Deploying IPv6 in Campus Networks

Last Updated: February 27, 2012

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Author

Shannon McFarland

Shannon McFarland, Corporate Consulting Engineer, Office of the CTO, Cisco Systems

Shannon McFarland, CCIE #5245, is a Corporate Consulting Engineer in the Office of the CTO and is focused on Enterprise IPv6 deployment, VDI, and Data Center technologies. Shannon has been responsible for the Enterprise IPv6 design and deployment effort at Cisco for the last 9 years. He has authored many technical papers and Cisco Validated Design guides, is a contributor to Cisco Press books, and is a frequent speaker at Cisco Live and other industry conferences. He co-authored a Cisco Press book, "IPv6 in Enterprise Networks". Prior to his time at Cisco corporate, Shannon was an SE in the Cisco Englewood, CO office. Shannon has been at Cisco for 11+ years.

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit http://www.cisco.com/go/designzone.

Deploying IPv6 in Campus Networks

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.

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 Design Zone for Campus
  http://www.cisco.com/en/US/netsol/ns815/networking_solutions_program_home.html
- Cisco IPv6
  http://www.cisco.com/ipv6
- *IPv6 in Enterprise Networks* by Shannon McFarland, Muninder Sambi, Nikhil Sharma, Sanjay Hooda (ISBN-10:1-58714-227-9; ISBN-13: 978-1-58714-227-7)
  http://www.ciscopress.com/bookstore/product.asp?isbn=1587142279
- "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
- *IPv6 Security* by Scott Hogg, Eric Vyncke (ISBN-10:1-58705-594-5; ISBN-13: 978-1-58705-594-2)
  http://www.ciscopress.com/bookstore/product.asp?isbn=1587055945

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

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.

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

- Dual-stack model (DSM)
- Hybrid model (HM)
- 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.

Deployment Models Overview

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

- DSM—End-to-end dual stack network
- HM—Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) and 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 29.

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

Conclusion, page 57 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

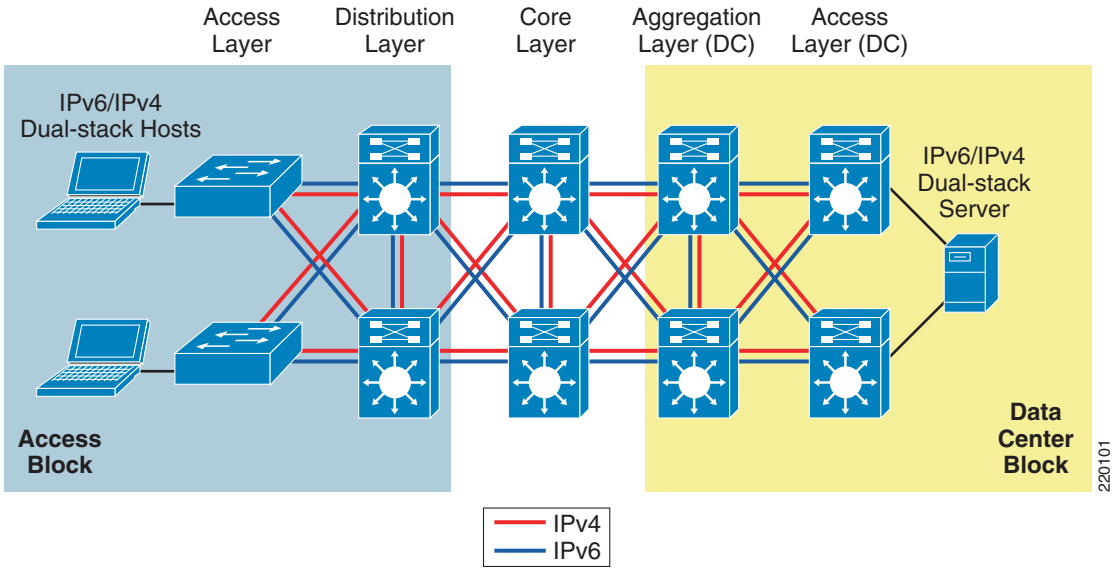

Tested Components

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 3750E/3560E	12.2(55)SE
	Catalyst 4500 Supervisor 6-E	12.2(53)SG
	Catalyst 4500 Supervisor 7-E	XE 3.1.1SG
	Catalyst 6500 Supervisor 720	12.2(33)SXI4
Host devices	Various clients	Microsoft Windows
Distribution layer	Catalyst 4500 Supervisor 6-E	12.2(46)SG
	Catalyst 4500 Supervisor 7-E	XE 3.1.1SG
	Catalyst 6500 Supervisor 720	12.2(33)SXI4
Core layer	Catalyst 6500 Supervisor 720	12.2(33)SXI4

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

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

The key aspect of HM 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. The HM leverages the ISATAP tunneling mechanisms on the hosts in the access layer to provide IPv6 addressing and off-link routing. The Microsoft Windows 7 (XP and Vista are also supported) 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.

Figure 2 shows the basic connectivity flow for HM.

Figure 2 Hybrid Model—Connectivity Flow

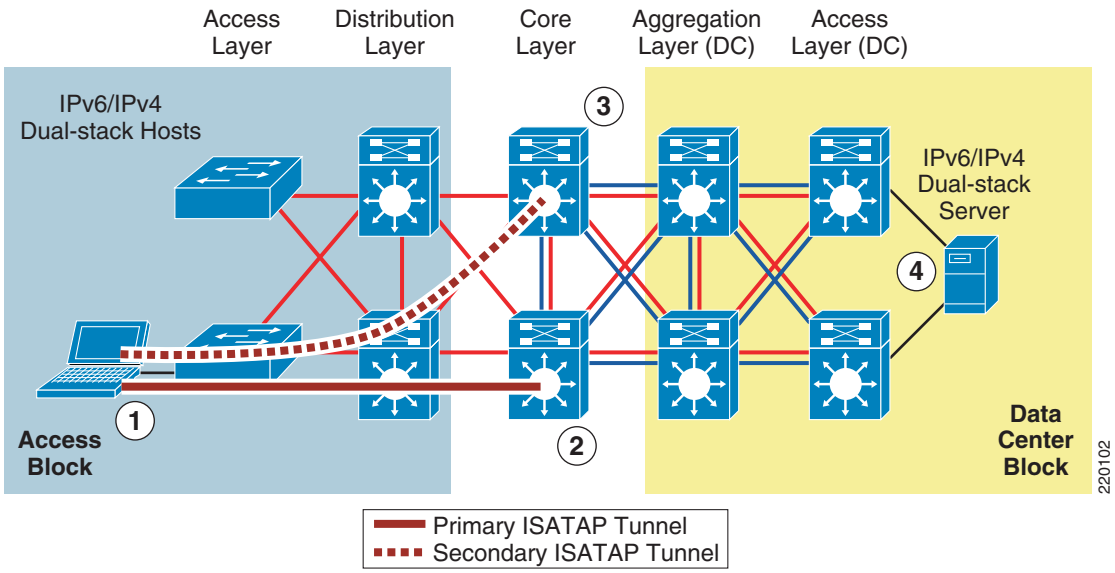

1. The host establishes an ISATAP tunnel to the core layer.
2. The core layer switches are configured with ISATAP tunnel interfaces and are the termination point for ISATAP tunnels established by the hosts.
3. Pairs of core layer switches are redundantly configured to accept ISATAP tunnel connections 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 Windows 7 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—ISATAP Tunnel Mapping

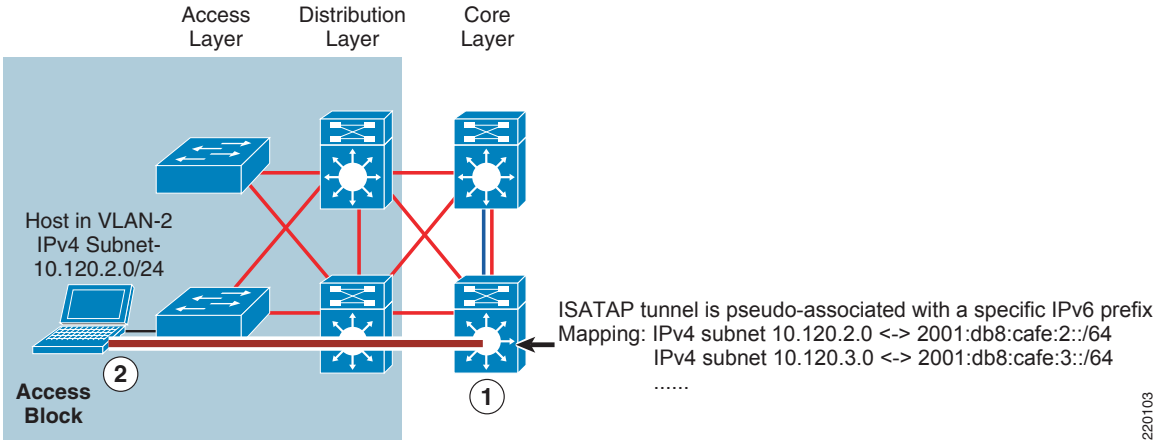

1. The core layer switch is configured with a loopback interface with the address of 10.122.10.2, which is used as the tunnel source for ISATAP, and is used only by users located on the 10.120.2.0/24 subnet.
2. The host in the access layer is connected to a port that is associated with a specific VLAN. 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 PowerShell, Microsoft Group Policy, 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.

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 HM 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—For years, it has been a common best practice 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 think about. It is important to remember that a single tunnel interface will have many tunnels established by ISATAP users and, based on Cisco testing, it is easy to scale the number of ISATAP connections into the thousands.

Solution Requirements

The following are the main solution requirements for HM 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.
- 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 HM.

Benefits and Drawbacks of This Solution

The primary benefit of HM 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, HM may be a suitable choice.

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

- The HM can be and most often is an operational nightmare to manage.
- IPv6 multicast is not supported within ISATAP tunnels.
- 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 57 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 HM. 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

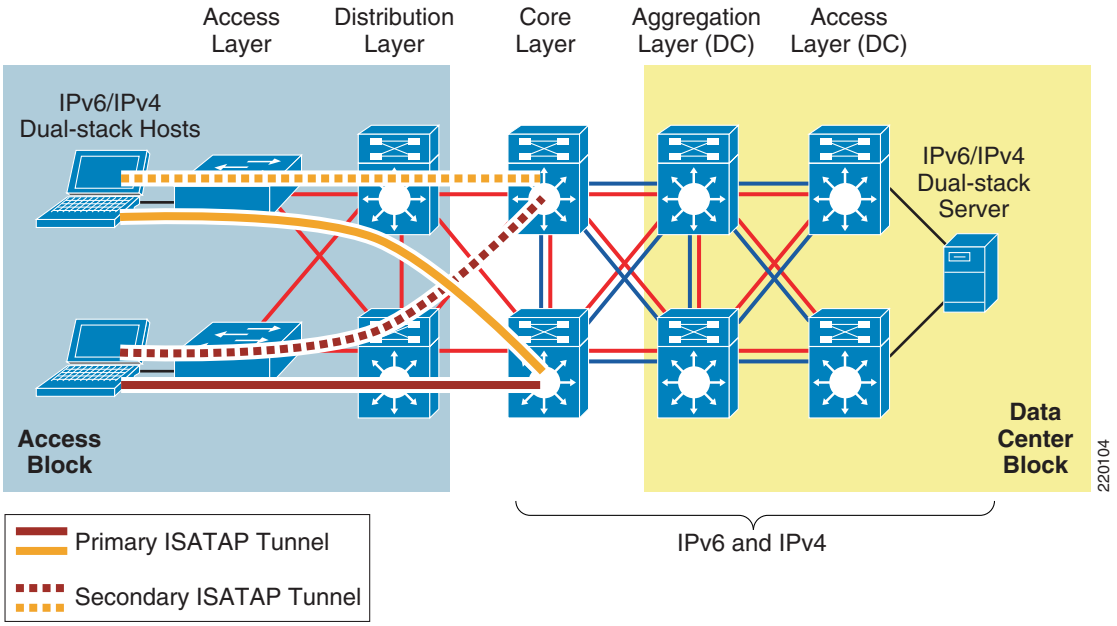

Tested Components

Table 2 lists the components used and tested in the HM configuration. It is important to note that the only Cisco Catalyst components that need to have IPv6 capabilities are those terminating ISATAP connections and the dual-stack links in the data center. Therefore, the software versions in the campus access and distribution layer roles are not relevant to this design.

Table 2 HM Tested Components

Campus Layer	Hardware	Software
Core layer	Catalyst 6500 Supervisor 720	12.2(33)SXI4

Service Block Model

Overview

SBM is significantly different compared to the other campus models discussed in this document. 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/solutions/ns340/ns414/ns742/ns815/landing_cNet_virtualization.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. Alternatively, Cisco routing platforms, such as the Cisco ISR series, can be used to terminate the ISATAP tunnels in smaller environments. 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 a few similarities between the SBM example given in this document and the HM. 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 HM). IPv4 routing is configured between the core layer and SBM switches to allow visibility to the SBM 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 HM, 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 HM design (lots of tunneling), but also adds the cost of additional equipment not found in HM. 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 HM and SBM, Cisco recommends to always try to deploy the DSM.

Conclusion, page 57 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 5 shows the ISATAP portion of the design and Figure 6 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 5 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 5 Service Block Model—Connecting the Hosts (ISATAP Layout)

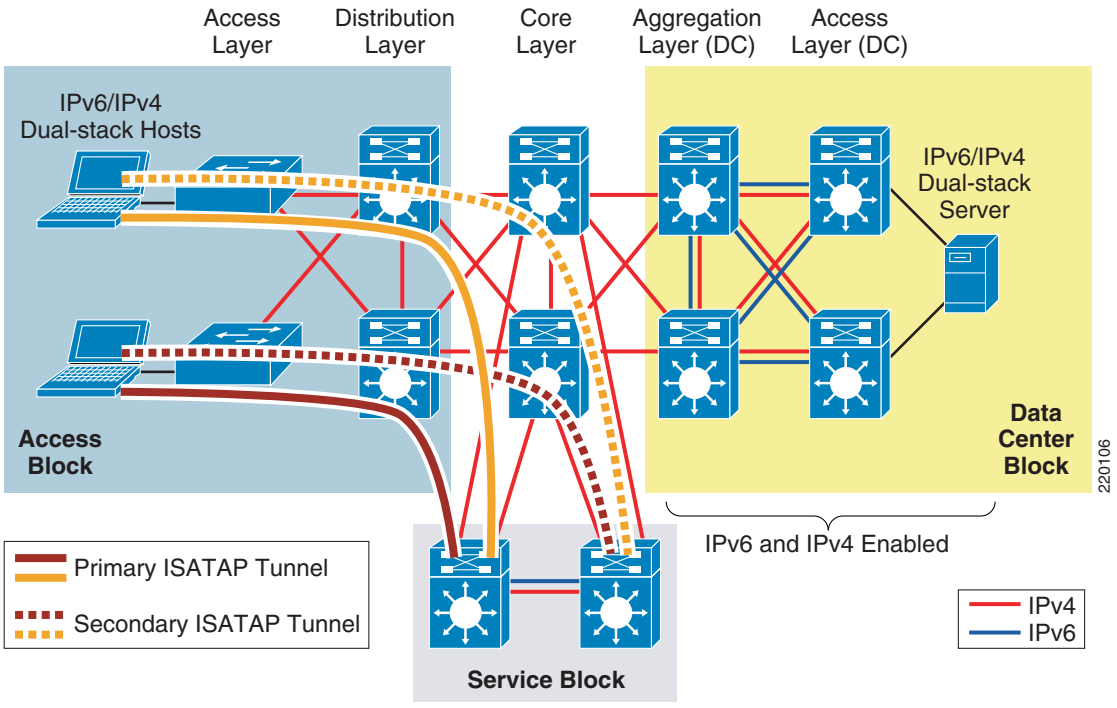

Figure 6 shows the redundant, manually-configured tunnels connecting the data center aggregation layer and the service blocks. Hosts located in the access layer can now reach IPv6 services in the data center access layer using IPv6.

Figure 6 Service Block Model—Connecting the Data Center (Manually-Configured Tunnel Layout)

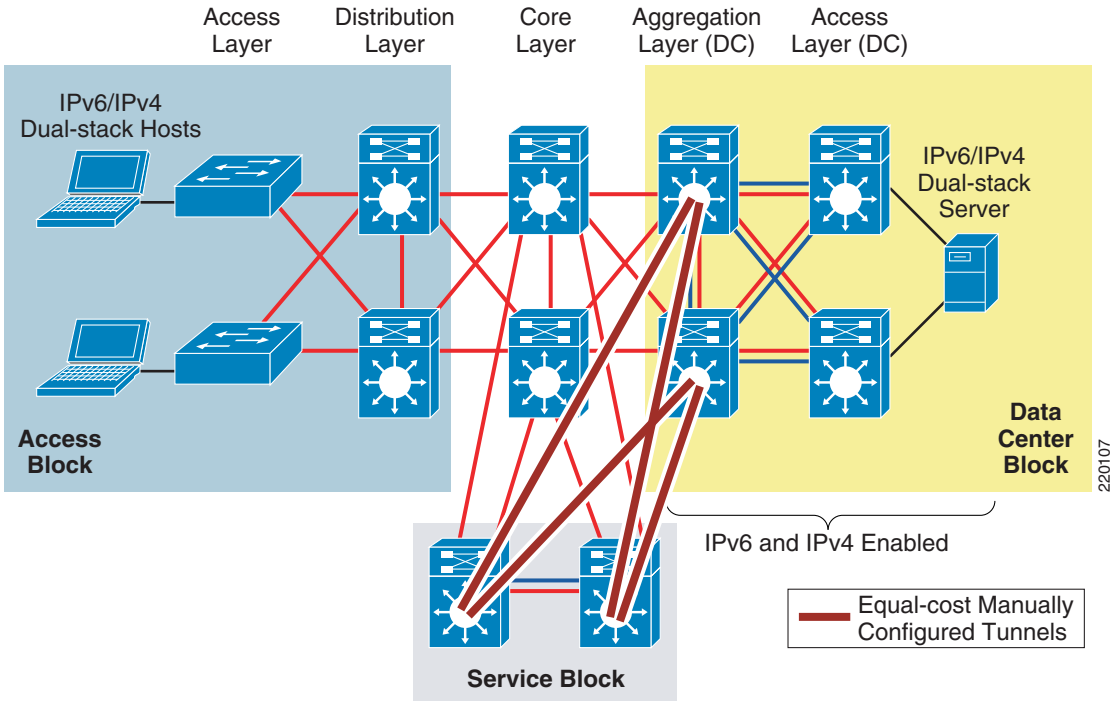

Tested Components

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

Table 3 SBM Tested Components

Campus Layer	Hardware	Software
Service block	Catalyst 6500 Supervisor 720	12.2(33)SXI4

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 at the following URL: http://www.cisco.com/en/US/netsol/ns815/networking_solutions_program_home.html.

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. specifically for links.

Table 4 breaks down the options related to the use of various prefix lengths on links.

Table 4 Prefix Link Considerations

64 Bits	Less than 64 Bits	Greater than 64 Bits
• Recommended by RFC3177 and IAB/IESG • Consistency makes management easy • This is required for SLAAC, SEND, Privacy extensions, and other configurations • Significant address space loss	• Enables more hosts per broadcast domain • Considered bad practice • 64 bits offers more space for hosts than the media can support efficiently	• Address space conservation • Special cases: – /126—Valid for p2p – /127—Valid for p2p with specific restrictions: draft-kohno-ipv6-prefixlen-pep and RFC3627 – /128 —loopback • Complicates management • Must avoid overlap with specific addresses: – Router Anycast (RFC3513) – Embedded RP (RFC3956)

- /64—On VLAN interfaces, it is recommended to use a /64 prefix because it is easy and consistent for address management. This is required for SLAAC, SEND, and privacy extension use.
- Less than 64—There are no real use cases where a site needs more addressing on a link than a /64 can provide and is considered a bad practice.
- Greater than 64—Some in the IPv6 community think that a /64 prefix for p2p links is a waste and even a security attack vector. The debate rages on regarding the use of various prefix lengths on p2p links and the reader is encouraged to balance the legalistic RFC stipulations with real-world deployment considerations. In many deployments it is common to use a /64 on VLANs (or links where hosts reside), /126 on p2p links, and /128 on loopbacks.

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. At the time this document was written, the IETF draft “Using 127-bit IPv6 Prefixes on Inter-Router Links” (draft-kohno-ipv6-prefixlen-p2p) provides the best guidance on the use of /127.

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 is supported with IPv6.

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 Cisco branch-LAN design best practice documents at:
http://www.cisco.com/en/US/netsol/ns816/networking_solutions_program_home.html.

Routing

The decision to run an IGP in the campus network was made 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 also be used. The IGP configurations for IPv6 can either be EIGRP or OSPFv3. These IGPs are interchanged in some sections to show the reader what the basic configuration looks like for either IGP.

As previously mentioned, every effort has been made to implement the current Cisco campus design best practices. Both the IPv4 and IPv6 IGPs 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.

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 using EIGRP for IPv4 when redundant paths for tunnels are needed, and EIGRP for IPv6 or 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 redundant Layer 3 switches are applicable in the HM 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 HM 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. In this document, configurations are shown with HSRP for IPv6.

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 HM 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 (HM) 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 7 illustrates the points where IPv6 QoS policies may be applied when using ISATAP in HM. 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 59 for more information about the Cisco campus QoS documentation.

Figure 7 QoS Policy Implementation—HM

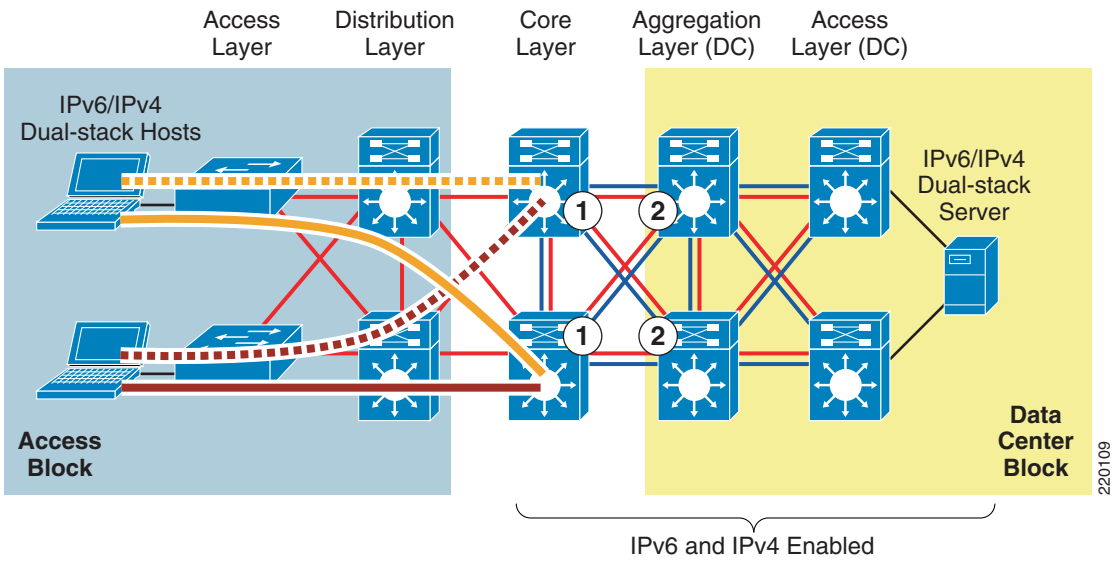

1. In HM, the first place to implement classification and marking is on the egress interfaces on the core layer switches. As was previously mentioned, the IPv6 packets have been tunneled from the hosts in the access layer to the core layer, and the IPv6 packets have not been "visible" in a decapsulated 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 polices may include trust (ingress), policing (ingress), and queuing (egress).

Figure 8 illustrates the points where IPv6 QoS policies may be applied in the SBM when ISATAP manually-configured tunnels are used.

Figure 8 QoS Policy Implementation—SBM (ISATAP and Manually-Configured Tunnels)

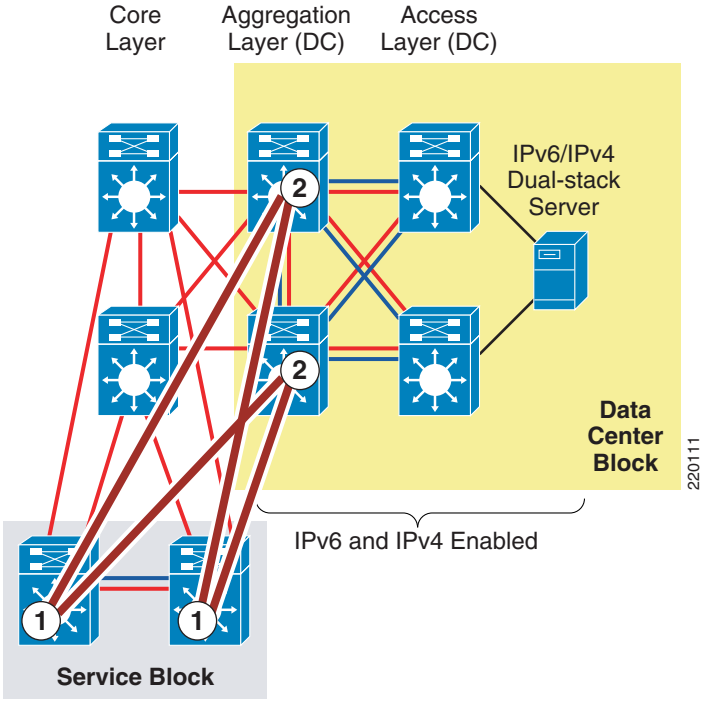

1. The SBM switches receive IPv6 packets coming from the ISATAP interfaces, which are now decapsulated, and can apply classification and marking policies on the egress manually-configured tunnel interfaces.
2. The upstream switches (aggregation layer and access layer) can now apply trust, policing, and 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/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/qos.html#wp1571584.

The DSM model is not shown here because the same recommendations for implementing QoS policies for IPv4 should also apply to IPv6.

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 59 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:

- 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
```

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
snmp-server host 2001:DB8:CAFE:100::60 version 3 auth jdoe
```

- 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).
  - HM—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.
  - 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/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/qos.html
- Enterprise QoS SRND— http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html

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:

001244: *Jun 13 09:33:22.426 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.

001245: *Jun 13 09:33:22.430 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/en/US/docs/switches/lan/catalyst6500/ios/12.2SXF/native/configuration/guide/dos.html

- First Hop Security—Features such as IPv6 Port-based ACL (PACL), Router Advertisement (RA) Guard, Neighbor Discovery (ND) Inspection, and SEcure Neighbor Discovery (SEND) are all mechanisms that can be deployed to protect the data and control plane of the first hop devices and attached hosts. The following shows an example of IPv6 PACL:

```
ipv6 access-list HOST_PACL
 remark Deny Rogue DHCP
 deny udp any eq 547 any eq 546
 remark Deny RA From Client
 deny icmp any any router-advertisement
 permit ipv6 any any
!
interface GigabitEthernet1/0/6
 ipv6 traffic-filter HOST_PACL in
```

  The following shows an example of RA Guard:

```
interface GigabitEthernet1/0/6
 ipv6 nd raguard
```

More information on IPv6 security can be found in Additional References, page 59.

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 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/docs/ios/ipv6/configuration/guide/ip6-multicast_ps10591_TSD_Products_Configuration_Guide_Chapter.html#wp1058805.

In this document, IPv6 multicast-enabled applications are supported in the DSM because no ISATAP configurations are used. The multicast-enabled application tested in this design is Windows Media Services with Embedded-RP and/or PIM-SSM groups. The multicast sources are running on Microsoft Windows Server 2008 servers located in the data center.

Several documents on cisco.com 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.

For more information, see:

- Cisco IPv6 Multicast: http://www.cisco.com/en/US/technologies/tk648/tk828/tk363/technologies_white_paper0900aecd8014d6dd.html
- Cisco IOS IPv6 Multicast: http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-multicast_ps10591_TSD_Products_Configuration_Guide_Chapter.html

Network and Address Management

Network Management

Many of the traditional management tools used today support IPv6 in much the same way as IPv4. In this document, the only considerations for management of the campus network are related to basic management services (Telnet, SSH, and SNMP). SNMP over IPv6 transport is supported on the latest versions of the software, depending on the Catalyst platform. Refer to the platform documentation for support for SNMP over IPv6 transport.

Address Management

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. However, there are many other places where manual address changes occur that general prefix does not help, such as ACLs, QoS policy definition, routing policies, etc.

More information on the general prefix feature can be found at the Cisco IOS IPv6 documentation page (see Additional References, page 59).

IPv6 address allocation to hosts located in the campus access layer can be assigned via SLAAC, static assignment or DHCPv6. DHCPv6 assignment is the predominant method that is desired by most enterprise campus administrators. Until support for DHCPv6 Relay was available on the Catalyst products, the administrator had no other choice but to rely on SLAAC as the primary means of allocating IPv6 addressing to hosts in the access layer.

Figure 9 shows that the placement of the DHCPv6 Relay is on the campus distribution layer switches, which is the same place as the IP helper used in DHCP for IPv4.

Figure 9 DHCPv6 Relay Placement in the Campus

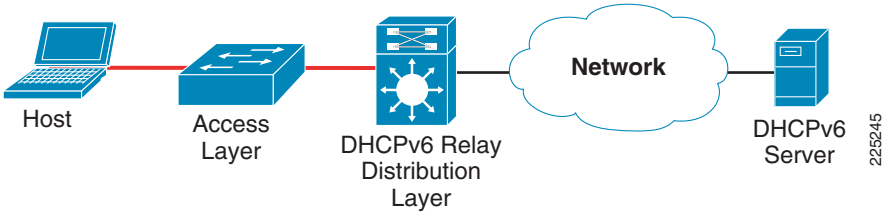

The configuration of the DHCPv6 Relay feature is straightforward. Implement the following configuration on the VLAN interface facing the access layer hosts:

```
interface Vlan2
 description ACCESS-DATA-2
 ipv6 address 2001:DB8:CAFE:2::A111:1010/64
 ipv6 nd prefix 2001:DB8:CAFE:2::/64 0 0 no-autoconfig
 ipv6 nd managed-config-flag
 ipv6 dhcp relay destination 2001:DB8:CAFE:11::9
```

The **ipv6 dhcp relay destination** command defines the unicast address of the DHCPv6 server. The **ipv6 nd managed-config-flag** command sets the “managed address configuration” flag in the RA so the host knows to use a stateful address configuration mechanism, such as DHCPv6. More information on

DHCPv6 Relay Agent can be found at: http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-dhcp_ps10591_TSD_Products_Configuration_Guide_Chapter.html.

Having DHCPv6 Relay Agent support in the network is only part of the equation. The client must support DHCPv6 (such as Microsoft Windows 7) and there must be a DHCPv6 server. When this document was written, Apple did not have DHCPv6 client support on its products. In designs where Apple products need addresses, SLAAC is still the method to be used.

Currently, there are three DHCPv6 servers that have been tested by Cisco in the campus:

- Cisco Network Registrar: http://www.cisco.com/en/US/products/sw/netmgtsw/ps1982/
- Cisco DHCPv6 Server in IOS: http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-dhcp_ps10591_TSD_Products_Configuration_Guide_Chapter.html#wp1054494
- Microsoft Windows Server 2008: http://technet2.microsoft.com/windowsserver2008/en/library/bab0f1a1-54aa-4cef-9164-139e8bcc44751033.mspx?mfr=true

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, refer to 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 Vl2
2001:DB8:CAFE:2:7DE5:E2B0:D4DF:97EC         16 000d.6084.2c7a  STALE Vl2
FE80::7DE5:E2B0:D4DF:97EC                   16 000d.6084.2c7a  STALE Vl2
```

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 HM are as follows:

- Access layer—There are no real scale or performance considerations for the access layer when using the HM. IPv6 is not supported in the access layer in the HM, 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—The same considerations as in the access layer.
- Core layer—There can be an impact on the core layer switches when using HM. There can be hundreds or more ISATAP tunnels terminating on the core layer switches in the HM. 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 HM 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 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 HM 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 (HM) 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.

Network Topology

The following diagrams are used as a reference for all DSM configuration examples. Figure 10 shows the physical port layout that is used for the DSM.

Figure 10 DSM Network Topology—Physical Ports

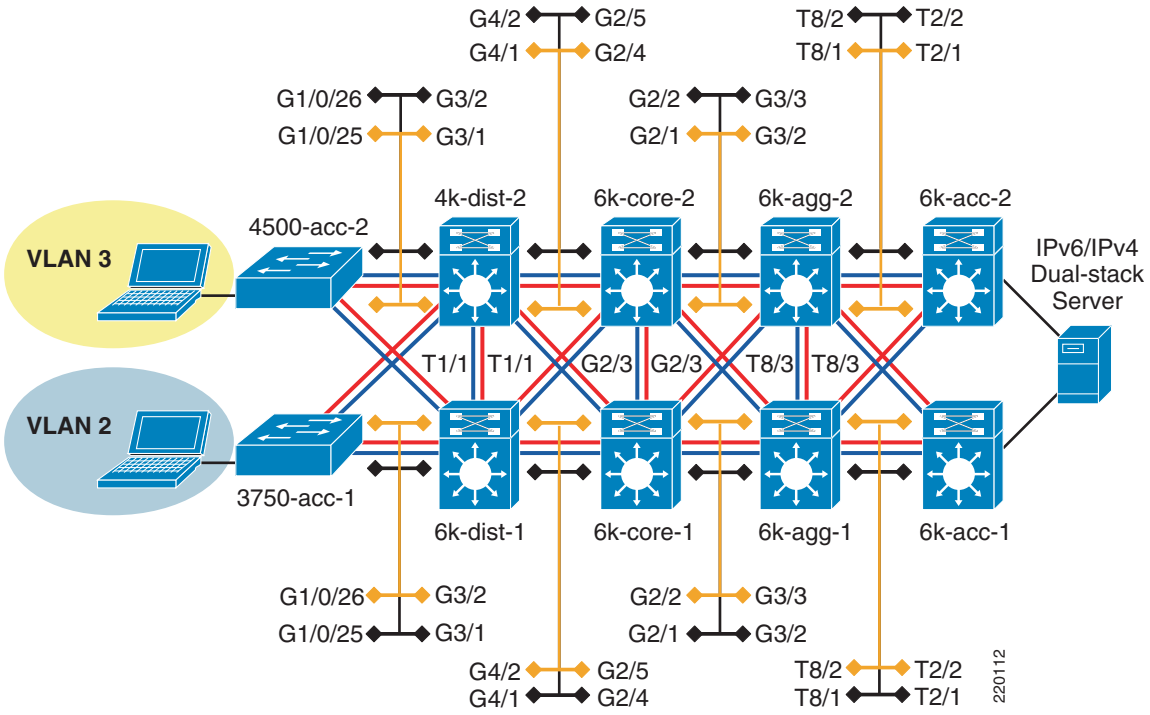

Figure 11 shows the IPv6 addressing plan for the DSM environment. To keep the diagram as simple to read as possible, the /48 prefix portion of the network is deleted. The IPv6 /48 prefix used in all the models in this paper is "2001:db8:cafe::/48".

Figure 11 DSM Network Topology—IPv6 Addressing

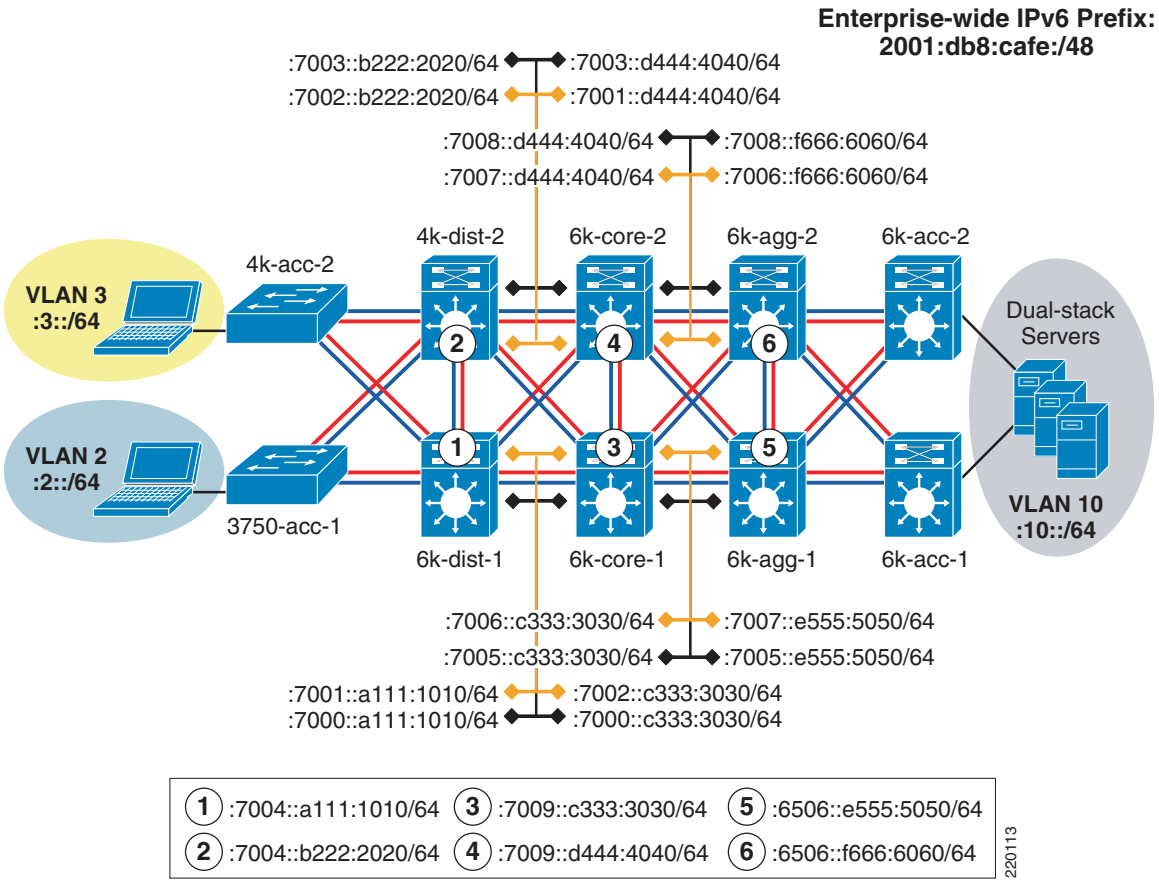

In addition to the physical interfaces, IPv6 addresses are assigned to loopback and VLAN interfaces. Table 6 shows the switch, interface, and IPv6 address for the interface.

Table 6 Switch, Interface, and IPv6 Addresses

Switch	Interface	IPv6 address
3750-acc-1	VLAN2	2001:db8:cafe:2::cac1:3750/64
4k-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
4k-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
 ipv6 address 2001:DB8:CAFE:7000::A111:1010/64   #Assign IPv6 address
 ipv6 nd suppress-ra                             #Disable RAs on this interface
```

- 6k-core-1:

```
ipv6 unicast-routing
ip cef distributed
ipv6 cef distributed
!
interface GigabitEthernet2/4
 description to 6k-dist-1
 ipv6 address 2001:DB8:CAFE:7000::C333:3030/64
 ipv6 nd suppress-ra
```

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 has 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, see: http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_55_se/configuration/guide/swsdm.html.

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
!                                        #VTP and STP configurations
!                                        #shown for completeness, but not
!                                        #specific to IPv6
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                  #Physical intf. to 6k-dist-1
 description TRUNK TO 6k-dist-1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 2
 switchport mode trunk
 switchport nonegotiate
!
interface Vlan2                                 #VLAN2 with IPv6 address used for mgmt.
 ipv6 address 2001:DB8:CAFE:2::CAC1:3750/64
```

- 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                   #6k-dist-1 is the STP root for
                                                        #VLAN2,3
!
vlan internal allocation policy descending
vlan dot1q tag native
!
vlan 2
 name ACCESS-DATA-2
!
vlan 3
 name ACCESS-DATA-3
!
interface GigabitEthernet3/1                        #Physical intf. to 3750-acc-1
 description to 3750-acc-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 2
 switchport mode trunk
 switchport nonegotiate
 no ip address
 spanning-tree guard root
!
interface Vlan2                                 #VLAN2 termination
                                                #point for trunked VLAN from 3750-acc-1
description ACCESS-DATA-2
 ipv6 address 2001:DB8:CAFE:2::A111:1010/64
 ipv6 nd prefix 2001:DB8:CAFE:2::/64 0 0 no-autoconfig
 ipv6 nd managed-config-flag             #Enable "managed" flag for DHCPv6
 ipv6 dhcp relay destination 2001:DB8:CAFE:11::9#Define the DHCPv6 server address
```

Although switch stacks are not used in any of the models discussed here, they are commonly used on the Catalyst 3750 series 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, refer to: http://www.cisco.com/en/US/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_55_se/configuration/guide/swstack.html.

Routing Configuration

As previously mentioned, the routing for the DSM is set up using EIGRP for both IPv4 and IPv6. The EIGRP configuration follows the recommended Cisco campus designs as much as possible. The configuration for EIGRP for IPv6 is shown for the 6k-dist-1 and 6k-core-1 switches.

- 6k-dist-1:

```
key chain eigrp
 key 100
   key-string 7 1111
!
interface Loopback0
 ip address 10.122.10.9 255.255.255.255              #Address used for RID on EIGRP
 ipv6 address 2001:DB8:CAFE:6507::A111:1010/128
 ipv6 eigrp 10
!
interface TenGigabitEthernet1/1
 description to 4k-dist-2
 ipv6 address 2001:DB8:CAFE:7004::A111:1010/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
!
interface GigabitEthernet4/1
 description to 6k-core-1
 ipv6 address 2001:DB8:CAFE:7000::A111:1010/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
!
interface GigabitEthernet4/2
 description to 6k-core-2
 ipv6 address 2001:DB8:CAFE:7001::A111:1010/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
!
interface Vlan2
 description ACCESS-DATA-2
 ipv6 address 2001:DB8:CAFE:2::A111:1010/64
 ipv6 eigrp 10
!
ipv6 router eigrp 10
 router-id 10.122.10.9                    #RID using Loopback0
 no shutdown
 passive-interface Vlan2                  #Do not establish adjacency over
                                          #VLANs
 passive-interface Vlan3
 passive-interface Loopback0
```

- 6k-core-1:

```
key chain eigrp
 key 100
   key-string 7 1111
```

```
!
interface Loopback0
 ip address 10.122.10.3 255.255.255.255
 ipv6 address 2001:DB8:CAFE:6507::C333:3030/128
 ipv6 eigrp 10
!
interface GigabitEthernet2/1
 description to 6k-agg-1
 ipv6 address 2001:DB8:CAFE:7005::C333:3030/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
!
interface GigabitEthernet2/4
 description to 6k-dist-1
 ipv6 address 2001:DB8:CAFE:7000::C333:3030/64
 ipv6 eigrp 10
 ipv6 hello-interval eigrp 10 1
 ipv6 hold-time eigrp 10 3
 ipv6 authentication mode eigrp 10 md5
 ipv6 authentication key-chain eigrp 10 eigrp
!
ipv6 router eigrp 10
 router-id 10.122.10.3
 no shutdown
 passive-interface Loopback0
```

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 59 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. All distribution pairs in the reference campus configuration are running HSRP for both IPv4 and IPv6. Optionally, GLBP can be used. The configuration for HSRP for IPv4 and IPv6 on the 6k-dist-1switch is shown below:

- 6k-dist-1:

```
interface Vlan2
 description ACCESS-DATA-2
 standby version 2                #Standby Version 2 is required for IPv6 support
 standby 1 ip 10.120.2.1
 standby 1 timers msec 250 msec 750
 standby 1 priority 110
 standby 1 preempt delay minimum 180
 standby 1 authentication ese
 standby 2 ipv6 autoconfig    #Allow the system to self-generate the IPv6
                              #virtual address
 standby 2 timers msec 250 msec 750
 standby 2 priority 110
 standby 2 preempt delay minimum 180
```

```
 standby 2 authentication ese
```

QoS Configuration

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 examples shown below are very basic and not meant to be a best practice recommendation.

The example shown below is doing classification/marking of Microsoft Windows Media Services on TCP 1755. The traffic is set to AF31 (DSCP 26) with basic policing and queuing applied.

Your own policies will differ from these. For the sake of brevity, not all interfaces are shown.

- Access Layer QoS Example (Catalyst 3750E)

```
mls qos map policed-dscp  46 to 8
mls qos map policed-dscp  26 to 28              #Policed DSCP QoS map 26 (AF31) to 28
mls qos srr-queue input threshold 1 40 60
mls qos srr-queue input threshold 2 40 60
mls qos srr-queue input buffers 33 67
mls qos srr-queue input priority-queue 1 bandwidth 33
mls qos srr-queue input dscp-map queue 1 threshold 3 46
mls qos srr-queue input dscp-map queue 2 threshold 1 8
mls qos srr-queue input dscp-map queue 2 threshold 2 0
mls qos srr-queue input dscp-map queue 2 threshold 3 26 28
mls qos srr-queue output dscp-map queue 1 threshold 3 46
mls qos srr-queue output dscp-map queue 2 threshold 1 28
mls qos srr-queue output dscp-map queue 2 threshold 3 26     #Queue/Drop threshold for
                                                            #AF31
mls qos srr-queue output dscp-map queue 3 threshold 3 0
mls qos srr-queue output dscp-map queue 4 threshold 3 8
mls qos queue-set output 1 threshold 1 100 100 100 100
mls qos queue-set output 1 threshold 2 10 70 100 100
mls qos queue-set output 1 threshold 3 100 100 100 100
mls qos queue-set output 1 threshold 4 100 100 100 100
mls qos queue-set output 1 buffers 33 25 37 5
mls qos
!

class-map match-all IPV6-STREAM
    match access-group name IPV6-STREAM-DATA
!
policy-map TEST-CLASS                         #Basic policy to police/Classify
    class IPV6-STREAM                            #Windows Media traffic
    police 37500000 80000 exceed-action policed-dscp-transmit
    set dscp af31
    class class-default
    set dscp default !
interface GigabitEthernet1/0/25
    description to Distribution Layer
    srr-queue bandwidth share 1 25 70 5
    srr-queue bandwidth shape 3 0 0 0
    priority-queue out
    mls qos trust dscp
!
interface GigabitEthernet1/0/7
 description to Access Port
 srr-queue bandwidth share 1 25 70 5
 srr-queue bandwidth shape 3 0 0 0
 priority-queue out
 mls qos trust device cisco-phone
```

```
 service-policy input TEST-CLASS        #Apply policy ingress from access port
!
ipv6 access-list IPV6-STREAM-DATA       #Example ACL for Windows Media
 permit tcp any any eq 1755
```

- 6k-dist-1

```
mls qos
!
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           #Queue 3 WRED Threshold 3
 wrr-queue cos-map 3 4 6
 wrr-queue cos-map 3 5 7
 mls qos trust dscp                #Trust DSCP markings
!
interface GigabitEthernet4/1
 description to 6k-core-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
```

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.

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    Gi1/0/25(dynamic), Gi1/0/26(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          Gi1/0/25, Gi1/0/26
```

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 12 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 12 DSM Topology—Routed Access Design

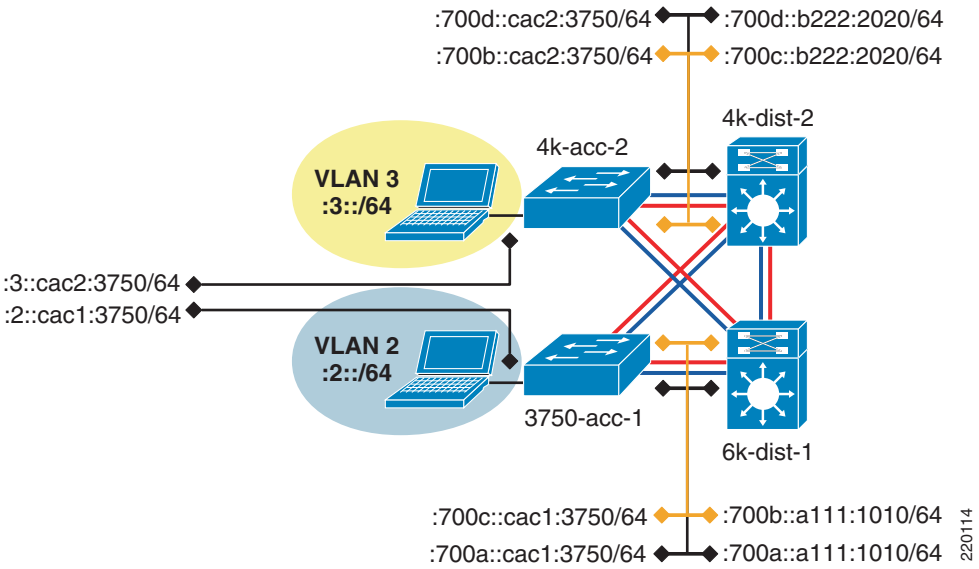

Figure 12 shows that the links between the access layer and distribution layer are now routed links instead of trunked Layer 2 links. IPv6 addressing and routing is configured on the new links, and the hosts in the VLANs use the IPv6 address of the VLAN interface on the access switch as the default gateway. Because the VLAN is terminated on the access layer switch, there is no need for a first hop redundancy protocol like HSRP. Availability is provided by routing between the access and distribution layers.

Note For those readers using OSPF in their network, the following IGP configuration is shown using OSPFv3. This is a sample of what the configurations would look like in the campus for OSPFv3 in routed access model. This is an effort to help the reader see the IGP configurations for both EIGRP for IPv6 and OSPFv3 in a campus environment.

The following configuration example shows the relevant configurations for the 3750-acc-1 and 6k-dist-1 switches.

- 3750-acc-1

```
ipv6 unicast-routing
!
interface GigabitEthernet1/0/25
 description To 6k-dist-1
 ipv6 address 2001:DB8:CAFE:700A::CAC1:3750/64
 ipv6 nd suppress-ra
 ipv6 ospf network point-to-point                #OSPFv3 is configured
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 ospf 1 area 2
 mls qos trust dscp
!
interface Vlan2
 ipv6 address 2001:DB8:CAFE:2::CAC1:3750/64      #VLAN2 on this switch becomes the
                                                 #first layer 3 hop for the hosts
                                                 #in VLAN2 - the link-local address
                                                 #on VLAN 2 will be the default
                                                 #gateway for the hosts
 ipv6 ospf 1 area 2
!
ipv6 router ospf 1
 router-id 10.120.2.1
 log-adjacency-changes
 auto-cost reference-bandwidth 10000
 area 2 stub no-summary                       #Per the Routed Access Design guide - the
                                              #area (area 2)  for the access layer
                                              #prefix is a totally stubby area

 passive-interface Vlan2
 timers spf 1 5
```

- 6k-dist-1

```
interface GigabitEthernet3/1
 description to 3750-acc-1
 ipv6 address 2001:DB8:CAFE:700A::A111:1010/64
 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

 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    #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 output of the **show ipv6 route** command 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            #4k-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 Design Guide—Routed Access and High Availability
  http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns815/landing_cHi_availability.html
- Routing in the Wiring Closet white paper
  http://www.cisco.com/en/US/netsol/ns340/ns394/ns147/ns17/networking_solutions_white_paper0900aecd804c6e73.shtml

Hybrid Model

Most of the campus network in the HM 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 HM uses dual-stack from the core layer into the data center. Those configurations are not relevant to the HM 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.

Network Topology

One difference in the HM 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 13 shows the network topology for the HM.

Figure 13 HM Network Topology

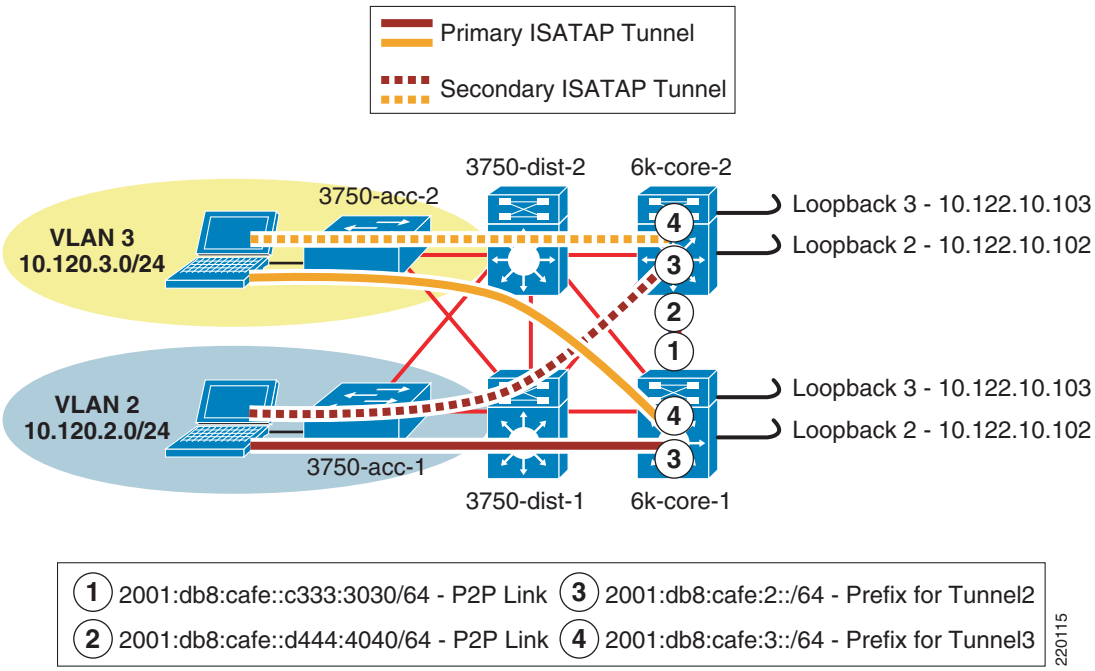

The topology is focused on the IPv4 addressing scheme in the access layer (used by the host to establish the ISATAP tunnel), core layer (used as the termination point by the host for ISATAP), and also the IPv6 addressing used in the core layer for both the p2p link and the ISATAP tunnel prefix. The configuration 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.

- 6k-core-1

```
interface GigabitEthernet1/1
 description to 3750-dist-1
 ip address 10.122.0.41 255.255.255.252
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 mls qos trust dscp
!
interface GigabitEthernet1/2
 description to 3750-dist-2
 ip address 10.122.0.45 255.255.255.252
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 mls qos trust dscp
!
interface GigabitEthernet2/3
```

```
 description to 6k-core-2
 ip address 10.122.0.21 255.255.255.252
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 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
 ip address 10.122.0.49 255.255.255.252
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 mls qos trust dscp
!
interface GigabitEthernet1/2
 description to 3750-dist-2
 ip address 10.122.0.53 255.255.255.252
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 mls qos trust dscp
!
interface GigabitEthernet2/3
 description to 6k-core-1
 ip address 10.122.0.22 255.255.255.252
 ip hello-interval eigrp 10 1
 ip hold-time eigrp 10 3
 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 performs 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 HM 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 HM 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 Windows 7 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                            #Tunnel interfaces disable the
                                                   #sending of RA's. This command
                                                   #re-enables RA's
 ipv6 ospf 1 area 2
 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
 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 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
                                                    #the route from 6k-core-2 is not
                                                    #preferred or equal to 6k-core-1
                                                    #Not required.
 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
```

Figure 14 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 14 HM—Preferred Route for 6k-core-1

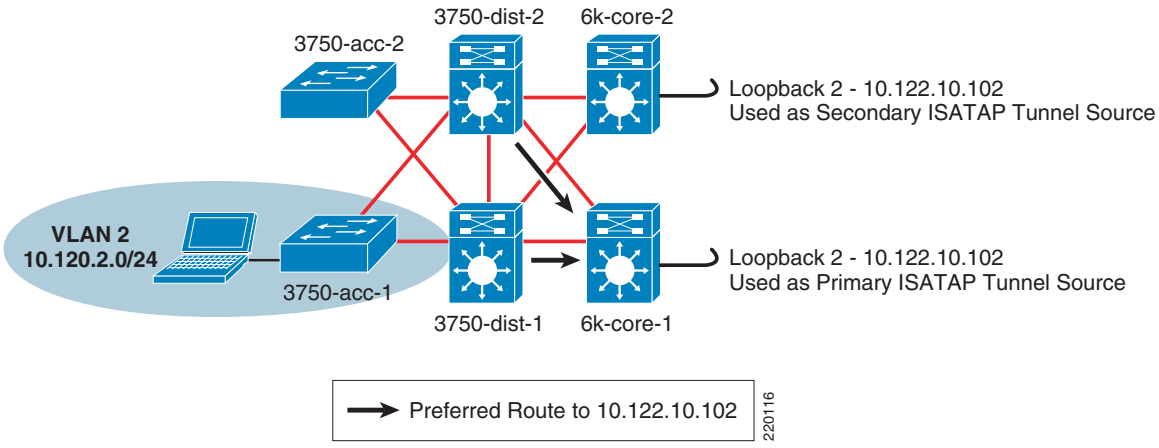

The routing table 10.122.10.102 on the distribution layer switches is as follows:

- 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-1#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
```

Figure 15 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 15 HM—Preferred Route for 6k-core-2 After Failure of 6k-core-1

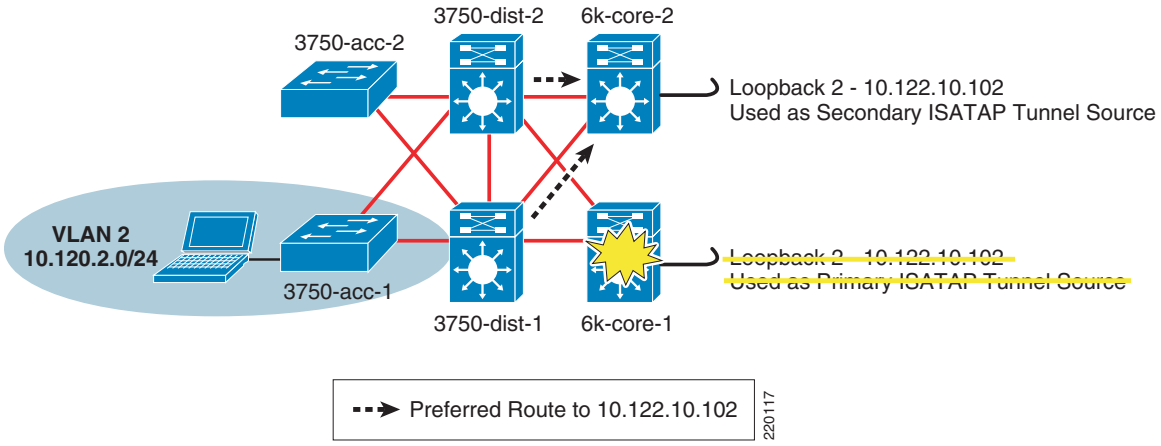

The updated routing table entry for 10.122.10.102 on the distribution layer switches is as follows:

- 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-1#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 HM 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 Microsoft Group Policy and Windows PowerShell can be used to run the command locally on the host at login or another predetermined time.

On the Microsoft Windows hosts 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 HM 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 10.122.10.10**2** where the "**2**" in "10**2**" signifies VLAN or Subnet **2**. The same would happen for VLAN **3** or 10.120.**3**.0/24 where the ISATAP router is 10.122.10.10**3**:

```
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" or "0200:5efe" (if using public IPv4 addresses). 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 PowerShell or login scripts 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 HM should match the existing IPv4 policies. As previously mentioned, the HM 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
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 29).

- 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 16, 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 HM. 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 HM 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 HM 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 HM configuration.

Also, the host configuration for the SBM is the same as HM because the ISATAP router addresses have to be reused in this example. Similar to the HM 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 16 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 16 SBM ISATAP Network Topology

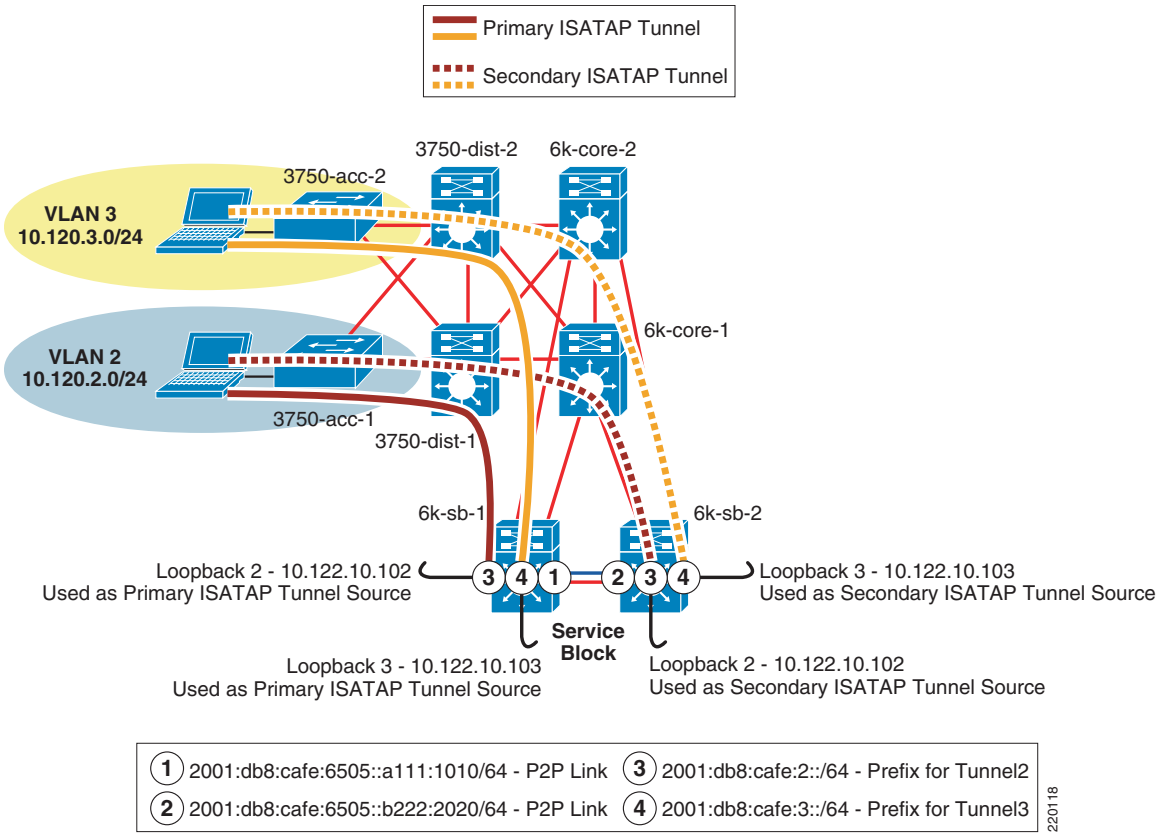

Figure 17 shows the manually-configured tunnel topology for the SBM. The topology diagram shows the loopback addresses on the service block switches (used as the tunnel source for configured tunnels) and the IPv6 addressing used on the manually-configured tunnel interfaces.

Figure 17 SBM Manually-Configured Tunnel Topology

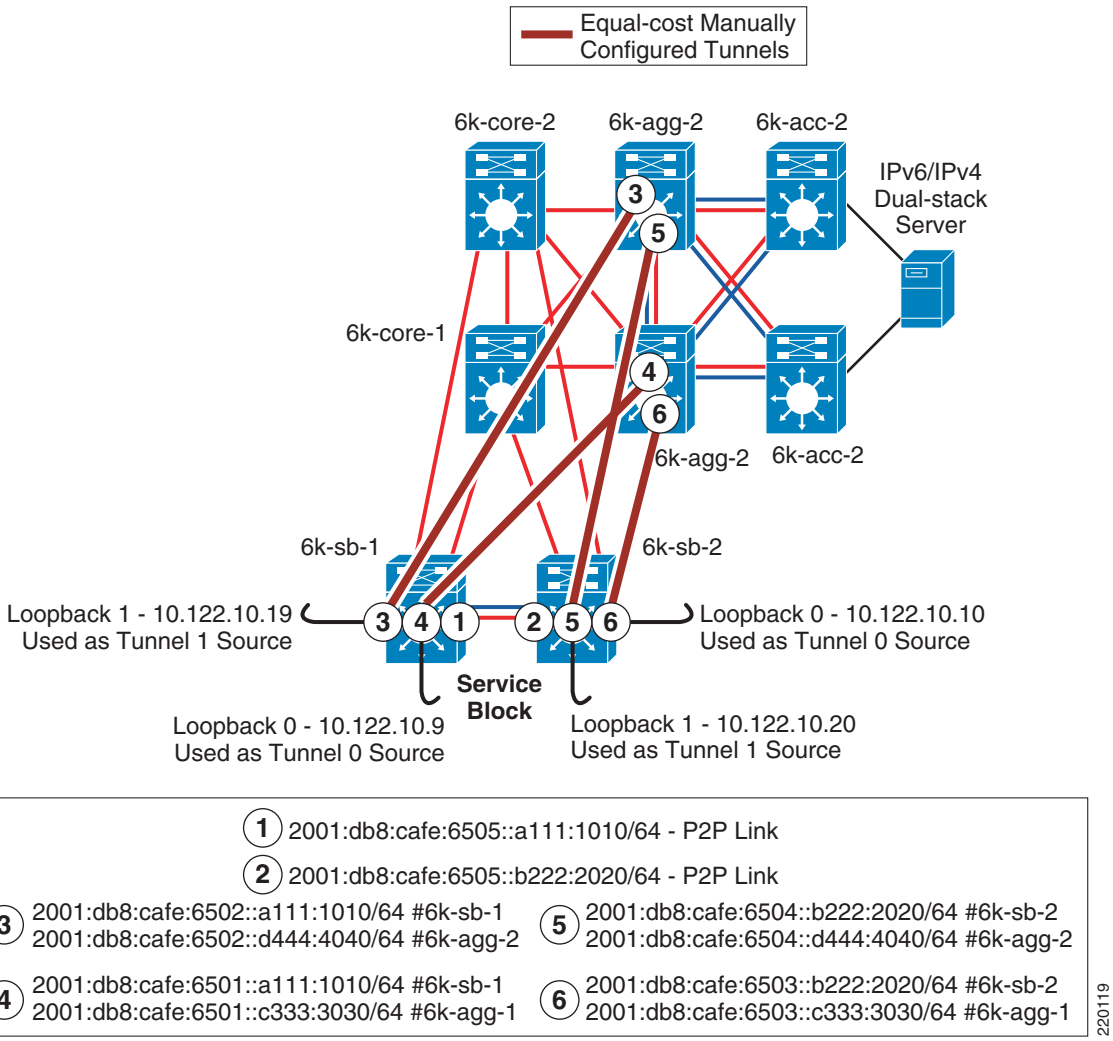

Physical Configuration

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.

- 6k-sb-1

```
interface GigabitEthernet4/1
 description to 6k-core-1
 ip address 10.122.0.78 255.255.255.252
 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
 mls qos trust dscp
!
```

```
interface GigabitEthernet4/2
 description to 6k-core-2
 ip address 10.122.0.86 255.255.255.252
 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
 mls qos trust dscp
!
interface GigabitEthernet4/3
 description to 6k-sb-2
 ip address 10.122.0.93 255.255.255.252
 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
 ipv6 address 2001:DB8:CAFE:6505::A111:1010/64        #p2p link between SBM switches
 ipv6 nd suppress-ra
 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
 ip address 10.122.0.82 255.255.255.252
 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
 mls qos trust dscp
!
interface GigabitEthernet4/2
 description to 6k-core-2
 ip address 10.122.0.90 255.255.255.252
 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
 mls qos trust dscp
!
interface GigabitEthernet4/3
 description to 6k-sb-1
 ip address 10.122.0.94 255.255.255.252
 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
 ipv6 address 2001:DB8:CAFE:6505::B222:2020/64          #p2p link between SBM switches
 ipv6 nd suppress-ra
 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 HM. 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 HM 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 DC 6k-agg-1
 no ip address
 ipv6 address 2001:DB8:CAFE:6501::A111:1010/64
 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 DC 6k-agg-2
 no ip address
 ipv6 address 2001:DB8:CAFE:6502::A111:1010/64
 ipv6 nd reachable-time 5000
 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
 ipv6 address 2001:DB8:CAFE:2::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 ospf 1 area 2
 tunnel source Loopback2
 tunnel mode ipv6ip isatap
!
interface Tunnel3
```

```
 description ISATAP VLAN3
 no ip address
 ipv6 address 2001:DB8:CAFE:3::/64 eui-64
 no ipv6 nd suppress-ra
 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
 ipv6 address 2001:DB8:CAFE:6503::B222:2020/64
 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
 ipv6 address 2001:DB8:CAFE:6504::B222:2020/64
 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
 ipv6 address 2001:DB8:CAFE:2::/64 eui-64
 no ipv6 nd suppress-ra
 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 nd suppress-ra
 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 the HM section apply to the SBM. Based on the example configuration shown in the case of HM, 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 HM 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)	No tunneling required No dependency on IPv4 (routing, QoS, HA, multicast, security, and management are separated) Superior performance and highest availability for IPv6 unicast and multicast Scalable	Requires IPv6 hardware-enabled campus switching equipment Operational challenges with supporting dual protocols—Training/management tools

Table 8 Benefits and Challenges of Various Models (continued)

Hybrid Model (HM)	Most of the existing IPv4-only campus equipment can be used (access and distribution layer) Per-user or per-application control for IPv6-service delivery Provides high-availability for IPv6 access over ISATAP tunnels	Tunneling is required; massive increase in operations and management Scale factors: • How many ISATAP tunnels are too many? • How many hosts per ISATAP tunnel are too many? • Ensure that the appropriate platform is used to support the total number of ISATAP connections IPv6 multicast is not supported Causes core layer to become an access layer for IPv6 tunnels Requires IPv6-enabled hosts with ISATAP configuration
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 Scale factors (see HM) 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. A similar document exists for IPv6 deployment in the branch: http://www.cisco.com/en/US/docs/solutions/Enterprise/Branch/BrchIPv6.html. Other documents will be published analyzing the deployment of IPv6 in the 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/designzone.

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
  - IPv6 in Enterprise Networks by Shannon McFarland, Muninder Sambi, Nikhil Sharma, Sanjay Hooda (ISBN-10:1-58714-227-9; ISBN-13: 978-1-58714-227-7) - http://www.ciscopress.com/bookstore/product.asp?isbn=1587142279
  - "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
  - IPv6 Security by Scott Hogg, Eric Vyncke (ISBN-10:1-58705-594-5; ISBN-13: 978-1-58705-594-2) - http://www.ciscopress.com/bookstore/product.asp?isbn=1587055945
  - Cisco IPv6— http://www.cisco.com/en/US/products/ps6553/products_ios_technology_home.html
  - Cisco Enterprise Design Zone— http://www.cisco.com/en/US/netsol/ns742/networking_solutions_program_category_home.html
  - Cisco Design Zone for Campus— http://www.cisco.com/en/US/netsol/ns815/networking_solutions_program_home.html
  - Enterprise QoS SRND— http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html
  - Catalyst 3750 Switch Software Configuration Guide, Release 12.2(55)SE—*Catalyst 3750-E and 3560-E Switch Software Configuration Guide* http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst3750e_3560e/software/release/12.2_55_se/configuration/guide/3750escg.html
  - Catalyst 4500 Switch Software Configuration Guide, Release 12.2(53)SG—http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst4500/12.2/53SG/configuration/config.html
  - Catalyst 6500 Switch Software Configuration Guide, Release 12.2(33)SXI—http://www.cisco.com/en/US/partner/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/book.html