GTRI
GLOBAL TECHNOLOGY RESOURCES INC.

CCNP Study Group
March 31, 2011

IPv6 Topics For ROUTE Exam

Scott Hogg

GTRI - Director of Advanced Technology Services

CCIE #5133, CISSP #4610

Internet Protocol Version 6

- IPv6 is the next generation computer network protocol for use on the Internet and within private networks.
- IPv6 is a standard defined by the Internet Engineering Task Force (IETF) and was first specified in the mid-90s.
- IPv6 is designed to replace IPv4 but IPv6 is a different protocol than IPv4 yet they can both coexist.
- IPv6 has taken many years to mature and get ready for mass deployment and now IPv6 is deployed on the global Internet.
- Organizations that connect to the Internet now need to learn about IPv6 and prepare their systems to communicate using this protocol.

IPv6 Header

IPv4 Header 20 bytes

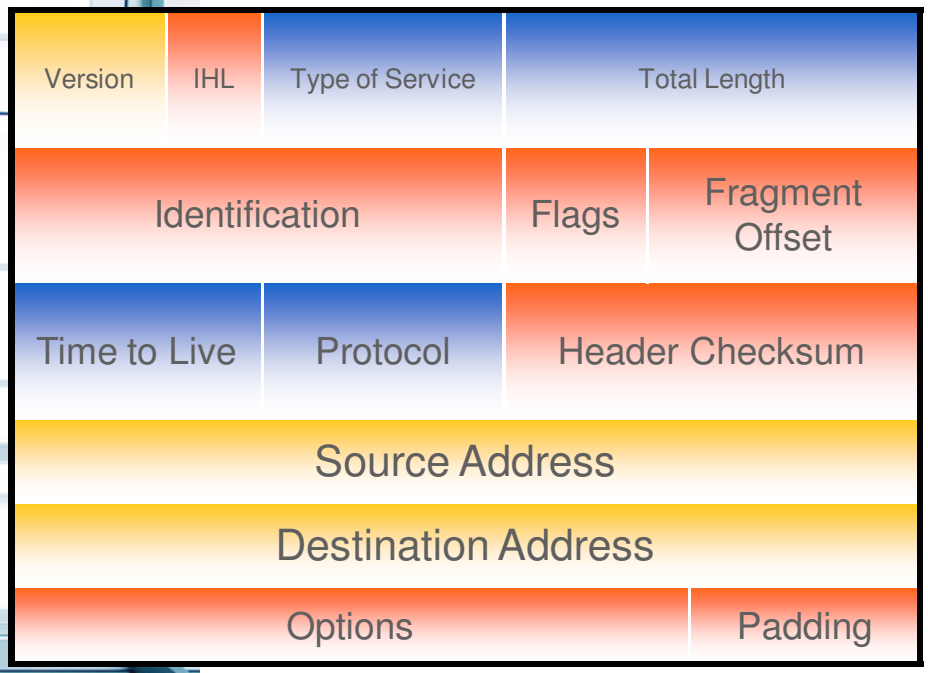

IPv6 Header, 40 bytes fixed

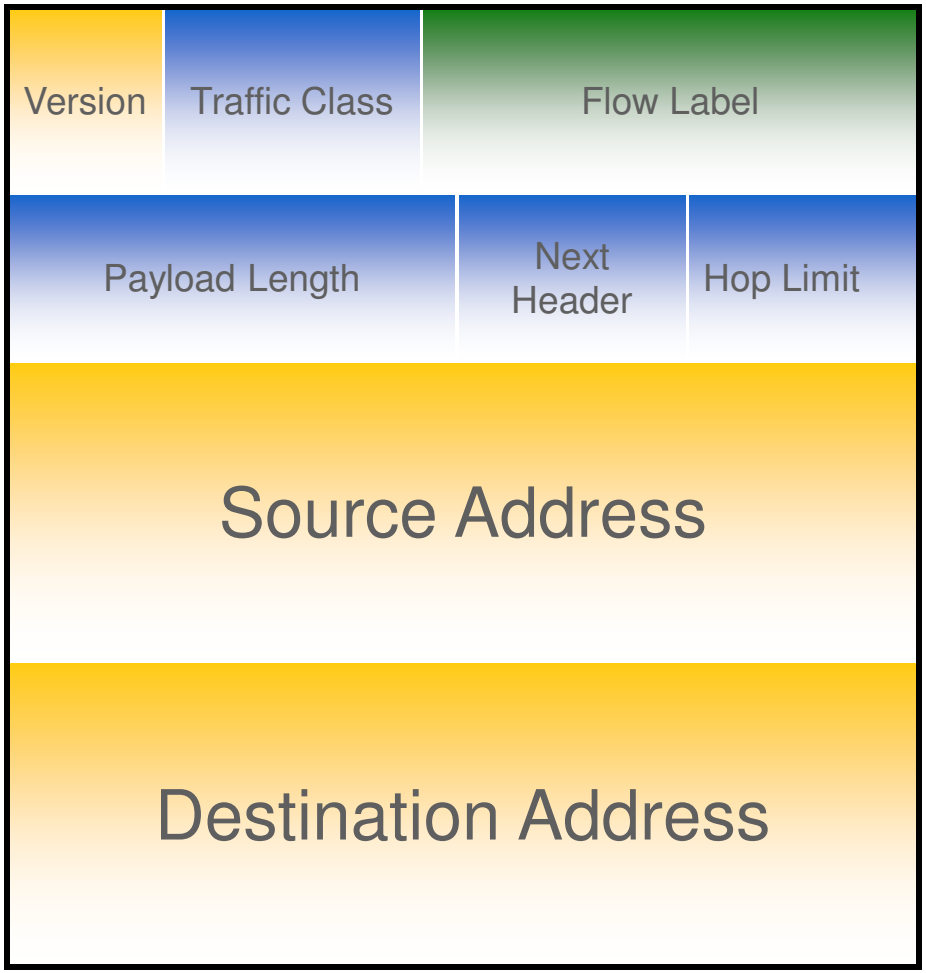

Legend

- **field's name kept from IPv4 to IPv6**
- **fields not kept in IPv6**
- **Name & position changed in IPv6**
- **New field in IPv6**

IPv6 Extension Headers

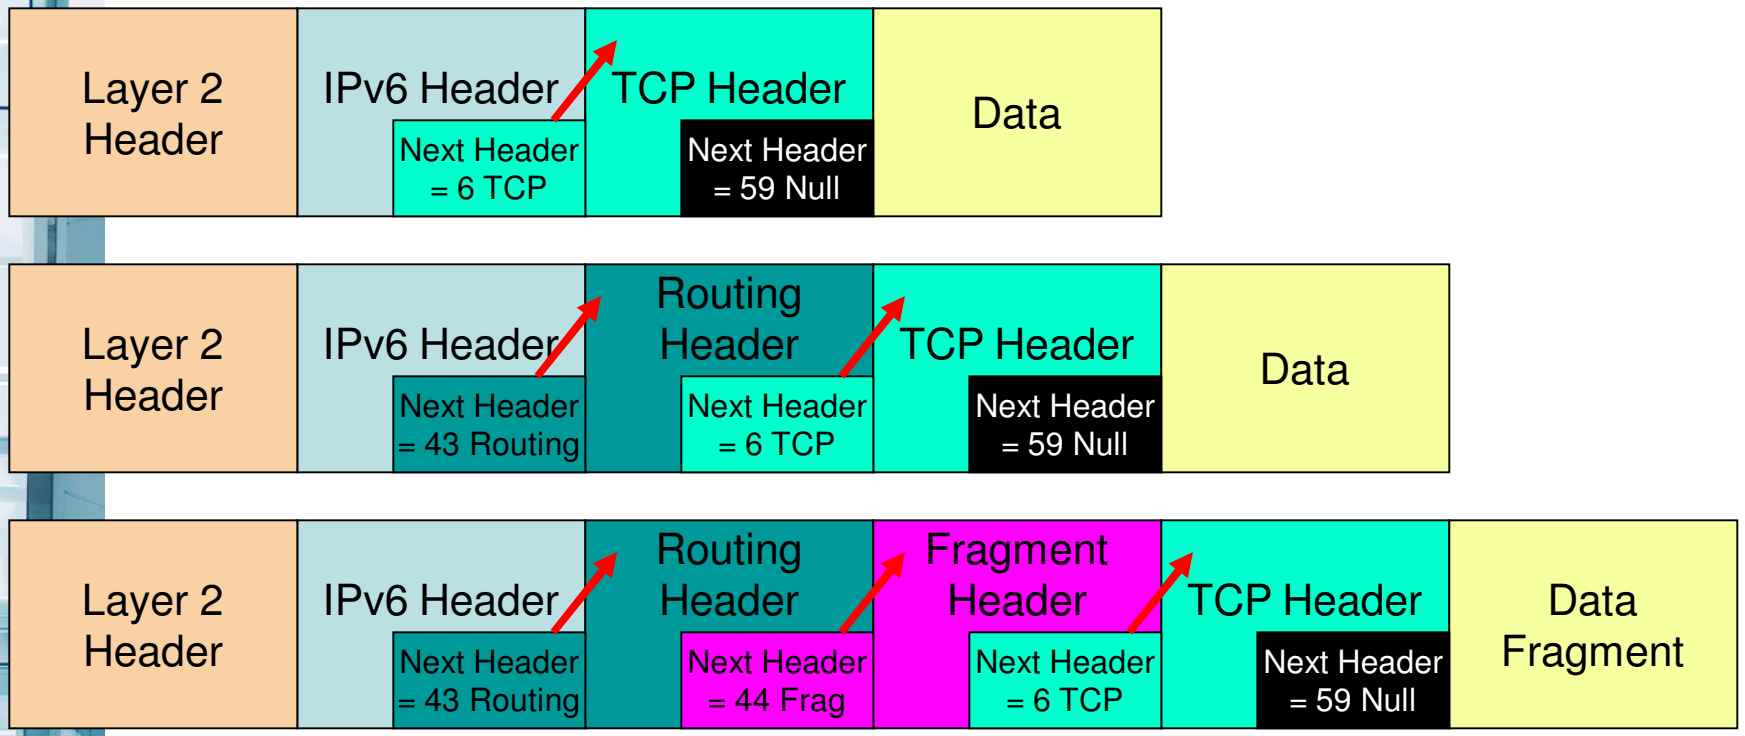

Addressing : Format Prefix

• Reserved (::0/128)	0000 0000
• Unassigned	0000 0001
• Reserved for NSAP Allocation	0000 001
• Reserved for IPX Allocation – later deprecated	0000 010
• Unassigned	0000 011
• Unassigned	0000 1
• Unassigned	0001
• Aggregatable Global Unicast Addresses (2000::/3)	001
• Provider-Based Unicast Address	010
• Unassigned	011
• Rsvd for Neutral-Interconnect-Based Unicast Addrs	100
• Unassigned	101
• Unassigned	110
• Unassigned	1110
• Unassigned	1111 0
• Unassigned	1111 10
• Unique Local (FC00::/7) (FC00::/8 & FD00::/8)	1111 110
• Unassigned	1111 1110 0
• Link Local Use Addresses (FE80::/10)	1111 1110 10
• Site Local Use Addresses (FEC0::/10) - deprecated	1111 1110 11
• Multicast Addresses (FF00::/8)	1111 1111

IPv6 Addressing Notation

- 128 bits get converted into more readable form
  - 0011 1111 1111 1110 1001 0000 1110 0000 0000 0000 0000 0011 0000 0000 0000 0000 / 0000 0000 0000 0000 0000 0000 0101 0000 0000 0000 0000 0000 0000 0000 0000 0000
- Convert bits to hex
  - 3FFE:90E0:0003:0000:0000:0050:0000:0000
- Reduce by removing leading zeros
  - 3FFE:90E0:3:0:0:50:0:0
- Use :: to consolidate multiple zeros – only once
  - 3FFE:90E0:3::50:0:0
  - or
  - 3FFE:90E0:3:0:0:50::
- Prefix format/notation
  - 3FFE:90E0:3::/64

IPv6 Address Types and Scope

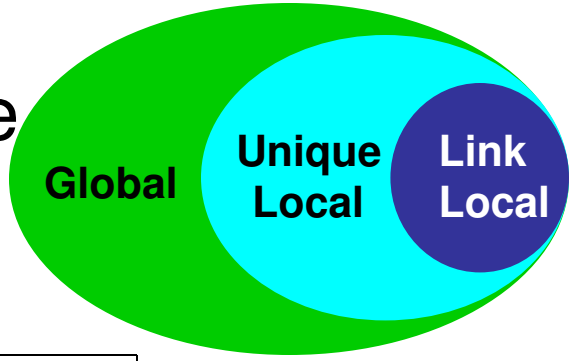

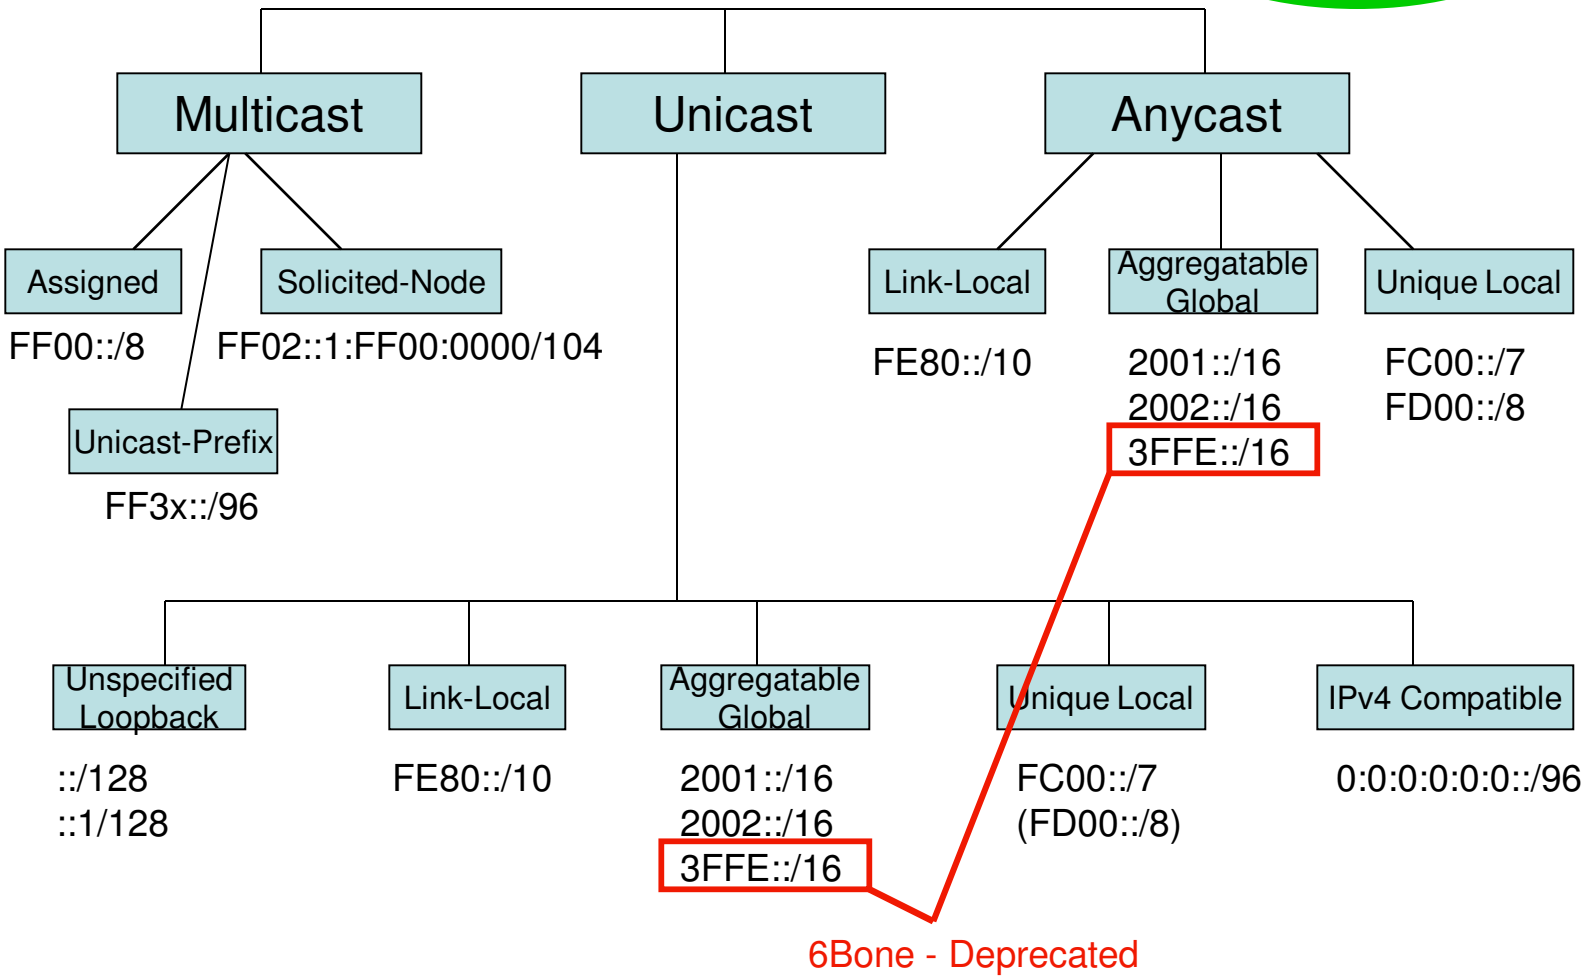

Original Address Allocation Policy

- The allocation process is under review by the Registries:
  - IANA allocates 2001::/16 to registries
  - Each registry gets a /23 prefix from IANA
  - Formerly, all ISP were getting a /35
  - 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)

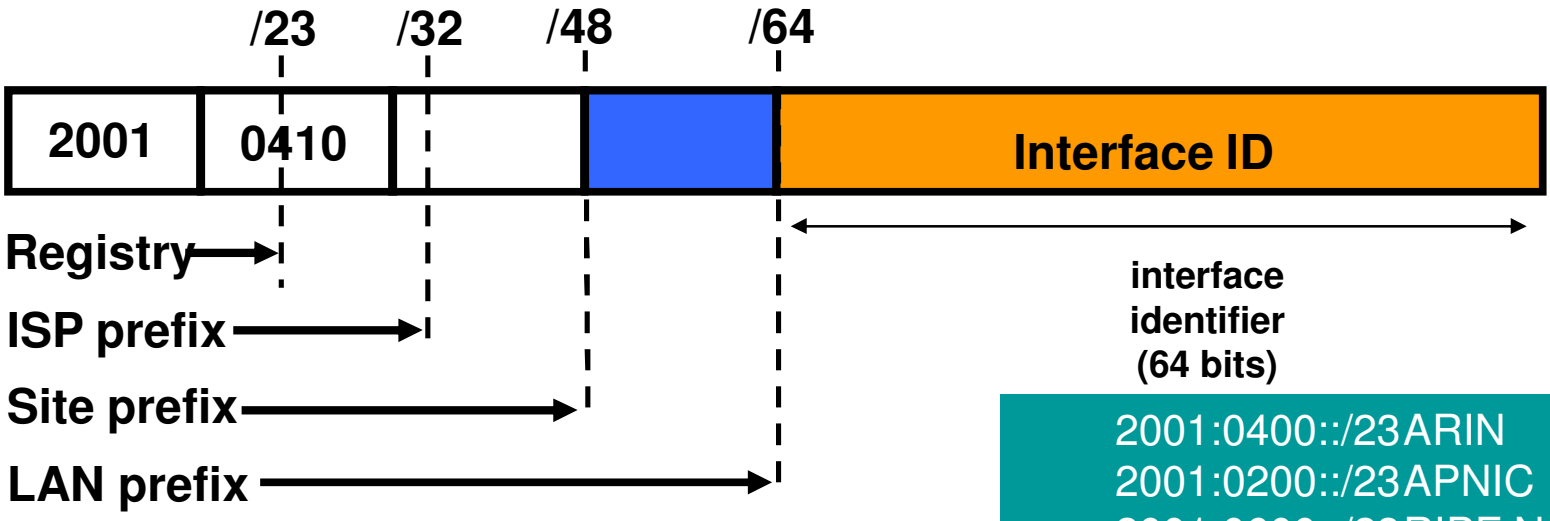

2001:0400::/23	ARIN
2001:0200::/23	APNIC
2001:0600::/23	RIPE NCC
2002::/16	6to4

IPv6 Addressing Plan

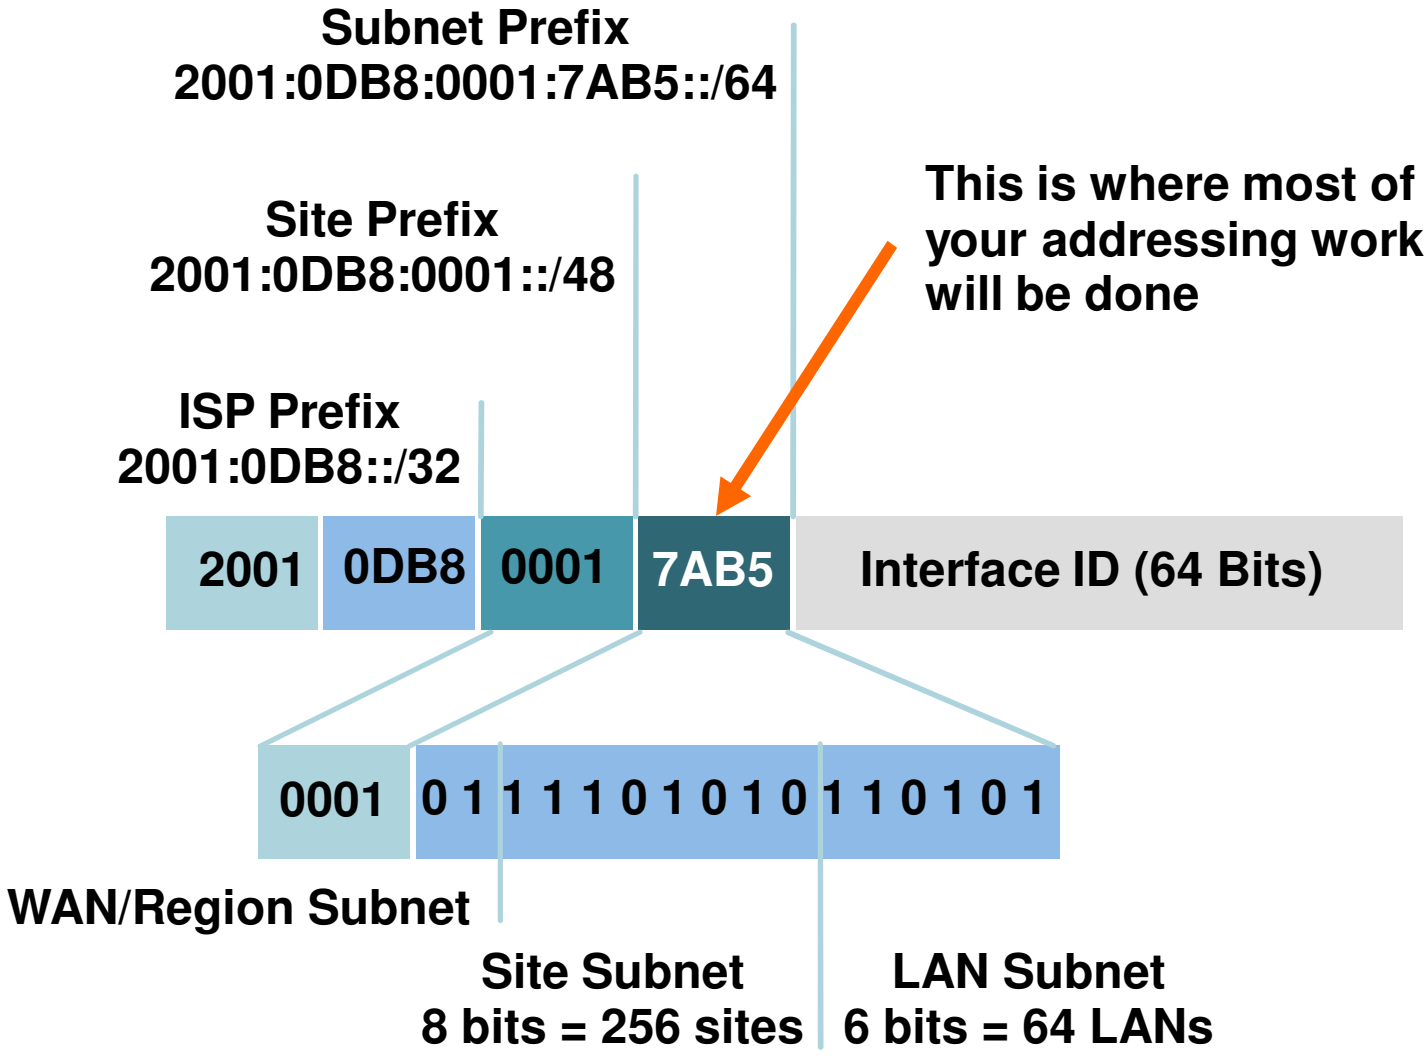

Additional IP Addressing Thoughts

- Large Enterprise Solution
  - /42 allocation
  - /48 for each campus/office
  - /56 for each building
  - /64 for each VLAN and P2P link
- IP Address Management Solutions are required
  - Excel spreadsheets are not a good plan
  - Get your staff to think/dream in hex

Multicast Addresses

- Flags Field: [0|R|P|T]
  - Bit 0 = reserved must be zero
  - Bit 4 = 0 if it is a well-known multicast address – Permanently assigned
  - Bit 4 = 1 if this is a non-permanently-assigned ("transient") – Temporarily assigned
- Scope Field:
  - 1 – Interface-Local – FF01
  - 2 – Link-Local – FF02
  - 4 – Admin-Local – FF04
  - 5 – Site-Local – FF05
  - 8 – Organization-Local – FF08
  - E – Global – FF0E

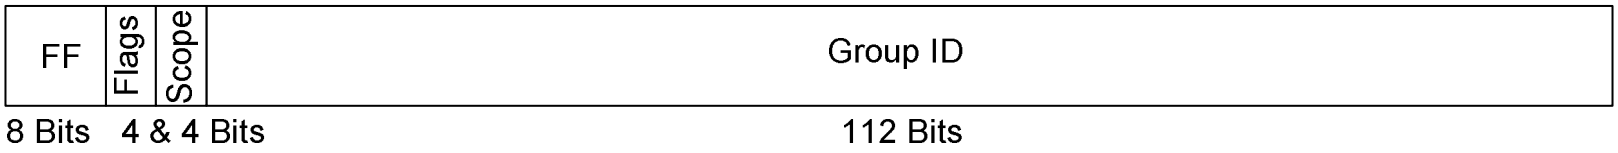

Multicast Addresses

- Permanently assigned multicast addresses
  - FF01:0:0:0:0:0:0:1 - All Nodes Address (FF01::1)
  - FF01:0:0:0:0:0:0:2 - All Routers Address (FF01::2)
  - FF02:0:0:0:0:0:0:1 - All Nodes Address (FF02::1)
  - FF02:0:0:0:0:0:0:2 - All Routers Address (FF02::2)
  - FF02:0:0:0:0:0:0:5 - OSPFIGP (FF02::5)
  - FF02:0:0:0:0:0:0:6 - OSPFIGP DR (FF02::6)
  - FF02:0:0:0:0:0:0:9 - RIPng Routers (FF02::9)
  - FF02:0:0:0:0:0:0:A - EIGRPv6 Routers (FF02::A)
  - FF02:0:0:0:0:0:0:B - Mobile-Agents (FF02::B)
  - FF02:0:0:0:0:0:0:D - All PIM Routers (FF02::D)
  - FF02:0:0:0:0:0:0:12 – VRRP (FF02::12)
  - FF02:0:0:0:0:0:0:16 - All MLDv2 capable Routers (FF02::16)
  - FF02:0:0:0:0:0:1:2 - All-dhcp-agents (FF02::1:2)
  - FF02:0:0:0:0:0:1:3 - Link Local multicast name resolution (FF02::1:3)
  - FF02:0:0:0:0:1:FF00:0000/104 - Solicited-Node Address for NDP
  - FF02:0:0:0:0:2:FF00::/104 - Node Information Queries (NIQ)
  - FF05:0:0:0:0:0:0:2 - All Routers Address (FF05::2)
  - FF05:0:0:0:0:0:1:3 - All DHCP servers (FF05::1:3)

Multicast on Ethernet Networks

- IPv4 multicast packets are sent to the Ethernet multicast address 01-00-5E-XX-XX-XX using Ethernet II type 0x0800
- IPv6 multicast packets are sent to the Ethernet multicast address 33-33-XX-XX-XX-XX using Ethernet II type 0x86DD
- The last 4 octets of the multicast address come from the last 4 octets of the IPv6 multicast group address
- For example:
  - FF02:0:0:0:0:0:0:1 - All Nodes Address (FF02::1)
  - 33-33-00-00-00-01 – Ethernet destination

Anycast Addresses

- Anycast is the method of sending a packet to the topologically “nearest” node from a group of potential receivers all using the same address
- Anycast is “one-to-one-of-many”
- Anycast addresses violate the “laws of uniqueness” of IP addresses
- Only routers and special servers should respond to an anycast address
- IPv6 anycast addresses look the same as an IPv6 unicast address
- Anycast is used for DNS root name servers
- Subnet Router Anycast Addresses (RFC 4291)
- Reserved Subnet Anycast Addresses (RFC 2526)

ICMPv6

- More powerful than ICMP for IPv4
- ICMPv6 uses IPv6 extension header #58 (RFC 2463)

Type	Description
1	Destination Unreachable
2	Packet to Big
3	Time exceeded
4	Parameter problem
128	Echo Request
129	Echo Reply
130	Multicast Listener Query – sent to ff02::1 (all nodes)
131	Multicast Listener Report
132	Multicast Listener Done – sent to ff02::2 (all routers)
133	Router Solicitation (RS) – sent to ff01::2 (all routers)
134	Router Advertisement (RA) – sent to ff01::1 (all nodes)
135	Neighbor Solicitation (NS) – sent to ff02:0:0:0:0:1:ff00::/104
136	Neighbor Advertisement (NA)
137	Redirect message

PING

MLD

Prefix Advertisement

Router Redirection

ARP Replacement

ICMPv6 Error Messages

ICMPv6 Error Messages:	Type 0 to 127
Destination Unreachable	Type 1
to Destination	Type 1, Code 0
Administratively Prohibited	Type 1, Code 1
Address Unreachable	Type 1, Code 3
Port Unreachable	Type 1, Code 4
Packet Too Big - Path MTU Discovery (PMTUD)	Type 2, Code 0
Time Exceeded	Type 3
Hop Limit Exceeded	Type 3, Code 0
Fragment reassembly time exceeded	Type 3, Code 1
Parameter Problem	Type 4
Erroneous Header Field	Type 4, Code 0
Unrecognized Next Header type	Type 4, Code 1
Unrecognized IPv6 option	Type 4, Code 2

Router Solicitations and Advertisements

Nodes send RSs (Type 133)
On bootup when they can't wait 200 seconds for the next RA

Source: FE80::/10
Link-Local address of Node
Destination: FF02::2 (all routers)

Data: Query to send RA

Routers send RAs (Type 134)
Every 200 seconds or
Responding to an RS message

Source: FE80::/10
Link-Local address of Router
Destination: FF02::1 (all nodes)

Data: Options, subnet prefix, lifetime, autoconfig flags (M&O bits)

RA Message Data

- RA messages contain valuable information for nodes to pull themselves up by their bootstraps and get on the network
  - Router Lifetime
  - Reachable Time
  - Retransmision Timer
  - Source Link-Layer Address
  - MTU size for the link
  - Prefix Information
  - Address Autoconfiguration Flag – A flag
  - Managed Address Configuration Flag - M flag
  - Other Stateful Configuration Flag - O flag

RA Message Data

- Address Autoconfiguration Flag – A flag
  - A = 0 = Node should not perform stateless address assignment
  - A = 1 = Node should perform stateless address assignment
- Managed Address Configuration Flag - M flag
  - When set to 1 it tells hosts to get stateful address through DHCPv6
- Other Stateful Configuration Flag - O flag
  - When set to 1 it tells host to use DHCPv6 to get other config info
  - M=0, O=0
    - End host uses stateless auto-configuration
  - M=0, O=1
    - DHCPv6 stateless, autoconfigured address but DNS and other information is added with DHCPv6 (i.e. from a router)
  - M=1, O=0
    - Not practical, no DNS info
  - M=1, O=1
    - End host uses DHCPv6 for all information

IPv6 Address Lifetimes

- Tentative : the address is in the process of being verified as unique
- Preferred : a node can send and receive unicast traffic to and from a preferred address
- Deprecated : the address is still valid, but using it for new communication is discouraged
- Invalid : the address can no longer send unicast traffic to or receive it from a node. An address enters this state after the valid lifetime expires

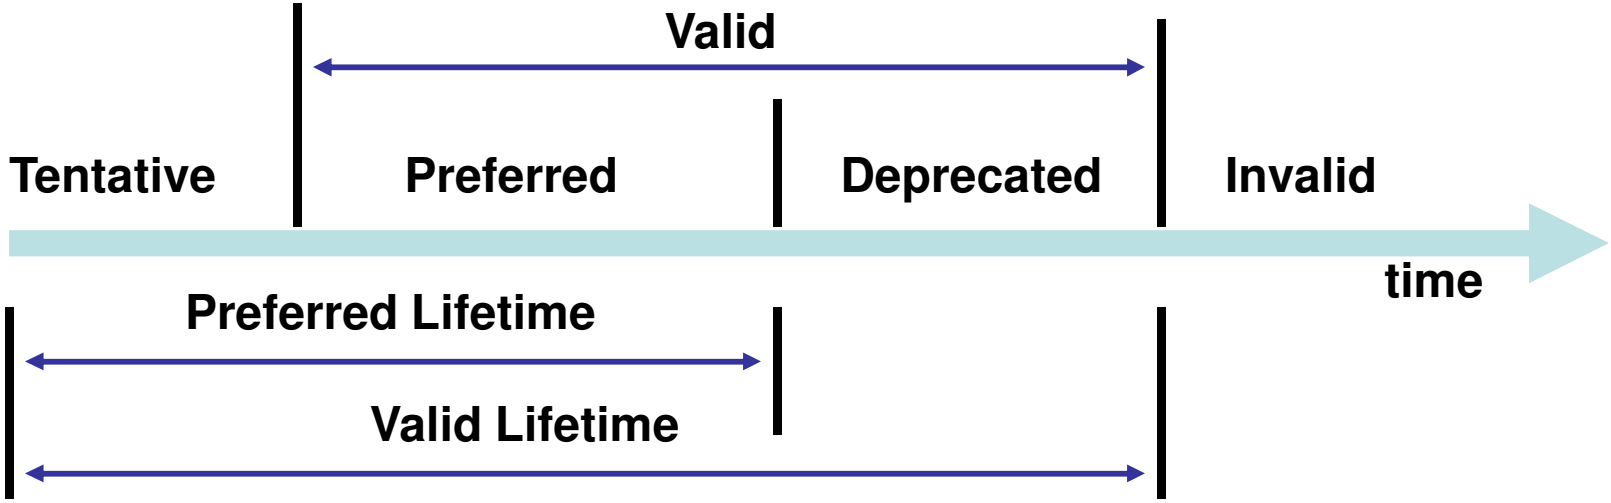

Neighbor Solicitations and Advertisements

Nodes send NSs (Type 135)
When sending IPv6 packet to Another node

Source: Unicast IPv6 Address
Destination: Solicited Node Multicast Address FF02::1:FFAA:BBCC

Data: Target link-layer address
Query: What is your link-layer address?

Routers send NAs (Type 136)
Responding to an NS message

Source: Unicast IPv6 Address
Destination: Unicast Address of Requestor or FF02::1 (all nodes)

Data: R/S/O Flags, Target's Link-layer address
Response: Here is my IPv6 and link-layer address.

Neighbor Cache State Machine

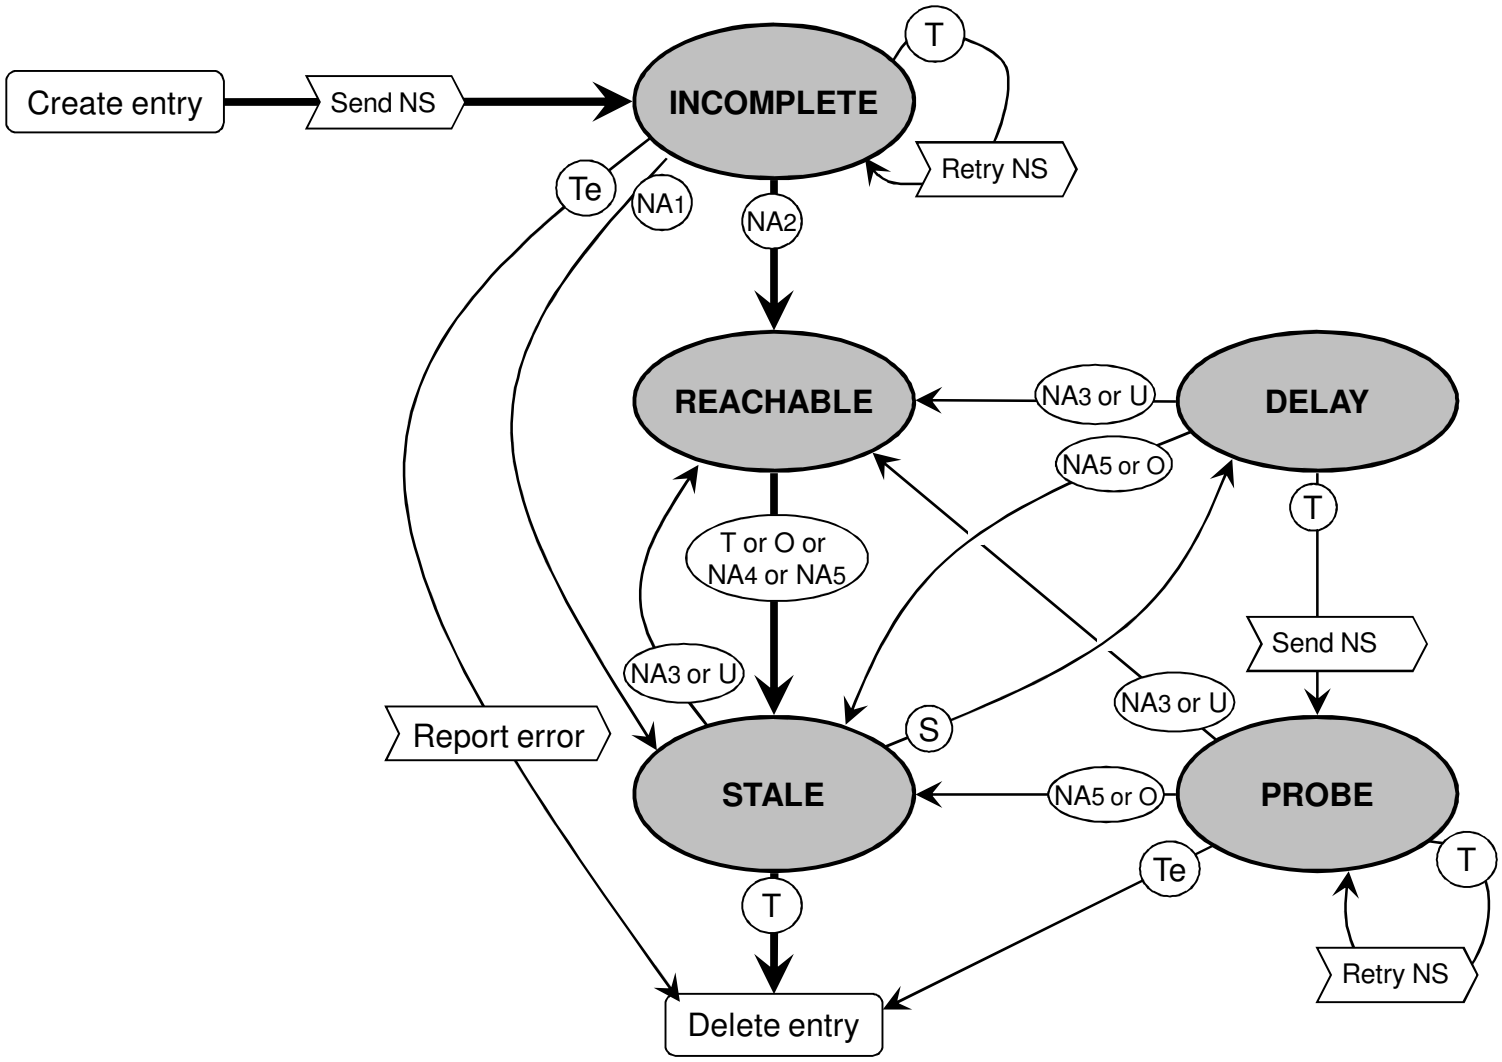

DHCPv6 (RFC 3315)

- Similar to DHCPv4, but hosts check RAs message from the local router first (M & O bits).
- Can be used for renumbering using Reconfigure-Init message.
- DHCPv6 provides DNS information, 16-bit option space
- DHCPv6 Messages: Solicit, Advertise, Request, Confirm, Renew, Rebind, Decline, Info-Request, Reply, Reconfigure
  - DHCPv6 clients listen on UDP port 546
  - DHCPv6 servers and relay agents listen on UDP port 547
  - All_DHCP_Relay_Agents_and_Servers (FF02::1:2) – link-local scope
  - All_DHCP_Servers (FF05::1:3) – site-local scope

DHCPv6 Messages

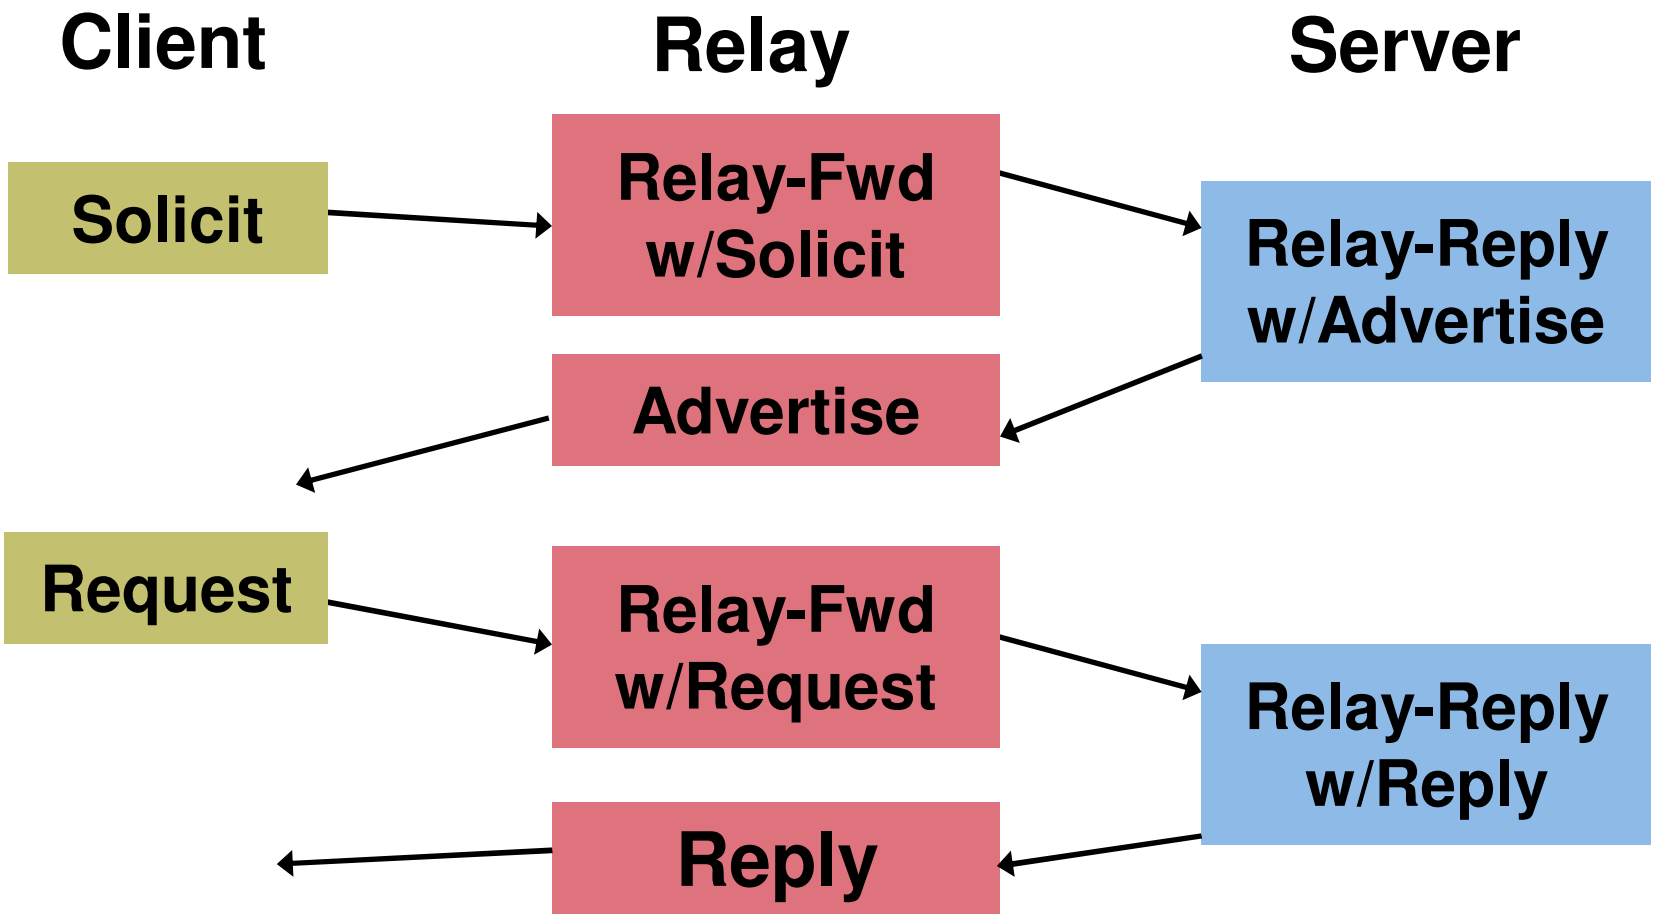

Stateless DHCPv6

- Stateless DHCPv6 Service for IPv6 (RFC 3736), Sometimes called DHCPv6lite
- Method to provide DNS server to client that isn't running stateful DHCPv6 client
- Like DHCPv4 DHCPINFORM/DHCPACK
- Router can perform as DHCPv6lite server

```
ipv6 dhcp pool V6POOL
 dns-server 2001:DB8:11::6
 domain-name example.com
interface Vlan11
 description User LAN - Dual-Protocol
 ipv6 address 2001:DB8:11::1/64
 ipv6 dhcp server V6POOL
```

Stateless Address Autoconfiguration

- If the Managed Address Configuration Flag – "M flag" and the Other Stateful Configuration Flag – "O flag" are both set to zero then this tells the end node to use stateless auto-configuration
- The host then uses the prefix information from the RA for the network portion of the node's IPv6 address
  - Statelessly configures its own Interface Identifier (Interface ID)
  - Using either EUI-64, privacy, or temporary addressing

IPv6 Interface Identifier – EUI-64

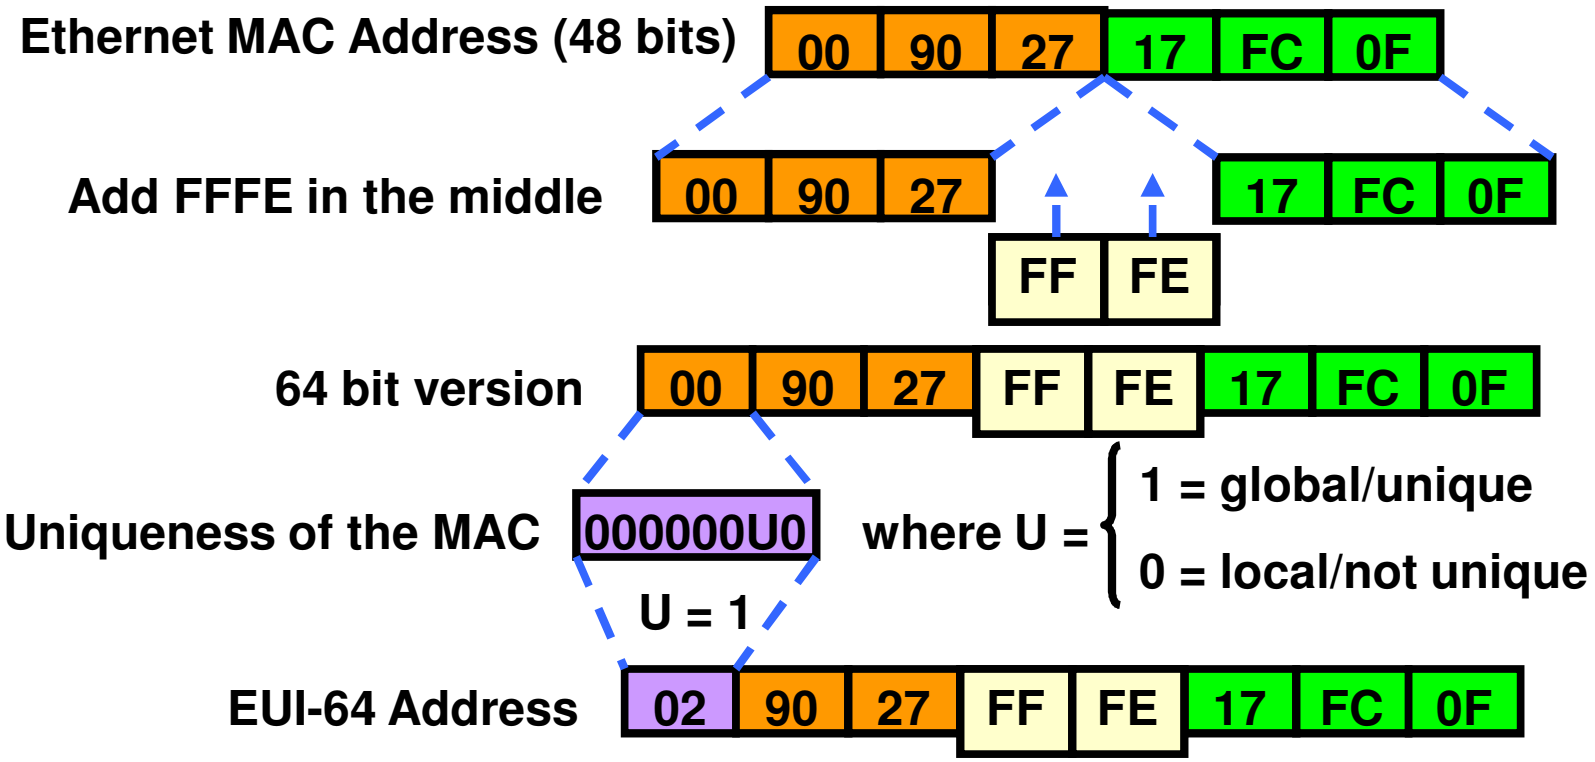

EUI-64 address is formed by inserting "FFFE" and "OR"ing a bit identifying the uniqueness of the MAC address.

Other IPv6 Features

- IPv6 requires every network link be capable of minimum MTU of 1280 bytes
- IPv6 routers don't fragment packets
  - Hosts perform their own Path MTU Discovery
  - MTU path discovery uses ICMP "packet too big" error messages
- Provider selection (based on policy, performance, cost, …)
- Host mobility (route to current location)
- Auto-readdressing (route to new address)
  - (Use IPv6's routing extension header)

IPv6 Transition Techniques

- Dual Stack
- Tunnel/Encapsulation
  - Configured Tunnels
  - Automatic Tunnels (Dynamic Tunneling)
    - 6to4
    - ISATAP
    - Teredo
    - Tunnel Broker with TSP
- Translation, Application Layer Gateways
  - Proxy
  - NAT-PT (deprecated), NAT64/DNS64, NAT66

Dual IP Stacks Model

- Dual-Stack Architecture – RFC 1933
- 4 different possibilities
- Choice of the IP version is based on name lookup, application or operating system preference
- IPv4 and IPv6 packets flow in Ethernet like “ships in the night”

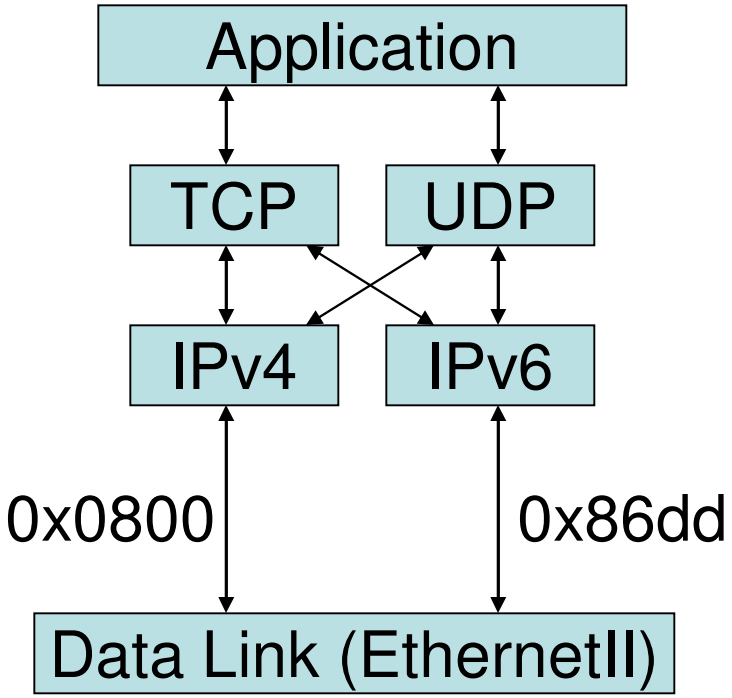

Sample Dual-Stack Router

```
ipv6 unicast-routing
ipv6 multicast-routing
ipv6 cef
!
interface Loopback0
 ip address 200.100.1.3 255.255.255.255
 ipv6 address 2001:db8:10:10::10/128
!
interface Ethernet 0
 ip address 192.168.100.1 255.255.255.0
 ipv6 address 2001:db8:1:1::1/64
!
ipv6 route ::/0 2001:db8:1:1::100
```

IPv6 Tunneling

- Manually configured or Automatic
- IPv6 PDUs encapsulated in IPv4 protocol 41

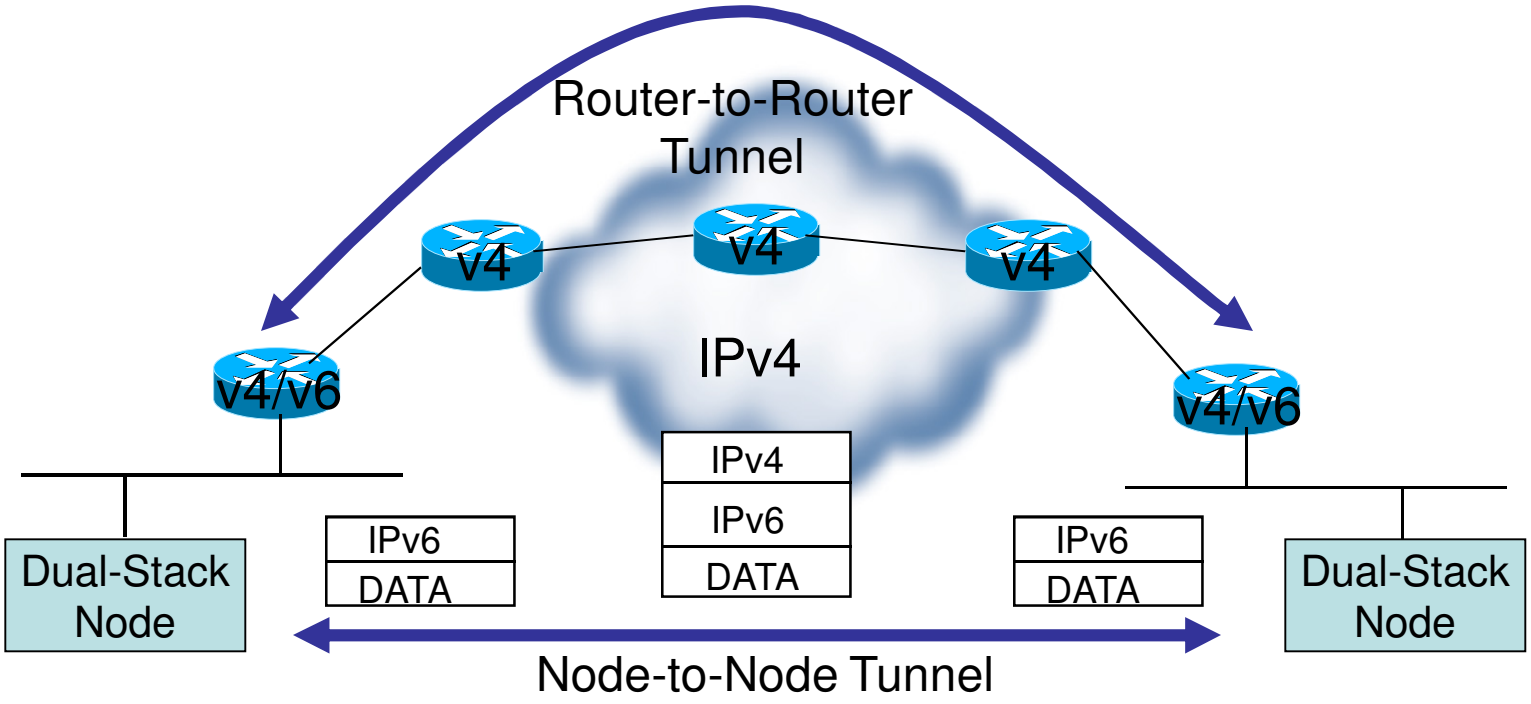

Tunnel Configuration

```
hostname Router1
interface Tunnel 0
 ipv6 address 2001:db8:100:1::1/126
 tunnel source 192.168.100.1
 tunnel destination 192.168.200.2
 tunnel mode ipv6ip
ipv6 route 2001:db8:c:1::/64 tunnel0

hostname Router2
interface Tunnel 0
 ipv6 address 2001:db8:100:1::2/126
 tunnel source 192.168.200.2
 tunnel destination 192.168.100.1
 tunnel mode ipv6ip
ipv6 route 2001:db8:d:2::/64 tunnel0
```

IPv6 Tunneling – 6to4

- Connection of Isolated IPv6 Domains via IPv4 Clouds Without Explicit Tunnels
- Inter-domain tunneling using IPv4 address as IPv6 site prefix IPv6 using IPv4 as a virtual link-layer
  - IPv6 VPN over IPv4 Internet (2002::/16 prefix)
  - Automatic tunneling approach - Minimal manual configuration
  - Uses globally unique prefix comprised of the unique 6to4 TLA and the globally unique IPv4 address of the exit router.
- 6to4 Relay is the gateway between the IPv6 and IPv4 worlds
  - No NAT can exist in the path
  - 6to4 Relay may be far away from end node
  - Security issues related to an open relay

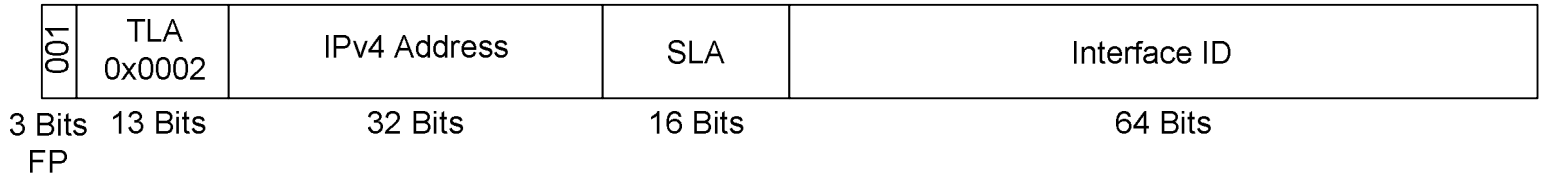

Automatic 6to4 Tunneling

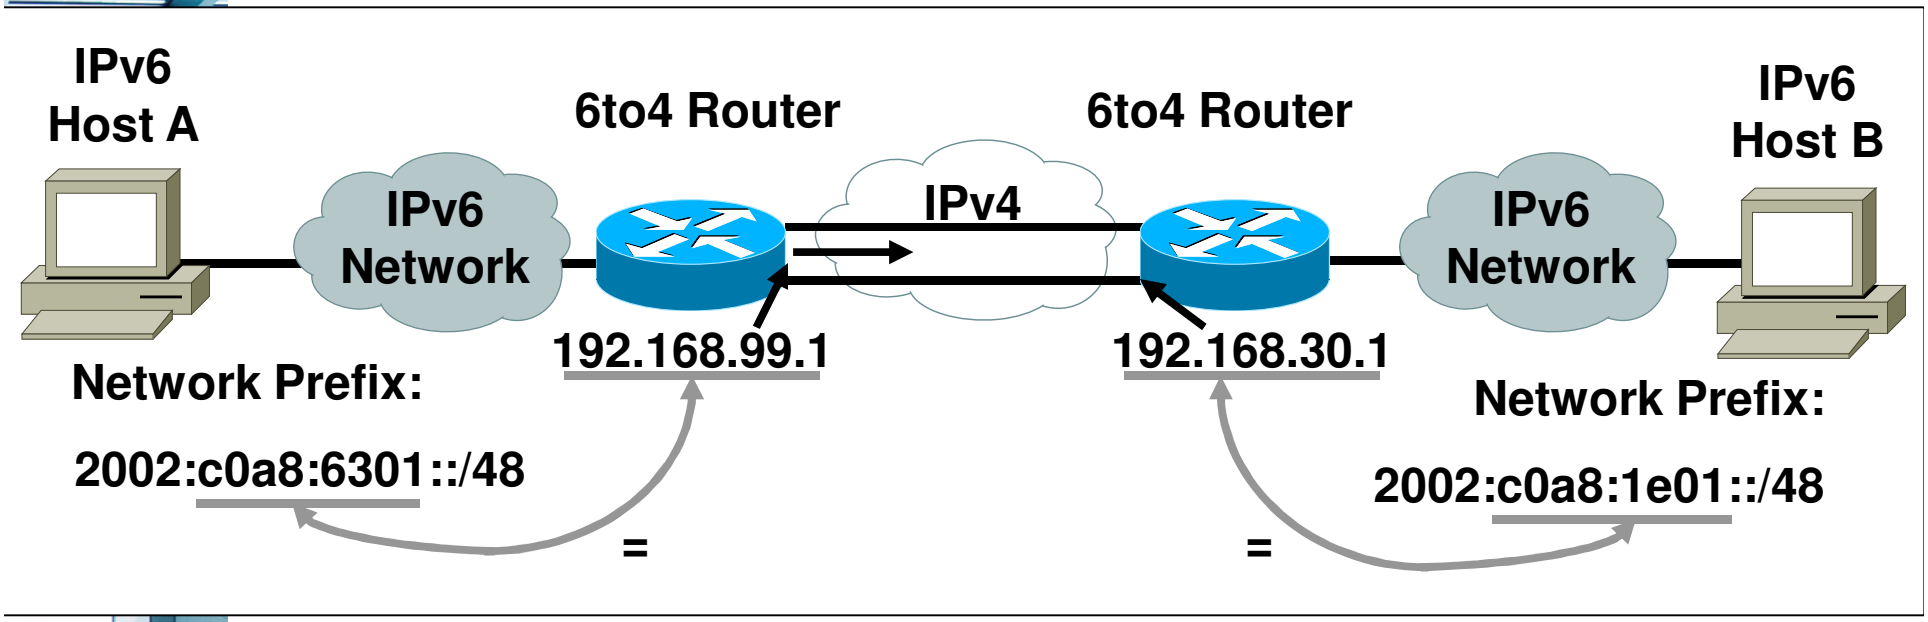

- 6to4:
  - Is an automatic tunnel method
  - Gives a prefix to the attached IPv6 network

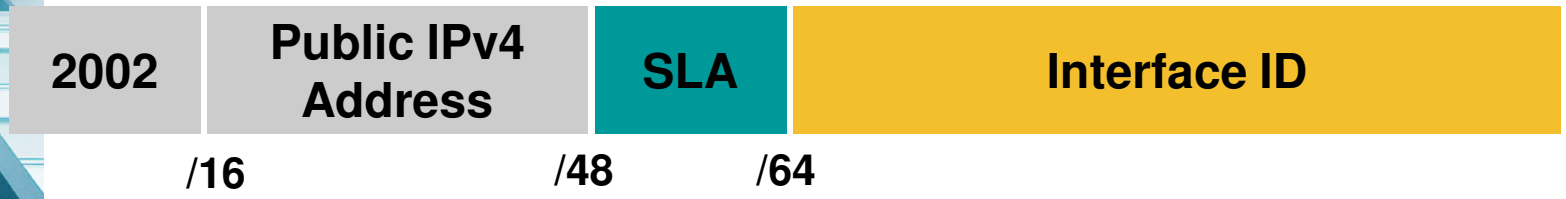

6-to-4 Configuration

```
hostname BorderRouter
interface Ethernet0
 ip address 200.168.100.1 255.255.255.0
interface Tunnel0
 no ip address
 ipv6 address 2002:c8a8:6401:1::1/128
 tunnel source Ethernet0
 tunnel mode ipv6ip 6to4
ipv6 route 2002::/16 Tunnel0
ipv6 route ::/0 2002:c8a8:c802:2::2
```

```
hostname 6to4RelayRouter
interface Ethernet0
 ip address 200.168.200.2 255.255.255.0
interface Tunnel0
 no ip address
 ipv6 address 2002:c8a8:c802:2::2/128
 tunnel source Ethernet0
 tunnel mode ipv6ip 6to4
ipv6 route 2002::/16 Tunnel0
```

IPv6 Tunneling – ISATAP

- Intra-Site Automatic Tunnel Addressing Protocol (RFC 4214)
- Automatic tunneling inside an enterprise
- ISATAP connections look like one flat network
- Creates a virtual IPv6 link over an IPv4 network
- Uses 5EFE just before the 32 bit IPv4 address bits – converted to hex
- Can use private address space
- Create a DNS “A” record for “isatap.domain.com” equal to router’s lo0
- Or “C:\>netsh interface ipv6 isatap set router <ip4addr>”
- Currently, ISATAP doesn’t support multicast

Prefix	00 00 5E FE	IPv4 Address
64 Bits	32 Bits	32 Bits

IPv6 Tunneling – ISATAP

```
interface Loopback0
 ip address 192.168.12.1 255.255.255.0
interface Tunnel0
 ipv6 address 2001:db8:c:1::/64 eui-64
 tunnel source loopback 0
 tunnel mode ipv6ip isatap
 no ipv6 nd suppress-ra
```

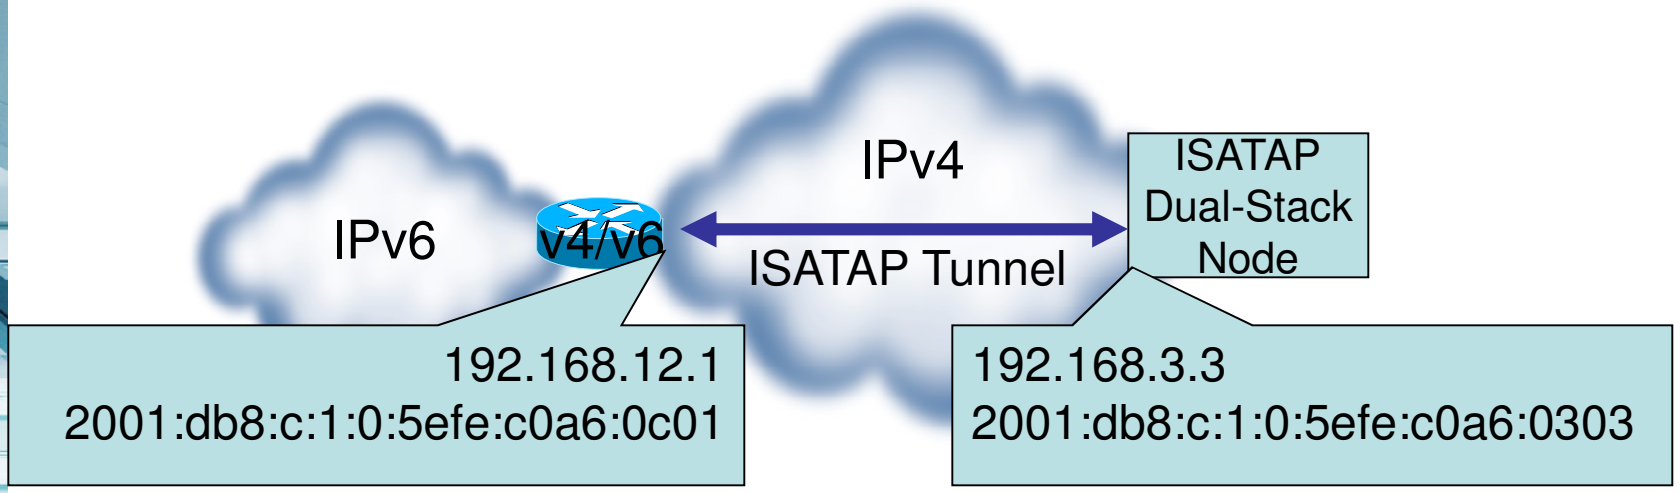

IPv6 Tunneling – Teredo

- Called Shipworm in earlier IETF drafts
- IPv4/UDP encapsulated IPv6 packets
- Works behind an IPv4 NAT
- Reduces MTU because of UDP encapsulation (port 3544)
- Uses Teredo server, Teredo relay, and a Teredo client
- External mapping of IPv4 address and port are discovered by the Teredo server (on the external side of the NAT)

IPv6 Tunneling – Tunnel Broker

- Tunnel Brokers use a web-based service to create a tunnel
- Connects an isolated host to IPv6 net of provider operating the tunnel broker
- Tunnel information is sent via http-ipv4
  - Tunnel managed by ISP
  - Sends scripts/configs to Dual Stack Router

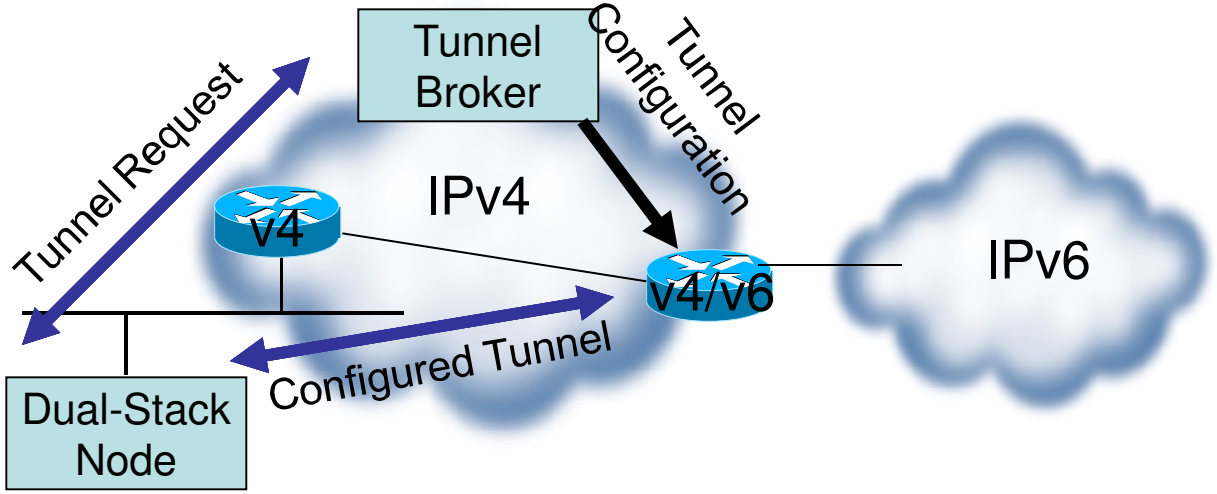

Dual Stack Where You Can, Tunnel Where You Must

Dual Stack Networks Run Both Protocols

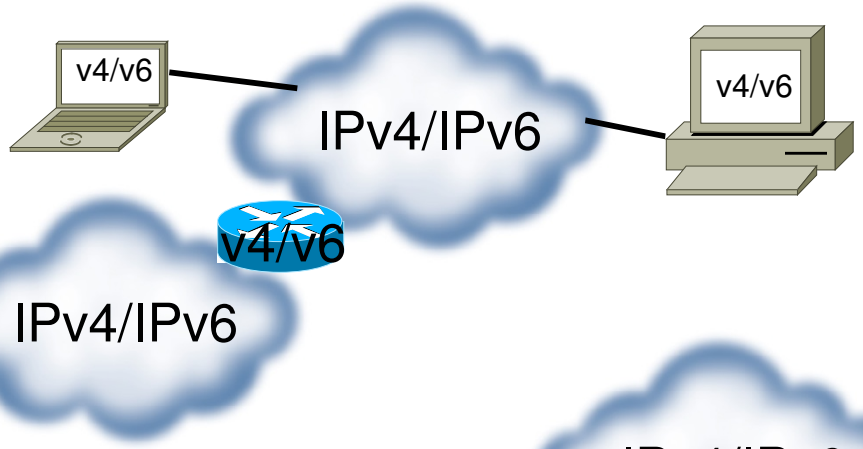

Tunneling IPv6 Across IPv4 Requires Careful Management

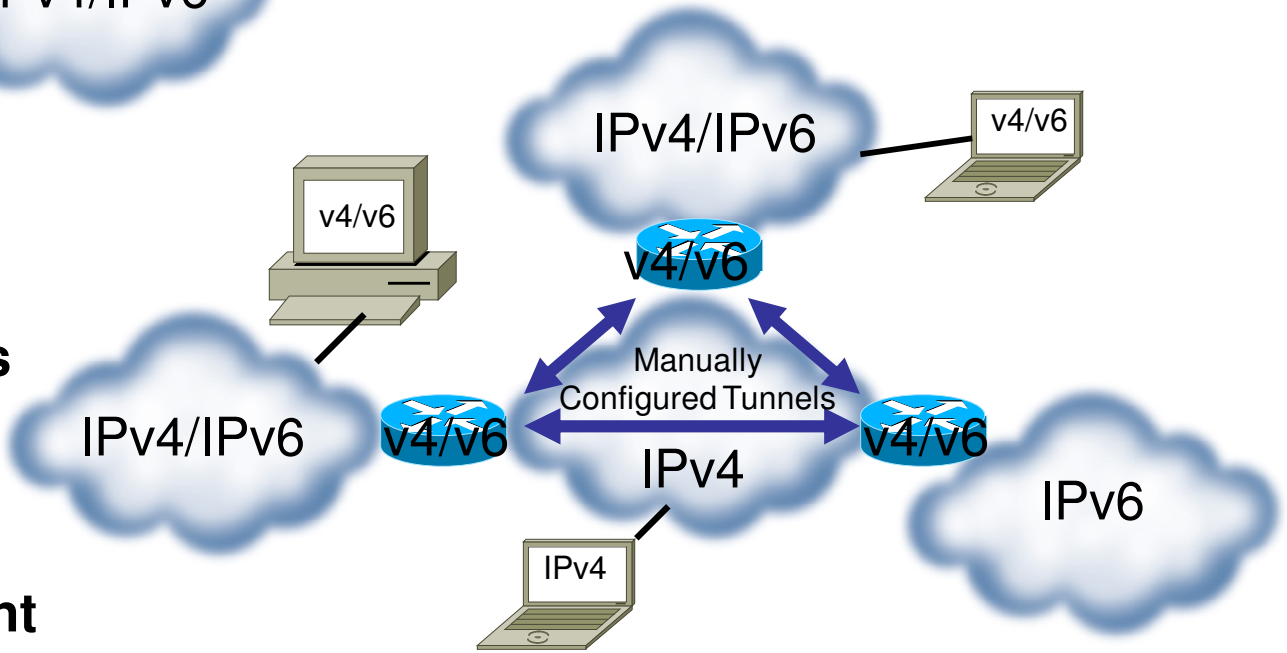

Translation/Proxy

- Application Layer Gateways (ALGs)
  - Application Layer Proxy
  - Terminate an IPv4 connection and initiate an IPv6 connection
  - Terminate an IPv6 connection and initiate an IPv4 connection
- Translation
  - NAT-PT (Now deprecated)
    - RFC 2766, RFC 4966
  - NAT64/DNS64
  - NAT66

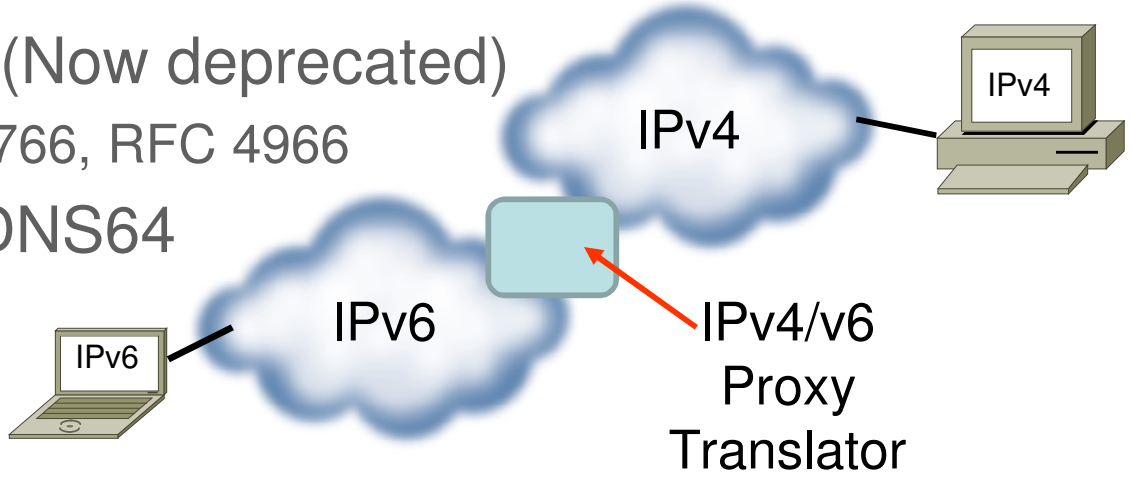

IPv4/IPv6 Proxy

- Can be performed by a Server Load Balancing (SLB) appliance
  - Citrix NetScaler 9.2
  - F5 BIG-IP Local Traffic Manager (LTM)
  - A10 AX Series Application Delivery Controller (ADC)
  - Brocade ServerIron ADX
  - Nginx
  - Apache 2.X mod_proxy.so

NAT64/DNS64

- Translation that uses NAT and DNS in combination
- Uses the IPv6 prefix 64:ff9b::/96 to map the 32-bit IPv4 DNS responses into an IPv6 address
- DNS64 and NAT64 share information about state of connections

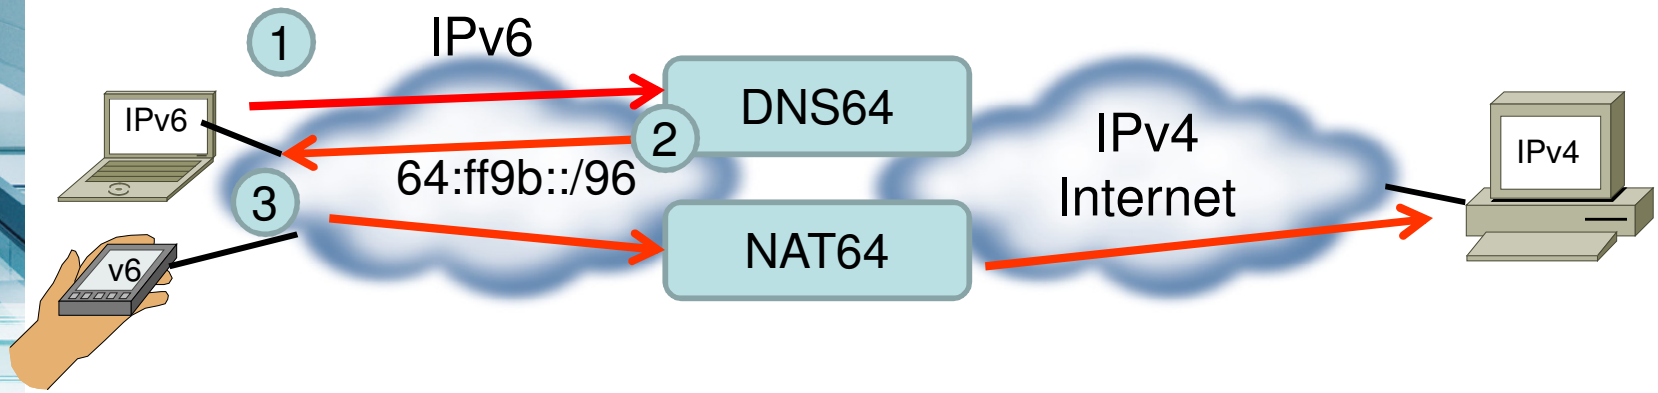

NAT64/DNS64

- Systems that can perform NAT64 and DNS64
  - Internet Systems Consortium (ISC) Address Family Transition Router (AFTR)
  - Viagenie Ecdysis LiveCD
- Systems that can perform Stateless NAT64
  - Cisco Routers (IOS-XE and CGSE with IOS-XR)
  - Juniper Routers
  - Juniper SSG 6.2.0R7
  - Brocade ServerIron ADX
  - A10 Networks AX Series Application Delivery Controller (ADC)
  - F5 Networks BIG-IP LTM
  - Citrix NetScaler 9.2

IPv6 Routing Protocols

- The key to scalable routing is to use hierarchical addressing
- IPv6-Capable Routing Protocols
  - RIPng
  - OSPFv3
  - Integrated IS-IS for IPv6
  - EIGRPv6 (12.4(6)T)
  - Multiprotocol-BGP
- IPv6 still uses longest-prefix matching
- IPv4 converges before IPv6
- Strive for IPv4/v6 Congruent Topologies
- Added CPU/Memory for running IPv4 and IPv6
  - Tuning can be done to speed convergence of both routing tables

RIPng

- RIPng (RFC 2080) is based on RIPv2 (RFC 2453)
- Distance-vector (routing by rumor)
- Classless
- Radius of 15 hops (hop-count based)
- Split-horizon
- Admin Distance (AD) 120
- Route Table Entries (RTEs) are 20 bytes each
- Uses link local address FF02::9 all-rip-routers multicast group, as the destination address for RIP updates
- Uses UDP port 521
- IPsec rather than MD5 for authentication

RIPng

```
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 ipv6 address FD00:68:8:8::8/128
!
interface Ethernet0/0
 ipv6 address 2001:78::8/64
 ipv6 enable
 ipv6 rip RIPNG enable
 ipv6 rip RIPNG default-information originate
!
interface Serial0/1
 ipv6 address 2001:68::8/64
 ipv6 address FD00:68::8/64
 ipv6 enable
 ipv6 rip RIPNG enable
!
ipv6 router rip RIPNG

# show ipv6 rip database
# show ipv6 route
# debug ipv6 rip
```

EIGRPv6

- EIGRPv6 uses a new Protocol Dependent Module (PDM) and IPv6 transport – 3 new TLVs
  - IPv6_REQUEST_TYPE (0X0401)
  - IPv6_METRIC_TYPE (0X0402)
  - IPv6_EXTERIOR_TYPE (0X0403)
- Same IPv6 protocol number 88 (Extension Header)
- Router ID stays 32 bits (manually configured unless there is an IPv4 interface on router)
- Hellos are sourced from the link-local address and destined to FF02::A (all EIGRP routers)
- “no auto-summary” is disabled by default. - Yeah!
- No split-horizon in the case of EIGRP for IPv6 (because IPv6 supports multiple prefixes per interface)
- Uses MD5 like for IPv4
  (IPsec authentication will be available soon)
- Same metrics (K1-K5)
- Available in 12.4(6)T

EIGRPv6

```
interface FastEthernet 0/0
 ipv6 enable
 ipv6 eigrp 10
 ipv6 bandwidth-percent eigrp <as#> <percent>
 ipv6 summary-address eigrp <as#> <ipv6-addr> [ad]
 ipv6 authentication mode eigrp <as#> md5
 ipv6 authentication key-chain eigrp <as#> <key-chain>
!
ipv6 router eigrp 10
 router-id 10.1.1.1
 no shutdown
 log-neighbor-changes
 log-neighbor-warnings [seconds]
 metric weights tos k1 k2 k3 k4 k5
!
show ipv6 eigrp interfaces
show ipv6 eigrp neighbors [detail]
show ipv6 eigrp topology [all-links]
show ipv6 eigrp traffic
```

OSPFv3

- Highly scalable link-state IGP
- OSPFv3 (RFC 5340) a bit different than OSPFv2 (RFC 2328)
- Fundamental OSPF mechanisms and algorithms unchanged - Packet and LSA formats are different
- Uses Extension Header Type 0x89
- Interfaces can have multiple IPv6 addresses
  - FF02::5 (all OSPF routers)
  - FF02::6 (all OSPF DRs)
- Neighbor Authentication done with IPsec (AH) (RFC 4552)
- IPv4 RIDs, Area IDs, and LSA IDs

OSPFv3

- OSPFv3 runs directly over IPv6 – extension header type 89
- Runs on a per-link basis rather than per-subnet
- OSPFv3 uses link-local IPv6 addresses
- OSPFv3 supports using multiple instances on a single link.
  - OSPFv3 can have multiple subnets on the same interface.
  - Multiple OSPFv3 instances (different AS numbers) can operate over the same link.
  - A single link can belong to multiple OSPFv3 areas.
  - This is due to the use of the "Instance ID" in the packet headers.

OSPFv3 & OSPFv2

- OSPFv3 operates in a ships-in-the-night method with OSPFv2 - they are mutually exclusive. The OSPFv2 and OSPFv3 topologies don't have to be congruent; they are independent of each other
- OSPFv3 and OSPFv2 can be run concurrently because each address family has a separate SPF
- One must run two different versions of OSPF (OSPFv2 for IPv4 routes and OSPFv3 for IPv6 routes) for dual-protocol routing
- Same Interface types: P2P, P2MP, Broadcast, NBMA, Virtual
- IPv4 and IPv6 control plane are different for OSPFv2 and OSPv3 respectively.
  - OSPFv2 IPv4 routes us IPv4 protocol 89 and OSPFv3 IPv6 routes use IPv6 protocol 89.
  - This reduces the risk of having one protocol negatively affect the other protocol's operation.

OSPFv2 LSA Types

- OSPFv3 Packet and LSA formats are different than OSPFv2 (OSPFv3 has two new LSA types)

LSA Name	LS Type code	Flooding scope	LSA Function code
Router LSA	0x2001	Area scope	1
Network LSA	0x2002	Area scope	2
Inter-Area-Prefix-LSA	0x2003	Area scope	3
Inter-Area-Router-LSA	0x2004	Area scope	4
AS-External-LSA	0x4005	AS scope	5
Group-membership-LSA	0x2006	Area scope	6
Type-7-LSA	0x2007	Area scope	7
Link-LSA	0x0008	Link-local scope	8
Intra-Area-Prefix-LSA	0x2009	Area scope	9

OSPFv3 Configuration

```
interface GigabitEthernet 0/0
 description Area 0 backbone interface
 ipv6 address 2001:db8:100:1::1/64
 ipv6 ospf network broadcast
 ipv6 ospf 100 area 0.0.0.0
!
interface GigabitEthernet 0/1
 description Area 1 interface
 ipv6 address 2001:db8:200:2::1/64
 ipv6 ospf network broadcast
 ipv6 ospf 100 area 0.0.0.1
!
ipv6 router ospf 100
 router-id 10.1.1.1
 area 0.0.0.0 range 2001:db8:100:1::/64
 area 0.0.0.1 range 2001:db8:200:2::/64
```

OSPFv3

```
2821-2#show ipv6 ospf
 Routing Process "ospfv3 100" with ID 2.2.2.2
 It is an area border router
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 LSA group pacing timer 240 secs
 Interface flood pacing timer 33 msecs
 Retransmission pacing timer 66 msecs
 Number of external LSA 0. Checksum Sum 0x000000
 Number of areas in this router is 2. 2 normal 0 stub 0 nssa
 Reference bandwidth unit is 100 mbps
    Area BACKBONE(0)
        Number of interfaces in this area is 1
        SPF algorithm executed 3 times
        Number of LSA 8. Checksum Sum 0x03E314
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
    Area 2
        Number of interfaces in this area is 1
        SPF algorithm executed 2 times
        Number of LSA 5. Checksum Sum 0x02DD13
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

OSPFv3

```
2821-2#show ipv6 ospf neighbor

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
1.1.1.1           1   FULL/DR         00:00:30    4               GigabitEthernet0/0

2821-2#show ipv6 ospf int gig 0/1
GigabitEthernet0/1 is up, line protocol is up
  Link Local Address FE80::213:C4FF:FE56:47B9, Interface ID 5
  Area 2, Process ID 100, Instance ID 0, Router ID 2.2.2.2
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 2.2.2.2, local address FE80::213:C4FF:FE56:47B9
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
  Index 1/1/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

OSPFv3

```
2821-1#show ipv6 ospf database

            OSPFv3 Router with ID (192.168.10.1) (Process ID 10)

                Router Link States (Area 0.0.0.0)

ADV Router      Age         Seq#       Fragment ID  Link count  Bits
192.168.10.1    1242        0x80000008 0            1           EB
192.168.10.33   1229        0x80000007 0            1           B

                Net Link States (Area 0.0.0.0)

ADV Router      Age         Seq#       Link ID    Rtr count
192.168.10.33   1229        0x80000004 2078       2

                Inter Area Prefix Link States (Area 0.0.0.0)

ADV Router      Age         Seq#       Prefix
192.168.10.1    1242        0x80000004 2001:DB8:12::/64
192.168.10.1    1242        0x80000004 2001:DB8:23::/64
192.168.10.1    1242        0x80000004 2001:DB8:30::/64
192.168.10.33   1470        0x80000004 2001:DB8:30::/64
192.168.10.33   1470        0x80000004 2001:DB8:12::/64
192.168.10.33   1470        0x80000004 2001:DB8:23::/64
```

OSPFv3

```
2821-3#sh ipv6 route
IPv6 Routing Table - 11 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, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OE2  ::/0 [110/1], tag 10
     via FE80::205:32FF:FE6F:1EE1, FastEthernet0/0
OI   2001:DB8:10::/64 [110/3]
     via FE80::205:32FF:FE6F:1EE1, FastEthernet0/0
B    2001:DB8:12::/64 [20/0]
     via FE80::205:32FF:FE6F:1EE1, FastEthernet0/0
C    2001:DB8:23::/64 [0/0]
     via ::, FastEthernet0/0
L    2001:DB8:23::3/128 [0/0]
     via ::, FastEthernet0/0
L    2001:DB8:23:0:205:32FF:FE34:6A80/128 [0/0]
     via ::, FastEthernet0/0
L    FE80::/10 [0/0]
     via ::, Null0
L    FF00::/8 [0/0]
     via ::, Null0
```

Optimized OSPF Configuration

```
service internal
!
ip routing protocol purge
interface
ip cef table loadinfo force
ip cef accounting non-recursive
ip cef linecard ipc service-timer
50
no tag-switching ip
!
process-max-time 50
!
interface POSx/y
 dampening
 pos ais-shut
 pos report lais
 pos report lrdi
 pos report pais
 pos report prdi
!
```

```
router ospf x
 ispf
 log-adjacency-changes detail
 timers throttle spf 50 50 5000
 timers throttle lsa all 0 20 5000
 timers lsa arrival 15
 timers pacing flood 15
 passive-interface Loopback0
!
ipv6 router ospf x
 log-adjacency-changes detail
 passive-interface Loopback0
 timers spf 0 1
 timers pacing flood 15
```

12.0(28)S

IS-IS for IPv6

- IS-IS is an extendible link-state intra-domain routing protocol
  - OSI standard ISO/IEC 10589:2002
- 2 level hierarchical architecture
  - Level 1 router – Responsible for intra area routing
  - Level 2 router – Responsible for inter area routing
  - Level 1/2 router – Responsible for intra/inter area routing
- RFC 5308 adds IPv6 address family support to IS-IS.
- RFC 5120 - M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)
  - RFC 5120 defines Multi-Topology concept for IS-IS and permits IPv4 and IPv6 topologies which are not identical.

IS-IS for IPv6

- New Tag/Length/Value (TLVs) added for IPv6
  - IPv6 Reachability TLV (0xEC, TLV type 236)—Describes network reachability (IPv6 routing prefix, metric information and option bits). Defines the External bit, Equivalent to IP Internal/External Reachability TLV's. The option bits indicate the advertisement of IPv6 prefix from a higher level, redistribution from other routing protocols. Equivalent to IP Internal/External Reachability TLVs described in RFC1195.
  - IPv6 Interface Address TLV (0xE8, TLV type 232)—Contains 128-bit address. Hello PDUs, must contain the link-local address but for LSP, must only contain the non-link-local address.
  - A new Network Layer Protocol Identifier (NLPID)—Allows IS-IS routers with IPv6 support to advertise IPv6 prefix payload using 0x8E value (IPv4 and OSI uses different values). This NLPID is advertised by IPv6 enabled routers

IS-IS for IPv6

- IS-IS can operate in several modes when it comes to ipv6: Single-Topology, Multi-Topology, or Multi-Instance. The following section provides the characteristics of each of these modes.
- Single-Topology
  - Single topology (default for all protocols supported). Potentially beneficial in saving resources (same topology and same SPF)
  - All routers must support the same address families (dual-stack, topologically congruent network). Adjacency checking should be disabled during migration.
  - Interface metrics apply to both IPv4 and IPv6

IS-IS for IPv6

- Multi-Topology
  - New TLVs attributes for Multi-Topology extensions.
  - Multi-topology TLV: contains one or more multi-topology ID in which the router participates. It is theoretically possible to advertise an infinite number of topologies. This TLV is included in IIH and the first fragment of a LSP.
  - MT Intermediate Systems TLV: this TLV appears as many times as the number of topologies a node supports. A MT ID is added to the extended IS reachability TLV type 22.
  - Multi-Topology Reachable IPv4 Prefixes TLV: this TLV appears as many times as the number of IPv4 announced by an IS for a give n MT ID. Its structure is aligned with the extended IS Reachability TLV Type 236 and add a MT ID.
  - Multi-Topology Reachable IPv6 Prefixes TLV: this TLV appears as many times as the number of IPv6 announced by an IS for a given MT ID. Its structure is aligned with the extended IS Reachability TLV Type 236 and add a MT ID.
  - Multi-Topology ID (MT ID) standardized and in use in Cisco IOS:
    - MT ID #0 – “standard” topology for IPv4/CLNS
    - MT ID #2 – IPv6 Routing Topology.

IPv6 for IS-IS

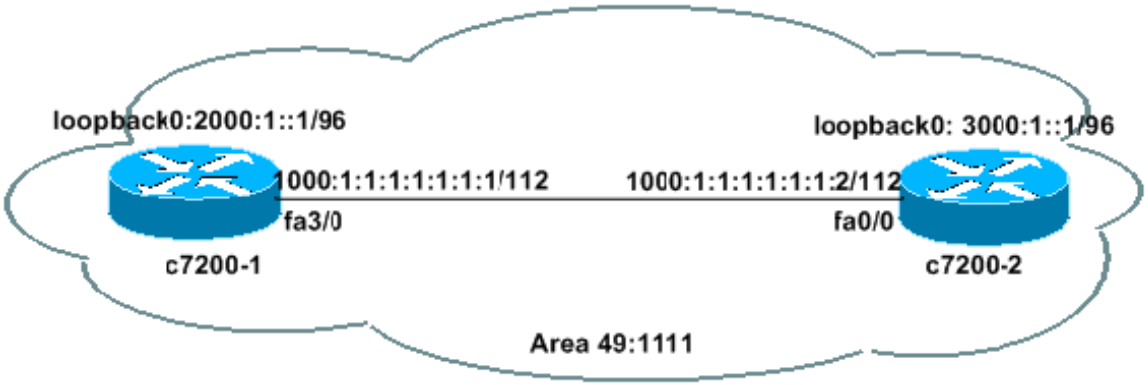


```
C7200-1#show run
ipv6 cef
ipv6 unicast-routing
!
interface loopback0
 no ip address
 ipv6 address 2000:1::1/96
!
ipv6 router isis alpha
!
interface fastethernet3/0
 no ip address
 ipv6 address 1000:1:1:1:1:1:1:1/112
 router isis alpha
!
router isis alpha
 net 49.1111.2220.3330.4440.00
```

IS-IS Single Topology

```
Router1#show isis database verbose level-1
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num    LSP Checksum    LSP Holdtime    ATT/P/OL
Router2.00-00         0x0000000B     0xAB35          1020            0/0/0
  Area Address: 49.0001
  NLPID:        0xCC 0x8E
  Hostname: Router2
  IP Address:   10.7.1.34
  Metric: 10         IP 10.7.1.32 255.255.255.252
  IPv6 Address: 2001:db8:FFFF::2
  Metric: 10         IPv6 2001:db8:FFFF::/64
  Metric: 10         IS Router2.01
```

```
router isis example-area
 net 49.0001.0000.0000.0001.00
!
interface FastEthernet0/1
 ip address 10.7.1.33 255.255.255.252
 ip router isis example-area
 ipv6 address 2001:db8:FFFF::1/64
 ipv6 enable
 ipv6 router isis example-area
```

```
Router1#show clns is-neighbors detail
System Id      Interface   State  Type Priority  Circuit Id          Format
Router2         Fa0/1        Up       L1L2 64/64     Router2.0
  Area Address(es): 49.0001
  IP Address(es):  10.7.1.34*
  IPv6 Address(es): FE80::2B0:4AFF:FE5C:ACA9
  Uptime: 00:01:25
  NSF capable
```

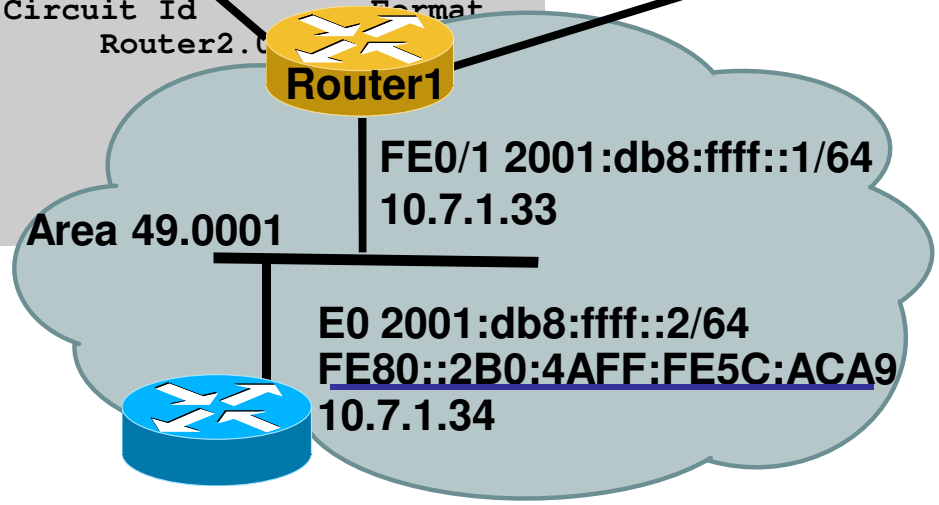

IS-IS Multi Topology

```
Router1#show isis database verbose level-1
IS-IS Level-1 Link State Database:
LSPID             LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
Router2.00-00     0x00000014   0x8B3E        1086          0/0/0
  Area Address: 49.0001
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname: Router2
  IP Address:   10.7.1.34
  Metric: 10          IP 10.7.1.32/30
  IPv6 Address: 2001:db8:FFFF::2
  Metric: 10          IPv6 (MT-IPv6) 2001:db8:FFFF::/64
  Metric: 10          IS (MT-IPv6) Router2.01
```

```
Router1#show clns is-neighbors detail
System Id  Interface State Type Priority Circuit Id  Format
Router2    Fa0/1     Up    L1L2 64/64    Router2.01  Phase V
  Area Address(es): 49.0001
  IP Address(es):  10.7.1.34*
  IPv6 Address(es): FE80::2B0:4AFF:FE5C:ACA9
  Uptime: 00:00:14
  NSF capable
  Topology: IPv4, IPv6
```

```
router isis example-area
 net 49.0001.0000.0000.0001.00
 metric-style wide transition
 !
 address-family ipv6
 multi-topology transition
```

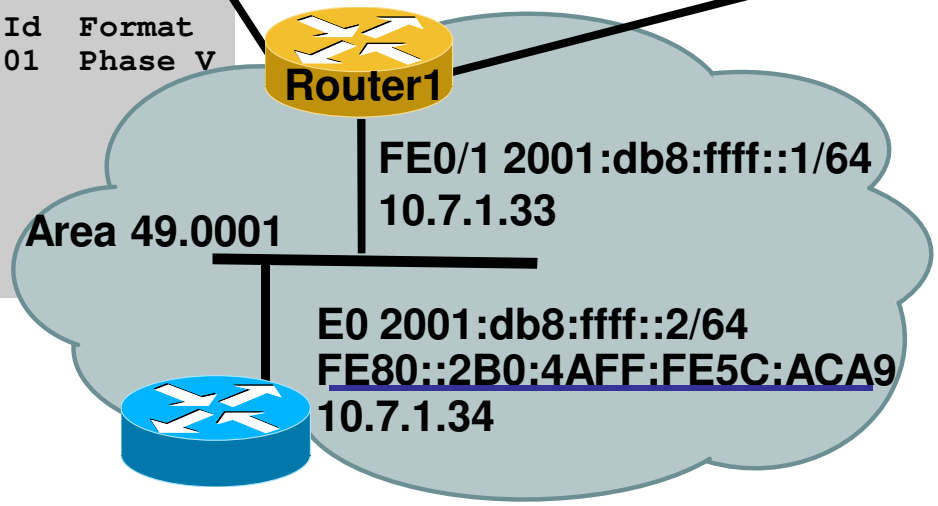

Example of Optimized ISIS Configuration

```
service internal
!
ip routing protocol purge interface
ip cef table loadinfo force
ip cef accounting non-recursive
ip cef linecard ipc service-timer 50
no tag-switching ip
!
process-max-time 50
!
interface POSx/y
 dampening
 pos ais-shut
 pos report lais
 pos report lrdi
 pos report pais
 pos report prdi
!
```

12.0(28)S

```
router isis
 net 49.0002.0000.0000.00
 is-type level-2-only
 ispf level-2
 metric-style wide
 fast-flood
 set-overload-bit on-startup 360
 max-lsp-lifetime 65535
 lsp-refresh-interval 65000
 spf-interval 5 1 100
 prc-interval 5 1 100
 lsp-gen-interval 5 1 100
 no hello padding
 log-adjacency-changes all
 passive-interface Loopback0
 !
 address-family ipv6
 multi-topology
 spf-interval 5 1 100
 prc-interval 5 1 100
 set-overload-bit on-startup 360
 exit-address-family
```

Multiprotocol BGP-4

- Multiprotocol Extensions for BGP-4 (RFC4760)
- New optional and non-transitive attributes (NLRIs)
  - MP_REACH_NLRI (attribute code: 14)
  - MP_UNREACH_NLRI (attribute code: 15)
- Address Family Identifier (AFI=2) for IPv6
- BGP TCP port 179 sessions can be over IPv4 or IPv6
  - MPBGP still relies upon a stable IGP
- Next-Hop attribute must be link-local or aggregatable global unicast IPv6 address
- 32-bit Router ID is required if no IPv4 interfaces
  - Still used as tie-breaker – sent in OPEN messages

Multiprotocol BGP-4

- Next hop reachability is VERY important in BGP
- If the Next Hop is inaccessible, the routes learned via BGP will not be installed in the routing table (In some case the routes will be rejected by BGP)
- 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 is also on a link-local address
- IPv6 NLRI in IPv6 (Global Unicast) works like IPv4 (3rd party NH not supported yet)

BGP-4 Configuration

```
router bgp 65500
 bgp log-neighbor-changes
 neighbor 2001:db8:100:10::8 remote-as 64900
 neighbor 172.16.1.2 remote-as 65500
!
 address-family ipv4
 neighbor 172.16.1.2 activate
 neighbor 172.16.1.2 prefix-list OUTFILTER out
 no neighbor 2001:db8:100:10::8 activate
 network 192.168.1.0
 no auto-summary
 no synchronization
 !
 address-family ipv6
 neighbor 2001:db8:100:10::8 activate
 neighbor 2001:db8:100:10::8 prefix-list FILTERIPV6 out
 network 2001:db8:100:30::/48
 no synchronization
!
ipv6 prefix-list FILTERIPV6 seq 10 permit 2001:db8:100:30::/48
ipv6 route 2001:db8:100:30::/48 null0
```

BGP-4 Configuration

- Link Local Peering

```
router bgp 100
 neighbor FE80::AABB:CCFF:FE01:F801%Ethernet0/0 remote-as
   200
 !
 address-family ipv6
 neighbor FE80::AABB:CCFF:FE01:F801%Ethernet0/0 activate
 neighbor FE80::AABB:CCFF:FE01:F801%Ethernet0/0 route-map
   SETNEXTHOP out
 redistribute connected
 no synchronization
!
route-map SETNEXTHOP permit 10
 set ipv6 next-hop 2001:db8:100:100::2
```

BGP-4 Configuration

```
Router#show bgp ipv6 2001:db8:100:100::/64
BGP routing table entry for 2001:db8:100:100::/64, version 71
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     1
  200
    2001:db8:100:100::2 (FE80::AABB:CCFF:FE01:F600) from
    FE80::AABB:CCFF:FE01:F600%Ethernet0/0 (200.11.11.1)
      Origin incomplete, metric 0, localpref 100, valid, external
  Local
    :: from 0.0.0.0 (200.14.14.1)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
```

BGP-4 Configuration

```
router bgp 65500
 bgp log-neighbor-changes
 neighbor 3ffe:1500:32c:74::8 remote-as 64900
 !
 address-family ipv6
 neighbor 3ffe:1500:32c:74::8 activate
 neighbor 3ffe:1500:32c:74::8 maximum-prefix 2500 80
 neighbor 3ffe:1500:32c:74::8 prefix-list FILTERIPV6 in
 network 2001:db8:31::/48
 no synchronization
!
ipv6 prefix-list FILTERIPV6 seq 10 permit 2001:500::/30 le 48
ipv6 prefix-list FILTERIPV6 seq 20 permit 2002::/16
ipv6 prefix-list FILTERIPV6 seq 30 permit 2000::/3 le 32
ipv6 prefix-list FILTERIPV6 seq 40 deny ::/0 le 128
!
ipv6 route 2001:db8:31::/48 null0
```

BGP-4 Configuration

```
interface Ethernet0
ipv6 address 5f00:0100:0:0:1::1 80
!
router bgp 100
no bgp default ipv4-unicast
neighbor 5f00:0100:0:0:2::1 remote-as 101
aggregate-address 2001:420:2000::/42 summary-only
!
address-family ipv6
neighbor 5f00:0100:0:0:2::1 activate
neighbor 5f00:0100:0:0:2::1 prefix-list BGP-IN in
neighbor 5f00:0100:0:0:2::1 prefix-list AGGREGATE out
network 5f00:0100:0:0:1::/40
exit-address-family
!
ipv6 prefix-list AGGREGATE seq 5 deny 3FFE:C00::/24 ge 25
ipv6 prefix-list AGGREGATE seq 10 permit ::/0 le 48
!
ipv6 prefix-list BGP-IN seq 5 deny 5F00::/8 le 128
ipv6 prefix-list BGP-IN seq 10 deny ::/0
ipv6 prefix-list BGP-IN seq 15 deny ::/1
ipv6 prefix-list BGP-IN seq 20 deny ::/2
ipv6 prefix-list BGP-IN n seq 25 deny ::/3 ge 4
ipv6 prefix-list BGP-IN seq 30 permit ::/0 le 128
```

Common IPv6 Routing Mistakes

- Need to define a Router-ID
  - 32-bit “IPv4 address”
- Forgot to “no shut” EIGRPv6 process
- Fat-finger prefix
- Forgot to enable protocol on the interface
- Tuning performance to improve IPv4 and IPv6 routing operation
- Issues with routing over a tunnel
- Avoid using link-local addresses for peers

IPv6 Multicast

- IPv6 multicast works similar to IPv4 multicast
  - Many of the multicast applications can work well with IPv4 or IPv6 sources and receivers
- The same Protocol Independent Multicast (PIM-SM) (RFC 4601) algorithm can operate on IPv4 or IPv6 networks
- Instead of IGMP for IPv4, IPv6 uses Multicast Listener Discovery (MLD)

Multicast Listener Discovery (MLD)

- MLD is equivalent to IGMP in IPv4
- MLD messages use ICMPv6
- MLD uses link-local source addresses
- MLD packets use “Router Alert” extension header (RFC 2711)
- Version number clarification:
  - MLDv1 (RFC 2710) ≈ IGMPv2 (RFC 2236)
  - MLDv2 (RFC 3810) ≈ IGMPv3 (RFC 3376)
- Only MIB available today is for MLDv1
- MLD Snooping (RFC 4541)
- Configuration on Cisco router
  - **`ipv6 mld snooping`**
  - **`ipv6 mld join-group FF05::1`**
  - **`show ipv6 mld [groups | interface | traffic]`**

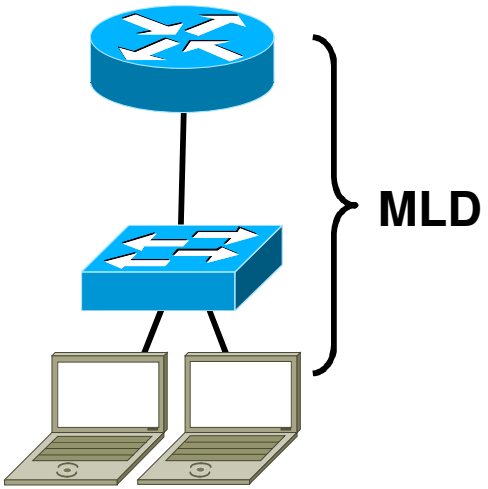

MLD Operations with ICMPv6

- General Query (ICMPv6 Type 130)
  - Sent to learn of listeners on the attached link – sent to ff02::1 (all nodes)
  - Sets the multicast address field to zero
  - Sent every 125 seconds (configurable)
- Report (ICMPv6 Type 131)
  - Contains the multicast address that the sender want to receive
- Leave/DONE (ICMPv6 Type 132)
  - Last host leaves—sends DONE to ff02::2 (all routers)
  - Router will respond with group-specific query (Type 130)
  - Router will use the last member query response interval (Default=1 sec) for each query
  - Query is sent twice and if no reports occur then entry is removed (2 seconds)

A Few Notes on Tunnels…

- PIM uses tunnels when RPs/sources are known
- Source registering (on first-hop router)
  - Uses virtual tunnel interface (appear in OIL for [S,G])
  - Created automatically on first-hop router when RP is known
  - Cisco IOS keeps tunnel as long as RP is known
  - Unidirectional (transmit only) tunnels
  - PIM Register-Stop messages are sent directly from RP to registering router (not through tunnel!)

PIM Tunnels (DR-to-RP)

```
branch#show ipv6 pim tunnel
Tunnel1*
  Type   : PIM Encap
  RP     : 2001:DB8:C003:1116::2
  Source: 2001:DB8:C003:111E::2
```

```
branch#show interface tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 2001:DB8:C003:111E::2 (Serial0/2),
destination 2001:DB8:C003:1116::2
  Tunnel protocol/transport PIM/IPv6, key disabled,
sequencing disabled
  Checksumming of packets disabled
  Tunnel is transmit only
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
… output truncated…
```

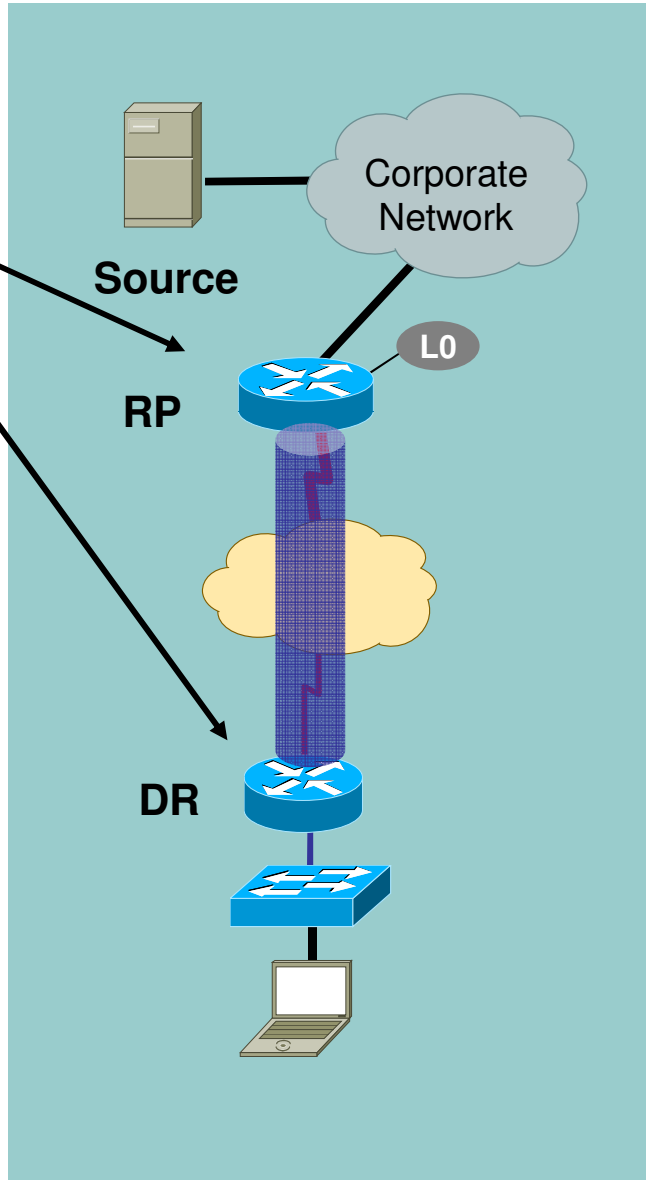

PIM Tunnels (RP)

- Source registering (on RP) → two virtual tunnels are created
  - 1 transmit only for registering sources locally connected to the RP
  - 1 receive only for decapsulation of incoming registers from remote designated routers
  - No one-to-one relationship between virtual tunnels on designated routers and RP!

PIM Tunnels (RP-for-Source)

```
RP-router#show ipv6 pim tunnel
Tunnel0*
  Type   : PIM Encap
  RP     : 2001:DB8:C003:1116::2
  Source: 2001:DB8:C003:1116::2
Tunnel1*
  Type   : PIM Decap
  RP     : 2001:DB8:C003:1116::2
  Source:  -
```

```
RP-router#show interface tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 2001:DB8:C003:1116::2
(FastEthernet0/0), destination 2001:DB8:C003:1116::2
  Tunnel protocol/transport PIM/IPv6, key disabled,
sequencing disabled
  Checksumming of packets disabled
  Tunnel is receive only
  … output truncated…
```

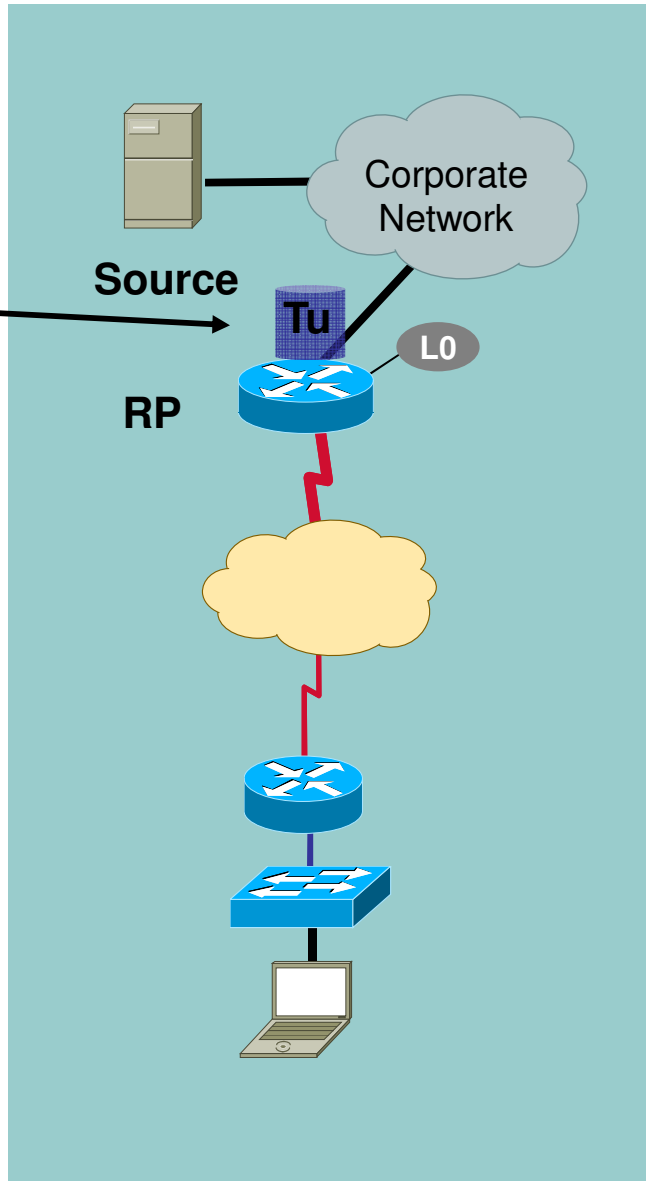

Multicast Interdomain Options

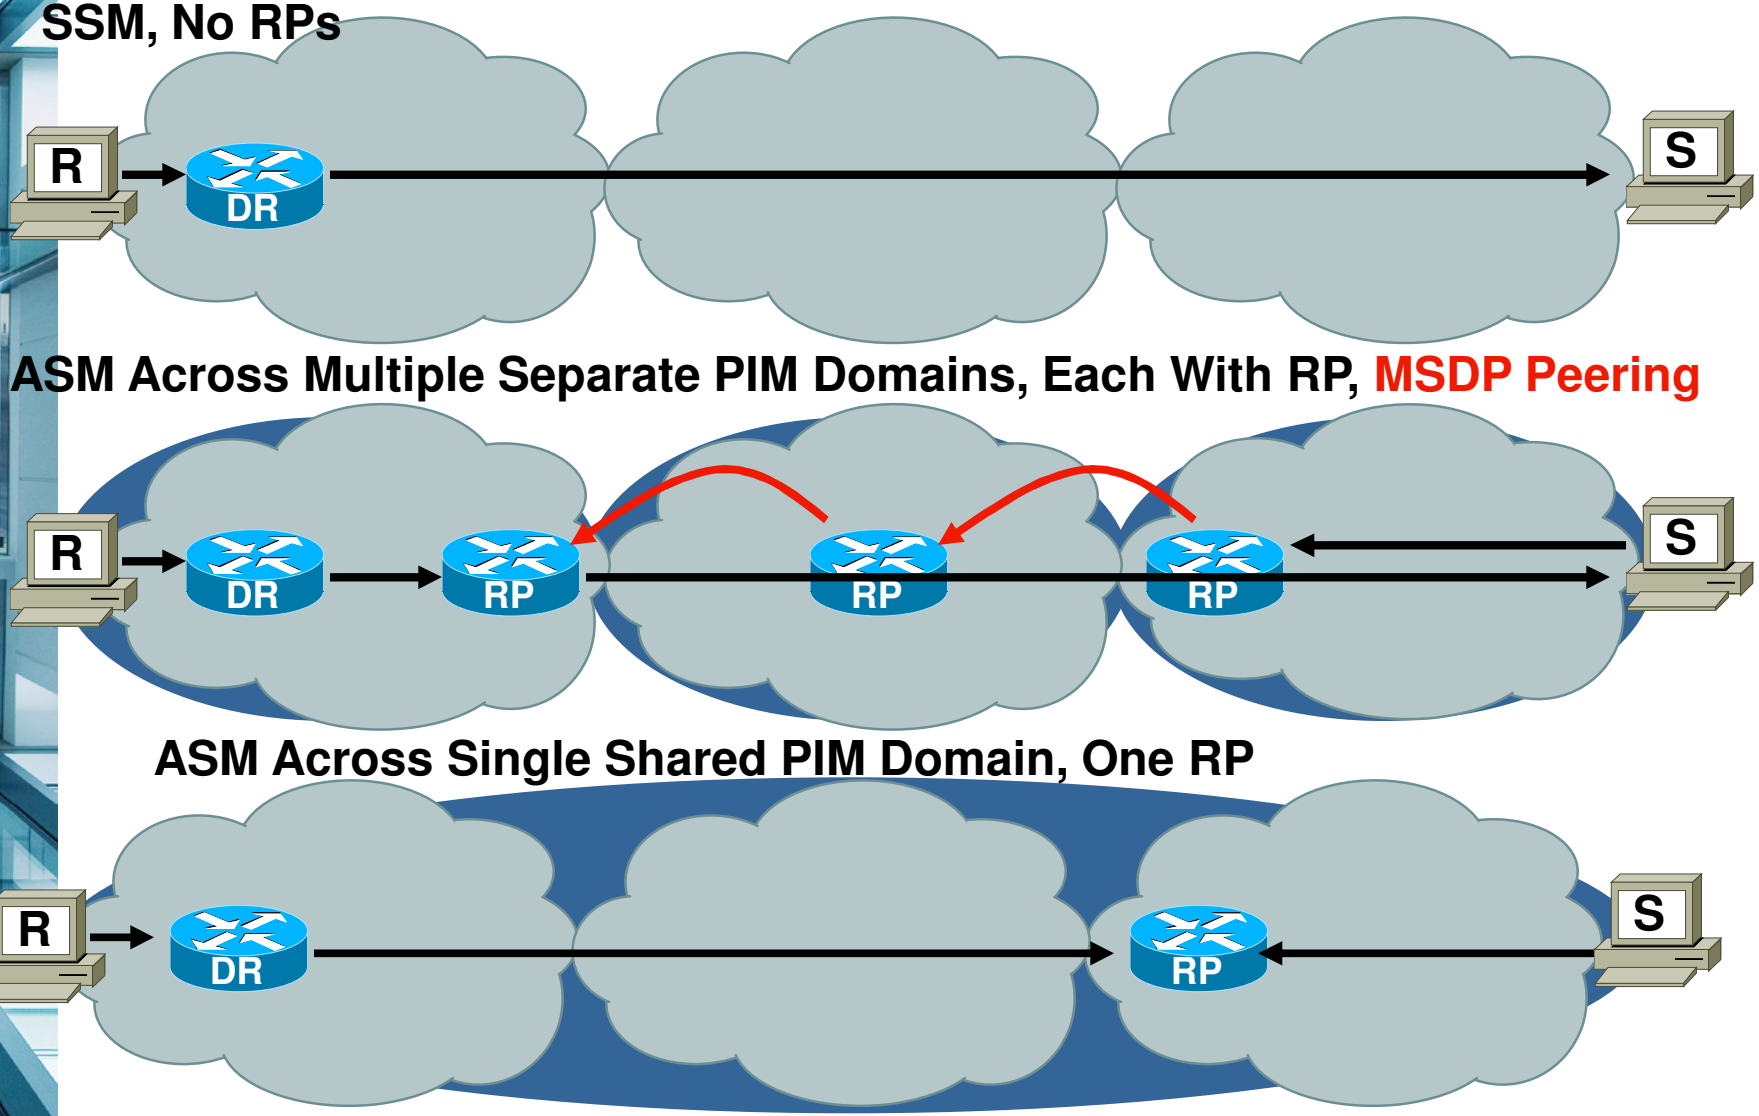

IPv6 Multicast Static RP

- Even easier than before
- 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
```

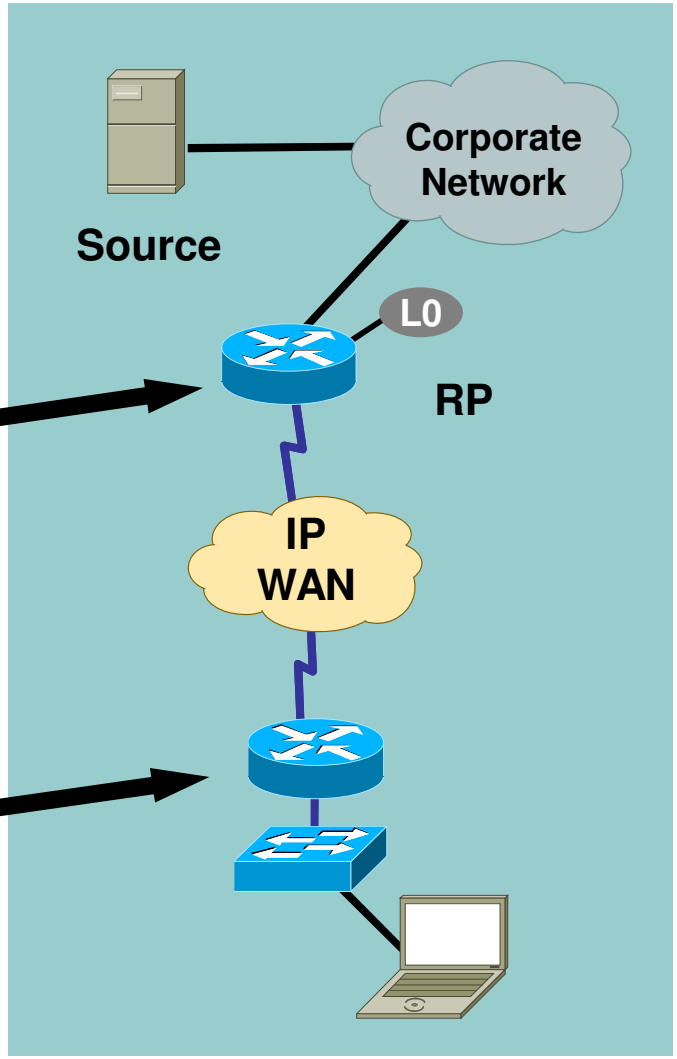

Bidirectional PIM (Bidir)

- 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

```
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
```

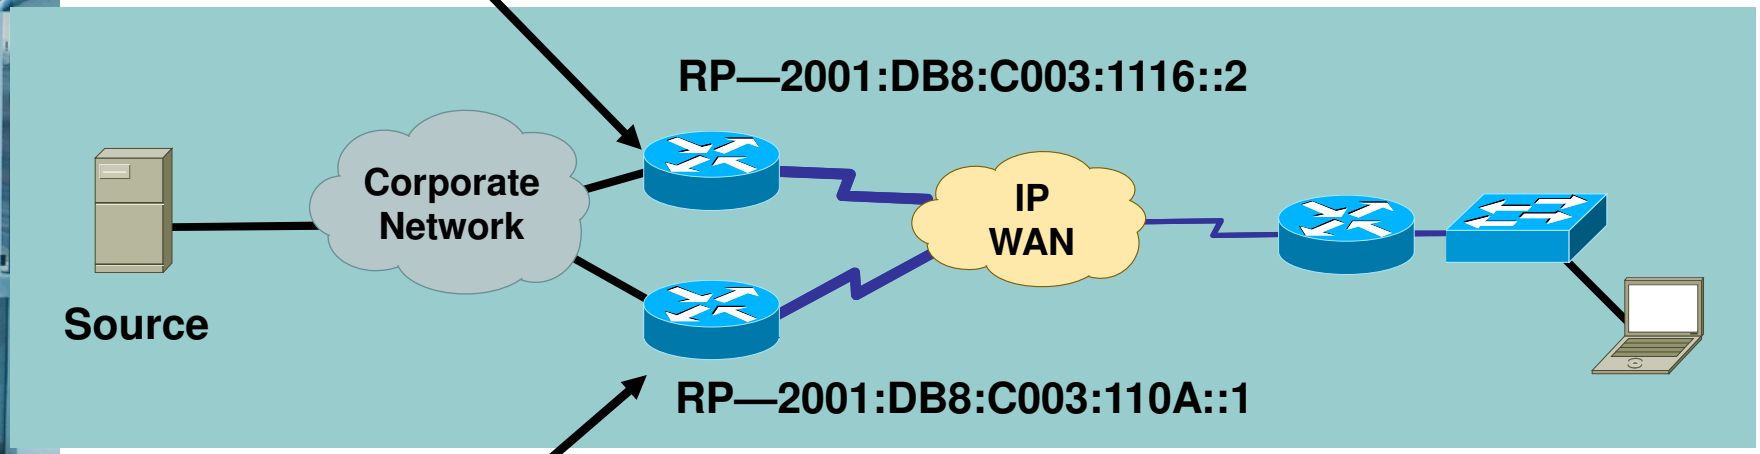


```
wan-bottom#sh run | incl ipv6 pim bsr

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

Embedded-RP Addressing

- Relies on a subset of RFC3306—IPv6 unicast-prefix-based multicast group addresses with special encoding rules:
  - Group address carries the RP address for the group!

8	4	4	4	4	8	64	32
FF	Flags	Scope	Rsvd	RPaddr	Plen	Network Prefix	Group ID

New address format defined :

Flags = 0RPT, R = 1, P = 1, T = 1 => RP address embedded

(0111=7)

Example Group: FF7E:0140:2001:0DB8:C003:111D:0000:1112

Embedded RP: 2001:0DB8:C003:111D::1

Embedded-RP Configuration

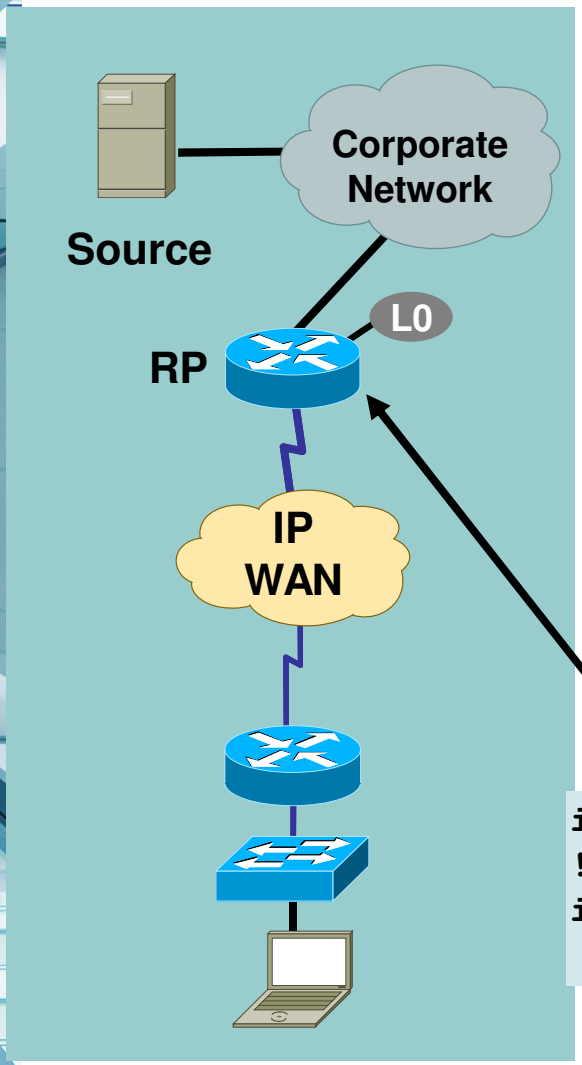

- RP to be used as an Embedded-RP needs to be configured with address/group range
- All other non-RP routers require no special configuration

```
ipv6 pim rp-address 2001:DB8:C003:111D::1 ERP
!
ipv6 access-list ERP
 permit ipv6 any FF7E:140:2001:DB8:C003:111D::/96
```

Embedded RP Operation

```
branch#show ipv6 pim group

FF7E:140:2001:DB8:C003:111D ::/96*
    RP       : 2001:DB8:C003:111D::1
    Protocol: SM
    Client   : Embedded
    Groups   : 1
    Info     : RPF: Se0/0.1,FE80::210:7FF:FEDD:40
```

```
branch#show ipv6 mroute active

Active IPv6 Multicast Sources - sending >= 4 kbps
Group: FF7E:140:2001:DB8:C003:111D:0:1112
  Source: 2001:DB8:C003:1109::2
    Rate: 21 pps/122 kbps(1sec), 124 kbps(last 100 sec)
```

```
branch#show ipv6 pim range | include Embedded

Embedded SM RP: 2001:DB8:C003:111D::1 Exp: never Learnt from : ::

 FF7E:140:2001:DB8:C003:111D::/96 Up: 00:00:24
```

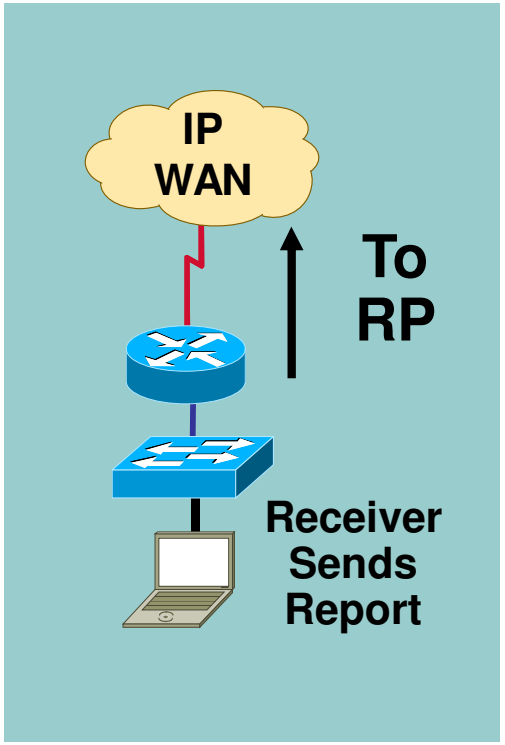

RP Redundancy

- 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

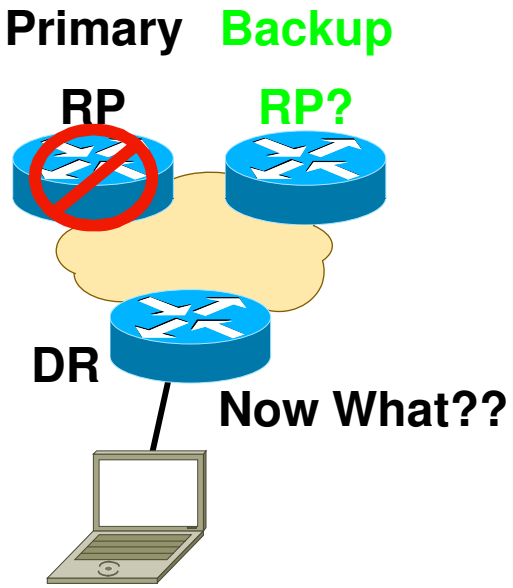

Source Specific Multicast (SSM)

- NO configuration required other than enabling
  - `ipv6 multicast-routing`
- SSM group ranges are automatically defined
- Requires MLDv2 on host or SSM Mapping feature

```
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
```

SSM-Mapping

```
core-1#show ipv6 mroute | begin 2001:DB8:CAFE:11::11
(2001:DB8:CAFE:11::11, FF33::DEAD), 00:01:20/00:03:06, flags: sT
  Incoming interface: GigabitEthernet3/3
  RPF nbr: FE80::20E:39FF:FEAD:9B00
  Immediate Outgoing interface list:
    GigabitEthernet5/1, Forward, 00:01:20/00:03:06
```

Static Mapping:

```
ipv6 multicast-routing
!
ipv6 mld ssm-map enable
ipv6 mld ssm-map static MAP 2001:DB8:CAFE:11::11
no ipv6 mld ssm-map query dns
!
ipv6 access-list MAP
 permit ipv6 any host FF33::DEAD
```

DNS Mapping (the default):

```
ipv6 multicast-routing
!
ipv6 mld ssm-map enable
!
ip domain multicast ssm-map.cisco.com
ip name-server 10.1.1.1
```

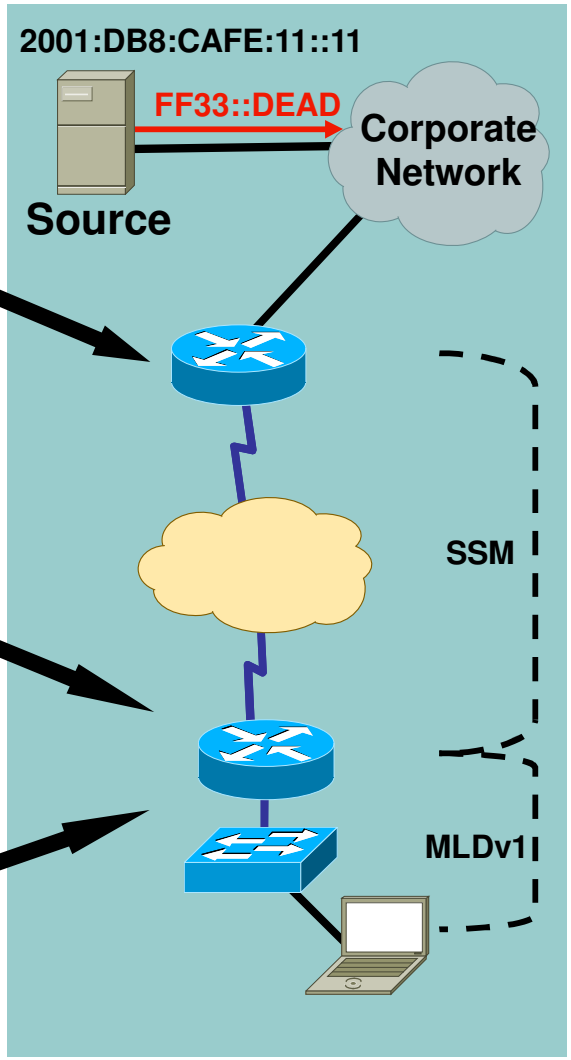

IPv6 Multicast Summary

- One size does NOT fit all
- Stick with basic PIM-SM if you want basic functions
- PIM-SSM is the way to go for one/few-to-many applications, but requires MLDv2 or SSM Mapping and the application 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

Policy-Based IPv6 Routing

- Select packets to be routing or marked in special ways
- Relies on Cisco Express Forwarding (CEF) to check input packets
- Packets can be checked for the:
  - Input interface
  - Source IPv6 address (using a prefix list or a standard or extended access list [ACL])
  - Destination IPv6 address (standard or extended ACL)
  - Protocol (extended ACL)
  - Source port and destination port (extended ACL)
  - Differentiated services code point (DSCP) (extended ACL)
  - Flow-label (extended ACL)
  - Fragment (extended ACL)
- Packets can be set for special routing:
  - IPv6 next hop
  - Output interface
  - Default IPv6 next hop
  - Default output interface

IPv6 General Prefix Config

- Provides an easy/fast way to deploy prefix changes
- Example:2001:db8:10::/48 = General Prefix
- Fill in interface specific fields after prefix
- “GP-1 ::1:0:0:0:1” = 2001:db8:10:1::1/64
  - **ipv6 unicast-routing**
  - **ipv6 cef**
  - **ipv6 general-prefix GP-1 2001:DB8:10::/48**
  - **!**
  - **interface GigabitEthernet0/1**
  - **ipv6 address GP-1 ::2/127**
  - **ipv6 cef**
  - **!**
  - **interface Vlan1**
  - **ipv6 address GP-1 ::1:0:0:0:1/64**
  - **ipv6 cef**

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

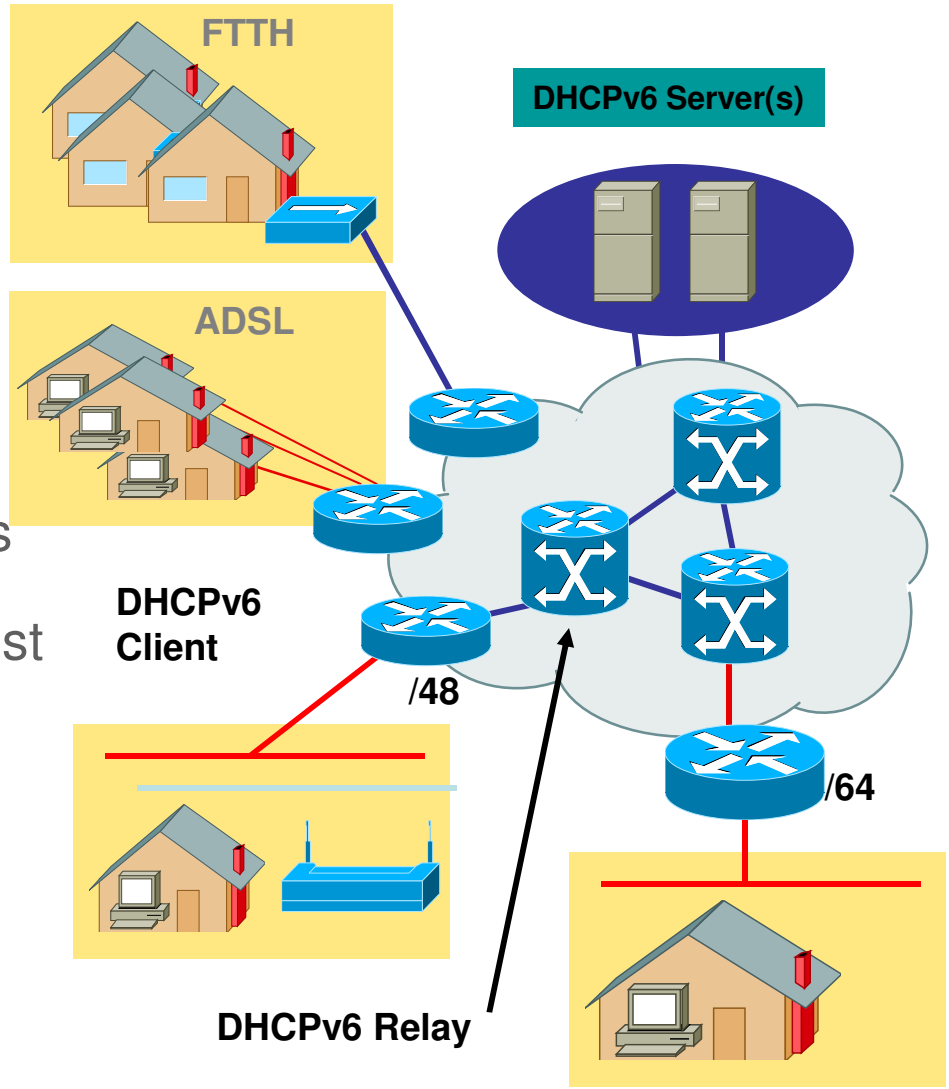

Prefix/Options Assignment

(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

AAA | DHCP | ND/DHCP

DHCPv6 Prefix Delegation

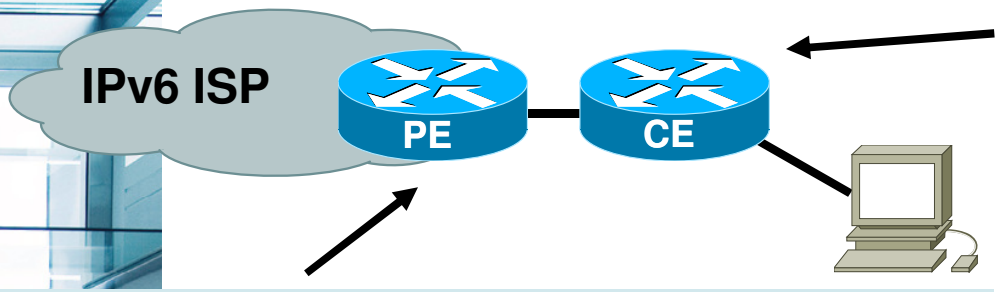


```
vpdn enable
!
vpdn-group pppoe
 accept-dialin
 protocol pppoe
 virtual-template 1
!
ipv6 dhcp pool FOO
 prefix-delegation 2001:7:7::/48 0003000100055FAF2C08
 prefix-delegation 2001:8:8::/48 0003000100055FAC1808
 dns-server 2001:4::1
 domain-name cisco.com
!
interface Virtual-Template1
 ipv6 enable
 no ipv6 nd suppress-ra
 ipv6 dhcp server FOO
 ppp authentication chap
!
interface FastEthernet1/0
 pppoe enable
```

```
vpdn enable
!
vpdn-group 1
 request-dialin
 protocol pppoe
!
interface FastEthernet0/1
 ipv6 address DH-PREFIX 0:0:0:1::/64 eui-
64
!
interface FastEthernet0/0
 pppoe enable
 pppoe-client dial-pool-number 1
!
interface Dialer1
 encapsulation ppp
 dialer pool 1
 dialer-group 1
 ipv6 address autoconfig
 ipv6 dhcp client pd DH-PREFIX
 ppp authentication chap callin
 ppp chap hostname dhcp
 ppp chap password 7 0300530816
!
ipv6 route ::/0 Dialer1
```

http://www.cisco.com/en/US/tech/tk872/technologies_white_paper09186a00801e199d.shtml

Default Gateway Redundancy

- Neighbor Unreachability Detection (NUD) can be used
  - Very rudimentary HA capability
    - `(config-if)#ipv6 nd reachable-time 5000`
  - Hosts use NUD “reachable time” to cycle to next known default gateway (30 seconds by default)
  - Use NUD as a last resort if HSRP or GLBP are not available
- HSRP for IPv6 and GLBP for IPv6 now exist in IOS 12.4
- VRRP for IPv6 exists on MDS 9000s but not released in IOS 12.4 yet

HSRP for IPv6

- Many similarities with HSRP for IPv4
- Changes occur in Neighbor Advertisement, Router Advertisement, and ICMPv6 redirects
- No need to configure GW on hosts
- RAs are sent from HSRP Active router
- 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 or specific debug
- MD5 authentication for HSRP for IPv6 (HSRP used to just use keyword)

HSRP for IPv6

```
interface FastEthernet0/1
 ipv6 address 2001:DB8:66:67::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
```

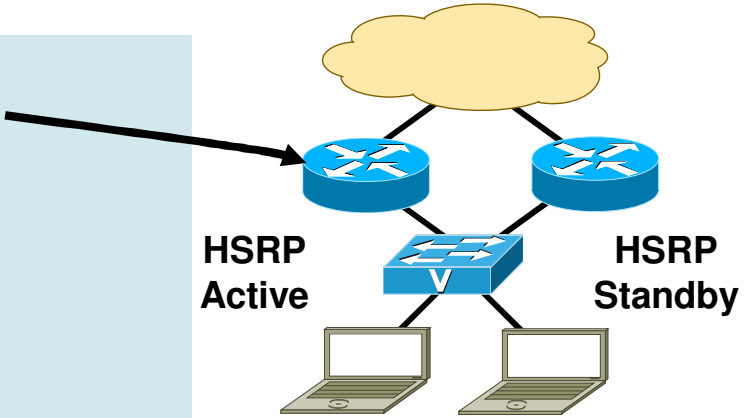

Host with GW of Virtual IP

```
#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
```

GLBP for IPv6

- Many similarities with GLBP for IPv4 (CLI, Load-balancing)
- Modification to Neighbor Advertisement, Router Advertisement
- GW is announced via RAs
- Virtual MAC derived from GLBP group number and virtual IPv6 Link-local address

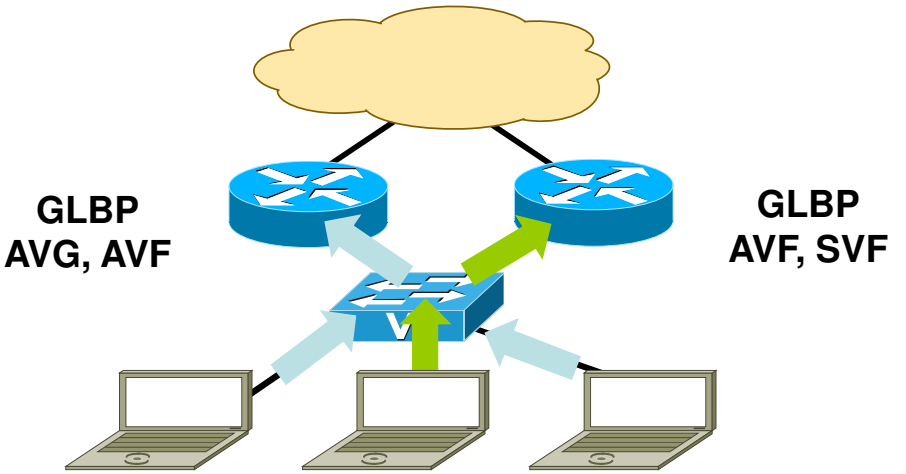


```
interface FastEthernet0/0
 ipv6 address 2001:DB8:1::1/64
 ipv6 cef
 glbp 1 ipv6 autoconfig
 glbp 1 timers msec 250 msec 750
 glbp 1 preempt delay minimum 180
 glbp 1 authentication md5 key-string cisco
```

AVG=Active Virtual Gateway
AVF=Active Virtual Forwarder
SVF=Standby Virtual Forwarder

IPv6 Quality of Service

- QoS is required for real time services
  - Minimize latency/jitter for VoIP
- 20-bit Flow Label - IDs of traffic flows (RFC3697)
  - Sequence of packets for which a source desires to label a flow
  - Flow classifiers have been based on 5-tuple: source/destination address, protocol type and port numbers of transport
  - Some of these fields may be unavailable due to fragmentation, encryption or locating them past extension headers.
  - Looking for classifier only into IP header
  - Only 3-tuple, flow label, source/destination address
- Drop Priority field to manage conflicts
- RSVP used by routers to deal with requests

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
- Keep an eye out for work being doing to leverage the flow label

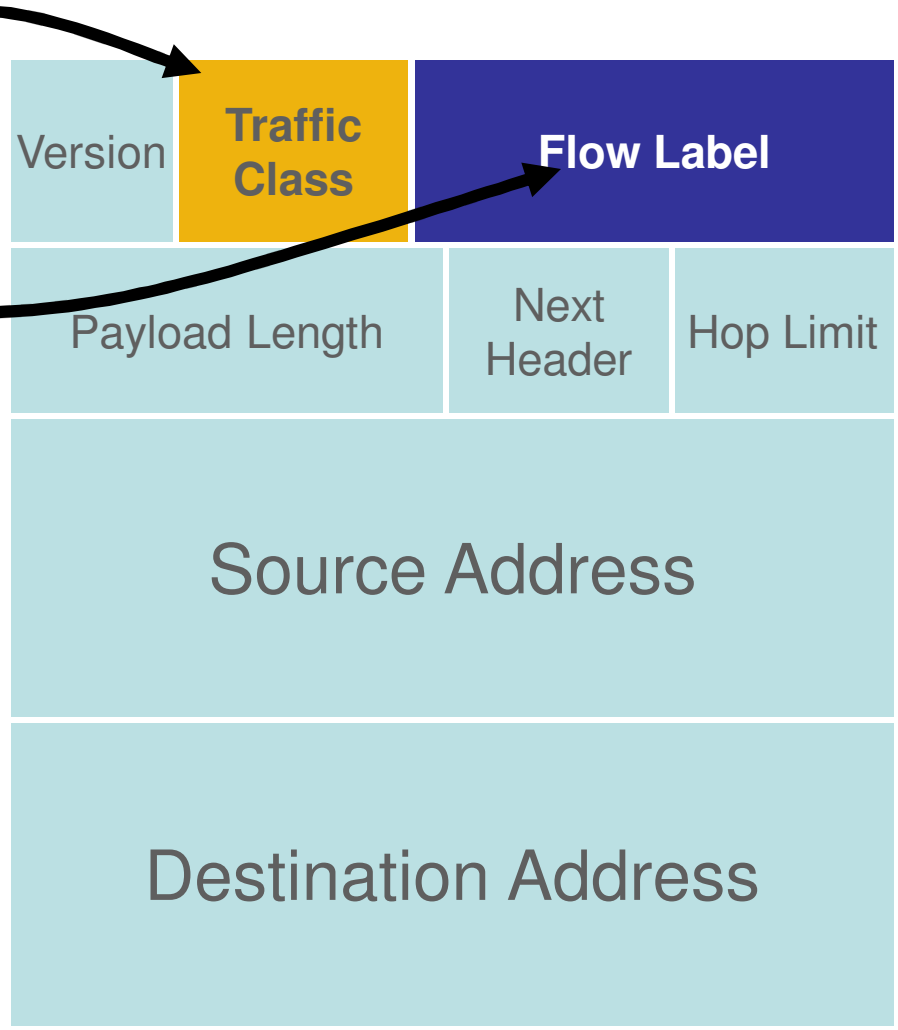

IPv6 QoS Syntax Changes

- 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 Config: IPv4/v6

```
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
```

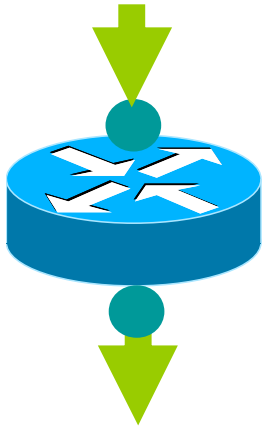

```
service-policy output RBR-WAN-EDGE
```

ACLs to Match for Both IPv4 and IPv6 Packets

Summary

- An IPv6-enabled Internet already exists
- An IPv6 transition is already underway in the Federal Government and other parts of the world.
- IPv6 infrastructure and Host OSs are ready now!
- Much of the infrastructure you have already purchased is IPv6 capable, it's just a matter of enabling (software upgrade)
- Service providers have initial IPv6 services and are continuing to work on their deployments
- You will be transitioning to IPv6 over the coming year – some perimeter IPv6 Internet connectivity already exists at your organization

IPv6 Books

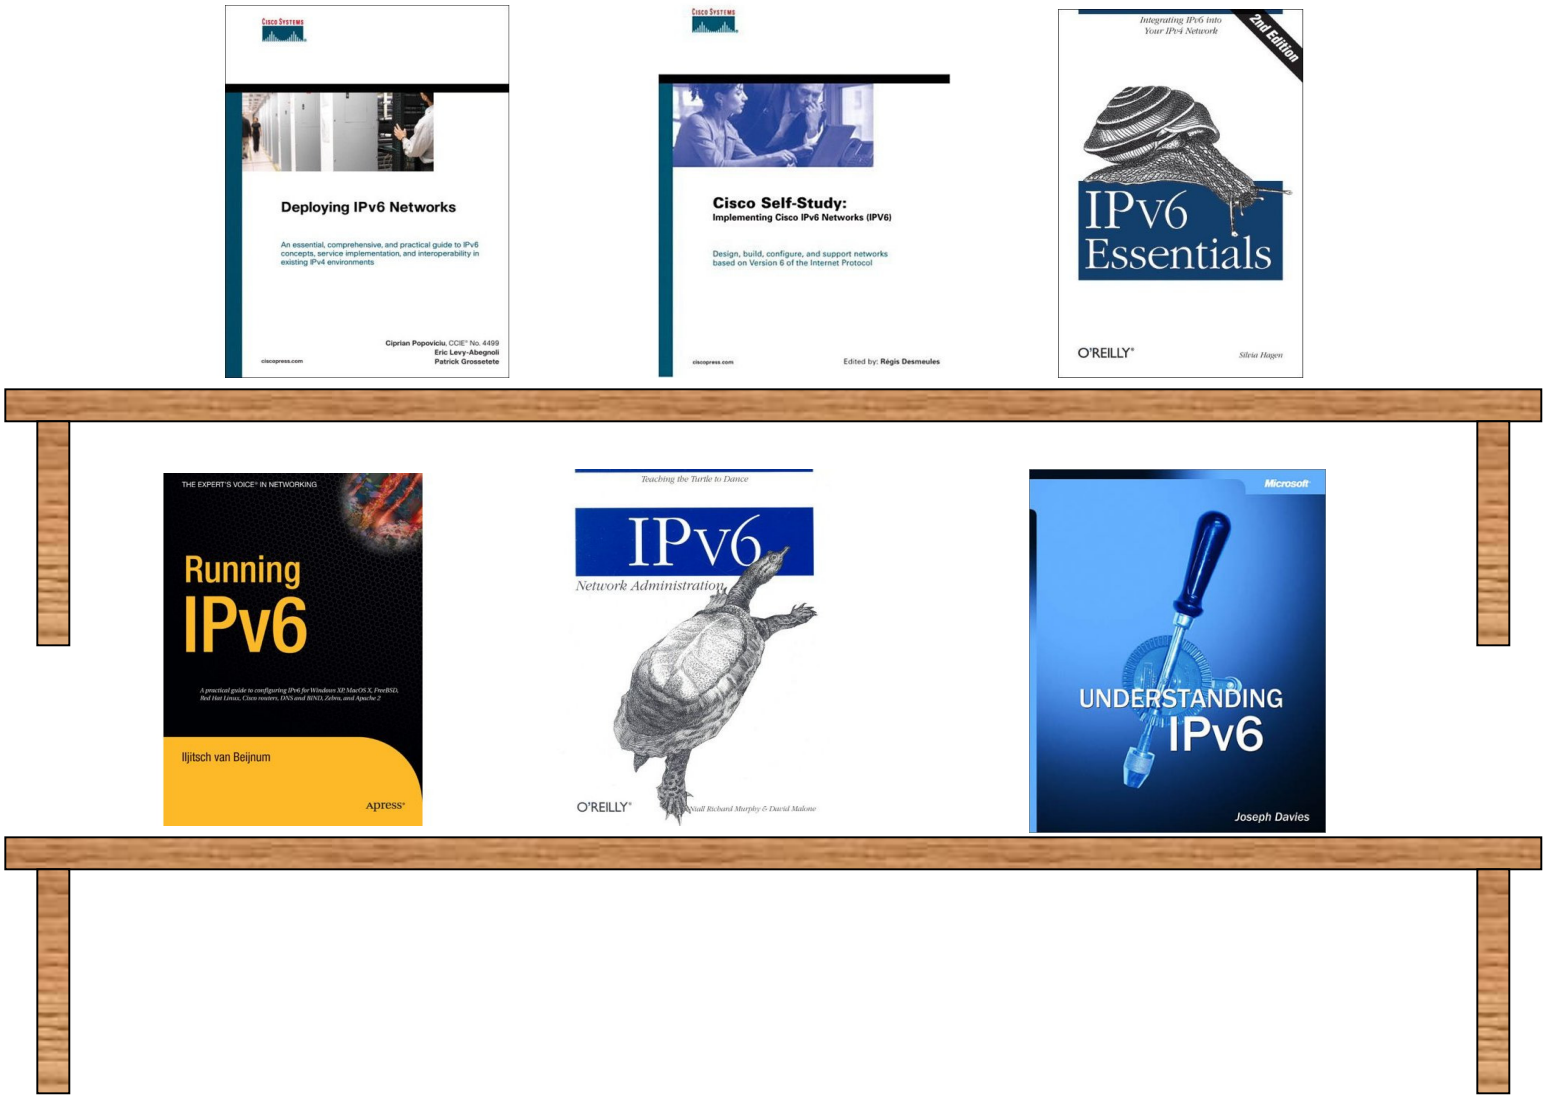

Transition Plan Books

- These books offer a higher-level business perspective on the transition to IPv6

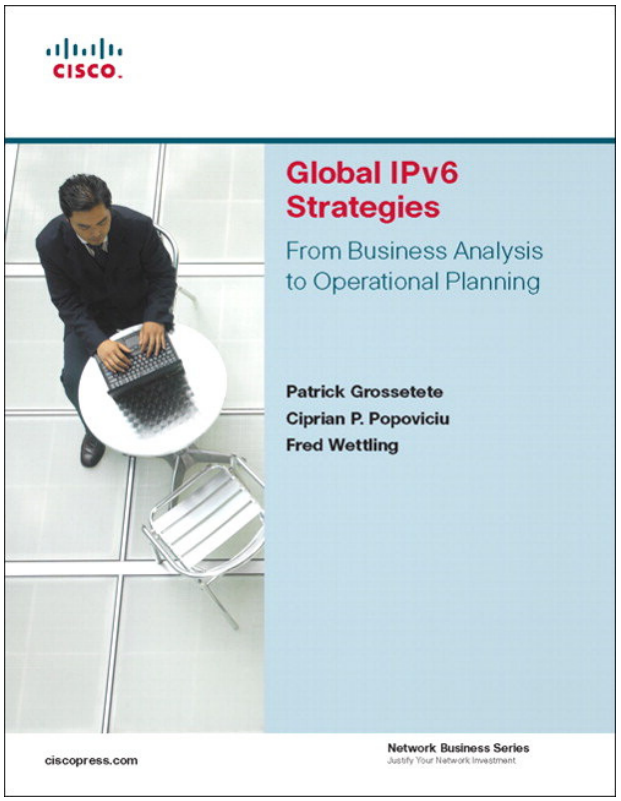

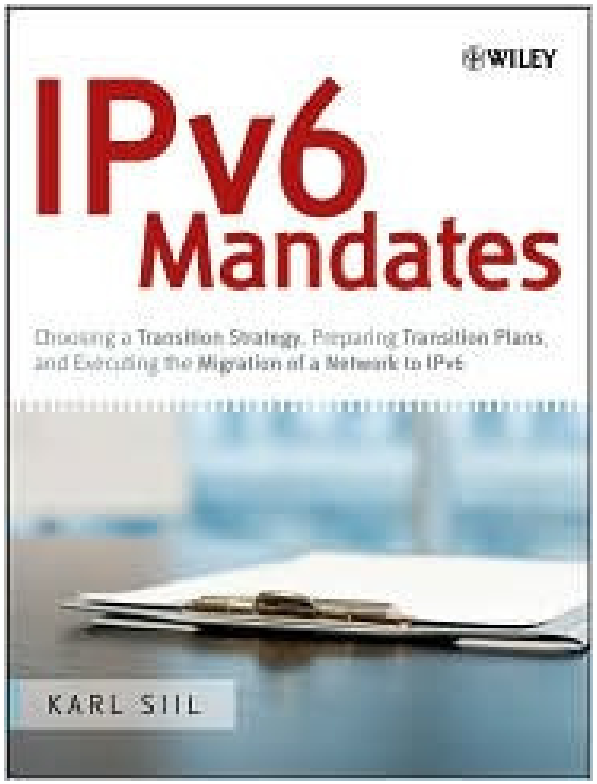

Yet another IPv6 Book

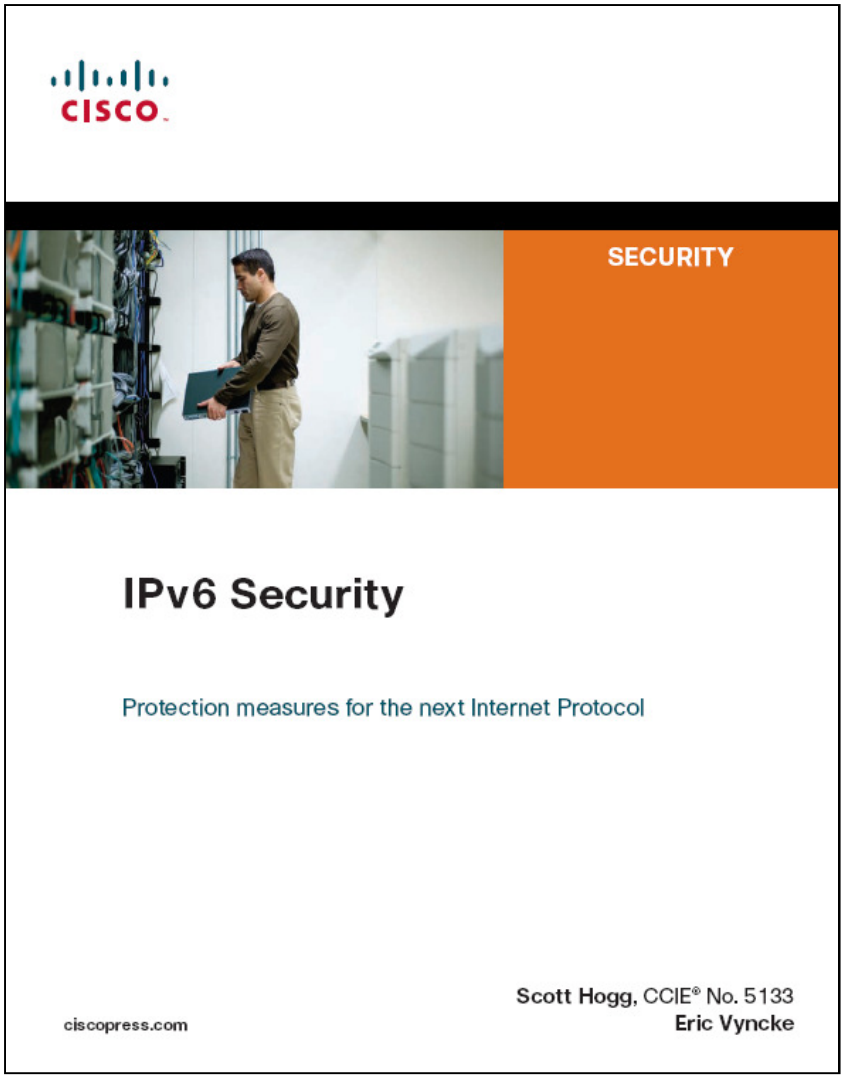

New Book for Enterprises

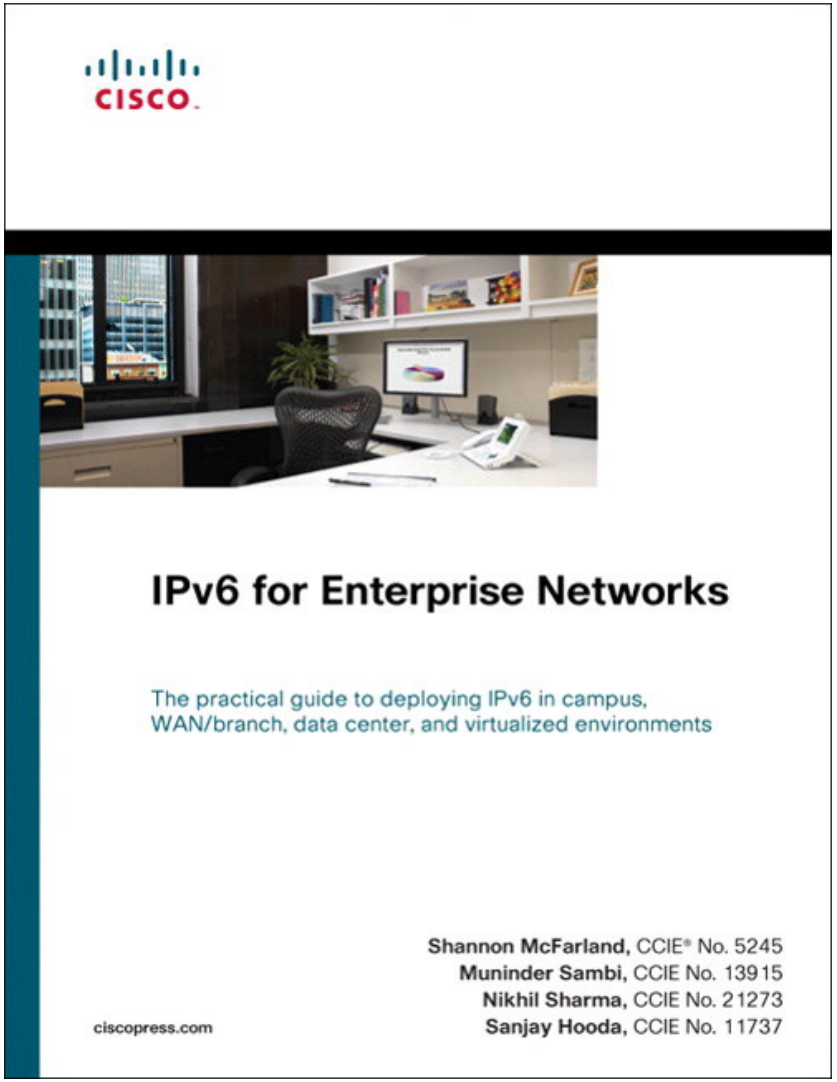

Rocky Mountain IPv6 Task Force

- Regional “chapter” of North American IPv6 Task Force and, therefore, the IPv6 Forum
- Our Charter
  - Provide Education on IPv6 and its benefits
  - Promotion of IPv6 technology
  - Research and Development and showcase IPv6 technology and services
  - Put on local IPv6-focused events
  - Work to further the use of IPv6 with a regional focus
- 4th Annual 2011 Rocky Mountain IPv6 Summit
  - April 25-27, 2011 – Denver
  - Download presentations from first 3 years of events
  - www.RMv6TF.org

NetworkWorld Blog

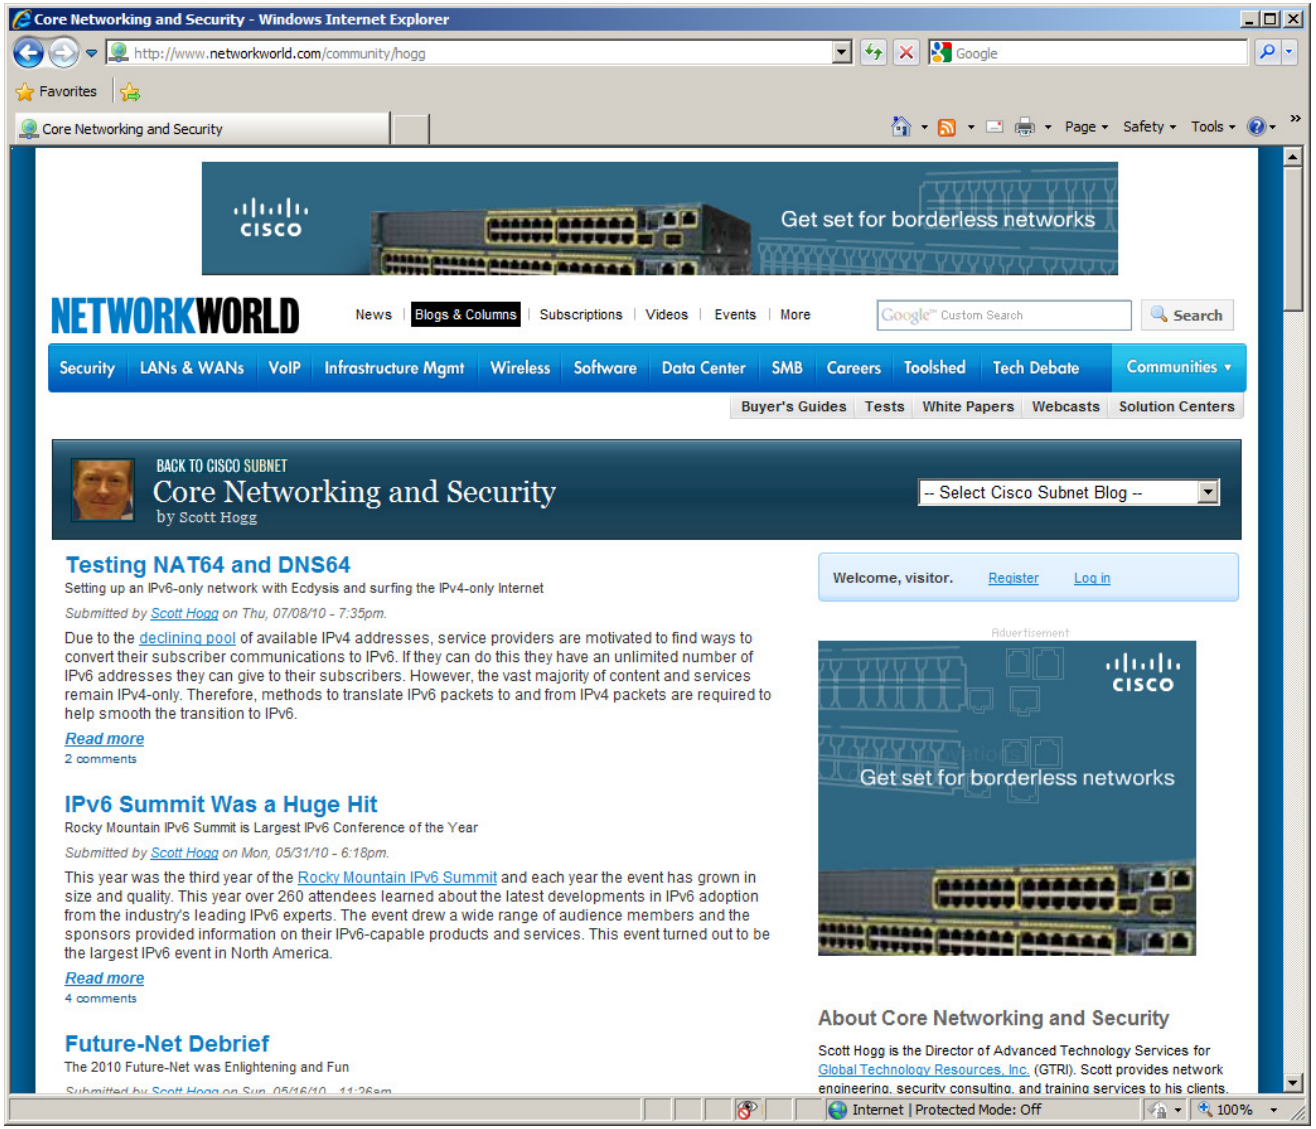

GTRI's IPv6 Transition Services

- IPv6 Inventory
  - Documentation of your current inventory and determination of IPv6 compatibility
  - Data gathering expertise (manual, data calls, automated utilities)
  - Cisco and GTRI automated tools
  - Inventory data aggregation and review
- IPv6 Training
  - Education for your teams to help them learn IPv6 technologies
  - Classroom and hands-on training
- IPv6 Impact Analysis
  - IPv6 Risk Assessment using OMB's own Risk Analysis Methodology
  - Custom-tailored transition planning for your IPv6 migration, tied to your enterprise architecture
- IPv6 Application Assessment
  - Software assessments leveraging COTS tools and our extensive experience
  - Review of your operating system constraints for IPv6 adoption
- IPv6 Experimentation and Testing
  - Systems testing in our IPv6 lab (DNS, routing, security, applications)
- IPv6 Deployment
  - Deployment of dual-stack and other IPv6 transition techniques
  - Dual Stack DNS servers and IPv6 security deployment

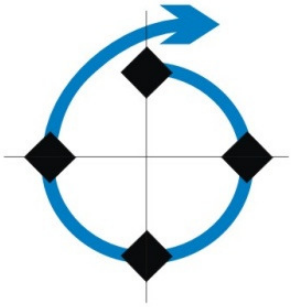

Question and Answer

Q:

&

A:

SHogg@GTRI.com
Scott@HoggNet.com

Mobile: 303-949-4865