Chapter 8: Implementing IPv6 in the Enterprise Network

CCNP ROUTE: Implementing IP Routing

Cisco | Networking Academy®
Mind Wide Open™

Chapter 8 Objectives

- Describe IPv6.
- Describe the basics of IPv6 addressing.
- Describe and configure IPv6 addresses.
- Describe and configure IPv6 routing.
- Describe and configure IPv6 tunneling.
- Describe and configure static and dynamic NAT-PT.

Introducing IPv6

- The ability to scale networks for future demands requires a limitless supply of IP addresses and improved mobility.
  - IPv6 combines expanded addressing with a more efficient and feature-rich header to meet these demands.
  - While it has many similarities to IPv4, IPv6 satisfies the increasingly complex requirements of hierarchical addressing that IPv4 does not support.

Explosion of New IP-Enabled Devices

- More and more IP-enabled devices are connecting.
  - Devices include cell phones, consumer products (blue ray players, TVs), etc.

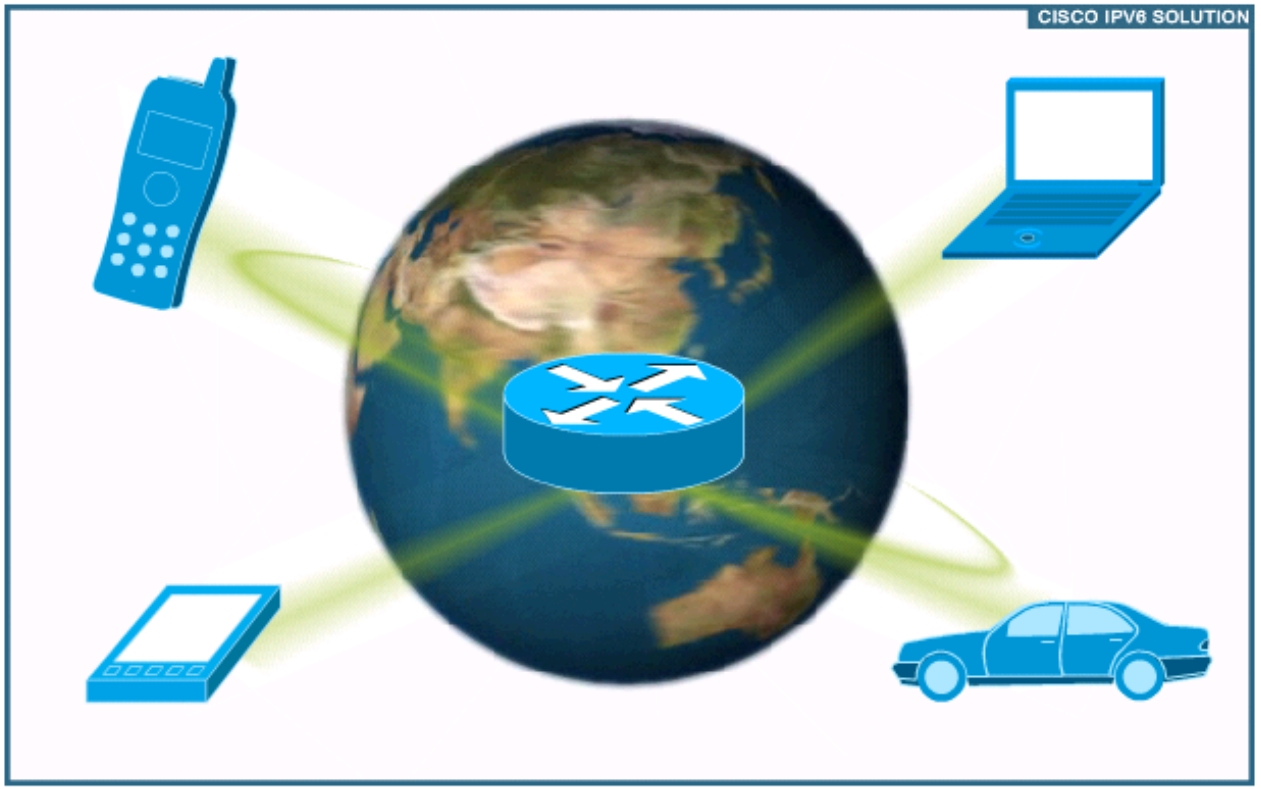

IPv4 Issues

- In January 2010, only 10% of the public IPv4 addresses remained unallocated.
  - It is estimated that this pool will have exhausted by the late 2011 (?).
- Internet routing table expansion
  - The Internet routing tables continue to grow which means Internet core routers require more processing power, memory, and overhead.
- Lack of true end-to-end model
  - IPv4 networks typically use NAT as the solution to address depletion.
  - However, NAT hides the true source address of traffic, which can cause other issues.

Features of IPv6

- **Larger address space**
  - IPv6 addresses are 128 bits, compared to IPv4's 32 bits.
    - There are enough IPv6 addresses to allocate more than the entire IPv4 Internet address space to everyone on the planet.
- **Elimination of public-to-private NAT**
  - End-to-end communication traceability is possible.
- **Elimination of broadcast addresses**
  - IPv6 now includes unicast, multicast, and anycast addresses.
- **Support for mobility and security**
  - Helps ensure compliance with mobile IP and IPsec standards.
- **Simplified header for improved router efficiency**

IPv6 Address Types

Address Type	Description	Topology
Unicast	"***One to One***" • An address destined for a single interface. • A packet sent to a unicast address is delivered to the interface identified by that address.	
Multicast	"***One to Many***" • An address for a set of interfaces (typically belonging to different nodes). • A packet sent to a multicast address will be delivered to all interfaces identified by that address.	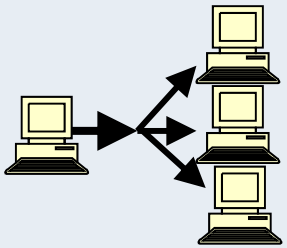
Anycast	"***One to Nearest***" (Allocated from Unicast) • An address for a set of interfaces. • In most cases these interfaces belong to different nodes. • A packet sent to an anycast address is delivered to the closest interface as determined by the IGP.	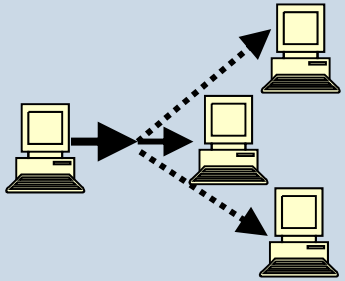

IPv4 Header vs. IPv6 Header

- The IPv4 header has 20 octets containing 12 basic header fields.
- The IPv6 header has 40 octets containing 8 fields.
- Three of these fields are identical in nature.
- Other fields serve similar functions as in IPv4.
- The remaining IPv4 fields no longer exist in IPv6.

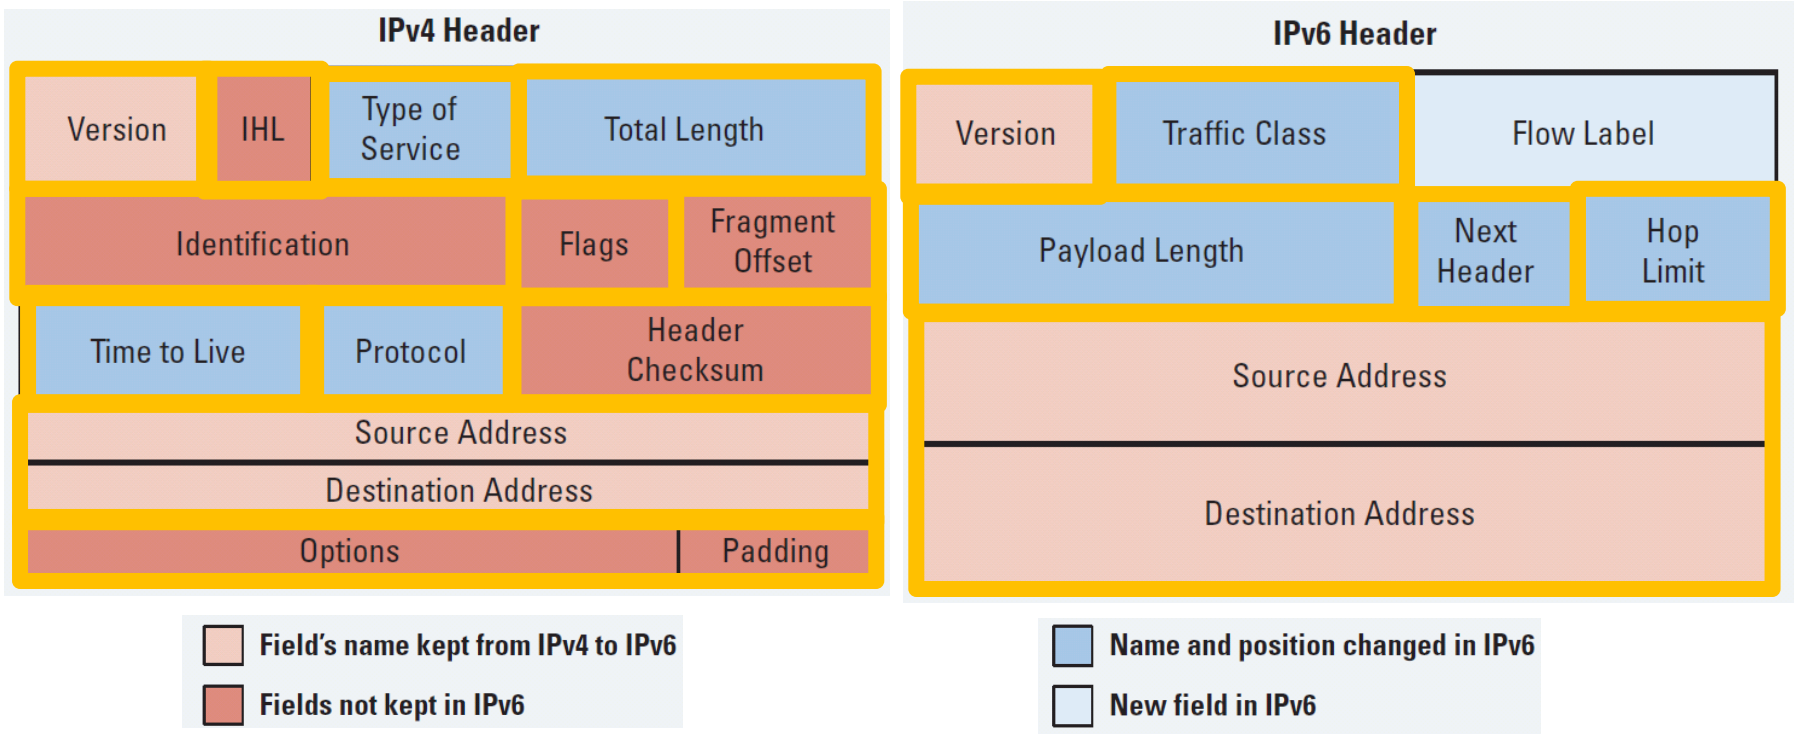

Protocol and Next Header Fields

- In IPv4 the Protocol field is used to identify the next level protocol (e.g., TCP, UDP, ICMP, …).
- In IPv6, this field is called the "Next Header" field and serves the same purpose.

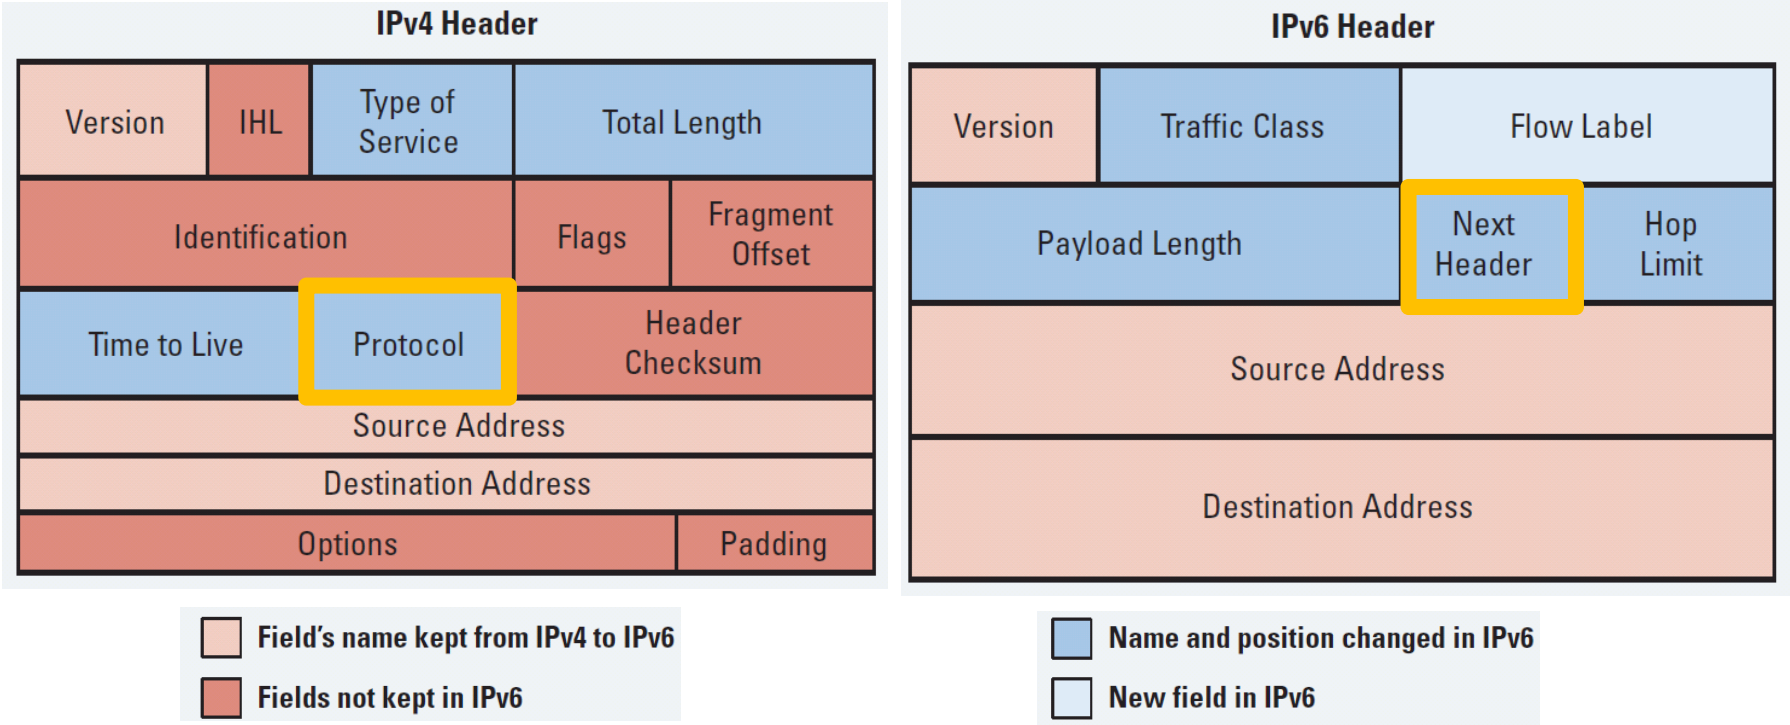

Extension Headers

- The Next Header field identifies what follows the Destination Address field:
  - Protocols:
    - TCP (protocol 6)
    - UDP (protocol 17)
    - ICMPv6 (protocol 58)
  - Extension header
- Extension headers make the handling of options more efficient.

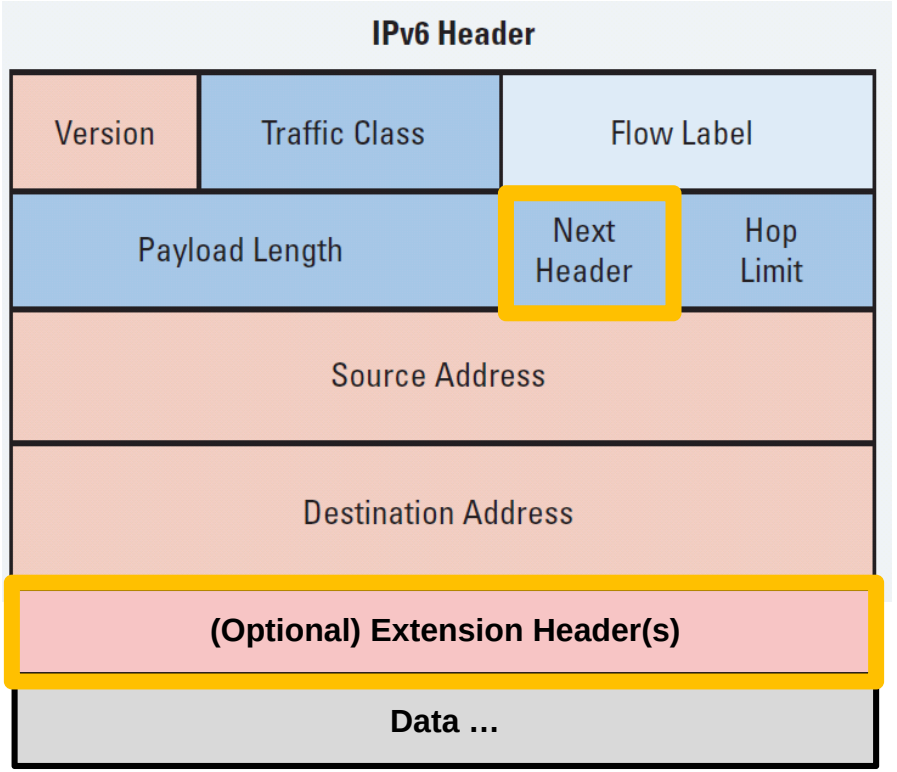

Extension Headers

- Multiple extension headers (called a chain) may be included in an IPv6 packet.
  - The number of extension headers is not fixed, so the total length of the extension header chain is variable.
- The destination node examines the first extension header (if any).
  - The contents determine whether or not the node should examine the next header.
  - Therefore, extension headers must be processed in the order they appear in the packet.

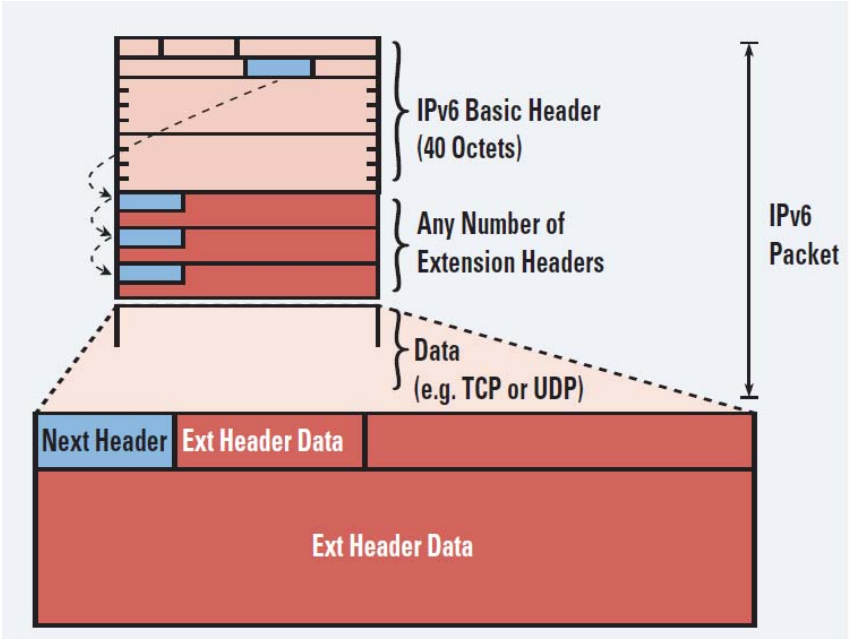

New IPv6 Features

- **Prefix renumbering**
  - IPv6 allows simplified mechanisms for address and prefix renumbering.
- **Multiple addresses per interface**
  - An IPv6 interface can have multiple addresses.
- **Link-local addresses**
  - IPv6 link-local addresses are used as the next hop when IGPs are exchanging routing updates.
- **Stateless autoconfiguration:**
  - DHCP is not required because an IPv6 device can automatically assign itself a unique IPv6 link-local address.
- **Provider-dependent or provider-independent addressing**

IPv6 Addressing Overview

- IPv6 increases the number of address bits by a factor of 4, from 32 to 128, providing a very large number of addressable nodes.

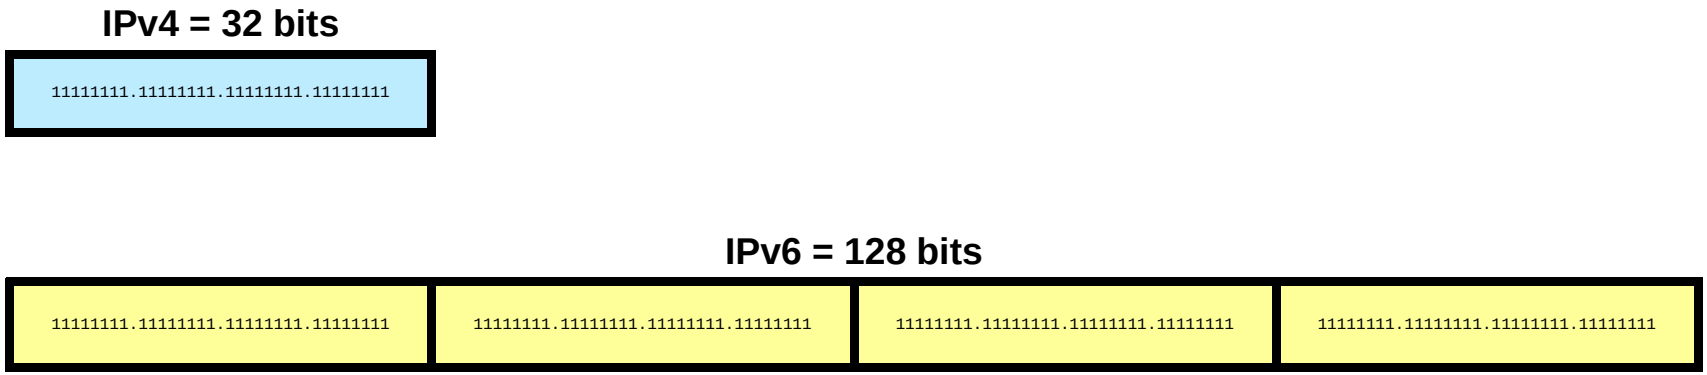

IPv6 Address Specifics

- The 128-bit IPv6 address is written using hexadecimal numbers.
  - Specifically, it consists of 8, 16-bit segments separated with colons between each set of four hex digits (16 bits).
  - Referred to as “coloned hex” format.
  - Hex digits are not case sensitive.
  - The format is **`x:x:x:x:x:x:x:x`**, where **`x`** is a 16-bit hexadecimal field therefore each **`x`** is representing four hexadecimal digits.
- An example address is as follows:
  - **`2035:0001:2BC5:0000:0000:087C:0000:000A`**

Abbreviating IPv6 Addresses

- Leading 0s within each set of four hexadecimal digits can be omitted.
  - **0**9C0 = 9C0
  - **0000** = **0**
- A pair of colons (“**::**”) can be used, *once* within an address, to represent any number (“a bunch”) of successive **0**s.

IPv6 Address Example

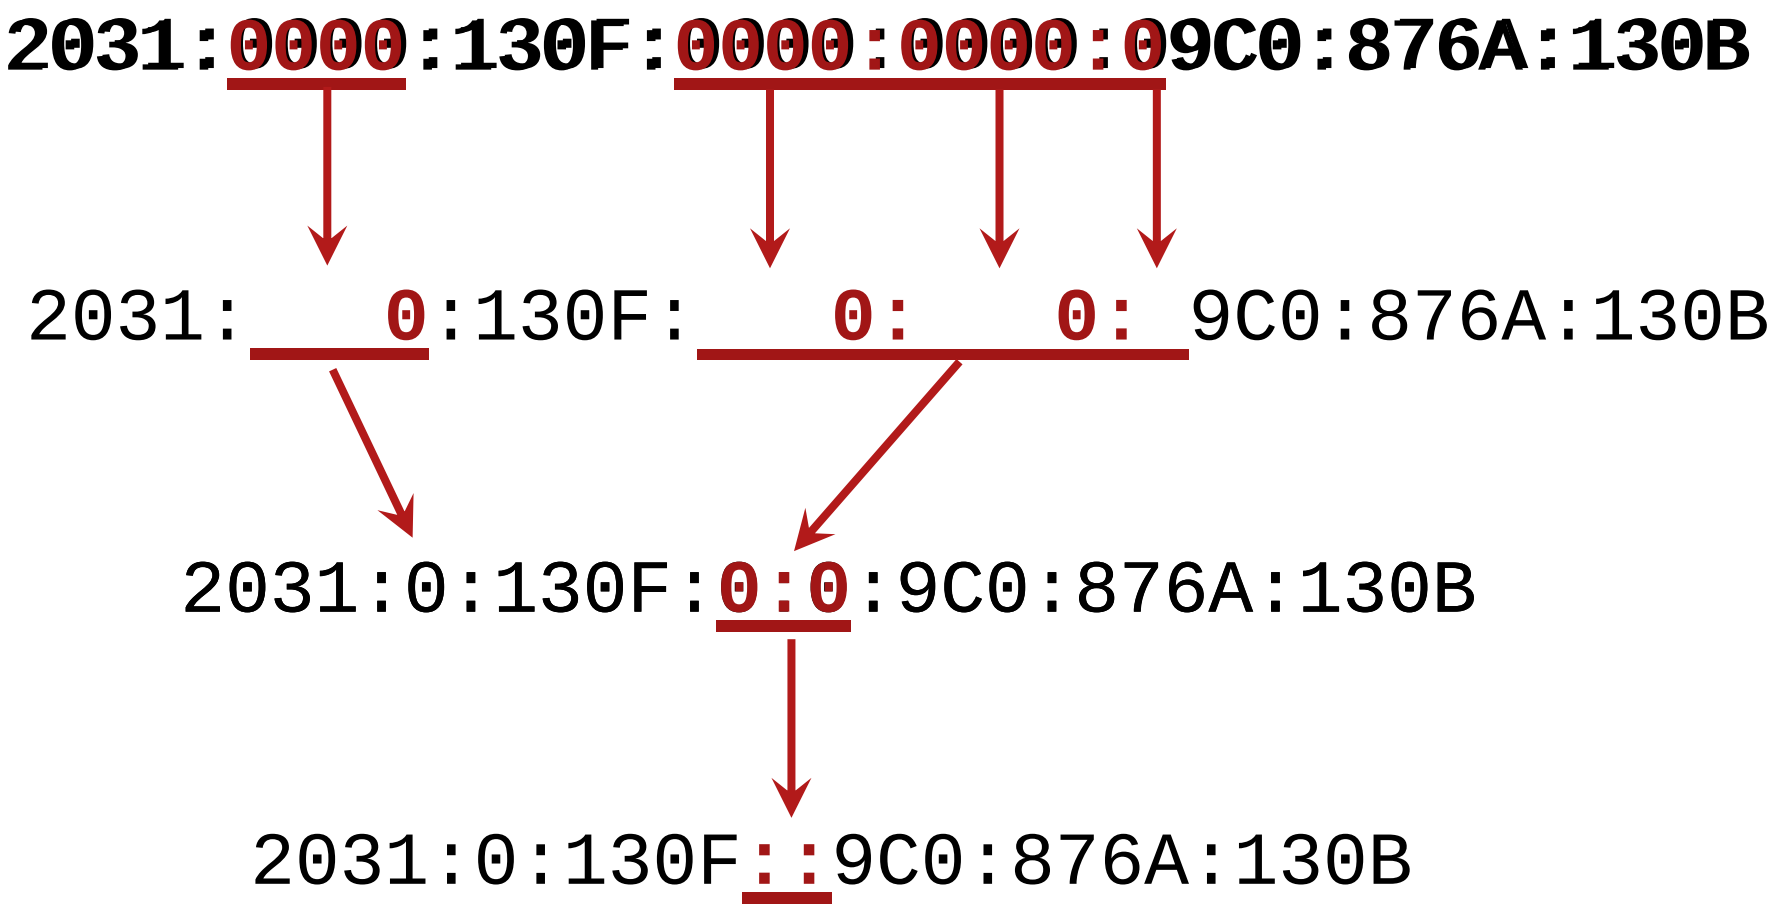

IPv6 Address Example

FF01:0000:0000:0000:0000:0000:0000:1

FF01:**0:0:0:0:0:0**:1 = FF01::1

E3D7:0000:0000:0000:51F4:00C8:C0A8:6420

= E3D7::51F4:C8:C0A8:6420

3FFE:0501:0008:0000:0260:97FF:FE40:EFAB

= 3FFE:501:8:0:260:97FF:FE40:EFAB

= 3FFE:501:8::260:97FF:FE40:EFAB

IPv6 Addressing in an Enterprise Network

- An IPv6 address consists of two parts:
  - A *subnet prefix* representing the network to which the interface is connected.
    - Usually 64-bits in length.
  - An *interface ID*, sometimes called a local identifier or a token.
    - Usually 64-bits in length.

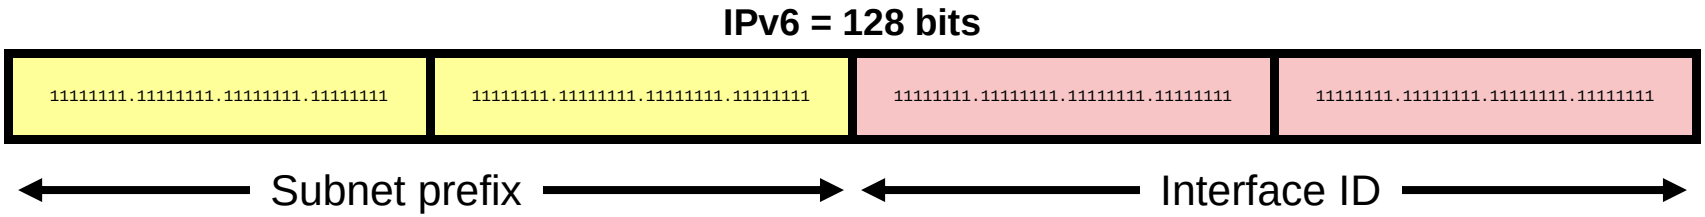

Subnet Prefix

- IPv6 uses the "*/prefix-length*" CIDR notation to denote how many bits in the IPv6 address represent the subnet.
- The syntax is *ipv6-address/prefix-length*
  - *ipv6-address* is the 128-bit IPv6 address
  - */prefix-length* is a decimal value representing how many of the left most contiguous bits of the address comprise the prefix.

For example:

fec0:0:0:1::1234/64

is really

fec0:0000:0000:0001:0000:0000:0000:1234/64

- The first 64-bits (**fec0:0000:0000:0001**) forms the address prefix.
- The last 64-bits (**0000:0000:0000:1234**) forms the Interface ID.

Interface Identifiers

- IPv6 addresses on a link must be unique.
  - Although they all share the same 64-bit subnet prefix they are made unique by the interface ID.
- Because the prefix length is fixed and well-known (64-bits), IPv6 hosts can automatically create a unique IPv6 address.
- For example, the following Layer 2 protocols can dynamically create the IPv6 address interface ID:
  - Ethernet (using the EUI-64 format discussed later)
  - PPP
  - HDLC
  - NBMA, Frame Relay

Special IPv6 Addresses

IPv6 Address	Description
::/0	• All routes and used when specifying a default static route. • It is equivalent to the IPv4 quad-zero (0.0.0.0).
::/128	• Unspecified address and is initially assigned to a host when it first resolves its local link address.
::1/128	• Loopback address of local host. • Equivalent to 127.0.0.1 in IPv4.
FE80::/10	• Link-local unicast address. • Similar to the Windows autoconfiguration IP address of 169.254.x.x.
FF00::/8	• Multicast addresses.
All other addresses	• Global unicast address.

IPv6 Address Scopes

- Address types have well-defined destination scopes:
  - **Link-local address**
  - **Global unicast address**
  - **Site-local address**

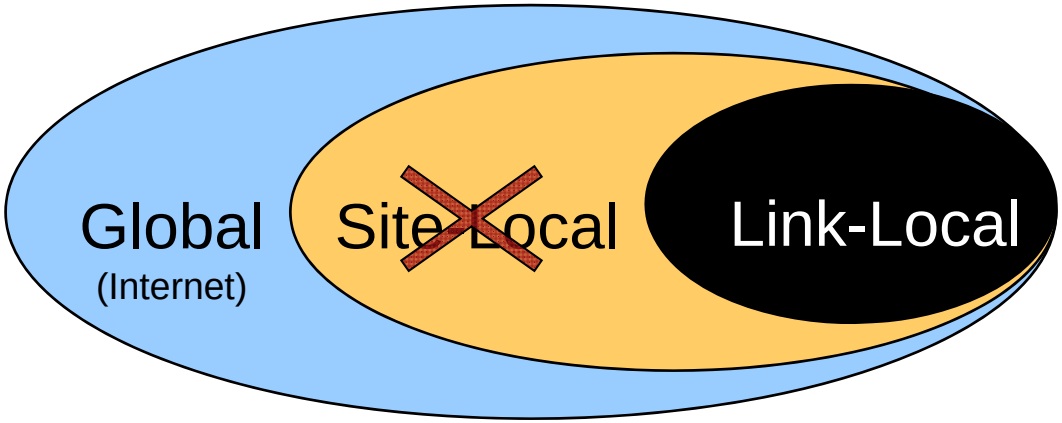

- **Note:**
  - Site-Local Address are no longer supported in RFC 3879.

Multiple IP Addresses per Interface

- An interface can have multiple IPv6 addresses simultaneously configured and enabled on it.
  - However, it must have a link-local address.
- Typically, an interface is assigned a link-local and one (or more) global IPv6 address.
  - For example, an Ethernet interface can have:
    - Link-local address `(e.g., FE80::21B:D5FF:FE5B:A408)`
    - Global unicast address `(e.g., 2001:8:85A3:4289:21B:D5FF:FE5B:A408)`
- Note:
  - An interface could also be configured to simultaneously support IPv4 and IPv6 addresses.
  - This creates a “dual-stacked” interface which is discussed later.

IPv6 Link-Local Address

- Link-local addresses are used for automatic address configuration, neighbor discovery, router discovery, and by many routing protocols.
- They are dynamically created using a link-local prefix of **FE80::/10** and a 64-bit interface identifier.
  - Unique only on the link, and it is not routable off the link.

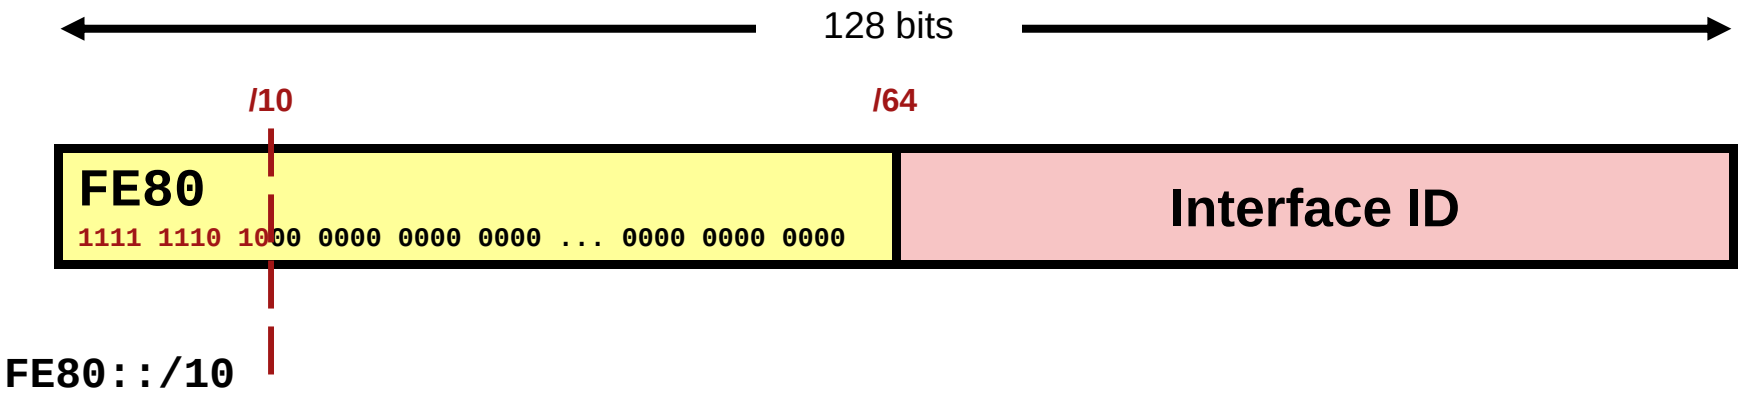

IPv6 Link-Local Address

- Link-local packets are unique only on the link, and are not routable off the link.
  - Packets with a link-local destination must stay on the link where they have been generated.
  - Routers that could forward them to other links are not allowed to do so because there has been no verification of uniqueness outside the context of the origin link.

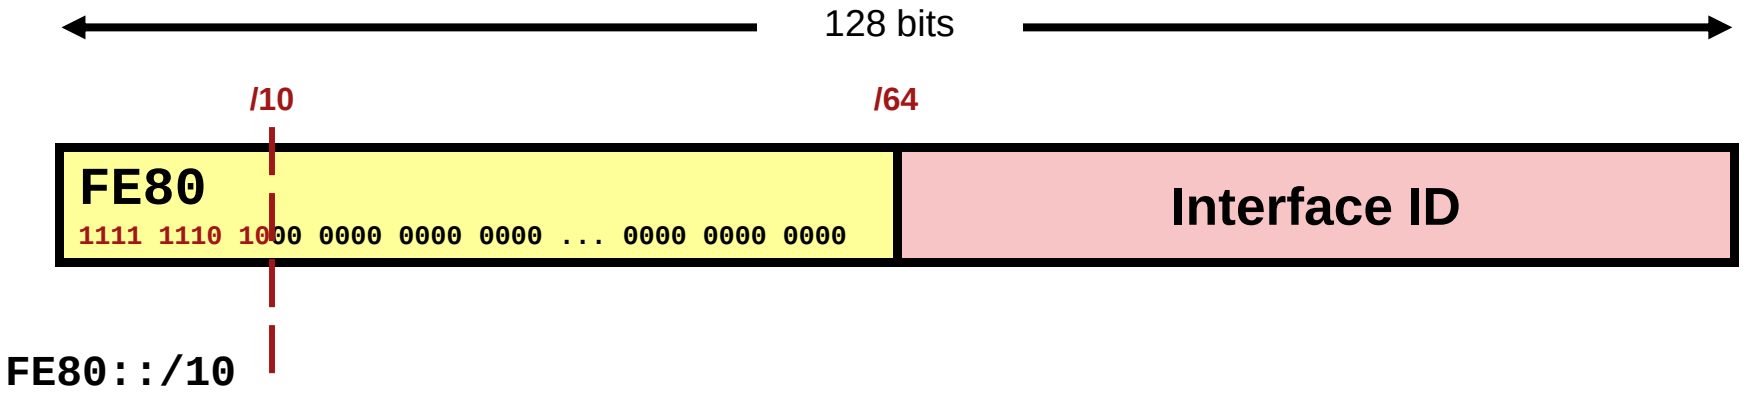

IPv6 Link-Local Address

- When communicating with a link-local address, the outgoing interface must be specified because every interface is connected to FE80::/10.
  - For example, if you ping the neighbor's link-local address, you will be asked to input the interface from which you wish to ping.

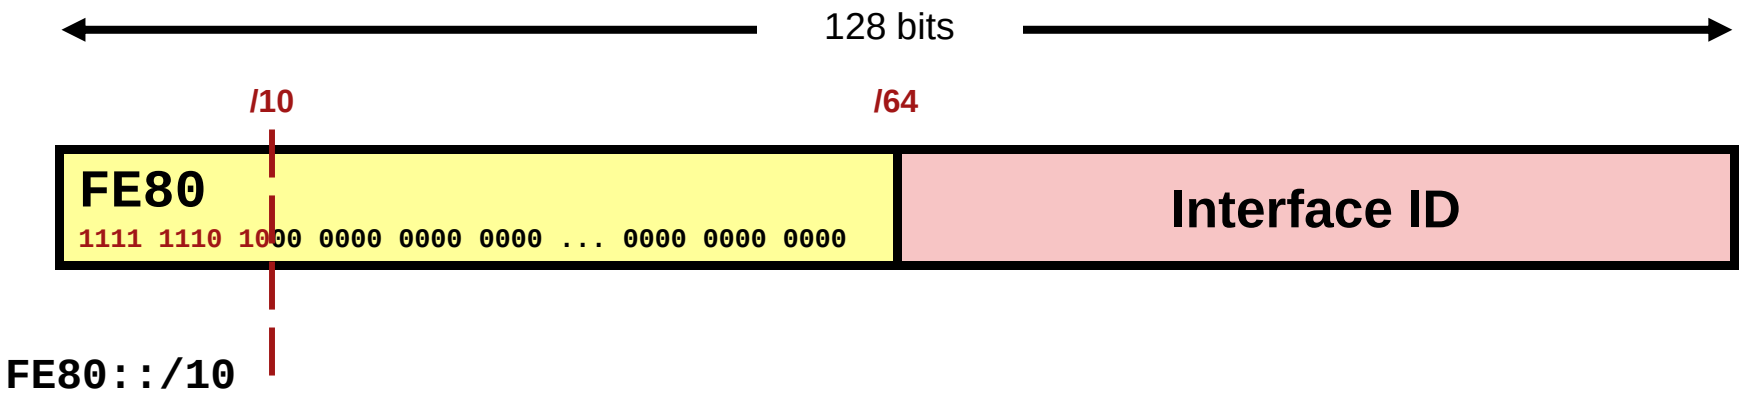

IPv6 Link-Local Address Example

```
R1# show ipv6 interface loopback 100
Loopback100 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::222:55FF:FE18:7DE8
  No Virtual link-local address(es):
  Global unicast address(es):
    2001:8:85A3:4290:222:55FF:FE18:7DE8, subnet is 2001:8:85A3:4290::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF18:7DE8
  MTU is 1514 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ICMP unreachables are sent
  ND DAD is not supported
  ND reachable time is 30000 milliseconds (using 31238)
  Hosts use stateless autoconfig for addresses.
R1#
```

IPv6 Global Unicast Address

- A global unicast address is an IPv6 address from the global public unicast prefix (2001::/16).
  - The structure enables aggregation of routing prefixes to reduce the number of routing table entries in the global routing table.
- Global unicast addresses are aggregated upward through organizations and eventually to the ISPs.

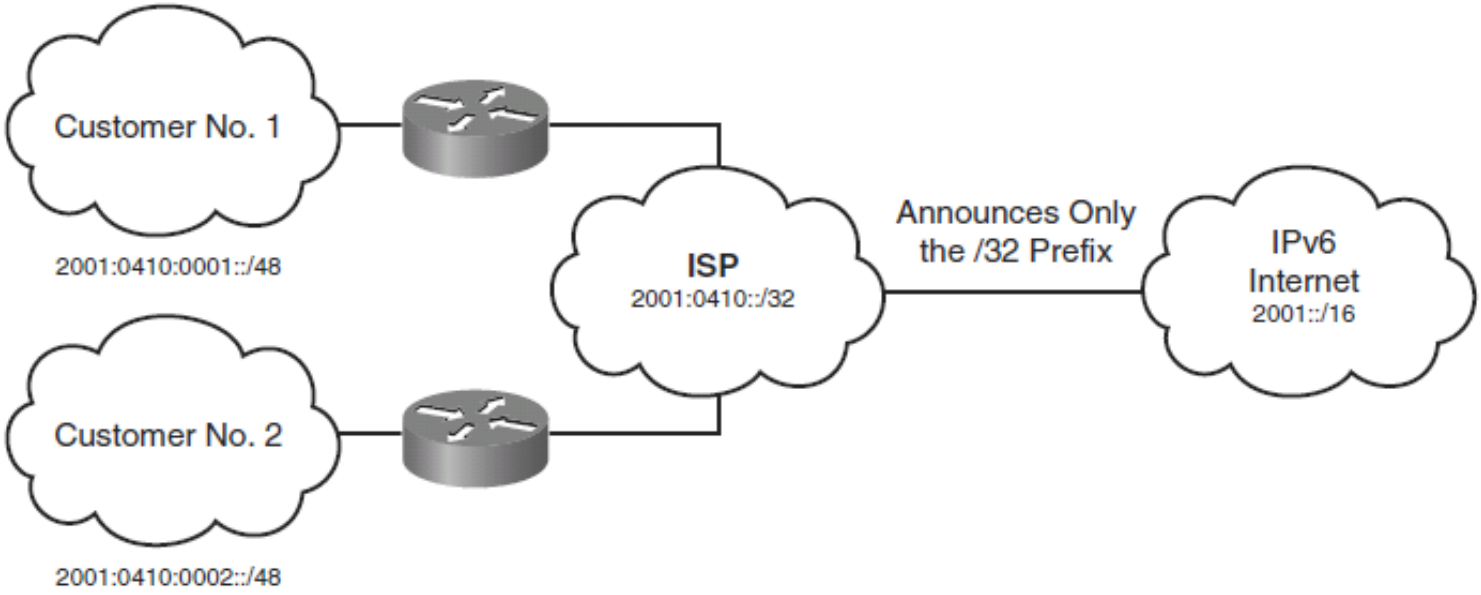

IPv6 Global Unicast Address

- The global unicast address typically consists of:
  - A 48-bit global routing prefix
  - A 16-bit subnet ID
  - A 64-bit interface ID (typically in EUI-64 bit format discussed later).

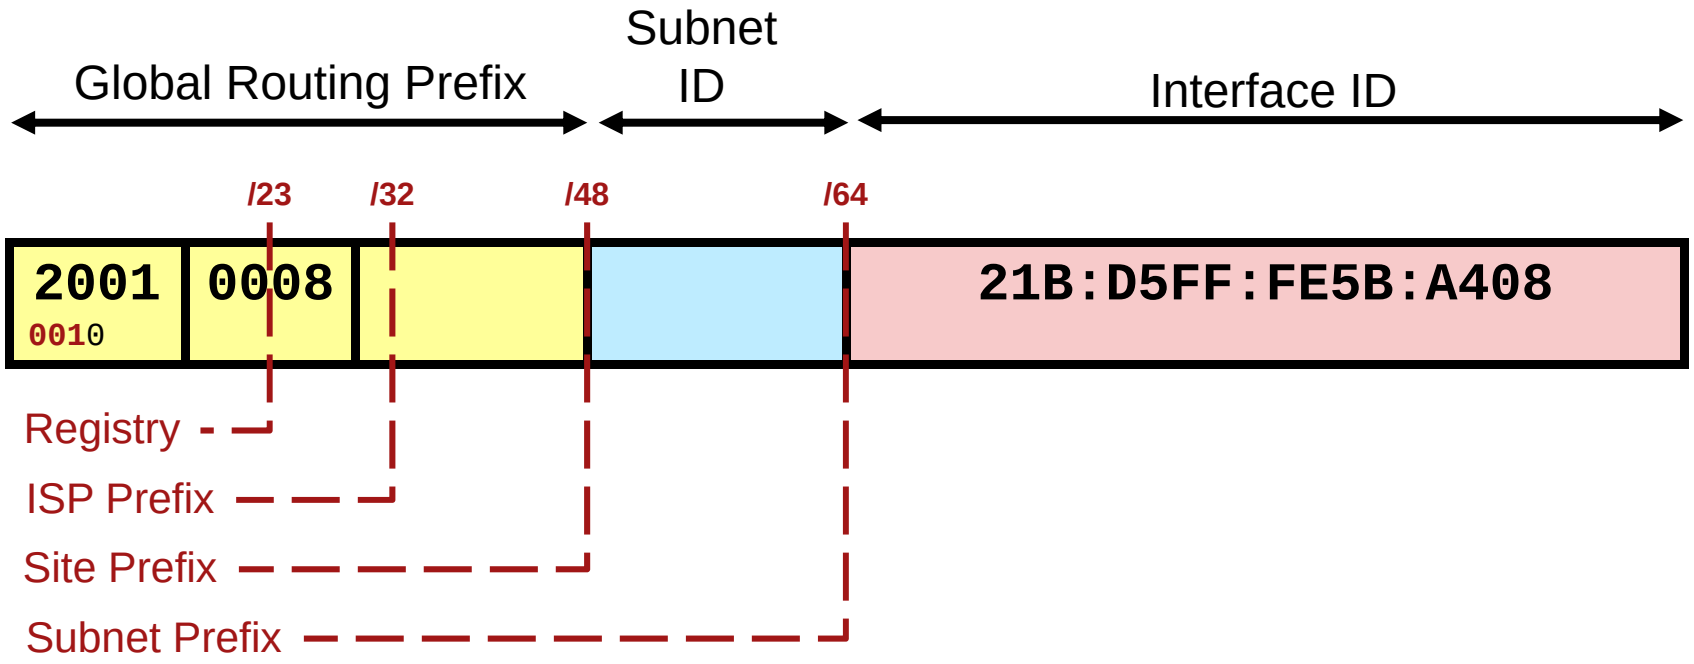

IPv6 Global Unicast Address

- The current IANA global routing prefix uses the range that start with binary **0010 (2000::/3).**
  - Addresses with a prefix of **2000::/3 (001)** to **E000::/3 (111)** are required to have 64-bit interface IDs in the extended universal identifier (EUI)-64 format.

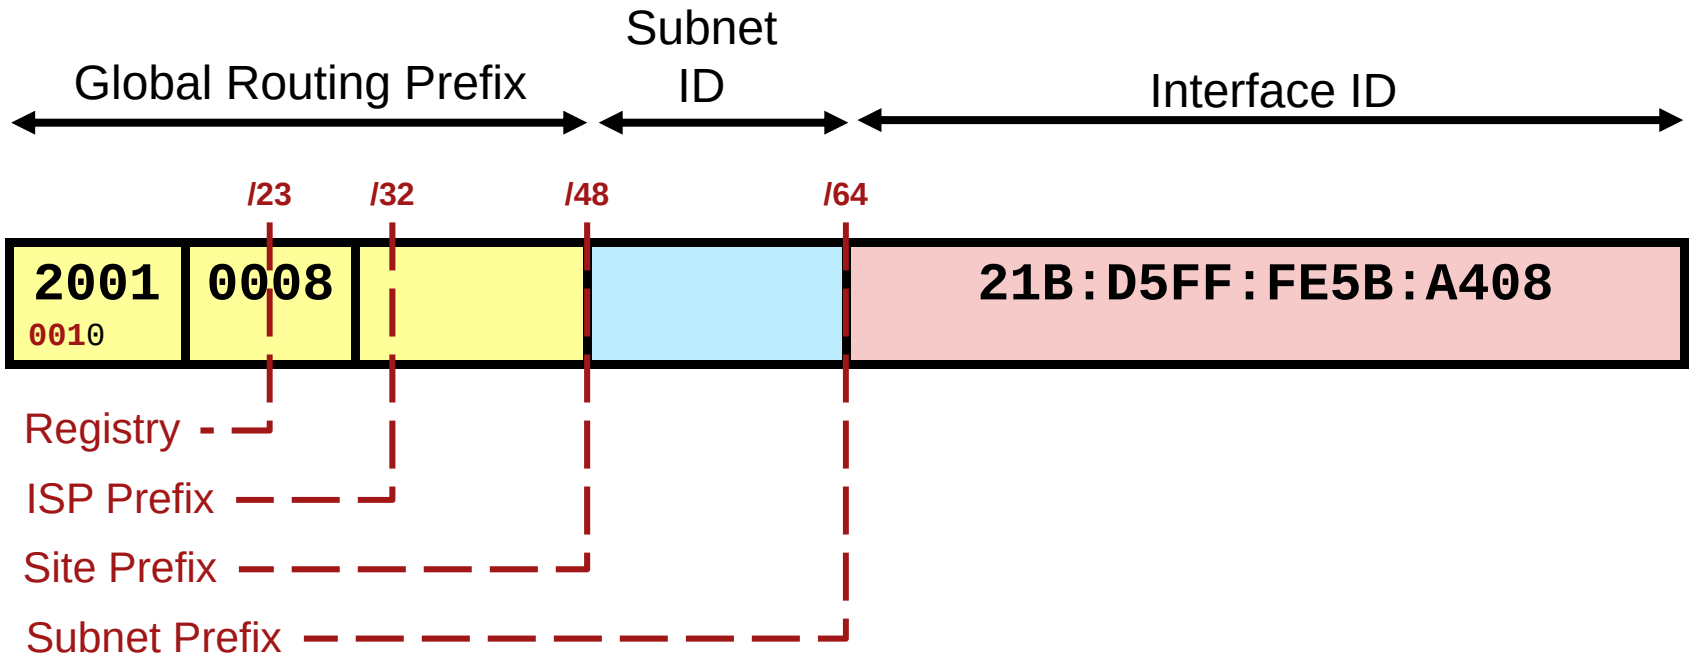

IPv6 Global Unicast Address

- The subnet ID can be used by an organization to create their own local addressing hierarchy.
  - This 16-bit field allows up to 65,536 individual subnets.

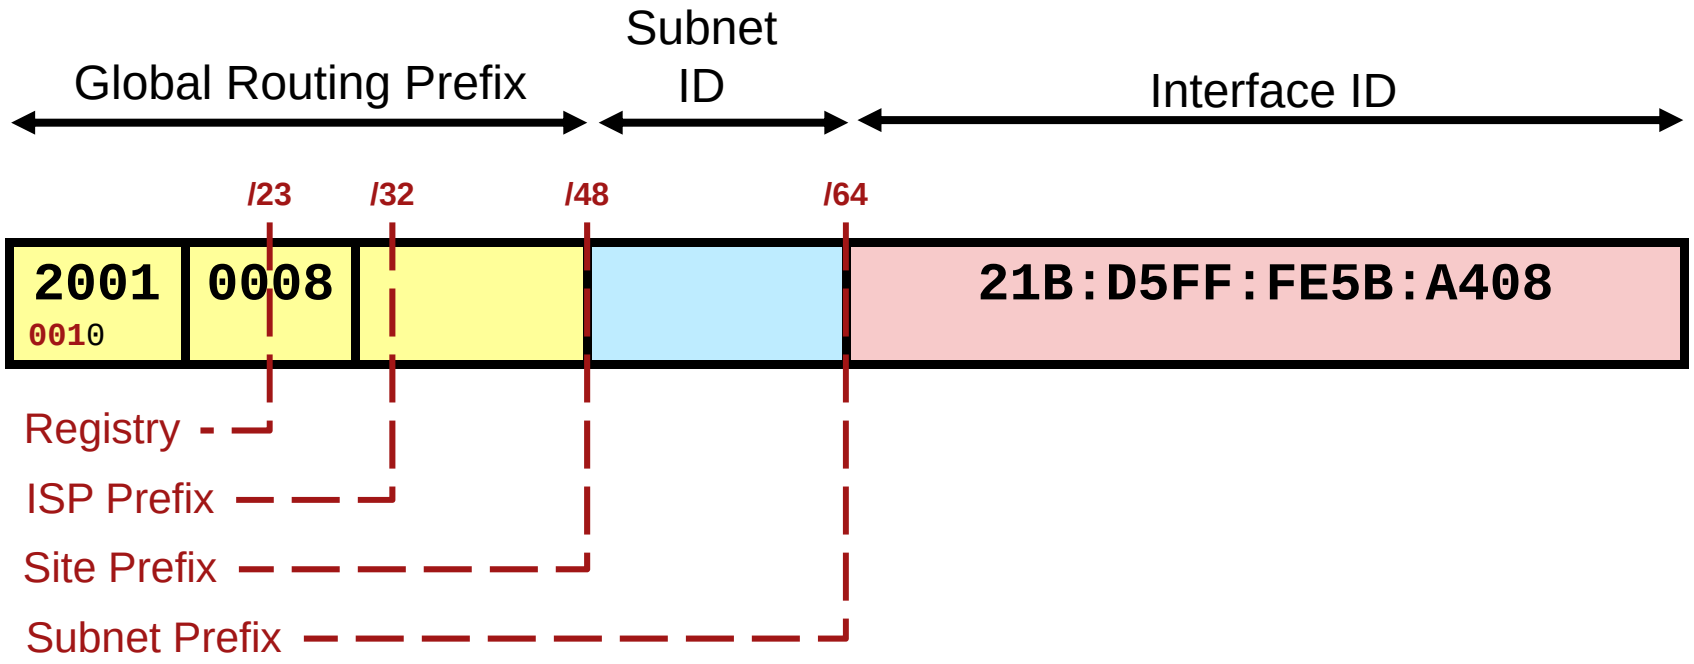

IPv6 Global Unicast Address Example

```
R1# show ipv6 interface loopback 100
Loopback100 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::222:55FF:FE18:7DE8
  No Virtual link-local address(es):
  Global unicast address(es):
    2001:8:85A3:4290:222:55FF:FE18:7DE8, subnet is 2001:8:85A3:4290::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF18:7DE8
  MTU is 1514 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ICMP unreachables are sent
  ND DAD is not supported
  ND reachable time is 30000 milliseconds (using 31238)
  Hosts use stateless autoconfig for addresses.
R1#
```

IPv6 Multicast Address

- Multicasting is at the core of many IPv6 functions and it is a replacement for the broadcast address.
- They are defined by the prefix **FF00::/8**.
  - An interface may belong to any number of multicast groups.

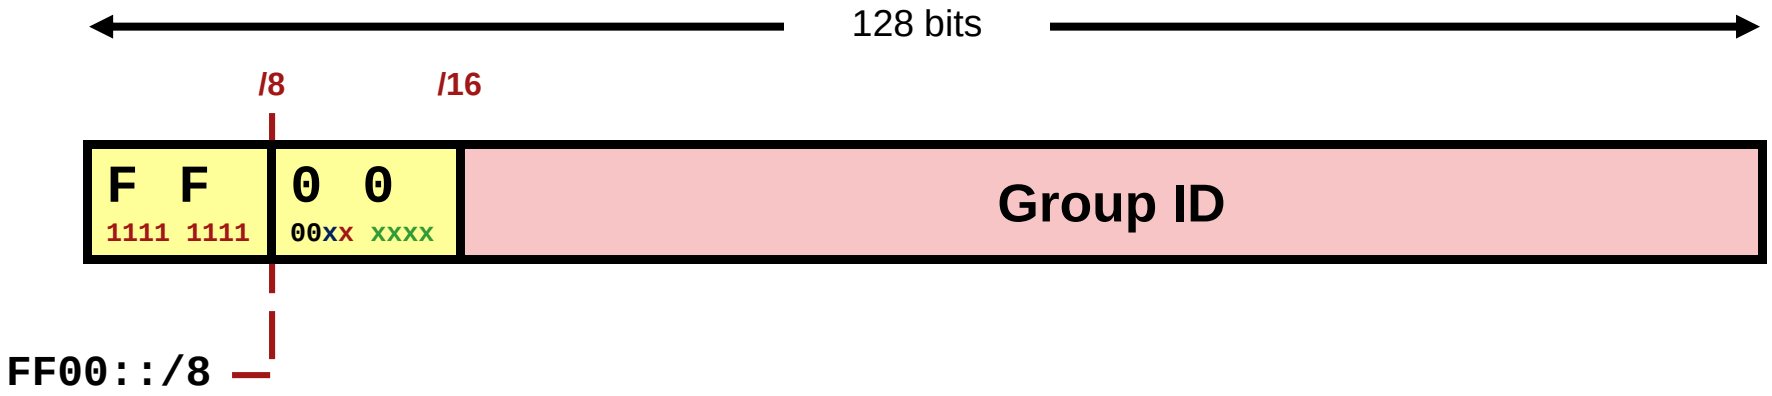

IPv6 Multicast Address

- The second octet of the address contains the prefix and transient (lifetime) flags, and the scope of the multicast address.

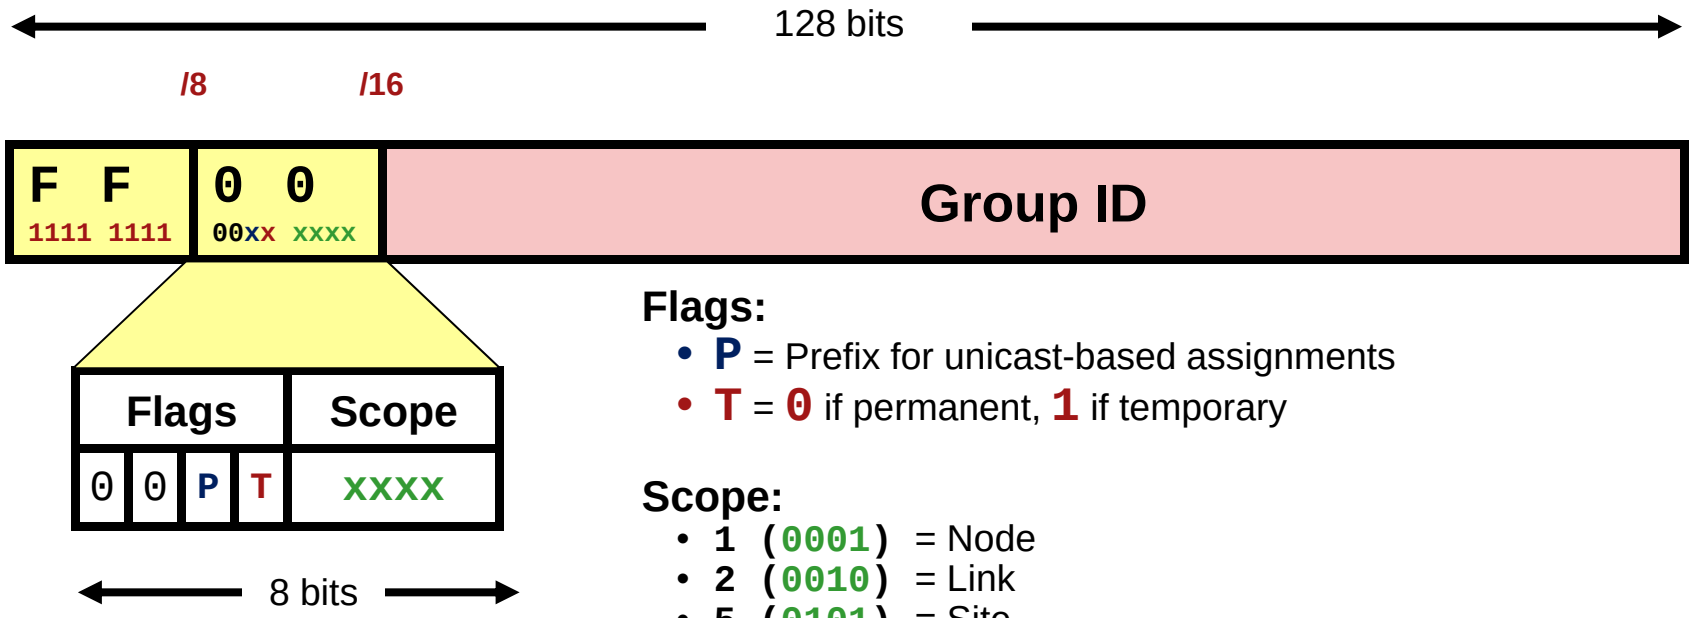

Flags:

- **P** = Prefix for unicast-based assignments
- **T** = **0** if permanent, **1** if temporary

Scope:

- **1 (0001)** = Node
- **2 (0010)** = Link
- **5 (0101)** = Site
- **8 (1000)** = Organization
- **E (1110)** = Global

IPv6 Multicast Address

- The multicast addresses **FF00::** to **FF0F::** have the **T** flag set to **0** and are therefore permanent and reserved.
- For example:
  - A multicast address starting with FF02::/16 is a permanent address.

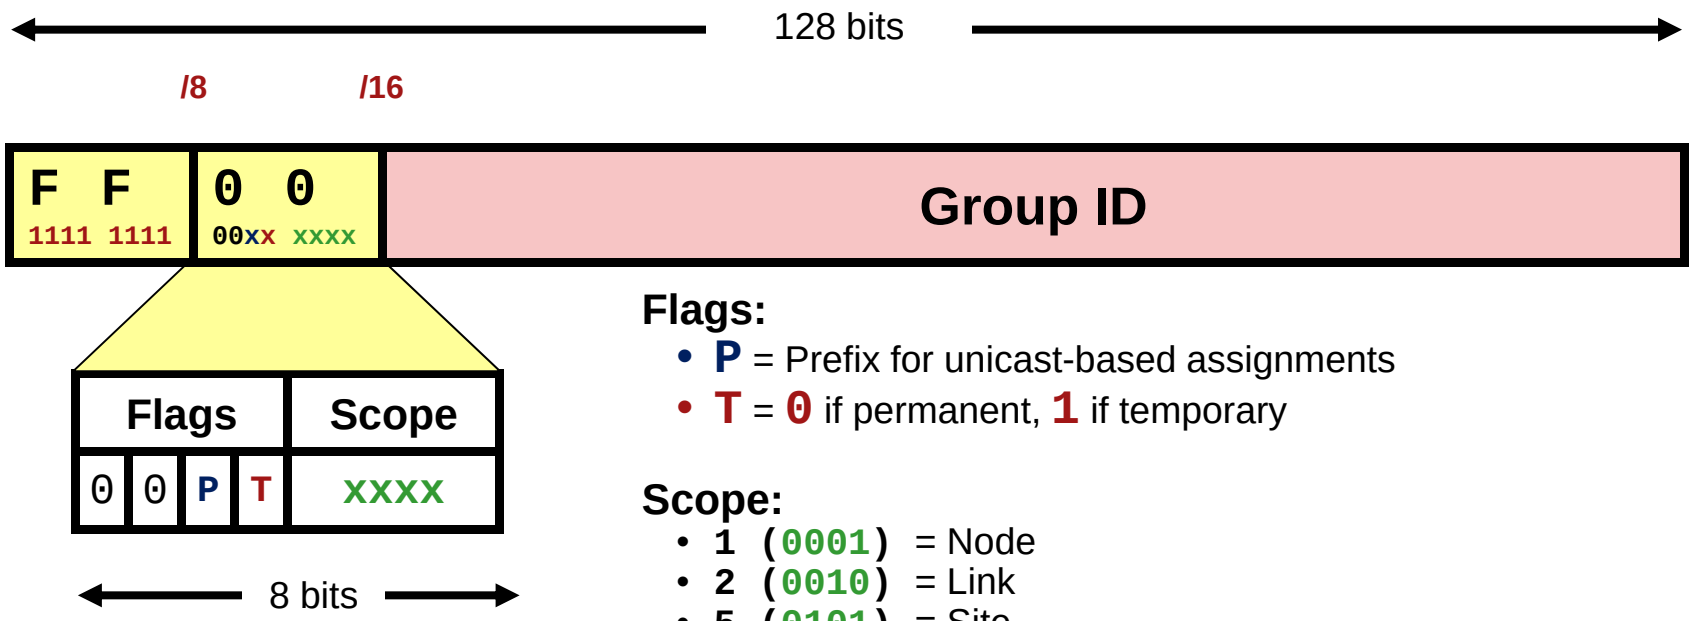

Flags:

- **P** = Prefix for unicast-based assignments
- **T** = **0** if permanent, **1** if temporary

Scope:

- **1** (**0001**) = Node
- **2** (**0010**) = Link
- **5** (**0101**) = Site
- **8** (**1000**) = Organization
- **E** (**1110**) = Global

Reserved IPv6 Multicast Addresses

Reserved Multicast Address	Description
FF02::1	• All nodes on a link (link-local scope).
FF02::2	• All routers on a link.
FF02::9	• All routing information protocol (RIP) routers on a link.
FF02::1:FFxx:xxxx	• All solicited-node multicast addresses used for host autoconfiguration and neighbor discovery (similar to ARP in IPv4). • The xx:xxxx is the far right 24 bits of the corresponding unicast or anycast address of the node.
FF05::101	• All Network Time Protocol (NTP) servers.

IPv6 Multicast Address Example

```
R1# show ipv6 interface loopback 100
Loopback100 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::222:55FF:FE18:7DE8
  No Virtual link-local address(es):
  Global unicast address(es):
    2001:8:85A3:4290:222:55FF:FE18:7DE8, subnet is 2001:8:85A3:4290::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF18:7DE8
  MTU is 1514 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ICMP unreachables are sent
  ND DAD is not supported
  ND reachable time is 30000 milliseconds (using 31238)
  Hosts use stateless autoconfig for addresses.
R1#
```

Solicited-Node Multicast Addresses

- The solicited-node multicast address (FF02::1:FF) is used for:
  - Neighbor discovery (ND) process
  - Stateless address autoconfiguration
- The Neighbor discovery (ND) process is used to:
  - Determine the local-link address of the neighbor.
  - Determine the routers on the link and default route.
  - Actively keep track of neighbor reachability.
  - Send network information from routers to hosts

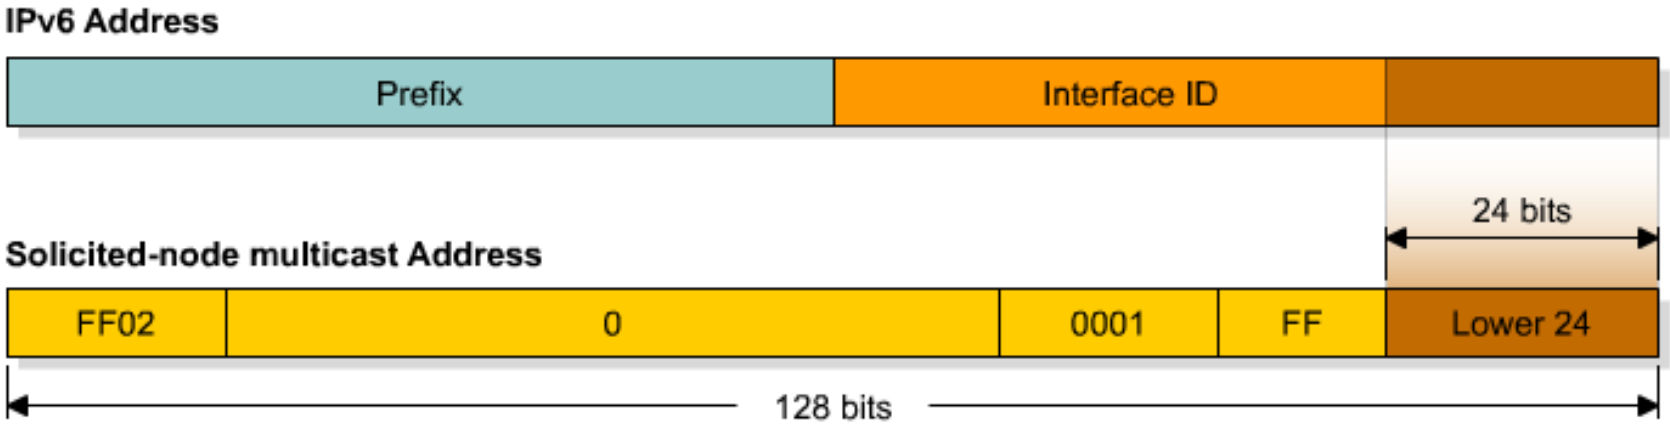

Solicited-Node Multicast Addresses

```
R1#sh ipv6 int e0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::200:CFF:FE3A:8B18
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FE3A:8B18
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
R1#
```

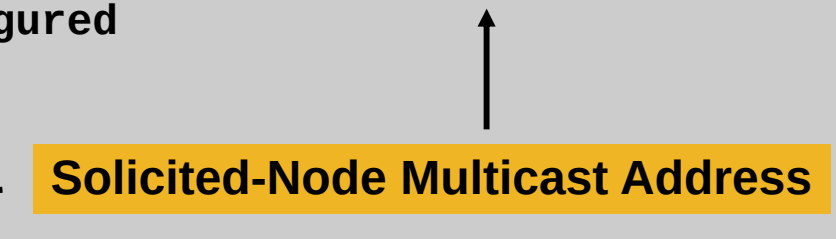

Neighbor Discovery ICMPv6 Packet Types

- Neighbor Discovery uses four ICMPv6 packet types:
  - Neighbor Solicitation and Neighbor Advertisement messages
  - Router Solicitation and Router Advertisement messages

ICMPv6 Message	Type	Description
Neighbor Solicitation (NS)	**135**	• Sent by a host to determine the link-layer address of a neighbor. • Used to verify that a neighbor is still reachable. • An NS is also used for Duplicate Address Detection (DAD).
Neighbor Advertisement (NA)	**136**	• A response to a NS message. • A node may also send unsolicited NA to announce a link-layer address change.
Router Advertisement (RA)	**134**	• RAs contain prefixes that are used for on-link determination or address configuration, a suggested hop limit value, MTU value, etc. • RAs are sent either periodically, or in response to a RS message.
Router Solicitation (RS)	**133**	• When a host is booting it send s out an RS requesting routers to immediately generate an RA rather than wait for their next scheduled time.

Neighbor Solicitation Example

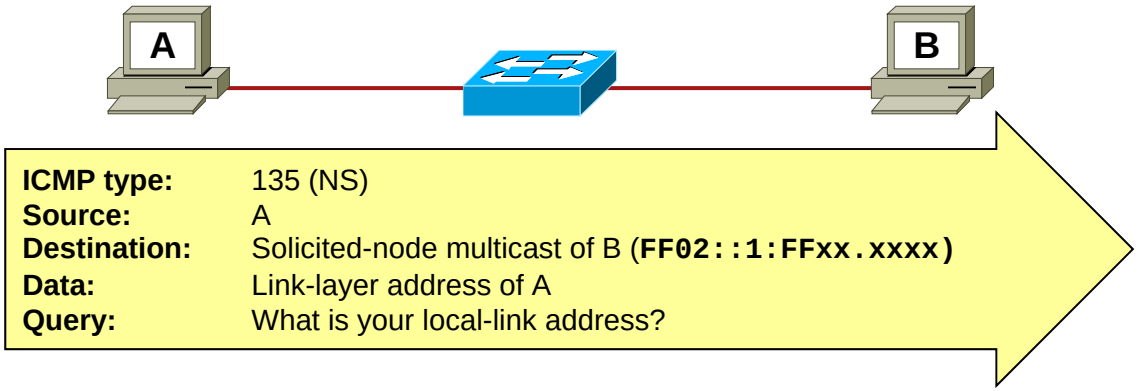

- ICMPv6 Neighbor Solicitation (NS) is similar to IPv4 ARP in that it is used when resolving an IPv6 address to a MAC address.
- For example, Host A needs to send a packet to Host B but needs the MAC address of host B.
  - Host A sends a Neighbor Solicitation (ICMPv6 message type 135) on the link.
  - The source address is the IPv6 address of the source node.

Neighbor Advertisement Example

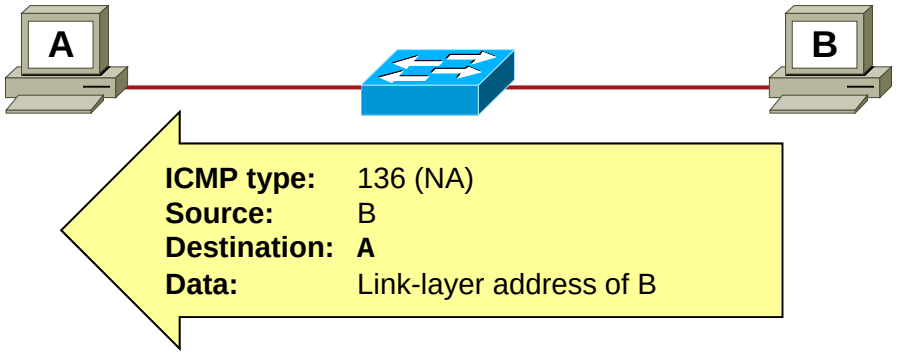

- Each destination node that receives the NS responds with an ICMPv6 message type 136, NA.
- The source address of this message is the IPv6 address of the responding node, and the destination address is the IPv6 address of the original source node (which sent the NS).
- The data portion includes the link-layer address of the destination node (even though the link-layer address is of course also included in the frame).

Solicited-Node Multicast Addresses

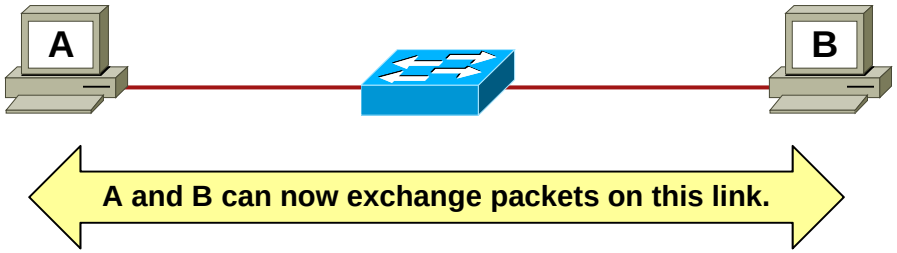

- The two devices can now communicate on the link because they know each other's link-layer addresses.

Solicited-Node Multicast Address Example

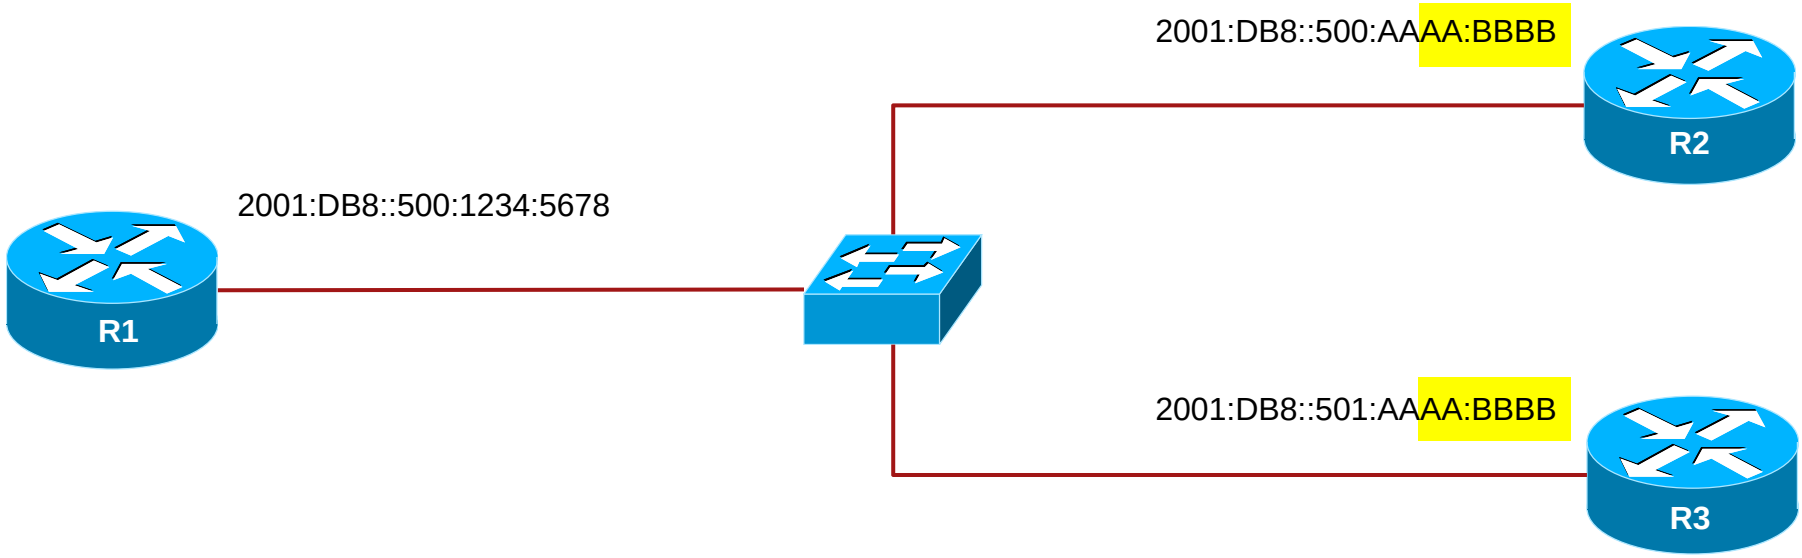

- What would happen if R1 wanted to exchange packets with R2?
- In this case R2 and R3 would have the same solicited-node multicast address of **FF02::1:FFAA:BBBB.**
  - Recall that a solicited-node address is **FF02::1:FFxx.xxxx** where the **xx:xxxx** is the far right 24 bits of the corresponding unicast or anycast address of the node.

Solicited-Node Multicast Address Example

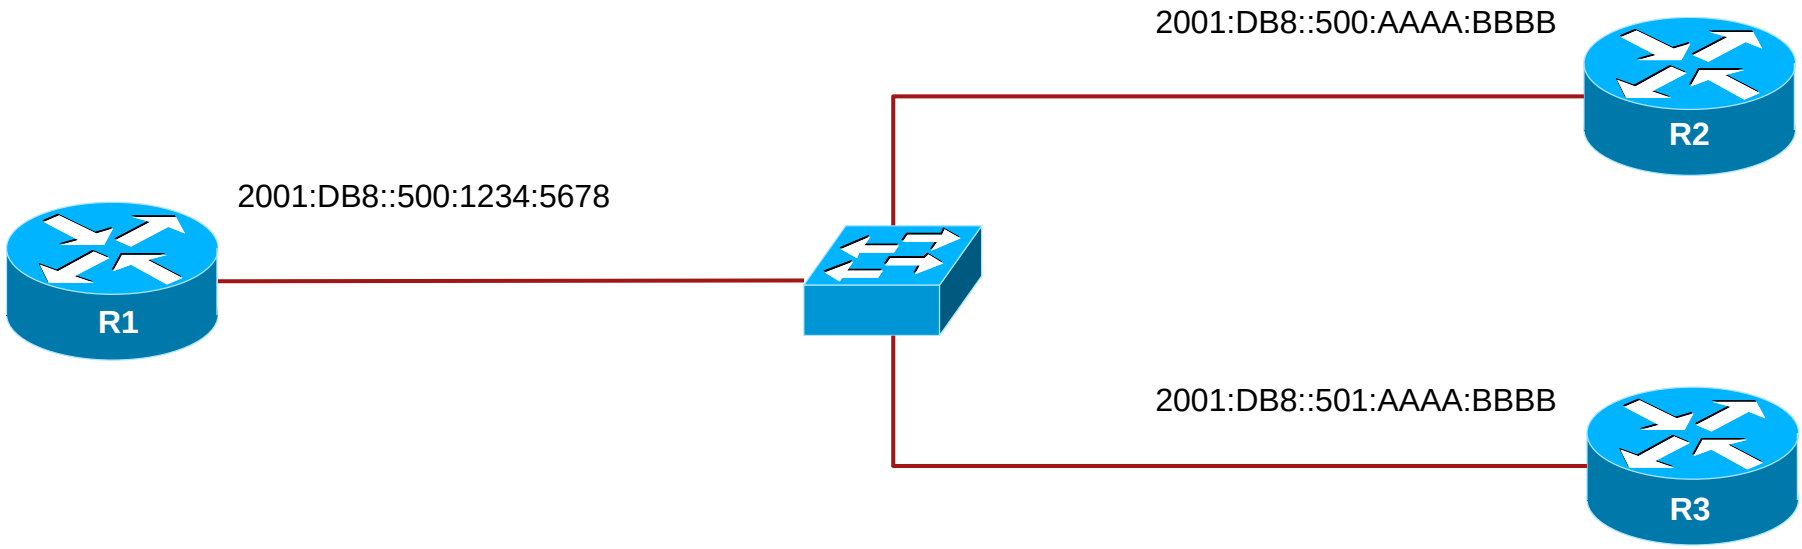

- When R1 desires to exchange packets with R2, R1 sends an NS message to the solicited-node multicast address of R2, (`FF02::1:FFAA:BBBB`).
- Along with other data, the NS message contains the “target address” which is the full IPv6 address that R1 is looking for (`2001:DB8::500:AAAA:BBBB`).

Solicited-Node Multicast Address Example

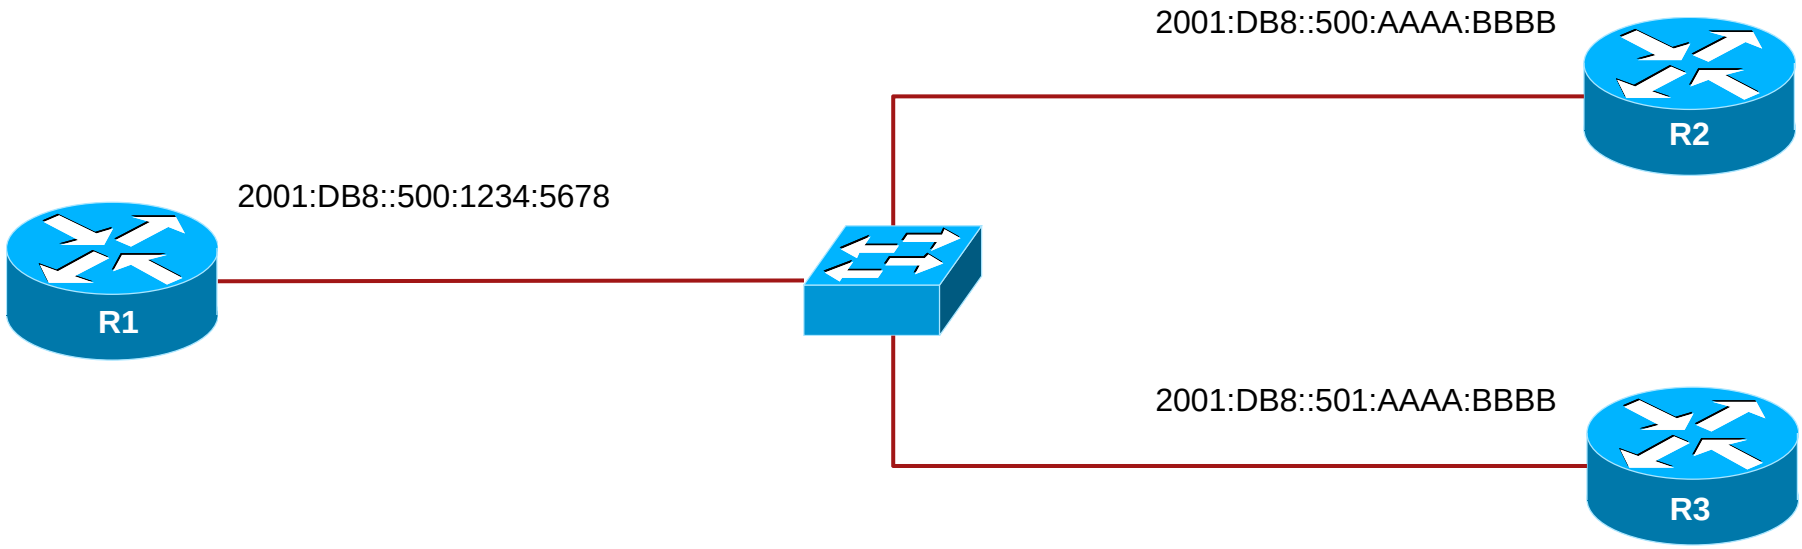

- Both R2 and R3 are listening to the same solicited-node multicast address (`FF02::1:FFAA:BBBB`)., so they both receive and process the packet.
  - R2 sees that the target address inside the packet is its own and responds with a neighbor advertisement (NA) that includes its MAC address.
  - R3 sees that the target address inside the packet is not its own and does not respond.

Stateless Autoconfiguration

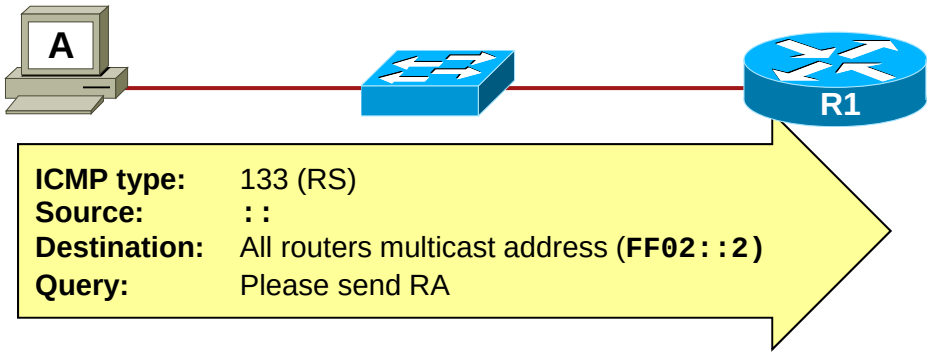

- Every IPv6 system (other than routers) is able to build its own unicast global address.
  - Enables new devices (e.g., cellular phones, wireless devices, home appliances, and home networks) to easily connect to the Internet.
- Stateless autoconfiguration uses the information in RA messages to configure hosts automatically.
- RAs are sent periodically, but a node can send out RS messages when it boots so that it doesn't have to wait for the next RA.

Stateless Autoconfiguration

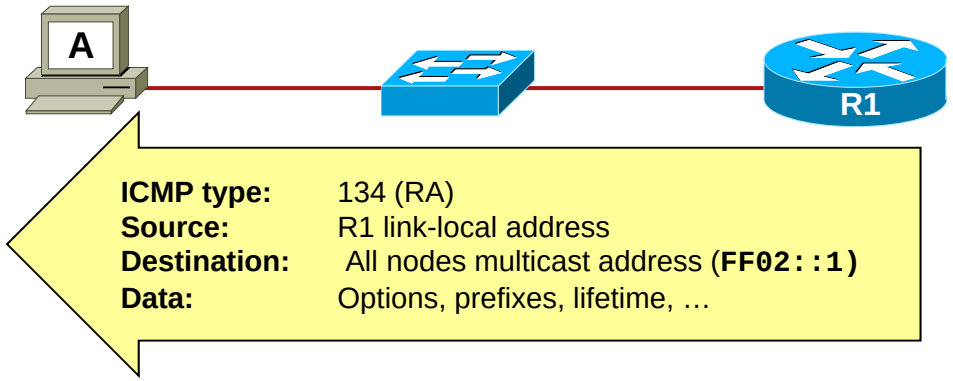

- All routers on the network reply to the RS immediately, with an RA sent to the all-nodes multicast address.
  - The prefix included in the RA is used as the /64 prefix for the host address.
  - The interface ID used is the EUI-64 format interface ID.

Stateless Autoconfiguration

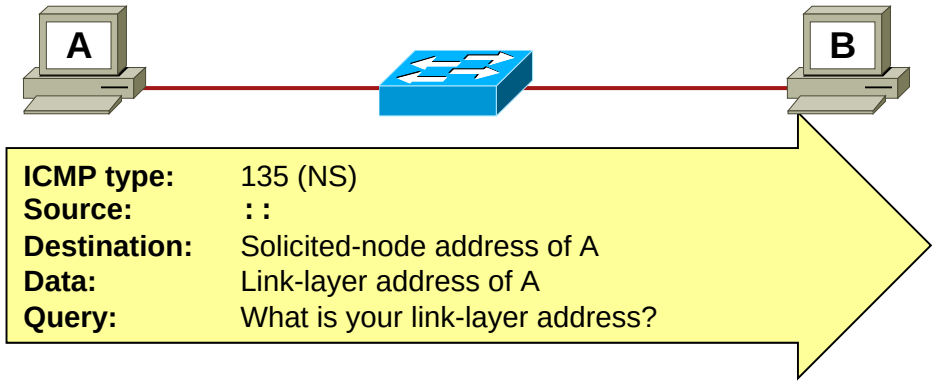

- The host now creates a link-local address and solicited-node address using the RA supplied by the router.
- Next it needs to verify that it's new IPv6 address is unique on the link using the Duplicate Address Detection (DAD) process.
  - DAD is used during the autoconfiguration process to ensure that no other device is using the autoconfiguration address.
- During the DAD phase, Host A sends an NS to query if another node on the link has the same IPv6 address.
  - If a node responds to the request, it means that the IPv6 address is already in use, and Host A needs to be manually configured.

IPv6 Routing

- IPv6 supports the following routing:
  - Static Routing
  - RIPng
  - OSPFv3
  - IS-IS for IPv6
  - EIGRP for IPv6
  - Multiprotocol BGP version 4 (MP-BGPv4)
- For each routing option above, the **`ipv6 unicast-routing`** command must be configured.

Static Routing

- Configured in the same way as IPv4.
- There is an IPv6-specific requirement per RFC 2461.
  - 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 basically means that using a global unicast address as a next-hop address with routing is not recommended.

RIPng

- Routing Information Protocol next generation (RIPng, RFC 2080) is a distance vector routing protocol for IPv6.
  - It's based on IPv4 RIP version 2 (RIPv2).
- It is similar to RIPv2 because:
  - The hop limit is still 15.
  - The administrative distance is still 120.
  - It still uses split horizon and poison reverse to prevent routing loops.
- Unlike RIPv2, RIPng is:
  - Used to transport IPv6 networks and prefixes.
    - It uses an IPv6 prefix and a next-hop IPv6 address.
  - Uses UDP port 520 (instead of UDP port 521).
  - Uses the multicast group FF02::9 (instead of 224.0.0.9).

OSPFv3

- The protocol implementation for IPv6 includes these characteristics:
  - Based on OSPF version 2 (OSPFv2), with enhancements
  - Distributes IPv6 prefixes
  - Runs directly over IPv6
  - Operates as “ships in the night” with OSPFv2
- This implementation adds these IPv6-specific attributes:
  - 128-bit addresses
  - Link-local address
  - Multiple addresses and instances per interface
  - Authentication (now uses IPsec)
  - OSPFv3 runs over a link rather than a subnet

OSPFv3

- Open Shortest Path First version 3 (OSPFv3 RFC 5340) is a link state routing protocol for IPv6.
  - It's based on OSPFv2.
- The following remained the same as OSPFv2:
  - Packet types (Hello, DBD, LSR, LSU, LSA)
  - Mechanisms for neighbor discovery and adjacency formation
  - LSA flooding and aging (but there are now 3 types of scopes)
  - SPF calculations
  - DR election procedure
  - Multi-area support (including NSSA)
  - Multiple topologies support (NBMA, point-to-multipoint, point-to-point and broadcast)
  - Router-ID is still a 32-bit address

OSPFv2 and OSPFv3 Differences

- Unlike OSPFv2, OSPFv3:
  - OSPFv3 runs over a link and is configured on an interface.
  - Uses the term “link” similarly to IPv4 OSPF's “subnet” or “network”.
  - IPv6 link-local addresses are required.
  - There are now three separate LSA flooding scopes: Link-local scope, Area scope, and AS scope.
  - Multiple OSPFv3 instances are supported on one interface.
  - Multicast addresses have changed.
  - Security is improved.

Link-Local Addresses Are Used

- OSPFv3 uses IPv6 link-local addresses to identify the OSPFv3 adjacency neighbors.
- Therefore, when configuring the **`ipv6 ospf neighbor`** command, the IPv6 address used must be the link-local address of the neighbor.

IPv4 to IPv6 Transition Mechanisms

- The transition from IPv4 to IPv6 does not require an upgrade on all nodes at the same time.
  - IPv4 and IPv6 will coexist for some time.
- A wide range of techniques are available for the period of transition between IPv4 and IPv6.
- These techniques can be grouped into three categories:
  - **Dual-stack techniques**
  - **Tunneling techniques**
  - **Translation techniques**

Dual-Stack Techniques

- Hosts and network devices run both IPv4 and IPv6 at the same time.
  - This technique is useful as a temporary transition, but it adds overhead and uses many resources.
- Cisco IOS Software is IPv6 ready.
  - As soon as IPv4 and IPv6 configurations are complete, the interface is dual stacked and it forwards both IPv4 and IPv6 traffic.
- Drawback of dual stacking includes:
  - The additional resources required to keep and process dual routing tables, routing protocol topology tables, etc.
  - The higher administrative overhead, troubleshooting, and monitoring, is more complex.

Dual-Stack Example


```
R1(config)# interface fa0/0
R1(config-if)# ip address 10.10.10.1 255.255.255.0
R1(config-if)# ipv6 address 2001:12::1/64
R1(config-if)# ^Z
R1#
```

- The FastEthernet 0/0 interface of R1 is dual stacked.
  - It is configured with an IPv4 and an IPv6 address.
  - Also notice that for each protocol, the addresses on R1 and R2 are on the same network.

Tunneling Techniques

- Isolated IPv6 networks are connected over an IPv4 infrastructure using tunnels.
- The edge devices are the only ones that need to be dual-stacked.
- Scalability may be an issue if many tunnels need to be created.
  - Tunnels can be either manually or automatically configured, depending on the scale required and administrative overhead tolerated.

Tunneling Techniques

- For IPv6, tunneling is an integration method in which an IPv6 packet is encapsulated within IPv4.
- This enables the connection of IPv6 islands without the need to convert the intermediary network to IPv6.

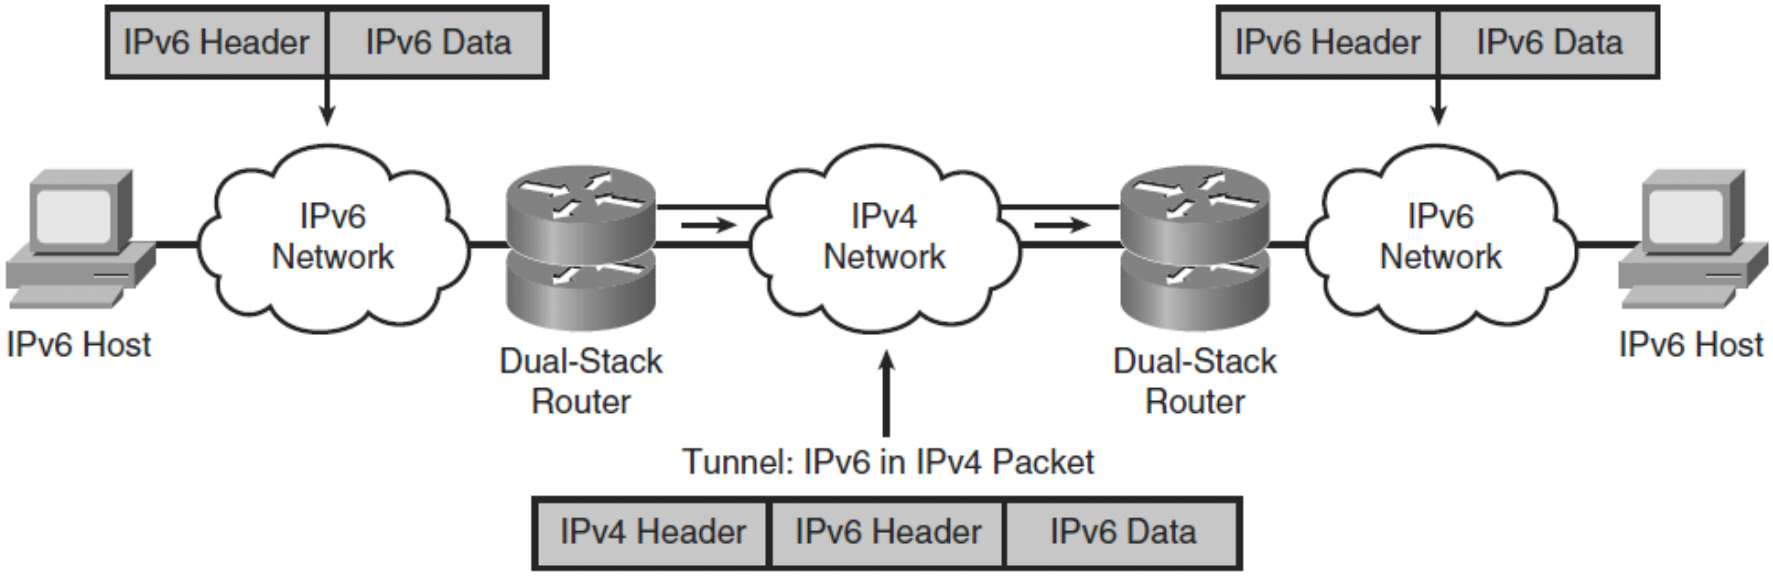

Tunneling Techniques

- In this example, the tunnel between sites is using:
  - IPv4 as the transport protocol (the protocol over which the tunnel is created).
  - IPv6 is the passenger protocol (the protocol encapsulated in the tunnel and carried through the tunnel).
  - GRE is used to create the tunnel, and is known as the tunneling protocol.

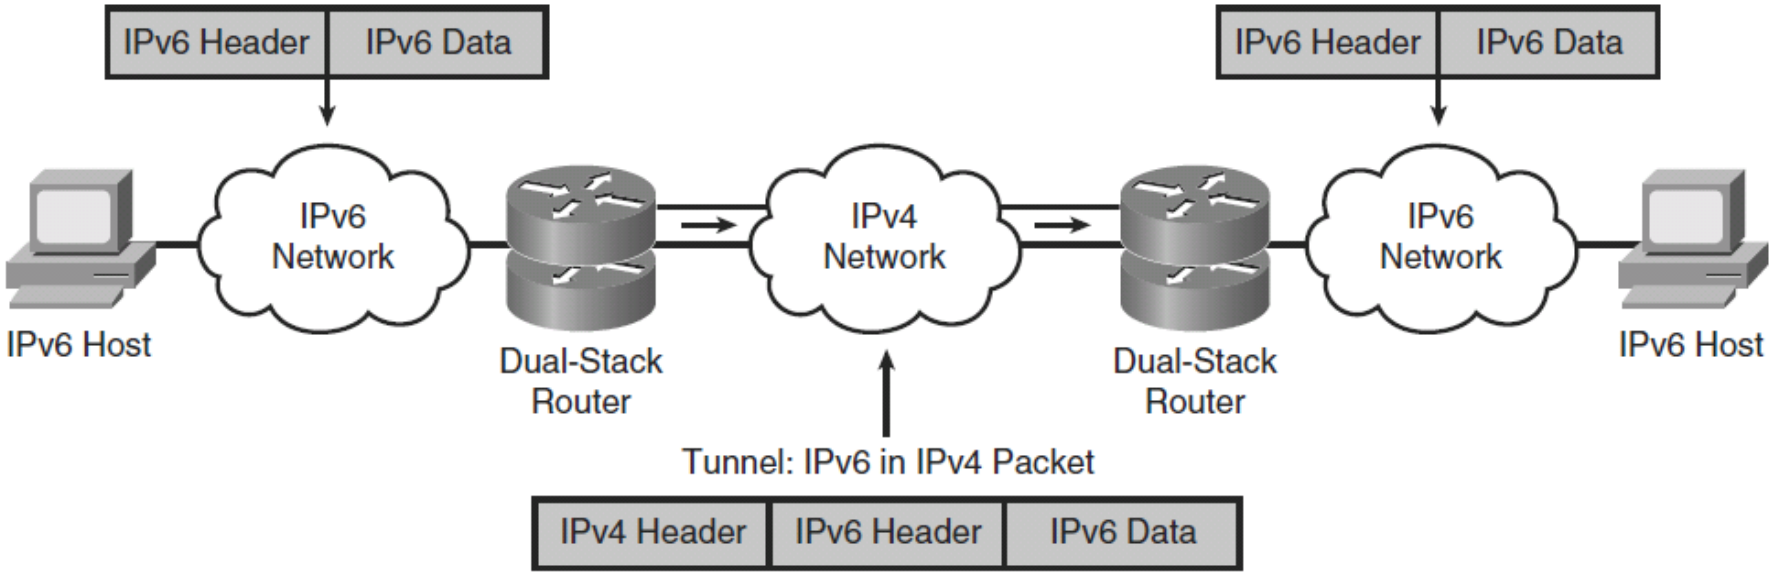

Translation Techniques

- Dual stack and tunneling techniques manage the interconnection of IPv6 domains.
- NAT-PT is an extension of NAT techniques and it provides protocol translation services for legacy equipment that cannot be upgraded to IPv6 and for some deployment scenarios.
- A router converts IPv6 packets into IPv4 packets and vise versa, allowing IPv6-only devices to communicate with IPv4-only devices.
  - Scalability may again be an issue because of the resources required on the translator device.

NAT-PT Example

- Node A is an IPv6 only node and wants to send an IPv6 datagram to node D and therefore forwards the packet to the NAT-PT router.
  - The NAT-PT router maintains a pool of globally routable IPv4 addresses that are assigned to IPv6 nodes dynamically as sessions are initiated.
- An advantage of NAT-PT is that no modifications are required on the hosts.

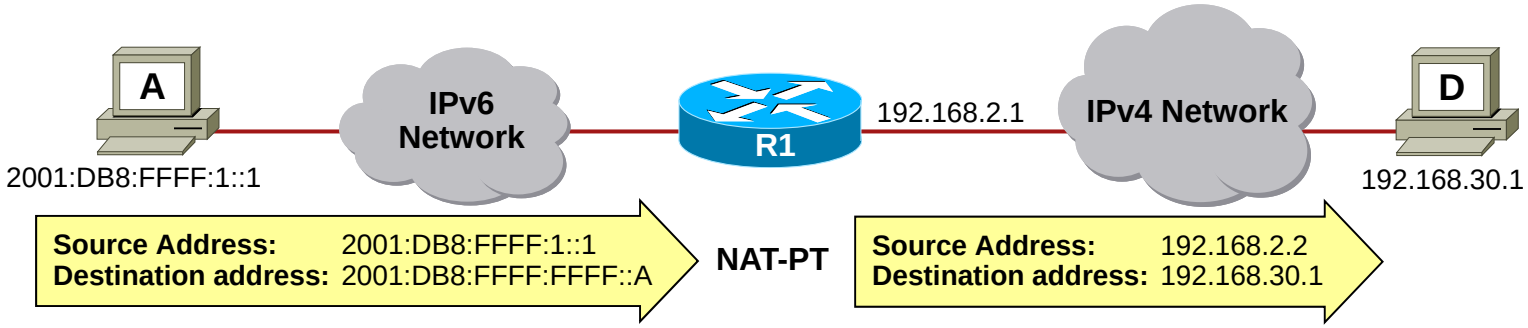

Types of Tunnels

- Tunnels can be created manually using:
  - Manual IPv6 tunnels
  - GRE IPv6 tunnels
- Tunnels can also be created automatically using:
  - IPv4-Compatible IPv6 Tunnels (now deprecated)
  - 6to4 tunnels
  - ISATAP Tunnels