



IPv6 Deployment Strategies

Version History

Version Number	Date	Notes
1	10/15/2001	This document was created.
2	11/13/2001	Update to the explanation of NAT along tunnel paths.
3	03/08/2002	Update to the “Related Documents” section.
4	12/23/2002	Change to the global aggregatable address format.

This solutions document provides information to help you plan to deploy Internet Protocol Version 6 (IPv6) in your network. The document introduces and compares the strategies available for the deployment of IPv6 and describes some of the tasks you need to complete before your deployment. The “Prerequisites” section lists sources for information on IPv6, and other IPv6 documentation and training available from Cisco. The “Related Documents” section lists additional solutions documents relevant to IPv6 deployment.

The document includes the following sections:

- [IPv6 Deployment Strategies Overview, page 2](#)
- [Prerequisites, page 3](#)
- [Planning to Deploy IPv6, page 4](#)
- [Identifying Requirements, page 5](#)
- [Selecting a Deployment Strategy, page 6](#)
- [Predeployment Tasks, page 23](#)
- [Related Documents, page 27](#)

IPv6 Deployment Strategies Overview

The continuous growth of the global Internet requires that its overall architecture evolve to accommodate the new technologies that support the growing numbers of users, applications, appliances, and services. IPv6 is designed to meet these requirements and allow a return to a global environment where the addressing rules of the network are again transparent to the applications.

The current IP address space is unable to satisfy the potential huge increase in the number of users or the geographical needs of the Internet expansion, let alone the requirements of emerging applications such as Internet-enabled personal digital assistants (PDAs), home area networks (HANs), Internet-connected automobiles, 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 that is crucial to the applications and services that are 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 the new applications. The need for always-on environments (such as 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 the 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 and quality of service (QoS) that are not always readily available throughout a NAT-based network.

Standards bodies for the wireless data services are preparing for the future, and IPv6 provides the end-to-end addressing required by these new environments for mobile phones and residential Voice over IP (VoIP) gateways. IPv6 provides the services, such as integrated autoconfiguration, QoS, security, and direct-path mobile IP, also required by these environments.

IPv6 provides the following benefits:

- Larger address space for global reachability and scalability
- Simplified header for routing efficiency and performance
- Deeper hierarchy and policies for network architecture flexibility
- Efficient support for routing and route aggregation
- Serverless autoconfiguration, easier renumbering, multihoming, and improved plug and play support
- Security with mandatory IP Security (IPSec) support for all IPv6 devices
- Improved support for Mobile IP and mobile computing devices (direct-path)
- Enhanced multicast support with increased addresses and efficient mechanisms

We are in the early stages in the deployment of IPv6, with few IPv6 applications in the market and the first router products needing to make trade-offs between the available IPv6 services. The initial focus of these products is on the migration and transition techniques required for the deployment, rather than on meeting the requirements for high levels of traffic.

Although the success of IPv6 will depend ultimately on the availability of applications that run over IPv6, a key part of the IPv6 design is its ability to integrate into and coexist with existing IPv4 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.

Cisco has been part of this activity, participating in the development of transition techniques and deployment strategies for its products that satisfy a range of customer and network requirements, whether you are a service provider or enterprise customer, and whether you are planning a trial deployment or deploying live in a controlled environment. Your selection of a deployment strategy, or strategies, will depend on your current network environment and on factors such as the forecast amount of IPv6 traffic and the availability of IPv6 applications on your end systems, and at your stage in the deployment.

This solutions document provides information to prepare for your transition from IPv4 to IPv6, from initial training and planning activities, through the selection of an appropriate strategy, to the tasks you need to complete before deployment. These tasks allow an ordered approach to your transition, from trial deployments to evaluate the products, to deployment in a controlled environment to test the network and application connectivity, and finally to full deployment across your network.

Prerequisites

Before beginning to plan to deploy IPv6, you should familiarize yourself with IPv6.

Much of the definition of IPv6 is under the control of the Internet Engineering Task Force (IETF). The IETF is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. The IETF has a working group for IPv6, and is very much involved with the definition of IPv6 through RFCs and Internet Drafts. Refer to the following sites for more information on the IETF and IPv6:

<http://www.ietf.org/html.charters/ipngwg-charter.html>

<http://playground.sun.com/ipv6/>

The IPv6 Forum was created in 1999 to promote and advocate the IPv6 protocols and their deployment. This forum now has over 100 members, with IPv6 Forum summit meetings held periodically around the world. The IPv6 Forum site provides information on available IPv6 resources and presentations, and on current deployments and implementations. It also lists the founding and current members. Refer to the following site for information on the IPv6 Forum:

<http://www.ipv6forum.com>

Cisco is a founding member of the IPv6 Forum, and has been involved with IPv6 since the creation of the IETF IPng Working Group. Refer to the following site on Cisco.com for more information on the Cisco involvement with IPv6:

<http://www.cisco.com/warp/public/732/ipv6/index.shtml>

This comprehensive site provides links to a variety of sources that let you do the following:

- Learn about Cisco IOS IPv6
- View Cisco IOS IPv6 technical documents
- View Cisco IOS IPv6 presentations
- View Cisco IOS IPv6 press kit
- Read Cisco IOS IPv6 articles
- Read about IPv6 early adopters
- Learn how to get IPv6 address space

Cisco has delivered the first versions of IPv6 on its router platforms. Refer to the Related Documents section for information on the IPv6 product documentation.

Cisco also is developing a comprehensive training program. The first of these IPv6 training courses, *Implementing IPv6 Networks*, is available now. This instructor-led training course covers the installation and configuration of IPv6 networks, and the integration of IPv6 and its coexistence with IPv4 networks. The course covers the following topics:

- IPv6 features (including IPv6 address types and formats, ICMPv6, neighbor discovery, security, and mobility)
- IPv6 routing protocol support
- IPv6 integration and coexistence strategies
- IPv6 host configuration (Solaris, Microsoft, and FreeBSD)
- Connecting to the IPv6 Internet

Planning to Deploy IPv6

Cisco favors a transition strategy from IPv4 to IPv6 that begins from the edges of the network and moves in toward the core. This strategy allows you to control the deployment cost and focus on the needs of the applications, rather than complete a full network upgrade to a native IPv6 network at this stage. Cisco IPv6 router products offer the features for a such an integration strategy. The various deployment strategies permit the first stages of the transition to IPv6 to happen now, whether as a trial of IPv6 capabilities or as the early controlled stages of major IPv6 network implementations.

Service Provider

As a network administrator for a service provider, you may want to evaluate and assess IPv6 now because your current 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 your customers. Using globally unique IPv6 addresses simplifies the mechanisms used for reachability and end-to-end security for networked devices, functionality that is crucial to the emerging applications such as Internet-enabled PDAs, HANs, Internet-connected automobiles, integrated telephony services, and distributed gaming.

You should look at the deployment of IPv6 in three key phases:

- Providing an IPv6 service at the customer access level
- Running IPv6 within the core infrastructure itself
- Interconnecting with other IPv6 service providers

Starting the deployment of IPv6 at the customer access level permits an IPv6 service to be offered now without a major upgrade to your core infrastructure and without an impact on current IPv4 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.

At the end of this initial evaluation and assessment stage, as support for IPv6 within the routers improves (particularly IPv6 high-speed forwarding), and as network management systems fully embrace IPv6, the network infrastructure can be upgraded to support IPv6. This upgrade path could involve use of dual-stack routers (a technique for running both IPv4 and IPv6 protocols in the same router), or eventually use of IPv6-only routers as the IPv6 traffic becomes predominant.

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

**Note**

The 6bone is a worldwide IPv6 test network, informally operated with oversight from the NGtrans (IPv6 Transition) Working Group of the IETF. Its current focus is testing of the transition and operational procedures required for the deployment of IPv6. Becoming a member of this 6bone community is one way of gaining valuable experience with IPv6.

Enterprise

As a network manager or operator for an enterprise, you may want to evaluate and assess IPv6 now because of your plans to introduce IPv6 applications within the network in the near future. Although it is not expected that a great number of IPv6-only applications will ship initially, some of the mobile IP offerings being introduced in the market perform and scale better using the direct-path features that will become available in an IPv6 infrastructure, rather than those available with IPv4.

You may also want to assess and evaluate IPv6 because of the end-to-end addressing, integrated autoconfiguration, QoS, and security required by the new environments for mobile phones, or you may want to expand your available address space for some new service such as an IP-based telephone system.

You may want to return to a global environment where the addressing rules of the network are more transparent to the applications, and reintroduce end-to-end security and QoS that are not readily available throughout IPv4 networks that use NAT and other techniques for address conversion, pooling, and temporary allocation.

Two key ways of evaluating and assessing IPv6 products and services are as follows:

- Set up an IPv6 domain and connect to an existing remote IPv6 network such as the 6bone
- Set up two or more IPv6 domains and interconnect these over your existing IPv4 infrastructures

The current IPv6 transition techniques supported in Cisco IOS software allow the assessment and test of the IPv6 products and applications in the environments described in an independent and isolated way such that there is no disruption to current business.

Identifying Requirements

You should identify your requirements for IPv6 because they will determine your selection of a deployment strategy.

Service Provider

To provide an IPv6 service at the customer level, as a network administrator for a service provider you should begin by deciding which areas and customers are most likely to want IPv6 services, and then identify the access routers that can be upgraded to be dual-stack (a technique for running both IPv4 and IPv6 protocols in the same router) so as to provide both an IPv4 and IPv6 service to these customer sites. Alternatively you could specify and install separate dual-stack access routers to provide solely an IPv6 service, thus minimizing the impact on your existing IPv4 services even further. Other activities consist of setting up a Domain Name Server (DNS) that supports both the existing IPv4 A records and the new IPv6 AAAA records, and, if there is a need for intercommunication between IPv6-only and IPv4-only hosts, operating one of the protocol translation mechanisms such as NAT-PT in the router or a TCP-UDP Relay.

Initially, these access routers should be interconnected over the existing IPv4 core routers or infrastructure using one of the available deployment strategies to carry IPv6 over IPv4: carrying IPv6 packets inside IPv4 packets (tunneling), running IPv6 over a dedicated Layer 2 technology (such as ATM), or forwarding IPv6 packets over Multiprotocol Label Switching (MPLS) backbones. Your choice of deployment strategy will determine your choice of an IPv4 or IPv6 routing protocol.

For high-level service providers, register for your own IPv6 address prefix using the relevant International Regional Internet Registry (RIR) Process. For intermediate and mid-level service providers, contact your high-level service provider. Alternatively, if you want to connect only to the IPv6 6bone for testing before formal registration, apply for a prefix from this 6bone community.

See the section “[Selecting a Deployment Strategy](#)” for a more detailed description of these deployment strategies, and for hints in helping to choose the correct strategy for your environment. See the section “[Predeployment Tasks](#)” for more information on IPv6 routing protocols, IPv6 addresses, and DNS requirements.

Enterprise

As a network manager or operator for an enterprise, you should begin by choosing the IPv6 applications and services you would like to offer through IPv6, and decide where you want to provide these services. Activities then consist of creating an IPv6 domain and configuring a DNS that supports both IPv4 and IPv6 records, and, if there is a need for intercommunication between IPv6-only and IPv4-only hosts, operating one of the protocol translation mechanisms such as NAT-PT in the router or a TCP-UDP Relay.

You should then identify the router or routers in the network that need to be dualstack. They will be part of the IPv6 domain, using IPv6 routing protocols to communicate with the IPv6 applications, and either IPv4 or IPv6 protocols to communicate outside of the domain. The protocol choice will be dependent on whether you are connecting directly to an IPv6 service provider, or using one of the available deployment strategies to carry the IPv6 traffic over the existing IPv4 infrastructure to a remote IPv6 network or domain. In both cases, apply for IPv6 addresses from the relevant service provider.

See the section “[Selecting a Deployment Strategy](#)” for a more detailed description of these deployment strategies, and for hints in helping to choose the correct strategy for your environment. See the section “[Predeployment Tasks](#)” for more information on IPv6 routing protocols, IPv6 addresses, and DNS requirements.

Selecting a Deployment Strategy

The key strategies used in deploying IPv6 at the edge of a network involve carrying IPv6 traffic over the IPv4 network, allowing isolated IPv6 domains to communicate with each other before the full transition to a native IPv6 backbone. It is also possible to run IPv4 and IPv6 throughout the network, from all edges through the core, or to translate between IPv4 and IPv6 to allow hosts communicating in one protocol to communicate transparently with hosts running the other protocol. All techniques allow networks to be upgraded and IPv6 deployed incrementally with little to no disruption of IPv4 services.

The four key strategies for deploying IPv6 are as follows:

- Deploying IPv6 over IPv4 tunnels: These tunnels encapsulate the IPv6 traffic within the 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 IPv4-compatible and 6to4.

- Deploying IPv6 over dedicated data links: This technique enables isolated IPv6 domains to communicate by using the same Layer 2 infrastructure as for IPv4, but with IPv6 using separate Frame Relay or ATM PVCs, separate optical links, or dense Wave Division Multiplexing (dWDM).
- Deploying IPv6 over MPLS backbones: This technique allows isolated IPv6 domains to communicate with each other, but over an MPLS IPv4 backbone. 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 IPv6 using dual-stack backbones: This technique allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. All 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.

Table 1 summarizes the primary use, benefits, and limitations for each strategy.

Table 1 Deployment Strategies: Primary Uses, Benefits, and Limitations

Deployment Strategy	Key User/ 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 networks.	Can demonstrate demand for IPv6 for minimal investment. Easy to implement over existing IPv4 infrastructures. Low cost, low risk.	Complex management and diagnostics due to the independence of the tunnel and link topologies.	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 metropolitan area networks (MANs) deploying ATM, Frame Relay, or dWDM.	Can provide end-to-end IPv6 with no impact on the IPv4 traffic and revenue.	Lack of IPv6-specific hardware acceleration and support for IPv6 network management in currently deployed hardware.	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.
IPv6 Using Dual-Stack Backbones	Small enterprise networks.	Easy to implement for small campus networks with a mixture of IPv4 and IPv6 applications.	Complex dual management of routing protocols. Major upgrade for large networks.	All routers are dual-stack with IPv4 and IPv6 addresses. Access to IPv6 DNS. Enough memory for both IPv4 and IPv6 routing tables.

In addition to the strategies for deploying IPv6 within your IPv4 environment, you also need protocol translation mechanisms (for example, a NAT-PT device to connect IPv6-only web browsers to IPv4-only web servers) or dual-stack servers (for example, an e-mail server that handles IPv4-only and IPv6-only mail clients) to allow communication between applications using IPv4 and applications using IPv6.

These mechanisms become increasingly important 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 the IPv4 and IPv6 protocols on the end system, or on a dedicated server, or on a router 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.

The following sections provide further information on IPv6 deployment strategies and protocol translation mechanisms:

- [Deploying IPv6 over IPv4 Tunnels](#)
- [Deploying IPv6 over Dedicated Data Links](#)
- [Deploying IPv6 over MPLS Backbones](#)
- [Deploying IPv6 Using Dual-Stack Backbones](#)
- [Protocol Translation Mechanisms](#)


Refer to RFC 2893 for general information on the transition mechanisms for IPv6 hosts and routers, and refer to RFC 2185 for general information on the routing aspects of IPv6 transition.

Deploying IPv6 over IPv4 Tunnels

Tunneling is the encapsulation of IPv6 traffic within IPv4 packets so that they can be sent over an IPv4 backbone, allowing isolated IPv6 end systems and routers to communicate without the need to upgrade the IPv4 infrastructure that exists between them. Tunneling is one of the key deployment strategies for both service providers and enterprises during the period of IPv4 and IPv6 coexistence. [Figure 1](#) shows the use of IPv6 over IPv4 tunnels.

Tunneling allows service providers to offer an end-to-end IPv6 service without major upgrades to the infrastructure and without impacting current IPv4 services. Tunneling allows enterprises to interconnect isolated IPv6 domains over their existing IPv4 infrastructures, or to connect to remote IPv6 networks such as the 6bone.

Figure 1 Deploying IPv6 over IPv4 Tunnels



A variety of tunnel mechanisms are available. These mechanisms include manually created tunnels such as IPv6 manually configured tunnels (RFC 2893) and IPv6 over IPv4 GRE tunnels, semiautomatic tunnel mechanisms such as that employed by tunnel broker services, and fully automatic tunnel mechanisms such as IPv4-compatible and 6to4. Manual and GRE tunnels are used between two points and require configuration of both the source and destination ends of the tunnel, whereas automatic tunnel mechanisms need only to be enabled and are more transient — they are set up and taken down as required, and last only as long as the communication.

IPv6 for Cisco IOS software supports IPv6 manually configured, IPv6 over IPv4 GRE, IPv4-compatible, and 6to4 tunnel mechanisms. Tunnel broker services are provided by service providers.

Other tunnel techniques, such as ISATAP and 6over4, are available for use over campus networks or for the transition of local nonrouter sites.

The ISATAP tunneling mechanism is very similar to 6to4 tunneling, with the IPv4 address embedded in the lower 32 bits rather than the upper 48 bits of the IPv6 address. Cisco plans to support ISATAP tunnels in the next phase of IPv6 for Cisco IOS software.

The 6over4 mechanism maps IPv6 multicast addresses into IPv4 multicast addresses, determining the endpoint of the tunnel using neighbor discovery. The mechanism emulates a virtual link layer or Ethernet within the site, but note that IPv4 multicast routing is a prerequisite. Cisco does not plan to support 6over4 within Cisco IOS software, and we recommend use of ISATAP tunneling when available, or use of native IPv6 routing within the campus.

[Table 2](#) summarizes the primary use, benefits, and limitations for each tunneling mechanism.

Table 2 Overlay Tunnel Mechanisms: Primary Uses, Benefits, and Limitations

Tunnel Mechanism	Primary Use	Benefits	Limitations	Requirements
IPv6 Manually Configured Tunnel	Stable and secure links for regular communication. Connection to 6bone.	Supported in IPv6 for Cisco IOS software now. DNS with support for IPv6 not required.	Tunnel between two points only. Large management overhead. No independently managed NAT.	ISP-registered IPv6 address. Dual-stack router.
IPv6 over IPv4 GRE Tunnel	Stable and secure links for regular communication.	Well known standard tunnel technique. Supported in IPv6 for Cisco IOS software now.	Tunnel between two points only. Management overhead. No independently managed NAT. Cannot use to connect to 6bone.	ISP-registered IPv6 address. Dual-stack router. Required by i/IS-IS for IPv6.
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 Cisco IOS software.
Automatic IPv4-Compatible Tunnel	Single hosts or small sites. Infrequent communication.	Supported in IPv6 for Cisco IOS software now.	Communication only with other IPv4-compatible sites. Does not scale well. No independently managed NAT.	IPv6 prefix (0::/96). Dual-stack router.
Automatic 6to4 Tunnel	Connection of multiple remote IPv6 domains. Frequent communication.	Easy to set up with no management overhead. Supported in IPv6 for Cisco IOS software now.	No independently managed NAT.	IPv6 prefix (2002::/16). Dual-stack router.
ISATAP Tunnels	Campus sites. Transition of nonrouted sites.	To be supported in the next phase of Cisco IOS software.	Not yet commercially available.	Dual-stack router.
6over4 Tunnels	Campus sites. Transition of nonrouted sites.	—	Not supported by Cisco IOS software.	—

All tunneling mechanisms require that the endpoints of the tunnel run both IPv4 and IPv6, that is, must 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. The design is very similar in concept to running IP and either IPX, DECnet, or AppleTalk on the same router, something Cisco IOS software has done since its inception.

Dual-stack end systems allow applications to migrate one at a time from an IPv4 to an IPv6 transport. Applications that are not upgraded (they support only the IPv4 stack) can coexist with upgraded applications on the same end system. Applications choose between using IPv4 or IPv6 based on name lookup; both the 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 and particular requirements of the communication.

It may be possible to protect the IPv6 over IPv4 tunnels using IPv4 IPSec by applying a crypto map to both the tunnel interface to encrypt outgoing traffic, and to the physical interface to decrypt the traffic flowing through. Note that it may not be possible to use in all environments due to the limitations of IPSec in IPv4. However, if possible, protecting tunnels in this way may have a substantial impact on performance, and you should balance this loss of performance against the security that can be achieved by careful configuration of your network.

The following sections describe each of the supported tunneling mechanisms in more detail, and, where relevant, provide cross references to other IPv6 documentation:

- [IPv6 Manually Configured Tunnel](#)
- [IPv6 over IPv4 GRE Tunnel](#)
- [Tunnel Broker](#)
- [Automatic IPv4-Compatible Tunnel](#)
- [Automatic 6to4 Tunnel](#)

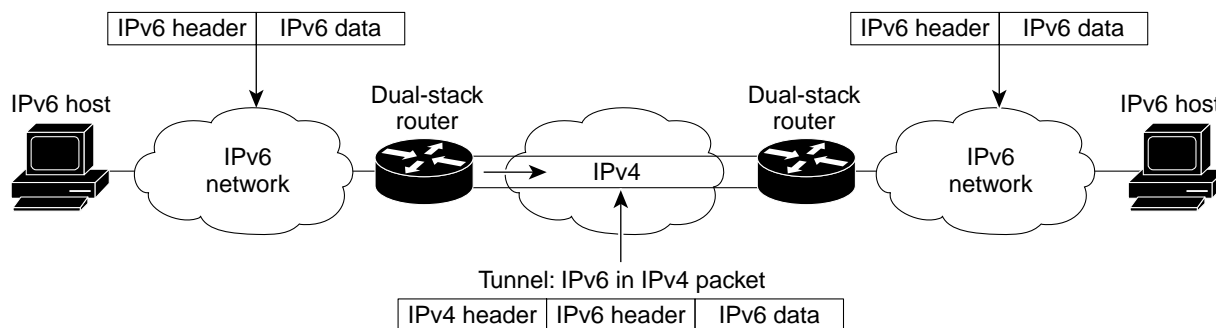
IPv6 Manually Configured Tunnel

A manually configured tunnel is equivalent to a permanent link between two IPv6 domains over an IPv4 backbone. The primary use is for stable connections that require regular secure communication between two edge routers or between an end system and an edge router, or for connection to remote IPv6 networks such as the 6bone. The edge routers and end systems, if they are at the end of the tunnel, must be dual-stack implementations.

At each end of the tunnel, you configure the IPv4 and IPv6 addresses of the dual-stack router on the tunnel interface, and identify the entry and exit (or source and destination) points using IPv4 addresses. For enterprises, your ISP provides you with the appropriate IPv6 address prefix for your site. Your ISP also provides you with the required destination IPv4 address for the exit point of the tunnel.

Figure 2 shows the configuration of a manually configured tunnel.

Figure 2 Manually Configured Tunnel



Because each tunnel exists between only two routers, adding routers means adding tunnels to cater for all the paths between the routers. Because each tunnel is independently managed, the more routers you have, the more tunnels you need, and the greater is the management overhead. As with other tunnel mechanisms, NAT, when applied to the outer IPv4 header, is allowed along the path of the tunnel only if the translation map is stable and preestablished.

Refer to RFC 2893 for further information on IPv6 manually configured tunnels. IPv6 for Cisco IOS software supports manually configured tunnels.

IPv6 over IPv4 GRE Tunnel


The IPv6 over IPv4 GRE tunnel uses the standard GRE tunneling technique that is designed to provide the services necessary to implement any standard point-to-point encapsulation scheme. As in IPv6 manually configured tunnels, GRE tunnels are links between two points, with a separate tunnel for each link. The tunnels are not tied to a specific passenger or transport protocol, but in this case carry IPv6 as the passenger protocol over GRE as the carrier protocol.

The primary use is for stable connections that require regular secure communication between two edge routers or between an edge router and an end system. The edge routers and, in the case described, the end systems must be dual-stack implementations.

Because i/IS-IS runs over a Layer 2 data link, tunneling techniques other than GRE cannot be used because i/IS-IS traffic cannot be distinguished from IPv6 traffic. GRE tunnels allow you to specify i/IS-IS as a passenger protocol, as you do for IPv6, and thus you can carry both i/IS-IS and IPv6 traffic at the same time over the same tunnel.

Figure 3 shows the configuration for an IPv6 over IPv4 GRE tunnel.

Figure 3 IPv6 over IPv4 GRE Tunnel



As with IPv6 manually configured tunnels, you configure the IPv4 and IPv6 addresses of the dual-stack router on the GRE tunnel interface, and identify the entry and exit (or source and destination) points of the tunnel using IPv4 addresses.

Also, as with manually configured tunnels, each GRE tunnel exists between only two routers, and thus adding routers means adding tunnels to cater for all the paths between the routers. Because each tunnel is independently managed, the more routers you have, the more tunnels you need, and the greater is the management overhead. As with other tunnel mechanisms, NAT, when applied to the outer IPv4 header, is allowed along the path of the tunnel only if the translation map is stable and preestablished.

IPv6 for Cisco IOS software supports IPv6 over IPv4 GRE tunnels. For further information, refer to the “Configuring Logical Interfaces” chapter of the *Cisco IOS Interface Configuration Guide*, Release 12.2.

Tunnel Broker

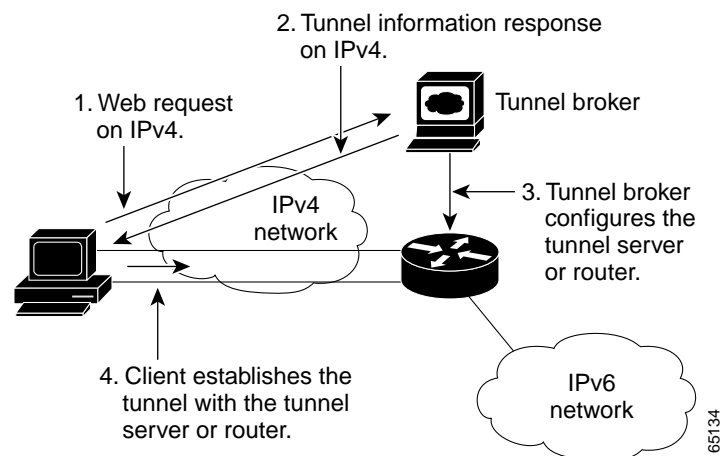
A tunnel broker service allows IPv6 applications on remote dual-stack end systems, or on IPv6 end systems connected to dual-stack routers, access to an IPv6 backbone. The tunnel broker service, using 6-over-4 tunnels to connect the end systems to the IPv6 backbone, automatically manages tunnel requests and configuration for the enterprise, rather than forcing the network administrator to manually configure tunnels.

65133

For instance, an enterprise could register the IPv4 address of the remote end system or router (using IPv4) with the service provider on a dedicated website. The service provider delivers a script that builds a tunnel to the IPv6 network, allocates an IPv6 address to the end system, and allocates a network prefix to the router to allow connectivity for the rest of the site. The tunnel broker manages the creation and deletion of the tunnel to the tunnel server, itself a dual-stack router that is connected to the IPv6 network.

Figure 4 shows the steps in the creation of a tunnel.

Figure 4 Tunnel Broker



The key limitation is that, by using this service, the end system or router is accepting a configuration change from a remote server, with the potential security implications of this activity.

Not all service providers offer a tunnel broker service, and not all available tunnel broker services support a script for routers from Cisco. Refer to the “other site” at the following URL for further information:

<http://www.ipv6.org>

Refer to RFCs 3051 and 3053 for further information on tunnel brokers.


Automatic IPv4-Compatible Tunnel

An automatic IPv4-compatible tunnel can be configured between edge routers or between an edge router and an end system. The edge routers and end systems must be dual-stack implementations.

An IPv4-compatible tunnel is one where the endpoints of the tunnel (the tunnel source and the tunnel destination) are automatically determined by the IPv4 address in the low-order 32 bits of the IPv4-compatible IPv6 address. This IPv4-compatible IPv6 address is a special IPv6 address with 0:0:0:0:0:0 in the high-order 96 bits and the IPv4 address in the low-order 32 bits.

Figure 5 shows the configuration of an IPv4-compatible tunnel.

Figure 5 IPv4-Compatible Tunnel



The IPv4-compatible tunnel is a transition mechanism that was defined early in the IPv6 development process, and its use in the future is under discussion in the IETF. Although it is an easy way to create tunnels for IPv6 over IPv4, it is a mechanism that does not scale well for large networks because each host requires an IPv4 address and an IPv6 address to be able to determine the endpoints of the tunnel. A further limitation is that all communication is always only between IPv4-compatible addresses. As with other tunnel mechanisms, NAT, when applied to the outer IPv4 header, is allowed along the path of the tunnel only if the translation map is stable and preestablished.

IPv6 for Cisco IOS software supports automatic IPv4-compatible tunnels.

Automatic 6to4 Tunnel

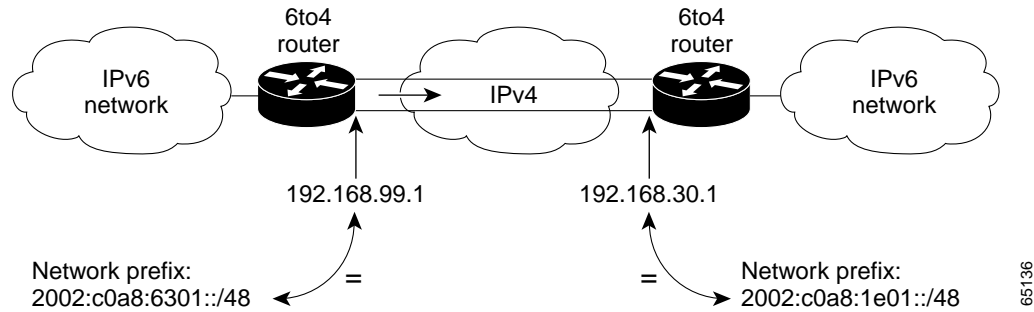
An automatic 6to4 tunnel allows isolated IPv6 domains to be connected over an IPv4 network and allows connections to remote IPv6 networks such as the 6bone. The key difference between this and manually configured tunnels is that the routers are not configured in pairs (and thus do not require manual configuration) because they treat the IPv4 infrastructure as a virtual nonbroadcast link, using an IPv4 address embedded in the IPv6 address to find the other end of the tunnel.

Each IPv6 domain requires a dual-stack router that identifies the IPv4 tunnel by a unique routing prefix in the IPv6 address (the IPv4 address of the tunnel destination is concatenated to the prefix 2002::/16). This unique routing prefix has been assigned permanently by the Internet Assigned Number Authority (IANA) for use in 6to4 schemes. Each site, even if it has just one public IPv4 address, has a unique routing prefix in IPv6. As with the manually configured and IPv4-compatible tunnel mechanisms, management of NAT needs to be linked with the management of the tunnel, and any independently managed NAT is not allowed along the path of the tunnel.

The simplest deployment scenario for 6to4 tunnels is to interconnect multiple IPv6 sites, each of which has at least one connection to a shared IPv4 network. This IPv4 network could be the global Internet or could be your corporate backbone. The key requirement is that each site has a 6to4 IPv6 address. As with other tunnel mechanisms, appropriate entries in a DNS that map between host names and IP addresses for both IPv4 and IPv6 allow the applications to choose the required address.

Figure 6 shows the configuration of a 6to4 tunnel for interconnecting 6to4 domains.

Figure 6 Interconnecting 6to4 Domains

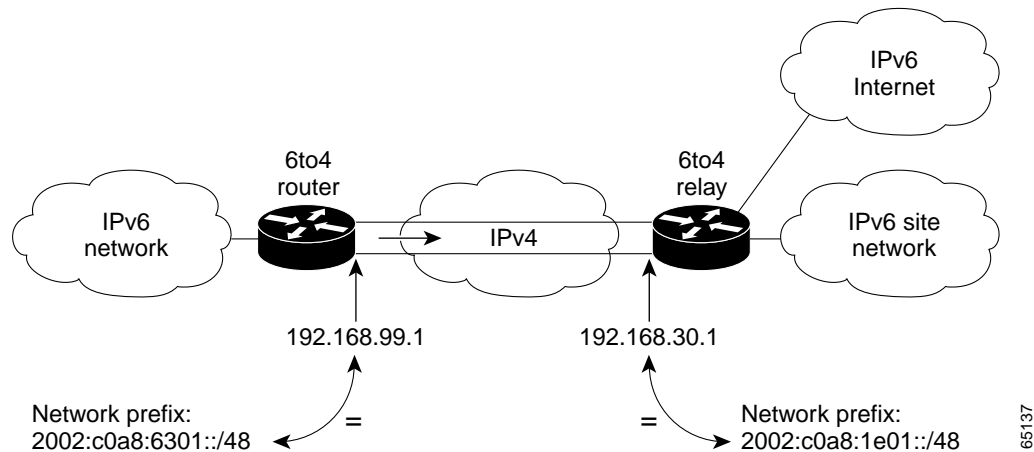


We recommend that each site have only one 6to4 address assigned to the external interface of the router. All sites need to run an IPv6 interior routing protocol such as RIPng for the IPv6 routing within the site, but exterior routing is handled by the relevant IPv4 exterior routing protocol.

As use of native IPv6 becomes more prevalent, the next stage is the use of 6to4 relay routers. These relay routers — standard routers but with both a 6to4 IPv6 address and a normal IPv6 address — provide a routing service between the native IPv6 domain, where a routing protocol is expected to be running, and the 6to4 domain, where there is no routing protocol. Communication between 6to4 sites and native IPv6 domains requires at least one relay router.

Figure 7 shows the configuration of a 6to4 tunnel for interconnecting 6to4 and native IPv6 domains.

Figure 7 Interconnecting 6to4 and Native IPv6 Domains Using Relay Routers



6to4 routers continue to run an IPv6 interior routing protocol for the IPv6 routing within the site, but participate in IPv6 interdomain routing by using a default IPv6 route that points to a specific relay router.

To avoid connectivity problems, you need to define routing and address filtering policies, based on network topology and traffic, that decide which IPv6 domains and which 6to4 prefixes the relay routers should advertise. Use of prefix lists should determine which parts of the IPv6 domain see the 2002::/16 prefix, and that any routing prefixes other than 2002::/16 are discarded.

Under discussion within the IETF is a method of simplifying the management of the default routes and configuring multiple relay routers using the 6to4 anycast address, thus allowing the 6to4 routers to discover the nearest available relay router automatically, and to increase the robustness of the 6to4 tunnel

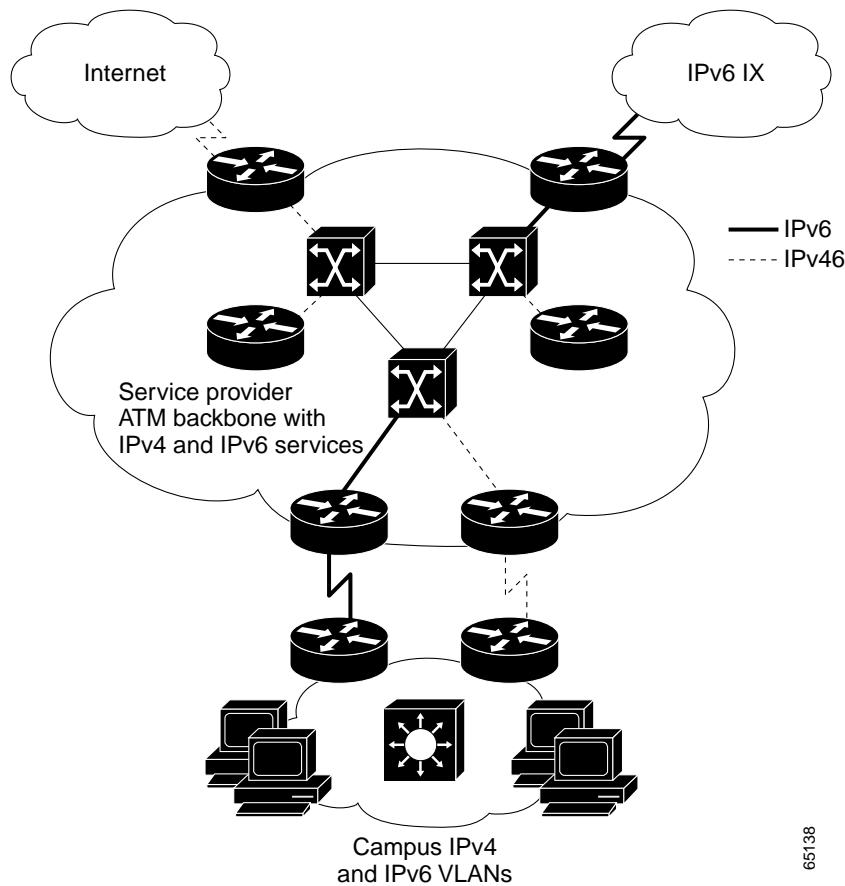
mechanism by allowing the 6to4 routers to switch automatically to another relay router in case of failure. IPv6 for Cisco IOS software is expected to support anycast addresses in this context as the method is finalized.

IPv6 for Cisco IOS software supports 6to4 tunnels.

Deploying IPv6 over Dedicated Data Links

Many WANs and MANs have been implemented by deploying Layer 2 technologies such as Frame Relay, ATM, or optical, and are beginning to use dWDM. Figure 8 shows a sample configuration for IPv6 over dedicated data links.

Figure 8 IPv6 over Dedicated Data Links



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. This configuration has the added benefit for the service provider of no loss in service or revenue for the IPv4 traffic.

IPv6 for Cisco IOS software supports IPv6 over dedicated data links.

Deploying IPv6 over MPLS Backbones

IPv6 over MPLS backbones enables isolated IPv6 domains to communicate with each other over an MPLS IPv4 core network. This implementation requires far fewer backbone infrastructure upgrades and lesser 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 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 using tunnels on the customer edge (CE) routers
- IPv6 over a circuit transport over MPLS
- IPv6 on the provider edge (PE) routers (known as 6PE)

The first of these strategies has no impact on and requires no changes to the MPLS provider (P) or PE routers because the strategy uses IPv4 tunnels to encapsulate the IPv6 traffic, thus appearing as IPv4 traffic within the network. The second of these strategies, only available on specific Cisco routers such as the Cisco 12000 and 7600 Internet routers, also requires no change to the core routing mechanisms. The last strategy requires changes to the PE routers to support a dual-stack implementation, but all the core functions remain IPv4.

[Table 3](#) summarizes the primary use, benefits, and limitations for each MPLS mechanism.

Table 3 *MPLS Mechanisms: Primary Uses, Benefits and Limitations*

MPLS 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.	Routers use IPv4-compatible or 6to4 addresses.	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.	Cisco 12000 or 7600 Internet routers in the core.
IPv6 on PE Routers	Internet and mobile service providers wanting to offer an IPv6 service.	Low cost and low risk upgrade to the PE routers. No impact on MPLS core.	No VPN or VRF support currently planned.	IPv6 software upgrade for PE routers.

IPv6 for Cisco IOS software currently supports the first two of these strategies, and Cisco has plans to support IPv6 provider edge routers in Phase II of its IPv6 for Cisco IOS software strategy.

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.

The following sections describe each mechanism in more detail, and provide cross references to the IPv6 documentation:

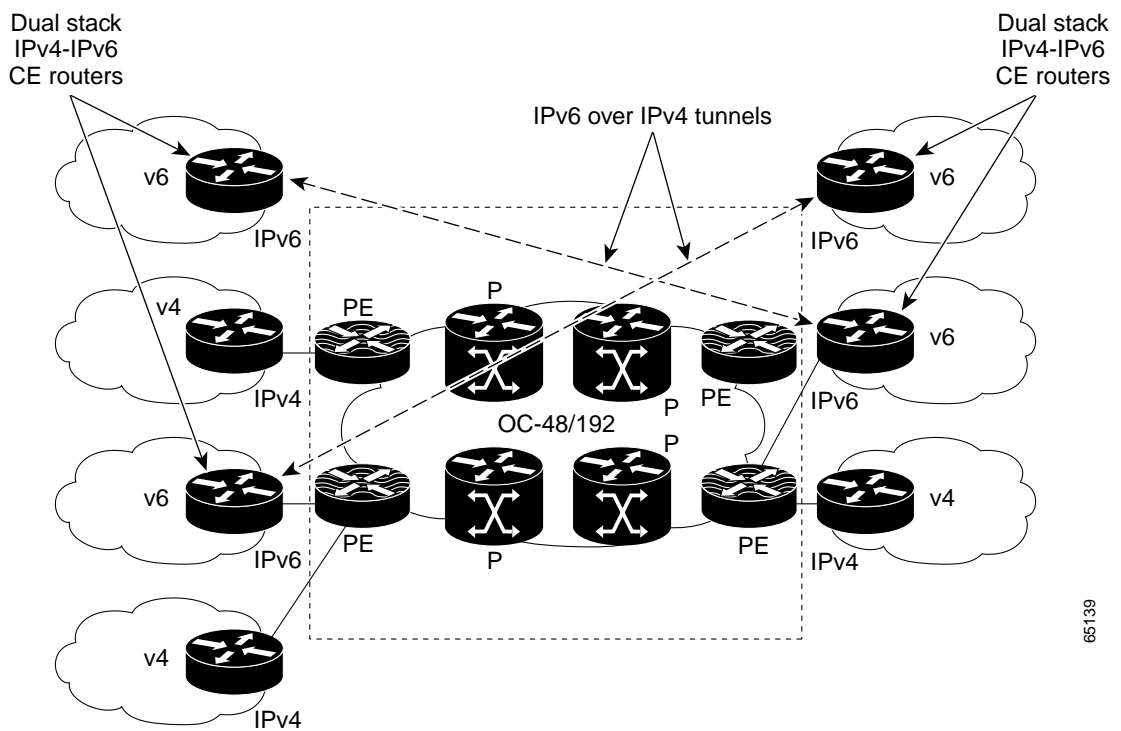
- [IPv6 Using Tunnels on the Customer Edge Routers](#)
- [IPv6 over a Circuit Transport over MPLS](#)
- [IPv6 on the Provider Edge Routers](#)

IPv6 Using Tunnels on the Customer Edge Routers

Using tunnels on the CE routers is the simplest way of deploying IPv6 over MPLS networks, having no impact on the operation or infrastructure of MPLS, and requiring no changes to either the P routers in the core or the PE routers connected to the customers.

Communication between the remote IPv6 domains uses standard tunneling mechanisms, running IPv6 over IPv4 tunnels in a similar way that MPLS VPNs support native IPv4 tunnels. The CE routers need to be upgraded to be dual-stack, and configured for IPv4-compatible or 6to4 tunnels, but communication with the PE routers is IPv4, and the traffic appears to the MPLS domain to be IPv4. The dual-stack routers use the IPv4-compatible or 6to4 address, rather than an IPv6 address supplied by the service provider. [Figure 9](#) shows the configuration using tunnels on the CE routers.

Figure 9 IPv6 Using Tunnels on the CE Routers




IPv6 over a Circuit Transport over MPLS

Using any 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 connected to the customers.

Communication between the remote IPv6 domains runs 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 (MPLS/AToM) or Ethernet over MPLS (EoMPLS), with the IPv6 routers connected through an ATM OC-3 or Ethernet interface, respectively. [Figure 10](#) shows the configuration for IPv6 over any circuit transport over MPLS.

Figure 10 IPv6 over a Circuit Transport over MPLS

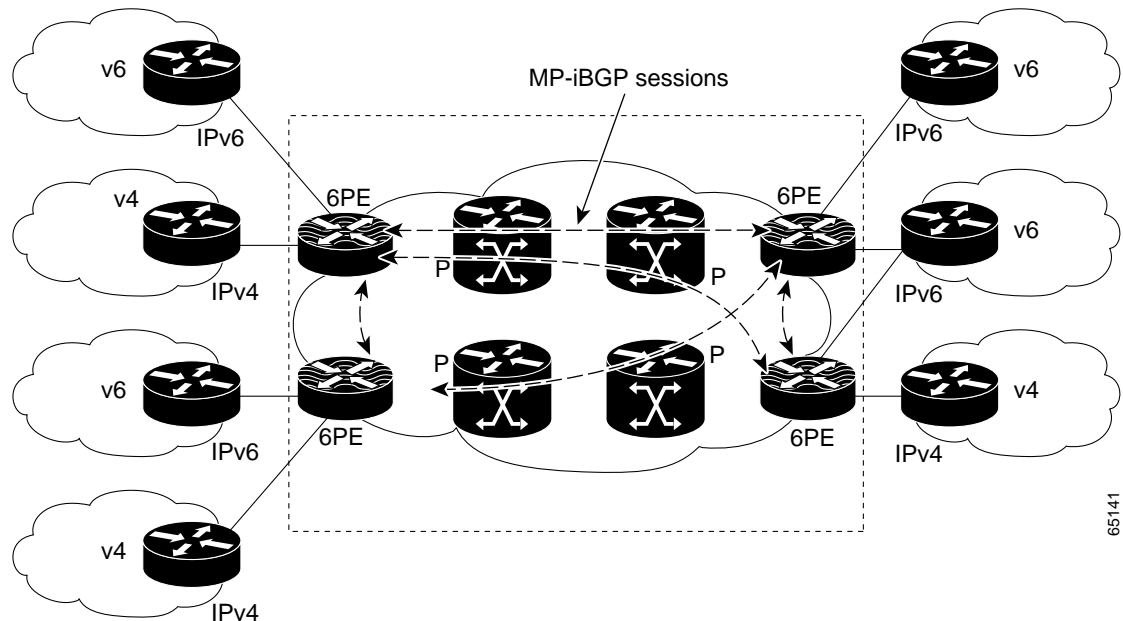


65140

IPv6 on the Provider Edge Routers

A further deployment strategy is to configure IPv6 on the MPLS PE routers. This strategy has a major advantage for service providers in that there is no need to upgrade either the hardware or software of the core network, and it thus eliminates the impact on the operation of or the revenue generated from the existing IPv4 traffic. The strategy maintains the benefits of the current MPLS features (for example, MPLS or VPN services for IPv4) while appearing to provide a native IPv6 service for enterprise customers (using ISP-supplied IPv6 prefixes). Figure 11 shows the configuration for IPv6 on the PE routers.

Figure 11 IPv6 on the Provider Edge Routers



65141

The IPv6 forwarding is done by label switching, eliminating the need for either IPv6 over IPv4 tunnels or for an additional Layer 2 encapsulation, allowing the appearance of a native IPv6 service to be offered across the network. The core network continues to run MPLS and any of the Cisco IOS software-supported IPv4 interior routing protocols, eliminating the requirement for upgrades to the hardware for native IPv6 forwarding and allowing the network to continue with current proven releases of Cisco IOS software.

Each PE router that must support IPv6 connectivity needs to be upgraded to be dual-stack (becoming a 6PE router) and configured to run MPLS on the interfaces connected to the core. 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 native 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 using the respective fast switching path (either Cisco Express Forwarding (CEF) or distributed CEF (dCEF) for IPv4 or CEF or dCEF for IPv6) 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 BGP, and shares a common IPv4 routing protocol (such as Open Shortest Path First (OSPF) or i/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 the label distribution protocol (LDP) 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 prefix of the destination through multiprotocol BGP-4.

The 6PE architecture allows support for IPv6 VPNs; Cisco IOS software may add support for VPN or VRF as the market requires.

Cisco plans to support 6PE routers in Phase II of its IPv6 for Cisco IOS software release strategy. Refer to the Internet-Draft *draft-ietf-ngtrans-bgp-tunnel-02.txt* for further information on 6PE routers.

Deploying IPv6 Using Dual-Stack Backbones

Using dual-stack backbones is a basic strategy for routing both IPv4 and IPv6. All routers in the network need to be upgraded to be dual-stack. IPv4 communication uses the IPv4 protocol stack (with forwarding of IPv4 packets based on routes learned through running IPv4-specific routing protocols), and IPv6 communication uses the IPv6 stack with routes learned through the IPv6-specific routing protocols.

The key requirements are that each site has an IPv6 unicast global prefix and appropriate entries in a DNS that map between host names and IP addresses for both IPv4 and IPv6. Applications choose between using IPv4 or IPv6 based on the response from the DNS resolver library, with the application selecting the correct address based on the type of IP traffic and particular requirements of the communication.

Today, dual-stack routing is a valid deployment strategy for specific network infrastructures with a mixture of IPv4 and IPv6 applications (such as 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, limitations to this approach are that 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.

Also, Cisco does not recommend an overall upgrade to a dual-stack network until there is a better parity between features and traffic levels. Although IPv6 for Cisco IOS software fully supports dual-stack, the current implementation of IPv6 requires enhancements to various services (for example, IPv6 multicast) before any network can be upgraded to dual-stack.

Protocol Translation Mechanisms

All of the integration strategies provide IPv6 end to end. Intercommunication between IPv4 and IPv6 requires some level of translation between the IPv4 and IPv6 protocols on the host or router, or dual-stack hosts, with an application-level understanding of which protocol to use.

A variety of protocol translation mechanisms are under consideration by the IETF NGTrans Working Group, as follows:

- Network Address Translation-Protocol Translation (NAT-PT)
- TCP-UDP Relay
- Bump-in-the-Stack (BIS)
- Dual Stack Transition Mechanism (DSTM)
- SOCKS-Based Gateway

These protocol translation mechanisms become more relevant as IPv6 becomes more prevalent, and even as IPv6 becomes the protocol of choice to allow legacy IPv4 systems to be part of the overall IPv6 network.

The 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 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 that requires extra protocol layers to be added to the IPv4 protocol stack.

[Table 4](#) provides a summary of the various translation mechanisms, with their primary use, benefits, and limitations.

Table 4 Protocol Translation Mechanisms: Primary Uses, Benefits, and Limitations

Translation Mechanism	Primary Use	Benefits	Limitations	Requirements
NAT-PT	IPv6-only hosts to IPv4-only hosts.	No dual stack. To be supported in IPv6 for Cisco IOS software Phase II.	No end-to-end IPsec. Dedicated server is single point of failure.	Dedicated server. DNS with support for IPv6.
TCP-UDP Relay	Translation between IPv6 and IPv4 on dedicated server.	Freeware. No changes to Cisco IOS software.	No end-to-end IPsec. Dedicated server is single point of failure.	Dedicated server. DNS with support for IPv6.
BIS	IPv4-only hosts communicating with IPv6-only hosts.	End-system implementation.	All stacks must be updated.	Updated IPv4 protocol stack.
DSTM	Dual-stack hosts (but with IPv6 address only).	Temporary IPv4 address allocated from pool.	No current support in Cisco IOS software.	Dedicated server to provide a temporary global IPv4 address.
SOCKS-Based IPv6/IPv4 Gateway	IPv6-only hosts to IPv4-only hosts.	Freeware. No changes to Cisco IOS software.	Additional software in the router.	Client and gateway software in the host and router.

The translation mechanisms that allow communication between IPv6-only and IPv4-only hosts, such as NAT-PT or BIS, use an algorithm called Stateless IP/ICMP Translator (SIIT). This mechanism translates, on a packet-by-packet basis, the headers in the IP packet between IPv4 and IPv6, and translates the

addresses in the headers between IPv4 and either IPv4-translated or IPv4-mapped IPv6 addresses. The mechanism assumes that each IPv6 host has a temporary IPv4 address assigned to it. SIIT is supported in IPv6 for Cisco IOS software as part of the NAT-PT implementation. Refer to RFC 2765 for further information on SIIT.

The following sections describe each protocol translation mechanism in more detail, and, where relevant, provide cross references to other IPv6 documentation:

- [NAT-PT](#)
- [TCP-UDP Relay](#)
- [BIS](#)
- [DSTM](#)
- [SOCKS-Based IPv6/IPv4 Gateway](#)

NAT-PT

The NAT-PT translation mechanism translates at the network layer between IPv4 and IPv6. An Application Level Gateway (ALG) translates between the IPv4 and IPv6 DNS requests and responses.

Its greatest use is where new hosts run only native IPv6 or the network has not implemented the dual-stack approach of IPv6 for Cisco IOS software. It has the same benefits as NAT for IPv4, and might be easier to introduce for IPv6 initially due to this familiarity and experience. However, NAT-PT also inherits the same limitations as NAT for IPv4, and makes fast rerouting difficult (ALGs are not as fast as IP routers). Also, the dedicated server is a single point of failure in the network. Although allowing security at an application level, NAT-PT inhibits end-to-end network security, and makes the merging of private-addressed networks extremely difficult.

Cisco plans to support NAT-PT in Phase II of its IPv6 for Cisco IOS software release strategy.

Refer to RFC 2766 for further information on NAT-PT.

TCP-UDP Relay

The TCP-UDP Relay mechanism is similar to NAT-PT in that it requires a dedicated server and DNS, but it translates at the transport layer rather than the network layer, with the DNS again providing the mapping between IPv4 and IPv6 addresses.

When the TCP relay server receives a request, it establishes separate connections at the transport level with both the source and destination IPv4 and IPv6 hosts, and then simply transfers data from one connection to the other. User Datagram Protocol (UDP) relays work in a similar manner.

The greatest use of this mechanism is for native IPv6 networks that want to access IPv4-only hosts, such as IPv4 web servers, but without the expense of upgrading either the IPv6 or IPv4 sides. The relay mechanism supports bidirectional traffic (multicast is not supported), but, as with NAT-PT, it allows application-level security but inhibits end-to-end network security, and makes the merging of private-addressed networks extremely difficult. Fast rerouting is difficult, and the dedicated server becomes a single point of failure in the network.

Implementations of the TCP-UDP relays are freely available from various locations. No changes are required to the Cisco IOS software.

BIS

The BIS mechanism is for communication between IPv4 applications on an IPv4-only host and IPv6-only hosts.

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. The translation mechanism is defined as part of SIIT.

This mechanism is for implementation on end systems only. An extension to the BIS mechanism allows dual-stack hosts to use the technique. Refer to RFC 2767 for further information.

DSTM

The DSTM translation mechanism is for dual-stack hosts in an IPv6 domain that have not yet had an IPv4 address assigned to the IPv4 side, but need to communicate with IPv4 systems or allow IPv4 applications to run on top of their IPv6 protocol stack. The mechanism requires a dedicated server that dynamically provides a temporary global IPv4 address for the duration of the communication (using DHCPv6), and uses dynamic tunnels to carry the IPv4 traffic within an IPv6 packet through the IPv6 domain.

DSTM becomes much more relevant as IPv6 becomes more prevalent and IPv4 addresses become scarce such that they need to be shared between hosts, and where the requirement is to carry IPv4 traffic over IPv6 or communicate between IPv6 hosts in an IPv6 domain and a few remote legacy IPv4 systems.

Support of DSTM within Cisco IOS software is under evaluation.

DSTM is at the Internet-Draft stage.

SOCKS-Based IPv6/IPv4 Gateway

The SOCKS-based IPv6/IPv4 gateway mechanism is for communication between IPv4-only and IPv6-only hosts. It consists of additional functionality in both the end system (client) and the dual-stack router (gateway) to permit a communications environment that relays two terminated IPv4 and IPv6 connections at the application layer.

This mechanism is based on the SOCKSv5 protocol, and inherits all the features of that protocol. Existing SOCKSv5 commands are unchanged, and the protocol maintains the end-to-end security between the client and the gateway, and the gateway and the destination.

The mechanism uses a feature called DNS Name Resolving Delegation to determine IPv6 addresses, delegating the name resolving to the gateway, thus requiring no change to existing DNSs.

Implementations of the SOCKS-based IPv6/IPv4 gateway are freely available from various locations. Refer to RFC 3089 for further information on the gateway and the locations of these sources.

Predeployment Tasks

Before deploying IPv6, you need to register for an IPv6 address (service provider) or request an IPv6 prefix from your service provider (enterprise), set up your DNS, decide on a policy for network management, and select required routing protocols.

IPv6 address registration is dependent on whether you are a service provider or enterprise, whether you are registering to be in a production environment or want to gain experience of IPv6 through the 6bone community, and the deployment strategy you selected. Your DNS must be configured with a server component to support IPv6, and with a resolver library that handles both IPv4 and IPv6 resource record types. Early network management of the dual-stack routers uses TFTP, ping, Telnet, and traceroute, with full support of IPv6 MIBs. Initial routing protocol support focuses on RIP, IS-IS, and multiprotocol BGP for IPv6, with support for OSPFv3 and Enhanced Interior Gateway Routing Protocol version 6 (EIGRPv6) planned for later releases of IPv6 for Cisco IOS software.

The following sections describe the predeployment tasks that you need to perform before starting your trial deployment:

- [IPv6 Address Requirements](#)
- [Domain Name Server Requirements](#)
- [Network Management](#)
- [Routing Protocols](#)

Refer to the *Cisco Statement of Direction for IPv6* at the following URL for a description of feature availability and time frames for release:

<http://www.cisco.com/warp/public/732/ipv6/index.shtml>

IPv6 Address Requirements

The process for IPv6 address allocation depends on whether you are a network administrator for a service provider or enterprise, whether you are registering to be in a production environment or want to gain experience of IPv6 through the 6bone community, and your choice of deployment strategy.

Refer to RFC 2471, *IPv6 Testing Address Allocation*, for information on the IPv6 testing address allocation, and RFC 2772 for information on the 6bone routing guidelines.

Service Providers

An aggregatable global address is an IPv6 address from the aggregatable global unicast prefix. The structure of aggregatable global unicast addresses enables strict aggregation of routing prefixes that limits the number of routing table entries in the global routing table. Aggregatable global addresses are used on links that are aggregated upward through organizations, and eventually to the Internet service providers (ISPs).

Aggregatable global IPv6 addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Except for addresses that start with binary 000, all global unicast addresses have a 64-bit interface ID. The current global unicast address allocation uses the range of addresses that start with binary value 001 (2000::/3).

Addresses with a prefix of 2000::/3 (001) through E000::/3 (111) are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format. The Internet Assigned Numbers Authority (IANA) allocates the IPv6 address space in the range of 2000::/16 to regional registries.

The aggregatable global address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID or Site-Level Aggregator (SLA). In the IPv6 aggregatable global unicast address format document (RFC 2374), the global routing prefix included two other hierarchically structured fields named Top-Level Aggregator (TLA) and Next-Level Aggregator (NLA). The IETF decided to remove the TLA and NLA fields from the RFCs because these fields are policy-based. Some existing IPv6 networks deployed before the change might still be using networks based on the older architecture.

A 16-bit subnet field called the Subnet ID could be used by individual organizations to create their own local addressing hierarchy and to identify subnets. A subnet ID is similar to a subnet in IPv4, except that an organization with an IPv6 subnet ID can support up to 65,535 individual subnets.

Registration as a production provider uses the RIR process. Service providers apply for a sub-TLA from the relevant regional registry: Asia-Pacific Network Information Centre (APNIC), American Registry for Internet Numbers (ARIN), or Réseaux IP Européens Network Coordination Centre (RIPE-NCC).

For APNIC registration, use the following URL:

<http://www.apnic.net/faq/IPv6-FAQ.html>

For ARIN registration, use the following URL:

<http://www.arin.net/regserv/ipv6/ipv6-regserv.html>

For RIPE-NCC registration, use the following URL:

<http://www.ripe.net/ripenc/mem-services/registration/ipv6/ipv6.html>

Service providers that want to be a member of the 6bone community can use the process specified at the following URL:

<http://www.6bone.net>

Enterprise

To run IPv6 in a production environment, request your IPv6 address from a suitable service provider. The service provider provides you with the 48-bit Public Topology external routing prefix of the IPv6 global unicast address, allowing you to create an appropriate address for your site as part of the router configuration process.



Note

Routers that are using automatic tunneling mechanisms have standard IPv6 prefixes to identify the tunnels (0::/96 for IPv4-compatible tunnels and 2002::/16 for 6to4 tunnels).

An alternative way to gain experience with IPv6 is to become a member of the 6bone community; that is, become an endsite of a 6bone service provider. You need to choose a provider that is reasonably local, because you must configure an IPv4 tunnel from your IPv6 router to the entry point of the 6bone network. Use the 6bone Registry to specify a short list of local providers. Determine the most suitable provider by checking the route to each using IPv4 pings and trace routes.

Contact the chosen provider, using information in the 6bone Registry entry, for an agreement to connect. Once you have an agreement to connect, create 6bone registry entries that provide information so that you can be connected, have addresses assigned, and be contacted in case of problems. The provider will supply the external routing prefix as in the production environment, and the router configuration process determines the rest of the address.

Domain Name Server Requirements

For dual-stack hosts, your selected DNS must provide resolver libraries that can handle IPv6 AAAA resource record types and IPv4 A record types, and must be capable of handling the cases where a query locates both IPv4 and IPv6 resource records. In this case, the DNS resolver library may return the IPv6

address, the IPv4 address, or both addresses to the application. The application then uses the IPv6 protocol or the IPv4 protocol, or makes a choice between the two based on the type of IP traffic and particular requirements of the communication.

IPv6 for Cisco IOS software currently supports DNS Client AAAA records over an IPv4 transport. DNS Client AAAA records over an IPv6 transport will be supported in Phase II of the IPv6 for Cisco IOS software strategy.

Your DNS should be running, or have equivalent capabilities of, Berkeley Internet Name Domain (BIND) version 9. This version provides an implementation of the major components of the DNS (DNS server, DNS resolver library, and verification tools) for IPv6.

Information on the latest release of BIND version 9 can be found at the following URL:

<http://www.isc.org/bind.html>

Network Management

The current dual-stack implementation in Cisco IOS software permits an interim network management solution, allowing applications such as TFTP, ping, Telnet, and traceroute to be run over either an IPv4 or IPv6 transport.

TFTP file downloading and uploading can be used to save the running configuration of the router to an IPv6 TFTP server. The **ping EXEC** command can accept a destination IPv6 address or IPv6 host name as an argument and send ICMPv6 echo request messages to the specified destination. The ICMPv6 echo reply messages are reported on the console. Extended ping functionality is also supported in IPv6. The Telnet client and server support IPv6 connections so that you can use Telnet to access the router or initiate Telnet connections from the router. The **traceroute EXEC** command accepts a destination IPv6 address or IPv6 host name as an argument and will generate IPv6 traffic to report each IPv6 hop used to reach the destination address.

IPv6 network management is supported by a series of Internet-Drafts for IP version-independent MIBs. The IPv6 for Cisco IOS Software product supports these MIBs, as well as Secure Shell (SSH) over IPv6 and Simple Network Management Protocol (SNMP) over IPv6. Further applications will be added as required.

Full management of IPv6 networks is dependent on IPv6 support within your particular network management system.

Routing Protocols

Selection of your routing protocols is dependent on the Cisco IOS software release deployed. This initial routing protocol support focuses on RIP for IPv6, i/IS-IS for IPv6, and multiprotocol extensions for BGP-4, with support for OSPFv3 and EIGRPv6 planned for later phases of IPv6 for Cisco IOS software.

IPv6 for Cisco IOS software supports RIP, based on RFC 2080, *RIPng for IPv6*, as the IGP. RIP in IPv6 functions in the same way and offers the same benefits as RIP in IPv4. IPv6 enhancements to RIP include support for IPv6 addresses and prefixes, and the use of the RIP router multicast group address FF02::9 as the destination address for RIP update messages.

IPv6 for Cisco IOS software supports i/IS-IS for IPv6 as currently defined in *draft-ietf-isis-ipv6-02*. i/IS-IS in IPv6 utilizes the same mechanisms described in RFC 1195 to add support for a new address family. This support is accomplished by adding two new type, length, and value (TLV) objects, and defining their use. Both IPv4 and IPv6 can be routed using a single intradomain routing protocol over the same topology.

IPv6 for Cisco IOS software also supports multiprotocol BGP, based on RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*, and RFC 2858, *Multiprotocol Extensions for BGP-4*, as the EGP. Multiprotocol BGP in IPv6 functions the same and offers the same benefits as multiprotocol BGP in IPv4. IPv6 enhancements to multiprotocol BGP include support for an IPv6 address family and Network Layer Reachability Information (NLRI) and next hop (the next router in the path to the destination) attributes that use IPv6 addresses and scoped addresses (the next hop attribute uses a global IPv6 addresses and potentially also a link-local address, when a peer is reachable on the local link).

Support for OSPFv3 and EIGRPv6 is planned for Phase III of IPv6 for Cisco IOS software. Although the fundamental mechanisms of the OSPFv2 implementation (for example: flooding, area support, SPF calculations, and so on) remain the same in OSPFv3, changes are necessary for OSPFv3 either because of changes in the protocol semantics between IPv4 and IPv6, or simply to handle the increased address size of IPv6. Refer to section 2 of RFC 2740 for a detailed explanation of the protocol modifications.

Related Documents

Refer to the following solutions documents for additional information on the deployment of IPv6:

- *IPv6: Connecting to the 6bone Using Manually Configured Tunnels*
- *IPv6: Connecting to the 6bone Using 6to4 Tunnels*
- *IPv6: Providing IPv6 Services over an IPv4 Backbone Using Tunnels*
- *Interconnecting IPv6 Domains Using Tunnels*

For releases 12.2(2)T through 12.2(11)T, and the 12.0 S and 12.0 ST releases, refer to the following documents for information on IPv6:

- *Start Here: Cisco IOS Software Release Specifics for IPv6 Features*
- *IPv6 for Cisco IOS Software, File 1 of 3: Overview*
- *IPv6 for Cisco IOS Software, File 2 of 3: Configuring*
- *IPv6 for Cisco IOS Software, File 3 of 3: Commands*

For release 12.2(13)T, refer to the following documents for information on IPv6:

- *Start Here: Cisco IOS Software Release Specifics for IPv6 Features*
- *Implementing IPv6 for Cisco IOS Software*
- *IPv6 for Cisco IOS Software Command Reference*



Note

The *Start Here: Cisco IOS Software Release Specifics for IPv6 Features* document details which IPv6 features are supported in each release of the S, ST, and T Cisco IOS software trains. All IPv6 features might not be supported in your Cisco IOS software release. We strongly recommend that you read the entire *Start Here: Cisco IOS Software Release Specifics for IPv6 Features* document before reading the other IPv6 for Cisco IOS Software feature documentation.

Refer to the Cisco Statement of Direction for IPv6 at the following URL for a description of feature availability and time frames for release:

<http://www.cisco.com/warp/public/732/ipv6/index.shtml>

