



Introducing IPv6 and Defining IPv6 Addressing



BSCI Module 8 Lessons 1 and 2

Objectives

- Explain the need for IPv6 address space.
- Explain how IPv6 deals with the limitations of IPv4.
- Describe the features of IPv6 addressing.
- Describe the structure of IPv6 headers in terms of format and extension headers.
- Show how an IPv6 address is represented.
- Describe the three address types used in IPv6.

IPv6 Advanced Features

Larger address space

- Global reachability and flexibility
- Aggregation
- Multihoming
- Autoconfiguration
- Plug-and-play
- End to end without NAT
- Renumbering

Simpler header

- Routing efficiency
- Performance and forwarding rate scalability
- No broadcasts
- No checksums
- Extension headers
- Flow labels

IPv6 Advanced Features (Cont.)

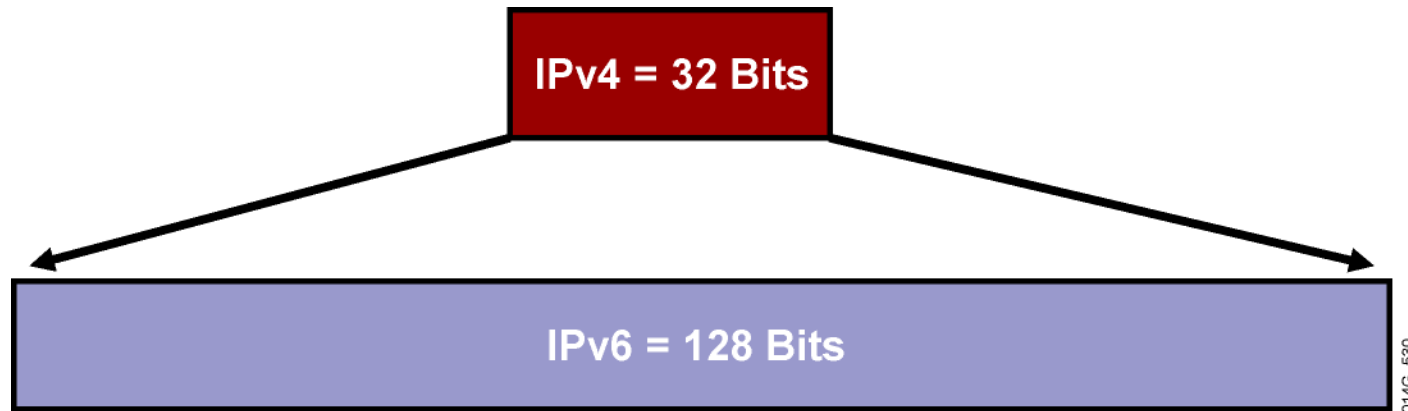
Mobility and security

- Mobile IP RFC-compliant
- IPSec mandatory (or native) for IPv6

Transition richness

- Dual stack
- 6to4 tunnels
- Translation

Larger Address Space




IPv4

- 32 bits or 4 bytes long
 - ≅ 4,200,000,000 possible addressable nodes

IPv6

- 128 bits or 16 bytes: four times the bits of IPv4
 - ≅ $3.4 * 10^{38}$ possible addressable nodes
 - ≅ 340,282,366,920,938,463,374,607,432,768,211,456
 - ≅ $5 * 10^{28}$ addresses per person


Larger Address Space Enables Address Aggregation




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

Simple and Efficient Header

IPv4 Header



IPv6 Header

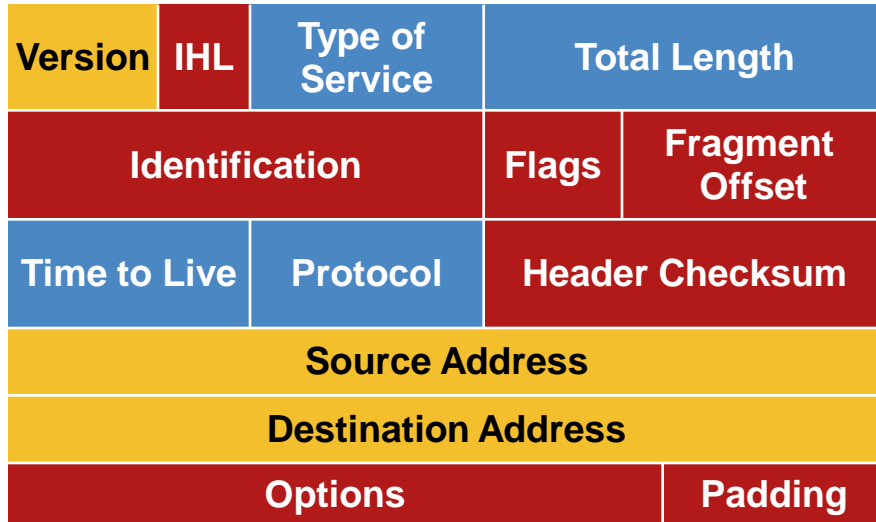


A simpler and more efficient header means:

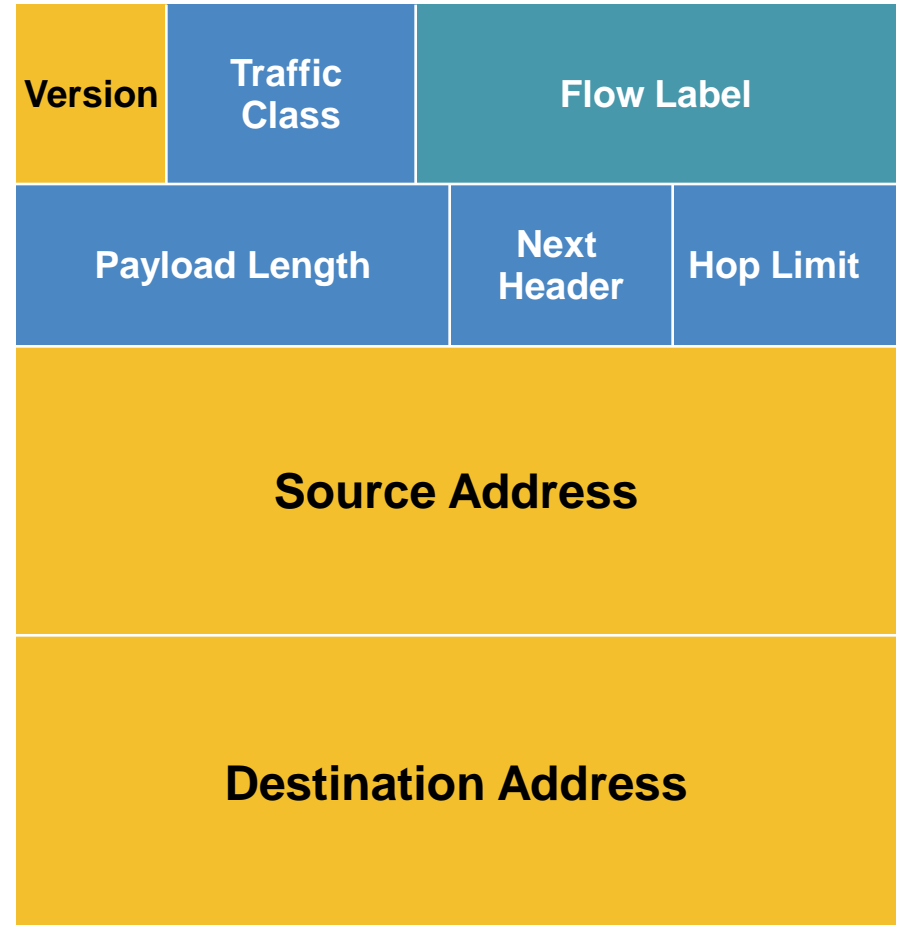
- 64-bit aligned fields and fewer fields
- Hardware-based, efficient processing
- Improved routing efficiency and performance
- Faster forwarding rate with better scalability

IPv4 and IPv6 Header Comparison

IPv4 Header



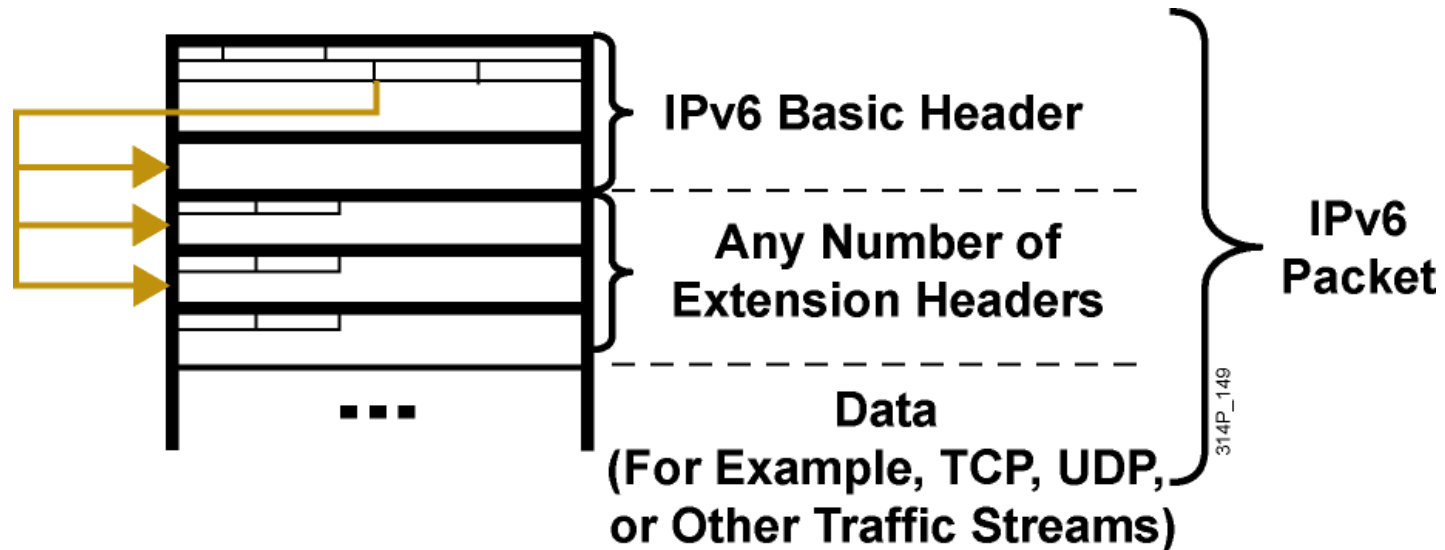
IPv6 Header



Legend

- Field's Name Kept from IPv4 to IPv6
- Fields Not Kept in IPv6
- Name and Position Changed in IPv6
- New Field in IPv6

IPv6 Extension Headers



Simpler and more efficient header means:

- IPv6 has extension headers.
- IPv6 handles the options more efficiently.
- IPv6 enables faster forwarding rate and end nodes processing.

IPv6 Address Representation

- **x:x:x:x:x:x:x**, where **x** is a 16-bit hexadecimal field
- Leading zeros in a field are optional:
2031:0:130F:0:0:9C0:876A:130B
- Successive fields of 0 can be represented as ::, but only once per address.

Examples:

2031:0000:130F:0000:0000:09C0:876A:130B

2031:0:130f::9c0:876a:130b


FF01:0:0:0:0:0:0:1 >>> FF01::1

0:0:0:0:0:0:0:1 >>> ::1

0:0:0:0:0:0:0:0 >>> ::

IPv6—Addressing Model

- Addresses are assigned to interfaces
 - Change from IPv4 mode:
- Interface “expected” to have multiple addresses
- Addresses have scope
 - Link Local
 - Unique Local
 - Global
- Addresses have lifetime
 - Valid and preferred lifetime



IPv6 Address Types

- Unicast

 - Address is for a single interface.

 - IPv6 has several types (for example, global and IPv4 mapped).

- Multicast

 - One-to-many

 - Enables more efficient use of the network

 - Uses a larger address range

- Anycast

 - One-to-nearest (allocated from unicast address space).

 - Multiple devices share the same address.

 - All anycast nodes should provide uniform service.

 - Source devices send packets to anycast address.

 - Routers decide on closest device to reach that destination.


 - Suitable for load balancing and content delivery services.

IPv6 Global Unicast (and Anycast) Addresses

The global unicast and the anycast share the same address format.

- 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, multicast).
- Every IPv6-enabled interface must contain at least one loopback (::1/128) and one link-local address.
- Optionally, every interface can have multiple unique local and global addresses.
- Anycast address is a global unicast address assigned to a set of interfaces (typically on different nodes).
- IPv6 anycast is used for a network multihomed to several ISPs that have multiple connections to each other.

IPv6 Global Unicast Addresses (Cont.)



- Global unicast and anycast addresses are defined by a global routing prefix, a subnet ID, and an interface ID.

IPv6 Unicast Addressing

- IPv6 addressing rules are covered by multiple RFCs.
Architecture defined by RFC 4291.
- Unicast: One to one
 - Global
 - Link local (FE80::/10)
- A single interface may be assigned multiple IPv6 addresses of any type: unicast, anycast, or multicast.

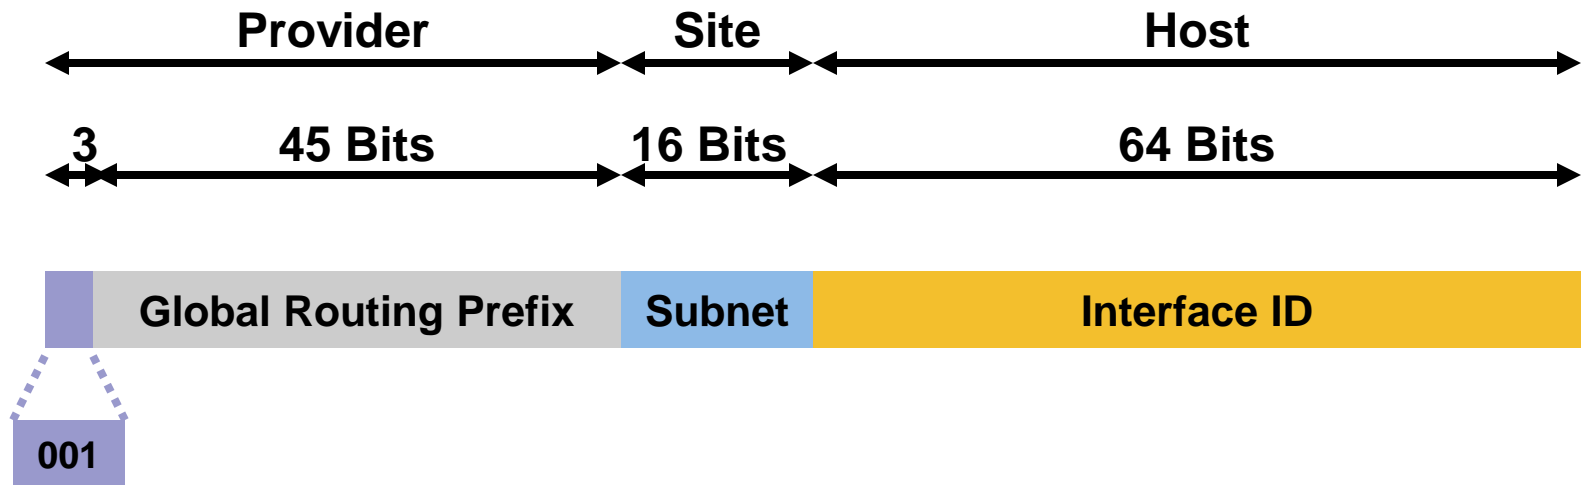


Implementing Dynamic IPv6 Addresses



BSCI Module 8

Aggregatable Global Unicast Addresses



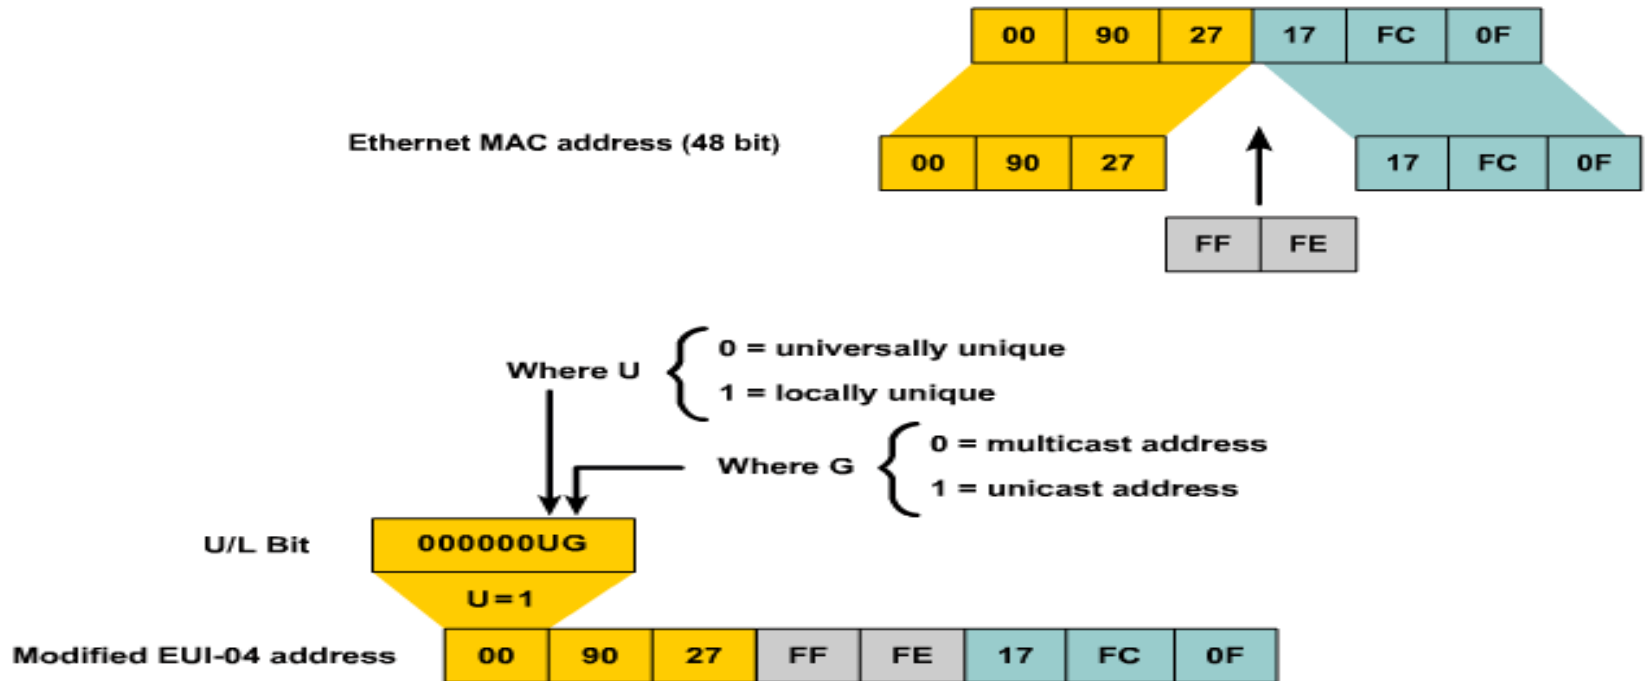
Aggregatable Global Unicast Addresses Are:

- Addresses for generic use of IPv6
- Structured as a hierarchy to keep the aggregation

IPv6 Interface ID

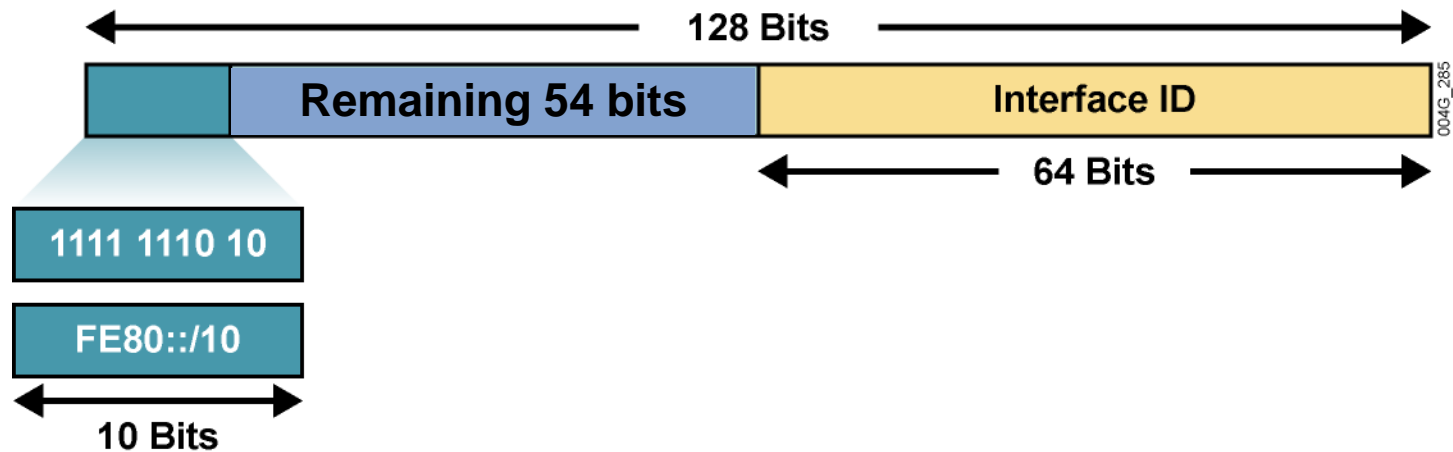
- Cisco uses the extended universal identifier (EUI)-64 format to do stateless autoconfiguration.
- This format expands the 48-bit MAC address to 64 bits by inserting “FFFE” into the middle 16 bits.
- To make sure that the chosen address is from a unique Ethernet MAC address, the universal/local (U/L bit) is set to 1 for global scope (0 for local scope).

MAC Address to EUI-64



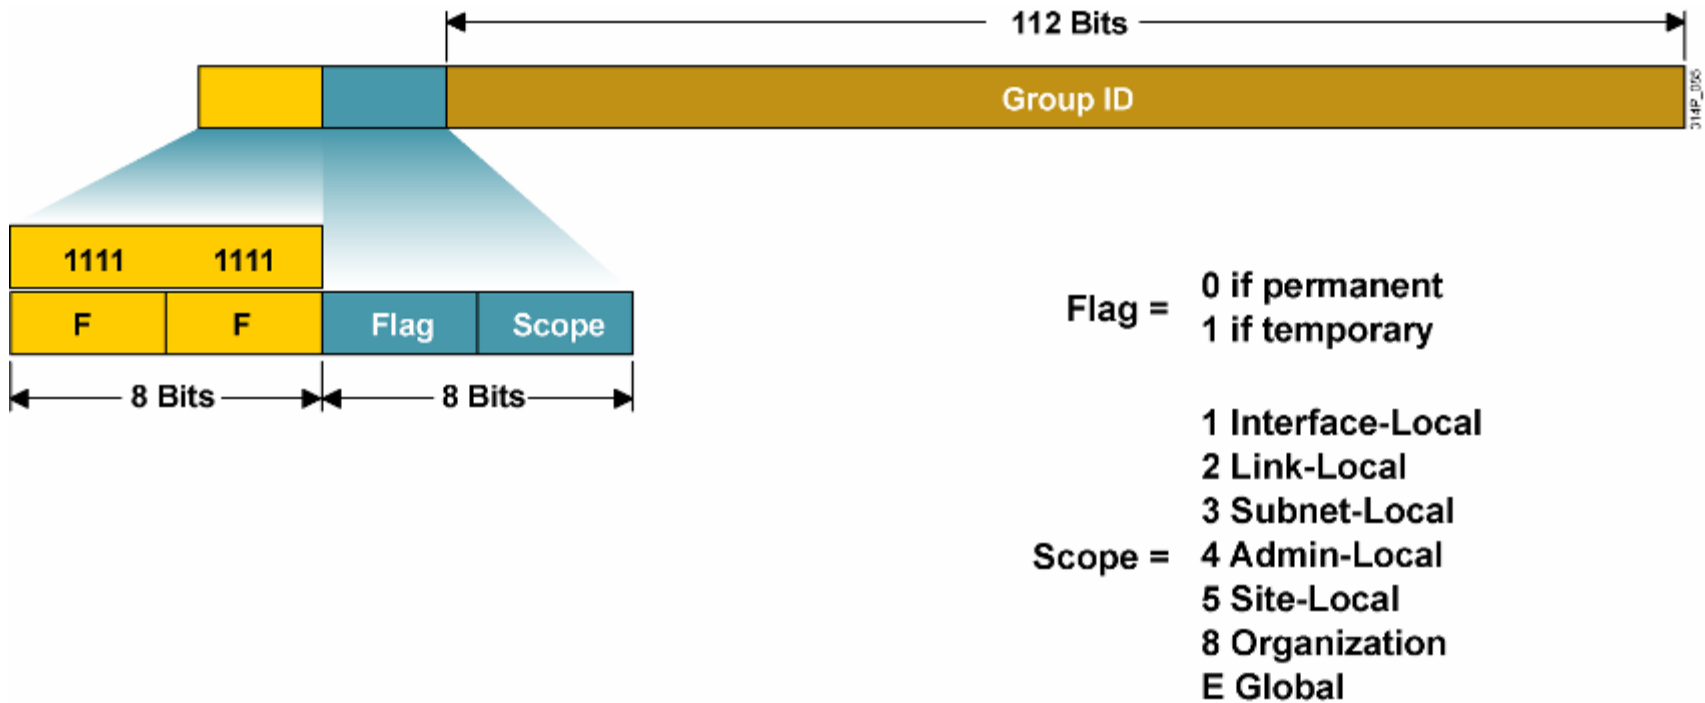
- This format expands the 48-bit Ethernet MAC address format to a 64-bit version by inserting "FFFE" in the middle of the 48 bits.
- This creates a 64-bit, unique interface identifier.
- The 7th bit (starting with the leftmost bit as "1") in an IPv6 interface identifier is referred to as the Universal/Local bit, or U/L bit.

Link-Local Address








- Mandatory address for communication between two IPv6 devices (similar to ARP but at Layer 3)
- Automatically assigned by router as soon as IPv6 is enabled
- Also used for next-hop calculation in routing protocols
- Only link specific scope
- Remaining 54 bits could be zero or any manual configured value

Multicasting




- Multicast is frequently used in IPv6 and replaces broadcast.

Examples of Permanent Multicast Addresses

	Meaning	Scope
FF02::1	All nodes 	Link-local
FF02::2	All routers 	Link-local
FF02::9	All RIP routers 	Link-local
FF02::1:FFXX:XXXX	Solicited-node 	Link-local
FF05::101	All NTP servers 	Site-local

Solicited-Node Multicast Address

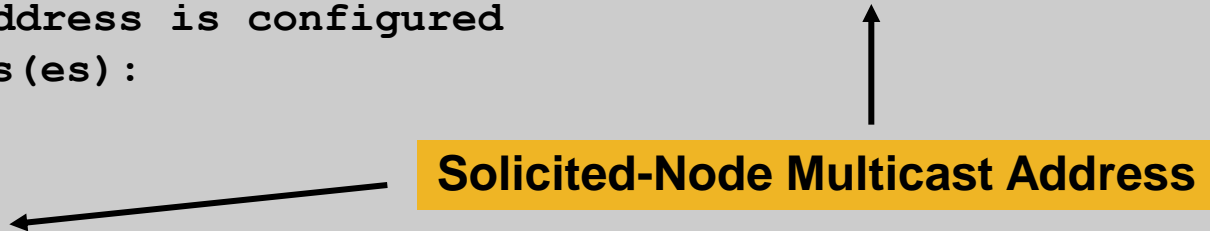
IPv6 Address



- Used in neighbor solicitation messages
- Multicast address with a link-local scope
- Solicited-node multicast consists of prefix + lower 24 bits from unicast, FF02::1:FF:


Router Interface

```
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#
```



Solicited-Node Multicast Address


Anycast



- An IPv6 anycast address is a global unicast address that is assigned to more than one interface.


Stateless Autoconfiguration

Interface Identifier ::2004:0FD1:9CAA:1002



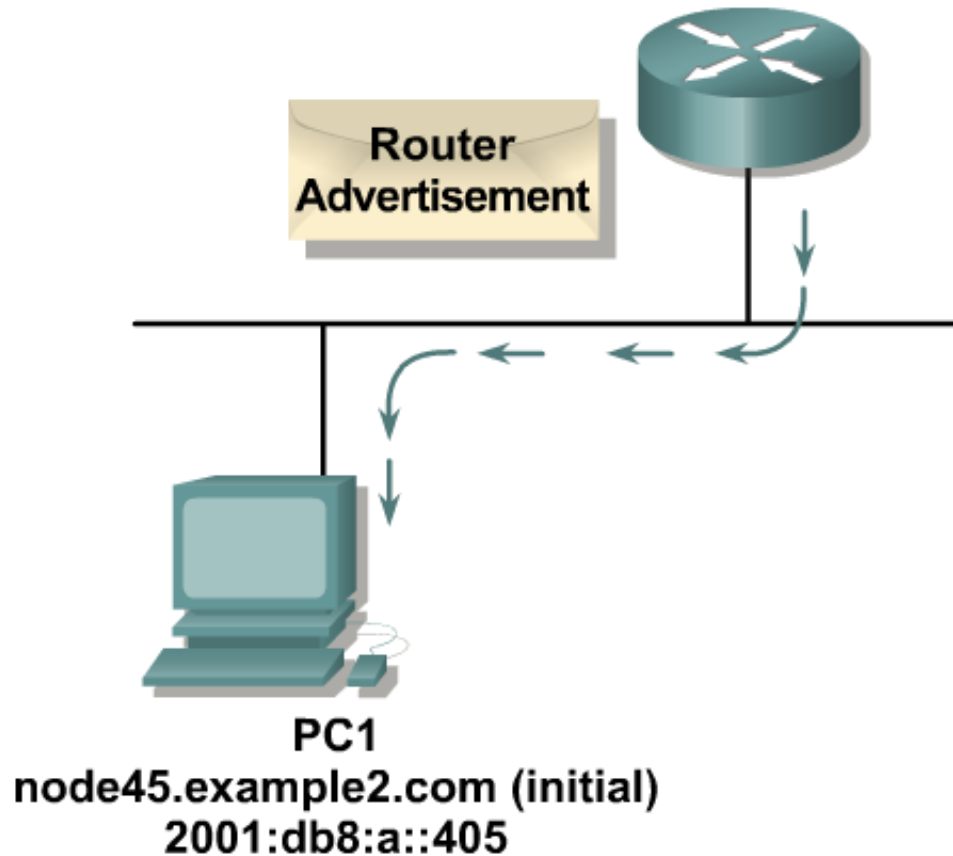
- A router sends network information to all the nodes on the local link.
- A host can autoconfigure itself by appending its IPv6 interface identifier (64-bit format) to the local link prefix (64 bits).
- The result is a full 128-bit address that is usable and guaranteed to be globally unique.

A Standard Stateless Autoconfiguration




- Stage 1: The PC sends a router solicitation to request a prefix for stateless autoconfiguration.

A Standard Stateless Autoconfiguration (Cont.)



- Stage 2: The router replies with a router advertisement.


IPv6 Mobility



IPv6 Routing Protocols



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

- **ipv6 unicast-routing** command is required to enable IPv6 before any routing protocol configured.

RIPng

- Same as IPv4:

 - Distance-vector, radius of 15 hops, split-horizon, and poison reverse

 - Based on RIPv2

- Updated features for IPv6:

 - IPv6 prefix, next-hop IPv6 address

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

 - Uses IPv6 for transport

 - Named RIPng

Integrated Intermediate System-to-Intermediate System (IS-IS)

- Same as for IPv4.
- Extensions for IPv6:
 - 2 new type-length-values (TLV):
 - IPv6 reachability (with 128 bits prefix)
 - IPv6 interface address (with 128 bits)
 - New protocol identifier
 - Not yet an IETF standard

Multiprotocol Border Gateway Protocol (MP-BGP) (RFC 2858)

- Multiprotocol extensions for BGPv4:
 - Enables protocols other than IPv4.
 - New identifier for the address family.
- IPv6 specific extensions:
 - Scoped addresses: NEXT_HOP contains a global IPv6 address and potentially a link-local address (only when there is a link-local reachability with the peer).
 - NEXT_HOP and NLRI (Network Layer Reachability Information) are expressed as IPv6 addresses and prefix in the multiprotocol attributes.

OSPF Version 3 (OSPFv3) (RFC 2740)

- Similar to OSPF for IPv4:
 - Same mechanisms, but a major rewrite of the internals of the protocol
- Updated features for IPv6:
 - Every IPv4-specific semantic is removed
 - Carry IPv6 addresses
 - Link-local addresses used as source
 - IPv6 transport
 - OSPF for IPv6 is currently an IETF proposed standard

OSPFv3



OSPFv3—Hierarchical Structure


- Topology of an area is invisible from outside of the area:

LSA flooding is bounded by area.

SPF calculation is performed separately for each area.

- Backbones must be contiguous.
- All areas must have a connection to the backbone:

Otherwise a virtual link must be used to connect to the backbone.



OSPFv3—Similarities with OSPFv2

- OSPFv3 is OSPF for IPv6 (RFC 2740):
 - Based on OSPFv2, with enhancements
 - Distributes IPv6 prefixes
 - Runs directly over IPv6
- OSPFv3 & v2 can be run concurrently, because each address family has a separate SPF (ships in the night).
- OSPFv3 uses the same basic packet types as OSPFv2:
 - Hello
 - Database description blocks (DDB)
 - Link state request (LSR)
 - Link state update (LSU)
 - Link state acknowledgement (ACK)

Enhanced Routing Protocol Support Differences from OSPFv2

- OSPFv3 has the same five packet types, but some fields have been changed.

Packet Type	Description
1	Hello
2	Database Description
3	Link State Request
4	Link State Update
5	Link State Acknowledgement

- All OSPFv3 packets have a 16-byte header verses the 24-byte header in OSPFv2.

Version	Type	Packet Length
Router ID		
Area ID		
Checksum	Autype	
Authentication		
Authentication		

Version	Type	Packet Length
Router ID		
Area ID		
Checksum	Instance ID	0

314P_110

OSPFv3—Differences from OSPFv2

OSPFv3 protocol processing per-link, not per-subnet:

- IPv6 connects interfaces to links.
- Multiple IPv6 subnets can be assigned to a single link.
- Two nodes can talk directly over a single link, even though they do not share a common subnet.
- The terms “network” and “subnet” are being replaced with “link”.
- An OSPF interface now connects to a link instead of a subnet.

OSPFv3—Differences from OSPFv2

Multiple OSPFv3 protocol instances can now run over a single link:

- This allows for separate autonomous systems, each running OSPF, to use a common link. A single link could belong to multiple areas.
- Instance ID is a new field that is used to have multiple OSPFv3 protocol instances per link.
- In order to have two instances talk to each other, they need to have the same instance ID. By default it is 0, and for any additional instance it is increased.

OSPFv3—Differences from OSPFv2

- Multicast addresses:

 - FF02::5 – Represents all SPF routers on the link local scope; equivalent to 224.0.0.5 in OSPFv2.

 - FF02::6 – Represents all DR routers on the link local scope; equivalent to 224.0.0.6 in OSPFv2.

- Removal of address semantics:

 - IPv6 addresses are no longer present in OSPF packet header (part of payload information).

 - Router LSA and network LSA do not carry IPv6 addresses.

 - Router ID, area ID, and link-state ID remains at 32 bits.

 - DR and BDR are now identified by their router ID and no longer by their IP address.

- Security:

 - OSPFv3 uses IPv6 AH and ESP extension headers instead of variety of mechanisms defined in OSPFv2.


LSA Overview

	LSA Function Code	LSA Type
Router-LSA	1	Ox2001
Network-LSA	2	Ox2002
Inter-Area-Prefix-LXA	3	Ox2003
Inter-Area-Router-LSA	4	Ox2004
AS-External-LSA	5	Ox4005
Group-Membership-LSA	6	Ox2006
Type-7-LSA	7	Ox2007
Link-LSA	8	Ox2008
Intra-Area-Prefix-LSA	9	Ox2009

Renamed

New

Larger Address Space Enables Address Aggregation



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

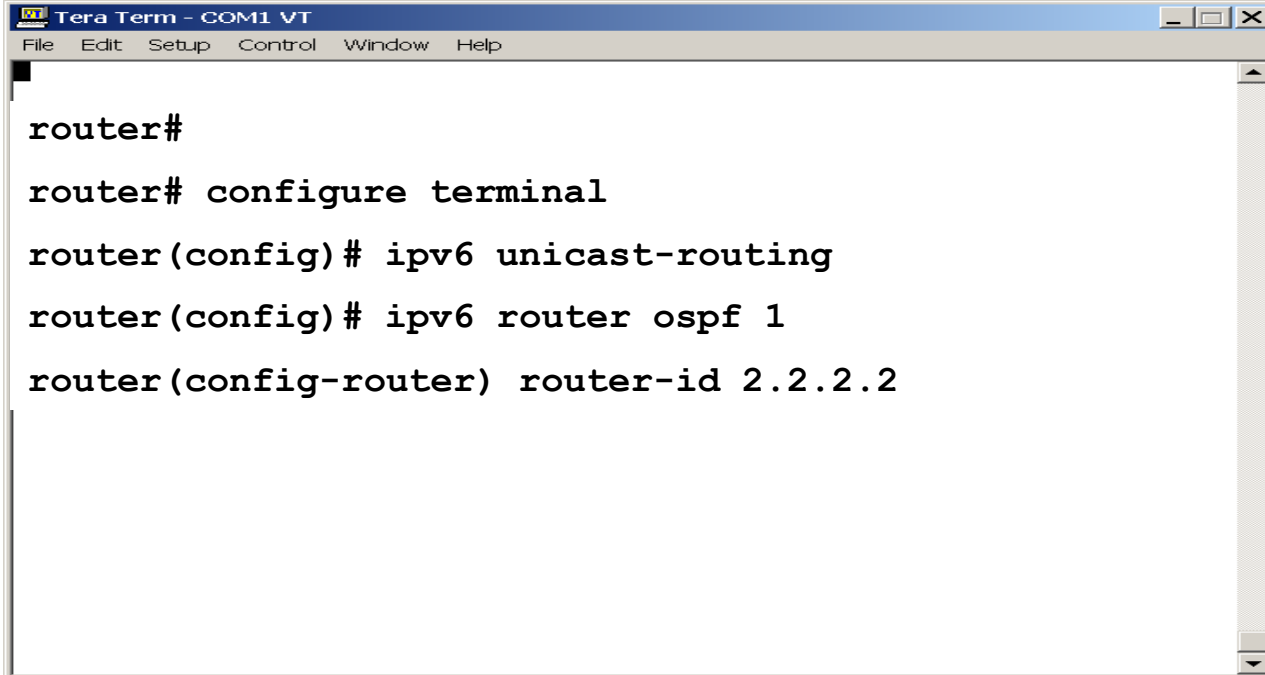
OSPFv3 Configuration



Configuring OSPFv3 in Cisco IOS Software

- Similar to OSPFv2:
 - Prefixing existing Interface and exec mode commands with “ipv6”
- Interfaces configured directly:
 - Replaces `network` command
- “Native” IPv6 router mode:
 - Not a submode of `router ospf` command

Enabling OSPFv3 Globally

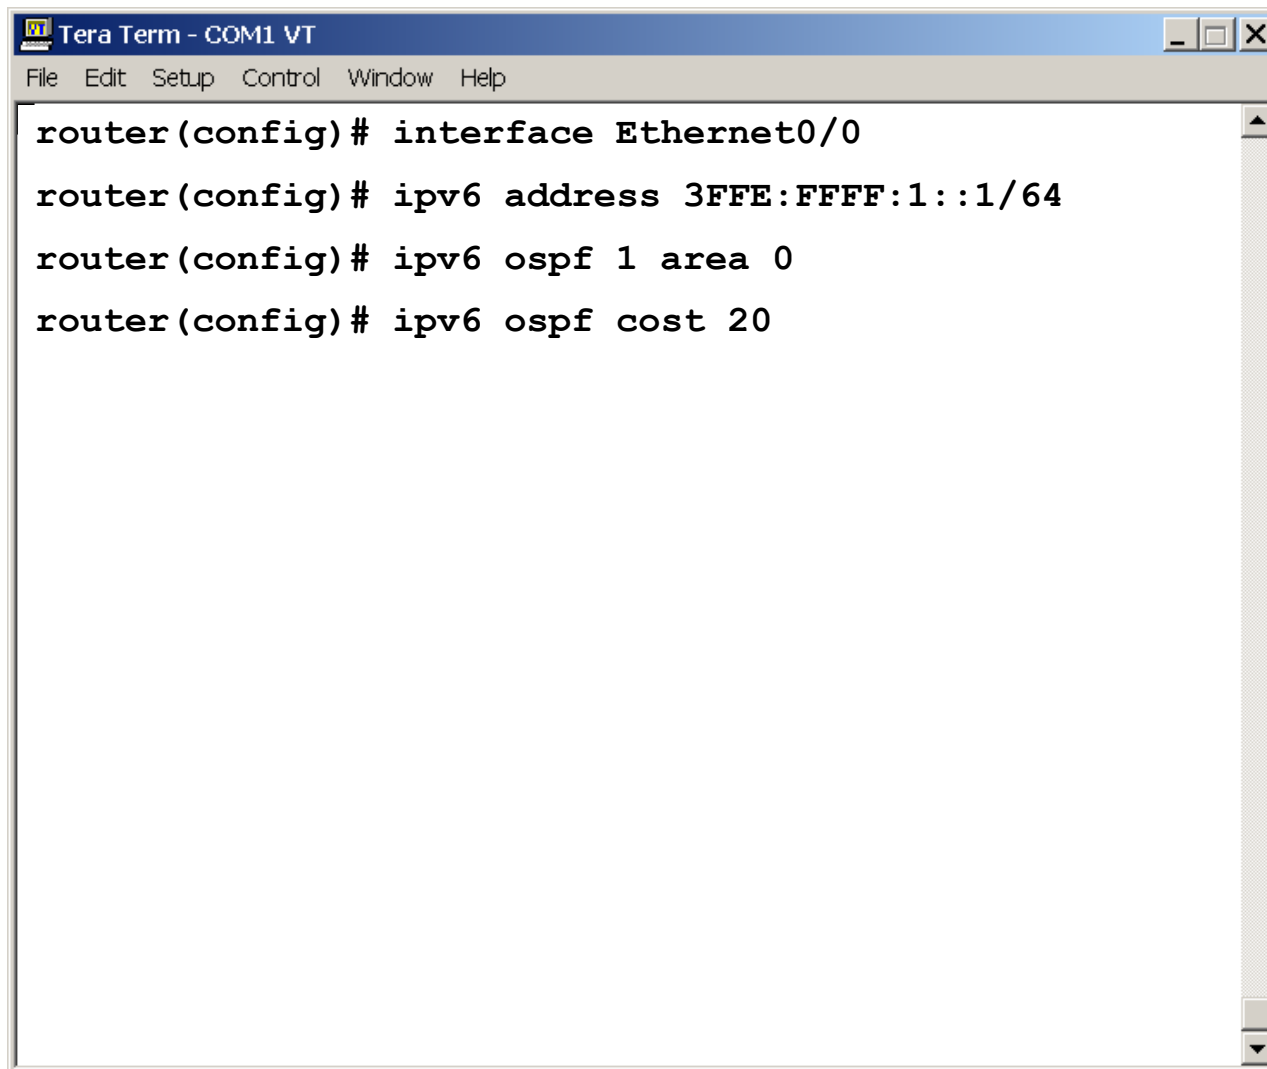


```
Tera Term - COM1 VT
File Edit Setup Control Window Help

router#
router# configure terminal
router(config)# ipv6 unicast-routing
router(config)# ipv6 router ospf 1
router(config-router) router-id 2.2.2.2
```

- Enables the forwarding of IPv6 unicast datagrams.
- Enables the OSPFv3 process number 1 on the router.
- Identifies 2.2.2.2 as the router-id for this router. It must be unique on each router


Enabling OSPFv3 on an Interface



```
Tera Term - COM1 VT
File Edit Setup Control Window Help
router(config)# interface Ethernet0/0
router(config)# ipv6 address 3FFE:FFFF:1::1/64
router(config)# ipv6 ospf 1 area 0
router(config)# ipv6 ospf cost 20
```


OSPFv3 Configuration Example

```
Router1#  
interface S1/1  
  ipv6 address  
2001:410:FFFF:1::1/64  
  ipv6 ospf 100 area 0  
  
interface S2/0  
  ipv6 address  
3FFE:B00:FFFF:1::2/64  
  ipv6 ospf 100 area 1  
  
  ipv6 router ospf 100  
    router-id 10.1.1.3  
  
Router2#  
interface S3/0  
  ipv6 address  
3FFE:B00:FFFF:1::1/64  
  ipv6 ospf 100 area 1  
  
  ipv6 router ospf 100  
    router-id 10.1.1.4
```



OSPFv3 Verification

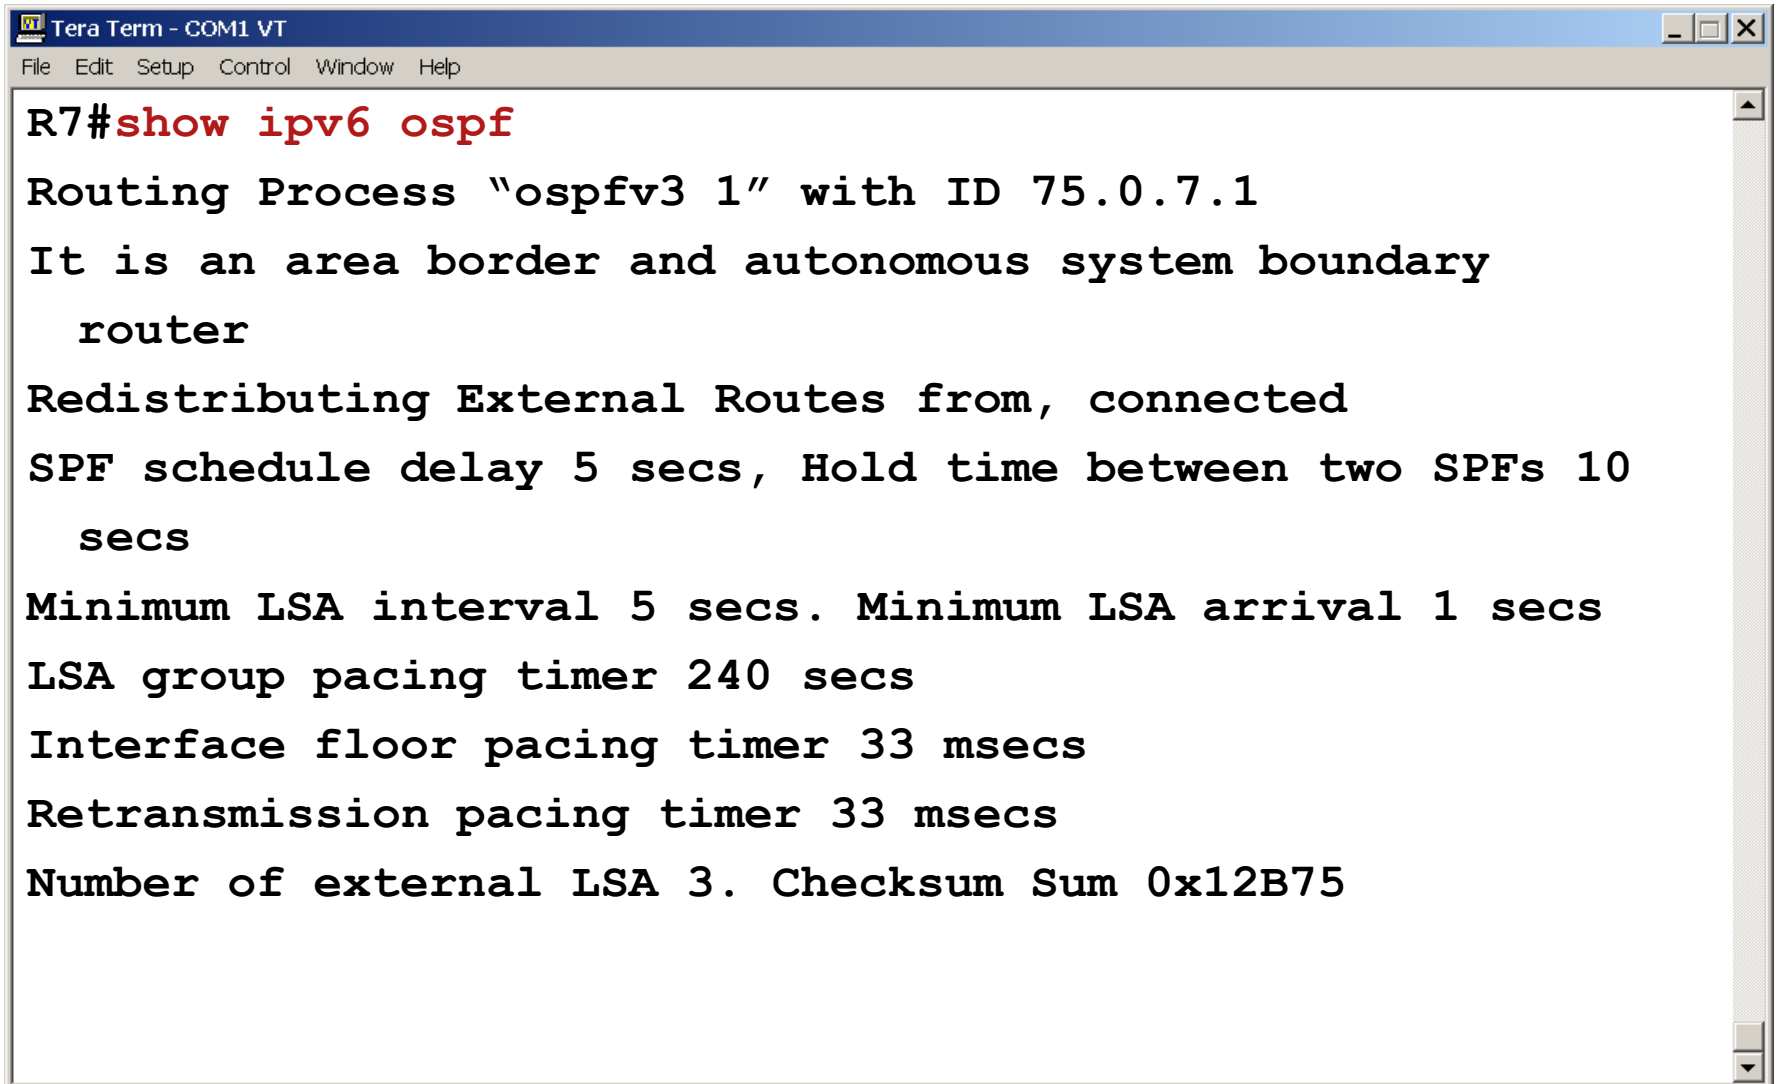


Verifying Cisco IOS OSPFv3

```
Tera Term - COM1 VT
File Edit Setup Control Window Help

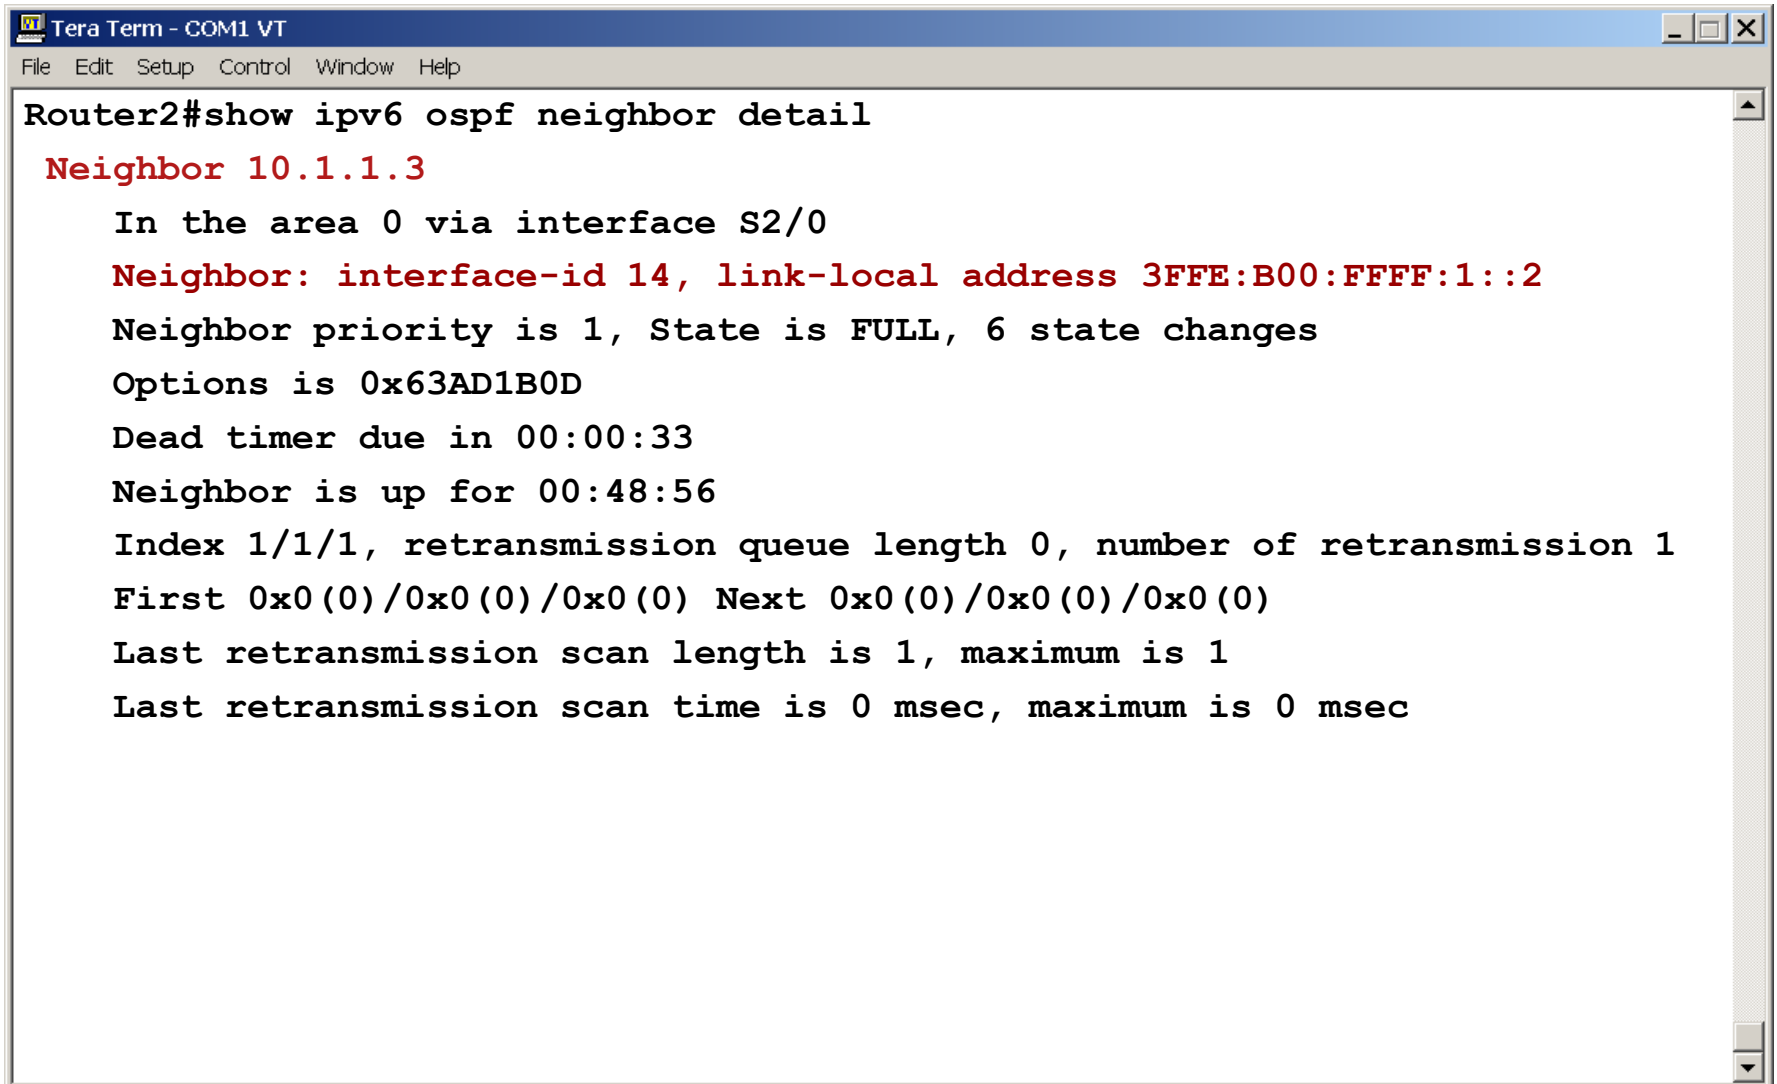
Router2#show ipv6 ospf int s 3/0
S3/0 is up, line protocol is up
  Link Local Address 3FFE:B00:FFFF:1::1, Interface ID 7
  Area 1, Process ID 100, Instance ID 0, Router ID 10.1.1.4
  Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
Retransmit 5
  Hello due in 00:00:02
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 3, maximum is 3
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.1.3
  Suppress hello for 0 neighbor(s)
```

show ipv6 ospf



```
Tera Term - COM1 VT
File Edit Setup Control Window Help
R7#show ipv6 ospf
Routing Process "ospfv3 1" with ID 75.0.7.1
It is an area border and autonomous system boundary
router
Redistributing External Routes from, connected
SPF schedule delay 5 secs, Hold time between two SPF's 10
secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
LSA group pacing timer 240 secs
Interface floor pacing timer 33 msec
Retransmission pacing timer 33 msec
Number of external LSA 3. Checksum Sum 0x12B75
```

show ipv6 ospf neighbor detail



```
Tera Term - COM1 VT
File Edit Setup Control Window Help

Router2#show ipv6 ospf neighbor detail
Neighbor 10.1.1.3
  In the area 0 via interface S2/0
  Neighbor: interface-id 14, link-local address 3FFE:B00:FFFF:1::2
  Neighbor priority is 1, State is FULL, 6 state changes
  Options is 0x63AD1B0D
  Dead timer due in 00:00:33
  Neighbor is up for 00:48:56
  Index 1/1/1, retransmission queue length 0, number of retransmission 1
  First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec
```

show ipv6 ospf database

Router Link States (Area 1)

ADV Router	Age	Seq#	Fragment	ID	Link count	Bits
26.50.0.1	1812	0x80000048	0	1	None	
26.50.0.2	1901	0x80000006	0	1	B	

Net Link States (Area 1)

ADV Router	Age	Seq#	Link	ID	Rtr count
26.50.0.1	57	0x8000003B	3		4

Inter Area Prefix Link States (Area 1)

ADV Router	Age	Seq#	Prefix
26.50.0.2	139	0x80000003	3FFE:FFFF:26::/64
26.50.0.2	719	0x80000001	3FFE:FFF:26::/64

Inter Area Router Link States (Area 1)

ADV Router	Age	Seq#	Link ID	Dest RtrID
26.50.0.2	772	0x80000001	1207959556	72.0.0.4
26.50.0.4	5	0x80000003	1258292993	75.0.7.1




Using IPv6 with IPv4



BSCI Module 8


IPv4-to-IPv6 Transition



Transition richness means:

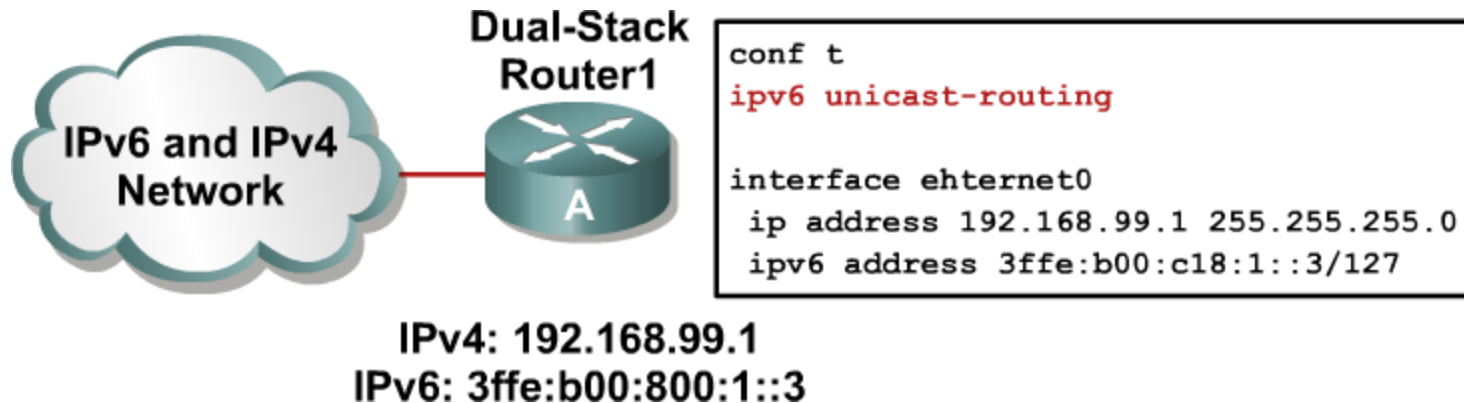
- No fixed day to convert; no need to convert all at once.
- Different transition mechanisms are available:
 - Smooth integration of IPv4 and IPv6
 - Use of dual stack or 6-to-4 tunnels
- Different compatibility mechanisms:
 - IPv4 and IPv6 nodes can communicate

Dual Stack




- Dual stack is an integration method where a node has “implementation and connectivity” to both an IPv4 and IPv6 network.

Cisco IOS Software Is IPV6-Ready: Cisco IOS Dual Stack




- If both IPv4 and IPv6 are configured on an interface, this interface is dual-stacked.

Tunneling




- Tunneling is an integration method where an IPv6 packet is encapsulated within another protocol, such as IPv4. This method of encapsulation is IPv4 protocol 41:
 - This includes a 20-byte IPv4 header with no options and an IPv6 header and payload.
 - This is considered dual stacking.

Cisco IOS Software Is IPv6-Ready: Overlay Tunnels




- Tunneling encapsulates the IPv6 packet in the IPv4 packet.

“Isolated” Dual-Stack Host




- Encapsulation can be done by edge routers between hosts or between a host and a router.

Cisco IOS Software Is IPv6-Ready: Configured Tunnel



- Configured tunnels require:
 - Dual-stack endpoints
 - IPv4 and IPv6 addresses configured at each end

Example: Cisco IOS Tunnel Configuration



Router1#	Router2#
<pre>interface Tunnel0 ipv6 address 2001:db8:1::1/64 tunnel source 192.0.2.1 tunnel destination 192.0.30.1 tunnel mode ipv6ip</pre>	<pre>interface Tunnel0 ipv6 address 2001:db8:1::2/64 tunnel source 192.0.30.1 tunnel destination 192.0.2.1 tunnel mode ipv6ip</pre>

Configuring static routes for 6-to-4 tunnels


A 6-to-4 tunnel establishes a transient link between IPv6 domains which are connected by an IPv4 backbone

- Create a tunnel interface
- Set tunnel mode with the **tunnel mode ipv6ip 6to4** command
- Create an IPv6 specific address
- Set the source interface for the tunnel
- Configure an Ipv6 static route

Converting IPv4 addresses to IPv6 for 6-4 Tunnel

- A 6-4 tunnel uses special addresses in the 2002::/16 address space
- The first 16 bits are the hexadecimal number 2002
- The next 32 bits are the original source address in hexadecimal form
- The IPv4 address of 172.16.12.1 is converted into AC10:0C01
- The complete address would be 2002:AC10:0c01:1::1/64

Configuration example for 6-to-4 tunnel between R1 and R3



```
R1
R1(config-if)# interface tunnel 0
R1(config-if)# tunnel mode 1pv6ip to 6to4
R1(config-if)# ipv6 address 2002:AC10:0C01:1::1/64
R1(config-if)# tunnel source s0/0/0
Exit
R1(config)# ipv6 route 2002::/16 tunnel 0
```

Verification


```
R1#ping 2002:AC10:1703:1::3
```

```
R3
R3(config-if)# interface tunnel 0
R3(config-if)# tunnel mode 1pv6ip to 6to4
R3(config-if)# ipv6 address 2002:AC10:1703:1::3/64
R3(config-if)# tunnel source s0/0/1
Exit
R3(config)# ipv6 route 2002::/16 tunnel 0
```

Verification

```
R3#ping 2002:AC10:1::1
```

IPv6 Tunnel Static Route Configuration



Router 1


```
Router1(config)# ipv6 unicast-routing
Router1(config)# ipv6 route FEC0::30/112 2002:AC10:1703:1::3
```

Router 2

```
Router2(config)# ipv6 unicast-routing
Router2(config)# ipv6 route FEC0::1:0/112 2002:AC10:C01:1::1
```

The next hop for both routers is the IPv6 address at the other end of the tunnel

Cisco IOS Software Is IPv6-Ready: 6-to-4 Tunneling

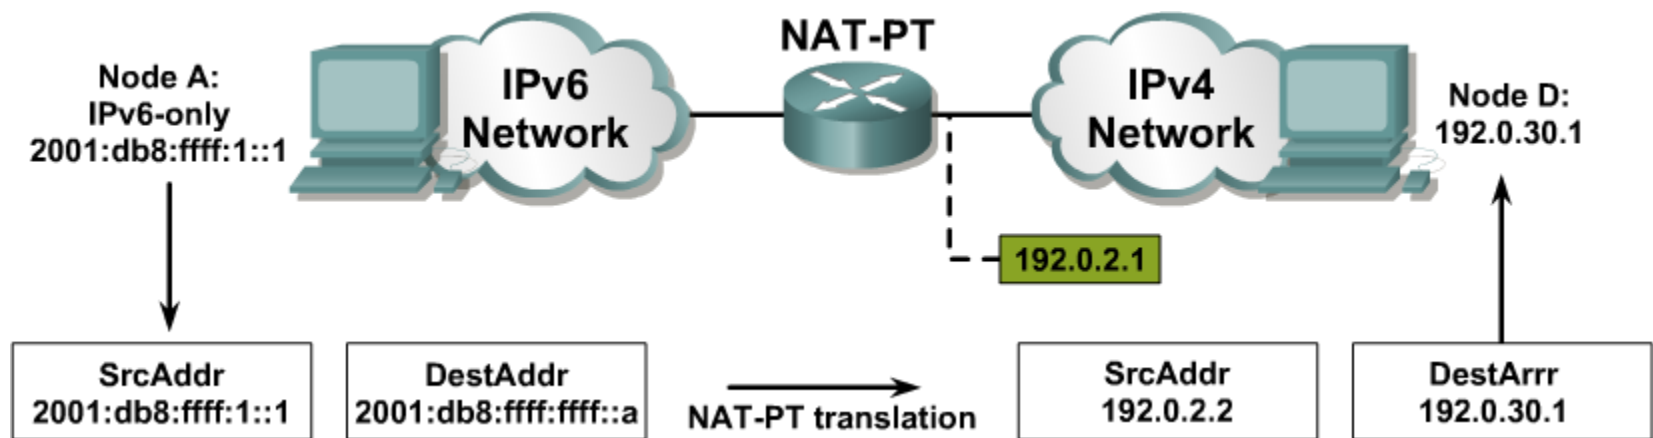


- 6-to-4:

Is an automatic tunnel method

Gives a prefix to the attached IPv6 network

Translation – NAT-PT



- NAT-Protocol Translation (NAT-PT) 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.

