Secure Data Center for Enterprise— Multi Data Center Sites Deployment of Cisco ASA Clustering with FirePOWER Services

Design and Implementation Guide—

Last Updated: May 19, 2015

About the Authors

Tom Hogue, Security Solutions Manager, Security Business Group, Cisco

Tom is the Data Center Security Solutions Manager at Cisco with over 30 years in developing integrated solutions with Cisco and previous roles in the industry. Tom led the development of the industry leading data center solutions such as the FlexPods, Vblocks, and Secure Multi-tenancy.

Tom Hogue

Bart McGlothin, Security Systems Architect, Security Business Group, Cisco

Bart is a Security Solutions Architect at Cisco with over 16 years of solutions experience. Bart leads Cisco's involvement with the National Retail Federation's Association for Retail Technology Standards Committee (ARTS) as a member of the ARTS board and Executive Committee. Prior to Cisco, Bart worked as the Network Architect at Safeway, Inc.

Bart McGlothin

Matt Kaneko, Security Systems Architect, Security Business Group, Cisco

Matt Kaneko is the solution technical lead for the Secure Data Center Solution team. In this role, Matt and his team work closely with product marketing teams of various business groups along with customer's feedback to create solution architecture. Prior to this role, Matt has worked as a Technical Marketing Manager for various Cisco Security Product lines including Cisco ASA Next Generation Firewall, Cisco Intrusion Protection System, Cisco AnyConnect, and associated Management products line.

Matt Kaneko

**Mike Storm, Sr. Technical Engineering Leader, Security Business Group, Cisco
CCIE Security #13847**

Mike leads the global security community at Cisco Systems for competitive architectures and insight. One of his primary disciplines is Security in the Data Center, developing architectures focused on tightly integrating Next-Generation Security Services with Data Center and Virtualization technologies for enterprise organizations. Storm has over 20 years in the networking and cyber security industry as an Enterprise Consultant and Technical Writer, as well as a Professional Speaker on such topics. Storm is the author of several relevant papers, including the Secure Data Center Design Field Guide and co-author of the current Secure Data Center CVD portfolio.

Mike Storm

C O N T E N T S

Introduction

Data centers are facing unprecedented change as businesses strive to accelerate their operations to capture new business opportunities. IT professionals are being asked to consolidate, virtualize, and achieve new levels of operational agility to respond to the increase in business demand for new IT services. As a result, IT organizations are finding challenges in scaling out their networks for workload mobility. They also find challenges in delivering on their operational service level agreements with legacy protocols. Even worse, they find that "bolt on" approaches to integrating services lead them to have significant data loss in their data centers. The industry has realized that existing architectures need improvements.

Primarily, the industry has an installed base of "distributed" fabrics that are primarily a three-tier architecture design. Recent evolutions in data center fabric architectures, as shown in Figure 1, are moving to *fabric-based* designs, leveraging existing switching platforms as well as platforms designed for new switching paradigms. These designs are optimized for horizontal scalability of both physical and virtual workloads. The data center will continue to evolve, as these data center fabrics become the building blocks for the "Internet of Everything" (IoE). IoE places demands based on the need for applications and services that are accessible anywhere and anytime. Every architectural approach must be able to deliver these applications and services anywhere and any time while ensuring that they are delivered in a secure and reliable manner.
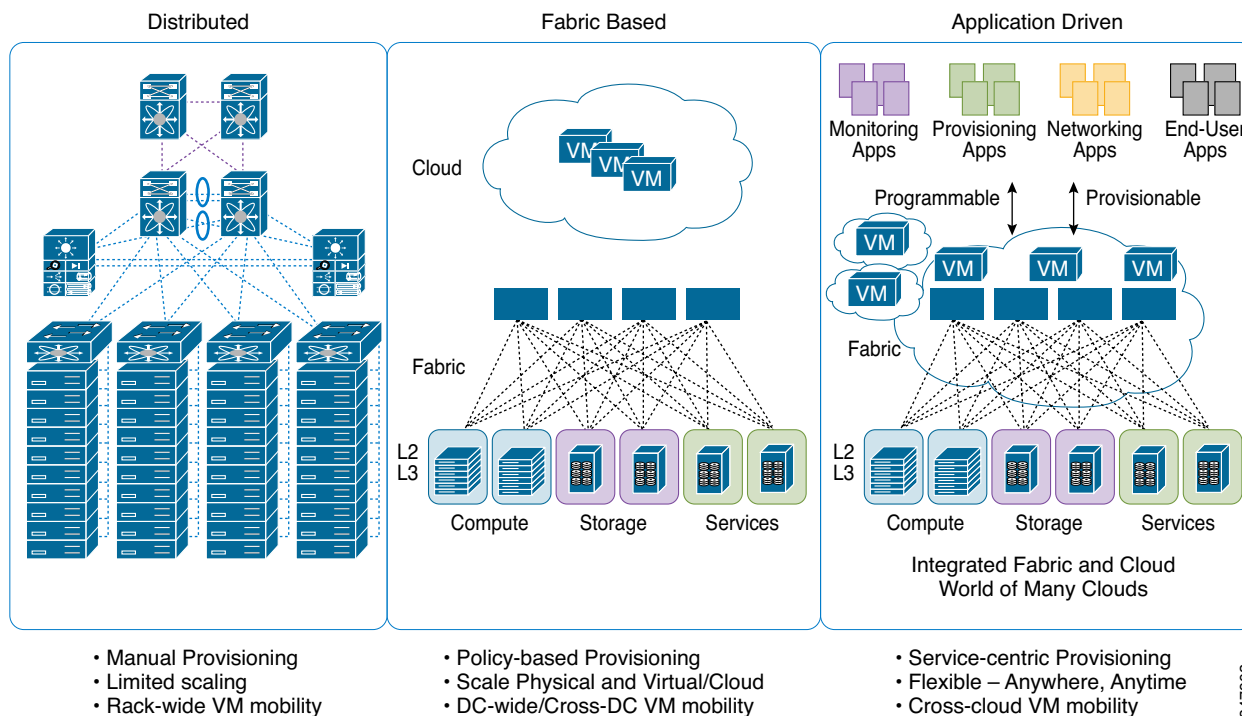
Figure 1   Evolution of Data Center Fabric Architectures

The Secure Data Center for the Enterprise solution brings together a portfolio of solutions to ensure the secure and reliable delivery of business applications. ASA Clustering with FirePOWER Services, which is part of the Secure Data Center for the Enterprise Solution portfolio, brings several key technologies, products, and associated architectures together in a design that provides application awareness to the data center fabric and network services. Following are some of the key features that the ASA Clustering with FirePOWER Services provides:

- Simplified operations
- Increased high availability
- Data loss protections
- Enterprise-wide consistent policies
- Enhanced security throughout the fabric
- Flexible scalability
- Efficient use of fabric resources
- Signature- and reputation-based protections
- Behavioral analysis for threat mitigation and remedy

While application security and delivery are key operational fundamentals in the data center, managing policies for the explosion of application workloads in the data center has created a significant operational challenge for customers. This solution provides guidance to address this challenge by leveraging a new approach to solve this problem by mapping users to data center assets in a way that provides consistency, simplification, and scalability of managing firewalls across the fabric.

Leveraging new technologies such as Cisco TrustSec, customers can now efficiently deploy proper security policies across the data center with policies based on the mapping of user roles to the various assets within the data center. In the past, customers relied on security policies to be enforced by the

data center border firewall. Now, enforcement of policies can be done at the nearest access switch or ingress port with the Secure Group Access Control List (SGACL) capability provided by TrustSec in addition to policy enforcement by the firewalls. This is a critical capability for the data center because this limits the amount of exposure and risk to the organization if the data center becomes compromised by an attack. Although TrustSec is a key enabling technology, customers can still choose to deploy the solution with or without TrustSec providing the secure separation between the various service level tiers, departments, or their preferred method of separating their workloads. The Secure Data Center for the Enterprise solution also provides flexibility and scalability with the firewall-clustering feature of the Cisco ASA 5585-X with FirePOWER services.

Secure Data Center for the Enterprise is a portfolio of solutions that includes the Cisco Secure Enclaves Architecture, Cisco Cyber Threat Defense for the Data Center, and Cisco Threat Management with NextGen IPS.

The Secure Enclaves Architecture solution provides customers with a comprehensive approach to the deployment and management of virtualized workloads being deployed on an integrated system such as a Cisco/NetApp FlexPod, a Cisco/EMC VPEX, or a Vblock from VCE.

The Cisco Cyber Threat Defense for the Data Center provides the behavioral analysis capability to the solution for a zero day mitigation response to new attacks in the data center. This solution uses the Lancope StealthWatch system that collects real-time network data with various strategically placed NetFlow collectors. The StealthWatch system performs behavioral analysis on the data to identify anomalies, which can indicate the presence of advanced persistent threats.

The Cisco Threat Management with NextGen IPS builds on top of ASA Clustering guides by showing customers how to integrate the FirePOWER NextGen IPS appliances into the architecture for higher levels of performance, and how the solution provides a comprehensive set of capabilities for a threat management system. The design guide takes a different approach by providing a perspective from a cyber attacker's point of view by looking at their *attack chain* where they develop their capabilities to execute a successful attack. Figure 2 illustrates the relationship among these solutions.

For additional content that lies outside the scope of this document, see the following URL: http://www.cisco.com/go/designzone.

Figure 2   ***Cisco Secure Data Center for the Enterprise Solution Portfolio***
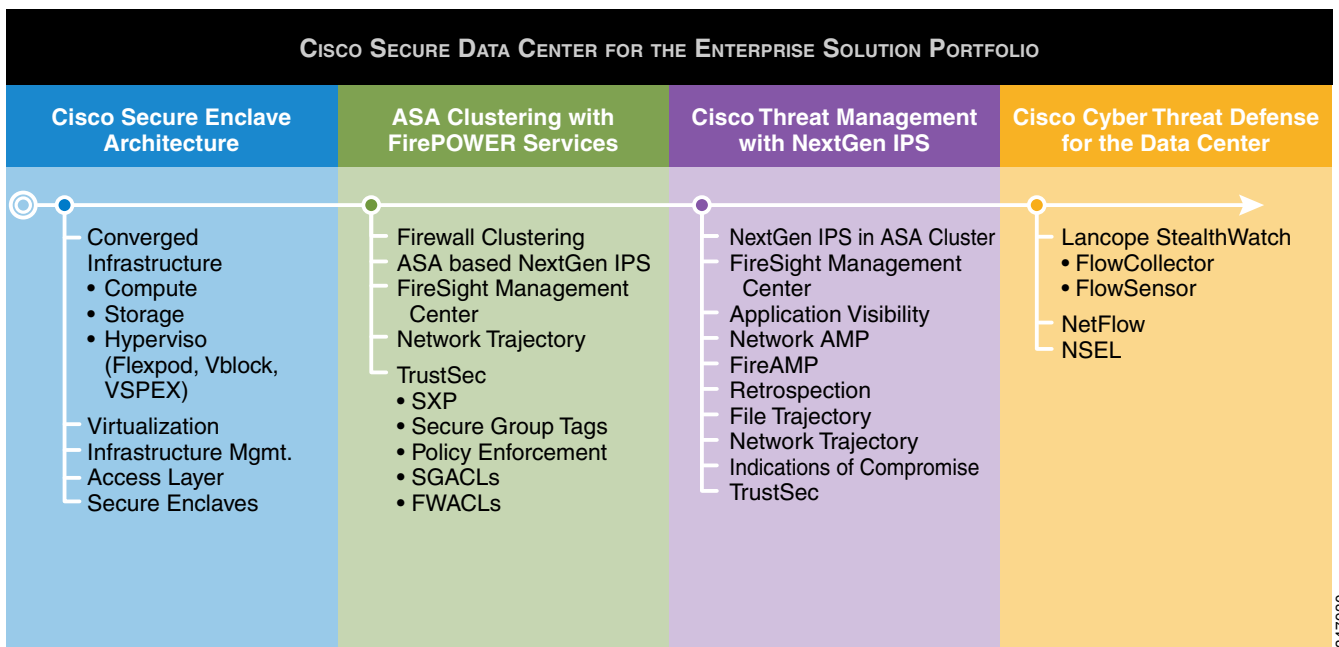
Design Overview

Goal of this Document

The Multi Sites Data Center deployment of ASA Clustering with FirePOWER Services Solution provides guidance for enterprises that are challenged with the exponential growth of data center resources and associated security policy complexity. Enterprises that want to protect against advanced data security threats can deploy a comprehensive set of security capabilities to address these needs. Figure 3 shows the architectural framework of the Secure Data Center Portfolio of products. Using Cisco's next generation firewalls, operating as a cluster with FirePOWER service modules, the goals of increased security capacity and simplicity can be jointly achieved.

Figure 3   Threat Management with FirePOWER Service Module
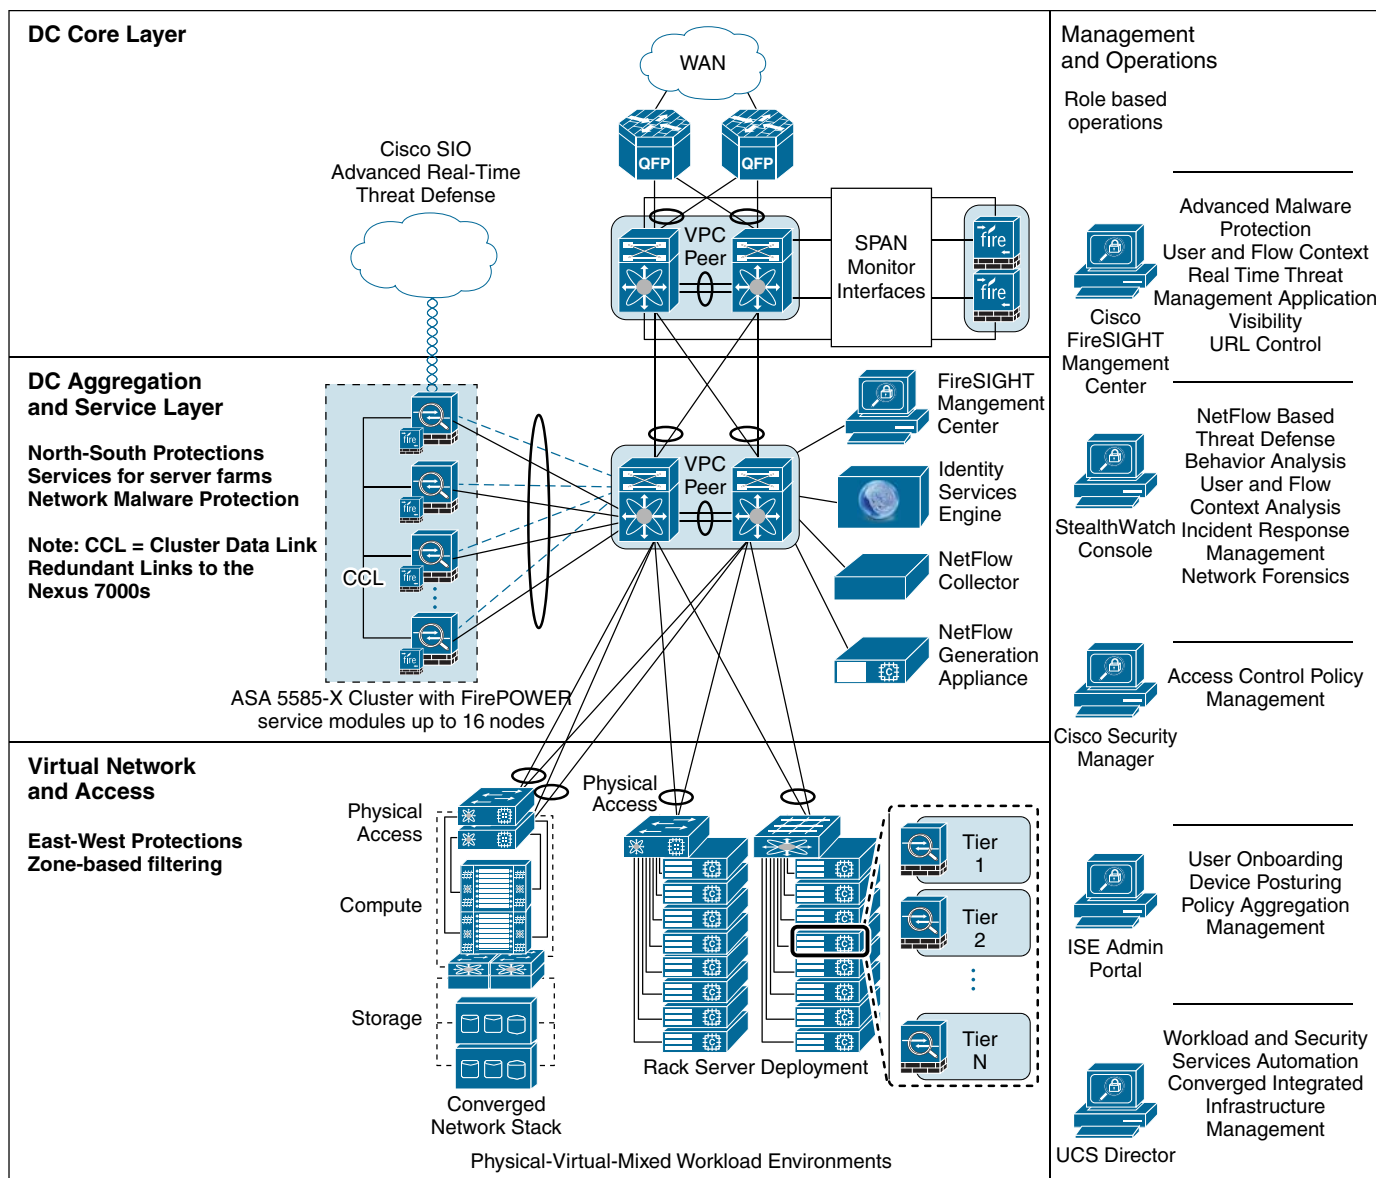
This document is specifically focused on providing implementation guidance for the ASA Clustering with FirePOWER Services solution. It is part of the Cisco Secure Data Center for the Enterprise portfolio of solutions. They provide the best protection available to address today's advanced data security threats. They contain design and implementation guidance for enterprises that want to deploy secure physical and virtualized workloads in their data centers. This solution builds on top of the Secure Data Center Single Site Clustering with TrustSec guide as a foundation, which should be treated as a pre-requisite for this implementation guide.

Intended Audience

This document is intended for, but not limited to, security architects, system architects, network design engineers, system engineers, field consultants, advanced services specialists, and customers who want to understand how to deploy a robust security architecture. This document details how specific use cases of the designs were implemented for validation. This implementation guide assumes that the reader is familiar with the basic concepts of IP protocols, quality of service (QoS), high availability (HA), and security technologies. This guide also assumes that the reader is aware of general system requirements and has knowledge of enterprise network and data center architectures.

Validated Components

Table 1 lists the validated components for the solution.

Table 1   *Validated Components*

Component	Role	Hardware	Release
Cisco Adaptive Security Appliance (ASA)	Data center firewall cluster	Cisco ASA 5585-SSP60 with additional interface card	Cisco ASA Software Release 9.3(2)
Cisco FirePOWER Service Module	Application inspection engines	Cisco FirePOWER SSP-60 Cisco FireSIGHT Management Center 3500	FirePOWER 5.4.0
Cisco Nexus 7000	Aggregation and FlexPod access switch	Cisco Nexus 7004 Cisco Nexus 7010	NX-OS version 6.2(8)

Solution Design Consideration

The Secure Data Center for the Enterprise solution is based on three key design principles:

- Provisioning
- Performance
- Protection

Provisioning

Provisioning of the Secure Data Center for the Enterprise solution, which includes automation and management, is achieved by leveraging the following four tightly integrated platforms that enable the various operational teams managing the data center to work from a common set of data, and to reduce human errors in the administration and deployment of services:

- Cisco Security Manager for managing the ASA 5585-X
- FireSIGHT Management Center for managing the NextGen IPS
- Cisco UCS Director for automating the integrated system (Vblocks, FlexPods) and third-party products
- Cisco Identity Services Engine (ISE) for policy aggregation using secure group tags (SGTs)
- Cisco Prime Network Services Controller for controlling the virtual network devices

Note The UCS Director interfaces with the Prime Network Services Controller directly, so references are not included in this document. For additional details, see the Secure Enclaves Architecture Solution Guide.

Additional management capability can be added by integrating the Data Center Network Manager platform for network management, although that product was not validated in this solution so as to limit the scope of the document.

Performance

As data centers have consolidated, more and more have updated their switching fabrics with 10G, 40G, and even some at 100G bandwidths. Firewalling capability must expand to keep from being a bottleneck in the data center. The solution achieves this with the use of the Cisco ASA 5585-X Firewall Cluster feature. By deploying the firewalls in a cluster configuration, the ASA 5585-X cluster can deliver from 10 Gbps to 256 Gbps of real-world mixed traffic throughput for ASA only filtering and 2.4 Gbps to 96 Gbps of simultaneous NGFW (AVC, and so on) and Next-Gen IPS throughput using 440-byte packet testing.

The ASA 5585-X Firewall Cluster also handles the expected asymmetric traffic flows found in a modern data center, eliminating packet loss and reducing the need for additional stateful load balancers in front of the firewalls. Fabric performance and data integrity are achieved by the use of virtual port channels (vPCs) and Link Aggregation Control Protocol (LACP).

Use case considerations for the scaling calculations of the integrated ASA FirePOWER Service Module solution should be based on the rates of the 440-byte test as the factor for sizing in the data center. While additional scale can be expected in certain situations, it is best to plan around worst-case scenario scaling. Additional scale may be achieved by adding more ASA units with FirePOWER Service Modules. Other design and deployment options could be considered for higher scaling while

requiring fewer ASA units by leveraging physical FirePOWER appliances in the context pairing solution.

For more information, see the following URL:
http://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-secure-data-center-portfolio/threat-mgmt-ips.pdf.

ASA Clustering

With Cisco ASA Software release 9.2, customers can combine up to sixteen Cisco ASA 5585-X Adaptive Security Appliance firewall modules to be joined in a single cluster for up to 256 Gbps of real-world throughput (640 Gbps max) and more than 50 million concurrent connections. Unlike competitive offerings, which experience significant declines in performance when placed into a cluster, the ASA Software clustering solution delivers a consistent scaling factor, irrespective of the number of units on the cluster. Unlike clustering on competitive platforms, which requires moderate to high changes to existing Layer 2 and Layer 3 networks, ASA Software uses the existing Cisco Virtual Switching System (VSS) and Cisco Virtual PortChannel (VPC)-based data center design and is built on standard Link Aggregation Control Protocol (LACP) and enhanced LACP, which allows up to 32 links (16 active links per switch) providing a common data plane for up to 16 devices to act as one. This enhancement, known as Spanned-EtherChannel or LACP, is supported directly in ASA software and on both Cisco Nexus and Cisco Catalyst 6000 platforms.

Note ASA EtherChannels now support up to 16 active links. With spanned EtherChannels, that functionality is extended to support up to 32 active links across the cluster when used with two switches in a vPC and when you disable dynamic port priority.

Adding a new ASA 5585-X into the cluster contributes to an increase of overall system throughput of about 70 to 80 percent of total processing capability of that unit. Throughput of an ASA 5585-X-SSP60 is 40 Gbps of optimal traffic and approximately 20 Gbps of mixed traffic. Maximum connections and connections per second have a scaling factor of 60 percent and 50 percent respectively.

Note Performance benchmarking is out of scope of the validation efforts for this design guide.

To protect high performance data centers from internal and external threats, the 16-unit cluster can be augmented by adding 16 FirePOWER services modules for up to 160 Gbps of Next-Gen IPS throughput. (See Table 2.)

Table 2   Firewall Performance Data for Cisco ASA Cluster-Real World Throughput

Platform	Single-Unit	2-Unit Cluster	4-Unit Cluster	8-Unit Cluster	16-Unit Cluster
Cisco ASA 5585X with SSP-10	2 Gbps	3.2 Gbps	6.4 Gbps	12.8 Gbps	25.6 Gbps
Cisco ASA 5585X with SSP-20	5 Gbps	8 Gbps	16 Gbps	32 Gbps	64 Gbps
Cisco ASA 5585X with SSP-40	10 Gbps	16 Gbps	32 Gbps	64 Gbps	128 Gbps
Cisco ASA 5585X with SSP-60	20 Gbps	32 Gbps	64 Gbps	128 Gbps	256 Gbps

ASA with FirePOWER Services Sizing per ASA

Table 3 compares the capabilities and capacities of the Cisco ASA with Cisco FirePOWER Security Services Processor (SSP) 10, 20, 40, and 60 hardware blades.

Table 3   Cisco ASA with FirePOWER SSP Hardware Blades

Feature	ASA 5585-X SSP-10	ASA 5585-X SSP-20	ASA 5585-X SSP-40	ASA 5585-X SSP-60
Maximum application control (AVC) or IPS throughput	4.5 Gbps	7 Gbps	10 Gbps	15 Gbps
Maximum AVC and IPS throughput	2 Gbps	3.5 Gbps	6 Gbps	10 Gbps
Maximum concurrent sessions	500,000	1,000,000	1,800,000	4,000,000
Maximum new connections per second	40,000	75,000	120,000	160,000
AVC or IPS sizing throughput [440-byte HTTP][1]	1.2 Gbps	2 Gbps	3.5 Gbps	6 Gbps
Supported applications	More than 3000			
URL categories	80+			
Number of URLs categorized	More than 280 million			
Centralized configuration, logging, monitoring, and reporting	Multi-device Cisco Security Manager and Cisco FireSIGHT Management Center			

1. 440-byte test measures greatest performance impact of multiple features.

Source:
http://www.cisco.com/c/en/us/products/security/asa-firepower-services/models-comparison.html#~ASA-5585-X

Note The ASA Context Pairing Solution using dedicated FirePOWER 8350 appliances is built to achieve up to 20 Gbps of inline AVC and NGIPS scanning per ASA in the cluster, which would be more appropriate for larger data center scaling requirements. See the following URL:
http://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-secure-data-center-portfolio/threat-mgmt-ips.pdf

Fabric Integration

To achieve the highest levels of reliability, performance, and deeper integration into the data center switching fabric, it is critical that the data center firewalls support advanced capabilities such as virtual port channel (vPC) for connecting into the fabric. vPC is a virtualization technology that presents both Cisco Nexus 7000 and 5000 Series paired devices as a unique Layer 2 logical node to access layer devices or endpoints. vPC belongs to the Multi-chassis EtherChannel [MCEC] family of technology. A vPC allows links that are physically connected to two Cisco Nexus 7000 Series devices to appear as a single port channel to a third device. vPC provides the following technical benefits:

- Eliminates Spanning Tree Protocol (STP) blocked ports
- Uses all available uplink bandwidth
- Allows dual-homed servers to operate in active-active mode
- Provides fast convergence on link or device failure

- Offers dual active/active default gateways for servers

vPC also leverages native split horizon/loop management provided by port channeling technology; a packet entering a port channel cannot immediately exit that same port channel.

By using vPC, users get the following immediate operational and architectural advantages:

- Simplified network design
- Highly resilient and robust Layer 2 network
- Enables seamless virtual machine mobility and server high-availability clusters
- Scales available Layer 2 bandwidth, increasing bisectional bandwidth
- Grows the size of the Layer 2 network

vPC leverages both hardware and software redundancy aspects as follows:

- vPC uses all port channel member links available so that if an individual link fails, the hashing algorithm redirects all flows to the remaining links.
- A vPC domain is composed of two peer devices. Each peer device processes half of the traffic coming from the access layer. If a peer device fails, the other peer device absorbs all the traffic with minimal convergence time impact.
- Each peer device in the vPC domain runs its own control plane, and both devices work independently. Any potential control plane issues stay local to the peer device and do not propagate or impact the other peer device.

From an STP standpoint, vPC eliminates STP blocked ports and uses all available uplink bandwidth. STP is used as a fail-safe mechanism and does not dictate the L2 path for vPC-attached devices.

Within a vPC domain, users can connect access devices in multiple ways:

- vPC-attached connections leveraging active/active behavior with port channel
- Active/standby connectivity using STP
- Single attachment without STP running on the access device

Components of vPC

Table 4 lists important terms you need to know to understand vPC technology. These terms are used throughout this guide and associated reference material. Figure 4 identifies these terms/components visually.

Table 4   vPC Terms

Term	Definition
vPC	The combined port channel between the vPC peers and the downstream device. A vPC is an L2 port type: switchport mode trunk or switchport mode access
vPC peer device	A vPC switch (one of a Cisco Nexus 7000 series pair)
vPC domain	Domain containing the two peer devices. Only two peer devices maximum can be part of the same vPC domain.
vPC member port	One of a set of ports (that is, port channels) that form a vPC (or port-member channel of a vPC).
vPC peer-link	Link used to synchronize the state between vPC peer devices. It must be a 10-Gigabit Ethernet link. vPC peer-link in an L2 trunk carrying vPC VLAN.

Table 4   ***vPC Terms (continued)***

vPC peer-keepalive link	The keepalive link between vPC peer devices; this link is used to monitor the liveness of the peer device.
vPC VLAN	VLAN carried over the vPC peer-link and used to communicate via vPC with a third device. As soon as a VLAN is defined on a vPC peer-link, it becomes a vPC VLAN.
non-vPC VLAN	A VLAN that is *not* part of any vPC and not present on vPC peer-link.
Orphan port	A port that belongs to a single attached device. vPC VLAN is typically used on this port
Cisco Fabric Services (CFS) protocol	"Underlying protocol running on top vPC peer-link providing reliable synchronization and consistency check mechanisms between the two peer devices."

Figure 4   ***vPC Components***

The use of vPCs for link aggregation with the ASAs can ensure that a proper data center internal zone deployment (redundant with vPC/vPC+) can be achieved by the ASA without compromising or changing the data center design, introducing new packet loss penalties or excessive risk that would otherwise not exist in the data center fabric.

To achieve this level of fabric integration, the firewall must be able to support all of the following (simultaneously):

- L2 transparent mode with VLAN bridging (re-tagging) in L2 mode

- Transparent traffic redirection through the firewall via switch trunks (with pruning)

- Dynamic LAG (LACP) to manage link consistency and prevent black hole ports. (Black hole ports are ports that are active but do not pass traffic.)

- Asymmetric flow handling on redundant links for multi-chassis link aggregation (vPC/vPC+)

- Source-Dest-IP Hash load balancing algorithm for LACP traffic distribution

Note   This is the default load-balancing algorithm on the Cisco Nexus and most all related switching products. It is not a global setting to be changed just to support a firewall, because it would affect non-security zones. To not experience traffic anomalies and/or dropped session states, this *must* match at both ends when using LACP.

Asymmetrical Data Flows vs. Symmetrical Data Flows

An end device connecting to a server in the data center is typically visualized as a single path for the IP packets to traverse across the network to get to the final destination. From the client PC or end device, the packet goes through several network devices such as switches, load balancers, firewalls, or intrusion protection devices before reaching the final destination. Packets that are part of a single communication stream and are more than a single packet are considered to be a connection stream or data flow. A connection stream that makes the entire journey through a single path is considered a *symmetrical* data flow. However, a proper large-scale data center uses load balancing in switches and appliances across the data center for various reasons such as simplified management, reduction of Spanning Tree, resource utilization, and others. The connection stream is likely to be broken up, and the packets traverse across different paths across the data center. When this occurs, it is considered to be an *asymmetrical* data flow. Legacy data center designs have required that firewall clusters maintain symmetrical data flows before entering the firewall cluster. (See Figure 5.)

Figure 5   *Legacy Load Balancing "Sandwich" Configurations Need Stateful Load Balancers*
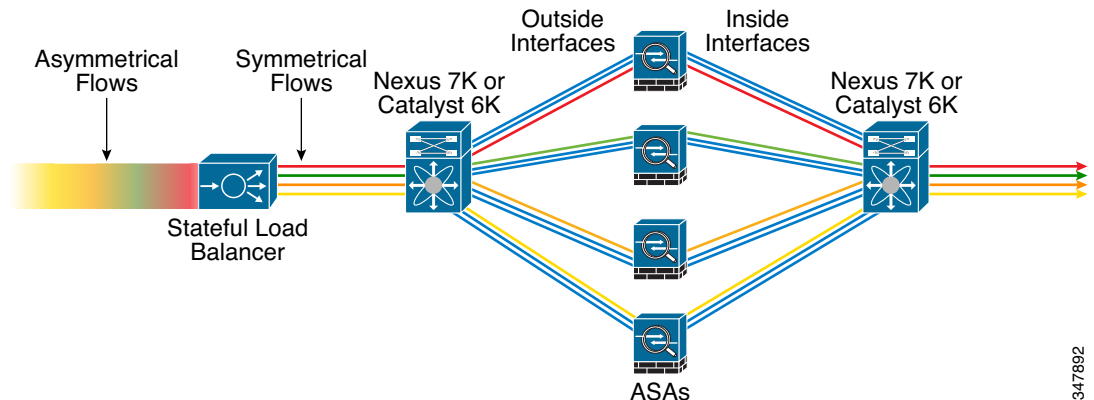
The asymmetrical and symmetrical data flows in the data center become extremely relevant in regard to firewall clustering with stateful firewalls. Cisco Nexus switches use innovative technology such as vPC and FabricPath in conjunction with LACP hashing algorithms (src-dst-ip is the default) to make asymmetric flows deterministic and highly reliable. When architected correctly, this eliminates the

need for protocols such as STP, as well as the latency, calculations, and convergence/re-convergence challenges that go along with it. The limiting factor to using many of these innovations has been the introduction of stateful devices, which by nature need to see every packet in a particular flow, in both directions, to ensure the secure nature/policy enforcement of that flow. To maintain the "stateful" deep packet inspection process working, you need to ensure that the individual connection flows are inspected by the same firewall connection process. If a firewall cluster is being fronted by a stateless load balancer, the connection flow may not be distributed to different firewalls. Previous implementations of firewall clustering were achieved by using external stateful load balancing and placing the firewalls in a "sandwich" configuration.

The load balancers would have to maintain the connection state so that packets belonging to a single connection are forwarded to the same ASA. There are several issues with this design:

- The total throughput is limited by the load balancers, which may not have sufficient bandwidth.

- The units in the cluster operate as individual devices and as such, they are managed separately, which increases complexity and operational overhead.

- Each of the firewalls in the load balancing group operates as a standalone unit, so there is no communication mechanism between them. Implementation of typical features such as NAT, AAA, and others are more difficult as a result. Each device manages its own NAT pool, once again creating a management challenge.

- To achieve a proper high availability (HA) capability, each firewall would require a stand-by unit attached to it so that in the event of a major failure, the traffic switches over to the standby unit.

- The traffic would not load balance between the two units, and packets would be lost if a switchover occurs.

- It does not have built-in mechanisms to support HA (also known as failover). Units do not monitor each, and do not share the connection state.

Some stateless load balancing mechanisms, such as policy-based routing (PBR), are readily available on high-end switches and routers. The load balancing decision is made per packet based on header information, so the router/switch does not maintain connection state across multiple packets. It is often implemented in hardware, so load balancing is done at line speed. The challenge is that each ASA is still managed individually and operates independently. As you can see, although the solution was the best available option for clustering firewalls, the solution was sub-optimal at best.

Protection

North-South Protection

The Secure Data Center for the Enterprise solution provides customers with two primary approaches to achieving north-south protection to the data center traffic flows. The more traditional approach uses VLANs and Layer 3 route points, while the other uses SGTs with Layer 3 route points. Customers can choose whether to use ASA 5585-X Layer 3 routed mode with multi-contexts as their route point; or they can operate the ASA 5585-X in Layer 2 transparent mode with multi-contexts, and choose to use Cisco Nexus 7004 virtual route forwarding (VRF) as their route point. 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. The IPS model provides an application-awareness capability that is more relevant for the typical traffic seen in the data center.

East-West Protection

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. At the time of this document, manually assigned SGTs on the Nexus 1000V port profiles is the method used in validation.

More information can be found in the Secure Enclaves Cisco Validated Design Guide at the following URL: http://www.cisco.com/go/designzone.

Figure 6 shows both north-south and east-west protection.

*Figure 6*   ***North-South and East-West Protection***
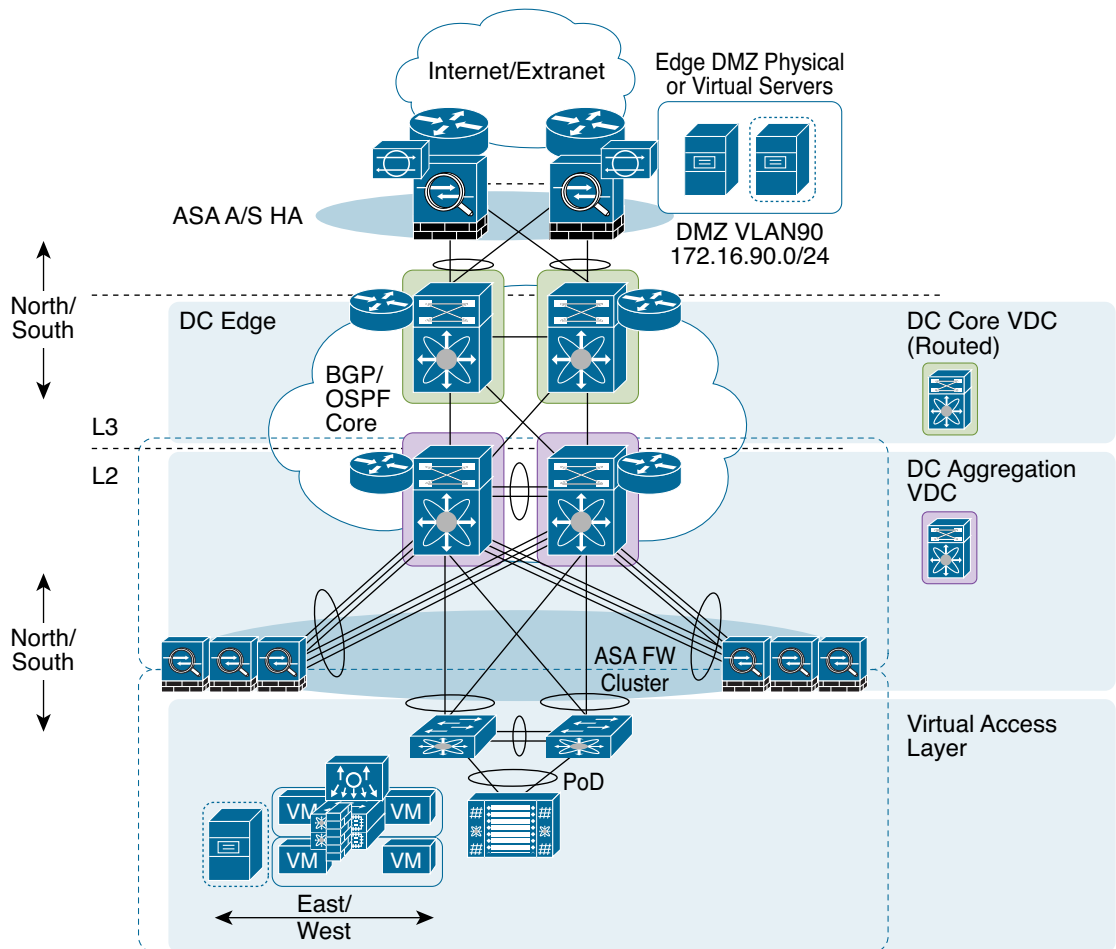
Cisco ASA 5585-X Next Generation Firewall Cluster

Figure 7 shows the Cisco ASA 5585-X cluster without external load balancers.

Figure 7   ***ASA 5585-X Cluster Without External Load Balancers***

The Cisco ASA 5585-X with clustering features addresses the challenges addressed above and provides a new way to deploy a scalable cluster that meets the business demand:

- A minimum of two Cisco ASA 5585-X firewalls can be configured up to a maximum of 16 Cisco ASA 5585-X firewalls.
- The cluster is managed as a single device for both configuration and monitoring.
- Asymmetrical connection flows are now supported as the firewall cluster redirects the connection flows to a proper inspection manager process.
- The cluster provides hardware backup to each of the other firewalls in the cluster, and each flow is backed up with a backup inspection process.
- The ASA 5585-X Cluster solution opens the opportunity to use readily available load balancing technologies such as PBR for L3 deployments and equal cost load balancing (ECLB) for L2 deployments in the switching fabric infrastructure.
- All ASA units share a system interface and appear as a single gateway in the network.

Consistent Configuration

When deploying the ASA Cluster, all of the ASAs must have the exact same configurations for the ASA system to work properly. In addition, they should be deployed in a consistent manner. This applies to using the same ports on each unit to connect to the fabric. Use the same ports for the Cluster Control Link to the switching fabric and the same with the Data links. When the ASA Cluster is deployed properly, the master unit of the cluster replicates its configuration to the other units in the cluster, and so the cluster must have a consistent deployment across all the units.

Cluster Control Link

Cluster Control Link (CCL) is a *backplane* network that is used by ASAs to exchange clustering protocol traffic (see Figure 8). Control plane traffic includes master election, configuration replication, statistics gathering, and so on. Data plane traffic includes state replication, ownership query, data packet forwarding, and so on. Deployment of the ASA Cluster requires that at least one hardware interface from each unit is dedicated to connect to the Nexus 7000 to form the CCL. Each ASA must be assigned a unique IP address on the CCL, and all IP addresses must be within the same subnet. This subnet is to be isolated from the rest of the network, which means that this subnet should contain no hosts other than the ASAs that are members of this cluster. Traffic within the subnet should not be routed to any other subnet, or vice versa. The CCL should be configured as a Layer 2 EtherChannel with LACP enabled, and connected to the Nexus 7000 with a port channel configured as a vPC.

It is important to point out that the clustered ASAs have the same port channel configuration because of the sync from the cluster, but the Nexus 7000s have different port channels configured because these are local and not spanned across the cluster.

Figure 8   ***Cluster Control Link***

ASA Port Channels for the CCL will be the same due to configuration replication by the Master Unit vPC Port Channels on the Nexus 7000 are not spanned

Cluster Control Link Sizing

Cisco recommends that the bandwidth of the CCL match at least the highest available bandwidth on the data interfaces. For example, if 2X 10GE ports are used as a data interface, the CCL also needs to support 20GB bandwidth. The reason is that the load balancing performed by the switches connecting to the cluster can be asymmetric and as such, it is possible that all traffic hits just one unit in the cluster, resulting in increased traffic.

Layer 2 Deployment Mode

In transparent mode, the ASA interfaces are grouped together in an EtherChannel to form a logical Ethernet link using Link Aggregation Control Protocol (LACP) (see Figure 9). Enabling LACP on the EtherChannel prevents data black holes caused by physical links losing the ability to pass traffic while still maintaining an active physical link status. LACP aggregates links that have the same capability

and that terminate in the same two end systems. The ASA and the switch should be configured to use ECLB to load balance traffic between the ASAs, and to use the same port channel load-balance hashing algorithm on the ASA and Nexus 7000. The ASA connects to the Nexus 7000, where the port channel connection should be mapped onto a member vPC link aggregation for resiliency and faster recovery times.

Figure 9   Layer 2 Deployment Data Plane
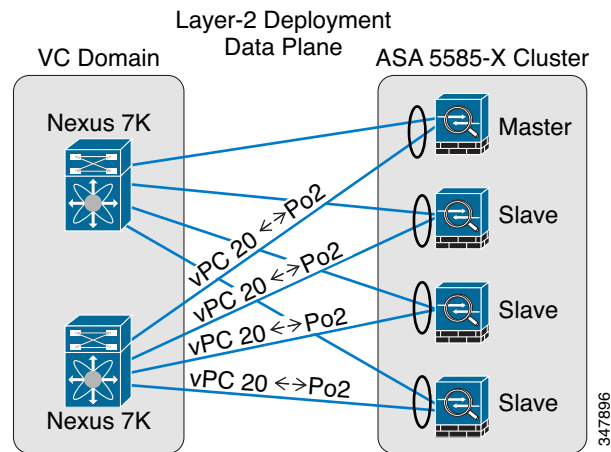
ASA Port Channels are "spanned" across units
to form single logical unit connecting to Nexus 7000 vPCs

Cluster Link Aggregation Control Protocol

Cisco has incorporated enhancements to the LACP, while maintaining complete interoperability with other devices complying with the IEEE 802.3ad LACP standard. These extensions are referred to as Cluster Link Aggregation Protocol (cLACP), and although the extensions are transparent in their use, they provide significant capabilities to the ASAs operating as a cluster. Designed to extend standard LACP to multiple devices, the implementation of cLACP is intended to be used on port channels that are connecting to switches operating in a vPC or VSS configuration.

The ASAs can operate in the cluster as a single unit by applying the cLACP "span-cluster" option in the port channel configuration. This enables the port channels across all of the units being used for the data connections to operate as if they are in the same port channel. All of the ASAs in the cluster share a single system IP and system MAC, so they appear as a single device in the network. In cluster Ethernet, Cisco strongly recommends that users configure a virtual MAC on the span-cluster port channel to make the port channel MAC stable in cluster. This prevents the cluster from becoming unstable because of the master unit leaving or joining the cluster. Because the configuration is replicated from the master to the slave units, the system's virtual MAC is persistent through any cluster configuration changes.

Automatic Management of Active Links

Typical EtherChannel deployments have a limitation of 8 active links. Beginning with Cisco NX-OS Release 5.1, you can bundle up to 16 active links into a port channel on the F Series module per switch. The ASA cluster can have up to 16 ASAs in the cluster, which requires 32 links connected to the switching fabric (see Figure 10). When the ASA is deployed in a vPC configuration, cLACP manages the 32 links so that the proper load balancing occurs and all 32 links can be active (as of ASA code release 9.2.1). This enables the ASA Cluster to achieve maximum throughput while being deployed

with optimal resiliency. When load balancing is enabled, cLACP assumes that the physical link connections between the ASA Cluster to the pair of Nexus 7000 vPC peers are balanced.

Figure 10 **ASA 5585-X Cluster**
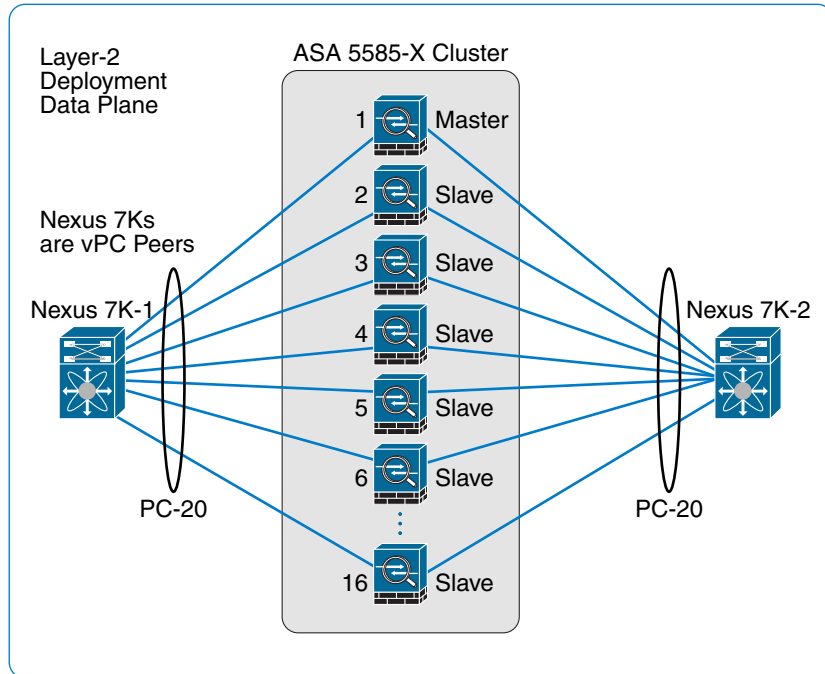
Firewall Modes

The ASA can be partitioned into multiple virtual devices, known as security contexts. Each context acts as an independent device, with its own security policy, interfaces, configuration, and administrators. Multiple contexts are similar to having multiple standalone devices.

Routed Firewall Mode

In routed firewall mode, the ASA is considered to be a router hop in the network. Routed firewall mode supports many interfaces. Each interface is on a different subnet. You can share interfaces between contexts. The ASA acts as a router between connected networks, and each interface requires an IP address on a different subnet. The ASA supports multiple dynamic routing protocols. However, Cisco recommends using the advanced routing capabilities of the upstream and downstream routers instead of relying on the ASA for extensive routing needs.

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, you can easily introduce a transparent firewall into an existing network. The Management and Data interfaces should not be 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 (or 250 bridge groups in 9.3+ software). Each bridge group can include up to four interfaces.

Note The transparent mode ASA does not pass CDP packets, or any packets that do not have a valid EtherType greater than or equal to 0x600. An exception is made for BPDUs and IS-IS, which are supported.

Mixed Firewall Mode

The ASA supports running multiple firewall modes independently for each security context in a multiple context mode (see Figure 11). This enables some of the security contexts to run in transparent mode while others are running in routed mode.

A critical benefit of the ASA supporting mixed firewall modes is that customers can now consolidate their Layer-2 and Layer-3 firewalls into a single ASA cluster.

Note that multiple context mode does not support the following features:

- RIP
- OSPFv3 (OSPFv2 is supported.)
- Multicast routing
- TLS Proxy and its derivatives, such as Phone Proxy and IME
- QoS
- Remote access VPN (site-to-site VPN is supported)

Figure 11   *Multiple Firewall Mode*
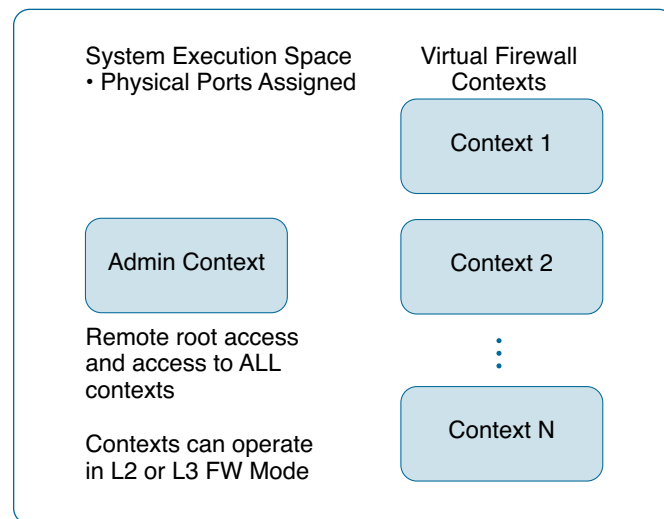
Cluster Configuration

Although the section above describes how a single ASA can be configured with multiple virtual firewalls, when put into a cluster, all ASA 5585-X units within the cluster share a single configuration. When configuration changes on the master unit are made, the changes are automatically replicated to

all slave units in the cluster. A configuration change directly made on slave units is prohibited.

Cluster Units Add/Removal

ASA units can be added or removed from a cluster at runtime. A small set of cluster bootstrap commands must be configured on a unit before it can be brought online. Cisco recommends that you deploy the cluster with an IP address scheme that accommodates future growth. For example, if the cluster is going to start with two units, use the following command:

```
ip local pool mgmt-pool 10.11.235.21-10.11.235.28 mask 255.255.255.0
```

As you add units to the cluster, the IP address scheme does not need to change because it is covered in the.21–.28 address range.

Management Network

All units in the cluster must be connected to a management network that is not the same as the CCL. The ASA-5585-X has dedicated management interfaces and Cisco highly recommends using these ports.

Each ASA is assigned a unique IP address, and a system IP is assigned to the master unit as its secondary IP address.

For inbound management traffic, an application such as Cisco Security Manager can access the master ASA by using the system IP address or individual ASA by its own IP address. For outbound traffic, such as TFTP or syslog, each ASA uses its own IP address to connect to the server. In multi-context mode, the same configuration applies to the admin context and any user contexts that allow remote management.

Cluster Roles For Connections (per Connection)

The following ASA roles are defined for each connection:

- Owner—The unit that receives the first packet in the connection assumes the *owner* role and each connection has only one owner. The owner maintains the TCP state and processes packets.

- Director—When the owner receives a new connection, it chooses a *director* based on a hash of the source/destination IP address and TCP/UDP ports, and sends a message to the director to register the new connection. The director unit handles owner lookup requests from forwarders and also maintains the connection state to serve as a backup if the owner fails. If packets arrive at any unit other than the owner, the unit queries the director about which unit is the owner so it can forward the packets. A connection has only one director.

- Backup director—Only needed when the owner and director are the same unit.

- Forwarder—A unit that forwards packets to the owner. If a *forwarder* receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder.

Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director (if you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required). For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

ASA Cluster Data Flows

The ASA needs to see packets from both directions of a connection to keep track of state and apply protocol conformance tests. The Cisco ASA 5585-X Clustering feature removes the need to have a stateful load balancer in front of the ASA Cluster. This solution still relies on an external load balancer to distribute traffic, which can easily be performed by the switches that connect to the ASA units; however, no assumption is made on how the traffic is distributed. In particular, it is not assumed that the external balancer will distribute traffic evenly or send all packets of a connection to the same ASA. Figure 12 shows a sample of an establishment of a new connection.

Figure 12   *Establishment of a New Connection*
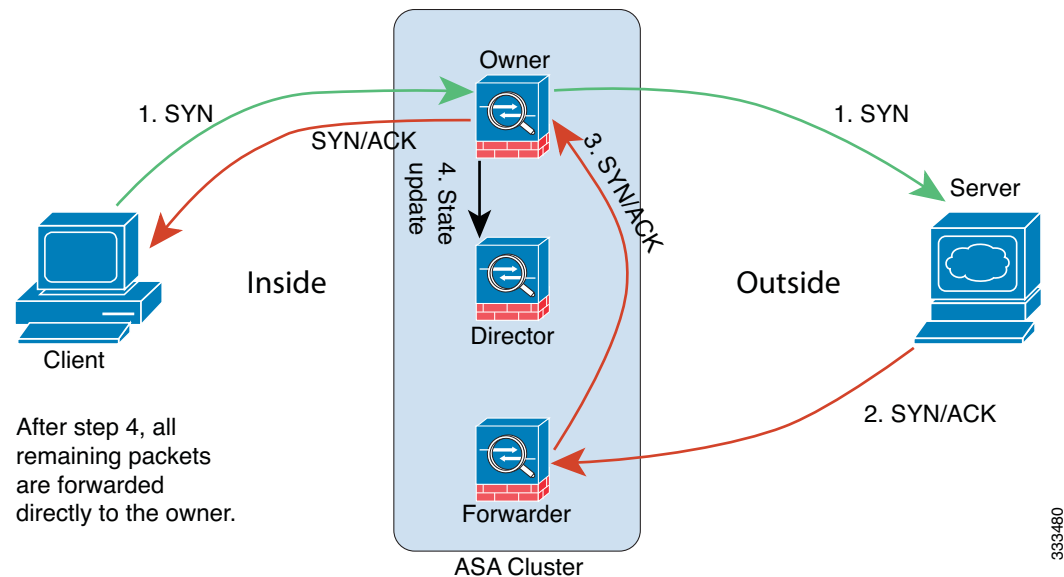
1. The SYN packet originates from the client and is delivered to an ASA (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.

2. The SYN-ACK packet originates from the server and is delivered to a different ASA (based on the load balancing method). This ASA is the forwarder.

3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.

4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.

5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.

6. Any subsequent packets delivered to the forwarder are forwarded to the owner.

7. If packets are delivered to any additional units, it queries the director for the owner and establishes a flow.

8. Any state change for the flow results in a state update from the owner to the director.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. TCP/UDP state information is replicated from owner to the backup. If the owner becomes unavailable, a switchover event is broadcast to remaining units, and a connection gets a new owner who retrieves the relevant state from the backup owner. Note that the selection of the new owner depends on how traffic is re-balanced by the external load balancer after failure. The backup owner is serving as a state repository. The state of some connections is lost permanently if more than one unit fails at the same time.

Unit Health Monitoring

The master unit monitors every unit in the cluster by sending keepalive messages over the cluster link. When the ASA interfaces are operating in spanned EtherChannel mode, the unit monitors the cLACP messages and reports a link status back to the master. With health monitoring enabled, the failed units are removed from the cluster automatically. If the master unit fails, another member of the cluster with the highest priority assumes the master role. From version 9.4(1), particular interface/s can be set to be exempt from the monitoring.

Connections Impact on Device Failure

Connections that do not require additional state above the TCP/UDP layer survive the switchover. Examples include SSH, HTTP, and SMTP. Other connections may get dropped because of the lack of a higher-level protocol state.

The following summarizes the centralized features that are dropped if the master unit fails:

- Site-to-site VPN
- The following application inspections:
    - DCERPC
    - NetBios
    - PPTP
    - RADIUS
    - RSH
    - SUNRPC
    - TFTP
    - XDMCP
- Dynamic routing (spanned EtherChannel mode only)
- Multicast routing (individual interface mode only)
- Static route monitoring
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Authentication and authorization for network access. Accounting is decentralized.
- URL filtering services

Cluster and Failover Are Mutually Exclusive

Failover and clustering cannot be enabled at the same time, which means:

- Failover from an active cluster to a standby cluster is not supported.
- Cluster of failover pair (that is, each cluster member is a failover pair) is not supported.

Syslog and NetFlow

Syslog and NetFlow are valuable tools for accounting, monitoring, and troubleshooting in such a high throughput environment. ASA units in the cluster generate syslogs independently. The syslog's HEADER field, which contains a timestamp and device ID, can be customized as required. A syslog collector uses the device ID to identify the syslog generator. The CLI is enhanced so that different ASAs can generate syslog with identical or different device ID. However, a per-unit NetFlow stream cannot be consolidated. The NetFlow collector handles each individual ASA separately.

NetFlow and Clustering

NetFlow is supported on both management and regular data interfaces. However, Cisco recommends using the management interfaces. When the NetFlow collector connection is configured on management-only interfaces, each ASA in the cluster uses its own per-unit source IP address and source port to send NetFlow packets. NetFlow may be used with both data interfaces in Layer-2 mode and Layer-3 mode.

For data interfaces in Layer-2 mode, each ASA in the cluster has the same source IP address but the source port is different. Although Layer-2 mode is designed to make a cluster appear as a single device, a NetFlow collector can differentiate between the different nodes in the cluster. For data interfaces in Layer-3 mode, NetFlow operates the same way as management-only interfaces. Each ASA node in the cluster establishes its own connection to the NetFlow collector(s) and advertises its templates independently. The collector uses the source IP address and source port of the packet to differentiate between the NetFlow exporters.

SNMP

An SNMP agent polls each individual ASA by its unique management IP address. The data returned represents the local state of that ASA. The agent needs to consolidate statistics by itself if it wants to present a system-wide view to the end user.

In failover, the engine ID is synced between the active and standby device. This is needed because when the standby switches to active, it takes over the active's IP address. The SNMP agent assumes a consistent IP-to engine-ID mapping. This is not needed because each unit in the cluster has its own unique management IP address. However, there is a system IP that floats with the master unit. If the SNMP agent uses the system IP to poll, and a new master is elected afterwards, the poll fails. The workaround is to always use a unit-unique management IP address.

Firewall Features With Special Behavior

The features described in this section are applied to the traffic/state that is local to an ASA unit, instead of the whole cluster.

QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each unit independently. For example, with "police output 56000 10500", the conform-rate

and conform-burst values are enforced on traffic going out a particular ASA. In the cluster with eight units and with traffic evenly distributed, the conform-rate actually becomes 56000*8=448000 from the system's point of view.

Threat detection works on each unit independently; for example, the Top N statistics are unit-specific. Port scanning detection, for example, does not work because scanning traffic is load balanced between all units, and no one sees it all.

There is no configuration sync or state sharing between IPS modules. Some IPS signatures require IPS to keep the state across multiple connections. For example, the port scanning signature is used when the IPS module detects that someone is opening many connections to one server but with different ports. In clustering, those connections are balanced between multiple ASA devices, each of which has its own IPS module. Because these IPS modules do not share state information, the cluster may not be able to detect port scanning.

Unsupported Features

The following features cannot be configured with clustering enabled, so the commands are rejected:

- TLS Proxy and its derivatives, such as Phone Proxy or IME
- Remote access VPN (SSL VPN and IPsec VPN)
- The following application inspections:
    - CTIQBE
    - GTP
    - H323, H225, and RAS
    - IPsec passthrough
    - MGCP
    - MMP
    - RTSP
    - SCCP (Skinny)
    - WAAS
    - WCCP
- Botnet traffic filter
- Auto update server
- DHCP client, server, and proxy
- VPN load balancing
- Failover
- ASA CX module

Cisco Security Manager

Cisco Security Manager still manages the policies on the ASAs as would be expected in a traditional deployment model. Security groups are downloaded from the ISE in the environment data after the Cisco Security Manager has established a secure connection by importing a PAC file from the ISE. As described above, the Cisco Security Manager issues a RADIUS request for the TrustSec environment

data, which includes the secure group table mapping secure group names to secure group numbers and are presented as secure group objects. After the environment data is downloaded, creating policies for the firewall is similar to creating extended ACLs.

Firewall Policies

The following policies allow you to configure security groups:

- AAA rules
- Access rules
- Inspection rules

Several firewall policies use extended ACL policy objects to define traffic matching criteria instead of incorporating a rule table directly in the policy. You can configure extended ACL policy objects to include security group specifications. You can then use these extended ACL objects in the following policies:

- Botnet traffic filter rules
- IPS, QoS, and connection rules (service policy rules)

NextGen IPS Overview

The ASA FirePOWER Module

The ASA FirePOWER module supplies next-generation firewall services, including Next-Generation IPS (NGIPS), Application Visibility and Control (AVC), URL filtering, and Advanced Malware Protection (AMP). The module can be used in single or multiple context mode, and in routed or transparent mode.

The module has a basic command line interface (CLI) for initial configuration and troubleshooting. Configuration of the security policy on the device is performed using a separate application, FireSIGHT Management Center, which is hosted on a separate FireSIGHT Management Center appliance or as a virtual appliance running on a VMware server. (FireSIGHT Management Center was formerly known as Defense Center.)

More information about the FirePOWER services module can be found in the *Cisco ASA FirePOWER Module Quick Start Guide* at the following URL:
http://www.cisco.com/c/en/us/td/docs/security/asa/quick_start/sfr/firepower-qsg.html

How the ASA FirePOWER Module Works with the ASA

The ASA FirePOWER module runs a separate application from the ASA. The module can be a hardware module (on the ASA 5585-X) or a software module (5512-X through 5555-X). As a hardware module, the device includes separate management and console ports, and extra data interfaces that are used directly by the ASA and not by the module itself.

The module can be configured in either a passive ("monitor only") or inline deployment.

- In a passive deployment, a copy of the traffic is sent to the device, but it is not returned to the ASA. Passive mode lets you see what the device would have done to traffic, and evaluate the content of the traffic, without impacting the network.

- In an inline deployment, the actual traffic is sent to the device, and the device's policy affects what happens to the traffic. After dropping undesired traffic and taking any other actions applied by policy, the traffic is returned to the ASA for further processing and ultimate transmission.

ASA FirePOWER Inline Mode

In inline mode, traffic goes through the firewall checks before being forwarded to the ASA FirePOWER module. When traffic is identified for ASA FirePOWER inspection on the ASA, traffic flows through the ASA and the module as follows:

1. Traffic enters the ASA.
2. Incoming VPN traffic is decrypted.
3. Firewall policies are applied.
4. Traffic is sent to the ASA FirePOWER module.
5. The ASA FirePOWER module applies its security policy to the traffic, and takes appropriate actions.
6. Valid traffic is sent back to the ASA; the ASA FirePOWER module might block some traffic according to its security policy, and that traffic is not passed on.
7. Outgoing VPN traffic is encrypted.
8. Traffic exits the ASA.

Figure 13 shows the traffic flow when using the ASA FirePOWER module in inline mode. In this example, the module blocks traffic that is not allowed for a certain application. All other traffic is forwarded through the ASA.

Figure 13   ASA FirePOWER Module Traffic Flow in the ASA

Note If you have a connection between hosts on two ASA interfaces, and the ASA FirePOWER service policy is configured for only one of the interfaces, all traffic between these hosts is sent to the ASA FirePOWER module, including traffic originating on the non-ASA FirePOWER interface (because the feature is bidirectional).

ASA FirePOWER Passive (Monitor-Only) Mode

The traffic flow in monitor-only mode is the same as it is for inline mode. The only difference is that the ASA FirePOWER module does not pass traffic back to the ASA. Instead, the module applies the security policy to the traffic and lets you know what it would have done if it were operating in inline

mode; for example, traffic might be marked "would have dropped" in events. You can use this information for traffic analysis and to help you decide whether inline mode is desirable.

To configure passive mode, include the monitor-only indication on the service policy that redirects traffic to the module.

Note You cannot configure both monitor-only mode and normal inline mode at the same time on the ASA. Only one type of security policy is allowed. In multiple context mode, you cannot configure monitor-only mode for some contexts, and regular inline mode for others.

Figure 14 shows the traffic flow when operating in passive mode.

Figure 14   *ASA FirePOWER Passive, Monitor-Only Mode*
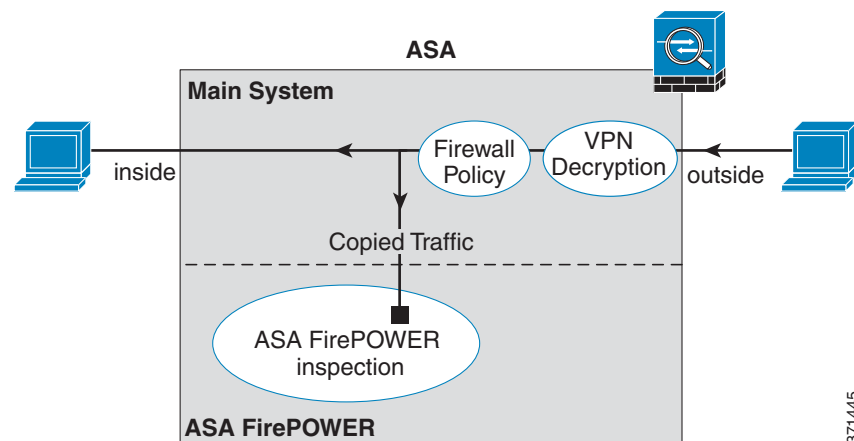
ASA FirePOWER Management Access

There are two separate layers of access for managing an ASA FirePOWER module: initial configuration (and subsequent troubleshooting), and policy management.

For initial configuration, you must use the CLI on the ASA FirePOWER module. To access the CLI, you can use the following methods:

- ASA 5585-X:
  - ASA FirePOWER console port—The console port on the module is a separate external console port.
  - ASA FirePOWER Management 1/0 interface using SSH—You can connect to the default IP address (192.168.45.45/24) or you can use Cisco Adaptive Security Device Manager (ASDM) to change the management IP address and then connect using SSH. The management interface on the module is a separate external Gigabit Ethernet interface.

Note You cannot access the ASA FirePOWER hardware module CLI over the ASA backplane using the **session** command.

- ASA 5512-X through ASA 5555-X:

- – ASA session over the backplane—If you have CLI access to the ASA, you can session to the module and access the module CLI.

- – ASA FirePOWER Management 0/0 interface using SSH—You can connect to the default IP address (192.168.45.45/24) or you can use ASDM to change the management IP address and then connect using SSH. These models run the ASA FirePOWER module as a software module. The ASA FirePOWER management interface shares the Management 0/0 interface with the ASA. Separate MAC addresses and IP addresses are supported for the ASA and ASA FirePOWER module. You must perform configuration of the ASA FirePOWER IP address within the ASA FirePOWER operating system (using the CLI or ASDM). However, physical characteristics (such as enabling the interface) are configured on the ASA. You can remove the ASA interface configuration (specifically the interface name) to dedicate this interface as an ASA FirePOWER-only interface. This interface is management-only.

After you perform initial configuration, configure the ASA FirePOWER security policy using FireSIGHT Management Center. Then configure the ASA policy for sending traffic to the ASA FirePOWER module using ASDM or Cisco Security Manager.

Guidelines for ASA FirePOWER

Compatibility with ASA Features

The ASA includes many advanced application inspection features, including HTTP inspection. However, the ASA FirePOWER module provides more advanced HTTP inspection than the ASA provides, as well as additional features for other applications, including monitoring and controlling application usage.

To take full advantage of the ASA FirePOWER module features, see the following guidelines for traffic that you send to the ASA FirePOWER module:

- • Do not configure ASA inspection on HTTP traffic.

- • Do not configure Cloud Web Security (ScanSafe) inspection. If you configure both ASA FirePOWER inspection and Cloud Web Security inspection for the same traffic, the ASA performs only ASA FirePOWER inspection.

- • Other application inspections on the ASA are compatible with the ASA FirePOWER module, including the default inspections.

- • Do not enable the Mobile User Security (MUS) server; it is not compatible with the ASA FirePOWER module.

- • If you enable failover, when the ASA fails over, any existing ASA FirePOWER flows are transferred to the new ASA. The ASA FirePOWER module in the new ASA begins inspecting the traffic from that point forward; old inspection states are not transferred.

ASA Clustering and Failover Guidelines

The FirePOWER modules do not support clustering directly, but can be used in a cluster as validated in this solution. You are responsible for maintaining consistent policies on the ASA FirePOWER modules in the cluster using FireSIGHT Management Center. Do not use different ASA-interface-based zone definitions for devices in the cluster.

When an ASA in the cluster fails, any existing ASA flows are transferred to the new ASA, however the FirePOWER module in the new ASA begins inspecting the traffic from that point forward; old inspection states are not transferred.

Model Guidelines

- Minimum software requirements are ASA Software 9.2(2)4 and ASA FirePOWER 5.3.1.
- Supported on the ASA 5585-X (as a hardware module) and 5512-X through ASA 5555-X (as a software module). See the *Cisco ASA Compatibility Matrix* for more information: http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html
- For the 5512-X through ASA 5555-X, you must install a Cisco solid state drive (SSD). For more information, see the ASA 5500-X hardware guide.

Additional Guidelines and Limitations

- See Compatibility with ASA Features: http://www.cisco.com/c/en/us/td/docs/security/asa/quick_start/sfr/firepower-qsg.html#96600.
- You cannot change the software type installed on the hardware module; if you purchase an ASA FirePOWER module, you cannot later install other software on it.
- You cannot configure both monitor-only mode and normal inline mode at the same time on the ASA. Only one type of security policy is allowed. In multiple context mode, you cannot configure monitor-only mode for some contexts, and regular inline mode for others.

Cisco TrustSec

The Cisco TrustSec solution (see Figure 15) provides the ability to create policies to map end users, or consumers, to data center assets, or servers and applications. Typical policies for securing the data center are the 5-tuple or even recently, context-based policies. These policies have been placed at the edge of the data center in a border-based architecture. TrustSec enables you to create policies that are much deeper than just roles-based or a 5-tuple-based approach, all while keeping a defense-in-depth architecture with enforcement points integrated throughout the fabric. 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.

Figure 15 **Cisco TrustSec**
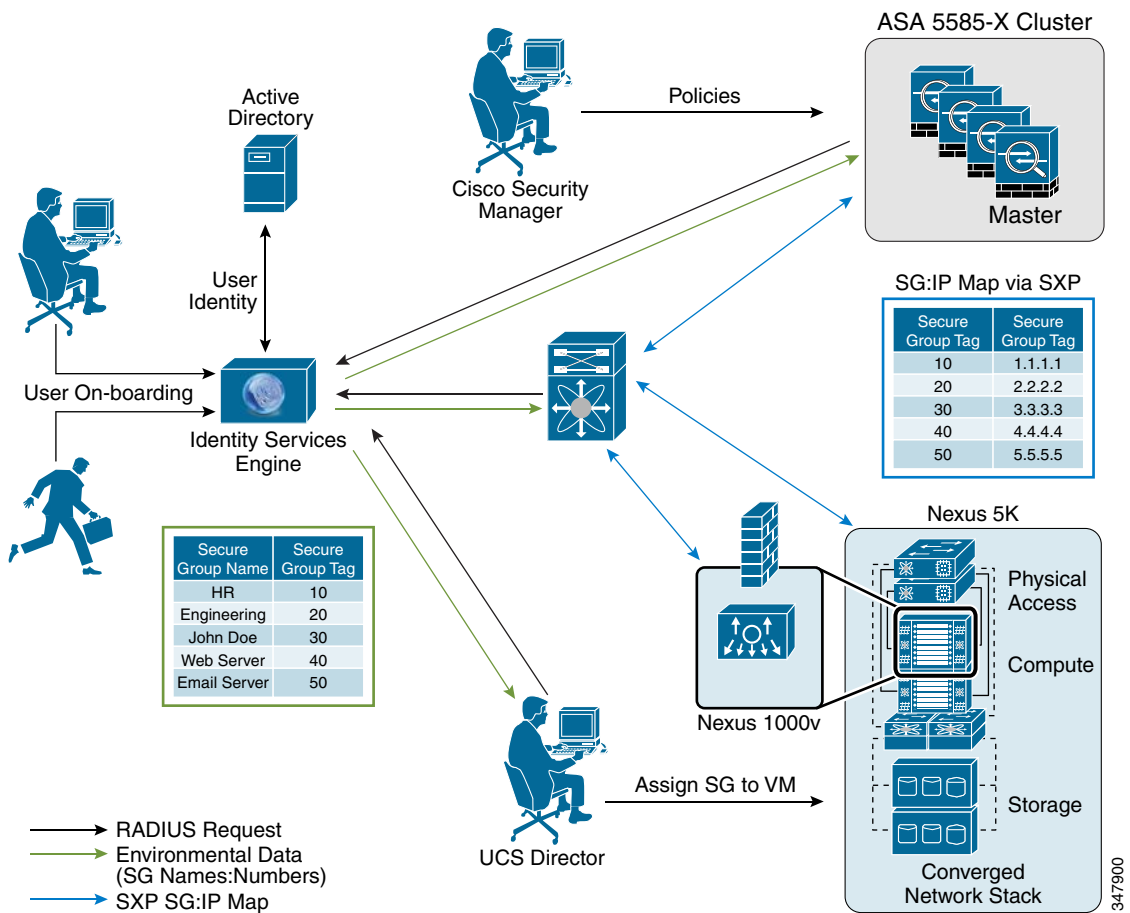
Cisco Identity Service Engine

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. In Figure 15, 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, such as secure groups, secure group names, and security group ACLs (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 lists (RBACLs). 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 an SGT, which 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 5 lists examples of secure group names and their respective SGTs.

Table 5   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.

In Figure 15, SXP flows between the Nexus 7000 and the ASA 5585-X, the Nexus 5000 access switch, and the Nexus 1000V. The Nexus 5000 and the Nexus 1000V act as "speakers" and send the SGT and IP address mapping back to the Nexus 7000 via SXP. The Nexus 7000 then sends the SGT-to IP address map to the ASA 5585-X, again via SXP.

Table 6 lists examples of SGTs and their respective IP addresses.

Table 6   Secure Group Tags and IP Addresses

Secure Group Tag	IP Address
10	1.1.1.1
20	2.2.2.2
30	3.3.3.3
40	4.4.4.4
50	5.5.5.5

SXP Compatibility and Caveats

Network Address Translation

NAT cannot be used for SXP peer communication. SXP conveys SGT-to-IP address mappings to enforcement points in the network. If the access layer switch belongs to a different NAT domain than the enforcing point, the SGT-to-IP address map it uploads is meaningless, and an SGT-to-IP address database lookup on the enforcement device yields nothing. This means it is not possible to apply identity-based (security-group-aware) ACLs on the enforcement device.

Through the Box SXP

Through-the-box transit SXP connections break if NAT is caused by TCP sequence number randomization and TCP option 19 stripping. To allow these connections, the following configuration is necessary:

```
class bypass
            set connection random-sequence-number disable
            set connection advanced-options sxp-tcp-map
tcp-map sxp-tcp-map
            tcp-options range 19 19 allow
```

Network Device Authorization

For network devices, management platforms, and network services, such as ASA firewalls, to join the TrustSec domain, they import a protected access credential (PAC) file from the ISE. Importing the PAC file to the network establishes a secure communication channel with the ISE. After the channel is established, the network device initiates a PAC-secure RADIUS transaction with the ISE and downloads Cisco TrustSec environment data; specifically, the ASA downloads the security group table. The security group table maps SGTs to security group names. Security group names are created on the ISE and provide user-friendly names for security groups.

Security Group ACLs

Security group ACLs (SGACLs, also known as RBACLs) are access control lists that restrict the operations that a user can perform based on the role of the user instead of the IP address or subnet mask alone. SGACLs allow you to control the access and permissions based on the SGTs that are assigned. The grouping of permissions into a role simplifies the management of security policy. As you add devices, you simply assign one or more security groups, and they immediately receive the appropriate permissions. You can modify the security groups to introduce new privileges or restrict current permissions. The SGACLs are enforced by the switches that have the SGACL enforcement enabled.

Configuring of the SGACL policies is done on the ISE and although it is not recommended, SGACLs can be manually provisioned on the switch. Any SGACL policy downloaded dynamically from the ISE overrides any conflicting locally-defined policy.

These are not to be confused with the secure group firewall access lists on the ASA. The SGACLs apply only to switches that are part of the TrustSec domain with *role-based enforcement* enabled.

ASA and TrustSec

Beginning with Cisco ASA Software Release 9.0.1, the ASA firewall gains Secure Group Firewall (SGFW) functionality. Policy in the firewall has been expanded to include source and destination security groups that are downloaded from the ISE in the environment data after the ASA has established a secure connection by importing a PAC file from the ISE. As described above and shown in Figure 15, the ASA issues a RADIUS request for the TrustSec environment data, which includes the secure group table mapping secure group names to secure group numbers. The ASA receives the secure

group numbers-to-secure group IP addresses mapping from the Nexus 7000 via SXP. If the PAC file downloaded from the ISE expires on the ASA and the ASA cannot download an updated security group table, the ASA continues to enforce security policies based on the last downloaded security group table until the ASA downloads an updated table.

Multi-Context Mode

Both single-context and multi-context modes are supported. Each context maintains its own configurations, databases, credentials, and environment data.

Firewall Mode

Both routed and transparent modes are supported. In transparent mode, each user context typically has an inside interface, an outside interface, and a management interface. You can assign an IP address to the management interface, or the inside and outside interfaces can be grouped into a bridge-group virtual interface (BVI) and you can assign an IP address to the BVI. This IP address should be used in the SXP communication with peer devices.

Clustering

Clustering is supported. The master unit contacts ISE and obtains environment data, which is then replicated to all units in the cluster via reliable messaging.

Security group-based policies are replicated as part of the configuration sync. The master unit establishes SXP connections and learns secure group-to-IP address mappings. This SXP mapping database is replicated to all units. Thus security group-based policies can be enforced on the slave units.

Scalability

Table 7 lists the number of IP-SGT mapped entries supported by ASA.

Table 7   Supported Number of IP-SGT Mapped Entries

| ASA Platform	SXP Connections	IP-SGT Mappings
ASA5585-SSP10	150	18,750
ASA5585-SSP20	250	25,000
ASA5585-SSP40	500	50,000
ASA5585-SSP60	1000	100,000

The ASA cluster configuration is performed via CLI or Cisco ASDM. Policies for the firewalls are managed via ADSM or CSM. Policies for the FirePOWER blades are managed in FireSIGHT Management Center. User/server device objects are managed in ISE along with TrustSec policy creation for remaining platforms. User accounts and authentication are linked to Active Directory.

Solution Component Implementation

This implementation guide focuses on how the ASA clustering feature can be extended between multiple data centers that are geographically separated. Overlay Transport Virtualization (OTV) was chosen to be the method of extending the internal VLANs across data centers.

Figure 16 portrays an overview of the lab deployment used for validation. The following sections show how each product was configured to match specific use cases desired in the validation.

Figure 16   Lab Overview

At each data center site, four Cisco ASA 5585-X SSP60 firewalls with integrated FirePOWER service modules are deployed as a single spanned cluster with Layer 2 mode contexts (multi-mode). The aggregation switch shown in Figure 16 above the clustered ASA firewall is running Layer 3, and switches and servers connected under the ASA are running Layer 2; therefore, the immediate next network hop from the server is the aggregation switch.

Multi-site Design Consideration

When designing the multi-site DC deployment, there are a few factors that administrators need to consider. You need to identify the existing design, such as Layer 3 and Layer 2, and which network and VLANs need to be extended among multiple sites. Often the OTV is used to communicate among the multiple sites. Dark fiber, or fiber that is currently unused but is a leasable and affordable dedicated fiber link owned or controlled by traditional carriers, is another method to establish the communication among data centers. This design includes a Layer 2 setup of a Cisco ASA Firewall connected to a Cisco Nexus 7000 aggregation switch as a default gateway. OTV is the choice for extending the L2 network among the data centers.

The design also includes a Hot Standby Router Protocol (HSRP) setup in the aggregation switch.

VLAN IP addresses were set up as follows:

Aggregation Switch 1 in Site 1

```
interface Vlan2001
  no shutdown
  ip address 10.11.1.254/24
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.1.1
```

Aggregation Switch 2 in Site 1

```
interface Vlan2001
  no shutdown
  ip address 10.11.1.253/24
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.1.1
```

Aggregation Switch 1 in Site 2

```
interface Vlan2001
  no shutdown
  ip address 10.11.1.252/24
  ip ospf passive-interface
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.1.1
```

Aggregation Switch 2 in Site 2

```
interface Vlan2001
  no shutdown
  ip address 10.11.1.251/24
ip ospf passive-interface
ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
  ip 10.11.1.1
```

Because both sites retain the same default gateways, filtering on the HRSP MAC address and VLAN access list on the OTV switch is required to prevent any network loops.

The design uses the Open Shortest Path First (OSPF) dynamic routing method. One of the caveats when running OSPF is that it is required to set one side of the data center to be a passive interface. See the OSPF configuration above where both vlan 2001 configurations in the site 2 aggregation switches are set as passive interfaces.

OTV Configuration

OTV configuration was performed via the console command line in the core, aggregation, and OTV switches.

For additional information of OTV, see the following resources:

- Cisco Nexus 7000 Series OTV Quick Start Guide—
  http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/OTV/quick_start_guide/b-Cisco-Nexus-7000-Series-OTV-QSG.html#reference_B0DE9A81C0FA44CA9D95347AA32E8E38

- Cisco Nexus 7000 Series NX-OS OTV Configuration Guide—
  http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/nx-os/OTV/config_guide/b_Cisco_Nexus_7000_Series_NX-OS_OTV_Configuration_Guide.html

Following are some guidelines for configuring OTV:

- OTV is currently supported by Nexus 7000 Series Switches with all M series or F3 line cards only.

- Beginning with Cisco NX-OS Release 6.2, OTV supports the coexistence of F1 or F2e Series modules with M1 or M2 Series modules in the same VDC.

- OTV requires the Transport Services license on the Nexus 7000.

- You can choose to extend multicast or unicast only mode.

- ASA routed mode is not supported as of ASA 9.3(2).

- OTV is compatible with a transport network configured only for IPv4.

- Enable IGMPv3 on the join interfaces.

- Ensure connectivity for the VLANs to be extended to the OTV edge device.

- If you configure VDCs, install the Advanced Services license and enter the desired VDC.

Figure 17 shows an illustration of the switch configuration described in the following sections.

Figure 17   ***Switch Configuration***

Configuring the Core Switches

First, make sure the core routers are configured for multicast communication.

Procedure

Step 1   The Core-1 router needs to support multicast connection among OTV and aggregation switches.

   a.   Enable the multicast feature：

```
feature pim
```

   b.   Set the rp address. This design used the following loopback IP address:

```
                    ip pim rp-address 10.11.255.1
```

   c. On Ethernet 3/8, 3/1 and 3/2, enable ip pim sparse-mode：

```
interface Ethernet3/8,e3/1,e3/2
  ip pim sparse-mode
  no shutdown
```

Step 2   Configure the Core-2 router the same as the Core-1 router.

   a. Set the rp address of Core-1 loopback

```
ip pim rp-address 10.11.255.1
```

   b. On Ethernet 3/8, 3/1, and 3/2, enable ip pim sparse-mode.

Configuring the OTV Switches

The next step is to configure the OTV switches. In this example, a virtual device context (VDC) is created for this function. See the configurations in the Appendix for VDC and allocated interfaces for the OTV switch.

Perform the following steps to configure OTV-1.

Procedure

Step 1   Enable the multicast feature and OTV feature:

```
feature pim, feature otv
```

Step 2   Configure the VLAN that needs to be advertised by OTV:

```
vlan 1,20-24,2000-2100,2201-2300,3001-3100,3150,3201-3400
```

Step 3   Configure the OTV site-vlan:

```
otv site-vlan 3150
```

Note   VLAN 3150 will be used for OTV purposes, so make sure to include this VLAN in the site-interface configurations in the OTV and aggregation switches.

Step 4   Create the OTV Layer 3 join-interface as follows:

```
interface port-channel150
  mtu 9216
  ip address 10.11.210.73/30
  ip ospf network point-to-point
  ip router ospf 5 area 0.0.0.0
  ip igmp version 3
```

Step 5   Set port-channel 150 to physical interfaces:

```
interface Ethernet3/7
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet3/8
  mtu 9216
  channel-group 150 mode active
  no shutdown
```

Note   Make sure to set channel-group mode to be active to enable LACP.

Step 6   Configure OTV Layer 2 site-interfaces:

```
interface port-channel151
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
```

Step 7   Set port-channel 151 to physical interfaces:

```
interface Ethernet3/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown

interface Ethernet3/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown
```

Note   Make sure to set channel-group mode to be active to enable LACP.

Step 8   Configure the OTV site-identifier:

```
otv site-identifier 0x1
```

Step 9   Configure the OTV VLAN ACL and MAC filtering.

In both data centers, HSRP provides default gateway redundancy for connected hosts. Each VLAN on the AGG-Edge that requires L3 connectivity to the rest of the network is configured with an HSRP gateway address. Because the VLANs are extended, there is a possibility of HSRP flapping between data centers. To enable hosts in the extended VLANs to use their local HSRP gateway, an IP gateway localization technique is used to keep HSRP protocol data units from getting forwarded on the overlay network. This technique uses a combination of VLAN access control lists (VACLs) and OTV MAC route filters in the OTV VDC.

The following VACL configuration is used on all OTV Edge devices to drop HSRP hello packets:

```
!
```

```
ip access-list ALL_IPs
10 permit ip any any
mac access-list ALL_MACs
10 permit any any
ip access-list HSRP_IP
10 permit udp any 224.0.0.2/32 eq 1985
mac access-list HSRP_VMAC
10 permit 0000.0c07.ac00 0000.0000.00ff any
vlan access-map HSRP_Localization 10
match mac address HSRP_VMAC
match ip address HSRP_IP
action drop
vlan access-map HSRP_Localization 20
match mac address ALL_MACs
match ip address ALL_IPs
action forward
vlan filter HSRP_Localization vlan-list 2000-2100, 3000-3100
!
```

Step 10 To prevent the HSRP MAC address and ASA contexts' virtual MAC address from flapping between sites, a route map is applied on the overlay interface to filter the virtual MAC address.

> ✎
>
> **Note** You can find the ASA virtual MAC address in the each contexts' inside and/or outside interfaces by executing *show interface inside*.

```
!
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 11 deny 78da.6ed9.767e ffff.ffff.ffff
mac-list OTV_HSRP_VMAC_deny seq 20 permit 0000.0000.0000 0000.0000.0000
route-map OTV_HSRP_filter permit 10
match mac-list OTV_HSRP_VMAC_deny
otv-isis default
vpn Overlay1
redistribute filter route-map OTV_HSRP_filter
!
```

Step 11 Configure OTV Overlay interface

```
interface Overlay1
  otv join-interface port-channel150
  otv control-group 239.1.0.1
  otv data-group 232.1.0.0/16
  otv extend-vlan 20-24, 2000-2100, 2201-2300, 3001-3100, 3201-3400
  no otv suppress-arp-nd
  no shutdown
```

Configure OTV-2 the same as OTV-1 except to use port-channel 250 for the join-interface and the Overlay 1 interface.

Configuring the Aggregation Switches

The final step is to configure aggregation switches.

To this for RAGG-1, perform the following steps.

Procedure

Step 1 Enter the core loopback IP address as the rp address:

```
ip pim rp-address 10.11.255.1 group-list 224.0.0.0/4
```

Step 2   Create port channel 150, 151, and 251 (RAGG-1); and 250,251, and 151 (RAGG-2).

Note the following:

- The MTU needs to be set for 9216 on the port-channel and applied physical interfaces

- Port-Channel 150 and 250 are set as the Layer 3 OTV join-interface.

- Port-Channel 151 and 251 are set as the Layer 2 OTV site-interface.

- Make sure that Port-channel 150 and 250 are configured with *ip pim sparse-mode* and *ip igmp version 3*

```
interface port-channel150
  mtu 9216
  ip address 10.11.210.74/30
  ip ospf network point-to-point
  no ip ospf passive-interface
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
  ip igmp version 3

interface port-channel151
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  vpc 151

interface port-channel251
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  vpc 251
```

Step 3   Set Port-Channel 150, 151, and 251 to physical interfaces:

```
interface Ethernet4/9
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown

interface Ethernet4/10
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 251 mode active
  no shutdown

interface Ethernet4/11
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet4/12
```

```
  mtu 9216
  channel-group 150 mode active
  no shutdown
```

Step 4 Configure all the associated VLANs including VLAN 3150, which is used for the OTV site-VLAN:

```
vlan 1,20-24,2000-2100,2201-2300,3001-3100,3150,3201-3400
```

For RAGG-2, configure accordingly to the diagram shown in Figure 17.

Following are some useful OTV commands to confirm the OTV local connection and adjacency.

```
DC2-1-DC2-OTV1# sh run otv

!Command: show running-config otv
!Time: Thu Mar 12 04:50:29 2015

version 6.2(8)
feature otv

otv site-vlan 3150


interface Overlay1
  otv join-interface port-channel150
  otv control-group 239.1.0.1
  otv data-group 232.1.0.0/16
  otv extend-vlan 20-24, 2000-2100, 2201-2300, 3001-3100, 3201-3400
  no otv suppress-arp-nd
  no shutdown
otv site-identifier 0x2


DC2-1-DC2-OTV1# sh otv

OTV Overlay Information
Site Identifier 0000.0000.0002

Overlay interface Overlay1

 VPN name             : Overlay1
 VPN state            : UP
 Extended vlans       : 20-24 2000-2100 2201-2300 3001-3100 3201-3400 (Total:506)
 Control group        : 239.1.0.1
 Data group range(s)  : 232.1.0.0/16
 Broadcast group      : 239.1.0.1
 Join interface(s)    : Po150 (10.12.210.73)
 Site vlan            : 3150 (up)
 AED-Capable          : Yes
 Capability           : Multicast-Reachable

DC2-1-DC2-OTV1# sh otv site

Dual Adjacency State Description
    Full     - Both site and overlay adjacency up
    Partial  - Either site/overlay adjacency down
    Down     - Both adjacencies are down (Neighbor is down/unreachable)
    (!)      - Site-ID mismatch detected

Local Edge Device Information:
    Hostname DC2-1-DC2-OTV1
    System-ID 8478.ac1d.2fc4
```

```
    Site-Identifier 0000.0000.0002
    Site-VLAN 3150 State is Up


Site Information for Overlay1:


Local device is AED-Capable
Neighbor Edge Devices in Site: 1


Hostname                 System-ID       Adjacency-       Adjacency-     AED-

                                         State            Uptime         Capable


-------------------------------------------------------------------------------------
DC2-1-DC2-OTV2           8478.ac1d.2fc5  Full             1d18h          Yes



DC2-1-DC2-OTV1# sh otv adjacency
Overlay Adjacency database

Overlay-Interface Overlay1  :
Hostname                      System-ID      Dest Addr     Up Time    State
RAGG-1-OTV-1                  0024.f719.2a42 10.11.210.73  1d18h      UP
RAGG-1-OTV-2                  0024.f719.2a43 10.11.210.77  1d18h      UP
DC2-1-DC2-OTV2                8478.ac1d.2fc5 10.12.210.77  1d18h      UP
DC2-1-DC2-OTV1#




RAGG-1-OTV-1# sh otv vlan
OTV Extended VLANs and Edge Device State Information (* - AED)
Legend:
(NA) - Non AED, (VD) - Vlan Disabled, (OD) - Overlay Down
(DH) - Delete Holddown, (HW) - HW: State Down
 (NFC) - Not Forward Capable

VLAN  Auth. Edge Device                    Vlan State            Overlay
----  -----------------------------------  ---------------------  ------
-
  20* RAGG-1-OTV-1                          active                Overlay1
2001  RAGG-1-OTV-2                          inactive(NA)          Overlay1
2002* RAGG-1-OTV-1                          active                Overlay1
2003  RAGG-1-OTV-2                          inactive(NA)          Overlay1
2004* RAGG-1-OTV-1                          active                Overlay1
2201  RAGG-1-OTV-2                          inactive(NA)          Overlay1
2202* RAGG-1-OTV-1                          active                Overlay1
2203  RAGG-1-OTV-2                          inactive(NA)          Overlay1
2204* RAGG-1-OTV-1                          active                Overlay1
3001  RAGG-1-OTV-2                          inactive(NA)          Overlay1
3002* RAGG-1-OTV-1                          active                Overlay1
3003  RAGG-1-OTV-2                          inactive(NA)          Overlay1
3004* RAGG-1-OTV-1                          active                Overlay1

3201  RAGG-1-OTV-2                          inactive(NA)          Overlay1
3202* RAGG-1-OTV-1                          active                Overlay1
3203  RAGG-1-OTV-2                          inactive(NA)          Overlay1
3204* RAGG-1-OTV-1                          active                Overlay1
```

Note   Every other VLANs are active and inactive; this is because OTV evenly distributes the VLANs between local OTV switches.

Cisco ASA Firewall Clustering

Initial configuration of the firewalls was performed via the console command line. After the first ASA was configured, additional firewalls were then added to the cluster. For additional information on cluster configuration options, refer to the following resources:

- Secure Data Center ASA Clustering with FirePOWER Services—
  http://www.cisco.com/c/dam/en/us/solutions/collateral/enterprise/design-zone-secure-data-center-portfolio/sdc_ig.pdf

- ASA 9.3 CLI configuration guide—
  http://www.cisco.com/c/en/us/td/docs/security/asa/asa93/configuration/general/asa-general-cli.html

- ASA Clustering within the VMDC Architecture—
  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/ASA_Cluster/ASA_Cluster.html

- Additional ASA configuration guides—
  http://www.cisco.com/c/en/us/support/security/asa-5500-series-next-generation-firewalls/products-installation-and-configuration-guides-list.html

The cluster in this validation consists of eight ASAs acting as a single unit. All units in the cluster are the same model with the same DRAM. The units used in the cluster are all running 9.3(2) software.

When deploying the ASA Cluster, all of the ASAs must have the exact same configurations for the ASA system to work properly. In addition, they should be deployed in a consistent manner. This applies to using the same type of ports on each unit to connect to the fabric. Use the same ports for the CCL to the switching fabric and the same type of ports used to connect the Data Plane links. When the ASA Cluster is deployed properly, the master unit of the cluster replicates its configuration to the other units in the cluster, and so the cluster must have a consistent deployment across all the units.

Keep in mind that these features are applied to each ASA unit, instead of the cluster as a whole:

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each unit independently. For example, if you configure policing on output, the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster with eight units and with traffic evenly distributed, the conform rate actually becomes eight times the rate for the cluster. QoS was not implemented in this validation.

- Threat detection—Threat detection works on each unit independently; for example, the top statistics are unit-specific. Port scanning detection, for example, does not work because scanning traffic is load-balanced between all units (when using **source-dest-ip-port** balancing), and one unit does not see all traffic.

- Resource management—Resource management in multiple context mode is enforced separately on each unit based on local usage.

- NGIPS FirePOWER modules — There is no configuration sync or state sharing between FirePOWER modules. More information on this is available below in the NGIPS section.

ASA Connectivity to Aggregation Switches

The ASA Cluster Data Plane interfaces were configured as a Spanned EtherChannel using a single port channel for both inside and outside VLAN interfaces. These channels connect to a pair of Nexus 7000s using a vPC. The EtherChannel aggregates the traffic across all the available active interfaces in the channel. A spanned EtherChannel accommodates both routed and transparent firewall modes per use case requirements. The EtherChannel inherently provides load balancing as part of basic operation

using cLACP. Figure 18 shows the connections and port channels implemented.

The complete configuration and implementation can be found in the *Secure Data Center ASA Clustering with FirePOWER Services Implementation Guide*.

Figure 18   *Cluster Connections*
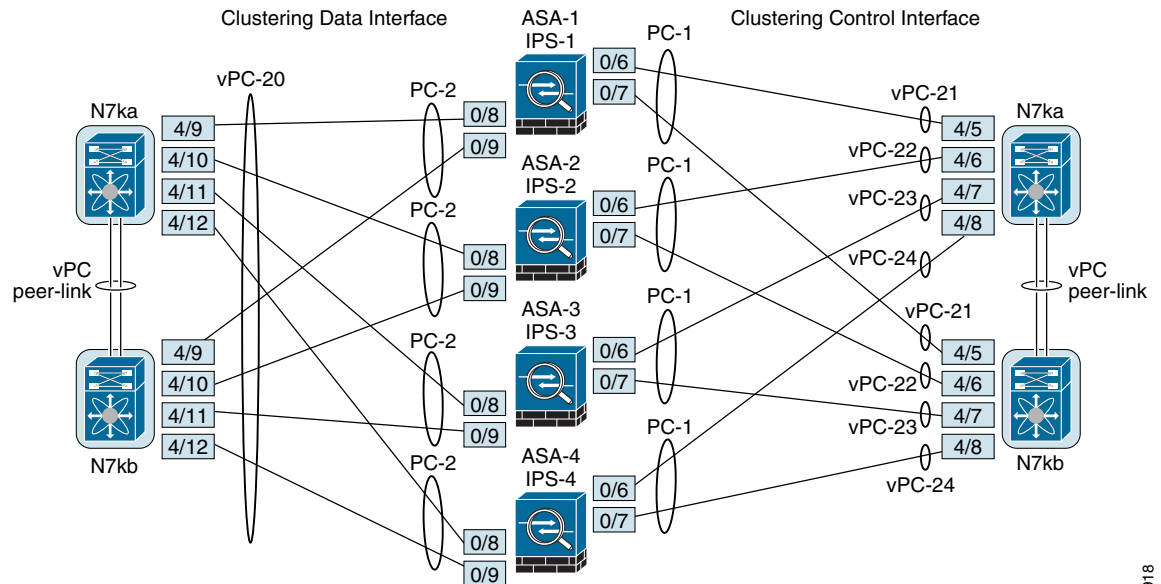
It is important to point out on the Clustering Control Plane that the clustered ASAs have a common port channel configuration because of the sync from the cluster, but the Nexus 7000s have different port channel identifiers configured per ASA because these are local and not spanned across the cluster. EtherChannels configured for the cluster control link are configured as discrete EtherChannels on the switch.

✎

Note   When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the ASA or the switch, or adding an additional switch to form a VSS or vPC) you should disable the health check feature. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the health check feature.

Configuring the Master Firewall

The interface type mode is the first item that must be specified before configuration of the ASAs. You must set the mode separately on each ASA that you want to add to the cluster. If the device is already configured for multiple context mode, configure this setting in the system execution space.

Procedure

Step 1   Configure the cluster interface mode for each unit using the console port:

```
ciscoasa(config)# cluster interface-mode spanned
```

The ASA firewall then clears all improper configurations and reboots.

Step 2   Next, configure the CCL interface. The CCL interface must be enabled before joining the cluster.

```
interface TenGigabitEthernet0/6
 channel-group 1 mode active
 no shutdown
!
interface TenGigabitEthernet0/7
 channel-group 1 mode active
 no shutdown
!
interface Port-channel1
 no shutdown
```

Step 3   For multi-mode, create or change to the Admin context. Be sure to assign the M0/0 interface.

Step 4   In the Admin context, configure the Cluster IP pool and then assign an IP address to the M0/0 interface specifying the cluster pool.

```
ip local pool mgmt-pool 10.11.235.21-10.11.235.28
!
interface Management0/0
management-only
nameif management
security-level 0
ip address 10.11.235.20 255.255.255.0 cluster-pool mgmt-pool
no shutdown
```

Step 5   Now the cluster wizard in ASDM can be started, or use the following configuration statements to create the master node of the cluster.

```
mtu cluster 9000
cluster group SJ-SD2
 key *****
 local-unit ASA-DC-1
 cluster-interface Port-channel1 ip 192.168.20.101 255.255.255.0
 priority 1
 console-replicate
 no health-check
 clacp system-mac auto system-priority 1
 enable
 conn-rebalance frequency 3
```

When completed, additional security contexts can be created and set as routed or transparent.

To configure multiple contexts, see the following URL:
http://www.cisco.com/c/en/us/td/docs/security/asa/asa93/configuration/general/asa-general-cli/ha-contexts.html

To configure routed or transparent mode, see the following URL:
http://www.cisco.com/c/en/us/td/docs/security/asa/asa93/configuration/general/asa-general-cli/intro-fw.html

Note   ASDM has a bug: you cannot add the IP address to the sub-interface (vlan) of the routed context for a spanned port channel. The IP address was able to be configured via the CLI. There were no problems creating a bridge interface on the transparent FW context.

Adding Slave Firewalls

Note   Be sure to upgrade ASA software version to match the cluster before adding to the cluster.

Each unit in the cluster requires a bootstrap configuration to join the cluster. Typically, the first unit configured in the cluster is the master unit. After you enable clustering, after an election period, the cluster elects a master unit. With only one unit in the cluster initially, that unit becomes the master unit. Subsequent units that you add to the cluster become slave units.

Add additional cluster members via the console, Telnet and SSH are not supported. Change the cluster interface mode to match the ASA cluster master, clear the configuration, configure the cluster control link interface, and join the cluster as a slave as follows:

```
changeto system
cluster interface-mode spanned force
clear configure cluster
mtu cluster 9000
interface TenGigabitEthernet0/6
    channel-group 1 mode active
    no shutdown
interface TenGigabitEthernet0/7
    channel-group 1 mode active
    no shutdown
interface Port-channel1
    no shutdown
cluster group SJ-SD2
    local-unit ASA-DC-#
! Where # is unique to identify the unit name
    priority #
!Where # is unique to set the priority
    cluster-interface Port-channel1 ip 192.168.20.104 255.255.255.0
    key *********
    enable as-slave noconfirm
```

You will be able to check the unit has successfully joined the cluster by executing the **show cluster info** commands.

If the unit did not join the cluster successfully, make sure all the clustering link interfaces are not shut and are able to communicate with master units.

The MTU command enables Jumbo-frame reservation, and should be added to the configuration manually because it is not synced via the cluster.

Firewall Contexts

The ASA cluster was partitioned into multiple virtual devices, known as security contexts. Each context acts as an independent device, with its own security policy, interfaces, configuration, and administrators. Multiple contexts are similar to having multiple standalone devices.

All ASA 5585-X units within the cluster share a single configuration. When configuration changes are made on the master unit, the changes are automatically replicated to all slave units in the cluster. A configuration change directly made on slave units is prohibited.

Two Transparent mode contexts were created and one Routed mode context was created. These were labeled as Enclaves 1 through 3 aligning with the Secure Enterprise Enclave (SEA) design guidance. Figure 19 shows the logical segmentation implemented.

**Figure 19** **Logical Topology**
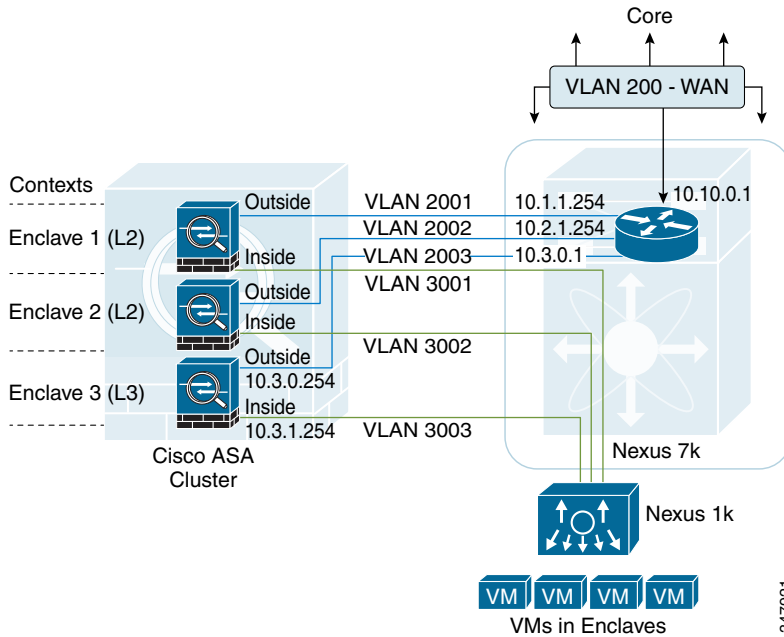
Management Network

All units in the cluster must be connected to a management network that is separate from the CCL. Use the dedicated management interfaces of each ASA as shown in Figure 20.

Each ASA is assigned a unique IP address, and a system IP is assigned to the master unit as its secondary IP address.

For inbound management traffic, an application such as Cisco Security Manager accesses the master ASA by using the system IP address or individual ASAs by their own IP address. For outbound traffic, such as SNMP or syslog, each ASA uses its own IP address to connect to the server. In multi-context mode, the same configuration applies to the admin context and any user contexts that allow remote management.

```
ip local pool mgmt-pool 10.11.235.21-10.11.235.28
!
interface Management0/0
 management-only
 nameif management
 security-level 0
 ip address 10.11.235.20 255.255.255.0 cluster-pool mgmt-pool
```

Figure 20   *Management Interface Connectivity*
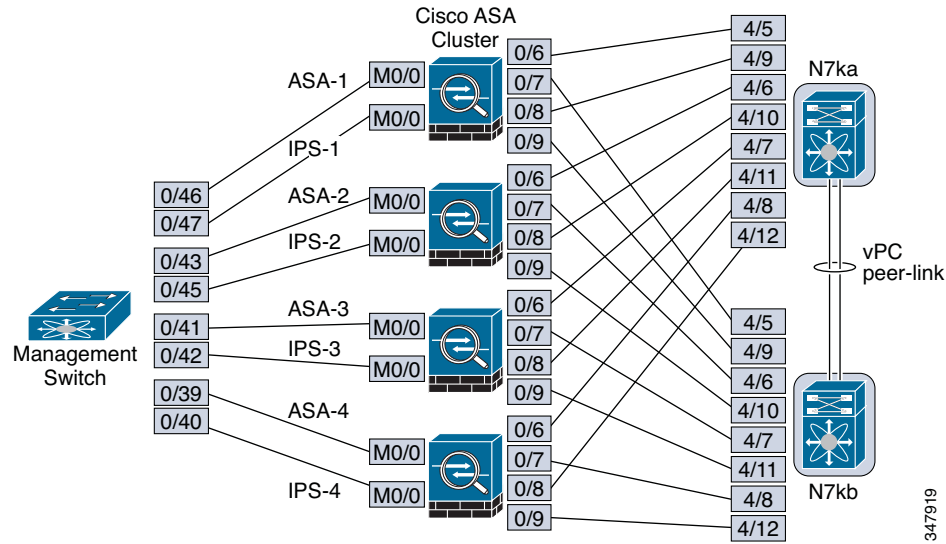
ASA and SXP for TrustSec

Each context maintains its own configurations, databases, credentials, and environment data. The master unit of the cluster contacts ISE via SXP connections from each context, and obtains the secure group-to-IP address mappings data, which is then replicated to all units in the cluster via reliable messaging, thus security group-based policies are enforced on the slave units as well. Security group-based policies are replicated as part of the configuration sync. Both routed and transparent firewall modes are supported. In this validation, the management interface was used for SXP communication with peer devices to keep it out of the normal flow of production traffic.

SXP flows from the Nexus 7000 to each of the contexts in the ASA cluster as shown in Figure 21. The Nexus 7000s act as the "speakers" and send the SGT and IP address mapping to the "listeners", which include each of the ASA cluster contexts, via SXP.

Figure 21 *TrustSec Communication*
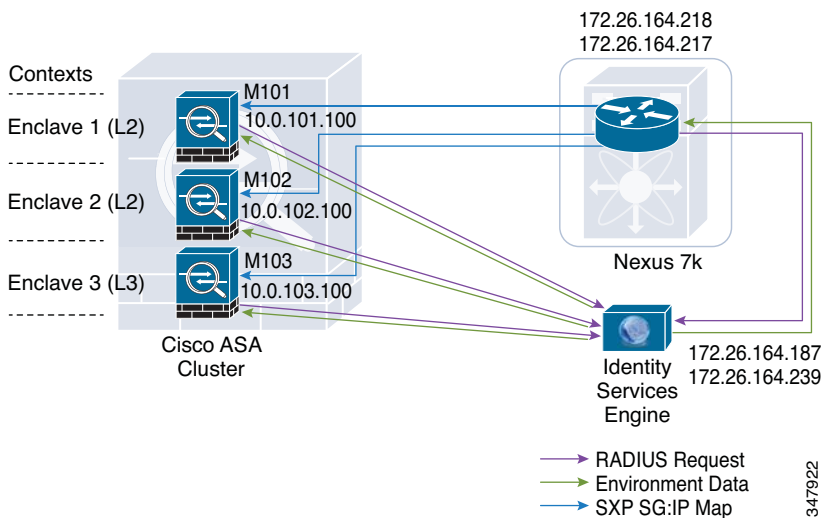
Configuration of ASA Context (Enclave 1)

```
cts server-group ISE-1
cts sxp enable
cts sxp default password *****
cts sxp default source-ip 10.11.101.50
cts sxp connection peer 10.11.255.11 source 10.11.101.50 password default mode local
listener
cts sxp connection peer 10.11.255.12 source 10.11.101.50 password default mode local
listener
```

Configurations of Nexus 7000

```
cts sxp enable
cts sxp default password 7 <removed>
cts sxp connection peer 10.11.101.50 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.101.100 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.102.50 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.102.100 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.103.50 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.103.100 source 10.11.255.11 password default mode
listener vrf default
```

Security Policy

Cisco ASA Software Release 9.0.1 and above is necessary for SGFW functionality. Policy in the firewall has been expanded to include source and destination security groups that are downloaded from the ISE in the environment data after the ASA has established a secure connection by importing a PAC file from the ISE. The ASA issues a RADIUS request for the TrustSec environment data, which includes the secure group table mapping secure group names to secure group numbers. The ASA receives the secure group numbers-to-secure group IP addresses mapping from the Nexus 7000 via SXP. If the PAC file downloaded from the ISE expires on the ASA and the ASA cannot download an updated security group table, the ASA continues to enforce security policies based on the last downloaded security group table until the ASA downloads an updated table.

As part of configuring the ASA to integrate with Cisco TrustSec, you must configure each context so that it can communicate with the ISE servers via RADIUS. The last configuration identifies the AAA server group that is used by Cisco TrustSec for environment data retrieval.

```
aaa-server ISE_Radius_Group protocol radius
aaa-server ISE_Radius_Group (management) host 10.11.230.111
 key *****
 radius-common-pw *****
aaa-server ISE_Radius_Group (management) host 10.11.230.112
 key *****
 radius-common-pw *****
!
cts server-group ISE_Radius_Group
```

When configuring access rules from ASDM and CSM, objects created in the PAC files are available as source and destination criteria.

Figure 22 shows the web interface for configuring the ASA.

Figure 22   Configuring the ASA
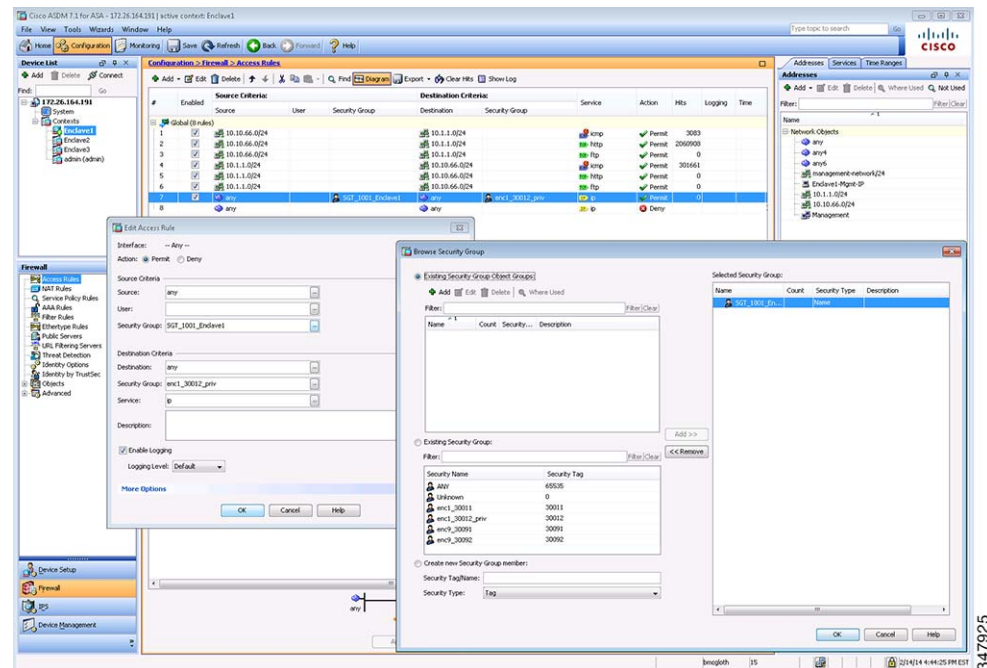
For more information on configuring the ASA to integrate with TrustSec, see the following URL:
http://www.cisco.com/c/en/us/td/docs/security/asa/asa91/configuration/general/asa_91_general_config/aaa_trustsec.html

ASA FirePOWER Mode

The ASA uses policies in each context to identify and divert traffic to the FirePOWER module for inspection. Policies to define traffic can be configured in ASDM, CSM, and the CLI.

The following example identifies HTTP traffic over port 80 to be directed for inspection:

```
class-map Port80Dest
  match port tcp eq www
class-map inspection_default
  match default-inspection-traffic
```

The following example identifies the global policy with HTTP inspection removed and the FirePOWER module to fail closed:

```
policy-map global_policy
 class inspection_default
   inspect dns preset_dns_map
   inspect ftp
   inspect ip-options
   inspect netbios
   inspect rsh
   inspect esmtp
   inspect sqlnet
   inspect sunrpc
   inspect tftp
   inspect xdmcp
 class global-class
   sfr fail-close
 !
service-policy global_policy global
```

Using ASDM, policies are defined under the Service Policy Rules menu option, as shown in Figure 23.

Figure 23   *ASDM Traffic Policy*

Figure 24 shows that the FirePOWER card fails closed.

Figure 24   *FirePOWER Module Operation Mode*

FirePOWER Installation and Configuration

Initial configuration of the FirePOWER appliances and FireSIGHT Management Center (formerly Defense Center) appliance was performed via the console command line where the management address and gateway were assigned following the steps in the *Quick Start Guide*. After each of the appliances were configured and accessible across the network, the remaining configuration was completed using the Web GUI. For additional information on configuration options, see the Sourcefire 3D System User, Installation, and Quick Start Guides for version 5.3 at the following URL: https://support.sourcefire.com/sections/10

Install Defense Center Appliance

The FireSIGHT Management Center appliance was set up as follows.

Procedure

Step 1   At the console, log into the appliance. Use *admin* as the username and *Sourcefire* as the password.

Step 2   At the admin prompt, run the following script:

```
sudo /usr/local/sf/bin/configure-network
```

Step 3   Follow the script's prompts. Configure IPv4 management settings and enter IPv4 addresses, including the netmask, in dotted decimal form.

```
10.11.236.21 255.255.255.0
```

Step 4 Confirm that your settings are correct.

If you entered settings incorrectly, type *n* at the prompt and press **Enter**. You can then enter the correct information. The console may display messages as your settings are implemented.

Step 5 Log out of the appliance.

For all FireSIGHT Management Centers, you must complete the setup process by logging into the FireSIGHT Management Center's web interface and specifying initial configuration options on a setup page. You must change the administrator password, specify network settings if you have not already, and accept the EULA.

The setup process also allows you to register and license devices. Before you can register a device, you must complete the setup process on the device itself, as well as add the FireSIGHT Management Center as a remote manager, or the registration fails.

Complete the Initial Setup on a FireSIGHT Management Center using its Web Interface—ACCESS: Admin

Procedure

Step 1 Direct your browser to *https://10.11.230.142/*, the IP address of the FireSIGHT Management Center's management interface:

The login page appears.

Step 2 Log in using *admin* as the username and *Sourcefire* as the password.

The setup page appears. Complete the following sections:

- Change Password
- Network Settings
- Time Settings
- Recurring Rule Update Imports
- Recurring Geolocation Updates
- Automatic Backups
- License Settings
- Device Registration
- End User License Agreement

For more specific steps, see the following URL:
http://www.cisco.com/c/dam/en/us/td/docs/security/sourcefire/3d-system/53/Sourcefire_3D_System_Installation_Guide_v53.pdf

Add Licenses to FireSIGHT Management Center

You can license a variety of features to create an optimal Sourcefire 3D System deployment for your organization. A FireSIGHT license on the FireSIGHT Management Center is required to perform host, application, and user discovery. Additional model-specific licenses allow your managed devices to perform a variety of functions. Because of architecture and resource limitations, not all licenses can be applied to all managed devices. Cisco recommends that you use the initial setup page to add the

licenses your organization has purchased. If you do not add licenses now, any devices you register during initial setup are added to the FireSIGHT Management Center as unlicensed; you must license each of them individually after the initial setup process is over. If you have not already obtained your licenses, click the link to navigate to *https://keyserver.sourcefire.com/* and follow the on-screen instructions. You need your license key (listed on the initial setup page), as well as the activation key previously emailed to the contact associated with your support contract.

To submit feature licenses for the device into FireSIGHT Management Center, you need:

- A license key from FireSIGHT Management Center that starts with 66:00:00:00:00:00:00 (00 being its MAC address)

- An activation key that was given to the specific device upon your order through the mail

With the above information, you can request a feature license key for the specific device.

Figure 25 shows a feature license successfully installed into the FireSIGHT Management Center.

Figure 25   License Added
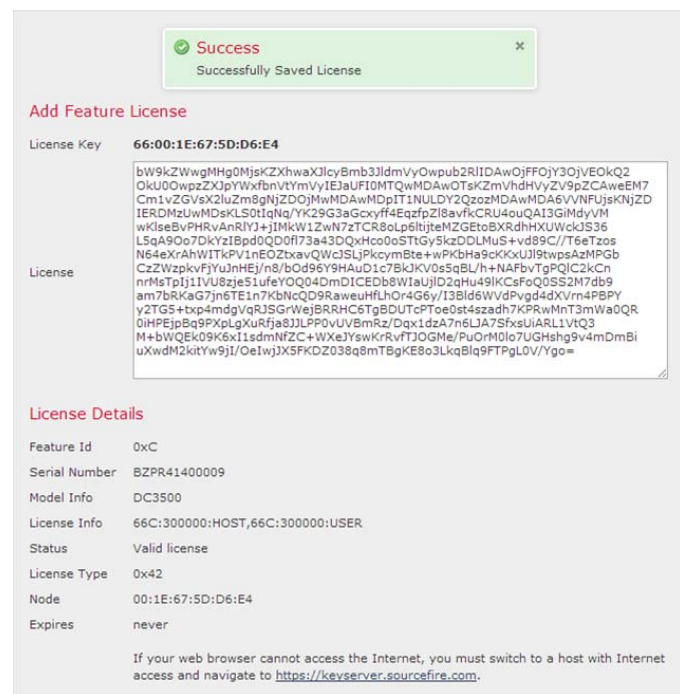
Add FirePOWER Appliances to FireSIGHT Management Center

You can add most pre-registered devices to the FireSIGHT Management Center during the initial setup process (see Figure 26). However, if a device and the FireSIGHT Management Center are separated by a NAT device, you must add it after the setup process completes.

You need the following:

- Host name or IP address of the device

- Registration key (a unique key string that you are assigned during the initial setup for the device)

- Access control policy

Figure 26   *FireSIGHT Management Center—Device Management*

To Add a Device to a FireSIGHT Management Center—ACCESS: Admin/Network Admin

Procedure

Step 1   Configure the device to be managed by the FireSIGHT Management Center.

Step 2   Select **Devices > Device Management**.

The Device Management page appears.

Step 3   From the Add drop-down menu, select **Add Device**.

The Add Device pop-up window appears, as shown in Figure 27.

Figure 27   *Add Device Pop-up Window*

Step 4 In the Host field, type the IP address or the hostname of the device you want to add.

The hostname of the device is the fully qualified domain name or the name that resolves through the local DNS to a valid IP address.

Step 5 In the Registration Key field, type the same registration key that you used when you configured the device to be managed by the FireSIGHT Management Center.

Step 6 Optionally, add the device to a device group by selecting the group from the Group drop-down list.

Step 7 From the Access Control Policy drop-down list, select an initial policy to apply to the device:

The Default Access Control policy blocks all traffic from entering your network.

Step 8 Select licenses to apply to the device.

Step 9 To allow the device to transfer packets to the FireSIGHT Management Center, select the Transfer Packets check box.

This option is enabled by default. If you disable it, you completely prohibit packet transfer to the FireSIGHT Management Center.

Step 10 Click **Register**.

The device is added to the FireSIGHT Management Center. Note that it may take up to two minutes for the FireSIGHT Management Center to verify the device's heartbeat and establish communication.

Security Policies

You can create multiple security policies and apply them to the appliance as a whole using the Default action, or apply policies to individual rules that can be tailored to the specific needs of an Enclave. (See Figure 28.)

Figure 28   *Creating Security Policies*

More information on configuring Intrusion policies can be found in the *Sourcefire 3D System User Guide*, Chapter 19, page 711.

Cisco TrustSec

Cisco Identity Service Engine

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

ISE installation was accomplished using the Cisco Identity Services Engine Installation and Upgrade Guides, available at the following URL:
http://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-guides-list.html

Installation was deployed using a pair of VMs as the laboratory utilizations are very low. Typical enterprise deployments should be on dedicated ISE hardware and scaled to meet enterprise requirements.

Each RADIUS client must be added to the ISE network devices, as shown in Figure 29. Within Cisco ISE, navigate to **Administration > Network Resources > Network Devices**.

Figure 29   *ISE Network Devices*
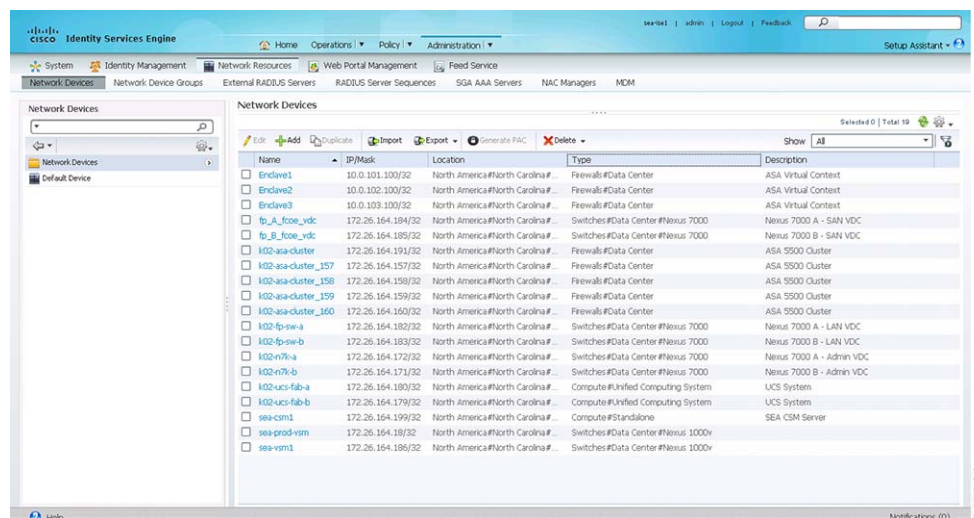
Add devices as follows.

Procedure

Step 1   Click **Add**.

Step 2   Enter the device name and an IP address.

Step 3   Under Network Device Group, select the Location and Device Type.

Step 4   Scroll down and check the box for Authentication Settings. Configure the shared secret.

Step 5   Scroll down and check the box for Security Group Access (SGA). Check the box to use the Device ID for SGA Identification, and configure the password to be used by the device during registration.

Step 6   Scroll down and check the box for Device Configuration Deployment. Fill in the exec mode username and password.

Step 7   This step is necessary for deploying the IP/hostname to SGT mapping.

Step 8   Click **Submit**.

Figure 30 shows the configuration of the ASA Enclave 1, and Figure 31 shows the configuration of the Nexus 7000.

Figure 30   ISE ASA Context Device

Figure 31   ISE Nexus 7000 Device
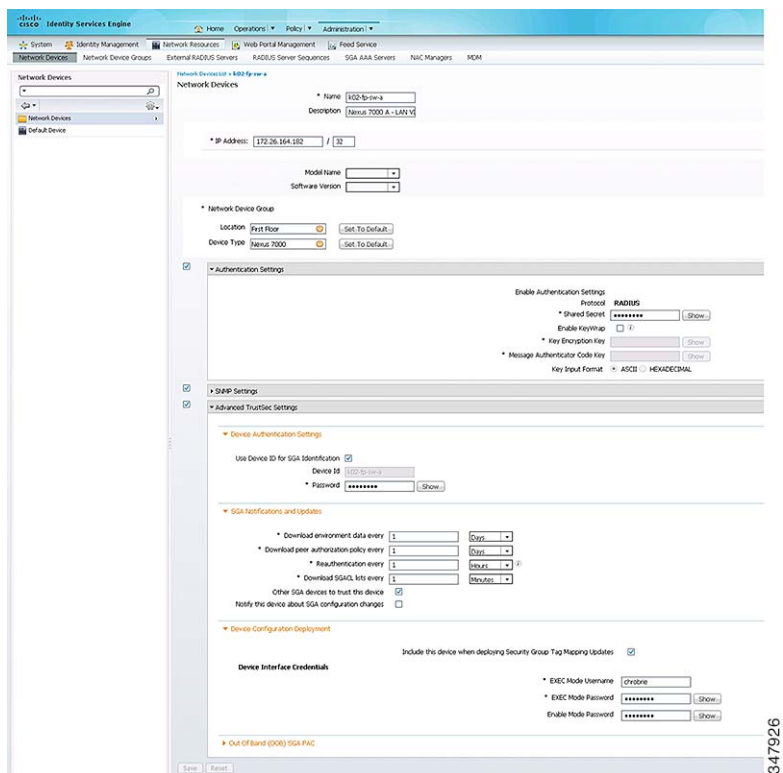
Each of the ASA contexts in the cluster is configured to communicate with the ISE server, as shown previously in the ASA section.

The Nexus 1000v and 7000 are both configured to join the Cisco TrustSec domain and receive the SGT PAC files. First configure the ISE RADIUS group as follows:

```
radius-server host 172.26.164.187 key 7 <removed> authentication accounting
radius-server host 172.26.164.239 key 7 <removed> authentication accounting
aaa group server radius ISE-Radius-Grp
    server 172.26.164.187
    server 172.26.164.239
    use-vrf management
    source-interface mgmt0
```

After the ISE RADIUS group is configured, configure the Authentication and Authorization actions:

```
aaa authentication dot1x default group ISE-Radius-Grp
aaa accounting dot1x default group ISE-Radius-Grp
aaa authorization cts default group ISE-Radius-Grp
```

Lastly configure the switch to join the Cisco TrustSec domain. This command invokes device registration with Cisco ISE and forces a PAC download. Make sure the device-id matches the name entry in Cisco ISE.

```
cts device-id k02-fp-sw-a password 7 <removed>
```

For more information on configuring TrustSec on Nexus devices, see the NX-OS Security Configuration Guide at the following URLs:

- http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-installation-and-configuration-guides-list.html

- http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/6_x/nx-os/security/configuration/guide/b_Cisco_Nexus_7000_NX-OS_Security_Configuration_Guide__Release_6-x/b_Cisco_Nexus_7000_NX-OS_Security_Configuration_Guide__Release_6-x_chapter_01101.html

Validation Testing

Summary of Tests Performed

These tests are designed to validate the integration of and general functionality of the Secure Data Center design. The common structure of the architecture is based on Cisco's integrated reference architectures.

Table 8 outlines the various tests conducted to validate the deployment.

Table 8   Test Scenarios

Test	Methodology
Physical Cisco ASA Failure and Recovery - Clustered Mode Fail Slave	In this failure scenario, Cisco manually removed and recovered power from one of the slave ASA nodes in the cluster.
Physical Cisco ASA Failure and Recovery - Clustered Mode Fail Master	In this failure scenario, Cisco manually removed and recovered power from the master ASA node in the cluster.
FirePOWER power recovery and FireSIGHT Management Center connectivity	In this failure scenario, Cisco manually removed the FirePOWER service module from the ASA. Cisco verified its recovery using FireSIGHT Management Center to regain the connectivity.
ASA Cluster Data Link Failures – Master and Slave	Fail and recover the following links: • Fail a data link to master • Fail both data links to master • Fail a data link to slave • Fail both data links to slave
ASA Cluster Control Link Failures – Master and Slave	Fail and recover the following links: • Fail a cluster link to master • Fail both cluster links to master • Fail a cluster link to slave • Fail both cluster links to slave
Create a policy to forward traffic to FirePOWER service module	In this test, Cisco created a policy in the ASA to set specific protocols to be inspected by the FirePOWER service module.
Sourcefire TCP flow checker	Test the outcome of enabling TCP checks per virtual interface and/or per intrusion policy.
Management Traffic Flows	Ensure centralized management access via private VLAN and firewall access control rules

Table 8   Test Scenarios (continued)

Asymmetric Traffic Flow Validation	Asymmetric traffic flows are introduced to the test bed. Ensure the ASA implementation properly manages these flows.
Validate Integrity of Sourcefire serviced flows	Validate integrity of flow and ability to enforce policy.
Validate OTV relation between two sites	• Disable OTV connection between two sites to see how ASA clustering functions.

Summary of Results

Table 9 lists the summary of results.

Table 9   Summary of Results

Test Description	Components	Result
Physical ASA Cluster failure and recovery (Fail Slave) (Fail Master)	ASA5585 ASDM and Spirent	No traffic interruption and notification syslog output with acceptable packet loss. When conducted over the multi-site environment, the role of master and slave accordingly changed by the cluster priority.
ASA link failure on data/clustering link	ASA5585 ASDM and Spirent	No traffic interruption and notification syslog output with acceptable packet loss. When conducted over the multi-site environment, the role of master and slave accordingly changed by the cluster priority.
FirePOWER failure and recovery	FirePOWER service module, FireSIGHT Management Center, Spirent and iPerformance	ASA clustering functions identify the link failure to disable the clustered device. Upon the recovery, units rejoined the cluster membership. When conducted over the multi-site environment, the role of master and slave accordingly changed by the cluster priority.
ASA management link failure	ASA5585 and ASDM	ASA unit with failed management interface leaves the cluster. When conducted over the multi-site environment, the role of master and slave accordingly changed by the cluster priority.
OTV link failure	Core Nexus 7000 and OTV switch	Upon link failure between two data center sites, ASA clustering loses the communication between the sites and forms an independent cluster group. All the sessions that were redirected to the other site lose the connection and the locally kept connection is sustained. Upon the re-establishment of the OTV link, cluster members were formed over two sites.

Conclusion

The Secure Data Center the Enterprise: The Multi Sites Data Center deployment of ASA Clustering with FirePOWER service is a Cisco Validated Design that enables customers to confidently integrate Cisco's security portfolio to respond to the increasing sophisticated attacks being targeted at the data center. This solution is made even stronger when customers also leverage the Secure Enclaves Architecture for securing the workloads, and leverage the Cyber Threat Defense for Data Center solution for enabling behavioral analysis, which provides zero day mitigation protections in the data center.

The test result showed that distributing ASA cluster membership among multiple sites extends its security and functionality along with FirePOWER services and ease of deployment.

References

- Cisco ASA Series CLI Configuration Guides—
  http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html

- Cisco ASA FirePOWER Module Quick Start Guide—
  http://www.cisco.com/c/en/us/td/docs/security/asa/quick_start/sfr/firepower-qsg.html

- Access Control Using Security Group Firewall (Aaron Woland, Cisco.com)—
  http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns744/docs/access_control_using_security.pdf

- Cisco TrustSec How-To Guide: Server-to-Server Segmentation Using SGA (Aaron Woland, Cisco.com)—
  http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns744/docs/howto_75_server_segmentation_sga.pdf

- Data Center Security Design Guide (Mike Storm)—
  http://www.cisco.com/en/US/netsol/ns750/networking_solutions_sub_program_home.html

- Cisco Adaptive Security Appliance Cluster Deployment Guide (Mason Harris, David Anderson, Mike Storm)—
  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/ASA_Cluster/ASA_Cluster.html

- OTV Solution guide and White Paper—
  http://www.cisco.com/c/en/us/solutions/data-center-virtualization/overlay-transport-virtualization-otv/index.html#~Overview

Appendix—Device Configurations

ASA Cluster Configurations

System Context

```
ASA Version 9.3(2) <system>
!
hostname ASA-MS-DC-Cluster1
enable password 8Ry2YjIyt7RRXU24 encrypted
no mac-address auto
!
interface GigabitEthernet0/0
 shutdown
!
interface GigabitEthernet0/1
 shutdown
!
interface GigabitEthernet0/2
 shutdown
!
interface GigabitEthernet0/3
 shutdown
!
interface GigabitEthernet0/4
 shutdown
!
interface GigabitEthernet0/5
 shutdown
!
interface Management0/0
!
interface Management0/0.2201
 description ** Enclave 1 Mgmt **
 vlan 2201
!
interface Management0/0.2202
 description ** Enclave 2 Mgmt **
 vlan 2202
!
interface Management0/0.2203
 description ** Enclave 3 Mgmt **
 vlan 2203
!
interface Management0/1
 shutdown
!
interface TenGigabitEthernet0/6
 description To RAGG-1 for Clst Control
 channel-group 1 mode active
!
interface TenGigabitEthernet0/7
 description To RAGG-2 for Clst Control
 channel-group 1 mode active
!
interface TenGigabitEthernet0/8
 description To RAGG-1 for Clst Data
 channel-group 2 mode active vss-id 1
!
```

```
interface TenGigabitEthernet0/9
 description To RAGG-2 for Clst Data
 channel-group 2 mode active vss-id 2
!
interface GigabitEthernet1/0
 shutdown
!
interface GigabitEthernet1/1
 shutdown
!
interface GigabitEthernet1/2
 shutdown
!
interface GigabitEthernet1/3
 shutdown
!
interface GigabitEthernet1/4
 shutdown
!
interface GigabitEthernet1/5
 shutdown
!
interface TenGigabitEthernet1/6
 shutdown
!
interface TenGigabitEthernet1/7
 shutdown
!
interface TenGigabitEthernet1/8
 shutdown
!
interface TenGigabitEthernet1/9
 shutdown
!
interface Port-channel1
 description Clustering Interface
 lacp max-bundle 8
!
interface Port-channel2
 description Cluster Spanned Data Link to RAGG vPC-20
 lacp max-bundle 8
 port-channel span-cluster vss-load-balance
!
interface Port-channel2.2001
 description Enclave1 outside
 vlan 2001
!
interface Port-channel2.2002
 description Enclave2 outside
 vlan 2002
!
interface Port-channel2.2003
 description Enclave3 outside
 vlan 2003
!
interface Port-channel2.2004
 shutdown
 no vlan
!
interface Port-channel2.3001
 description Enclave1 inside
 vlan 3001
!
interface Port-channel2.3002
```

```
 description Enclave2 inside
 vlan 3002
!
interface Port-channel2.3003
 description Enclave3 inside
 vlan 3003
!
interface Port-channel2.3004
 shutdown
 no vlan
!
class default
  limit-resource Mac-addresses 65535
  limit-resource All 0
  limit-resource ASDM 5
  limit-resource SSH 5
  limit-resource Telnet 5
!

boot system disk0:/asa932-smp-k8.bin
ftp mode passive
clock timezone PST -8
clock summer-time PDT recurring
cluster group MS-SJ1
 local-unit ASA-DC-5
 cluster-interface Port-channel1 ip 192.168.20.101 255.255.255.0
 priority 1
 health-check holdtime 30
 clacp static-port-priority
 clacp system-mac auto system-priority 1
 enable
pager lines 24
mtu cluster 1500
no failover
asdm image disk0:/asdm-732.bin
no asdm history enable
arp timeout 14400
no arp permit-nonconnected
ssh stricthostkeycheck
console timeout 0
!
tls-proxy maximum-session 1000
!

admin-context admin
context admin
  allocate-interface Management0/0
  config-url disk0:/admin.cfg
!

context Enclave1
  description North Enclave 1
  allocate-interface Management0/0.2201 Mgmt2201
  allocate-interface Port-channel2.2001 outside
  allocate-interface Port-channel2.3001 inside
  config-url disk0:/Enclave1.cfg
!

context Enclave2
  description North Enclave 2
  allocate-interface Management0/0.2202 Mgmt2202
  allocate-interface Port-channel2.2002 outside
  allocate-interface Port-channel2.3002 inside
  config-url disk0:/Enclave2.cfg
```

```
!

prompt hostname context
no call-home reporting anonymous
Cryptochecksum:dac426eebd9d37320d9ba9165678ab4f
: end
ASA-MS-DC-Cluster1(config)#
```

Admin Context

```
ASA Version 9.3(2) <context>
!
hostname admin
enable password 8Ry2YjIyt7RRXU24 encrypted
names
ip local pool mgmt-pool 10.11.235.25-10.11.235.32
!
interface Management0/0
 management-only
 nameif management
 security-level 0
 ip address 10.11.235.20 255.255.255.0 cluster-pool mgmt-pool
!
pager lines 24
logging enable
logging asdm informational
no logging message 106015
no logging message 313001
no logging message 313008
no logging message 106023
no logging message 710003
no logging message 106100
no logging message 302015
no logging message 302014
no logging message 302013
no logging message 302018
no logging message 302017
no logging message 302016
no logging message 302021
no logging message 302020
flow-export destination management 10.11.230.154 2055
flow-export template timeout-rate 2
mtu management 1500
icmp unreachable rate-limit 1 burst-size 1
no asdm history enable
arp timeout 14400
route management 0.0.0.0 0.0.0.0 10.11.235.1 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
user-identity default-domain LOCAL
no snmp-server location
no snmp-server contact
crypto ipsec security-association pmtu-aging infinite
telnet timeout 5
no ssh stricthostkeycheck
```

```
ssh timeout 5
ssh key-exchange group dh-group1-sha1
no threat-detection statistics tcp-intercept
!
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 rsh
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect xdmcp
  inspect netbios
  inspect tftp
  inspect ip-options
 class class-default
  flow-export event-type all destination 10.11.230.154
!
service-policy global_policy global
Cryptochecksum:179ad8cb56829bacd32b15f0d7c63678
: end
```

Enclave1 Context

```
ASA Version 9.3(2) <context>
!
firewall transparent
hostname Enclave1
domain-name cisco-x.com
enable password 8Ry2YjIyt7RRXU24 encrypted
names
ip local pool Enclave1-pool 10.11.101.1-10.11.101.32 mask 255.255.255.192
!
interface BVI1
 description Enclave1
 ip address 10.11.1.251 255.255.255.0
!
interface Mgmt2201
 management-only
 shutdown
 nameif management
 security-level 0
 ip address 10.11.101.50 255.255.255.192 cluster-pool Enclave1-pool
!
interface outside
 shutdown
 nameif outside
 bridge-group 1
 security-level 0
!
interface inside
 shutdown
 nameif inside
```

```
 bridge-group 1
 security-level 100
!
dns server-group DefaultDNS
 domain-name cisco-x.com
access-list inside_access_in extended permit tcp security-group name ServersV3001 any
security-group name ClientV201 any eq www inactive
access-list inside_access_in extended permit tcp security-group name ServersV3001 any
security-group name ClientV201 any eq https inactive
access-list inside_access_in extended permit ip any any
access-list outside_access_in extended permit tcp security-group name ClientV201 any
security-group name ServersV3001 any eq www inactive
access-list outside_access_in extended permit tcp security-group name ClientV201 any
security-group name ServersV3001 any eq https inactive
access-list outside_access_in extended permit ip any any
access-list outside_access_in extended permit icmp any any
pager lines 24
logging enable
logging asdm informational
mtu management 1500
mtu outside 1500
mtu inside 1500
icmp unreachable rate-limit 1 burst-size 1
no asdm history enable
arp timeout 14400
access-group outside_access_in in interface outside
access-group inside_access_in in interface inside
route management 0.0.0.0 0.0.0.0 10.11.101.62 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-1 protocol radius
aaa-server ISE-1 (management) host 10.11.230.111
 key *****
 radius-common-pw *****
cts server-group ISE-1
cts sxp enable
cts sxp default password *****
cts sxp default source-ip 10.11.101.50
cts sxp connection peer 10.11.255.11 source 10.11.101.50 password default mode local
listener
cts sxp connection peer 10.11.255.12 source 10.11.101.50 password default mode local
listener
user-identity default-domain LOCAL
http 10.11.0.0 255.255.0.0 management
no snmp-server location
no snmp-server contact
crypto ipsec security-association pmtu-aging infinite
telnet timeout 5
ssh stricthostkeycheck
ssh 10.11.0.0 255.255.0.0 management
ssh timeout 5
ssh key-exchange group dh-group1-sha1
no threat-detection statistics tcp-intercept
!
class-map Port80Dest
 match port tcp eq www
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 rsh
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect xdmcp
  inspect netbios
  inspect tftp
  inspect ip-options
 class Port80Dest
  sfr fail-close
!
service-policy global_policy global
Cryptochecksum:982e2038ac7a2028c80d2d716fd6d8b4
: end
```

Enclave 2 Context

```
ASA Version 9.3(2) <context>
!
firewall transparent
hostname Enclave2
domain-name cisco-x.com
enable password 8Ry2YjIyt7RRXU24 encrypted
names
ip local pool NEnclave2-pool 10.11.102.1-10.11.102.32 mask 255.255.255.192
!
interface BVI1
 description NEnclave2
 ip address 10.11.2.251 255.255.255.0
!
interface Mgmt2202
 management-only
 shutdown
 nameif management
 security-level 0
 ip address 10.11.102.50 255.255.255.192 cluster-pool NEnclave2-pool
!
interface outside
 shutdown
 nameif outside
 bridge-group 1
 security-level 0
!
interface inside
 shutdown
 nameif inside
 bridge-group 1
 security-level 100
!
dns server-group DefaultDNS
 domain-name cisco-x.com
access-list inside_access_in extended permit tcp any any
```

```
access-list inside_access_in extended permit icmp any any
access-list inside_access_in extended permit udp any any
access-list outside_access_in extended permit icmp any any
access-list outside_access_in extended permit tcp any any
access-list outside_access_in extended permit udp any any
pager lines 24
logging asdm informational
mtu management 1500
mtu outside 1500
mtu inside 1500
icmp unreachable rate-limit 1 burst-size 1
no asdm history enable
arp timeout 14400
access-group outside_access_in in interface outside
route management 0.0.0.0 0.0.0.0 10.11.102.62 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-1 protocol radius
aaa-server ISE-1 (management) host 10.11.230.111
 key *****
 radius-common-pw *****
cts server-group ISE-1
cts sxp enable
cts sxp default password *****
cts sxp default source-ip 10.11.102.50
cts sxp connection peer 10.11.255.11 source 10.11.102.50 password default mode local
listener
cts sxp connection peer 10.11.255.12 source 10.11.102.50 password default mode local
listener
user-identity default-domain LOCAL
http 10.11.0.0 255.255.0.0 management
no snmp-server location
no snmp-server contact
crypto ipsec security-association pmtu-aging infinite
telnet timeout 5
ssh stricthostkeycheck
ssh 10.11.0.0 255.255.0.0 management
ssh timeout 5
ssh key-exchange group dh-group1-sha1
no threat-detection statistics tcp-intercept
!
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 rsh
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect xdmcp
```

```
  inspect netbios
  inspect tftp
  inspect ip-options
!
service-policy global_policy global
Cryptochecksum:b0ad2052ad735fc67a158d3f4ec7ea1a
: end
```

Site 1

Nexus 7000 Core Configuration

```
version 6.2(8)
power redundancy-mode ps-redundant

switchname RCORE-1
no system admin-vdc
vdc RCORE-1 id 1
  limit-resource module-type f2e
  cpu-share 5
  allocate interface Ethernet3/1-48
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 2
  limit-resource monitor-session-erspan-dst minimum 0 maximum 23
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 768
  limit-resource u4route-mem minimum 96 maximum 96
  limit-resource u6route-mem minimum 24 maximum 24
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
  limit-resource monitor-session-inband-src minimum 0 maximum 1
  limit-resource anycast_bundleid minimum 0 maximum 16
  limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
  limit-resource monitor-session-extended minimum 0 maximum 12

cfs eth distribute
feature ospf
feature ospfv3
feature pim
feature interface-vlan
feature netflow
feature hsrp
feature lacp
feature vpc
feature lldp

username admin password 5 $1$BTPb8ZVP$C4bTjTDXJSTxOz6LNPaRx1  role network-admin
username chambers password 5 $1$GB.lb2ps$lUo/OuRGJMuY5i7/84TQZ/  role network-ad
min
username bmcgloth password 5 $1$0HujTJuK$477la6o8qxQSMyE4HyEJ20  role network-ad
min
username cstought password 5 $1$NS2iVdxz$pvyKC1aq/QBBGVoy6zGey1  role network-ad
min
no password strength-check
ip domain-lookup
system default switchport
copp profile strict
snmp-server user admin network-admin auth md5 0x4dfc136c6c373a7bbc09598492a38c0c
 priv 0x4dfc136c6c373a7bbc09598492a38c0c localizedkey
snmp-server user cstought network-admin auth md5 0xd39a78f6b3a3ab9eec078678f0a7e
```

```
613 priv 0xd39a78f6b3a3ab9eec078678f0a7e613 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 172.26.129.252 prefer
ntp server 172.28.189.1
ntp source-interface  loopback0

ip route 0.0.0.0/0 10.11.211.30 name Default
ip route 10.10.96.0/20 10.11.211.40
ip pim rp-address 10.11.255.1 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
vlan 1,211,540
vlan 211
  name WAN
vlan 540
  name linktosaccess

vrf context management
  ip route 0.0.0.0/0 10.11.236.1

interface Vlan1
  no shutdown
  no ip redirects
  ip address 10.11.210.29/30
  no ipv6 redirects
  ip router ospf 5 area 0.0.0.0

interface Vlan211
  no shutdown
  no ip redirects
  ip address 10.11.211.11/24
  no ipv6 redirects
  ip router ospf 5 area 0.0.0.0
  hsrp 211
    preempt
    ip 10.11.211.10

interface Vlan540
  description link to SACCESS-1
  no shutdown
  ip address 10.11.210.41/30
  ip router ospf 5 area 0.0.0.0

interface port-channel1
  switchport mode trunk
  spanning-tree port type network

interface Ethernet3/1
  description RAGG-1 Port T4/1
  no switchport
  ip address 10.11.210.13/30
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet3/2
  description RAGG-2 Port T4/1
  no switchport
  ip address 10.11.210.17/30
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
```

```
  no shutdown

interface Ethernet3/3
  switchport mode trunk
  spanning-tree port type network
  channel-group 1

interface Ethernet3/4
  switchport mode trunk
  spanning-tree port type network
  channel-group 1

interface Ethernet3/5
  switchport access vlan 211
  spanning-tree port type edge

interface Ethernet3/6
  switchport access vlan 211
  spanning-tree port type edge

interface Ethernet3/7
  switchport access vlan 211
  spanning-tree port type edge

interface Ethernet3/8
  description Private Metro WAN to DC2-RCORE-1
  no switchport
  ip address 192.168.11.13/30
  ip ospf cost 200
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet3/9

interface Ethernet3/10

interface Ethernet3/11

interface Ethernet3/12

interface Ethernet3/13
  switchport monitor

interface Ethernet3/14
  switchport monitor

interface Ethernet3/15

interface Ethernet3/16

interface Ethernet3/17
  switchport access vlan 211

interface Ethernet3/18

interface Ethernet3/19

interface Ethernet3/20

interface Ethernet3/21

interface Ethernet3/22
```

```
interface Ethernet3/23

interface Ethernet3/24

interface Ethernet3/25

interface Ethernet3/26

interface Ethernet3/27

interface Ethernet3/28
  switchport access vlan 211

interface Ethernet3/29

interface Ethernet3/30

interface Ethernet3/31

interface Ethernet3/32
  switchport access vlan 211

interface Ethernet3/33

interface Ethernet3/34

interface Ethernet3/35

interface Ethernet3/36

interface Ethernet3/37

interface Ethernet3/38

interface Ethernet3/39

interface Ethernet3/40

interface Ethernet3/41

interface Ethernet3/42

interface Ethernet3/43

interface Ethernet3/44

interface Ethernet3/45
  description SACI-Leaf-1 Port 1/3
  no switchport
  mtu 9000
  ip router ospf 5 area 0.0.0.100
  no shutdown

interface Ethernet3/45.10
  mtu 9000
  encapsulation dot1q 10
  ip address 10.11.210.53/30
  ip router ospf 5 area 0.0.0.100
  no shutdown

interface Ethernet3/46
  description SACI-Leaf-2 Port 1/3
  no switchport
  mtu 9000
```

```
      ip address 10.11.210.57/30
      ip router ospf 5 area 0.0.0.100
      no shutdown

    interface Ethernet3/47

    interface Ethernet3/48
      description SACCESS-1 Routed
      switchport mode trunk
      switchport trunk allowed vlan 540

    interface mgmt0
      vrf member management
      ip address 10.11.236.11/24

    interface loopback0
      ip address 10.11.255.1/32
      ip router ospf 5 area 0.0.0.0
    clock timezone PST -8 0
    clock summer-time PDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60
    cli alias name wr copy run start
    line console
    line vty
    boot kickstart bootflash://sup-1/n7000-s2-kickstart.6.2.8.bin sup-1
    boot system bootflash://sup-1/n7000-s2-dk9.6.2.8.bin sup-1
    router ospf 5
      router-id 10.11.255.1
      area 0.0.0.100 nssa default-information-originate
      default-information originate always
    no system default switchport shutdown
    monitor session 16
      source interface Ethernet3/1 both
      source interface Ethernet3/2 both
      destination interface Ethernet3/13
      no shut
    no system auto-upgrade epld
```

Nexus 7000 Aggregation Switch Configuration

```
    version 6.2(8)
    poweroff module 1
    poweroff module 2
    power redundancy-mode ps-redundant

    hostname RAGG-1
    no system admin-vdc
    vdc RAGG-1 id 1
      limit-resource module-type f2 f2e
      cpu-share 5
      allocate interface Ethernet4/1-48
      limit-resource vlan minimum 16 maximum 4094
      limit-resource monitor-session minimum 0 maximum 2
      limit-resource monitor-session-erspan-dst minimum 0 maximum 23
      limit-resource vrf minimum 2 maximum 4096
      limit-resource port-channel minimum 0 maximum 768
      limit-resource u4route-mem minimum 96 maximum 96
      limit-resource u6route-mem minimum 24 maximum 24
      limit-resource m4route-mem minimum 58 maximum 58
      limit-resource m6route-mem minimum 8 maximum 8
      limit-resource monitor-session-inband-src minimum 0 maximum 1
      limit-resource anycast_bundleid minimum 0 maximum 16
```

```
     limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
     limit-resource monitor-session-extended minimum 0 maximum 12
vdc OTV-1 id 2
     limit-resource module-type m1 m1xl m2xl f2e
     cpu-share 5
     allocate interface Ethernet3/7-8,Ethernet3/11-12
     boot-order 1
     limit-resource vlan minimum 16 maximum 4094
     limit-resource monitor-session minimum 0 maximum 2
     limit-resource monitor-session-erspan-dst minimum 0 maximum 23
     limit-resource vrf minimum 2 maximum 4096
     limit-resource port-channel minimum 0 maximum 768
     limit-resource u4route-mem minimum 8 maximum 8
     limit-resource u6route-mem minimum 4 maximum 4
     limit-resource m4route-mem minimum 8 maximum 8
     limit-resource m6route-mem minimum 5 maximum 5
     limit-resource monitor-session-inband-src minimum 0 maximum 1
     limit-resource anycast_bundleid minimum 0 maximum 16
     limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
     limit-resource monitor-session-extended minimum 0 maximum 12
vdc OTV-2 id 3
     limit-resource module-type m1 m1xl m2xl f2e
     cpu-share 5
     allocate interface Ethernet3/9-10,Ethernet3/13-14
     boot-order 1
     limit-resource vlan minimum 16 maximum 4094
     limit-resource monitor-session minimum 0 maximum 2
     limit-resource monitor-session-erspan-dst minimum 0 maximum 23
     limit-resource vrf minimum 2 maximum 4096
     limit-resource port-channel minimum 0 maximum 768
     limit-resource u4route-mem minimum 8 maximum 8
     limit-resource u6route-mem minimum 4 maximum 4
     limit-resource m4route-mem minimum 8 maximum 8
     limit-resource m6route-mem minimum 5 maximum 5
     limit-resource monitor-session-inband-src minimum 0 maximum 1
     limit-resource anycast_bundleid minimum 0 maximum 16
     limit-resource monitor-session-mx-exception-src minimum 0 maximum 1
     limit-resource monitor-session-extended minimum 0 maximum 12

cfs eth distribute
feature ospf
feature pim
feature private-vlan
feature udld
feature interface-vlan
feature dot1x
feature hsrp
feature lacp
feature cts
cts device-id RAGG-1 password 7 1Oihmrdyq!
cts sxp default password 7 1Oihmrdyq!
cts sxp connection peer 10.11.101.50 source 10.11.255.11 password default mode l
istener vrf default
cts sxp connection peer 10.11.101.100 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.102.50 source 10.11.255.11 password default mode l
istener vrf default
cts sxp connection peer 10.11.102.100 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.103.50 source 10.11.255.11 password default mode l
istener vrf default
cts sxp connection peer 10.11.103.100 source 10.11.255.11 password default mode
listener vrf default
cts sxp connection peer 10.11.230.241 source 10.11.255.11 password default mode
```

```
speaker vrf default
cts sxp connection peer 10.11.236.33 source 10.11.255.11 password default mode l
istener vrf default
cts sxp connection peer 10.11.236.34 source 10.11.255.11 password default mode l
istener vrf default
cts sxp connection peer 10.11.255.1 source 10.11.255.11 password default mode li
stener vrf default
cts sxp connection peer 10.11.255.2 source 10.11.255.11 password default mode li
stener vrf default
feature vpc

logging level private-vlan 3
username admin password 5 $1$Oi.sBfur$yc1wX3aTeA3UzZdf3GsVu1  role network-admin
username mkaneko password 5 $1$ktErVJU/$s2/FWJX1hL6OjgReHnoK10  role network-ope
rator
username mkaneko role network-admin
username chambers password 5 $1$ZrHC9lwM$g3xggPHRYGndylVfYWNQ3/  role network-ad
min
username bmcgloth password 5 $1$gXbx3OcJ$MdgsXlVniRpl.uY3Rp/w90  role network-ad
min
username ISEServer password 5 $1$rFesQx9j$8aARna9IDBjddo83FPAc61  role network-a
dmin
no password strength-check
ip domain-lookup
radius-server host 10.11.230.111 key 7 "1Oihmrdyq!" pac authentication accountin
g
aaa group server radius aaa-private-sg
aaa group server radius CTS-RADIUS
    server 10.11.230.111
copp profile strict
snmp-server user admin network-admin auth md5 0xaa4c9c11831d1baa960fbb1b013158b9
 priv 0xaa4c9c11831d1baa960fbb1b013158b9 localizedkey
snmp-server user mkaneko network-operator auth md5 0x48963d2d44706040ee514dbdcc0
f5e83 priv 0x48963d2d44706040ee514dbdcc0f5e83 localizedkey
snmp-server user bmcgloth network-admin auth md5 0x88ab82d413b64a8dd22659b60843a
8e9 priv 0x88ab82d413b64a8dd22659b60843a8e9 localizedkey
snmp-server user chambers network-admin auth md5 0xaa4c9c11831d1baa960fbb1b01315
8b9 priv 0xaa4c9c11831d1baa960fbb1b013158b9 localizedkey
snmp-server user ISEServer network-admin auth md5 0xaa4c9c11831d1baa960fbb1b0131
58b9 priv 0xaa4c9c11831d1baa960fbb1b013158b9 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
aaa authentication dot1x default group CTS-RADIUS
aaa authorization cts default group CTS-RADIUS

ip route 10.11.3.0/24 10.11.103.206 name Enclave3
ip route 10.11.103.192/30 10.11.103.206 name Enclave3-bridge
ip pim rp-address 10.11.255.1 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
vlan 1-2,10,20,242,2000-2002,2004,3001-3002,3004,3150
vlan 2
  name AltNative
vlan 20
  name ASA-Cluster-Control
vlan 242
  name VMware-blade
vlan 2001
  name Enclave1-North
vlan 2002
  name Enclave2-North
vlan 2004
```

```
    name FPS-outside
vlan 3001
  name Enclave1-South
vlan 3002
  name Enclave2-South
vlan 3004
  name FPS54-Inside

route-map Enclave3 permit 10
  description Enclave3 Subnets
  match interface Vlan2003
vrf context management
  ip route 0.0.0.0/0 10.11.236.1
vpc domain 100
  role priority 10
  peer-keepalive destination 10.11.236.32 source 10.11.236.31
  peer-gateway


interface Vlan1
  no ip redirects
  no ipv6 redirects

interface Vlan10
  description <RAGG1&2 interface>
  no shutdown
  no ip redirects
  ip address 10.11.210.45/30
  ip router ospf 5 area 0.0.0.0

interface Vlan20
  description <** ASA Cluster control **>
  no shutdown
  no ip redirects

interface Vlan242
  no shutdown
  ip address 10.11.242.254/24
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.242.1

interface Vlan2001
  no shutdown
  ip address 10.11.1.254/24
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.1.1

interface Vlan2002
  no shutdown
  ip address 10.11.2.254/24
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.2.1

interface Vlan2003
  no shutdown
  ip address 10.11.103.202/29
  ip router ospf 5 area 0.0.0.0
  hsrp 1
```

```
      preempt
      ip 10.11.103.201

interface Vlan2004
  no shutdown
  no ip redirects
  ip address 10.11.4.254/24
  no ipv6 redirects
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.4.1

interface port-channel10
  description <<vPC peer-link>>
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
  spanning-tree port type network
  vpc peer-link

interface port-channel13
  description <<VPC Peer SACCESS-3>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  vpc 13

interface port-channel14
  description <<VPC Peer SACCESS-4>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  vpc 14

interface port-channel20
  description ASA Cluster Data Link
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2001-2100,3001-3100
  vpc 20

interface port-channel21
  description <<ASA-5-Control>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  no lacp graceful-convergence
  vpc 21

interface port-channel22
  description <<ASA-6-Control>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  no lacp graceful-convergence
  vpc 22

interface port-channel23
```

```
    description <<ASA-7-Control>>
    switchport
    switchport access vlan 20
    spanning-tree port type edge
    no lacp graceful-convergence
    vpc 23

interface port-channel24
    description <<ASA-8-Control>>
    switchport
    switchport access vlan 20
    spanning-tree port type edge
    no lacp graceful-convergence
    vpc 24

interface port-channel111
    description <<VPC Peer UCS Fabric A>>
    switchport
    switchport mode trunk
    switchport trunk native vlan 242
    switchport trunk allowed vlan 242,3001-3100
    spanning-tree port type normal
    vpc 111

interface port-channel112
    description <<VPC Peer UCS Fabric B>>
    switchport
    switchport mode trunk
    switchport trunk native vlan 242
    switchport trunk allowed vlan 242,3001-3100
    spanning-tree port type normal
    vpc 112

interface port-channel150
    mtu 9216
    ip address 10.11.210.74/30
    ip ospf network point-to-point
    no ip ospf passive-interface
    ip router ospf 5 area 0.0.0.0
    ip pim sparse-mode
    ip igmp version 3

interface port-channel151
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
    switchport trunk allowed vlan add 3150,3201-3400
    mtu 9216
    vpc 151

interface port-channel251
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
    switchport trunk allowed vlan add 3150,3201-3400
    mtu 9216
    vpc 251

interface Ethernet4/1
    description RCORE-1 port T3/1
    ip address 10.11.210.14/30
    ip router ospf 5 area 0.0.0.0
    ip pim sparse-mode
    no shutdown
```

```
interface Ethernet4/2
  no shutdown

interface Ethernet4/3
  description RCORE-2 port T3/1
  ip address 10.11.210.22/30
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet4/4
  no shutdown

interface Ethernet4/5
  description NGA-DC-1 port 1
  switchport
  switchport monitor
  no shutdown

interface Ethernet4/6

interface Ethernet4/7

interface Ethernet4/8

interface Ethernet4/9
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown

interface Ethernet4/10
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 251 mode active
  no shutdown

interface Ethernet4/11
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet4/12
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet4/13
  description <<VPC Peer ASA5:T6>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  channel-group 21 mode active
  no shutdown

interface Ethernet4/14
  description <<VPC Peer ASA6:T6>
```

```
    switchport
    switchport access vlan 20
    spanning-tree port type edge
    channel-group 22 mode active
    no shutdown

interface Ethernet4/15
    description <<VPC Peer ASA7:T6>
    switchport
    switchport access vlan 20
    spanning-tree port type edge
    channel-group 23 mode active
    no shutdown

interface Ethernet4/16
    description <<VPC Peer ASA8:T6>
    switchport
    switchport access vlan 20
    spanning-tree port type edge
    channel-group 24 mode active
    no shutdown

interface Ethernet4/17
    description <<VPC Peer ASA-5:T8>>
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 2001-2100,3001-3100
    spanning-tree port type edge
    channel-group 20 mode active
    no shutdown

interface Ethernet4/18
    description <<VPC Peer ASA-6:T8>>
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 2001-2100,3001-3100
    spanning-tree port type edge
    channel-group 20 mode active
    no shutdown

interface Ethernet4/19
    description <<VPC Peer ASA-7:T8>>
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 2001-2100,3001-3100
    spanning-tree port type edge
    channel-group 20 mode active
    no shutdown

interface Ethernet4/20
    description <<VPC Peer ASA-8:T8>>
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 2001-2100,3001-3100
    spanning-tree port type edge
    channel-group 20 mode active
    no shutdown

interface Ethernet4/21
    no shutdown
```

```
interface Ethernet4/22
  no shutdown

interface Ethernet4/23
  no shutdown

interface Ethernet4/24
  no shutdown

interface Ethernet4/25
  description <<VPC Peer F-UCS-1:E1/17>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 242
  switchport trunk allowed vlan 242,3001-3100
  spanning-tree port type normal
  channel-group 111 mode active
  no shutdown

interface Ethernet4/26
  description <<VPC Peer SACCESS-3:E1/45>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  channel-group 13 mode active
  no shutdown

interface Ethernet4/27
  description <<VPC Peer F-UCS-1:E1/18>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 242
  switchport trunk allowed vlan 242,3001-3100
  spanning-tree port type normal
  channel-group 111 mode active
  no shutdown

interface Ethernet4/28
  description <<VPC Peer SACCESS-3:E1/46>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  channel-group 13 mode active
  no shutdown

interface Ethernet4/29
  description <<VPC Peer F-UCS-2:E1/17>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 242
  switchport trunk allowed vlan 242,3001-3100
  spanning-tree port type normal
  channel-group 112 mode active
  no shutdown

interface Ethernet4/30
  description <<VPC Peer SACCESS-4:E1/45>>
  switchport
  switchport mode trunk
```

```
    switchport trunk native vlan 2
    switchport trunk allowed vlan 3001-3100
    spanning-tree port type normal
    channel-group 14 mode active
    no shutdown

interface Ethernet4/31
  description <<VPC Peer F-UCS-2:E1/18>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 242
  switchport trunk allowed vlan 242,3001-3100
  spanning-tree port type normal
  channel-group 112 mode active
  no shutdown

interface Ethernet4/32
  description <<VPC Peer SACCESS-4:E1/46>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  channel-group 14 mode active
  no shutdown

interface Ethernet4/33

interface Ethernet4/34

interface Ethernet4/35

interface Ethernet4/36

interface Ethernet4/37

interface Ethernet4/38

interface Ethernet4/39
  no shutdown

interface Ethernet4/40
  no shutdown

interface Ethernet4/41
  description <<VPC Peer RAGG1-RAGG2:4/41>>
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
  channel-group 10 mode active
  no shutdown

interface Ethernet4/42
  description <<VPC Peer RAGG1-RAGG2:4/42>>
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
  channel-group 10 mode active
  no shutdown

interface Ethernet4/43
  description <<VPC Peer RAGG1-RAGG2:4/43>>
  switchport
  switchport mode trunk
```

```
      switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
      channel-group 10 mode active
      no shutdown

    interface Ethernet4/44
      description <<VPC Peer RAGG1-RAGG2:4/44>>
      switchport
      switchport mode trunk
      switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
      channel-group 10 mode active
      no shutdown

    interface Ethernet4/45

    interface Ethernet4/46

    interface Ethernet4/47

    interface Ethernet4/48

    interface mgmt0
      vrf member management
      ip address 10.11.236.31/24

    interface loopback0
      ip address 10.11.255.11/32
      ip router ospf 5 area 0.0.0.0
    cli alias name wr copy run start
    cli alias name bye end | exit
    line console
    line vty
    boot kickstart bootflash://sup-1/n7000-s2-kickstart.6.2.8.bin sup-1
    boot system bootflash://sup-1/n7000-s2-dk9.6.2.8.bin sup-1
    boot kickstart bootflash://sup-2/n7000-s2-kickstart.6.2.8.bin sup-2
    boot system bootflash://sup-2/n7000-s2-dk9.6.2.8.bin sup-2
    ip radius source-interface loopback0
    router ospf 5
      router-id 10.11.236.31
      redistribute static route-map Enclave3
    monitor session 5
      description NGA-DC-1
      source interface Ethernet4/1 both
      source interface Ethernet4/3 both
      destination interface Ethernet4/5
      no shut
    no system auto-upgrade epld
```

Nexus 7000 OTV Switch Configuration

```
    version 6.2(8)
    hostname OTV-1

    feature ospf
    feature pim
    feature otv
    feature private-vlan
    feature udld
    feature interface-vlan
    feature dot1x
    feature hsrp
    feature lacp
```

```
feature cts

username admin password 5 $1$8HGBNjSS$pHtE5zzU/XeYZs3ySrOs00  role vdc-admin
ip domain-lookup
aaa group server radius aaa-private-sg
ip access-list ALL_IPs
  10 permit ip any any
mac access-list ALL_MACs
  10 permit any any
ip access-list HSRP_IP
  10 permit udp any 224.0.0.2/32 eq 1985
mac access-list HSRP_VMAC
  10 permit 0000.0c07.ac00 0000.0000.00ff any
vlan access-map HSRP_Localization 10
        match mac address HSRP_VMAC
        match ip address HSRP_IP
        action drop
vlan access-map HSRP_Localization 20
        match mac address ALL_MACs
        match ip address ALL_IPs
        action forward
vlan filter HSRP_Localization vlan-list 2000-2100,3000-3100
snmp-server user admin vdc-admin auth md5 0x5517d5610c4b3a30ef75f43a51e1ed1e pri
v 0x5517d5610c4b3a30ef75f43a51e1ed1e localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vlan 1,20-24,2000-2100,2201-2300,3001-3100,3150,3201-3400

otv site-vlan 3150
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 11 deny 78da.6ed9.767e ffff.ffff.ffff
mac-list OTV_HSRP_VMAC_deny seq 20 permit 0000.0000.0000 0000.0000.0000
route-map OTV_HSRP_filter permit 10
  match mac-list OTV_HSRP_VMAC_deny
vrf context management


interface Vlan1

interface port-channel150
  mtu 9216
  ip address 10.11.210.73/30
  ip ospf network point-to-point
  ip router ospf 5 area 0.0.0.0
  ip igmp version 3

interface port-channel151
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216

interface Overlay1
  otv join-interface port-channel150
  otv control-group 239.1.0.1
  otv data-group 232.1.0.0/16
  otv extend-vlan 20, 2001-2004, 2201-2204, 3001-3004, 3201-3204
  no otv suppress-arp-nd
  no shutdown
```

```
interface Ethernet3/7
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet3/8
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet3/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown

interface Ethernet3/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown
cli alias name wr copy run start
line vty
router ospf 5
  router-id 10.11.210.73
otv-isis default
  vpn Overlay1
    redistribute filter route-map OTV_HSRP_filter
otv site-identifier 0x1
```

Site 2

Nexus 7000 Core Configuration

```
version 6.2(8)
hostname RCORE

feature ospf
feature pim
feature udld

username admin password 5 $1$jPRi9VwB$/i1VZ7G2wV3Tw8KH16Lvv0  role vdc-admin
username bmcgloth password 5 $1$U8x9uEkR$9cCECGxSSyqek76TPCw64.  role vdc-operat
or
ip domain-lookup
snmp-server user admin vdc-admin auth md5 0x73f60e2123697e8140f2c6a60a970adb pri
v 0x73f60e2123697e8140f2c6a60a970adb localizedkey
snmp-server user bmcgloth vdc-operator auth md5 0x6fce1cac3c392093e76ae4dcdfe0b2
f4 priv 0x6fce1cac3c392093e76ae4dcdfe0b2f4 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
```

```
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

ip route 0.0.0.0/0 10.12.211.30 name default
ip pim rp-address 10.11.255.1 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
vlan 1

vrf context management

interface port-channel1
  no ip redirects
  ip address 10.12.210.29/30
  no ipv6 redirects
  ip router ospf 5 area 0.0.0.0

interface Ethernet3/1
  description RAGG-1 Port T3/9
  ip address 10.12.210.13/30
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet3/2
  description RAGG-2 Port T3/9
  ip address 10.12.210.17/30
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface Ethernet3/3
  description 10Gig LINK to RCORE-2
  channel-group 1
  no shutdown

interface Ethernet3/4
  description 10Gig LINK to RCORE-2
  channel-group 1
  no shutdown

interface Ethernet3/5

interface Ethernet3/6

interface Ethernet3/7

interface Ethernet3/8
  description Private Metro WAN to DC1-RCORE-1
  ip address 192.168.11.14/30
  ip router ospf 5 area 0.0.0.0
  ip pim sparse-mode
  no shutdown

interface loopback0
  ip address 10.12.255.1/32
  ip router ospf 5 area 0.0.0.0
logging source-interface loopback 0
cli alias name wr copy run start
line vty
router ospf 5
  router-id 10.12.255.1
  area 0.0.0.0 range 10.12.210.12/30
  area 0.0.0.0 range 10.12.210.16/30
  area 0.0.0.0 range 10.12.210.28/30
  area 0.0.0.0 range 10.12.255.1/32
```

```
        area 0.0.0.0 range 192.168.11.12/30
        timers throttle spf 10 100 5000
        auto-cost reference-bandwidth 10000
```

Nexus 7000 Aggregation Switch Configuration

```
version 6.2(8)
hostname RAGG

cfs eth distribute
feature ospf
feature pim
feature private-vlan
feature udld
feature interface-vlan
feature dot1x
feature hsrp
feature lacp
feature glbp
feature cts
feature vpc

username admin password 5 $1$qsFbAVQi$6bFYSs1ZYXsxO7UGW6ei60  role vdc-admin
ip domain-lookup
aaa group server radius aaa-private-sg
snmp-server user admin vdc-admin auth md5 0x2c3fc61d1fae7e68f289c0b70ab86c0c pri
v 0x2c3fc61d1fae7e68f289c0b70ab86c0c localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

ip pim rp-address 10.11.255.1 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
vlan 1-2,10,20,242,2000-2004,3001-3002,3004,3150
vlan 2
  name AltNative
vlan 20
  name ASA-Cluster-Control
vlan 242
  name VMware-blade
vlan 2001
  name Enclave1-North
vlan 2002
  name Enclave2-North
vlan 2003
  name Enclave3-North
vlan 2004
  name FPS-outside
vlan 3001
  name Enclave1-South
vlan 3002
  name Enclave2-South
vlan 3004
  name FPS54-Inside

vrf context management
  ip route 0.0.0.0/0 10.11.246.1
vpc domain 100
  peer-switch
```

```
    role priority 11
    peer-keepalive destination 10.11.255.12 source 10.11.255.11 vrf default
    peer-gateway

interface Vlan1
  no ip redirects

interface Vlan10
  description <RAGG1&2 interface>
  no shutdown
  no ip redirects
  ip address 10.12.210.45/30
  ip router ospf 5 area 0.0.0.0

interface Vlan20
  description <** ASA Cluster control **>
  no shutdown
  no ip redirects

interface Vlan242
  no shutdown
  ip address 10.11.242.252/24
  ip ospf passive-interface
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.242.1

interface Vlan2001
  no shutdown
  ip address 10.11.1.252/24
  ip ospf passive-interface
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.1.1

interface Vlan2002
  no shutdown
  ip address 10.11.2.252/24
  ip ospf passive-interface
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.2.1

interface Vlan2003
  no shutdown
  ip address 10.11.103.202/29
  ip ospf passive-interface
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.103.201

interface Vlan2004
  no shutdown
  no ip redirects
  ip address 10.11.4.252/24
  ip ospf passive-interface
  ip router ospf 5 area 0.0.0.0
  hsrp 1
    preempt
    ip 10.11.4.1
```

```
interface port-channel10
  description <<vPC peer-link>>
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
  spanning-tree port type network
  vpc peer-link

interface port-channel13
  description <<VPC Peer SACCESS-3>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  vpc 13

interface port-channel20
  description ASA Cluster Data Link
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2001-2100,3001-3100
  spanning-tree port type edge trunk
  vpc 20

interface port-channel25
  description <<ASA-9-Control>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  no lacp graceful-convergence
  vpc 25

interface port-channel26
  description <<ASA-10-Control>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  no lacp graceful-convergence
  vpc 26

interface port-channel27
  description <<ASA-11-Control>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  no lacp graceful-convergence
  vpc 27

interface port-channel28
  description <<ASA-12-Control>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  no lacp graceful-convergence
  vpc 28

interface port-channel111
  description <<VPC Peer UCS Fabric A>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 242
```

```
    switchport trunk allowed vlan 242,3001-3100
    spanning-tree port type normal
    vpc 111

interface port-channel112
    description <<VPC Peer UCS Fabric B>>
    switchport
    switchport mode trunk
    switchport trunk native vlan 242
    switchport trunk allowed vlan 242,3001-3100
    spanning-tree port type normal
    vpc 112

interface port-channel150
    mtu 9216
    ip address 10.12.210.74/30
    ip ospf network point-to-point
    no ip ospf passive-interface
    ip router ospf 5 area 0.0.0.0
    ip pim sparse-mode
    ip igmp version 3

interface port-channel151
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
    switchport trunk allowed vlan add 3150,3201-3400
    mtu 9216
    vpc 151

interface port-channel251
    switchport
    switchport mode trunk
    switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
    switchport trunk allowed vlan add 3150,3201-3400
    mtu 9216
    vpc 251

interface port-channel341
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 2,2004,3004
    vpc 41

interface port-channel342
    switchport
    switchport mode trunk
    switchport trunk native vlan 2
    switchport trunk allowed vlan 2,2004,3004
    vpc 42

interface Ethernet3/9
    description RCORE-1 Port T3/1
    ip address 10.12.210.14/30
    ip router ospf 5 area 0.0.0.0
    ip pim sparse-mode
    no shutdown

interface Ethernet3/10
    description RCORE-2 Port T3/1
    ip address 10.12.210.22/30
    ip router ospf 5 area 0.0.0.0
    ip pim sparse-mode
```

```
  no shutdown

interface Ethernet3/11
  description <<VPC Peer RAGG1-RAGG2:3/11>>
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
  channel-group 10 mode active
  no shutdown

interface Ethernet3/12
  description <<VPC Peer RAGG1-RAGG2:3/12>>
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10,20,242,2001-2100,3001-3100,3150
  channel-group 10 mode active
  no shutdown

interface Ethernet3/13
  description <<VPC Peer ASA9:T6>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  channel-group 25 mode active
  no shutdown

interface Ethernet3/14
  description <<VPC Peer ASA10:T6>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  channel-group 26 mode active
  no shutdown

interface Ethernet3/15
  description <<VPC Peer ASA-9:T8>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge trunk
  channel-group 28 mode active
  no shutdown

interface Ethernet3/16
  description <<VPC Peer ASA12:T6>>
  switchport
  switchport access vlan 20
  spanning-tree port type edge
  channel-group 28 mode active
  no shutdown

interface Ethernet3/17
  description <<VPC Peer ASA-9:T8>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2001-2100,3001-3100
  spanning-tree port type edge trunk
  channel-group 20 mode active
  no shutdown

interface Ethernet3/18
  description <<VPC Peer
  switchport
  switchport mode trunk
```

```
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2001-2100,3001-3100
  spanning-tree port type edge
  channel-group 20 mode active
  no shutdown

interface Ethernet3/19
  description <<VPC Peer ASA-11:T8>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2001-2100,3001-3100
  spanning-tree port type edge
  channel-group 20 mode active
  no shutdown

interface Ethernet3/20
  description <<VPC Peer ASA-12:
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 2001-2100,3001-3100
  spanning-tree port type edge
  channel-group 20 mode active
  no shutdown

interface Ethernet3/21
  no shutdown

interface Ethernet3/22
  description <<VPC Peer SACCE
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  channel-group 13 mode active
  no shutdown

interface Ethernet3/23
  description <<VPC Peer SACCESS-3:E1/45>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  channel-group 13 mode active
  no shutdown

interface Ethernet3/24
  description <<VPC Peer SACCESS-3:E1/47>>
  switchport
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 3001-3100
  spanning-tree port type normal
  channel-group 13 mode active
  no shutdown

interface Ethernet3/25
  mtu 9216
  channel-group 150 mode active
  no shutdown
```

```
interface Ethernet3/26
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet3/27
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown

interface Ethernet3/28
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 251 mode active
  no shutdown

interface loopback0
  ip address 10.11.255.11/32
  ip router ospf 5 area 0.0.0.0
logging source-interface loopback 0
cli alias name wr copy run sta
line vty
router ospf 5
  router-id 10.11.255.11
  area 0.0.0.0 range 10.12.210.12/30
  area 0.0.0.0 range 10.12.210.20/30
  area 0.0.0.0 range 10.12.255.11/32
  timers throttle spf 10 100 5000
  auto-cost reference-bandwidth 10000
```

Nexus 7000 OTV Switch Configuration

```
version 6.2(8)
switchname DC2-OTV1

feature ospf
feature pim
feature otv
feature private-vlan
feature udld
feature interface-vlan
feature dot1x
feature hsrp
feature lacp
feature cts

logging level private-vlan 3
username admin password 5 $1$MfKL6NL7$O6JTthUrnS7aF2a4Rkhgv1  role vdc-admin
ip domain-lookup
aaa group server radius aaa-private-sg
ip access-list ALL_IPs
  10 permit ip any any
mac access-list ALL_MACs
```

```
  10 permit any any
ip access-list HSRP_IP
  10 permit udp any 224.0.0.2/32 eq 1985
mac access-list HSRP_VMAC
  10 permit 0000.0c07.ac00 0000.0000.00ff any
vlan access-map HSRP_Localization 10
        match mac address HSRP_VMAC
        match ip address HSRP_IP
        action drop
vlan access-map HSRP_Localization 20
        match mac address ALL_MACs
        match ip address ALL_IPs
        action forward
vlan filter HSRP_Localization vlan-list 2000-2100,3000-3100
snmp-server user admin vdc-admin auth md5 0x487057e9a1ba61a20fd65ab9040f7e05 pri
v 0x487057e9a1ba61a20fd65ab9040f7e05 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vlan 1,20-24,2000-2100,2201-2300,3001-3100,3150,3201-3400

otv site-vlan 3150
mac-list OTV_HSRP_VMAC_deny seq 10 deny 0000.0c07.ac00 ffff.ffff.ff00
mac-list OTV_HSRP_VMAC_deny seq 11 deny 78da.6ed9.767e ffff.ffff.ffff
mac-list OTV_HSRP_VMAC_deny seq 20 permit 0000.0000.0000 0000.0000.0000
route-map OTV_HSRP_filter permit 10
  match mac-list OTV_HSRP_VMAC_deny
vrf context management


interface Vlan1

interface port-channel150
  mtu 9216
  ip address 10.12.210.73/30
  ip ospf network point-to-point
  ip router ospf 5 area 0.0.0.0
  ip igmp version 3

interface port-channel151
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216

interface Overlay1
  otv join-interface port-channel150
  otv control-group 239.1.0.1
  otv data-group 232.1.0.0/16
  otv extend-vlan 20, 2001-2004, 2201-2204, 3001-3004, 3201-3204
  no otv suppress-arp-nd
  no shutdown

interface Ethernet3/33

interface Ethernet3/34
```

```
interface Ethernet3/35

interface Ethernet3/36

interface Ethernet4/7
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet4/8
  mtu 9216
  channel-group 150 mode active
  no shutdown

interface Ethernet4/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown

interface Ethernet4/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 20-24,2000-2100,2201-2300,3001-3100
  switchport trunk allowed vlan add 3150,3201-3400
  mtu 9216
  channel-group 151 mode active
  no shutdown
cli alias name wr copy run start
line vty
router ospf 5
  router-id 10.12.210.73
otv-isis default
  vpn Overlay1
    redistribute filter route-map OTV_HSRP_filter
otv site-identifier 0x2
```