the machines are added to the catalog (e.g., CreateResourceand Get-MachineInfo operations).	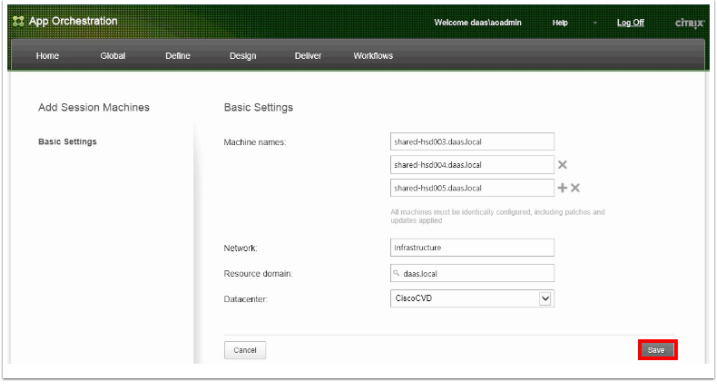
When the operation is complete, the Session Machine Catalog will list the available machines.	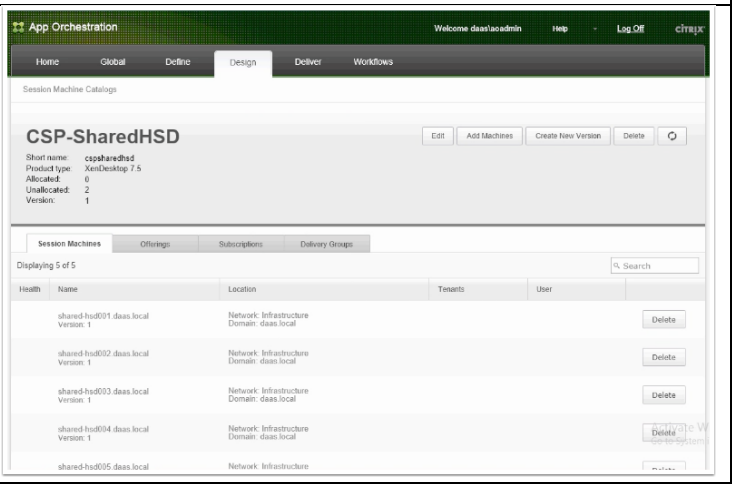

Setting up CPSM Reporting on Secondary SQL Server

Reporting for CloudPortal Services Manager delivers usage and billing reports to your customers and application vendors. It includes standard reports to support standard provisioned services and a data warehouse. The Reporting service communicates directly with the SQL Server Reporting Services web service.

For a large deployment, separate SQL servers are typically used for hosting the reporting and billing databases. CloudPortal Services Manager Report Mailer gathers anonymous usage data and emails usage reports to the Citrix license monitor.

Configure Reporting Service Prerequisites on the CPSM SQL Server

The following steps were performed on the secondary SQL server (SQL02) in this CVD to modify the Report Server configuration file . These steps are part of the Reporting (Data Warehouse) server prerequisites (see http://support.citrix.com/proddocs/topic/ccps-11/ccp-10-system-requirements-roles.html.)

Instructions	Visual
On the SQL02 server: Use the SQL Server 2012 Reporting Services Configuration Manager wizard to step through the configuration of the reporting database. These steps are part of the Reporting (Data Warehouse) Server pre-requisites: http://support.citrix.com/proddocs/topic/ccps-11/ccp-10-system-requirements-roles.html	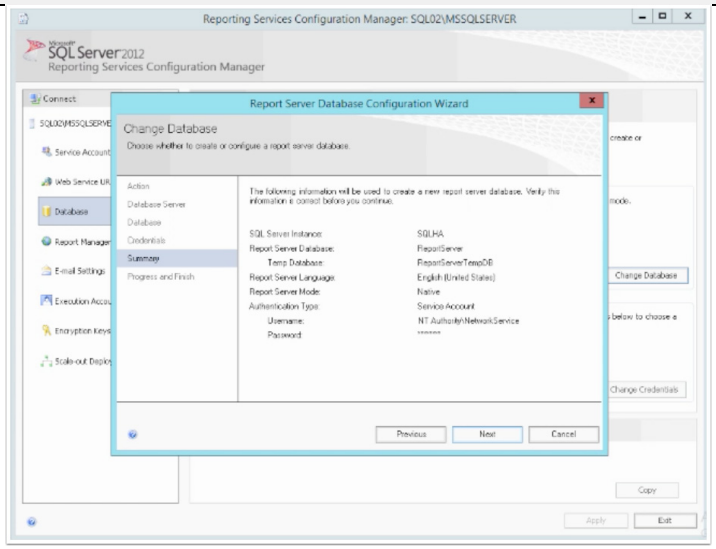
Edit the `rsreportserver.config` file and add `<RSWindowsNegotiate/>` as AuthenticationType. Adding the tag `<RSWindowsNegotiate/>` allows the CPSM Web portal to call the reporting services web service.	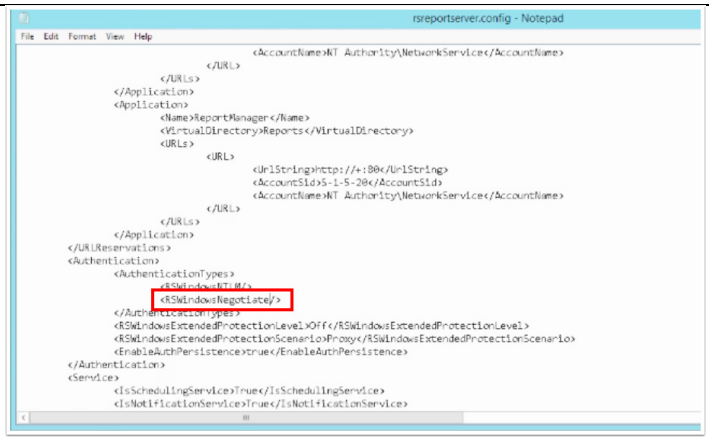

Install Server Roles for Reporting on SQL server

The following steps are performed on a CPSM server (CPSM02) to install server roles for reporting on a SQL server.

Instructions	Visual
Run **setup.exe** to launch the CloudPortal Services Manager console. From the Select Deployment Task dialog, select **Install CloudPortal Services Manager**. On the Install CloudPortal Services Manager dialog, select **Deploy Server Roles & Primary Location**. When prompted, read and accept the user license and click **Next**.	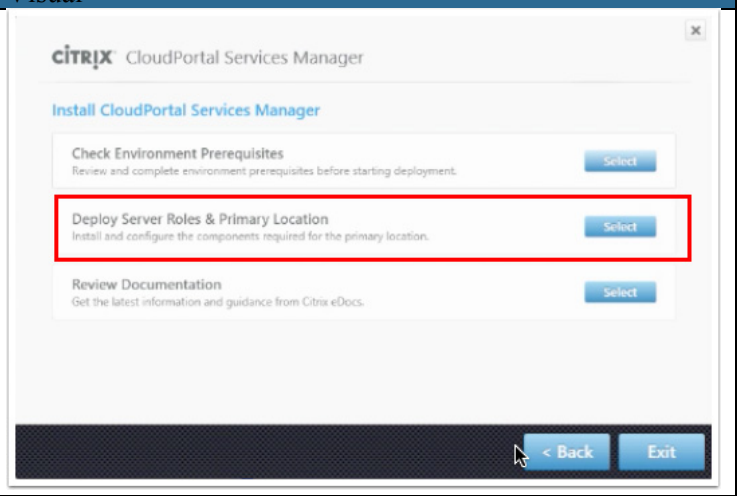
On the Select Server Roles dialog, enable the checkbox for **Reporting** (note: the Configuration Tool is already checked, and should be left enabled). Click **Next**.	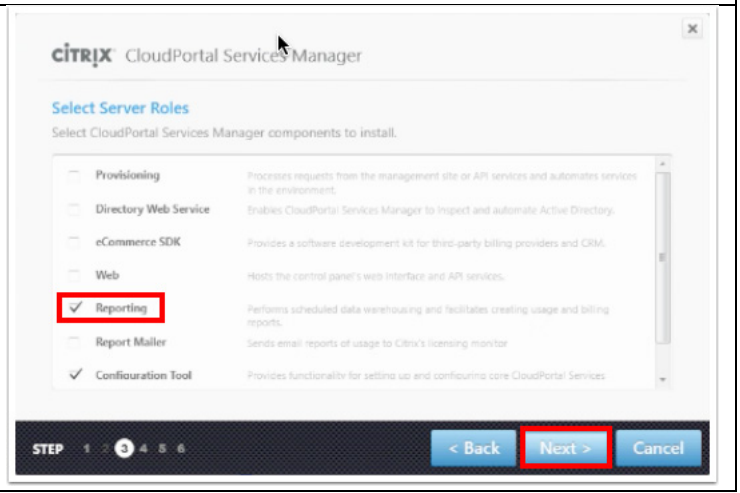

Review the prerequisites and then click **Install**. When the installation completes, click **Finish**.	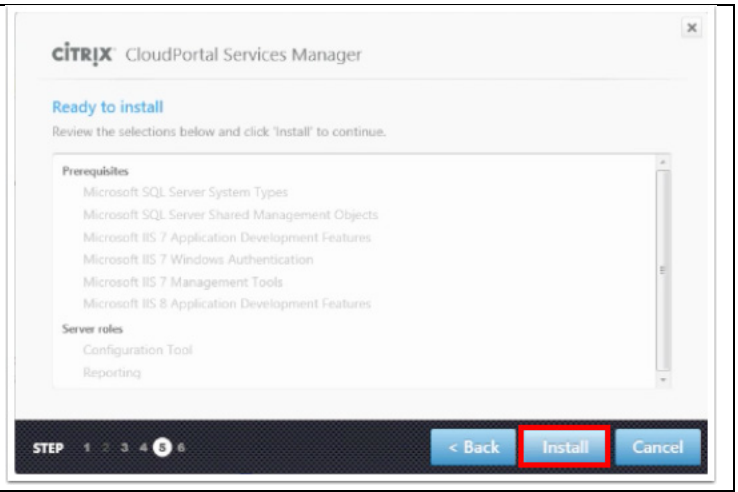

Configure Server Roles for Reporting on SQL Server

CloudPortal Services Manager uses a 12-step process to configure server roles and services for SQL server reporting.

Instructions	Visual
(1) From Deploy Server Roles & Primary Location dialog, select **Configure Server Roles**.	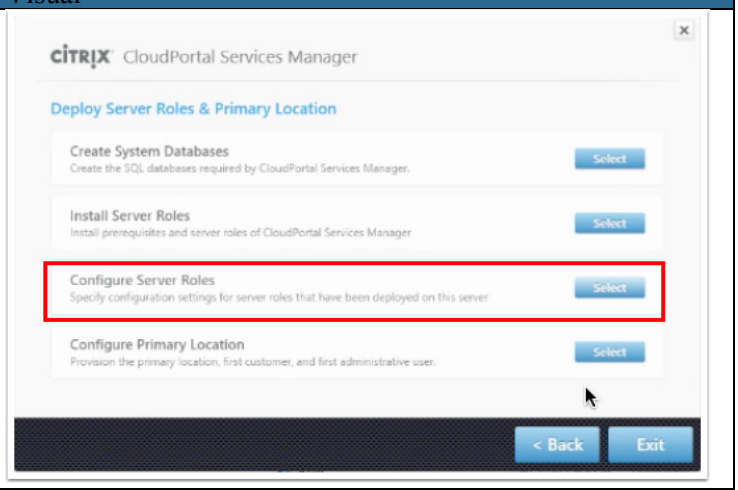

(2) Load the deployment file and click **Next**. Select the Configuration Task by enabling the checkbox for **Reporting**. Click **Next**.	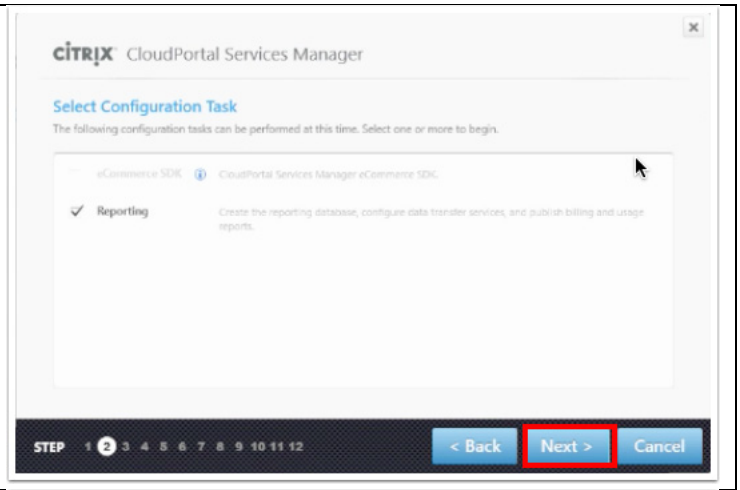
(3) Enter Reporting Database Credentials. Click Next.	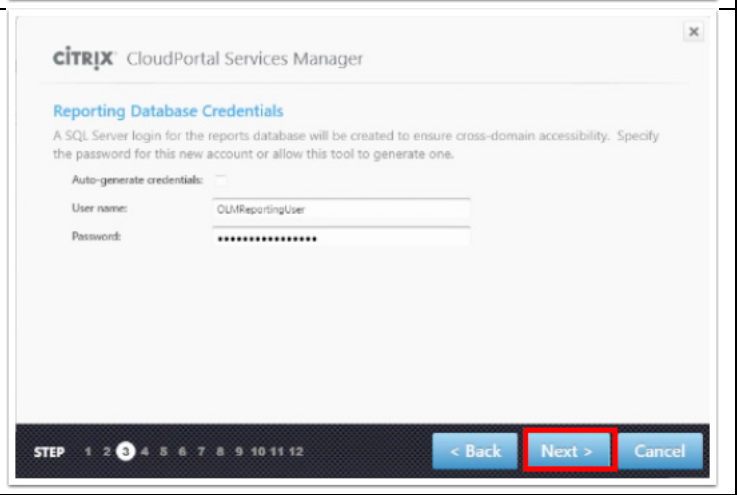
(4) Configure the Mail Server as **cpsm02.daas.local**. Click **Next**.	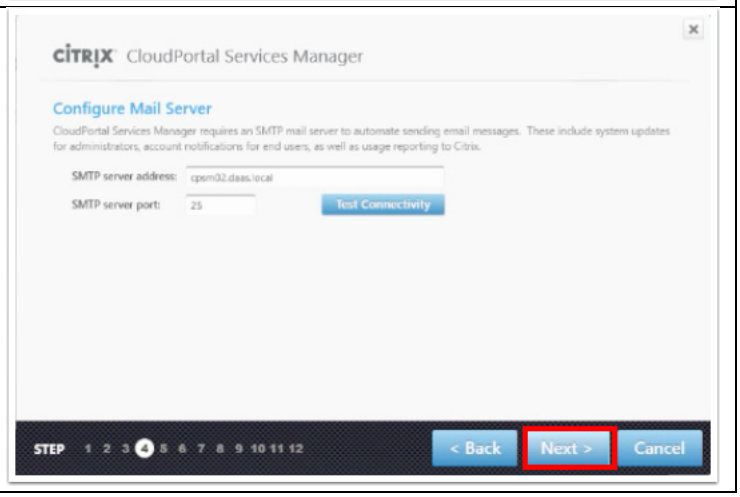

(5) The setup previews service packages for import into the control panel. Click **Next**.	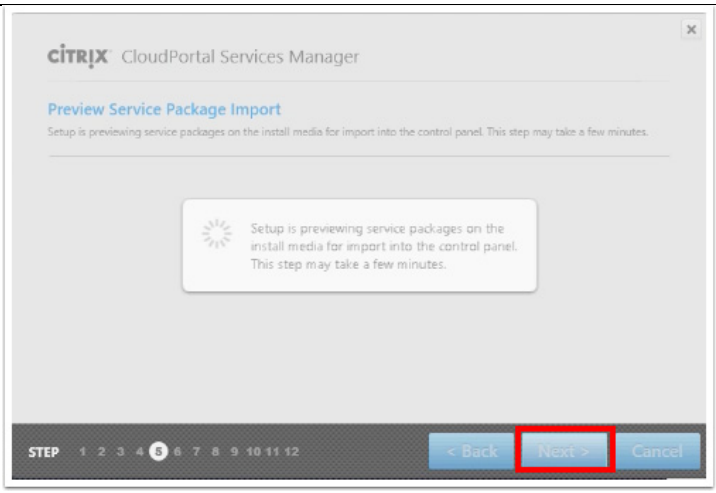
(6) On the Configure Reporting Database dialog, set the server address to an SQL server. This CVD uses a separate SQL server (**SQLHA**) for reporting. . · Click · **Test Connection** to test the connection to the database. Click **Next**.	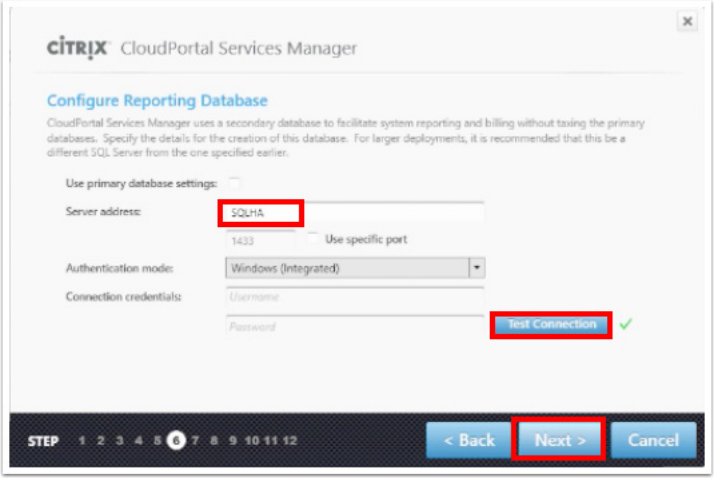
(7) On the Configure Data Transfer Service dialog, enable the checkbox for **Create if it doesn't exist**. Click **Next**.	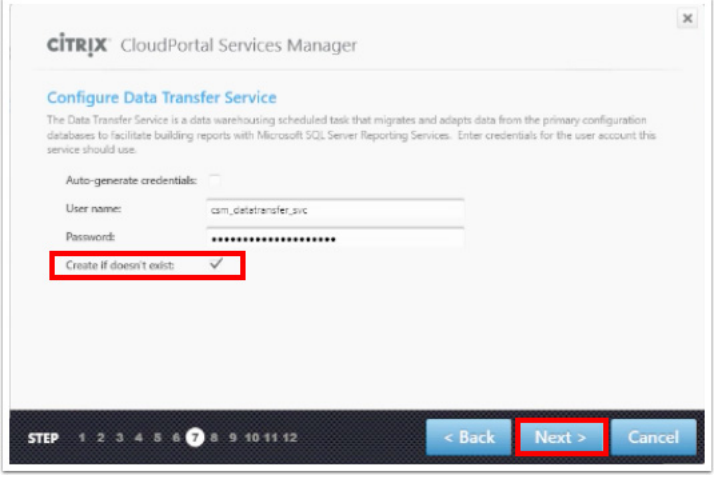

(8) On the Data Transfer Notifications dialog, enter email addresses for notification. Click **Next**.	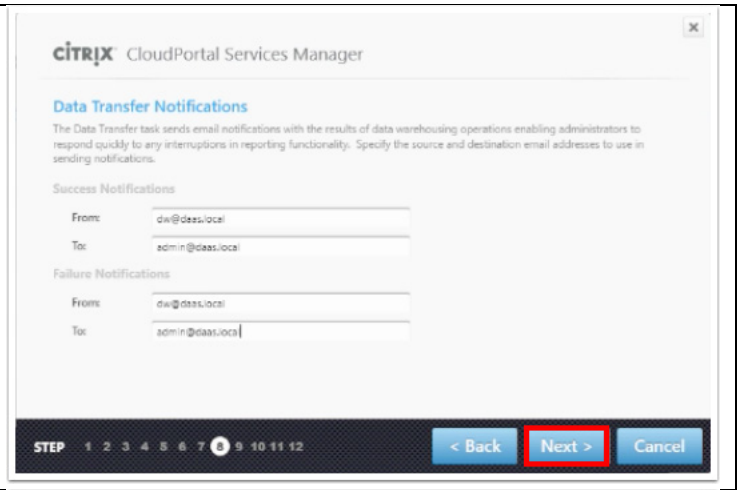
(9) Enter information on the Specify Reporting Services Details dialog. Click **Test Connection** to verify the configuration. Click **Next**.	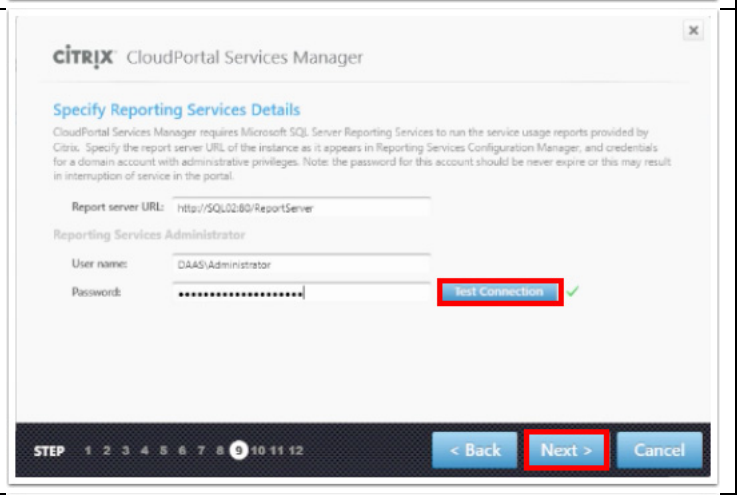
(10)Accept the defaults on the Data Warehouse Service dialog. Click **Next**.	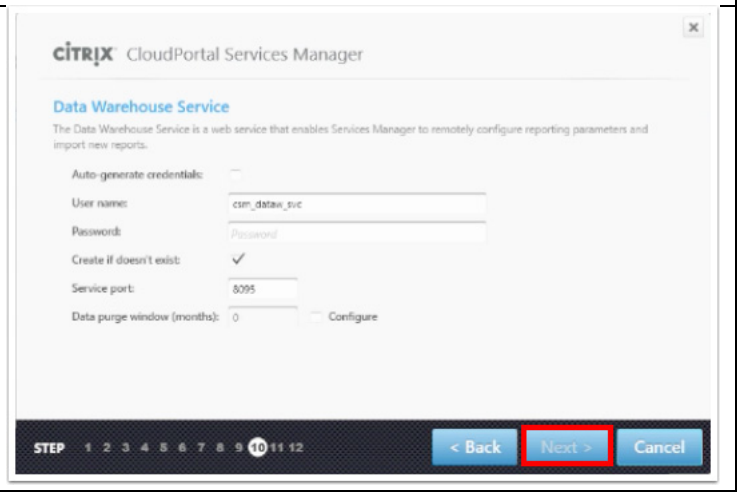

(11)Review the Summary. Click **Commit**.	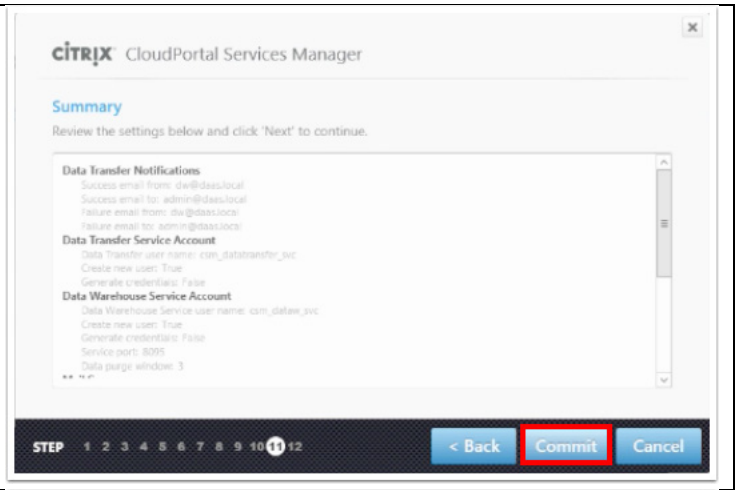
(12)The configuration is applied. When the configuration completes, click **Finish** to continue.	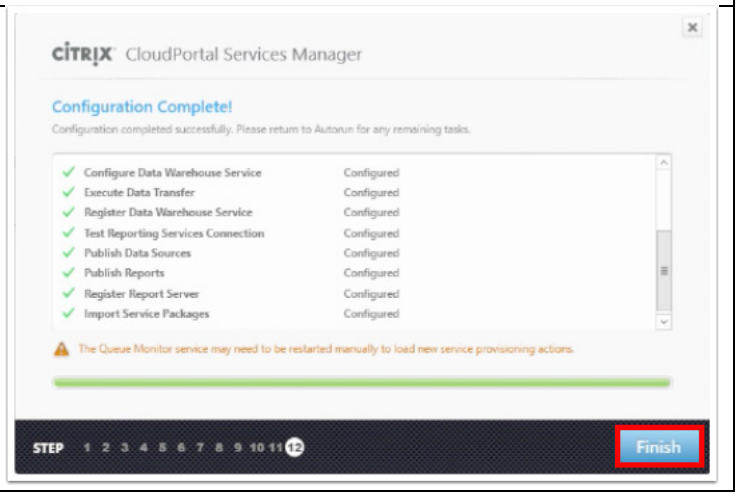

Validate the Reporting Configuration

The following steps were performed to validate the reporting configuration.

Instructions	Visual
On the SQL02 server: Select the Task Manager from the Windows Server Manager Tools panel. Select the CPSM category in the task scheduler library, and then select the **Reporting Data Transfer** task. Select **Run**.	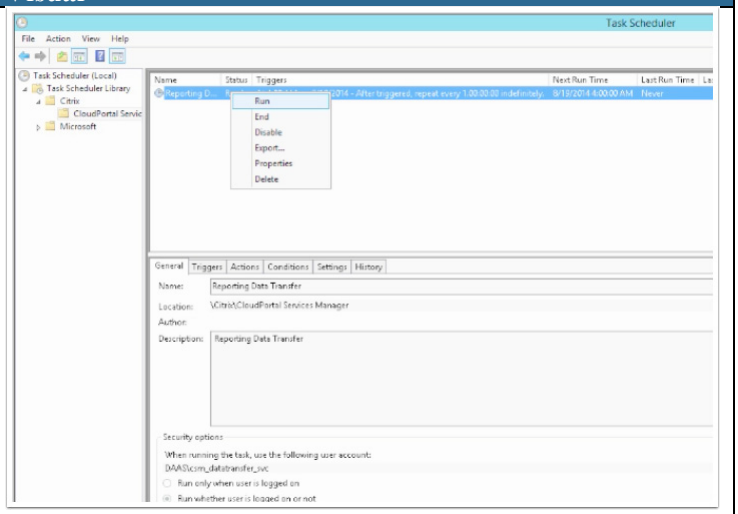

Select **Billing** from the SQL Server Reporting Service. You should see customer, distributor and reseller detail options.	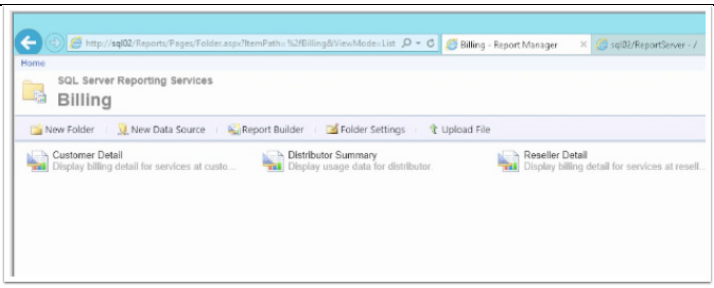
Next, validate the Data Warehouse. Log into the CloudPortal Services Manager interface as **cspadmin_CSP**. Select **Reports-> Configuration->Data Warehouse** and select **Transfer** to move data from views to the Data Warehouse.	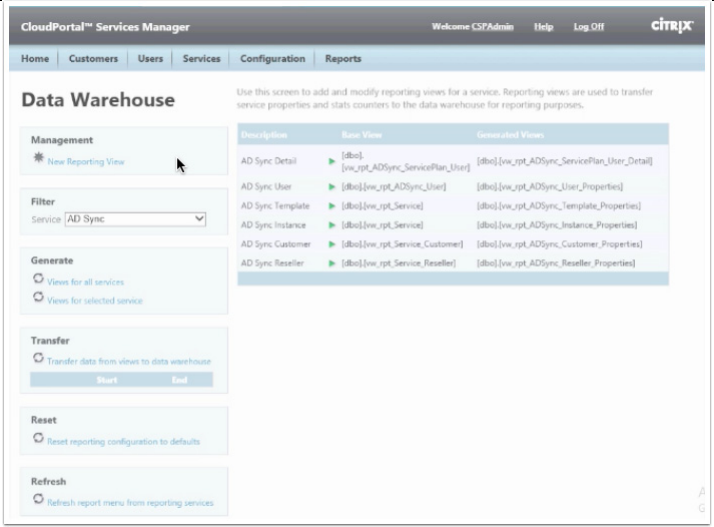
Confirm the SQL02 server connection: Select **Server Connections** from the CloudPortal Services Manager menu. Both Data Warehouse and Reporting should be displayed as connecting to the SQL02 server. Click the Test icon to test these two connections.	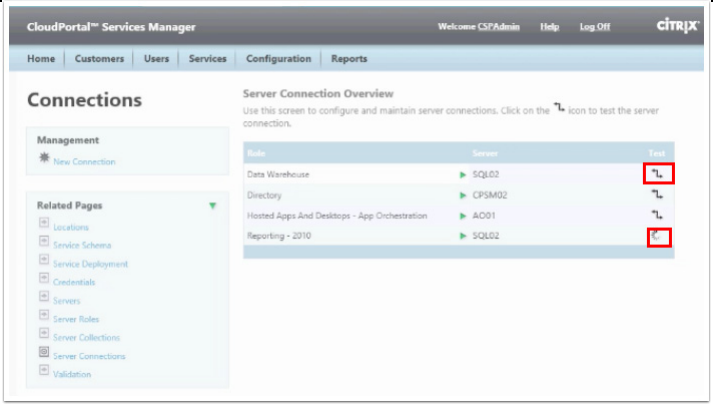

Expand the connection for Reporting and change Credentials to: **DAAS\csm_dataw_svc** Retest the connection.	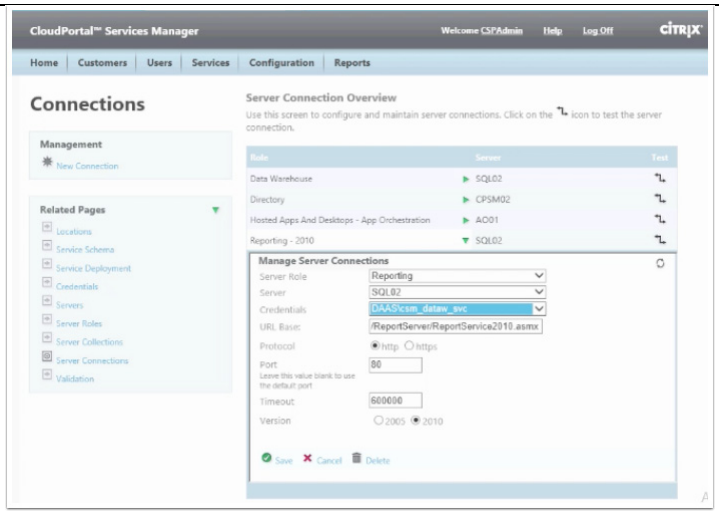
Go to the CloudPortal Services Manager **Credentials** page, and delete the credentials for the Administrator. Edit **csm_dataw_svc** to encrypt the password. Click **Update**.	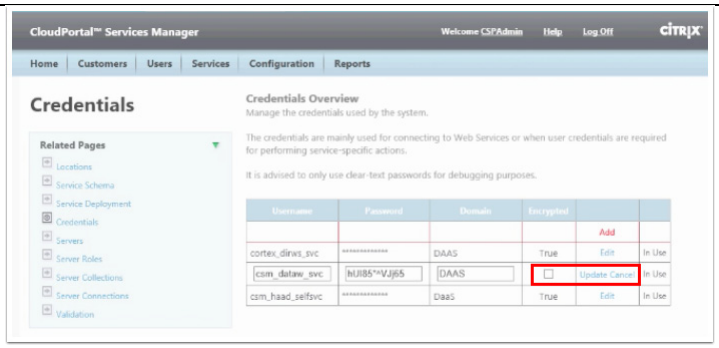
Go to the CloudPortal Services Manager **View Reports** page. Expand the **Billing** item to display the customer, distributor, and reseller detail options.	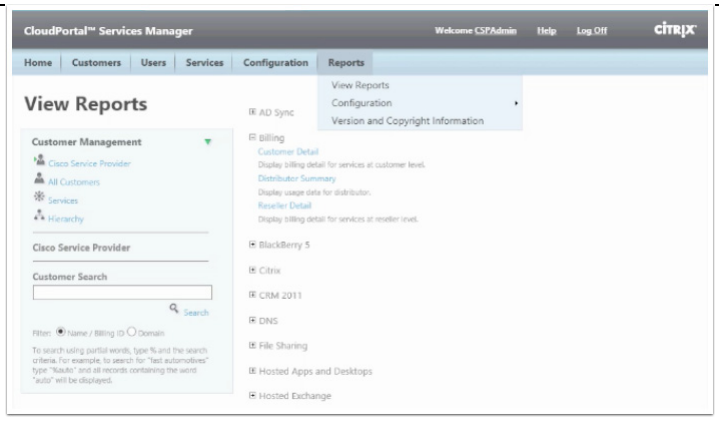

Configuring High Availability for the CPSM Databases

The following steps were performed to configure the CPSM databases for AlwaysOn availability.

Instructions	Visual

Backup up the database to a file.	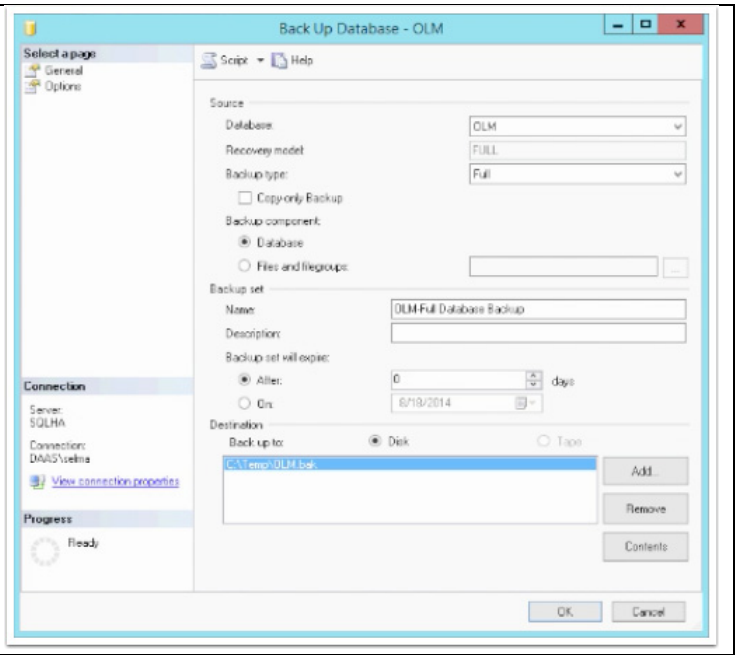
Add the OLM and OLMReports databases to the Availability Group in PVS. Specify **Full data sync**.	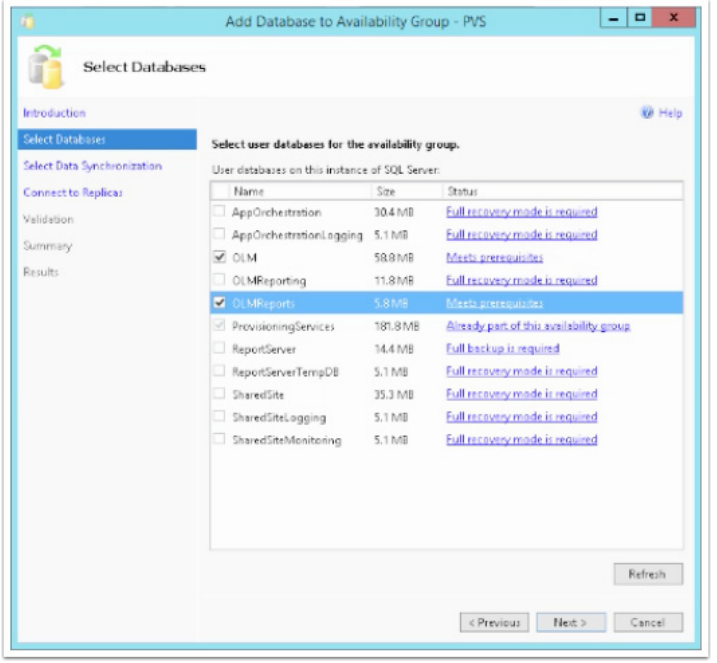

Configuring Offerings for Shared Site Tenants

This section documents the steps performed in this CVD to configure offerings for shared site tenants. Key steps include:

- Create a desktop offering name “CSP-HSD” that uses Private Delivery Group isolation
- Assign shared apps (Microsoft Excel, Word, Outlook, PowerPoint, etc.) to the CSP-HSD offering
- Test the connection to the HaaD service and publish the CSP-HSD offering as the default offering
- Allow the Reseller to provision apps within the CSP-HSD offering

Create a Desktop Offering (“CSP-HSD”)

Using the Citrix App Orchestration web console, create a desktop offering (“CSP-HSD”) that uses Private Delivery Group isolation].

Instructions	Visual
	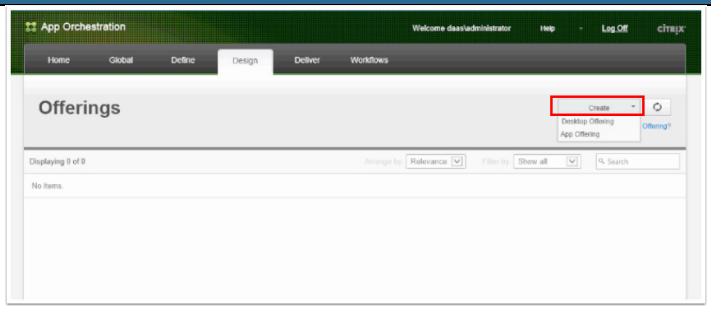
Complete the **Advanced Settings**: • Isolation: Private Delivery Group • Name: CSP-HSD Click **Save.**	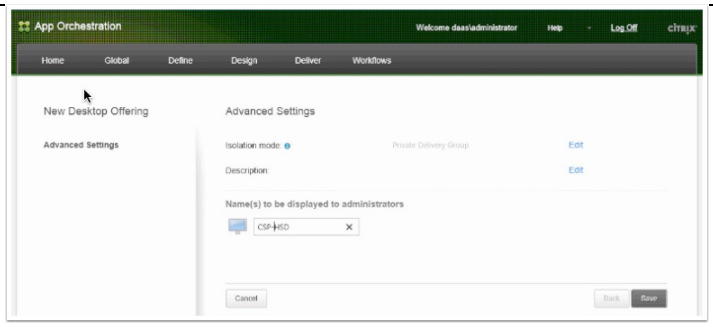
Select the CSP-HSD offering and click **Edit**. The **Edit Offering** dialog appears. Edit the **Basic Settings**, and click **Save Offering**. Make sure the Name, Display Name, and Descriptions can easily be understood once imported into CPSM.	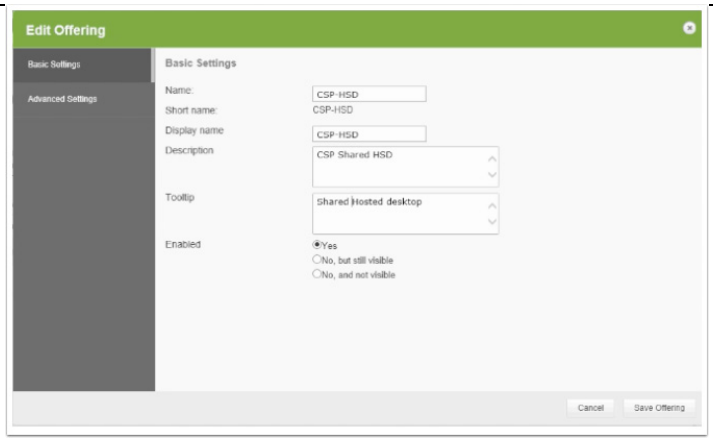

Assign Shared Apps to the CSP-HSD Offering

Use the Citrix App Orchestration web console to assign shared apps (such as Microsoft Excel and Word) to the CSP-HSD offering.

Instructions	Visual

In the Citrix App Orchestration web console, select **Design->Offerings**, and then select **App Offering** from the **Create** dropdown menu.	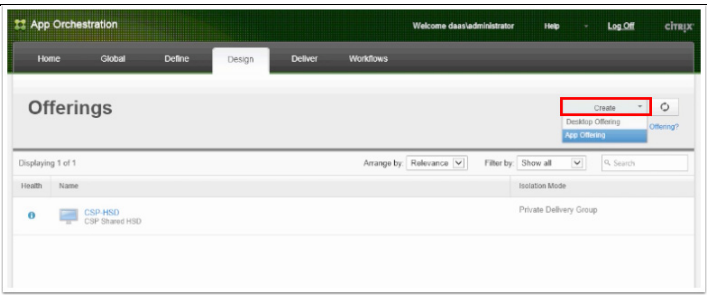
On the **New App Offering** dialog, enable the checkboxes for the following Installed Apps: Adobe Reader, Microsoft Access, Excel, Picture Manager, Outlook, PowerPoint, and Word. Select the **System Apps** tab to add any system applications. Click **Next**.	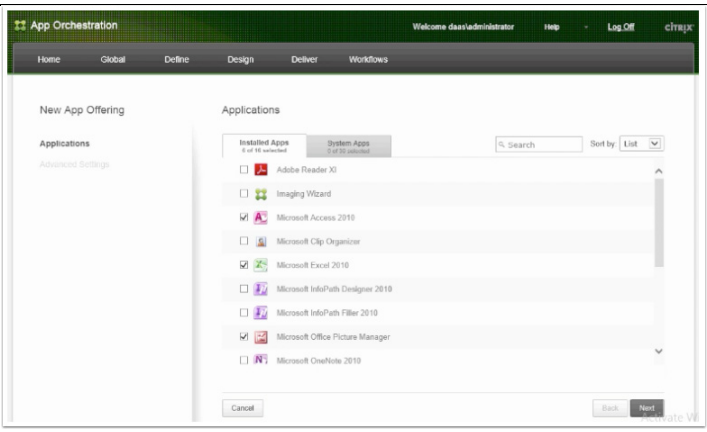
On the **Advanced Settings** panel, change the names of the apps to prepend “Shared”. Click **Save**.	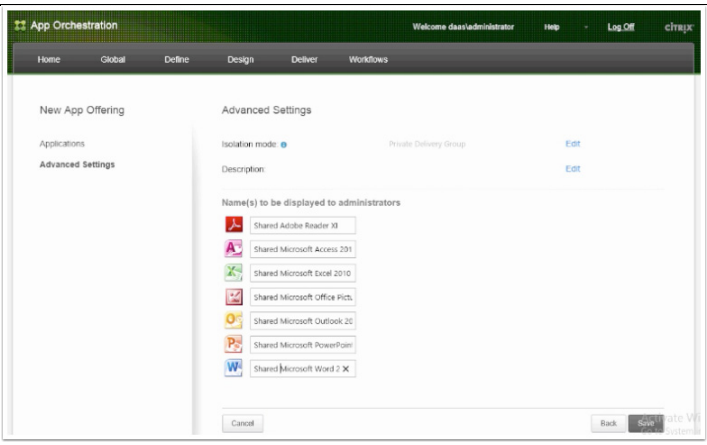
The apps are displayed on the Offerings page.	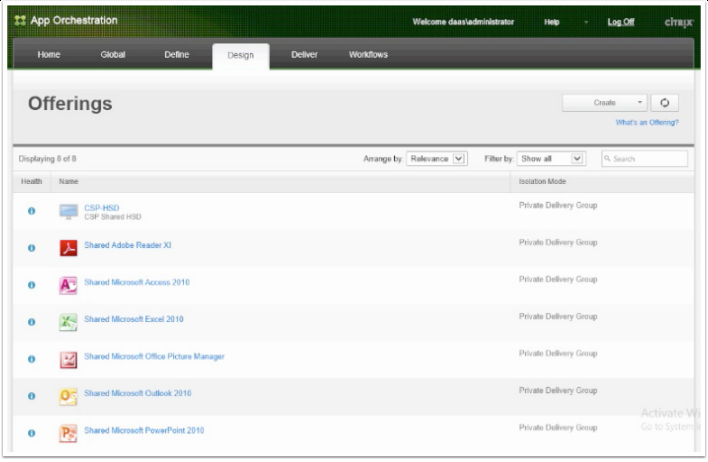

Select each app and modify the offering's Display name, prepending "Shared" to match the Name.	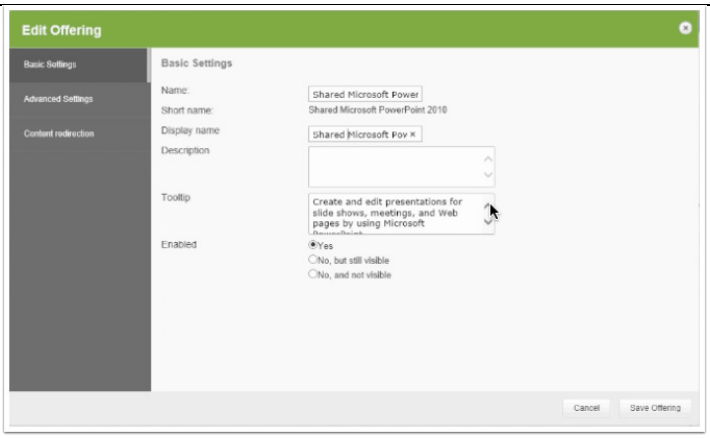

Test Connection to the HaaD Service and Publish CSP-HSD Offering

Using CloudPortal Services Manager, test the connection to the HaaD service and publish the CSP-HSD offering as the default offering.

Instructions	Visual
Log in to the CloudPortal Services Manager portal as the service provider administrator, **cspadmin_csp**. Select **Configuration->System Manager-> Server Connections**. Selected **Hosted Apps And Desktops – App Orchestration**, and click on the test connection icon.	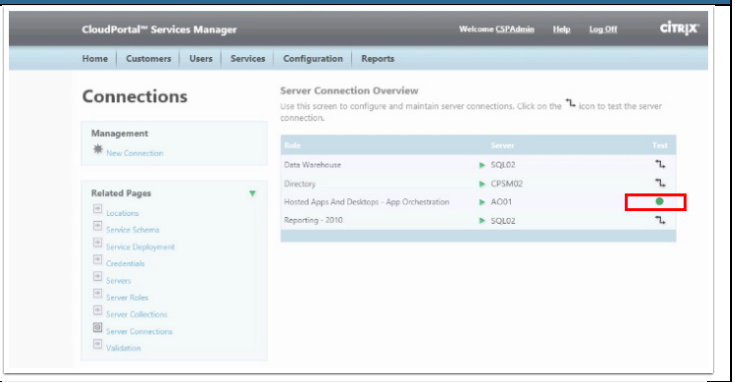
Select **Services->Hosted Apps and Desktops->Offering Management**. The app offerings are displayed.	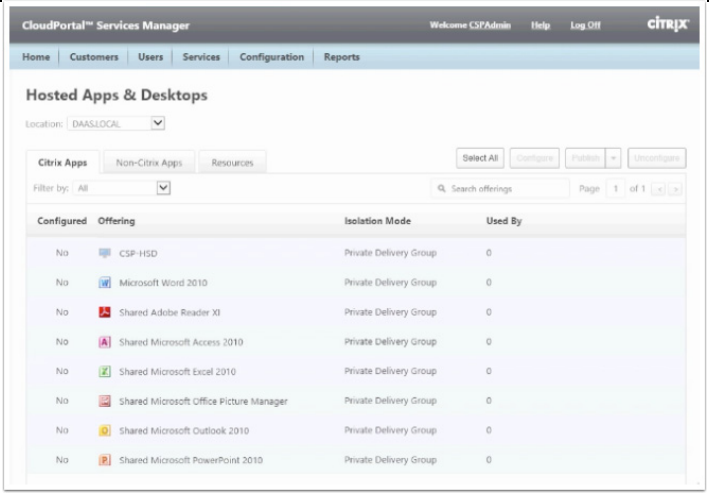

Provision apps within the CSP-HSD offering

Allow the Reseller to provision apps within the CSP-HSD offering.

Instructions	Visual
In CloudPortal Services Manager, select the CSP-HSD offering and click **Configure**. Check the box to allocate as the default application, and click **Save.**	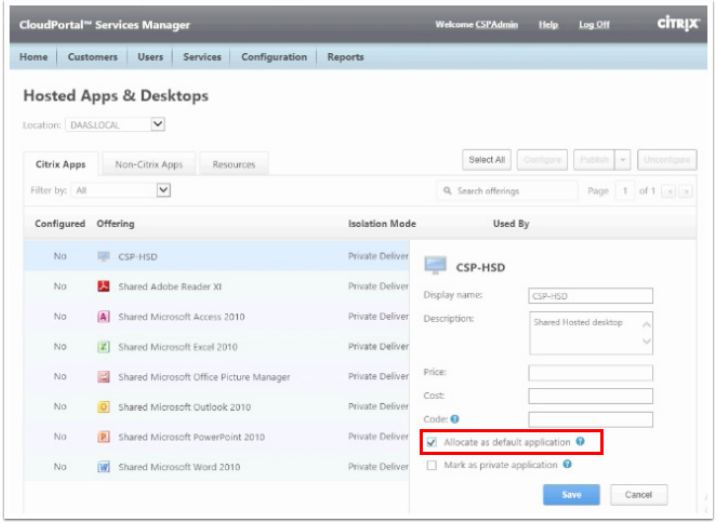
Click **Publish**, and select **Publish** in the popup to confirm publication to everyone.	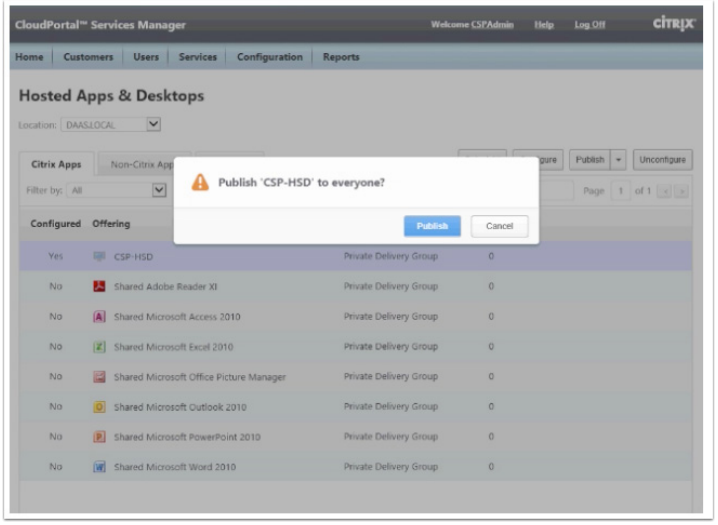

Select the Shared Microsoft Outlook offering. Configure, and click **Save**. Repeat for other shared offerings (Shared Microsoft Excel, Word, etc.)	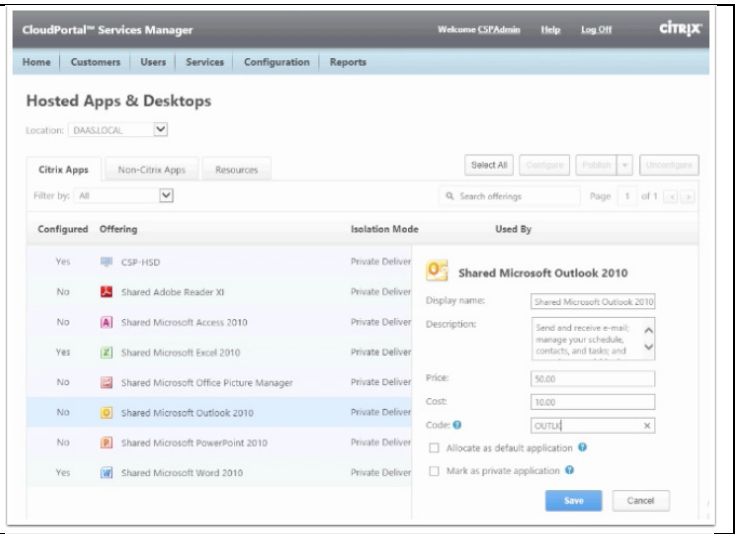
On the **Customer Services** page for the Cisco Service Provider, select **Reseller** to assign services that a reseller can provision. Click **Search Offerings** and enable the apps by clicking the checkboxes in the Enable column. Click **Apply changes**.	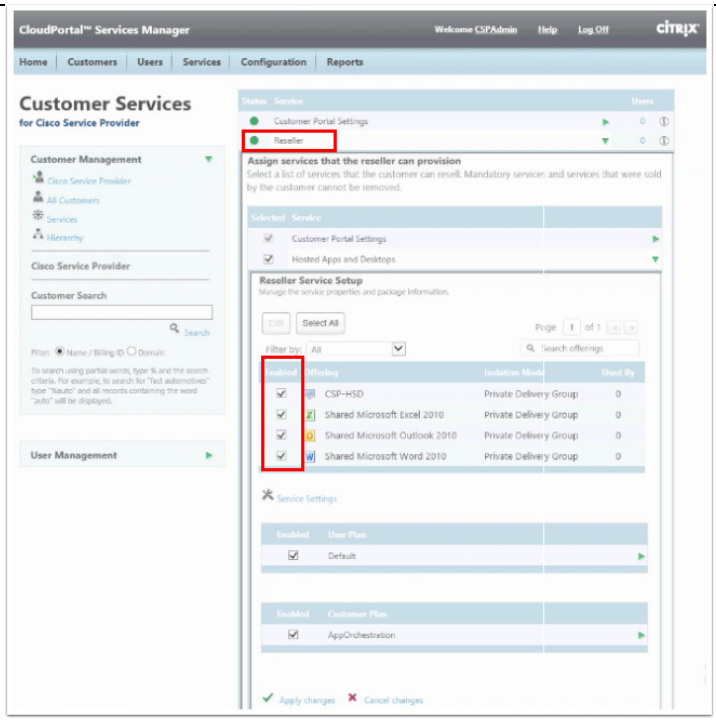

Click **Provision**. A green status icon indicates success.	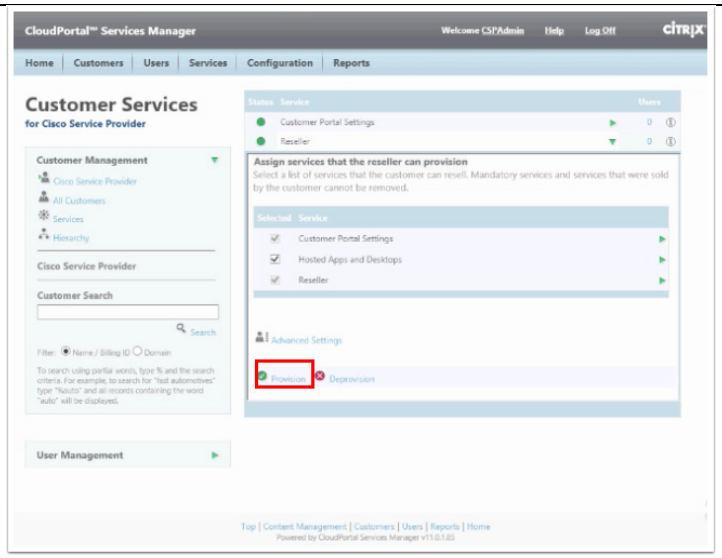

Configuring Example Shared Delivery Site Tenants

The steps in this section show how offerings are provisioned to two different shared delivery site tenants (shared delivery site tenants use private delivery group isolation). This section also describes how self-service workflows can be established for provisioning, which allows an on-site administrator to approve site-specific provisioning requests. (The sample customer “Install Test Customer” is configured to represent the use case of an on-site administrator to approving user self-service provisioning requests.)

This section includes the following tasks:

- Configuring a service plan for two example customers
- Setting up two users within the customer “Install Test Customer”
- Setting up a workflow approval for the manager to approve user subscription requests
- Impersonating a user that subscribes to offerings
- Impersonating the user’s manager to see and approve the request

Instructions	Visual
Launch the **CloudPortal Services Manager** console and login as the CSP Administrator. Display the **Customers** page. Select the customer “Install Test Customer.”	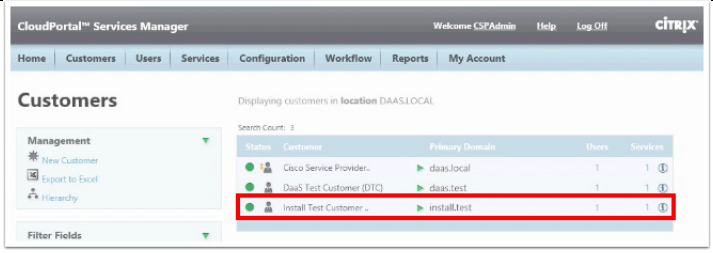

On the **Customer Services** page, select **Hosted Apps and Desktops** to configure the service plan to be applied to this customer. Enable the appropriate offerings for this tenant (in this case Hosted Shared Desktops and shared Microsoft applications are enabled). Click **Provision**.	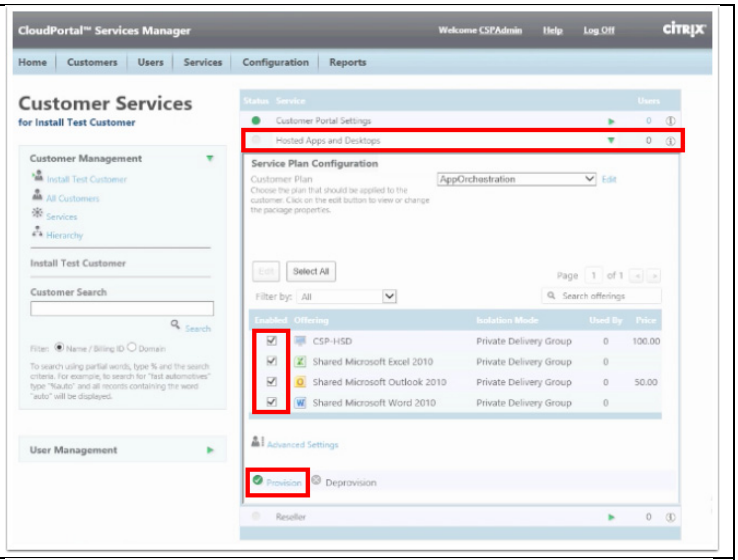
On the **Customers** page, select the customer "DaaS Test Customer."	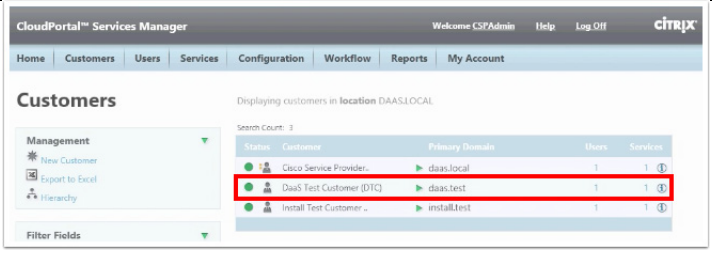
On the **Customer Services** page, select **Hosted Apps and Desktops** to configure the service plan to be applied to this customer. Enable offerings for this tenant. Click **Provision**.	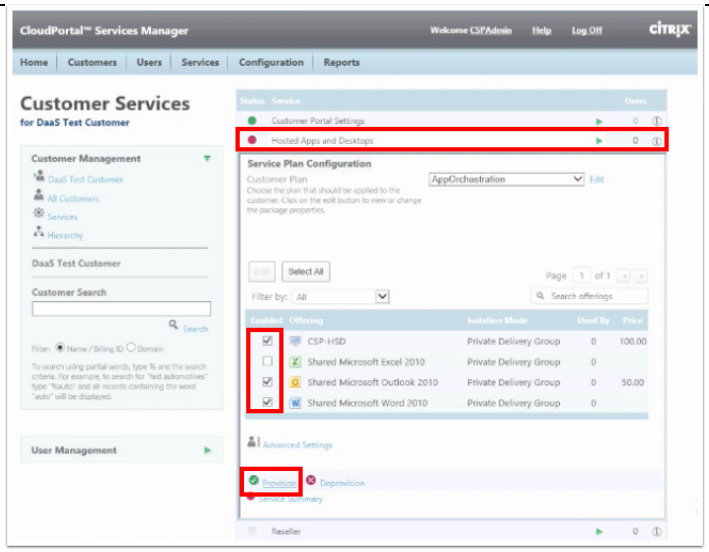

From the Configuration tab, select **Workflow Setup** from the System Manager menu.	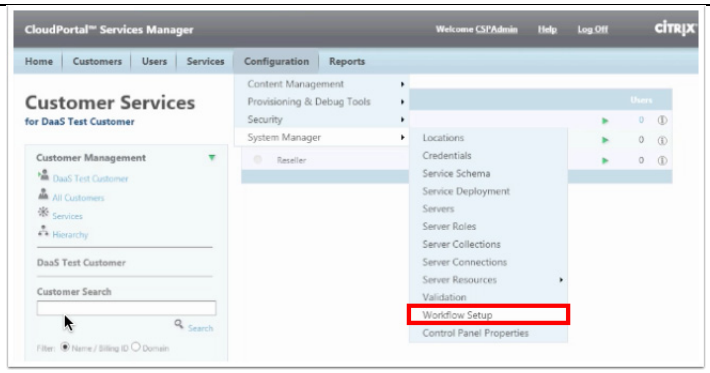
Click **Enable** to enable workflow approval.	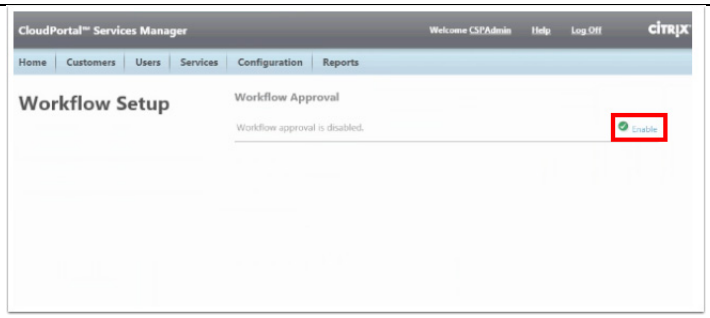
On the **Workflow Setup** page, uncheck Group approval. (**Manager approval** remains selected.) Enter the appropriate email information for approval notifications. Verify that the URL matches the current CPSM web portal URL for the deployment; if not, make the appropriate correction. Click **Save**.	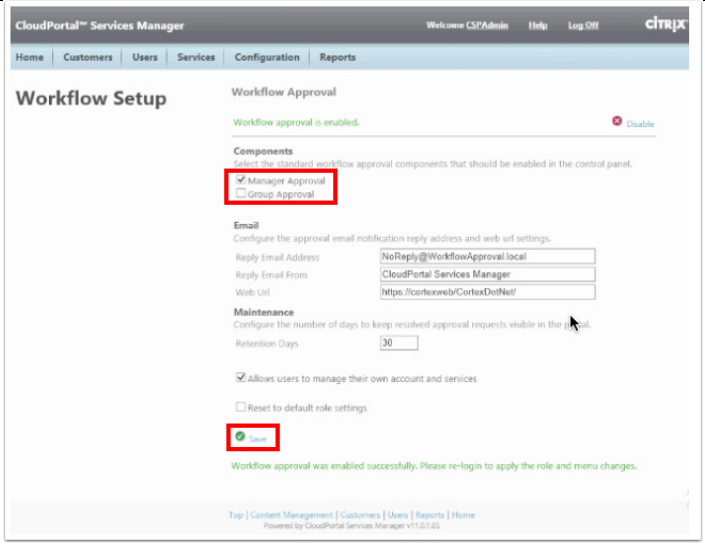
Log off and then log back in as the same service provider administrator for the newly enabled Workflow menu group to appear. Next, configure a Workflow Manager to approve self-service Hosted Apps and Desktops provisioning, From the Workflow tab, select **Workflow Managers** from the Configuration menu.	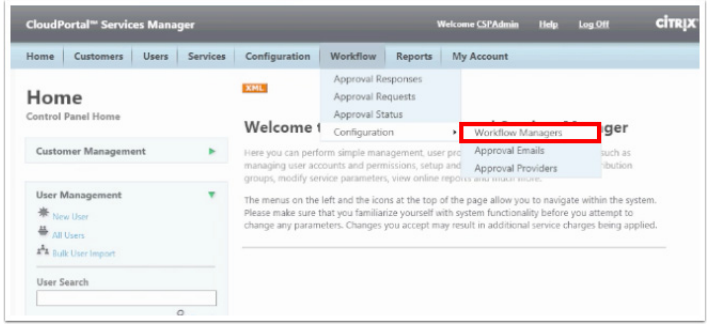

Specify **Hosted Apps and Desktops** as the user services filter. Enable provisioning by clicking the **Provision** radio button. Click **Save**.	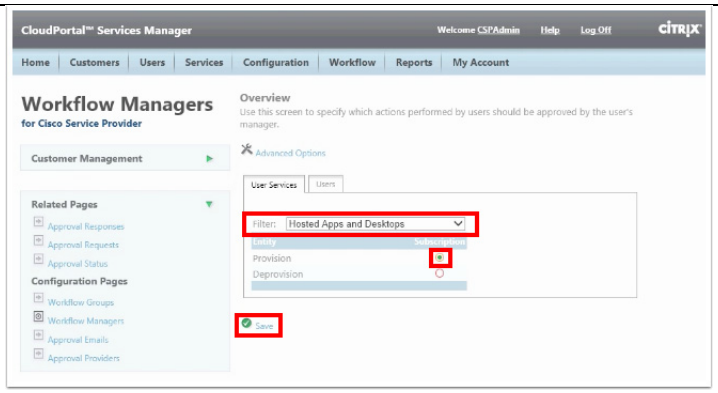
Creating users will demonstrate how the user self-service provisioning and manager approval process works. On the **Customers** page, select "Install Test Customer," which displays users for this customer. Select the administrator (Install Admin) and click on **Impersonate** in the **User Functions** menu. (Impersonating Install Admin enables creating a new user for this customer.)	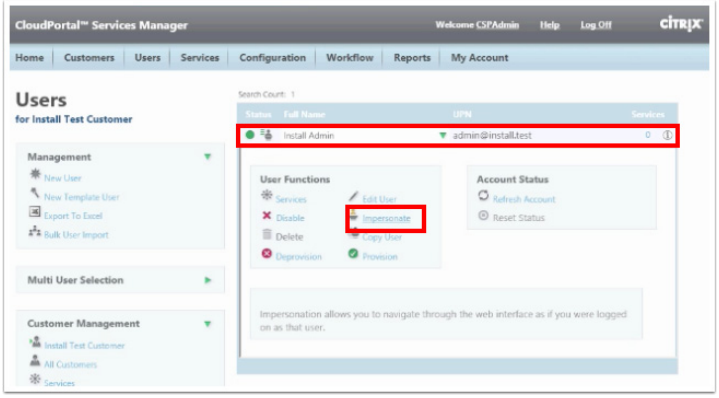
Note the ID in the banner changes to "CSPAdmin impersonating Install Admin". Click on **New User** and create a sample user (UserOne). Under Additional Properties, enter Install Admin as the Manager. Click **Provision** at bottom of the form. Log Off to stop impersonating Install Admin.	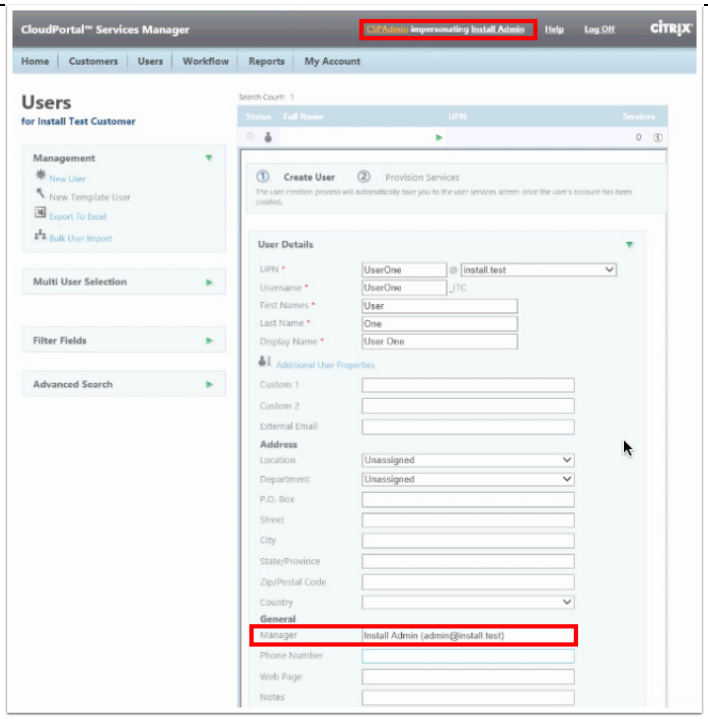

From the **Users** page, select User One and click **Impersonate** in the **User Functions** menu.	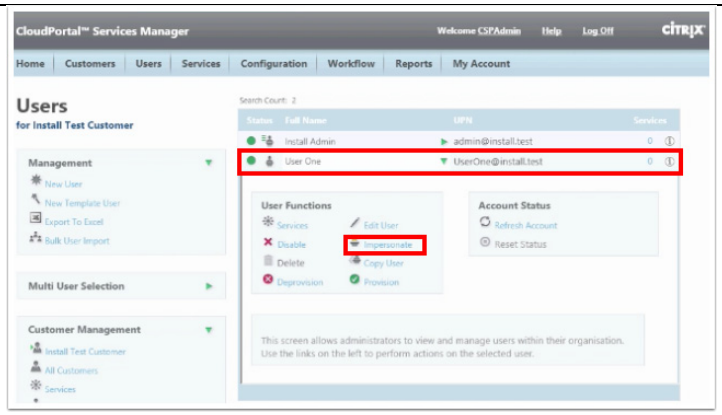
As the user, select **Services** from the **MyAccount** tab. On the **User Services** page, subscribe to offerings from those listed for the Hosted Apps and Desktops service. Click **Provision**. Log Off to stop impersonating User One. Note that workflow approval has not yet been enabled for Install Test Customer, so User One was provisioned without having to go through a workflow approval process.	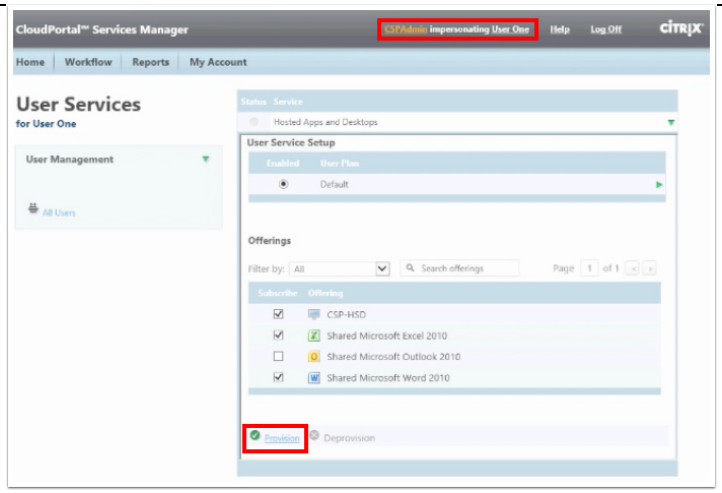
Choose **Workflow Managers** from the Workflow -> Configuration menu. On the **User Services** tab, select **Hosted Apps and Desktops** as the filter. Click the **Provision** radio button. Click **Save**. (This step configures a workflow manager at the Install Test Customer management level).	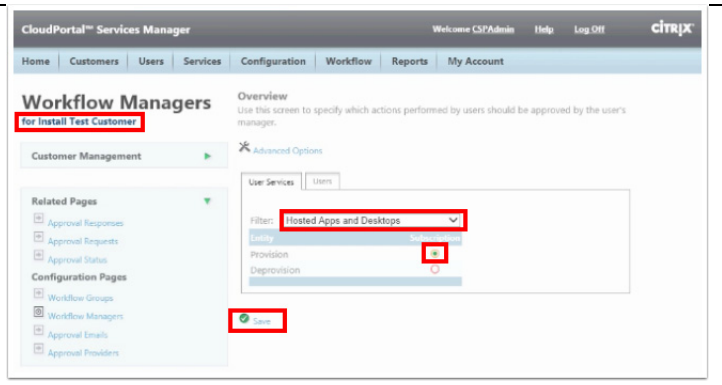

Create a new user (User Two) for the Install Test Customer. Under **Additional Properties**, enter Install Admin as the **Manager**. Impersonate the user and subscribe to offerings, as with the other new user.	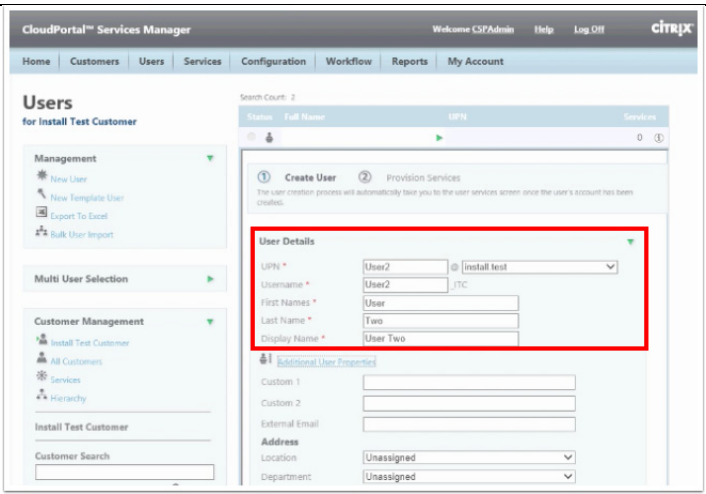
Next, impersonate the customer administrator to approve the pending provisioning request for User Two. From the Users menu, select “Install Admin” and click Impersonate. Choose Workflow -> Approval Requests, and select the pending request. Click **Request Details**.	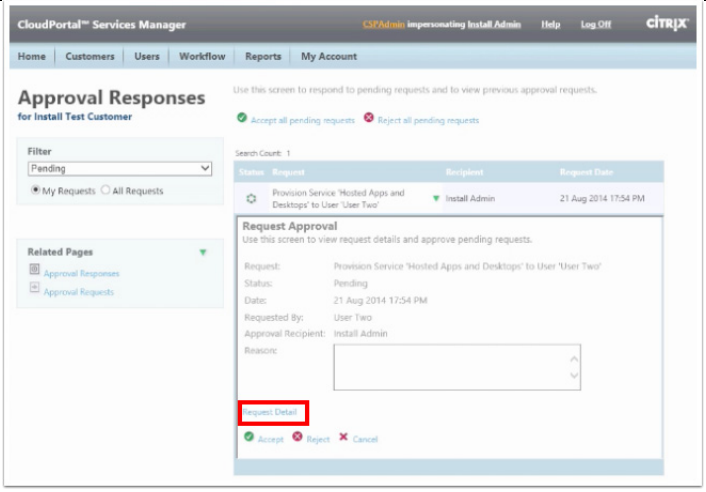
Clicking on **Request Details** causes a pop-up to appear with additional information. Since workflow approval was enabled for this customer, User Two is now subject to the manager’s approval process before service offerings can actually be provisioned.	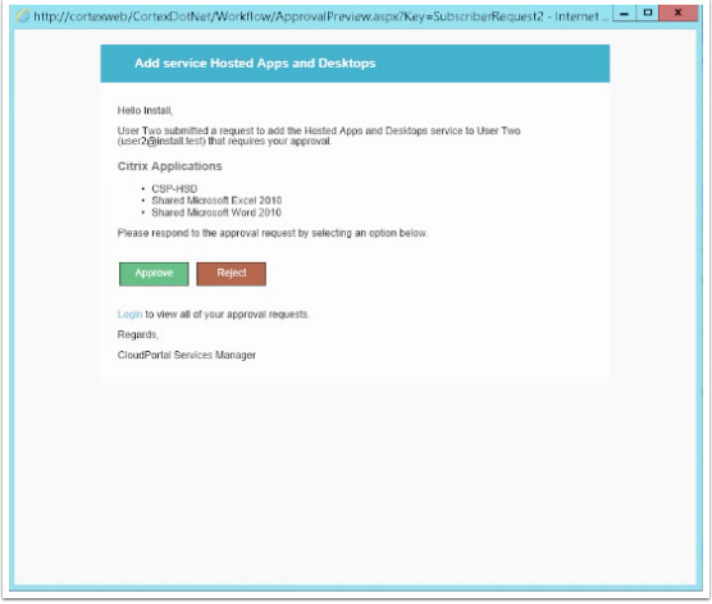

As the CSP administrator, select the DaaS Test Customer and enable the default offerings for users within this customer. From the **Customers** menu, select the DaaS Test Customer. Select **Services** from the User Functions menu. Select the offerings from those listed for the Hosted Apps and Desktops service. Click **Provision.** This step demonstrates how the CSP administrator can manage any customer's users directly.	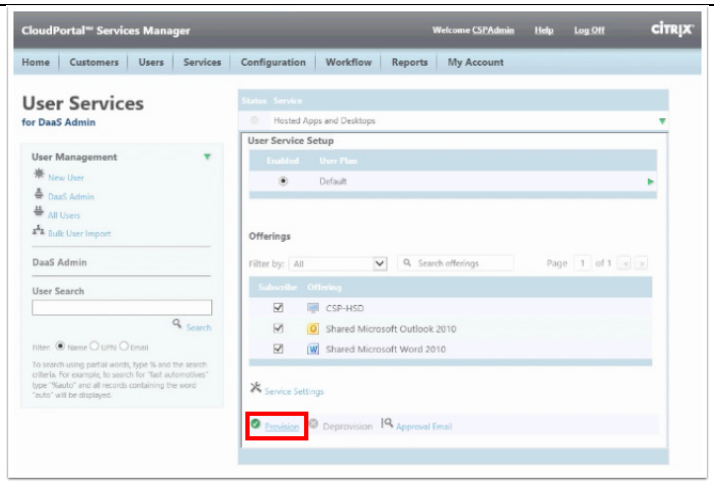

Configuring a Session Machine Catalog for Shared Delivery Site VDI Users

The steps in this section set up a Session Machine catalog for users in a shared delivery site that require Server VDI desktops using XenDesktop (rather than Hosted Shared Desktops via XenApp).

Instructions	Visual
Launch the App Orchestration Web Console and select **Session Machine Catalogs** from the Design tab. Select "**Created externally**" from the Use Machines dropdown list. The HSD catalog for shared delivery site tenants is visible in the session machine catalog listing.	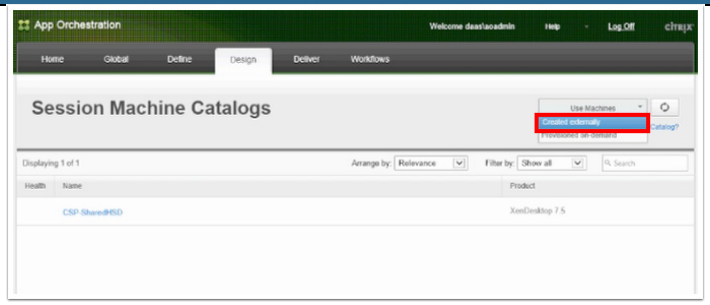
Configure Basic Settings for the new Server VDI Session Machine Catalog. Name the catalog (CSP-SVDI) and check the radio button for a Single-User OS. Click **Next**.	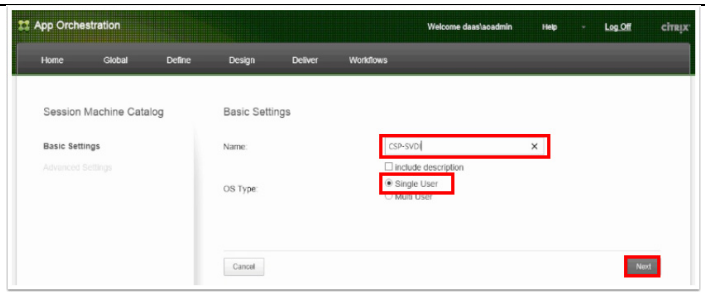
On **Advanced Settings**, select **Static** for the type of desktop. Click **Next**.	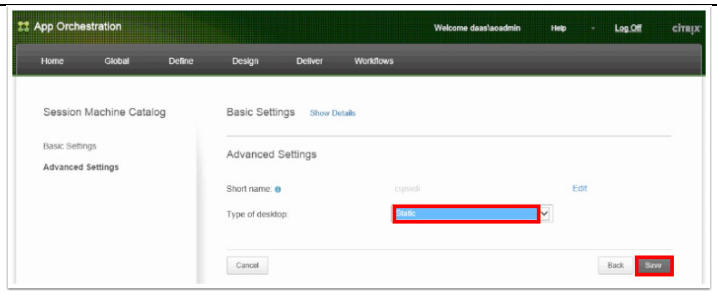

The Session Machine Catalogs page now shows two catalogs, one for shared SVDI and one for shared HSD sessions. Select the SVDI catalog (CSP-SVDI).	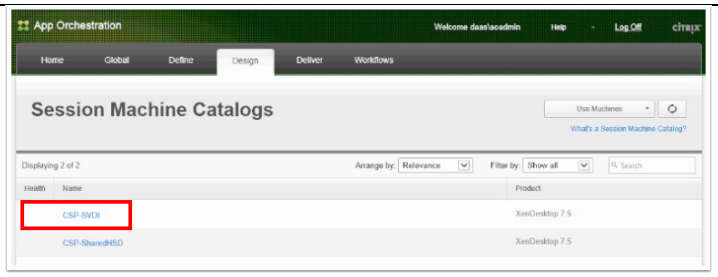
On the StoreFront server in the Shared Delivery Site Domain (e.g., DaaS-SF01 or DaaS-SF02), launch the Citrix **StoreFront console**. From the Storefront console, expand the **Authentication** panel. Select add/remove Authentication Methods and check **Domain pass-through**. Click **OK**.	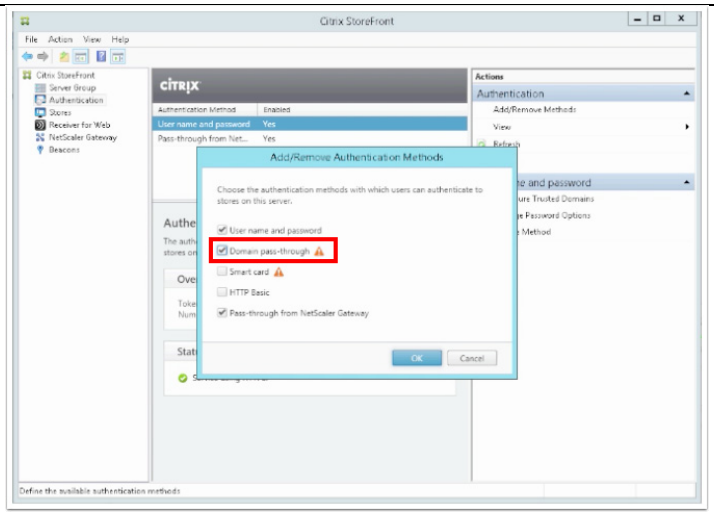

Validation with Citrix Receiver

This section illustrates how a user connects to the shared delivery site, which validates the installation of the Citrix software components and the configuration of the shared delivery site. The load balancer runs Citrix Receiver, simulating a user from a shared delivery site tenant domain connecting to the site.

Instructions	Visual
From Stores panel on the Storefront console, select the **SharedSiteStore**. Note the URL for the load balancer.	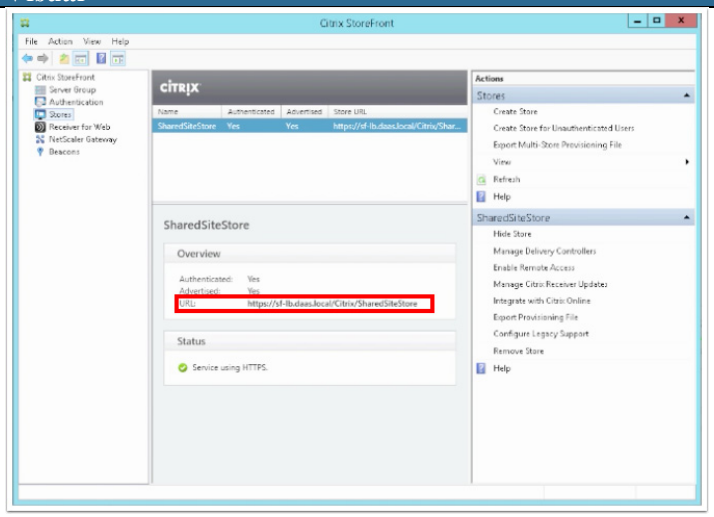

If the external StoreFront URL is not yet available, connect to the internal load balanced address, in this case: https://sf-lb.dass.local/citrix/sharedsite.	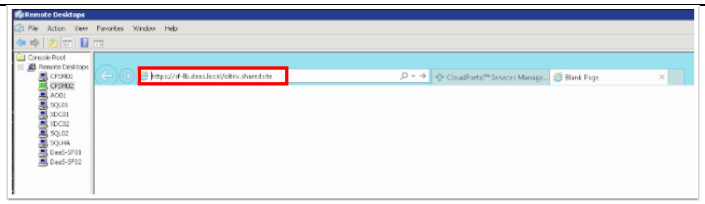
The installation screen for Citrix Receiver appears. Click **Install**.	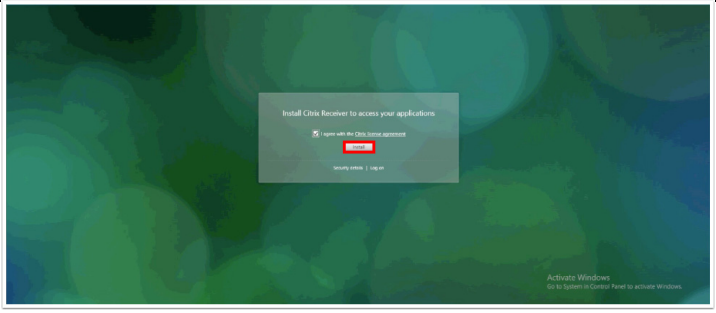
On the CloudPortal Services Manager server, run the Citrix Receiver setup program, **CitrixReceiverWeb.exe**. The Setup screen appears. Click **Install**.	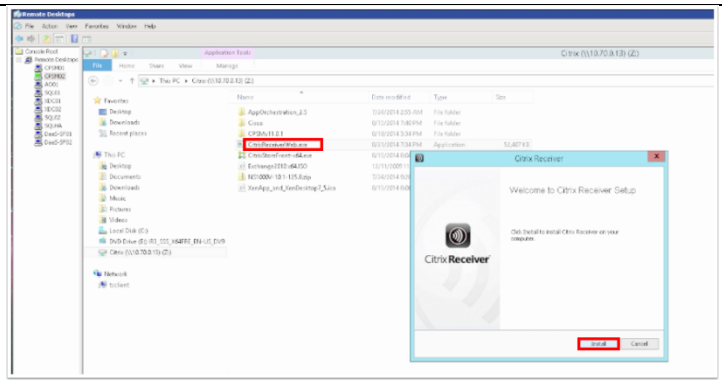
Log into Citrix Receiver using a test user and customer domain defined previously (e.g., UserOne@install.test).	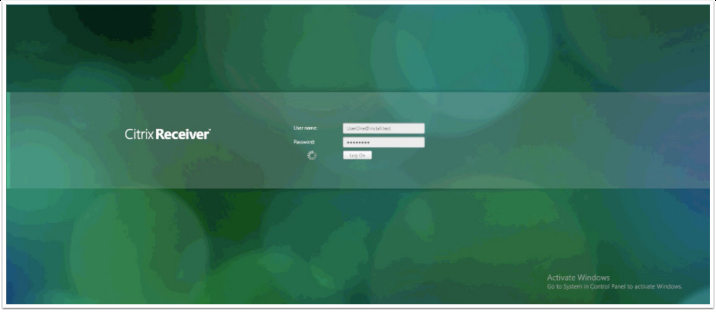
A hosted shared desktop (HSD) session launches, running a session offering defined in the session machine catalog (CSP-HSD).	

The user can then access Windows applications to which the user subscribed when configuring the offering. This access validates the successful configuration of the shared delivery site.	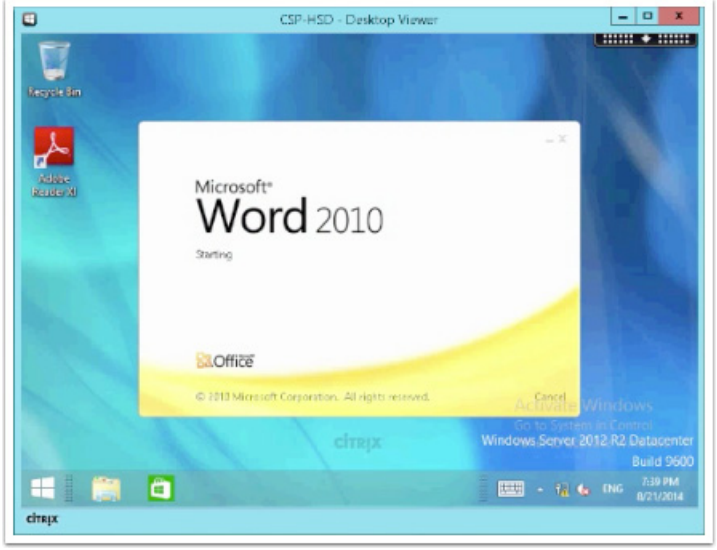

Configuring DaaS Delivery for a Private Delivery Site

Access to published apps and desktops is controlled through the App Orchestration Delivery Groups and Delivery Sites. StoreFront Server Groups, which can be private or shared, manage desktop and application offerings for subscribers.

The following sections illustrate the following tasks in configuring service delivery for a new private tenant (Private Tenant2) in this CVD:

- Install and Configure the Zero Trust Agent (ZTA)
- Configure CloudPortal Services Manager (server roles and location)
- Configure Delivery Sites
- Configure a Session Machine Catalog
- Configure StoreFront Server Groups
- Configure Offerings

Install the Zero Trust Agent (ZTA)

This section illustrates installing and configuring the Zero Trust Agent for a new private delivery site (domain **fi2.local**). The Zero Trust Agent (also known as a domain agent) allows the App Orchestration configuration server to orchestrate resources in a private domain without requiring Active Directory trusts between the App Orchestration domain and the target orchestrated domain.

For more information, see the Citrix support document Deploying the ZeroTrust Agent in App Orchestration 2.5.

Install and Configure the Domain Agent

Instructions	Visual
On the server where you are installing the domain agent (FI2-CPSM, in this CVD), log on with a domain user account that is also a local administrator on the server. From the App Orchestration installation media, launch Setup.exe to run the Citrix App Orchestration Install Center and then click **App Orchestration Domain Agent**. When prompted, click **Start** and accept the End User License Agreement. The Citrix App Orchestration Domain Agent Setup installs the agent software.	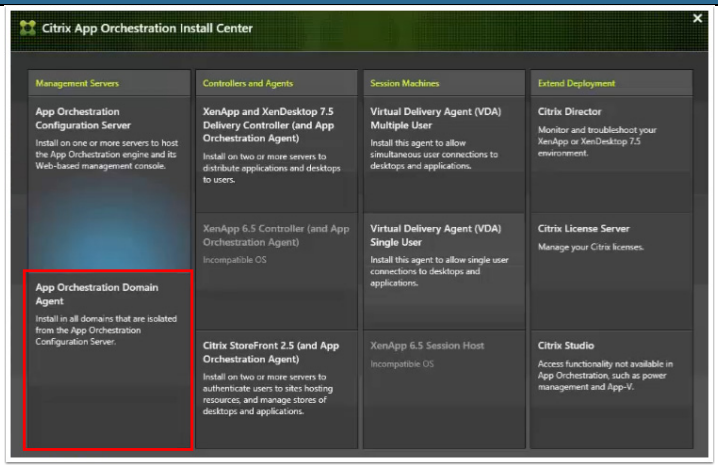
When the installation completes, it prompts you to launch the Citrix App Orchestration Server Configuration wizard. Click Close to continue.	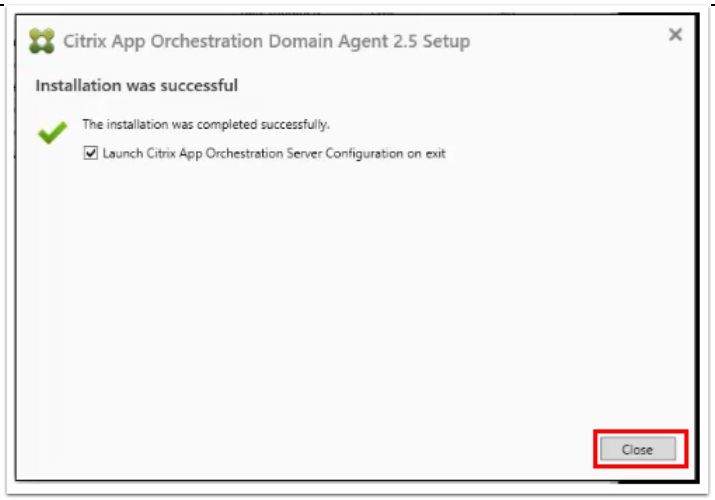
From the Citrix App Orchestration Server Configuration wizard, select **Create a new domain**.	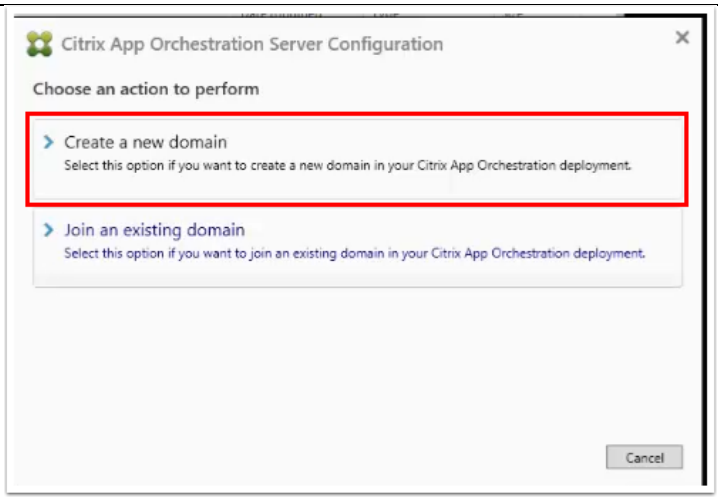

Enter the information for the new domain: Configuration server — FQDN of one of the configuration servers, a user account that is defined as a Full Admin in AO, and the password for the configuration server user account. Domain — The name of the new domain to be defined (**fi2.local**), and the domain type. Service user — name and password of the user account to be used within this domain. This account must have Full Control permissions for the App Orchestration root OU in this domain. Click **Next**.	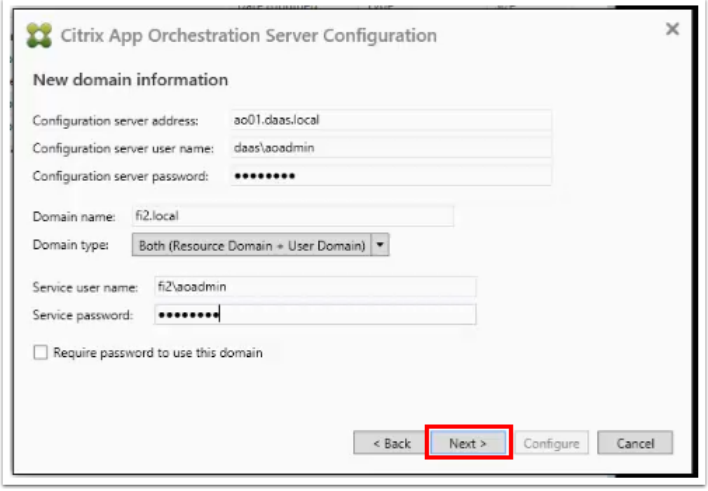 **Note:** if the certificate of the App Orchestration server or issuer is not in the trust store of this Zero Trust Agent machine, this step will fail.
The **Ready to configure** screen appears, summarizing the information for configuring the new domain in App Orchestration. Click **Configure**.	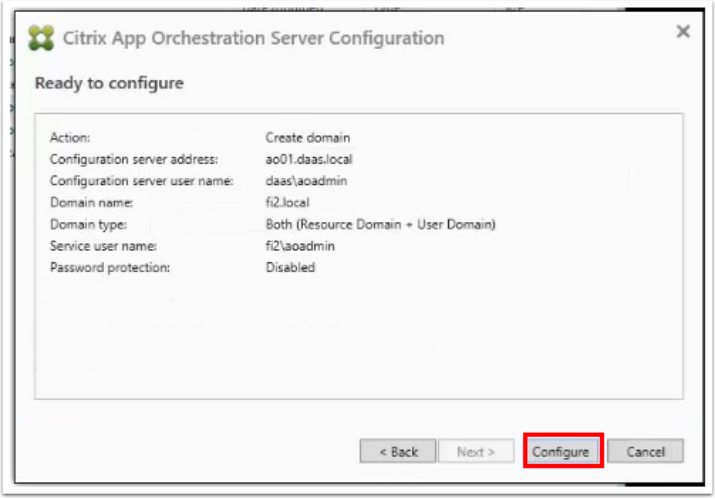
When finished, a notification appears confirming the configuration is successful. Click **Close** to continue.	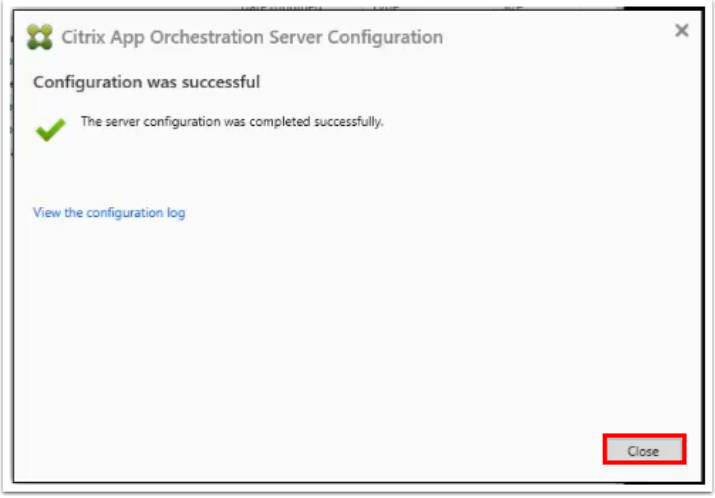

Verify New Domain Configuration

The following steps verify that the new domain is configured successfully.

Instructions	Visual
Log on to the App Orchestration web console and click **Define > Domains**.	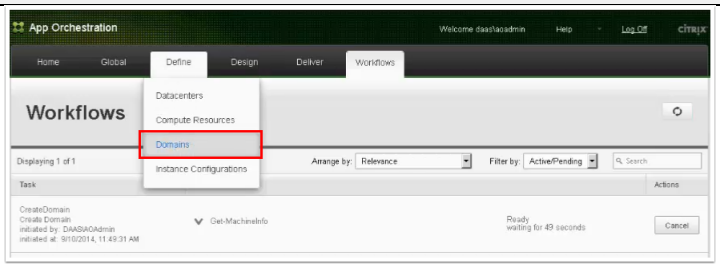
Before using the new domain (**fi2.local**, which has a scope of Private), it should be verified that all appropriate workflows have successfully completed.	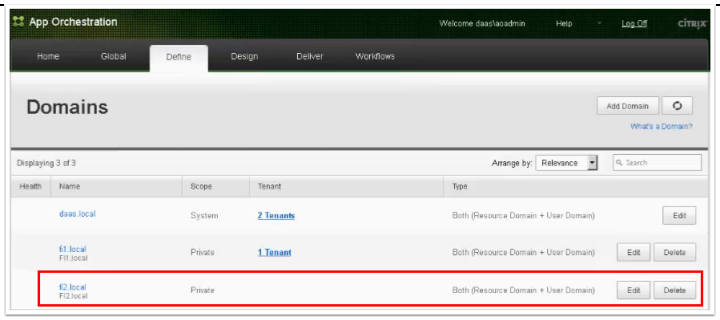

Configure CloudPortal Services Manager

This section illustrates configuring CloudPortal Services Manager for the new private delivery site. Specific subtasks include:

Configure server roles for Web Directory Services and Provisioning

Configure the new location in Active Directory and the database

For more information, see the CloudPortal Services Manager 11.0 documentation.

Configure Server Roles

The following steps configure the server roles for a new remote location.

Instructions	Visual
Log in to the FI2-CPSM server as Admin-User. Start CloudPortal Services Manager Installer, from the **Configure Remote Location** page of the Setup Tool, select **Configure Server Roles**.	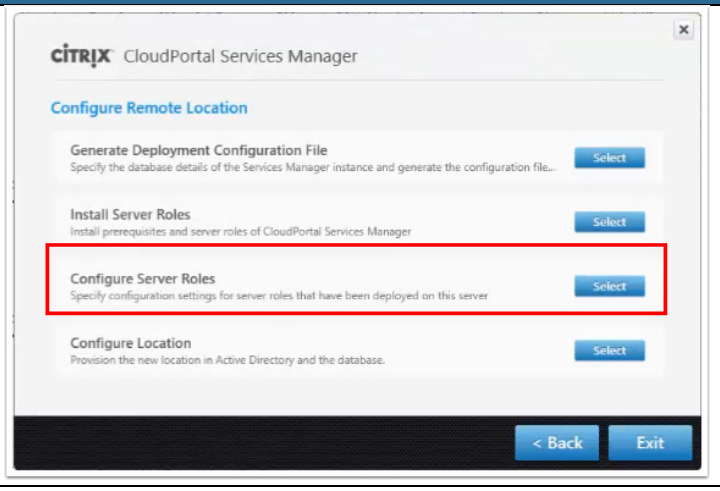

Select **Configure Server Roles**.	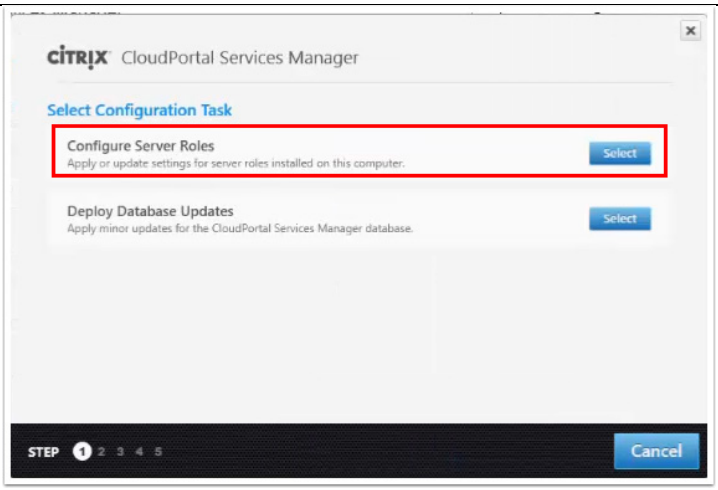
Enter the path to the configuration file (created earlier in the installation process).	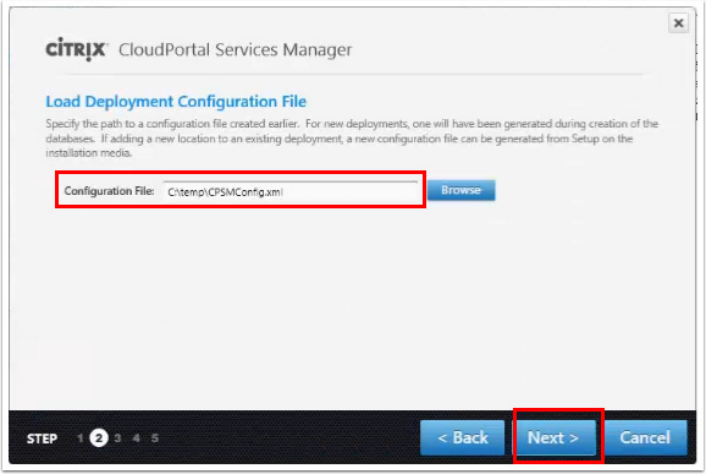
Enable the checkbox for **Directory Web Service**. Click **Next**.	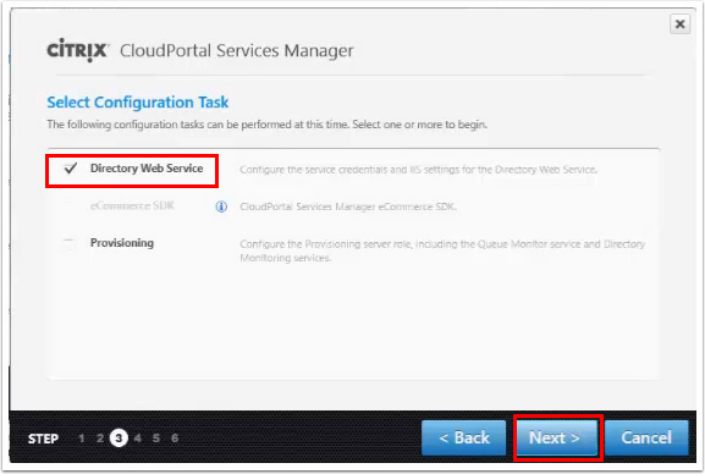

Accept defaults (from previously specified configuration file). Click **Next**.	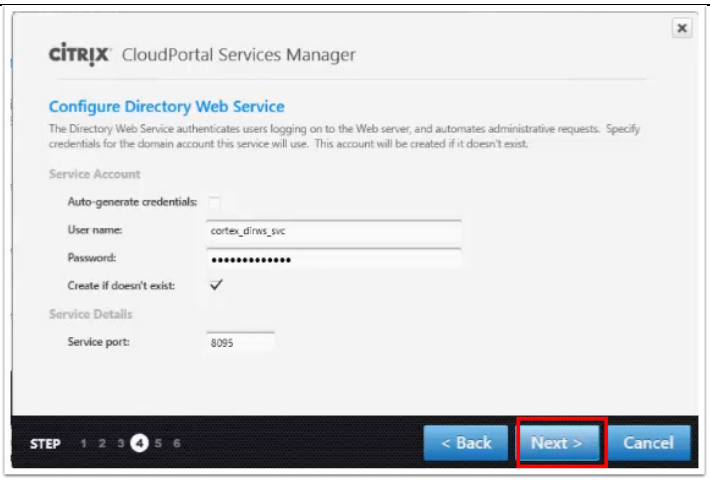
Confirm the information on the **Summary** page. Click **Commit**.	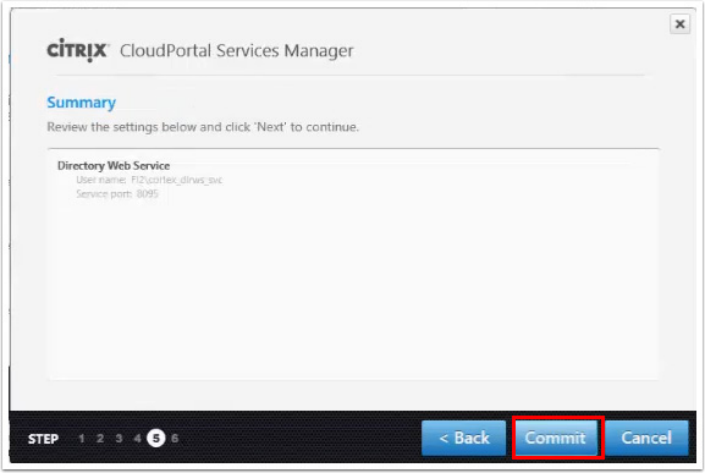
When the configuration completes, click **Finish** to continue.	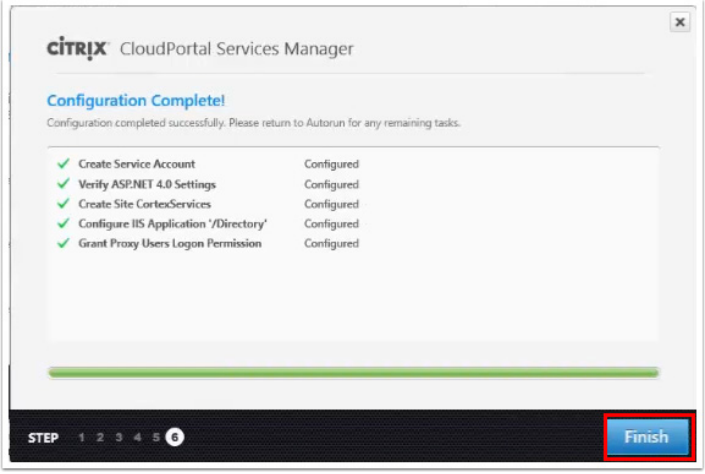

Repeat, this time configuring Provisioning server role. Select **Configure Server Roles**.	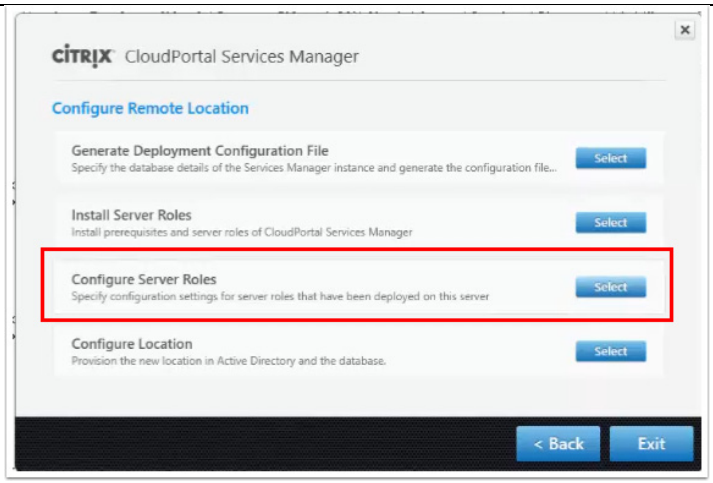
Select **Configure Server Roles**.	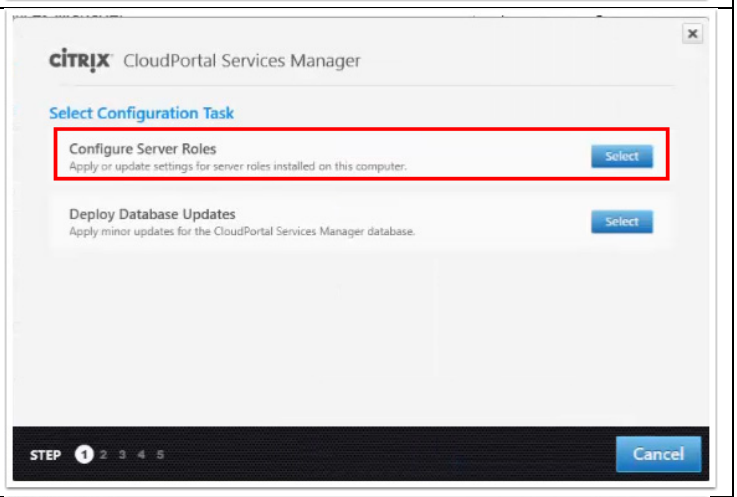
Enter the path to the configuration file (created earlier in the installation process).	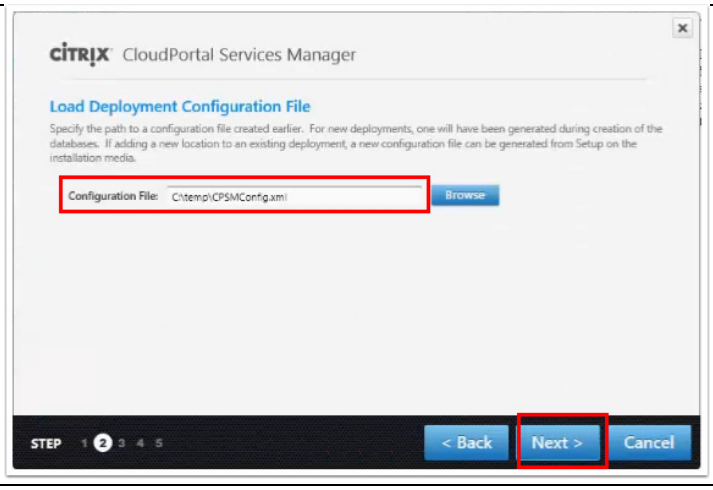

Enable the checkbox for **Provisioning**. Click **Next**.	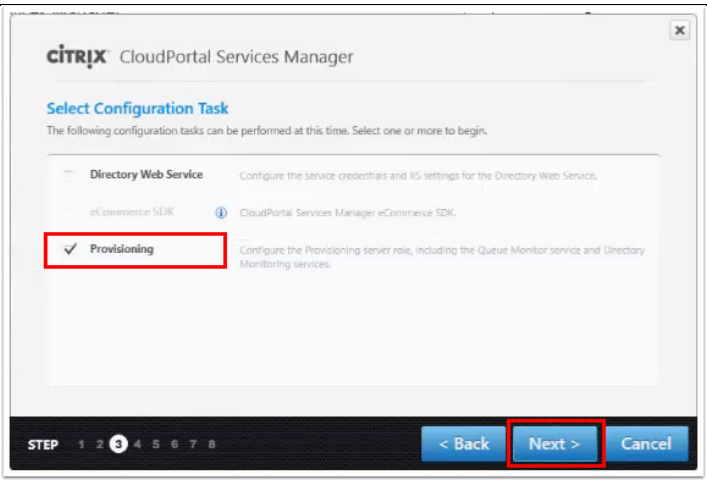
Accept defaults, and then click **Next**.	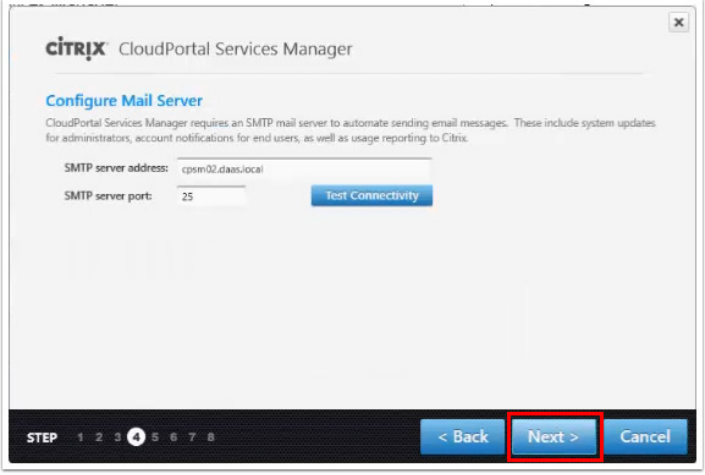
Accept defaults, and then click **Next**.	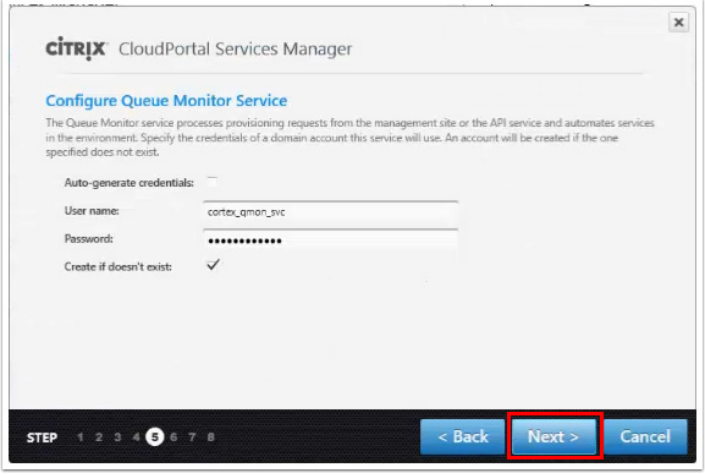

Confirm the information in the **Summary**. Click **Commit**.	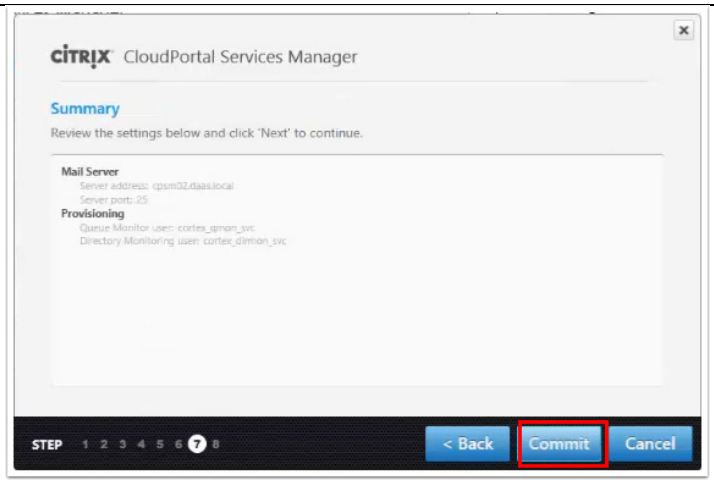
When the configuration completes, click **Finish** to continue.	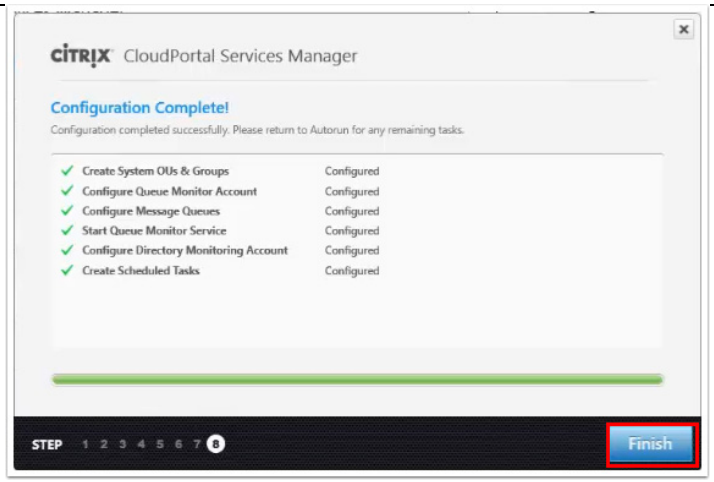

Configure location

Next, provision the location for this new private delivery site in Active Directory and the database. This task associates the new location with the existing Services Manager instance.

Instructions	Visual

In CloudPortal Services Manager, click **Configure Location**.	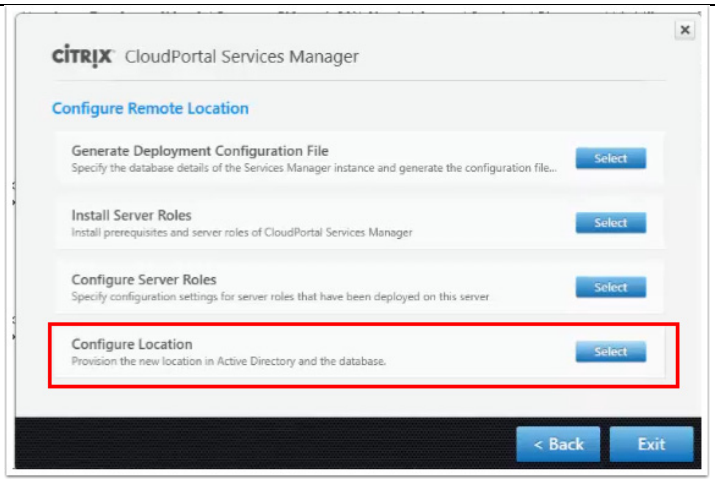
Enter the path to the configuration file (created earlier in the installation process).	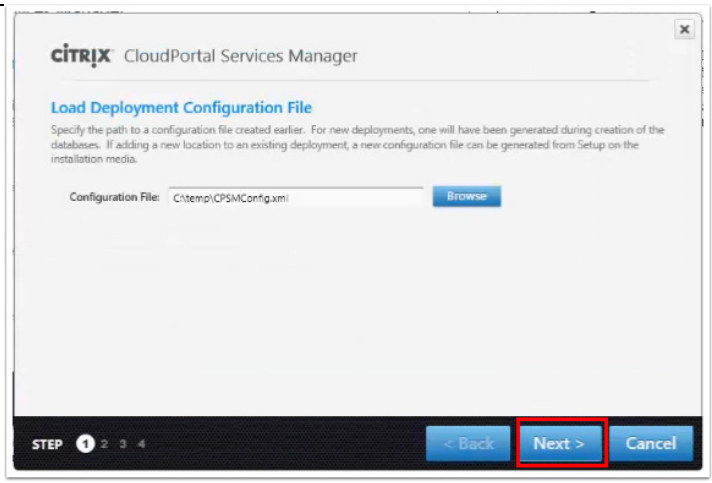
On the **Specify Location Name Details** page, enter the name and description for the private delivery site, and the OU name and display name for the new tenant. Note: content is loaded from previously specified configuration file. If this is not prepared in the file, enter the appropriate values (e.g., Location Settings->Name). Click **Next**.	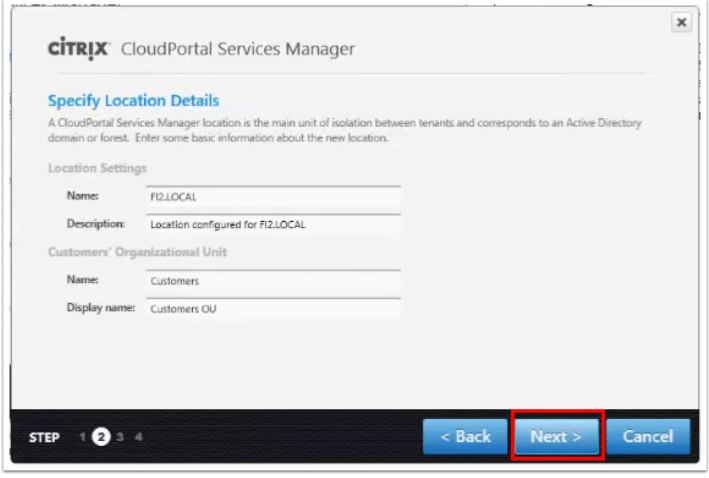

Confirm the information in the **Summary**. Click **Commit**.	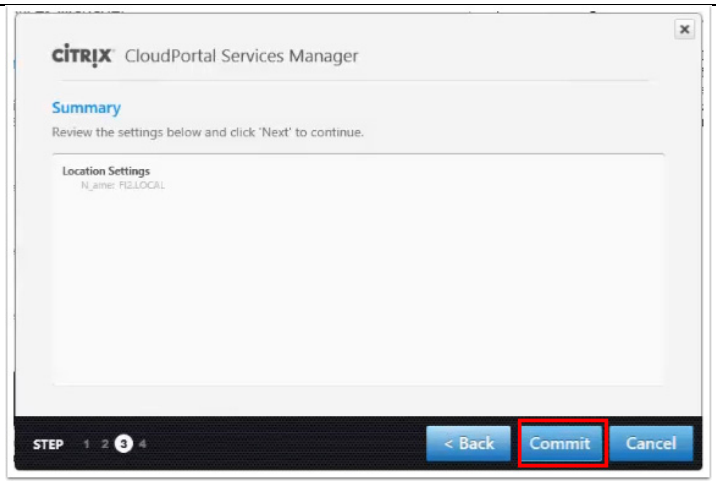
When the configuration completes, click **Finish**.	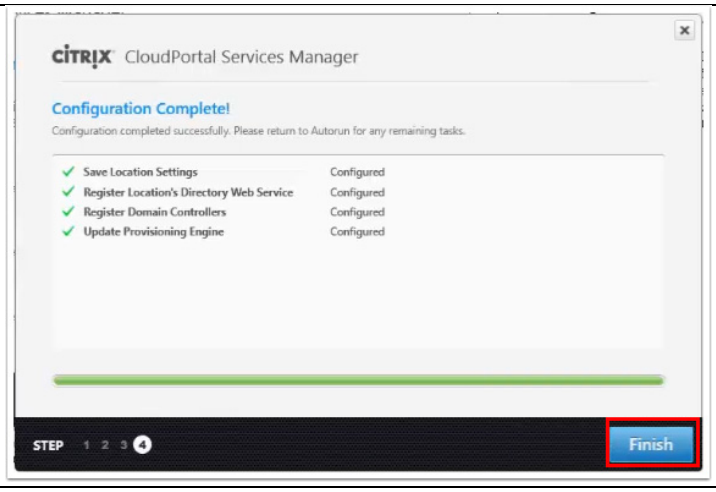

Configure Hosted Apps and Desktops Service

Using CloudPortal Services Manager, configure the server, credentials, and server connections/roles for the Hosted Apps and Desktops service.

Instructions	Visual
Log in to CloudPortal Services Manager, and display the **Service Deployment** page. Filter results by **Active Delivery Location Services** and use the Location Filter of **FI2.LOCAL**. Click on **Hosted Apps and Desktops** to configure the settings for this service. Click on **Customer Plans**. Then, click on **App Orchestration** to configure the settings for this customer plan.	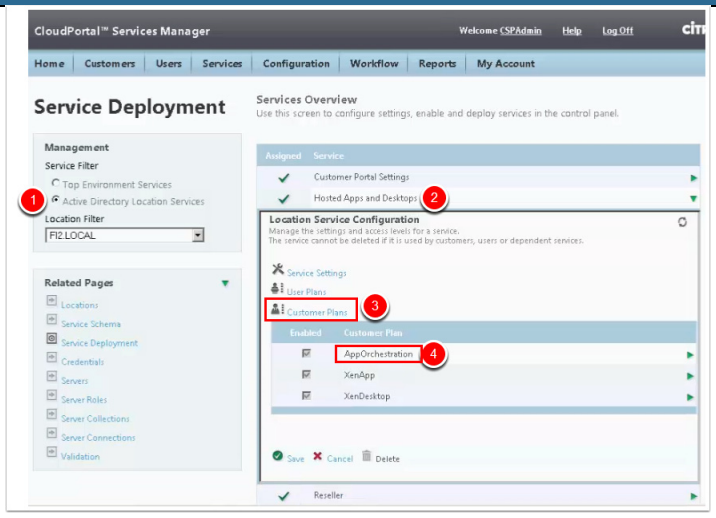

The **Configure Service Settings** panel is displayed. Expand the **App Orchestration** entry, and select **Private Server Group** for the **App Orchestration Store Front Isolation** option. Click **Apply changes**. Then, click **Save**.	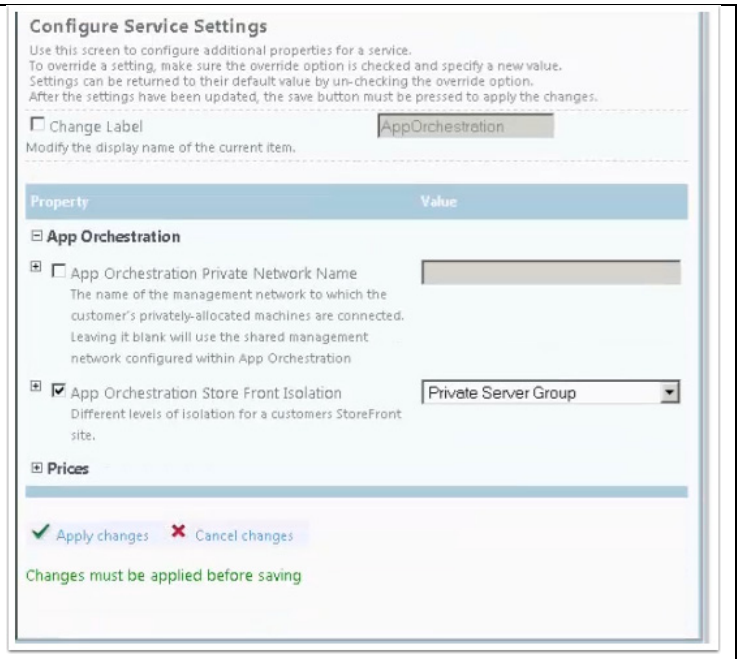
Next, display the **Credentials** page in CloudPortal Services Manager. Click **Add** to add a new credential. The username **csm_haad_selfsvc** is created automatically when installing the CPSM AO Tool in the shared environment on the App Orchestration server. Enter a password for this user and the domain name. (In this CVD, the domain name is **daas**.) Click **Add**.	
Next, display the **Servers** page in CloudPortal Services Manager. Click **Add a server**. Enter the server name (**AO01.daas.local** in this CVD), and click **Add Server**.	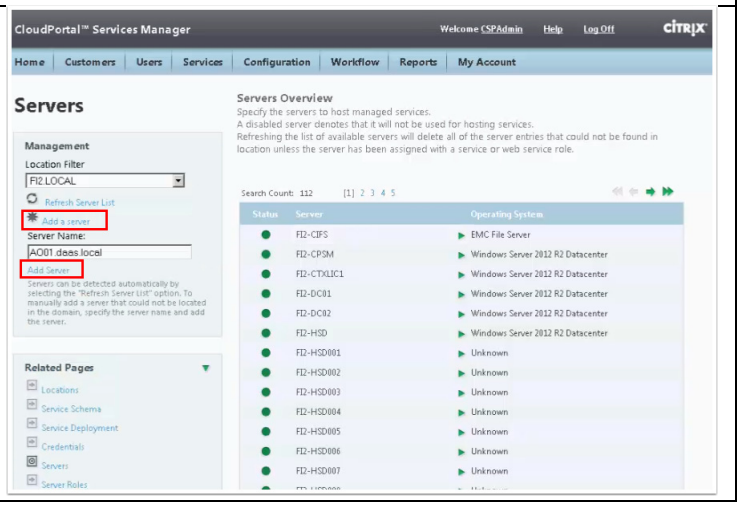

The server AO01.daas.local appears in the list of servers. Click on **AO01.daas.local** to expand this entry. Set the “Alias” to the FQDN of the App Orchestration server (this is critical since communication between CPSM and AO will otherwise fail since certificates are used for securing communication). Scroll down to the bottom of the **Server Setup** panel, and click **Save**.	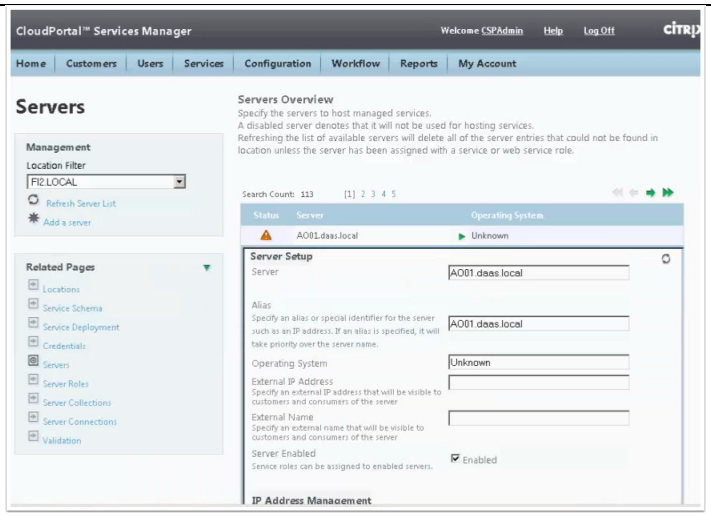
Next, display the **Server Roles** page in CloudPortal Services Manager. Click on the AO01.daas.local entry to expand it. Enable the checkbox for **Hosted Apps And Desktops**. Scroll down, and click **Save**.	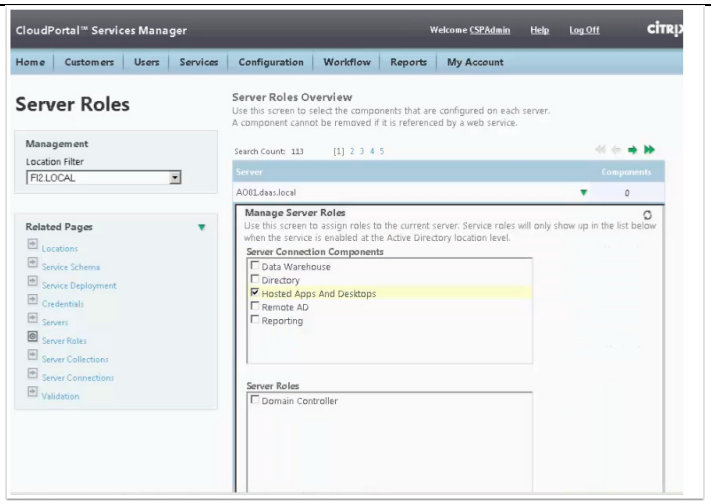
Next, display the **Server Connections** page in CloudPortal Services Manager. Click **New Connection**. Then, in the **Manage Server Connections** panel, select **Hosted Apps and Desktops** as the Server Role. Enter the previously created username in the **Credentials** field, **Port 443**, and select https as the **Protocol** (note that these selections are required). Enable the radio button for **App Orchestration**. Click **Save**.	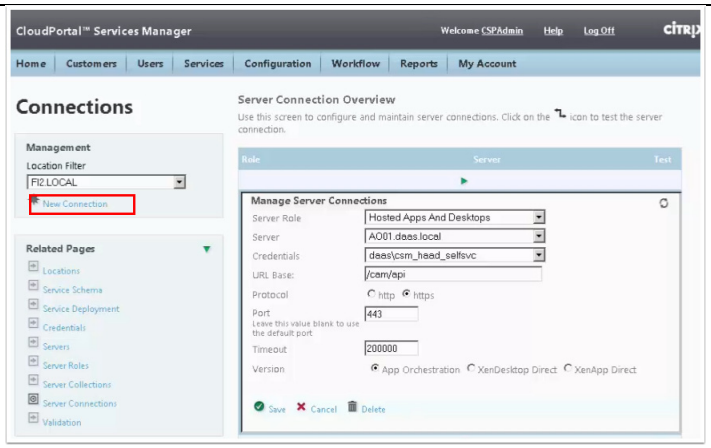
An entry for **Hosted Apps And Desktops – App Orchestration** appears in the Server Connection Overview. Click on the test icon to test the connection. If the connection is successful, the icon changes to a green dot.	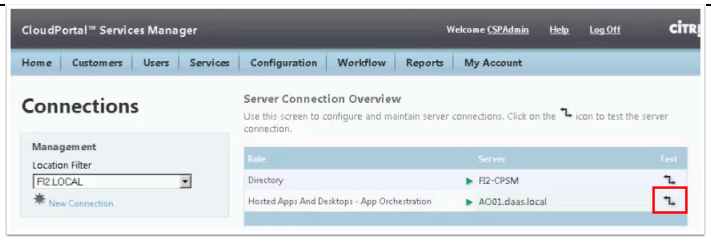

Next, display the **Service Deployment** page in CloudPortal Services Manager. Click on the entry for **Hosted Apps and Desktops** to expand that entry, and then click **Service Settings**.	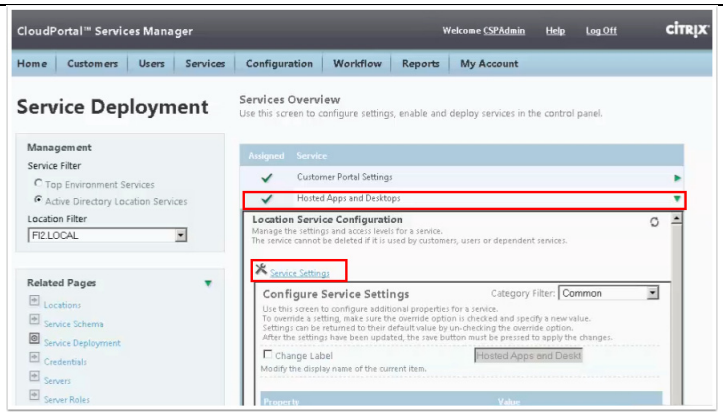
In the **Configure Service Settings** panel, scroll down and expand the **App Orchestration** entry. Select **FarmIso2** as the App Orchestration Datacenter. Click **Apply changes**.	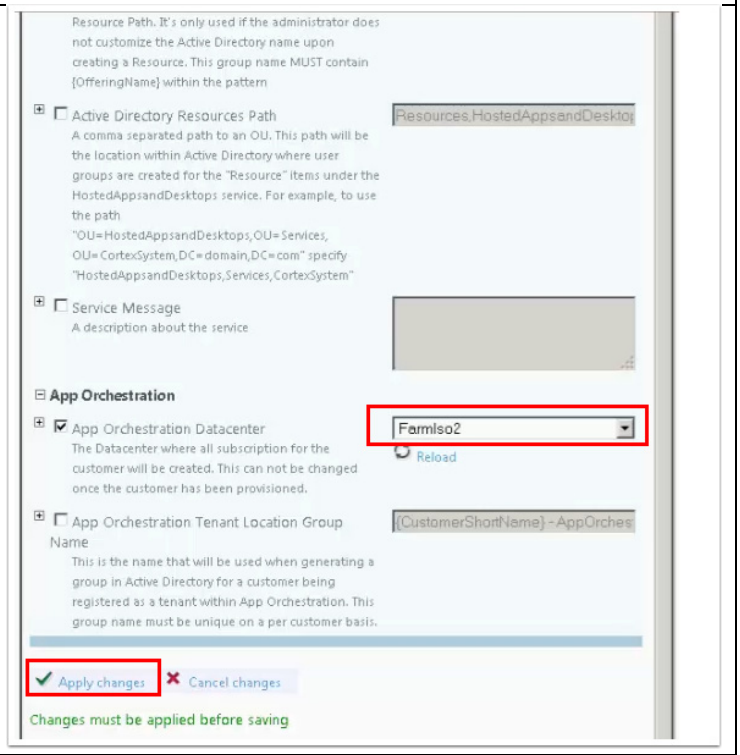
On the Location Service Configuration panel, click **Save**.	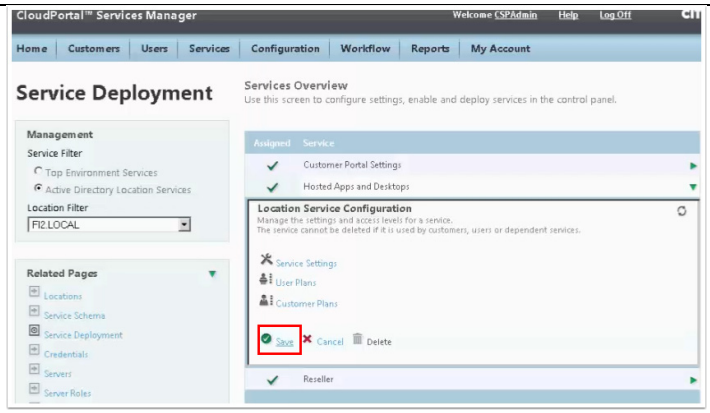

Configure Delivery Sites

App Orchestration is used to create and configure a new private Delivery Site named **FI2-PrivateDS**. A Delivery Site consists of at least two Delivery Controllers (a primary Controller and a backup Controller). In this CVD, servers **fi2-xdc01** and **fi2-xdc02** are used as the Delivery Controllers for the private Delivery Site.

Instructions	Visual
In the App Orchestration web console, display the Delivery Sites. Click **Create new** from the **Manage** menu.	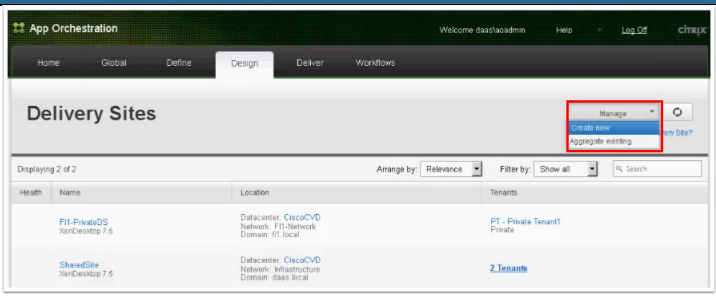
Fill in the **Basic Settings** information. In this example, the new Delivery site is named **FI2-PrivateDS**, and uses XenDesktop 7.5. Click **Next**.	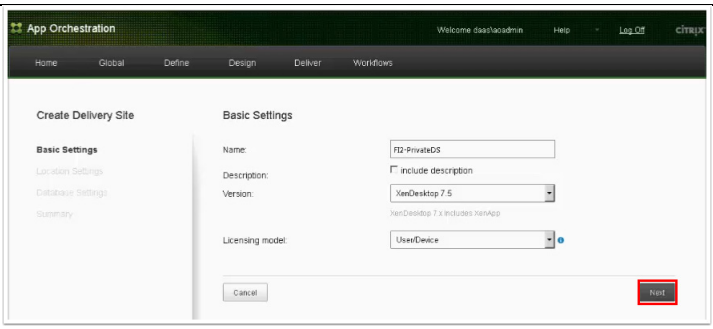
Fill in the **Location Settings** information. Enter the machine name of the two servers to deploy as Delivery Controllers to this Site (**fi2-xdc01**, **fi2-xdc02**), and specify the resource domain (**fi2-local**) and datacenter (**CiscoCVD**) in which they reside. Also enter the administrator group (**SiteAdmins**), user name (**siteadmin**), and password for the Site administrator. Click **Next**.	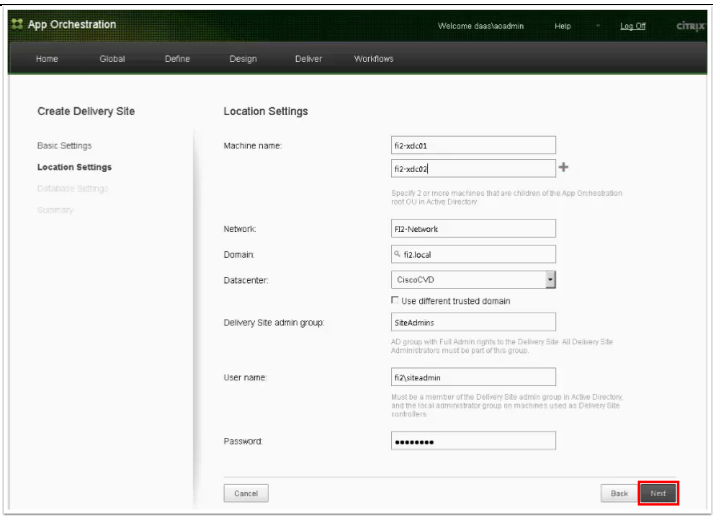

Fill in the **Database Settings**. Enter the database server, database name, and database administrator and password. Click **Next**.	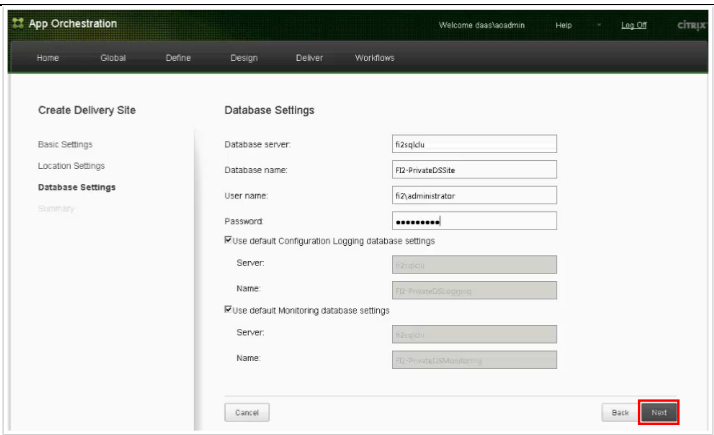
A summary is displayed. Confirm the information you entered is correct. Click **Save.**	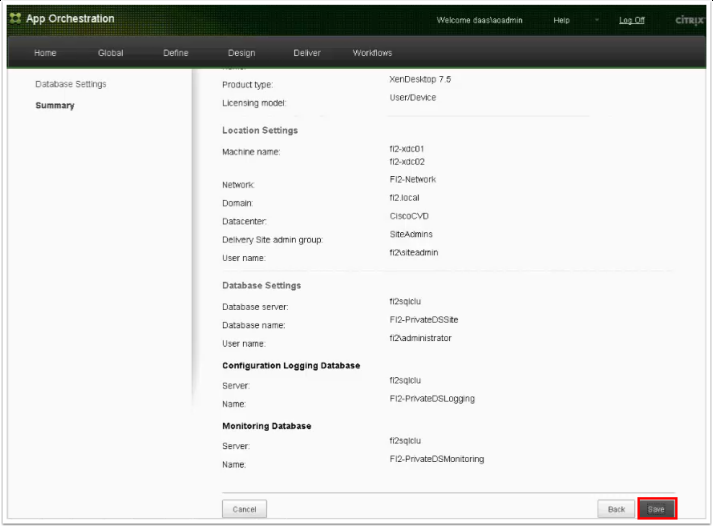
App Orchestration creates the Delivery Site and joins the Delivery Controllers to it. When it completes, the new Delivery Site is displayed in the list of Delivery Sites.	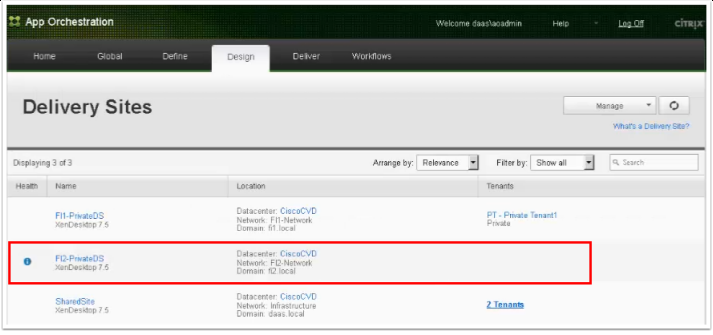

Configure StoreFront Server Groups

This section describes configuration steps required to create a StoreFront Server Group and specify the servers to add to it. A Server Group consists of at least two StoreFront servers. StoreFront is Active/Active using load balancing. In this CVD, a StoreFront **FI2-PrivateSF** is created using servers **fi2-sf01** and **fi2-sf02**.

Instructions	Visual

In the App Orchestration web console, display the StoreFront Server Groups. Click the **Add StoreFront Server Group** button.	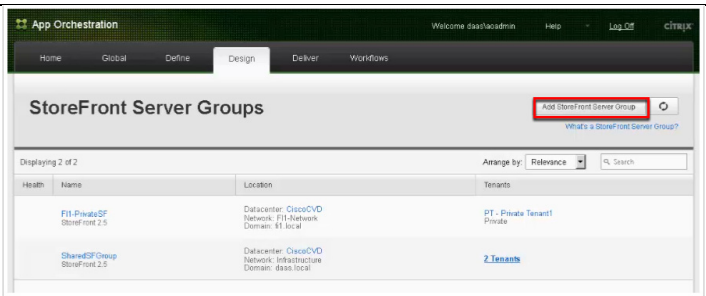
Fill in the **Basic Settings** information. Enter the name for the StoreFront (**FI2-PrivateSF**), SSL certificate information, and load balancer URL. Click **Next**.	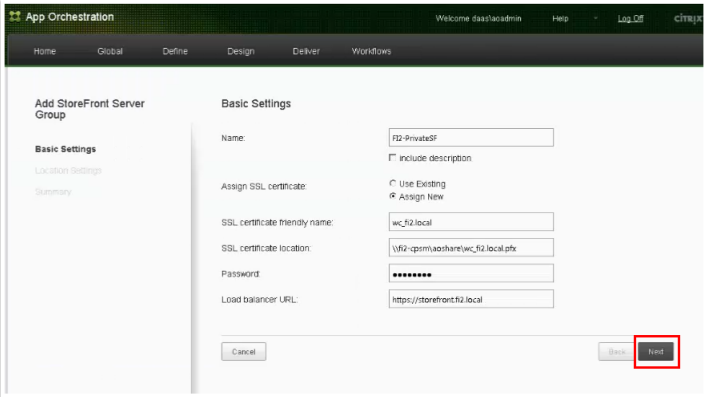
Fill in the **Location Settings** information. Enter the names of the two StoreFront servers (**fi2-sf01**, **fi2-sf02**) and the network (**FI2-Network**), resource domain (**fi2.local**), and datacenter (**CiscoCVD**) in which they reside. Click **Next**.	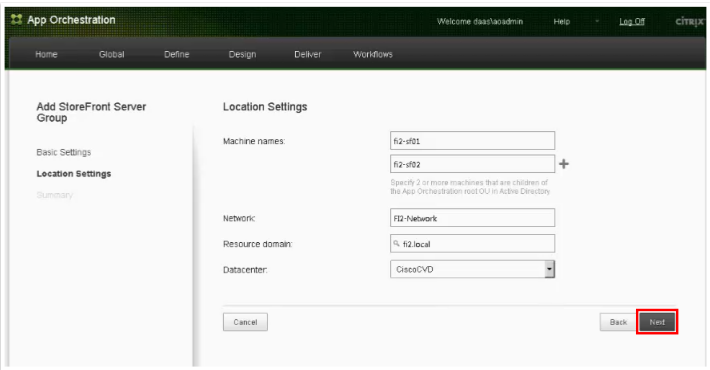
A summary is displayed. Confirm the information you entered is correct. Click **Save.**	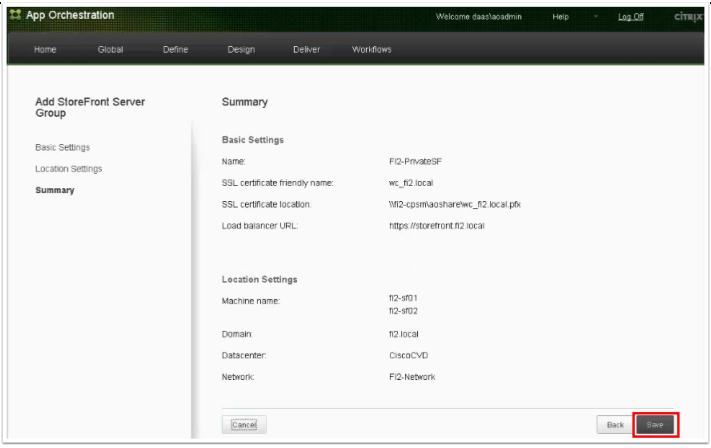

App Orchestration creates the new server group and adds the StoreFront servers to it. When this process completes, the new StoreFront Server Group is displayed in the list of Server Groups.	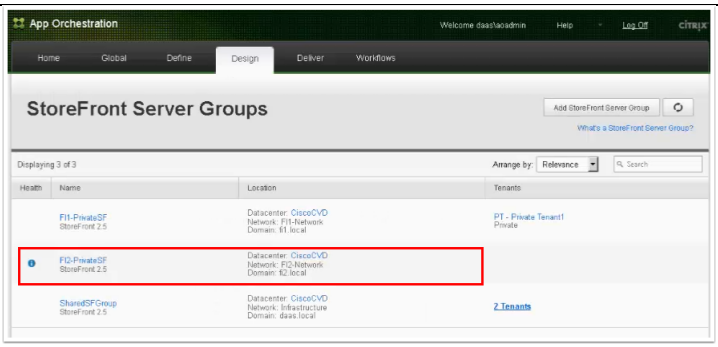

Configure a Session Machine Catalog

App Orchestration is used to create and populate two Session Machine Catalogs. The first, **FI2-IsolatedHSD**, is used for delivering shared desktops to users in the private delivery site. The second, **FI2-IsolatedSVDI**, is used for delivering VDI sessions in the private delivery site.

Instructions	Visual
In the App Orchestration web console, display the current Session Machine Catalogs. Click **Created externally** from **Use Machines** menu.	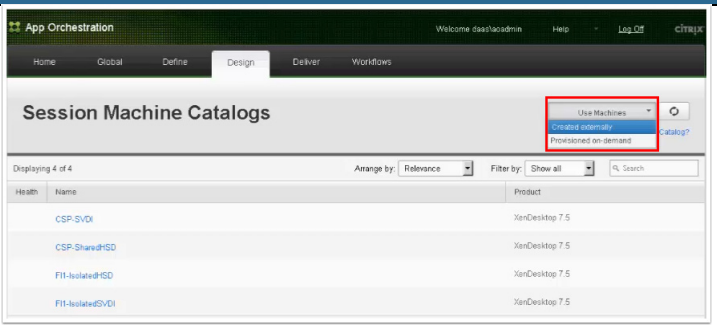
First, create the catalog for HSD. Fill in the **Basic Settings** information, specifying a name (**FI2-IsolatedHSD**) for the Session Machine Catalog and enabling the **Multi User** radio button. Click **Next**.	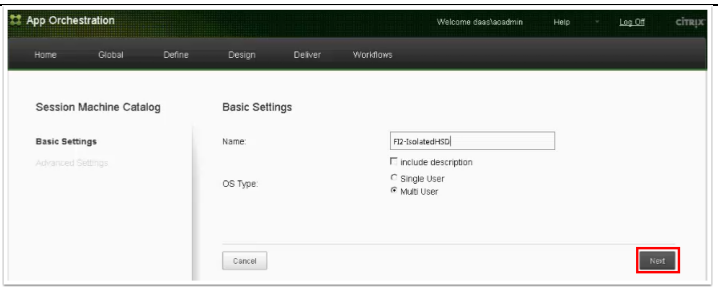
Fill in the **Advanced Settings** information. Change the **Number of users allowed per machine** to 50. Note: this value of 50 is an example value only; this value strongly depends on the sizing of the VMs. Click **Save**.	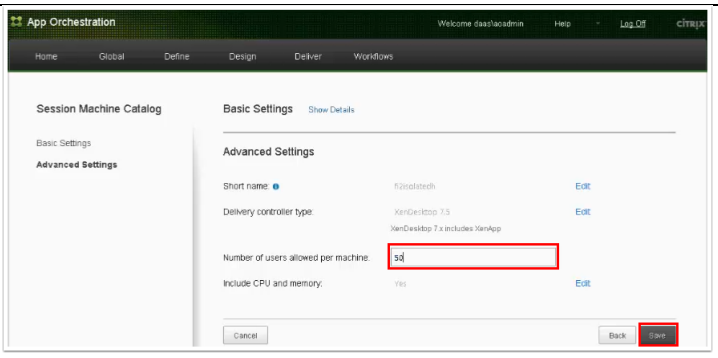

Repeat, this time creating the catalog for VDI sessions. Click **Created externally** from **Use Machines** menu.	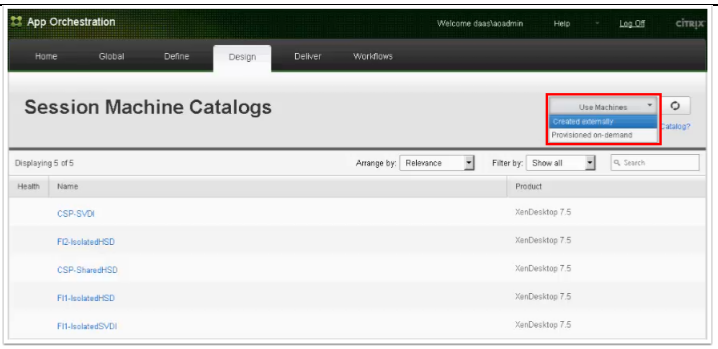
Fill in the **Basic Settings** information, specifying a name (**FI2-IsolatedSVDI**) for the Session Machine Catalog and enabling the **Single User** radio button. Click **Next**.	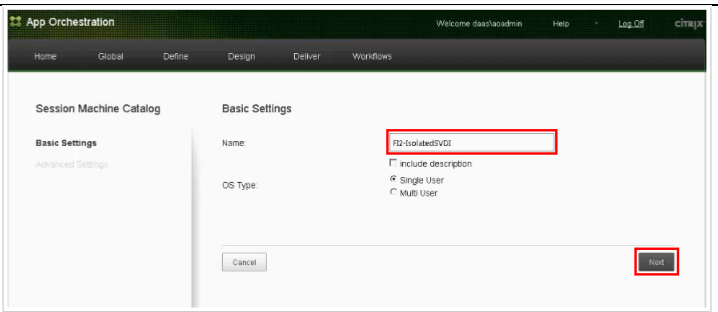
Accept the defaults on the **Advanced Settings** dialog. Click **Save**.	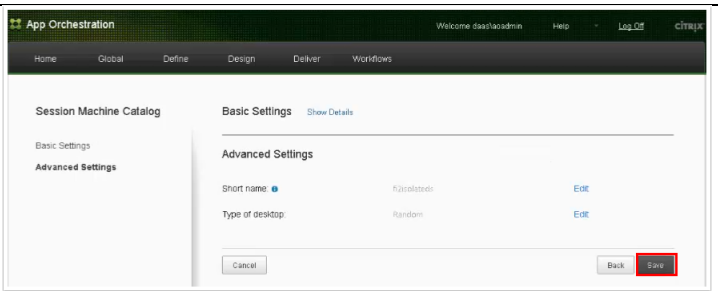
Now that the catalogs are created, add session machines to them. Select the **FI2-IsolatedHSD** entry in the Session Machine Catalogs list.	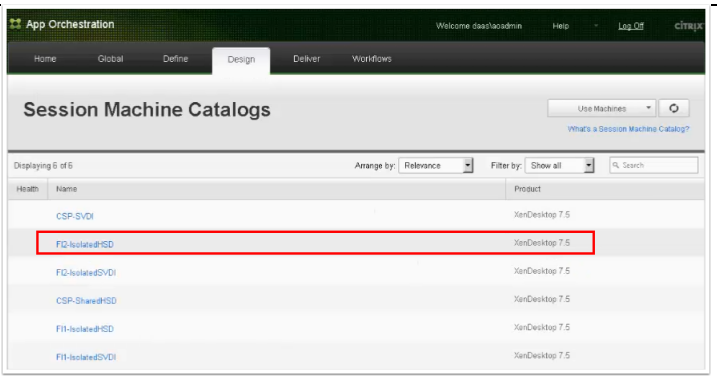
Click the **Add Machines** button.	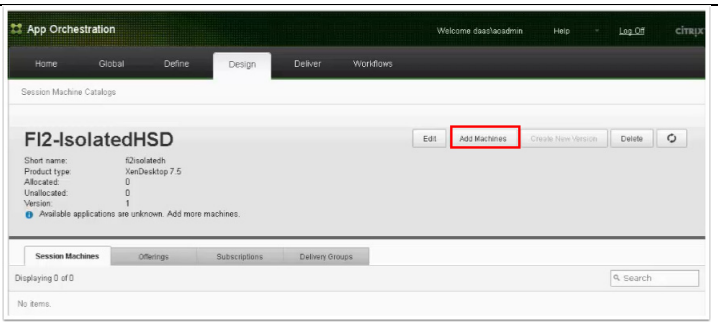

Enter information in the **Basic Settings** dialog, including the machine names, network, and resource domain and datacenter in which these session machines reside. Click **Save**.	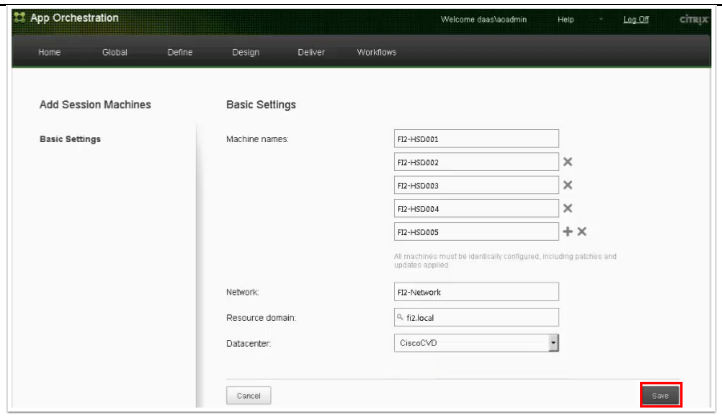
Repeat, this time adding machines to the SVDI catalog. Select the **FI2-IsolatedSVDI** entry from the Session Machine Catalogs list.	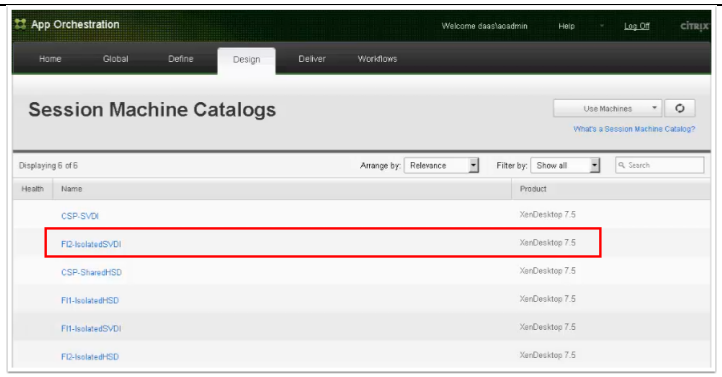
Click the **Add Machines** button.	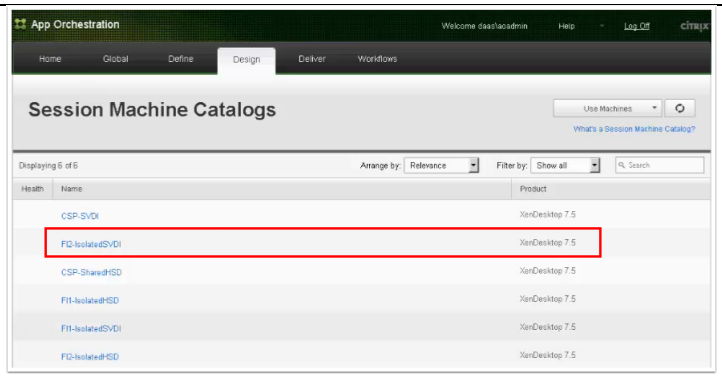
Enter information in the **Basic Settings** dialog, including the machine names. Click **Save**.	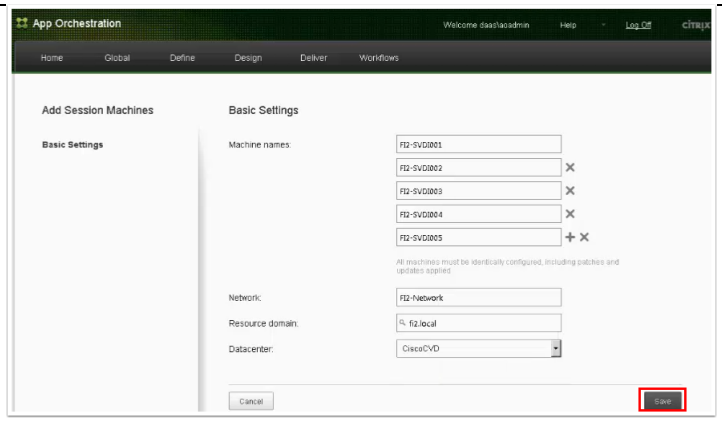

Configure Offerings

This section describes how to make applications and desktops (hosted on the Session Machines) available for tenant subscription. In this example, new offerings are added to the FI2-IsolatedHSD Session Machine Catalog. Then, a new customer (tenant) is added and given access to these offerings.

Create New Offering

To create offerings, App Orchestration is used to specify the applications and desktops you want to include and the isolation level at which you want to provide the offering to tenants. In this example, a new offering is created in the FI2-IsolatedHSD session machine catalog.

Instructions	Visual
In the App Orchestration web console, click **Design->Offerings** to display the current offerings. Then, from the **Create** menu, select **Desktop Offering.**	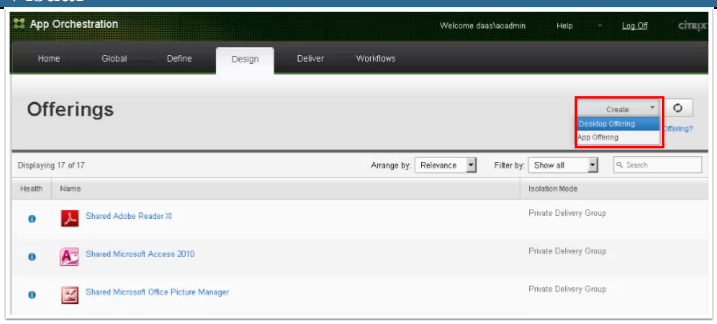
The Session Machine Catalog is displayed. Enable the radio button for the **FI2-IsolatedHSD** Session Machine Catalog. Click **Next**.	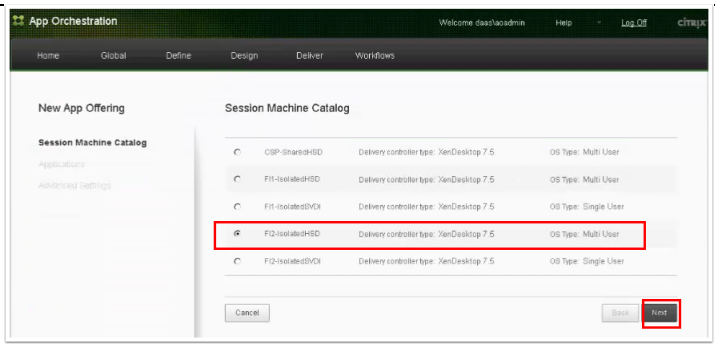
Scroll through the list of applications, and enable the checkboxes for the applications you want to include in this offering. Click **Next**.	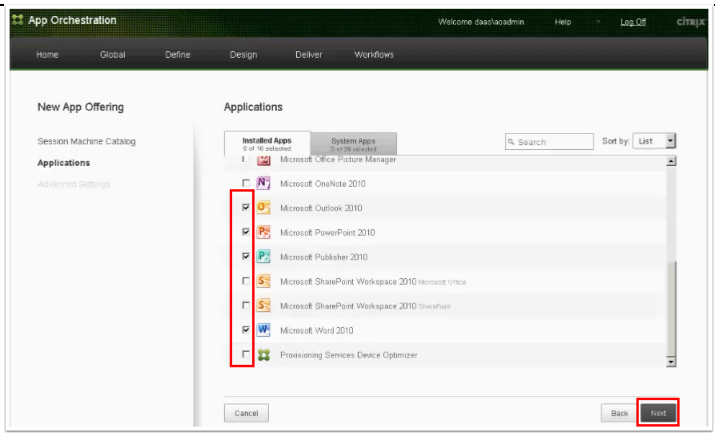

In the **Advanced Settings**, first edit the **Isolation mode** and change it to **PrivateDelivery Site**. Then, edit the names to be displayed to administrators. For convenience and clarity, “FI2 HSD” is prepended to the name of each application. Click **Save**.	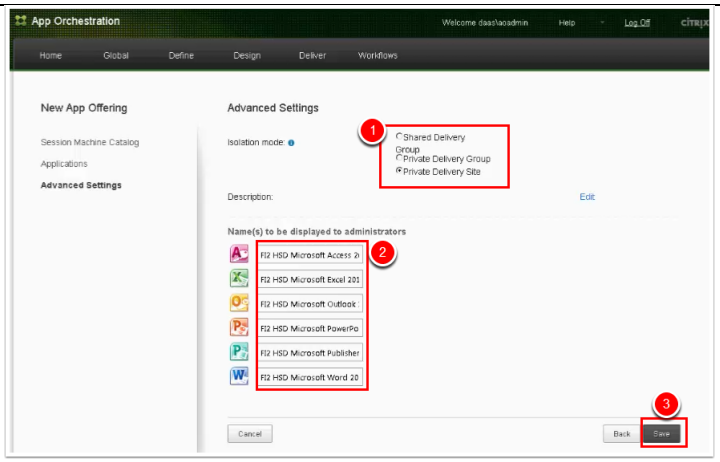
The newly created offerings display in the list of Offerings.	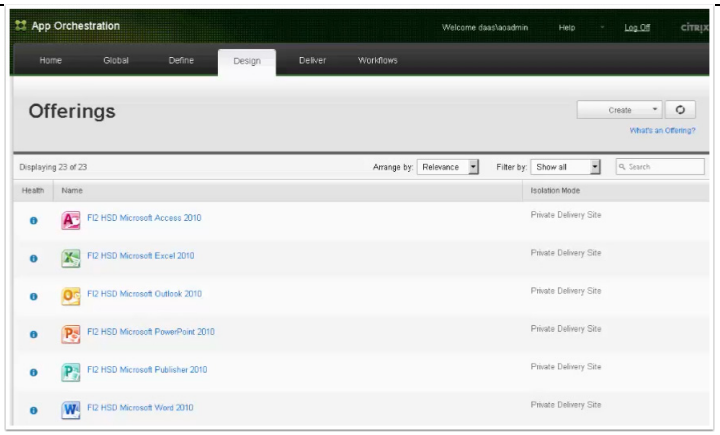

Using CloudPortal Services Manager, view the offerings for the FI2.LOCAL domain.

Instructions	Visual
From the **Services** tab of CloudPortal Services Manager, select **Hosted Apps and Desktops** and then **Offering Management.**	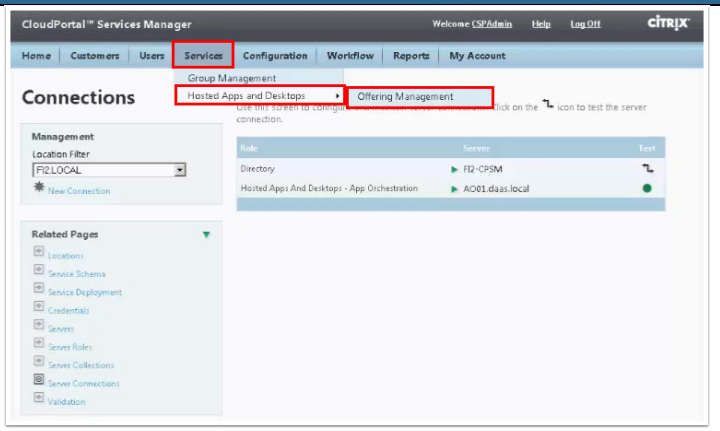

Select the FI2.LOCAL from the **Location** menu to display the hosted apps and desktops for this location. Note that there are two offerings for Excel. Because they have the same name, it's not readily apparent which offering is for the FI2.Local location/HSD.	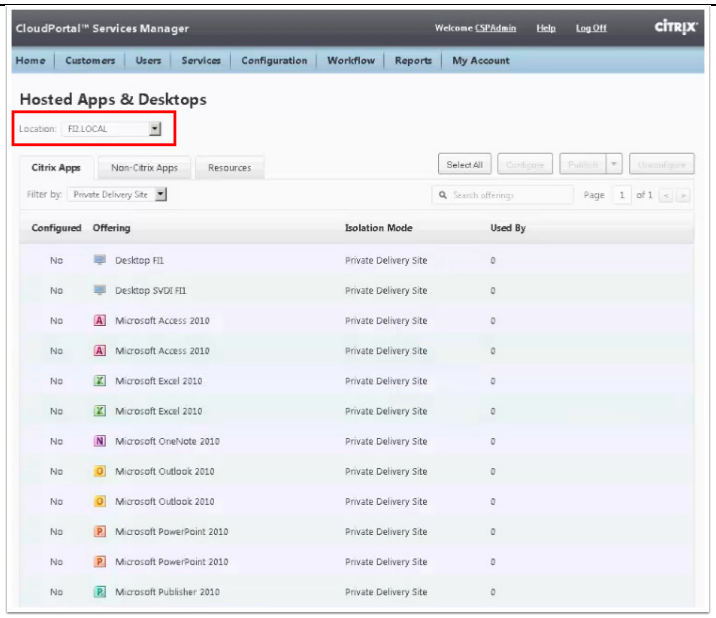

Edit Offering

In this section, the newly created offerings are edited so that they have names that are more easily associated with the FI2-IsolatedHSD catalog. This is also useful for transparency in later CloudPortal Services Manager use.

Instructions	Visual
In the App Orchestration web console **Offerings** list, select the **FI2 HSD Microsoft Access 2010** offering.	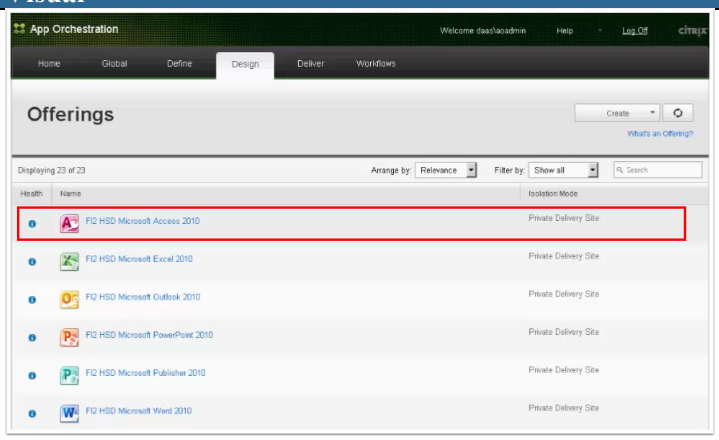
Click on the **Edit** button to edit this offering.	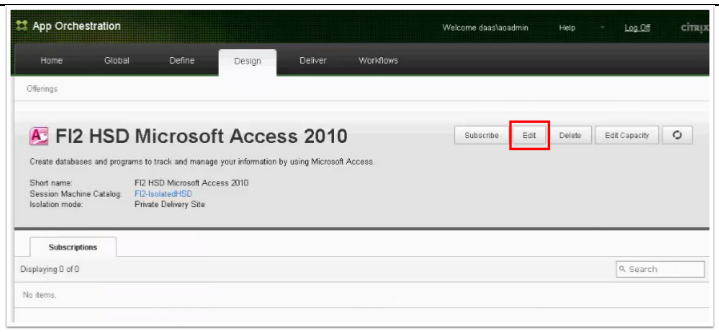

Edit the settings for this offering. In this CVD, the Display name was changed to include “FI2 HSD”. For example, “Microsoft Access 2010” was changed to “FI2 HSD Microsoft Access 2010”. Click **Save Offering**. Repeat this step for all other offerings. In this CVD, that includes Microsoft Excel, Outlook, PowerPoint, Publisher, and Word.	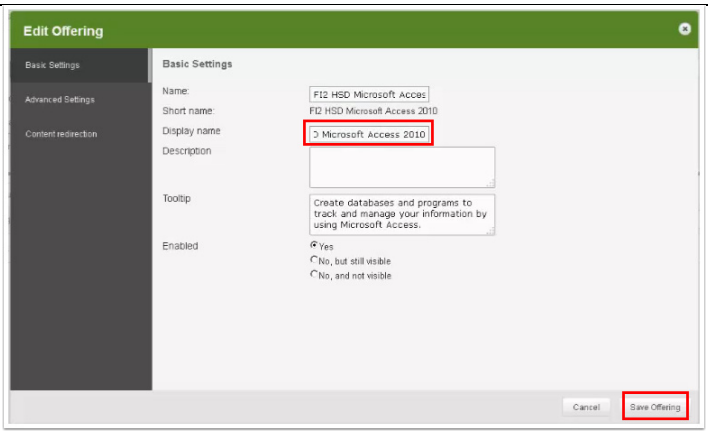

Create New Customer

This section illustrated creating a new customer (tenant). To add a tenant, you create the tenant, create an administrator for this tenant, and provision services.

Instructions	Visual
In CPSM, select **Customers** from the **Customers** tab to display the current list of customers.	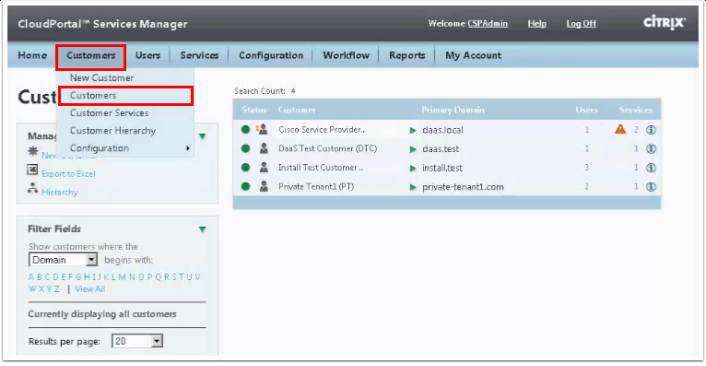
Click **New Customer**.	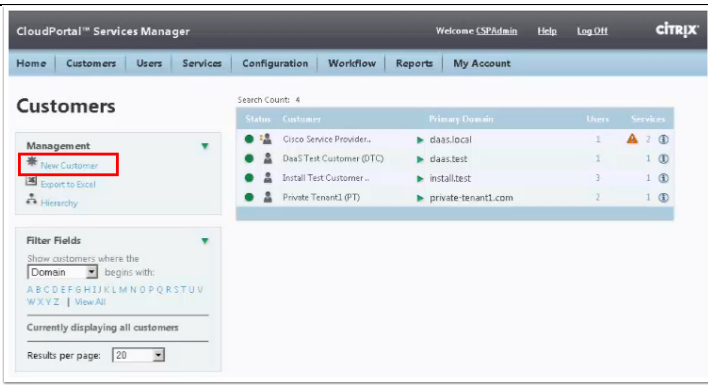

Enter **Customer Details**. In the **Domain Management** section, enter the domain name for this client. Click **Update**. Then, click **Provision**.	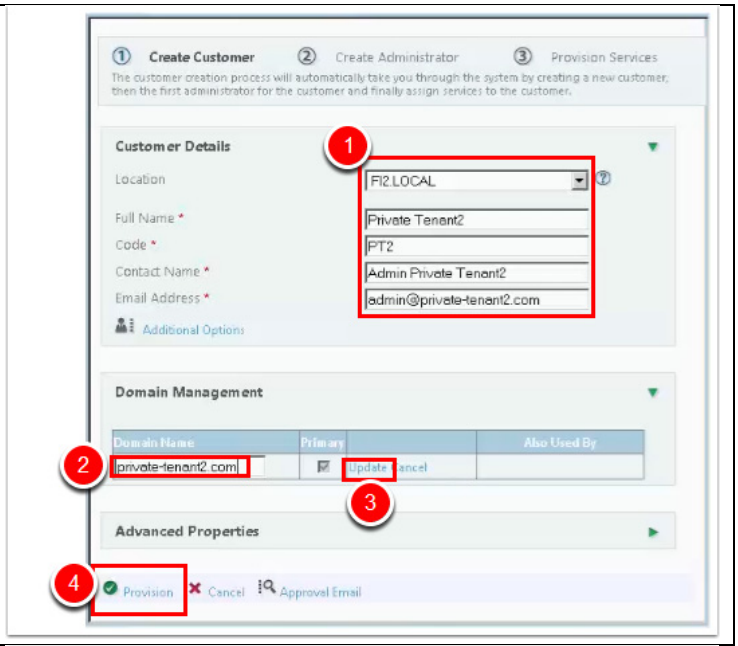
Next, create the administrator for this new customer. Enter user details and password. Click **Provision**.	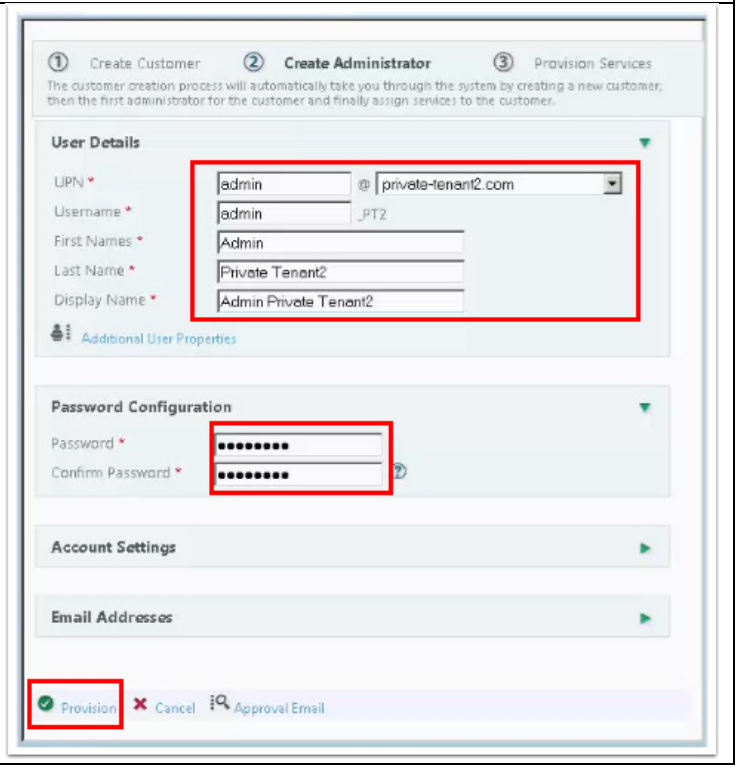

Click on **Hosted Apps and Desktops** to configure.	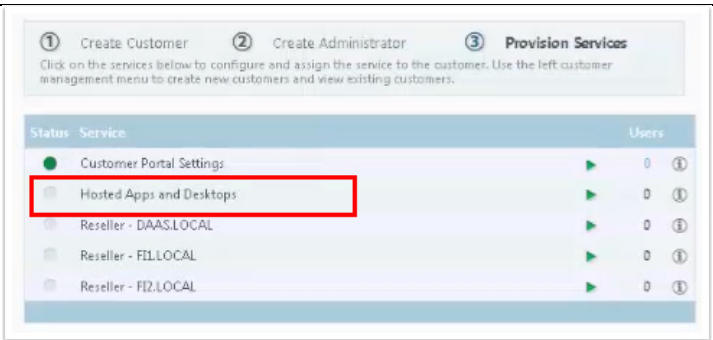
Click on **Advanced Settings**.	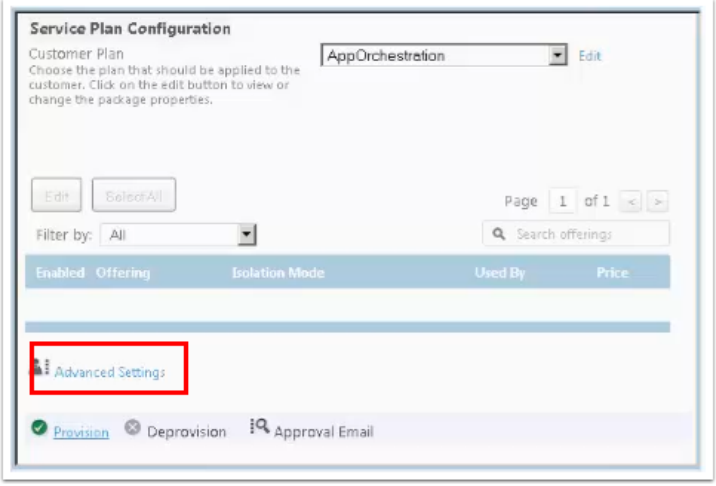
Expand the **App Orchestration** item, and click the checkbox for **App Orchestration Store Front Isolation**. Choose **Private Server Group** from the menu. Click **Apply Changes**.	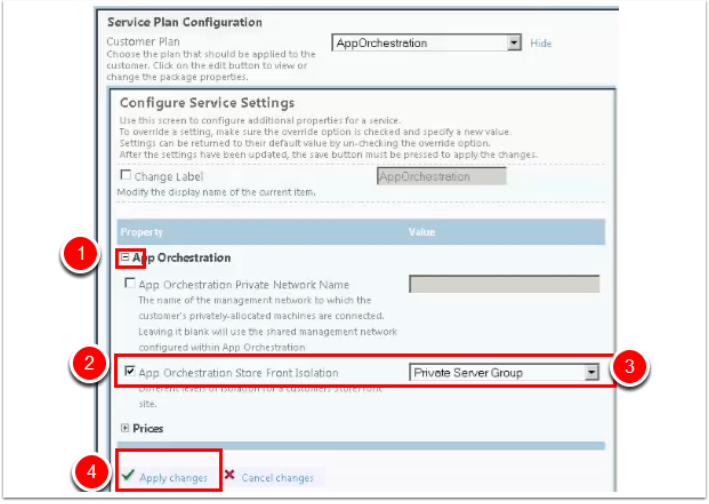

Click **Provision**.	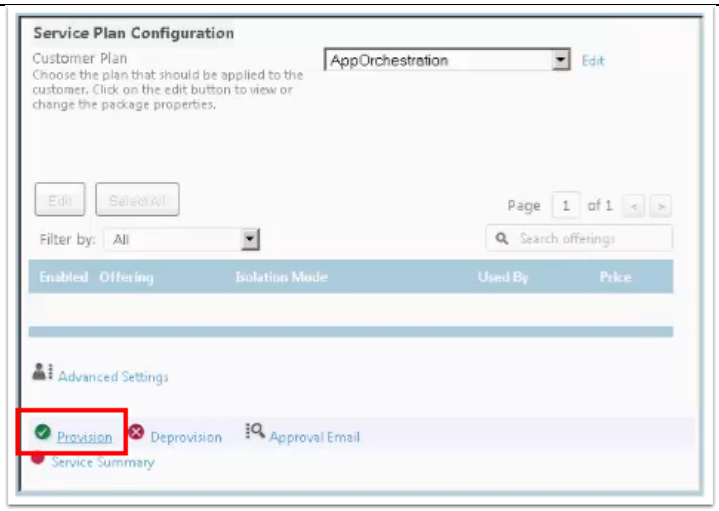
When the provisioning for hosted apps and desktops completes successfully, the status light turns green.	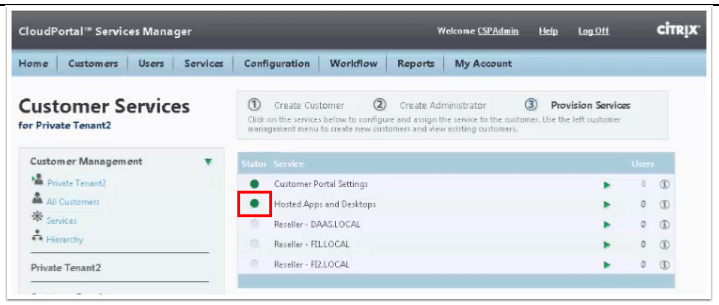

Configure Private Tenant

After the new tenant is created, use App Orchestration to configure it. In this example, we configure the new tenant to use private network isolation.

Instructions	Visual
In the App Orchestration web console, click **Deliver->Tenants** to display the list of tenants. Click on the entry for **PT2 – Private Tenant2**.	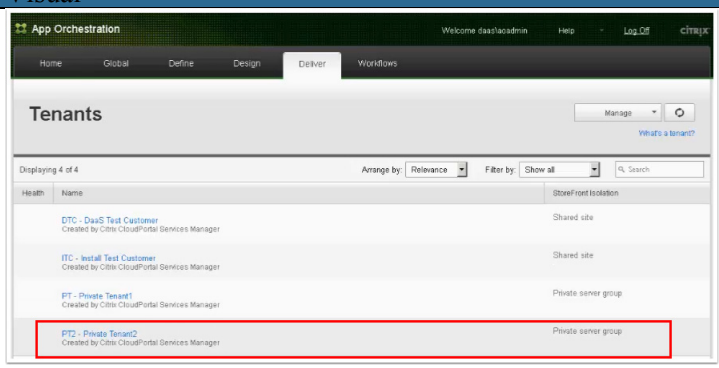

Click on the **Edit** button.	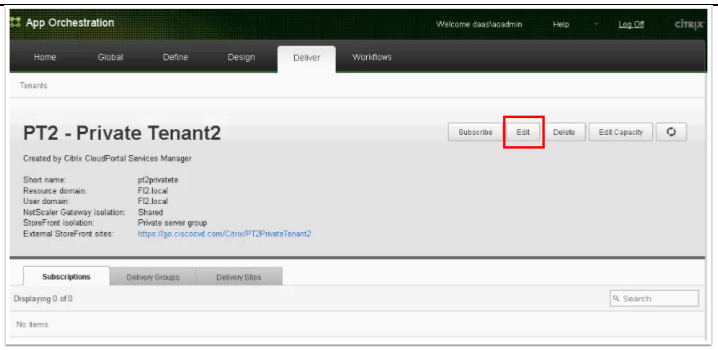
On the **Network Isolation** tab, enable the checkbox for **Use private network isolation** and enter the name for the **Private management network**. Click **Save**.	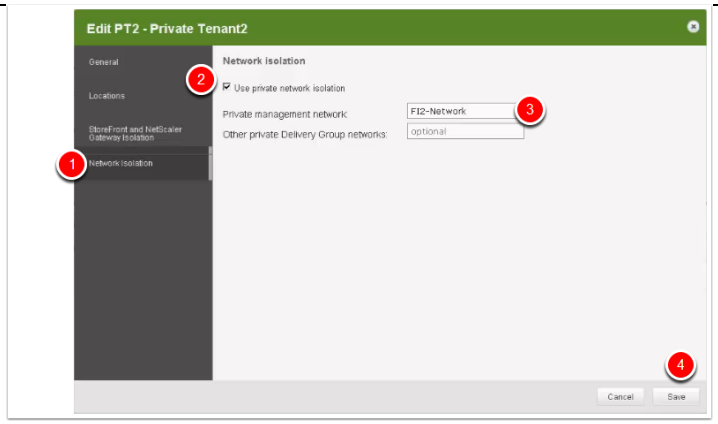
In CPSM, Select **Hosted Apps and Desktops** and then **Offering Management** from the **Services** tab.	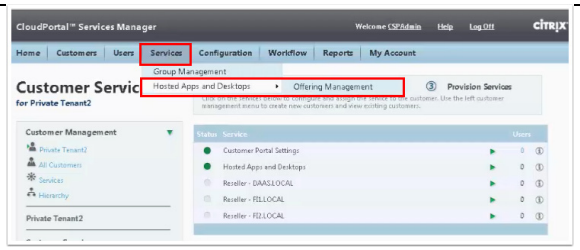
Select **FI2.LOCAL** from the **Location** menu to display the hosted apps and desktops for that location. Then click on the entry for **FI2 HSD Microsoft Access 2010** to edit that offering.	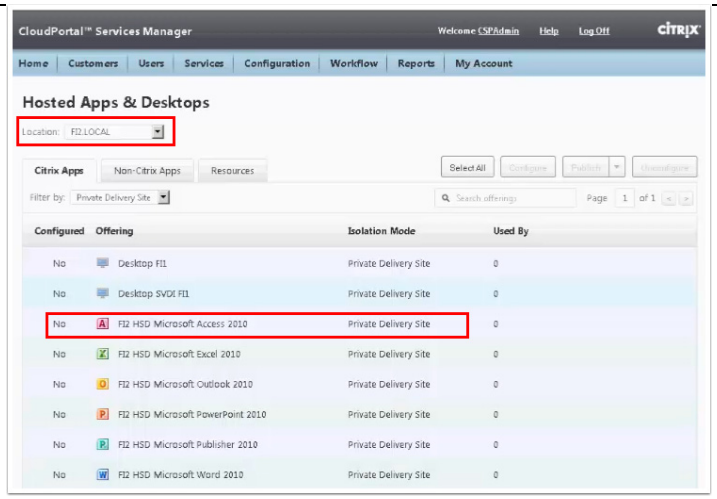

Enable the checkbox for **Mark as private application**. Enter the customer name ("Private Tenant2" in this example). Click **Save**.	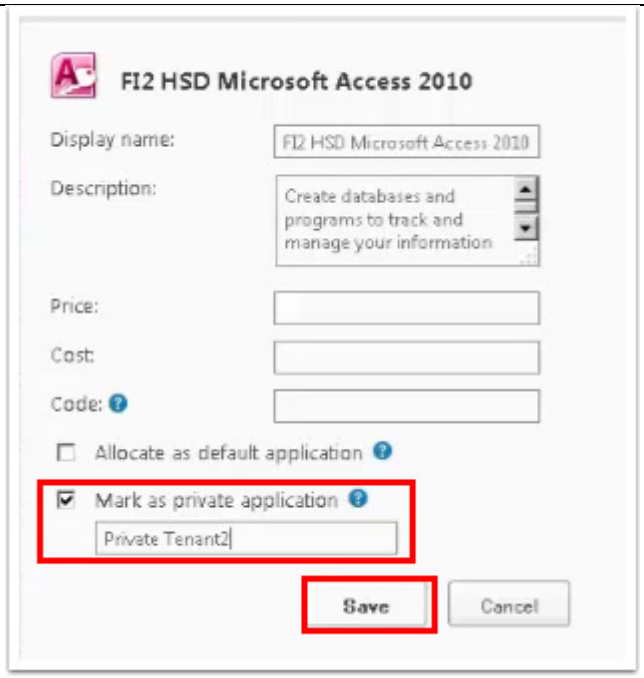
Now, the entry for the FI2 HSD Microsoft Access 2010 offering shows that it is configured and used by Private Tenant2. Repeat this configuration for other offerings. In this CVD, that includes Microsoft Excel, Outlook, PowerPoint, Publisher, and Word.	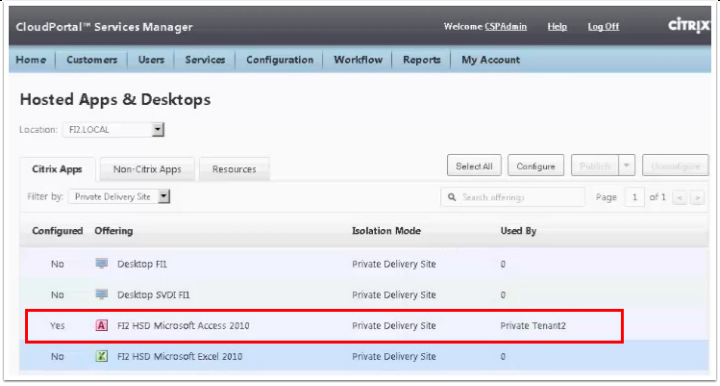

Configure Customer Services for Private Tenant

Before provisioning the offerings to the customer, the following steps confirm that the offerings are provisioned to the customer.

Instructions	Visual

Display the customers in location FI2.LOCAL. For the customer Private Tenant2 (PT2), click **Services**.	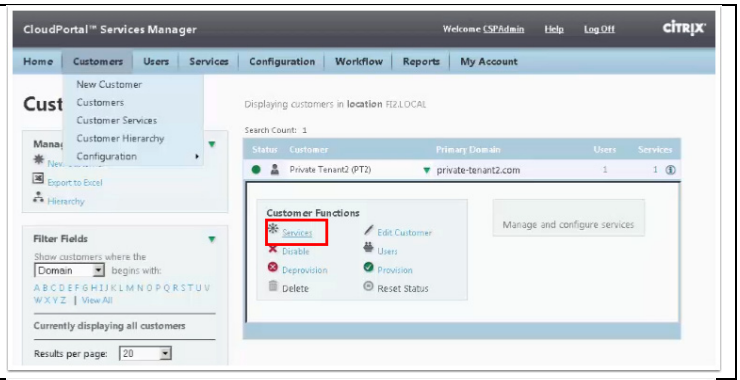
The Customer Services for Private Tenant2 are displayed. Select **Hosted Apps and Desktops**.	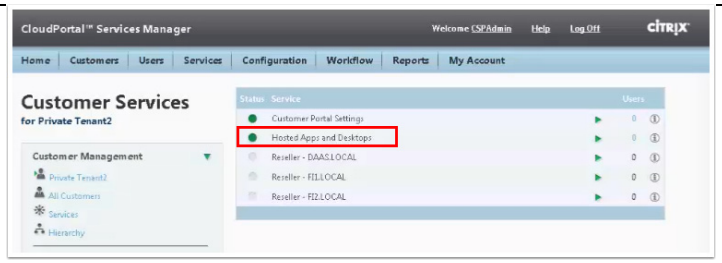
The Service Plan Configuration is displayed. Verify the offerings are enabled and using Isolation Mode **Private Delivery Site**. Click **Provision**.	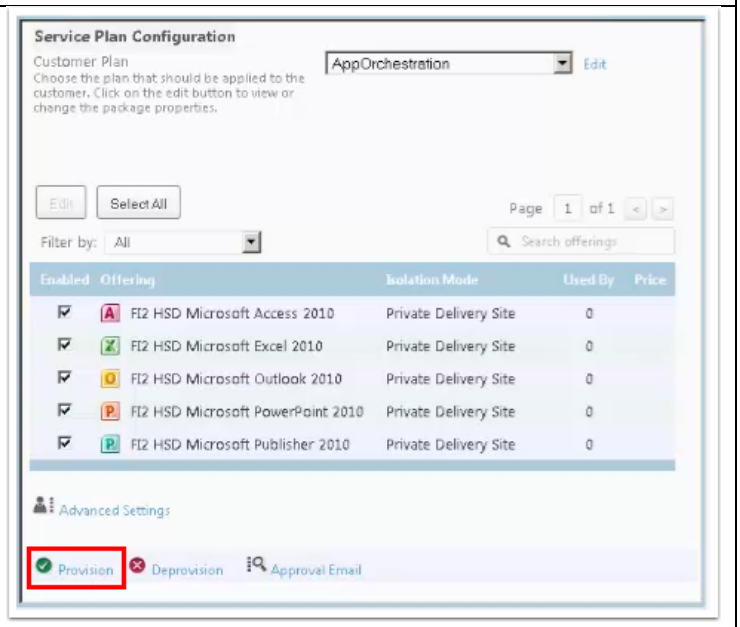

Configure Users for Private Tenant

Next, add and configure a user for Private Tenant2.

Instructions	Visual
In CPSM, display the Customer Services for Private Tenant2. Click **Private Tenant2**.	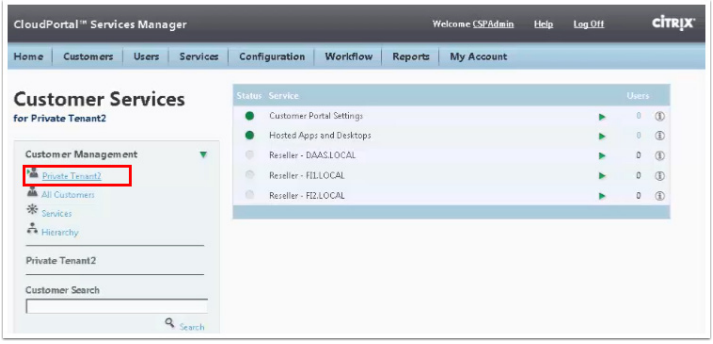
Information about Private Tenant2 is displayed. Click **Users**.	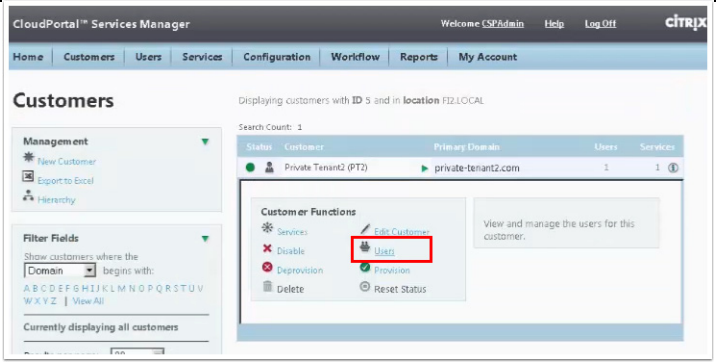
Click **New User** to start adding a new user.	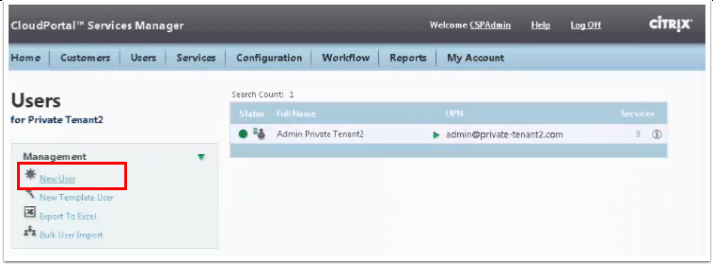

Enter user details and password information. Click **Provision**.	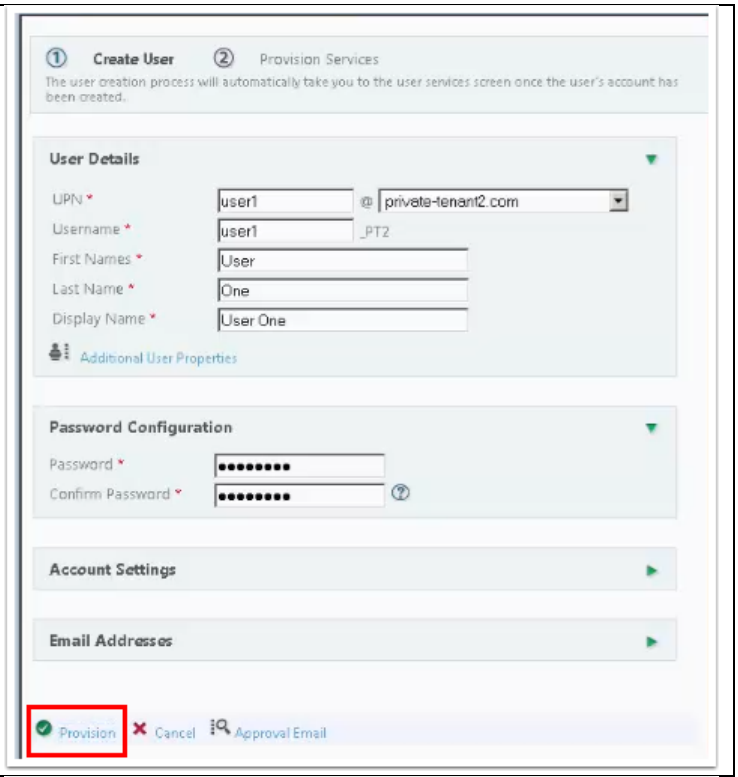
The User Service Setup screen is displayed. Enable the checkboxes for the offerings this user will subscribe to. Click **Provision**.	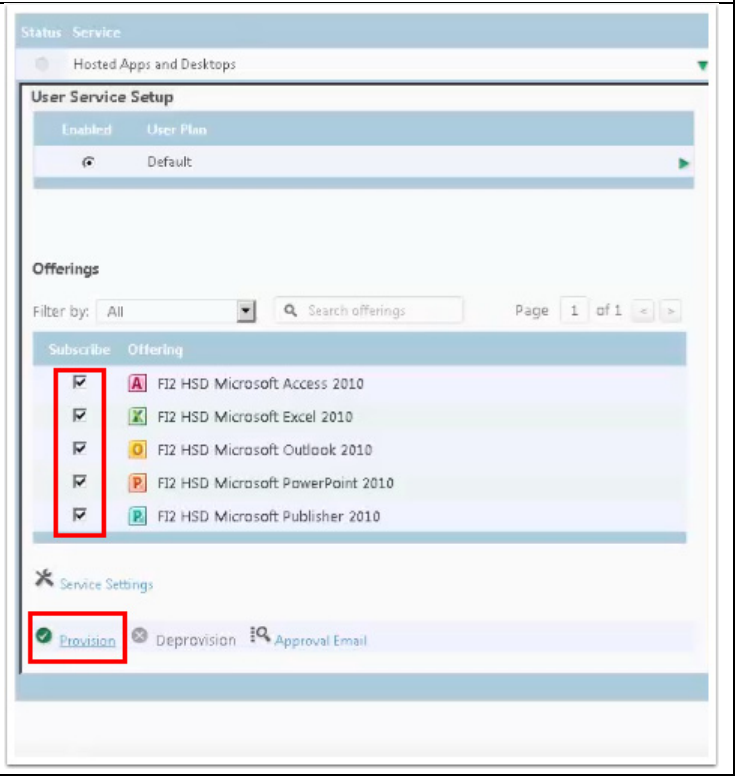

Test Setup, Configurations, and Results

In this project, we tested the functionality of provisioning 10 separate tenants with differing tenant models ranging from a shared infrastructure and workload server model to completely isolated tenants running on their own dedicated hardware and infrastructure.

Login VSI

We utilized Login VSI version 3.7 to configure a test that allowed us to have 10 separate tenants with up to 2000 users login simultaneously. We used the custom command .csv file option to have separate users from separate domains able to participate in a single test.

Since this was not a performance load test, no CPU or VSIMax graphs will be included in this paper.

Cisco UCS Test Configuration for Testing Tenant Offering Access

For our testing we provisioned our users as follows.

- 4 tenants with 50 users each to share the common infrastructure.
- 1 tenant with 100 users to share the common infrastructure.
- 2 tenants with 150 users each to share the common infrastructure.
- 1 tenant with 200 users in a private isolation model with dedicated hosts and infrastructure.
- 1 tenant with 500 users running on shared infrastructure but dedicated hardware.
- 1 tenant with 700 users in a private isolation model with dedicated hosts and infrastructure.

Hardware components

- 2 X Cisco UCS B200-M3 (E5-2650v2) blade servers with 256 GB of memory (16 GB X 16 DIMMS @ 1866 MHz) Infrastructure Servers
- 14 X Cisco UCS B200-M3 (E5-2680 @ 2.8 GHz) blade servers with 256 GB of memory (16 GB X 16 DIMMS @ 1866 MHz) to run SVDI and RDS workloads on.
- 1X VIC1240 Converged Network Adapter/Blade (B200 M3)
- 2 X Cisco UCS 6248UP Fabric Interconnects
- 2 X Cisco Nexus 5548UP Access Switches
- 1 X EMC VNX5600 system with 32 x 600GB SAS drives, 24 x 2TB Near Line SAS Drives, 3 x 100GB Flash Drives (Fast Cache) including hot spares.

Software components

- Cisco UCS firmware 2.2(2c)
- Cisco UCS Director 4.2Cisco Nexus 1000V virtual distributed switch
- VMware ESXi 5.5 Hosts
- Citrix XenDesktop 7.5 Server VDIl Desktops and RDS Hosted Shared Desktops
- Citrix Provisioning Server 7.1
- Citrix User Profile Manager
- Microsoft Windows Server 2012 R2 64 bit, 1vCPU, 1.5 GB RAM, 24 GB hard disk/VM

- Microsoft Windows Server 2012 R2 64 bit , 4 vCPU, 16GB RAM, 50 GB hard disk/VM

Testing Methodology and Success Criteria

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

The testing results focused on pure functionality of the Citrix DaaS Provisioning process, new tenant on-boarding and the ability of 2000 users spread across 10 different tenants being able to login simultaneously.

The figure below shows the Citrix Studio results during the test of the shared DaaS tenants. Tenants 3-7 have 2 Delivery Groups each with 90% of their users on Hosted Shared Desktops and 10% on Server VDI desktops.

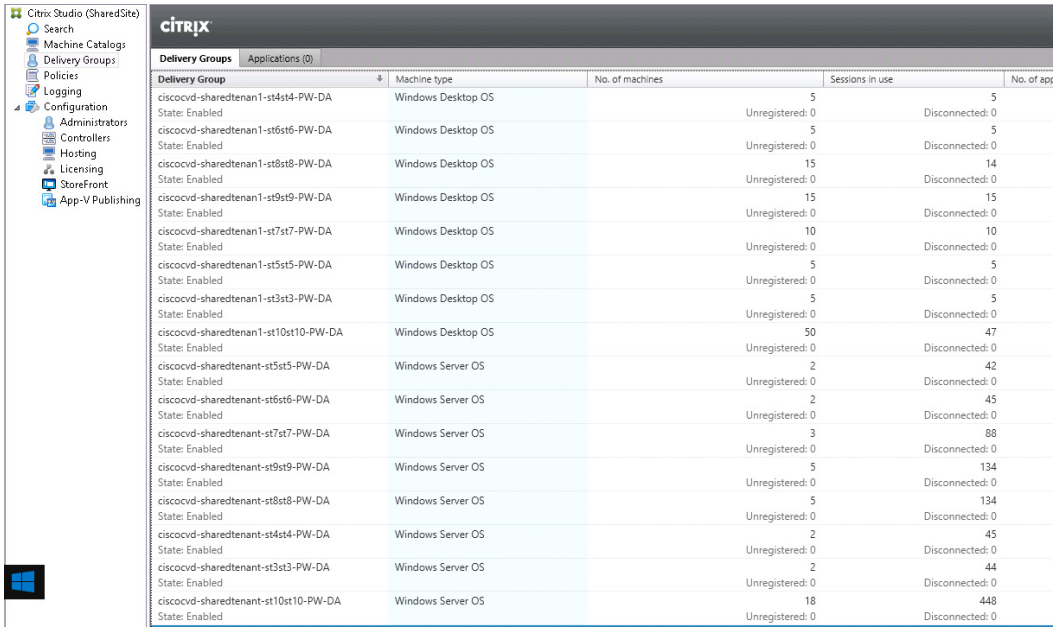

The figure below shows the Studio results for our Private Isolated Tenant 1 model. They had a total of 200 users with 90% using Hosted Shared Desktops and 10% using Server VDI Desktops.

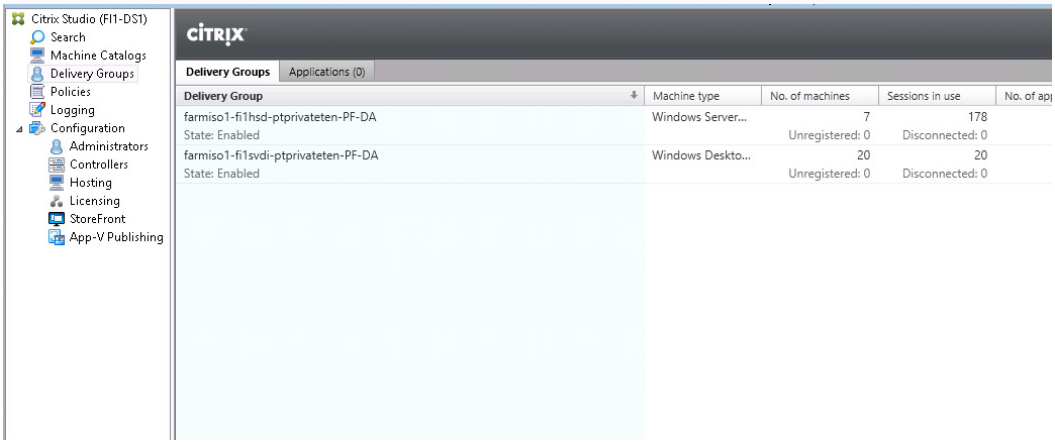

The figure below shows the Studio results for our Private Isolated Tenant 2 model. They had a total of 700 users with 90% using Hosted Shared Desktops and 10% using Server VDI Desktops.

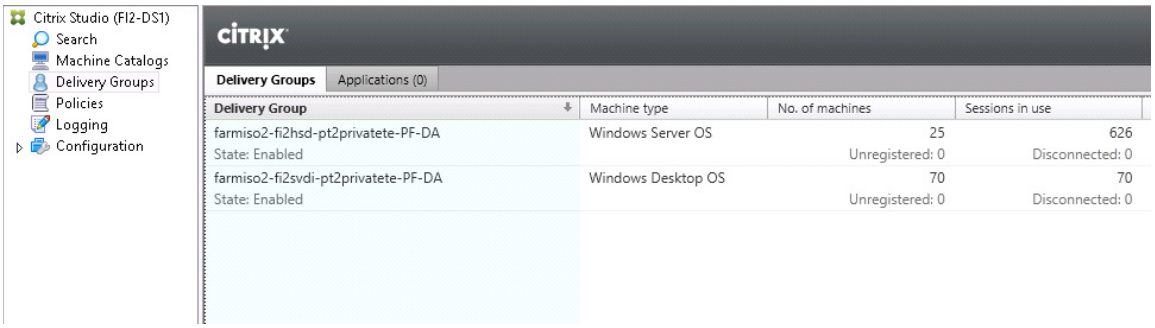

Scalability Considerations andGudelines

There are many factors to consider when you begin to scale beyond 2000 Users, two chassis 8 mixed workload SVDI/HSD host server configuration, which this reference architecture has successfully tested. In this section we give guidance to scale beyond the 2000 user system.

Cisco UCS System Scalability

As our results indicate, we have proven linear scalability in the Cisco UCS Reference Architecture as tested.

- Cisco UCS 2.2(2c) management software supports up to 20 chassis within a single Cisco UCS domain on our second generation Cisco UCS Fabric Interconnect 6248 and 6296 models. Our single UCS domain can grow to 160 blades.
- With Cisco UCS 2.2(2c) management software, released in May 2014, each UCS 2.2(2c) Management domain is extensibly manageable by UCS Central, our new manager of managers, vastly increasing the reach of the UCS system.

- As scale grows, the value of the combined UCS fabric, Nexus physical switches and Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100% of the time.
- To accommodate the Cisco Nexus 5500 upstream connectivity in the way we describe in the LAN and SAN Configuration section, we need two Ethernet uplinks to be configured on the Cisco UCS Fabric interconnect. And based on the number of uplinks from each chassis, we can calculate number of desktops can be hosted in a single UCS domain. Assuming eight links per chassis, four to each 6248, scaling beyond 10 chassis would require a pair of Cisco UCS 6296 fabric interconnects.
- A 25,000 virtual desktop building block, managed by a single UCS domain, with its support infrastructure services can be built out from the RA described in this study with eight links per chassis and 152 Cisco UCS B200 M3 Servers and 8 infrastructure blades configured per the specifications in this document in 20 chassis.

Of course, the backend storage has to be scaled accordingly, based on the IOP considerations as described in the EMC scaling section. Please refer the EMC section that follows this one for scalability guidelines.

Citrix DaaS Scalability Considerations

Citrix Desktop-as-a Service (DaaS) environments can scale to large numbers of tenants and users, especially given the simplicity of provisioning and managing the provider environment with the combined solution of Citrix App Orchestration and CloudPortal Services Manager.

For this CVD, the environment was configured to support, overall, 2000 users and 10 tenants. The distribution of Hosted Shared Desktops (HSD) and Server VDI desktops (SVDI) followed a ratio of 90% HSD and 10% SVDI, which is common is many real-world customer deployments.

When implementing Citrix DaaS with Citrix XenDesktop 7.5, system architects should carefully consider a number of factors, including the following:

- The number and types of tenants and whether each tenant uses a shared or private delivery site
- Types of desktops that will be deployed and the quantity of each type (HSD and SVDI)
- Types of storage in the environment
- Data protection requirements
- Write cache size and placement for pooled desktops

When designing a solution based on this CVD environment, it's recommended that solution architects follow the best practices used in this CVD's implementation:

- Use an N+1 schema for the virtualization host servers to provide resiliency. In the CVD environment, all host servers were configured with N+1 redundancy.
- Configure network adapters on all Provisioning Servers using static IP addresses. Separate management networks for these servers should also be configured.

EMC VNX5600 Storage Guidelines for Mixed Desktops Virtualization Workload

Sizing VNX storage system to meet virtual desktop IOPS requirement is a complicated process. When an I/O reaches the VNX storage, it is served by several components such as Data Mover (NFS), backend dynamic random access memory (DRAM) cache, FAST Cache, and disks. To reduce the complexity, EMC recommends using a building block approach to scale to thousands of virtual desktops.

For more information on storage sizing guidelines to implement your end-user computing solution on VNX unified storage systems, refer to the EMC VSPEX sizing tool portal.

Cisco Reference Doucments

Cisco Unified Computing System Manager Home Page

http://www.cisco.com/en/US/products/ps10281/index.html

Cisco UCS B200 M3 Blade Server Resources

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

Cisco Nexus 5500 Series Switches Resources

http://www.cisco.com/en/US/products/ps9670/index.html

Download Cisco UCS Manager and Blade Software Version 2.2(2c)

https://software.cisco.com/download/release.html?mdfid=283612660&softwareid=283655658&release=2.2%283e%29&relind=AVAILABLE&rellifecycle=&reltype=latest

Cisco UCS Director Resources

http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-director/tsd-products-support-series-home.html

Citrix Reference Documents

Citrix Product Downloads
http://www.citrix.com/downloads/xendesktop.html

Citrix Knowledge Center

http://support.citrix.com

Citrix documentation
http://www.citrix.com/edocs

XenApp 7.5 and XenDesktop 7.5 Documentation
http://support.citrix.com/proddocs/topic/xenapp-xendesktop/cds-xenapp-xendesktop-75-landing.html

Citrix Provisioning Services
http://support.citrix.com/proddocs/topic/provisioning-7/pvs-provisioning-7.html

Citrix XenDesktop 7.5 Reviewer's Guide (describes how to set up an evaluation)
http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/xendesktop-reviewers-guide.pdf

Citrix App Orchestration 2.5 Documentation
http://support.citrix.com/proddocs/topic/app-orchestration/cao-app-orchestration-25-landing.html

Known Issues for App Orchestration 2.5

Terminology in App Orchestration 2.5
Getting Started with Citrix App Orchestration 2.5
Configuring SSL for App Orchestration 2.5
Deploying the Zero Trust Agent in App Orchestration 2.5

CloudPortal Services Manager Documentation
http://support.citrix.com/proddocs/topic/cloudportal/ccps-services-manager.html

Known Issues for CloudPortal Services Manager

Login VSI
http://www.loginvsi.com/documentation/

EMC References

- EMC VSPEX End User Computing Solution Overview
- EMC VSPEX End-User Computing: Citrix XenDesktop 7 and VMware vSphere for up to 2,000 Virtual Desktops – Design Guide
- EMC VSPEX End-User Computing: Citrix XenDesktop 7 and VMware vSphere for up to 2,000 Virtual Desktops – Implementation Guide
- EMC VSPEX End-User Computing: Citrix XenDesktop 7 and Microsoft Hyper-V for up to 2,000 Virtual Desktops – Design Guide
- EMC VSPEX End-User Computing: Citrix XenDesktop 7 and Microsoft Hyper-V for up to 2,000 Virtual Desktops – Implementation Guide

VMware References

VMware vCenter Server

http://www.vmware.com/products/vcenter-server/

VMware VSphere

http://www.vmware.com/products/vsphere/

Appendix A Nexus 5548UP Configurations

N5548UP-A Configuration

```
version 5.1(3)N2(1)
feature fcoe
hostname N5KA-EXC
feature npiv
no feature telnet
no feature http-server
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
username admin password 5 $1$6sb7/rvC$ds.AAWPP3DFjJu.7VBd35/  role network-admin
no password strength-check
ip domain-lookup
logging event link-status default
ip access-list VLAN77
  10 permit ip 10.77.0.0/24 any
  20 deny ip any any
class-map type qos class-fcoe
class-map type qos match-all VLAN77
  match access-group name VLAN77
class-map type queuing class-fcoe
  match qos-group 1
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 VLAN77
  class VLAN77
  class class-default
class-map type network-qos class-fcoe
  match qos-group 1
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 jumbo
```

```
  class type network-qos class-default
    mtu 9216
    multicast-optimize
system qos
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type qos input fcoe-default-in-policy
  service-policy type network-qos jumbo
slot 2
  port 1-8 type ethernet
  port 9-16 type fc
snmp-server user admin network-admin auth md5 0xcb74bb1d86558f73d4d8688660b53c49
 priv 0xcb74bb1d86558f73d4d8688660b53c49 localizedkey
vrf context management
spanning-tree port type edge default
spanning-tree vlan 1-3967 priority 24576
service dhcp
ip dhcp relay
vpc domain 210
  role priority 1
  peer-keepalive destination 10.60.0.9
vsan database
  vsan 1 name "default"
device-alias database
  device-alias name CH1-BL5 pwwn 20:00:00:25:b5:03:00:05
  device-alias name CH1-BL6 pwwn 20:00:00:25:b5:03:00:03
  device-alias name CH1-BL7 pwwn 20:00:00:25:b5:03:00:02
  device-alias name CH2-BL5 pwwn 20:00:00:25:b5:03:00:04
  device-alias name CH2-BL6 pwwn 20:00:00:25:b5:03:00:01
  device-alias name CH2-BL7 pwwn 20:00:00:25:b5:03:00:00

  device-alias name Infra1-A pwwn 20:00:00:25:b5:03:00:1f
  device-alias name Infra2-A pwwn 20:00:00:25:b5:03:00:1e
  device-alias name vnx5600_SPA_22 pwwn 50:06:01:63:08:64:2f:88
  device-alias name vnx5600_SPB_23 pwwn 50:06:01:6b:08:64:2f:88
  device-alias name SP-CH2-BL1_vHBA1 pwwn 20:00:00:25:b5:f2:20:03
  device-alias name SP-Shared-2_vHBA1 pwwn 20:00:00:25:b5:f2:20:05
  device-alias name SP-Shared-3_vHBA1 pwwn 20:00:00:25:b5:f2:20:07
  device-alias name SP-Shared-4_vHBA1 pwwn 20:00:00:25:b5:f2:20:01
  device-alias name SP-Shared-5_vHBA1 pwwn 20:00:00:25:b5:f2:20:09
  device-alias name SP-Shared-6_vHBA1 pwwn 20:00:00:25:b5:f2:20:0b
  device-alias name SP-Shared-7_vHBA1 pwwn 20:00:00:25:b5:f2:20:0d
  device-alias name SP-FarmIso-T1-1_vHBA1 pwwn 20:00:00:25:b5:f2:20:11
  device-alias name SP-FarmIso-T1-2_vHBA1 pwwn 20:00:00:25:b5:f2:20:0f
  device-alias name SP-FarmIso-T1-3_vHBA1 pwwn 20:00:00:25:b5:f2:20:13
  device-alias name SP-FarmIso-T2-1_vHBA1 pwwn 20:00:00:25:b5:f2:20:15
  device-alias name SP-FarmIso-T2-2_vHBA1 pwwn 20:00:00:25:b5:f2:20:17
  device-alias name SP-FarmIso-T2-3_vHBA1 pwwn 20:00:00:25:b5:f2:20:19
  device-alias name SP-FarmIso-T2-4_vHBA1 pwwn 20:00:00:25:b5:f2:20:1b

device-alias commit

fcdomain fcid database
  vsan 1 wwn 21:00:00:24:ff:68:1d:dc fcid 0x940000 dynamic
  vsan 1 wwn 21:00:00:24:ff:68:2d:d4 fcid 0x940001 dynamic
  vsan 1 wwn 20:1f:54:7f:ee:f8:02:80 fcid 0x940002 dynamic
  vsan 1 wwn 20:20:54:7f:ee:f8:02:80 fcid 0x940003 dynamic
  vsan 1 wwn 20:00:00:25:b5:03:00:03 fcid 0x940011 dynamic
```

```
!            [CH1-BL6]
  vsan 1 wwn 20:00:00:25:b5:03:00:02 fcid 0x940012 dynamic
!            [CH1-BL7]
  vsan 1 wwn 20:00:00:25:b5:03:00:00 fcid 0x940013 dynamic
!            [CH2-BL7]
  vsan 1 wwn 20:00:00:25:b5:03:00:05 fcid 0x940014 dynamic
!            [CH1-BL5]
  vsan 1 wwn 20:00:00:25:b5:03:00:01 fcid 0x940015 dynamic
!            [CH2-BL6]
  vsan 1 wwn 20:4c:54:7f:ee:77:74:c0 fcid 0x940018 dynamic
  vsan 1 wwn 50:06:01:6b:08:64:2f:88 fcid 0x9400ef dynamic
!            [vnx5600_SPB_23]
  vsan 1 wwn 50:06:01:63:08:64:2f:88 fcid 0x9401ef dynamic
!            [vnx5600_SPA_22]

  vsan 1 wwn 20:00:00:25:b5:f2:20:01 fcid 0x940019 dynamic
!            [SP-Shared-4_vHBA1]
  vsan 1 wwn 20:4b:54:7f:ee:77:74:c0 fcid 0x940016 dynamic
  vsan 1 wwn 20:00:00:25:b5:f2:20:1b fcid 0x940017 dynamic
!            [SP-FarmIso-T2-4_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:03 fcid 0x94001a dynamic
!            [SP-CH2-BL1_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:05 fcid 0x94001b dynamic
!            [SP-Shared-2_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:07 fcid 0x94001c dynamic
!            [SP-Shared-3_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:09 fcid 0x94001d dynamic
!            [SP-Shared-5_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:0b fcid 0x94001e dynamic
!            [SP-Shared-6_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:0d fcid 0x94001f dynamic
!            [SP-Shared-7_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:0f fcid 0x940020 dynamic
!            [SP-FarmIso-T1-2_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:11 fcid 0x940021 dynamic
!            [SP-FarmIso-T1-1_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:13 fcid 0x940022 dynamic
!            [SP-FarmIso-T1-3_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:15 fcid 0x940023 dynamic
!            [SP-FarmIso-T2-1_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:17 fcid 0x940024 dynamic
!            [SP-FarmIso-T2-2_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:19 fcid 0x940025 dynamic
!            [SP-FarmIso-T2-3_vHBA1]
  vsan 1 wwn 20:00:00:25:b5:f2:20:1d fcid 0x940026 dynamic
  vsan 1 wwn 20:00:00:25:b5:f2:20:21 fcid 0x940027 dynamic
  vsan 1 wwn 20:00:00:25:b5:f2:20:1f fcid 0x940028 dynamic
  vsan 1 wwn 21:00:00:24:ff:57:fa:c1 fcid 0x94002b dynamic


interface Vlan1
  no shutdown


interface Vlan70
  no shutdown
  description SP Pod Management
  no ip redirects
```

```
  ip address 10.70.0.2/24
  hsrp version 2
  hsrp 70
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.70.0.1

interface Vlan71

  no shutdown
  description SP Pod Infrastructure
  no ip redirects
  ip address 10.71.0.2/21
  hsrp version 2
  hsrp 71
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.71.0.1
  ip dhcp relay address 10.71.0.10

interface Vlan72
  no shutdown
  description SP Pod Storage
  no ip redirects
  ip address 10.72.0.2/24
  hsrp version 2
  hsrp 72
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.72.0.1


interface Vlan73
  no shutdown
  description SP Pod VMotion
  no ip redirects
  ip address 10.73.0.2/24
  hsrp version 2
  hsrp 73
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.73.0.1

interface Vlan74
  no shutdown
  description SP Pod Shared Tenant
  no ip redirects
  ip address 10.74.0.2/24
  hsrp version 2
  hsrp 74
    preempt delay minimum 240
    priority 130
    timers  1  3
```

```
    ip 10.74.0.1

interface Vlan75
  no shutdown
  description SP Pod Server Isolation Tenant
  no ip redirects
  ip address 10.75.0.2/24
  hsrp version 2
  hsrp 75
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.75.0.1

interface Vlan76
  no shutdown
  description SP Pod Farm Isolation Tenant 1
  no ip redirects
  ip address 10.76.0.2/24
  hsrp version 2
  hsrp 76
    preempt delay minimum 240
    priority 130

    timers  1  3
    ip 10.76.0.1
  ip dhcp relay address 10.76.0.10

interface Vlan77
  no shutdown
  description SP Pod Farm Isolation Tenant 2
  no ip redirects
  ip address 10.77.0.2/24
  hsrp version 2
  hsrp 77
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.77.0.1
  ip dhcp relay address 10.77.0.10

interface Vlan78
  description SP Pod DMZ
  ip address 10.78.0.2/24
  hsrp version 2
  hsrp 78
    preempt delay minimum 240

    priority 130
    timers  1  3
    ip 10.78.0.1

interface Vlan79
  no shutdown
  description UCS Director PXE
  ip address 10.79.0.2/24
  hsrp version 2
  hsrp 79
```

```
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.79.0.1


interface port-channel1
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type network
  vpc peer-link

interface port-channel15
  description VNX-DM3-0
  switchport mode trunk
  switchport trunk allowed vlan 72
  spanning-tree port type edge trunk
  vpc 15

interface port-channel16
  description VNX-DM2-0
  switchport mode trunk
  switchport trunk allowed vlan 72
  spanning-tree port type edge trunk
  vpc 16

interface port-channel17
  description FI-A-Uplink-Launcher
  switchport mode trunk

  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  vpc 17

interface port-channel20
  description FI-B-Uplink-Launcher
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  vpc 20

interface port-channel21
  description FI-A-Uplink-SP
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,801
  spanning-tree port type edge trunk
  vpc 21

interface port-channel24
  description FI-B-Uplink-SP
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk

  vpc 24

interface port-channel25
  description FI-A-Uplink-HyperV
```

```
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  vpc 25

interface port-channel28
  description FI-B-Uplink-HyperV
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  vpc 28

vsan database
  vsan 4094 interface vfc13
  vsan 4094 interface vfc14

interface fc2/9
  no shutdown

interface fc2/10
  no shutdown

interface fc2/11

  no shutdown

interface fc2/12
  no shutdown

interface fc2/13
  no shutdown

interface fc2/14
  no shutdown

interface fc2/15
  switchport trunk allowed vsan 1
  no shutdown

interface fc2/16
  no shutdown

interface Ethernet1/1
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  channel-group 1 mode active


interface Ethernet1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  channel-group 1 mode active

interface Ethernet1/3
  switchport access vlan 60
  speed 1000
  duplex full
```

```
interface Ethernet1/4
  speed 1000


interface Ethernet1/6
  description TOR-Local60
  switchport access vlan 60
  speed 1000

interface Ethernet1/8
  speed 1000

interface Ethernet1/10
  description VNX-Mgmt
  switchport access vlan 70
  speed 1000

interface Ethernet1/11
  description SP-FI-Mgmt
  switchport access vlan 70
  speed 1000

interface Ethernet1/15
  switchport mode trunk
  switchport trunk allowed vlan 72
  spanning-tree port type edge
  channel-group 15 mode active

interface Ethernet1/16
  switchport mode trunk

  switchport trunk allowed vlan 72
  channel-group 16 mode active

interface Ethernet1/17
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 17 mode active

interface Ethernet1/18
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 17 mode active

interface Ethernet1/19
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 20 mode active

interface Ethernet1/20
  switchport mode trunk

  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 20 mode active
```

```
interface Ethernet1/21
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,801
  spanning-tree port type edge trunk
  channel-group 21 mode active

interface Ethernet1/22
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,801
  spanning-tree port type edge trunk
  channel-group 21 mode active

interface Ethernet1/23
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 24 mode active

interface Ethernet1/24

  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 24 mode active

interface Ethernet1/25
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 25 mode active

interface Ethernet1/26
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 25 mode active

interface Ethernet1/27
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 28 mode active


interface Ethernet1/28
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 28 mode active
interface Ethernet1/31

interface Ethernet1/32

interface Ethernet2/6
  speed 1000

interface mgmt0
```

```
  ip address 10.60.0.8/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin
ip routing event-history general size large
ip routing event-history summary size large
interface fc2/9
interface fc2/10
interface fc2/11

interface fc2/12
interface fc2/13
interface fc2/14
interface fc2/15
interface fc2/16
!Full Zone Database Section for vsan 1

zone name B200M3-CH1-BL5-FC0 vsan 1
    member pwwn 20:00:00:25:b5:03:00:05
!               [CH1-BL5]


zone name B200M3-CH2-BL5-FC0 vsan 1
    member pwwn 20:00:00:25:b5:03:00:04
!               [CH2-BL5]

zone name UCS-VNX5600A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:1b
!               [SP-FarmIso-T2-4_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]
    member pwwn 20:00:00:25:b5:f2:20:13
!               [SP-FarmIso-T1-3_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:0f
!               [SP-FarmIso-T1-2_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:0b
!               [SP-Shared-6_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:07

!               [SP-Shared-3_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:03
!               [SP-CH2-BL1_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:17
!               [SP-FarmIso-T2-2_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:01
!               [SP-Shared-4_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:09
!               [SP-Shared-5_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:05
!               [SP-Shared-2_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:11
!               [SP-FarmIso-T1-1_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:0d
!               [SP-Shared-7_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:15
```

```
!               [SP-FarmIso-T2-1_vHBA1]
    member pwwn 20:00:00:25:b5:f2:20:19
!               [SP-FarmIso-T2-3_vHBA1]

zone name SP-CH2-BL1_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:03
!               [SP-CH2-BL1_vHBA1]

    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name SP-Shared-4_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:01
!               [SP-Shared-4_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name SP-Shared-1_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:03
!               [SP-CH2-BL1_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name B200M3-CH1-BL6-FC0 vsan 1
    member pwwn 20:00:00:25:b5:03:00:03

!               [CH1-BL6]
    member pwwn 21:00:00:24:ff:50:ab:b5
!               [VNX5600-P1]

zone name B200M3-CH1-BL7-FC0 vsan 1
    member pwwn 20:00:00:25:b5:03:00:02
!               [CH1-BL7]
    member pwwn 21:00:00:24:ff:50:ab:b5
!               [VNX5600-P1]

zone name B200M3-CH2-BL6-FC0 vsan 1
    member pwwn 20:00:00:25:b5:03:00:01
!               [CH2-BL6]
    member pwwn 21:00:00:24:ff:57:fa:c0
!               [VNX560024TB-P1]

zone name B200M3-CH2-BL7-FC0 vsan 1
    member pwwn 20:00:00:25:b5:03:00:00
!               [CH2-BL7]
    member pwwn 21:00:00:24:ff:50:ab:b5
!               [VNX5600-P1]

zone name SP-Shared-2_ZONE_A vsan 1

    member pwwn 20:00:00:25:b5:f2:20:05
!               [SP-Shared-2_vHBA1]
```

```
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name SP-Shared-3_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:07
!               [SP-Shared-3_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name SP-Shared-5_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:09
!               [SP-Shared-5_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]


zone name SP-Shared-6_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:0b
!               [SP-Shared-6_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name SP-Shared-7_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:0d
!               [SP-Shared-7_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name SP-FarmIso-T1-2_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:0f
!               [SP-FarmIso-T1-2_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]


_[K
zone name SP-FarmIso-T1-1_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:11
!               [SP-FarmIso-T1-1_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!               [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!               [vnx5600_SPB_23]

zone name SP-FarmIso-T1-3_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:13
```

```
!                [SP-FarmIso-T1-3_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!                [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!                [vnx5600_SPB_23]

zone name SP-FarmIso-T2-1_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:15
!                [SP-FarmIso-T2-1_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!                [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88

!                [vnx5600_SPB_23]

zone name SP-FarmIso-T2-2_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:17
!                [SP-FarmIso-T2-2_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!                [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!                [vnx5600_SPB_23]

zone name SP-FarmIso-T2-3_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:19
!                [SP-FarmIso-T2-3_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!                [vnx5600_SPA_22]
    member pwwn 50:06:01:6b:08:64:2f:88
!                [vnx5600_SPB_23]

zone name SP-FarmIso-T2-4_ZONE_A vsan 1
    member pwwn 20:00:00:25:b5:f2:20:1b
!                [SP-FarmIso-T2-4_vHBA1]
    member pwwn 50:06:01:63:08:64:2f:88
!                [vnx5600_SPA_22]

    member pwwn 50:06:01:6b:08:64:2f:88
!                [vnx5600_SPB_23]

zone name B200M4-CH1-BL7-FC0 vsan 1
    member pwwn 20:00:00:25:b5:00:00:3f
!                [B200M4-CH1-BL7]
    member pwwn 21:00:00:24:ff:57:fa:c0
!                [VNX560024TB-P1]

zone name B200M4-CH1-BL6 vsan 1
    member pwwn 20:00:00:25:b5:03:00:03
!                [CH1-BL6]
    member pwwn 21:00:00:24:ff:57:fa:c0
!                [VNX560024TB-P1]

zoneset name UCS-VNX5600A vsan 1
    member B200M3-CH1-BL8-FC0
    member B200M3-CH2-BL8-FC0
    member B200M3-CH1-BL1-FC0
    member B200M3-CH1-BL2-FC0
    member B200M3-CH1-BL3-FC0
```

```
    member B200M3-CH1-BL4-FC0
    member B200M3-CH2-BL1-FC0

    member B200M3-CH2-BL2-FC0
    member B200M3-CH2-BL3-FC0
    member B200M3-CH2-BL4-FC0
    member B200M3-CH1-BL5-FC0
    member B200M3-CH2-BL5-FC0
    member SP-CH2-BL1_ZONE_A
    member SP-Shared-4_ZONE_A
    member SP-Shared-1_ZONE_A
    member B200M3-CH1-BL6-FC0
    member B200M3-CH1-BL7-FC0
    member B200M3-CH2-BL6-FC0
    member B200M3-CH2-BL7-FC0
    member SP-Shared-2_ZONE_A
    member SP-Shared-3_ZONE_A
    member SP-Shared-5_ZONE_A
    member SP-Shared-6_ZONE_A
    member SP-Shared-7_ZONE_A
    member SP-FarmIso-T1-2_ZONE_A
    member SP-FarmIso-T1-1_ZONE_A
    member SP-FarmIso-T1-3_ZONE_A
    member SP-FarmIso-T2-1_ZONE_A
    member SP-FarmIso-T2-2_ZONE_A
    member SP-FarmIso-T2-3_ZONE_A

    member SP-FarmIso-T2-4_ZONE_A
    member B200M4-CH1-BL7-FC0
    member B200M4-CH1-BL6

zoneset activate name UCS-VNX5600A vsan 1
```

N5548UP-B Configuration

```
version 5.1(3)N2(1)
feature fcoe
hostname N5KB-EXC
feature npiv
no feature telnet
no feature http-server
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
username admin password 5 $1$VfkcAdcR$ioaEijZgf.0./cIoNDVUz.  role network-admin
no password strength-check
ip domain-lookup
ip access-list VLAN77
  10 permit ip 10.77.0.0/24 any
  20 deny ip any any
_[7m--More--_[27m
```

```
class-map type qos class-fcoe
class-map type qos match-all VLAN77
  match access-group name VLAN77
class-map type queuing class-fcoe
  match qos-group 1
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 VLAN77
  class VLAN77
  class class-default
class-map type network-qos class-fcoe
  match qos-group 1
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 jumbo
  class type network-qos class-default
    mtu 9216
    multicast-optimize
system qos

  service-policy type queuing output fcoe-default-out-policy
  service-policy type qos input fcoe-default-in-policy
  service-policy type network-qos jumbo
  service-policy type queuing input fcoe-default-in-policy
slot 2
  port 1-8 type ethernet
  port 9-16 type fc
snmp-server user admin network-admin auth md5 0xc97a3a05a3ee4e5eccd57e453bfe93e8
 priv 0xc97a3a05a3ee4e5eccd57e453bfe93e8 localizedkey
vrf context management
  ip route 0.0.0.0/0 10.60.0.1
spanning-tree port type edge default
spanning-tree vlan 1-3967 priority 28672
service dhcp
ip dhcp relay
vpc domain 210
  role priority 2
  peer-keepalive destination 10.60.0.8
vsan database
  vsan 1 name "default"
  vsan 501 name "Fabric_B"
device-alias database
  device-alias name CH1-BL5 pwwn 20:00:00:25:b5:03:00:15
  device-alias name CH1-BL6 pwwn 20:00:00:25:b5:03:00:13
  device-alias name CH1-BL7 pwwn 20:00:00:25:b5:03:00:12
  device-alias name CH2-BL5 pwwn 20:00:00:25:b5:03:00:14
  device-alias name CH2-BL6 pwwn 20:00:00:25:b5:03:00:11
  device-alias name CH2-BL7 pwwn 20:00:00:25:b5:03:00:10
  device-alias name SRA-FC2 pwwn 21:00:00:24:ff:68:1d:dd
  device-alias name SRB-FC2 pwwn 21:00:00:24:ff:68:2d:d5
  device-alias name Infra1-B pwwn 20:00:00:25:b5:03:00:0f
  device-alias name Infra2-B pwwn 20:00:00:25:b5:03:00:0e
  device-alias name VNX5600-P2 pwwn 21:00:00:24:ff:50:ab:b4
  device-alias name B200M4-CH1-BL7 pwwn 20:00:00:25:b5:00:00:2f
```

```
  device-alias name VNX560024TB-P1 pwwn 21:00:00:24:ff:57:fa:c1
  device-alias name vnx5600_SPA_20 pwwn 50:06:01:62:08:64:2f:88
  device-alias name vnx5600_SPB_21 pwwn 50:06:01:6a:08:64:2f:88
  device-alias name SP-CH2-BL1_vHBA2 pwwn 20:00:00:25:b5:f2:20:02
  device-alias name SP-Shared-2_vHBA2 pwwn 20:00:00:25:b5:f2:20:04
  device-alias name SP-Shared-3_vHBA2 pwwn 20:00:00:25:b5:f2:20:06
  device-alias name SP-Shared-4_vHBA2 pwwn 20:00:00:25:b5:f2:20:00
  device-alias name SP-Shared-5_vHBA2 pwwn 20:00:00:25:b5:f2:20:08
  device-alias name SP-Shared-6_vHBA2 pwwn 20:00:00:25:b5:f2:20:0a
  device-alias name SP-Shared-7_vHBA2 pwwn 20:00:00:25:b5:f2:20:0c
  device-alias name SP-FarmIso-T1-1_vHBA2 pwwn 20:00:00:25:b5:f2:20:10
  device-alias name SP-FarmIso-T1-2_vHBA2 pwwn 20:00:00:25:b5:f2:20:0e
  device-alias name SP-FarmIso-T1-3_vHBA2 pwwn 20:00:00:25:b5:f2:20:12
  device-alias name SP-FarmIso-T2-1_vHBA2 pwwn 20:00:00:25:b5:f2:20:14
  device-alias name SP-FarmIso-T2-2_vHBA2 pwwn 20:00:00:25:b5:f2:20:16

  device-alias name SP-FarmIso-T2-3_vHBA2 pwwn 20:00:00:25:b5:f2:20:18
  device-alias name SP-FarmIso-T2-4_vHBA2 pwwn 20:00:00:25:b5:f2:20:1a

device-alias commit

fcdomain fcid database
  vsan 1 wwn 21:00:00:24:ff:68:1d:dd fcid 0x730000 dynamic
!            [SRA-FC2]
  vsan 1 wwn 21:00:00:24:ff:68:2d:d5 fcid 0x730001 dynamic
!            [SRB-FC2]
  vsan 1 wwn 20:1f:54:7f:ee:83:42:00 fcid 0x730002 dynamic
  vsan 1 wwn 20:20:54:7f:ee:83:42:00 fcid 0x730003 dynamic
  !            [Infra1-B]
  vsan 1 wwn 20:00:00:25:b5:03:00:0e fcid 0x730006 dynamic
!            [Infra2-B]
   vsan 1 wwn 20:00:00:25:b5:03:00:05 fcid 0x73000e dynamic
  vsan 1 wwn 20:00:00:25:b5:03:00:14 fcid 0x73000f dynamic
!            [CH2-BL5]
  vsan 1 wwn 21:00:00:24:ff:50:ab:b4 fcid 0x730010 dynamic
!            [VNX5600-P2]
  vsan 1 wwn 20:00:00:25:b5:03:00:13 fcid 0x730011 dynamic
!            [CH1-BL6]
  vsan 1 wwn 20:00:00:25:b5:03:00:12 fcid 0x730012 dynamic
!            [CH1-BL7]
  vsan 1 wwn 20:00:00:25:b5:03:00:10 fcid 0x730013 dynamic
!            [CH2-BL7]
  vsan 1 wwn 20:00:00:25:b5:03:00:15 fcid 0x730014 dynamic
!            [CH1-BL5]
  vsan 1 wwn 20:00:00:25:b5:03:00:11 fcid 0x730015 dynamic

!            [CH2-BL6]
  vsan 1 wwn 21:00:00:24:ff:57:fa:c0 fcid 0x730016 dynamic
  vsan 1 wwn 50:06:01:6a:08:64:2f:88 fcid 0x7300ef dynamic
!            [vnx5600_SPB_21]
  vsan 1 wwn 50:06:01:62:08:64:2f:88 fcid 0x7301ef dynamic
!            [vnx5600_SPA_20]
  vsan 1 wwn 20:4b:54:7f:ee:76:cd:00 fcid 0x730017 dynamic
  vsan 1 wwn 20:4c:54:7f:ee:76:cd:00 fcid 0x730018 dynamic
  vsan 1 wwn 20:00:00:25:b5:f2:20:1a fcid 0x730019 dynamic
!            [SP-FarmIso-T2-4_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:00 fcid 0x73001a dynamic
!            [SP-Shared-4_vHBA2]
```

```
  vsan 1 wwn 20:00:00:25:b5:f2:20:02 fcid 0x73001b dynamic
!            [SP-CH2-BL1_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:04 fcid 0x73001c dynamic
!            [SP-Shared-2_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:06 fcid 0x73001d dynamic
!            [SP-Shared-3_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:08 fcid 0x73001e dynamic
!            [SP-Shared-5_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:0a fcid 0x73001f dynamic
!            [SP-Shared-6_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:0c fcid 0x730020 dynamic

!            [SP-Shared-7_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:0e fcid 0x730021 dynamic
!            [SP-FarmIso-T1-2_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:10 fcid 0x730022 dynamic
!            [SP-FarmIso-T1-1_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:12 fcid 0x730023 dynamic
!            [SP-FarmIso-T1-3_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:14 fcid 0x730024 dynamic
!            [SP-FarmIso-T2-1_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:16 fcid 0x730025 dynamic
!            [SP-FarmIso-T2-2_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:18 fcid 0x730026 dynamic
!            [SP-FarmIso-T2-3_vHBA2]
  vsan 1 wwn 20:00:00:25:b5:f2:20:1c fcid 0x730027 dynamic
  vsan 1 wwn 20:00:00:25:b5:f2:20:20 fcid 0x730028 dynamic
  vsan 1 wwn 20:00:00:25:b5:f2:20:1e fcid 0x730029 dynamic
  vsan 1 wwn 21:00:00:24:ff:57:fa:c1 fcid 0x73002a dynamic
!            [VNX560024TB-P1]
  vsan 1 wwn 20:00:00:25:b5:00:00:2f fcid 0x73002b dynamic
!            [B200M4-CH1-BL7]


interface Vlan1

  no shutdown


interface Vlan70
  no shutdown
  description SP Pod Management
  no ip redirects
  ip address 10.70.0.3/24
  hsrp version 2
  hsrp 70
    preempt delay minimum 240

    priority 130
    timers  1  3
    ip 10.70.0.1

interface Vlan71
  no shutdown
  description SP Pod Infrastructure
  no ip redirects
  ip address 10.71.0.3/21
  hsrp version 2
```

```
  hsrp 71
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.71.0.1
  ip dhcp relay address 10.71.0.10

interface Vlan72
  no shutdown
  description SP Pod CIFS/NFS Storage
  no ip redirects
  ip address 10.72.0.3/24
  hsrp version 2

  hsrp 72
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.72.0.1

interface Vlan73
  no shutdown
  description SP Pod VMotion
  no ip redirects
  ip address 10.73.0.3/24
  hsrp version 2
  hsrp 73
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.73.0.1

interface Vlan74
  no shutdown
  description SP Pod Shared Tenants
  no ip redirects
  ip address 10.74.0.3/24

  hsrp version 2
  hsrp 74
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.74.0.1
  ip dhcp relay address 10.71.0.10

interface Vlan75
  no shutdown
  description SP Pod Server Iso Tenant
  no ip redirects
  ip address 10.75.0.3/24
  hsrp version 2
  hsrp 75
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.75.0.1
```

```
interface Vlan76
  no shutdown
  description SP Pod Farm Iso Tenant 1

  no ip redirects
  ip address 10.76.0.3/24
  hsrp version 2
  hsrp 76
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.76.0.1
  ip dhcp relay address 10.76.0.10

interface Vlan77
  no shutdown
  description SP Pod Farm Iso Tenant 2
  no ip redirects
  ip address 10.77.0.3/24
  hsrp version 2
  hsrp 77
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.77.0.1
  ip dhcp relay address 10.77.0.10


interface Vlan78
  no shutdown
  description SP Pod DMZ
  ip address 10.78.0.3/24
  hsrp version 2
  hsrp 78
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.78.0.1

interface Vlan79
  no shutdown
  description UCS Director PXE
  ip address 10.79.0.3/24
  hsrp version 2
  hsrp 79
    preempt delay minimum 240
    priority 130
    timers  1  3
    ip 10.79.0.1


interface port-channel1
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type network
  vpc peer-link

interface port-channel15
```

```
  description VNX-DM3-1
  switchport mode trunk
  switchport trunk allowed vlan 72
  spanning-tree port type edge trunk
  vpc 15

interface port-channel16
  description VNX-DM2-1
  switchport mode trunk
  switchport trunk allowed vlan 72
  spanning-tree port type edge trunk
  vpc 16

interface port-channel21
  description FI-A-Uplink-SP
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,801
  spanning-tree port type edge trunk
  vpc 21

interface port-channel24
  description FI-B-Uplink-SP

  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  vpc 24

interface port-channel25
  description FI-A-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  vpc 25

interface port-channel28
  description FI-B-Uplink
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  vpc 28


interface fc2/9
  no shutdown

interface fc2/10
  no shutdown

interface fc2/11

  no shutdown

interface fc2/12
  no shutdown

interface fc2/13
  no shutdown
```

```
interface fc2/14
  no shutdown

interface fc2/15
  no shutdown

interface fc2/16
  no shutdown

interface Ethernet1/1
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  channel-group 1 mode active

interface Ethernet1/2

  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  channel-group 1 mode active

interface Ethernet1/8
  speed 1000

interface Ethernet1/9
  description JumpBox-Local80
  switchport access vlan 80
  speed 1000

interface Ethernet1/10
  description JumpBox-Local70
  switchport access vlan 70
  speed 1000

interface Ethernet1/11
  description SP-FI-Mgmt
  switchport access vlan 70
  speed 1000

interface Ethernet1/12
  description Launcher-FI-Mgmt
  switchport access vlan 801
  speed 1000

interface Ethernet1/13
  description VNX5600
  switchport access vlan 70
  speed 1000

interface Ethernet1/14
  description VNX5600
  switchport access vlan 70
  speed 1000

interface Ethernet1/15
  switchport mode trunk
  switchport trunk allowed vlan 72
  spanning-tree port type edge
```

```
  channel-group 15 mode active

interface Ethernet1/16

  switchport mode trunk
  switchport trunk allowed vlan 72
  channel-group 16 mode active

interface Ethernet1/17
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 17 mode active

interface Ethernet1/18
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 17 mode active

interface Ethernet1/19
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 20 mode active

interface Ethernet1/20

  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 20 mode active

interface Ethernet1/21
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,801
  spanning-tree port type edge trunk
  channel-group 21 mode active

interface Ethernet1/22
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,801
  spanning-tree port type edge trunk
  channel-group 21 mode active

interface Ethernet1/23
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 24 mode active


interface Ethernet1/24
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 24 mode active
```

```
interface Ethernet1/25
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 25 mode active

interface Ethernet1/26
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 25 mode active

interface Ethernet1/27
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 28 mode active

interface Ethernet1/28
  switchport mode trunk
  switchport trunk allowed vlan 1,60-66,70-82,701,801
  spanning-tree port type edge trunk
  channel-group 28 mode active

interface Ethernet1/31


interface Ethernet1/32

interface Ethernet2/5

interface Ethernet2/6

interface mgmt0
  ip address 10.60.0.9/24
line console
line vty
boot kickstart bootflash:/n5000-uk9-kickstart.5.1.3.N2.1.bin
boot system bootflash:/n5000-uk9.5.1.3.N2.1.bin
interface fc2/9
interface fc2/10
interface fc2/11
interface fc2/12
interface fc2/13
interface fc2/14
interface fc2/15
interface fc2/16

!Full Zone Database Section for vsan 1
zone name B200M3-CH1-BL8-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:0f
!               [Infra1-B]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH2-BL8-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:0e
!               [Infra2-B]
```

```
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH1-BL1-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:0d
!               [VDI1-B]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH1-BL2-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:0c
!               [VDI2-B]
    member pwwn 21:00:00:24:ff:50:ab:b4

!               [VNX5600-P2]

zone name B200M3-CH1-BL3-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:0b
!               [VDI3-B]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH1-BL4-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:0a
!               [VDI4-B]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH2-BL1-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:19
!               [VDI5-B]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH2-BL2-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:18
!               [VDI6-B]

    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH2-BL3-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:17
!               [VDI7-B]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH2-BL4-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:16
!               [VDI8-B]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH2-BL5-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:14
!               [CH2-BL5]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]
```

```
zone name B200M3-CH1-BL5-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:15

!               [CH1-BL5]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name SP-CH2-BL1_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:02
!               [SP-CH2-BL1_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-Shared-4_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:00
!               [SP-Shared-4_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-Shared-1_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:02
!               [SP-CH2-BL1_vHBA2]

    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-Shared-2_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:04
!               [SP-Shared-2_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name B200M3-CH1-BL6-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:13
!               [CH1-BL6]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name B200M3-CH1-BL7-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:12
!               [CH1-BL7]
    member pwwn 21:00:00:24:ff:50:ab:b4

!               [VNX5600-P2]

zone name B200M3-CH2-BL6-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:11
!               [CH2-BL6]
    member pwwn 21:00:00:24:ff:50:ab:b4
```

```
!               [VNX5600-P2]

zone name B200M3-CH2-BL7-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:10
!               [CH2-BL7]
    member pwwn 21:00:00:24:ff:50:ab:b4
!               [VNX5600-P2]

zone name SP-Shared-3_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:06
!               [SP-Shared-3_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-Shared-5_ZONE_B vsan 1

    member pwwn 20:00:00:25:b5:f2:20:08
!               [SP-Shared-5_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-Shared-6_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:0a
!               [SP-Shared-6_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-Shared-7_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:0c
!               [SP-Shared-7_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]


zone name SP-FarmIso-T1-2_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:0e
!               [SP-FarmIso-T1-2_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-FarmIso-T1-1_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:10
!               [SP-FarmIso-T1-1_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]
```

```
zone name SP-FarmIso-T1-3_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:12
!               [SP-FarmIso-T1-3_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]


_[K
zone name SP-FarmIso-T2-1_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:14
!               [SP-FarmIso-T2-1_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-FarmIso-T2-2_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:16
!               [SP-FarmIso-T2-2_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name SP-FarmIso-T2-3_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:18
!               [SP-FarmIso-T2-3_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88

!               [vnx5600_SPB_21]

zone name SP-FarmIso-T2-4_ZONE_B vsan 1
    member pwwn 20:00:00:25:b5:f2:20:1a
!               [SP-FarmIso-T2-4_vHBA2]
    member pwwn 50:06:01:62:08:64:2f:88
!               [vnx5600_SPA_20]
    member pwwn 50:06:01:6a:08:64:2f:88
!               [vnx5600_SPB_21]

zone name B200M4-CH1-BL7-FC1 vsan 1
    member pwwn 20:00:00:25:b5:00:00:2f
!               [B200M4-CH1-BL7]
    member pwwn 21:00:00:24:ff:57:fa:c1
!               [VNX560024TB-P1]

zone name B200M4-CH1-BL6-FC1 vsan 1
    member pwwn 20:00:00:25:b5:03:00:13
!               [CH1-BL6]
    member pwwn 21:00:00:24:ff:57:fa:c1
!               [VNX560024TB-P1]

zoneset name UCS-VNX5600-B vsan 1

    member B200M3-CH1-BL8-FC1
```

```
    member B200M3-CH2-BL8-FC1
    member B200M3-CH1-BL1-FC1
    member B200M3-CH1-BL2-FC1
    member B200M3-CH1-BL3-FC1
    member B200M3-CH1-BL4-FC1
    member B200M3-CH2-BL1-FC1
    member B200M3-CH2-BL2-FC1
    member B200M3-CH2-BL3-FC1
    member B200M3-CH2-BL4-FC1
    member B200M3-CH2-BL5-FC1
    member B200M3-CH1-BL5-FC1
    member SP-CH2-BL1_ZONE_B
    member SP-Shared-4_ZONE_B
    member SP-Shared-1_ZONE_B
    member SP-Shared-2_ZONE_B
    member B200M3-CH1-BL6-FC1
    member B200M3-CH1-BL7-FC1
    member B200M3-CH2-BL6-FC1
    member B200M3-CH2-BL7-FC1
    member SP-Shared-3_ZONE_B
    member SP-Shared-5_ZONE_B
    member SP-Shared-6_ZONE_B

    member SP-Shared-7_ZONE_B
    member SP-FarmIso-T1-2_ZONE_B
    member SP-FarmIso-T1-1_ZONE_B
    member SP-FarmIso-T1-3_ZONE_B
    member SP-FarmIso-T2-1_ZONE_B
    member SP-FarmIso-T2-2_ZONE_B
    member SP-FarmIso-T2-3_ZONE_B
    member SP-FarmIso-T2-4_ZONE_B
    member B200M4-CH1-BL7-FC1
    member B200M4-CH1-BL6-FC1

zoneset activate name UCS-VNX5600-B vsan 1


N5KB-EXC-HyperV#  exit
```