Introduction to IPv6

Villayat Muhammad

Service Provider Solution Architect

Why IPv6?

A Need for IPv6?

- IETF IPv6 WG began in early 90s, to solve addressing growth issues, but

  CIDR, NAT,…were developed

- IPv4 32 bit address = 4 billion hosts

  ~40% of the IPv4 address space is still unused which is different from unallocated

  The rising of Internet connected device and appliance will eventually deplete the IPv4 address space

- IP is everywhere

  Data, voice, audio and video integration is a reality

  Regional registries apply a strict allocation control

- So, only compelling reason: More IP addresses

Why Not NAT

- It was created as a temp solution
- NAT breaks the end-to-end model
- Growth of NAT has slowed down growth of transparent applications
- No easy way to maintain states of NAT in case of node failures
- NAT break security
- NAT complicates mergers, double NATing is needed for devices to communicate with each other

IPv6 Technology

IPv4 and IPv6 Header Comparison

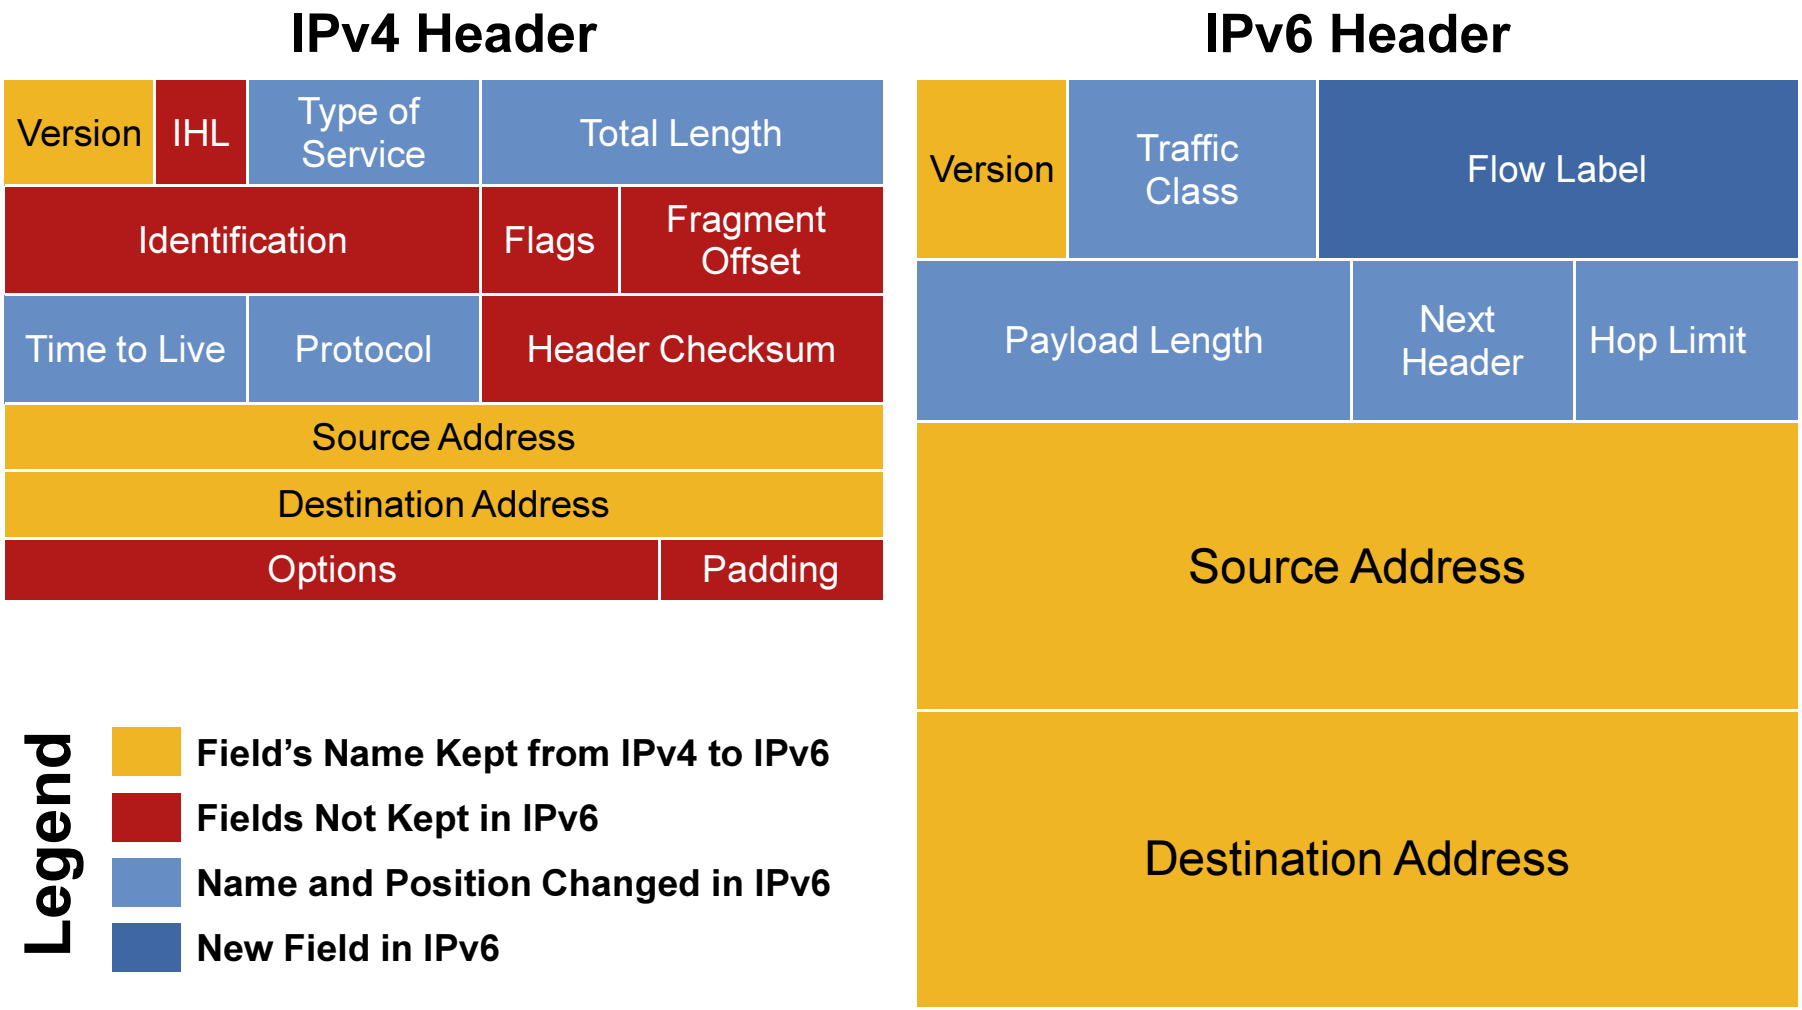

IPv6 Header New Field—Flow Label (RFC3697)

20-Bit Flow Label Field to Identify Specific Flows Needing Special QoS

- Flow classifiers had been based on 5-tuple: Source/destination address, protocol type and port numbers of transport
- Some of these fields may be unavailable due to fragmentation, encryption or locating them past extension headers
- With flow label, each source chooses its own flow label values; routers use source addr + flow label to identify distinct flows
- Flow label value of 0 used when no special QoS requested (the common case today)

IPv6 Header

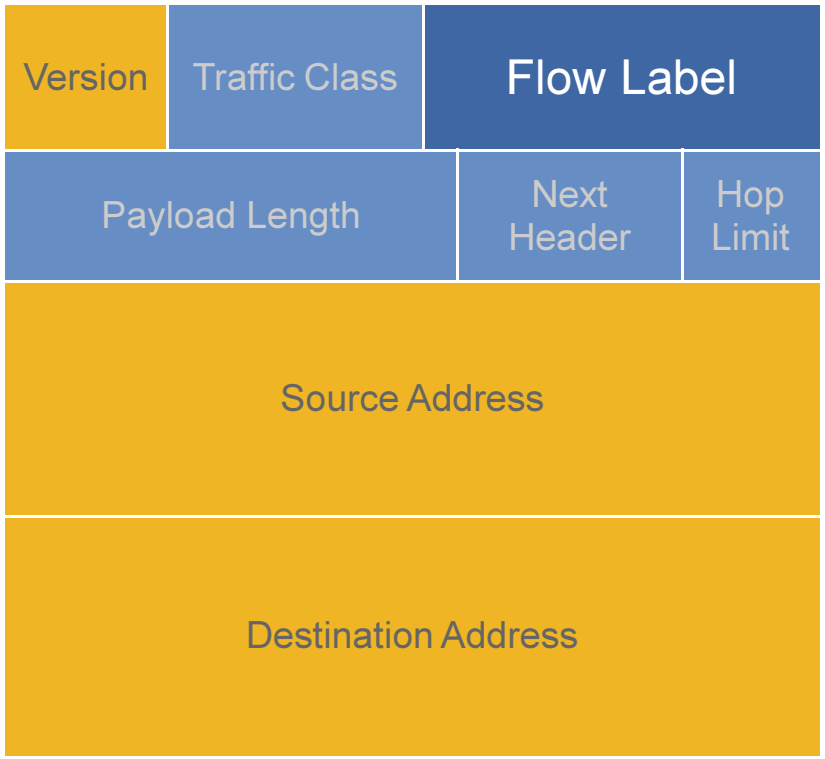

Extension Headers

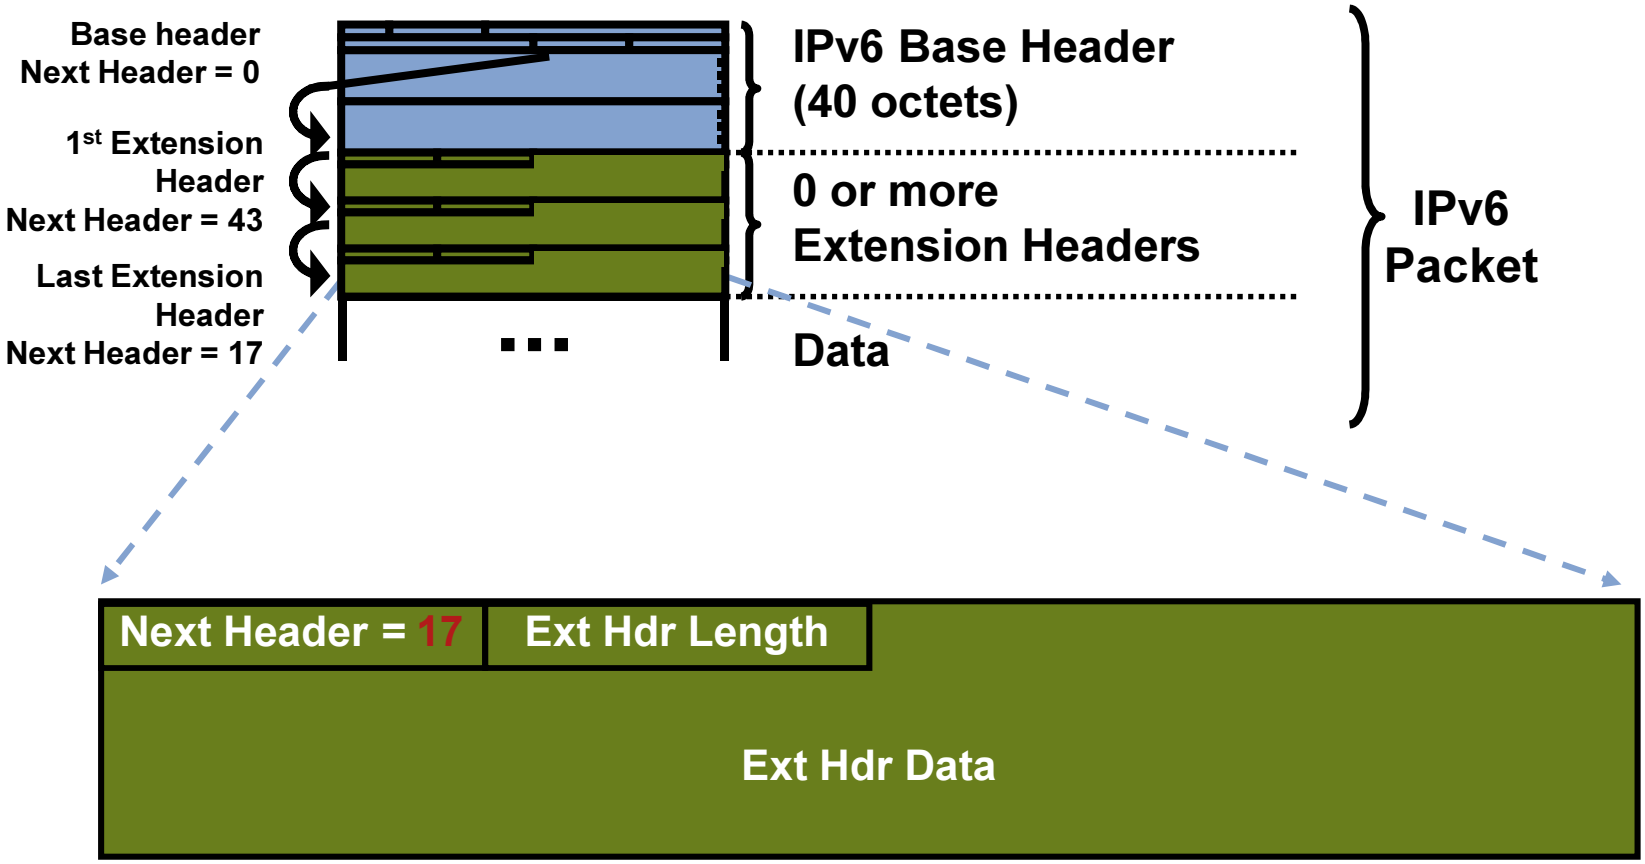

MTU Issues

- Minimum link MTU for IPv6 is 1280 octets (vs. 68 octets for IPv4)

  => on links with MTU < 1280, link-specific fragmentation and reassembly must be used

- Implementations are expected to perform path MTU discovery to send packets bigger than 1280
- Minimal implementation can omit PMTU discovery as long as all packets kept ≤ 1280 octets

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

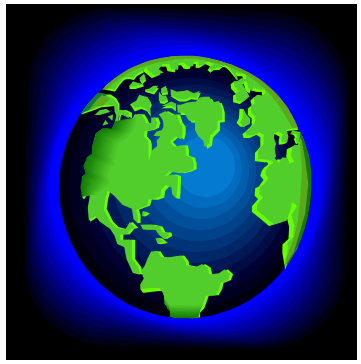

World's population is approximately 6.5 billion

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

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

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

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:adfc**::/64 ≠ 2001:db8:12:adfc::/64

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

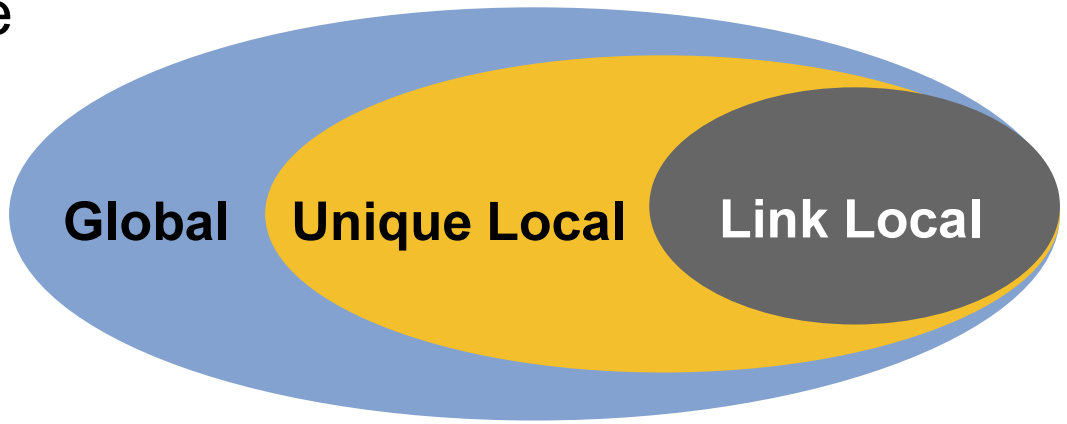

Types of IPv6 Addresses

- 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

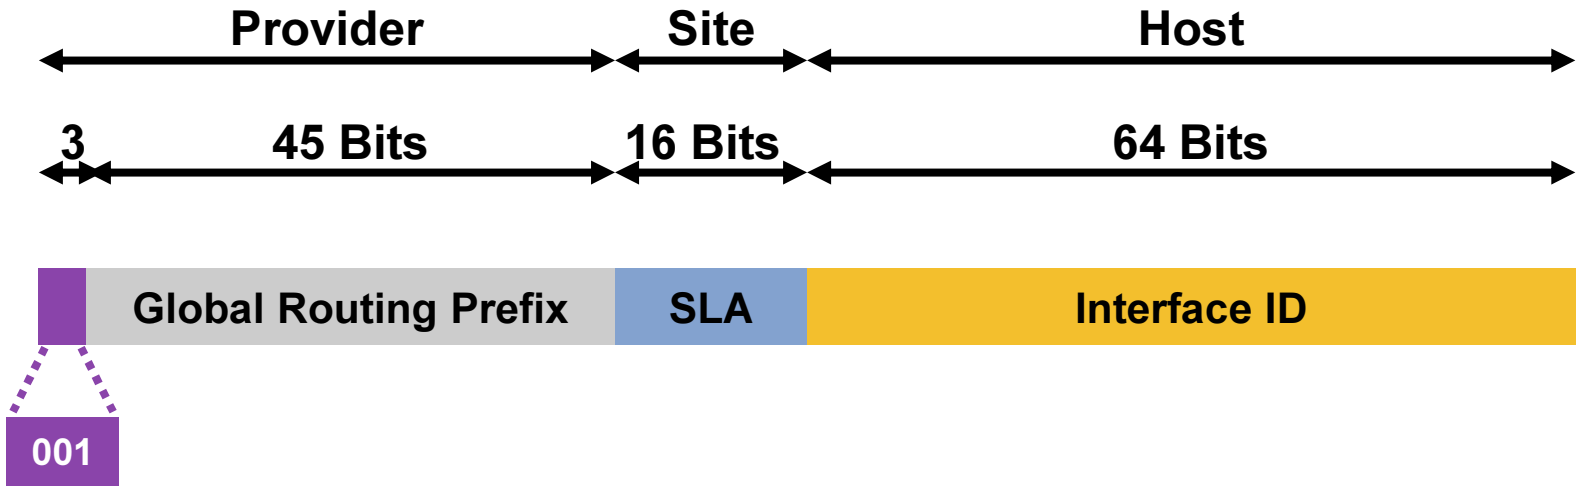

Aggregatable Global Unicast Addresses Are:

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

Unique-Local

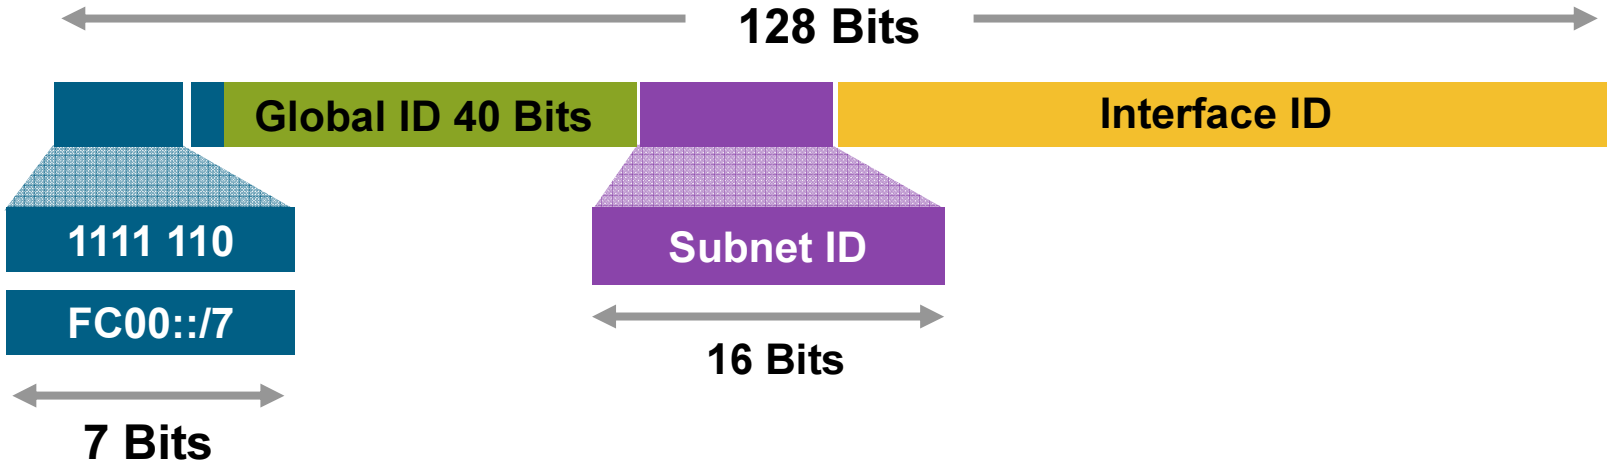

Unique-Local Addresses Used for:

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

Link-Local

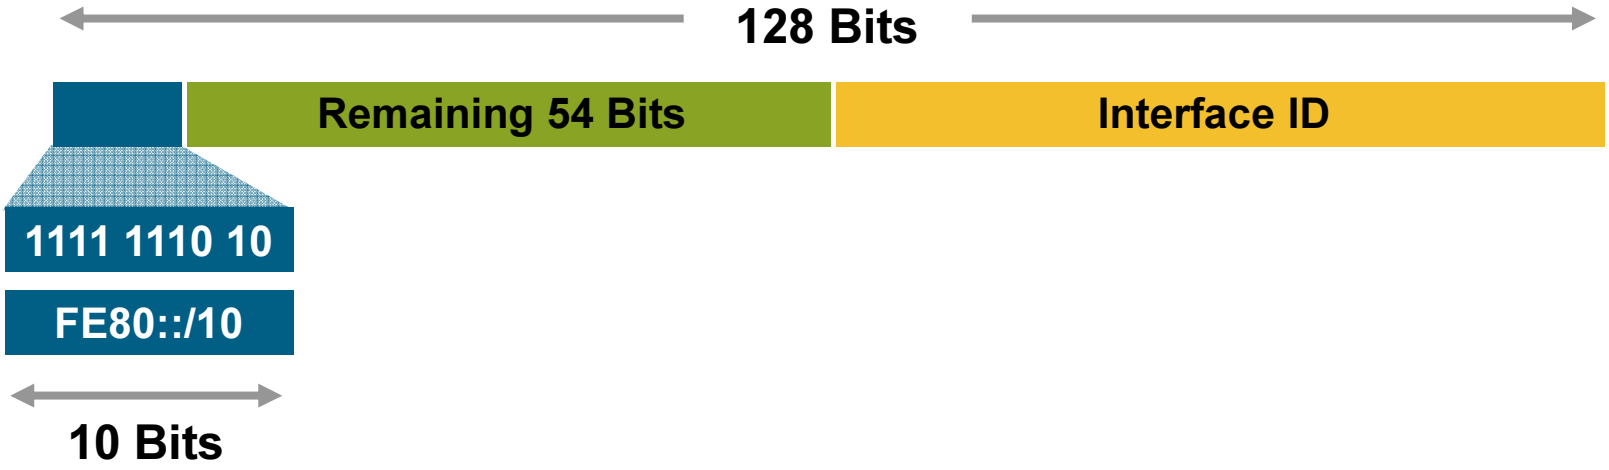

Link-Local Addresses Used for:

- 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 Multicast Address

- IP multicast address has a prefix FF00::/8 (1111 1111); the second octet defines the lifetime and scope of the multicast address

8-bit	4-bit	4-bit	112-bit
1111 1111	Lifetime	Scope	Group-ID

Lifetime	
0	If Permanent
1	If Temporary

Scope	
1	Node
2	Link
5	Site
8	Organization
E	Global

Some Well Known Multicast Addresses

Address	Scope	Meaning
FF01::1	Node-Local	All Nodes
FF02::1	Link-Local	All Nodes
FF01::2	Node-Local	All Routers
FF02::2	Link-Local	All Routers
FF05::2	Site-Local	All Routers
FF02::1:FFXX:XXXX	Link-Local	Solicited-Node

- Note that 02 means that this is a permanent address and has link scope
- More details at http://www.iana.org/assignments/ipv6-multicast-addresses

Solicited-Node Multicast Address

- For each unicast and anycast address configured there is a corresponding solicited-node multicast
- This is specially used for two purpose, for the replacement of ARP, and DAD
- 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:

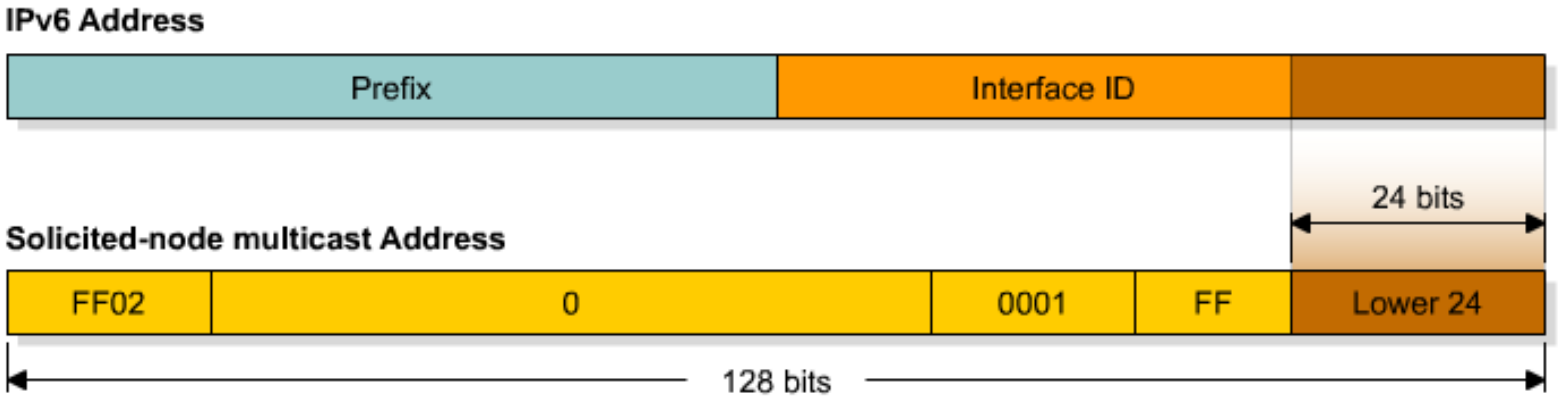

Router Interface

```
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:FF3A: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#
```

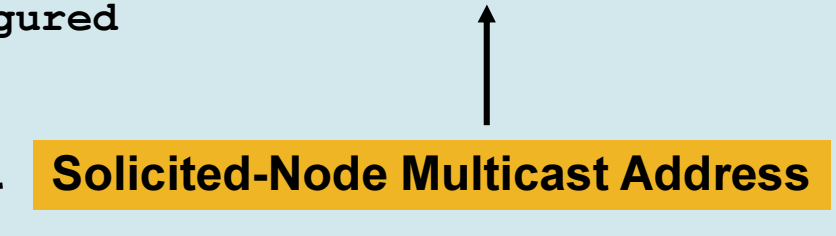

IPv6 Prefix Allocation Hierarchy and Policy Example

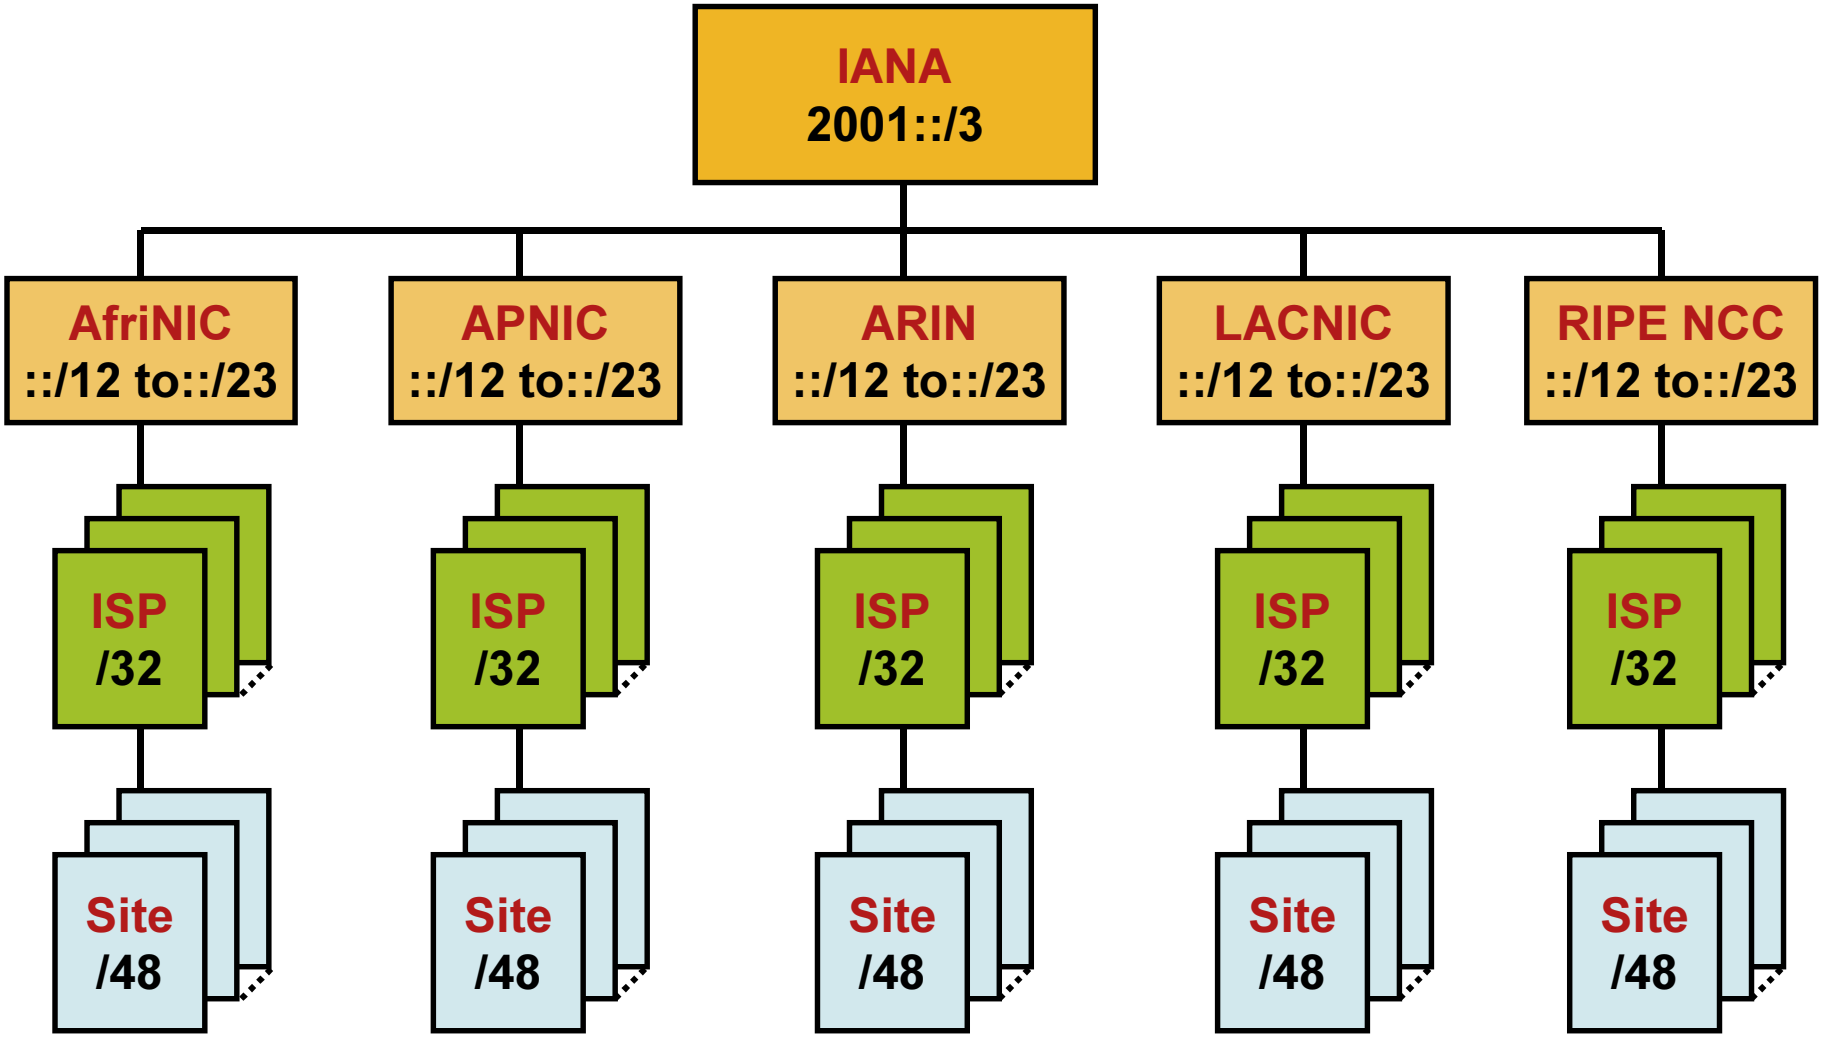

IPv6 Address Allocation Process

- 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)
  - Auto-generated pseudo-random number (to address privacy concerns)
  - Assigned via DHCP
  - Manually configured

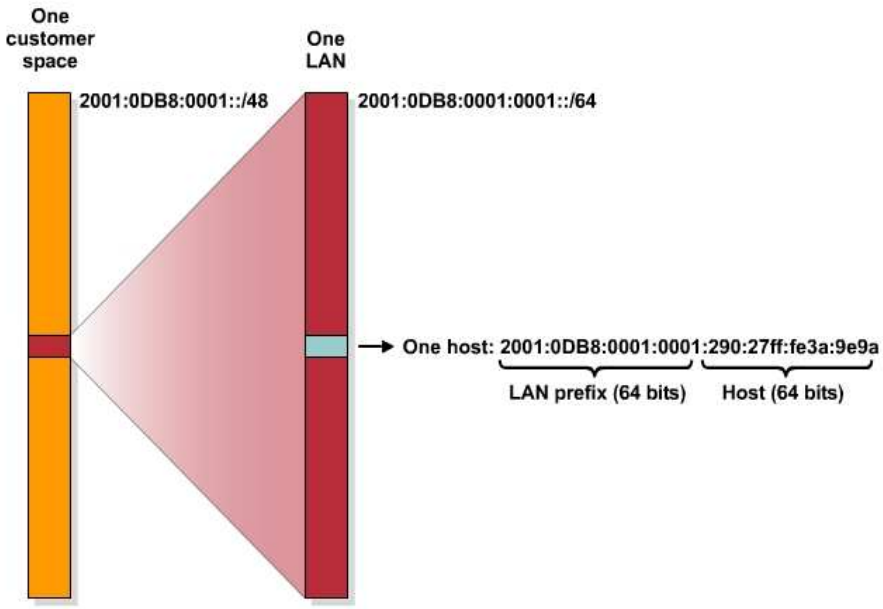

IPv6 Interface Identifier

- 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

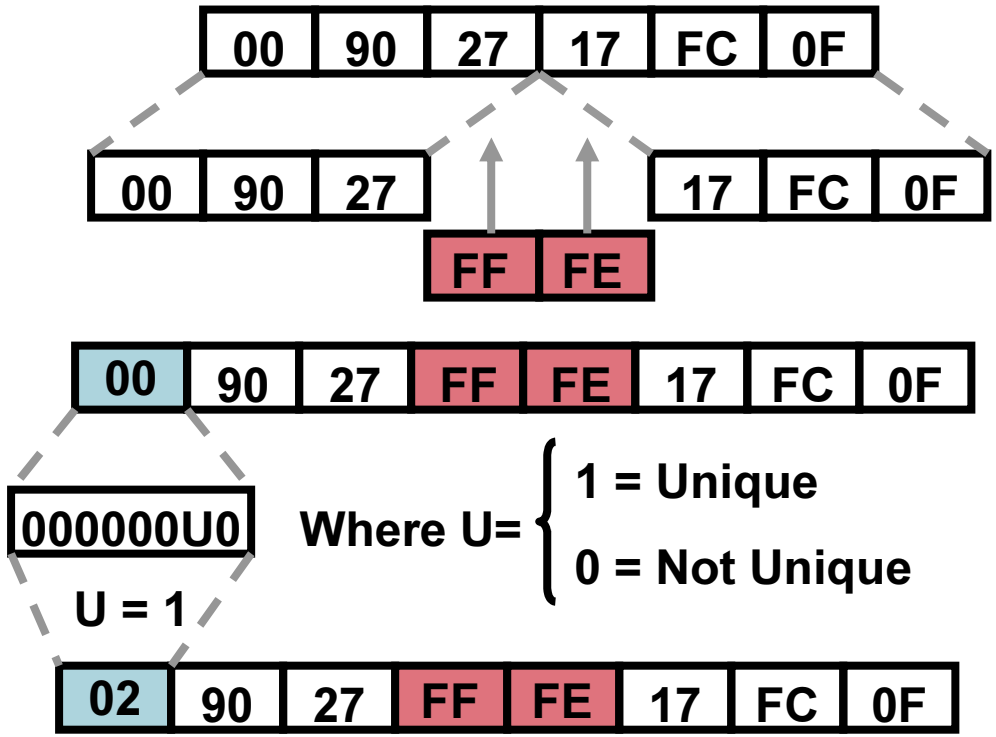

ICMPv6 and Neighbor Discovery

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

- 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

Duplicate Address Detection

```
ICMP type = 135
Src = 0   (::)
Dst = Solicited-node multicast of A
Data = link-layer address of A
Query = what is your link address?
```

- Duplicate Address Detection (DAD) uses neighbor solicitation to verify the existence of an address to be configured

Redirect

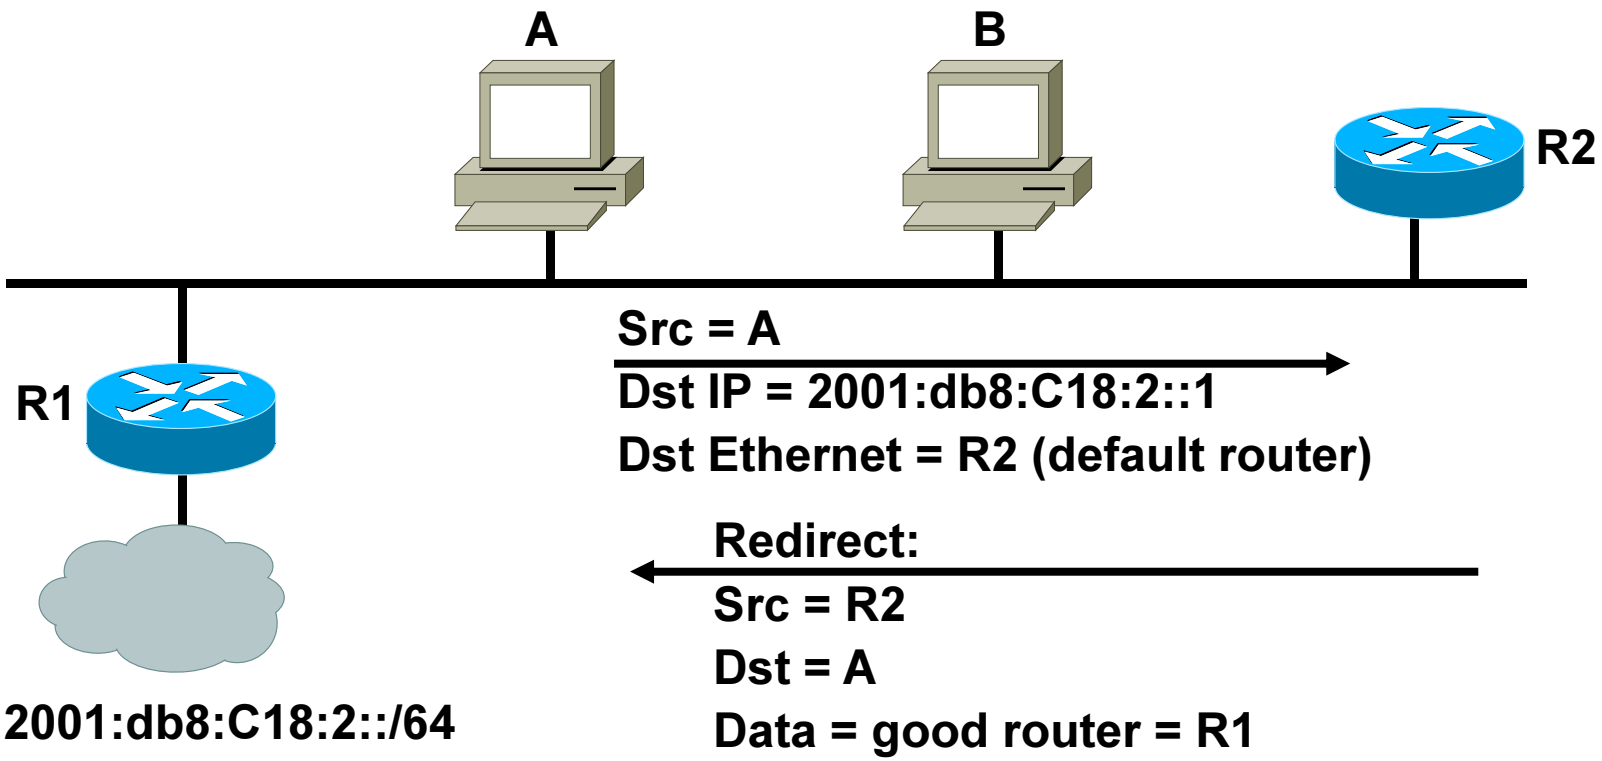

- Redirect is used by a router to signal the reroute of a packet to a better router

Autoconfiguration

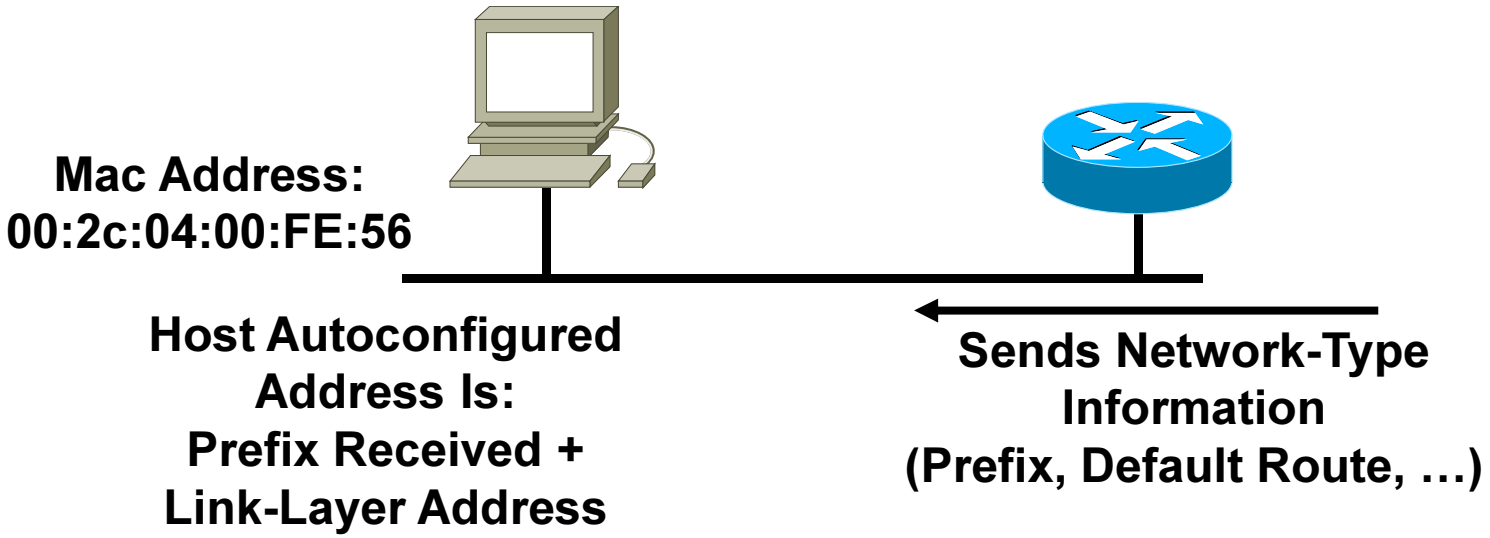

Larger Address Space Enables:

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

DHCP and DNS for IPv6

IPv6 and DNS

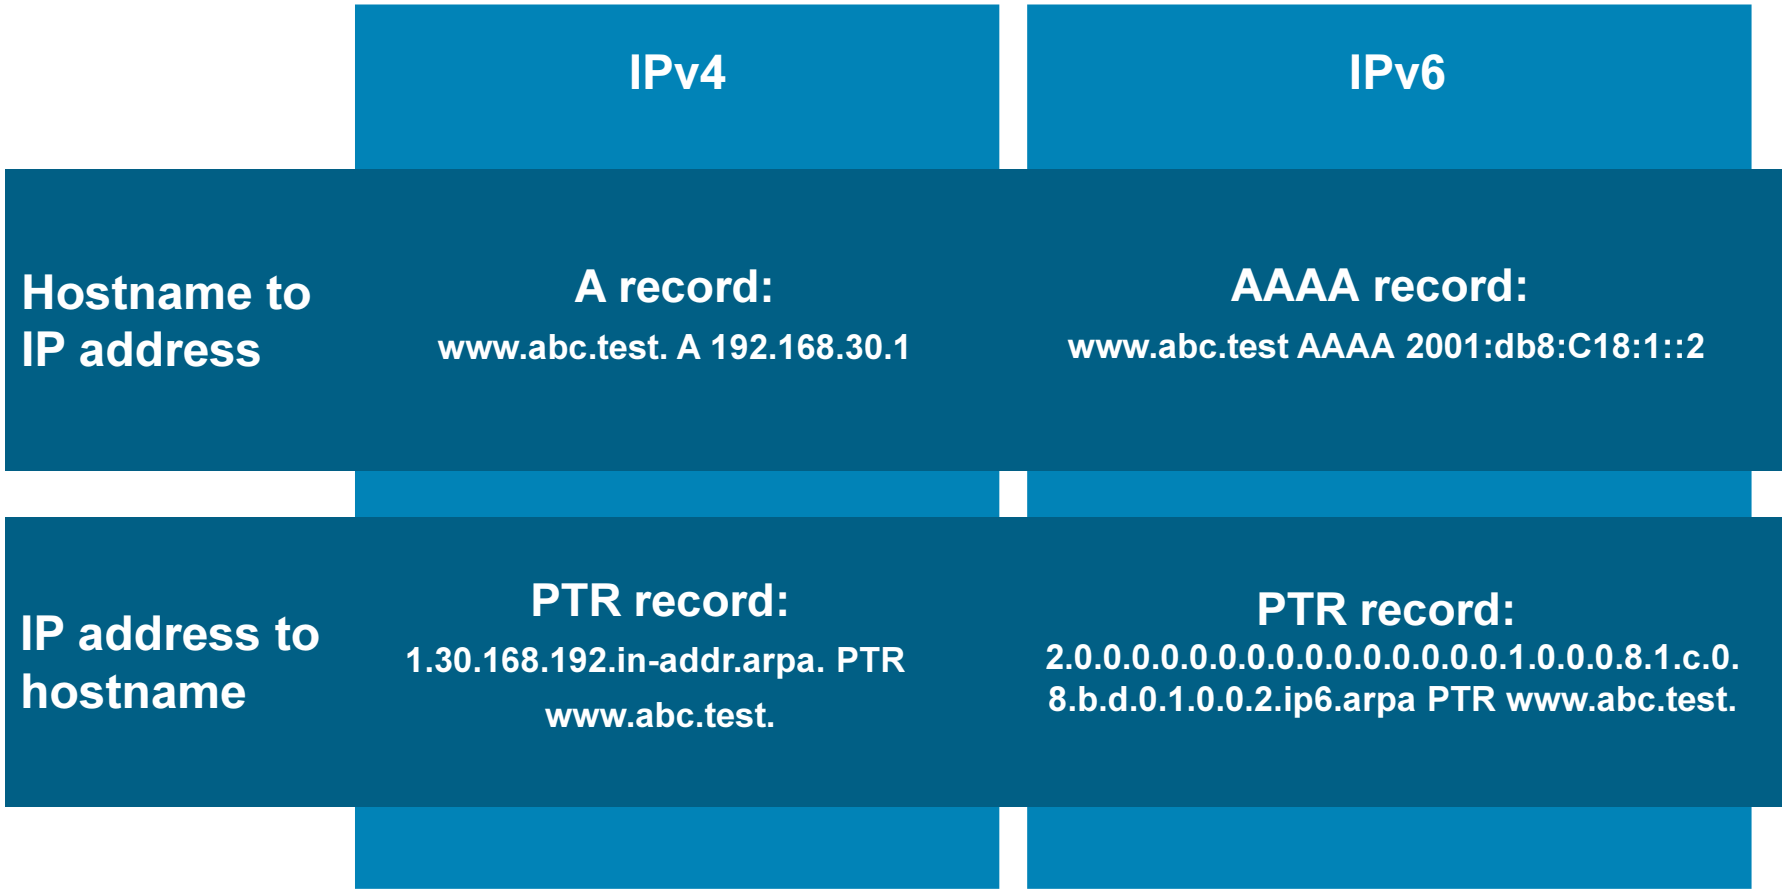

DHCPv6

- Updated version of DHCP for IPv4
- Supports new addressing
- Can be used for renumbering
- DHCP Process is same as in IPv4, but,
- Client first detect the presence of routers on the link
- If found, then examines router advertisements to determine if DHCP can be used
- If no router found or if DHCP can be used, then

  DHCP Solicit message is sent to the All-DHCP-Agents multicast address

  Using the link-local address as the source address

DHCPv6 Operation

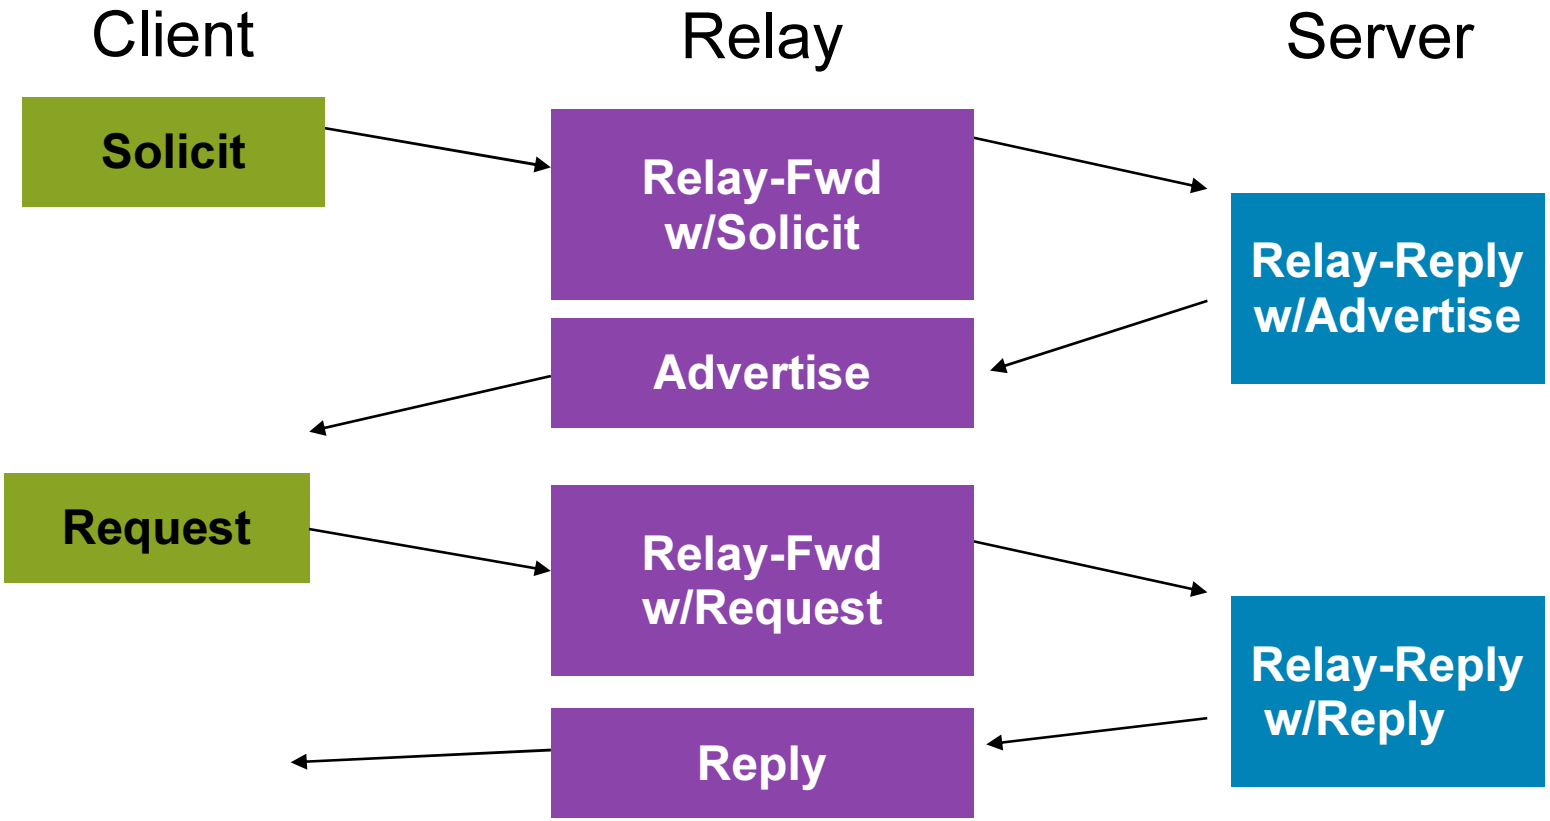

- All_DHCP_Relay_Agents_and_Servers (FF02::1:2)
- All_DHCP_Servers (FF05::1:3)
- DHCP Messages: Clients listen UDP port 546; servers and relay agents listen on UDP port 547

IPv6 Configurations

IOS IPv6 Addressing Examples (1)

Manual Interface Identifier

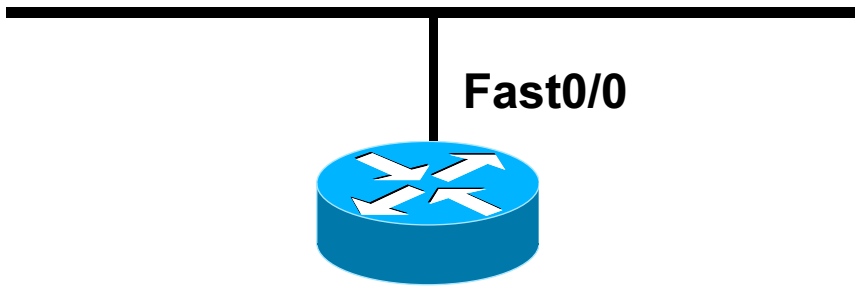


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 ip pim sparse-mode
 duplex auto
 speed auto
 ipv6 address 2006:1::1/64
 ipv6 enable
 ipv6 nd ra-interval 30
 ipv6 nd prefix 2006:1::/64 300 300
!
```

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

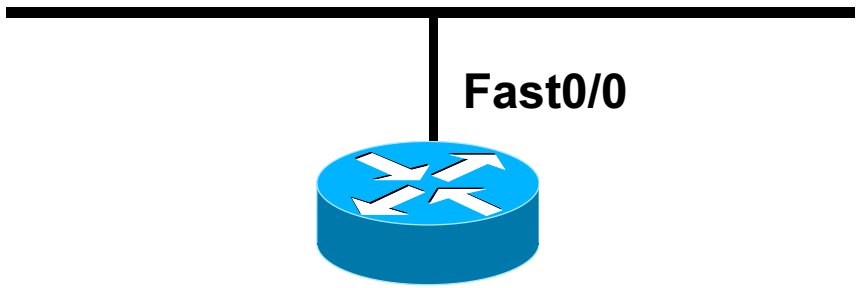


```
ipv6 unicast-routing
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 ip pim sparse-mode
 duplex auto
 speed auto
 ipv6 address 2006:1::/64 eui-64
 ipv6 enable
 ipv6 nd ra-interval 30
 ipv6 nd prefix 2006:1::/64 300 300
!
```

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)
```

IPv6 Routing

Static Routing

Static Routing

*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
```

RIPng (RFC 2080)

Enhanced Routing Protocol Support RIPng Overview RFC 2080

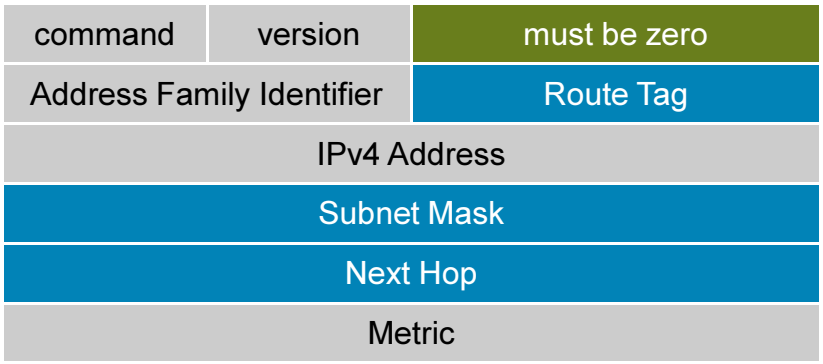

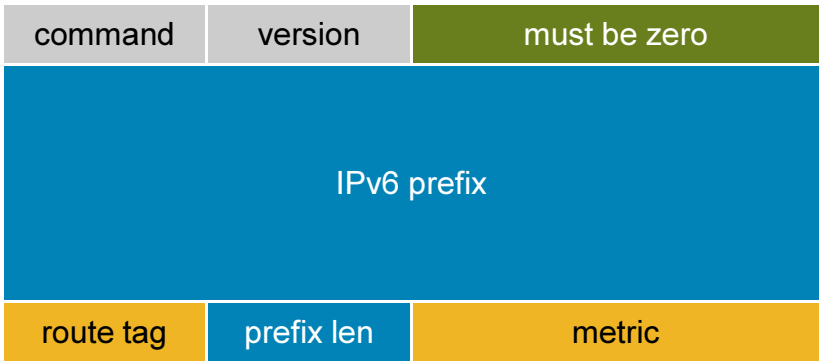

- Similar characteristics as IPv4

  Distance-vector, hop limit of 15, split-horizon, multicast based (FF02::9), UDP port (521) etc.

- Updated features for IPv6

  IPv6 prefix & prefix len

- Special Handling for the NH

  Route tag and prefix len for NH is all 0. Metric will have 0xFF; NH must be link local

Enhanced Routing Protocol Support
RIPng Configuration and Display

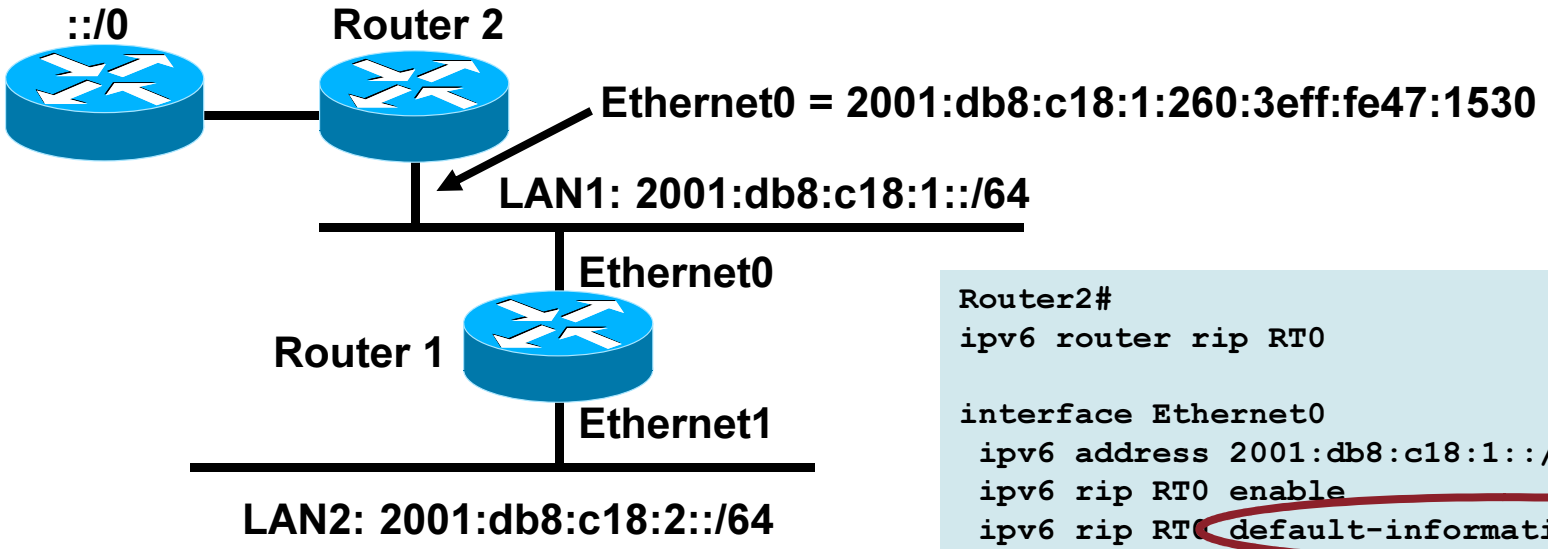


```
Router2#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 2001:db8:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 ipv6 rip RT0 default-information originate
```

```
Router1#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 2001:db8:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 Interface Ethernet1
 ipv6 address 2001:db8:c18:2::/64 eui-64
 ipv6 rip RT0 enable
```

```
Router2# debug ipv6 rip
RIPng: Sending multicast update on Ethernet0 for RT0
       src=FE80::260:3eff:fe47:1530
       dst=FF02::9 (Ethernet0)
       sport=521, dport=521, length=32
       command=2, version=1, mbz=0, #rte=1
       tag=0, metric=1, prefix=::/0
```

Multicast All RIP-Routers

Link-Local src Address

Access-List

Cisco IOS Standard Access Lists

When Used for Traffic Filtering, IPv6 Standard Access Control Lists (ACL) Offers the Following Functions:

- Can filter traffic based on source and destination address
- Can filter traffic inbound or outbound on a specific interface
- Can add priority to the ACL
- Implicit "deny all" at the end of access list

IPv6 Access-List Example

- Filtering outgoing traffic from site-local source addresses

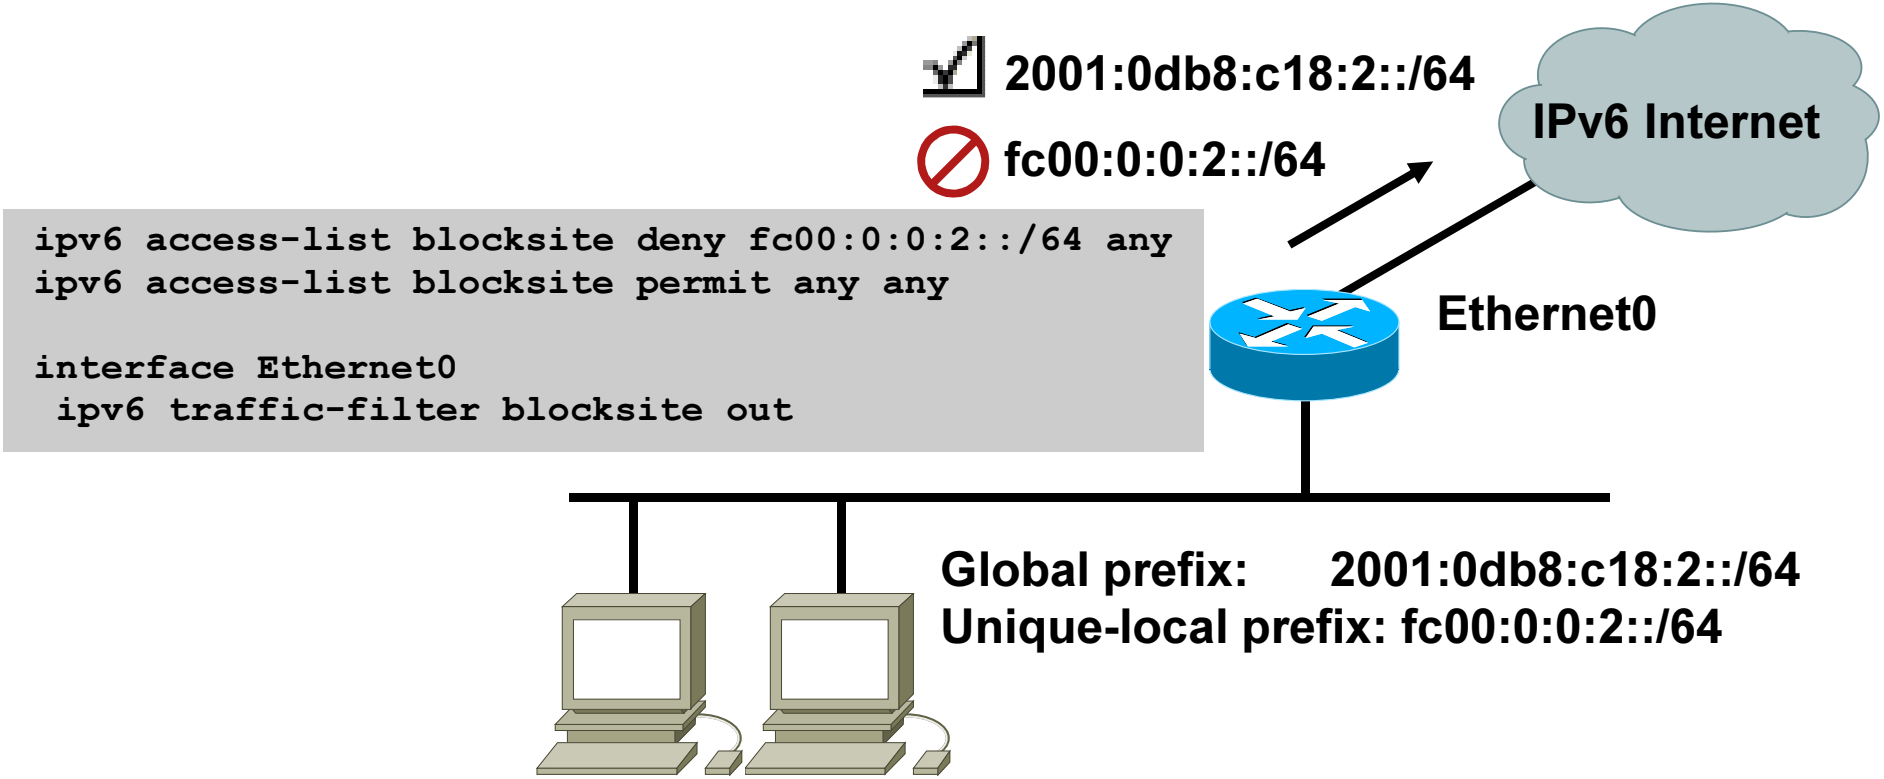

Deployment

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

Dual Stack Approach

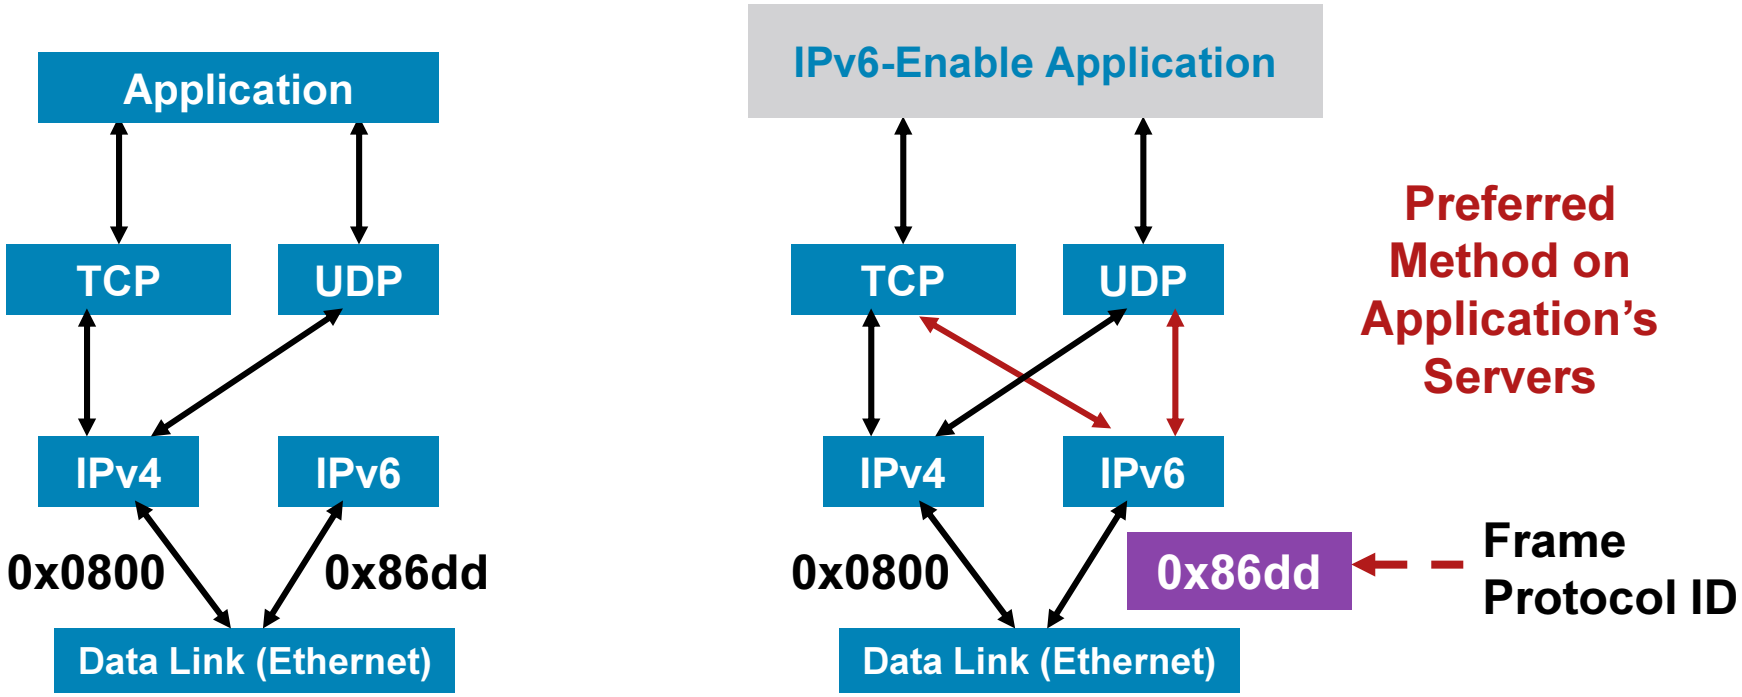

Dual Stack Node Means:

- Both IPv4 and IPv6 stacks enabled
- Applications can talk to both
- Choice of the IP version is based on name lookup and application preference

Cisco IOS Dual Stack Configuration

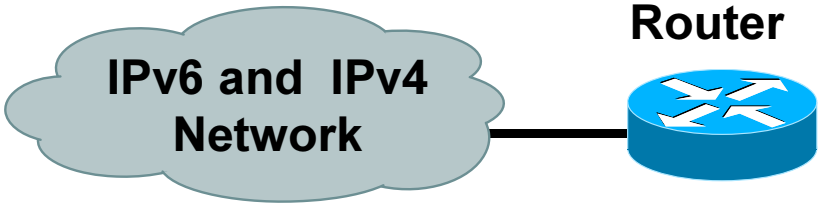


```
router#
ipv6 unicast-routing

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

IPv4: 192.168.99.1

IPv6: 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

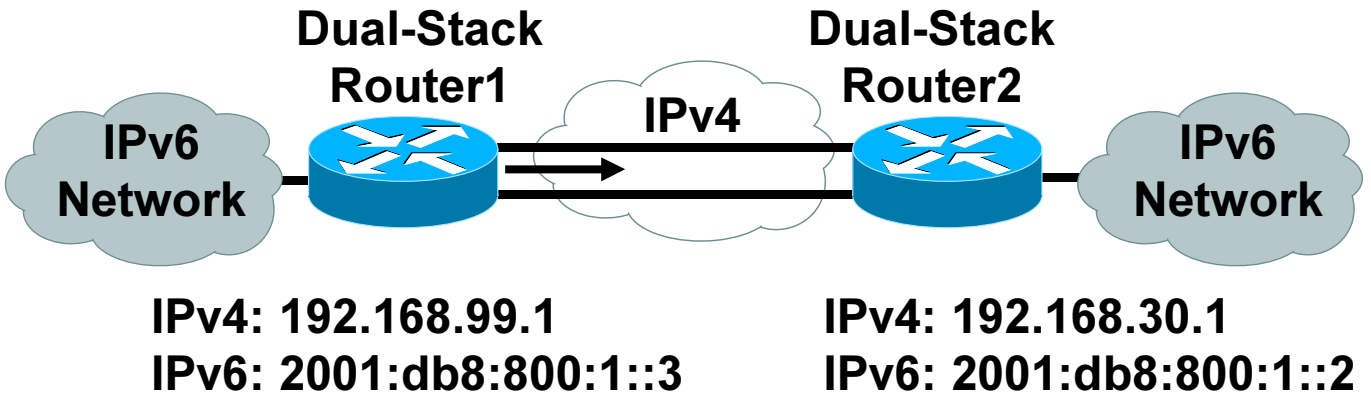


```
router1#

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

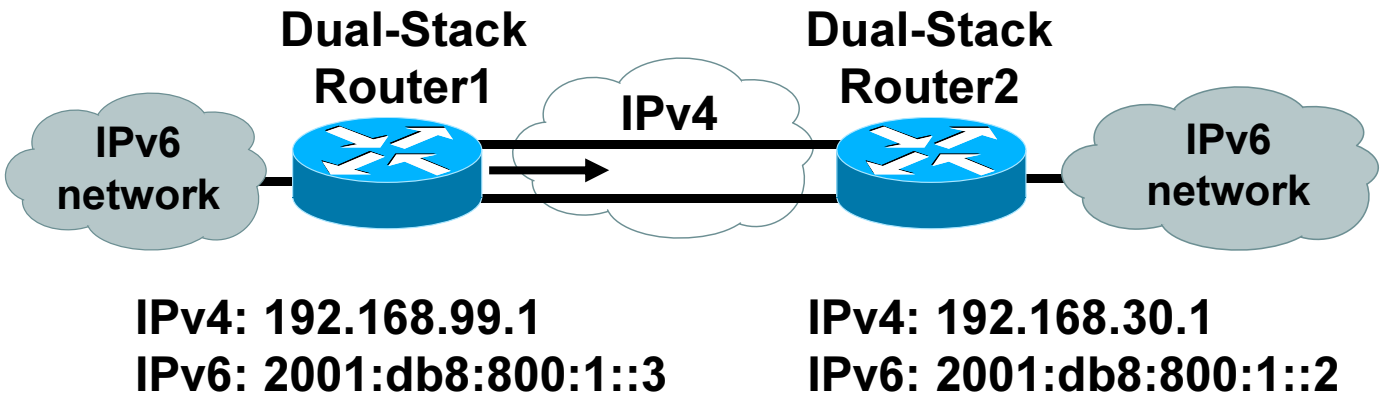


```
router1#

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
```

6to4 Tunneling

Automatic 6to4 Tunnels

- Automatic 6to4 tunnel allows isolated IPv6 domains to connect over an IPv4 network
- Unlike the manual 6to4 the tunnels are not point-to-point, they are multipoint tunnels
- IPv4 is embedded in the IPv6 address is used to find the other end of the tunnel
- Address format is 2002:IPv4 address::

Automatic 6to4 Tunnel (RFC 3056)

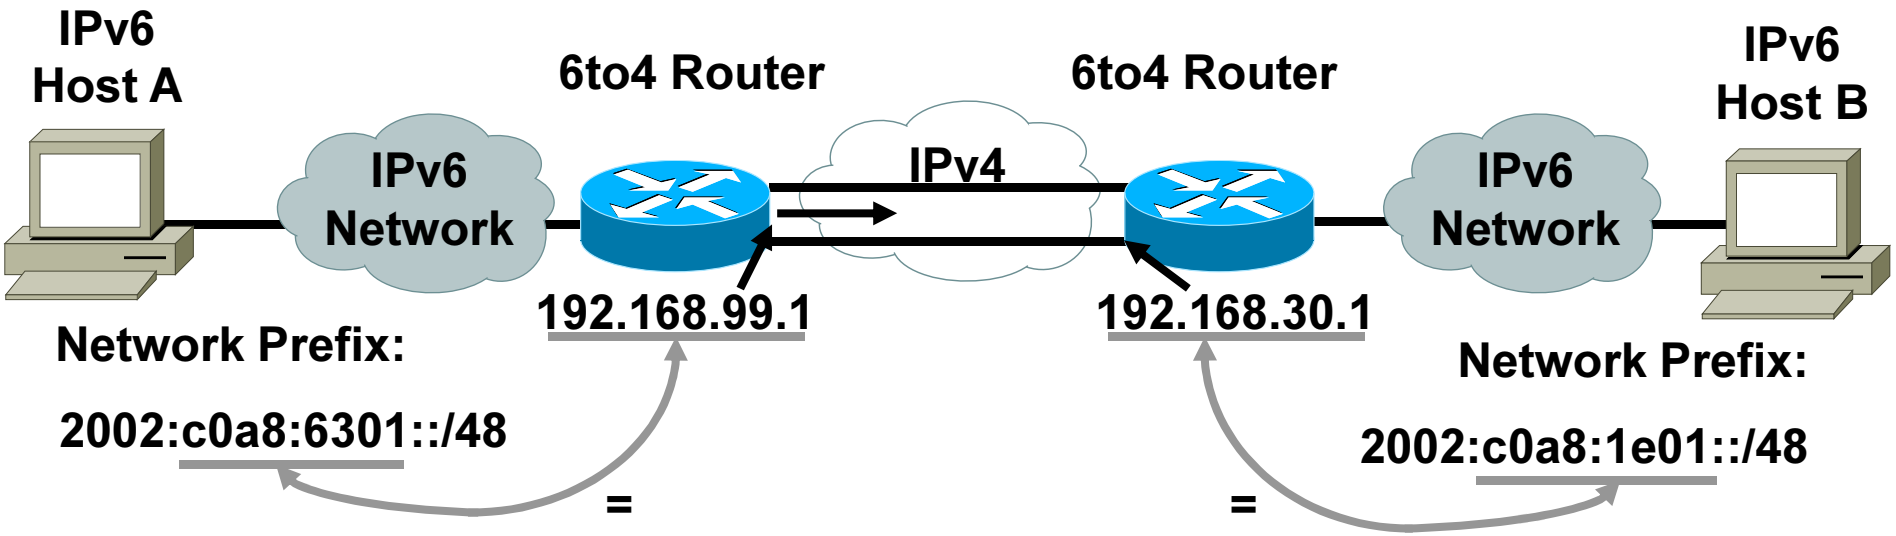

6to4:

- Is an automatic tunnel method
- Gives a prefix to the attached IPv6 network

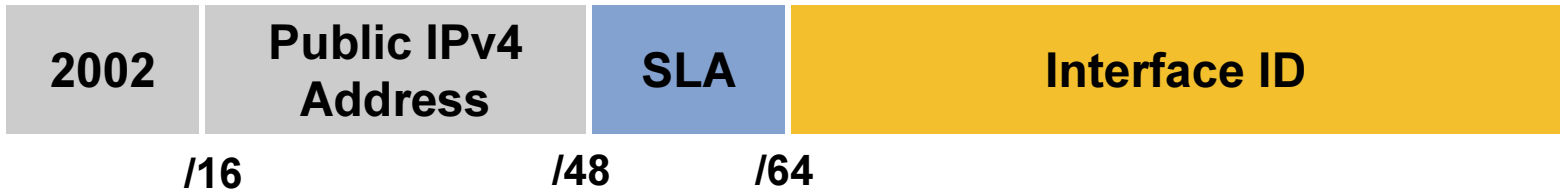

Automatic 6to4 Tunnel (RFC 3056)

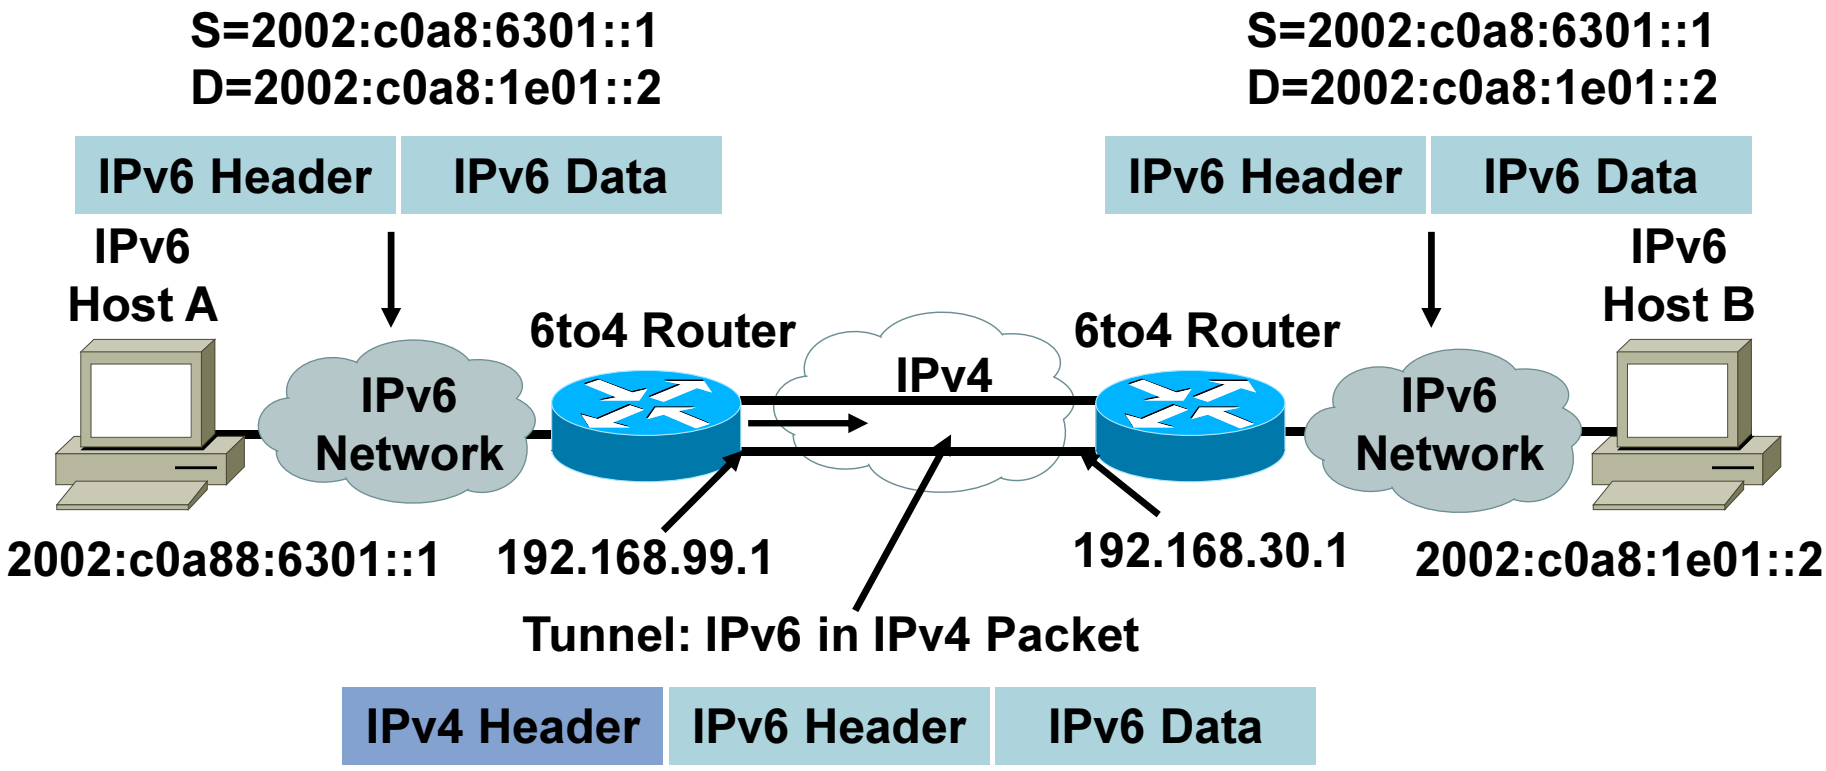

S(v4)=192.168.99.1
D(v4)=192.168.30.1
S(v6)=2002:c0a8:6301::1
D(v6)=2002:c0a8:1e01::2

Automatic 6to4 Configuration

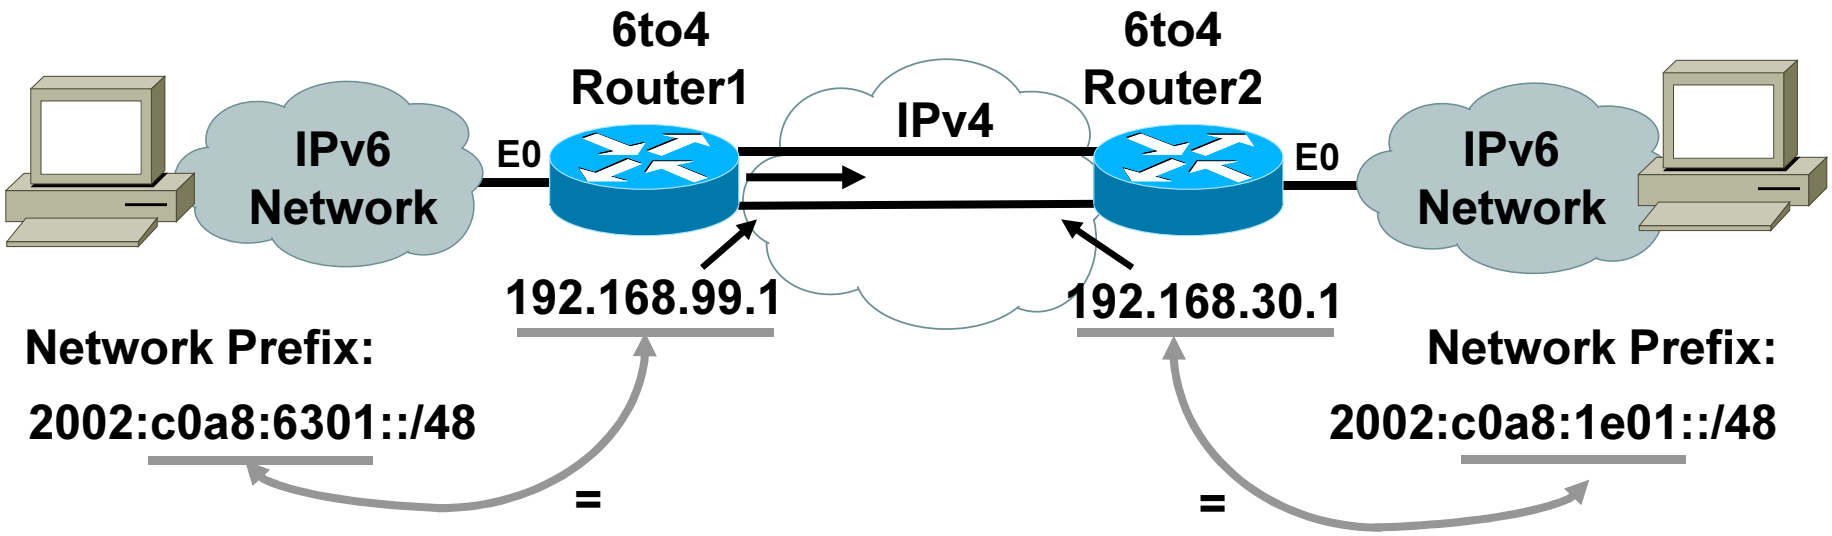


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
```

```
router2#
interface Ethernet0
 ipv6 address 2002:c0a8:1e01:1::/64 eui-64
Interface Ethernet1
 ip address 192.168.30.1 255.255.0.0
interface Tunnel0
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
```

Automatic 6to4 Relay

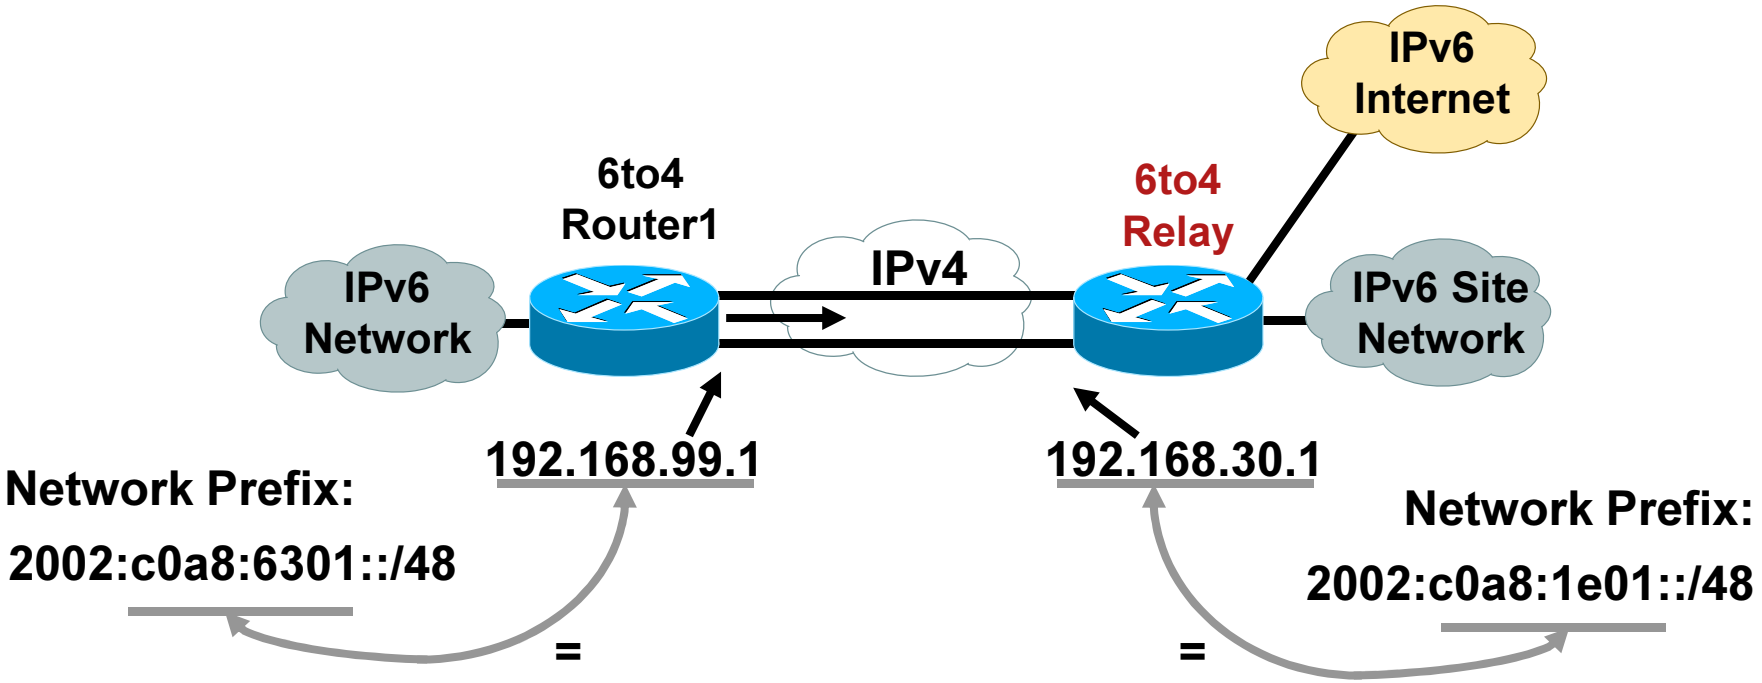

6to4 Relay:

- Is a gateway to the rest of the IPv6 Internet
- Is a default router

Automatic 6to4 Relay Configuration

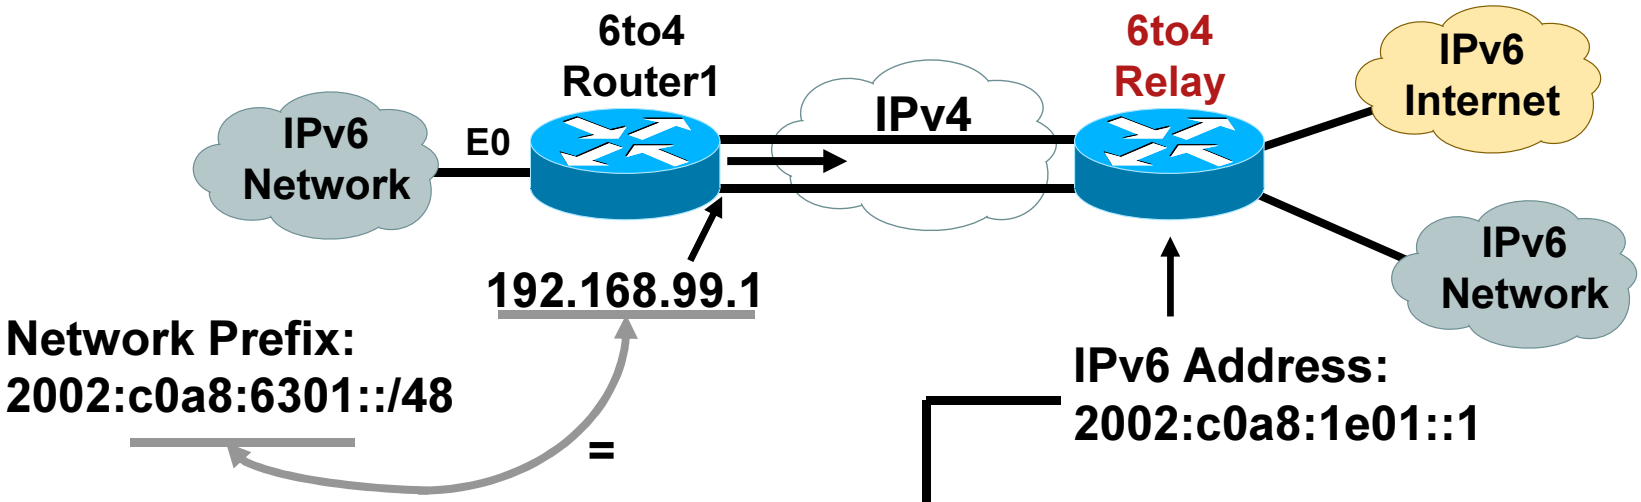


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 no ip address
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
ipv6 route ::/0 2002:c0a8:1e01::1
```

ISATAP Tunneling

Intrasite Automatic Tunnel Address Protocol

- RFC 4214
- To deploy a router is identified that carries ISATAP services
- ISATAP routers need to have at least one IPv4 interface and 0 or more IPv6 interface
- DNS entries are created for each of the ISATAP routers IPv4 addresses
- Hosts will automatically discover ISATAP routers and can get access to global IPv6 network
- Host can apply the ISATAP service before all this operation but its interface will only have a link local v6 address until the first router appears

Intrasite Automatic Tunnel Address Protocol

Use IANA's OUI 00-00-5E and
Encode IPv4 Address as Part of EUI-64

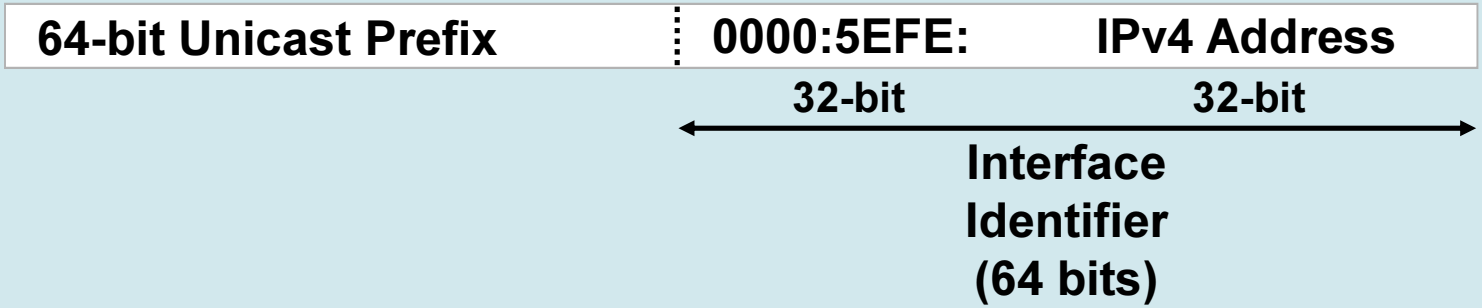

- ISATAP is used to tunnel IPv4 within as administrative domain (a site) to create a virtual IPv6 network over a IPv4 network
- Supported in Windows XP Pro SP1 and others

IPv6 Campus ISATAP Configuration

- Supported in Windows XP Pro SP1 and others
- ISATAP connections look like one flat network
- Create DNS “A” record for “ISATAP” = 10.1.1.1
- Use Static Config if DNS use is not desired:

```
C:\>netsh interface ipv6 isatap set
router 10.1.1.1
```

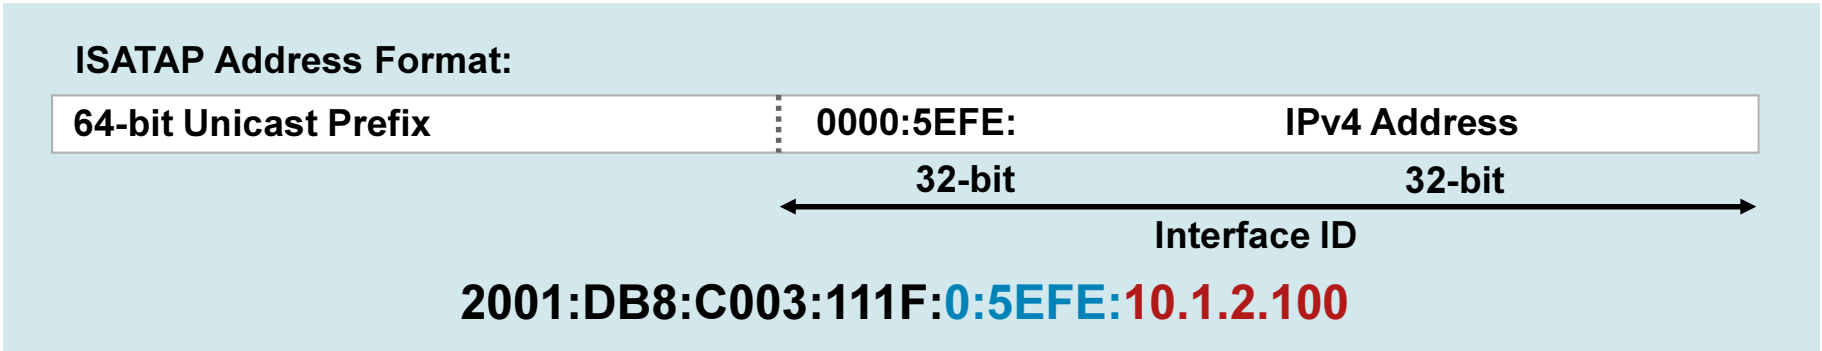

Client Configuration (Linux): ISATAP Tunnels

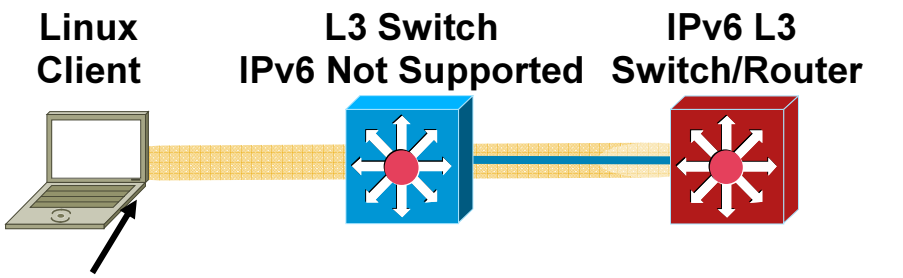

- IPv6-enabled
- Requires Kernel support for ISATAP—USAGI
- Modified IProute package—USAGI
- Must configure ISATAP router—not automatic

Host IP Router IP

```
# ip tunnel add is0 mode isatap 10.1.1.100 v4any 30.1.1.1 ttl 64
# ip link set is0 up
```

Automatic Advertisement of ISATAP Prefix

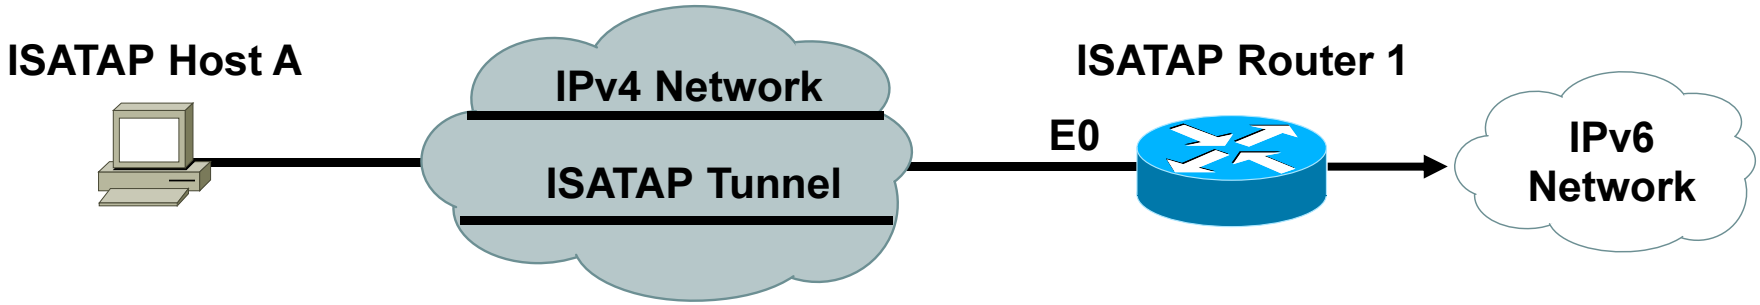

ICMPv6 Type 133 (RS)
IPv4 Source: 206.123.20.100
IPv4 Destination: 206.123.31.200
IPv6 Source: fe80::5efe:ce7b:1464
IPv6 Destination: fe80::5efe:ce7b:1fc8
Send me ISATAP Prefix

ICMPv6 Type 134 (RA)
IPv4 Source: 206.123.31.200
IPv4 Destination: 206.123.20.100
IPv6 Source: fe80::5efe:ce7b:1fc8
IPv6 Destination: fe80::5efe:ce7b:1464
ISATAP Prefix: 2001:db8:ffff :2::/64

Automatic Address Assignment of Host and Router

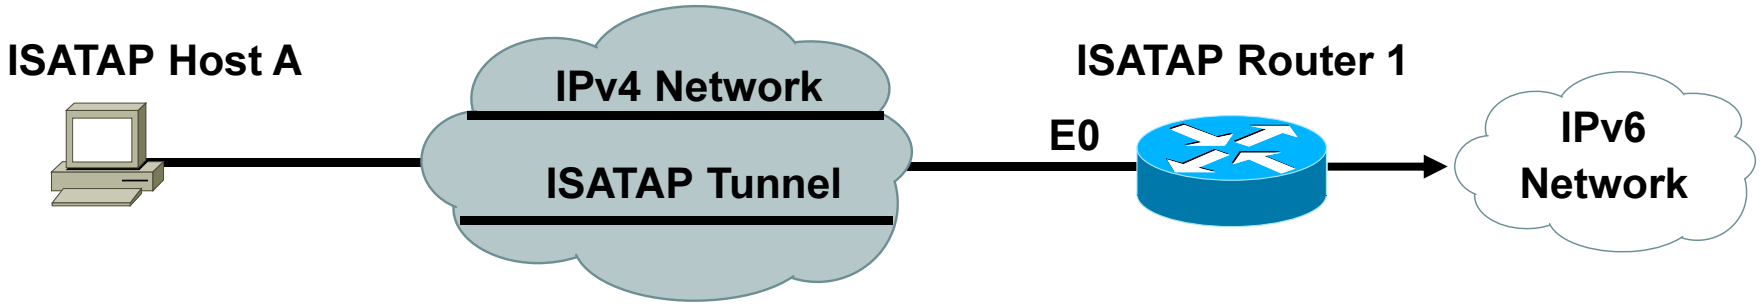

- ISATAP host A receives the ISATAP prefix 2001:db8:ffff:2::/64 from ISATAP Router 1
- When ISATAP host A wants to send IPv6 packets to 2001:db8:ffff:2::5efe:ce7b:1fc8, ISATAP host A encapsulates IPv6 packets in IPv4. The IPv4 packets of the IPv6 encapsulated packets use IPv4 source and destination address.

Automatic Configuring ISATAP

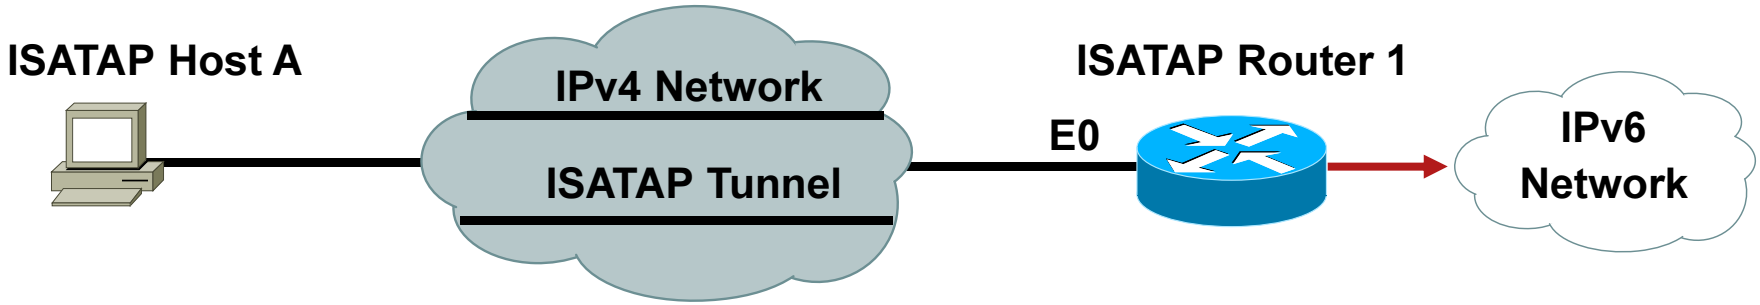


```
ISATAP-router1#
!
interface Ethernet0
 ip address 206.123.31.200 255.255.255.0
!
interface Tunnel0
 ipv6 address 2001:db8:ffff:2::/64 eui-64
 no ipv6 nd suppress-ra
 tunnel source Ethernet0
 tunnel mode ipv6ip isatap
```

- The tunnel source command must point to an interface with an IPv4 address configured
- Configure the ISATAP IPv6 address, and prefixes to be advertised just as you would with a native IPv6 interface
- The IPv6 address has to be configured as an EUI-64 address since the last 32 bits in the interface identifier is used as the IPv4 destination address

Conclusion

- IPv6 is real!
- Start now rather than later
  - Purchase for the future
  - Start moving legacy application towards IPv6 support
  - Test, test and then test some more!
- Integration can be done per Application (Dual Stack or Tunneled)
- Microsoft Vista and Longhorn have IPv6 enabled by default and preferred over IPv4
- Enterprise and SP Deployment Scenarios:
  - ISP IPv6 Deployment Scenarios in Broadband Access Networks (RFC 4779)
  - Scenarios and Analysis for Introducing IPv6 into ISP Networks (RFC 4029)
  - IPv6 Enterprise Network Scenarios (RFC 4057)
  - Procedures for Renumbering an IPv6 Network without a Flag Day (RFC 4192)

Complete The Evaluation

- Win FLIP Video; give us your evaluation forms.
- Winner will be announced at the end of each day