Cisco Secure Enclave Datacenter Solution for EMC VSPEX

Last Updated: May 2, 2014

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Shivakumar Shastri, Technical Marketing Engineer, Server Access Virtualization Business Unit (SAVBU), Cisco Systems.Shivakumar has over 18 years of experience in multiple areas of IT infrastructure.

Acknowledgments

The author would like to acknowledge the following for their support and contribution to the design, validation and creation of this Cisco Validated Design (CVD):

- Chris O Brien—Cisco
- Mehul Bhatt—Cisco
- Vijay Durairaj—Cisco
- Matt Kaneko—Cisco
- Shankar Varanasy—Cisco
- Bathumalai Krishnan—Cisco
- Prashanto Kochavara—EMC

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.

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

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

Cisco Secure Enclave Datacenter Solution for EMC VSPEX

Introduction

Security is about safeguarding access to data. Appropriate security can be ensured only with complete visibility into and control of elements within the system. Visibility shows whether authorized users alone are accessing permitted data while granular and efficient controls provide necessary levers to prevent unauthorized access. A major consideration is the recent change in what defines a network, which goes beyond physical walls and other traditional boundaries to make up the extended network - Endpoints and mobile devices. These factors, an understanding of the changing nature of threat vectors and inevitability of security breaches has led to a design with a hard to corrupt core and adaptable modules. The solution, based on the VSPEX integrated system, can be augmented to address business, compliance and application requirements. This document presents a reliable introduction of Cisco® products and services along with partner compliments, including EMC and Lancope, in the data center to create a flexible, functional and secure application environment that can be readily automated.

Objective

The purpose of this document is to present an IT security framework that conforms to established design principles and to provide details about the solution stemming from this framework - Cisco Secure Enclaves Architecture (SEA). This document considers both the design and the composition of components to develop a coherent security solution that takes into account both hardware and software at every level of a Cisco and EMC integrated infrastructure stack (VSPEX). The goal of the design is to provide appropriate security that provides desirable levels of performance and fault tolerance with ease of management at a competitive price.

Audience

This document is intended to provide technical direction to channel partners and end-user customers interested in making security an integral part of their IT infrastructure. The need for security is even greater when IT resources are shared among groups of people whose data cannot be shared. This design and future implementations arising from it address the challenges and requirements of such a shared platform.

Problem Statement

Any shared platform, including Cloud, opens up access to key resources such as Infrastructure, Users and Applications. Security is all about ensuring confidentiality of data. Ensuring the consistent and correct delivery of data on a shared platform comes with increased risk and complexity. Further, most computing platforms are designed to meet performance and functional requirements with little or no attention to trustworthiness. The trend towards optimal use of IT resources through virtualization has resulted in an environment in which true and implied security accorded by physical separation has essentially vanished. System consolidation efforts have also accelerated the movement toward co-hosting on integrated platforms, and the likelihood of compromise is increased in a highly shared environment. This situation presents a need for enhanced security and an opportunity to create a framework and platform that instills trust. Lack of confidence that such a trust environment can be delivered with ease and maintained with resilient resource management is a major obstacle to the physical consolidation of applications and wider adoption of cloud-computing service models.

Many enterprises and IT service providers are developing cloud service offerings for public and private consumption. Regardless of whether the focus is on public or private cloud services, these efforts share several common objectives:

- Cost-effective use of capital IT resources through co-hosting
- Better service quality through virtualization features
- Increased operational efficiency and agility through automation

One essential characteristic of cloud architecture is the capability to pool resources and deliver these resources to tenants while conforming to declared SLA's. Power savings brought about by consolidation also contributes to reduced total cost of ownership (TCO). Achieving these goals can have a positive impact on profitability, productivity, and product quality.

Enabling enterprises to migrate such environments to cloud architecture requires the capability to provide customer confidentiality while delivering the management and flexibility benefits of shared resources. Both private and public cloud providers must secure all customer data, communication, and application environments from unauthorized access. Such separation, with regulatory compliance measures, must be complete and consistent to instill confidence and achieve widespread adoption.

Business Benefits

The Cisco Secure Enclaves architecture helps evolve the current converged infrastructure offering of Cisco (VSPEX) by simplifying and standardizing the delivery of Cisco application and security services on architecturally consistent platforms. This approach is a logical extension of these data center building blocks, advancing the benefits of standardization beyond the infrastructure to the applications and services required. This design provides the following features that facilitate a uniform approach to IT in the data center:

- Flexible consumption model, allowing customer requirements to be met from both application and business perspectives
- Automation of well-known and well-understood resource pools of networking, computing, and storage resources with security constructs.
- Onboarding of services and applications.
- Platform hardening and automation of security operations such as:
  - Configuration
  - Auditing
  - Patching
  - Responses
- Operation compliance and certifications

Infrastructure as a Service (IaaS), from the provider perspective, consists of a set of modular building blocks of underlying resources assembled systematically based on services requested and overlaid with security. Services may be introduced either through dedicated appliances or through virtual appliance implementations on shared general-purpose computing resources. The main design objective is to help ensure that applications in this environment meet their subscribed to service-level agreements (SLAs), including confidentiality requirements, by using pretested and validated IT infrastructure components to prevent inefficiency and inaccuracy.

Note Performance benchmarking is out of scope for this validation effort.

Security Philosophy—The Reference Monitor

The three most basic and necessary characteristics of the components that enforce security and instill trust are as follows:

- The mechanism must be protected from modification by unauthorized methods and users.
- The mechanism must not be allowed to be bypassed.
- The mechanism must be simple to understand and monitor.

These core requirements together help ensure the trustworthiness of the enforcement module - the monitor.

Design Principles

Design principles are rules and guidelines instituted to help ensure, inform, and support the way in which an architectural implementation fulfills its mission. They provide a means to tie components and methods to the business objectives: protection, performance, and provisioning.

In a security platform, trust is paramount and must not be misplaced. In a system consisting of components of varying levels of trustworthiness, the assumption is that the overall trustworthiness of the system is not better than the least trustworthy subcomponent. Security is enforced through access control, which requires complete visibility into that which is being secured. Relevant principles, the rationale for inclusion in an enclave, and the scope of an enclave are summarized here.

- **Least-common mechanism:** This principle states the need to globalize common and shared modules (in the enforcement domain). It has the effect of reducing duplicates, which can result in fewer opportunities for compromise. It also has the advantage of less overhead due to fewer number of instances with potential for better performance when implemented in hardware. Another positive effect of this principle is ease of maintenance.
- **Reduced sharing:** In the user domain, no computer resource should be shared between components or subjects unless it is necessary to do so. This approach helps prevent both inadvertent and deliberate encroachment. When information needs to be shared, it should be done only if sharing has been explicitly requested and granted.
- **Efficient mediated access:** As with most IT systems, development of secure systems includes interaction between hardware and software mechanisms. In a hierarchically constructed system with hardware constituting the lowest layer, when possible the most efficient choice is to allocate an access mediation mechanism to the hardware. Although hardware implementations provide greater performance, software equivalents provide flexibility, which is crucial in devising an adaptable solution. The principle of efficient mediated access strikes a balance between two possibilities by stipulating access control functions be allocated to the lowest possible level (closer to hardware) that still meets flexibility requirements.

These design principles, while appearing to be contradictory, are complementary when their respective scopes or domains are clearly defined. The first two principles are relevant in different spaces: the enforcement and user domains. Such principles are brought together to achieve a protected platform that can perform as desired and be provisioned with ease and correctness when required. Adopting global and dedicated appliances such as Cisco ASA firewalls and Cisco NGA NetFlow devices enable desired levels of performance for the most critical elements (least-common mechanism) while also conforming to the reference monitor tenet of preserving the fidelity of the enforcement module. Management of Cisco ASA with Cisco Security Manager and authentication and authorization services provided by Cisco ISE software demonstrate the flexibility brought about by a centralized and global policy and configuration engine. User-domain abstractions are encapsulated in fenced containers (enclaves) automated through Cisco UCS Director, providing efficient mediated access.

After implementation, the platform needs to be able to continuously enforce data protection at every stage of the life of information, through encryption, to help ensure integrity. Techniques used to deploy the model components should be repeatable for easy and correct construction. The orchestration capabilities of Cisco UCS Director are used for this purpose. Eventually, every engineered system is a work in progress and so must take into account planned upgrades and maintenance. A system composed of simple and essential components contributes to reduced complexity and better understanding. Other criteria that facilitate easy adoption include less intrusive and more intuitive interfaces with clear user expectations of security risk.

Architectural Overview

The Cisco Secure Enclaves design uses the common components of VSPEX Integrated Systems along with additional services integration to address business and application requirements. These functional requirements promote uniqueness and innovation in the integrated computing stack that augment the original design to support these prerequisites. The result is a region, or enclave, and more likely multiple enclaves, in the integrated infrastructure designed and built to address the unique workload activities and business objectives of an organization.

The VSPEX Integrated System combines Cisco Unified Computing System™ (Cisco UCS®) and Cisco Nexus® switching platforms with storage from EMC's next generation VNX 5400 array. The result is a standardized infrastructure and the foundation to rapidly deliver data center applications, virtualized desktops, and cloud computing services. Figure 1 illustrates the Cisco structural elements currently used in VSPEX The enclave infrastructure foundation is formed using a subset of these components.

Figure 1 Cisco Integrated Components

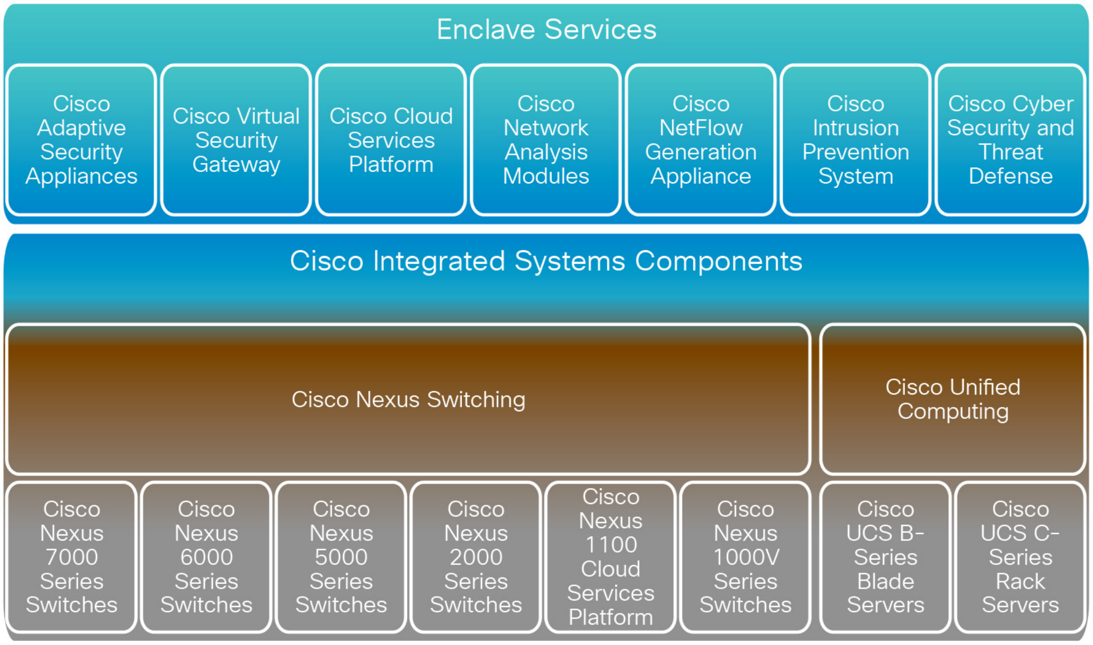
Enclave Services
Cisco Adaptive Security Appliances
Cisco Virtual Security Gateway
Cisco Cloud Services Platform
Cisco Network Analysis Modules
Cisco NetFlow Generation Appliance
Cisco Intrusion Prevention System
Cisco Cyber Security and Threat Defense
Cisco Integrated Systems Components
Cisco Nexus Switching
Cisco Unified Computing
Cisco Nexus 7000 Series Switches
Cisco Nexus 6000 Series Switches
Cisco Nexus 5000 Series Switches
Cisco Nexus 2000 Series Switches
Cisco Nexus 1100 Cloud Services Platform
Cisco Nexus 1000V Series Switches
Cisco UCS B-Series Blade Servers
Cisco UCS C-Series Rack Servers

For more information about Cisco Integrated Systems, go to:

http://www.cisco.com/c/en/us/solutions/data-center-virtualization/integrated-systems/index.html

The enclave strategy is a logical extension of the foundational platforms found in Cisco's converged infrastructure stacks. The enclave maintains the traditional design pillars associated with the shared computing stack architectures, which provide service assurance and enterprise-class availability in the data center. This foundation is readily extended to include organic and supplementary security services enabled or attached to this base as the application workloads or business initiatives require. Figure 2 shows a generic physical layout of Cisco Integrated Systems components that constitute the foundation of the enclave model.

Figure 2 Cisco Converged Infrastructure Physical Components

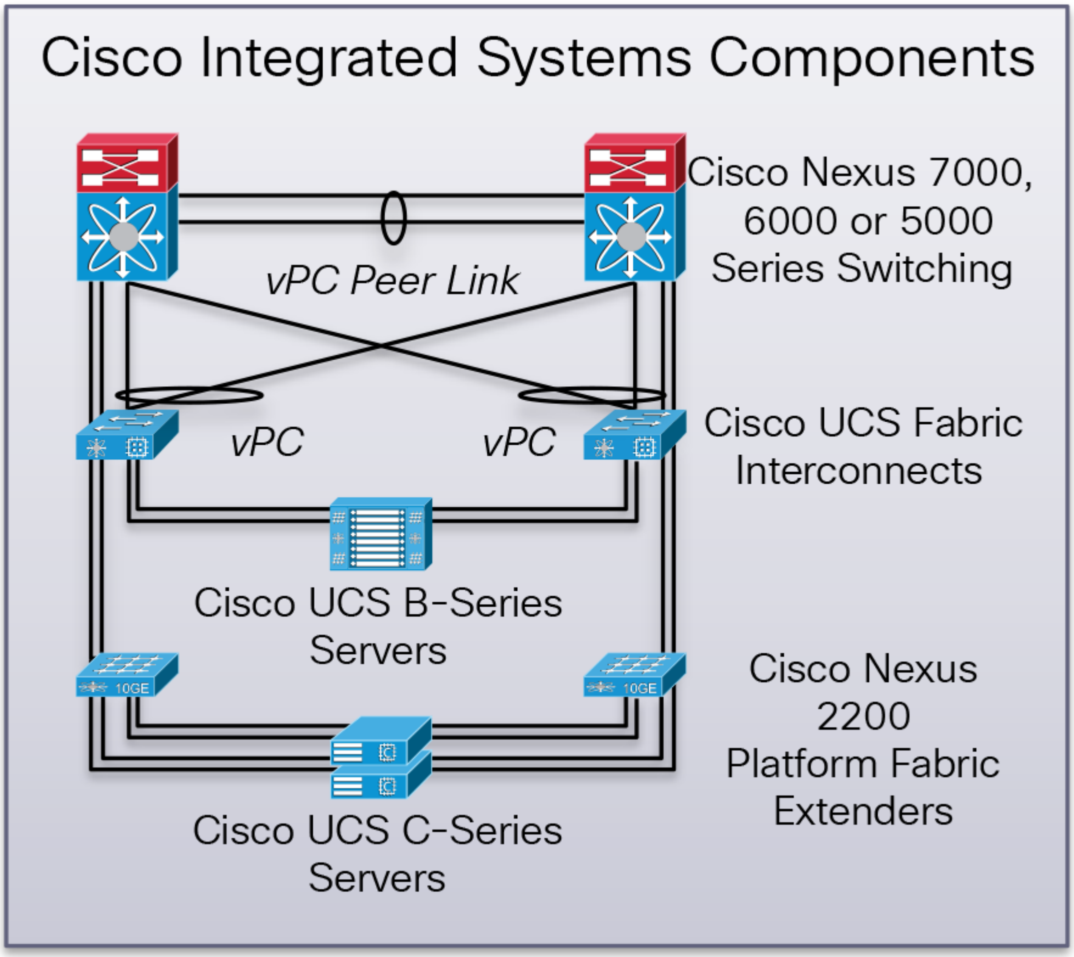
Cisco Integrated Systems Components
Cisco Nexus 7000, 6000 or 5000 Series Switching
vPC Peer Link
vPC
vPC
Cisco UCS Fabric Interconnects
Cisco UCS B-Series Servers
10GE
10GE
Cisco Nexus 2200 Platform Fabric Extenders
Cisco UCS C-Series Servers

Figure 3 shows the extension of Cisco Integrated Systems to include features and functions beyond the foundational elements. Access controls, visibility, and threat defense are all elements that can be uniformly introduced into the system as required. The main feature of the enclave is the extensibility of the architecture to integrate current and future technologies within and upon its underpinnings, expanding the value of the infrastructure stack to address current and future application requirements.

Figure 3 Cisco Secure Enclaves Architecture Structure

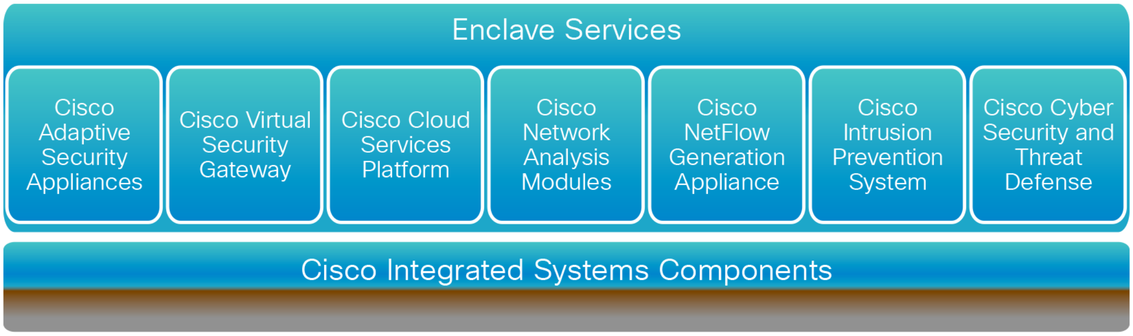
Enclave Services
Cisco Adaptive Security Appliances
Cisco Virtual Security Gateway
Cisco Cloud Services Platform
Cisco Network Analysis Modules
Cisco NetFlow Generation Appliance
Cisco Intrusion Prevention System
Cisco Cyber Security and Threat Defense
Cisco Integrated Systems Components

The augmentation of the converged infrastructure stacks can be both physical and virtual. Figure 4 and Figure 5 illustrate the addition of physical Cisco Adaptive Security Appliances (ASA) Next-Generation Firewall Services and NetFlow

The Cisco ASA 5585 provides advanced stateful firewall and VPN concentrator functionality in one device, and for some models, integrated services modules such as IPS. The ASA includes many advanced features, such as multiple security contexts (similar to virtualized firewalls), clustering (combining multiple firewalls into a single firewall), transparent (Layer 2) firewall or routed (Layer 3) firewall operation, advanced inspection engines, VPN support, Cisco TrustSec and many more features.The enclave design uses the Security Group Firewall (SGFW) functionality of the ASA to enforce policy to and between servers in the data center. The SGFW objects are centrally defined in the Cisco Identity Services Engine (ISE) and used by the security operations team to create access policies. The Cisco ASA simply has the option to use the source and destination security groups to make decisions.

Figure 4 shows a more traditional Cisco ASA high-availability pair deployment model in which the Cisco Nexus switches provide a connection point for the appliances. Cisco offers a number of Cisco ASA models to address the organization's specific scale requirements.

Figure 4 Physical Extension of Cisco Integrated Systems with Cisco ASA Security Services

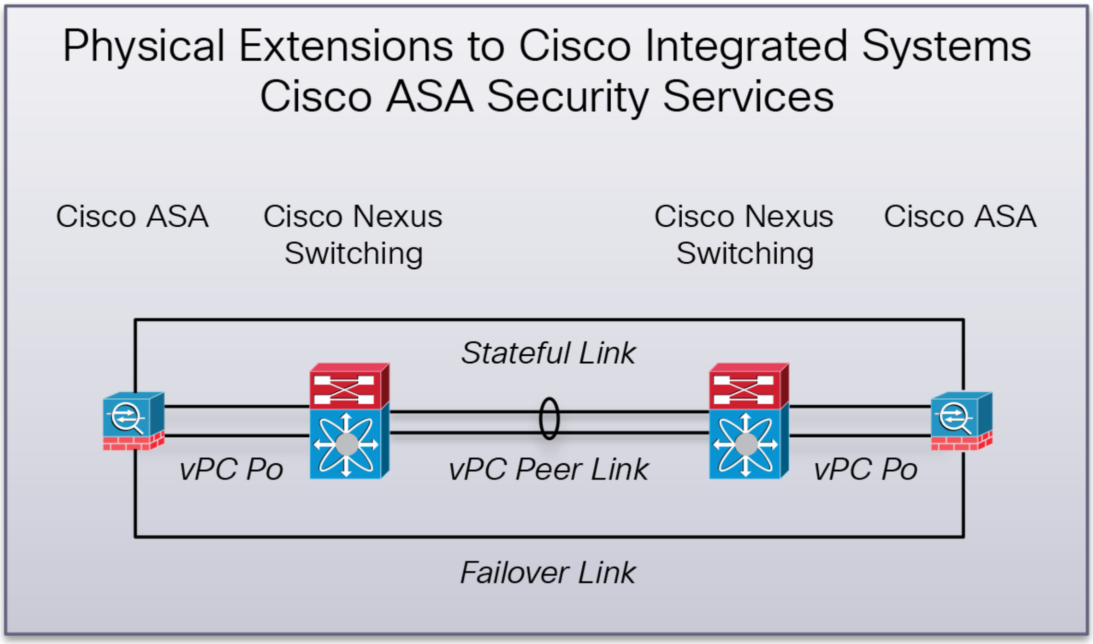

In addition to the Cisco ASA platforms, the integrated stack readily supports other services. For example, the Cisco NetFlow Generation Appliance (NGA 3240) introduces a highly scalable, cost-effective architecture for cross-device flow generation. The Cisco NGA generates, unifies, and exports flow data, empowering network operations, engineering, and security teams to boost network operations excellence, enhance services delivery, implement accurate billing, and harden network security.The NGA is a promiscuous device and can accept mirrored traffic from any source to create NetFlow records for export. The export target in this design is the cyber threat detection system - Lancope StealthWatch.

The use of threat defense systems allows an organization to address compliance and other mandates, network and data security concerns as well as monitoring and visibility of the data center. Cyber threat defense addresses several use cases including:

- Detecting advanced security threats that have breached the perimeter security boundaries
- Uncovering Network & Security Reconnaissance
- Malware and BotNet activity
- Data Loss Prevention

Figure 6 shows the deployment of Cisco NGA on the stack to provide these services, accepting mirrored traffic from various sources of the converged infrastructure as Cisco NetFlow source data.

Figure 5 Physical Extensions of Cisco Integrated Systems with NetFlow Offload Services

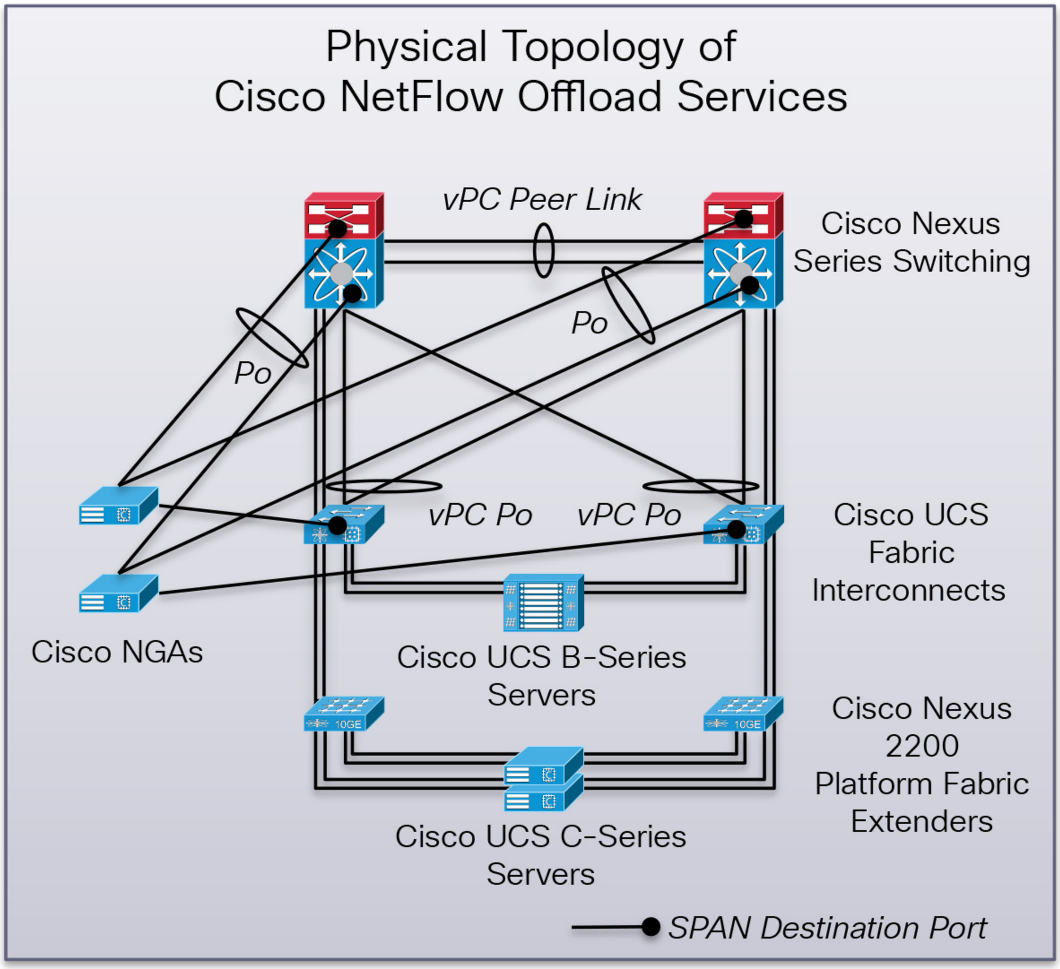
Physical Topology of
Cisco NetFlow Offload Services
vPC Peer Link
Cisco Nexus Series Switching
Po
Po
vPC Po
vPC Po
Cisco UCS Fabric Interconnects
Cisco NGAs
Cisco UCS B-Series Servers
10GE
10GE
Cisco Nexus 2200 Platform Fabric Extenders
Cisco UCS C-Series Servers
SPAN Destination Port

The strategic value of the enclave framework is in the capability of the structure to adapt to an organization's needs. Supporting other physical appliance-based services beyond the Cisco ASA and NGA platforms is certainly feasible. "The Enclave" section of this document describes the blending of physical and virtual services to construct one or multiple unique regions.

The management of the enclave can be performed by individual domain managers or unified through Cisco UCS Director (Figure 7). Cisco UCS Director offers converged infrastructure management and extensions to control additions to the stack. Either model allows an organization to maintain traditional policy roles associated with computing, networking, and storage resources and security groups. The enclave framework currently uses the following domain management platforms:

- Cisco UCS Manager

- Cisco Prime™ Network Services Controller (NSC)
- Cisco Identity Services Engine (ISE)
- Cisco Security Manager
- Lancope StealthWatch Management Console

Figure 6 illustrates the extension of Cisco Integrated Systems to include features and functions beyond the foundational elements. Access controls, visibility, and threat defense are all elements that can be uniformly introduced into the system as required. The main feature of the enclave framework is the extensibility of the architecture to integrate current and future technologies within and upon its underpinnings, expanding the value of the infrastructure stack to address current and future application requirements.

Figure 6 Cisco Secure Enclave Architecture Management Structure

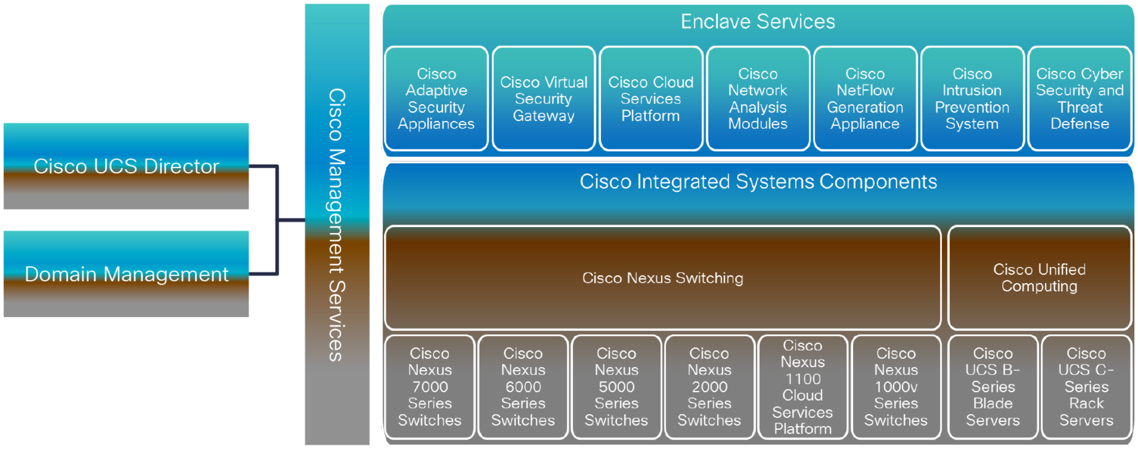
Cisco UCS Director
Domain Management
Cisco Management Services
Enclave Services
Cisco Adaptive Security Appliances
Cisco Virtual Security Gateway
Cisco Cloud Services Platform
Cisco Network Analysis Modules
Cisco NetFlow Generation Appliance
Cisco Intrusion Prevention System
Cisco Cyber Security and Threat Defense
Cisco Integrated Systems Components
Cisco Nexus Switching
Cisco Unified Computing
Cisco Nexus 7000 Series Switches
Cisco Nexus 6000 Series Switches
Cisco Nexus 5000 Series Switches
Cisco Nexus 2000 Series Switches
Cisco Nexus 1100 Cloud Services Platform
Cisco Nexus 1000v Series Switches
Cisco UCS B-Series Blade Servers
Cisco UCS C-Series Rack Servers

Figure 7 VSPEX Building Blocks

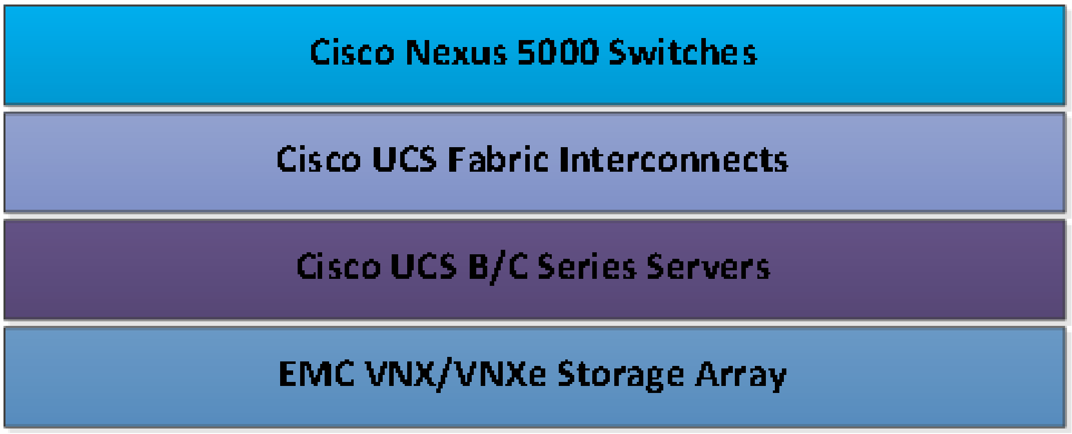
Cisco Nexus 5000 Switches
Cisco UCS Fabric Interconnects
Cisco UCS B/C Series Servers
EMC VNX/VNXe Storage Array

Cisco VSPEX Secure Enclave Data Center uses the VSPEX converged integrated platform as its foundation. The VSPEX infrastructure solution from Cisco and EMC leverage previously integrated components for validation to expedite IT infrastructure and application deployment while simultaneously reducing cost, complexity and risk. The VSPEX platform contains Cisco Nexus Networking, Cisco Unified Computing System™ (Cisco UCS®) and EMC next generation VNX or VNXe series storage systems. One especially significant benefit of the VSPEX reference architecture is the flexibility in sizing and performance it accords while ensuring compatibility to suit most customer requirements. Flexibilities come from both hardware choices as well as operating systems or hypervisors that are supported.

The Cisco VSPEX Secure Enclave Datacenter design augments inherent capabilities of the VSPEX platform with critical security services to address specific business and application requirements of a multi-tenant enterprise. The result is a region or enclave and more likely multiple enclaves, built on VSPEX to address the unique workload activities and business objectives of an organization.

Cisco VSPEX Secure Enclave Data Center is developed using the following technologies:

VSPEX infrastructure from Cisco and EMC.

- VMware vSphere
- Cisco Adaptive Security Appliance (ASA)
- Cisco NetFlow Generation Appliance (NGA)
- Cisco Virtual Security Gateway (VSG)
- Cisco Identity Services Engine (ISE)
- Cisco Network Analysis Module
- Cisco UCS Director
- Lancope StealthWatch System

The Enclave

The enclave is a distinct logical entity that encompasses essential constructs including security along with application or customer-specific resources to deliver a trusted platform that meets SLAs. The modular construction and automated delivery help make the enclave a scalable and securely separated layer of abstraction that conforms with the design philosophy. The use of multiple enclaves delivers increased isolation, addressing disparate requirements of the converged infrastructure stack.

Figure 8 provides a conceptual view of the enclave that defines an enclave in relation to an n-tier application.

The enclave provides the following functions:

- Access control point for the secure region (public)
- Access control within and between application tiers (private)
- Cisco Cyber Security and Threat Defense operations to expose and identify malicious traffic
- Cisco TrustSec® security using secure group access control to identify server roles and enforce security policy
- Out-of-band management for centralized administration of the enclave and its resources
- Optional load-balancing capabilities

Figure 8 Cisco Secure Enclave Model

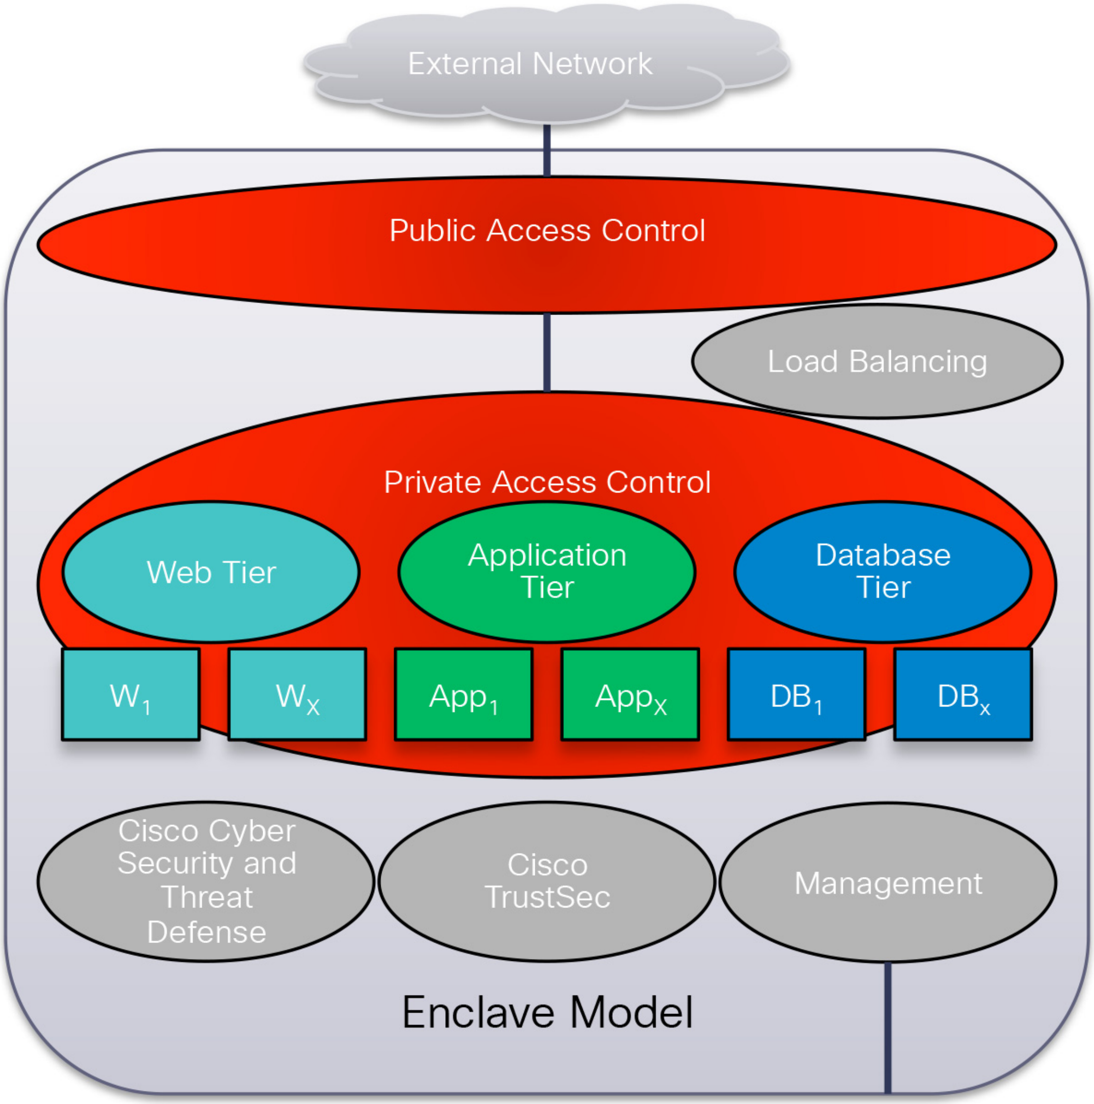

The components that form the enclave may vary in form-factor and be physical or virtual, and the requirements for functions may be based on business or application needs, but the structure is consistent in its form and manageability. The next sections discuss the enclave model to provide a better understanding of the system, its components, and their roles. The topics discussed include:

- Host topology
- Enclave topology
- Traffic patterns

Note The Cisco Secure Enclaves architecture is hypervisor independent. The details provided in this document address a VMware vSphere deployment. Future efforts could address other virtualization platforms and bare-metal instances.

Host Topology

The Cisco UCS Manager resides on a pair of Cisco UCS 6200 Series Fabric Interconnects using a clustered, active-standby configuration for high availability. The software gives administrators a single interface for compute device discovery, provisioning, inventory, configuration, event and performance monitoring and auditing. Cisco UCS Manager service profiles and templates support versatile role and policy-based management, and system configuration information can be exported to configuration management databases (CMDBs) to facilitate processes based on IT Infrastructure Library (ITIL) concepts.

Compute nodes are deployed in a Cisco UCS environment by applying Cisco UCS service profiles. Service profiles let server, network, and storage administrators treat Cisco UCS servers as raw computing capacity to be allocated and reallocated as needed. The profiles define server I/O properties, personalities and firmware revisions and are stored in the Cisco UCS 6200 Series Fabric Interconnects. Using service profiles, administrators can provision infrastructure resources in minutes instead of days, creating a more dynamic environment and more efficient use of server capacity.

Each service profile consists of a server software definition and the server's LAN and SAN connectivity requirements. When a service profile is deployed to a server, Cisco UCS Manager automatically configures the server, adapters, fabric extenders for rack servers, and fabric interconnects to match the configuration specified in the profile. The automatic configuration of servers, network interface cards (NICs), host bus adapters (HBAs), LAN and SAN switches, lowers the risk of human error, improves consistency, and decreases server deployment times.

Service profiles benefit both virtualized and non-virtualized environments in the Cisco Secure Enclave deployment. The profiles increase the mobility of non-virtualized servers, such as when moving workloads from server to server or taking a server offline for service or upgrade. Profiles can also be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing virtual machine mobility. Figure 9 shows the uniform deployment of VMware ESXi within the enclave framework.

The main features include:

- The VMware ESXi host resides in a Cisco converged infrastructure.
- The VMware ESXi host is part of a larger VMware vSphere High Availability (HA) and Distributed Resource Scheduler (DRS) cluster
- Cisco virtual interface cards (VICs) offer multiple virtual PCI Express (PCIe) adapters for the VMware ESXi host for further traffic isolation and specialization.
- Six Ethernet-based virtual network interface cards (vNICs) with specific roles associated with the enclave system, enclave data, and core services traffic are created:
  - vmnic0 and vmnic1 for the Cisco Nexus 1000V system uplink support management, VMware vMotion, and virtual service control traffic.
  - vmnic2 and vmnic3 support data traffic originating from the enclaves.
  - vmnic4 and vmnic5 carry core services traffic.
- Private VLANs isolate traffic to the virtual machines within an enclave, providing core services such as Domain Name System (DNS), Microsoft Active Directory, Domain Host Configuration Protocol (DHCP), and Microsoft Windows updates.
- Two virtual host bus adapters (vHBAs) for multihoming to available block-based storage.
- Four VMkernal ports are created to support the following traffic types:
  - vmknic0 supports VMware ESXi host management traffic.
  - vmknic1 supports VMware vMotion traffic.

- Two VMknics (vmknic2 and vmknic3) provide the Virtual Extensible LAN (VXLAN) tunnel endpoint (VTEP) to support traffic with path load balancing through the Cisco UCS fabric.
- Additional Network File System (NFS) and Small Computer System Interface over IP (iSCSI) VMknics can be assigned to individual enclaves to support application and segmentation requirements. These VMknics use the PortChannel dedicated to enclave data.

Note A maximum of 256 VMkernal NICs are available per VMware ESXi host.

- Cisco Nexus 1000V is deployed on the VMware ESXi host with the following elements:
  - PortChannels created for high availability and load balancing
  - Segmentation of traffic through dedicated vNICs, VLANs, and VXLANs

Figure 9 *VMware ESXi Uniform Host Topology*

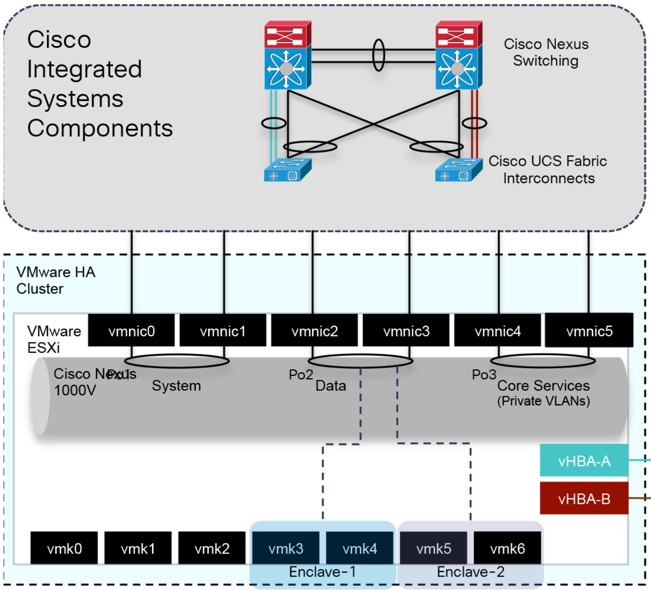

- vmknic0 is the ESXi management interface and will use the System Port Channel 01
- vmknic1 is the vMotion vmknic and will use the System Port Channel 01
- vmknic2 supports VXLAN functionality and will use the Data Port Channel 02
- The Core Services Port Channel 03 provides traffic isolation for virtual machines within each enclave (not shown)
- Fabric A and B VSANs will be shared across the production cluster
- All ESXi hosts boot off of infrastructure dedicated LUNs
- Enclaves may require one or more additional VMkernal NICs to provide isolated storage access
- Maximum of 256 Vmkernal NICs per host

The enclave architecture is not restricted to virtualized server platforms. Bare-metal servers persist in many organizations to address various performance and compliance requirements. To address bare-metal operating systems within an enclave (Figure 10), the following features were enabled:

- Cisco UCS fabric failover provides fabric-based high availability.
  - This feature precludes the use of host-based link aggregation or bonding.
  - Cisco VICs to provide multiple virtual PCIe adapters to the host for further traffic isolation and specialization.
- Ethernet-based vNICs with specific roles associated with the enclave system, enclave data, and core services traffic are created:
  - vnic-a and vnic-b support data traffic originating from the host. Two vNICs were defined to allow host-based bonding. One vNIC is required.
  - vcore supports core services traffic.
- Private VLANs isolate traffic to the virtual machines within an enclave, providing core services such as DNS, Microsoft Active Directory, DHCP, and Microsoft Windows Updates.

- Two virtual HBAs provide multihoming to available block-based storage.
- Dedicated VLANs per enclave for bare-metal server connections.

Figure 10 *Bare-Metal Host Topology*

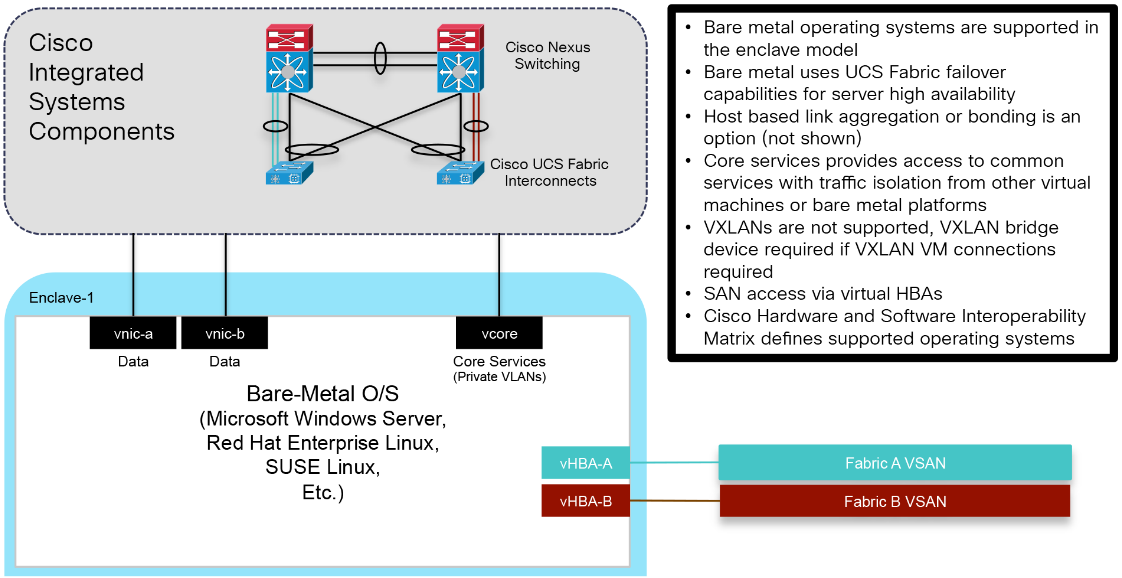
Cisco Integrated Systems Components
Cisco Nexus Switching
Cisco UCS Fabric Interconnects
Enclave-1
vnic-a
vnic-b
vcore
Data
Data
Core Services (Private VLANs)
Bare-Metal O/S (Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux, Etc.)
vHBA-A
vHBA-B
Fabric A VSAN
Fabric B VSAN

- Bare metal operating systems are supported in the enclave model
- Bare metal uses UCS Fabric failover capabilities for server high availability
- Host based link aggregation or bonding is an option (not shown)
- Core services provides access to common services with traffic isolation from other virtual machines or bare metal platforms
- VXLANs are not supported, VXLAN bridge device required if VXLAN VM connections required
- SAN access via virtual HBAs
- Cisco Hardware and Software Interoperability Matrix defines supported operating systems

Enclave Topology

The network fabric and services, integral to the VSPEX solution, combines the previously defined storage and compute devices into a cohesive system. This combination, along with other modular pieces outlined, create an efficient, consistent and secure application platform—the Enclave. The instantiation of this design can be further standardized by using automation tools such as Cisco UCS Director as a delivery mechanism. This section describes an enclave model and their components and capabilities.

Figure 11 depicts an enclave using two VLANs, with one or more VXLANs used at the virtualization layer. The VXLAN solution provides logical isolation within the hypervisor and removes the scale limitations associated with VLANs. The enclave is constructed as follows:

- Two VLANs are consumed on the physical switch for the entire enclave.
- The Cisco Nexus Series Switch provides the policy enforcement point and default gateway (SVI 2001).
- Cisco ASA provides the security group firewall for traffic control enforcement.
- Cisco ASA provides virtual context bridging for two VLANs (VLANs 2001 to 3001 in the figure).
- VXLAN is supported across the infrastructure for virtual machine traffic.
- Consistent security policy is provided through universal security group tags (SGTs):
  - The import of the Cisco ISE protected access credential (PAC) file establishes a secure communication channel between Cisco ISE and the device.
  - Cisco ISE provides SGTs to Cisco ASA, and Cisco ASA defines security group access control lists (SGACLs).
  - Cisco ISE provides SGTs and download-able SGACLs to the Cisco Nexus switch.

  - Cisco ISE provides authentication and authorization across the infrastructure.
- An SGT is assigned on the Cisco Nexus 1000V port profile.
- The Cisco Nexus 1000V propagates IP address-to-SGT mapping across the fabric through the SGT Exchange Protocol (SXP) for SGTs assigned to the enclave.
- The Cisco VSG for each enclave provides Layer 2 firewall functions.
- Load-balancing services are optional but readily integrated into the model.
- Dedicated VMknics are available to meet dedicated NFS and iSCSI access requirements.

Figure 11 Enclave Model: Single VLAN with VXLAN (Cisco ASA Transparent Mode)

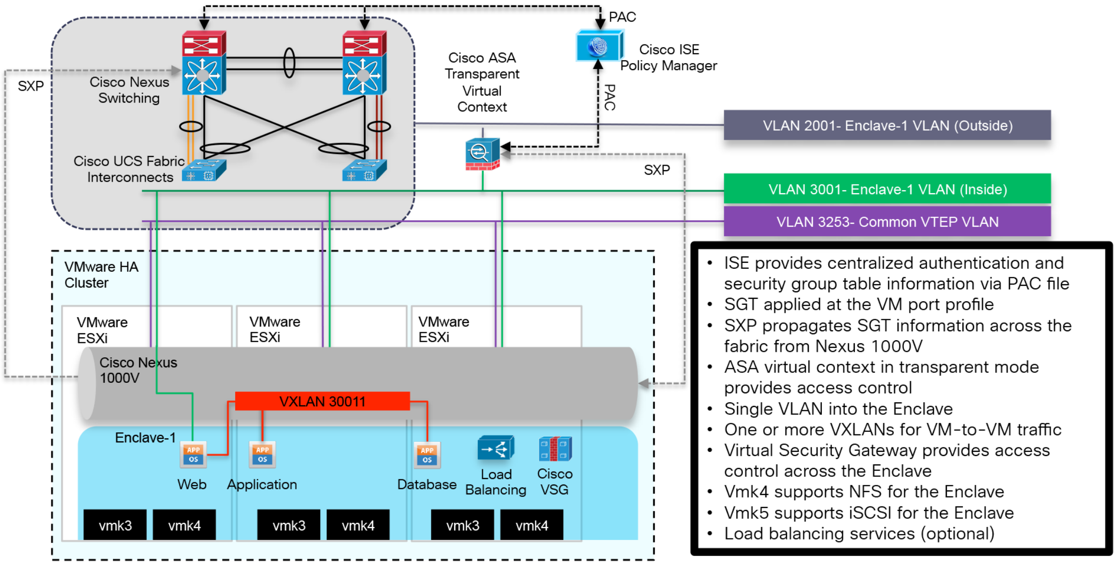

- ISE provides centralized authentication and security group table information via PAC file
- SGT applied at the VM port profile
- SXP propagates SGT information across the fabric from Nexus 1000V
- ASA virtual context in transparent mode provides access control
- Single VLAN into the Enclave
- One or more VXLANs for VM-to-VM traffic
- Virtual Security Gateway provides access control across the Enclave
- Vmk4 supports NFS for the Enclave
- Vmk5 supports iSCSI for the Enclave
- Load balancing services (optional)

Figure 12 illustrates the logical structure of another enclave on the same shared infrastructure employing the Cisco ASA routed virtual context as the default gateway for the web server. The construction of this structure is identical to the previously documented enclave except for the firewall mode of operation.

Security Services

Firewall

Firewalls are the primary control point for access between two distinct network segments, commonly referred to as inside, outside, public or private. The Cisco Secure Enclave Architecture uses two categories of firewalls- zone or edge, for access control into, between and within the enclave. The enclave model promotes security "proximity" meaning where possible, traffic within an enclave will stay local. The use of multiple policy enforcement points promotes optimized traffic paths while catering to security needs in a scalable manner.

Cisco Virtual Security Gateway

The Cisco Virtual Security Gateway (VSG) protects traffic within the enclave, enforcing security policy at the VM level by applying policy based on VM or network based attributes. Typically this traffic is considered "east - west" in nature. The reality is any traffic into a VM is subject to the VSG security policy. The enclave model calls for a single VSG instance per enclave allowing the security operations team to develop granular security rules based on the application and associated business requirements.

The Cisco Nexus 1000v Virtual Ethernet Module (VEM) will redirect the initial packet destined to a VM to the VSG where policy evaluation occurs. The redirection of traffic occurs using vPath when the virtual service is defined on the port profile of the VM. The VEM encapsulates the packet and forwards it to the VSG assigned to the enclave. The Cisco VSG processes the packet and forwards the result to the vPath on the VEM where the policy decision is cached and enforced for subsequent packets. The vPath will maintain the cache until the flow is reset (RST), finished (FIN) or a timeout occurs.

Note The Cisco Virtual Security Gateway may be deployed adjacent to the Nexus 1000v VEM or across a number of Layer 3 hops.

Cisco Adaptive Security Appliances

The edge of the enclave is protected using Cisco's Adaptive Security Appliance (ASA). The Cisco ASA 5585 can be partitioned into multiple security contexts (<250) allowing each enclave to have a dedicated virtual ASA to apply access control, intrusion prevention, and antivirus policy. The primary role of each ASA enclave context is to control access between the "inside and outside" network segments. This traffic is typically referred to as "north - south".The Cisco ASA supports Cisco TrustSec. Cisco TrustSec is an intelligent role-based access solution providing secure network access based on Secure Group Tags (SGT). Contextual data pertaining to the "who, what, where, when, and how," is captured by a SGT at ingress to a domain and the policy (SGACL) enforced at the egress point for a distributed architecture.Cisco TrustSec in the enclave architecture uses Security Group Tag (SGT) assignment on the Nexus 1000v and the ASA as a Security Group Firewall (SGFW) to enforce the role based access control policy.

The Cisco Identity Services Engine (ISE) is a required component in the CiscoTrustSec implementation providing centralized definitions of the SGTs to IP mapping. A Protected Access Credential (PAC) file secures the communication between the ISE and ASA platforms and allows for the ASA to download the security group table. This table contains SGT to security group names translation. The security operations team can then create access rules based on the object tags (SGTs), simplifying policy configuration in the data center.

The SGT is assigned at the VM port profile on the Nexus 1000v. The SGT assignment is propagated to the ASA via the Security eXchange Protocol (SXP). SXP is a secure conversation between the two devices a speaker and listener. The ASA may perform both roles but in this design it is strictly a listener learning and acting as a SGFW. If the IP to SGT mapping is part of a security group policy the ASA enforces the rule.

Threat Context via Cisco Identity Services Engine (ISE)

In order to provide some context, the Lancope StealthWatch system employs the services of the Cisco Identity Services Engine. The ISE can provide device and user information, offering more information for the security operations team to use during the process of threat analysis and potential response. In addition to the device profile and user identity, the ISE can provide time, location, and network data to create a contextual identity to who and what is on the network.

Network Telemetry via NetFlow

NetFlow was developed by Cisco to collect network traffic information and enable monitoring of the network. The data collected by NetFlow provides insight into specific traffic flows in the form of records. The enclave framework uses several methods to reliably collect NetFlow data and provide a full picture of the Data Center environment including:

- NetFlow Generation Appliances (NGA)
- Direct NetFlow Sources
- Cisco ASA 5500 NetFlow Secure Event Logging (NSEL)

The effectiveness of any monitoring system is dependent on the completeness of the data it captures. With that in mind, the enclave model does not recommend using sampled NetFlow. Ideally the NetFlow records should reflect traffic in its entirety. To that end the physical Nexus switches are relieved of NetFlow responsibilities and implement line-rate SPAN. The NGA are connected to SPAN destination ports on the Nexus switches and UCS Fabric Interconnects. The NGA devices are promiscuous supporting up to 40Gbps of mirrored traffic to create NetFlow records for export to the Lancope StealthWatch Flow Collectors.

Direct NetFlow sources generate and send flow records directly to the Lancope FlowCollectors. The Nexus 1000v virtual distributed switch provides this functionality for the virtual access layer of the enclave. It is recommended to enable Netflow on the Nexus 1000v interfaces. In larger environments where the limits of the Nexus 1000v NetFlow resources are reached, NetFlow should be enabled on VM interfaces with data sources.

Another source of direct flow data is the Cisco ASA 5585. The Cisco ASA generates NSEL records. These records differ from traditional NetFlow but are fully supported by the Lancope StealthWatch system. In fact, the records include the "action" permit or deny taken by the ASA on the flow as well as NAT that adds another layer of depth to the telemetry of the CTD system.

Unified Visibility, Analysis and Context Through Lancope StealthWatch

The Lancope StealthWatch system collects, organizes and analyzes all of the incoming data points to provide a cohesive view into the workings of the enclave. The StealthWatch Management Console (SMC) is the central point of control supporting millions of flows. The primary SMC dashboard offers insight into network reconnaissance, malware propagation, command and control traffic, data exfiltration, and internal host reputation.

Cyber Threat Defense

Cyber threats are attacks focused on seizing information related to sensitive data, money or ideas. The Cisco Cyber Threat Defense Solution provides greater visibility into threats by identifying suspicious network traffic patterns within the network providing security analysts the contextual information necessary to discern the level of threat presented by these suspicious patterns. The CTD solution employs three primary components to provide this crucial visibility:

- Network Telemetry via NetFlow
- Threat Context via Cisco Identity Services Engine (ISE)
- Unified Visibility, Analysis and Context via Lancope StealthWatch

Figure 12 Cisco Secure Enclave Cyber Threat Defense Model

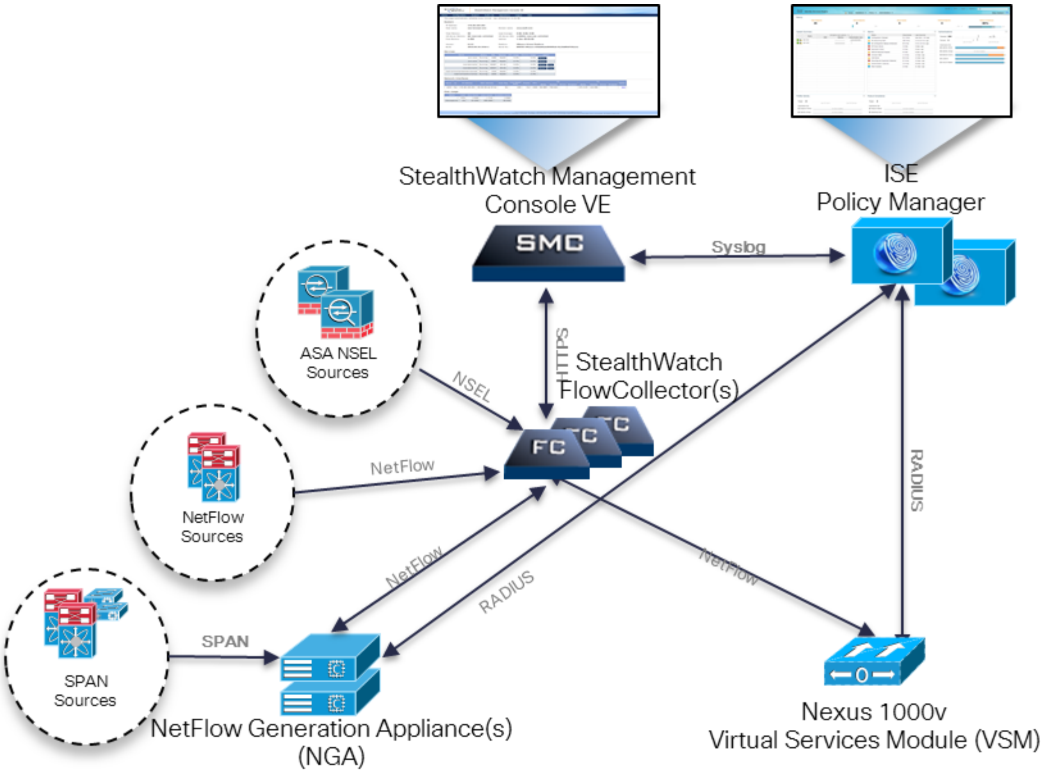

Enclave Management

The enclave management network is a dedicated VLAN providing centralized access, visibility, and control of resources within the system. The management network supports domain- and element-level management to provide comprehensive administration of the shared resources that compose the Cisco Secure Enclave architecture. The administrative interfaces and open APIs available in this management portfolio provide the foundation for delivery of cohesive service lifecycle orchestration with Cisco UCS Director.

Figure 13 shows the management services currently used in the design and their associated host platforms. In practice, multiple Cisco UCS servers, depending on the resource requirements of the installation, are deployed as part of the dedicated VMware vSphere HA and DRS cluster for management. The Cisco Nexus 1110 appliances, which support a number of virtual services blades (VSBs) and contains two instances of the Cisco Nexus 1000V Virtual Supervisor Module (VSM) may also be used for the VSPEX PoD servicing the production enclaves. The Cisco VSG VSB provides Layer 2 security services to the virtual machines in the environment.

Figure 13 Enclave Management Services and Positioning

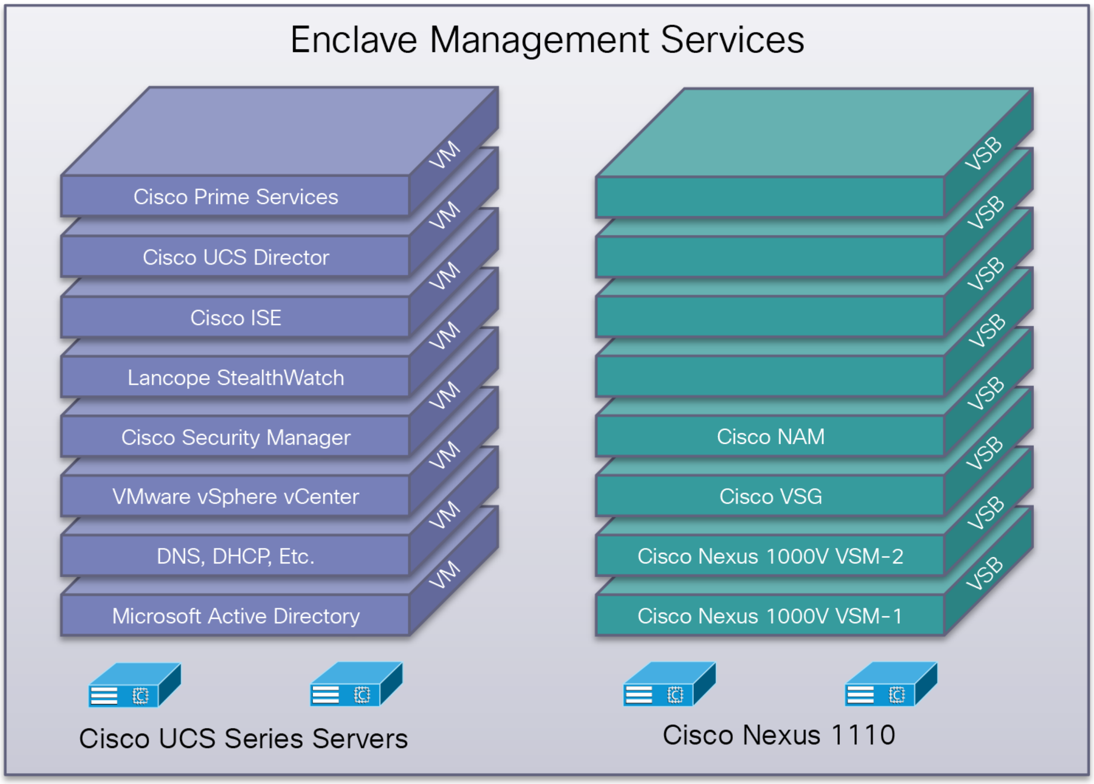

The enclave framework is not restricted to the management services listed in Figure 13.

The communication between the management domain, the hardware infrastructure, and the enclaves is established through traditional paths as well as through the use of private VLANs on the Cisco Nexus 1000V and Cisco UCS fabric interconnects. The use of dedicated out-of-band management VLANs for the hardware infrastructure, including Cisco Nexus switching and the Cisco UCS fabric, is a common best practice. The enclave model suggests the use of a single isolated private VLAN that is maintained between the bare-metal and virtual environments. This private isolated VLAN allows all virtual machines and bare-metal servers to converse with the services in the management domain, which is a promiscuous region. The private VLAN feature enforces separation between servers within a single enclave and between enclaves.

Figure 14 shows the logical construction of this private VLAN environment, which supports directory, DNS, Microsoft Windows Server Update Services (WSUS), and other common required services for an organization.

Figure 14 Private VLANs Providing Secure Access to Core Services

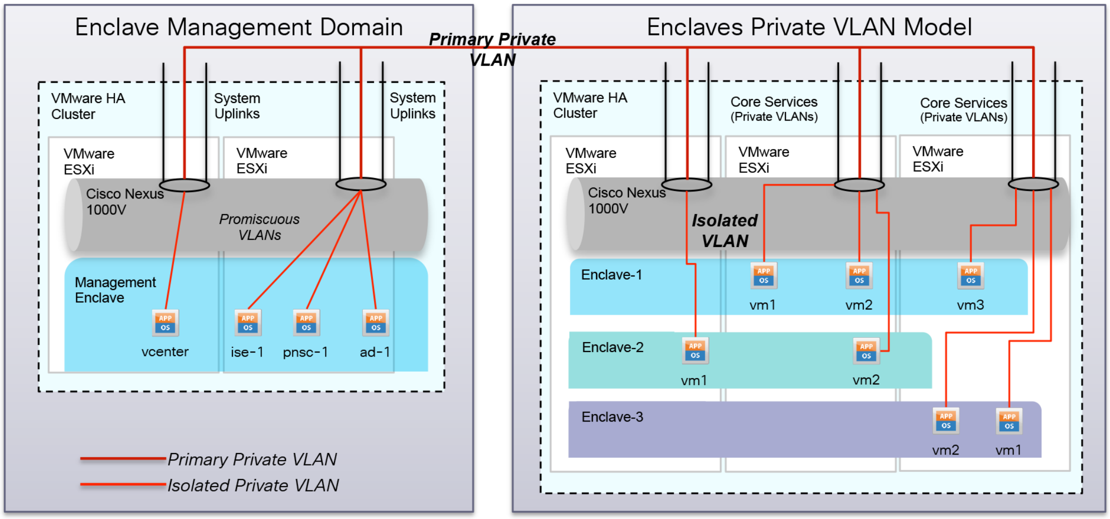

Figure 15 illustrates that the virtual machine connection points to the management domain and the data domain. As illustrated, the traffic patterns are completely segmented through the use of traditional VLANs, VXLANs, and isolated private VLANs. The figure also shows the use of dedicated PCIe devices and logical PortChannels created on the Cisco Nexus 1000V to provide load balancing, high availability, and additional traffic separation.

Figure 15 Enclave Virtual Machine Connections

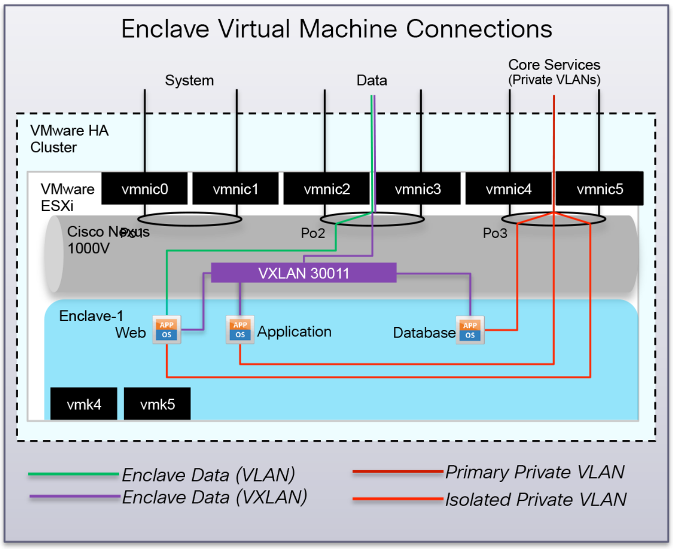

- Core Services use the Isolated Private VLAN
- VM-to-VM traffic over Isolated VLAN is not possible
- Nexus 1000V port channel dedicated to core services supporting only the Isolated Private VLAN
- VXLAN traffic uses the Nexus 1000V port channel dedicated to Enclave data traffic
- Public VLAN uses the Nexus 1000V port channel dedicated to Enclave data traffic

Management Services

The Cisco VSPEX Secure Enclave Data Center employs numerous domain level managers to provision, organize and coordinate the operation of the enclaves on the shared infrastructure. The domain level managers employed during the validation are listed in Table 1 and Table 2. Table 1 describes the role of the management product while Table 2 indicates the positioning of that product within the architecture.

Table 1 Cisco VSPEX Secure Enclave Data Center Management Platforms

Product	Role
Cisco Unified Computing System Manager (UCSM)	Provides administrators a single interface for performing server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection.
Microsoft Active Directory, DNS, DHCP, WSUS.	Microsoft directory services provided centralized authentication and authorization for users and computers. DNS Services are centralized for TCP/IP name translation. DHCP provides automated IP address assignment this is coordinated with the DNS records. Windows Update Services are provided and defined and applied via AD Group Policy. This service maintains the Windows operating systems currency.
VMware vSphere vCenter	Provides centralized management of the vSphere ESXi hosts, virtual machines and enabling of VMware features such as vMotion and DRS cluster services.
Lancope StealthWatch System	Ingests and processes NetFlow records providing unique insight into network transactions, allowing for greater understanding of the network and fine grained analysis of security incidents under its watch.
Cisco Identity Services Engine	Provides user and device identity and context information to create policies that govern authorized network access. ISE is the policy control point of the Cisco TrustSec deployment allowing for centralized object based security.
Cisco Prime Network Services Controller	Provides centralized device and security policy management of the Cisco Virtual Security (VSG) and other virtual services.
EMC VNX Unisphere	Provides comprehensive browser based management of VNX arrays with support for Fiber Channel (block based) or File based (NFS/CIFS) storage.
Cisco Nexus 1000v Virtual Supervisor Module for VMware vSphere	Provides a comprehensive and extensible architectural platform for virtual machine (VM) and cloud networking
Cisco Virtual Security Gateway	Delivers security, compliance, and trusted access for virtual data center and cloud computing environments

Table 2 Cisco VSPEX Secure Enclave Data Center Management Platform Positioning

Product	Positioned
Microsoft Active Directory, DNS, DHCP, WSUS.	VMware vSphere Management Cluster – Common Infrastructure (C.I)
VMware vSphere vCenter	VMware vSphere Management Cluster (C.I)
Lancope StealthWatch System	VMware vSphere Management Cluster (C.I)
Cisco Identity Services Engine	VMware vSphere Management Cluster (C.I)
Cisco Prime Network Services Controller	VMware vSphere Management Cluster (C.I)
EMC VNX Unisphere	VNX Storage System.

Cisco Nexus 1000v Virtual Supervisor Module	VMware vSphere Management Cluster (C.I)
Cisco Virtual Security Gateway	VMware vSphere Management Cluster (C.I)

Cisco UCS Director provides a central user portal for managing the environment and enables the automation of the manual tasks associated with the provisioning and subsequent operation of the enclave.

Figure 16 Cisco UCS Director Key Features

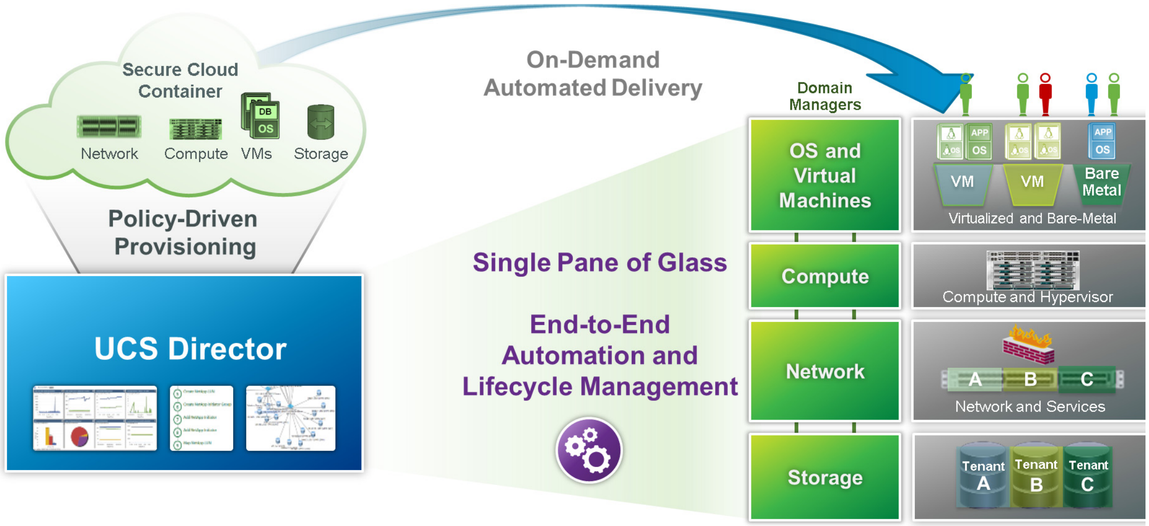

Cisco UCS Director can directly or indirectly manage the enclave components. Ideally, the north bound APIs of the various management domains are used but UCS Director may also directly access devices to create the Enclave environment.

Figure 17 Cisco UCS Director Management Structure

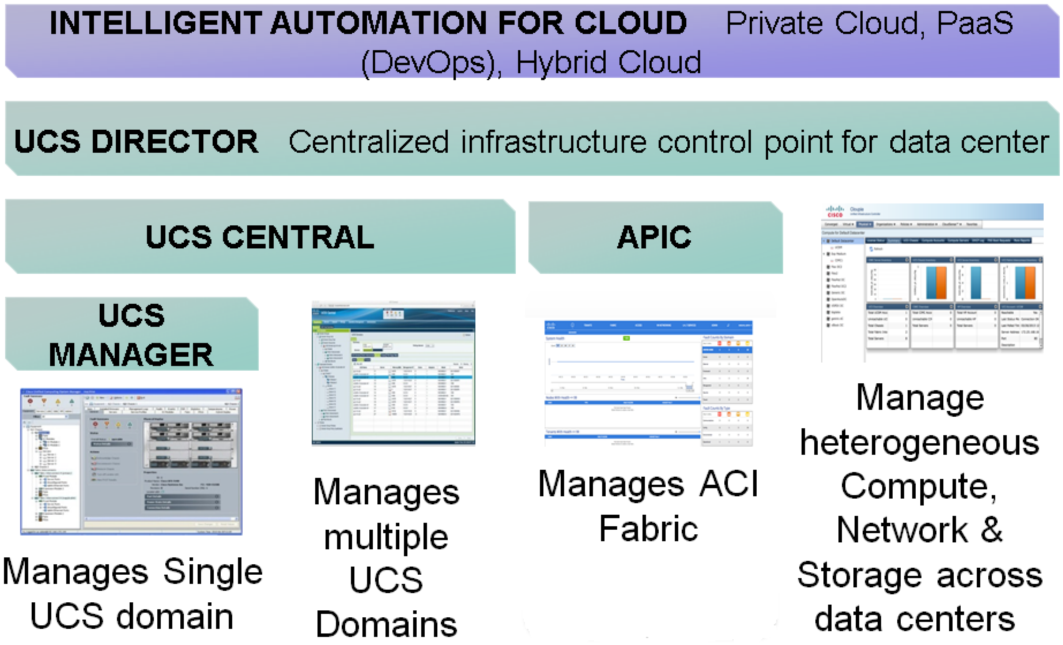
INTELLIGENT AUTOMATION FOR CLOUD Private Cloud, PaaS (DevOps), Hybrid Cloud
UCS DIRECTOR Centralized infrastructure control point for data center
UCS CENTRAL
APIC
UCS MANAGER
Manages Single UCS domain
Manages multiple UCS Domains
Manages ACI Fabric
Manage heterogeneous Compute, Network & Storage across data centers

Note The Cyber Threat Defense components are not directly accessed since the protections are overlays encompassing the entire infrastructure. Figure 16 shows the interfaces that Cisco UCS Director employs.

Figure 18 Cisco UCS Director Enclave Control Framework

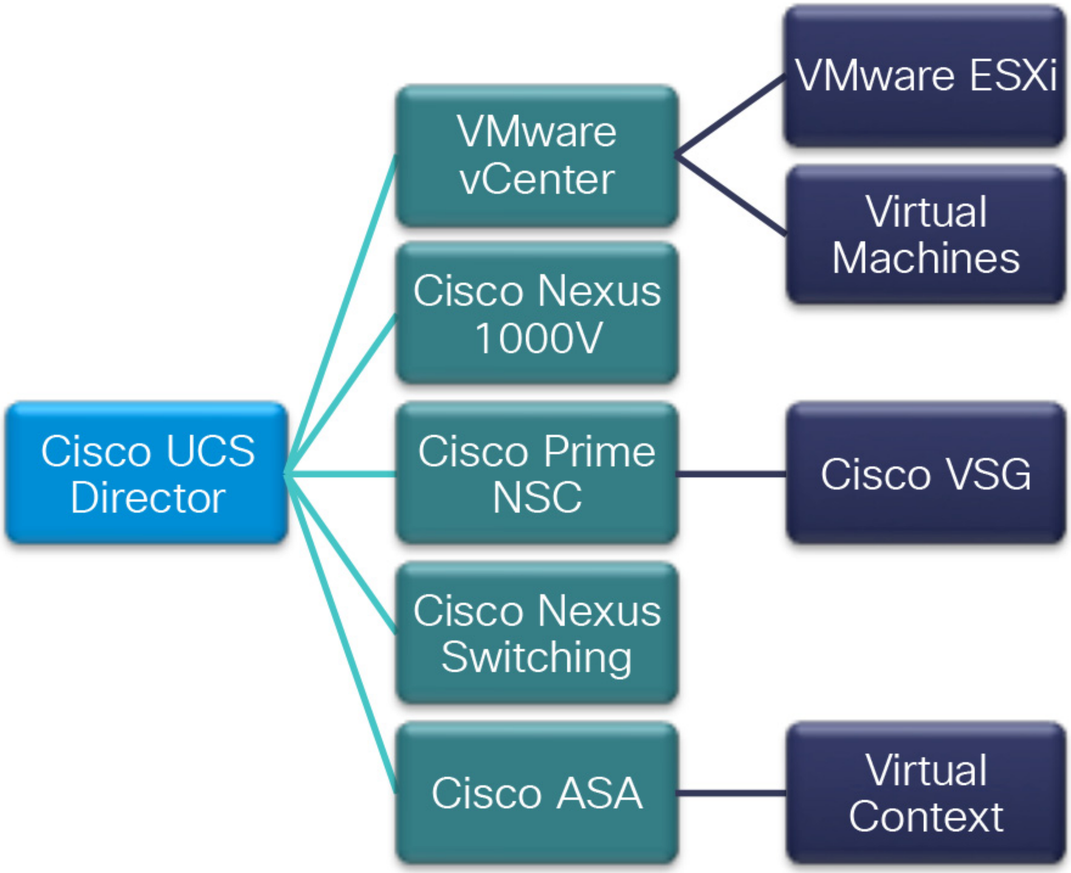

The instantiation of multiple enclaves on the VSPEX Data Center platform through Cisco UCS Director offers operational efficiency and consistency to the organization. Figure 18 illustration points out the hierarchical structure of a Cisco UCS Director lead approach to automate the infrastructure through a single pane of glass.

Traffic Patterns

Traffic patterns in the shared infrastructure can be divided into two categories: north-south (client to server) and east-west (server to server); see Figure 17.

Client-to-Server (North-South) Flows

North-south traffic flows are either ingress or egress flows from the enclave perspective. In this deployment, the ASA 5585 operates in Layer 2 transparent mode. Using the Layer 2 transparent mode of the ASA 5585-X makes for a simpler deployment model with less impact to the existing network architecture when deploying the solution. North-south traffic flows represent an increased risk of including malicious traffic, so Cisco recommends that customers consider identifying some or all of the traffic to be monitored by the Cisco IPS module in the ASA 5585-X NextGen firewall. This traffic traverses the enclave and is exposed to any number of services in its path, including firewalls, load balancers, intrusion detection, and network analysis devices in the enclave. In a multiple-enclave environment, traffic between enclaves (enclave-to-enclave traffic) is also considered north-south in nature and therefore is subject to each enclave's independent policies as it progresses.

Server-to-Server (East-West) Flows

East-west traffic refers to the communication between servers within an enclave. East-west protection in the virtualization layer or in the Secure Enclaves is achieved using the Cisco Virtual Security Gateway (VSG) along with the Cisco Nexus 1000V Virtual Ethernet Switch. The Cisco Nexus 1000V communicates with the VSG using a message bus called vPath to provide efficient policy enforcement as well as service chaining to ensure the expected traffic flows through the virtualized appliances. The Cisco Nexus 1000V provides additional capability such as the ability to apply an SGT to the virtual machine at the time of the provisioning and deployment of the virtual machine. The SGT can be assigned manually or automatically with the use of the Cisco UCS Director in future releases.

Figure 19 Enclave Traffic Patterns

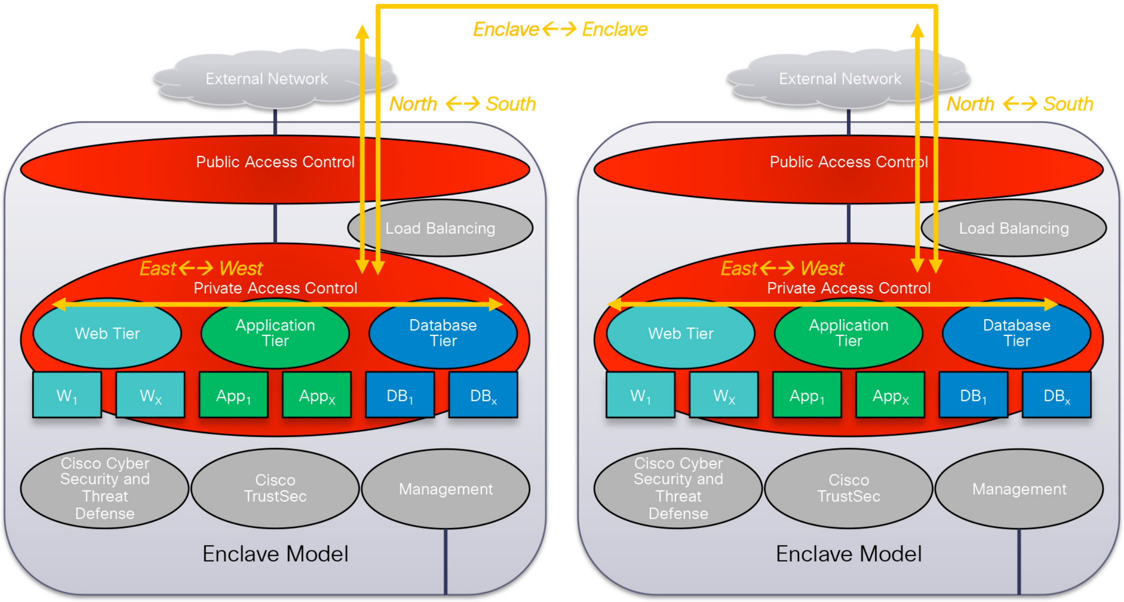

Design Considerations

The key to developing a robust design is clearly defining the requirements and applying a proven methodology based on sound design principles of:

- Protection
- Performance
- Provisioning
- Availability
- Service assurance

A framework that provides secure administrative and user-domain protection with application-level service assurance through quality of service (QoS) delivered by dedicated high-performance appliances on a highly available platform constitute the foundation on which complementary products can be deployed to provide customer-specific features. Another essential component of such a design is automation of resource provisioning to provide operational efficiency and help ensure implementation accuracy.

Protection

Given the borderless nature of users and access methods currently in use, the security mechanism needs to be ubiquitous (defense in breadth) and deep (defense in depth) to eliminate both circumvention and penetration, which can lead to intrusion. Cisco ASA firewalls serve as the first line of defense for both ingress and egress traffic to and from the enclave. Cisco VSG integrated into the Cisco Nexus 1000V is a virtual firewall that provides distinct trust zones on shared computing infrastructure for east-west traffic between virtual machines. Together, the Cisco ASA firewall and Cisco VSG provide protection against perimeter attacks as well as internal attacks. This comprehensive protection safeguards against disruptions to critical administrative functions of the cloud infrastructure so that valuable shared user-domain resources are protected.

Access to resources such as virtual machines, network bandwidth, data, and storage needs to be curtailed at both the logical and physical levels, where possible, to impose necessary controls. Protection against denial-of-service (DoS) attacks and unauthorized access leading to data loss is delivered through the threat analysis and zero-day protection features of Lancope StealthWatch.

Preservation of user-space confidentiality through encryption and other means at multiple levels through use of access controls, virtual storage controllers, VLAN segmentation, firewall rules, and intrusion protection should be employed where possible. Data protection through continuous encryption of data in flight and at rest is essential for integrity. Cisco TrustSec SGT support on most Cisco devices is crucial to enabling proper access control in a distributed manner for a scalable and secure platform. Assessing the efficacy of the implementation and adapting to defend against new and evolving threats require continuous and comprehensive visibility into the operations of the network and its components. Cisco NetFlow implementations along with flow analysis by Lancope StealthWatch are invaluable in this area. Together, Cisco NetFlow and Cisco TrustSec deliver visibility and control, which are essential for an open and secure platform.

The enforcement module with Cisco TrustSec and Cisco NetFlow extends its reach into underlying networking, computing, storage, and management components within this architecture to provide features and capabilities that together provide a trusted environment.

Performance

Delivery of security involves consumption of computing resources. Providing security in a consistently responsive manner requires even more resources from potentially a shared pool that may be needed by hosted applications. Although these resources may appear as overhead, user confidentiality is essential. Thus, organizations need to implement secure services that are scalable with the least overhead. Previously discussed design principles such as the least-common mechanism and efficient mediated access are shown to provide guidance in devising a platform that delivers on these needs. The goal is to provide sufficient performance to help ensure that the necessary security checks are performed within the permitted time and before the user experience becomes a concern.

To this end, the high throughput of Cisco ASA firewalls and Cisco NetFlow appliances with managed device capabilities removes overhead from the underlying shared platform used by hosted applications. This design methodology also leads to a scalable cluster of firewalls and Cisco NetFlow devices that can grow to accommodate the needs of the enterprise with a secure and separate enforcement module that conforms to the reference monitor. Virtual PortChannel (vPC) technology on both physical and virtual implementations of Cisco Nexus switches provides link-layer resiliency and additional bandwidth as needed.

Figure 17 introduces the Cisco ASA platforms to a secure enclave as a clustered service. The Cisco ASA cluster model scales up to a maximum of eight nodes managed as a single unit. In clustered mode, every member of the cluster is capable of forwarding every traffic flow and can be active for all flows.

Note Performance benchmarking is out of scope of this validation effort. However, documents that speak to performance of individual components is included in the reference section.

Note Currently, Cisco ASA clustering is supported only for the Cisco Nexus 7000 Series switching platforms because of the need for Cisco Link Aggregation Control Protocol (LACP) support.

Figure 20 Physical Extensions of Cisco Integrated Systems with Clustered ASA Security Services

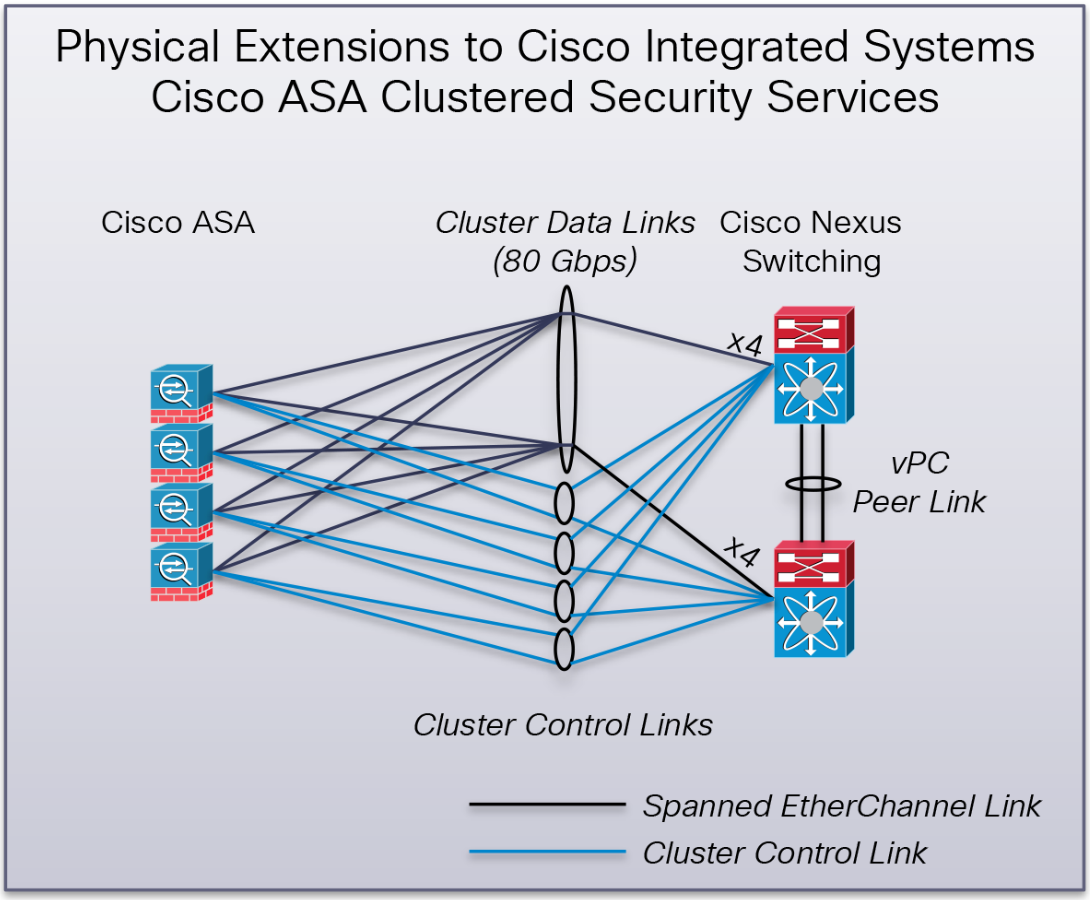
Physical Extensions to Cisco Integrated Systems
Cisco ASA Clustered Security Services
Cisco ASA
Cluster Data Links (80 Gbps)
Cisco Nexus Switching
x4
vPC Peer Link
x4
Cluster Control Links
Spanned EtherChannel Link
Cluster Control Link

Provisioning: Ease of Management

Provisioning design includes such features as rapid and correct provisioning for easy adoption and for effective management of resources. Domain and element management provides comprehensive administration of the shared resources that compose the secure cloud architecture. The demarcation point for managing components in this design is defined by individual programmable interfaces delivered by Cisco and partner products. The administrative interfaces and APIs in this portfolio address infrastructure components such as VMware vCenter, Cisco UCS Manager, Cisco Data Center Network Manager (DCNM), and storage managers. These element managers and their associated open APIs provide the foundation for delivery of cohesive service lifecycle orchestration with Cisco UCS Director. At a logical level, Cisco UCS Director integrates infrastructure components into a single management pane (Figure 21).

Figure 21 Cisco UCS Director Abstracts Infrastructure Layers into a Single Management Pane

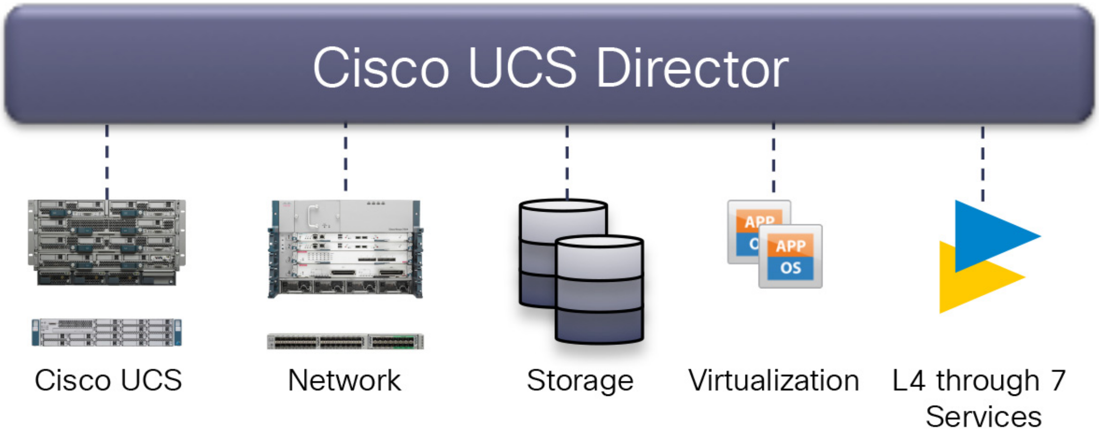

High Availability

High availability helps ensure that systems and data are available and accessible to authorized users when they are needed by introducing redundancy at every layer of the infrastructure: computing, network, and storage. One other desirable outcome of the elimination of single points of failure is a setup that allows planned maintenance with little or no disruption in most cases. Each layer has its own way of providing a highly available configuration that works transparently with adjacent layers.

Service Assurance

Service assurance requires available controls and components to help ensure that the SLAs expected from the platform, including security, are met. The components and features necessary to deliver agreed-on system performance pertaining to underlying components such as computing, networking, and storage resources during both steady-state and non-steady-state environments are covered. For example, the network and Cisco UCS blade architectures can provide detailed bandwidth guarantees using QoS; resource pools in VMware help balance and guarantee CPU and memory resources, while comparable features at the storage level support declared I/O operations per second (IOPS) deliverables.

Deployment

Base Platform

This Document assumes that you have followed the procedure detailed in the link below to build the base VSPEX platform:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/ucs_vspex_vmwpc_51.html

Physical devices added to the above VSPEX infrastructure include a pair each of Cisco ASA 5585 firewalls with TrustSec capability and Cisco NGA3240 Netflow generation devices. Following is the high-level architecture diagram for all devices in this solution. Common infrastructure management services and applications such as Active Directory, DNS, DHCP, SMTP and NTP and applications including VMware vCenter, Cisco UCS Director with the Bare-Metal agent, Cisco Prime Network Services Controller (PNSC), Cisco Identity Services Engine (ISE) and Lancope's StealthWatch network analysis tool and Cisco Nexus 1000v Virtual Supervisor Module (VSM) are hosted external to the PoD (VSPEX) as shown below. Common Infrastructure components need to be highly redundant to ensure un-interrupted service as the applications residing in this space are shared and critical to the operation of the entire Data Center which could include multiple such converged stacks. The focus is on using a validated converged infrastructure (VSPEX) with tested security components to ensure compatibility and other requirements including necessary automation with Cisco UCS Director.

Figure 22 High-Level Architecture

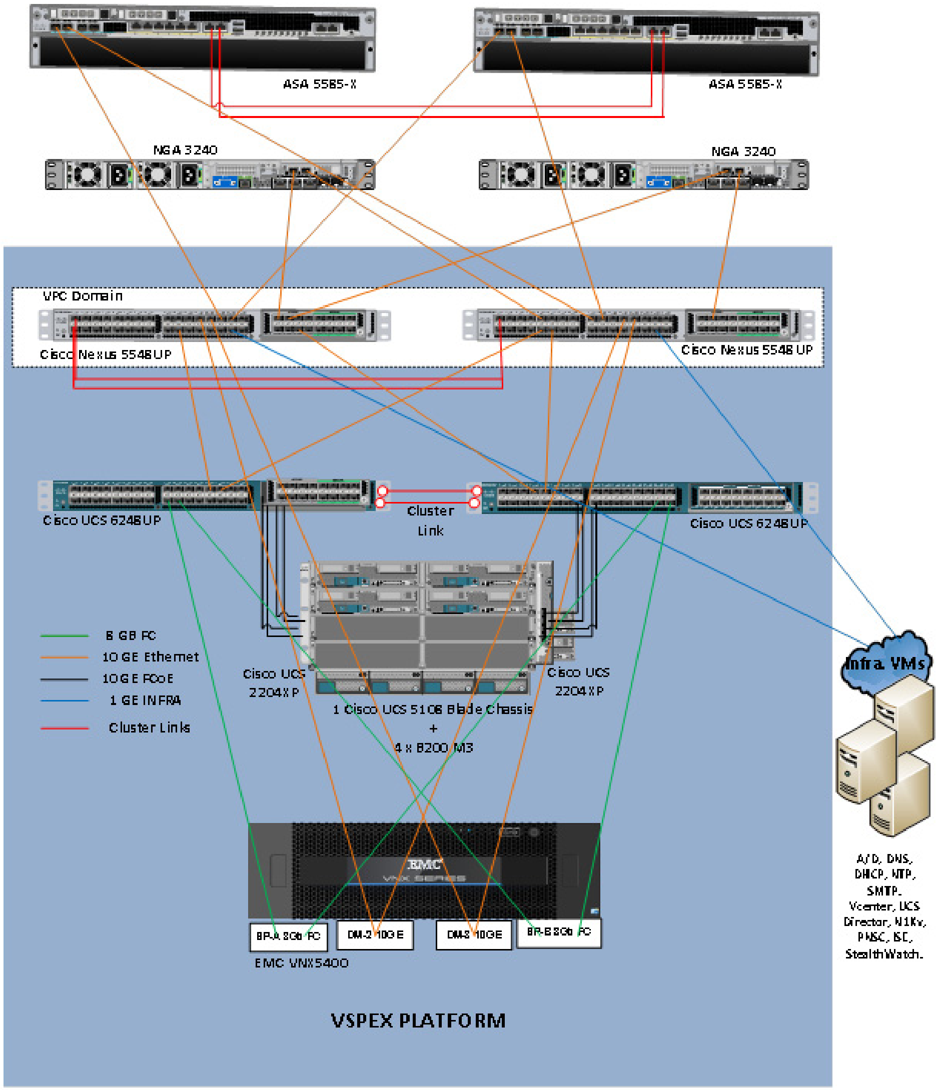

Figure 23 Cabling Diagram

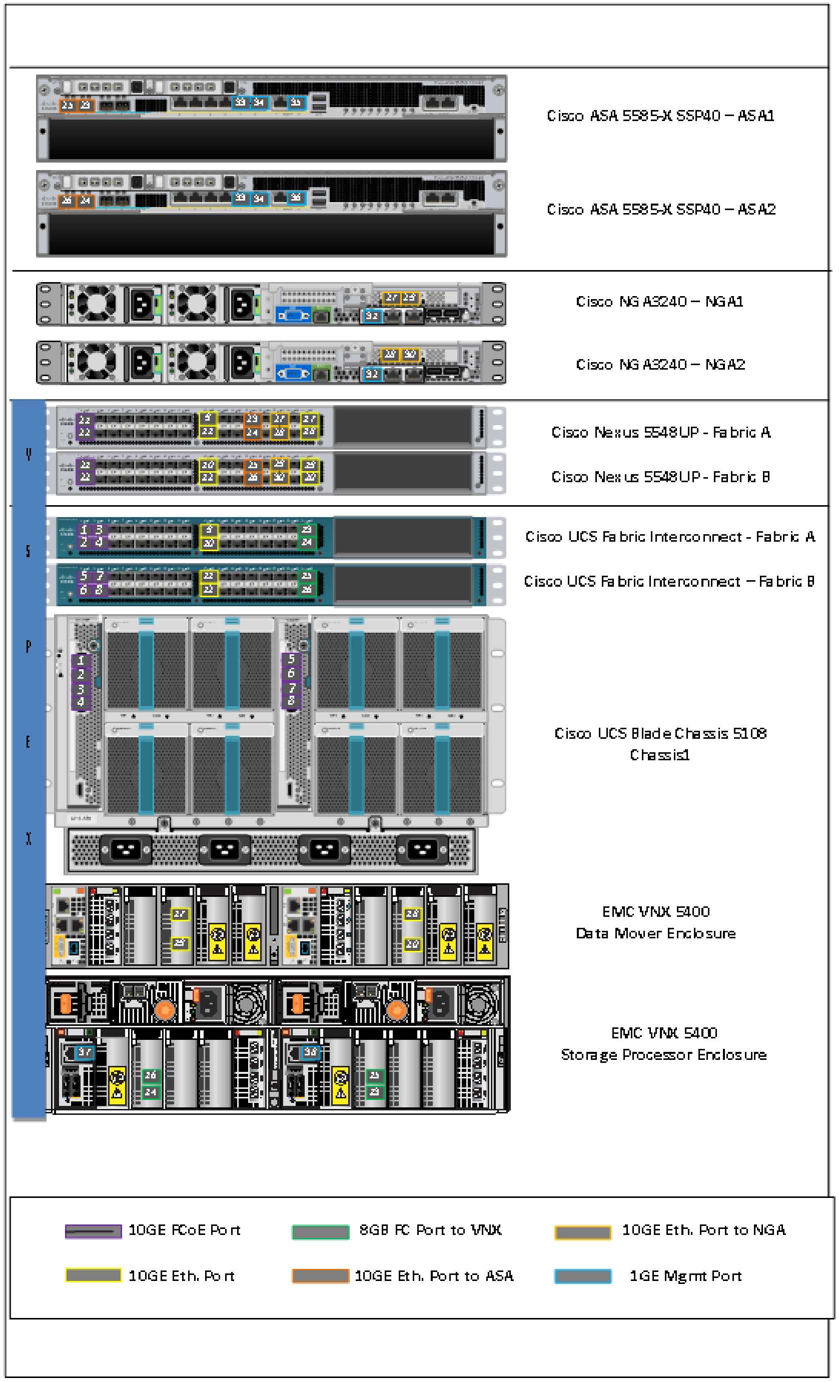
Cisco ASA 5585-X SSP40 – ASA1
Cisco ASA 5585-X SSP40 – ASA2
Cisco NGA3240 – NGA1
Cisco NGA3240 – NGA2
Cisco Nexus 5548UP - Fabric A
Cisco Nexus 5548UP - Fabric B
Cisco UCS Fabric Interconnect - Fabric A
Cisco UCS Fabric Interconnect – Fabric B
Cisco UCS Blade Chassis 5108
Chassis1
EMC VNX 5400
Data Mover Enclosure
EMC VNX 5400
Storage Processor Enclosure
V
S
P
E
X
10GE FCoE Port
8GB FC Port to VNX
10GE Eth. Port to NGA
10GE Eth. Port
10GE Eth. Port to ASA
1GE Mgmt Port

Base Platform Variations

Cisco Unified Computing System

The Cisco Unified Computing System configuration is based upon the tested and recommended practices of VSPEX platform. The enclave architecture will build off of this baseline deployment to instantiate new Service Profiles and the objects required for their instantiation.

Cisco UCS Service Profile

Service profiles are the central concept of Cisco UCS. Each service profile serves a specific purpose: ensuring that the associated server hardware has the configuration required to support the applications it will host.

The service profile maintains configuration information about the server hardware, interfaces, fabric connectivity, and server and network identity. This information is stored in a format that you can manage through Cisco UCS Manager. All service profiles are centrally managed and stored in a database on the fabric interconnect.

Every server must be associated with a service profile. The VSPEX Data Center baseline service profiles were used to build the enclave environment. Modifications were made in regards to the QoS policy of the service profiles, as well as, the number of VNICs instantiated on a given host.

Whether Cisco UCS controls the CoS for a vNIC or not strictly depends on the Host Control field of the QoS Policy, which is assigned to that particular vNIC. Referring to Figure 24, the QoS_N1k policy allows full host control. Since Full is selected and if the packet has a valid CoS assigned by the Nexus 1000v, then UCS trusts the CoS settings assigned at the host level. Otherwise, Cisco UCS uses the CoS value associated with the priority selected in the Priority drop-down list, in this case Best Effort. The None selection indicates that the UCS will assign the CoS value associated with the Priority Class given in the QoS policy, disregarding any of the settings implemented at the host level by the Nexus 1000v.

Figure 24 Cisco UCS QoS Policy—Allow Host Control

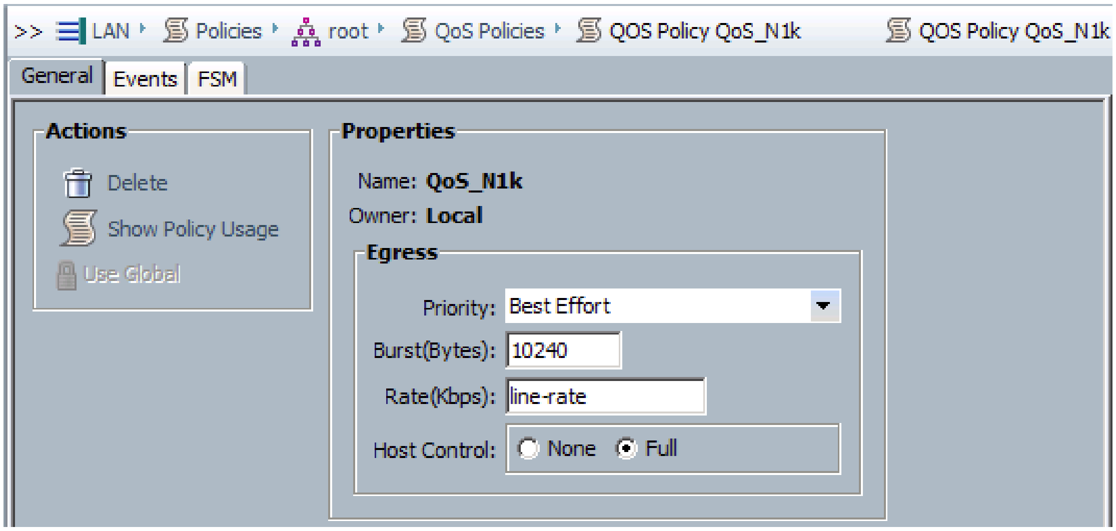

Name: QoS_N1k
Owner: Local

Priority: Best Effort
Burst(Bytes): 10240
Rate(Kbps): line-rate
Host Control: ○ None ● Full

The vNIC template uses the QoS policy to defer classification of traffic to the host or in the enclave architecture the Cisco Nexus 1000v. Figure 25 is a sample vNIC template where the QoS Policy and MTU are defined for any Service Profile using this template.

Figure 25 Service Profile vNIC Template Example

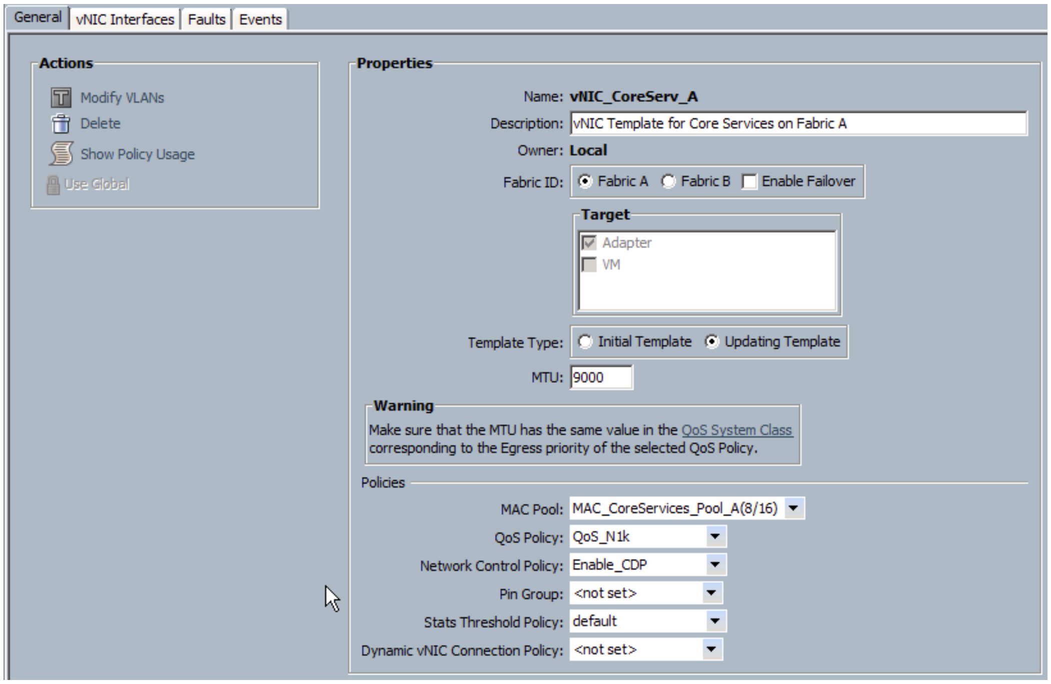

Properties

Name: **vNIC_CoreServ_A**

Description: vNIC Template for Core Services on Fabric A

Owner: **Local**

Fabric ID: ● Fabric A ○ Fabric B □ Enable Failover

Target

☑ Adapter

□ VM

Template Type: ○ Initial Template ● Updating Template

MTU: 9000

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

Policies

MAC Pool: MAC_CoreServices_Pool_A(8/16)

QoS Policy: QoS_N1k

Network Control Policy: Enable_CDP

Pin Group: <not set>

Stats Threshold Policy: default

Dynamic vNIC Connection Policy: <not set>

Figure 26 illustrates all of the vNIC templates defined for the production servers in the enclave VMware DRS cluster. Each template uses the QoS_N1k QoS policy and an MTU of 9000. The naming standard also indicates there is fabric alignment of the vNIC to Fabric Interconnect A or B. Figure 27 is the example adapter summary for the enclave service profile.

Figure 26 Cisco UCS vNIC Templates for Enclave Production Servers

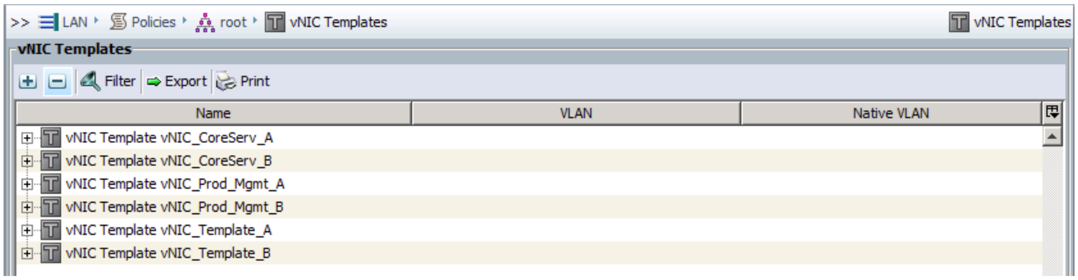

vNIC Templates

Name	VLAN	Native VLAN
vNIC Template vNIC_CoreServ_A		
vNIC Template vNIC_CoreServ_B		
vNIC Template vNIC_Prod_Mgmt_A		
vNIC Template vNIC_Prod_Mgmt_B		
vNIC Template vNIC_Template_A		
vNIC Template vNIC_Template_B		

Figure 27 Cisco UCS ESXi Host Service Profile

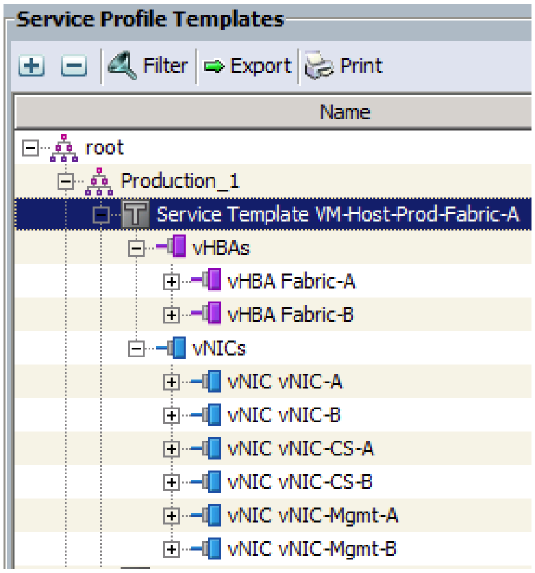

Service Profile Templates

- root
  - Production_1
    - Service Template VM-Host-Prod-Fabric-A
      - vHBAs
        - vHBA Fabric-A
        - vHBA Fabric-B
      - vNICs
        - vNIC vNIC-A
        - vNIC vNIC-B
        - vNIC vNIC-CS-A
        - vNIC vNIC-CS-B
        - vNIC vNIC-Mgmt-A
        - vNIC vNIC-Mgmt-B

User Management

The Cisco UCS domain is configured to use the radius services of the ISE for user management, centralizing authentication and authorization policy in the organization. The following configurations were put in place to achieve this goal.

- Create a radius provider
- Create a radius group
- Define an authentication domain
- Revise the Native Authentication policy

The following screen shots step through the Cisco UCS integration of ISE radius services.

Note The figures include the Cisco UCS navigation path.

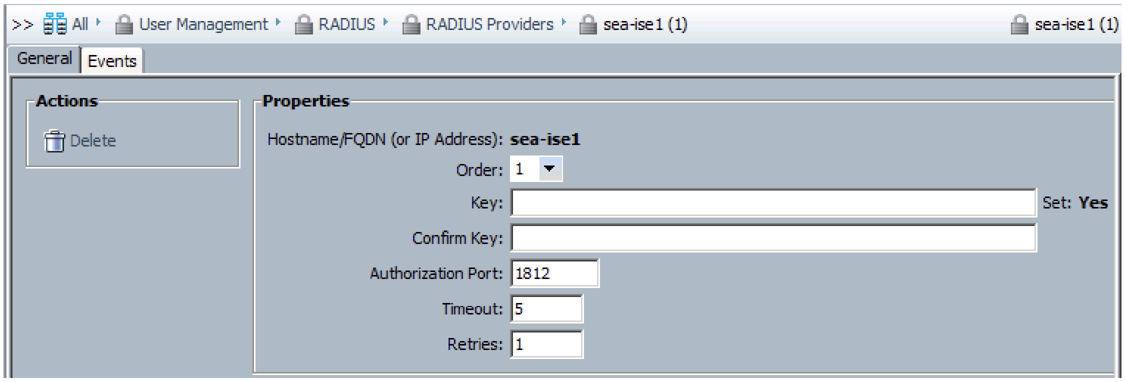

>> All ▸ User Management ▸ RADIUS ▸ RADIUS Providers ▸ sea-ise1 (1)

sea-ise1 (1)

Hostname/FQDN (or IP Address): **sea-ise1**
Order: 1
Key:
Set: **Yes**
Confirm Key:
Authorization Port: 1812
Timeout: 5
Retries: 1

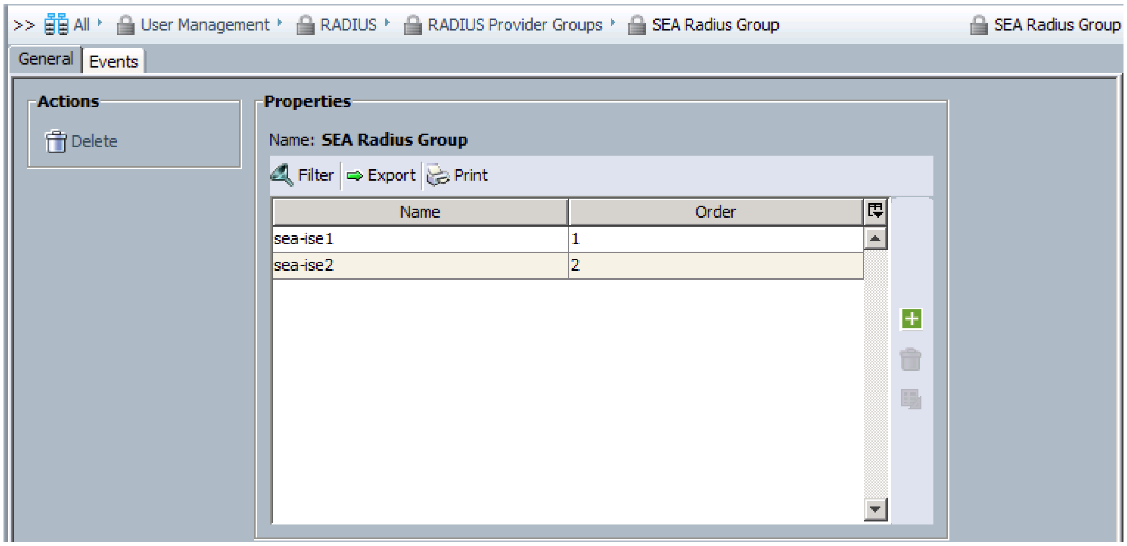
>> All ▸ User Management ▸ RADIUS ▸ RADIUS Provider Groups ▸ SEA Radius Group
SEA Radius Group
General
Events
Actions
Delete
Properties
Filter
Export
Print

Name: **SEA Radius Group**

Name	Order
sea-ise1	1
sea-ise2	2

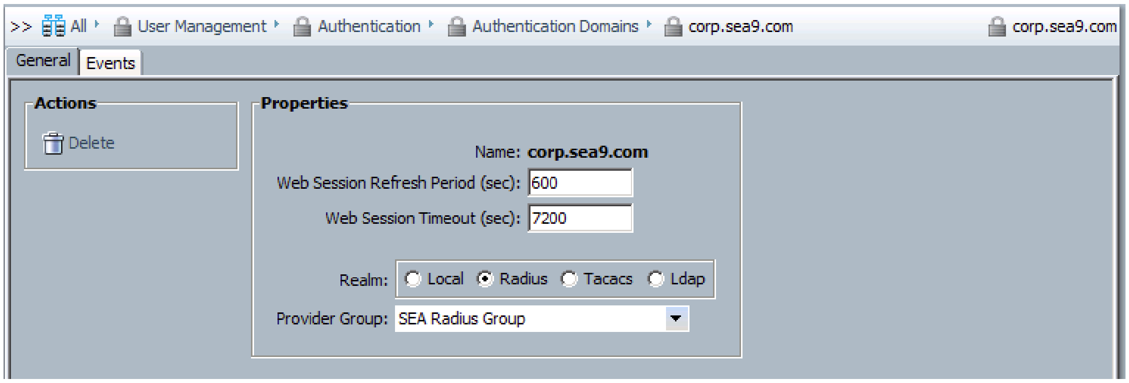
>> All ▸ User Management ▸ Authentication ▸ Authentication Domains ▸ corp.sea9.com
corp.sea9.com
General
Events
Actions
Delete
Properties

Name: **corp.sea9.com**

Web Session Refresh Period (sec): 600

Web Session Timeout (sec): 7200

Realm: ○ Local ● Radius ○ Tacacs ○ Ldap

Provider Group: SEA Radius Group

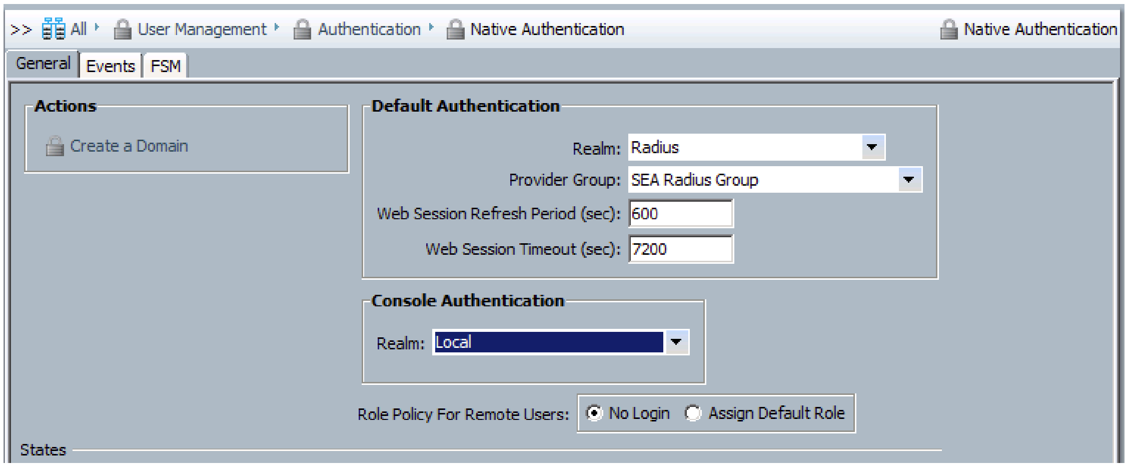
>> All ▸ User Management ▸ Authentication ▸ Native Authentication
Native Authentication
General
Events
FSM
Actions
Create a Domain
States

Default Authentication

Realm: Radius

Provider Group: SEA Radius Group

Web Session Refresh Period (sec): 600

Web Session Timeout (sec): 7200

Console Authentication

Realm: Local

Role Policy For Remote Users: ● No Login ○ Assign Default Role

VMware vSphere

ESXi

The ESXi hosts are uniform in their deployment employing the FC boot practices established in the VSPEX PoD. The Cisco UCS service profile is altered to provides 6 vmnics for use by the hypervisor as described in the previous section. The following sample from one of the ESXi hosts reflects the Cisco UCS VNIC construct and MTU settings provided by the Cisco Nexus 1000v.

ESXi Host Example

```
~ # esxcfg-nics -l
Name    PCI          Driver              Link Speed    Duplex MAC Address  MTU Description
vmnic0  0000:06:00.00 enic      Up   40000Mbps Full   00:25:b5:02:0a:04 9000   Cisco Systems Inc
Cisco VIC Ethernet NIC
vmnic1  0000:07:00.00 enic      Up   40000Mbps Full   00:25:b5:02:0b:04 9000   Cisco Systems Inc
Cisco VIC Ethernet NIC
vmnic2  0000:08:00.00 enic      Up   40000Mbps Full   00:25:b5:02:5a:04 9000   Cisco Systems Inc
Cisco VIC Ethernet NIC
vmnic3  0000:09:00.00 enic      Up   40000Mbps Full   00:25:b5:02:5b:04 9000   Cisco Systems Inc
Cisco VIC Ethernet NIC
vmnic4  0000:0a:00.00 enic      Up   40000Mbps Full   00:25:b5:02:3a:04 9000   Cisco Systems Inc
Cisco VIC Ethernet NIC
vmnic5  0000:0b:00.00 enic      Up   40000Mbps Full   00:25:b5:02:3b:04 9000   Cisco Systems Inc
Cisco VIC Ethernet NIC
```

The vmknics vmko, vmk1 and vmk2 are provisioned for infrastructure services management, vMotion and VXLAN VTEP respectively.

```
Interface  Port Group/DVPort   IP Family IP Address                    Netmask       Broadcast      MAC
Address       MTU     TSO MSS   Enabled Type

vmk2       704              IPv4      192.168.1.90                 255.255.255.0   192.168.1.255
00:50:56:62:27:83 1500    65535     true    STATIC

vmk1       128              IPv4      10.10.10.90                  255.255.255.0   10.10.10.255
00:50:56:6b:b7:17 9000    65535     true    STATIC

vmk0       34               IPv4      10.29.131.90                 255.255.255.0   10.29.131.255
00:25:b5:02:3b:04 1500    65535     true    STATIC
```

Each enclave has a dedicated NFS mount-point made available to it from the EMC VNX storage array. Secure separation at the disk level is provided by carving out a separate storage pool for each of the NFS area. One method to ensure access control is exercised within this setup is to map a minimum of two ESX hosts per VMware cluster and apply host IP based read/write access from the storage side based on target IPs. However, this approach requires allocating a minimum of two blade resources to each enclave/cluster and this could lead to stranded capacity. This design aims to create VMware clusters with more resources and enclaves within to accord greater flexibility and lower stranded capacity. In taking this approach, IP based storage side access control is relinquished for hypervisor and TrustSec based control. Another add-on could be VMware vShield to further enforce access restrictions to shared NFS space.

DRS for Virtual Service Nodes

The VMware DRS cluster provides affinity controls and rules for VM and ESXi host alignment. In the enclave design, virtual services are retained within the production cluster to manage traffic patterns and offer the performance inherent to locality. To avoid a single point of failure, the ESXi host, and DRS cluster setting were modified and placement policies created.

Two virtual machine DRS Groups were created indicating the primary and secondary members of HA pairs. In this example, the Primary VSG and Secondary VSG are instantiated and VSGs were assigned to each group as appropriate. The DRS production cluster ESXi host resources were "split" into two categories based on the naming standard of odd and even ESXi hosts.

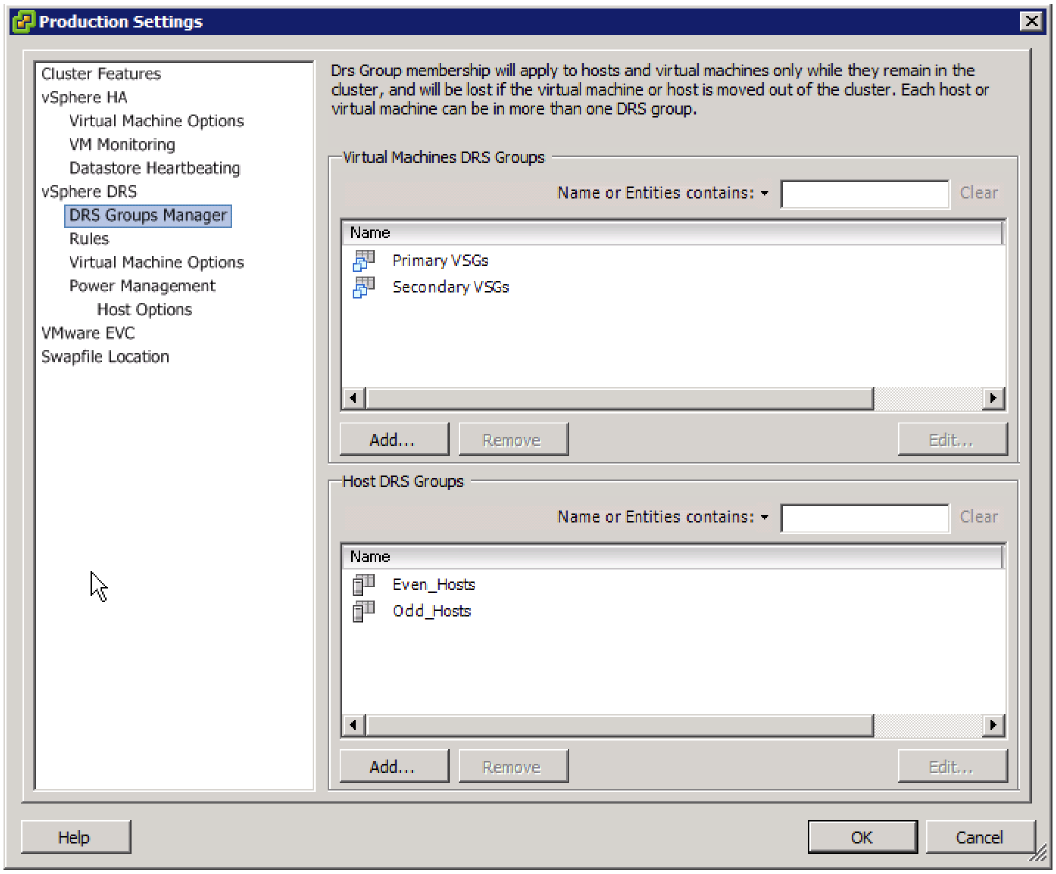

Drs Group membership will apply to hosts and virtual machines only while they remain in the cluster, and will be lost if the virtual machine or host is moved out of the cluster. Each host or virtual machine can be in more than one DRS group.

Virtual Machines DRS Groups

Name or Entities contains:

Name
Primary VSGs
Secondary VSGs

Host DRS Groups

Name or Entities contains:

Name
Even_Hosts
Odd_Hosts

Two DRS virtual machine rules were created defining the acceptable positioning of VSG services on the DRS cluster. As shown, the previously created DRS cluster VM and Host groups are used to define two distinct placement policies in the cluster, essentially removing the ESXi host as a single point of failure for the identified services (VMs).

Storage

In this deployment, the need is for flexibility in resourcing the enclave at the virtual level while preventing un-authorized data access. To this end, ESXi boot LUNs are grouped in a separate Fiber Channel (FC) storage pool and shared by all ESXi hosts within the PoD. Data, also on the SAN through Network File System (NFS), is mapped as separate mount-points. To make sure there is a secure separation at the disk level, one data storage pool is created for each enclave data NFS mount-point. All virtual Machines (VMs) within the enclave share this NFS pool for data. Since each enclave will span all ESXi hosts, this leads to a situation where all NFS mount-points will be visible at the host level to each of the ESXi hosts within the PoD. However, user access controls at the hypervisor level (VMware) coupled with TrustSec (SGT) control from the security layer ensure users in one enclave will not have access to NFS space from another enclave. Further access controls may be exercised through VMware vShield if desired.

Figure 28 Storage Layout

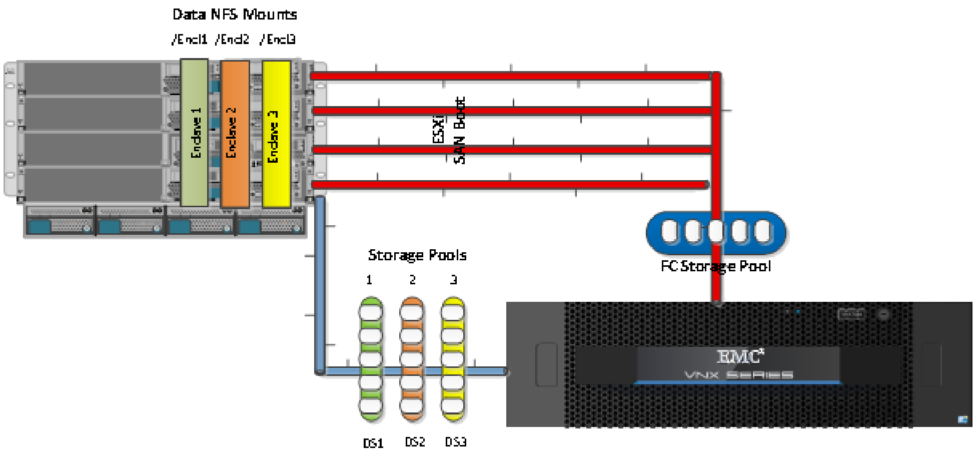

The storage model followed is one that corresponds to the NFS variant in the referenced VSPEX CVD. As mentioned, the difference is that enclave common data storage on each NFS mount is carved out of a separate storage pool. The procedure to setup storage pools for both FC based boot volumes and File based data is the same as in the aforementioned base VSPEX platform CVD. System access controls at the time of creating NFS exports on VNX via Unisphere should list IPs of all target ESXi hosts for the "Root Hosts" and "Access Hosts" fields to allow complete access since each enclave spans all ESXi hosts within the cluster.

Security on the VNX

The EMC VNX 5400 storage array provides several layers of security including at the user access and logging and auditing levels. A Virtual Data Mover (VDM), which is a logical network abstraction on top of physical Data Movers, provides for additional network end-points to facilitate IP based separation for NFS mounts.

Highly Secure Storage Management Network Topology

This configuration, provides a very high level of security for a company's storage systems. Potential threats are reduced to a breach of physical resources.In addition, enabling IP filtering for the VNX domain limits the management of the storage systems to a single Windows server, namely the Unisphere Client/Server management station. IP filtering allows each storage system or domain to have a list of trusted client IP addresses. The storage system will accept management connections only from trusted clients. IP filtering does not affect other traffic, such as Event Monitor polls, email notifications, or SNMP. The IP filtering configuration can be found through http://<SP IP address>/setup pages or through the naviseccli security-trusted client switch.

Figure 29 Highly Secure Storage Management Network Topology

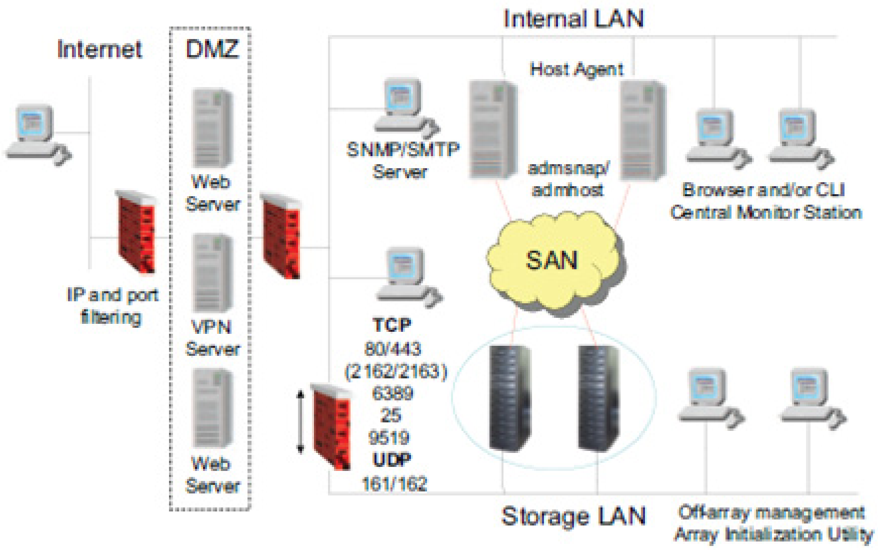

This configuration provides two layers of authentication. First, the user must have valid Windows credentials to log in to the management station. Second, the user must have valid Unisphere credentials to manage the storage system. The trade-off with this configuration is the loss of flexibility in terms of management options. Neither the ability to manage from anywhere in the enterprise nor the ability to centrally monitor the entire network is available. Also, remote support (call-home) of the storage system by using the ESRS IP Client is not possible in this environment. Note that ESRS IP Client can still send notifications to EMC Customer Service. The Unisphere architecture is very flexible in its ability to integrate into several secure environments.

Encryption

The storage management server provides 256-bit (128-bit is also supported) symmetric encryption of all data passed between it and the client components that communicate with it, as listed in Ports used by Unisphere components on VNX for block on page 38 (Web browser, Secure CLI), as well as all data passed between storage management servers. The encryption is provided using SSL/TLS and uses the RSA encryption algorithm, which provides the same level of cryptographic strength as is employed in e-commerce. Encryption protects the transferred data from prying eyes-whether on the local LANs behind the corporate firewalls, or if the storage systems are being remotely managed over the Internet.

Defense in Depth

Because the behavior of the vast majority of the open network ports on VNX for file is governed by network standards, there are no additional steps available for VNX for file to protect these ports other than disabling their associated services and closing the ports. Disabling services such as portmap will hinder the general operations of VNX for file, and in some cases, the impact will be severe. However, the notion of defense in depth dictates that any potential vulnerability is addressed with additional protections to control who may access the ports. This may be done with firewalls in the network environment (external to VNX for file) or by enabling the IP tables functionality on the Control Station.In addition, the VNX for file Data Mover provides two powerful mechanisms for controlling network connectivity:

- Packet Reflect
- Virtual local area networks (VLANs)

Packet Reflect makes sure that outbound (reply) packets always exit through the same interfaces through which the inbound (request) packets entered. Because majority of the network traffic on a Data Mover, including all file system I/O, is initiated by the client, the Data Mover uses Packet Reflect to reply to client requests. With Packet Reflect, there is no need to determine the route to send the reply packets. Packet Reflect is enabled by default.

Communication Security

VLANs are logical networks that function independently of the physical network configuration. For example, VLANs enable you to put all of a department's computers on the same logical subnet, which can increase security and reduce network broadcast traffic.Configuring and Managing Networking on VNX provides additional information about Packet Reflect and VLANs as well as how to configure these features.

Please see "Security Configuration Guide for VNX P/N 300-015-128 Rev 01 and P/N 300-013-510 Rev 03 for more details on security features on the VNX Series of arrays.

NFS Security Settings

Although generally regarded as a vulnerable file-sharing protocol, you can make NFS more secure by using the following configuration settings:

- Defining read-only access for some (or all) hosts
- Limiting root access to specific systems or subnets
- Hiding export and mount information if a client does not have mount permissions for the file system corresponding to that entry

In addition, if strong authentication is required, you can configure Secure NFS, which uses Kerberos. Configuring NFS on VNX describes how to configure these settings. All NFS exports are displayed by default. To hide NFS exports, you must change the value of the forceFullShowmount for mount facility parameter.

The deployment details provide example configurations necessary to achieve enclave functionality.

Cisco Nexus 5000

The Cisco Nexus 5000 Series (http://www.cisco.com/en/US/products/ps9670/index.html), part of the Cisco Nexus Family of data center class switches, delivers an innovative architecture that simplifies data center transformation. These switches deliver high performance, standards-based Ethernet and FCoE that enables the consolidation of LAN, SAN, and cluster network environments onto a single Unified Fabric. Backed by a broad group of industry-leading complementary technology vendors, the Cisco Nexus 5000 Series is designed to meet the challenges of next-generation data centers, including dense multisocket, multicore, virtual machine-optimized deployments, where infrastructure sprawl and increasingly demanding workloads are commonplace.

The Cisco Nexus 5000 Series is built around two custom components: a unified crossbar fabric and a unified port controller application-specific integrated circuit (ASIC). Each Cisco Nexus 5000 Series Switch contains a single unified crossbar fabric ASIC and multiple unified port controllers to support fixed ports and expansion modules within the switch.

The unified port controller provides an interface between the unified crossbar fabric ASIC and the network media adapter and makes forwarding decisions for Ethernet, Fibre Channel, and FCoE frames. The ASIC supports the overall cut-through design of the switch by transmitting packets to the unified crossbar fabric before the entire payload has been received. The unified crossbar fabric ASIC is a single-stage, nonblocking crossbar fabric capable of meshing all ports at wire speed. The unified crossbar fabric offers superior performance by implementing QoS-aware scheduling for unicast and multicast traffic. Moreover, the tight integration of the unified crossbar fabric with the unified port controllers helps ensure low latency lossless fabric for ingress interfaces requesting access to egress interfaces.

ISE Integration

The Identity Services Engine provisioned assumes the following personas:

- Administration Node
- Policy Service Node
- Monitoring Node

The ISE provides RADIUS services to the Nexus 5548UP switches. Each switch is configured as Network as follows:

```
radius-server key 7 "K1kmN0gy"
radius distribute
radius-server host 10.29.133.31 key 7 "K1kmN0gy" authentication accounting
radius commit
aaa group server radius aaa-private-sg
    server 10.29.133.31
    use-vrf management
    source-interface mgmt0
ip radius source-interface mgmt0
aaa authentication login default group aaa-private-sg
aaa authentication dot1x default group aaa-private-sg
aaa accounting dot1x default group aaa-private-sg
aaa authorization cts default group aaa-private-sg
aaa accounting default group aaa-private-sg
no aaa user default-role
```

Cisco TrustSec

Cisco TrustSec enables an access-control framework based on contextual information - who, what, where, when and how. This role-based approach is independent of topology and based on privileges accorded to Secure Group Tags (SGT) on the source and destination devices rather than on network IP addresses. The Cisco Nexus 5500 platform supports TrustSec but cannot act as an SXP "listener". This means it cannot aggregate and advertise via SXP the IP to SGT mappings learned from the Nexus 1000v. In light of this, the Nexus 1000v will implement an SXP connection to each ASA virtual context directly to advertise the CTS tag to IP information.

Private VLANs

The use of private VLANs allows for the complete isolation of control and management traffic within an Enclave. The Cisco Nexus 5548UP supports private VLANs and used the following structure during validation. In this sample, VLAN 3171 is the primary VLAN and 3172 is an isolated VLAN carried across the infrastructure.

```
vlan 3171
  name core-services-primary
  private-vlan primary
  private-vlan association 3172
vlan 3172
  name core-services-isolated
  private-vlan isolated
```

Port Profiles

A Port profile is a mechanism for simplified configuration of interfaces. A port profile can be assigned to multiple interfaces giving them all the same configuration. Changes to the port profile can be propagated automatically to the configuration of any interface assigned to it. Port profiles that are configured as uplinks, can be assigned by the server administrator to physical ports (vmnic or a pnic in VMware). Port profiles that are not configured as uplinks can be assigned to a VM virtual port.

```
interface port-channel1
  description VPC-Peerlink
  switchport mode trunk
  switchport trunk native vlan 131
  switchport trunk allowed vlan
1,10-41,98-99,131,133,201-219,666,2001-2019,3001-3019,3170-3173,3175-3179,32
50-3251,3253-3255
  spanning-tree port type network
  speed 10000
  vpc peer-link

interface port-channel18
  description to FI-A
  switchport mode trunk
  switchport trunk native vlan 131
  switchport trunk allowed vlan
1-2,10-41,98-99,131,133,201-219,666,2001-2019,3001-3019,3170-3173,3175-3179,
3250-3251,3253-3255
  spanning-tree port type edge trunk
  speed 10000
  vpc 18

interface port-channel19
  description to FI-B
  switchport mode trunk
  switchport trunk native vlan 131
  switchport trunk allowed vlan
1-2,10-41,98-99,131,133,201-219,666,2001-2019,3001-3019,3170-3173,3175-3179,
3250-3251,3253-3255
  spanning-tree port type edge trunk
  speed 10000
  vpc 19

interface port-channel20
  description <<** Po20 tp sea-asa1 Po1 **>>
  switchport mode trunk
  switchport trunk allowed vlan 200,666,2001-2135,3001-3135
  spanning-tree port type normal
  speed 10000
  vpc 20

interface port-channel21
  description <<** Po21 tp sea-asa2 Po1 **>>
  switchport mode trunk
  switchport trunk allowed vlan 200,666,2001-2135,3001-3135
  spanning-tree port type normal
  speed 10000
  vpc 21

interface port-channel31
  description to VNX5400-DM2
  untagged cos 5
  switchport access vlan 10
  switchport trunk native vlan 10
  speed 10000
  vpc 31

interface port-channel32
  description to VNX5400-DM3
  untagged cos 5
  switchport access vlan 10
  switchport trunk native vlan 10
  speed 10000
  vpc 32
```

Monitoring and SPAN

The ability to monitor network traffic within the Nexus platform is key to ensure the efficient operation of the solution. The design calls for the use of Switched Port Analyzer (SPAN) as well as NetFlow services to provide visibility.

Switched Port Analyzer (SPAN) sources refer to the interface from which traffic can be monitored. SPAN sources send a copy of the traffic to a destination port. The network analyzer, which is attached with destination port, analyzes the traffic that passes through source port.

```
interface Ethernet1/17
  switchport mode trunk
  switchport trunk native vlan 131
  switchport trunk allowed vlan
1-2,10-41,98-99,131,133,201-219,666,2001-2019,30
01-3019,3170-3173,3175-3179,3250-3251,3253-3255
  channel-group 18 mode active
```

```
interface Ethernet1/18
  switchport mode trunk
  switchport trunk native vlan 131
  switchport trunk allowed vlan 
1-2,10-41,98-99,131,133,201-219,666,2001-2019,30
01-3019,3170-3173,3175-3179,3250-3251,3253-3255
  channel-group 19 mode active

monitor session 1
  description <<** SPAN of Po20 to NGAs **>>
  source interface port-channel20 rx
  destination interface Ethernet1/27
```

NetFlow

NetFlow technology efficiently provides accounting for various applications such as network traffic accounting, usage-based network billing, network planning, as well as Denial Services monitoring capabilities, network monitoring, outbound marketing, and data mining capabilities for both Service Provider and Enterprise organizations. The NetFlow architecture consists of flow records, flow exports and flow monitors. NetFlow consumes hardware resources such as TCAM and CPU in the switching environment. It is also not a recommended practice to use NetFlow sampling as this provides an incomplete view of network traffic.

To avoid NetFlow resource utilization in the Nexus switch and potential "blind spots" the NetFlow service is offloaded to dedicated devices, namely the Cisco NetFlow Generation Appliances (NGA). The NGAs consume SPAN traffic from the Nexus 5548UP. Please see the Cisco NetFlow Generation Appliance section for details on its implementation in the design.

Cisco Nexus 1000V

The Cisco Nexus 1000V switch is a software switch on a server that delivers Cisco VN-Link services to virtual machines hosted on that server. It takes advantage of the VMware vSphere framework to offer tight integration between server and network environments and help ensure consistent, policy-based network capabilities to all servers in the data center. It allows policy to move with a virtual machine during live migration, ensuring persistent network, security, and storage compliance, resulting in improved business continuance, performance management, and security compliance. Last but not least, it aligns management of the operational environment for virtual machines and physical server connectivity in the data center, reducing the total cost of ownership (TCO) by providing operational consistency and visibility throughout the network. It offers flexible collaboration between the server, network, security, and storage teams while supporting various organizational boundaries and individual team autonomy.

Architecture Overview

In addition to the security features offered by the Cisco Nexus 1000V, the Cisco virtual distributed switch supports VMware's vShield Edge technology. To achieve Enhanced Secure Multi-Tenancy, it is important to carve off one or more isolated Layer 2 adjacent segments on the Cisco Nexus 1000V for each tenant. These VLAN segments are further secured at the perimeter via vShield Edge which allows certain centralized services such as DNS or AD to be readily consumed by tenant virtual machines within the data center.

For more information, see: http://www.cisco.com/en/US/products/ps9902/index.html. The Cisco Nexus 1010 Virtual Services Appliance hosts the Cisco Nexus 1000V Virtual Supervisor Module (VSM) and supports the Cisco Nexus 1000V Network Analysis Module (NAM) Virtual Service Blade to provide a comprehensive solution for virtual access switching. The Cisco Nexus 1010 provides dedicated hardware for the VSM, making the virtual access switch deployment much easier for the network administrator.

For more information on Nexus 1000V, see: http://www.cisco.com/en/US/partner/products/ps10785/index.html.

For more information on Cisco VN-Link technologies see: http://www.cisco.com/en/US/netsol/ns894/index.html.

The following section describes the implementation of the Cisco Nexus 1000v VSM and VEMs in the enclave architecture.

Note You will need a Cisco Nexus 1000v Advanced Edition license for the TrustSec feature (show switch edition). After installing the necessary license as follows, please enable the "cts" function:

1. Download advanced edition license to bootflash: directory on the VSMs
   a. Copy ftp://<user>@<server>/<lic_file>.lic bootflash:
   b. Enter vrf ( ): <hit return key>
   c. Password: <password>
2. Enter configuration mode (config t)
   a. Svs switch edition advanced
3. Confirm update to editionshow switch edition
4. Enable cts feature
   a. config t
   b. Feature cts
5. Confirm 'cts' feature is enabled by running 'show feature'.

SVS Domain

```
interface mgmt0
  ip address 10.29.133.37/24

interface control0
  ip address 10.29.131.42/24

SVS domain config:
  Domain id:    100
  Control vlan:   NA
  Packet vlan:    NA
  L2/L3 Control mode: L3
  L3 control interface: control0
  Status: Config push to VC successful.
  Control type multicast: No

svs connection vcenter:
    ip address: 10.29.133.31
    remote port: 80
    protocol: vmware-vim https
    certificate: default
    datacenter name: SEA_DC
    admin:
    max-ports: 8192
    DVS uuid: de 03 2c 50 bf 8c f8 bb-b6 72 e7 f3 06 be f1 b5
    config status: Enabled
    operational status: Connected
    sync status: Complete
    version: VMware vCenter Server 5.1.0 build-799731
    vc-uuid: 6080265F-6ACA-4B05-8608-468A3F1342E9
```

A Cisco Nexus 1000v DVS (sea-prod-vsm) is created with a unique SVS domain to support the new production enclave environment. This new virtual distributed switch will associate with the baseline VSPEX VMware vCenter Server. The control0 interface on a unique VLAN is used to provide ESXi host isolation from the remaining management network. All VEM to VSM communication will occur over this dedicated VLAN. The `svs mode L3 interface control0` command assigns communication between the VSM and VEM across the control interface.

ISE Integration

The ISE provides RADIUS services to each of the Nexus 1000v VSM which are configured as network devices in the ISE tool.

```
radius-server key 7 "K1kmN0gy"
radius-server host 10.29.133.39 key 7 "K1kmN0gy" authentication accounting
aaa group server radius ise-radius-grp
    server 10.29.133.37
    use-vrf management
    source-interface mgmt0

ip radius source-interface mgmt0
```

The following AAA commands were used:

```
aaa authentication login default group ise-radius-grp
aaa authorization cts default group ise-radius-grp
aaa accounting default group ise-radius-grp
no aaa user default-role
```

Virtual Extensible LAN

Virtual Extensible LAN (VXLAN) allows organizations to scale beyond the 4000 VLAN limit present in traditional switching environments by encapsulating frames MAC frames in IP. This approach allows a single overlay VLAN to support multiple VXLAN segments, simultaneously addressing VLAN scale issues and network segmentation requirements.

In the enclave architecture, the use of VXLAN is enabled via the segmentation feature and the Unicast-only mode was validated. Unicast-only mode distributes a list of IP addresses associated with a particular VXLAN to all Cisco Nexus 1000v VEM. Each VEM requires at least one IP/MAC address pair to terminate VXLAN packets. This IP/MAC address pair is known as the VXLAN Tunnel End Point

(VTEP) IP/MAC addresses. The distribution MAC feature enables the VSM to distribute a list of MAC to VTEP associations. he combination of these two features eliminates unicast flooding as all MAC addresses are known to all VEMs under the same VSM.

- Feature segmentation
- Segment mode Unicast-only
- Segment distribution MAC

The IP/MAC address that the VTEP uses is configured when you enter the capability vxlan command. You can have a maximum of four VTEPs in a single VEM. The production Cisco Nexus 1000v uses VLAN 3253 to support VXLAN traffic. The Ethernet uplink port-profile supporting traffic originating from the enclaves will support the VXLAN VLAN. Notice the MTU of the uplink is large enough to accommodate the additional VXLAN encapsulation header of 50 bytes.

```
port-profile type ethernet enclave-data-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk allowed vlan 10,2001-2003,3001-3003,3254-3255
  switchport trunk native vlan 10
  mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 10,2001-2003,3001-3003,3254-3255
  state enabled
```

The VXLAN vethernet port profile uses that capability VXLAN to enable the VXLAN functionality on the VMKNIC on the Cisco Nexus 1000v VEM.

```
port-profile type vethernet VXLAN-VTEP
  vmware port-group
  switchport mode access
  switchport access vlan 3253
  capability vxlan
  no shutdown
  state enabled
```

To create VXLAN segments IDs or domains it is necessary to construct bridge domains in the Nexus 1000v configuration. The bridge domains are referenced by Virtual Machine port profiles requiring VXLAN services. As the naming standard dictates there are two VXLAN segments for each of the enclaves. The segment ID is assigned by the administrator. The current version of the Nexus 1000v supports up to 2048 VXLAN bridge domains.

```
bridge-domain bd-enclave-1
  segment id 30011
bridge-domain bd-enclave-2
  segment id 30021
bridge-domain bd-enclave-1-2
  segment id 30012
```

```
bridge-domain bd-enclave-2-2
  segment id 30022
```

The Cisco Nexus 1000v VXLAN enabled port profile referencing the previously defined bridged domains. The following is an example of the port group availability in VMware vCenter:

```
port-profile type vethernet enc1-vxlan1
  vmware port-group
  inherit port-profile enc-base
  switchport access bridge-domain bd-enclave-1
  state enabled
port-profile type vethernet enc2-vxlan1
  vmware port-group
  inherit port-profile enc-base
  switchport access bridge-domain bd-enclave-2
  state enabled
port-profile type vethernet enc1-vxlan2
  vmware port-group
  inherit port-profile enc-base
  switchport access bridge-domain bd-enclave-1-2
  state enabled
port-profile type vethernet enc2-vxlan2
  vmware port-group
  inherit port-profile enc-base
  switchport access bridge-domain bd-enclave-2-2
  state enabled
```

Visibility

The following Cisco Nexus 1000v features were enabled to provide virtual access visibility, awareness and to support cyber threat defense technologies.

SPAN

The Cisco Nexus 1000v supports the mirroring of traffic within the virtual distributed switch as well as externally to third party network analysis devices or probes. Each of these capabilities has been implemented with the Secure Enclave architecture to advance understanding of traffic patterns and performance of the environment.

NetFlow

The Nexus 1000v supports NetFlow. The data may be exported to the Lancope StealthWatch system for analysis. As shown below the NetFlow feature is enabled. The destination of the flow records is defined as "sjc-export-1" which is the Lancope Cyber Thread Defense (CTD) solution. The flow record "SJC-sea-enclaves" defines the interesting parameters to be captured with each flow and indicates the "sjc-export-1" as the collector.

```
feature netflow


flow exporter nf-export-1
  description <<** SEA Lancope Flow Collector  **>>
  destination 172.26.164.240 use-vrf management
  transport udp 2055
  source mgmt0
  version 9
    option exporter-stats timeout 300
    option interface-table timeout 300
flow monitor sea-enclaves
  record netflow-original
  exporter nf-export-1
  timeout inactive 15
  timeout active 60
```

The NetFlow monitor definition is applied to the port profile. In this example, the port-profile associated with a data base will be the subject of NetFlow exports.

```
port-profile type vethernet encl1-db
  vmware port-group
  inherit port-profile enc-base
  switchport mode access
  switchport access vlan 3001
  cts sgt 4
  ip flow monitor SJC-sea-enclaves input
  vservice node encl1-vsg profile encl1_db
  org root/Enclave1
  no shutdown
  description <<** SJC: Enclave 1 Data DB**>>
  state enabled
```

The validated version of the Cisco Nexus 1000v supports up to 32 NetFlow monitors and 256 instances. An instance being the application of the monitor to a port-profile. If resource availability is a concern, it is suggested that the monitoring focus on data sources such as data base profiles and critical enclaves within the architecture.

For more information on the Cyber Threat Defense system implemented for the Secure Enclave architecture please visit the Cisco Cyber Threat Defense for the Data Center Solution: Cisco Validated Design at
www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-security/ctd-first-look-design-guide.pdf

TrustSec

The Cisco Nexus 1000v supports the Cisco TrustSec architecture by implementing the SGT Exchange Protocol (SXP). The SXP protocol is used to propagate the IP addresses of virtual machines and their corresponding SGTs up to the upstream Cisco TrustSec-capable switches or Cisco ASA firewalls. The SXP protocol is a secure communication between the speaker (Nexus 1000v) and listener devices.

The following configuration describes the enabling of the CTS feature on the Nexus 1000v. The feature is enabled with device tracking. CTS device tracking allows the switch to capture the IP address and associated SGT assigned at the port profile of the virtual machine.

```
feature cts
cts device tracking
```

The SXP configuration can be optimized by configuring a default password and source IP address associated with any SXP connection. The SXP connection definition in this example points to the Nexus 7000 switches that are configured as listeners. In a VSPEX configuration when the Nexus 7000 switches are involved, it is recommended to use these as SXP listeners. The Cisco Nexus 7000 switches will act as a CTS IP-to-SGT aggregation point and can be configured to transmit (speak) the CTS mapping information to other CTS infrastructure devices such as the Cisco ASA.The Cisco Nexus 5000 and the Cisco Nexus 1000V currently act only as "speakers" and send the SGT and IP address mapping via SXP. A network element capable of "listening" such as the Cisco Nexus 7000 can then sends the SGT-to IP address map to the ASA 5585-X, again through SXP:

```
cts sxp enable
cts sxp default password 7 K1kmN0gy
cts sxp default source-ip 10.29.133.37
cts sxp connection peer 10.29.133.48 password default mode listener vrf
management
cts sxp connection peer 10.29.133.49 password default mode listener vrf
management
```

Switches such as the Cisco Nexus 5000 do not support the SXP listener role. In this scenario, the Cisco Nexus 1000v will "speak" directly to each ASA virtual context providing SGT to IP mapping information for use in the access control service policies.

Figure 30 *Cisco Nexus 1000v TrustSec SXP Example*

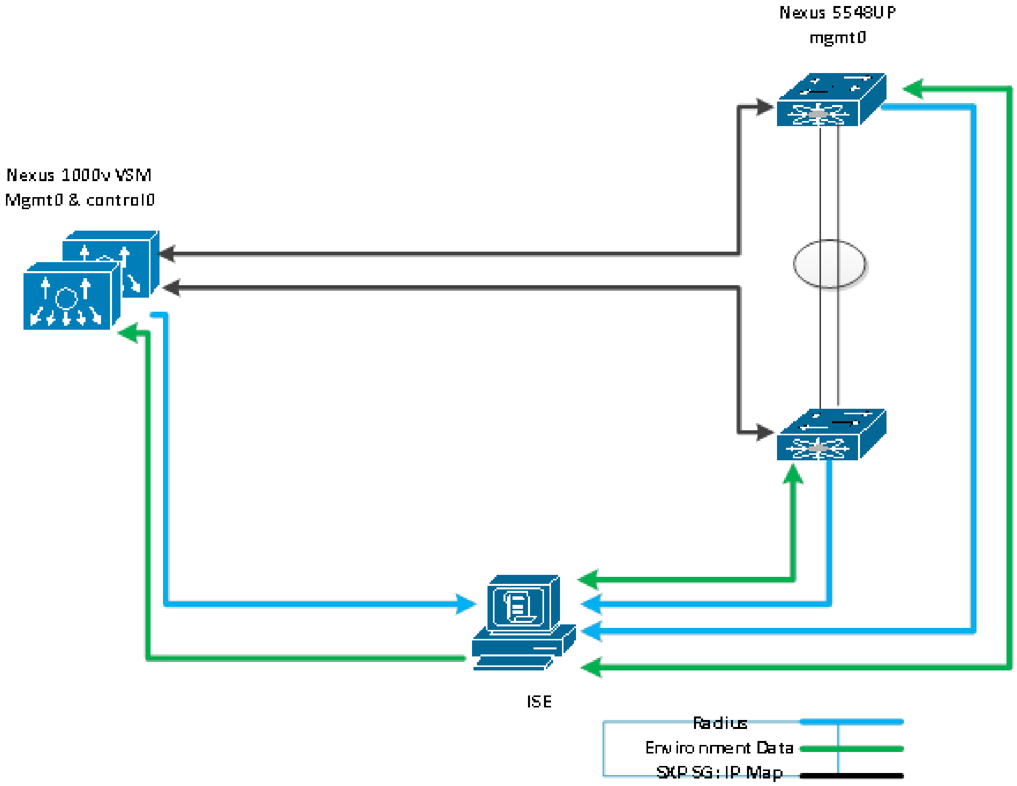

Private VLANs

The private VLAN configuration on the Nexus 1000v supports the isolation of enclave management traffic. This configuration requires enabling the feature and definition of two VLANs. In this example, VLAN 3172 is the primary VLAN supporting the isolated VLAN 3172.

```
feature private-vlan
vlan 3171
  name core-services-primary
  private-vlan primary
  private-vlan association 3172
vlan 3172
  name core-services-isolated
  private-vlan isolated
```

The private VLAN construct is then applied to a vethernet port profile. The sample below indicates the use of the private VLAN for core services traffic. Traffic such as Active Directory, DNS, Windows Update Services. It is important to remember that virtual machines connected to an isolated private VLAN cannot communicate with other VMs on the same segment.

```
port-profile type vethernet pvlan_core_services
  vmware port-group
  switchport mode private-vlan host
  switchport private-vlan host-association 3171 3172
  service-policy type qos input Platinum
  no shutdown
  state enabled
```

The Cisco Nexus 1000v management VSM defines a promiscuous port profile allowing isolated traffic on the production VSM to communicate with virtual machines using the core_services profile.

```
port-profile type vethernet core_services
  vmware port-group
  switchport mode private-vlan promiscuous
  switchport access vlan 3171
  switchport private-vlan mapping 3171 3172
  ip flow monitor sea-enclaves input
  no shutdown
  state enabled
```

Port Profiles

The Cisco Nexus 1000v production VSM uses three unique Ethernet type port profiles for uplink transport. This is accomplished by defining six VNICs on the ESXi UCS service profile. The VNICs are deployed in parallel offering connectivity to either Cisco UCS Fabric A or B. The Cisco Nexus 1000v VEM provides host based port aggregation of these VNICs creating port channels. The segmentation and availability of the enclave is enhanced by using dedicated vNICs with the HA features of Cisco Nexus 1000v port channeling. The system-uplink port profile supports all of the VLANs required for control and management services. The MTU is set to 9000 requiring jumbo enforcement at the edge and enabling across the infrastructure.

```
port-profile type ethernet system-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 131,133,3254-3255
  system mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 131
  state enabled
```

The enclave port profile uplinks support traffic directly associated with the enclaves. This includes NFS, iSCSI and enclave data flows. It is important to understand that these VLANs to capture the limits of the environment.

```
port-profile type ethernet enclave-data-uplink
  vmware port-group
  switchport mode trunk
  switchport trunk native vlan 10
  system mtu 9000
  switchport trunk allowed vlan 10,2001-2003,3001-3003,3254-3255
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 10,2001-2003,3001-3003
  state enabled
```

The core-uplinks port profile supports the private VLANs, primary and isolated, that offer complete isolation of management traffic to all enclaves in the architecture. The port-channel created in the design is dedicated to only these two VLANs.

The show port-channel summary command for a single VEM module, ESXi host, captures the three port channel uplinks created.

N1KV(config)# show port-channel summary | in Eth4

```
2     Po2(SU)     Eth      NONE      Eth4/1(P)    Eth4/2(P)
3     Po3(SU)     Eth      NONE      Eth4/3(P)    Eth4/4(P)
4     Po4(SU)     Eth      NONE      Eth4/5(P)    Eth4/6(P)
```

Virtual Service Integration (Virtual Security Gateway)

Integration of virtual services into the Cisco Nexus 1000v environment requires that the switch register with the Cisco Prime Network Services Controller (PNSC). The registration process requires the presence of a policy-agent file, the PNSC IP address and a shared secret for secure communication between the VSM and controller. The following sample details the policy-agent configuration in the enclave environment.

```
vnm-policy-agent
  registration-ip 10.29.131.45
  shared-secret **********
  policy-agent-image bootflash:/vnmc-vsmpa.2.0.0.38.bin
  log-level crit
```

The Cisco Nexus 1000v allows for the global definition of vservice specific attributes that can be inherited by the instantiated services. The global VSG qualities are defined below. The bypass asa-traffic command indicates that traffic will bypass an ASA Cisco Nexus 1000v. The ASA Cisco Nexus 1000v is not part of this design, this command is unnecessary.

```
vservice global type vsg
  tcp state-checks invalid-ack
  tcp state-checks seq-past-window
  no tcp state-checks window-variation
! This refers to the ASA Nexus 1000v platform which is not in this design
  bypass asa-traffic
```

The instantiation of a vservice in the Nexus 1000v requires the network administrator to define the service node, and bind the security profile to the port-profile. In this example, the VSG service node is named enc1-vsg. The vPath communication will occur at Layer 2 between the VEM and vPath interface of the VSG. The IP address of the VSG is resolved via ARP and data (vPath) traverses VLAN 3254. In this example, if the VSG should fail the traffic will be not be permitted to flow.

```
vservice node encl1-vsg type vsg
  ip address 111.111.111.111
  adjacency l2 vlan 3254
  fail-mode close
```

vPath is an encapsulation technique that will add 74 bytes when used in L2 mode or 94 bytes if using L3 mode. To avoid fragmentation in a Layer 2 implementation, ensure the outgoing uplinks support the required MTU. If it is a Layer 3 enabled vPath packets will be dropped and ICMP error messages sent to the traffic source.

The port profile enc1-web uses the previously described service node. The **vservice** command binds a specific Cisco VSG (enc1-vsg) and security profile (enc1_web) to the port profile. This enables vPath to redirect the traffic to the Cisco VSG. The **org** command defines the tenant with the PNSC where the firewall is enabled.

```
port-profile type vethernet encl1-web
  vservice node encl1-vsg profile encl1_web
  org root/Enclave1
  no shutdown
  description <<** Enclave 1 Data WEB **>>
  state enabled
```

vTracker

The vTracker feature on the Cisco Nexus 1000V switch provides information about the virtual network environment. vTracker provides various views that are based on the data sourced from the vCenter, the Cisco Discovery Protocol (CDP), and other related systems connected with the virtual switch. vTracker enhances troubleshooting, monitoring, and system maintenance. Using vTracker show commands, you can access consolidated network information across the following views:

- Module-View—VTracker showing information about a server module
- upstream-View—VTracker information showing from upstream switch
- Vlan-View—VTracker showing information vlan usage by Virtual machines

- Vm-View—VTracker showing information about a virtual machine
- VMotion-View—VTracker showing information about VM migration

For example, the show vtracker module-view provides visibility into the ESXi pNICS defined as vNICS on the UCS system:

```
N1KV# show vtracker module-view pnic


--------------------------------------------------------------------------------
Mod  EthIf    Adapter  Mac-Address    Driver    DriverVer         FwVer
           Description
--------------------------------------------------------------------------------
3    Eth3/1    vmnic0   0050.5652.0a04 enic     2.1.2.22          2.1(3a)
           Cisco Systems Inc Cisco VIC Ethernet NIC

3    Eth3/2    vmnic1   0050.5652.0b04 enic     2.1.2.22          2.1(3a)
           Cisco Systems Inc Cisco VIC Ethernet NIC

3    Eth3/3    vmnic2   0050.5652.5a04 enic     2.1.2.22          2.1(3a)
           Cisco Systems Inc Cisco VIC Ethernet NIC

3    Eth3/4    vmnic3   0050.5652.5b04 enic     2.1.2.22          2.1(3a)
           Cisco Systems Inc Cisco VIC Ethernet NIC

3    Eth3/5    vmnic4   0050.5652.3a04 enic     2.1.2.22          2.1(3a)
           Cisco Systems Inc Cisco VIC Ethernet NIC

3    Eth3/6    vmnic5   0050.5652.3b04 enic     2.1.2.22          2.1(3a)
           Cisco Systems Inc Cisco VIC Ethernet NIC

4    Eth4/1    vmnic0   0050.5652.0a05 enic     2.1.2.22          2.1(3a)
           Cisco Systems Inc Cisco VIC Ethernet NIC
```

Security Package

Data Center wide infrastructure and security related services external to the converged infrastructure include Active Directory (AD), Dynamic Host Configuration Protocol (DHCP), Domain Name Services (DNS), VMware vCenter Server, Cisco UCS Director Server, Cisco UCS Director Bare-Metal Server, Cisco Nexus 1000V Virtual Supervisor Modules (VSM), Cisco Prime Network Services Controller (PNSC), Cisco Identity Services Engine (ISE) and Lancope StealthWatch.

Deployment details of enterprise services specific to this solution such as VSG with PNSC, ISE with TrustSec and ASA, NGA 3240 devices with NetFlow and Lancope StealthWatch and UCS Director is discussed below.

Prerequisites (same as for base platform):

- Working Active Directory, DNS and DHCP infrastructure components.
- An ftp server reachable by UCS Director server for patching.
- Routable networking for resources accessed by end users such as VMs.
- VMware vCenter 5.1 update 1
- SMTP email server for status emails.

Cisco Virtual Security Gateway

The Cisco Virtual Security Gateway is a virtual appliance that works with the Cisco Nexus 1000v to consistently enforce security policies in virtualized environments. The Cisco VSG operates at Layer 2 creating zones based segmentation. he Enclave architecture uses the VSG to secure "east-west" traffic patterns.

Figure 31 describes the flow and how segregation of duties and ownership is maintained for provisioning the Virtual Security Gateway. The security, network and server administrators each have a role in the process. his section of the document will focus on the security administrator role as the network Nexus 1000v configuration is covered in Virtual Service Integration (Virtual Security Gateway) section and the assignment of a port group to a virtual machines is a well-known operation.

Figure 31 Cisco VSG Deployment Process

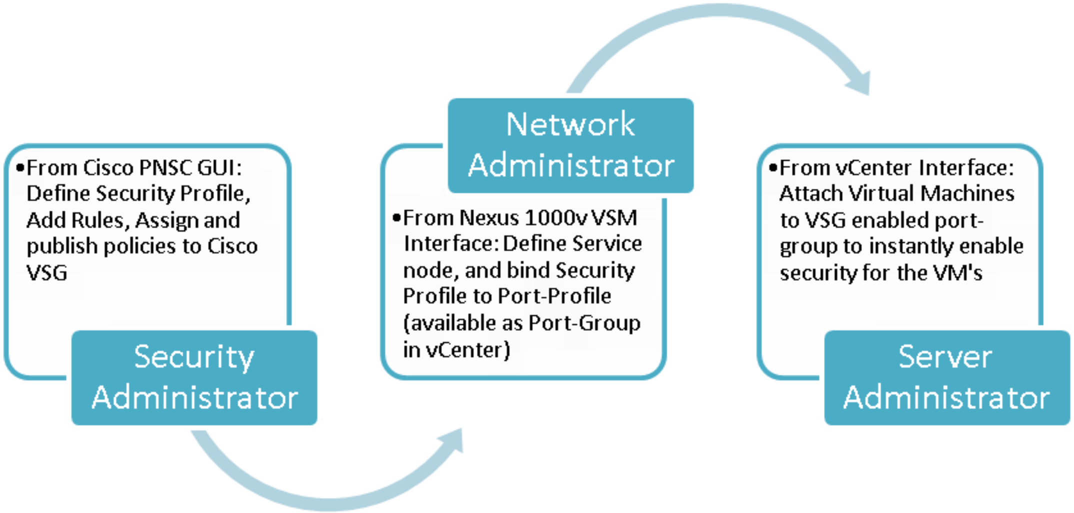
•From Cisco PNSC GUI: Define Security Profile, Add Rules, Assign and publish policies to Cisco VSG
Security Administrator
Network Administrator
•From Nexus 1000v VSM Interface: Define Service node, and bind Security Profile to Port-Profile (available as Port-Group in vCenter)
•From vCenter Interface: Attach Virtual Machines to VSG enabled port-group to instantly enable security for the VM's
Server Administrator

It is not recommended to use VMware High Availability (HA) or fault-tolerance with the Cisco VSG. It is recommended to use an HA pair of VSGs and VMware DRS groups as described in the DRS for Virtual Service Nodes section of this document. In situations where neither the primary nor the standby Cisco VSG is available to vPath, configure the failure mode as Fail Open or Fail Close as dictated by the security requirements of the Enclave.

Figure 32 Virtual Gateway (VSG) and Prime Network Services Controller (PNSC)
NEED GRAPHIC

PNSC Install and Configuration

http://www.cisco.com/c/en/us/td/docs/net_mgmt/virtual_network_mgmt_center/3-0/quick-start-guide/b_30_Quick_Start_Guide.pdf

Cisco TrustSec

Cisco TrustSec provides an access control solution that builds upon an existing identity-aware infrastructure to ensure data confidentiality between network devices and to integrate security access services on one platform Cisco TrustSec SGTs allow you to map user roles to server roles in the data center to ensure that proper access is granted where it is needed and denied where it is not needed.

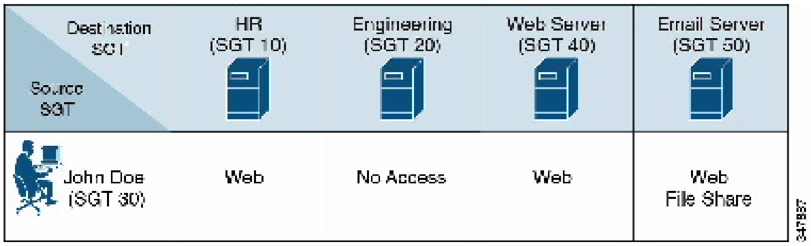

Destination SGT Source SGT	HR (SGT 10)	Engineering (SGT 20)	Web Server (SGT 40)	Email Server (SGT 50)
John Doe (SGT 30)	Web	No Access	Web	Web File Share

The availability and propagation of this information enables security solutions across networks at the access, distribution, and core layers of the network. Cisco TrustSec provides an additional capability by providing an alternate method for deploying secure separation. Using the TrustSec SGTs and the advance policy capability, you can also leverage TrustSec at the data center virtualization layer to enable separation for your secure containers. Further details and comprehensive information about and deploying TrustSec Solutions can be found at http://www.cisco.com/go/trustsec

Cisco Adaptive Security Appliance

Transparent Firewall Mode

A security context can be operated in transparent mode, which acts like a Layer 2 firewall that appears to be a "bump in the wire" or a "stealth firewall", and is not seen as a router hop to connected devices. The ASA connects to the same network between its interfaces. Because the firewall is not a routed hop in this mode, it can be introduced into an existing network. The Management (VLAN 133) and Data interfaces are not connected to the same switch because the Data interface receives the MAC address table updates with a minimum of a 30-second delay for security reasons. At least one bridge group is required per context, but each context can support up to eight bridge groups. Each bridge group can include up to four interfaces. The Intrusion Prevention System (IPS) is not installed or tested in this setup. Please refer to basic configuration of transparent mode in the following document:

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/SDC/DG/SDC_DesignGuide/SDC_DG_2013-11-25_v10.html#pgfId-405501

The Cisco ASA 5585-X is a high-performance, 2-slot chassis, with the firewall Security Services Processor (SSP) occupying the bottom slot, and the IPS Security Services Processor (IPS SSP) in the top slot of the chassis. The ASA includes many advanced features, such as multiple security contexts, clustering, transparent (Layer 2) firewall or routed (Layer 3) firewall operation, advanced inspection engines, and many more features. The VSPEX Data Center readily supports the ASA platform to provide security services and the enclave design.

It should be noted that the Secure Enclave validation effort has resulted in a number of Cisco Validated Designs that detail the security implementation of the Cisco ASA platforms. The Design Zone for Secure Data Center Portfolio page: http://www.cisco.com/c/en/us/solutions/enterprise/design-zone-secure-data-center-portfolio/index.html references these documents:

- Cisco Secure Data Center for Enterprise Solution Design Guide at http://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-security/sdc-dg.pdf

  This guide includes design and implementation guidance specifically focused on single site clustering with Cisco TrustSec.

- Cisco Secure Data Center for Enterprise (Implementation Guide) at http://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-security/sdc-ig.pdf

  This document is focused on providing implementation guidance for the Cisco Single Site Clustering with IPS and TrustSec solution.

- Cisco Cyber Threat Defense for the Data Center Solution: First Look Guide at http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns744/docs/ctd-first-look-design-guide.pdf

  This guide provides design details and guidance for detecting threats already operating in an internal network or data center.

Please see Appendix—A for a complete configuration on the ASA 5585-X.

Cisco Identity Services Engine

ISE Integration

Two Identity Services Engines can be provisioned in a primary secondary configuration for high availability. Each ISE assumes the following personas:

- Administration Node
- Policy Service Node
- Monitoring Node

The ISE provides RADIUS services.

Cisco Identity Service Engine

Cisco Identity Services Engine (ISE) is an access control system. It provides authentication, authorization, and accounting (AAA) services for a variety of external actors. In the CTS architecture, it has the role of authentication and authorization server. The ISE provides several key roles to the implementation of TrustSec in the data center:

- End-user authentication
- TrustSec device enrollment and authorization (switches, firewalls, management platforms)
- Establishment and central management of SGTs
- Establishment and management of roles-based policies

- Propagates environment data (secure groups, secure group names, SGACLs)
- Manages change of authorizations (CoAs)

The ISE performs other functions, but these are of most interest and relevance to the Secure Data Center for the Enterprise solution.

Secure Group Tags

The Cisco ISE enables end-to-end policies enforced on the basis of role-based access-control (RBACL). Device and user credentials acquired during authentication are used to classify packets by security groups. Every packet entering the Cisco TrustSec domain is tagged with a secure group tag (SGT). The SGT identifies the packet as belonging to either a user or an asset in the data so that policy enforcement can be applied to the packet at the appropriate enforcement point or be processed by advance processing in the ASA 5585-X. Tagging helps trusted intermediaries identify the source identity of the packet and enforce security policies along the data path. An SGT is assigned to a device through IEEE 802.1X authentication, web authentication, or MAC authentication bypass (MAB), which happens with a RADIUS vendor-specific attribute. An SGT can be assigned statically to a particular IP address or to a switch interface. An SGT is passed along dynamically to a switch or access point after successful authentication.

Table 3 lists examples of secure group names and their respective SGTs.

Table 3 Secure Group Names and their SGTs

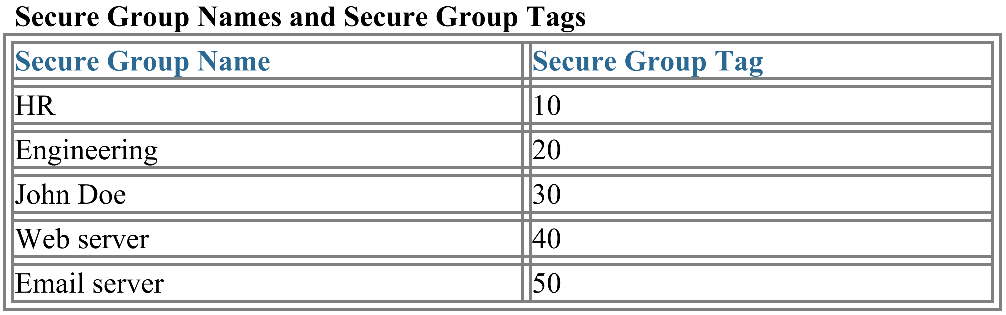

Secure Group Names and Secure Group Tags

Secure Group Name	Secure Group Tag
HR	10
Engineering	20
John Doe	30
Web server	40
Email server	50

SGT Exchange Protocol

SGT Exchange Protocol (SXP) is a protocol developed for Cisco TrustSec to propagate the IP-to-SGT mapping database across network devices that do not have SGT-capable hardware support to hardware that supports SGTs and security group access lists. Typically, SXP is conceived as the protocol between the switches that is used to map SGTs to IP addresses. SXP, a control plane protocol, passes IP-SGT mappings from authentication points (such as legacy access layer switches) to upstream switches and authenticated devices in the network. The SXP connections are point-to-point and use TCP as the underlying transport protocol. SXP uses the well-known TCP port number 64999 when initiating a connection.

Cisco ISE in the Enclave

The enclave uses this for auth_c and auth_z functionality across the system as well as role-based identities for enhanced security. Figure 33 summarizes the two node ISE pair deployed for validation. These are virtual machines deployed via OVF on the VMware vSphere enabled management domain. Notice these two engines support all ISE roles, for larger deployments these personas can be distributed among multiple virtual machines.

Figure 33 Identity Services Engine Nodes

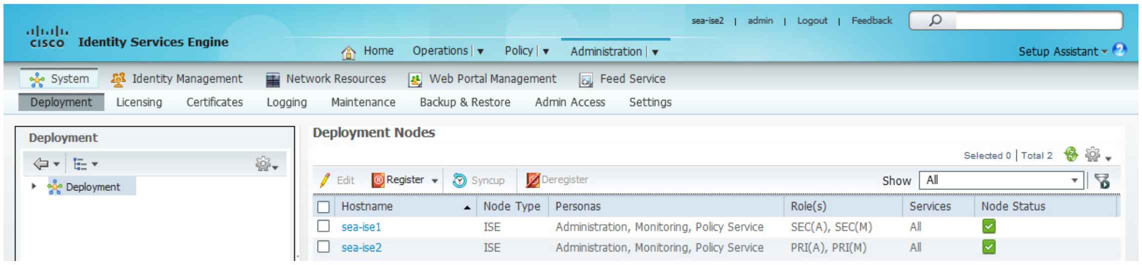

Deployment Nodes

Selected 0 | Total 2

□	Hostname	Node Type	Personas	Role(s)	Services	Node Status
□	sea-ise1	ISE	Administration, Monitoring, Policy Service	SEC(A), SEC(M)	All	✓
□	sea-ise2	ISE	Administration, Monitoring, Policy Service	PRI(A), PRI(M)	All	✓

The remaining sections of this document capture the configurations to address administrative and policy functionality implemented in the enclave. The primary areas of focus include:

- Network Resources
- Identity Management
- Policy Elements
- Authentication Policy
- Authorization Policy

Note The ISE is a powerful tool and the configuration and capabilities captured in this document are simply scratching the surface. It is recommended that readers use the reference documents to fully explore the ISE platform.

Administering Network Resources

A network device such as a switch or a router is an authentication, authorization, and accounting (AAA) client through which AAA service requests are sent to Cisco ISE. The Cisco ISE only supports network devices defined to it. A network device that is not defined in Cisco ISE cannot receive AAA services from Cisco ISE. There are two primary steps to register the device create Network Device Group details and define the device.

Network Device Groups

Network Device Groups (NDGs) that contain network devices. NDGs logically group network devices based on various criteria such as geographic location, device type, and the relative place in the network. Figure 34 illustrates the two forms necessary to complete during NDG creation and a sample from the lab environment. These conditions can be used later to refine device authentication rules.

Figure 34 Network Device Group Types and Locations

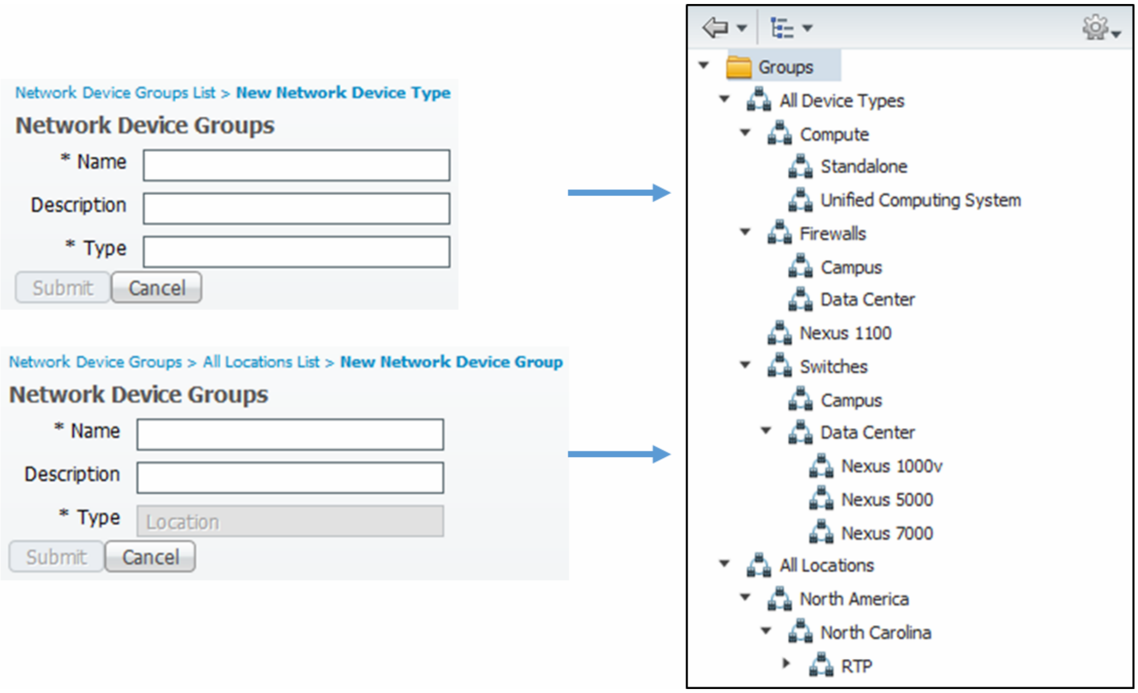

Network Device Groups List > New Network Device Type

Network Device Groups

* Name

Description

* Type

Network Device Groups > All Locations List > New Network Device Group

Network Device Groups

* Name

Description

* Type: Location

- Groups
  - All Device Types
    - Compute
      - Standalone
      - Unified Computing System
    - Firewalls
      - Campus
      - Data Center
    - Nexus 1100
    - Switches
      - Campus
      - Data Center
        - Nexus 1000v
        - Nexus 5000
        - Nexus 7000
  - All Locations
    - North America
      - North Carolina
        - RTP

Network Devices

Figure 35 summarizes the Network Device definitions and required elements. Figure 36 is the expanded view of the default radius authentication settings for the device. These fields should correspond to the radius definitions provided in each of the network elements definition. The name should be identical to the hostname of the device.

Figure 35 Network Device Definition

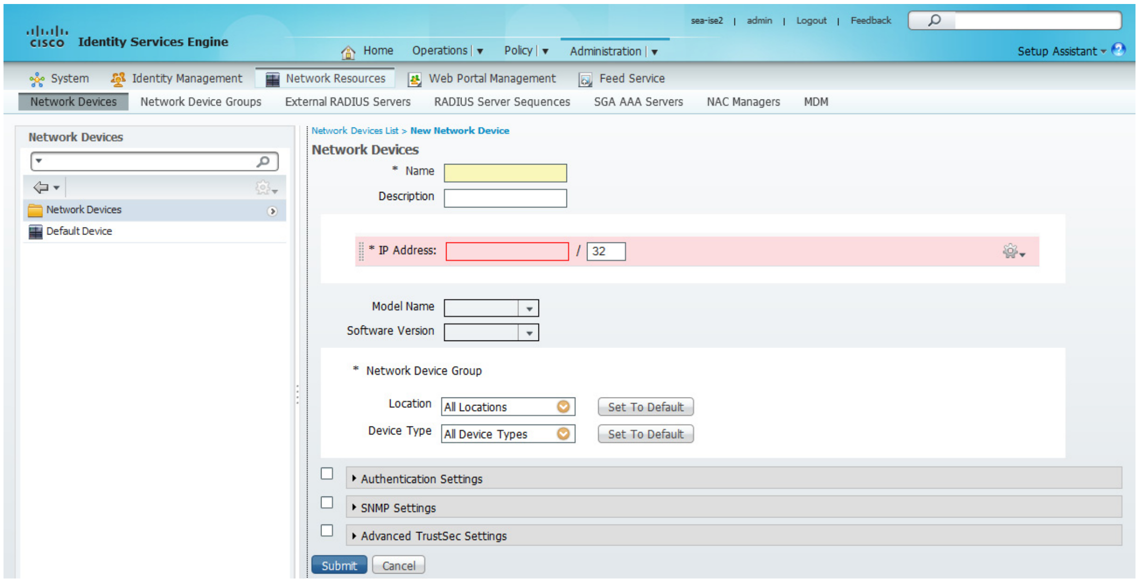

Network Devices List > New Network Device

Network Devices

* Name

Description

* IP Address: / 32

Model Name

Software Version

* Network Device Group

Location: All Locations — Set To Default

Device Type: All Device Types — Set To Default

☐ Authentication Settings

☐ SNMP Settings

☐ Advanced TrustSec Settings

Figure 36 Authentication Settings—Radius (Default)

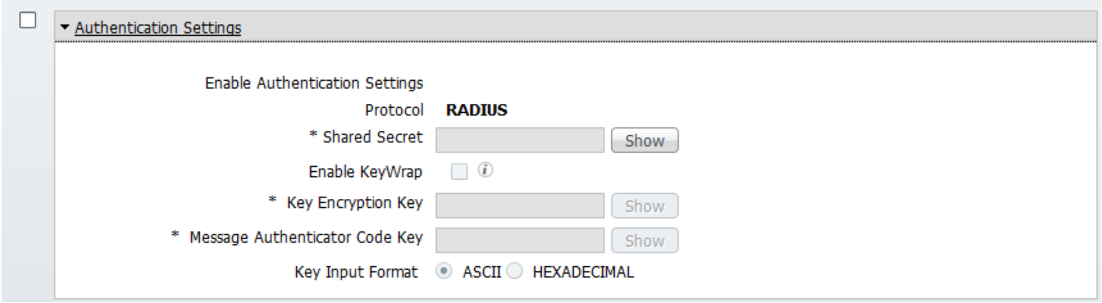

☐ ▾ Authentication Settings

Enable Authentication Settings

Protocol **RADIUS**

* Shared Secret [ ] Show

Enable KeyWrap ☐ ⓘ

* Key Encryption Key [ ] Show

* Message Authenticator Code Key [ ] Show

Key Input Format ● ASCII ○ HEXADECIMAL

Figure 37 is the form for enabling Cisco TrustSec for a particular device. This section defines the Security Group Access attributes for the newly added network device. he PAC file is generated from this page to secure communications between the ISE and the network device.

Figure 37 Advanced TrustSec Settings

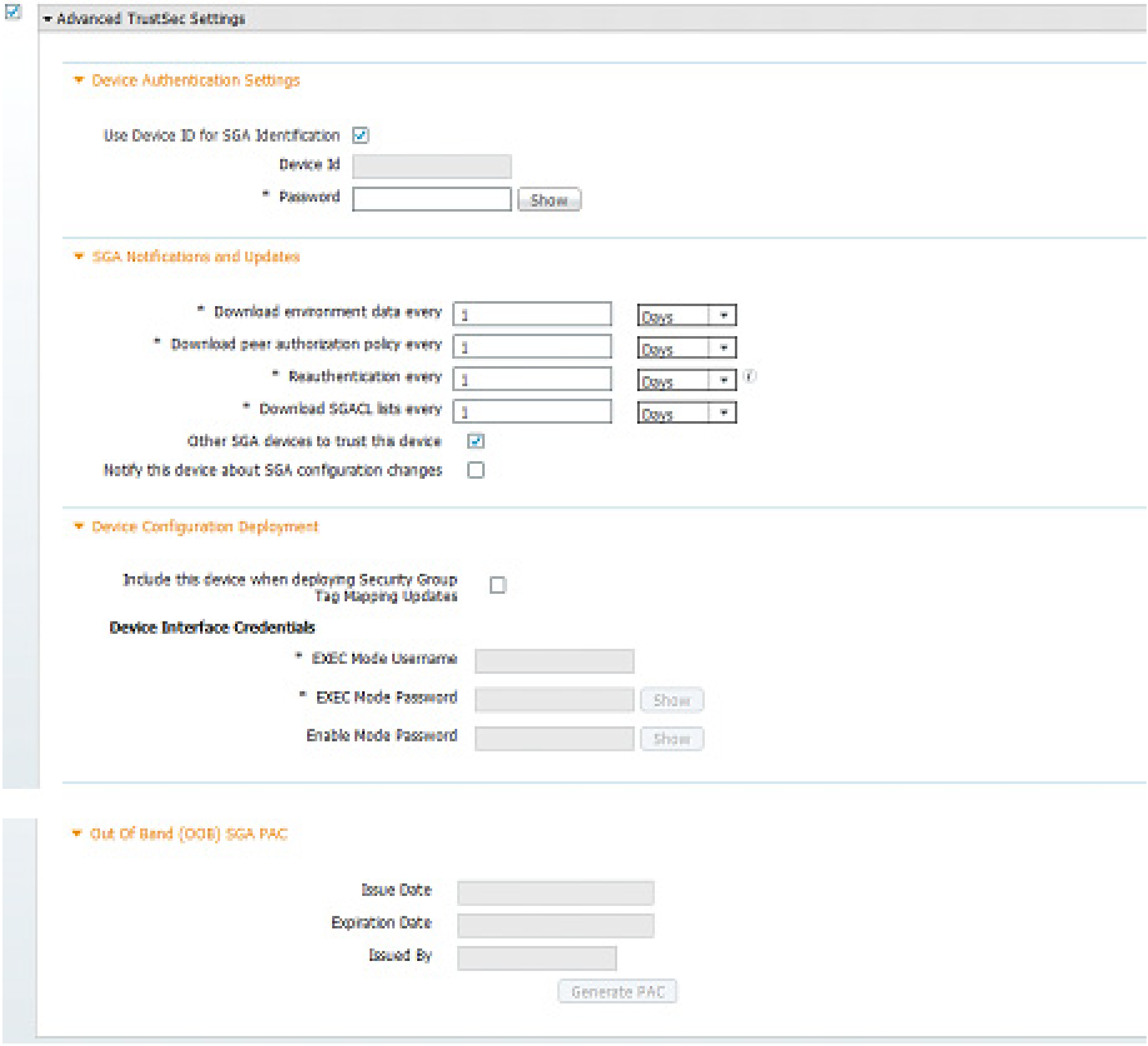

☑ ▾ Advanced TrustSec Settings

▾ Device Authentication Settings

Use Device ID for SGA Identification ☑

Device Id [ ]

* Password [ ] Show

▾ SGA Notifications and Updates

* Download environment data every [1] [Days ▾]

* Download peer authorization policy every [1] [Days ▾]

* Reauthentication every [1] [Days ▾] ⓘ

* Download SGACL lists every [1] [Days ▾]

Other SGA devices to trust this device ☑

Notify this device about SGA configuration changes ☐

▾ Device Configuration Deployment

Include this device when deploying Security Group Tag Mapping Updates ☐

Device Interface Credentials

* EXEC Mode Username [ ]

* EXEC Mode Password [ ] Show

Enable Mode Password [ ] Show

▾ Out Of Band (OOB) SGA PAC

Issue Date [ ]

Expiration Date [ ]

Issued By [ ]

Generate PAC

Administering Identity Management

External Identity Sources

The Cisco ISE can store or reference internal or external user information for authentication and authorization. The following example will document the use of Microsoft Active Directory as the singles source of truth for valid users in the organization. Using a single source of truth minimizes risk as data and its concurrency is maintained in a single repository promoting accuracy and operation efficiency.

Connection

The connection to the Active Directory external identity store is established by providing Domain and a locally significant name to the data source. Figure 38 shows the connection between the ISE active standby pair and the CORP domain. After joining the domain the Cisco ISE can access user, group and device data.

Figure 38 Cisco ISE Active Directory Connection Example

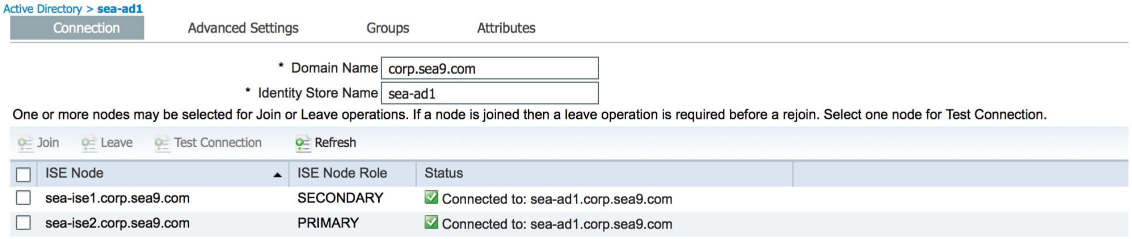

* Domain Name: corp.sea9.com

* Identity Store Name: sea-ad1

One or more nodes may be selected for Join or Leave operations. If a node is joined then a leave operation is required before a rejoin. Select one node for Test Connection.

□	ISE Node	ISE Node Role	Status
□	sea-ise1.corp.sea9.com	SECONDARY	✓ Connected to: sea-ad1.corp.sea9.com
□	sea-ise2.corp.sea9.com	PRIMARY	✓ Connected to: sea-ad1.corp.sea9.com

Groups

The Active Directory connection allows the Cisco ISE to use the repositories group construct. These groups can be referenced for authentication rules. For example, Figure 39 shows four groups defined in AD being used by the ISE.

Figure 39 Cisco ISE Active Directory Group Reference Example

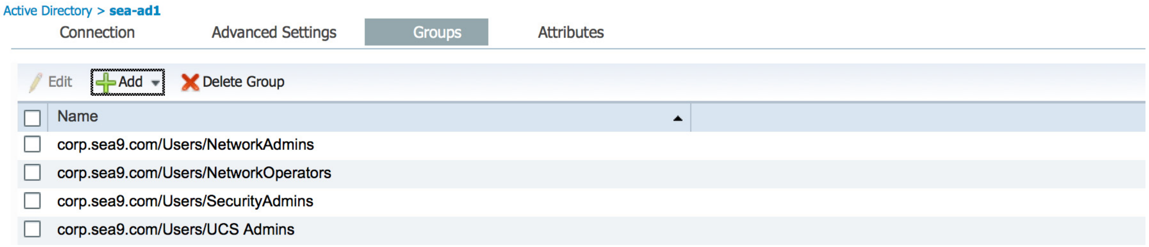

□	Name
□	corp.sea9.com/Users/NetworkAdmins
□	corp.sea9.com/Users/NetworkOperators
□	corp.sea9.com/Users/SecurityAdmins
□	corp.sea9.com/Users/UCS Admins

Figure 40 is a snippet of the form to add these groups to the Cisco ISE. Notice the groups previously selected.

Figure 40 Cisco ISE Select Directory Groups Form Example

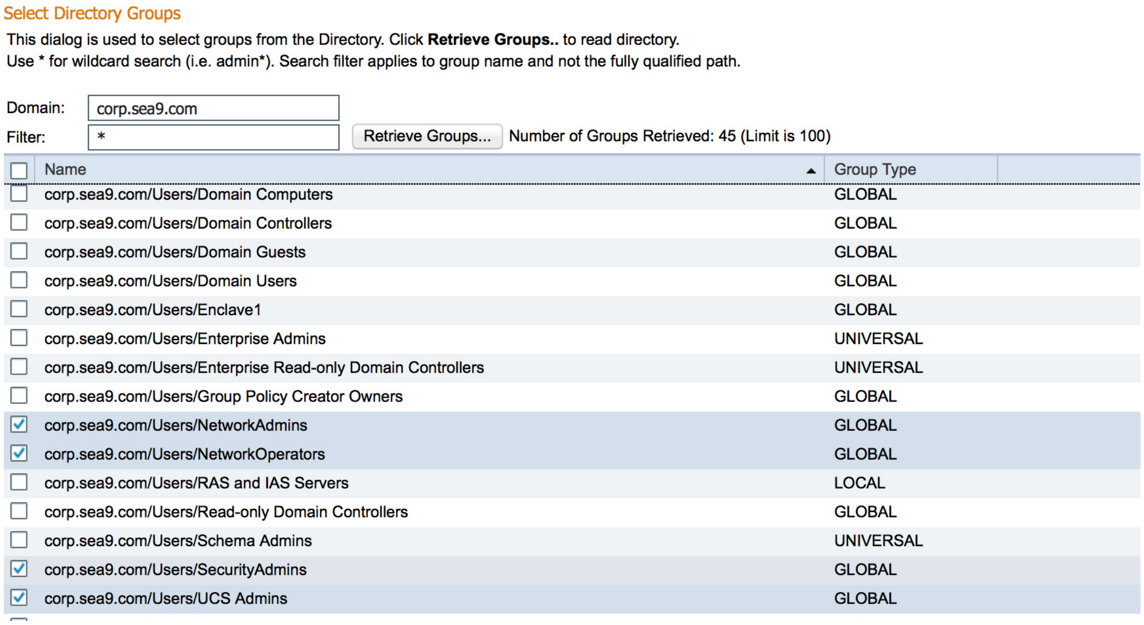

Select Directory Groups

This dialog is used to select groups from the Directory. Click **Retrieve Groups..** to read directory.
Use * for wildcard search (i.e. admin*). Search filter applies to group name and not the fully qualified path.

Domain: corp.sea9.com

Filter: *

Number of Groups Retrieved: 45 (Limit is 100)

☐	Name	Group Type
☐	corp.sea9.com/Users/Domain Computers	GLOBAL
☐	corp.sea9.com/Users/Domain Controllers	GLOBAL
☐	corp.sea9.com/Users/Domain Guests	GLOBAL
☐	corp.sea9.com/Users/Domain Users	GLOBAL
☐	corp.sea9.com/Users/Enclave1	GLOBAL
☐	corp.sea9.com/Users/Enterprise Admins	UNIVERSAL
☐	corp.sea9.com/Users/Enterprise Read-only Domain Controllers	UNIVERSAL
☐	corp.sea9.com/Users/Group Policy Creator Owners	GLOBAL
☑	corp.sea9.com/Users/NetworkAdmins	GLOBAL
☑	corp.sea9.com/Users/NetworkOperators	GLOBAL
☐	corp.sea9.com/Users/RAS and IAS Servers	LOCAL
☐	corp.sea9.com/Users/Read-only Domain Controllers	GLOBAL
☐	corp.sea9.com/Users/Schema Admins	UNIVERSAL
☑	corp.sea9.com/Users/SecurityAdmins	GLOBAL
☑	corp.sea9.com/Users/UCS Admins	GLOBAL

Identity Source Sequences

Identity source sequences define the order in which Cisco ISE looks for user credentials in the different databases available to it. Cisco ISE supports the following identity sources:

- Internal Users
- Guest Users
- Active Directory
- LDAP
- RSA
- RADIUS Token Servers
- Certificate Authentication Profiles

The ISE uses a first match policy across the identity sources for authentication and authorization purposes.

AD Sequence

The Active Directory service sequence is added referencing the previously joined domain. This sequence will be used during authentication policy creation. Figure 41 illustrates the addition of an "AD_Sequence" using the previously joined AD domain as an identity source.

Figure 41 Cisco ISE Identity Source Sequence Example

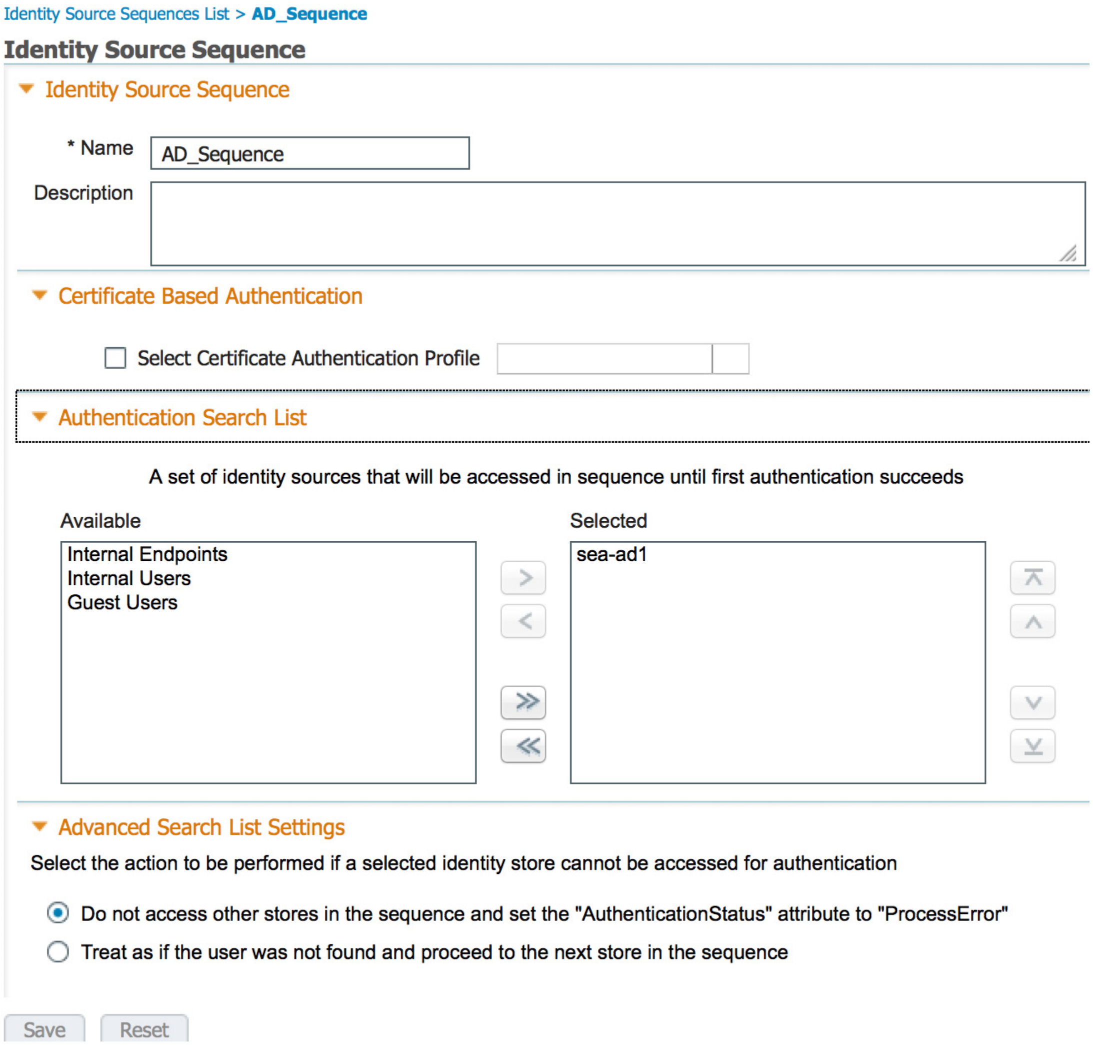

Identity Source Sequences List > **AD_Sequence**

Identity Source Sequence

▾ Identity Source Sequence

* Name: AD_Sequence

Description:

▾ Certificate Based Authentication

□ Select Certificate Authentication Profile

▾ Authentication Search List

A set of identity sources that will be accessed in sequence until first authentication succeeds

Available	Selected
Internal Endpoints Internal Users Guest Users	sea-ad1

▾ Advanced Search List Settings

Select the action to be performed if a selected identity store cannot be accessed for authentication

● Do not access other stores in the sequence and set the "AuthenticationStatus" attribute to "ProcessError"

○ Treat as if the user was not found and proceed to the next store in the sequence

Policy Elements—Results

This following policy elements were defined in the Secure Enclave architecture:

- Authorization Profiles
- Security Group Access
- Authorization Profiles

Policy elements are the components that construct the policies associated with authentication, authorization, and secure group access. Authorization profiles define policy components related to permissions. Authorization profiles are used when creating authorization policies. The authorization profile returns permission attributes when the RADIUS request is accepted.

Figure 42 captures some of the default and custom authorization profiles used during validation. Figure 43 details the UCS_Admins profile that upon authentication the UCS admin role is assigned through the cisco-av-pair radius attribute value. Note that the cisco-av-pair value varies based on the Cisco device type, please refer to device specific documentation for the proper syntax.

Figure 42 Policy Element Results—Authorization Profiles Example

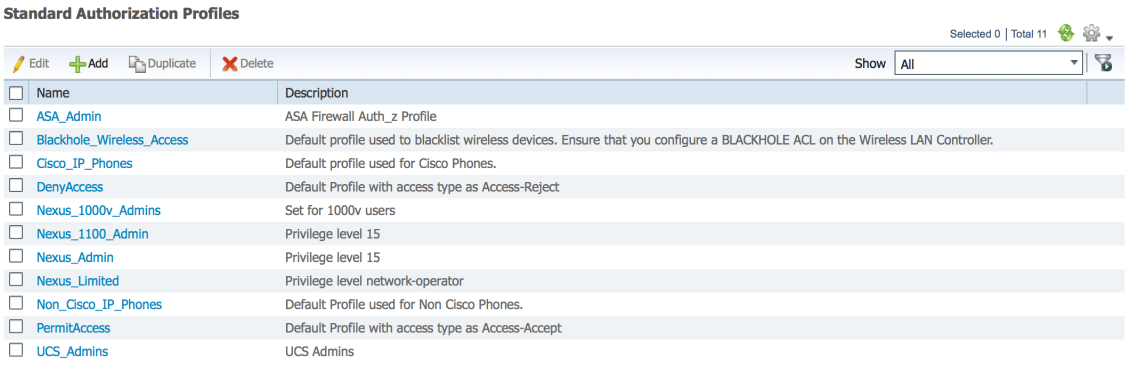

Standard Authorization Profiles

	Name	Description
□	ASA_Admin	ASA Firewall Auth_z Profile
□	Blackhole_Wireless_Access	Default profile used to blacklist wireless devices. Ensure that you configure a BLACKHOLE ACL on the Wireless LAN Controller.
□	Cisco_IP_Phones	Default profile used for Cisco Phones.
□	DenyAccess	Default Profile with access type as Access-Reject
□	Nexus_1000v_Admins	Set for 1000v users
□	Nexus_1100_Admin	Privilege level 15
□	Nexus_Admin	Privilege level 15
□	Nexus_Limited	Privilege level network-operator
□	Non_Cisco_IP_Phones	Default Profile used for Non Cisco Phones.
□	PermitAccess	Default Profile with access type as Access-Accept
□	UCS_Admins	UCS Admins

Figure 43 Cisco ISE Authorization Profile Example

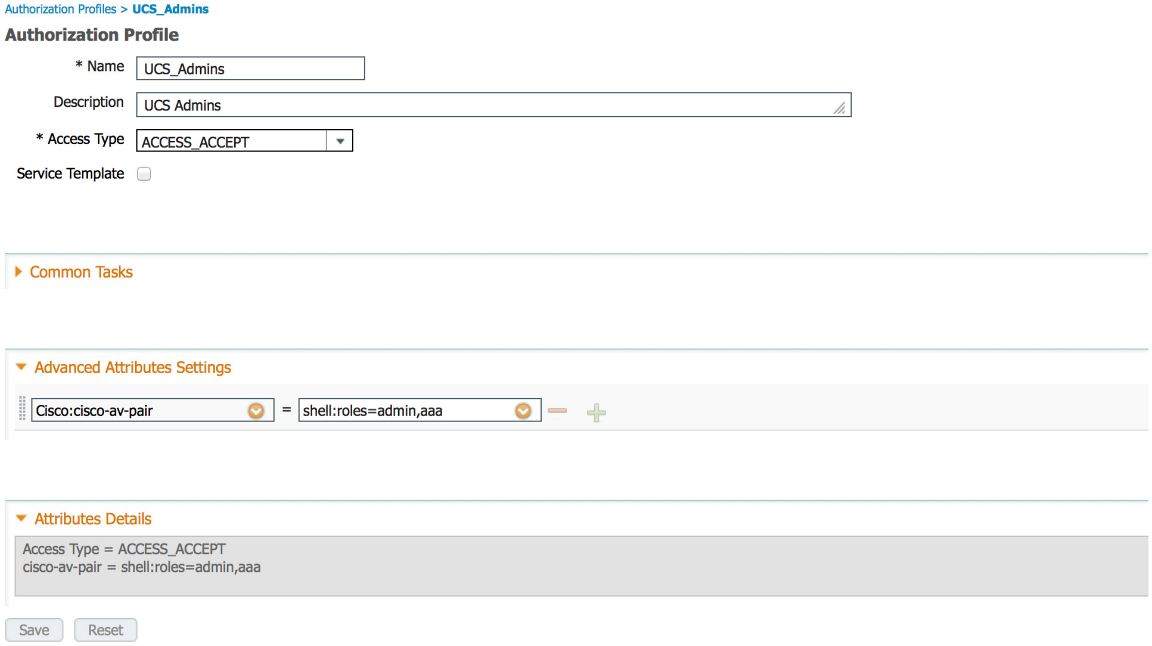

Authorization Profiles > **UCS_Admins**

Authorization Profile

* Name: UCS_Admins

Description: UCS Admins

* Access Type: ACCESS_ACCEPT

Service Template: □

Advanced Attributes Settings

Cisco:cisco-av-pair = shell:roles=admin,aaa

Attributes Details

Access Type = ACCESS_ACCEPT
cisco-av-pair = shell:roles=admin,aaa

Secure Group Access—Security Groups

Packets within the Secure Enclave architecture are tagged to support role-based security policy. The Cisco ISE contains the tag definitions that can be auto-generated or manually assigned. Figure 44 is a sample of the tags used in the Enclave validation effort. Notice that each enclave role (app, db or web) has a unique tag. Figure 45 captures the import process form which allows for bulk create of SGT information on the ISE platform.

Figure 44 Cisco ISE Security Groups Example

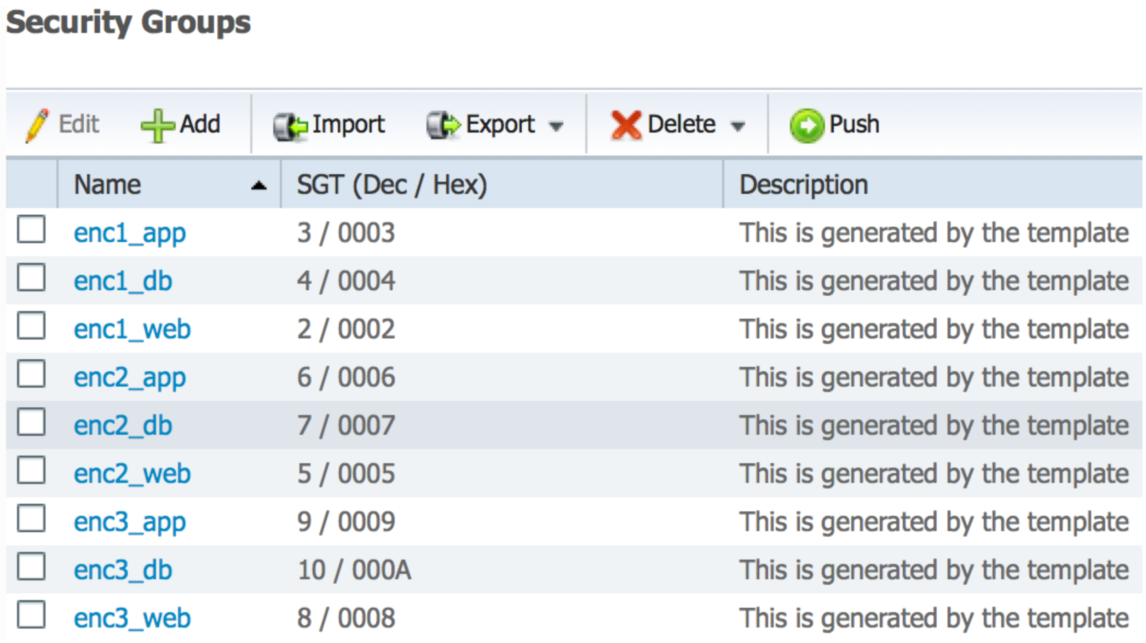

	Name	SGT (Dec / Hex)	Description
□	enc1_app	3 / 0003	This is generated by the template
□	enc1_db	4 / 0004	This is generated by the template
□	enc1_web	2 / 0002	This is generated by the template
□	enc2_app	6 / 0006	This is generated by the template
□	enc2_db	7 / 0007	This is generated by the template
□	enc2_web	5 / 0005	This is generated by the template
□	enc3_app	9 / 0009	This is generated by the template
□	enc3_db	10 / 000A	This is generated by the template
□	enc3_web	8 / 0008	This is generated by the template

Figure 45 Security Groups Form—Import Process

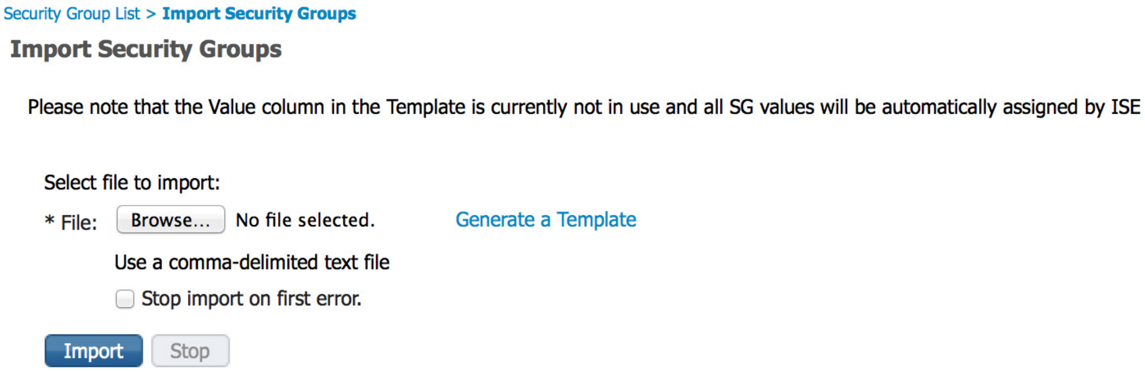

Security Group List > **Import Security Groups**

Import Security Groups

Please note that the Value column in the Template is currently not in use and all SG values will be automatically assigned by ISE

Select file to import:

* File: Browse... No file selected. Generate a Template

Use a comma-delimited text file

□ Stop import on first error.

Authentication Policy

The Cisco ISE authentication policy defines the acceptable communication protocol and identity source for network device authentication. This policy is built using conditions or device attributes previously defined such as device type or location as well as the acceptable network protocol. Figure 46 shows the authentication policy associated with the Cisco UCS system. Essentially the rule states that if the device type is UCS and the communication is using the password authentication protocol (Pap_ASCII) use the identity source defined in the AD_Sequence.

Figure 46 Cisco ISE Authentication Policy Example

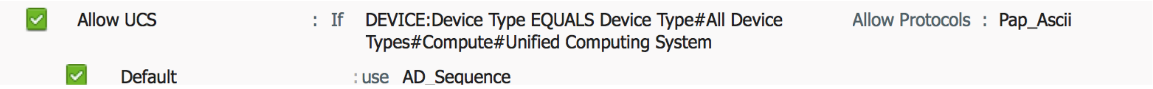

☑	Allow UCS	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Compute#Unified Computing System	Allow Protocols	:	Pap_Ascii
☑	Default	: use	AD_Sequence			

Figure 47 illustrates the definition of multiple ISE authentication policies each built to meet the specific needs of the network device and the overall organization.

Figure 47 Cisco ISE Authentication Policies

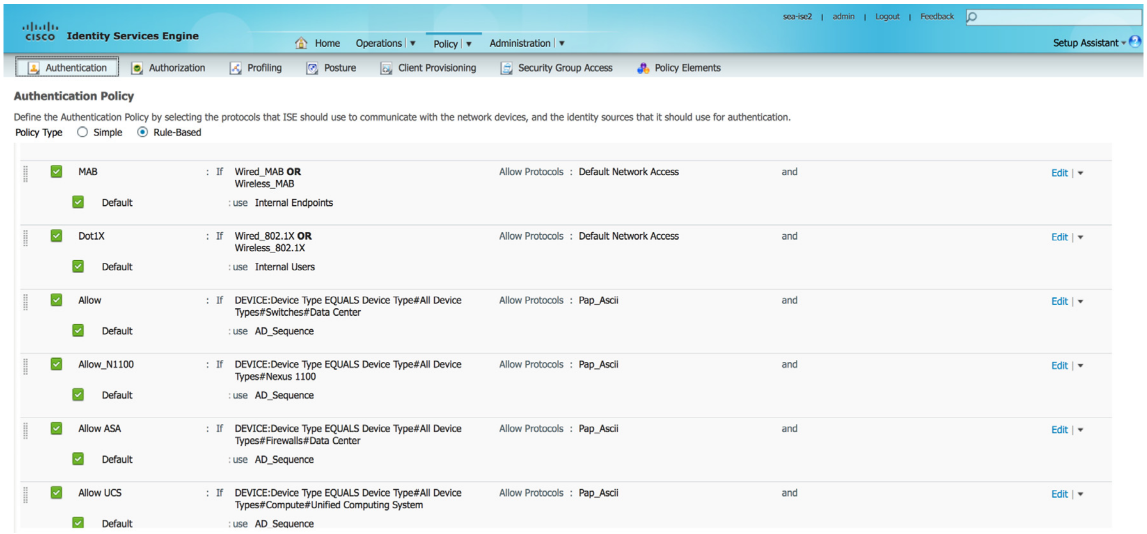

Authentication Policy

Define the Authentication Policy by selecting the protocols that ISE should use to communicate with the network devices, and the identity sources that it should use for authentication.

Policy Type ○ Simple ● Rule-Based

☑	MAB	: If	Wired_MAB **OR** Wireless_MAB	Allow Protocols	:	Default Network Access	and	Edit | ▾
☑	Default	: use	Internal Endpoints					
☑	Dot1X	: If	Wired_802.1X **OR** Wireless_802.1X	Allow Protocols	:	Default Network Access	and	Edit | ▾
☑	Default	: use	Internal Users					
☑	Allow	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Switches#Data Center	Allow Protocols	:	Pap_Ascii	and	Edit | ▾
☑	Default	: use	AD_Sequence					
☑	Allow_N1100	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Nexus 1100	Allow Protocols	:	Pap_Ascii	and	Edit | ▾
☑	Default	: use	AD_Sequence					
☑	Allow ASA	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Firewalls#Data Center	Allow Protocols	:	Pap_Ascii	and	Edit | ▾
☑	Default	: use	AD_Sequence					
☑	Allow UCS	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Compute#Unified Computing System	Allow Protocols	:	Pap_Ascii	and	Edit | ▾
☑	Default	: use	AD_Sequence					

Authorization Policy

The ISE authorization policy enables the organization to set specific privileges and access rights based on any number of conditions. If the conditions are met a permission level or authorization profile is assigned to the user and applied to the network device being accessed. For example, in Figure 48 the UCS Admins authorization policy has a number of conditions that must be met including location, access protocol and Active Directory group membership before the UCS_Admins authorization profiles permissions are assigned to that user session. The Cisco ISE allows organizations to capture the context of a user session and make decisions more intelligently. Figure 49 shows that multiple authorization policies are supported.

Figure 48 Cisco ISE Authorization Policy Example

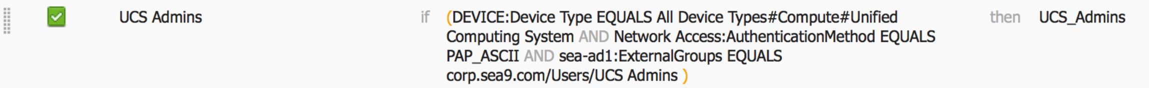

☑	UCS Admins	if	(DEVICE:Device Type EQUALS All Device Types#Compute#Unified Computing System AND Network Access:AuthenticationMethod EQUALS PAP_ASCII AND sea-ad1:ExternalGroups EQUALS corp.sea9.com/Users/UCS Admins)	then	UCS_Admins

Figure 49 *Cisco ISE Authorization Policies*

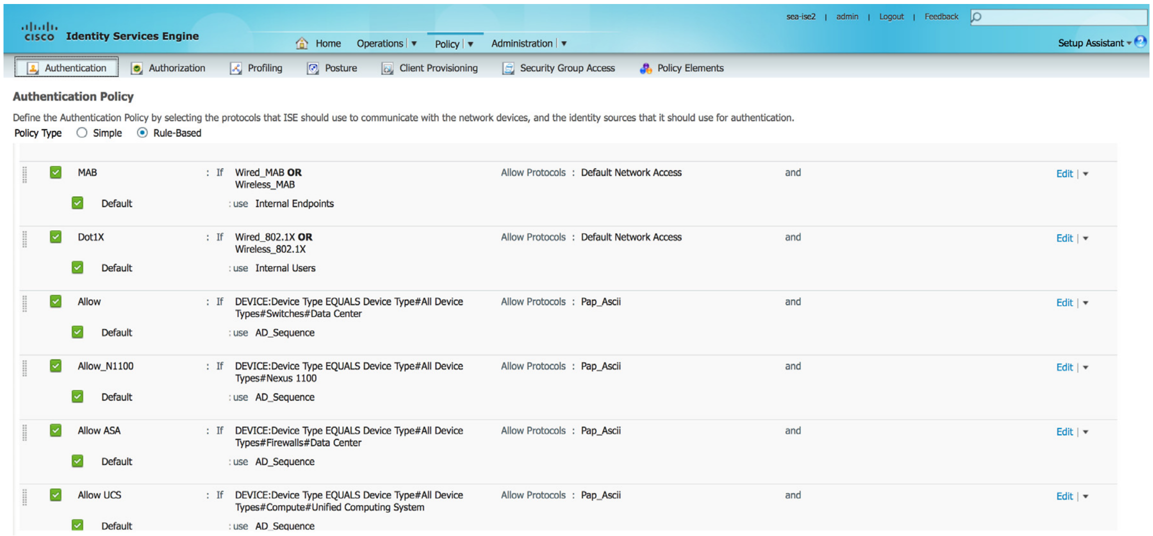

Authentication Policy

Define the Authentication Policy by selecting the protocols that ISE should use to communicate with the network devices, and the identity sources that it should use for authentication.

Policy Type ○ Simple ● Rule-Based

☑ MAB	: If	Wired_MAB **OR** Wireless_MAB	Allow Protocols : Default Network Access	and	Edit ▾
☑ Default	: use	Internal Endpoints			
☑ Dot1X	: If	Wired_802.1X **OR** Wireless_802.1X	Allow Protocols : Default Network Access	and	Edit ▾
☑ Default	: use	Internal Users			
☑ Allow	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Switches#Data Center	Allow Protocols : Pap_Ascii	and	Edit ▾
☑ Default	: use	AD_Sequence			
☑ Allow_N1100	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Nexus 1100	Allow Protocols : Pap_Ascii	and	Edit ▾
☑ Default	: use	AD_Sequence			
☑ Allow ASA	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Firewalls#Data Center	Allow Protocols : Pap_Ascii	and	Edit ▾
☑ Default	: use	AD_Sequence			
☑ Allow UCS	: If	DEVICE:Device Type EQUALS Device Type#All Device Types#Compute#Unified Computing System	Allow Protocols : Pap_Ascii	and	Edit ▾
☑ Default	: use	AD_Sequence			

NetFlow

A flow is identified as a unidirectional stream of packets between a given source and destination. NetFlow is a Cisco application that measures the IP network traffic attributes of a traffic flow as it traverses the Cisco device. NetFlow was initially created to measure network traffic characteristics such as bandwidth, application performance, and utilization; and has historically been used for billing and accounting, network capacity planning, and availability monitoring. NetFlow is a reporting technology. As traffic traverses a device, the device gathers information about the traffic flow and reports on the information after the flow has occurred. NetFlow reporting has tremendous security applications as well, including the ability to provide non-repudiation, anomaly detection, and investigative capabilities.The Cisco Cyber Threat Defense for the Data Center Solution uses NetFlow Version 9. NetFlow Version 9 completely separates the collection and export process and allows the customization of the NetFlow collection. Using this approach, the Cisco Cyber Threat Defense for the Data Center Solution captures NetFlow data across the infrastructure to maximize the security monitoring potential by collecting packet fields such as TCP flags, Time To Live (TTL) values, and protocol.

Some data center network devices support NetFlow via software, rather than hardware support, such as the Cisco Nexus 1000V. Give some consideration to a software device's current utilization when deploying software-supported NetFlow services, because enabling NetFlow can affect device performance. An alternative to using NetFlow natively on the Cisco Nexus 1000V is to use a Lancope StealthWatch FlowSensor VE to generate NetFlow records from the virtual access layer.

NetFlow Generation Appliance NGA3240

The Cisco NetFlow Generation Appliance (NGA), a purpose-built, high-performance solution for flow visibility in multi-gigabit data centers, can restore flow visibility in these environments in a scalable and affordable manner.The Cisco NGA has four 10G monitoring interfaces and up to four independent flow caches and flow monitors. This means that the Cisco NGA can receive up to 40 Gigabit of data and support various combinations of data ports, record templates, and export parameters. This is important to consider when placing the NGA inside the data center. The NGA can be placed to receive data from the physical access, aggregation, and core layers. The objective is to ensure complete visibility of all traffic within the data center, as well as traffic that is leaving the data center. Traffic within the virtual environment (VM-to-VM traffic) can be monitored using the Nexus 1000V, while traffic entering and leaving the data center can be monitored using edge devices such as the ASA and Cisco Nexus 5000.

Cisco NetFlow Generation Appliance

Each NetFlow Generation Appliance is configured to accept SPAN traffic from up to four different ten Gigabit Ethernet data ports. These promiscuous ports can be easily setup using the NGA Quick Setup web form as shown in Figure 50. The quick setup pane configures setup to a single collector.

Figure 50 Cisco NetFlow Generation Appliance—Quick Setup Form

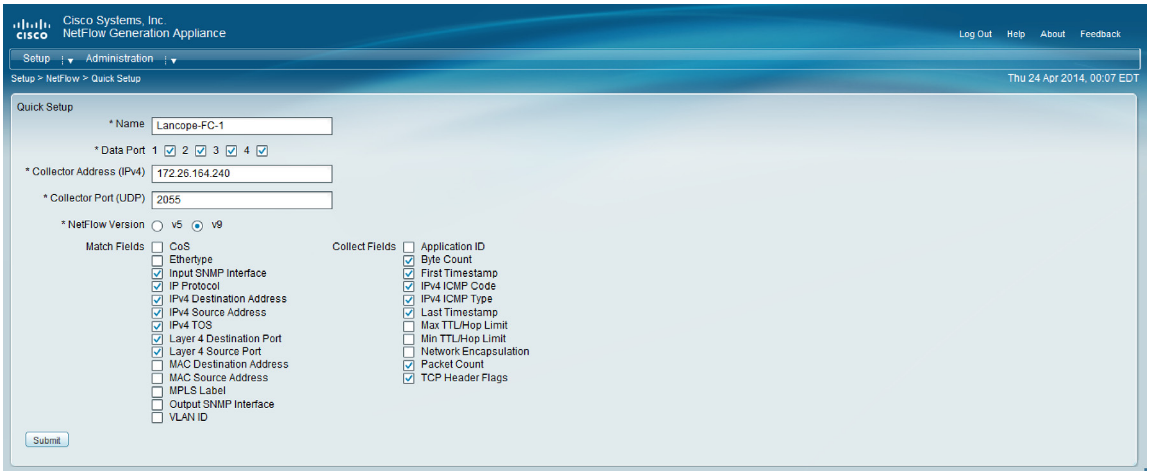
Cisco Systems, Inc.
NetFlow Generation Appliance
Log Out
Help
About
Feedback
Setup
Administration
Setup > NetFlow > Quick Setup
Thu 24 Apr 2014, 00:07 EDT

Quick Setup

* Name: Lancope-FC-1

* Data Port: 1 ☑ 2 ☑ 3 ☑ 4 ☑

* Collector Address (IPv4): 172.26.164.240

* Collector Port (UDP): 2055

* NetFlow Version: ○ v5 ● v9

Match Fields:
- ☐ CoS
- ☐ Ethertype
- ☑ Input SNMP Interface
- ☑ IP Protocol
- ☑ IPv4 Destination Address
- ☑ IPv4 Source Address
- ☑ IPv4 TOS
- ☑ Layer 4 Destination Port
- ☑ Layer 4 Source Port
- ☐ MAC Destination Address
- ☐ MAC Source Address
- ☐ MPLS Label
- ☐ Output SNMP Interface
- ☐ VLAN ID

Collect Fields:
- ☐ Application ID
- ☑ Byte Count
- ☑ First Timestamp
- ☑ IPv4 ICMP Code
- ☑ IPv4 ICMP Type
- ☑ Last Timestamp
- ☐ Max TTL/Hop Limit
- ☐ Min TTL/Hop Limit
- ☐ Network Encapsulation
- ☑ Packet Count
- ☑ TCP Header Flags

The following screenshots capture a single NGA configuration used in the enclave validation effort. The NGA redirects all traffic to the Lancope Flow Collector at 172.26.164.240. Figure 51 describes the collector defined using the quick form.

Figure 51 NGA Flow Collector Definition Example

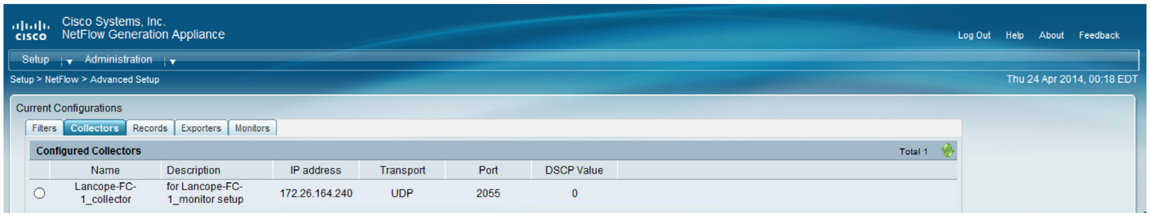

Configured Collectors — Total 1

	Name	Description	IP address	Transport	Port	DSCP Value
○	Lancope-FC-1_collector	for Lancope-FC-1_monitor setup	172.26.164.240	UDP	2055	0

Figure 52 details the NetFlow record being sent to the collector.

Figure 52 NGA NetFlow Record Definition Example

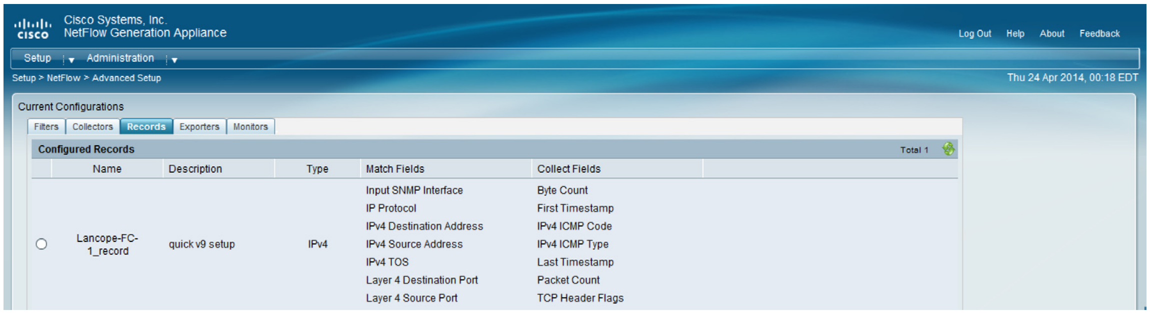

	Name	Description	Type	Match Fields	Collect Fields
○	Lancope-FC-1_record	quick v9 setup	IPv4	Input SNMP Interface IP Protocol IPv4 Destination Address IPv4 Source Address IPv4 TOS Layer 4 Destination Port Layer 4 Source Port	Byte Count First Timestamp IPv4 ICMP Code IPv4 ICMP Type Last Timestamp Packet Count TCP Header Flags

The export details are set to their defaults.

Figure 53 NGA NetFlow Exporter Definition Example

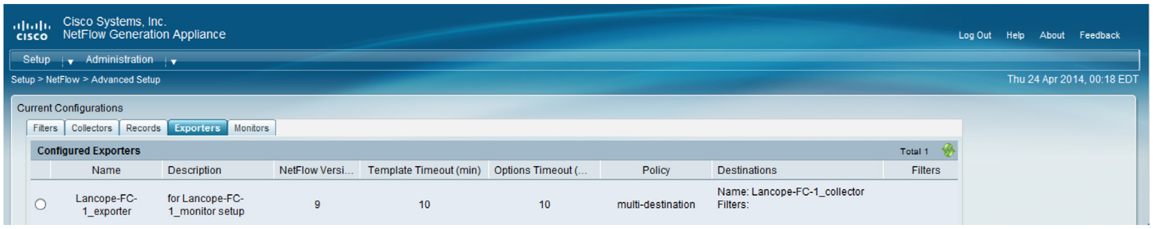

	Name	Description	NetFlow Versi...	Template Timeout (min)	Options Timeout (...	Policy	Destinations	Filters
○	Lancope-FC-1_exporter	for Lancope-FC-1_monitor setup	9	10	10	multi-destination	Name: Lancope-FC-1_collector Filters:	

Figure 54 shows the result of the quick setup implemented in the enclave architecture. The Lancope monitor is created with all four data ports of mirrored traffic being sent to the Lancope flow collector.

Figure 54 NGA Monitor Definition Example

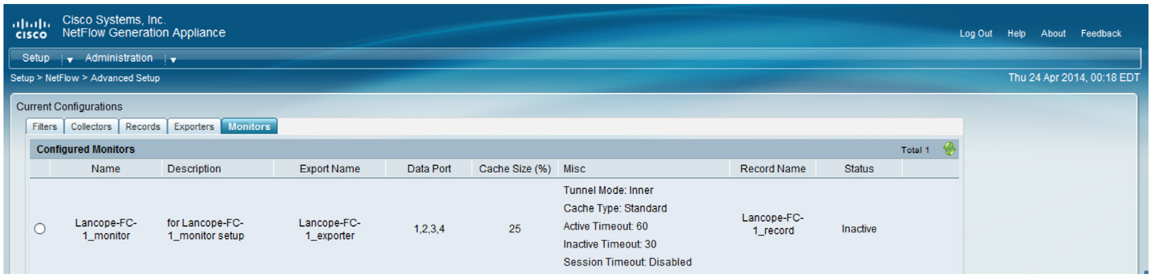

	Name	Description	Export Name	Data Port	Cache Size (%)	Misc	Record Name	Status
○	Lancope-FC-1_monitor	for Lancope-FC-1_monitor setup	Lancope-FC-1_exporter	1,2,3,4	25	Tunnel Mode: Inner Cache Type: Standard Active Timeout: 60 Inactive Timeout: 30 Session Timeout: Disabled	Lancope-FC-1_record	Inactive

Lancope StealthWatch System

The Lancope StealthWatch Management Console (SMC) accepts and parses syslog messages from any ISE node to collect identity information. The Lancope SMC must be configured as a remote logging target on the ISE monitoring node with the RADIUS Accounting, Profiler, and Administrative and Operational Audit Logging categories set to log to the SMC target. Further details about this and a NetFlow quick setup procedure on the NGA3240, NSEL configuration for ASA device enablement and Switched Port Analyzer (SPAN) for the Nexus 5000/7000 is detailed in the following document:

http://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-security/ctd-first-look-design-guide.pdf

Automation Using Cisco UCS Director

The following section outlines pre-requisites to install and setup a working instance of Cisco UCS Director. The intent is to leverage the automation features of Cisco UCS Director for correct and consistent deployment of enclaves. The construct that is applicable is "Application Container Template" which will allow for defining a base three-tier application container with web, application and database instances and a firewall (VSG) for access control. The section below will conclude with a procedure to create an application container template. The understanding is that the same procedure could be followed to deploy customized enclaves when necessary support from Cisco UCS Director for all underlying components is available.

Cisco UCS Director Installation and Configuration

Down load Cisco UCS Director 4.1 version OVF template from:
http://software.cisco.com/cisco/pub/software/portal/select.html?&i=!m&mdfid=284775897

Proceed with download of all relevant patches as well. This includes cucsd_patch_4_1 HOTFIX and everything following it until cucsd_patch_4_1_0_3.zip. You will need baremetal agent software (cucsd_BMA_4_1_0_0_GA.zip) if building baremetal instances. In this case, we will be installing only VMWare virtual instances.

1. Through the vSphere client, connect to an ESX host within your common infrastructure.
2. Select File, Deploy OVF Template, and choose Browse to navigate to the location of the downloaded CUCSD OVF file.

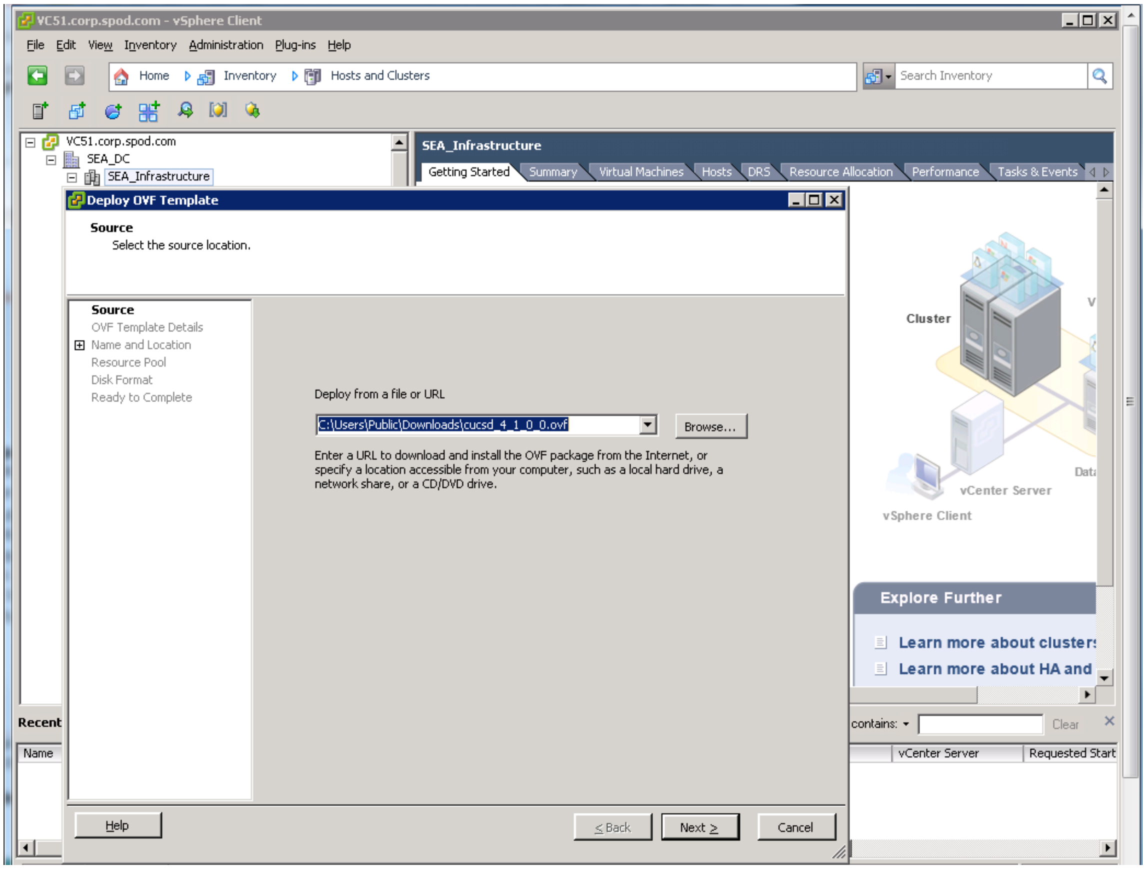

Source
Select the source location.

Deploy from a file or URL
C:\Users\Public\Downloads\cucsd_4_1_0_0.ovf

Enter a URL to download and install the OVF package from the Internet, or specify a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

3. Select the OVF file and click Open, then click Next.
4. Click Next on the OVF Template Details page.
5. Read the terms of the End User License Agreement, and click Accept, then click Next.
6. Leave CUCSD-4.1.0.0 as the VM name and choose the SEA_Infrastructure cluster under SEA_DC. Datacenter as the Inventory Location, and click Next.
7. Select SEA_Infrastructure cluster and click Next.
8. Choose the storage location for Datastore and click Next.
9. Click Next for the Disk Format page.
10. Choose the ProdNet for destination network and click Next.

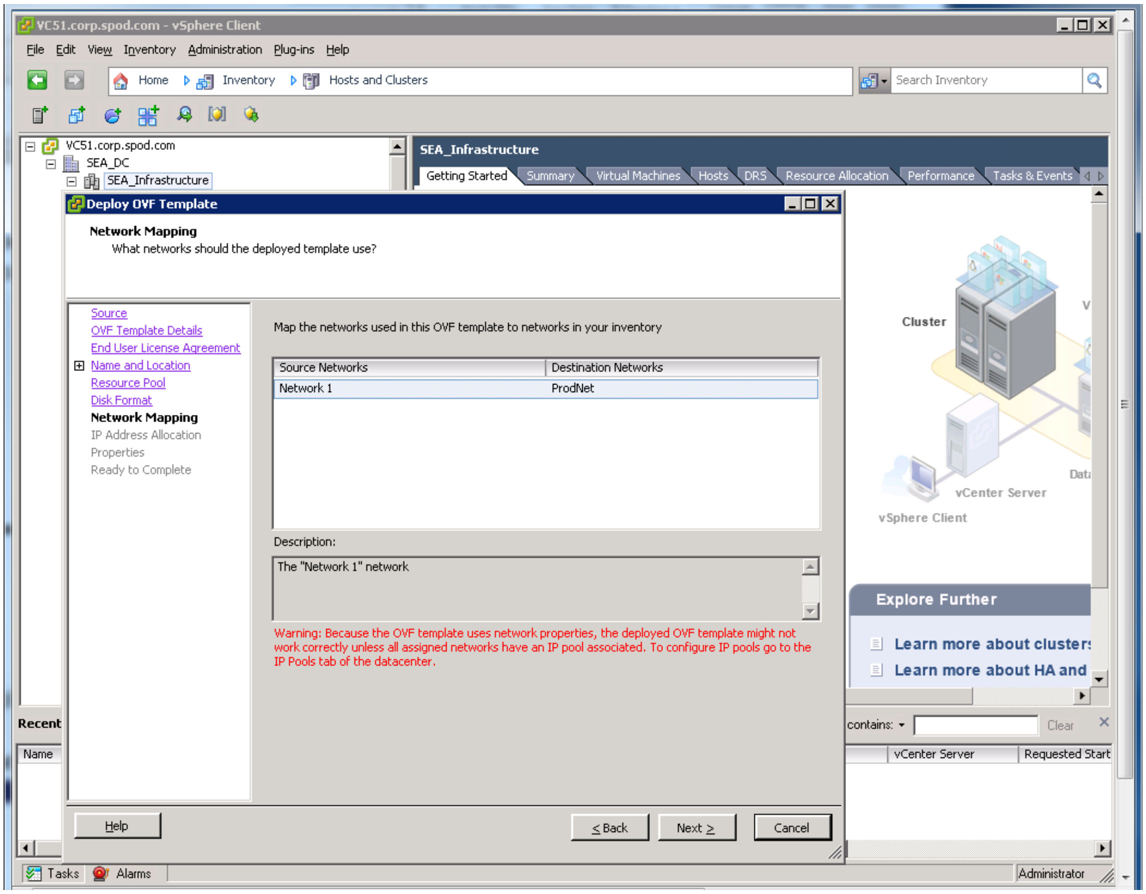

Deploy OVF Template

Network Mapping

What networks should the deployed template use?

Map the networks used in this OVF template to networks in your inventory

Source Networks	Destination Networks
Network 1	ProdNet

Description:

The "Network 1" network

Warning: Because the OVF template uses network properties, the deployed OVF template might not work correctly unless all assigned networks have an IP pool associated. To configure IP pools go to the IP Pools tab of the datacenter.

11. Leave DHCP selected for now on the IP allocation page and click Next.
12. Click Finish. The import will begin and the progress of the import will be displayed on the screen.
13. Click Close once this operation completes

Initial Setup

In this section we will configure the UCS Director Virtual Machine on VMware.

Note Upgrade the reserved resources for the newly created VM.

1. Right-click the CUSD VM and click Edit Settings.

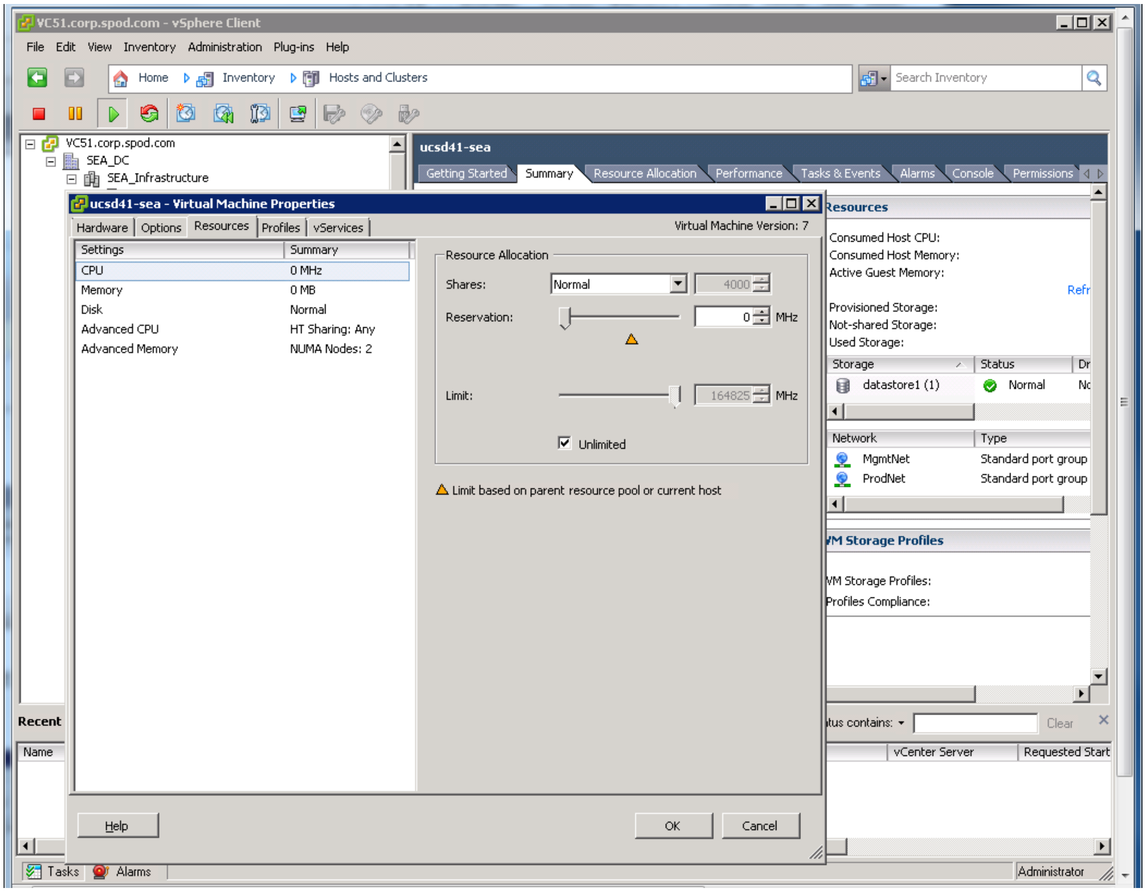

Settings	Summary
CPU	0 MHz
Memory	0 MB
Disk	Normal
Advanced CPU	HT Sharing: Any
Advanced Memory	NUMA Nodes: 2

Resource Allocation

Shares: Normal 4000

Reservation: 0 MHz

Limit: 164825 MHz

☑ Unlimited

⚠ Limit based on parent resource pool or current host

2. Select the Resources tab.

3. Select CPU, and change the Reservation to 4000 MHz, then select Memory, and change Reservation to 4000MB.

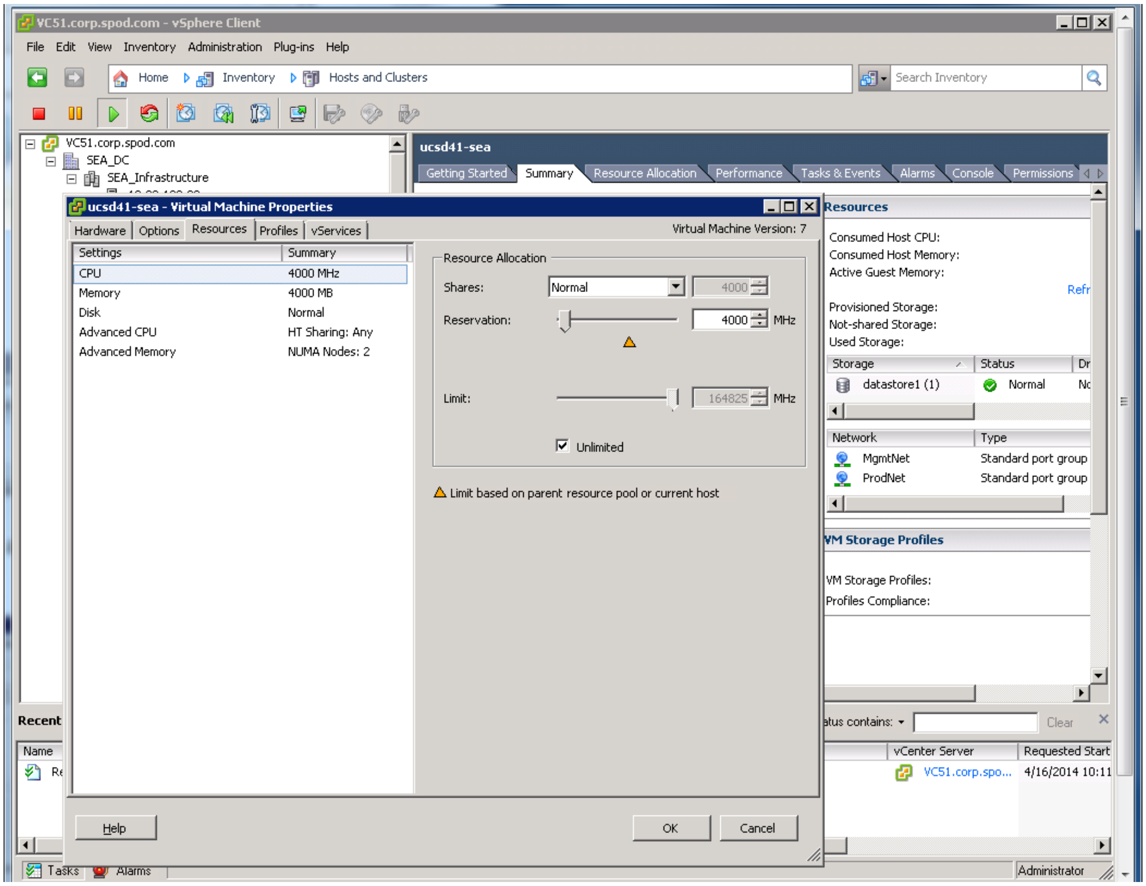

VC51.corp.spod.com - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home Inventory Hosts and Clusters
Search Inventory
VC51.corp.spod.com
SEA_DC
SEA_Infrastructure
ucsd41-sea
Getting Started Summary Resource Allocation Performance Tasks & Events Alarms Console Permissions
Recent
Name
Tasks Alarms
Administrator
Help OK Cancel

ucsd41-sea - Virtual Machine Properties

Hardware | Options | Resources | Profiles | vServices

Virtual Machine Version: 7

Settings	Summary
CPU	4000 MHz
Memory	4000 MB
Disk	Normal
Advanced CPU	HT Sharing: Any
Advanced Memory	NUMA Nodes: 2

Resource Allocation

Shares: Normal 4000

Reservation: 4000 MHz

Limit: 164825 MHz

☑ Unlimited

⚠ Limit based on parent resource pool or current host

Resources

Consumed Host CPU:
Consumed Host Memory:
Active Guest Memory:

Provisioned Storage:
Not-shared Storage:
Used Storage:

Storage	Status	Dr
datastore1 (1)	Normal	Nc

Network	Type
MgmtNet	Standard port group
ProdNet	Standard port group

VM Storage Profiles

VM Storage Profiles:
Profiles Compliance:

vCenter Server	Requested Start
VC51.corp.spo...	4/16/2014 10:11

4. Click the Options tab, and select VMware Tools, then click Synchronize guest time with host, then click OK to save the changes.

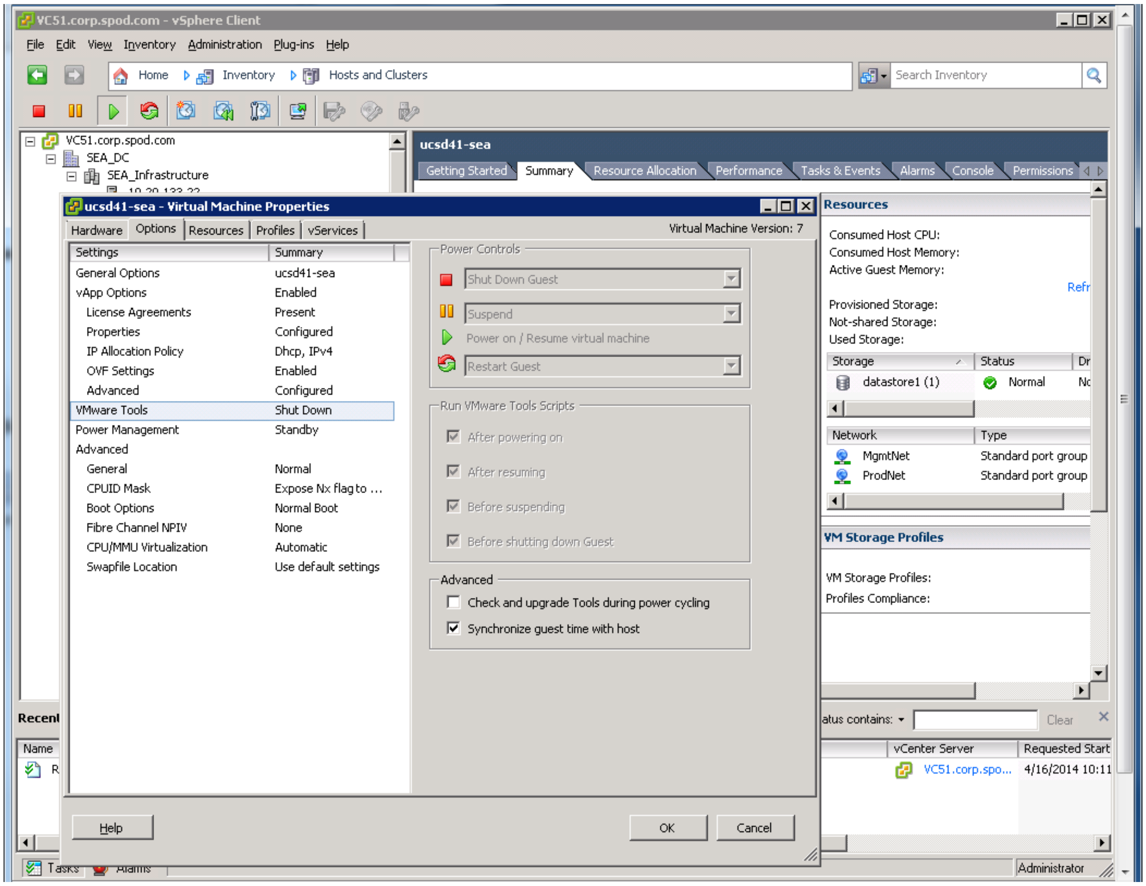

ucsd41-sea - Virtual Machine Properties

Hardware | Options | Resources | Profiles | vServices

Virtual Machine Version: 7

Settings	Summary
General Options	ucsd41-sea
vApp Options	Enabled
License Agreements	Present
Properties	Configured
IP Allocation Policy	Dhcp, IPv4
OVF Settings	Enabled
Advanced	Configured
VMware Tools	Shut Down
Power Management	Standby
Advanced	
General	Normal
CPUID Mask	Expose Nx flag to ...
Boot Options	Normal Boot
Fibre Channel NPIV	None
CPU/MMU Virtualization	Automatic
Swapfile Location	Use default settings

Power Controls

- Shut Down Guest
- Suspend
- Power on / Resume virtual machine
- Restart Guest

Run VMware Tools Scripts

☑ After powering on

☑ After resuming

☑ Before suspending

☑ Before shutting down Guest

Advanced

☐ Check and upgrade Tools during power cycling

☑ Synchronize guest time with host

Resources

Consumed Host CPU:
Consumed Host Memory:
Active Guest Memory:

Provisioned Storage:
Not-shared Storage:
Used Storage:

Storage	Status
datastore1 (1)	Normal

Network	Type
MgmtNet	Standard port group
ProdNet	Standard port group

VM Storage Profiles

VM Storage Profiles:
Profiles Compliance:

vCenter Server	Requested Start
VC51.corp.spo...	4/16/2014 10:11

5. Right-click the CUCSD-4.1.0.0 VM, select Power, Power on

6. Right-click the CUCSD-4.1.0.0 VM, select Open Console to configure Cisco UCS Director. Wait for the boot script to run to help you configure a static IP.

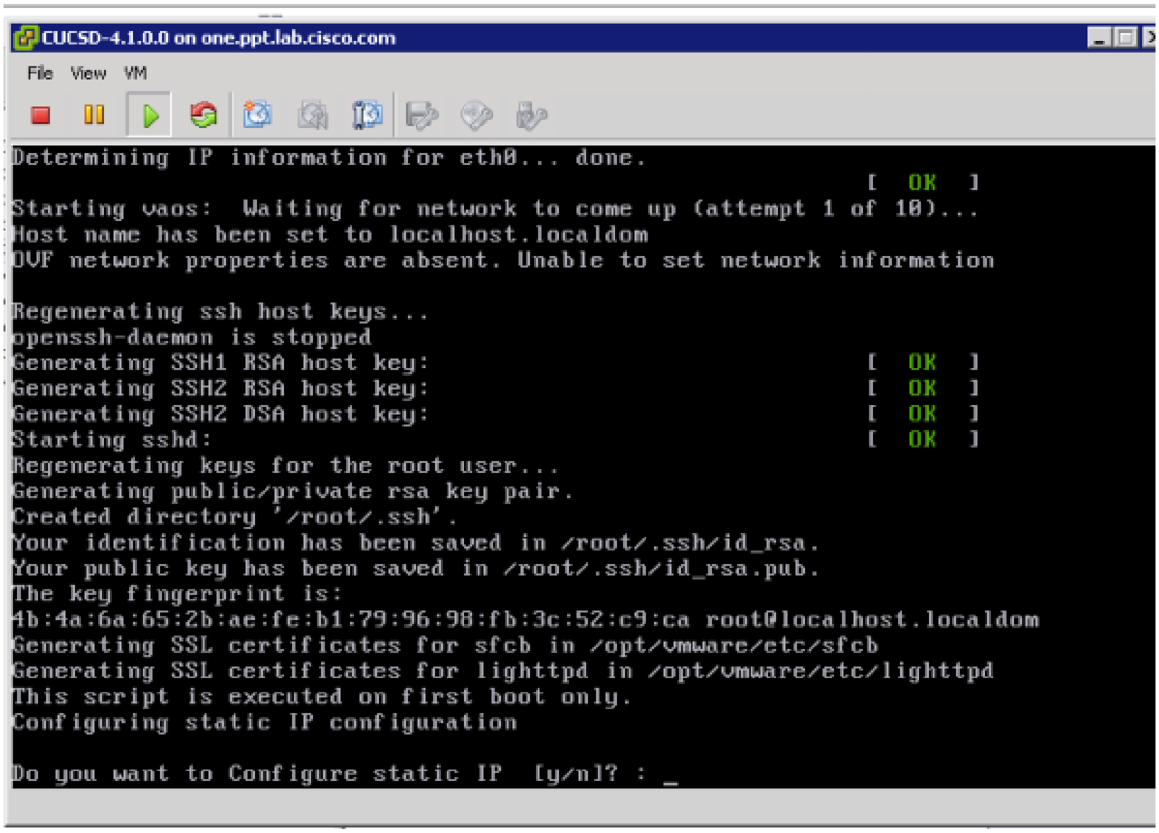

```
Determining IP information for eth0... done.
                                                           [  OK  ]
Starting vaos:  Waiting for network to come up (attempt 1 of 10)...
Host name has been set to localhost.localdom
OVF network properties are absent. Unable to set network information

Regenerating ssh host keys...
openssh-daemon is stopped
Generating SSH1 RSA host key:                              [  OK  ]
Generating SSH2 RSA host key:                              [  OK  ]
Generating SSH2 DSA host key:                              [  OK  ]
Starting sshd:                                             [  OK  ]
Regenerating keys for the root user...
Generating public/private rsa key pair.
Created directory '/root/.ssh'.
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
4b:4a:6a:65:2b:ae:fe:b1:79:96:98:fb:3c:52:c9:ca root@localhost.localdom
Generating SSL certificates for sfcb in /opt/vmware/etc/sfcb
Generating SSL certificates for lighttpd in /opt/vmware/etc/lighttpd
This script is executed on first boot only.
Configuring static IP configuration

Do you want to Configure static IP  [y/n]? : _
```

7. For "do you want to configure static ip?" Enter n for no as we will configure the IP after boot. The boot process will take a few minutes.

8. After boot completes, you will see a login screen. (A DHCP assigned address might exist if you have DHCP running on this subnet.) Hit Enter to select login. Log in as ID "shelladmin" and password of "changeme"

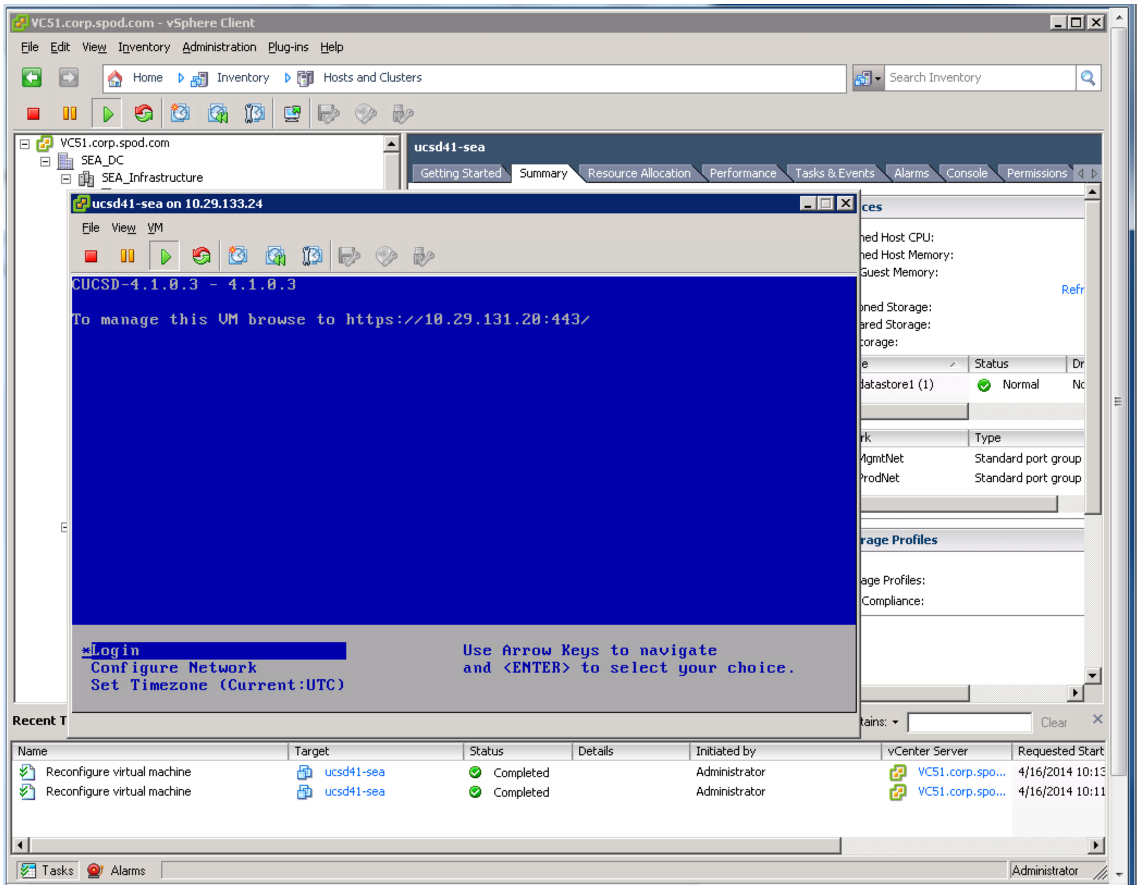
VC51.corp.spod.com - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home Inventory Hosts and Clusters
Search Inventory
VC51.corp.spod.com
SEA_DC
SEA_Infrastructure
ucsd41-sea
Getting Started Summary Resource Allocation Performance Tasks & Events Alarms Console Permissions
ucsd41-sea on 10.29.133.24
File View VM
Recent T
Clear
Tasks Alarms Administrator

```
CUCSD-4.1.0.3 - 4.1.0.3

To manage this VM browse to https://10.29.131.20:443/

*Login                              Use Arrow Keys to navigate
 Configure Network                  and <ENTER> to select your choice.
 Set Timezone (Current:UTC)
```

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start
Reconfigure virtual machine	ucsd41-sea	Completed		Administrator	VC51.corp.spo...	4/16/2014 10:13
Reconfigure virtual machine	ucsd41-sea	Completed		Administrator	VC51.corp.spo...	4/16/2014 10:11

9. At the Select prompt, input 14 to Configure Network Interface.

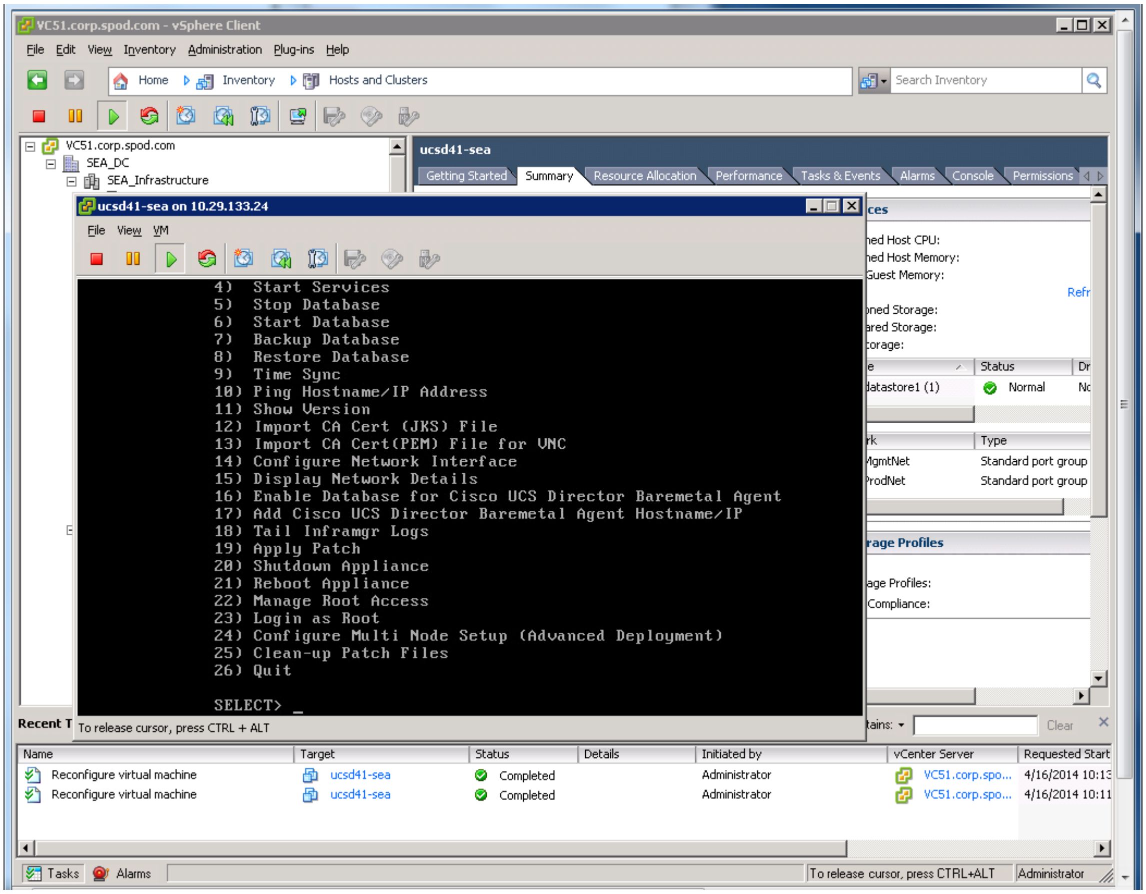
VC51.corp.spod.com - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home Inventory Hosts and Clusters
Search Inventory
VC51.corp.spod.com
SEA_DC
SEA_Infrastructure
ucsd41-sea
Getting Started Summary Resource Allocation Performance Tasks & Events Alarms Console Permissions
ucsd41-sea on 10.29.133.24
File View VM

```
4) Start Services
5) Stop Database
6) Start Database
7) Backup Database
8) Restore Database
9) Time Sync
10) Ping Hostname/IP Address
11) Show Version
12) Import CA Cert (JKS) File
13) Import CA Cert(PEM) File for VNC
14) Configure Network Interface
15) Display Network Details
16) Enable Database for Cisco UCS Director Baremetal Agent
17) Add Cisco UCS Director Baremetal Agent Hostname/IP
18) Tail Inframgr Logs
19) Apply Patch
20) Shutdown Appliance
21) Reboot Appliance
22) Manage Root Access
23) Login as Root
24) Configure Multi Node Setup (Advanced Deployment)
25) Clean-up Patch Files
26) Quit

SELECT> _
```

To release cursor, press CTRL + ALT

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start
Reconfigure virtual machine	ucsd41-sea	Completed		Administrator	VC51.corp.spo...	4/16/2014 10:13
Reconfigure virtual machine	ucsd41-sea	Completed		Administrator	VC51.corp.spo...	4/16/2014 10:11

10. For "Do you want to configure DHCP/Static IP?" Enter S for static.
11. Enter Eth0 for the interface you will configure.
12. Enter y for question if you want to configure Static IP for eth0.
13. Enter the IP, Netmask, Gateway and DNS server. Note the UCS Director server has dual interfaces with one leg on a routable network (vlan 133) and another (eth1) on a private production vlan (131).
14. Review the information and enter y to continue
15. Hit Enter to return to menu and input 1 to change the password
16. Hit Enter to return to main menu.
17. Enter 26 to quit. Notice the web URL to connect to https://<assigned IP>:443
18. For some browsers you may need to add the web URL to trusted sites to display correctly. Open the browser, and input the URL to UCSD. For I.E, click Tools, Internet Options, Security tab, Trusted Sites, Sites, and the address for your UCSD system and hit Add then Close. Hit F5 to refresh browser

Configuring the Admin Account

1. Connect to the URL for your CUSCSD system via the IP address you assigned.
2. Login with user "admin" and password of "admin" and select Login.

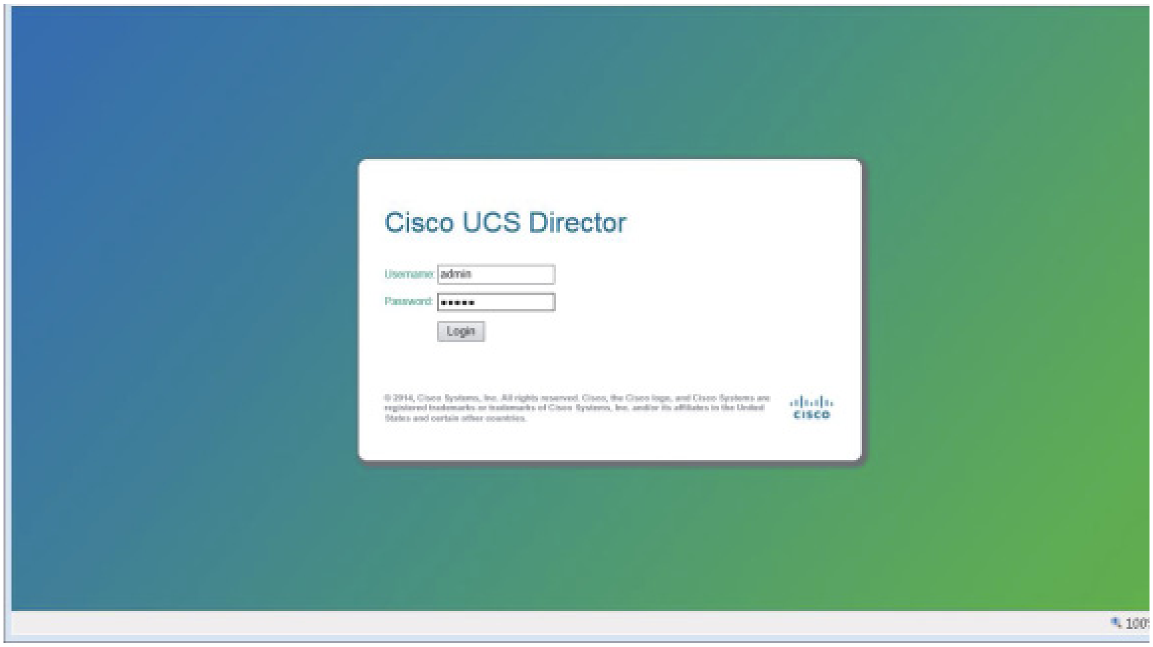

Cisco UCS Director

Username: admin

Password: •••••

3. Click OK to temporarily ignore the popup information message for login profile.
4. Click Administration on the menu bar and choose Users and Groups
5. Click the Login Users tab, highlight admin, and click Change Password to input new password then click Save, then OK.

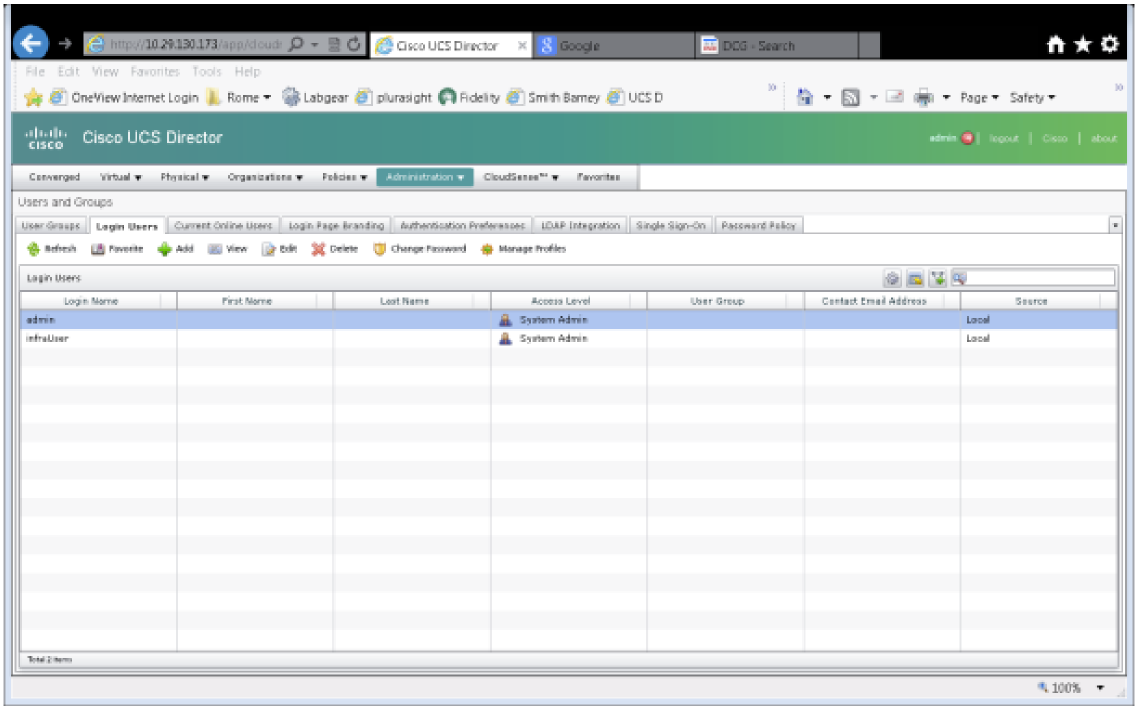

Users and Groups

Login Users

Login Name	First Name	Last Name	Access Level	User Group	Contact Email Address	Source
admin			System Admin			Local
infraUser			System Admin			Local

Total 2 items

6. With 'admin' still selected, click on 'edit' and Input user email address, click Save, and then OK.
7. Go to Administration, System and select Mail Setup tab.

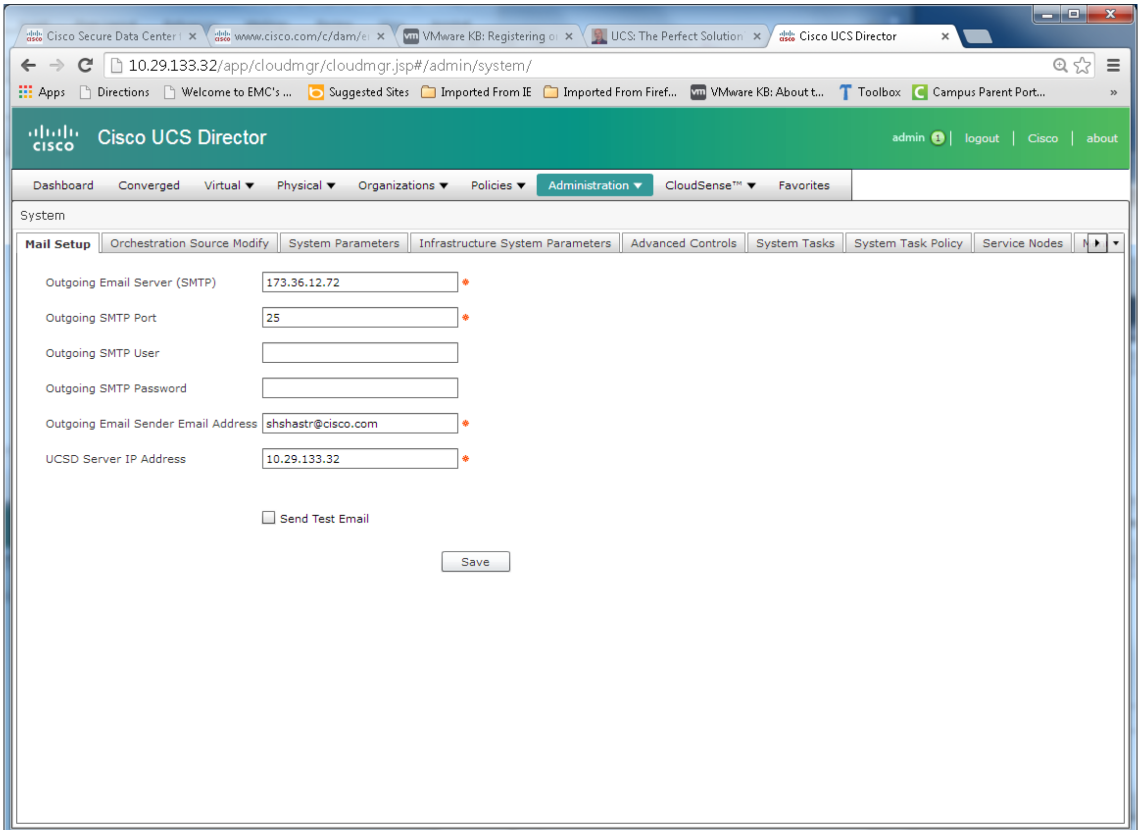

Outgoing Email Server (SMTP)	173.36.12.72 *
Outgoing SMTP Port	25 *
Outgoing SMTP User	
Outgoing SMTP Password	
Outgoing Email Sender Email Address	shshastr@cisco.com *
UCSD Server IP Address	10.29.133.32 *

□ Send Test Email

Save

8. Input SMTP server ip address or hostname if you have working DNS
9. Input correct SMTP port (25 is default)
10. Input Outgoing Email sender address.
11. Input the Server IP Address of the UCSD server
12. Click the Send Test Email box
13. Enter Test Email Address
14. Click the Save button and validate that you get a "Successfully update mail settings. Test email Succeeded." Message then click OK

Installing Licenses

1. Install the license by choosing Administration, License, then the License Keys tab for Update License.

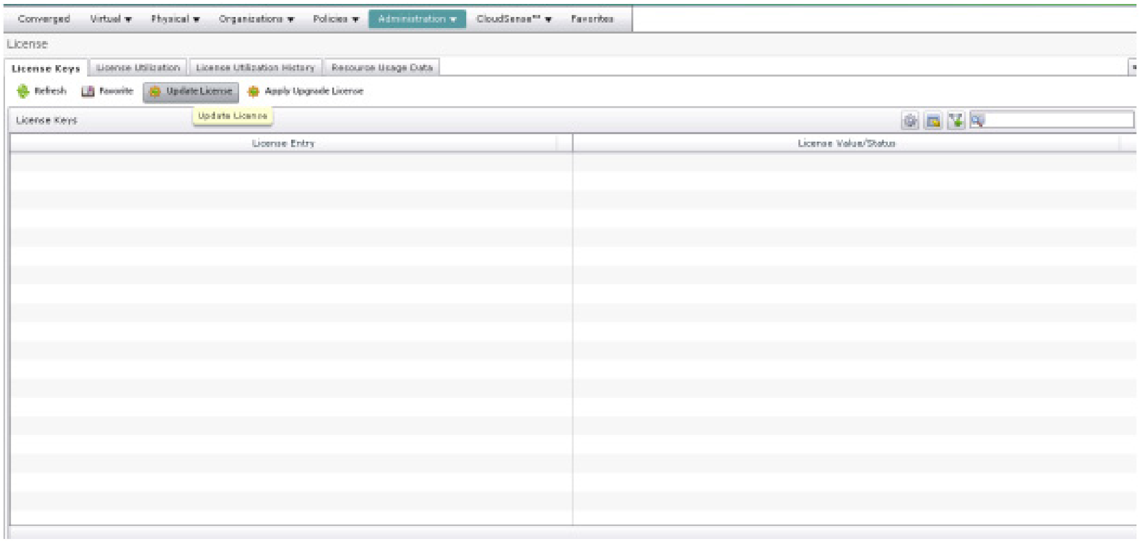

2. Click Browse and choose the license file that you received from Cisco. Select the file and hit Open then select Upload.

3. After upload complete message, click OK, and then Submit.

4. Select the License Keys tab, click Refresh and validate you have a minimum of the base license.

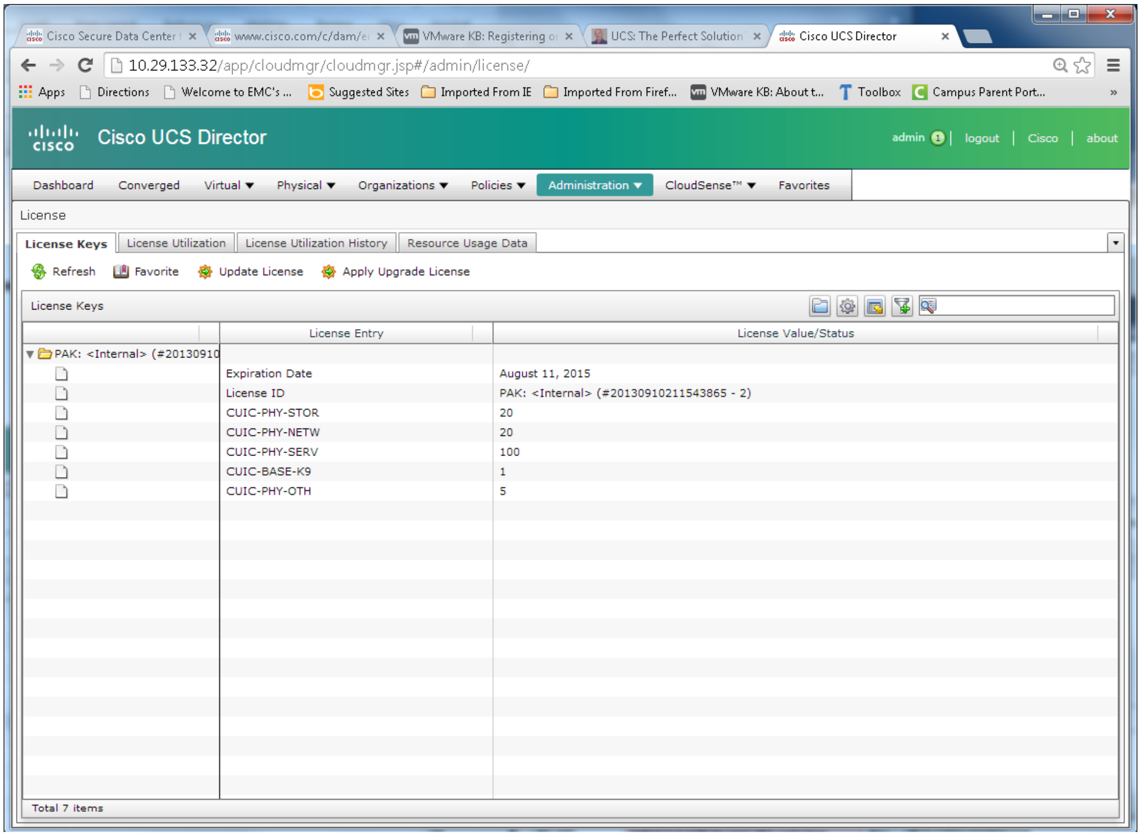

	License Entry	License Value/Status
▼ PAK: <Internal> (#20130910		
	Expiration Date	August 11, 2015
	License ID	PAK: <Internal> (#20130910211543865 - 2)
	CUIC-PHY-STOR	20
	CUIC-PHY-NETW	20
	CUIC-PHY-SERV	100
	CUIC-BASE-K9	1
	CUIC-PHY-OTH	5

Total 7 items

Create a Converged Pod

A Pod is a collection of physical and virtual resources that can be manages together. We will create a site and a pod that will contain our VSPEX resources.

1. Add a site name by clicking Administration, Physical Accounts, then the Site Management tab.
2. Click Add and input site name and contact name then click submit. Hit OK to successfully added message.
3. Click Converged from the main menu and click Add

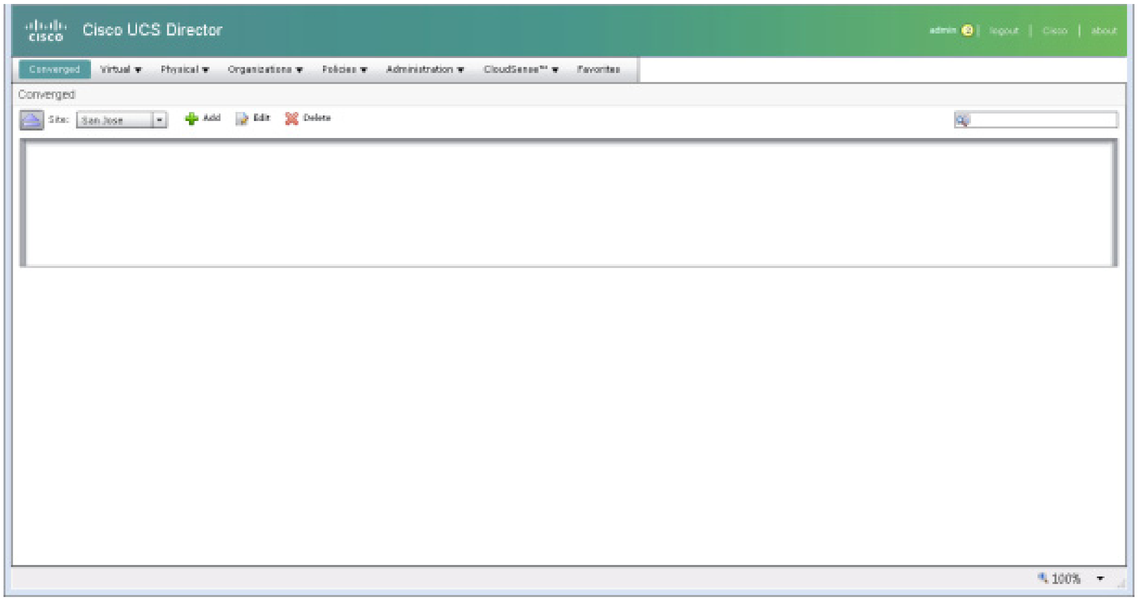

Site: San Jose

4. Enter a Pod Name, Site and select type VSPEX. Click Add.

Adding EMC VNX Storage

1. Select Administration and then Physical Accounts.

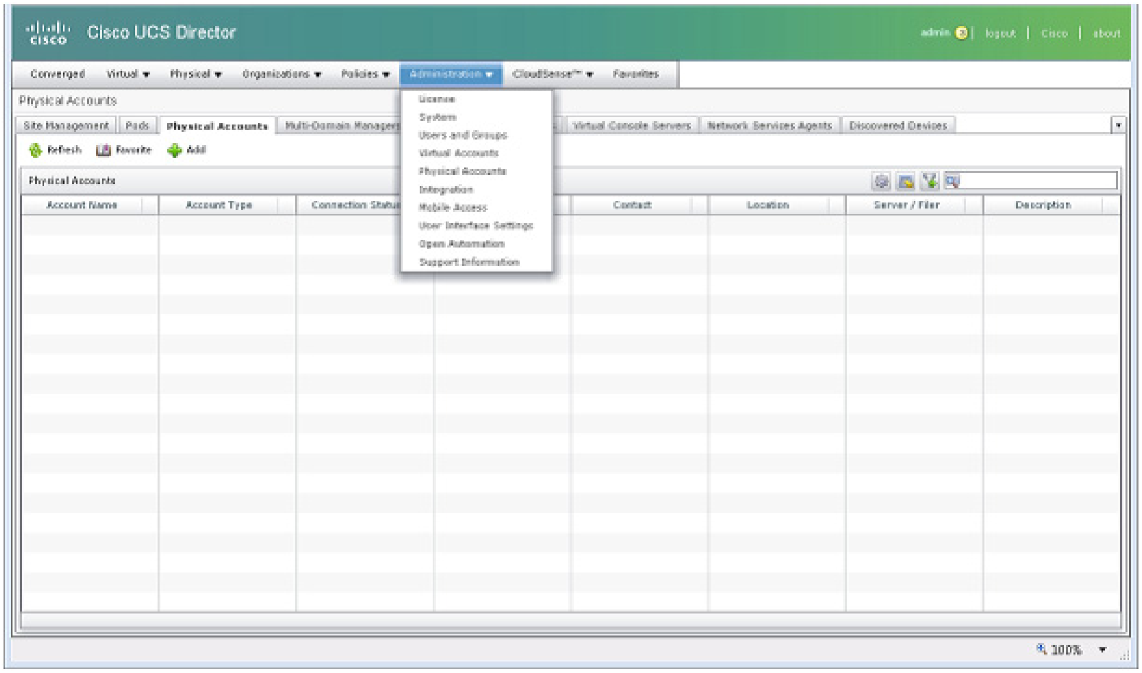

2. Click on Add.

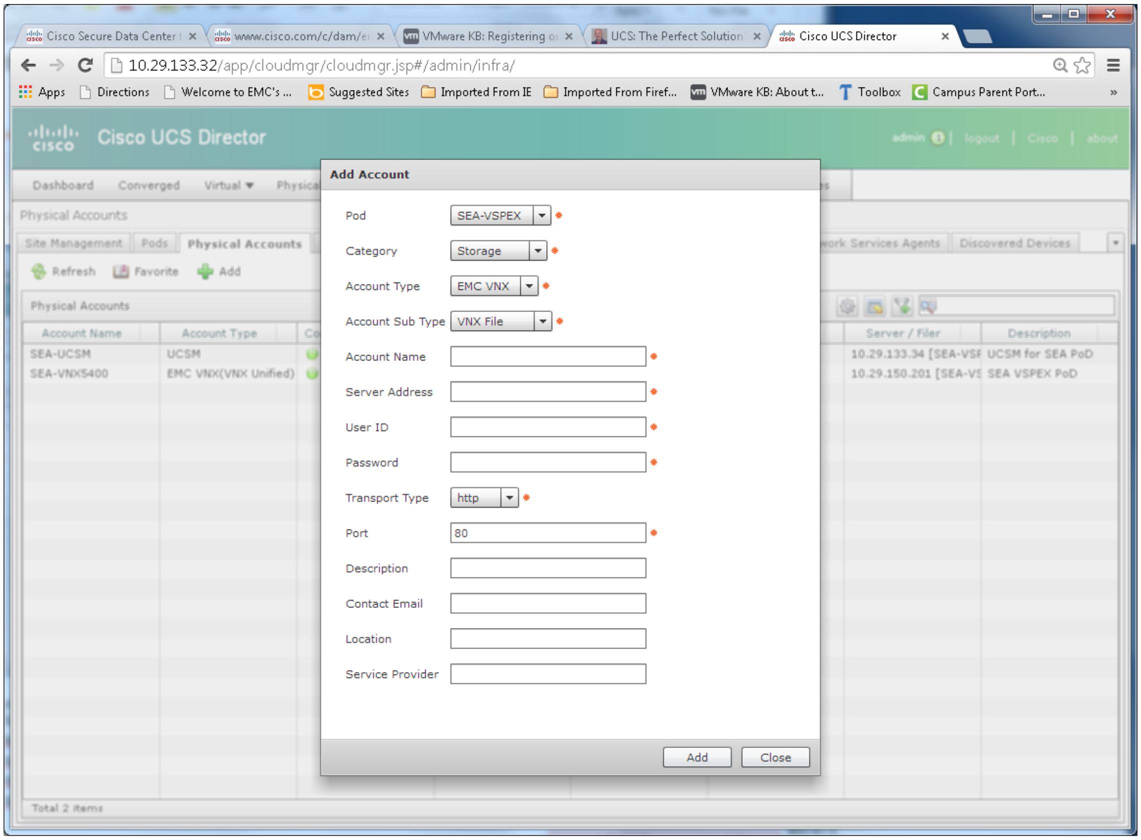

10.29.133.32/app/cloudmgr/cloudmgr.jsp#/admin/infra/

Account Name	Account Type	Server / Filer	Description
SEA-UCSM	UCSM	10.29.133.34 [SEA-VSF	UCSM for SEA PoD
SEA-VNX5400	EMC VNX(VNX Unified)	10.29.150.201 [SEA-VS	SEA VSPEX PoD

Add Account

Pod	SEA-VSPEX *
Category	Storage *
Account Type	EMC VNX *
Account Sub Type	VNX File *
Account Name	*
Server Address	*
User ID	*
Password	*
Transport Type	http *
Port	80 *
Description	
Contact Email	
Location	
Service Provider	

Add Close

3. Make sure to select the Pod Name for the Pod you created in the previous step

4. Select Storage for Category Type
5. Select EMC VNX for the Account Type.
6. Enter Account name, VNX data mover IP address.
7. Enter login and password for VNX IP.
8. Select HTTPS for the Transport Type.
9. Optionally add a Description, Contact Email, Location and Service Provider.
10. Click Add, then OK.
11. Once the account has been added, select the newly added account from the list and choose Test Connection
12. A window will appear that displays Connection Successful. Click Close.

Adding Cisco Nexus Switches

In this section we will add our Cisco switches to our Pod. Repeat these steps for both of your Cisco Nexus 5000 Switches.

1. Select Administration and then Physical Accounts. Click the Manage Network Elements tab.

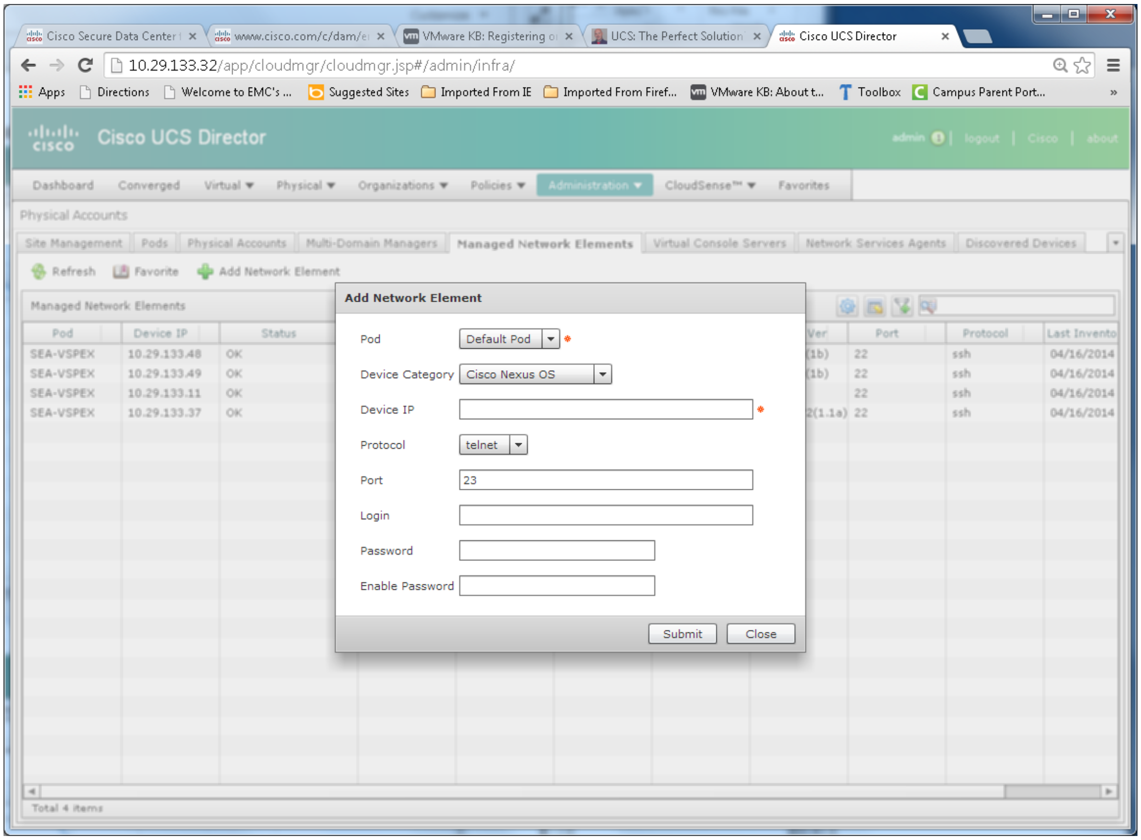

Pod	Device IP	Status	Ver	Port	Protocol	Last Invento
SEA-VSPEX	10.29.133.48	OK	(1b)	22	ssh	04/16/2014
SEA-VSPEX	10.29.133.49	OK	(1b)	22	ssh	04/16/2014
SEA-VSPEX	10.29.133.11	OK		22	ssh	04/16/2014
SEA-VSPEX	10.29.133.37	OK	2(1.1a)	22	ssh	04/16/2014

Add Network Element

Pod: Default Pod *
Device Category: Cisco Nexus OS
Device IP: *
Protocol: telnet
Port: 23
Login:
Password:
Enable Password:

2. Click on Add Network Element.
3. Select Pod you created.
4. Select Device Category Cisco Nexus OS.

5. Enter switch management IP for the Device IP.
6. Select SSH for the Protocol.
7. Enter 22 for the Port.
8. Enter admin for the Login.
9. Enter switch admin password for Password.
10. Enter password again for Enable Password field.
11. Click Submit. (Note: It can take a few minutes to complete this operation)
12. When the account has been added, repeat for the other switch.
13. Select the newly added switches and choose Test Connection. Click Close.

Add VMware Virtual Account

The VCenter server needs to be added to our converged Pod in order for Cisco UCS Director to manage our VMware infrastructure.

1. Click Administration, Virtual Accounts and then the Virtual Accounts tab.
2. Select Add.
3. On the Add Cloud popup, Select cloud type VMware.
4. A second add screen will appear with VMware selected for cloud type.

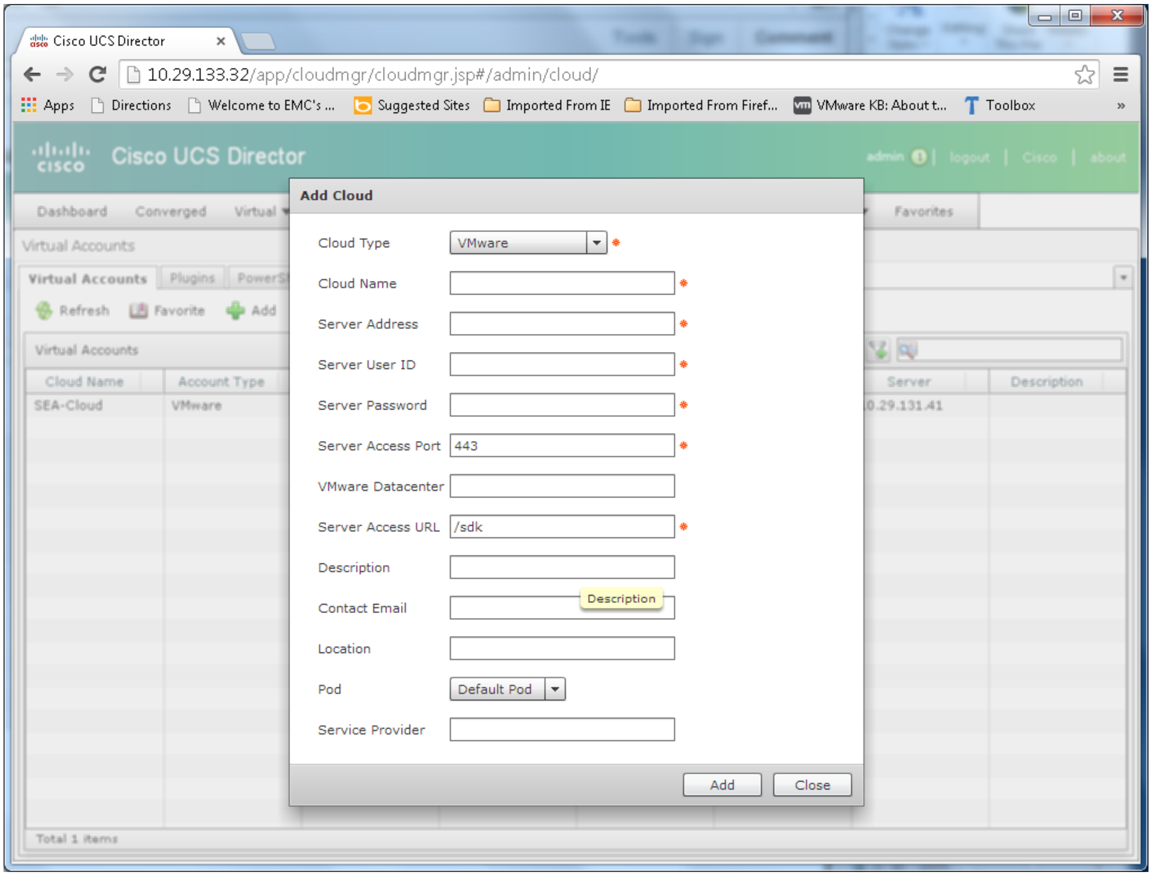

Add Cloud

Cloud Type: VMware *
Cloud Name: *
Server Address: *
Server User ID: *
Server Password: *
Server Access Port: 443 *
VMware Datacenter:
Server Access URL: /sdk *
Description:
Contact Email:
Location:
Pod: Default Pod
Service Provider:

Add | Close

5. Input cloud name. (Example: SEA-Cloud)
6. input vCenter server IP address.
7. Input vcenter login and password for connectivity.
8. Leave the server access URL set to /sdk.
9. For Pod, input the pod you created (Example SEA-VSPEX).
10. Click Add.
11. It can take a few minutes for Cisco UCS Director to complete the query of the Vcenter objects and the connection status to change to success. Highlight the account, and click Test Connectivity, click Close.

Create Local Users and Groups

With Cisco UCS Director, you can use local accounts and/or Windows Active directory accounts. Here we will go through steps necessary to create a group and users within the group. You can use these for production or test purposes prior to rollout.

Adding PNSC

1. Select Administration, Multi-Domain Managers and then "Add".

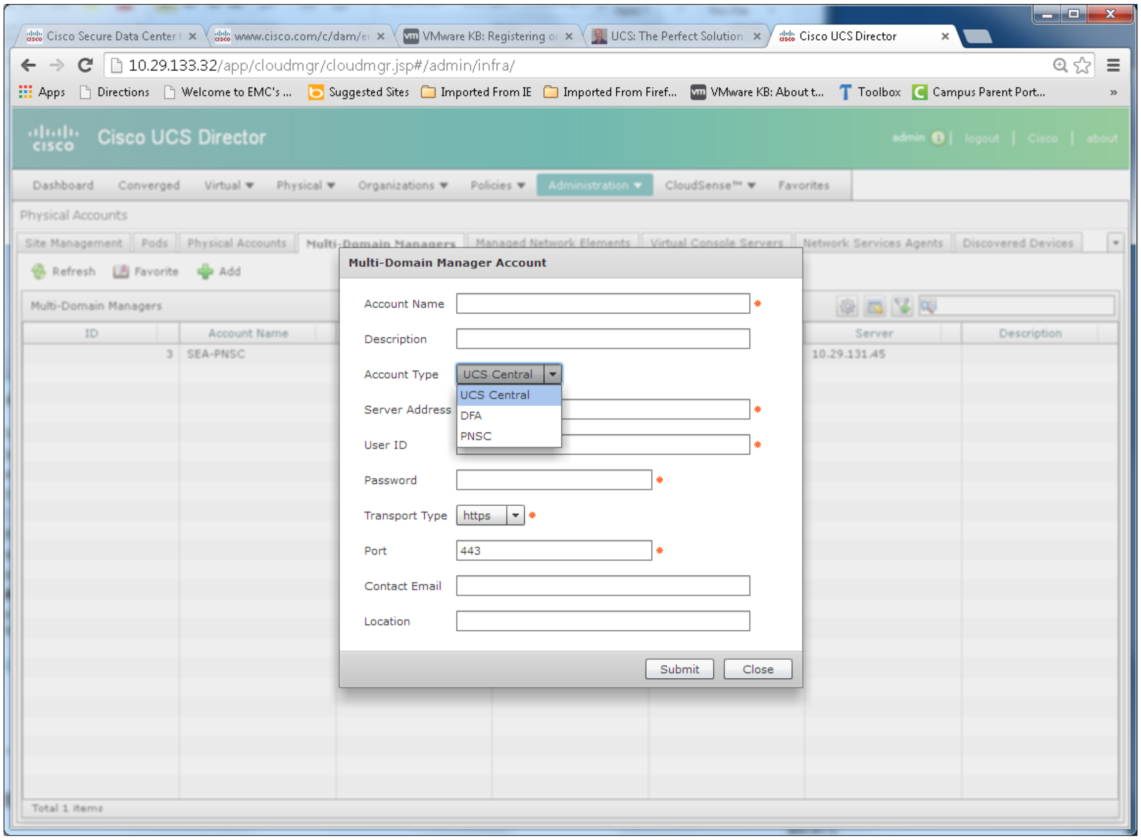

Multi-Domain Manager Account

Account Name
Description
Account Type: UCS Central (UCS Central / DFA / PNSC)
Server Address
User ID
Password
Transport Type: https
Port: 443
Contact Email
Location

Submit | Close

2. After providing an account name, pick a type of PNSC from the drop-down.
3. Provide login and password for the PNSC management service.

4. Protocol of https with default port of 443 and hit submit.

Adding ASA 5585 Firewall Device

1. From Administration, Managed Network Elements tab, select Add.
2. In the pop-up, pick the pod you wish to add the ASA to with Device Category of ASA.

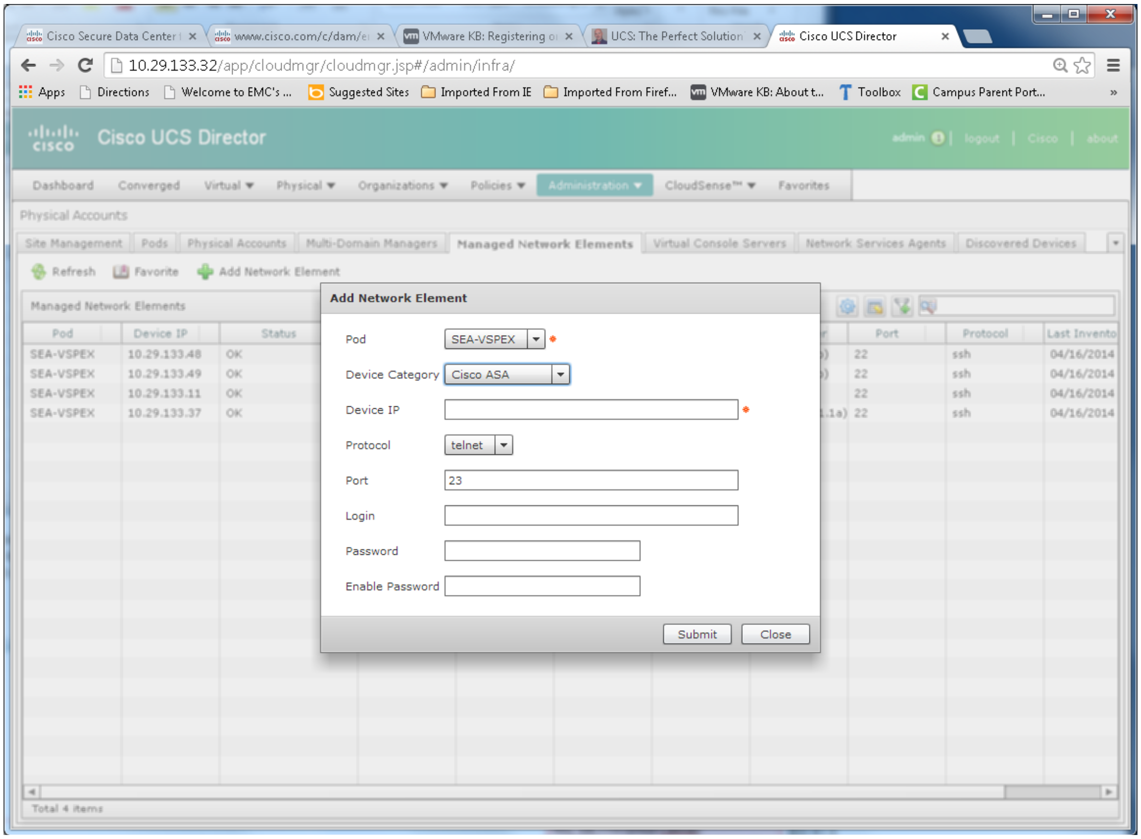

Pod	Device IP	Status	Port	Protocol	Last Invento
SEA-VSPEX	10.29.133.48	OK	22	ssh	04/16/2014
SEA-VSPEX	10.29.133.49	OK	22	ssh	04/16/2014
SEA-VSPEX	10.29.133.11	OK	22	ssh	04/16/2014
SEA-VSPEX	10.29.133.37	OK	22	ssh	04/16/2014

Add Network Element

Pod	SEA-VSPEX *
Device Category	Cisco ASA
Device IP	*
Protocol	telnet
Port	23
Login	
Password	
Enable Password	

3. Provide IP and login credentials and hit submit for the ASA to be added to Cisco UCS Director.

Adding Application Container Template

This construct allows for grouping of components including VM's in a three tiered application with load-balancers, VSG and other supported components for easy and quick deployment. The following are a few prerequisites before deploying this feature:

1. Upload gateway ova/ovf:
   a. First create a user with Service End-User type by selecting.
   b. Administration, Users and Groups and then Login Users.
   c. Select Add.

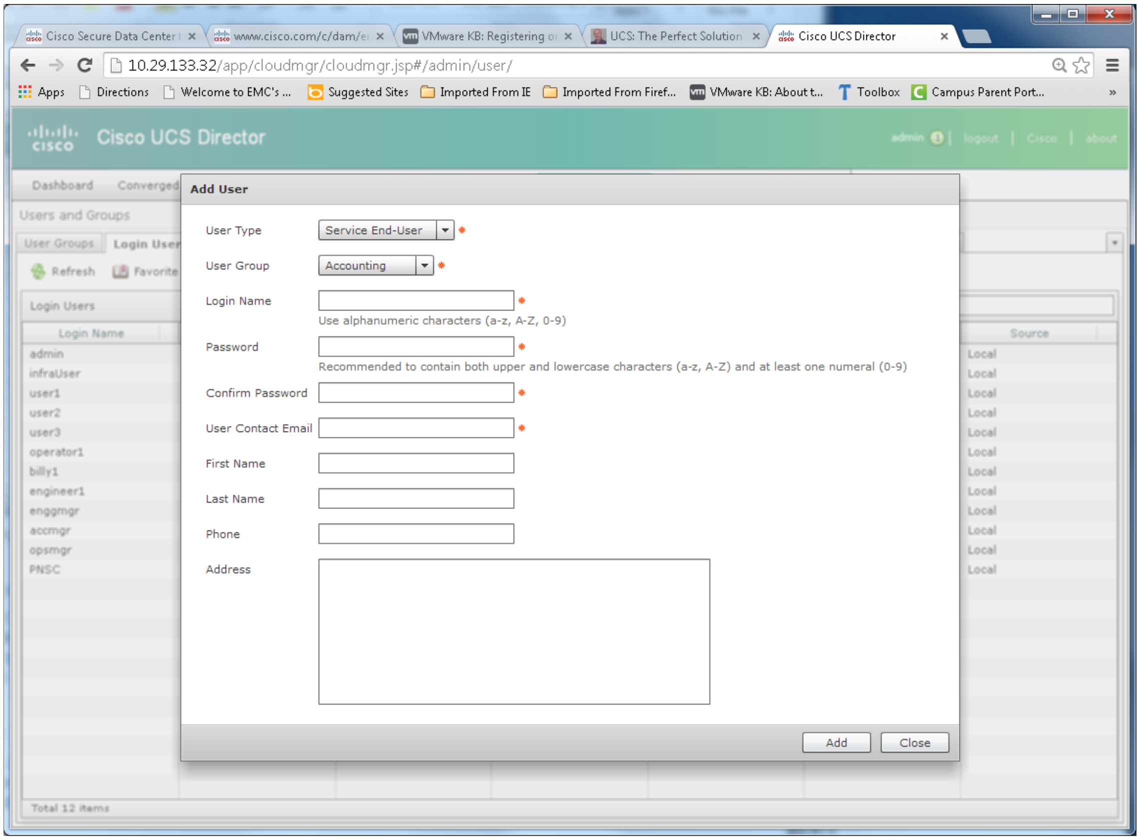

Add User

User Type: Service End-User *

User Group: Accounting *

Login Name: *
Use alphanumeric characters (a-z, A-Z, 0-9)

Password: *
Recommended to contain both upper and lowercase characters (a-z, A-Z) and at least one numeral (0-9)

Confirm Password: *

User Contact Email: *

First Name:

Last Name:

Phone:

Address:

d. Logout and log back in with the newly created user credentials for required access to upload files. Select Services, then click the Upload tab.

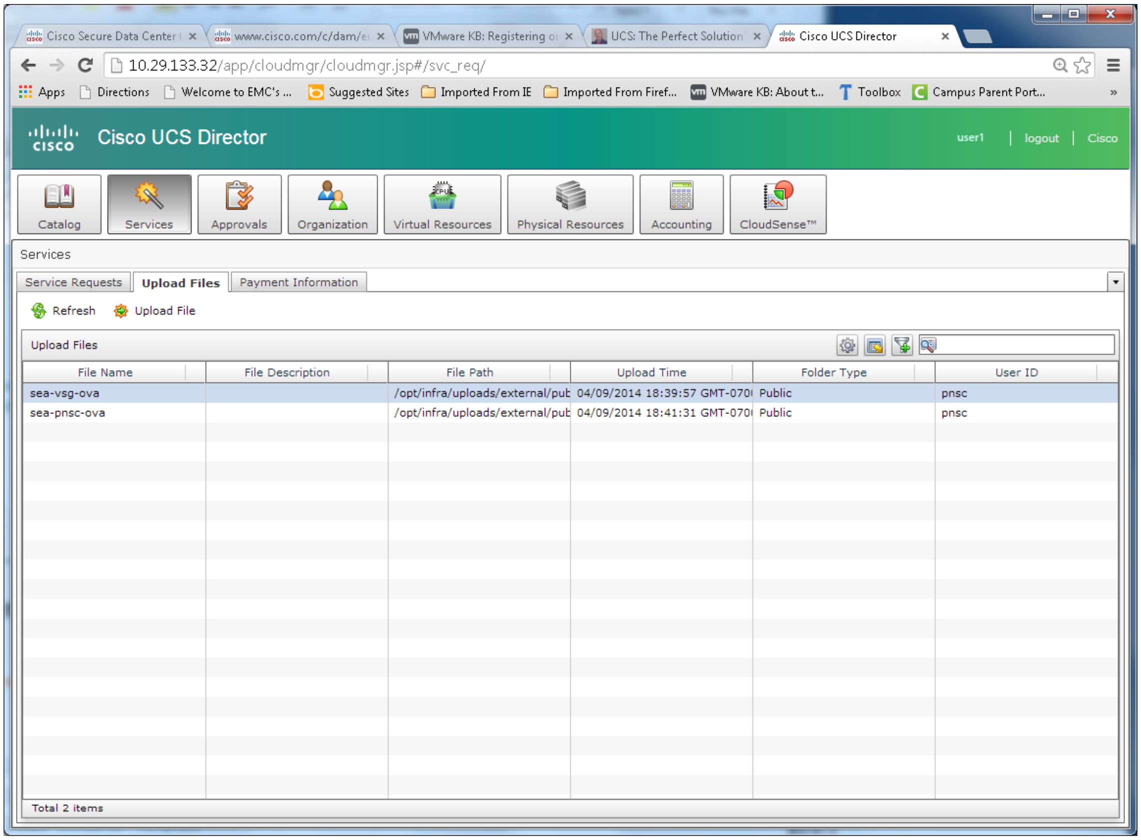

File Name	File Description	File Path	Upload Time	Folder Type	User ID
sea-vsg-ova		/opt/infra/uploads/external/pub	04/09/2014 18:39:57 GMT-070	Public	pnsc
sea-pnsc-ova		/opt/infra/uploads/external/pub	04/09/2014 18:41:31 GMT-070	Public	pnsc

Total 2 items

e. Select Upload File and leave the Folder type as Public for easy access. Provide a file name and browse to the location where the files (VSG and PNSC) reside and make the selection. Hit submit.

f. Perform this twice for each of the two files to be uploaded.

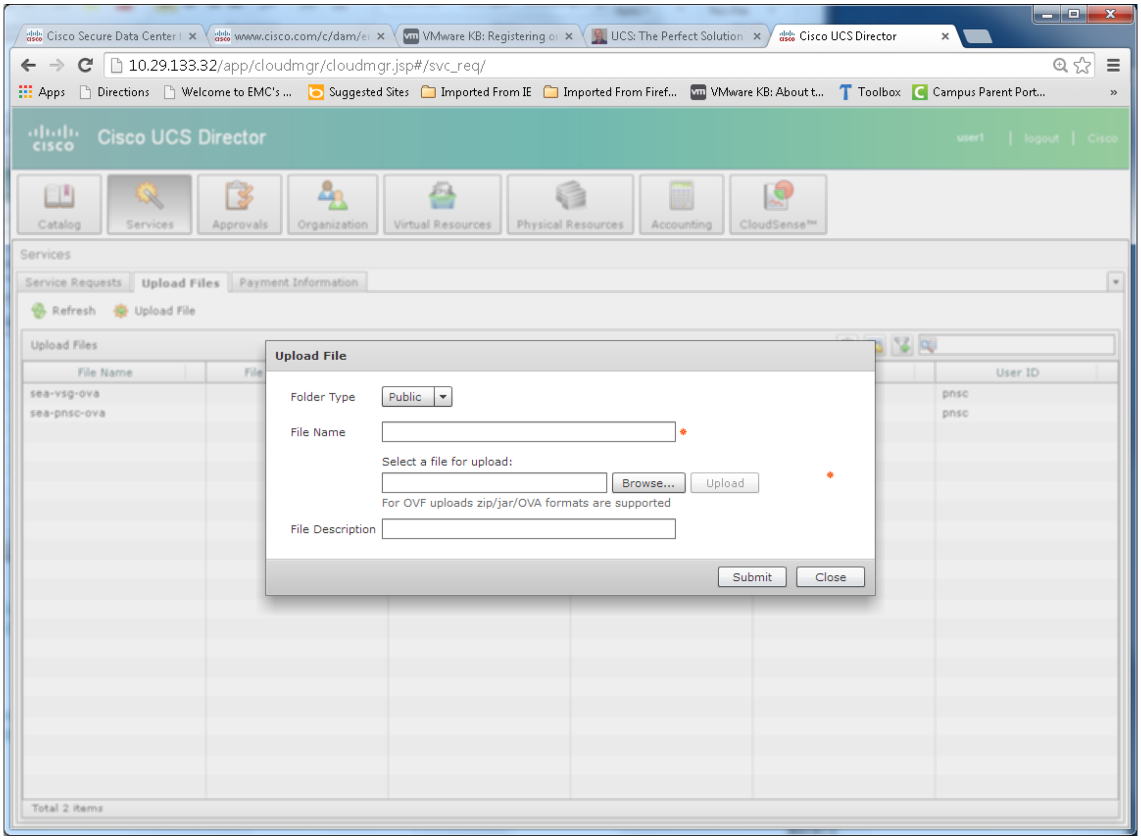

File Name	User ID
sea-vsg-ova	pnsc
sea-pnsc-ova	pnsc

Upload File

Folder Type: Public

File Name: *

Select a file for upload: *

Browse... Upload

For OVF uploads zip/jar/OVA formats are supported

File Description:

Submit Close

2. Create PNSC firewall policy:

   Use a firewall policy to enforce network traffic on a Cisco VSG. VSG is a firewall used inside an enclave for inter-VM access control. The policy engine uses the policy as a filter for network traffic received by the VSG.

   a. Select Physical, Network and then the PNSC account in the left and right panes to bring up the "view Details option as shown:

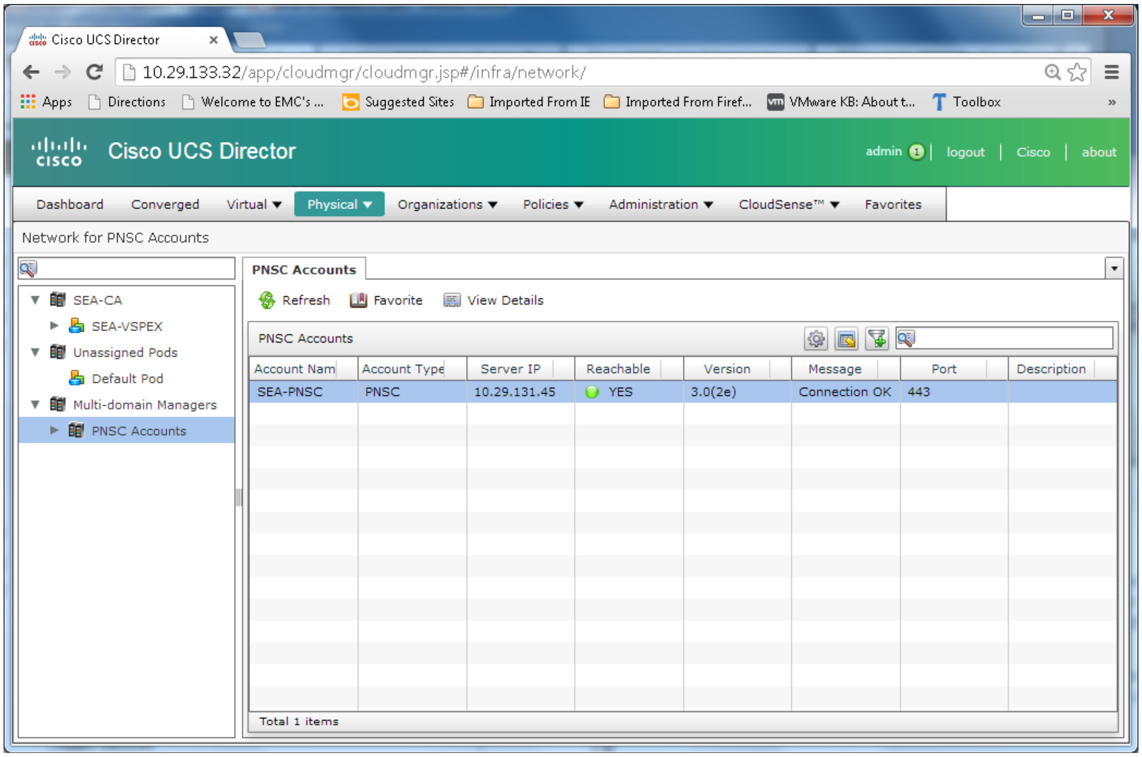

Account Nam	Account Type	Server IP	Reachable	Version	Message	Port	Description
SEA-PNSC	PNSC	10.29.131.45	● YES	3.0(2e)	Connection OK	443	

Total 1 items

b. Selecting the "View Details" option brings up the following screen.

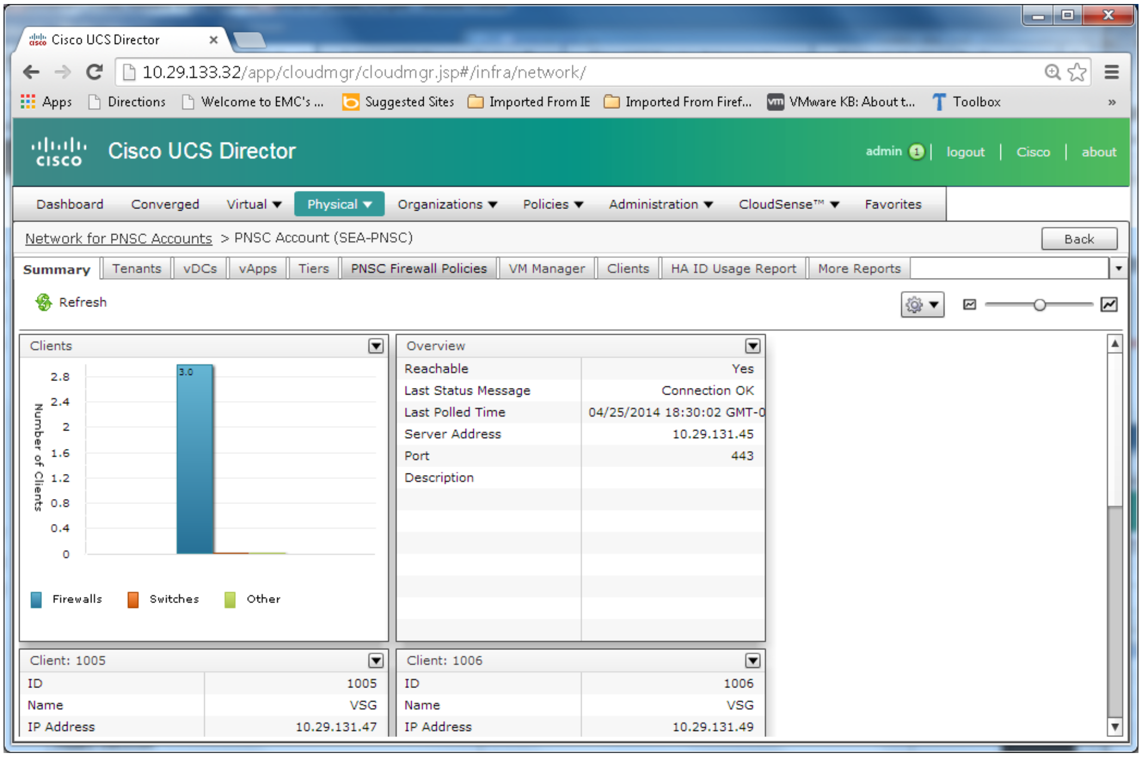

Reachable	Yes
Last Status Message	Connection OK
Last Polled Time	04/25/2014 18:30:02 GMT-0
Server Address	10.29.131.45
Port	443
Description	

Client: 1005

ID	1005
Name	VSG
IP Address	10.29.131.47

Client: 1006

ID	1006
Name	VSG
IP Address	10.29.131.49

c. Select "PNSC Firewall Policies" tab and then click Add.

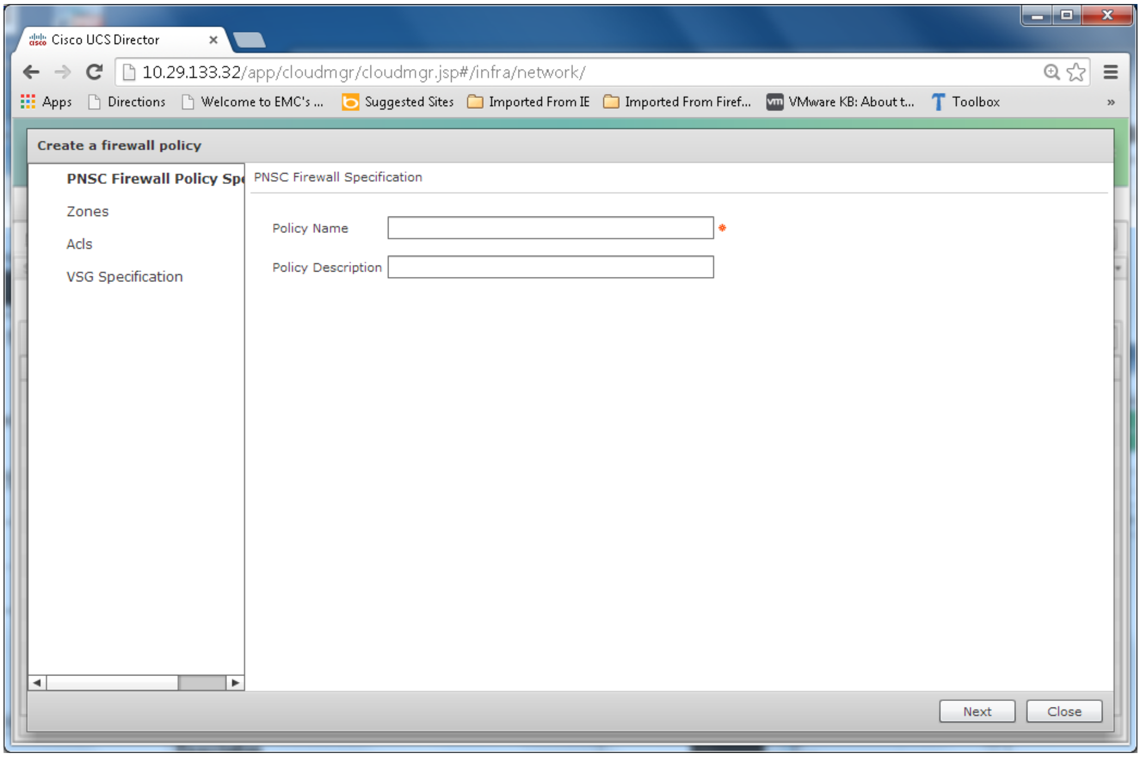

d. Enter a unique Policy Name and click "Next" followed by clicking the "+" to add a new PNSC zone. Select "+" again on this screen to add a new zone condition.

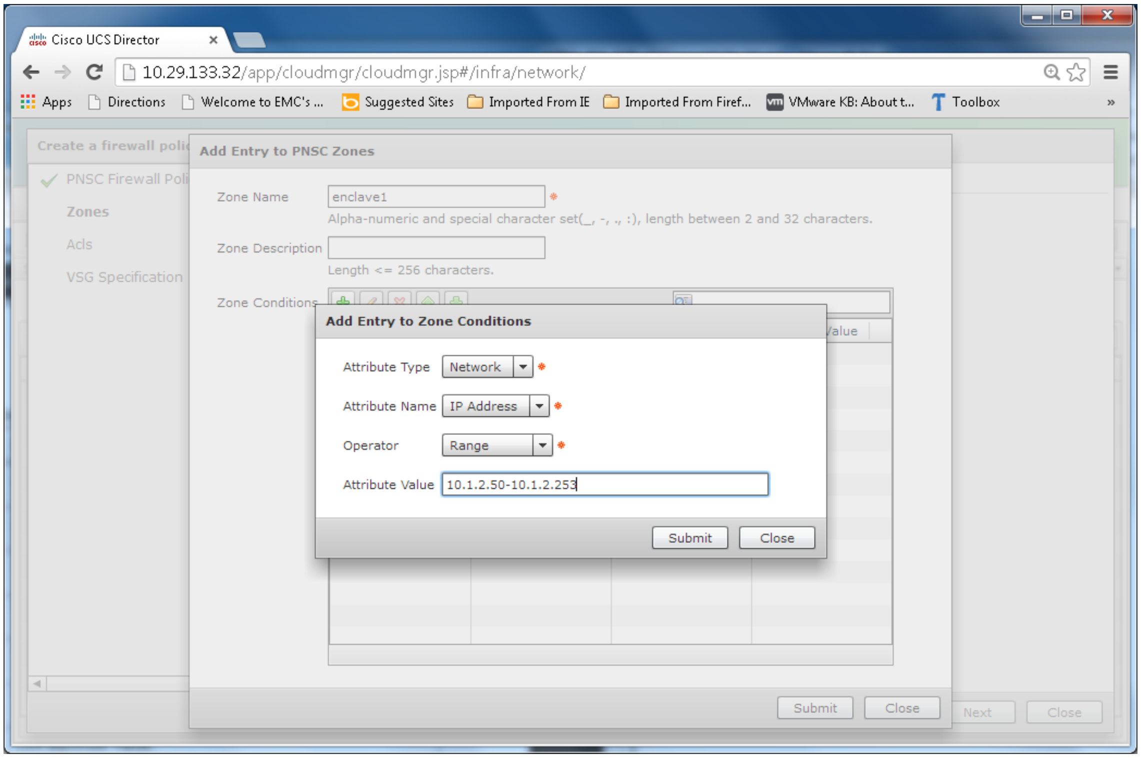

Zone Name: enclave1

Alpha-numeric and special character set(_, -, ., :), length between 2 and 32 characters.

Zone Description:

Length <= 256 characters.

Attribute Type: Network

Attribute Name: IP Address

Operator: Range

Attribute Value: 10.1.2.50-10.1.2.253

e. Click "Submit" twice. Click "Next", select "+" to add a new entry to the zone.

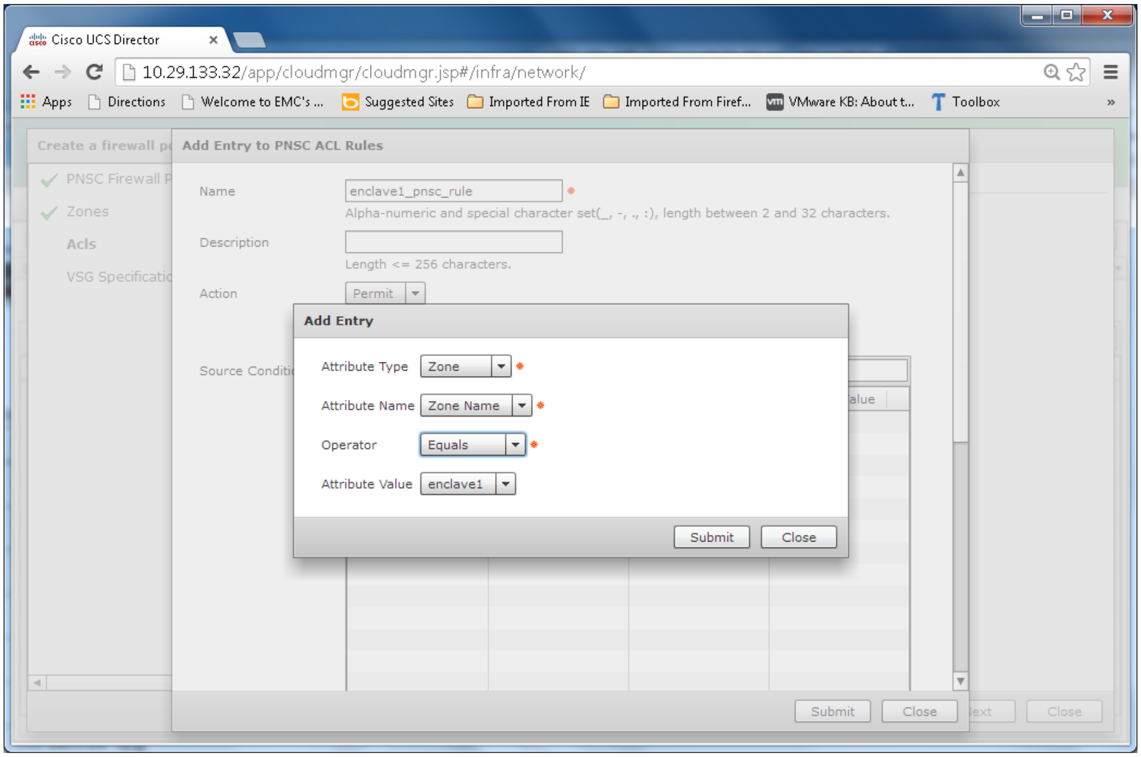

Add Entry to PNSC ACL Rules

Name: enclave1_pnsc_rule *
Alpha-numeric and special character set(_, -, ., :), length between 2 and 32 characters.

Description:
Length <= 256 characters.

Action: Permit

Add Entry

Attribute Type: Zone *

Attribute Name: Zone Name *

Operator: Equals *

Attribute Value: enclave1

f. Click "Submit" twice and click "Next". Provide necessary input and "submit".

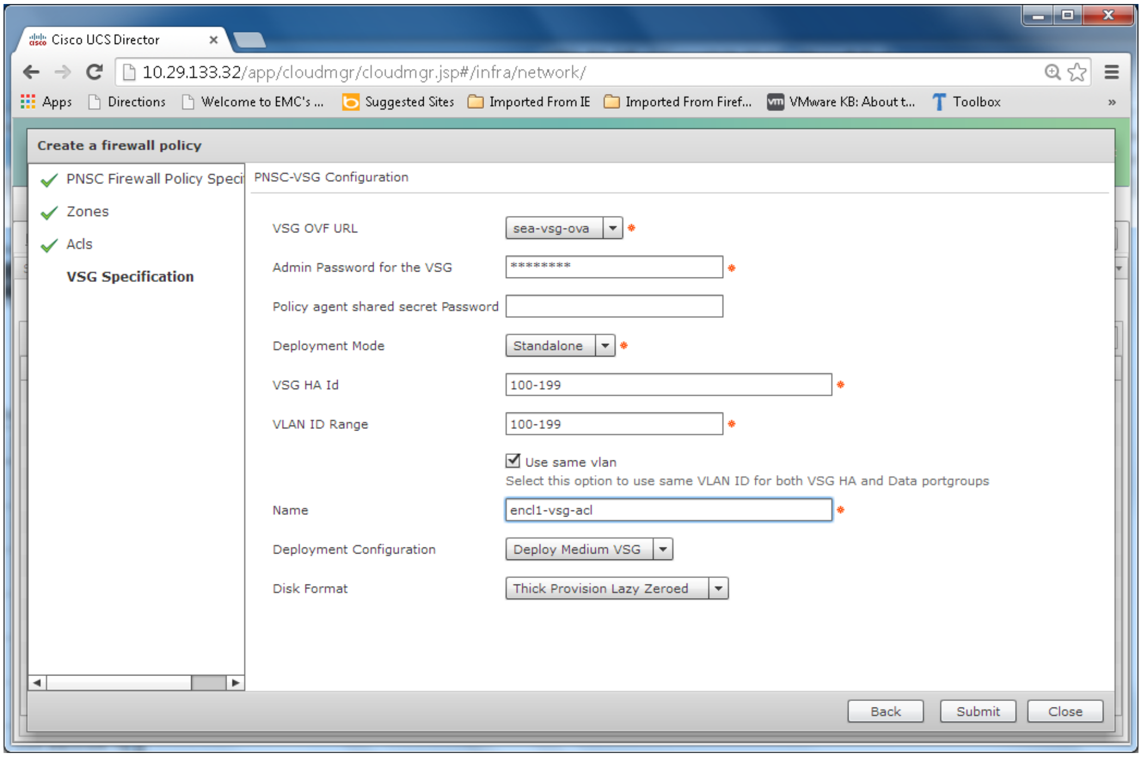

Create a firewall policy

✓ PNSC Firewall Policy Speci
✓ Zones
✓ Acls
VSG Specification

PNSC-VSG Configuration

VSG OVF URL: sea-vsg-ova *

Admin Password for the VSG: ******** *

Policy agent shared secret Password:

Deployment Mode: Standalone *

VSG HA Id: 100-199 *

VLAN ID Range: 100-199 *

☑ Use same vlan
Select this option to use same VLAN ID for both VSG HA and Data portgroups

Name: encl1-vsg-acl *

Deployment Configuration: Deploy Medium VSG

Disk Format: Thick Provision Lazy Zeroed

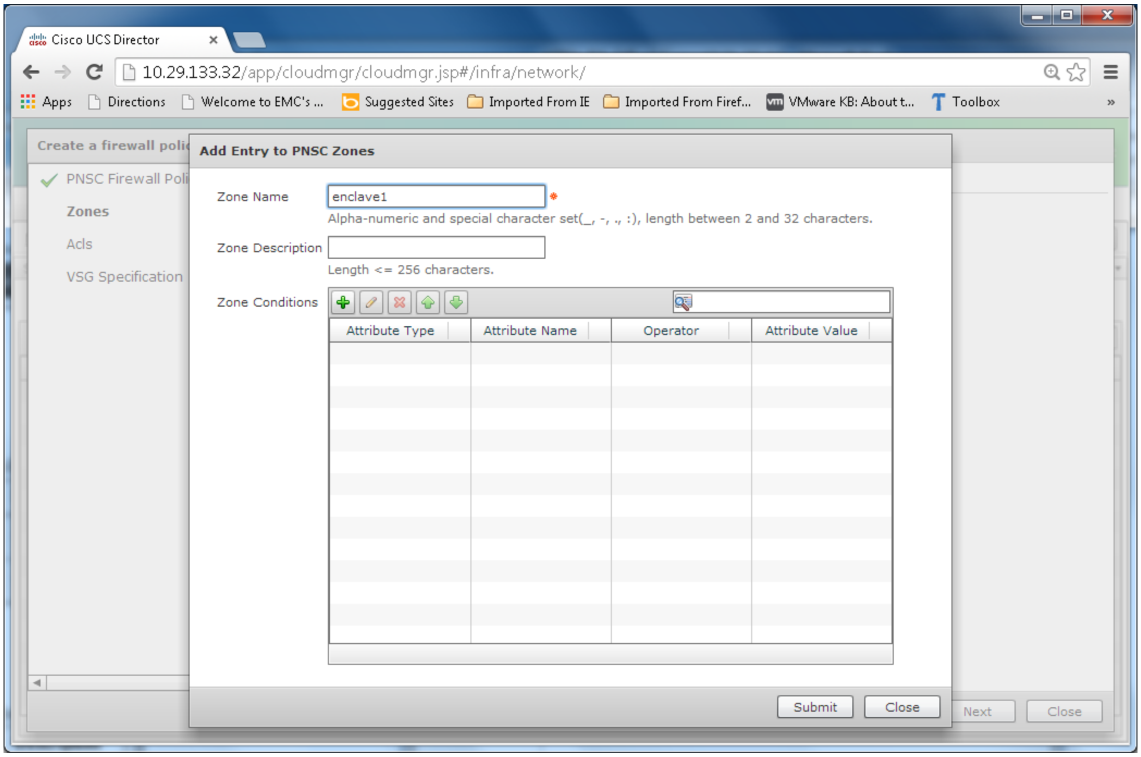

Add Entry to PNSC Zones

Zone Name: enclave1 *

Alpha-numeric and special character set(_, -, ., :), length between 2 and 32 characters.

Zone Description:

Length <= 256 characters.

Zone Conditions

Attribute Type	Attribute Name	Operator	Attribute Value

Submit Close

3. Create Virtual Infrastructure Policy:

   The virtual infrastructure policy defines which VM to use and what type of container to provision.

   a. Select Application Container from the Policy drop down. Select the Virtual Infrastructure Policy tab and then click "+" to add a new policy.

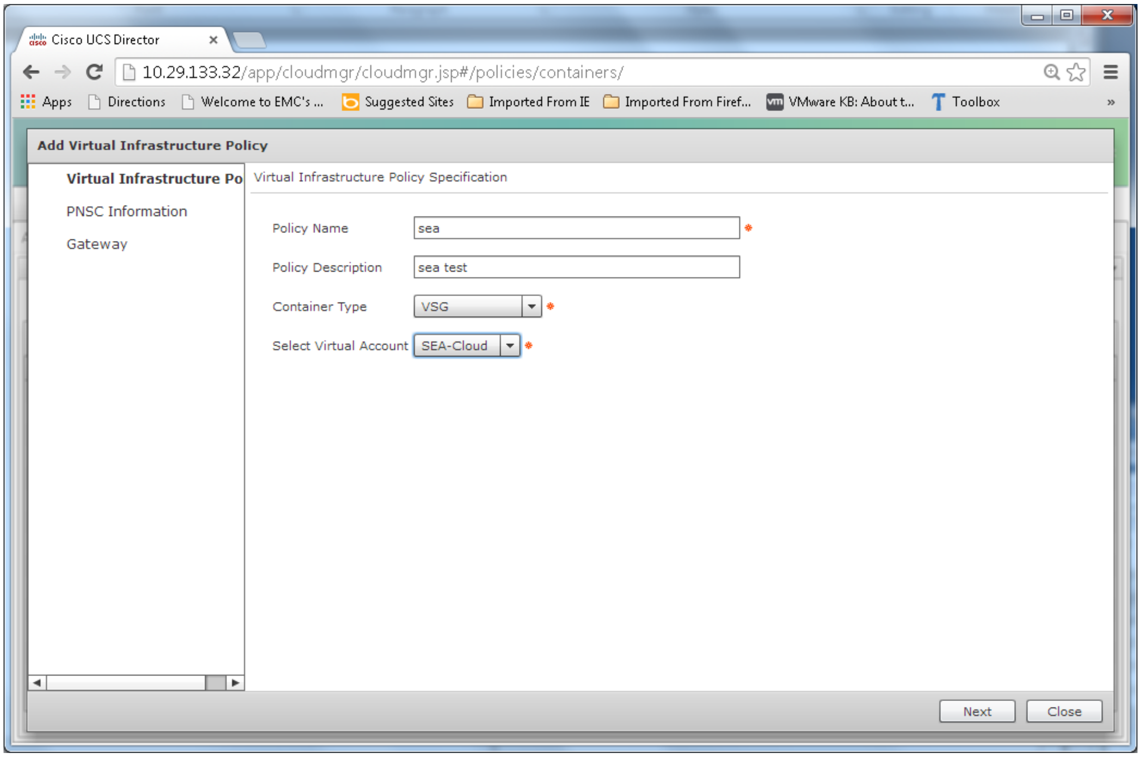

Add Virtual Infrastructure Policy

Virtual Infrastructure Policy Specification

Policy Name: sea *

Policy Description: sea test

Container Type: VSG *

Select Virtual Account: SEA-Cloud *

b. Click "Next" and select the created PNSC account and Policy and click "Next".

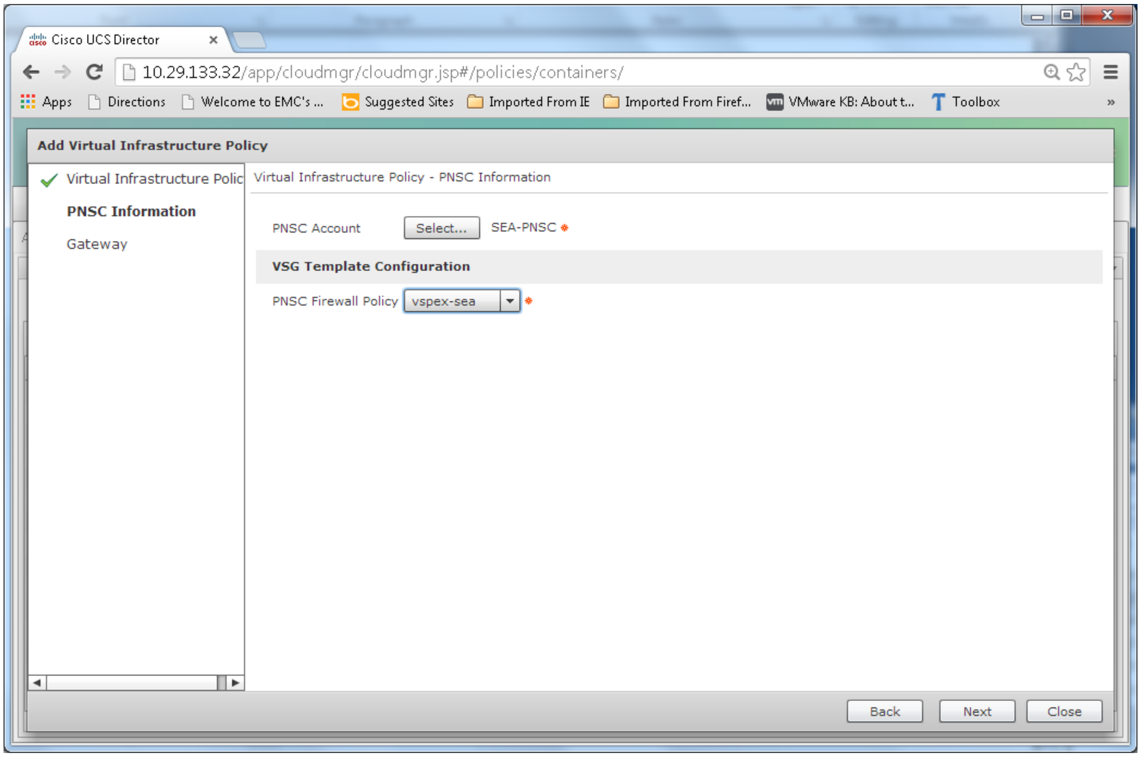

Add Virtual Infrastructure Policy

Virtual Infrastructure Policy - PNSC Information

PNSC Account: SEA-PNSC *

VSG Template Configuration

PNSC Firewall Policy: vspex-sea *

c. Click Next and populate the fields. Click Submit.

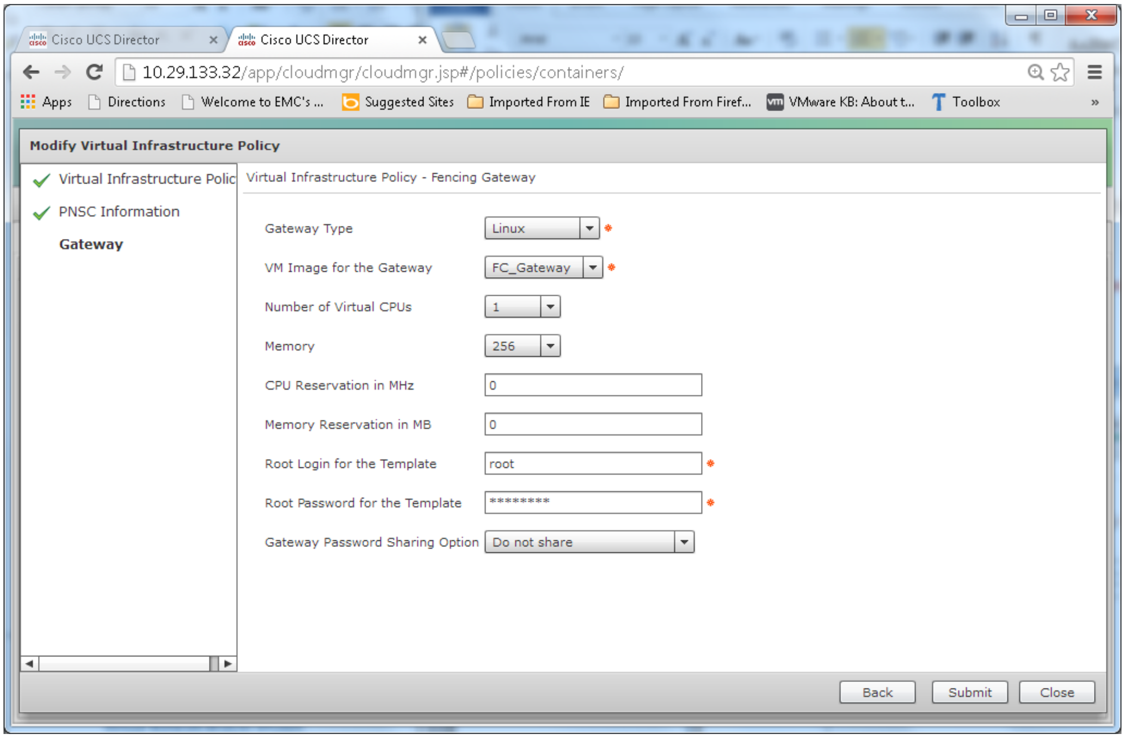

Modify Virtual Infrastructure Policy

- ✓ Virtual Infrastructure Polic
- ✓ PNSC Information
- **Gateway**

Virtual Infrastructure Policy - Fencing Gateway

Gateway Type	Linux *
VM Image for the Gateway	FC_Gateway *
Number of Virtual CPUs	1
Memory	256
CPU Reservation in MHz	0
Memory Reservation in MB	0
Root Login for the Template	root *
Root Password for the Template	******** *
Gateway Password Sharing Option	Do not share

4. Create Application Container Template:

Create the required VMWare images to serve as a templates and upload to vCenter inventory. In our case, the FC_Gateway linux image serves as a template to build the three-tiered application with VSG.

a. Select "Application container Templates" tab from Pthe olicy - "Application Container" option on the main screen.

b. Click "Add Template".

c. Provide a name for the new template, click Next and select the Virtual Infrastructure Policy recently created and click "Next".

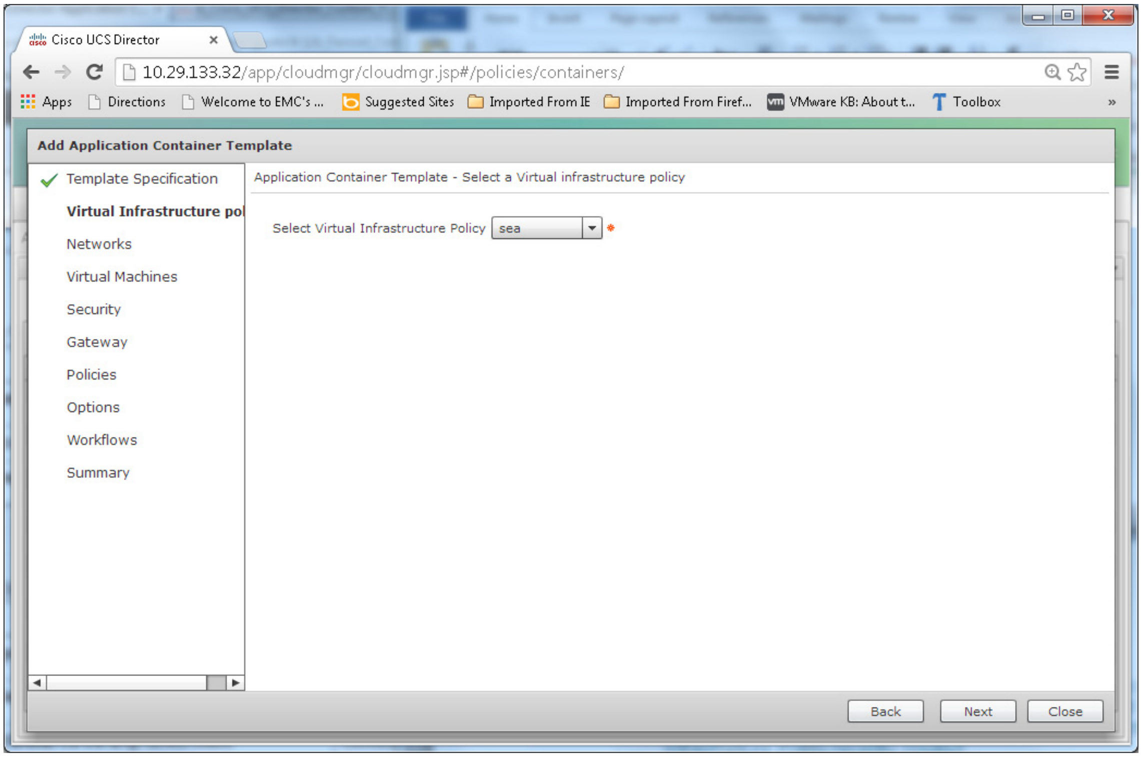

Add Application Container Template

Application Container Template - Select a Virtual infrastructure policy

Select Virtual Infrastructure Policy: sea *

d. Click "Submit" in the "Add Entry to Networks" screen.

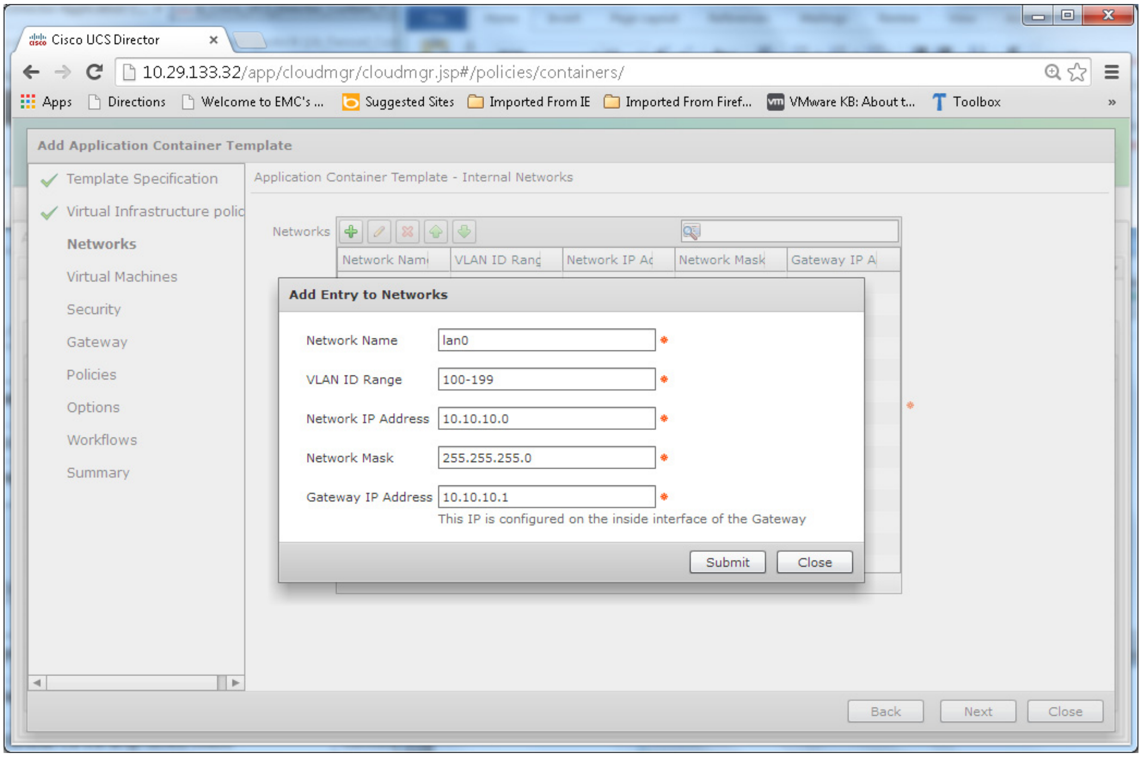

Add Application Container Template

Application Container Template - Internal Networks

Add Entry to Networks

Network Name	lan0 *
VLAN ID Range	100-199 *
Network IP Address	10.10.10.0 *
Network Mask	255.255.255.0 *
Gateway IP Address	10.10.10.1 * This IP is configured on the inside interface of the Gateway

e. Select "next" to get to the "Virtual Machines" tab where we need to add virtual machine entry and hit "submit" in the end.

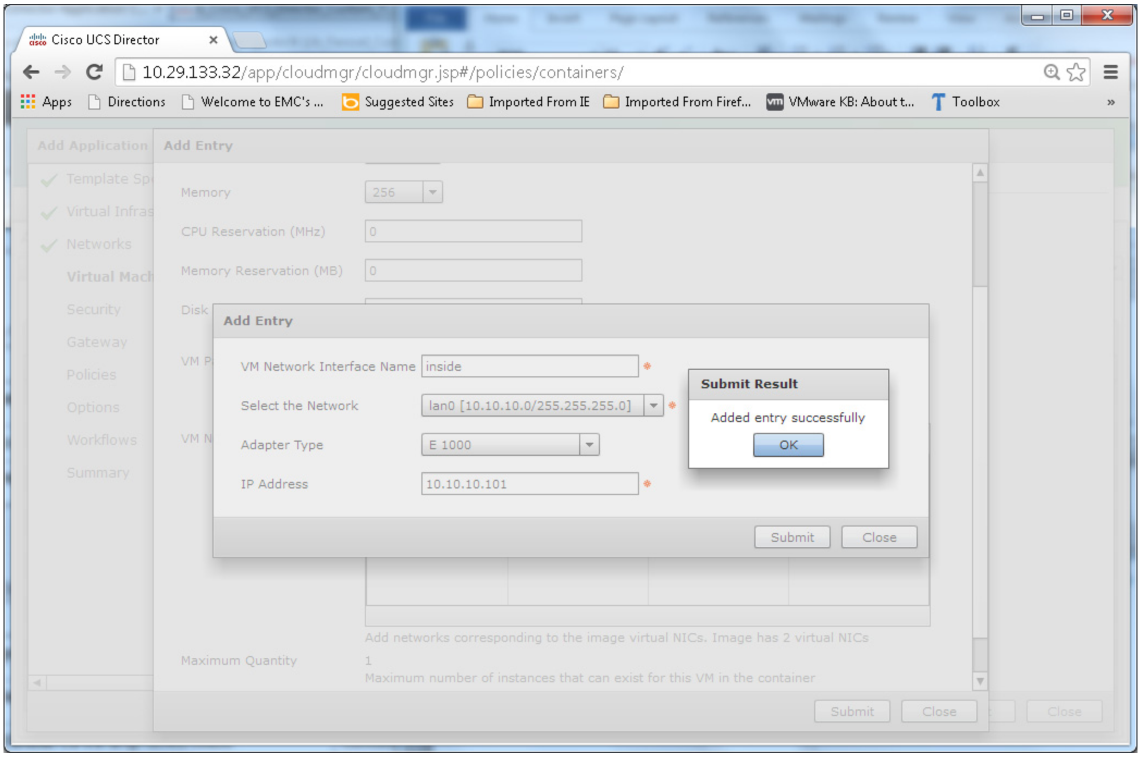

10.29.133.32/app/cloudmgr/cloudmgr.jsp#/policies/containers/

Memory: 256

CPU Reservation (MHz): 0

Memory Reservation (MB): 0

Add networks corresponding to the image virtual NICs. Image has 2 virtual NICs

Maximum Quantity: 1

Maximum number of instances that can exist for this VM in the container

Add Entry

VM Network Interface Name: inside

Select the Network: lan0 [10.10.10.0/255.255.255.0]

Adapter Type: E 1000

IP Address: 10.10.10.101

Submit Result

Added entry successfully

f. Add the required number of virtual machines. In this case, we added one each of web, app and database VMs for a three-tiered application with VSG firewall within the enclave.

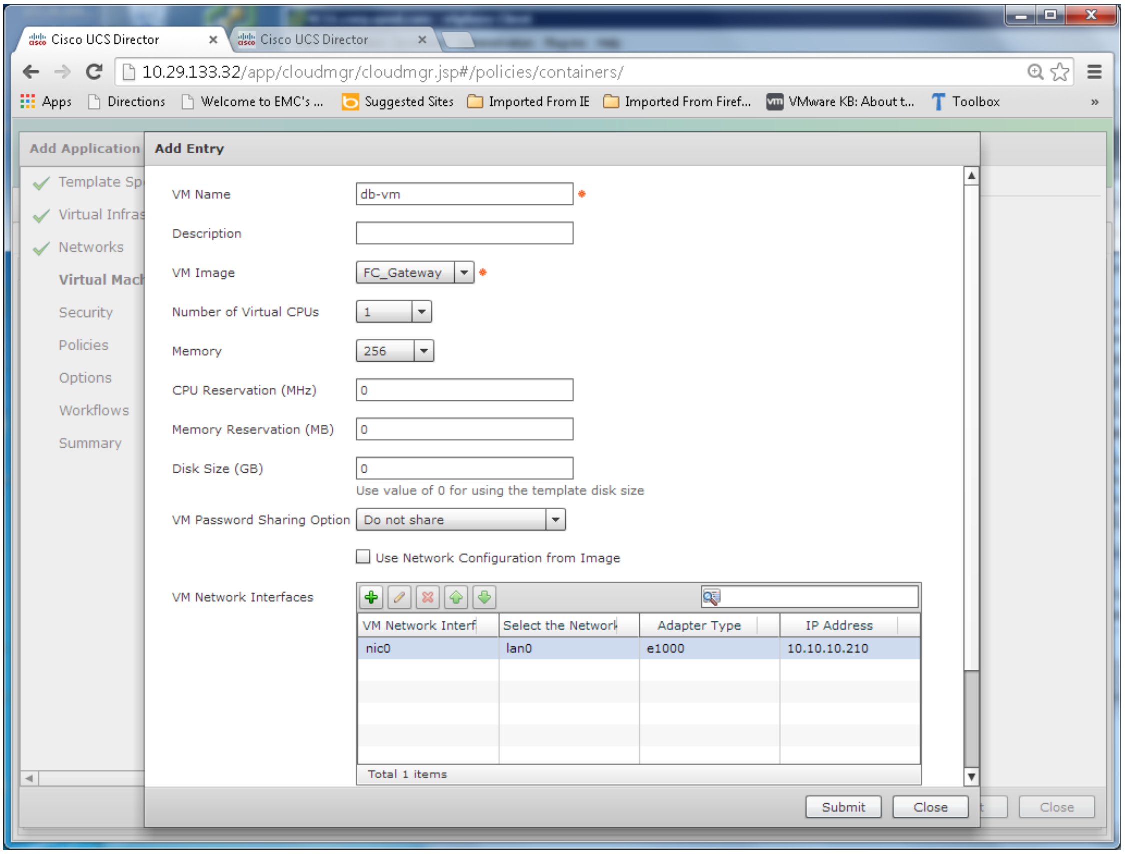

Add Entry

VM Name	db-vm *
Description	
VM Image	FC_Gateway *
Number of Virtual CPUs	1
Memory	256
CPU Reservation (MHz)	0
Memory Reservation (MB)	0
Disk Size (GB)	0 Use value of 0 for using the template disk size
VM Password Sharing Option	Do not share
	□ Use Network Configuration from Image

VM Network Interfaces

VM Network Interf	Select the Network	Adapter Type	IP Address
nic0	lan0	e1000	10.10.10.210

Total 1 items

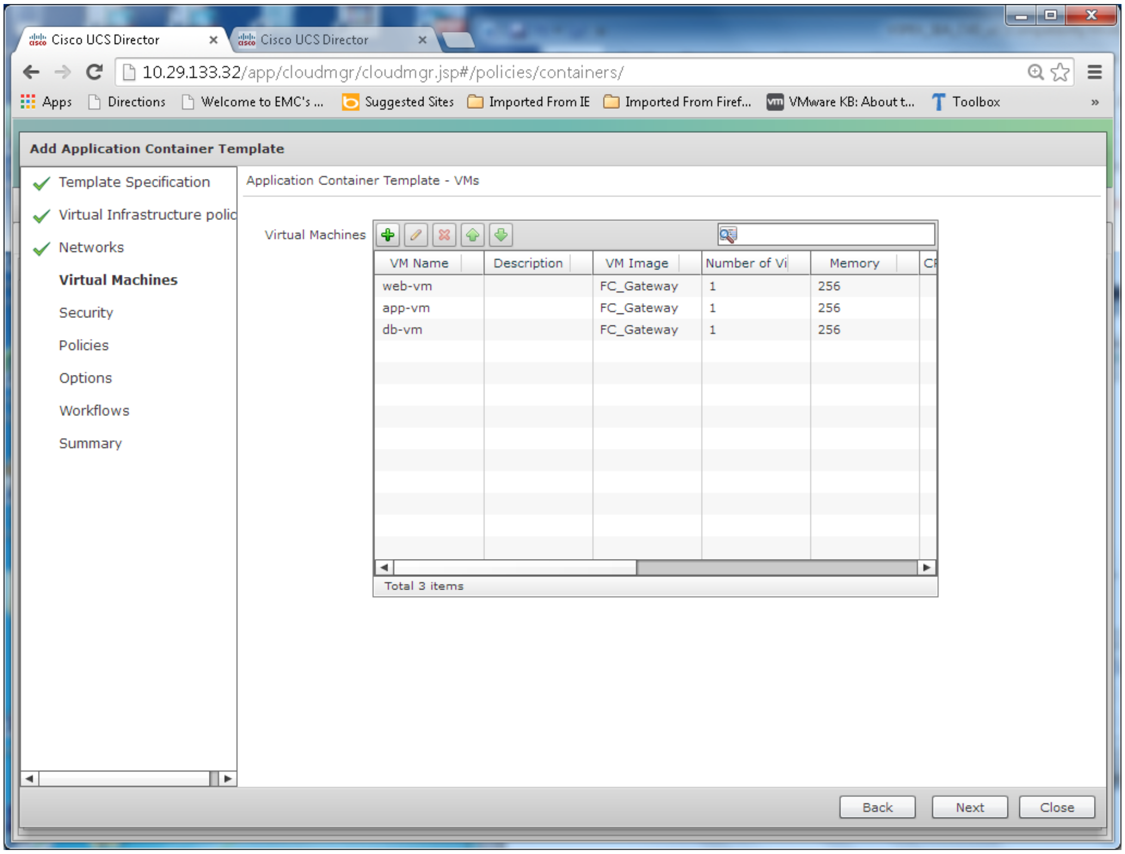

Add Application Container Template

Application Container Template - VMs

Virtual Machines

VM Name	Description	VM Image	Number of Vi	Memory	CI
web-vm		FC_Gateway	1	256	
app-vm		FC_Gateway	1	256	
db-vm		FC_Gateway	1	256	

Total 3 items

g. Click Next twice to get to the Security tab. Make the following selections:

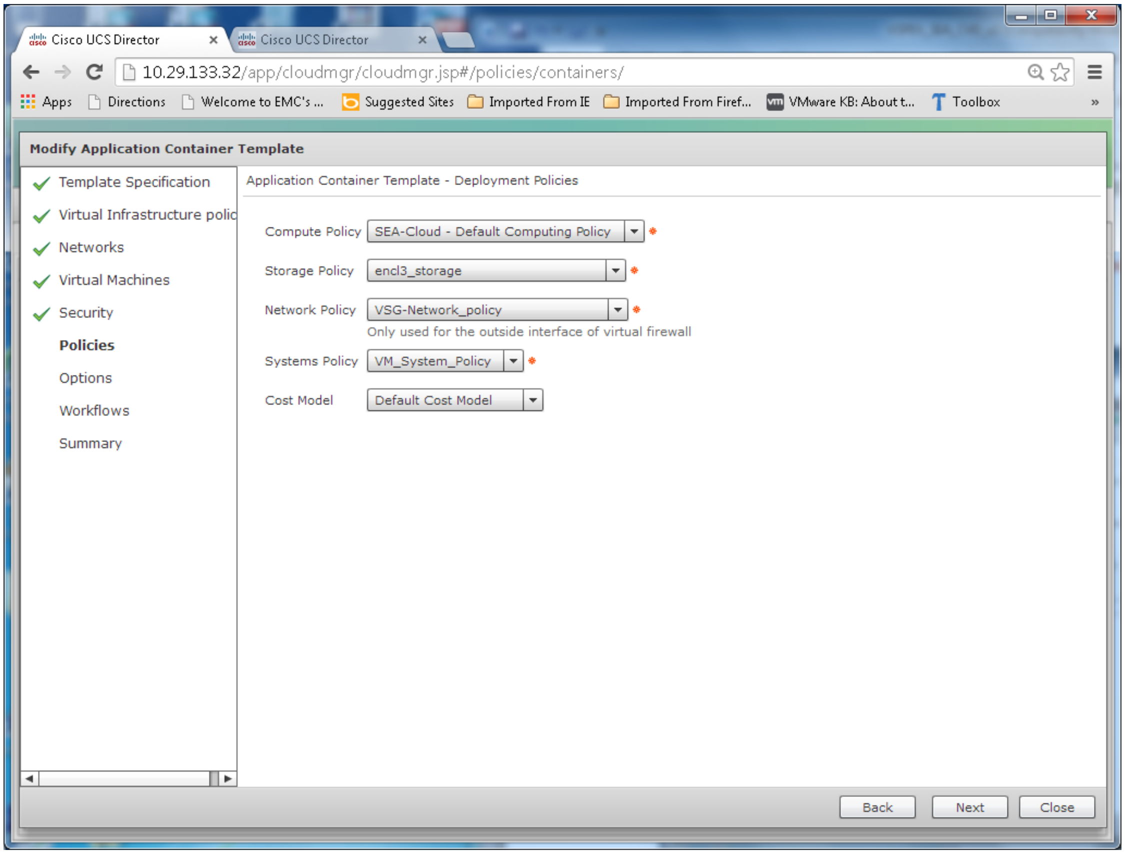

Modify Application Container Template

Application Container Template - Deployment Policies

Compute Policy: SEA-Cloud - Default Computing Policy *

Storage Policy: encl3_storage *

Network Policy: VSG-Network_policy *

Only used for the outside interface of virtual firewall

Systems Policy: VM_System_Policy *

Cost Model: Default Cost Model

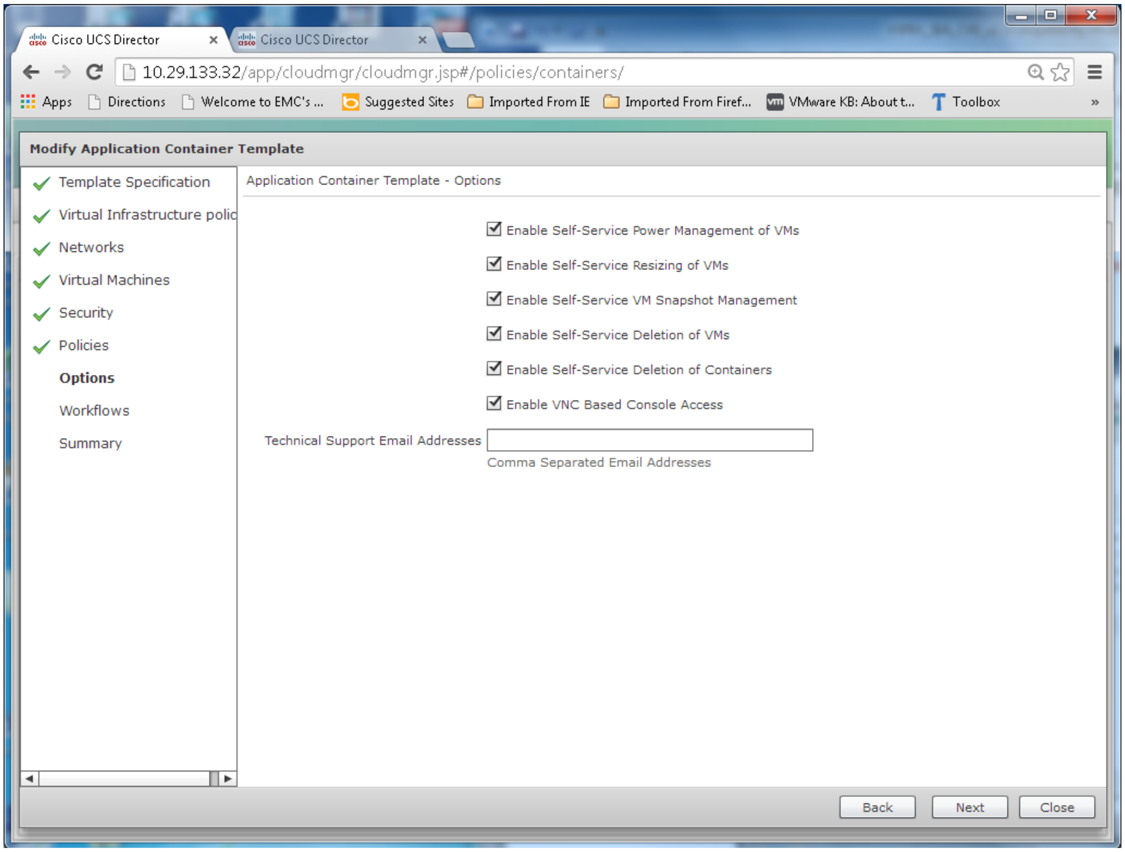

Modify Application Container Template

- ✓ Template Specification
- ✓ Virtual Infrastructure polic
- ✓ Networks
- ✓ Virtual Machines
- ✓ Security
- ✓ Policies
- **Options**
- Workflows
- Summary

Application Container Template - Options

☑ Enable Self-Service Power Management of VMs

☑ Enable Self-Service Resizing of VMs

☑ Enable Self-Service VM Snapshot Management

☑ Enable Self-Service Deletion of VMs

☑ Enable Self-Service Deletion of Containers

☑ Enable VNC Based Console Access

Technical Support Email Addresses

Comma Separated Email Addresses

h. Select the "VSG Container Setup" workflow and click Select and Next.

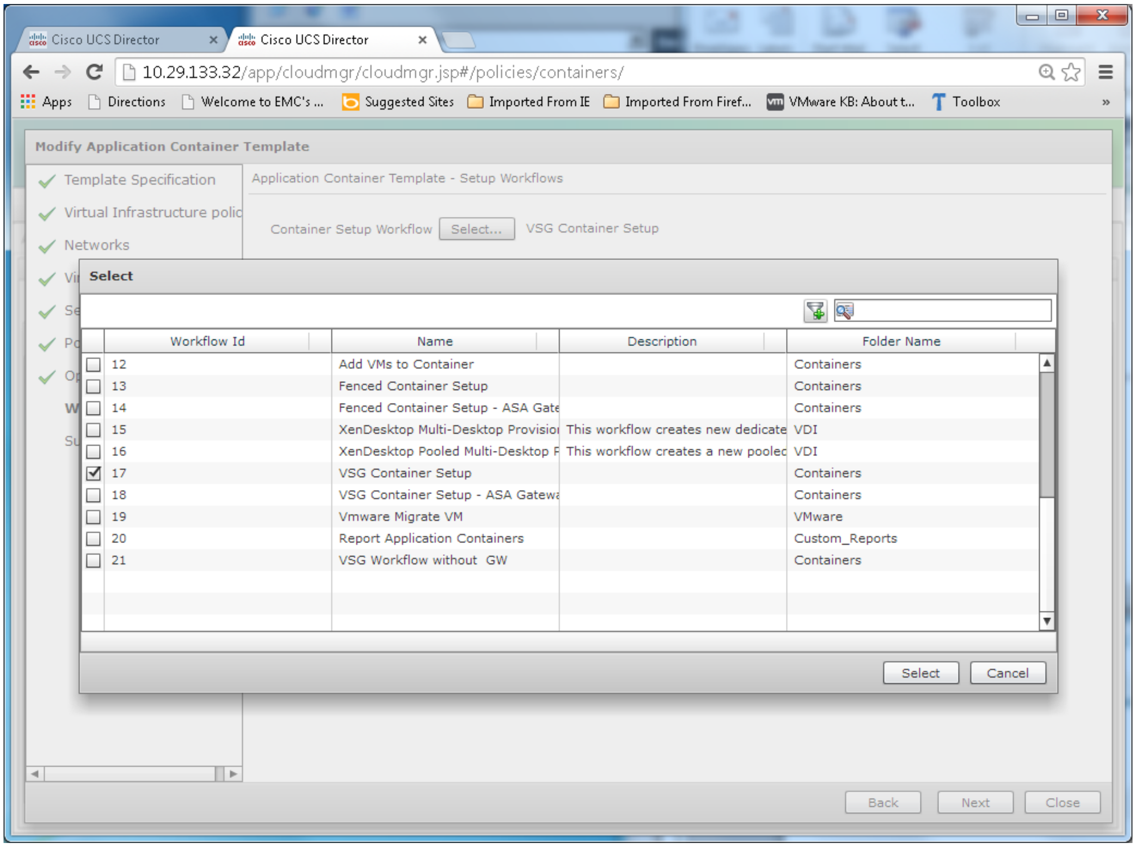

Modify Application Container Template

✓ Template Specification
✓ Virtual Infrastructure polic
✓ Networks

Application Container Template - Setup Workflows

Container Setup Workflow Select... VSG Container Setup

Select

	Workflow Id	Name	Description	Folder Name
□	12	Add VMs to Container		Containers
□	13	Fenced Container Setup		Containers
□	14	Fenced Container Setup - ASA Gate		Containers
□	15	XenDesktop Multi-Desktop Provisioı	This workflow creates new dedicate	VDI
□	16	XenDesktop Pooled Multi-Desktop F	This workflow creates a new poolec	VDI
☑	17	VSG Container Setup		Containers
□	18	VSG Container Setup - ASA Gatewa		Containers
□	19	Vmware Migrate VM		VMware
□	20	Report Application Containers		Custom_Reports
□	21	VSG Workflow without GW		Containers

i. A summary page will display the selections made for confirmation:

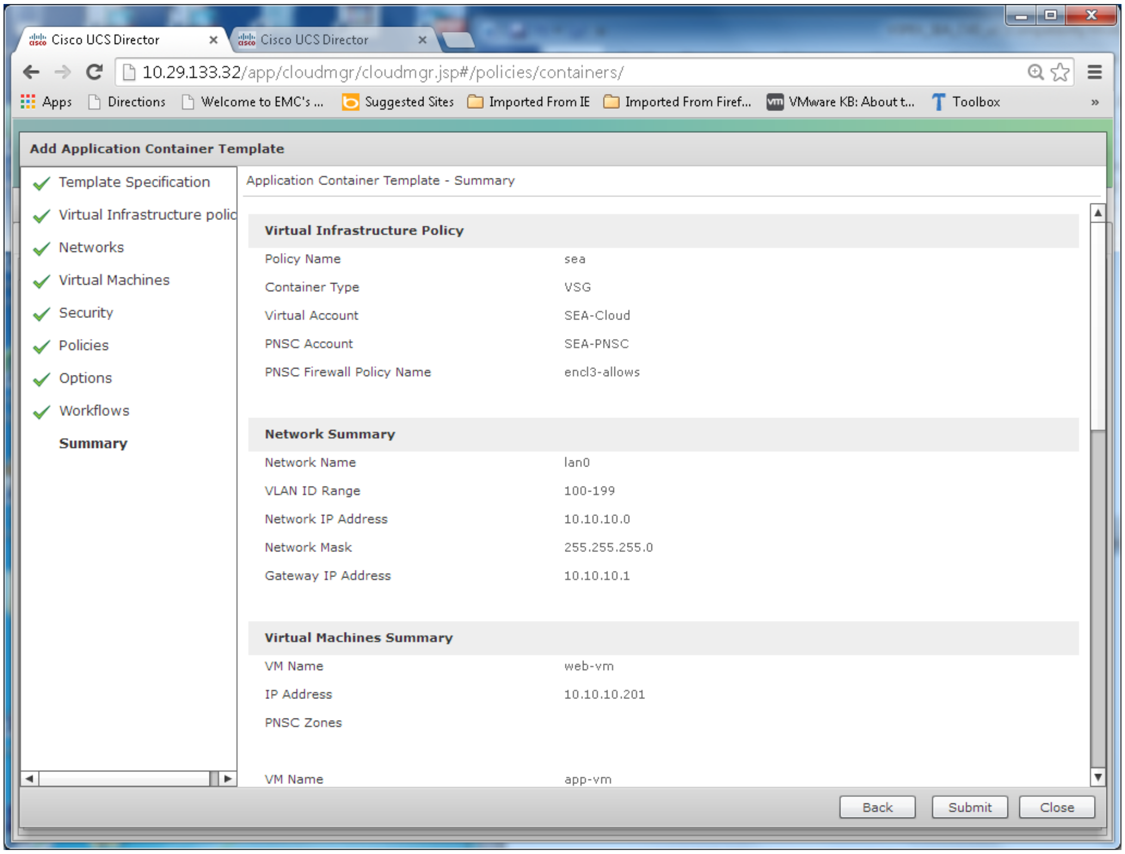

Add Application Container Template

Application Container Template - Summary

Virtual Infrastructure Policy

Policy Name	sea
Container Type	VSG
Virtual Account	SEA-Cloud
PNSC Account	SEA-PNSC
PNSC Firewall Policy Name	encl3-allows

Network Summary

Network Name	lan0
VLAN ID Range	100-199
Network IP Address	10.10.10.0
Network Mask	255.255.255.0
Gateway IP Address	10.10.10.1

Virtual Machines Summary

VM Name	web-vm
IP Address	10.10.10.201
PNSC Zones	
VM Name	app-vm

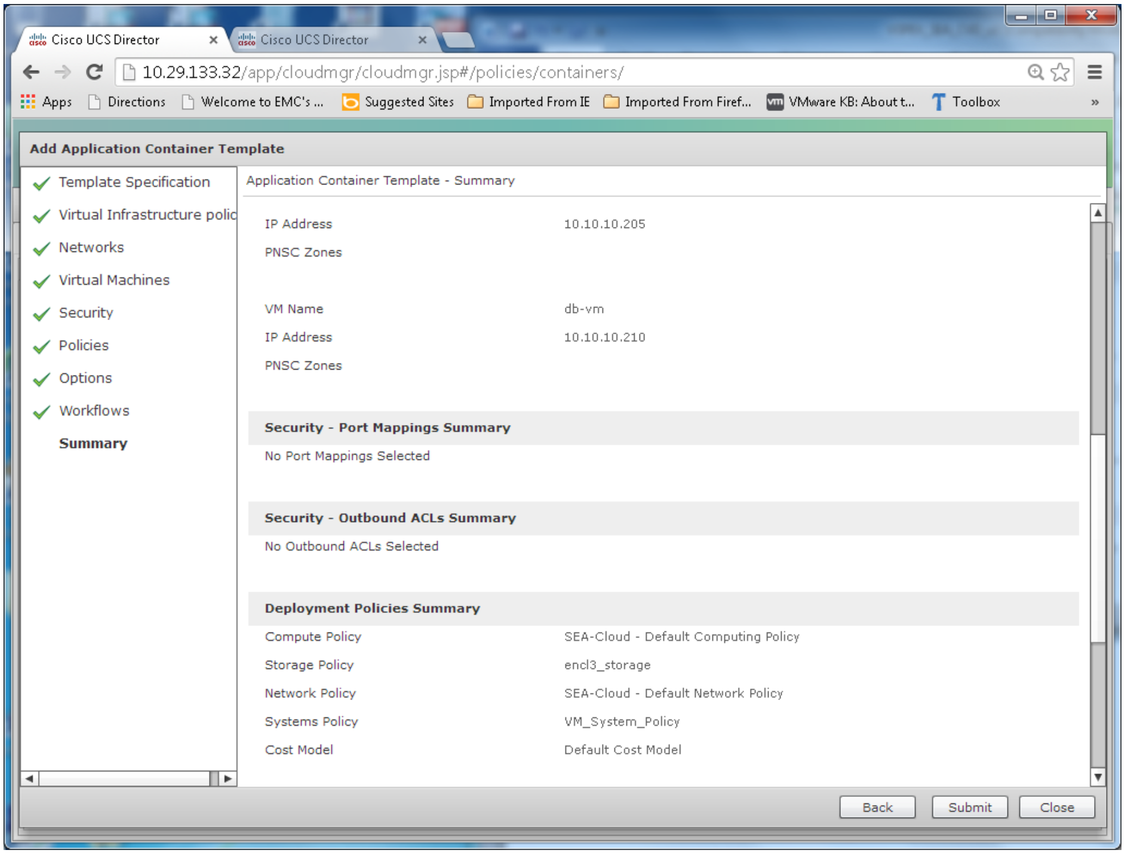

Add Application Container Template

- ✓ Template Specification
- ✓ Virtual Infrastructure polic
- ✓ Networks
- ✓ Virtual Machines
- ✓ Security
- ✓ Policies
- ✓ Options
- ✓ Workflows
- **Summary**

Application Container Template - Summary

IP Address	10.10.10.205
PNSC Zones	
VM Name	db-vm
IP Address	10.10.10.210
PNSC Zones	

Security - Port Mappings Summary

No Port Mappings Selected

Security - Outbound ACLs Summary

No Outbound ACLs Selected

Deployment Policies Summary

Compute Policy	SEA-Cloud - Default Computing Policy
Storage Policy	encl3_storage
Network Policy	SEA-Cloud - Default Network Policy
Systems Policy	VM_System_Policy
Cost Model	Default Cost Model

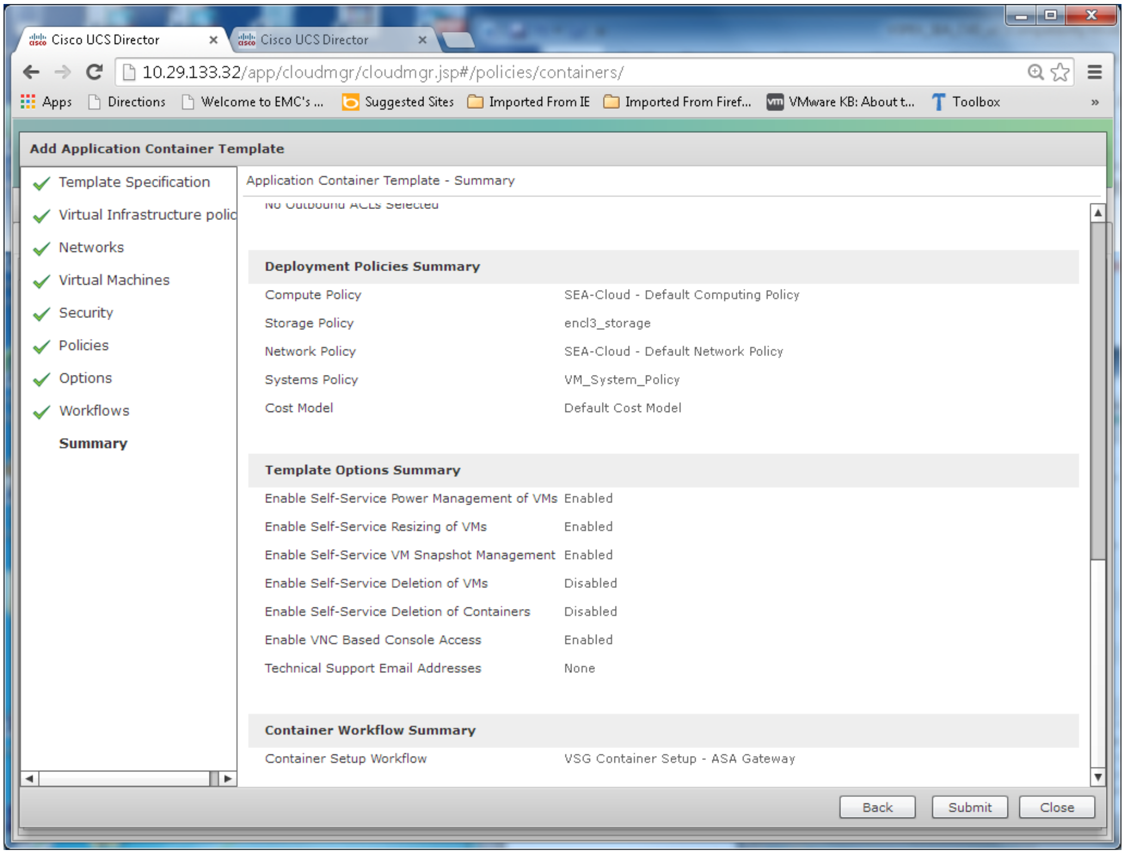

Add Application Container Template

Application Container Template - Summary

No Outbound ACLs Selected

Deployment Policies Summary	
Compute Policy	SEA-Cloud - Default Computing Policy
Storage Policy	encl3_storage
Network Policy	SEA-Cloud - Default Network Policy
Systems Policy	VM_System_Policy
Cost Model	Default Cost Model

Template Options Summary	
Enable Self-Service Power Management of VMs	Enabled
Enable Self-Service Resizing of VMs	Enabled
Enable Self-Service VM Snapshot Management	Enabled
Enable Self-Service Deletion of VMs	Disabled
Enable Self-Service Deletion of Containers	Disabled
Enable VNC Based Console Access	Enabled
Technical Support Email Addresses	None

Container Workflow Summary	
Container Setup Workflow	VSG Container Setup - ASA Gateway

The template (shown above), when deployed, creates a three-tiered application container with a VSG. Once deployed, elements within the container/enclave may be altered as required including adding or deleting VMs from the Cisco UCS Director level. This method may be used for fast instantiation of necessary enclaves with supported security features.

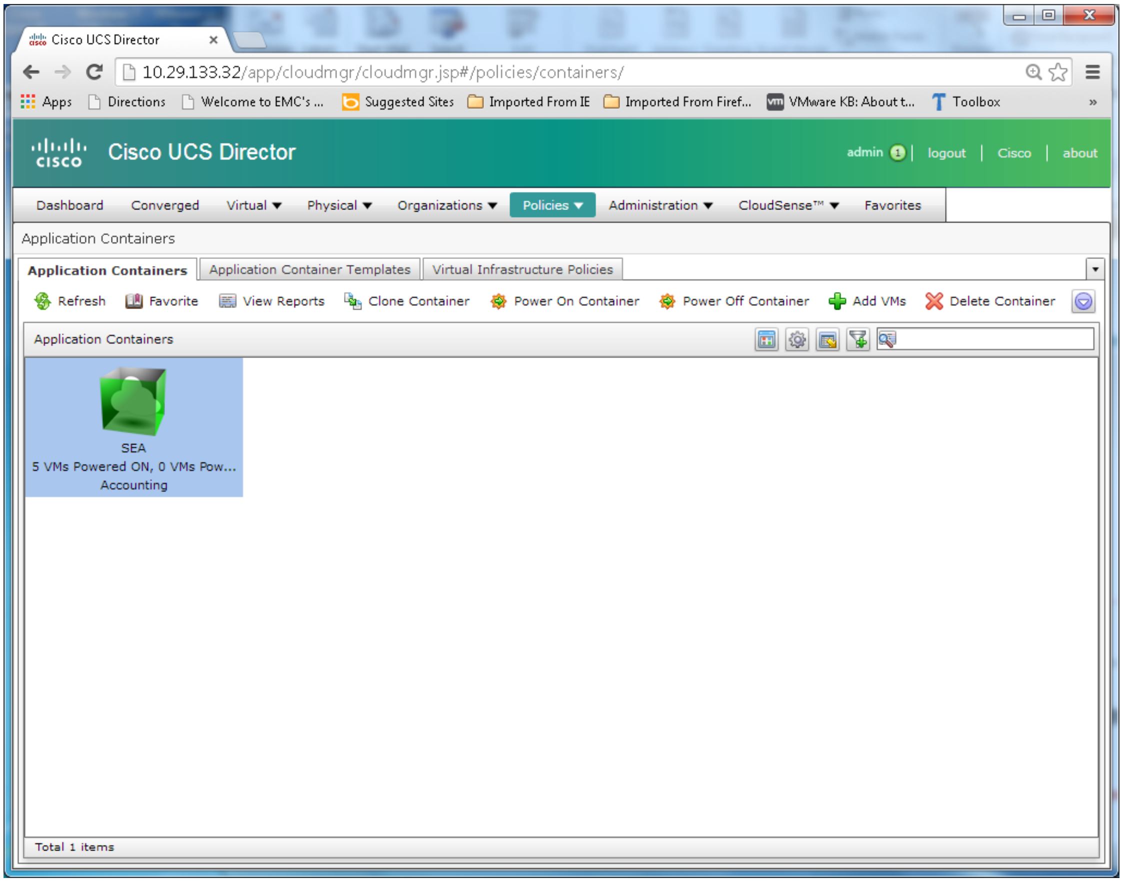

Note The vCenter view of enclave created is detailed in the following section. Please note the"Enclave_3" cluster with a three-tiered container and VSG.

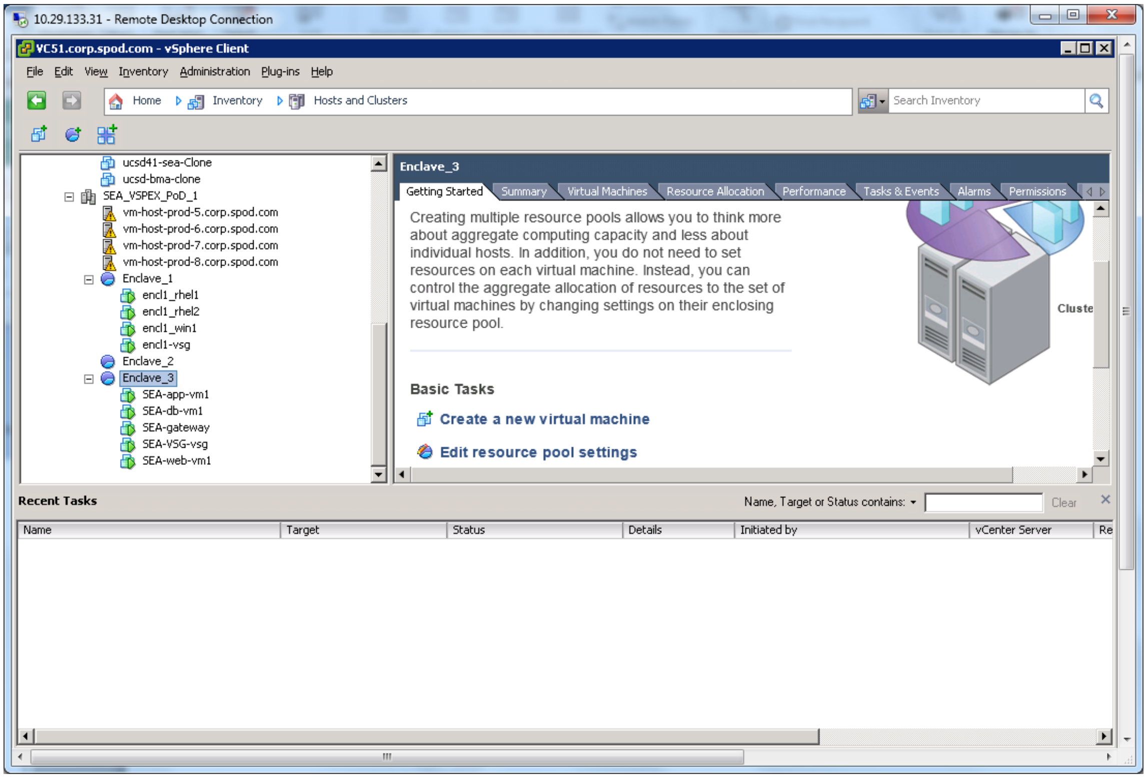
10.29.133.31 - Remote Desktop Connection
VC51.corp.spod.com - vSphere Client
File Edit View Inventory Administration Plug-ins Help
Home Inventory Hosts and Clusters
Search Inventory
Getting Started Summary Virtual Machines Resource Allocation Performance Tasks & Events Alarms Permissions
Cluste
Recent Tasks
Name, Target or Status contains:
Clear
Name Target Status Details Initiated by vCenter Server Re

- ucsd41-sea-Clone
- ucsd-bma-clone
- SEA_VSPEX_PoD_1
  - vm-host-prod-5.corp.spod.com
  - vm-host-prod-6.corp.spod.com
  - vm-host-prod-7.corp.spod.com
  - vm-host-prod-8.corp.spod.com
  - Enclave_1
    - encl1_rhel1
    - encl1_rhel2
    - encl1_win1
    - encl1-vsg
  - Enclave_2
  - Enclave_3
    - SEA-app-vm1
    - SEA-db-vm1
    - SEA-gateway
    - SEA-VSG-vsg
    - SEA-web-vm1

Enclave_3

Creating multiple resource pools allows you to think more about aggregate computing capacity and less about individual hosts. In addition, you do not need to set resources on each virtual machine. Instead, you can control the aggregate allocation of resources to the set of virtual machines by changing settings on their enclosing resource pool.

Basic Tasks

Create a new virtual machine

Edit resource pool settings

Note The following is a screenshot of the Cisco UCS Director workflow steps and the time it took to build the enclave/container. It took about ten minutes to create the container with VMs as specified. Customizations may be applied to the VMs as needed. Appendix—C provides a build report generated by this process with details on the container created.

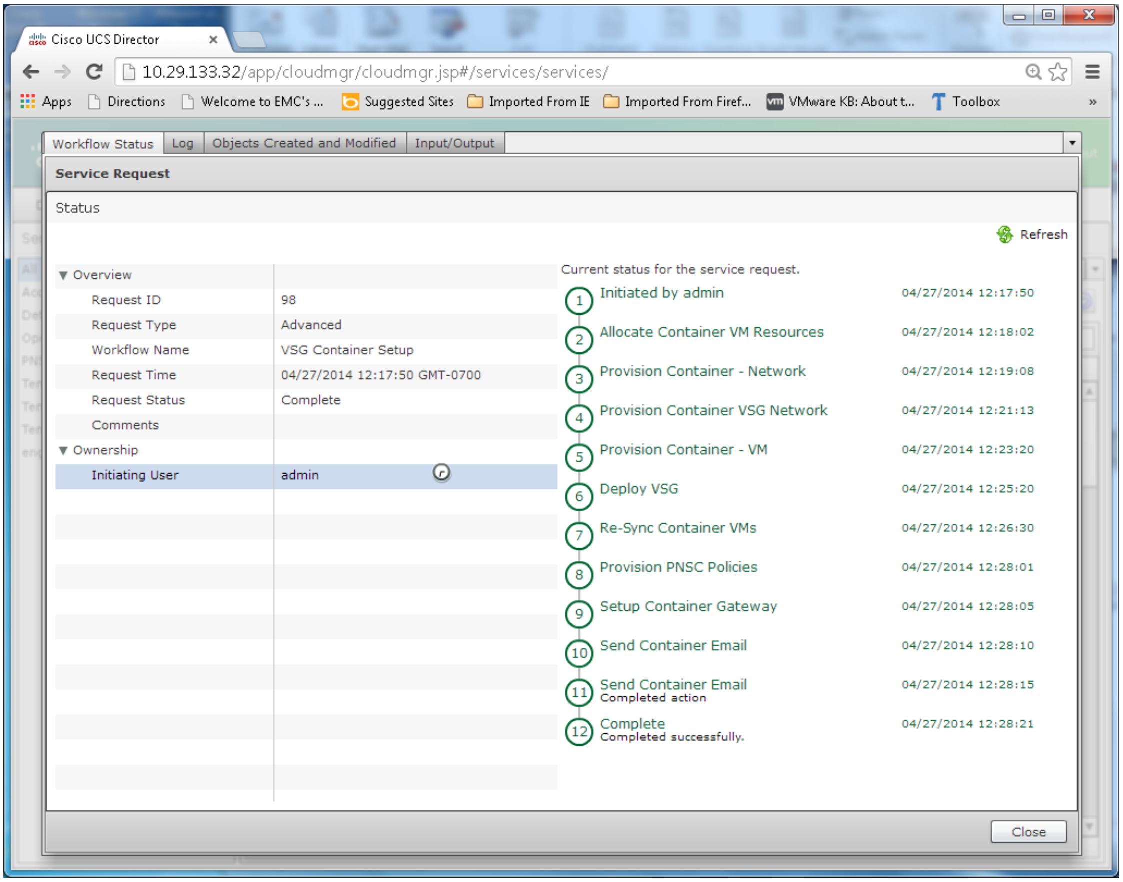

Service Request

Status

▼ Overview	
Request ID	98
Request Type	Advanced
Workflow Name	VSG Container Setup
Request Time	04/27/2014 12:17:50 GMT-0700
Request Status	Complete
Comments	
▼ Ownership	
Initiating User	admin

Current status for the service request.

1	Initiated by admin	04/27/2014 12:17:50
2	Allocate Container VM Resources	04/27/2014 12:18:02
3	Provision Container - Network	04/27/2014 12:19:08
4	Provision Container VSG Network	04/27/2014 12:21:13
5	Provision Container - VM	04/27/2014 12:23:20
6	Deploy VSG	04/27/2014 12:25:20
7	Re-Sync Container VMs	04/27/2014 12:26:30
8	Provision PNSC Policies	04/27/2014 12:28:01
9	Setup Container Gateway	04/27/2014 12:28:05
10	Send Container Email	04/27/2014 12:28:10
11	Send Container Email Completed action	04/27/2014 12:28:15
12	Complete Completed successfully.	04/27/2014 12:28:21

Bill of Materials

Software Revisions

Table 4 details the software revisions of various components used in the solution validation.

Table 4 Software Revisions

Component		Software	Count
Network	Nexus 5548UP	NX-OS - 6.0(2)N1(2a)	2
	Nexus 5000		2
	Nexus 1000v	4.2(1)SV2(2.2)	2
Compute	Cisco UCS Fabric Interconnect 6248	2.1(3a)	2
	Cisco UCS Fabric Extender - 2232	2.1(3a)	2

	Cisco UCS C220-M3	2.1(3a)	2
	Cisco UCS B200-M3	2.1(3a)	4
	VMware ESXi	5.1u1	X
	Cisco eNIC Driver	2.1.2.38	
	Cisco fNIC Driver	1.5.0.45	
	VMware vCenter	5.1u1	1
Services	Cisco Virtual Security Gateway (VSG)	4.2(1)VSG1(1)	X
	Cisco UCS Manager (UCSM)	2.1(3)	1
	Cisco Network Analysis Module (NAM) VSB	5.1(2)	1
	Cisco NetFlow Generation Appliance (NGA)	1.0(2)	2
	Cisco Identity Services Engine (ISE)	1.2	2
	Lancope StealthWatch	6.3	
	Cisco Adaptive Security Appliance (ASA) 5585	9.1(2)	2
	Lancope StealthWatch FlowCollector	6.3	
Management	Cisco UCS Director	4.1 patch 3	1
	Lancope StealthWatch Management Console	6.3	
	Cisco Security Manager (CSM)	4.4	1
	Cisco Prime Network Services Controller	3.0(2e)	1
Storage	EMC VNX 5400	05.33.000.3.700 8.1.0-34700	1

Conclusion

The Cisco Secure Enclaves architecture uses the common components of Cisco VSPEX Integrated Systems with additional services integrated to address business and application security requirements. These functional requirements promote uniqueness and innovation in the integrated computing stack, augmenting their original design with support for essential services such as security and manageability. The result is a region, or enclave, and more likely multiple enclaves, within the integrated infrastructure designed and built to appropriately address the unique workload activities and business goals of an organization. This design and the validation discussed here describes the benefits of secure enclaves in Cisco's integrated stacks.

References

- Cisco Secure Enclaves Architecture Design Guide

  http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-manager/whitepaper-c07-731204.html
- Gaining Visibility and Context Through NetFlow Security Event Logging: http://www.cisco.com/en/US/solutions/collateral/ns1015/ns1238/guide_c07-728135.pdf
- Gain Visibility into the Data Center with the Cisco NetFlow Generation Appliance: http://www.cisco.com/en/US/solutions/collateral/ns1015/ns1238/guide_c07-728136.pdf
- Lancope NetFlow Bandwidth Calculator: http://www.lancope.com/resource-center/netflow-bandwidth-calculator-stealthwatch-calculator/
- Cisco NetFlow Performance Analysis:

  http://www.cisco.com/en/US/solutions/collateral/ns341/ns524/ns562/ns583/net_implementation_white_paper0900aecd80308a66.pdf
- Cisco Cyber Threat Defense for the Data Center Solution:

  http://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-security/ctd-first-look-design-guide.pdf
- Cisco Cyber Security and Threat Defense Solution 1.1 Design and Implementation Guide: http://www.cisco.com/en/US/solutions/collateral/ns1015/ns1238/cyber_threat_defense_design_guide.pdf
- Cisco Secure Data Center for Enterprise Design Guide:

  http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/SDC/DG/SDC_DesignGuide/SDC_DG_2013-11-25_v10.html#pgfId-407304
- Cisco TrustSec Solution 2.0 Design and Implementation Guide: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/TrustSec_2.0/trustsec_2.0_dig.pdf
- Design Principles for Security, by Terry V. Benzel, Cynthia E. Irvine, Timothy E. Levin, Ganesha Bhaskara, Thuy D. Nguyen, and Paul C. Clark:

  www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA476035
- Observations on the Effects of Defense in Depth on Adversary Behavior in Cyber Warfare, by Dorene L. Kewley and John Lowry:

  http://www.bbn.com/resources/pdf/USMA_IEEE02.pdf

- Security Configuration Guide for VNX P/N 300-015-128 Rev 01 and P/N 300-013-510 Rev 03.
- EMC Unified Storage and Multi-tenancy – Technology Concepts and Business Considerations.

  http://www.emc.com/collateral/hardware/white-papers/h8094-unified-storage-multitenancy-wp.pdf
- EMC Multi-tenant File Storage Solution:

  http://www.emc.com/collateral/white-papers/h12051-wp-multi-tenant-file-storage.pdf
- Cisco Validated Designs: http://www.cisco.com/go/designzone
- PNSC Install and configuration:

  http://www.cisco.com/c/en/us/td/docs/net_mgmt/virtual_network_mgmt_center/3-0/quick-start-guide/b_30_Quick_Start_Guide.pdf
- Cisco UCS Director Application Containers Guide, Release 4.1

  http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-director/fenced-container-guide/4-1/b_Fenced_Container_Guide_4_1.pdf

Appendix—A

```
ASA Configuration

: Saved
: Written by shshastr at 07:19:25.204 UTC Sat Apr 26 2014
!
ASA Version 9.1(4) <context>
!
hostname sea-asa1
enable password H3HbV9UKYOtF4uUK encrypted
xlate per-session deny tcp any4 any4
xlate per-session deny tcp any4 any6
xlate per-session deny tcp any6 any4
xlate per-session deny tcp any6 any6
xlate per-session deny udp any4 any4 eq domain
xlate per-session deny udp any4 any6 eq domain
xlate per-session deny udp any6 any4 eq domain
xlate per-session deny udp any6 any6 eq domain
passwd 2KFQnbNIdI.2KYOU encrypted
names
!
interface Management0/0
 management-only
 nameif management
 security-level 100
 ip address 10.29.133.11 255.255.255.0 standby 10.29.133.12
!
pager lines 24
logging enable
logging asdm informational
mtu management 1500
icmp unreachable rate-limit 1 burst-size 1
icmp permit any management
asdm history enable
arp timeout 14400
```

```
route management 0.0.0.0 0.0.0.0 10.29.133.1 1
timeout xlate 3:00:00
timeout pat-xlate 0:00:30
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat
0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect
0:02:00
timeout sip-provisional-media 0:02:00 uauth 0:05:00 absolute
timeout tcp-proxy-reassembly 0:01:00
timeout floating-conn 0:00:00
aaa-server ISE_Radius_Group protocol radius
aaa-server ISE_Radius_Group (management) host 10.29.133.39
 key K1kmN0gy
 radius-common-pw K1kmN0gy
user-identity default-domain LOCAL
aaa authentication enable console ISE_Radius_Group LOCAL
aaa authentication serial console ISE_Radius_Group LOCAL
aaa authentication ssh console ISE_Radius_Group LOCAL
aaa authentication http console ISE_Radius_Group LOCAL
aaa accounting enable console ISE_Radius_Group
aaa accounting serial console ISE_Radius_Group
aaa accounting ssh console ISE_Radius_Group
aaa local authentication attempts max-fail 3
aaa authorization exec authentication-server
http server enable
http 0.0.0.0 0.0.0.0 management
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
crypto ipsec security-association pmtu-aging infinite
telnet timeout 5
ssh 0.0.0.0 0.0.0.0 management
ssh timeout 5
ssh version 2
ssh key-exchange group dh-group1-sha1
no threat-detection statistics tcp-intercept
username admin password 6.mxRGn0QfHzx5wl encrypted privilege 15
username chrobrie password 8AqnDReSq7.GIgYf encrypted privilege 15
!
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum client auto
  message-length maximum 512
policy-map global_policy
 class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect rsh
  inspect rtsp
  inspect esmtp
  inspect sqlnet
  inspect skinny
  inspect sunrpc
  inspect xdmcp
  inspect sip
  inspect netbios
  inspect tftp
  inspect ip-options
 class class-default
  user-statistics accounting
!
service-policy global_policy global
Cryptochecksum:f2ca21fb96c022461fe9968526139be6
: end
```

Appendix—B

```
EMC VNX 5400 Storage Array Setup


Storage Pools:

[nasadmin@VNX5400 setup_backend]$ nas_pool -l
id      inuse   acl     name                      storage system
51      y       0       SEA-NFS-Pool 1            FNM00130702649
52      y       0       SEA-NFS-Pool 2            FNM00130702649
55      y       0       SEA-NFS-Pool 3            FNM00130702649
56      y       0       SEA-NFS-Pool 4            FNM00130702649



Filesystems:

[nasadmin@VNX5400 setup_backend]$ nas_fs  -l
id      inuse type acl   volume    name                server
1         n    1   0     10        root_fs_1
28        y    1   0     194       Encl1_DS            1
30        y    1   0     440       Encl2_DS            1
32        y    1   0     580       Encl3_DS            1
34        y    1   0     596       Encl4_DS            1


Server Mount-Points:

[nasadmin@VNX5400 setup_backend]$ server_mountpoint server_2 -l
server_2 :
/Encl1_DS
/Encl2_DS
/Encl3_DS
/Encl4_DS


NFS Export Details:

[nasadmin@VNX5400 setup_backend]$ server_export server_2 -P nfs -list
```

```
export "/Encl4_DS" rw=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
root=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
access=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93

export "/Encl3_DS" rw=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
root=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
access=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93

export "/Encl2_DS" rw=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
root=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
access=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
export "/Encl1_DS" rw=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
root=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93
access=10.10.10.90:10.10.10.91:10.10.10.92:10.10.10.93



Details about a Filesystem:

[nasadmin@VNX5400 setup_backend]$ nas_fs -i Encl1_DS
id        = 28
name      = Encl1_DS
acl       = 0
in_use    = True
type      = uxfs
worm      = off
volume    = v194
pool      = SEA-NFS-Pool 1
member_of = root_avm_fs_group_51
rw_servers= server_2
ro_servers=
rw_vdms   =
ro_vdms   =
auto_ext  = no,thin=no
fast_clone_level = 1
deduplication   = Off
thin_storage    = True
tiering_policy  = Auto-Tier/Highest Available Tier
compressed= False
mirrored  = False
stor_devs =
FNM00130702649-0018,FNM00130702649-0017,FNM00130702649-001A,FNM00130702649-0
019,FNM00130702649-001C,FNM00130702649-001B,FNM00130702649-001E,FNM001307026
49-001D,FNM00130702649-0020,FNM00130702649-001F
disks     = d31,d32,d33,d34,d35,d36,d37,d38,d40,d39
 disk=d31   stor_dev=FNM00130702649-0018 addr=c0t1l4         server=server_2
 disk=d31   stor_dev=FNM00130702649-0018 addr=c16t1l4        server=server_2
 disk=d32   stor_dev=FNM00130702649-0017 addr=c0t1l5         server=server_2
 disk=d32   stor_dev=FNM00130702649-0017 addr=c16t1l5        server=server_2
 disk=d33   stor_dev=FNM00130702649-001A addr=c0t1l6         server=server_2
 disk=d33   stor_dev=FNM00130702649-001A addr=c16t1l6        server=server_2
 disk=d34   stor_dev=FNM00130702649-0019 addr=c0t1l7         server=server_2
 disk=d34   stor_dev=FNM00130702649-0019 addr=c16t1l7        server=server_2
 disk=d35   stor_dev=FNM00130702649-001C addr=c0t1l8         server=server_2
 disk=d35   stor_dev=FNM00130702649-001C addr=c16t1l8        server=server_2
 disk=d36   stor_dev=FNM00130702649-001B addr=c0t1l9         server=server_2
 disk=d36   stor_dev=FNM00130702649-001B addr=c16t1l9        server=server_2
 disk=d37   stor_dev=FNM00130702649-001E addr=c0t1l10        server=server_2
 disk=d37   stor_dev=FNM00130702649-001E addr=c16t1l10       server=server_2
 disk=d38   stor_dev=FNM00130702649-001D addr=c0t1l11        server=server_2
 disk=d38   stor_dev=FNM00130702649-001D addr=c16t1l11       server=server_2
 disk=d40   stor_dev=FNM00130702649-0020 addr=c0t1l12        server=server_2
 disk=d40   stor_dev=FNM00130702649-0020 addr=c16t1l12       server=server_2
 disk=d39   stor_dev=FNM00130702649-001F addr=c0t1l13        server=server_2
 disk=d39   stor_dev=FNM00130702649-001F addr=c16t1l13       server=server_2


Block Storage Pool Details:

Pool Name:  SEA-NFS-Pool 3
Pool ID:  5
Raid Type:  r_5
Percent Full Threshold:  70
Description:
Disk Type:  SAS
State:  Ready
Status:  OK(0x0)
Current Operation:  None
Current Operation State:  N/A
Current Operation Status:  N/A
Current Operation Percent Completed:  0
Raw Capacity (Blocks):  2814421510
Raw Capacity (GBs):  1342.021
User Capacity (Blocks):  2247810048
User Capacity (GBs):  1071.839
Consumed Capacity (Blocks):  121872384
Consumed Capacity (GBs):  58.113
Available Capacity (Blocks):  2125937664
Available Capacity (GBs):  1013.726
Percent Full:  5.422
Total Subscribed Capacity (Blocks):  1133097984
Total Subscribed Capacity (GBs):  540.303
Percent Subscribed:  50.409
Oversubscribed by (Blocks):  0
Oversubscribed by (GBs):  0.000
Disks:
Bus 0 Enclosure 1 Disk 10
Bus 0 Enclosure 1 Disk 12
Bus 0 Enclosure 1 Disk 14
Bus 0 Enclosure 1 Disk 11
Bus 0 Enclosure 1 Disk 13
LUNs:  44, 49, 48, 52, 43, 51, 47, 45, 46, 50
```

Appendix—C

Cisco UCS Director Application Container Build Report:

Container Name:	SEA
Container Template:	SEAs
Group:	Accounting
Created:	Sun Apr 27 19:17:49 UTC 2014
Leased Until:	
Service Request:	98 (Initiated by admin)

Virtual Machines

Cloud Name	VMID	VM Name	Status	IP Address	Provisioned Time
SEA-Cloud	81	SEA-db-vm1	ON	10.10.10.41	Sun Apr 27 19:22:51 UTC 2014
SEA-Cloud	80	SEA-app-vm1	ON	10.10.10.31	Sun Apr 27 19:22:21 UTC 2014
SEA-Cloud	78	SEA-gateway	ON	10.29.131.200	Sun Apr 27 19:21:17 UTC 2014
SEA-Cloud	82	SEA-VSG-vsg	ON		Sun Apr 27 19:23:28 UTC 2014
SEA-Cloud	79	SEA-web-vm1	ON	10.10.10.21	Sun Apr 27 19:21:49 UTC 2014

VM Name:	**SEA-db-vm1**
VM ID:	81
VM Type:	Application VM
OS:	Red Hat Enterprise Linux 4 (32-bit)
Hostname:	SEA-db-vm1
Status:	ON (poweredOn)
Disk Size:	0.91 GB Committed, 0.08 GB Uncommitted
Memory:	0.25 GB (0 Reserved)
CPU:	1 X 2.2 GHz (0 Reserved)

Network Interfaces:	**Adaptor Name** **IP Address** **MAC Address** **Network Name**
	Network adapter 1 10.10.10.41,fe80::250:56ff:feac:31a7 00:50:56:ac:31:a7 SEA-pg-lan0
PNSC Zones:	
Console Access:	vnc://vm-host-prod-5.corp.spod.com:5921 Password: ******

VM Name:	**SEA-app-vm1**
VM ID:	80
VM Type:	Application VM
OS:	Red Hat Enterprise Linux 4 (32-bit)
Hostname:	SEA-app-vm1
Status:	ON (poweredOn)
Disk Size:	0.91 GB Committed, 0.08 GB Uncommitted
Memory:	0.25 GB (0 Reserved)
CPU:	1 X 2.2 GHz (0 Reserved)
Network Interfaces:	**Adaptor Name** **IP Address** **MAC Address** **Network Name**
	Network adapter 1 10.10.10.31,fe80::250:56ff:feac:552c 00:50:56:ac:55:2c SEA-pg-lan0
PNSC Zones:	
Console Access:	vnc://vm-host-prod-5.corp.spod.com:5944 Password: ******

VM Name:	**SEA-gateway**
VM ID:	78
VM Type:	External Gateway
OS:	Red Hat Enterprise Linux 4 (32-bit)
Hostname:	SEA-gateway
Status:	ON (poweredOn)
Disk Size:	0.91 GB Committed, 0.08 GB Uncommitted
Memory:	0.25 GB (0 Reserved)
CPU:	1 X 2.2 GHz (0 Reserved)
Network Interfaces:	**Adaptor Name** **IP Address** **MAC Address** **Network Name**
	Network adapter 1 10.29.131.200,fe80::250:56ff:feac:64a7 00:50:56:ac:64:a7 N1kv_L3
	Network adapter 2 10.10.10.1,fe80::250:56ff:feac:72e0 00:50:56:ac:72:e0 SEA-pg-lan0
Console Access:	vnc://vm-host-prod-6.corp.spod.com:5924
	Password: ******

VM Name:	**SEA-VSG-vsg**
VM ID:	82
VM Type:	VSG
OS:	Other 2.6.x Linux (64-bit)
Hostname:	
Status:	ON (poweredOn)
Disk Size:	0.09 GB Committed, 2.9 GB Uncommitted
Memory:	2 GB (2 Reserved)
CPU:	1 X 2.2 GHz (1.5 Reserved)

Network Interfaces:

Adaptor Name
IP Address
MAC Address
Network Name

Network adapter 1

00:50:56:ac:5d:5b
SEA-vsgdata

Network adapter 2

00:50:56:ac:70:4d
N1kv_L3

Network adapter 3

00:50:56:ac:7b:e2
SEA-vsgha

Console Access: vnc://vm-host-prod-6.corp.spod.com:5946

Password: ******

VM Name:	**SEA-web-vm1**
VM ID:	79
VM Type:	Application VM
OS:	Red Hat Enterprise Linux 4 (32-bit)
Hostname:	SEA-web-vm1
Status:	ON (poweredOn)
Disk Size:	0.91 GB Committed, 0.08 GB Uncommitted
Memory:	0.25 GB (0 Reserved)
CPU:	1 X 2.2 GHz (0 Reserved)

Network Interfaces:

Adaptor Name	IP Address	MAC Address	Network Name
Network adapter 1	10.10.10.21,fe80::250:56ff:feac:5669	00:50:56:ac:56:69	SEA-pg-lan0

PNSC Zones:

Console Access: vnc://vm-host-prod-7.corp.spod.com:5951

Password: ******

Port Mappings

There are no port mappings specified for this container.

Event History

Event ID	Time	Type	User VM	Decription
10849	Sun Apr 27 19:17:49 UTC 2014	Container Added		Container SEA created.