Advanced IPv6 Deployment & Services

BRKRST-3305

Prerequisites: Session Abstract

- This session will cover how an ISP can deploy IPv6, how an ISP provide IPv6 connectivity to its customers while still running IPv4. IPv6 service should be added to an existing IPv4 network without any interruption of V4 services. We will look at current SP topologies and protocols and evaluate best methodologies for introducing IPv6. We will evaluate existing transition mechanisms in the context of existing v4 deployment scenarios. Finally we will discuss MPLS based networks pure IP network deployments, and in that context discuss different protocols when deploying dual stack. Session will cover OSPFv3, ISIS, BGP architectural consideration when deploying IPV6.
- Attendee must have a solid foundation of IPv6 basics (addressing, routing), MPLS, IPv4 networks and provisioning

Agenda

- SP Architecture
  - Pure IP Networks
  - MPLS networks
- Enterprise Architecture
- Address Allocation in SP & Enterprise
- Routing Deployment – IGP & BGP
- Routing Protocols Co-existence & Convergence

SP Architecture

Pure IP Networks

ISP Deployment Activities

- Several Market segments
  - IX, Carriers, Regional ISP, Wireless
- ISP have to get an IPv6 prefix from their Regional Registry
  - http://www.arin.net
- Large carriers are running trial networks but
  - Plans are largely driven by customer’s demand
- Regional ISP focus on their specific markets
  - Japan is leading the worldwide deployment
  - Target is Home Networking services (dial, DSL, Cable, Ethernet-to-the-Home,…)
- No easy Return on Investment (RoI) computation

A Today's Network Infrastructure

- Service Providers core infrastructure are basically following 2 paths.
  - MPLS with its associated services
    - MPLS/VPN, L2 services over MPLS, TE, QoS,…
  - Native IPv4 core with associated services
    - L2TPv3, QoS, Multicast,…
- IP services portfolio
  - Enterprise: Lease Lines
  - Home Users/SOHO: ADSL, ETTH, Dial
  - Data Center: Web hosting, servers,…
- Next – The Integration of IPv6 services

Service Provider networks

- Major routing information is ~320K via BGP
- Largest known IGP routing table is ~6–7K
- Total of 327K
- 6K/327K ~ 2% of IGP routes in an ISP network
- A very small factor but has a huge impact on network convergence!

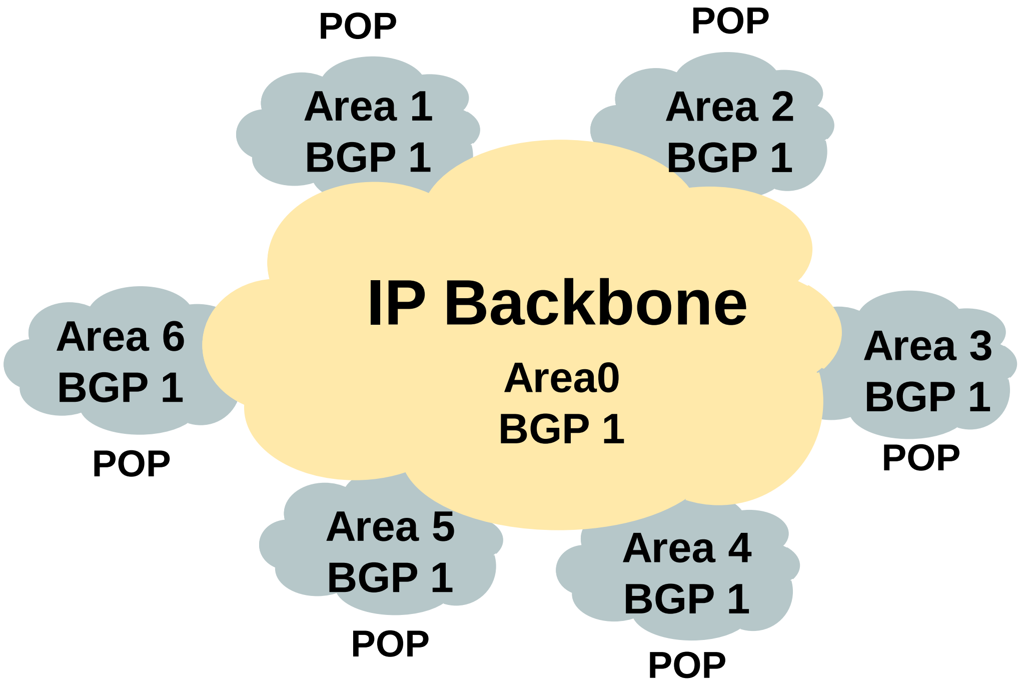

Service Provider networks

- You can reduce the IGP size to approx the number of exit routers in your network
- This will bring really fast convergence
- Optimized where you must and summarize where you can
- Stops unnecessary flapping

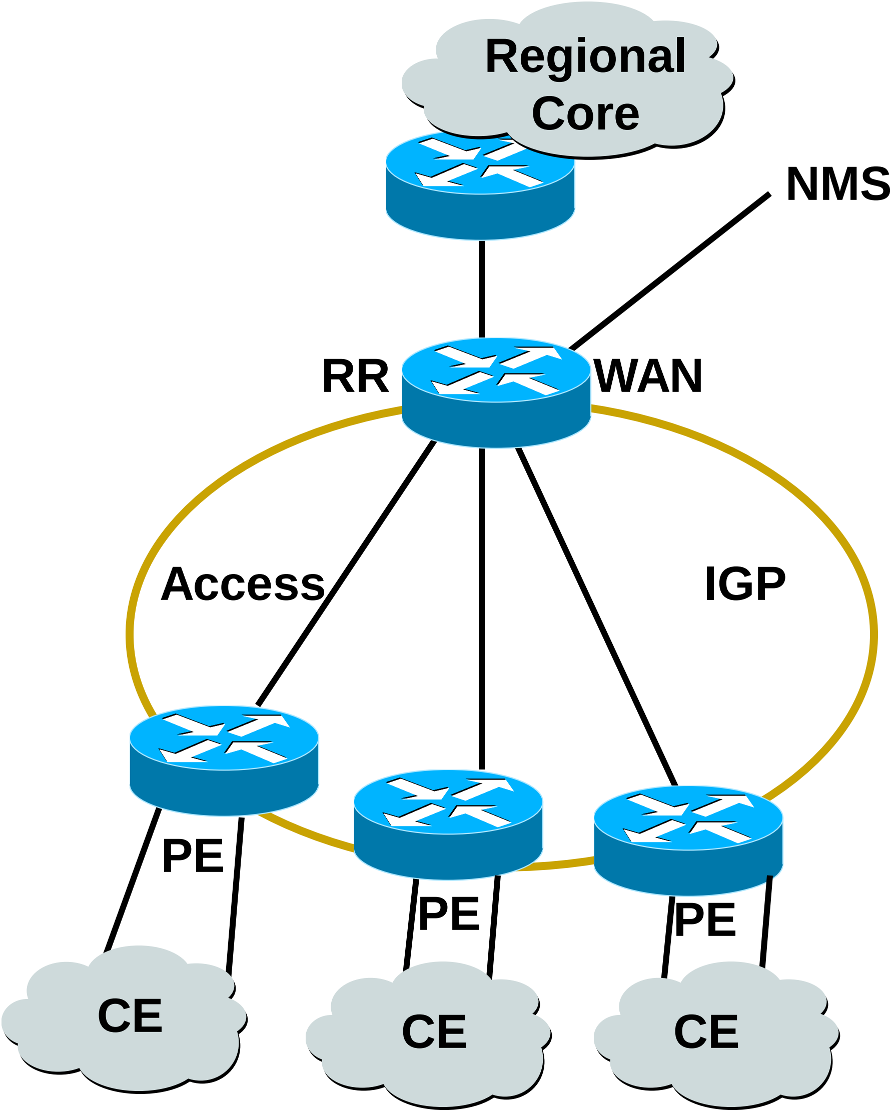

Addressing

- The link between PE-CE needs to be known for management purpose
- BGP next-hop-self should be done on all access routers—unless PE-CE are on shared media (rare case)
- This will cut down the size of the IGP
- For PE-CE link do redistributed connected in BGP
- These connected subnets should ONLY be sent through RR to NMS for management purpose; this can be done through BGP communities

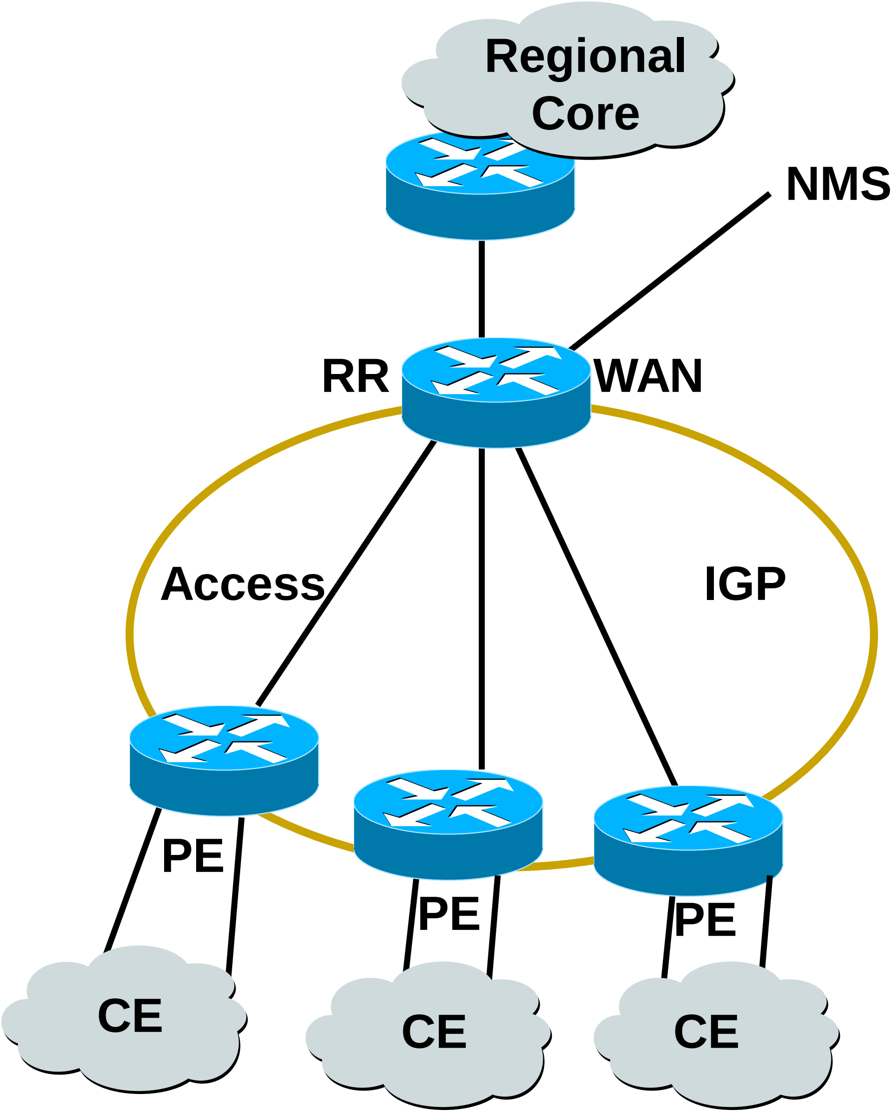

Addressing

- Divide the address into two parts

  1. Physical links
  2. Loopback interfaces

- Physical address should be in a contagious block
- Loopback should be from public address space
- Optimal path to the next hop is necessary

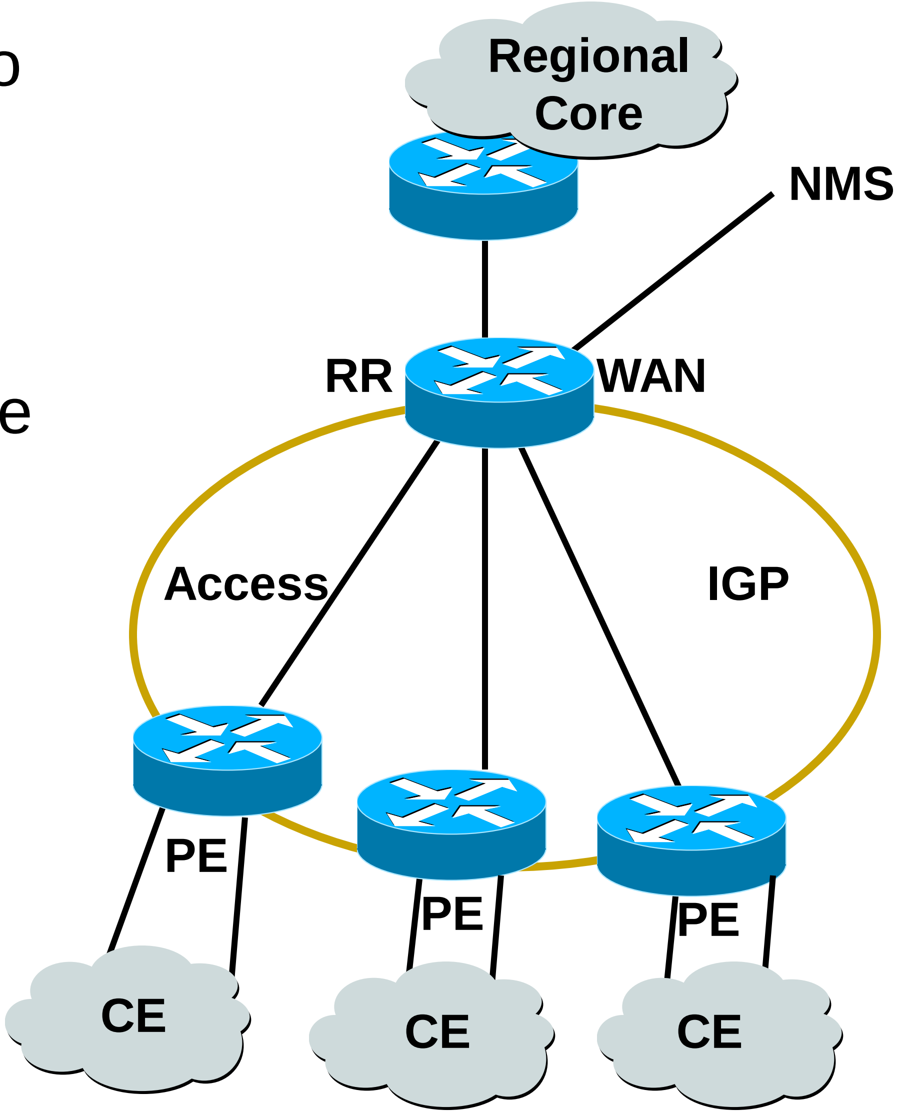

Addressing

- Assign ::/56 per pop for physical links
- Once out grow add another contiguous ::/56
- When assigning address to another POP keep few contiguous address open
- Summarize pop address at the WAN routers
- Leak loopback as specific
- Current trend within ISP's, are public address for loopback and public or private for infrastructure

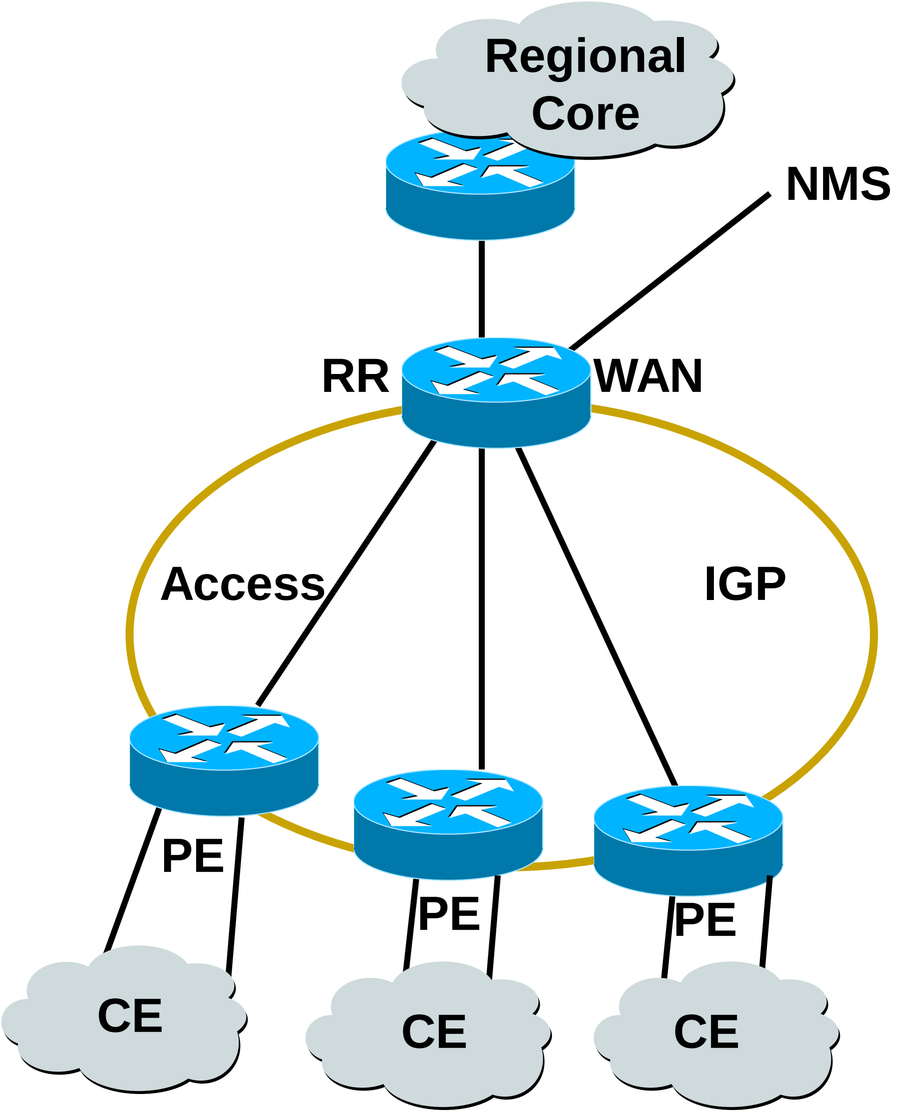

SP Architecture

MPLS Networkers

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 Provider Edge Router (6PE) over MPLS

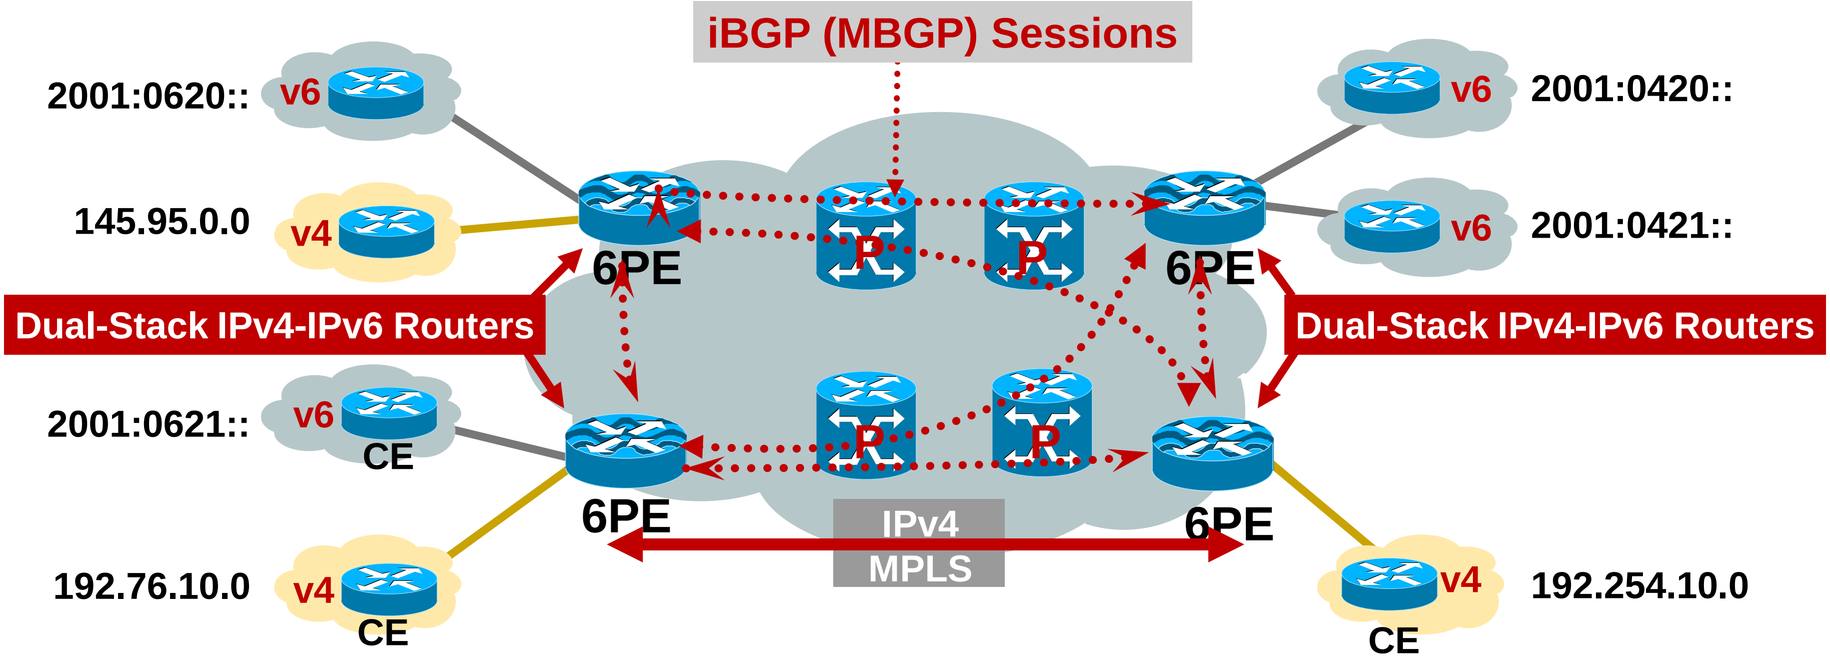

- IPv4 or MPLS core infrastructure is IPv6-unaware
- PEs are updated to support dual stack/6PE
- IPv6 reachability exchanged among 6PEs via iBGP (MBGP)
- IPv6 packets transported from 6PE to 6PE inside MPLS

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

6PE Routing/Label Distribution

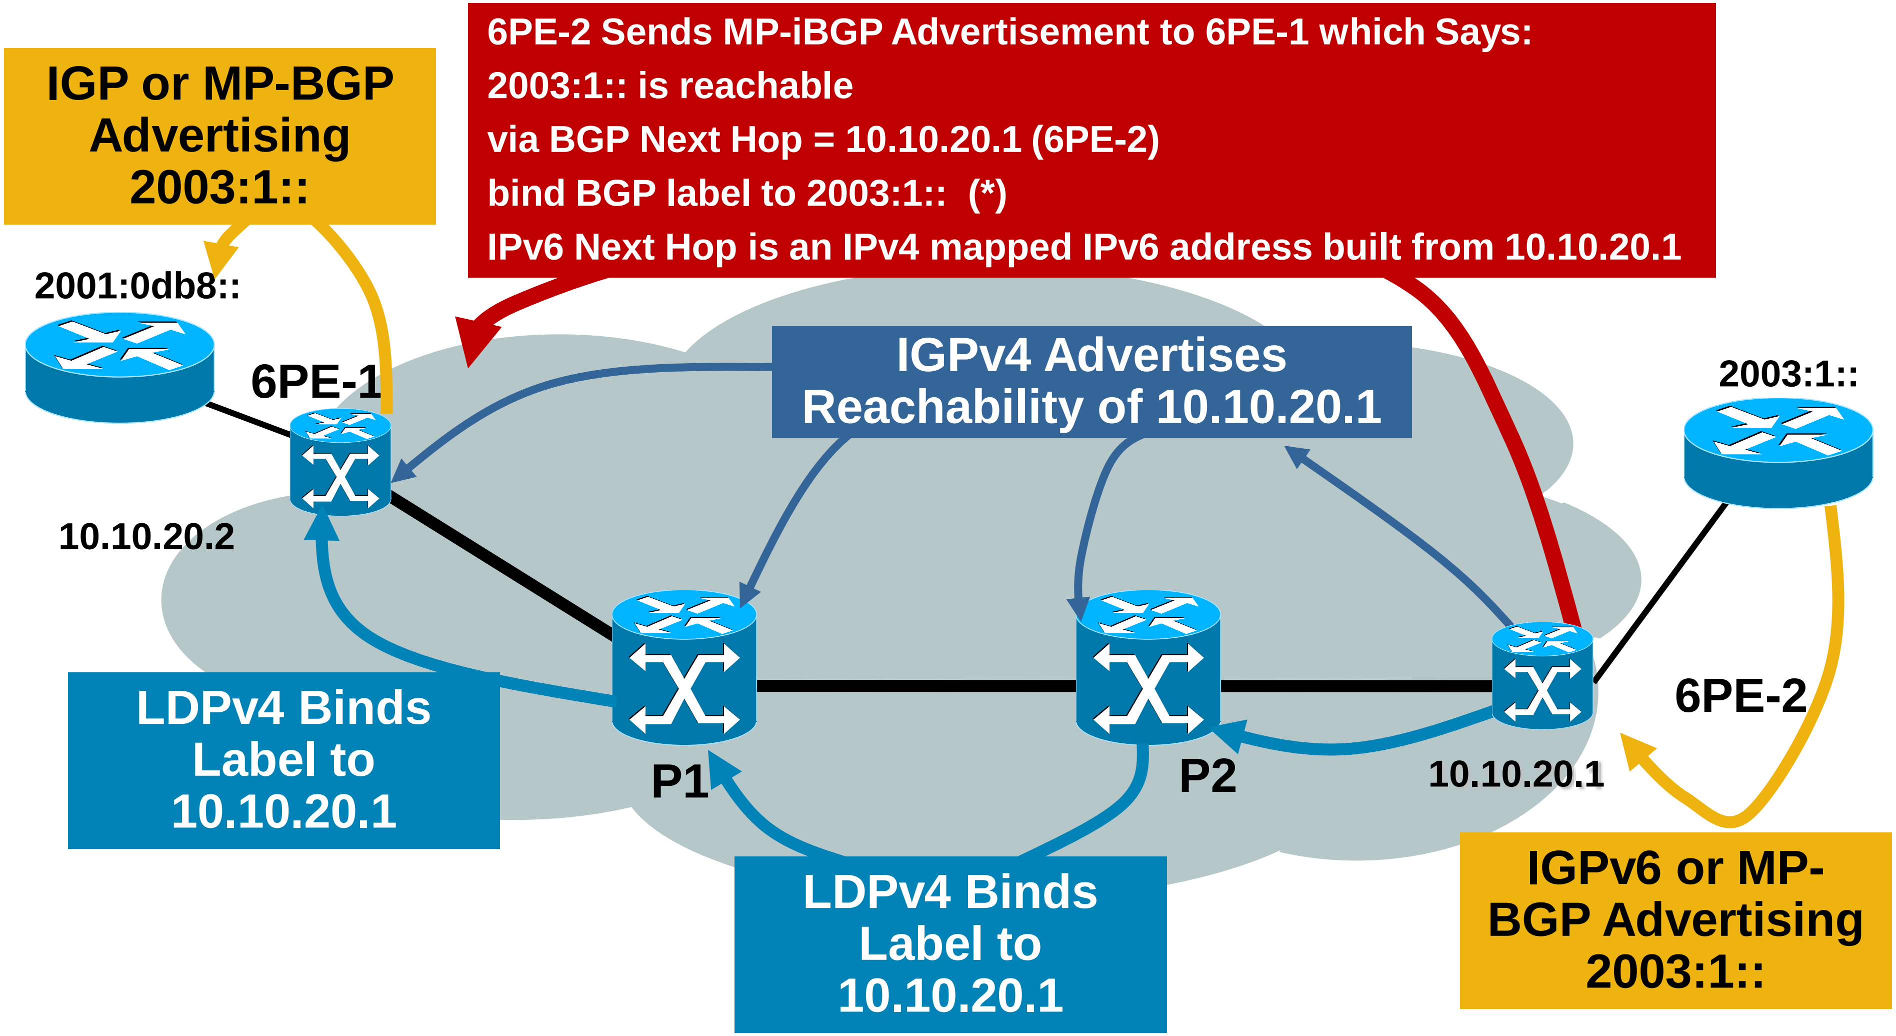

6PE-2 Sends MP-iBGP Advertisement to 6PE-1 which Says:
2003:1:: is reachable
via BGP Next Hop = 10.10.20.1 (6PE-2)
bind BGP label to 2003:1:: (*)
IPv6 Next Hop is an IPv4 mapped IPv6 address built from 10.10.20.1

IGP or MP-BGP Advertising 2003:1::

IGPv4 Advertises Reachability of 10.10.20.1

LDPv4 Binds Label to 10.10.20.1

LDPv4 Binds Label to 10.10.20.1

IGPv6 or MP-BGP Advertising 2003:1::

6PE Configuration

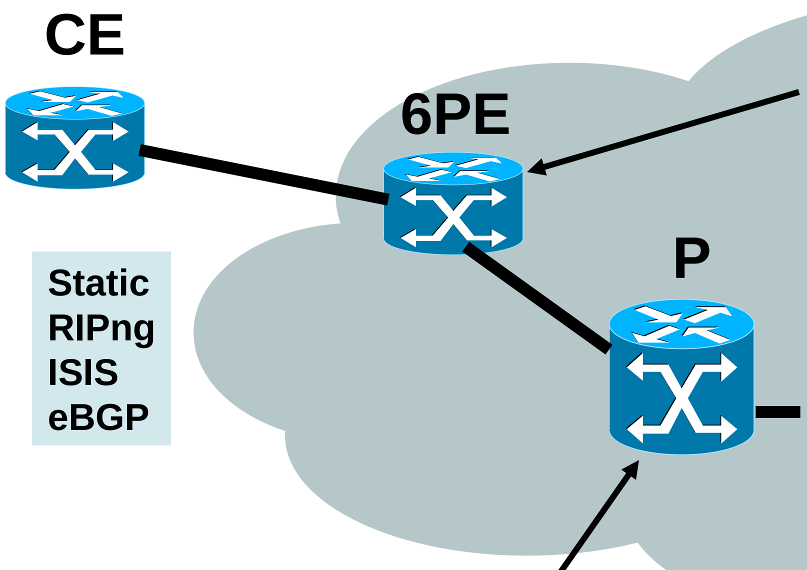

Static
RIPng
ISIS
eBGP

```
ipv6 cef
mpls label protocol ldp
mpls ldp router-id loopback0
!
interface Loopback0
 ip address 10.10.20.2 255.255.255.255
 ipv6 address 2003::/64 eui-64
!
router bgp 100
 no synchronization
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 10.10.20.1 remote-as 100
 neighbor 10.10.20.1 update-source Loopback0
 !
 address-family ipv6
 neighbor 10.10.20.1 activate
 neighbor 10.10.20.1 send-label
 redistribute connected
 redistribute rip ripv6CE1
exit-address-family
!
```

```
ip cef
mpls label protocol ldp
tag-switching tdp router-id loopback0
!
interface Serial2/0
 ip address 10.10.10.2 255.255.255.252
 ip router isis
 mpls label protocol ldp
 tag-switching ip
!
```

Note: send-label will cause flap on peer

Why Cisco IOS IPv6 VPN Provider Edge (6VPE)?

- For VPN customers, IPv6 VPN service is exactly the same as IPv4 VPN service
- Current 6PE is “like VPN” but this is NOT VPN, i.e., global reachability
- For ISP offering MPLS/VPN for IPv4 that wish to add IPv6 services as well

  No modification on the MPLS core

  Support both IPv4 and IPv6 VPNs concurrently on the same interfaces

  Configuration and operations of IPv6 VPNs exactly like IPv4 VPNs

6VPE Deployment

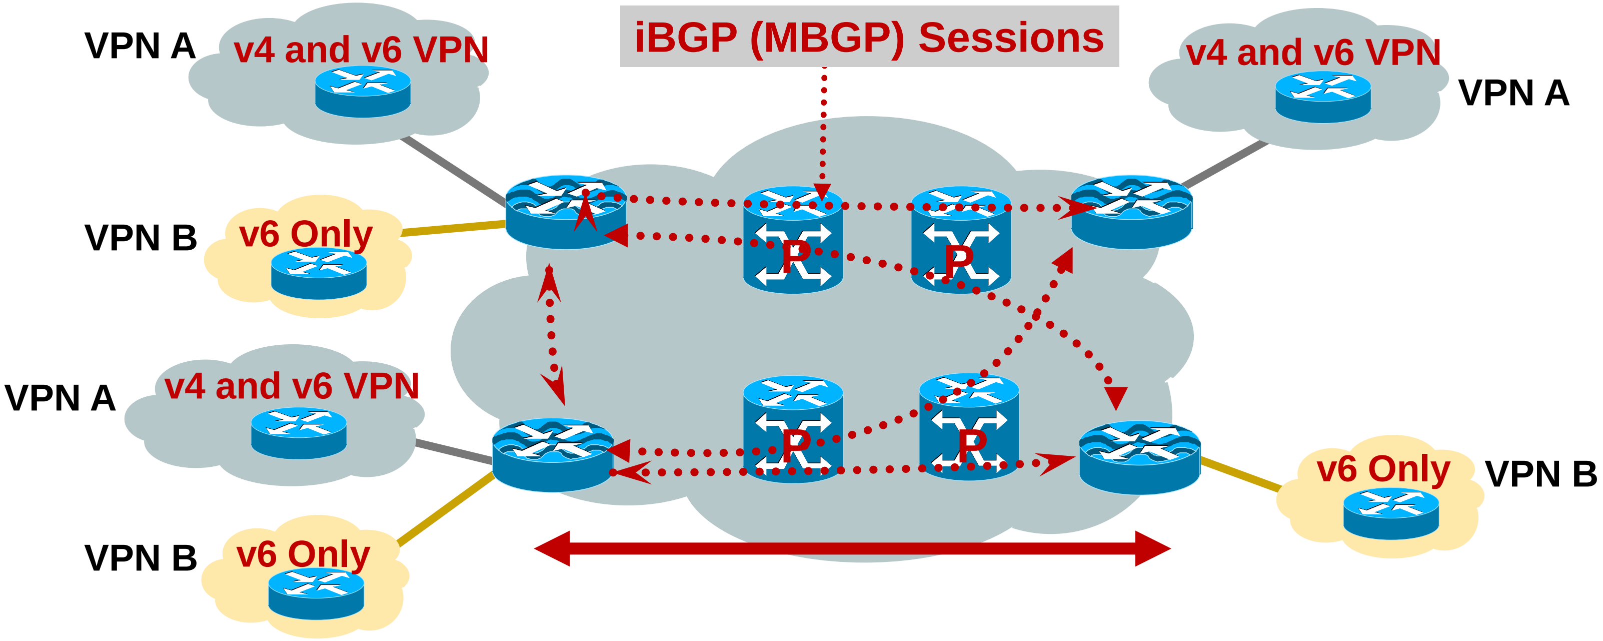

- IPv6 VPN can coexist with IPv4 VPN—same coverage
- 6VPE is added only when and where the service is required
- 6VPE—An implementation of <draft-ietf-bgp-ipv6-vpn> over MPLS/IPv4
- Standards work going forward—<draft-ietf-l3vpn-bgp-ipv6-xx.txt>

6VPE Configuration Example

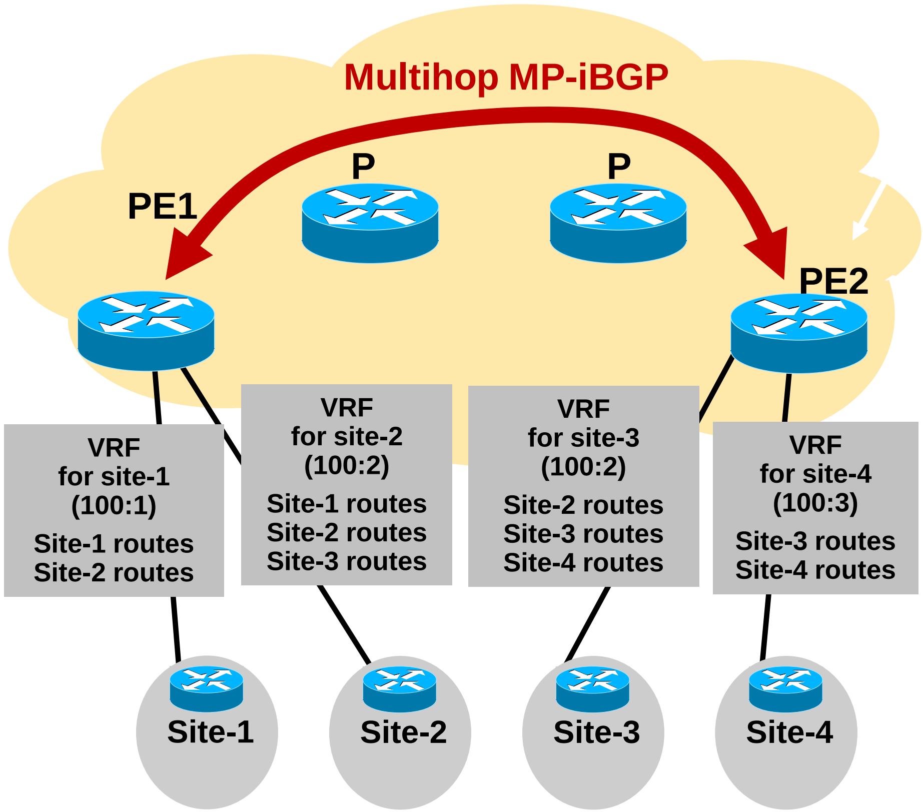

```
vrf definition SITE-3
 rd 100:2
 address-family ipv6
  route-target export 100:2
  route-target import 100:2
  route-target import 100:3
  route-target export 100:3
!
vrf definition SITE-4
 rd 100:3
 address-family ipv6
  route-target export 100:3
  route-target import 100:3
!
interface Serial4/6
 vrf forwarding SITE-3
 ipv6 address 2001:DB8:3::1/64
!
interface Serial4/7
 vrf forwarding SITE-4
 ipv6 address 2001:DB8:4::1/64
```

6VPE Configuration Example (Cont.)

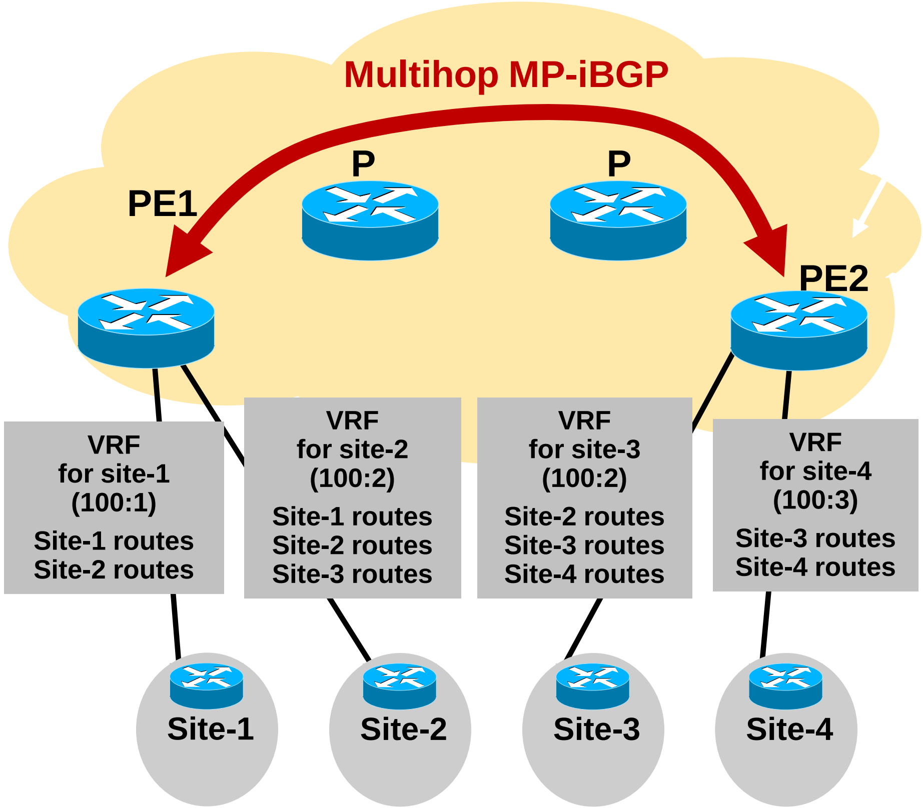

```
router bgp 100
 no bgp default ipv4-unicast
 neighbor 6.6.6.6 remote-as 100
 neighbor 6.6.6.6 update-source loopback0
!
address-family vpnv6
 neighbor 6.6.6.6 activate
 neighbor 6.6.6.6 send-community-extended
exit-address-family
!
address-family ipv6 vrf SITE-4
 neighbor 2001:DB8:4::2 remote-as 65504
 neighbor 2001:DB8:4::2 activate
exit-address-family
 !
address-family ipv6 vrf SITE-3
 neighbor 2001:DB8:3::2 remote-as 65503
 neighbor 2001:DB8:3:2 activate
exit-address-family
```

Enterprise Architecture

IPv6 Coexistence

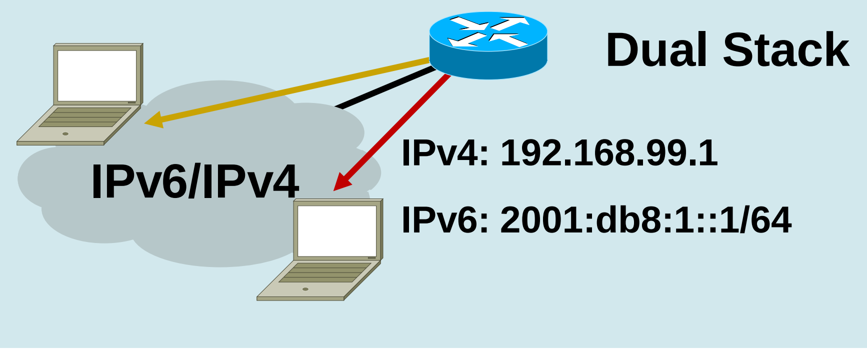

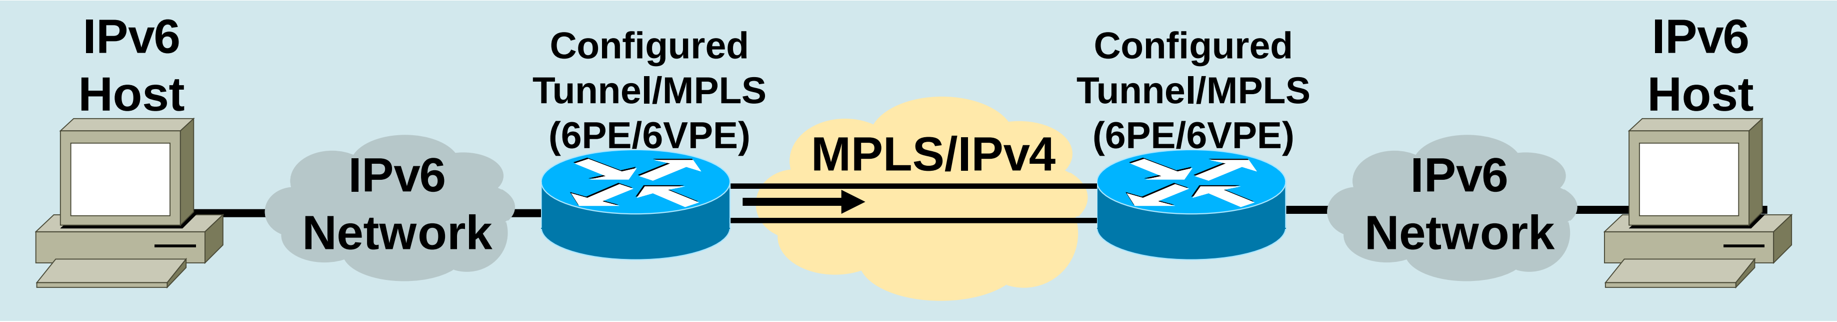

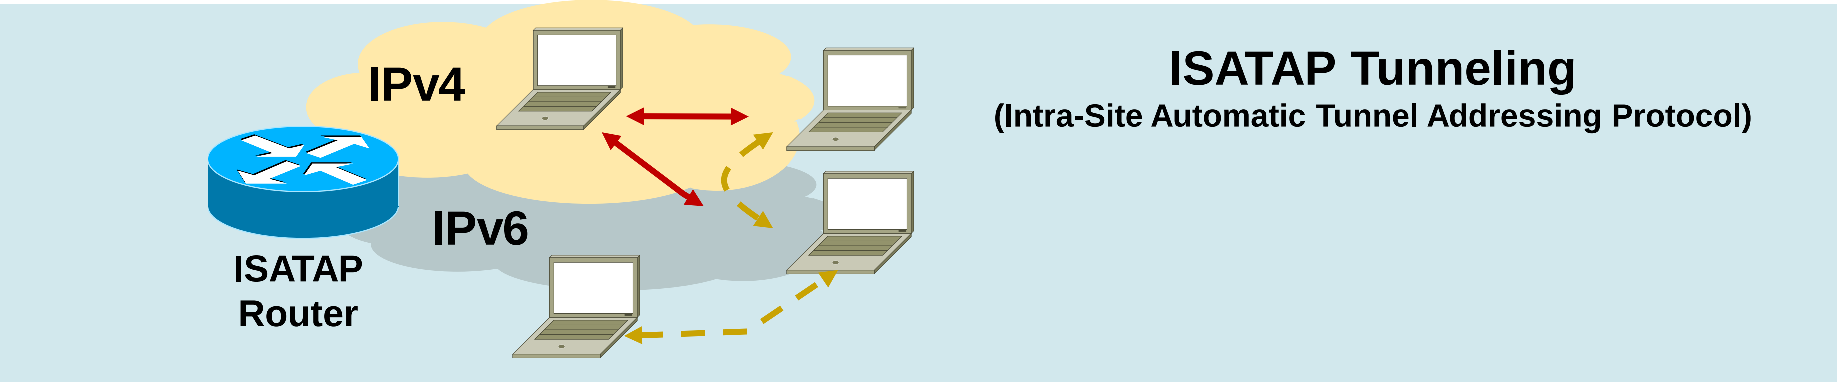

ISATAP Tunneling

(Intra-Site Automatic Tunnel Addressing Protocol)

Campus IPv6 Deployment
Three Major Options

- Dual-stack – The way to go for obvious reasons: performance, security, QoS, Multicast and management
  - Layer 3 switches should support IPv6 forwarding in hardware
- Hybrid – Dual-stack where possible, tunnels for the rest, but all leveraging the existing design/gear
  - Pro – Leverage existing gear and network design (traditional L2/L3 and Routed Access)
  - Con – Tunnels (especially ISATAP) cause unnatural things to be done to infrastructure (like Core acting as Access layer) and ISATAP does not support IPv6 multicast
- IPv6 Service Block – A new network block used for interim connectivity for IPv6 overlay network
  - Pro – Separation, control and flexibility (still supports traditional L2/L3 and Routed Access)
  - Con – Cost (more gear), does not fully leverage existing design, still have to plan for a real dual-stack deployment and ISATAP does not support IPv6 multicast

Campus IPv6 Deployment Options

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
  - Intelligent services on WLAN
- Requires robust control plane for both IPv4 and IPv6
  - Variety of routing protocols—The same ones in use today with IPv4
- Requires support for IPv6 multicast, QoS, infrastructure security, etc…

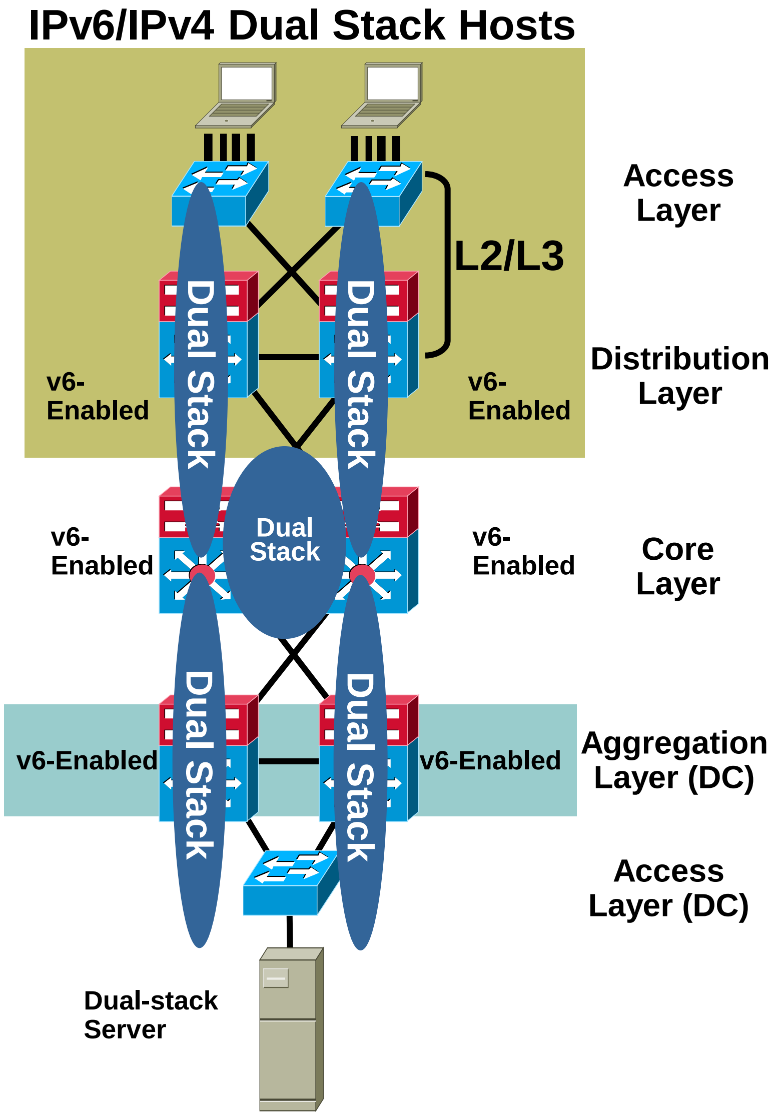

Campus IPv6 Deployment Options
Hybrid Model

- Offers IPv6 connectivity via multiple options
  - Dual-stack
  - Configured tunnels – L3-to-L3
  - ISATAP – Host-to-L3
- Leverages existing network
- Offers natural progression to full dual-stack design
- May require tunneling to less-than-optimal layers (i.e. Core layer)
- ISATAP creates a flat network (all hosts on same tunnel are peers)
  - Create tunnels per VLAN/subnet to keep same segregation as existing design (not clean today)
- Provides basic HA of ISATAP tunnels via old Anycast-RP idea
- ISATAP does not support IPv6 Multicast
- Configured tunnels do support IPv6 Multicast

Hybrid Model

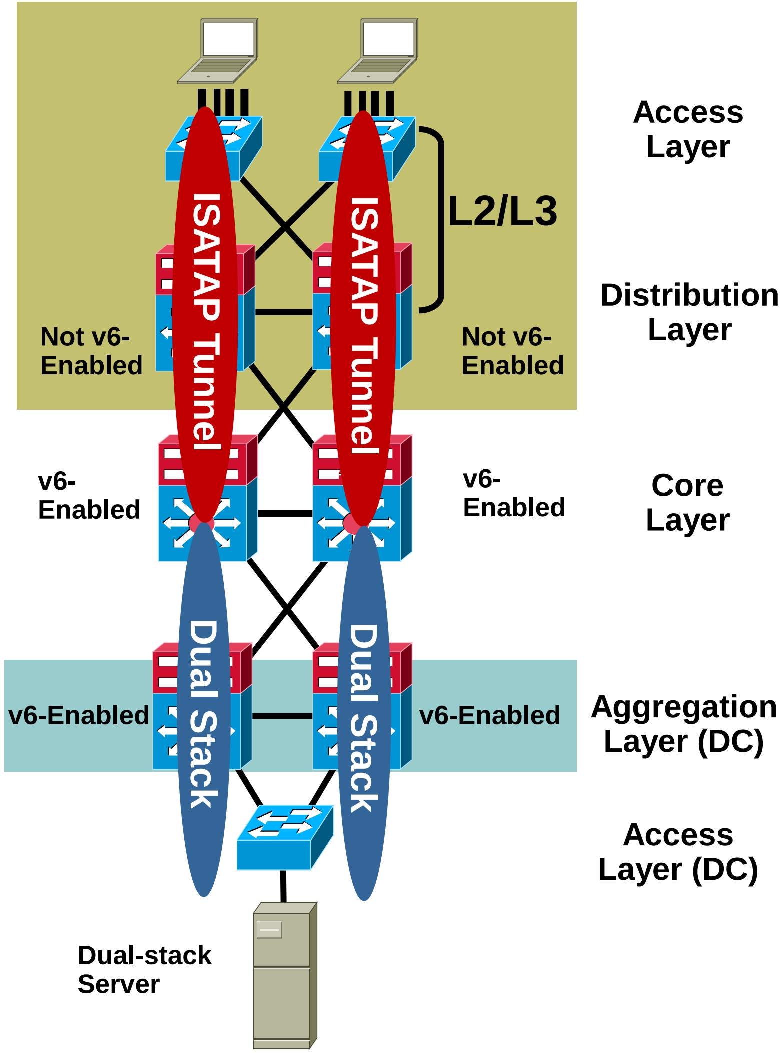

Hybrid Model Examples

Hybrid Model Example #1

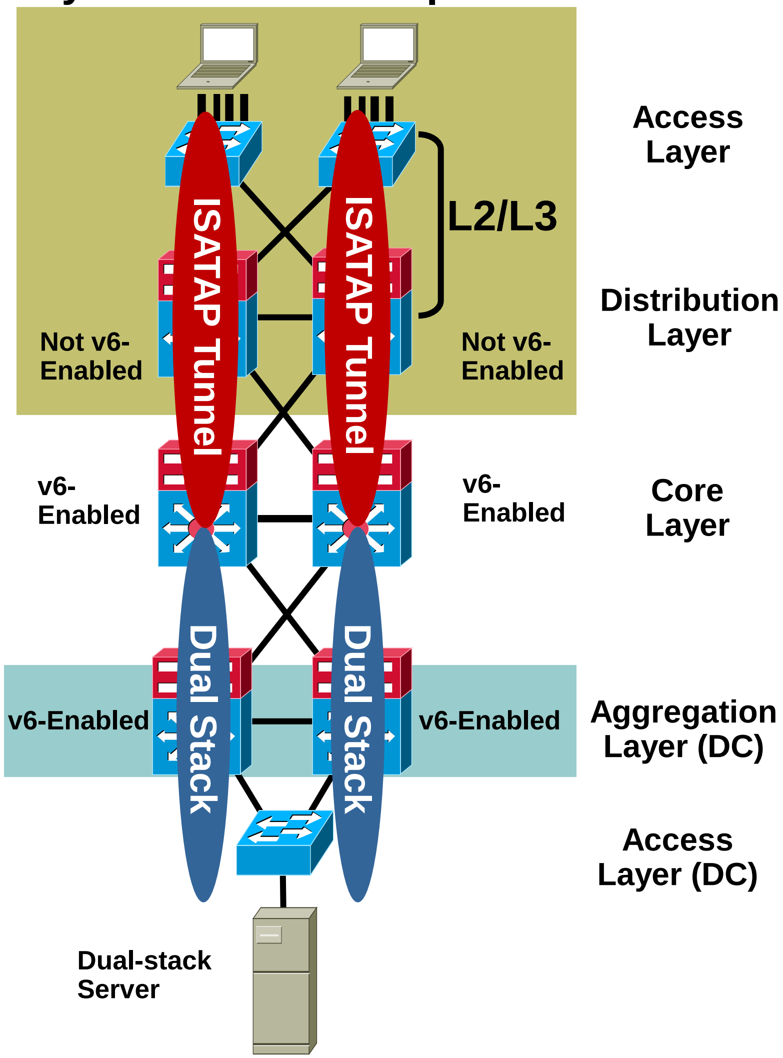

Hybrid Model Example #2

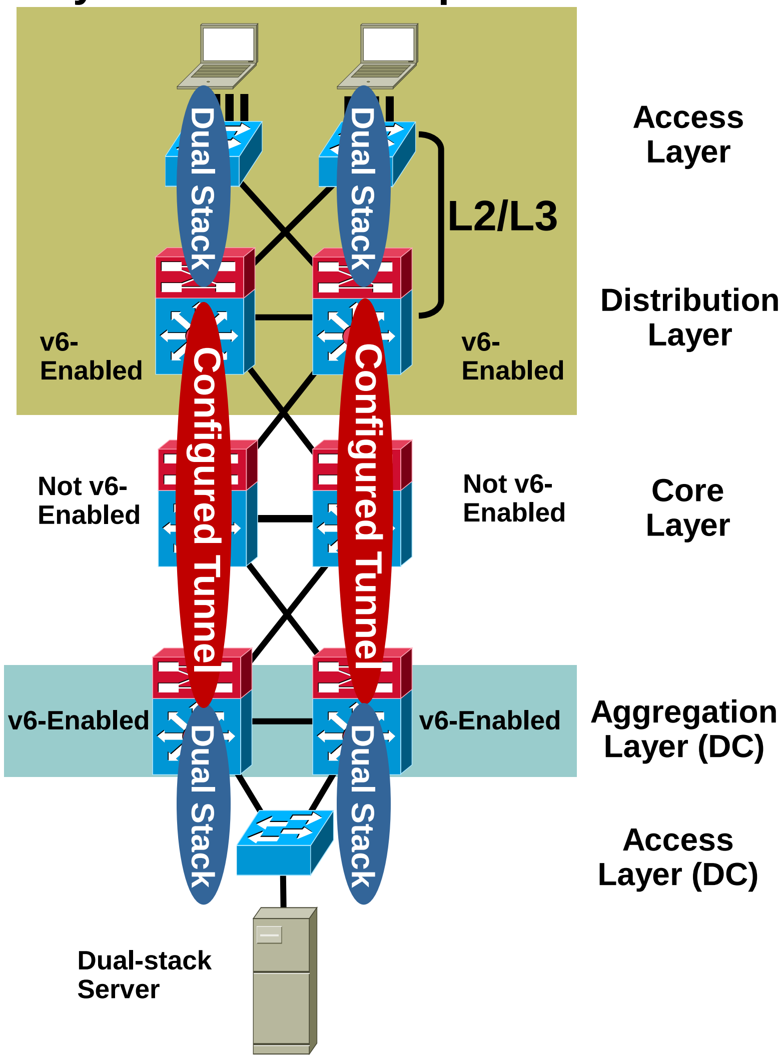

Highly Available ISATAP Design

Topology

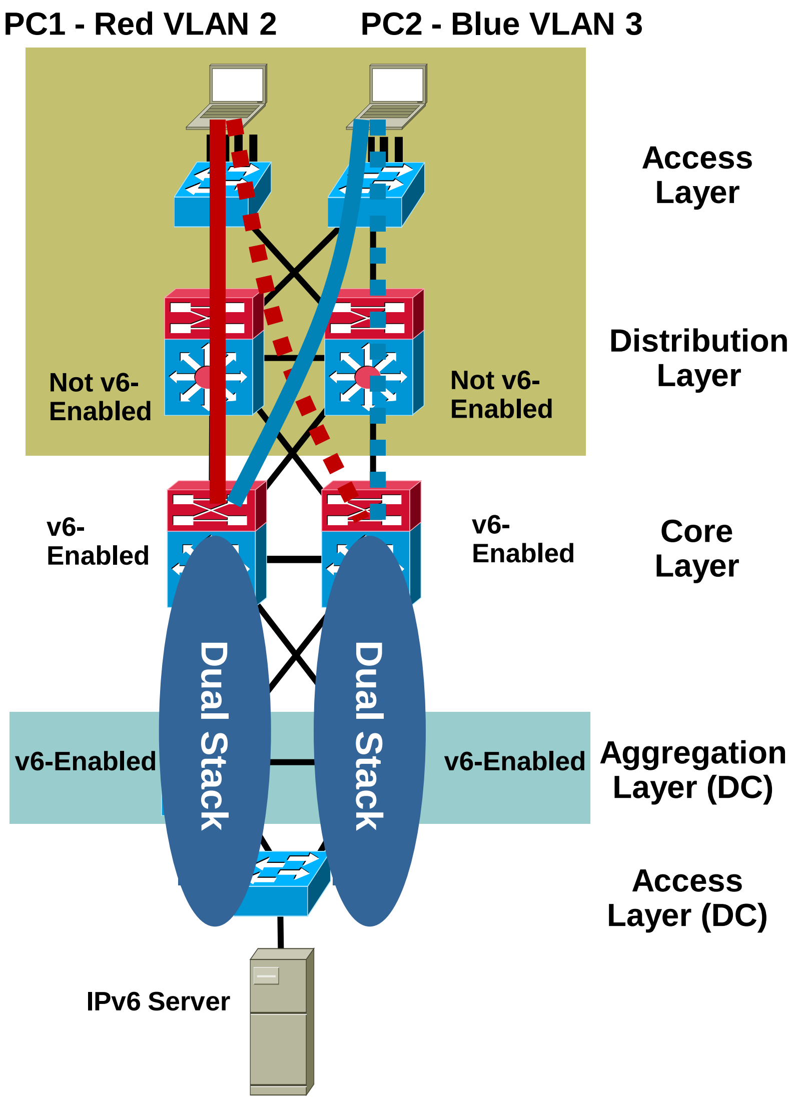

- ISATAP tunnels from PCs in Access layer to Core switches
- Redundant tunnels to Core or Service block
- Use IGP to prefer one Core switch over another (both v4 and v6 routes) - deterministic
- Preference is important due to the requirement to have traffic (IPv4/IPv6) route to the same interface (tunnel) where host is terminated on - Windows XP/2003
- In this example dual-stack is used from Data Center to Core

IPv6 Campus ISATAP Configuration

ISATAP Client Configuration

Windows XP/Vista Host

```
interface Tunnel3
 ipv6 address 2001:DB8:CAFE:3::/64 eui-64
 no ipv6 nd suppress-ra
 ipv6 ospf 1 area 2
 tunnel source Loopback3
 tunnel mode ipv6ip isatap
!
interface Loopback3
 description Tunnel source for ISATAP-VLAN3
 ip address 10.122.10.103 255.255.255.255
```

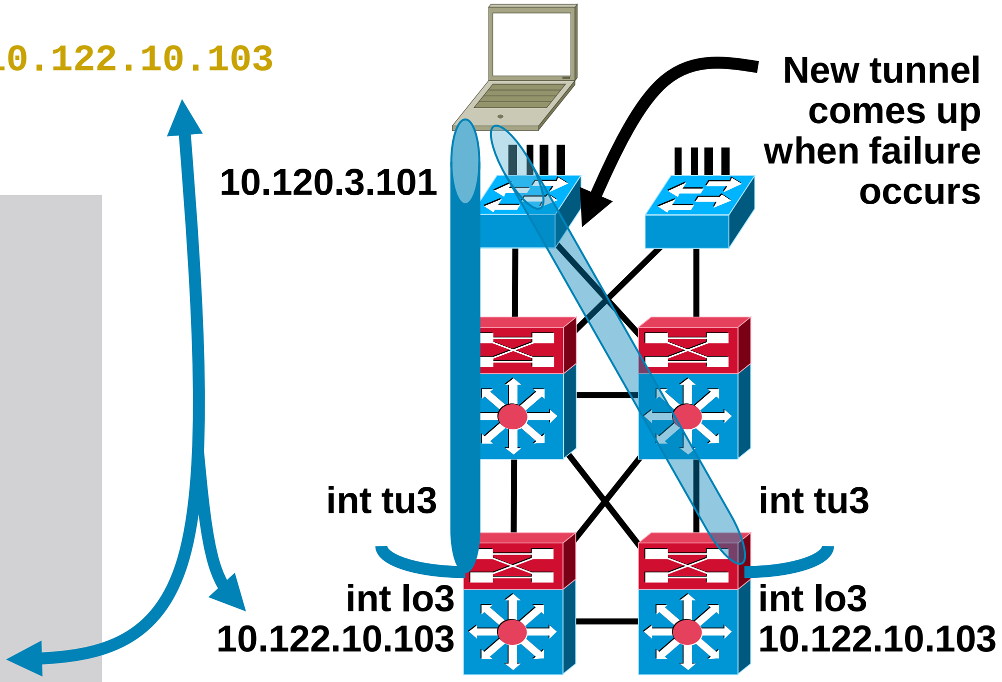

2001:db8:cafe:3:0:5efe:10.120.3.101

fe80::5efe:10.122.10.103%2

Campus IPv6 Deployment Options

IPv6 Service Block – An Interim Approach

- Provides ability to rapidly deploy IPv6 services without touching existing network
- Provides tight control of where IPv6 is deployed and where the traffic flows (maintain separation of groups/locations)
- Offers the same advantages as Hybrid Model without the alteration to existing code/configurations
- Configurations are very similar to the Hybrid Model
  - ISATAP tunnels from PCs in Access layer to Service Block switches (instead of core layer – Hybrid)
- 1) Leverage existing ISP block for both IPv4 and IPv6 access
- 2) Use dedicated ISP connection just for IPv6 – Can use IOS FW or PIX/ASA appliance

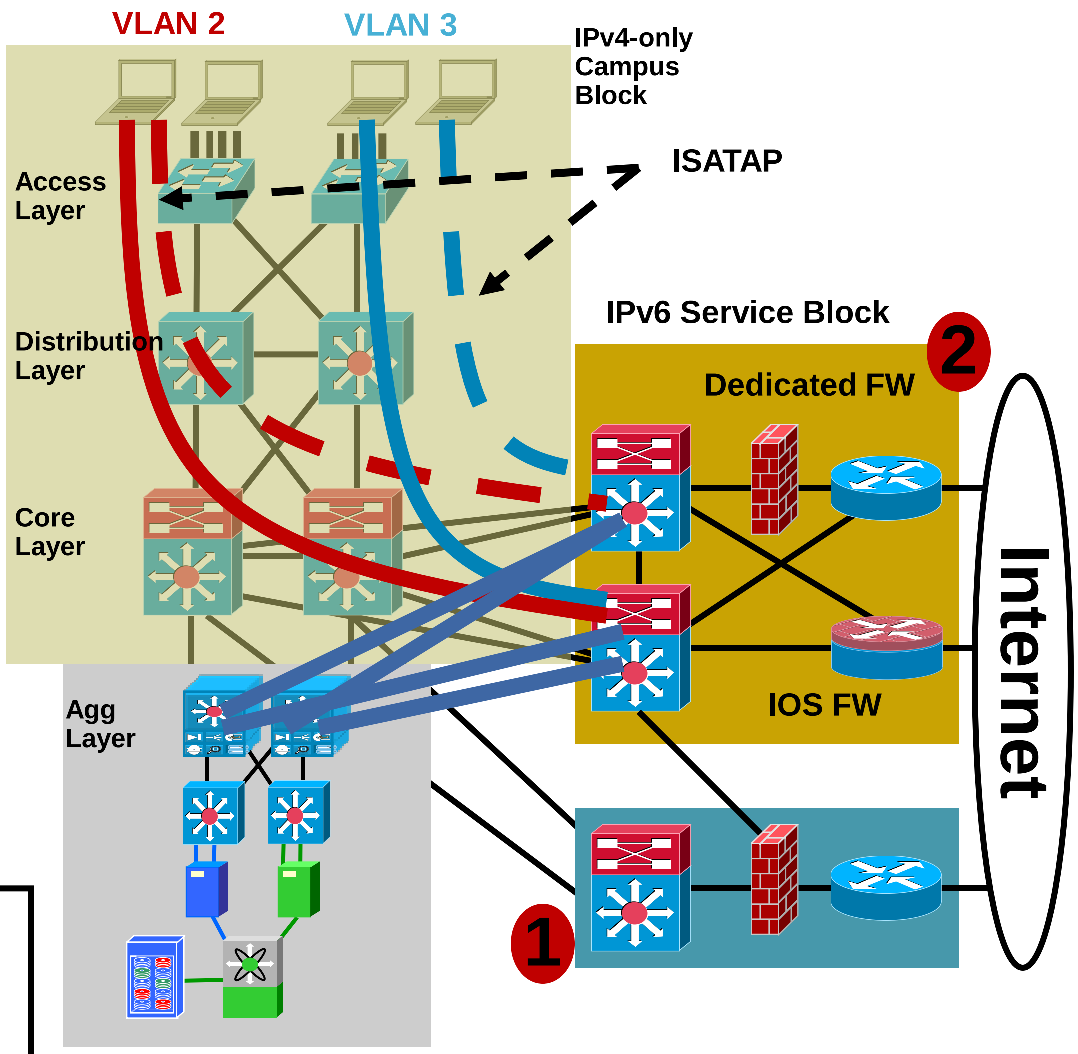

Address Allocations

SP & Enterprise

Allocation Recommendations

- IANA allocates from 2001::/16 or shorter to regional registries
- Each regional registry's allocation is a ::/23 or shorter
- ISP allocations from the regional registry is a ::/36 (immediate allocation) or ::/32 (initial allocation) or shorter with justification (Example: FT recently acquired a /19)
- The policy expectation is that an ISP allocates a ::/48 prefix to each customer, longer prefixes (but shorter than /64) for home users
- Link prefix length is no longer than /64 with the exception of point-to-point where /127 can be used (not encouraged)

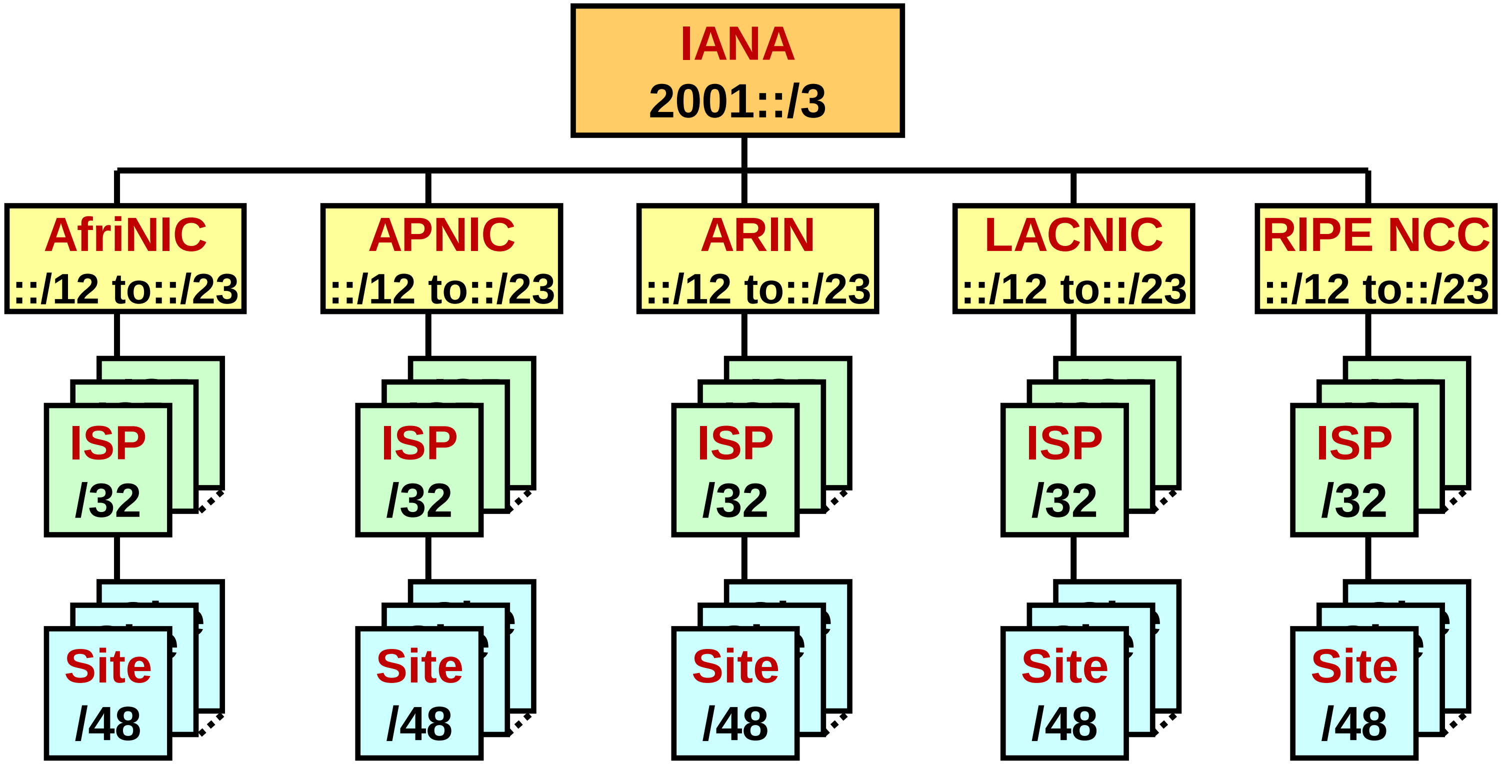

SP IPv6 Address Allocation

- SP addressing scheme

  Usually SP get the address allocated by the local registry via IANA

  The block is usually /32 but exception can be made for a bigger ISP

- SP usually assign addresses for Consumers. There are 2 types:
- Fixed allocation:

  Cable customers, DSL customers, ETTH etc

- Mobile allocation:

  Mobile customers

Cisco IOS IPv6 Broadband Access Solutions

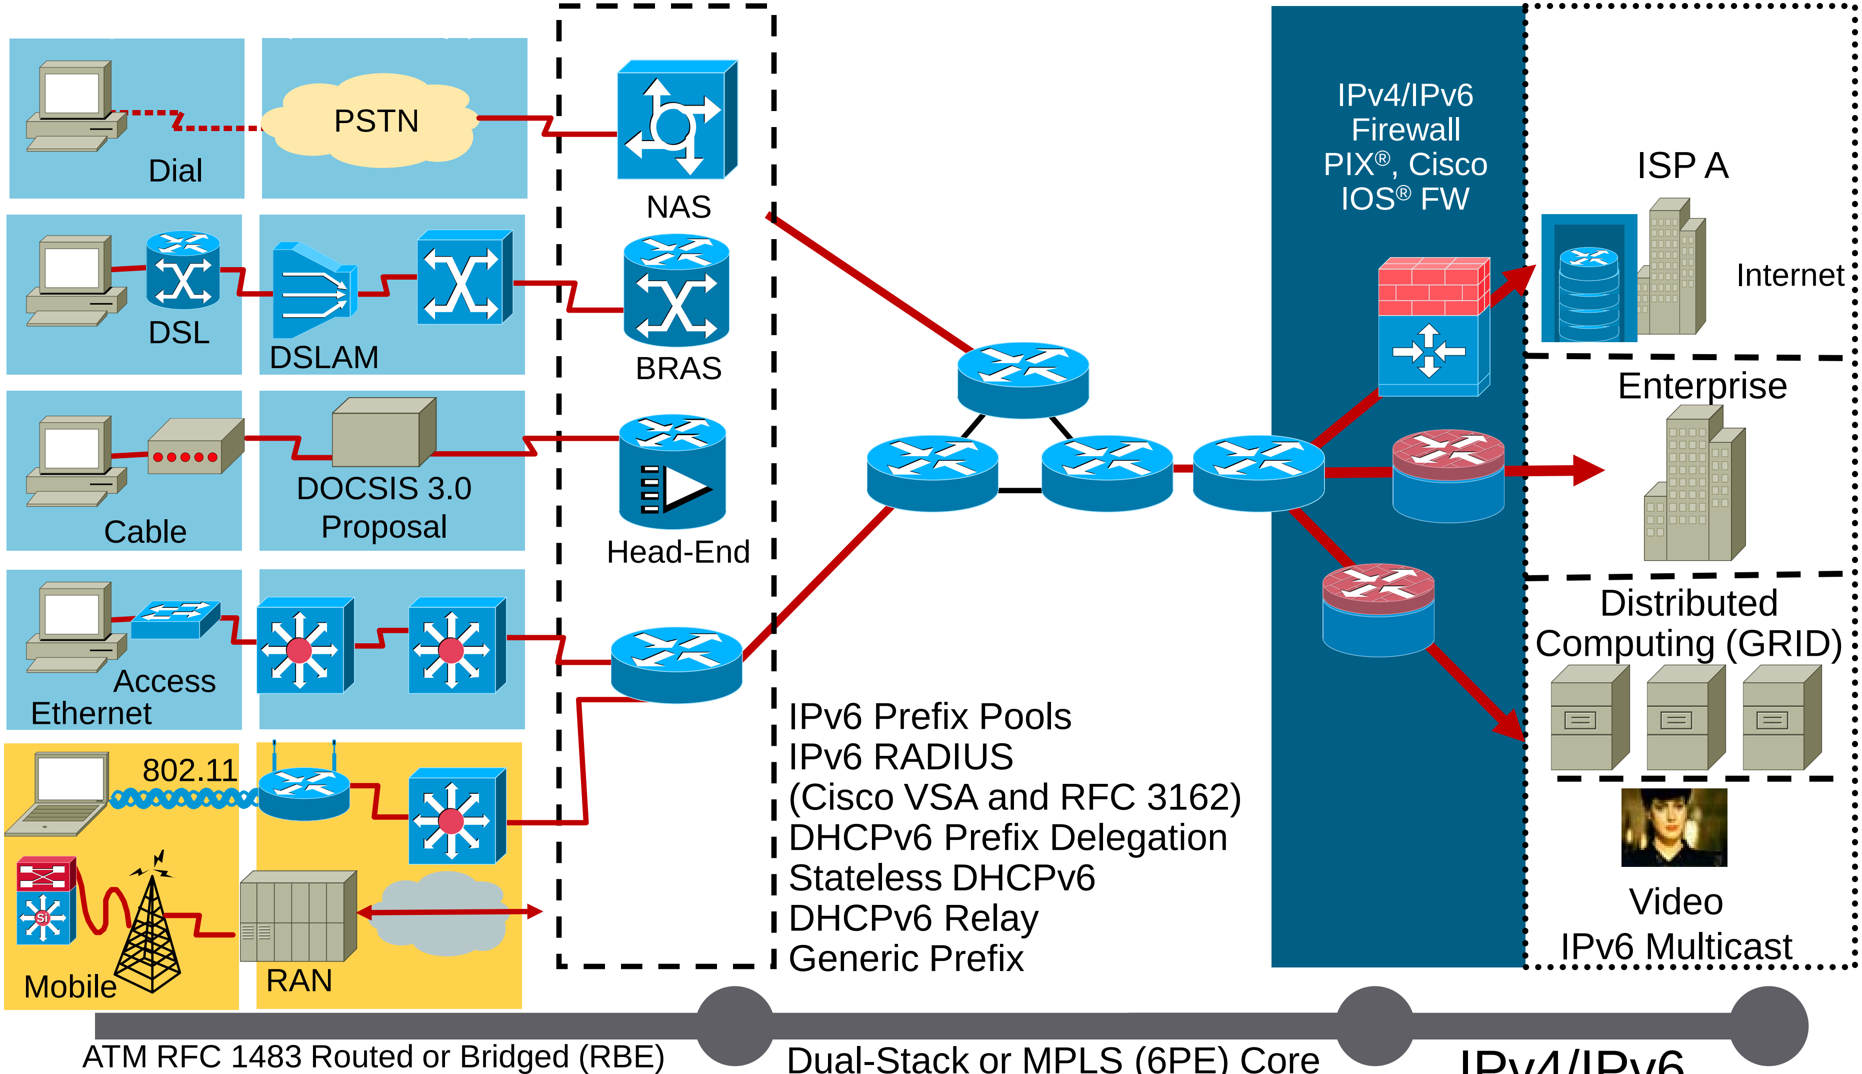

IPv6 Prefix Pools
IPv6 RADIUS
(Cisco VSA and RFC 3162)
DHCPv6 Prefix Delegation
Stateless DHCPv6
DHCPv6 Relay
Generic Prefix

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

Dual-Stack or MPLS (6PE) Core

IPv4/IPv6

IPv6 prefix-pools

- Normal prefix pools:
  ipv6 prefix-pool foo 3ffe:c00:1::/48 64

  A Separate /64 is assigned each user/interface. The prefix is advertised in RA's and a route is installed in the RIB.

- Shared prefix pools:
  ipv6 prefix-pool foo 3ffe:c00:2::/64 128 shared

  /64 prefix is shared between all users of the pool. The same /64 prefix is advertised in RA's out all interfaces. The user gets an /128 based on the prefix and his Interface-Identifier. A route in the RIB is installed only for the /128.

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

Policy Implementation

- Give Home/SOHO a permanent /64 – single link
- Give Home/SOHO a permanent /48
- Short-lived /64 from a prefix-pool

  A Separate /64 is assigned to each user/interface. The prefix is advertised in RA’s and a route is installed in the RIB.

- Short-lived /128 from a shared prefix-pool

  /64 prefix is shared between all users of the pool. The same /64 prefix is advertised in RA’s out all interfaces. The user gets an /128 based on the prefix and his Interface-Identifier. A route in the RIB is installed only for the /128.

- For some users set the Interface-ID explicitly

Give home users a permanent /64 – single link

- Use: for single PC or network with only one link
- AAA static prefix attribute. Interface-Id attribute to specify the complete address
- CPE: single PC, proxy RA, or configured router

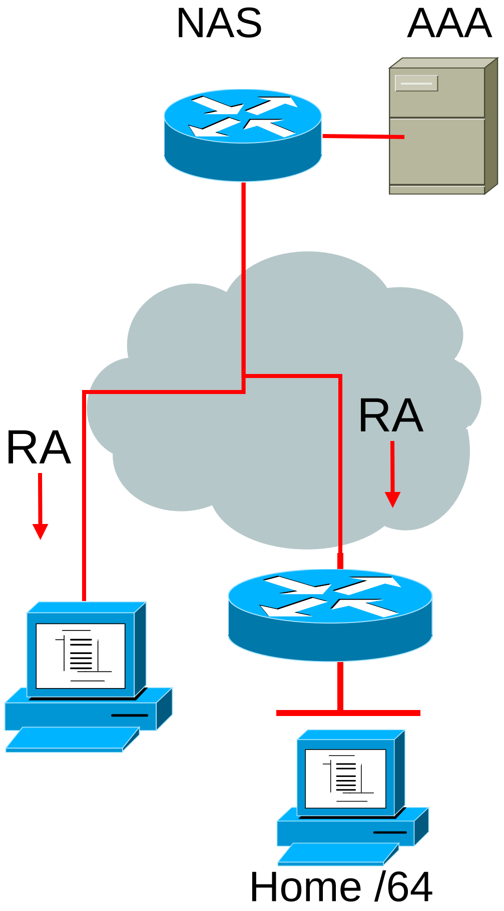

AAA config:
Auth-Type = Local, Password = “foo”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “ipv6:prefix=3ffe:c00::/64
Framed-Interface-Id = 0:0:0:1

Give home users a permanent /48

- Use: whole site -supports multiple links
- AAA prefix-attribute
- Use DHCP-PD to configure the CPE

```
interface Atm 0
pvc 1/23
 encapsulation aal5mux ppp dialer
dialer pool-member 1
!
interface dialer1
 ipv6 dhcp client pd DH-PREFIX
!
interface FastEthernet0
   ipv6 address DH-PREFIX 0:0:0:1::/64 eui-64
!
```

Auth-Type = Local, Password = “foo2”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “ipv6:prefix=3ffe:c00::/64

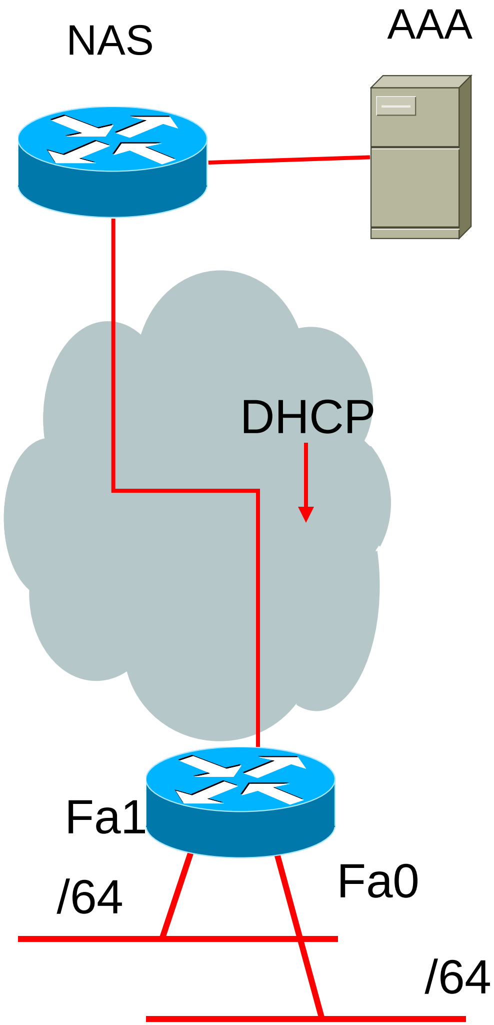

Address Assignment – short-lived /64

- Use: for single PC or very simple network
- NAS: IPv6 prefix pool
- CPE: Proxy-RA/multi-link subnet/bridging
  Renumbering issues

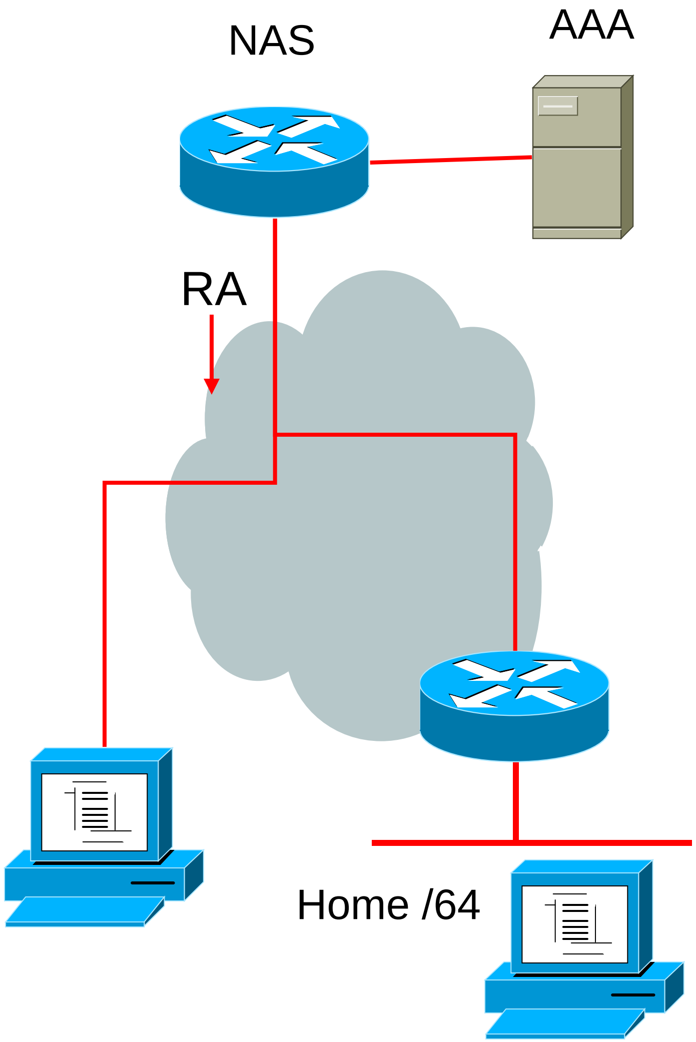

AAA config:
Auth-Type = Local, Password = “foo”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “addr-pool=“foo”

Address Assignment – short-lived /128

- Use: for single PC only. Allows one address
- /64 prefix shared between all users of the pool
- AAA interface-id attribute can be used to specify complete address
- NAS: IPv6 shared prefix pools
- CPE: Single PC

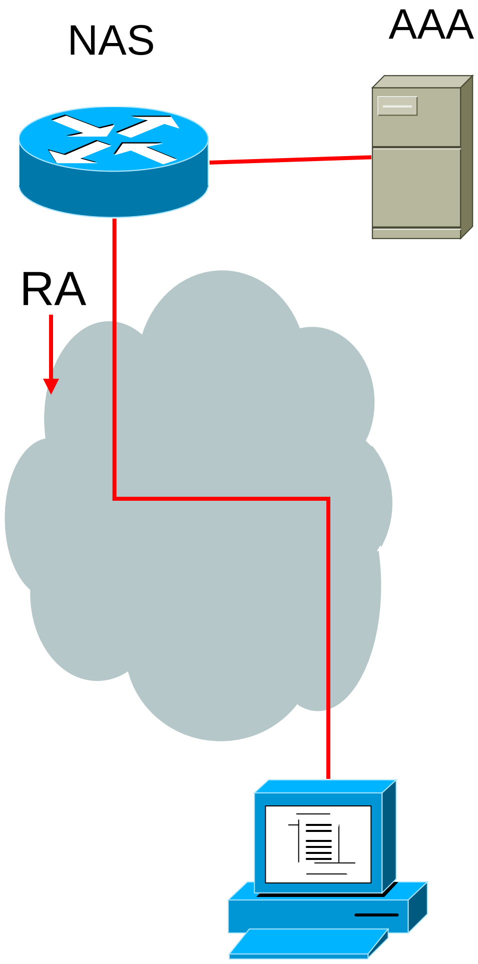

AAA config:
Auth-Type = Local, Password = “foo”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “addr-pool=“foo-shared”

IPv6 on Broadband Infrastructure Requirements

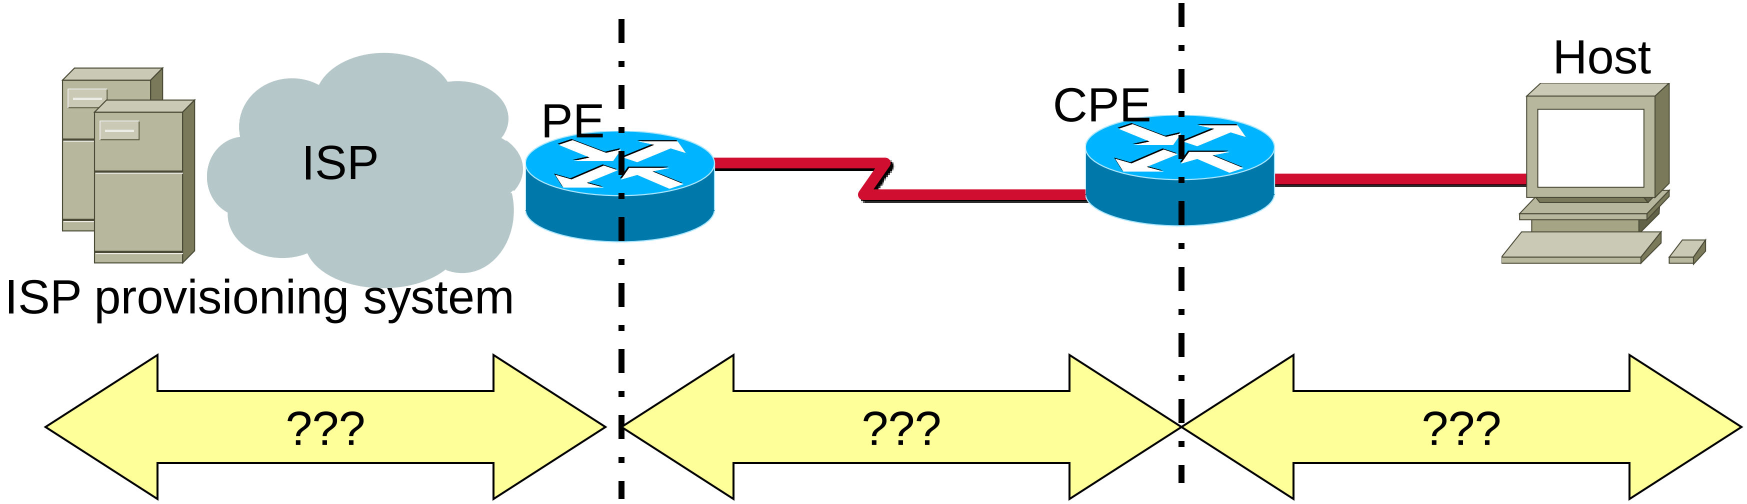

How do we get the configuration information and prefixes from the ISP provisioning system, to the PE, from the PE to the user CPE, and from the CPE to the end user hosts?
Routes for delegated prefixes/addresses also need to be injected into the ISP's routing system.

Prefix Delegation

Assignment of variable length prefixes
Independent of end user topology
Media independent
Additional Informations (DNS, NTP, SMTP, POP, etc)

Large Scale Deployment Suggested solution

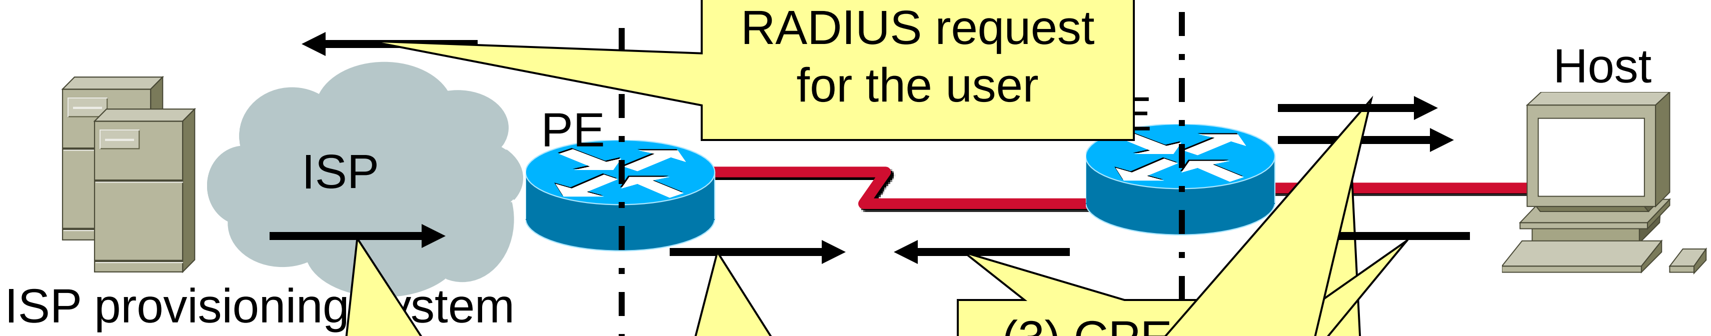

(1) PE sends RADIUS request for the user

(2) RADIUS
respond
user's pre

(3) CPE
D
wit

(4) PE sends
DHCP REPLY

(
configures addresses
the prefix on its
terfaces, and
-bit is set to on.

AA

/DHCP

(6) Host
the prefi

INFOR

(7) CPE sends a DHCP REPLY containing request options. Note that the CPE is configured as a DHCP client upstream, and as a DHCP server downstream. The DHCP downstream server acts as a cache, and uses the options received on the upstream interface.

The PE
configu
E can auto-
erent from the prefix
assigned to the user.

Enterprise IPv6 Address Allocation

- Enterprise addressing scheme
  - Get you own address from local registry via IANA OR
  - Get it via Service Providers
- Unique local address if the network does not need to go on the Internet
- Usually get a block of /48 unless a justification for a larger block is made
- PI address for multihoming

Provider-Independent Addresses

Provider Independent Proposal:
http://www.arin.net/policy/proposals/2005_1.html

- Driven mainly by enterprises
- Adopted (April 2006) because there is no consensus on Multihoming for IPv6 (NANOG rejected the IETF shim proposal)
- The possible impact is still debated but it seems we will just have to deal with it. Lack of PI could however slow down IPv6 adoption.
- BGP can only control routing table growth if routes are aggregated
- Number of multi-homed sites increasing quickly (>10,000)
- The IPv6 address space is very large
- Routing table growth could be problematical with the capability of the current hardware and protocols

Link Level – Prefix Length Considerations

64 bits

- Recommended by RFC3177 and IAB/IESG
- Consistency makes management easy
- MUST for SLAAC
- Significant Address space loss

< 64 bits

- Enables more hosts per broadcast domain
- Considered bad practice
- 64 bits offers more space for hosts than the media can support efficiently

> 64 bits

- Address space conservation
- Special cases:
  /126 – valid for p2p
  /127 – not valid for p2p (RFC3627)
  /128 – loopback
- Complicates management
- Must avoid overlap with specific addresses:
  Router Anycast (RFC3513)
  Embedded RP (RFC3956)
  ISATAP addresses

Interface-ID Selection
Network Devices

- Reconnaissance for network devices – the search for something to attack
- Use random 64-bit interface-IDs for network devices

  2001:DB8:CAFE:2::1/64 – Common IID

  2001:DB8:CAFE:2::9A43:BC5D/64 – Random IID

  2001:DB8:CAFE:2::A001:1010/64 – Semi-random IID
- Operational management challenges with this type of numbering scheme
- EUI-64 remains the easiest form of select interface-ID

Routing Deployments

IGP & BGP

Routing in SP network

- Prefixes coming into the SP network could be:
  - SP's owned pefixes assigned by the SP to the consumer
  - Enterprise owned prefixes from their allocated block
- Options for SP to provide Transit services
  - The transit routing can be done via BGP as in IPv4
  - The MPLS based SP can provide 6PE & 6VPE services
- Purist Provider
  - Cable providers (usually no MPLS)
  - Tunnel at the edge using GRE, L2TP
  - 6to4?

IPv6 Challenges to Router Performance

Addressing Driven

- Forwarding challenges—lookup not impacted as much as originally thought, different size prefixes typically see little difference in forwarding performance
- Control plane challenges—routing table sizes:

  IPv6 supports multiple addresses per interface
  (not the most significant concern at this time but
  it could be in the future)

  IPv6 can have a lot more prefixes due to a significantly larger address space

The Questions Are the Same as for IPv4… Almost

- Is one routing protocol better than any other routing protocol?
- Define "Better"

- Converges faster?
- Uses less resources?
- Easier to troubleshoot?
- Easier to configure?
- Scales to a larger number of routers, routes, or neighbors?
- More flexible?
- Degrades more gracefully?
- And so on

IPv6 IGP Selection—In Theory

In Theory:

- The similarity between the IPv6 and IPv4 routing protocols leads to similar behavior and expectations
- To select the IPv6 IGP, start by using the IPv4 IGP rules of thumb

IPv6 IGP Selection—In Practice

- In practice:
  - The IPv6 IGP implementations might not be fully optimized yet so there is a bit more uncertainty
  - Not all knobs for Fast Convergence might be available
  - No significant operational experience with large scale IPv6 networks

Conclusions

- Same topology considerations as for IPv4
- Convergence time

  There are HW and SW dependencies

  The average convergence time is 100% larger than IPv4, as IPv6 converges after IPv4

  Not all knobs are available. Ex: Fast Hellos for OSPFv3 -> Bidirectional Forwarding Detection (BFD) instead in the future.

  Test tools still need to improve

IGP deployment

- ALL IPv6 IGP runs over link local addressing
- Global Prefixes may not need to be assigned on the interface
- This reduces the size of the routing table
- SNMP can be used to manage the links
- Router needs to have one IPv4 & one IPv6 loopback assigned.
- SNMP polling can be done over IPv4 or IPv6.
- Infrastructure security by reducing routes & using link-local
- Core will act as a transit point.
- This makes the network more scalable

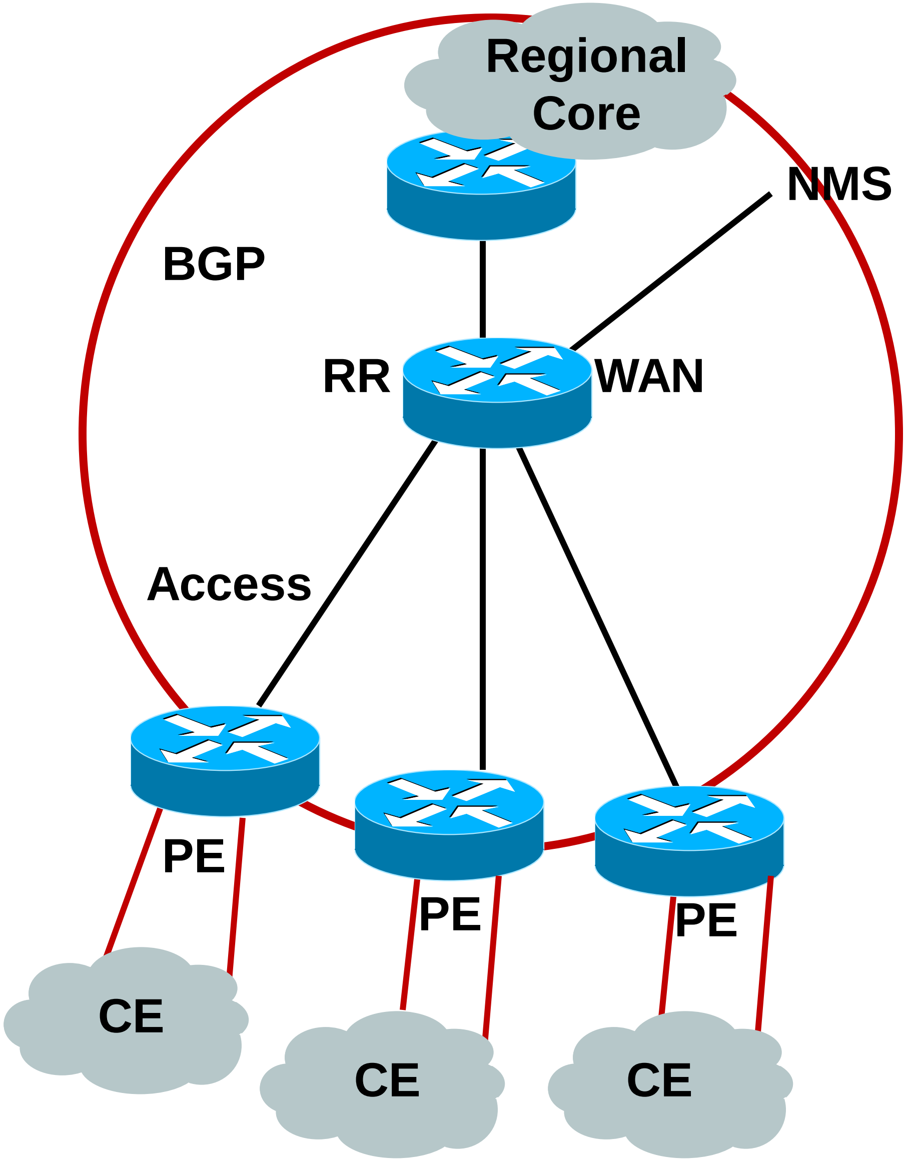

Routing Deployments

ISISv6

Integrated IS-IS for IPv6—Overview

- IETF draft: draft-ietf-isis-ipv6-06.txt
- Two TLVs added to support IPv6:

  IPv6 Reachability TLV (0xEC)—Describes network reachability (IPv6 routing prefix, metric information and option bits). 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)—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)

Integrated IS-IS—IPv4 and IPv6

- 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

- Multi-topology (draft-ietf-isis-wg-multi-topology)

  Independent IPv4 and IPv6 topologies

  Independent interface metrics

- Transition mode available—both types of TLVs are advertised

IS-IS Single Topology Example

```
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     Router2.01       Phase V
  Area Address(es): 49.0001
  IP Address(es):  10.7.1.34*
  IPv6 Address(es): FE80::2B0:4AFF:FE5C:ACA9
  Uptime: 00:01:25
  NSF capable
```

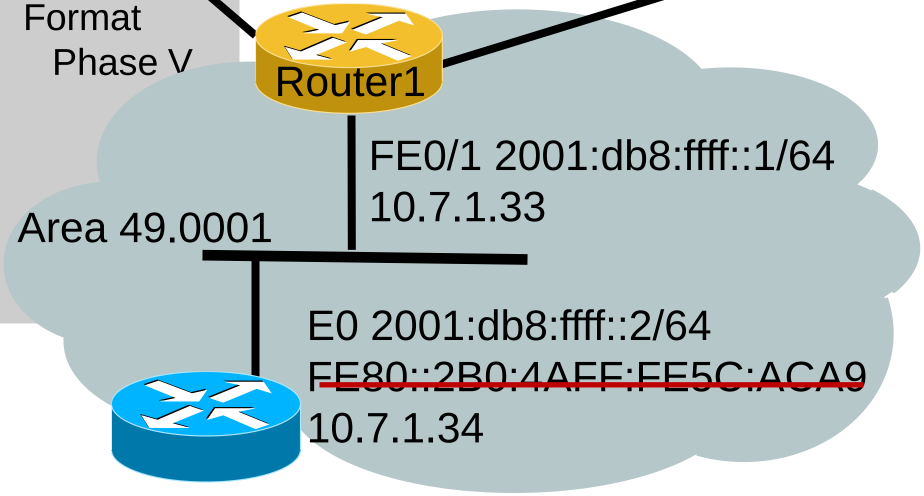

IS-IS Multi Topology Example

```
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
```

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

```
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
```

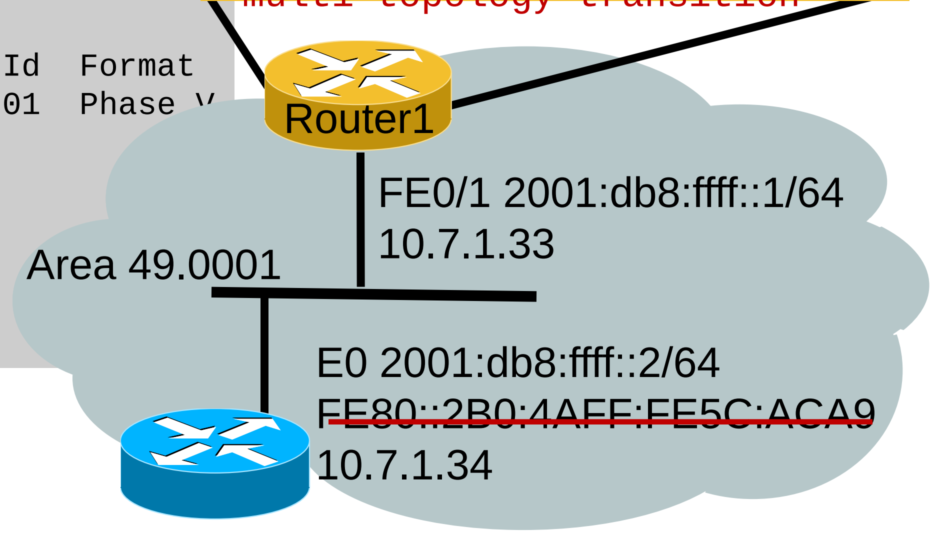

Routing Deployments

OSPFv3

Similarities with OSPFv2

- OSPFv3 is based on OSPFv2:
  - Runs directly over IPv6 (port 89)
  - Uses the same basic packet types
  - Neighbor discovery and adjacency formation mechanisms are identical (all OSPF routers FF02::5, all OSPF DRs FF02::6)
  - LSA flooding and aging mechanisms are identical
  - Same interface types (P2P, P2MP, broadcast, NBMA, virtual)
- OSPFv3 and OSPFv2 are independent processes and run as ships in the night

V2, V3 Differences

OSPFv3 Is Running per Link Instead of per Node (and IP Subnet)

- A link by definition is a medium over which two nodes can communicate at link layer
- Regardless of assigned prefixes, two devices can communicate using link-local addresses therefore OSPFv3 is running per link instead of per IP prefix
- Multiple IPv6 prefixes can be assigned to the same link

V2, V3 Differences (Cont.)

Support of Multiple Instances per Link

- New field (instance) in OSPF packet header allows running multiple instances per link
- Instance ID should match before packet is being accepted
- Useful for traffic separation, multiple areas per link

V2, V3 Differences (Cont.)

Address Semantic Changes in LSA

- Router and network LSA carry only topology information
- Router LSA can be split across multiple LSAs; link state ID in LSA header is a fragment ID
- Intra-area prefixes are carried in a new LSA payload called intra-area-prefix-LSAs
- Prefixes are carried in the payload of inter-area and external LSA

V2, V3 Differences (Cont.)

Generalization of Flooding Scope

- In OSPFv3 there are three flooding scopes for LSAs (link-local scope, area scope, AS scope) and they are coded in the LS type explicitly
- In OSPFv2 initially only area and AS wide flooding was defined; later opaque LSAs introduced link local scope, as well

V2, V3 Differences (Cont.)

Explicit Handling of Unknown LSA

- The handling of unknown LSA is coded via U-bit in LS type
- When U bit is set, the LSA is flooded within the corresponding flooding scope, as if it was understood
- When U bit is not set, the LSA is flooded within the link local scope
- In v2 unknown LSA were discarded

V2, V3 Differences (Cont.)

Authentication Is Removed from OSPF

- Authentication in OSPFv3 has been removed and OSPFv3 relies now on IPv6 authentication header since OSPFv3 runs over IPv6
- Autype and authentication field in the OSPF packet header therefore have been suppressed

V2, V3 Differences (Cont.)

OSPF Packet Format Has Been Changed

- The mask field has been removed from hello packet
- IPv6 prefix are only present in payload of link state update packet

V2, V3 Differences (Cont.)

Two New LSAs Have Been Introduced

- Link-LSA has a link local flooding scope and has three purposes

  Carry IPv6 link local address used for NH calculation

  Advertise IPv6 global address to other routers on the link (used for multi-access link)

  Convey router options to DR on the link

- Intra-area-prefix-LSA to advertise router’s IPv6 address within the area

LSA Types

	LSA Function Code	LSA Type
Router-LSA	1	0x2001
Network-LSA	2	0x2002
Inter-Area-Prefix-LXA	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	0x0008
Intra-Area-Prefix-LSA	9	0x2009

New

OSPFv3 Configuration Example

```
Router2#
interface POS3/0
 ipv6 address 2001:db8:FFFF:1::1/64
 ipv6 enable
 ipv6 ospf 100 area 1

ipv6 router ospf 100
   router-id 10.1.1.4
```

```
Router1#
interface POS1/1
 ipv6 address 2001:db8:EEEE:1::1/64
 ipv6 enable
 ipv6 ospf 100 area 0

interface POS2/0
 ipv6 address 2001:db8:FFFF:1::2/64
 ipv6 enable
 ipv6 ospf 100 area 1

 ipv6 router ospf 100
   router-id 10.1.1.3
```

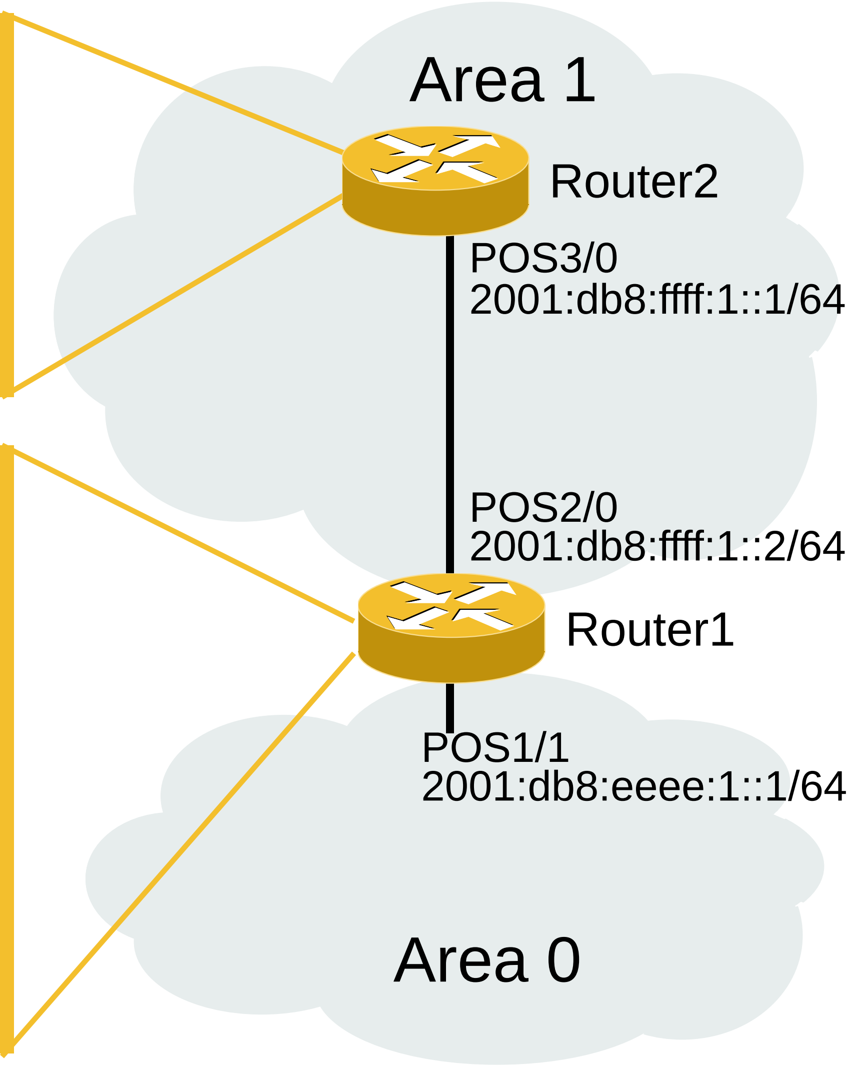

OSPFv3 Configuration Example (Cont.)

```
Router2#show ipv6 ospf int pos 3/0
POS3/0 is up, line protocol is up
  Link Local Address FE80::290:86FF:FE5D:A000, Interface ID 7
  Area 1, Process ID 100, Instance ID 0, Router ID 10.1.1.4
  Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 3, maximum is 3
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.1.3
  Suppress hello for 0 neighbor(s)
```

OSPFv3 Configuration Example (Cont.)

```
Router2#show 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:db8:EEEE:1::/64 [110/2]
     via FE80::2D0:FFFF:FE60:DFFF, POS3/0
C   2001:DB8:FFFF:1::/64 [0/0]
     via ::, POS3/0
L   2001:DB8:FFFF:1::1/128 [0/0]
     via ::, POS3/0
L   FE80::/10 [0/0]
     via ::, Null0
L   FF00::/8 [0/0]
     via ::, Null0
```

OSPFv3 Future Developments

- OSPFv3 must be developed to support other capabilities besides unicast IPv6 routing:
  - IPv6 unicast and multicast
  - IPv4 unicast and multicast
  - Multi-topologies within each address family
- This is work in progress in terms of standardization, with implementations to follow:
  - The complete solution is offered through MT support for multiple address families: draft-ietf-ospf-mt-ospfv3
  - An intermediary solution is proposed where distinct instances of OSPFv3 are used for each address family. Each AF/Instance will have its own adjacencies*, databases and SPF calculations thus operating as ships in the night: draft-ietf-ospfv3-af-alt.

*For All AFs, the Adjacencies Are Built over IPv6

Routing Deployments

BGP4+

BGP for IPv6

- BGP can carry IPv6 prefixes without changing its current transport mchanism which is IPv4
- Link Local peering can also be used for more secure peering
- Few things need to be considered with link local peering

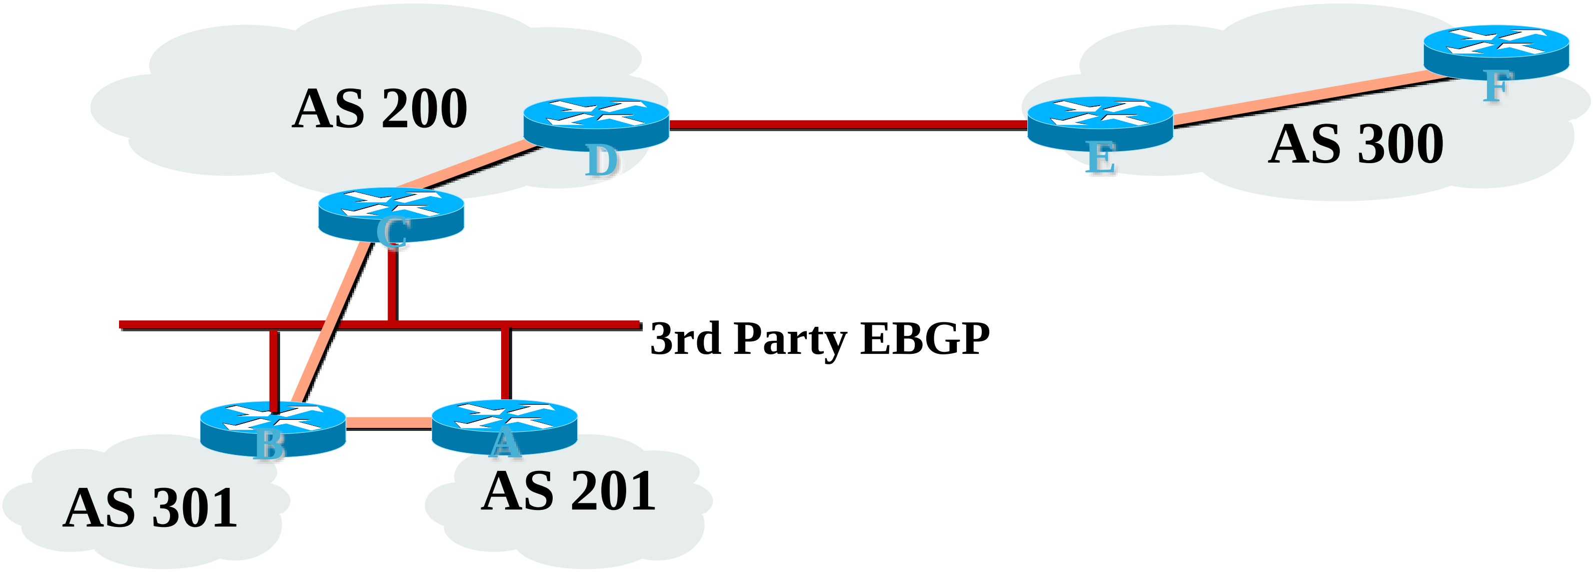

BGP-4 Extensions for IPv6

BGP-4 Carries Only 3 Pieces of Information Which Are Truly IPv4 Specific:

- NLRI in the UPDATE message contains an IPv4 prefix
- NEXT_HOP path attribute in the UPDATE message contains an IPv4 address
- BGP Identifier is in the OPEN message and AGGREGATOR attribute

BGP-4 Extensions for IPv6

To Make BGP-4 Available for Other Network Layer Protocols, RFC 2858 (Obsoletes RFC 2283) Defines Multiprotocol Extensions for BGP-4:

- Enables BGP-4 to carry information of other protocols (MPLS, IPv6, etc.)
- 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 changed 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 sent within the OPEN message

BGP-4 Configurations for IPv6 Non-Link Local Peering

Router A

```
router bgp 100
 bgp log-neighbor-changes
 neighbor 2001:100:3:4::1 remote-as 100
 neighbor 200.10.10.1 remote-as 200
 !
 address-family ipv6
 neighbor 2001:100:3:4::1 activate
 neighbor 200.10.10.1 activate
 neighbor 200.10.10.1 route-map SETNH
out
 redistribute connected
!
route-map SETNH permit 10
 set ipv6 next-hop 2001:100:3:1::1
```

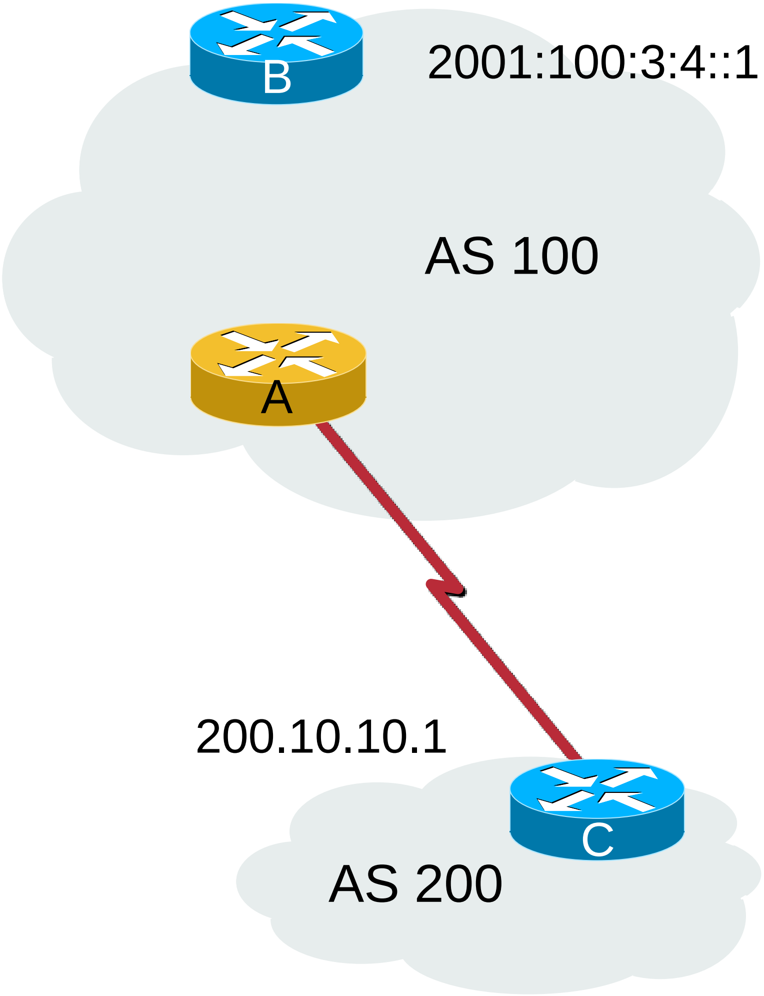

BGP-4 Configurations for IPv6 Link Local Peering

Router A

```
router bgp 200
 neighbor FE80::A8BB:CCFF:FE01:F600%Ethernet0/0 remote-as 100
 !
 address-family ipv6
 neighbor FE80::A8BB:CCFF:FE01:F600%Ethernet0/0 activate
 neighbor FE80::A8BB:CCFF:FE01:F600%Ethernet0/0 route-map SETNH out
 redistribute connected
 no synchronization
!
route-map SETNH permit 10
 set ipv6 next-hop 2001:100:1:1::2
```

Future CLI

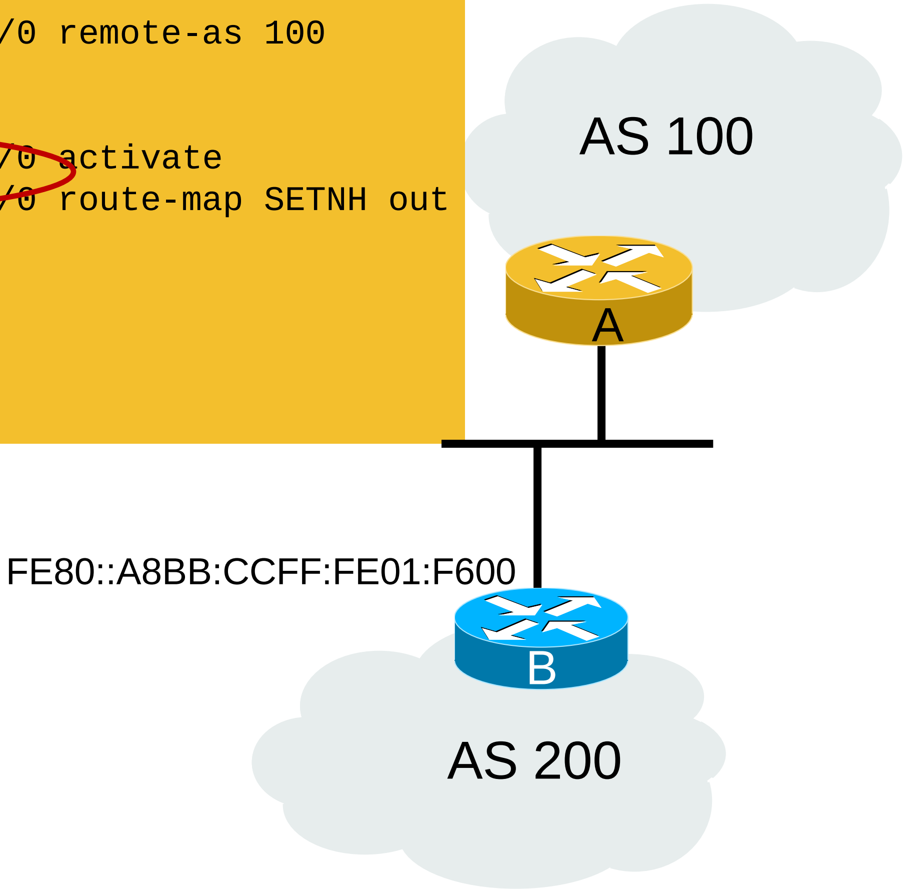

BGP-4 for IPv6 « Show Command »

- Show bgp IPv6

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

Routing Protocols Co-existence & Convergence

The Questions Are Almost the Same as for IPv4

- Most likely the IPv6 IGP will not be deployed in a brand new network and just by itself
- Most likely the IPv4 services are more important at first since they are generating most of the revenue
- Redefine “better”

- What is the impact on the convergence of IPv4?
- Are the resources optimally shared?
- Are the topologies going to be congruent?
- Etc.

Co-existence—Convergence Considerations

At First, the IPv6 IGP Convergence Might Be Less Important than the Impact of IPv6 on the Convergence of the Existent IPv4 Infrastructure

- What IGPs coexist better?
- What IPv6 IGP impacts IPv4 the least (hopefully not at all)?

Nothing Is for Free

- Resources **will be shared** between the two IGPs and they **will compete** for processor cycles in a way that reflects their relative configuration
- This has implications on:
  - Expected convergence behavior
  - Single process/topology vs Multi process/topology selection
  - Resources (Memory, CPU) planning

Coexistence—Resources Considerations

- With the exception of ISIS single topology, the IPv4 and IPv6 routing processes claim their own memory and processing resources for maintaining adjacencies, databases and related calculations
- It is important to define the IPv6 network design in order to understand the new resource requirements (memory) and the new operational parameters (max CPU) for the network devices

Coexistence—Topology Considerations

- The IPv4 and IPv6 topologies can be:
  - Congruent
    - Dual-stack deployment
  - Non-Congruent
    - Not all network devices are supporting the necessary IPv6 features so they must be avoided during migration
- Non-congruent is not necessarily bad, even though it might be more difficult to manage and troubleshoot. Strive for congruent topologies.

Convergence Considerations

The IGPs Will Compete over Processor Cycles Based on Their Relative Tuning

- If you configure the IPv4 and IPv6 IGPs the same way (aggressively tuned for fast convergence), naturally expect a doubling of their stand alone operation convergence time
- If the IPv6 IGP is operating under default settings, the convergence time for the optimally tuned IPv4 IGP is not significantly affected

OSPFv3 Fast Convergence

- Following Techniques/tools are available for fast convergence in OSPFv3
  - Carrier Delays **Detect**
  - Hello/dead timers (Fast Hellos) (not available) **Detect**
  - Bi-Directional Forwarding Detection—(BFD) **Detect**
  - LSA packet pacing **Propagate**
  - Interface event dampening - **Propagate**
  - Exponential throttle timers for LSA & SPF **Process**
  - MinLSArrival Interval **Process**
  - Incremental SPF (not available) **Process**
- Techniques/tools for Resiliency
  - Stub router (e.g., max-metric) (not available)
  - Cisco NSF (RFC 4811,4812,4813) (not available)
  - Graceful Restart (ONLY RFC 3623)

ISIS Fast Convergence

- Following Techniques/tools are available for fast convergence in ISIS

  Carrier Delays **Detect**

  Hello/dead timers (Fast Hellos) **Detect**

  Bi-Directional Forwarding Detection—(BFD) **Detect**

  LSP pacing **Propagate**

  Interface event dampening - **Propagate**

  Exponential throttle timers for LSA & SPF **Process**

  PRC-interval **Process**

  Incremental SPF **Process**

- Techniques/tools for Resiliency

  Cisco NSF

  Graceful Restart

Summary

- In summary we learned:
- Address allocation in both SP and Enterprise networks
- SP & Enterprise Architecture
- IPv6 Routing deployment techniques
- Co-existence & Convergence of Routing protocols

Complete Your Online Session Evaluation

- Give us your feedback and you could win fabulous prizes. Winners announced daily.
- Receive 20 Cisco Preferred Access points for each session evaluation you complete.
- Complete your session evaluation online now (open a browser through our wireless network to access our portal) or visit one of the Internet stations throughout the Convention Center.

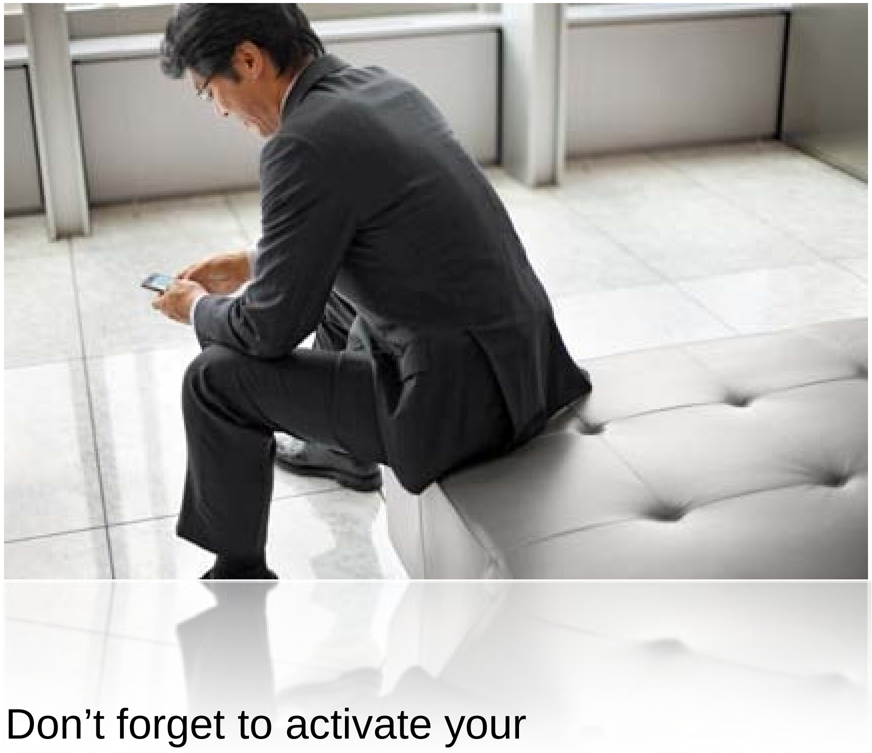

Don’t forget to activate your Cisco Live and Networkers Virtual account for access to all session materials, communities, and on-demand and live activities throughout the year. Activate your account at any internet station or visit www.ciscolivevirtual.com.