IPv6 for Enterprise Networks

The practical guide to deploying IPv6 in campus, WAN/branch, data center, and virtualized environments

Shannon McFarland, CCIE® No. 5245
Muninder Sambi, CCIE No. 13915
Nikhil Sharma, CCIE No. 21273
Sanjay Hooda, CCIE No. 11737

ciscopress.com

IPv6 for Enterprise Networks

Shannon McFarland
Muninder Sambi
Nikhil Sharma
Sanjay Hooda

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

- **Networks:** Assess what types of networks are present in the enterprise. Possible network types include IP, wireless, DSL, VoIP, and CPE. A complete network equipment inventory should be done.
- **Network services:** Determine what network services are running today in the enterprise. For example, DNS, AAA, DHCP, NTP, and so on.
- **Network management:** Assess how the network would be managed after IPv6 is deployed on the network. Some tools that can be used for network management include NetFlow, MIBS, SNMP, and so on.
- **Network applications:** Assess how network applications such as VoIP, databases, and so on interact with IPv6. A software and operation system inventory should be done.
- **Other IP-based/aware services:** Other services, such as location and mobility, will use IPv6 for communications.

Design

Develop an overall IPv6 design for various impacted areas. The IPv6 design plan should be standards-based and should provide as much IPv4 feature parity as applicable to support a smooth transition. The design should take into account any new networks and services as well as traffic growth that an enterprise foresees. The following sections describe areas that can be considered when developing an IPv6 design.

IPv6 Addressing Plan

Understand addressing requirements within the enterprise. As described in this section, the addressing plan should outline the enterprise's IP addressing needs for the next few years and the address allocation, management, and acquisition processes. For example:

- Examine addressing needs for an enterprise's own infrastructure, its intranet, its extranet, sites not managed by the organization, and services (for example, Layer 3 VPNs) that it supports or offers, and forecast IP address usage.
- Develop a plan of IPv6 address allocation that outlines how IPv6 addresses should be allocated to support infrastructure and end users and provide efficient and scalable network routing. As part of the allocation, decide whether an enterprise uses Stateless Address Autoconfiguration (SLAAC) or Stateful Configuration for the hosts.
- Determine the management of IPv6 privacy extensions. This enables hosts to use different IPv6 source addresses with time; for example, using a different IPv6 address daily.
- Determine the impact of IPv6 routing, its integration, and changes with the existing IPv4 routing.
- Allocate IPv6 addresses to external connections to the Internet and other sites. This means working with Internet service providers (ISP) on IPv6 peering. Plan end-to-end site connectivity using IPv6 as well.

Chapter 12 provides additional information and a sample of an IPv6 addressing plan.

Transition Mechanisms

Because IPv4 and IPv6 can coexist for a substantial period of time during the transition to IPv6, an organization needs to consider the various transition mechanisms that can facilitate the transition to IPv6 while coexisting with the IPv4 network environment. When selecting a transition mechanism, an organization needs to consider its current network environment, IPv6 traffic forecast, IPv6-capable devices or applications, and IPv6 deployment plan. Transition mechanisms fall into three major categories: dual-stack, tunneling, and translation, as explained in Chapter 3, "Common IPv6 Coexistence Mechanisms." These transition mechanisms also need to be included in the pilot to understand their behavior on the network and provide a learning experience.

Network Services

Network services will be impacted as you deploy IPv6. Understanding these changes is key to having an operational IPv6 network. Network services include Domain Name System (DNS); Dynamic Host Configuration Protocol (DHCP); authentication, authorization, and accounting (AAA); and Network Time Protocol (NTP). For example, an organization might need to decide whether to support autoconfiguration or DHCPv6 or both. It can consider implementing dual-stack DNS to support IPv6 address queries as well as existing IPv4 queries. As part of the pilot, these network services also need to be deployed, especially if their IPv4 counterparts exist in the network today to provide end-to-end services.

Security

In addition to the similar security threats that exist in the IPv4 world, an enterprise needs to protect new threats that arise during the transition to IPv6. For example, existing firewalls and network intrusion detection systems do not provide IPv6 detection or filtering capabilities and might need to be upgraded. Malicious users might be able to tunnel IPv6 traffic through these security devices undetected. Some of the automation features that reduce operational overhead can increase vulnerabilities.

For example, malicious users can spoof solicitation, advertisement, and binding messages. Applications that use automated tunneling can traverse firewalls and therefore expose the network to the outside world. To minimize these problems, mechanisms and policies need to be developed to secure an IPv6 network.

The pilot should include how to secure the network using IPv6, such as using IPsec with IPv6.

New Features in IPv6

IPv6 introduces several new features such as Stateless Address Autoconfiguration (SLAAC), "built in" IPsec, Mobile IPv6, flow label, and more. The pilot presents the perfect opportunity to evaluate these features. Each enterprise can use some or all of these features. A deep understanding of these features will help in determining IPv6 policies for an enterprise.

As an example, SLAAC enables hosts to obtain an IPv6 address automatically when connected to a routed IPv6 network. However, this might be unsuitable for an application, where a stateful configuration using DHCPv6 or static IPv6 address might have to be used.

Scalability and Reliability

As enterprises introduce IPv6 to their existing networks, they must ensure that the IPv6 design itself is a scalable and reliable design. In addition, they must ensure that the IPv6 solution does not introduce negative impacts to the existing network environment.

Dual-stack implies that there would be two protocol stacks running on devices. The overhead that dual-stack brings to devices needs to be looked at. Running dual-stack could, for example, increase CPU utilization on certain devices.

In terms of scalability, the number of routes and neighbors needs to be addressed. This can be worked out with the vendor to determine IPv6 scalability on the device.

Service Level Agreements

Develop service level agreements (SLA) that reflect any new IPv6 policy and include transition mechanisms because they often impact the ability to adhere to a tight SLA. Commonly, enterprises will try to match their existing SLAs and simply add IPv6 to them. This often fails because of the enterprise and service provider(s) having to, for a period of time, use transition mechanisms that introduce additional latency, loss, and jitter to the network.

Lessons Learned and Implementation

After the pilot, it is important to understand the results. Data gathered during the pilot provides information on how an enterprise can migrate to IPv6. The pilot also provides the opportunity for the enterprise to learn IPv6 as a technology. Some features will prove beneficial to use, whereas some features will need to be turned off, depending on the results of the pilot.

Going forward, organizations will need to develop their own IPv6 compliance standard for each device and application in the enterprise. These standards should be based on best industry standards. Some standards/certifications that an enterprise can follow are

USGv6 or the IPv6-ready logo. For those devices and applications that need to transition to IPv6, do the following:

- **Assess when they will be IPv6-ready.** This step might require contacting vendors for the availability of their IPv6 product road map.
- **Identify the impact to the supporting services and customers.** This might require establishing a relationship between the systems integrators and vendors. A vendor might provide a product or service that is IPv6 capable. A system integrator would provide different products or applications that would be IPv6 capable. It is important to understand the interaction between applications and devices in an IPv4/IPv6 environment.

After a successful pilot, develop a plan to deploy the IPv6 transition throughout the enterprise. The results of the test plan will play a key role in deciding which IPv6 features are to be used, how they are to be configured, and how they need to be managed.

Start by developing a list of projects to be implemented. Within that list, identify dependencies and prioritize the projects and tasks appropriately. Leverage the IPv6 testing/lab environment to verify architecture, designs, business rules, and so on.

Next, deploy IPv6 during an infrastructure refresh cycle and ensure that all devices procured or developed are IPv6 capable and, more importantly, have the features that the organization needs based on the outcome of the pilot.

Client/Server IPv6 Migration Scenarios

There is certainly no shortage of technology options when deploying IPv6. Selecting the right path will depend on the current environment in terms of end-user devices and operating systems, router models and versions, key applications, budget and resources, and schedule constraints.

The following sections review some starting points of migration scenarios. In reviewing these, the diagram illustrated in Figure 5-1 will be used as a baseline and depicts an all IPv4-only network. Three basic scenarios are discussed. Depending on the complexity of the network, business policies, and various other factors, any of the scenarios can be followed.

Tthis figure has a client with IPv4 applications, IPv4 sockets API and TCP/IPv4 stack, and a server with a comparable configuration. The interconnection network is split into access networks for the client and server, respectively, and a core or backbone network.

This basic diagram illustrates a pair-wise, client/server connection. For a given use case, this could represent an internal client accessing an internal server through an all-internal access and core network as well as over the Internet.

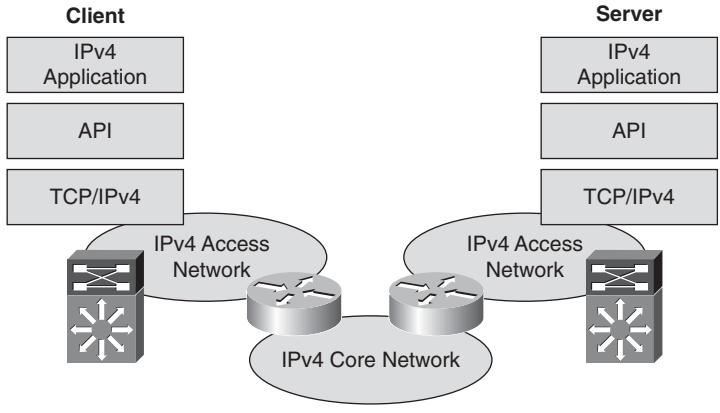

Figure 5-1 *Typical IPv4 Client/Server Model Using IPv4*

IPv6 Core Deployment: "Start at the Core"

This scenario starts by deploying IPv6 in the backbone or core network and slowly moves the IPv6 deployment toward the edges of the network, as illustrated in Figure 5-2.

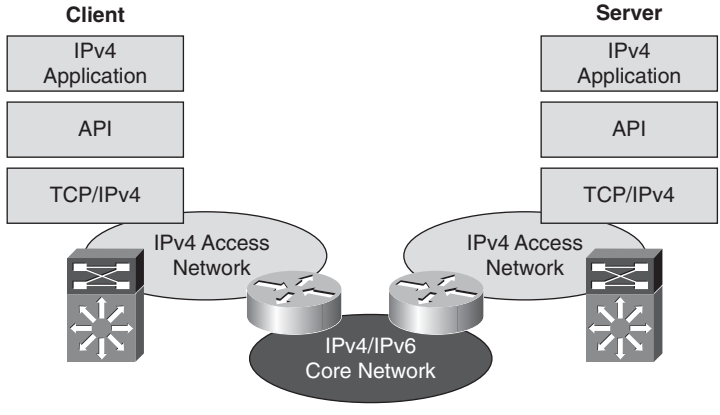

Figure 5-2 *IPv6 Core Deployment Scenario*

This topology requires enabling core routers to support dual-stack IPv4 and IPv6 routing and routing protocols. By enabling a dual-stack core, the enterprise can learn the protocol, routing differences, and many of the other differences with IPv6 without interrupting production IPv4 traffic between endpoints. This is a great option if there is no clear timeline for IPv6 deployment on the endpoints themselves. The core network could be an internal backbone or an IPv6 ISP network, in which case you might need to use 6PE or 6VPE mechanisms.

Localized IPv6 Server-Side Deployment

The localized server-side scenario involves upgrading servers and application hosts to dual-stack implementations. With the server still able to support IPv4 communications and applications, end clients connect as before through IPv4. However, the server would be able to serve IPv6 clients as well when they are dual-stack or IPv6-only enabled. At this point, unless the client is IPv6 enabled, either the client-to-server connection is IPv4-only or a translation or proxy device would have to be deployed. This option works best when a focus needs to be placed on the data center applications, operating systems, and networking components and their support of IPv6. This often occurs before any client connectivity over IPv6 is supported. Many enterprises will start on the server-side simply because it is a more complex area of deployment, and the operating system, applications, and supporting network infrastructure might require a more comprehensive set of IPv6 capabilities than any other area of the network such as the campus. One point to note in this topology is that certain network services, for example, DNS, will be enabled only when end-to-end IPv6 support is available. Figure 5-3 illustrates this configuration.

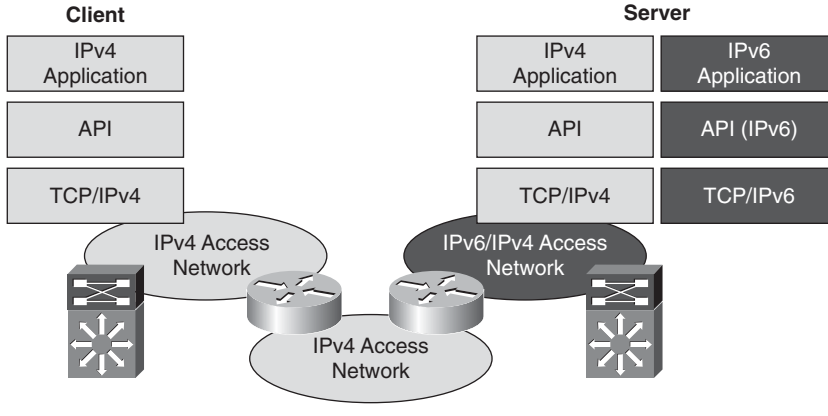

Figure 5-3 *Server-Side IPv6 Deployment Scenario*

Client-Side Deployment

Another scenario is to enable clients for dual-stack operation as well as to access network routers. This is more of a preparation step just like in the server-side scenario. Because end-to-end IPv6 connectivity is not enabled, the clients won't really be accessing an IPv6-enabled service unless they use a tunneling or translation mechanism. This option enables the endpoints to be dual-stack enabled and then, at some point, the end-to-end dual-stack configuration can be made (as seen in Figure 5-5). Existing IPv4 client devices would be supplemented with IPv6 applications, API, and TCP/IP stack. Figure 5-4 illustrates this configuration.