IPv6 Integration and Coexistence Strategies for Next-Generation Networks

Mallik Tatipamula and Patrick Grossetete, Cisco Systems
Hiroshi Esaki, University of Tokyo

Abstract

IPv6 has been designed, among other things, to provide an expanded address space to satisfy the future networking requirements. In this article we analyze and discuss important aspects of IPv6 deployment scenarios, and propose the system architecture coexisting and integrating with IPv4/MPLS networks. We investigate on various IPv6 deployment strategies along with network design examples, comparing these techniques. Then the IPv6 deployment in service provider environments is proposed.

Introduction

The continuous growth of the global Internet requires that its overall architecture evolve to accommodate new technologies to support the growing numbers of users, applications, appliances, and services. IPv6 is designed to satisfy these requirements and allow a return to a global end-to-end environment where the addressing rules of the network are again transparent to applications. The current IP address space is unable to satisfy the potentially huge increase in number of users or the geographical needs of Internet expansion, let alone the requirements of emerging applications such as Internet-enabled personal digital assistants (PDAs), home area networks (HANs), Internet-connected transportation, integrated telephony services, and distributed gaming. IPv6 quadruples the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides more than enough globally unique IP addresses for every network device on the planet. The use of globally unique IPv6 addresses simplifies the mechanisms used for reachability and end-to-end security for network devices, functionality crucial to the applications and services driving the demand for the addresses. The lifetime of IPv4 has been extended using techniques such as address reuse with translation and temporary use allocations. Although these techniques appear to increase the address space and satisfy the traditional client/server setup, they fail to meet the requirements of new applications. The need for always-on environments (e.g., residential Internet through broadband, cable modem, or Ethernet to the home) to be contactable precludes these IP address conversion, pooling, and temporary allocation techniques, and the "plug-and-play" required by consumer Internet appliances further increases address requirements. The flexibility of the IPv6 address space provides the support for private addresses but should reduce the use of Network Address Translation (NAT) because global addresses are widely available. IPv6 reintroduces end-to-end security that is not always readily available throughout a NAT-based network.

We are in the early stages of IPv6 deployment, with fewer IPv6 applications than IPv4 on the market and networking products needing to make trade-offs between available IPv6 services. Although the success of IPv6 will depend ultimately on the innovative applications that run over IPv6, a key part of IPv6 design is its ability to integrate into and coexist with existing IP networks. It is expected that IPv4 and IPv6 hosts will need to coexist for a substantial time during the steady migration from IPv4 to IPv6, and the development of transition strategies, tools, and mechanisms has been part of the basic IPv6 design from the start. Selection of a deployment strategy, or strategies, will depend on current network environment, and factors such as the forecast amount of IPv6 traffic and the availability of IPv6 applications on end systems, and the stage in deployment. This article summarizes various strategies for IPv6 integration/coexistence along with network design examples. We propose a system architecture coexisting and integrating with IPv4/multiprotocol label switching (MPLS)

0163-6804/04/$20.00 © 2004 IEEE

networks. We discuss briefly IPv6 network design consideration for service provider network environments along with a comparison of these deployment strategies.

IPv6 Integration and Coexistence Strategies

The successful market adoption of any new technology depends on its easy integration with the existing infrastructure without significant disruption of services. Several Internet Engineering Task Force (IETF) working groups (e.g., IPv6, v6ops) have been active in defining strategies for the deployment of IPv6. The following deployment scenarios are being discussed in this article:

- Dual-stack backbones
- IPv6 over IPv4 tunnels
- Protocol translation mechanisms
- Dedicated data links
- MPLS backbones

This article briefly revisits and compares the first three deployment scenarios, covered in detail in [1]. IPv6 over dedicated data links and MPLS backbone scenarios are proposed and discussed in this article.

Transition Mechanism Overview

Focusing on the primary goal, to enable IPv6 applications on hosts to communicate, many network designers recommend deploying IPv6 at the edge first, where the applications and hosts reside, and then moving toward the network core to reduce the cost, operational instability, and impact of integration. Also, the migration of IPv6 into the edge or user site is relatively easier, as major operating systems (e.g., Microsoft, Linux) are already IPv6-capable.

The key strategies used in deploying IPv6 at the edge of a network involve carrying IPv6 traffic over an IPv4 network infrastructure, allowing isolated IPv6 domains to communicate with each other before the full transition to a native IPv6 backbone. Then later, when it is time to plan a full upgrade, it is possible to run both IPv4 and IPv6 throughout the network, from all edges through the core. Additionally, a mechanism may be required to translate between IPv4-only and IPv6-only devices to allow hosts supporting only one protocol to communicate transparently with hosts running the other. All techniques allow networks to be upgraded and IPv6 deployed *incrementally* with little or no disruption of IPv4 services. The four key techniques for deploying IPv6 are as follows.

Deploying dual-stack backbones [1, 2]: This technique allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. All routers (e.g., access customer premises equipment, aggregation and core routers) in the network need to be upgraded to be dual-stack, with IPv4 communication using the IPv4 protocol stack and IPv6 communication using the IPv6 stack. Routing protocols for both IP versions must be selected and configured adequately; interior gateway protocol (IGP) selection is between a "ship in the night" solution (e.g., OSPFv2 for IPv4 and OSPFv3 for IPv6), or an integrated solution (e.g., IS-IS), mandating IPv4 and IPv6 topologies to be aligned.

Deploying IPv6 over IPv4 tunnels [1, 2]: These tunnels encapsulate IPv6 traffic within IPv4 packets, and are primarily for communication between isolated IPv6 sites or connection to remote IPv6 networks over an IPv4 backbone. The techniques include using manually configured tunnels, generic routing encapsulation (GRE) tunnels, semiautomatic tunnel mechanisms such as tunnel broker services, and fully automatic tunnel mechanisms such as 6to4 for the wide area network (WAN) and intrasite automatic tunnel addressing protocol (ISATAP) for the campus environment. This is an easy scenario for network managers who want to get familiar with IPv6 technology.

Deploying IPv6 over dedicated data links: This technique enables IPv6 domains to communicate by using the same layer 2 infrastructure used for IPv4, but with IPv6 using separate frame relay or asynchronous transfer mode (ATM) permanent virtual circuits (PVCs), separate optical links, or lambdas in dense wavelength-division multiplexing (DWDM).

Deploying IPv6 over MPLS backbones: This technique allows IPv6 domains to communicate with each other, but over an IPv4 MPLS backbone without modifying the core infrastructure. Multiple techniques are available at different points in the network, but each requires little change to the backbone infrastructure or reconfiguration of the core routers because forwarding is based on labels rather than the IP header itself.

Deploying IPv4/IPv6 Dual Stack

Dual-stack backbone [1, 2] is a basic strategy for routing both IPv4 and IPv6. Applications, that are not upgraded to support the IPv6 stack, can coexist with upgraded applications on the same end system. A new application programming interface (API) has been defined to support both IPv4 and IPv6 addresses and Domain Name Service (DNS) requests. Applications choose between using IPv4 or IPv6 based on name lookup; both IPv4 and IPv6 addresses may be returned from the DNS, with the application (or the system according to the rules defined in the IETF document "Default Address Selection for IPv6") selecting the correct address based on the type of IP traffic.

With dual-stack backbone deployment, all routers in the network need to be upgraded to be dual-stack. Today, dual-stack routing is a valid deployment strategy for specific network infrastructures with a mixture of IPv4 and IPv6 applications (e.g., on a campus or an aggregation point of presence), requiring both protocols to be configured. However, apart from the obvious need to upgrade all routers in the network, network managers selecting this approach must be aware that all the routers require a dual addressing scheme to be defined, require dual management of the IPv4 and IPv6 routing protocols, and must be configured with enough memory for both the IPv4 and IPv6 routing tables.

Dual-stack backbone is a basic strategy for routing both IPv4 and IPv6. Applications, that are not upgraded to support the IPv6 stack, can coexist with upgraded applications on the same end system.

Mechanism	Primary use	Benefits	Limitations	Requirements
IPv6 manually configured tunnels	Stable and secure links for regular communication Connection to Internet IPv6	Well-known standard tunnel technique demonstrated for years on the 6Bone Tunnel endpoints can be secured using IPv4 IPsec	Tunnel between two points only. Large management overhead.	ISP registered IPv6 address. Dual stack router
IPv6 over IPv4 GRE tunnel	Stable and secure links for regular communication	Well-known standard tunnel technique Tunnel endpoints can be secured using IPv4 IPsec	Tunnel between two points only. Management overhead. GRE tunnel implementation is rarely available on hosts.	ISP-registered IPv6 address. Dual stack router. Required with IS-IS for IPv6 is configured over a tunnel
Tunnel broker	Standalone isolated IPv6 end systems	Tunnel set up and managed by ISP	Potential security implications.	Tunnel broker service must know how to create and send a script for software
Automatic IPv4 compatible tunnel	Single hosts or small sites Infrequent communication	Automatic tunnel	Communication only with other IPv4-compatible sites Does not scale well as it only offers the same address space as IPv4, nearly deprecated as 6to4 is a preferred solution.	IPv6 prefix (0::/96) Dual stack router IPv4 addresses required to each host.
Automatic 6to4 tunnel	Connection of multiple remote IPv6 domains Frequent communication	Easy to set up with no management overhead	When communicating with the IPv6 Internet, return path selection is not optimized Potential security issue if not secured through IPsec (either IPv4 or IPv6)	IPv6 prefix (2002::/16). Dual stack router
ISATAP tunnels	Campus sites Transition of nonrouted sites	Ease IPv6 deployment for a sparse IPv6 host population on a campus	May not offer the best performance path compared to native IPv6 layer 3 switch Does not offer a solution for IPv6 multicast traffic	ISATAP implementation on IPv6 hosts and router Dual stack router
6over4 tunnels	Campus sites Transition of non-routed sites	Ease IPv6 deployment for a sparse IPv6 host population on a campus	Deprecated, replaced by ISATAP Requires IPv4 multicast	N/A

Table 1. *Comparison of various tunneling mechanisms.*

DEPLOYING IPV6 OVER IPV4 TUNNELS

Tunneling is one of the key deployment strategies for both service providers and enterprises during the period of IPv4 and IPv6 coexistence, as has been demonstrated over the 6Bone for years.

A variety of tunnel mechanisms are available for deploying IPv6. These mechanisms [1, 2] include manually created tunnels such as IPv6 configured tunnels and IPv6 over IPv4 GRE tunnels, semiautomatic tunnel mechanisms such as those employed by tunnel broker services, and fully automatic tunnel mechanisms such as ISATAP and 6to4 tunnels.

All tunneling mechanisms require that the endpoints of the tunnel run in dual-stack mode. The dual-stack routers run both IPv4 and IPv6 protocols simultaneously and thus can interoperate directly with both IPv4 and IPv6 end systems and routers.

Not all transition strategies will be applicable to all situations and all networks. Because it is expected that, at least initially, most customers might be interested in tunneling IPv6 over their existing IPv4 networks, this section compares the following IPv6 tunneling techniques to be used over IPv4 networks.

- IPv6 manually configured tunnel [1, 2]
- IPv6 over IPv4 GRE tunnel [1, 3]
- Automatic IPv4-compatible tunnel [1, 2]
- Automatic 6to4 tunnel [4]
- ISATAP tunnel [5]
- Teredo tunnel [6]

Table 1 summarizes the features of all tunnel mechanisms listed above. Each mechanism has pros and cons. However, one of the important facts based on observation of Table 1 is that even without manual configuration, we can oper-

Mechanism	Primary use	Benefits	Limitations	Requirements
NAT-PT	IPv6 only hosts to IPv4 only hosts	No dual stack	No end-to-end IPSec Dedicated server is single point of failure NAT-PT requires an ALG for application that embeds an IP address	Dedicated server DNS with IPv6 support
TCP-UDP relay	Translation between TCP/UDPv6 and TCP/UDPv4 sessions	Freeware	No end-to-end IPSec Dedicated server is single point of failure	Dedicated server. DNS with IPv6 support
BIS	IPv4 only hosts communicating with IPv6 only hosts	End system implementation	All stacks must be updated	Updated IPv4 protocol stack
SOCKS-based IPv6/IPv4 gateway	IPv6 only hosts to IPv4 only hosts	Freeware	Requires additional software in the gateway	Client and gateway software in the host and router

Table 2. *A comparison of protocol translation mechanisms.*

ate IPv6 end stations and campuses over IPv4 cloud using the above tunneling mechanisms. To secure the IPv6 over IPv4 tunnels configuration, network managers can configure IPsec either for IPv4 or IPv6 on the endpoint routers.

IPv6/IPv4 Translation Mechanisms

All of these integration strategies provide IPv6 end to end. However, some organizations or individuals might not want to implement any of these IPv6 transition strategies. And some organizations or individuals might install only IPv6 in their nodes or networks, but might not implement dual-stack. Even if some nodes or networks do install dual-stack, these nodes might not have IPv4 addresses to be used with the dual-stack nodes. Under these circumstances, intercommunication between IPv6-only hosts and IPv4-only hosts require some level of translation between the IPv6 and IPv4 protocols on the host or router, or dual-stack hosts, with an application-level understanding of which protocol to use. For example, an IPv6-only network might still want to be able to access IPv4-only resources, such as IPv4-only Web servers.

A variety of IPv6-to-IPv4 translation mechanisms [1] are under consideration by the IETF v6Ops Working Group, as follows:

- NAT-Protocol Translation (NAT-PT)
- TCP-UDP relay
- Bump-in-the-stack (BIS)
- SOCKS-based gateway

These protocol translation mechanisms become more relevant as IPv6 becomes more prevalent, and as IPv6 becomes the protocol of choice to allow legacy IPv4 systems to be part of the overall IPv6 network. The translation mechanisms tend to fall into two categories; those that require no changes to either the IPv4 or IPv6 hosts, and those that do. An example of the former is the TCP-UDP relay mechanism [1] that runs on a dedicated server and sets up separate connections at the transport level with IPv4 and IPv6 hosts, and then simply transfers information between the two. An example of the latter is the BIS mechanism [1] that requires extra protocol layers to be added to the IPv4 protocol stack. In the BIS mechanism, three extra layers (name resolver extension, address mapper, and translator) are added to the IPv4 protocol stack between the application and network layers. Whenever an application needs to communicate with an IPv6 only host, the extra layers map an IPv6 address into the IPv4 address of the IPv4 host.

In addition to the strategies for deploying IPv6 within an IPv4 environment, one also needs protocol translation mechanisms, such as NAT-PT [1], to allow communication between applications using IPv4 and those using IPv6 (e.g., to enable IPv6-only Web browsers to communicate with IPv4-only Web servers or dual-stack), but one drawback — well known from NAT users — is the need for dedicated application layered gateways (ALGs) when an application payload embeds an IP address.

The SOCKS-based IPv4/IPv6 gateway mechanism [1] is based on a mechanism that relays two "terminated" IPv4 and IPv6 connections at the application layer. It consists of additional functionality in both the end system (client) and dual-stack router (gateway) to enable communication between IPv4 and IPv6 nodes. This mechanism is based on the SOCKSv5 protocol and inherits all its features.

These translation mechanisms may be helpful as IPv6 deployment moves from the testing to the actual usage phase, and more relevant as application developers decide that continuing to support IPv4 is not cost effective. Eventually, as IPv6 becomes the protocol of choice, these mechanisms will allow legacy IPv4 systems to be part of the overall IPv6 network. The mechanisms translate between IPv4 and IPv6 on end systems, dedicated servers, and routers within the IPv6 network, and together with dual-stack hosts provide a full set of tools for the incremental deployment of IPv6 with no disruption to the IPv4 traffic. Table 2 provides a comparison of translation mechanisms.

Routers attached to the ISP WANs or MANs can be configured to use the same Layer 2 infrastructure as for IPv4, but to run IPv6, for example, over separate ATM or Frame Relay PVCs or separate optical lambda.

Deploying IPv6 over Dedicated Data Links

Many WANs and metropolitan-area networks (MANs) have been implemented by deploying layer 2 technologies such as frame relay, ATM, or optical, and some are beginning to use DWDM. Figure 1 shows a sample configuration for IPv6 over dedicated data links.

Routers attached to Internet service provider (ISP) WANs or MANs can be configured to use the same layer 2 infrastructure as IPv4, but to run IPv6, for example, over separate ATM or frame relay PVCs or separate optical lambda. This configuration has the added benefit for the service provider of not jeopardizing IPv4 revenue and traffic by the integration of IPv6 even utilizing the layer 2 infrastructure.

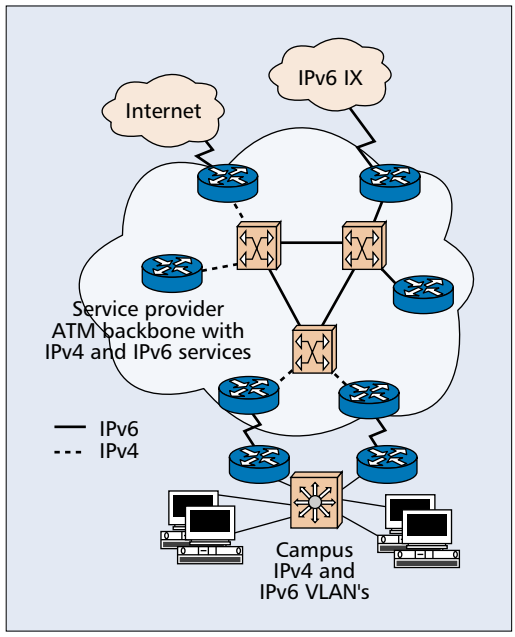

Figure 1. *IPv6 deployment over a dedicated data link.*

Deploying IPv6 over MPLS Backbone

IPv6 over MPLS backbone enables IPv6 domains to communicate with each other over an IPv4 MPLS core network. This implementation requires far fewer backbone infrastructure upgrades and no reconfiguration of core routers, because forwarding is based on labels rather than the IP header itself, providing a very cost-effective strategy for the deployment of IPv6. Additionally, the inherent virtual private network (VPN) and traffic engineering (TE) services available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure supporting IPv4 VPNs and MPLS-TE.

A variety of deployment strategies are available or under development, as follows:

- IPv6 tunnels on customer edge (CE) routers
- Layer 2 circuit transport over MPLS
- IPv6 on provider edge (PE) routers (6PE)
- Adding IPv6 MPLS VPNs to 6PE (6VPE)
- Native IPv6 MPLS-based backbone (MPLS control plane is IPv6-based)

As shown in Fig. 2, the first of these strategies has no impact on and requires no changes to the MPLS core consisting of provider (P) and PE routers. It is because this strategy uses IPv4 tunnels on dual-stack CE routers, as previously discussed, to encapsulate the IPv6 traffic, thus appearing as IPv4 traffic within the MPLS network. The second strategy requires no change to the core routing mechanisms. The third and fourth strategies require changes to the PE routers to support a dual-stack implementation, but all the core functions of P routers remain IPv4. A final strategy would be to run a native IPv6 MPLS core, but this strategy would require a full network upgrade to all P and PE routers, with dual control planes for IPv4 and IPv6. Table 3 provides a comparison of these strategies for transporting IPv6 over an MPLS backbone. The following sections describe each mechanism in more detail.

IPv6 over Layer 2 Circuit Transport over MPLS

Using any layer 2 circuit transport for deploying IPv6 over MPLS networks has no impact on the operation or infrastructure of MPLS. It requires no changes to either the P routers in the core or the PE routers (to support one of the layer 2 circuit transport over MPLS mechanisms) connect-

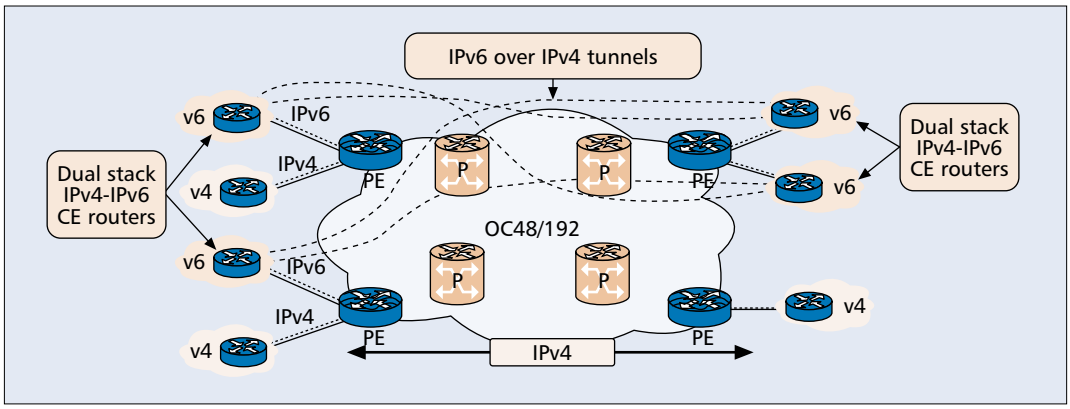

Figure 2. *IPv6 deployment using tunnels on the CE routers.*

Mechanism	Primary use	Benefits	Limitations	Requirements
IPv6 using tunnels on CE routers	Enterprise customers wanting to use IPv6 over existing MPLS services	No impact on MPLS infrastructure	Scalability issue when the number of tunnels grow between CEs	Dual-stack CE routers
IPv6 over a circuit transport over MPLS	Service providers with ATM or Ethernet links to CE routers	Fully transparent IPv6 communication	No mix of IPv4 and IPv6 traffic	Need layer 2 transport layer over MPLS
IPv6 Provider Edge router (6PE) over MPLS	Internet and mobile service providers wanting to offer an IPv6 service	Low-cost and low-risk upgrade to the PE routers, and no impact on MPLS core	Applicable to MPLS infrastructure only	Software upgrade for PE routers
IPv6 VPN Provider Edge router (6VPE) over MPLS	Internet and mobile service providers wanting to offer IPv6 VPN services	Low-cost and low-risk upgrade to the PE routers and no impact on MPLS core	Applicable to MPLS infrastructure although the implementation could be done for other tunneling techniques. IPv6 address leakage on the global routing table must be well controlled	VPN or VRF support

Table 3. *A comparison of various IPv6 over MPLS backbone transition mechanisms.*

ed to the customers. Communications between the remote IPv6 domains run native IPv6 protocols over a dedicated link, where the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled using any transport over MPLS (AToM) or Ethernet over MPLS (EoMPLS) [7], with the IPv6 routers connected through an ATM or Ethernet interface, respectively.

Figure 3 shows an example of IPv6 deployment over any circuit transport over MPLS.

IPv6 on the Provider Edge Routers

Another deployment strategy is to configure IPv6 on the MPLS PE routers [8]. This strategy has a major advantage for service providers in that there is no need to upgrade either the hardware or software of the P routers in the MPLS core network, and it thus eliminates the impact on the operation of or revenue generated from existing IPv4 traffic. The strategy maintains the benefits of the current IPv4 MPLS features (e.g., MPLS-TE or VPNs), while appearing to provide a native IPv6 service for enterprise customers (using ISP-supplied IPv6 prefixes). The 6PE architecture allows support for IPv6 VPNs. Figure 4 shows an example of IPv6 deployment on the PE routers.

The IPv6 forwarding is done by label switching, eliminating the need for either IPv6 over IPv4 tunnels or additional layer 2 encapsulation, allowing the appearance of a native IPv6 service to be offered across the network and scaling as IPv6 service users grow since techniques such as route reflectors can be configured later.

Each PE router that must support IPv6 con-

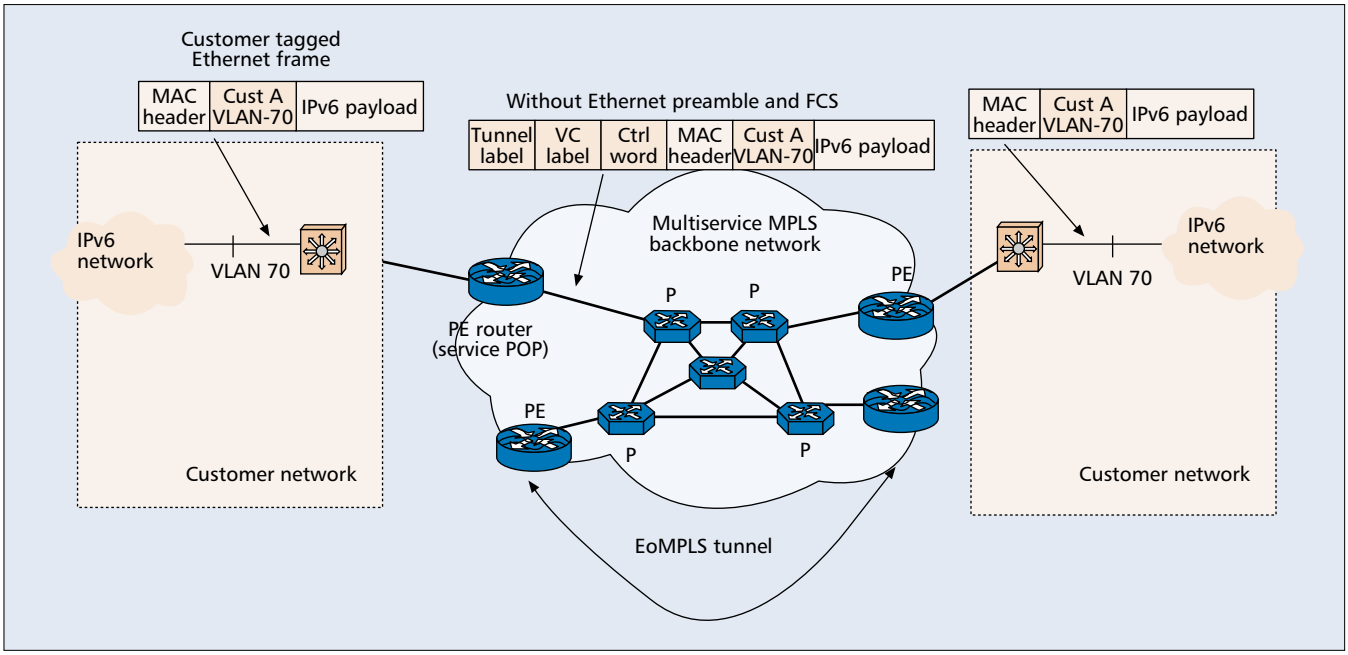

Figure 3. *IPv6 over "Ethernet over MPLS."*

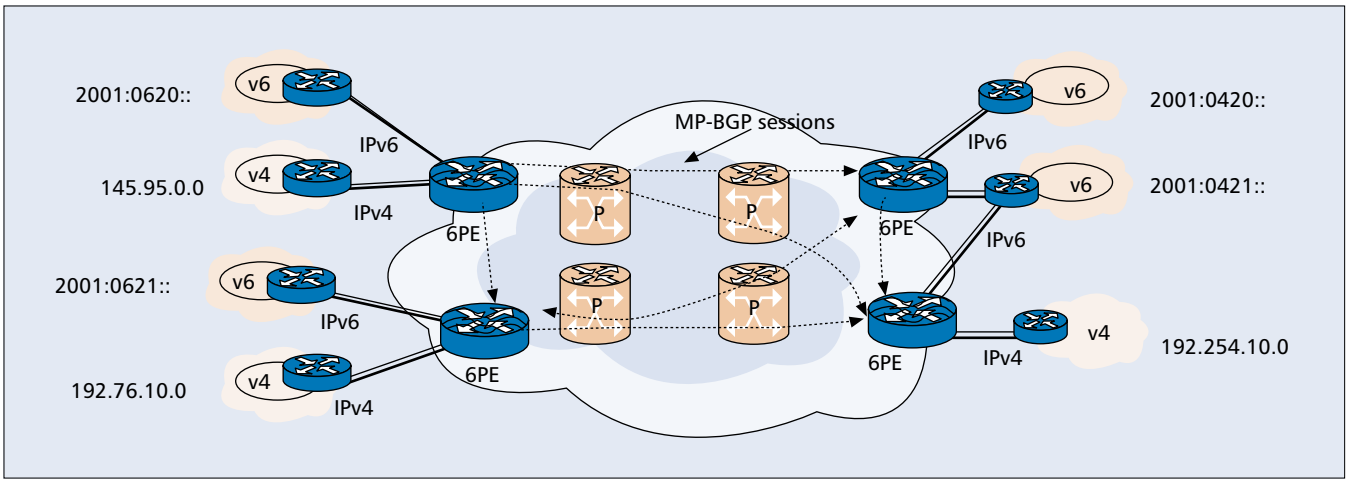

Figure 4. *IPv6 on provider edge routers.*

nectivity needs to be upgraded to be dual-stack (becoming a 6PE router) and configured to run MPLS on the interfaces connected to the core P routers. Depending on the site requirements, each router can be configured to forward IPv6 or IPv6 and IPv4 traffic on the interfaces to the CE routers, thus providing the ability to offer only native IPv6 or both IPv6 and IPv4 services. The 6PE router exchanges either IPv4 or IPv6 routing information through any of the supported routing protocols, depending on the connection, and switches IPv4 and IPv6 traffic over the native IPv4 and IPv6 interfaces not running MPLS.

The 6PE router exchanges reachability information with the other 6PE routers in the MPLS domain using multiprotocol Border Gateway Protocol (BGP), and shares a common IPv4 routing protocol, such as Open Shortest Path First (OSPF) or integrated Intermediate System to Intermediate System (IS-IS), with the other P and PE devices in the domain. The 6PE routers encapsulate IPv6 traffic using two levels of MPLS labels. The top label is distributed by a label distribution protocol (LDP) or tag distribution protocol (TDP) used by the devices in the core to carry the packet to the destination 6PE using IPv4 routing information. The second or bottom label is associated with the IPv6 address prefix of the destination through multiprotocol BGP-4, enabling load balancing to be performed.

IPv6 VPNs Provider Edge Routers over MPLS Backbone

Service providers who offer MPLS/VPN services to their customers may look forward to adding IPv6 VPN services to their portfolio. A VPN is said to be an IPv6 VPN [9] when a CE router turns on native IPv6 over an interface or subinterface to the PE router. Adding IPv6 VPN capability to a 6PE router, named 6VPE for IPv6 VPN Provider Edge Router over MPLS, is an option that enables an ISP to deliver similar services to IPv4. Similar to IPv4 VPN routes distribution, BGP and its extensions are used to distribute routes from an IPv6 VPN site to all other 6VPE routers connected to a site of the

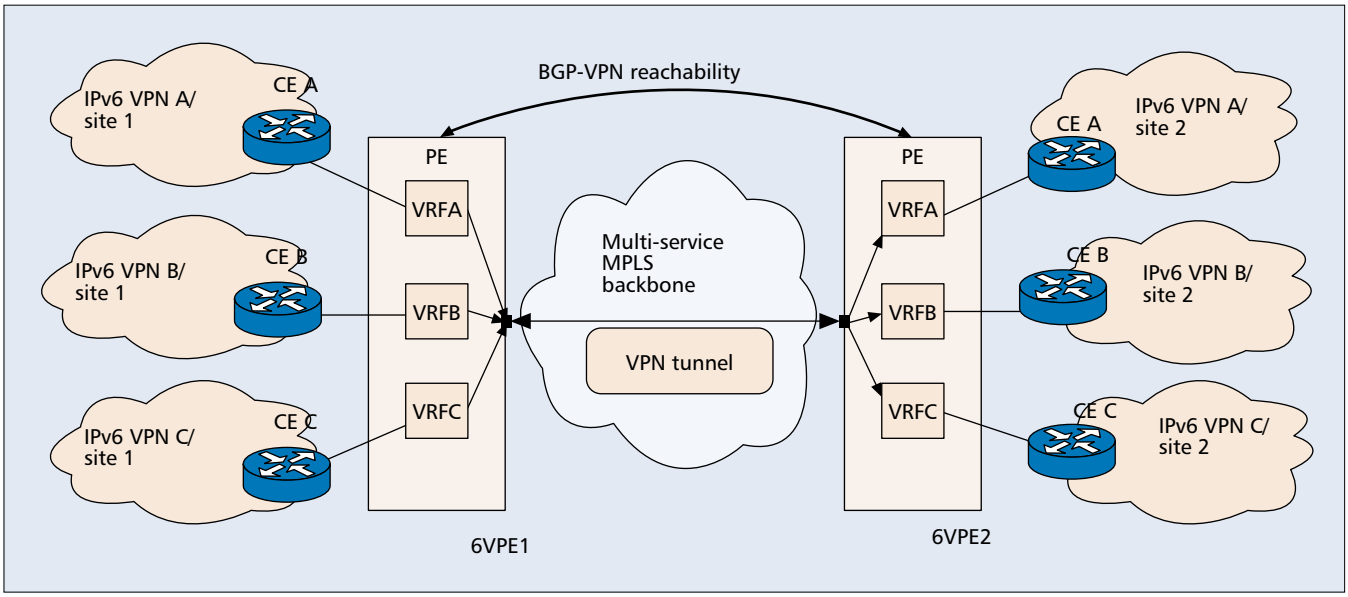

Figure 5. *IPv6 MPLS VPN architecture.*

Deployment strategy	Key user and primary use	Benefits	Limitations	Requirements
IPv6 over IPv4 tunnels	Service provider wanting to offer initial IPv6 service. Enterprise wanting to interconnect IPv6 domains or link to remote IPv6 network	Can demonstrate demand for minimal investment Easy to implement over existing IPv4 infrastructure Low cost and low risk	Complex management and diagnostics due to the independence of the tunnel and link topology	Access to IPv4 through dual stack router with IPv4 and IPv6 addresses. Access to IPv6 DNS
IPv6 over dedicated data links	Service provider WANs or MANs deploying ATM, Frame Relay or DWDM	Can provide end to end IPv6 with no impact on IPv4 traffic and revenue		Access to the WAN through dual stack router with IPv4 and IPv6 addresses. Access to IPv6 DNS
IPv6 over MPLS backbones	Mobile or greenfield service providers, or current regional service providers deploying MPLS	Integrates IPv6 over MPLS, thus no hardware or software upgrades required to the core	Implementation required to run MPLS; high management overhead	Minimum changes to the customer edge (CE) or provider edge (PE) routers, depending on the technique
Dual-stack backbones	Small enterprise networks Service providers' infrastructure Enterprise WAN infrastructure Campus infrastructure	Easy to implement for small campus network with a mixture of IPv4 and IPv6 applications Able to provide similar services (multicast, QoS) for both IPv4 and IPv6	Complex management of routing protocols. Major upgrade for large networks	Networking devices must be dual-stack-capable IPv6 entries on DNS Network design must apply to both IP versions with enough memory for routing tables

Table 4. *A comparison of all deployment or transition mechanisms.*

same IPv6 VPN. PEs use VPN routing and forwarding tables (VRFs) to separately maintain the reachability information and forwarding information of each IPv6 VPN, as shown in Fig. 5.

When a 6VPE1 router receives an IPv6 packet from CE A, it looks up the packet's IPv6 destination address in the VRF A. This enables it to find a VPN-IPv6 route. The VPN-IPv6 route will have an associated MPLS label and an associated BGP next hop. This MPLS label is imposed on the IPv6 packet. 6VPE1 directly pushes another label, top label binded by LDP/IGPv4 to the IPv4 address of BGP next hop to reach 6VPE2 through MPLS cloud, on the label stack of the labeled IPv6 VPN packet. This topmost imposed label corresponds to the label switched path (LSP) starting on 6VPE1 and ending on 6VPE2. As mentioned above, the bottom label is bound to the IPv6 VPN prefix via BGP.

All the P routers in the backbone network switch the VPN packet based only on the top label in the stack, which points toward the 6VPE2 router. Because of the normal MPLS forwarding rules, the P routers never look beyond the first label and are thus completely unaware of the second label or the IPv6 VPN packet carried across the backbone network.

The egress PE router, 6VPE2, receives the labeled IPv6 VPN packet, drops the first label, and performs a lookup on the second label, which uniquely identifies the target VRF A and sometimes even the outgoing interface on the 6VPE2. A lookup is performed in the target VRF A, and the IPv6 packet is sent toward the proper CE router in IPv6 domain or site.

IPv6 Network Design Considerations

For IPv6 deployment, when network designers favor an integration strategy for IPv6 that begins from the edges of the network and move in toward the core, this allows control over deployment cost and focus on the needs of the applications, rather than complete full upgrade to a native IPv6 network at this stage. The various deployment strategies permit the first stages of the transition to IPv6 to happen now, whether as a trial of IPv6 capabilities or the early controlled stages of major IPv6 network implementations. Table 4 compares various deployment strategies in terms of key users/primary use, benefits, limitations and requirements for each strategy.

Deploying IPv6 in a Service Provider Network Environment

As a network administrator for a service provider, one may want to evaluate and assess IPv6 now because current allocated IP address space may not be able to satisfy the potential huge increase in the number of users or the demand for new technologies from end customers can open new business opportunities for the service provider. Using globally unique IPv6 addresses raise the opportunity to create new business models, add revenues and increases the portfolio of services. Specified for the Internet

It is expected that IPv4 and IPv6 hosts will need to coexist for a substantial time during the steady migration from IPv4 to IPv6, and the development of transition strategies, tools, and mechanisms has been part of the basic IPv6 design from the start.

future of our next generations, IPv6 can be used for reachability and end-to-end security for networked devices, functionality crucial to emerging environments such as Internet- enabled PDAs and HANs, Internet-connected automobiles, integrated telephony services, and distributed gaming.

One should look at the deployment of IPv6 with the following three key phases, focusing on a business model that will help the management to see the added value of the project. Here, it is highly recommended that the IPv6 service be an IPv4/IPv6 dual stack service, when the ISP operators have enough experience in IPv6 dual stack operation.

Providing an IPv6 service (including IPv4 and IPv6 dual stack service) at the customer access level: Starting the deployment of IPv6 at the customer access level permits an IPv6 service to be offered now without a major upgrade to the core infrastructure and without an impact on current IPv4/MPLS services. This approach allows an evaluation of IPv6 products and services before full implementation in the network, and an assessment of the future demand for IPv6 without substantial investment at this early stage.

Running IPv6 (including IPv4 and IPv6 dual stack service) within the core infrastructure itself: At the end of this initial evaluation and assessment stage, and as network management systems fully embrace IPv6, the network infrastructure can be upgraded to support IPv6.

Interconnecting with other IPv6 service providers: Interconnections with other IPv6 service providers or the 6BONE allow further assessment and evaluation of IPv6, and a better understanding of the requirements for IPv6.

Conclusions

It is expected that IPv4 and IPv6 hosts will need to coexist for a substantial time during the steady migration from IPv4 to IPv6, and the development of transition strategies, tools, and mechanisms has been part of the basic IPv6 design from the start. In this article tunneling, translation, and dual-stack mechanisms are briefly revisited and compared, as they play a key role in this integration and coexistence of IPv4 and IPv6. At the same time, given the vast deployment of layer 2 infrastructure, this article proposes techniques for transporting IPv6 over dedicated links that service providers can consider utilizing their existing layer 2 infrastructure. Also, for service providers that have deployed MPLS infrastructure in the backbone, we have proposed and compared several techniques for transporting IPv6 over an MPLS backbone along with network design examples. A comparison of various IPv6 deployment strategies based on these transition mechanisms are being examined along with network design consideration for service provider environments.

References

[1] D. G. Waddington and F. Chang, "Realizing the Transition to IPv6," *IEEE Commun. Mag.*, June 2002.

[2] R. Gilligan *et al.*, "Transition Mechanisms for IPv6 Hosts and Routers," RFC 2893.

[3] A. Conta *et al.*, "Generic Packet Tunneling in IPv6," RFC 2473.

[4] B. Carpenter *et al.*, "Connection of IPv6 Domains via IPv4 Clouds," RFC 3056.

[5] F. Templin *et al.*, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)," IETF draft, draft-ietf-ngtrans-isatap-14.txt, work in progress, Aug. 2003.

[6] C. Huitema, "Teredo: Tunneling IPv6 over UDP through NATs," IETF draft, draft-huitema-v6ops-teredo-00.txt, work in progress, June 2003.

[7] L. Martini *et al.*, "Transport of Layer 2 Frames over MPLS," IETF draft, draft-martini-l2circuit-trans-mpls-11.txt, work in progress, Apr. 2003.

[8] J. De Clercq *et al.*, "Connecting IPv6 Islands across IPv4 Clouds with BGP," IETF draft, draft-ooms-v6ops-bgp-tunnel-00.txt, work in progress, Oct. 2002.

[9] J. De Clercq *et al.*, "BGP-MPLS VPN extension for IPv6 VPN," IETF draft, draft-ietf-l3vpn-bgp-ipv6-01.txt, work in progress, Aug. 2003.

Additional Reading

[1] A. Durand *et al.*, "IPv6 Tunnel Broker," RFC 3053.

Biographies

MALLIK TATIPAMULA [SM] (mtatipam@cisco.com) received a B.Tech. in electronics and communications engineering from Regional Engineering College, Warangal, India, and an M.Tech in communication systems and high frequency technologies from Indian Institute of Technology, Chennai. He is currently a senior product manager for advanced technologies in the Routing Technologies group at Cisco Systems. His expertise includes VoIP, mobile wireless, optical, IPv6, and GMPLS technologies. He closely works with service providers and national research networks around the world on deploying these advanced technologies in their next-generation networks. He has been with Cisco since 1998. Previously he worked at Motorola as principal engineer, responsible for design of next-generation wireless networks. From 1993 to 1997 he was with Nortel Networks, Ottawa, Canada, as senior member of scientific staff, working on Nortel's optical and wireless products. He has over 12 years of experience in telecom and networking. He has served on technical program committees of several leading IEEE and SPIE international conferences. He has delivered invited talks, and offered tutorials and short courses at leading conferences.

PATRICK GROSSETETE is a senior product manager in the Internet Technologies Division of Cisco Systems, responsible for Cisco IOS IPv6 and switching infrastructure (CEF) components. He is member of the IPv6 Forum Technical Directorate and manages Cisco participation in the IPv6 Forum as a founding member. In 1994 he joined Cisco as consulting engineer, and was then promoted to field distinguished engineer where he was closely involved with Cisco customers on network designs and new technologies introduction, particularly focusing on switching technologies. Before joining Cisco Systems, he worked for Digital Equipment Corporation in France from 1980 to 1994 as a network consultant and post-sales support. He was involved in the deployment of DECnet/OSI, learning important lessons for the ongoing IPv6 deployment period. He received computer degrees from Control Data Research Institute in 1979.

HIROSHI ESAKI received B.E. and M.E. degrees from Kyushu University, Fukuoka, Japan, in 1985 and 1987, respectively. He received a Ph.D. from the University of Tokyo in 1998. In 1987 he joined the Research and Development Center, Toshiba Corporation, where he engaged in research on ATM systems. Since 1998 he has worked at the University of Tokyo as an associate professor, and for the WIDE project as a board member. He was at Bellcore in New Jersey as a residential researcher from 1990 to 1991, and was engaged in research on high-speed computer communications. From 1994 to 1996 he was at the Center for Telecommunications Research of Columbia University, New York, as a visiting scholar. He is currently interested in high-speed internet architecture, including MPLS technology, mobile computing, and IPv6.