IPv6 ISP Routing and Design Considerations

Khalid Raza

sraza@cisco.com

Agenda

- Why IPv6
- Today's networks and services
- IPv6 Routing Protocols Overview
    - EIGRP
    - ISIS
    - OSPFv3
    - BGPv4
- Dual Stack
    - IGP considerations
- ISP network architecture
- MPLS

Why IPv6

Brief Review

IPv6 Adoption Drivers

- Address Space
  - IPv4 Global Address Space depletion
  - IPv4 Private Address Space depletion
- New Services and Applications
  - Design or Re-Design services leveraging IPv6 resources
  - New Service models
  - Improve manageability of existing services and networks
- Government Policies
  - Mandates
  - Tax Incentives

Global IPv4 Address Space Exhaustion

IP Address Allocation History
Full discussion at: www.cisco.com/ipj
The Internet Protocol Journal
Volume 8, Number 3, September 2005

- Consumption is accelerating despite increasingly intense conservation efforts

  PPP / DHCP (temporal address sharing)

  CIDR (classless inter-domain routing)

  NAT (network address translation)

  plus some address reclamation

- Growth is occurring in all regions

  While growth as seen in the routing system is strongest in Asia, the allocation growth is strongest in Europe

- **"ARIN Board Advises Internet Community on Migration to IPv6"**

  Posted: Mon, 21 May 2007

  (http://www.arin.net/announcements/20070521.html)

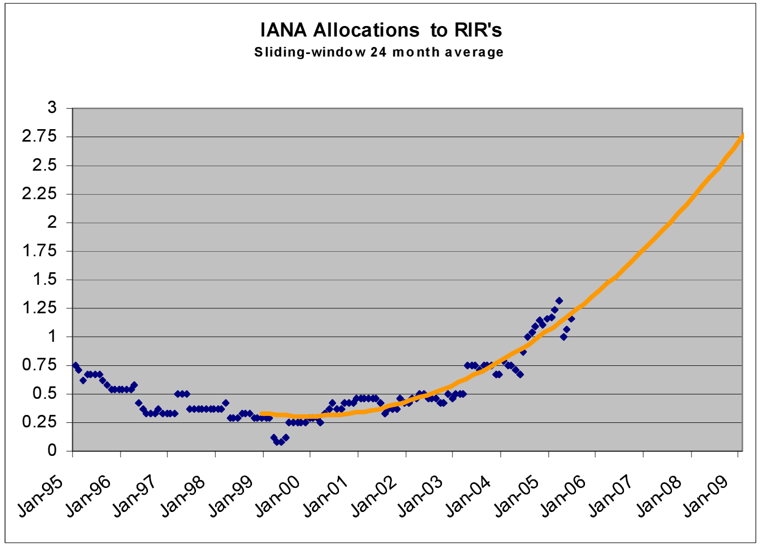

Projection based on IANA* data from 2000

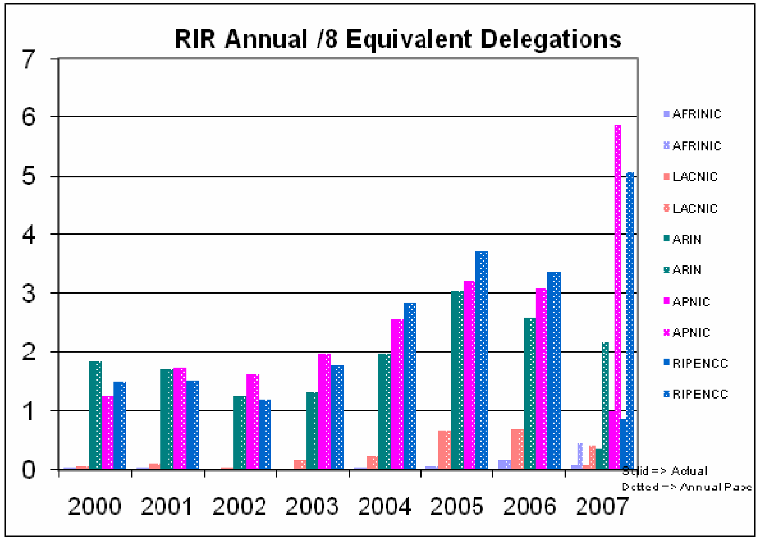

Private IPv4 Address Space Exhaustion

- RFC 1918 is not sufficient for large environments
  - MSOs
  - Mobile providers (fixed/mobile convergence)
  - Large Enterprises
- ARIN turned down the request to increase RFC 1918 private address space; the guideline is to use global addresses instead. This leads to an accelerated depletion of the global address space

New Services—Content Delivery

NTT-East rolled out native IPv6 multicast services instead of IPv4 offering IP TV, Music and games: http://www.ipv6style.jp/en/action/20040902/index.shtml

New Applications

Today, all O.S. Are Dual-Stack

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

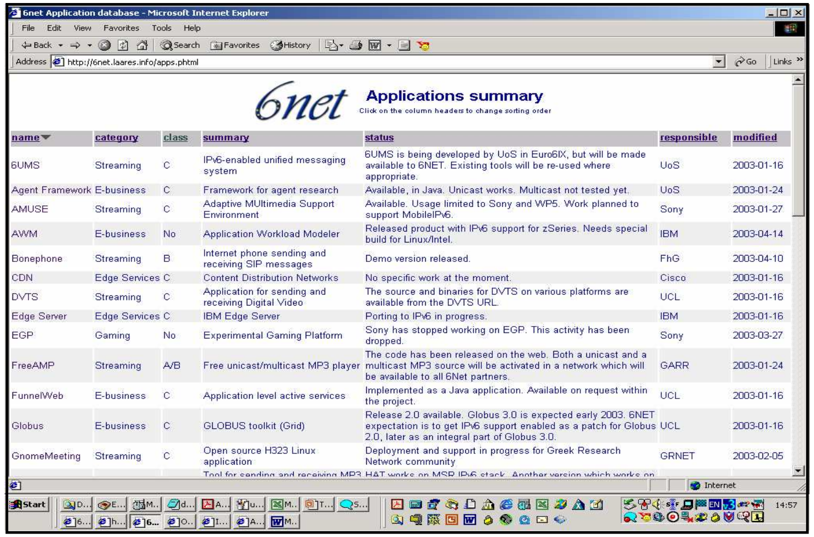

name	category	class	summary	status	responsible	modified
6UMS	Streaming	C	IPv6-enabled unified messaging system	6UMS is being developed by UoS in Euro6IX, but will be made available to 6NET. Existing tools will be re-used where appropriate.	UoS	2003-01-16
Agent Framework	E-business	C	Framework for agent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

Call for Applications—Protocol Agnostic

New Generation of Internet Appliances

National Strategies on IPv6

Various Approaches

Government encouragement + financial and legal incentives

Government encouragement + support for research projects

Lead by example

Adjust government acceptance policies

http://www.cisco.com/en/US/products/ps6553/products_white_paper0900aecd8032b2ad.shtml

Today's networks and services

ISP Deployment Activities

- Several Market segments
  - IX, Carriers, Regional ISP, Wireless
- ISP have to get an IPv6 prefix from their Regional Registry
  - http://www.ripe.net/ripencc/mem-services/registration/ipv6/ipv6allocs.html
- 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

Allocation Recommendations

- Beyond the ISP boundary there are allocation recommendations and not policies
- It is recommended that ISP assign /48 to their customers:

  http://www.ripe.net/ripe/docs/ipv6policy.html

  http://www.icann.org/announcements/ipv6-report-06sep05.htm

- In practice, it is most likely that the allocations below ISP will be between /48 and /64.

What to Expect?

No Rules or Strong Guidelines Have Been Posted for Address Planning and Design—Hence the Unknown Challenges:

- We don’t know what are the consequences of relaxed allocation policies
- There is a lot of experimentation with the many available bits

http://www.ietf.org/internet-drafts/draft-vandevelde-v6ops-addcon-03.txt

IPv6 Routing Protocols

Overview

IPv6 Routing Protocols

Ipv6 Is an Evolutionary Not a Revolutionary Step and This Is Very Clear in the Case of Routing Which Saw Minor Changes Even Though Most of the Routing Protocols Were Completely Rebuilt:

- IPv6 EIGRP ⇔ EIGRP
- IPv6 IS-IS
- OSPFv3 ⇔ OSFPv2
- IPv6 extensions for BGP

EIGRP for IPv6

Cisco Proprietary Protocol

EIGRP for IPv6 Overview

- Just another protocol module (IP, IPX, Appletalk) with three new TLVs:
  - IPv6_REQUEST_TYPE (0X0401)
  - IPv6_METRIC_TYPE (0X0402)
  - IPv6_EXTERIOR_TYPE (0X0403)
- Other similarities
  - Same protocol number 88
  - Router ID stays 32 bits (must be configured explicitly if there is no IPv4 interface on the router)
  - Uses MD5 like for IPv4 (IPsec authentication will be available soon)
  - Same metrics

EIGRP for IPv6 Specific Features

Several IPv6 Specific Differences with Respect to IPv4:

- Hellos are sourced from the link-local address and destined to FF02::A (all EIGRP routers); this means that neighbors do not have to share the same global prefix (with the exception of explicitly specified neighbors where traffic is unicasted)
- Automatic summarization is disabled by default for IPv6 (unlike IPv4)
- No split-horizon in the case of EIGRP for IPv6 (because IPv6 supports multiple prefixes per interface)

EIGRP for IPv6—Configuration and Display

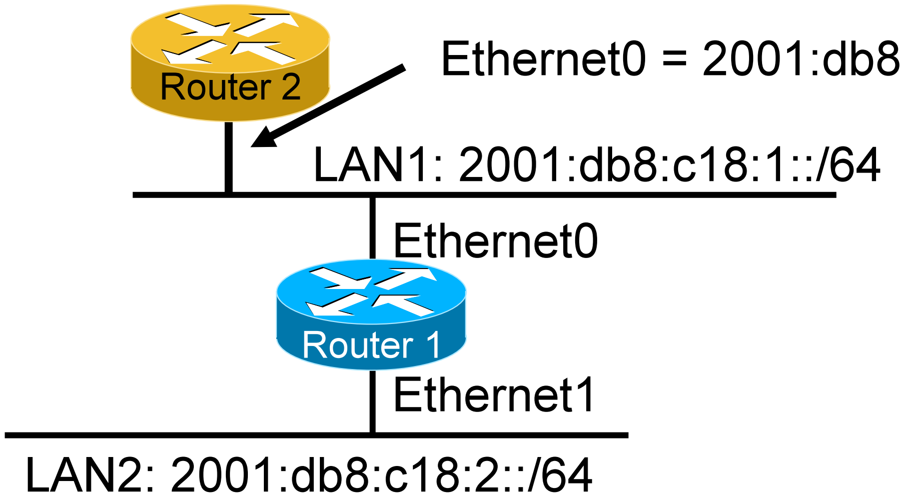


```
Router2#
ipv6 router eigrp 100
 Router-id 10.10.10.1

interface Ethernet0
 ipv6 address 2001:db8:c18:1::/64 eui-64
 ipv6 enable
 ipv6 eigrp 100
```

```
Router1#show ipv6 eigrp neighbor
IPv6-EIGRP neighbors for process 100
H   Address                  Interface     Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)        Cnt Num
0   FE80::260:3eff:fe47:1530   E0           14 00:01:43    1     4500  0   1
```

```
Router1#show ipv6 eigrp topology all-links
IPv6-EIGRP Topology Table for AS(100)/ID(10.10.10.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 2001:db8:c18:1::/64, 1 successors, FD is 28160, serno 1
        via Connected, Ethernet0
        via FE80::260:3eff:fe47:1530 (30720/28160), Ethernet0
```

Neighbor identified by Link-local address

ISIS

IS-IS fundamentals

- Fundamentally 3 types of packet
  1. Hello
  2. Link State Packets
  3. Sequence number packets

ISIS Fundamentals

- TLV's

TLV	Type
IP Internal Reachability information	128
Protocol Supported	129
IP interface address	132
IP External Reachability Information	130
Extended IP Reachability	135
IPv6 Reachability	236
IPv6 Interface Address	232

ISIS Fundamentals

- IP Internal Reachability information (TLV 128)

  Stores lists of directly connected IP prefixes

  Each prefix is assigned a metric value, which is that of the link over which the IP prefix is configured

  Value is 12 bytes: 4 bytes metric, 4 byte IP prefix, 4 byte mask

ISIS Fundamentals

- Protocol Supported TLV 129
  - Defines the supported layer 3 protocols
  - This appears in LSP number 0
  - Currently supported protocols

CLNP	0x81
IPv4	0xCC
IPv6	0x8E

ISIS Fundamentals

- IP Interface Address TLV 132
- Depending upon the type of packet its function differs slightly

  IP interface address of the outgoing interface

  One or more IP address associated with the router, in current Cisco IOS this is highest loopback address of the router

ISIS Fundamentals

- IP External Reachability Information TLV 130

  All routes redistributed into ISIS, source is from other routing protocols

  Only allowed in level 2 but we allow is for L1 as well

  We only support default metric

ISIS Fundamentals

- Extended IP Reachability TLV 135

  Delay, Expense and Error were for QOS for CLNP packet header, these were optional path selection service for network

  TLV is used to extend the metric from 6 to 32 bits

0	**I/E**	**Default metric**
S	**R**	**Delay Metric**
S	**R**	**Expense Metric**
S	**R**	**Error Metric**
IP Address		
Subnet Mask		

Default metric		
U/D	**Sub TLV**	**Prefix Length**
Prefix (0-4 bytes)		

IPv6 New TLV's

- IPv6 Reachability TLV 236

  Defines both IPv6 Internal and External reachability information

  Metric is still 32 bits

  U: Up/Down

  X: External origin bit

  S: Sub-TLV present

  Prefix length: Length of prefix 8 bits

  Prefix: Number of octet is calculated depending on the prefix length

IPv6 New TLV's

- IPv6 address TLV 232

  Modified to carry IPv6 address

  For hello PDU interface address must use link local IPv6 address assigned to the interface

  For LSP non-link local address must be used

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

Multi Instance ISIS

- Multiple instances enhances the ability to isolate the resources associated with instance
  - Within a router
  - Across the network
- You can have instance specific prioritization
- Flooding prioritization can give network wide resource allocation

OSPFv3

RFC 2740

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-LSA	3	0x2003
Inter-Area-Router-LSA	4	0x2004
AS-External-LSA	5	0x4005
Group-membership-LSA	6	0x2006
Type-7-LSA	7	0x2007
Link-LSA	8	0x2008
Intra-Area-Prefix-LSA	9	0x2009

New

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
```

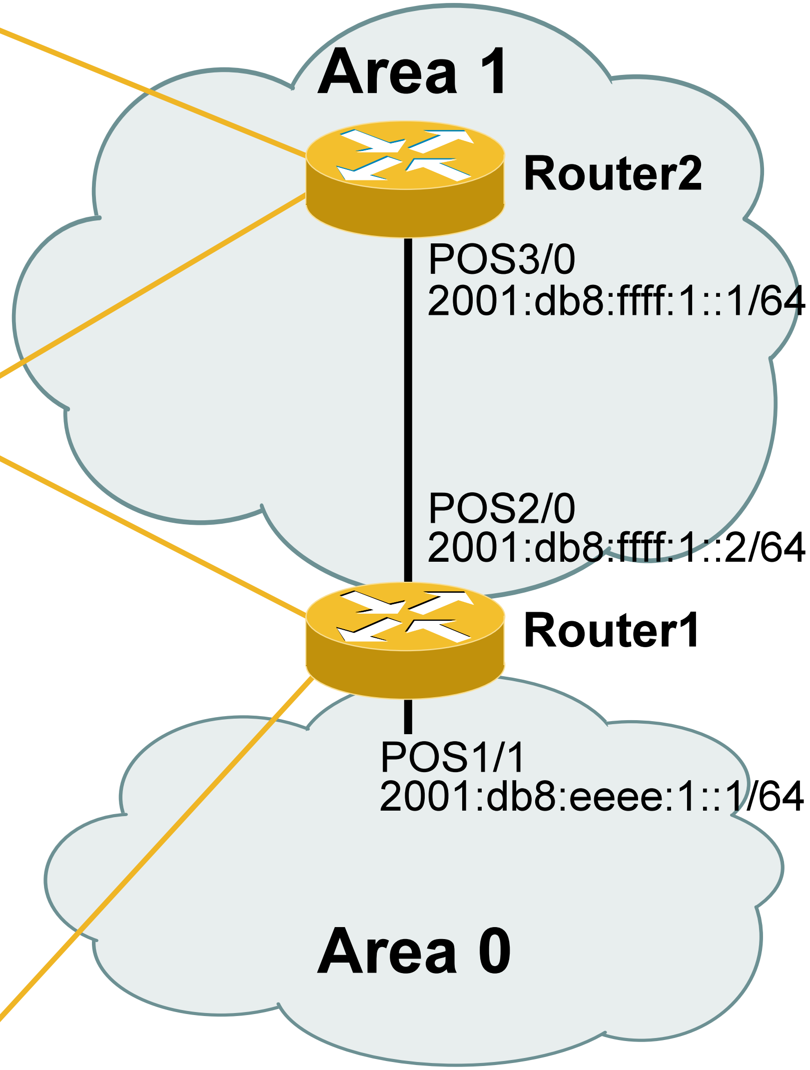

MP-BGP

RFC 2545

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

Selecting an IPv6 IGP

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?
- Operational experience

IPv6 IGP Selection…in Theory

In **Theory**:

- The similarity between the IPv6 and IPv4 routing protocols leads to similar behaviour 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 Fast Convergence optimizations might be available
- Operational experience with large scale IPv6 networks is being developed

Conclusions

- Same topology considerations as for IPv4
- Convergence time
  - When comparing apples to apples the convergence times are very similar
  - Other tools are also leveraged: Bidirectional Forwarding Detection (BFD)
  - There are HW and SW dependencies

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?
- How are the resources shared between the two protocols?
- Are the topologies going to be congruent?
- …

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

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

Some Preliminary Thoughts…

- Requiring parity is natural but that does come at a cost so it is worth considering if parity is needed or worth the cost
- The IPv6 IGP configuration can change as deployment progresses from trial to large scale production
- Most commonly IPv6 is deployed on a service by service basis so it is possible to have different convergence requirements for IPv6 than IPv4

Revisiting ISIS Single Topology

Multi Process/Topo	Single Process/Topo*
• Clear separation of the two control planes • Non-congruent topologies are very common if not desired in deployments	• Requires less resources • Might provide a more deterministic co-existence of IPv4 and IPv6

*Today most IPv6 IGPs are distinct from their IPv4 counterparts and will run as ships in the night. The only exception was ISIS.

ISP infrastructure

Service Provider networks

- Major routing information more then few 100 K via BGP
- Large IGP routing table are ~6–7K expected to grow farther
- 6K/117K ~ 5% of IGP routes in an ISP network
- A very small factor but has a huge impact on network convergence!

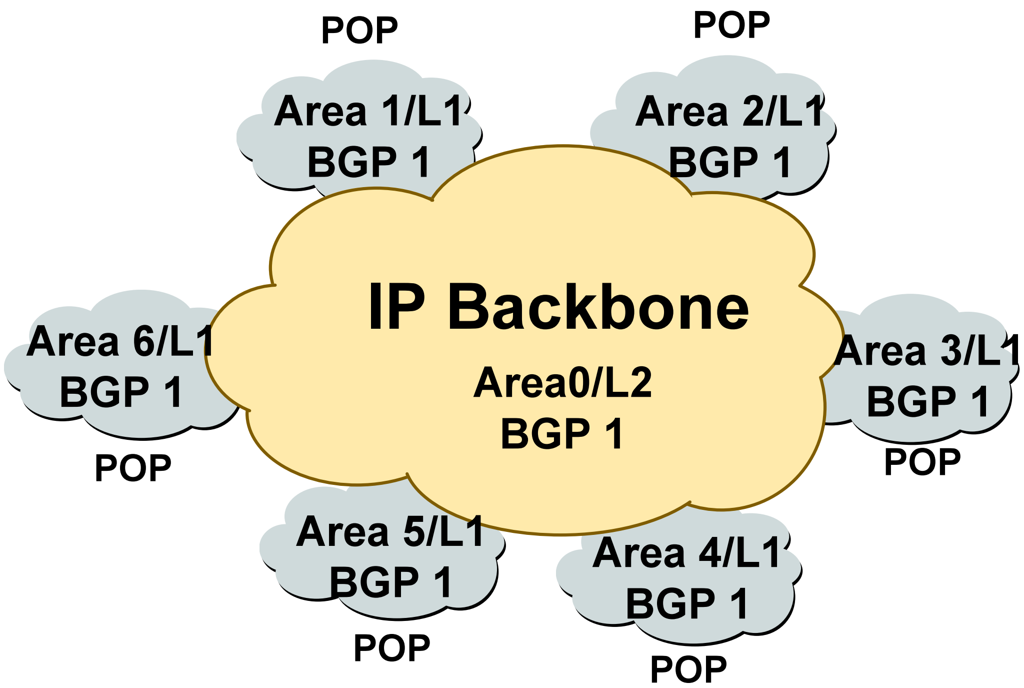

Service Provider networks

- 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

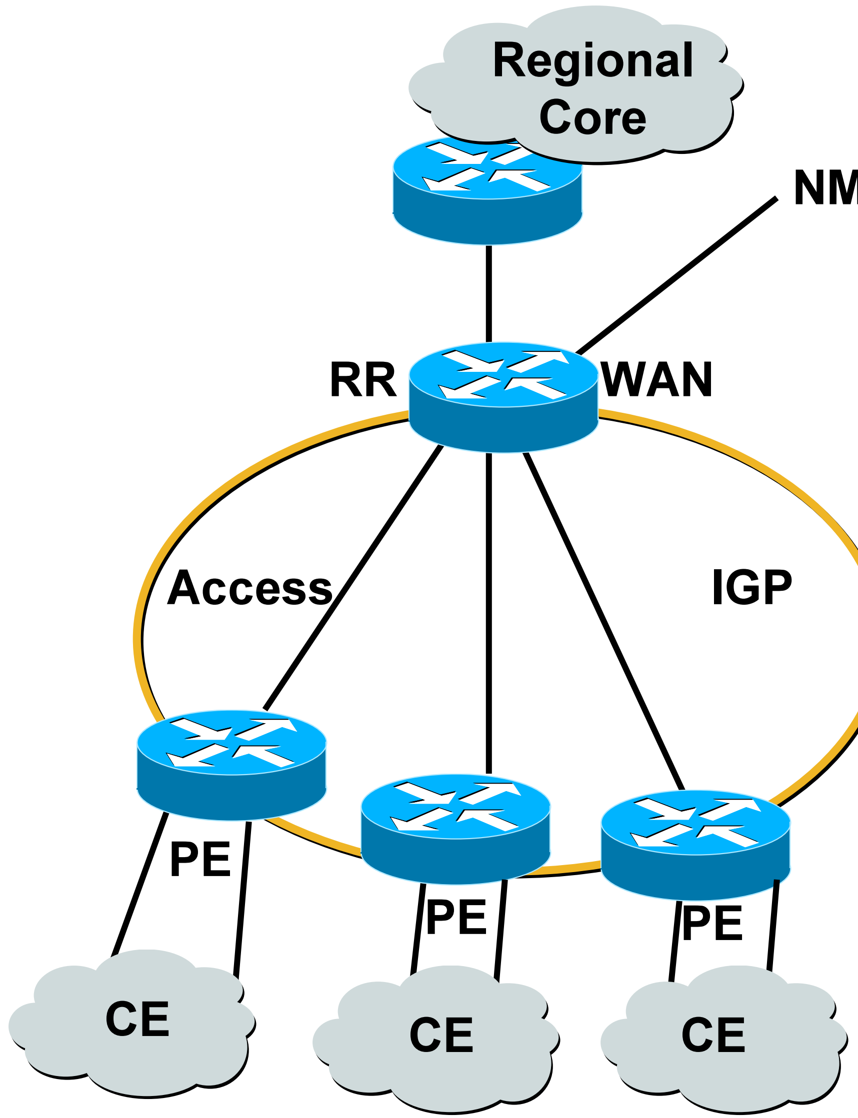

Addressing

- 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

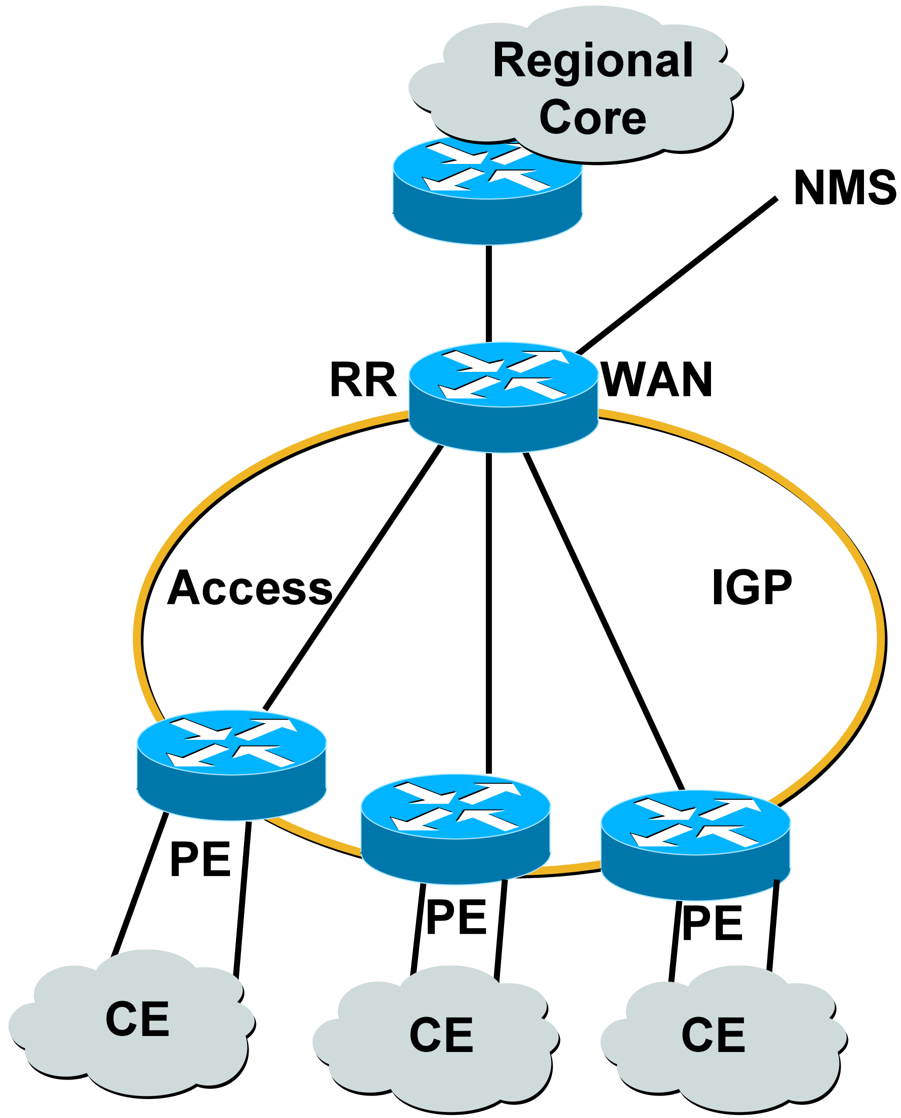

Addressing

- Divide the address into two parts
  1. Physical links
  2. Loopback interfaces
- Physical address should be a contagious
- Loopback should be from public address space
- Optimal path to the next hop is necessary

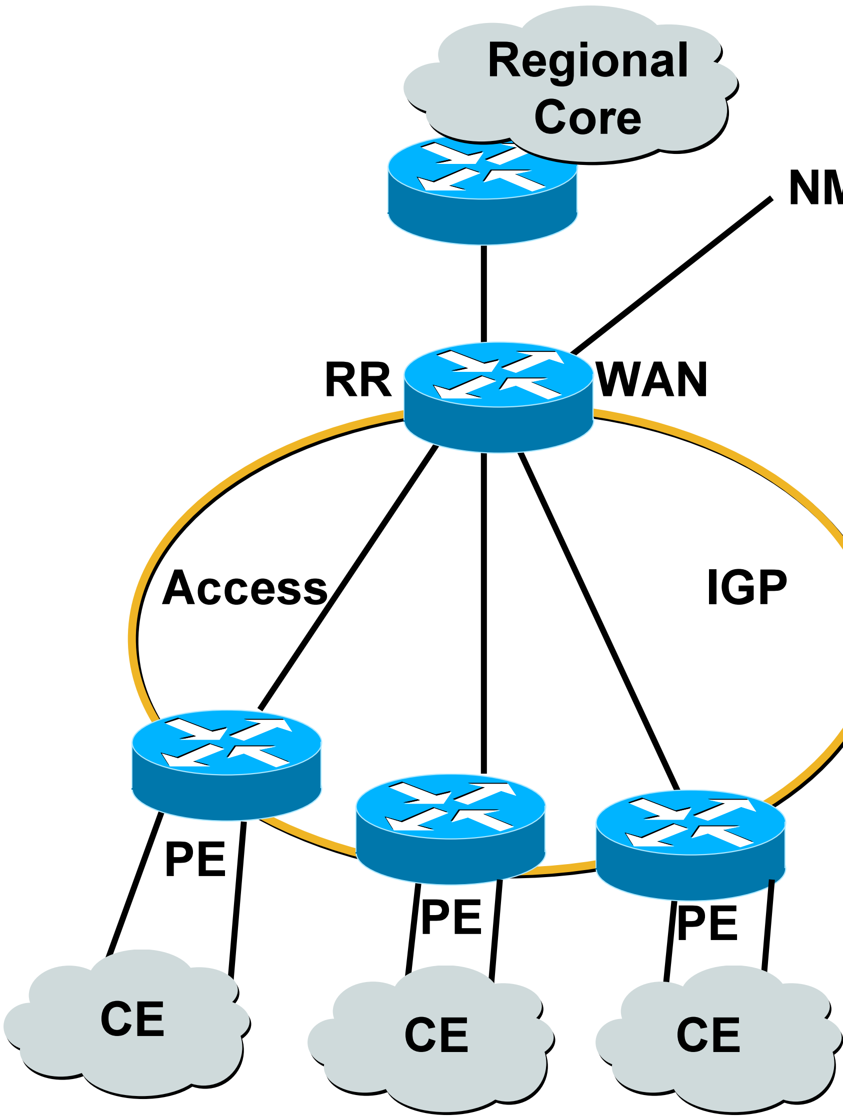

Addressing

- 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

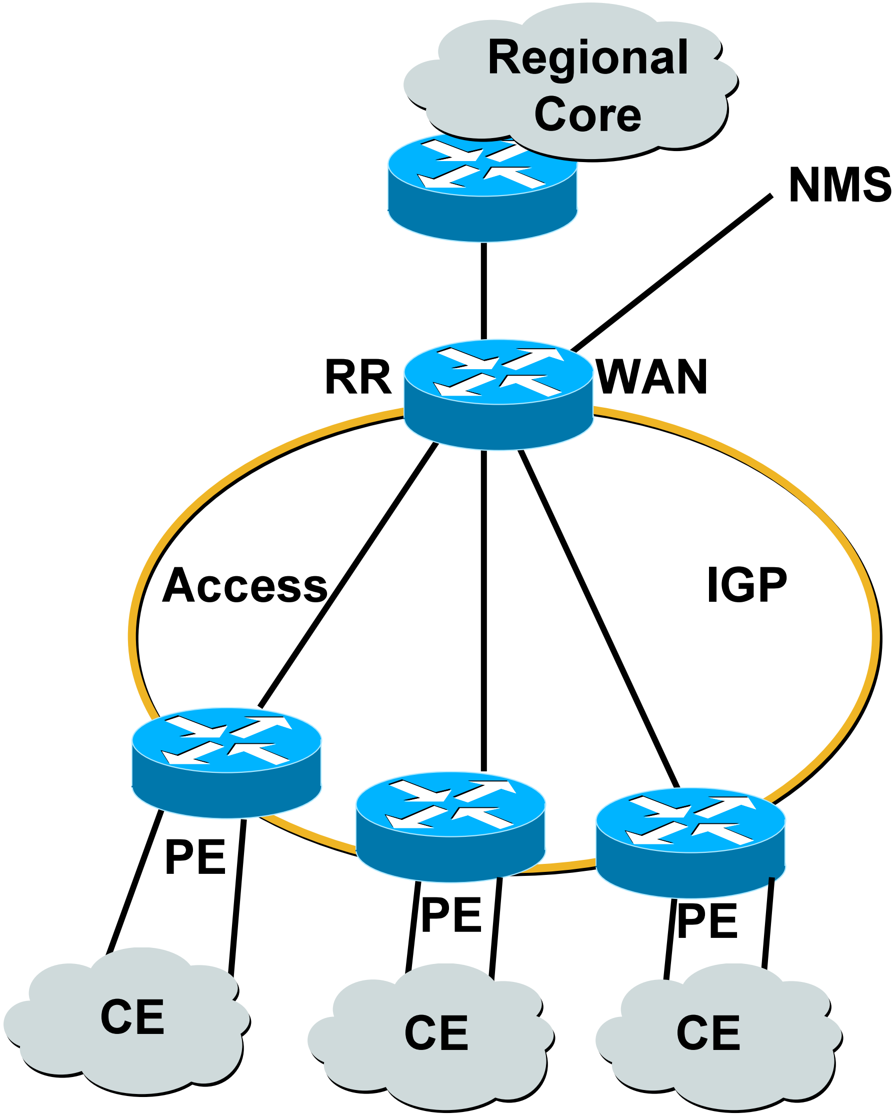

IPv6 over MPLS

IPv6 over MPLS
Deployment Scenarios

- **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/VPN, MPLS/QoS, MPLS/TE, ATM + IP switching, L2 services (AToM, EoMPLS)**
- **MPLS can be used to facilitate IPv6 integration**
- **Multiple approaches for IPv6 over MPLS:**
  - **IPv6 CE-to-CE IPv6 over IPv4 Tunnels**
  - **IPv6 over “Circuit_over_MPLS”**
  - ~~**Native IPv6 MPLS**~~
  - **IPv6 Provider Edge Router (6PE) over MPLS**

IPv6 Tunnels configured on CE

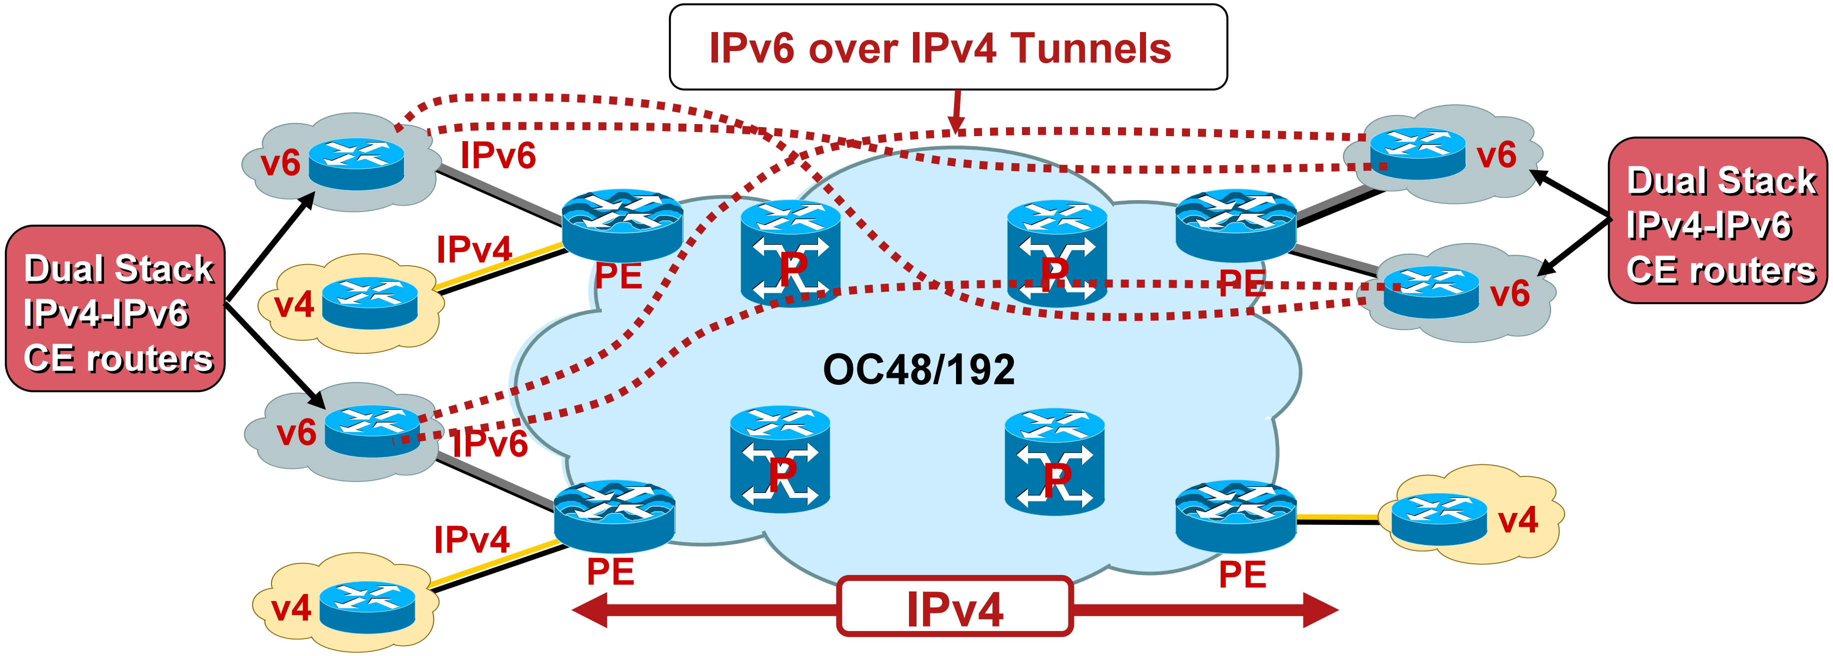

- No impact on existing IPv4 or MPLS Core (IPv6 unaware)
- Only CEs have to be IPv6-aware (Dual stack)
- Mesh of IPv6 over IPv4 Tunnels CE-to-CE
- Overhead: IPv4 header + MPLS header
- MPLS/VPN support IPv4-native and IPv6 tunnels
- Service Provider can’t delegate his IPv6 prefix to the CE routers

IPv6 over "Circuit_over_MPLS"

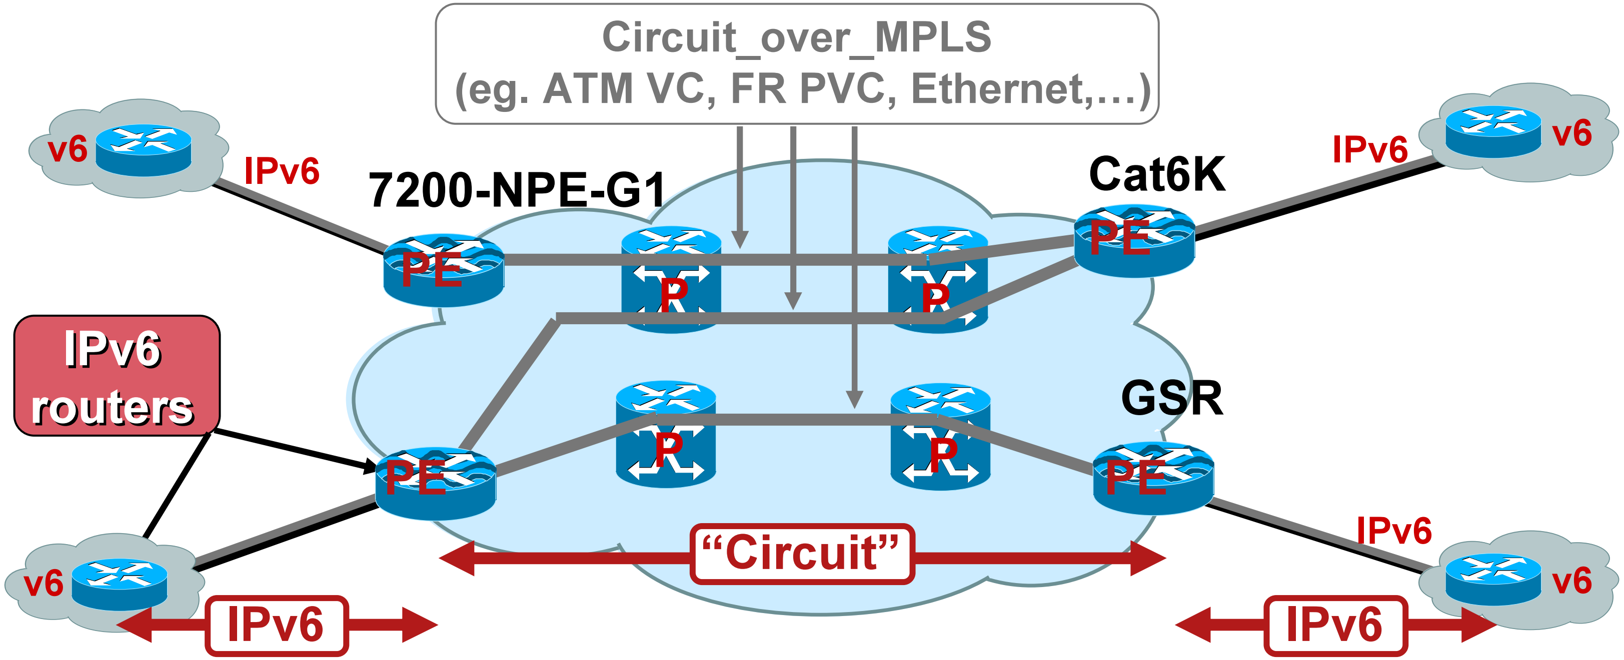

- No impact on existing IPv4 or MPLS Core (IPv6 unaware)
- Edge MPLS Routers need to support "Circuit_over_MPLS"
- Mesh of "Circuit_Over_MPLS" PE-to-PE
- PE routers can also be regular IPv6 Routers (IPv6 over ATM, IPv6 over FR, IPv6 over Ethernet,…) to aggregate Customer's IPv6 routers

Minimum Infrastructure Upgrade for 6PE

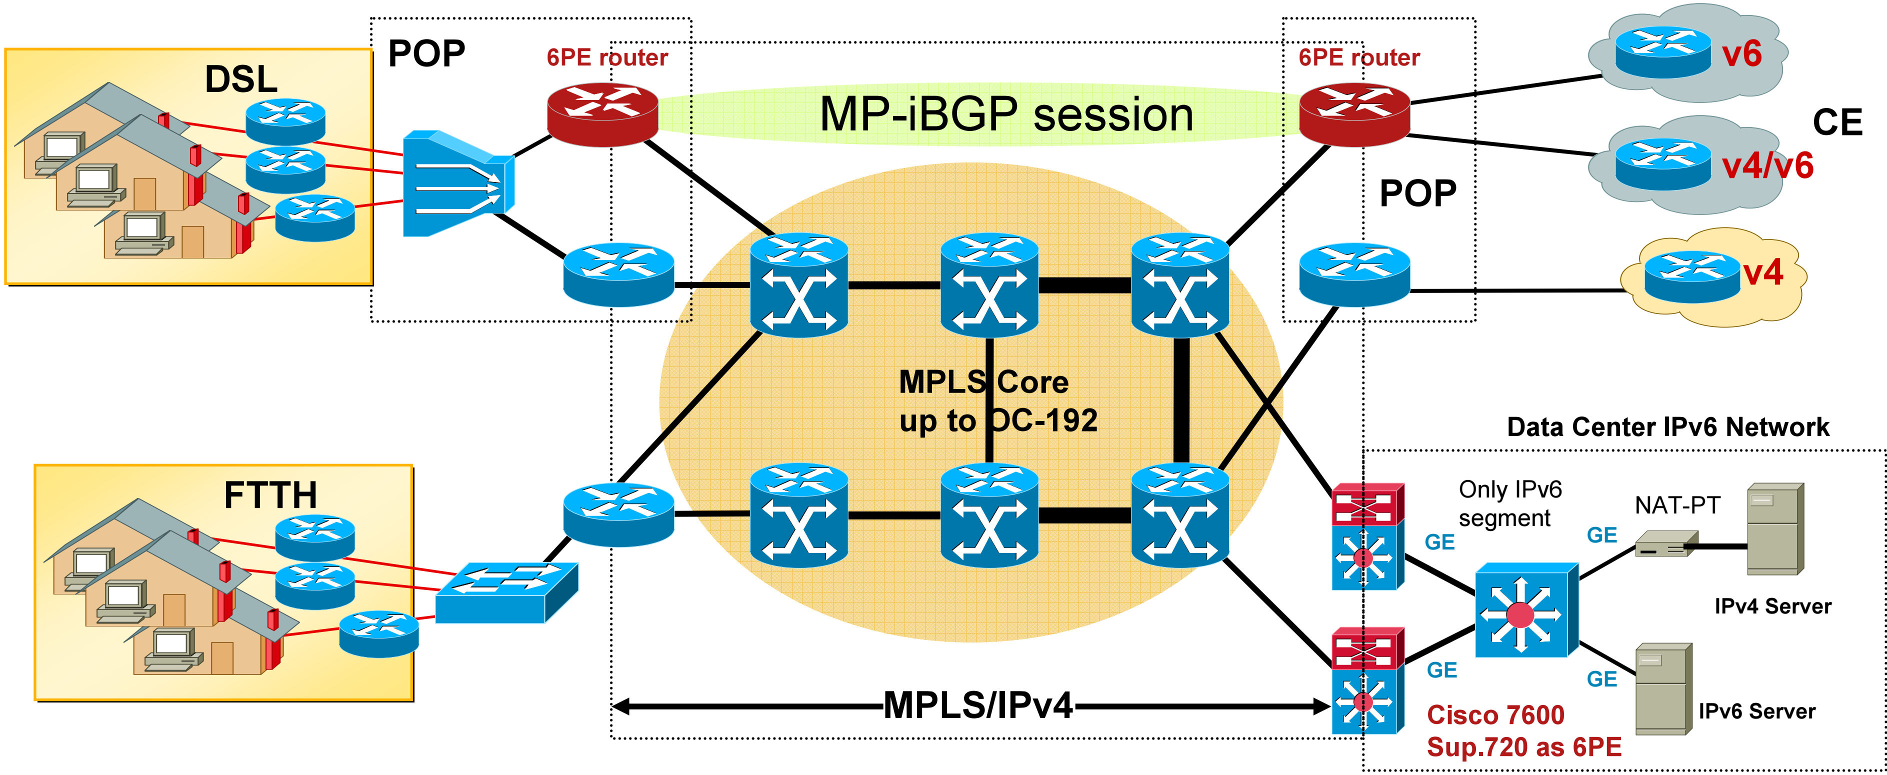

- **MPLS/IPv4 Core Infrastructure is IPv6-unaware**
- **PEs are updated to support Dual Stack/6PE**
- **IPv6 reachability exchanged among 6PEs via iBGP (MP-BGP)**
- **IPv6 packets transported from 6PE to 6PE inside MPLS**

6PE Overview

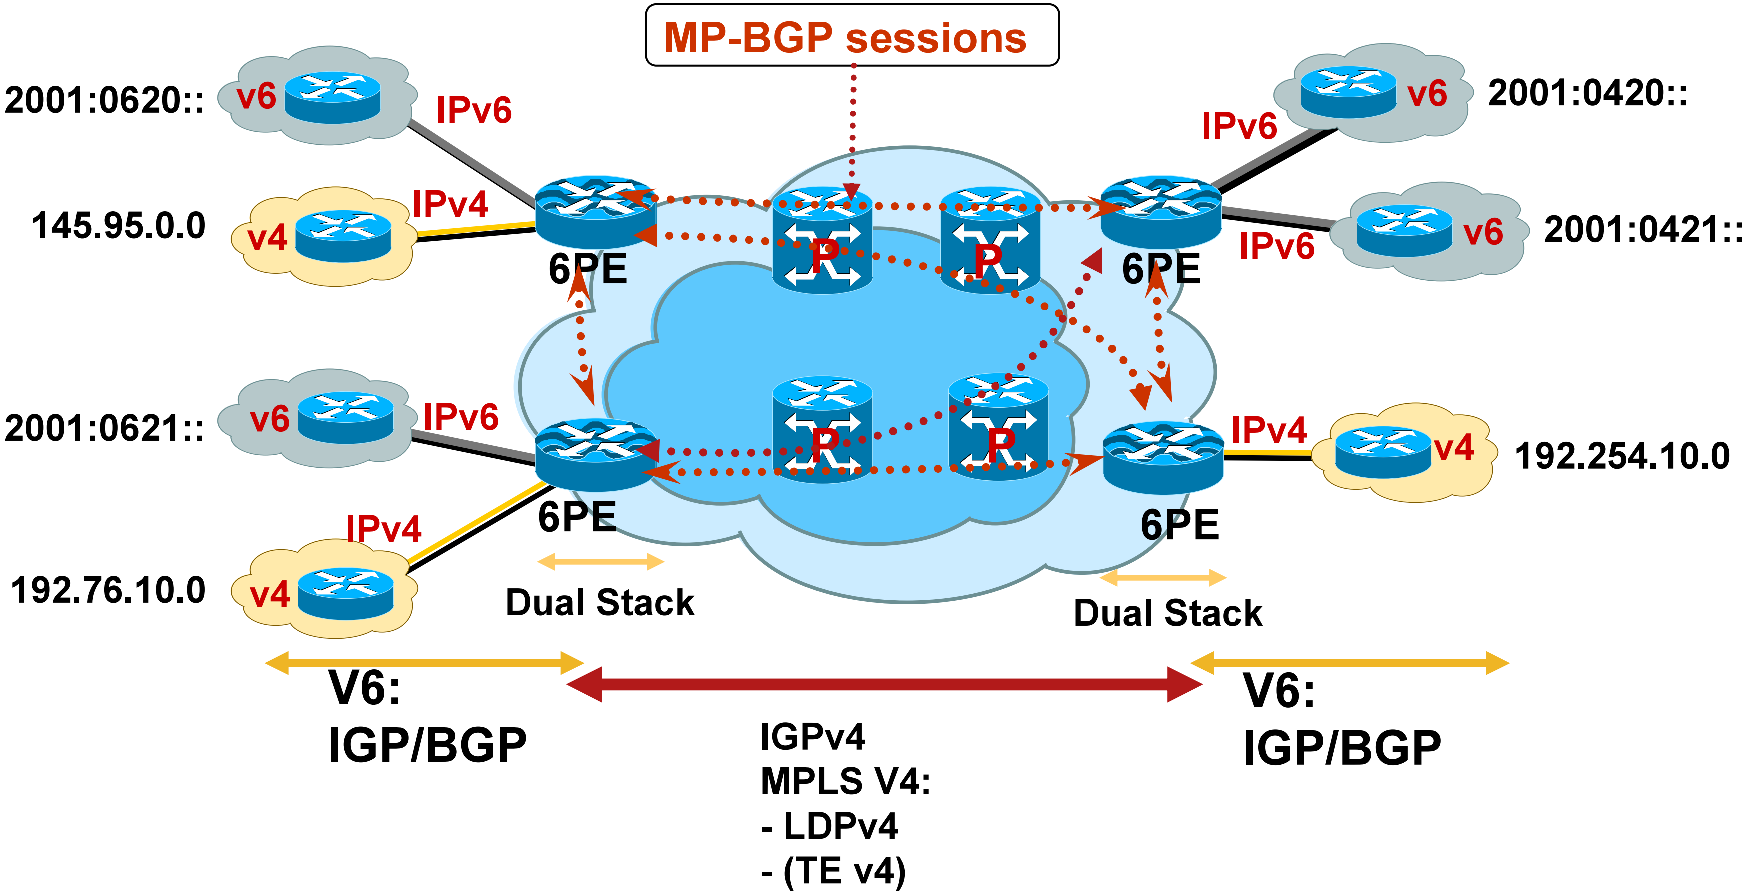

IPv6 unaware
No core upgrade

6PE Routing

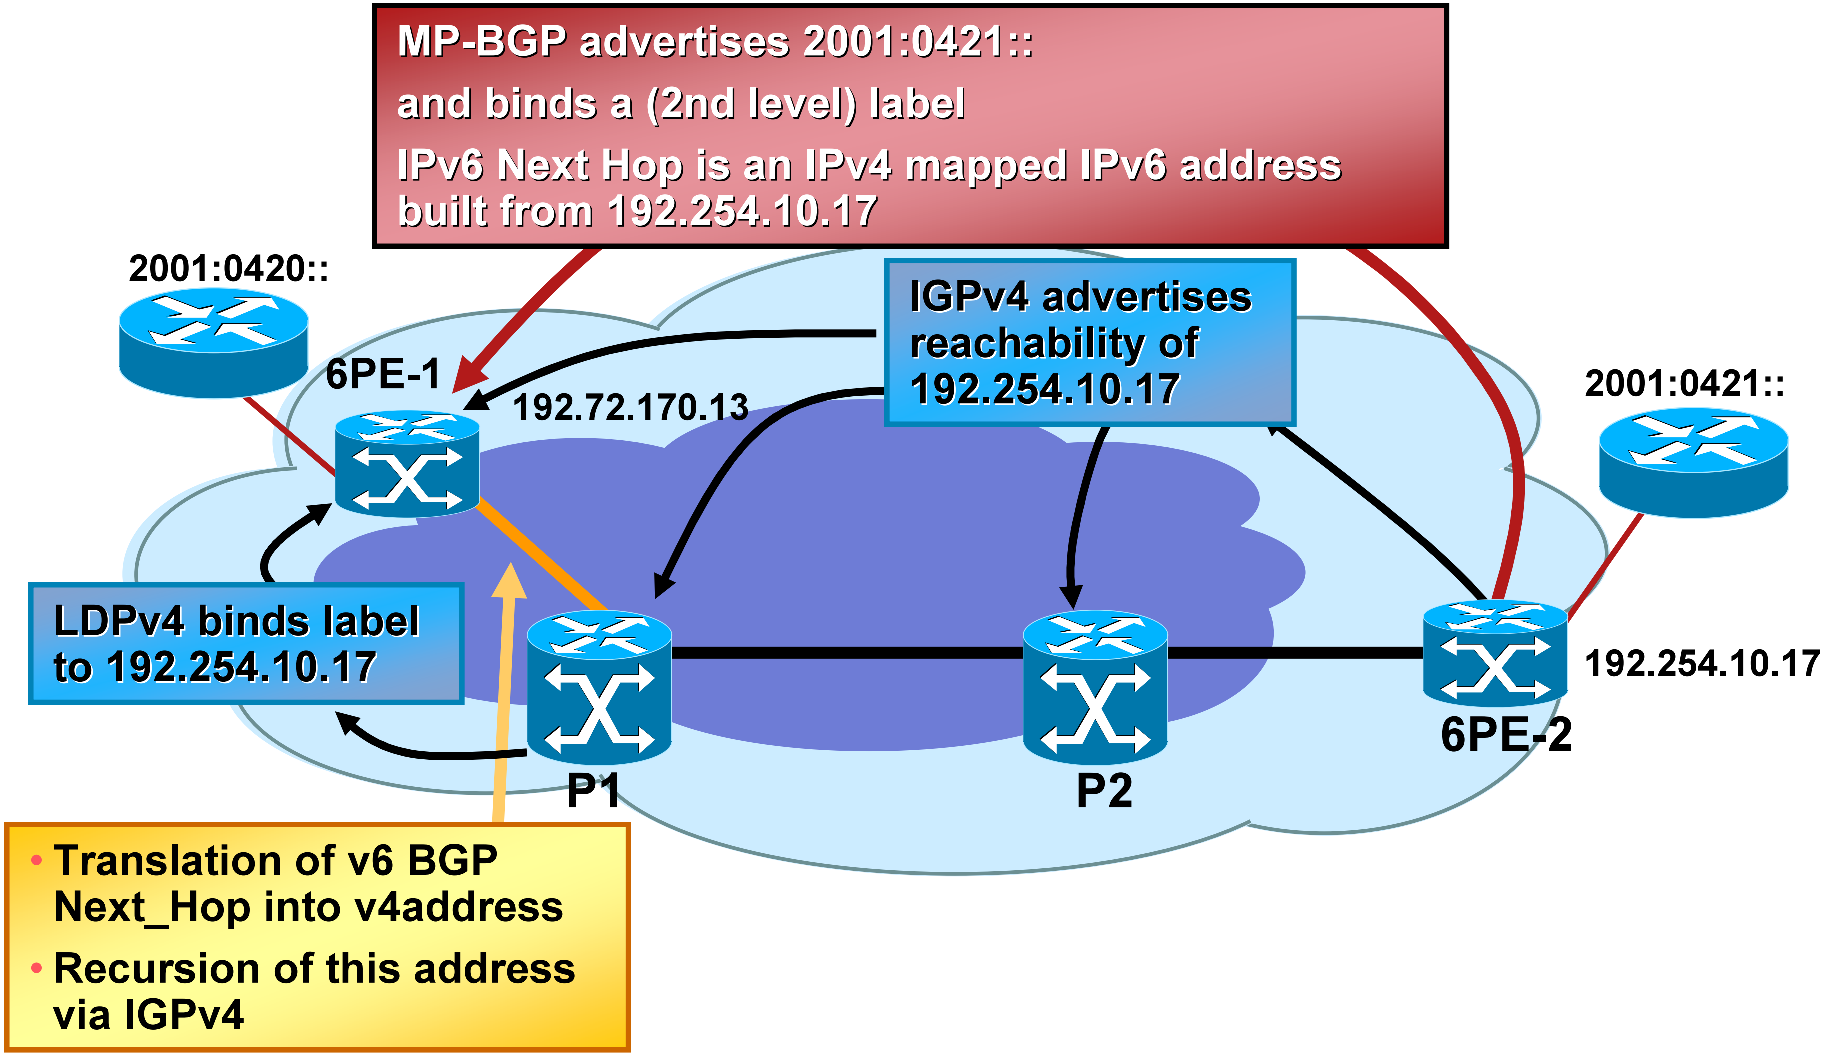
MP-BGP advertises 2001:0421::
and binds a (2nd level) label
IPv6 Next Hop is an IPv4 mapped IPv6 address built from 192.254.10.17
2001:0420::
6PE-1
192.72.170.13
IGPv4 advertises reachability of 192.254.10.17
2001:0421::
LDPv4 binds label to 192.254.10.17
192.254.10.17
6PE-2
P1
P2
• Translation of v6 BGP Next_Hop into v4address
• Recursion of this address via IGPv4

6PE Forwarding

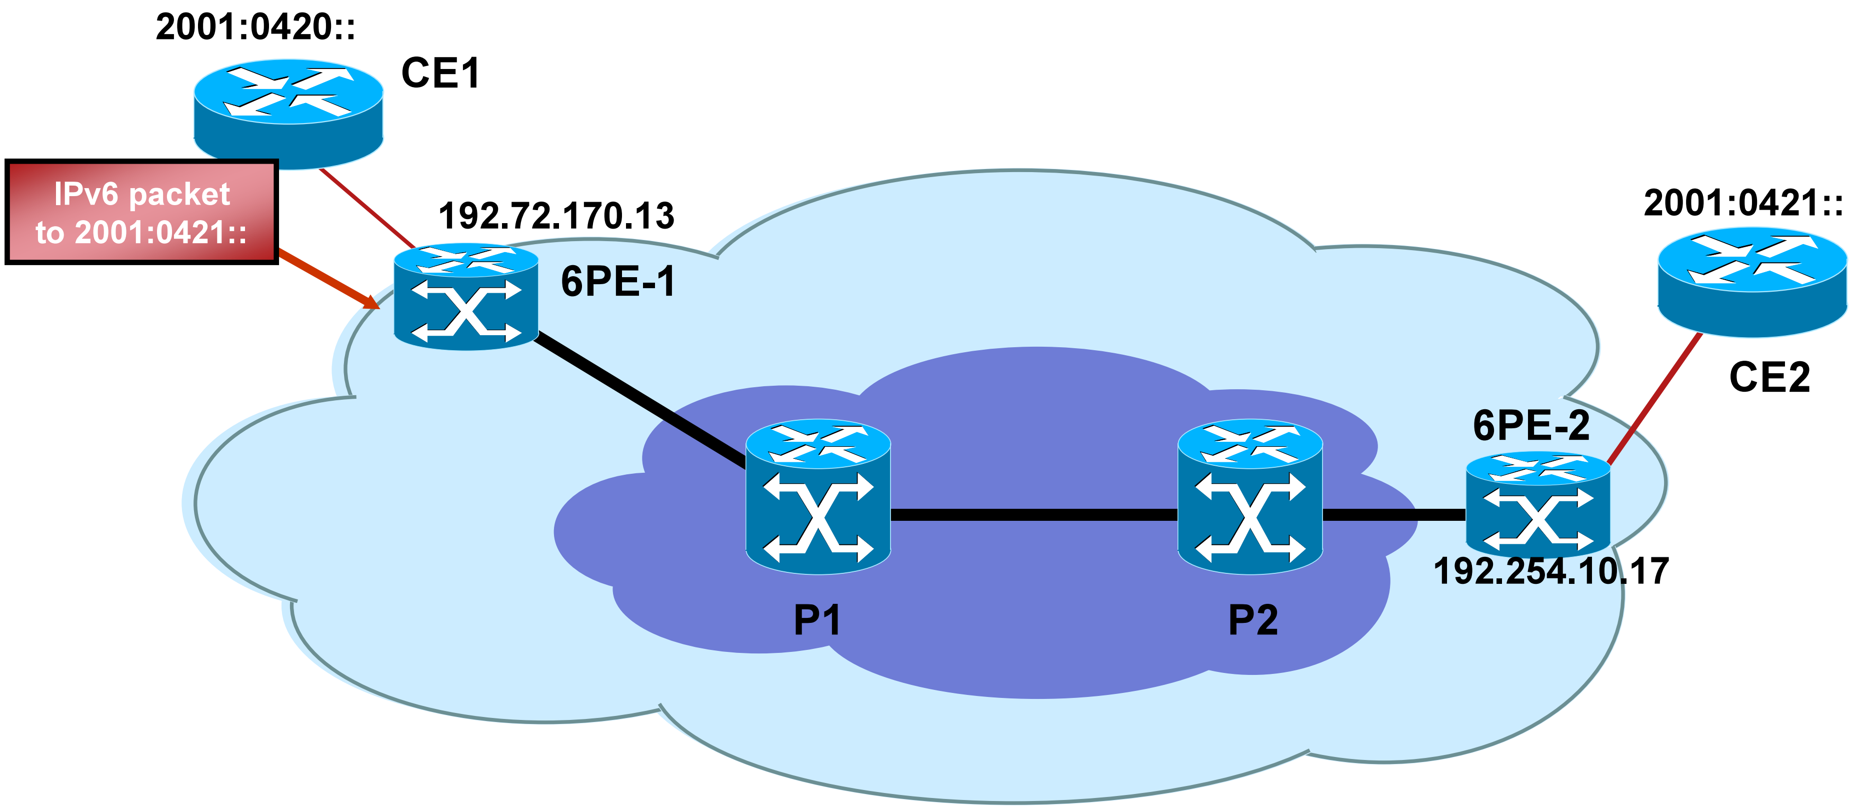

6PE Forwarding (6PE-1)

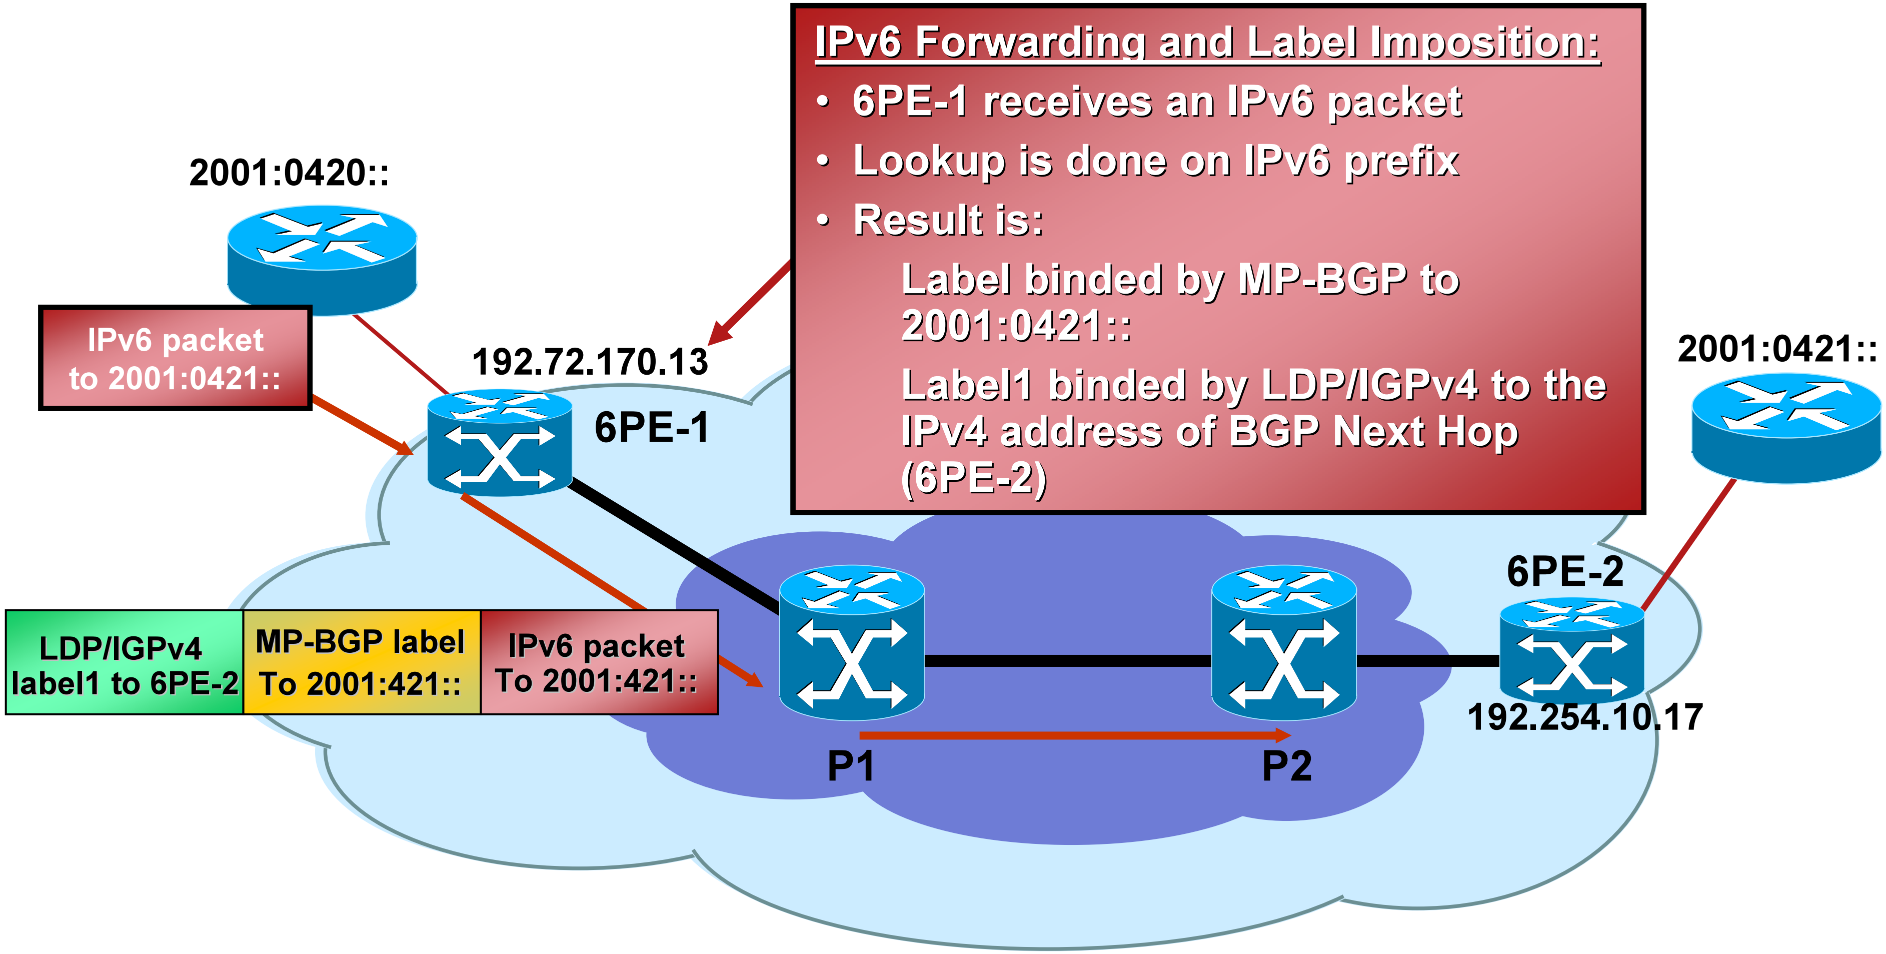

6PE Forwarding (P1)

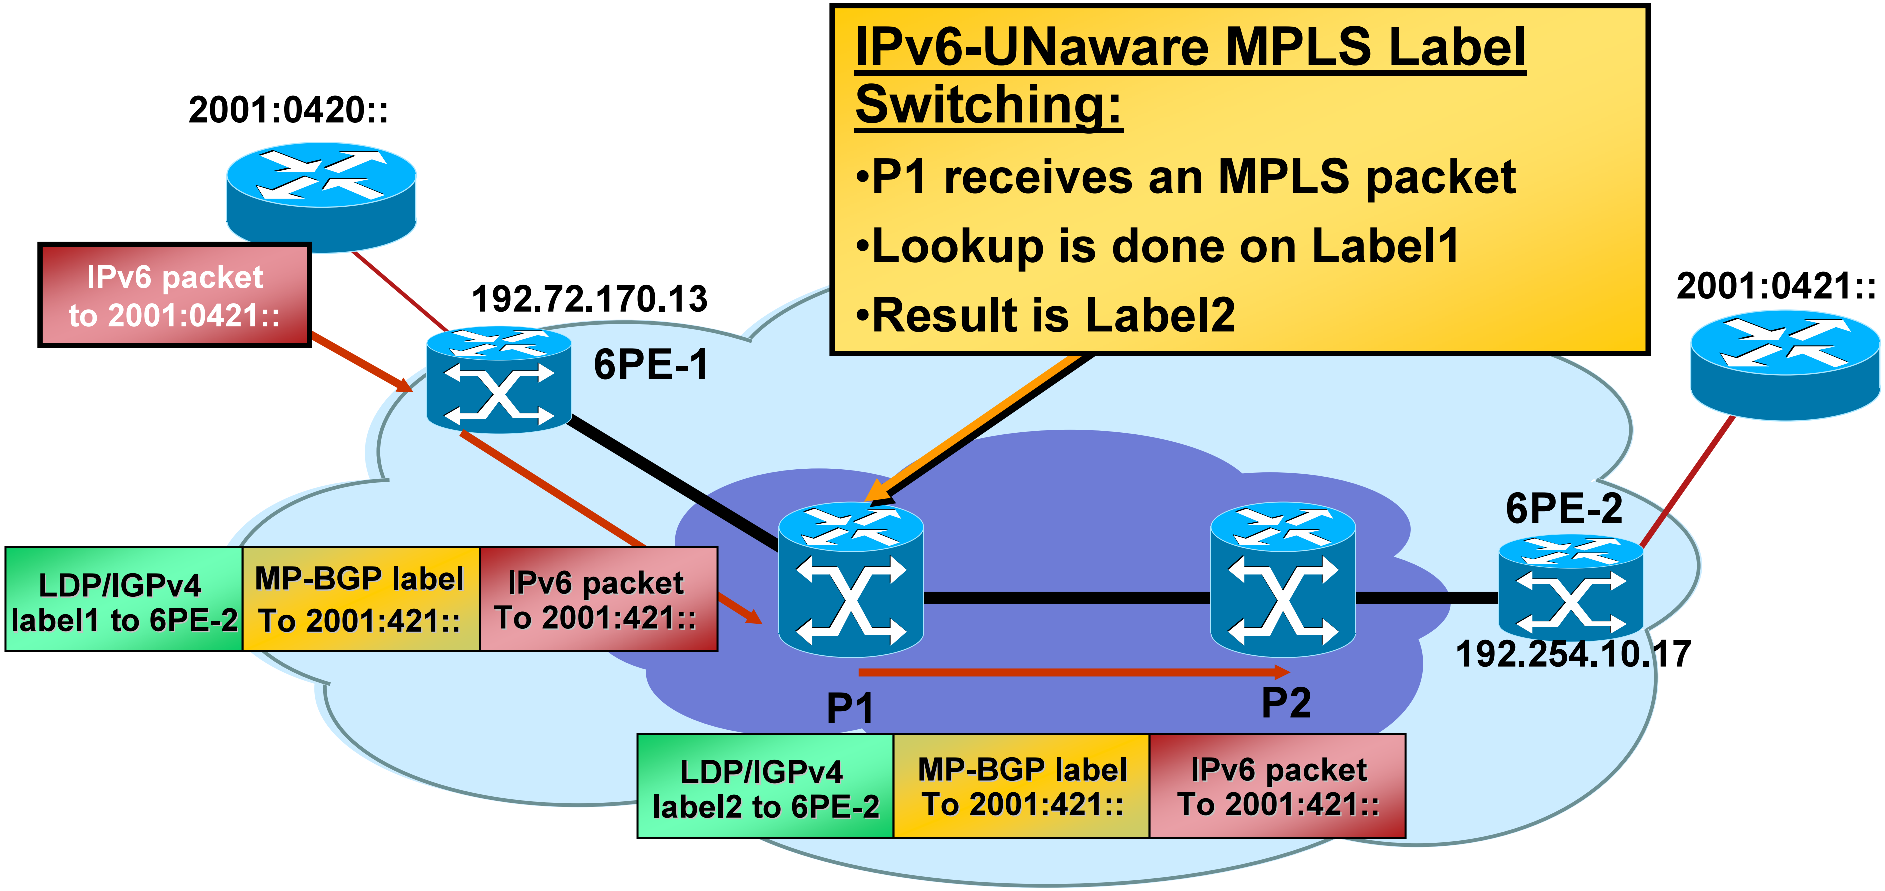

6PE Forwarding (P2)

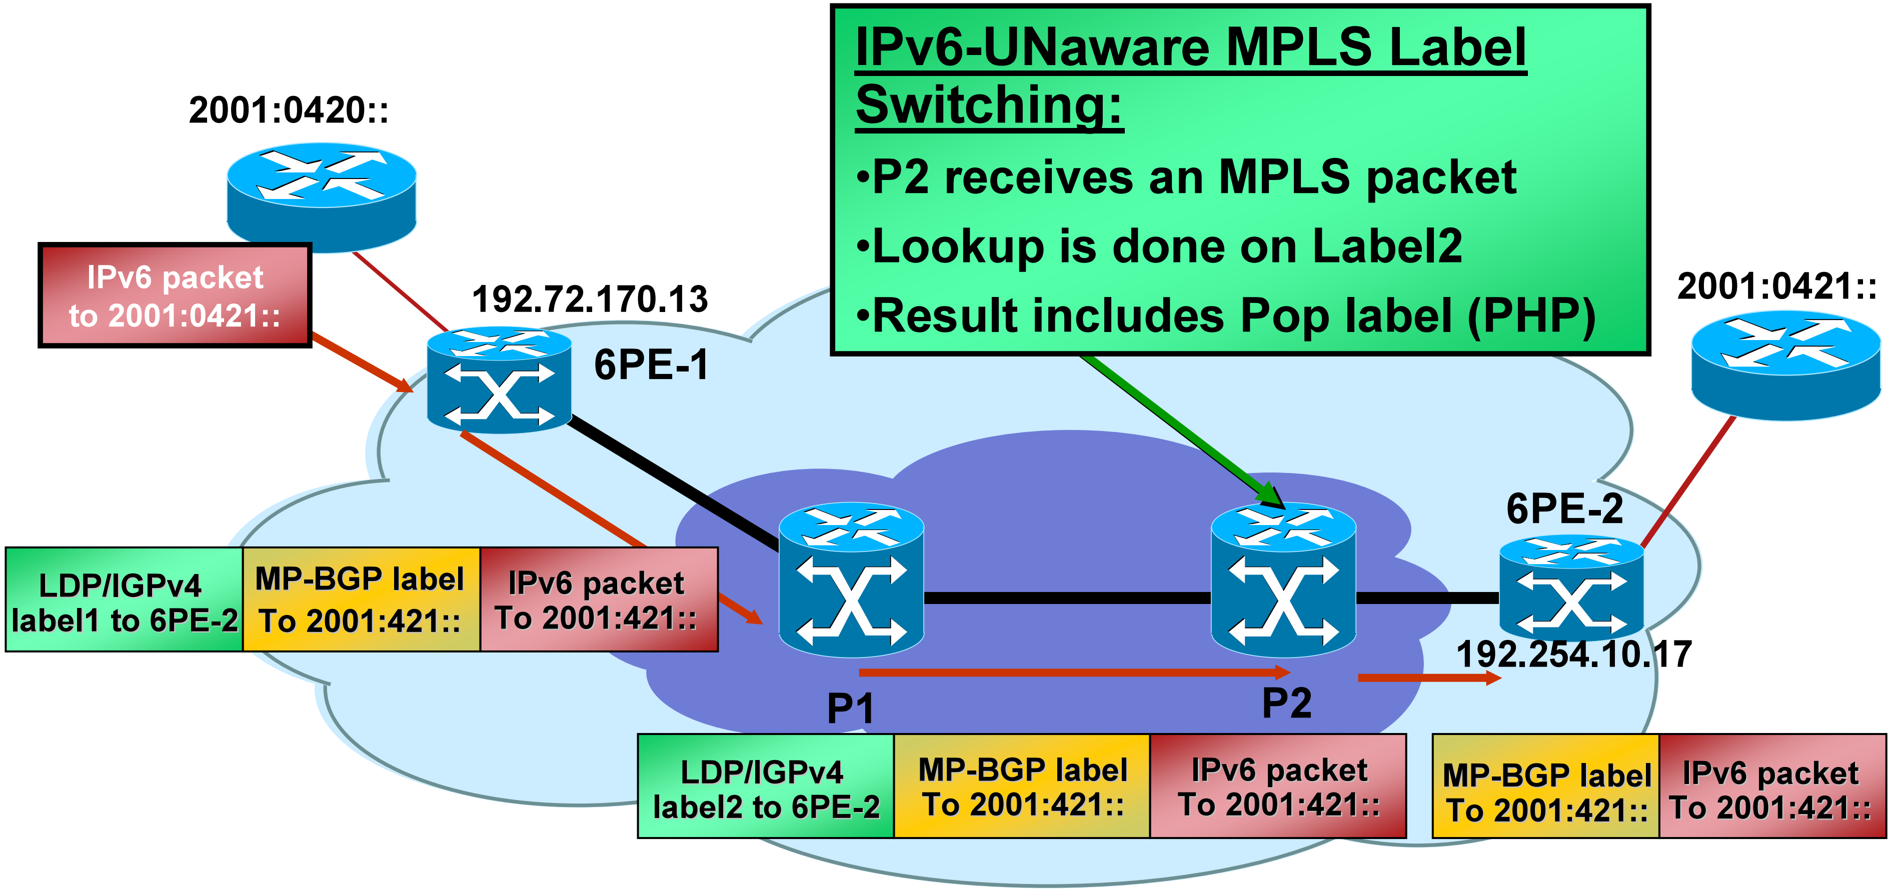

6PE Forwarding (6PE-2)

MPLS Label Pop and IPv6 Forwarding :

- 6PE-2 receives an MPLS packet
- Lookup is done on Label
- Result is:

  Pop the label & switch

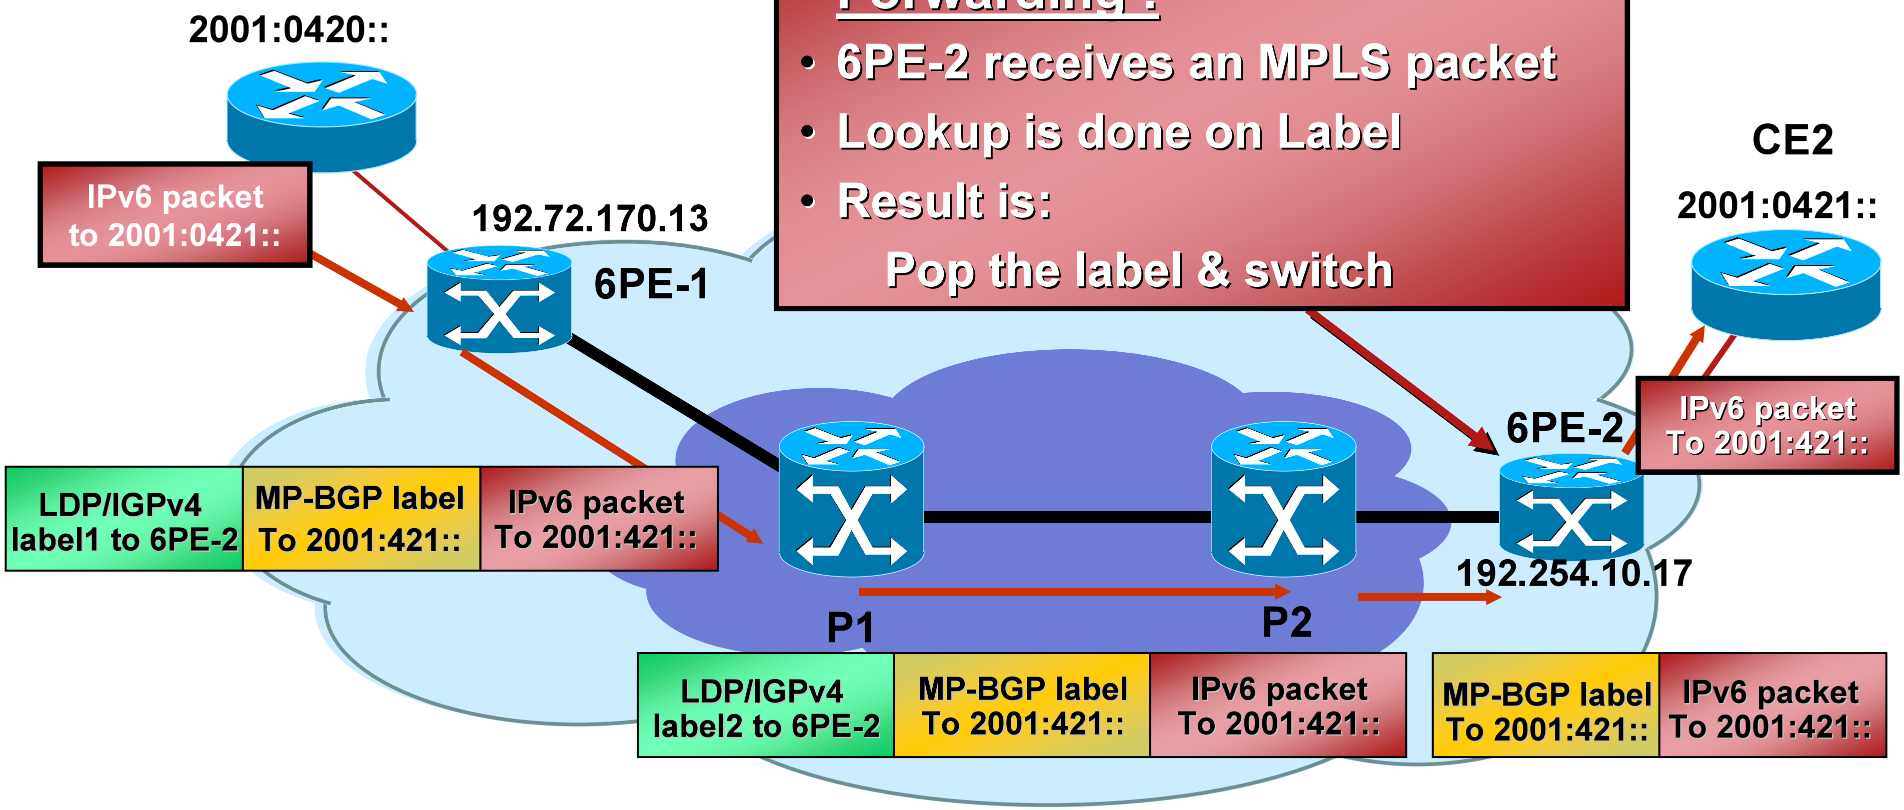

6PE Benefits

For SPs already running MPLS, 6PE approach has many benefits:

- Core Infrastructure needs no upgrade and no configuration change
- Upgrade only on the required edge routers (ie upgrade of existing PEs to 6PE, or add separate 6PEs)
- IPv6 supported simultaneously with existing MPLS services (MPLS v4_VPNs, QoS, ATM, v4 Internet, …)

➔ 6PE allows IPv6 to be deployed over existing MPLS Multiservice infrastructure with marginal operational impact/cost /risk

6PE Benefits

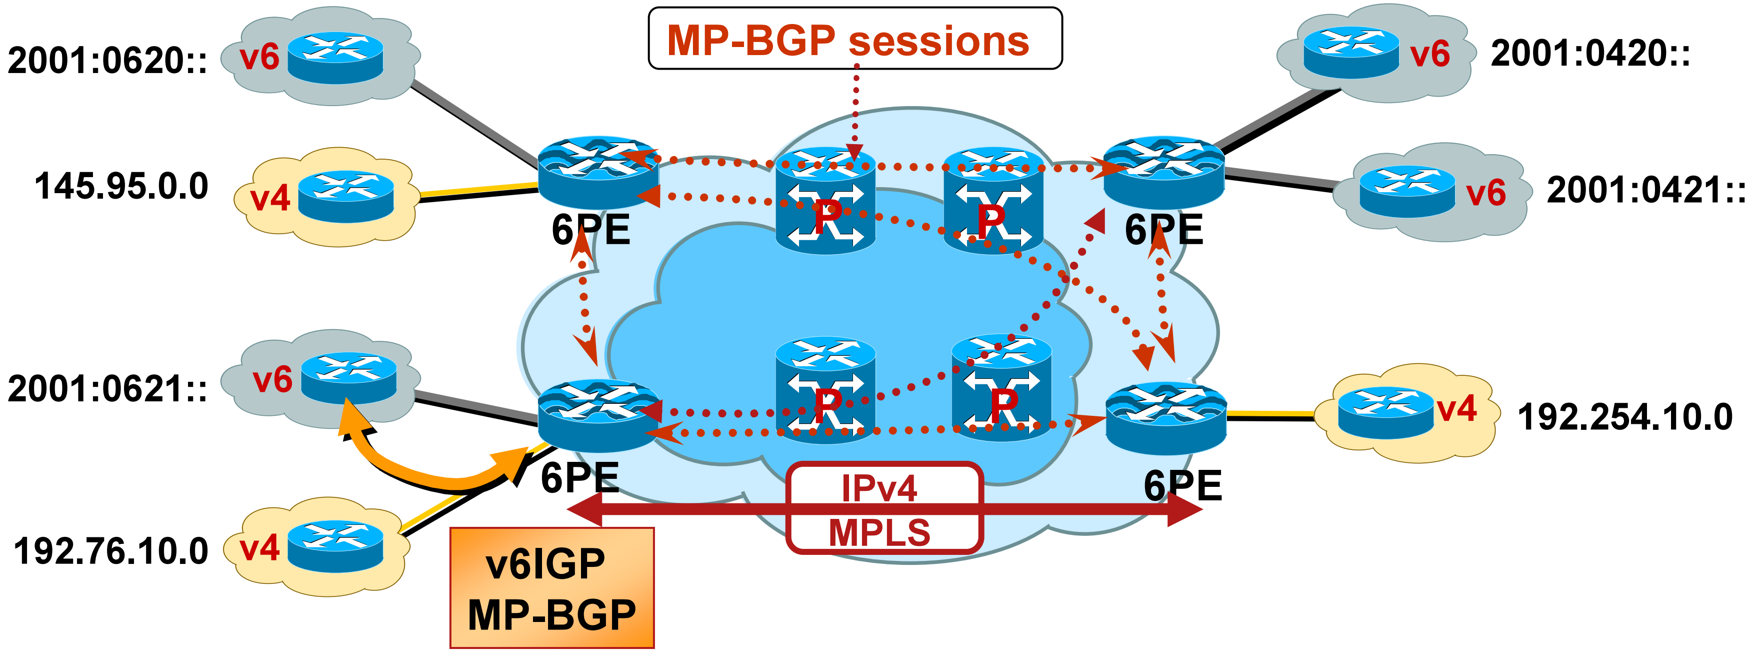

IPv6 CE only has a single Routing Peer (PE) regardless of how many remote IPv6 CEs it communicates with

No change on an IPv6 CE when remote CEs are added/removed (reachability automatically learnt)

No tunnel/"circuit" to be configured

➔6PE offers scalable and flexible solution (benefits are analogous to RFC2547bis layer 3 VPN solution for IPv4)

6PE Benefits

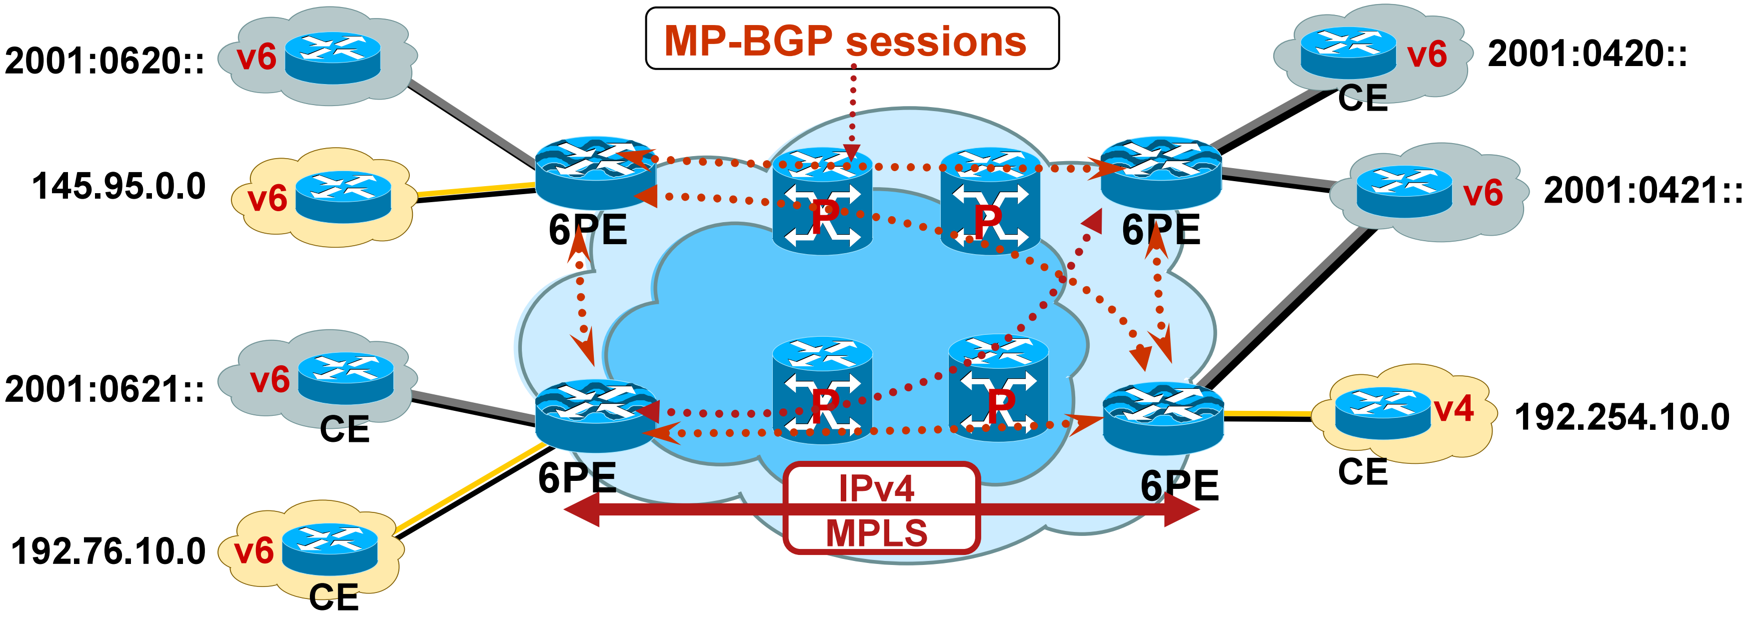

→6PE solution can be easily extended to support same VPN services for IPv6 as currently supported for IPv4 with RFC2547bis (isolation, Internet access, QoS…)

6PE Cons

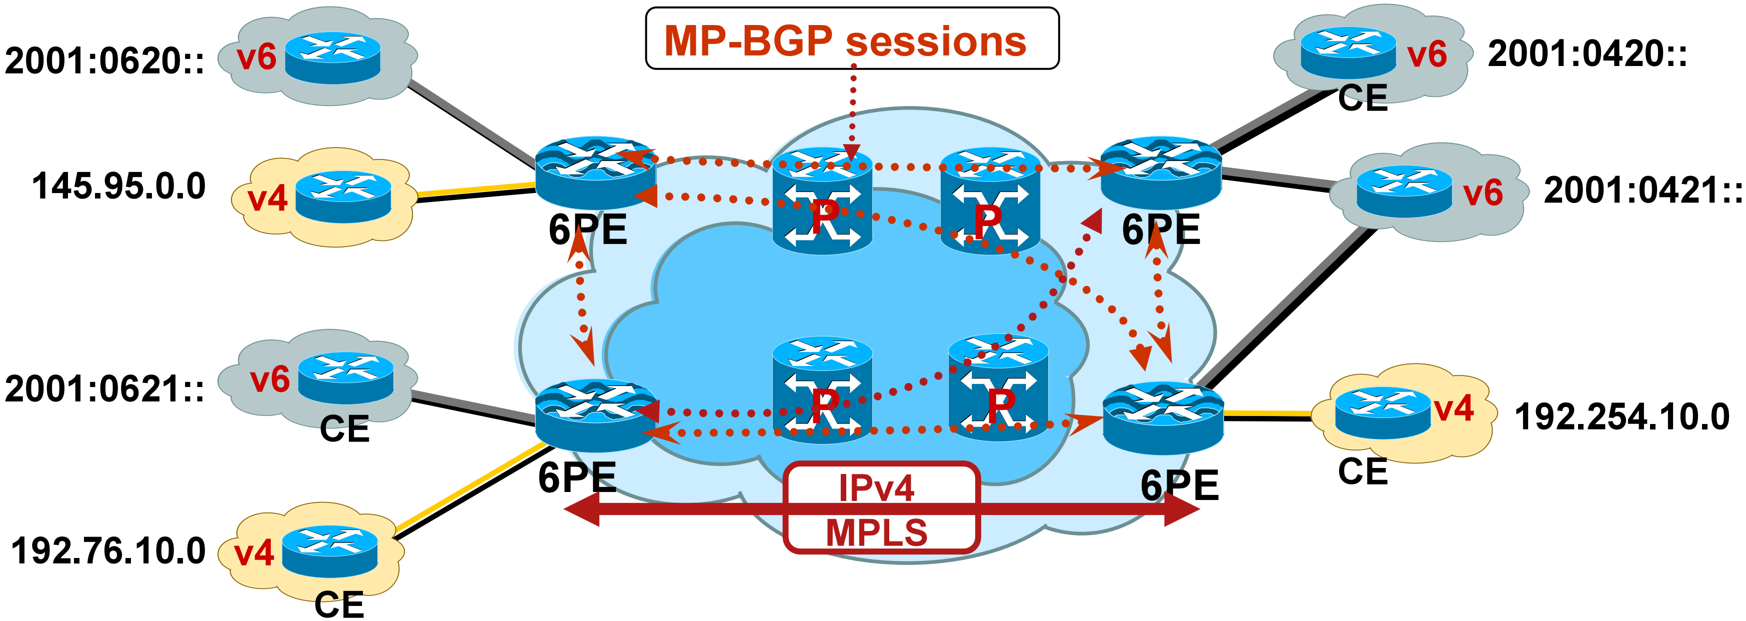

- Only makes sense where network already runs MPLS
- Requires knowledge of MPLS and BGP technologies
- Requires dual-stack and software upgrade on existing PE or deployment of dedicated 6PE routers

Why Cisco IOS 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 VPN’s concurrently on the same interfaces
    - Configuration and operations of IPv6 VPN’s exactly like IPv4 VPN’s

Cisco IOS 6VPE Architecture

- Similar to IPv4-VPN per RFC2547bis
- 6VPE definition: Each site is IPv6 capable and connect to SP backbone over an IPv6 interface via a 6VPE
- Sites may be both IPv4 & IPv6 capable
- 6VPE is a dual stack router
- 6VPE maintains IPv6 VRFs and each IPv6 VPN have their own address space
- Routes exchanged via MP-BGP : using NLRI (afi=2, safi=128 + label). NH encoding « a la 6PE » when applicable. Capability negotiation « as usual ».
- Forwarding over MPLS
- RD used to create disting routes to common IPv6 prefixes
- RT similar to IPv4

6VPE Deployment

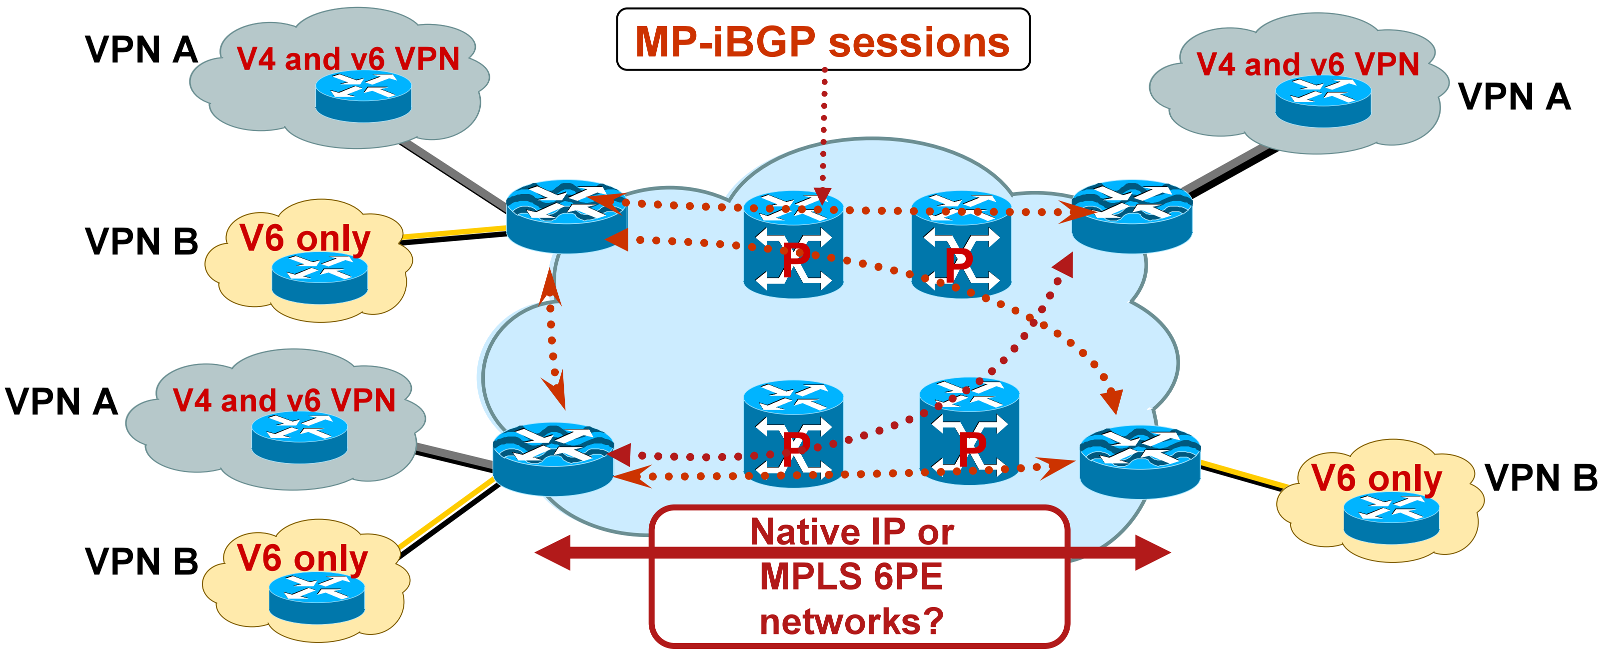

- **Customer Connectivity Deployed on a Shared Infrastructure with the Same Policies as a Private Network**
- **IPv6 VPN can co-exist with IPv4 VPN – same coverage**
- **6VPE is added only when and where the service is required**

Q and A

Recommended Reading

- Continue your Networkers at Cisco Live learning experience with further reading from Cisco Press
- Check the Recommended Reading flyer for suggested books

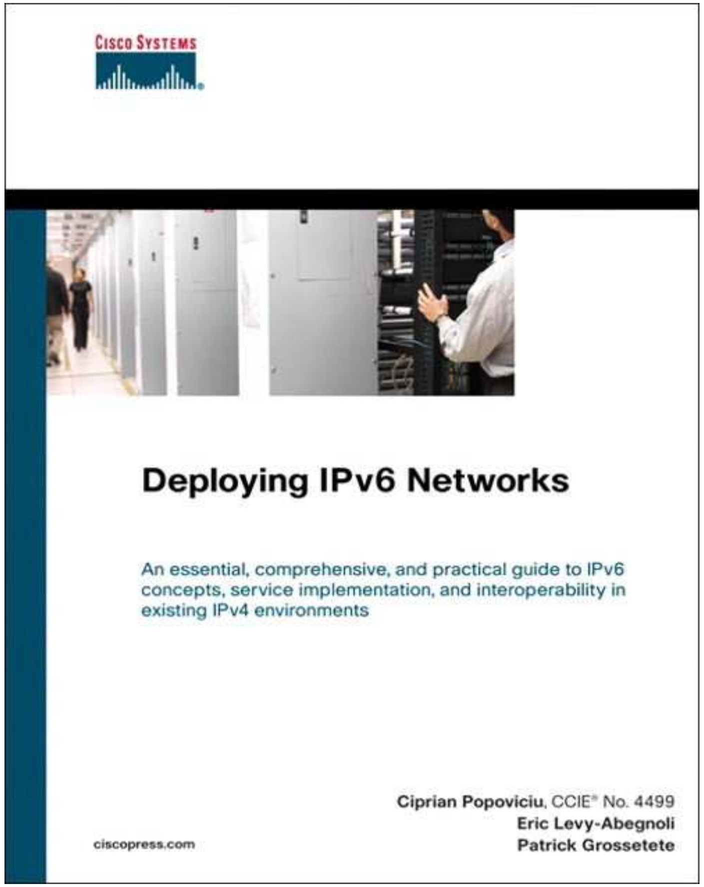

Available Onsite at the Cisco Company Store

Complete Your Online Session Evaluation

- Win fabulous prizes; give us your feedback
- Receive ten Passport Points for each session evaluation you complete
- Go to the Internet stations located throughout the Convention Center to complete your session evaluation
- Winners will be announced daily at the Internet stations