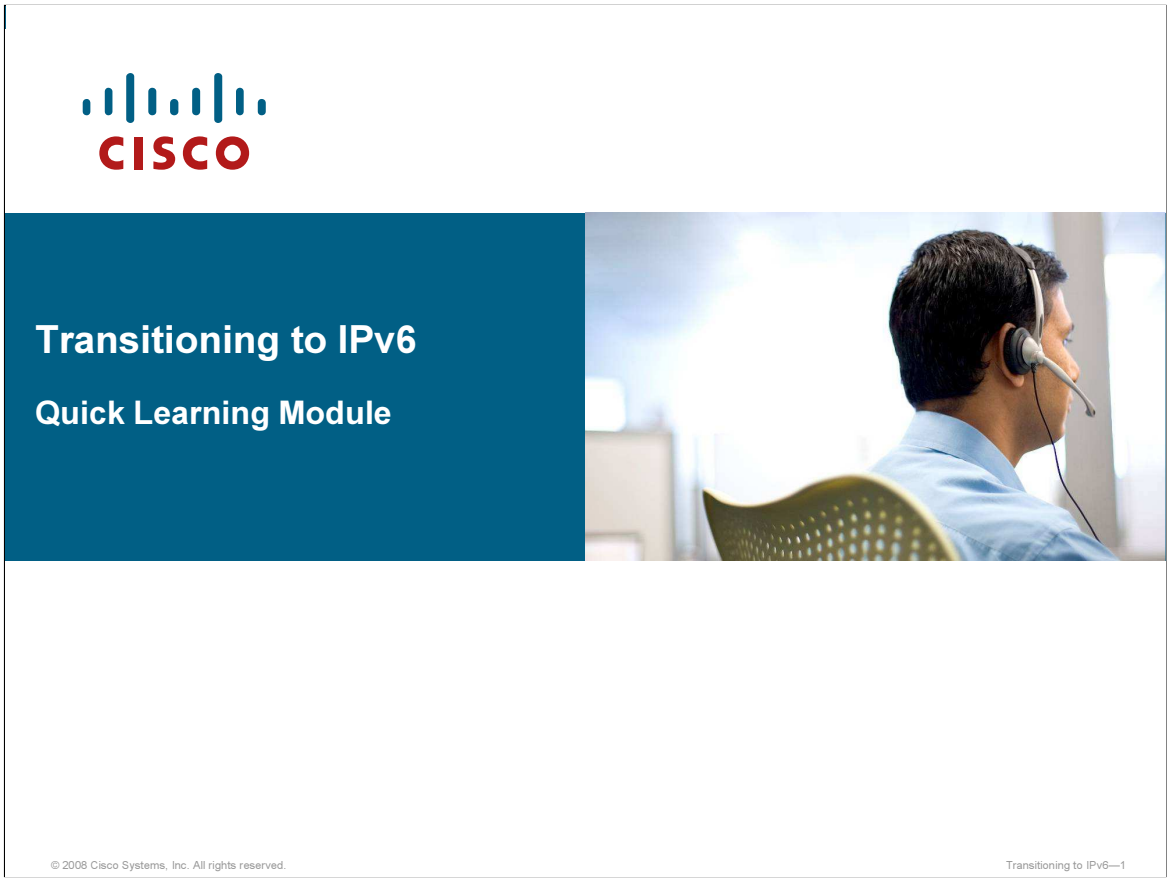

Welcome to the “Transitioning to IPv6” Quick Learning Module. Quick Learning Modules are byte-sized chunks of learning that explain the purpose, operation, and configuration of Cisco technology features. This content is designed to supplement your learning experience and exam preparation. It is not meant to replace any in-classroom training or online learning modules. We hope that you enjoy this training.

Objectives

Upon completing this module, you should be able to explain the format of IP version 6 (IPv6) addresses and the components that are required to run IPv6, explain the impact of IPv6 on network routing, and configure basic IPv6 parameters. These abilities include being able to meet these objectives:

- Explain the need for IPv6
- Describe the format of the IPv6 address
- Explain the methods that are used to assign an IPv6 address
- Explain how IPv6 affects common routing protocols and the necessary modifications that you need to make to these protocols
- Explain transition strategies for implementing IPv6
- Configure IPv6 with Routing Information Protocol next generation (RIPng) through an IP version 4 (IPv4) network

Upon completing this module, you should be able to explain the format of IPv6 addresses and the components that are required to run IPv6, explain the impact of IPv6 on network routing, and configure basic IPv6 parameters. These abilities include being able to meet these objectives:

•Explain the need for IPv6

•Describe the format of the IPv6 address

•Explain the methods that are used to assign an IPv6 address

•Explain how IPv6 affects common routing protocols and the necessary modifications you need to make to these protocols

•Explain transition strategies for implementing IPv6

•Configure IPv6 with Routing Information Protocol next generation (RIPng) through an IP version 4 (IPv4) network

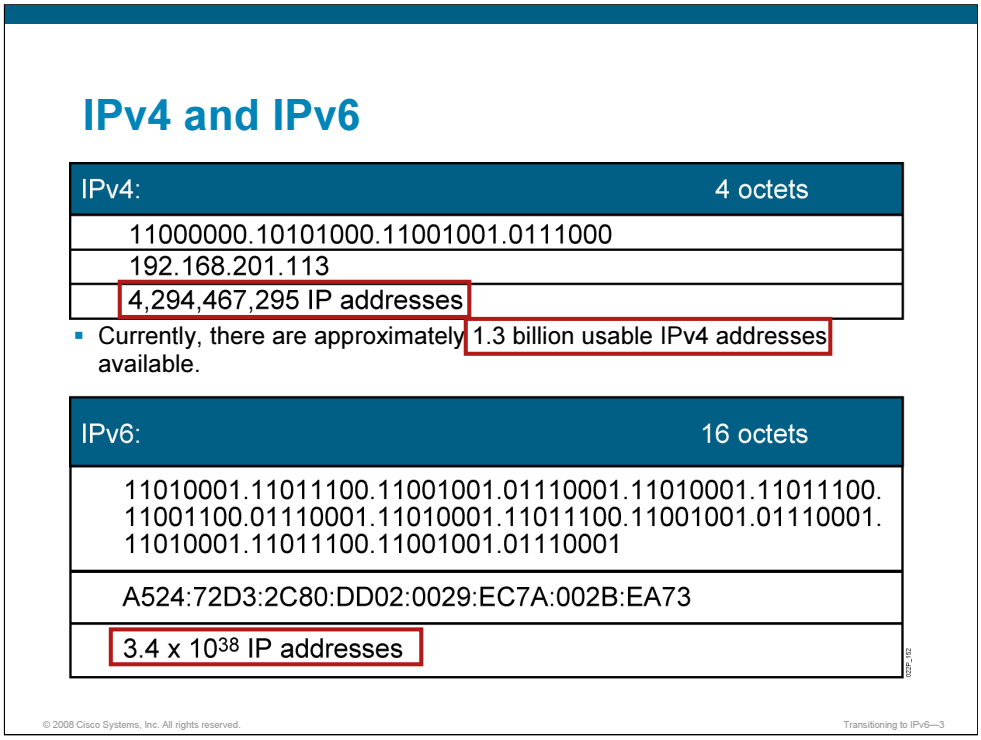

The IPv4 address space provides approximately 4.3 billion addresses. Of that address space, approximately 3.7 billion addresses are actually assignable; the other addresses are reserved for special purposes such as multicasting, private address space, loopback testing, and research. Based on figures as of January 1, 2007, about 2.4 billion of these available addresses are currently assigned to either end users or Internet service providers (ISPs). That leaves roughly 1.3 billion addresses still available from the IPv4 address space.

An IPv6 address is a 128-bit binary value, which can be displayed as 32 hexadecimal digits, as shown in the figure. It provides 3.4 times 10 to the 38th IP addresses. This version of IP addressing should provide sufficient addresses for future Internet growth needs.

In addition to its technical and business potential, IPv6 offers a virtually unlimited supply of IP addresses. Because of its generous 128-bit address space, IPv6 generates a virtually unlimited stock of addresses—enough to allocate more than the entire IPv4 Internet address space to everyone on the planet.

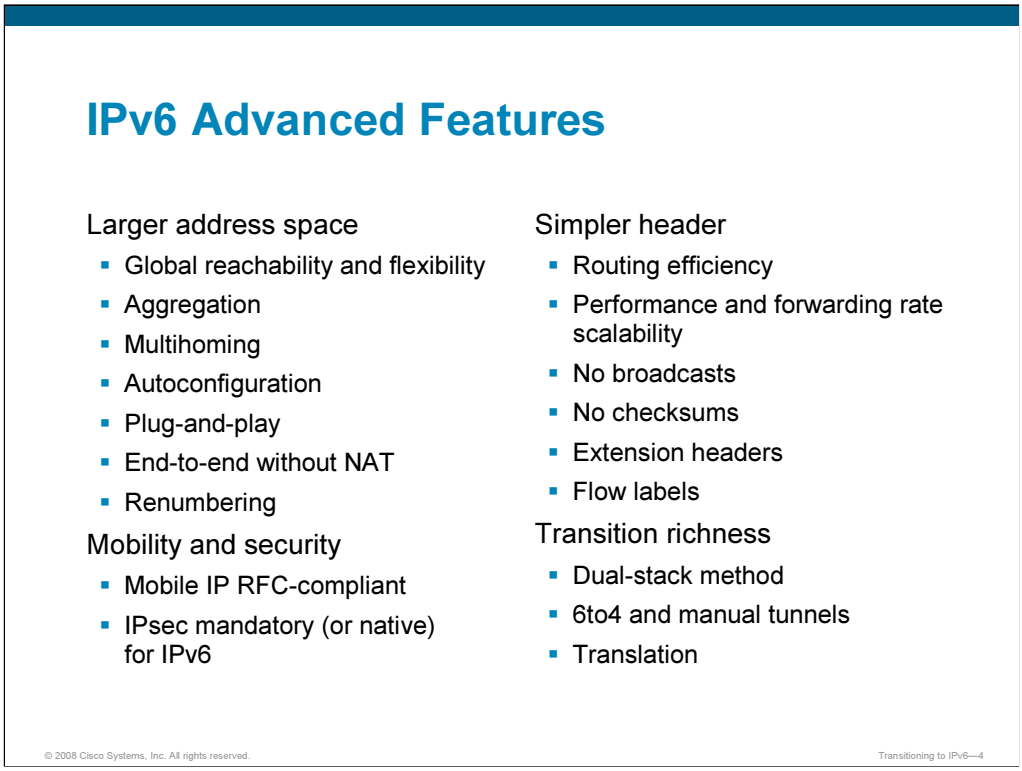

IPv6 is a powerful enhancement to IPv4. Several features in IPv6 offer functional improvements, including the following:

- **Larger address space, including:**
  - Improved global reachability and flexibility
  - The aggregation of prefixes that are announced in routing tables
  - Multihoming to several ISPs
  - Autoconfiguration that can include data link layer addresses in the address space
  - Plug-and-play options
  - Public-to-private readdressing end-to-end without Network Address Translation (NAT)
  - Simplified mechanisms for address renumbering and modification
- **A simpler header, including**:
  - Better routing efficiency for performance and forwarding-rate scalability
  - No broadcasts and thus no potential threat of broadcast storms
  - No requirement for processing checksums
  - Simpler and more efficient extension header mechanisms
  - Flow labels for per-flow processing with no need to open a transport inner packet to identify the various traffic flows
- **Mobility and security,** to help comply with mobile IP and IP Security (IPsec) standards functionality. Mobility enables people with mobile network devices—many with wireless connectivity—to move around within networks.
  - Mobile IP is an Internet Engineering Task Force (IETF) standard that is available for both IPv4 and IPv6. The standard enables mobile devices to move without breaks in established network connections. Because IPv4 does not automatically provide this kind of mobility, you must add it with additional configurations.
  - In IPv6, mobility is built in, which means that any IPv6 node can use mobility when necessary. The routing headers of IPv6 make mobile IPv6 much more efficient for end nodes than mobile IPv4.
  - IPsec is the IETF standard for IP network security, available for both IPv4 and IPv6. Although the functionalities are essentially identical in both environments, IPsec is mandatory in IPv6. IPsec is enabled on every IPv6 node and is available for use, making the IPv6 Internet more secure. IPsec also requires keys for each party, which implies a global key deployment and distribution.
- **Finally, transition richness:** There are several ways to incorporate existing IPv4 capabilities with the added features of IPv6.
  - One approach is to implement a dual-stack method, with both IPv4 and IPv6 configured on the interface of a network device.
  - Tunneling is another technique that should become more prominent as the adoption of IPv6 grows. There is a variety of IPv6-over-IPv4 tunneling methods. Some methods require manual configuration, while others are dynamic.
  - Cisco IOS Release 12.3(2)T and later also include Network Address Translation-Protocol Translation (NAT-PT) between IPv6 and IPv4. This translation allows direct communication between hosts that use different versions of the IP protocol.

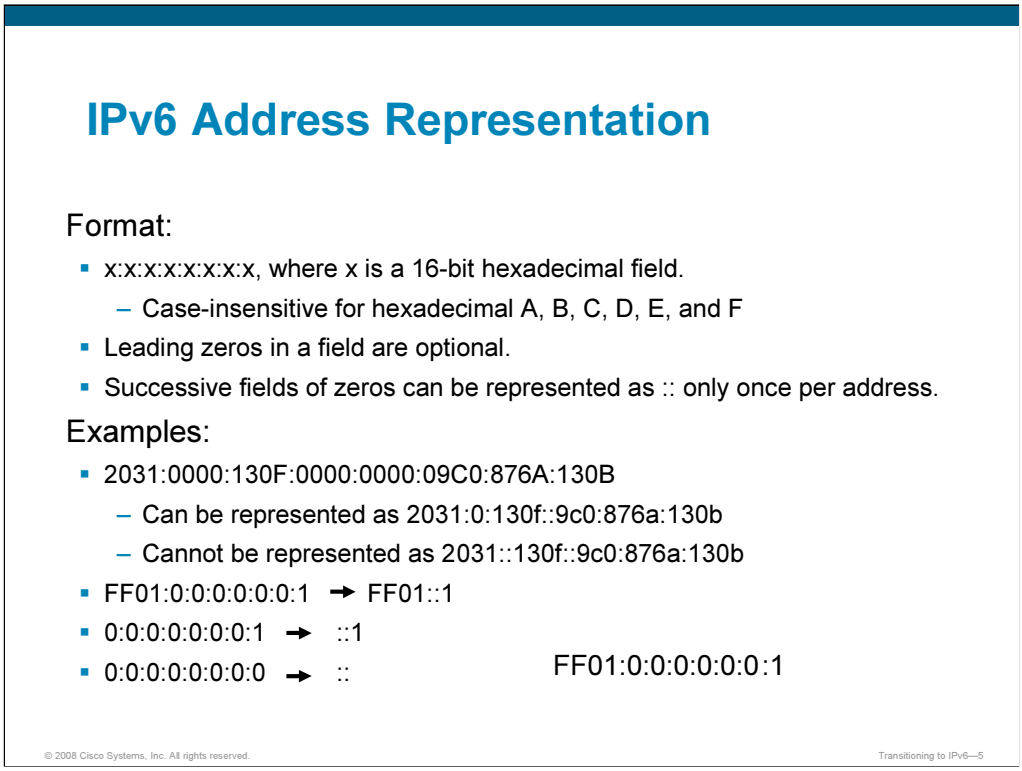

Colons separate entries in a series of 16-bit hexadecimal fields that represent IPv6 addresses. The hexadecimal digits A, B, C, D, E, and F that are represented in IPv6 addresses are not case-sensitive.

IPv6 does not require explicit address string notation. Use the following guidelines for IPv6 address string notations:

•The leading zeros in a field are optional, so that 09C0 equals 9C0 and 0000 equals 0.

•Successive fields of zeros can be represented as "::" only once in an address.

•An unspecified address is written as "::" because it contains only zeros.

•Using the "::" notation greatly reduces the size of most addresses. For example, FF01:0:0:0:0:0:0:1 becomes FF01::1.

•An address parser identifies the number of missing zeros by separating the two parts and entering 0 until the 128 bits are complete. If two "::" notations are placed in the address, there is no way to identify the size of each block of zeros.

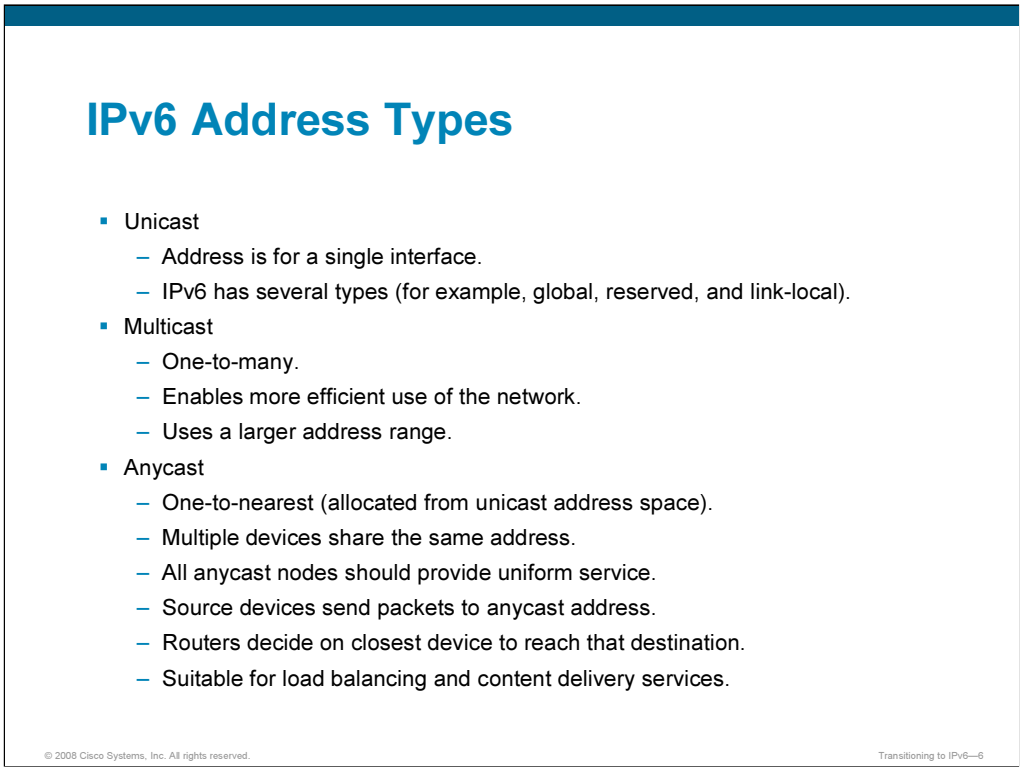

Broadcasting in IPv4 results in a number of problems. Broadcasting generates interrupts in every computer on the network and, in some cases, triggers malfunctions that can completely halt an entire network. This disastrous network event is known as a “broadcast storm.”

In IPv6, broadcasting does not exist. IPv6 replaces broadcasts with multicasts and anycasts. Multicast enables efficient network operation by using a number of functionally specific multicast groups to send requests to a limited number of computers on a network. The multicast groups prevent most of the problems that are related to broadcast storms in IPv4.

The range of multicast addresses in IPv6 is larger than in IPv4. For the near future, allocation of multicast groups is not being limited.

IPv6 also defines a new type of address known as an anycast address. An anycast address identifies a list of devices or nodes; therefore, an anycast address identifies multiple interfaces. Anycast addresses are like a cross between unicast and multicast addresses. Unicast sends packets to one specific device with one specific address, and multicast sends a packet to every member of a group. Anycast addresses send a packet to any one member of the group of devices with the anycast address assigned.

For efficiency, a packet that is sent to an anycast address is delivered to the closest interface—as defined by the routing protocols in use—that is identified by the anycast address, so anycast can also be thought of as a “one-to-nearest” type of address. Anycast addresses are syntactically indistinguishable from global unicast addresses because anycast addresses are allocated from the global unicast address space.

Note that there is little experience with widespread, arbitrary use of Internet anycast addresses, and there are some known complications and hazards when using them in their full generality. Until more experience has been gained and solutions have been agreed upon for those problems, the following restrictions are imposed on IPv6 anycast addresses: An anycast address MUST NOT be used as the source address of an IPv6 packet. And an anycast address MUST NOT be assigned to an IPv6 host, that is, it may be assigned to an IPv6 router only.

IPv6 Unicast Addressing

- These are types of IPv6 unicast addresses.
  - Global: Starts with 2000::/3 and assigned by IANA
  - Reserved: Used by the IETF
  - Private: Link local (starts with FE80::/10)
  - Loopback (::1)
  - Unspecified (::)
- A single interface may be assigned multiple IPv6 addresses of any type: unicast, anycast, or multicast.
- IPv6 addressing rules are covered by multiple RFCs.
  - Architecture defined by RFC 4291

There are several basic types of IPv6 unicast addresses: global, reserved, private (including link-local), loopback, and unspecified.

The IPv6 global unicast address is equivalent to the IPv4 global unicast address. A global unicast address is an IPv6 address from the global unicast prefix. The structure of global unicast addresses enables the aggregation of routing prefixes, which limits the number of routing table entries in the global routing table. Global unicast addresses used on links are aggregated upward through organizations and eventually to the ISPs.

The IETF reserves a portion of the IPv6 address space for various uses, both present and future. Reserved addresses represent 1/256th of the total IPv6 address space. Some of the other types of IPv6 addresses come from this block.

A block of IPv6 addresses is set aside for private addresses, just as is in IPv4. These private addresses are local only to a particular link or site and are never routed outside of a particular company network. Private addresses have a first octet value of “FE” in hexadecimal notation, with the next hexadecimal digit being a value from 8 to F. These addresses are further divided into types based on their scope.

- The concept of link-local addresses is new to IPv6. These addresses have a smaller scope than site-local addresses; they refer only to a particular physical link (such as a physical network). Routers do not forward datagrams using link-local addresses, not even within an organization; they are only for local communication on a particular physical network segment.
- These addresses are used for link communications such as automatic address configuration, neighbor discovery, and router discovery. Many IPv6 routing protocols also use link-local addresses.

Just as in IPv4, a provision is made for a special loopback IPv6 address for testing; datagrams sent to this address “loop back” to the sending device. However, in IPv6 there is just one address, not a whole block, for this function. The loopback address is 0:0:0:0:0:0:0:1, which is normally expressed using zero compression as “::1”.

In IPv4, an IP address of all zeros has a special meaning; it refers to the host itself, and is used when a device does not know its own address. In IPv6, this concept is formalized, and the all-zeros address (0:0:0:0:0:0:0:0) is named the “unspecified” address. It is typically used in the source field of a datagram that is sent by a device that seeks to have its IP address configured. You can apply address compression to this address; because the address is all zeros, the address becomes just “::”.

IPv6 Global Unicast (and Anycast) Addresses

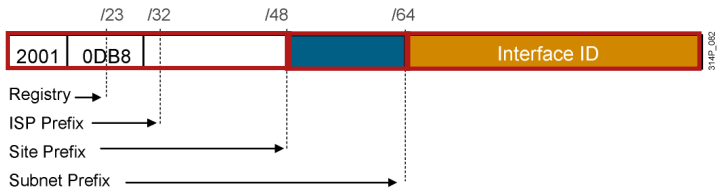

IPv6 has the same address format for global unicast and for anycast addresses.

- Uses a global routing prefix—a structure that enables aggregation upward, eventually to the ISP.
- A single interface may be assigned multiple addresses of any type (unicast, anycast, and multicast).
- Every IPv6-enabled interface contains at least one loopback (::1/128) and one link-local address.
- Optionally, every interface can have multiple unique local and global addresses.

Global unicast addresses are defined by a global routing prefix, a subnet ID, and an interface ID. The IPv6 unicast address space encompasses the entire IPv6 address range, with the exception of FF00::/8 (1111 1111), which is used for multicast addresses. The current global unicast address assignment by the Internet Assigned Numbers Authority (IANA) uses the range of addresses that start with binary value 001 (2000::/3), which is 1/8 of the total IPv6 address space and is the largest block of assigned block addresses.

Addresses with a prefix of 2000::/3 (001) through E000::/3 (111) are required to have 64-bit interface identifiers in the extended universal identifier (EUI)-64 format.

The IANA is allocating the IPv6 address space in the ranges of 2001::/16 to the registries.

The global unicast address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID. Individual organizations can use a 16-bit subnet field known as a "Subnet ID" to create their own local addressing hierarchy and to identify subnets. This field allows an organization to use up to 65,535 individual subnets. For more information, refer to RFC 3587, *IPv6 Global Unicast Address Format*, which replaces RFC 2374.

Link-Local Addresses

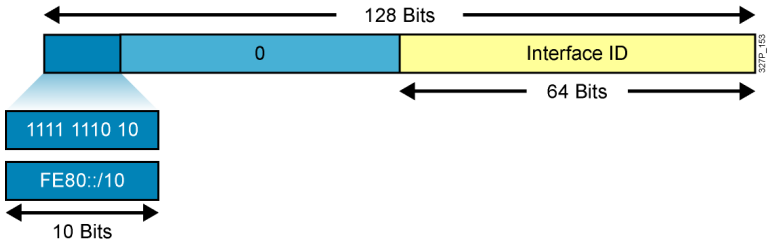

- Link-local addresses have a scope limited to the link and are dynamically created on all IPv6 interfaces by using a specific link-local prefix FE80::/10 and a 64-bit interface identifier.
- Link-local addresses are used for automatic address configuration, neighbor discovery, and router discovery. Link-local addresses 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.
- When communicating with a link-local address, you must specify the outgoing interface because every interface is connected to FE80::/10.

IPv6 is defined on most of the current data link layer protocols, including the following:

•Ethernet*

•PPP

•High-Level Data Link Control (HDLC)

•FDDI

•Token Ring

•Attached Resource Computer network (ARCnet)

•Nonbroadcast multiaccess (NBMA)

•ATM

•Frame Relay

•And IEEE 1394

An RFC describes the behavior of IPv6 in each of these specific data link layers, but Cisco IOS software does not necessarily support all of them. The data link layer defines how IPv6 interface identifiers are created and how neighbor discovery deals with data link layer address resolution.

Larger Address Space Enables Address Aggregation

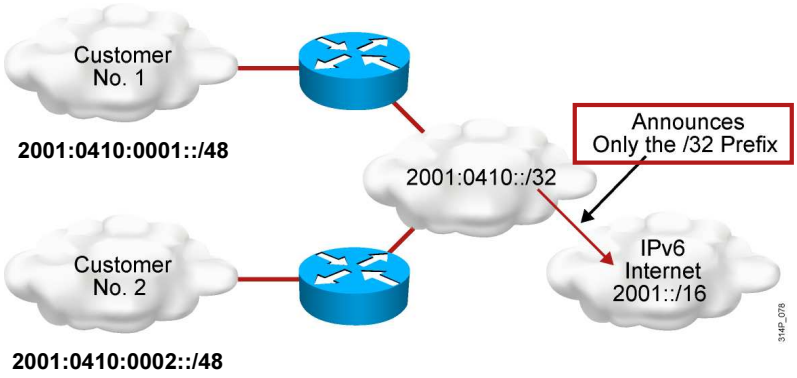

Address aggregation provides the following benefits:

- Aggregation of prefixes announced in the global routing table
- Efficient and scalable routing
- Improved bandwidth and functionality for user traffic

IPv6 allocates large numbers of addresses to ISPs and organizations. An ISP aggregates all of the prefixes of its customers into a single prefix and announces the single prefix to the IPv6 Internet. The increased address space is sufficient to allow organizations to define a single prefix for their entire network.

Aggregation of customer prefixes results in an efficient and scalable routing table. Scalable routing is necessary to expand broader adoption of network functions. Scalable routing also improves network bandwidth and functionality for user traffic that connects the various devices and applications.

Internet usage—both now and in the future—can include the following elements:

•A large increase in the number of broadband consumers with high-speed connections that are always on

•Users who spend more time online and are generally willing to spend more money on communications services and high-value searchable offerings

•Home networks with expanded network applications such as wireless VoIP, home surveillance, and advanced services such as real-time video on demand (VoD)

•And massive scalable games with global participants and media-rich e-learning, providing learners with on-demand remote labs or lab simulations

Assigning IPv6 Global Unicast Addresses

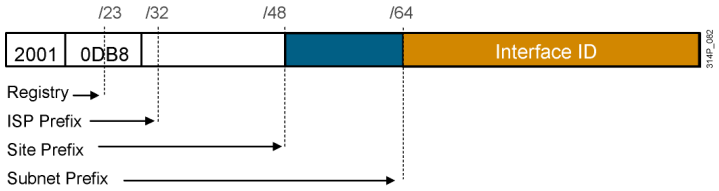

- Static assignment
  - Manual interface ID assignment
  - EUI-64 interface ID assignment
- Dynamic assignment
  - Stateless autoconfiguration
  - DHCPv6 (stateful)

```
RouterX(config-if) ipv6 address 2001:DB8:2222:7272::72/64
```

```
RouterX(config)# interface ethernet 0
RouterX(config-if)# ipv6 address 2001:0DB8:0:1::/64 eui-64
```


Interface identifiers in IPv6 addresses are used to identify interfaces on a link. They can also be thought of as the "host portion" of an IPv6 address. Interface identifiers are required to be unique on a specific link. Interface identifiers are always 64 bits and can be dynamically derived from a Layer 2 media and encapsulation.

Here are several ways to assign an IPv6 address to a device:

•Static assignment using a manual interface ID

•Static assignment using an EUI-64 interface ID

•Stateless autoconfiguration

•DHCP for IPv6 (DHCPv6)

One way to statically assign an IPv6 address to a device is to manually assign both the prefix (or network) and interface ID (or host) portion of the IPv6 address. To configure an IPv6 address on a Cisco router interface and enable IPv6 processing on that interface, use the **ipv6 address** command in interface configuration mode.

The example shows how to enable IPv6 processing on the interface and configure an address based on the directly specified bits.

Another way to statically assign an IPv6 address is to configure the prefix portion of the IPv6 address and derive the interface ID from the Layer 2 MAC address of the device, known as the EUI-64 interface ID.

To configure an IPv6 address for an interface and enable IPv6 processing on the interface using an EUI-64 interface ID in the low-order 64 bits of the address (or host), use the **ipv6 address eui-64** command in interface configuration mode.

The example assigns an IPv6 address to Ethernet interface 0 and uses an EUI-64 interface ID in the low order 64 bits of the address.

Autoconfiguration, as the name implies, is a mechanism that automatically configures the IPv6 address of a node. In IPv6, it is assumed that both PCs and non-PC devices are connected to a network. The autoconfiguration mechanism was introduced to enable plug-and-play networking of these devices and to help reduce administrative overhead.

DHCP for IPv6 enables DHCP servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes. It automatically allocates reusable network addresses and additional configuration with flexibility. This protocol is a stateful counterpart to IPv6 stateless address autoconfiguration (under RFC 2462), and can be used separately or concurrently with IPv6 stateless address autoconfiguration to obtain configuration parameters.

IPv6 EUI-64 Interface Identifier

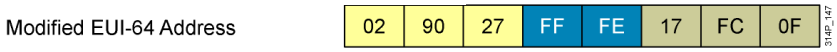

- Cisco can use the EUI-64 format for interface identifiers.
- This format expands the 48-bit MAC address to 64 bits by inserting “FFFE” into the middle 16 bits.
- To indicate that the chosen address is from a unique Ethernet MAC address, the U/L bit is set to 1 for global scope (0 for local scope).

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

Interface identifiers in global unicast and other IPv6 address types must be 64 bits long and can be constructed in the 64-bit EUI-64 format. The EUI-64 format interface ID is derived from the 48-bit data link layer (MAC) address by inserting the hexadecimal number FFFE between the upper three bytes (or the Organizational Unique Identifier [OUI] field) and the lower three bytes (or serial number) of the data link layer address. To indicate that the chosen address is from a unique Ethernet MAC address, the seventh 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.

The interface identifier for stateless autoconfiguration in an Ethernet environment uses the modified EUI-64 format. This format expands the 48-bit Ethernet MAC address format to a 64-bit version by inserting "FFFE" in the middle of the 48 bits. This creates a 64-bit version.

The seventh bit (starting with the leftmost bit as “1”) in an IPv6 interface identifier is referred to as the Universal/Local bit, or U/L bit. This bit identifies whether this interface identifier is locally unique on the link or that it is universally unique. In the case where the interface identifier is created from an Ethernet MAC address, it is assumed that the MAC address is universally unique and, therefore, the interface identifier is universally unique.

The rationale of the U/L bit is for future use of the upper-layer protocols to uniquely identify a connection, even in the context of a change in the leftmost part of the address. However, this is not yet used.

The eighth bit (starting with leftmost bit as “1”), also known as the “G” bit, is a group/individual bit for managing groups.

IPv6 Routing Protocols

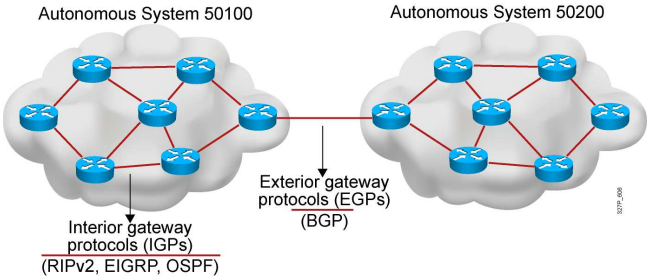

- IPv6 routing types
  - Static
  - RIPng (RFC 2080)
  - OSPFv3 (RFC 2740)
  - IS-IS for IPv6
  - MP-BGP4 (RFC 2545/2858)
  - EIGRP for IPv6
- The **ipv6 unicast-routing** command is required to enable IPv6 before any routing protocol is configured.

IPv6 uses longest-prefix match routing just like IPv4 classless interdomain routing (CIDR). Many of the common routing protocols have been modified to handle longer IPv6 addresses and different header structures. The updated routing protocols shown in the figure are currently available.

You can use and configure IPv6 static routing in the same way that you would with IPv4. There is an IPv6-specific requirement per RFC 2461 that a router must be able to determine the link-local address of each of its neighboring routers to ensure that the target address of a redirect message identifies the neighbor router by its link-local address. This requirement means that using a global unicast address as a next-hop address with IPv6 routing is not recommended.

The Cisco IOS global command to enable IPv6 is **ipv6 unicast-routing**. You must enable IPv6 unicast routing before an IPv6-capable routing protocol, or an IPv6 static route, will work.

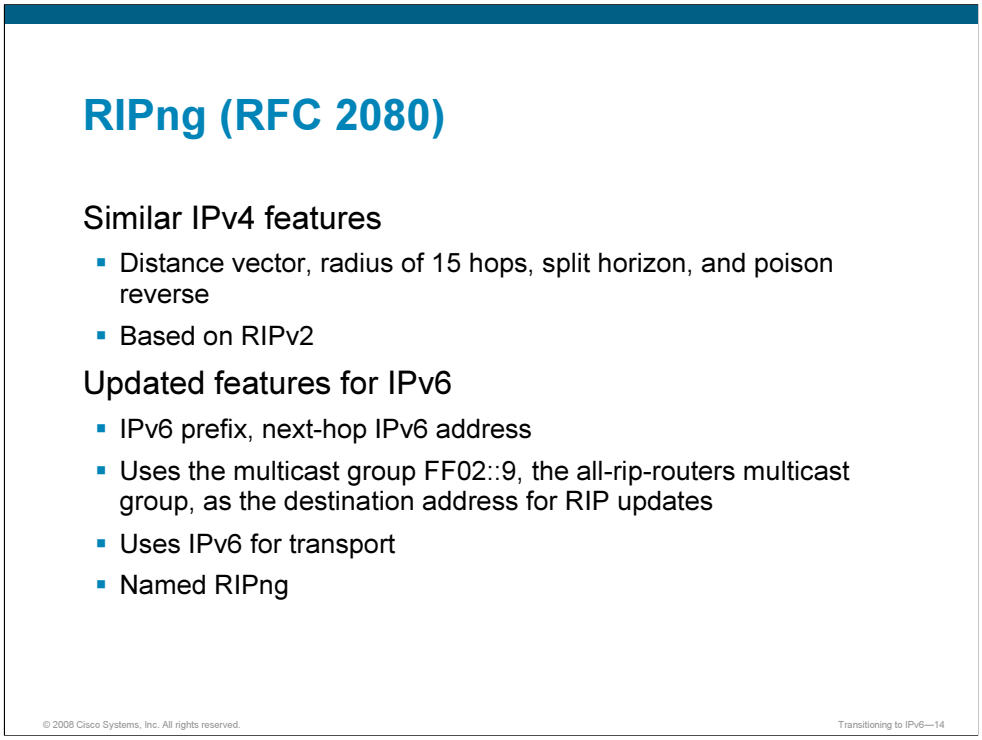

Routing Information Protocol next generation (RIPng) is a distance vector routing protocol with a limit of 15 hops that uses split horizon and poison reverse to prevent routing loops. RIPng includes the following features:

•Is based on and is similar to IPv4 RIP version 2

•Uses IPv6 for transport

•Includes the IPv6 prefix and next-hop IPv6 address

•Uses the multicast group FF02::9, the all-RIP-routers multicast group, as the destination address for RIP updates

•Sends updates on User Datagram Protocol (UDP) port 521

•Is supported by Cisco IOS Release 12.2(2)T and later

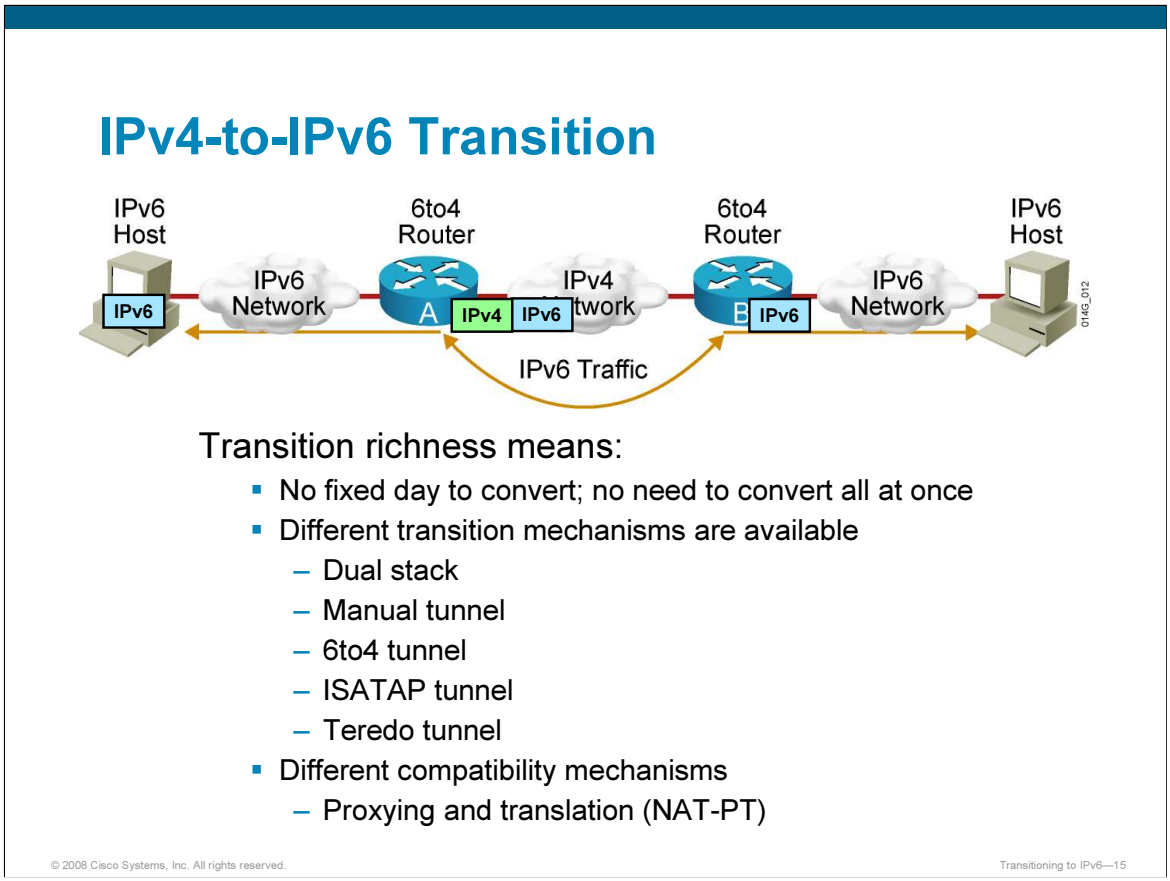

The transition from IPv4 does not require upgrades on all nodes at the same time. Many transition mechanisms enable smooth integration of IPv4 and IPv6. Other mechanisms that allow IPv4 nodes to communicate with IPv6 nodes are available. All of these mechanisms are applied to different situations.

The three most common techniques to transition from IPv4 to IPv6 are as follows:

- Dual stack is an integration method where the node has connectivity to both an IPv4 and IPv6 network. As a result, the node and its corresponding routers have two protocol stacks.
- There are several tunneling techniques available:
    - Manual IPv6-over-IPv4 tunneling is an integration method where an IPv6 packet is encapsulated within IPv4. This requires dual-stack routers.
    - Dynamic 6to4 tunneling is a method that automatically connects IPv6 islands through an IPv4 network, typically the Internet. The 6to4 tunneling method dynamically applies a valid, unique IPv6 prefix to each IPv6 island, enabling fast deployment of IPv6 in a corporate network without address retrieval from ISPs or registries.
    - Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunneling is an automatic overlay tunneling mechanism that uses the underlying IPv4 network as a link layer for IPv6. ISATAP tunnels allow individual IPv4 or IPv6 dual-stack hosts within a site to communicate with other such hosts on a virtual link, creating an IPv6 network using the IPv4 infrastructure.
    - Teredo tunneling is an IPv6 transition technology that provides host-to-host automatic tunneling instead of gateway tunneling. It passes unicast IPv6 traffic when dual-stacked hosts (that is, hosts that are running both IPv6 and IPv4) are located behind one or multiple IPv4 Network Address Translators.
- Proxying and translation is a translation mechanism that sits between an IPv6 network and an IPv4 network. The job of the translator is to translate IPv6 packets into IPv4 packets and vice versa.

Cisco IOS Dual Stack

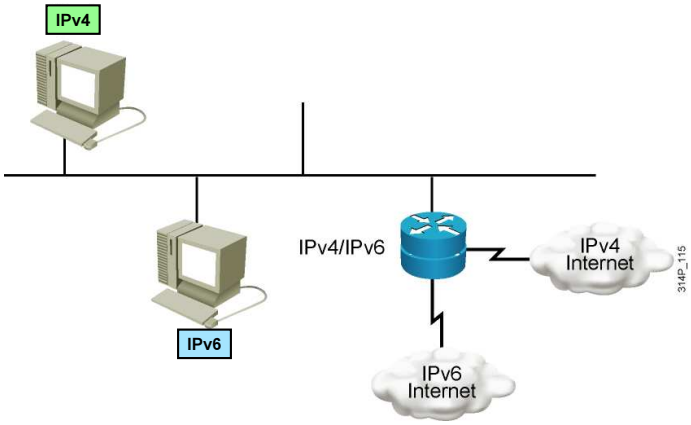

Dual stack is an integration method in which a node has implementation and connectivity to both an IPv4 and IPv6 network.

Dual stack is an integration method where a node has connectivity to both an IPv4 and IPv6 network, thus the node has two stacks. You can configure this on the same interface or on multiple interfaces. Features of the dual-stack method are as follows:

•A dual-stack node chooses which stack to use based on the destination address. A dual-stack node should prefer IPv6 when it is available. The dual-stack approach to IPv6 integration, in which nodes have both IPv4 and IPv6 stacks, is one of the most commonly used integration methods. Old IPv4-only applications continue to work as before. New and modified applications take advantage of both IP layers.

•A new application programming interface (API) supports both IPv4 and IPv6 addresses and DNS requests. This new API replaces the “get host by name” and “get host by address” calls. A converted application can make use of both IPv4 and IPv6. An application can be converted to the new API while still using only IPv4.

•Experience in porting IPv4 applications to IPv6 suggests that, for most applications, there is a minimal change in some localized places inside the source code. This technique is well known and has been applied in the past for other protocol transitions. It enables gradual application upgrades, one by one, to IPv6.

Cisco IOS Dual Stack (Cont.)

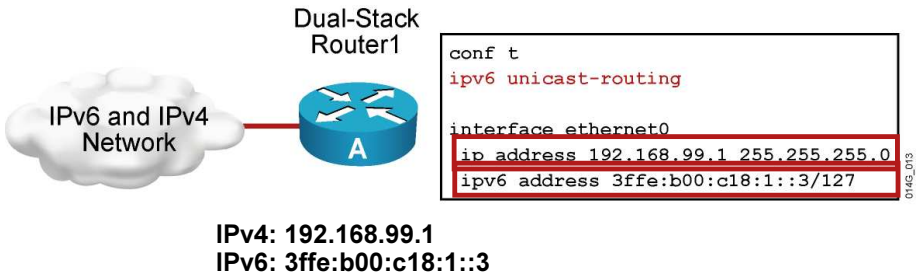

IPv4: 192.168.99.1
IPv6: 3ffe:b00:c18:1::3

When both IPv4 and IPv6 are configured on an interface, the interface is considered dual-stacked.

Cisco IOS Software Release 12.2(2)T and later are IPv6-ready. As soon as you configure basic IPv4 and IPv6 on an interface, it is dual-stacked and forwards IPv4 and IPv6 traffic. Using IPv6 on a Cisco IOS router requires that you use the global configuration command **ipv6 unicast-routing**. This command enables forwarding of IPv6 datagrams.

Note that you must configure all interfaces that forward IPv6 traffic with an IPv6 address using the interface command **ipv6 address**.

IPv6 Tunneling

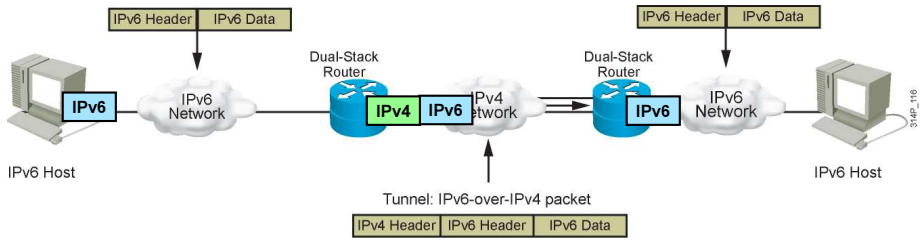

Tunneling is an integration method in which an IPv6 packet is encapsulated within another protocol, such as IPv4. This method of encapsulation is IPv4.

- Includes a 20-byte IPv4 header with no options and an IPv6 header and payload
- Requires dual-stack routers

Tunneling is an integration method where an IPv6 packet is encapsulated within another protocol, such as IPv4. When IPv4 encapsulates the IPv6 packet, a protocol type of 41 is specified in the IPv4 header, and the packet has the following characteristics:

•It includes a 20-byte IPv4 header with no options and an IPv6 header and payload.

•It requires dual-stack routers. This process connects IPv6 islands without needing to also convert an intermediary network to IPv6. Tunneling presents these two issues:

- •The maximum transmission unit (MTU) is effectively decreased by 20 octets if the IPv4 header does not contain any optional field.
- •A tunneled network is often difficult to troubleshoot. Tunneling is an intermediate integration and transition technique that should not be considered a final solution. A native IPv6 architecture should be the ultimate goal.

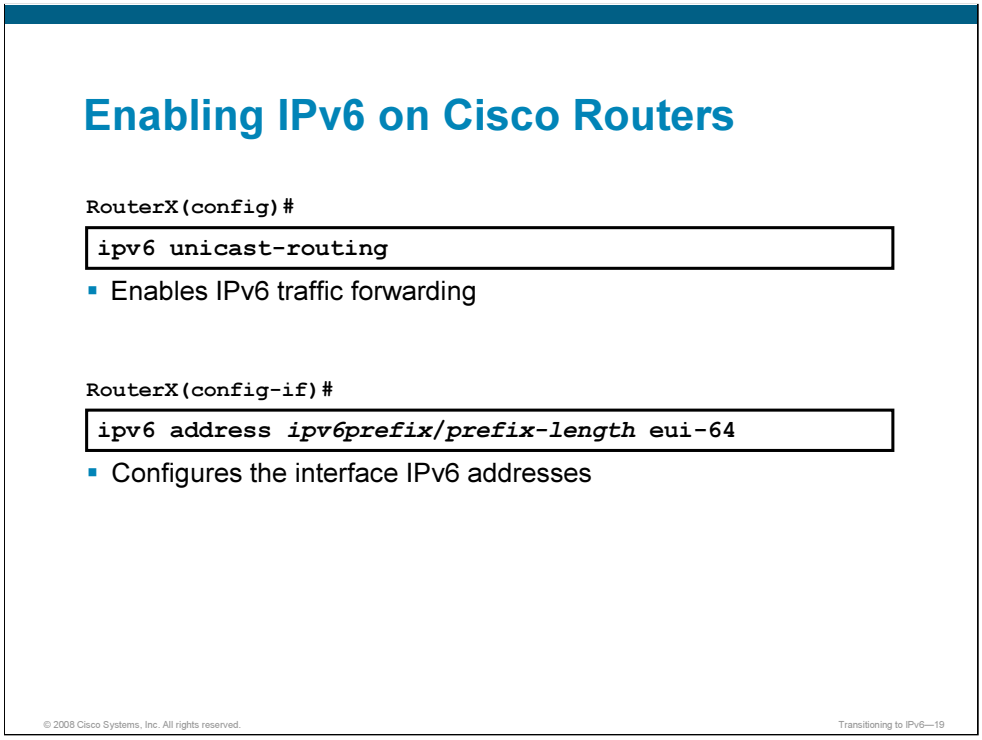

There are two basic steps to activate IPv6 on a router. First, IPv6 traffic forwarding must be activated, then each interface where IPv6 is required must be configured.

By default, IPv6 traffic forwarding is disabled on a Cisco router. To activate IPv6 traffic forwarding between interfaces, the global command **ipv6 unicast-routing** must be configured. This enables the forwarding of unicast IPv6 traffic.

IPv6 is enabled on a per-interface basis.

The **ipv6 address** command configures a global IPv6 address. The link-local address is automatically configured when an address is assigned to an interface. You must specify the entire 128-bit IPv6 address or use the 64-bit prefix by using the **eui-64** option.

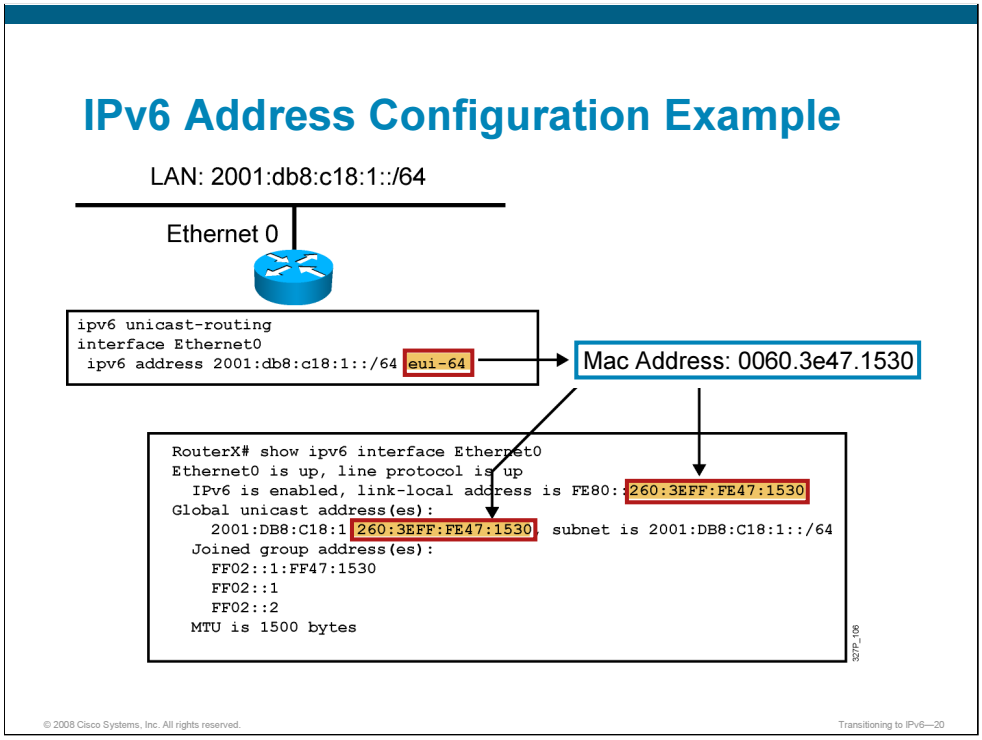

You can completely specify the IPv6 address or compute the host identifier (which is the rightmost 64 bits) from the EUI-64 identifier of the interface. In this example, the IPv6 address of the interface is configured using the EUI-64 format.

Alternatively, you can completely specify the entire IPv6 address to assign a router interface an address using the **ipv6 address** command in interface configuration mode.

Note that configuring an IPv6 address on an interface automatically configures the link-local address for that interface.

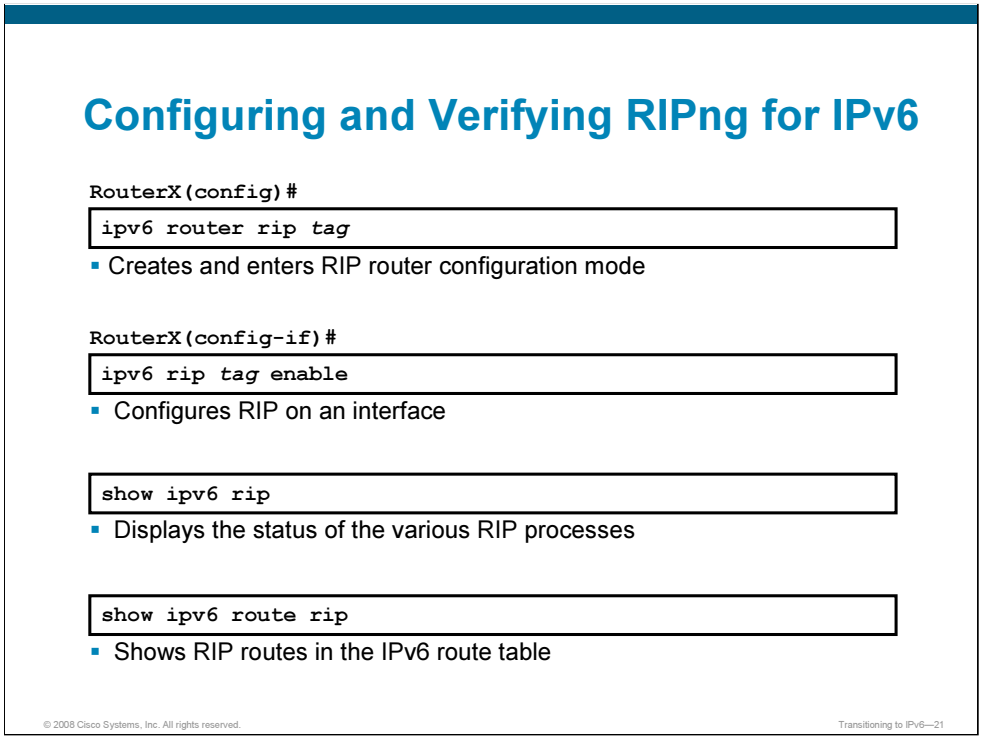

This figure shows examples of command syntax used to configure RIPng. The syntax is similar, if not identical, to their IPv4 counterparts.

For RIPng, instead of using the **network** command to identify which interfaces should run RIPng, use the command **ipv6 rip** *tag* **enable** in interface configuration mode to enable RIPng on an interface. The *tag* parameter that you use for the **ipv6 rip enable** command must match the *tag* parameter in the **ipv6 router rip** command.

Note that enabling RIP on an interface dynamically creates a “router rip” process, if necessary.

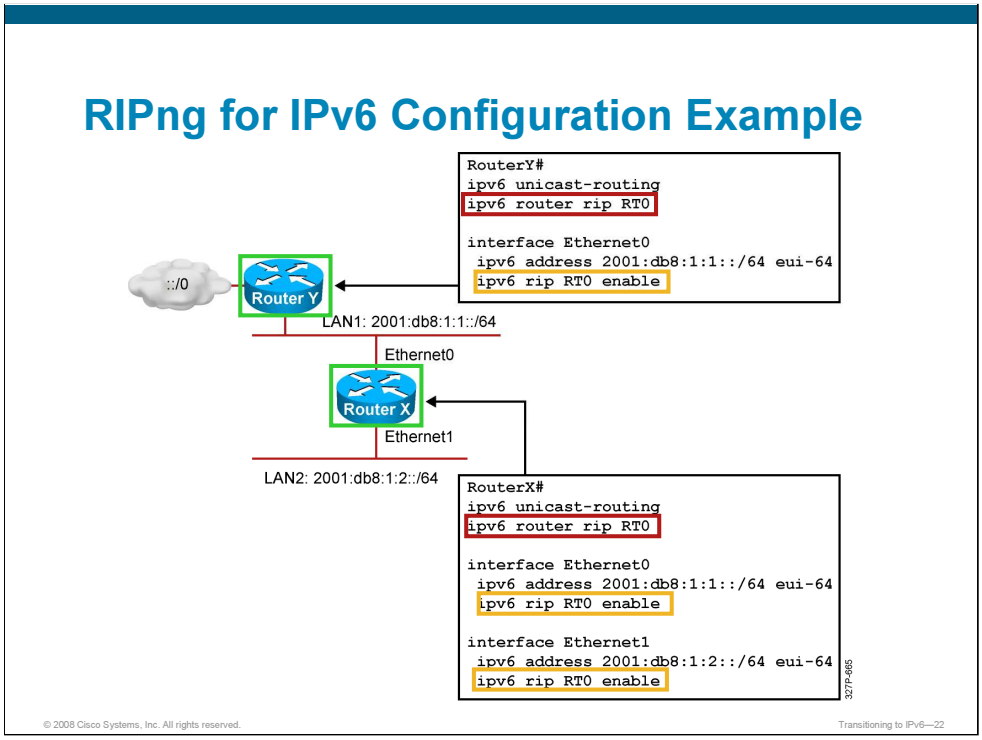

The example shows a network of two routers. Router Y is connected to the default network. On both router X and router Y, “RT0” is a tag that identifies the RIPng process. RIPng is enabled on the first Ethernet interface of router Y using the **ipv6 rip RT0 enable** command. Router X shows that RIPng is enabled on both Ethernet interfaces using the **ipv6 rip RT0 enable** command.

Summary

- IPv6 offers many additional benefits to IPv4, including a larger address space, easier address aggregation, and integrated security.
- An IPv6 address is 128 bits long and is made up of a 48-bit global prefix, a 16-bit subnet ID, and a 64-bit interface identifier.
- There are several ways to assign IPv6 addresses: statical, stateless autoconfiguration, and DHCPv6.
- Cisco supports all of the major IPv6 routing protocols: RIPng, OSPFv3, and EIGRP.
- Transitioning from IPv4 to IPv6 requires dual stacks, tunneling, and possibly NAT-PT.
- Use the **ipv6 unicast-routing** command to enable IPv6 and the **ipv6 address** *ipv6-address/prefix-length* command to assign interface addresses and enable an IPv6 routing protocol.

In summary, IPv6 offers many additional benefits to IPv4, including a larger address space, easier address aggregation, and integrated security.

An IPv6 address is 128 bits long and is made up of a 48-bit global prefix, a 16-bit subnet ID, and a 64-bit interface identifier.

There are several ways to assign IPv6 addresses: statical, stateless autoconfiguration, and DHCPv6.

Cisco supports all of the major IPv6 routing protocols, including RIPng, OSPFv3, and EIGRP.

Transitioning from IPv4 to IPv6 requires dual stacks, tunneling, and possibly NAT-PT.

Use the **ipv6 unicast-routing** command to enable IPv6 and the **ipv6 address** command to assign interface addresses and enable an IPv6 routing protocol.

Quiz Question 1

Which address type from IPv4 was eliminated in IPv6?

a) unicast
b) multicast
c) everycast
d) broadcast

The correct answer is d.

Quiz Question 2

How can you condense consecutive sets of zeros in an IPv6 address?

a) with the ":::" symbol
b) by eliminating leading zeros
c) by replacing four consecutive zeros with a single zero
d) with the "::" symbol

The correct answer is d.

Thank you for taking the Transitioning to IPv6 Quick Learning Module.