

Federal Agencies and the Transition to IPv6

Introduction

Because of the federal mandate to transition from IPv4 to IPv6, IT departments must include IPv6 as a core element of their current and future IT strategies.

Although migrating from IPv4 to IPv6 may be strategically and technically challenging, many organizations are actually more prepared to begin the transition than they may realize: much of the current IPv4 infrastructure is IPv6 capable, and approximately one-third of the deployed desktop systems are IPv6 capable. Additionally, a key part of the IPv6 design is its ability to integrate into and coexist with existing IPv4 networks, which means that—as anticipated—IPv4 can and will coexist for the foreseeable future. Because there is no single day on which a planned outage will occur and by which all network changes must be complete, with careful planning, federal agencies can conduct business as usual throughout the long transition process.

This white paper introduces several specific transition mechanisms and campus designs that can help federal agencies with the planning process. The mechanisms and campus designs discussed in this document are based on the Cisco® High Availability Campus Design reference architectures (http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html).

Transition Mechanisms


Federal agencies can use a number of mechanisms to guide their transition from IPv4 to IPv6. Because each transition mechanism involves different implementation technologies and features, they offer varying degrees of effectiveness and success depending on the current state of the network. The most highly recommended mechanisms are dual stack, manually configured tunnels, and Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunnels.

Dual Stack

Dual stack (Figure 1) runs both IPv4 and IPv6 protocol stacks on a router in parallel, making it similar to the multiprotocol network environments of the past, which often ran Internetwork Packet Exchange (IPX), AppleTalk, IP, and other protocols concurrently. The technique of deploying IPv6 using dual-stack backbones allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. The IPv4 communication uses the IPv4 protocol stack, and the IPv6 communication uses the IPv6 stack.

As a transition strategy, dual stack is ideal for campus networks with a mixture of IPv4 and IPv6 applications. The technology is not difficult to implement, though it may require router upgrades because all routers must be dual stack with IPv4 and IPv6 addresses. Dual stack also requires access to the IPv6 Domain Name System (DNS) and enough memory for both IPv4 and IPv6 routing tables. Given that the dual management of protocols can introduce challenges such as memory and CPU exhaustion and additional security requirements, such potential challenges should be addressed proactively to help ensure as smooth and cost-effective a solution as possible.

Figure 1. Dual-Stack Example



Manually Configured Tunnels

Tunneling is the encapsulation of IPv6 traffic within IPv4 packets so that they can be sent over an IPv4 backbone, thus allowing isolated IPv6 end systems and routers to communicate without the need to upgrade the IPv4 infrastructure that exists between them. Manually configured tunneling is one of many tunneling techniques currently available (Figure 2).

Manually configured tunnels are used primarily as stable links for regular communication; they use standards-based security mechanisms such as those provided by IP Security (IPsec) to help ensure the security of that communication. With manually configured tunnels, an IPv6 address is manually configured on a tunnel interface, and the IPv4 addresses are manually configured at the tunnel source and the tunnel destination. Because they are configured one-to-one between well-known endpoints, manually configured tunnels make traffic information available for each endpoint; when IPsec is used, they also provide extra security against injected traffic.

This transition strategy is easily deployed over existing IPv4 infrastructures. However, it is important to note that because manually configured tunnels require configuration at both ends of the tunnel—and because the tunnel can be between two points only—they have a larger management overhead when multiple tunnels are implemented. The independence of the tunnel and link topologies also requires additional diagnostics. For this reason, manually configured tunnels are ideal for networks with a limited number of required tunnels.

Figure 2. Manually Configured Tunnel Example



ISATAP Tunnels

Intra-site Automatic Tunnel Addressing Protocol (ISATAP) is another tunneling technique (Figure 3). It differs from manually configured tunneling in that ISATAP tunnels are automatically defined rather than statically defined. Also, ISATAP tunnels are primarily used between hosts and routers whereas manually configured tunnels are used between routers, as described in the “Manually Configured Tunnels” section.

With ISATAP, a tunnel is created from a host, such as a PC, to a router or Layer 3 switch. The tunnel is established by using the IPv4 addresses of both the host and the router. The tunnel is “automatic” in that the host establishes the tunnel only when it needs to; the host is also able to discover the tunnel destination (that is, the router’s IPv4 address) dynamically through DNS or a local definition. Because both IPv4 and tunneled IPv6 packets are transported over a single


common IPv4 infrastructure, IPv4-dependent applications can continue to use IPv4 while newer IPv6-capable applications can be deployed immediately.

As a transition strategy, ISATAP is ideal for campus and branch sites because ISATAP allows organizations to activate IPv6 connectivity on the existing IPv4-only network while the infrastructure is gradually migrated to integrate native IPv6 capabilities. Thus, the immediate effect on the IPv4 support infrastructure is reduced to the configuration of one ISATAP router, though Cisco Systems® recommends the deployment of at least two ISATAP routers for high availability.

Additionally, because ISATAP allows native IPv6 connectivity to be activated first in the backbone, other parts of the IPv4 infrastructure can preserve their investment as they naturally evolve to support IPv6. ISATAP islands can also be created to allow gradual evolution to native IPv6 capabilities within different parts of the organization without blocking end-to-end IPv6 service deployments.

Although it has many benefits, ISATAP does not support IPv6 multicast, so this would not be the most appropriate transition strategy for organizations requiring that capability.

Figure 3. ISATAP Example



Campus Design Models

There are three campus design models that Cisco generally recommends for deploying IPv6, each of which uses one or more of the transition mechanisms described in the preceding sections. Cisco defines these three campus design models as dual stack, hybrid, and service block.

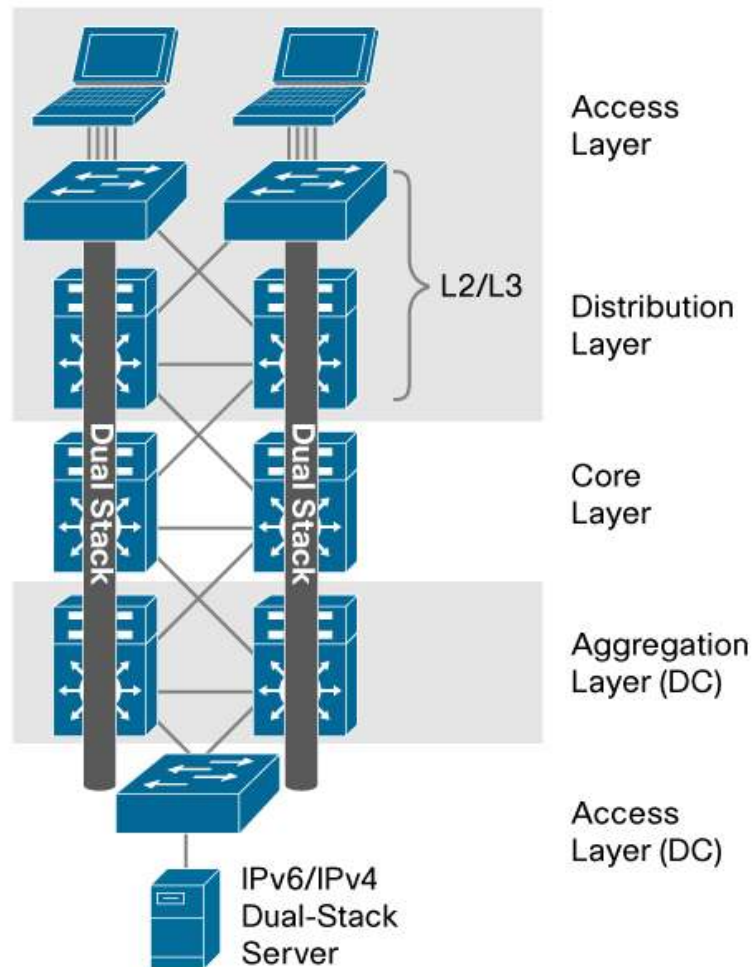
Dual Stack

The Cisco dual-stack campus design model uses the dual-stack transition strategy in its design (Figure 4).

To implement the dual-stack campus design model, each site should have an IPv6 unicast global or unique local prefix and appropriate entries in a DNS that map between host names and IP addresses for both IPv4 and IPv6. As previously noted, all routers in the network also need to be upgraded to dual stack so that IPv4 and IPv6 can run together over the same backbone. IPv4 communication uses the IPv4 protocol stack, with forwarding of IPv4 packets based on routes learned through running IPv4-specific routing protocols, whereas IPv6 communication uses the IPv6 stack with routes learned through the IPv6-specific routing protocols. Applications choose between IPv4 or IPv6 based on the response from the DNS, with the application selecting the correct address based on the type of IP traffic and the particular requirements of the communication.

Today, dual stack is a valid campus design model for specific network infrastructures with a mixture of IPv4 and IPv6 applications such as on a campus or an aggregation point of presence. A consideration for this approach—in addition to ensuring that network routers are IPv6 capable—is that the routers require a dual addressing scheme to be defined. Additionally, the IPv4 and IPv6 routing protocols require dual management and must both be configured with enough memory and CPU performance to support both the IPv4 and IPv6 protocols.

Figure 4. Campus Dual-Stack Example



Hybrid

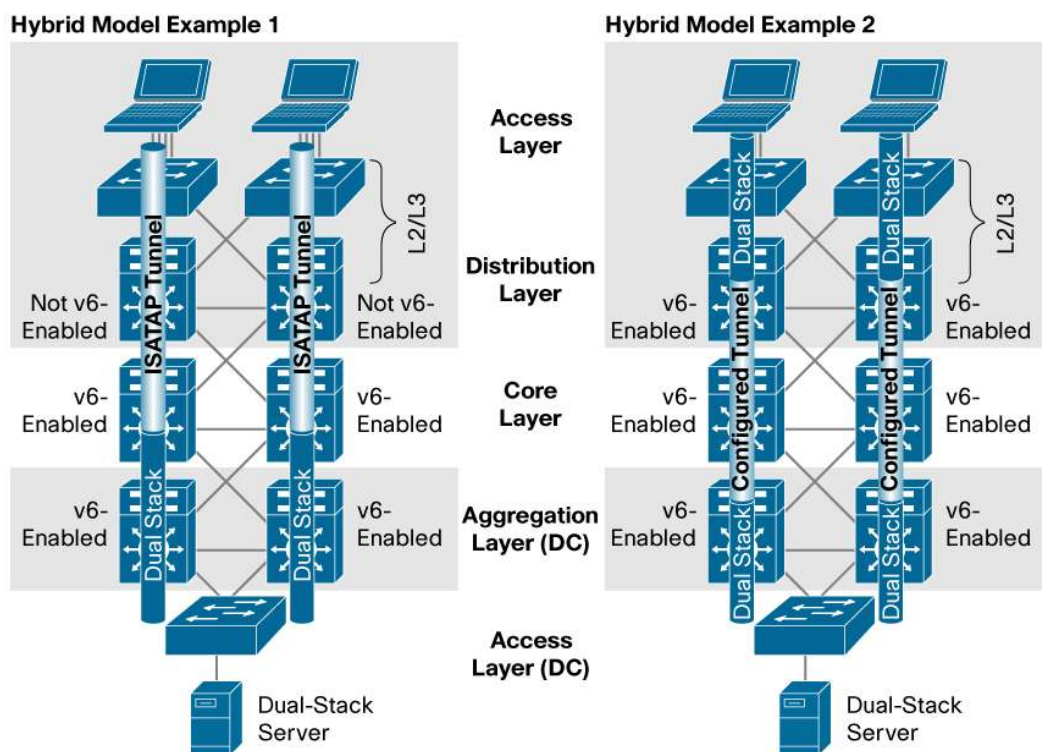
The Cisco hybrid campus design model (Figure 5) is a mix, or hybrid, that offers IPv6 connectivity using any or all of the aforementioned IPv6 transition mechanisms (dual stack, manually configured, and ISATAP). The particular combination of mechanisms depends on the hardware capabilities of the network.

The hybrid model allows organizations to use their existing network design and network equipment and offers a natural progression to full dual-stack design. It also offers a tremendous amount of flexibility for IPv6 deployment, as it allows IPv6 services to be delivered over a large-scale environment—or over a network that may have just a few IPv6-capable devices.

Depending on the hardware capability of the campus network components (mainly routers and switches), a transition mechanism can be deployed to provide access to IPv6-capable services. For example, if the campus distribution layer does not have Layer 3 IPv6 hardware capabilities, ISATAP tunnels can be used from the hosts to the core layer. The core layer, which may have IPv6 capabilities, would then use dual stack into the data center area of the network (see Figure 5, example 1). Alternatively, if the distribution and data center aggregation layers have Layer 3 IPv6 hardware capabilities but the core layer does not, dual stack can be used from the hosts to the distribution layer, and then manually configured tunnels can be used from the distribution to the data center aggregation layer.

There are some issues to be aware of with this campus design model, such as the fact that the tunnels—especially ISATAP—cause the infrastructure to function in “unnatural” ways; for example, with the core acting as the access layer.

Figure 5. Campus Hybrid Example



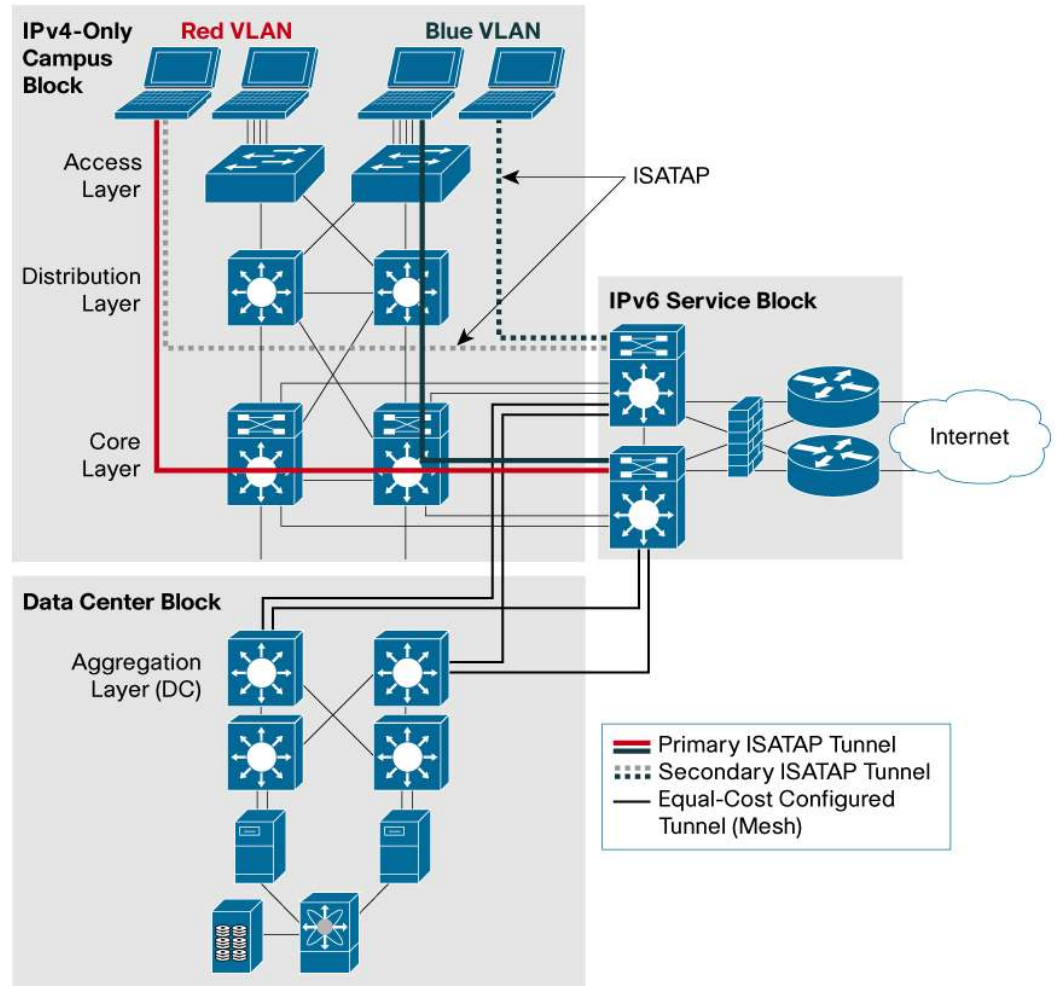
Service Block

The Cisco service block campus design model (Figure 6) offers interim connectivity for the IPv6 overlay network during the transition process. It is similar to the hybrid model in that it offers the same flexibility options and allows the deployment of the various transition mechanisms based on the application, location, or performance and scalability requirements. However, with the hybrid model, the existing network infrastructure devices are used for tunnel termination; with the service block model, all transition mechanisms terminate in the newly created service block. Also, whereas the hybrid campus design uses the network design as it currently exists, the service block campus design adds a new network layer (or block) to the existing network design that is used solely for IPv6 and the various transition mechanisms. In this way, it provides the capability to rapidly deploy IPv6 services without touching the existing network.

An important benefit of the service block campus design model is that as hardware is upgraded, the tunnels can be replaced with dual-stack connections, eventually freeing the service block equipment for other purposes. Consider, for example, an IPv6-capable application that is hosted in the data center, with hosts in the access layer of the campus that need access to that application. ISATAP tunnels can be used from the hosts, and manually configured tunnels can be used from the data center aggregation layer, both terminating on the service block switches. As network components such as core layer switches become dual-stack capable, the tunnels can be replaced with dual-stack links. Eventually, the service block can be removed entirely and the network equipment repurposed.

Note that because the service block campus design does not fully use the existing infrastructure design, it is more costly because organizations are required to purchase additional equipment. Furthermore, this model is not a substitute for a real dual-stack deployment, which means that organizations still must make appropriate plans for such a deployment.

Figure 6. Campus Service Block Example



For More Information

To learn more about IPv6 transition mechanisms—dual stack, manually configured tunnels, ISATAP tunnels, and other mechanisms that Cisco supports—or for additional information about campus designs that Cisco recommends, please visit <http://www.cisco.com/ipv6>.

Cisco will also publish a series of IPv6 implementation documents at <http://www.cisco.com/ipv6> and <http://www.cisco.com/go/srnd> as its work on IPv6 reference models continues.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)