IPv6 Tutorial

SANOG8
Karachi, Pakistan
August 2nd, 2006

Khalid Raza, Cisco Distinguished Engineer
Salman Asadullah, Technical Leader
Cisco Systems

Agenda

- IPv6 Merits, Motivation & Planning
- IPv6 Addressing, Headers & Basics
- IPv6 Addressing Planning & Assignments
- IPv6 & DNS
- IPv6 Network Management
- IPv6 Routing Protocols
- Enterprise Deployment
  - Campus
  - WAN
  - S2S VPN
  - Remote Access
- Service Provider Deployment
  - Core
  - Access
- IPv6 Services
  - Multicast
  - QoS
  - Security
  - Mobility

IPv6 Merits, Motivation & Planning

Drivers for IPv6

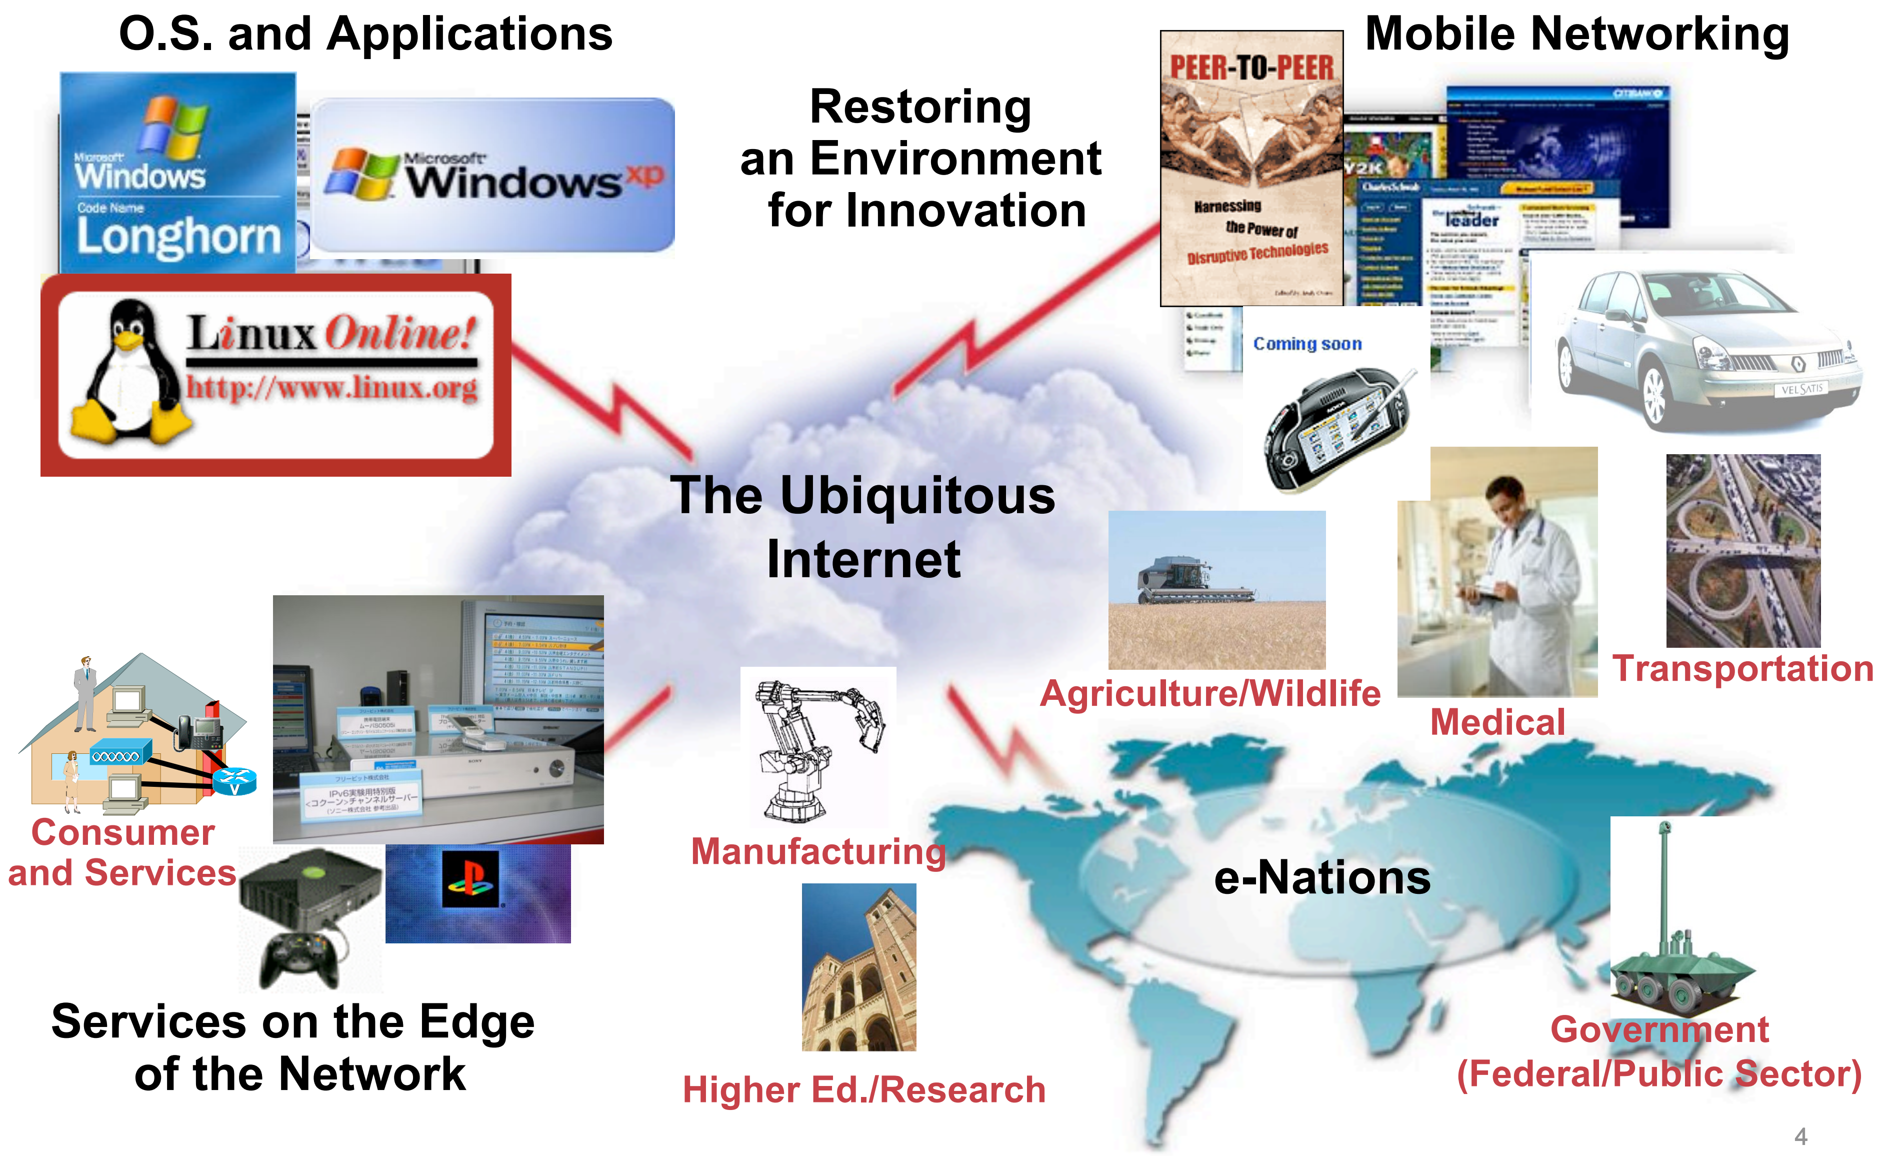

IPv6 Vertical Activity

Higher Ed./Research

- Media services
- Collaboration
- Mobility

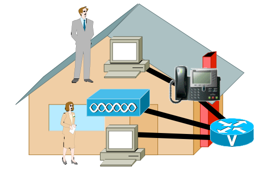

Consumer

- Set-top boxes
- Gaming
- Appliances
- Voice/video
- Security monitoring

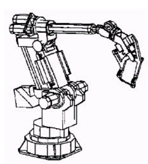

Manufacturing

- Embedded devices
- Industrial Ethernet
- IP-enabled components

Government (Federal/Public Sector)

- DoD
- WIN-T
- FCS
- JTRS
- GIG-BE

Transportation

- Telematics
- Traffic control
- Hotspots
- Transit services

Agriculture/Wildlife

- Animal tags
- Imagery
- Botanical
- Weather

Medical

- Home care
- Imaging
- Mobility

IPv6 for the Military

- Soldiers
- Weapons
- Sensors
- Command/control
- Logistics

- Massive address space (billions)
- Mobile IP
- Security/encryption
- Simplified management
- Inter-service interoperability

FCS (Future Combat Systems)

WIN-T (Warfighter Information Network—Tactical)

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**
  - **BUT**
- **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!**

A Need for IPv6?

- Internet Population
  - ~600M users in Q4 CY2002, ~945M by end CY 2004 – only 10-15%of the total population
  - How to address the future Worldwide population? (~9B in CY 2050)
  - Emerging Internet countries need address space, eg: China uses nearly 2 class A (11/2002), ~20 class A needed if every student (320M) has to get an IP address
- Mobile Internet introduces new generation of Internet devices
  - PDA (~20M in 2004), Mobile Phones (~1.5B in 2003), Tablet PC
  - Enable through several technologies, eg: 3G, 802.11,…
- Transportation – Mobile Networks
  - 1B automobiles forecast for 2008 – Begin now on vertical markets
  - Internet access on planes, eg. Lufthansa – train, eg. Narita express
- Consumer, Home and Industrial Appliances

IP Address Allocation History

1981 - IPv4 protocol published

1985 ~ 1/16 of total space

1990 ~ 1/8 of total space

1995 ~ 1/3 of total space

2000 ~ 1/2 of total space

2002.5 ~ 2/3 of total space

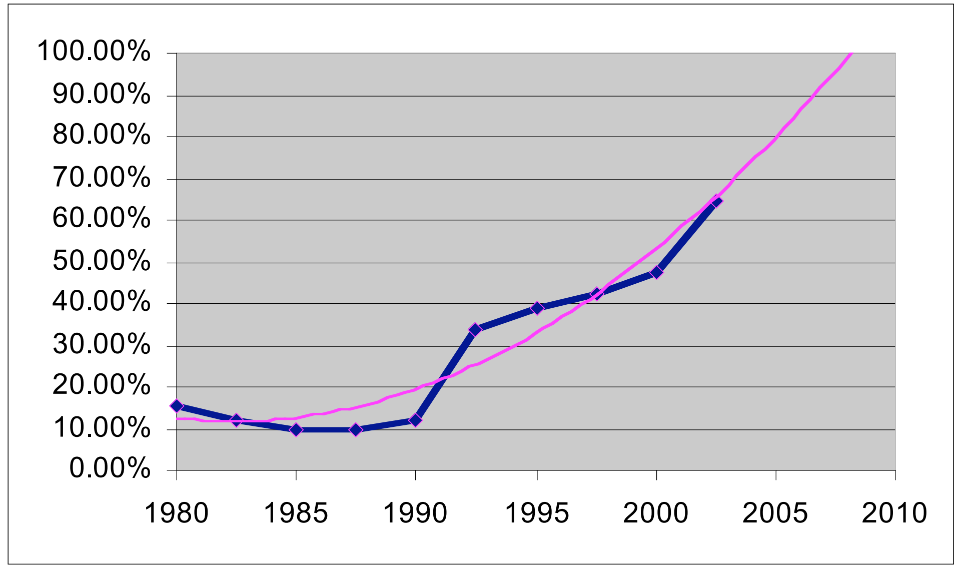

- **This despite increasingly intense conservation efforts**
  - **PPP / DHCP address sharing** **NAT (network address translation)**
  - **CIDR (classless inter-domain routing)** **plus some address reclamation**
- **Theoretical limit of 32-bit space: ~4 billion devices**
  Practical limit of 32-bit space: ~250 million devices (RFC 3194)

IPv6: Addressing Customer Problems

- NAT overlap
  - Acquisitions and mergers with overlapping private addressing (address space collisions)
  - How to access resources without massive renumbering
- Address constraints prohibit new services
  - Large customers with address shortages (private and public space)
  - Route table runaway
  - Makes managing existing services difficult
- New service and application requirements
  - Key for both enterprise and service providers trying to launch new services to users and subscribers
  - Facing large increase in IP-enabled devices (NAT kills peer-to-peer applications)
  - Scalable peer-to-peer communications (SIPv6, VoIPv6, etc…)
  - Ability to finally use multicast to its full potential

NAT Overlap

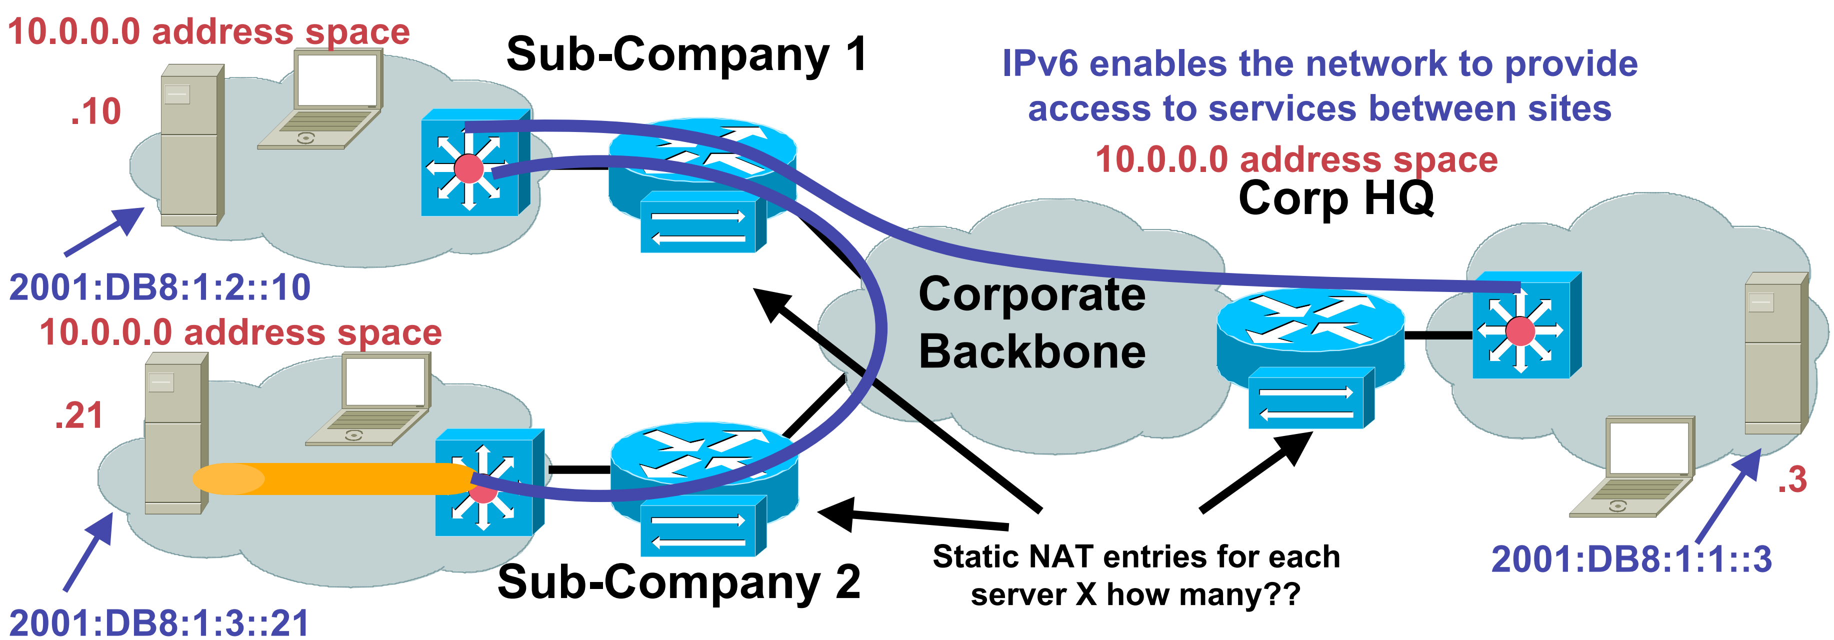

- **Merger and acquisition complexity force many to leave existing IPv4 address space in place vs. full integration/consolidation**
- **When server-to-server or client-to-server service is required then single/double static NAT translations are often required**
- **IPv6 can be deployed to enable service access per site and/or per application**

New Applications and Services

Today, all O.S. Are Dual-Stack

As soon as the infrastructure is IPv6 capable…IPv6 integration can follow a non-disruptive "per application" model

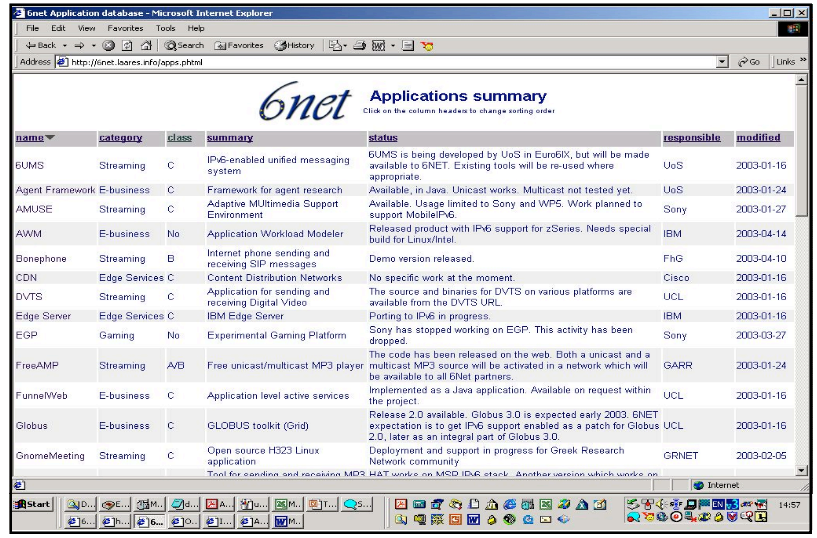

name	category	class	summary	status	responsible	modified
6UMS	Streaming	C	IPv6-enabled unified messaging system	6UMS is being developed by UoS in Euro6IX, but will be made available to 6NET. Existing tools will be re-used where appropriate.	UoS	2003-01-16
Agent Framework	E-business	C	Framework for agent research	Available, in Java. Unicast works. Multicast not tested yet.	UoS	2003-01-24
AMUSE	Streaming	C	Adaptive MUltimedia Support Environment	Available. Usage limited to Sony and WP5. Work planned to support MobileIPv6.	Sony	2003-01-27
AWM	E-business	No	Application Workload Modeler	Released product with IPv6 support for zSeries. Needs special build for Linux/Intel.	IBM	2003-04-14
Bonephone	Streaming	B	Internet phone sending and receiving SIP messages	Demo version released.	FhG	2003-04-10
CDN	Edge Services	C	Content Distribution Networks	No specific work at the moment.	Cisco	2003-01-16
DVTS	Streaming	C	Application for sending and receiving Digital Video	The source and binaries for DVTS on various platforms are available from the DVTS URL.	UCL	2003-01-16
Edge Server	Edge Services	C	IBM Edge Server	Porting to IPv6 in progress.	IBM	2003-01-16
EGP	Gaming	No	Experimental Gaming Platform	Sony has stopped working on EGP. This activity has been dropped.	Sony	2003-03-27
FreeAMP	Streaming	A/B	Free unicast/multicast MP3 player	The code has been released on the web. Both a unicast and a multicast MP3 source will be activated in a network which will be available to all 6Net partners.	GARR	2003-01-24
FunnelWeb	E-business	C	Application level active services	Implemented as a Java application. Available on request within the project.	UCL	2003-01-16
Globus	E-business	C	GLOBUS toolkit (Grid)	Release 2.0 available. Globus 3.0 is expected early 2003. 6NET expectation is to get IPv6 support enabled as a patch for Globus 2.0, later as an integral part of Globus 3.0.	UCL	2003-01-16
GnomeMeeting	Streaming	C	Open source H323 Linux application	Deployment and support in progress for Greek Research Network community	GRNET	2003-02-05

Call for Applications—Protocol Agnostic

New Generation of Internet Appliances

How Do we Get There from Here?

- **IPv4 & IPv6 will coexist for the foreseeable future**
  - **No D-Day / Flag Day.**
- **Education & Careful Planning are crucial.**
  - **How long does it take in your environment?**
- **IPv4 & IPv6 implementations must be scalable, reliable, secure and feature rich.**

Strategy that reflects this …

Starting with Edge upgrades enable IPv6 service offerings now

IPv6 Integration

- **Many ways to deliver IPv6 services to End Users, Most important is End to End IPv6 traffic forwarding**
- **Service Providers and Enterprises may have different deployment needs and mechanisms but basic steps are common**
  - **Definition of an IPv6 addressing scheme**
  - **Selection of the IPv6 routing protocol(s)**
  - **DNS server ready to register AAAA record**
  - **IPv6 devices management rules**
  - **Security rules for IPv6 access**

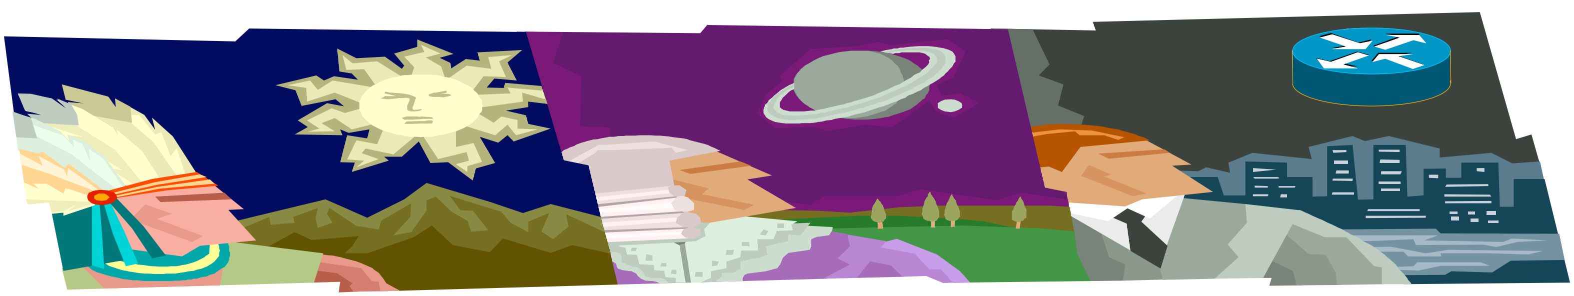

Transition & Integration Richness

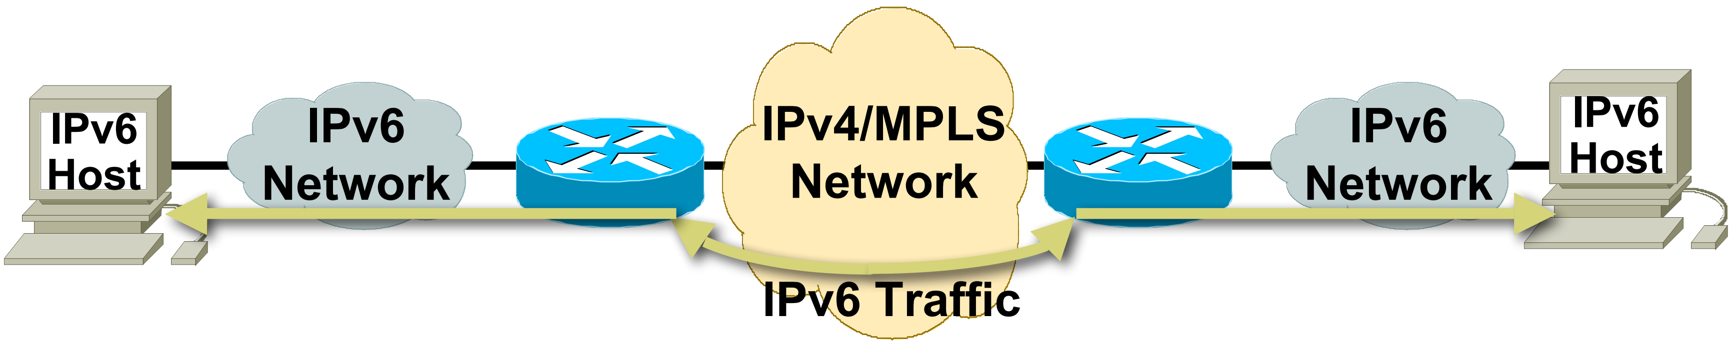

- Transition richness means:
  - No fixed day to convert
  - No need to convert all at once
  - Different transition mechanisms are available
  - Smooth integration of IPv4 and IPv6
  - Different compatibility mechanisms
  - IPv4 and IPv6 nodes can talk to each other

IPv4-IPv6 Transition / Co-Existence

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

IPv6 Addressing, Header & Basics

Expanded Address Space
IPv6 Addressing

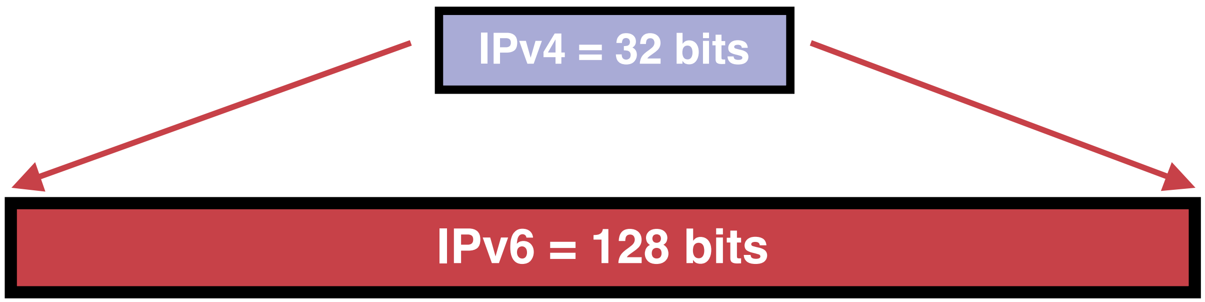

For IPv4, We have 32 bits

= ~ 4,200,000,000 possible addressable nodes.

For IPv6

Some wanted fixed-length, 64-bit addresses

Some wanted variable-length, up to 160 bits

Settled for 128 bits

=340,282,366,920,938,463,463,374,607,431,768,211,456 nodes

Expanded Address Space
IPv6 Addressing

- **IPv6 Address Format**

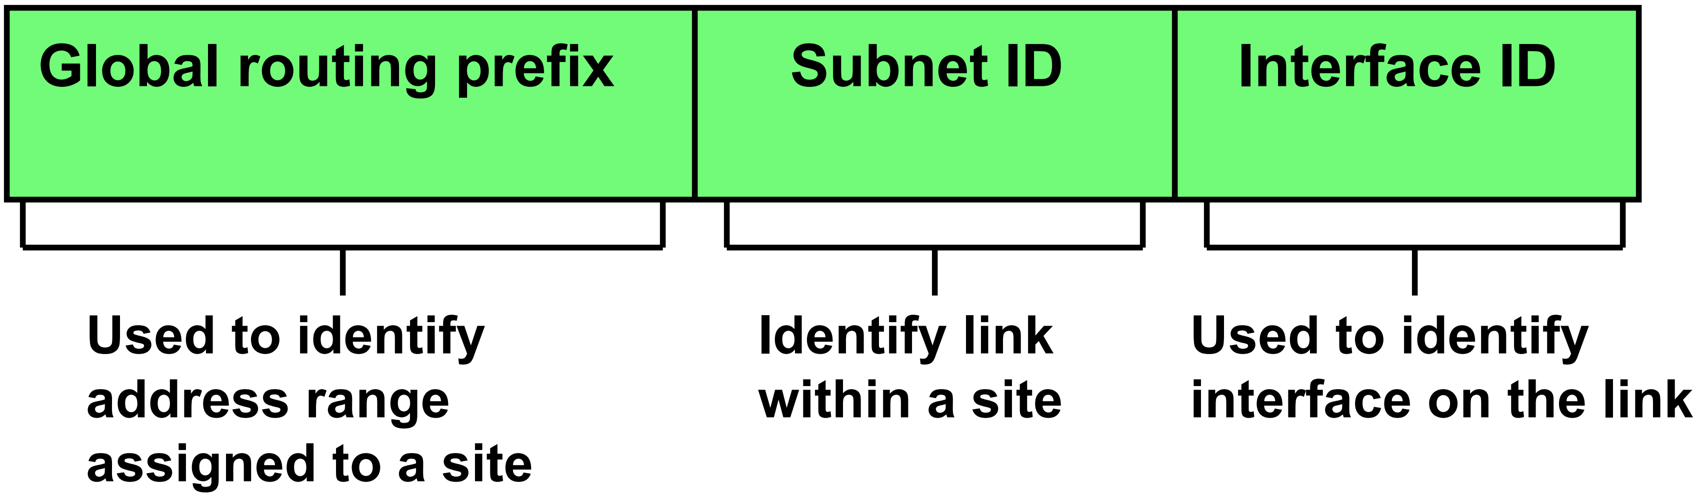

Expanded Address Space
IPv6 Addressing

- **Representation**
  - **16 bit hexadecimal numbers**
  - **Numbers are separated by (:)**
  - **Hex numbers are not case sensitive**
  - **Example:**
  - **2003:0000:130F:0000:0000:087C:876B:140B**

Expanded Address Space
IPv6 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 3ef8:ca62:12FE::/48

Addressing

- **Prefix representation**

Hex	Binary	Number of bits
3ef	001111101111	16
CA6	110010100110	16
12	0001 0010	8

How to get an IPv6 Address?

- **How to get address space?**

  Real IPv6 address space now allocated by APNIC, ARIN and RIPE NCC (Registries) to ISP

APNIC	2001:0200::/23 & 2001:0C00::/23
ARIN	2001:0400::/23
RIPE NCC	2001:0600::/23 - 2001:0B00::/23

- **IXCs** 2001:0700::/23
- **6Bone** 3FFE::/16
- **6to4 tunnels** 2002::/16
- **Mostly, Enterprises get their IPv6 address space from their ISP ☺**

Expanded Address Space
IPv6 Address Representation

- **16-bit fields in case insensitive colon hexadecimal representation**

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

- **Leading zeros in a field are optional:**

  2031:0:130F:0:0:9C0:876A:130B

- **Successive fields of 0 represented as ::, but only once in an address:**

  2031:0:130F::9C0:876A:130B

  2031::130F::9C0:876A:130B

- **IPv4-compatible address representation**

  0:0:0:0:0:0:192.168.30.1 = ::192.168.30.1 = ::C0A8:1E01

- **Loopback address representation**

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

- **Unspecified address representation**

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

Expanded Address Space
IPv6 Addressing

- **IPv6 addressing rules are covered by multiples RFC's**
  - **Architecture defined by RFC 3513**
- **Address types are:**
  - **Unicast: one to one (global, link local, site local, compatible)**
  - **Anycast: one to nearest (allocated from unicast)**
  - **Multicast: one to many**
  - **Reserved**
- **A single interface may be assigned multiple IPv6 addresses of any type (unicast, anycast, multicast)**
  - **No broadcast address - > use multicast**

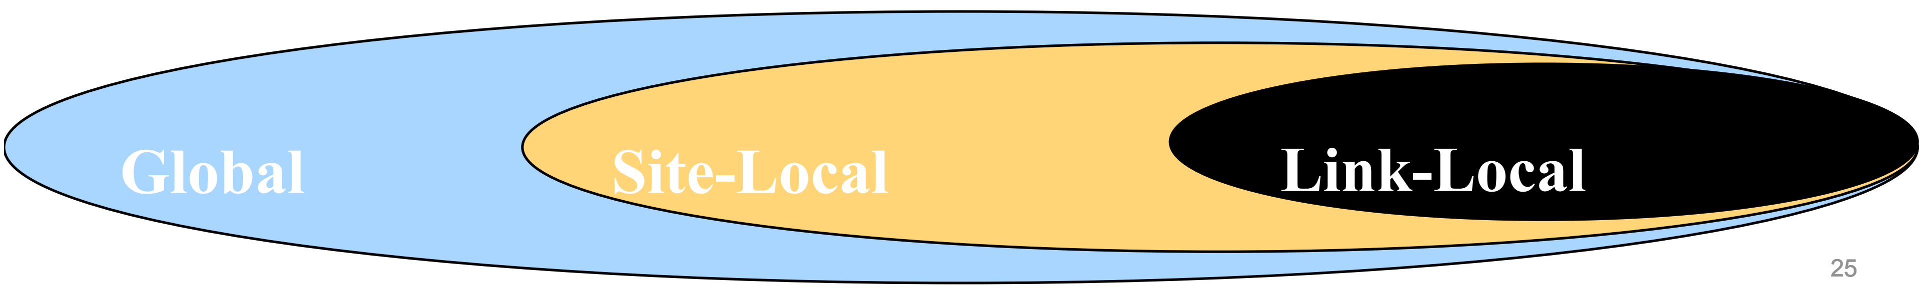

Expanded Address Space
IPv6 Address Range Reserved or Assigned

Of the Full Address Space

- 2000::/3 (001) is for aggregatable global unicast address
- FE80::/10 (1111 1110 10) is for link-local
- FEC0::/10 (1111 1110 11) is for site-local
- FF00::/8 (1111 1111) is for multicast
- ::/8 is reserved for the "unspecified address"
- Other values are currently unassigned (approx. 7/8th of total)

Expanded Address Space
Aggregatable Global Unicast Addresses

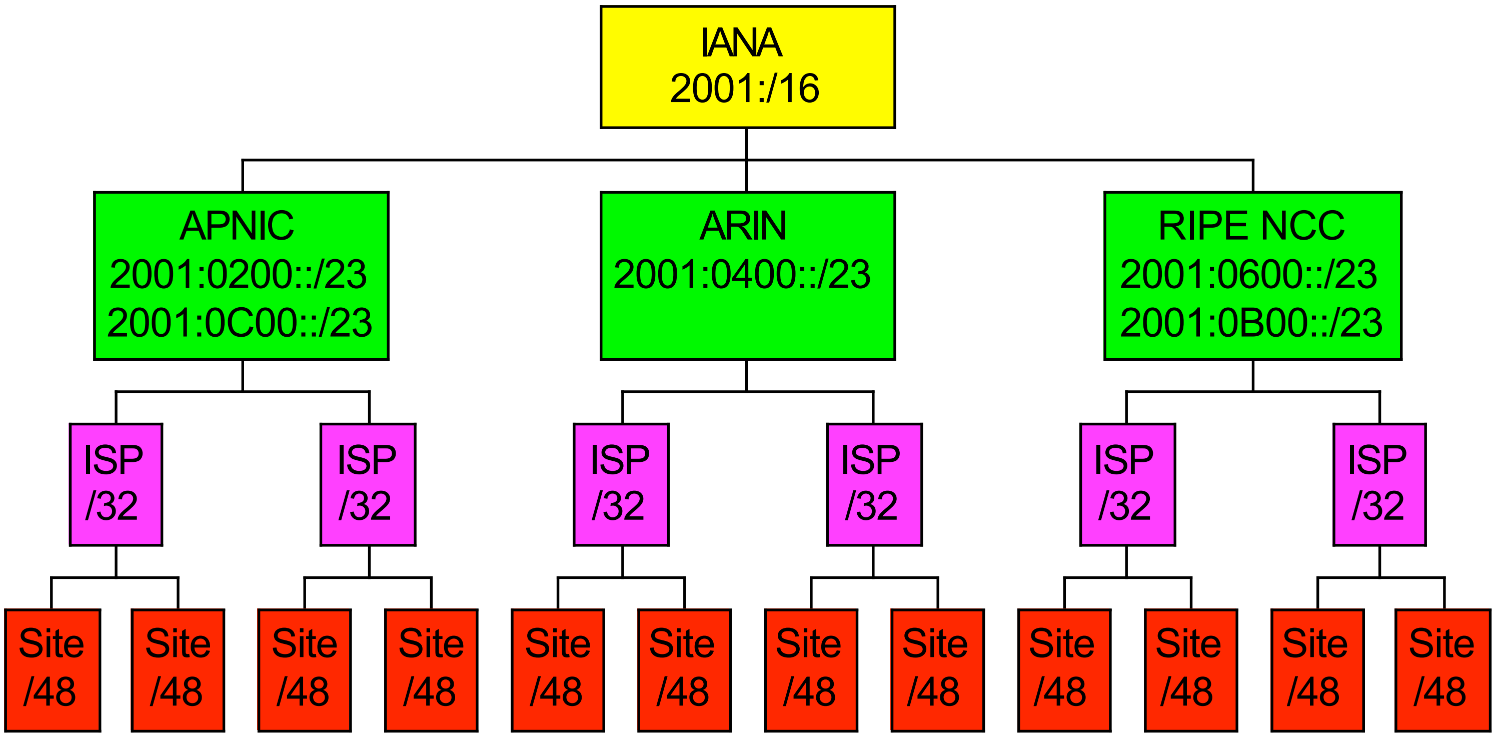

Expanded Address Space
Aggregatable Global Unicast Addresses

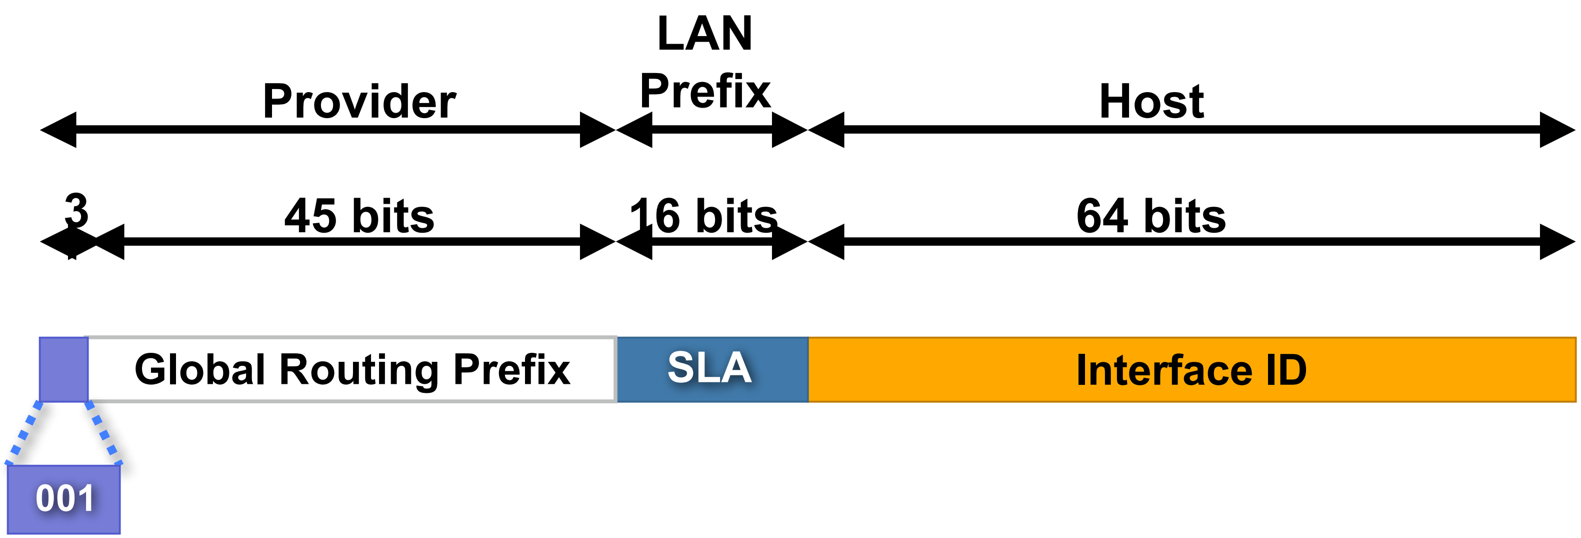

- Aggregatable global unicast addresses are:
  - Addresses for generic use of IPv6
  - Structured as a hierarchy to keep the aggregation
- See draft-ietf-ipngwg-addr-arch-v3-07

Expanded Address Space
Address Allocation

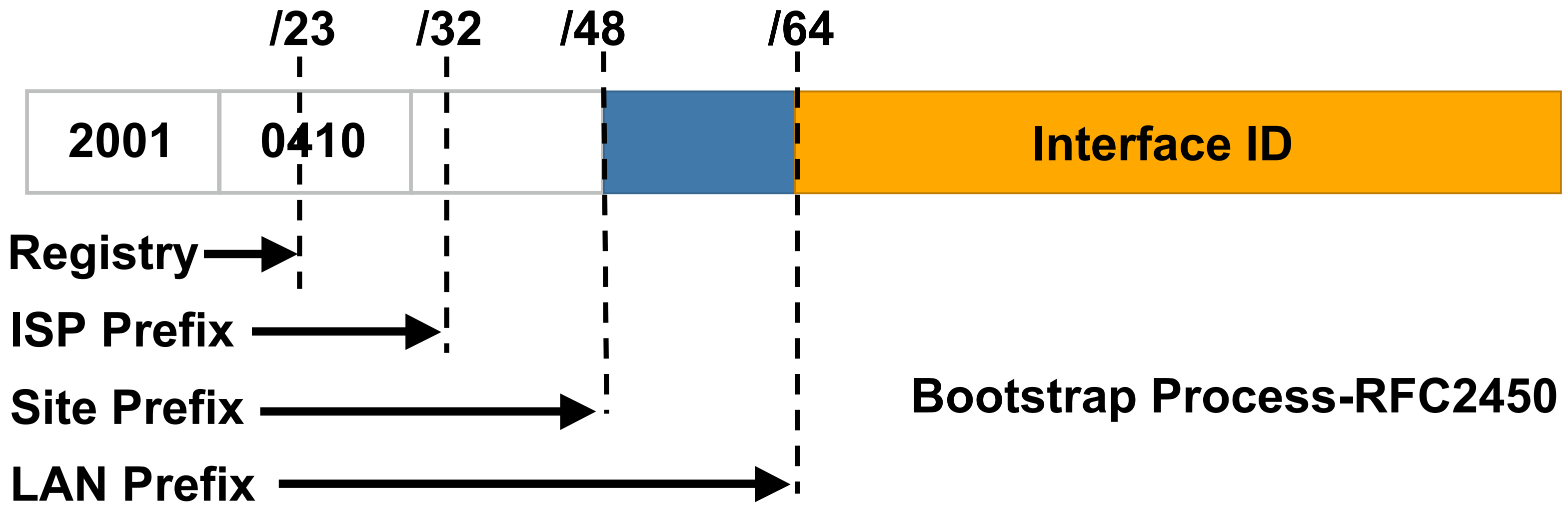

Bootstrap Process-RFC2450

- **The allocation process is under reviewed by the registries:**
  - Each registry gets a /23 prefix from IANA
  - Up to now, all ISP were getting a /32
  - With the new proposal, registry allocates a /36 (immediate allocation) or /32 (initial allocation) prefix to an IPv6 ISP
  - Policy is that an ISP allocates a /48 prefix to each end customer
  - IPv6 address allocation and assignment global policy
  - ftp://ftp.cs.duke.edu/pub/narten/ietf/global-ipv6-assign-2002-04-25.txt

Expanded Address Space
Hierarchical Addressing & Aggregation

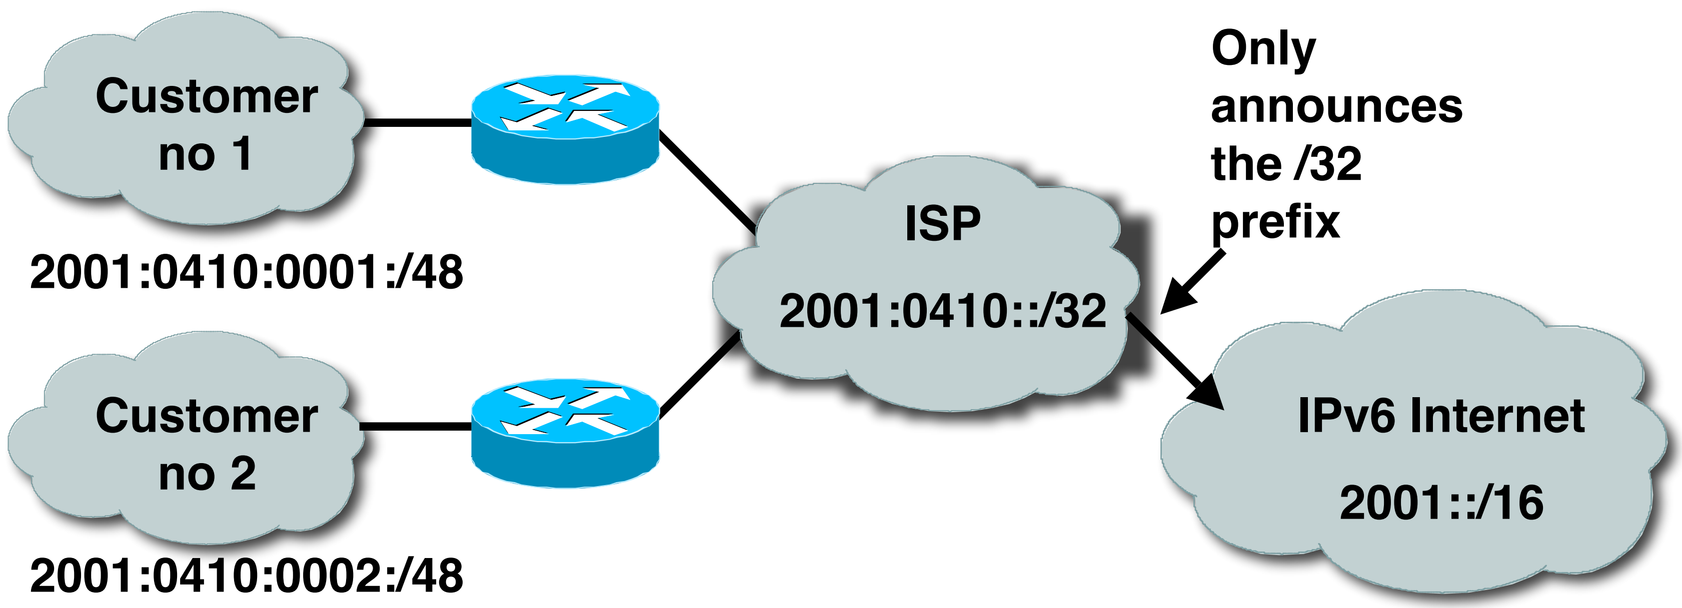

Larger address space enables:

- **Aggregation of prefixes announced in the global routing table.**
- **Efficient and scalable routing.**

Expanded Address Space
Hierarchical Addressing & Aggregation

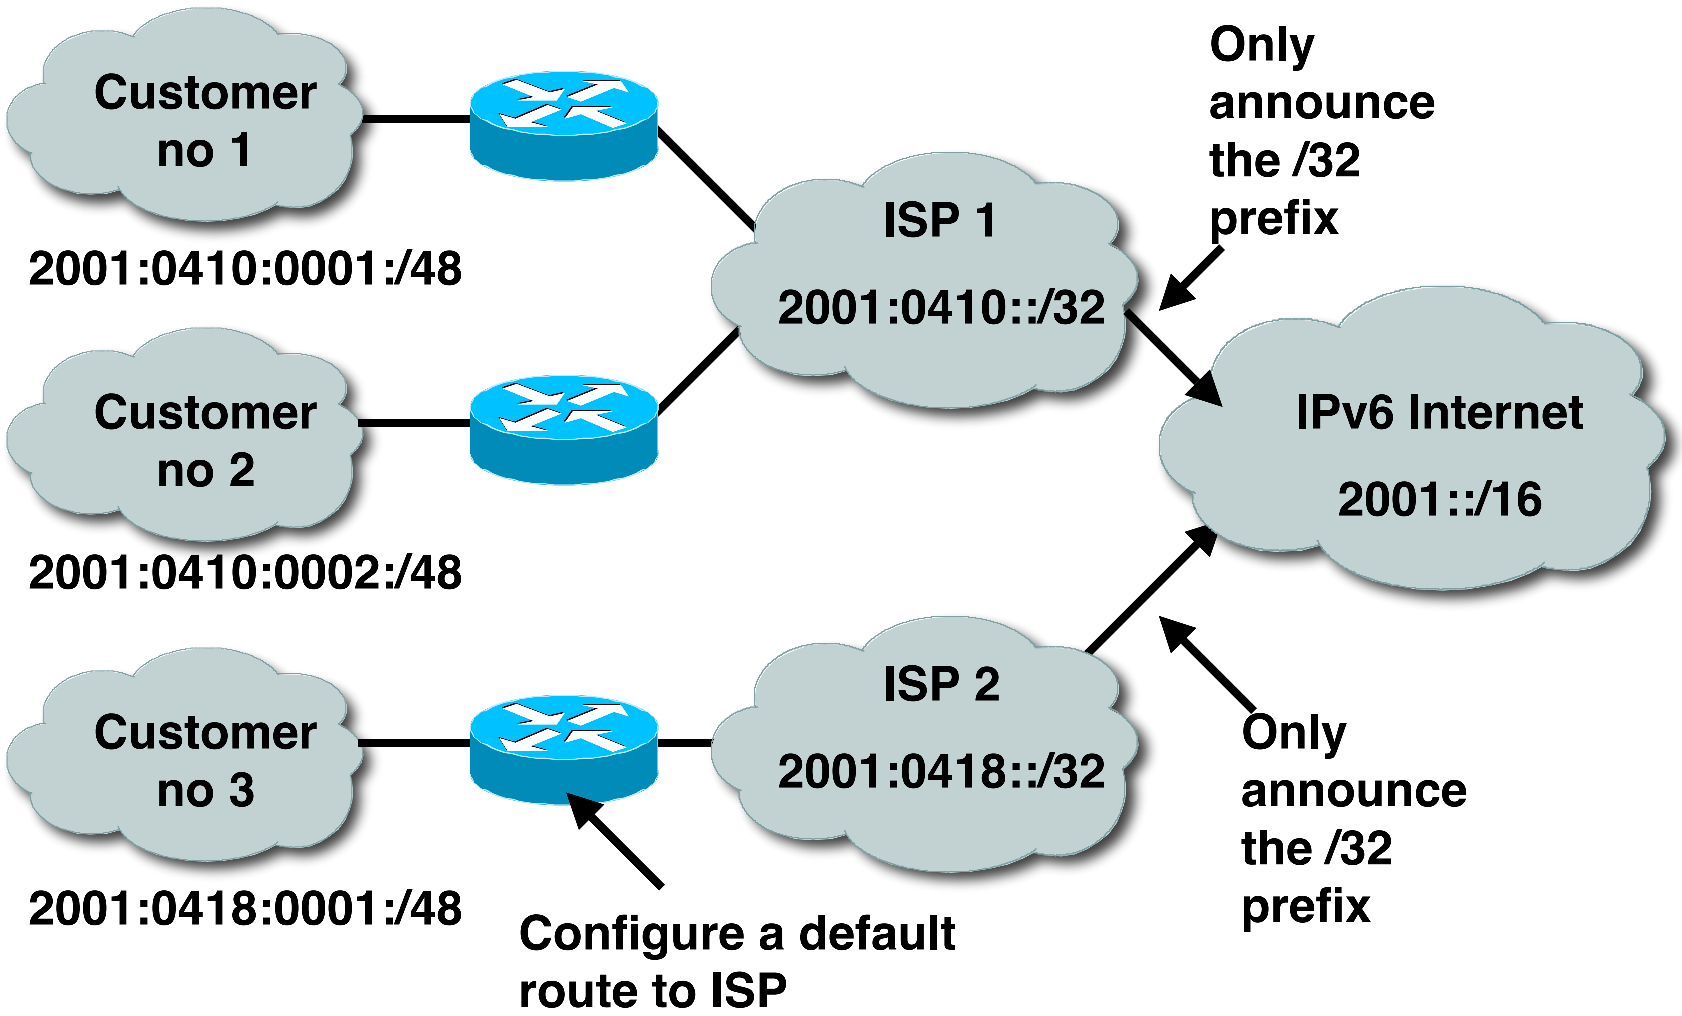

Expanded Address Space
Link-Local and Site-Local Unicast Addresses

- **Link-local addresses for use during auto-configuration and when no routers are present:**

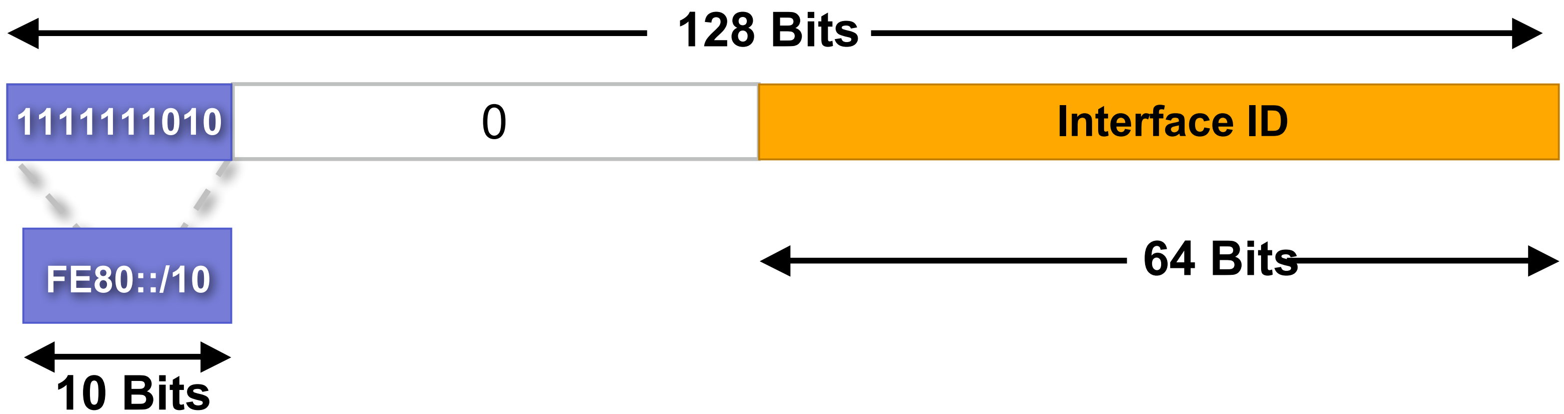

- **Site-local addresses equivalent of IPv4 private addresses**

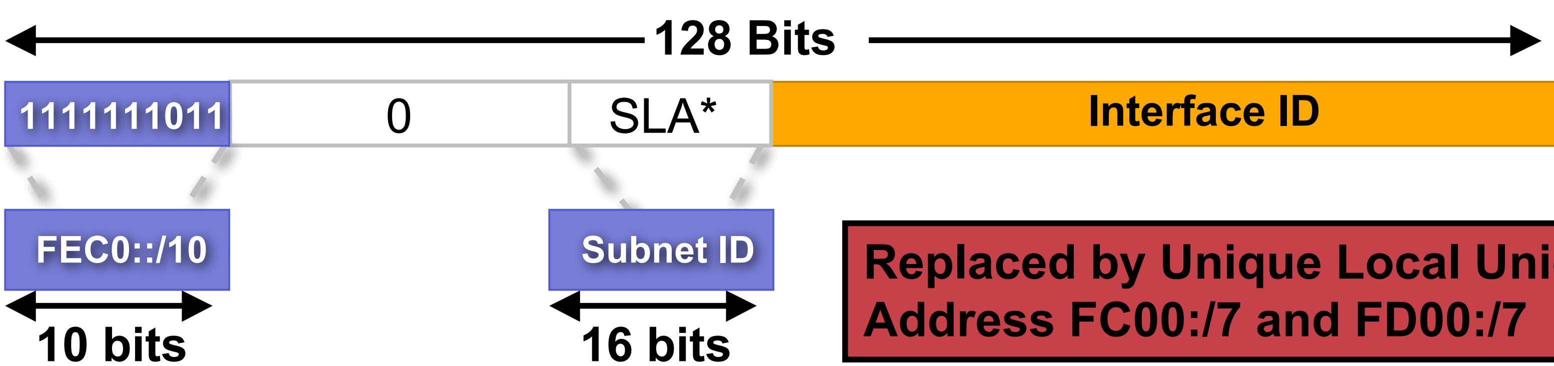

Replaced by Unique Local Unicast Address FC00:/7 and FD00:/7

Expanded Address Space
IPv4-Compatible & Mapped IPv6 Address

80 bits	16 bits	32 bits
0000…………………………….0000	0000	IPv4 Address

IPv4 Compatible IPv6 Addresses

80 bits	16 bits	32 bits
0000…………………………….0000	FFFF	IPv4 Address

IPv4 Mapped IPv6 Address

0:0:0:0:0:0:192.168.30.1 where X=0000 for Compatible , X=FFFF for Mapped
::192.168.30.1(compatible) and ::FFFF:192.168.30.1 (mapped)
::C0AB:1E01 (compatible) and ::FFFF:C0AB:1E01 (mapped)

Expanded Address Space
IPv6 eui-64

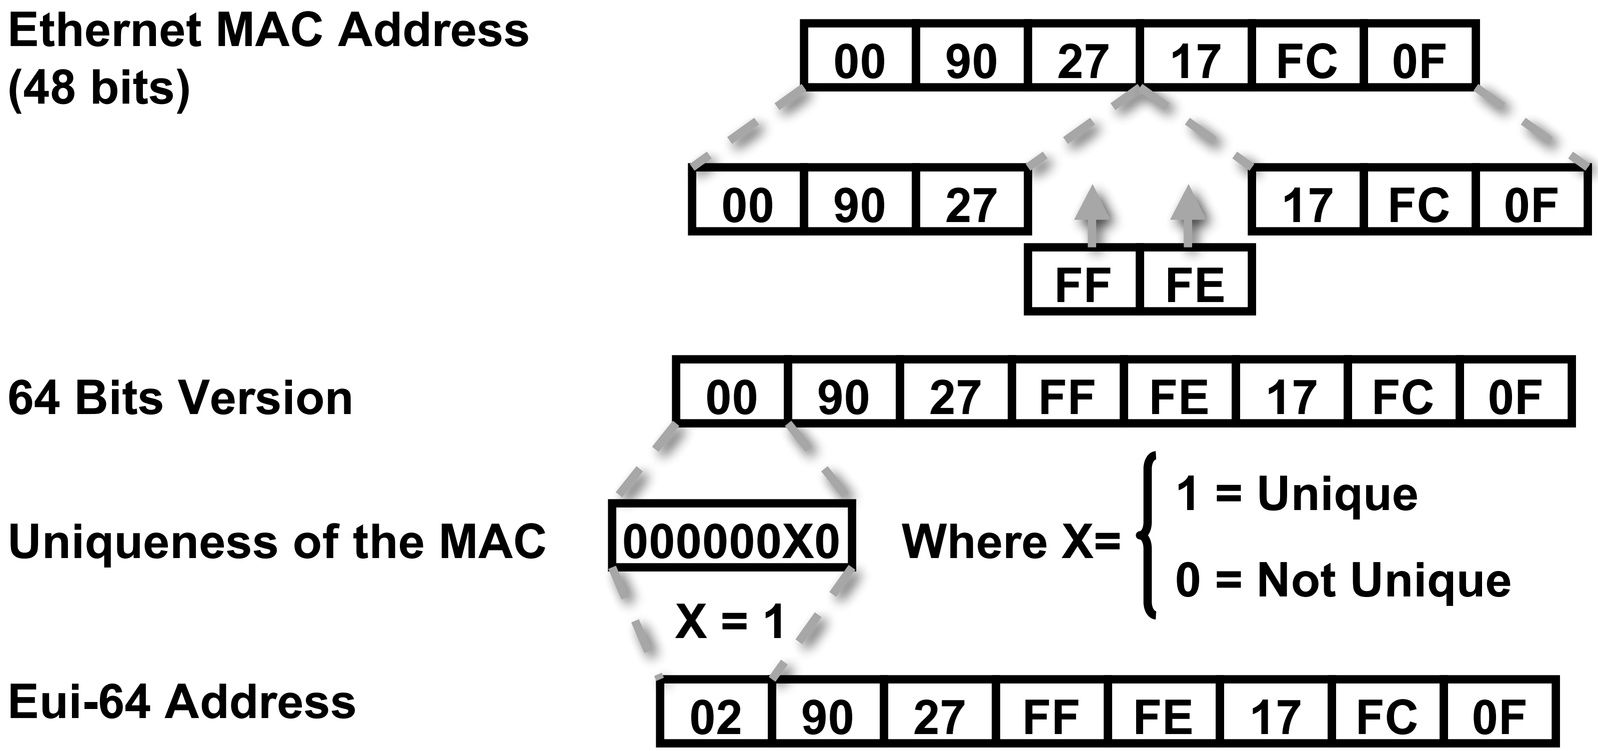

- Eui-64 address is formed by inserting "FFFE" and ORing a bit identifying the uniqueness of the MAC address

Expanded Address Space
Multicast Addresses (RFC 3306)

11111111	00PT	Scope	Res.	Plen	Prefix	Group ID
8	4	4	8	8	64	32

- **T of Lifetime Flag**

  0 if permanent,

  1 if temporary

- **P Flag**

  0—address not assigned on prefix

  1—prefix based assignment

  If P = 1:

  Plen—length of network prefix

  Prefix—network prefix, at most 64 bits

Scope :

0 reserved
1 interface-local scope
2 link-local scope
4 admin-local scope
5 site-local scope
8 organization-local scope
E global scope
F reserved

Update from RFC 2373
See also RFC 3307

Expanded Address Space
Multicast Assigned Addresses (RFC 3306)

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

Expanded Address Space
6to4 and ISATAP Addresses

- 6to4 (RFC 3056) – WAN Tunneling

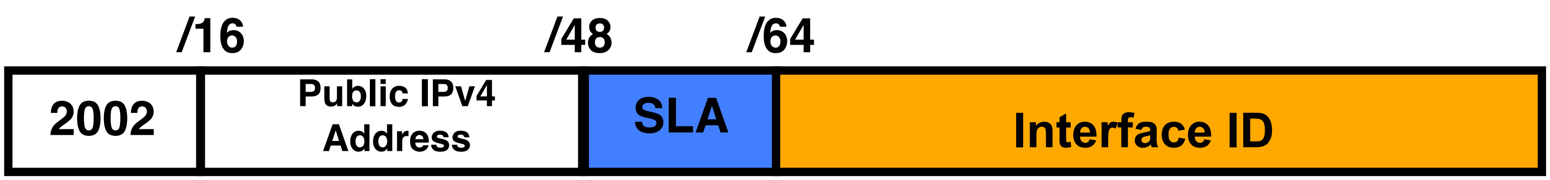

- ISATAP (Draft) – Campus Tunneling

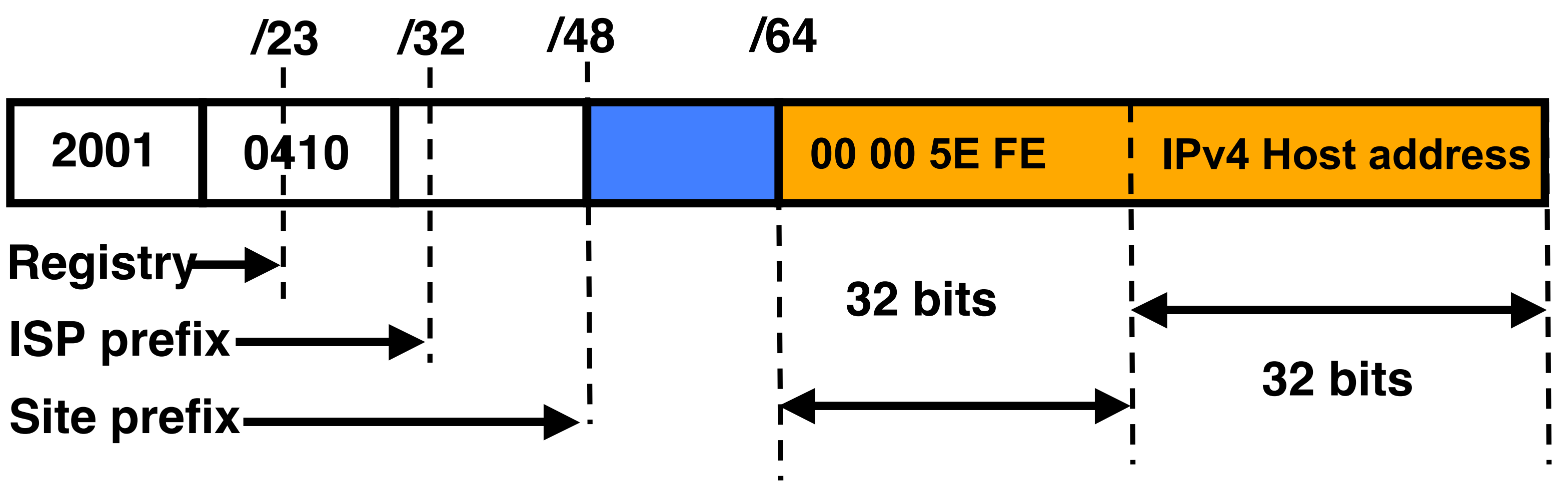

IPv4 & IPv6 Header Comparison

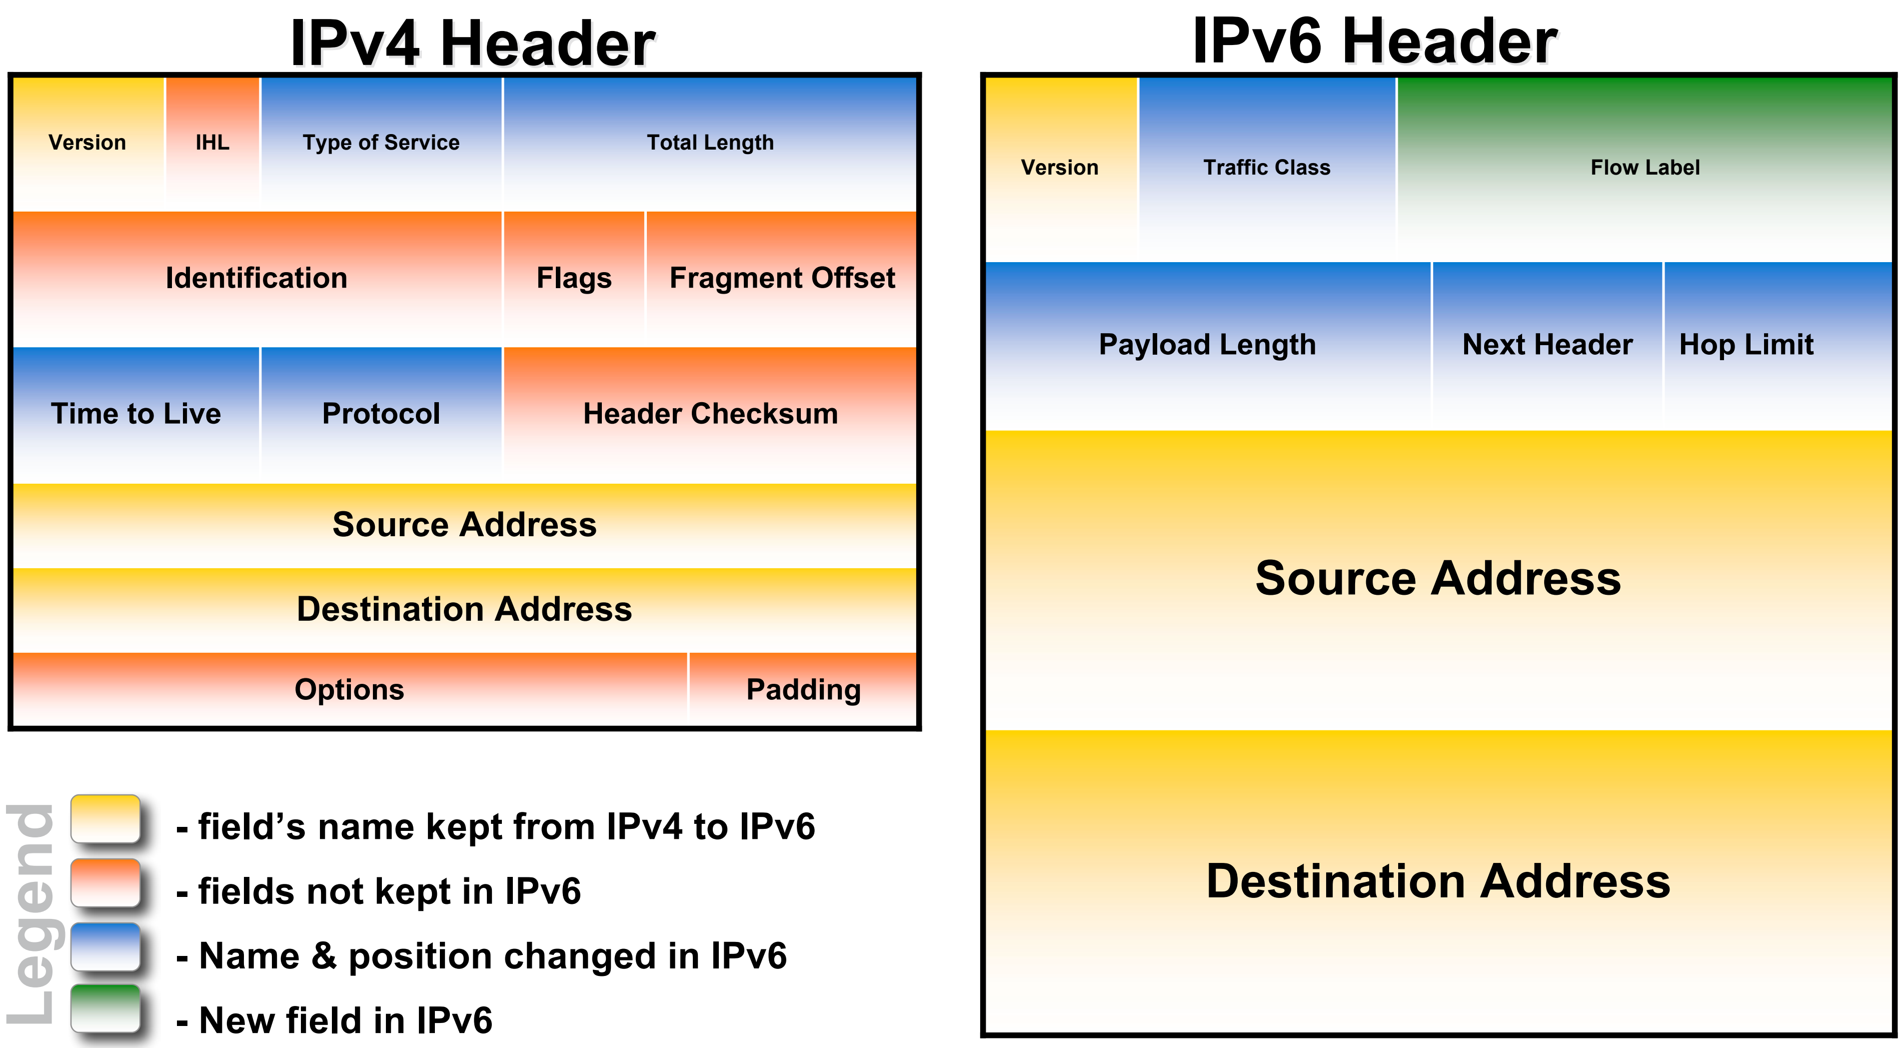

IPv4 & IPv6 Header Comparison Fields Renamed

- **Version:** A 4-bit field that contains the number 6 instead of 4.
- **Traffic Class:** An 8-bit field that is similar to the TOS field in IPv4. It tags the packet with a traffic class that can be used in differentiated services. These functionalities are the same as in IPv4.
- **Payload Length:** This is similar to the Total Length in IPv4, except it does not include the 40 byte header.
- **Hop Limit:** Like TTL field, decrements by one for each router.
- **Next Header:** Similar to the Protocol field in IPv4. The value in this field tells you what type of information follows, e.g. TCP, UDP, Extension Header.

IPv4 & IPv6 Header Comparison Fields Removed

- Header Length field is removed, because it is used to define the length of the header, since IPv4 header is variable length. IPv6 has a fixed header length so there is no need for this field
- Fragmentation field is used to split the packets into smaller segments over a network to accommodate smaller packet size interfaces. IPv6 does not do fragmentation, from operational experience, loss of one fragment cause complete retransmission which is very inefficient. IPv6 host learns the path MTU through MTU path discovery process. If sending host wants to do fragmentation it will do it through extension headers
- Identification field is used to identify the datagram from the source. Along with source IP address this is used to uniquely identify the datagram as it leaves the source. This is helpful in reassembling the fragmented packets. No fragmentation is done in IPv6 so no need for identification, also no need for flag
- Checksum not needed because both media access and upper layer protocol (UDP and TCP) have the checksum. IP is best-effort, plus removing checksum helps expedite packet processing

IPv4 & IPv6 Header Comparison Field Added

- 20-bit Flow Label field to identify specific flows needing special QoS
  - 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)

http://www.ietf.org/internet-drafts/draft-ietf-ipv6-flow-label-07.txt

Header Format Simplification
IPv6 Extension Headers

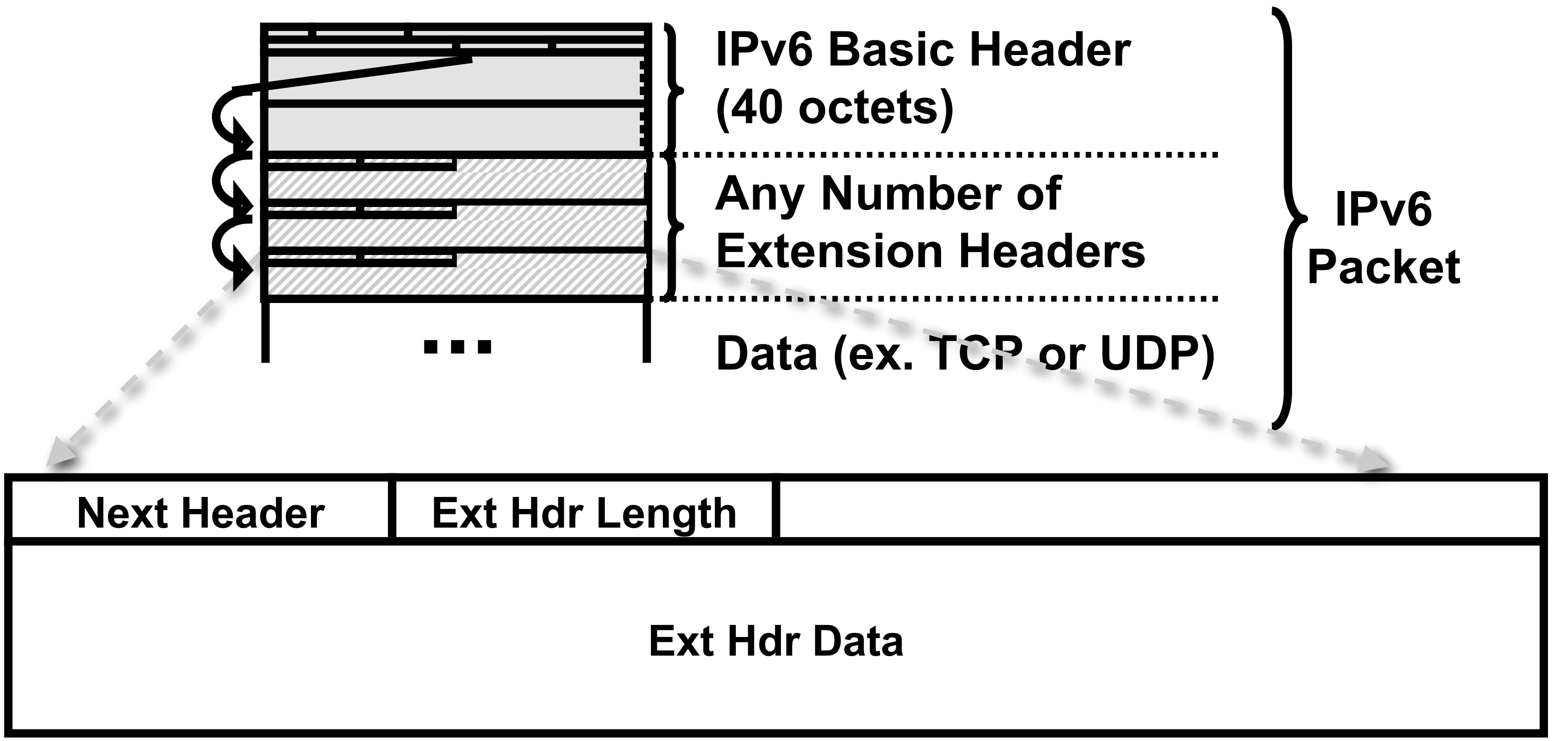

The value of this field determines the type of information following the basic IPv6 header; it can be transport layer packet, such as tcp/udp or can be extension header (6 types of extension headers)

Next header field of the previous header identifies the next extension header

Extension headers are optional following the IPv6 basic header

Each extension header is 8 octets(64 bits) aligned

Together all extension headers form a chained list of headers

IPv6 Extension Header Types
Hop By Hop Header (Protocol 0)

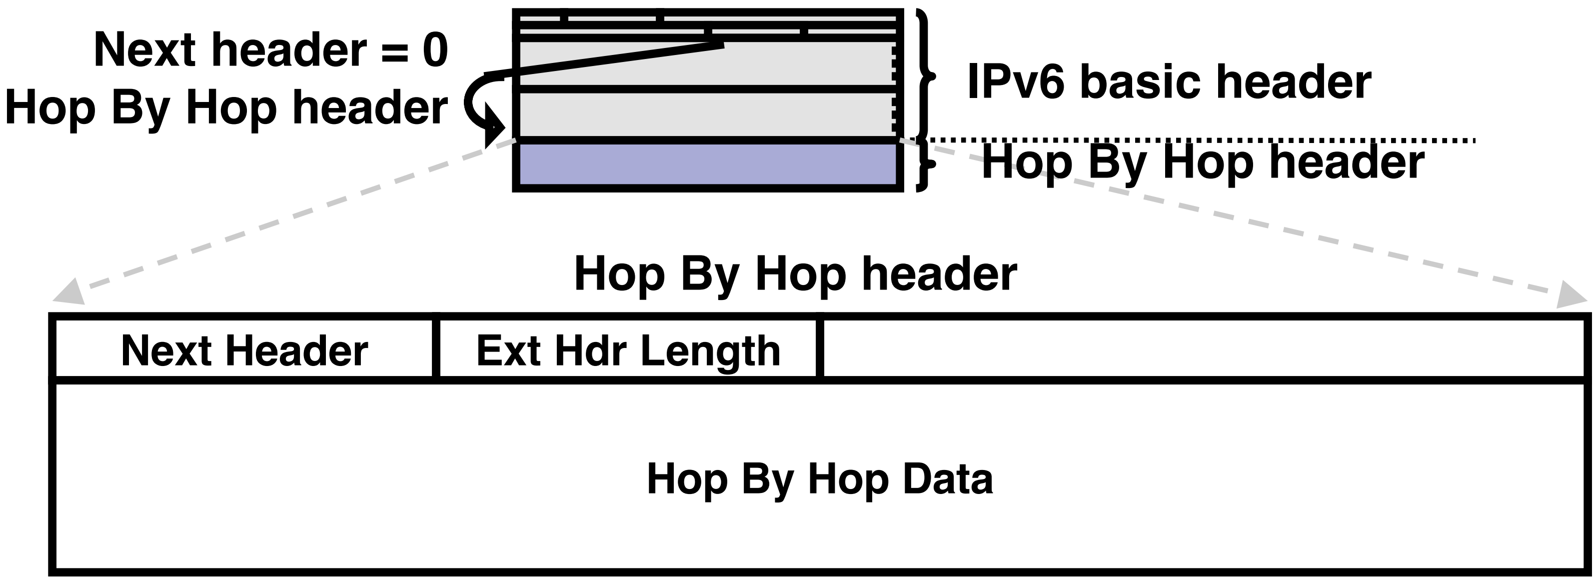

Read and Processed by every node and router along the delivery path.

When Presents, Follows Immediately after the Basic IPv6 Packet Header

Used for router alerts, An example of applying this option would RSVP, because each router needs to look at it.

IPv6 Extension Header Types
Destination Option Header (Protocol 60)

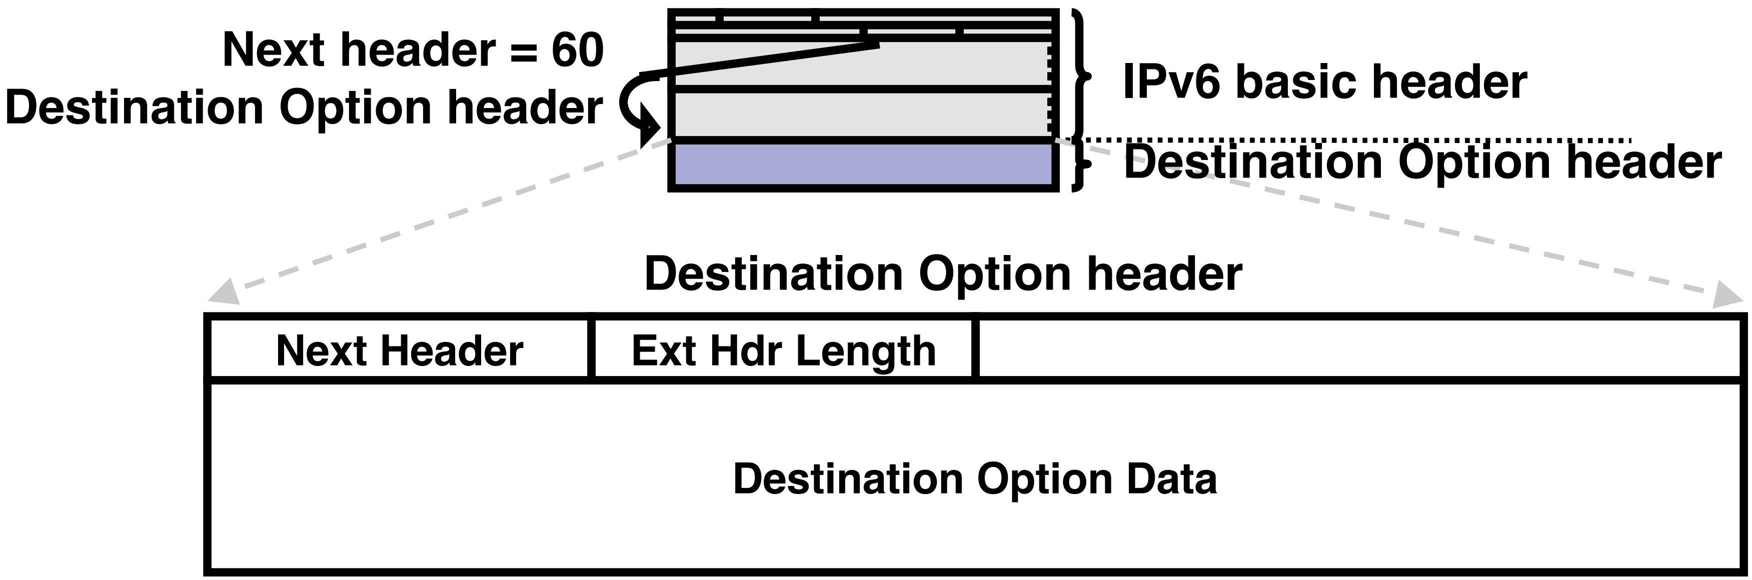

Carries optional information that is specifically targeted to packet's destination address.

The Mobile IPv6 uses this option to exchange registration messages between mobile nodes and the home agent.

IPv6 Extension Header Types
Routing Header (Protocol 43)

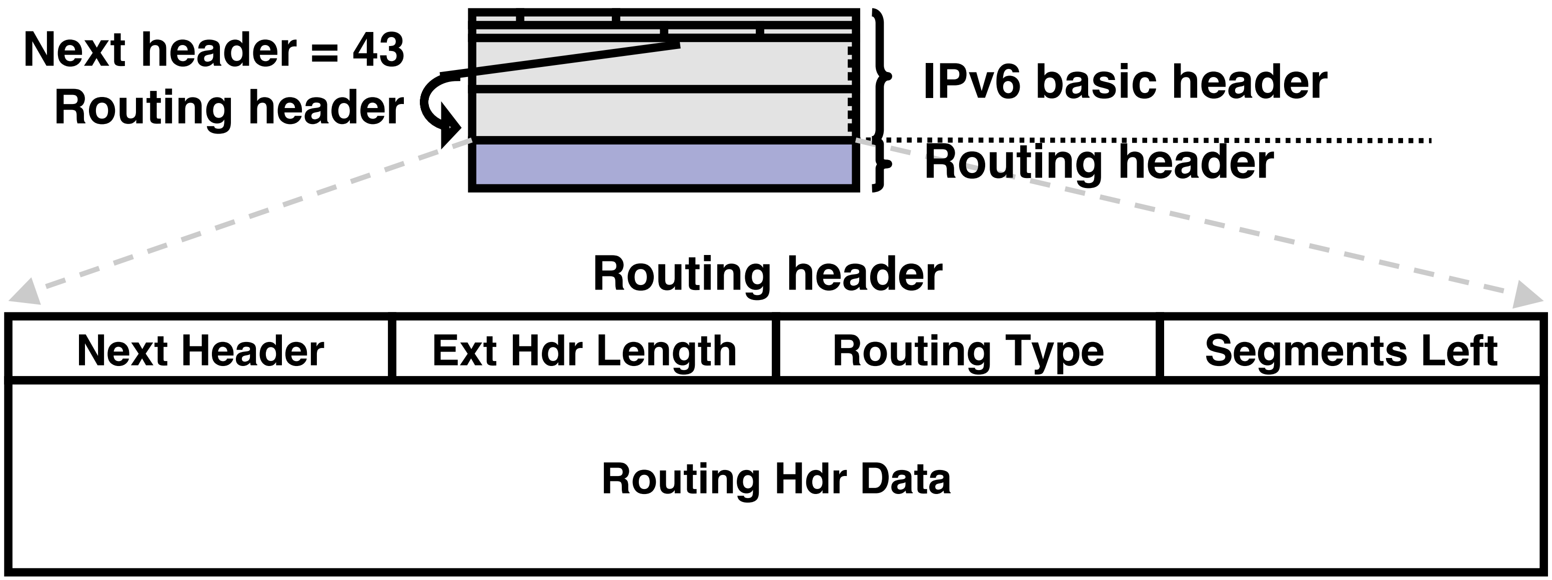

Routing header forces the routing through a list of intermediate routers.

This is similar to the "Loose Source Route" option in IPv4.

IPv6 Extension Header Types Fragment Header (Protocol 44)

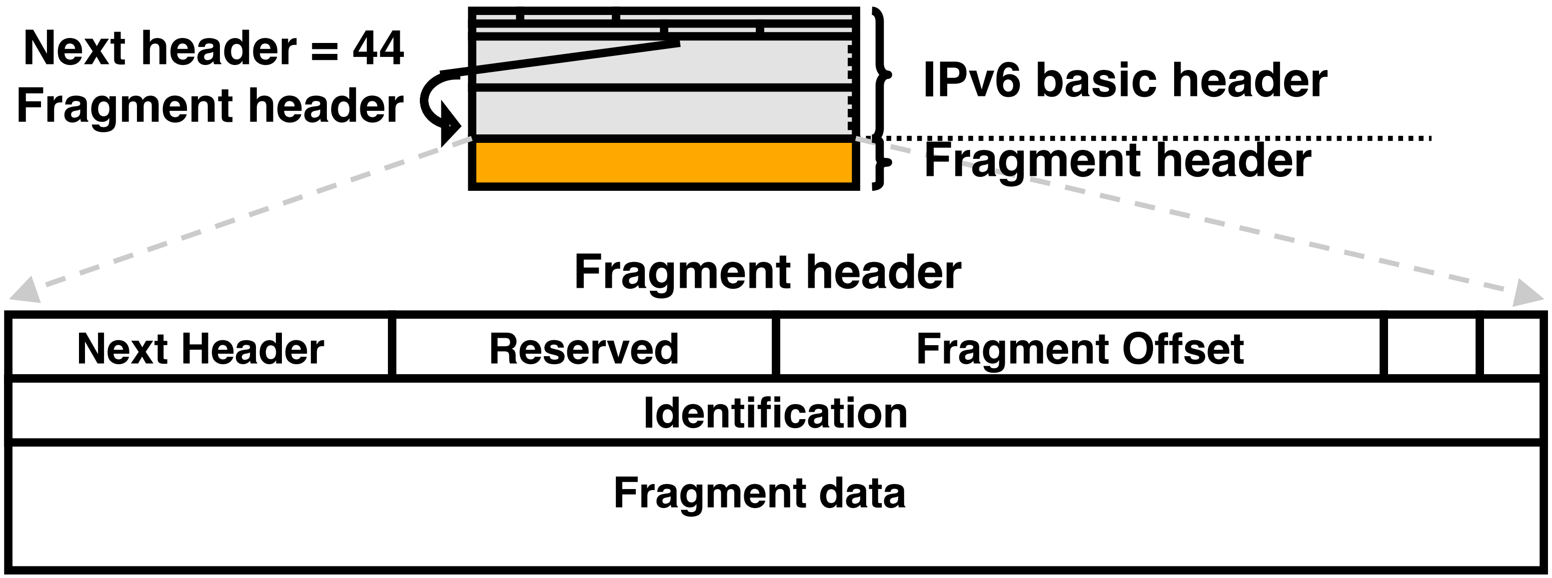

Used by Source When Packet Is Fragmented

Fragment Header Is Used in Each Fragmented Packet

Fragment Offset: Identifies the position of the specific fragment in the full original packet.

Identification: A number to identify fragments of the same original packet.

Fragment offset: Used by destination node to reassemble the packet back to it's original form.

IPv6 Extension Header Types
IPSec Authentication Header (Protocol 51)

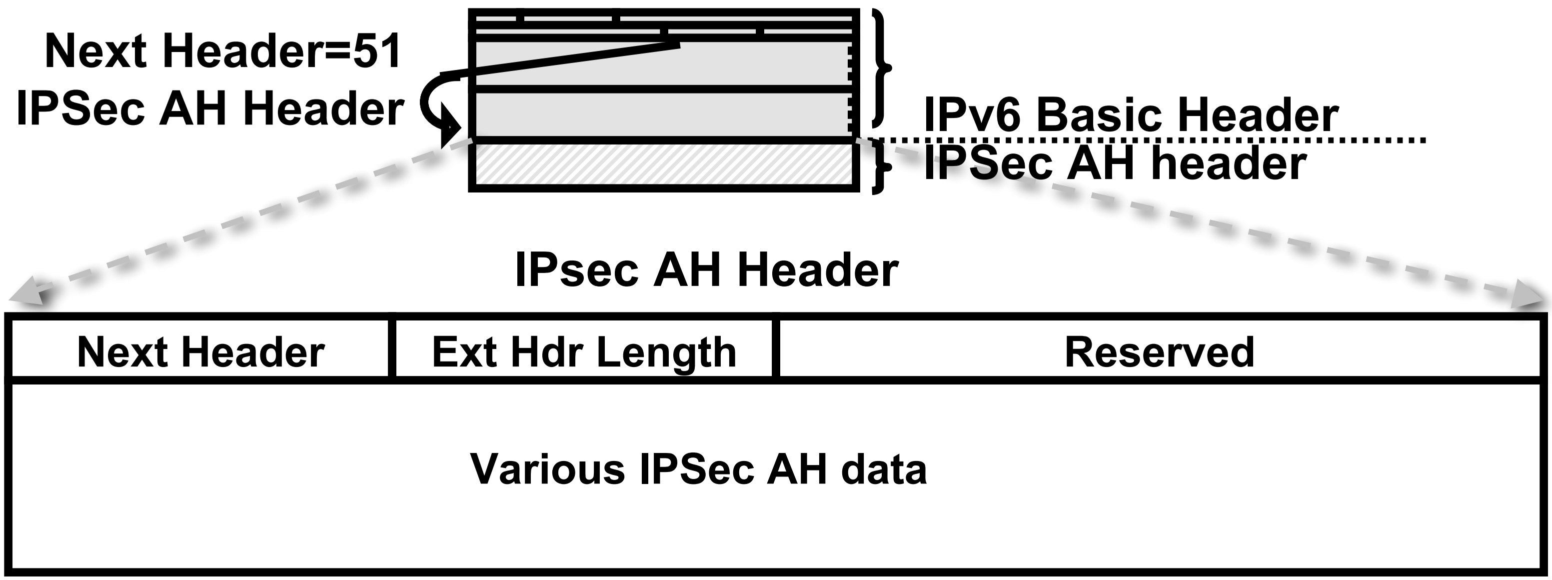

IPSec Authentication Header (AH) provides:

- Confidentiality
- Integrity
- Authentication of the source

IPv6 Extension Header Types
IPSec ESP (Protocol 50)

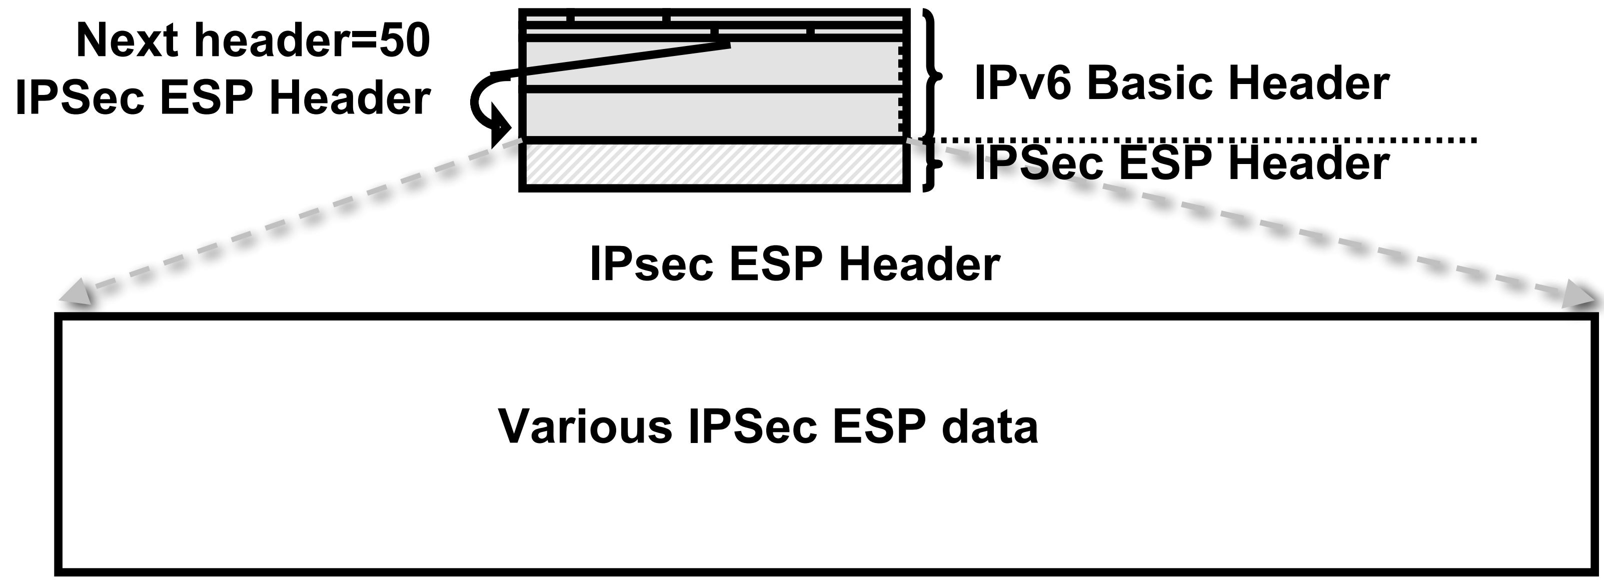

IPSec Encapsulating Security Payload (ESP) provides:

- Confidentiality
- Integrity
- Authentication of the source

Upper Layer Header
User Datagram Protocol (Protocol 17)

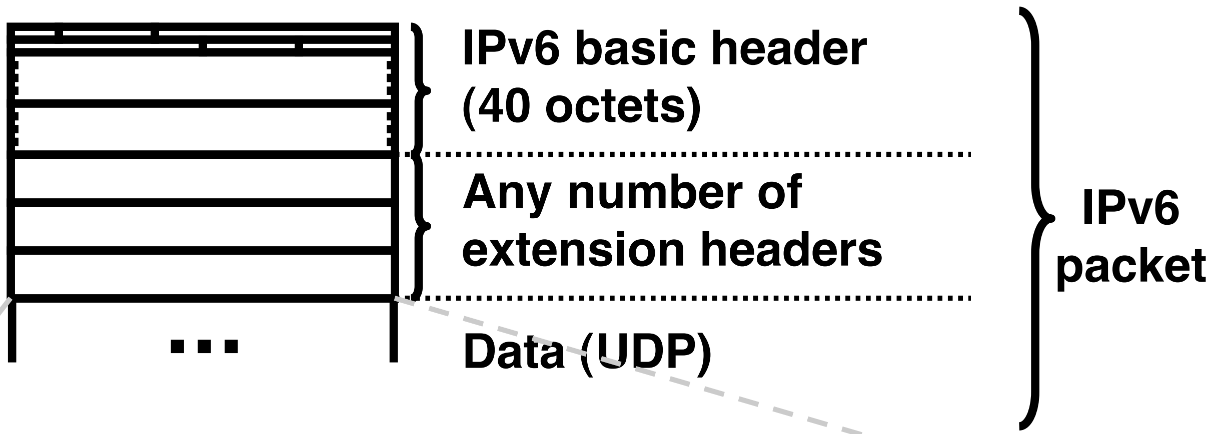

UDP Packet

Source Port	Destination Port
Length	UDP Checksum
UDP Data Portion	

- Upper layer (UDP, TCP, ICMPv6) checksum must be computed
- These Are the Typical Headers Used Inside a Packet to Transport Data
- This could be UDP (Protocol 17), TCP (Protocol 6) or ICMPv6 (Protocol 58)

Upper Layer Header
ICMPv6 (Protocol 58)

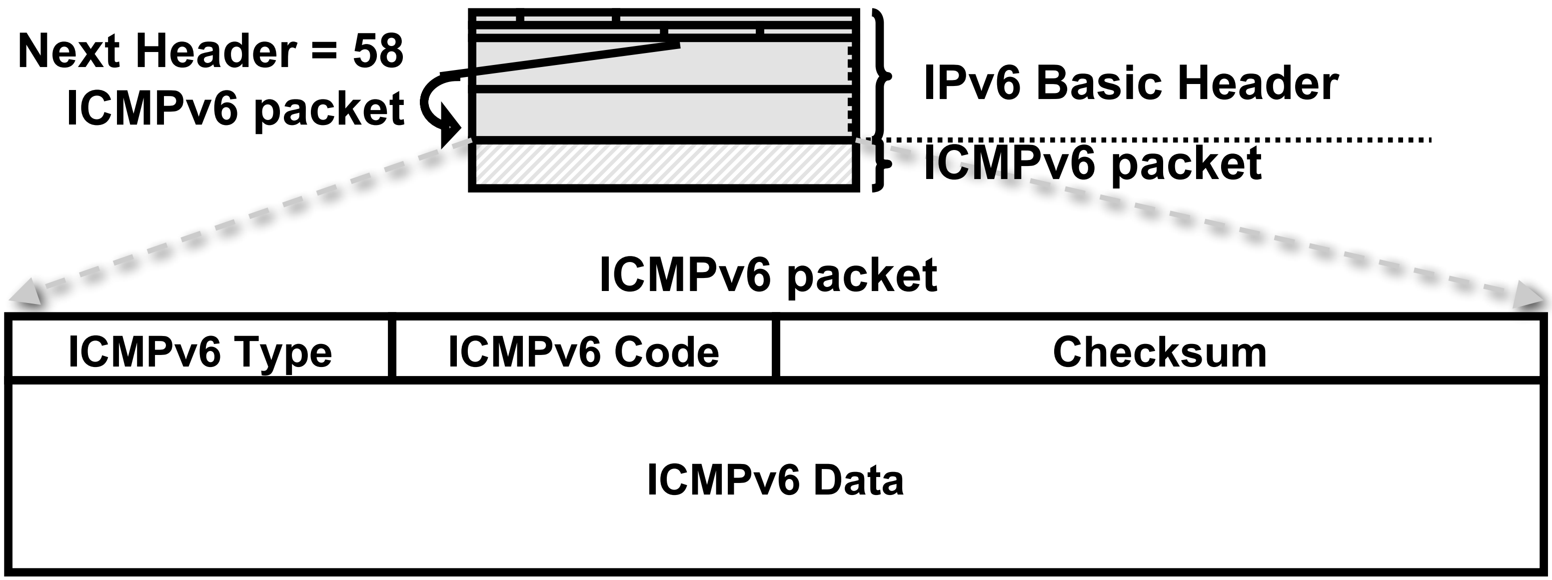

ICMPv6 is similar to IPv4: Provides diagnostic and error messages.

Additionally it's used for neighbor discover, path MTU discovery and Mcast listener discovery (MLD)

Header Format Simplification
Path MTU Discovery

- **Definitions:**
  - **Link MTU is link's maximum transmission unit.**
  - **Path MTU is the minimum MTU of all the links in a path between a source and a destination**
- **Minimum link MTU for IPv6 is 1280 octets (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 octets:**
  - **For each destination, start by assuming MTU of first-hop link**
  - **If a packet reaches a link in which it cannot fit, will invoke ICMP "packet too big" message to source, reporting the link's MTU; MTU is cached by source for specific destination**

Header Format Simplification
Path MTU Discovery

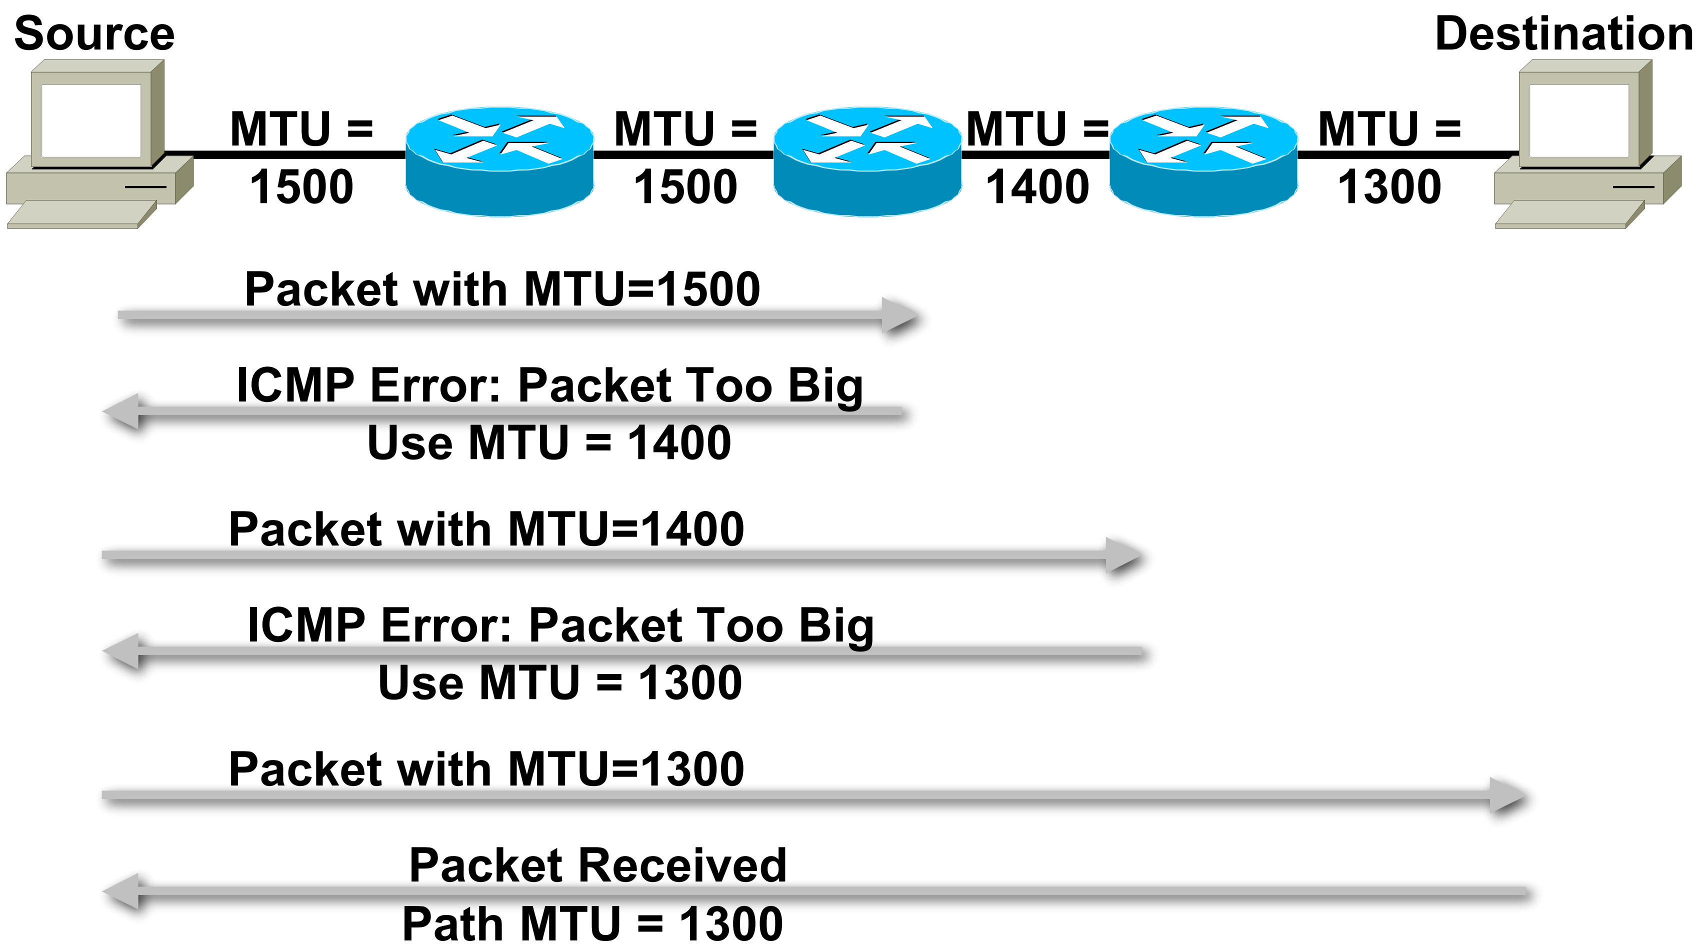

- **Minimum link MTU for IPv6 is 1280 octets (versus 68 octets for IPv4)**

Header Format Simplification Neighbour Discovery (RFC 2463)

Protocol built on top of ICMPv6 (RFC 2463)

Combination of IPv4 protocols (ARP, ICMP, IGMP,…)

- Uses ICMP messages and solicited-node multicast addresses
- Determines the link-layer address of a neighbor on the same link
- Finds neighbor routers
- Verifies the reachability of neighbors
- Comprised of different message types:
  - Neighbor Solicitation (NS)/Neighbor Advertisment(NA)
  - Router Soliciation(RS)/Router Advertisement(RA)
  - Redirect
  - Renumbering

Neighbor Solicitation & Advertisement

Neighbor Solicitation:
ICMP type = 135
Src = A
Dst = Solicited-node multicast Address
Data = link-layer address of A
Query = what is your link-layer 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

IPv6 Auto-Configuration

- **Stateless (RFC2462)**

  Router solicitation are sent by booting nodes to request RAs for configuring the interfaces.

  Host autonomously configures its own Link-Local address.

- **Stateful**

  DHCPv6

- **Renumbering**

  Hosts renumbering is done by modifying the RA to announce the old prefix with a short lifetime and the new prefix.

  Router renumbering protocol (RFC 2894), to allow domain-interior routers to learn of prefix introduction / withdrawal

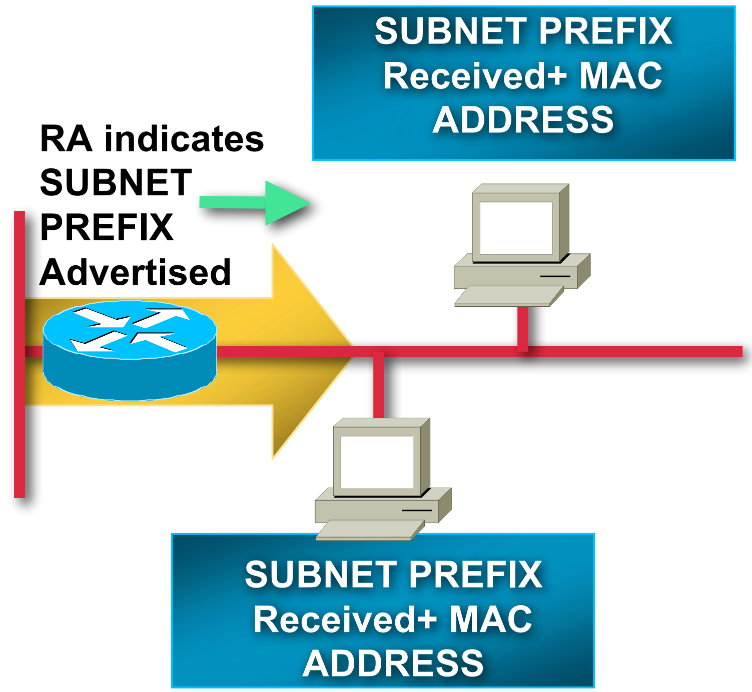

At boot time, an IPv6 host build a Link-Local address, then its global IPv6 address(es) from RA

Stateless Autoconfiguration

1 - ICMP Type = 133 (RS)

Src = Link-local Address (FE80::/10)

Dst = All-routers multicast Address (FF02::2)

Query= please send RA

2 - ICMP Type = 134 (RA)

Src = Link-local Address (FE80::/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.

Duplicate Address Detection (DAD)

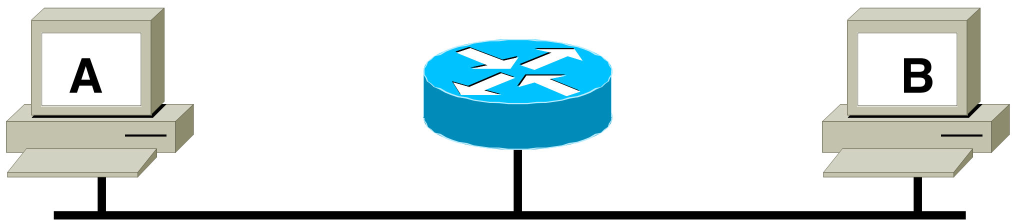

1.Host A boots up and assigns it self LINK LOCAL ADDRESS (FF80::/10)

2.Host A sends RS (ICMP Type 133)

3.Host A receives RA (ICMP Type 134) with subnet prefix (2001:0410:1/64)

Now the Host A wants to assign itself a unique global unicast address 2001:0410:1::34:123A . Before it does that it sends out DAD request to all nodes on the link by doing the following:

4.Host A sends NS (ICMP Type 135) with:

Source address (::)

Destination address FF02::1:FF34:123A (solicited-node Mcast address for 2001:0410:1::34:123A)

5.If Host A does not receive a reply back it will assign itself 2001:0410:1::34:123A

Redirect

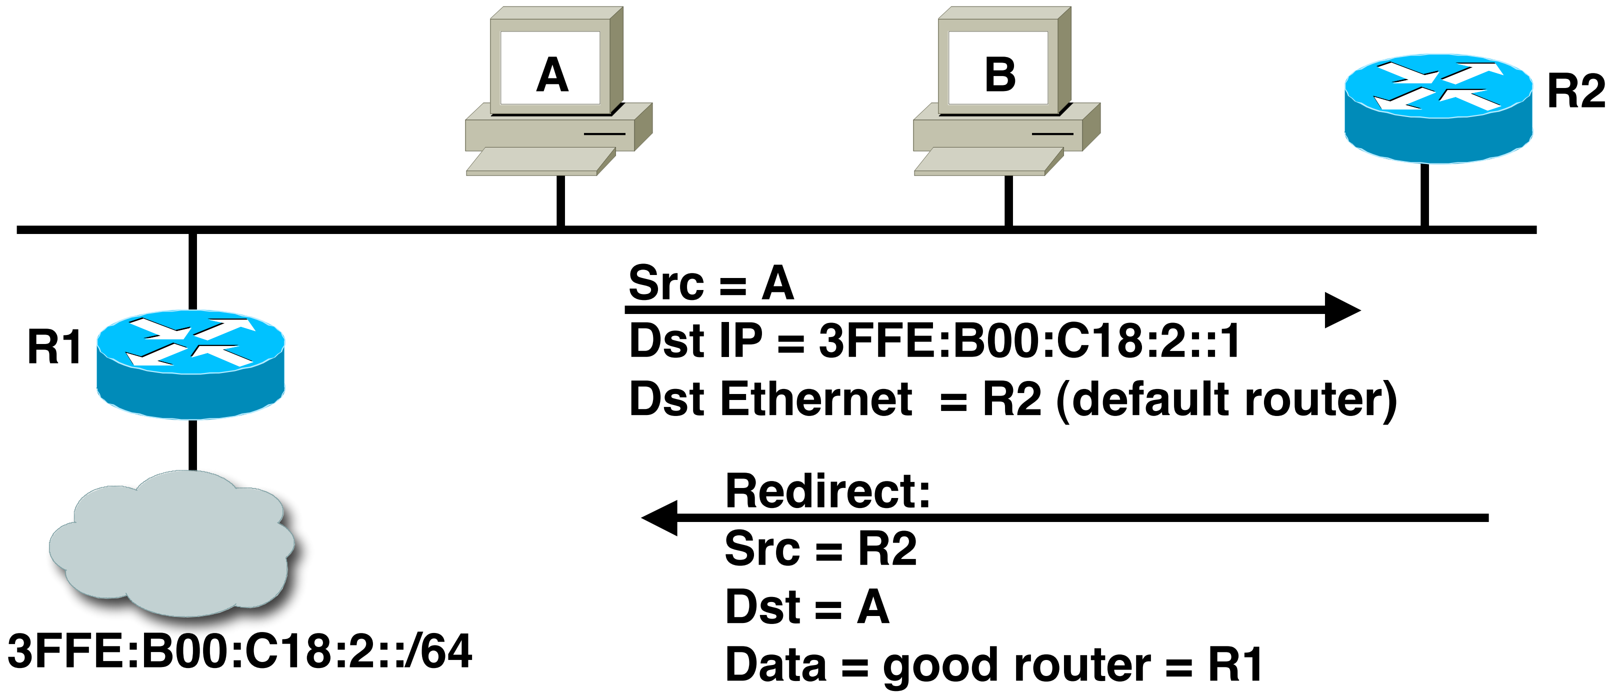

Redirect is used by a router to signal the reroute of a packet to a better router.

Renumbering

RA packet definitions:

- ICMP Type = 134
- Src = Router Link-local Address
- Dst = All-nodes multicast address
- Data= 2 prefixes:
  - Current prefix (to be deprecated) with short lifetime
  - New prefix (to be used) with normal lifetime

Renumbering - Modify the RA to announce the old prefix with a short lifetime and the new prefix.

IPv6 Address Configuration

LAN: 3ffe:b00:c18:1::/64

Ethernet0

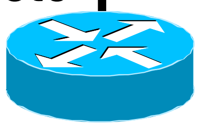

```
IPv6 unicast-routing

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
```

MAC address: 0060.3e47.1530

```
router# show ipv6 interface Ethernet0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::260:3EFF:FE47:1530
Global unicast address(es):
    3FFE:B00:C18:1:260:3EFF:FE47:1530, subnet is 3FFE:B00:C18:1::/64
  Joined group address(es):
    FF02::1:FF47:1530
    FF02::1
    FF02::2
  MTU is 1500 bytes
```

Sample Configuration


```
3640-a#sh run
ipv6 unicast-routing
interface Ethernet0/0
ipv6 address 3FFE:ABCD:ABCD:1::/64 eui-64
 ipv6 address FEC0::1/64

3640-a#sho int
Ethernet0/0 is up, line protocol is up
address is 0010.7bc7.3440

3640-a#show ipv6 int
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is
FE80::210:7BFF:FEC7:3440
  Global unicast address(es):
    3FFE:ABCD:ABCD:1:210:7BFF:FEC7:3440,
subnet is 3FFE:ABCD:ABCD:1::/64
    FEC0::1, subnet is FEC0::/64
```

```
3640-b#show run
ipv6 unicast-routing
interface Ethernet0/0
ipv6 address 3FFE:ABCD:ABCD:1::/64 eui-64
 ipv6 address FEC0::2/64

3640-b#show int
Ethernet0/0 is up, line protocol is up
address is 0010.7bc7.38c0

3640-b#show ipv6 int e0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is
FE80::210:7BFF:FEC7:38C0
  Global unicast address(es):
    3FFE:ABCD:ABCD:1:210:7BFF:FEC7:38C0,
subnet is 3FFE:ABCD:ABCD:1::/64
    FEC0::2, subnet is FEC0::/64
```

Sample Configuration (cont.)


```
3640-a#ping ipv6
3FFE:ABCD:ABCD:1:210:7BFF:FEC7:38C0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to
3FFE:ABCD:ABCD:1:210:7BFF:FEC7:38C0,
timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-
trip min/avg/max = 1/2/4 ms

3640-a#ping FEC0::2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FEC0::2,
timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-
trip min/avg/max = 1/2/4 ms
```

```
3640-b#show run
ipv6 unicast-routing
interface Ethernet0/0
ipv6 address 3FFE:ABCD:ABCD:1::/64 eui-64
 ipv6 address FEC0::2/64

3640-b#show int
Ethernet0/0 is up, line protocol is up
address is 0010.7bc7.38c0

3640-b#show ipv6 int e0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is
FE80::210:7BFF:FEC7:38C0
  Global unicast address(es):
    3FFE:ABCD:ABCD:1:210:7BFF:FEC7:38C0,
subnet is 3FFE:ABCD:ABCD:1::/64
    FEC0::2, subnet is FEC0::/64
```

IPv6 Addressing Planning & Assignments

IPv6 Technology Comparison

Service	IPv4 Solution	IPv6 Solution
Addressing Range	32-bit, Network Address Translation	128-bit, Multiple Scopes
Autoconfiguration	DHCP	Stateless, Reconfiguration, DHCP
Security	IPSec	IPSec Mandated, Works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality-of-Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
Multicast	IGMP/PIM/MBGP	MLD/PIM/MBGP, Scope Identifier

IPv6 Addressing

- IPv6 addressing rules are covered by multiples RFC's
  - Architecture defined by RFC 3513
- Address types are:
  - Unicast: one to one
    - Global
    - Link local (FE80)
    - ~~Site local (FEC0)~~ DEPRECATED—Replaced by Unique Local (RFC 4193)
    - Anycast: one to nearest (allocated from unicast)
  - Multicast (FF): one to many
  - Reserved
- A single interface may be assigned multiple IPv6 addresses of any type (unicast, anycast, multicast)
  - No broadcast address. Now uses multicast

Interface Address Set

- Loopback (Required) — (Only assigned to a single interface per node)
- Link local (Required) — (Required on all interfaces)
- Unique local (Optional) — (Addressing valid only within a site)
- Auto-configured 6to4 (Optional) — (If IPv4 public is address available)
- Solicited node Multicast (Required) — (Required for neighbor discovery - DAD)
- All node multicast (Required)
- Global (Optional) — (Globally routed prefix – Does not mean globally available)

Aggregatable Global Unicast Addresses

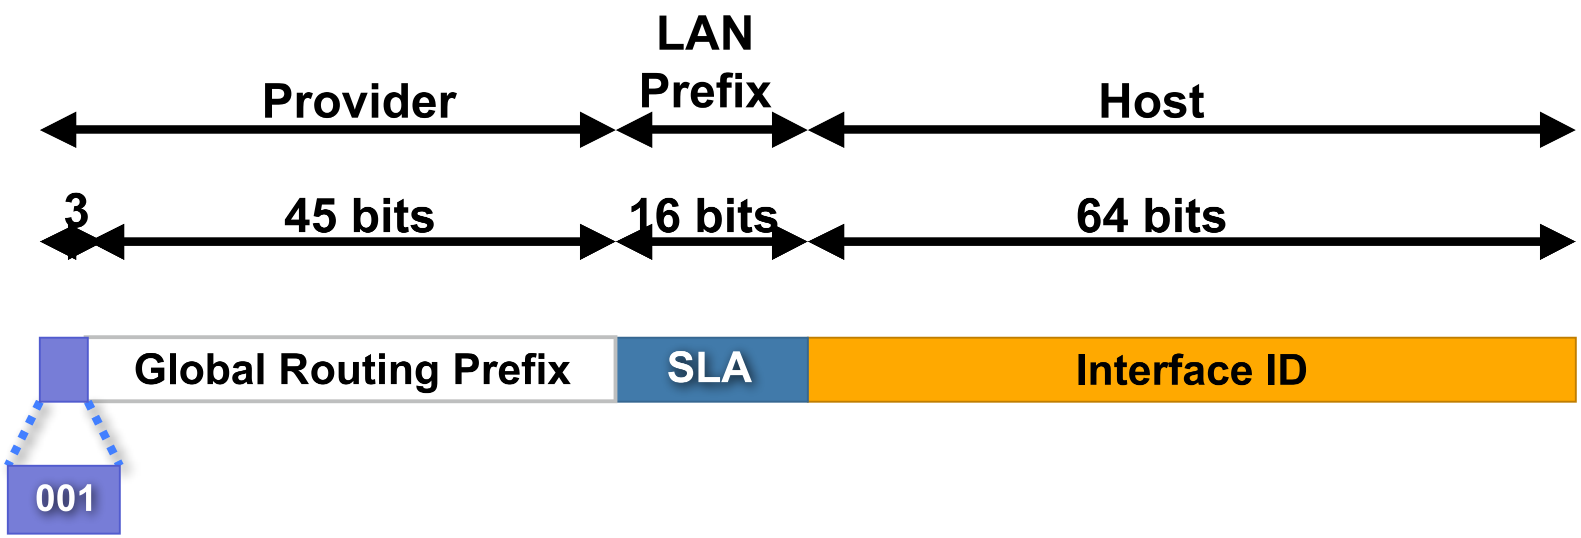

- Aggregatable global unicast addresses are:
  - Addresses for generic use of IPv6
  - Structured as a hierarchy to keep the aggregation
- See draft-ietf-ipngwg-addr-arch-v3-07

Address Allocation Policy

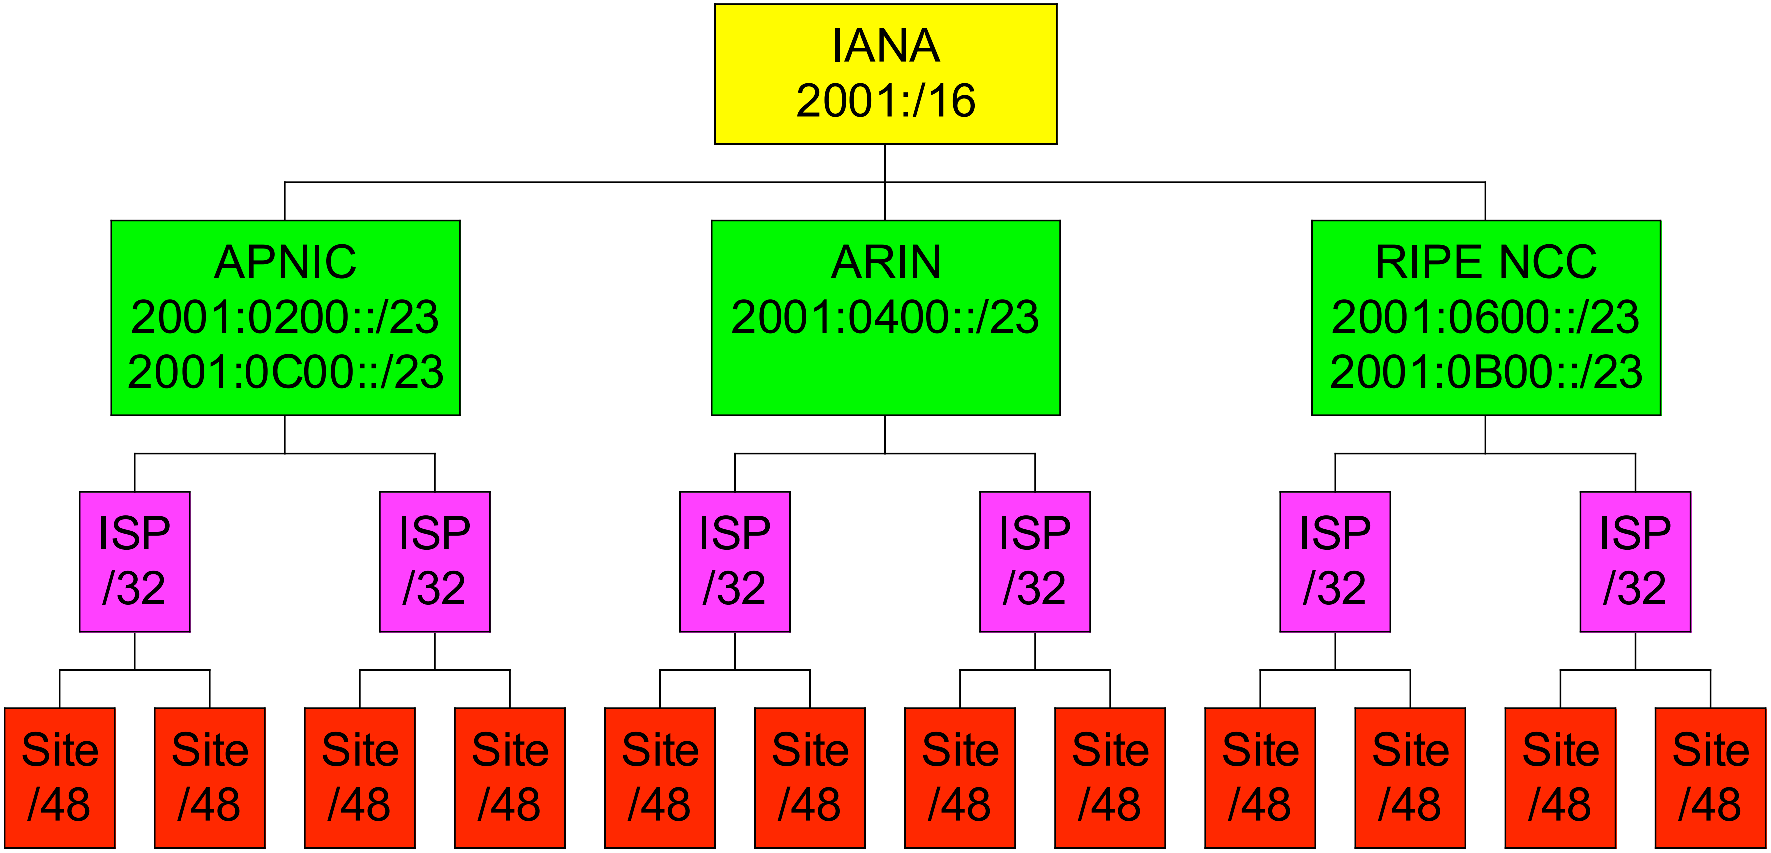

Address Allocation Policy

Administered by IANA to Regional Registries: ARIN, APNIC, RIPE, LACNIC

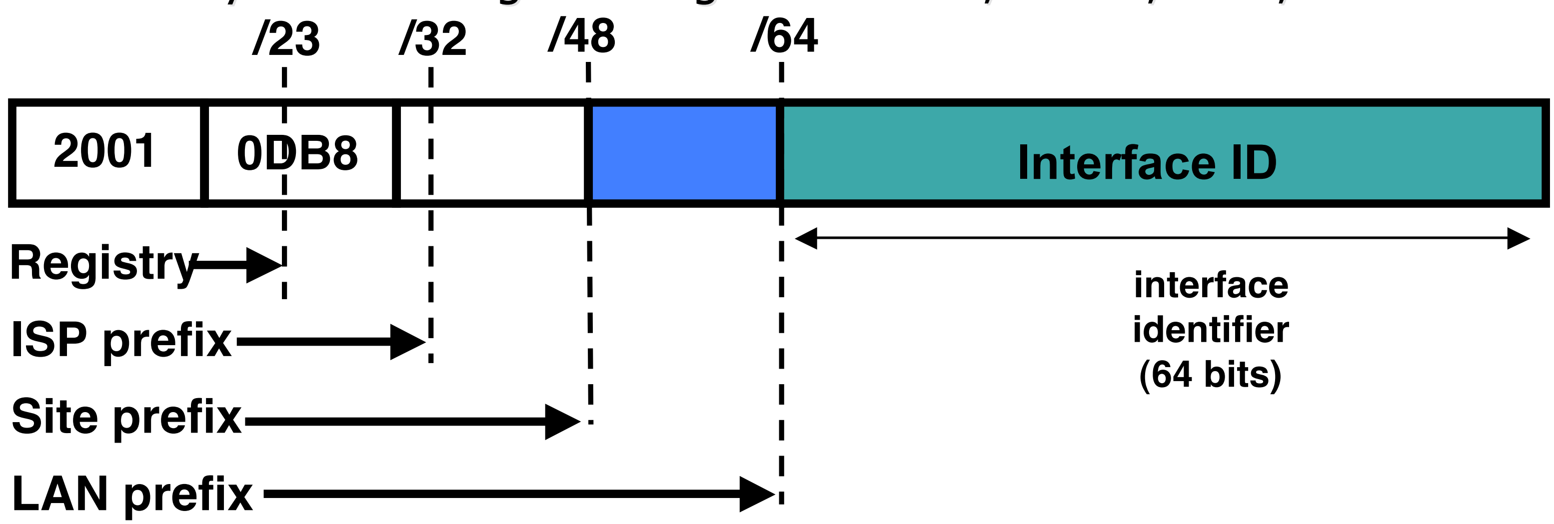

The allocation process is under review by the Registries:

- IANA has allocated 2001::/16 to the registries
- Each registry gets a /23 prefix from IANA
- Larger allocation done on specific request, eg. /20 recently allocated to one ISP in Europe
- With the new policy, Registry allocates a /32 prefix to an IPv6 ISP
- Then the ISP allocates a /48 prefix to each customer (or potentially /64)
- http://www.ripe.net/ipv6/global-ipv6-assign-2002-04-25.html

IPv6 Addressing

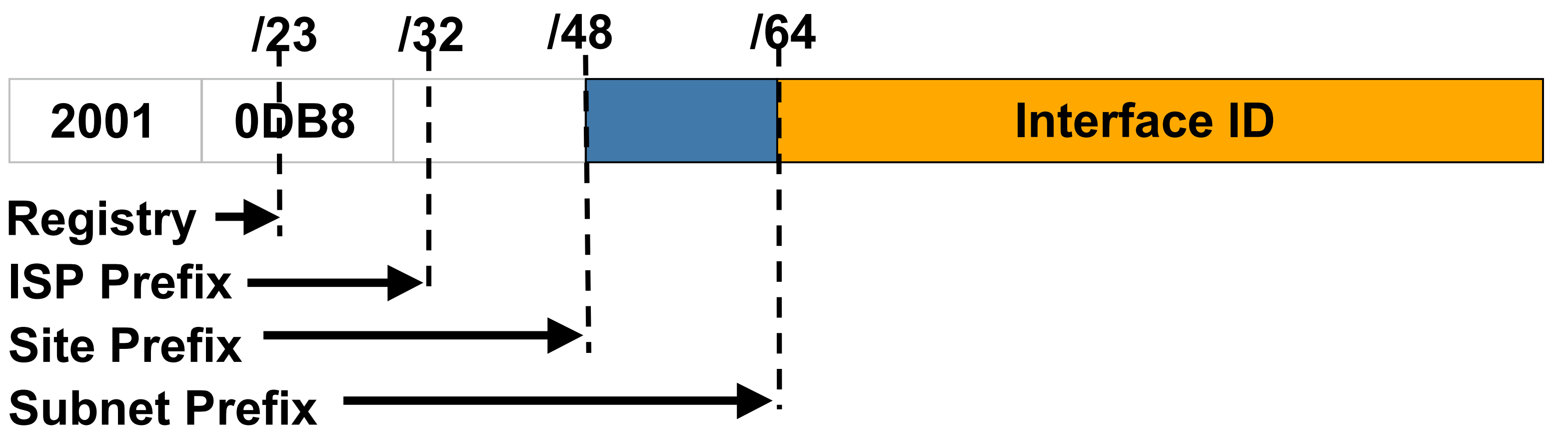

Represented as:

x:x:x:x:x:x:x:x where x is a 16-bit hexadecimal field

- 2001:0DB8:C003:0001:0000:0000:0000:BEEF
- 2001:DB8:C003:1:0:0:0:BEEF
- 2001:DB8:C003:1::BEEF

- 0:0:0:0:0:0:0:1 --> ::1 - Loopback address

Hierarchical Addressing & Aggregation

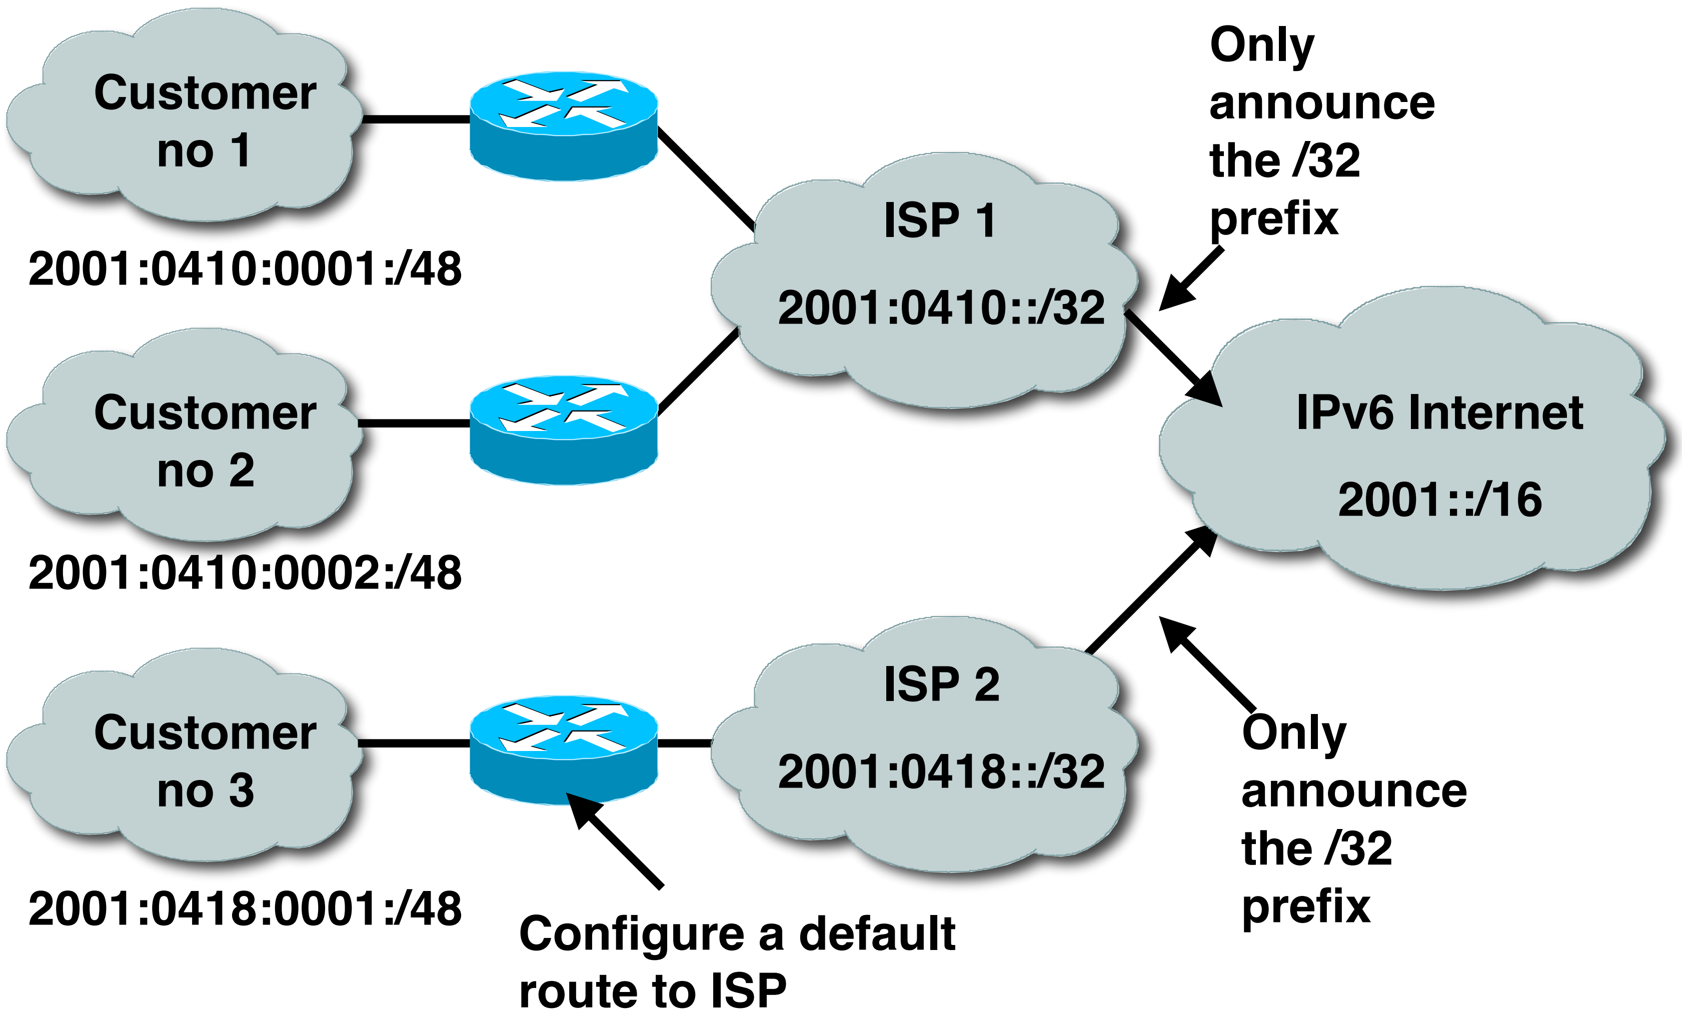

First Steps

- Talk with your service provider(s) about getting your IPv6 prefix(s) and what kind of services they plan to support
- Start a pilot or lab network to gain familiarity with IPv6 and YOUR applications
- Include IPv6 in your investment strategy for new operating systems, networking gear, deployment and management
- Understand the reasons why you are going this route

Hierarchical Addressing and Aggregation

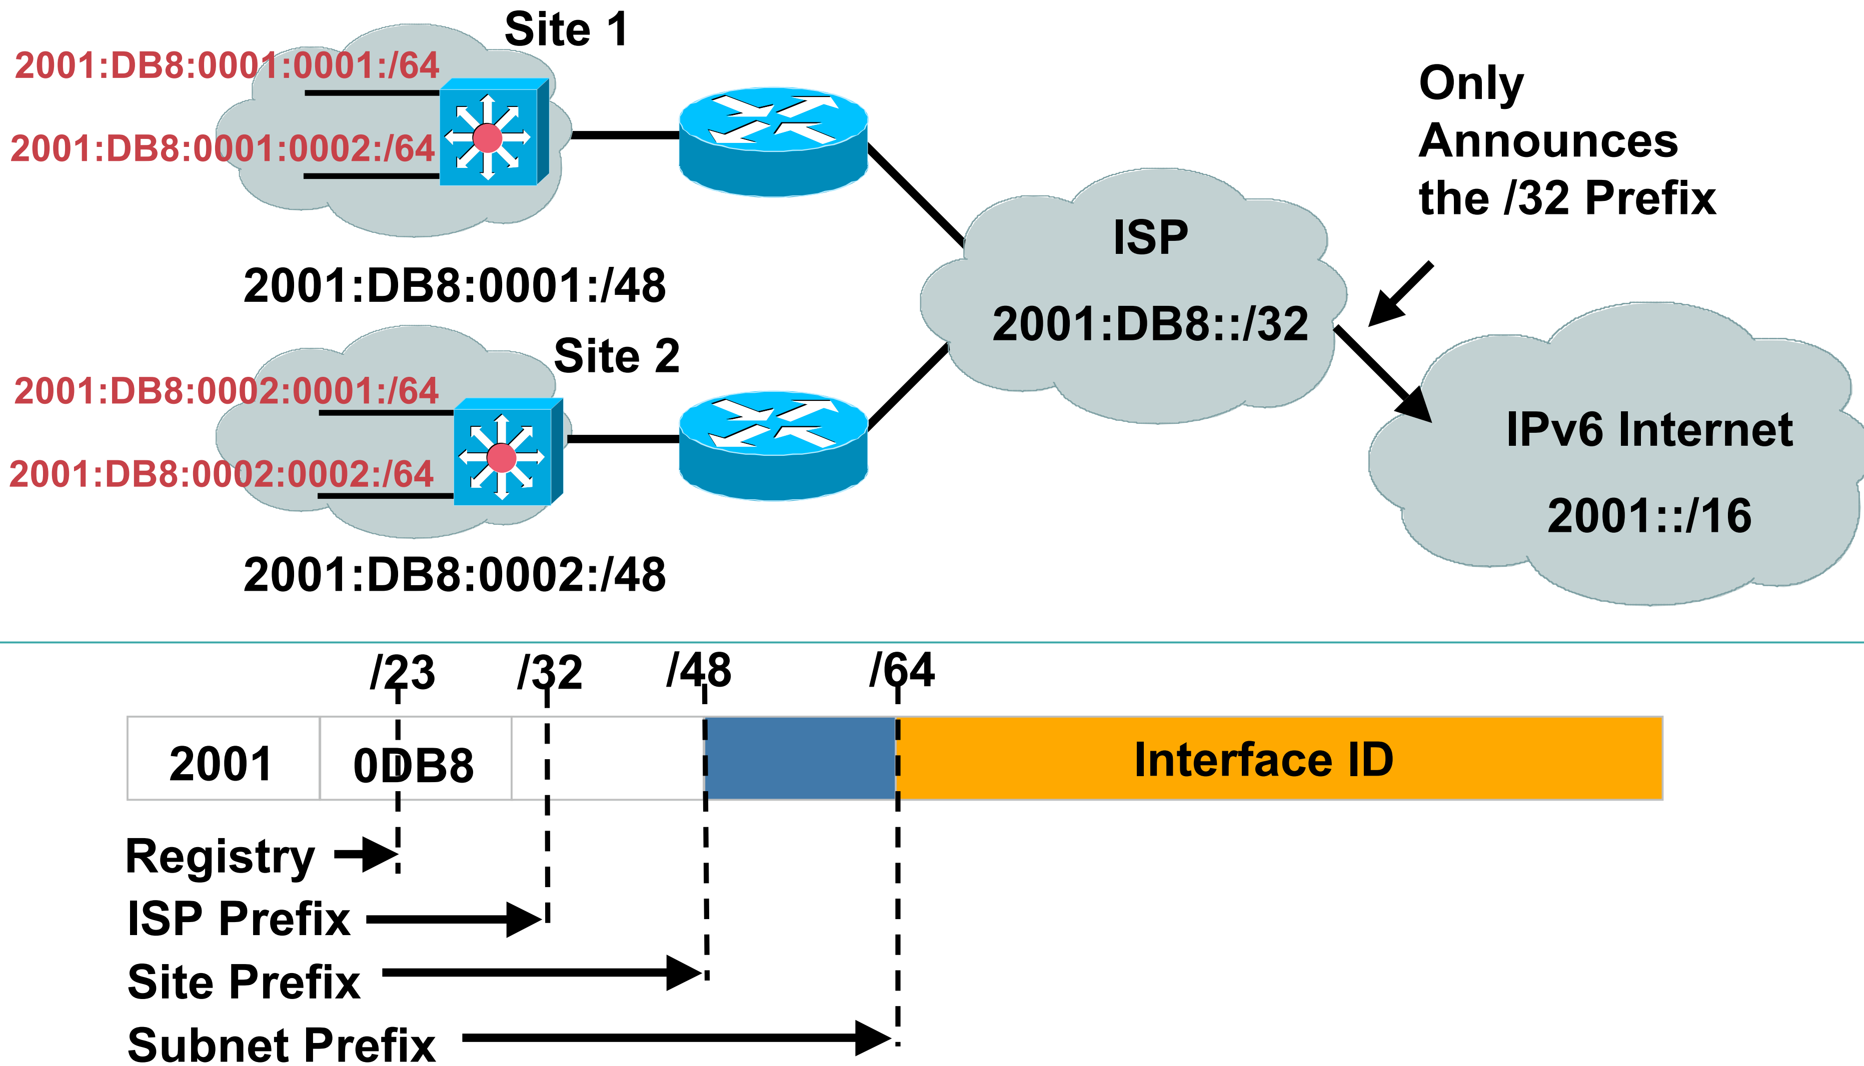

IPv6 & DNS

DNS Basics

- DNS is a database managing Resource Records (RR)

  - stockage of RR from various types - IPV4 and IPV6:

    -Start of Authority (SoA)

    -Name Server

    -Address - A and AAAA

    -Pointer - PTR

- DNS is an IP application

  -It uses either UDP or TCP on top of IPv4 or IPv6

- References

  RFC3596 : DNS Extensions to Support IP Version 6

  RFC3363 : Representing Internet Protocol Version 6 Addresses in Domain Name system (DNS)

  RFC3364: Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)

DNS Services

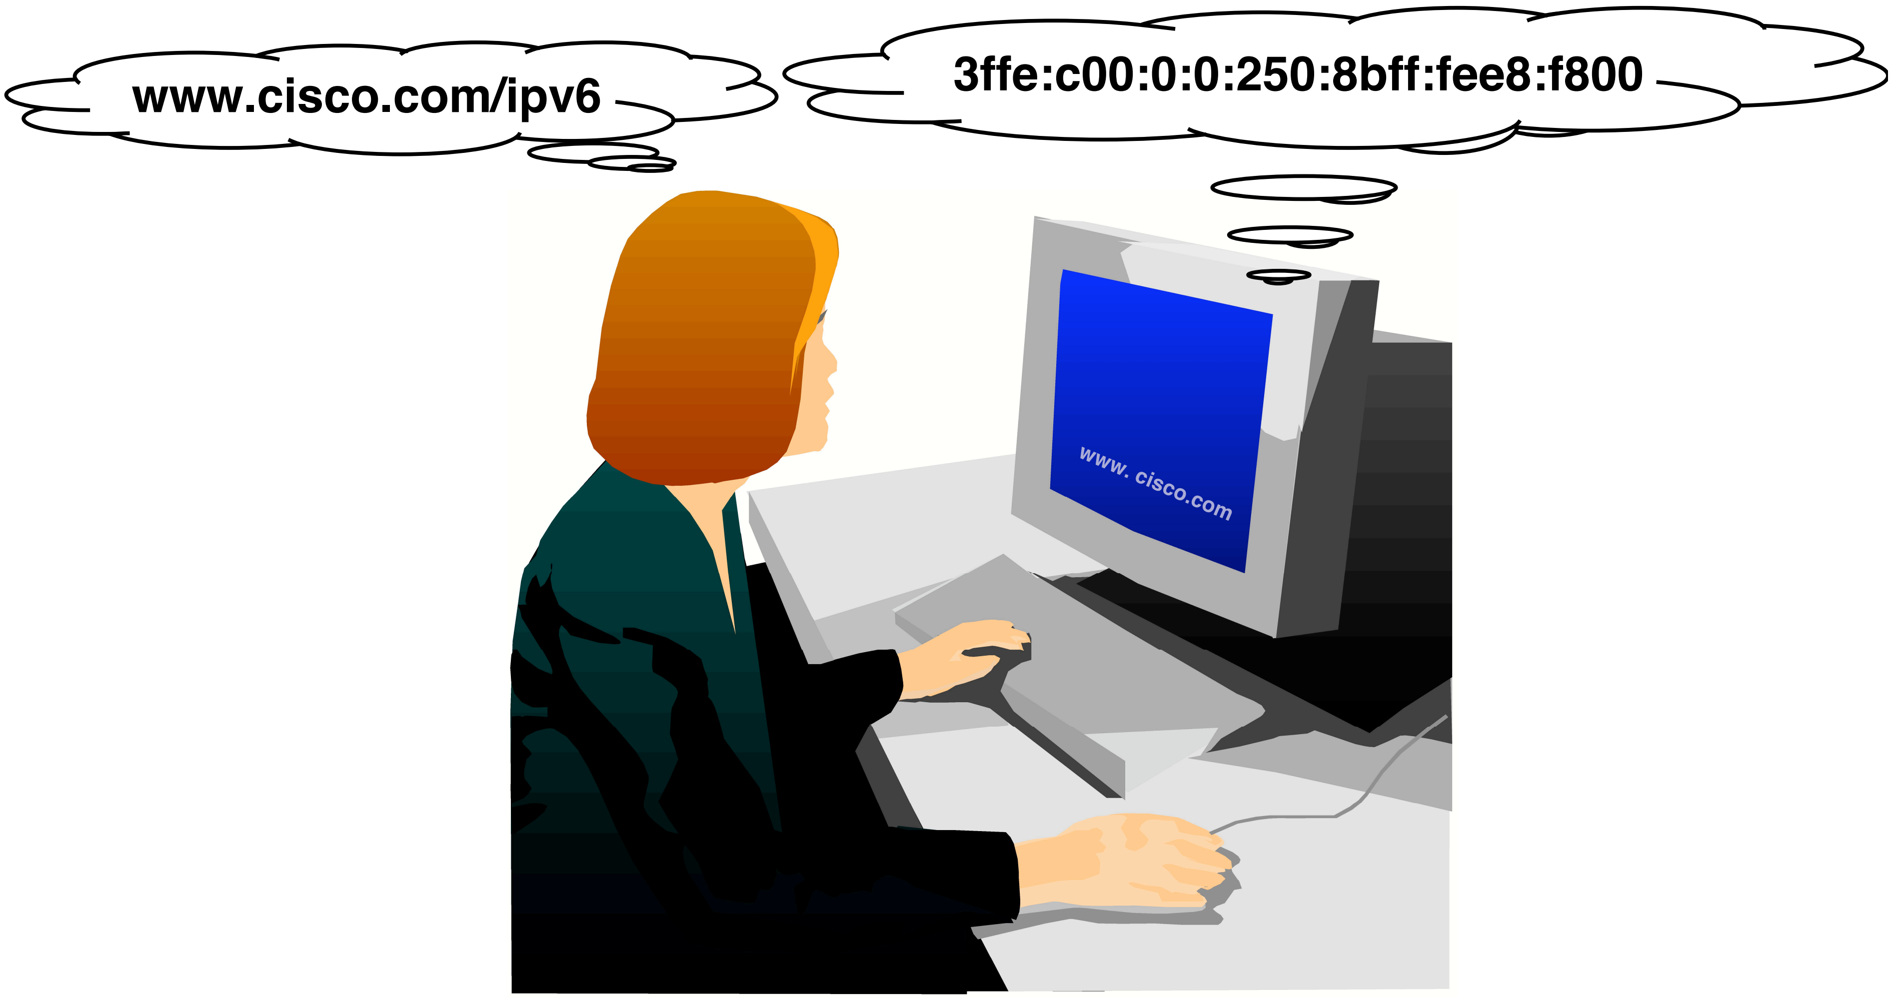

IPv6 and DNS

	IPv4	IPv6
Hostname to IP address	**A record:** www.abc.test. A 192.168.30.1	**AAAA record:** www.abc.test AAAA 3FFE:B00:C18:1::2
IP address to hostname	**PTR record:** 1.30.168.192.in-addr.arpa. PTR www.abc.test.	**PTR record:** 2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0. 0.0.b.0.e.f.f.3.ip6.arpa PTR www.abc.test.

Potential IPv4/IPv6 DNS issue (1)

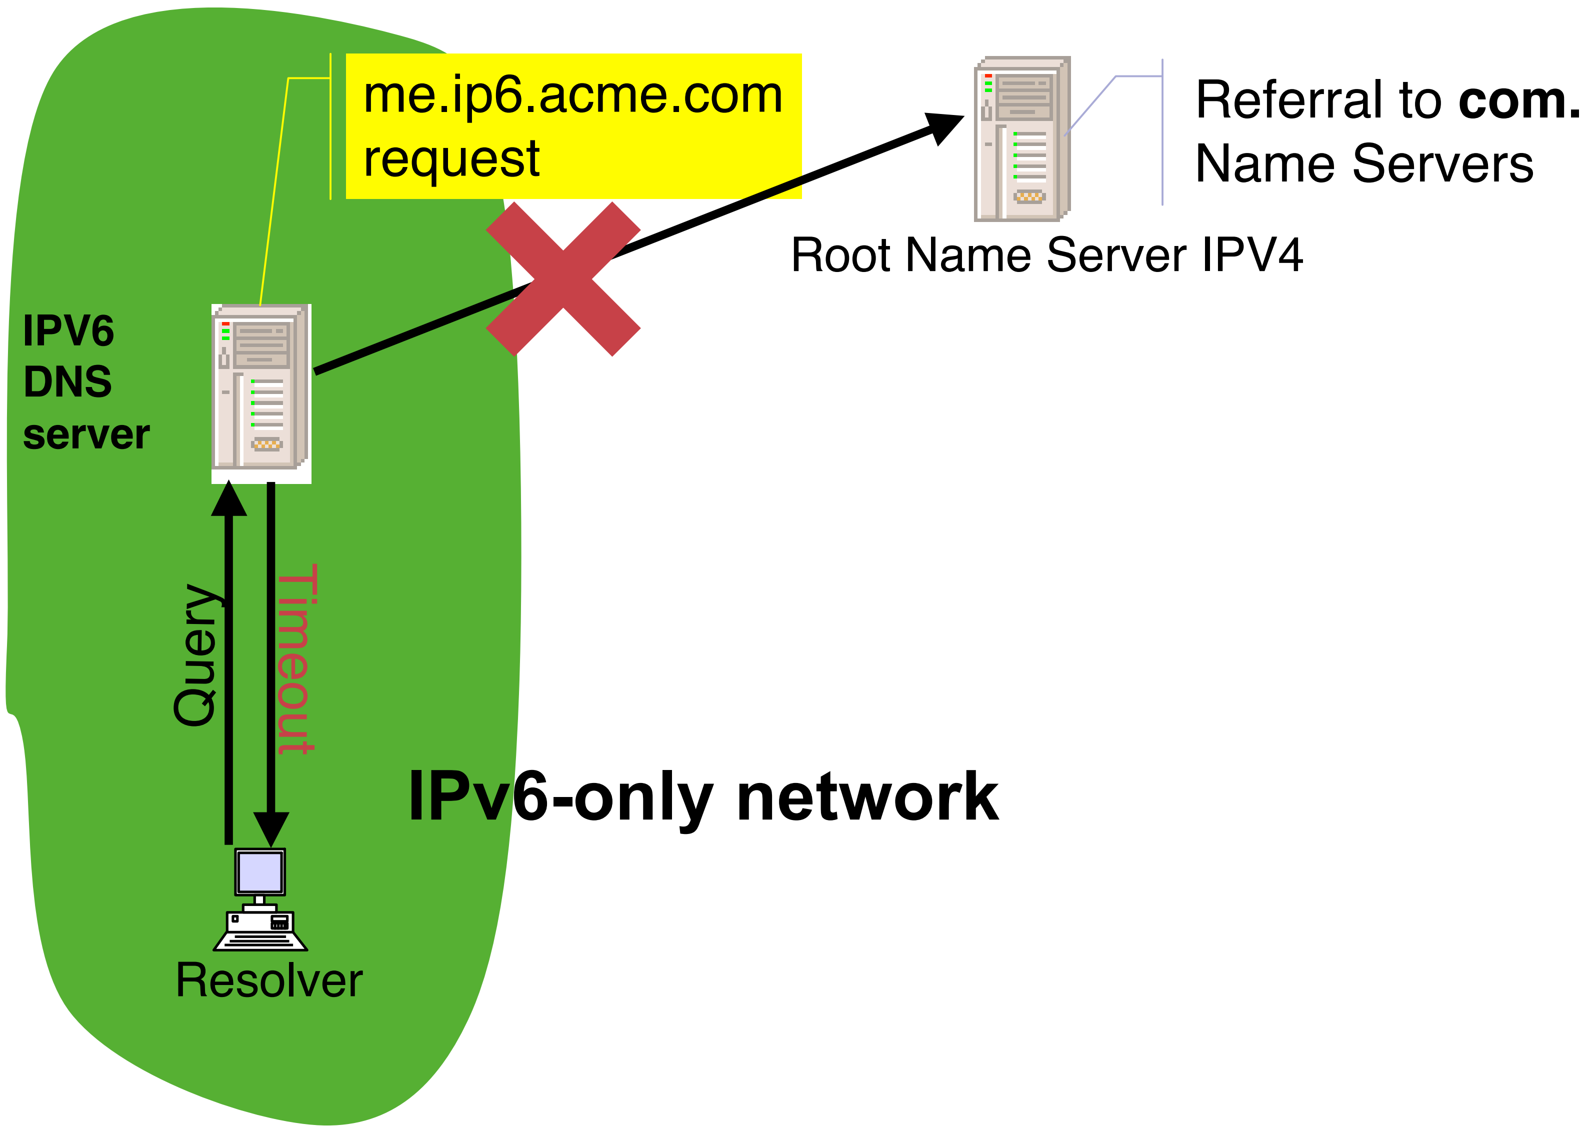

Potential IPv4/IPv6 DNS issue (2)

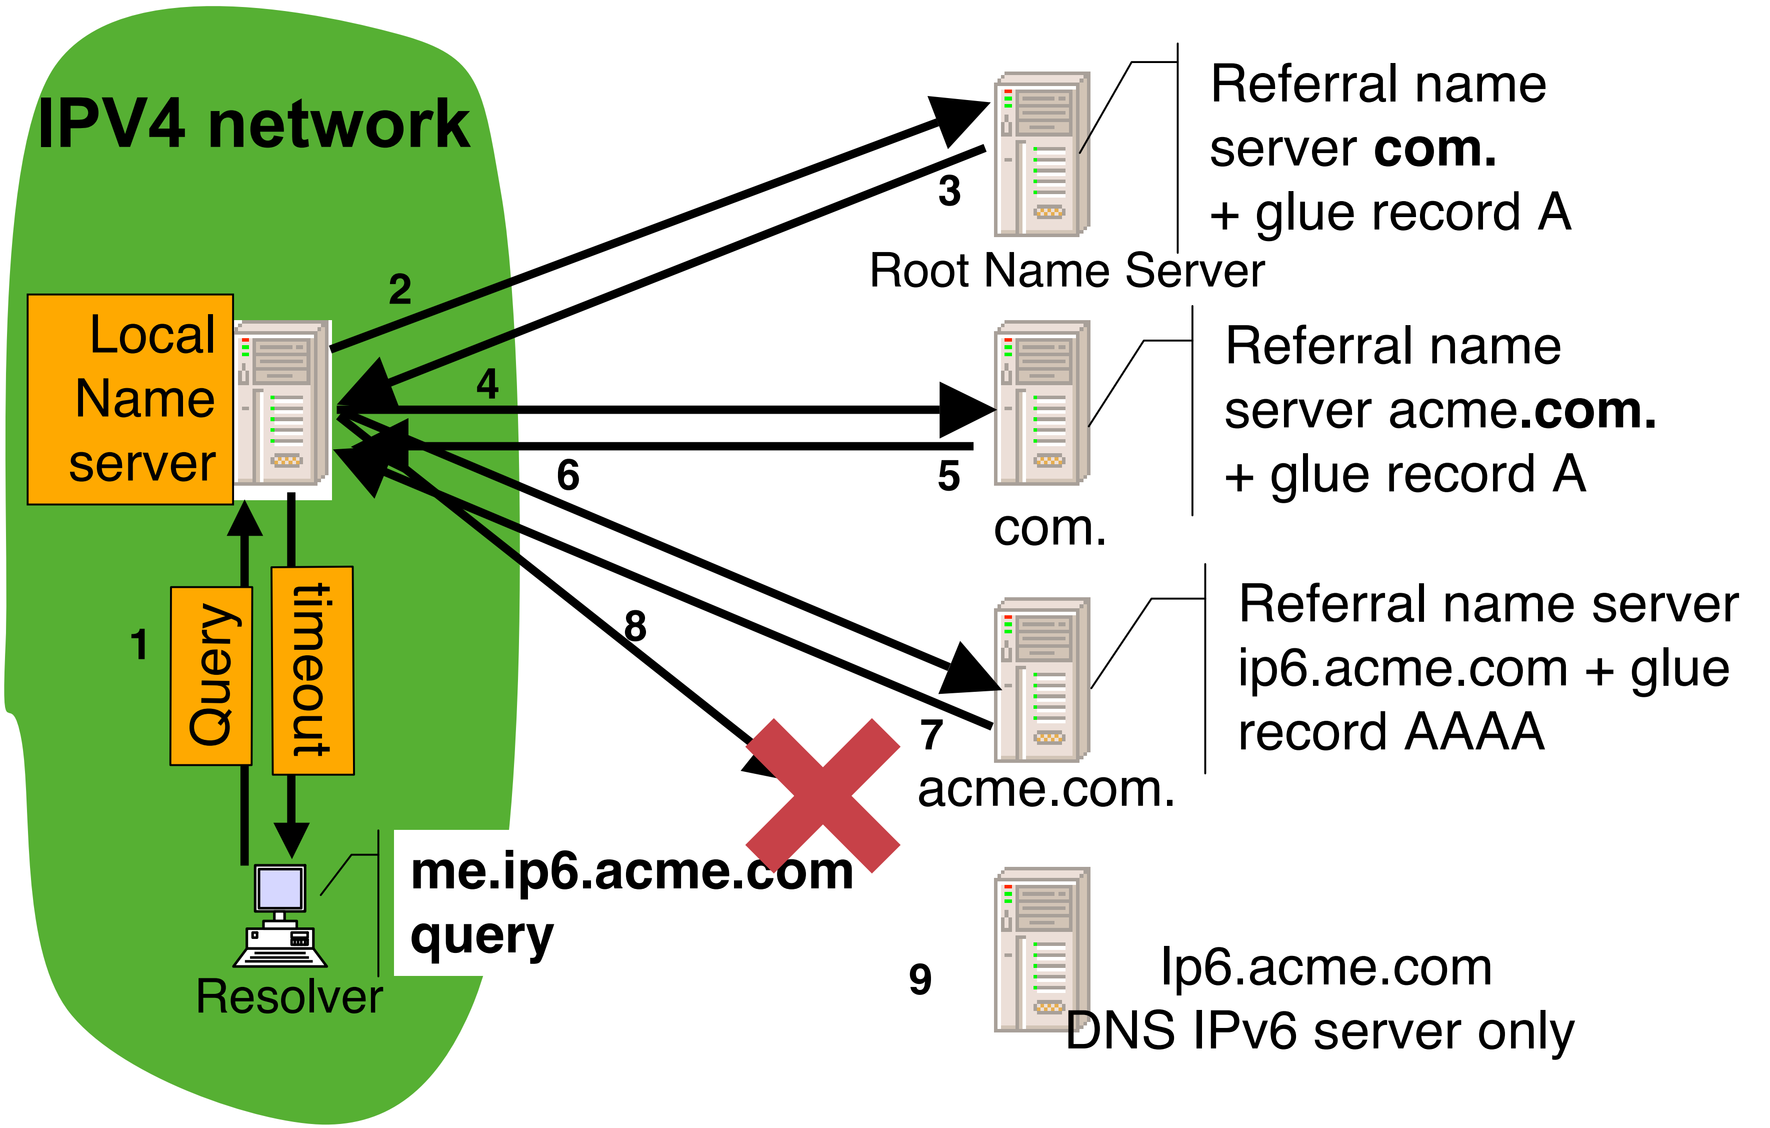

IPv6 Network Management

Network Management Differentiation

Instrumentation

- **MIB's, Netflow records which gives statistics about the IPv6 traffic.**

Transport

- **You can certainly do SNMP, SysLog over IPv6 but as you still have to manage IPv4, it may increase the complexity for operations.**

Applications

- **Products such as CiscoWorks LMS 2.5, CNR 6.2 do support IPv6 and offer specific features such as topology mapping, user tracking, address management and etc.**

SNMP and IPv6

- **MIBs:**

  First rewritten as separate IPv6 MIBs

  RFC 3291 defines representations of addresses in MIBs: IPv4, IPv6, DNS

  Current versions extend original MIBs for new address forms:

IP:	**RFC 2011 - draft-ietf-ipv6-rfc2011-update (RFC Editor publication queue)**
TCP:	**RFC 2012 - RFC 4022**
UDP:	**RFC 2013 - RFC 4113**
IP Forwarding:	**RFC 2096 - draft-ietf-ipv6-rfc2096-update (RFC Editor publication queue)**
BGP:	**draft-ietf-idr-bgp4-mibv2-05.txt**

Tools for SNMPv6

- HP OpenView
- CiscoWorks
  - Basic support over IPv4 today: IPv6 addresses, basic MIBs, configuration
  - Subsequent phases will use IPv6 transport and provide additional functions
- NetFlow Collector v5
- Other tools (from 6NET D6.2.4):
  - Argus, Cricket – network monitors
  - IPv6 Lan Dynamic Topology Discovery
  - IPv6 Management Gateway – manage IPv6 nodes with IPv4 management platform
  - net-snmp
  - network weathermap
  - See http://www.6net.org/publications/deliverables/#wp6 for details

How to manage an IPv6 network ?

- Evolution follows the IPv6 deployment
  1. Integration of IPv6 in existing network
  2. Parity between v4/v6 - Dual stack IPv6 networks
  3. IPv6 only

     This is not yet the main case …

     Important to think / know IPv4 could be removed

IPv6 Routing Protocols

Routing in IPv6

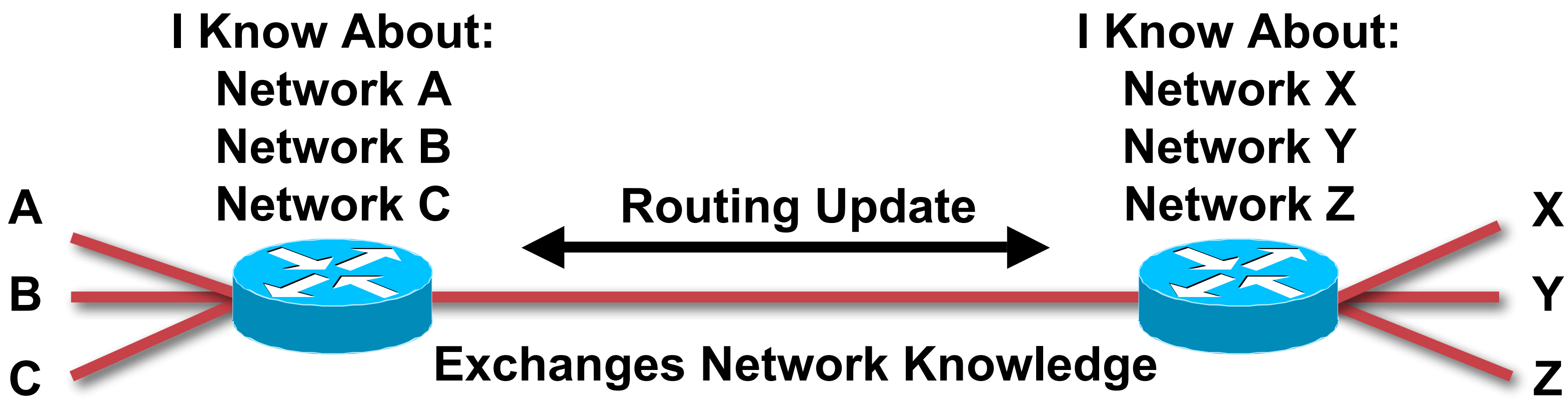

Routing protocols still:

- **Exchange NLRI**
- **Optimal path selection**
- **Loop-free routing**
- **Longest-prefix match routing algorithm.**

*****Additional memory may be used to maintain two route tables…***

Routing in IPv6

- As in IPv4, IPv6 has 2 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-ipv6-02)
    - OSPFv3 (RFC 2740)
- EGP : MP-BGP4 (RFC 2858 and RFC 2545)
- Cisco IOS supports all of them
    - Pick one that meets your objectives

Default Routing Example

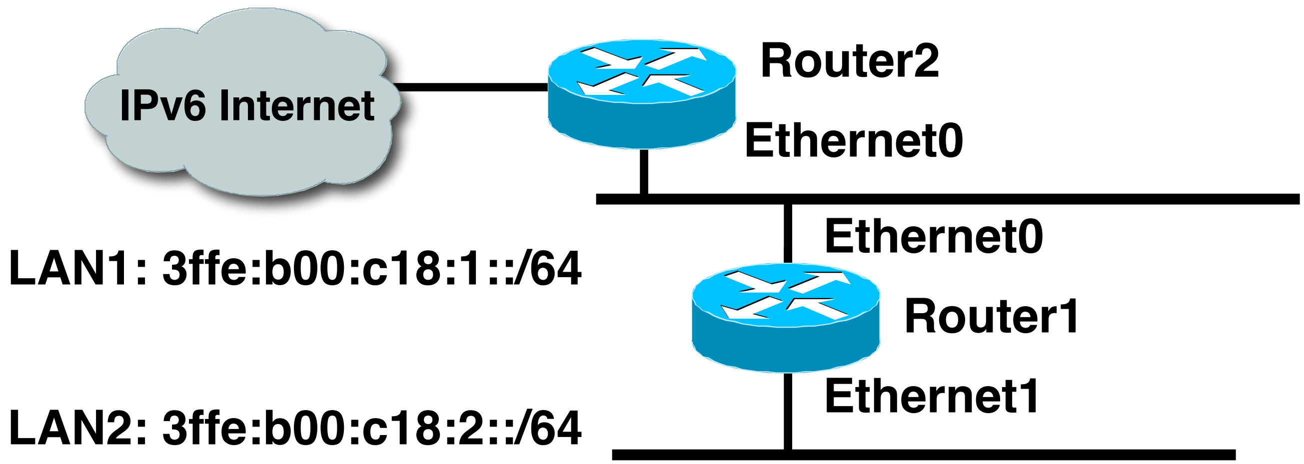


```
Router1#
ipv6 unicast-routing

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
 ipv6 nd prefix-advertisement 3ffe:b00:c18:1::/64
43200 43200 onlink autoconfig

interface Ethernet1
 ipv6 address 3ffe:b00:c18:2::/64 eui-64
 ipv6 nd prefix-advertisement 3ffe:b00:c18:2::/64
43200 43200 onlink autoconfig

ipv6 route ::/0 3ffe:b00:c18:1:260:3eff:fe47:1530
```

Default route to Router2 E0

IPv6 Router Configuration

```
PMO_7200-1#wr t
:
interface Loopback3
 no ip address
 ipv6 address 3FFE:1100:0:CC00::1/64
 ipv6 enable
!
interface POS4/0
 no ip address
 ipv6 address 2001:420:1921:6801::/64 eui-64
 ipv6 enable
 ipv6 rip 7206-1 enable
 clock source internal
!
ipv6 router rip 7206-1
```

Show IPv6 Interface Command

```
PMO_7200-1#show ipv6 interface pos4/0
POS4/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::230:96FF:FE07:F000
  Global unicast address(es):
    2001:420:1921:6801:230:96FF:FE07:F000, subnet is 2001:420:1921:6801::/64
  Joined group address(es):
    FF02::1 ← Link Local All Nodes Mcast
    FF02::2 ← Link Local All Routers Mcast
    FF02::9 ← RIPng Mcast
    FF02::1:FF07:F000 ← Link Local Solicited Node Mcast (Remember DAD ☺)
  MTU is 4470 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
  Hosts use stateless autoconfig for addresses.
```

IPv6 Routing Table

```
PMO_7200-1#show ipv6 route
IPv6 Routing Table - 8 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
R   2001:410:1921:6801::/64 [120/3]
     via FE80::202:7EFF:FE37:1CFF, POS4/0
C   2001:420:1921:6801::/64 [0/0]
     via ::, POS4/0
L   2001:420:1921:6801:230:96FF:FE07:F000/128 [0/0]
     via ::, POS4/0
R   3FFE:B00:C18:1::/64 [120/2]
     via FE80::202:7EFF:FE37:1CFF, POS4/0
C   3FFE:1100:0:CC00::/64 [0/0]
     via ::, Loopback3
L   3FFE:1100:0:CC00::1/128 [0/0]
     via ::, Loopback3
L   FE80::/10 [0/0]
     via ::, Null0
L   FF00::/8 [0/0]
     via ::, Null0
```

RIPng
(RFC 2080)

Enhanced Routing Protocol Support
RIPng Overview

- RIPng for IPv6, RFC 2080
- Same as IPv4:
  - Distance-vector, radius of 15 hops, split-horizon & etc.
  - 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

Enhanced Routing Protocol Support RIPng Configuration and Display

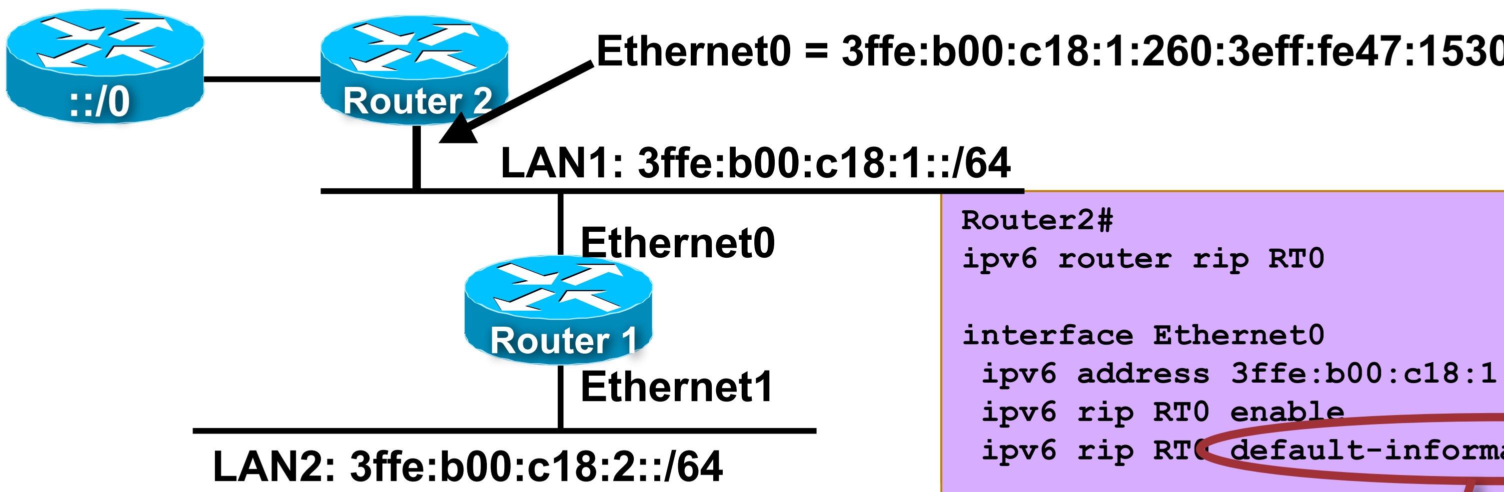


```
Router2#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 ipv6 rip RT0 default-information originate
```

```
Router1#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 Interface Ethernet1
 ipv6 address 3ffe:b00: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

I/IS-IS for IPv6

Enhanced Routing Protocol Support Integrated IS-IS for IPv6 Overview

- **2 Tag/Length/Values added to introduce IPv6 routing**
- **IPv6 Reachability TLV (0xEC)**

  Describes network reachability such as IPv6 routing prefix, metric information and some option bits. The option bits indicates the advertisement of IPv6 prefix from a higher level, redistribution from other routing protocols.

  Equivalent to IP Internal/External Reachability TLV's described in RFC1195

- **IPv6 Interface Address TLV (0xE8)**

  Contains 128 bit address

  For Hello PDUs, must contain the link-local address (FE80::/10)

  For LSP, must only contain the non link-local address

- **A new Network Layer Protocol Identifier (NLPID) is defined**

  Allowing IS-IS routers with IPv6 support to advertise IPv6 prefix payload using 0x8E value (IPv4 & OSI uses different values)

Enhanced Routing Protocol Support
I/IS-IS for IPv6-Only Configuration Example

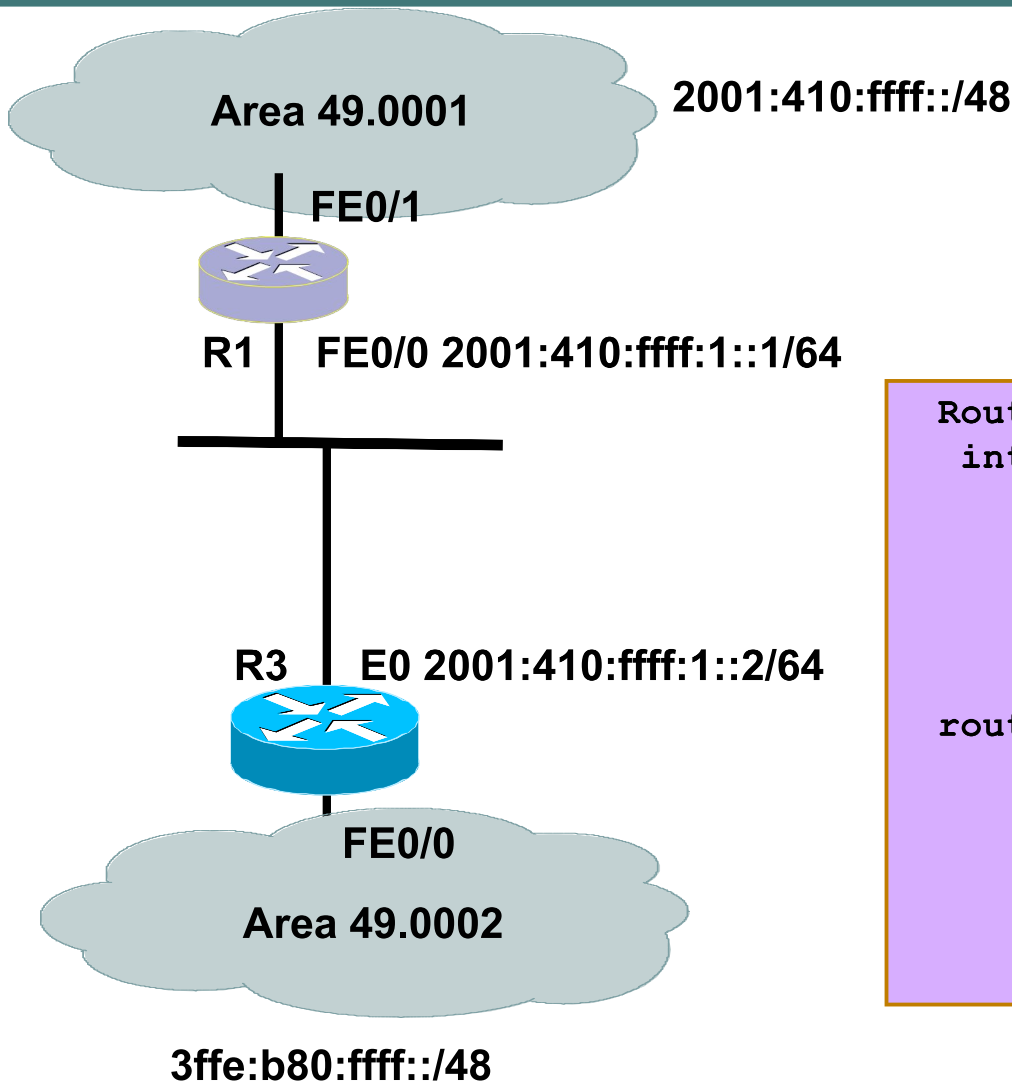


```
Router1#
 interface fastethernet0/0
     ipv6 address 2001:410:ffff:1::1/64
     ipv6 router isis
     isis circuit-type level-2-only


router isis
      address-family ipv6
      redistribute static
      exit-address-family
      net 49.0001.1921.6801.0001.00
```

Enhanced Routing Protocol Support Cisco IOS I/IS-IS Dual IP Configuration

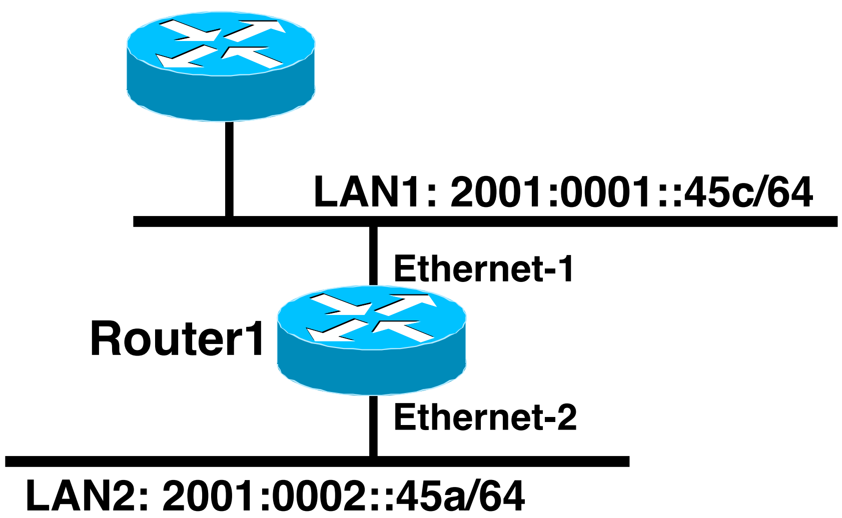


```
Router1#
 interface ethernet-1
     ip address 10.1.1.1 255.255.255.0
     ipv6 address 2001:0001::45c/64
     ip router isis
     ipv6 router isis

  interface ethernet-2
     ip address 10.2.1.1 255.255.255.0
     ipv6 address 2001:0002::45a/64
     ip router isis
     ipv6 router isis

  router isis
      address-family ipv6
      redistribute static
      exit-address-family
      net 49.0001.1921.6801.0001.00
      redistribute static
```

Dual IPv4/IPv6 Configuration Redistributing Both IPv6 Static Routes and IPv4 Static Routes

Enhanced Routing Protocol Support
Cisco IOS I/IS-IS Display

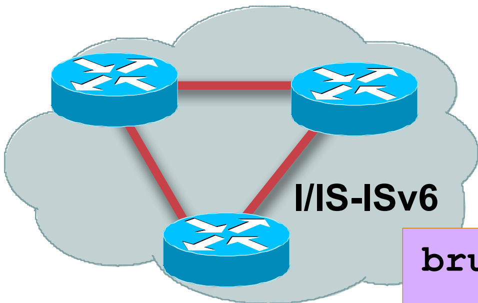


```
brum-45c#sho ipv6 rou is-is
IPv6 Routing Table - 14 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
Timers: Uptime/Expires

I1  2001:45A:1000::/64 [115/20]
     via FE80::210:7BFF:FEC2:ACCC, Ethernet1, 00:10:12/never
I1  2001:72B:2000::/64 [115/10]
     via FE80::210:7BFF:FEC2:ACCC, Ethernet1, 00:05:19/never
I1  2002:49::/64 [115/10]
     via FE80::210:7BFF:FEC2:ACCC, Ethernet1, 00:05:19/never
```

OSPFv3
(RFC 2740)

Enhanced Routing Protocol Support
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 such as hello, database description blocks (DDB), link state request (LSR), link state update (LSU) and link state advertisements (LSA)
- Neighbor discovery & adjacency formation mechanism are identical
- RFC compliant NBMA and point to multipoint topology modes are supported. Also supports other modes from Cisco such as point to point and broadcast
- LSA flooding and aging mechanisms are identical

Enhanced Routing Protocol Support Differences from OSPFv2

- **OSPF Packet Type**
- **OSPFv3 will have the same 5 packet type but some fields have been changed**

Packet type	Descrption
1	Hello
2	Database Description
3	Link State Request
4	Link State Update
5	Link State Acknowledgment

- **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

Enhanced Routing Protocol Support Differences from OSPFv2

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

  IPv6 connects interfaces to links.

  Multiple IP subnets can be assigned to a single link.

  Two nodes can talk directly over a single even they do not share and common subnet.

  The term "network" and "subnet" is being replaced with "link".

  An OSPF interface now connects to a link instead of a subnet.

- **Multiple OSPFv3 protocol instances can now run over a single link**

  This allows for separate ASes, each running OSPF, to use a common link. Single link could belong to multiple areas

  Instance ID is a new field that is used to have multiple OSPFv3 protocol instance per link.

  In order to have 2 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.

Enhanced Routing Protocol Support Differences from OSPFv2

- Uses link local addresses

  To identify the OSPFv3 adjacency neighbors

- Two New LSA Types

  Link-LSA (LSA Type 0x2008)

  There is one Link-LSA per link. This LSA advertises the router's link-local address, list of all IPv6 prefixes and options associated with the link to all other routers attached to the link

  Intra-Area-Prefix-LSA (LSA Type 0x2009)

  Carries all IPv6 prefix information that in IPv4 is included in Router-LSAs and Network-LSAs

- Two LSAs are Renamed

  Type-3 summary-LSAs, renamed to “Inter-Area-Prefix-LSAs”

  Type-4 summary LSAs, renamed to “Inter-Area-Router-LSAs”

Enhanced Routing Protocol Support
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, 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 & ESP extension headers instead of variety of mechanisms defined in OSPFv2

LSA Types

	LSA function code	LSA type
Router-LSA	1	0x2001
Network-LSA	2	0x2002
Inter-Area-Prefix-LSA	3	0x2003
Inter-Area-Router-LSA	4	0x2004
AS-External-LSA	5	0x4005
Group-membership-LSA	6	0x2006
Type-7-LSA	7	0x2007
Link-LSA	8	0x2008
Intra-Area-Prefix-LSA	9	0x2009

NEW

Enhanced Routing Protocol Support
OSPFv3 Configuration Example

Do it again ..

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

interface POS2/0
 ipv6 address 3FFE: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 3FFE:B00:FFFF:1::1/64
 ipv6 enable
 ipv6 ospf 100 area 1

ipv6 router ospf 100
   router-id 10.1.1.4
```

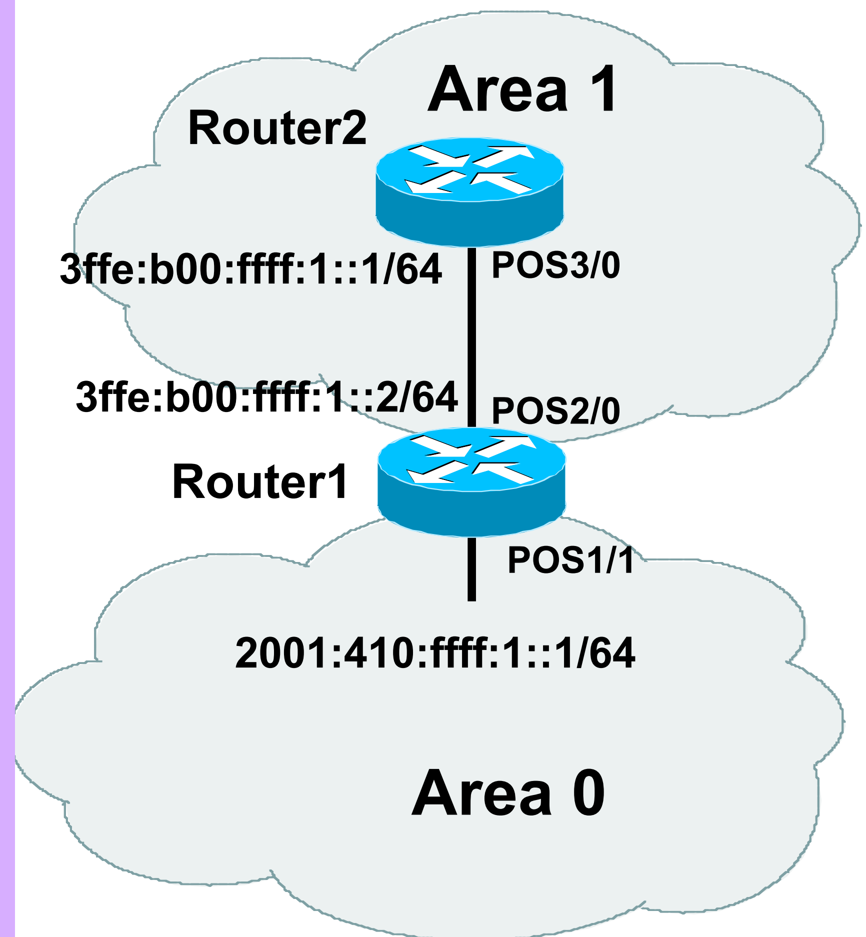

Enhanced Routing Protocol Support
Cisco IOS OSPFv3

```
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)
```

Enhanced Routing Protocol Support
Cisco IOS OSPFv3

```
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
```

Enhanced Routing Protocol Support
Cisco IOS OSPFv3

```
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   3FFE:B00:FFFF:1::/64 [0/0]
     via ::, POS3/0
L   3FFE: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 Extensions for IPv6

- BGP-4 carries only 3 pieces of information which is truly IPv4 specific:
  - NLRI in the UPDATE message contains an IPv4 prefix
  - NEXT_HOP path attribute in the UPDATE message contains a IPv4 address
  - BGP Identifier is in the OPEN message & AGGREGATOR attribute
- To make BGP-4 available for other network layer protocols, RFC 2858 (obsoletes RFC 2283) defines multi-protocol extensions for BGP-4
  - Enables BGP-4 to carry information of other protocols e.g MPLS,IPv6
  - New BGP-4 optional and non-transitive attributes:
    - MP_REACH_NLRI
    - MP_UNREACH_NLRI
  - Protocol independent NEXT_HOP attribute
  - Protocol independent NLRI attribute

BGP-4 Extensions for IPv6

- New optional and non-transitive BGP attributes:
  - MP_REACH_NLRI (Attribute code: 14)
    - “Carry the set of reachable destinations together with the next-hop information to be used for forwarding to these destinations” (RFC2858)
  - MP_UNREACH_NLRI (Attribute code: 15)
    - Carry the set of unreachable destinations
- Attribute 14 and 15 contains one or more Triples:
  - Address Family Information (AFI)
  - Next-Hop Information (must be of the same address family)
  - NLRI

BGP-4 Extensions for IPv6

- **Address Family Information (AFI) for IPv6**
  - AFI = 2 (RFC 1700)
  - Sub-AFI = 1 Unicast
  - Sub-AFI = 2 (Mulitcast for RPF check)
  - Sub-AFI = 3 for both Unicast and Mulitcast
  - Sub-AFI = 4 Label
  - Sub-AFI= 128 VPN

BGP-4 Extensions for IPv6

- Next-hop contains a global IPv6 address or potentially a link local (for iBGP update this has to be change to global IPv6 address with route-map)
- The value of the length of the next hop field on MP_REACH_NLRI attribute is set to 16 when only global is present and is set to 32 if link local is present as well
- Link local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised)

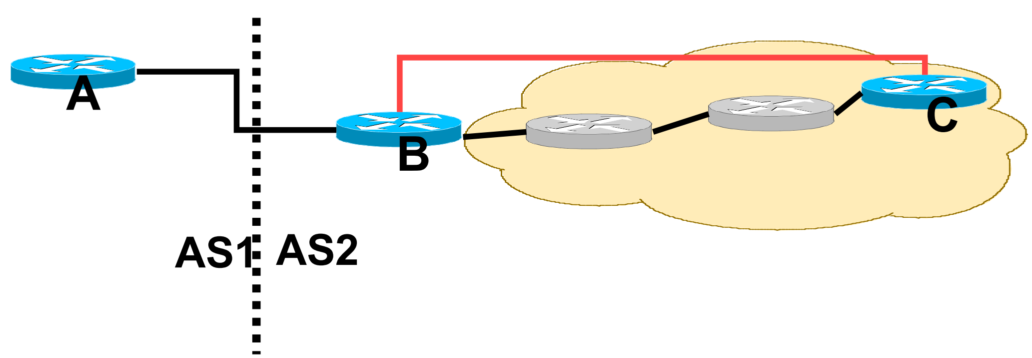

BGP-4 Extensions for IPv6

- **TCP Interaction**

  BGP-4 runs on top of TCP

  This connection could be setup either over IPv4 or IPv6

- **Router ID**

  When no IPv4 is configured, an explicit bgp router-id needs to be configured

  This is needed as a BGP Identifier, this is used as a tie breaker, and is send within the OPEN message

BGP-4 Configurations for IPv6 Non Link Local Peering

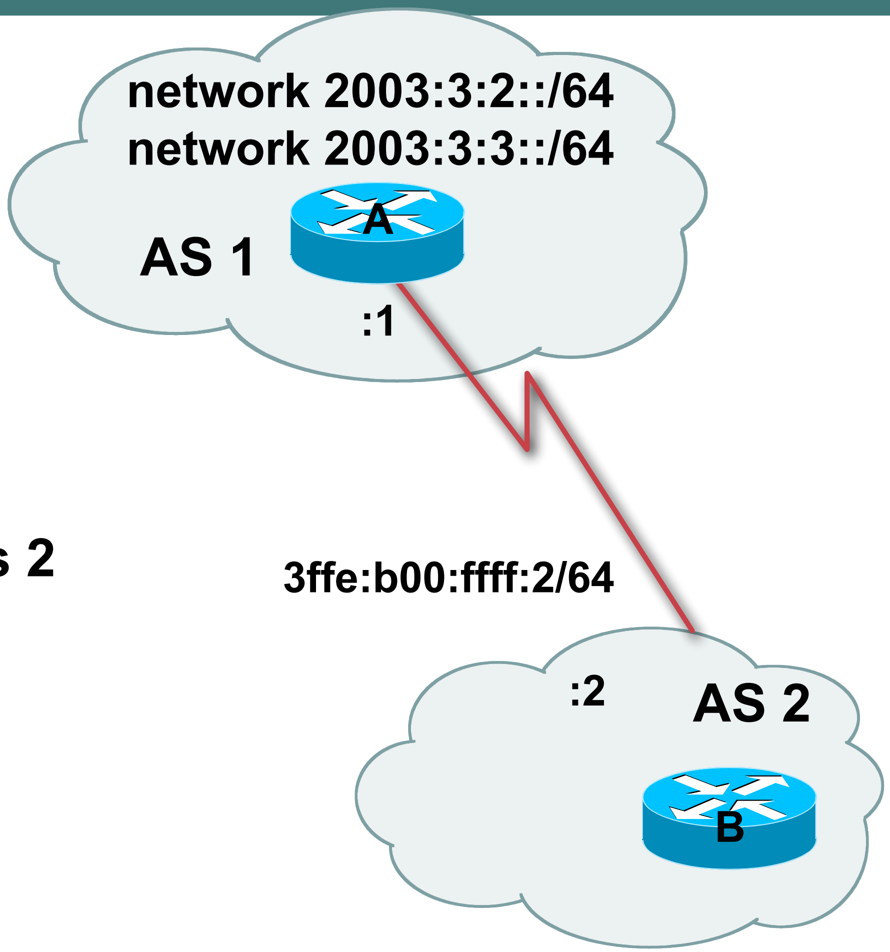

Router A

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

BGP-4 Configurations for IPv6 Link Local Peering

Router A

Interface e2
ipv6 address 2001:412:ffco:1::1/64

router bgp 1
no bgp default ipv4 unicast
bgp router-id 1.1.1.1
neighbor fe80::260:3eff:c043:1143 remote-as 2
neighbor fe80::260:3eff:c043:1143 update source e0
address-family ipv6
neighbor fe80::260:3eff:c043:1143 activate
neighbor fe80::260:3eff:c043:1143 route-map next-hop out

route-map next-hop
set ipv6 next-hop 2001:412:ffco:1::1

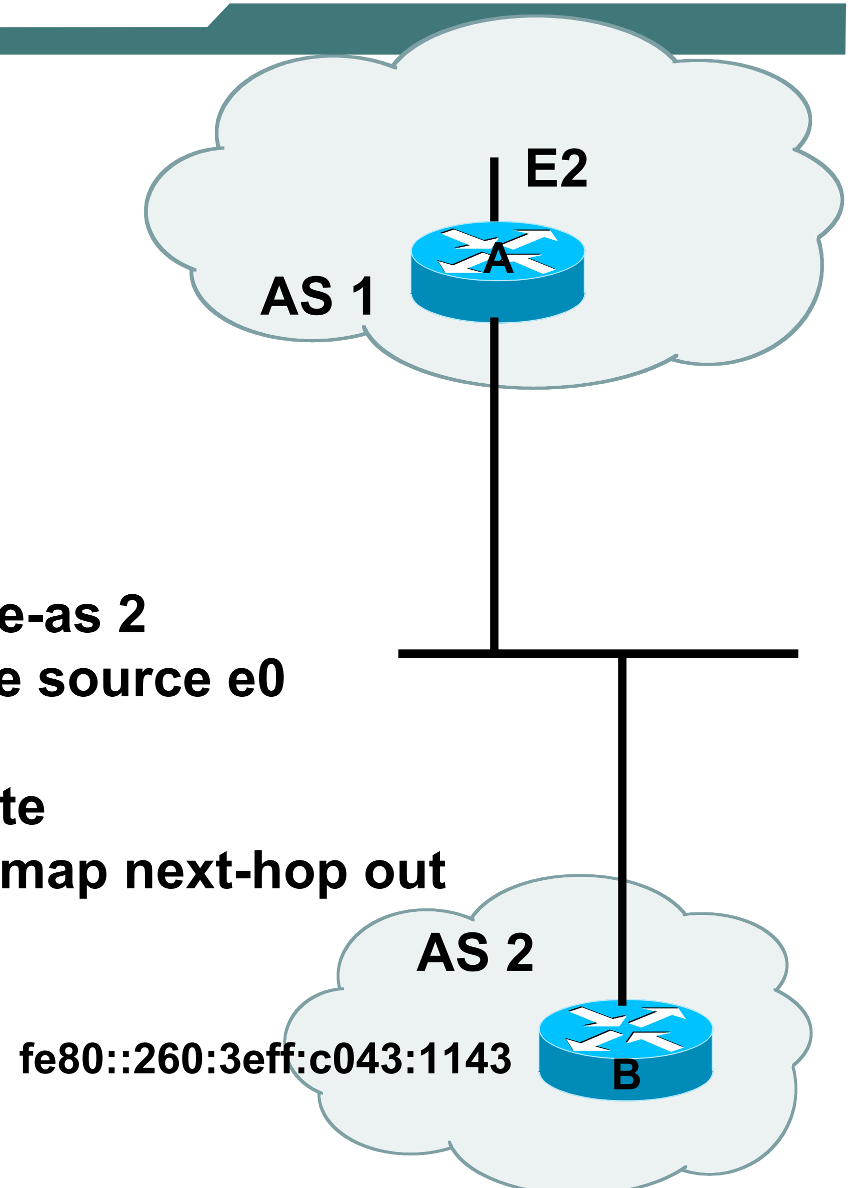

BGP Configurations

Carrying IPv4 inside IPv6 peering

```
router bgp 1
neighbor 3ffe:b00:ffff:2::2 remote-as 2
address-family ipv6
neighbor 3ffe:b00:ffff:2::2 activate
neighbor 3ffe:b00:ffff:2::2 route-map IPv4 in

route-map ipv4 permit 10
Set ip next-hop 131.108.1.1
```

BGP-4 for IPv6 « Show Command »

Show bgp ipv6 summary

Displays summary information regarding the state of the BGP neighbors

```
RouterA# show bgp ipv6 summary
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 69046, main routing table version 69046
92 network entries and 92 paths using 17756 bytes of memory
826 BGP path attribute entries using 43108 bytes of memory
703 BGP AS-PATH entries using 19328 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
745 BGP filter-list cache entries using 8940 bytes of memory
BGP activity 22978/18661 prefixes, 27166/22626 paths, scan interval 15 secs

Neighbor       V    AS   MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3FFE:B00:FFFF:2::2
               4    2    84194   14725    69044    0    0  3d08h          92
```

Neighbor Information

BGP Messages Activity

IPv6 Integration & Transition

Start Here: Cisco IOS Software Release Specifics for IPv6 Features
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/ftipv6s.htm

Transition & Integration Richness

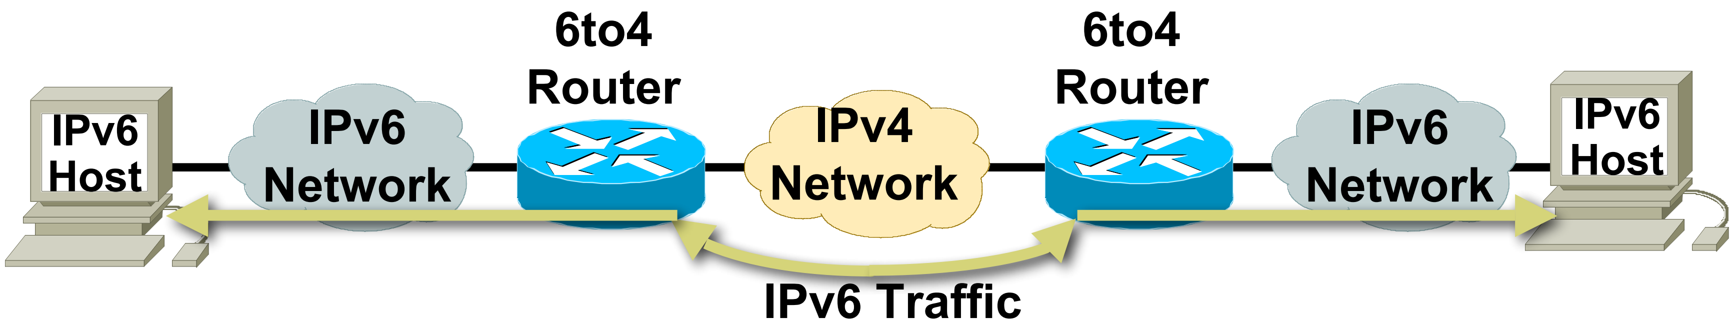

- **Transition richness means:**
  - **No fixed day to convert**
  - **No need to convert all at once**
  - **Different transition mechanisms are available**
  - **Smooth integration of IPv4 and IPv6**
  - **Different compatibility mechanisms**
  - **IPv4 and IPv6 nodes can talk together**

IPv4-IPv6 Transition / Co-Existence

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

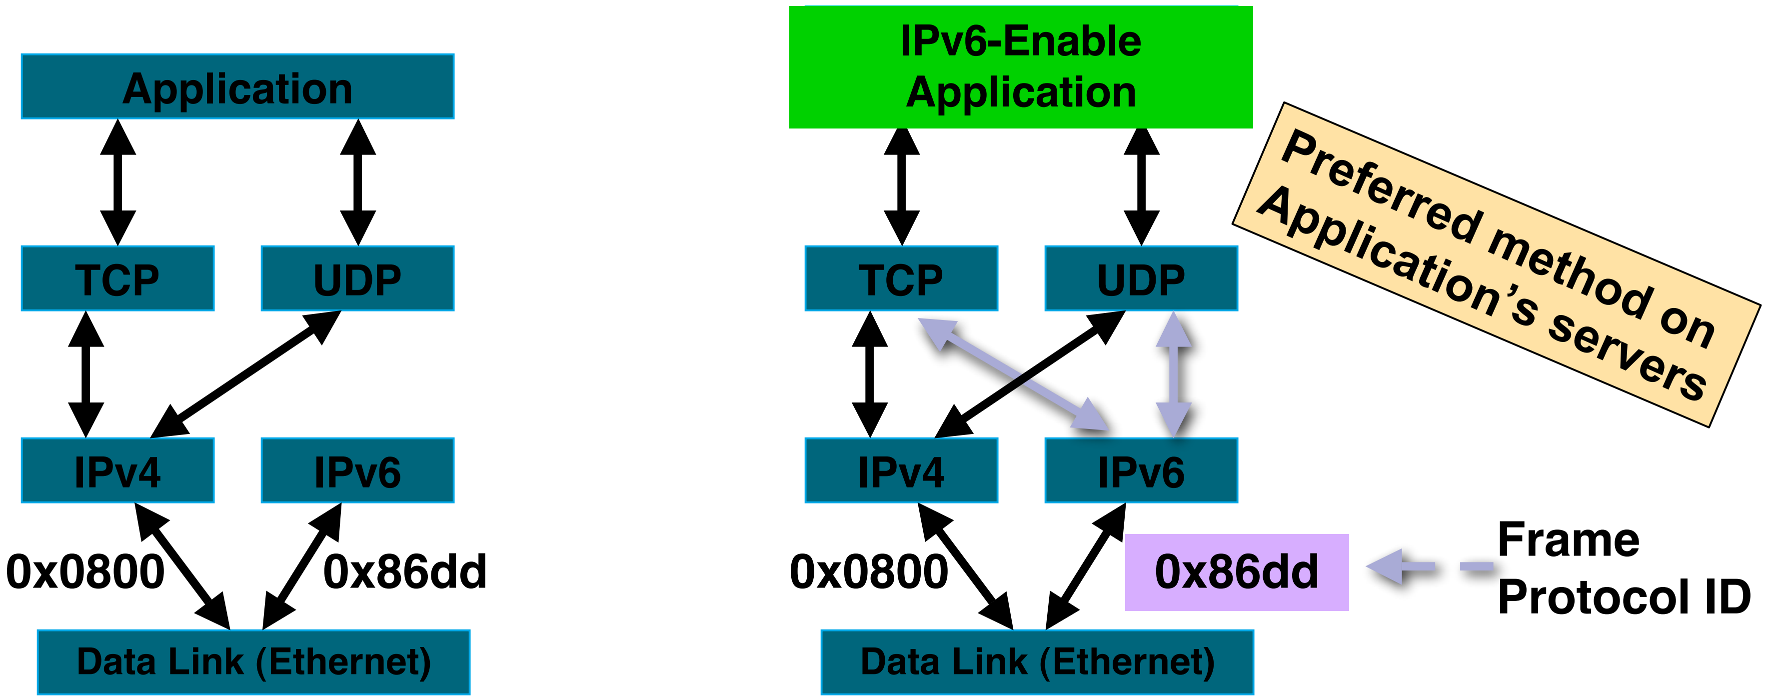

- Dual stack node means:
  - Both IPv4 and IPv6 stacks enabled
  - Applications can talk to both
  - Choice of the IP version is based on name lookup and application preference

Host Running IPv4 Stack

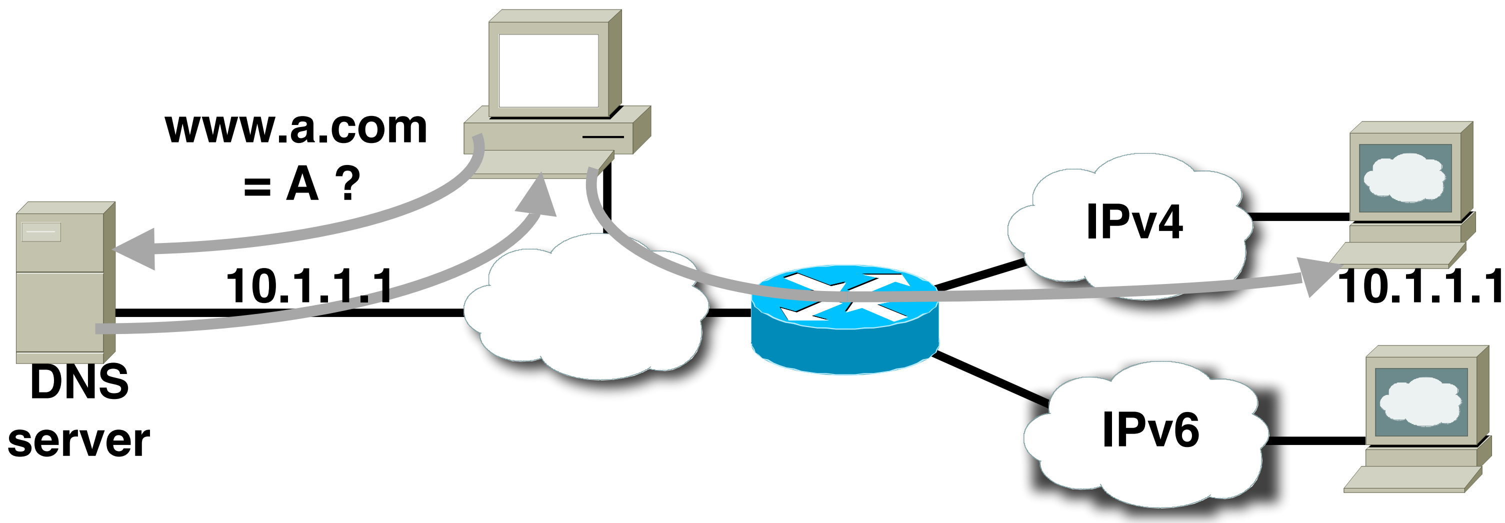

Without IPv6, an application:

Asks the DNS for the IPv4 address

And connects to the IPv4 address

Host Running IPv6 Stack

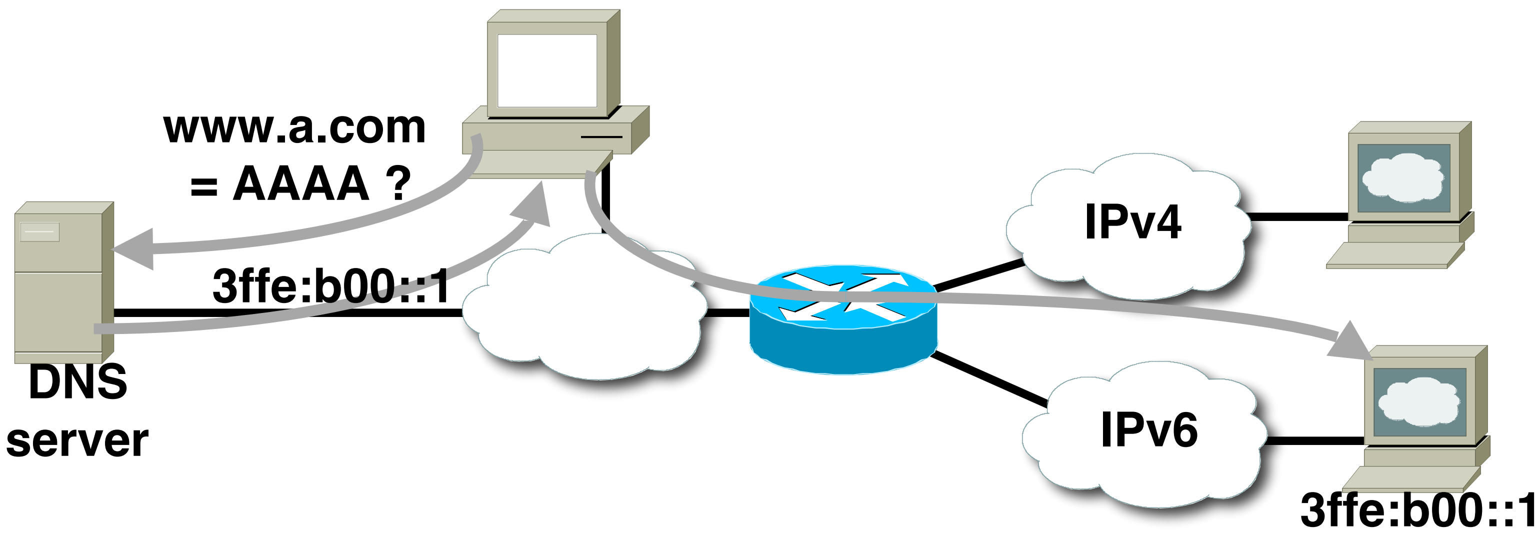

In an IPv6-only case, an application:

Asks the DNS for the IPv6 address

And then connects to the IPv6 address

Host Running Dual Stack

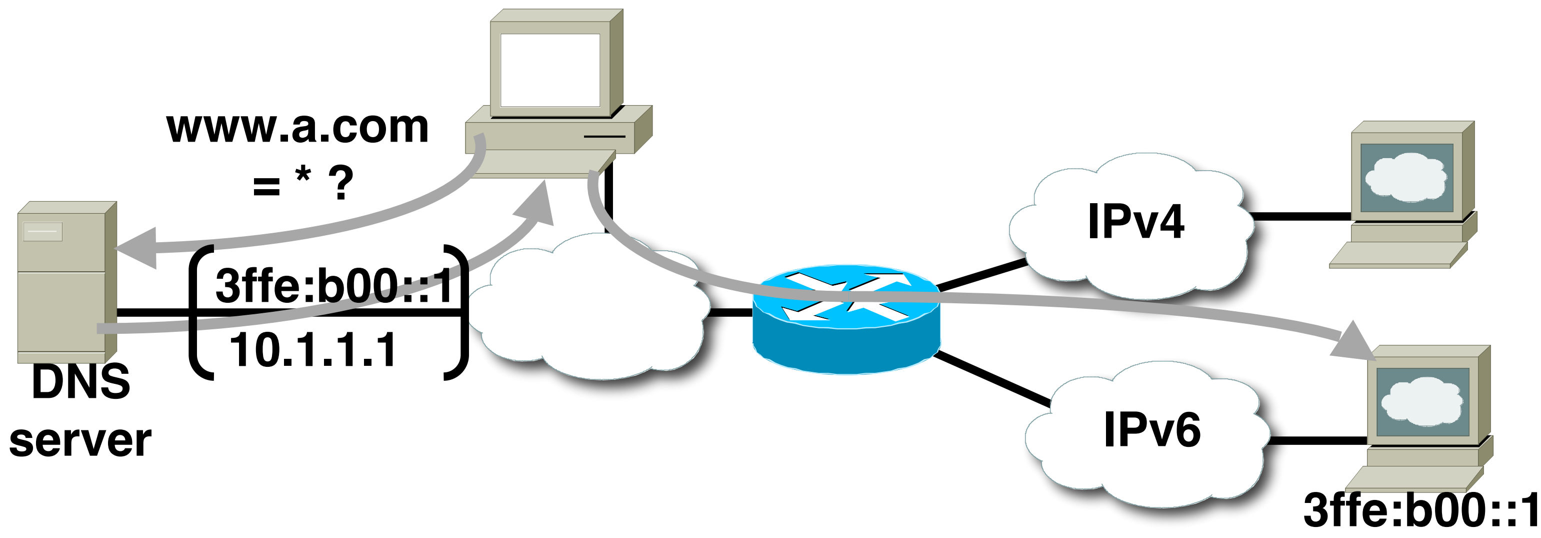

In a dual stack case, an application that:

- 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

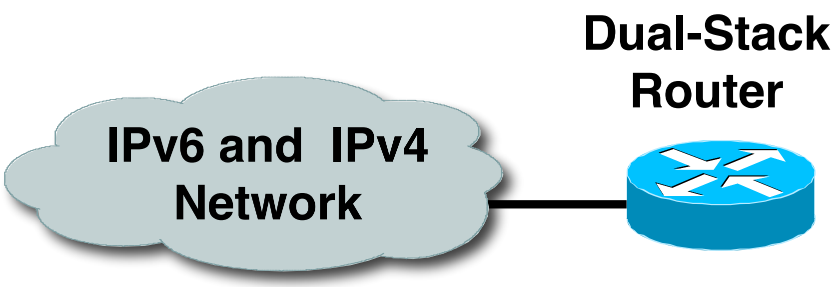


```
router#
ipv6 unicast-routing

interface Ethernet0
 ip address 192.168.99.1 255.255.255.0
 ipv6 address 2001:410:213:1::/64 eui-64
```

IPv4: 192.168.99.1

IPv6: 2001:410:213:1::/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,…

Using Tunnels for IPv6 Deployment

- Many techniques are available to establish a tunnel:
  - Manually Configured
    - Manual Tunnel (RFC 2893)
    - GRE (RFC 2473)
  - Semi-automated
    - Tunnel broker
  - Automatic
    - Compatible IPv4 (RFC 2893): Deprecated
    - 6over4: Deprecated
    - 6to4 (RFC 3056)
    - ISATAP

IPv6 over IPv4 Tunnels

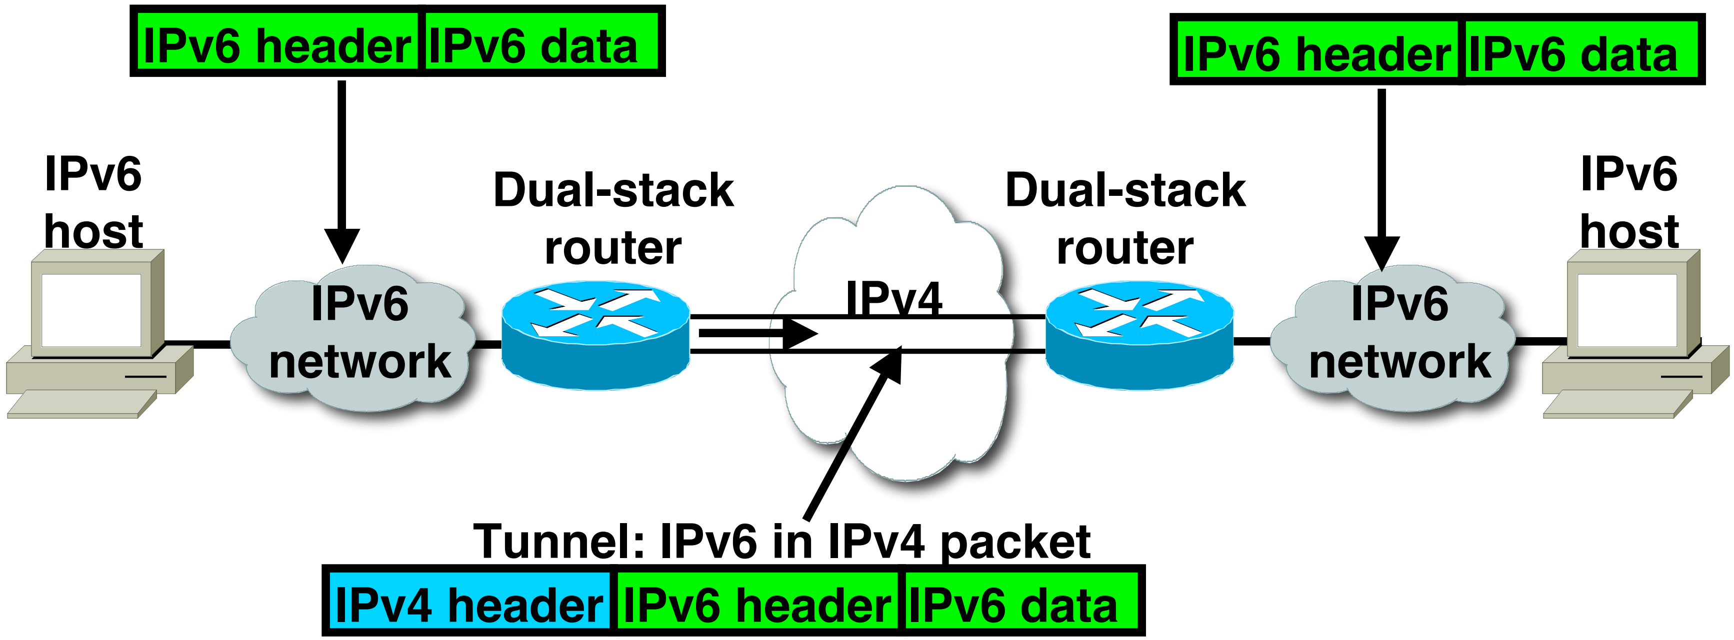

Tunneling is encapsulating the IPv6 packet in the IPv4 packet (IPv4 protocol type = 41).

IPv6 over IPv4 Tunnels

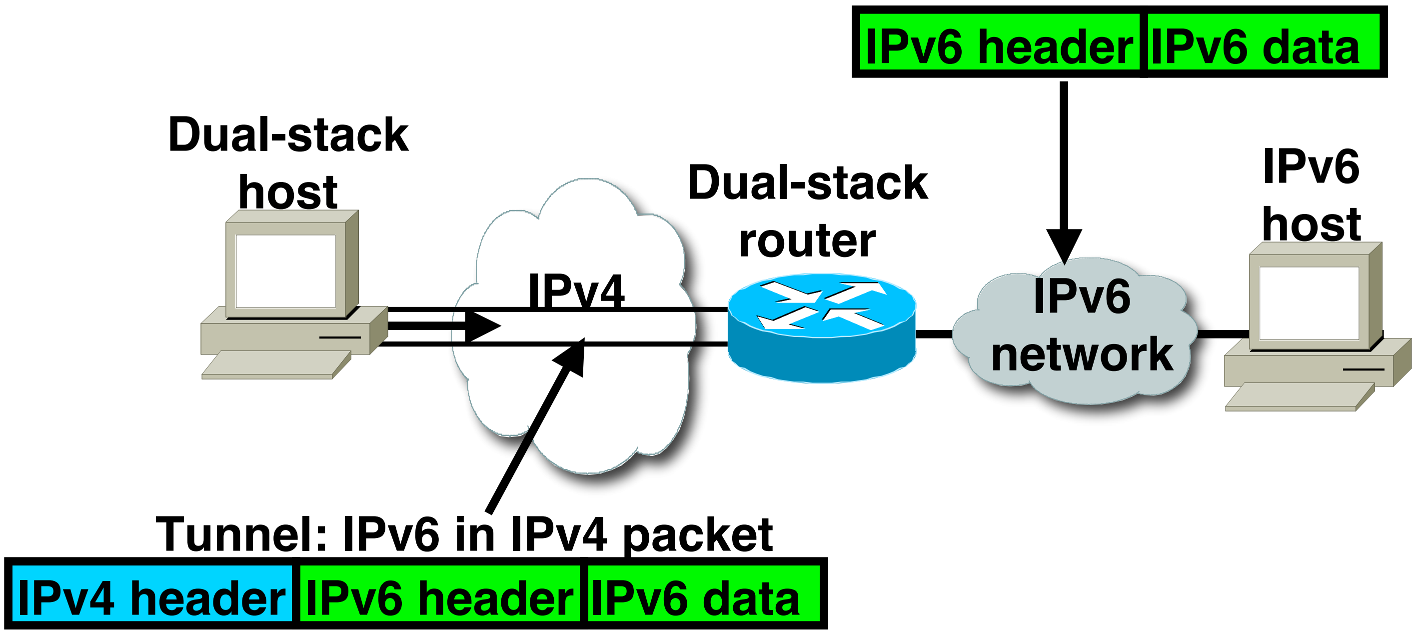

- Tunneling can be used by routers and hosts.

Manually Configured
Manual Tunnel (RFC 2893)

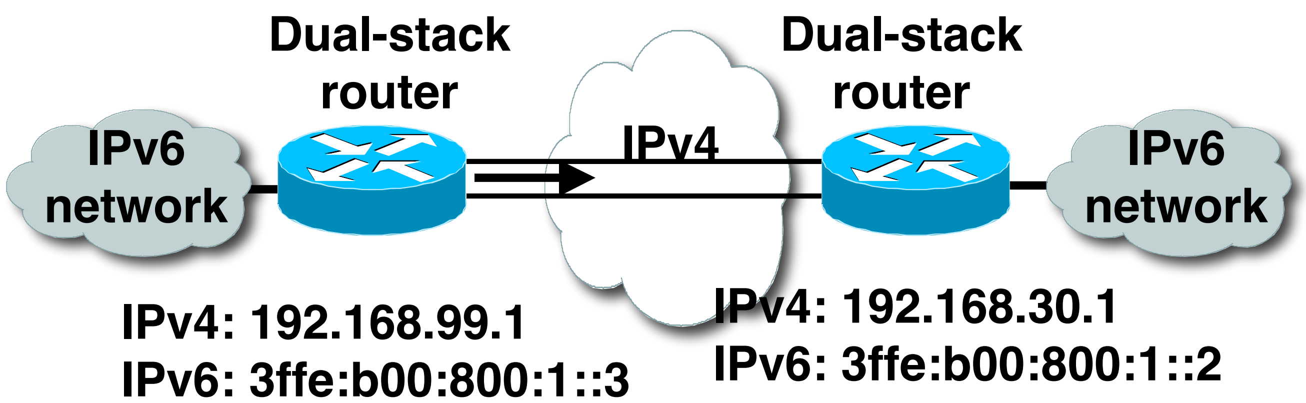

Manually configured tunnels require:

Dual stack end points.

Both IPv4 and IPv6 addresses configured at each end.

Manually Configured
Manual Tunnel Configuration

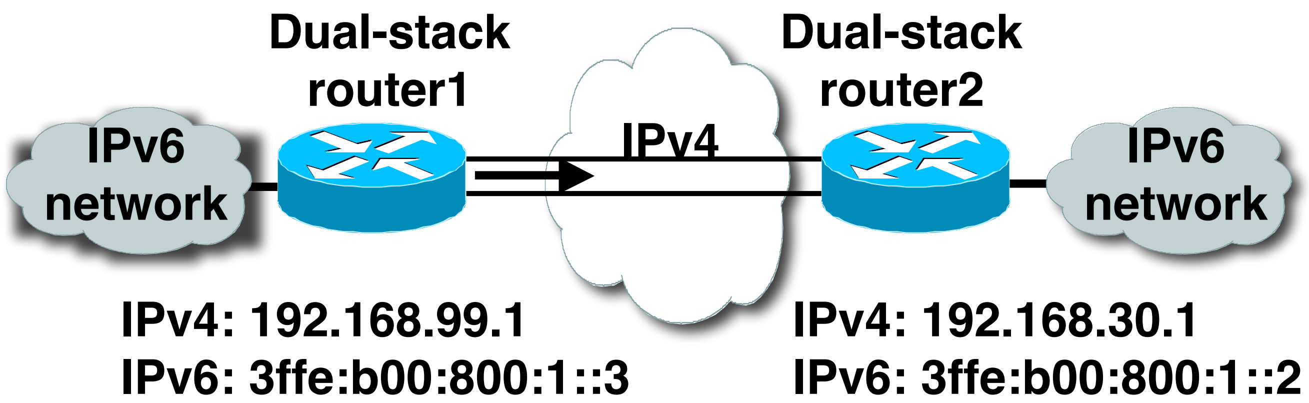


```
router1#

interface Tunnel0
 ipv6 enable
 ipv6 address 3ffe:b00: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 3ffe:b00:c18:1::2/127
 tunnel source 192.168.30.1
 tunnel destination 192.168.99.1
 tunnel mode ipv6ip
```

Manually Configured IPv6 Over GRE Tunnel

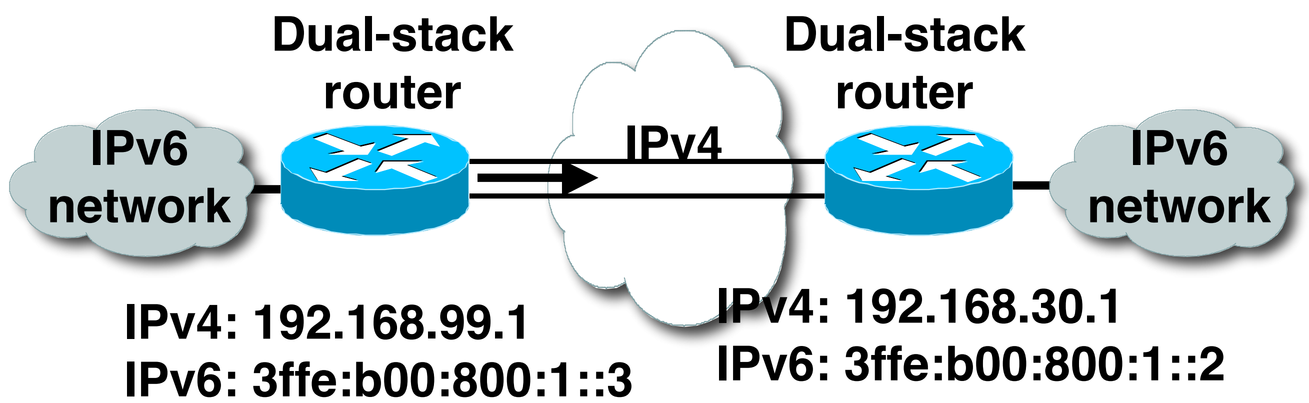

GRE Tunnel require:

Dual stack end points

Both IPv4 and IPv6 addresses configured at each end

Provide secure point to point secure tunnels

GRE Tunnels can be used simultaneously within a network to carry both IPv6 packets and IS-IS link layer messages between IS-IS routers

Manually Configured GRE Tunnel Configuration

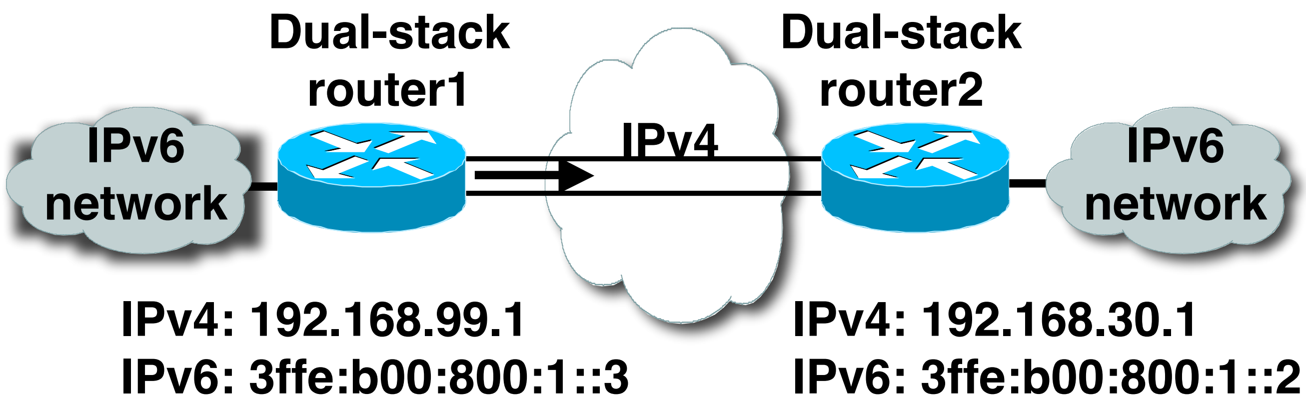


```
router1#

interface Tunnel0
 ipv6 enable
 ipv6 address 3ffe:b00: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 3ffe:b00:c18:1::2/128
 tunnel source 192.168.30.1
 tunnel destination 192.168.99.1
 tunnel mode gre ipv6
```

Semi Automated Tunnel Broker (RFC 3053)

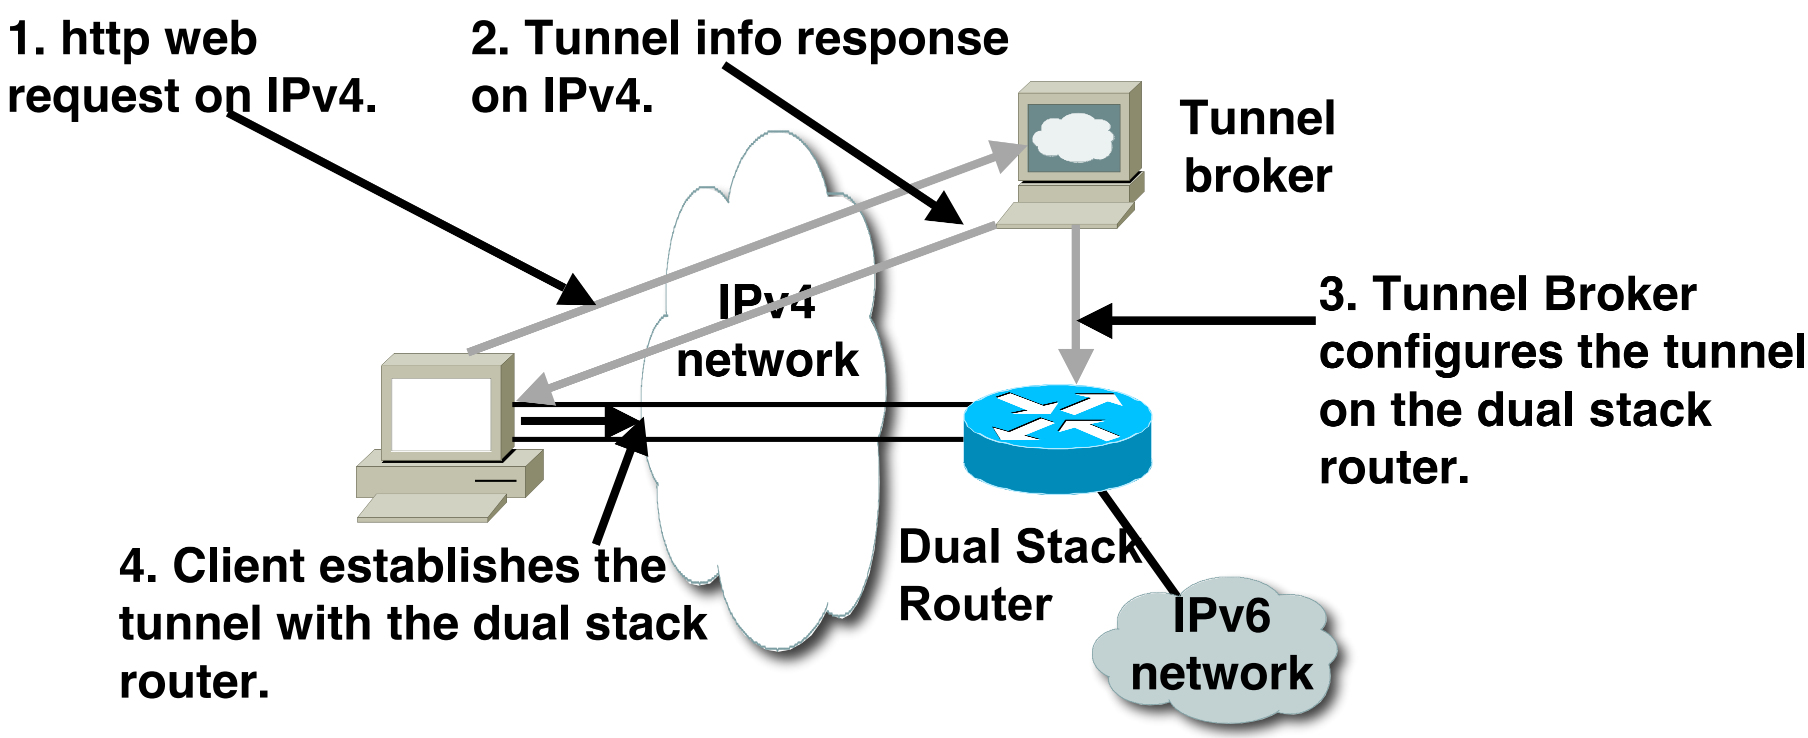

Tunnel broker is a external system rather than a router. Cisco does not support tunnel brokers.

However several tunnel broker implementations available on the Internet uses Cisco routers for their operation ☺

See www. 6bone.net to get information about tunnel brokers available on the internet.

Automatic
6to4 Tunnel (RFC 3056)

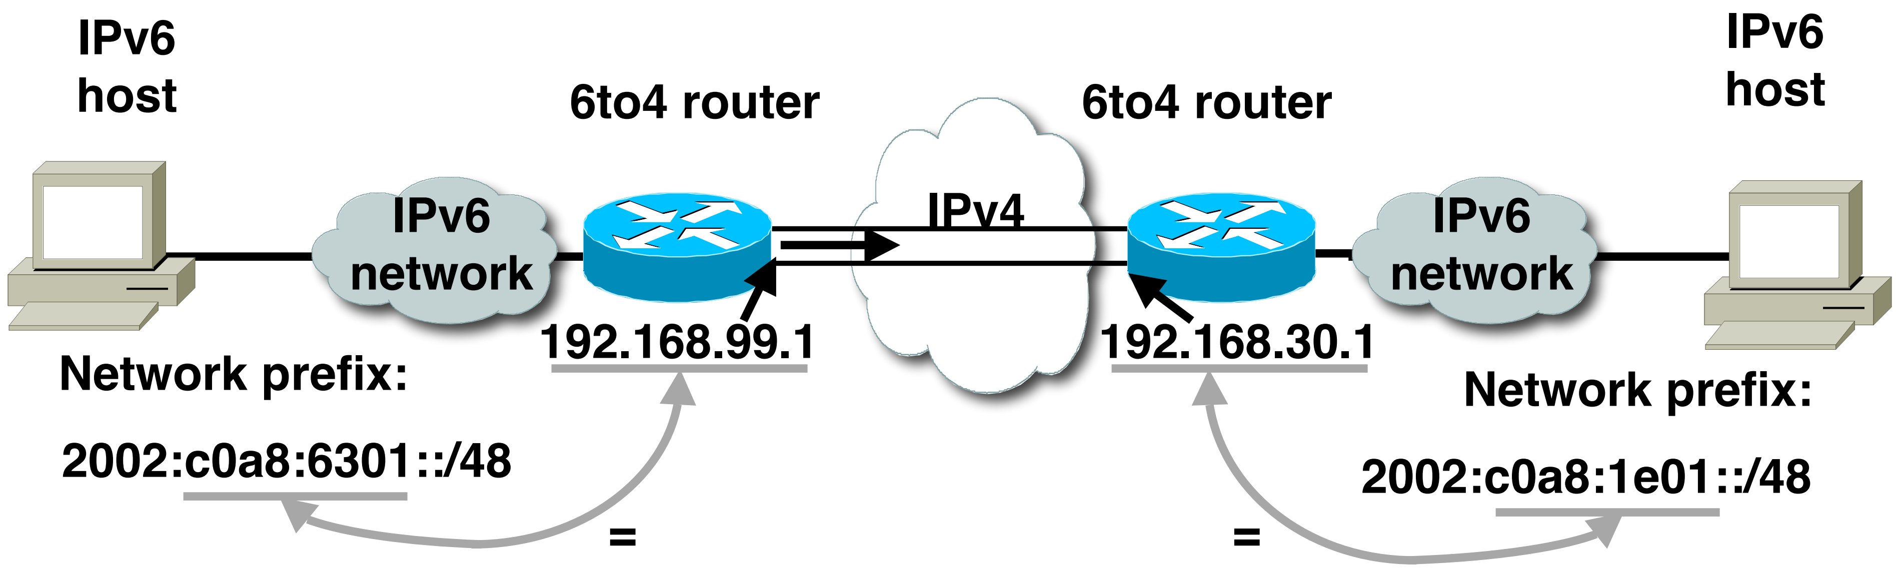

6to4:

Is an automatic tunnel method

Gives a prefix to the attached IPv6 network.

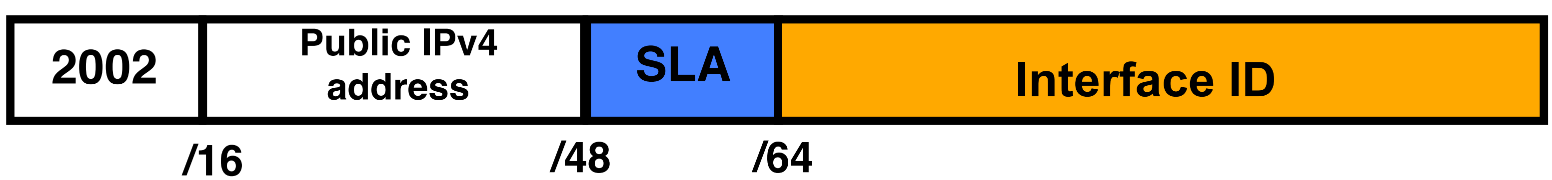

Automatic 6to4 Tunnel (RFC 3056)

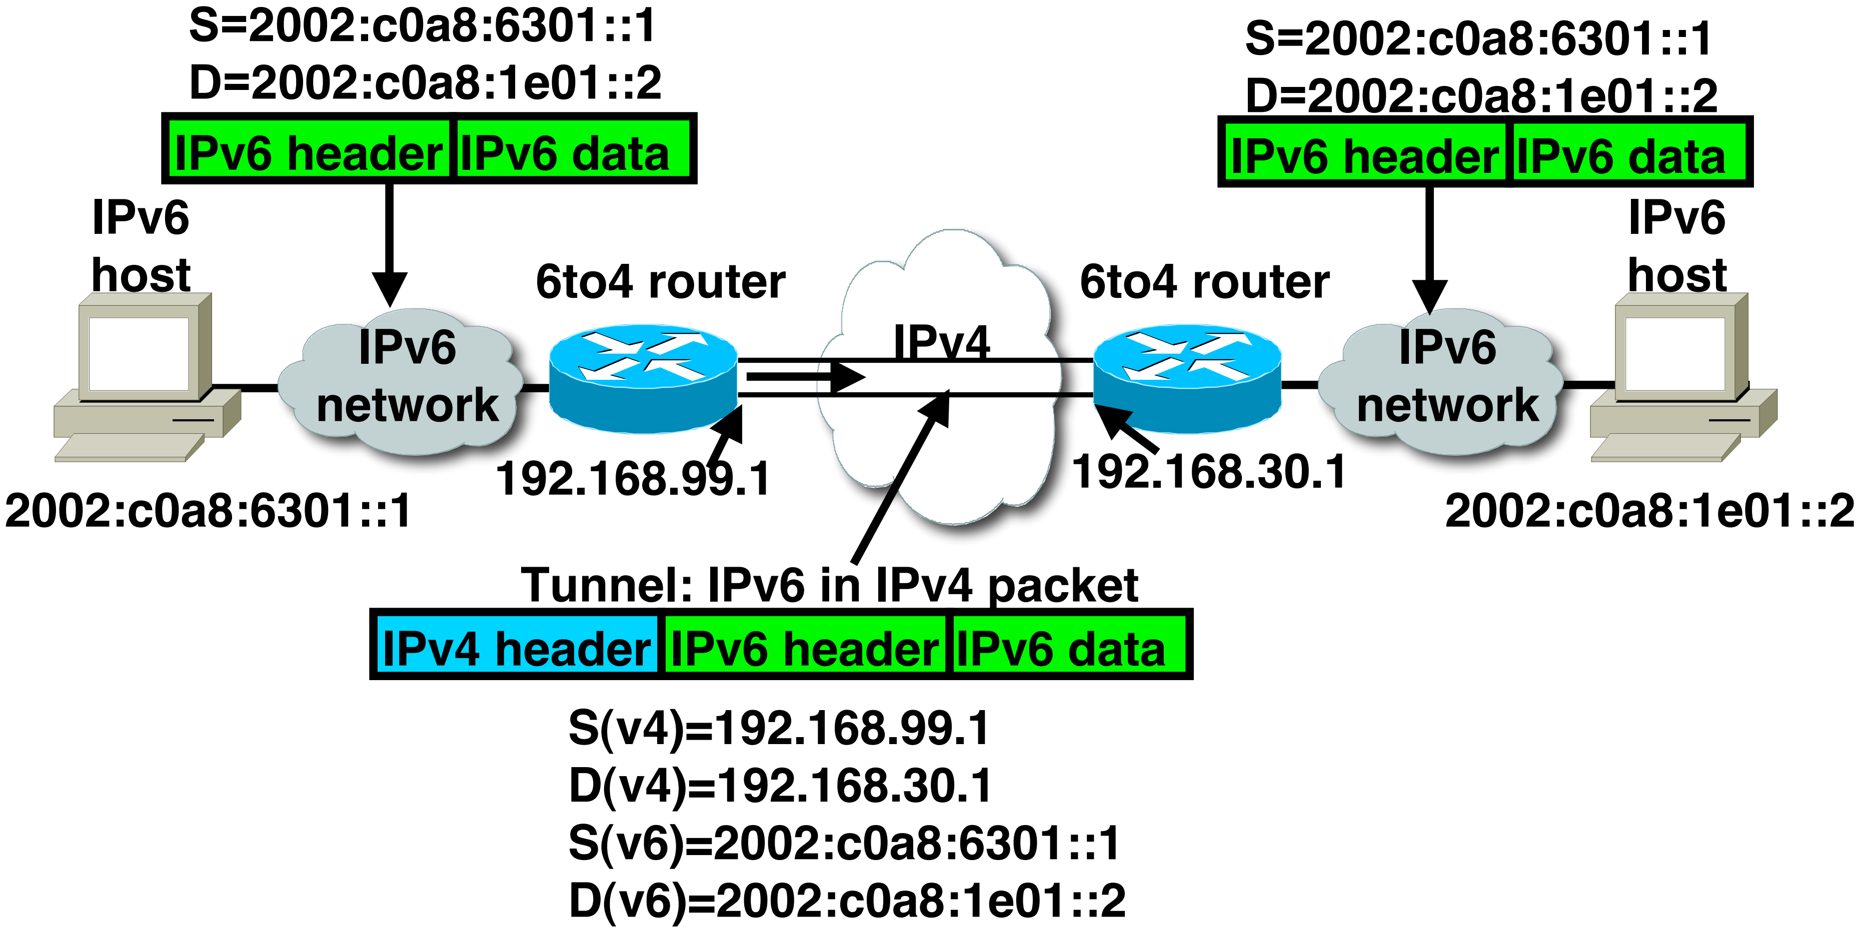

Automatic
6to4 Configuration

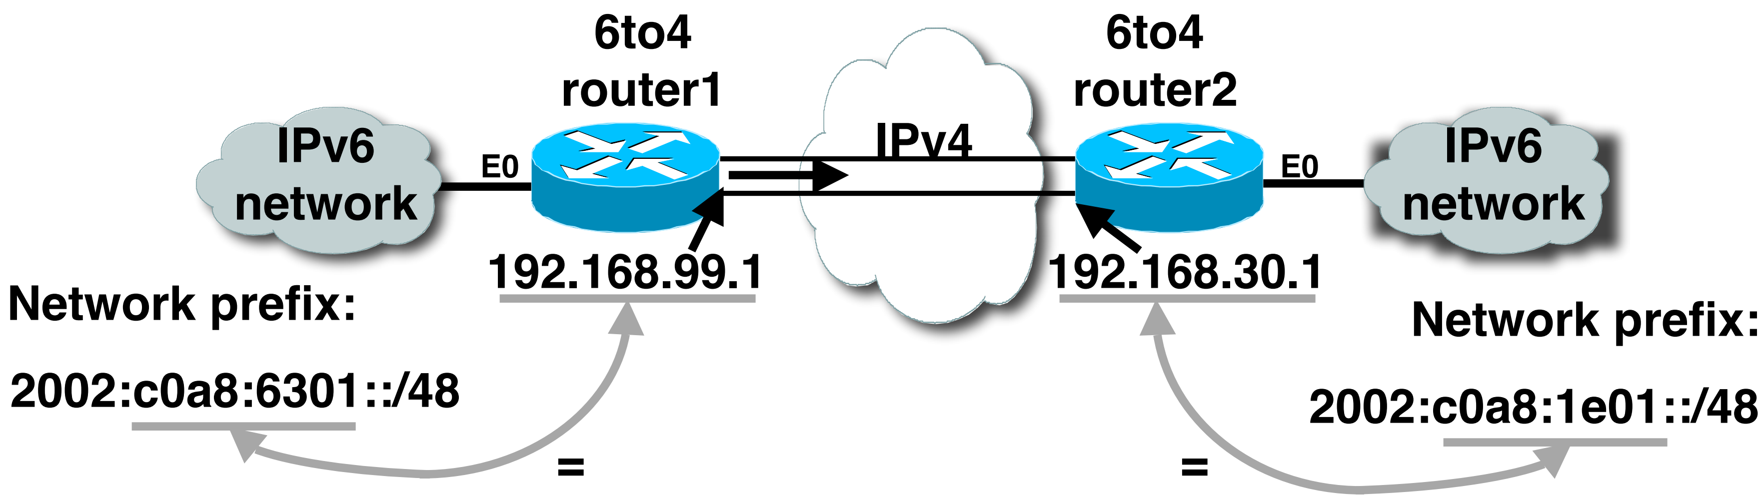


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

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

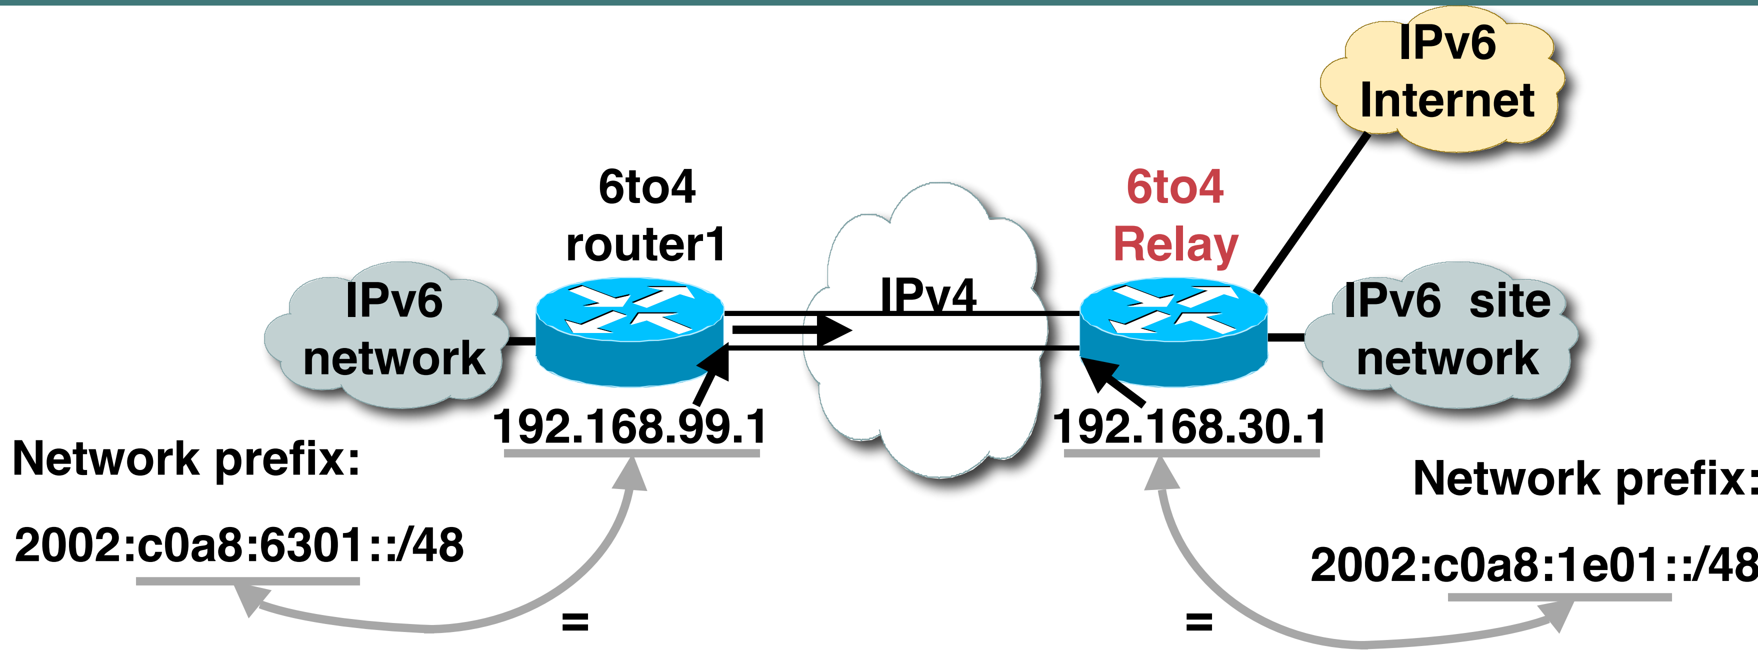

6to4 relay:

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

Automatic
6to4 Relay Configuration

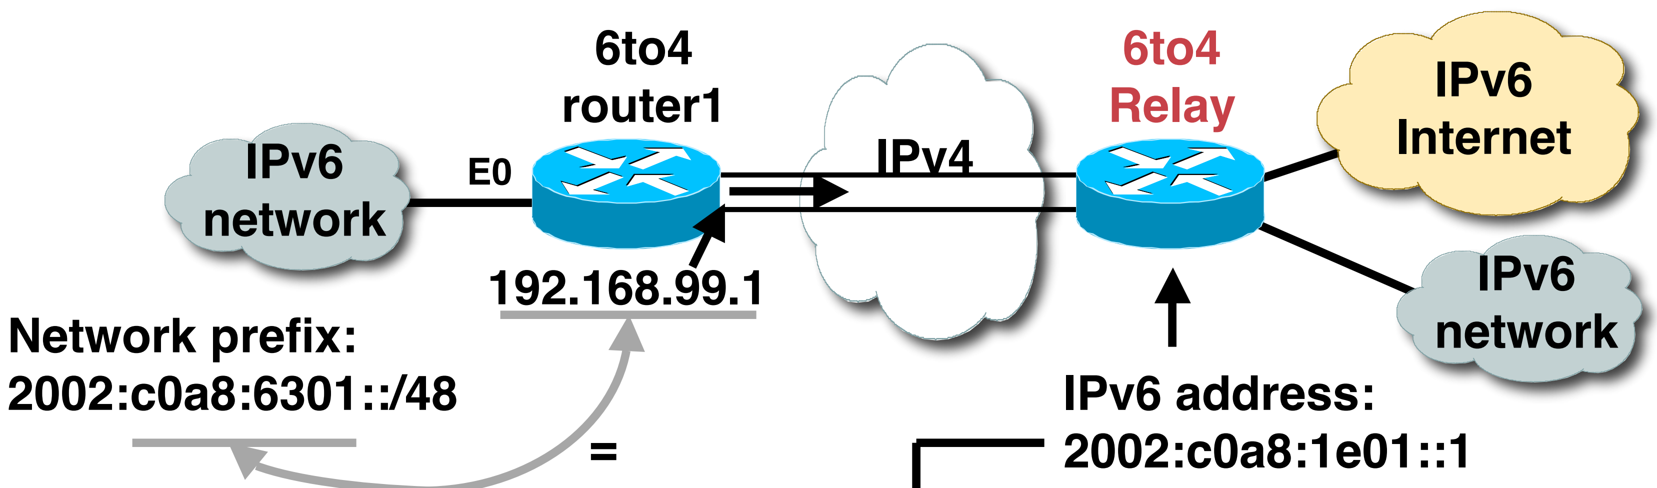


```
router1#
interface Ethernet0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
Interface Ethernet1
  ip address 192.168.99.1 255.255.0.0
interface Tunnel0
 no ip address
 ipv6 unnumbered Ethernet0
 tunnel source Ethernet1
 tunnel mode ipv6ip 6to4

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

Automatic
Intrasite Automatic Tunnel Address Protocol

Use IANA‘s OUI 00-00-5E & encode IPv4 address as part of EUI-64

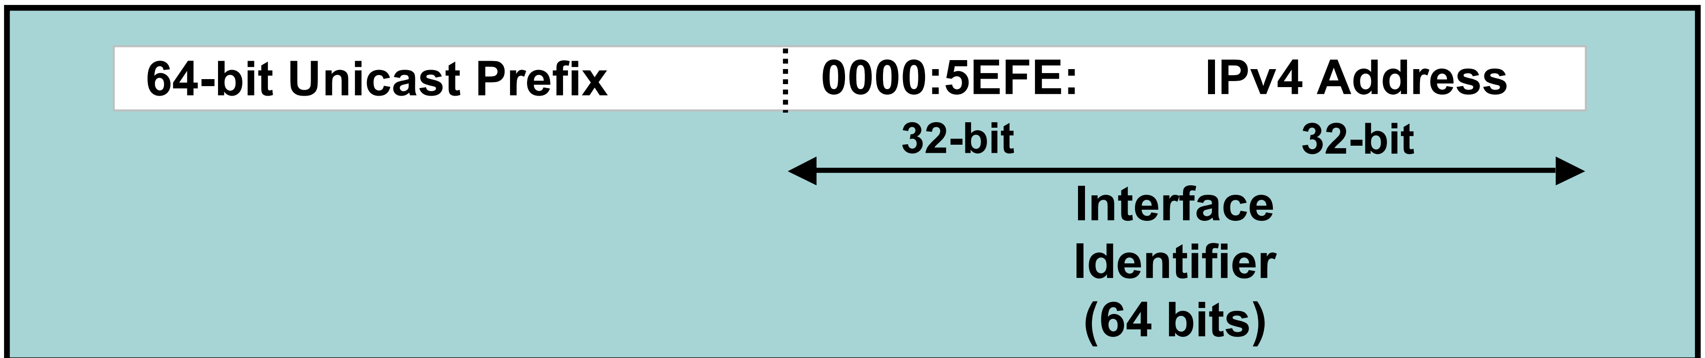

- create a virtual IPv6 network over a IPv4 network
- Supported in Windows XP Pro SP1 and others

draft-ietf-ngtrans-isatap-11
draft-ietf-ngtrans-isatap-scenario-01

Automatic Advertisement of ISATAP Prefix

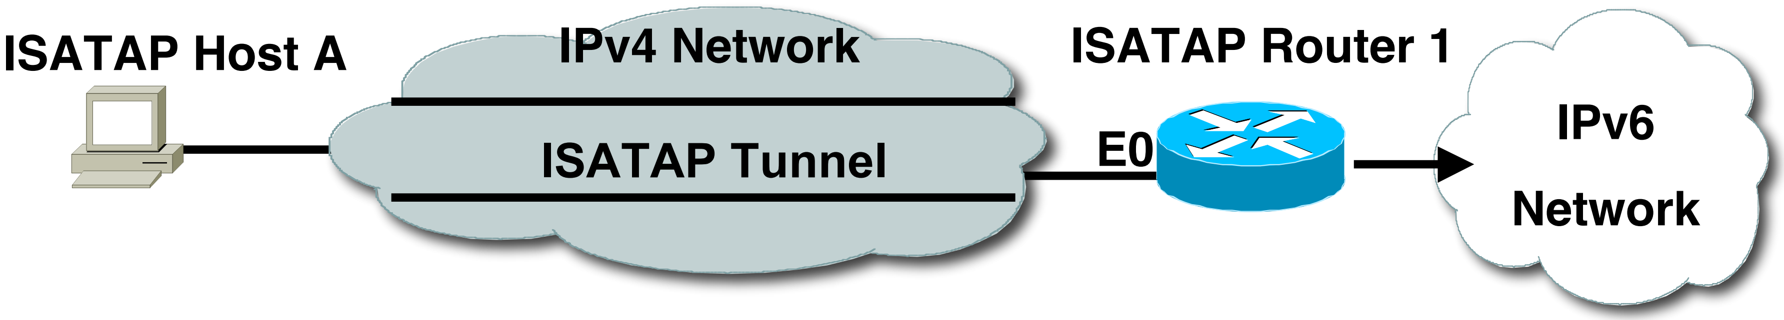

ICMPv6 Type 133 (RS)
IPv4 Source: 206.123.20.100
IPv4 Destination: 206.123.31.200
IPv6 Source: fe80::5efe:ce7b:1464
IPv6 Destination: fe80::5efe:ce7b:1fc8
Send me ISATAP Prefix

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: 3ffe:b00:ffff :2::/64

Automatic
Address Assignment of Host & Router

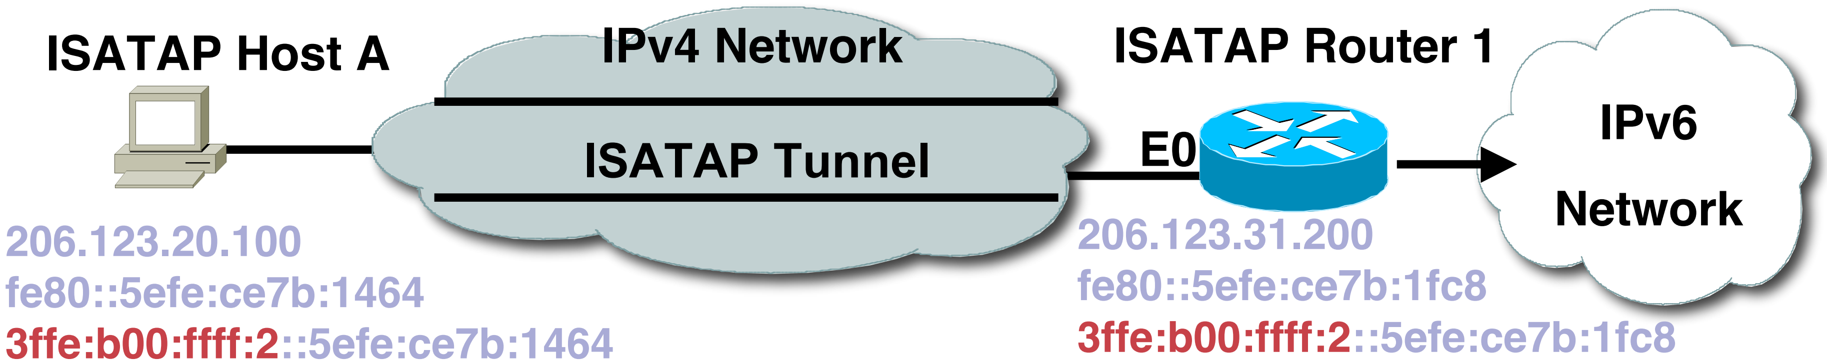

ISATAP host A receives the ISATAP prefix 3ffe:b00:ffff:2::/64 from ISATAP Router 1

When ISATAP host A wants to send IPv6 packets to 3ffe:b00: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

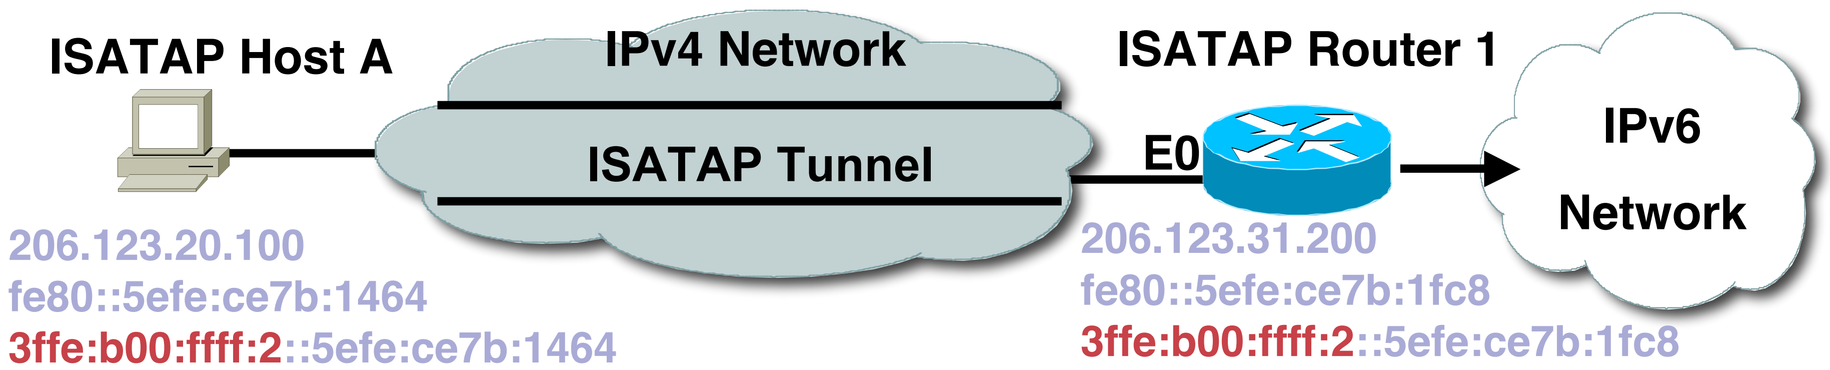


```
ISATAP-router1#
!
interface Ethernet0
 ip address 206.123.31.200 255.255.255.0
!
interface Tunnel0
 ipv6 address 3ffe:b00:ffff:2::/64 eui-64
 no ipv6 nd suppress-ra
 tunnel source Ethernet0
 tunnel mode ipv6ip isatap
```

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

Translation Techniques
NAT-PT for IPv6

- **NAT-PT (Network Address Translation - Protocol Translation) - RFC 2766**
- **NAT-PT allows native IPv6 hosts & applications to communicate with native IPv4 hosts and applications, and vice versa.**
- **Support for ICMP and DNS embedded translation**
- **Easy-to-use transition and co-existence solution**
- **Enable applications to cross the protocol barrier**

NAT-PT Concept

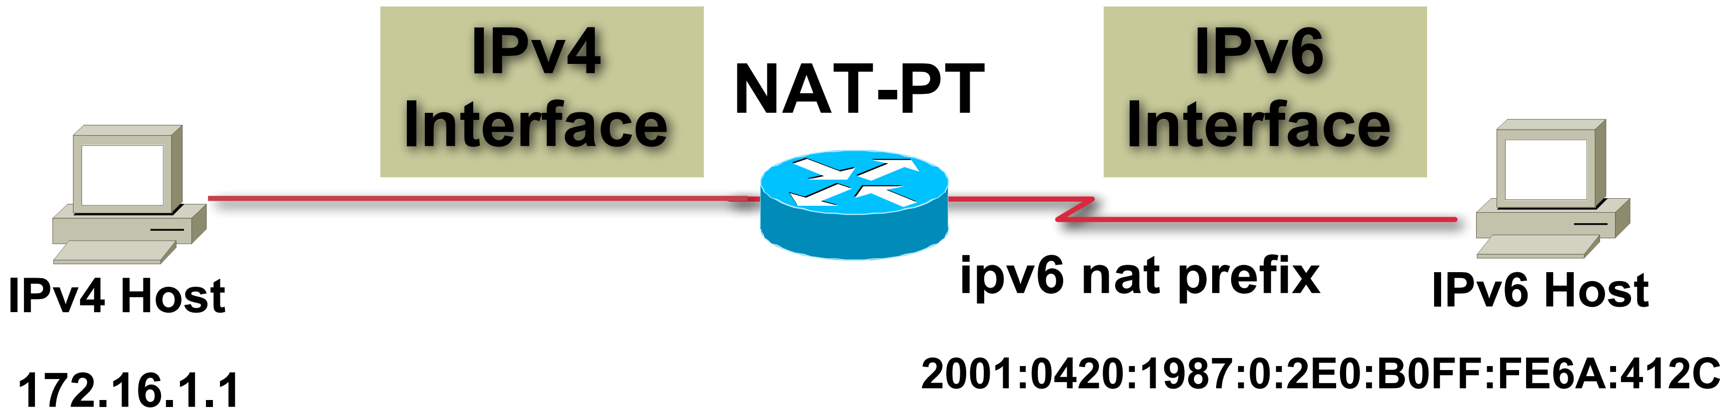

- **PREFIX is a 96-bit field that allows routing back to the NAT-PT device**

NAT-PT Packet Flow

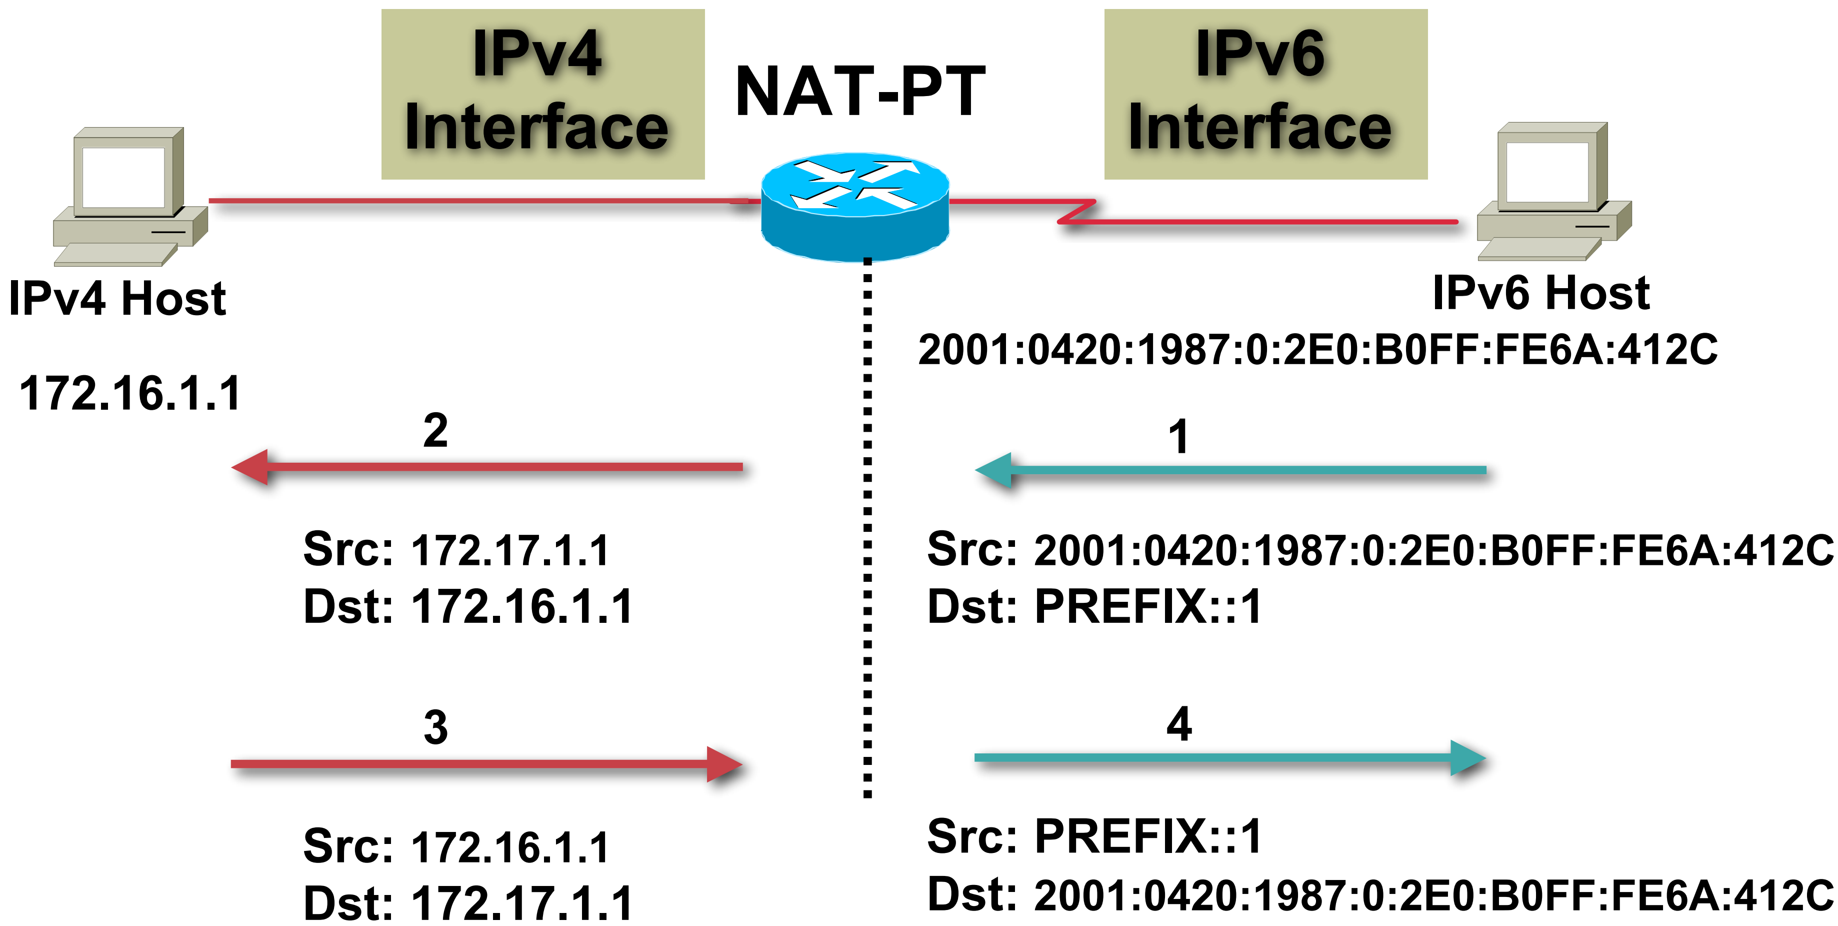

PREFIX is a 96-bit field that allows routing back to the NAT-PT device

ENTERPRISE DEPLOYMENT

Start Here: Cisco IOS Software Release Specifics for IPv6 Features
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/ftipv6s.htm

IPv6 Coexistence in the Enterprise

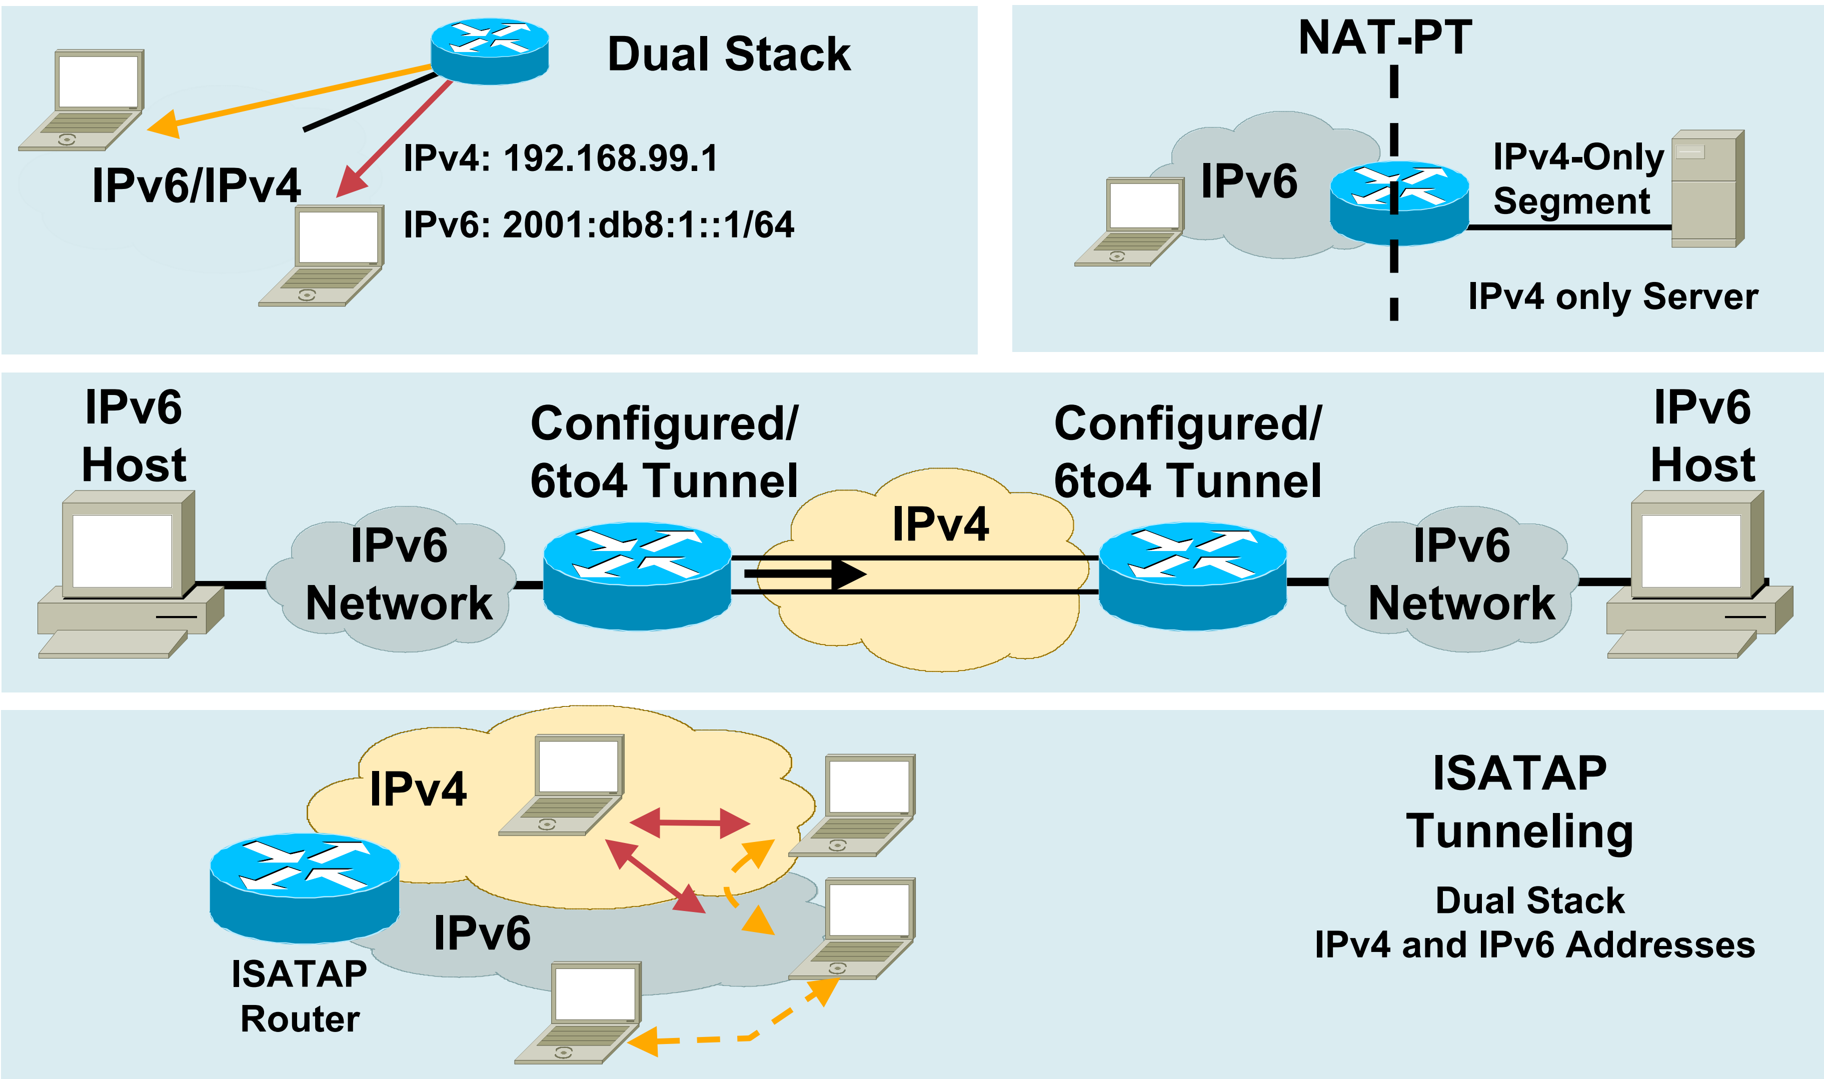

ENTERPRISE DEPLOYMENT: CAMPUS

Campus IPv6 Deployment

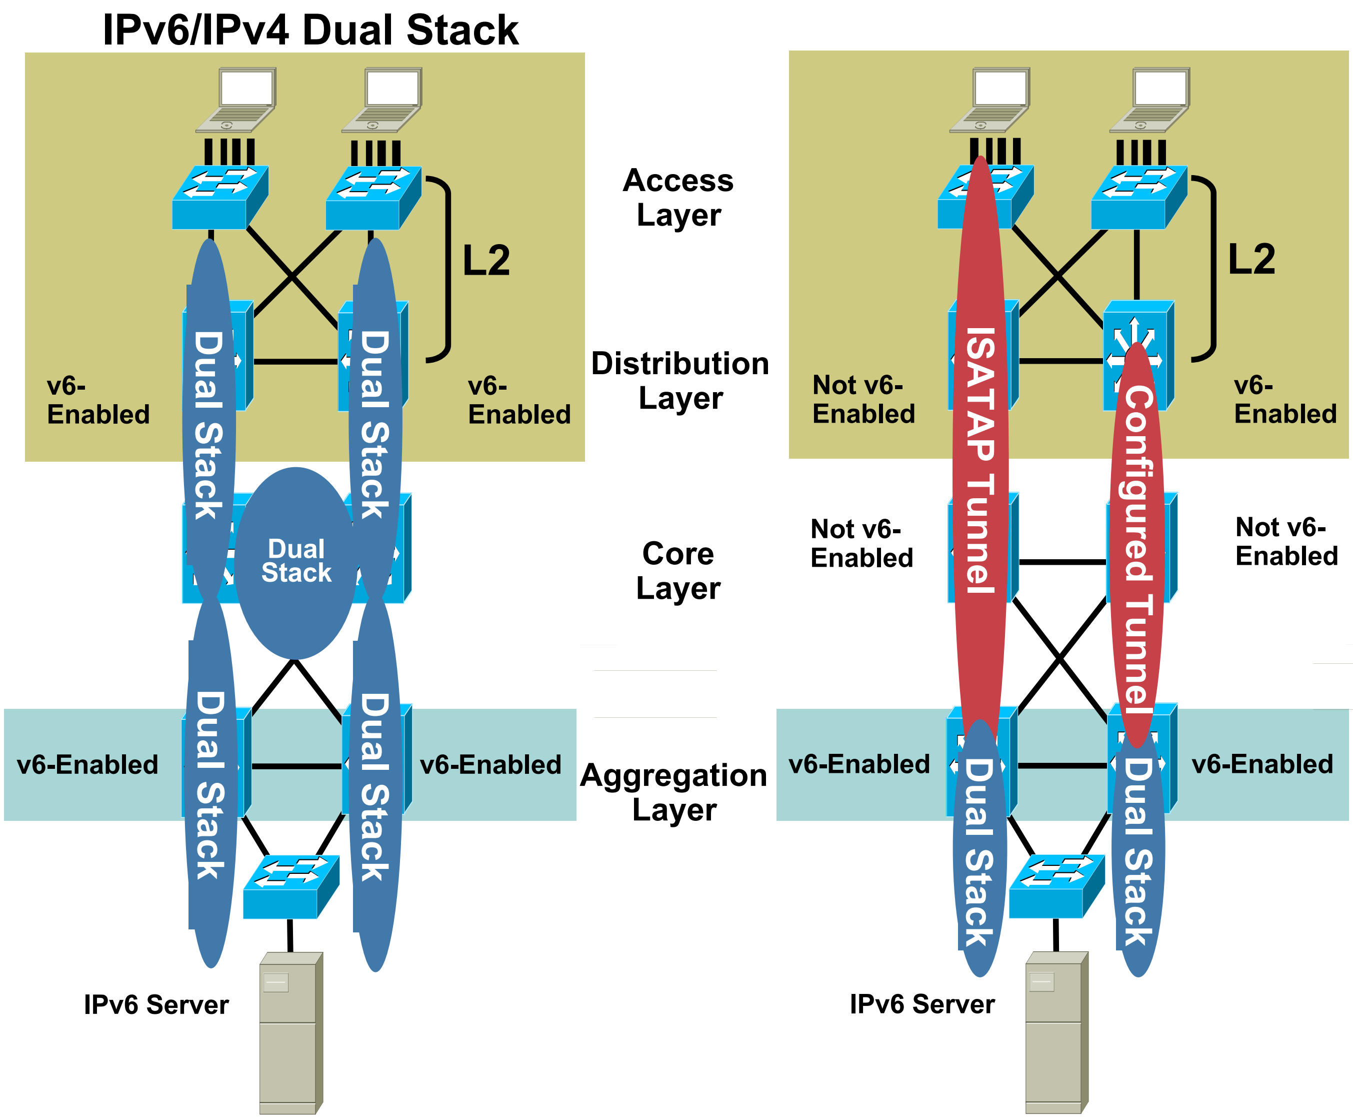

- **Dual Stack and Tunnels in use depending on platform support**
- **Configured between IPv6-enabled L3-switches**
- **ISATAP between clients and a L3-switch**

  ISATAP –

  Intra-Site Automatic Tunnel Addressing Protocol

IPv6 on a Campus: Dual-Stack IPv4-IPv6

- Requires switching/routing platforms to support hardware based forwarding for IPv4 and IPv6
- IPv6 is transparent on L2 switches except for multicast - MLD snooping
    - IPv6 management—Telnet/SSH/HTTP/SNMP
- Requires robust control plane for both IPv4 and IPv6
    - Variety of routing protocols—The same ones in use today with IPv4
- IPv6 multicast, QoS, infrastructure security, etc…
- IPv4 and IPv6 control planes and data planes must not impact each other

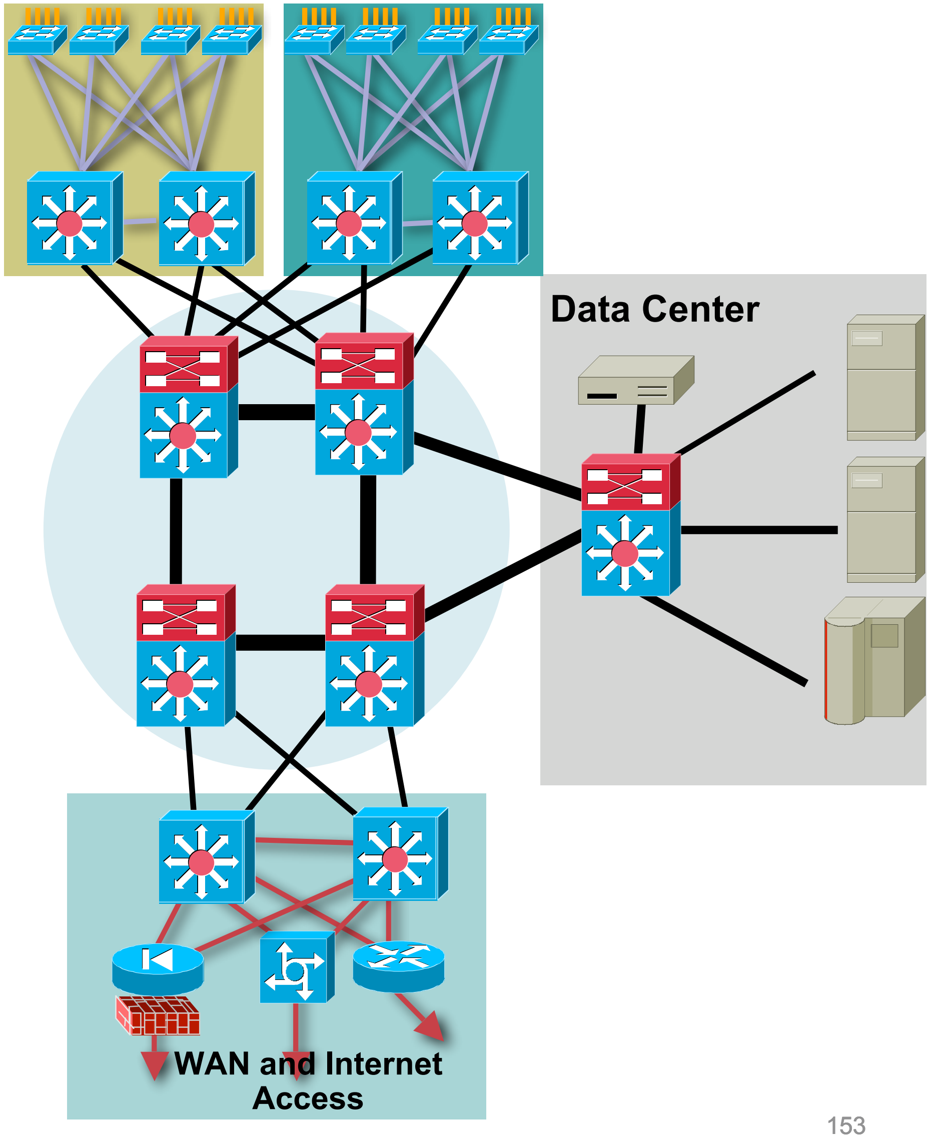

Distribution Layer: Dual Stack

```
ipv6 unicast-routing
ipv6 multicast-routing
ipv6 cef
!
interface GigabitEthernet1/1
 description To 6k-core-right
 ipv6 address 2001:DB8:C003:1105::1/127
 ipv6 ospf 1 area 0
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
!
interface GigabitEthernet1/2
 description To 6k-core-left
 ipv6 address 2001:DB8:C003:1106::3/127
 ipv6 ospf 1 area 0
 ipv6 ospf hello-interval 1
 ipv6 ospf dead-interval 3
 ipv6 cef
```

```
interface Vlan10
 description Data VLAN for Access
 ipv6 address 2001:DB8:C003:1102::1/64
 ipv6 nd prefix 2001:DB8:C003:1102::/64
86400 86400
 ipv6 nd reachable-time 5000
 ipv6 ospf 1 area 1
 ipv6 cef
!
ipv6 router ospf 1
 router-id 10.122.0.25
 log-adjacency-changes
 passive-interface Vlan10
 timers spf 1 1
```

- **Optional: lower valid/preferred lifetimes from defaults (2592000/604800)—in seconds to match DHCPv4 lease times**
- **Optional: lower Neighbor Unreachability Detection (NUD) from 30 seconds (faster failover until HSRP is available - IPv6 HSRP is now available)**

IPv6 Campus ISATAP Configuration

- **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
  ```
- **Currently ISATAP does not support multicast!!**

No Configuration Change on Non-v6 Switches

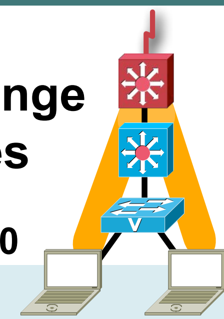


```
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 description ISATAP address for Access Layer
 ip address 10.1.1.1 255.255.255.255
!
interface GigabitEthernet2/10
 ipv6 address 2001:DB8:C003:111C::2/64
 ipv6 cef
!
interface Tunnel0
 ipv6 address 2001:DB8:C003:111F::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 cef
 tunnel source Loopback0
 tunnel mode ipv6ip isatap
```

ISATAP Address Format:

64-bit Unicast Prefix	0000:5EFE	IPv4 Address
	32-bit	32-bit
	Interface ID	

2001:DB8:C003:111F:0:5EFE:10.1.2.100

ENTERPRISE DEPLOYMENT: WAN

WAN Deployment

- Cisco WAN routers support IPv6
- Dual-stack is recommended due to ease of deployment, security advantage and performance
- Support for every media/WAN type you want to use (Frame Relay, leased-line, broadband, MPLS, etc…)

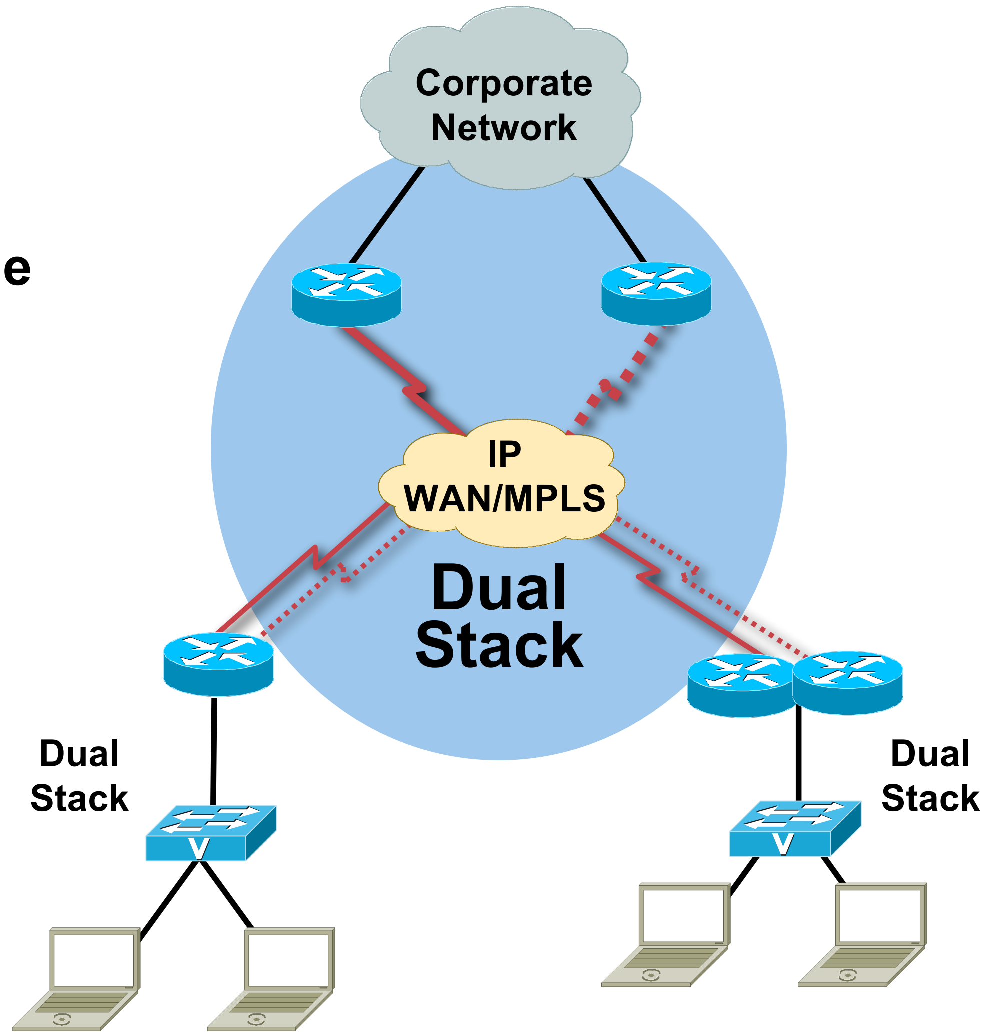

HSRP for IPv6

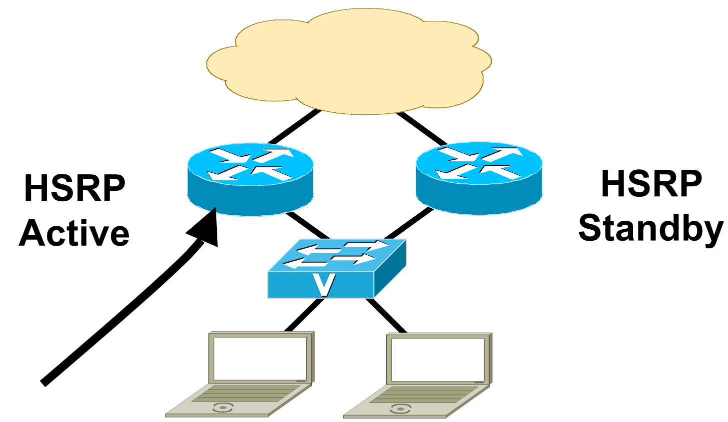

- Basically the same as HSRP for IPv4
- Changes occur in Neighbor Advertisement, Router Advertisement, and ICMPv6 redirects
- Virtual MAC derived from HSRP group number and virtual IPv6 Link-local address
- IPv6 Virtual MAC range:

  0005.73A0.0000 - 0005.73A0.0FFF (4096 addresses)
- HSRP IPv6 UDP Port Number 2029 (IANA Assigned)
- No HSRP IPv6 secondary address
- HSRP IPv6 specific debug

```
interface FastEthernet0/1
 ipv6 address 2001:66:67::2/64
 ipv6 cef
 standby version 2
 standby 1 ipv6 autoconfig
 standby 1 timers msec 250 msec 800
 standby 1 preempt
 standby 1 preempt delay minimum 180
 standby 1 authentication md5 key-string cisco
 standby 1 track FastEthernet0/0
```

Host with GW of Virtual IP

```
#route -A inet6 | grep ::/0 | grep eth2
::/0      fe80::207:85ff:fef3:2f60        UGDA  1024  3        0 eth2
::/0      fe80::205:9bff:febf:5ce0        UGDA  1024  0        0 eth2
::/0      fe80::5:73ff:fea0:1             UGDA  1024  0        0 eth2
```

OTHER TRANSITION TYPES

Configured Tunnel

Building the Tunnel

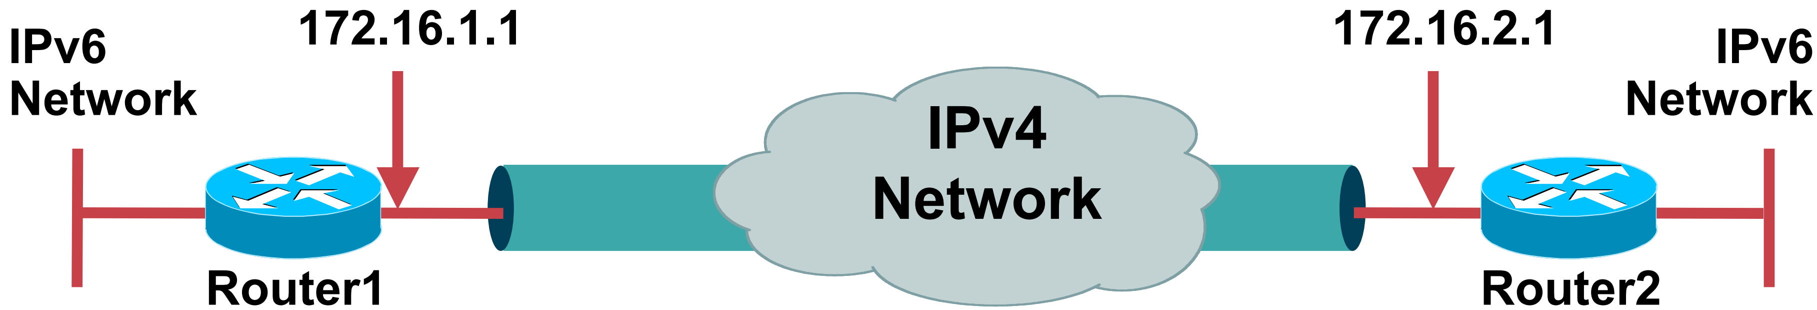


```
ipv6 unicast-routing
ipv6 cef
!
interface Tunnel0
 no ip address
 ipv6 cef
 ipv6 address 2001:DB8:C003:3::1/64
 ipv6 rip v6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.2.1
 tunnel mode ipv6ip
!
interface FastEthernet0/0
 ipv6 address 2001:DB8:C003:1::1/64
 ipv6 cef
!
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.252
```

```
ipv6 unicast-routing
ipv6 cef
!
interface Tunnel0
 no ip address
 ipv6 cef
 ipv6 address 2001:DB8:C003:3::2/64
 ipv6 rip v6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.1.1
 tunnel mode ipv6ip
!
interface FastEthernet0/0
 ipv6 address 2001:DB8:C003:2::1/64
 ipv6 cef
!
interface FastEthernet0/1
 ip address 172.16.2.1 255.255.255.252
```

Internet Key Exchange (IKE) Policy
Configured Tunnel (Static Maps)

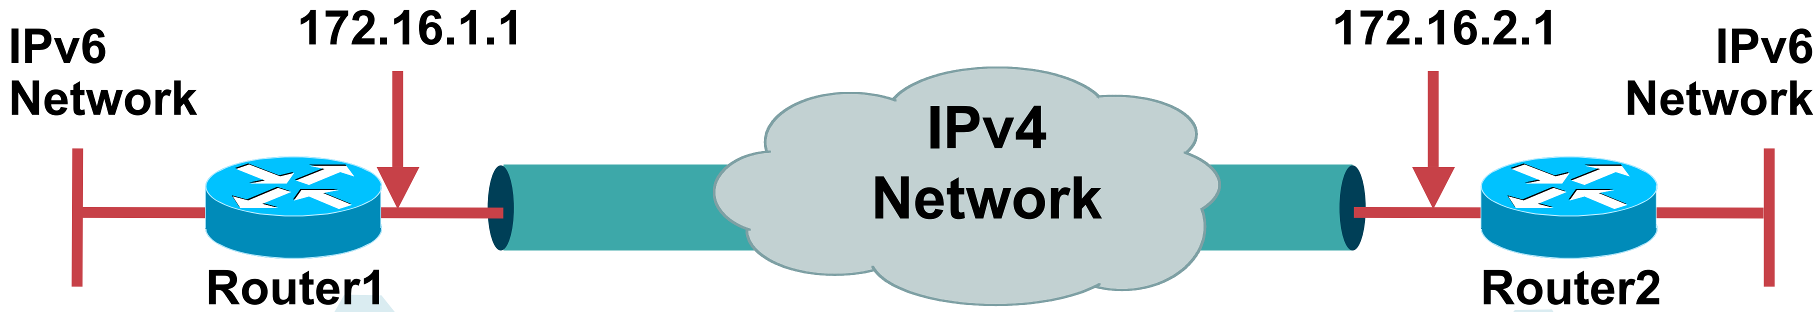

Router1:

```
crypto isakmp policy 1
   encr 3des
   authentication pre-share
   group 2
!
crypto isakmp key CISCO address
   172.16.2.1
```

Router2:

```
crypto isakmp policy 1
   encr 3des
   authentication pre-share
   group 2
!
crypto isakmp key CISCO address
   172.16.1.1
```

IPSec Policy
Configured Tunnel (Static Maps)

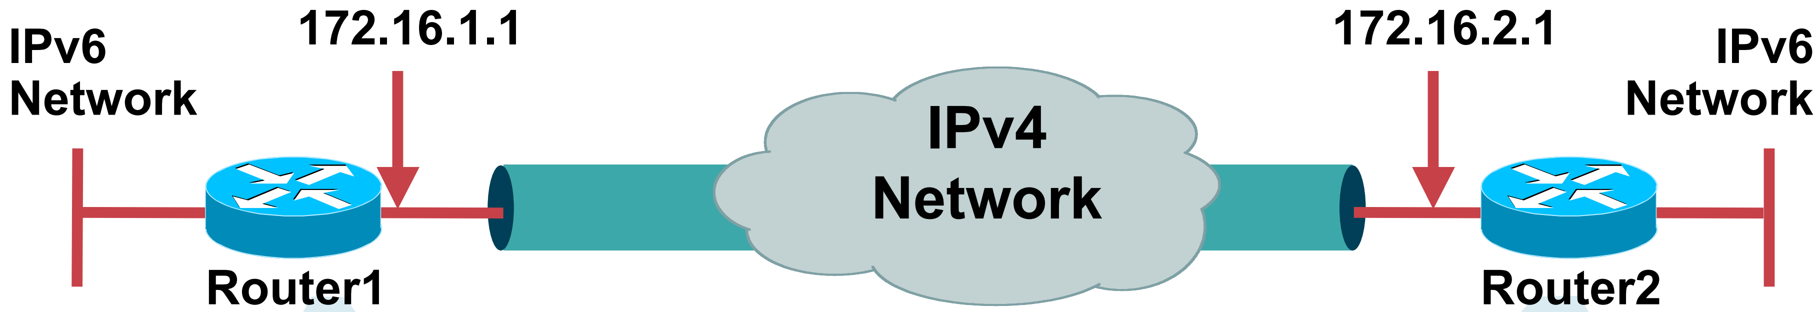


```
crypto ipsec transform-set STRONG esp-
3des esp-sha-hmac
!
crypto map STATIC-MAP local-address
FastEthernet0/1
!
crypto map STATIC-MAP 1 ipsec-isakmp
 set peer 172.16.2.1
 set transform-set STRONG
 set pfs group2
 match address VPN-TO-R2
!
ip access-list extended VPN-TO-R2
 permit 41 host 172.16.1.1 host
172.16.2.1
```

```
crypto ipsec transform-set STRONG esp-
3des esp-sha-hmac
!
crypto map STATIC-MAP local-address
FastEthernet0/1
!
crypto map STATIC-MAP 1 ipsec-isakmp
 set peer 172.16.1.1
 set transform-set STRONG
 set pfs group2
 match address VPN-TO-R1
!
ip access-list extended VPN-TO-R1
 permit 41 host 172.16.2.1 host
172.16.1.1
```

Apply VPN Configuration
Configured Tunnel

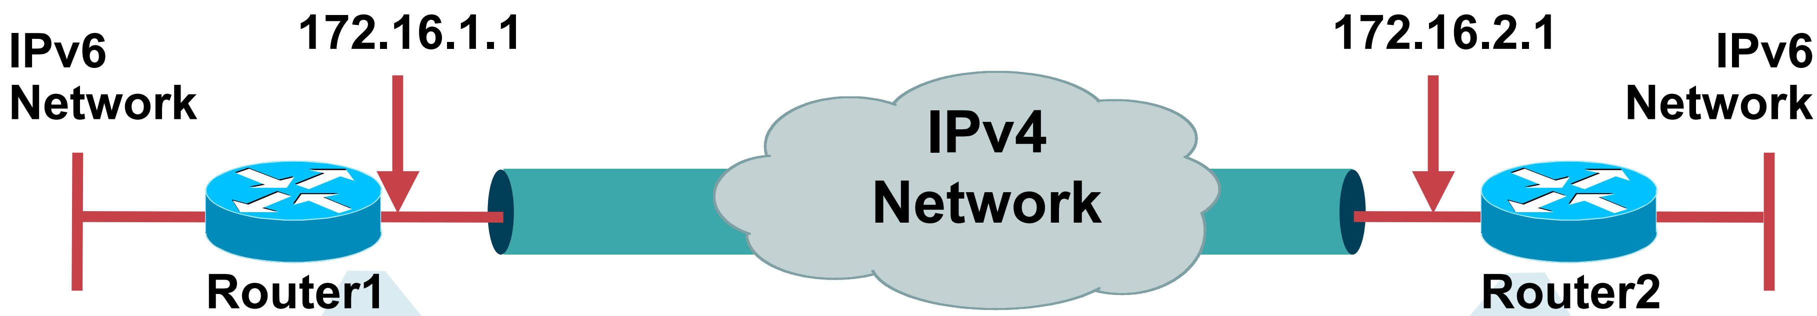

Router1:

```
interface Tunnel0
 no ip address
 ipv6 address 2001:DB8:C003:3::1/64
 ipv6 cef
 ipv6 mtu 1400
 ipv6 rip V6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.2.1
 tunnel mode ipv6ip
!
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.252
 crypto map STATIC-MAP
!
ip route 172.16.2.1 255.255.255.255
172.16.1.2
```

Router2:

```
interface Tunnel0
 no ip address
 ipv6 address 2001:DB8:C003:3::2/64
 ipv6 cef
 ipv6 mtu 1400
 ipv6 rip V6 enable
 tunnel source FastEthernet0/1
 tunnel destination 172.16.1.1
 tunnel mode ipv6ip
!
interface FastEthernet0/1
 ip address 172.16.2.1 255.255.255.252
 crypto map STATIC-MAP
!
ip route 172.16.1.1 255.255.255.255
172.16.2.2
```

IPv6 IPSec Example
IKE/IPSec Policies

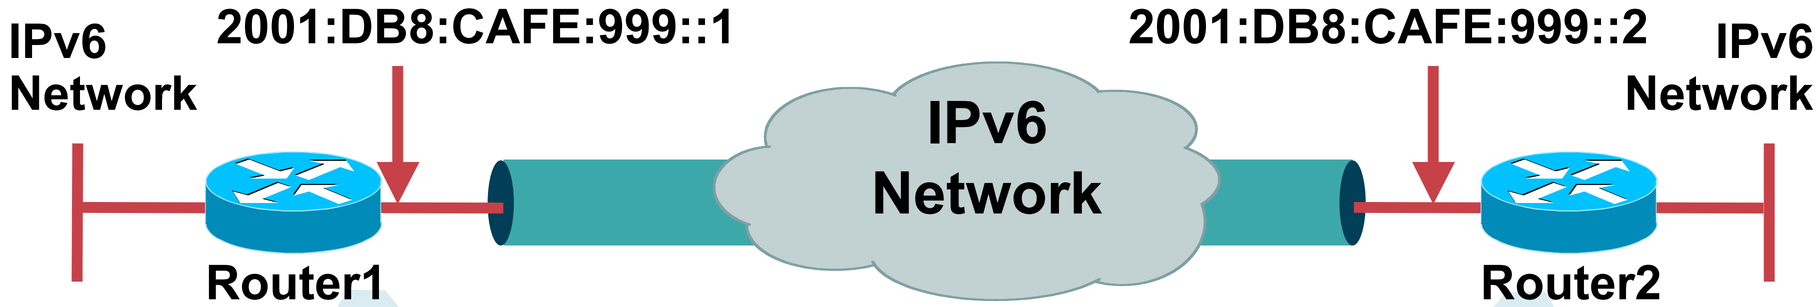


```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key CISCOKEY address ipv6
2001:DB8:CAFE:999::2/128
crypto isakmp keepalive 30 30
!
crypto ipsec transform-set v6STRONG esp-
3des esp-sha-hmac
!
crypto ipsec profile v6PRO
 set transform-set v6STRONG
```

```
crypto isakmp policy 1
 authentication pre-share
crypto isakmp key CISCOKEY address ipv6
2001:DB8:CAFE:999::1/128
crypto isakmp keepalive 30 30
!
crypto ipsec transform-set v6STRONG esp-
3des esp-sha-hmac
!
crypto ipsec profile v6PRO
 set transform-set v6STRONG
```

IPv6 IPSec Example

Tunnels

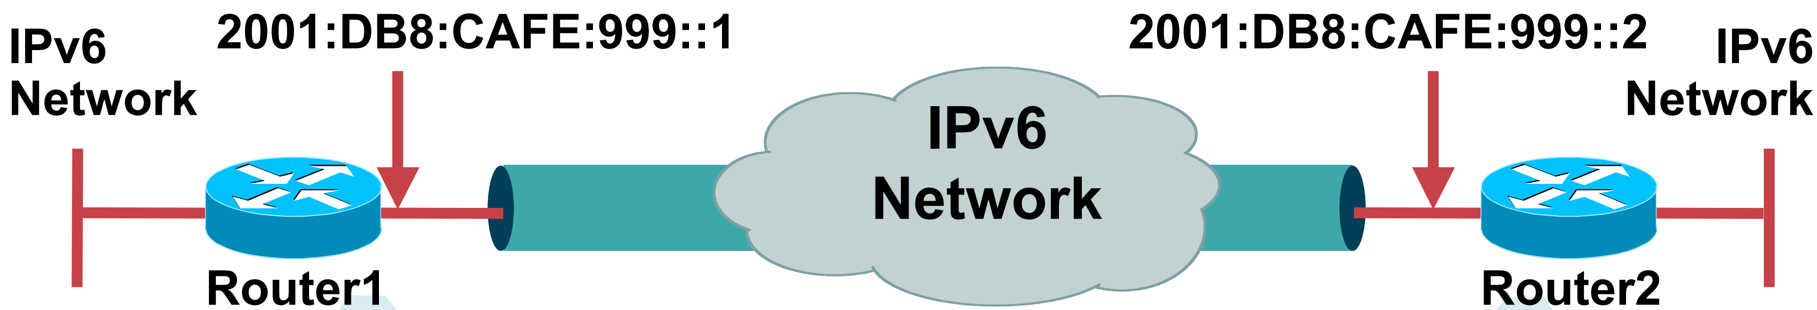


```
interface Tunnel0
 ipv6 address 2001:DB8:CAFE:F00D::/64
eui-64
 ipv6 ospf 1 area 0
 tunnel source Serial2/0
 tunnel destination
2001:DB8:CAFE:999::2
 tunnel mode ipsec ipv6
 tunnel protection ipsec profile v6PRO
!
interface Ethernet0/0
 ipv6 address 2001:DB8:CAFE:100::1/64
 ipv6 ospf 1 area 1
!
interface Serial2/0
 ipv6 address 2001:DB8:CAFE:999::1/64
```

```
interface Tunnel0
 ipv6 address 2001:DB8:CAFE:F00D::/64
eui-64
 ipv6 ospf 1 area 0
 tunnel source Serial2/0
 tunnel destination
2001:DB8:CAFE:999::1
 tunnel mode ipsec ipv6
 tunnel protection ipsec profile v6PRO
!
interface Ethernet0/0
 ipv6 address 2001:DB8:CAFE:200::1/64
 ipv6 ospf 1 area 2
!
interface Serial2/0
 ipv6 address 2001:DB8:CAFE:999::2/64
```

IPv6 IPSec Example

Show Output

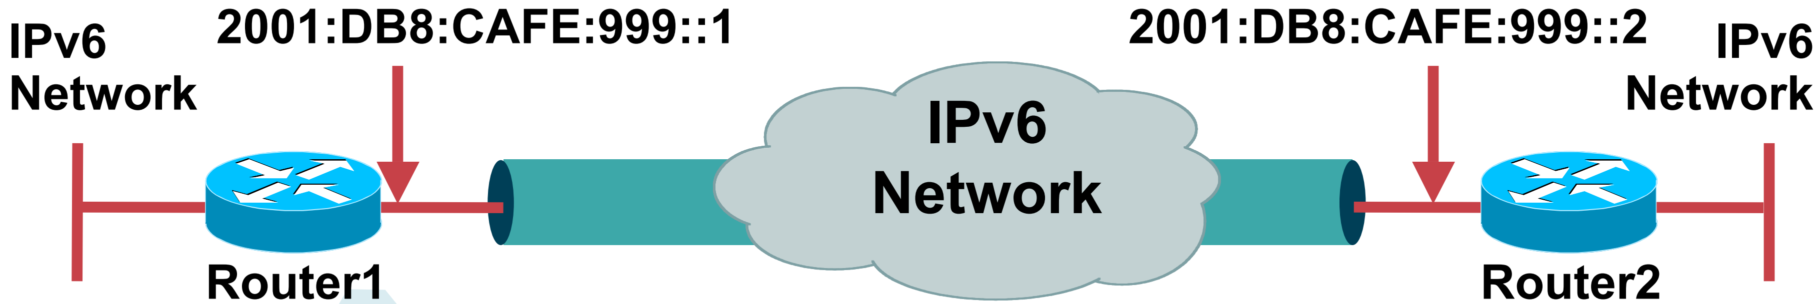


```
Router1#show crypto engine connections active
Crypto Engine Connections
   ID Intfc  Type  Algorithm           Encrypt  Decrypt IP-Address
    3 Tu0    IPsec 3DES+SHA                  0       17 2001:DB8:CAFE:999::1
    4 Tu0    IPsec 3DES+SHA                 16        0 2001:DB8:CAFE:999::1
 1006 Tu0    IKE   SHA+DES                   0        0 2001:DB8:CAFE:999::1

Router1#show crypto sessions
Crypto session current status
Interface: Tunnel0
Session status: UP-ACTIVE
Peer: 2001:DB8:CAFE:999::2 port 500
  IKE SA: local 2001:DB8:CAFE:999::1/500
          remote 2001:DB8:CAFE:999::2/500 Active
  IPSEC FLOW: permit 41 ::/0 ::/0
        Active SAs: 2, origin: crypto map
```

Configured Tunnels vs. Automatic Tunnels

	Configured	ISATAP	6to4
Manual Configuration per Client (Router-Side)	YES	NO	NO
Manual Configuration per Client (Client-Side)	YES	NO	NO
IPv6 Multicast Support	YES	NO	NO
Broad Client OS Support	YES	NO	YES
Optimal for Remote Access Clients	NO	YES	YES

*GRE must be used if ISIS is used as the routing protocol

Legacy Services (IPv4 Only)

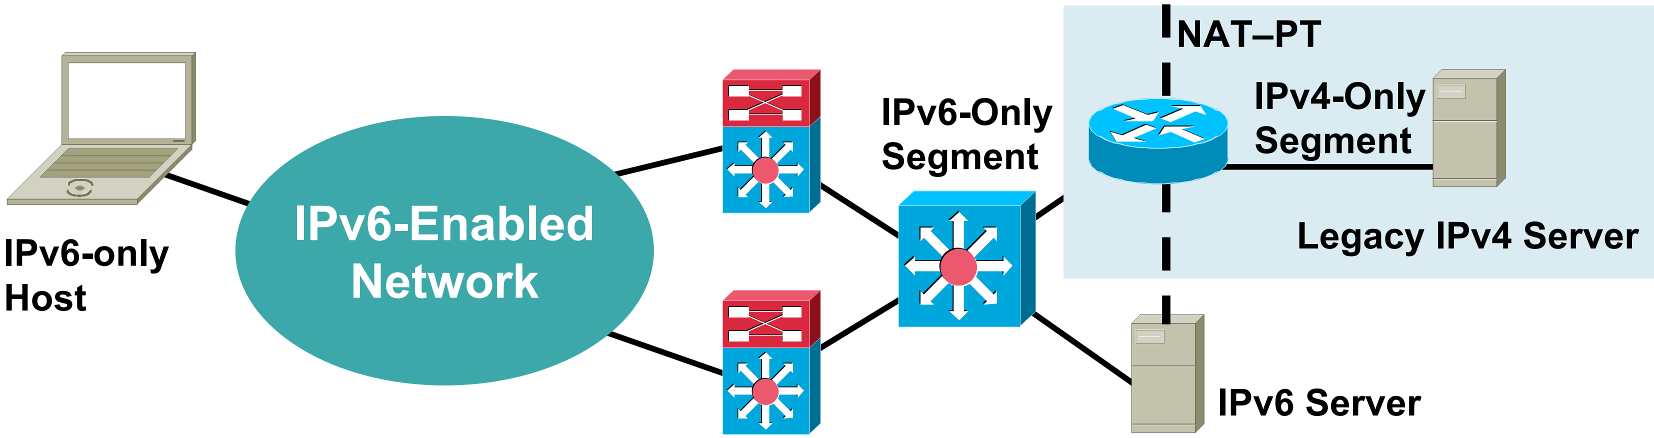

- Many of the non-routing/switching products do not yet support IPv6 (i.e., content switching modules)
- NAT-PT (Network Address Translation–Protocol Translation) as an option to front-end IPv4-only server—NOTE: NAT-PT IS BEING MOVED TO EXPERIMENTAL
- Place NAT-PT box as close to IPv4 only server as possible
- Be VERY aware of performance and manageability issues

Configuring Cisco IOS NAT-PT

- **NAT-PT enables communication between IPv6-only and IPv4-only nodes**
- **CEF switching in 12.3(14)T**

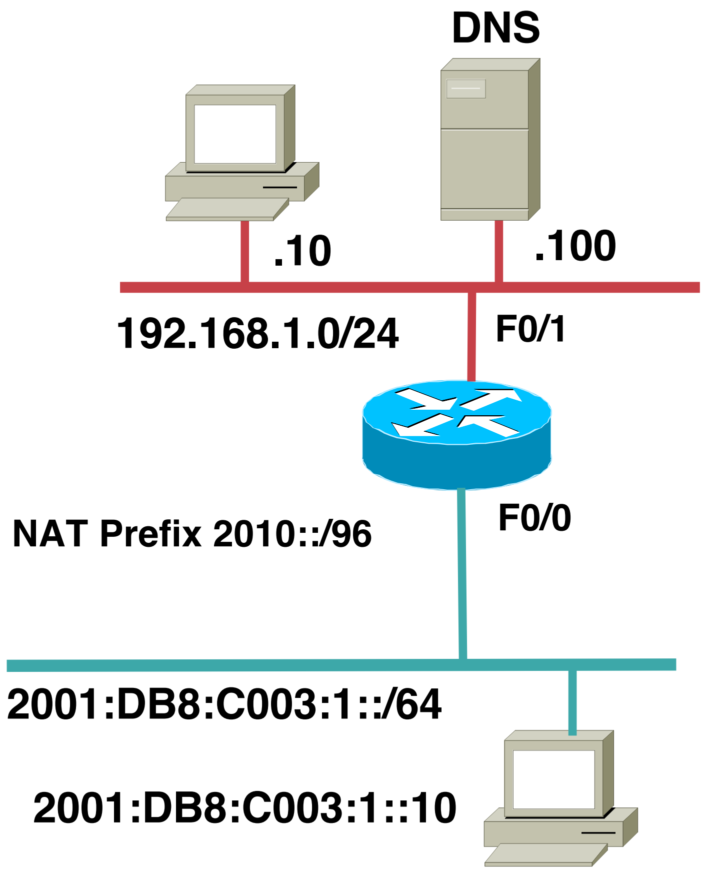


```
interface FastEthernet0/0
   ipv6 address 2001:DB8:C003:1::1/64
   ipv6 cef
   ipv6 nat
!
interface FastEthernet0/1
   ip address 192.168.1.1 255.255.255.0
   ipv6 nat prefix 2010::/96
   ipv6 nat
!
ipv6 nat v4v6 source 192.168.1.100 2010::100
!
ipv6 nat v6v4 source route-map MAP1 pool V4POOL
ipv6 nat v6v4 pool V4POOL 192.168.2.1
192.168.2.10 prefix-length 24
!
route-map MAP1 permit 10
 match interface FastEthernet0/1
```

ENTERPRISE DEPLOYMENT: REMOTE ACCESS

IPv6 for Remote Devices

- Remote nodes can use a VPN client or router to establish connectivity back to enterprise
- Possible over IPv4 today, not possible over IPv6 today (key management is still in progress)
- How could we allow access to IPv6 services at central site or Internet in a secure fashion?
  - Enabling IPv6 traffic inside the Cisco VPN client tunnel
  - Allow remote host to establish a v6-in-v4 tunnel either automatically or manually
    - ISATAP—Intra Site Automatic Tunnel Addressing Protocol
    - Configured—Static configuration for each side of tunnel
  - Same split-tunneling issues exists

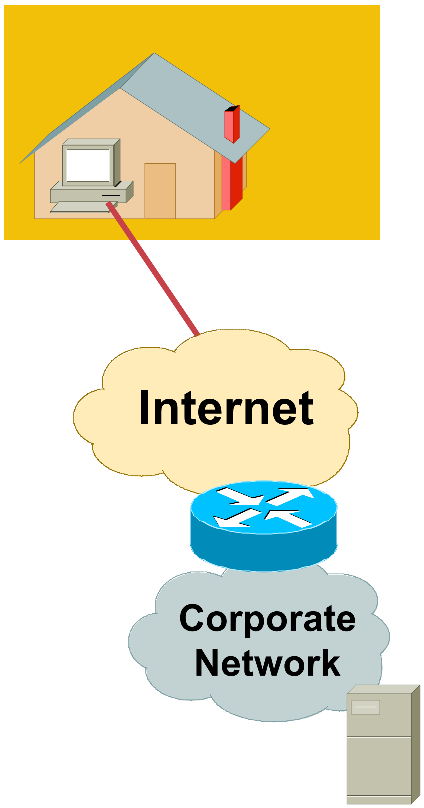

IPv6-in-IPv4 Tunnel Example

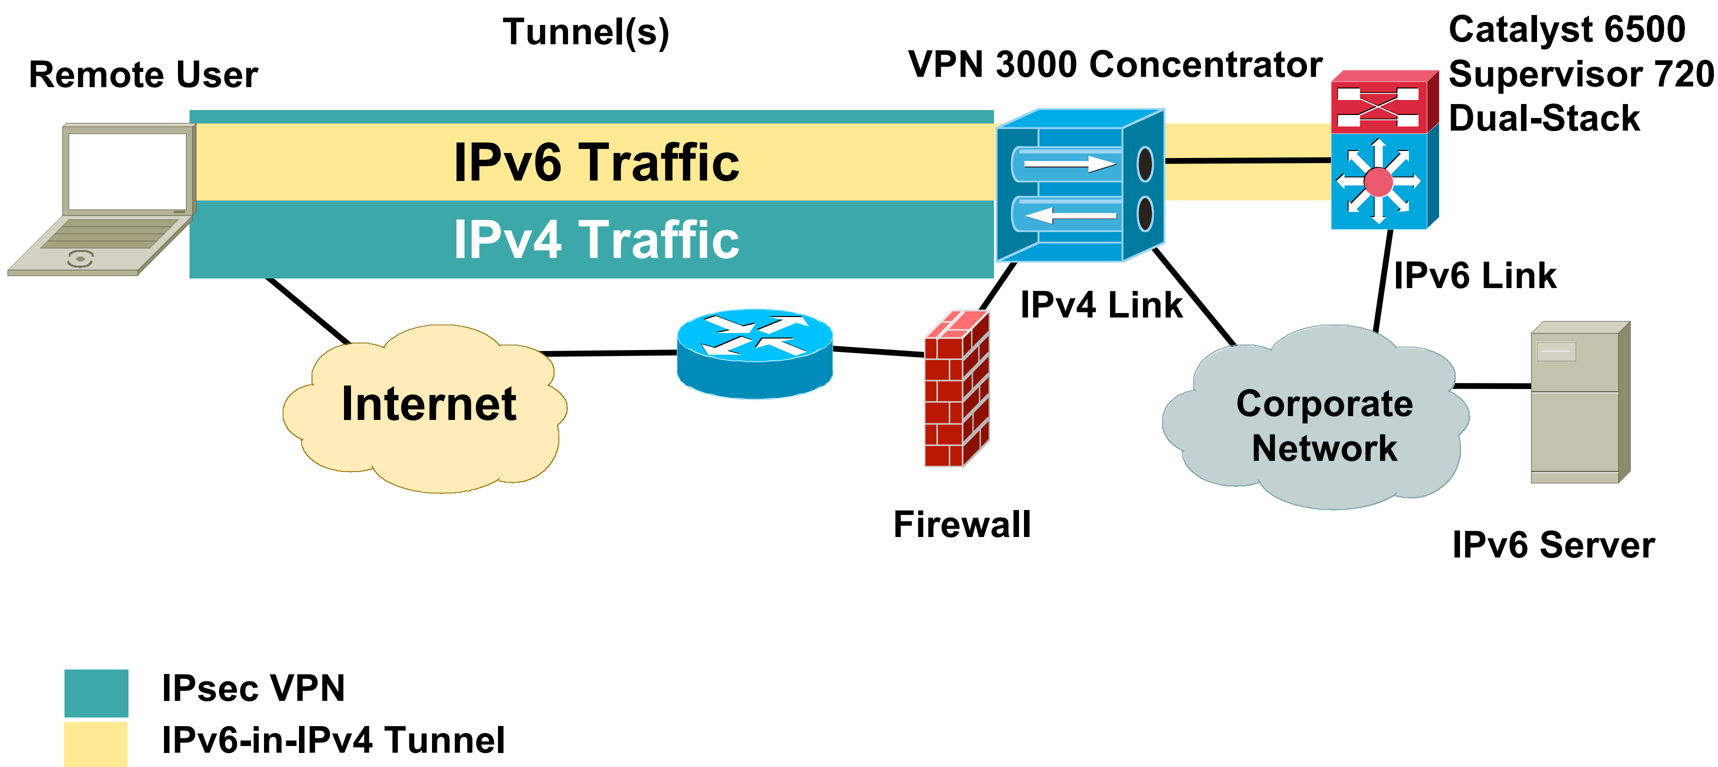

Note: The VPN Concentrator could be replaced with a VPN-enabled Cisco IOS Router or PIX®

Split Tunneling

- Ensure that the IPv6 traffic is properly routed through the IPv4 IPSec tunnel
- IPv6 traffic MAY take a path via the clear (unencrypted) route
- This is bad if YOU ARE UNAWARE THAT IT IS HAPPENING

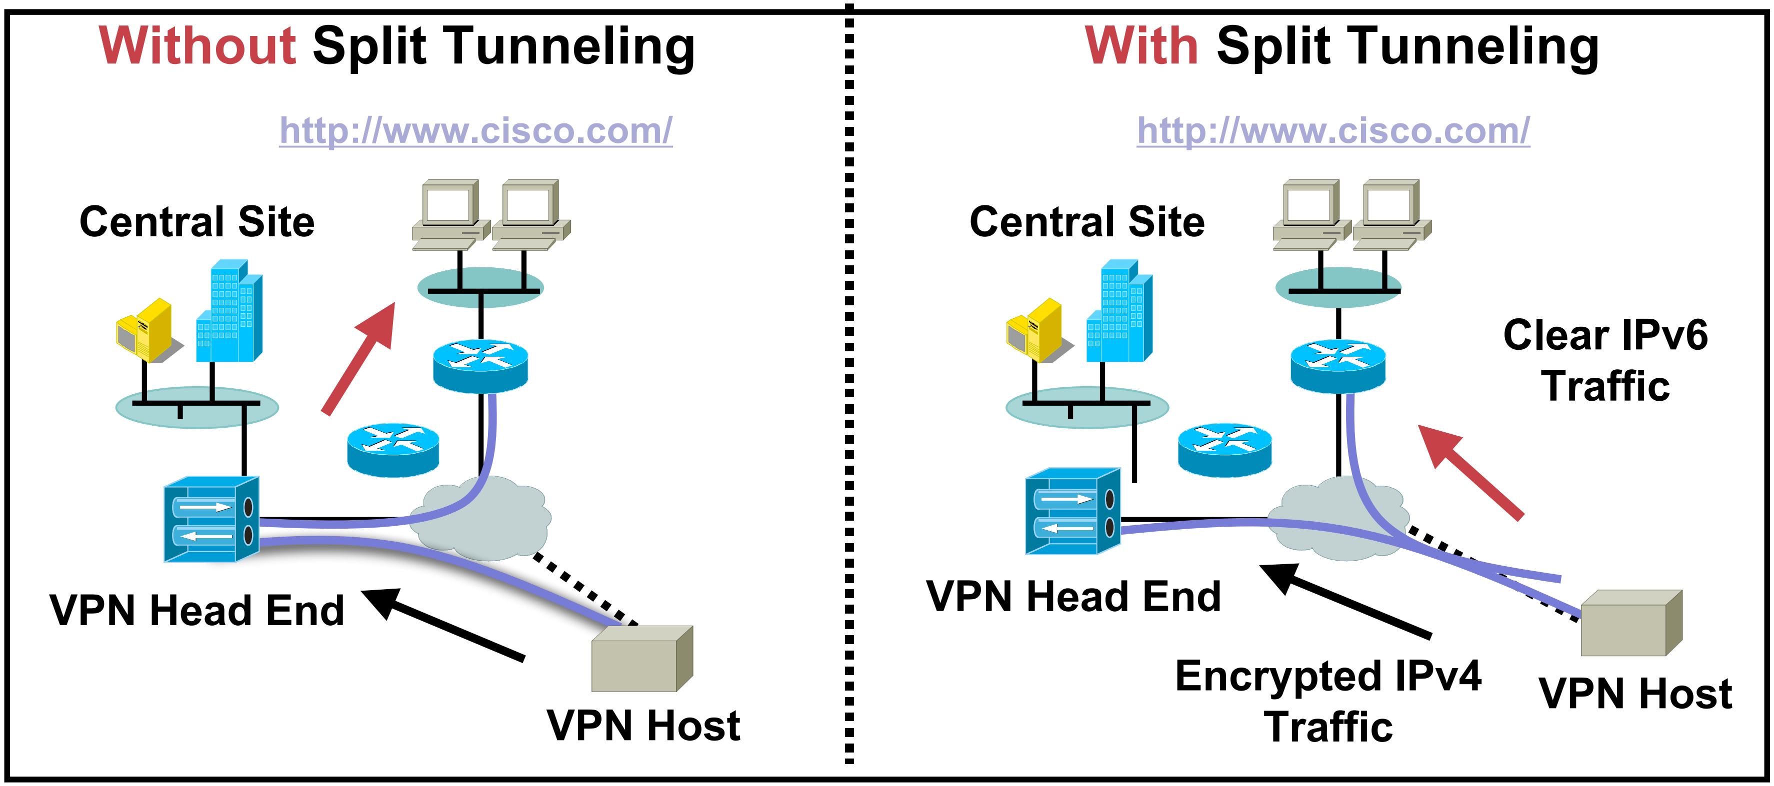

Considerations

- Cisco IOS® version supporting IPv6 configured/ ISATAP tunnels
  - Configured—12.3(1)M/12.3(2)T/12.2(14)S and above (12.4M/12.4T)
  - ISATAP—12.3(1)M, 12.3(2)T, 12.2(14)S and above (12.4M/12.4T)
  - Catalys t® 6500 with Sup720—12.2(17a)SX1—HW forwarding
- Be aware of the security issues if split-tunneling is used
  - Attacker can come in IPv6 interface and jump on the IPv4 interface (encrypted to enterprise)
- Remember that the IPv6 tunneled traffic is still encapsulated as a tunnel WHEN it leaves the VPN device
- Allow IPv6 tunneled traffic across any access lists (Protocol 41)

Required Stuff: Client Side

- **Client operating system with IPv6**
  - Microsoft Windows XP SP1 (Supports Configured/ISATAP)
  - Linux (7.3 or higher)—USAGI port required for ISATAP
  - Mac OS X (10.2 or higher)—Currently need a VPN device on client network
  - SunOS (8 or higher)—Currently need a VPN device on client network
  - See reference slide for links/OS listing

- **Cisco VPN Client 4.0.1 and higher for configured/ISATAP**
- **Cisco VPN Client 3.x for configured ONLY**
- **Cisco HW VPN Client 3002—recommended for Mac/Sun clients until virtual adapter support is available**

IPv6 Using Cisco VPN Client

Example: Client Configuration (Windows XP): ISATAP

- Microsoft Windows XP (SP1 or higher)
- IPv6 must be installed

```
C:\>ipv6 install
```

- XP will automatically attempt to resolve the name "ISATAP"
  - Local host name
  - Hosts file—SystemRoot\system32\drivers\etc
  - DNS name query
  - NetBIOS and Lmhosts
- Manual ISATAP router entry can be made

```
netsh interface ipv6 isatap set router 20.1.1.1
```

- Key fact here is that NO additional configuration on the client is needed again!
- USE PREVIOUS ISATAP CONFIGURATIONS SHOWN FOR ROUTER-SIDE

Note: ISATAP is supported on some versions of Linux/BSD (manual router entry is required)

Does It Work?

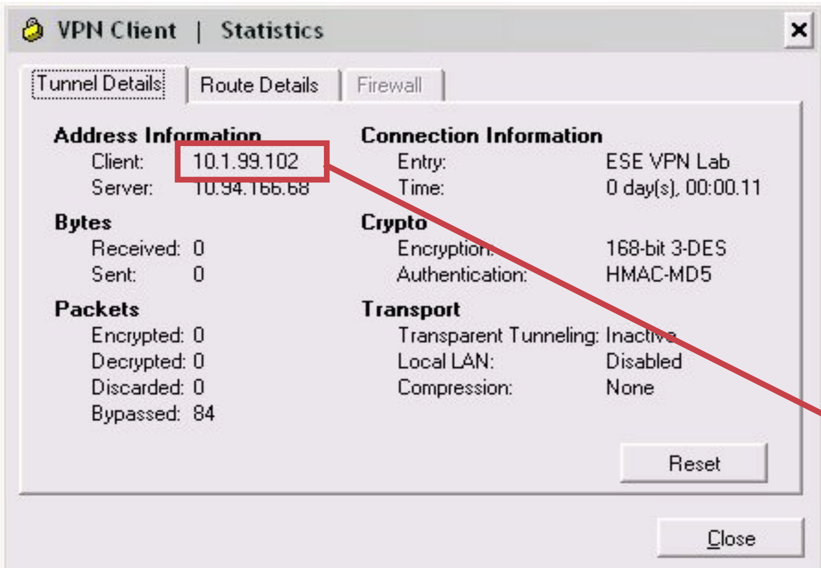

Windows XP Client **VPN 3000** **Catalyst 6500/Sup 720 Dual-Stack**

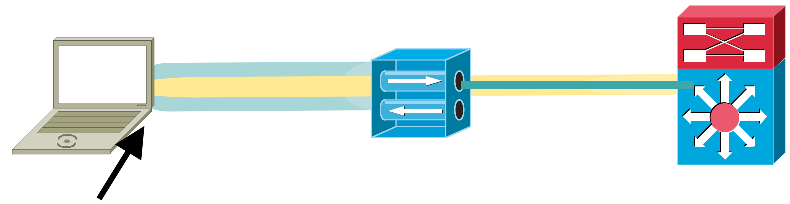

10.1.99.102—VPN Address

2001:DB8:c003:1101:0:5efe:10.1.99.102—IPv6 address

```
Interface 2: Automatic Tunneling Pseudo-Interface

Addr Type  DAD State  Valid Life   Pref. Life   Address
---------  ---------- ------------ ------------ -----------------------------
Public     Preferred   29d23h56m5s  6d23h56m5s 2001:db8:c003:1101:0:5efe:10.1.99.102
Link       Preferred   infinite     infinite   fe80::5efe:10.1.99.102
```

```
netsh interface ipv6>show route
Querying active state...

Publish  Type       Met  Prefix                   Idx  Gateway/Interface Name
-------  --------  ----  ------------------------ ---  ---------------------
no       Autoconf     9  2001:db8:c003:1101::/64    2  Automatic Tunneling Pseudo-Interface
no       Manual       1  ::/0                       2  fe80::5efe:20.1.1.1
```

SERVICE PROVIDER DEPLOYMENT

Start Here: Cisco IOS Software Release Specifics for IPv6 Features
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/ftipv6s.htm

IPv6 in the SP: What Does It Do for Me?

- **Benefits for the ISP (short term):**
  - Expanded private use address pool for internal devices
  - Ability to acquire large enough address blocks to avoid impeding rollout/subscriber-growth business plans
  - Not lose existing or new customers due to lack of support
- **Benefits for the ISP (long term):**
  - Reduction in 'application failure' related support calls caused by IPv4/NAT
  - Ability to remove customer-managed infrastructure component (NAT) from the path, improving application support
  - Ability to deploy new service offerings into the home without dealing with translation issues and address constraints

Today's Network Infrastructure

- Service Providers core infrastructure are basically following two paths
  - MPLS with its associated services
    - MPLS/VPN, L2 services over MPLS, QoS,…
  - Native IPv4 core with associated services
    - L2TPv3, QoS, Multicast,…
- IP services portfolio—Access
  - Enterprise: Lease lines
  - Home Users/SOHO: ADSL, FTTH, Dial
  - Data Center: Web hosting, servers,…
- Next step—The integration of IPv6 services

Note: Don't classify IPv6 tunneled traffic as "undetermined" (Protocol 41)

IPv6 Tunnels & Native IPv6

- ISP scenario

  Configured Tunnels or Native IPv6 between IPv6 Core Routers

  Configured Tunnels or Native IPv6 to IPv6 Enterprise's Customers

  MP-BGP4 Peering with other 6Bone users

  Connection to an IPv6 IX

  6to4 relay service

- Enterprise/Home scenario

  6to4 tunnels between sites, use 6to4 Relay to connect to the IPv6 Internet

  Configured tunnels between sites or to 6Bone users

  ISATAP tunnels or Native IPv6 on a Campus

Use the most appropriate

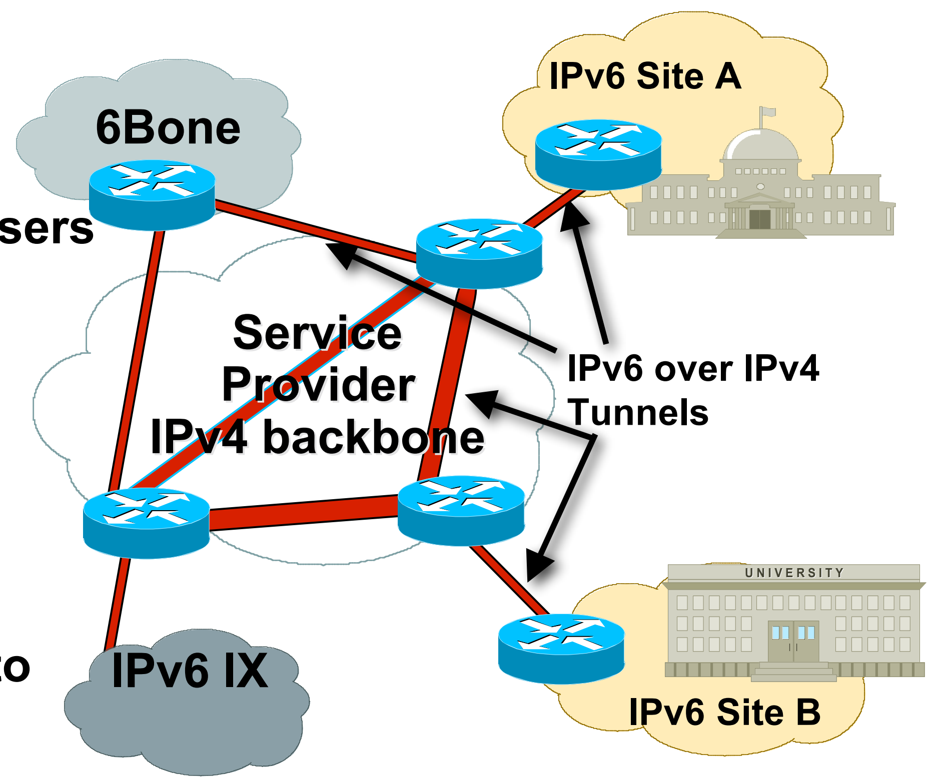

Native IPv6 over Dedicated Data Link

- ISP scenario

  Dedicated Data Links between Core routers

  Dedicated Data Links to IPv6 Customers

  Connection to an IPv6 IX

- Enterprise scenario

  Experimental LAN segment, eg. Dedicated Ethernet or VLAN

  Between Campus over a MAN Infrastructure

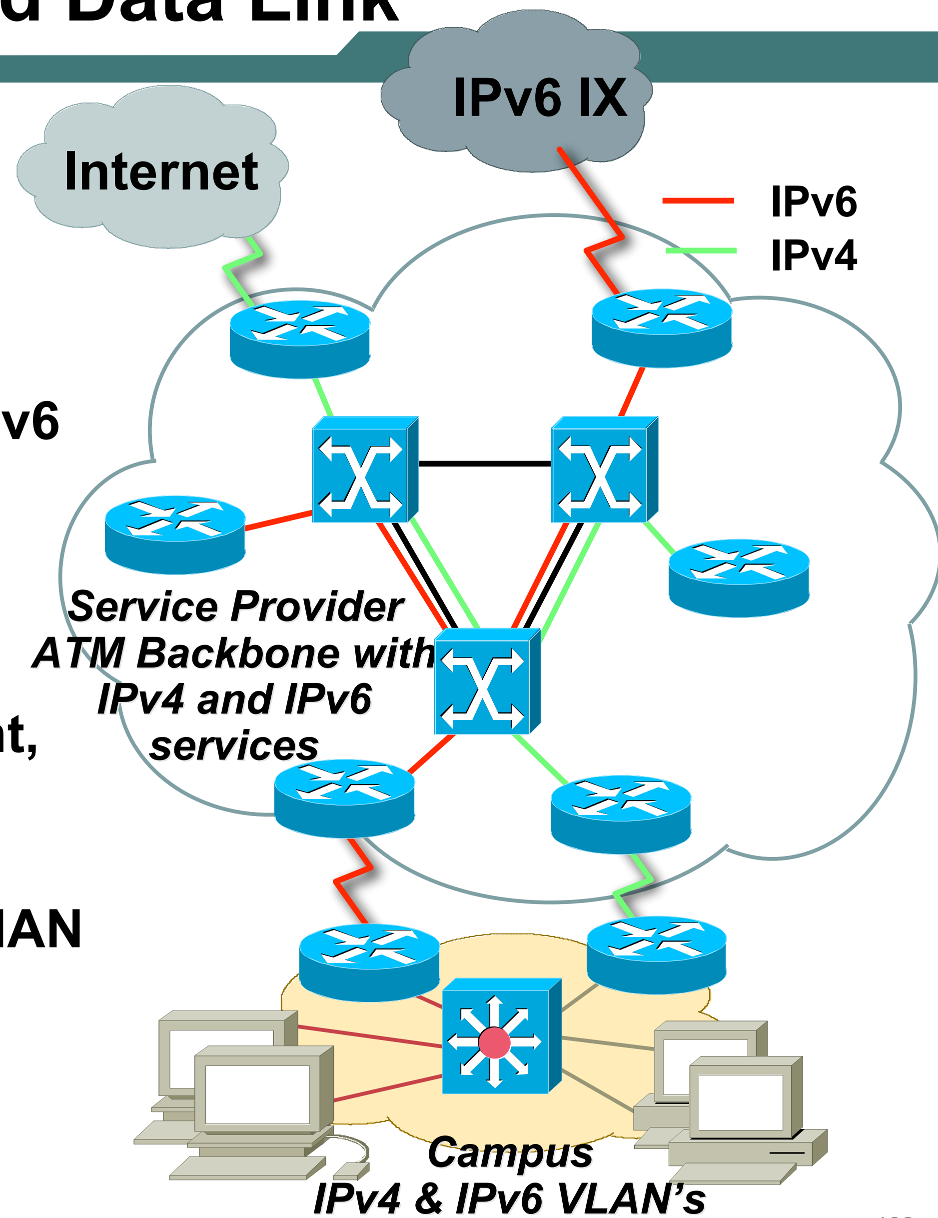

Dual-Stack IPv4-IPv6

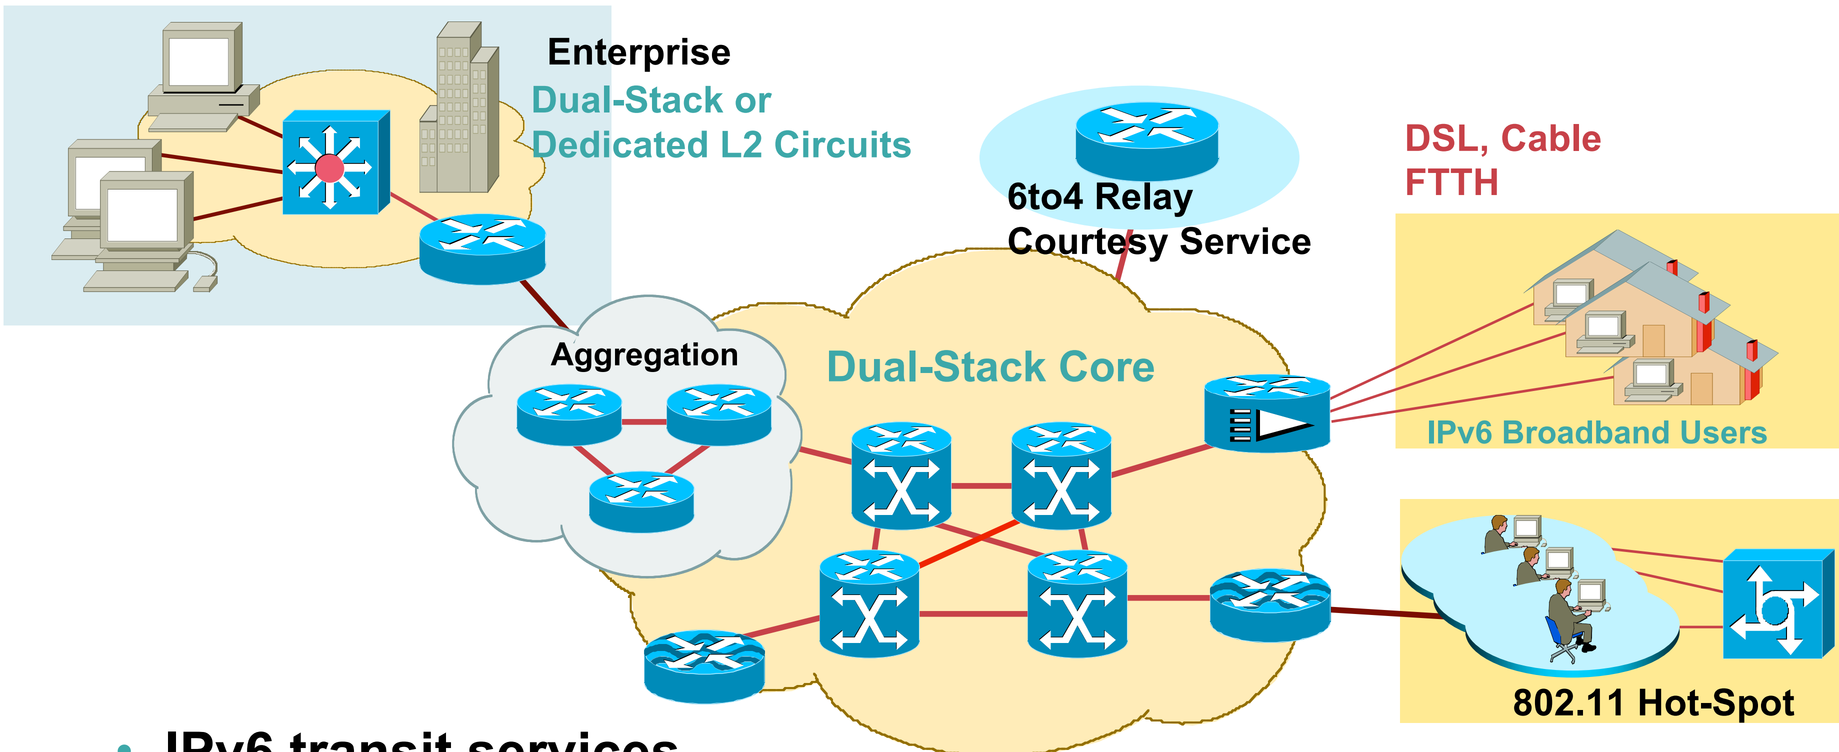

- IPv6 transit services
- IPv6 enabled on Core routers
- Enterprise and consumer IPv6 access
- Additional services
  - IPv6 multicast for streaming

IPv6 over MPLS

- Many service providers have already deployed MPLS in their IPv4 backbone for various reasons
- MPLS can be used to facilitate IPv6 integration
- Multiple approaches for IPv6 over MPLS:

  IPv6 over L2TPv3

  IPv6 over EoMPLS/AToM

  IPv6 CE-to-CE IPv6 over IPv4 Tunnels

  IPv6 Provider Edge Router (6PE) over MPLS

  IPv6 VPN Provider Edge (6VPE) over MPLS

  Native IPv6 over MPLS

IPv6 Tunnels configured on CE

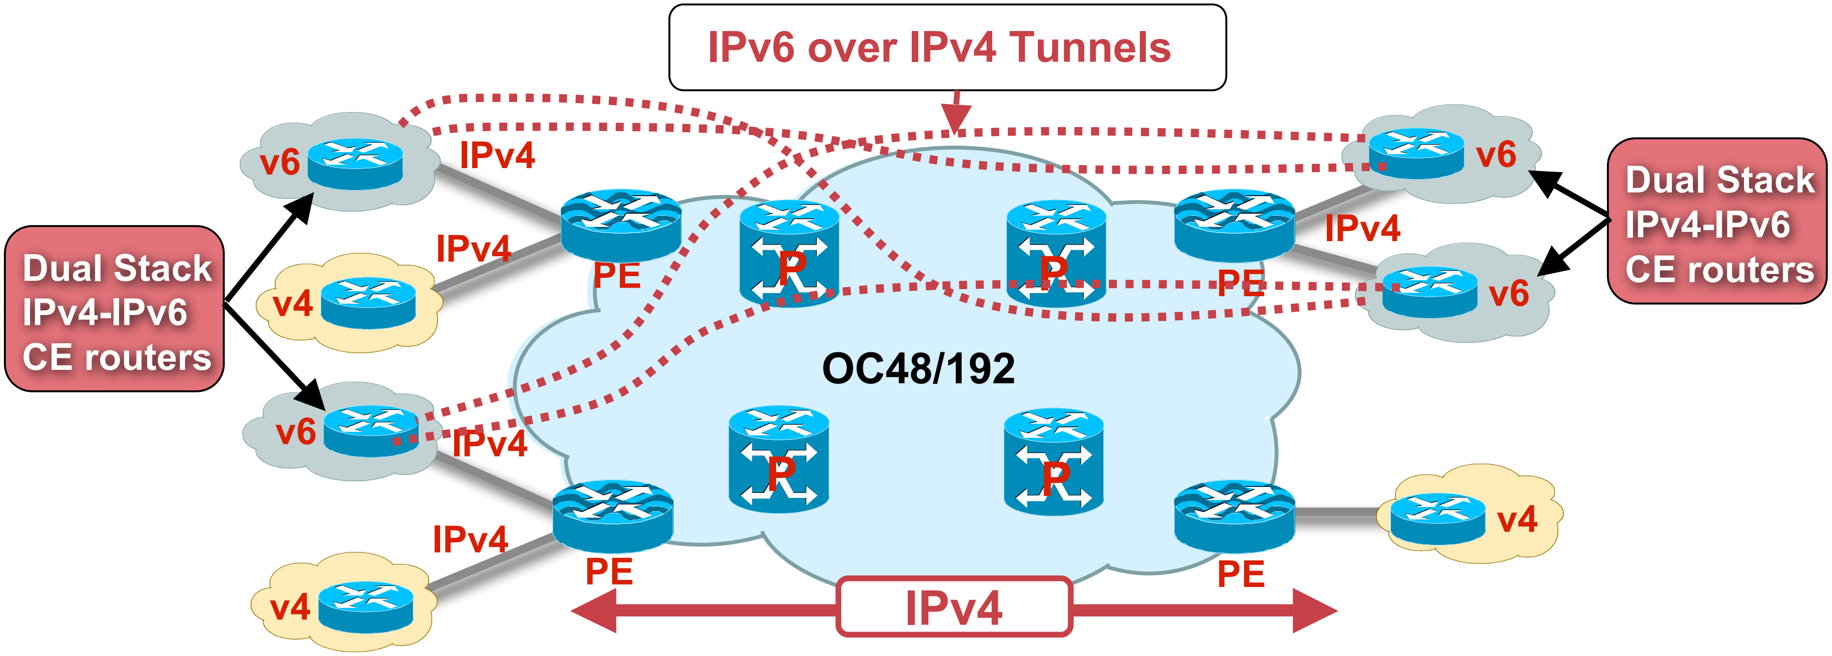

- No impact on existing IPv4 or MPLS Core (IPv6 unaware)
- Only CEs have to be IPv6-aware (Dual stack)
- Mesh of IPv6 over IPv4 Tunnels CE-to-CE
- Overhead: IPv4 header + MPLS header
- MPLS/VPN support IPv4-native and IPv6 tunnels
- Service Provider can’t delegate his IPv6 prefix to the CE routers

IPv6 over “Circuit_over_MPLS”

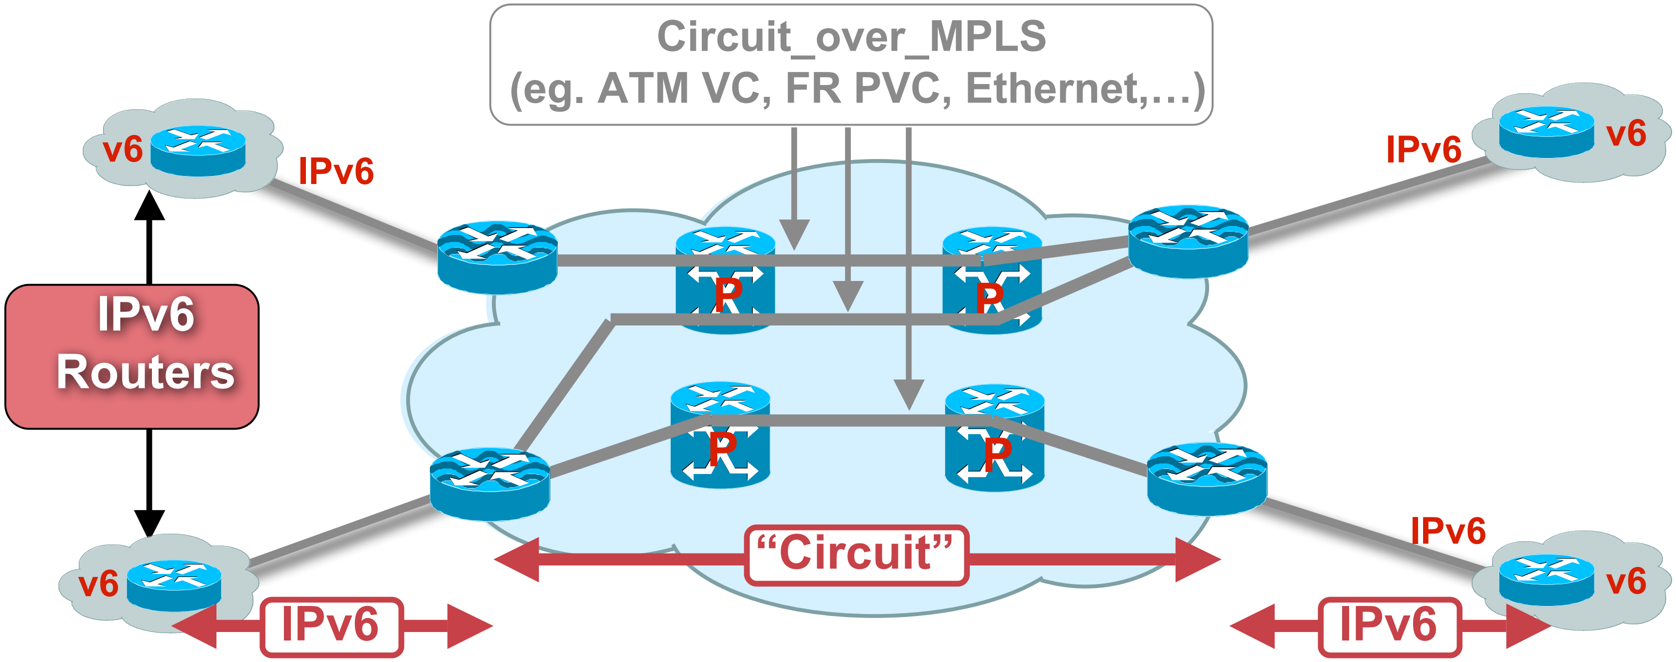

- No impact on existing IPv4 or MPLS Core (IPv6 unaware)
- Edge MPLS Routers need to support “Circuit_over_MPLS” (AToM)
- Mesh of “Circuit_Over_MPLS” PE-to-PE
- PE routers (IPv6 over ATM, IPv6 over FR, IPv6 over Ethernet,…) to aggregate Customer’s IPv6 routers

Native MPLS Support of IPv6

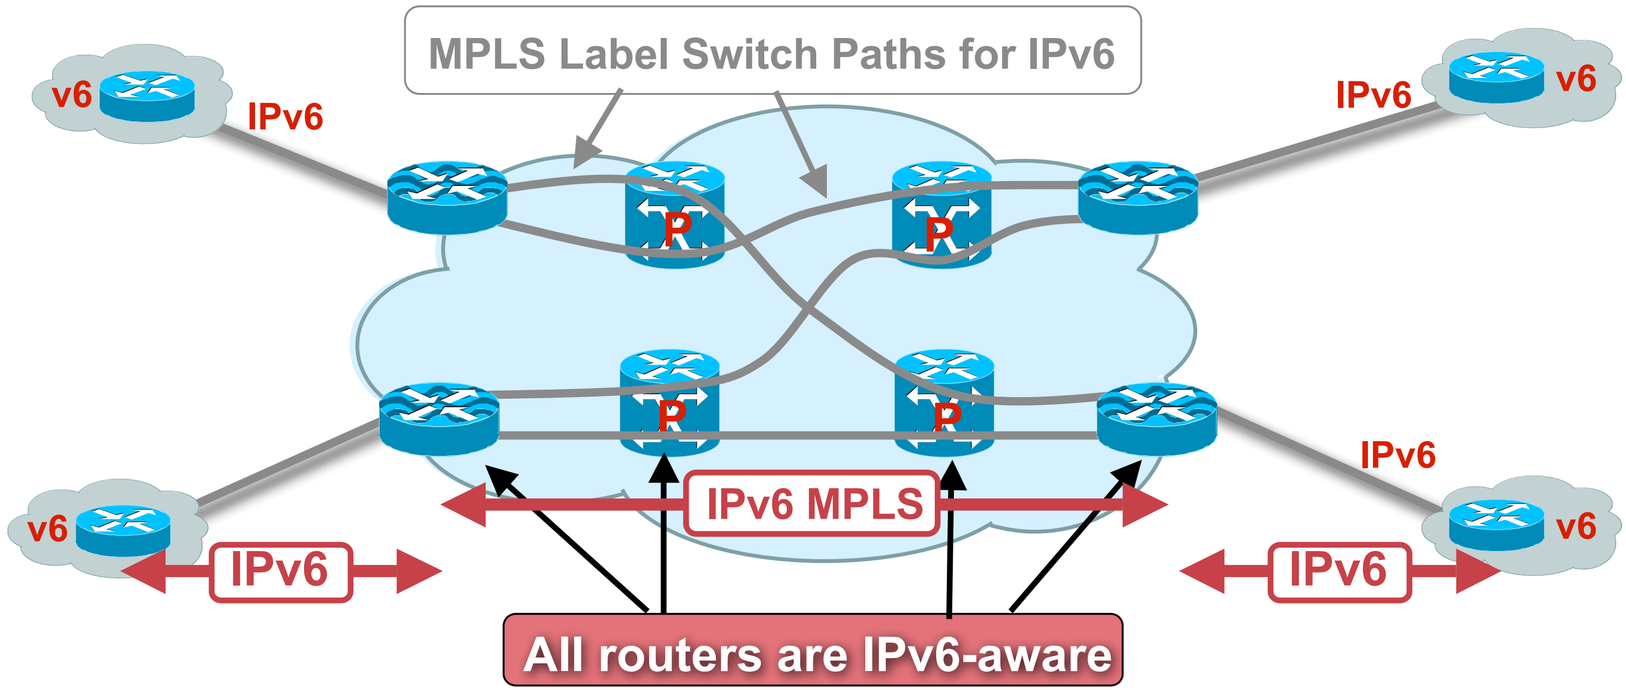

- Core Infrastructure requires full Control Plane upgrade to IPv6
  - IPv6 Routing in core
  - IPv6 Label Distribution Protocol in core
- Dual Control Plane management if IPv4 and IPv6 services

6PE Overview

IPv6 Provider Edge Router (6PE) over MPLS

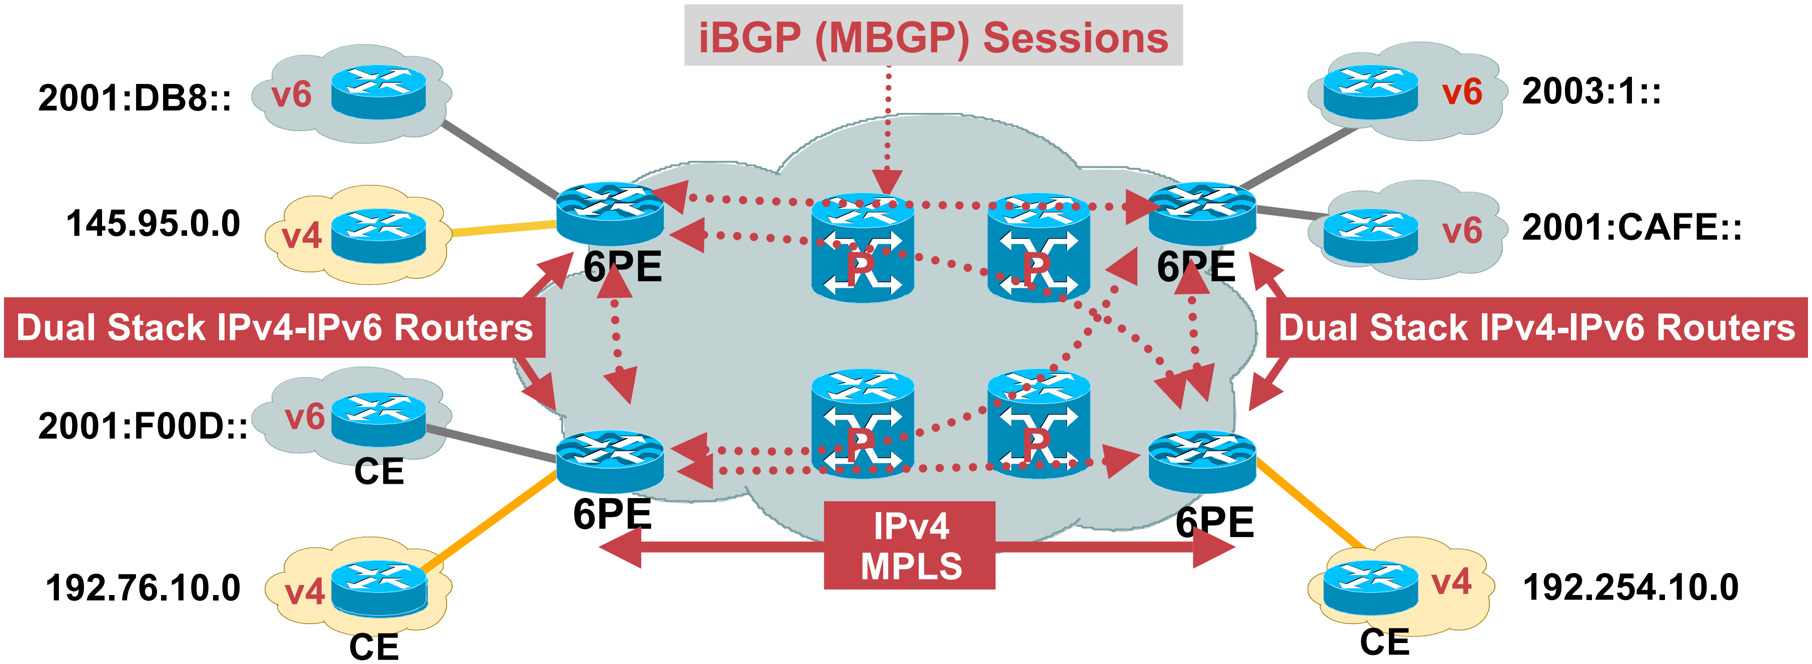

- IPv6 global connectivity over and IPv4-MPLS core
- Transitioning mechanism for providing unicast IP
- PEs are updated to support dual stack/6PE
- IPv6 reachability exchanged among 6PEs via iBGP (MBGP)
- IPv6 packets transported from 6PE to 6PE inside MPLS

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/iosip_an.htm

6PE Routing/Label Distribution

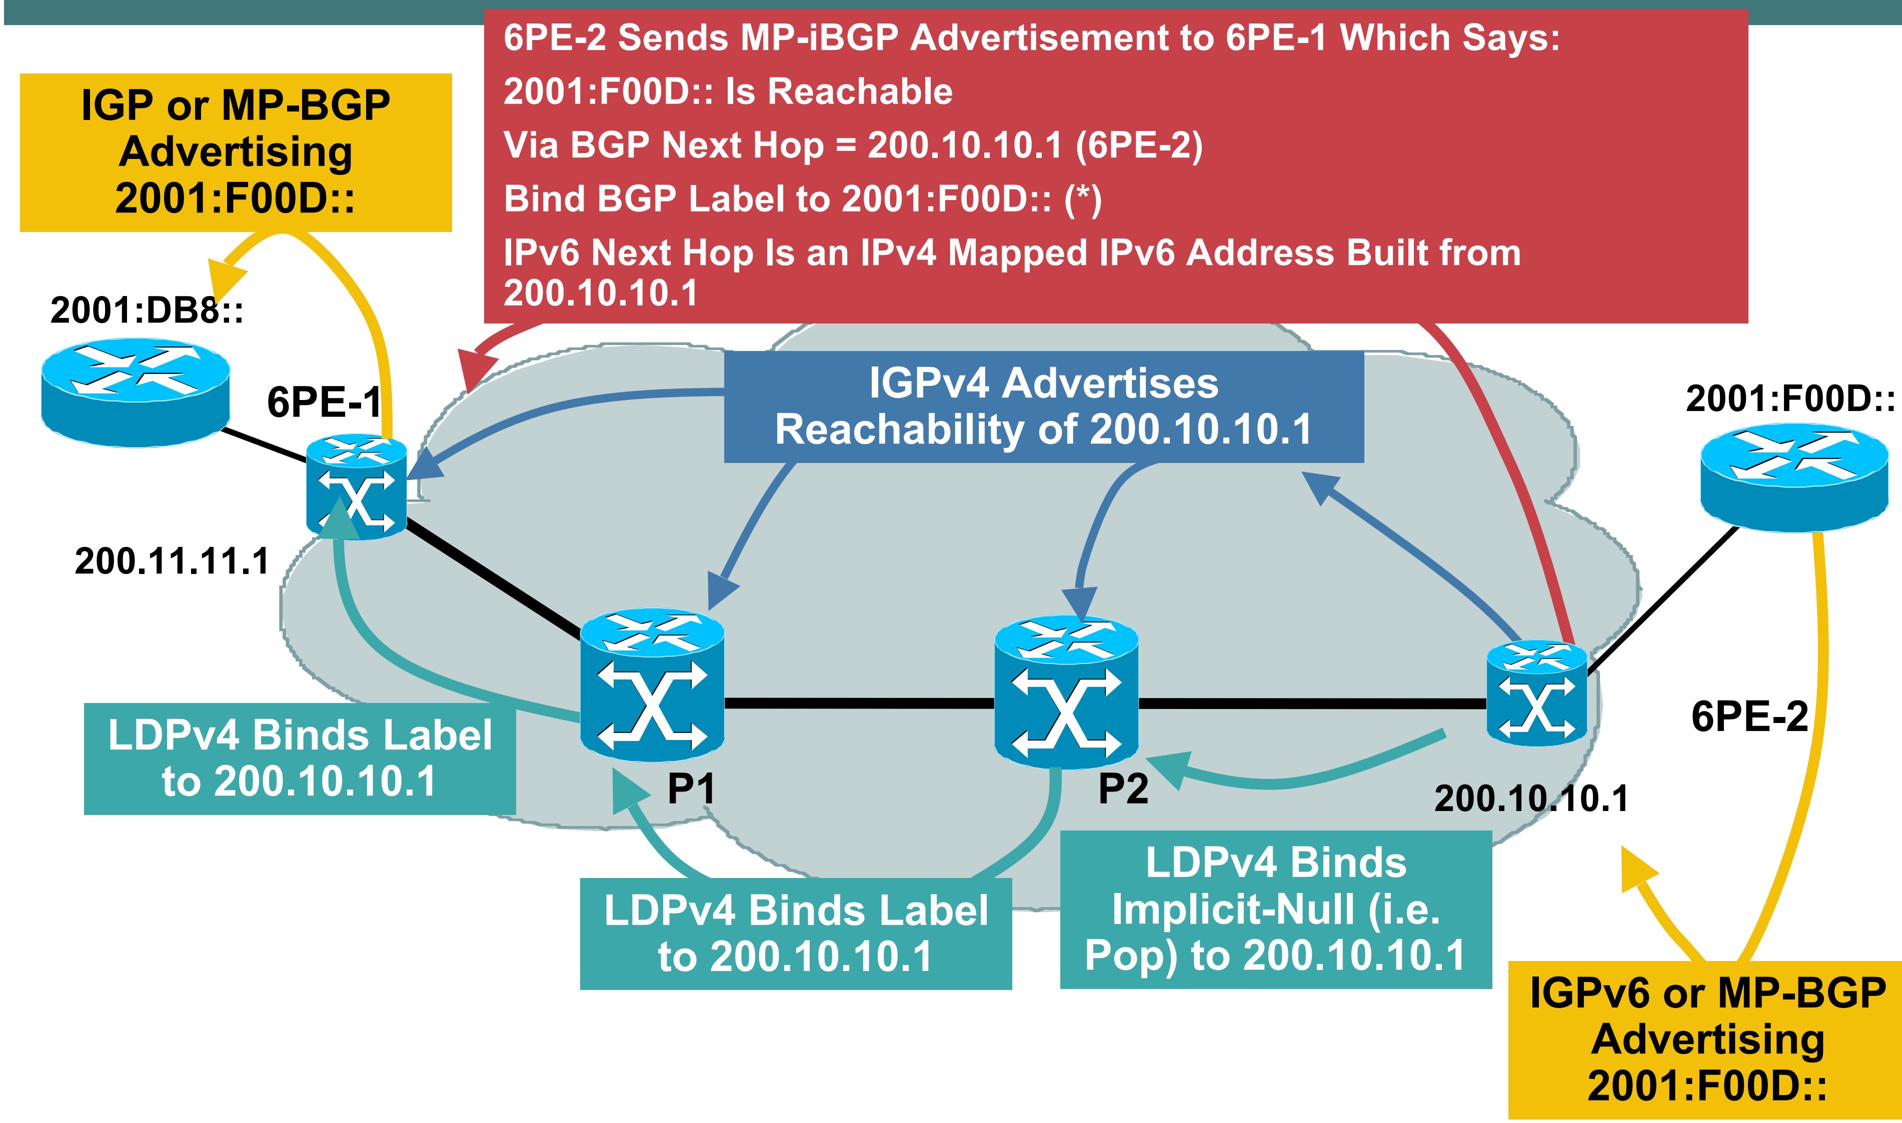

6PE Forwarding (6PE-1)

IPv6 Forwarding and Label Imposition:

- 6PE-1 receives an IPv6 packet
- Lookup is done on IPv6 prefix
- Result is:
  - Label binded by MP-BGP to 2001:F00D::
  - Label1 binded by LDP/IGPv4 to the IPv4 address of BGP next hop (6PE-2)

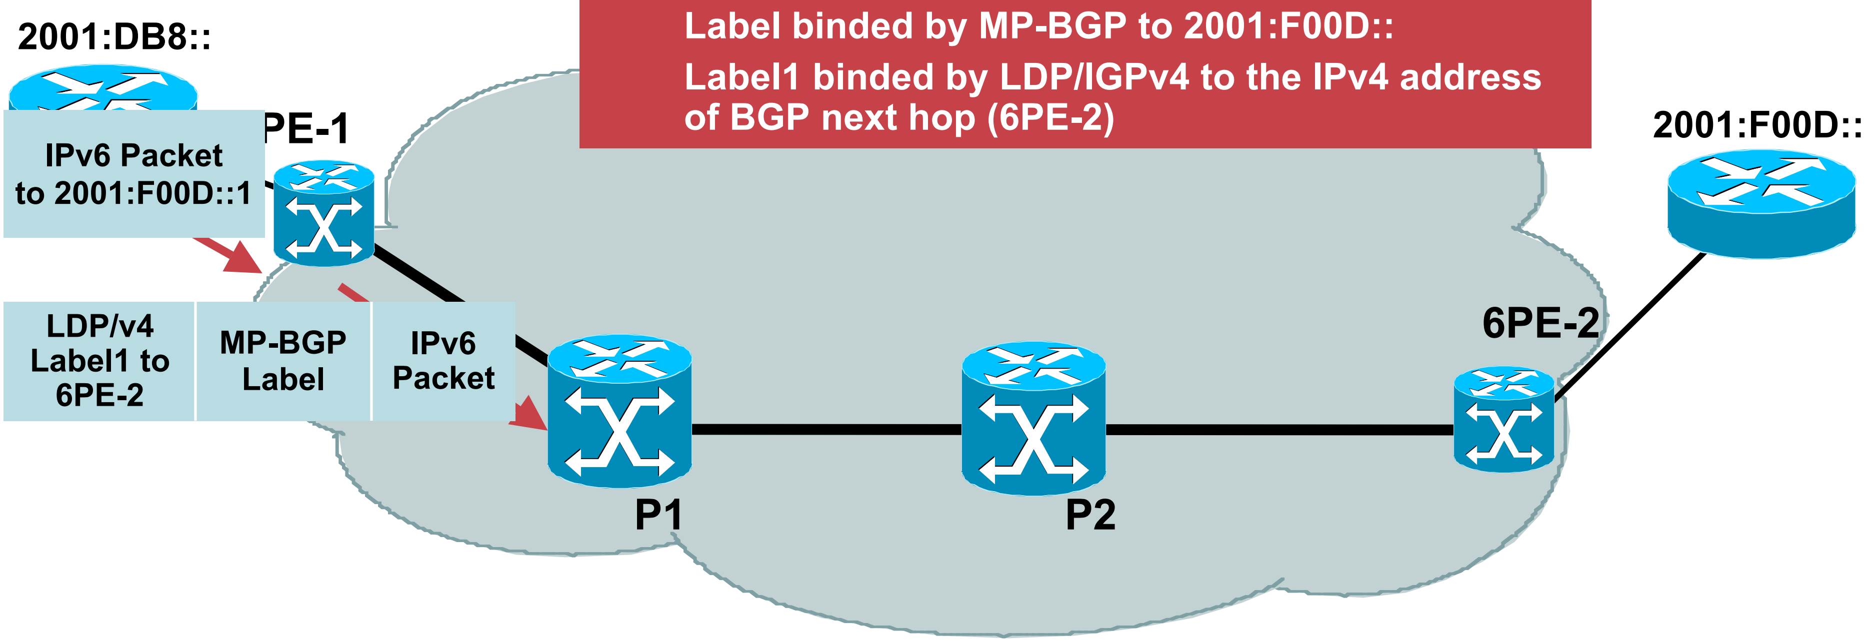

6PE Forwarding (P1)

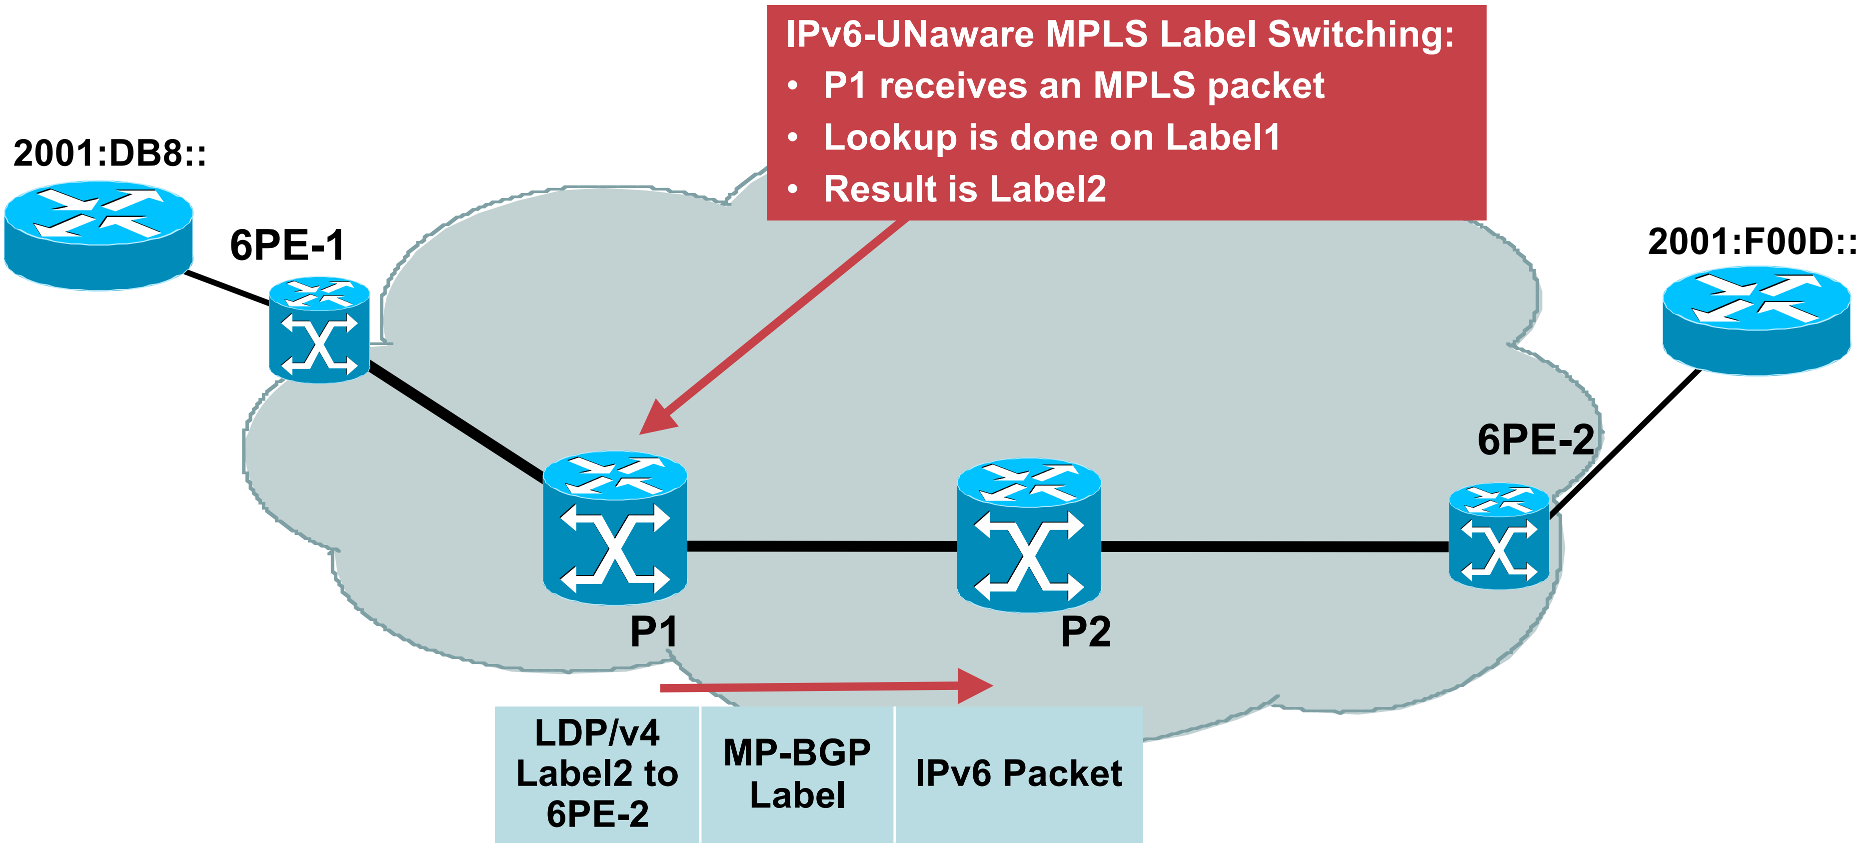

6PE Forwarding (P2)

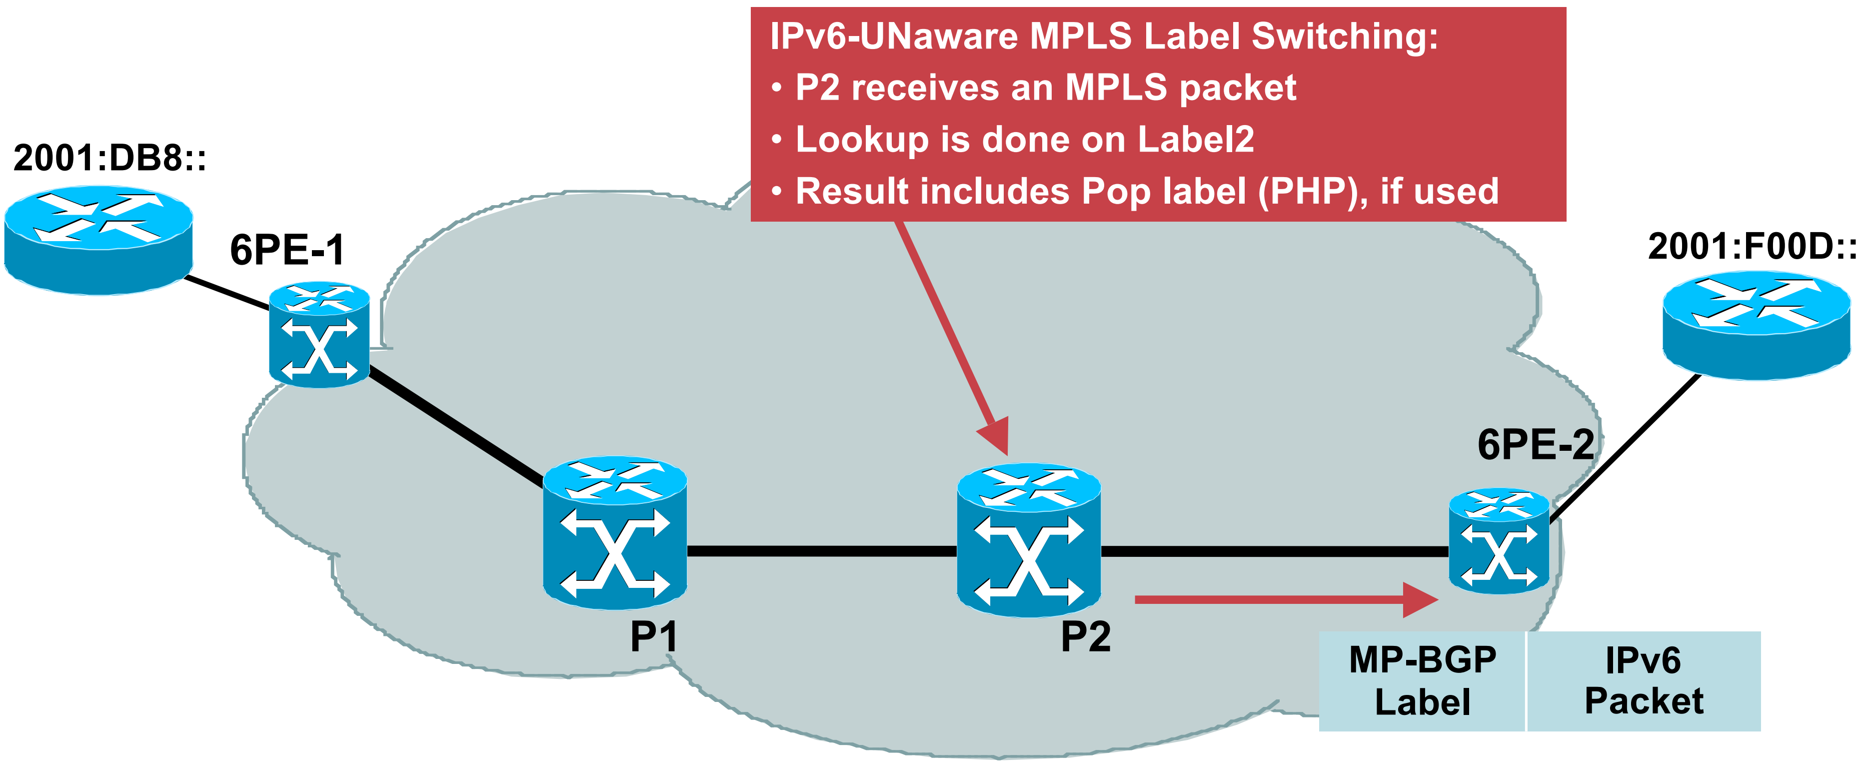

6PE Forwarding (6PE-2)

- MPLS label forwarding:
- 6PE-2 receives an MPLS packet
- Lookup is done on label
- Result is:
  Pop label and do IPv6 lookup on v6 destination

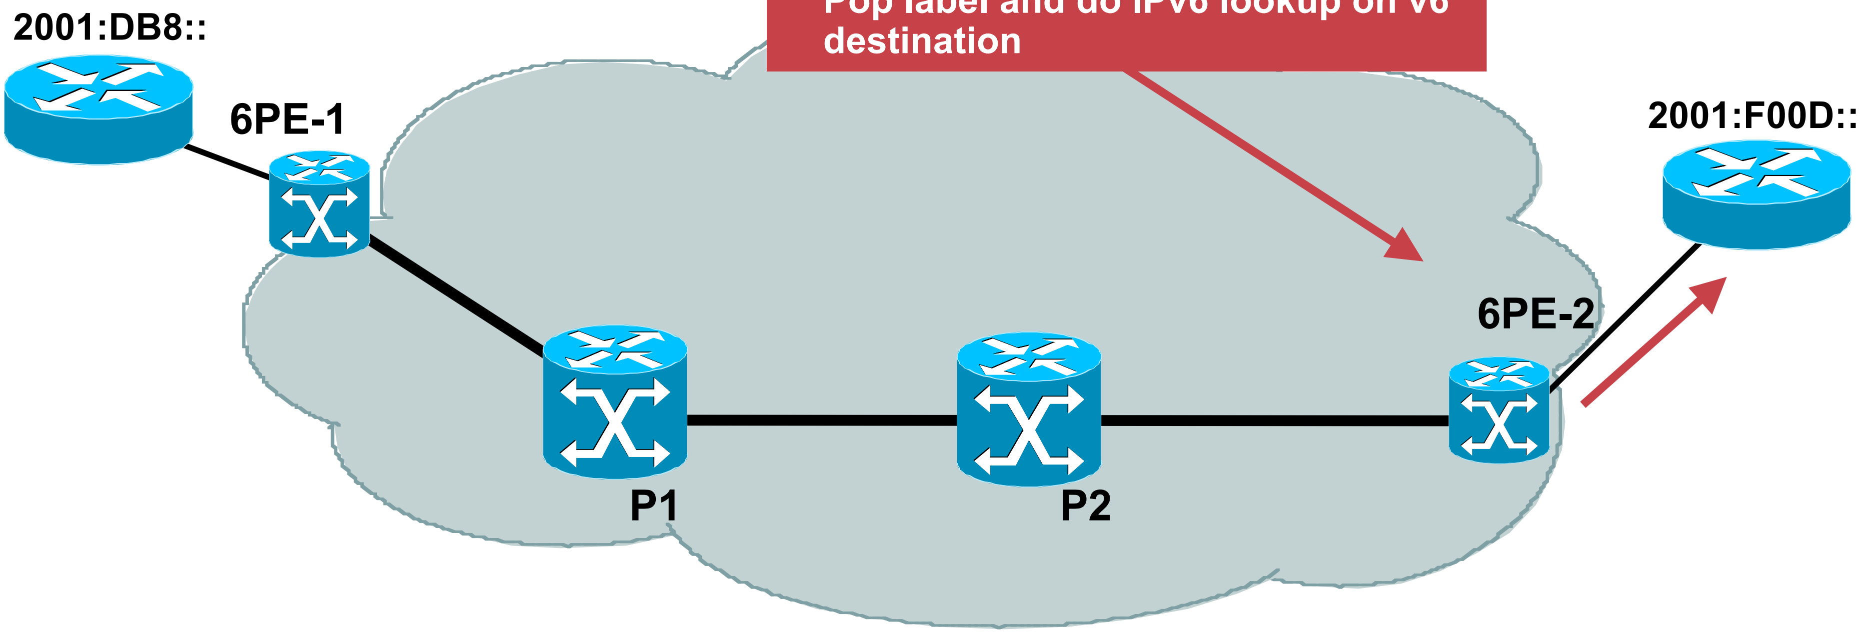

6PE-1 Configuration

```
ipv6 cef
!
mpls label protocol ldp
!
router bgp 100
 no synchronization
 no bgp default ipv4 unicast
 neighbor 2001:DB8:1::1 remote-as 65014
 neighbor 200.10.10.1 remote-as 100
 neighbor 200.10.10.1 update-source Loopback0
 !
 address-family ipv6
 neighbor 200.10.10.1 activate
 neighbor 200.10.10.1 send-label
 neighbor 2001:DB8:1::1 activate
 redistribute connected
 no synchronization
 exit-address-family
```

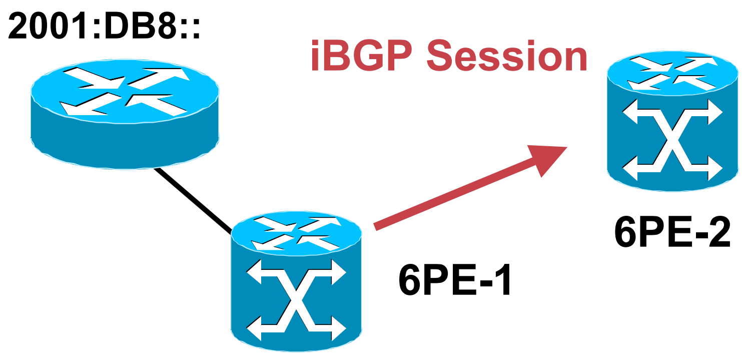

2001:DB8:1::1 Is the Local CE

200.10.10.1 Is the Remote 6PE

Send Labels Along with IPv6 Prefixes by Means of MP-BGP Note: Will Cause Session to Flap

6PE Show Output

```
6PE-1#show ip route 200.10.10.1
Routing entry for 200.10.10.1/32
 Known via "isis", distance 115, metric 20, type level-2
[snip]
  * 10.12.0.1, from 200.10.10.1, via FastEthernet1/0
   Route metric is 20, traffic share count is 1
```

```
6PE-1#show ipv6 route
B  2001:F00D::/64 [200/0]
   via ::FFFF:200.10.10.1, IPv6-mpls
```

```
6PE-1#show ipv6 cef internal #hidden command
.. OUTPUT TRUNCATED ..
2001:F00D::/64,
  nexthop ::FFFF:200.10.10.1
 fast tag rewrite with F0/1, 10.12.0.1, tags imposed {17 28}
```

Other Useful Output:

```
show bgp ipv6 neighbors
show bgp ipv6 unicast
show mpls forwarding #more on this later
```

6PE Benefits/Drawbacks

- Core network (Ps) untouched (no HW/SW upgrade, no configuration change)
- IPv6 traffic inherits MPLS benefits (wire-rate, fast re-route, TE, etc.)
- Incremental deployment possible (i.e., only upgrade the PE routers which have to provide IPv6 connectivity)
- Each site can be v4-only, v4VPN-only, v4+v6, v4VPN+v6
- P routers won't be able to send ICMP messages (TTL expired, traceroute)

Application Note—IPv6 over MPLS (Cisco® 6PE)

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/iosip_an.htm

"IPv6 Over MPLS" presentation:

http://www.cisco.com/warp/public/732/Tech/ipv6/docs/IPV6overMPLS.pdf

6VPE Overview

6VPE Deployment

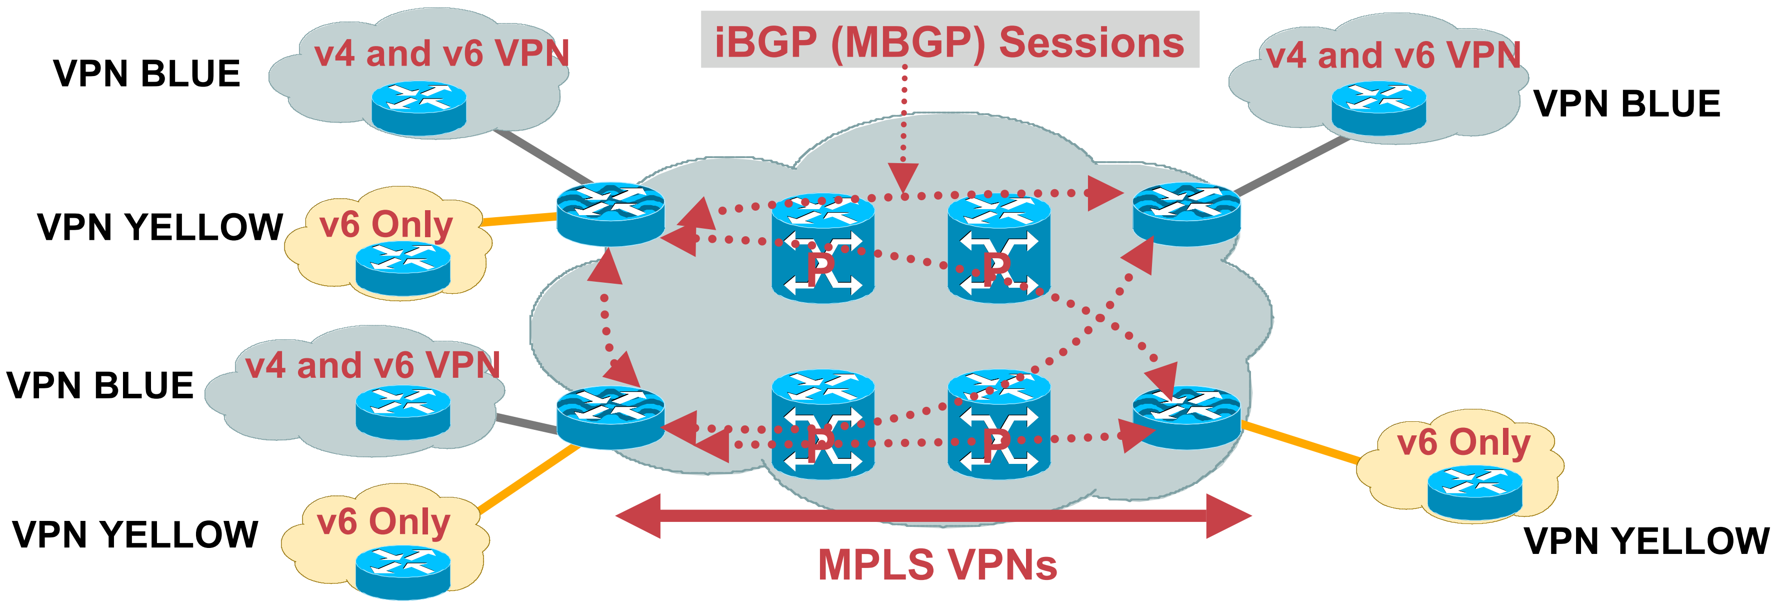

- 6VPE ~ IPv6 + BGP-MPLS IPv4 VPN + 6PE
- Cisco 6VPE is an implementation of <draft-ietf-bgp-ipv6-vpn> over MPLS/IPv4
- VPNv6 address:
  - Address including the 64 bits route distinguisher and the 128 bits IPv6 address
- MP-BGP VPNv6 address-family:
  - AFI “IPv6” (2), SAFI “VPN” (128)
- VPN IPv6 MP_REACH_NLRI
  - With VPNv6 next-hop (192bits) and NLRI in the form of <length, IPv6-prefix, label>
- Encoding of the BGP next-hop

6VPE Example Design
Addressing/Routing

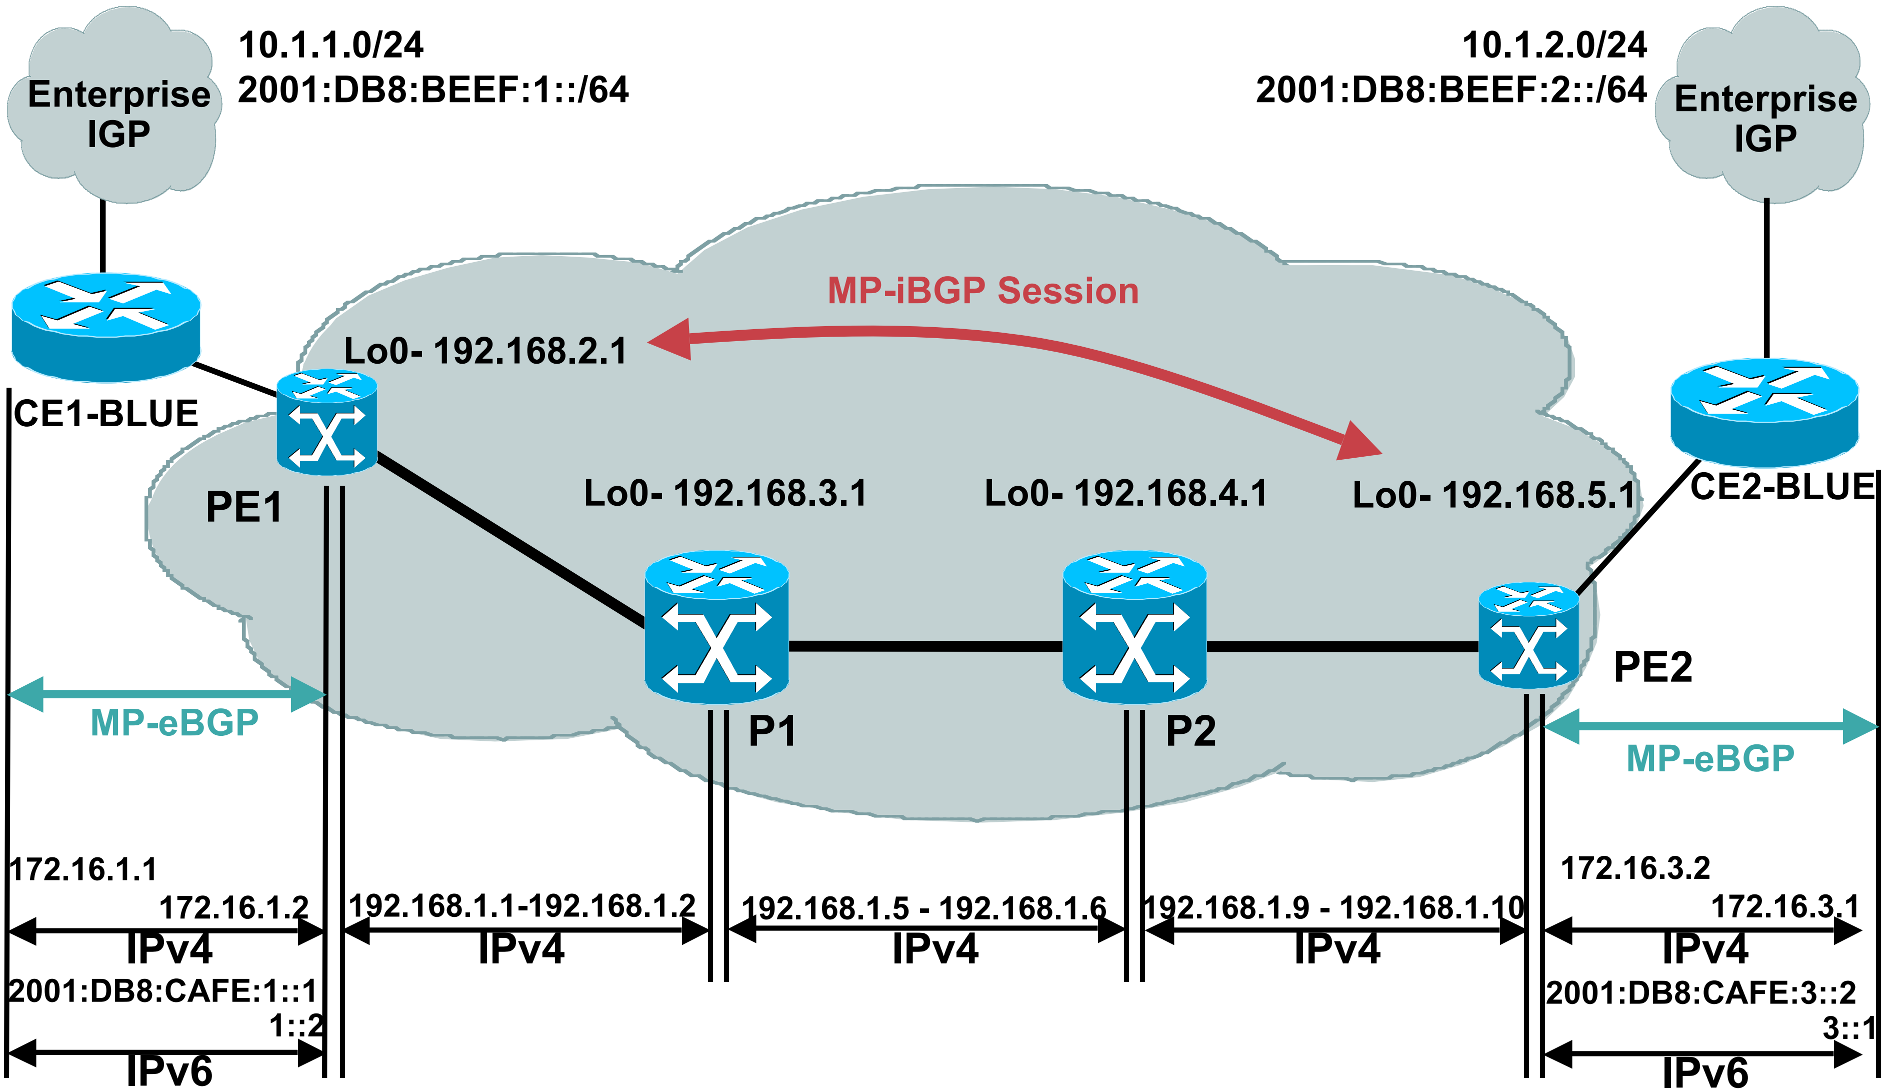

6VPE Configuration Example
CE1-BLUE to PE1

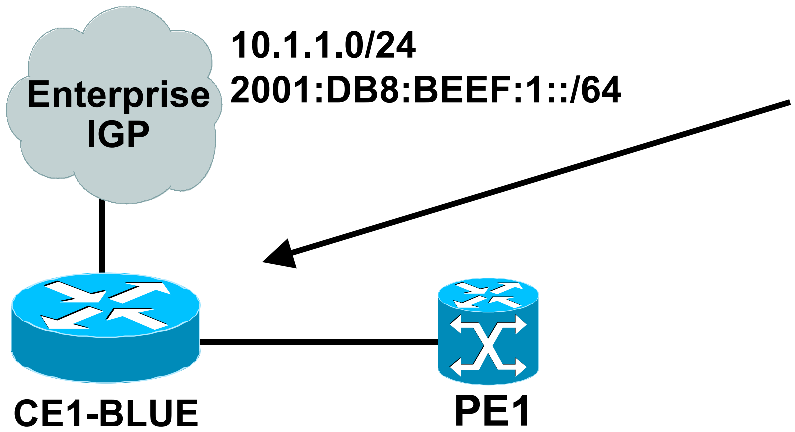


```
ipv6 unicast-routing
ipv6 cef
!
interface Ethernet0/0
 description to PE1
 ip address 172.16.1.1 255.255.255.0
 ipv6 address 2001:DB8:CAFE:1::1/64
!
interface Ethernet1/0
 description to BLUE LAN
 ip address 10.1.1.1 255.255.255.0
 ipv6 address 2001:DB8:BEEF:1::1/64
 ipv6 rip BLUE enable
```

```
router bgp 500
 bgp log-neighbor-changes
 neighbor 2001:DB8:CAFE:1::2 remote-as 100
 neighbor 172.16.1.2 remote-as 100
 !
 address-family ipv4
 redistribute connected
 redistribute eigrp 100
 no neighbor 2001:DB8:CAFE:1::2 activate
 neighbor 172.16.1.2 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv6
 neighbor 2001:DB8:CAFE:1::2 activate
 redistribute connected
 redistribute rip BLUE
 no synchronization
 exit-address-family
!
ipv6 router rip BLUE
 redistribute bgp 500
```

6VPE Configuration Example

PE1 Connections

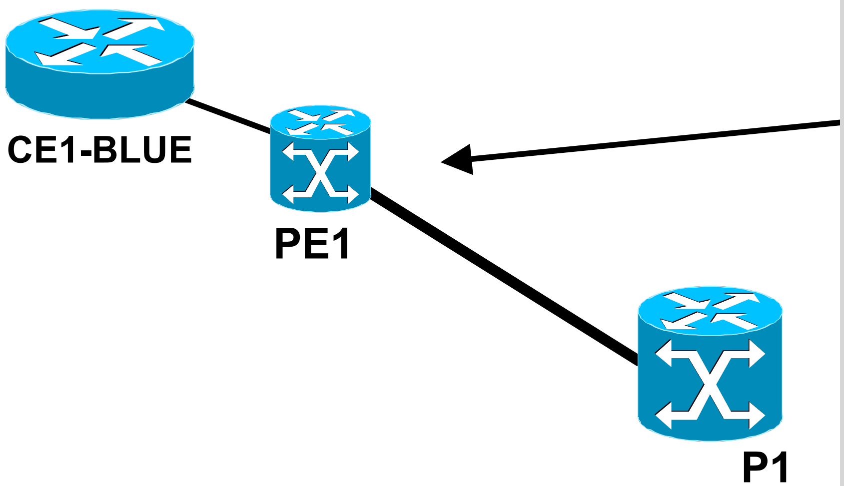

- **Standard MPLS configuration between PE-P**
- **Running IGP in the cloud (OSPF)**

```
ipv6 unicast-routing
ipv6 cef
mpls ldp router-id Loopback0
!
interface Loopback0
 ip address 192.168.2.1 255.255.255.255
!
interface Ethernet0/0
 description to CE1-BLUE
 vrf forwarding BLUE
 ip address 172.16.1.2 255.255.255.0
 ipv6 address 2001:DB8:CAFE:1::2/64
!
interface Ethernet2/0
 description to P1
 ip address 192.168.1.1 255.255.255.252
 mpls ip
!
router ospf 1
 log-adjacency-changes
 redistribute connected subnets
 passive-interface Loopback0
 network 192.168.1.0 0.0.0.255 area 0
```

6VPE Configuration Example

PE1 VRF Definitions

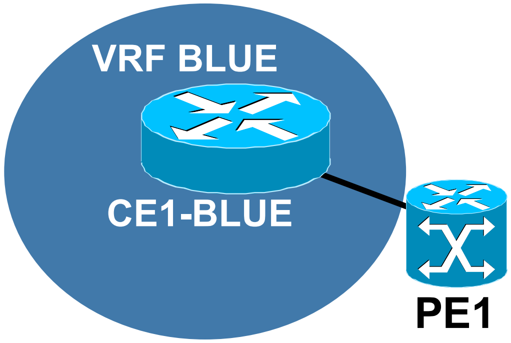

VRF
BLUE
(RD - 200:1)

PE1/PE2 Will Hold

CE1 Routes
CE2 Routes

```
vrf definition BLUE
 rd 200:1
 !
 route-target export 200:1
 route-target import 200:1
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
```

- **Migration commands available for VPNv4 to multi-protocol VRF**

```
(config)#vrf upgrade-cli multi-af-mode {common-
 policies | non-common-policies} [vrf <name>]
```

- **This command forces migration from old CLI for IPv4 VRF to new VRF multi-AF CLI**

6VPE Configuration Example

PE1 BGP Setup

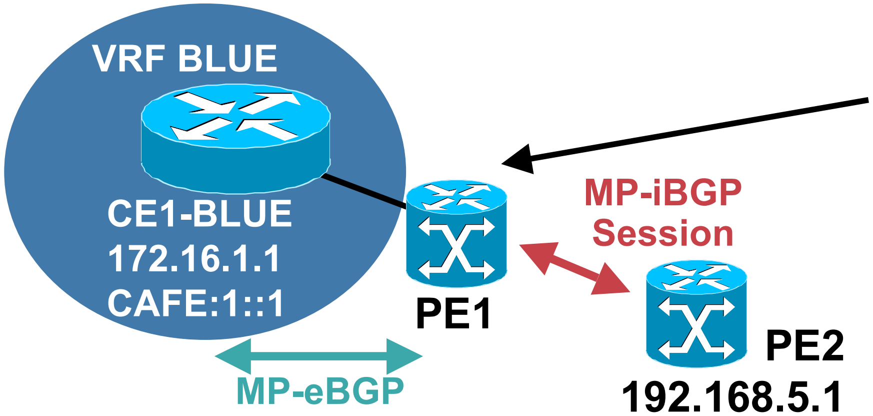


```
router bgp 100
 bgp log-neighbor-changes
 neighbor 192.168.5.1 remote-as 100
 neighbor 192.168.5.1 update-source
Loopback0
 !
 address-family ipv4
 neighbor 192.168.5.1 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family vpnv4
 neighbor 192.168.5.1 activate
 neighbor 192.168.5.1 send-community
extended
 exit-address-family
```

```
 address-family vpnv6
 neighbor 192.168.5.1 activate
 neighbor 192.168.5.1 send-community
extended
 exit-address-family
!
 address-family ipv4 vrf BLUE
 redistribute connected
 neighbor 172.16.1.1 remote-as 500
 neighbor 172.16.1.1 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv6 vrf BLUE
 neighbor 2001:DB8:CAFE:1::1 remote-as
500
 neighbor 2001:DB8:CAFE:1::1 activate
 redistribute connected
 no synchronization
 exit-address-family
```

6VPE Configuration Example
P Connections


```
mpls ldp router-id Loopback0
!
interface Loopback0
 ip address 192.168.3.1 255.255.255.255
!
interface Ethernet0/0
 description to PE1
 ip address 192.168.1.2 255.255.255.252
 mpls ip
!
interface Ethernet1/0
 description to P2
 ip address 192.168.1.5 255.255.255.252
 mpls ip
!
router ospf 1
 log-adjacency-changes
 redistribute connected subnets
 passive-interface Loopback0
 network 192.168.1.0 0.0.0.255 area 0
```

```
mpls ldp router-id Loopback0
!
interface Loopback0
 ip address 192.168.4.1 255.255.255.255
!
interface Ethernet0/0
 description to P1
 ip address 192.168.1.6 255.255.255.252
 mpls ip
!
interface Ethernet1/0
 description to PE2
 ip address 192.168.1.9 255.255.255.252
 mpls ip
!
router ospf 1
 log-adjacency-changes
 redistribute connected subnets
 passive-interface Loopback0
 network 192.168.1.0 0.0.0.255 area 0
```

IPv6 Routing Tables
CE1-CE2

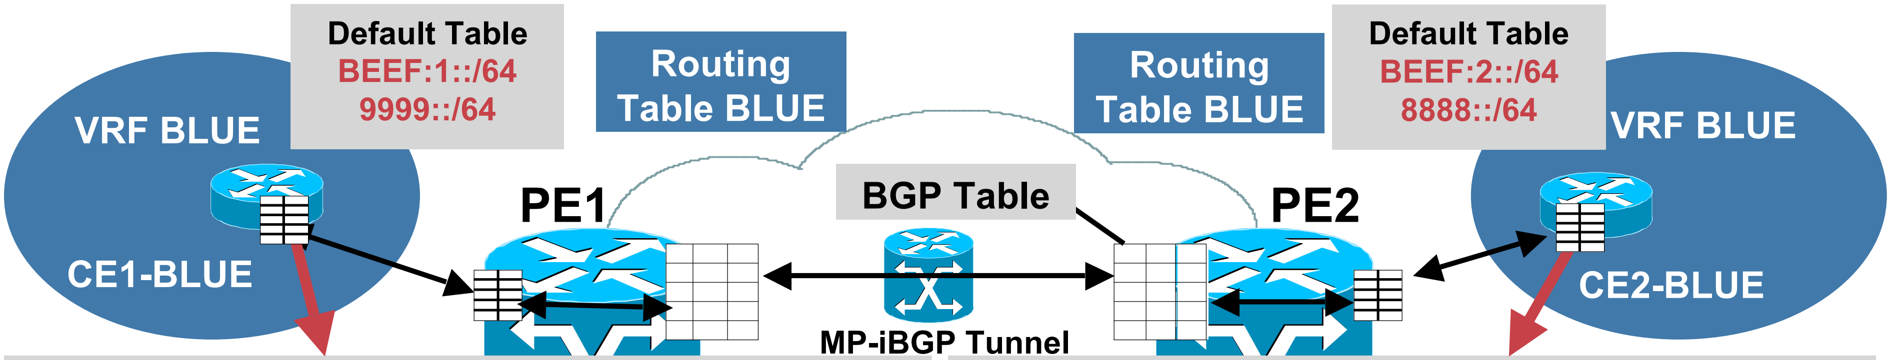


```
ce1-blue#show ipv6 route
C   2001:DB8:BEEF:1::/64 [0/0]
    via Ethernet1/0, directly connected
L   2001:DB8:BEEF:1::1/128 [0/0]
    via Ethernet1/0, receive
B   2001:DB8:BEEF:2::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F600, Ethernet0/0
C   2001:DB8:CAFE:1::/64 [0/0]
    via Ethernet0/0, directly connected
L   2001:DB8:CAFE:1::1/128 [0/0]
    via Ethernet0/0, receive
B   2001:DB8:CAFE:3::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F600, Ethernet0/0
B   8888::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F600, Ethernet0/0
R   9999::/64 [120/2]
    via FE80::A8BB:CCFF:FE01:9000, Ethernet1/0
L   FF00::/8 [0/0]
    via Null0, receive
```

```
ce2-blue#show ipv6 route
B   2001:DB8:BEEF:1::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F901, Ethernet0/0
C   2001:DB8:BEEF:2::/64 [0/0]
    via Ethernet1/0, directly connected
L   2001:DB8:BEEF:2::1/128 [0/0]
    via Ethernet1/0, receive
B   2001:DB8:CAFE:1::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F901, Ethernet0/0
C   2001:DB8:CAFE:3::/64 [0/0]
    via Ethernet0/0, directly connected
L   2001:DB8:CAFE:3::1/128 [0/0]
    via Ethernet0/0, receive
R   8888::/64 [120/2]
    via FE80::A8BB:CCFF:FE02:5800, Ethernet1/0
B   9999::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F901, Ethernet0/0
L   FF00::/8 [0/0]
    via Null0, receive
```

IPv6 Routing Tables

PE1-PE2

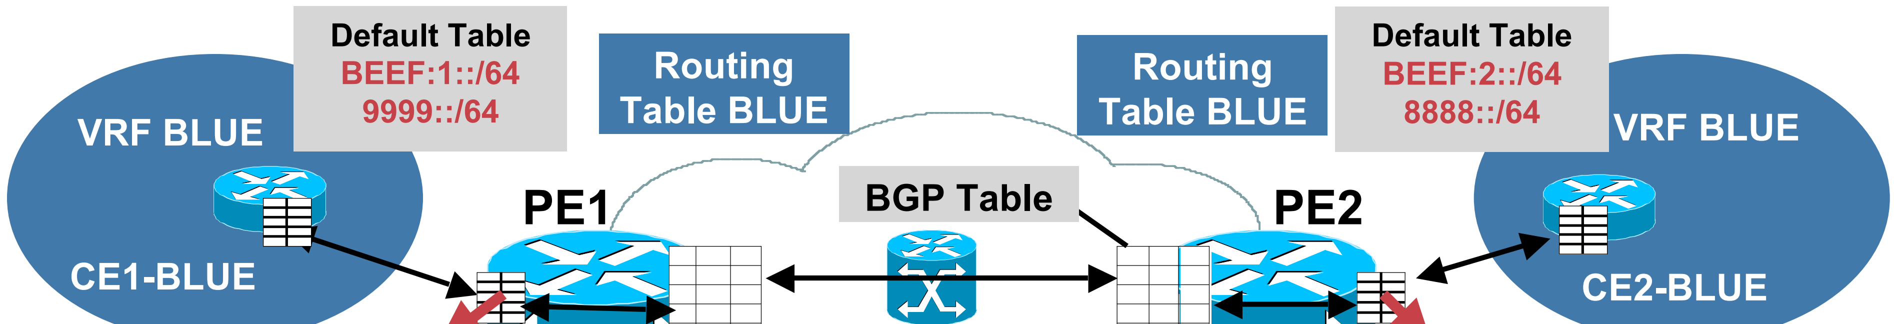


```
pe1#show ipv6 route vrf BLUE
B   2001:DB8:BEEF:1::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F400, Ethernet0/0
B   2001:DB8:BEEF:2::/64 [200/0]
    via 192.168.5.1%Default-IP-Routing-Table,
indirectly connected
C   2001:DB8:CAFE:1::/64 [0/0]
    via Ethernet0/0, directly connected
L   2001:DB8:CAFE:1::2/128 [0/0]
    via Ethernet0/0, receive
B   2001:DB8:CAFE:3::/64 [200/0]
    via 192.168.5.1%Default-IP-Routing-Table,
indirectly connected
B   8888::/64 [200/2]
    via 192.168.5.1%Default-IP-Routing-Table,
indirectly connected
B   9999::/64 [20/2]
    via FE80::A8BB:CCFF:FE01:F400, Ethernet0/0
L   FF00::/8 [0/0]
    via Null0, receive
```

```
pe2#show ipv6 route vrf BLUE
B   2001:DB8:BEEF:1::/64 [200/0]
    via 192.168.2.1%Default-IP-Routing-Table,
indirectly connected
B   2001:DB8:BEEF:2::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:FA00, Ethernet1/0
B   2001:DB8:CAFE:1::/64 [200/0]
    via 192.168.2.1%Default-IP-Routing-Table,
indirectly connected
C   2001:DB8:CAFE:3::/64 [0/0]
    via Ethernet1/0, directly connected
L   2001:DB8:CAFE:3::2/128 [0/0]
    via Ethernet1/0, receive
B   8888::/64 [20/2]
    via FE80::A8BB:CCFF:FE01:FA00, Ethernet1/0
B   9999::/64 [200/2]
    via 192.168.2.1%Default-IP-Routing-Table,
indirectly connected
L   FF00::/8 [0/0]
    via Null0, receive
```

IPv6 Routing Tables
PE1 BGP Next-Hop

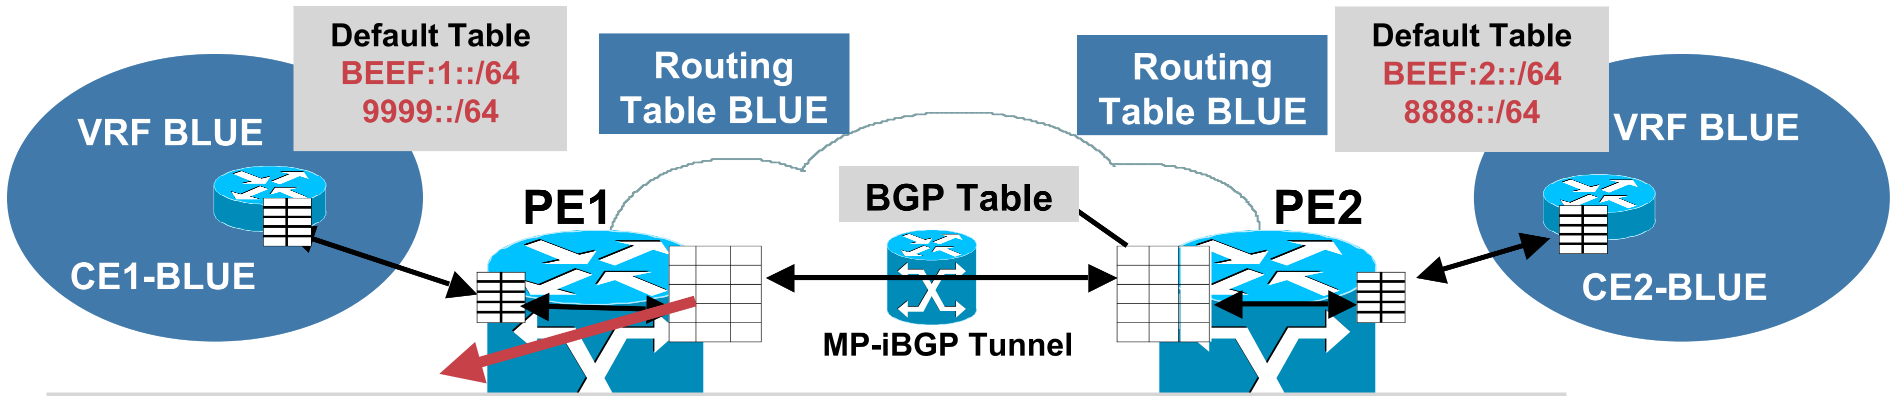


```
pe1#show bgp vpnv6 unicast all  #OUTPUT SHORTENED FOR CLARITY
Network                    Next Hop               Metric LocPrf Weight Path
Route Distinguisher: 200:1 (default for vrf BLUE)
*> 2001:DB8:BEEF:1::/64
                           2001:DB8:CAFE:1::1
                                                       0              0 500 ?
*>i2001:DB8:BEEF:2::/64
                           ::FFFF:192.168.5.1
                                                       0     100      0 506 ?
*>i2001:DB8:CAFE:3::/64
                           ::FFFF:192.168.5.1
                                                       0     100      0 ?
*>i8888::/64               ::FFFF:192.168.5.1
                                                       2     100      0 506 ?
*> 9999::/64               2001:DB8:CAFE:1::1
                                                       2              0 500 ?
```

IPv4-Mapped IPv6 Address (IPv4-Based LSP Setup)

MPLS Forwarding
PE1

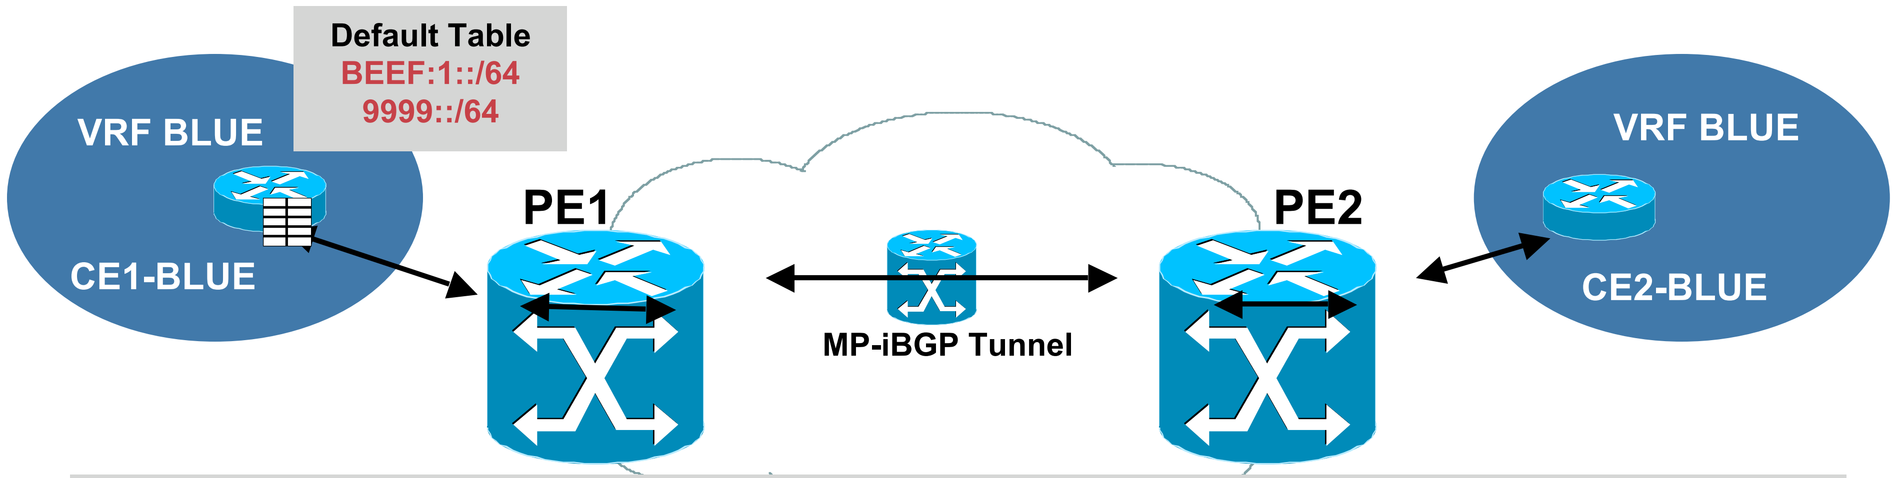


```
pe1#show mpls forwarding
Local   Outgoing      Prefix                 Bytes Label   Outgoing    Next Hop
Label   Label or VC   or Tunnel Id           Switched      interface
16      Pop Label     192.168.1.4/30         0             Et2/0       192.168.1.2
17      16            192.168.1.8/30         0             Et2/0       192.168.1.2
18      Pop Label     192.168.3.1/32         0             Et2/0       192.168.1.2
19      18            192.168.4.1/32         0             Et2/0       192.168.1.2
20      19            192.168.5.1/32         0             Et2/0       192.168.1.2
21      No Label      10.1.1.0/24[V]         0             Et0/0       172.16.1.1
22      Aggregate     172.16.1.0/24[V]       570           BLUE
25      No Label      2001:DB8:BEEF:1::/64[V]    \
                                             570           Et0/0       FE80::A8BB:CCFF:FE01:F400
26      Aggregate     2001:DB8:CAFE:1::/64[V]    \
                                             35456         BLUE
27      No Label      9999::/64[V]           570           Et0/0       FE80::A8BB:CCFF:FE01:F400
```

A Look at Forwarding

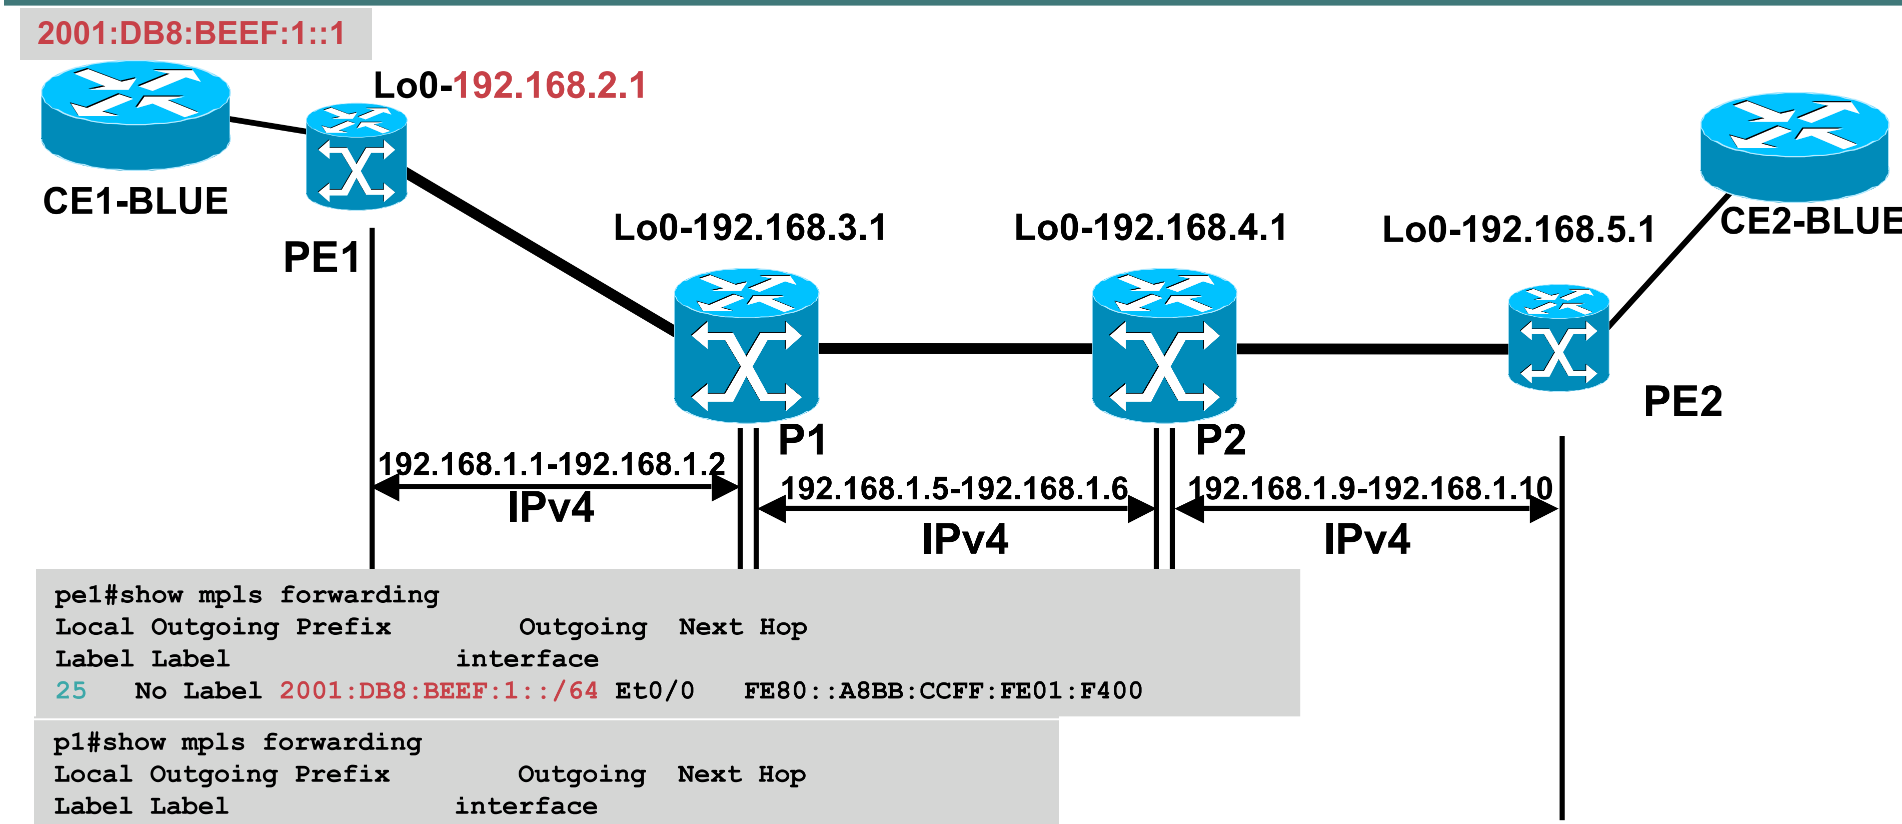


```
pe1#show mpls forwarding
Local Outgoing Prefix          Outgoing  Next Hop
Label Label                   interface
25    No Label 2001:DB8:BEEF:1::/64 Et0/0    FE80::A8BB:CCFF:FE01:F400
```

```
p1#show mpls forwarding
Local Outgoing Prefix          Outgoing  Next Hop
Label Label                   interface
17    Pop Label    192.168.2.1/32  Et0/0    192.168.1.1
```

```
p2#show mpls forwarding
Local Outgoing Prefix          Outgoing  Next Hop
Label Label                   interface
18    17        192.168.2.1/32  Et0/0    192.168.1.5
```

```
pe2#sh ipv cef vrf BLUE
2001:DB8:BEEF:1::/64
 nexthop 192.168.1.9 Ethernet0/0 label 18 25
```

6VPE Summary

- 6VPE simply adds IPv6 support to current IPv4 MPLS VPN offering
- For end-users: v6-VPN is same as v4-VPN services (QoS, hub and spoke, internet access, etc.)
- For operators:
  - Same configuration operation for v4 and v6 VPN
  - No upgrade of IPv4/MPLS core (IPv6 unaware)
- Cisco 6VPE is an implementation of <draft-ietf-bgp-ipv6-vpn> over MPLS/IPv4
- <draft-ietf-l3vpn-bgp-ipv6-xx>
  - BGP-MPLS VPN extension for IPv6 VPN
  - Generic for operations over any tunneling technique (MPLS, IPsec, L2TPv3, GRE)

SERVICE PROVIDER—ACCESS

Drivers for IPv6 in Broadband

- **Network Management:** The most striking aspect of Broadband Access Services is the large number of users that imply a larger number of devices to be managed by providers. Even the private IPv4 address space will be unable to withstand the expected needs. IPv6 is seen as the answer to this problem.
- **New Services:** The current business models for Network Access Provider (wholesale model) avoid handling users at Layer 3 at the access layer. These models do not scale for services such as Multicast. IPv6 offers the address resources needed to deploy such services optimally.
- **Prepare for the Future:** Build an infrastructure that would be ready for the new services and IP enabled appliances.

Broadband Home and IPv6 – a Must!

IPv6 Multicast-Based Multimedia Services (NTT-East Example)

- **NTT-East rolled out native IPv6 multicast services instead of IPv4 offering IPTV, music and games:**

  http://www.ipv6style.jp/en/action/20040902/index.shtml

- **The IPv6 solution is scaleable since it allows for the replication to be performed at the access layer**

Cisco IOS IPv6 Broadband Access Solutions

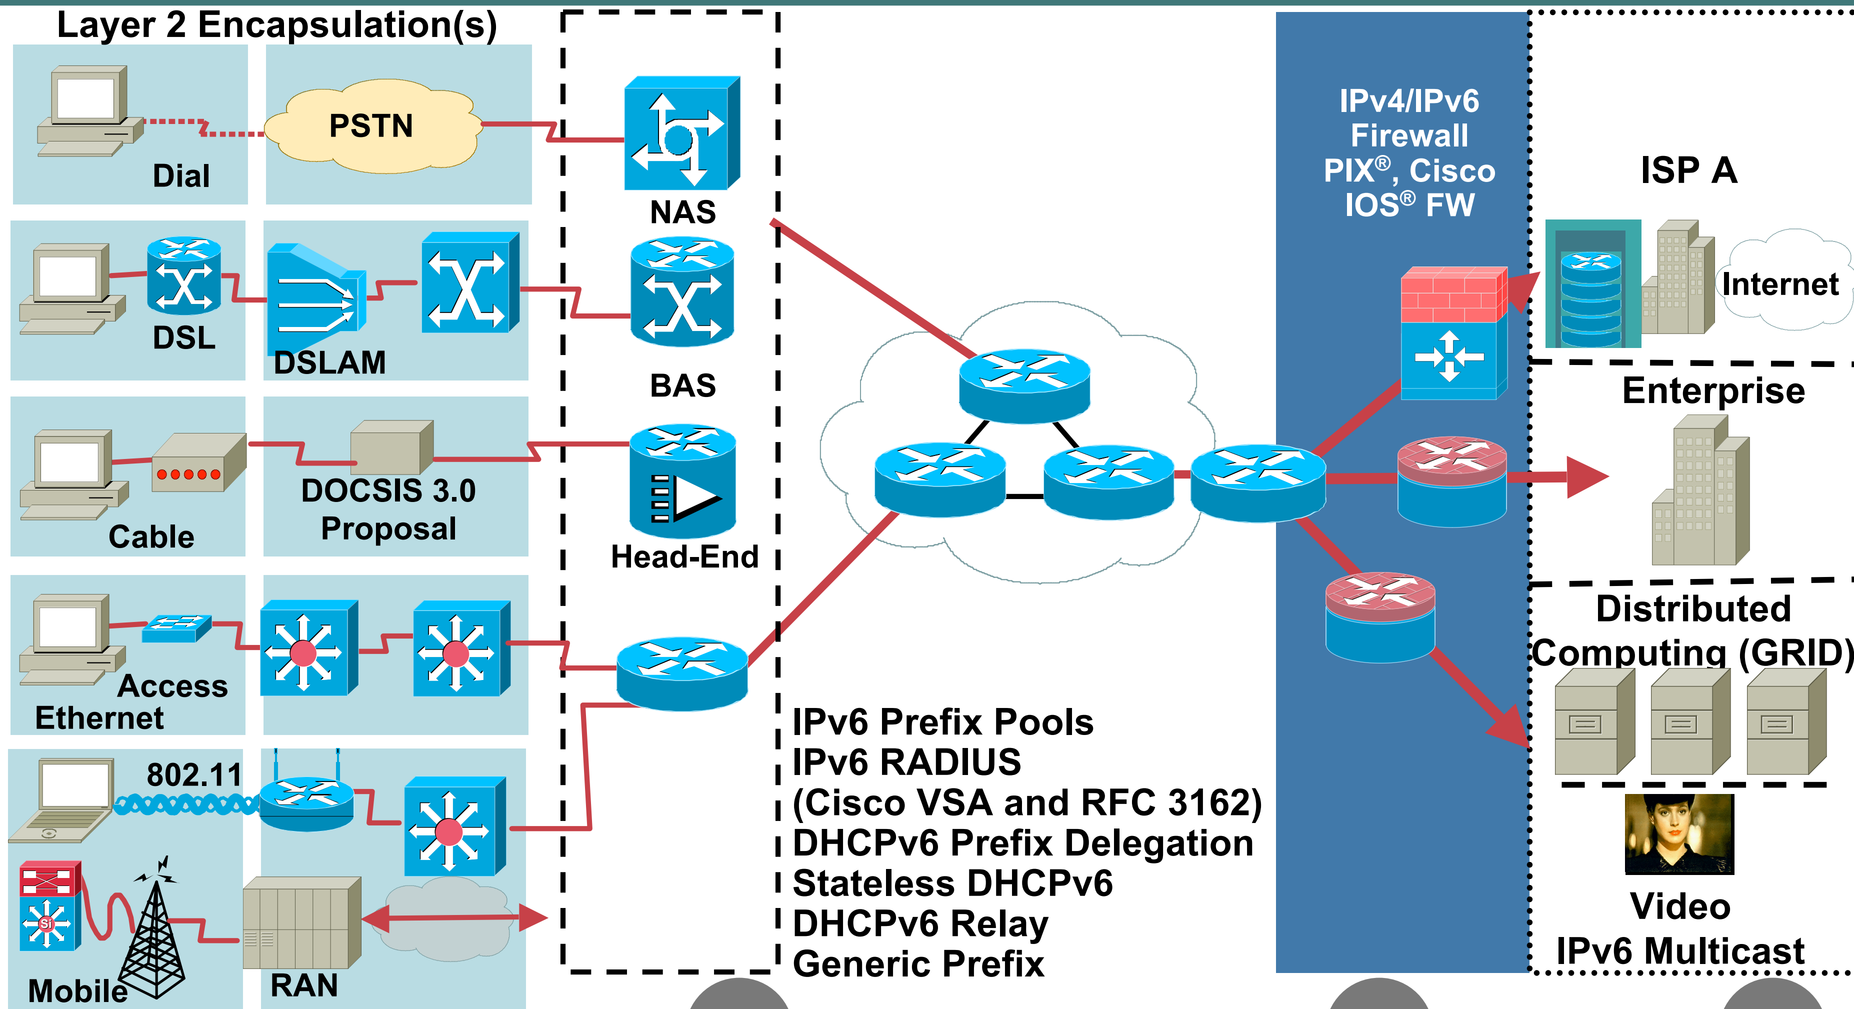

IPv6 Prefix Pools
IPv6 RADIUS
(Cisco VSA and RFC 3162)
DHCPv6 Prefix Delegation
Stateless DHCPv6
DHCPv6 Relay
Generic Prefix

ATM RFC 1483 Routed or Bridged (RBE)
PPP, PPPoA, PPPoE, Tunnel (Cable)

Dual-Stack or MPLS (6PE) Core

IPv4/IPv6

Two Broadband Access Models Today

- **Network access provider = internet service provider**

- **Network access provider # internet service provider**

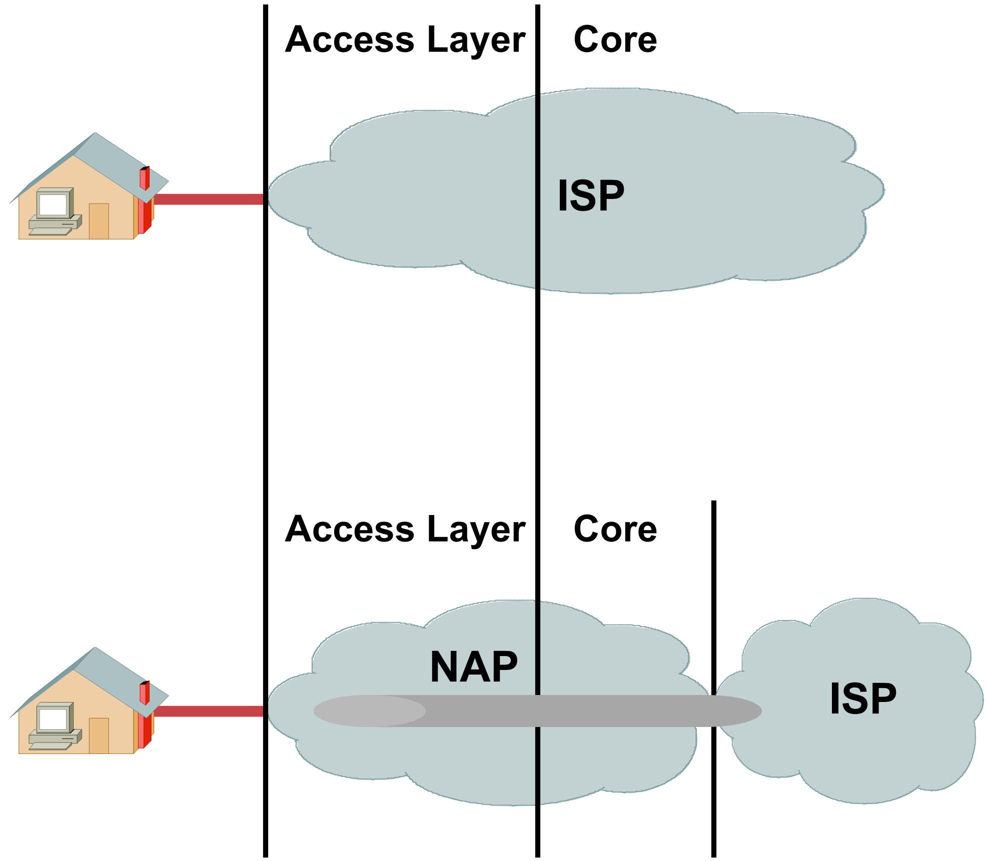

xDSL, ETTH and WLAN Networks

Point-to-Point Model

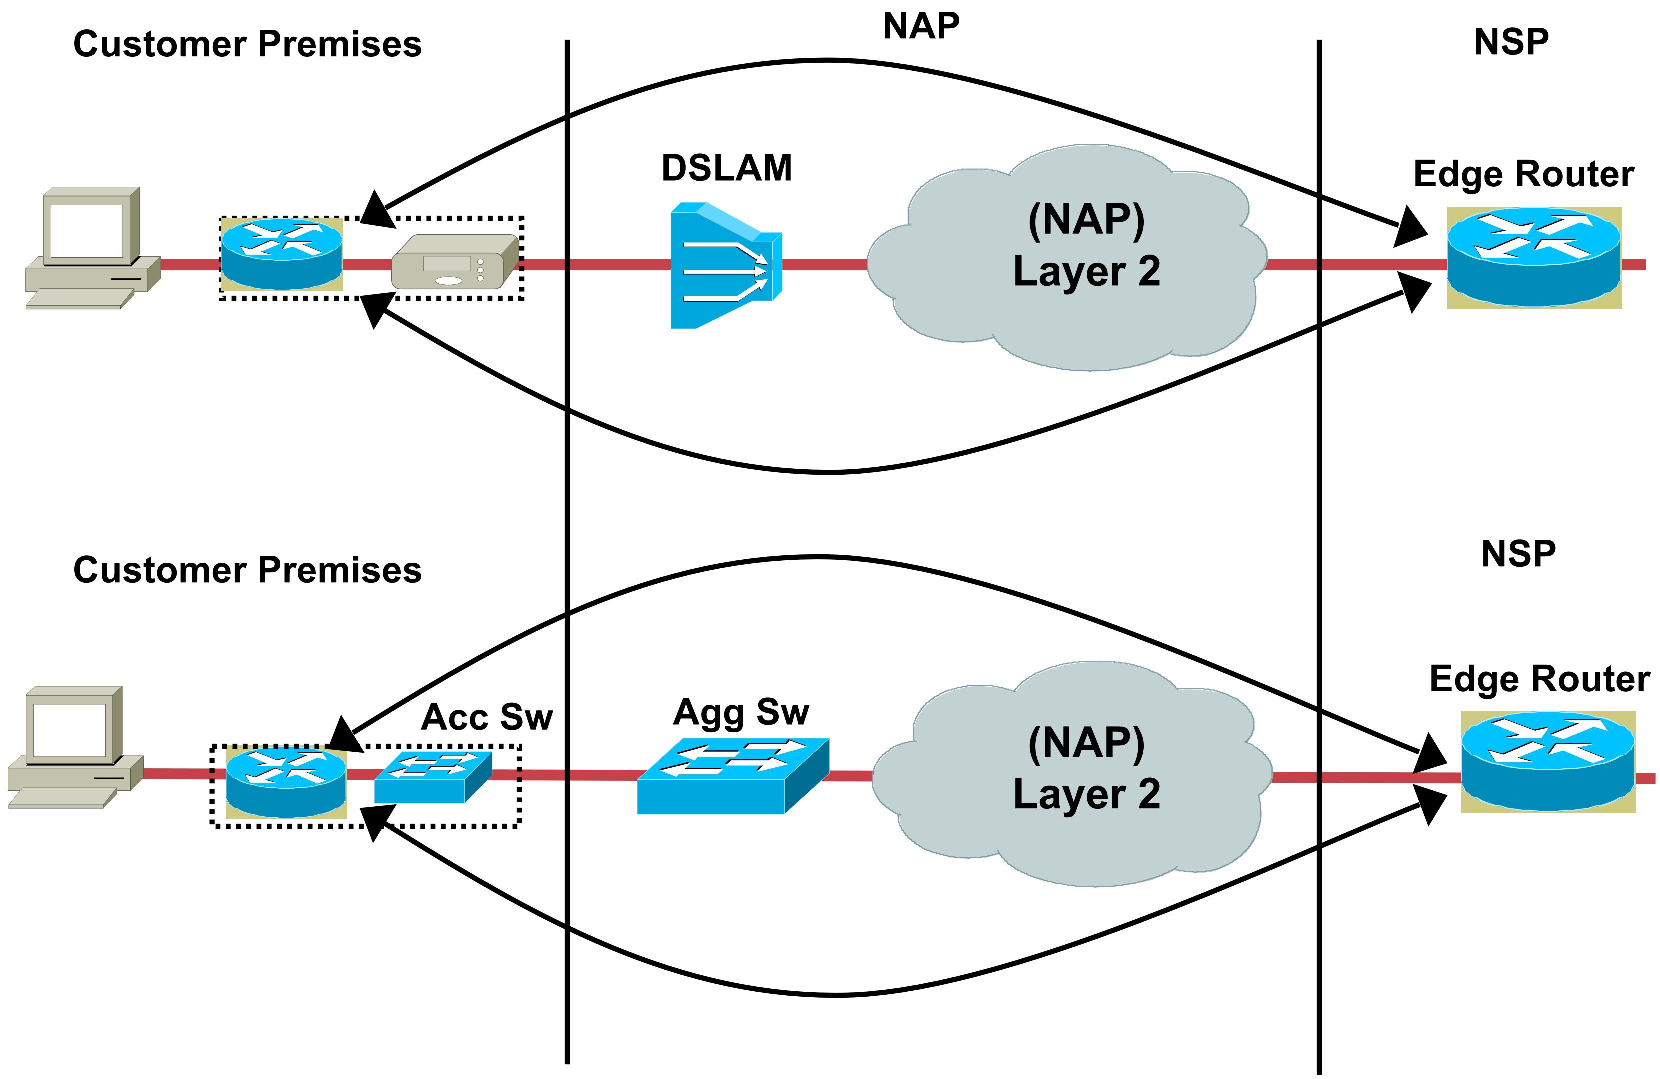

L2TPv2 Access Aggregation (LAA) Model

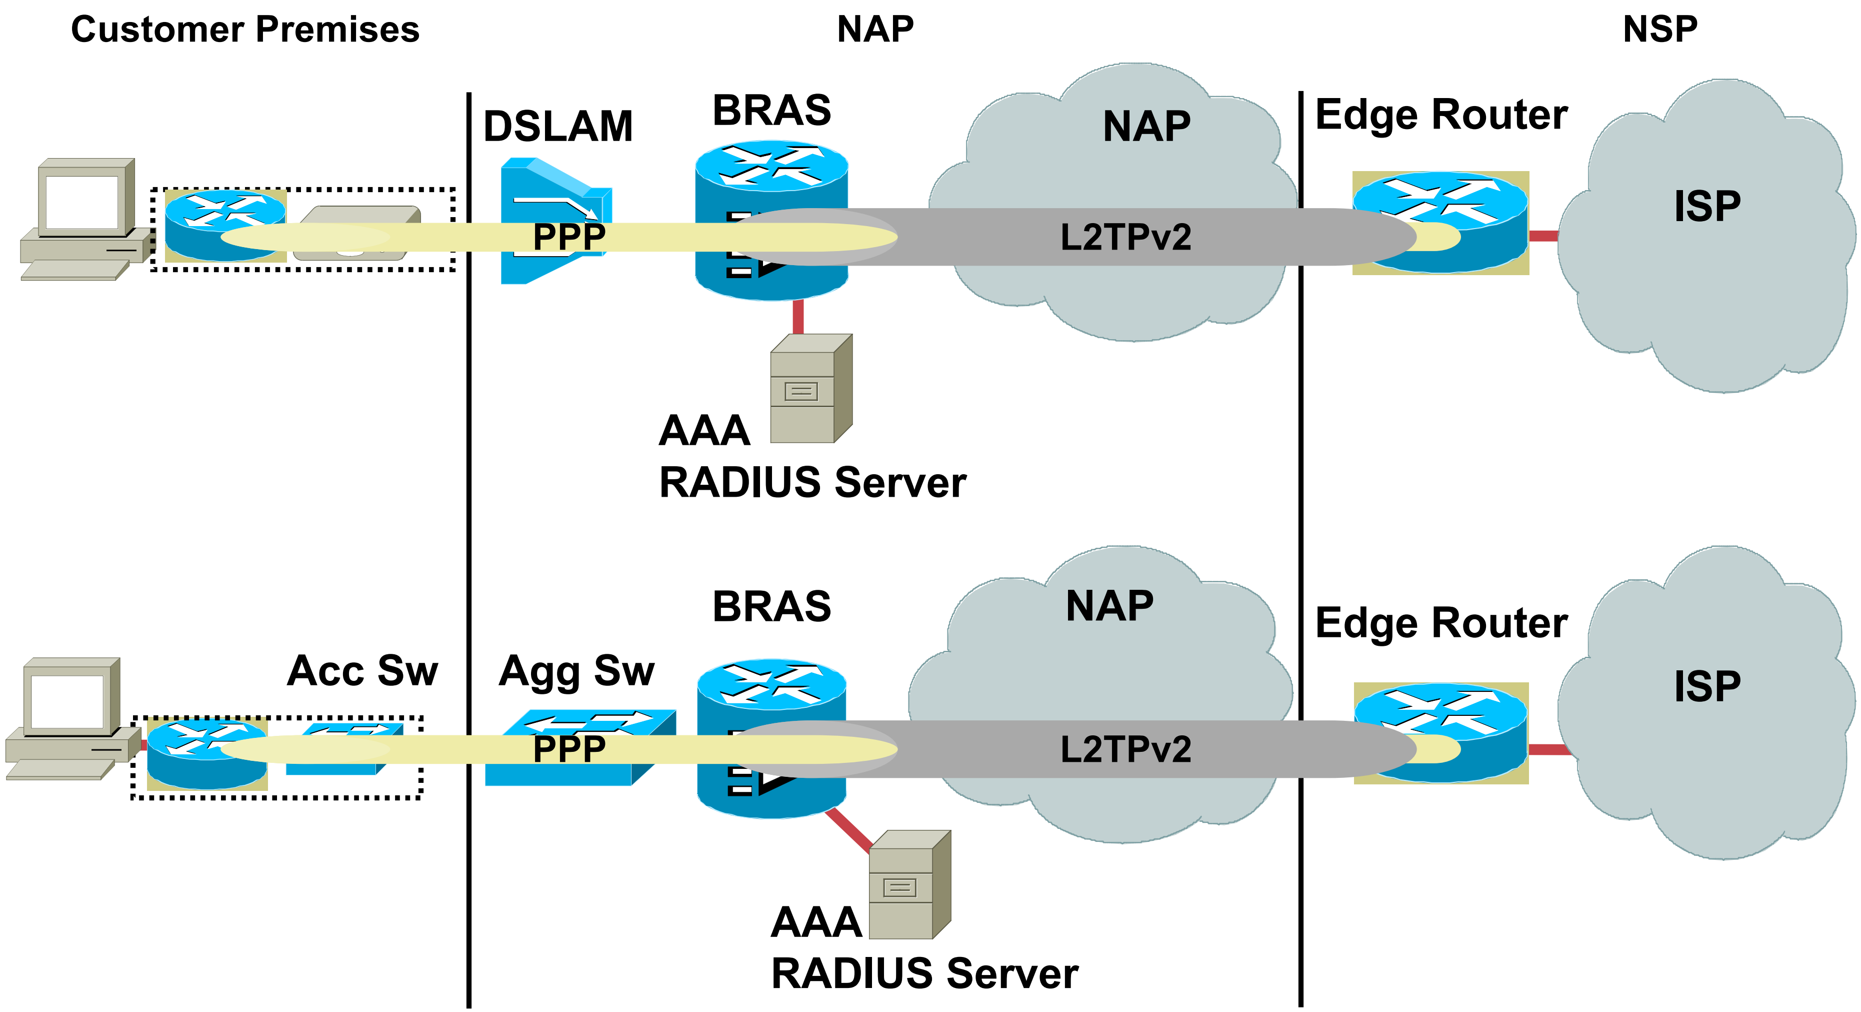

PPP Terminated Aggregation (PTA) Model

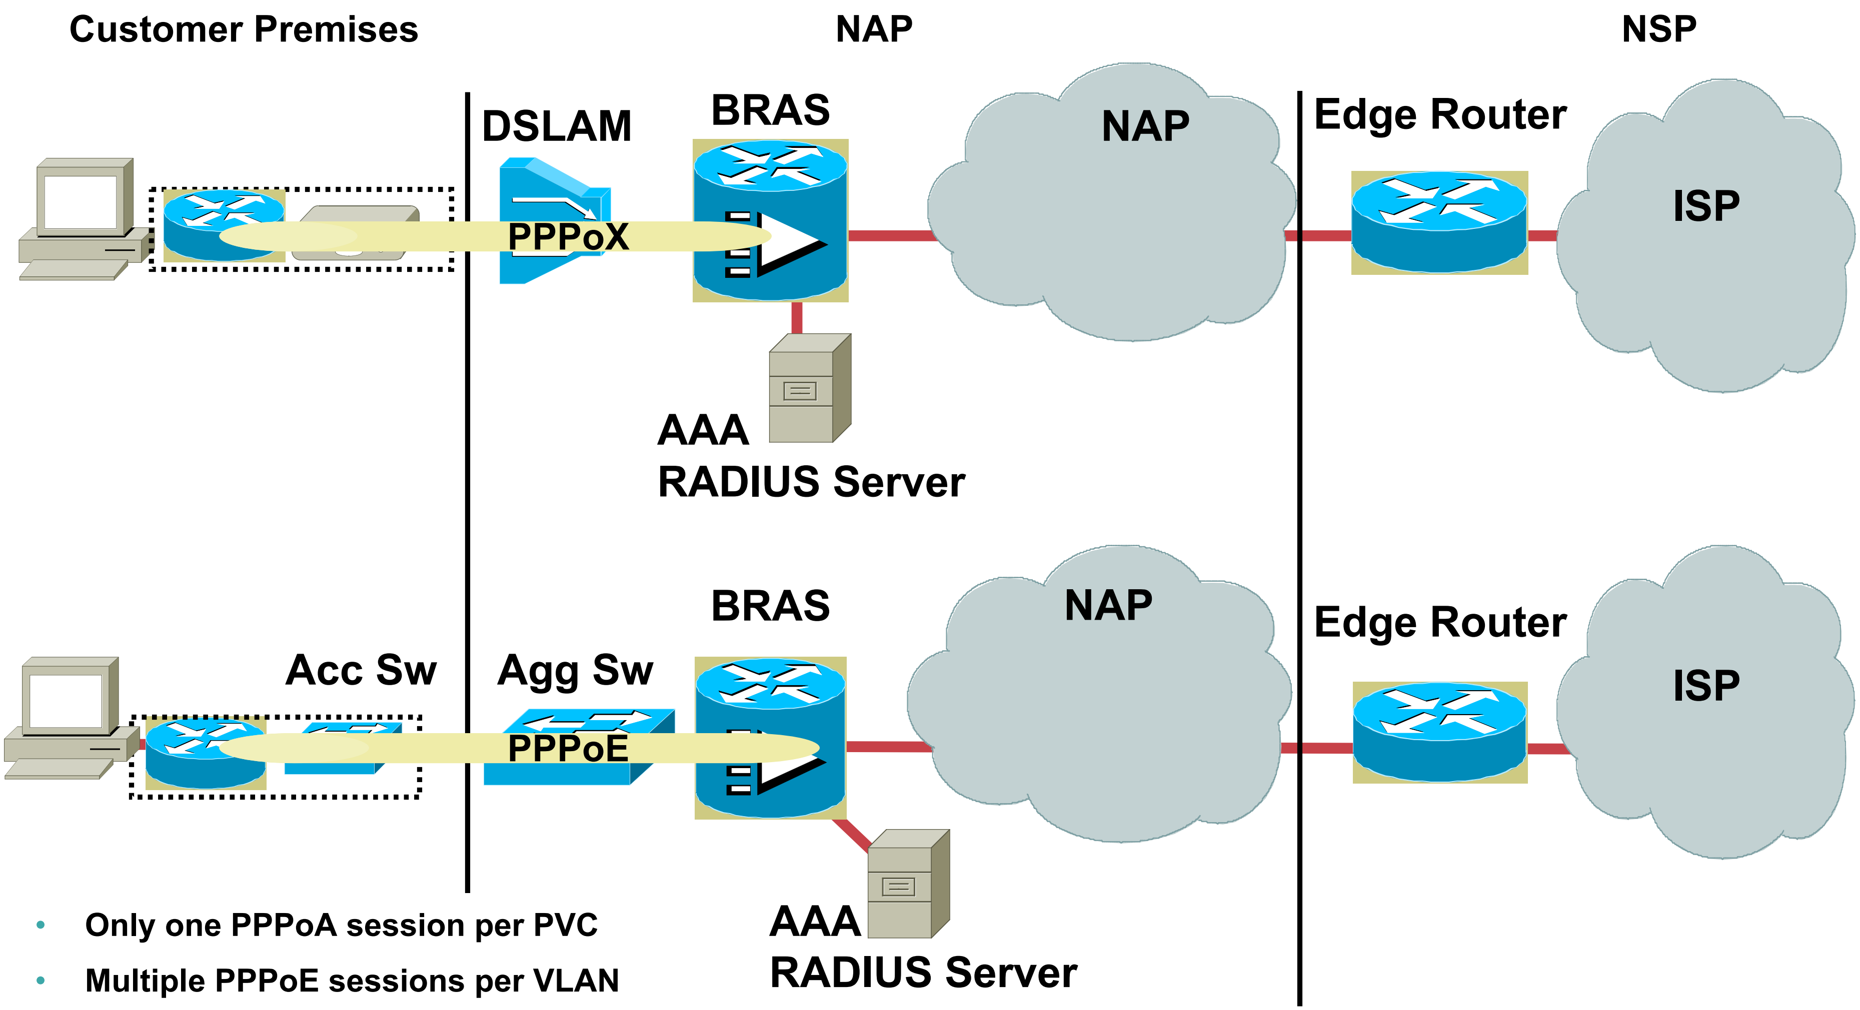

- Only one PPPoA session per PVC
- Multiple PPPoE sessions per VLAN
- The PPPoE sessions can be initiated by the hosts or the CPE

Hybrid: IPv4 LAA Model and IPv6 PTA Model

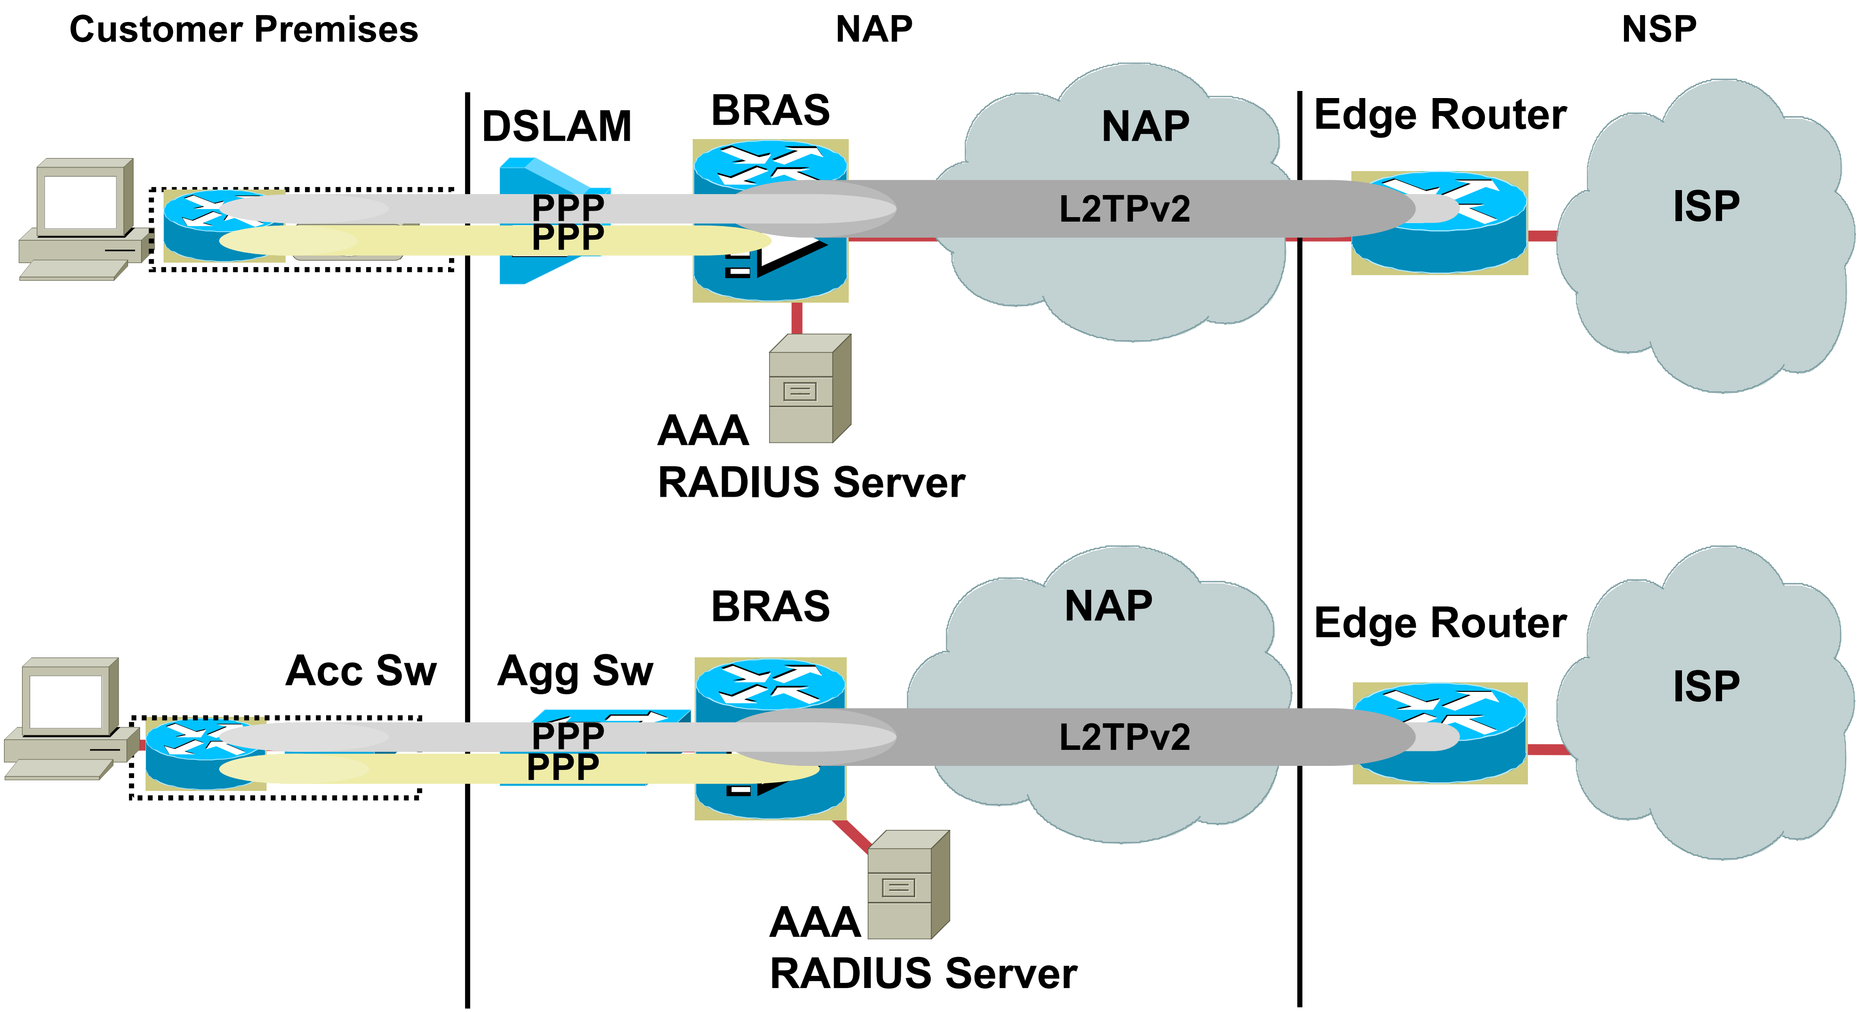

IPv6 RBE

Different Then IPv4 RBE:

- Pick out the 0x86DD type and route the traffic
- Enabled per PVC, IPv6 address is configured per PVC, each PVC supports a different subnet

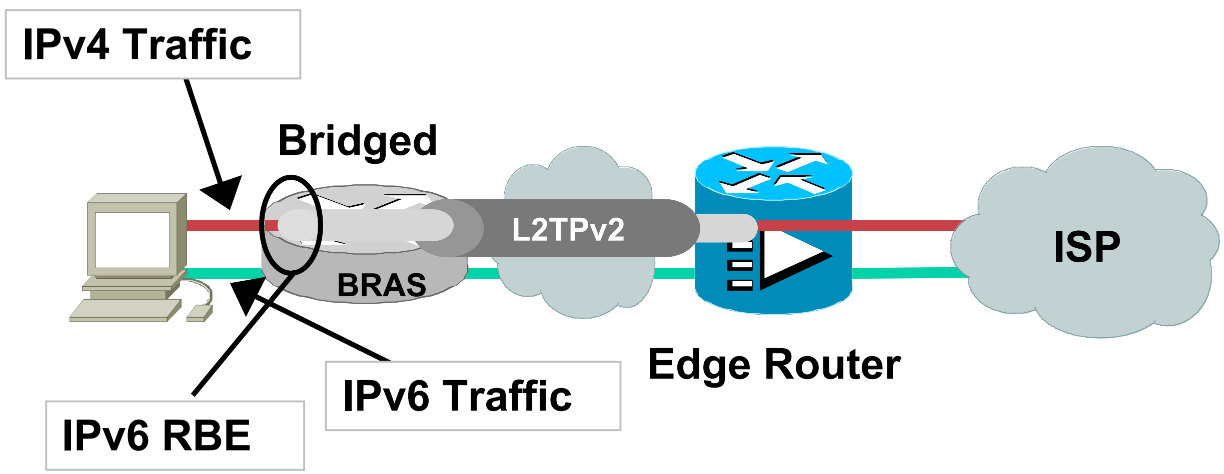

Cable Networks

Drivers for IPv6 in Cable

- **Use IPv6 for managing large number of devices on the network**
  - **Exponential growth in number of IP-enabled devices connected to CMTS**
  - **Cable MSOs in the US would like to use IPv6 to manage CM/MTA**
  - **Currently RFC1918 addresses assigned to CM for management**
- **RFC 1918 provides 16 million 10.net addresses, plus:**
  - **1M addresses under 172.16.0.0/12**
  - **65K addresses under 192.168/16**
- **Moreover, address utilization efficiency for large numbers decreases with topology hierarchies***
  - **6.5M addresses for 4M CMs**
  - **Only 61.5% efficient use**
  - **Density of only 9.8M CMs exhausts all 16M RFC1918 addresses**

***See HD Ratio, RFC1715 and RFC3194**

IPv6 Deployment Models for Cable

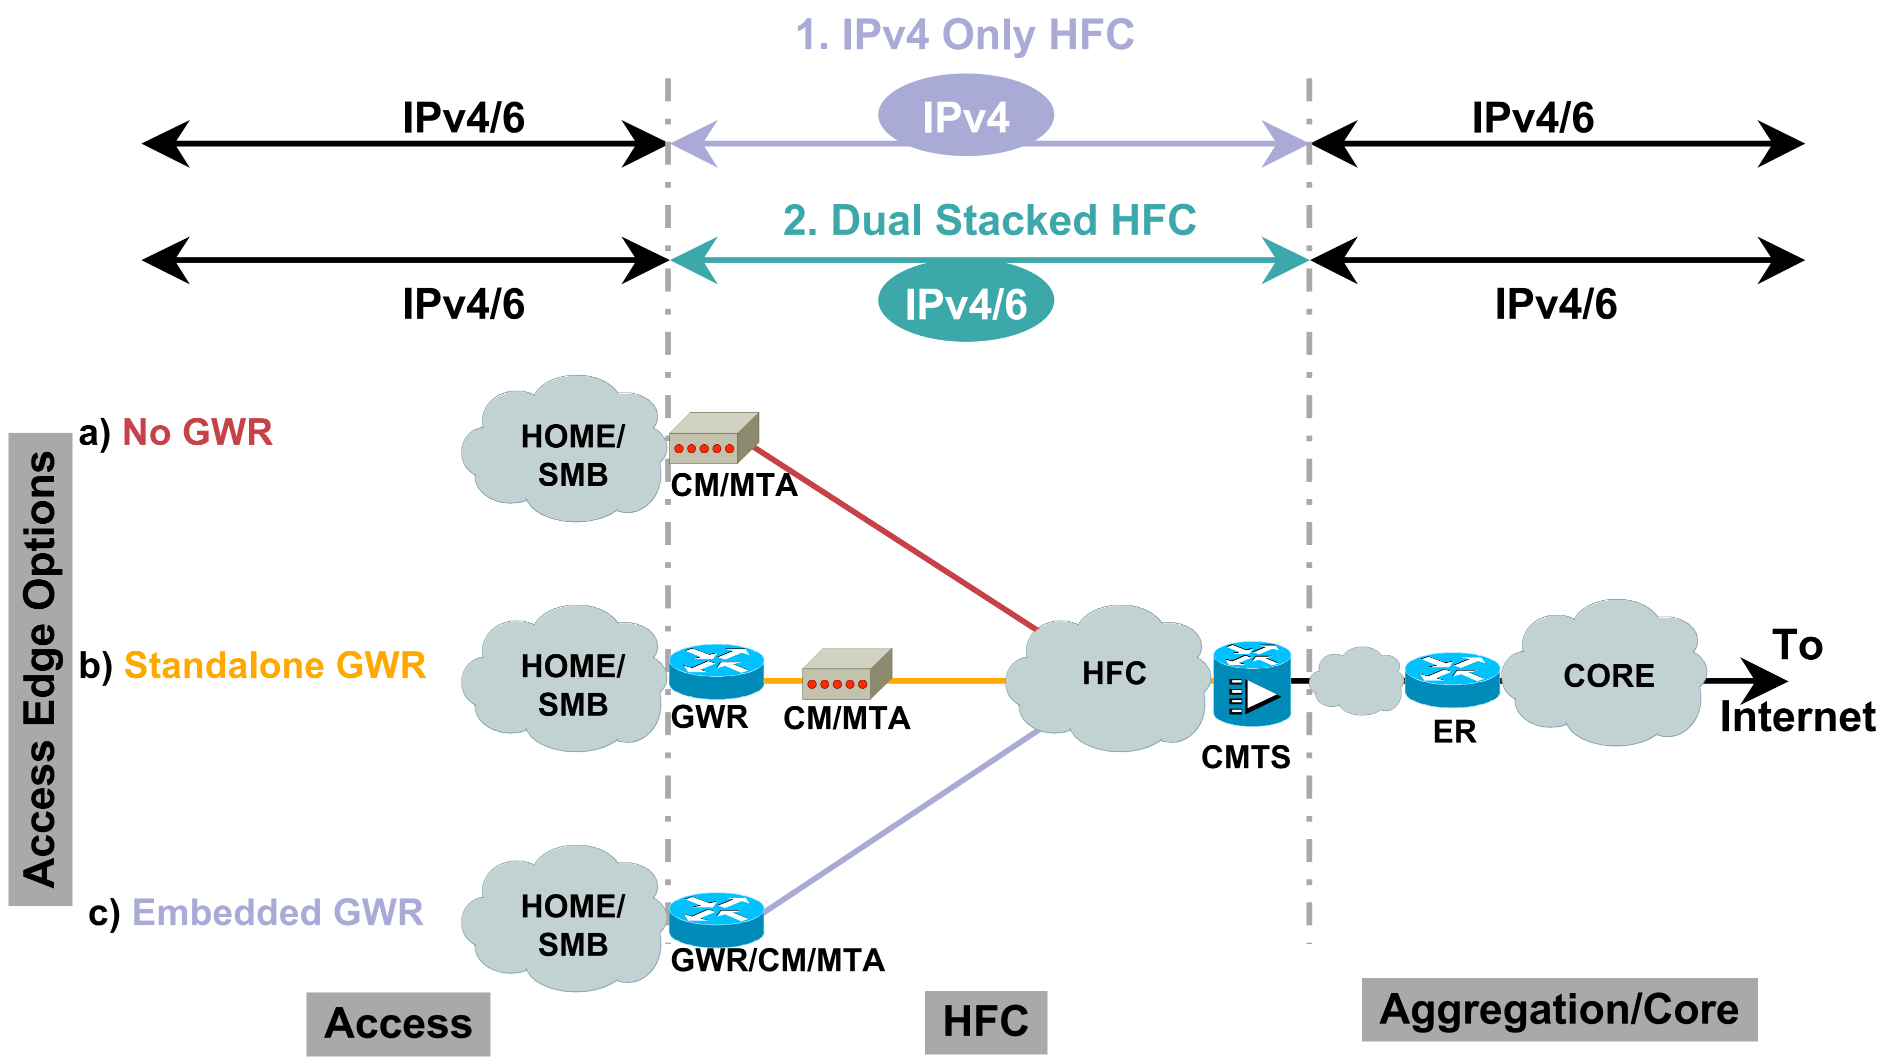

IPv6 Deployment Challenges in Cable

- Problems with Neighbor Discovery (ND) on CM and CMTS, due to lack of IGMPv3/MLDv2 or v1 snooping support
- No way to classify IPv6 traffic on the CM and CMTS. Cannot provide appropriate QoS to traffic, everything sent as Best Effort (BE).
- Changes needed in the DOCSIS RFI specification to support native IPv6 deployment over cable

Addressed in DOCSIS 3.0 Standardization

IPv6 Deployment Models for DOCSIS 3.0

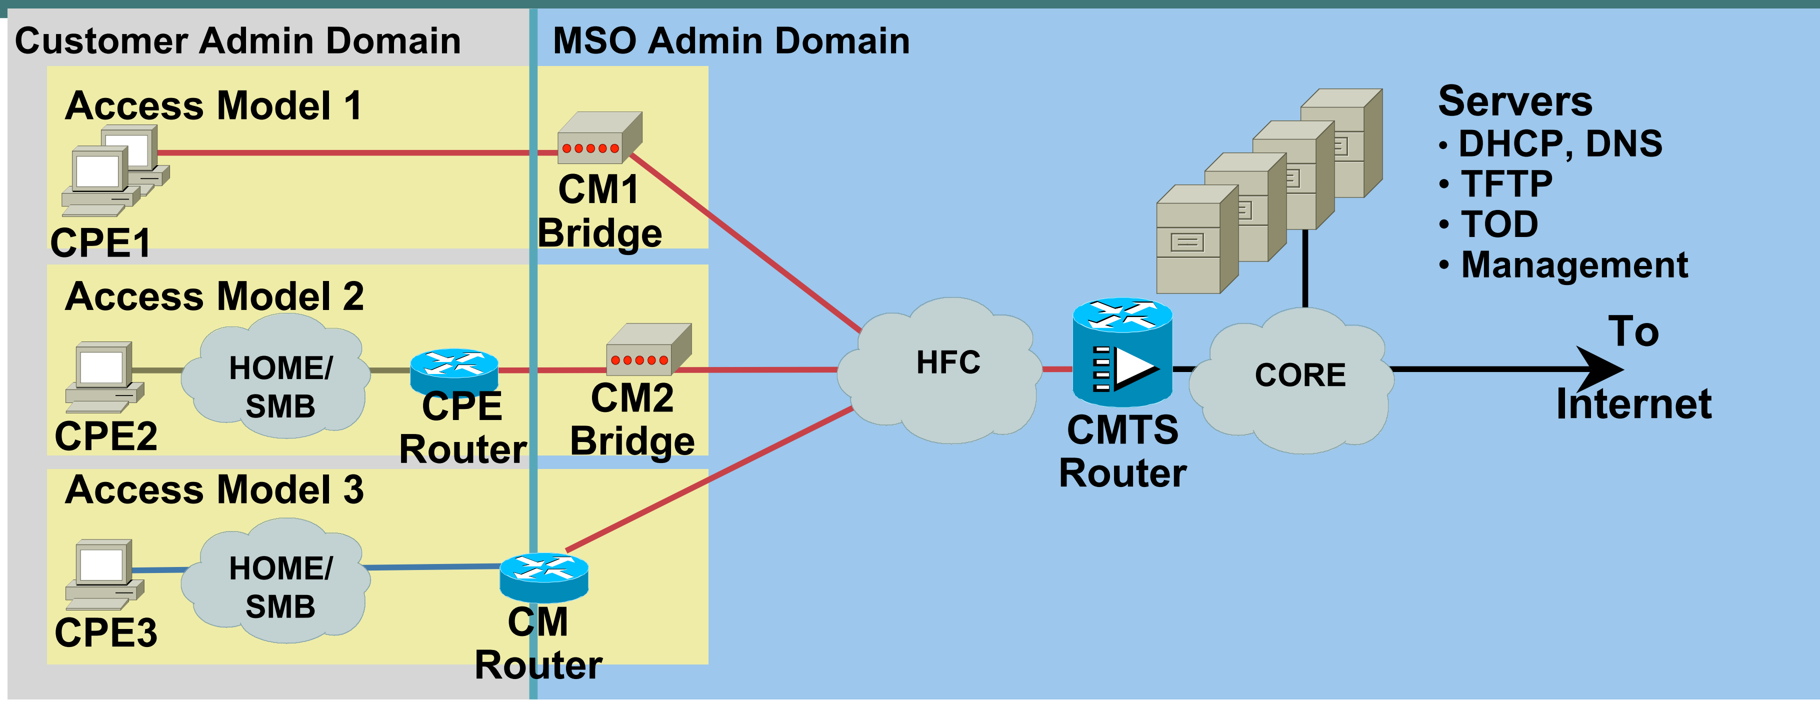

Management Prefix: 2001:DB8:FFFF:0::/64
Service Prefix: 2001:DB8:FFFE:0::/64
Customer 2 Prefix: 2001:DB8:2::/48
Customer 3 Prefix: 2001:DB8:3::/48

HFC Link; Assigned 2001:DB8:FFFF:0::/64 (Mgmt) and 2001:DB8:FFFE:0::/64 (Serv)
Customer 2 Premises Link; Assigned 2001:DB8:2:0::/64
Customer 3 Premises Link; Assigned 2001:DB8:3:0::/64

Routers Span Customer and MSO Administrative Domains

Provisioning in IPv6 Broadband Environments

IPv6 Address Allocation Guidelines

"…recommends the assignment of /48 in the general case, /64 when it is known that one and only one subnet is needed…"

RFC3177
IAB/IESG Recommendations on IPv6 Address Allocations to Sites

Note: /128 Assignment Can Be Used When It Is Absolutely Known That One and Only One Device Is Connecting

DHCPv6 Overview

- **Operational model based on DHCPv4**
- **Details are different**
  - Client uses link-local address for message exchanges
  - Server can assign multiple addresses per client through identity associations
  - Clients and servers identified by DUID
  - Address assignment
  - Prefix delegation
  - Message exchanges similar, but will require new protocol engine
  - Server-initiated configuration, authentication part of the base specification
  - Extensible option mechanism
  - Relay-agents
- **Allows both stateful and stateless configuration**
- **RFC 3315 (DHCPv6)**
  - Additional options:
    - DNS configuration—RFC 3646
    - Prefix delegation—RFC 3633
    - NTP servers
    - Stateless DHCP for IPv6—RFC 3736

DHCPv6 Operation

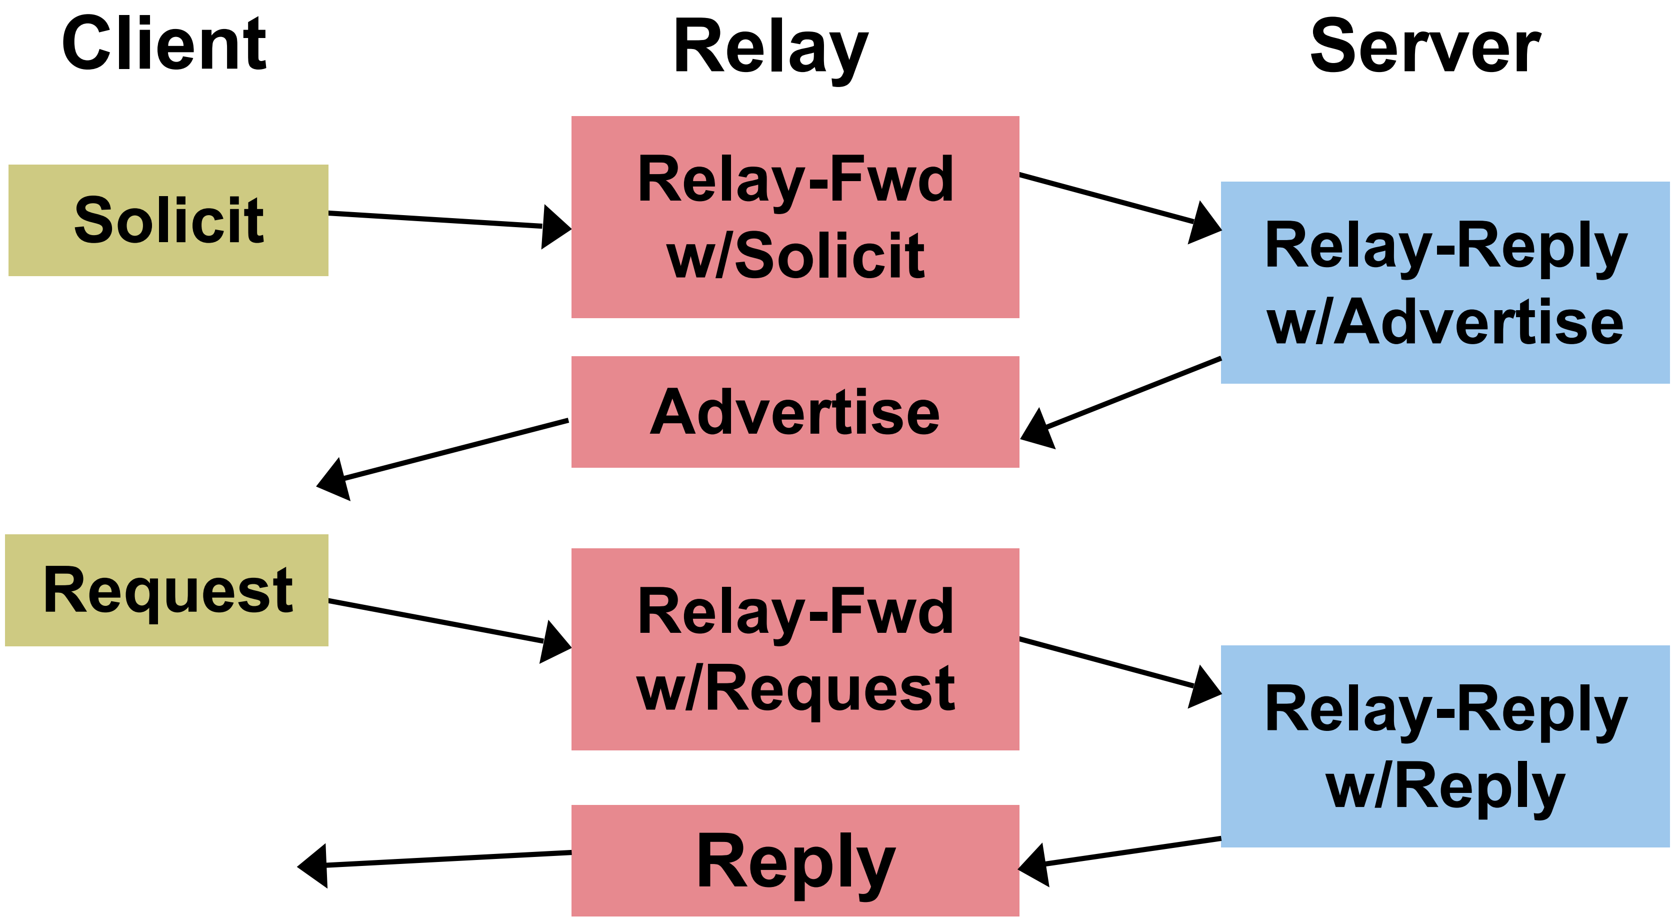

- All_DHCP_Relay_Agents_and_Servers (FF02::1:2)
- All_DHCP_Servers (FF05::1:3)
- DHCP Messages: Clients listen UDP port 546. Servers and relay agents listen on UDP port 547

DHCPv6 PD: RFC 3633

- **Media independence**
  - E.g., ADSL, FTTH
  - Only knows identity of requesting router
- **Leases for prefixes**
- **Flexible deployments**
  - Client/relay/server model
- **Requesting router includes request for prefixes in DHCP configuration request**
- **Delegating router assigns prefixes in response along with other DHCP configuration information**

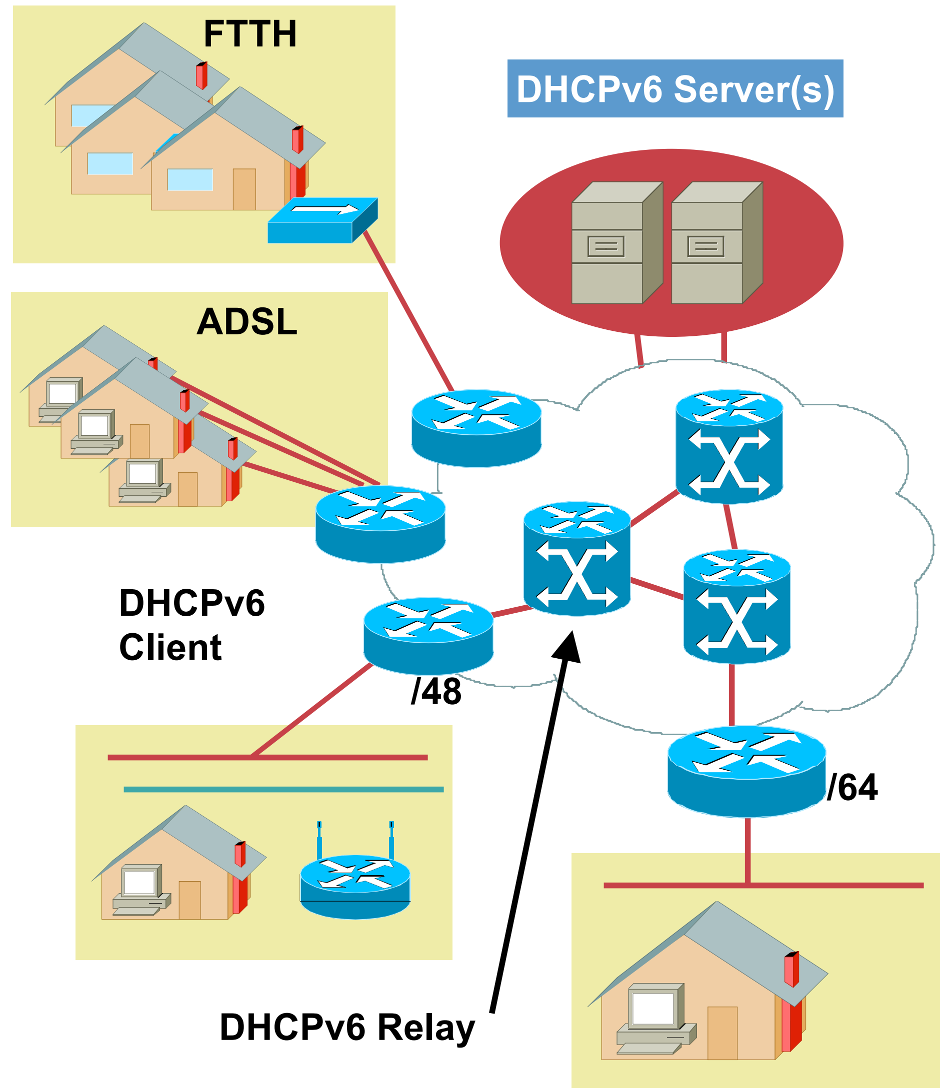

Router Advertisement

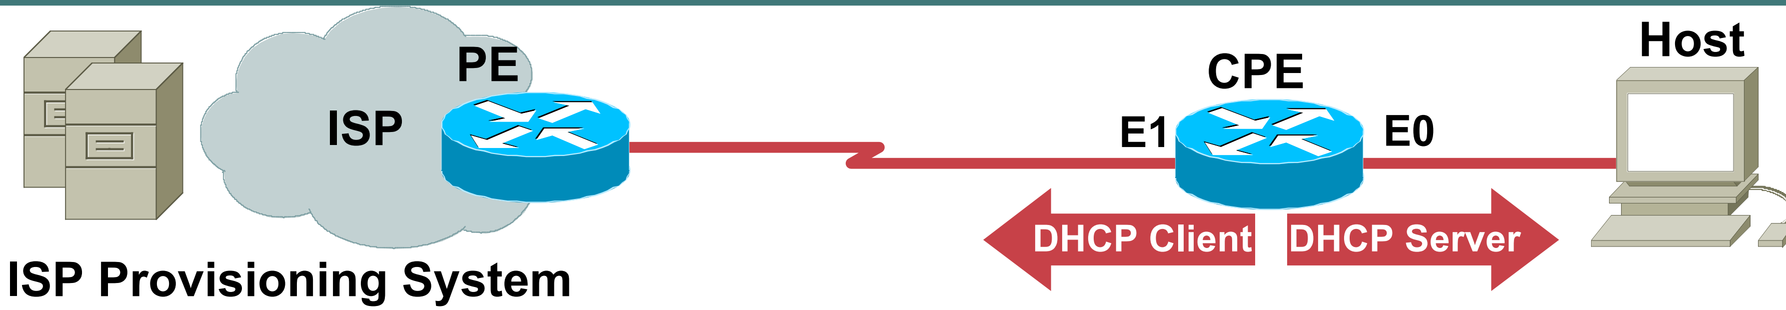

Source of RA	User of RA	A Bit		M/O Bits	
		A	Operation	M/O	Operation
PE	CPE E1	0	Don’t Do Stateless Address Assignment	11	Use Dhcpv6 for Address + Other Config. (I.E. Stateful Dhcpv6)
CPE Router	Host	1	Do Stateless Address Assignment	01	Use Dhcpv6 for Other Config. (I.E. Stateless Dhcpv6)

Stateless (RFC2462)
RS Are Sent by Booting Nodes to Request RAs for Configuring the Interfaces; Host Autonomously Configures Its Own Link-Local Address

Prefix/Options Assignment

1. CPE Sends DHCP Solicit with ORO = PD
2. PE Sends RADIUS Request for the User
3. RADIUS Responds with User's Prefix(es)
4. PE Sends DHCP REPLY with Prefix Delegation Options
5. CPE Configures Addresses from The Prefix on Its Downstream Interfaces, and Sends an RA. O-bit Is Set to On
6. Host Configures Addresses Based on the Prefixes Received In the RA. As the O-bit Is on, It Sends a DHCP Information-request Message, with an ORO = DNS
7. CPE Sends a DHCP REPLY Containing Request Options

PE/CE IPv6 Debugs

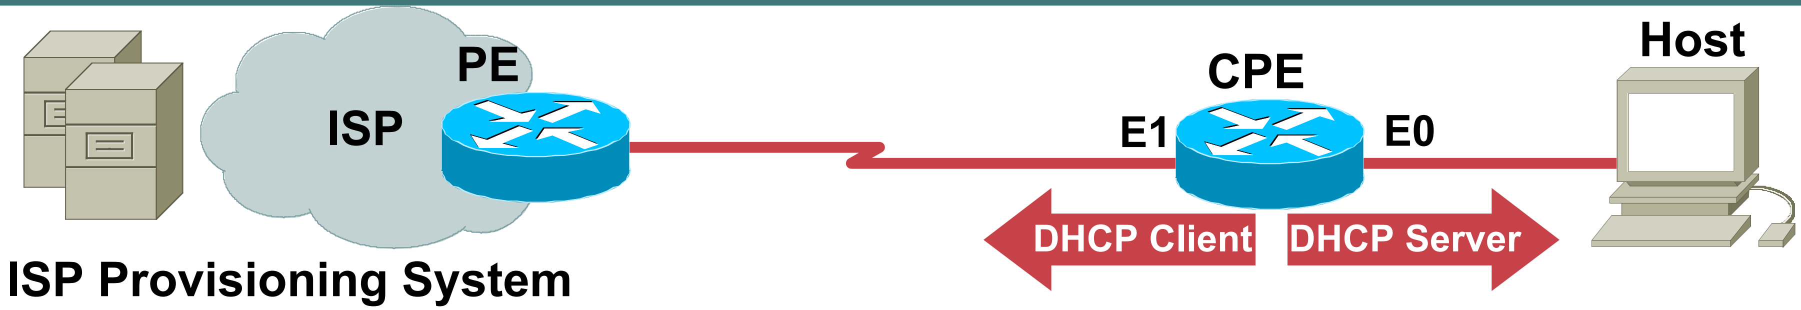

debug ipv6 nd
debug ipv6 dhcp detail
debug ipv6 dhcp relay

PE#show debug
Generic IPv6:
ICMP Neighbor Discovery events debugging is on
IPv6 DHCP debugging is on (detailed)
IPv6 DHCP relay debugging is on

PE Configuration


```
!
hostname PE_Router
!
interface GigabitEthernet3/1
ipv6 address 2001:420:3800:800:0:1:0:1/96
 ipv6 enable
 ipv6 nd ra-interval 5
 ipv6 nd prefix default no-advertise
 ipv6 nd managed-config-flag
 ipv6 nd other-config-flag
 ipv6 rip PE_Router enable
 ipv6 mld static-group FF0E:0:0:1::1000
 ipv6 dhcp relay destination 2001:420:8:1:5::2 GigabitEthernet0/1
!
interface GigabitEthernet0/1
ip address 10.89.240.235 255.255.255.248
 ip pim sparse-mode
 negotiation auto
 ipv6 address 2001:420:3800:800::12/124
 ipv6 enable
 ipv6 router isis
 ipv6 mld static-group FF0E:0:0:1::1000
 hold-queue 2048 in
!
```

PE Debugs: ND-SOLICIT

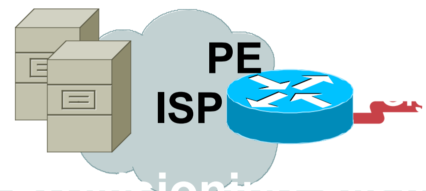

ISP Provisioning System

*Feb 15 21:35:16.946: ICMPv6-ND: Received NS for FE80::207:EFF:FE03:6E65 on GigE3/1 from ::
[DAD request from CPE for Link-local Address]
*Feb 15 21:35:17.650: ICMPv6-ND: Sending RA to FF02::1 on GigE3/1
*Feb 15 21:35:17.650: ICMPv6-ND: MTU = 1500
*Feb 15 21:35:17.934: ICMPv6-ND: Received NA for FE80::207:EFF:FE03:6E65 on GigE3/1 from FE80::207:EFF:FE03:6E65
[CPE assigns Link-local Address and sends NA]

*Feb 15 21:35:19.862: IPv6 DHCP: Received SOLICIT from FE80::207:EFF:FE03:6E65 on GigE3/1
*Feb 15 21:35:19.862: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:19.862: src FE80::207:EFF:FE03:6E65 (GigE3/1)
*Feb 15 21:35:19.862: dst FF02::1:2
*Feb 15 21:35:19.862: type SOLICIT(1), xid 13518535
*Feb 15 21:35:19.862: option ELAPSED-TIME(8), len 2
*Feb 15 21:35:19.862:   elapsed-time 0
*Feb 15 21:35:19.862: option CLIENTID(1), len 10
*Feb 15 21:35:19.862:   0003000100070E036E65
*Feb 15 21:35:19.862: option IA-NA(3), len 12
*Feb 15 21:35:19.862:   IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:19.862: option IA-PD(25), len 12
*Feb 15 21:35:19.862:   IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:19.862: option ORO(6), len 4
*Feb 15 21:35:19.862:   DNS-SERVERS,DOMAIN-LIST

PE Debugs: RELAY-FORWARD w/ SOLICIT

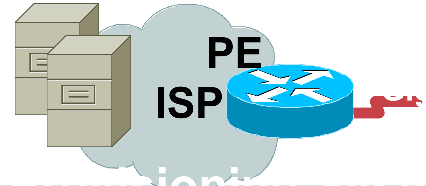

*Feb 15 21:35:19.862: IPv6 DHCP_RELAY: Relaying SOLICIT from FE80::207:EFF:FE03:6E65 on GigE3/1 [PE received SOLICIT request from CPE]
*Feb 15 21:35:19.862: IPv6 DHCP_RELAY: to 2001:420:8:1:5::2 via GigabitEthernet0/1
*Feb 15 21:35:19.862: IPv6 DHCP: Sending RELAY-FORWARD to 2001:420:8:1:5::2 on GigabitEthernet0/1 next hop FE80::201:97FF:FE39:2070 [Forwarding the SOLICIT message to DHCPv6 server]

```
*Feb 15 21:35:19.862: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:19.862:  src 2001:420:8:1:1::2
*Feb 15 21:35:19.862:  dst 2001:420:8:1:5::2 (GigabitEthernet0/1)
*Feb 15 21:35:19.862:  type RELAY-FORWARD(12), hop 0
*Feb 15 21:35:19.862:  link 2001:420:8:1:6:1:1:1
*Feb 15 21:35:19.862:  peer FE80::207:EFF:FE03:6E65
*Feb 15 21:35:19.862:  option RELAY-MSG(9), len 64
*Feb 15 21:35:19.862:   type SOLICIT(1), xid 13518535
*Feb 15 21:35:19.862:   option ELAPSED-TIME(8), len 2
*Feb 15 21:35:19.862:    elapsed-time 0
*Feb 15 21:35:19.862:   option CLIENTID(1), len 10
*Feb 15 21:35:19.862:    0003000100070E036E65
*Feb 15 21:35:19.862:   option IA-NA(3), len 12
*Feb 15 21:35:19.862:    IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:19.862:   option IA-PD(25), len 12
*Feb 15 21:35:19.862:    IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:19.862:   option ORO(6), len 4
*Feb 15 21:35:19.862:    DNS-SERVERS,DOMAIN-LIST
*Feb 15 21:35:19.862:  option INTERFACE-ID(18), len 4
*Feb 15 21:35:19.862:   0x00000007
```

PE Debugs: RELAY-REPLY w/ ADVERTISE


```
*Feb 15 21:35:19.866: IPv6 DHCP: Received RELAY-REPLY from 2001:420:8:1:5::2 on GigabitEthernet0/1
[PE received ADVERTISE from DHCPv6 server]
*Feb 15 21:35:19.866: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:19.866:  src 2001:420:8:1:5::2 (GigabitEthernet0/1)
*Feb 15 21:35:19.866:  dst 2001:420:8:1:1::2
*Feb 15 21:35:19.866:  type RELAY-REPLY(13), hop 0
*Feb 15 21:35:19.866:  link 2001:420:8:1:6:1:1:1
*Feb 15 21:35:19.866:  peer FE80::207:EFF:FE03:6E65
*Feb 15 21:35:19.866:  option INTERFACE-ID(18), len 4
*Feb 15 21:35:19.866:   0x00000007
*Feb 15 21:35:19.866:  option RELAY-MSG(9), len 206
*Feb 15 21:35:19.866:   type ADVERTISE(2), xid 13518535
*Feb 15 21:35:19.866:   option CLIENTID(1), len 10
*Feb 15 21:35:19.866:    0003000100070E036E65
*Feb 15 21:35:19.866:   option SERVERID(2), len 14
*Feb 15 21:35:19.866:    0001000143BF22B6080020E8FAC0
*Feb 15 21:35:19.866:   option IA-NA(3), len 40
*Feb 15 21:35:19.866:    IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:19.866:    option IAADDR(5), len 24
*Feb 15 21:35:19.866:     IPv6 address 2001:420:8:1:6:1:1:EBF1
*Feb 15 21:35:19.866:     preferred 604800, valid 1209600
*Feb 15 21:35:19.866:   option IA-PD(25), len 41
*Feb 15 21:35:19.866:    IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:19.866:    option IAPREFIX(26), len 25
*Feb 15 21:35:19.866:     preferred 604800, valid 1209600, prefix 2001:420:8::/48
*Feb 15 21:35:19.866:   option DNS-SERVERS(23), len 16
*Feb 15 21:35:19.866:    2001:420:3800:801:A00:20FF:FEE5:63E3
*Feb 15 21:35:19.866:   option DOMAIN-LIST(24), len 14
*Feb 15 21:35:19.866:    v6.cisco.com
```

PE Debugs: ADVERTISE

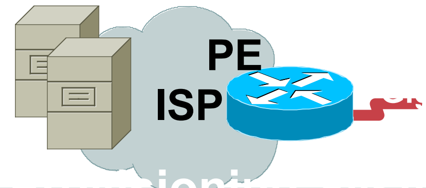


```
*Feb 15 21:35:19.866: IPv6 DHCP: Sending ADVERTISE to FE80::207:EFF:FE03:6E65 on GigE3/1 [PE forwards ADVERTISE message to CPE]
*Feb 15 21:35:19.866: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:19.866:  src FE80::21A:C4FF:FE29:1155
*Feb 15 21:35:19.866:  dst FE80::207:EFF:FE03:6E65 (GigE3/1)
*Feb 15 21:35:19.866:  type ADVERTISE(2), xid 13518535
*Feb 15 21:35:19.866:  option CLIENTID(1), len 10
*Feb 15 21:35:19.866:   0003000100070E036E65
*Feb 15 21:35:19.866:  option SERVERID(2), len 14
*Feb 15 21:35:19.866:   0001000143BF22B6080020E8FAC0
*Feb 15 21:35:19.866:  option IA-NA(3), len 40
*Feb 15 21:35:19.866:   IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:19.866:   option IAADDR(5), len 24
*Feb 15 21:35:19.866:    IPv6 address 2001:420:8:1:6:1:1:EBF1
*Feb 15 21:35:19.866:    preferred 604800, valid 1209600
*Feb 15 21:35:19.866:  option IA-PD(25), len 41
*Feb 15 21:35:19.866:   IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:19.866:   option IAPREFIX(26), len 25
*Feb 15 21:35:19.866:    preferred 604800, valid 1209600, prefix 2001:420:8::/48
*Feb 15 21:35:19.866:  option DNS-SERVERS(23), len 16
*Feb 15 21:35:19.866:   2001:420:3800:801:A00:20FF:FEE5:63E3
*Feb 15 21:35:19.866:  option DOMAIN-LIST(24), len 14
*Feb 15 21:35:19.866:   v6.cisco.com
```

PE Debugs: REQUEST


```
*Feb 15 21:35:20.938: IPv6 DHCP: Received REQUEST from FE80::207:EFF:FE03:6E65 on GigE3/1
[PE received REQUEST from CPE]
*Feb 15 21:35:20.938: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:20.938:  src FE80::207:EFF:FE03:6E65 (GigE3/1)
*Feb 15 21:35:20.938:  dst FF02::1:2
*Feb 15 21:35:20.938:  type REQUEST(3), xid 13530568
*Feb 15 21:35:20.938:  option ELAPSED-TIME(8), len 2
*Feb 15 21:35:20.938:   elapsed-time 0
*Feb 15 21:35:20.938:  option CLIENTID(1), len 10
*Feb 15 21:35:20.938:   0003000100070E036E65
*Feb 15 21:35:20.938:  option IA-NA(3), len 40
*Feb 15 21:35:20.938:   IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:20.938:   option IAADDR(5), len 24
*Feb 15 21:35:20.938:    IPv6 address 2001:420:8:1:6:1:1:EBF1
*Feb 15 21:35:20.938:    preferred 0, valid 0
*Feb 15 21:35:20.938:  option IA-PD(25), len 12
*Feb 15 21:35:20.938:   IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:20.938:  option ORO(6), len 4
*Feb 15 21:35:20.938:   DNS-SERVERS,DOMAIN-LIST
*Feb 15 21:35:20.938:  option SERVERID(2), len 14
*Feb 15 21:35:20.938:   0001000143BF22B6080020E8FAC0
```

PE Debugs: RELAY-FORWARD w/ REQUEST


```
**Feb 15 21:35:20.938: IPv6 DHCP: Sending RELAY-FORWARD to 2001:420:8:1:5::2 on
GigabitEthernet0/1 next hop FE80::201:97FF:FE39:2070 [PE forwards REQUEST to DHCPv6 server]
*Feb 15 21:35:20.938: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:20.938:  src 2001:420:8:1:1::2
*Feb 15 21:35:20.938:  dst 2001:420:8:1:5::2 (GigabitEthernet0/1)
*Feb 15 21:35:20.938:  type RELAY-FORWARD(12), hop 0
*Feb 15 21:35:20.938:  link 2001:420:8:1:6:1:1:1
*Feb 15 21:35:20.938:  peer FE80::207:EFF:FE03:6E65
*Feb 15 21:35:20.938:  option RELAY-MSG(9), len 110
*Feb 15 21:35:20.938:   type REQUEST(3), xid 13530568
*Feb 15 21:35:20.938:   option ELAPSED-TIME(8), len 2
*Feb 15 21:35:20.938:    elapsed-time 0
*Feb 15 21:35:20.938:   option CLIENTID(1), len 10
*Feb 15 21:35:20.938:    0003000100070E036E65
*Feb 15 21:35:20.938:   option IA-NA(3), len 40
*Feb 15 21:35:20.938:    IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:20.938:    option IAADDR(5), len 24
*Feb 15 21:35:20.938:     IPv6 address 2001:420:8:1:6:1:1:EBF1
*Feb 15 21:35:20.938:     preferred 0, valid 0
*Feb 15 21:35:20.938:   option IA-PD(25), len 12
*Feb 15 21:35:20.938:    IAID 0x00020001, T1 0, T2 0
*Feb 15 21:35:20.938:   option ORO(6), len 4
*Feb 15 21:35:20.938:    DNS-SERVERS,DOMAIN-LIST
*Feb 15 21:35:20.938:   option SERVERID(2), len 14
*Feb 15 21:35:20.938:    0001000143BF22B6080020E8FAC0
```

PE Debugs: RELAY-REPLY w/ REPLY


```
*Feb 15 21:35:20.942: IPv6 DHCP: Received RELAY-REPLY from 2001:420:8:1:5::2 on GigabitEthernet0/1
[PE received REPLY from DHCPv6 server]
*Feb 15 21:35:20.942: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:20.942:  src 2001:420:8:1:5::2 (GigabitEthernet0/1)
*Feb 15 21:35:20.942:  dst 2001:420:8:1:1::2
*Feb 15 21:35:20.942:  type RELAY-REPLY(13), hop 0
*Feb 15 21:35:20.942:  link 2001:420:8:1:6:1:1:1
*Feb 15 21:35:20.942:  peer FE80::207:EFF:FE03:6E65
*Feb 15 21:35:20.942:  option INTERFACE-ID(18), len 4
*Feb 15 21:35:20.942:   0x00000007
*Feb 15 21:35:20.942:  option RELAY-MSG(9), len 206
*Feb 15 21:35:20.942:   type REPLY(7), xid 13530568
*Feb 15 21:35:20.942:   option CLIENTID(1), len 10
*Feb 15 21:35:20.942:    0003000100070E036E65
*Feb 15 21:35:20.942:   option SERVERID(2), len 14
*Feb 15 21:35:20.942:    0001000143BF22B6080020E8FAC0
*Feb 15 21:35:20.942:   option IA-NA(3), len 40
*Feb 15 21:35:20.942:    IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:20.942:    option IAADDR(5), len 24
*Feb 15 21:35:20.942:     IPv6 address 2001:420:8:1:6:1:1:EBF1
*Feb 15 21:35:20.942:     preferred 604800, valid 1209600
*Feb 15 21:35:20.942:   option IA-PD(25), len 41
*Feb 15 21:35:20.942:    IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:20.942:    option IAPREFIX(26), len 25
*Feb 15 21:35:20.942:     preferred 604800, valid 1209600, prefix 2001:420:8::/48
*Feb 15 21:35:20.942:   option DNS-SERVERS(23), len 16
*Feb 15 21:35:20.942:    2001:420:3800:801:A00:20FF:FEE5:63E3
*Feb 15 21:35:20.942:   option DOMAIN-LIST(24), len 14
*Feb 15 21:35:20.942:    v6.cisco.com
```

PE Debugs: REPLY

*Feb 15 21:35:20.942: IPv6 DHCP: Sending REPLY to FE80::207:EFF:FE03:6E65 on GigE3/1 [PE forwards REPLY message to CPE]
*Feb 15 21:35:20.942: IPv6 DHCP: detailed packet contents
*Feb 15 21:35:20.942: src FE80::21A:C4FF:FE29:1155
*Feb 15 21:35:20.942: dst FE80::207:EFF:FE03:6E65 (GigE3/1)
*Feb 15 21:35:20.942: type REPLY(7), xid 13530568
*Feb 15 21:35:20.942: option CLIENTID(1), len 10
*Feb 15 21:35:20.942: 0003000100070E036E65
*Feb 15 21:35:20.942: option SERVERID(2), len 14
*Feb 15 21:35:20.942: 0001000143BF22B6080020E8FAC0
*Feb 15 21:35:20.942: option IA-NA(3), len 40
*Feb 15 21:35:20.942: IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:20.942: option IAADDR(5), len 24
*Feb 15 21:35:20.942: IPv6 address 2001:420:8:1:6:1:1:EBF1
*Feb 15 21:35:20.942: preferred 604800, valid 1209600
*Feb 15 21:35:20.942: option IA-PD(25), len 41
*Feb 15 21:35:20.942: IAID 0x00020001, T1 302400, T2 483840
*Feb 15 21:35:20.942: option IAPREFIX(26), len 25
*Feb 15 21:35:20.942: preferred 604800, valid 1209600, prefix 2001:420:8::/48 [DHCP-PD]
*Feb 15 21:35:20.946: option DNS-SERVERS(23), len 16
*Feb 15 21:35:20.946: 2001:420:3800:801:A00:20FF:FEE5:63E3
*Feb 15 21:35:20.946: option DOMAIN-LIST(24), len 14
*Feb 15 21:35:20.946: v6.cisco.com

PE Debugs: ND

*Feb 15 21:35:20.970: ICMPv6-ND: Received NS for 2001:420:8:1:6:1:1:EBF1 on GigE3/1 from :: [DAD Request from CPE]
*Feb 15 21:35:21.490: ICMPv6-ND: Sending RA to FF02::1 on GigE3/1
*Feb 15 21:35:21.490: ICMPv6-ND: MTU = 1500
*Feb 15 21:35:21.974: ICMPv6-ND: Received NA for 2001:420:8:1:6:1:1:EBF1 on GigE3/1 from 2001:420:8:1:6:1:1:EBF1 [CPE Assigns Address & Sends NA to PE]

*Feb 15 21:35:24.866: ICMPv6-ND: DELAY -> PROBE: FE80::207:EFF:FE03:6E65
*Feb 15 21:35:24.866: ICMPv6-ND: Sending NS for FE80::207:EFF:FE03:6E65 on GigE3/1 [PE sends NS to CPE]
*Feb 15 21:35:24.878: ICMPv6-ND: Received NA for FE80::207:EFF:FE03:6E65 on GigE3/1 from FE80::207:EFF:FE03:6E65 [CPE responds with NA to PE]
*Feb 15 21:35:24.878: ICMPv6-ND: PROBE -> REACH: FE80::207:EFF:FE03:6E65
*Feb 15 21:35:26.102: ICMPv6-ND: Sending RA to FF02::1 on GigE3/1
*Feb 15 21:35:26.102: ICMPv6-ND: MTU = 1500
*Feb 15 21:35:28.942: ICMPv6-ND: Received NS for FE80::21A:C4FF:FE29:1155 on GigE3/1 from FE80::207:EFF:FE03:6E65 [PE receives NS from CPE for it's Link-local]
*Feb 15 21:35:28.942: ICMPv6-ND: Sending NA for FE80::21A:C4FF:FE29:1155 on GigE3/1 [PE send NA to CPE]
*Feb 15 21:35:30.302: ICMPv6-ND: Sending RA to FF02::1 on GigE3/1
*Feb 15 21:35:30.302: ICMPv6-ND: MTU = 1500

CPE Router Configuration

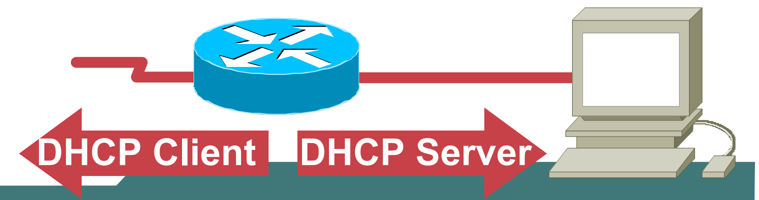


```
ip dhcp pool CPEv4
  network 192.168.51.0 255.255.255.0
  dns-server 80.10.0.1
  domain-name cisco.com
  default-router 80.10.0.1
!
ip multicast-routing
ipv6 unicast-routing
ipv6 dhcp pool v6transfer-pool
 dns-server 2001:420:3800:801:A00:20FF:FEE5:63E3
 domain-name v6.cisco.com
!
interface Ethernet0
 ip address 192.168.51.1 255.255.255.0
 ip pim sparse-mode
 ip virtual-reassembly
 load-interval 30
 ipv6 address v6Prefix 0:0:0:1::/64 eui-64
 ipv6 enable
 ipv6 nd other-config-flag
 ipv6 nd ra interval 5
 ipv6 dhcp server v6transfer-pool
 hold-queue 2048 out
```

```
interface Ethernet1
 ip pim sparse-mode
 ip virtual-reassembly
 load-interval 30
 ipv6 address autoconfig default
 ipv6 enable
 ipv6 nd ra suppress
 ipv6 dhcp client pd v6Prefix
 ipv6 rip RIP enable
 no keepalive
 hold-queue 2048 in
!
ip pim rp-address 10.89.240.226
!
ipv6 router rip RIP
 redistribute connected
```

CPE Router: ND

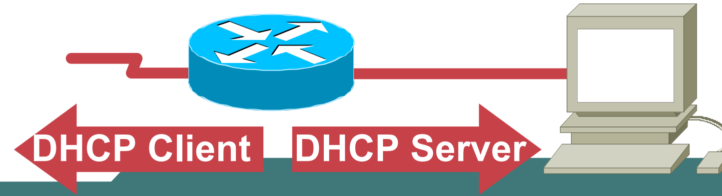

*Mar 2 02:44:54.349: ICMPv6-ND: Received RA from FE80::21A:C4FF:FE29:1155 on Ethernet1
*Mar 2 02:44:54.349: ICMPv6-ND: Selected new default router FE80::21A:C4FF:FE29:1155 on Eth1
*Mar 2 02:44:54.353: ICMPv6-ND: checking DHCP
*Mar 2 02:44:54.353: ICMPv6-ND: stateless DHCP
*Mar 2 02:44:54.357: ICMPv6-ND: statefull DHCP
*Mar 2 02:44:54.357: ICMPv6-ND: M bit set; checking prefix delegation DHCP
*Mar 2 02:44:54.357: ICMPv6-ND: O bit set; [Since M and O bit are set, do statefull DHCPv6]

*Mar 2 02:45:02.709: ICMPv6-ND: Sending NS for FE80::207:EFF:FE03:6E65 on Ethernet1 [DAD Request for Linklocal Address]
*Mar 2 02:45:03.709: ICMPv6-ND: DAD: FE80::207:EFF:FE03:6E65 is unique.
*Mar 2 02:45:03.709: ICMPv6-ND: Sending NA for FE80::207:EFF:FE03:6E65 on Ethernet1
*Mar 2 02:45:03.709: ICMPv6-ND: Linklocal FE80::207:EFF:FE03:6E65 on Ethernet1, Up
*Mar 2 02:45:03.717: ICMPv6-ND: Address FE80::207:EFF:FE03:6E65/10 is up on Ethernet1

*Mar 2 02:45:04.221: ICMPv6-ND: Received RA from FE80::21A:C4FF:FE29:1155 on Ethernet1
*Mar 2 02:45:04.225: ICMPv6-ND: checking stateless DHCP
*Mar 2 02:45:04.225: ICMPv6-ND: O bit set;
*Mar 2 02:45:06.509: ICMPv6-ND: Prefix Information change for 2001:420:8::/48 [DHCP-PD Prefix]
*Mar 2 02:45:06.509: ICMPv6-ND: Adding prefix 2001:420:8::/48 to Ethernet0
*Mar 2 02:45:06.513: ICMPv6-ND: Sending NS for 2001:420:8:1:7::1 on Ethernet0
*Mar 2 02:45:06.513: ICMPv6-ND: Prefix Information change for 2001:420:8:1:6:1:1:EBF1/128
*Mar 2 02:45:06.517: ICMPv6-ND: Adding prefix 2001:420:8:1:6:1:1:EBF1/128 to Ethernet1
*Mar 2 02:45:06.517: ICMPv6-ND: Sending NS for 2001:420:8:1:6:1:1:EBF1 on Ethernet1
*Mar 2 02:45:07.517: ICMPv6-ND: DAD: 2001:420:8:1:6:1:1:EBF1 is unique.
*Mar 2 02:45:07.517: ICMPv6-ND: Sending NA for 2001:420:8:1:6:1:1:EBF1 on Ethernet1
*Mar 2 02:45:07.517: ICMPv6-ND: Address 2001:420:8:1:6:1:1:EBF1/128 is up on Ethernet1

CPE Router: ND

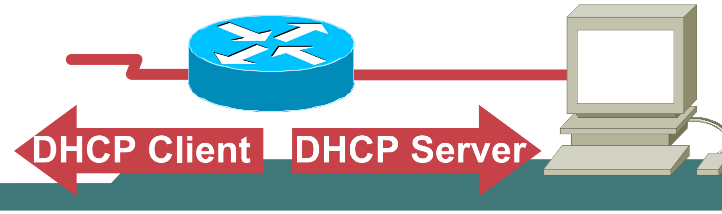


```
*Mar 2 02:45:07.193: ICMPv6-ND: Request to send RA for FE80::207:EFF:FE03:6E64
*Mar 2 02:45:07.193: ICMPv6-ND: Sending RA from FE80::207:EFF:FE03:6E64 to FF02::1 on Ether0
*Mar 2 02:45:07.193: ICMPv6-ND: Prefix = 2001:420:8:1::/64 onlink autoconfig
*Mar 2 02:45:07.193: ICMPv6-ND:          1209600/604800 (valid/preferred)
*Mar 2 02:45:07.513: ICMPv6-ND: DAD: 2001:420:8:1:7::1 is unique.
*Mar 2 02:45:07.513: ICMPv6-ND: Sending NA for 2001:420:8:1:7::1 on Ethernet0
*Mar 2 02:45:07.513: ICMPv6-ND: Address 2001:420:8:1:7::1/80 is up on Ethernet0

*Mar 2 02:45:07.717: ICMPv6-ND: STALE -> DELAY: FE80::21A:C4FF:FE29:1155
*Mar 2 02:45:10.353: ICMPv6-ND: Received NS for FE80::207:EFF:FE03:6E65 on Ether1 from
FE80::21A:C4FF:FE29:1155 [PE to CPE]
*Mar 2 02:45:10.353: ICMPv6-ND: Sending NA for FE80::207:EFF:FE03:6E65 on Ether1 [CPE to PE]

*Mar 2 02:45:12.717: ICMPv6-ND: DELAY -> PROBE: FE80::21A:C4FF:FE29:1155
*Mar 2 02:45:12.717: ICMPv6-ND: Sending NS for FE80::21A:C4FF:FE29:1155 on Ether1 [CPE to PE]
*Mar 2 02:45:12.733: ICMPv6-ND: Received NA for FE80::21A:C4FF:FE29:1155 on Ether1 from
FE80::21A:C4FF:FE29:1155 [PE to CPE]
*Mar 2 02:45:12.737: ICMPv6-ND: PROBE -> REACH: FE80::21A:C4FF:FE29:1155
```

CPE Router: SOLICIT

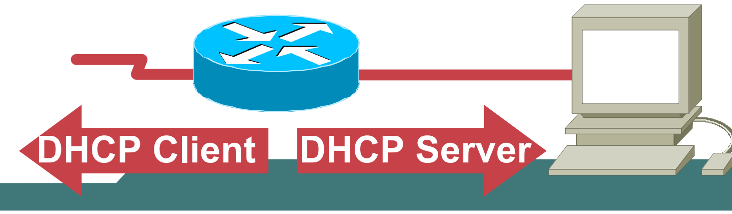


```
*Mar 2 03:39:22.613: IPv6 DHCP: Sending SOLICIT to FF02::1:2 on Ethernet1
*Mar 2 03:39:22.613: IPv6 DHCP: detailed packet contents
*Mar 2 03:39:22.613:  src FE80::207:EFF:FE03:6E65
*Mar 2 03:39:22.613:  dst FF02::1:2 (Ethernet1) [All_DHCP_Relay_Agents_and_Servers Address]
*Mar 2 03:39:22.613:  type SOLICIT(1), xid 16585219
*Mar 2 03:39:22.617:  option ELAPSED-TIME(8), len 2
*Mar 2 03:39:22.617:   elapsed-time 0
*Mar 2 03:39:22.617:  option CLIENTID(1), len 10
*Mar 2 03:39:22.617:   0003000100070E036E65
*Mar 2 03:39:22.617:  option IA-NA(3), len 12
*Mar 2 03:39:22.617:   IAID 0x00020001, T1 0, T2 0
*Mar 2 03:39:22.617:  option IA-PD(25), len 12
*Mar 2 03:39:22.617:   IAID 0x00020001, T1 0, T2 0
*Mar 2 03:39:22.621:  option ORO(6), len 4
*Mar 2 03:39:22.621:   DNS-SERVERS,DOMAIN-LIST
```

CPE Router: ADVERTISE

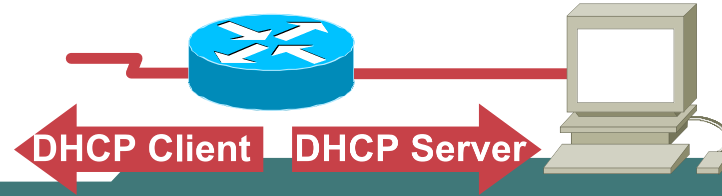


```
*Mar 2 03:39:22.657: IPv6 DHCP: Received ADVERTISE from FE80::21A:C4FF:FE29:1155 on Ether1
*Mar 2 03:39:22.657: IPv6 DHCP: detailed packet contents
*Mar 2 03:39:22.657:  src FE80::21A:C4FF:FE29:1155 (Ethernet1) [Link-local Address of PE]
*Mar 2 03:39:22.657:  dst FE80::207:EFF:FE03:6E65 [Link-local Address of CPE Ether1]

*Mar 2 03:39:22.657:  type ADVERTISE(2), xid 16585219
*Mar 2 03:39:22.657:  option CLIENTID(1), len 10
*Mar 2 03:39:22.657:   0003000100070E036E65
*Mar 2 03:39:22.661:  option SERVERID(2), len 14
*Mar 2 03:39:22.661:   0001000143BF22B6080020E8FAC0
*Mar 2 03:39:22.661:  option IA-NA(3), len 40
*Mar 2 03:39:22.661:   IAID 0x00020001, T1 302400, T2 483840
*Mar 2 03:39:22.661:   option IAADDR(5), len 24
*Mar 2 03:39:22.661:    IPv6 address 2001:420:8:1:6:1:1:EBF1
*Mar 2 03:39:22.661:    preferred 604800, valid 1209600
*Mar 2 03:39:22.665:  option IA-PD(25), len 41
*Mar 2 03:39:22.665:   IAID 0x00020001, T1 302400, T2 483840
*Mar 2 03:39:22.665:   option IAPREFIX(26), len 25
*Mar 2 03:39:22.665:    preferred 604800, valid 1209600, prefix 2001:420:8::/48
*Mar 2 03:39:22.669:  option DNS-SERVERS(23), len 16
*Mar 2 03:39:22.669:   2001:420:3800:801:A00:20FF:FEE5:63E3
*Mar 2 03:39:22.669:  option DOMAIN-LIST(24), len 14
*Mar 2 03:39:22.669:   v6.cisco.com
```

CPE Router: REQUEST

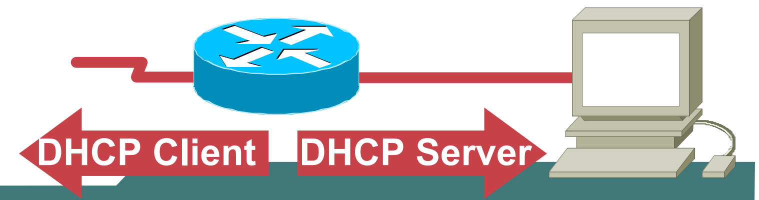


```
*Mar 2 03:39:23.741: IPv6 DHCP: Sending REQUEST to FF02::1:2 on Ethernet1
*Mar 2 03:39:23.741: IPv6 DHCP: detailed packet contents
*Mar 2 03:39:23.745:  src FE80::207:EFF:FE03:6E65 [Link-local Address of CPE Ether1]
*Mar 2 03:39:23.745:  dst FF02::1:2 (Ethernet1) [All_DHCP_Relay_Agents_and_Servers Address]
*Mar 2 03:39:23.745:  type REQUEST(3), xid 16596644
*Mar 2 03:39:23.745:  option ELAPSED-TIME(8), len 2
*Mar 2 03:39:23.745:   elapsed-time 0
*Mar 2 03:39:23.745:  option CLIENTID(1), len 10
*Mar 2 03:39:23.749:   0003000100070E036E65
*Mar 2 03:39:23.749:  option IA-NA(3), len 40
*Mar 2 03:39:23.749:   IAID 0x00020001, T1 0, T2 0
*Mar 2 03:39:23.749:   option IAADDR(5), len 24
*Mar 2 03:39:23.749:    IPv6 address 2001:420:8:1:6:1:1:EBF1
*Mar 2 03:39:23.749:    preferred 0, valid 0
*Mar 2 03:39:23.749:  option IA-PD(25), len 12
*Mar 2 03:39:23.753:   IAID 0x00020001, T1 0, T2 0
*Mar 2 03:39:23.753:  option ORO(6), len 4
*Mar 2 03:39:23.753:   DNS-SERVERS,DOMAIN-LIST
*Mar 2 03:39:23.753:  option SERVERID(2), len 14
*Mar 2 03:39:23.753:   0001000143BF22B6080020E8FAC0
```

CPE Router: REPLY

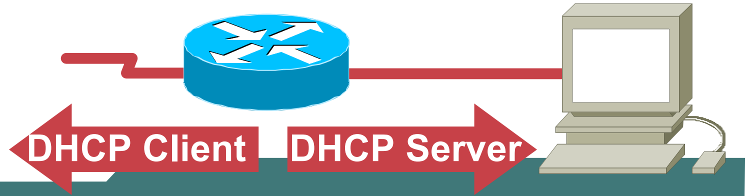


```
*Mar 2 03:39:23.797: IPv6 DHCP: Received REPLY from FE80::21A:C4FF:FE29:1155 on Ether1
*Mar 2 03:39:23.797: IPv6 DHCP: detailed packet contents
*Mar 2 03:39:23.797:  src FE80::21A:C4FF:FE29:1155 (Ethernet1) [Link-local Address of PE]
*Mar 2 03:39:23.797:  dst FE80::207:EFF:FE03:6E65 [Link-local Address of CPE Ether1]
*Mar 2 03:39:23.801:  type REPLY(7), xid 16596644
*Mar 2 03:39:23.801:  option CLIENTID(1), len 10
*Mar 2 03:39:23.801:   0003000100070E036E65
*Mar 2 03:39:23.801:  option SERVERID(2), len 14
*Mar 2 03:39:23.801:   0001000143BF22B6080020E8FAC0
*Mar 2 03:39:23.801:  option IA-NA(3), len 40
*Mar 2 03:39:23.801:   IAID 0x00020001, T1 302400, T2 483840
*Mar 2 03:39:23.801:   option IAADDR(5), len 24
*Mar 2 03:39:23.805:    IPv6 address 2001:420:8:1:6:1:1:EBF1
*Mar 2 03:39:23.805:    preferred 604800, valid 1209600
*Mar 2 03:39:23.805:  option IA-PD(25), len 41
*Mar 2 03:39:23.805:   IAID 0x00020001, T1 302400, T2 483840
*Mar 2 03:39:23.805:   option IAPREFIX(26), len 25
*Mar 2 03:39:23.805:    preferred 604800, valid 1209600, prefix 2001:420:8::/48
*Mar 2 03:39:23.809:  option DNS-SERVERS(23), len 16
*Mar 2 03:39:23.809:   2001:420:3800:801:A00:20FF:FEE5:63E3
*Mar 2 03:39:23.809:  option DOMAIN-LIST(24), len 14
*Mar 2 03:39:23.809:   v6.cisco.com
```

CPE Router: Ethernet Interfaces

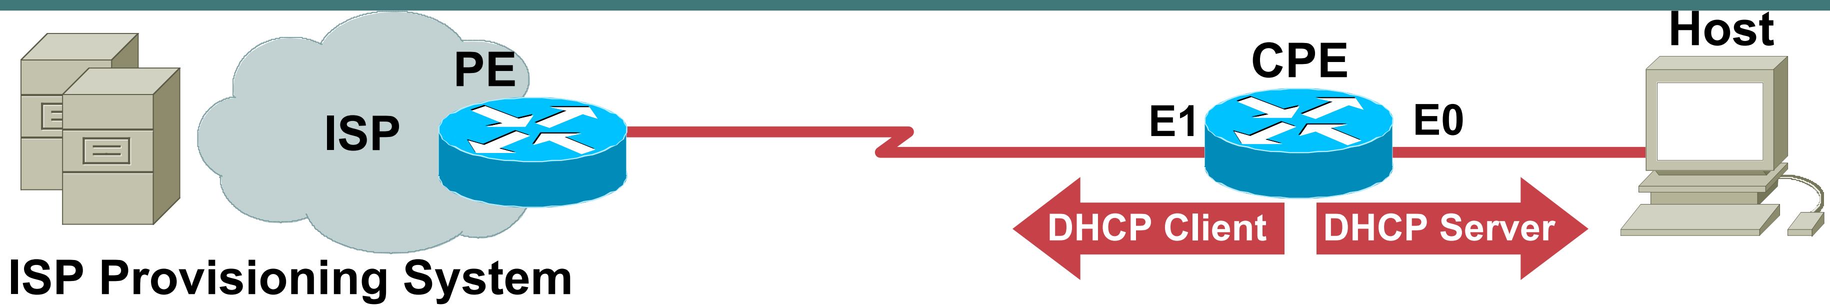

CPE Router#show ipv6 interface e1
cable-modem0 is up, line protocol is up
 IPv6 is enabled, link-local address is FE80::207:EFF:FE03:6E65
 No Virtual link-local address(es):
 Global unicast address(es):
  2001:420:8:1:6:1:1:EBF1, subnet is 2001:420:8:1:6:1:1:EBF1/128 [CAL/PRE] [Address assigned by DHCPv6]
   valid lifetime 1121384 preferred lifetime 516584

CPE Router#show ipv6 interface e0
Ethernet0 is up, line protocol is up
 IPv6 is enabled, link-local address is FE80::207:EFF:FE03:6E64
 No Virtual link-local address(es):
 Global unicast address(es):
  2001:420:8:1:7::1, subnet is 2001:420:8:1::/64 [CAL/PRE] [Address assigned using DHCP-PD]
   valid lifetime 1121385 preferred lifetime 516585

Provisioning Tools

AAA/RADIUS

- RADIUS attributes and IPv6 (RFC3162)—Cisco IOS 12.3(4)T
- RADIUS Server support requires an upgrade (supporting RFC3162) Few RADIUS solutions support RFC3162 functionality today
- Prefix pools and pool names are configurable through AAA
- The following RADIUS attributes as described in RFC 3162 are supported for IPv6: Framed-Interface-Id, Framed-IPv6-Prefix, Login-IPv6-Host, Framed-IPv6-Route, Framed-IPv6-Pool
- IPv6 AAA/RADIUS configuration

    http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/ipv6a_wp.htm

RADIUS Configuration with Permanently Assigned /64:

```
Auth-Type = Local, Password = "foo"
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = "ipv6:prefix=2001:DB8:1:1::/64"
```

Interface Identifier Attribute (Framed-Interface-Id) Can Be Used:

```
Interface-Id = "0:0:0:1",
```

CNR 6.2 - DHCPv6 Supports

- Links and prefixes—similar to DHCPv4's networks and scopes. These define the network topology—each link can have one or more prefixes. Links are optional.
- Policies and options—allows attributes and options to be assigned to links, prefixes, and clients
- VPN support—allows for multiple numbering spaces
- Client classing—allows for clients to be classified and prefixes to be selected based on known clients or packet based expressions
- Static reservations—allows for clients to receive predetermined addresses
- Statistics collection—allows for monitoring the server's activities
- Logging—allows for monitoring the server's activities

IPv6 BB Summary

- Existing IPv4 BB networks can implement/integrate IPv6

  http://www.cisco.com/en/US/products/ps6553/products_data_sheet09186a008011b68d.html

- ISP IPv6 Deployment Scenarios in Broadband Access Networks IETF draft covers ETTH, DSL, WLAN, PLC and Cable:

  <draft-ietf-v6ops-bb-deployment-scenarios-05.txt>

- Some issues in order to deploy native IPv6 in BB Cable networks which are being addressed in DOCSIC 3.0 standardization. These issues are highlighted in:

  <draft-ietf-v6ops-bb-deployment-scenarios-05.txt>

  <draft-mule-cablelabs-docsis3-ipv6-00.txt>

IPv6 SERVICES

IPv6 MULTICAST

IPv4 and IPv6 Multicast Comparison

Service	IPv4 Solution	IPv6 Solution
Addressing Range	32-bit, Class D	128-bit (112-bit Group)
Routing	Protocol Independent, All IGPs and MBGP	Protocol Independent, All IGPs and MBGP with v6 mcast SAFI
Forwarding	~~PIM-DM~~, PIM-SM, PIM-SSM, PIM-bidir, PIM-BSR	PIM-SM, PIM-SSM, PIM-bidir, PIM-BSR
Group Management	IGMPv1, v2, v3	MLDv1, v2
Domain Control	Boundary, Border	Scope Identifier
Interdomain Solutions	MSDP Across Independent PIM Domains	Single RP Within Globally Shared Domains

IPv6 Multicast Addresses (RFC 3513)

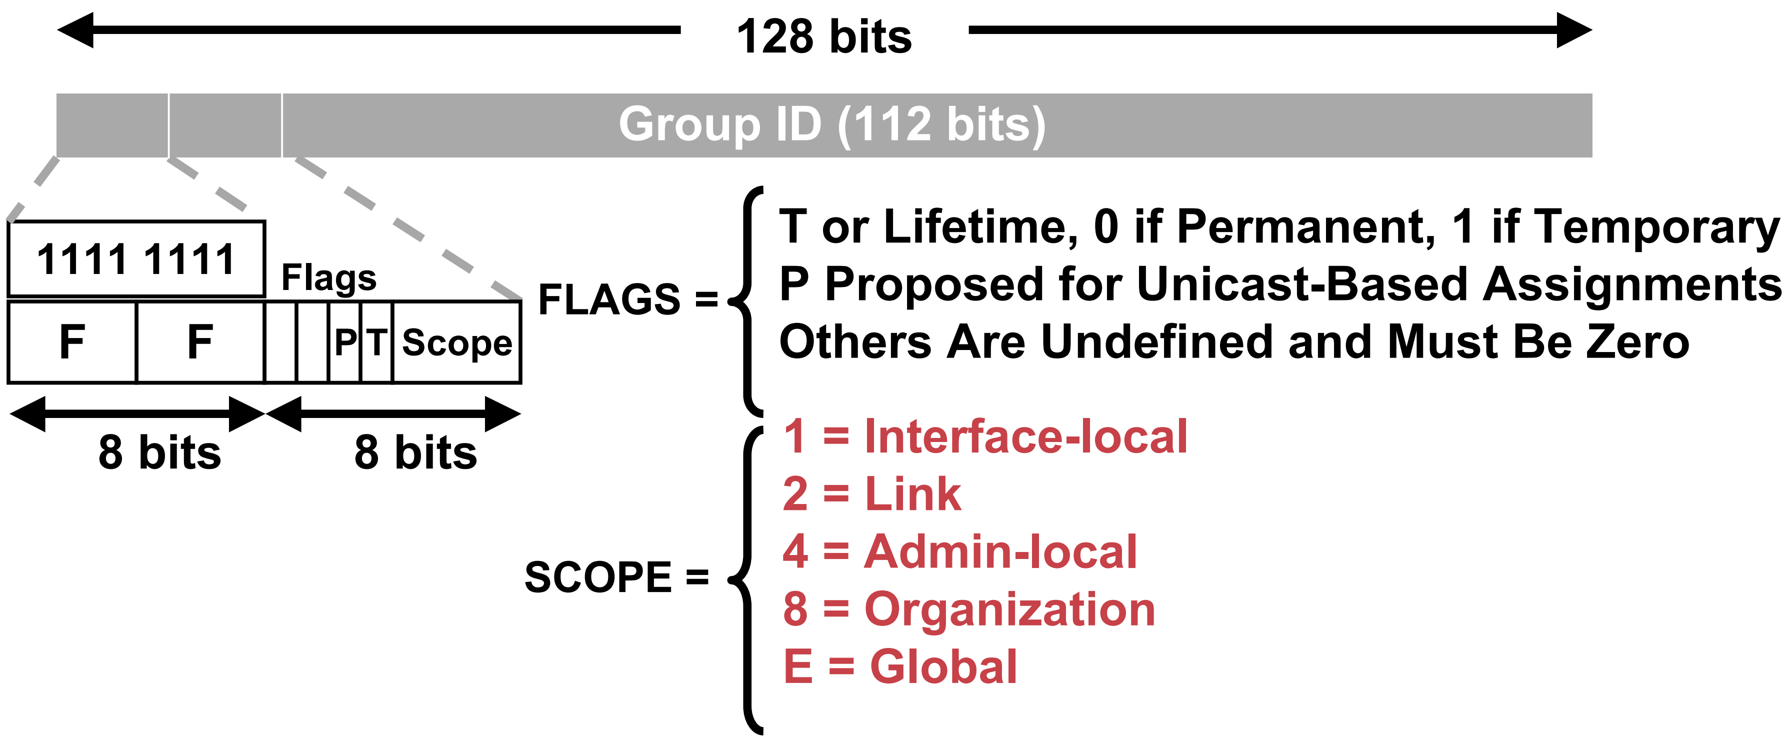

IPv6 Unicast-Based Multicast Addresses (RFC3306)

- Solves the old IPv4 address assignment problem: How can I get global IPv4 multicast addresses?
- In IPv6, if you own an IPv6 unicast address prefix you implicitly own an RFC3306 IPv6 multicast address prefix:

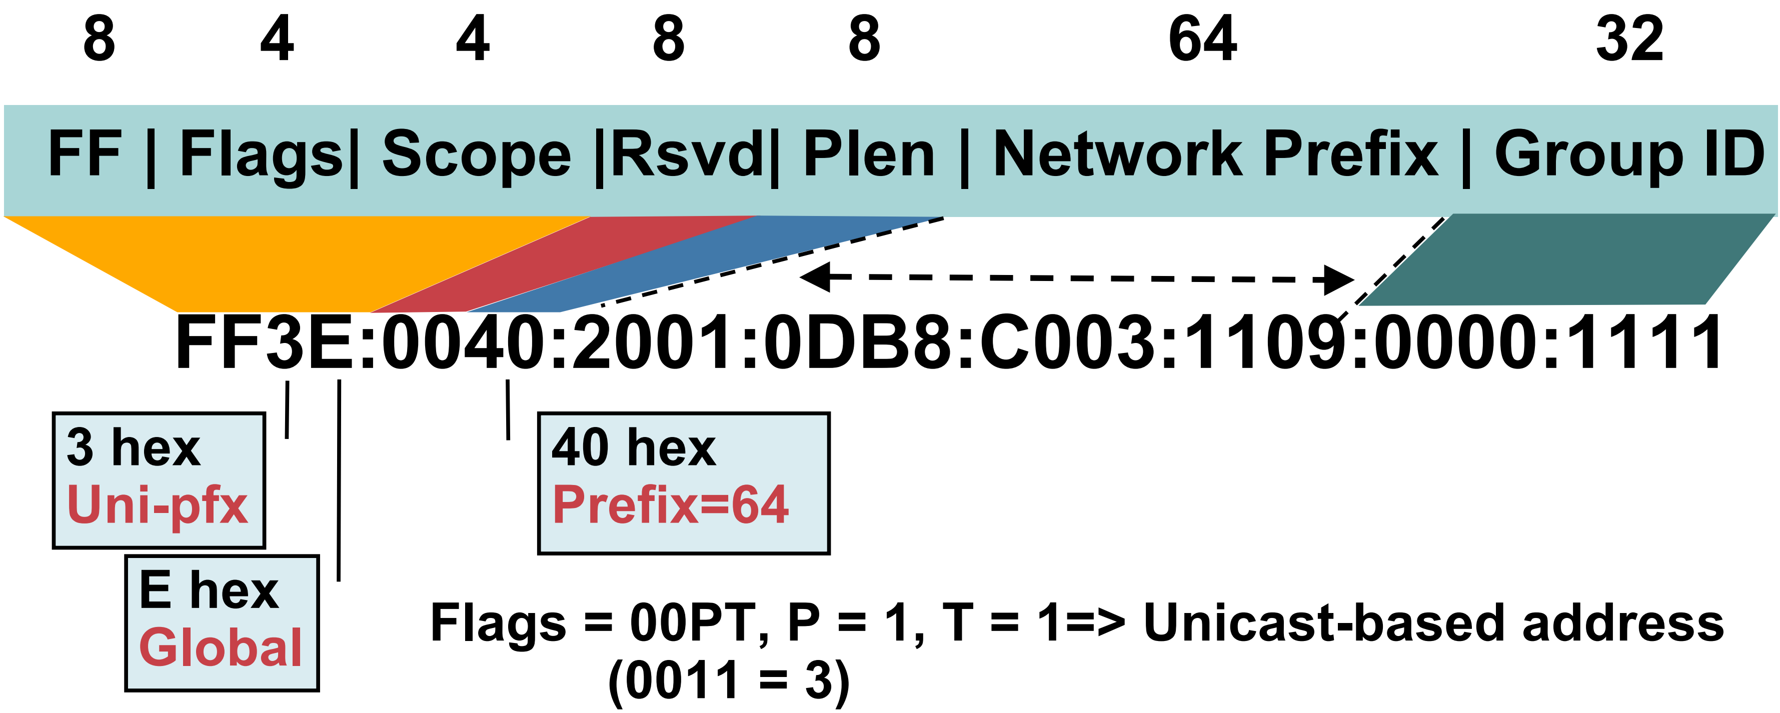

Multicast Listener Discovery: MLD
Multicast Host Membership Control

- MLD is equivalent to IGMP in IPv4
- MLD messages are transported over ICMPv6
- MLD uses link local source addresses
- MLD packets use "Router Alert" option in IPv6 header (RFC2711)
- Version number confusion:
  - MLDv1 (RFC2710) like IGMPv2 (RFC2236)
  - MLDv2 (RFC3810) like IGMPv3 (RFC3376)
- Only MIB available today is for MLDv1
- MLD snooping

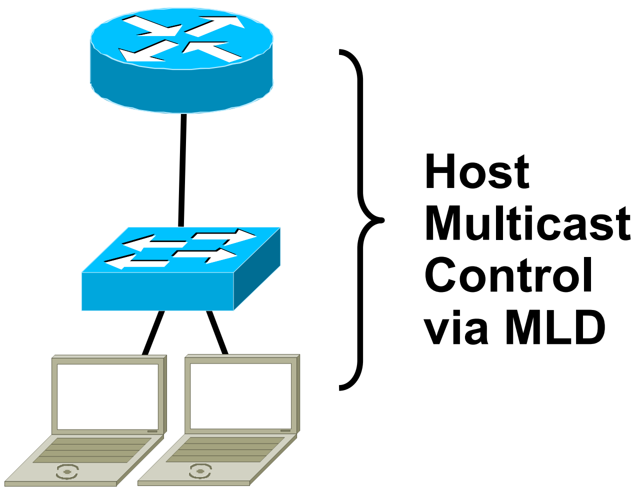

Multicast Interdomain Options

With and Without Rendezvous Points (RP)

SSM, No RPs

ASM Across Multiple Separate PIM Domains, Each With RP, MSDP Peering

ASM Across Single Shared PIM Domain, One RP

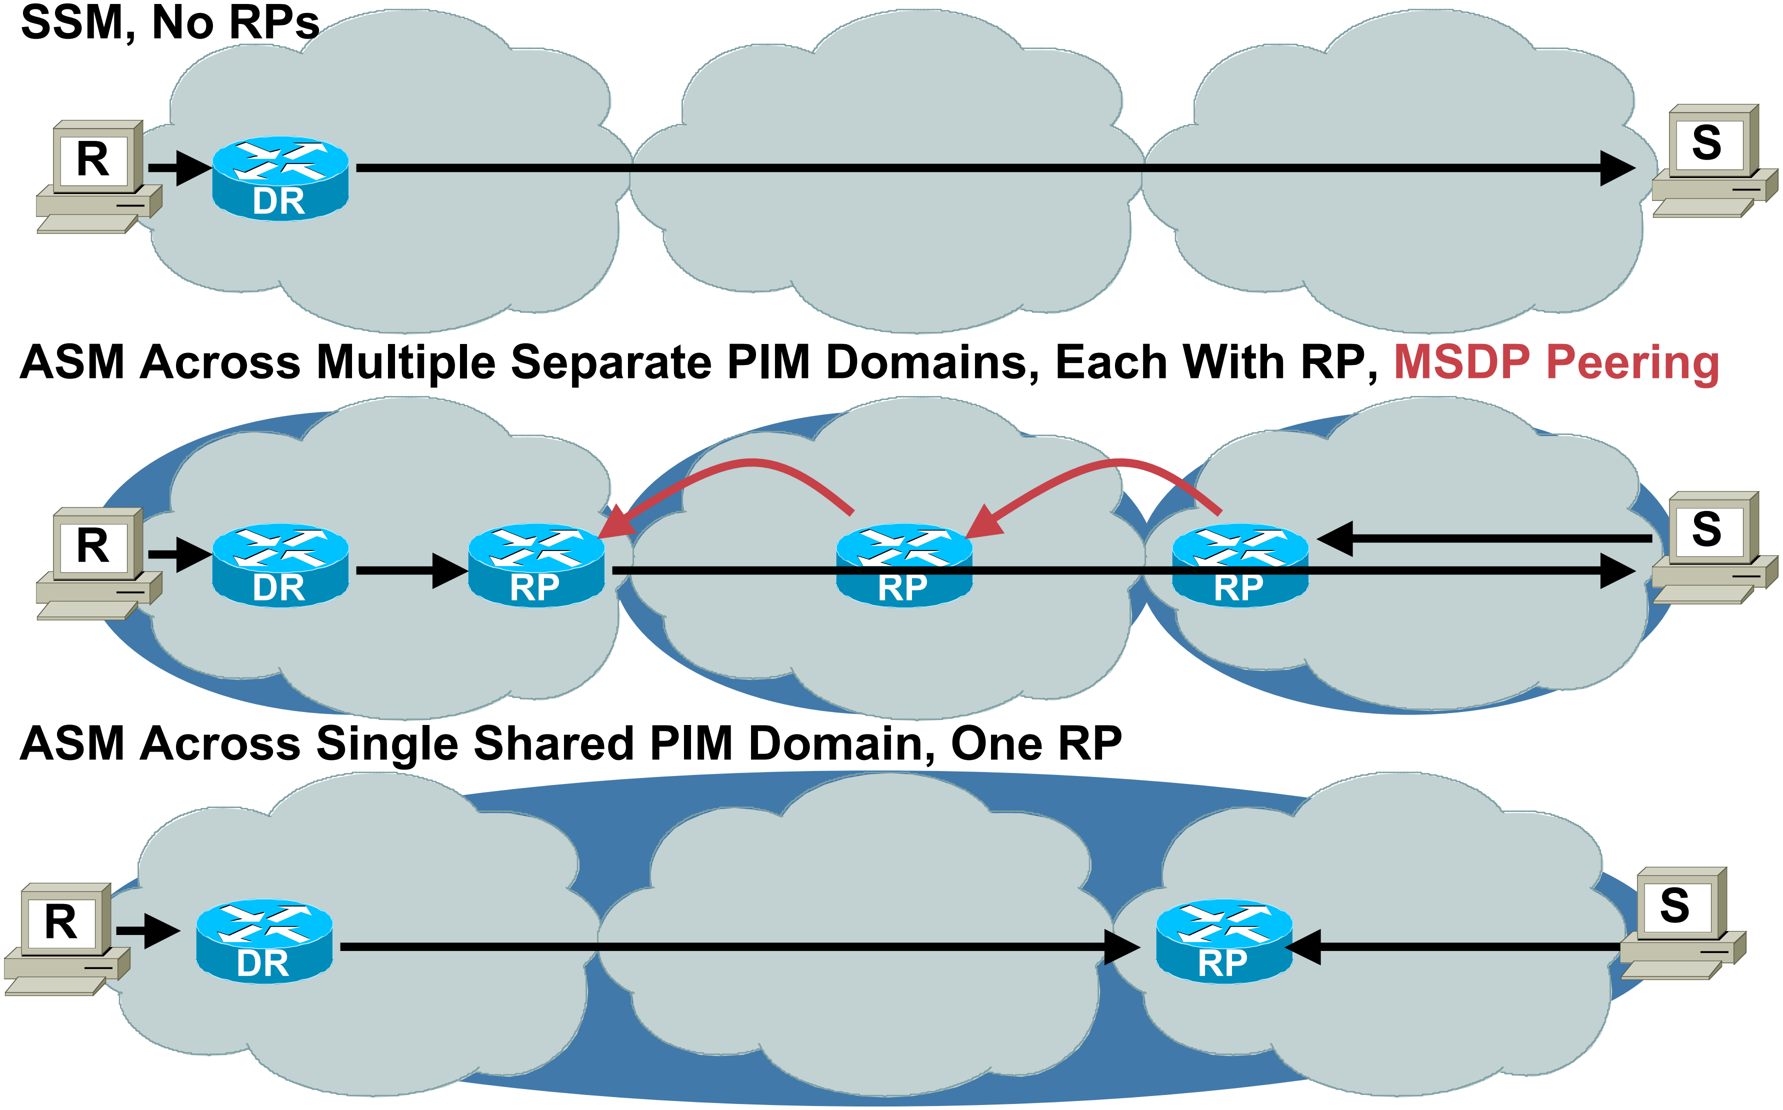

Source Specific Multicast (SSM)

- **NO configuration required other than enabling**

  `ipv6 multicast-routing`

- **SSM group ranges are automatically defined**
- **Very few applications support MLDv2…yet**

```
router#show ipv6 pim range-list
config SSM Exp: never Learnt from : ::
  FF33::/32 Up: 1d00h
  FF34::/32 Up: 1d00h
  FF35::/32 Up: 1d00h
  FF36::/32 Up: 1d00h
  FF37::/32 Up: 1d00h
  FF38::/32 Up: 1d00h
  FF39::/32 Up: 1d00h
  FF3A::/32 Up: 1d00h
  FF3B::/32 Up: 1d00h
  FF3C::/32 Up: 1d00h
  FF3D::/32 Up: 1d00h
  FF3E::/32 Up: 1d00h
  FF3F::/32 Up: 1d00h
```

Rendezvous Point (RP) Deployment Types

- Static RP
  - For PIM-SM and Bidir-PIM
  - Provides group-to-RP mapping, no RP redundancy (yet)
- Boot Strap Router (BSR)
  - Provides group-to-RP mapping AND RP redundancy
- Embedded-RP
  - Easy to deploy
  - Group-to-RP mapping only, no RP redundancy (yet)
  - PIM-SM only (today), no Bidir-PIM

IPv6 Multicast Static RP

- Easier than before as PIM is auto-enabled on every interface

```
ipv6 multicast-routing
!
interface Loopback0
 description IPV6 IPmc RP
 no ip address
 ipv6 address 2001:DB8:C003:110A::1/64
!
ipv6 pim rp-address 2001:DB8:C003:110A::1/64
```

```
ipv6 multicast-routing
!
ipv6 pim rp-address 2001:DB8:C003:110A::1/64
```

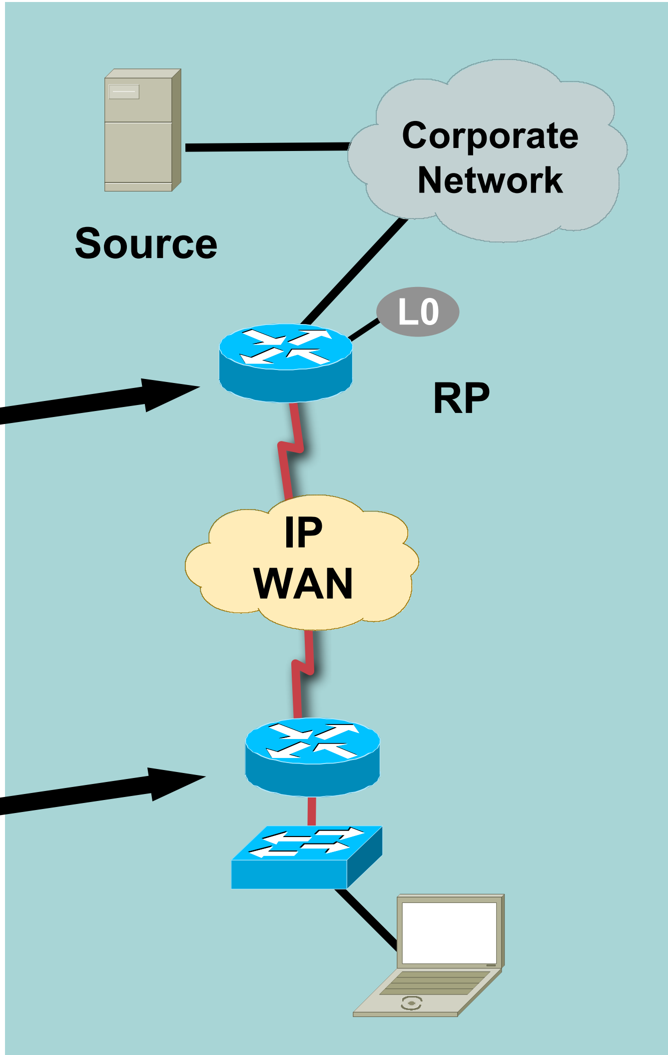

Bidirectional PIM (Bidir)

- The same many-to-many model as before
- Configure Bidir RP and range via the usual `ip pim rp-address` syntax with the optional `bidir` keyword

```
!
ipv6 pim rp-address 2001:DB8:C003:110A::1 bidir
!
#show ipv6 pim range | include BD

Static BD RP: 2001:DB8:C003:110A::1 Exp: never Learnt from : ::
```

IPv6 Multicast PIM BSR: Configuration

```
wan-top#sh run | incl ipv6 pim bsr

ipv6 pim bsr candidate-bsr 2001:DB8:C003:1116::2
ipv6 pim bsr candidate-rp 2001:DB8:C003:1116::2
```

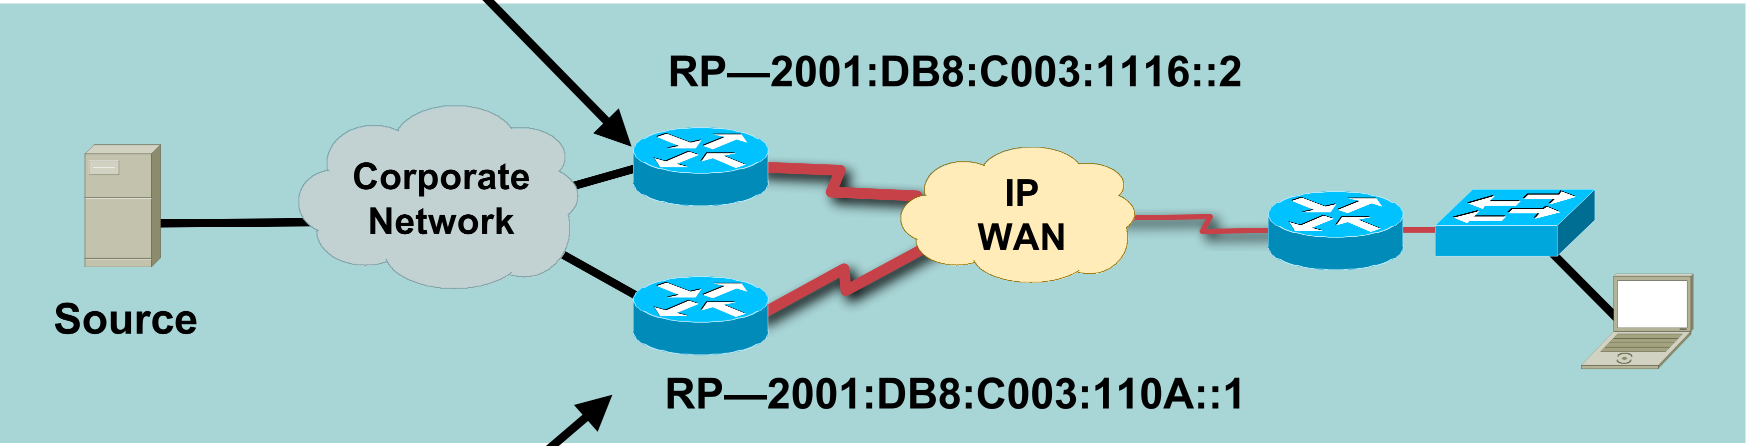


```
wan-bottom#sh run | incl ipv6 pim bsr

ipv6 pim bsr candidate-bsr 2001:DB8:C003:110A::1
ipv6 pim bsr candidate-rp 2001:DB8:C003:110A::1
```

Embedded-RP

- **PIM-SM protocol operations with embedded-RP:**
  - **Intradomain transition into embedded-RP is easy:**
    - **Non-supporting routers simply need to be configured statically or via BSR for the embedded-RPs!**
- **Embedded-RP is just a method to learn ONE RP address for a multicast group:**
  - **It can not replace RP-redundancy as possible with BSR or MSDP/Anycast-RP**
- **Embedded-RP does not (yet) support Bidir-PIM**

RP Redundancy

Potential Anycast RP Alternatives

- **draft-ietf-pim-anycast-rp-xx.txt**
  - Most simple protocol doing exactly what MSDP needs to do in one mesh-group: PIM-SM register messages are unicast forwarded between the redundant RPs
  - (Almost) no operational differences to MSDP for Anycast-RP
- **Prefix-length/Anycast-RP (a.k.a. PriorityCast)**
  - Solution without any new protocol (in that way similar to embedded-RP)—a.k.a. most simple solution?
  - Could support PIM-SM and Bidir-PIM, IPv4 and IPv6
  - Work being done now

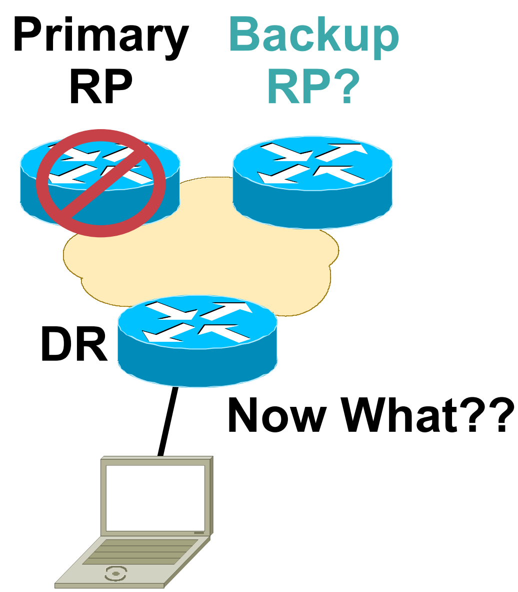

IPv6 Multicast Summary

- One size does NOT fit all
- PIM-SSM is the way to go for one/few-to-many applications, but requires MLDv2 or SSM Mapping and the app to support SSM operation
- Embedded-RP is simple to deploy, but does not currently provide for RP redundancy (in the works)
- PIM-BSR provides for easier RP deployment than static RP and provides for RP redundancy (albeit slow), but is a bit more complicated
- Cisco is working on scalable and highly-available RP deployment methods

IPv6 QoS

IPv6 QoS: Header Fields

- **IPv6 traffic class**

  Exactly the same as TOS field in IPv4

- **IPv6 flow label (RFC 3697)**

  A new 20-bit field in the IPv6 basic header which:

  - **Labels packets belonging to particular flows**
  - **Can be used for special sender requests**

  Per RFC, Flow Label must not be modified by intermediate routers

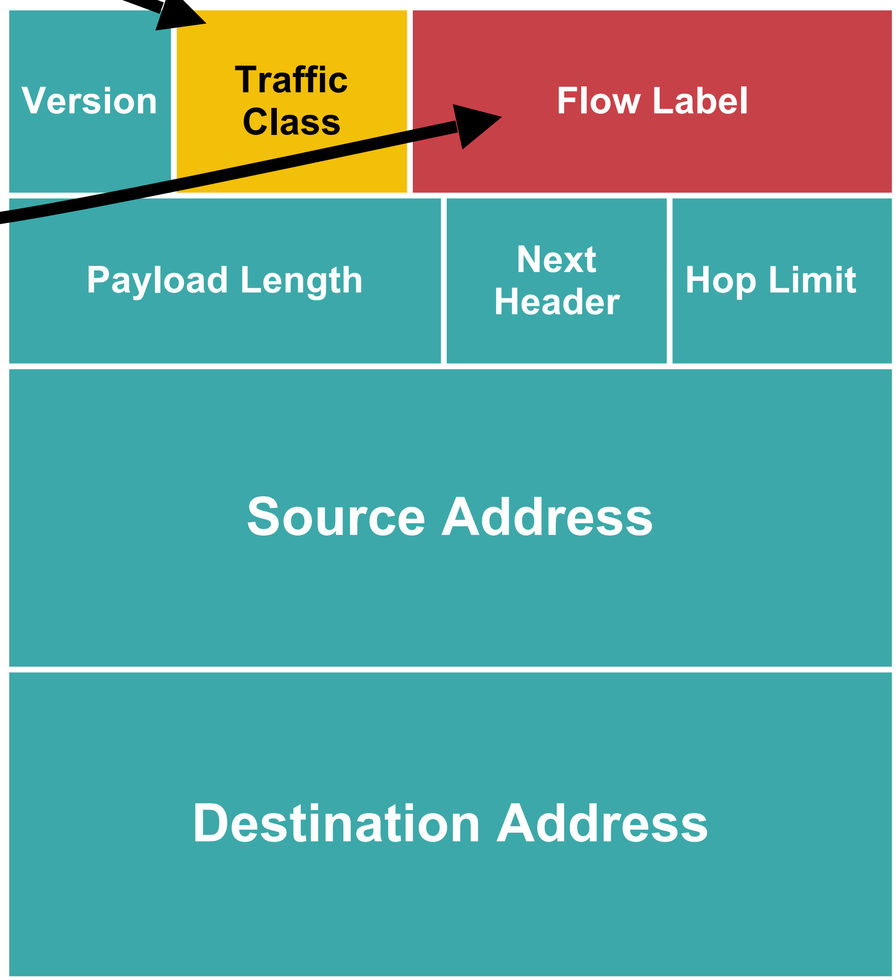

IPv6 QoS Syntax Changes

- **IPv4 syntax has used "ip" following match/set statements**

  Example: `match ip dscp, set ip dscp`

- **Modification in QoS syntax to support IPv6 and IPv4**

  `New match criteria`

  `match dscp—Match DSCP in v4/v6`

  `match precedence—Match Precedence in v4/v6`

  `match protocol ipv6—Match on IPv6 Protocol`

  `New set criteria`

  `set dscp—Set DSCP in v4/v6`

  `set precedence—Set Precedence in v4/v6`

- **Additional support for IPv6 does not always require new Command Line Interface (CLI)**

  Example—WRED

Simple QoS Example: IPv4 and IPv6

```
class-map match-any BRANCH-BULK-DATA
 match access-group name BULK-DATA-IPV6
 match access-group name BULK-DATA
class-map match-all BULK-DATA
 match  dscp af11
!
policy-map RBR-WAN-EDGE
 class BULK-DATA
  bandwidth percent 4
  random-detect
!
policy-map RBR-LAN-EDGE-IN
 class BRANCH-BULK-DATA
  set dscp af11
!
ip access-list extended BULK-DATA
 permit tcp any any eq ftp
 permit tcp any any eq ftp-data
!
ipv6 access-list BULK-DATA-IPV6
 permit tcp any any eq ftp
 permit tcp any any eq ftp-data
```

ACL Match To Set DSCP (If Packets Are Not Already Marked)

```
service-policy input RBR-LAN-EDGE-IN
```

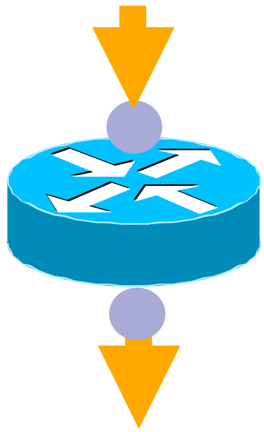

```
service-policy output RBR-WAN-EDGE
```

ACLs to Match for Both IPv4 and IPv6 Packets

IPv6 QoS: Support

- Cisco's current IPv6 QoS implementation supports:
  - Packet classification
  - Queuing—(does support LLQ)—excluding PQ/CQ
  - Traffic shaping
  - WRED
  - Class-based packet marking
  - Policy-based packet marking
- Cisco's current IPv6 QoS implementation does not support:
  - Compressed Real-Time Protocol (CRTP)
  - Network-Based Application Recognition (NBAR)
  - Committed Access Rate (CAR)
  - Priority Queuing (PQ)
  - Custom Queuing (CQ)

IPv6 SECURITY

www.cisco.com/security_services/ciag/documents/v6-v4-threats.pdf

IPv6 Security

- RFC "mandates" privacy and encryption
- Same IPSec you already know
- Two security extension headers defined; all implementations required to support (IPSec)
    - Authentication Header (AH)
    - Encapsulating Security Payload (ESP)
    - Key distribution protocols are under development
    - Support for manual key configuration required
- IPv6 Security is MORE THAN IPSec!
- New concept of privacy addressing
    - On by default in Microsoft XP SP1+
    - Randomly generated address
- Nearly impossible to perform successful network scans

IPv6 Protocol Challenges

- **Inherits many challenges found in IPv4**

  Same applications

  Same TCP, UDP layers

- **Many new features**

  Autoconfiguration (router advertisements)

  ND—Neighbor Discovery (altering ICMPv6 packets)

  DAD—Multiple (bad) addresses

  Mobile IPv6—binding update, etc.

IPv6 Security Considerations

- If all hosts are performing encryption, what happens to…
  - Intrusion detection
  - Intrusion prevention (inline filtering)
  - Virus protection
  - Deep packet inspection
  - Proxies
- The real world will likely implement…
  - Decoupling of end to end encryption (terminate connections on a bulk encryption device)
  - Use of authentication headers providing packet integrity, but not encryption
  - Extensive use of personal (host-based) firewalls and host-based IDS (Cisco Security Agent) to augment network-based security tools

IPv6 Transition Mechanism Challenges

- **Dual stack**
  - Consider security for both protocols
  - Cross v4/v6 abuse
  - Resiliency (shared resources)
- **Tunnels**
  - Bypass firewalls (protocol 41)
  - Relayed DoS attacks from v6 to v4 and vice versa
- **Translation mechanisms**
  - Prevent end-to-end network and transport layer security

Basic IPv6 Packet Filtering
(Access Control List)

When Used for Traffic Filtering, IPv6 Access Control Lists (ACL) Offers the Same Level of Support as in IPv4

- **Every IPv6 ACL has `implicit permit icmp any any nd-na` and permit icmp any any nd-ns**
- **Implicit `deny all` at the end of access list**

```
interface FastEthernet0/0
 ipv6 address 2001:DB8:C003:1101::1/64
 ipv6 traffic-filter V6FILTER in
!
ipv6 access-list V6FILTER
 permit tcp any host 2001:DB8:C003:1102::10 eq web
!
```

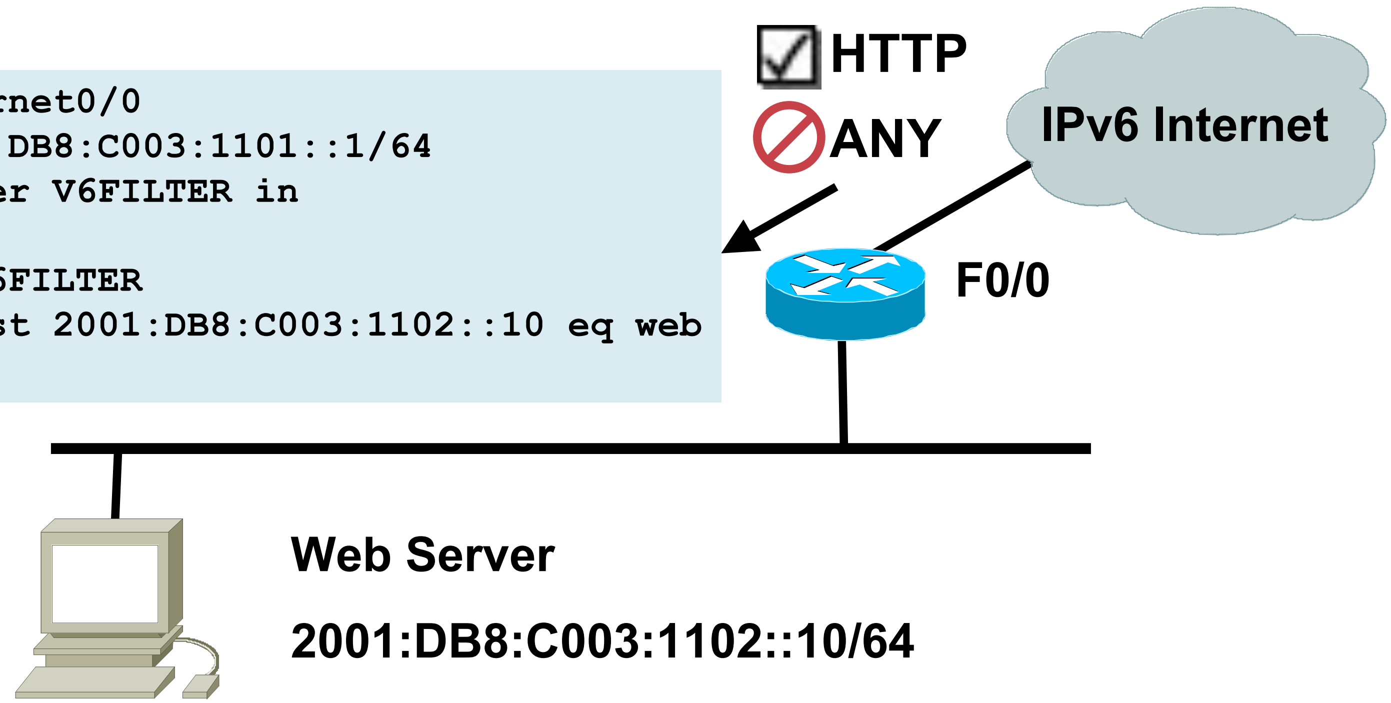

Cisco IOS IPv6 Firewall Feature Set

Example: Nothing New from IPv4

- **Cisco IOS Firewall released 12.3(7)T**

```
ipv6 unicast-routing
ipv6 cef
!
ipv6 inspect audit-trail
ipv6 inspect max-incomplete low 150
ipv6 inspect max-incomplete high 250
ipv6 inspect one-minute low 100
ipv6 inspect one-minute high 200
ipv6 inspect name V6FW tcp timeout 300
ipv6 inspect name V6FW udp
ipv6 inspect name V6FW icmp
!
interface FastEthernet0/0
ipv6 address 2001:DB8:C003:1112::2/64
ipv6 cef
ipv6 traffic-filter EXAMPLE in
ipv6 inspect V6FW in
!
ipv6 access-list EXAMPLE
permit tcp any host 2001:DB8:C003:1113::2 eq www
permit tcp any host 2001:DB8:C003:1113::2 eq ftp
deny ipv6 any any log
```

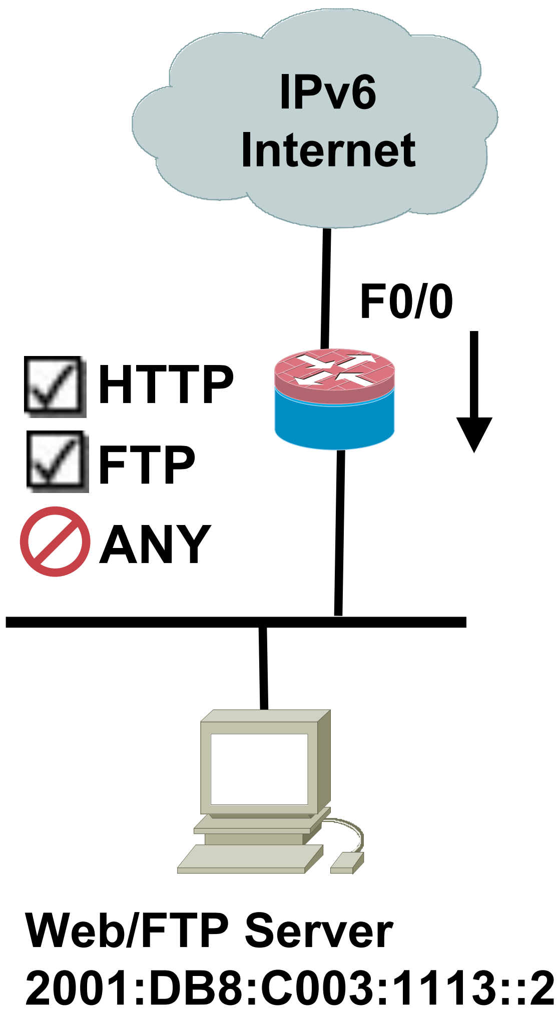

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/ps5761/index.html

PIX 7.0: ACL
Very Similar to Cisco IOS

```
interface Ethernet0
 nameif outside
 ipv6 address 2001:db8:c000:1051::37/64
 ipv6 enable
interface Ethernet1
 nameif inside
 ipv6 address 2001:db8:c000:1052::1/64
 ipv6 enable

ipv6 unicast-routing

ipv6 route outside ::/0 2001:db8:c000:1051::1

ipv6 access-list SECURE permit tcp any host 2001:db8:c000:1052::7 eq telnet
ipv6 access-list SECURE permit icmp6 any 2001:db8:c000:1052::/64

access-group SECURE in interface outside
```

PIX 7.0 and Stateful Inspection

```
pixA# show conn
4 in use, 7 most used
ICMP out fe80::206:d7ff:fe80:2340:0 in fe80::209:43ff:fea4:dd07:0
  idle 0:00:00 bytes 16
UDP out 2001:db8:c000:1051::138:53 in 2001:db8:c000:1052::7:50118
  idle 0:00:02 flags -
TCP out 2001:200:0:8002:203:47ff:fea5:3085:80 in
  2001:db8:c000:1052::7:11009 idle 0:00:14 bytes 8975 flags UfFRIO
TCP out 2001:db8:c000:1051::1:11008 in 2001:db8:c000:1052::7:23
  idle 0:00:04 bytes 411 flags UIOB
```

MOBILE IPv6 (MIPv6)

Mobile IPv6 Benefits

- IPv6 address space enables mobile IP deployment in any kind of large environment
- No foreign agent needed in MIPv6
  - Infrastructure does not need an upgrade to accept Mobile IPv6 Nodes
- IPv6 auto-configuration simplifies Mobile Node (MN) CoA (Care of Address) assignment
- MIPv6 take benefits of IPv6 protocol itself
  - E.g.. option headers, neighbor discovery
- Optimized routing—avoids triangular routing
  - Scale easier but network management challenges
- MN's work transparently even with other nodes that do not support mobility
  - Albeit without route optimization

Mobile IPv6: a Native Extension of IPv6

Unfragmented Packet Example:

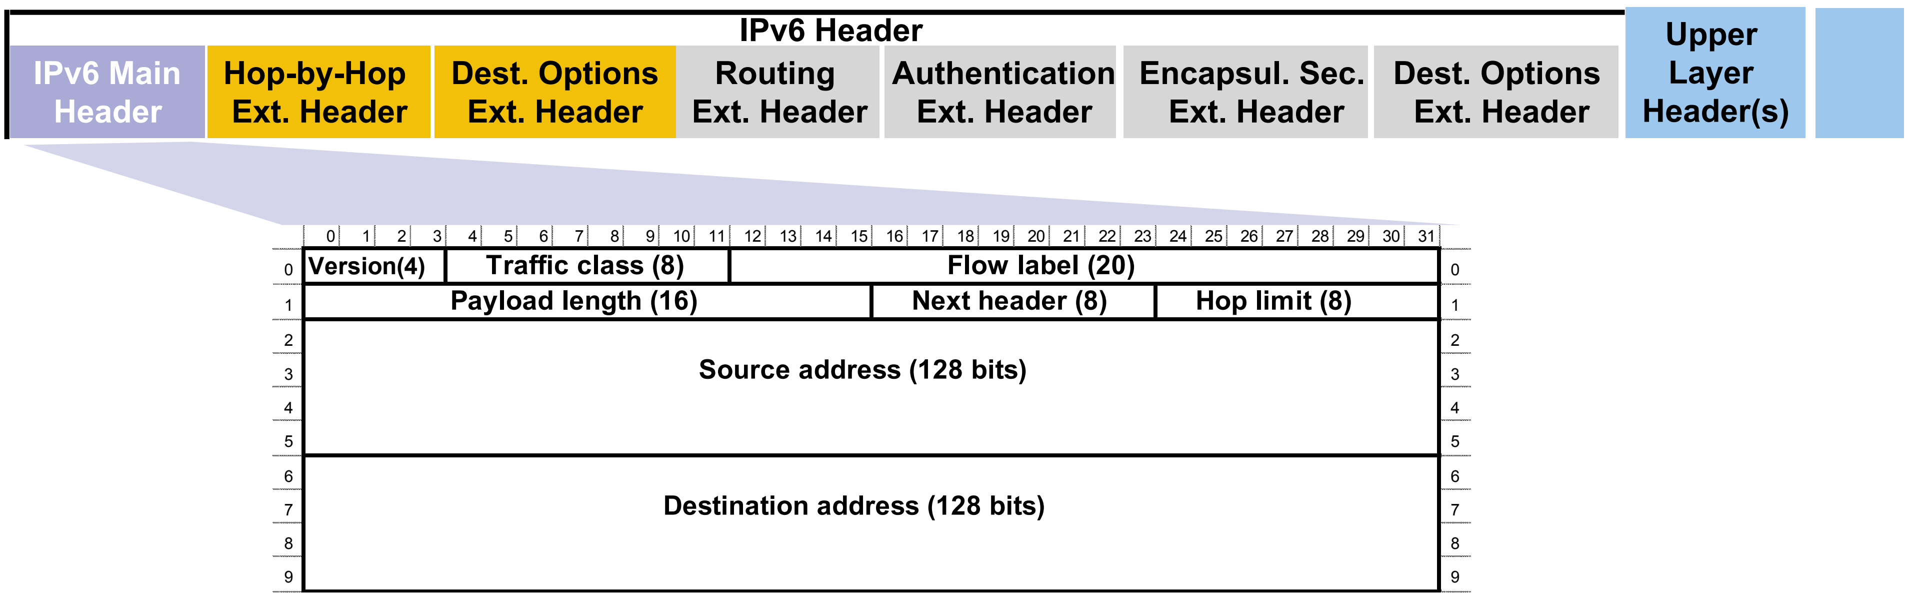

- Take advantage of the IPv6 packet structure as defined in RFC 2460
- Create new extension header—mobility header
- Add new routing header type
- Add new destination option

Mobile IPv6: Key Components

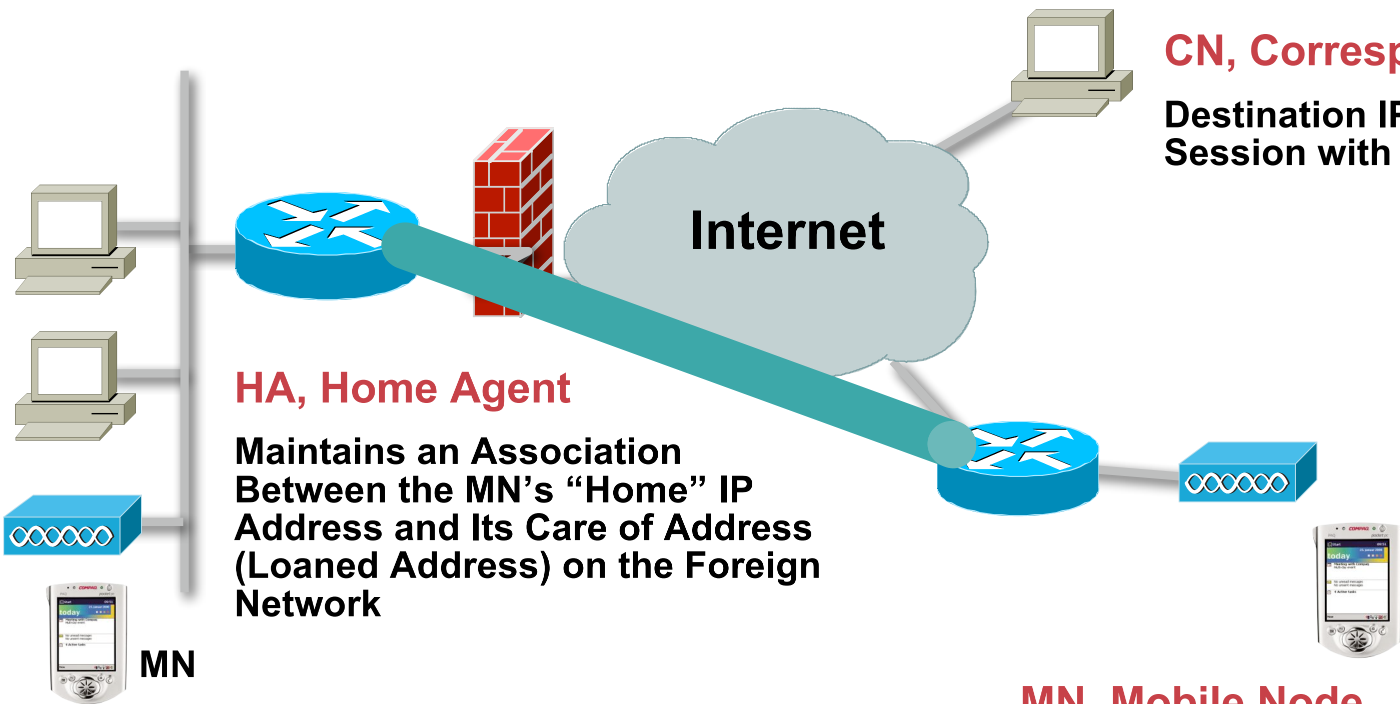

CN, Correspondent Node

Destination IP Host in Session with a Mobile Node

HA, Home Agent

Maintains an Association Between the MN's "Home" IP Address and Its Care of Address (Loaned Address) on the Foreign Network

MN, Mobile Node

An IP Host that Maintains Network Connectivity Using Its "Home" IP Address, Regardless of which Link (or Network) It Is Connected to

MIPv6 Operations: MN on its Home Network

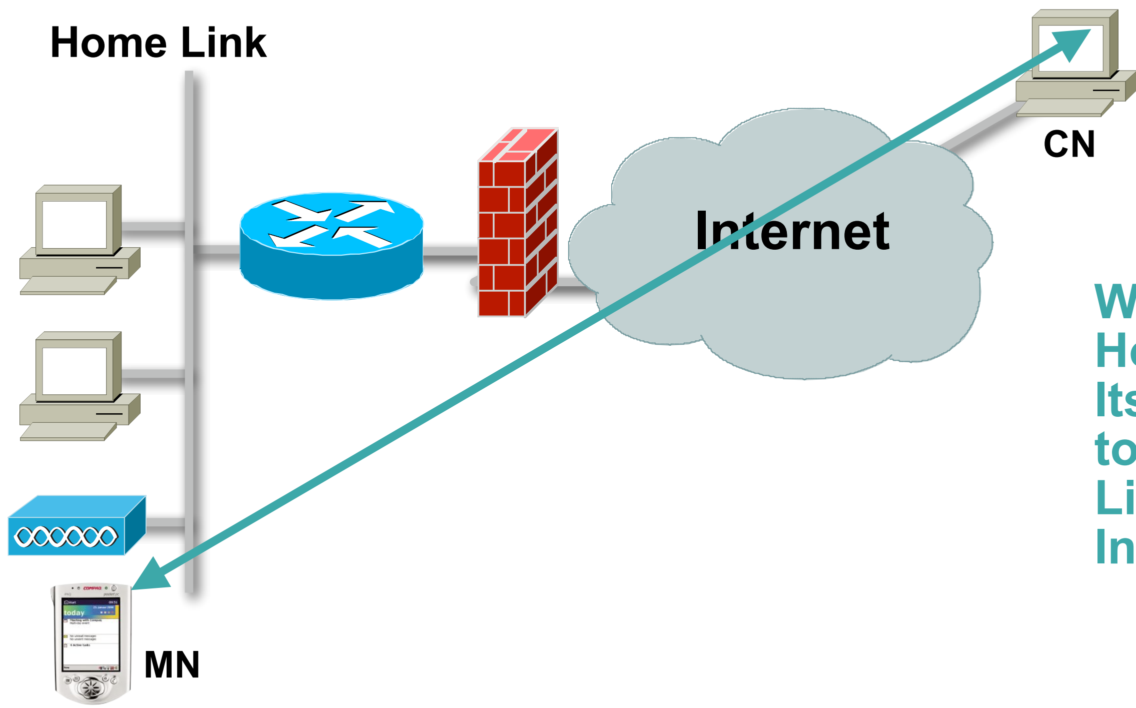

While a Mobile Node Is at Home, Packets Addressed to Its Home Address Are Routed to the Mobile Node's Home Link, Using Conventional Internet Routing Mechanisms

- A Mobile Node (MN) is always expected to be addressable at its home address, whether it is currently attached to its home link or is away from home
- The “home address” is an IP address assigned to MN within its home subnet prefix on its home link

MIPv6 Operations: MN Moving to a New Link

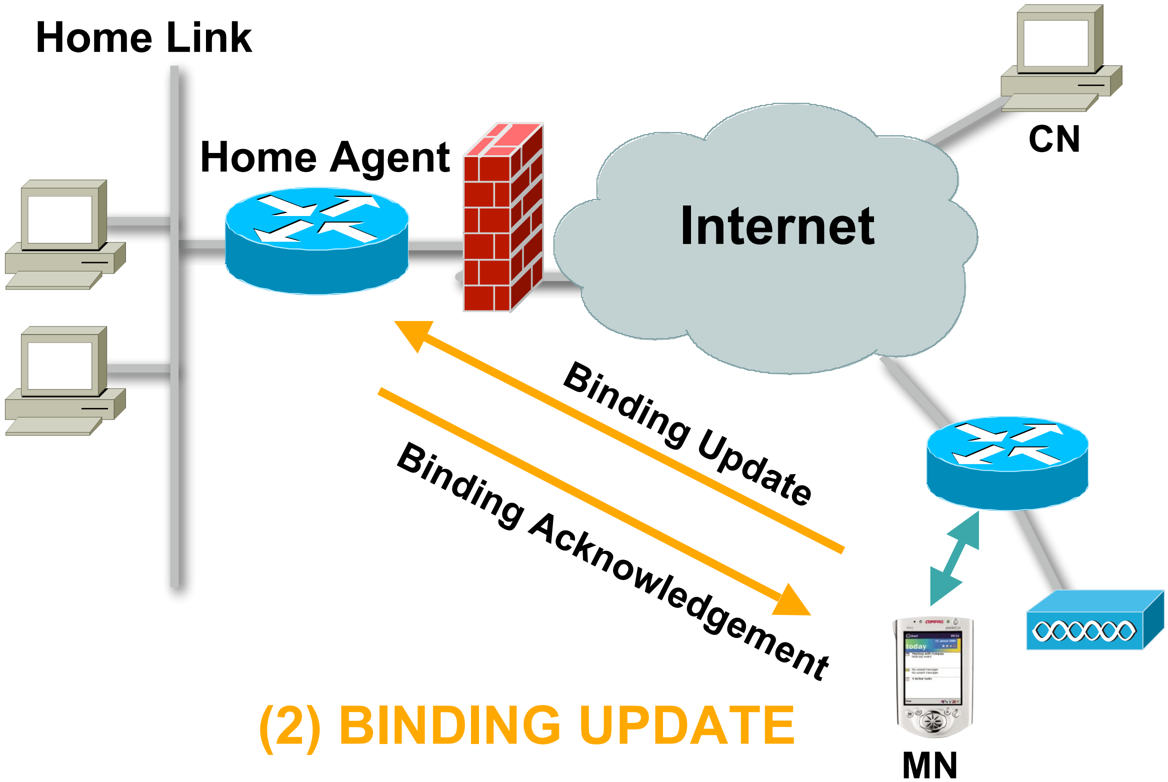

(1) CARE OF ADDRESS

MN Obtains an IPv6 Address in the Visited Network Through Stateless or Stateful Auto-Configuration

(2) BINDING UPDATE

While away from Home, a MN Registers Its Primary Care-of Address with a Router on Its Home Link, Requesting this Router to Function as the "Home Agent" for the MN

Packet Forwarding
Bidirectional Tunneling Mode

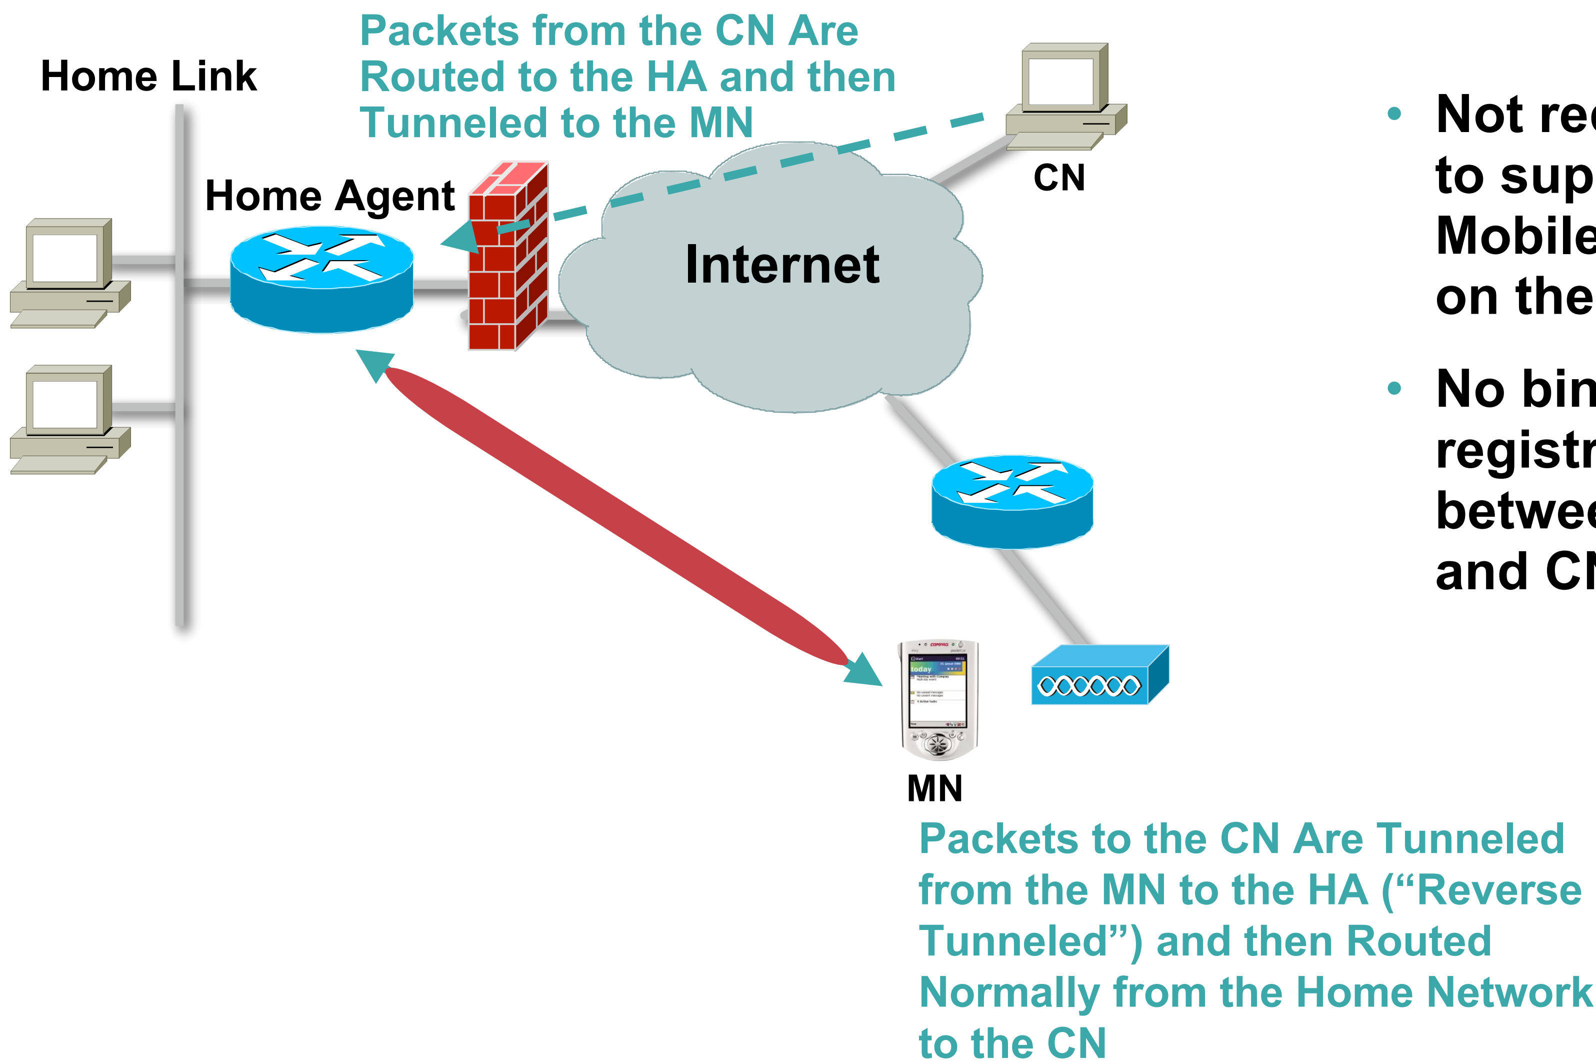

- Not required to support Mobile IPv6 on the CN
- No binding registration between MN and CN

Packet Forwarding
Route Optimization Mode

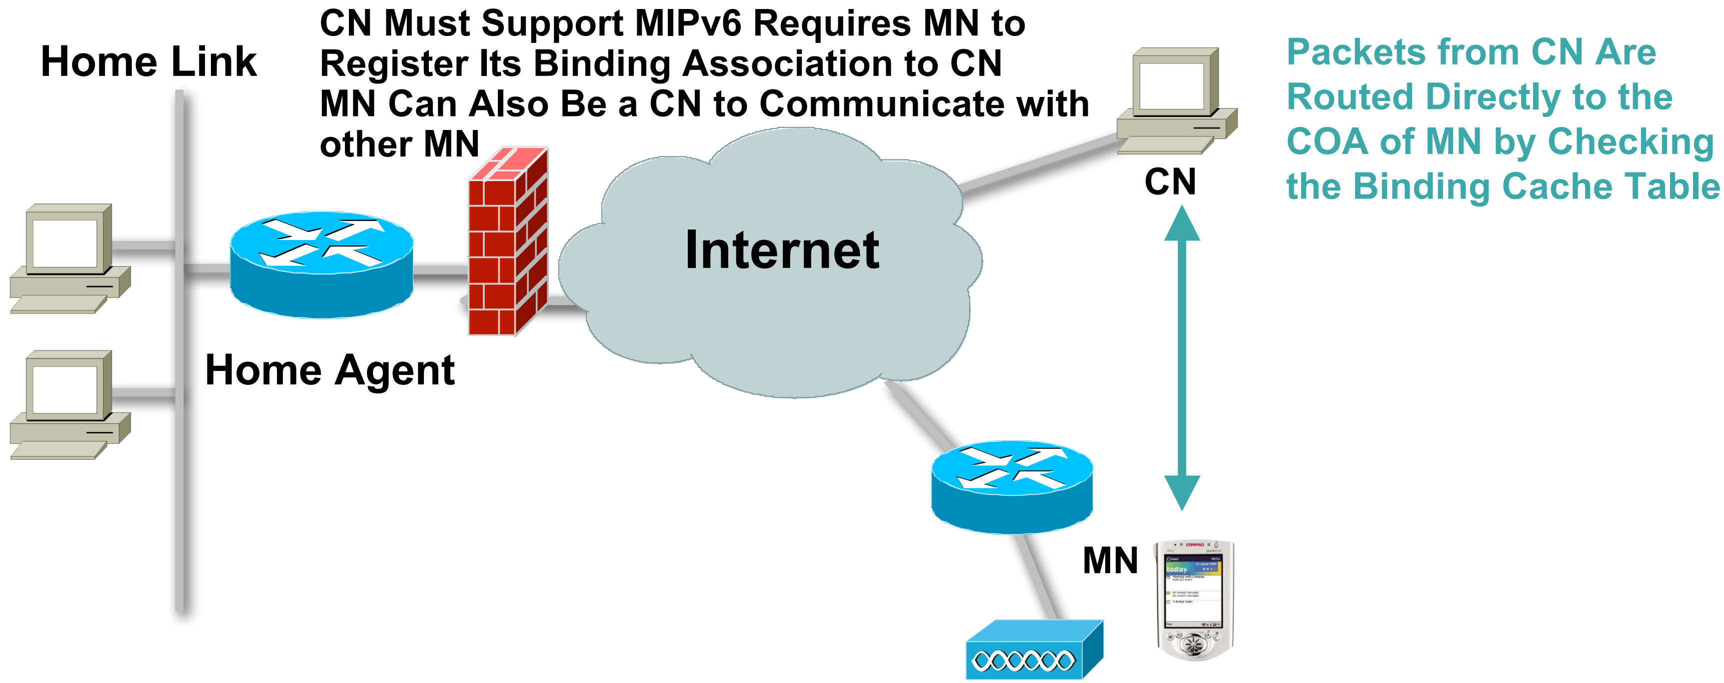

Traffic Is Going through HA until the Return Routability Procedure Is Performed

Signaling via HA, and Home Registrations Still Keep HA Informed

When Routing Packets Directly to CN, MN Sets the Source Address in the Packet's IPv6 Header to Its Current CoA

Cisco IOS Mobile IPv6 Home Agent Technology Preview

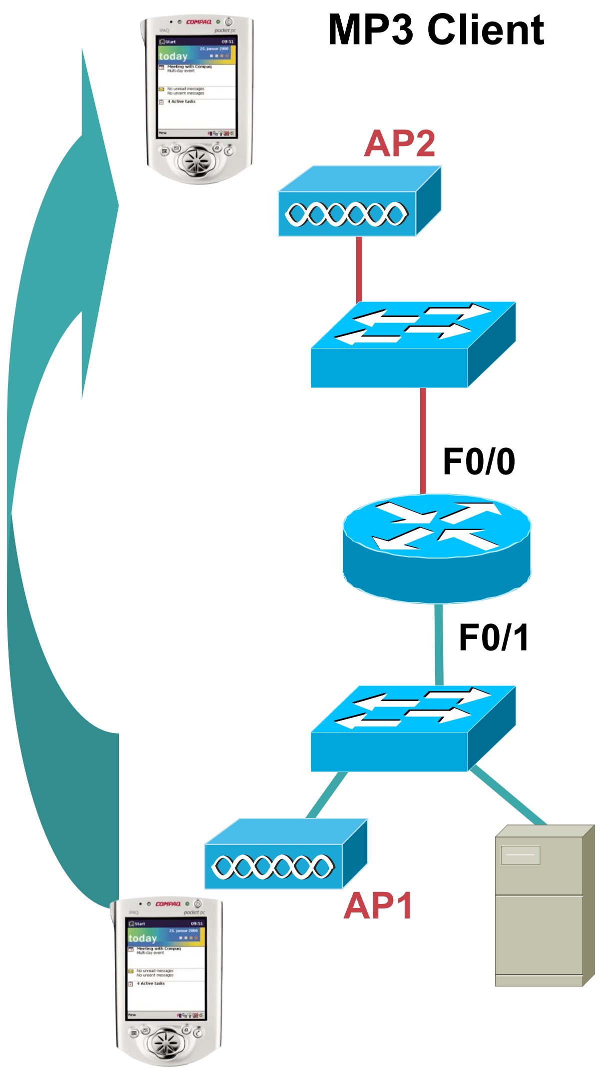

- **MIPv6 home agent technology preview release built on IETF MIPv6 draft 24**
  - IPSec support planned for a later stage
  - IETF MIPv6 standard completed: RFC 3775
  - IPSec for MN: RFC 3776
  - Binding update can be filtered by source address using ACL
- **Only available for testing and experiment**
  - Tested with BSD, Linux and Windows MIPv6 client

```
ipv6 unicast-routing
ipv6 cef
!
interface FastEthernet0/0
 ipv6 address 2001:DB8:C003:1101::45A/64
 ipv6 mobile home-agent run
!
interface FastEthernet0/1
 ipv6 address 2001:DB8:C003:1102::45C/64
 ipv6 mobile home-agent run
```

Mobile IPv6 at Cisco

- MIPv6 RFC 3775/3776 provides a number of security features
  - Protection of Binding Updates both to home agents, correspondent nodes
  - Protection of mobile prefix discovery through the use of IPSec extension headers
  - Protection of the mechanisms that MIPv6 uses for transporting data packets
- Home agent
  - In field trial since CY ’01
  - RFC3775 compliant
    - http://www.ietf.org/internet-drafts/draft-ietf-mobileip-ipv6-24.txt
    - Some issues from TAHI—Dynamic HA Address Discovery, Mobile Prefix Discovery—Worked-on/Fixed.
  - Available from Cisco IOS 12.3(14)T
  - Enhanced ACL—routing type filtering capability—planned
  - Light authentication planned
- Mobile IPv6 is part of the planned IPv6 rollouts
  - http://www.cisco.com/warp/public/732/Tech/ipv6/ipv6_learnabout.shtml
  - http://www.cisco.com/warp/public/732/Tech/ipv6/

Reference Materials

- “Deploying IPv6 Networks” by Patrick Grossetete, Eric Levy-Abegnoli, Ciprian Popoviciu—Cisco Press (ISBN: 1587052105)
- “Understanding IPv6” by Joseph Davies—Microsoft Press (ISBN: 0735612455)
- “IPv6 Essentials” by Silvia Hagen—O’Reilly & Associates (ISBN: 0596001258)
- www.cisco.com/go/ipv6—CCO IPv6 main page
- www.cisco.com/go/srnd—CISCO NETWORK DESIGN CENTRAL
- www.cisco.com/go/fn—select “Feature” and search for “IPv6”, then select “IPv6 for Cisco IOS Software”
- www.ietf.org
- www.hs247.com
- www.ipv6forum.com
- www.ipv6.org
- www.nav6tf.org/
- www.usipv6.com
- www.6net.org

Conclusion

- **Start now rather than later**
  - Purchase for the future and test, test and then test some more
  - Start moving legacy application towards IPv6 support
- **Things we did not talk about, but they are very important to consider**
  - ISP multihoming solutions (Multi6 WG)—"Goals for IPv6 Site-Multihoming Architectures" (RFC 3582)—http://www.ietf.org/html.charters/multi6-charter.html
  - Other transition methods such as EoMPLS, L2TPv3
- **Things to consider:**
  - Don't assume your favorite vendor/app/gear has an IPv6 plan
  - Full parity between IPv4 and IPv6 is still a ways off
- **Enterprise and SP deployment sScenarios**
  - http://www.ietf.org/internet-drafts/draft-ietf-v6ops-bb-deployment-scenarios-05.txt
  - 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)
- **What we hope to discuss next time**
  - Native IPv6 MPLS/recommendations for an IPv6-enabled SP core
- **What would you like to see/hear next time?**

Q AND A