TOMORROW starts here.

Cisco live!

Service Provider IPv6 Deployment

BRKSPG-3300

Salman Asadullah
IPv6 Forum Fellow
Cisco Distinguished Engineer
Broadband Forum Ambassador
Co-Chair IPv6 Ready Logo Program

Cisco live!

Prerequisites: Session Abstract

- This session focuses on SP IPv6 deployment techniques in core networks which will help network designers and administrators understand IPv6 operation and implementation options for native IPv4 and MPLS core environments. This session will also shed light on IPv6 Multihoming, addressing and Cisco Carrier-Grade IPv6 (CGv6) Solution considerations in core networks.
- Attendee must have a solid foundation of IPv6 basics (Protocol, Addressing, Routing), MPLS and Multicast.

Agenda

- **SP IPv6 Integration Strategy**
- IPv6 in Core Networks and Deployment Models
  - Native IPv4 Environments
  - MPLS Environments
- IPv6 Addressing Considerations
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- IPv6 Multi-homing Considerations
  - Goals and Solutions
  - LISP6
- Carrier-Grade IPv6 Solution – CGv6
  - CGN, 6rd, DS-Lite, AFT64, MAP
- Conclusion

Growing Internet Challenge & Evolution

Moving to 3 IP Address Families: Public IPv4, Private IPv4, IPv6

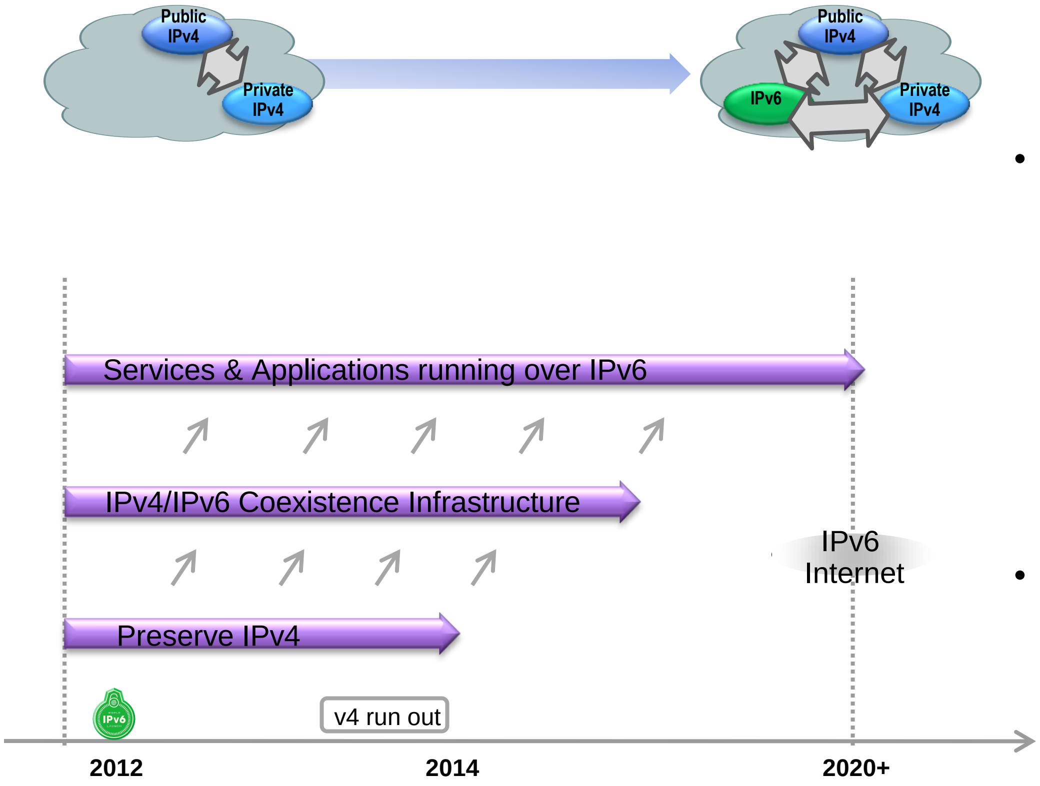

- **2012: Mandates take effect – Globalisation - WorldIPv6Launch - Massive Mobile deployment.** Transition to IPv6 forces customers to acquire product or managed services to sustain business and customer reach
- **2014: IPv6 is mainstream** Customers without transition infrastructure experience reduced service levels, diminished customer reach

IPv6 Market Adoption: 6lab.cisco.com/stats

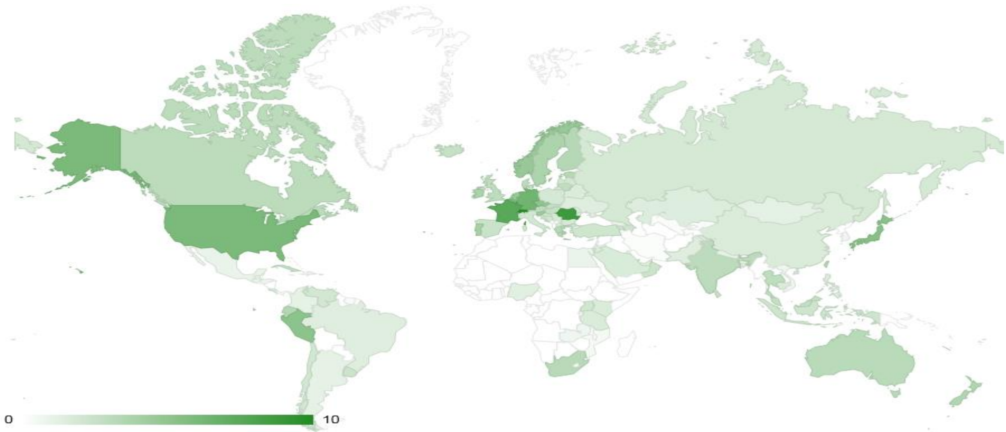

"If we have data, let's look at data. If all we have, are opinions, let's go with mine."

Jim Barksdale,

Former Netscape CEO

- ~80 % of Internet Core transit is IPv6 enabled
- ~ 35% of global Internet content/Web pages are reachable over IPv6
- ~1% of Internet users have IPv6
  - Great disparities across countries
- Progress has been encouraging: Google has recently reported that 3% of their traffic is now IPv6; Comcast and AT&T are showing double digit percentages of traffic. T-Mobile and Verizon has made great strides in getting IPv6 in their mobile networks, and the IPv6-disabled-by-default Windows XP goes end of life this year.

IPv6 in the SP: Drivers

- **External Drivers**
  - World IPv6 Launch :: June 6th, 2012
  - SP customers that need access to IPv6 resources (for development or experimentation purposes)
  - SP customers that need to interconnect their IPv6 sites
  - SP customers that need to interface with their own customers over IPv6 (ex: contractors for DoD)
- **Internal Drivers**
  - Handle some problems that are hard to fix with IPv4 (ex: managing large number of devices such as Cell phones, set-tops, IP cameras, sensors, etc.)
  - Public IPv4 address exhaustion
  - Private IPv4 address exhaustion
- **Strategic Drivers**
  - Long term expansion plans and service offering strategies
  - Preparing for new services and gaining competitive advantage

IPv6 Integration and Co-Existence

- Many ways to deliver IPv6 services to End Users, Most important is End to End IPv6 traffic forwarding as applications are located at the edge
- SPs may have different deployment needs and mechanisms but basic steps are common
  - –IPv6 Addressing Scheme
  - –Routing Protocol(s)
  - –IPv6 Services - QoS, Multicast, DNS, …
  - –Security
  - –Network Management
- Resources are shared between the two protocols for both Control and Forwarding Plane. Evaluate processor utilisation and memory needs
- Most vendors have good IPv6 HW forwarding performance

Agenda

- SP IPv6 Integration Strategy
- **IPv6 in Core Networks and Deployment Models**
  - IPv6 in Native IPv4 Environments
  - IPv6 in MPLS Environments
- IPv6 Addressing Considerations
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- IPv6 Multi-homing Considerations
  - Goals and Solutions
  - LISP6
- Carrier-Grade IPv6 Solution – CGv6
  - CGN, 6rd, DS-Lite, AFT64, MAP
- Conclusion

IPv6 Deployment Options — CORE

- SP Core Infrastructures – 2 Basic Paths
  - Native IPv4 core with associated services
  - L2TPv3, QoS, Multicast, …
  - MPLS with its associated services
  - MPLS/VPN, L2 services over MPLS, QoS, …
- IPv6 in Native IPv4 Environments
  - Tunnelling IPv6-in-IPv4
  - Native IPv6 with Dedicated Resources
  - Dual-Stack IPv4 and IPv6
- IPv6 in MPLS Environments
  - 6PE
  - 6VPE

Agenda

- SP IPv6 Integration Strategy
- IPv6 in Core Networks and Deployment Models
  - **IPv6 in Native IPv4 Environments**
  - IPv6 in MPLS Environments
- IPv6 Addressing Considerations
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- IPv6 Multi-homing Considerations
  - Goals and Solutions
  - LISP6
- Carrier-Grade IPv6 Solution – CGv6
  - CGN, 6rd, DS-Lite, AFT64, MAP
- Conclusion

Tunnelling IPv6 in IPv4

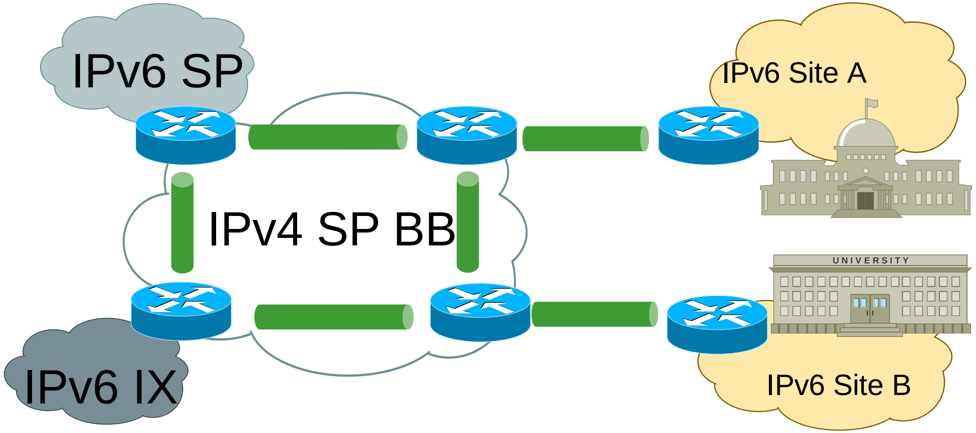

- Tunnelling Options
  - Manual Tunnels (RFC 2893), GRE Tunnels (RFC 2473), L2TPv3, …
- SP Scenarios
  - Configured Tunnels in Core
  - Configured Tunnels or Native IPv6 to IPv6 Enterprise’s Customers
  - MP-BGP4 Peering with other IPv6 users
  - Connection to an IPv6 IX

Native IPv6 over Dedicated Data Link

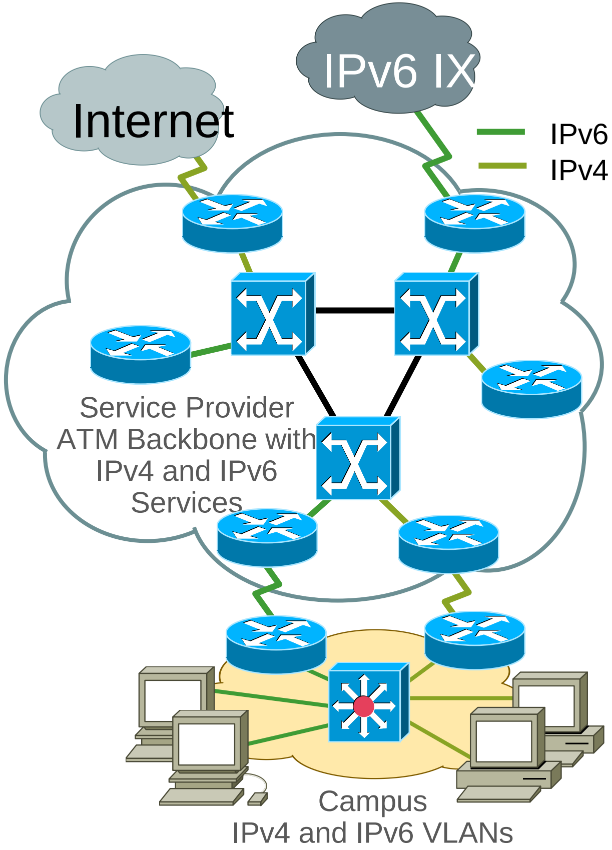

- ISP Scenario
  - –Dedicated Data Links between Core routers
  - –Dedicated Data Links to IPv6 Customers
  - –Connection to an IPv6 IX

Dual Stack IPv4 and IPv6

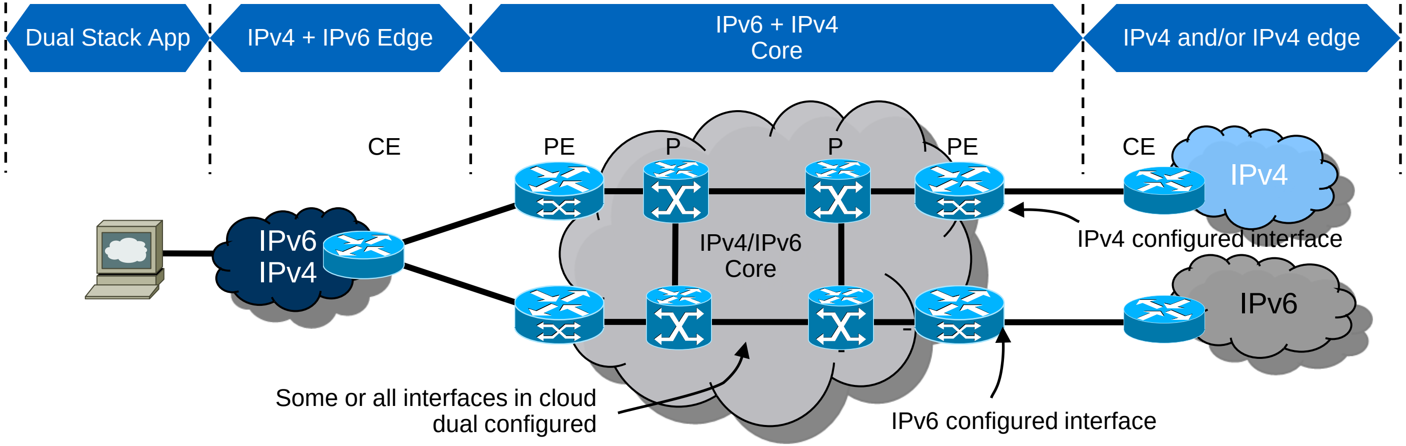

- All P + PE routers are capable of IPv4+IPv6 support
- Two IGPs supporting IPv4 and IPv6
- Memory considerations for larger routing tables
- Native IPv6 multicast support
- All IPv6 traffic routed in global space
- Good for content distribution and global services (Internet)

Agenda

- SP IPv6 Integration Strategy
- IPv6 in Core Networks and Deployment Models
  - IPv6 Native IPv4 Environments
  - **IPv6 in MPLS Environments**
- IPv6 Addressing Considerations
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- IPv6 Multi-homing Considerations
  - Goals and Solutions
  - LISP6
- Carrier-Grade IPv6 Solution – CGv6
  - CGN, 6rd, DS-Lite, AFT64, MAP
- Conclusion

IPv6 over MPLS

- Many ways to deliver IPv6 services to end users
  - Most important is end-to-end IPv6 traffic forwarding
- 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 MPLS




MPLS 6PE Overview

IPv6 Provider Edge Router (6PE) over MPLS

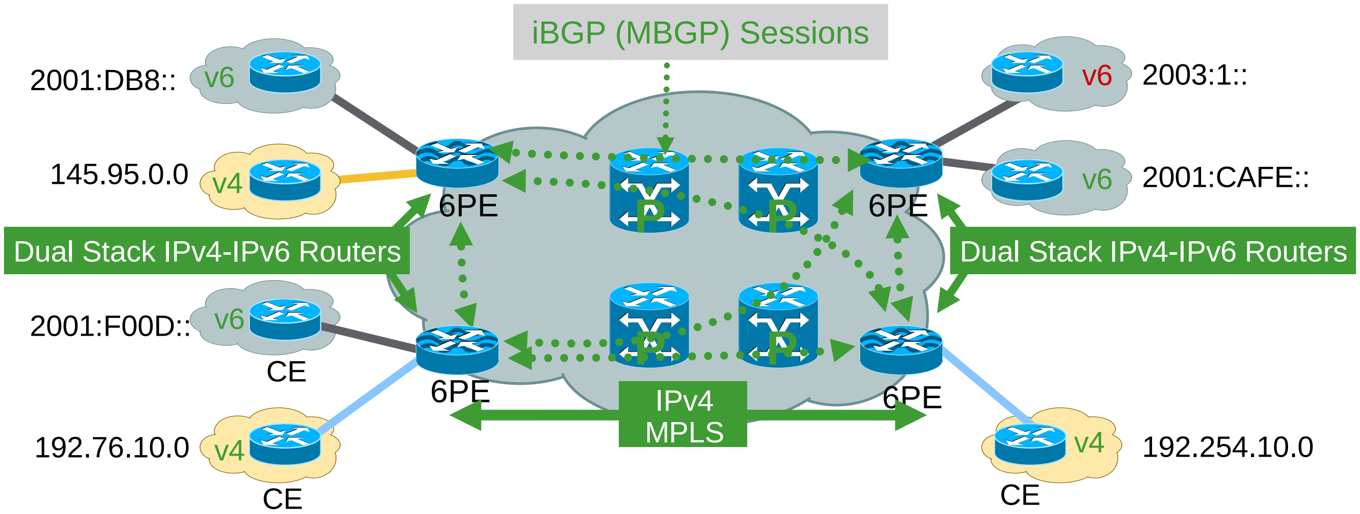

- IPv6 global connectivity over and IPv4-MPLS core
- Transitioning mechanism for providing unicast IP
- PEs are updated to support dual stack/6PE
- IPv6 reachability exchanged among 6PEs via iBGP (MBGP)
- IPv6 packets transported from 6PE to 6PE inside MPLS

6PE Routing/Label Distribution

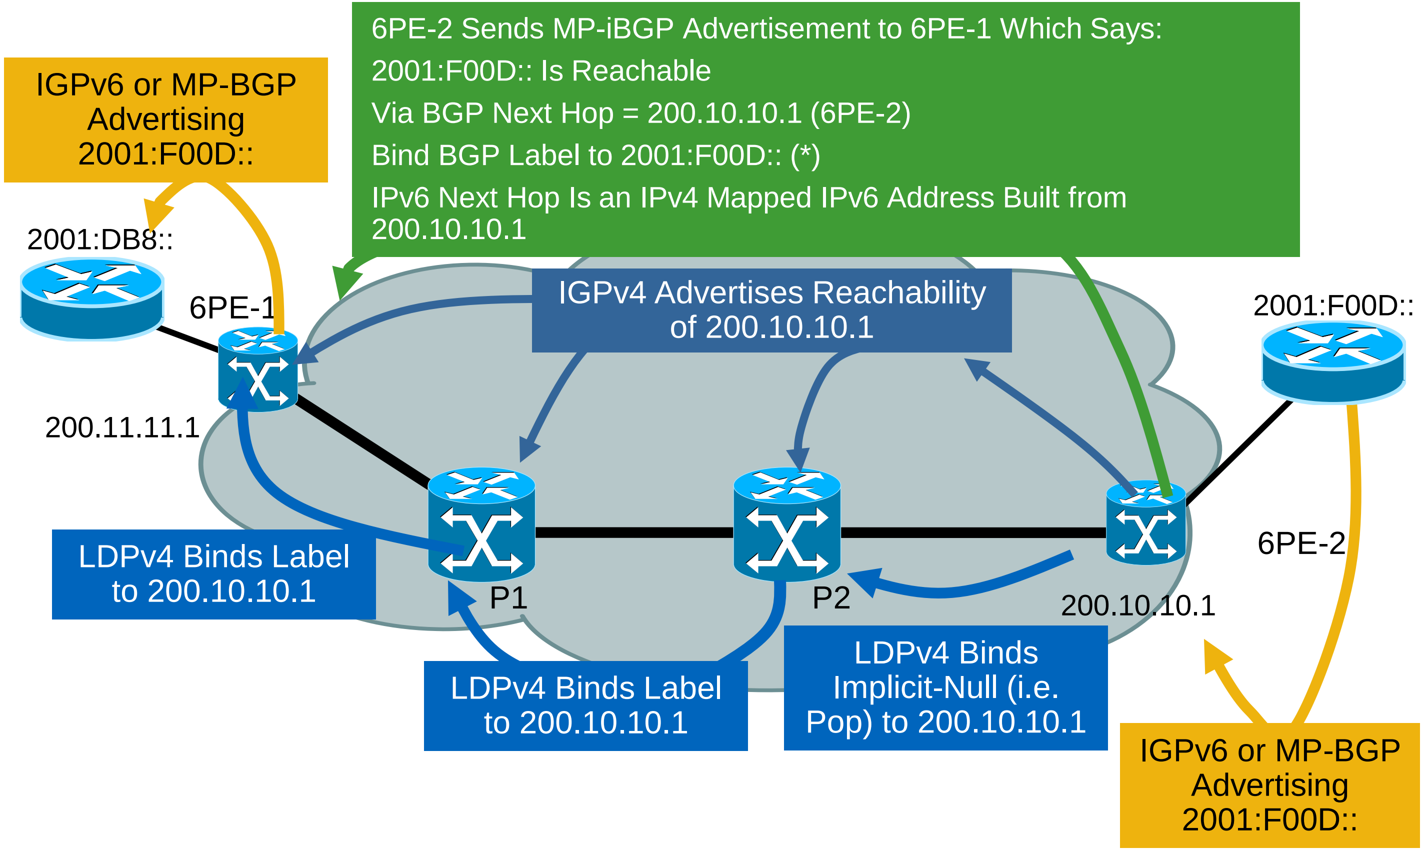

6PE Forwarding (6PE-1)

IPv6 Forwarding and Label Imposition:

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

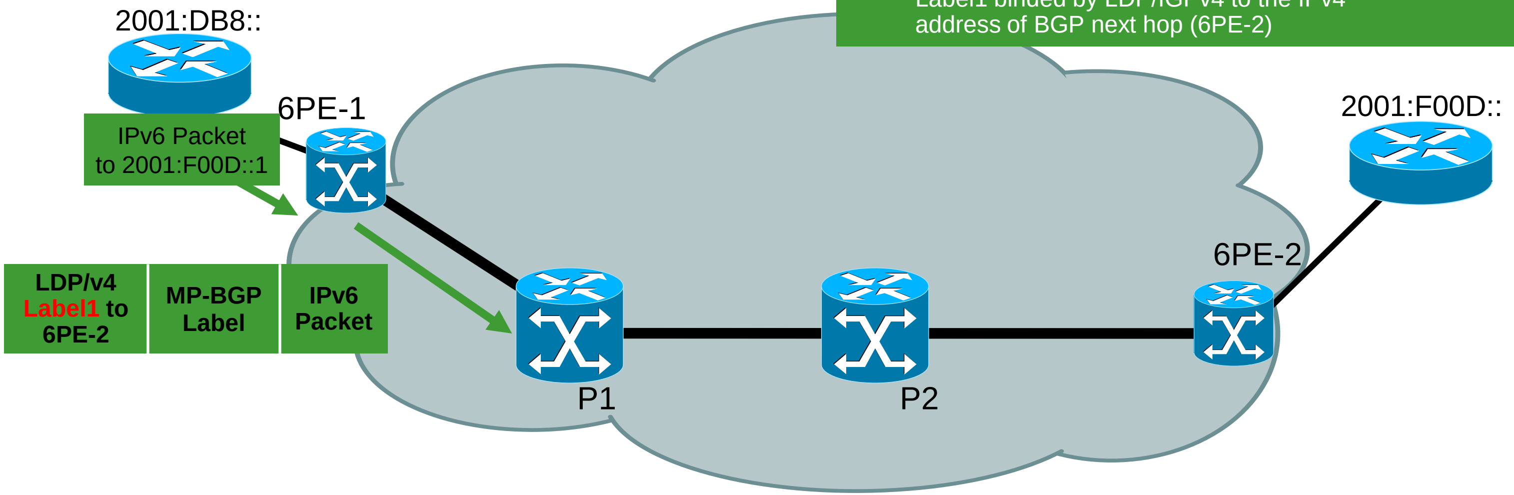

6PE Forwarding (P1)

IPv6-UNaware MPLS Label Switching:

- P1 receives an MPLS packet
- Lookup is done on Label1
- Result is Label2

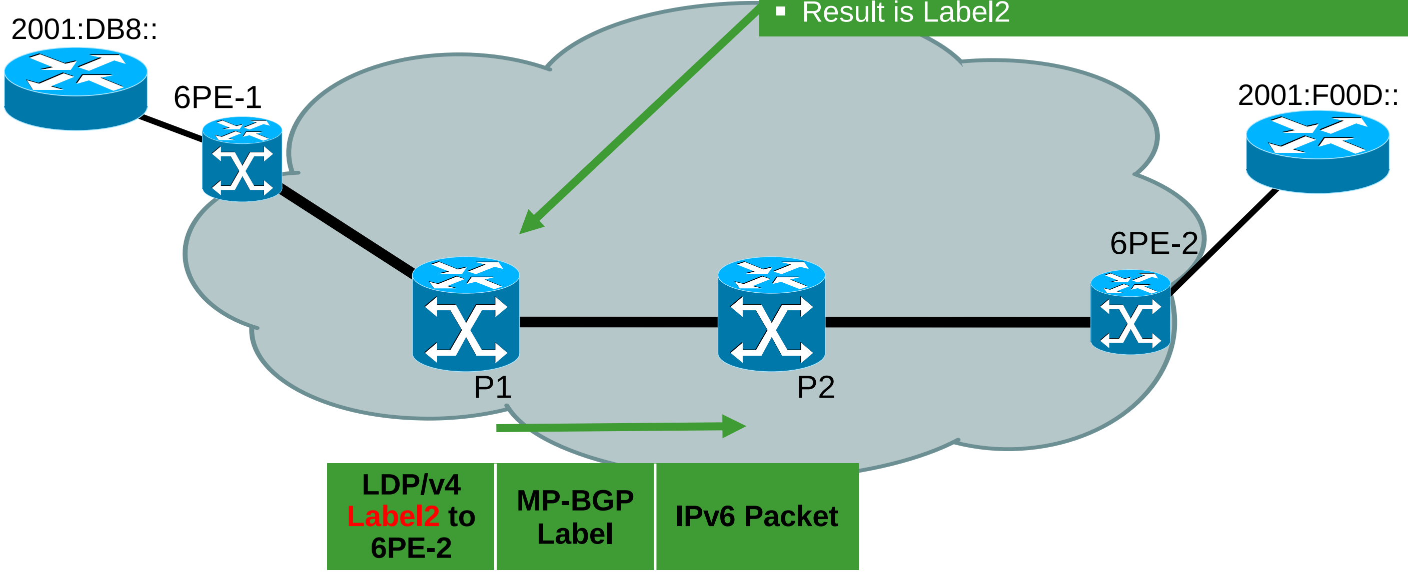

6PE Forwarding (P2)

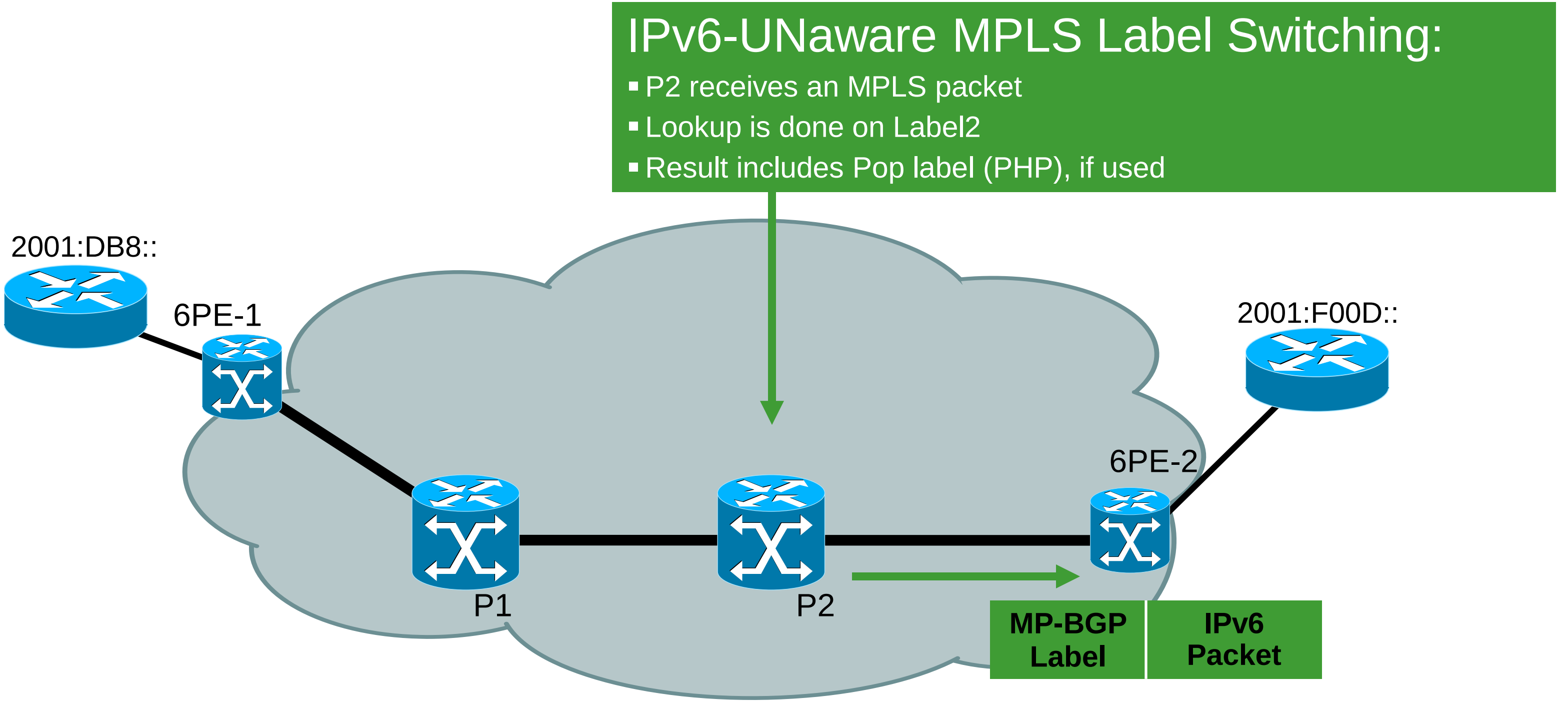

6PE Forwarding (6PE-2)

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

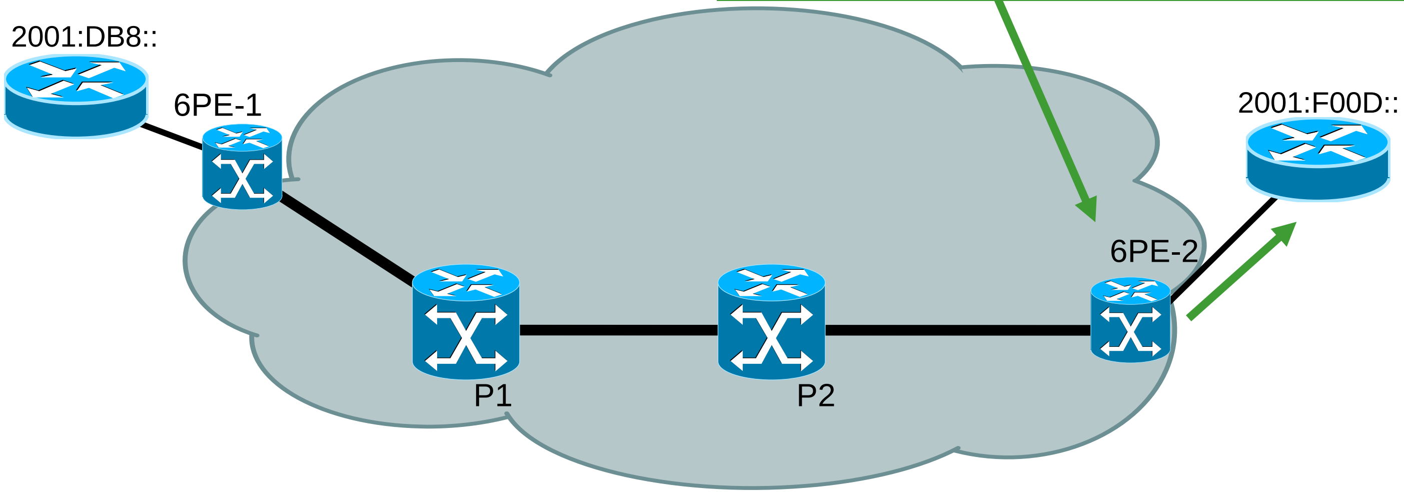

6PE-1 Configuration

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

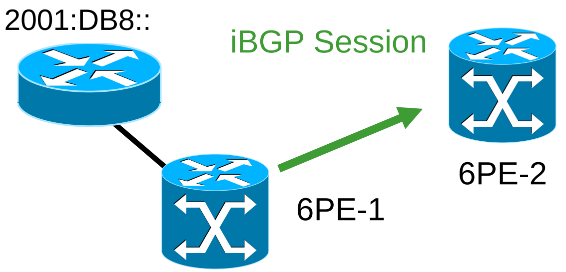

2001:DB8:1::1 Is the Local CE
200.10.10.1 Is the Remote 6PE

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

6PE Show Output

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

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

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

Other Useful Output:

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

6PE Summary

- Core network (Ps) untouched
- IPv6 traffic inherits MPLS benefits (fast re-route, TE, etc.)
- Incremental deployment possible (i.e., only upgrade the PE routers which have to provide IPv6 connectivity)
- Each site can be v4-only, v4VPN-only, v4+v6, v4VPN+v6
- P routers won't be able to send ICMPv6 messages (TTL expired, trace route)
- Scalability issues arise as a separate RIB and FIB is required for each connected customer
- Good solution only for SPs with limited devices in PE role
- Cisco 6PE Documentation/Presentations:

  http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_data_sheet09186a008052edd3.html




MPLS 6VPE Overview

6VPE over MPLS

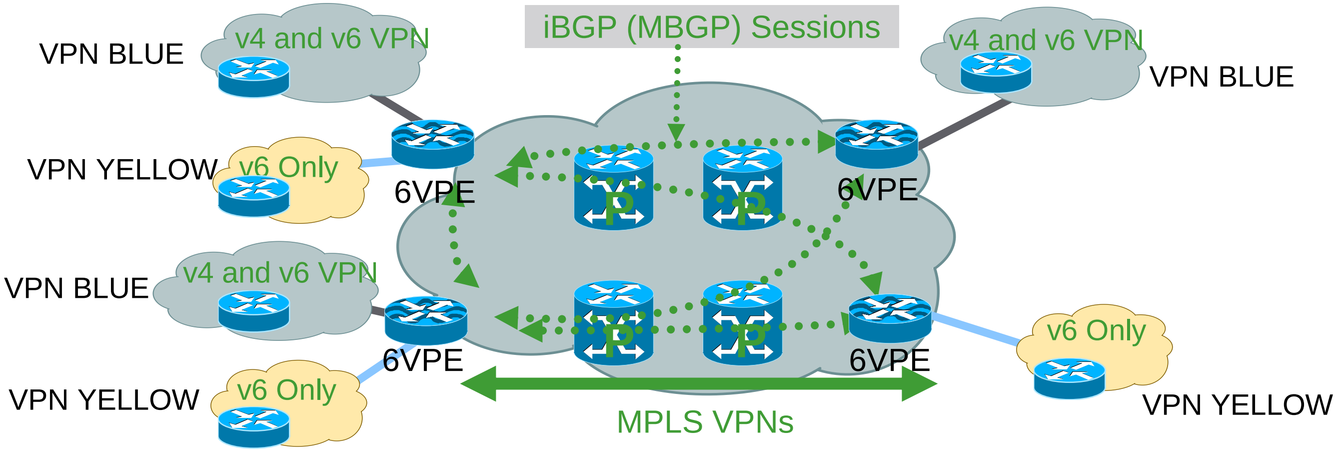

- 6VPE ~ IPv6 + BGP-MPLS + IPv4VPN + 6PE
- Cisco 6VPE (RFC4659)
- VPNv6 address:
  - Address including the 64 bits route distinguisher and the 128 bits IPv6 address

- MP-BGP VPNv6 address-family:
  - AFI “IPv6” (2), SAFI “VPN” (128)
- VPN IPv6 MP_REACH_NLRI
  - With VPNv6 next-hop (192bits) and NLRI in the form of <length, IPv6-prefix, label>
- Encoding of the BGP next-hop

Implementing 6VPE

- Implementing 6VPE is no different than implementing VPNv4 in the network
- Define VRFs if not already done for IPv4
- Associate interfaces to VRFs
- Add VPNv6 address family under MP-BGP
- Send extended communities for the prefixes under the new address family
- Consider redundancy and downtime while migrating from 6PE to 6VPE

6VPE Example Design
Addressing/Routing

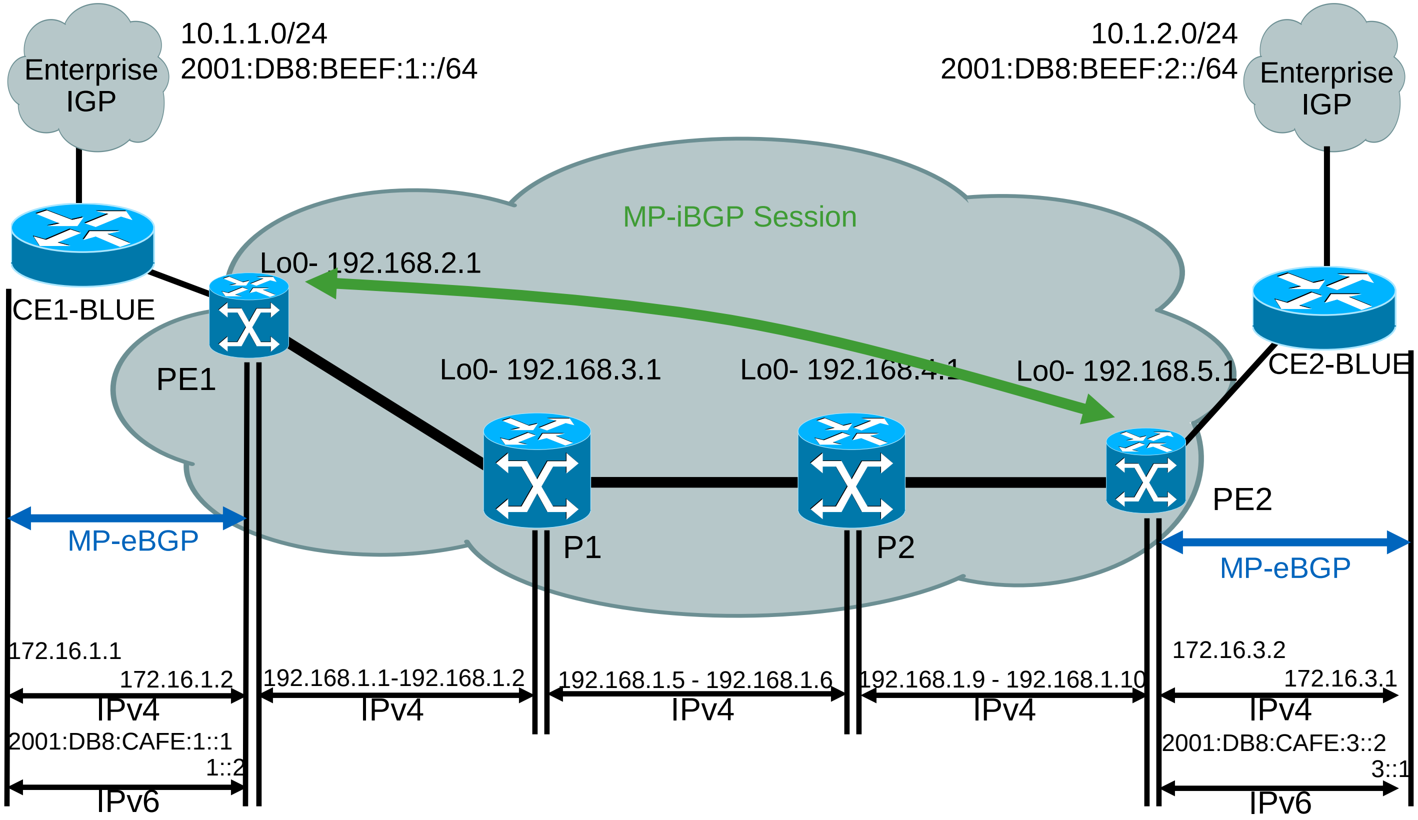

6VPE Configuration Example

CE1-BLUE to PE1

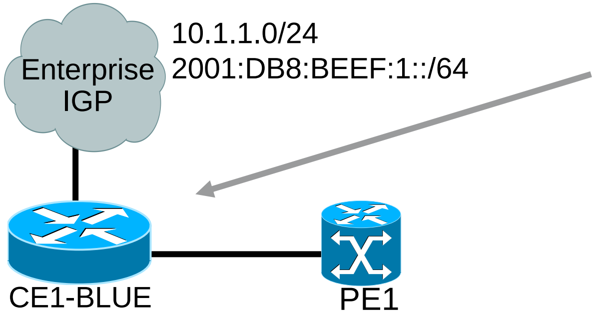


```
ipv6 unicast-routing
ipv6 cef
!
interface Ethernet0/0
 description to PE1
 ip address 172.16.1.1
255.255.255.0
 ipv6 address
2001:DB8:CAFE:1::1/64
!
interface Ethernet1/0
 description to BLUE LAN
 ip address 10.1.1.1 255.255.255.0
 ipv6 address
2001:DB8:BEEF:1::1/64
 ipv6 rip BLUE enable
```

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

6VPE Configuration Example

PE1 Connections

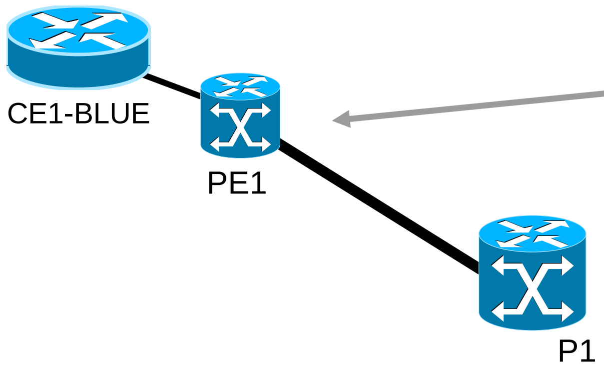


```
ipv6 unicast-routing
ipv6 cef
mpls ldp router-id Loopback0
!
interface Loopback0
 ip address 192.168.2.1 255.255.255.255
!
interface Ethernet0/0
 description to CE1-BLUE
 vrf forwarding BLUE
 ip address 172.16.1.2 255.255.255.0
 ipv6 address 2001:DB8:CAFE:1::2/64
!
interface Ethernet2/0
 description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

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

6VPE Configuration Example

PE1 VRF Definitions

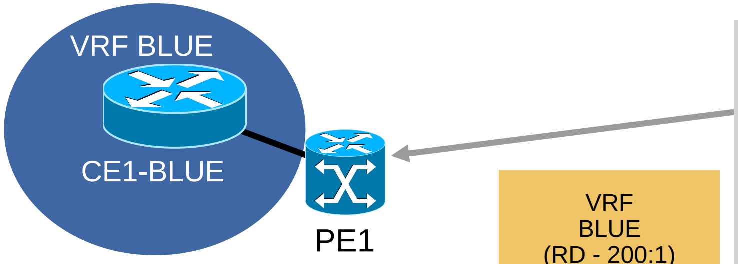

VRF
BLUE
(RD - 200:1)
PE1/PE2 Will Hold
CE1 Routes
CE2 Routes

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

- Migration commands available for VPNv4 to multi-protocol VRF
  - (config)#vrf upgrade-cli multi-af-mode {common-policies | non-common-policies} [vrf <name>]
- This command forces migration from old CLI for IPv4 VRF to new VRF multi-AF CLI

6VPE Configuration Example

PE1 BGP Setup

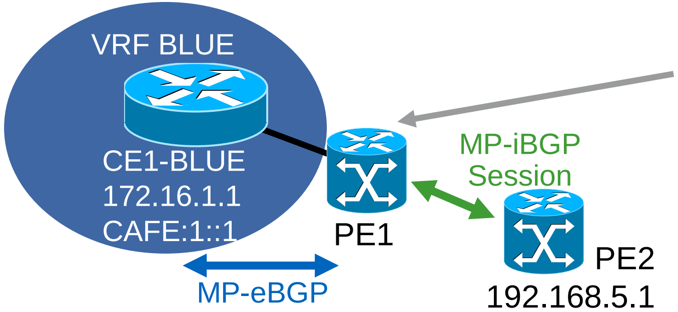


```
router bgp 100
 bgp log-neighbor-changes
 neighbor 192.168.5.1 remote-as 100
 neighbor 192.168.5.1 update-source
Loopback0
 !
 address-family ipv4
 neighbor 192.168.5.1 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family vpnv4
 neighbor 192.168.5.1 activate
 neighbor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

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

6VPE Configuration Example

P Connections

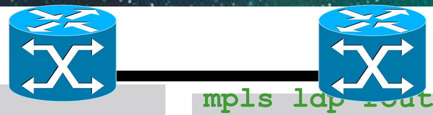


```
mpls ldp router-id Loopback0
!
interface Loopback0
 ip address 192.168.3.1
255.255.255.255
!
interface Ethernet0/0
 description to PE1
 ip address 192.168.1.2
255.255.255.252
mpls ip
!
interface Ethernet1/0
 description to P2
 ip address 192.168.1.5
255.255.255.252
mpls ip
!
router ospf 1
 log-adjacency-changes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

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

IPv6 Routing Tables

CE1-CE2

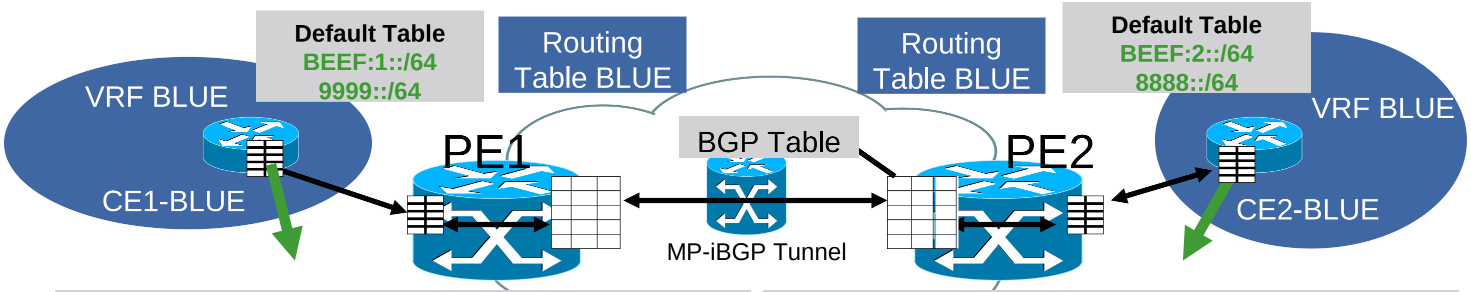


```
ce1-blue#show ipv6 route
C   2001:DB8:BEEF:1::/64 [0/0]
    via Ethernet1/0, directly connected
L   2001:DB8:BEEF:1::1/128 [0/0]
    via Ethernet1/0, receive
B   2001:DB8:BEEF:2::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F600, Ethernet0/0
C   2001:DB8:CAFE:1::/64 [0/0]
    via Ethernet0/0, directly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

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

IPv6 Routing Tables

PE1-PE2

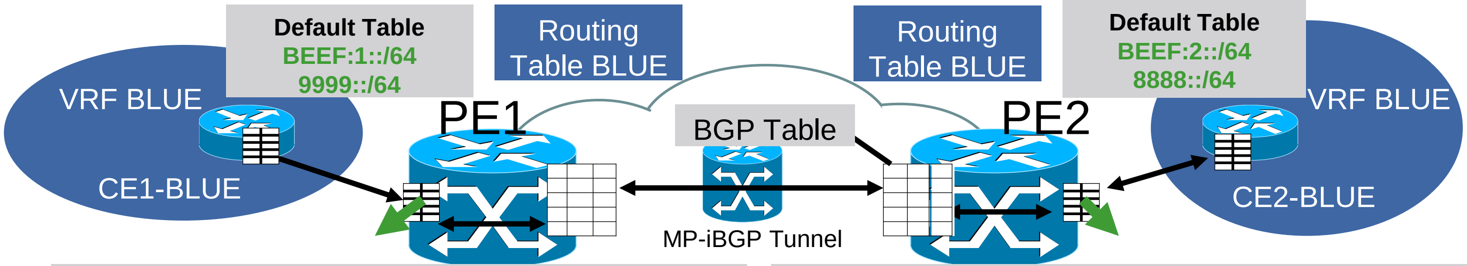


```
pe1#show ipv6 route vrf BLUE
B   2001:DB8:BEEF:1::/64 [20/0]
    via FE80::A8BB:CCFF:FE01:F400, Ethernet0/0
B   2001:DB8:BEEF:2::/64 [200/0]
    via 192.168.5.1%Default-IP-Routing-Table,
indirectly connected
C   2001:DB8:CAFE:1::/64 [0/0]
    via Ethernet0/0, directly connected
L   2001:DB8:CAFE:1::2/128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

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

IPv6 Routing Tables

PE1 BGP Next-Hop

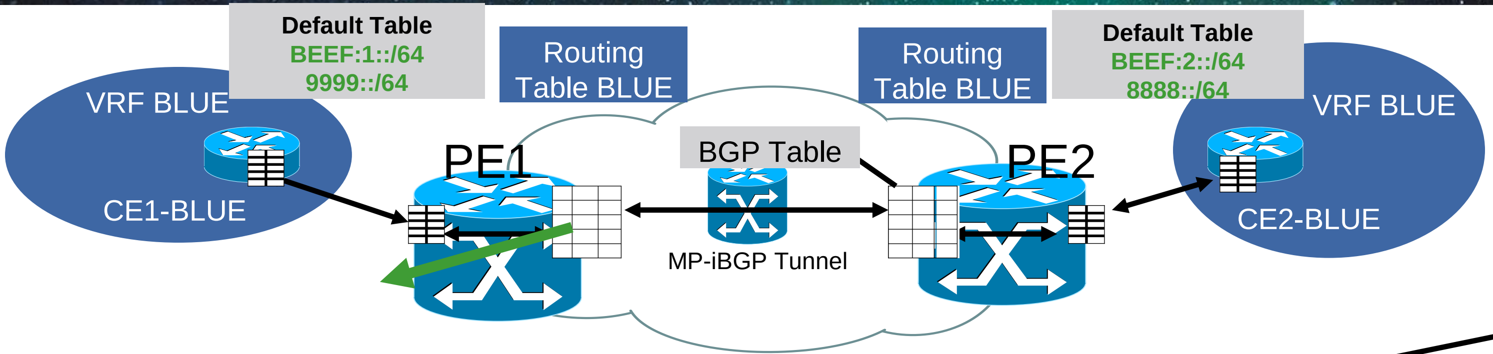

IPv4-Mapped IPv6 Address (IPv4-Based LSP Setup)

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

MPLS Forwarding

PE1

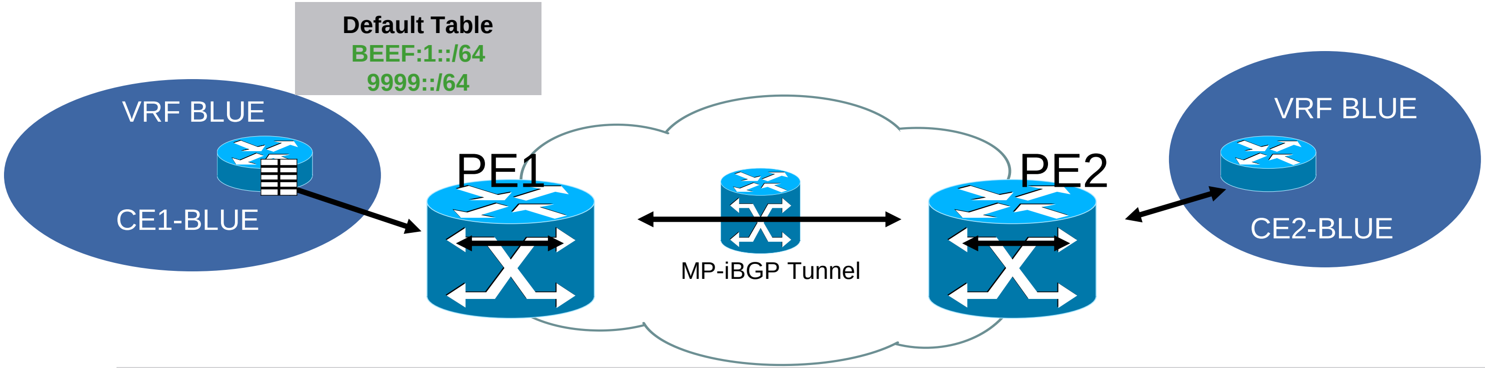


```
pe1#show mpls forwarding
Local   Outgoing     Prefix            Bytes Label   Outgoing    Next Hop
Label   Label or VC  or Tunnel Id      Switched      interface
16      Pop Label    192.168.1.4/30    0             Et2/0       192.168.1.2
17      16           192.168.1.8/30    0             Et2/0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

A Look at Forwarding

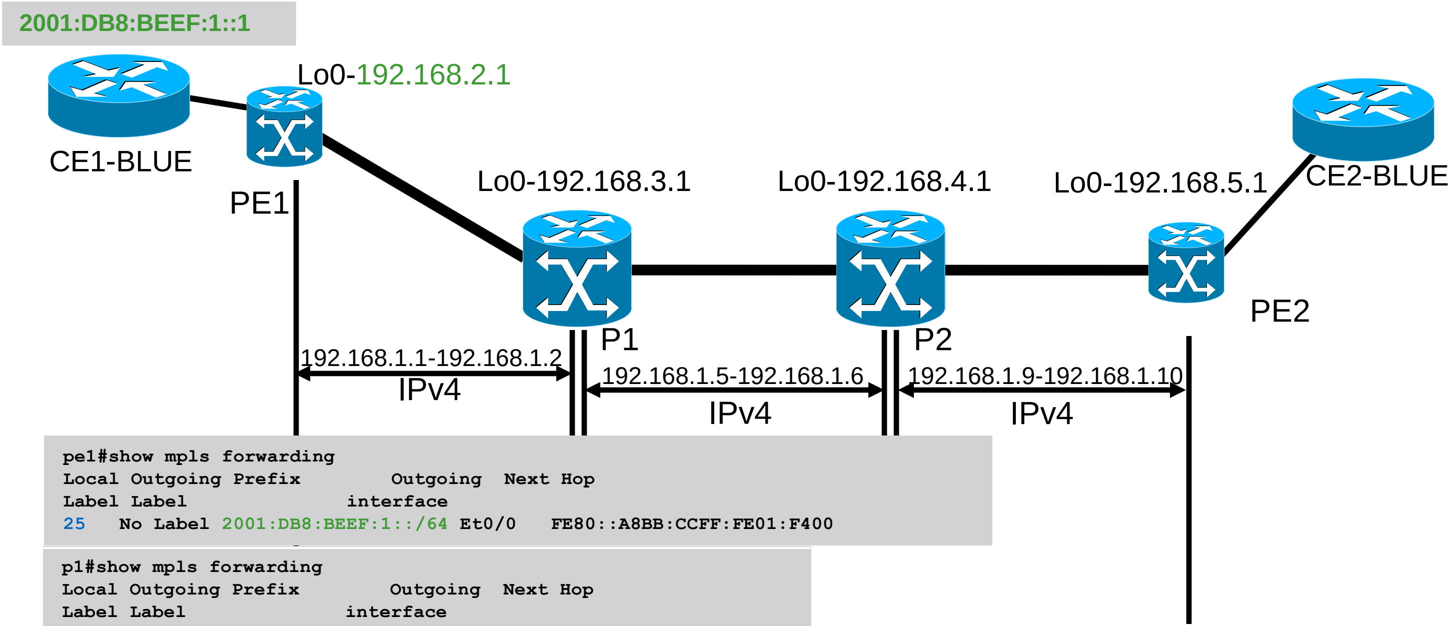


```
pe1#show mpls forwarding
Local Outgoing Prefix          Outgoing  Next Hop
Label Label                 interface
25    No Label 2001:DB8:BEEF:1::/64 Et0/0    FE80::A8BB:CCFF:FE01:F400
```

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

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

```
pe2#sh ipv6 cef vrf BLUE
2001:DB8:BEEF:1::/64
 nexthop 192.168.1.9 Ethernet0/0 label 1825
```

6VPE Summary

- RFC4659: BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
- 6VPE simply adds IPv6 support to current IPv4 MPLS VPN offering
- For end-users: v6-VPN is same as v4-VPN services (QoS, hub and spoke, internet access, etc.)
- For operators:
  - Same configuration operation for v4 and v6 VPN
  - No upgrade of IPv4/MPLS core (IPv6 unaware)
  - Consider redundancy and downtime while migrating from 6PE to 6VPE
- Cisco 6VPE Documentation: http://www.cisco.com/en/US/docs/net_mgmt/ip_solution_center/5.2/mpls_vpn/user/guide/ipv6.html




6PE/6VPE Route Reflector (RR) Considerations

Scaling 6PE and 6VPE with RR

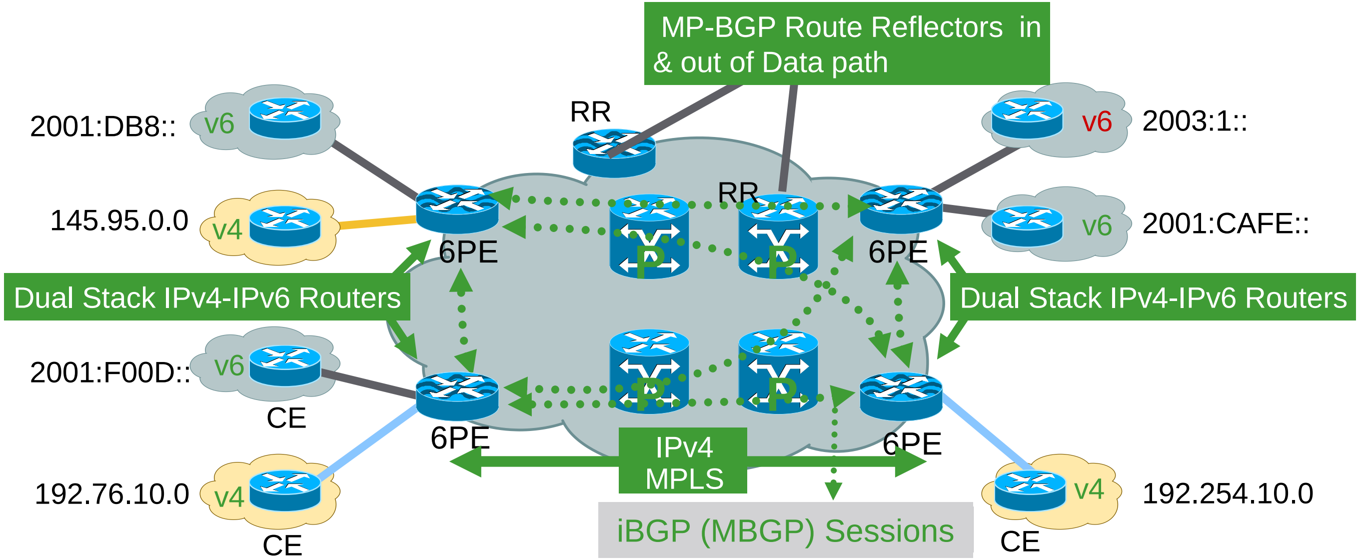

- Similar to IPv4 BGP, Route reflectors and Confederation can help improve scalability with 6PE and 6VPE deployments
- Dedicated Route Reflectors or data path RRs in BGP/6PE/6VPE deployment
- Route reflector per address family or route reflector for multiple address families
- 6PE-RR must support the "IPv6+label“ functionality
- 6VPE-RR must support the "IPv6+label" and extended communities functionality

Agenda

- SP IPv6 Integration Strategy
- IPv6 in Core Networks and Deployment Models
  - Native IPv4 Environments
  - MPLS Environments
- **IPv6 Addressing Considerations**
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- IPv6 Multi-homing Considerations
  - Goals and Solutions
  - LISP6
- Carrier-Grade IPv6 Solution – CGv6
  - CGN, 6rd, DS-Lite, AFT64, MAP
- Conclusion

BGP Route Table Today – Granular IPv4 Prefixes

- Lack of proper route aggregation
- Exponential growth of Internet
- Improper Address/subnet planning
- Allocation policies from SP and Internet registries
- Traffic engineering via routing

Provider Assigned (PA) Addresses

- SPs assigns address blocks to customers from its own address space received from its upstream provider
- Aggregation happens towards upstream provider
- Minimizing the Internet routing table size
- PA Cons:
  - Customers are locked with the SPs
  - Renumbering burden
  - Multi-homing and Traffic Engineering problems

IPv6 PA Allocation Hierarchy

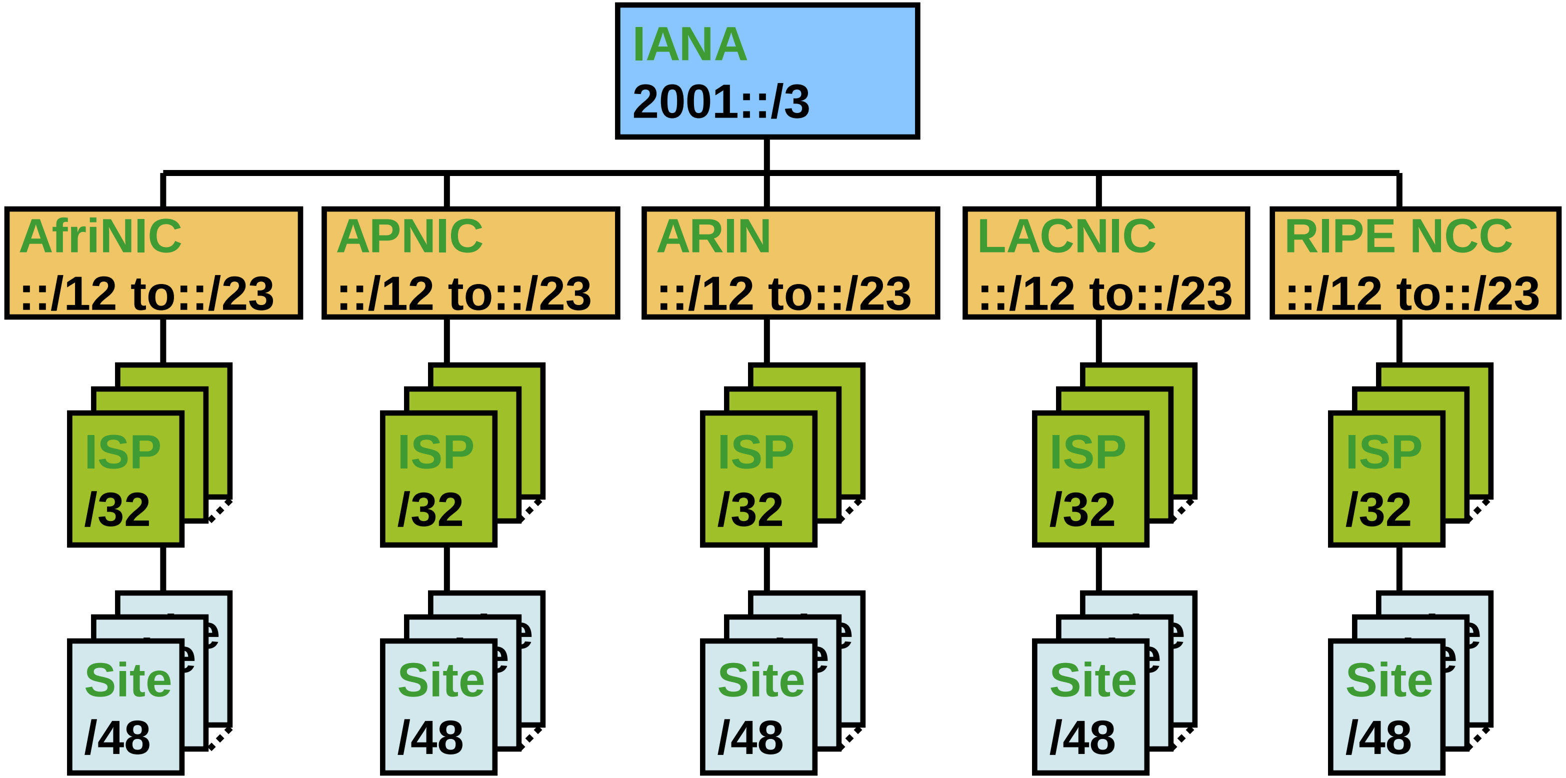

Provider Independent (PI) Addresses

- RIRs directly assigns address space to end customers
- Customers are no longer “locked” to a SP
  - No need for re-addressing
- Multi homing is straightforward
- PI Cons:
  - Larger Internet routing table due to lack of efficient aggregation
  - Memory and CPU needs on BGP speakers

IPv6 PI Allocation Hierarchy

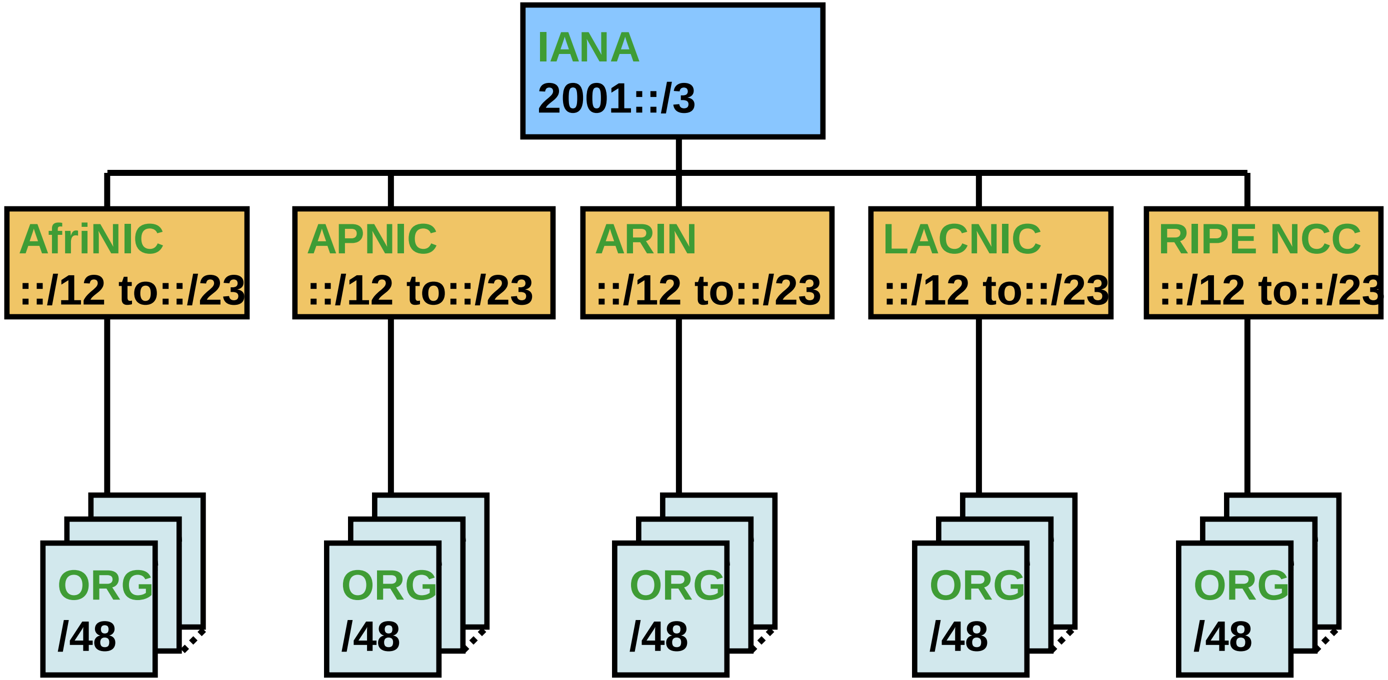

Infrastructure Addressing (ULA vs. Global)

- What type of addressing should I deploy internal to my network?
  - It depends …
- Unique Local Addresses (ULA) – FC00::/7
  - Prefix FC00::/7 is reserved by IANA for ULA (bit 8 determines if locally or centrally assigned, so ULA or ULA-Central)
- Global-only – 2000::/3
  - Recommended approach, but breaks topology hiding
- ULA + Global
  - Allows for the best of both worlds BUT at a price – much more address management with DHCP, DNS, routing and security

Link Level – Prefix Length Considerations

64 bits	< 64 bits	> 64 bits
▪ Recommended by RFC3177 and IAB/IESG ▪ Consistency makes management easy ▪ MUST for SLAAC (MSFT DHCPv6 also) ▪ Significant address space loss (18.466 Quintillion)	▪ Considered bad practice ▪ 64 bits offers more space for hosts than the media can support efficiently	▪ Address space conservation ▪ Special cases: /126—valid for p2p /127—valid for p2p /128—loopback ▪ Complicates management ▪ Must avoid overlap with specific addresses: Router Anycast (RFC3513) Embedded RP (RFC3956) ISATAP addresses

Prefix Sizing Considerations (RFC 5375)

- The IPv6 specifications prescribe a /64 prefix length for IPv6 unicast addresses.
- Using a different prefix length other than a /64 subnet in IPv6 breaks the operation of the following technologies and more:
    - Neighbour Discovery (ND)
    - Secure Neighbour Discovery (SEND) [RFC3971]
    - Privacy-extensions [RFC4941]
    - PIM-SM with Embedded-RP [RFC3956]

Prefix Allocation Practices

- Many SPs offer /48, /52, /56, /60 or /64 prefixes
- Enterprise customers receives a /48
- Large Enterprises receive one or more /48 prefixes
- Small business customers receive /52 or /56
- BB customers using DHCP-PD receives /56 or /60

Prefix Delegation

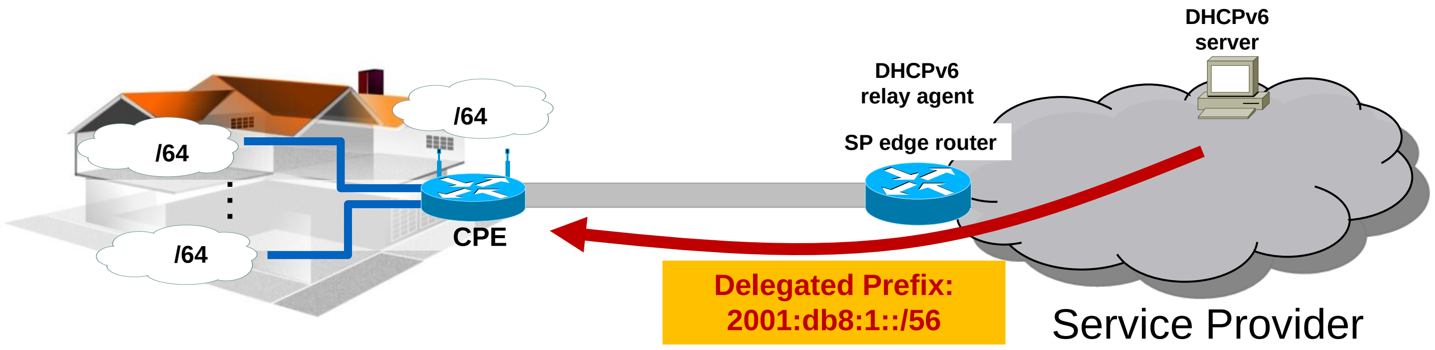

- Prefix Delegation w/ DHCPv6 PD (RFC3633)
- Prefix Allocation methods

    Local Pool or Radius

Agenda

- SP IPv6 Integration Strategy
- IPv6 in Core Networks and Deployment Models
  – Native IPv4 Environments
  – MPLS Environments
- IPv6 Addressing Considerations
  – Provider Assigned (PA) vs. Provider Independent (PI)
  – **IPv6 Addressing Assignment - Case Study**
- IPv6 Multi-homing Considerations
  – Goals and Solutions
  – LISP6
- Carrier-Grade IPv6 Solution – CGv6
  – CGN, 6rd, DS-Lite, AFT64, MAP
- Conclusion

SP IPv6 Address Assignment – Case Study

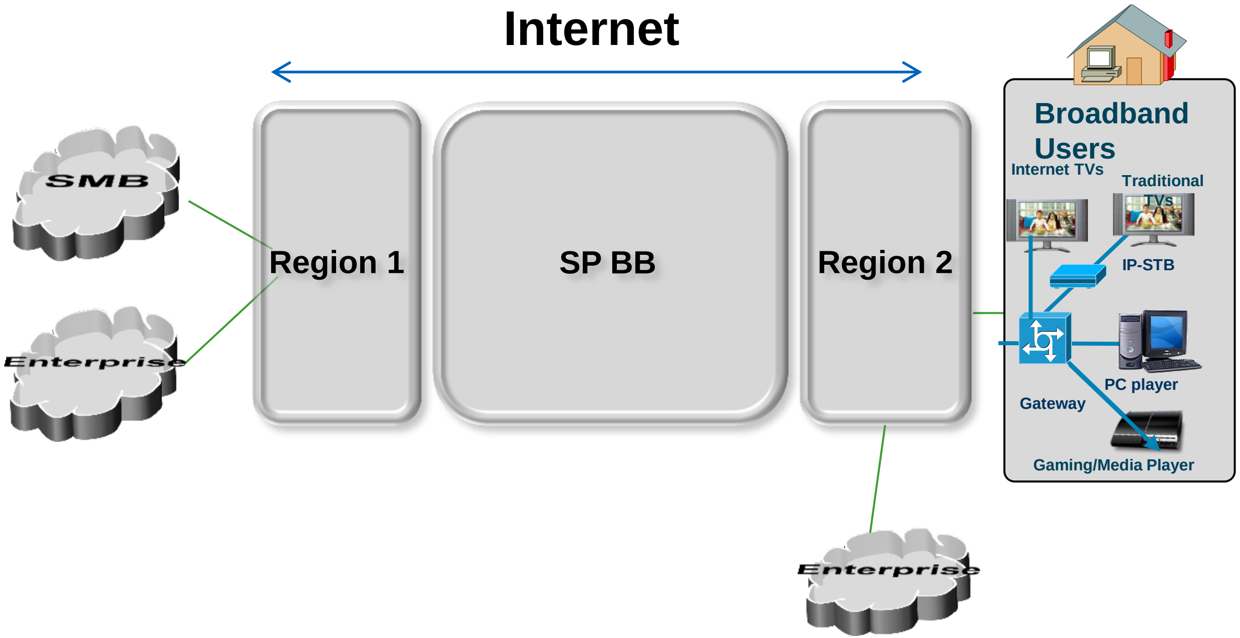

SP IPv6 Address Assignment – Case Study

- SP Addressing Needs
  - Infrastructure addresses
  - Enterprise customers
  - Small Medium Business (SMB) customers
  - Broadband customers
- Aggregation is one of the Main Concerns
  - Adequate address space buffers facilitate aggregation for possible future growth

SP IPv6 Address Assignment – Case Study

- SP receives 2001:DB8::/32 and 2001:DB9::/48 from RIR
- SP carves /32 into 16 /36 addresses

2001:db8:0000::/36	2001:db8:1000::/36
2001:db8:2000::/36	2001:db8:3000::/36
2001:db8:4000::/36	2001:db8:5000::/36
2001:db8:6000::/36	2001:db8:7000::/36
2001:db8:8000::/36	2001:db8:9000::/36
2001:db8:A000::/36	2001:db8:B000::/36
2001:db8:C000::/36	2001:db8:D000::/36
2001:db8:E000::/36	2001:db8:F000::/36

SP IPv6 Address Assignment – Case Study

- SP utilises the /48 for infrastructure needs (2001:DB9:abcd::/48)
- SP further carves the infrastructure /48 into 16 /52s and assigns one or multiple /52s for each regional network depending on the need

Allocated block	Reserved for future use
2001:db9:abcd:1000::/52(SP BB)	2001:db9:abcd:1000::/52(SP BB)
2001:db9:abcd:2000::/52(Reg 1)	2001:db9:abcd:3000::/52(Reg 1)
2001:db9:abcd:4000::/52(Reg 2)	2001:db9:abcd:5000::/52(Reg 2)
2001:db9:abcd:6000::/52	2001:db9:abcd:7000::/52
2001:db9:abcd:8000::/52	2001:db9:abcd:9000::/52
2001:db9:abcd:A000::/52	2001:db9:abcd:B000::/52
2001:db9:abcd:C000::/52	2001:db9:abcd:D000::/52
2001:db9:abcd:E000::/52	2001:db9:abcd:F000::/52

SP IPv6 Address Assignment – Case Study

SP Regional Network

- Receives a /52 for infrastructure addresses
- Receives a sequential chunk of /36s for customer address space
- Utilises a /36 for Enterprise customers (2^12 /48s)
- Reserves a subsequent /36 for Enterprise customers

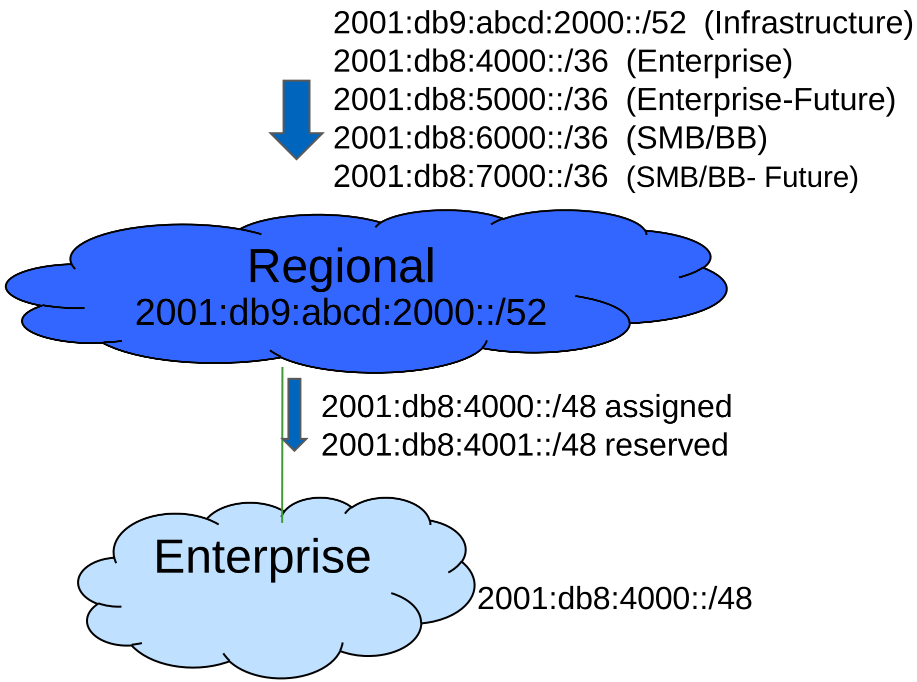

SP IPv6 Address Assignment – Case Study

SMB/Broadband Customer

- SP Reg2 uses a /36 for SMB and BB customer space (2^20 /56s or 2^24 /60s)
- SMB receives a /56 and allocates /64 prefix for each network segment
- BB user gets a /56 or /60 PD prefix and allocates /64 per network segment:
  2001:db8:6000:1000::/60
  16 /64 addresses per BB user

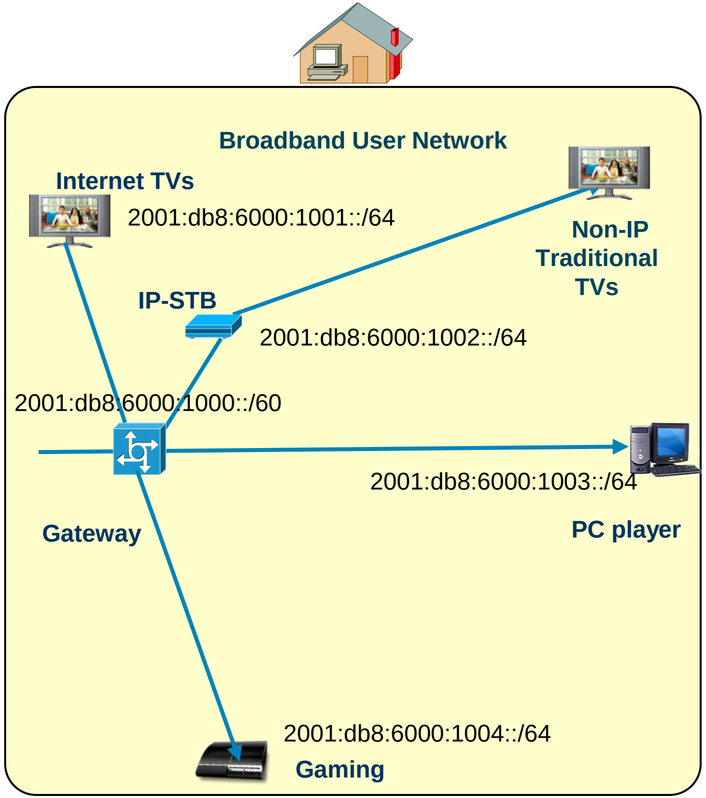

Agenda

- SP IPv6 Integration Strategy
- IPv6 in Core Networks and Deployment Models
  - Native IPv4 Environments
  - MPLS Environments
- IPv6 Addressing Considerations
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- **IPv6 Multi-homing Considerations**
  - Goals and Solutions
  - LISP6
- Carrier-Grade IPv6 Solution – CGv6
  - CGN, 6rd, DS-Lite, AFT64, MAP
- Conclusion

Multi-Homing Goals

- Redundancy: Insulate from failures in upstream providers
- Load Sharing: Concurrent use of multiple transit provider and distribute traffic load
- Simple and scalable solution
- Minimal impact on existing network devices

IPv6 Multi-Homing Solutions

- Traditional Multi-homing (PI)
  - –Advertise address space to multiple transit providers
  - –Longer prefixes to prefer a ISP and/or Load balancing
  - –Large BGP routing table
- Dual Address blocks from upstream ISPs (PA)
  - –Receive and utilise address blocks from both ISPs
  - –Overhead and renumbering problems
- NPTv6/Proxy RFC 6296
  - –Address independence without PI
  - –NAT again?
- Multi-homing without NAT or PI
  - –SHIM6 – Host based (RFC 5533)
  - –LISP – Network based

Network Based Multi-homing with Locator/ID Separation Protocol (LISP)

- LISP decouples Internet addresses into EIDs/RLOCs.
- Network-based map-n-encap protocol implemented mostly on network edge routers
- Reduces the size and dynamic properties of the core routing tables
  - LISP can also be used as IPv6 transition mechanism
  - LISP Attributes:
  - **EID (Endpoint Identifier) is the IP address of a host – just as it is today**
  - **RLOC (Routing Locator) is the IP address of the LISP router for the host**
  - **EID-to-RLOC mapping is the distributed architecture that maps EIDs to RLOCs**
  - LISP Advantages:
    - Network-based solution
    - Address family agnostic
    - No host/DNS changes
    - No new addressing to site devices; minimal configuration changes
    - Incrementally deployable; interoperable with existing Internet

LISP: Network based Multi-Homing Solution

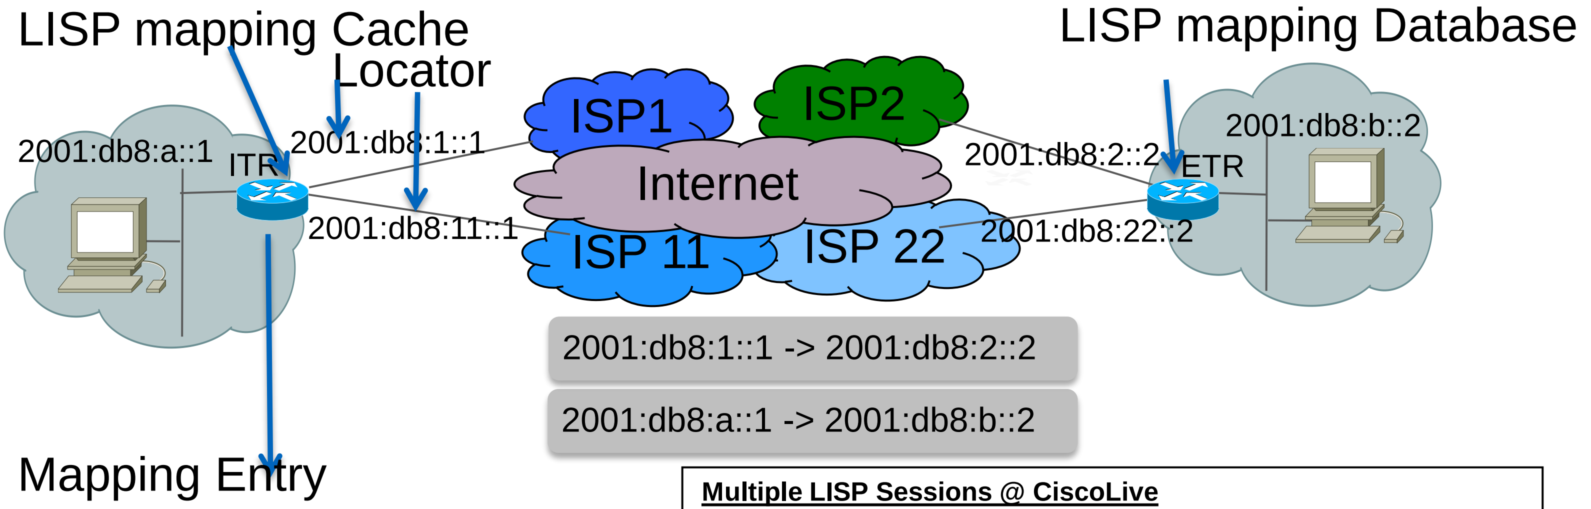

EID-prefix: 2001:db8:b::/48

Locator-set:

2001:db8:2::2, priority: 1, weight: 50 (D1)

2001:db8:22:2, priority: 1, weight: 50 (D2)

Multiple LISP Sessions @ CiscoLive

BRKRST-3045 LISP - A Next Generation Networking Architecture
LABCRS-2100 Deploying LISP in the Enterprise Data Centre
LTRCRS-2005 Deploying Multi-Tenant VPN using LISP
TECIPM-3191 Advanced - LISP Techtorial
BRKRST-3046 Advanced LISP – What's in for me?

Agenda

- SP IPv6 Integration Strategy
- IPv6 in Core Networks and Deployment Models
  - Native IPv4 Environments
  - MPLS Environments
- IPv6 Addressing Considerations
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- IPv6 Multi-homing Considerations
  - Goals and Solutions
  - LISP6
- **Carrier-Grade IPv6 Solution – CGv6**
  - CGN, softwires, 6rd, DS-Lite, AFT64, MAP
- Conclusion

Carrier-Grade IPv6 Solutions – CGv6

- CGN
- Softwires
- 6rd
- DS-Lite
- AFT64
- MAP

Public IPv4 Exhaustion with CGN Solution

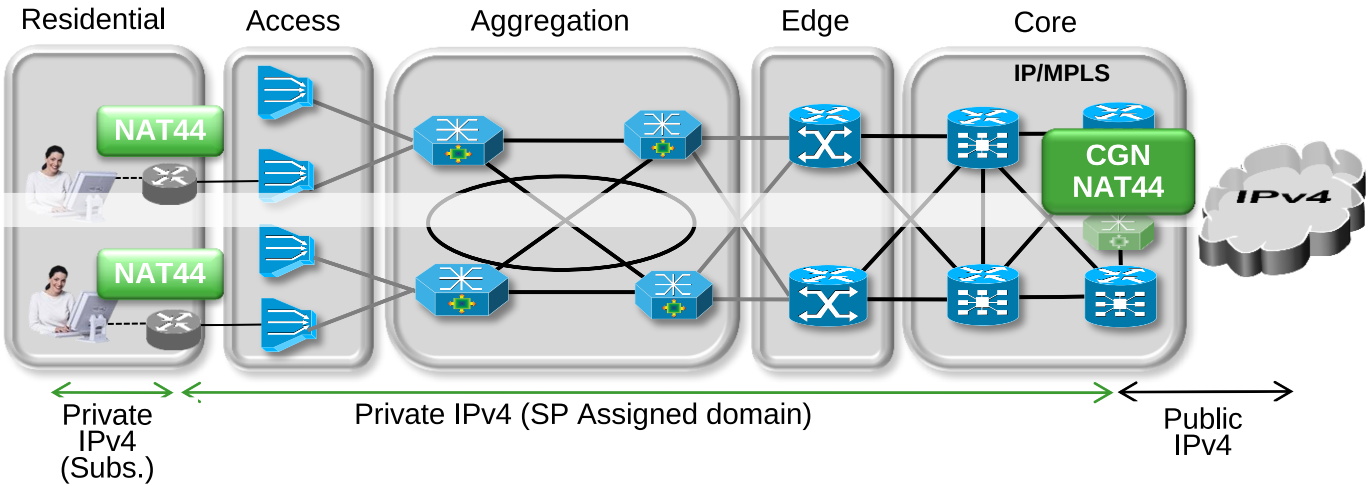

- Short-term solution to public IPv4 exhaustion issues without any changes on RG and SP Access/Aggregation/Edge infrastructure
- Subscriber uses NAT44 (i.e. IPv4 NAT) in addition to the SP using CGN with NAT44 within its network
- CGN NAT44 multiplexes several customers onto the same public IPv4 address
- CGN performance and capabilities should be analysed in planning phase
- Long-term solution is to have IPv6 deployed

IPv6 over L2TP Softwires

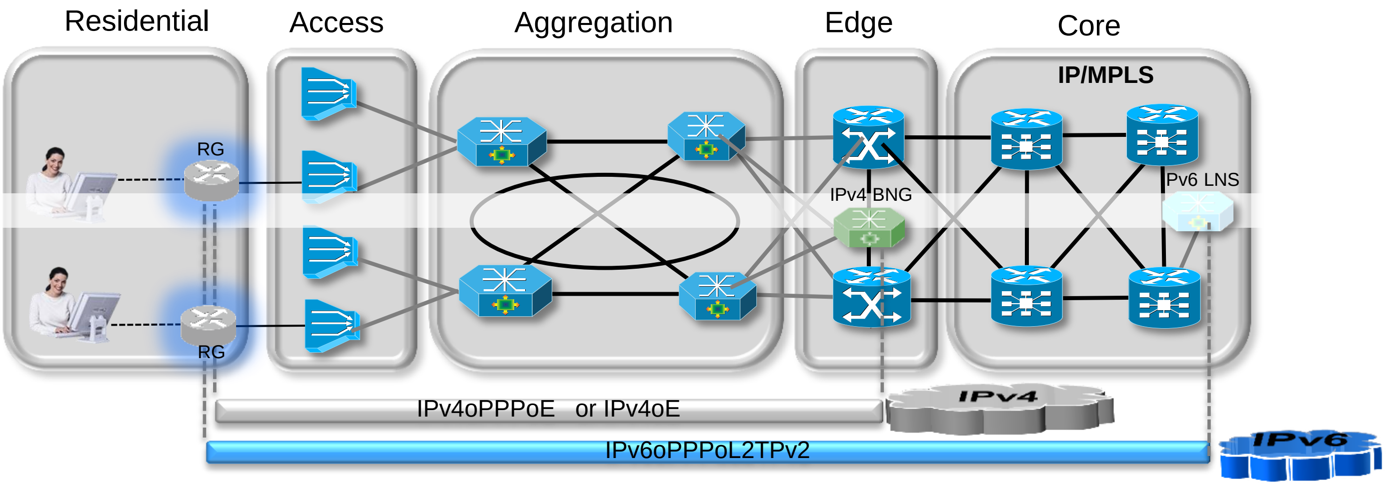

- Dual-Stack IPv4/IPv6 service on RG LAN side
- PPPoE or IPv4oE Termination on IPv4-only BNG
- L2TPv2 softwire between RG and IPv6-dedicated L2TP Network Server (LNS)
- Stateful architecture on LNS, offers dynamic control and granular accounting of IPv6 traffic
- Limited investment & impact on existing infrastructure

Broadband Forum WT-242: Getting to Dual Stack

IPv6 over IPv4 via 6rd (RFC 5569)

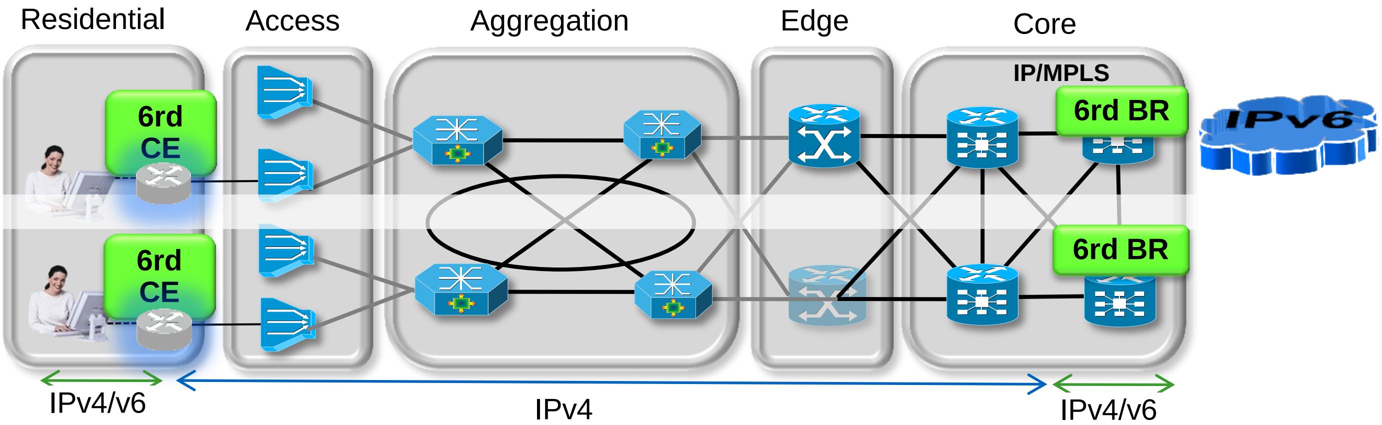

- Introduction of two Components: 6rd CE (Customer Edge) and 6rd BR (Border Relay)
- Automatic Prefix Delegation on 6rd CE
- Simple, stateless, automatic IPv6-in-IPv4 encap and decap functions on 6rd (CE & BR)
- IPv6 traffic automatically follows IPv4 Routing
- 6rd BRs addressed with IPv4 anycast for load-balancing and resiliency
- Limited investment & impact on existing infrastructure

IPv4 via IPv6 Using DS-Lite (w/NAT44)

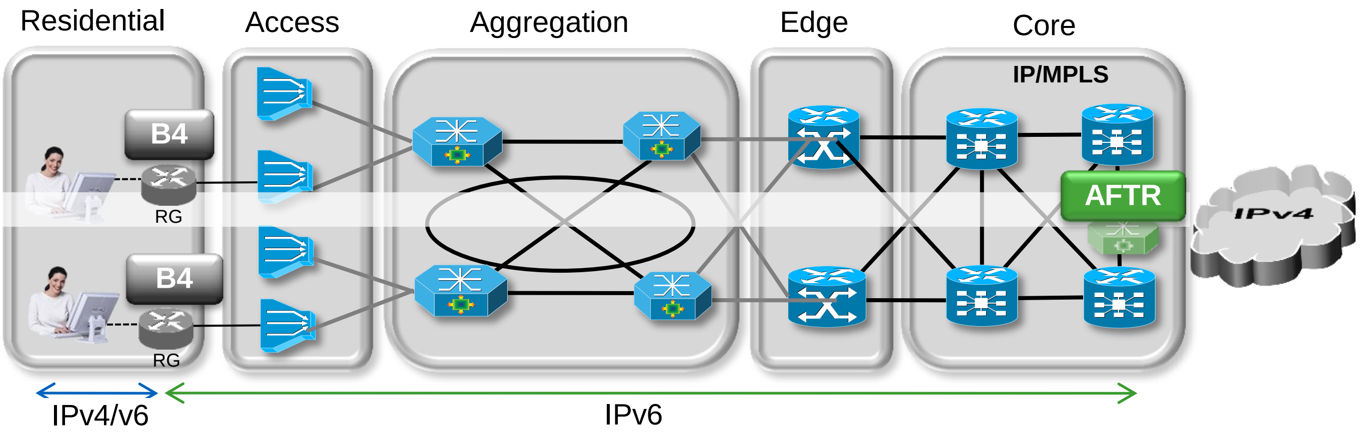

- Access, Aggregation, Edge and Core migrated to IPv6. NMS/OSS and network services migrated to IPv6 as well (DNS, DHCP)
- IPv4 Internet service still available and overlaid on top of IPv6-only network.
- Introduction of two Components: B4 (Basic Bridging Broadband Element) and AFTR (Address Family Transition Router)
  - B4 typically sits in the RG
  - AFTR is located in the Core infrastructure
- <u>Assumption:</u> IPv4 has been phased out, IPv6 only Access/aggregation network

Connecting IPv6-Only with IPv4-Only: AFT64

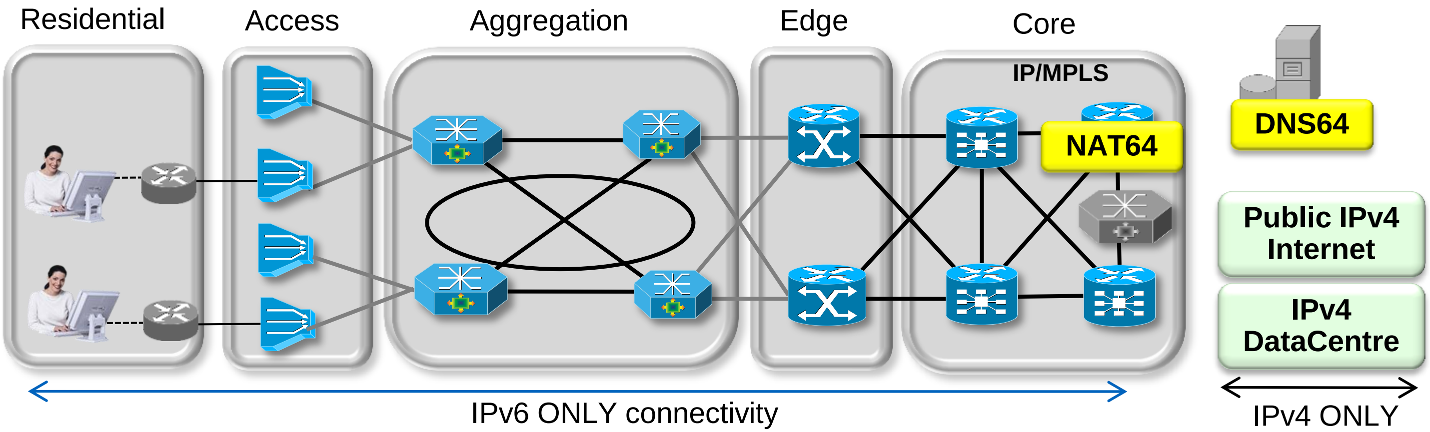

- AFT64 technology is only applicable in case where there are IPv6 only end-points that need to talk to IPv4 only end-points (AFT64 for going from IPv6 to IPv4)
- AFT64:= “stateful v6 to v4 translation” or “stateless translation”, ALG still required
- Key components includes NAT64 and DNS64
- <u>Assumption:</u> Network infrastructure and services have fully transitioned to IPv6 and IPv4 has been phased out

Mapping of Address/Port using Translation (MAP-T)

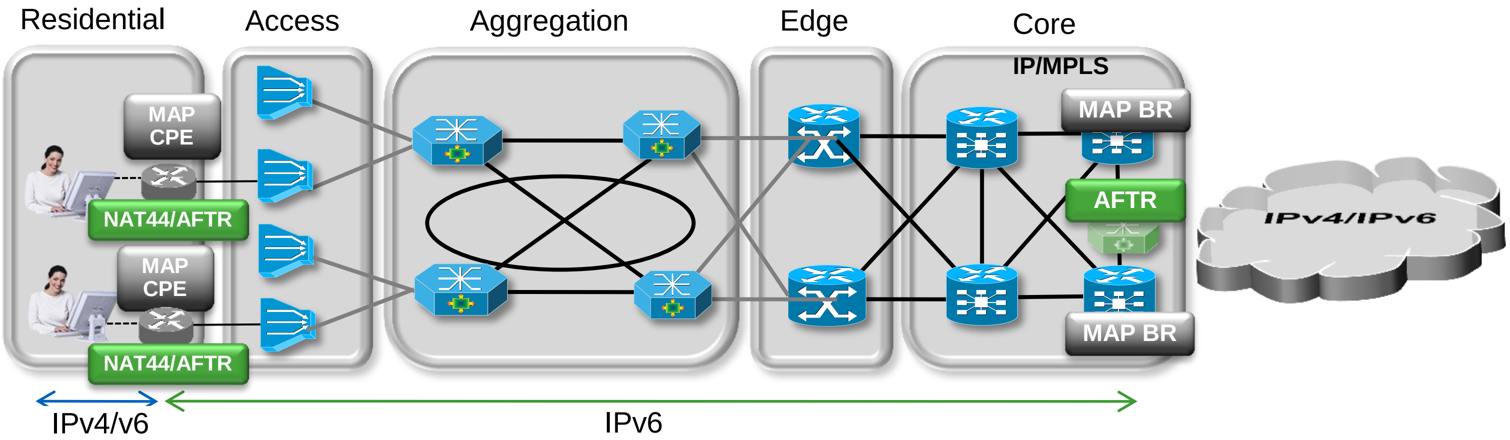

- Mapping algorithm for port indexing and IPv4<->IPv6 address
- Stateless: better scalability/no redundancy issues
- MAP CPE: NAT46 to- and from-core
- MAP stateless Border Relay (MAP BR)
- Backwards compatible with statefull NAT64
- CPE uses combination of IPv6 prefix + DHCPv6 MAP Option to derive CPE's IPv4 address + port set id
- CPE derives its TCP/UDP port range by decoding the port-set id using MAP algorithm

Mapping of Address/Port using Encapsulation (MAP-E)

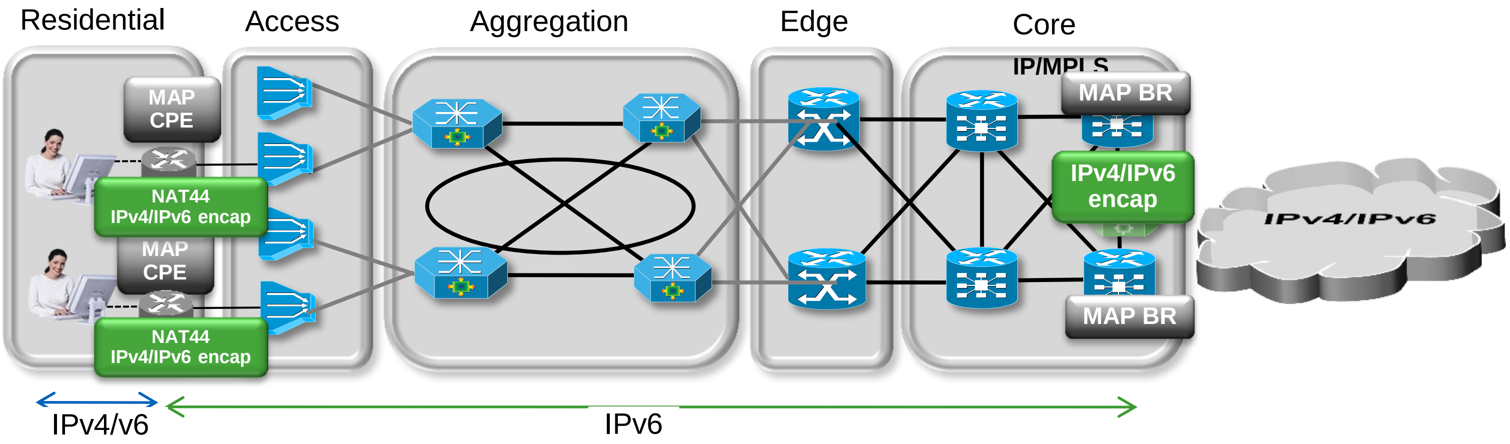

- Mapping algorithm for port indexing and IPv4<->IPv6 address
- Stateless: better scalability/no redundancy issues
- MAP CPE: IPv4 to IPv6 encap/decap to- and from-core
- MAP stateless Border Relay (MAP BR)
- Backwards compatible with DS-Lite AFTR core
- CPE uses combination of IPv6 prefix + DHCPv6 MAP Option to derive CPE's IPv4 address + port set id
- CPE derives its TCP/UDP port range by decoding the port-set id using MAP algorithm

Agenda

- SP IPv6 Integration Strategy
- IPv6 in Core Networks and Deployment Models
  - Native IPv4 Environments
  - MPLS Environments
- IPv6 Addressing Considerations
  - Provider Assigned (PA) vs. Provider Independent (PI)
  - IPv6 Addressing Case Study
- IPv6 Multi-homing Considerations
  - Goals and Solutions
  - LISP6
- Carrier-Grade IPv6 Solution – CGv6
  - CGN, 6rd, DS-Lite, AFT64, MAP
- **Conclusion**

Conclusion

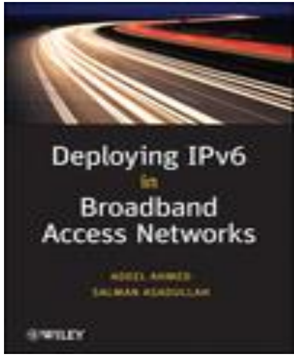

- Start now rather than later:
  - Multiple technology adoption scenarios available!!!
  - Purchase for the future and test, test and then test some more
  - Start moving legacy application towards IPv6 support
  - Don’t assume your favorite vendor/app/gear has an IPv6 plan
  - Full parity between IPv4 and IPv6 is still a ways off
- Deploying IPv6 in Broadband Access Networks - Adeel Ahmed, Salman Asadullah, John Wiley & Sons Publications®
- Deploying IPv6 Networks – Ciprian P., Patrick G., Eric L., Cisco Press®
- IPv6 Security - Scott Hogg, Eric Vyncke, Cisco Press®
- IPv6 for Enterprise Networks - Shannon McFarland et. al, Cisco Press®
- www.cisco.com/go/ipv6 - CCO IPv6 Main Page
- www.cisco.com/go/srnd - Cisco Network Design Central
- www.ietf.org
- www.ipv6forum.org

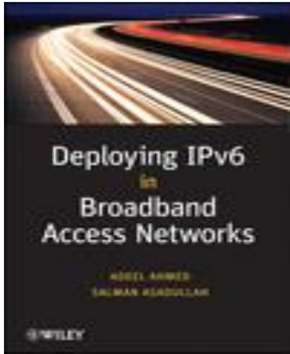

Cisco IPv6 Book Authors

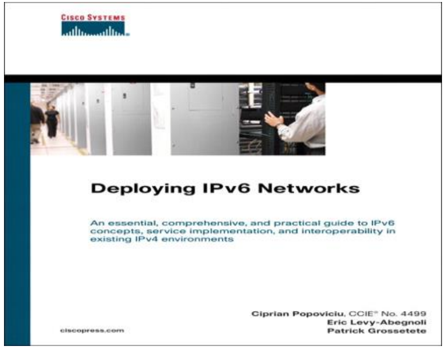

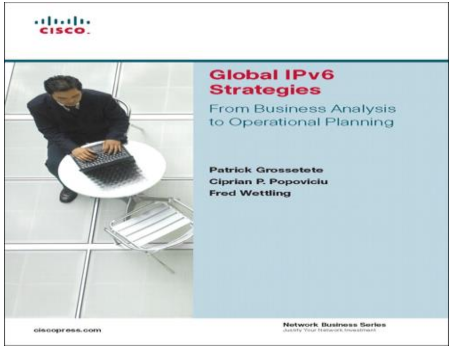

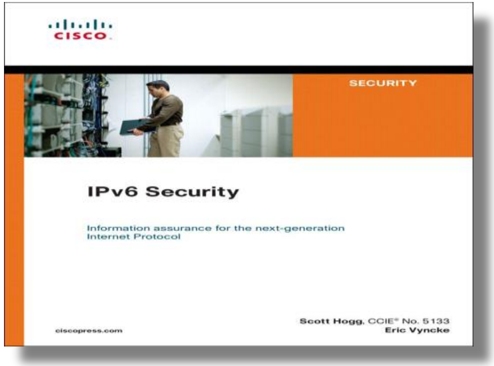

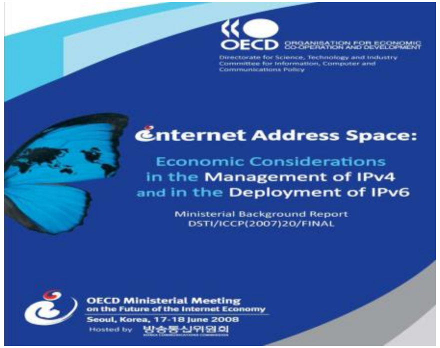

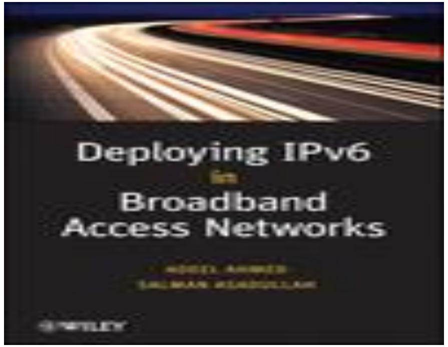

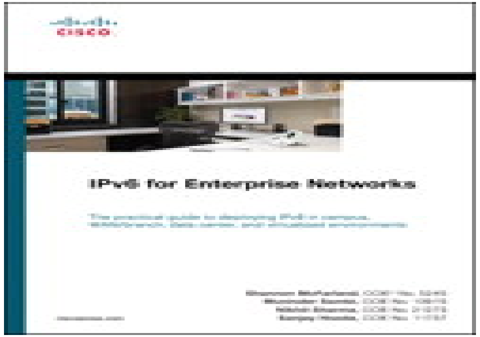

Join Cisco IPv6 Support Community!

- **Free** for anyone with Cisco.com registration
- Get **timely** answers to your technical questions
- Find **relevant** technical documentation
- Engage with over 200,000 **top technical experts**
- **Seamless** transition from discussion to TAC Service Request *(Cisco customers and partners only)*
- Visit the Cisco Support Community booth in the World of Solutions for more information

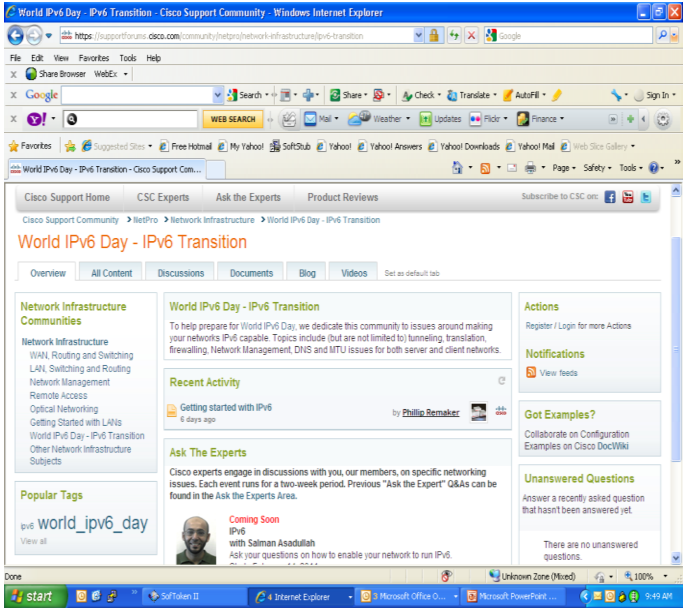

The Cisco Support Community is your one-stop community destination from Cisco for sharing current, real-world technical support knowledge with peers and experts.

supportforums.cisco.com
supportforums.cisco.mobi

Cisco live!




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

CISCO™