



Deploying IPv6 in Campus Networks

This document guides customers in their planning or deployment of IPv6 in campus networks. This document does not introduce campus design fundamentals and best practices, IPv6, transition mechanisms, or IPv4-to-IPv6 feature comparisons. [Document Objectives, page 3](#) provides additional information about the purpose of this document and references to related documents.

Contents

Introduction	3
Document Objectives	3
Document Format and Naming Conventions	3
Deployment Models Overview	4
Dual-Stack Model	4
Overview	4
Benefits and Drawbacks of This Solution	4
Solution Topology	5
Tested Components	5
Hybrid Model	6
Overview	6
Hybrid Model—Example 1	6
Overview	6
Solution Requirements	9
Benefits and Drawbacks of This Solution	9
Solution Topology	10
Tested Components	10
Hybrid Model—Example 2	11
Overview	11
Benefits and Drawbacks of This Solution	11



Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Copyright © 2006 Cisco Systems, Inc. All rights reserved.

Solution Topology	12
Tested Components	12
Service Block Model	13
Overview	13
Benefits and Drawbacks of This Solution	13
Solution Topology	14
Tested Components	15
General Considerations	16
Addressing	16
Physical Connectivity	17
VLANs	17
Routing	18
High Availability	18
QoS	20
Security	23
Multicast	27
Management	28
Scalability and Performance	29
Dual-Stack Model—Implementation	31
Network Topology	31
Physical/VLAN Configuration	33
Routing Configuration	35
High-Availability Configuration	37
QoS Configuration	37
Multicast Configuration	38
Routed Access Configuration	40
Hybrid Model—Example 1 Implementation	43
Network Topology	43
Physical Configuration	44
Tunnel Configuration	45
QoS Configuration	51
Infrastructure Security Configuration	52
Service Block Model—Implementation	52
Network Topology	52
Physical Configuration	54
Tunnel Configuration	56
QoS Configuration	59
Infrastructure Security Configuration	59
Conclusion	60

Future Work	61
Additional References	61
Appendix—Configuration Listings	63
Dual-Stack Model (DSM)	63
3750-acc-1	63
3750-acc-2	68
6k-dist-1	71
6k-dist-2	78
6k-core-1	84
Dual-Stack Model (DSM)—Routed Access	94
3750-acc-1	94
6k-dist-1	99
6k-dist-2	104
Hybrid Model Example 1 (HME1)	110
6k-core-1	110
6k-core-2	116
Service Block Model (SBM)	121
6k-sb-1	122
6k-sb-2	127

Introduction

Document Objectives

The reader must be familiar with the Cisco campus design best practices recommendations as well as the basics of IPv6 and associated transition mechanisms. The prerequisite knowledge can be acquired through many documents and training opportunities available both through Cisco and the industry at large. Following are a few recommended information resources for these areas of interest:

- Cisco Solution Reference Network Design (SRND) Campus Guides—
http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html#anchor2
- Cisco IPv6 CCO website—<http://www.cisco.com/go/ipv6>
- Catalyst 6500 Series Cisco IOS Software Configuration Guide, 12.2SX—
<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/index.htm>
- Catalyst 3750 Switch Software Configuration Guide, 12.2(25)SEE—
<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12225see/scg/index.htm>
- “Deploying IPv6 Networks” by Ciprian P. Popoviciu, Eric Levy-Abegnoli, Patrick Grossetete (ISBN-10:1-58705-210-5; ISBN-13:978-1-58705-210-1)—
<http://www.ciscopress.com/bookstore/product.asp?isbn=1587052105&rl=1>
- go6 IPv6 Portal—IPv6 Knowledge Center— http://wiki.go6.net/index.php?title=Main_Page

- 6NET—Large-Scale International IPv6 Pilot Network—<http://www.6net.org/>
- IETF IPv6 Working Group—<http://www.ietf.org/html.charters/ipv6-charter.html>
- IETF IPv6 Operations Working Group—<http://www.ietf.org/html.charters/v6ops-charter.html>

Document Format and Naming Conventions

This document provides a brief overview of the various campus IPv6 deployment models and general deployment considerations, and also provides the implementation details for each model individually.

In addition to any configurations shown in the general considerations and implementation sections, the full configurations for each campus switch can be found in [Appendix—Configuration Listings, page 66](#).

The following abbreviations are used throughout this document when referring to the campus IPv6 deployment models:

- Dual-stack model (DSM)
- Hybrid model example 1 (HME1)
- Hybrid model example 2 (HME2)
- Service block model (SBM)

User-defined properties such as access control list (ACL) names and quality of service (QoS) policy definitions are shown in ALL CAPS to differentiate them from command-specific policy definitions.



Note

The applicable commands in each section below are in **red** text.

Deployment Models Overview

This section provides a high-level overview of the following three campus IPv6 deployment models and describes their benefits applicability:

- DSM
- Hybrid Model
 - HME1—Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) + dual-stack
 - HME2—Manually-configured tunnels + dual-stack
- SBM—Combination of ISATAP, manually-configured tunnels, and dual-stack

Dual-Stack Model

Overview

DSM is completely based on the dual-stack transition mechanism. A device or network on which two protocol stacks have been enabled at the same time operates in dual-stack mode. Examples of previous uses of dual-stack include IPv4 and IPX, or IPv4 and Apple Talk co-existing on the same device.

Dual-stack is the preferred, most versatile way to deploy IPv6 in existing IPv4 environments. IPv6 can be enabled wherever IPv4 is enabled along with the associated features required to make IPv6 routable, highly available, and secure. In some cases, IPv6 is not enabled on a specific interface or device because of the presence of legacy applications or hosts for which IPv6 is not supported. Inversely, IPv6 may be enabled on interfaces and devices for which IPv4 support is no longer needed.

The tested components area of each section of this paper gives a brief view of the common requirements for the DSM to be successfully implemented. The most important consideration is to ensure that there is hardware support of IPv6 in campus network components such as switches. Within the campus network, link speeds and capacity often depend on such issues as the number of users, types of applications, and latency expectations. Because of the typically high data rate requirements in this environment, Cisco does not recommend enabling IPv6 unicast or multicast layer switching on software forwarding-only platforms. Enabling IPv6 on software forwarding-only campus switching platforms may be suitable in a test environment or small pilot network, but certainly not in a production campus network.

Benefits and Drawbacks of This Solution

Deploying IPv6 in the campus using DSM offers several advantages over the hybrid and service block models. The primary advantage of DSM is that it does not require tunneling within the campus network. DSM runs the two protocols as “ships-in-the-night”, meaning that IPv4 and IPv6 run alongside one another and have no dependency on each other to function except that they share network resources. Both IPv4 and IPv6 have independent routing, high availability (HA), QoS, security, and multicast policies. Dual-stack also offers processing performance advantages because packets are natively forwarded without having to account for additional encapsulation and lookup overhead.

Customers who plan to or have already deployed the Cisco routed access design will find that IPv6 is also supported because the network devices support IPv6 in hardware. Discussion on implementing IPv6 in the routed access design follows in [Dual-Stack Model—Implementation, page 33](#).

The primary drawback to DSM is that network equipment upgrades might be required when the existing network devices are not IPv6-capable.

[Conclusion, page 63](#) summarizes the benefits and challenges of the various campus design models in a tabular format.

Solution Topology


[Figure 1](#) shows a high-level view of the DSM-based deployment in the campus networks. This example is the basis for the detailed configurations that are presented later in this document.



Note

The data center block is shown here for reference only and is not discussed in this document. A separate document will be published to discuss the deployment of IPv6 in the data center.

Figure 1 Dual-Stack Model Example



Tested Components

Table 1 lists the components that were used and tested in the DSM configuration.

Table 1 DSM Tested Components

Campus Layer	Hardware	Software
Access layer	Cisco Catalyst 3750	Advanced IP Services–12.2(25)SED1
	Catalyst 6500 Supervisor 32 or 720	Advanced Enterprise Services SSH–12.2(18)SXF5
Host devices	Various laptops—IBM, HP, and Apple	Microsoft Windows XP SP2, Vista RC1, Apple Mac OS X 10.4.7, and Red Hat Enterprise Linux WS
Distribution layer	Catalyst 6500 Supervisor 32 or 720	Advanced Enterprise Services SSH–12.2(18)SXF5
Core layer	Catalyst 6500 Supervisor 720	Advanced Enterprise Services SSH–12.2(18)SXF5

Hybrid Model

Overview

The hybrid model strategy is to employ two or more independent transition mechanisms with the same deployment design goals. Flexibility is the key aspect of the hybrid approach in which any combination of transition mechanisms can be leveraged to best fit a given network environment.

The hybrid model adapts as much as possible to the characteristics of the existing network infrastructure. Transition mechanisms are selected based on multiple criteria, such as IPv6 hardware capabilities of the network elements, number of hosts, types of applications, location of IPv6 services, and network infrastructure feature support for various transition mechanisms.

The following are the three main IPv6 transition mechanisms leveraged by this model:

- Dual-stack—Deployment of two protocol stacks: IPv4 and IPv6
- ISATAP—Host-to-router tunneling mechanism that relies on an existing IPv4-enabled infrastructure
- Manually-configured tunnels—Router-to-router tunneling mechanism that relies on an existing IPv4-enabled infrastructure

The following two sections discuss the hybrid model in the context of two specific examples:

- HME1—Focuses on using ISATAP to connect hosts located in the access layer to the core layer switches plus dual-stack in the core layer and beyond
- HME2—Focuses on using manually-configured tunnels between the distribution layer and the data center aggregation layer plus dual-stack in the access-to-distribution layer

The subsequent sections provide a high-level discussion of these models. Later in the document, the HME1 implementation is discussed in detail.

Hybrid Model—Example1

Overview

HME1 provides hosts with access to IPv6 services even when the underlying network infrastructure may not support IPv6 natively.

The key aspect of HME1 is the fact that hosts located in the campus access layer can use IPv6 services when the distribution layer is not IPv6-capable or enabled. The distribution layer switch is most commonly the first Layer 3 gateway for the access layer devices. If IPv6 capabilities are not present in the existing distribution layer switches, the hosts cannot gain access to IPv6 addressing (stateless autoconfiguration or DHCP for IPv6) router information, and subsequently cannot access the rest of the IPv6-enabled network.

Tunneling can be used on the IPv6-enabled hosts to provide access to IPv6 services located beyond the distribution layer. Example 1 leverages the ISATAP tunneling mechanisms on the hosts in the access layer to provide IPv6 addressing and off-link routing. The Microsoft Windows XP and Vista hosts in the access layer need to have IPv6 enabled and either a static ISATAP router definition or DNS “A” record entry configured for the ISATAP router address.




Note

The configuration details are shown in [Network Topology, page 46](#).

[Figure 2](#) shows the basic connectivity flow for HME1.

Figure 2 Hybrid Model Example 1—Connectivity Flow




1. The host establishes an ISATAP tunnel to the core layer.
2. The core layer switches are configured with ISATAP tunnel interfaces and are the termination point for ISATAP tunnels established by the hosts.
3. Pairs of core layer switches are redundantly configured to accept ISATAP tunnel connections to provide high availability of the ISATAP tunnels. Redundancy is available by configuring both core layer switches with loopback interfaces that share the same IPv4 address. Both switches use this redundant IPv4 address as the tunnel source for ISATAP. When the host connects to the IPv4 ISATAP router address, it connects to one of the two switches (this can be load balanced or be configured to have a preference for one switch over the other). If one switch fails, the IPv4 Interior Gateway Protocol (IGP) converges and uses the other switch, which has the same IPv4 ISATAP address as the primary. The failover takes as long as the IGP convergence time + the Neighbor Unreachability Detection (NUD) time expiry. With Microsoft Vista configurations, basic load balancing of the ISATAP routers (core switches) can be implemented. For more information on the Microsoft implementation of ISATAP on Windows platforms, see the following URL:
<http://www.microsoft.com/downloads/details.aspx?FamilyId=B8F50E07-17BF-4B5C-A1F9-5A09E2AF698B&displaylang=en>
4. The dual-stack configured server accepts incoming and/or establishes outgoing IPv6 connections using the directly accessible dual-stack-enabled data center block.

One method to help control where ISATAP tunnels can be terminated and what resources the hosts can reach over IPv6 is to use VLAN or IPv4 subnet-to-ISATAP tunnel matching.

If the current network design has a specific VLAN associated with ports on an access layer switch and the users attached to that switch are receiving IPv4 addressing based on the VLAN to which they belong, a similar mapping can be done with IPv6 and ISATAP tunnels.

Figure 3 illustrates the process of matching users in a specific VLAN and IPv4 subnet with a specific ISATAP tunnel.

Figure 3 Hybrid Model Example 1—ISATAP Tunnel Mapping



1. The core layer switch is configured with a loopback interface with the address of 10.122.10.2, which is used as the tunnel source for ISATAP, and is used only by users located on the 10.120.2.0/24 subnet.
2. The host in the access layer is connected to a port that is associated with a specific VLAN. In this example, the VLAN is “VLAN-2”. The host in VLAN-2 is associated with an IPv4 subnet range (10.120.2.0/24) in the DHCP server configuration.

The host is also configured for ISATAP and has been statically assigned the ISATAP router value of 10.122.10.2. This static assignment can be implemented in several ways. An ISATAP router setting can be defined via a command on the host (**netsh interface ipv6 isatap set router 10.122.10.2**—details provided later in the document), which can be manually entered or scripted via a Microsoft SMS Server, Windows Scripting Host, or a number of other scripting methods. The script can determine to which value to set the ISATAP router by examining the existing IPv4 address of the host. For instance, the script can analyze the host IPv4 address and determine that the value “2” in the 10.120.2.x/24 address signifies the subnet value. The script can then apply the command using the ISATAP router address of 10.122.10.2, where the “2” signifies subnet or VLAN 2. The 10.122.10.2 address is actually a loopback address on the core layer switch and is used as the tunnel endpoint for ISATAP.



Note

Configuration details on the method described above can be found in [Network Topology, page 46](#).

A customer might want to do this for the following reasons:

- **Control and separation**—If a security policy is in place that disallows certain IPv4 subnets from accessing a specific resource, and ACLs are used to enforce the policy. What happens if HME1 is implemented without consideration for this policy? If the restricted resources are also IPv6 accessible, those users who were previously disallowed access via IPv4 can now access the protected resource via IPv6. If hundreds or thousands of users are configured for ISATAP and a single ISATAP tunnel interface is used on the core layer device, controlling the source addresses via ACLs would be very difficult to scale and manage. If the users are logically separated into ISATAP tunnels in the same way they are separated by VLANs and IPv4 subnets, ACLs can be easily deployed to permit or deny access based on the IPv6 source, source/destination, and even Layer 4 information.
- **Scale**—It has been a common best practice for years to control the number of devices within each single VLAN of the campus networks. This practice has traditionally been enforced for broadcast domain control. Although IPv6 and ISATAP tunnels do not use broadcast, there are still scalability considerations to consider. Based on customer deployment experiences, it was concluded that it was better to spread fewer hosts among a greater number of tunnel interfaces than it was to have a greater

number of hosts across a single or a few tunnel interfaces. The optimal number of hosts per ISATAP tunnel interface is not known, but this is most likely not a significant issue unless thousands of hosts are deployed in an ISATAP configuration. Nevertheless, continue to watch for documents from Cisco (<http://www.cisco.com/ipv6>) and independent test organizations on ISATAP scalability results and best practices.

Solution Requirements

The following are the main solution requirements for HME1 strategies:

- IPv6 and ISATAP support on the operating system of the host machines
- IPv6/IPv4 dual-stack and ISATAP feature support on the core layer switches

As mentioned previously, numerous combinations of transition mechanisms can be used to provide IPv6 connectivity within the enterprise campus environment, such as the following two alternatives to the requirements listed above:

- Using 6to4 tunneling instead of ISATAP if multiple host operating systems such as Linux, FreeBSD, Sun Solaris, and Mac OS X are in use within the access layer. The reader should research the security implications of using 6to4.
- Terminating tunnels at a network layer different than the core layer, such as the data center aggregation layer.



Note

The 6to4 and non-core layer alternatives are not discussed in this document and are listed only as secondary options to the deployment recommendations for the HME1.

Benefits and Drawbacks of This Solution

The primary benefit of HME1 is that the existing network equipment can be leveraged without the need for upgrades, especially the distribution layer switches. If the distribution layer switches currently provide acceptable IPv4 service and performance and are still within the depreciation window, HME1 may be a suitable choice.

It is important to understand the drawbacks of the hybrid model, specifically with HME1:

- It is not yet known how much the ISATAP portion of the design can scale. Questions such as the following still need to be answered:
 - How many hosts should terminate to a single tunnel interface on the switch?
 - How much IPv6 traffic within the ISATAP tunnel is too much for a specific host? Tunnel encapsulation/decapsulation is done by the CPU on the host.
- IPv6 multicast is not supported within ISATAP tunnels. This is a limitation that needs to be resolved within RFC 4214.
- Terminating ISATAP tunnels in the core layer makes the core layer appear as an access layer to the IPv6 traffic. Network administrators and network architects design the core layer to be highly optimized for the role it plays in the network, which is very often to be stable, simple, and fast. Adding a new level of intelligence to the core layer may not be acceptable.

As with any design that uses tunneling, considerations that must be accounted for include performance, management, security, scalability, and availability. The use of tunnels is always a secondary recommendation to the DSM design.

[Conclusion, page 63](#) summarizes the benefits and challenges of the various campus design models in a tabular format.

Solution Topology


Figure 4 shows a high-level view of the campus HME1. This example is the basis for the detailed configurations that follow later in this document.



Note

The data center block is shown here for reference purpose only and is not discussed in this document. A separate document will be published to discuss the deployment of IPv6 in the data center.

Figure 4 Hybrid Model Example 1



Tested Components

[Table 2](#) lists the components used and tested in the HME1 configuration.

Table 2 HME1 Tested Components

Campus Layer	Hardware	Software
Access layer	Catalyst 3750	Advanced IP Services—12.2(25)SED1
Host devices	Various laptops—IBM, HP	Microsoft Windows XP SP2, Vista RC1
Distribution layer	Catalyst 3750	Advanced IP Services—12.2(25)SED1
	Catalyst 4500 Supervisor 5	Enhanced L3 3DES—12.2.25.EWA6

Table 2 HME1 Tested Components (continued)

	Catalyst 6500 Supervisor 2/MSFC2	Advanced Enterprise Services SSH—12.2(18)SXF5
Core layer	Catalyst 6500 Supervisor 720	Advanced Enterprise Services SSH—12.2(18)SXF5

Hybrid Model—Example 2

Overview

HME2 provides access to IPv6 services by bridging the gap with the core layer support of IPv6. In this example, dual-stack is supported in the access/distribution layers and also in the data center access and aggregation layers. Common reasons why the core layer might not be enabled for IPv6 are either that the core layer does not have hardware-based IPv6 support at all, or has limited IPv6 support but with low performance capabilities.

The configuration uses manually-configured tunnels exclusively from the distribution-to-aggregation layers. Two tunnels from each switch are used for redundancy and load balancing. From an IPv6 perspective, the tunnels can be viewed as virtual links between the distribution and aggregation layer switches. On the tunnels, routing and IPv6 multicast are configured in the same manner as with a dual-stack configuration. QoS differs only in that **mls qos trust dscp** statements apply to the physical interfaces connecting to the core versus the tunnel interfaces. This configuration should be considered for any non-traditional QoS configurations on the core that may impact tunneled or IPv6 traffic because the QoS policies on the core would not have visibility into the IPv6 packets. Similar considerations apply to the security of the network core. If special security policies exist in the core layer, those policies need to be modified (if supported) to account for the tunneled traffic crossing the core.

For more information about the operation and configuration of manually-configured tunnels, refer to [Additional References, page 64](#).

Benefits and Drawbacks of This Solution

HME2 is a good model to use if the campus core is being upgraded or has plans to be upgraded, and access to IPv6 services is required before the completion of the core upgrade.

Like most traffic in the campus, IPv6 should be forwarded as fast as possible. This is especially true when tunneling is used because there is an additional step of processing involved in the encapsulation and decapsulation of the IPv6 packets. Cisco Catalyst platforms such as the Catalyst 6500 Supervisor 32 and 720 forward tunneled IPv6 traffic in hardware.

In many networks, HME2 has less applicability than HME1, but is nevertheless discussed in the model overview section as another option. HME2 is not shown in the configuration/implementation section of this document because the implementation is relatively straightforward and mimics most of the considerations of the dual-stack model as it applies to routing, QoS, multicast, infrastructure security, and management.


As with any design that uses tunneling, considerations that must be accounted for include performance, management (lots of static tunnels are difficult to manage), scalability, and availability. The use of tunnels is always a secondary recommendation to the DSM design.

[Conclusion, page 63](#) summarizes the benefits and challenges of the various campus design models in a tabular format.

Solution Topology

Figure 5 provides a high-level perspective of HME2. As previously mentioned, the access/distribution layers fully support IPv6 (in either a Layer 2 access or Layer 3 routed access model), and the data center access/aggregation layers support IPv6 as well. The core layer does not support IPv6 in this example. A redundantly-configured pair of manually-configured tunnels is used between the distribution and aggregation layer switches to provide IPv6 forwarding across the core layer.

Figure 5 Hybrid Model Example 2



Tested Components

Table 3 lists the components used and tested in the HME2 configuration.

Table 3 HME2 Tested Components

Campus Layer	Hardware	Software
Access layer	Catalyst 3750	Advanced IP Services—12.2(25)SED1
	Catalyst 6500 Supervisor 32 or 720	Advanced Enterprise Services SSH—12.2(18)SXF5
Host devices	Various laptops—IBM, HP and Apple	Microsoft Windows XP SP2, Vista RC1, Apple Mac OS X 10.4.7, and Red Hat Enterprise Linux WS
Distribution layer	Catalyst 6500 Supervisor 32 or 720	Advanced Enterprise Services SSH—12.2(18)SXF5

Table 3 HME2 Tested Components (continued)

Core layer	Catalyst 6500 Supervisor 2/MSFC2	Advanced Enterprise Services SSH—12.2(18)SXF5
Data center aggregation layer	Catalyst 6500 Supervisor 720	Advanced Enterprise Services SSH—12.2(18)SXF5

Service Block Model

Overview

SBM is the most different of the various campus models discussed in this paper. Although the concept of a service block-like design is not a new concept, the SBM does offer unique capabilities to customers facing the challenge of providing access to IPv6 services in a short time. A service block-like approach has also been used in other design areas such as Cisco Network Virtualization (http://www.cisco.com/en/US/netsol/ns658/networking_solutions_package.html), which refers to this concept as the “Services Edge”. The SBM is unique in that it can be deployed as an overlay network without any impact to the existing IPv4 network, and is completely centralized. This overlay network can be implemented rapidly while allowing for high availability of IPv6 services, QoS capabilities, and restriction of access to IPv6 resources with little or no changes to the existing IPv4 network.

As the existing campus network becomes IPv6 capable, the SBM can become decentralized. Connections into the SBM are changed from tunnels (ISATAP and/or manually-configured) to dual-stack connections. When all the campus layers are dual-stack capable, the SBM can be dismantled and re-purposed for other uses.

The SBM deployment is based on a redundant pair of Catalyst 6500 switches with a Supervisor 32 or Supervisor 720. The key to maintaining a highly scalable and redundant configuration in the SBM is to ensure that a high-performance switch, supervisor, and modules are used to handle the load of the ISATAP, manually-configured tunnels, and dual-stack connections for an entire campus network. As the number of tunnels and required throughput increases, it may be necessary to distribute the load across an additional pair of switches in the SBM.

There are many similarities between the SBM example given in this document and the combination of the HME1 and HME2 examples. The underlying IPv4 network is used as the foundation for the overlay IPv6 network being deployed. ISATAP provides access to hosts in the access layer (similar to HME1). Manually-configured tunnels are used from the data center aggregation layer to provide IPv6 access to the applications and services located in the data center access layer (similar to HME2). IPv4 routing is configured between the core layer and SMB switches to allow visibility to the SMB switches for the purpose of terminating IPv6-in-IPv4 tunnels. In the example discussed in this paper, however, the extreme case is analyzed where there are no IPv6 capabilities anywhere in the campus network (access, distribution, or core layers). The SBM example used in this document has the switches directly connected to the core layer via redundant high-speed links.

Benefits and Drawbacks of This Solution

From a high-level perspective, the advantages to implementing the SBM are the pace of IPv6 services delivery to the hosts, the lesser impact on the existing network configuration, and the flexibility of controlling the access to IPv6-enabled applications.

In essence, the SBM provides control over the pace of IPv6 service rollout by leveraging the following:

- Per-user and/or per-VLAN tunnels can be configured via ISATAP to control the flow of connections and allow for the measurement of IPv6 traffic use.
- Access on a per-server or per-application basis can be controlled via ACLs and/or routing policies at the SBM. This level of control allows for access to one, a few, or even many IPv6-enabled services while all other services remain on IPv4 until those services can be upgraded or replaced. This enables a “per service” deployment of IPv6.
- Allows for high availability of ISATAP and manually-configured tunnels as well as all dual-stack connections.
- Flexible options allow hosts access to the IPv6-enabled ISP connections, either by allowing a segregated IPv6 connection used only for IPv6-based Internet traffic or by providing links to the existing Internet edge connections that have both IPv4 and IPv6 ISP connections.
- Implementation of the SBM does not disrupt the existing network infrastructure and services.

As mentioned in the case of HME1 and HME2, there are drawbacks to any design that relies on tunneling mechanisms as the primary way to provide access to services. The SBM not only suffers from the same drawbacks as the HME designs (lots of tunneling), but also adds the cost of additional equipment not found in HME1 or HME2. More switches (the SBM switches), line cards to connect the SBM and core layer switches, and any maintenance or software required represent additional expenses.

Because of the list of drawbacks for HME1, HME2, and SBM, Cisco recommends to always try to deploy the DSM.

[Conclusion, page 63](#) summarizes the benefits and challenges of the various campus design models in a tabular format.

Solution Topology

Two portions of the SBM design are discussed in this document. [Figure 6](#) shows the ISATAP portion of the design and [Figure 7](#) shows the manually-configured tunnel portion of the design. These views are just two of the many combinations that can be generated in a campus network and differentiated based on the goals of the IPv6 design and the capabilities of the platforms and software in the campus infrastructure.

As mentioned previously, the data center layers are not specifically discussed in this document because a separate document will focus on the unique designs and challenges of the data center. This document presents basic configurations in the data center for the sake of completeness. Based on keeping the data center portion of this document as simple as possible, the data center aggregation layer is shown as using manually-configured tunnels to the SBM and dual-stack from the aggregation layer to the access layer.

[Figure 6](#) shows the redundant ISATAP tunnels coming from the hosts in the access layer to the SBM switches. The SBM switches are connected to the rest of the campus network by linking directly to the core layer switches via IPv4-enabled links. The SBM switches are connected to each other via a dual-stack connection that is used for IPv4 and IPv6 routing and HA purposes.

Figure 6 Service Block Model—Connecting the Hosts (ISATAP Layout)





Figure 7 shows the redundant, manually-configured tunnels connecting the data center aggregation layer and the service blocks. Hosts located in the access layer can now reach IPv6 services in the data center access layer using IPv6. Refer to [Conclusion, page 63](#) for the details of the configuration.

Figure 7 Service Block Model—Connecting the Data Center (Manually-Configured Tunnel Layout)



Tested Components

Table 4 lists the components used and tested in the SBM configuration.

Table 4 SBM Tested Components

Campus Layer	Hardware	Software
Access layer	Catalyst 3750	Advanced IP Services—12.2(25)SED1
	Catalyst 6500 Supervisor 32 or 720	Advanced Enterprise Services SSH—12.2(18)SXF5
Host devices	Various laptops—IBM, HP	Microsoft Windows XP SP2, Vista RC1
Distribution layer	Catalyst 3750	Advanced IP Services—12.2(25)SED1
	Catalyst 4500 Supervisor 5	Enhanced L3 3DES—12.2.25.EWA6
	Catalyst 6500 Supervisor 2/MSFC2	Advanced Enterprise Services SSH—12.2(18)SXF5
Core layer	Catalyst 6500 Supervisor 720	Advanced Enterprise Services SSH—12.2(18)SXF5
Service block	Catalyst 6500 Supervisor 32 or 720	Advanced Enterprise Services SSH—12.2(18)SXF5

General Considerations

Many considerations apply to all the deployment models discussed in this document. This section focuses on the general ones that apply to deploying IPv6 in a campus network regardless of the deployment model being used. If a particular consideration must be understood in the context of a specific model, this model is called out along with the consideration. Also, the configurations for any model-specific considerations can be found in the implementation section of that model.

All campus IPv6 models discussed in this document leverage the existing campus network design as the foundation for providing physical access, VLANs, IPv4 routing (for tunnels), QoS (for tunnels), infrastructure security (protecting the tunnels), and availability (device, link, trunk, and routing). When dual-stack is used, nearly all design principles found in Cisco campus design best practice documents are applicable to both IPv4 and IPv6.

It is critical to understand the Cisco campus best practice recommendations before jumping into the deployment of the IPv6 campus models discussed in this document.

The Cisco campus design best practice documents can be found under the “Campus” section at the following URL: <http://www.cisco.com/go/srnd>.

Addressing

As mentioned previously, this document is not an introductory document and does not discuss the basics of IPv6 addressing. However, it is important to discuss a few addressing considerations for the network devices.

In most cases, using a /64 prefix on a point-to-point (p2p) link is fine. IPv6 was designed to have a large address space and even with poor address management in place, the customer should not experience address constraints.

Some network administrators think that a /64 prefix for p2p links is a waste. There has been quite a bit of discussion within the IPv6 community about the practice of using longer prefixes for p2p links. For network administrators who want to more tightly control the address space, it is safe to use a /126 prefix on p2p links in much the same way as /30 is used with IPv4.

RFC 3627 (<http://www.ietf.org/rfc/rfc3627.txt>) discusses the reasons why the use of a /127 prefix is harmful and should be discouraged.

In general, Cisco recommends using either a /64 or /126 on p2p links.

Efforts are being made within IETF to better document the address assignment guidelines for varying address types and prefix links. IETF work within the IPv6 operations working group can be tracked at the following URL: <http://www.ietf.org/html.charters/v6ops-charter.html>.

The p2p configurations shown in this document use /64 prefixes.

Physical Connectivity

Considerations for physical connectivity with IPv6 are the same as with IPv4, with the addition of the following three elements:

- Ensuring that there is sufficient bandwidth for both existing and new traffic
This is an important factor for the deployment of any new technology, protocol, or application.
- Understanding how IPv6 deals with the maximum transmission unit (MTU) on a link
This document is not an introductory document for basic IPv6 protocol operation or specifications. Cisco recommends reading the following documentation for more information on MTU and fragmentation in IPv6. A good starting point for understanding MTU and Path MTU Discovery (PMTUD) for IPv6 is with RFC 2460 and RFC 1981 at the following URLs:
 - <http://www.ietf.org/rfc/rfc2460.txt>
 - <http://www.ietf.org/rfc/rfc1981.txt>
- IPv6 over wireless LANs (WLANs)
IPv6 should operate correctly over WLAN access points in much the same way as IPv6 operates over Layer 2 switches. However, the reader must consider IPv6 specifics in WLAN environments include managing WLAN devices (APs and controllers) via IPv6, and controlling IPv6 traffic via AP or controller-based QoS, VLANs, and ACLs. IPv6 must be supported on the AP and/or controller devices to take advantage of these more intelligent services on the WLAN devices.

Cisco supports the use of IPv6-enabled hosts that are directly attached to Cisco IP Phone ports, which are switch ports and operate in much the same way as plugging the host directly into a Catalyst Layer 2 switch.

In addition to the above considerations, Cisco recommends that a thorough analysis of the existing traffic profiles, memory, and CPU utilization on both the hosts and network equipment, and also the Service Level Agreement (SLA) be completed before implementing any of the IPv6 models discussed in this document.

VLANs

VLAN considerations for IPv6 are the same as for IPv4. When dual-stack configurations are used, both IPv4 and IPv6 traverse the same VLAN. When tunneling is used, IPv4 and the tunneled IPv6 (protocol 41) traffic traverse the VLAN. The use of private VLANs is not included in any of the deployment models discussed in this document and it was not tested, but will be included in future campus IPv6 documents.

The use of IPv6 on data VLANs that are trunked along with voice VLANs (behind IP Phones) is fully supported.

For the current VLAN design recommendations, see the references to the Cisco campus design best practice documents in [Additional References, page 64](#).

Routing

Choosing an IGP to run in the campus network is based on a variety of factors such as platform capabilities, IT staff expertise, topology, and size of network. In this document, the IGP for IPv4 is EIGRP, but OSPFv2 for IPv4 can be also used. OSPFv3 for IPv6 is used for the IGP within the campus.

**Note**

At the time of this writing, EIGRP for IPv6 is available in Cisco IOS, but has not yet been implemented in the Catalyst platforms. Future testing and documentation will reflect design and configuration recommendations for both EIGRP and OSPFv3 for IPv6. For the latest information, watch the links on CCO at the following URL: <http://www.cisco.com/go/ipv6>.

As previously mentioned, every effort has been made to implement the current Cisco campus design best practices. Both the IPv4 and IPv6 IGP have been tuned according to the current best practices where possible. It should be one of the top priorities of any network design to ensure that the IGPs are tuned to provide a stable, scalable, and fast converging network.

One final consideration to note for OSPFv3 is that at the time of this writing, the use of IPsec for OSPFv3 has not been implemented in the tested Cisco Catalyst platforms. IPsec for OSPFv3 is used to provide authentication and encryption of OSPFv3 neighbor connections and routing updates. More information on IPsec for OSPFv3 can be found at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/sa_ospf3.htm#wp1160900

High Availability

Many aspects of high availability (HA) are not applicable to or are outside the scope of this document. Many of the HA requirements and recommendations are met by leveraging the existing Cisco campus design best practices. The following are the primary HA components discussed in this document:

- Redundant routing and forwarding paths—These are accomplished by leveraging EIGRP for IPv4 when redundant paths for tunnels are needed, and OSPFv3 for IPv6 when dual-stack is used, along with the functionality of Cisco Express Forwarding.
- Redundant Layer 3 switches for terminating ISATAP and manually-configured tunnels—These are applicable in the HME1, HME2, and SBM designs. In addition to having redundant hardware, it is important to implement redundant tunnels (ISATAP and manually-configured). The implementation sections illustrate the configuration and results of using redundant tunnels for HME1 and SBM designs.
- High availability of the first-hop gateways—In the DSM design, the distribution layer switches are the first Layer 3 devices to the hosts in the access layer. Traditional campus designs use first-hop redundancy protocols such as Hot Standby Routing Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), or Virtual Router Redundancy Protocol (VRRP) to provide first-hop redundancy.

**Note**

At the time of this writing, HSRP and GLBP are available for IPv6 in Cisco IOS, but have not yet been implemented in the Catalyst.

To deal with the lack of a first-hop redundancy protocol in the campus platforms, a method needs to be implemented to provide some level of redundancy if a failure occurs on the primary distribution switch. Neighbor Discovery for IPv6 (RFC 2461) implements the use of Neighbor Unreachability Detection (NUD). NUD is a mechanism that allows a host to determine whether a router (neighbor) in the host default gateway list is unreachable. Hosts receive the NUD value, which is known as the “reachable time”, from the routers on the local link via regularly advertised router advertisements (RAs). The default reachable time is 30 seconds.

NUD is used when a host determines that the primary gateway for IPv6 unicast traffic is unreachable. A timer is activated, and when the timer expires (reachable time value), the neighbor begins to send IPv6 unicast traffic to the next available router in the default gateway list. Under default configurations, it should take a host no longer than 30 seconds to use the next gateway in the default gateway list.

Cisco recommends that the reachable time be adjusted to 5000 msecs (5 seconds) on the VLANs facing the access layer via the **(config-if)#ipv6 nd reachable-time 5000** command. This value allows the host to fail to the secondary distribution layer switch in no more than 5 seconds. Recent testing has shown that hosts connected to Cisco Catalyst switches that use the recommended campus HA configurations along with a reachable time of 5 seconds rarely notice a failover of IPv6 traffic that takes longer than 1 second. Remember that the reachable time is the maximum time that a host should take to move to the next gateway.


One issue to note with NUD is that Microsoft Windows XP and 2003 hosts do not use NUD on ISATAP interfaces. This means that if the default gateway for IPv6 on a tunnel interface becomes unreachable, it may take a substantial amount of time to reestablish the tunnel to another tunnel and gateway. Microsoft Windows Vista and Windows Server codename “Longhorn” allow for NUD on ISATAP interfaces to be enabled. **netsh interface ipv6 set interface interface_Name_or_Index nud=enabled** can be enabled on the host directly.

The NUD value should be adjusted only on links/VLANs where hosts reside. Switches that support a real first-hop redundancy protocol such as HSRP or GLBP for IPv6 do not need to have the reachable time adjusted.

This is an overly simplistic explanation of the failover decision process because the operation of how a host determines the loss of a neighbor is quite involved, and is not discussed at length in this document. More information on how NUD works can be found at the following URL:
<http://www.ietf.org/rfc/rfc2461.txt>.

Figure 8 shows a dual-stack host in the access layer that is receiving IPv6 RAs from the two distribution layer switches. HSRP, GLBP, or VRRP for IPv6 first-hop redundancy are not being used on the two distribution switches. Adjustments to the NUD mechanism can allow for crude decision-making by the host when a first-hop gateway is lost.

Figure 8 Host Receiving an Adjusted NUD Value from Distribution Layer



- Both distribution layer switches are configured with a reachable time of 5000 msecs on the VLAN interface for the host.

```
interface Vlan2
  description ACCESS-DATA-2
  ipv6 address 2001:DB8:CAFE:2::A111:1010/64
  ipv6 nd reachable-time 5000
```

The new reachable time is sent via the RA on the next interface.

2. The host receives the RA from the distribution layer switches and modifies the local “reachable time” to the new value. On a Windows host that supports IPv6, the new reachable time can be seen by running the following:

```
netsh interface ipv6 show interface [[interface=<string>]
```

QoS


With DSM, it is easy to extend or leverage the existing IPv4 QoS policies to include the new IPv6 traffic traversing the campus network. Cisco recommends that the QoS policies be implemented to be application- and/or service-dependent instead of protocol-dependent (IPv4 or IPv6). If the existing QoS policy has specific classification, policing, and queuing for an application, that policy should treat equally the IPv4 and IPv6 traffic for that application.

Special consideration should be provided to the QoS policies for tunneled traffic. QoS for ISATAP-tunneled traffic is somewhat limited. When ISATAP tunnels are used, the ingress classification of IPv6 packets cannot be made at the access layer, which is the recommended location for trusting or classifying ingress traffic. In the HME1 and SBM designs, the access layer has no IPv6 support. Tunnels are being used between the hosts in the access layer and either the core layer (HME1) or the SBM switches, and therefore ingress classification cannot be done.

QoS policies for IPv6 can be implemented after the decapsulation of the tunneled traffic, but this also presents a unique challenge. Tunneled IPv6 traffic cannot even be classified after it reaches the tunnel destination, because ingress marking cannot be done until the IPv6 traffic is decapsulated (ingress classification and marking are done on the physical interface and not the tunnel interface). Egress classification policies can be implemented on any IPv6 traffic now decapsulated and being forwarded by the switch. Trust, policing, and queuing policies can be implemented on upstream switches to properly deal with the IPv6 traffic.

Figure 9 illustrates the points where IPv6 QoS policies may be applied when using ISATAP in HME1. The dual-stack links shown have QoS policies that apply to both IPv4 and IPv6 and are not shown because those policies follow the Cisco campus QoS recommendations. Refer to [Additional References, page 64](#) for more information about the Cisco campus QoS documentation.


Figure 9 QoS Policy Implementation—HME1



1. In HME1, the first place to implement classification and marking is on the egress interfaces on the core layer switches. As was previously mentioned, the IPv6 packets have been tunneled from the hosts in the access layer to the core layer, and the IPv6 packets have not been “visible” in a decapsulated state until the core layer. Because QoS policies for classification and marking cannot be applied to the ISATAP tunnels on ingress, the first place to apply the policy is on egress.
2. The classified and marked IPv6 packets (see item 1) can now be examined by upstream switches (for example, aggregation layer switches), and the appropriate QoS policies can be applied on ingress. These policies may include trust (ingress), policing (ingress), and queuing (egress).

Figure 10 illustrates the points where IPv6 QoS policies may be applied in the SBM when ISATAP manually-configured tunnels are used.

Figure 10 QoS Policy Implementation—SBM (ISATAP and Manually-Configured Tunnels)



1. The SBM switches receive IPv6 packets coming from the ISATAP interfaces, which are now decapsulated, and can apply classification and marking policies on the egress manually-configured tunnel interfaces.
2. The upstream switches (aggregation layer and access layer) can now apply trust, policing, and queuing policies after the IPv6 packets leave the manually-configured tunnel interfaces in the aggregation layer.



Note

At the time of the writing of this document, the capability for egress per-user microflow policing of IPv6 packets on the Catalyst 6500 Supervisor 32/720 is not supported. When this capability is supported, classification and marking on ingress can be combined with per-user microflow egress policing on the same switch. In the SBM design, as of the release of this document, the policing of IPv6 packets must take place on ingress, and the ingress interface must not be a tunnel. For more information, see the PFC3 QoS documentation at the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/qos.htm>

The DSM model is not shown here because the same recommendations for implementing QoS policies for IPv4 should also apply to IPv6. Also, the HME2 QoS considerations are the same as those for Figure 10 and are not shown for the sake of brevity.

The key consideration as far as Modular QoS CLI (MQC) is concerned is the removal of the “ip” keyword in the QoS “match” and “set” statements. Modification in the QoS syntax to support IPv6 and IPv4 allows for a new configuration criteria, as shown in Table 5.

Table 5 New Configuration Criteria

IPv4-Only QoS Syntax	IPv4/IPv6 QoS Syntax
match ip dscp	match dscp

Table 5 *New Configuration Criteria (continued)*

match ip precedence	match precedence
set ip dscp	set dscp
set ip precedence	set precedence

There are QoS features that work for both IPv6 and IPv4, but require no modification to the CLI (for example, WRED, policing, and WRR).

The implementation section for each model does not go into great detail on QoS configuration in relation to the definition of classes for certain applications, the associated mapping of DSCP values, and the bandwidth and queuing recommendations. Cisco provides an extensive collection of QoS recommendations for the campus, which is available on CCO, as well as the Cisco Press book *End-to-End QoS Network Design*.

Refer to [Additional References, page 64](#) for more information about the Cisco campus QoS recommendations and Cisco Press books.

Security

Many of the common threats and attacks on existing IPv4 campus networks also apply to IPv6. Unauthorized access, spoofing, routing attacks, virus/worm, denial of service (DoS), and man-in-the-middle attacks are just a few of the threats to both IPv4 and IPv6.

With IPv6, many new threat possibilities do not apply at all or at least not in the same way as with IPv4. There are inherent differences in how IPv6 handles neighbor and router advertisement and discovery, headers, and even fragmentation. Based on all these variables and possibilities, the discussion of IPv6 security is a very involved topic in general, and detailed security recommendations and configurations are outside the scope of this document. There are numerous efforts both within Cisco and the industry to identify, understand, and resolve IPv6 security threats. This document points out some possible areas to address within the campus and gives basic examples of how to provide protection for IPv6 dual-stack and tunneled traffic.



Note

The examples given in this document are in no way meant to be recommendations or guidelines, but rather intended to challenge the reader to carefully analyze their own security policies as they apply to IPv6 in the campus.

The following are general security guidelines for network device protection that apply to all campus models:

- Make reconnaissance more difficult through proper address planning for campus switches:
 - Addressing of campus network devices (L2 and L3 switches) should be well-planned. Common recommendations are to devise an addressing plan so that the 64-bit interface-ID of the switch is a value that is random across all the devices. An example of a bad interface-ID for a switch is if VLAN 2 has an address of 2001:db8:cafe:2::1/64 and VLAN 3 has an address of 2001:db8:cafe:3::1/64, where ::1 is the interface-ID of the switch. This is easily guessed and allows for an attacker to quickly understand the common addressing for the campus infrastructure devices. Another choice is to randomize the interface-ID of all the devices in the campus. Using the VLAN 2 and VLAN 3 examples from above, a new address can be

constructed by using an address such as 2001:db8:cafe:2::a010:f1a1 for VLAN 2 and 2001:db8:cafe:3::c801:167a for VLAN 3, where “a010:f1a1” is the interface-ID of VLAN 2 for the switch.

The addressing consideration described above introduces real operational challenges. For the sake of easing operational management of the network devices and addressing, the reader should balance the security aspects of randomizing the interface-IDs with the ability to deploy and manage the devices via the randomized addresses.

- Control management access to the campus switches
 - All the campus switches for each model have configurations in place to help protect access to the switch for management purposes. All switches have loopback interfaces configured for management and routing purposes. The IPv6 address for the loopback interfaces uses the previously-mentioned addressing approach of avoiding well-known interface-ID values. In this example, the interface-ID is using “::A111:1010”.

```
interface Loopback0
ipv6 address 2001:DB8:CAFE:6507::A111:1010/128
no ipv6 redirects
```

To more tightly restrict access to a particular switch via IPv6, an ACL is used to permit access to the management interface (line vty) by way of the loopback interface. The permitted source network is from the enterprise IPv6 prefix. To make ACL generation more scalable for a wide range of network devices, the ACL definition can permit the entire enterprise prefix as the primary method for controlling management access to the device instead of filtering to a specific interface on the switch. The IPv6 prefix used in this enterprise site (example only) is 2001:db8:cafe::/48.

```
ipv6 access-list MGMT-IN
remark Permit MGMT only to Loopback0
permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::A111:1010
deny ipv6 any any log-input
!
line vty 0 4
session-timeout 3
access-class MGMT-IN-v4 in
password 7 08334D400E1C17
ipv6 access-class MGMT-IN in #Apply IPv6 ACL to restrict access
logging synchronous
login local
exec prompt timestamp
transport input ssh #Accept access to VTY via SSH
```

- The security requirements for running Simple Network Management Protocol (SNMP) are the same as with IPv4. If SNMP is needed, a choice should be made on the SNMP version and then access control and authentication/encryption.

In the campus models discussed in this document, SNMPv3 (AuthNoPriv) is used to provide polling capabilities for the Cisco NMS servers located in the data center. Following is an example of the SNMPv3 configuration used in the campus switches in this document:

```
snmp-server contact John Doe - ipv6rocks@cisco.com
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server user jdoe IPv6-ADMIN v3 auth md5 cisco1234
```

- Control access via HTTP—At the time of this writing, Cisco Catalyst switches do not support the use of IPv6 HTTP ACLs to control access to the switch. This is very important because switches that currently use “ip http access-class” ACLs for IPv4 do not have the same level of protection for IPv6. This means that subnets or users that were previously denied access via HTTP/HTTPS for IPv4 now have access to the switch via IPv6.

- IPv6 traffic policing—Traffic policing can be considered a QoS and/or security function. There may be existing requirements to police traffic either on an aggregate or per-user microflow basis. In the campus models discussed in this document, certain places are appropriate for implementing IPv6 policing, specifically per-user microflow policing:
 - DSM—The per-user microflow policing of IPv6 traffic is performed against ingress traffic on the Catalyst 6500 distribution layer switches (ideal).
 - HME1—The per-user microflow policing of IPv6 traffic is performed against ingress traffic (from the hosts in the campus access layer) on the Catalyst 6500 data center aggregation layer switches. This is not ideal; it is preferred to perform ingress microflow policing on the core layer switches, but in this model, the ingress policing cannot be applied to tunnel interfaces, so it has to be done at the next layer.
 - HME2—The per-user microflow policing of IPv6 traffic is performed against ingress traffic on the Catalyst 6500 distribution layer switches (ideal).
 - SBM—The per-user microflow policing of IPv6 traffic is a challenge in the specific SBM example discussed in this document. In the SBM, the service block switches are Catalyst 6500s and have PFC3 cards. The Catalyst 6500 with PFC3 supports ingress per-user microflow policing, but does not currently support IPv6 egress per-user microflow policing. In the SBM example in this document, IPv6 passes between the ISATAP and manually-configured tunnel interface on the service block switches. Because ingress policing cannot be applied to either ISATAP tunnels or manually-configured tunnel interfaces, there are no applicable locations to perform policing in the service block.

A basic example of implementing IPv6 per-user microflow policing follows. In this example, a downstream switch has been configured with a QoS policy to match IPv6 traffic and to set specific DSCP values based on one of the Cisco-recommended QoS policy configurations. The configuration for this particular switch (shown below) is configured to perform policing on a per-user flow basis (based on IPv6 source address in this example). Each flow is policed to 5 Mbps and is dropped if it exceeds the profile.

```

mls qos
!
class-map match-all POLICE-MARK
  match access-group name V6-POLICE-MARK
!
policy-map IPv6-ACCESS
  class POLICE-MARK
    police flow mask src-only 5000000 8000 conform-action transmit exceed-action drop
  class class-default
    set dscp default
!
ipv6 access-list V6-POLICE-MARK
  permit ipv6 any any
!
interface GigabitEthernet3/1
  mls qos trust dscp
  service-policy input IPv6-ACCESS

```

**Note**

This example is not based on the Cisco campus QoS recommendations but is shown as an informational illustration of how the configuration for per-user microflow policing might look.

More information on microflow policing can be found at the following URLs:

- <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/qos.htm>
- Enterprise QoS SRND—<http://www.cisco.com/go/srnd>

**Note**

At the time of this writing, the Catalyst 6500 Supervisor 32 and 720 do not support IPv6 per-user microflow policing and IPv6 multicast routing in hardware if enabled together. The supervisor supports policing in hardware or IPv6 multicast routing/forwarding in hardware, but not at the same time. If **ipv6 multicast-routing** command is already configured on the switch and an IPv6 per-user microflow policing policy is applied, the system returns a message indicating that the IPv6 packets are software switched. Inversely, if an IPv6 per-user microflow policing policy is applied to an interface on the switch and **ipv6 multicast-routing** command is enabled, the same message appears (see below).

Following is an example of the warning message:

```
006256: *Aug 31 08:23:22.426 mst:
%FM_EARL7-2-IPV6_PORT_QOS_MCAST_FLOWMASK_CONFLICT: IPv6 QoS Micro-flow
policing configuration on port GigabitEthernet3/1 conflicts for flowmask with IPv6 multicast
hardware forwarding on SVI interface Vlan2, IPv6 traffic on the SVI interface may be switched in
software
```

```
006257: *Aug 31 08:23:22.430 mst: %FM_EARL7-4-FEAT_QOS_FLOWMASK_CONFLICT:
Features configured on interface Vlan2 conflict for flowmask with QoS configuration on switch port
GigabitEthernet3/1, traffic may be switched in software
```

- **Control Plane Policing (CoPP)**—In the context of the campus models discussed in this document, CoPP applies only to the Catalyst 6500 Supervisor 32/720. CoPP protects the Multiswitch Feature Card (MSFC) by preventing DoS or unnecessary traffic from negatively impacting MSFC resources. Priority is given to important control plane/management traffic. The Catalyst 6500 with PFC3 supports CoPP for IPv6 traffic. The configuration of CoPP is based on a wide variety of factors and no single deployment recommendation can be made because the specifics of the policy are determined on a case-by-case basis.

More information on CoPP can be found at the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/dos.htm>

- **Control ingress traffic from the access layer**—Filtering which prefixes are allowed to source traffic. This is most commonly done on ingress on the VLAN interface of the distribution layer switches (DSM), but can also be applied to the ingress on the ISATAP tunnel interfaces in the HME1 or SBM. Controlling IPv6 traffic based on source prefix can help protect the network against basic spoofing.

An example of a basic ACL permitting only the IPv6 prefix for a VLAN is as follows:

```
ipv6 access-list VLAN2-v6-INGRESS
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
permit icmp 2001:DB8:CAFE:2::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
permit ipv6 2001:DB8:CAFE:2::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
interface Vlan2
  ipv6 traffic-filter VLAN2-v6-INGRESS in
```

**Note**

Cisco IOS IPv6 ACLs contain implicit permit entries for IPv6 neighbor discovery. If **deny ipv6 any any** is configured, the implicit neighbor discovery entries are overridden. It is important to remember that if a manually-configured catch-all deny statement is used for logging purposes, the following two permit entries must be added back in:

```
permit icmp any any nd-na
permit icmp any any nd-ns
```

In the VLAN2-v6-INGRESS example given above, a more permissive entry (**permit icmp FE80::/16 any**) is made to account for the neighbor discovery requirement as well as any other ICMPv6 services that are needed for link operation on VLAN2. There are RFCs, drafts, and IPv6 deployment books that specifically discuss the various ICMPv6 types that should or should not be blocked. Refer to [Additional References, page 64](#) for links to the IETF and Cisco Press book that discusses the filtering of ICMPv6 packets.

- Block the use of Microsoft Teredo—Teredo is used to provide IPv6 support to hosts that are located behind Network Address Translation (NAT) gateways. Teredo introduces several security threats that need to be thoroughly understood. Until well-defined security recommendations can be made for Teredo in campus networks, the reader may want to ensure that Teredo is disabled on Microsoft Windows XP SP2 and Vista be configured to disable Teredo. As a backup precaution, the reader may also want to consider configuring ACLs (which can be done at the access layer or further upstream, such as at the border routers) to block UDP port 3544 to prevent Teredo from establishing a tunnel outside the campus network. Information on Teredo can be found at the following URLs:

- <http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/teredo.mspx>
- <http://www.microsoft.com/technet/community/columns/cableguy/cg1005.mspx#EVF>

More information on IPv6 security can be found in [Additional References](#).

Multicast

IPv6 multicast is an important service for any enterprise network design, and IPv6 multicast requirements may cause the reader to re-consider the models discussed in this document. The most important issue to understand with IPv6 multicast and the various models is that IPv6 multicast is not supported over ISATAP. This is not a limitation of equipment or software, but rather a shortcoming of the ISATAP tunneling mechanism (RFC4214).

One of the most important factors to consider in IPv6 multicast deployment is to ensure that host/group control is handled properly in the access layer. Multicast Listener Discovery (MLD) in IPv6 is the equivalent to Internet Group Management Protocol (IGMP) in IPv4. Both are used for multicast group membership control. MLD Snooping is the feature that enables Layer 2 switches to control the distribution of multicast traffic only to the ports that have listeners. Without it, multicast traffic meant for only a single receiver (or group of receivers) is flooded to all ports on an access layer switch. In the access layer, it is important that the switches support MLD Snooping for MLD version 1 and/or version 2 (this is only applicable when running dual-stack at the access layer).

**Note**

At the time of the writing of this document, there were very few host implementations of MLDv2. Various Linux and BSD implementations support MLDv2, as does Microsoft Windows Vista. MLDv2 is important in PIM-SSM-based deployments. The use of MLDv2 with PIM-SSM is an excellent design combination for a wide variety of IPv6 multicast deployments. Some hosts do not yet support MLDv2. Cisco IOS provides a feature called SSM-mapping that map MLDv1 reports to MLDv2 reports to be used by PIM-SSM. More information can be found at the following URL:

http://www.cisco.com/en/US/partner/products/ps6350/products_configuration_guide_chapter09186a00801d6618.html#wp1290106

In this document, IPv6 multicast-enabled applications are supported in the DSM and HME2 models because no ISATAP configurations are used in either model. The multicast-enabled applications tested in this design are Windows Media Services and VideoLAN Media Client (VLC) using Embedded-RP and PIM-SSM groups. The multicast sources are running on Microsoft Windows Server 2003, Longhorn, and Red Hat 4.0 servers located in the data center access layer.

Several documents on CCO and within the industry discuss IPv6 multicast in detail. No other configuration notes are made in this document except for generic references to the commands to enable IPv6 multicast and requirements for Embedded-RP definition. For more information, see the Cisco IPv6 Multicast webpage at the following URL:

http://www.cisco.com/en/US/products/ps6594/products_ios_protocol_group_home.html.

Management

Management tools and instrumentation for IPv6 are under development and have a long way to go. Many of the traditional management tools used today support IPv6 as well. In this document, the only considerations for management of the campus network are related to basic management services (Telnet, SSH, and SNMP). All the IPv6-enabled devices in the various campus models discussed are manageable over IPv6 via the above-mentioned services except SNMP. At the time of this writing, the Catalyst switches discussed do not yet support SNMP over IPv6 transport. However, the management of IPv6-specific MIBs/Traps/Informs is supported on the Catalyst platforms using SNMP over IPv4 transport.

Another area of management that the reader must thoroughly research is that of address management. Anyone who analyzes IPv6 even at an elementary level understands the size and potential complexity of deploying and managing the IPv6 address structure. Deploying large hexadecimal addresses on many network devices should, at some point, be automated or at least more user-friendly than it is today. Several efforts are underway within the industry to provide recommendations and solutions to the address management issues. Cisco is in the forefront of this effort.

Today, one way to help with the deployment of address prefixes on a campus switch is through the use of the “general prefix” feature. The general prefix feature allows the customer to define a prefix or prefixes in the global configuration of the switch with a user-friendly name. That user-friendly name can be used on a per-interface basis to replace the usual IPv6 prefix definition on the interface. Following is an example of how to use the general prefix feature:

- Define the general prefix:


```
6k-agg-1(config)#ipv6 general-prefix ESE-DC-1 2001:DB8:CAFE::/48
```
- Configure the general prefix named “ESE-DC-1” on a per-interface basis:


```
6k-agg-1(config-if)#ipv6 address ESE-DC-1 ::10:0:0:F1A1:6500/64
```

- Verify that the general prefix was correctly assigned to the interface:

```
6k-agg-1#show ipv6 interface vlan 10
Vlan10 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::211:BCFF:FEC0:C800
Description: VLAN-SERVERFARM-WEB
Global unicast address(es):
  2001:DB8:CAFE:10::F1A1:6500, subnet is 2001:DB8:CAFE:10::/64
```



Note The use of the general prefix feature is useful where renumbering is required, because changing the general prefix value can renumber a router quickly.

More information on the general prefix feature can be found at the Cisco IOS IPv6 documentation page (see [Additional References, page 64](#)).

Cisco supports the management of IPv6-enabled network devices via a variety of network management products to include DNS, DHCPv6, device management, and monitoring; and also network management, troubleshooting, and reporting. For more information on the various Cisco Network Management solutions, see the following URL:

<http://www.cisco.com/en/US/products/sw/netmgtsw/index.html>.

Scalability and Performance

This document is not meant to analyze scalability and performance information for the various platforms tested. The discussion of scale and performance is more focused on general considerations when planning and deploying IPv6 in the campus versus a platform-specific view.

In general, the reader should understand the link, memory, and CPU utilization of the existing campus network. If any of these aspects are already stressed, adding IPv6 or any new technology, feature, or protocol into the design is a recipe for disaster. However, it is common to see in IPv6 implementations a change in traffic utilization ratios on the campus network links. As IPv6 is deployed, IPv4 traffic utilization is very often reduced as users leverage IPv6 as the transport for applications that were historically IPv4-only. There is an increase in overall network utilization that usually derives from control traffic for routing and also tunnel overhead when ISATAP or manually-configured tunnels are used.

Scalability and performance considerations for the DSM are as follows:

- Routed access design (access layer)—One of the primary scalability considerations is that of running two protocols on the access (routed access) or distribution layer switch. In the routed access or distribution layer, the switch must track both IPv4 and IPv6 neighbor information. Similar to Address Resolution Protocol (ARP) in IPv4, neighbor cache exists for IPv6. The primary consideration here is that with IPv4, there is usually a one-to-one mapping of IPv4 address-to-MAC address; but with IPv6, there can be several mappings for multiple IPv6 addresses that the host may have (for example, link-local, unique-local, and multiple global addresses) to a single MAC address in the neighbor cache of the switches. Following is an example of ARP and neighbor cache entries on a Catalyst 6500 located in the distribution layer for a host with the MAC address of “000d.6084.2c7a”.

ARP entry for host in the distribution layer:

```
Internet 10.120.2.200          2    000d.6084.2c7a  ARPA    Vlan2
```

IPv6 neighbor cache entry:

```
2001:DB8:CAFE:2:2891:1C0C:F52A:9DF1      4    000d.6084.2c7a  STALE V12
2001:DB8:CAFE:2:7DE5:E2B0:D4DF:97EC     16   000d.6084.2c7a  STALE V12
```

FE80::7DE5:E2B0:D4DF:97EC

16 000d.6084.2c7a STALE V12

The neighbor cache shows that there are three entries listed for the host. The first address is one of two global IPv6 addresses assigned (optional) and reflects the global IPv6 address generated by the use of IPv6 privacy extensions. The second address is another global IPv6 address (optional) that is assigned by stateless autoconfiguration (it can also be statically defined or assigned via DHCPv6), and the third address is the link-local address (mandatory) generated by the host. The number of entries can decrease to a minimum of one (link-local address) to a multitude of entries for a single host, depending on the address types used on the host.

It is very important to understand the neighbor table capabilities of the routed access and distribution layer platforms being used to ensure that the tables are not being filled during regular network operation. Additional testing is planned to understand whether recommendations should be made to adjust timers to time out entries faster, to rate limit neighbor advertisements, and to better protect the access layer switch against DoS from IPv6 neighbor discovery-based attacks.

Another consideration is with IPv6 multicast. As mentioned previously, it is important to ensure that MLD Snooping is supported at the access layer when IPv6 multicast is used to ensure that IPv6 multicast frames at Layer 2 are not flooded to all the ports.

- Distribution layer—In addition to the ARP/neighbor cache issues listed above, there are two other considerations for the distribution layer switches in the DSM:
 - IPv6 routing and forwarding must be performed in hardware.
 - It is imperative that the processing of ACL entries be performed in hardware. IPv6 ACLs in the distribution layer are primarily used for QoS (classification and marking of ingress packets from the access layer), for security (controlling DoS, snooping and unauthorized access for ingress traffic in the access layer), and for a combination of QoS and security to protect the control plane of the switch from attack.
- Core layer—The considerations for scale and performance are the same as with the distribution layer.

Scalability and performance considerations for the HME1 are as follows:

- Access layer—There are no real scale or performance considerations for the access layer when using the HME1. IPv6 is not supported in the access layer in the HME1, so there is not much to discuss. Link utilization is the only thing to consider because there may be an additional amount of traffic (tunneled IPv6 traffic) present on the links. As mentioned previously, however, as IPv6 is deployed there may be a replacement of link utilization ratios from IPv4 to IPv6 as users begin to use IPv6 for applications that were historically IPv4-only.
- Distribution layer—Same considerations as in the access layer.
- Core layer—There can be an impact on the core layer switches when using HME1. There can be hundreds or more ISATAP tunnels terminating on the core layer switches in the HME1. The reader should consult closely with partners and Cisco account teams to ensure that the existing core layer switches can handle the number of tunnels required in the design. If the core layer switches are not going to be able to support the number of tunnels coming from the access layer, it might be required to either plan to move to the DSM or use the SBM instead of HME1 so that dedicated switches can be used just for tunnel termination and management until DSM can be supported. Three important scale and performance factors for the core layer are as follows:
 - Control plane impact for the management of ISATAP tunnel interfaces. This can be an issue if there is a one-to-one mapping between the number of VLANs to the number of ISATAP tunnels. In large networks, this mapping results in a substantial number of tunnels that the CPU must track. The control plane management of virtual interfaces is done by the CPU.

- Control plane impact for the management of route tables associated with the prefixes associated with the ISATAP tunnels.
- Link utilization—There is an increase in link utilization coming from the distribution layer (tunneled traffic) and a possible increase in link utilization by adding IPv6 (now dual-stack) to the links from the core layer to the data center aggregation layers.

Scalability and performance considerations for the HME2 are as follows:

- Access layer—Same considerations as with the access layer in the DSM.
- Distribution layer—In the HME2, dual-stack is used for the access layer and manually-configured tunnels are used to traverse the core and terminate in the data center aggregation layer. The scale and performance considerations for the access layer are the same as with the DSM distribution layer. The considerations for manually-configured tunnels are similar to those for the core layer in the HME1. However, there should be tunnels only between the distribution pair and the total number of data center aggregation layer switches. In some cases, this is as few as two tunnels per distribution switch. In some cases, this can be hundreds of tunnels. In either case, if the Catalyst platform used supports IPv6 tunneling in hardware, even hundreds of tunnels do not cause performance or scale issues.
- Core layer—The core layer is IPv4-only in the HME2 and requires no specific scale or performance comments.

Scalability and performance considerations for the SBM are as follows:

- Access layer—The access layer is IPv4-only in the SBM and requires no specific scale or performance considerations.
- Distribution layer—The distribution layer is IPv4-only in the SBM and requires no specific scale or performance considerations.
- Core layer—The core layer is IPv4-only in the SBM and requires no specific scale or performance considerations.
- Service block—Most of the considerations found in the core layer of HME1 apply to the service block switches. The one difference is that the service block is terminating both ISATAP and manually-configured tunnels on the same switch pair. The advantage with the SBM is that the switch pair is dedicated for tunnel termination and can have additional switches added to the service block to account for more tunnels, and therefore can allow for a larger tunnel-based deployment. Adding more switches for scale is difficult to do in a core layer (HME1) because of the central role the core has in connecting the various network blocks (access, data center, WAN, and so on).

Dual-Stack Model—Implementation


This section is focused on the configuration of the DSM. The configurations are divided into specific areas such as VLAN, routing, and HA configuration. Many of these configurations such as VLANs and physical interfaces are not specific to IPv6. VLAN configurations for the DSM are the same for IPv4 and IPv6, but are shown for completeness. An example configuration is shown for only two switches (generally the pair in the same layer or a pair connecting to each other), and only for the section being discussed; for example, routing or HA. The full configuration for each switch in the campus network can be found in [Appendix—Configuration Listings, page 66](#).

All commands that are applicable to the section covered are in **BOLD**.

Network Topology


The following diagrams are used as a reference for all DSM configuration examples. [Figure 11](#) shows the physical port layout that is used for the DSM.

Figure 11 DSM Network Topology—Physical Ports



[Figure 12](#) shows the IPv6 addressing plan for the DSM environment. To keep the diagram as simple to read as possible, the /48 prefix portion of the network is deleted. The IPv6 /48 prefix used in all the models in this paper is “2001:db8:cafe::/48”.

Figure 12 DSM Network Topology—IPv6 Addressing



In addition to the physical interfaces, IPv6 addresses are assigned to loopback and VLAN interfaces. Table 6 shows the switch, interface, and IPv6 address for the interface.

Table 6 Switch, Interface, and IPv6 Addresses

Switch	Interface	IPv6 address
3750-acc-1	VLAN2	2001:db8:cafe:2::cac1:3750/64
3750-acc-2	VLAN3	2001:db8:cafe:3::cac2:3750/64
6k-dist-1	Loopback0	2001:db8:cafe:6507::a111:1010/128
	VLAN2	2001:db8:cafe:2::a111:1010/64
	VLAN3	2001:db8:cafe:3::a111:1010/64
6k-dist-2	Loopback0	2001:db8:cafe:6507::b222:2020/128
	VLAN2	2001:db8:cafe:2::b222:2020/64
	VLAN3	2001:db8:cafe:3::b222:2020/64
6k-core-1	Loopback0	2001:db8:cafe:6507::c333:3030/128
6k-core-2	Loopback0	2001:db8:cafe:6507::d444:4040/128
6k-agg-1	Loopback0	2001:db8:cafe:6507::e555:5050/128
	VLAN10	2001:db8:cafe:10::e555:5050/64

Table 6 Switch, Interface, and IPv6 Addresses (continued)

6k-agg-2	Loopback0	2001:db8:cafe:6507::f666:6060/128
	VLAN10	2001:db8:cafe:10::f666:6060/64
6k-acc-1	VLAN10	2001:db8:cafe:10::dca1:6506/64
6k-acc-2	VLAN10	2001:db8:cafe:10::dca2:6506/64

Physical/VLAN Configuration

Physical p2p links are configured in much the same way as IPv4. The following example is the p2p interface configuration for the link between 6k-dist-1 and 6k-core-1.

- 6k-dist-1:

```

ipv6 unicast-routing           #Globally enable IPv6 unicast routing
ip cef distributed             #Ensure IP CEF is enabled (req. for
                               #IPv6 CEF to run).
ipv6 cef distributed           #Globally enable IPv6 CEF.
!
interface GigabitEthernet4/1
  description to 6k-core-1
  dampening
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7000::A111:1010/64 #Assign IPv6 address
  no ipv6 redirects            #Disable IPv6 redirects
  ipv6 nd suppress-ra          #Disable RAs on this interface
  ipv6 cef                     #Enable IPv6 CEF for this intf.

```

- 6k-core-1:

```

ipv6 unicast-routing
ip cef distributed
ipv6 cef distributed
!
interface GigabitEthernet2/4
  description to 6k-dist-1
  dampening
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7000::C333:3030/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef

```

Although not required, it is a good practice to disable the sending of RAs on p2p links. The RAs are not needed on p2p links that are statically defined. It is also important to note that depending on the platform and version of code, it may not be necessary to enable **ipv6 cef** on a per-interface basis. It is shown here because it is used as a safeguard to ensure IPv6 CEF is enabled on the interface. However, newer versions of code do this automatically when IPv6 unicast routing is enabled globally and IPv6 is enabled on the interface.

On the Catalyst 3750 and 3560 switches, it is required to enable the correct Switch Database Management (SDM) template to allow the ternary content addressable memory (TCAM) to be used for different purposes. The 3750-acc-1 and 3750-acc-2 have been configured with the “dual-ipv4-and-ipv6” SDM template using the **sdm prefer dual-ipv4-and-ipv6 default** command. For more information about the **sdm prefer** command and associated templates, refer to the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12225see/scg/swsdm.htm#>

The access layer uses a single VLAN per switch; voice VLANs are not discussed. The VLANs do not span access layer switches and are terminated at the distribution layer. The following example is of the 3750-acc-1 and 6k-dist-1 VLAN2 configuration.

- 3750-acc-1:

```

vtp domain ese-dc
vtp mode transparent
!
!
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree portfast bpduguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
vlan 2
  name ACCESS-DATA-2
!
interface GigabitEthernet1/0/25
  description TRUNK TO 6k-dist-1
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 2
  switchport mode trunk
  switchport nonegotiate
  load-interval 30
!
interface Vlan2
  ipv6 address 2001:DB8:CAFE:2::CAC1:3750/64
  no ipv6 redirects

```

#VTP and STP configurations
#shown for completeness, but not
#specific to IPv6

#VLAN2 - Data VLAN for 3750-acc-1

#Physical intf. to 6k-dist-1

#VLAN2 with IPv6 address used for mgmt.

- 6k-dist-1:

```

vtp domain ese-dc
vtp mode transparent
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 2-3 priority 24576
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
vlan 2
  name ACCESS-DATA-2
!
vlan 3
  name ACCESS-DATA-3
!
interface GigabitEthernet3/1
  description to 3750-acc-1
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 2
  switchport mode trunk
  switchport nonegotiate
  no ip address

```

#6k-dist-1 is the STP root for
#VLAN2,3

#VLAN2 defined for 3750-acc-1

#VLAN3 defined for 3750-acc-2

#Physical intf. to 3750-acc-1

```

load-interval 30
spanning-tree guard root
!
interface Vlan2                                #VLAN2 intf is VLAN termination
                                                #point for trunked VLAN from 3750-acc-1
description ACCESS-DATA-2
ipv6 address 2001:DB8:CAFE:2::A111:1010/64    #IPv6 address and prefix used for
                                                #Stateless autoconfiguration for VLAN2

no ipv6 redirects

```

Although stacks are not used in any of the models discussed here, they are commonly used on the Catalyst 3750 and 3560 in the access layer. IPv6 is supported in much the same way as IPv4 when using switch stacks. For more information on IPv6 with switch stacks, see the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12225see/scg/swipv6.htm#wp1090623>

Routing Configuration

As previously mentioned, the routing for the DSM is set up using EIGRP for IPv4 and OSPFv3 for IPv6. The OSPFv3 configuration follows the recommended Cisco campus designs as much as possible. Certain features are available for OSPFv2 that are still being integrated in OSPFv3. These features are predominately focused on fast convergence and authentication. The configuration for OSPFv3 is shown for the 6k-dist-1 and 6k-core-1 switches.

- 6k-dist-1:

```

interface Loopback0
ip address 10.122.10.9 255.255.255.255        #Address used for RID on OSPFv3
ipv6 address 2001:DB8:CAFE:6507::A111:1010/128
ipv6 ospf 1 area 0                            #Part of area 0
!
interface TenGigabitEthernet1/1
description to 6k-dist-2
ipv6 address 2001:DB8:CAFE:7004::A111:1010/64
ipv6 cef
ipv6 ospf network point-to-point            #Defining network type as p2p
ipv6 ospf hello-interval 1                  #Lower hello/dead intervals as a fail-safe
                                                #for lower convergence times. p2p links do
                                                #not use hello/dead as the primary
                                                #detector of link/node failures.

ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2
!
interface GigabitEthernet4/1
description to 6k-core-1
ipv6 address 2001:DB8:CAFE:7000::A111:1010/64
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
!
interface GigabitEthernet4/2
description to 6k-core-2
ipv6 address 2001:DB8:CAFE:7001::A111:1010/64
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
!
interface Vlan2

```

```

description ACCESS-DATA-2
ipv6 address 2001:DB8:CAFE:2::1/64 anycast      #One way to provide gateway
                                                #redundancy for the mgmt interface
                                                #on the 3750-dist-1. This anycast
                                                #address is used on both
                                                #distribution layer switches

ipv6 address 2001:DB8:CAFE:2::A111:1010/64
ipv6 nd reachable-time 5000                    #Lower the NUD-reachable time to provide
                                                #basic first-hop redundancy for the hosts
                                                #in VLAN2. This command is used on both
                                                #distribution layer switches' VLAN
                                                #interfaces.

ipv6 cef
ipv6 ospf 1 area 2                             #Per the campus best practices - Access
                                                #VLAN is not in area 0

!
ipv6 router ospf 1                             #RID using Loopback0
router-id 10.122.10.9
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 range 2001:DB8:CAFE:2::/64 cost 10      #Summarize the /64 prefix for
                                                #VLAN2/3 into area 0

area 2 range 2001:DB8:CAFE:3::/64 cost 10
passive-interface Vlan2                       #Do not establish adjacency over
                                                #VLAN2/3 with 6k-dist-2

passive-interface Vlan3
passive-interface Loopback0
timers spf 1 5                                #Lower the SPF throttling timer for
                                                #convergence time improvement. In campus
                                                #networks these values have been tested as
                                                #low as 1 and 1 without issue. Adjusting
                                                #these values should be well understood
                                                #and consistent #between all switches in
                                                #the campus

```

- 6k-core-1:

```

interface Loopback0
ip address 10.122.10.3 255.255.255.255
ipv6 address 2001:DB8:CAFE:6507::C333:3030/128
ipv6 ospf 1 area 0
!
interface GigabitEthernet2/1
description to 6k-agg-1
ipv6 address 2001:DB8:CAFE:7005::C333:3030/64
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
!
interface GigabitEthernet2/4
description to 6k-dist-1
ipv6 address 2001:DB8:CAFE:7000::C333:3030/64
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
!
ipv6 router ospf 1
router-id 10.122.10.3
log-adjacency-changes
auto-cost reference-bandwidth 10000

```

```
passive-interface Loopback0
timers spf 1 5
```

It is important to read and understand the implications of modifying various IGP timers. The campus network should be designed to converge as fast as possible. The campus network is also capable of running much more tightly-tuned IGP timers than in a branch or WAN environment. The routing configurations shown are based on the Cisco campus recommendations. The reader should understand the context of each command and the timer value selection before pursuing the deployment in a live network. Refer to [Additional References, page 64](#) for links to the Cisco campus design best practice documents.

High-Availability Configuration

The HA design in the DSM consists of running two of each switches (applicable in the distribution, core, and data center aggregation layers) and ensuring that the IPv4 and IPv6 routing configurations are tuned and completely fault-tolerant. The reader should investigate the impact of running single-chassis dual-supervisor deployments if this model is a requirement. Some HA features such as Non-Stop Forwarding (NSF) and Stateful Switchover (SSO) may not support IPv6 across all Catalyst platforms. Check for the availability of NSF/SSO with IPv6 for the platforms and software of interest.

QoS Configuration

The QoS configurations for the DSM are the same as those for IPv4. The policies for classification, marking, queuing, and policing vary greatly based on the customer service requirements. The types of queuing and number of queues supported also vary between platform-to-platform and line card-to-line card. The basic configuration for the 6k-dist-1 is shown and is meant to be for reference *only*. For the sake of brevity, not all interfaces are shown.

- 6k-dist-1

```
mls qos
!
interface TenGigabitEthernet1/1
description to 6k-dist-2
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
!
interface GigabitEthernet3/1
description to 3750-acc-1
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
```

```

wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
!
interface GigabitEthernet4/1
description to 6k-core-1
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp

```

Multicast Configuration

IPv6 multicast is fully supported in the DSM. Although IPv6 multicast design is outside the scope of this document, configurations are shown for IPv6 multicast on the 3750-acc-1, 6k-dist-1, 6k-core-1, and 6k-agg-1 (acting as RP) switches. Most of the configuration examples are trivial, but are shown from the access layer to the aggregation layer for operational consistency.

- 3750-acc-1

```

ipv6 mld snooping #Globally enable MLD snooping (see note below)

```

- 6k-dist-1

```

ipv6 multicast-routing #Globally enable IPv6 multicast routing

```

- 6k-core-1

```

ipv6 multicast-routing

```

- 6k-agg-1

```

ipv6 multicast-routing
!
ipv6 pim rp-address 2001:DB8:CAFE:10::e555:5050 ERP #Embedded-RP is being used
#which requires the local
#definition of the RP.
#This command line states
#that this switch (v6 address
#on VLAN10) is the RP for any
#group permitted in the ACL
#ERP
!
ipv6 access-list ERP #ACL to permit Embedded-RP group range
#FF7E:140:2001:DB8:CAFE:10::/96
permit ipv6 any FF7E:140:2001:DB8:CAFE:10::/96 log-input

```

The first thing to understand is the lack of CLI input required to enable IPv6 multicast when using PIM-SSM or Embedded-RP. If PIM-SSM is used exclusively, it is only required to enable “ipv6 multicast-routing” globally, which automatically enables PIM on all IPv6-enabled interfaces. This is a dramatic difference from what is required with IPv4 multicast.

In the example above, the Layer 2 switch (3750-acc-1) needs to have IPv6 multicast awareness to control the distribution of multicast traffic only on ports that are actively listening. This is accomplished by enabling MLD Snooping. With MLD Snooping enabled on the 3750-acc-1 switch and with IPv6 multicast routing enabled on the 6k-dist-1 (and 6k-dist-2) switch, it can be seen that the 3750-acc-1 can see both distribution layer switches as locally attached multicast routers.

```
3750-acc-1#sh ipv6 mld snooping mrouter
Vlan      ports
-----
 2      Gil/0/25(dynamic), Gil/0/26(dynamic)
```

When a group is active on the access layer switch, information about the group can be displayed:

```
3750-acc-1#show ipv6 mld snooping address
Vlan      Group          Type      Version      Port List
-----
 2          FF35::1111     mld       v2           Gil/0/25, Gil/0/26
```

On the 6k-dist-1, information about PIM, multicast route, RPF, and groups can be viewed in much the same way as with IPv4. Following is the output of an active group using PIM-SSM (FF35::1111). This stream is coming in from the 6k-core-1 switch and going out the VLAN2 (3750-acc-1) interface:

```
6k-dist-1#show ipv6 mroute                                     #"show ipv6 pim topology" can also be used
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(2001:DB8:CAFE:11:2E0:81FF:FE2C:9332, FF35::1111), 19:58:58/never, flags: sTI
  Incoming interface: GigabitEthernet4/1
  RPF nbr: FE80::215:C7FF:FE24:7440
  Immediate Outgoing interface list:
    Vlan2, Forward, 19:58:58/never
```

Routed Access Configuration

The primary change to the campus implementation when using the routed access design applies to the access and distribution layer configurations. With the routed access design, the access layer performs routing where the previous (traditional) design had the access layer as a Layer 2-only component and the first Layer 3 component was in the distribution layer. This guide is not meant to discuss the advantages and disadvantages of the routed access design. However, the failover performance improvements realized along with the important fact that spanning tree is not an active component, make this design attractive to many customers. Because of customer demand, performance, and operational advantages with the routed access design, this paper discusses implementing IPv6 in this design.

Extending the DSM to now be a routed access design is quite easy. The removal of dependency on a redundant first-hop protocol is also a major improvement in the access layer. Basically, the access layer switches enable IPv6 routing and change the trunk links to routed links, and the distribution layer switches remove the trunks and VLANs for the access layer.

Figure 13 shows the updated DSM topology that has the routed access component included. Because nothing has changed upstream of the distribution layer, this diagram includes only the changed layers, which are the access and distribution layers.

Figure 13 DSM Topology—Routed Access Design




Figure 13 shows that the links between the access layer and distribution layer are now routed links instead of trunked Layer 2 links. IPv6 addressing and routing is configured on the new links, and the hosts in the VLANs use the IPv6 address of the VLAN interface on the access switch as the default gateway. The following configuration example shows the relevant configurations for the 3750-acc-1 and 6k-dist-1 switches.

- 3750-acc-1

```

ipv6 unicast-routing                               #Globally enable IPv6 unicast routing
!
interface GigabitEthernet1/0/25
description To 6k-dist-1
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
ipv6 address 2001:DB8:CAFE:700A::CAC1:3750/64      #Link is now a routed link
ipv6 nd suppress-ra
ipv6 ospf network point-to-point                  #OSPFv3 is configured in order to
                                                  #establish a peer relationship with
                                                  #6k-dist-1

ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2                               #Link is in area 2
no ipv6 redirects
mls qos trust dscp
!
interface Vlan2
load-interval 30
ipv6 address 2001:DB8:CAFE:2::CAC1:3750/64       #VLAN2 on this switch becomes the
                                                  #first layer 3 point for the hosts
                                                  #in VLAN2 - the link-local address

```

```

#on VLAN 2 will be the default
#gateway for the hosts
#VLAN2 is in area 2

ipv6 ospf 1 area 2
no ipv6 redirects
!
ipv6 router ospf 1
router-id 10.120.2.1
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 stub no-summary

passive-interface Vlan2
timers spf 1 5

```

#Per the Routed Access Design guide - the
#area (area 2) for the access layer
#prefix is a totally stubby area

- 6k-dist-1

```

interface GigabitEthernet3/1
description to 3750-acc-1
dampening
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:700A::A111:1010/64 #Link is now a routed link
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2 #Link is in area 2
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
!
ipv6 router ospf 1
router-id 10.122.10.9
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 stub no-summary

area 2 range 2001:DB8:CAFE:2::/64 cost 10 #Per the Routed Access Design guide - the
#area (area 2) for the access layer
#prefix is a totally stubby area
#Send a summary into area 0 for
#prefix "2" in area 2

area 2 range 2001:DB8:CAFE:3::/64 cost 10
area 2 range 2001:DB8:CAFE:7004::/64 cost 10
area 2 range 2001:DB8:CAFE:700A::/64 cost 10
area 2 range 2001:DB8:CAFE:700B::/64 cost 10
passive-interface Loopback0
timers spf 1 5

```

The “show ipv6 route” output for the 3750-acc-1 shows a default route coming from the two distribution layer switches (the default is injected by the upstream switches where the Internet edge connects to the core layer):

```
3750-acc-1#show ipv6 route
IPv6 Routing Table - 13 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OI  ::/0 [110/11]
     via FE80::213:5FFF:FE1F:F840, GigabitEthernet1/0/26          #6k-dist-2
     via FE80::215:C7FF:FE25:9580, GigabitEthernet1/0/25          #6k-dist-1
C   2001:DB8:CAFE:2::/64 [0/0]
     via ::, Vlan2
L   2001:DB8:CAFE:2::CAC1:3750/128 [0/0]
     via ::, Vlan2
```

**Note**

This output is only a snippet.

The other configuration change that is made in the DSM when using the routed access design is with IPv6 multicast. Now that the access layer switch is actually routing, the switch needs to be configured to support PIM of whatever variety is used in the rest of the network. The previous multicast configurations shown for the 6k-dist-1 would work for generic PIM-SSM or PIM-SM with Embedded-RP. It is important to note that the customer needs to validate which access layer platforms have IPv6 multicast routing support and in which code version.

Additional information on the Cisco routed access design can be found at the following URLs:

- Campus SRNDs—Routed Access and High Availability—
http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html#anchor2
- Routed Access Q&A—
<http://www.cisco.com/en/US/netsol/ns340/ns394/ns147/ns17/netqa0900aec8045965a.html>
- Routing in the Wiring Closet white paper—
http://www.cisco.com/en/US/netsol/ns340/ns394/ns147/ns17/networking_solutions_white_paper0900aec804c6e73.shtml


Hybrid Model—Example 1 Implementation

Most of the campus network in the HME1 is IPv4-only. The IPv6 part of the campus network begins in the core layer. This section shows the core layer configuration as well as the basic ISATAP configuration on the host. As mentioned previously, the HME1 uses dual-stack from the core layer into the data center. Those configurations are not relevant to the HME1 model because the configurations are the same as those in the DSM. As with the DSM implementation section, configuration snippets for each aspect of the deployment are shown in this section. The full configurations can be found in [Appendix—Configuration Listings, page 66](#).

Network Topology

One difference in the HME1 topology is that the distribution layer is using a pair of Catalyst 3750 switches instead of the Catalyst 6500. This is not because of any particular issue or recommendation, but just the way the test lab is configured. Figure 14 shows the network topology for the HME1.

Figure 14 HME1 Network Topology



The topology is focused on the IPv4 addressing scheme in the access layer (used by the host to establish the ISATAP tunnel), core layer (used as the termination point by the host for ISATAP), and also the IPv6 addressing used in the core layer for both the p2p link and the ISATAP tunnel prefix. The configuration shows that the ISATAP access high availability is accomplished by using redundantly configured loopback interfaces that share the same IPv4 address between both core switches. To maintain prefix consistency for the ISATAP hosts in the access layer, the same prefix is used on both the primary and backup ISATAP tunnels.

Physical Configuration

The configurations for both core layer switches are shown and include only the distribution and core layer-facing interfaces. Configurations for the IPv4 portion of the distribution and access layer are based on existing campus design best practices and are not discussed in this section. Full configurations that include the IPv4 settings are available in [Appendix—Configuration Listings, page 66](#).

- 6k-core-1

```
interface GigabitEthernet1/1
  description to 3750-dist-1
  dampening
  ip address 10.122.0.41 255.255.255.252
  no ip redirects
  no ip proxy-arp
```

```

ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
interface GigabitEthernet1/2
  description to 3750-dist-2
  dampening
  ip address 10.122.0.45 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  load-interval 30
  carrier-delay msec 0
  mls qos trust dscp
!
interface GigabitEthernet2/3
  description to 6k-core-2
  dampening
  ip address 10.122.0.21 255.255.255.252
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE::c333:3030/64          #p2p link between core switches
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  mls qos trust dscp

```

- 6k-core-2

```

interface GigabitEthernet1/1
  description to 3750-dist-1
  dampening
  ip address 10.122.0.49 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  load-interval 30
  carrier-delay msec 0
  mls qos trust dscp
!
interface GigabitEthernet1/2
  description to 3750-dist-2
  dampening
  ip address 10.122.0.53 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  load-interval 30
  carrier-delay msec 0
  mls qos trust dscp
!
interface GigabitEthernet2/3
  description to 6k-core-1
  dampening
  ip address 10.122.0.22 255.255.255.252
  ip hello-interval eigrp 10 1

```

```

ip hold-time eigrp 10 3
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE::d444:4040/64
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
mls qos trust dscp

```

Tunnel Configuration

The ISATAP configuration at the tunnel level is relatively straightforward, but the potentially confusing part relates to the high availability design for the ISATAP tunnels. The basic configuration of ISATAP on a host consists of enabling IPv6 and configuring the ISATAP router name or IPv4 address. By default, Microsoft Windows XP and Vista perform a DNS query of “isatap.domain.com”, where “domain.com” is the local domain name. If a DNS “A” record for “isatap” has been configured, the host begins to establish an ISATAP tunnel to that address. This default configuration works fine until something happens to the ISATAP router or the path to that router. All the configurations discussed in this paper include the ability to provide fault tolerance of IPv6 services as optimally as possible.

Providing high availability for ISATAP is crucial in the HME1 environment. Several methods provide redundancy of the ISATAP routers. The method discussed in this guide uses the two core layer switches to provide very fast failover of the ISATAP tunnels. The other method commonly used relies on DNS. Although the DNS method is faster to implement, it is also the most limiting in the overall IPv6 campus design and is the slowest for failover.

It is important to ensure that the tunnel destination (from the host point of view) is redundant across both core switches, and to ensure that both IPv4 and IPv6 routing is configured properly.



Note

At the time of this writing, the Catalyst 6500 using a single loopback for multiple tunnel source commands processes the tunneled traffic in software. The system generates a warning message to inform the user of this. The HME1 design does not suffer from this issue because each tunnel is using a separate loopback for control and scale purposes.

One question commonly asked is: Should I have deterministic routing from the distribution layer (IPv4) to one ISATAP router or is there any value in load balancing? The following considerations apply:

- The only host operating system that supports the outbound load balancing of ISATAP tunnels is Microsoft Windows Vista and it needs to be configured.
- Using customer deployments and detailed testing as a baseline, it has been found that there are few to no benefits in load balancing ISATAP hosts to the ISATAP routers. Testing shows that load balancing from the host side when using redundant IPv6 prefixes for the ISATAP tunnels causes return routability issues. The core layer switches in this example are capable of taking the load for all the ISATAP tunnels in this design. If the primary core layer switch fails, the secondary can take all the tunnels with no issue. Load balancing in this design provides no improvement in performance, load, or availability and further complicates the management for the operator because troubleshooting the flow of traffic for ISATAP is made even more difficult. Implementing a design that is deterministic for ISATAP eases the burden of traffic management and troubleshooting as well as eliminating the return routability issues.

To maintain low convergence times for ISATAP tunnels when a core layer switch fails, it is important to provide redundant and duplicated tunnel addresses across both core switches. When this is done, only one ISATAP router address or name is needed on the host, and DNS round-robin is not required. The following steps describe this process:

1. Both core layer switches are configured with the same loopback address (for example, 10.122.10.102). Loopback interfaces are used for their stable state and are perfect for tunnel termination.
2. Both core layer switches are configured with a single ISATAP tunnel that uses the loopback as a source (for example, Loopback2—10.122.10.102). The ISATAP IPv6 prefix is the same on both switches, so that no matter on which switch the host is terminated, it uses the same prefix for connectivity.
3. Both core layer switches are configured to advertise the loopback address via the IPv4 IGP. The primary switch (6k-core-1) uses default IGP metrics for the loopback address. The secondary switch (6k-core-2) alters the IGP metric (delay value on EIGRP) to make the loopback address on this switch to be less preferred. Again, Cisco recommends having a deterministic flow for the tunnels because load balancing between the tunnels using the same prefix is not desirable.
4. Both core layer switches are configured to advertise the ISATAP IPv6 prefix via the IPv6 IGP. The primary switch (6k-core-1) uses the default IGP metrics for the IPv6 prefix on the ISATAP tunnel. The secondary switch (6k-core-2) alters the IGP metric (cost value on OSPFv3) to make the ISATAP prefix on this switch to be less preferred. This is optional. It is used in this document because a deterministic flow for both IPv4 (see step 3) and IPv6 is desired.
5. The host is configured with a manually-defined ISATAP router address or name (which correlates to a DNS “A” record).

To keep the ISATAP tunnels, HA, and routing configurations simple to understand, they are shown together.

For the sake of simplicity, the only configuration shown is that of the tunnels for VLAN2. The tunnels for VLAN3 are the same except for addressing specifics.

The following configurations illustrate the five steps described above.

- 6k-core-1

```
interface Loopback2
  description Tunnel source for ISATAP-VLAN2
  ip address 10.122.10.102 255.255.255.255           #Address that will be used as the
                                                    #ISATAP tunnel2 source

!
interface Tunnel2
  description ISATAP VLAN2
  no ip address
  no ip redirects
  ipv6 address 2001:DB8:CAFE:2::/64 eui-64         #Tunnel prefix used for ISATAP
                                                    #hosts connecting to this tunnel.
                                                    #Interface-ID address for this
                                                    #switch will be generated using
                                                    #EUI-64
  no ipv6 nd suppress-ra                          #Tunnels interfaces disable the
                                                    #sending of RA's. This command
                                                    #re-enables RA's on this
                                                    #interface.

  ipv6 cef
  ipv6 ospf 1 area 2                               #Just like the VLAN in the DSM,
                                                    #this interface is not part of
                                                    #area 0
  tunnel source Loopback2                         #Tunnel2 uses loopback2 as the
                                                    #source
```

```

tunnel mode ipv6ip isatap                                #Define the tunnel as ISATAP
!
router eigrp 10
  passive-interface Loopback0
  passive-interface Loopback1
  passive-interface Loopback2
  passive-interface Loopback3
  network 10.0.0.0                                       #Covers Loopback2 interface and ensures
                                                         #that the 10.122.10.102 address is
                                                         #advertised to the rest of the network

  no auto-summary
  eigrp router-id 10.122.10.9
!
ipv6 router ospf 1
  router-id 10.122.10.9
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  area 2 range 2001:DB8:CAFE:2::/64 cost 10             #Advertise a summary for the prefix on
                                                         #Tunnel2 - just like a VLAN prefix
                                                         #would be sent in the DSM

  area 2 range 2001:DB8:CAFE:3::/64 cost 10
  passive-interface Loopback0
  passive-interface Loopback2
  passive-interface Loopback3
  passive-interface Tunnel2
  passive-interface Tunnel3
  timers spf 1 5

```

- 6k-core-2

```

interface Loopback2
  description Tunnel source for ISATAP-VLAN2
  ip address 10.122.10.102 255.255.255.255
  delay 1000                                             #Delay adjusted for EIGRP (IPv4)
                                                         #in order to adjust preference
                                                         #for the 10.122.10.102 host
                                                         #route. This ensures that
                                                         #6k-core-2 is SECONDARY to 6k-core-1
!
interface Tunnel2
  description ISATAP VLAN2
  no ip address
  no ip redirects
  ipv6 address 2001:DB8:CAFE:2::/64 eui-64
  no ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf 1 area 2
  tunnel source Loopback2
  tunnel mode ipv6ip isatap
!
router eigrp 10
  passive-interface Loopback0
  passive-interface Loopback1
  passive-interface Loopback2
  passive-interface Loopback3
  network 10.0.0.0
  no auto-summary
  eigrp router-id 10.122.10.10
!
ipv6 router ospf 1
  router-id 10.122.10.10
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  area 2 range 2001:DB8:CAFE:2::/64 cost 20             #Cost for prefix adjusted so that

```

```

area 2 range 2001:DB8:CAFE:3::/64 cost 20
passive-interface Loopback0
passive-interface Loopback2
passive-interface Loopback3

passive-interface Tunnel2
passive-interface Tunnel3
timers spf 1 5

```


```

#the route from 6k-core-2 is not
#preferred or equal to 6k-core-1
#Not required.

```

Figure 15 shows the IPv4 routing view from the distribution layer switches to the ISATAP tunnels interfaces (loopbacks on core switches). Loopback2 on 6k-core-1 is set as the primary ISATAP router address for the host. As shown in the previous IPv4 IGP configuration, 6k-core-2 is configured to have the host route of 10.122.10.102 have a higher delay and therefore is not preferred. When a packet arrives from the host in VLAN2 for the ISATAP router (10.122.10.102), a lookup is performed in the distribution layer switch for 10.122.10.102 and the next hop for that address is 6k-core-1.

Figure 15 HME1—Preferred Route for 6k-core-1



The routing table 10.122.10.102 on the distribution layer switches is as follows:

- 3750-dist-1 (route output shortened for brevity)

```

3750-dist-1#show ip route | b 10.122.10.102/32
D      10.122.10.102/32
[90/130816] via 10.122.0.41, 00:09:23, GigabitEthernet1/0/27

```

#3750-dist-1
#has only one
#route for
#10.122.10.102
#which is via
#10.122.0.41
#6k-core-1)

- 3750-dist-2

```

3750-dist-2#show ip route | b 10.122.10.102/32
D      10.122.10.102/32
[90/130816] via 10.122.0.45, 00:10:03, GigabitEthernet1/0/27


```

#3750-dist-2
#has only one
#route for
#10.122.10.102
#which is via
#10.122.0.45

#6k-core-1)

Figure 16 shows that 6k-core-1 has failed and therefore the route to loopback2 (10.122.10.102) is no longer available. When the 6k-core-1 route is removed, the new route for 10.122.10.102 is used and packets are then forwarded to 6k-core-2.

Figure 16 HME1—Preferred Route for 6k-core-2 After Failure of 6k-core-1



The updated routing table entry for 10.122.10.102 on the distribution layer switches is as follows:

- 3750-dist-1 (route output shortened for brevity)

```
3750-dist-1#show ip route | b 10.122.10.102/32
D      10.122.10.102/32
       [90/258816] via 10.122.0.49, 00:00:08, GigabitEthernet1/0/28
```

- 3750-dist-2

```
3750-dist-2#show ip route | b 10.122.10.102/32
D      10.122.10.102/32
       [90/258816] via 10.122.0.53, 00:00:08, GigabitEthernet1/0/28
```

The following two ways enable the host for ISATAP communication in the HME1 environment:

- Manual definition of the ISATAP IPv4 router address
- Manual definition of the ISATAP IPv4 DNS name (requires DNS record entries)

Using the ISATAP IPv4 router address method is straightforward but difficult to scale without some kind of script or host management tools. As previously mentioned, various tools such as Windows Scripting Host and Microsoft SMS Server can be used to run the command locally on the host at login or another predetermined time.

On the Microsoft Windows XP or Windows Vista host in VLAN 2, ISATAP is enabled and the IPv4 ISATAP router address is defined (IPv6 has already been enabled on the host). As previously mentioned, the HME1 design maps the host in a VLAN/subnet to a specific ISATAP router address. Here the host is in VLAN 2, which is in the 10.120.2.0/24 subnet and is therefore configured to use the ISATAP router of 10.122.10.102 where the “2” in “102” signifies VLAN or Subnet 2. The same would happen for VLAN 3 or 10.120.3.0/24 where the ISATAP router is 10.122.10.103:

```
C:\>netsh interface ipv6 isatap set router 10.122.10.102 enabled
Ok.
```

The following command can be used to verify that the address has been accepted:

```
C:\>netsh interface ipv6 isatap show router
```

```

Router Name       : 10.122.10.102
Use Relay        : enabled
Resolution Interval : default

```

The host has successfully established an ISATAP connection to the primary core layer switch (6k-core-1) and received a valid prefix (2001:db8:cafe:2:0:5efe:10.120.2.101). ISATAP uses the IPv4 address on the host as the right-most 32-bit portion of the 64-bit interface ID. ISATAP “pads” the left-most 32-bits of the 64-bit interface ID with “0000:5efe”. The IPv4 address (10.120.2.101) is used as the tunnel source on the host side of the tunnel, and loopback2 (10.122.10.102) on the core layer switches is used as the tunnel destination (previously configured ISATAP router address) for the host.

The tunnel adapter automatic tunneling pseudo-interface is as follows:

```

Connection-specific DNS Suffix . :
IP Address. . . . . : 2001:db8:cafe:2:0:5efe:10.120.2.101
IP Address. . . . . : fe80::5efe:10.120.2.101%2
Default Gateway . . . . . : fe80::5efe:10.122.10.102%2

```

Using the ISATAP IPv4 router name method is also straightforward, but requires DNS entries. It is also difficult to scale without some kind of script or host management tools. As previously mentioned, various tools such as Windows Scripting Host and Microsoft SMS Server can be used to run the command locally on the host at time of login or another predetermined time.

In this example, a name is used for the ISATAP router instead of an ISATAP IPv4 address. The default DNS name that ISATAP tries to resolve is “isatap” along with the domain suffix. For example, if this host is in domain “cisco.com”, the host attempts to resolve “isatap.cisco.com”. The user has the capability to alter this name similarly to altering the address selection.

```

C:\>netsh interface ipv6 isatap set router vlan2-isatap enabled
Ok.

```

```

C:\>netsh interface ipv6 isatap show router
Router Name       : vlan2-isatap
Use Relay        : enabled
Resolution Interval : default

```

On the DNS server, the following entries were made for the two VLANs shown in this document:

- vlan2-isatap—Host (A) 10.122.10.102
- vlan3-isatap—Host (A) 10.122.10.103

QoS Configuration

The QoS policies for HME1 should match the existing IPv4 policies. As previously mentioned, the HME1 model presents a challenge with respect to where the IPv6 packets are classified and marked. The IPv6 packets are encapsulated within ISATAP tunnels all the way from the host in the access layer to the core layer, and IPv6 QoS policies cannot see the packets inside the tunnel. The first point where the IPv6 packets can have policies applied is at the egress interfaces of the core layer switches. The following configuration is meant as a simple example only and is not based on Cisco campus QoS recommendations. In this policy, class maps are used to match against IPv6 access lists as listed in [Table 7](#).

Table 7 IPv6 QoS—Class Map, Match ACL, and DSCP Setting

Application	Access Group Name	DSCP Setting
FTP	BULK-APPS	AF11

Table 7 IPv6 QoS—Class Map, Match ACL, and DSCP Setting (continued)

Telnet	TRANSACTIONAL-APPS	AF21
SSH	TRANSACTIONAL-APPS	AF21
ALL OTHERS	N/A	0 (default)

The policy is applied on egress interfaces (upstream from the access layer). Upstream switches can trust these DSCP settings and also apply queuing and policing as appropriate (see [Dual-Stack Model—Implementation, page 33](#)).

- 6k-core-1

```

mls qos
!
class-map match-all CAMPUS-BULK-DATA
  match access-group name BULK-APPS
class-map match-all CAMPUS-TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-APPS
!
policy-map IPv6-ISATAP-MARK
  class CAMPUS-BULK-DATA
    set dscp af11
  class CAMPUS-TRANSACTIONAL-DATA
    set dscp af21
  class class-default
    set dscp default
!
ipv6 access-list BULK-APPS
  permit tcp any any eq ftp
  permit tcp any any eq ftp-data
!
ipv6 access-list TRANSACTIONAL-APPS
  permit tcp any any eq telnet
  permit tcp any any eq 22
!
interface GigabitEthernet2/1
  description to 6k-agg-1
  mls qos trust dscp
  service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet2/2
  description to 6k-agg-2
  mls qos trust dscp
  service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet2/3
  description to 6k-core-1
  mls qos trust dscp
  service-policy output IPv6-ISATAP-MARK

```

Infrastructure Security Configuration

In addition to the security configurations discussed in [Addressing, page 18](#), the customer may want to further tighten IPv6 access control for ISATAP tunnels at the access layer. An access list can be applied to either a host port or an uplink/trunk port at the access layer. It is easier to manage the ACL at the uplink rather than configuring ACLs on each host port.

One access list that can be used is an ACL to permit tunnels from the hosts on the access switch to the ISATAP router address for that VLAN. For example, the following ACL permits the ISATAP tunnels (via protocol 41) only if their destination is 10.122.10.102 (the ISATAP router address previously configured). Again, this ACL can be applied on a specific host port on input (**ip access-group 100 in**) or an uplink trunk or routed port (**ip access-group 100 out**).

```
access-list 100 remark Permit approved IPv6-Tunnels
access-list 100 permit 41 any host 10.122.10.102
access-list 100 deny 41 any any log-input
access-list 100 permit ip any any
```

Service Block Model—Implementation

The ISATAP deployment on the SBM is nearly identical to that of HME1. Both models deploy a redundant pair of switches used to provide fault tolerant termination of ISATAP tunnels coming from the hosts in the access layer. The only difference between the SBM and HME1 is that the SBM is using a new set of switches that are dedicated to terminating connections (ISATAP, configured tunnels, or dual-stack) while the HME1 uses the existing core layer switches for termination.

This section is focused on the configuration of the interfaces on the service block switches (physical and logical) as well as the data center aggregation layer tunnel interfaces (show only for completeness). The entire IPv4 network is the same as the one described in the HME1 configuration.

Also, the host configuration for the SBM is the same as HME1 because the ISATAP router addresses have to be reused in this example. Similar to the HME1 configuration section, the loopback, tunnel, routing, and high availability configurations are all presented.

Network Topology

To keep the diagrams simple to understand, the topology is separated into two parts: the ISATAP topology and the manually-configured tunnel topology.

[Figure 17](#) shows the ISATAP topology for the SBM. The topology is focused on the IPv4 addressing in the access layer (used by the host to establish the ISATAP tunnel), the service block (used as the termination point for the ISATAP tunnels), and also the IPv6 addressing used in the service block for both the p2p link and the ISATAP tunnel prefix. The configuration shows that the ISATAP availability is accomplished by using loopback interfaces that share the same IPv4 address between both SBM switches. To maintain prefix consistency for the ISATAP hosts in the access layer, the same prefix is used on both the primary and backup ISATAP tunnels.

Figure 17 SBM ISATAP Network Topology





Figure 18 shows the manually-configured tunnel topology for the SBM. The topology diagram shows the loopback addresses on the service block switches (used as the tunnel source for configured tunnels) and the IPv6 addressing used on the manually-configured tunnel interfaces.

Figure 18 SBM Manually-Configured Tunnel Topology



①	2001:db8:cafe:6505::a111:1010/64 - P2P Link	
②	2001:db8:cafe:6505::b222:2020/64 - P2P Link	
③	2001:db8:cafe:6502::a111:1010/64 #6k-sb-1 2001:db8:cafe:6502::d444:4040/64 #6k-agg-2	⑤ 2001:db8:cafe:6504::b222:2020/64 #6k-sb-2 2001:db8:cafe:6504::d444:4040/64 #6k-agg-2
④	2001:db8:cafe:6501::a111:1010/64 #6k-sb-1 2001:db8:cafe:6501::c333:3030/64 #6k-agg-1	⑥ 2001:db8:cafe:6503::b222:2020/64 #6k-sb-2 2001:db8:cafe:6503::c333:3030/64 #6k-agg-1

220119

Physical Configuration

The configurations for both service block switches are shown, including the core layer-facing interfaces. Configurations for the IPv4 portion of the above topology is shown only for the service block switches. All other IPv4 configurations are based on existing campus design best practices and are not discussed in this section. Full configurations that include the IPv4 setup are available in [Appendix—Configuration Listings, page 66](#).

- 6k-sb-1

```
interface GigabitEthernet4/1
  description to 6k-core-1
  dampening
  ip address 10.122.0.78 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
```

```

ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
interface GigabitEthernet4/2
  description to 6k-core-2
  dampening
  ip address 10.122.0.86 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  mls qos trust dscp
!
interface TenGigabitEthernet1/1
  description to 6k-sb-2
  dampening
  ip address 10.122.0.93 255.255.255.252
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:6505::A111:1010/64           #p2p link between SBM
                                                         #switches
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  mls qos trust dscp

```

- 6k-sb-2

```

interface GigabitEthernet4/1
  description to 6k-core-1
  dampening
  ip address 10.122.0.82 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  mls qos trust dscp
!
interface GigabitEthernet4/2
  description to 6k-core-2
  dampening
  ip address 10.122.0.90 255.255.255.252

```

```

no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
interface TenGigabitEthernet1/1
  description to 6k-sb-1
  dampening
  ip address 10.122.0.94 255.255.255.252
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:6505::B222:2020/64
  #p2p link between SBM
  #switches

no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
mls qos trust dscp

```

Tunnel Configuration

The tunnel and routing configuration for ISATAP is exactly the same as for HME1. The configurations are shown below, but to avoid repeating information presented in previous sections, none of the configurations for the ISATAP tunneling and routing are explained (see the HME1 example explanations).

The manually-configured tunnel configurations are shown for the service block switches. The tunnel configurations for the data center aggregation switches (6k-agg-1/6k-agg-2) are identical to the service block except for address specifics.

- 6k-sb-1

```

interface Loopback0
  description Tunnel source for 6k-agg-1
  ip address 10.122.10.9 255.255.255.255
!
interface Loopback1
  description Tunnel source for 6k-agg-2
  ip address 10.122.10.19 255.255.255.255
!
interface Loopback2
  description Tunnel source for ISATAP-VLAN2
  ip address 10.122.10.102 255.255.255.255
!
interface Loopback3
  description Tunnel source for ISATAP-VLAN3

```

```

ip address 10.122.10.103 255.255.255.255
!
interface Tunnel0
  description tunnel to 6k-agg-1
  no ip address
  ipv6 address 2001:DB8:CAFE:6501::A111:1010/64
  no ipv6 redirects
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  tunnel source Loopback0
  tunnel destination 10.122.10.1 #10.122.10.1 is loopback0 on 6k-agg-1
  tunnel mode ipv6ip
!
interface Tunnel1
  description tunnel to 6k-agg-2
  no ip address
  ipv6 address 2001:DB8:CAFE:6502::A111:1010/64
  no ipv6 redirects
  ipv6 nd reachable-time 5000
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  tunnel source Loopback1
  tunnel destination 10.122.10.2 #10.122.10.2 is loopback0 on 6k-agg-2
  tunnel mode ipv6ip
!
interface Tunnel2
  description ISATAP VLAN2
  no ip address
  no ip redirects
  ipv6 address 2001:DB8:CAFE:2::/64 eui-64
  no ipv6 redirects
  no ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf 1 area 2
  tunnel source Loopback2
  tunnel mode ipv6ip isatap
!
interface Tunnel3
  description ISATAP VLAN3
  no ip address
  no ip redirects
  ipv6 address 2001:DB8:CAFE:3::/64 eui-64
  no ipv6 redirects
  no ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf 1 area 2
  tunnel source Loopback3
  tunnel mode ipv6ip isatap
!
ipv6 router ospf 1
  router-id 10.122.10.9
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  area 2 range 2001:DB8:CAFE:2::/64 cost 10
  area 2 range 2001:DB8:CAFE:3::/64 cost 10
  passive-interface Loopback0
  passive-interface Loopback1
  passive-interface Loopback2

```

```

passive-interface Loopback3
passive-interface Tunnel2
passive-interface Tunnel3
timers spf 1 5

```

- 6k-sb-2

```

interface Loopback0
  description Tunnel source for 6k-agg-1
  ip address 10.122.10.10 255.255.255.255
!
interface Loopback1
  description Tunnel source for 6k-agg-2
  ip address 10.122.10.20 255.255.255.255
!
interface Loopback2
  description Tunnel source for ISATAP-VLAN2
  ip address 10.122.10.102 255.255.255.255
  delay 1000
!
interface Loopback3
  description Tunnel source for ISATAP-VLAN3
  ip address 10.122.10.103 255.255.255.255
  delay 1000
!
interface Tunnel0
  description tunnel to 6k-agg-1
  no ip address
  load-interval 30
  ipv6 address 2001:DB8:CAFE:6503::B222:2020/64
  no ipv6 redirects
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf priority 255
  ipv6 ospf 1 area 0
  tunnel source Loopback0
  tunnel destination 10.122.10.11
  tunnel mode ipv6ip
!
interface Tunnel1
  description tunnel to 6k-agg-2
  no ip address
  load-interval 30
  ipv6 address 2001:DB8:CAFE:6504::B222:2020/64
  no ipv6 redirects
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  tunnel source Loopback1
  tunnel destination 10.122.10.12
  tunnel mode ipv6ip
!
interface Tunnel2
  description ISATAP VLAN2
  no ip address
  ip access-group 100 in
  no ip redirects
  load-interval 30
  ipv6 address 2001:DB8:CAFE:2::/64 eui-64
  no ipv6 redirects

```

```

no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback2
tunnel mode ipv6ip isatap
!
interface Tunnel3
description ISATAP VLAN3
no ip address
no ip redirects
load-interval 30
ipv6 address 2001:DB8:CAFE:3::/64 eui-64
no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback3
tunnel mode ipv6ip isatap
!
ipv6 router ospf 1
router-id 10.122.10.10
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 range 2001:DB8:CAFE:2::/64 cost 20
area 2 range 2001:DB8:CAFE:3::/64 cost 20
passive-interface Loopback0
passive-interface Loopback1
passive-interface Loopback2
passive-interface Loopback3
passive-interface Tunnel2
passive-interface Tunnel3
timers spf 1 5

```

QoS Configuration

The same QoS configurations and discussions from [Dual-Stack Model—Implementation, page 33](#) and [Conclusion, page 63](#) apply to the SBM. Based on the example configuration shown in the case of HME1, the only change relates to the interfaces where the classification and marking policies are applied. In the SBM, the service policy is applied to the egress on the manually-configured tunnels towards 6k-agg-1 and 6k-agg-2.

As an example for 6k-sb-1, the service policy would be applied to Tunnel0 and Tunnel1:

```

interface Tunnel0
description tunnel to 6k-agg-1
service-policy output IPv6-ISATAP-MARK
!
interface Tunnel1
description tunnel to 6k-agg-2
service-policy output IPv6-ISATAP-MARK

```

Infrastructure Security Configuration

The security considerations and configurations discussed in [Dual-Stack Model—Implementation, page 33](#) and [Infrastructure Security Configuration, page 54](#) apply directly to the SBM.

Conclusion

This document analyzes various architectures for providing IPv6 services in campus networks. The models discussed are certainly not the only ways to deploy IPv6 in this environment, but they provide options that can be leveraged based on environment, deployment schedule, and targeted services specifics.

Table 8 summarizes the benefits and challenges with each of the models discussed in this document.

Table 8 *Benefits and Challenges of Various Models*

Model	Benefit	Challenge
Dual-stack model (DSM)	<p>No tunneling required</p> <p>No dependency on IPv4 (routing, QoS, HA, multicast, security, and management are separated)</p> <p>Superior performance and highest availability for IPv6 unicast and multicast</p> <p>Scalable</p>	<p>Requires IPv6 hardware-enabled campus switching equipment</p> <p>Operational challenges with supporting dual protocols— Training/management tools</p>
Hybrid model example 1 (HME1)	<p>Most of the existing, IPv4-only campus equipment can be used (access and distribution layer)</p> <p>Per-user or per-application control for IPv6-service delivery</p> <p>Provides high-availability for IPv6 access over ISATAP tunnels</p>	<p>Tunneling is required; increase in operations and management</p> <p>Unknown scale factors:</p> <ul style="list-style-type: none"> • How many ISATAP tunnels are too many? • How many hosts per ISATAP tunnel are too many? • How much performance is sacrificed for host-to-service access? <p>IPv6 multicast is not supported</p> <p>Causes core layer to become an access layer for IPv6 tunnels</p> <p>Requires IPv6-enabled hosts with ISATAP configuration</p>

Table 8 Benefits and Challenges of Various Models (continued)

Hybrid model example 2 (HME2)	Core layer not required to be dual-stack Temporary solution when core is being upgraded or has IPv6 hardware support limitations IPv6 multicast is supported (not as high performing or scalable as DSM) Provides high availability for IPv6 connectivity over configured tunnels	Tunneling is required; increase in operations and management In large networks, the use of many p2p manually-configured tunnels can be hard to scale and manage Requires IPv6-enabled hosts
Service block model (SBM)	Highly reduced time-to-delivery for IPv6-enabled services Requires no changes to existing campus infrastructure Per-user or per-application control for IPv6-service delivery Provides high-availability for IPv6 access over ISATAP tunnels Provides high-availability for IPv6 connectivity over configured tunnels	New IPv6 hardware capable, campus switches are required Tunneling is required (extensively)—increase in operations and management Many unknown scale factors (see HME1) IPv6 multicast is not supported on the ISATAP tunnels Requires IPv6-enabled hosts + ISATAP configuration

Future Work

This document is one of several in a series focused on providing basic IPv6 implementation guidance for the enterprise customers. Similar documents will be published analyzing the deployment of IPv6 in the branch, WAN, data center, and enterprise edge.

This document is a “living document”; changes will be made to it as features mature. It is the goal, however, to fully integrate IPv6 into all enterprise architecture design guides where IPv6 will become another baseline component. This will provide one place to go to learn the latest design best practices for every area of the enterprise instead of reading about various technologies or designs in separate papers. The enterprise architecture design guides can be found at the following URL: <http://www.cisco.com/go/srnd>.

Additional References

Many notes and disclaimers in this document discuss the need to fully understand the technology and protocol aspects of IPv6. There are many design considerations associated with the implementation of IPv6 that include security, QoS, availability, management, IT training, and application support.

The following references are a few of the many that provide more details on IPv6, Cisco design recommendations, products and solutions, and industry activity.

- Cisco-specific links

- “Deploying IPv6 Networks” by Ciprian P. Popoviciu, Eric Levy-Abegnoli, Patrick Grossetete (ISBN-10:1-58705-210-5; ISBN-13:978-1-58705-210-1)—
<http://www.ciscopress.com/bookstore/product.asp?isbn=1587052105&rl=1>
- Cisco IPv6 CCO home page—<http://www.cisco.com/ipv6>
- Cisco Solution Reference Network Design (SRND) Guides—<http://www.cisco.com/go/srnd>
- Cisco Solution Reference Network Design (SRND) Campus Guides—
http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html#anchor2
- Enterprise QoS SRND—
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration_09186a008049b062.pdf
- Cisco IOS IPv6 Configuration Guide, Release 12.4—
http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cg/hipv6_c/index.htm
- Catalyst 6500 Series Cisco IOS Software Configuration Guide, 12.2SX—
<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/index.htm>
- Catalyst 3750 Switch Software Configuration Guide, 12.2(25)SEE—
<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12225see/scg/index.htm>
- Cisco Network Virtualization—
http://www.cisco.com/en/US/netsol/ns658/networking_solutions_package.html
- Cisco IOS IPv6 Traffic Filter Configurations—
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/v6_tffw.htm
- Securing Cisco Routers Online Training and Documentation—
http://www.cisco.com/web/about/security/security_services/ciag/workforce_development/securing_cisco_routers.html
- Microsoft IPv6 links
 - Microsoft IPv6 Home—
<http://www.microsoft.com/technet/itsolutions/network/ipv6/default.mspx>
 - Microsoft–Cisco ISATAP white paper—
<http://www.microsoft.com/downloads/details.aspx?FamilyId=B8F50E07-17BF-4B5C-A1F9-5A09E2AF698B&displaylang=en>
 - Microsoft TechNet–Teredo Overview—
<http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/teredo.mspx>
- IPv6 industry links
 - National Security Agency–Security for IPv6 on Cisco Routers—
<http://www.nsa.gov/notices/notic00004.cfm?Address=/snac/routers/I33-002R-06.pdf>
 - IETF IPv6 Ops Working Group—<http://www.ietf.org/html.charters/v6ops-charter.html>
 - go6 IPv6 Portal–IPv6 Knowledge Center—http://wiki.go6.net/index.php?title=Main_Page
 - 6NET–Large-Scale International IPv6 Pilot Network—<http://www.6net.org/>
 - IETF IPv6 Working Group—<http://www.ietf.org/html.charters/ipv6-charter.html>
 - IETF IPv6 Operations Working Group—<http://www.ietf.org/html.charters/v6ops-charter.html>
 - Internet Protocol, Version 6 (IPv6) Specification—<http://www.ietf.org/rfc/rfc2460.txt>
 - Neighbor Discovery for IPv6—<http://www.ietf.org/rfc/rfc2461.txt>

- IPv6 Stateless Address Autoconfiguration—<http://www.ietf.org/rfc/rfc2462.txt>
- Transmission of IPv6 Packets over Ethernet Networks—<http://www.ietf.org/rfc/rfc2464.txt>
- Transition Mechanisms for IPv6 Hosts and Routers—<http://www.ietf.org/rfc/rfc2893.txt>
- Privacy Extensions for Stateless Address Autoconfiguration in IPv6—
<http://www.ietf.org/rfc/rfc3041.txt>
- Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)—
<http://www.ietf.org/rfc/rfc4214.txt>
- IPv6 Addressing Architecture—<http://www.ietf.org/rfc/rfc4291.txt>
- Internet Control Message Protocol (ICMPv6) for Internet Protocol Version 6 (IPv6) Specification—<http://www.ietf.org/rfc/rfc4443.txt>
- North American IPv6 Task Force—IPv6 Security Technology Paper—
http://www.ipv6forum.org/dl/white/NAv6TF_Security_Report.pdf

Appendix—Configuration Listings

This section contains the full configurations of the switches used in the three models discussed (DSM, HME1, and SBM). Some switch configurations are shown only once because their configuration is identical across all models in this paper.

For the sake of brevity, unused or shutdown interfaces are removed from the configurations.

Dual-Stack Model (DSM)

3750-acc-1

```

version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service sequence-numbers
!
hostname 3750-acc-1
!
logging count
logging buffered 8192 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$RPWW$KR0qkU1kW0a7Pd1svvhgB.
!
username cisco privilege 15 secret 5 $1$KNSF$3XavijOuoVCq68RASxTYD1
no aaa new-model
clock timezone mst -7
switch 1 provision ws-c3750g-24ts
vtp domain ese-dc
vtp mode transparent
udld enable

udld message time 7

```

```

ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
ip telnet source-interface Vlan2
no ip domain-lookup
ip domain-name cisco.com
ip dhcp smart-relay
!
ip dhcp snooping vlan 2
ip dhcp snooping database flash:dhcp.txt
ip dhcp snooping database timeout 10
ip dhcp snooping
ip ftp source-interface Vlan2
ip ftp username cisco
ip ftp password 7 120B0419150E1E
ip tftp source-interface Vlan2
ip ssh time-out 60
ip ssh authentication-retries 2
ip ssh source-interface Vlan2
ip ssh version 2
ip arp inspection vlan 2
ip arp inspection validate src-mac
ip arp inspection log-buffer entries 100
ip arp inspection log-buffer logs 20 interval 120
login block-for 30 attempts 3 within 200
login delay 2
ipv6 mld snooping
!
mls qos map policed-dscp 0 10 18 24 25 34 to 8
mls qos map cos-dscp 0 8 16 24 32 46 48 56
mls qos srr-queue output cos-map queue 1 threshold 3 5
mls qos srr-queue output cos-map queue 2 threshold 1 2 4
mls qos srr-queue output cos-map queue 2 threshold 2 3
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
mls qos srr-queue output cos-map queue 3 threshold 3 0
mls qos srr-queue output cos-map queue 4 threshold 3 1
mls qos srr-queue output dscp-map queue 1 threshold 3 46
mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22 25 32 34 36
mls qos srr-queue output dscp-map queue 2 threshold 1 38
mls qos srr-queue output dscp-map queue 2 threshold 2 24 26
mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
mls qos srr-queue output dscp-map queue 3 threshold 3 0
mls qos srr-queue output dscp-map queue 4 threshold 1 8
mls qos srr-queue output dscp-map queue 4 threshold 3 10 12 14
mls qos queue-set output 1 threshold 2 70 80 100 100
mls qos queue-set output 1 threshold 4 40 100 100 100
mls qos
!
crypto pki trustpoint TP-self-signed-3669881984
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-3669881984
  revocation-check none
  rsakeypair TP-self-signed-3669881984
!
!
crypto ca certificate chain TP-self-signed-3669881984
  certificate self-signed 01
    30820299 30820202 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
    57312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
    69666963 6174652D 33363639 38383139 38343124 30220609 2A864886 F70D0109
    02161533 3735302D 65646765 2D312E63 6973636F 2E636F6D 301E170D 39333033
    30313030 30313533 5A170D32 30303130 31303030 3030305A 3057312F 302D0603
    55040313 26494F53 2D53656C 662D5369 676E6564 2D436572 74696669 63617465

```

```

2D333636 39383831 39383431 24302206 092A8648 86F70D01 09021615 33373530
2D656467 652D312E 63697363 6F2E636F 6D30819F 300D0609 2A864886 F70D0101
01050003 818D0030 81890281 8100BF7D E793E21B 6C1F75C2 16AAFF9F C110D038
2D6D1DC9 04758DB8 7B3AD4C9 9F36A3B1 54983BEC 10FEA2D4 151D2783 5765C58A
A95E6364 CCBCF7F9 F4750437 AB8C00BF EFD54E88 6F650A5C 9563A309 247C6070
DC38A870 DEC5D4DA 765AD7A4 274B7649 36D876CA E28CF66C 77335F90 949DF258
E3E019BD 5E6801EC 9E15F980 C13D0203 010001A3 75307330 0F060355 1D130101
FF040530 030101FF 30200603 551D1104 19301782 15333735 302D6564 67652D31
2E636973 636F2E63 6F6D301F 0603551D 23041830 168014E0 18682CE1 D6A3EF2C
32A7C8D5 4DAFB9AA F11A0030 1D060355 1D0E0416 0414E018 682CE1D6 A3EF2C32
A7C8D54D AFB9AAF1 1A00300D 06092A86 4886F70D 01010405 00038181 00755909
C99DEB7F E05FC2A3 482558FA 33C292AE 7E4543E3 5BD6F32F 1D671B97 BC45B73B
85E954ED 7FC58F90 23A38132 24216CDB C978B3DD 9BCBC48E 519D01BF F4CEBB82
07834C2D D82CA163 8E638214 5B5C277D 5E7DD52E 56172675 BD563769 590E4DC6
39AD9BF8 CDBBA241 E5E2C666 5CAE912E 40DC2150 A1CE39B4 D8101D33 A8
quit
!
!
errdisable recovery cause link-flap
errdisable recovery interval 60
no file verify auto
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree portfast bpduguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
vlan 2
 name ACCESS-DATA-2
!
class-map match-all DVLAN-PC-VIDEO
 match access-group name DVLAN-PC-VIDEO
class-map match-all DVLAN-Transactional-Data
 match access-group name DVLAN-Transactional-Data
class-map match-all DVLAN-Mission-Critical-Data
 match access-group name DVLAN-Mission-Critical-Data
class-map match-all DVLAN-Bulk-Data
 match access-group name DVLAN-Bulk-Data
!
!
policy-map DATA
 class DVLAN-PC-VIDEO
  set dscp af41
  police 48000 8000 exceed-action policed-dscp-transmit
 class DVLAN-Mission-Critical-Data
  set dscp 25
  police 5000000 8000 exceed-action policed-dscp-transmit
 class DVLAN-Transactional-Data
  set dscp af21
  police 5000000 8000 exceed-action policed-dscp-transmit
 class DVLAN-Bulk-Data
  set dscp af11
  police 5000000 8000 exceed-action policed-dscp-transmit
!
!
!
interface Null0
 no ip unreachable
!
interface GigabitEthernet1/0/5
 description to PC-DATA-ONLY

```

```

switchport access vlan 2
switchport mode access
switchport port-security maximum 3
switchport port-security
switchport port-security aging time 2
switchport port-security violation restrict
switchport port-security aging type inactivity
ip arp inspection limit rate 100
service-policy input DATA
load-interval 30
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
spanning-tree portfast
spanning-tree bpduguard enable
ip verify source
ip dhcp snooping limit rate 100
!
interface GigabitEthernet1/0/25
description TRUNK TO 6k-dist-1
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 2
switchport mode trunk
switchport nonegotiate
switchport port-security aging time 10
ip arp inspection trust
load-interval 30
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
no cdp enable
ip dhcp snooping limit rate 10
ip dhcp snooping trust
!
interface GigabitEthernet1/0/26
description TRUNK TO 6k-dist-2
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 2
switchport mode trunk
switchport nonegotiate
switchport port-security aging time 10
ip arp inspection trust
load-interval 30
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
no cdp enable
ip dhcp snooping limit rate 10
ip dhcp snooping trust
!
interface Vlan2
ip address 10.120.2.4 255.255.255.0
no ip redirects
no ip unreachable
no ip proxy-arp
no ip route-cache
no ip mroute-cache
ipv6 address 2001:DB8:CAFE:2::CAC1:3750/64
no ipv6 redirects
!
ip default-gateway 10.120.2.1
ip classless
no ip http server
no ip http secure-server

```

```

!
!
ip access-list extended DVLAN-Bulk-Data
  permit tcp any any eq 143
  permit tcp any any eq 220
ip access-list extended DVLAN-Mission-Critical-Data
  permit tcp any any range 3200 3203
  permit tcp any any eq 3600
  permit tcp any any range 2000 2002
ip access-list extended DVLAN-PC-VIDEO
  permit udp any any range 16384 32767
ip access-list extended DVLAN-Transactional-Data
  permit tcp any any eq 1352
ip access-list extended MGMT-IN-v4
  permit tcp 10.120.0.0 0.0.255.255 any log-input
  permit tcp 10.121.0.0 0.0.255.255 any log-input
  permit tcp 10.122.0.0 0.0.255.255 any log-input
  deny ip any any log-input
!
logging source-interface Vlan2
logging 10.121.11.9
no cdp run
ipv6 route ::/0 2001:DB8:CAFE:2::1
!
snmp-server contact John Doe - ipv6rocks@cisco.com
snmp-server group IPv6-ADMIN v3 auth write v1default
!
ipv6 access-list MGMT-IN
  remark Permit MGMT only to VLAN2
  permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:2::CAC1:3750 log-input
  deny ipv6 any any log-input
!
control-plane
!
banner login ^C
Unauthorized access to this device and/or network is prohibited.
^C
!
line con 0
  session-timeout 3
  password 7 06140E2F4B4B1B
  logging synchronous
  login local
  transport output telnet ssh
line vty 0 4
  session-timeout 3
  password 7 06140E2F4B4B1B
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local
  exec prompt timestamp
  transport input telnet ssh
line vty 5 15
  session-timeout 3
  password 7 0101070A5C0E14
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local
  exec prompt timestamp
  transport input telnet ssh
!
!
end

```

3750-acc-2

```

version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service sequence-numbers
!
hostname 3750-acc-2
!
logging count
logging buffered 8192 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$2Nvb$4.Zz3eHzz2KgCFzc12pLh.
!
username cisco privilege 15 secret 5 $1$/S8z$0ek0VBA.QWuKR8R6U1.Ly0
no aaa new-model
clock timezone mst -7
switch 1 provision ws-c3750g-24ts
vtp domain ese-dc
vtp mode transparent
udld enable

udld message time 7

ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
ip telnet source-interface Vlan3
no ip domain-lookup
ip domain-name cisco.com
ip dhcp smart-relay
!
ip dhcp snooping vlan 3
ip dhcp snooping database flash:dhcp.txt
ip dhcp snooping database timeout 10
ip dhcp snooping
ip ftp source-interface Vlan3
ip ftp username cisco
ip ftp password 7 111B180B101719
ip tftp source-interface Vlan3
ip ssh time-out 60
ip ssh authentication-retries 2
ip ssh source-interface Vlan3
ip ssh version 2
ip arp inspection vlan 3
ip arp inspection validate src-mac
ip arp inspection log-buffer entries 100
ip arp inspection log-buffer logs 20 interval 120
login block-for 30 attempts 3 within 200
login delay 2
ipv6 mld snooping
!
mls qos map policed-dscp 0 10 18 24 25 34 to 8
mls qos map cos-dscp 0 8 16 24 32 46 48 56
mls qos srr-queue output cos-map queue 1 threshold 3 5
mls qos srr-queue output cos-map queue 2 threshold 1 2 4
mls qos srr-queue output cos-map queue 2 threshold 2 3
mls qos srr-queue output cos-map queue 2 threshold 3 6 7

```

```

mls qos srr-queue output cos-map queue 3 threshold 3 0
mls qos srr-queue output cos-map queue 4 threshold 3 1
mls qos srr-queue output dscp-map queue 1 threshold 3 46
mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22 25 32 34 36
mls qos srr-queue output dscp-map queue 2 threshold 1 38
mls qos srr-queue output dscp-map queue 2 threshold 2 24 26
mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
mls qos srr-queue output dscp-map queue 3 threshold 3 0
mls qos srr-queue output dscp-map queue 4 threshold 1 8
mls qos srr-queue output dscp-map queue 4 threshold 3 10 12 14
mls qos queue-set output 1 threshold 2 70 80 100 100
mls qos queue-set output 1 threshold 4 40 100 100 100
mls qos
!
errdisable recovery cause link-flap
errdisable recovery interval 60
no file verify auto
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree portfast bpduguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
vlan internal allocation policy ascending
vlan dot1q tag native
!
vlan 3
 name ACCESS-DATA-3
!
class-map match-all DVLAN-PC-VIDEO
 match access-group name DVLAN-PC-VIDEO
class-map match-all DVLAN-Transactional-Data
 match access-group name DVLAN-Transactional-Data
class-map match-all DVLAN-Mission-Critical-Data
 match access-group name DVLAN-Mission-Critical-Data
class-map match-all DVLAN-Bulk-Data
 match access-group name DVLAN-Bulk-Data
!
!
policy-map DATA
 class DVLAN-PC-VIDEO
  set dscp af41
  police 48000 8000 exceed-action policed-dscp-transmit
 class DVLAN-Mission-Critical-Data
  set dscp 25
  police 5000000 8000 exceed-action policed-dscp-transmit
 class DVLAN-Transactional-Data
  set dscp af21
  police 5000000 8000 exceed-action policed-dscp-transmit
 class DVLAN-Bulk-Data
  set dscp af11
  police 5000000 8000 exceed-action policed-dscp-transmit
!
interface Null0
 no ip unreachable
!
interface GigabitEthernet1/0/5
 description to PC-DATA-ONLY
 switchport access vlan 3
 switchport mode access
 switchport port-security maximum 3
 switchport port-security
 switchport port-security aging time 2

```

```

switchport port-security violation restrict
switchport port-security aging type inactivity
ip arp inspection limit rate 100
service-policy input DATA
load-interval 30
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
spanning-tree portfast
spanning-tree bpduguard enable
ip verify source
ip dhcp snooping limit rate 100
!
interface GigabitEthernet1/0/25
description TRUNK TO 6k-dist-1
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 3
switchport mode trunk
switchport nonegotiate
switchport port-security aging time 10
ip arp inspection trust
load-interval 30
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
no cdp enable
spanning-tree guard loop
ip dhcp snooping limit rate 10
ip dhcp snooping trust
!
interface GigabitEthernet1/0/26
description TRUNK TO 6k-dist-2
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 3
switchport mode trunk
switchport nonegotiate
switchport port-security aging time 10
ip arp inspection trust
load-interval 30
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
no cdp enable
spanning-tree guard loop
ip dhcp snooping limit rate 10
ip dhcp snooping trust
!
interface Vlan3
ip address 10.120.3.4 255.255.255.0
no ip redirects
no ip unreachable
no ip proxy-arp
no ip route-cache
no ip mroute-cache
ipv6 address 2001:DB8:CAFE:3::CAC2:3750/64
no ipv6 redirects
!
ip default-gateway 10.120.3.1
ip classless
no ip http server
no ip http secure-server
!
ip access-list extended DVLAN-Bulk-Data
permit tcp any any eq 143

```

```

    permit tcp any any eq 220
ip access-list extended DVLAN-Mission-Critical-Data
    permit tcp any any range 3200 3203
    permit tcp any any eq 3600
    permit tcp any any range 2000 2002
ip access-list extended DVLAN-PC-VIDEO
    permit udp any any range 16384 32767
ip access-list extended DVLAN-Transactional-Data
    permit tcp any any eq 1352
ip access-list extended MGMT-IN-v4
    permit tcp 10.120.0.0 0.0.255.255 any log-input
    permit tcp 10.121.0.0 0.0.255.255 any log-input
    permit tcp 10.122.0.0 0.0.255.255 any log-input
    deny ip any any log-input
!
logging source-interface Vlan3
logging 10.121.11.9
no cdp run
ipv6 route ::/0 2001:DB8:CAFE:3::1
!
snmp-server contact John Doe - ipv6rocks@cisco.com
snmp-server group IPv6-ADMIN v3 auth write v1default
!
ipv6 access-list MGMT-IN
    remark Permit MGMT only to VLAN2
    permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:3::CAC2:3750 log-input
    deny ipv6 any any log-input
!
control-plane
!
banner login ^C
Unauthorized access to this device and/or network is prohibited.
^C
!
line con 0
    session-timeout 3
    password 7 00161208035E19
    logging synchronous
    login local
    transport output telnet ssh
line vty 0 4
    session-timeout 3
    password 7 0519070126495C
    ipv6 access-class MGMT-IN in
    logging synchronous
    login local
    exec prompt timestamp
    transport input telnet ssh
line vty 5 15
    session-timeout 3
    password 7 071D2042490C0B
    ipv6 access-class MGMT-IN in
    logging synchronous
    login local
    exec prompt timestamp
    transport input telnet ssh
!
end

```

6k-dist-1

```

upgrade fpd auto

```

```

version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service sequence-numbers
service counters max age 5
!
hostname 6k-dist-1
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$EZB7$nbqNvH9QH4qo3zpGsZog6/
!
username cisco privilege 15 secret 5 $1$VwbR$aXCTEAusOPhpmcfig7nOn01
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
!
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 121A0C041104
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip cef load-sharing full
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
    key-string 7 1111
!
redundancy
  mode sso
  main-cpu

```

```

    auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 2-3 priority 24576
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
vlan 2
    name ACCESS-DATA-2
!
vlan 3
    name ACCESS-DATA-3
!
interface Loopback0
    ip address 10.122.10.9 255.255.255.255
    no ip redirects
    no ip unreachable
    no ip proxy-arp
    ipv6 address 2001:DB8:CAFE:6507::A111:1010/128
    no ipv6 redirects
    ipv6 ospf 1 area 0
!
interface Null0
    no ip unreachable
!
interface TenGigabitEthernet1/1
    description to 6k-dist-2
    dampening
    ip address 10.122.0.93 255.255.255.252
    no ip redirects
    no ip unreachable
    no ip proxy-arp
    ip hello-interval eigrp 10 1
    ip hold-time eigrp 10 3
    ip authentication mode eigrp 10 md5
    ip authentication key-chain eigrp 10 eigrp
    load-interval 30
    carrier-delay msec 0
    ipv6 address 2001:DB8:CAFE:7004::A111:1010/64
    no ipv6 redirects
    ipv6 nd suppress-ra
    ipv6 cef
    ipv6 ospf network point-to-point
    ipv6 ospf hello-interval 1
    ipv6 ospf dead-interval 3
    ipv6 ospf 1 area 2
    wrr-queue bandwidth 5 25 70
    wrr-queue queue-limit 5 25 40
    wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
    wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
    wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
    wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
    wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
    wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
    wrr-queue cos-map 1 1 1

```

```

wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
!
interface GigabitEthernet3/1
description to 3750-acc-1
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 2
switchport mode trunk
switchport nonegotiate
no ip address
load-interval 30
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable
spanning-tree guard root
!
interface GigabitEthernet3/2
description to 3750-acc-2
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 3
switchport mode trunk
switchport nonegotiate
no ip address
load-interval 30
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable
spanning-tree guard root
!

```

```

interface GigabitEthernet4/1
  description to 6k-core-1
  dampening
  ip address 10.122.0.78 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  ip summary-address eigrp 10 10.120.0.0 255.255.0.0 5
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7000::A111:1010/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  wrr-queue bandwidth 30 70
  wrr-queue queue-limit 40 30
  wrr-queue random-detect min-threshold 1 40 80
  wrr-queue random-detect min-threshold 2 70 80
  wrr-queue random-detect max-threshold 1 80 100
  wrr-queue random-detect max-threshold 2 80 100
  wrr-queue cos-map 1 1 1
  wrr-queue cos-map 1 2 0
  wrr-queue cos-map 2 1 2 3 4
  wrr-queue cos-map 2 2 6 7
  mls qos trust dscp
!
interface GigabitEthernet4/2
  description to 6k-core-2
  dampening
  ip address 10.122.0.86 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  ip summary-address eigrp 10 10.120.0.0 255.255.0.0 5
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7001::A111:1010/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  wrr-queue bandwidth 30 70
  wrr-queue queue-limit 40 30
  wrr-queue random-detect min-threshold 1 40 80
  wrr-queue random-detect min-threshold 2 70 80
  wrr-queue random-detect max-threshold 1 80 100
  wrr-queue random-detect max-threshold 2 80 100
  wrr-queue cos-map 1 1 1
  wrr-queue cos-map 1 2 0
  wrr-queue cos-map 2 1 2 3 4
  wrr-queue cos-map 2 2 6 7

```

```

mls qos trust dscp
!
interface Vlan2
description ACCESS-DATA-2
ip address 10.120.2.2 255.255.255.0
ip helper-address 10.121.10.7
no ip redirects
no ip unreachablees
no ip proxy-arp
ipv6 address 2001:DB8:CAFE:2::1/64 anycast
ipv6 address 2001:DB8:CAFE:2::A111:1010/64
ipv6 traffic-filter VLAN2-v6-INGRESS in
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf 1 area 2
arp timeout 200
standby 1 ip 10.120.2.1
standby 1 timers msec 250 msec 800
standby 1 priority 110
standby 1 preempt delay minimum 180
standby 1 authentication ese
!
interface Vlan3
description ACCESS-DATA-3
ip address 10.120.3.2 255.255.255.0
ip helper-address 10.121.10.7
no ip redirects
no ip unreachablees
no ip proxy-arp
ipv6 address 2001:DB8:CAFE:3::1/64 anycast
ipv6 address 2001:DB8:CAFE:3::A111:1010/64
ipv6 traffic-filter VLAN3-v6-INGRESS in
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf 1 area 2
arp timeout 200
standby 1 ip 10.120.3.1
standby 1 timers msec 250 msec 800
standby 1 priority 110
standby 1 preempt delay minimum 180
standby 1 authentication ese
!
router eigrp 10
passive-interface Vlan2
passive-interface Vlan3
passive-interface Loopback0
network 10.0.0.0
no auto-summary
eigrp router-id 10.122.10.9
!
ip classless
!
no ip http server
!
ip access-list extended MGMT-IN-v4
remark Permit v4MGMT only to Lo0
permit tcp 10.120.0.0 0.0.255.255 any log-input
permit tcp 10.121.0.0 0.0.255.255 any log-input
permit tcp 10.122.0.0 0.0.255.255 any log-input
deny ip any any log-input
!
logging source-interface Loopback0

```

```

logging 10.121.11.9
ipv6 router ospf 1
router-id 10.122.10.9
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 range 2001:DB8:CAFE:2::/64 cost 10
area 2 range 2001:DB8:CAFE:3::/64 cost 10
passive-interface Vlan2
passive-interface Vlan3
passive-interface Loopback0
timers spf 1 5
!
snmp-server contact John Doe - ipv6rocks@cisco.com
snmp-server group IPv6-ADMIN v3 auth write v1default
!
ipv6 access-list MGMT-IN
remark Permit MGMT only to Loopback0
permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::A111:1010 log-input
deny ipv6 any any log-input
!
ipv6 access-list VLAN2-v6-INGRESS
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX CAFE:2::/64
permit icmp 2001:DB8:CAFE:2::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX CAFE:2::64
permit ipv6 2001:DB8:CAFE:2::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
ipv6 access-list VLAN3-v6-INGRESS
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX CAFE:3::/64
permit icmp 2001:DB8:CAFE:3::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX CAFE:3::64
permit ipv6 2001:DB8:CAFE:3::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C
Unauthorized access to this device and/or network is prohibited!
^C
!
line con 0
session-timeout 3
password 7 06140E2F4B4B1B
logging synchronous
login local
transport output none
line vty 0 4
session-timeout 3
access-class MGMT-IN-v4 in
password 7 08334D400E1C17
ipv6 access-class MGMT-IN in
logging synchronous
login local
exec prompt timestamp
transport input telnet ssh
!

```

```
no cns aaa enable
end
```

6k-dist-2

```
upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-dist-2
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$0UKB$YN.zu/IQt0ivlkqLsdxtY0
!
username cisco privilege 15 secret 5 $1$SUYV$ShOXyooFT..fL/4tyOnAN1
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 01100F175804
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip cef load-sharing full
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
```

```

key chain eigrp
  key 100
    key-string 7 1111
  !
redundancy
  mode sso
  main-cpu
  auto-sync running-config
  !
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 2-3 priority 28672
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
  !
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
  !
vlan 2
  name ACCESS-DATA-2
  !
vlan 3
  name ACCESS-DATA-3
  !
  !
interface Loopback0
  ip address 10.122.10.10 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ipv6 address 2001:DB8:CAFE:6507::B222:2020/128
  no ipv6 redirects
  ipv6 ospf 1 area 0
  !
interface Null0
  no ip unreachable
  !
interface TenGigabitEthernet1/1
  description to 6k-dist-1
  dampening
  ip address 10.122.0.94 255.255.255.252
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7004::B222:2020/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 2
  wrp-queue bandwidth 5 25 70

```

```

wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
!
interface GigabitEthernet3/1
description to 3750-acc-1
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 2
switchport mode trunk
switchport nonegotiate
no ip address
load-interval 30
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable
!
interface GigabitEthernet3/2
description to 3750-acc-1
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 3
switchport mode trunk
switchport nonegotiate
no ip address
load-interval 30
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2

```

```

wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable
!
interface GigabitEthernet4/1
description to 6k-core-1
dampening
ip address 10.122.0.82 255.255.255.252
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
ip summary-address eigrp 10 10.120.0.0 255.255.0.0 5
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7002::B222:2020/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
interface GigabitEthernet4/2
description to 6k-core-2
dampening
ip address 10.122.0.90 255.255.255.252
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
ip summary-address eigrp 10 10.120.0.0 255.255.0.0 5
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7003::B222:2020/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80

```

```

wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
interface Vlan2
description ACCESS-DATA-2
ip address 10.120.2.3 255.255.255.0
ip helper-address 10.121.10.7
no ip redirects
no ip unreachablees
no ip proxy-arp
ipv6 address 2001:DB8:CAFE:2::1/64 anycast
ipv6 address 2001:DB8:CAFE:2::B222:2020/64
ipv6 traffic-filter VLAN2-v6-INGRESS in
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf 1 area 2
arp timeout 200
standby 1 ip 10.120.2.1
standby 1 timers msec 250 msec 800
standby 1 priority 105
standby 1 preempt delay minimum 180
standby 1 authentication ese
!
interface Vlan3
description ACCESS-DATA-3
ip address 10.120.3.3 255.255.255.0
ip helper-address 10.121.10.7
no ip redirects
no ip unreachablees
no ip proxy-arp
ipv6 address 2001:DB8:CAFE:3::1/64 anycast
ipv6 address 2001:DB8:CAFE:3::B222:2020/64
ipv6 traffic-filter VLAN3-v6-INGRESS in
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf 1 area 2
arp timeout 200
standby 1 ip 10.120.3.1
standby 1 timers msec 250 msec 800
standby 1 priority 105
standby 1 preempt delay minimum 180
standby 1 authentication ese
!
router eigrp 10
passive-interface Vlan2
passive-interface Vlan3
passive-interface Loopback0
network 10.0.0.0
no auto-summary
eigrp router-id 10.122.10.10
!
ip classless
!
no ip http server
!
ip access-list extended MGMT-IN-v4
remark Permit v4MGMT only to Lo0

```

```

permit tcp 10.120.0.0 0.0.255.255 any log-input
permit tcp 10.121.0.0 0.0.255.255 any log-input
permit tcp 10.122.0.0 0.0.255.255 any log-input
deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
ipv6 router ospf 1
router-id 10.122.10.10
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 range 2001:DB8:CAFE:2::/64 cost 10
area 2 range 2001:DB8:CAFE:3::/64 cost 10
passive-interface Vlan2
passive-interface Vlan3
passive-interface Loopback0
timers spf 1 5
!
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - ipv6rocks@cisco.com
!
!
ipv6 access-list MGMT-IN
remark Permit MGMT only to Loopback0
permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::B222:2020 log-input
deny ipv6 any any log-input
!
ipv6 access-list VLAN2-v6-INGRESS
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
permit icmp 2001:DB8:CAFE:2::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::64
permit ipv6 2001:DB8:CAFE:2::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
ipv6 access-list VLAN3-v6-INGRESS
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::/64
permit icmp 2001:DB8:CAFE:3::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::64
permit ipv6 2001:DB8:CAFE:3::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C

Unauthorized access to this device and network is prohibited.

^C
!
line con 0
session-timeout 3
password 7 1317161C0C0916
logging synchronous
login local
transport output none

```

```

line vty 0 4
  session-timeout 3
  access-class MGMT-IN-v4 in
  password 7 071D2042490C0B
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local
  exec prompt timestamp
  transport input telnet ssh
!
no cns aaa enable
end

```

6k-core-1

```

upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-core-1
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$hHvm$ln0YNSdCb7lgHMmFz9Bi0
!
username cisco privilege 15 password 7 120B0419150E1E
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 071D2042490C0B
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc

```

```

vtp mode transparent
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
    key-string 7 1111
!
redundancy
  mode sso
  main-cpu
  auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
interface Loopback0
  ip address 10.122.10.3 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ipv6 address 2001:DB8:CAFE:6507::C333:3030/128
  no ipv6 redirects
  ipv6 ospf 1 area 0
!
interface Null0
  no ip unreachable
!
interface GigabitEthernet2/1
  description to 6k-agg-1
  dampening
  ip address 10.122.0.26 255.255.255.252
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7005::C333:3030/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0

```

```

wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp
!
interface GigabitEthernet2/2
description to 6k-agg-2
dampening
ip address 10.122.0.34 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7006::C333:3030/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp
!
interface GigabitEthernet2/3
description to 6k-core-2
dampening
ip address 10.122.0.21 255.255.255.252
no ip redirects
no ip unreachablees
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7009::C333:3030/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1

```

```

wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp
!
interface GigabitEthernet2/4
description to 6k-dist-1
dampening
ip address 10.122.0.77 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7000::C333:3030/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp
!
interface GigabitEthernet2/5
description to 6k-dist-2
dampening
ip address 10.122.0.81 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7002::C333:3030/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp

```

```

!
router eigrp 10
  passive-interface Loopback0
  network 10.0.0.0
  no auto-summary
  eigrp router-id 10.122.10.3
!
ip classless
!
no ip http server
ip http path bootflash:
!
ip access-list extended MGMT-IN-v4
  remark Permit v4MGMT only to Lo0
  permit tcp 10.120.0.0 0.0.255.255 any log-input
  permit tcp 10.121.0.0 0.0.255.255 any log-input
  permit tcp 10.122.0.0 0.0.255.255 any log-input
  deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
ipv6 router ospf 1
  router-id 10.122.10.3
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  passive-interface Loopback0
  timers spf 1 5
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - IPv6 ipv6rocks@cisco.com
!
ipv6 access-list MGMT-IN
  remark Permit MGMT only to Loopback0
  permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::C333:3030 log-input
  deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C
Unauthorized access to this device and/or network is prohibited.
^C
!
line con 0
  session-timeout 3
  exec-timeout 0 0
  password 7 095E4F071E0005
  logging synchronous
  login local
  transport output telnet ssh
line vty 0 4
  session-timeout 3
  access-class MGMT-IN-v4 in
  exec-timeout 30 0
  password 7 03165A05010A33
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local
  exec prompt timestamp
  transport input telnet ssh
!
no cns aaa enable
end

```

```

6k-core-2
upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-core-2
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$W031$3oHUmVfyHXW7fwdMWEBq01
!
username cisco privilege 15 secret 5 $1$vVK6$2OGixA1lDE.ufYBqWUrru/
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 111B180B101719
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
  key-string 7 1111
!
redundancy
mode sso
main-cpu

```

```

    auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
interface Loopback0
 ip address 10.122.10.4 255.255.255.255
 no ip redirects
 no ip unreachable
 no ip proxy-arp
 ipv6 address 2001:DB8:CAFE:6507::D444:4040/64
 no ipv6 redirects
 ipv6 ospf 1 area 0
!
interface Null0
 no ip unreachable
!
interface GigabitEthernet2/1
 description to 6k-agg-1
 dampening
 ip address 10.122.0.30 255.255.255.252
 no ip redirects
 no ip unreachable
 no ip proxy-arp
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 ip authentication mode eigrp 10 md5
 ip authentication key-chain eigrp 10 eigrp
 load-interval 30
 carrier-delay msec 0
 ipv6 address 2001:DB8:CAFE:7007::D444:4040/64
 no ipv6 redirects
 ipv6 nd suppress-ra
 ipv6 cef
 ipv6 ospf network point-to-point
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 ospf 1 area 0
 wrr-queue bandwidth 30 70
 wrr-queue queue-limit 40 30
 wrr-queue random-detect min-threshold 1 40 80
 wrr-queue random-detect min-threshold 2 70 80
 wrr-queue random-detect max-threshold 1 80 100
 wrr-queue random-detect max-threshold 2 80 100
 wrr-queue cos-map 1 1 1
 wrr-queue cos-map 1 2 0
 wrr-queue cos-map 2 1 2 3 4
 wrr-queue cos-map 2 2 6 7
 mls qos trust dscp
!
interface GigabitEthernet2/2
 description to 6k-agg-2
 dampening
 ip address 10.122.0.38 255.255.255.252

```

```

no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7008::D444:4040/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
interface GigabitEthernet2/3
description to 6k-core-1
dampening
ip address 10.122.0.22 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7009::D444:4040/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
interface GigabitEthernet2/4
description to 6k-dist-1

```

```

dampening
ip address 10.122.0.85 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7001::D444:4040/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
interface GigabitEthernet2/5
description to 6k-dist-2
dampening
ip address 10.122.0.89 255.255.255.252
no ip redirects
no ip unreachableables
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7003::D444:4040/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!

```

```

!
router eigrp 10
  passive-interface Loopback0
  network 10.0.0.0
  no auto-summary
  eigrp router-id 10.122.10.4
!
ip classless
!
no ip http server
!
ip access-list extended MGMT-IN-v4
  remark Permit v4MGMT only to Lo0
  permit tcp 10.120.0.0 0.0.255.255 any log-input
  permit tcp 10.121.0.0 0.0.255.255 any log-input
  permit tcp 10.122.0.0 0.0.255.255 any log-input
  deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
ipv6 router ospf 1
  router-id 10.122.10.4
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  passive-interface Loopback0
  timers spf 1 5
!
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - IPv6 ipv6rocks@cisco.com
!
ipv6 access-list MGMT-IN
  remark Permit MGMT only to Loopback0
  permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::D444:4040 log-input
  deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C
Unauthorized access to this device and/or network is prohibited.
^C
!
line con 0
  session-timeout 3
  exec-timeout 0 0
  password 7 120B0419150E1E
  logging synchronous
  login local
  transport output telnet ssh
line vty 0 4
  session-timeout 3
  access-class MGMT-IN-v4 in
  exec-timeout 30 0
  password 7 105C0817021200
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local
  exec prompt timestamp
  transport input telnet ssh
!
no cns aaa enable
end

```

Dual-Stack Model (DSM)—Routed Access

This section contains configurations only for access layer (3750-acc-1) and distribution layer (6k-dist-1/6k-dist-2).

3750-acc-1

```

version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service sequence-numbers
!
hostname 3750-acc-1
!
logging count
logging buffered 8192 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$RPWW$KR0qkU1kW0a7Pd1svvhgB.
!
username cisco privilege 15 secret 5 $1$KNSF$3XavijOuoVCq68RASxTYD1
no aaa new-model
clock timezone mst -7
switch 1 provision ws-c3750g-24ts
vtp domain ese-dc
vtp mode transparent
udld enable

udld message time 7

ip subnet-zero
no ip source-route
ip routing
ip icmp rate-limit unreachable 2000
ip telnet source-interface Vlan2
no ip domain-lookup
ip domain-name cisco.com
ip dhcp smart-relay
!
ip dhcp snooping vlan 2
ip dhcp snooping database flash:dhcp.txt
ip dhcp snooping database timeout 10
ip dhcp snooping
ip ftp source-interface Vlan2
ip ftp username cisco
ip ftp password 7 120B0419150E1E
ip tftp source-interface Vlan2
ip ssh time-out 60
ip ssh authentication-retries 2
ip ssh source-interface Vlan2
ip ssh version 2
ip arp inspection vlan 2
ip arp inspection validate src-mac
ip arp inspection log-buffer entries 100

```

```

ip arp inspection log-buffer logs 20 interval 120
login block-for 30 attempts 3 within 200
login delay 2
ipv6 mld snooping
ipv6 unicast-routing
!
mls qos map policed-dscp 0 10 18 24 25 34 to 8
mls qos map cos-dscp 0 8 16 24 32 46 48 56
mls qos srr-queue output cos-map queue 1 threshold 3 5
mls qos srr-queue output cos-map queue 2 threshold 1 2 4
mls qos srr-queue output cos-map queue 2 threshold 2 3
mls qos srr-queue output cos-map queue 2 threshold 3 6 7
mls qos srr-queue output cos-map queue 3 threshold 3 0
mls qos srr-queue output cos-map queue 4 threshold 3 1
mls qos srr-queue output dscp-map queue 1 threshold 3 46
mls qos srr-queue output dscp-map queue 2 threshold 1 16 18 20 22 25 32 34 36
mls qos srr-queue output dscp-map queue 2 threshold 1 38
mls qos srr-queue output dscp-map queue 2 threshold 2 24 26
mls qos srr-queue output dscp-map queue 2 threshold 3 48 56
mls qos srr-queue output dscp-map queue 3 threshold 3 0
mls qos srr-queue output dscp-map queue 4 threshold 1 8
mls qos srr-queue output dscp-map queue 4 threshold 3 10 12 14
mls qos queue-set output 1 threshold 2 70 80 100 100
mls qos queue-set output 1 threshold 4 40 100 100 100
mls qos
!
key chain eigrp
key 100
key-string 7 1111
!
crypto pki trustpoint TP-self-signed-3669881984
enrollment selfsigned
subject-name cn=IOS-Self-Signed-Certificate-3669881984
revocation-check none
rsaкеypair TP-self-signed-3669881984
!
crypto ca certificate chain TP-self-signed-3669881984
certificate self-signed 01
30820299 30820202 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
57312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
69666963 6174652D 33363639 38383139 38343124 30220609 2A864886 F70D0109
02161533 3735302D 65646765 2D312E63 6973636F 2E636F6D 301E170D 39333033
30313030 30313533 5A170D32 30303130 31303030 3030305A 3057312F 302D0603
55040313 26494F53 2D53656C 662D5369 676E6564 2D436572 74696669 63617465
2D333636 39383831 39383431 24302206 092A8648 86F70D01 09021615 33373530
2D656467 652D312E 63697363 6F2E636F 6D30819F 300D0609 2A864886 F70D0101
01050003 818D0030 81890281 8100BF7D E793E21B 6C1F75C2 16AAF9F9 C110D038
2D6D1DC9 04758DB8 7B3AD4C9 9F36A3B1 54983BEC 10FEA2D4 151D2783 5765C58A
A95E6364 CCBCF7F9 F4750437 AB8C00BF EFD54E88 6F650A5C 9563A309 247C6070
DC38A870 DEC5D4DA 765AD7A4 274B7649 36D876CA E28CF66C 77335F90 949DF258
E3E019BD 5E6801EC 9E15F980 C13D0203 010001A3 75307330 0F060355 1D130101
FF040530 030101FF 30200603 551D1104 19301782 15333735 302D6564 67652D31
2E636973 636F2E63 6F6D301F 0603551D 23041830 168014E0 18682CE1 D6A3EF2C
32A7C8D5 4DAFB9AA F11A0030 1D060355 1D0E0416 0414E018 682CE1D6 A3EF2C32
A7C8D54D AFB9AAF1 1A00300D 06092A86 4886F70D 01010405 00038181 00755909
C99DEB7F E05FC2A3 482558FA 33C292AE 7E4543E3 5BD6F32F 1D671B97 BC45B73B
85E954ED 7FC58F90 23A38132 24216CDB C978B3DD 9BCBC48E 519D01BF F4CEBB82
07834C2D D82CA163 8E638214 5B5C277D 5E7DD52E 56172675 BD563769 590E4DC6
39AD9BF8 CDBBA241 E5E2C666 5CAE912E 40DC2150 A1CE39B4 D8101D33 A8
quit
!
errdisable recovery cause link-flap
errdisable recovery interval 60
no file verify auto

```

```

!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree portfast bpduguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
vlan 2
  name ACCESS-DATA-2
!
class-map match-all DVLAN-PC-VIDEO
  match access-group name DVLAN-PC-VIDEO
class-map match-all DVLAN-Transactional-Data
  match access-group name DVLAN-Transactional-Data
class-map match-all DVLAN-Mission-Critical-Data
  match access-group name DVLAN-Mission-Critical-Data
class-map match-all DVLAN-Bulk-Data
  match access-group name DVLAN-Bulk-Data
!
!
policy-map DATA
  class DVLAN-PC-VIDEO
    set dscp af41
    police 48000 8000 exceed-action policed-dscp-transmit
  class DVLAN-Mission-Critical-Data
    set dscp 25
    police 5000000 8000 exceed-action policed-dscp-transmit
  class DVLAN-Transactional-Data
    set dscp af21
    police 5000000 8000 exceed-action policed-dscp-transmit
  class DVLAN-Bulk-Data
    set dscp af11
    police 5000000 8000 exceed-action policed-dscp-transmit
!
!
interface Null0
  no ip unreachable
!
interface GigabitEthernet1/0/5
  description to PC-DATA-ONLY
  switchport access vlan 2
  switchport mode access
  switchport port-security maximum 3
  switchport port-security
  switchport port-security aging time 2
  switchport port-security violation restrict
  switchport port-security aging type inactivity
  ip arp inspection limit rate 100
  service-policy input DATA
  load-interval 30
  srr-queue bandwidth share 1 70 25 5
  srr-queue bandwidth shape 3 0 0 0
  priority-queue out
  spanning-tree portfast
  spanning-tree bpduguard enable
  ip verify source
  ip dhcp snooping limit rate 100
!
interface GigabitEthernet1/0/25
  description To 6k-dist-1
  no switchport
  dampening
  ip address 10.120.0.2 255.255.255.252

```

```

no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
ipv6 address 2001:DB8:CAFE:700A::CAC1:3750/64
ipv6 nd suppress-ra
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2
no ipv6 redirects
mls qos trust dscp
no cdp enable
!
interface GigabitEthernet1/0/26
description To 6k-dist-2
no switchport
dampening
ip address 10.120.0.10 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
srr-queue bandwidth share 1 70 25 5
srr-queue bandwidth shape 3 0 0 0
priority-queue out
ipv6 address 2001:DB8:CAFE:700C::CAC1:3750/64
ipv6 nd suppress-ra
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2
no ipv6 redirects
mls qos trust dscp
no cdp enable
!
interface Vlan2
ip address 10.120.2.1 255.255.255.0
ip helper-address 10.121.10.7
no ip redirects
no ip unreachable
no ip proxy-arp
no ip mroute-cache
load-interval 30
ipv6 address 2001:DB8:CAFE:2::CAC1:3750/64
ipv6 ospf 1 area 2
no ipv6 redirects
!
router eigrp 10
passive-interface Vlan2
network 10.120.0.0 0.0.255.255

```

```

no auto-summary
eigrp router-id 10.120.2.1
eigrp stub connected
!
ip classless
no ip http server
no ip http secure-server
!
ip access-list extended DVLAN-Bulk-Data
  permit tcp any any eq 143
  permit tcp any any eq 220
ip access-list extended DVLAN-Mission-Critical-Data
  permit tcp any any range 3200 3203
  permit tcp any any eq 3600
  permit tcp any any range 2000 2002
ip access-list extended DVLAN-PC-VIDEO
  permit udp any any range 16384 32767
ip access-list extended DVLAN-Transactional-Data
  permit tcp any any eq 1352
ip access-list extended MGMT-IN-v4
  permit tcp 10.120.0.0 0.0.255.255 any log-input
  permit tcp 10.121.0.0 0.0.255.255 any log-input
  permit tcp 10.122.0.0 0.0.255.255 any log-input
  deny ip any any log-input
!
logging source-interface Vlan2
logging 10.121.11.9
no cdp run
!
ipv6 router ospf 1
  router-id 10.120.2.1
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  area 2 stub no-summary
  passive-interface Vlan2
  timers spf 1 5
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - IPv6 ipv6rocks@cisco.com
!
ipv6 access-list MGMT-IN
  remark Permit MGMT only to VLAN2
  permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:2::CAC1:3750 log-input
  deny ipv6 any any log-input
!
control-plane
!
banner login ^C
Unauthorized access to this device and/or network is prohibited.
^C
!
line con 0
  session-timeout 3
  exec-timeout 0 0
  password 7 06140E2F4B4B1B
  logging synchronous
  login local
  transport output telnet ssh
line vty 0 4
  session-timeout 3
  password 7 06140E2F4B4B1B
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local

```

```

exec prompt timestamp
transport input telnet ssh
line vty 5 15
session-timeout 3
password 7 0101070A5C0E14
ipv6 access-class MGMT-IN in
logging synchronous
login local
exec prompt timestamp
transport input telnet ssh
!
end

```

6k-dist-1

```

upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service sequence-numbers
service counters max age 5
!
hostname 6k-dist-1
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$EZB7$nbqNvH9QH4qo3zpzGzZog6/
!
username cisco privilege 15 secret 5 $1$VwBR$aXCTEAusOPhpmcfg7nOn01
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 121A0C041104
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent

```

```

mls ip cef load-sharing full
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
    key-string 7 1111
!
redundancy
  mode sso
  main-cpu
    auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 2-3 priority 24576
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
interface Loopback0
  ip address 10.122.10.9 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ipv6 address 2001:DB8:CAFE:6507::A111:1010/128
  no ipv6 redirects
  ipv6 ospf 1 area 0
!
interface Null0
  no ip unreachable
!
interface TenGigabitEthernet1/1
  description to 6k-dist-2
  dampening
  ip address 10.120.0.13 255.255.255.252
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7004::A111:1010/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3

```

```

ipv6 ospf 1 area 2
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
!
interface GigabitEthernet3/1
description to 3750-acc-1
dampening
ip address 10.120.0.1 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:700A::A111:1010/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable
spanning-tree guard root
!
interface GigabitEthernet3/2
description to 3750-acc-2
dampening
ip address 10.120.0.5 255.255.255.252
no ip redirects
no ip unreachable

```

```

no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:700B::A111:1010/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable
spanning-tree guard root
!
interface GigabitEthernet4/1
description to 6k-core-1
dampening
ip address 10.122.0.78 255.255.255.252
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
ip summary-address eigrp 10 10.120.0.0 255.255.0.0 5
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7000::A111:1010/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4

```

```

wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
interface GigabitEthernet4/2
description to 6k-core-2
dampening
ip address 10.122.0.86 255.255.255.252
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
ip summary-address eigrp 10 10.120.0.0 255.255.0.0 5
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7001::A111:1010/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
router eigrp 10
passive-interface Loopback0
network 10.120.0.0 0.0.255.255
network 10.122.0.0 0.0.0.255
distribute-list DEFAULT out GigabitEthernet3/1
distribute-list DEFAULT out GigabitEthernet3/2
no auto-summary
eigrp router-id 10.122.10.9
!
ip classless
!
no ip http server
!
ip access-list standard DEFAULT
permit 0.0.0.0
!
ip access-list extended MGMT-IN-v4
remark Permit v4MGMT only to Lo0
permit tcp 10.120.0.0 0.0.255.255 any log-input
permit tcp 10.121.0.0 0.0.255.255 any log-input
permit tcp 10.122.0.0 0.0.255.255 any log-input
deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
ipv6 router ospf 1
router-id 10.122.10.9
log-adjacency-changes

```

```

auto-cost reference-bandwidth 10000
area 2 stub no-summary
area 2 range 2001:DB8:CAFE:2::/64 cost 10
area 2 range 2001:DB8:CAFE:3::/64 cost 10
area 2 range 2001:DB8:CAFE:7004::/64 cost 10
area 2 range 2001:DB8:CAFE:700A::/64 cost 10
area 2 range 2001:DB8:CAFE:700B::/64 cost 10
passive-interface Loopback0
timers spf 1 5
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - ipv6rocks@cisco.com
!
ipv6 access-list MGMT-IN
 remark Permit MGMT only to Loopback0
 permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::A111:1010 log-input
 deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C
Unauthorized access to this device and/or network is prohibited!
^C
!
line con 0
 session-timeout 3
 password 7 06140E2F4B4B1B
 logging synchronous
 login local
 transport output none
line vty 0 4
 session-timeout 3
 access-class MGMT-IN-v4 in
 password 7 08334D400E1C17
 ipv6 access-class MGMT-IN in
 logging synchronous
 login local
 exec prompt timestamp
 transport input telnet ssh
!
no cns aaa enable
end

```

6k-dist-2

```

upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-dist-2
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin

```

```

logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$0UKB$YN.zu/IQt0ivlkqLsdxtY0
!
username cisco privilege 15 secret 5 $1$SUYV$ShOXyooFT..fL/4tyOnAN1
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 01100F175804
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip cef load-sharing full
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
    key-string 7 1111
!
redundancy
  mode sso
  main-cpu
  auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 2-3 priority 28672
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000

```

```

!
interface Loopback0
 ip address 10.122.10.10 255.255.255.255
 no ip redirects
 no ip unreachableables
 no ip proxy-arp
 ipv6 address 2001:DB8:CAFE:6507::B222:2020/128
 no ipv6 redirects
 ipv6 ospf 1 area 0
!
interface Null0
 no ip unreachableables
!
interface TenGigabitEthernet1/1
 description to 6k-dist-1
 dampening
 ip address 10.120.0.14 255.255.255.252
 no ip redirects
 no ip unreachableables
 no ip proxy-arp
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 ip authentication mode eigrp 10 md5
 ip authentication key-chain eigrp 10 eigrp
 load-interval 30
 carrier-delay msec 0
 ipv6 address 2001:DB8:CAFE:7004::B222:2020/64
 no ipv6 redirects
 ipv6 nd suppress-ra
 ipv6 cef
 ipv6 ospf network point-to-point
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 ospf 1 area 2
 wrr-queue bandwidth 5 25 70
 wrr-queue queue-limit 5 25 40
 wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
 wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
 wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
 wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
 wrr-queue cos-map 1 1 1
 wrr-queue cos-map 2 1 0
 wrr-queue cos-map 3 1 4
 wrr-queue cos-map 3 2 2
 wrr-queue cos-map 3 3 3
 wrr-queue cos-map 3 4 6
 wrr-queue cos-map 3 5 7
 mls qos trust dscp
!
interface GigabitEthernet3/1
 description to 3750-acc-1
 dampening
 ip address 10.120.0.9 255.255.255.252
 no ip redirects
 no ip unreachableables
 no ip proxy-arp
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 ip authentication mode eigrp 10 md5
 ip authentication key-chain eigrp 10 eigrp
 load-interval 30
 carrier-delay msec 0

```

```

ipv6 address 2001:DB8:CAFE:700C::B222:2020/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable
!
interface GigabitEthernet3/2
description to 3750-acc-1
dampening
ip address 10.120.0.17 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:700D::B222:2020/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 2
wrr-queue bandwidth 5 25 70
wrr-queue queue-limit 5 25 40
wrr-queue random-detect min-threshold 1 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 2 80 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 3 50 60 70 80 90 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 60 70 80 90 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 4
wrr-queue cos-map 3 2 2
wrr-queue cos-map 3 3 3
wrr-queue cos-map 3 4 6
wrr-queue cos-map 3 5 7
mls qos trust dscp
no cdp enable

```

```

!
interface GigabitEthernet4/1
  description to 6k-core-1
  dampening
  ip address 10.122.0.82 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7002::B222:2020/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  wrr-queue bandwidth 30 70
  wrr-queue queue-limit 40 30
  wrr-queue random-detect min-threshold 1 40 80
  wrr-queue random-detect min-threshold 2 70 80
  wrr-queue random-detect max-threshold 1 80 100
  wrr-queue random-detect max-threshold 2 80 100
  wrr-queue cos-map 1 1 1
  wrr-queue cos-map 1 2 0
  wrr-queue cos-map 2 1 2 3 4
  wrr-queue cos-map 2 2 6 7
  mls qos trust dscp
!
interface GigabitEthernet4/2
  description to 6k-core-2
  dampening
  ip address 10.122.0.90 255.255.255.252
  no ip redirects
  no ip proxy-arp
  ip hello-interval eigrp 10 1
  ip hold-time eigrp 10 3
  ip authentication mode eigrp 10 md5
  ip authentication key-chain eigrp 10 eigrp
  load-interval 30
  carrier-delay msec 0
  ipv6 address 2001:DB8:CAFE:7003::B222:2020/64
  no ipv6 redirects
  ipv6 nd suppress-ra
  ipv6 cef
  ipv6 ospf network point-to-point
  ipv6 ospf hello-interval 1
  ipv6 ospf dead-interval 3
  ipv6 ospf 1 area 0
  wrr-queue bandwidth 30 70
  wrr-queue queue-limit 40 30
  wrr-queue random-detect min-threshold 1 40 80
  wrr-queue random-detect min-threshold 2 70 80
  wrr-queue random-detect max-threshold 1 80 100
  wrr-queue random-detect max-threshold 2 80 100
  wrr-queue cos-map 1 1 1
  wrr-queue cos-map 1 2 0
  wrr-queue cos-map 2 1 2 3 4
  wrr-queue cos-map 2 2 6 7
  mls qos trust dscp

```

```

!
router eigrp 10
  passive-interface Loopback0
  network 10.120.0.0 0.0.255.255
  network 10.122.0.0 0.0.0.255
  distribute-list DEFAULT out GigabitEthernet3/1
  distribute-list DEFAULT out GigabitEthernet3/2
  no auto-summary
  eigrp router-id 10.122.10.10
!
ip classless
!
no ip http server
!
ip access-list standard DEFAULT
  permit 0.0.0.0
!
ip access-list extended MGMT-IN-v4
  remark Permit v4MGMT only to Lo0
  permit tcp 10.120.0.0 0.0.255.255 any log-input
  permit tcp 10.121.0.0 0.0.255.255 any log-input
  permit tcp 10.122.0.0 0.0.255.255 any log-input
  deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
ipv6 router ospf 1
  router-id 10.122.10.10
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  area 2 stub no-summary
  area 2 range 2001:DB8:CAFE:2::/64 cost 10
  area 2 range 2001:DB8:CAFE:3::/64 cost 10
  area 2 range 2001:DB8:CAFE:7004::/64 cost 10
  area 2 range 2001:DB8:CAFE:700A::/64 cost 10
  area 2 range 2001:DB8:CAFE:700B::/64 cost 10
  passive-interface Loopback0
  timers spf 1 5
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - ipv6rocks@cisco.com
!
ipv6 access-list MGMT-IN
  remark Permit MGMT only to Loopback0
  permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::B222:2020 log-input
  deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C

Unauthorized access to this device and network is prohibited.

^C
!
line con 0
  session-timeout 3
  password 7 1317161C0C0916
  logging synchronous
  login local
  transport output telnet ssh
line vty 0 4

```

```

session-timeout 3
access-class MGMT-IN-v4 in
password 7 071D2042490C0B
ipv6 access-class MGMT-IN in
logging synchronous
login local
exec prompt timestamp
transport input telnet ssh
!
no cns aaa enable
end

```

Hybrid Model Example 1 (HME1)

Configurations are shown for the core layer only because all other layers are IPv4 only and those configurations match the IPv4 configurations in the DSM (Slight address changes in distribution layer).

6k-core-1

```

upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-core-1
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$hHvm$ln0YNSdCb7lzgHmFz9Bi0
!
username cisco privilege 15 password 7 120B0419150E1E
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 071D2042490C0B
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress

```

```

ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
  key-string 7 1111
!
redundancy
  mode sso
  main-cpu
  auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
class-map match-all CAMPUS-BULK-DATA
  match access-group name BULK-APPS
class-map match-all CAMPUS-TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-APPS
!
policy-map IPv6-ISATAP-MARK
  class CAMPUS-BULK-DATA
    set dscp af11
  class CAMPUS-TRANSACTIONAL-DATA
    set dscp af21
  class class-default
    set dscp default
!
interface Loopback0
  ip address 10.122.10.3 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ipv6 address 2001:DB8:CAFE:6507::C333:3030/128
  no ipv6 redirects
  ipv6 ospf 1 area 0
!
interface Loopback2
  description Tunnel source for ISATAP-VLAN2
  ip address 10.122.10.102 255.255.255.255
  no ip redirects
  no ip unreachable

```

```

no ip proxy-arp
!
interface Loopback3
description Tunnel source for ISATAP-VLAN3
ip address 10.122.10.103 255.255.255.255
no ip redirects
no ip unreachable
no ip proxy-arp
!
interface Tunnel2
description ISATAP VLAN2
no ip address
no ip redirects
ipv6 address 2001:DB8:CAFE:2::/64 eui-64
ipv6 traffic-filter SOURCE-ISATAP-2 in
no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback2
tunnel mode ipv6ip isatap
!
interface Tunnel3
description ISATAP VLAN3
no ip address
no ip redirects
ipv6 address 2001:DB8:CAFE:3::/64 eui-64
ipv6 traffic-filter SOURCE-ISATAP-3 in
no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback3
tunnel mode ipv6ip isatap
!
interface Null0
no ip unreachable
!
interface GigabitEthernet1/1
description to 3750-dist-1
dampening
ip address 10.122.0.41 255.255.255.252
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
interface GigabitEthernet1/2
description to 3750-dist-2
dampening
ip address 10.122.0.45 255.255.255.252
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
interface GigabitEthernet2/1
description to 6k-agg-1

```

```

dampening
ip address 10.122.0.26 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7005::C333:3030/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp
service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet2/2
description to 6k-agg-2
dampening
ip address 10.122.0.34 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7006::C333:3030/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp
service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet2/3
description to 6k-core-2
dampening
ip address 10.122.0.21 255.255.255.252

```

```

no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7009::C333:3030/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 30 70
wrr-queue threshold 1 40 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4 6 7
wrr-queue cos-map 2 2 5
mls qos trust dscp
!
router eigrp 10
  passive-interface Loopback0
  passive-interface Loopback2
  passive-interface Loopback3
  network 10.0.0.0
  no auto-summary
  eigrp router-id 10.122.10.3
!
ip classless
!
no ip http server
!
ip access-list extended MGMT-IN-v4
  remark Permit v4MGMT only to Lo0
  permit tcp 10.120.0.0 0.0.255.255 any log-input
  permit tcp 10.121.0.0 0.0.255.255 any log-input
  permit tcp 10.122.0.0 0.0.255.255 any log-input
  deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
ipv6 router ospf 1
  router-id 10.122.10.3
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  area 2 range 2001:DB8:CAFE:2::/64 cost 10
  area 2 range 2001:DB8:CAFE:3::/64 cost 10
  passive-interface Loopback0
  passive-interface Loopback2
  passive-interface Loopback3
  passive-interface Tunnel2
  passive-interface Tunnel3
  timers spf 1 5
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - IPv6 ipv6rocks@cisco.com
!
ipv6 access-list SOURCE-ISATAP-2

```

```

remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
permit icmp 2001:DB8:CAFE:2::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::64
permit ipv6 2001:DB8:CAFE:2::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
ipv6 access-list SOURCE-ISATAP-3
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::/64
permit icmp 2001:DB8:CAFE:3::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::64
permit ipv6 2001:DB8:CAFE:3::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
ipv6 access-list BULK-APPS
permit tcp any any eq ftp
permit tcp any any eq ftp-data
!
ipv6 access-list TRANSACTIONAL-APPS
permit tcp any any eq telnet
permit tcp any any eq 22
!
ipv6 access-list MGMT-IN
remark Permit MGMT only to Loopback0
permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::C333:3030 log-input
deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C
Unauthorized access to this device and/or network is prohibited.
^C
!
line con 0
session-timeout 3
exec-timeout 0 0
password 7 095E4F071E0005
logging synchronous
login local
transport output telnet ssh
line vty 0 4
session-timeout 3
access-class MGMT-IN-v4 in
exec-timeout 30 0
password 7 03165A05010A33
ipv6 access-class MGMT-IN in
logging synchronous
login local
exec prompt timestamp
transport input telnet ssh
!
no cns aaa enable
end

```

6k-core-2

```

upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-core-2
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$W031$3oHUmVfyHXW7fwdMWEBq01
!
username cisco privilege 15 secret 5 $1$vVK6$2OGixA11DE.ufYBqWUrru/
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 111B180B101719
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
ipv6 multicast-routing
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
  key-string 7 1111
!
redundancy

```

```

mode sso
main-cpu
  auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
class-map match-all CAMPUS-BULK-DATA
  match access-group name BULK-APPS
class-map match-all CAMPUS-TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-APPS
!
policy-map IPv6-ISATAP-MARK
  class CAMPUS-BULK-DATA
    set dscp af11
  class CAMPUS-TRANSACTIONAL-DATA
    set dscp af21
  class class-default
    set dscp default
!
interface Loopback0
ip address 10.122.10.4 255.255.255.255
no ip redirects
no ip unreachable
no ip proxy-arp
ipv6 address 2001:DB8:CAFE:6507::D444:4040/64
no ipv6 redirects
ipv6 ospf 1 area 0
!
interface Loopback2
description Tunnel source for ISATAP-VLAN2
ip address 10.122.10.102 255.255.255.255
no ip redirects
no ip unreachable
no ip proxy-arp
delay 1000
!
interface Loopback3
description Tunnel source for ISATAP-VLAN3
ip address 10.122.10.103 255.255.255.255
no ip redirects
no ip unreachable
no ip proxy-arp
delay 1000
!
interface Tunnel2
description ISATAP VLAN2
no ip address
no ip redirects
ipv6 address 2001:DB8:CAFE:2::/64 eui-64
ipv6 traffic-filter SOURCE-ISATAP-2 in
no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef

```

```

    ipv6 ospf 1 area 2
    tunnel source Loopback2
    tunnel mode ipv6ip isatap
    !
interface Tunnel3
    description ISATAP VLAN3
    no ip address
    no ip redirects
    ipv6 address 2001:DB8:CAFE:3::/64 eui-64
    ipv6 traffic-filter SOURCE-ISATAP-3 in
    no ipv6 redirects
    no ipv6 nd suppress-ra
    ipv6 cef
    ipv6 ospf 1 area 2
    tunnel source Loopback3
    tunnel mode ipv6ip isatap
    !
interface Null0
    no ip unreachable
    !
interface GigabitEthernet1/1
    description to 3750-dist-1
    dampening
    ip address 10.122.0.49 255.255.255.252
    no ip redirects
    no ip proxy-arp
    ip hello-interval eigrp 10 1
    ip hold-time eigrp 10 3
    load-interval 30
    carrier-delay msec 0
    mls qos trust dscp
    !
interface GigabitEthernet1/2
    description to 3750-dist-2
    dampening
    ip address 10.122.0.53 255.255.255.252
    no ip redirects
    no ip proxy-arp
    ip hello-interval eigrp 10 1
    ip hold-time eigrp 10 3
    load-interval 30
    carrier-delay msec 0
    mls qos trust dscp
    !
interface GigabitEthernet2/1
    description to 6k-agg-1
    dampening
    ip address 10.122.0.30 255.255.255.252
    no ip redirects
    no ip unreachable
    no ip proxy-arp
    ip hello-interval eigrp 10 1
    ip hold-time eigrp 10 3
    ip authentication mode eigrp 10 md5
    ip authentication key-chain eigrp 10 eigrp
    load-interval 30
    carrier-delay msec 0
    ipv6 address 2001:DB8:CAFE:7007::D444:4040/64
    no ipv6 redirects
    ipv6 nd suppress-ra
    ipv6 cef
    ipv6 ospf network point-to-point
    ipv6 ospf hello-interval 1
    ipv6 ospf dead-interval 3

```

```

ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
service-policy output IPv6-ISATAP-MARK

!
interface GigabitEthernet2/2
description to 6k-agg-2
dampening
ip address 10.122.0.38 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7008::D444:4040/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
service-policy output IPv6-ISATAP-MARK

!
interface GigabitEthernet2/3
description to 6k-core-1
dampening
ip address 10.122.0.22 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:7009::D444:4040/64
no ipv6 redirects

```

```

ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
wrr-queue bandwidth 30 70
wrr-queue queue-limit 40 30
wrr-queue random-detect min-threshold 1 40 80
wrr-queue random-detect min-threshold 2 70 80
wrr-queue random-detect max-threshold 1 80 100
wrr-queue random-detect max-threshold 2 80 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 1 2 0
wrr-queue cos-map 2 1 2 3 4
wrr-queue cos-map 2 2 6 7
mls qos trust dscp
!
router eigrp 10
  passive-interface Loopback0
  passive-interface Loopback2
  passive-interface Loopback3
  network 10.0.0.0
  no auto-summary
  eigrp router-id 10.122.10.4
!
ip classless
!
no ip http server
!
ip access-list extended MGMT-IN-v4
  remark Permit v4MGMT only to Lo0
  permit tcp 10.120.0.0 0.0.255.255 any log-input
  permit tcp 10.121.0.0 0.0.255.255 any log-input
  permit tcp 10.122.0.0 0.0.255.255 any log-input
  deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
ipv6 router ospf 1
  router-id 10.122.10.4
  log-adjacency-changes
  auto-cost reference-bandwidth 10000
  area 2 range 2001:DB8:CAFE:2::/64 cost 20
  area 2 range 2001:DB8:CAFE:3::/64 cost 20
  passive-interface Loopback0
  passive-interface Loopback2
  passive-interface Loopback3
  passive-interface Tunnel2
  passive-interface Tunnel3
  timers spf 1 5
!
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - IPv6 ipv6rocks@cisco.com
!
ipv6 access-list SOURCE-ISATAP-2
  remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
  permit icmp 2001:DB8:CAFE:2::/64 any
  remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::64
  permit ipv6 2001:DB8:CAFE:2::/64 any
  remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
  permit icmp FE80::/10 any
  remark DENY ALL OTHER IPv6 PACKETS AND LOG

```

```

deny ipv6 any any log-input
!
ipv6 access-list SOURCE-ISATAP-3
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::/64
permit icmp 2001:DB8:CAFE:3::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::64
permit ipv6 2001:DB8:CAFE:3::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
ipv6 access-list BULK-APPS
permit tcp any any eq ftp
permit tcp any any eq ftp-data
!
ipv6 access-list TRANSACTIONAL-APPS
permit tcp any any eq telnet
permit tcp any any eq 22
!
ipv6 access-list MGMT-IN
remark Permit MGMT only to Loopback0
permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::D444:4040 log-input
deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C
Unauthorized access to this device and/or network is prohibited.
^C
!
line con 0
session-timeout 3
exec-timeout 0 0
password 7 120B0419150E1E
logging synchronous
login local
transport output telnet ssh
line vty 0 4
session-timeout 3
access-class MGMT-IN-v4 in
exec-timeout 30 0
password 7 105C0817021200
ipv6 access-class MGMT-IN in
logging synchronous
login local
exec prompt timestamp
transport input telnet ssh
!
no cns aaa enable
end

```

Service Block Model (SBM)

Configurations are shown for the service block switches only (6k-sb-1/6k-sb-2) as all other configurations for the access, distribution and core layers all use the IPv4 configurations shown in the DSM section.

6k-sb-1

```

upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-sb-1
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$EZB7$nbqNvH9QH4qo3zpGsZog6/
!
username cisco privilege 15 secret 5 $1$VwbR$aXCTEAusOPhpmcfcg7nOn01
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 121A0C041104
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com
ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip cef load-sharing full
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
  key-string 7 1111
!
redundancy

```

```

mode sso
main-cpu
  auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
class-map match-all CAMPUS-BULK-DATA
  match access-group name BULK-APPS
class-map match-all CAMPUS-TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-APPS
!
policy-map IPv6-ISATAP-MARK
  class CAMPUS-BULK-DATA
    set dscp af11
  class CAMPUS-TRANSACTIONAL-DATA
    set dscp af21
  class class-default
    set dscp default
!
interface Loopback0
  description Tunnel source for 6k-agg-1
  ip address 10.122.10.9 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ipv6 address 2001:DB8:CAFE:6507::A111:1010/128
  no ipv6 redirects
  ipv6 ospf 1 area 0
!
interface Loopback1
  description Tunnel source for 6k-agg-2
  ip address 10.122.10.19 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
!
interface Loopback2
  description Tunnel source for ISATAP-VLAN2
  ip address 10.122.10.102 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
!
interface Loopback3
  description Tunnel source for ISATAP-VLAN3
  ip address 10.122.10.103 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
!
interface Tunnel0
  description tunnel to 6k-agg-1
  no ip address

```

```

ipv6 address 2001:DB8:CAFE:6501::A111:1010/64
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
tunnel source Loopback0
tunnel destination 10.122.10.1
tunnel mode ipv6ip
service-policy output IPv6-ISATAP-MARK
!
interface Tunnel1
description tunnel to 6k-agg-2
no ip address
ipv6 address 2001:DB8:CAFE:6502::A111:1010/64
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
tunnel source Loopback1
tunnel destination 10.122.10.2
tunnel mode ipv6ip
service-policy output IPv6-ISATAP-MARK
!
interface Tunnel2
description ISATAP VLAN2
no ip address
no ip redirects
ipv6 address 2001:DB8:CAFE:2::/64 eui-64
ipv6 traffic-filter SOURCE-ISATAP-2 in
no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback2
tunnel mode ipv6ip isatap
!
interface Tunnel3
description ISATAP VLAN3
no ip address
no ip redirects
ipv6 address 2001:DB8:CAFE:3::/64 eui-64
ipv6 traffic-filter SOURCE-ISATAP-3 in
no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback3
tunnel mode ipv6ip isatap
!
interface TenGigabitEthernet1/1
description to 6k-sb-2
dampening
ip address 10.122.0.93 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3

```

```

ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:6505::A111:1010/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
mls qos trust dscp
service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet4/1
description to 6k-core-1
dampening
ip address 10.122.0.78 255.255.255.252
ip access-group 101 in
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
interface GigabitEthernet4/2
description to 6k-core-2
dampening
ip address 10.122.0.86 255.255.255.252
ip access-group 101 in
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
router eigrp 10
passive-interface Loopback0
passive-interface Loopback1
passive-interface Loopback2
passive-interface Loopback3
network 10.0.0.0
no auto-summary
eigrp router-id 10.122.10.9
!
ip classless
!
no ip http server
!
ip access-list extended MGMT-IN-v4
remark Permit v4MGMT only to Lo0
permit tcp 10.120.0.0 0.0.255.255 any log-input
permit tcp 10.121.0.0 0.0.255.255 any log-input
permit tcp 10.122.0.0 0.0.255.255 any log-input
deny ip any any log-input

```

```

!
logging source-interface Loopback0
logging 10.121.11.9
access-list 101 permit 41 10.120.2.0 0.0.0.255 host 10.122.10.102
access-list 101 permit 41 10.120.3.0 0.0.0.255 host 10.122.10.103
access-list 101 permit 41 host 10.122.10.1 host 10.122.10.9
access-list 101 permit 41 host 10.122.10.2 host 10.122.10.19
access-list 101 permit 41 host 10.122.10.11 host 10.122.10.10
access-list 101 permit 41 host 10.122.10.12 host 10.122.10.20
access-list 101 deny 41 any any
access-list 101 permit ip any any
ipv6 router ospf 1
router-id 10.122.10.9
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 range 2001:DB8:CAFE:2::/64 cost 10
area 2 range 2001:DB8:CAFE:3::/64 cost 10
passive-interface Loopback0
passive-interface Tunnel2
passive-interface Tunnel3
timers spf 1 5
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - ipv6rocks@cisco.com
!
ipv6 access-list SOURCE-ISATAP-2
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
permit icmp 2001:DB8:CAFE:2::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::64
permit ipv6 2001:DB8:CAFE:2::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
ipv6 access-list SOURCE-ISATAP-3
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::/64
permit icmp 2001:DB8:CAFE:3::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::64
permit ipv6 2001:DB8:CAFE:3::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input
!
ipv6 access-list BULK-APPS
permit tcp any any eq ftp
permit tcp any any eq ftp-data
!
ipv6 access-list TRANSACTIONAL-APPS
permit tcp any any eq telnet
permit tcp any any eq 22
!
ipv6 access-list MGMT-IN
remark Permit MGMT only to Loopback0
permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::A111:1010 log-input
deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C
Unauthorized access to this device and/or network is prohibited!

```

```

^C
!
line con 0
  session-timeout 3
  password 7 06140E2F4B4B1B
  logging synchronous
  login local
  transport output none
line vty 0 4
  session-timeout 3
  access-class MGMT-IN-v4 in
  password 7 08334D400E1C17
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local
  exec prompt timestamp
  transport input telnet ssh
!
no cns aaa enable
end

```

6k-sb-2

```

upgrade fpd auto
version 12.2
no service pad
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
service internal
service sequence-numbers
service counters max age 5
!
hostname 6k-sb-2
!
boot system flash sup-bootdisk:s3223-adventerprisek9_wan-mz.122-18.SXF5.bin
logging buffered 64000 debugging
logging rate-limit 5
no logging console
enable secret 5 $1$0UKB$YN.zu/IQt0ivlkqLsdxtY0
!
username cisco privilege 15 secret 5 $1$SUYV$ShOXyooFT..fL/4tyOnAN1
no aaa new-model
clock timezone mst -7
ip subnet-zero
no ip source-route
ip icmp rate-limit unreachable 2000
!
ip ftp source-interface Loopback0
ip ftp username cisco
ip ftp password 7 01100F175804
ip tftp source-interface Loopback0
no ip bootp server
ip telnet source-interface Loopback0
ip ssh time-out 30
ip ssh authentication-retries 2
ip ssh source-interface Loopback0
ip ssh version 2
no ip domain-lookup
ip domain-name cisco.com

```

```

ipv6 unicast-routing
ipv6 mfib hardware-switching replication-mode ingress
udld enable

udld message time 7

vtp domain ese-dc
vtp mode transparent
mls ip cef load-sharing full
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
mls qos
mls rate-limit unicast ip icmp unreachable acl-drop 0
no mls acl tcam share-global
mls cef error action freeze
!
key chain eigrp
  key 100
    key-string 7 1111
!
!
redundancy
  mode sso
  main-cpu
  auto-sync running-config
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
environment temperature-controlled
system flowcontrol bus auto
diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy descending
vlan dot1q tag native
vlan access-log ratelimit 2000
!
class-map match-all CAMPUS-BULK-DATA
  match access-group name BULK-APPS
class-map match-all CAMPUS-TRANSACTIONAL-DATA
  match access-group name TRANSACTIONAL-APPS
!
policy-map IPv6-ISATAP-MARK
  class CAMPUS-BULK-DATA
    set dscp af11
  class CAMPUS-TRANSACTIONAL-DATA
    set dscp af21
  class class-default
    set dscp default
!
interface Loopback0
  description Tunnel source for 6k-agg-1
  ip address 10.122.10.10 255.255.255.255
  no ip redirects
  no ip unreachable
  no ip proxy-arp
  ipv6 address 2001:DB8:CAFE:6507::B222:2020/128
  no ipv6 redirects
  ipv6 ospf 1 area 0
!
interface Loopback1

```

```

description Tunnel source for 6k-agg-2
ip address 10.122.10.20 255.255.255.255
no ip redirects
no ip unreachable
no ip proxy-arp
!
interface Loopback2
description Tunnel source for ISATAP-VLAN2
ip address 10.122.10.102 255.255.255.255
no ip redirects
no ip unreachable
no ip proxy-arp
delay 1000
!
interface Loopback3
description Tunnel source for ISATAP-VLAN3
ip address 10.122.10.103 255.255.255.255
no ip redirects
no ip unreachable
no ip proxy-arp
delay 1000
!
interface Tunnel0
description tunnel to 6k-agg-1
no ip address
load-interval 30
ipv6 address 2001:DB8:CAFE:6503::B222:2020/64
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
tunnel source Loopback0
tunnel destination 10.122.10.11
tunnel mode ipv6ip
service-policy output IPv6-ISATAP-MARK
!
interface Tunnel1
description tunnel to 6k-agg-2
no ip address
load-interval 30
ipv6 address 2001:DB8:CAFE:6504::B222:2020/64
no ipv6 redirects
ipv6 nd reachable-time 5000
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
tunnel source Loopback1
tunnel destination 10.122.10.12
tunnel mode ipv6ip
service-policy output IPv6-ISATAP-MARK
!
interface Tunnel2
description ISATAP VLAN2
no ip address
ip access-group 100 in
no ip redirects
load-interval 30
ipv6 address 2001:DB8:CAFE:2::/64 eui-64
ipv6 traffic-filter SOURCE-ISATAP-2 in

```

```

no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback2
tunnel mode ipv6ip isatap
!
interface Tunnel3
description ISATAP VLAN3
no ip address
no ip redirects
load-interval 30
ipv6 address 2001:DB8:CAFE:3::/64 eui-64
ipv6 traffic-filter SOURCE-ISATAP-3 in
no ipv6 redirects
no ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf 1 area 2
tunnel source Loopback3
tunnel mode ipv6ip isatap
!
interface TenGigabitEthernet1/1
description to 6k-sb-1
dampening
ip address 10.122.0.94 255.255.255.252
no ip redirects
no ip unreachable
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
ipv6 address 2001:DB8:CAFE:6505::B222:2020/64
no ipv6 redirects
ipv6 nd suppress-ra
ipv6 cef
ipv6 ospf network point-to-point
ipv6 ospf hello-interval 1
ipv6 ospf dead-interval 3
ipv6 ospf 1 area 0
mls qos trust dscp
service-policy output IPv6-ISATAP-MARK
!
interface GigabitEthernet4/1
description to 6k-core-1
dampening
ip address 10.122.0.82 255.255.255.252
ip access-group 101 in
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
interface GigabitEthernet4/2
description to 6k-core-2
dampening
ip address 10.122.0.90 255.255.255.252

```

```

ip access-group 101 in
no ip redirects
no ip proxy-arp
ip hello-interval eigrp 10 1
ip hold-time eigrp 10 3
ip authentication mode eigrp 10 md5
ip authentication key-chain eigrp 10 eigrp
load-interval 30
carrier-delay msec 0
mls qos trust dscp
!
router eigrp 10
passive-interface Loopback0
passive-interface Loopback1
passive-interface Loopback2
passive-interface Loopback3
network 10.0.0.0
no auto-summary
eigrp router-id 10.122.10.10
!
ip classless
!
no ip http server
!
ip access-list extended MGMT-IN-v4
remark Permit v4MGMT only to Lo0
permit tcp 10.120.0.0 0.0.255.255 any log-input
permit tcp 10.121.0.0 0.0.255.255 any log-input
permit tcp 10.122.0.0 0.0.255.255 any log-input
deny ip any any log-input
!
logging source-interface Loopback0
logging 10.121.11.9
access-list 101 permit 41 10.120.2.0 0.0.0.255 host 10.122.10.102
access-list 101 permit 41 10.120.3.0 0.0.0.255 host 10.122.10.103
access-list 101 permit 41 host 10.122.10.1 host 10.122.10.9
access-list 101 permit 41 host 10.122.10.2 host 10.122.10.19
access-list 101 permit 41 host 10.122.10.11 host 10.122.10.10
access-list 101 permit 41 host 10.122.10.12 host 10.122.10.20
access-list 101 deny 41 any any
access-list 101 permit ip any any
ipv6 router ospf 1
router-id 10.122.10.10
log-adjacency-changes
auto-cost reference-bandwidth 10000
area 2 range 2001:DB8:CAFE:2::/64 cost 20
area 2 range 2001:DB8:CAFE:3::/64 cost 20
passive-interface Loopback0
passive-interface Tunnel2
passive-interface Tunnel3
timers spf 1 5
!
snmp-server group IPv6-ADMIN v3 auth write v1default
snmp-server contact John Doe - ipv6rocks@cisco.com
!
ipv6 access-list SOURCE-ISATAP-2
remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
permit icmp 2001:DB8:CAFE:2::/64 any
remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:2::/64
permit ipv6 2001:DB8:CAFE:2::/64 any
remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
permit icmp FE80::/10 any
remark DENY ALL OTHER IPv6 PACKETS AND LOG
deny ipv6 any any log-input

```

```

!
ipv6 access-list SOURCE-ISATAP-3
  remark PERMIT ICMPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::/64
  permit icmp 2001:DB8:CAFE:3::/64 any
  remark PERMIT IPv6 PACKETS FROM HOSTS WITH PREFIX 2001:DB8:CAFE:3::/64
  permit ipv6 2001:DB8:CAFE:3::/64 any
  remark PERMIT ALL ICMPv6 PACKETS SOURCED BY HOSTS USING THE LINK-LOCAL PREFIX
  permit icmp FE80::/10 any
  remark DENY ALL OTHER IPv6 PACKETS AND LOG
  deny ipv6 any any log-input
!
ipv6 access-list BULK-APPS
  permit tcp any any eq ftp
  permit tcp any any eq ftp-data
!
ipv6 access-list TRANSACTIONAL-APPS
  permit tcp any any eq telnet
  permit tcp any any eq 22
!
ipv6 access-list MGMT-IN
  remark Permit MGMT only to Loopback0
  permit tcp 2001:DB8:CAFE::/48 host 2001:DB8:CAFE:6507::B222:2020 log-input
  deny ipv6 any any log-input
!
control-plane
!
dial-peer cor custom
!
banner login ^C

Unauthorized access to this device and network is prohibited.

^C
!
line con 0
  session-timeout 3
  password 7 021405550C031D
  logging synchronous
  login local
  transport output none
line vty 0 4
  session-timeout 3
  access-class MGMT-IN-v4 in
  password 7 071D2042490C0B
  ipv6 access-class MGMT-IN in
  logging synchronous
  login local
  exec prompt timestamp
  transport input telnet ssh
!
no cns aaa enable
end

```

