IPv6 Lab and Techtorial

Hinwoto

htjun@cisco.com

Agenda (Lecture)

- IPv6 Addressing
- IPv6 Neighbor Discovery
- IPv6 Configuration on Hosts and Routers
- IPv6 Static Routing
- IPv6 OSPFv3
- IPv6 BGP
- IPv6 Deployment Techniques

Agenda (Labs)

- Lab1: IPv6 Addressing
- Lab2: IPv6 Neighbor Discovery
- Lab3: IPv6 Static Routing
- Lab4: IPv6 OSPFv3
- Lab5: IPv6 BGP
- Lab6: IPv6 Tunneling (Manual) *Bonus*

Prerequisites for this Session

- A network engineer
- Or wannabe network engineer
- Love Cisco
- Not Customer Service Officer, Not Account Manager,
  Not Spy, Not etc…..
- Have Own Laptop ☺

IPv6 Technology

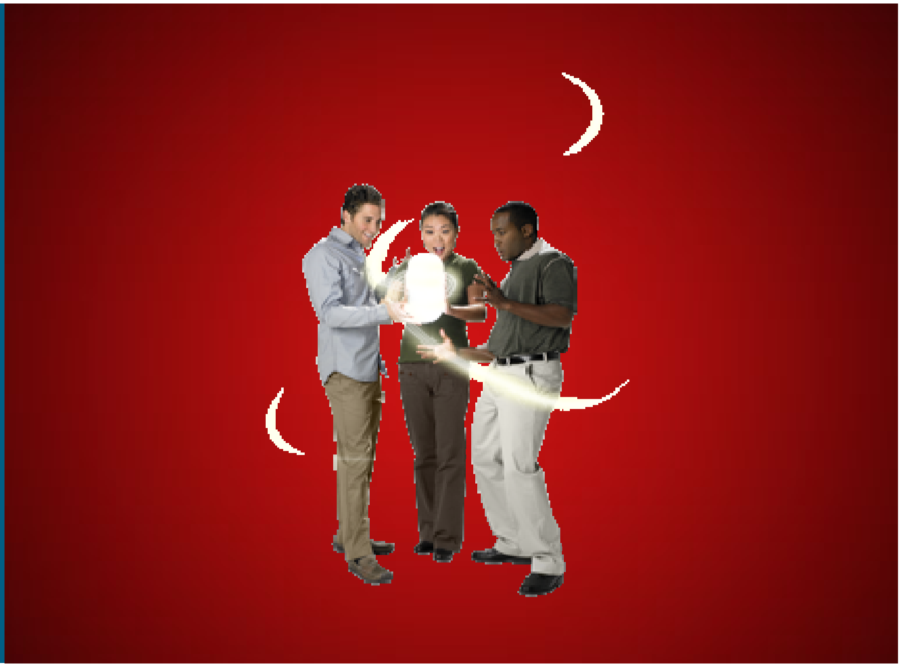

IPv4 and IPv6 Header Comparison

IPv4 Header

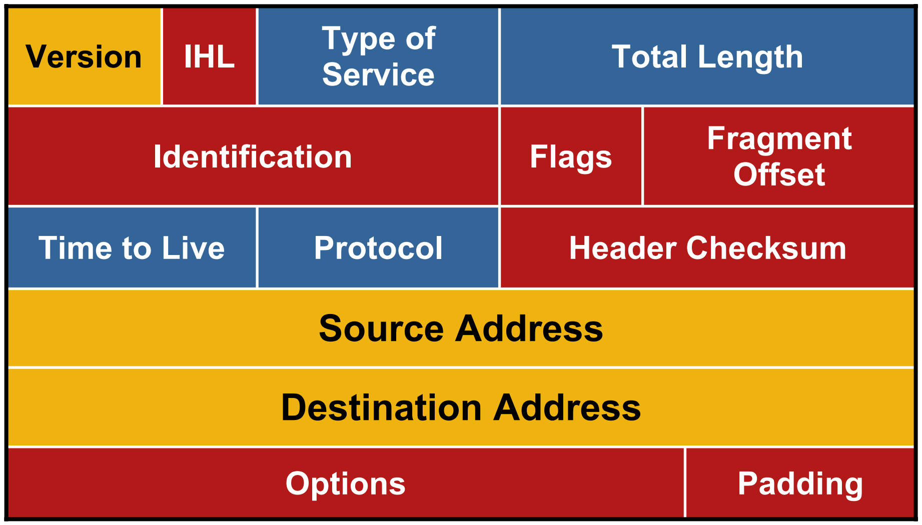

IPv6 Header

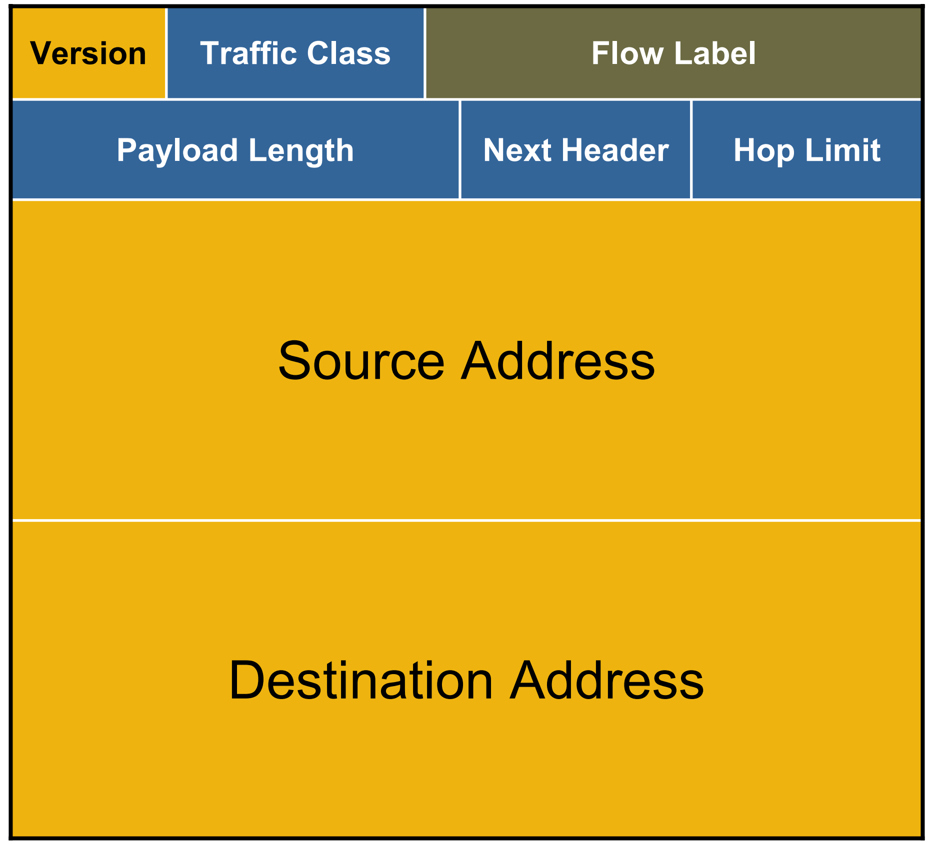

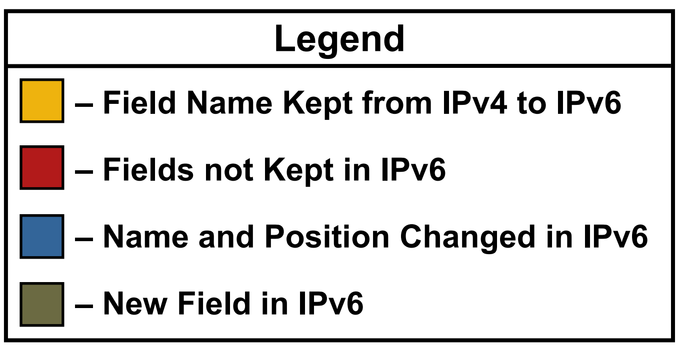

IPv6 Addressing

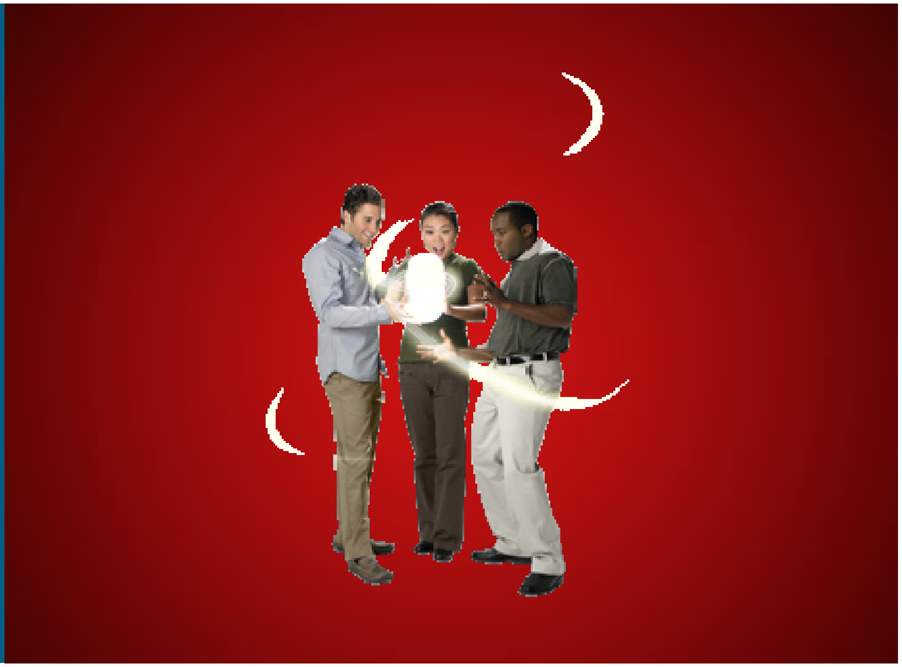

IPv6 Addressing

IPv4 32-bits

IPv6 128-bits

2^{32} = 4,294,967,296

2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456

$2^{128} = 2^{32} * 2^{96}$

2^{96} = 79,228,162,514,264,337,593,543,950,336 times the number of possible IPv4 Addresses (79 trillion trillion)

IPv6 Addressing

$$\frac{2^{128}}{6.5 \text{ Billion}}$$ = 52 Trillion Trillion IPv6 addresses per person

World's population is approximately 6.5 billion

$$\frac{\text{52 Trillion Trillion}}{\text{100 Billion}}$$ = 523 Quadrillion (523 thousand trillion) IPv6 addresses for every human brain cell on the planet!

Typical brain has ~100 billion brain cells (your count may vary)

Addressing Format

Representation

- 16-bit hexadecimal numbers
- Numbers are separated by (:)
- Hex numbers are not case sensitive
- Abbreviations are possible

    Leading zeros in contiguous block could be represented by (::)

    Example:

        2001:0db8:0000:130F:0000:0000:087C:140B

        2001:0db8:0:130F::87C:140B

    Double colon only appears once in the address

Addressing

Prefix Representation

- Representation of prefix is just like CIDR
- In this representation you attach the prefix length
- Like v4 address:

  198.10.0.0/16

- V6 address is represented the same way:

  2001:db8:12::/48

- Only leading zeros are omitted. Trailing zeros are not omitted.

  2001:0db8:0012::/48 = 2001:db8:12::/48

  2001:db8:**1200**::/48 ≠ 2001:db8:12::/48

IPv6 Address Representation

- Loopback address representation

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

  Same as 127.0.0.1 in IPv4

  Identifies self

- Unspecified address representation

  0:0:0:0:0:0:0:0=> ::

  Used as a placeholder when no address available

  (Initial DHCP request, Duplicate Address Detection, DAD)

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

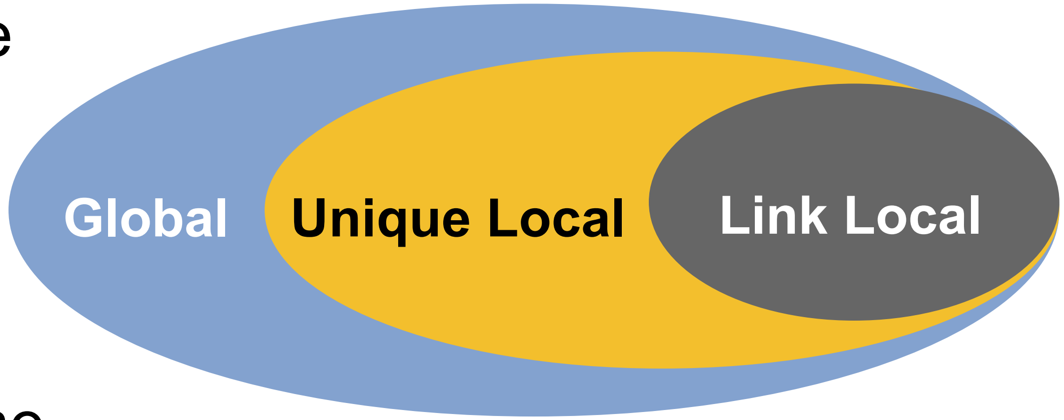

Addressing

Some Special Addresses

Type	Binary	Hex
Aggregatable Global Unicast Address	001	2 or 3
Link-Local Unicast Address	1111 1110 10	FE80::/10
Unique Local Unicast Address	1111 1100 1111 1101	FC00::/7 FC00::/8(Registry) FD00::/8 (No Registry)
Multicast Address	1111 1111	FF00::/8

Types of IPv6 Addresses

- Unicast

  Address of a single interface. One-to-one delivery to single interface.

- Multicast

  Address of a set of interfaces. One-to-many delivery to all interfaces in the set.

- Anycast

  Address of a set of interfaces. One-to-one-of-many delivery to a single interface in the set that is closest.

- No more broadcast addresses

Aggregatable Global Unicast Addresses

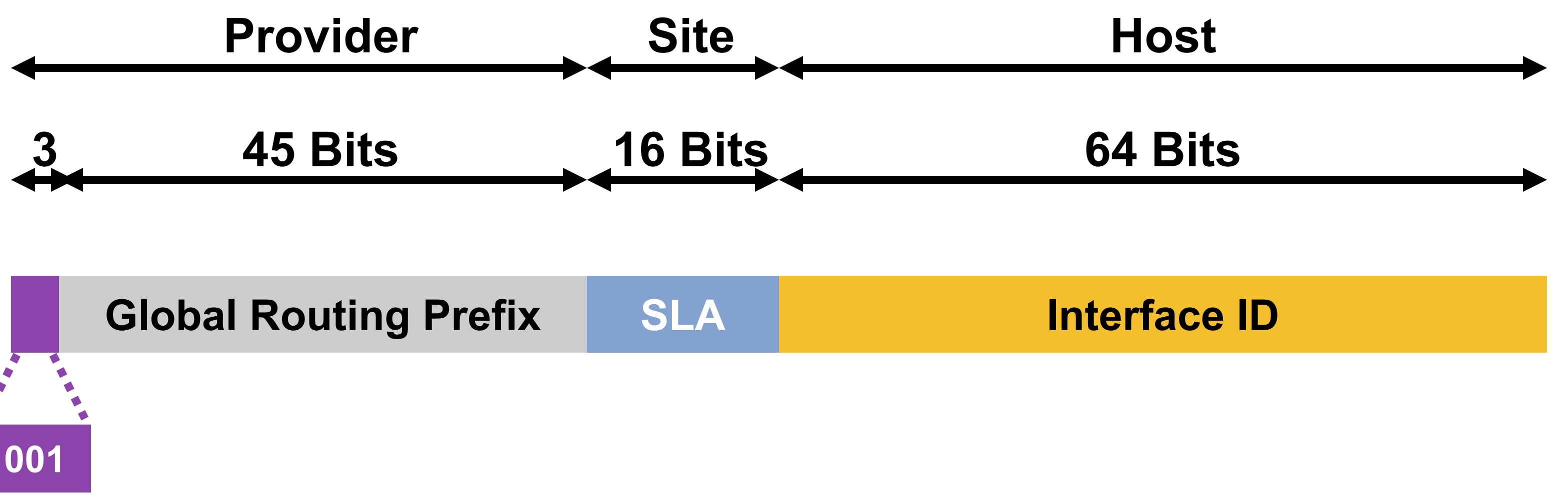

Aggregatable Global Unicast Addresses Are:

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

Unique-Local

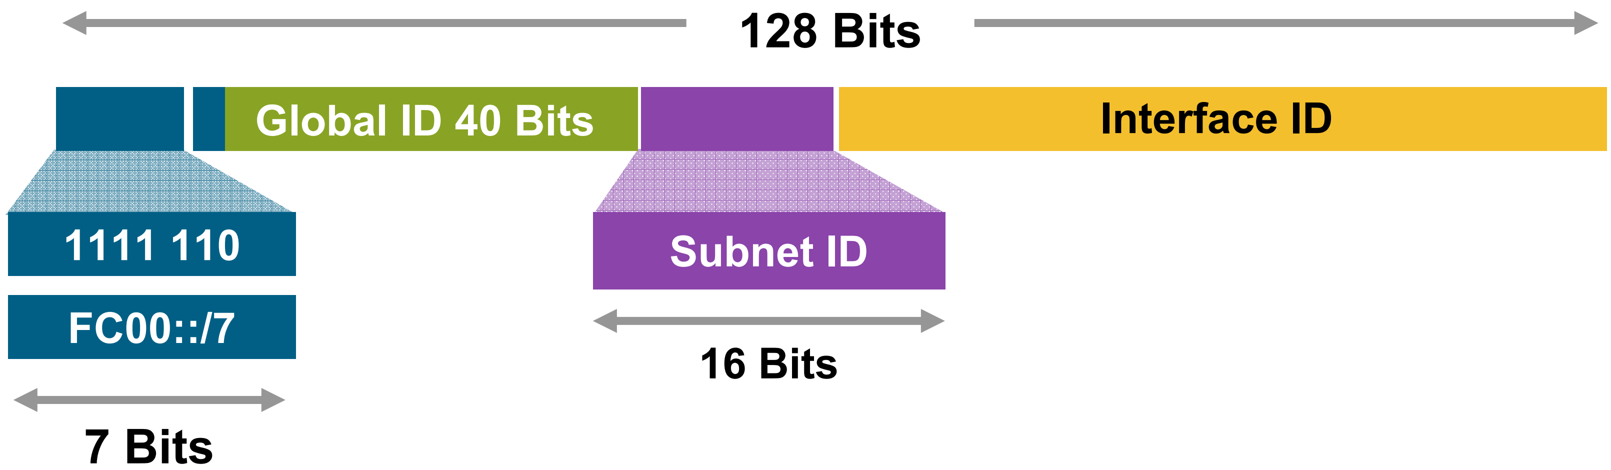

Unique-Local Addresses Used for:

- Local communications
- Inter-site VPNs
- Not routable on the Internet

Link-Local

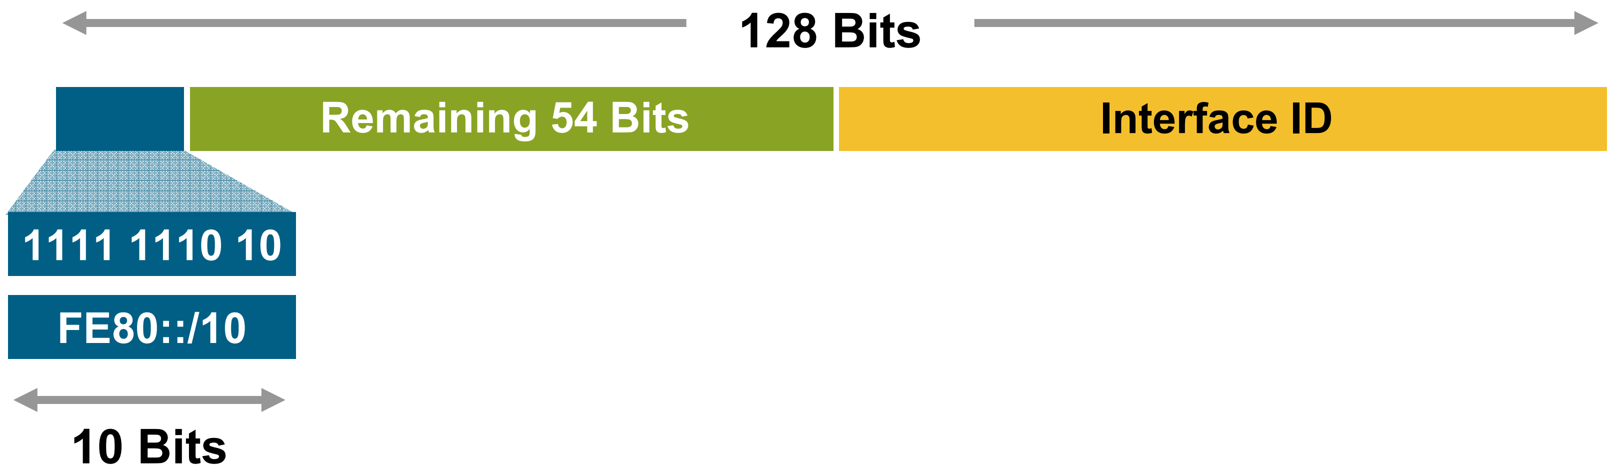

Link-Local Addresses Used for:

- Mandatory address for communication between two IPv6 device (like 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

IPv6 Address Allocation Process

Partition of Allocated IPv6 Address Space

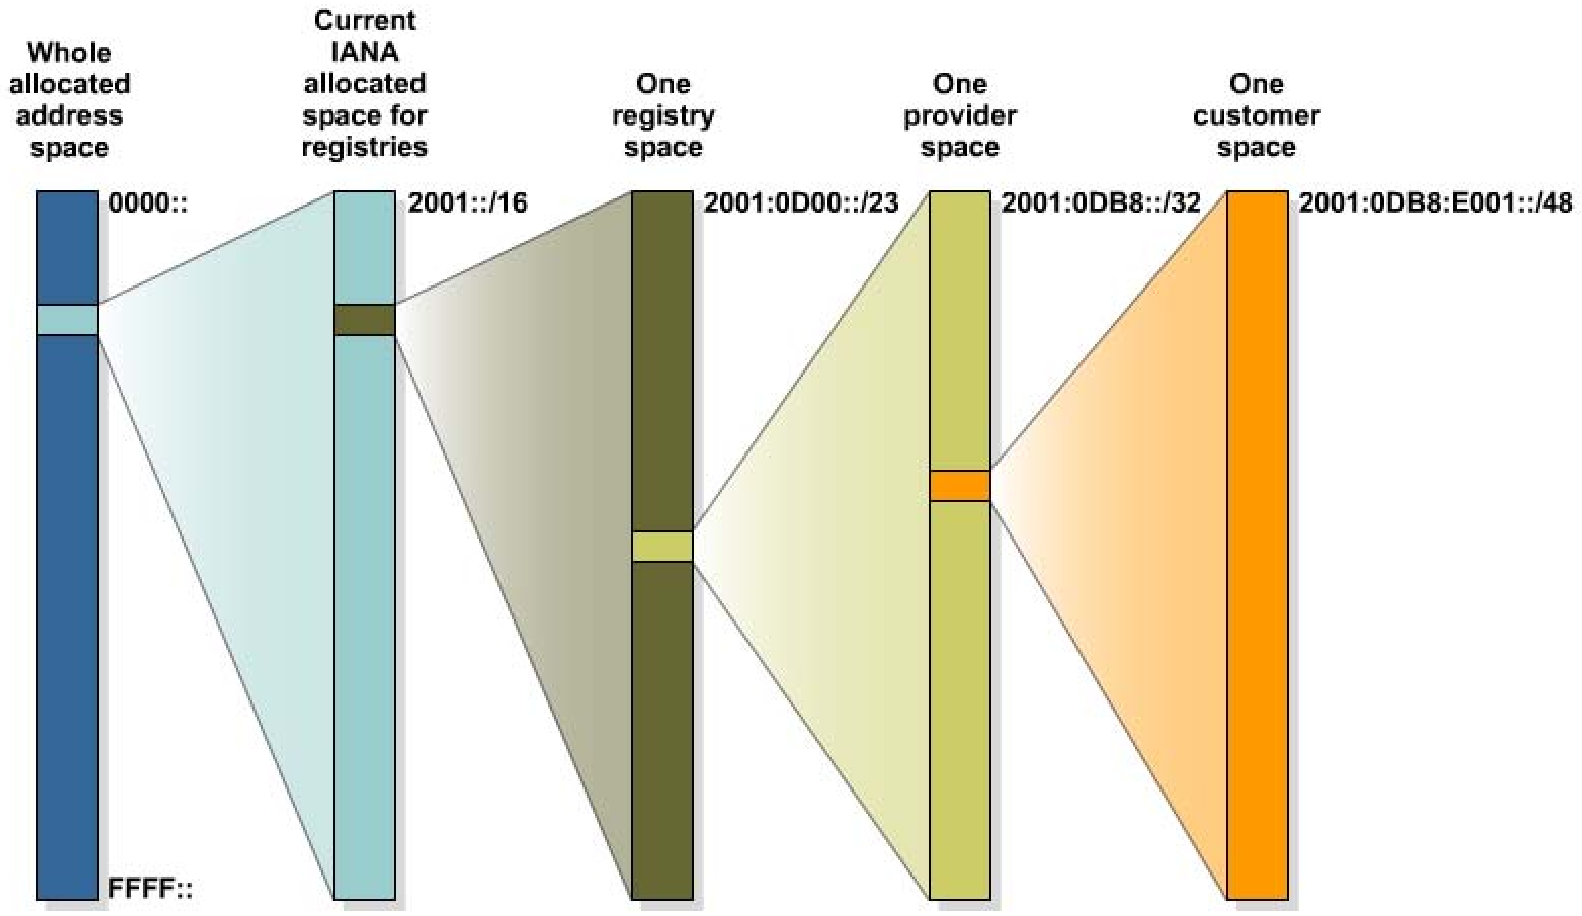

IPv6 Address Allocation Process

Partition of Allocated IPv6 Address Space (Cont.)

- Lowest-order 64-bit field of unicast address may be assigned in several different ways:
  - –Auto-configured from a 64-bit EUI-64, or expanded from a 48-bit MAC address (e.g., Ethernet address)
  - –Autogenerated pseudo-random number (to address privacy concerns)
  - –Assigned via DHCP
  - –Manually configured

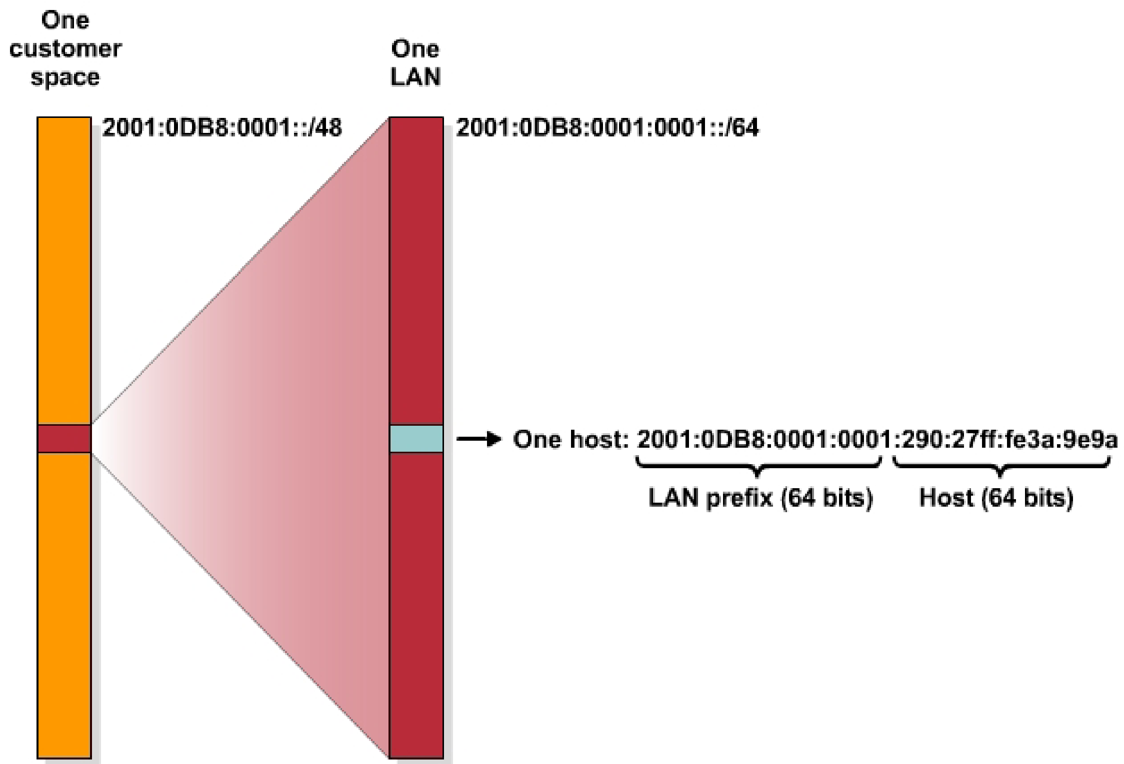

IPv6 Interface Identifier

- Cisco uses the EUI-64 format to do stateless auto-configuration
- 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” bit) is set to 1 for global scope and 0 for local scope

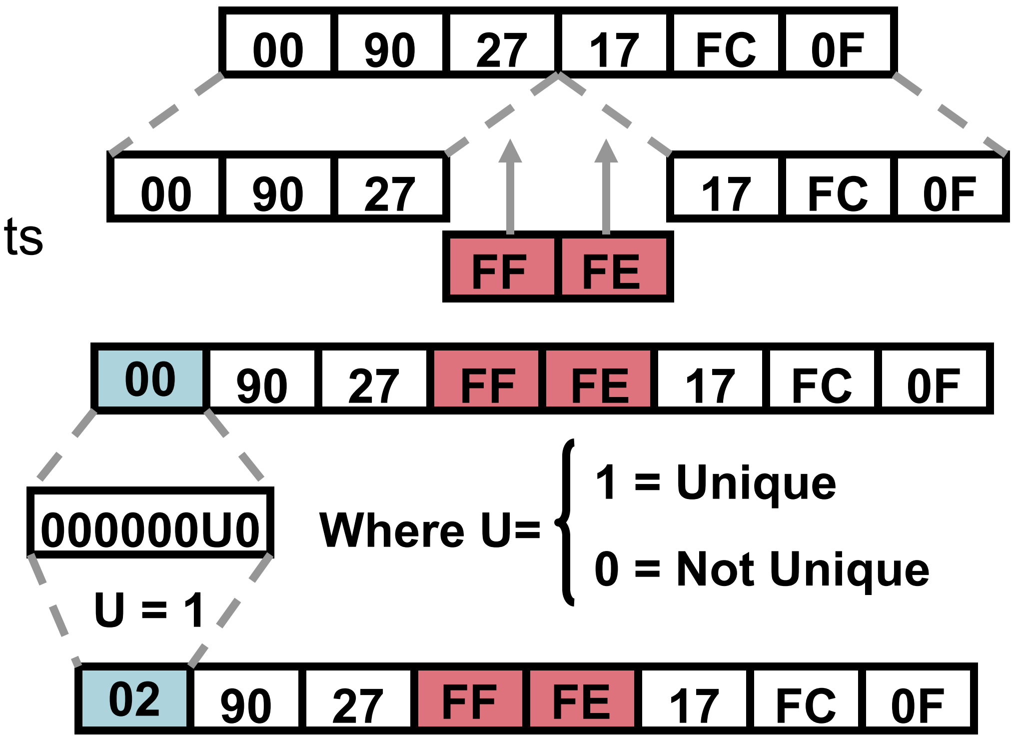

ICMPv6 and Neighbor Discovery

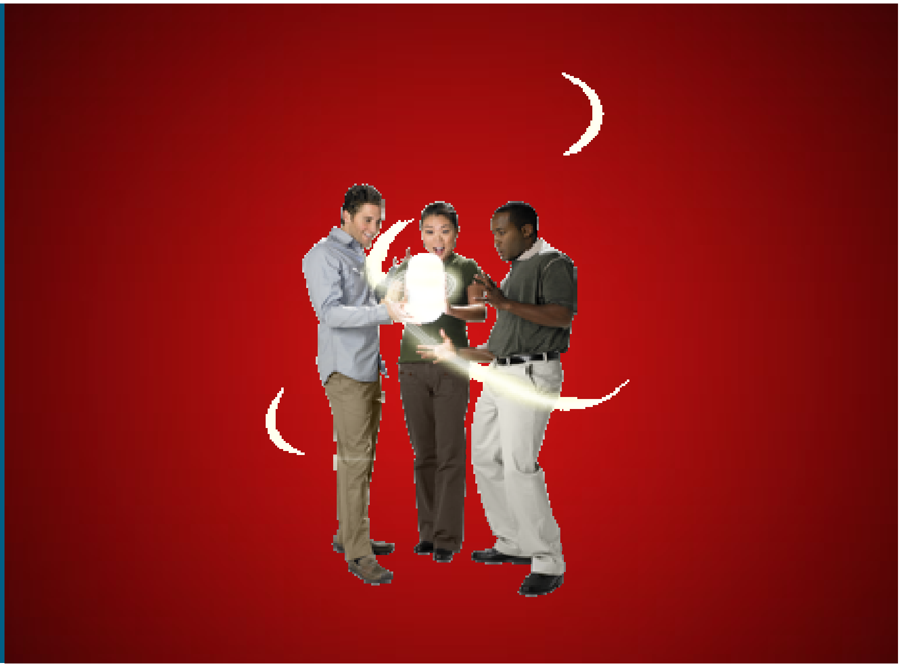

ICMPv6

- Internet Control Message Protocol version 6
- RFC 2463
- Modification of ICMP from IPv4
- Message types are similar
(but different types/codes)
    - Destination unreachable (type 1)
    - Packet too big (type 2)
    - Time exceeded (type 3)
    - Parameter problem (type 4)
    - Echo request/reply (type 128 and 129)

Neighbor Discovery

- Replaces ARP, ICMP (redirects, router discovery)
- Reachability of neighbors
- Hosts use it to discover routers, auto configuration of addresses
- Duplicate Address Detection (DAD)

Neighbor Discovery

- Neighbor discovery uses ICMPv6 messages, originated from node on link local with hop limit of 255
- Consists of IPv6 header, ICMPv6 header, neighbor discovery header, and neighbor discovery options
- Five neighbor discovery messages
  1. **Router solicitation (ICMPv6 type 133)**
  2. **Router advertisement (ICMPv6 type 134)**
  3. **Neighbor solicitation (ICMPv6 type 135)**
  4. **Neighbor advertisement (ICMPv6 type 136)**
  5. Redirect (ICMPV6 type 137)

Router Solicitation and Advertisement

1—ICMP Type = 133 (RS)

Src = link-local address (FE80::1/10)

Dst = all-routers multicast address (FF02::2)

Query = please send RA

2—ICMP Type = 134 (RA)

Src = link-local address (FE80::2/10)

Dst = all-nodes multicast address (FF02::1)

Data = options, subnet prefix, lifetime, autoconfig flag

- Router Solicitations (RS) are sent by booting nodes to request RAs for configuring the interfaces
- Routers send periodic Router Advertisements (RA) to the all-nodes multicast address

Neighbor Solicitation and Advertisement

Neighbor Solicitation
ICMP type = 135

Src = A
Dst = Solicited-node multicast of B
Data = link-layer address of A
Query = what is your link address?

Neighbor Advertisement
ICMP type = 136
Src = B
Dst = A
Data = link-layer address of B

A and B Can Now Exchange Packets on this Link

Autoconfiguration

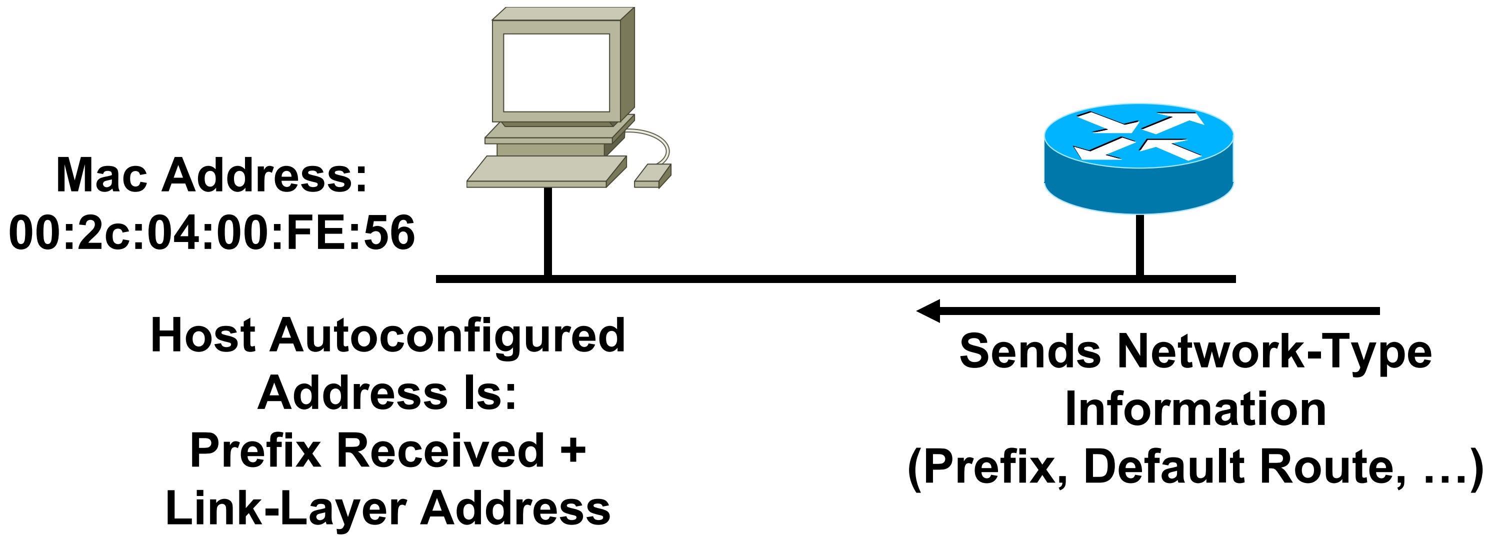

Larger Address Space Enables:

- The use of link-layer addresses inside the address space
- Autoconfiguration with “no collisions”
- Offers “plug-and-play”

Renumbering

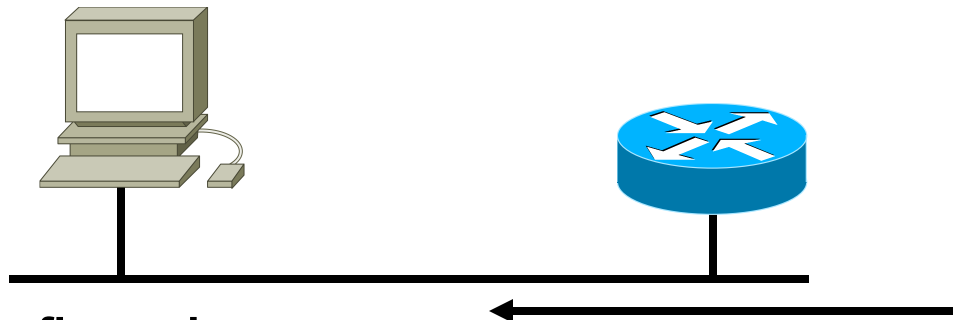

Mac Address:
00:2c:04:00:FE:56

Host Autoconfigured Address Is:
New Prefix Received + Link-Layer Address

Sends New Network-Type Information (Prefix, Default Route, …)

Data = Two Prefixes:
Current Prefix (to Be Deprecated), with Short Lifetimes
New Prefix (to Be Used), with Normal Lifetimes

Larger Address Space Enables:

- Renumbering, using autoconfiguration and multiple addresses

Renumbering (Cont.)

Router Configuration after Renumbering:

```
interface Ethernet0
  ipv6 nd prefix 2001:db8:c18:1::/64 43200 0
  ipv6 nd prefix 2001:db8:c18:2::/64 43200 43200
```

or:

```
interface Ethernet0
  ipv6 nd prefix 2001:db8:c18:1::/64 at Jul 31 2008 23:59 Jul 20 2008 23:59
  ipv6 nd prefix 2001:db8:c18:2::/64 43200 43200
```

New Network Prefix: 2001:db8:c18:2::/64
Deprecated Prefix: 2001:db8:c18:1::/64

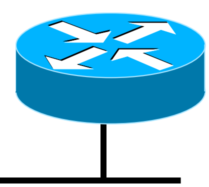

Router Advertisements

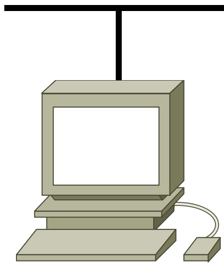

Autoconfiguring IPv6 Hosts

Host Configuration:

```
deprecated address 2001:db8:c18:1:260:8ff:fede:8fbe
preferred address 2001:db8:c18:2:260:8ff:fede:8fbe
```


Configuring IPv6

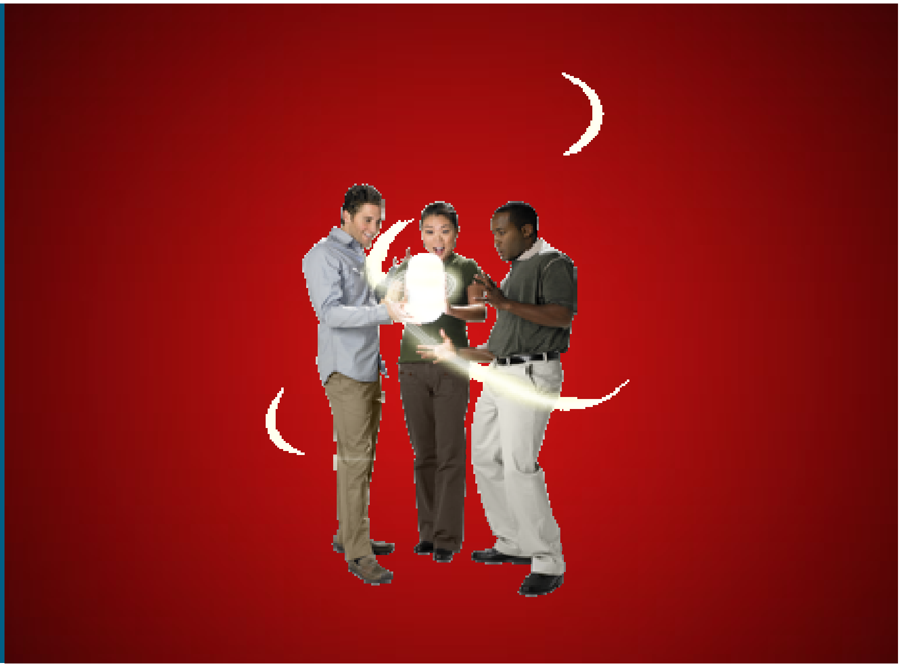

IOS IPv6 Addressing Examples (1)

Manual Interface Identifier

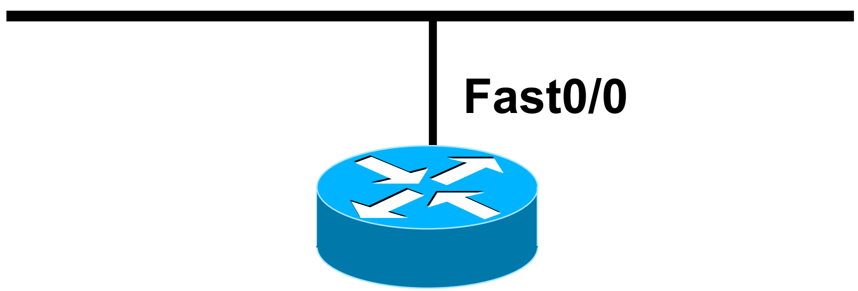


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 ip pim sparse-mode
 duplex auto
 speed auto
 ipv6 address 2006:1::1/64
 ipv6 enable
 ipv6 nd ra-interval 30
 ipv6 nd prefix 2006:1::/64 300 300
!
```

IOS IPv6 Addressing Examples (1)

Manual Interface Identifier

```
r1#sh ipv6 int fast0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::207:50FF:FE5E:9460
  Global unicast address(es):
    2006:1::1, subnet is 2006:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
    FF02::1:FF5E:9460
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND
  ND
  ND
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 30 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
r1#
```

MAC Address : 0007.505e.9460

```
r1#sh int fast0/0
FastEthernet0/0 is up, line protocol is up
  Hardware is AmdFE, address is 0007.505e.9460 (bia 0007.505e.9460)
```


IOS IPv6 Addressing Examples (2)

EUI-64 Interface Identifier

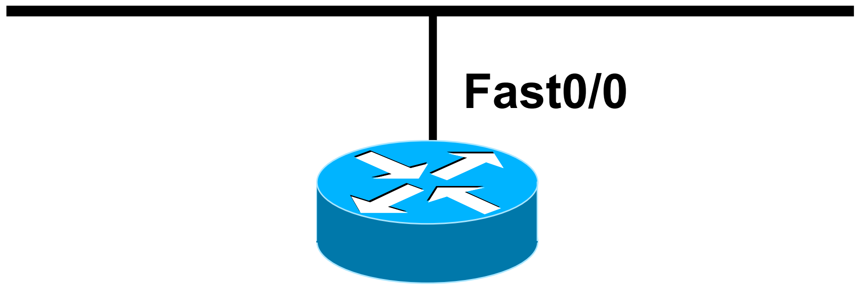


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 ip pim sparse-mode
 duplex auto
 speed auto
 ipv6 address 2006:1::/64 eui-64
 ipv6 enable
 ipv6 nd ra-interval 30
 ipv6 nd prefix 2006:1::/64 300 300
!
```


IOS IPv6 Addressing Examples (2)

EUI-64 Interface Identifier

```
r1#sh ipv6 int fast0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::207:50FF:FE5E:9460
  Global unicast address(es):
    2006:1::207:50FF:FE5E:9460, subnet is 2006:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF5E:9460
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICM
  ND
  ND
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 30 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
r1#
```

MAC Address : 0007.505e.9460

```
r1#sh int fast0/0
FastEthernet0/0 is up, line protocol is up
  Hardware is AmdFE, address is 0007.505e.9460 (bia 0007.505e.9460)
```

Configuring IPv6 on Hosts

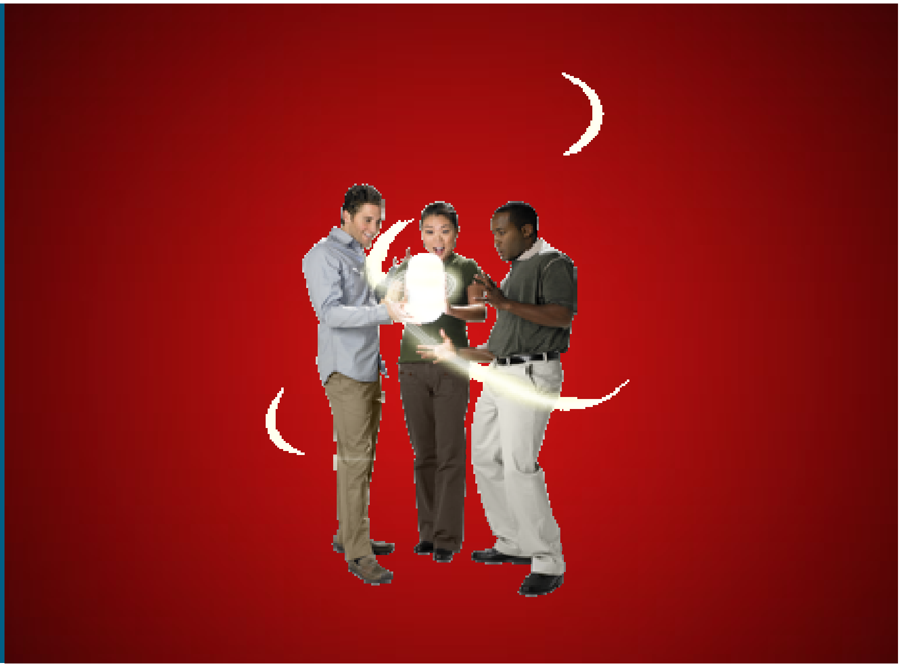

Microsoft (Windows)

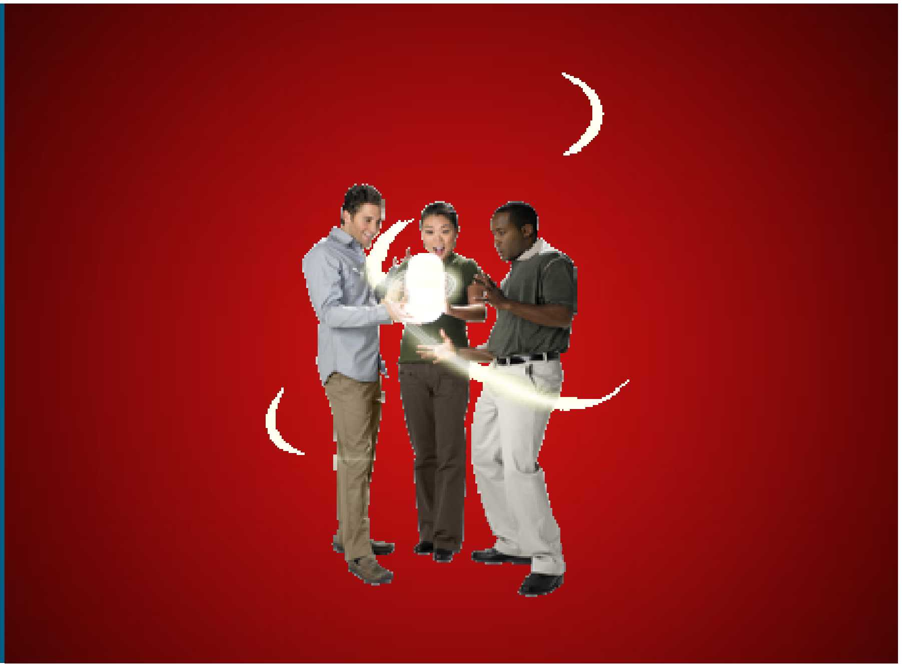

Client Configuration (Windows XP): Dual-Stack

- Required
  - Microsoft Windows XP (SP1 or higher)
  - Microsoft Windows Server 2003
- IPv6 must be installed
  - C:\>ipv6 install
- Have network (routers/switches) configured for IPv6
  - Stateless autoconfiguration and/or DHCPv6


```
C:\>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection 1:
        Connection-specific DNS Suffix  . :
        IP Address. . . . . . . . . . . . : 10.1.1.100
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        IP Address. . . . . . . . . . . . : 2001:DB8:C003:1122:203:ffff:fe81:d6da
        IP Address. . . . . . . . . . . . : fe80::203:ffff:fe81:d6da%4
        Default Gateway . . . . . . . . . : 10.1.1.1
                                            fe80::201:42ff:fe2d:9580
```

IPv6 Access Lists

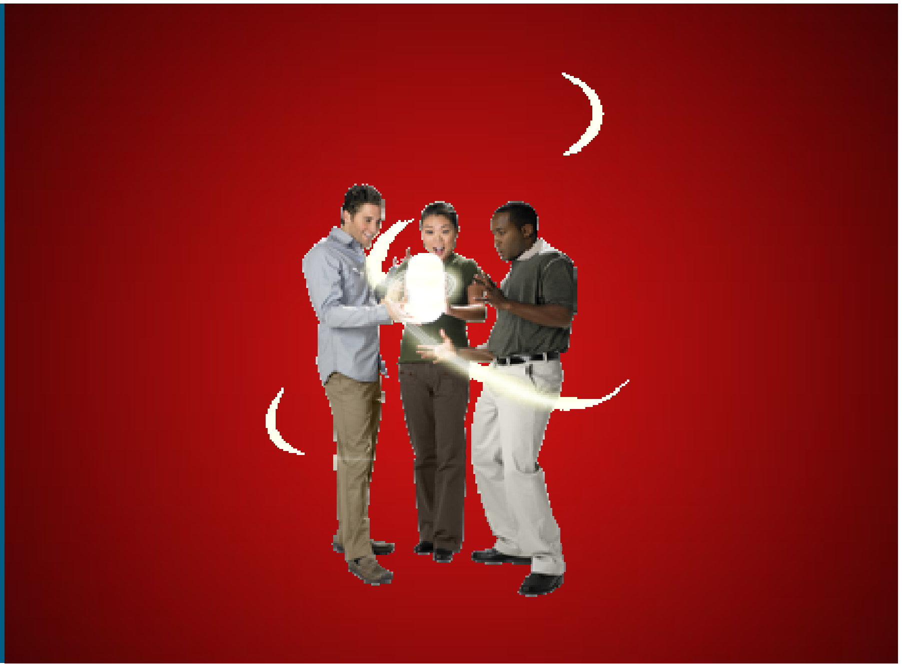

IPv6 Access-List Example

- Filtering outgoing traffic from site-local source addresses

```
ipv6 access-list blocksite deny fc00:0:0:2::/64 any
ipv6 access-list blocksite permit any any

interface Ethernet0
 ipv6 traffic-filter blocksite out
```

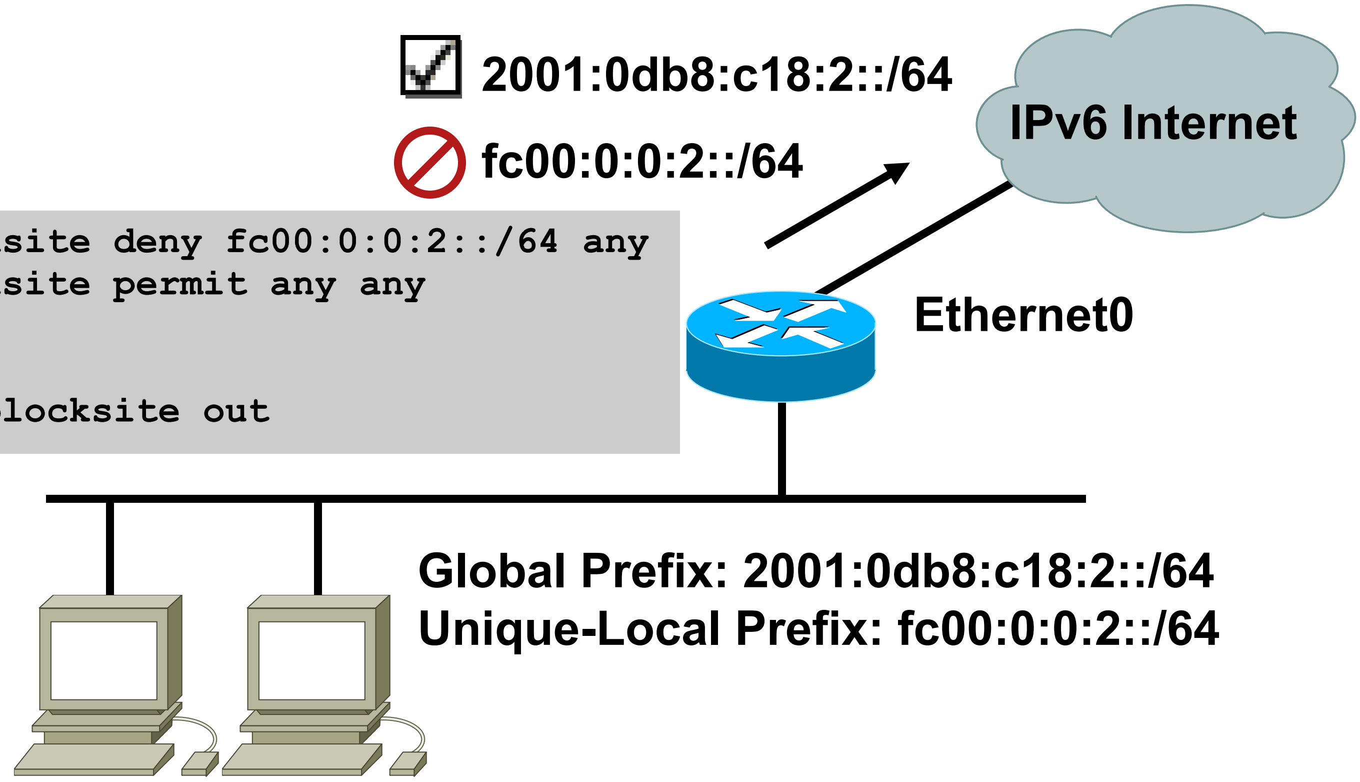

Routing in IPv6

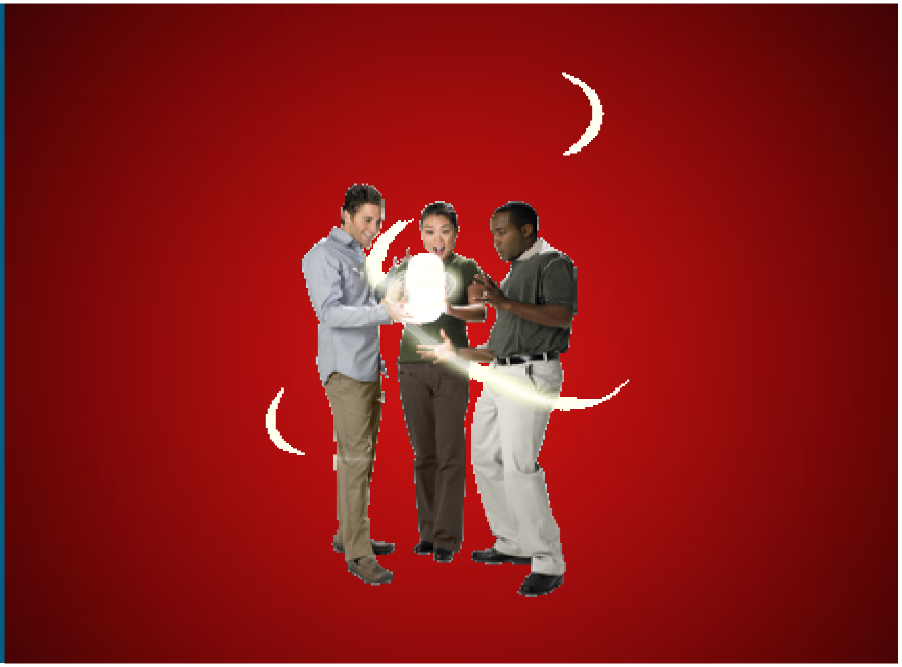

Routing in IPv6

- As in IPv4, IPv6 has two families of routing protocols: IGP and EGP, and still uses the longest-prefix match routing algorithm
- IGP
  - RIPng (RFC 2080)
  - Cisco EIGRP for IPv6
  - Integrated IS-ISv6 (draft-ietf-isis-ipv6-07)
  - OSPFv3 (RFC 2740)—(draft-ietf-ospf-ospfv3-update-19)
- EGP: MP-BGP4 (RFC 2858 and RFC 2545)
- Cisco IOS supports all of them
  - Pick one that meets your objectives

IPv6 Default and Static Routing

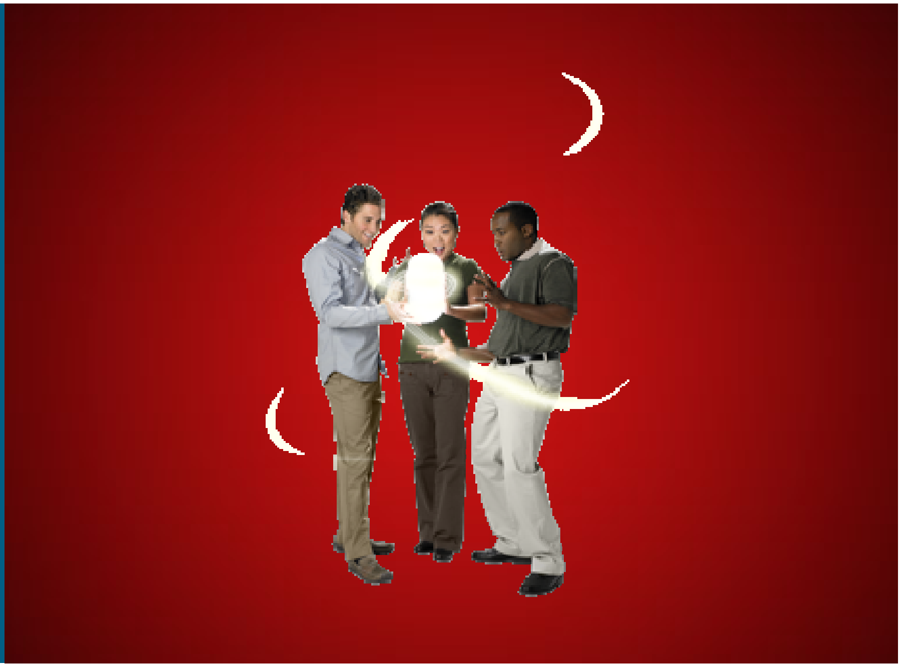

Default and Static Routing

- Similar to IPv4. Need to define the next hop/interface.
- Default route denoted as ::/0

*ipv6 route ipv6-prefix/**prefix-length {ipv6-address** | **interface-type interface-number [ipv6-address]} [administrative-distance] [administrative-multicast-distance** | unicast | multicast**] [**tag **tag]***

Examples:

- Forward packets for network 2001:DB8::0/32 through 2001:DB8:1:1::1 with an administrative distance of 10

```
Router(config)# ipv6 route 2001:DB8::0/32 2001:DB8:1:1::1 10
```

- Default route to 2001:DB8:1:1::1

```
Router(config)# ipv6 route ::/0 2001:DB8:1:1::1
```


OSPFv3 (RFC 2740)

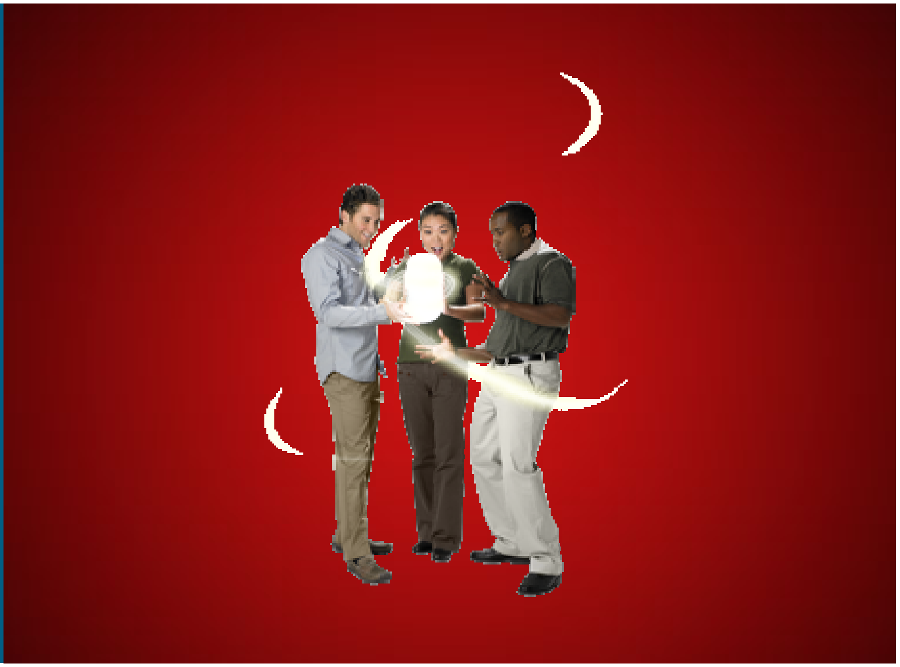

OSPFv3 Configuration Example

```
Router1#
interface POS1/1
 ipv6 address 2001:410:FFFF:1::1/64
 ipv6 enable
 ipv6 ospf 100 area 0

interface POS2/0
 ipv6 address 2001:B00:FFFF:1::2/64
 ipv6 enable
 ipv6 ospf 100 area 1

 ipv6 router ospf 100
   router-id 10.1.1.3

Router2#
interface POS3/0
 ipv6 address 2001:B00:FFFF:1::1/64
 ipv6 enable
 ipv6 ospf 100 area 1

ipv6 router ospf 100
   router-id 10.1.1.4
```

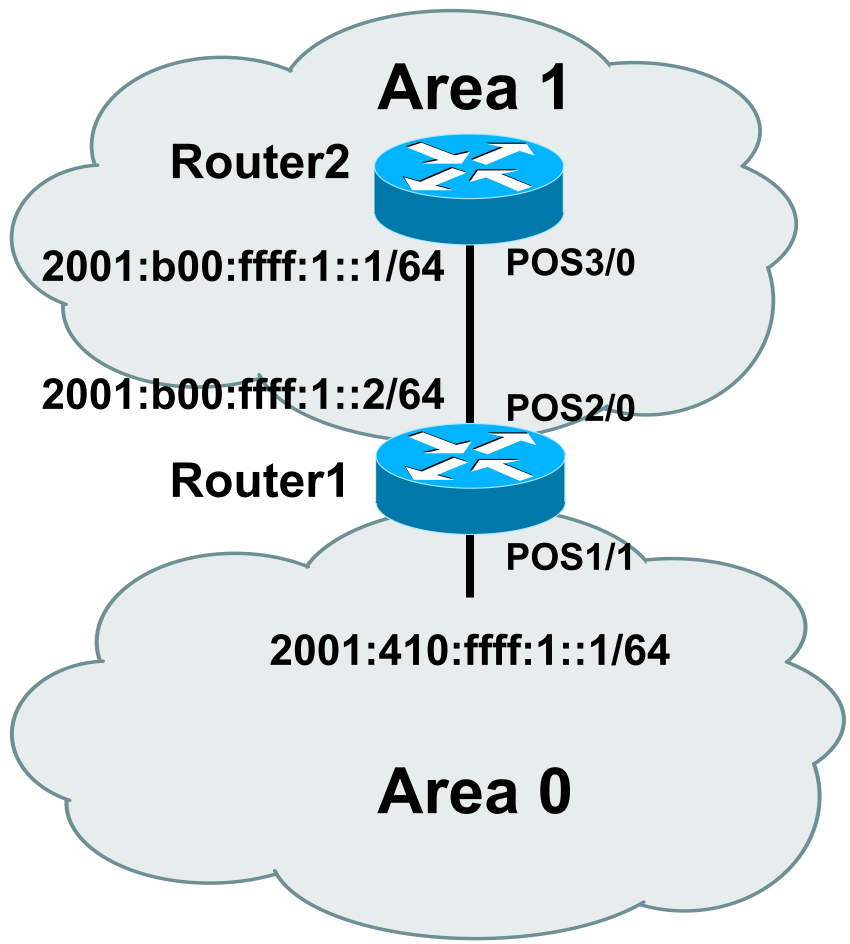

OSPFv3 Show Commands

```
Router2#sh ipv6 ospf int pos 3/0
POS3/0 is up, line protocol is up
  Link Local Address FE80::290:86FF:FE5D:A000, 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)
```

OSPFv3 Show Commands

```
Router2#sh ipv6 ospf  neighbor detail
 Neighbor 10.1.1.3
    In the area 1 via interface POS3/0
    Neighbor: interface-id 8, link-local address FE80::2D0:FFFF:FE60:DFFF
    Neighbor priority is 1, State is FULL, 12 state changes
    Options is 0x630C34B9
    Dead timer due in 00:00:33
    Neighbor is up for 00:49:32
    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 2, maximum is 2
    Last retransmission scan time is 0 msec, maximum is 0 msec
```


OSPFv3 Show Commands

```
Router2#sh ipv6 route
IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
OI   2001:410:FFFF:1::/64 [110/2]
      via FE80::2D0:FFFF:FE60:DFFF, POS3/0
C    2001:B00:FFFF:1::/64 [0/0]
      via ::, POS3/0
L    2001:B00:FFFF:1::1/128 [0/0]
      via ::, POS3/0
L    FE80::/10 [0/0]
      via ::, Null0
L    FF00::/8 [0/0]
      via ::, Null0
```


BGP-4 Extensions for IPv6 (RFC 2545)

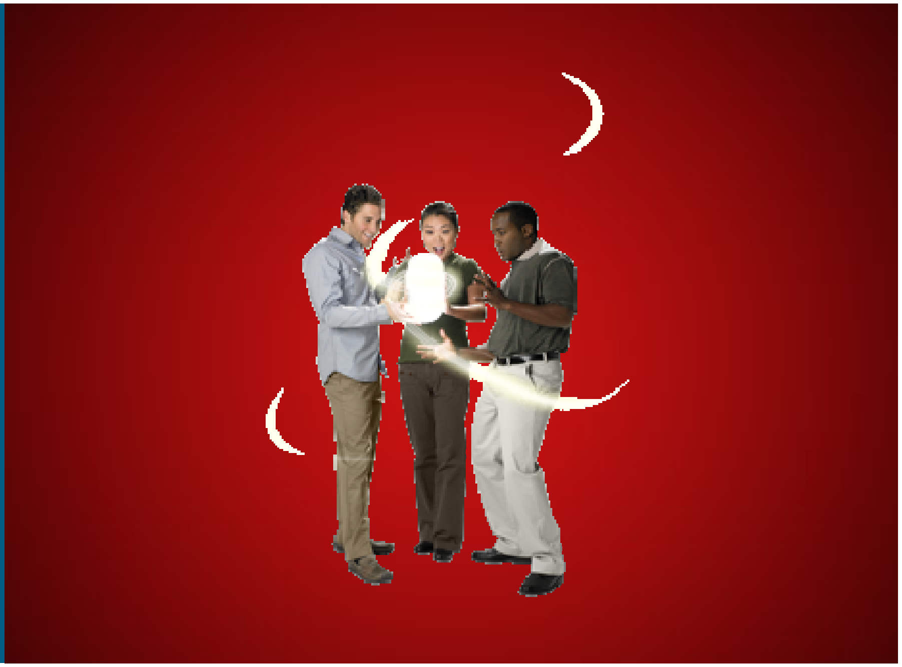

BGP-4 Configurations for IPv6 Non-Link-Local Peering

Router A

```
router bgp 1
no bgp default ipv4 unicast
bgp router-id 1.1.1.1
neighbor 2006:b00:ffff:2::2 remote-as 2
      address-family ipv6
      neighbor 2006:b00:ffff:2::2 activate
       network 2003:3:2::/64
       network 2003:3:3::/64
```

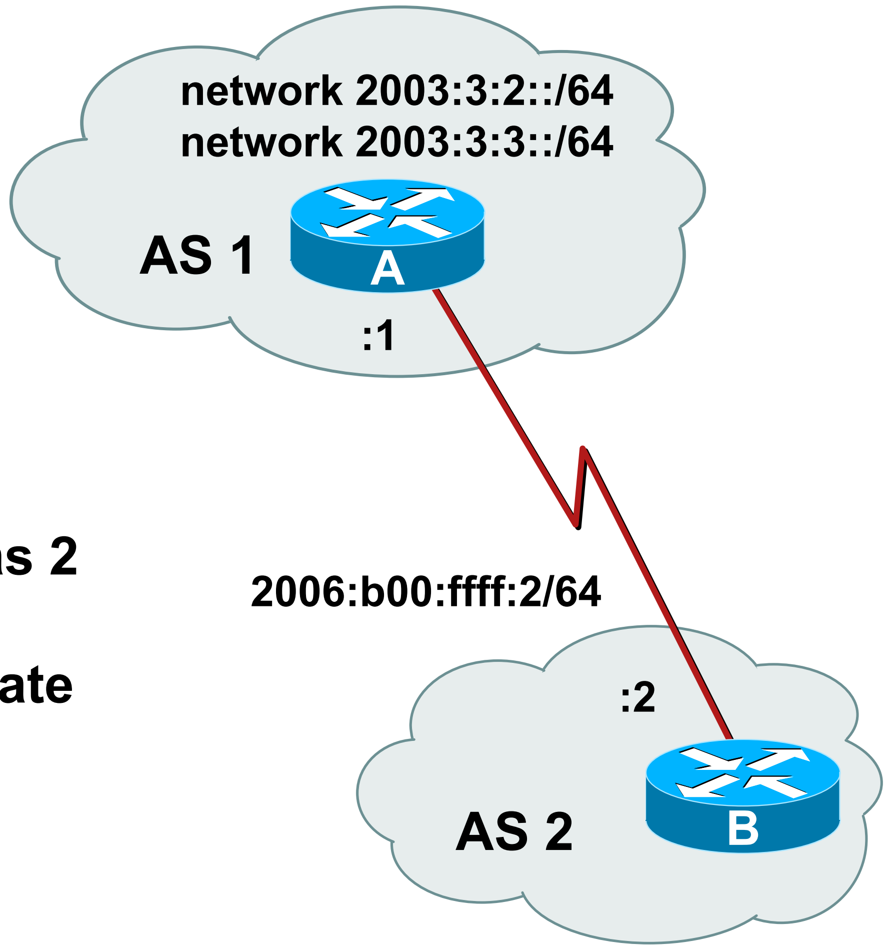

BGP-4 for IPv6 « Show Command »

```
RouterA#show bgp ipv6 2001:100:1:1::/64
BGP routing table entry for 2001:100:1:1::/64, version 71
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     1
  100
    2001:100:1:1::1 (FE80::A8BB:CCFF:FE01:F600) from FE80::A8BB:CCFF:FE01:F600%Ethernet0/0
(200.11.11.1)
      Origin incomplete, metric 0, localpref 100, valid, external
  Local
    :: from 0.0.0.0 (200.14.14.1)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
```


Deployment

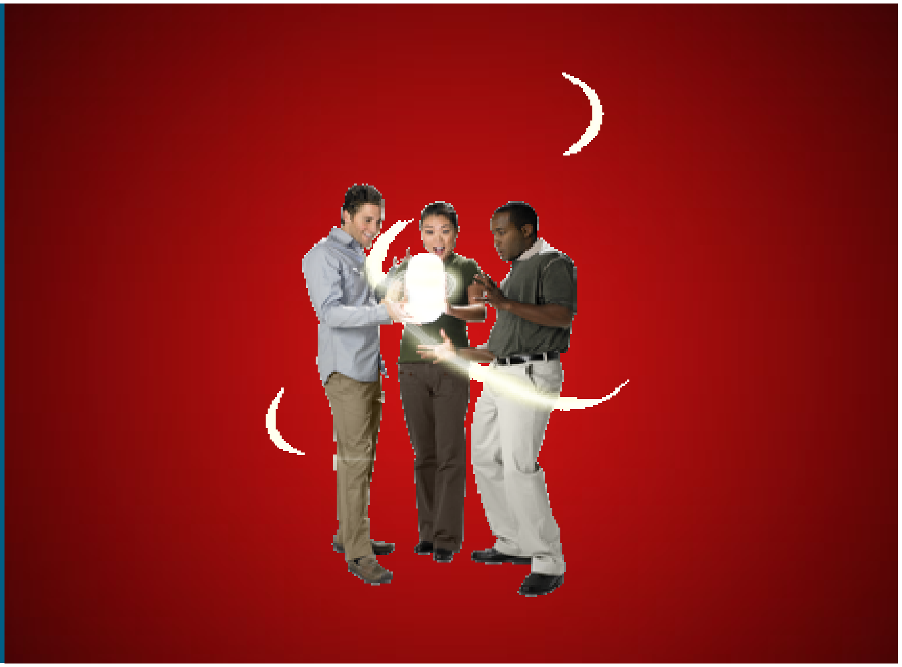

IPv4-IPv6 Transition/Coexistence

- A wide range of techniques have been identified and implemented, basically falling into three categories:
  1. Dual-stack techniques, to allow IPv4 and IPv6 to co-exist in the same devices and networks
  2. Tunneling techniques, to avoid order dependencies when upgrading hosts, routers, or regions
  3. Translation techniques, to allow IPv6-only devices to communicate with IPv4-only devices
- Expect all of these to be used, in combination

Host-Running Dual Stack

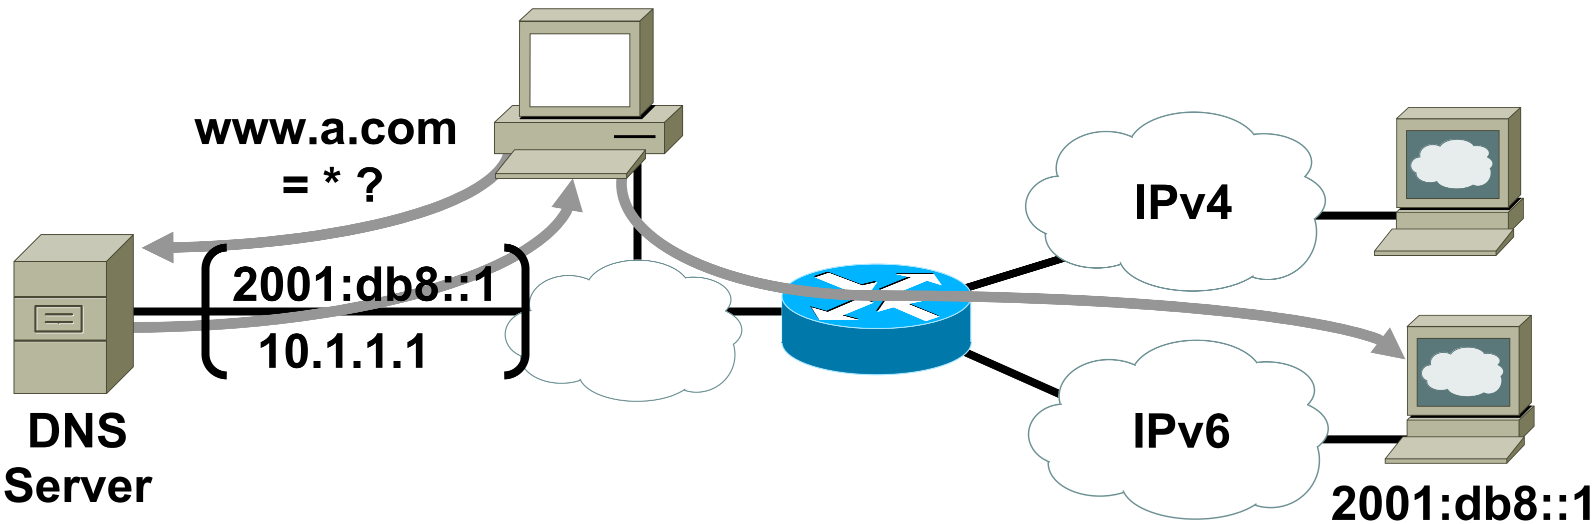

In a Dual-Stack Case, an Application that:

- Is IPv4 and IPv6-enabled
- Asks the DNS for all types of addresses
- Chooses one address and, for example, connects to the IPv6 address

Cisco IOS Dual-Stack Configuration

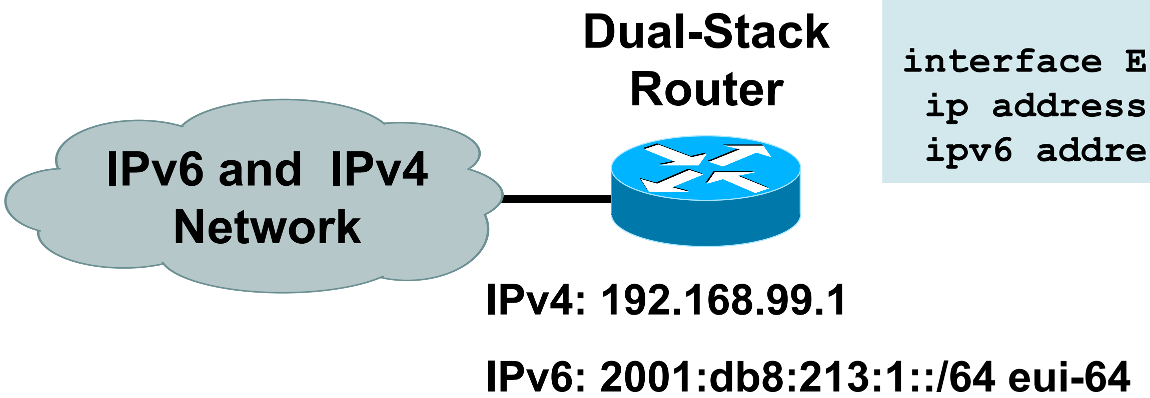


```
router#
ipv6 unicast-routing

interface Ethernet0
 ip address 192.168.99.1 255.255.255.0
 ipv6 address 2001:db8:213:1::/64 eui-64
```

Cisco IOS® Is IPv6-Enable:

- If IPv4 and IPv6 are configured on one interface, the router is dual-stacked
- Telnet, ping, traceroute, SSH, DNS client, TFTP, etc.

Tunneling

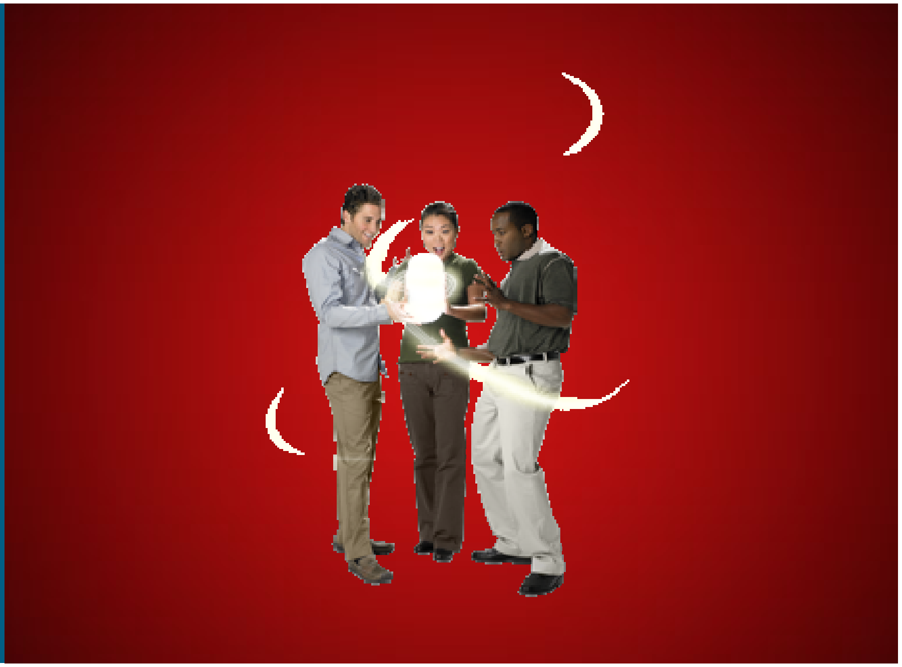

Tunneling

Many Ways to Do Tunneling

- Some ideas same as before

  GRE, MPLS, IP

- Native IP over data link layers

  ATM PVC, dWDM Lambda, Frame Relay PVC, serial, Sonet/SDH, Ethernet

- Some new techniques

  Automatic tunnels using IPv4, compatible IPv6 address, 6to4, ISATAP

Manually-Configured GRE Tunnel

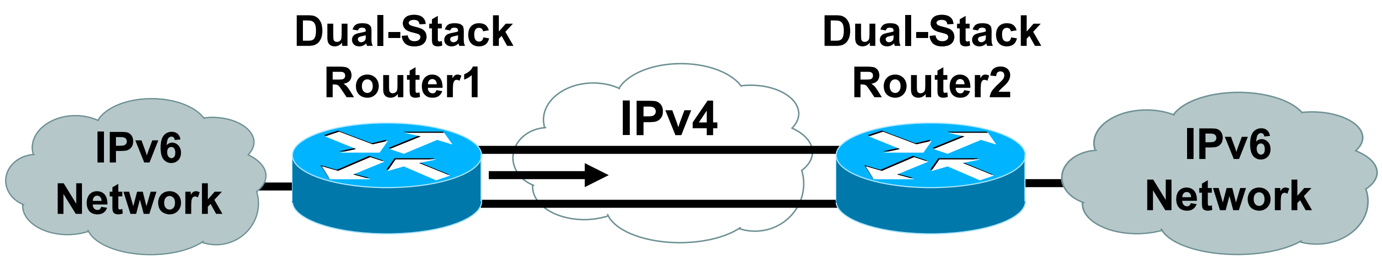

IPv4: 192.168.99.1
IPv6: 2001:db8:800:1::3

IPv4: 192.168.30.1
IPv6: 2001:db8:800:1::2

```
router1#

interface Tunnel0
 ipv6 enable
 ipv6 address 2001:db8:c18:1::3/128
 tunnel source 192.168.99.1
 tunnel destination 192.168.30.1
 tunnel mode gre ipv6
```

```
router2#

interface Tunnel0
 ipv6 enable
 ipv6 address 2001:db8:c18:1::2/128
 tunnel source 192.168.30.1
 tunnel destination 192.168.99.1
 tunnel mode gre ipv6
```

Manually-Configured IPv6 over IPv4 Tunnel

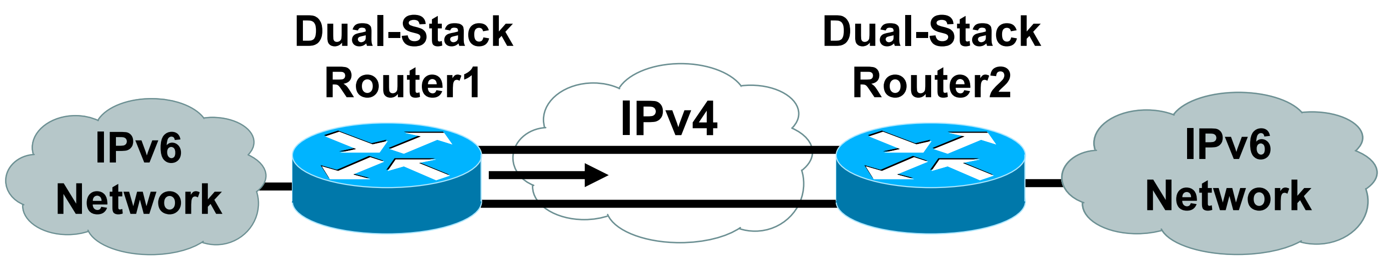

IPv4: 192.168.99.1
IPv6: 2001:db8:800:1::3

IPv4: 192.168.30.1
IPv6: 2001:db8:800:1::2

```
router1#

interface Tunnel0
 ipv6 enable
 ipv6 address 2001:db8:c18:1::3/127
 tunnel source 192.168.99.1
 tunnel destination 192.168.30.1
 tunnel mode ipv6ip
```

```
router2#

interface Tunnel0
 ipv6 enable
 ipv6 address 2001:db8:c18:1::2/127
 tunnel source 192.168.30.1
 tunnel destination 192.168.99.1
 tunnel mode ipv6ip
```


Labs

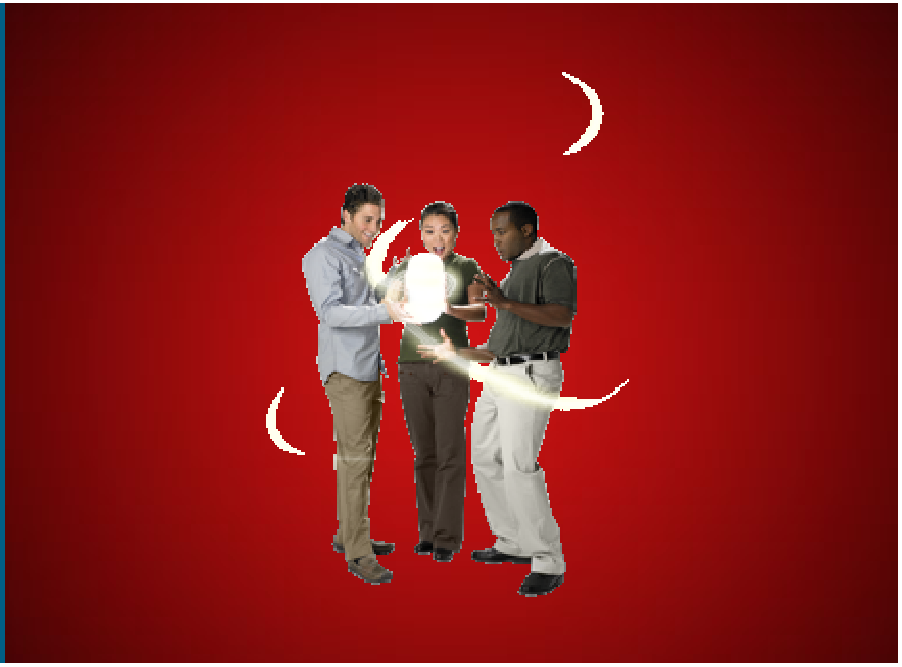

Network Diagram—IPv6 Lab

Sample Topology for a Single Pod

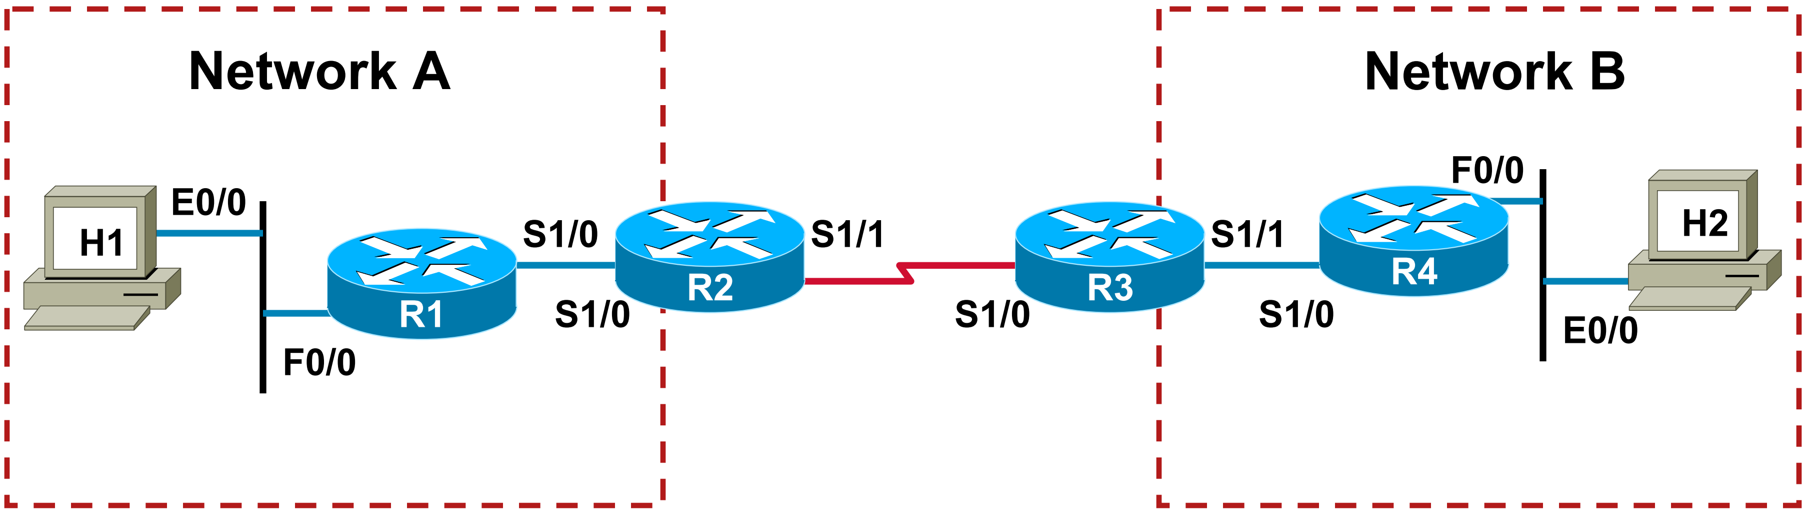

All Routers/Hosts Basic Configuration

Configure the Following on Your Router and Host:

- For example: R1 and H1

```
hostname R1 / H1
!
no ip domain-lookup
!
line con 0
no login
exec-timeout 0 0
privilege level 15
```

Lab 1: IPv6 Addressing

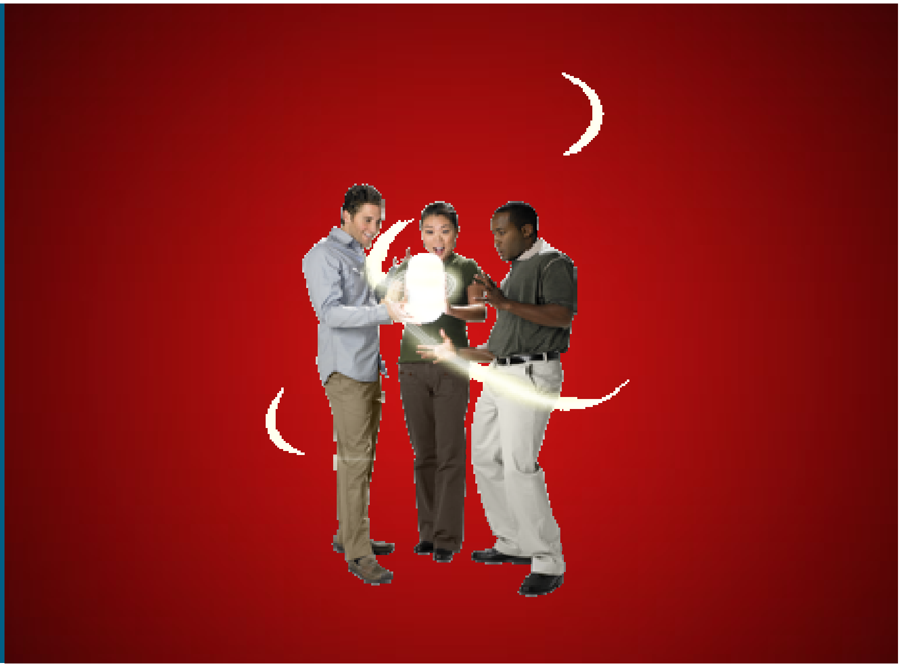

Lab 1: Task Summary

- Enable IPv6 on all four routers (R1 to R4)
- Configure hosts (H1 and H2) to use stateless address auto-configuration for IPv6
- Configure IPv6 addresses on Ethernet and serial interfaces using addresses listed in **Table 1**
- Use show commands to view IPv6 configuration and addresses
- Use ping to verify IPv6 connectivity

Lab 1: Configuring IPv6 Interfaces (1)

Table 1: IPv6 Prefixes for Ethernet and Serial Interfaces

Router Number	Global IPv6 Prefix	Unique-Local IPv6 Prefix
R1 (S1/0)	2001:DB8:1122:12::1/64	fc00:1:2:12::1/64
R1 (F0/0)	2001:DB8:1234:1::/64 eui-64	
R2 (S1/0)	2001:DB8:1122:12::2/64	fc00:1:2:12::2/64
R2 (S1/1)	2004:DB8:2233:23::2/64	
R3 (S1/1)	2001:DB8:3344:34::3/64	fc00:3:4:34::3/64
R3 (S1/0)	2004:DB8:2233:23::3/64	
R4 (F0/0)	2001:DB8:1234:4::/64 eui-64	
R4 (S1/0)	2001:DB8:3344:34::4/64	fc00:3:4:34::4/64

Lab 1: Configuring IPv6 Interfaces (2)

Task: Configuring the IPv6 Ethernet Interface

- On routers R1 and R2, go to global configuration mode

    #configure terminal

- Enable IPv6 unicast routing on the router

    (config)#ipv6 unicast-routing

- Configure the f0/0 interface

    (config)#interface f0/0

- Enable IPv6 on the interface

    (config-if)#ipv6 enable

    (config-if)#no shutdown

- Quit the configure mode

    (config-if)#end

- Verify that the Ethernet interface is configured

    #show ipv6 interface f0/0

- Identify the type of addresses that are configured

Lab 1: Configuring IPv6 Interfaces (3)

Task: Finding the Link-Local Address of a Host

- On hosts H1 and H2, go to global configuration mode

  #configure terminal

- Configure the host F0/0 interface

  (config)#interface F0/0

- Enable IPv6 on the interface

  (config-if)#ipv6 enable

  (config-if)#ipv6 address autoconfig

  (config-if)#no shutdown

- Quit the configure mode

  (config-if)#end

- Verify that the interface is configured

  #show ipv6 interface F0/0

Lab 1: Configuring IPv6 Interfaces (4)

Task: Verifying IPv6 Link-Local Connectivity

- On routers R1 and R4, enable console debugging of ICMP and Neighbor Discovery (ND) IPv6 packets

  #debug ipv6 icmp

  #debug ipv6 nd

- On the router, ping the local Ethernet interface of host specifying the link-local address in the ping command destination address

  #ping fe80::**<link local address of Host>**

  #ping ipv6 fe80::**<link local address of Host>**

- On the router, verify the list of IPv6 neighbors

  #show ipv6 neighbors

Lab 1: Configuring IPv6 Interfaces (5)

Task: Configuring a Static IPv6 Address with EUI-64

- On routers R1 and R4, configure a global IPv6 address on f0/0 interface using the /64 prefix of your subnet and the EUI-64 format
- Subnet = router number, e.g., **R1 = 2001:DB8:1234:1::/64** & **R2 = 2001:DB8:1234:2::/64)**

```
#configure terminal
(config)#interface f0/0
(config-if)#ipv6 address 2001:DB8:1234:<router #>::/64 eui-64
(config)#end
```

- Verify the configuration
- On the routers, disable all console debugging

```
#undebug all
```


Lab 1: Configuring IPv6 Interfaces (6)

Task: Configuring IPv6 addresses on Ethernet1/0 and Serial2/0 Interfaces

- On all four routers, configure a global and unique-local address on other interfaces by using the IPv6 prefix assigned. Use the router number for the host part. **(Refer to table 1 for details)**
- For example, on R1:

```
#configure terminal
(config)#interface f0/1
(config-if)#ipv6 enable
(config-if)#ipv6 address fc00:1:2:12::<router #>/64
(config-if)#ipv6 address 2001:DB8:1122:12::<router #>/64
(config-if)#no shutdown
```

- Verify connectivity by using ping to next-hop router

```
#ping <global or unique-local address of next-hop router>
#show ipv6 neighbors
```


Lab 2: Using Neighbor Discovery

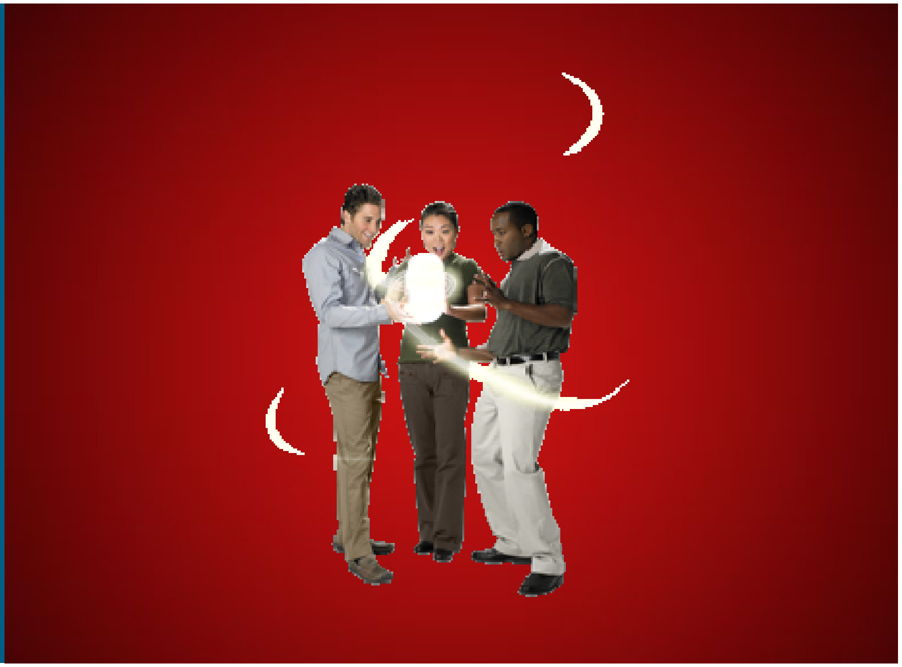

Lab 2: Task Summary

- Enable Router Advertisements (RA) for global and unique-local IPv6 prefixes on f0/0 (R1 and R2)
- Use show commands and debugs (debug ipv6 nd) to view IPv6 configuration
- Use ping to verify IPv6 connectivity
- Renumber the E0/0 interfaces on R1 and R2 by deprecating the old IPv6 prefixes and announcing a new IPv6 prefix

  Pick a new IPv6 global prefix (2003:DB8:ffff:**<router #>**::/64)

  Deprecate the old prefix; hint: change the valid and preferred lifetimes

Lab 2: Using Neighbor Discovery (1)

Task: Configuring Router Advertisements for Global Addresses

- Look at the configuration of hosts H1 and H2 and determine if they have received a global address?

  #show ipv6 interface F0/0

- Enable the debugging mode for IPv6 Neighbor Discovery (ND)

  #debug ipv6 nd

- On routers R1 and R4, go to configuration mode and then under Ethernet 0/0 interface

  #configure terminal

  (config)#interface f0/0

- Enable RAs by using the ND command with the subnet prefix assigned to your LAN. Because infinite lifetime is not desired, use five minutes (120 seconds) for lifetime (both preferred and valid).

  (config-if)#ipv6 nd prefix 2001:DB8:1234:**<router #>**::/64 120 120

- Verify that the hosts now have an IPv6 address that was automatically configured with this subnet prefix. Note that the previously configured link-local address is still present and valid. Also look at the valid and preferred lifetimes, are they decrementing?

  #show ipv6 interface f0/0

- Verify connectivity using the ping command on the router to the host using the newly assigned address of the host as the ping destination

  #ping **<global address of Host>**

Lab 2: Using Neighbor Discovery (2)

Task: Configuring Router Advertisements for Unique-Local Addresses

- Configure a unique-local address on f0/0 interface of Routers R1 and R4 by using the subnet prefix for your LAN. Use EUI-64 format.

  #configure terminal

  (config)#interface f0/0

  (config-if)#ipv6 address fc00:0:0:**<router #>**::/64 eui-64

- Enable RAs by using the ND command with subnet prefix assigned to your LAN for unique-local addresses. Use five minutes (120 seconds) for the lifetime.

  (config-if)# ipv6 nd prefix fc00:0:0:**<router #>**::/64 120 120

- Change the RA interval to 30 seconds

  (config-if)# ipv6 nd ra interval 30

- Verify that hosts H1 and H2 now have an IPv6 address configured with this subnet prefix

  #show ipv6 interface f0/0

- Verify connectivity to the host using the ping command on the router to the host using the newly assigned address of the host as the ping destination address

  #ping **<unique-local address of Host>**

Lab 2: Using Neighbor Discovery (3)

Task: Renumbering the Local Network on the Router

- On routers R1 and R2, configure the new address of workgroup f0/0 interface by using the new global subnet prefix assigned to your LAN. Use EUI-64 format.

  #configure terminal

  (config)#interface f0/0

  (config-if)#ipv6 address 2003:DB8:ffff:**<router #>**::/64 eui-64

- Enable RAs by using the ND command with the new subnet prefix assigned to your LAN. Use five minutes (240 seconds) for the lifetime.

  (config-if)# ipv6 nd prefix 2003:DB8:ffff:**<router #>**::/64 240 240

- Verify that hosts H1 and H2 have a new address with the new prefix

  #show ipv6 interface ethernet0/0

Lab 2: Using Neighbor Discovery (4)

- Modify the Neighbor Advertisement (NAs) for 2001:DB8:1234:**<router #>**::/64 (the old prefix) by making the valid and preferred lifetimes equal to 60 and 0 respectively

  #configure terminal

  (config)#interface ethernet0/0

  (config-if)# ipv6 nd prefix 2001:DB8:1234:**<router #>**::/64 20 0

- Verify that the hosts deprecate the use of the “old” prefix and prefers the “new” one

  #show ipv6 interface ethernet0/0

- Verify connectivity to the host by using the ping command on the router using the new assigned address of host

  #ping **<new global address of Host>**

  #undebug all

Lab 3: IPv6 Static Routing

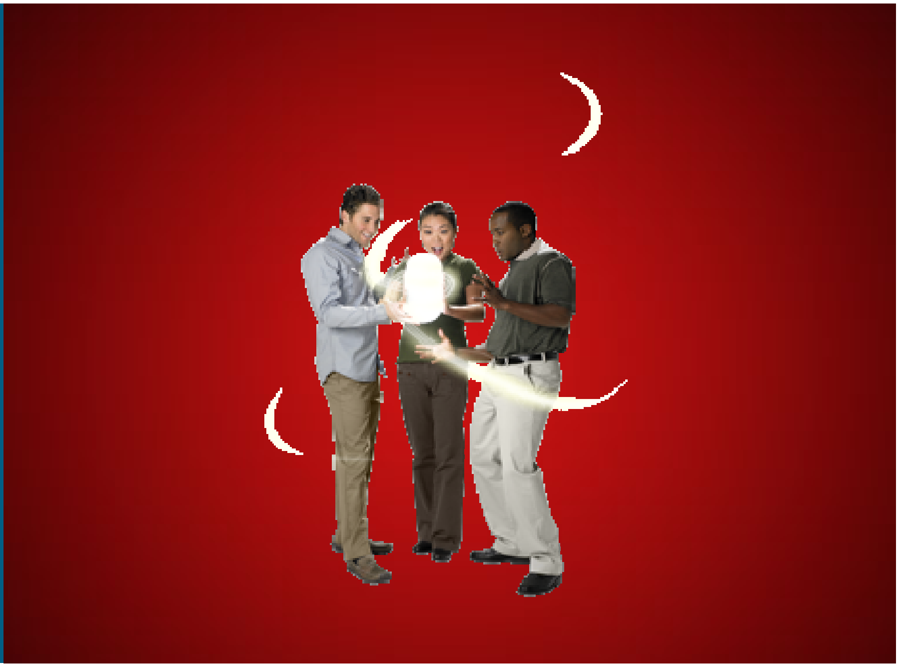

Lab 3: Task Summary

- Configure IPv6 default static routes on R1 and R4
- Configure IPv6 static routes on R3 and R4
- Use ping to verify IPv6 connectivity between R1 and R4
- Use show commands to view IPv6 routing table

Lab 3: IPv6 Static Routing (1)

Task: Configuring IPv6 Static Routing on Routers

- Configure static default routes on R1 and R4 pointing to R2 and R3 respectively (see table 1 for addressing details)
- On R1:

  (config)#ipv6 route ::/0 2001:DB8:1122:12::2
- On R4:

  (config)#ipv6 route ::/0 2001:DB8:3344:34::3

  Configure static routes on R2 pointing to R4's E1/0 network; R3 pointing to R1's E1/0 network. (see table 1 for addressing details)
- On R2:

  (config)#ipv6 route 2001:DB8:3344:34::/64 2004:DB8:2233:23::3
- On R3:

  (config)#ipv6 route 2001:DB8:1122:12::/64 2004:DB8:2233:23::2

Lab 3: IPv6 Static Routing (2)

Task: Verifying IPv6 Connectivity Using Static Routes

- Verify the routing table

  #show ipv6 route

- Verify IPv6 connectivity using the ping command using the peer’s IPv6 address

  #ping **<your peer’s Global IPv6 Address>**

- Can you ping from R1 to R4 and vice versa? If not, troubleshoot why the ping is not working.
- Hint: use “debug ipv6 icmp” and “debug ipv6 packet” on R2 and R3 to see if traffic is being forwarded

Lab 4: Routing with OSPFv3

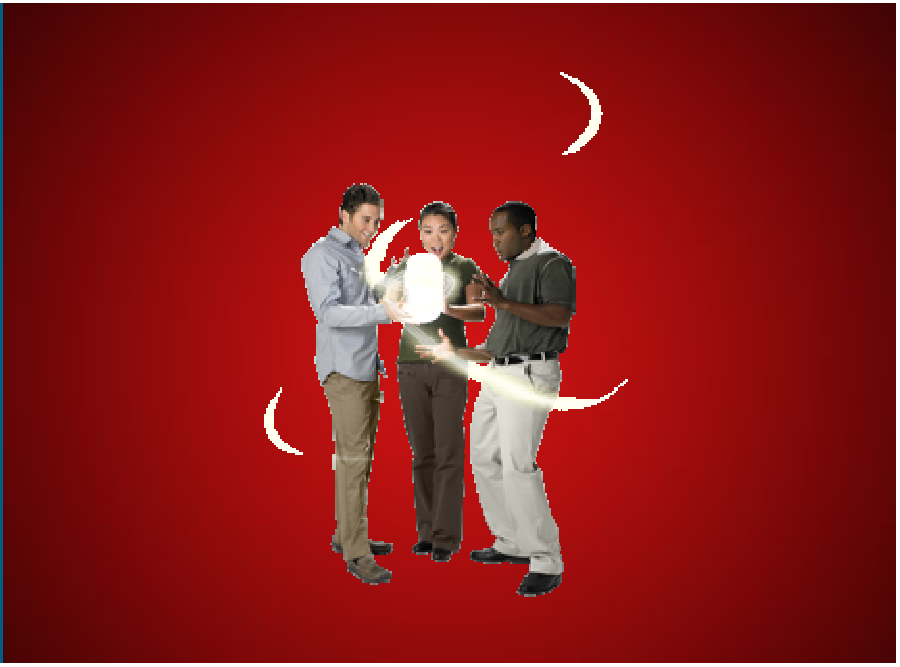

Lab 4: Task Summary

- Enable debugs for IPv6 routing and IPv6 OSPFv3

    #debug ipv6 routing

    #debug ipv6 ospf event

- Configure OSPFv3 on all four routers

    Look at topology diagram for details

    Use 192.168.30.xx for router-id (where xx is the router number)

- Remove all IPv6 static and default routes on all routers
- Use ping to verify IPv6 connectivity between H1 and H2
- Use show commands to view IPv6 routing table

    #show ipv6 route

    #show ipv6 ospf neighbor

    #show ipv6 ospf database

- Answer review questions at the end of this section

OSPFv3 Configuration Example

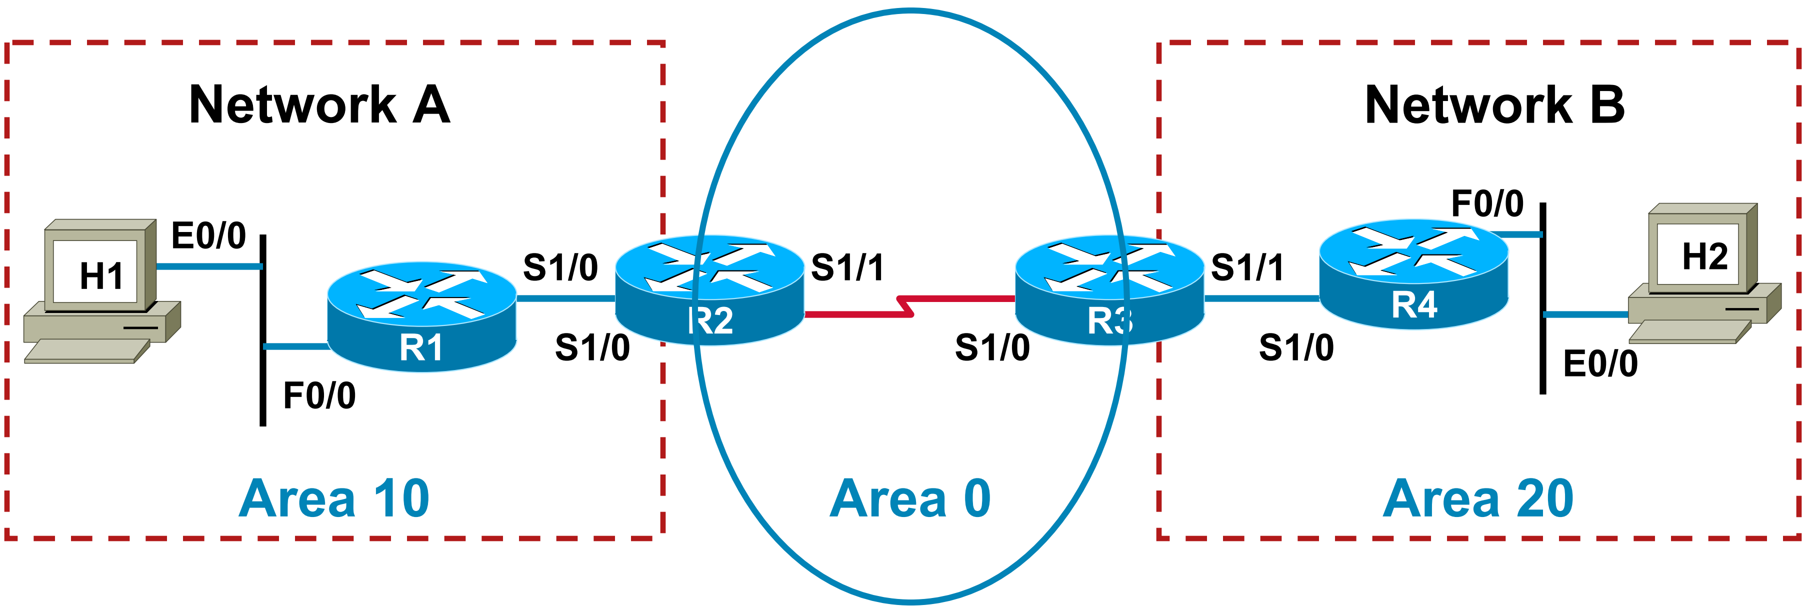

Lab 4: Routing with OSPFv3 (1)

Task: Configuring OSPFv3

- Enable the debugging mode for IPv6 routing and IPv6 OSPFv3

  #debug ipv6 routing

  #debug ipv6 ospf event

- Configure an OSPFv3 process 1 on the routers with a router-ID

  (config)#ipv6 router ospf 1

  (config-rtr)#router-id 192.168.30.xx **(where xx is your router #)**

- Enable OSPFv3 on each interface on all four routers

  For example on R1:

  (config)#interface ethernet0/0

  (config-if)#ipv6 ospf 1 area **10**

  (config-if)#interface ethernet1/0

  (config-if)#ipv6 ospf 1 area **10**

OSPFv3 Configuration—Router R1 (2)

```
Router R1

ipv6 unicast-routing

interface Ethernet0/0
 no ip address
 ipv6 ospf 1 area 10

interface Ethernet1/0
 no ip address
 ipv6 ospf 1 area 10

ipv6 router ospf 1
 router-id 192.168.30.1
```


OSPFv3 Configuration—Router R2 (3)

```
Router R2

ipv6 unicast-routing

interface Ethernet1/0
 no ip address
 ipv6 ospf 1 area 10

interface Serial2/0
 no ip address
 ipv6 ospf 1 area 0


ipv6 router ospf 1
 router-id 192.168.30.2
```


Lab 4: Routing with OSPFv3 (4)

Task: Configuring OSPFv3

- Verify that the routing table is updated by entries coming from the other workgroup routers, OSPFv3 Neighbor, OSPFv3 database

  #show ipv6 route

  #show ipv6 ospf neighbor

  #show ipv6 ospf database

- Do you see the OSPFv3 routes in the routing table on R2 and R3? Why or why not?
- Remove the static routes on all 4 routers

  #no ipv6 route **<network> <next hop>**

- Verify IPv6 connectivity using the ping command on Host. Use the remote host's global IPv6 address as the destination address of the ping command

  #ping **<Global IPv6 address of the remote host>**

  #undebug all

- On R4 define a loopback4 with address 2040::4/64 and do a redistribute connected under ipv6 router ospf. Check the database of R1 and R2 with show ipv6 ospf database external. Do you see the loopback address?
- On R1 and R2 under ipv6 router ospf configure "area 10 stub". Do you see the external routes now? Why not?

Lab 5: Routing with BGP

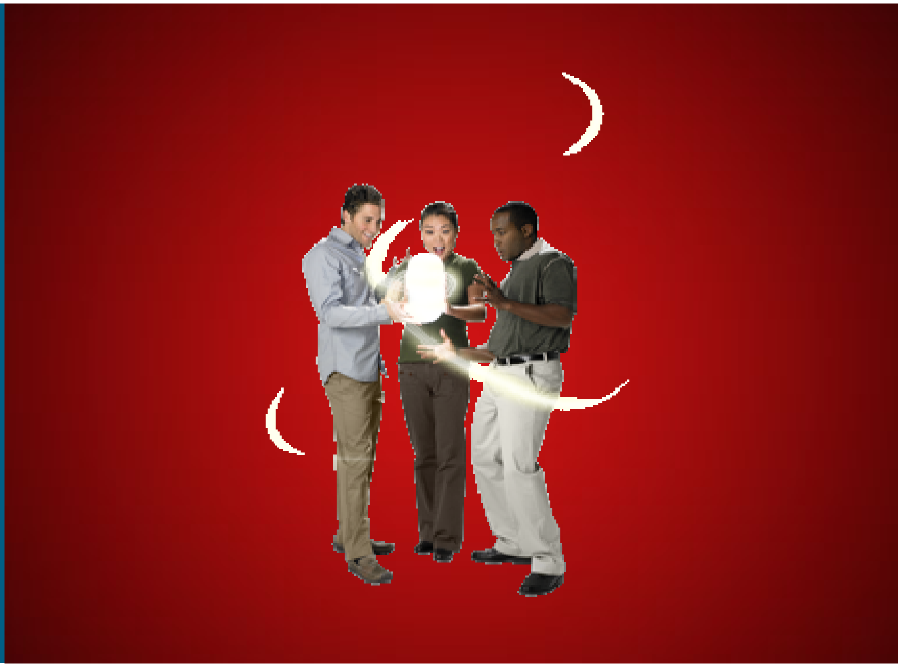

Lab 5: Task Summary

- Enable debugging for BGP

  #debug bgp ipv6 unicast
- Configure **BGP** AS number and router-id on all four routers using **table 2**
- Activate iBGP peering between R1 and R2 and R3 and R4 using the interface global IPv6 address. Announce the global prefix and unique-local prefix.
- Activate **eBGP** peering between R2 and R3 using their global IPv6 addresses
- On R2 and R3, filter the unique-local prefix (fc00::/10) so that you do not receive them from the peers
- Use show commands to view the BGP configuration

  #show bgp ipv6 unicast summary

  #show bgp ipv6 unicast neighbors **<IP address>** advertised-routes

  #show bgp ipv6 unicast neighbors **<IP address>** routes

  #show ipv6 route bgp
- Use show bgp ipv6 unicast to see the routes in the BGP table
- Answer review questions at the end of this section

BGP Configuration Example

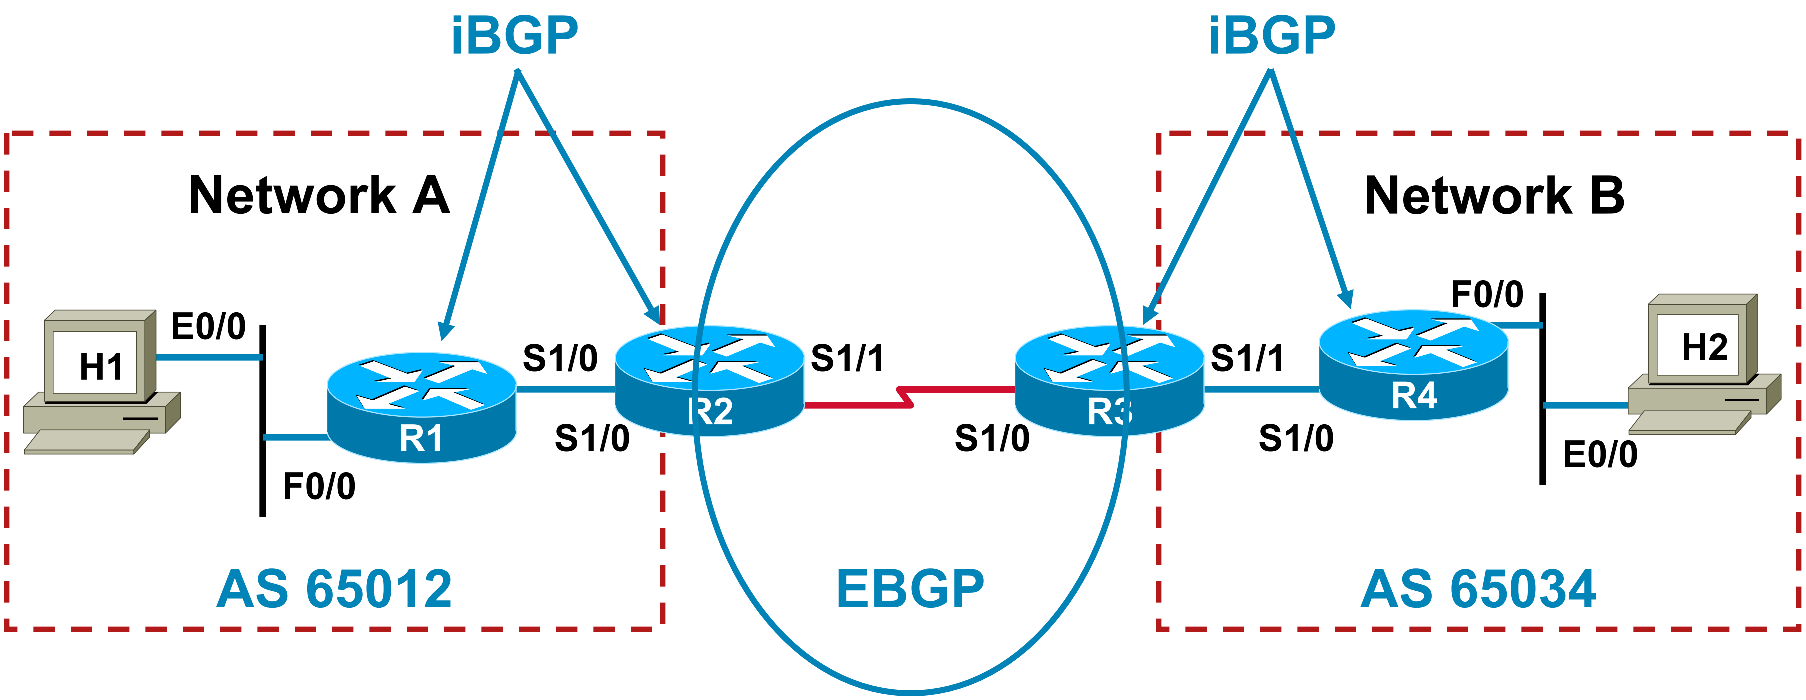

Lab 5: Routing with BGP (1)

Task: Configuring BGP

- Refer to **table 2** for the router address and autonomous system (AS) number
- Enable the debugging of BGP

  #debug bgp ipv6

- Define your router as a BGP router with your AS number

  (config)#router bgp **<AS number>**

- Because IPv4 has not been configured yet and BGP protocol uses the IPv4 address as an identifier, a specific command to identify the router-id is needed. Use **table 2** to find the router-id. Then configure the router BGP router-id.

  (config-router)#bgp router-id **<your router-id>**

- Deactivate the IPv4 default peering

  (config-router)#no bgp default ipv4-unicast

Lab 5: Routing with BGP (2)

Table 2: Assigned AS Number and Router ID

Router Number	AS Number	Router ID
R1	65012	192.168.60.1
R2	65012	192.168.60.2
R3	65034	192.168.60.3
R4	65034	192.168.60.4

Lab 5: Routing with BGP (3)

- Activate **iBGP** peering between R1 and R2 and R3 and R4 using the Ethernet 1/0 global IPv6 address. Announce your global prefix and your unique-local prefix.

  (config-router)#neighbor **<peer's S1/0 global IPv6 address >** remote-as **<AS number>**

  (config-router)#neighbor **<peer's S1/0 global IPv6 address >** update-source S1/0

  (config-router)#address-family ipv6

  (config-router-af)#neighbor **<peer's S1/0 global IPv6 address>** activate

  (config-router-af)#network **<F0/0 global network prefix>**/64

  (config-router-af)#network **<S1/0 unique local prefix>**/64

  R1 and R4 only

  (config-router-af)#exit

- Look at the BGP routing table to determine if you received BGP routes from your neighbor

  #show bgp ipv6 unicast summary

  #show bgp ipv6 unicast neighbors **<IP address>** advertised-routes

  #show bgp ipv6 unicast neighbors **<IP address>** routes

  #show ipv6 route bgp

Lab 5: Routing with BGP (4)

- Activate **eBGP** peering between R2 and R3 using their global IPv6 addresses

  (config-router)#neighbor **<peer's global IPv6 address>** remote-as **<AS number>**

  (config-router)#neighbor **<peer's global IPv6 address>**

  (config-router)#address-family ipv6

  (config-router-af)#neighbor **<peer's global IPv6 address>** activate

  (config-router-af)#end

- Look at the BGP routing table to determine if you received BGP routes from your neighbor

  #show bgp ipv6 unicast summary

  #show bgp ipv6 unicast neighbors **<IP address>** advertised-routes

  #show bgp ipv6 unicast neighbors **<IP address>** routes

  #show ipv6 route bgp

Lab 5: Routing with BGP (5)

Task: Filtering BGP Routes

- On R2 and R3, filter the unique-local prefix (fc00::/10) so that you do not receive them from the peers

  (config)#ipv6 prefix-list lab5 deny fc00::/10 le 128

  (config)#ipv6 prefix-list lab5 permit ::/0 le 64

  (config)#router bgp **<your AS number>**

  (config-router)#address-family ipv6

  (config-router-af)#neighbor **<peer's global IPv6 address>** prefix-list lab5 in

  (config-router-af)#exit

- On R1 and R4, verify that you do not receive any unique-local routes via BGP from the remote peer. Only the global IPv6 prefixes should be received.

  #clear bgp ipv6 unicast **<peer AS number>** in

  #show bgp ipv6 unicast

  #show bgp ipv6 unicast neighbors **<IP address>** advertised-routes

  #show bgp ipv6 unicast neighbors **<IP address>** routes

  #show ipv6 route bgp

Lab 6: Manual Tunneling in IPv6

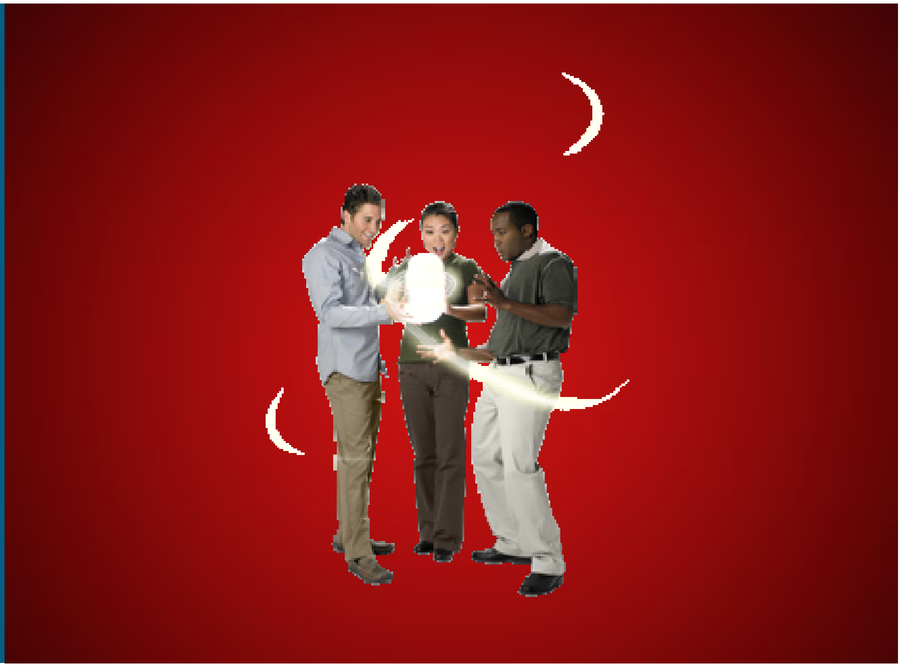

Lab 6: Task Summary

- Disable IPv6 routing protocols on all routers
- Remove the IPv6 addresses on serial2/0 interface of R2 and R3
- Configure IPv4 address on serial2/0 interface of R2 and R3, use 192.168.30.<router xx>/24
- Configure tunnel interfaces on R2 and R3 and configure them for manual tunneling
- Enable RIPng on both tunnel endpoints and other relevant interfaces on all four routers
- Enable EIGRPv6 over the tunnel
- Enable mutual redistribution between EIGRPv6 and RIPng on R2 and R3
- Use show commands to view the IPv6 routing table and tunnel interface statistics
- Use ping to verify IPv6 connectivity between H1 and H2
- Answer review questions at the end of section 7

Manual Tunnel Configuration Example

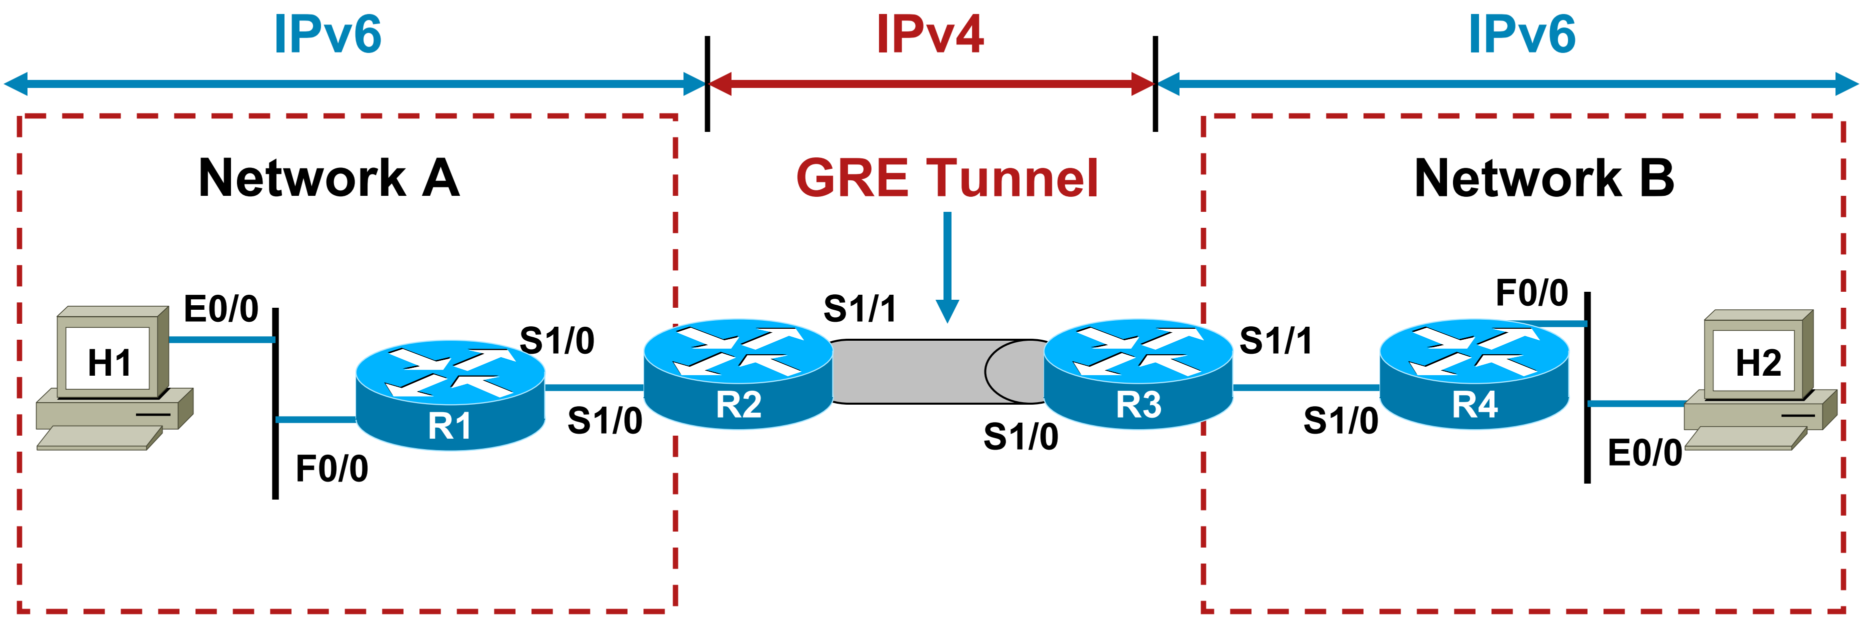

Lab 6: Using Tunnels (Manual-1)

Task: Configuring the Tunnel

- First disable both OSPFv3 and BGP on all routers
- Remove the IPv6 address on Serial interface between R2 and R3

  (config)#no ipv6 router ospf 1

  (config)#no router bgp <your AS>

  (config)#interface Serialxx

  (config-if)#no ipv6 address

  (config-if)#no ipv6 enable

- Configure IPv4 address on interface between R2 and R3

  (config)#interface xxx

  (config-if)#ip address 192.168.30.**<router #>** 255.255.255.0

Lab 6: Using Tunnels (Manual-2)

- Configure a tunnel interface using tunnel0

  (config)#interface tunnel 0

- Use an IPv6 unnumbered address from the S1/0 @R2, S1/1@R3

  (config-if)#ipv6 unnumbered S1/0@R2 (S1/1@R3)

- Identify the source and destination address of the tunnel, followed by the tunnel mode

  (config-if)#tunnel source Serial1/0@R2 (S1/1@R3)

  (config-if)#tunnel destination **<peer’s IPv4 address>**

  (config-if)#tunnel mode gre ip

- Enable RIPng on all four routers

  (config)#ipv6 router rip lab

  (config-rtr)#redistribute connected

- Enable RIPng on Ethernet 0/0 and Ethernet 1/0 interfaces on R1 and R4, and Ethernet 1/0 interfaces on R2 and R3

  (config-if)#ipv6 rip lab enable

Lab 6: Using Tunnels (Manual-3)

- Configure OSPFv3 on the routers R2 and R3 with a Router-ID

  (config)#ipv6 router ospf 1

  (config-rtr)#router-id 192.168.30.xx **(where xx is your router #)**

- Enable OSPFv3 on tunnel interfaces on R2 and R3

  (config)#interface tunnel 0

  (config-if)#ipv6 ospf 1 are 0

- Enable mutual redistribution on R2 and R3 between RIPng and OSPFv3

  (config)#ipv6 router ospf 1

  (config)#redistribute rip lab

  (config)#ipv6 router rip lab

  (config)#redistribute ospf 1

Lab 6: Using Tunnels (Manual-4)

- Verify that RIPng updates are carried over the tunnel

  #show ipv6 route

- Look at Tunnel 0 debugs and statistics

  #debug tunnel

  #show interface tunnel 0 accounting

- Verify connectivity between R1 & R4 using the ping command

  #ping **<global IPv6 address of Ethernet0/0 of peer>**

- Verify connectivity between H1 and H2 using the ping command

  #ping **<global IPv6 address of peer Host>**

Thank You

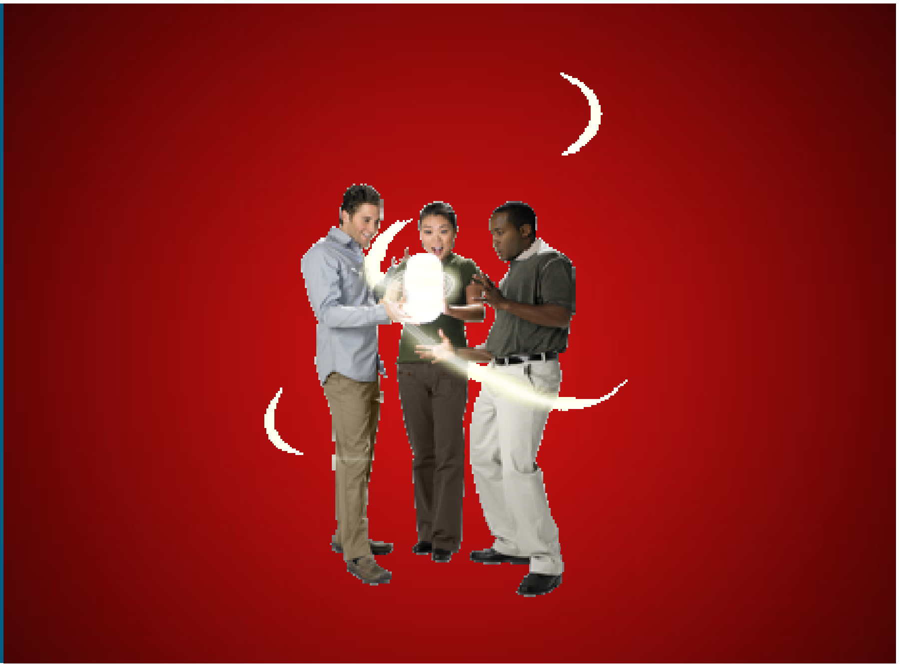

Recommended Reading

Available Online at the Cisco Press