Tutorial: IPv6 Technology Overview Part II

Speaker: Byju Pularikkal, Cisco Systems, Inc

Date: 11th November 2011

Tutorial-1: Agenda

- Structure of IPv6 Protocol
  - IPv4 and IPv6 Header Comparison
  - IPv6 Extension Headers
- IPv6 Addressing
  - Addressing Format
  - Types of IPv6 addresses
- ICMPv6 and Neighbor Discovery
  - Router Solicitation & Advertisement
  - Neighbor Solicitation & Advertisement
  - Duplicate Address Detection
- Multicast in IPv6
- DHCP & DNS for IPv6
  - DNS with IPv6
  - DHCPv6 Overview

Tutorial-2: Agenda

- Routing in IPv6
  - RIPng
  - OSPFv3
  - IS-IS for IPv6
  - BGP-4 Extensions for IPv6
- IPv6 Transition Mechanisms
  - Dual Stack Technique
  - 6 to 4 Tunneling
  - ISATAP
  - 6RD
  - Dual-stack Lite
  - 6PE
  - 6VPE

IPv6 Routing

Dynamic Routing in IPv6

- Routing in IPv6 is unchanged from IPv4
  - Still has two families of routing protocols: IGP and EGP
  - Still uses the longest-prefix match routing algorithm
- IGP
  - RIPng (RFC 2080)
  - Integrated IS-IS for IPv6 (RFC 5308)
  - OSPFv3 (RFC 5340)
- EGP
  - MP-BGP4 (RFC 4760) and Using MP-BGP for IPv6 (RFC 2545)

RIPng (RFC 2080)

RIP Next Generation (RIPng)

- Referred to as RIP Next Generation, distance vector protocol
- For the SP industry not recommended, limited use in Enterprise environments
- ISPs do not use RIP in any form unless there is absolutely no alternative
- RIPng was used in the early days of the IPv6 test network
  - Superior routing protocols such as ISIS, OSPF and BGP rapidly replaced RIPng

RIPng Overview

- Similar to RIPv2
  - Distance-vector, Hop limit of 15, split-horizon,
  - All RIP routers is FF02::9, UDP port (521)
- Updated features for IPv6
  - Prefix length added, address-family and subnet mask fields removed
- Special Handling for the NH
  - One NH entry per group of prefixes

RIPng Header Comparison

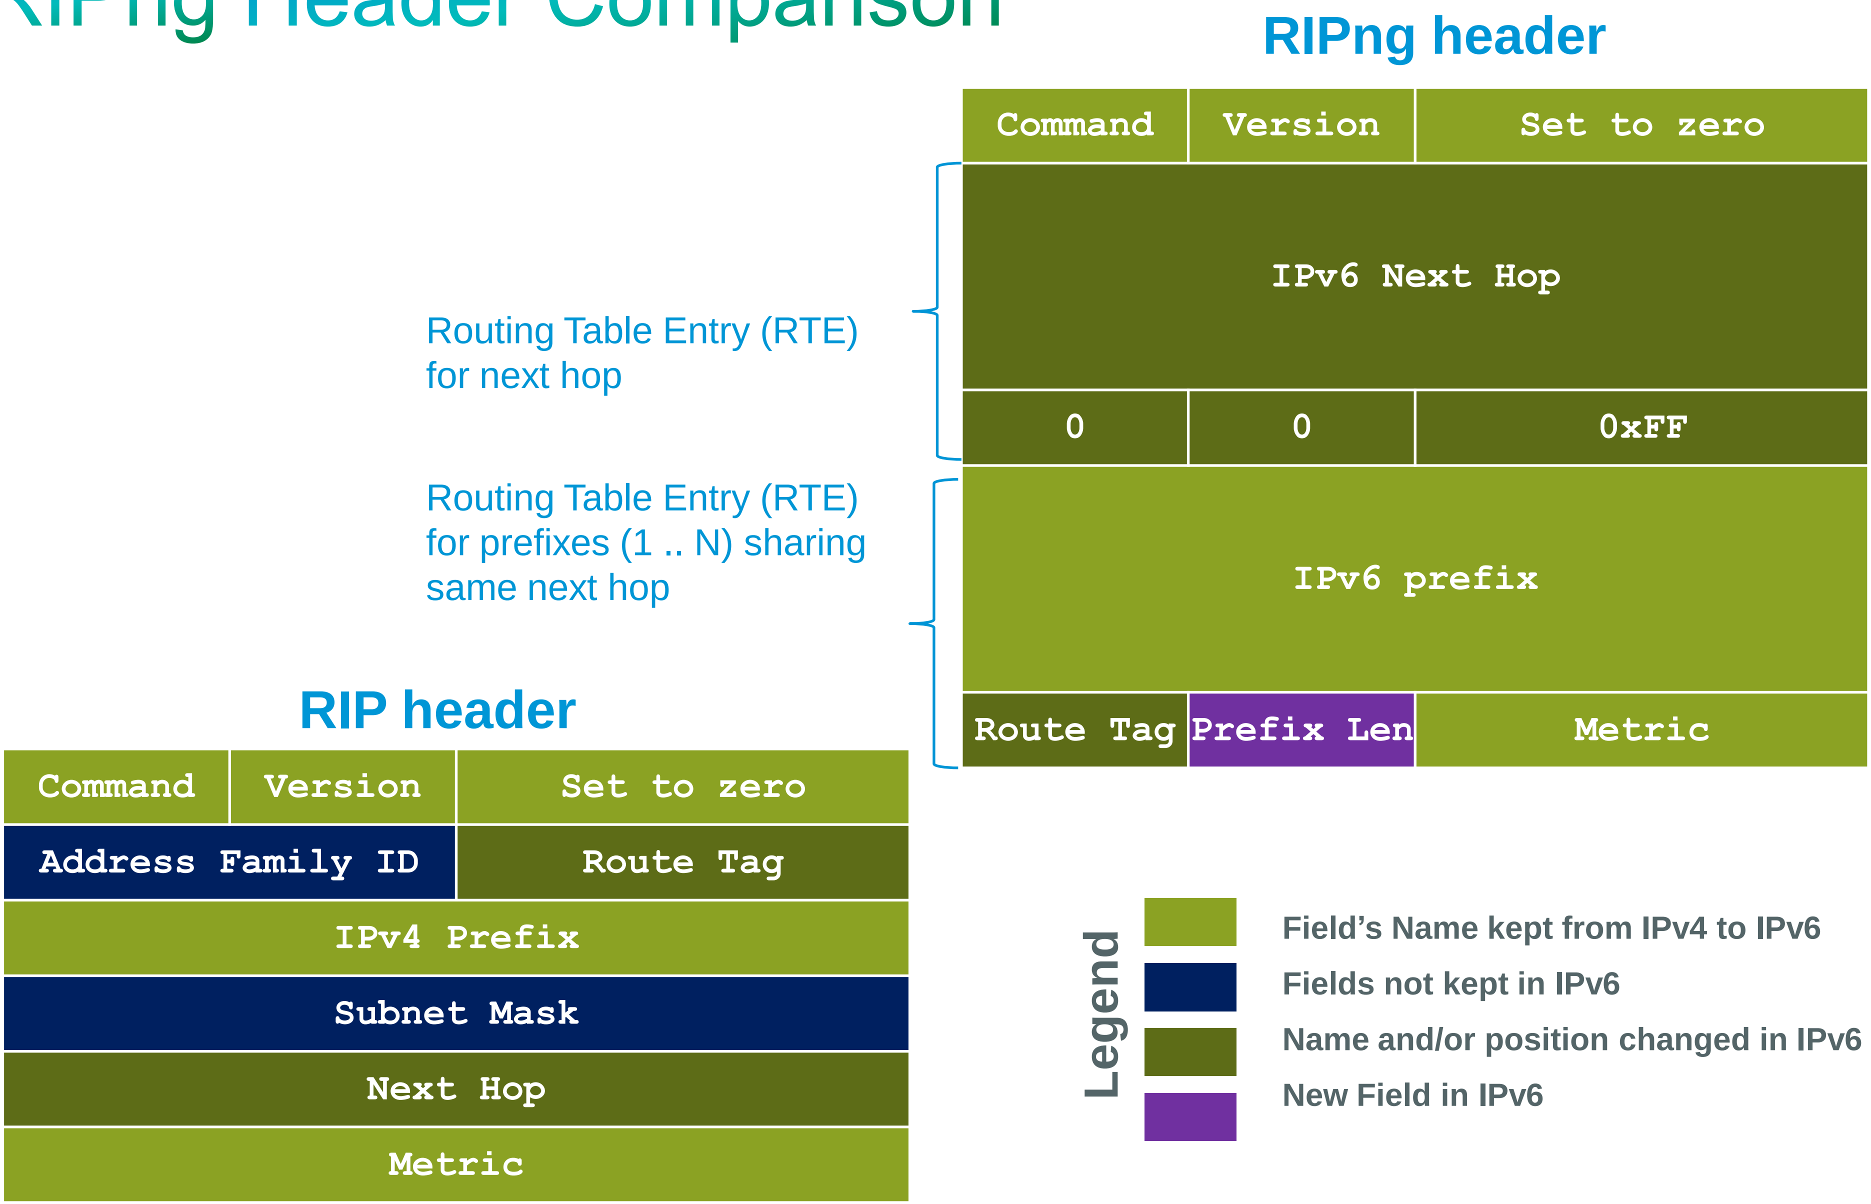

OSPFv3 (RFC 2740)

OSPFv3 Overview

- OSPFv3 is OSPF for IPv6 (RFC 5340)
- Based on OSPFv2 with enhancements
- Distributes IPv6 prefixes only
- Runs directly over IPv6
- Ships-in-the-night with OSPFv2

OSPFv3 Overview

- OSPFv3 has same 5 packet types some fields have been changed

Packet Type	Description
1	Hello
2	Database description
3	Link state request
4	Link state update
5	Link state acknowledgement

OSPFv3 Differences from OSPFv2

- OSPFv3 packets have a 16 byte header versus the 24 byte header in OSPFv2

OSPFv2

Version	Type	Packet Length
Router ID		
Area ID		
Checksum		Authtype
Authentication		
Authentication		

OSPFv3

Version	Type	Packet Length	
Router ID			
Area ID			
Checksum		Instance ID	0

Legend

- Field's Name kept from IPv4 to IPv6
- Fields not kept in IPv6
- Name and/or position changed in IPv6
- New Field in IPv6

OSPFv3 Differences from OSPFv2

- Uses link local addresses
  - To identify the OSPFv3 adjacency neighbors
- Two New LSA Types
  - Link-LSA (LSA Type 0x2008)
  - There is one Link-LSA per link. This LSA advertises the router's link-local address, list of all IPv6 prefixes and options associated with the link to all other routers attached to the link
  - Intra-Area-Prefix-LSA (LSA Type 0x2009)
  - Carries all IPv6 prefix information that in IPv4 is included in Router-LSAs and Network-LSAs
- Two LSAs are renamed
  - Type-3 summary-LSAs, renamed to “Inter-Area-Prefix-LSAs”
  - Type-4 summary LSAs, renamed to “Inter-Area-Router-LSAs”

OSPFv3 Differences from OSPFv2

- Multicast Addresses
  - FF02::5 – Represents all SPF routers on the link local scope, Equivalent to 224.0.0.5 in OSPFv2
  - FF02::6 – Represents all DR routers on the link local scope, Equivalent to 224.0.0.6 in OSPFv2
- Removal of Address Semantics
  - IPv6 addresses are no longer present in OSPF packet header (Part of payload information)
  - Router LSA, Network LSA do not carry IPv6 addresses
  - Router ID, Area ID and Link State ID remains at 32 bits
  - DR and BDR are now identified by their Router ID and no longer by their IP address
- Security
  - OSPFv3 uses IPv6 AH & ESP extension headers instead of variety of mechanisms defined in OSPFv2

OSPFv3 LSA Types

LSA Description	LSA Code	LSA Type	Bits Set=1
Router LSA	1	0x2001	S1
Network LSA	2	0x2002	S1
Inter-Area-Prefix-LSA	3	0x2003	S1
Inter-Area-Router-LSA	4	0x2004	S1
AS-External-LSA	5	0x4005	S2
Deprecated	6	0x2006	S1
NSSA-LSA	7	0x2007	S1
Link-LSA	8	0x0008	
Intra-Area-Prefix-LSA	9	0x2009	S1

S2	S1	Flooding Scope
0	**0**	Link-Local Scoping - Flooded only on originating link
0	**1**	Area Scoping - Flooded only in originating area
1	**0**	AS Scoping - Flooded throughout AS
1	**1**	Reserved

LSA Type Format

U Bit	LSA Handling
0	Treat the LSA as if it had link-local flooding scope
1	Store and flood the LSA as if the type is understood

ISIS for IPv6

IS-IS Standards History

- IETF ISIS for Internets Working Group
- ISO 10589 specifies OSI IS-IS routing protocol for CLNS traffic
  - Type/Length/Value (TLV) options to enhance the protocol
  - A Link State protocol with a 2 level hierarchical architecture
- RFC 1195 added IP support, also known as Integrated IS-IS (I/IS-IS)
  - I/IS-IS runs on top of the Data Link Layer
  - Requires CLNP to be configured
- RFC5308 adds IPv6 address family support to IS-IS
- RFC5120 defines Multi-Topology concept for IS-IS
  - Permits IPv4 and IPv6 topologies which are not identical
  - Allows gradual roll out of IPv6 across backbone without impacting IPv4

Integrated IS-IS for IPv6

- 2 new TLVs added to introduce IPv6 routing
  - ➢IPv6 Reachability TLV (0xEC) – 236
    - ✓ Describes network reachability such as IPv6 routing prefix, metric information and some option bits.
    - ✓ Option bits indicates advertisement of IPv6 prefix from a higher level, redistribution from other protocols
    - ✓ Equivalent to IP Internal/External Reachability TLV's described in RFC1195 – TLV 128 & 130
  - ➢IPv6 Interface Address TLV (0xE8) - 232
    - ✓ Contains 128 bit address
    - ✓ For Hello PDUs, must contain the link-local address (FE80::/10)
    - ✓ For LSP, must only contain the non link-local address
- A new Network Layer Protocol Identifier (NLPID) is defined in TLV 129
  - ➢Allowing IS-IS routers with IPv6 support to advertise IPv6 prefix payload using 0x8E value
  - ➢IPv4 & OSI uses 0xCC, CLNP is 0x81

Single Topology IS-IS

- A single SPF runs per level for OSI, IPv4 and IPv6
  - All routers in an area must run the same set of protocols [IPv4-only, IPv6-only, IPv4-IPv6]
  - L2 routers may not be configured similarly but no routing hole must exist

Single Topology SPF rules

- If IS-IS is used for both IPv4 and IPv6 in an area
  - Then both protocols must support the same topology within this area
- All interfaces configured with IS-ISv6 must support IPv6
- All interfaces configured with IS-IS for both protocols must support both of them
  - IPv6 configured tunnel (protocol type = 41) will not work, GRE should be used in this configuration
- Otherwise, consider Multi-Topology IS-IS (separate SPF)

Single SPF IS-IS for IPv6 restrictions

- IS-IS for IPv6 uses the same SPF for both IPv4 and IPv6
  - ➢Not suitable for an existing IPv4 IS-IS network where customer wants to turn on scattered IPv6 support
- If using IS-IS for both IPv4 and IPv6 then the IPv4 and IPv6 topologies MUST match exactly
  - ➢You cannot run IS-IS IPv6 on some interfaces, IS-IS IPv4 on others.
- ISIS will only form adjacencies with similarly-configured routers
  - ➢An IS-IS IPv6-only router will not form an adjacency with an IS-IS IPv4/IPv6 router
    - ✓Exception is over L2-only interface
- Cannot join two IPv6 areas via an IPv4-only area
  - ➢L2 adjacencies will form OK but IPv6 traffic will black-hole in the IPv4 area

Multi-Topology IS-IS (MT IS-IS)

- IS-IS for IPv6 assumes that the IPv6 topology is the same as the IPv4 topology
  - Single SPF running, multiple address families
  - Some networks may follow this rule, others may not
- Multi-Topology IS-IS solves the problem of different IPv4 and IPv6 Topologies
  - New TLV attributes introduced
  - New Multi-Topology ID #2 for IPv6 Routing Topology
- Router maintains two topologies
  - ISO/IPv4 Routing Topology
  - IPv6 Routing Topology

Multi-Topology IS-IS extensions

- There are several new TLVs attributes for Multi-Topology extensions
  - Multi-topology TLV contains one or more multi-topology ID in which the router participates. 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
  - Multi-Topology Reachable IPv4 Prefixes TLV this TLV appears as many times as the number of IPv4 announced by an IS for a given 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
- Multi-Topology ID Values
  - MT ID #0 – standard topology for IPv4/CLNS
  - MT ID #2 – IPv6 Routing Topology

Multi-Topology IS-IS Restrictions

- Not compatible with the previous single SPF model
  - ➢New TLV are used to transmit and advertise IPv6 capabilities.
  - ➢All routers that run IS-IS for IPv6 need to enable multi-topology within the network
- IPv4, IPv6, IPv4/IPv6 may be configured on the interface for either Level-1, Level-2 or Level-1-2
  - ➢If IPv4 and IPv6 are configured on the same interface, then the IS-IS level must match
  - ➢For example, IPv4 cannot be be Level-1 while IPv6 is configured to run ISIS level-2 on the same interface

Multi-Topology IS-IS Restrictions (Contd..)

- All routers on a LAN or point to point must have at least one common supported topology (IPv4 or IPv6) when operating in Multi-Topology IS-IS mode
  - A router that is not operating in MT IS-IS IPv6 mode cannot form adjacency with MT IS-IS IPv6 router,. However, if IPv4 is the common topology between two routers, an adjacency should be formed.
- Wide metric is required to be enabled globally on the Autonomous System to run

Multi-Topology IS-IS Example

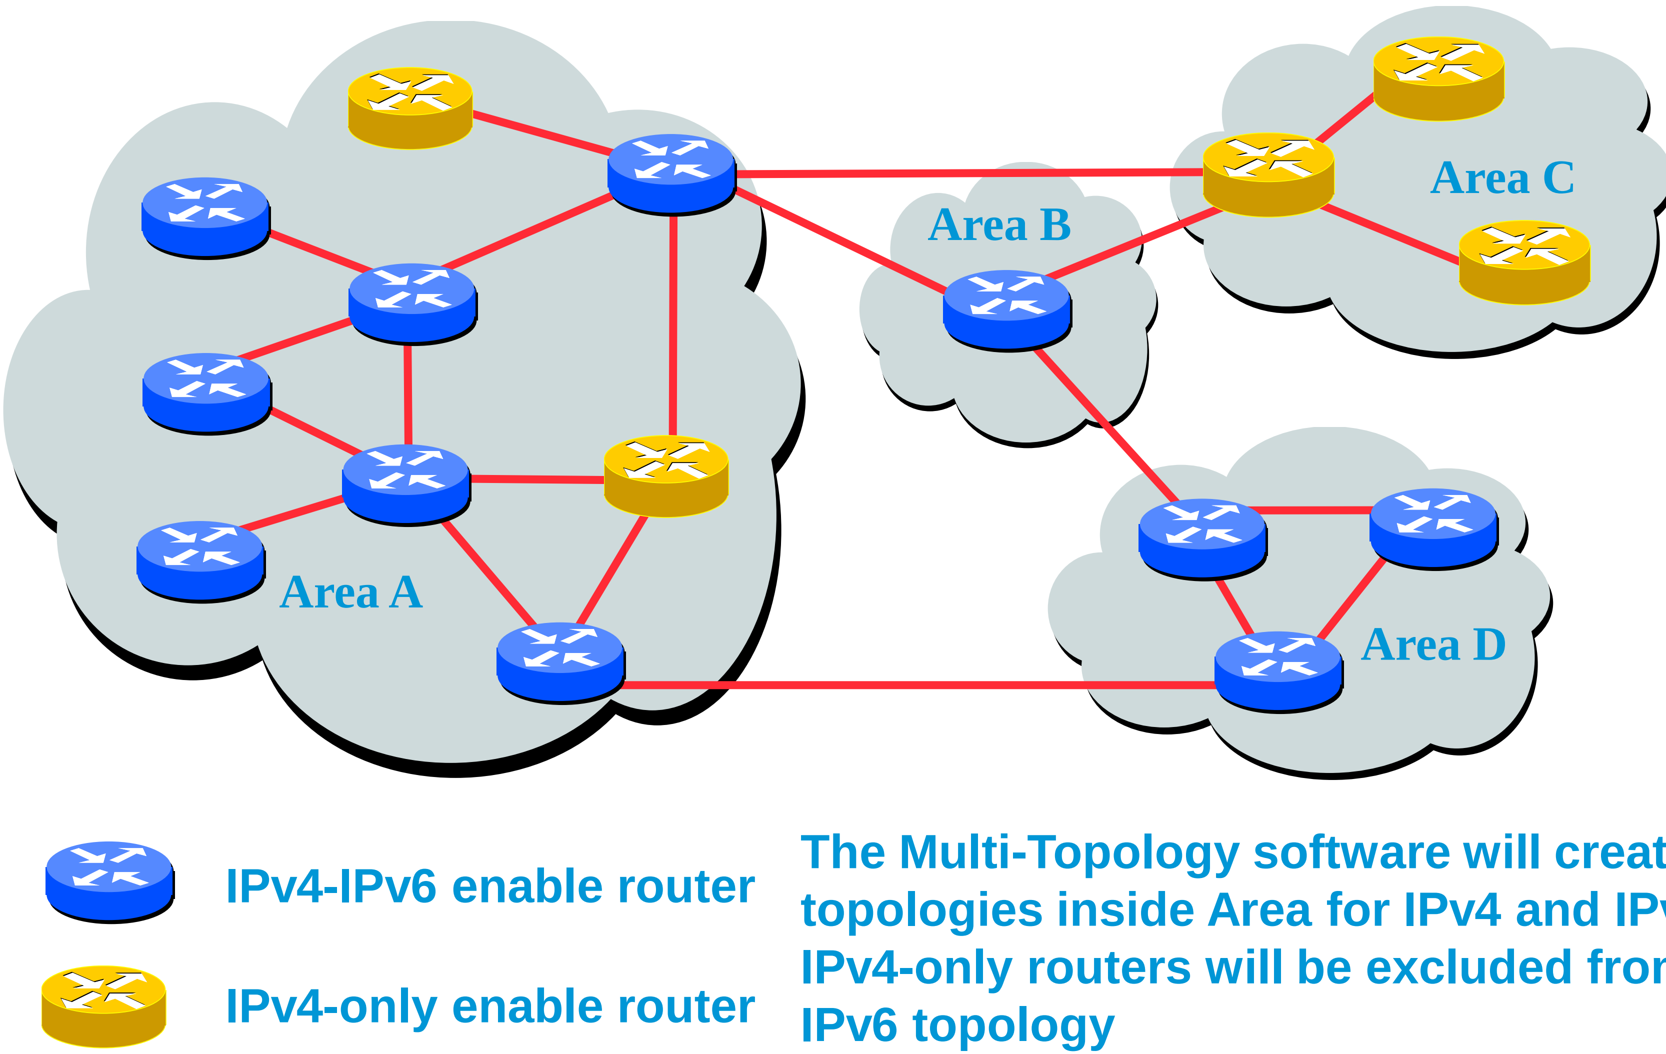

The Multi-Topology software will create two topologies inside Area for IPv4 and IPv6. IPv4-only routers will be excluded from the IPv6 topology

MP-BGP4 for IPv6

BGP-4 Extensions for IPv6 (MP-BGP)

- BGP-4 carries only 3 pieces of IPv4 specific information
  - NLRI in the UPDATE message contains an IPv4 prefix
  - NEXT_HOP path attribute in the UPDATE message contains a IPv4 address
  - BGP Identifier in the OPEN message & AGGREGATOR attribute
- RFC 4760 defines multi-protocol extensions for BGP-4 to support protocols other than IPv4
  - New BGP-4 optional and non-transitive attributes:
    - ✓MP_REACH_NLRI
    - ✓MP_UNREACH_NLRI
  - Protocol independent NEXT_HOP attribute
  - Protocol independent NLRI attribute

MP-BGP IPv6 Support

- 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" (RFC4760)
  - 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), Sub AFI (SAFI)
  - Next-Hop Information (must be of the same address family)
  - NLRI

AFI	Meaning
1	IPv4
2	IPv6

SAFI	Meaning
1	NLRI used for unicast
2	NLRI used for multicast
3	NLRI used for unicast and multicast
4	NLRI with MPLS labels
64	Tunnel SAFI
65	VPLS
66	BGP MDT
128	MPLS-labeled VPN address (VPNv4, VPNv6)

MP-BGP for IPv6 Considerations

- TCP Interaction
  - BGP-4 runs over a TCP (179) session using IPv4 or IPv6
  - The NLRI BGP carried (IPv4, IPv6, MPLS) is agnostic of the session protocol
- Router ID
  - If IPv4 session is not used, a BGP router-id must still exist in a 32 bit dotted decimal notation
  - The RID does not have to be in valid IPv4 format. For example, 0.0.0.1 is valid
  - The sole purpose of RID is for identification
  - In BGP it is used as a tie breaker and is sent within the OPEN message
- Next-hop contains a global IPv6 address (or potentially a link local address)
- Link local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised)

BGP Peering Address

- Two options are available for configuring BGP peering
- Using link local addressing
  - ISP uses FE80:: addressing for BGP neighbours
  - Deployable but not recommended
    - There are plenty of IPv6 addresses
    - Unnecessary configuration complexity
- Using global unicast addresses
  - As with IPv4
  - Recommended option

BGP Considerations

- Rules for constructing the NEXTHOP attribute:
  - When two peers share a common subnet, the NEXTHOP information is formed by a global address and a link local address
  - Redirects in IPv6 are restricted to the usage of link local addresses

IPv6 Transition Mechanisms

Dual Stack Technique

IPv6 using Dual Stack Backbone

- 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)

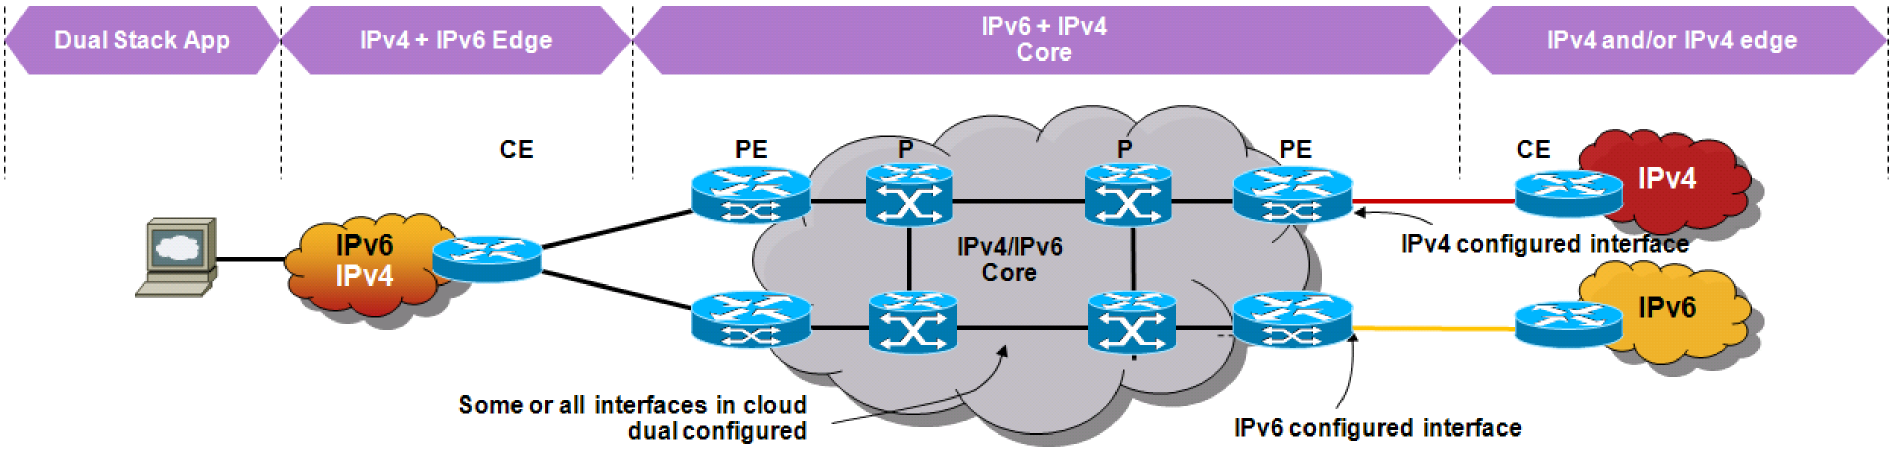

Application Dual Stack Approach

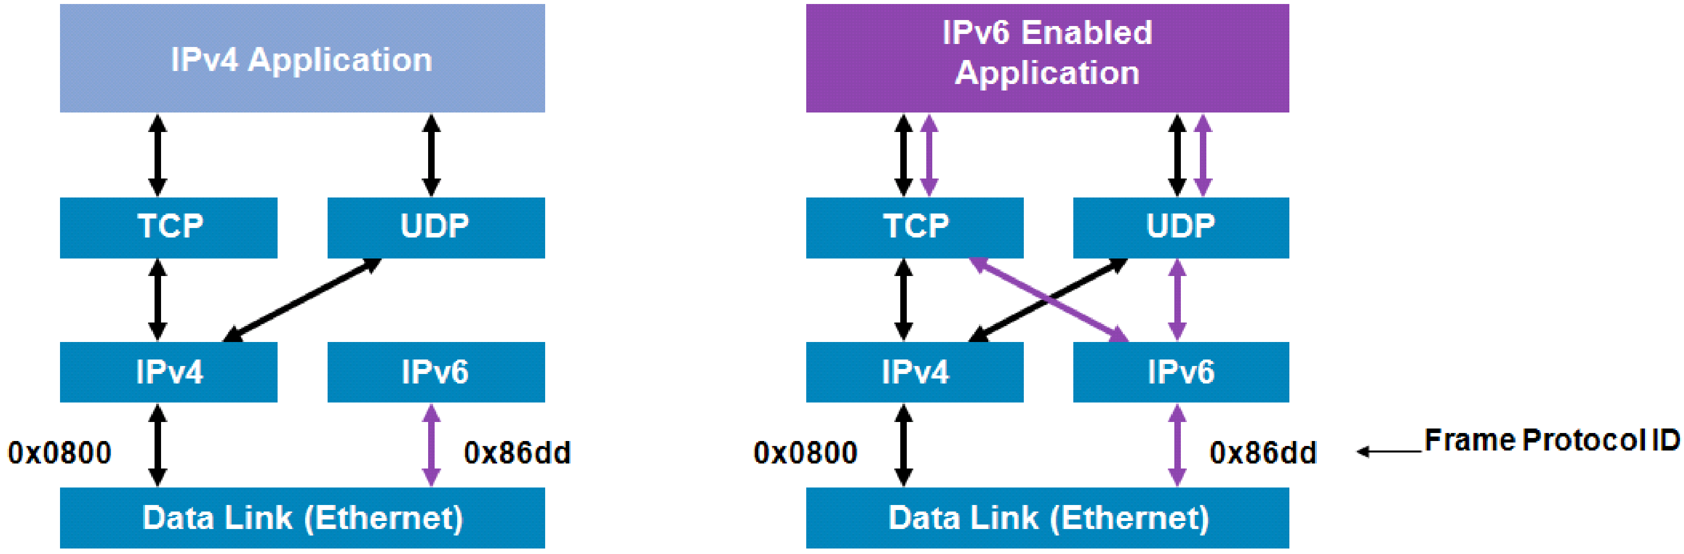

- Dual stack in a device means
  - Both IPv4 and IPv6 stacks enabled
  - Applications can talk to both
  - Choice of the IP version is based on DNS and application preference
- Dual stack at edge does not necessarily mean dual stack backbone

Tunneling

Tunneling

- There are two general types of tunneling
- Manually configured tunneling of IPv6 over IPv4
  - IPv6 packets are encapsulated in IPv4 packets to be carried over IPv4 routing infrastructures.
  - These are point-to-point tunnels that need to be configured manually
- Automatic tunneling of IPv6 over IPv4
  - IPv6 nodes can use different types of addresses, such as 6 to 4 or ISATAP addresses, to dynamically tunnel IPv6 packets over an IPv4 routing infrastructure
  - These special IPv6 unicast addresses carry an IPv4 address in some parts of the IPv6 address fields.

Manual Tunnels

Manual Tunnels

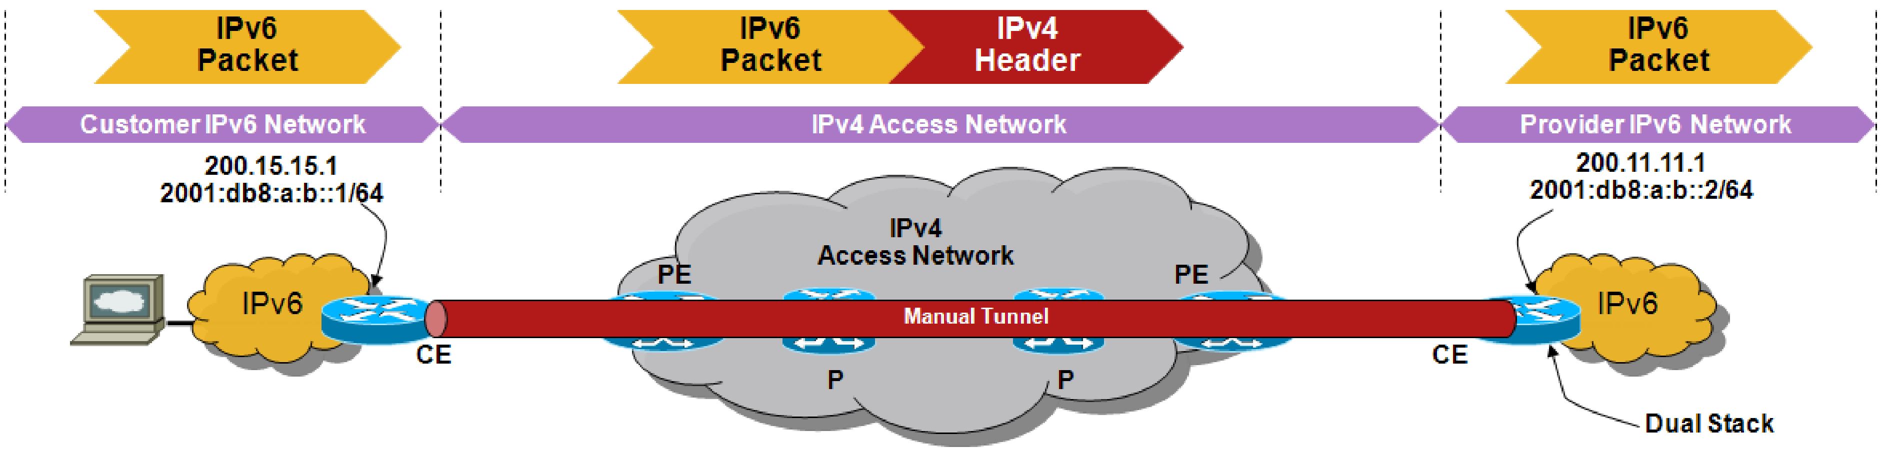

- One of the first transition mechanisms developed for IPv6
  - Static P2P tunnel, IP protocol type = 41, no additional header
- Terminates on dual stack end points
  - IPv4 end point address must be routable
  - IPv6 prefix configured on tunnel interface
- Difficult to scale and manage
  - For link few sites in fixed long term topology
  - Use across IPv4 access network to reach IPv6 Provider

IPv6 over GRE Tunnel

- Similar to Manual Tunnel (RFC 2893)
  - But can transport non IP packets
  - Hence can be used to support ISIS across the tunnel
- GRE header uses 0x86DD to identify IPv6 payload
- Similar scale and management issues

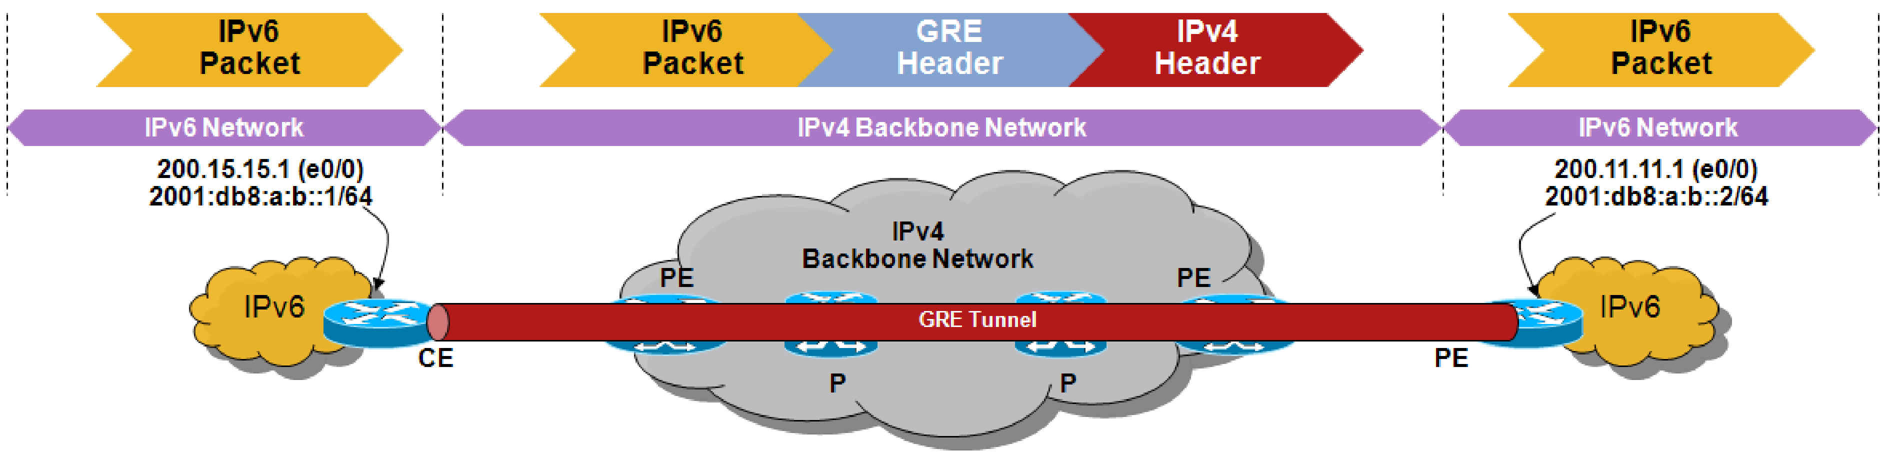

6 to 4 Tunneling

Automatic 6 to 4 Tunnels

- Automatic tunnel method using 2002:IPv4::/48 IPv6 range
  - IPv4 embedded in IPv6 format eg:- **`2002:c80f:0f01::`** = **`200.15.15.1`**
- No impact on existing IPv4 or MPLS Core (IPv6 unaware)
- Tunnel endpoints have to be IPv6 and IPv4 aware (Dual stack)
- Transition technology – not for long term use
- No multicast support, Static Routing
- Intrinsic linkage between destination IPv6 Subnet and IPv4 gateway interface
  - IPv4 Gateway = Tunnel End point

Automatic 6to4 Tunnel (RFC 3056)

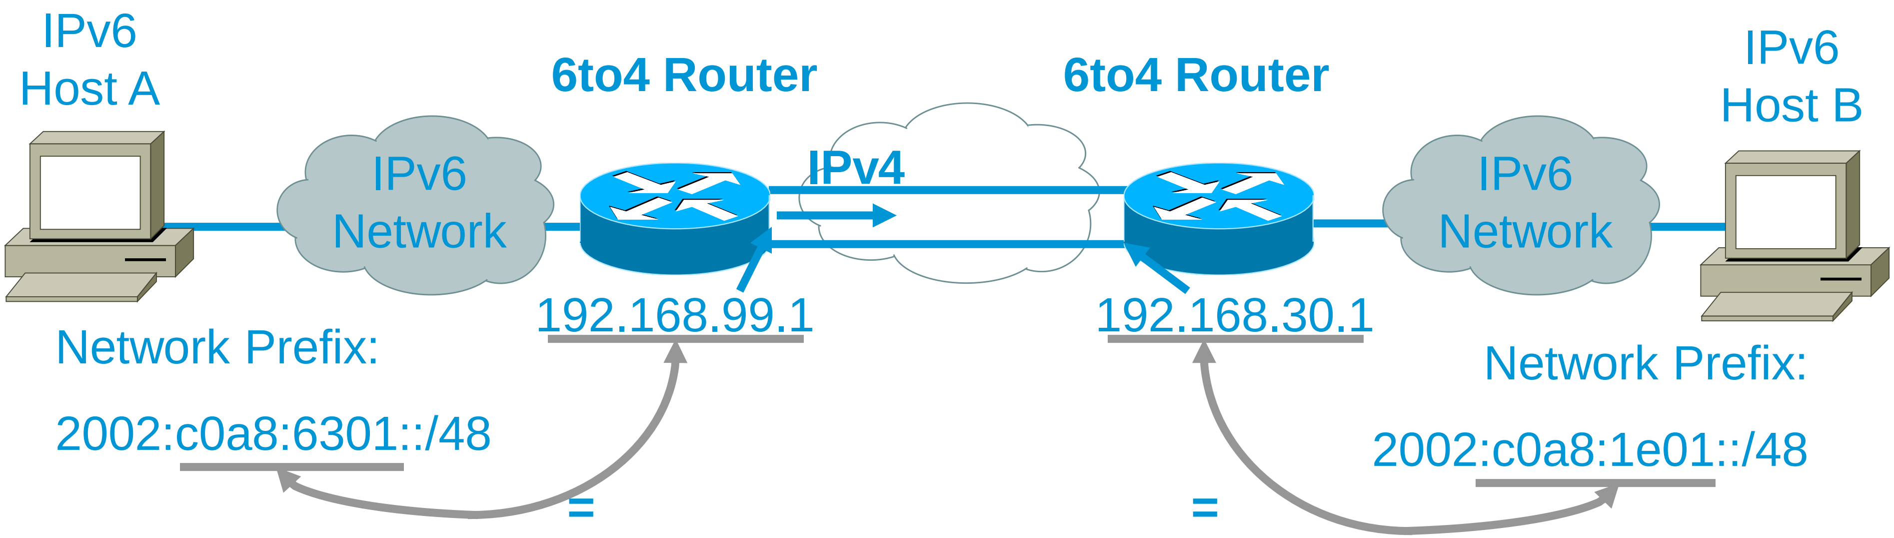

6to4:

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

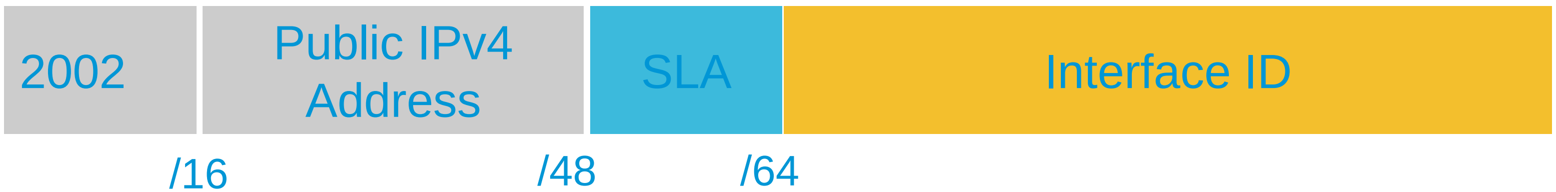

Automatic 6to4 Tunnel (RFC 3056)

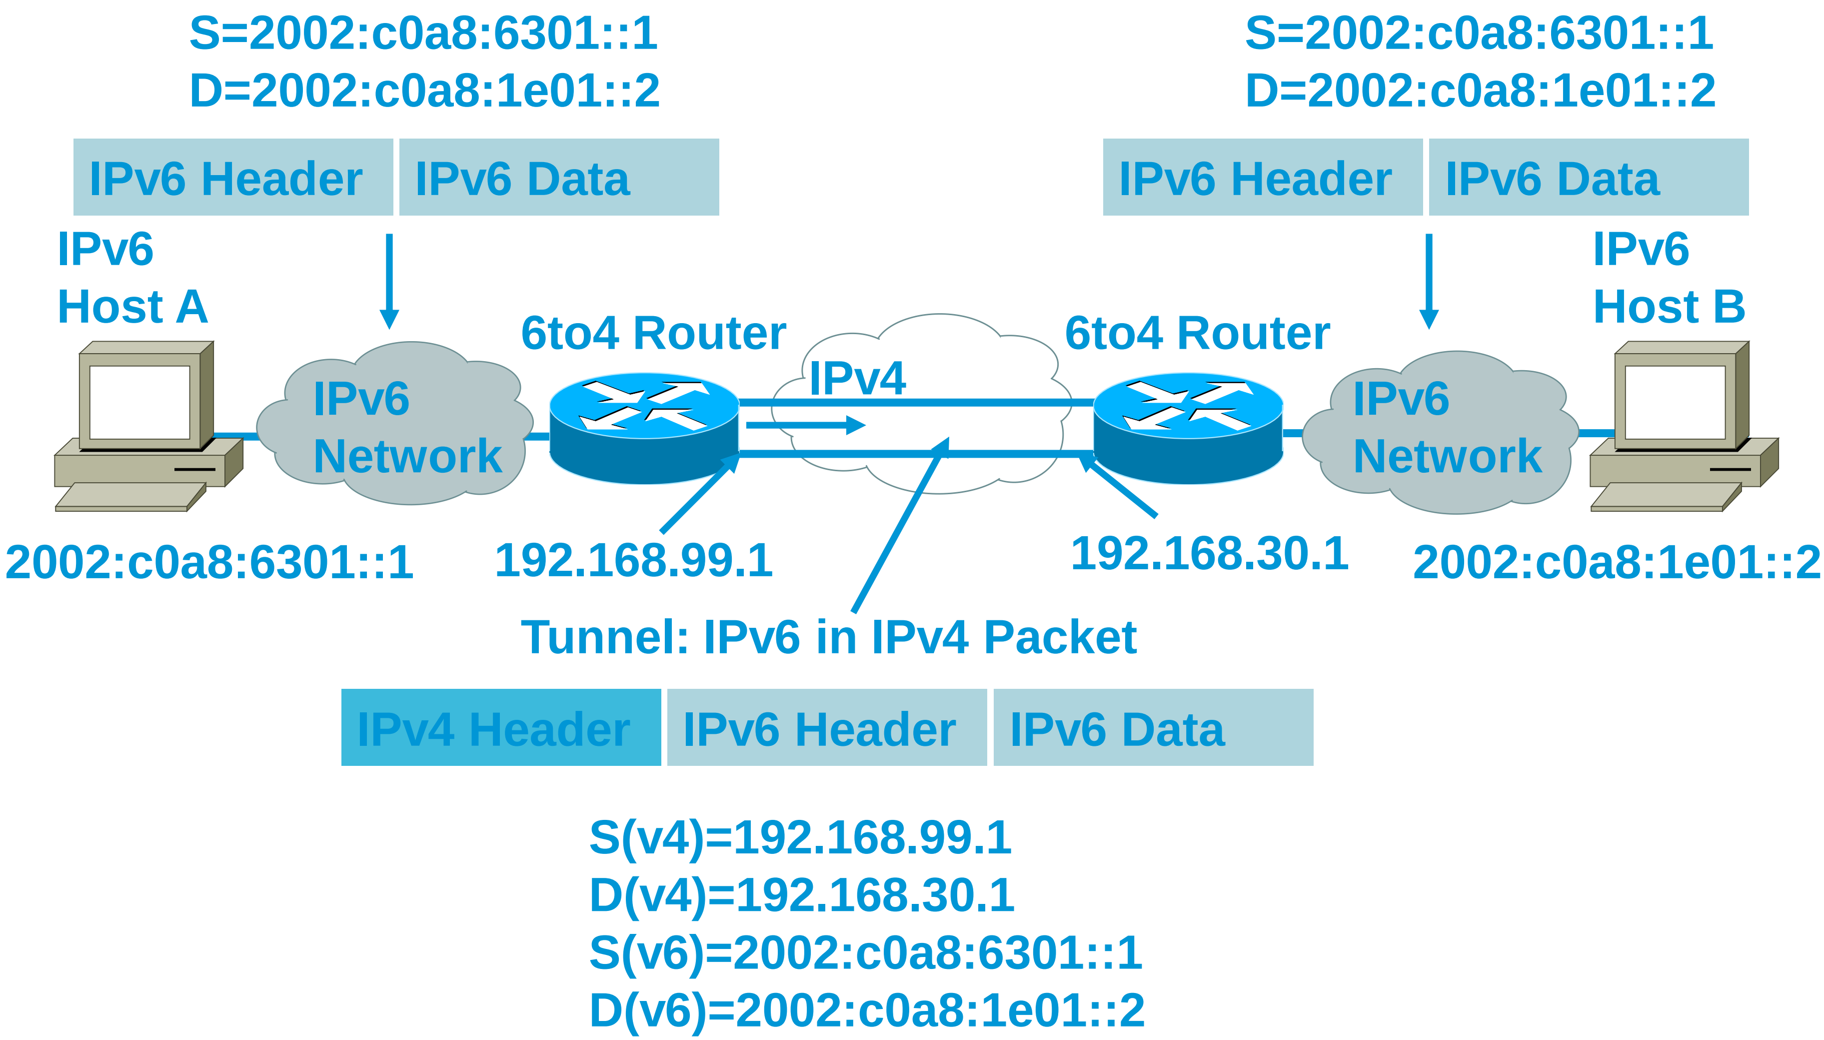

6 to 4 Relay Service

- 6 to 4 relay allows access to IPv6 global network
- Can use tunnel Anycast address 192.88.99.1
  - 6 to 4 router finds closest 6-to-4 relay router
  - Return path could be asymmetric
- Default route to IPv6 Internet

Automatic 6to4 Relay

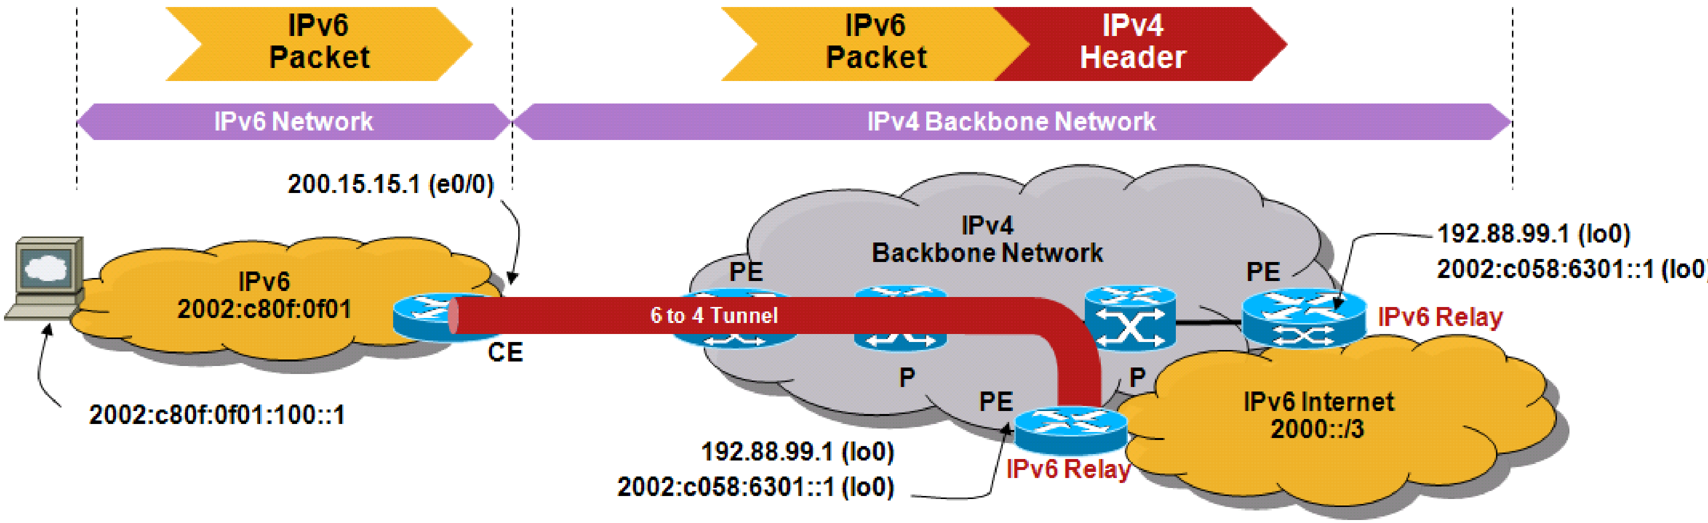

- 6 to 4 relay allows access to IPv6 global network
- Can use tunnel Anycast address 192.88.99.1
  - 6 to 4 router finds closest 6-to-4 relay router
  - Return path could be asymmetric
- Default route to IPv6 Internet

ISATAP Tunneling

ISATAP Address Format

- ISATAP hosts use a special IPV6 address format
- Interface ID carries information
  - Rightmost 32 bits contains the host IPv4 address
  - Leftmost 32 bits contains “0000:5efe”
- Global prefix provided by ISATAP router
  - Interface ID portion remain static for all packets
  - Link-Local addresses used for solicitation of global address

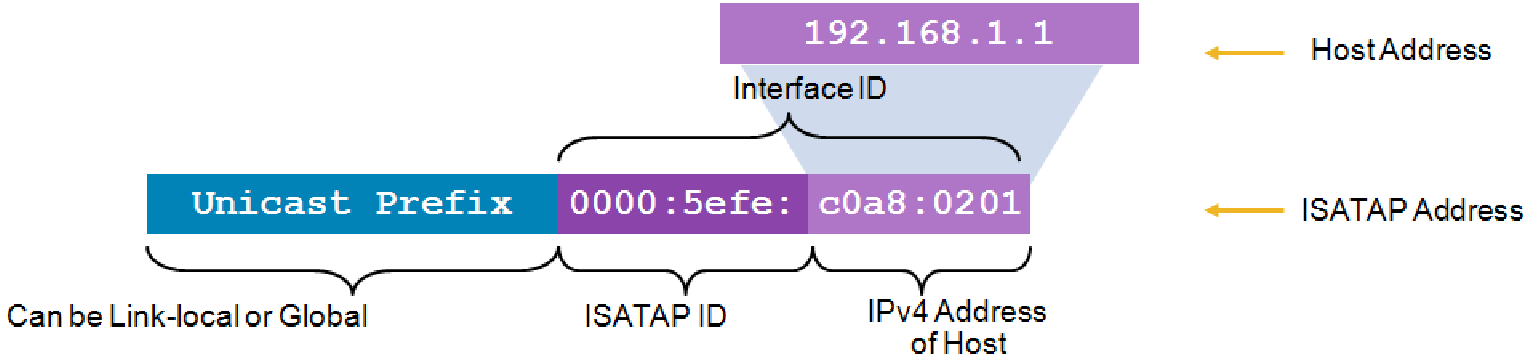

Automatic Address Assignment of Host and Router

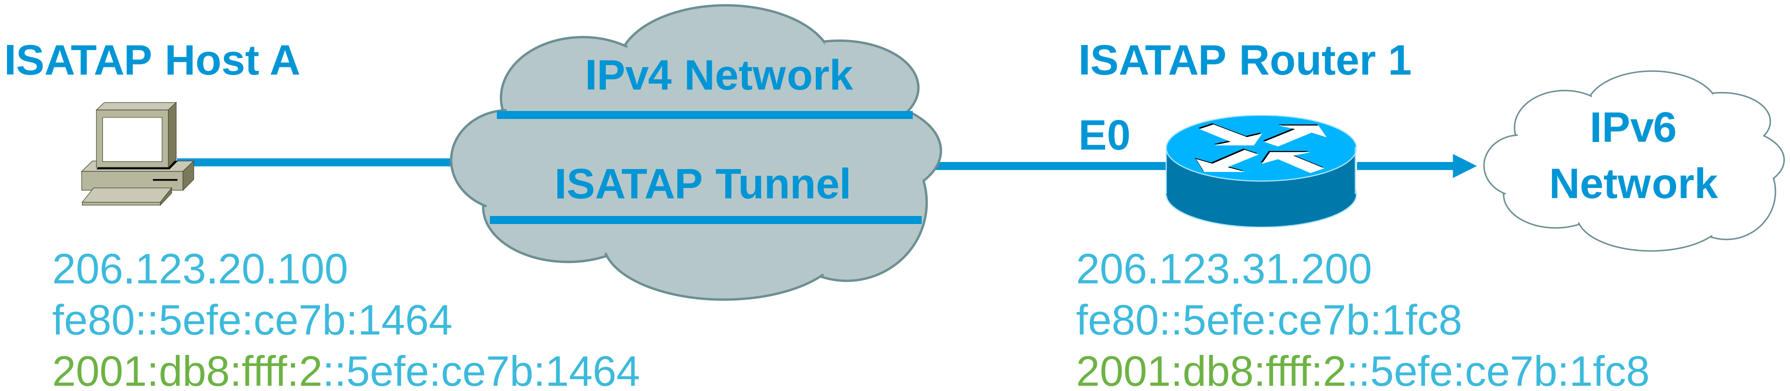

- ISATAP host A receives the ISATAP prefix 2001:db8:ffff:2::/64 from ISATAP Router 1
- When ISATAP host A wants to send IPv6 packets to 2001:db8:ffff:2::5efe:ce7b:1fc8, ISATAP host A encapsulates IPv6 packets in IPv4. The IPv4 packets of the IPv6 encapsulated packets use IPv4 source and destination address.

IPv6 Rapid Deployment (6rd)

6RD Overview

- 6rd is a tunnelling method specified in RFC 5969
  - Superset of 6to4 tunnelling [RFC3056]
  - 6rd utilises an SP's own IPv6 address prefix - avoids well-known prefix (2002::/16)
- Method of incrementally deploying IPv6 to end sites in an SP network
  - SP access and aggregation infrastructure remains IPv4
  - End site is provided a dual stack service
  - Access/Aggregation between SP and end sites looks like multipoint network
- End sites share a common IPv6 prefix allocated by SP
- 6rd primarily supports IPv6 deployment to
  - A customer site (residential gateway)
  - To an individual IPv6 host acting as a CE

6RD Tunnels

- Native dual-stack IP service to the end site
- Simple, stateless, automatic IPv6-in-IPv4 encap and decap functions
- Embedded IPv4 address needs to match IPv4 address in Tunnel header for security
- IPv6 traffic automatically follows IPv4 Routing (IPv4 address used as tunnel endpoint)
- BRs placed at IPv6 edge, addressed via anycast for load-balancing and resiliency

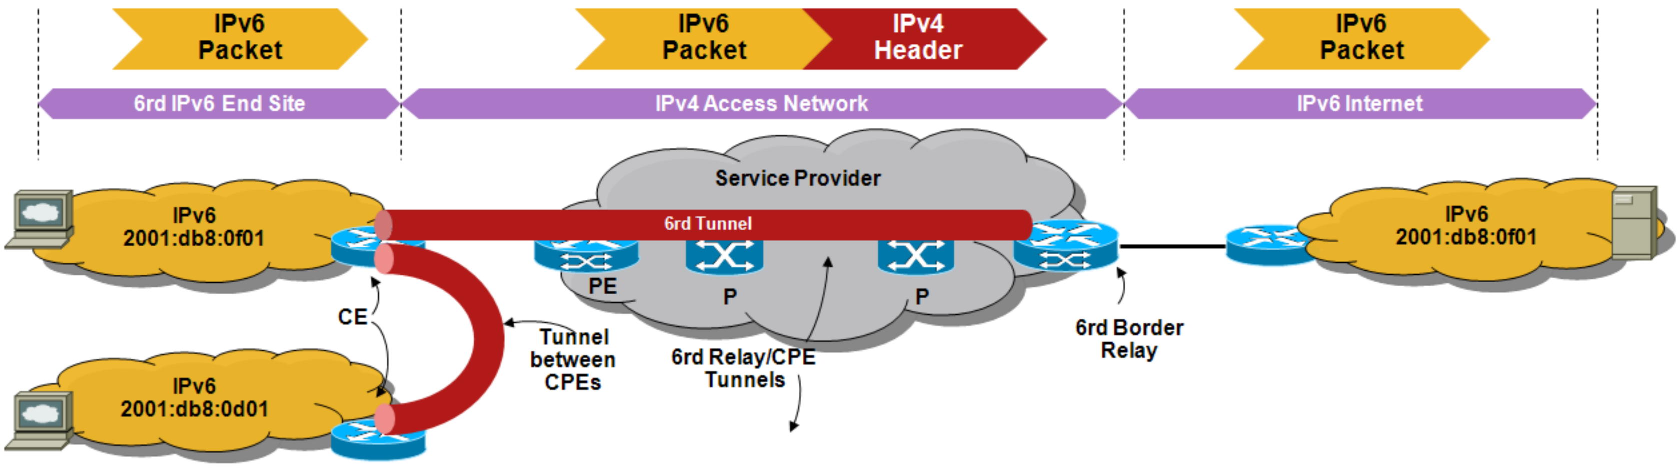

6RD Logical NBMA Behaviour

- 6rd views the IPv4 network as an NBMA link layer for IPv6
- Border Relay serves has a single multipoint interface
  - No per user state, serves all users in 6rd domain

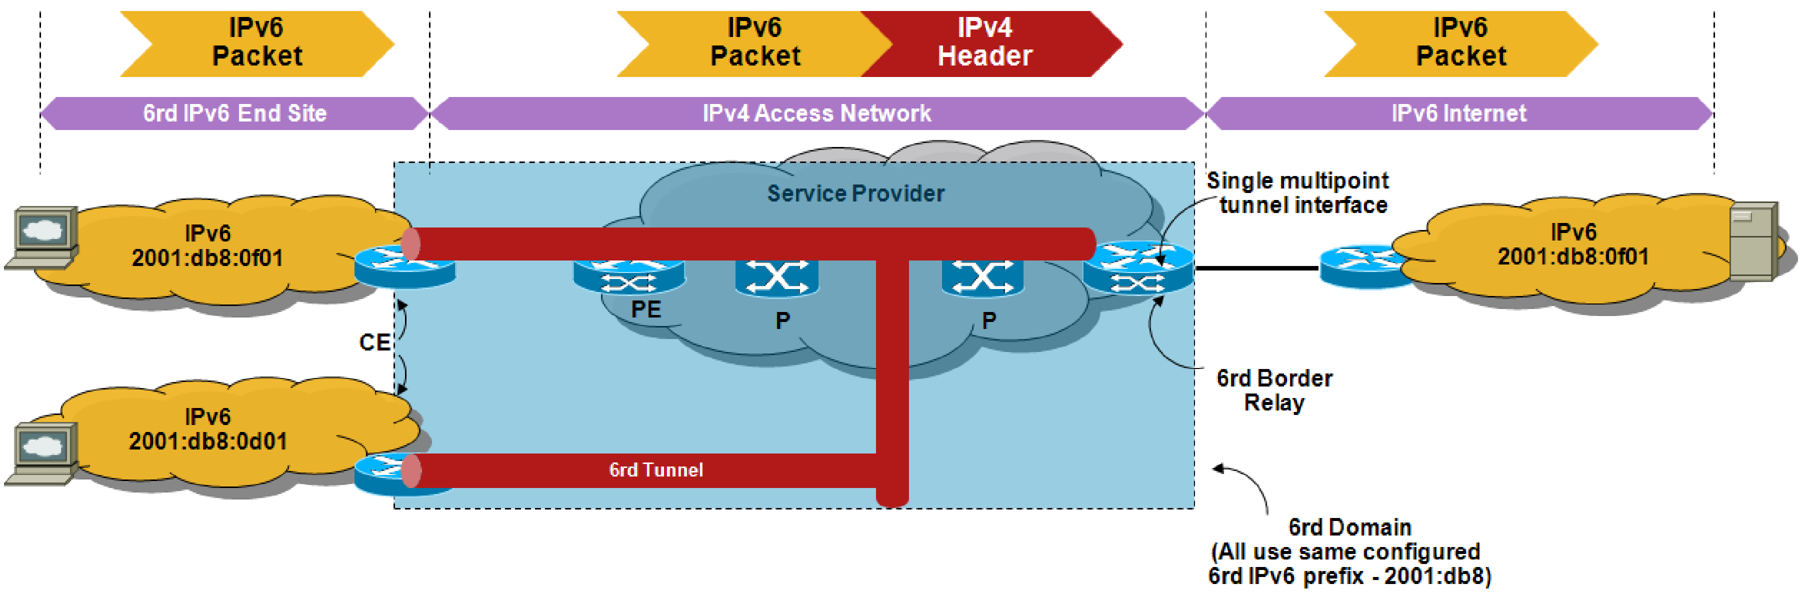

6RD Delegated Prefix

SP v6 Prefix: 2001:B000::/32

v4 common prefix 10.1.0.0/16

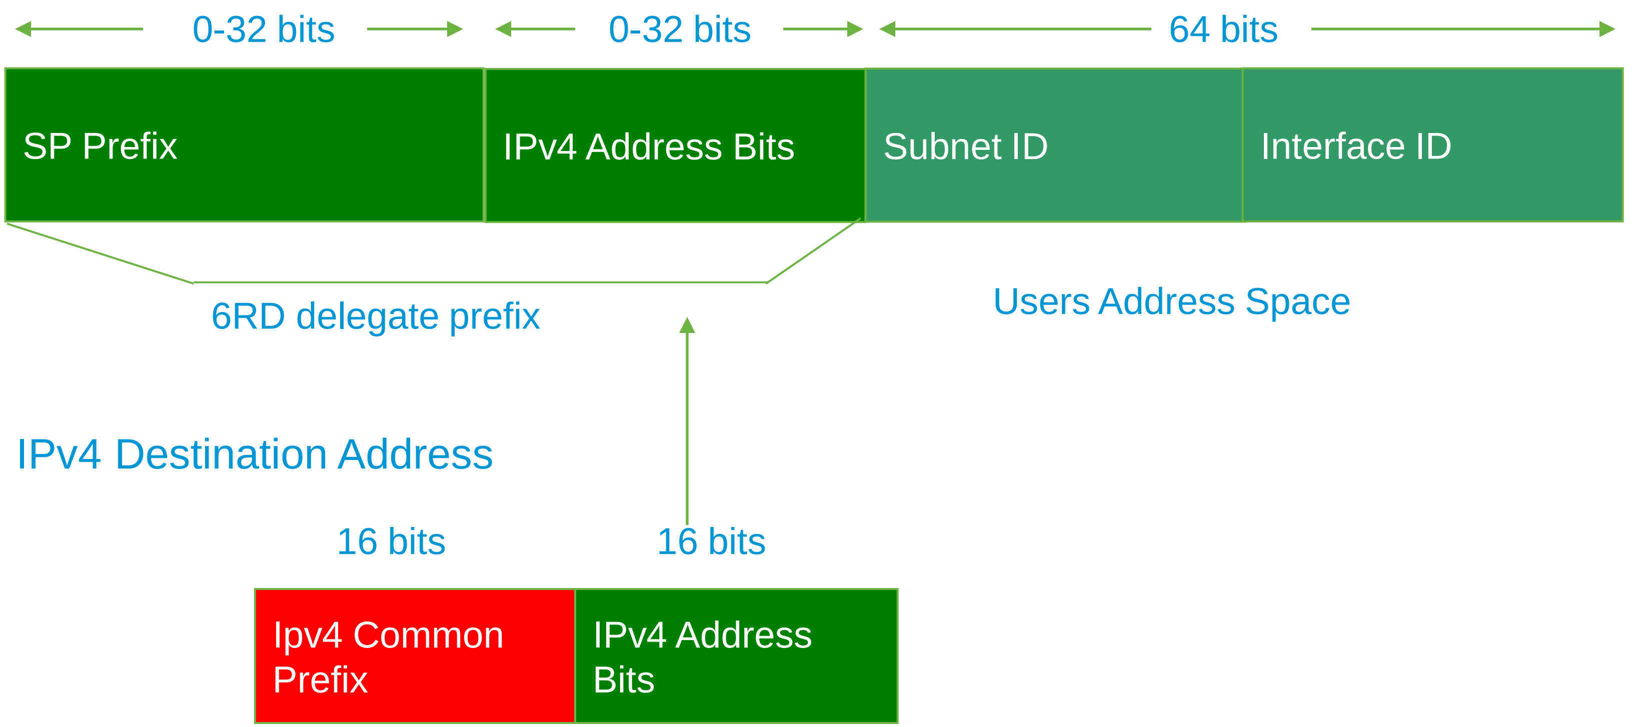

Destination Dynamically Computed Example

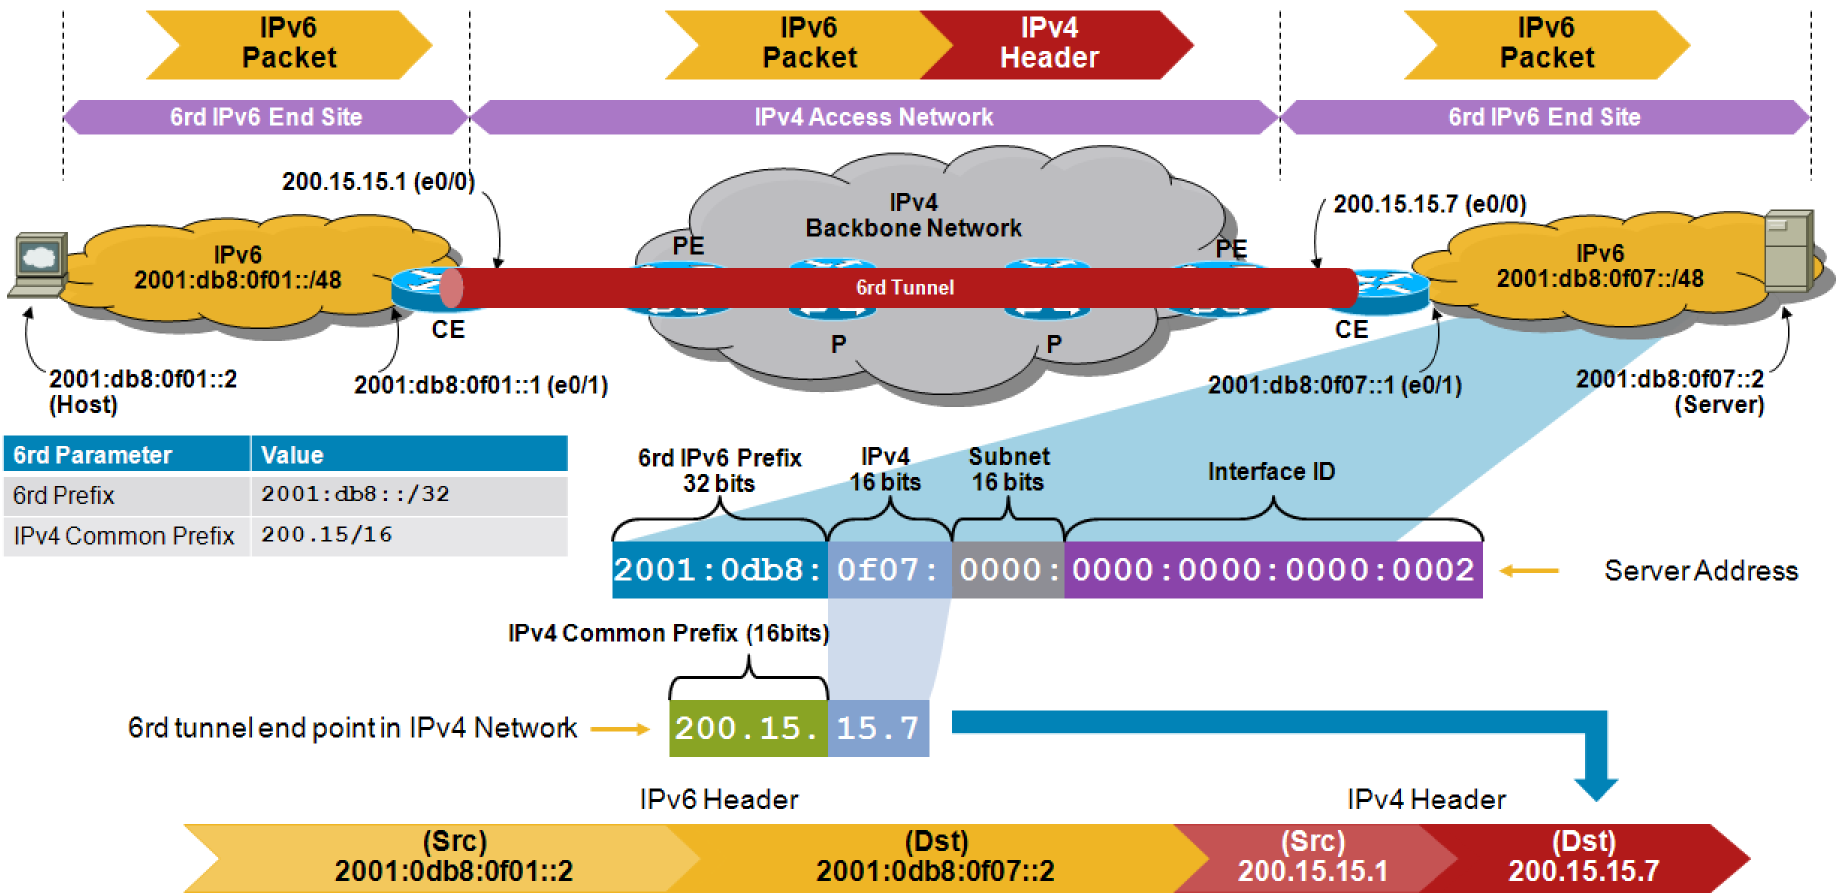

6rd Parameter	Value
6rd Prefix	2001:db8::/32
IPv4 Common Prefix	200.15/16

Internet Access through 6RD Border Relay

- 6rd Border Relay allows access to IPv6 global Internet
  - If IPv6 destination outside of 6rd prefix then tunnel packet to border relay
- Can use tunnel pre-provisioned IPv4 Anycast address
  - 6rd CE router finds closest 6rd BR router based on IGP
- Default route to IPv6 Internet, usually pre-provisioned by the ISP

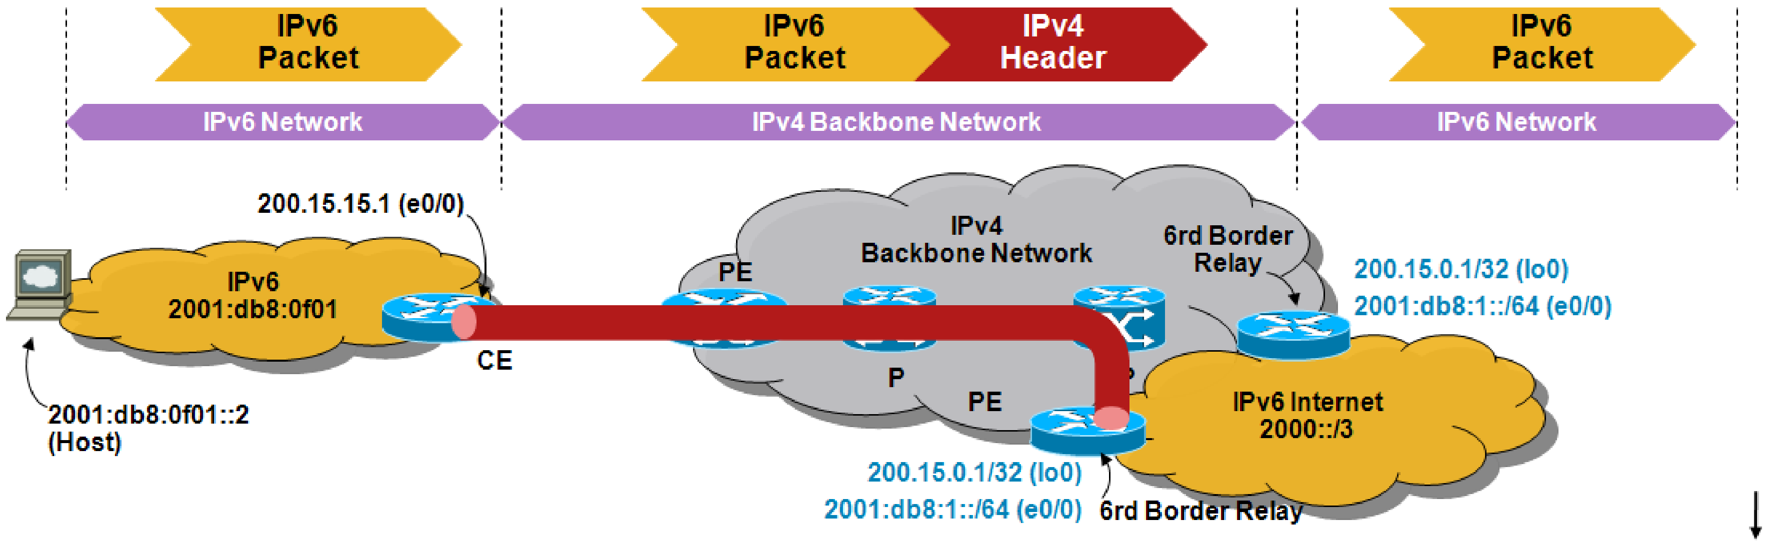

Dual Stack Lite

Dual Stack Lite Overview

- It uses IPv6-only links between the provider and the customer, but does not use NAT64 translation
- When a device in the customer network sends an IPv4 packet to an external destination, the IPv4 packet is encapsulated in an IPv6 packet for transport into the provider network
- At the Provider Edge, the packet is decapsulated and NAT44 is performed
- Tunneling IPv4 over IPv6 is far simpler than translation, so the performance and redundancy concerns are eliminated

Dual-Stack Lite Overview

- If a simple mapping between inside IPv4 address / port was performed on outgoing packets, as is done with regular NAT44, the LSN would have no way to differentiate between overlapping RFC1918 IPv4 addresses in different customer networks
- Therefore an additional element is added to the address mapping: The source address of the encapsulating IPv6 packet (the address of the customer end of the IPv6 link) is added to the inside IPv4 source address and port
- Because the IPv6 address is unique to each customer, the combination of IPv6 source address + IPv4 source address + port makes the mapping unambiguous.

Dual-Stack Lite Overview

IPv6 packet
IPv6 src: IPv6 address of home gateway (IGD)
IPv6 dst: IPv6 address of tunnel concentrator, discovered with DHCPv6
IPv4 src: 192.168.1.3
IPv4 dst: www.nanog.org (198.108.95.21)
IPv4 src port: 1001
IPv4 dst port: 80

IPv4 packet
IPv4 src: from the pool of the ISP
IPv4 dst: www.nanog.org (198.108.95.21)
IPv4 src port: 45673
IPv4 dst port: 80

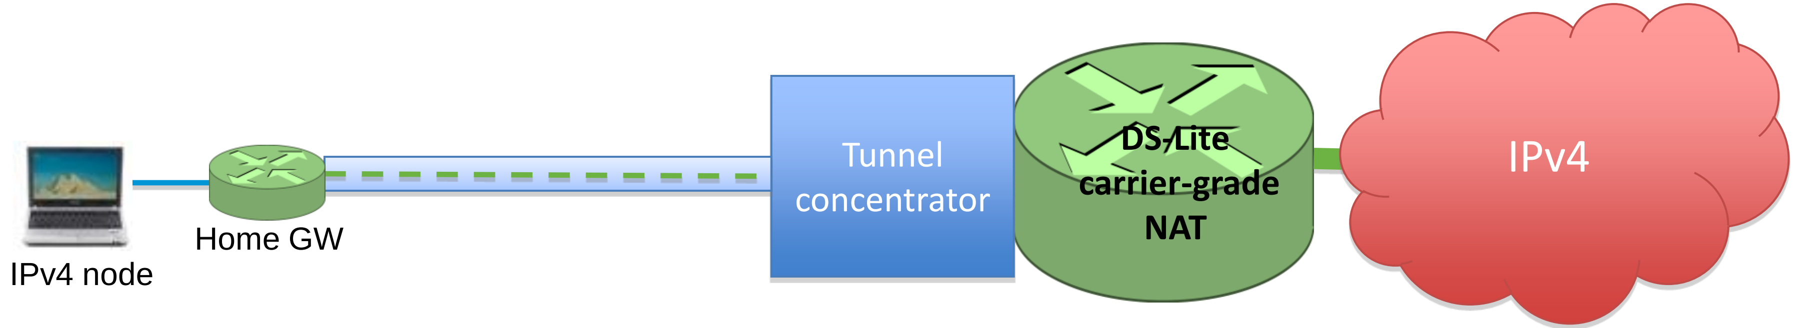

DS-Lite CGN route needs to maintain the mapping:
IN:
IPv6 src: IPv6 address of IGD + 192.168.1.3 + port1001
OUT:
IPv4 src address: from pool of the ISP + port: 45673

DS-Lite: IPv4 & IPv6 Internet Access

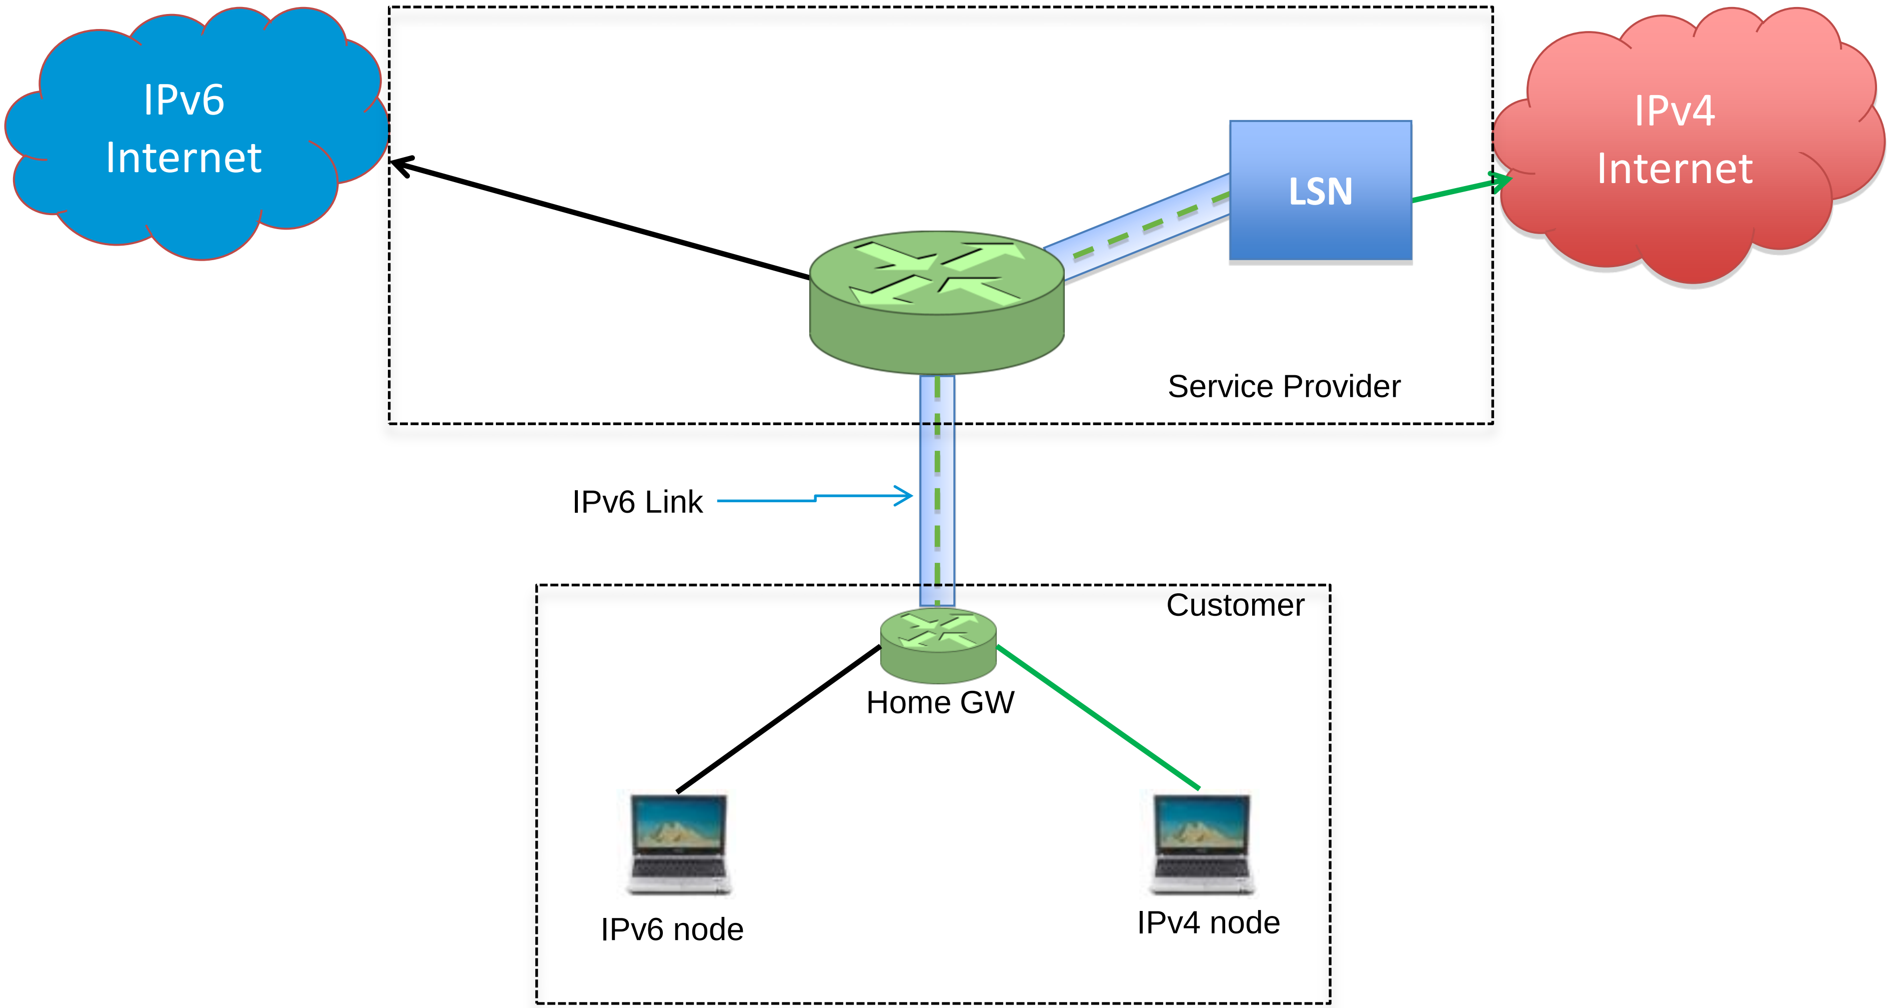

IPv6 packets are routed normally , while IPv4 packets are encapsulated in IPv6 and routed to the LSN

6PE Overview

IPv6 Provider Edge Router (6PE) over MPLS

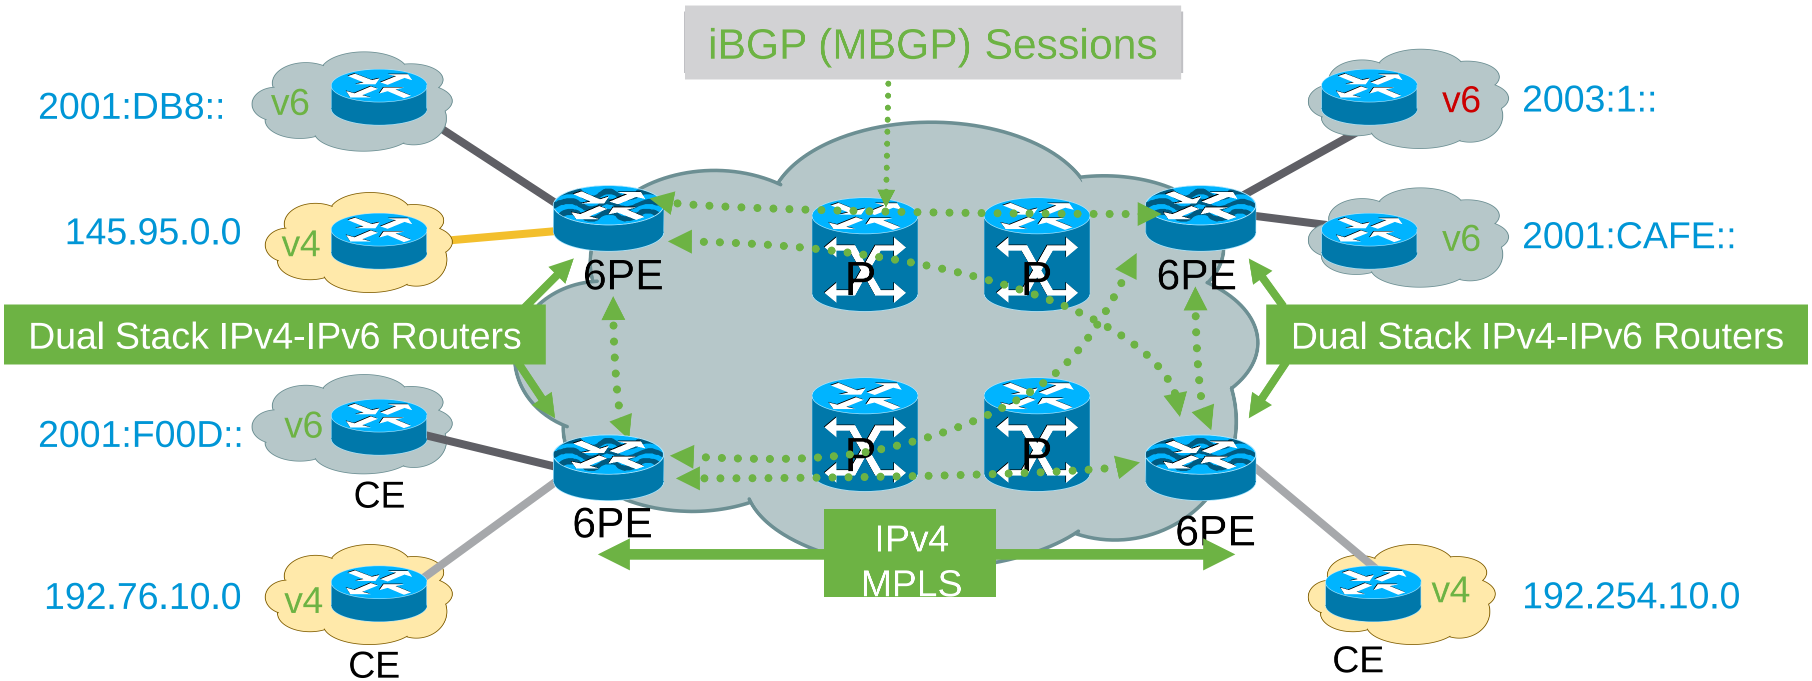

- IPv6 global connectivity over and IPv4-MPLS core
- Transitioning mechanism for providing unicast IP
- PEs are updated to support dual stack/6PE
- IPv6 reachability exchanged among 6PEs via iBGP (MBGP)
- IPv6 packets transported from 6PE to 6PE inside MPLS

6PE Routing/Label Distribution

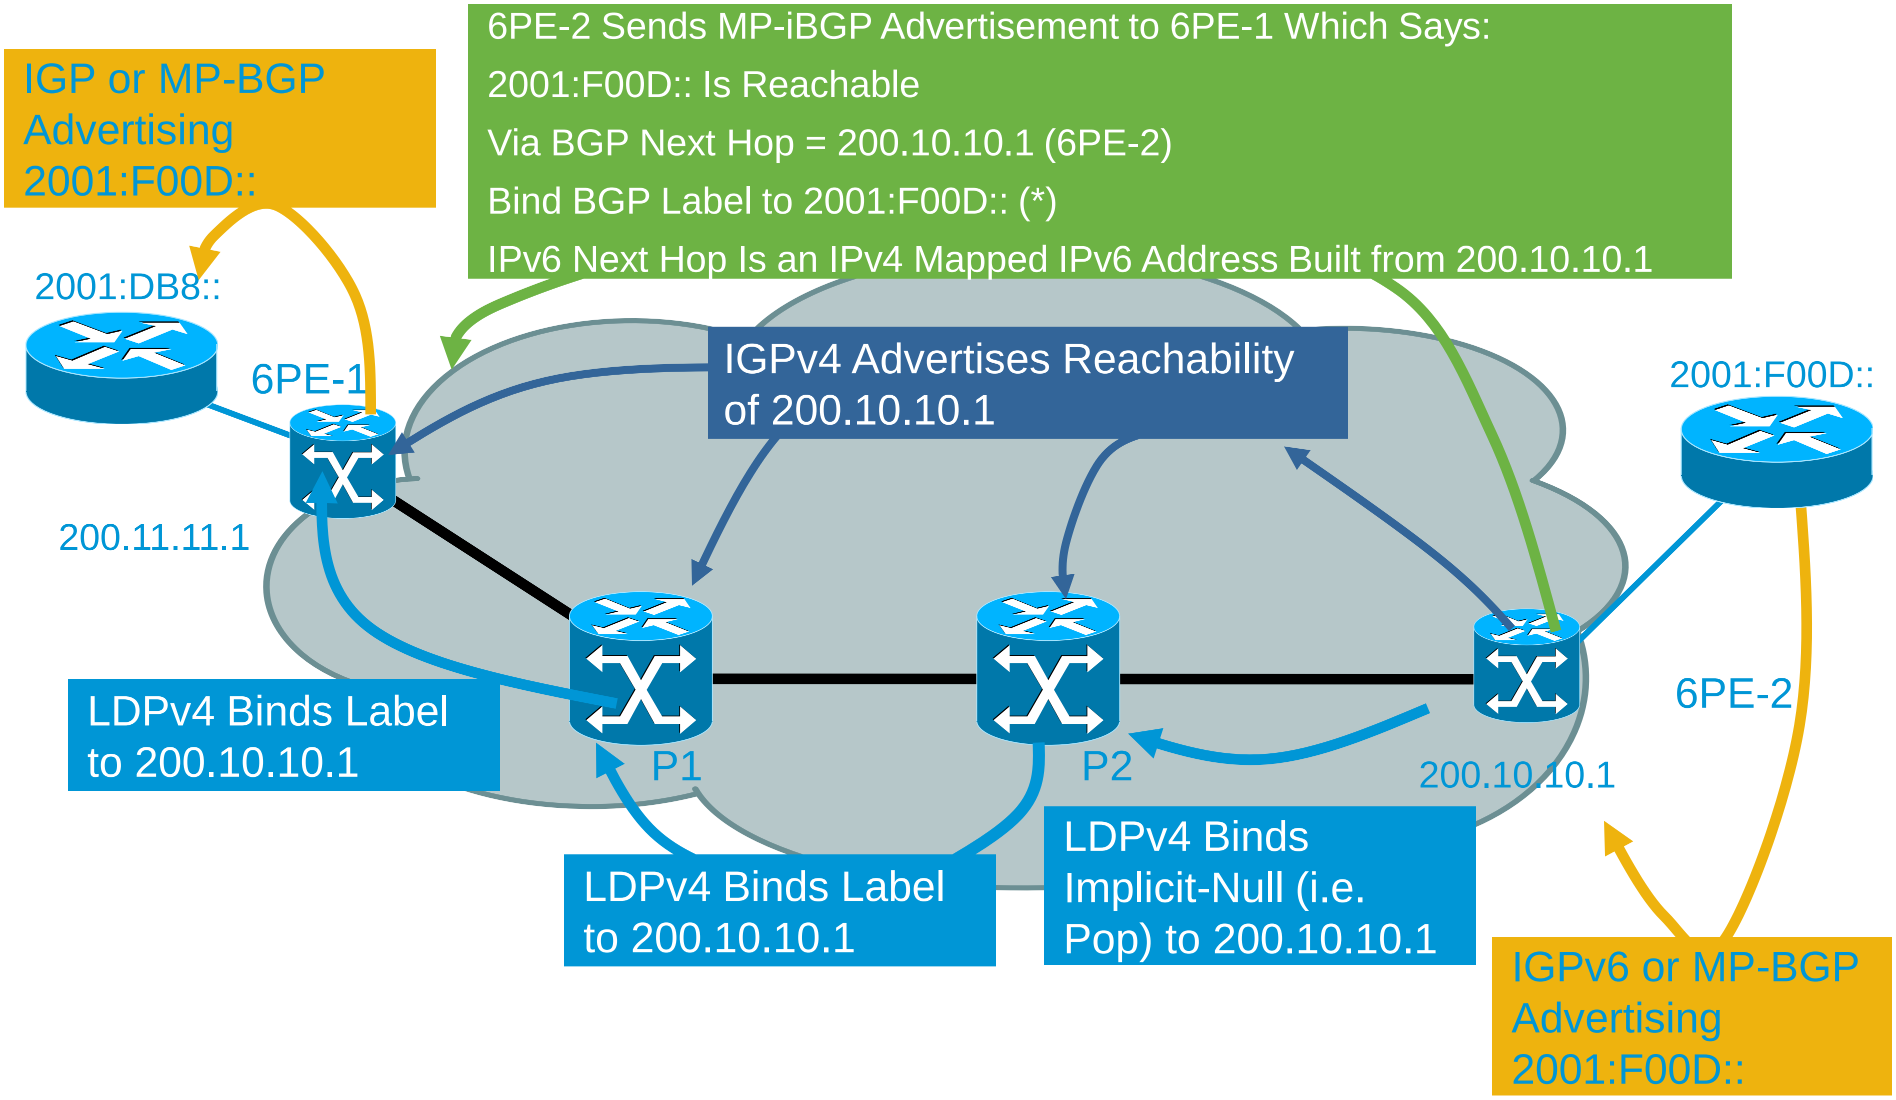

6PE Forwarding (6PE-1)

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)

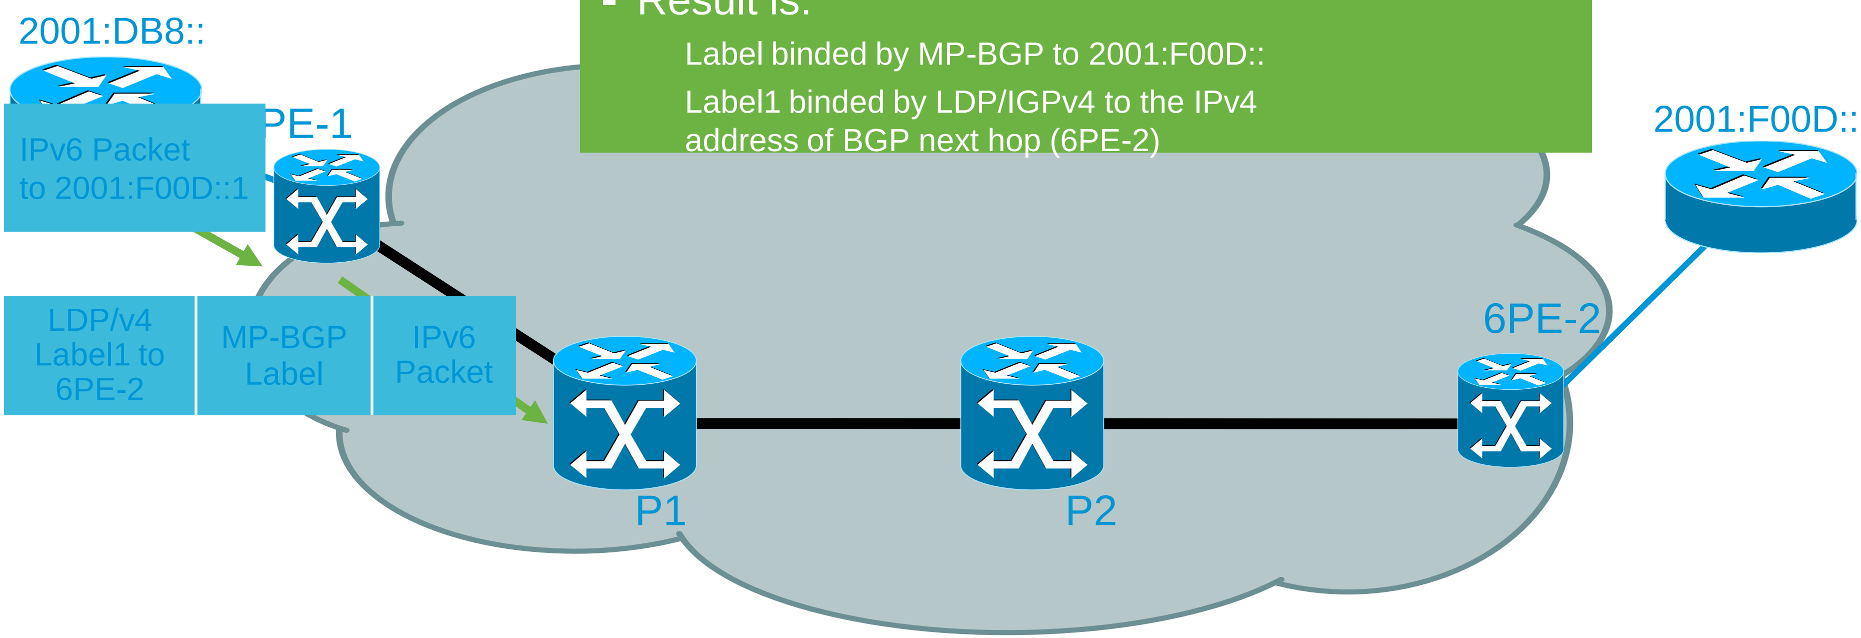

6PE Forwarding (P1)

IPv6-UNaware MPLS Label Switching:

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

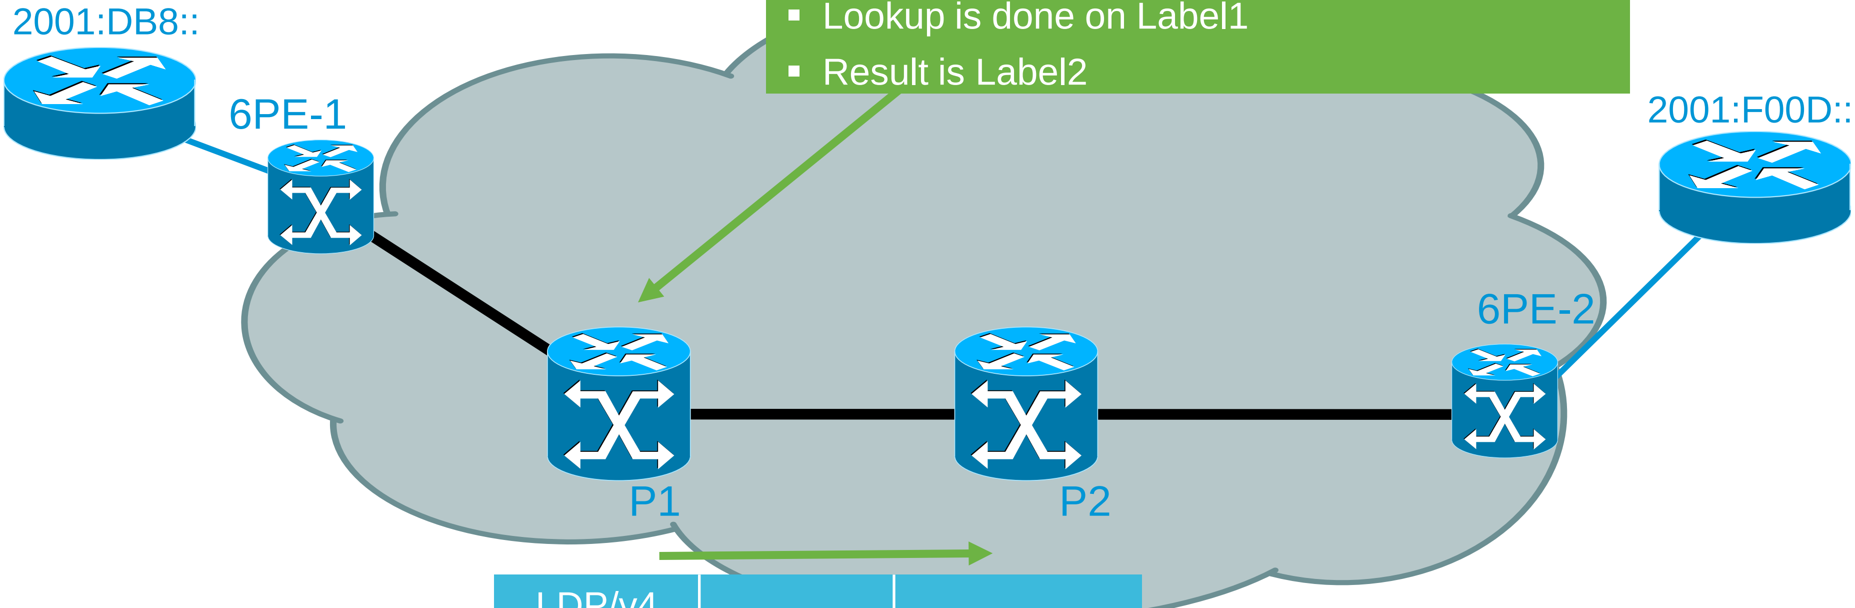

LDP/v4 Label2 to 6PE-2	MP-BGP Label	IPv6 Packet

6PE Forwarding (P2)

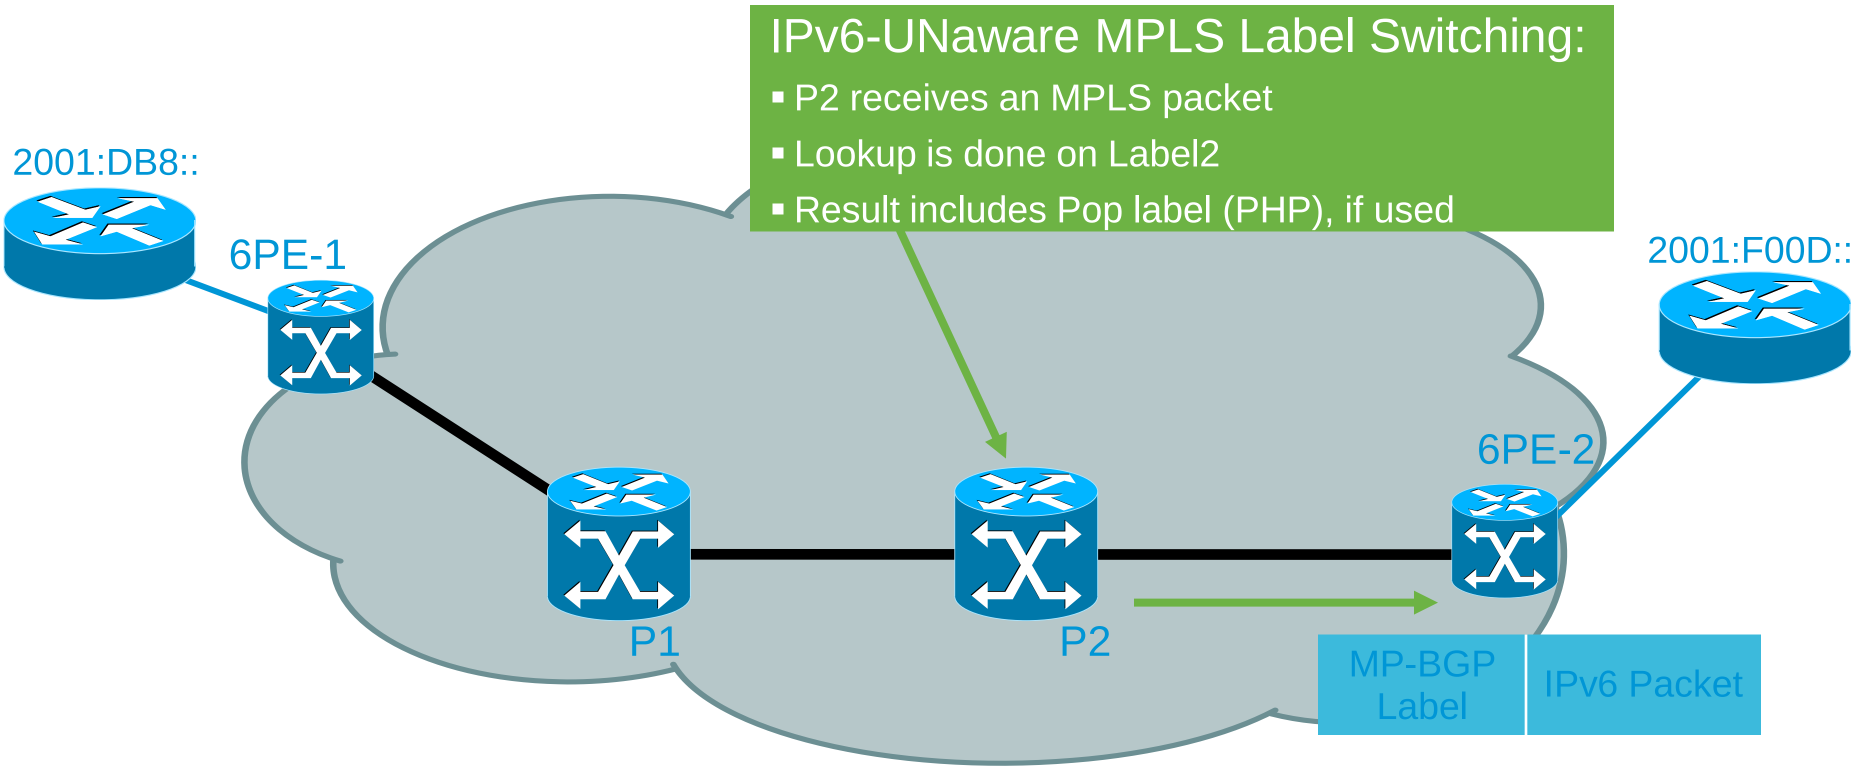

6PE Forwarding (6PE-2)

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

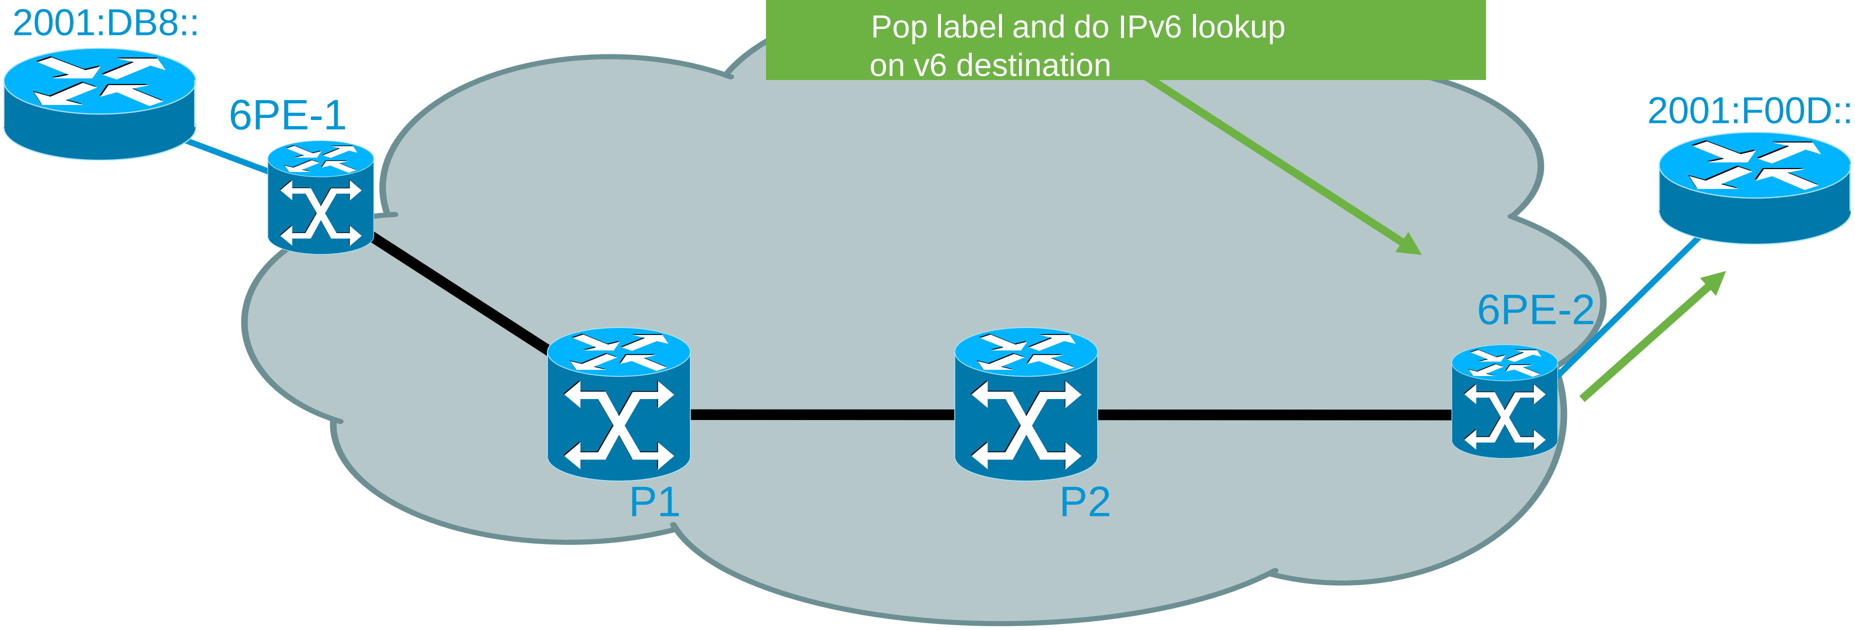

6PE Benefits/Drawbacks

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

6VPE Overview

6VPE Deployment

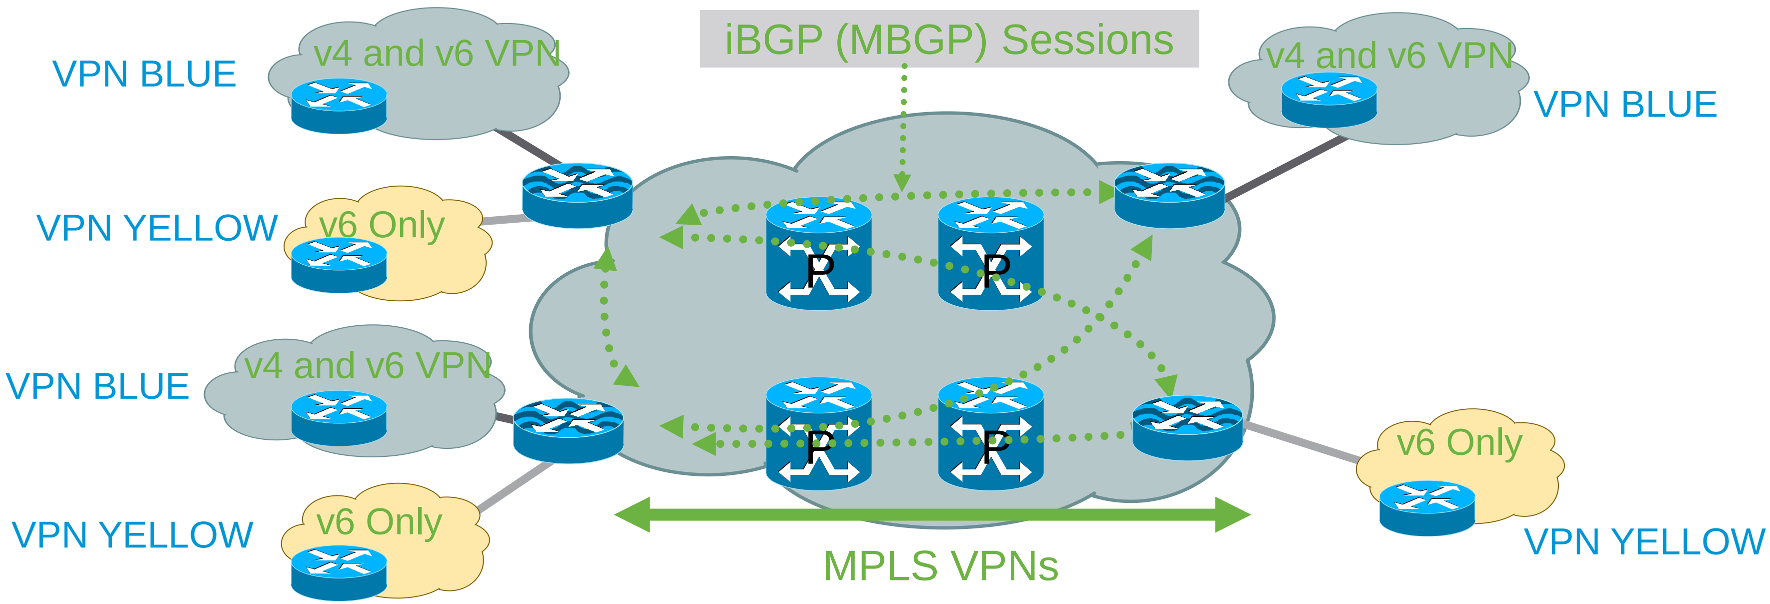

- 6VPE ~ IPv6 + BGP-MPLS IPv4 VPN + 6PE
- VPNv6 address:

  Address including the 64 bits route distinguisher and the 128 bits IPv6 address

- MP-BGP VPNv6 address-family:

  AFI “IPv6” (2), SAFI “VPN” (128)

- VPN IPv6 MP_REACH_NLRI

  With VPNv6 next-hop (192bits) and NLRI in the form of <length, IPv6-prefix, label>

- Encoding of the BGP next-hop

6VPE Example Design
Addressing/Routing

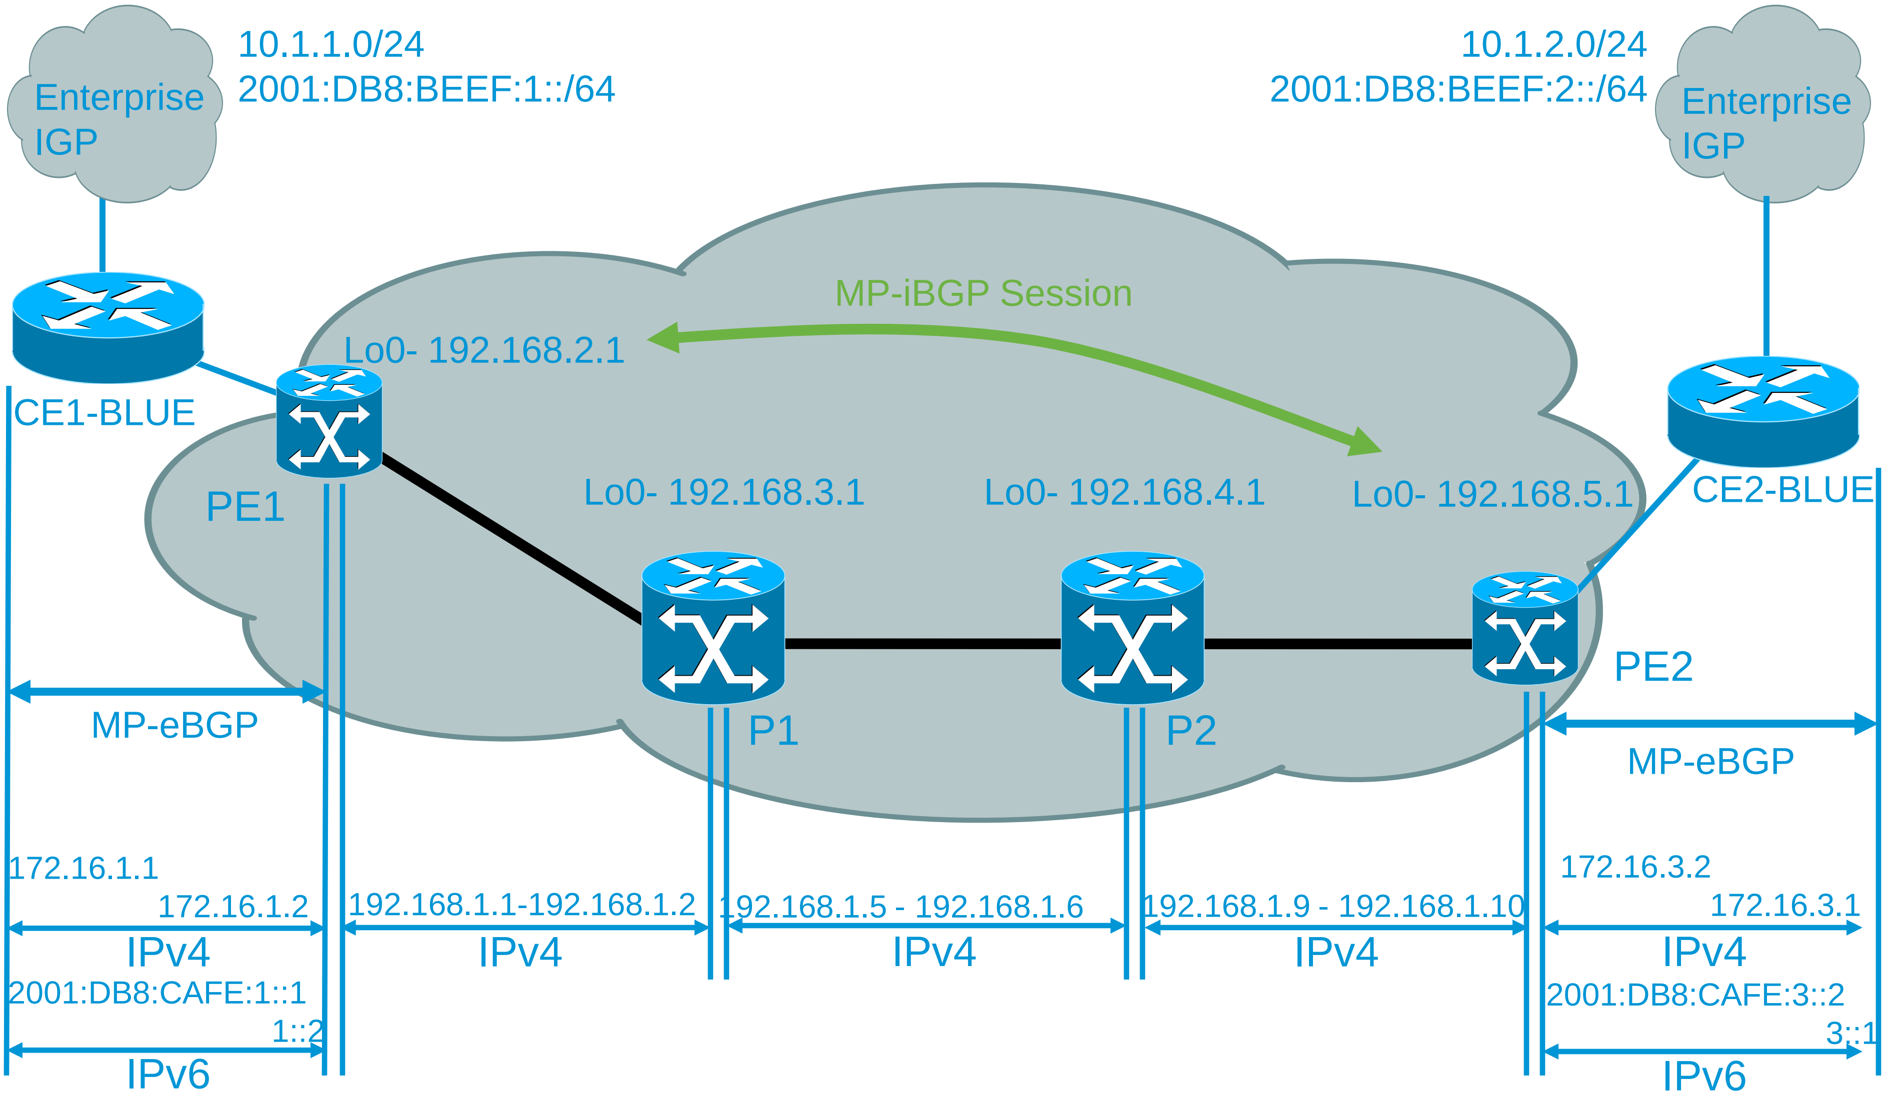

Q & A

Thank You!