Deploy Cisco UCS X210c Compute Node with Cisco Intersight Management Mode for VDI

Last Updated: July 6, 2021

Contents

Executive summary

The beating heart of innovation, applications are the direct expression of consumer demand and your presence. Apps have forced a whole new way of thinking on IT. The old rules no longer apply. It's not on premises vs. cloud. It's both. It's massive scale and granular control. It's always-on availability enabled by modular components that can be molded and shaped to the needs of your applications. It's a new game where developers are free to write their own rules but where IT must free itself from the burden of managing massive complexity in order to keep up. With the X-Series platform, Virtual Desktop Infrastructure architects will have flexible options to create a robust solution for all VDI User profiles from task workers to high end graphics workstations.

Tomorrow thinks differently, wanting hardware that thinks like software. The Cisco UCS® X-Series with Intersight is a modular system managed from the cloud. It is designed be shaped to meet the needs of modern applications and improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design.

Designed to be managed exclusively from the cloud

- Simplify with cloud-operated infrastructure
- Simplify with an adaptable system designed for modern applications
- Simplify with a system engineered for the future

Overview

This section describes the Cisco® components used in the architecture.

Cisco Intersight platform

The Cisco Intersight™ platform is a software-as-a-service (SaaS) infrastructure lifecycle management solution that delivers simplified configuration, deployment, maintenance, and support. With the Cisco Intersight platform, customers get all the benefits of SaaS delivery and the full lifecycle management of Cisco Intersight connected distributed servers and third-party storage systems across data centers, remote sites, branch offices, and edge environments (Figure 1).

The Cisco Intersight platform is designed to be modular, so customers can adopt services based on their individual requirements. The platform significantly simplifies IT operations by bridging applications with infrastructure, providing visibility and management from bare-metal servers and hypervisors to serverless applications, thereby reducing costs and mitigating risk. This unified SaaS platform uses OpenAPI to provide a unified interface that natively integrates with third-party platforms and tools.

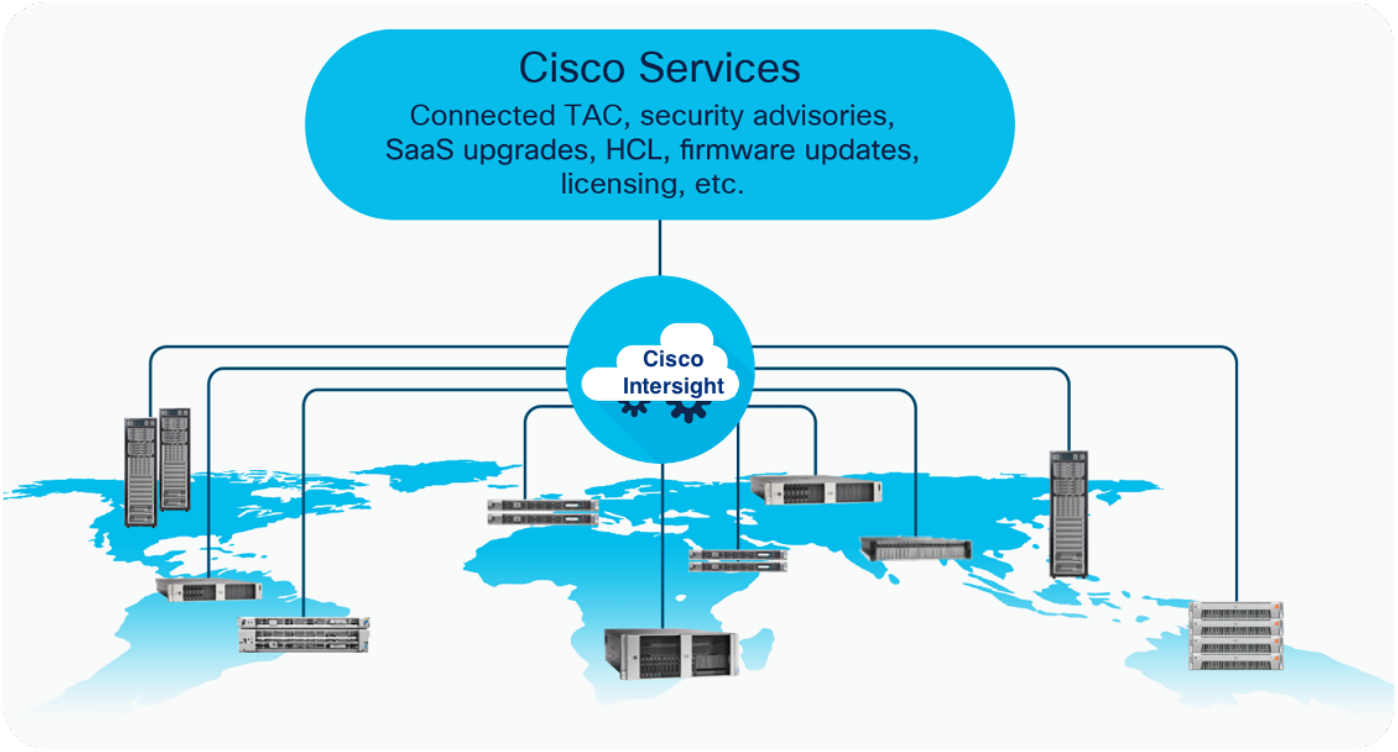

Figure 1.
Cisco Intersight overview

The main benefits of Cisco Intersight infrastructure services are summarized here:

- Simplify daily operations by automating many daily manual tasks.
- Combine the convenience of a SaaS platform with the capability to connect from anywhere and manage infrastructure through a browser or mobile app.
- Stay ahead of problems and accelerate trouble resolution through advanced support capabilities.
- Gain global visibility of infrastructure health and status along with advanced management and support capabilities.
- Upgrade to add workload optimization and Kubernetes services when needed.

The Cisco Unified Computing System™ (Cisco UCS®) is a next-generation data center platform that unites computing, networking, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency; lossless 25 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. It is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain.

Cisco UCS platform

The main components of Cisco UCS are as follows:

- Computing: The system is based on an entirely new class of computing system that incorporates blade servers based on Intel® Xeon® Scalable family processors.
- Network: The system is integrated on a low-latency, lossless, 25 Gigabit Ethernet unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables needed, and by decreasing the power and cooling requirements.
- Virtualization: The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- Storage access: The system provides consolidated access to local storage, SAN storage, and network-attached storage (NAS) over the unified fabric. With storage access unified, Cisco UCS can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and Small Computer System Interface over IP (iSCSI) protocols. This capability provides customers with choice for storage access and investment protection. In addition, server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management and helping increase productivity.

Cisco UCS is designed to deliver these benefits:

- Reduced TCO and increased business agility
- Increased IT staff productivity through just-in-time provisioning and mobility support
- A cohesive, integrated system that unifies the technology in the data center; the system is managed, serviced, and tested as a whole
- Scalability through a design for hundreds of discrete servers and thousands of virtual machines and the capability to scale I/O bandwidth to match demand
- Industry standards supported by a partner ecosystem of industry leaders

Cisco UCS 6400 Series Fabric Interconnects

The Cisco UCS 6400 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system (Figure 2 and Figure 3). The Cisco UCS 6400 Series offer line-rate, low-latency, lossless 10, 25, 40, and 100 Gigabit Ethernet; FCoE; and Fibre Channel functions.

The Cisco UCS 6400 Series provide the management and communication backbone for the Cisco UCS X-Series X9508, Cisco UCS B-Series Blade Servers, 5108 Blade Server Chassis, C-Series Rack Servers, and S-Series Storage Servers. All servers attached to a Cisco UCS 6400 Series Fabric Interconnect become part of a single, highly available management domain. In addition, by supporting a unified fabric, Cisco UCS 6400 Series Fabric Interconnects provide both the LAN and SAN connectivity for all servers within the system's domain.

From a networking perspective, the Cisco UCS 6400 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10, 25, 40, and 100 Gigabit Ethernet ports, switching capacity of 3.82 Tbps for the Cisco UCS 6454 Fabric Interconnect, 7.42 Tbps for the Cisco UCS 64108 Fabric Interconnect, and 200 Gbps bandwidth between the 6400 Series Fabric Interconnect and the Cisco UCS Fabric Extender for each Server Chassis, independent of packet size and enabled services. The product family supports Cisco low-latency, lossless 10, 25, 40, and 100 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the server through the fabric interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Figure 2.
Cisco UCS 6400 Series Fabric Interconnects: Cisco UCS 6454 Fabric Interconnect front view

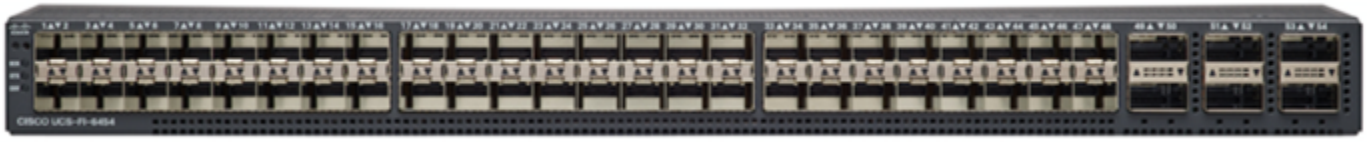

Figure 3.
Cisco UCS 6400 Series Fabric Interconnects: Cisco UCS 6454 Fabric Interconnect rear view

Cisco UCS X-Series Modular System

The Cisco UCS® X-Series Modular System simplifies your data center, adapting to the unpredictable needs of modern applications while also accommodating traditional scale-out and enterprise workloads. It reduces the number of server types you need to maintain, helping improve operational efficiency and agility by reducing complexity. Powered by the Cisco Intersight cloud-operations platform, it shifts users' IT focus from administrative details to business outcomes with a hybrid-cloud infrastructure that is assembled from the cloud, shaped to users' workloads, and continuously optimized.

Cisco UCS X9508 Chassis

The Cisco UCS X-Series Modular System begins with the Cisco UCS X9508 Chassis (Figure 4), engineered to be adaptable and future ready. The X-Series is a standards-based open system designed to be deployed and automated quickly in a hybrid cloud environment.

With a midplane-free design, I/O connectivity for the X9508 Chassis is accomplished with front-loading vertically oriented computing nodes that intersect with horizontally oriented I/O connectivity modules in the rear of the chassis. A unified Ethernet fabric is supplied with the Cisco UCS 9108 Intelligent Fabric Modules. In the future, Cisco UCS X-Fabric Technology interconnects will supply other industry-standard protocols as standards emerge. Interconnections can easily be updated with new modules.

The Cisco UCS X9508 Chassis provides these features and benefits:

- The seven-rack-unit (7RU) chassis has eight front-facing flexible slots. These can house a combination of computing nodes and a pool of future I/O resources, which may include graphics processing unit (GPU) accelerators, disk storage, and nonvolatile memory.
- Two Cisco UCS 9108 Intelligent Fabric Modules at the top of the chassis connect the chassis to upstream Cisco UCS 6400 Series Fabric Interconnects. Each Intelligent Fabric Module offers these features:
  - The module provides up to 100 Gbps of unified fabric connectivity per computing node.
  - The module provides eight 25-Gbps Small Form-Factor Pluggable 28 (SFP28) uplink ports.
  - The unified fabric carries management traffic to the Cisco Intersight cloud-operations platform, FCoE traffic, and production Ethernet traffic to the fabric interconnects.
- At the bottom of the chassis are slots ready to house future I/O modules that can flexibly connect the computing modules with I/O devices. Cisco calls this connectivity Cisco UCS X-Fabric technology, because “X” is commonly used as a variable, signifying a system that can evolve with new technology developments.
- Six 2800-watt (W) power supply units (PSUs) provide 54 volts (V) of power to the chassis with N, N+1, and N+N redundancy. A higher voltage allows efficient power delivery with less copper wiring needed and reduced power loss.
- Efficient, 4 x 100–mm, dual counter-rotating fans deliver industry-leading airflow and power efficiency. Optimized thermal algorithms enable different cooling modes to best support the network environment. Cooling is modular, so future enhancements can potentially handle open- or closed-loop liquid cooling to support even higher-power processors.

Figure 4.
Cisco UCS 9508 X-Series Chassis, front (left) and back (right)

Since Cisco first delivered the Cisco Unified Computing System in 2009, our goal has been to simplify the data center. We pulled management out of servers and into the network. We simplified multiple networks into a single unified fabric. And we eliminated network layers in favor of a flat topology wrapped into a single unified system. With the Cisco UCS X-Series Modular System, the simplicity is extended even further:

- Simplify with cloud-operated infrastructure. We move management from the network into the cloud so that you can respond at the speed and scale of your business and manage all your infrastructure.

You can shape Cisco UCS X-Series Modular System resources to workload requirements with the Cisco Intersight cloud-operations platform. You can integrate third-party devices, including storage from NetApp, Pure Storage, and Hitachi. In addition, you gain intelligent visualization, optimization, and orchestration for all your applications and infrastructure.

- Simplify with an adaptable system designed for modern applications. Today's cloud-native, hybrid applications are inherently unpredictable. They are deployed and redeployed as part of an iterative DevOps practice. Requirements change often, and you need a system that doesn't lock you in to one set of resources when you find that you need a different set. For hybrid applications, and for a range of traditional data center applications, you can consolidate your resources on a single platform that combines the density and efficiency of blade servers with the expandability of rack servers. The result is better performance, automation, and efficiency.
- Simplify with a system engineered for the future. Embrace emerging technology and reduce risk with a modular system designed to support future generations of processors, storage, nonvolatile memory, accelerators, and interconnects. Gone is the need to purchase, configure, maintain, power, and cool discrete management modules and servers. Cloud-based management is kept up-to-date automatically with a constant stream of new capabilities delivered by the Cisco Intersight SaaS model.
- Support a broader range of workloads. A single server type supporting a broader range of workloads means fewer different products to support, reduced training costs, and increased flexibility.

Cisco UCS X210c Series Servers

The Cisco UCS X-Series Modular System simplifies your data center, adapting to the unpredictable needs of modern applications while also accommodating traditional scale-out and enterprise workloads. It reduces the number of server types that you need to maintain, helping to improve operational efficiency and agility by reducing complexity. Powered by the Cisco Intersight cloud operations platform, it shifts your thinking from administrative details to business outcomes with hybrid cloud infrastructure that is assembled from the cloud, shaped to your workloads, and continuously optimized.

The Cisco UCS X210c M6 Compute Node is the first computing device integrated into the Cisco UCS X-Series Modular System. Up to eight computing nodes can reside in the 7RU Cisco UCS X9508 Chassis, offering one of the highest densities of computing, I/O, and storage resources per rack unit in the industry. The Cisco UCS X210c harnesses the power of the latest Third-Generation (3rd Gen) Intel Xeon Scalable processors (Ice Lake). It includes the following features:

- CPU: Install up to two 3rd Gen Intel Xeon Scalable processors with up to 40 cores per processor and 1.5 MB of Level 3 cache per core.
- Memory: Install up to thirty-two 256-GB DDR4 3200-MHz DIMMs for up to 8 TB of main memory. Configuring up to sixteen 512-GB Intel Optane™ persistent-memory DIMMs can yield up to 12 TB of memory.
- Storage: Install up to six hot-pluggable solid-state disks (SSDs) or Non-Volatile Memory Express (NVMe) 2.5-inch drives with a choice of enterprise-class RAID or pass-through controllers with four lanes each of PCIe Gen 4 connectivity and up to two M.2 SATA drives for flexible boot and local storage capabilities.

- Modular LAN-on-motherboard (mLOM) virtual interface card (VIC): The Cisco UCS VIC 14425 occupies the server's mLOM slot, enabling up to 50-Gbps unified fabric connectivity to each of the chassis Intelligent Fabric Modules for 100-Gbps connectivity per server.
- Optional mezzanine VIC: The Cisco UCS VIC 14825 can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric technology that is planned for future I/O expansion. An included bridge card extends this VIC's two 50-Gbps network connections through Intelligent Fabric Module connectors, bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per server).
- Security: The server supports an optional Trusted Platform Module (TPM). Additional features include a secure boot field-programmable gateway (FPGA) and Anti-Counterfeit Technology 2 (ACT2) provisions

Cisco Nexus 93180YC-FX Switch

The Cisco Nexus® 93180YC-FX Switch (Figure 5) provides a flexible line-rate Layer 2 and Layer 3 feature set in a compact form factor. Designed with Cisco Cloud Scale technology, it supports highly scalable cloud architectures. With the option to operate in Cisco NX-OS or Application Centric Infrastructure (Cisco ACI™) mode, it can be deployed across enterprise, service provider, and Web 2.0 data centers.

- Architectural flexibility
  - Deploy top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures.
  - Leaf-node support for Cisco ACI architecture is on the roadmap.
  - Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support.
- Comprehensive feature set
  - Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability.
  - Cisco ACI ready infrastructure helps users take advantage of automated policy-based systems management.
  - Virtual Extensible LAN (VXLAN) routing provides network services.
  - Robust traffic-flow telemetry provides line-rate data collection.
  - Real-time buffer utilization reporting per port and per queue lets you monitor traffic microbursts and application traffic patterns.
- Highly available and efficient design
  - The switch uses a high-density, nonblocking architecture.
  - Easily deploy the switch in either a hot-aisle or cold-aisle configuration.
  - Redundant, hot-swappable power supplies and fan trays help protect your system.
- Simplified operations
  - Power-on autoprovisioning (POAP) support simplifies software upgrades and configuration-file installation.

- An intelligent API offers switch management through remote procedure calls (RPCs), JavaScript Object Notation (JSON), or XML over HTTP/HTTPS infrastructure.
- Python scripting provides programmatic access to the switch CLI.
- The switch supports hot and cold patching and online diagnostics.

- Investment protection
  - A Cisco 25 Gigabit Ethernet bidirectional transceiver allows reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 25 Gigabit Ethernet support for 1 and 10 Gigabit Ethernet access connectivity for data centers migrating access switching infrastructure to faster speeds.

The Cisco Nexus 93180YC-X includes the following:

- 1.8 Tbps of bandwidth in a 1RU form factor
- 48 fixed 1, 10, and 25 Gigabit Ethernet Enhanced SFP (SFP+) ports
- 6 fixed 40 and 100 Gigabit Ethernet Quad SFP+ (QSFP+) ports for uplink connectivity
- Latency of less than 2 microseconds
- Front-to-back or back-to-front airflow configurations
- 1+1 redundant hot-swappable 80 Plus Platinum-certified power supplies
- Hot swappable 3+1 redundant fan trays

Figure 5.
Cisco Nexus 93180YC-FX Switch

Solution design: Deploy fabric in Cisco Intersight managed mode

This section provides an overview of the infrastructure setup, software and hardware requirements, and some of the design details. This section does not cover the design details and configuration of components such as Cisco Nexus and Cisco MDS switches and storage array systems because their design and configuration conform to various Cisco Validated Designs for converged infrastructure and are covered widely elsewhere. This document focuses on the design elements and performance of the Intel platform for application in virtual desktop infrastructure (VDI).

Physical architecture

The architecture deployed in this solution is highly modular and follow Cisco Validated Design implementation principals for converged infrastructure. Although each customer's environment may vary in its exact configuration, after the architecture described in this document has been built, it can easily be scaled as requirements and demands change. It can be scaled both up (adding resources within a Cisco UCS domain) and out (adding Cisco UCS domains).

Figure 6 shows the physical architecture.

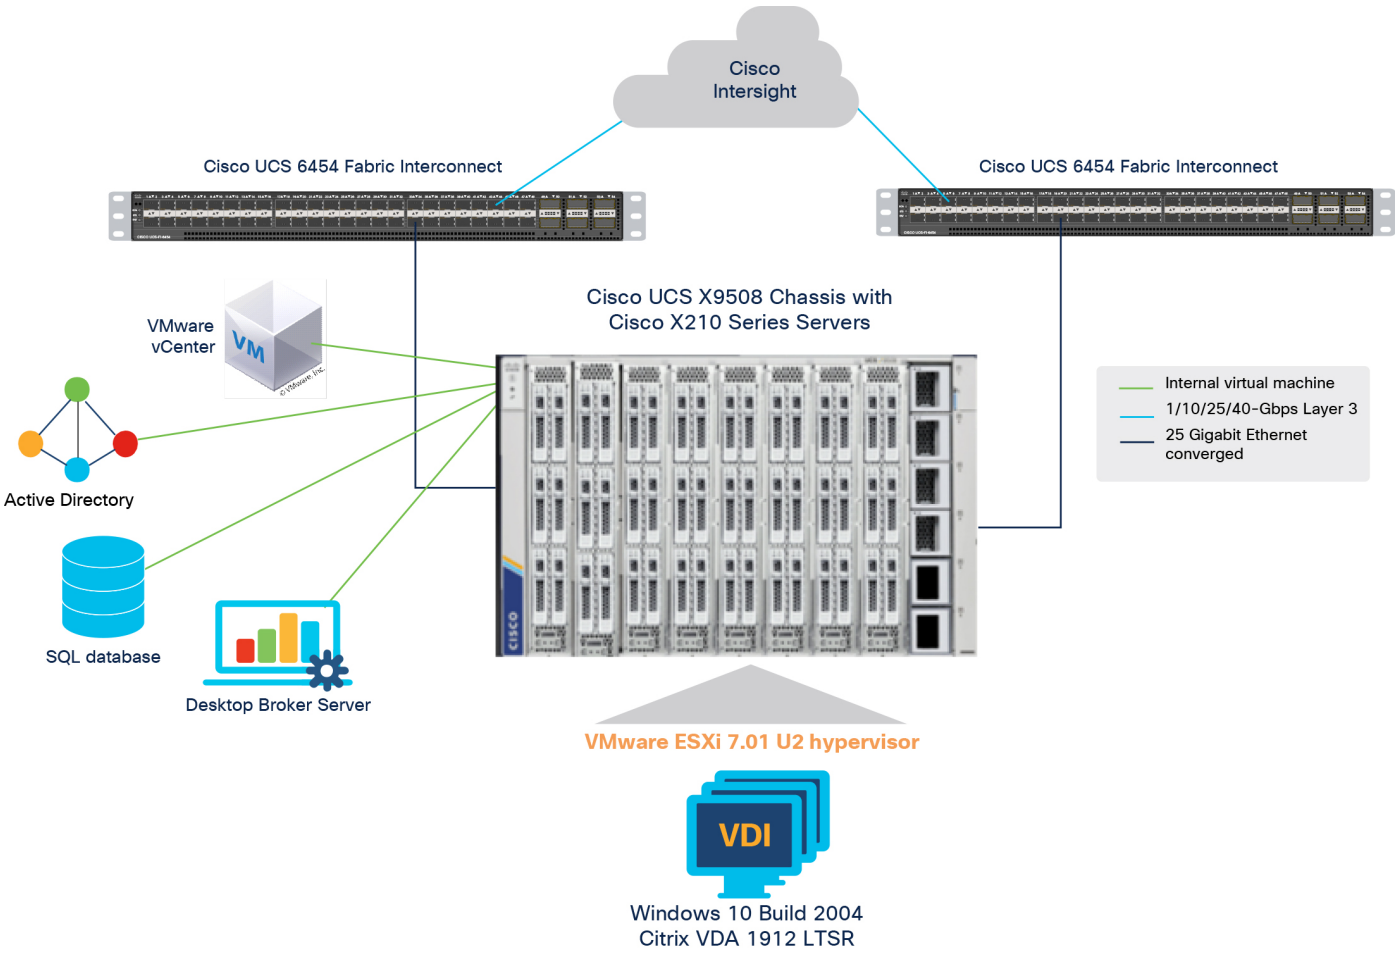

Figure 6.
Physical architecture

The following components are deployed:

- Two Cisco Nexus 93180YC-FX Switches
- Two Cisco UCS 6454 Fabric Interconnects
- Four Cisco UCS X210c Compute Nodes with 1 TB of DRAM memory

Logical architecture

The logical architecture (Figure 7) is configured to match the Cisco Validated Design standards to help ensure consistent VDI performance across all our solutions.

For desktop virtualization, the deployment includes Citrix 1912 Long-Term Service Release (LTSR) CU2 running on VMware vSphere ESXi 7.0.2.

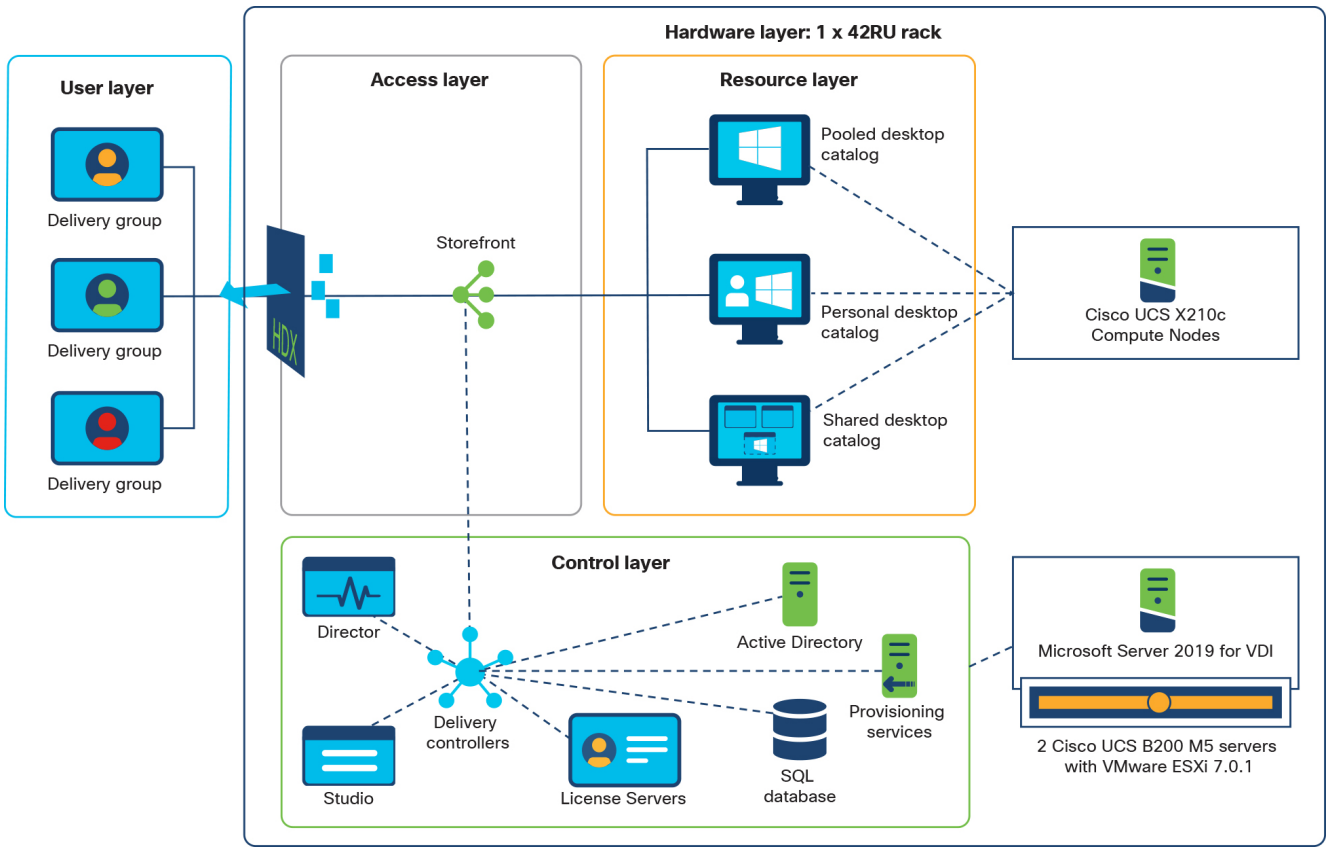

Figure 7.
Logical architecture

Table 1 lists the software and firmware versions used in the solution presented here.

Table 1. Software and firmware versions

Component	Version
Cisco UCS component firmware	Release 4.1(5g) bundle
Cisco Intersight management mode	Release 4.1(5g) bundle
Cisco UCS X210c M6 nodes	Release 4.1(5g) bundle
Cisco UCS VIC 14425	Release 4.1(5g) bundle
VMware vCenter Server Appliance 7.0.2	Release 17694817
VMware vSphere 7.0.2	Release 17867351
Citrix Virtual Apps & Desktops 1912 LTSR CU2	Release 1912.2000
Citrix Provisioning Services (PVS)	Release 1912.2000
Citrix Virtual Delivery Agent (VDA)	Release 1912.2000
Microsoft FSLogix for profile management	FSLogix_Apps_2.9.7654.46150

Deploy the Cisco Intersight platform

The Cisco Intersight platform provides an integrated and intuitive management experience for resources in the traditional data center and at the edge. With flexible deployment options to address complex security needs, getting started with Cisco Intersight is quick and easy.

To configure Cisco Intersight to use Cisco Intersight managed mode, follow these steps:

1. If you do not already have a Cisco Intersight account, you need to set up a new account in which to claim your Cisco UCS deployment. Start by connecting to https://intersight.com.

   If you have an existing Cisco Intersight account, connect to https://intersight.com and sign in with your Cisco ID, select the appropriate account, and skip to step 6.

2. Click Create an account.
3. Sign in with your Cisco ID.
4. Read, scroll through, and accept the end-user license agreement. Click Next.
5. Enter an account name and click Create.
6. Choose ADMIN > Targets. Click Claim a New Target.

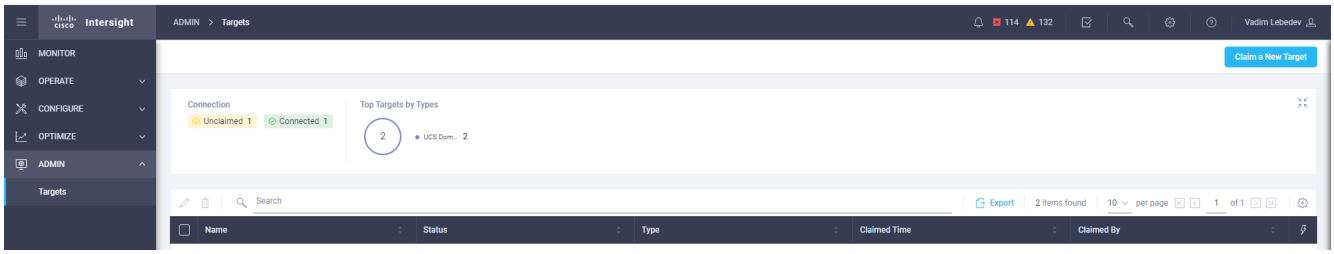

7. Select Cisco UCS Domain (Intersight Managed) and click Start.

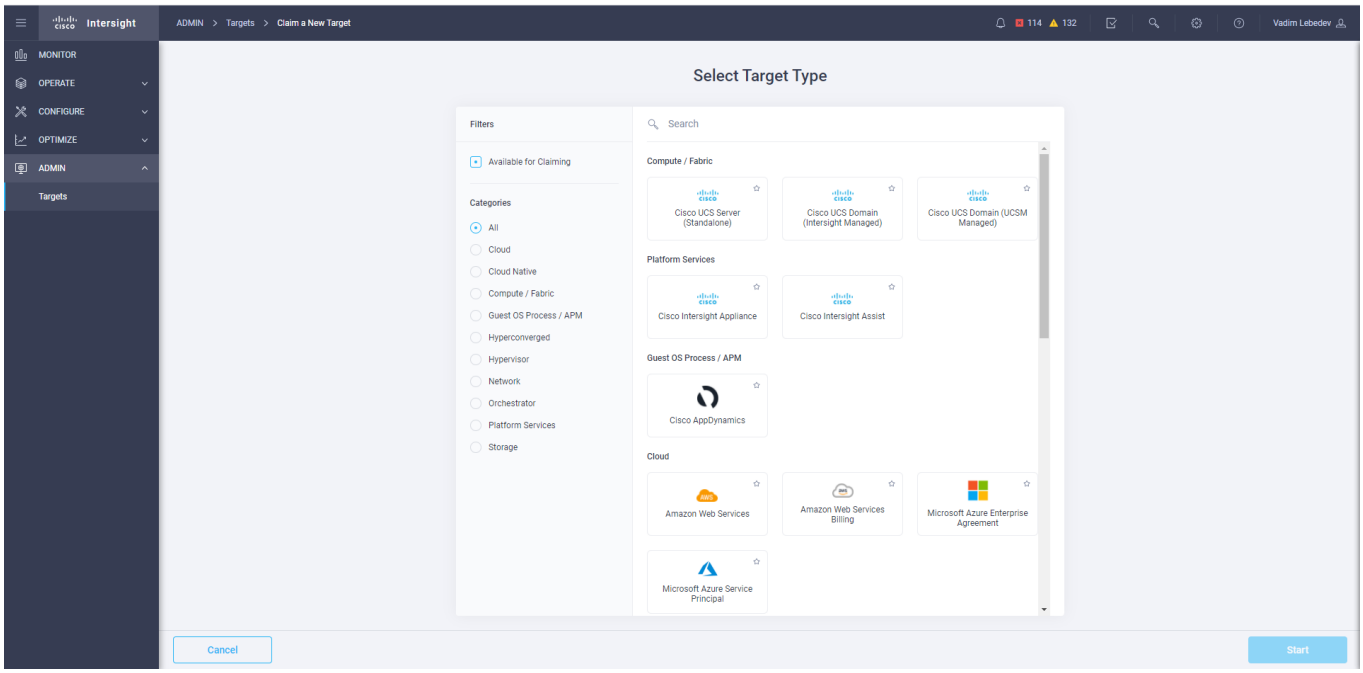

8. Fill in the Device ID and Claim Code fields and click Claim.

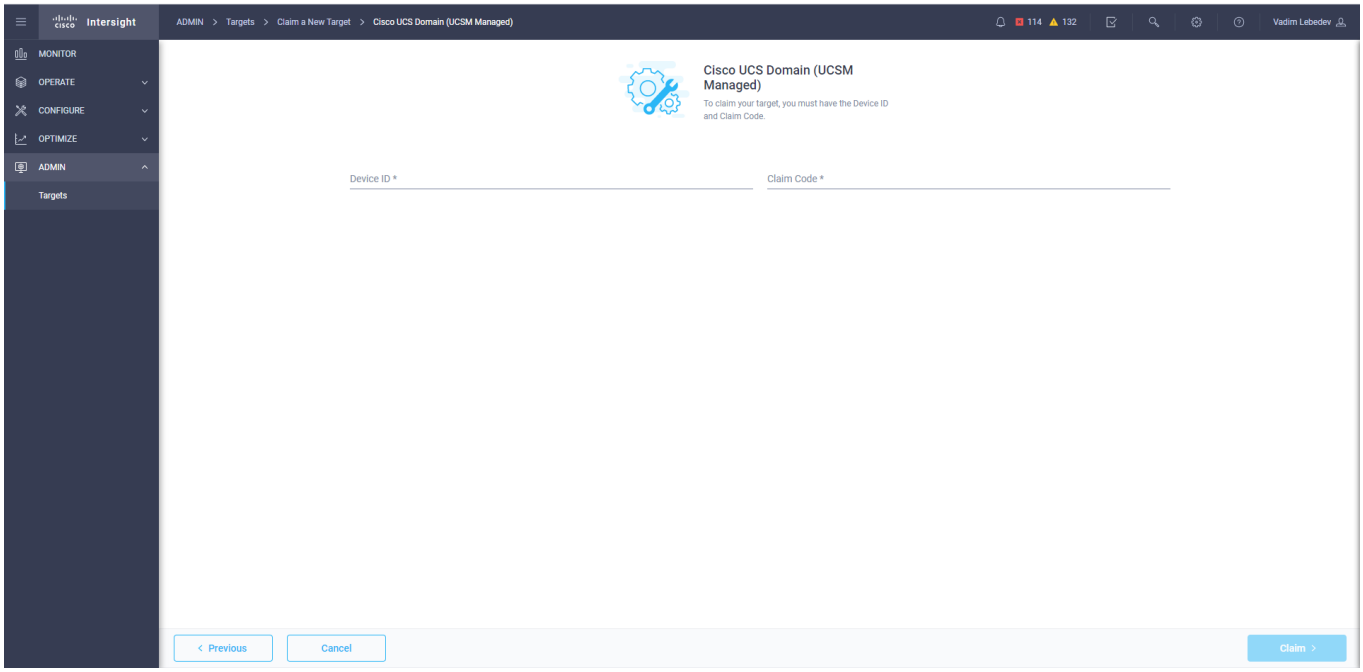

Note: To obtain the device ID and claim code, connect to Cisco Intersight and choose Admin > All > Device Connector. The device ID and claim code are on the right.

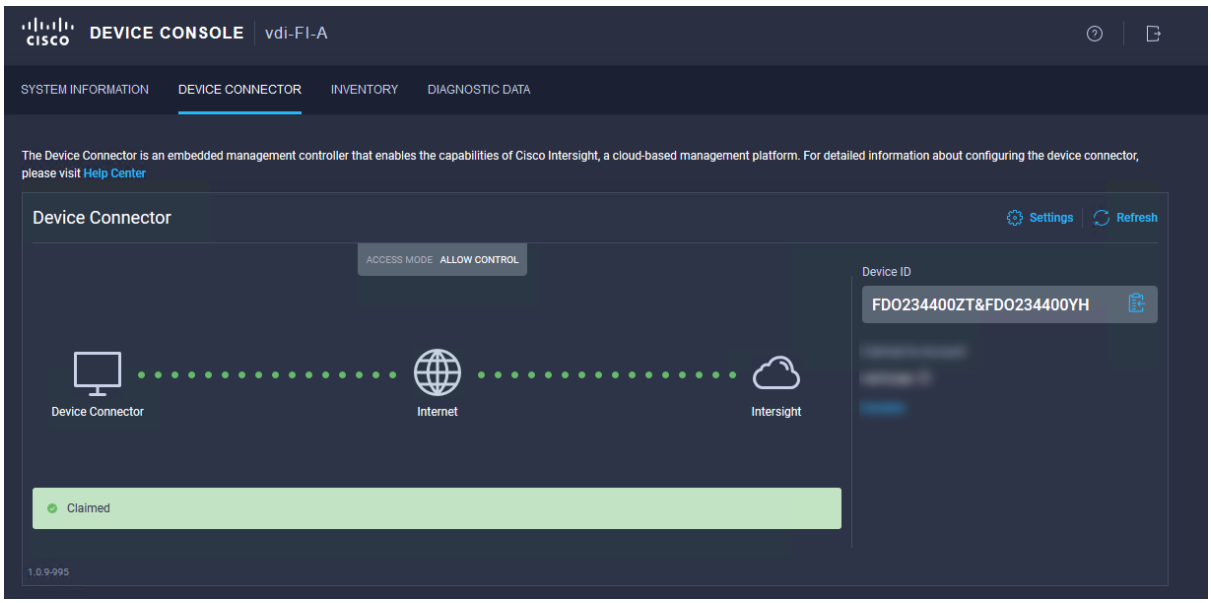

9. Verify that the target is visible in the list of available targets.

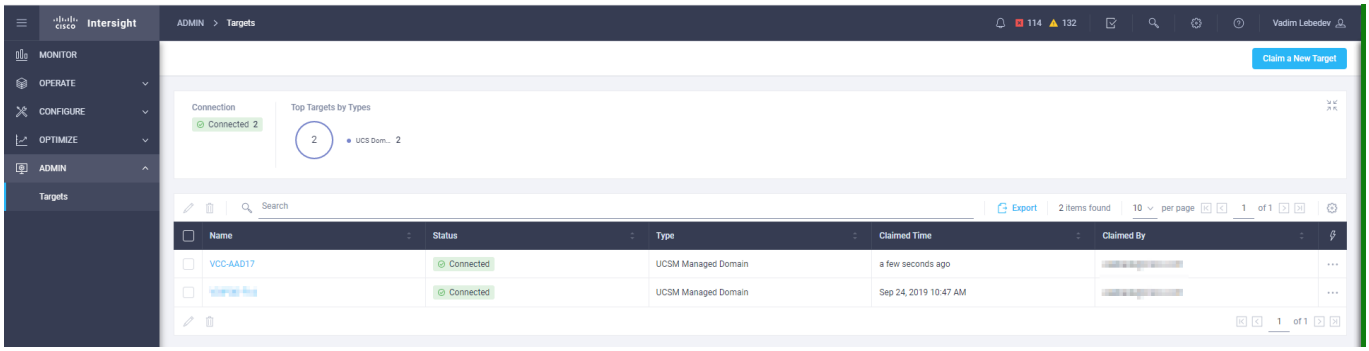

10. In the Cisco Intersight window, click the Settings button and choose Licensing. If this is a new account, all servers connected to the Cisco UCS domain will appear under the Base license tier. If you have purchased Cisco Intersight licenses and have them in your Cisco Smart Account, click Register and follow the prompts to register this Cisco Intersight account to your Cisco Smart Account. Cisco Intersight also offers a one-time 90-day trial of Premier licensing for new accounts. Click Start Trial and then Start to begin this evaluation. The remainder of this section will assume Premier licensing. Intersight Essentials licensing is required to run the X Series platform.

Create a Cisco UCS domain profile

Next you need to create a Cisco UCS domain profile to configure the fabric interconnect ports and discover connected chassis. A domain profile is composed of several policies. Table 2 lists the policies required for the solution described in this document.

Table 2. Policies required for domain profile

Policy	Description
VLAN multicast policy for each VLAN policy	Network connectivity
VSAN configuration policy for fabric A	Simple Network Management Protocol (SNMP)
VSAN configuration policy for fabric B	System quality of service (QoS)
Port configuration policy for fabric A	Switch control
Port configuration policy for fabric B	
Network Time Protocol (NTP) policy	
Syslog policy	

Follow these steps:

1. Create VLAN configuration policy.
    - Under Policies, select Create Policy. Then select VLAN and give the VLAN a name (for example, **VLAN_Config_Policy**).
    - Click Add VLANs to add your required VLANs.
    - Click Multicast Policy to add or create a multicast policy for your VLAN policy.
    - Click Create.

Note: If you will be using the same VLANs on fabric interconnect A and fabric interconnect B, you can use the same policy for both.

2. Create VSAN configuration policy.
    - Under Policies, select Create Policy. Then select VSAN and give the VSAN a name (for example, **VSAN_Config_Policy_A**).
    - Click Add VSAN and add the required VSAN names and IDs.
    - Add the corresponding FCoE VLAN ID.

Note: According to best practices, the VSAN IDs should be different for each fabric, so you should create two separate policies.

3. Create port configuration policy.
    - Under Policies, select Create Policy.
    - Select Port and Start.
    - Give the port policy a name.
    - Click Next.
    - If you need Fibre Channel, use the slider to define Fibre Channel ports. Then click Next.

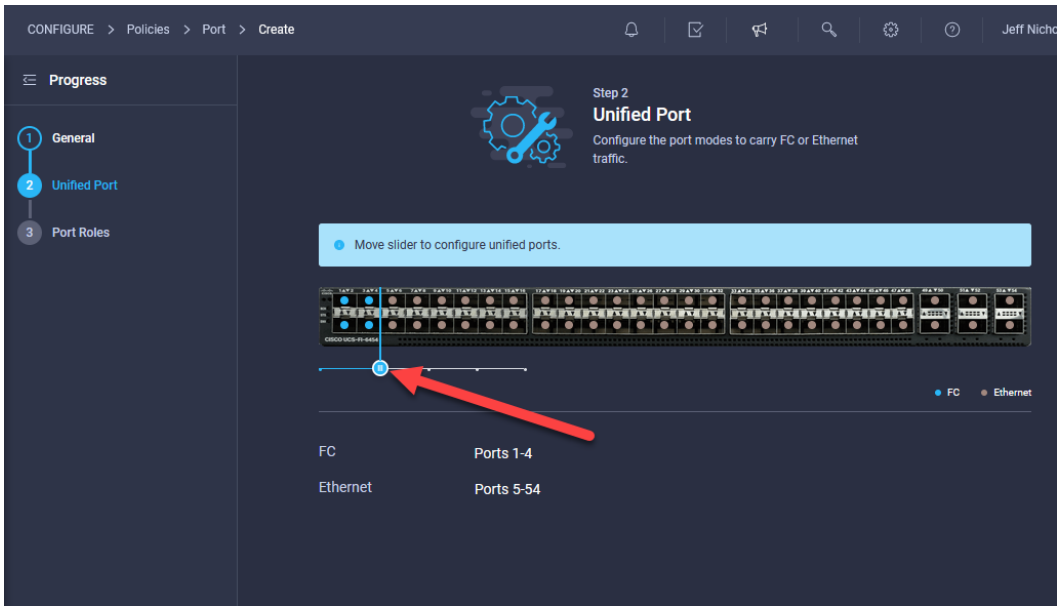

    - Define the port roles: server ports for chassis and server connections, Fibre Channel ports for SAN connections, or network uplink ports.
    - Select ports 1 through 4 and click Configure.

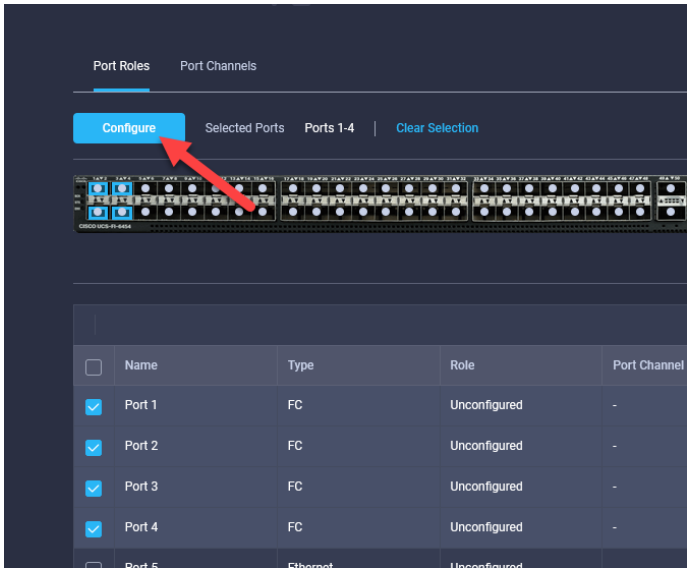

- From the drop-down menu, choose FC Uplink and define the admin speed and VSAN ID (the solution described here uses 32 GBps for the speed and VSAN ID **400** for fabric interconnect A and 401 for fabric interconnect B).

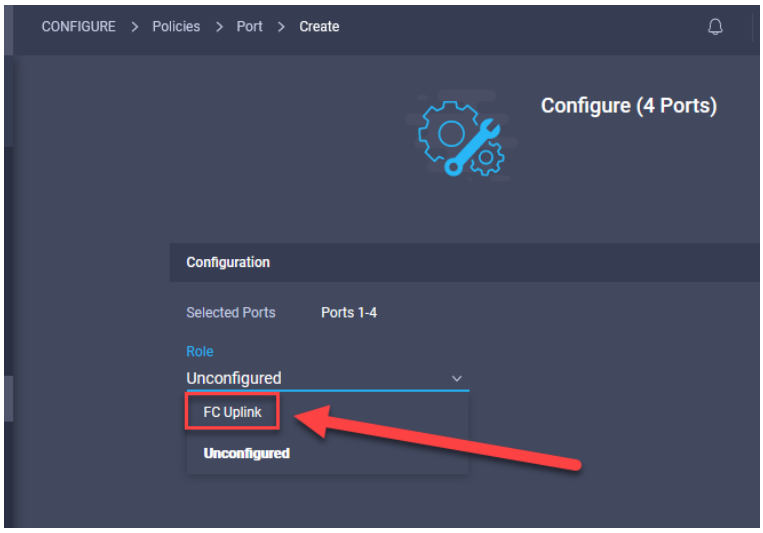

- To configure server ports, select the ports that have chassis or rack-mounted servers plugged into them and click Configure. From the drop-down menu, choose Server.

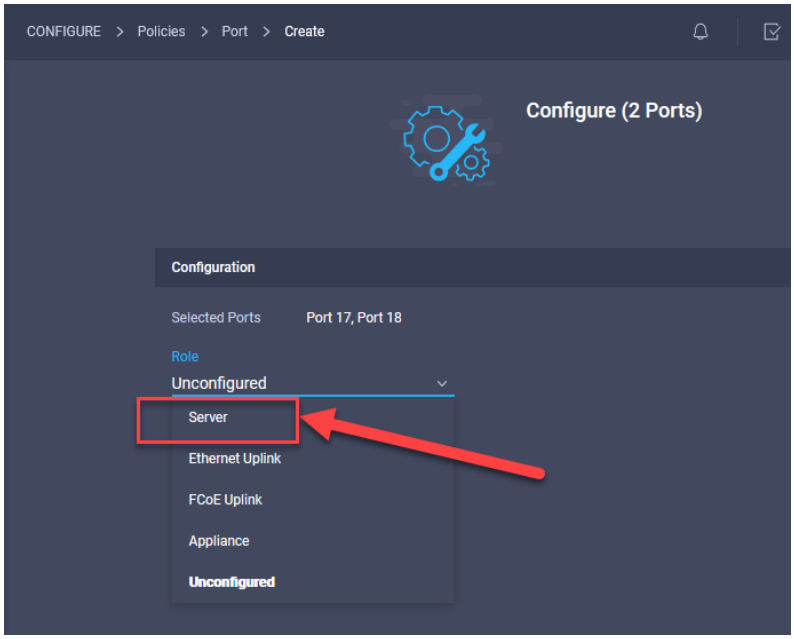

- To configure network uplink ports, select the ports connected to the upstream network switches and click Configure. From the drop-down menu, choose Ethernet Uplink.

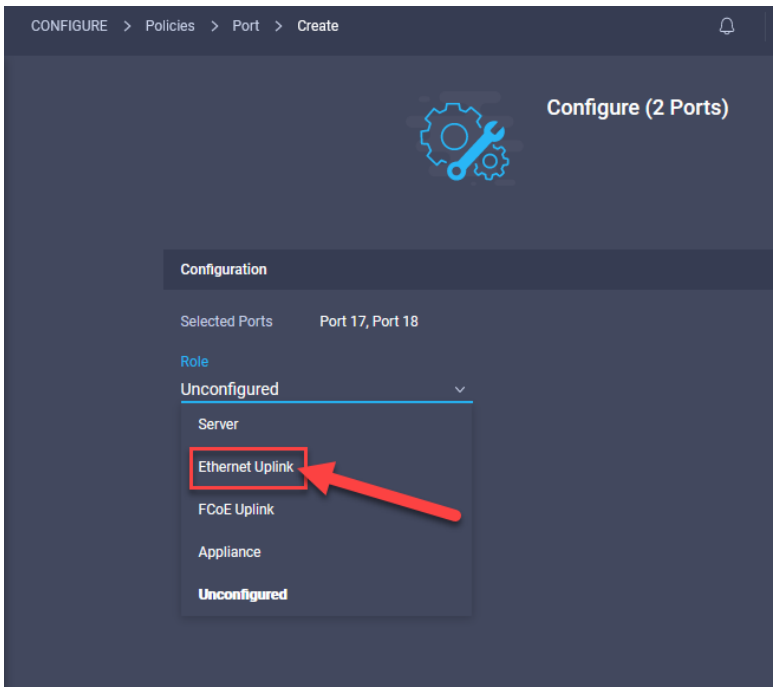

- Click Save to save the port policy.

Repeat the preceding steps to create a port policy for Fabric Interconnect B.

4. Create NTP policy.
   - Under Policies, select Create Policy.
   - Select NTP and Start.
   - Give the NTP policy a name.
   - Click Next.
   - Define the name or IP address for the NTP servers.
   - Define the correct time zone.
   - Click Create.
5. Create syslog policy.
   - Under Policies, select Create Policy.
   - Select Syslog and Start.
   - Give the syslog policy a name.
   - Click Next.
   - Define the syslog severity level that triggers a report.
   - Define the name or IP address for the syslog servers.
   - Click Create.

Note: You do not need to enable the syslog server.

6. Create the domain profile.
    - Click Profiles in the side menu.
    - Select UCS Domain Profiles at the top of the screen.
    - Click Create UCS Domain Profile.

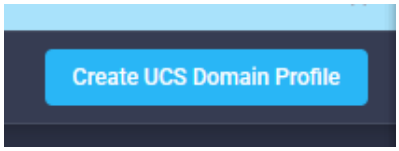

    - Give the profile a name (for example, **vdi_dom_prof**) and click Next.
    - Select the fabric interconnect domain pair created when you claimed your fabric interconnects.
    - Under VLAN & VSAN Configuration, click Select Policy to select the policies created earlier. (Be sure that you select the appropriate policy for each side of the fabric.)

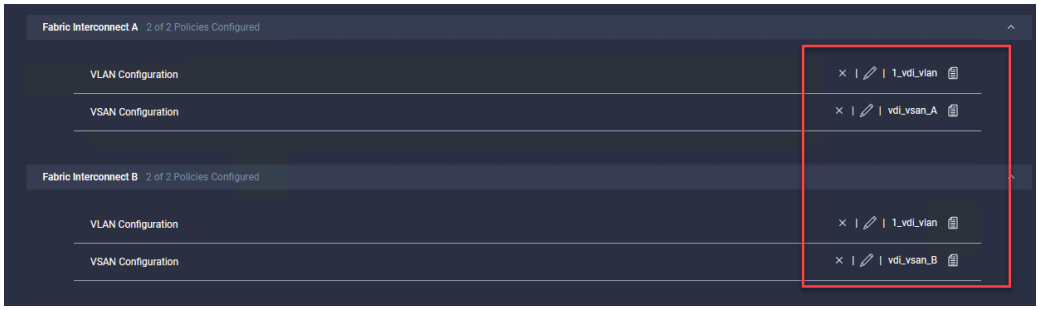

    - Under Ports Configuration, select the port configuration policies created earlier. (Be sure to assign the correct fabric policy to the appropriate side: that is, assign port_FIA to fabric interconnect A and assign port_FIB to fabric interconnect B)
    - Under UCS Domain Configuration, select all the policies you created earlier.

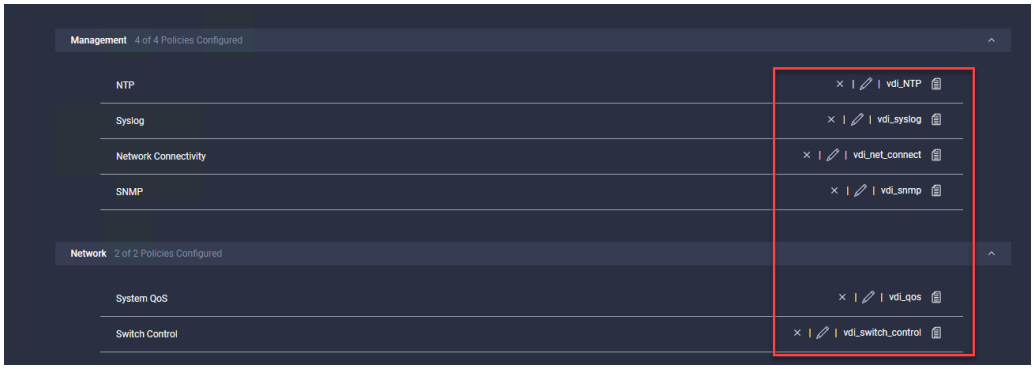

    - Click Next.
    - Click Deploy.

Create a UCS server profile

Much like Cisco UCS Manager, Cisco Intersight managed mode lets you to deploy server profiles that consist of a series of policies. This capability allows detailed management of the entire environment.

Figure 8 shows a list of the policies that make up a server profile.

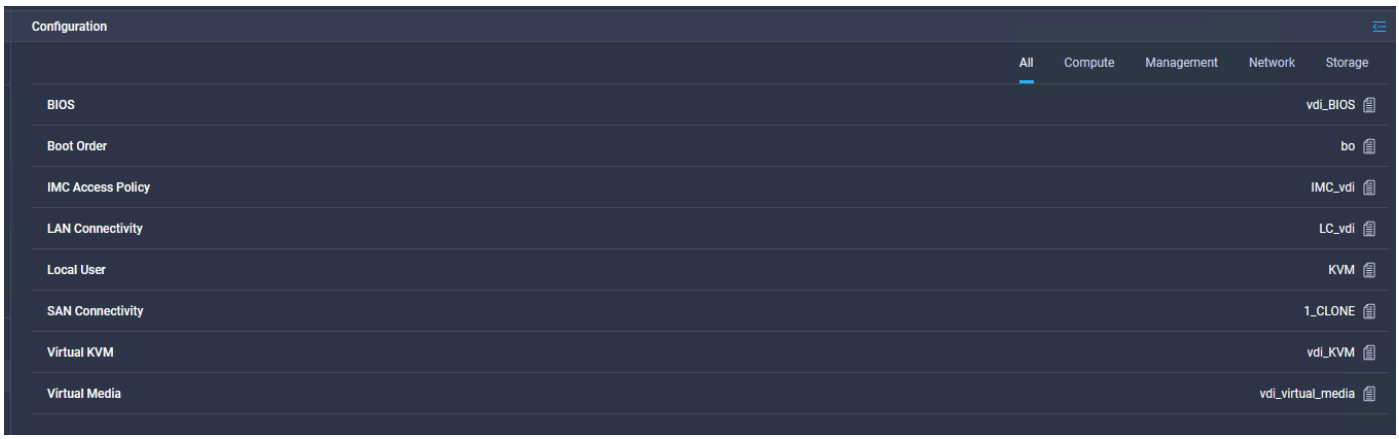

Figure 8.
Polices in a server profile

1. For the VDI testing described in this document, use platform-default for all settings except those in the Processor section.

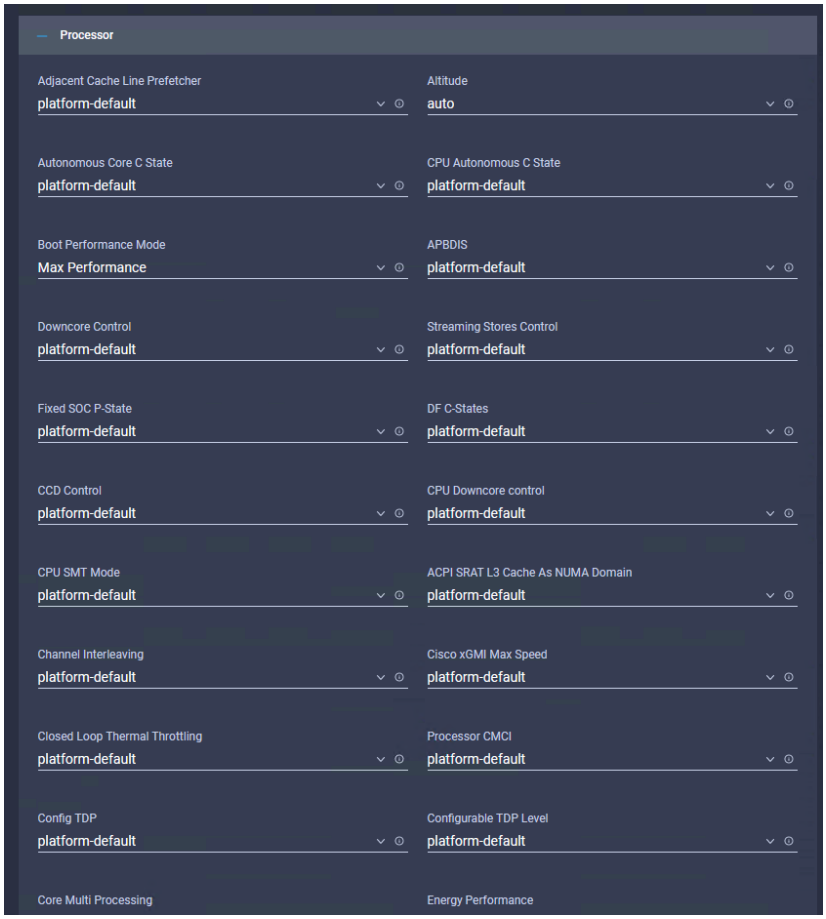

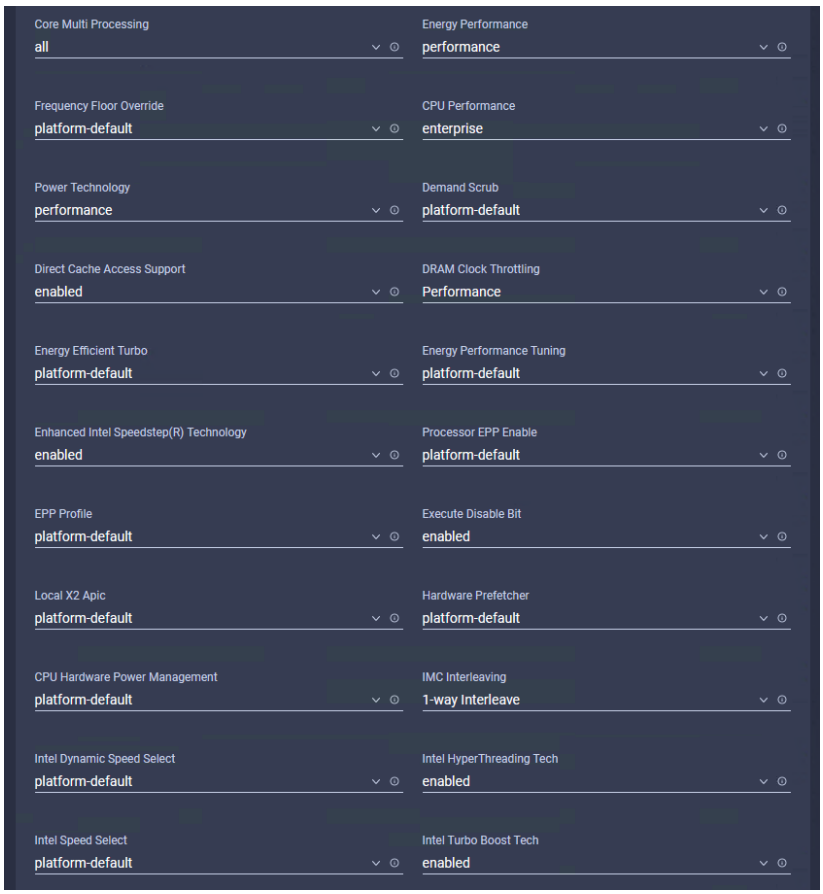

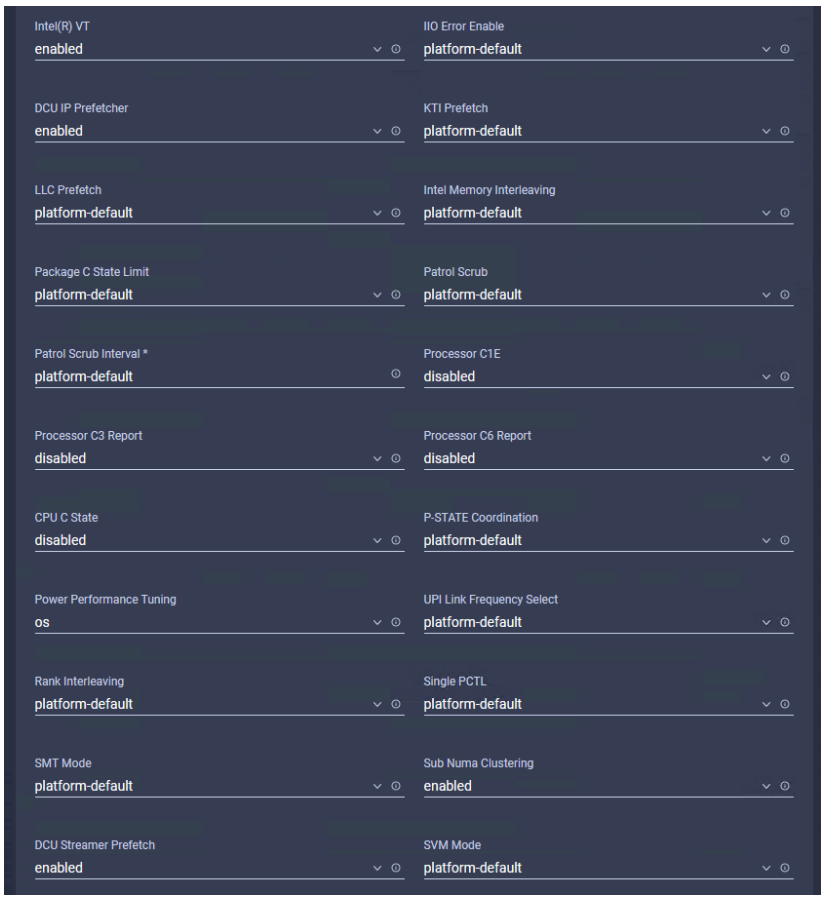

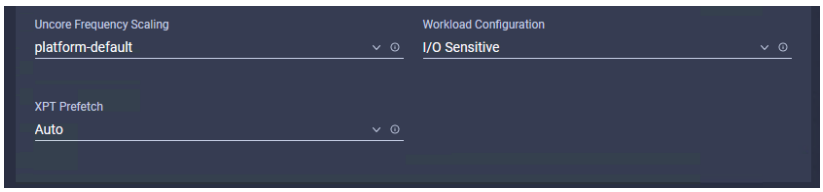

2. Define a boot order (M.2 cards were used in this testing).

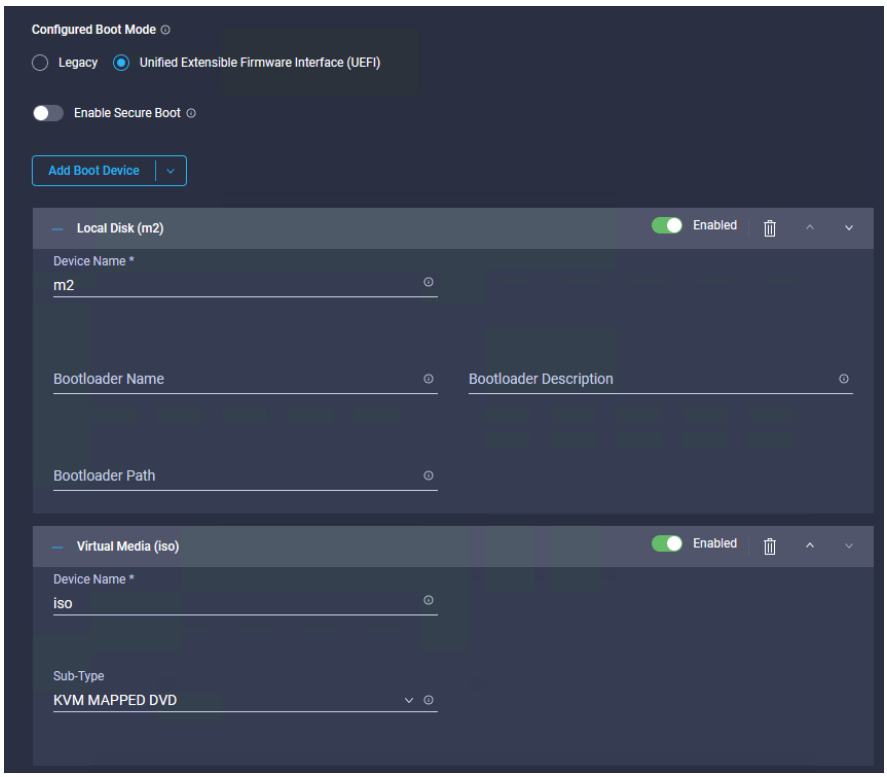

3. Define the Cisco Integrated Management Controller (IMC) access policy and IP address pool for the keyboard, video, and mouse (KVM).

4. For this policy, create an IP address pool under Pools.

5. Create a LAN connectivity policy. (This policy will have a sub-policy for every virtual NIC [vNIC] you need. The testing described here uses two vNICs in the LAN connectivity policy.)

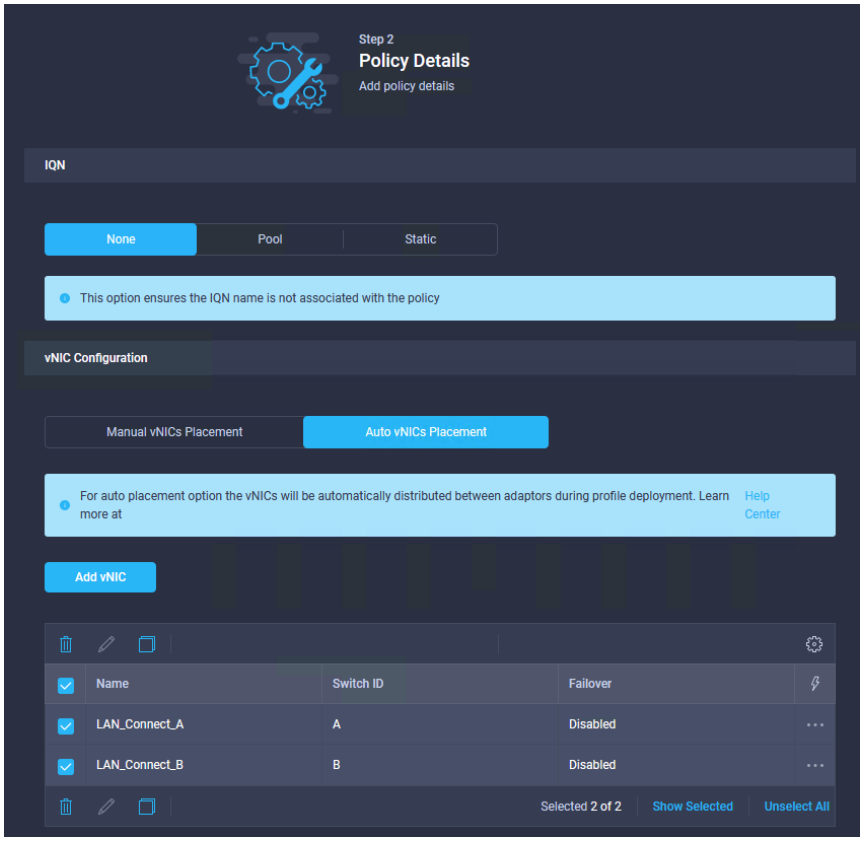

6. For vNIC Configuration, select Auto vNICs Placement.
7. Click Add vNIC. This policy requires you to create the following pools and policies:
    - Create the MAC address pool. (Create this pool under Pools. Create two pools: one for side A and one for side B.)

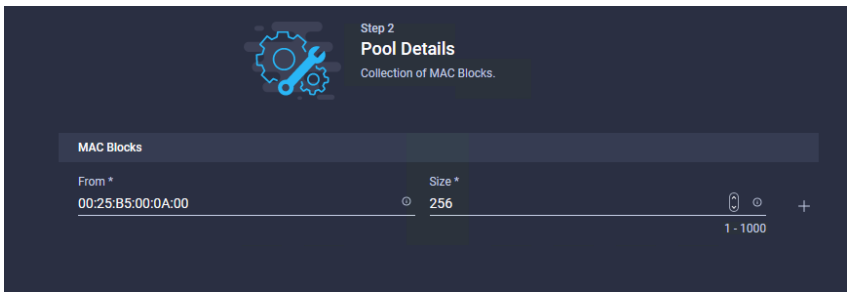

    - Create Ethernet network group policy.

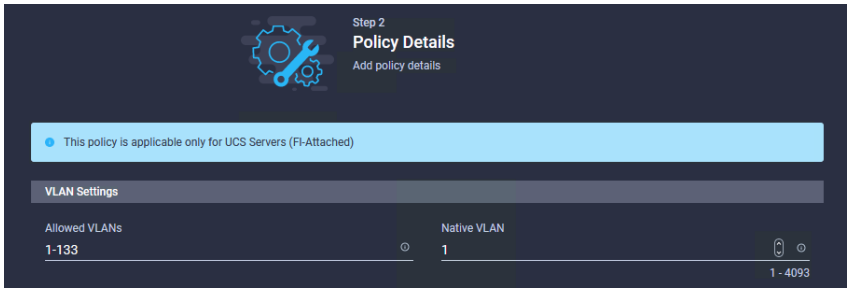

- Create Ethernet network control policy.

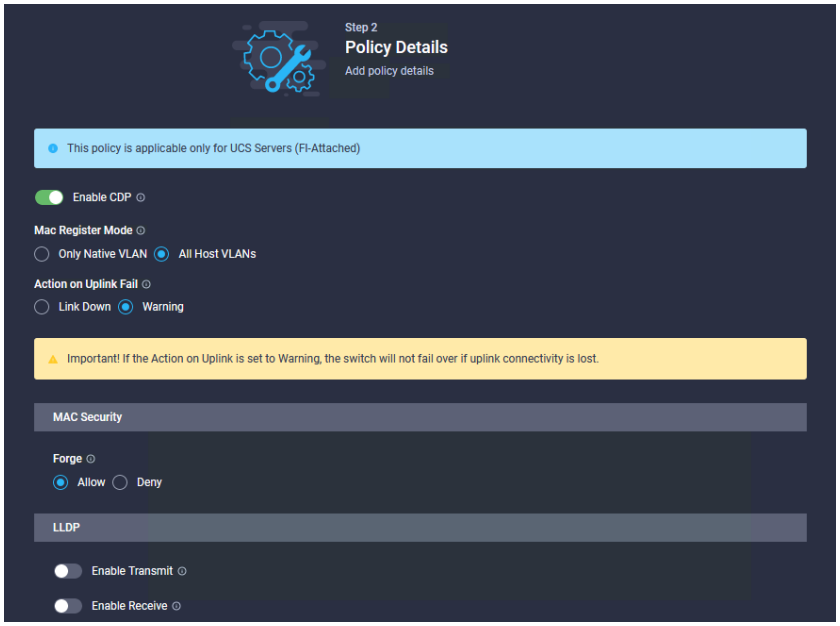

- Create Ethernet QoS policy.

- Create Ethernet adapter policy.

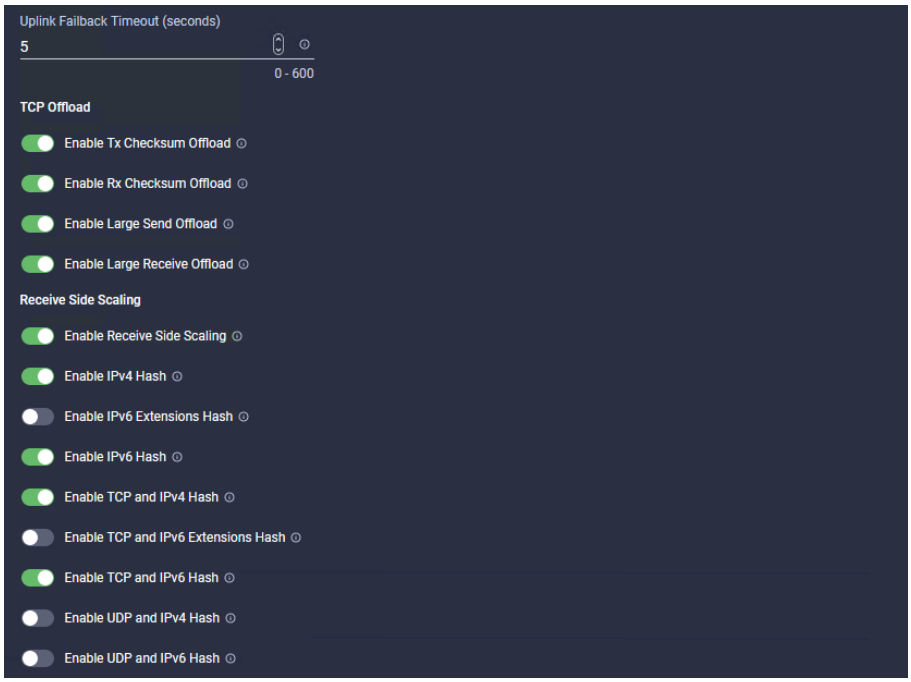

8. Create local-user policy. (This policy allows the use of the KVM console.)

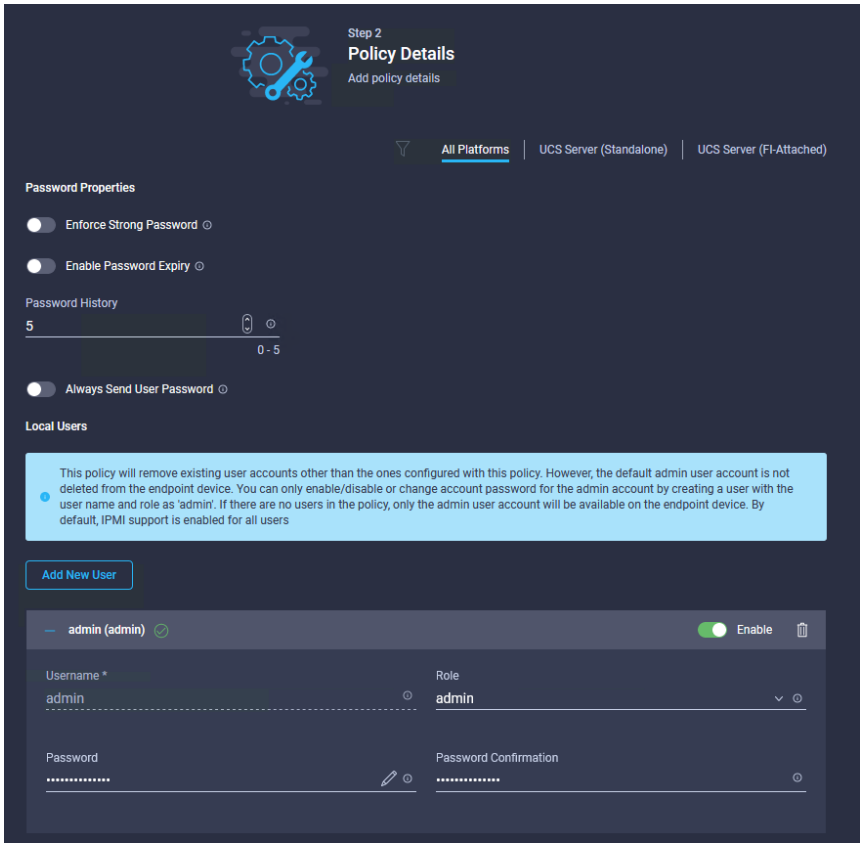

9. Create SAN connectivity policy. (This policy requires server subpolicies and pools.)

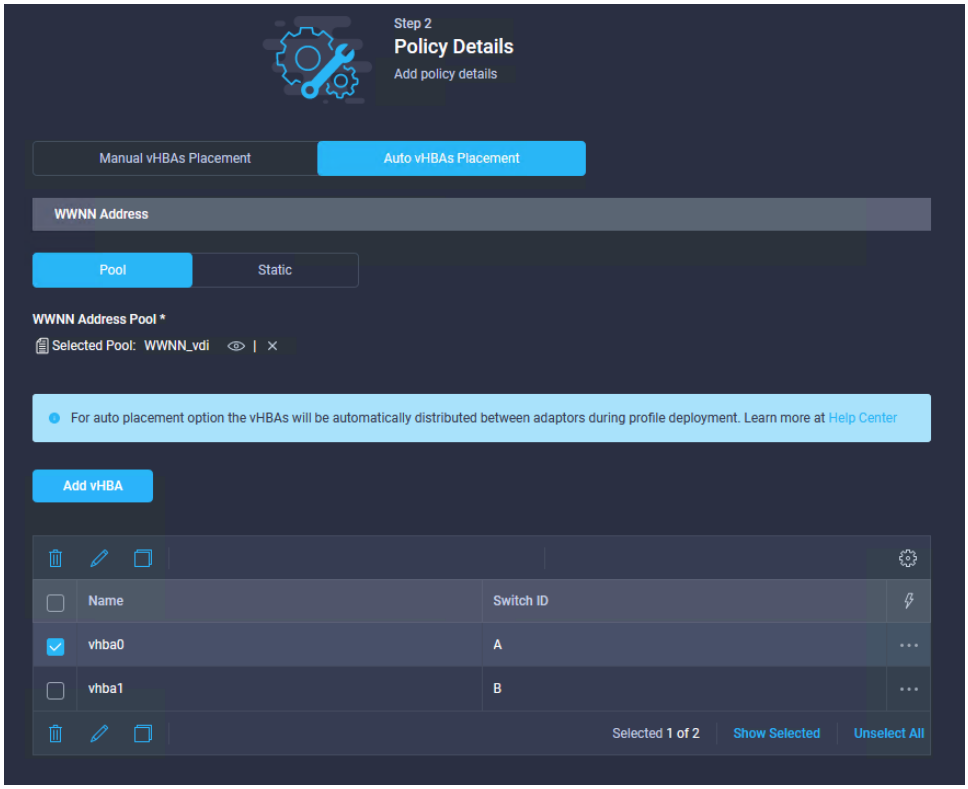

10. Create the World Wide Node Name (WWNN) pool.

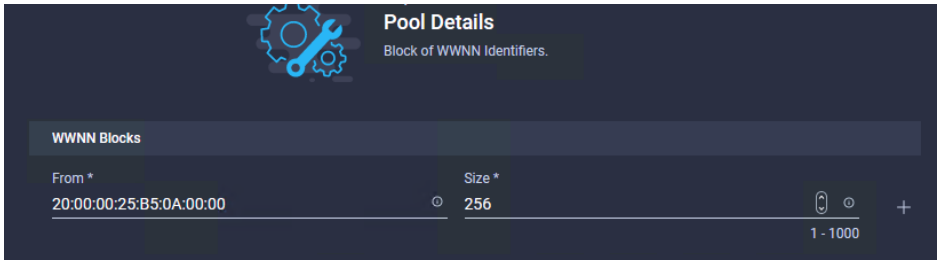

11. Create the World Wide Port Name (WWPN) pool. (Create a pool for each side of the fabric: for example, **wwpn_a** and **wwpn_b**.)

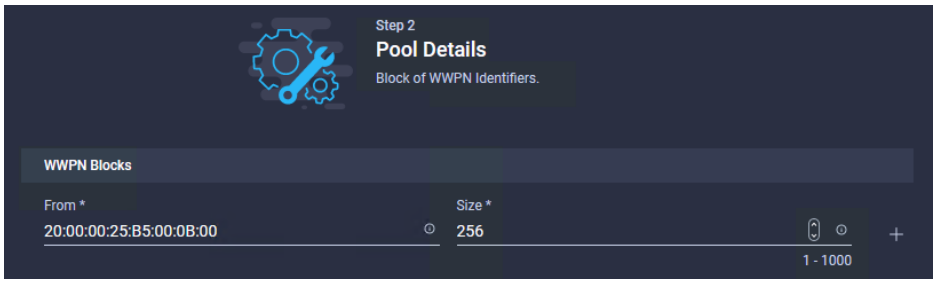

12. Click Add vHBA. (Perform this step twice: once for each side of the fabric.)
    - Select the WWPN pool created earlier.
    - Name the vHBA and define its type (fc-initiator is used in the testing described here).

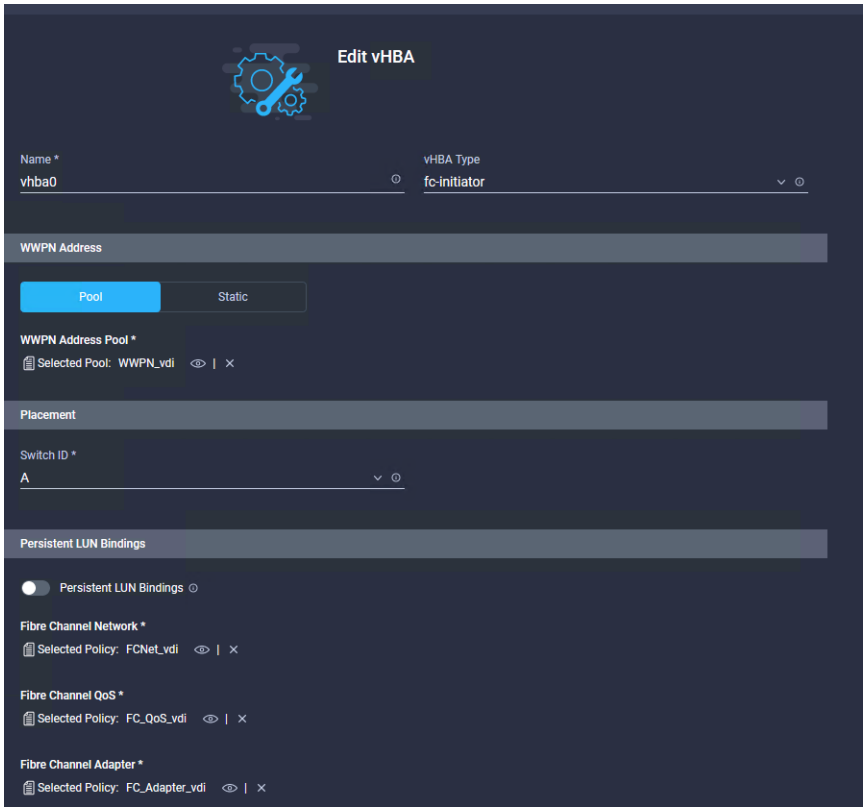

    - Select the switch ID.

- Create the Fibre Channel network policy. (Be sure that the VSAN IDs match the fabric side you are configuring.)

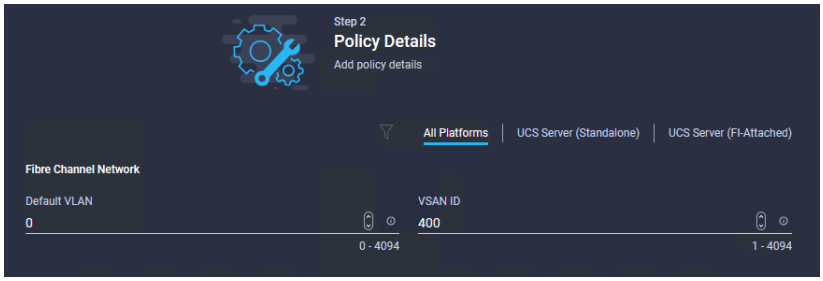

- Create the Fibre Channel QoS policy.

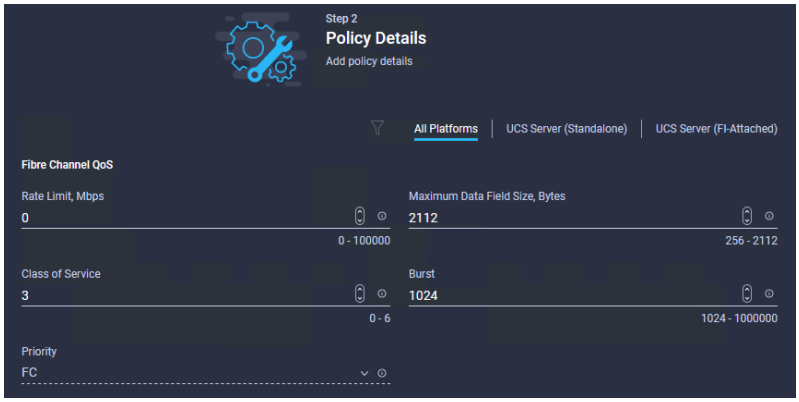

- Create the Fibre Channel adapter policy.

13. Create a virtual KVM policy.

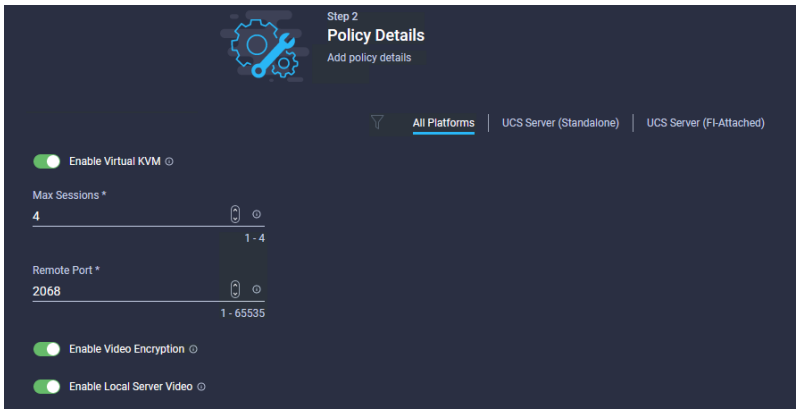

14. Create a virtual media policy if needed for the installation.

15. After you have created all these polices and pools, step through the Create Server Profile wizard and select each policy or pool as appropriate. After this process is complete, the service profile can be cloned, converted to a template, or directly assigned to a computing node.

Configure the VDI master target

You must first install the virtual machines for the master targets with the software components needed to build the golden images. Additionally, all available security patches for the Microsoft operating system and for Microsoft Office should be installed.

Four main steps are needed to prepare the master virtual machines: install the OS and VMware tools, install application software, install the Citrix Virtual Desktop Agent, and optimize the image with the Citrix Optimizer OS optimization tool.

Note: Citrix Optimizer 2.7, the optimization tool, includes customizable templates to enable or disable Windows system services and features using Citrix's recommendations and best practices across multiple systems. Because most Windows system services are enabled by default, you can use the optimization tool to easily disable unnecessary services and features to improve the performance of your virtual desktops.

Note: The images contain the basic features needed to run the Login VSI workload.

The master target virtual machine described in this document was configured as listed in Table 3.

Table 3. VDI virtual machines configuration

Configuration	VDI virtual machines
Operating system	Microsoft Windows 10 64-bit
Virtual CPU amount	2
Memory amount	3 GB
Network	VMXNET3 VDI
Virtual disk (vDisk) size	32 GB
Additional software used for testing	• Microsoft Office 2016 • Login VSI 4.1.40 (Knowledge Worker Workload)

Cisco UCS X210c Compute Node VDI testing

This section reports the testing used to evaluate the Cisco UCS X210c Compute Node and Cisco UCS X9508 Chassis solution.

Test strategy

To evaluate the Cisco UCS X210c Compute Node and Cisco UCS X9508 Chassis solution, we created a strategy to test for the optimal price-to-performance ratio for the mainstream user personas for VDI (also referred to as end-user computing) known as knowledge workers.

We used Windows 10 virtual desktops using nonpersistent Citrix Provisioning Services.

We evaluated the Intel Xeon processor 6348 using the Login VSI Knowledge Worker test workload shown in Table 4, in benchmark mode.

Table 4. User type and delivery mechanism combinations tested

Delivery mechanism	Task worker	Knowledge worker	Power user
Citrix Provisioning Services (VDI)	Not tested	Tested	Not tested

We started with our knowledge of the performance of the second-generation Intel Xeon Scalable Gold processors and compared their benchmark performance to that of the third-generation processors to guide our initial selections for evaluation. Our plan was to identify a processor for each user type that delivered the best price-to-performance ratio for both the VDI and Remote Desktop Session Host (RDSH) delivery modalities for Citrix Virtual Apps & Desktops.

Test methodology

For each user type and processor combination tested, we created a Citrix Virtual Apps & Desktops virtual machine with specifications as shown in Table 5.

Table 5. Tested user type and OS combinations

Combination	vCPU	Memory	vNIC
Knowledge worker: Windows 10	2 vCPUs	3 GB of memory	1 x 10-GB vNIC
Knowledge worker: Windows Server 2019	8 vCPUs	32 GB of memory	1 x 10-GB vNIC

We installed the chosen processor candidate in a Cisco UCS X210c Compute Node and ran Login VSI benchmark mode tests at calculated maximum user densities to determine the actual maximum user density per server. The maximum recommended user density is some number of users that complete the Login VSI workload with all attempted users active and logged off without triggering Login VSImax. In addition, CPU utilization on the host should not exceed 90 percent during the test.

We used the maximum recommended user density achieved to determine server loading in a server maintenance or failure scenario: typically, N – 1. We expect that customers would run their environment only at this load in those cases.

We compared performance and price per user at the maximum recommended user density to determine the best processor for the user type. We used Windows 10 and Windows Server 2019 across all processor testing.

Test data

This section presents the data from the test runs for the processor selected for the user type: in this case, knowledge workers

Knowledge workers are individuals in an organization who use a large number of applications to perform their duties. Examples of knowledge workers are sales and marketing professionals, business development managers, healthcare clinicians, and project managers.

In some cases, these workers can be served by RDSH server sessions or published applications. In most cases, organizations provide a medium-capability Windows 10 virtual desktop to these users.

We tested both use cases using the Login VSI Knowledge Worker workload in benchmark mode. You can find additional information about Login VSI and all the workloads we tested for this document here.

In addition to the Login VSI test suite, we measured host utilization by gathering data from ESXTOP. We also captured performance monitoring (perfmon) data from sample RDSH server virtual machines during the full server load tests.

Windows 10 and Citrix Virtual Apps & Desktops 1912 LTSR single-server synopsis: Intel Xeon Scalable Gold processor 6348

The test results are summarized here and in Figures 9, 10, and 11.

- Operating system: Windows 10 64-bit (2004) with Citrix optimizations
- 2 vCPUs; 3 GB of RAM
- Number of users: 280 users running Login VSI Knowledge Worker workload with Windows 10
- No VSImax; Login VSI baseline = 664 ms

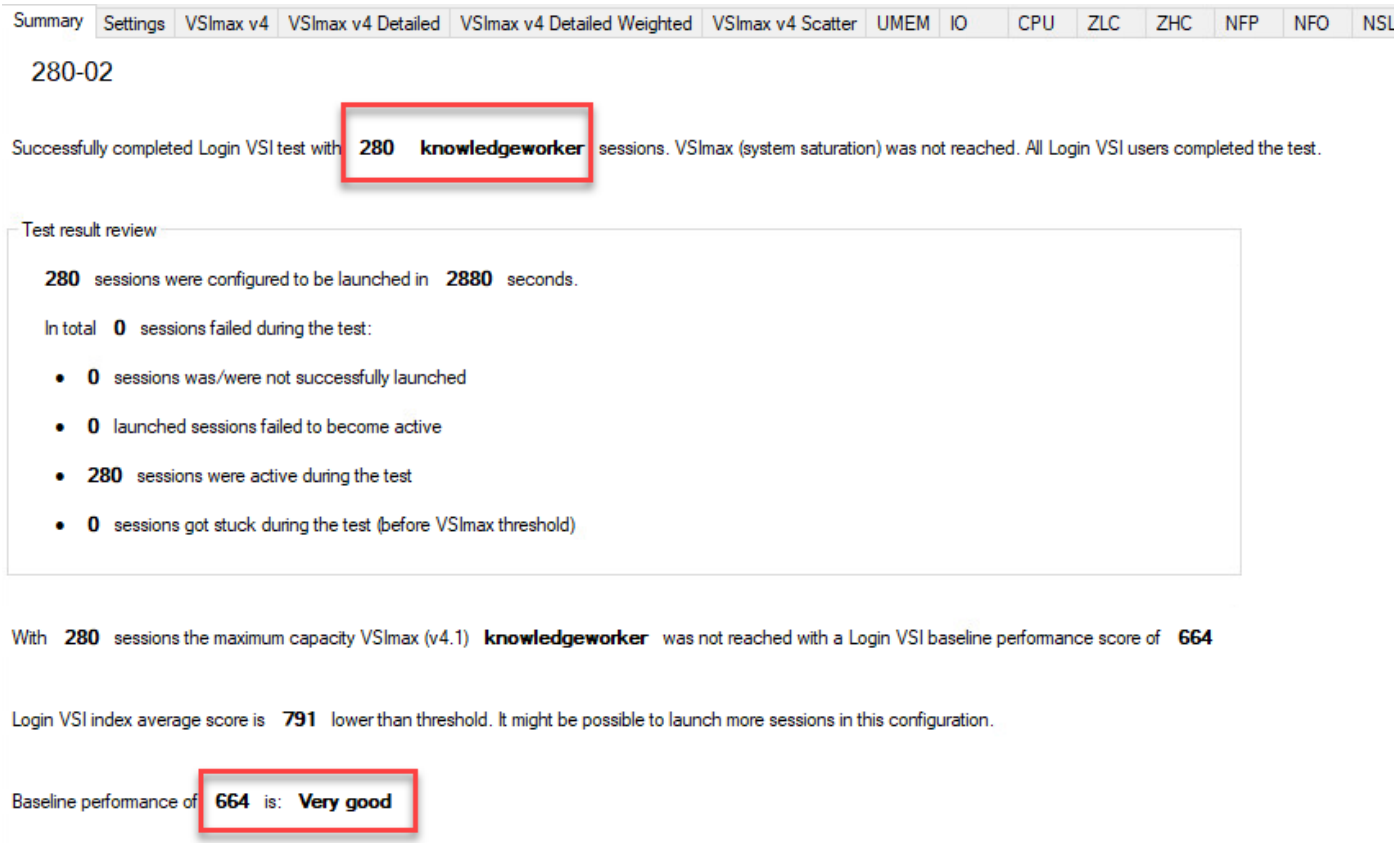

Figure 9.
Login VSI end-user experience summary

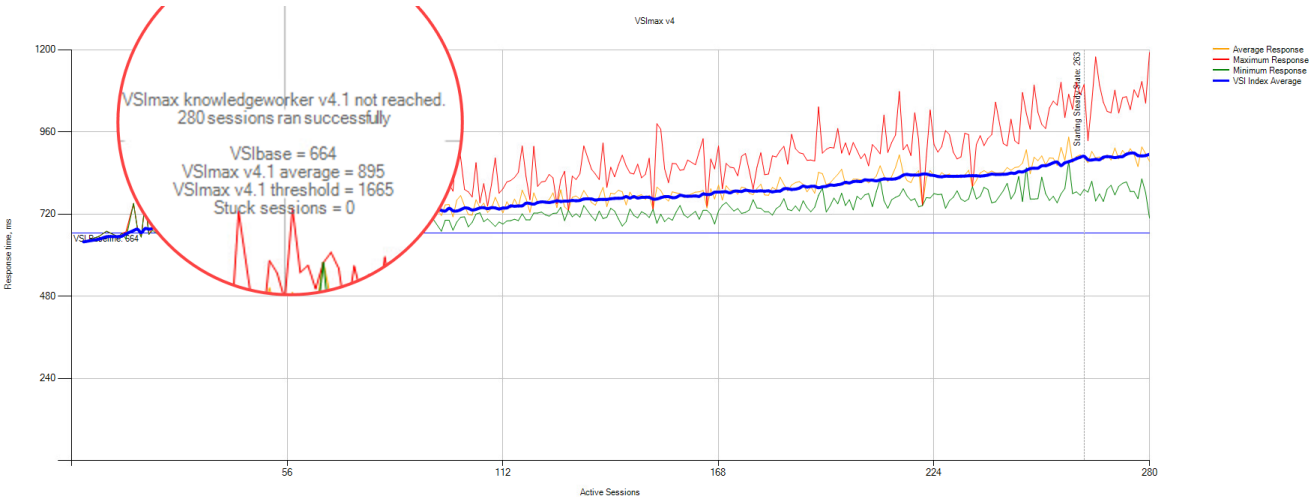

Figure 10.
Login VSI end-user experience performance

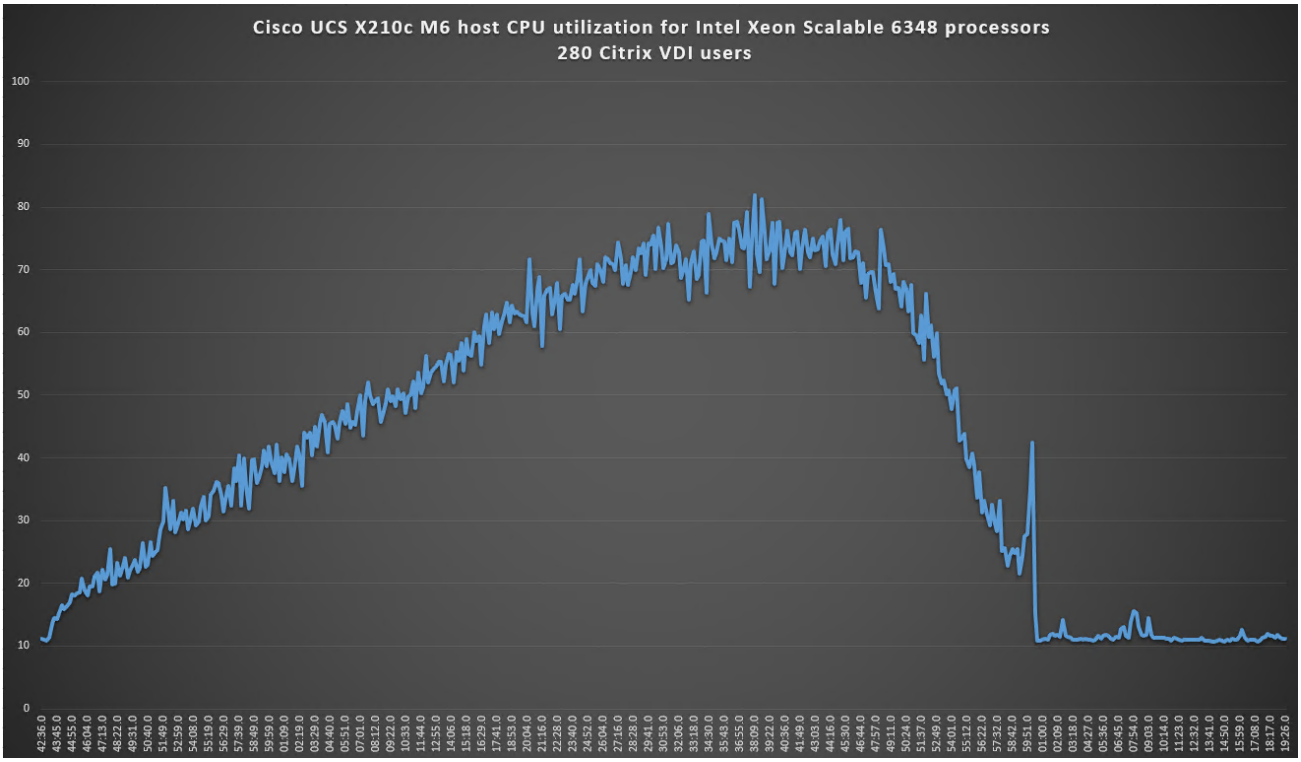

Figure 11.
VMware ESXi host CPU utilization percentage during testing

Windows Server 2019 and Cirix Virtual Apps & Desktops 1912 LTSR single-server synopsis: Intel Xeon Scalable Gold processor 6348

The test results are summarized here and in Figures 12, 13, and 14.

- Operating system: Windows Server 2019 64-bit (2004) with Citrix optimizations
- 8 vCPUs; 32 GB of RAM
- Number of users: 420 users running Login VSI Knowledge Worker workload with Windows Server 2019
- No VSImax; Login VSI baseline = 600 ms

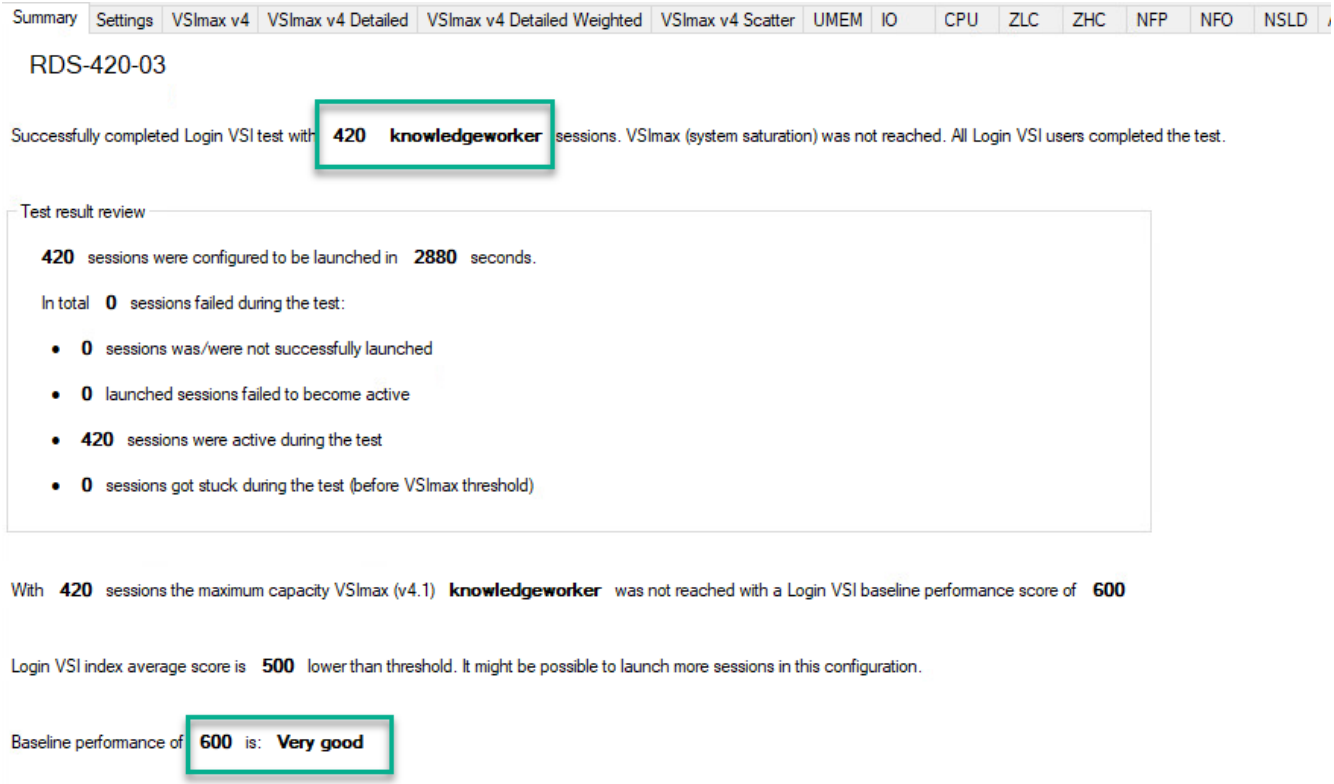

Figure 12.
Login VSI end-user experience summary

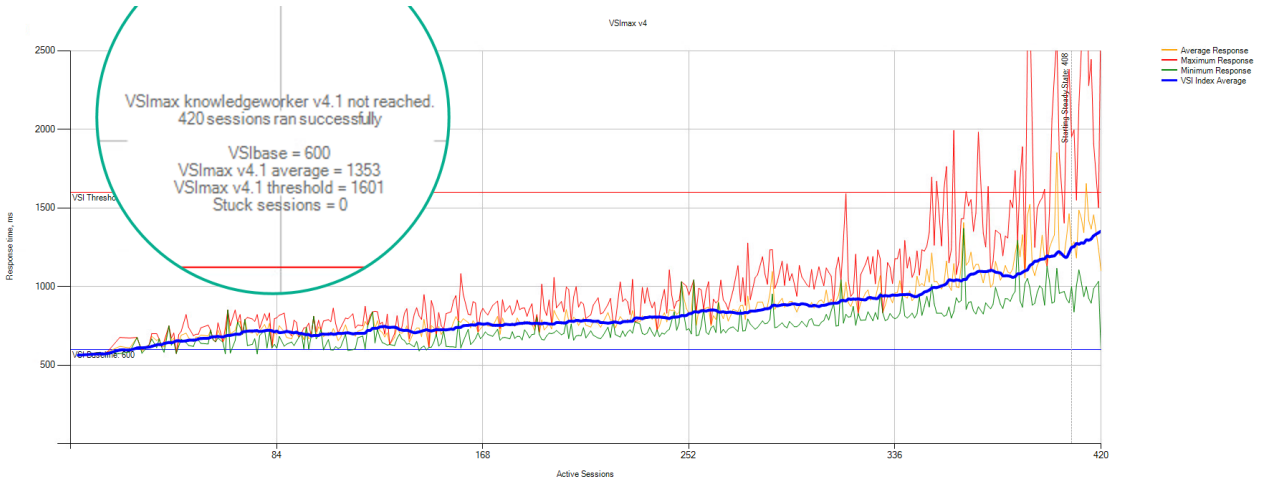

Figure 13.
Login VSI end-user experience performance

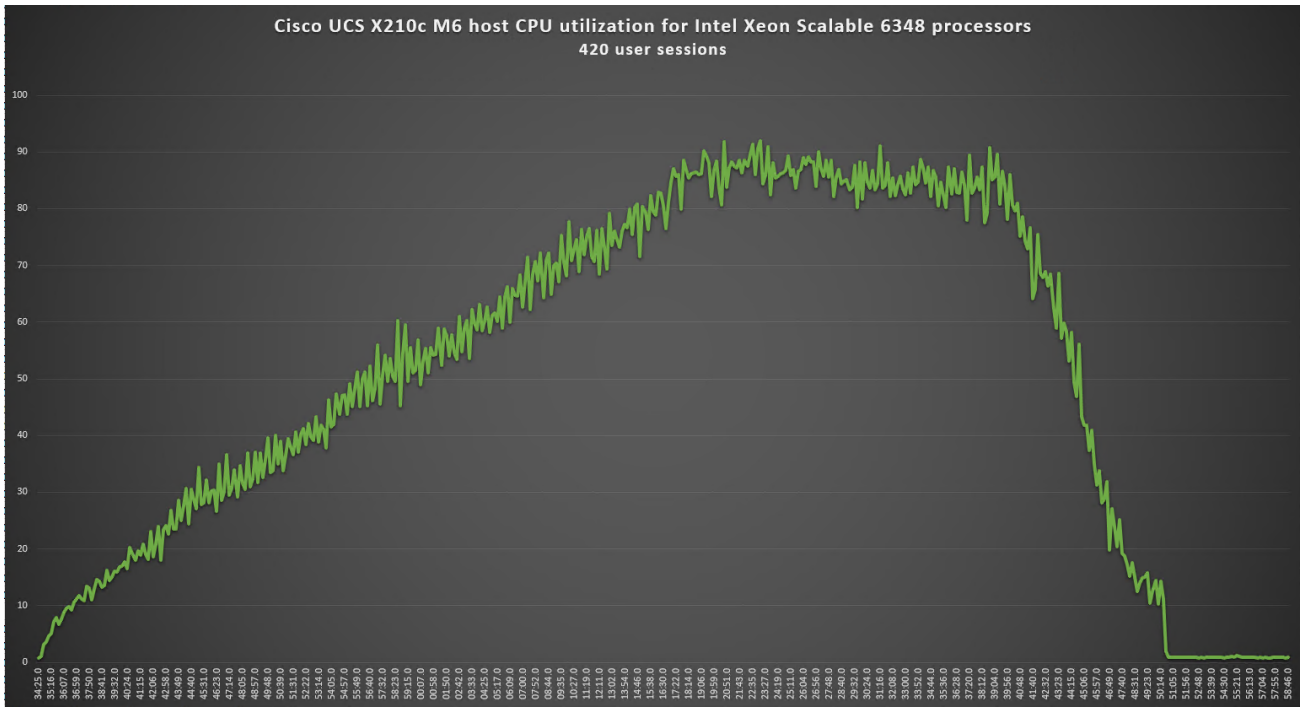

Figure 14.
VMware ESXi host CPU utilization percentage during testing

Conclusion

Cisco delivers a highly capable platform for enterprise end-user computing deployments with the Cisco UCS X210c M6 Compute Node with Intel Xeon CPUs.

Integrating the Cisco Intersight platform into your environment provides global visibility of infrastructure health and status along with a constantly growing list of advanced management and support capabilities.

In this study we saw a significant increase in user density and performance in our knowledge worker workloads when compared to our latest CVDs running on B200 M5 blade servers.

For more information

Consult the following references for additional information about the topics discussed in this document.

Products and solutions

- Cisco Intersight platform: https://www.intersight.com
- Cisco Unified Computing System: http://www.cisco.com/en/US/products/ps10265/index.html
- Cisco UCS 6454 Fabric Interconnect: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html
- Cisco UCS X9508 Chassis: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/datasheet-c78-2472574.html
- Cisco UCS X210c Compute Node: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/datasheet-c78-2465523.html
- Cisco UCS adapters: http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html
- Cisco Nexus 9000 Series Switches: http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html

Interoperability matrixes

- Cisco UCS Hardware Compatibility List: https://ucshcltool.cloudapps.cisco.com/public/

Cisco Validated Designs for VDI

- Deployment guide for FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_vmware_vs_7_design.html

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA 221102.7 07/21