Understanding the Essentials Series

www.cisco.com/go/abc

Preface 5
Introduction. The Rationale for a New Version of IP 7
Network Address Translation 8
Limitations of NAT 8
Meeting Future Network Requirements 9
Evolution of Internet Protocol Version 6 9
Chapter 1. Features and Benefits of Using IPv6 11
Larger Address Space for Global Reachability and Scalability 11
Simplified Header for Efficient Packet Handling 11
Hierarchical Network Architecture for Routing Efficiency 13
Multihoming 13
Support for Routing Protocols 13
Routing Information Protocol 13
Open Shortest Path First Protocol Version 3 13
IS-IS Protocol 14
Multiprotocol Border Gateway Protocol+ 14
Autoconfiguration and “Plug-and-Play” Support 14
Easier Renumbering 14
Elimination of Need for Network Address Translation and Application’s Layered Gateway 14
Embedded Security with Mandatory IPSec Implementation 15
Enhanced Support for Mobile IP and Mobile Computing Devices 15
Increased Number of Multicast Addresses 15
Multicast Scope Address 15
Quality of Service 16
Chapter 2. IPv6 Header Format 17
IPv6 Header Fields 17
Description of IPv6 Header Fields 17
IPv6 Extension Headers 18
Order of Extension Headers 19
Routing Header 20
Fragment Header 20
ICMPv6 Packet 21
Chapter 3. IPv6 Addressing Architecture 23
IPv6 Address Format 23
IPv6 Address Prefix 24
IPv6 Address Types 24
IPv6 Address Assignment 24
IPv6 Unicast Address 25
What Is an IPv6 Global Unicast Address? 25
What Is an IPv6 Site-Local Unicast Address? 26
What Is an IPv6 Link-Local Unicast Address? 27
What Is an IPv4-Compatible IPv6 Address? 28
What Is an IPv4-Mapped IPv6 Address? 28
IPv6 Anycast Address 29

Chapter 3. IPv6 Addressing Architecture (cont.)
IPv6 Multicast Address 29
Multicast Group Membership Requirement for IPv6 Nodes 30
What Is an IPv6 Solicited-Node Multicast Address? 30
Special IPv6 Addresses 31
What Is an IPv6 Unspecified Address? 31
What Is an IPv6 Loopback Address? 31
IPv6 Address Allocation 31
6BONE Network Address Allocation 32
How Is an IPv6 Address Represented in a URL? 32
How Many IP Addresses Does an IPv6 Host Require? 32
How Many IP Addresses Does an IPv6 Router Require? 32
Chapter 4. Operation of IPv6 33
Neighbor Discovery 33
What Is IPv6 Neighbor Solicitation? 33
What Is IPv6 Neighbor Advertisement? 34
IPv6 Router Discovery 34
What Is IPv6 Router Advertisement? 34
What Is IPv6 Router Solicitation? 35
IPv6 Redirect Message 35
Stateless Autoconfiguration 36
Renumbering of IPv6 Nodes 36
How Does Duplicate Address Detection Work? 36
Path Maximum Transmission Unit Discovery 36
How Does IPv6 Path MTU Discovery Work? 37
Dynamic Host Configuration Protocol Version 6 37
IPv6 Domain Name System Operation 38
Using AAAA Records for DNS Resolution 38
Chapter 5. Integration and Coexistence Strategies 39
Transition Mechanisms 39
Using IPv4-IPv6 Protocol Dual Stack Devices 40
Deploying IPv6 Using Dual Stack Backbones 41
Deploying IPv6 over IPv4 Tunnels 42
Tunneling Requirements 42
Tunneling and Security 42
IPv6 Tunnel Mechanisms 43
IPv6 Manually Configured Tunnel 43
IPv6 over IPv4 GRE Tunnel 44
Automatic IPv4-Compatible Tunnel 45
Automatic 6to4 Tunnel 45
ISATAP Tunnel 47
Teredo Tunnel 48
Deploying IPv6 over Dedicated Data Links 48
Deploying IPv6 over MPLS Backbones 49
Deploying IPv6 Using Tunnels on the Customer Edge Routers 50
Deploying IPv6 over a Circuit Transport over MPLS 51
Deploying IPv6 on the Provider Edge Routers 51

Chapter 5. Integration and Coexistence Strategies (cont.)
Protocol Translation Mechanisms 52
Stateless IP/ICMP Translator 53
Network Address Translation-Protocol Translation 53
TCP-UDP Relay 54
Bump-in-the-Stack 54
Dual Stack Transition Mechanism 55
SOCKS-Based IPv6/IPv4 Gateway 55
Deployment of Translation Mechanisms 55
Chapter 6. IPv6 Network Design Considerations 57
Deploying IPv6 in a Service Provider Network Environment 57
Deploying IPv6 in an Enterprise Network Environment 57
IPv6 Support from Cisco 58
Appendix A 59
Bibliography and Reference Resources 59
Cisco Statement of Direction for IPv6 59
Cisco Technical Documentation 59
Books 59
White Papers and other Documentation 59
RFCs and Drafts 60
Rationale and Case for IPv6 60
Protocols 60
IPv6 Address Types 60
IPv6 Autoconfiguration and Renumbering 60
IPv6 Link Layer 60
IPv6 Routing Protocol Support 61
IPv6 Integration and Transition Mechanisms 61
IPv6 Deployment 61
Other Web References 61
IPv6 Host Configuration 62
IPv6 Address Allocation 62
IPv6 Address Registries 62
Current Sub-TLA Allocations 62
Appendix B 63
Glossary 63
Appendix C 67
Review Questions 67
Appendix D 72
Answers to the Review Questions 72

Preface

The ABCs of IP Version 6 is intended for network professionals with good IP version 4 (IPv4) networking skills and knowledge. This document is ideal for anyone, including account managers and system engineers, who is required to analyze IPv6 network requirements and develop strategies for the deployment of IPv6 networks.

We have kept the technical content in this document as generic as possible. Where appropriate we have provided more details on certain technologies or strategies, based on Cisco's product implementation of IPv6. We have purposely omitted topological and configuration discussion and examples from this document.

Although you are encouraged to read the chapters in this document sequentially, you could choose to read the chapters you are most interested in. You will find the review quiz in Appendix C at the end of the document useful in reinforcing your learning. In addition to the resources listed in the appendixes, you can also find information on IPv6 implementation details including the roadmap, software configuration, and Statement of Direction at www.cisco.com/ipv6.

To obtain more in-depth IPv6 training, please consult the Learning Locator on **www.cisco.com** or send e-mail to **abcios@cisco.com**.

Thanks to Steve Deering, Patrick Grossetete, Tony Hain, Ole Troan, Florent Parent, Kevin Flood, Neville Fleet, Simon Pollard, and Yatman Lai for their contribution and technical review.

Casimir Sammanasu

Cisco IOS Learning Services

Rationale for a New Version of IP

IP version 6 is a new IP protocol designed to replace IP version 4, the Internet protocol that is predominantly deployed and extensively used throughout the world.

The current version of IP has not been substantially changed since RFC 791, *Internet Protocol DARPA Internet Program Protocol Specification* was published in 1981. IPv4 has proven to be robust, easily implemented, and interoperable, and has stood the test of scaling an internetwork to a global utility the size of the Internet today.

However, the initial design did not anticipate the following conditions:

- Recent exponential growth of the Internet and the impending exhaustion of the IPv4 address space
- Growth of the Internet and the ability of Internet backbone routers to maintain large routing tables
- Need for simpler autoconfiguration and renumbering
- Requirement for security at the IP level
- Need for better support for real-time delivery of data—also called quality of service (QoS)

Note: Features such as IP Security (IPSec) and QoS have been specified for both versions of IP.

Though the 32-bit address space of IPv4 supports about 4 billion IP devices, the IPv4 addressing scheme is not optimal, as described by Christian Huitema in RFC 3194, *The Host-Density Ratio for Address Assignment Efficiency: An Update on the H Ratio.*

IPv4 will soon reach the stage where you will have to choose between new capabilities or a larger network, but not both. So, we need a new protocol to provide new and enhanced features in addition to solving the IP address exhaustion problem.

Network Address Translation

Emerging countries are facing the IPv4 address crunch more strongly than Europe or the United States. Although the use of NAT has delayed the IPv4 address exhaustion, the use of NAT introduces some complications that can be overcome only with a new IP protocol.

In IPv4 networks, NAT is typically used to connect internal networks by translating packets between an internal network, which uses the private address space, as described in RFC 1918 *Address Allocation for Private Internets,* and the Internet. NAT uses only a few global (external) addresses even in a large internal network.

Limitations of NAT

Note that the use of NAT only delays the time of exhaustion of the IPv4 addresses but does not solve the real large-scale growth problem, because IP is now widely adopted as the application's convergence layer for non-computing devices. Additionally, use of NAT has many implications, as identified in RFC 2775, *Internet Transparency,* and RFC 2993, *Architectural Implications of NAT*. Some of these problems follow and can be solved only with a new protocol, such as IPv6:

- With IPv4, only the endpoints handle the connection and the underlying layers do not handle any connection. However, when NAT is used, it breaks the end-to-end connection model of IP.
- Because NAT must handle the translation of addresses and ports, NAT requires the network to keep the states of the connections. In case of failure of the NAT device or the links near the NAT device, the need to keep the state of the connections in NAT makes fast rerouting difficult.
- NAT also inhibits the implementation of end-to-end network security. The integrity of the IP header is protected by some cryptographic functions. This header cannot be changed between the origin of the packet, which protects the integrity of the header and the final destination, where the integrity of the received packet is checked. Any translation of parts of the headers along the path will break the integrity check
- With applications that are not "NAT-friendly," more than just port and address mapping is necessary to forward the packet through the NAT device. NAT must embed complete information of all the applications to accomplish this goal, especially in the case of dynamically allocated ports with rendezvous ports, embedded IP addresses in application protocols, security associations, and so on. Every new deployment of a non-NAT-friendly application will require the upgrading of the NAT device.
- When different networks that are using the same private address space, such as 10.0.0.0/8, need to be combined or connected, as in the case of a merger, an address space collision will result. Though techniques such as renumbering or twice-NAT can resolve this collision, these techniques are very difficult and will increase the complications of NAT.
- The ratio of internal/reachable to external addresses mapping must be large to make NAT effective. However, when there are many servers inside, the same protocol cannot be multiplexed on the same port using the NAT external address. For example, two internal servers using the same port (80) cannot use the same external outside address without changing the port number. Each inside server that must be reachable from the outside will start using one external address. Because there are many protocols that make nodes as servers and consume many external addresses, NAT is not quite as useful if the number of inside servers is large.

Meeting Future Network Requirements

Though the exhaustion of IPv4 addresses is the primary reason for the development of a new protocol, the designers of IPv6 added other new features and some critical improvements to IPv4.

IPv6 is designed to meet the user, application, and service requirements, and allow a return to a simpler environment where the operation of the network is again transparent to the applications.

The anticipated rollout of wireless data services has been identified as a key IPv6 driver. The wireless industry standardization bodies, for example, the 3rd Generation Partnership Project (www.3gpp.org), Universal Mobile Telecommunication System (www.umts-forum.org), and Mobile Wireless Internet Forum (www.mwif.org) are considering IPv6 as the foundation for future IP services. Today, IPv6 services are available over IEEE 802.11 from some "hot-spot" locations.

The overall market adoption of IPv6 will be determined by the ability of the architecture to best accommodate Internet growth, new IP applications, and services. All these factors underscore the original rationale behind definition of IPv6 and the market drivers.

Evolution of Internet Protocol Version 6

IPv5 is an experimental resource reservation protocol intended to provide QoS, defined as the Internet Stream Protocol or ST. ST is not a replacement of IP, but uses an IP version number (number five), because it uses the same link-layer framing as IPv4. Resource reservation is now done using other protocols (for example, resource reservation protocol (RSVP)). IPv5/ST protocol is documented in RFC 1190, *Experimental Internet Stream Protocol, Version 2 (ST-II)* and RFC 1819, *Internet Stream Protocol Version 2 (ST2) Protocol Specification - Version ST2+*.

The original proposal for IPv6 proposed in RFC 1752, *The Recommendation for the IP Next Generation Protocol* was the *Simple Internet Protocol Plus (SIPP)* with a larger (128 bit) address space. The main author of SIPP was Steve Deering, now a Cisco Fellow. Following that proposal, the IETF started a working group and the first specification came in late 1995 with RFC 1883, *Internet Protocol, Version 6 (IPv6) Specification*. RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*, by Steve Deering (Cisco) and Rob Hinden (Nokia), obsoletes RFC 1883 and is the present standard for IPv6.

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 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.

The ABCs of IP Version 6 document discusses the following topics in detail:

1. Features and Benefits of using IPv6
2. IPv6 Header Format
3. IPv6 Addressing Architecture
4. Operation of IPv6
5. Integration and Coexistence Strategies
6. IPv6 Network Design Considerations

Chapter 1

Features and Benefits of Using IPv6

In addition to meeting the anticipated future demand for globally unique IP addresses, IPv6 provides the following benefits to network and IT professionals:

- Larger address space for global reachability and scalability
- Simplified header format for efficient packet handling
- Hierarchical network architecture for routing efficiency
- Support for widely deployed routing protocols
- Autoconfiguration and plug-and-play support
- Elimination of need for network address translation (NAT) and application's layered gateway (ALG)
- Embedded security with mandatory IPSec implementation
- Enhanced support for Mobile IP and Mobile Computing Devices
- Increased number of multicast addresses

Larger Address Space for Global Reachability and Scalability

The availability of an almost unlimited number of IP addresses is the most compelling benefit of implementing IPv6 networks. Compared to IPv4, IPv6 increases the number of address bits by a factor of 4, from 32 bits to 128 bits. The 128 bits provide approximately 3.4x1038 addressable nodes, enough to allocate about 1030 addresses per person on this planet. Figure 1 shows the general format of an IPv6 address.

Figure 1: IPv6 Address Format

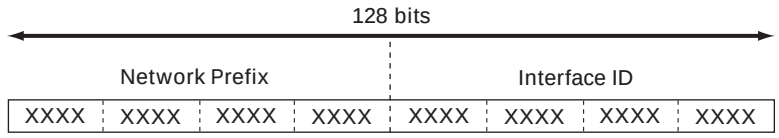

XXXX = 0000 through FFFF

$3.4x10^{38}$ = ~340,282,366,920,938,463,374,607,432,768,211,456 IPv6 Addresses

The ability to provide a unique address for each network device enables end-to-end reachability, which is especially important for residential IP telephony. IPv6 also provides full support for application protocols without requiring special processing at the edges of the networks, eliminating the problems associated with NAT.

Simplified Header for Efficient Packet Handling

Although the increase in the number of bits in the IPv6 address results in an increase in IPv6 header size, the IPv6 header format is simpler compared to the IPv4 header. The basic IPv4 header size is only 20 octets, but the

variable length of the Options field adds to the total size of the IPv4 packet. The IPv6 header has a fixed size of 40 octets. Although 6 of the 12 IPv4 header fields have been removed in IPv6, some IPv4 fields have been carried over with modified names, and some new fields have been added to improve efficiency and introduce new features. As shown in Figure 2, the Header Length (IHL), Identification, Flags, Fragment Offset, Header Checksum, and Padding fields have been removed from the IPv6 header.

This removal results in faster processing of the basic IPv6 header, but routing efficiency and overall performance are dependent on the option headers treatment and lookup algorithms a given device must run. Also, all fields in the IPv6 header are 64 bits, taking advantage of the current generation of 64-bit processors.

Figure 2: Comparison of IPv4 and IPv6 Headers

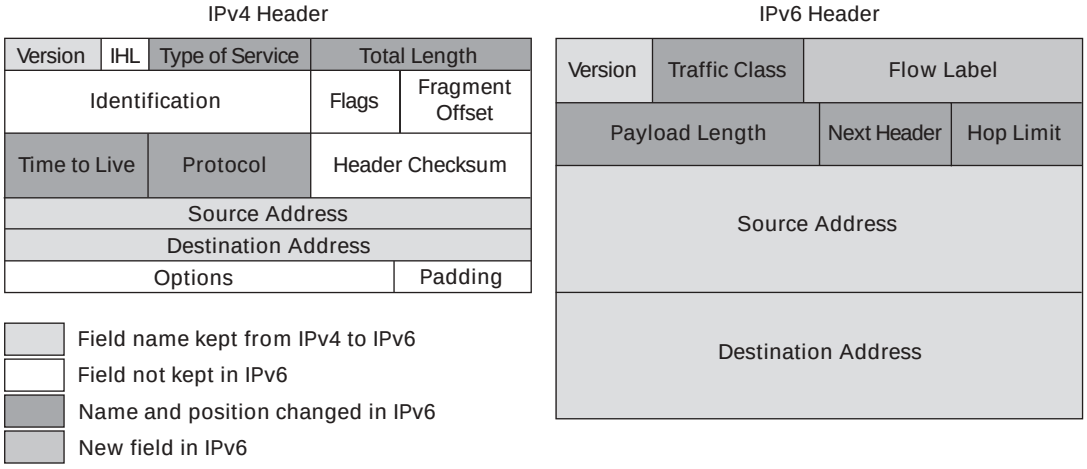

Fragmentation is now managed differently and does not need the fields in the basic IP header. Routers no longer do fragmentation in IPv6, which removes the processing issues caused by routers managing IPv4 fragmentation. Because checksum has been removed and the routers along the path of an IPv6 packets need not recalculate checksum every time, routing efficiency is improved in IPv6.

In IPv6 networks, fragmentation is handled by the source device with the help of path maximum transmission unit (MTU) discovery protocol.

The checksum has been removed at the IP layer because most link-layer technologies already do checksum and error control. And because the relative reliability of the link layer is very good, IP header checksum was considered unnecessary and not very useful. In addition to the error detection handled by the link layer technologies, the transport layer that handles end-to-end connection has a checksum that enables error detection.

However, this removal forces the upper-layer optional checksums, like User Datagram Protocol (UDP), to become mandatory in IPv6, whereas the UDP transport layer uses an optional checksum in IPv4.

The Options field of IPv4 is changed in IPv6 and is now managed by an extension header chain. The majority of the other fields were either not changed or changed only slightly. In addition to a smaller number of fields, the header is 64 bits aligned to enable fast processing by current processors.

Hierarchical Network Architecture for Routing Efficiency

The availability of a very large addressing space and network prefixes provides a flexible network architecture. This flexibility allows an organization to use only one prefix for the entire network of the organization.

A larger address space allows the allocation of large address blocks to Internet service providers (ISP) and to other organizations. This allocation in turn, allows the ISP to aggregate the prefixes of all its customers into a single prefix and announce this one prefix to the IPv6 Internet.

The larger IPv6 address space also enables the use of multiple levels of hierarchy inside the address space. Each level helps to aggregate the traffic at that level and enhance the allocation of addresses in a hierarchical format. The implementation of multiple levels in the address hierarchy permits flexibility and new functionalities, such as the scoping of addresses. The hierarchical network architecture of IPv6 allows the ISPs to use aggregation of network prefixes to provide efficient and scalable routing.

The hierarchical addressing structure is designed to reduce the size of Internet routing tables. Without a good hierarchical addressing scheme, routers will have to store large routing tables. Though classless interdomain routing (CIDR) in IPv4 solves this problem with the use of route aggregation, it is neither scalable nor efficient.

Multihoming

Though multihoming allows a network to be connected to two or more ISPs and is desirable for high reliability, it is difficult to connect a network to multiple providers in IPv4 because such connection breaks any kind of aggregation in the global routing table. The availability of a much larger address space in IPv6 enables the use of multiple simultaneous prefixes for a network, without breaking the global routing table.

However, redundancy and load sharing for multihomed networks, scalability of the global routing table, and a simple and operationally manageable multihoming guidelines still need to be defined. IPv6 multihoming capabilities and application impacts are under study in the IETF Multi6 working group.

Support for Routing Protocols

To enable scalable routing, IPv6 supports existing Interior Gateway Protocols (IGPs) and Exterior Gateway Protocols (EGPs). Similar to IPv4, IPv6 uses the longest prefix match for a routing algorithm.

Routing Information Protocol

The Routing Information Protocol Next-Generation (RIPng) protocol explained in RFC 2080, *RIPng for IPv6*, functions the same and offers the same benefits as RIP-2 (RFC 1721, *RIP Version 2 Protocol Analysis*) in IPv4. IPv6 enhancements to RIPng include support for IPv6 addresses and prefixes, including next hop IPv6. RIPng uses the all-RIP routers multicast group address FF02::9 as the destination address for RIP update messages. RIPng uses IPv6 for transport of the protocol messages.

Open Shortest Path First Protocol Version 3

Although most of the algorithms of OSPFv2 are the same in OSPFv3, some changes have been made in OSPFv3, particularly to handle the increased address size in IPv6 and the fact that OSPF runs directly over IP. Because OSPFv2 is heavily dependent on the IPv4 address for its operation, changes were necessary in OSPFv3 protocol to support IPv6, as outlined in RFC 2740, *OSPF for IPv6*. Some of the notable changes include platform-independent implementation, protocol processing per-link rather than per-node processing, explicit support for multiple instances per link, and changes in authentication and packet format.

IPv6 OSPF is now an IETF proposed standard. Like RIPng, IPv6 OSPFv3 uses IPv6 for transport and uses link-local addresses as source address.

IS-IS Protocol

The IS-IS routing protocol is an IGP protocol and IPv6 IS-IS is an IETF draft. New IPv6 routing capability has been added to the existing IS-IS protocol. Internet Draft draft-ietf-isis-ipv6-02.txt specifies a method for exchanging IPv6 routing information using the IS-IS routing protocol utilizing the same mechanisms described in RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*. This is accomplished by adding 2 new type-length-values (TLVs)—"IPv6 Reachability" (128 bits) and "IPv6 Interface Address" (128 bits)—and a new IPv6 protocol identifier.

Multiprotocol Border Gateway Protocol+

Multiprotocol BGP in IPv6 is an EGP that functions the same and offers the same benefits as multiprotocol BGP in IPv4. RFC 2858, *Multiprotocol Extensions for BGP-4* describes multiprotocol extensions for BGP4 defined as new attributes. RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Interdomain Routing* describes the enhancements to multiprotocol BGP that 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. These attributes use IPv6 addresses and scoped addresses. The next hop attribute uses a global IPv6 address and potentially also a link-local address, when a peer is reachable on the local link.

Autoconfiguration and "Plug-and-Play" Support

The address autoconfiguration feature is built into the IPv6 protocol to facilitate intranet-wide address management, enabling large number of IP hosts to easily discover the network and get new and globally unique IPv6 address associated with their location. The autoconfiguration feature enables "plug-and-play" Internet deployment of new consumer devices, such as cell phones, wireless devices, home appliances, and so on. As a result, network devices could connect to the network without manual configuration and without any servers, such as DHCP servers.

A router on the local link will send network-type information, such as the prefix of the local link and the default route in its router advertisements. The router provides this information to all the nodes on the local link. As a result, a host can autoconfigure itself by appending its 48-bit link layer address (MAC address) in an extended universal identifier EUI-64-bit format to the 64 bits of the local link prefix advertised by the router.

Easier Renumbering

In IPv6 networks, the autoconfiguration feature makes renumbering of an existing network simpler and relatively easier. The router sends the new prefix from the new upstream provider in its router announcements. The hosts in the network will automatically pick the new prefix from the router advertisements and then use it to create their new addresses. As a result, the transition from provider A to B becomes more manageable for network operators.

Elimination of Need for Network Address Translation and Application's Layered Gateway

With the availability of a large number of IPv6 addresses to provide globally unique IP addresses for all IP devices, there is no need for translating hundreds of internal IP addresses into a few global IP addresses. The

elimination of the need for deploying NAT boxes in networks will also eliminate other problems associated with the deployment of NAT. Particularly, the elimination of NAT provides end-to-end transparency in networks, reduces network complexity, and helps to reduce network operational costs for enterprises and ISPs.

Embedded Security with Mandatory IPSec Implementation

While the use of IPSec is optional in IPv4, IPSec is mandatory in IPv6 and is part of the IPv6 protocol suite. Therefore, network implementers could enable IPSec in every IPv6 node, potentially making the networks more secure.

IPv6 provides security extension headers, making it easier to implement encryption, authentication, and virtual private networks (VPNs). Because IPv6 offers globally unique addresses and security, IPv6 can provide end-to-end security services such as access control, confidentiality and data integrity without the need for additional firewalls that might introduce additional problems, including performance bottlenecks.

Enhanced Support for Mobile IP and Mobile Computing Devices

In IPv6, mobility is built in and any IPv6 node can use mobility as needed. Mobility is becoming an important and critical feature in networks. Mobile IP is an IETF standard allowing mobile devices to move around without breaking their existing connections. In IPv4, the mobility function must be added as a new feature. Mobility support in IPv6 is discussed in the latest version of Internet Draft draft-ietf-mobileip-ipv6-17.txt.

IPv6 packets addressed to the home address of a mobile node are transparently routed to its care-of address through the caching of the binding of its home address with its care-of address. This binding allows any packets destined for the mobile node to be directed to it at this care-of address. Mobile IPv6 defines four new IPv6 destination options: binding update option, binding acknowledgement option, binding request option, and home address option.

The routing headers in IPv6 make Mobile IPv6 much more efficient for end devices than Mobile IPv4. The use of the routing header for Mobile IP, rather than IP encapsulation, enables Mobile IP to avoid triangle routing, making it much more efficient in IPv6 than in IPv4.

Note: The authentication of the binding update between the mobile node and correspondent node is still under discussion at the IETF.

Increased Number of Multicast Addresses

One of the salient features of IPv6 is that it does not use broadcasts at all. The functions previously supported by IPv4 broadcasts such as router discovery and router solicitation requests are handled by IPv6 multicast. Multicast allows IP packets such as a video stream to be sent to multiple destinations at the same time, saving network bandwidth. Multicast improves the efficiency of a network by limiting the broadcast requests to a smaller number of only interested nodes. IPv6 uses specific multicast group addresses for its various functions. Thus, IPv6 multicast prevents the problems caused by broadcast storms in IPv4 networks.

Multicast Scope Address

IPv4 networks use administratively scoped IP multicast addresses as described in RFC 2365, *Administratively Scoped IP Multicast*, to allow packets to be addressed to a specific range of multicast addresses (for example, 239.0.0.0 to 239.255.255.255). By specifying a multicast scope, the packets are prevented from crossing the configured administrative boundaries. IPv4 uses one broadcast address for a particular scoped zone or IP multicast boundary, and the broadcasts are received by all hosts in this scoped zone.

IPv6 uses a 4-bit Scope ID to specify address ranges reserved for multicast addresses for each scope. Thus, only those hosts in a specified scope address range configured to listen to a specific multicast address receive the multicast. However, a host can be a member of several workgroups and can listen to several multicast addresses at the same time.

IPv6 provides a larger range of multicast addresses compared to IPv4. So, allocation of addresses for multicast groups will not be limited for the foreseeable future.

Quality of Service

QoS in IPv6 is handled in the same way it is currently handled in IPv4. Support for class of service is available through the Traffic Class field compliant with the IETF Differentiated Services (DiffServ) model.

However, IPv6 header has a new field named Flow label which can contain a label identifying a specific flow, such as video stream or videoconference. The source node generates this flow label. Having a flow label enables QoS devices in the path to take appropriate actions based on this label. But, the existence of the flow label itself is not a feature of QoS.

Chapter 2

IPv6 Header Format

The IPv6 header is simpler and more efficient than the IPv4 header. The simplified IPv6 header helps to reduce processing costs.

This chapter discusses the major differences between the IPv4 and IPv6 header formats and how the IPv6 header has been simplified. The following topics are covered in this chapter:

- IPv6 header fields
- IPv6 extension headers
- IPv6 fragmentation header
- IPv6 routing header
- IPv6 ICMP packet

IPv6 Header Fields

The basic IPv6 packet header consists of 8 fields as shown in Figure 3.

Figure 3: IPv6 Header Fields

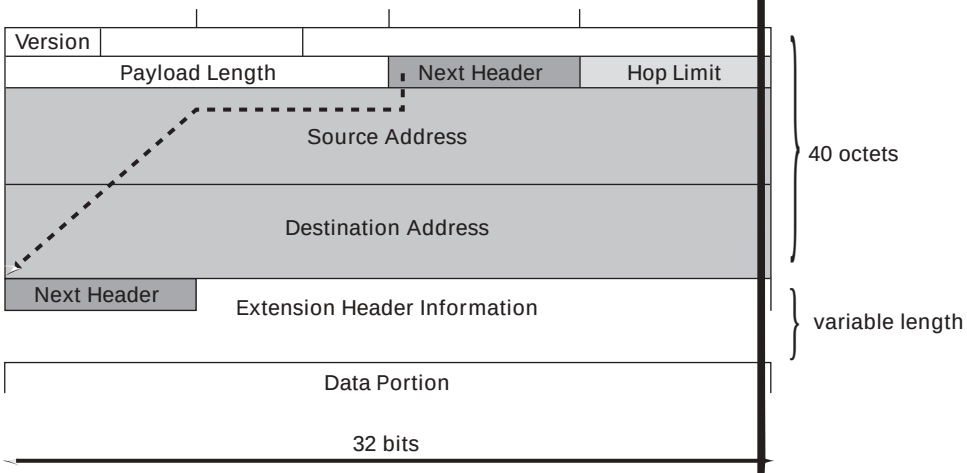

The IPv6 header contains the fields described in the following sections:

Description of IPv6 Header Fields

Version Number: The version is a 4-bit field as in IPv4. The field contains the number 6 for IPv6, instead of the number 4 for IPv4.

Traffic Class: The Traffic Class field is an 8-bit field similar to the type of service (ToS) field in IPv4. The Traffic Class field tags the packet with a traffic class that can be used in Differentiated Services. The functionalities are the same in IPv4 and IPv6.

Flow Label: The 20-bit Flow Label field is a new field in IPv6. The Flow Label field can be used to tag packets of a specific flow to differentiate the packets at the network layer. Hence, the Flow Label field enables identification of a flow and per-flow processing by the routers in the path. With this label, a router need not check deep into the packet to identify the flow, because this information is available in the IP packet header. The Flow Label allows applications on the end system to easily differentiate the traffic at the IP layer making it easier to provide QoS for packets that have been encrypted by IPSec.

For further information regarding the proposal for the implementation of the flow label, refer to the Internet Draft *draft-ietf-ipv6-flow-label-01.txt.*

Payload Length: Similar to the Total Length field in IPv4, the Payload Length field indicates the total length of the data portion of the packet.

Next Header: Similar to the Protocol field in the IPv4 packet header, the value of the Next Header field in IPv6 determines the type of information following the basic IPv6 header. The type of information following the basic IPv6 header can be a transport layer packet, such as a TCP or UDP packet, or an Extension Header, as shown in Figure 4.

IPv6 uses a different approach to manage optional information in the header. It defines extension headers that form a chain of headers linked together by the Next Header field, contained in each extension header. This mechanism provides more efficiency in the processing of extension headers, enables a faster forwarding rate, and leaves the router with less processing work for each packet. All extension headers are daisy-chained, each one pointing to the next one, until they reach the transport layer data.

Hop Limit: Similar to the Time to Live field in the IPv4 packet header, the value of the Hop Limit field specifies the maximum number of routers (hops) that an IPv6 packet can pass through before the packet is considered invalid. Each router decrements the value by one. Because there is no checksum in the IPv6 header, the router can decrement the value without needing to recalculate the checksum, which saves processing resources.

Source Address: The IPv6 source address field is similar to the Source Address field in the IPv4 packet header, except that the field contains a 128-bit source address for IPv6 instead of a 32-bit source address for IPv4.

Destination Address: The IPv6 destination address field is similar to the Destination Address field in the IPv4 packet header, except that the field contains a 128-bit destination address for IPv6 instead of a 32-bit destination address for IPv4.

IPv6 Extension Headers

Following the eight fields of the basic IPv6 packet header are the optional extension headers and the data portion of the packet. If present, each extension header is aligned to 64 bits. There is no fixed number of extension headers in an IPv6 packet. Together, the extension headers form a chain of headers that can be parsed for information such as TCP/UDP port.

The Next Header field of the previous header identifies the extension header. Typically, the final extension header has a Next Header field of a transport layer protocol, such as TCP or UDP. Figure 4 shows the IPv6 extension header format.

Figure 4: IPv6 Extension Header

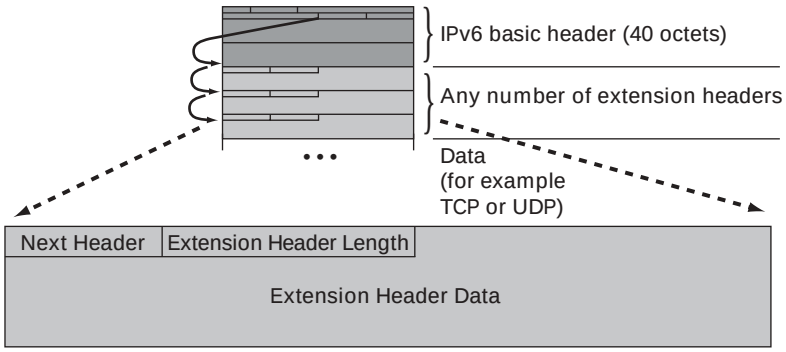

Order of Extension Headers

There are many types of extension headers. When multiple extension headers are used in the same packet, the order of the headers should be as follows:

1. **Hop-by-Hop Options header**. Used for the Router Alert (RSVP and MLDv1) and the Jumbogram, this header (value=0) is processed by all hops in the path of a packet. When present, the hop-by-hop options header always follows immediately after the basic IPv6 packet header.

2. **Destination Options header**. This header (value=60) can follow any hop-by-hop options header, in which case the destination options header is processed at the final destination and also at each visited address specified by a routing header. Alternatively, the destination options header can follow any Encapsulating Security Payload (ESP) header, in which case the destination options header is processed only at the final destination. For example, mobile IP uses this header.

3. **Routing header**. This header (value=43) is used for source routing and Mobile IPv6.

4. **Fragment header**. This header is used when a source must fragment a packet that is larger than the maximum transmission unit (MTU) for the path between itself and a destination device. The Fragment header is used in each fragmented packet.

5. **Authentication header and Encapsulating Security Payload header**. The Authentication header (value=51) and the ESP header (value=50) are used within IPSec to provide authentication, integrity, and confidentiality of a packet. These headers are identical for both IPv4 and IPv6.

6. **Upper-Layer header**. The upper-layer (transport) headers are the typical headers used inside a packet to transport the data. The two main transport protocols are TCP (value=6) and UDP (value=17).

Note: The source node should follow this order, but destination nodes should be prepared to receive in any order.

Routing Header

The Routing Header is one of the IPv6 extension headers and is identified by a value of 43 in the Next Header field. A routing header can appear either as the first extension header after the IPv6 basic header, or after another extension header.

As in any extension header, the first field of the routing header is the Next Header field, which identifies the type of header following the routing header. The second field is the length of the routing header. The "routing type" identifies the type of routing header used. The "segments left" identifies the number of intermediate routers that are in the data portion of the routing header. The routing header with routing type = 0 forces the routing through a list of intermediate routers. This action is similar to the "Loose Source Route" option in IPv4.

Figure 5 shows the use of the routing type 0 of the routing header and the routing path based on the intermediate routers R2 and R5. As in "Loose Source Route" in IPv4, the complete list of routers in the path is not necessary.

Figure 5: IPv6 Routing Header

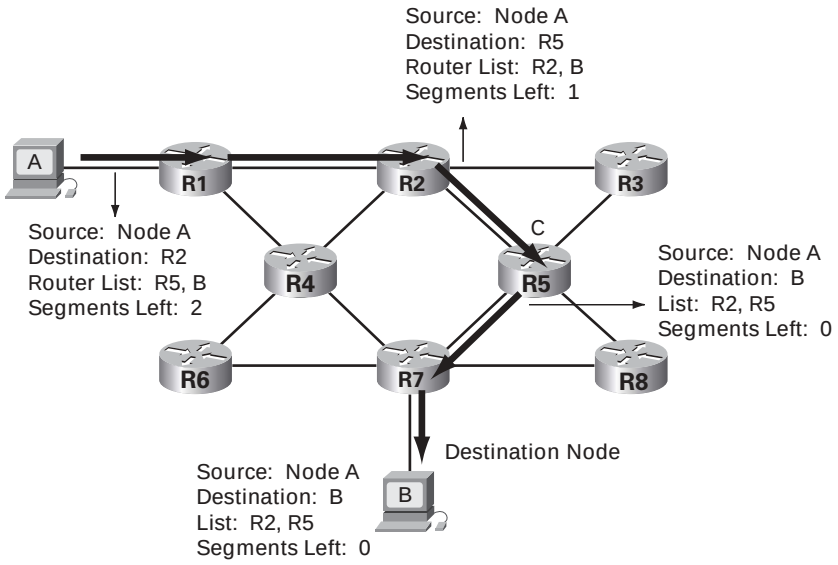

The way the routing header and the destination address of the IPv6 packet interact is new in IPv6. Upon receiving the packet, each intermediate router in the list will process the routing header by swapping the destination address to the next router in the list. The number of segments left is decremented and the packet is sent to the new destination. The final destination node (B) will receive a routing header where the number of segment left is zero. Because B is the final destination, it will process the next header following the routing header.

Fragment Header

IPv6 does not support fragmentation by routers. The source node does the fragmentation when the path MTU is not big enough. In IPv4, path MTU is optional and seldom used.

The fragment header is used when a node has to send a packet larger than the path MTU. In this situation, the source node slices the packet into fragments and sends each fragment in a separate packet and identifies the

fragments by adding the fragment header in the IP header of the packets.

The fields of the fragment header look like the fragment fields in the IPv4 header and include the following:

- A fragment offset that identifies the position of the specific fragment in the full original IP packet
- An identification number that identifies fragments belonging to the same original packet

The destination node then reassembles the packet by concatenating the received fragments in an order given by the fragment offset.

ICMPv6 Packet

Internet Control Message Protocol (ICMP) in IPv6 functions the same as ICMP in IPv4 (RFC 792). ICMPv6 generates error messages, such as ICMP destination unreachable messages and informational messages like ICMP echo request and reply messages.

Additionally, ICMP packets in IPv6 are used in the IPv6 neighbor discovery process, path MTU discovery, and the Multicast Listener Discovery (MLD) protocol for IPv6. IPv6 routers use MLD to discover multicast listeners (nodes that want to receive multicast packets destined for specific multicast addresses) on directly attached links. MLDv1 is described in RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*, which is based on version 2 of the Internet Group Management Protocol (IGMP) for IPv4. MLDv2 (draft) is similar to IGMPv3.

A value of 58 in the Next Header field of the basic IPv6 packet header identifies an IPv6 ICMP packet. ICMP packets in IPv6 are like a transport layer packet in the sense that the ICMP packet follows all the extension headers and is the last piece of information in the IPv6 packet.

Within IPv6 ICMP packets, the ICMPv6 Type and ICMPv6 Code fields identify IPv6 ICMP packet specifics, such as the ICMP message type. The value in the Checksum field is derived from the fields in the IPv6 ICMP packet and the IPv6 header. The ICMPv6 Data field contains error or diagnostic information relevant to IP packet processing.

Similar to ICMPv4, ICMPv6 is often blocked by security policies implemented in corporate firewalls because of attacks based on ICMP. However, ICMPv6 has the capability to use IPSec authentication and encryption. These security services decrease the possibilities of an attack based on ICMPv6. Figure 6 shows the IPv6 ICMP packet format.

Figure 6: IPv6 ICMP Packet

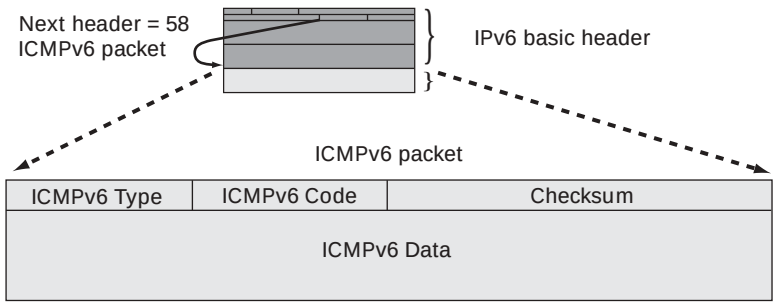

IPv6 Address Prefix

The IPv6 prefix is part of the address that represents the left-most bits that have a fixed value and represent the network identifier. IPv6 prefix is represented using the IPv6-prefix/prefix-length format just like an IPv4 address represented in the classless interdomain routing (CIDR) notation. The IPv6-prefix variable must conform to RFC 2373.

The /prefix-length variable is a decimal value that indicates the number of high-order contiguous bits of the address comprising the prefix, which is the network portion of the address. For example, 1080:6809:8086:6502::/64 is an acceptable IPv6 prefix. If the address ends in a double colon, the trailing double colon can be omitted. So, the same address can be written as 1080:6809:8086:6502/64. In either case, the prefix length is written as a decimal number 64 and represents the left-most bits of the IPv6 address.

IPv6 Address Types

There is a major difference in the IP address requirements between an IPv4 node and an IPv6 node. An IPv4 node typically uses one IP address; but an IPv6 node requires more than one IP address.

There are three major types of IPv6 addresses as follows:

- **Unicast**—An address for a single interface. A packet that is sent to a unicast address is delivered to the interface identified by that address.
- **Anycast**—An address for a set of interfaces that typically belong to different nodes. A packet sent to an anycast address is delivered to the closest interface—as defined by the routing protocols in use—identified by the anycast address.
- **Multicast**—An address for a set of interfaces (in a given scope) that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address (in a given scope).

IPv6 Address Assignment

An IPv6 address is assigned to a single interface, not a node. But, a single interface could be assigned multiple IPv6 addresses. Hence, it is easy to identify a node by any of its unicast addresses. The following are notable exceptions to these general rules:

- Multiple interfaces can have a single unicast address assigned to them when they are used for load sharing over multiple physical interfaces. The same is true when multiple physical interfaces are treated as a single interface at the Internet layer.
- Routers using unnumbered interfaces on point-to-point links are not assigned IPv6 addresses, because the interfaces do not function as a source or destination for IP datagrams.

IPv6 Unicast Address

A unicast address is an address for a single interface. A packet that is sent to a unicast address is delivered to the interface identified by that address. Although implementation details might depend on specific vendors, the Cisco IOS software supports the following IPv6 unicast address types:

- Global unicast address
- Site-local unicast address
- Link-local unicast address
- IPv4-mapped IPv6 address
-

What Is the Structure of a Global Unicast Address?

A *fixed prefix* of 2000::/3 (001) indicates a global IPv6 address. Addresses with a prefix of 2000::/3 (001) through E000::/3 (111), excluding the FF00::/8 (1111 1111) multicast addresses, are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format. The Internet Assigned Numbers Authority (IANA) is allocating the IPv6 address space in the ranges of 2001::/16 to the registries.

The following section describes the new addressing scheme recommended by the IETF in the Internet Draft draft-ietf-ipngwg-addr-arch-v3-07.txt.

The global unicast address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID. In the *IPv6 aggregatable global unicast address format* document (RFC 2374), the global routing prefix included two other hierarchically structured fields called Top-Level Aggregator and Next-Level Aggregator. Because these fields were policy-based, the IETF decided to remove the fields from the RFCs. However, some existing IPv6 networks deployed in the early days 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. This field allows an organization to use up to 65,535 individual subnets.

Use of EUI-64 Format in IPv6 Addresses

The 64-bit interface identifier in an IPv6 address is used to identify a unique interface on a link. A link is a network medium over which network nodes communicate using the link layer. The interface identifier may also be unique over a broader scope. In many cases, an interface identifier will be the same as or based on the link-layer (MAC) address of an interface. As in IPv4, a subnet prefix in IPv6 is associated with one link.

Interface identifiers used in global unicast and other IPv6 address types must be 64 bits long and constructed in the EUI-64 format. The EUI-64 format interface ID is derived from the 48-bit link-layer (MAC) address by inserting the hex number FFFE between the upper three bytes (OUI field) and the lower 3 bytes (serial number) of the link layer address. To make sure that the chosen address is from a unique Ethernet MAC address, the 7th bit in the high-order byte is set to 1 (equivalent to the IEEE G/L bit) to indicate the uniqueness of the 48-bit address.

What Is an IPv6 Site-Local Unicast Address?

Site-local unicast addresses are similar to the private addresses such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16 used in IPv4 networks. Private addresses can be used to restrict communication to a specific domain, or to assign addresses in a site that is not connected to the global Internet, without requiring a globally unique prefix. IPv6 routers must not advertise routes or forward packets that have site-local source or destination addresses outside the site boundary. If the site requires global connectivity in the future, a global unicast prefix must be assigned to that site. The site-local addressing plan initially defined for site-local addressing can be directly applied using the global unicast prefix.

A site-local unicast address shown in Figure 8 is an IPv6 unicast address that uses the prefix range FEC0::/10 (1111 1110 11) and concatenates the subnet identifier (the16-bit Subnet ID field) with the interface ID in the EUI-64 format. The site-local unicast address range uses 1/1024 of the total address space.

Figure 8: IPv6 Site-Local Unicast Address Format

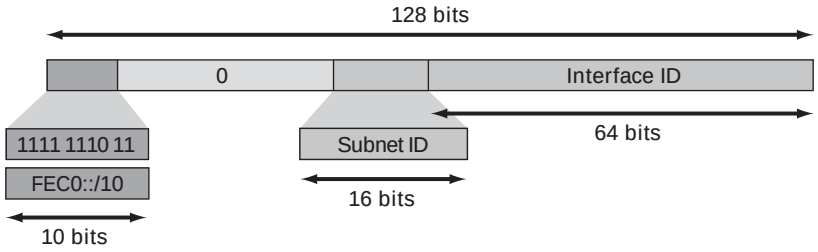

What Is an IPv6 Link-Local Unicast Address?

A link-local unicast address is an IPv6 unicast address that is automatically configured on an IPv6 node interface by using the link-local prefix FE80::/10 (1111 1110 11) and the interface ID in the EUI-64 format.

Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process discussed in Chapter 4. Link-local addresses are typically used to connect devices on the same local link network without the need for global addresses. Hence, link-local addresses are useful only in the context of the local link network.

Nodes on a local link can use link-local addresses to communicate with each other without the need for a router. IPv6 nodes do not need site-local or globally unique addresses to communicate. IPv6 routers must not forward to other links packets that have link-local source or destination addresses. FE80::/10 is the link-local unicast address range and uses 1/1024 of the IPv6 address space. Figure 9 shows the structure of a link-local address.

Figure 9: IPv6 Link-Local Unicast Address Format

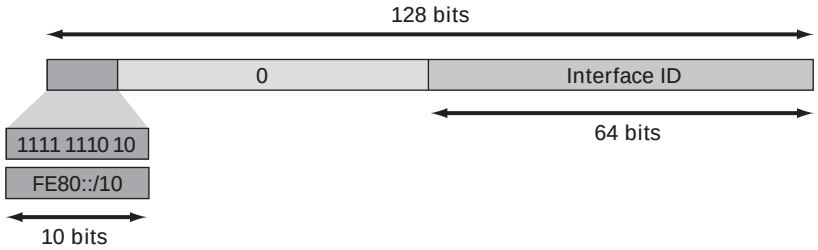

What Is an IPv4-Compatible IPv6 Address?

The IPv4-compatible IPv6 address is used in IPv6 transition mechanisms to tunnel IPv6 packets dynamically over IPv4 infrastructures. The IPv4-compatible IPv6 address is a type of IPv6 unicast address that embeds an IPv4 address in the low-order 32 bits and zeros in the high-order 96 bits of the IPv6 address.

The format of an IPv4-compatible IPv6 address is 0:0:0:0:0:0:A.B.C.D or ::A.B.C.D. The entire 128-bit IPv4-compatible IPv6 address is used as the IPv6 address of a node and the IPv4 address embedded in the low-order 32-bits is used as the IPv4 address of the node. IPv4-compatible IPv6 addresses are assigned to nodes that support both the IPv4 and IPv6 protocol stacks and are used in automatic tunnels discussed in Chapter 5. Figure 10 shows the structure of an IPv4-compatible IPv6 address and a few acceptable representations for the address.

Figure 10: IPv4-Compatible IPv6 Address Format

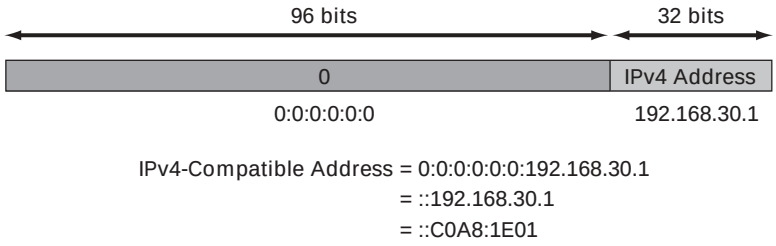

What Is an IPv4-Mapped IPv6 Address?

The IPv4-mapped IPv6 address is another type of IPv6 unicast address that embeds an IPv4 address in the low-order 32 bits, zeros in the high-order 80 bits, and ones in bits 81 through 96 of the IPv6 address. This address type is used to represent the address of an IPv4 node as an IPv6 address. On a dual stack node, an IPv6 application sending traffic to a destination represented by an IPv4-mapped IPv6 address will send IPv4 packets to that destination. Figure 11 shows the structure of an IPv4-mapped IPv6 address.

Figure 11: IPv4-Mapped IPv6 Address Format

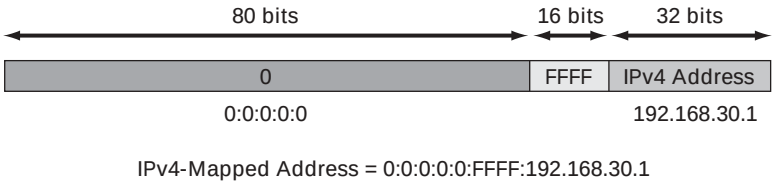

IPv6 Anycast Address

The anycast address is a global unicast address that is assigned to a set of interfaces that typically belong to different nodes. Hence an anycast address identifies multiple interfaces. A packet sent to an anycast address is delivered to the closest interface—as defined by the routing protocols in use—identified by the anycast address. Anycast addresses are syntactically indistinguishable from global unicast addresses because anycast addresses are allocated from the global unicast address space.

Note: Anycast addresses must not be used as the source address of an IPv6 packet.

IPv6 Multicast Address

An IPv6 multicast address (Figure 12) is an IPv6 address that has a prefix of FF00::/8 (1111 1111). The multicast address range uses 1/256 of the total IPv6 address space. An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address.

The second octet following the prefix defines the lifetime and scope of the multicast address. A permanent multicast address has a lifetime parameter equal to 0; a temporary multicast address has a lifetime parameter equal to 1. A multicast address that has the scope of a an interface, link, subnet, admin, site, organization, or a global scope has a scope parameter of 1, 2, 3, 4, 5, 8, or E, respectively. The IPv6 addressing scheme is designed to support millions of multicast group addresses.

Figure 12: IPv6 Multicast Address Format

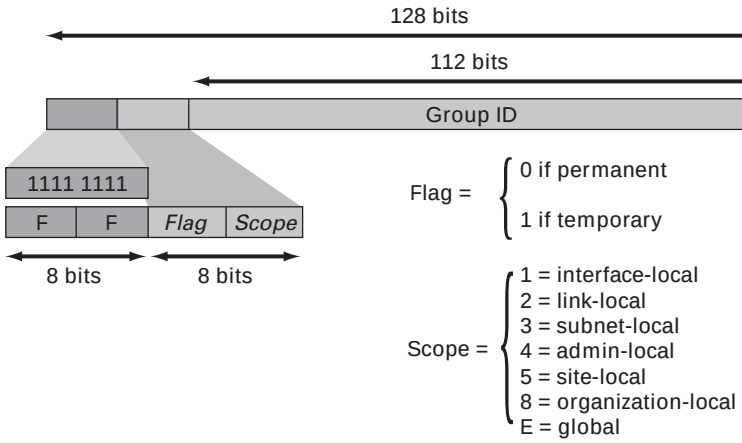

Within the reserved multicast address range of FF00:: to FF0F::, the following addresses are assigned to identify specific functions:

FF01::1—All Nodes within the node-local scope (that is, only for that host)

FF02::1—All Nodes on the local link (link-local scope).

FF01::2—All Routers within the node-local scope

FF02::2—All Routers on the link-local scope

FF05::2—All Routers in the site (site-local scope)

FF02::1:FFXX:XXXX—Solicited-Node multicast address, where XX:XXXX represent the last 24 bits of the IPv6 address of node.

Note that the time-to-live (TTL) field is not used in IPv6 multicast.

The correlation between an IPv6 multicast address and the Ethernet address is that the last 32 bits of the IPv6 multicast address are added to the 33:33: prefix for multicast Ethernet. A host sending a packet to an IPv6 multicast address uses this newly formed multicast Ethernet address to reach the destination on the local link.

Multicast Group Membership Requirement for IPv6 Nodes

IPv6 nodes, both hosts and routers, are required to join (receive packets destined for) the following multicast groups:

- All-nodes multicast group FF02:0:0:0:0:0:0:1 (scope is link-local)
- Solicited-node multicast group FF02:0:0:0:0:1:FF00:0000/104 for each of its assigned unicast and anycast addresses

Additionally, IPv6 routers must also join the all-routers multicast group FF02:0:0:0:0:0:0:2 (scope is link-local).

What Is an IPv6 Solicited-Node Multicast Address?

Solicited-node multicast addresses are used in neighbor solicitation messages to help with neighbor discovery, which is discussed in Chapter 4. The solicited-node multicast address is a multicast group address that corresponds to an IPv6 unicast or anycast address. An IPv6 node must join the associated solicited-node multicast group for every unicast and anycast address it has been assigned. The IPv6 solicited-node multicast address has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6 unicast or anycast address, as shown in Figure 13.

Figure 13: IPv6 Solicited-Node Multicast Address Format

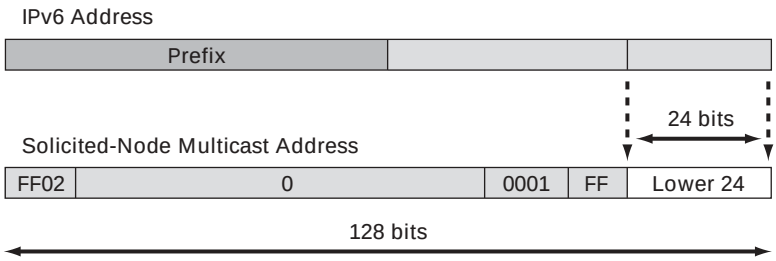

For example, the solicited-node multicast address corresponding to the IPv6 address 2037::01:800:200E:8C6C is FF02::1:FF0E:8C6C.

Special IPv6 Addresses

In addition to all the unicast addresses described in this section, IPv6 also supports unspecified and loopback addresses.

What Is an IPv6 Unspecified Address?

An unspecified IPv6 address is a special address used as a placeholder by a node when it does not have an address. For example, a node uses the unspecified address upon startup, when a node does not have an assigned address and the node requests an address from a dynamic host configuration server (DHCP) server, or when the duplicate address detection packet is sent. The unspecified address 0:0:0:0:0:0:0:0 is also represented by 0::0 or more commonly by ::/128.

The IPv6 unspecified address cannot be assigned to any interface and should not be used as destination addresses in IPv6 packets or the IPv6 routing header.

What Is an IPv6 Loopback Address?

IPv6 loopback address identifies the local interface in the IP stack. It is similar to the 127.0.0.1 loopback address in IPv4. The IPv6 loopback address is 0:0:0:0:0:0:0:1 or is simply represented by ::1.

The IPv6 loopback address cannot be assigned to a physical interface and the IPv6 routers do not forward packets that have the IPv6 loopback address as their source or destination address.

IPv6 Address Allocation

The Internet Assigned Numbers Authority (IANA) allocates 2001::/16 to registries from the full address space. From IANA, each registry gets a /23 prefix within the 2001::/16 space, as follows:

- 2001:0200::/23 and 2001:0C00::/23 allocated to Asia Pacific Network Information Centre (APNIC) for use in Asia.
- 2001:0400::/23 allocated to American Registry for Internet Numbers (ARIN) for use in the Americas.
- 2001:0600::/23 and 2001:0800::/23 allocated to Reseaux IP Europeens—Network Coordination Center (RIPE NCC) for use in Europe and the Middle East.

The registries then allocate an initial /32 prefix to the IPv6 ISPs and the ISPs allocate a /48 prefix (out of the /32) to each customer or site. The /48 prefix of site could be further allocated to each LAN using a /64 prefix for a maximum of 64 bits ID hosts in each LAN. Each site could subnet the site into a maximum of 65,535 LANs. A site should make an address plan prior to beginning allocation of its /48 space.

In order to receive a /32 prefix address block from a registry, an ISP must have an exterior routing protocol peering with at least 3 other ISPs and either have at least 40 customers or demonstrate a clear intent to provide an IPv6 service within 12 months.

For the latest information about allocation of IPv6 address space to the registries by IANA, refer to the URL at http://www.iana.org/assignments/ipv6-tla-assignments.

6BONE Network Address Allocation

The 6BONE is a worldwide network of IPv6 networks using IPv6 links carrying IPv6 traffic over WAN or LAN over IPv4 tunnels on the current Internet. The 6BONE is a testbed used to test new protocols, implementations, transition mechanisms, and operational procedures. The 6BONE collaborative project is overseen by the IETF.

The current allocations of the 6BONE address starts at 3ffe:0000::/16, where any pseudo Top-Level Aggregator (pTLA) receives a /28 prefix. This prefix is inside the 3ffe:0800::/28 range and allows for a maximum of 2048 pTLAs. An end site receives a /48 from its upstream provider and a LAN within a site is assigned a /64 prefix from that site prefix.

The 6BONE topology is a hierarchy of provider networks. The 6BONE address allocation by the IANA and the 6BONE policy is defined in RFC 2921, *6BONE pTLA and pNLA Formats (pTLA).*

How Is an IPv6 Address Represented in a URL?

In a URL, the colon is already used to indicate an optional port number, as shown in this URL example: http://www.abc.test:8080/index.html. So, the colon cannot be used to indicate an IPv6 address in a URL. If a URL contains two colons, a URL parser must be able to differentiate between the colon of a port number and the colon inside an IPv6 address. This is impossible because of the use of the compression technique.

To identify the IPv6 address while still keeping the colon, the address must be enclosed between brackets, as shown in the following example:

http://**[2001:1:4F3A::206:AE14]**:8080/index.html

The use of IPv6 addresses inside a URL is cumbersome and is recommended only for diagnostic purposes, or when no naming service is available. Otherwise, it is advisable to use only fully qualified domain names.

How Many IP Addresses Does an IPv6 Host Require?

An IPv6 node requires the following IPv6 addresses for proper operation:

- Link-local address for each interface
- Assigned unicast address(es)
- Loopback address
- All-nodes multicast address
- Solicited–node multicast address for each of its assigned unicast and anycast addresses
- Multicast addresses of all other groups to which the host belongs
- Site-local address, if used

How Many IP Addresses Does an IPv6 Router Require?

An IPv6 router requires the following IPv6 addresses for proper operation:

- All the required node addresses
- All-router multicast addresses
- Subnet-router anycast addresses for the interfaces configured to act as forwarding interfaces
- Other anycast configured addresses
- Specific multicast addresses for routing protocols

Chapter 4

Operation of IPv6

The IPv6 neighbor discovery protocol and the internet message control protocol (ICMP) are critical to the operation of IPv6.

The following topics are covered in this chapter:

- Neighbor discovery
- Router discovery
- Stateless autoconfiguration
- Path maximum transfer unit (MTU) discovery
- Dynamic Host Configuration Protocol Version 6 (DHCPv6)
- Domain Name Server (DNS) Operation

Neighbor Discovery

The neighbor discovery protocol enables IPv6 nodes and routers to:

- Determine the link-layer address of a neighbor on the same link.
- Find neighboring routers.
- Keep track of neighbors.

The IPv6 neighbor discovery process uses IPv6 ICMP (ICMPv6) messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the reachability of a neighbor, and keep track of neighbor routers. Every IPv6 node is required to join the multicast groups corresponding to its unicast and anycast addresses.

The IPv6 neighbor discovery process uses the following mechanisms for its operation:

- Neighbor solicitation
- Neighbor advertisement

What Is IPv6 Neighbor Solicitation?

Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link. This function is similar to the ARP in IPv4, but avoids broadcasts used in IPv4 ARP messages, where all nodes receive unnecessary broadcast requests that do not concern them.

The source node takes the right-most 24 bits of the IPv6 address of the destination node and sends a neighbor solicitation message, which has a value of 135 in the Type field of the ICMP packet header, to the solicited-node multicast group address on the local link. The destination node will respond with its link-layer address. To send a neighbor solicitation message, the source node must first identify the IPv6 unicast address of the destination node using a naming service mechanism such as DNS.

Neighbor solicitation message is also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. Figure 14 shows how the neighbor solicitation message is used to determine the link-layer address of a neighbor.

Figure 14: Neighbor Solicitation Message

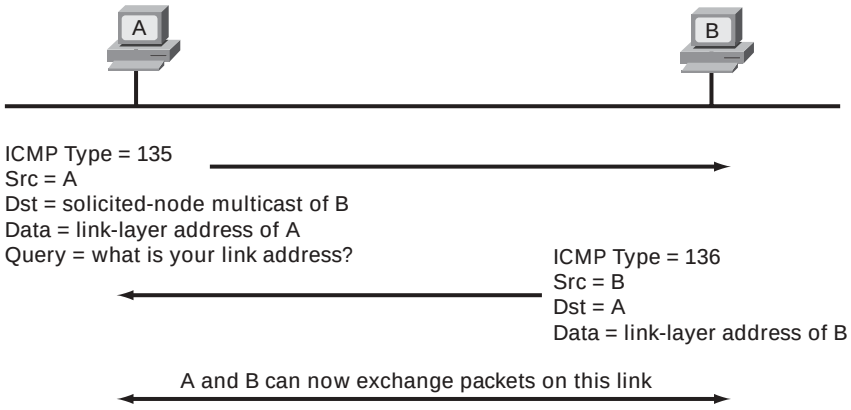

What Is IPv6 Neighbor Advertisement?

The IPv6 neighbor advertisement message is a response to the IPv6 neighbor solicitation message. After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message on the local link with a value of 136 in the Type field of the ICMP packet header. After receiving the neighbor advertisement, the source node and destination node can communicate.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link.

IPv6 Router Discovery

IPv6 router discovery is a process used by IPv6 nodes to discover the routers on the local link. The IPv6 router discovery process is similar to ICMP router discovery in IPv4, except for one major difference described later in this section.

The IPv6 router discovery process uses the following messages:

- Router advertisements
- Router solicitations

What Is IPv6 Router Advertisement?

Router advertisement messages are periodically sent out on each configured interface of an IPv6 router. Router advertisements are also sent out in response to router solicitation messages from IPv6 nodes on the link. The router advertisements are sent to the all-nodes link-local multicast address (FF02 ::1) or the unicast IPv6 address of a node that sent the router solicitation messages.

Router advertisement has a value of 134 in the Type field of the ICMP packet header and contains the following information in the message:

- Whether nodes could use address autoconfiguration
- Flags to indicate the type of autoconfiguration (stateless or stateful) that can be completed
- One or more on-link IPv6 prefixes that nodes on the local link could use to automatically configure their IPv6 addresses

- Lifetime information for each prefix included in the advertisement
- Whether the router sending the advertisement should be used as a default router and, if so, the amount of time (in seconds) the router should be used as a default router
- Additional information for hosts, such as the hop limit and maximum transmission unit (MTU) a host should use in packets that it originates

The IPv6 nodes on the local link receive the router advertisement messages and use the information to keep the information about default router and prefix lists and other configuration parameters updated. Figure 15 shows an example of the router advertisement.

Figure 15: Router Advertisement

RA packet definitions:
ICMP Type = 134
Src = router link-local address
Dst = all-nodes multicast address
Data = options, prefix, lifetime, autoconfig flag

What Is IPv6 Router Solicitation?

When a host does not have a configured unicast address, for example at system startup, it sends a router solicitation message. A router solicitation is helpful, because it enables the host to autoconfigure itself quickly without having to wait for the next scheduled router advertisement message. A router solicitation message has a value of 133 in the Type field of the ICMP packet header.

The source address used in a router solicitation messages is usually the unspecified IPv6 address (0:0:0:0:0:0:0:0). If the host has a configured unicast address, the unicast address of the interface sending the router solicitation message is used as the source address in the message.

The destination address in the router solicitation messages is the all-routers multicast address (FF02::2) with the link-local scope. When a router advertisement is sent in response to a router solicitation, the destination address used in the router advertisement message is the unicast address of the source of the router solicitation message.

Note: A router solicitation is sent at boot time and only three times afterward to avoid flooding of router solicitation packets in the absence of a router on the network.

IPv6 Redirect Message

As with IPv4, an IPv6 redirect message is sent by a router only to help with the reroute of a packet to a better router. The node receiving the redirect message will then readdress the packet to a better router. Routers send redirect messages only for unicast traffic, only to the originating nodes, and to be processed by the nodes.

Stateless Autoconfiguration

Stateless autoconfiguration is a key feature of IPv6. It enables serverless basic configuration of the IPv6 nodes and easy renumbering. Stateless autoconfiguration uses the information in the router advertisement messages to configure the node. The prefix included in the router advertisement is used as the /64 prefix for the node address. For Ethernet, the remaining 64 bits are obtained from the interface ID in EUI-64 format. Thus, an IPv6 node can autoconfigure itself with a globally unique IPv6 address by appending its link-layer address (EUI-64 format) to the local link prefix (64 bits).

Renumbering of IPv6 Nodes

Renumbering of IPv6 nodes is possible with the help of router advertisements. Router advertisement messages contain both the old prefix and the new prefix. A decrease in the lifetime value of the old prefix alerts the nodes to use the new prefix, while still keeping their current connections intact with the old prefix. During this period of time, nodes have two unicast addresses in use. When the old prefix is no longer usable, the router advertisements will include only the new prefix.

If stateless autoconfiguration is not used for renumbering, other ways of renumbering should be used. Autoconfiguration greatly helps the renumbering process. Renumbering requires changes to the DNS entries and the introduction of new IPv6 DNS records. Renumbering of a whole site also requires that all the routers be renumbered. A router renumbering protocol has been proposed at the IETF.

Stateless autoconfiguration does not address the issue of finding the DNS server for DNS resolution or registering the computer in the DNS space. These issues are being discussed at the IETF.

How Does Duplicate Address Detection Work?

IPv6 also provides a safety mechanism to detect duplicate addresses in the network and prevent any address collision. IPv6 uses neighbor solicitation to detect if another node on the link has the same IPv6 address.

Duplicate address detection is used during the autoconfiguration process to ensure that no other node is using the autoconfigured address.

Path Maximum Transmission Unit Discovery

Because IPv6 routers do not handle fragmentation, fragmentation is done by the originating node or source node of a packet, when necessary. The path MTU discovery process is critical to handling of fragmentation by the hosts in IPv6 networks. IPv6 uses the path MTU discovery to find the maximum MTU in a path between the source and the destination. The source node starts the path MTU discovery process before actually sending the packets. When the path MTU of every link along a given data path in an IPv6 network is not large enough to accommodate the size of the packets, the source node fragments the packet and resends it.

As in IPv4, path MTU discovery in IPv6 allows a node to dynamically discover and adjust to differences in the MTU size of every link along a given data path. In IPv4, the minimum link MTU size is 68 octets and the recommended minimum is 576 octets, which is the minimum reassembly buffer size. So, any IPv4 packet must be at least 68 octets in length.

In IPv6, the minimum link MTU is 1280 octets, but the recommended MTU value for IPv6 links is 1500 octets. The maximum packet size supported by the basic IPv6 header is 64,000 octets. Larger packets called jumbograms could be handled using a hop-by-hop extension header option.

How Does IPv6 Path MTU Discovery Work?

The IPv6 source node sends a packet equal in size to the maximum MTU of its link layer. In the example shown in Figure 16, an MTU size of 1500 is used. The packet is forwarded through the network up to the destination, unless it encounters a smaller MTU in the path. When the packet encounters a link with a smaller MTU, the router sends the source node an ICMP error message of type 2, named "packet too big." The content of the ICMP packet includes the MTU size of the next link, which is smaller than the size of the packet (for examples 1400 and 1300 for the last two links in the Figure 16).

The source IPv6 node then resends a packet equal to the size of the received maximum MTU. This process repeats until the packet reaches the destination. In this example, the path MTU of the last link is 1300.

Figure 16: Path MTU Discovery

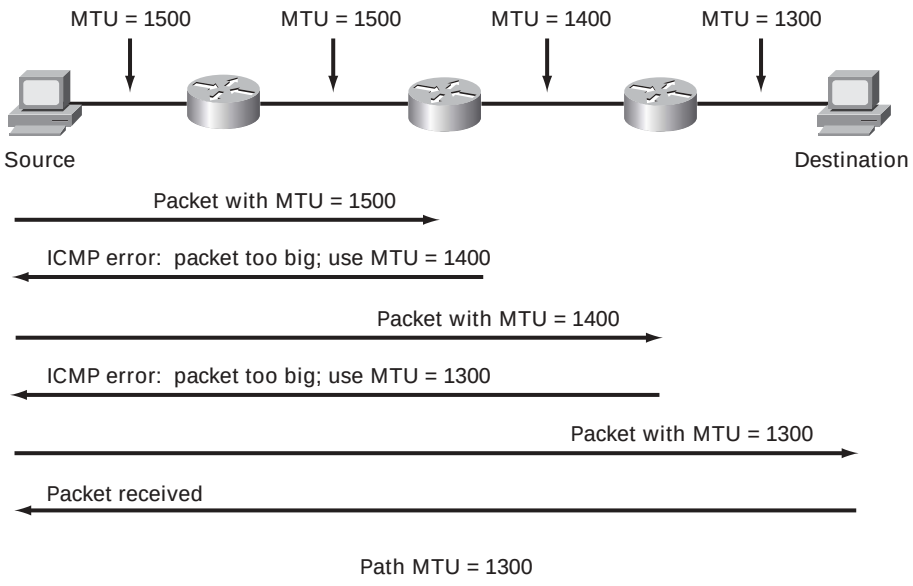

Dynamic Host Configuration Protocol Version 6

The process for acquiring configuration data for a client is similar to that in IPv4. However, DHCPv6 uses multicast for many of its messages. Initially, the client must first detect the presence of routers on the link using neighbor discovery messages. If a router is found, then the client examines the router advertisements to determine if DHCP should be used. If the router advertisements enable use of DHCP on that link or if no router is found, then the client starts a DHCP solicitation phase to find a DHCP server.

The following are the benefits of DHCPv6:

- Provides more control than serverless/stateless auto-configuration
- Used in a routerless environment, using only servers
- Used concurrently with stateless auto-configuration
- Used for renumbering

- Used for automatic domain name registration of hosts using dynamic DNS
- Used to delegate IPv6 prefix to leaf customer-premise equipment (CPE) routers

IPv6 Domain Name System Operation

IPv6 introduces new DNS record types for IPv6 addresses that are supported in the DNS name-to-address and address-to-name lookup processes. DNS query is possible over an IPv4 transport or an IPv6 transport. But, DNS root servers are not yet reachable through an IPv6 transport. The record types are as follows:

AAAA record—Also known as a "quad A" record, this record maps a host name to an IPv6 address. This record is equivalent to an A record in IPv4 and uses the format: www.abc.test AAAA 3FFE:B00:C18:1::2. The IETF has decided to use this record for host name-to-IP address resolution.

A6 record

Chapter 5

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. The Internet consists of hundreds of thousands of IPv4 networks and millions of IPv4 nodes. The challenge lies in making the integration and transition as transparent as possible to the end users.

The approximate time line for IPv6 deployment in various sectors is expected to be as follows:

- **1996-2002:** As is the case with any new technology, the early adopters including technology enthusiasts and academic institutions were first to deploy IPv6 networks. To support the early adopters, IPv6 for Cisco IOS software has been available for early field trial (EFT) since 1996.
- **2001-2005:** Porting of existing applications to IPv6, a critical requirement for the adoption of IPv6, started during the latter part of the year 2001. This process is expected to take more than three years to complete.
- **2001-2005:** Internet service providers started deploying IPv6 during the latter part of the year 2001 to be able to provide IPv6 services to their customers. The ISP deployment phase is expected to last longer than three years.
- **2003-2010:** Consumer adoption of IPv6 services is dependent on the availability of applications such as distributed gaming and peer-to-peer computing and is expected to become popular in the year 2003 and last longer than five years.
- **2003-2010:** Similar to consumer adoption of IPv6 services, enterprises are waiting for the full availability of applications and are expected to start deploying IPv6 in the year 2003 and beyond.

The IETF IPv6 working group has designed several strategies for the deployment of IPv6. The following transition strategies are covered in this chapter:

- Deploying IPv6 over dual stack backbones
- Deploying IPv6 over IPv4 tunnels
- Deploying IPv6 over dedicated data links
- Deploying IPv6 over multiprotocol label switching (MPLS) backbones
- Deploying IPv6 using protocol translation mechanisms

Transition Mechanisms

Network designers recommend deploying IPv6 at the edge first and then moving towards the network core to reduce the cost and operational impacts of the integration. 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 dual-stack backbones**—This technique allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. All routers or a portion of them (for example, access CPE routers and aggregation routers are dual stack, but core routers stay as they are) 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.
- **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 6to4 for the WAN and intra-site automatic tunnel addressing protocol (ISATAP) for the campus environment.
- **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 ATM permanent virtual circuits (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 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.

Using IPv4-IPv6 Protocol Dual Stack Devices

Dual stack backbone is a basic strategy for routing both IPv4 and IPv6 and requires network devices such as routers and end systems running both IPv4 and IPv6 protocol stacks. Dual stack end systems allow applications to migrate one at a time from an IPv4 to an IPv6 transport. Applications that are not upgraded to support IPv6 stack can coexist with upgraded applications on the same end system.

As shown in Figure 18, new and upgraded applications simply make use of both the IPv4 and IPv6 protocol stacks. A new application-programming interface (API) has been defined to support both IPv4 and IPv6 addresses and DNS requests. An application can be upgraded to the new API and still use only the IPv4 protocol stack.

Figure 18: IPv4-IPv6 Dual Stack

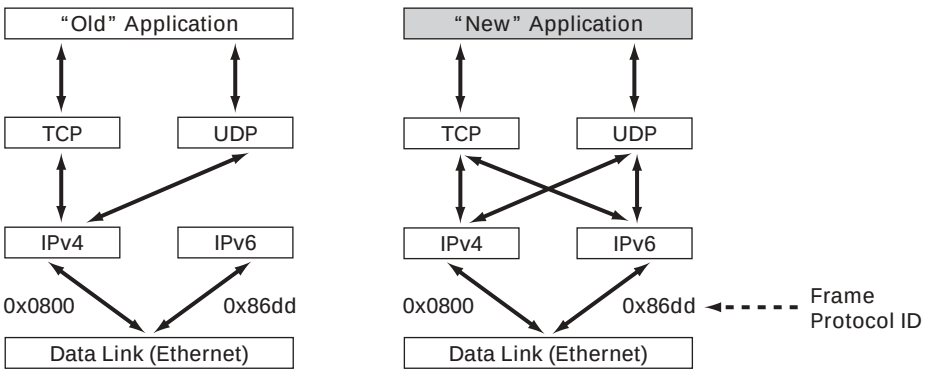

Applications choose between using IPv4 or IPv6 protocol based on name lookup; both the IPv4 and IPv6 addresses may be returned from the DNS, with the application (or the system according to the rules defined in the IETF document Default Address Selection for IPv6) selecting the correct address based on the type of IP traffic and particular requirements of the communication.

An application that supports dual IPv4 and IPv6 protocol stacks requests all available addresses for the destination host name (for example, www.a.com) from a DNS server. The DNS server replies with all available addresses (both IPv4 and IPv6 addresses) for www.a.com. The application chooses an address—in most cases, IPv6 addresses are the default choice—and connects the source node to the destination using the IPv6 protocol stack. Figure 19 shows an example of IPv4 and IPv6 dual stack operation.

Figure 19: IPv4-IPv6 Dual Stack Operation

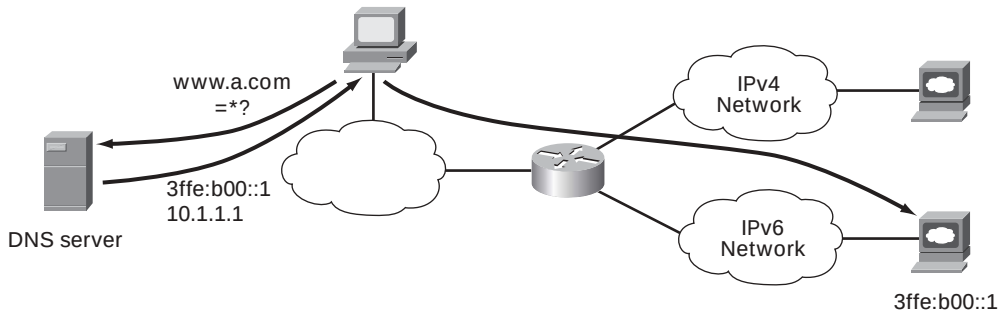

Deploying IPv6 Using Dual Stack Backbones

With the dual stack backbone deployment, 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.

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.

Deploying IPv6 over IPv4 Tunnels

Tunneling encapsulates IPv6 traffic within IPv4 packets so 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 20 shows the use of IPv6 over IPv4 tunnels.

Figure 20: IPv6 over IPv4 Tunnels

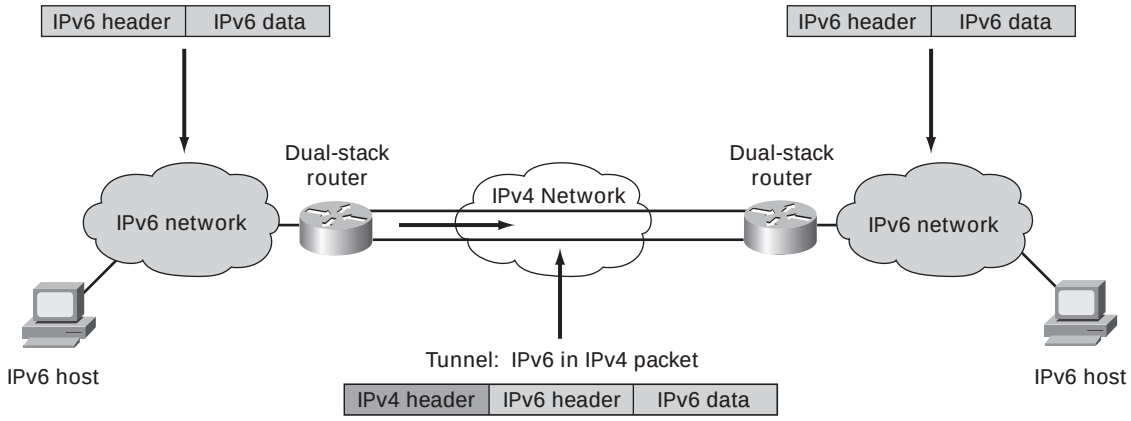

For example, tunneling allows service providers to offer an end-to-end IPv6 service without major upgrades to the infrastructure and without impacting current IPv4 services and allows enterprises to interconnect isolated IPv6 domains over their existing IPv4 infrastructures, or to connect to remote IPv6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

A variety of tunnel mechanisms are available for deploying IPv6. These mechanisms include manually created tunnels such as IPv6 manually configured tunnels (RFC 2893) and IPv6 over IPv4 GRE tunnels, semiautomatic tunnel mechanisms, and fully automatic tunnel mechanisms such as IPv4-compatible and 6to4 tunnels. Other tunneling techniques, such as ISATAP on campus, 6over4, and tunnel broker service (provided by service providers) are also available.

Tunneling Requirements

All tunneling mechanisms require that the endpoints of the tunnel run both IPv4 and IPv6 protocol stacks, that is, endpoints 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 dual-stack approach is similar to running IP and either IPX, DECnet, or AppleTalk on the same router, something Cisco IOS® Software has done since its inception.

For proper operation of the 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.

Tunneling and Security

It is possible to protect the IPv6 traffic 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.

Because protecting tunnels in this way may negatively impact performance, design considerations should balance this loss of performance against the security that can be achieved by careful configuration of the network.

Note: If a middle device between the two endpoints of the tunnel filters out IPv4 protocol 41, which is the IPv6 traffic in IPv4 encapsulation, the tunnel will not work.

IPv6 Tunnel Mechanisms

Not all transition strategies are 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 discusses the details about the following IPv6 tunneling techniques to be used over IPv4 networks.

- IPv6 Manually Configured Tunnel
- IPv6 over IPv4 GRE Tunnel
- Automatic IPv4-Compatible Tunnel
- Automatic 6to4 Tunnel
- ISATAP Tunnel
- Teredo Tunnel

IPv6 Manually Configured Tunnel

The primary use of a configured tunnel is to provide stable and secure connections for regular 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 used as tunnel endpoints must be dual-stack devices. Manual tunnels are used between two points and require configuration of both the source and destination addresses of the tunnel, whereas automatic tunnel mechanisms need to be only enabled and are more transient.

Because each tunnel is independently managed, the more tunnel endpoints you have, the more tunnels you need, and the greater is the management overhead. As with other tunnel mechanisms, network address translation (NAT) is not allowed along the path of the tunnel. Figure 21 shows the configuration of a manually configured tunnel.

Figure 21: Manually Configured Tunnel

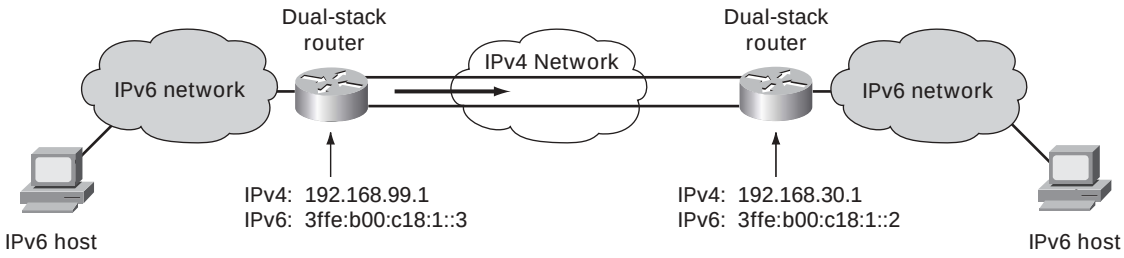

Refer to RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers* for further information on IPv6 manually configured tunnels.

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 manually configured tunnels, these tunnels are links between two points, with a separate tunnel for each link. The GRE tunnels are not tied to a specific passenger or transport protocol, but in this case carry IPv6 traffic as the passenger protocol over GRE as the carrier protocol.

Similar to the manual tunnels, the GRE tunnels are used between two points and require configuration of both the source and destination addresses of the tunnel. The edge routers and end systems used as tunnel end points must be dual stack devices.

Because the integrated IS-IS routing protocol runs over a Layer 2 data link, tunneling techniques other than GRE cannot be used. 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. Figure 22 shows how an IPv6 packet is carried over a GRE tunnel.

Figure 22: IPv6 over GRE Tunnel

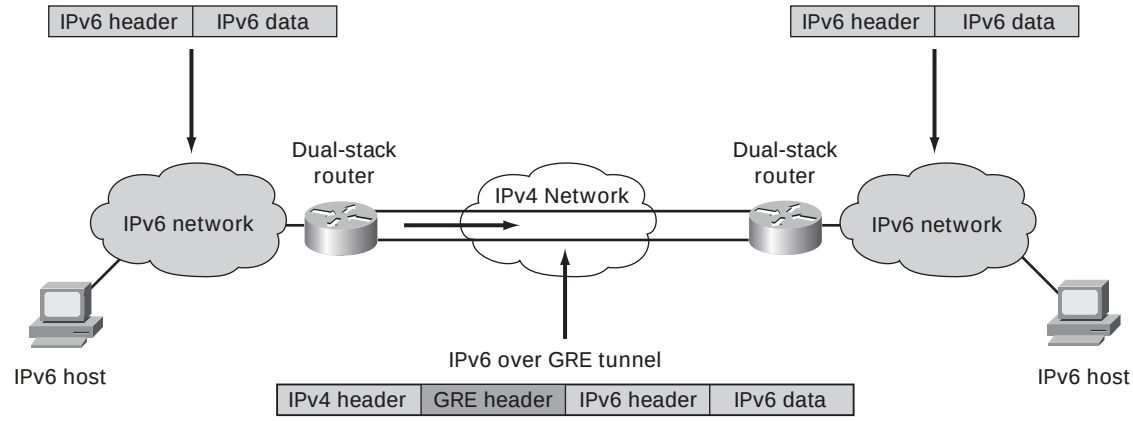

As with manually configured tunnels, you configure the IPv4 and IPv6 addresses of the dual-stack router on the GRE tunnel interface, and identify the entry and exit (or source and destination) points of the tunnel, using IPv4 addresses.

Because each GRE tunnel is independently managed, the more tunnel endpoints you have, the more tunnels you need, and the greater is the management overhead. As with other tunnel mechanisms, network address translation (NAT) is not allowed along the path of the tunnel.

Automatic IPv4-Compatible Tunnel

The automatic IPv4-compatible tunnel is an IPv6 over IPv4 tunnel mechanism, which uses an IPv4-compatible IPv6 address. An IPv4-compatible IPv6 address is the concatenation of zeros in the left-most 96 bits and an IPv4 address embedded in the last 32 bits. For example, ::192.168.99.1 is an IPv4-compatible IPv6 address.

Although an automatic tunnel can be configured between end systems, edge routers, or an edge router and an end system, the automatic IPv4-compatible tunnel has mainly been used to establish connection between routers.

Unlike a manually configured tunnel, the automatic IPv4-compatible tunnel technique constructs tunnels with remote nodes on the fly. Manual configuration of the endpoints of the tunnel is not required because the tunnel source and the tunnel destination are automatically determined by the IPv4 address. The automatic tunnels are set up and taken down as required, and last only as long as the communication. Figure 23 shows the configuration of an automatic IPv4-compatible tunnel.

Figure 23: Automatic IPv4-Compatible Tunnel

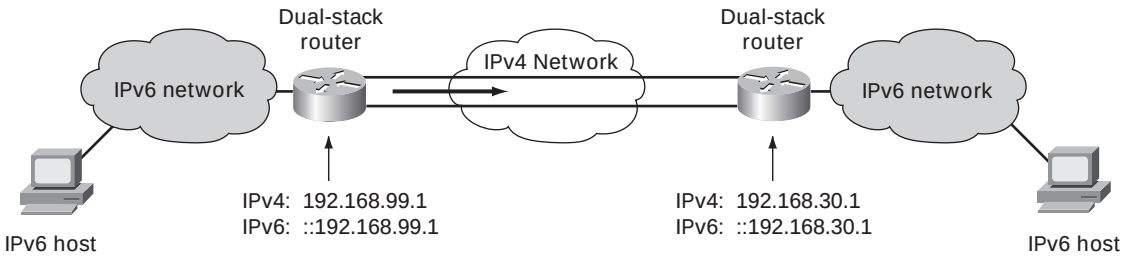

Although an easy way to create tunnels, the IPv4-compatible tunnel mechanism does not scale well for IPv6 networks deployment, because each host requires an IPv4 address removing the benefit of the large IPv6 addressing space. The IPv4-compatible tunnel is largely replaced by the 6to4 (RFC 3056, *Connection of IPv6 Domains via IPv4 Clouds*) automatic tunnel mechanism. Hence, the use of IPv4-compatible tunnel as a transition mechanism is nearly deprecated.

For further information on IPv4-compatible tunnels, refer to RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers.*

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 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 6to4 tunnel treats 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 automatically builds the IPv4 tunnel using a unique routing prefix 2002::/16 in the IPv6 address with the IPv4 address of the tunnel destination concatenated to the unique routing prefix. The key requirement is that each site has a 6to4 IPv6 address. Each site, even if it has just one public IPv4 address, has a unique routing prefix in IPv6. Figure 24 shows the configuration of a 6to4 tunnel interconnecting 6to4 domains.

Figure 24: Automatic 6to4 Tunnel

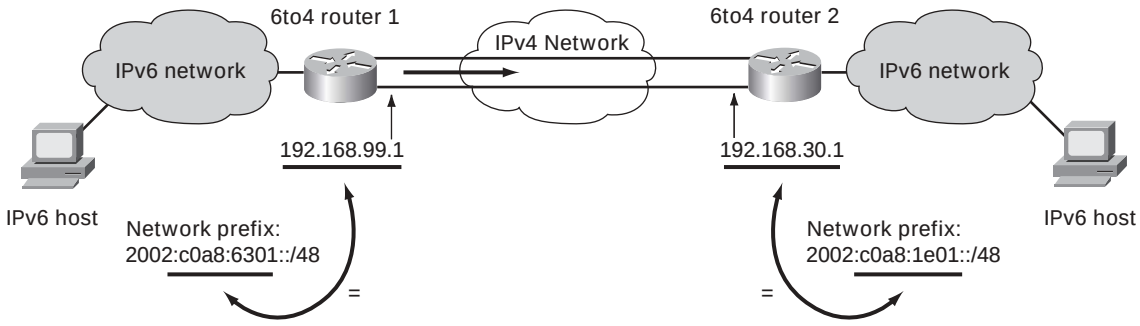

We recommend that each site have only one 6to4 address assigned to the external interface of the router. All sites need to run an IPv6 interior routing protocol, such as routing information protocol next generation (RIPng) for routing IPv6 within the site; exterior routing is handled by the relevant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

For further information on 6to4 tunnels, refer to RFC 3056, *Connection of IPv6 Domains via IPv4 Clouds.*

6to4 Relay Routers

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.

6to4 enables the edge router to forward packets to any destination with a 2002::/16 prefix. However, other IPv6 destinations are unreachable, unless one of the 6to4 edge routers, specified as a 6to4 relay, offers traffic forwarding to the IPv6 Internet.

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. Figure 25 shows the use of a 6to4 relay router for interconnecting 6to4 and native IPv6 domains.

Figure 25: 6to4 Relay Router

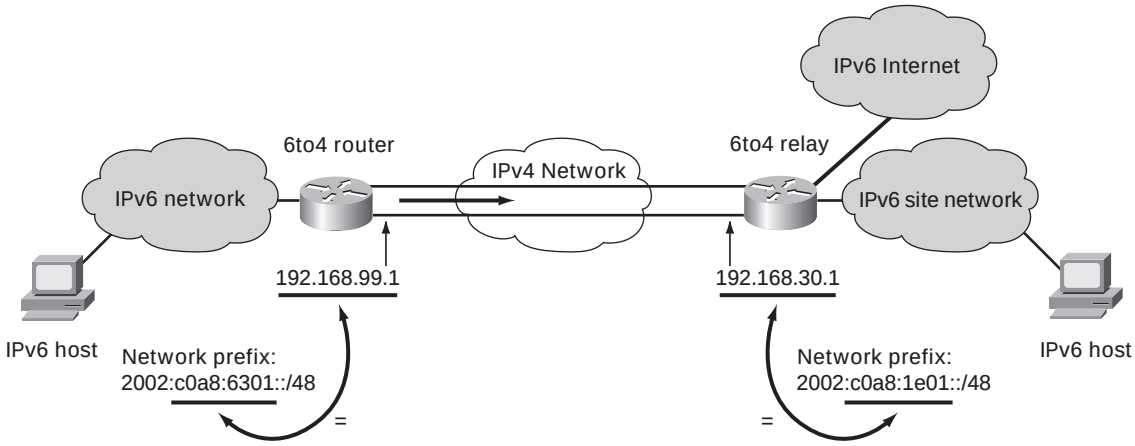

Note: The IPv4 addresses shown in Figure 27 are private and must not be used by a 6to4 Relay for a real Internet connection. Instead a global unicast addresses must be used to forward packets to the Internet.

ISATAP Tunnel

ISATAP is an IPv6 transition mechanism similar to 6to4 tunnels that enables incremental deployment of IPv6 by treating the site IPv4 infrastructure as a nonbroadcast multiaccess (NBMA) link layer.

The ISATAP transition mechanism enables a simple and scalable large-scale incremental deployment of IPv6 for nodes within the existing IPv4 network of a site without incurring aggregation-scaling issues and without the requirement for site-wide deployment of special IPv4 services such as multicast.

ISATAP tunnels are available for use over campus networks or for the transition of local sites. ISATAP supports IPv6 routing within both the site-local and global IPv6 routing domains and automatic IPv6 tunneling across portions of an IPv4 network of a site without any native IPv6 support. ISATAP also supports automatic tunneling within sites that use nonglobally unique IPv4 address assignments combined with network address translation (NAT). All ISATAP nodes are dual stacked.

ISATAP uses a 64-bit network prefix from which the ISATAP addresses are formed. The 64-bit interface identifier is formed by concatenating 0000:5EFE and the IPv4 address of the dual-stack node (192.168.99.1). For example, 3FFE:0B00:0C18:0001:0:5EFE.192.168.99.1 is an ISATAP address. Because ISATAP tunneling typically occurs only within the boundaries of a site, the embedded IPv4 address need not be globally unique. Figure 26 shows an example of the ISATAP tunnel mechanism.

Figure 26: ISATAP Tunnel

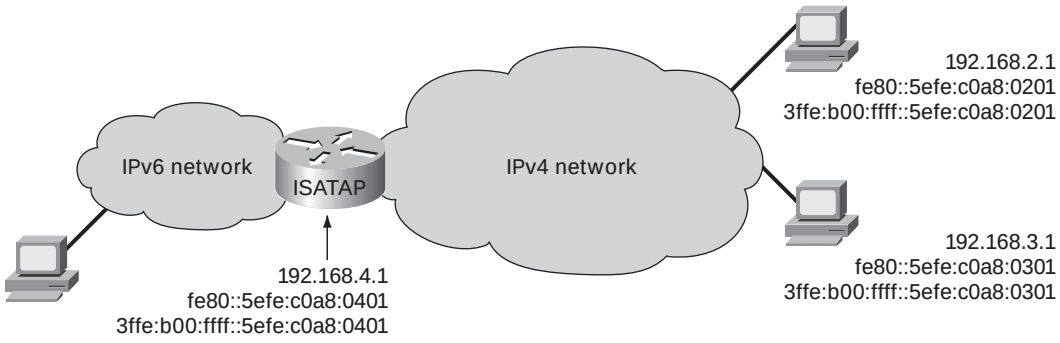

The 6to4 and ISATAP transition mechanisms provide IPv6 connectivity for a node under three typical scenarios: an ISP or an enterprise network provides IPv6 connectivity; the node has access to at least one global IPv4 address; or the enterprise network has deployed an ISATAP router. However, if a node is part of a private network behind a NAT device that is not participating in 6to4, these tunneling mechanisms cannot be used.

For further information on the ISTAP tunnel, refer to the document *Intra-Site Automatic Tunnel Addressing Protocol* (draft-ietf-ngtrans-isatap-04.txt).

Teredo Tunnel

The Teredo (also known as Shipworm) service is a tunnel mechanism that provides IPv6 connectivity to nodes located behind one or more IPv4 NATs by tunneling IPv6 packets over the User Datagram Protocol (UDP) through NAT devices. The Teredo service is defined for the case where the NAT device cannot be upgraded to offer native IPv6 routing or act as a 6to4 router.

Teredo tunnels use Teredo servers and Teredo relays. The Teredo servers are stateless, and manage a small fraction of the traffic between Teredo clients, while the Teredo relays act as IPv6 routers between the Teredo service and the native IPv6 Internet. The Teredo network consists of a set of Teredo clients, servers, and relays. The Teredo network does not require configuration for the Teredo clients. The clients are assigned specially formed IPv6 address prefix, and Teredo servers and relays use globally unique IPv4 addresses.

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 27 shows a sample configuration for IPv6 over dedicated data links.

Figure 27: IPv6 Deployment over Dedicated Data Links

...ne ISP WANs or MA... ...
... run IPv6, for example, over separate ...
configuration has the added benefit for the service

Deploying IPv6 over MPLS Backbones

IPv6 over MPLS backbones enables isolated IPv6 ...
IPv4 core network. This implementation require...
reconfiguration of core routers, because forwar...
providing a very cost-effective strategy for the d...

Additionally, the inherent VPN and traffic engineeri...
networks to be combined into VPNs or extranets ov...

A variety of deployment strategies are available or...

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 strategy, only applicable 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. Another 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.

Deploying 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 using manually configured 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 6to4 addresses or an IPv6 prefix assigned from a distant provider, rather than an IPv6 address supplied by the service provider. Figure 28 shows an example for the deployment of IPv6 using tunnels on the CE routers.

Figure 28: IPv6 Deployment Using Tunnels on the Customer Edge Routers

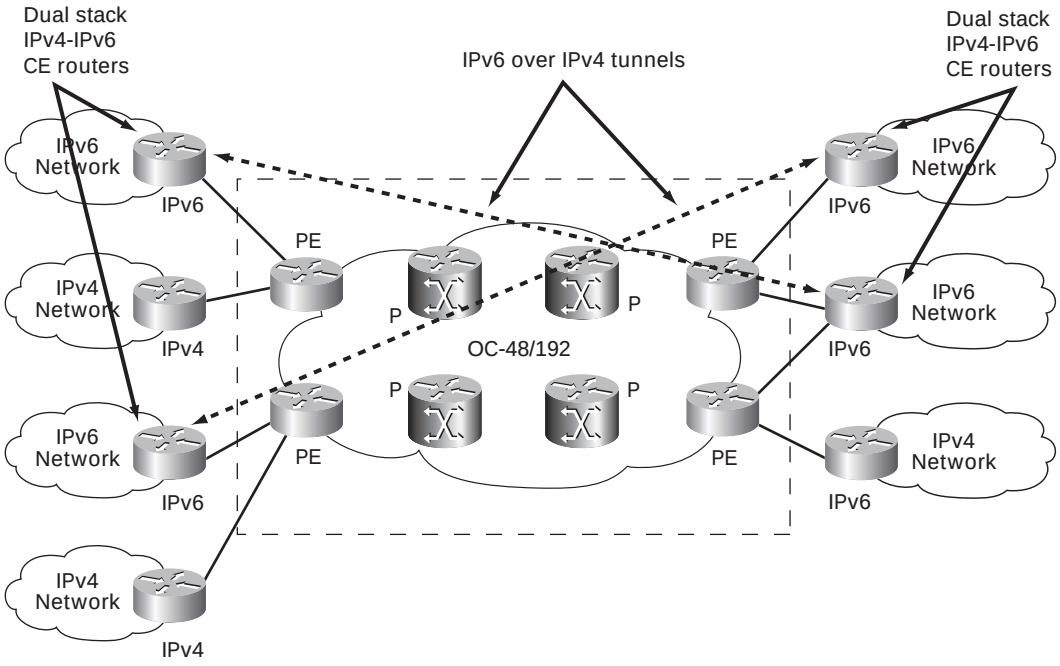

Deploying IPv6 over a Circuit Transport over MPLS

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

Figure 29: IPv6 Deployment over a Circuit Transport over MPLS

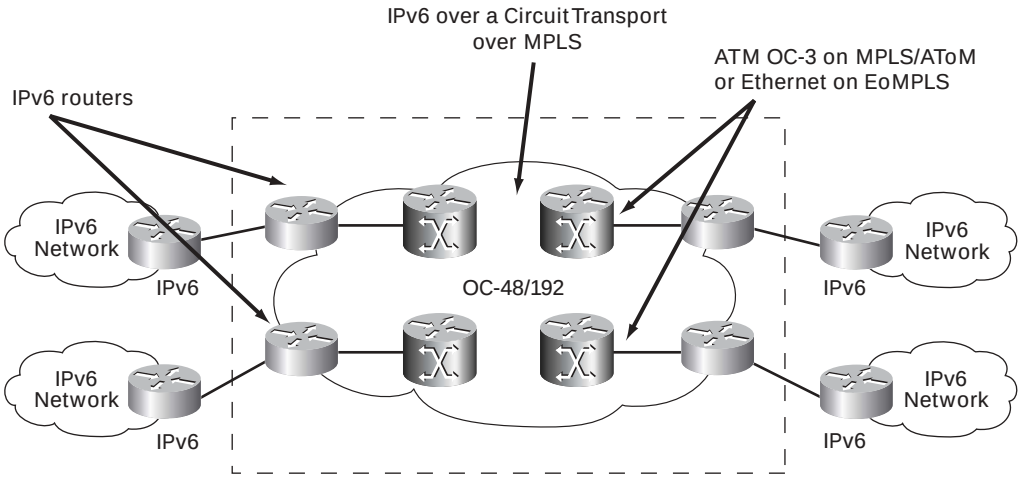

Deploying IPv6 on the Provider Edge Routers

Another 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). The 6PE architecture allows support for IPv6 VPNs. Figure 30 shows an example of IPv6 deployment on the PE routers.

Figure 30: IPv6 Deployment on the Provider Edge Routers

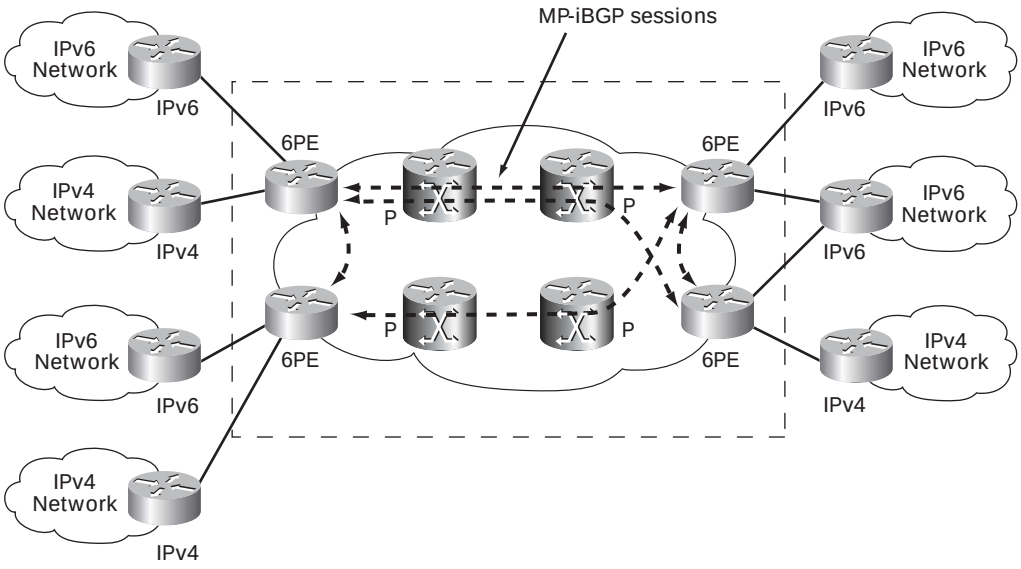

The IPv6 forwarding is done by label switching, eliminating the need for either IPv6 over IPv4 tunnels or for an additional Layer 2 encapsulation, allowing the appearance of a native IPv6 service to be offered across the network.

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 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 OSPF or integrated 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 prefix of the destination through multiprotocol BGP4.

Refer to the Internet-Draft *draft-ietf-ngtrans-bgp-tunnel-04.txt* for further information on 6PE routers.

Protocol 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 requires 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 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 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 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.

Stateless IP/ICMP Translator

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 algorithm 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. This algorithm does not include a mechanism that allows IPv6 hosts to acquire an IPv4 address or route packets to and from that address, but assumes that each IPv6 host has a temporary IPv4 address assigned to it.

For further information on SIIT, refer to RFC 2765, *Stateless IP/ICMP Translation Algorithm (SIIT).*

The following sections describe each protocol translation mechanism in more detail.

Network Address Translation-Protocol Translation

NAT-PT will enable IPv6-only ISPs to interconnect with IPv4 hosts and applications. NAT-PT will be a valuable transition mechanism, when most of the Internet consists of IPv6 network domains.

The NAT-PT translation mechanism (RFC 2766) translates at the network layer between IPv4 and IPv6 addresses and allows native IPv6 hosts and applications to communicate with native IPv4 hosts and applications. An

Application Level Gateway (ALG) translates between the IPv4 and IPv6 DNS requests and responses. Figure 31 shows an example of the use of NAT-PT for deploying IPv6.

Figure 31: Deployment of IPv6 Using NAT-PT

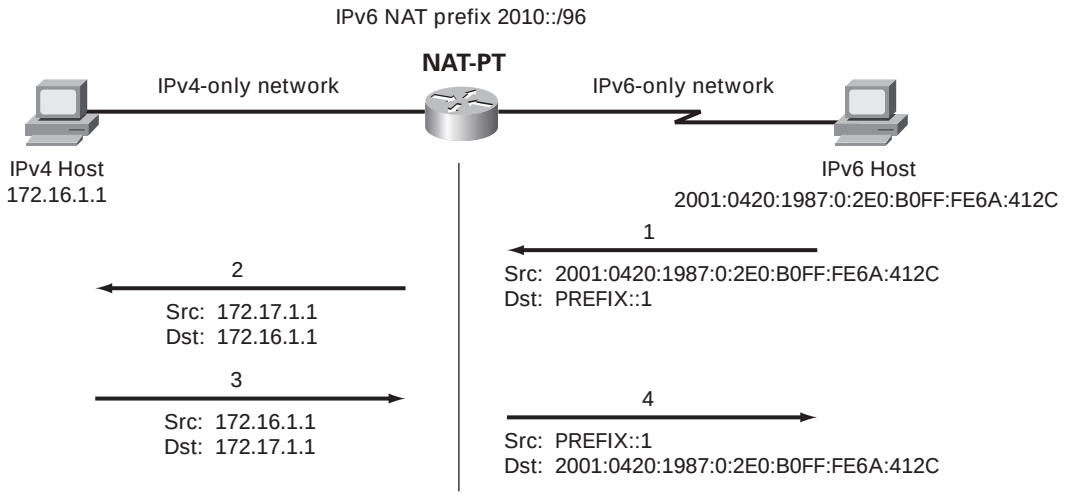

Although familiarity with NAT implementation might encourage people to consider NAT-PT as a protocol translation mechanism, NAT-PT has the same limitations as IPv4 NAT. In addition to the single point of failure, the reduced performance of an ALG, coupled with limitations on the kinds of applications that work, decreases the overall value and utility of the network. NAT-PT also inhibits the ability to deploy security at the IP layer.

For further information on NAT-PT, refer to RFC 2766, *Network Address Translation—Protocol Translation (NAT-PT).*

TCP-UDP Relay

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

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. Implementations of the TCP-UDP relays are freely available from various locations.

For further information on TCP-UDP Relay, refer to RFC3142, *An IPv6-to-IPv4 Transport Relay Translator.*

Bump in the Stack

The BIS mechanism is used 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, *Dual Stack Hosts using the "Bump-In-the-Stack" Technique (BIS)* for further information.

Dual-Stack Transition Mechanism

The DSTM translation mechanism is used 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.

The DSTM 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.

For further information on DSTM, refer to the Internet Draft *draft-ietf-ngtrans-dstm-07.txt*.

SOCKS-Based IPv6/IPv4 Gateway

The SOCKS-based IPv6/IPv4 gateway mechanism is used 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.

Refer to RFC 3089, *A SOCKS-based IPv6/IPv4 Gateway Mechanism* for further information on the gateway and the locations of these sources.

Deployment of Translation Mechanisms

In addition to the strategies for deploying IPv6 within your IPv4 environment, protocol translation mechanisms (for example, NAT-PT or application level gateways) are needed to allow communication between applications using IPv4 and applications using IPv6 (for example, to enable IPv6-only web browsers to communicate with IPv4-only web servers).

These 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 the IPv4 and IPv6 protocols 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.

Refer to RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers*, for general information on the transition mechanisms for IPv6 hosts and routers, and refer to RFC 2185, *Routing Aspects of IPv6 Transition*, for general information on the routing aspects of IPv6 transition.

Chapter 6

IPv6 Network Design Considerations

For IPv6 deployment, 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 to focus on the needs of the applications, rather than complete a full upgrade to a native IPv6 network at this stage. Cisco IPv6 router products offer the features for 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.

Deploying IPv6 in a Service Provider Network Environment

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 personal digital assistants (PDAs), home-area networks (HANs), Internet-connected automobiles, integrated telephony services, and distributed gaming.

We recommend that you should look at the deployment of IPv6 in three key phases:

- Providing an IPv6 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 services. This approach allows an evaluation of the 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 within the core infrastructure itself: 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.
- Interconnecting with other IPv6 service providers: 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.

Deploying IPv6 in an Enterprise Network Environment

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 network address translation (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 current business is not disrupted.

IPv6 Support from Cisco

Cisco Systems is one of the founding members of the IPv6 Forum (www.ipv6forum.com). Cisco has taken a leading role in the definition and implementation of the IPv6 architecture within the IETF and continues to lead the industry efforts for standardization. Cisco employees Steve Deering and Tony Hain are the cochairs of the IETF IPv6 working group and the next-generation transition (Ngtrans) working group, respectively. Many of the IPv6 standards are already published by the IETF, although enhancements are still being made.

Cisco has decided to implement IPv6 features in the Cisco IOS Software and its products in three phases. You will find IPv6 implementation details, including the Cisco roadmap for IPv6 and the Statement of Direction, at **www.cisco.com/ipv6**.

The early field trial version of the IPv6 for Cisco IOS Software has been available freely for more than three years. Cisco IOS IPv6 software has been extensively deployed in the 6BONE network (www.6BONE.net) for test purposes over several years. Also, the Cisco 6BONE router has been operational as a major 6BONE hub for more than 5 years with more than 70 tunnels to other companies.

IPv6 for Cisco IOS Software is available for all Cisco router platforms, from the low-end Cisco 800 series routers to high-end platforms that include the Cisco 12000 Internet routers. Since Cisco IOS Software Release 12.2(2)T, Cisco officially provides worldwide technical support. Additional information on Cisco IOS IPv6 is available at www.cisco.com/ipv6.

Cisco plans to release a set of solutions documents covering the deployment of IPv6. This section will be updated as these documents become available. Also, if you need IPv6 for Cisco IOS Software configuration information for use with Cisco equipment, refer to the Cisco documentation on cisco.com.

Appendix A

Bibliography and Reference Resources

This section provides information on documents, books, and RFCs that served as resources for this document and additional resources for IPv6 that you might find useful, including the technical documents from Cisco and the Cisco reference page on IPv6. It also includes information about relevant RFCs and drafts.

Cisco Statement of Direction for IPv6

http://www.cisco.com/warp/public/732/tech/ipv6/

Cisco Technical Documentation

IPv6 for Cisco IOS Software feature documentation (Cisco.com) for IPv6 overview, configuration, and command reference information:
http://www.cisco.com/univercd/cc/td/doc/product/software/

IPv6 integrated solutions documents (ISDs) for detailed information about the various IPv6 transition mechanisms:
http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/ipv6_sol/index.htm

Books

Marcus Gonglaves and Kitty Niles, *IPv6 Networks*, McGraw Hill, New York, NY 1998

Microsoft Windows 2000 Server, *Introduction to IP Version 6*, white paper

Cisco Training Guide: *Implementing IPv6 Networks*

White Papers and Other Documents

Alcatel Technical Paper: *The Move to IPv6*

Cisco: Engineering Training for IPv6

Cisco: The Internet Protocol Journal

Glocom Platform Tech Reviews: Nobuo Ikeda and Hajime Yamada, *Is IPv6 Necessary?*

Hirahara.ourfamily.com: *Technical – IPv6 and Subnetting IPv4*

IP Infusion White Paper: *IPv6 Network Processing*

IPv6 Forum IPv6 Tutorial: Jordi Palet, *ICMPv6 & Neighbor Discovery*

Nortel Networks White Paper: *Building the Foundation of the Multimedia Wireless Internet _The migration to IPv6 and MPLS for UMTS*

The O'Relly Network: *Introduction to IPv6*

RFCs and Drafts

Rationale and Case for IPv6

The Recommendation for the IP Next-Generation Protocol: RFC 1752

The Case for IPv6: draft-iab-case-for-ipv6-06.txt

Classless interdomain routing (CIDR): RFC 1519

The H Ratio for Address Assignment Efficiency: RFC 1715

Architectural Implications of NAT: RFC 2993

Internet Transparency: RFC 2775

Protocols

Internet Protocol Version 6 (IPv6) Specification: RFC 2460

Path MTU Discovery for IP Version 6: RFC 1981

IP Version 6 Management Information Base for TCP: RFC 2452

Internet Control Message Protocol (ICMPv6) for the IPv6 Specification: RFC 2463

IPv6 Address Types

IP Version 6 Addressing Architecture: RFC 2373

An IPv6 Aggregatable Global Unicast Address Format: RFC 2374

Format for Literal IPv6 Addresses in URL's: RFC 2732

IPv6 Multicast Address Assignments: RFC 2375

IPv6 Autoconfiguration and Renumbering

Neighbor Discovery for IP Version 6 (IPv6): RFC 2461

IPv6 Stateless Address Autoconfiguration: RFC 2462

Router Renumbering for IPv6: RFC 2894

IPv6 Link Layer

Transmission of IPv6 Packets over Ethernet Networks: RFC 2464

Transmission of IPv6 Packets over FDDI Networks: RFC 2467

Transmission of IPv6 Packets over Token Ring Networks: RFC 2470

Transmission of IPv6 over IPv4 Domains without Explicit Tunnels: RFC 2529

Transmission of IPv6 Packets over ARCnet Networks: RFC 2497

IP Version 6 over PPP: RFC 2472

IPv6 over Non-broadcast Multiple Access (NBMA) Networks: RFC 2491

IPv6 over ATM Networks: RFC 2492

Transmission of IPv6 Packets over Frame Relay Networks Specification: RFC 2590

IPv6 Routing Protocol Support

RIPng for IPv6: RFC 2080

OSPF for IPv6: RFC 2740

Routing IPv6 with IS-IS: draft-ietf-isis-ipv6-02.txt

Multiprotocol Extensions for BGP-4: RFC 2858

Use of BGP-4 Multiprotocol Extensions for IPv6 Interdomain Routing: RFC 2545

IPv6 Integration and Transition Mechanisms

Transmission of IPv6 over IPv4 Domains Without Explicit Tunnels (6over4): RFC 2529

Connection of IPv6 Domains via IPv4 Clouds (6to4): draft-ietf-ngtrans-6to4-07.txt

Network Address Translation-Protocol Translation (NAT-PT): RFC 2766

Transition Mechanisms for IPv6 Hosts and Routers: RFC 2893

Generic Packet Tunneling in IPv6: RFC 2473

Connection of IPv6 Domains via IPv4 Clouds: RFC 3056

On overview of the introduction of IPv6 in the Internet: draft-ietf-ngtrans-introduction-to-ipv6-transition-04.txt

IPv6 Tunnel Broker: draft-ietf-ngtrans-broker-04.txt

IPv6 Deployment

6BONE pTLA and pNLA Formats (pTLA): RFC 2921

6BONE Backbone Routing Guidelines: RFC 2772

Other Web References

IETF: **http://www.ietf.org/html.charters/ipv6-charter.html**

6BONE: **http://www.6bone.net**

6TAP exchange: **http://www.6tap.net**

Freenet6: **http://www.freenet6.net**

IPv6 Forum: **http://www.ipv6forum.com**

http://playground.sun.com/ipv6

http://www.hs247.com

IPv6 Host Configuration

Solaris IPv6

Solaris IPv6: **http://www.sun.com/software/solaris/ipv6/**

Microsoft IPv6

www.microsoft.com/ipv6

FreeBSD IPv6

Kame: **http://www.kame.net**

IPv6 Address Allocation

Proposed TLA and NLA Assignment Rules: RFC 2450

IPv6 Address Allocation and Assignment Global Policy:

http://www.ripe.net/ripencc/mem-services/registration/ipv6/global-ipv6-assign-2002-04-25.html

Efficient method for address plan: draft-ietf-ipngwg-ipaddressalloc-01.txt

Allocation policy: *ftp://ftp.ripe.net/ripe/docs/ripe-196.txt*

IPv6 Address Registries

RIPE NCC (Europe and Middle East): **http://www.ripe.net**

ARIN (Americas): **http://www.arin.net**

APNIC (Asia): **http://www.apnic.net**

Current Sub-TLA Allocations

http://www.ripe.net/ripencc/mem-services/registration/ipv6/ipv6allocs.html

Appendix B

Glossary

6BONE—An IPv6 testbed that consists of IPv6 networks. The 6BONE is a worldwide informal collaborative project, informally operated with oversight from the IPv6 Working Group of the IETF. Though it started as a virtual network using IPv6 tunnels or encapsulation over IPv4 networks, it is slowly migrating to native links for IPv6 transport.

6to4 tunnel—An IPv6 automatic tunneling technique where the tunnel endpoint is determined by the globally unique IPv4 address embedded in a 6to4 address. A 6to4 address is a combination of the prefix 2002::/16 and a globally unique 32-bit IPv4 address. (IPv4-compatible addresses are not used in 6to4 tunneling.)

6to4 relay—A 6to4 border router that offers traffic forwarding to the IPv6 Internet for other 6to4 border routers. A 6to4 relay forwards packets to any destination that has a 2002::/16 prefix.

A6 record—A Domain Name System (DNS) record that stores IPv6 numbers used to represent a 128-bit IPv6 address. When an IPv6-aware application wants to look up the name of an IPv6 server, it could request an A6 record from the DNS server. The A6 record is not the preferred record for name resolution with IPv6, because it has been set aside for experimental purpose.

AAAA—A Domain Name System (DNS) record that stores IPv6 numbers used to represent a 128-bit IPv6 address. The AAAA records are used to resolve host names. This operation is similar to the process where applications request the A record in IPv4. The AAAA record is the preferred record for name resolution with IPv6.

anycast address—An identifier for a set of interfaces that typically belong to different nodes. A packet sent to an anycast address is delivered to the closest interface—as defined by the routing protocols in use—identified by the anycast address. See also global unicast address, IPv6 multicast address, link-local address, site-local address, and solicited-node multicast address.

APNIC—Asia Pacific Network Information Centre. The regional Internet registry (RIR) responsible for assigning IP addresses to the countries in the Asia Pacific region.

ARIN—The American Registry for Internet Numbers. The regional Internet registry (RIR) responsible for assigning IP addresses to the countries in the North and South American regions.

automatic IPv6 tunnel—An IPv6 tunneling technique (to be deprecated soon), where the tunnel source and tunnel destination are automatically determined by using the IPv4 address in the low-order 32 bits of IPv6 addresses using the specially assigned 6to4 IPv6 prefix 2002::/16. The host or router at each end of an IPv6 automatic tunnel must support both the IPv4 and IPv6 protocol stacks. Automatic tunnels can be configured between border routers or between a border router and a host. See also IPv4-compatible IPv6 address and manually configured IPv6 tunnel.

BIS—Bump-in-the-Stack. Translation mechanism used for communication between IPv4 applications on an IPv4-only host and IPv6-only hosts. It uses a snooping module and an automatically allocated IPv4 address from a pool and works like a self-translator.

CE router—Customer edge router is a router that is part of a customer MPLS network and interfaces to a provider edge (PE) router.

DSTM—Dual-Stack Transition Mechanism. A translation mechanism for dual stack hosts in an IPv6 domain that do not have an IPv4 routing infrastructure, but need to communicate with IPv4 systems or allow IPv4 applications to run on top of their IPv6 protocol stack. DSTM operation is based on the use of IPv4-over-IPv6 tunnels and the temporal allocation of a global IPv4 address to hosts requiring such communication.

global unicast address—An IPv6 unicast address similar to a typical IPv4 address. It enables aggregation of routing prefixes in order to limit the number of routing table entries in the global routing table. See also anycast address, IPv6 multicast address, link-local address, and site-local address.

GRE tunnel—A manually configured tunnel, particularly suitable for use with the IS-IS protocol. The GRE tunnel is not tied to a specific passenger or transport protocol, but in this case carries IPv6 traffic as the passenger protocol over GRE as the carrier protocol. Generic routing encapsulation is a network protocol that allows any arbitrary passenger protocol to be sent over any carrier protocol.

IANA—Internet Assigned Numbers Authority. Responsible for assigning unique parameter values to Internet protocols.

IETF—Internet Engineering Task Force. International group of network researchers, designers, operators, and vendors responsible for the design and engineering of TCP/IP and the global Internet.

IPv4-compatible IPv6 address—An IPv6 unicast address that has zeros in the high-order 96 bits of the address and an IPv4 address in the low-order 32 bits of the address. The format of an IPv4-compatible IPv6 address is 0:0:0:0:0:0:A.B.C.D or ::A.B.C.D, where A.B.C.D represents the IPv4 address. The entire 128-bit IPv4-compatible IPv6 address is used as the IPv6 address of a node, and the IPv4 address embedded in low-order 32-bits is used as the IPv4 address of the node. IPv4-compatible IPv6 addresses are assigned to nodes that support both the IPv4 and IPv6 protocol stacks, and are used in automatic tunneling. See also anycast address, automatic IPv6 tunnel, IPv6 multicast address, link-local address, and site-local address.

IPv6 multicast address—An IPv6 address with a prefix of FF00::/8. An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes. A packet sent to a multicast address is delivered to all interfaces identified by the multicast address. See also global unicast address, anycast address, link-local address, site-local address, and solicited-node multicast address.

ISATAP—A transition mechanism used for deploying IPv6, particularly in the campus network environment. ISATAP enables incremental deployment of IPv6 by treating the IPv4 infrastructure of the site as a nonbroadcast multiaccess (NBMA) link layer.

link—Links are networks arbitrarily segmented by a network administrator in order to provide a multilevel, hierarchical routing structure while shielding the subnetwork from the addressing complexity of attached networks. Similar to a subnetwork in IPv4. A subnetwork prefix is associated with one link, but multiple subnetwork prefixes may be assigned to the same link.

link-local address—An IPv6 unicast address that has a scope limited to the local link (local network). Link-local addresses are automatically configured on all IPv6 interfaces by using a specific prefix for link-local addresses (FE80::/10) and adding the interface ID in the modified EUI-64 format. Link-local addresses are used by the neighbor discovery protocol and the router discovery protocol. They are also used by many routing protocols. Link-local addresses can serve as a way to connect devices on the same local network without needing global addresses. See also global unicast address, anycast address, IPv6 multicast address, site-local address, and solicited-node multicast address.

manually configured IPv6 tunnel—An IPv6 tunneling technique where a manually configured IPv6 address is configured on a tunnel interface and manually configured IPv4 addresses are assigned to the tunnel source and the tunnel destination. The host or router at each end of a configured tunnel must support both the IPv4 and IPv6 protocol stacks. Manually configured tunnels can be configured between border routers or between a border router and a host. See also automatic IPv6 tunnel.

MPLS—Multiprotocol Label Switching. A switching technique that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on preestablished IP routing information.

NAT-PT—Network address translation-protocol translation. A translation mechanism that translates at the network layer between IPv4 and IPv6 addresses and allows native IPv6 hosts and applications to communicate with native IPv4 hosts and applications. An Application Level Gateway (ALG) translates between the IPv4 and IPv6 DNS requests and responses.

NLA—Next Level Aggregator as originally described in the IPv6 network hierarchy. An IPv6 service provider below the Top Level Aggregator service provider. The NLA field of 24 bits could support up to 16 million Site Level Aggregators. The NLA is no longer part of the IPv6 RFCs. See TLA and SLA.

pTLA—pseudo Top Level Aggregator. As originally described in the IPv6 network hierarchy, used with the 6BONE network, a testbed network of IPv6 networks. See TLA.

SIIT—Stateless IP/ICMP Translator. An algorithm that 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.

RIPE NCC— Reseaux IP Europeens_Network Coordination Center (RIPE NCC). The regional Internet registry (RIR) responsible for assigning IP addresses to the countries in Europe and the Middle East

site-local address—Address that is useful only in the context of the site and is similar to the private addresses in IPv4. Its scope is limited to this context. When configured, a site-local address uses a specific prefix (FEC0::/10) and concatenates the subnet ID as a 16-bit field and then the interface ID in the modified EUI-64 format. See also anycast address, global unicast address, IPv6 multicast address, link-local address, and solicited-node multicast address.

SLA—Site Level Aggregator. As originally described in the IPv6 network hierarchy, an IPv6 service provider below the Next Level Aggregator service provider. The SLA field of 16 bits could support up to 65,535 subnets within a site. See NLA and TLA.

solicited-node multicast address—An IPv6 address that has the prefix FF02:0:0:0:0:1:FF00:0000/104 concatenated with the 24 low-order bits of a corresponding IPv6 unicast or anycast address. The solicited-node multicast address is a multicast group that corresponds to an IPv6 unicast or anycast address. Solicited-node multicast addresses are used in neighbor solicitation messages. See also anycast address, global unicast address, IPv6 multicast address, link-local address, and site-local address.

TCP-UDP Relay—Translation mechanism similar to NAT-PT. It requires a dedicated server and DNS; it translates at the transport layer rather than the network layer, with the DNS again providing the mapping between IPv4 and IPv6 addresses.

Teredo tunnel—The Teredo (also known as Shipworm) service is a tunnel mechanism that provides IPv6 connectivity to nodes located behind one or more IPv4 NATs by tunneling IPv6 packets over UDP through NATs.

TLA—Top Level Aggregator. As originally described in the IPv6 network hierarchy, a service provider at the top of an IPv6 network hierarchy. The TLA is responsible for maintaining the upper levels of the IPv6 network routing hierarchy. The TLA field of 13 bits supports up to 8192 TLAs. The TLA is no longer part of the IPv6 RFCs.

Appendix C

Review Questions

The following review questions will help you assess how well you have learned the technical information provided in the ABCs of IP Version 6 document. Appendix D of this document provides the answers to these review questions.

1. Why is Network Address Translation (NAT) *not* an ideal solution to solve the IP address exhaustion problem?
   a. NAT breaks the end-to-end security model of IP.
   b. NAT protects network devices and data from possible external intruders.
   c. NAT conserves global IP addresses by using private address space within a network.
   d. NAT dynamically allocates global addresses to internal network devices to allow communication with the Internet.
2. How many bits are supported in the IPv6 address scheme?
   a. 32 bits
   b. 64 bits
   c. 96 bits
   d. 128 bits
3. What IPv4 header fields have been removed from the IPv6 header?
   a. Version, fragmentation fields, Header Checksum, and Padding.
   b. Version, Header Length, fragmentation fields, and Header Checksum.
   c. Header Length, fragmentation fields, Header Checksum, and Flow Label.
   d. Header Length, fragmentation fields, Header Checksum, and Padding.
4. Which one of the following statements is true about IPv6 routing protocols?
   a. IPv6 RIP protocol is named RIP-2.
   b. IPv6 IS-IS protocol is currently an IETF standard.
   c. IPv6 OSPF protocol is a proposed IETF standard.
   d. IPv6 OSPF protocol is currently not an IETF standard.
5. Which one of the following statements is *not* true about routing in IPv6?
   a. Both IPv4 and IPv6 support the same routing protocols.
   b. IPv6 RIP updates are sent to the all-rip-routers multicast group address FF02::9.
   c. IPv6 does not use the longest-prefix match for routing algorithm.
   d. BGP-4+ NEXT_HOP and NLRI are expressed as IPv6 addresses and prefix.

6. Which one of the following statements is true about IPv6 autoconfiguration?

   a. IPv6 mandates the use of DHCP servers in all IPv6 networks.

   b. Lack of collisions in IPv6 networks makes autoconfiguration possible.

   c. Larger address space enables IPv6 hosts to autoconfigure themselves.

   d. IPv6 devices come with preset IPv6 addresses and need no configuration.

7. Which one of the following statements is true about broadcasts and multicasts in IPv6?

   a. Broadcast is not used in IPv6.

   b. Broadcast can completely hang up an IPv6 network.

   c. Broadcast is the basic mechanism used for various operations in IPv6.

   d. Broadcast can interrupt all the nodes on a LAN network.

8. Which one of the following statements is not a feature of IPv6?

   a. Autoconfiguration.

   b. Automatic QoS support.

   c. Easier renumbering.

   d. Larger address space.

9. Which one of the following fields is a new field in the IPv6 header?

   a. Destination Address.

   b. Source Address.

   c. Flow Label.

   d. Version.

10. Which one of the following statements is not true about IPv6 feature set?

    a. UDP checksum is mandatory in IPv6.

    b. Any IPv6 node can use the mobility feature.

    c. With built-in IPSec support, IPv6 supports end-to-end to security.

    d. Multihoming implementation is more difficult in IPv6 than in IPv4.

11. Which one of the following statements is the correct IPv6 link-local address prefix?

    a. 2001:1

    b. 2002:1

    c. FE80::/10

    d. FEC0::/10

12. The IPv6 link-local addresses:
    a. Consist of 96 zeros at the left-most fields of the address.
    b. Consist of the link-local prefix, 16-bit subnet ID field, and the interface ID in EUI-64 format.
    c. Serve as a way to connect devices between two networks without needing global addresses.
    d. Are automatically configured on all interfaces using the link-local prefix and the interface ID in the EUI-64 format.
13. Which one of the following statements is not true about IPv6 header fields?
    a. IPv6 Next Header field is similar to the Protocol field in the IPv4 header.
    b. IPv6 Traffic Class field is similar to the Type of Service field in the IPv4 header.
    c. IPv6 Hop Limit field makes the computing of checksum very efficient.
    d. The value in the Next Header field determines the type of information following the basic IPv6 header.
14. IPv6 handles extension headers more efficiently by:
    a. Looking at every individual header field to allow accurate processing.
    b. Daisy-chaining the extension header fields to allow faster processing.
    c. Daisy-chaining the routing header fields to allow faster processing.
    d. Ignoring the extension header field, if Routing Header is not present.
15. Which one of the following addresses is a valid IPv6 address?
    a. 2001:1:0:4F3A:206:AE14
    b. 2001:1:0:4F3A:0:206:AE14
    c. 2001:1:0:4F3A::206:AE14
    d. 2001:1::4F3A:206::AE14
16. Which of the following is not a required address for an IPv6 node?
    a. All-nodes multicast address.
    b. Link-local address for each interface.
    c. Specific multicast address for routing protocols.
    d. Solicited-node multicast address for each of the assigned unicast and anycast addresses.
17. Which one of the following statements is not true about 6BONE network?
    a. 6BONE is a testbed network of IPv6 networks.
    b. 6BONE topology consists of a hierarchy of providers.
    c. 6BONE network pTLA prefixes are in the 2001::/16 range.
    d. 6BONE allows only registry-assigned and 6BONE addresses.

18. The neighbor discovery process helps to determine the:
    a. Link-layer address of a neighbor on the same link.
    b. Multicast address of a neighbor on a different link.
    c. IPv6 address of the nearest router on a different link
    d. IPv6 address of a neighbor on the same link.
19. IPv6 neighbor solicitation:
    a. Could be used to verify the reachability of a neighbor.
    b. Is sent at boot time to promptly receive router advertisements.
    c. Is similar to Reverse Address Resolution Protocol (RARP) used in IPv4.
    d. Is periodically sent as an advertisement to the all-nodes multicast address.
20. What is the minimum maximum transmission unit (MTU) supported in IPv6?
    a. 68 octets
    b. 576 octets
    c. 1280 octets
    d. 1500 octets
21. Which one of the following statements best describes correct IPv6 operation?
    a. Fragmentation is only done by the originating host.
    b. Fragmentation is only done by the originating router.
    c. Fragmentation is handled exactly as in IPv4.
    d. ICMPv6 messages are not used in the fragmentation process.
22. Which one of the following statements best describes DHCPv6?
    a. DHCPv6 can be used only in stateful autoconfiguration.
    b. DHCPv6 can be used only in stateless/serverless autoconfiguration.
    c. DHCPv6 cannot be used with automatic domain name registration.
    d. DHCPv6 can also be concurrently used with stateless autoconfiguration.
23. Which IPv6 DNS record is recommended for Hostname-to-IP address translation for DNS?
    a. A
    b. AAAA
    c. A6
    d. PTR

24. Which of the following statements is not true about tunnel endpoints in a configured tunnel for IPv6?

    a. Tunnel endpoints must be configured with dual stack.

    b. Both IPv4 and IPv6 addresses are configured on tunnel endpoints.

    c. Configuration of the tunnel endpoints changes dynamically.

    d. Tunnel endpoints could be an edge router and an end system.

25. Which one of the following statements describes the major difference between the manually configured tunnel and the IPv4-compatible tunnel?

    a. Manually configured tunnel is a static tunnel, but the IPv4-compatible tunnel is an automatic tunnel.

    b. The manually configured tunnel does not scale at all, but the IPv4-compatible tunnel scales significantly.

    c. The manually configured tunnel helps to conserve IPv6 addresses, but the IPv4-compatible tunnel helps to conserve IPv4 addresses.

    d. Although the manually configured tunnel uses IPv4 addresses, the IPv4-compatible tunnel uses only IPv6 addresses.

26. Which one of the following statements describes the major difference between the IPv4-compatible tunnel and the 6to4 tunnel?

    a. The IPv4-compatible tunnel is a static tunnel, but the 6to4 tunnel is an automatic tunnel.

    b. The IPv4-compatible tunnel is typically used only between two IPv6 domains, but the 6to4 tunnel is used to connect multiple IPv6 domains.

    c. The deployment of the IPv4-compatible tunnel requires a special code on the edge routers, but the 6to4 tunnel does not require any special code.

    d. For the IPv4-compatible tunnel, the ISP assigns only IPv4 addresses for each domain, but for the 6to4 tunnel, the ISP assigns only IPv6 addresses for each domain.

27. Which one of the following statements best describes the operation of a 6to4 relay?

    a. A 6to4 relay is not a router, but is a gateway to the IPv6 internet.

    b. A 6to4 relay is used to forward packets only to other 6to4 routers.

    c. A 6to4 relay is used to forward packets only to the IPv6 internet.

    d. A 6to4 relay is used to forward packets to other 6to4 routers and to the IPv6 internet.

Appendix D

Answers to the Review Questions

1a, 2d, 3d, 4c, 5c, 6c, 7a, 8b, 9c, 10d, 11c, 12d, 13c, 14b, 15c, 16c, 17c, 18a, 19a, 20c, 21a, 22d, 23b, 24c, 25a, 26b, 27d

Please direct comments to abcios@cisco.com

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux
Cedex 9
France
http://www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
http://www.cisco.com
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at http://www.cisco.com/go/offices.

Argentina · Australia · Austria · Belgium · Brazil · Canada · Chile · China · Colombia · Costa Rica · Croatia · Czech Republic · Denmark · Dubai, UAE
Finland · France · Germany · Greece · Hong Kong · Hungary · India · Indonesia · Ireland · Israel · Italy · Japan · Korea · Luxembourg · Malaysia
Mexico · The Netherlands · New Zealand · Norway · Peru · Philippines · Poland · Portugal · Puerto Rico · Romania · Russia · Saudi Arabia · Singapore
Slovakia · Slovenia · South Africa · Spain · Sweden · Switzerland · Taiwan · Thailand · Turkey · Ukraine · United Kingdom · United States · Venezuela

Copyright ® 2002, Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, Cisco IOS, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries. All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0206R)