N.T.U.A Network Management Center

6PE:
IPv6 over MPLS

Presented by:
Andreas Polyrakis
Dimitrios Kalogeras

CISCO EXPO, January 2005

RoadMap

- Introduction
  - The Increasing Need for IPv6
  - MPLS & MPLS VPNs Basic
  - 6PE motivation
- The 6PE Approach
  - Applicability & Requirements
  - Label Distribution
  - Packet Switching
  - Router Interactions
  - Configuration Examples
- Advanced Topics: 6PE over CSC
  - The SEEREN case study
  - 6PE Troubleshooting
- Conclusions

What is 6PE?

- A method to provide IPv6 connectivity to your customers without upgrading the entire network to IPv6
- Requires MPLS
- Similar to MPLS VPNs

- Necessary Background
  - Basic IPv6
  - Basic MPLS / MPLS VPNs

The Boost of IPv6

- IPv6 is here for quite some time but…
  - NAT extended the life of IPv4
  - Not motivation to deploy IPv6
  - Vendors did not support IPv6
- However, in the past few years…
  - New types of applications (P2P) – NAT is no longer sufficient
  - Enormous increase of internet users – DSL
  - Not only PCs connect to the internet – 3G mobile devices
- IPv6 is now boosted
  - Vendors now support IPv6
  - ISPs deploy IPv6 services

IPv6 Provisioning in Existing Networks

- Upgrade the entire network to dual stack
  - native IPv6 provisioning
  - operational costs and risks (planning, upgrades, lack of expertise) L
- In pure IPv4 backbones
  - Connect IPv6 sites (clients, upstream, server farms) though a mesh of GRE tunnels
    - *Not scalable L*
    - *Big operational costs (OPEX) L*
- In IPv4/MPLS backbones
  - Connect IPv6 sites through a mesh of p2p L2 MPLS VPNs (ATOM or L2 Interworking)
    - *Similar to GRE tunnels, not scalable L*
    - *ATOM: Same media at endpoints L*
    - *L2 Interworking: Many restrictions apply, too L*
  - Upgrade some edge routers to dual stack and use 6PE

MPLS Simplified

- Packet Forwarding based on a Label imposed at the ingress point
- Terminology
  - CE: Customer Edge Router (C1,C2)
  - PE: Provider Edge Router (Ra, Rd)
  - P: Provider Core Router (Rb, Rc)
- Label Creation & Distribution
  - Routers in the MPLS domain create labels for all ipv4 IGP routes
  - Label information is exchanged (with a Label Distribution Protocol – LDP or BGP)
  - Labels + Routing Table è Label Forwarding Information Base (LFIB)

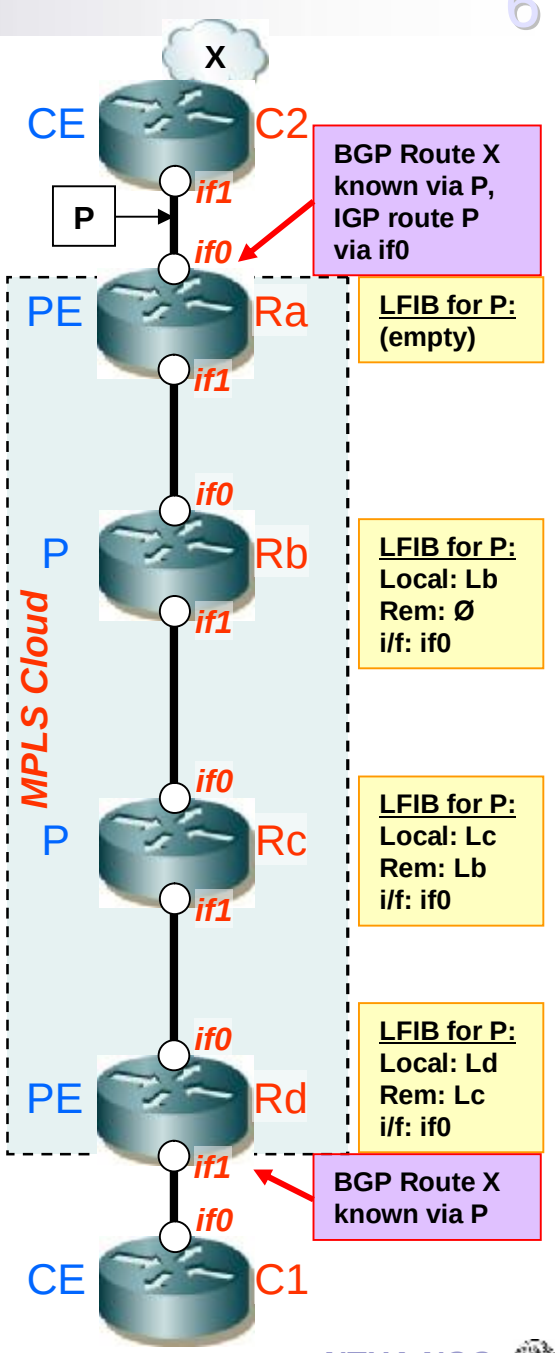

MPLS Simplified #2

- **Packet Forwarding**
  - **One label is imposed on ingress based on destination IP and LFIB**
    - *IGP routes exist in LFIB*
    - *BGP routes don't exist in LFIB; Choose the label for the BGP-next-hop of the route*
  - **Labels are swapped while packet is forwarded**
  - **Penultimate Hop (in most cases) performs Penultimate Hop Popping (PHP); in this case the last hop receives an IP packet**
  - **Last hop forwards the packet to the appropriate egress interface**

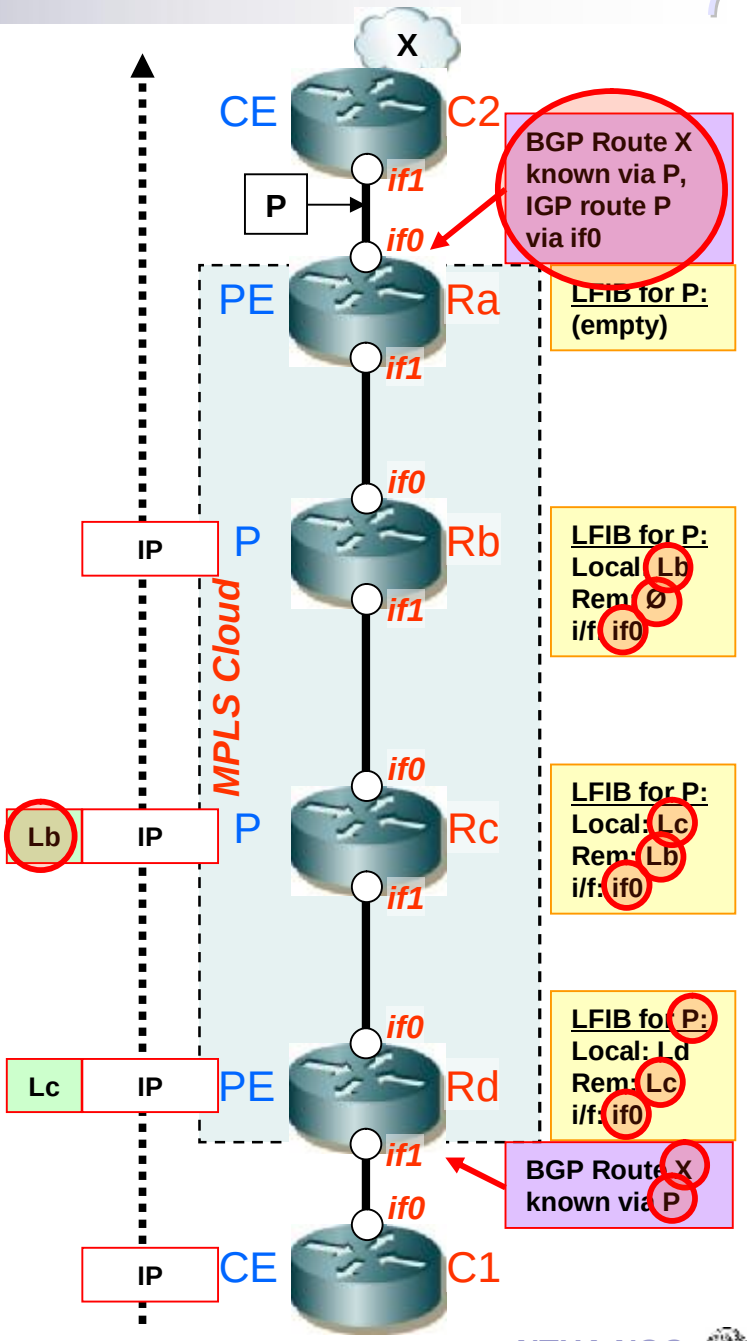

MPLS VPNs Simplified

- VPN definition
  - if0@Ra and if1@Rd belong to the same VPN
- MP-BGP exchanges VPN information
  - VPN Routes + VPN Label for each PE
- Ingress VPN interface: On packet arrival lookup the VPN table:
  - Find the Label for the VPN route
  - Find the egress router
- Push two labels:
  - Inner: The label of the VPN route
  - Outer: The label towards the egress router
- Forwarding: Swap exterior label
- Penultimate Hop Popping (PHP)
- Last Hop: Packet received with interior label only
- Inner label à VPN & egress interface
- The last label is popped, the IP packet is forwarded to the CE router

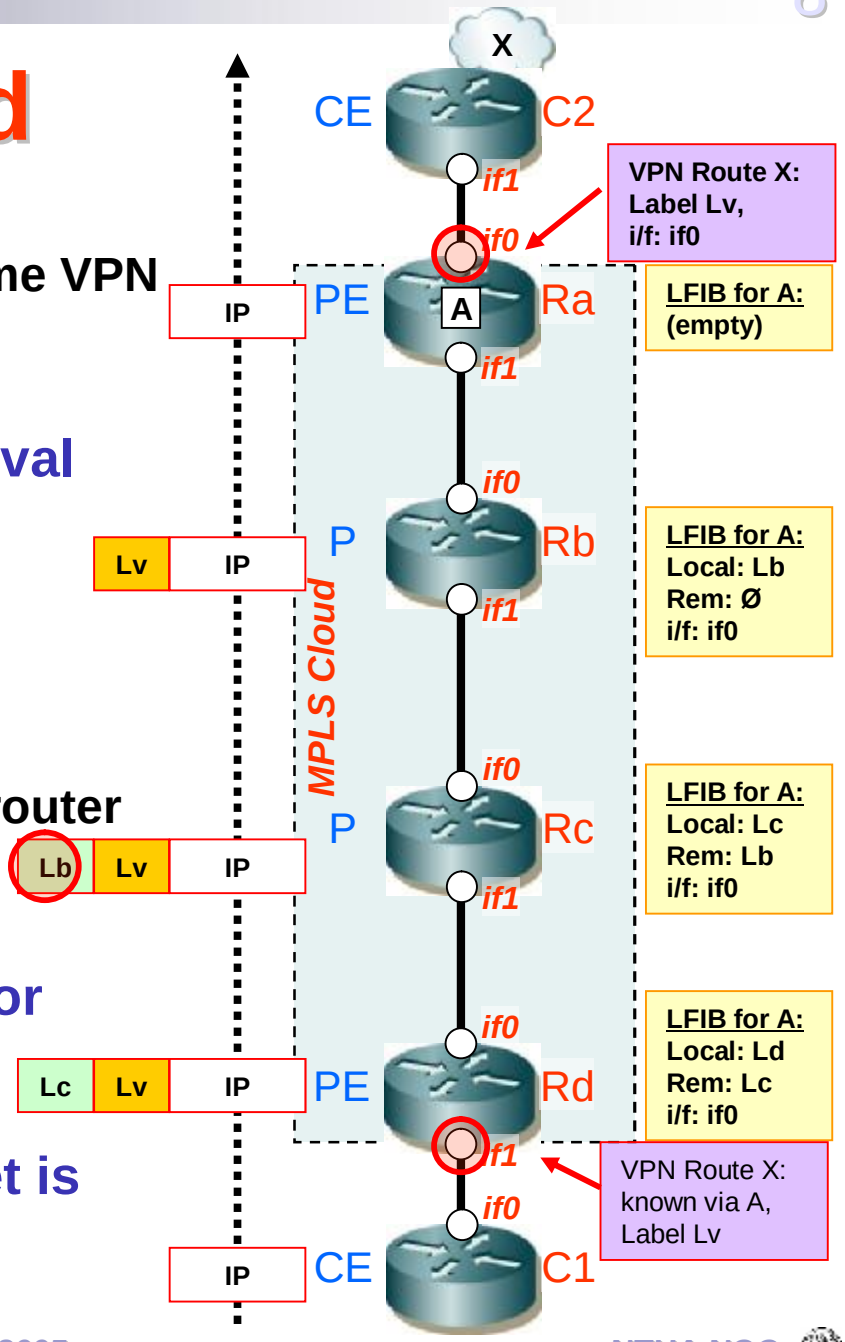

Motivation for 6PE

If Ra,Rd,C1,C2 were dual stack

- Could the same be done with an IPv6 route?
  - i.e: Find the egress router
  - Use two labels:
    - *The outer to reach the egress router*
    - *The inner to declare that this packet must be looked up in the IPv6 routing table*
- Yes!

  Issues:
  - How Rd will find the eggress router Ra?
    - *Not with IPv4 IGPs (cannot carry IPv6)*
    - *Not with IPv6 IGPs (are not supported by the core)*

      è *MP-BPG*
  - How will Rd determine the inner label?

    è *MP-BGP again*

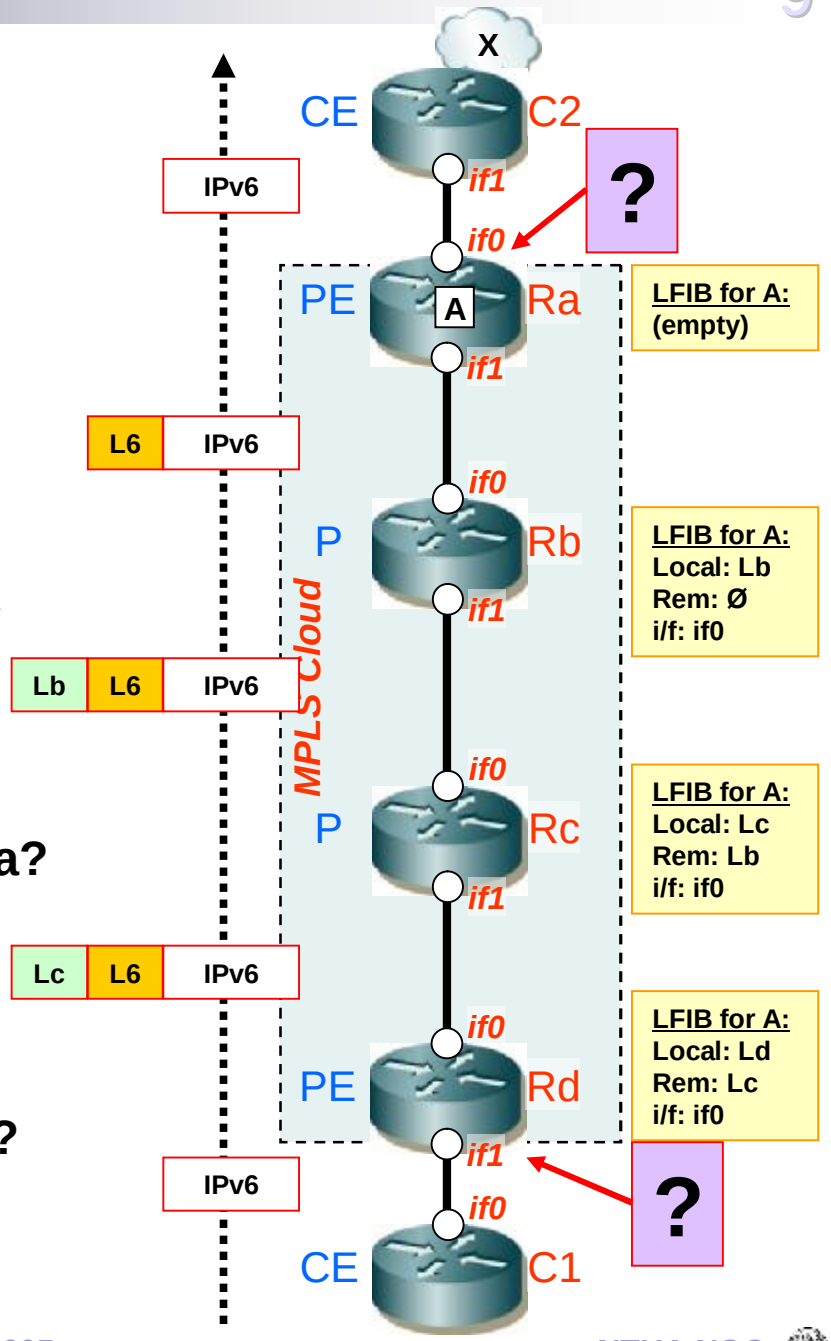

The 6PE Approach

- **What is 6PE: A method to provide IPv6 connectivity to your customers without upgrading the entire network to IPv6**

- **6PE requirements:**
  - **The ISP has to upgrade the Provider Edge (PE) routers to support IPv6 and MP-BGP**
  - **Core (P) routers do not need any change in terms of configuration or software.**

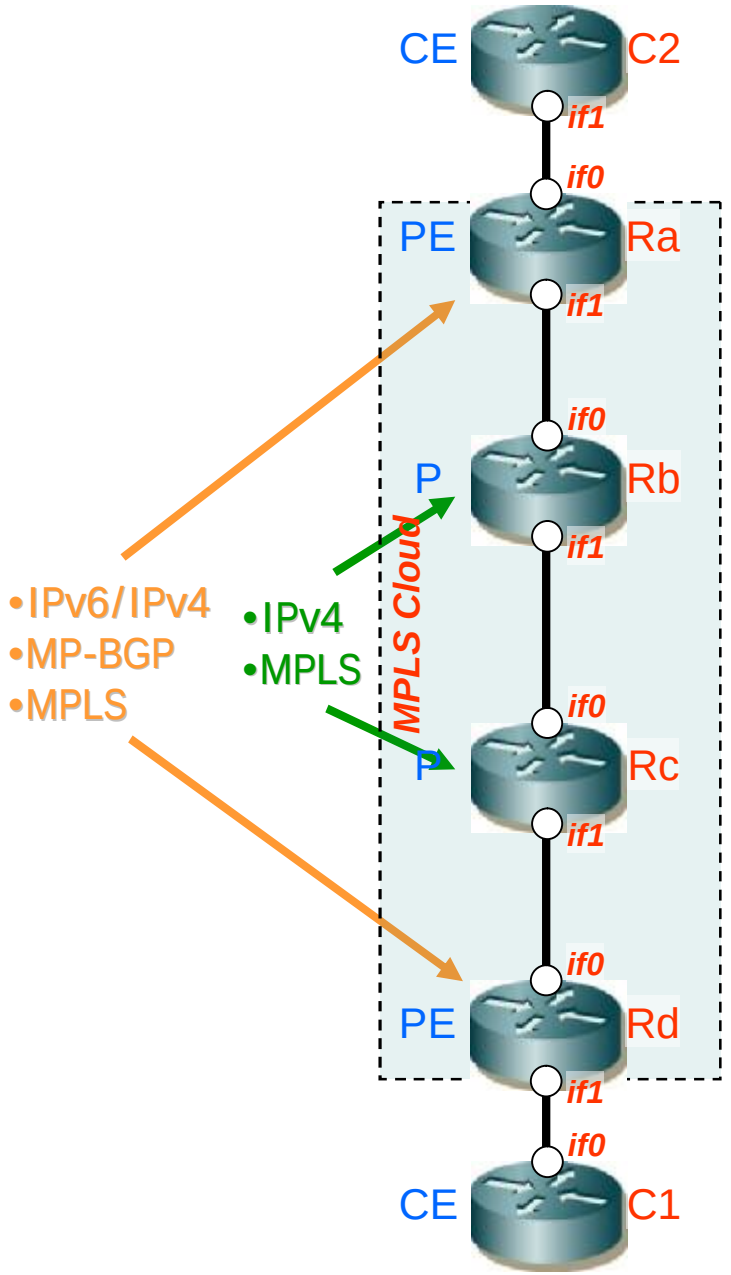

The 6PE Applicability

- 6PE is typically deployed by ISPs that have MPLS core network and (possible) supports MPLS VPN (or other) services.

- IPv6 services are requested by a medium number of customers.
  - If the IPv6 customers are limited, a L2 tunnelled solution may be preferable.
  - If the IPv6 customers are require most of the access routers to become 6PE, ISP may consider to upgrade to whole network.

- The ISP wants to avoid either to fully upgrade the core network or to deploy IPv6-over-IPv4 tunnels.

Configuration Overview

- Step 1: Upgrade PEs to Dual Stack
- Step2: Populate the PE's IPv6 Routing Table
  - Establish IPv6 route exchange with the CE
- Step3: Configure MP-BGP
  - Exchange IPv6 Routes
  - Exchange IPv6 Labels

Step 1: Upgrade PEs to Dual Stack

- Upgrade to an IOS that supports 6PE:
  - IPv6
  - BGP-MP
- Enable IPv6
- Optionally, assign IPv6 addresses to some interfaces

```
PE1#
ipv6 unicast-routing
ipv6 cef
```

Step 2: Populate the v6 Routing Table

- PE Router is Dual Stack
- IPv6 routing table exists
- IPv6 routing table can be populated with:
  - Routes learned through eBGP between CE and PE (typically)
  - Routes of connected interfaces
  - Routes learned through an IPv6 capable IGP protocol (eg ISIS)
  - Static or default routes

F These routes (IPv6 Table) need to be exchanged with the other PEs

Step 2: Configuration Commands

```
CE1#
interface Ethernet0/0
 no ip address
 ipv6 address 2001:648:2000::1/64
 ipv6 rip RTP enable
!
interface serial0/0
 description *** LINK TO PE ***
 no ip address
 ipv6 enable
 ipv6 rip RTR enable
!
ipv6 router rip RTR
```

```
PE1#
interface Loopback0
 ip address 194.22.15.2
    255.255.255.255
 ip router isis
 ipv6 address 2001:648:2fff:05::/128
 ipv6 rip RIPv6 enable
!
interface serial1/0
 description *** Link to CE1 ***
 no ip address
 ipv6 enable
 ipv6 rip RIPv6 enable
 ipv6 rip RIPv6 default-information
    originate
!
ipv6 router rip RIPv6
```

Step 3: Establish MP-BGP

- What is MultiProtocol BGP?
  - BGP extension that allows the exchange of multiple Routing Tables
    - *IPv4 unicast*
    - *IPv4 multicast*
    - *IPv6*
    - *…*
  - "Address-families"
- Labels can also be distributed via MP-BGP
  - Carrier-Supporting-Carrier
  - InterAS VPNs
  - 6PE
  - …

Step 3: Label Distribution in 6PE

- 6PE routers establish MP-BGP sessions
  - running over IPv4
- IPv6 prefixes are exchanged through MP-BGP
- BGP Next Hop field is the IPv4-mapped IPv6 address of the 6PE router
  - Mapping: X.Y.Z.W è ::FFFF:X.Y.Z.W
- An "Aggregated IPv6 Label" is sent for IPv6 routes
  - This will be used as the "inner" label

Note: Labels are made and exchanged for the ipv4 loopbacks of the PE routers, too (typical MPLS + LDP)

Step 3: Configuration Commands

- include “sent-label” in the ipv6 address family
- Must be done in both directions PEs
- In a Mesh of PEs, it must be done with all PEs

```
PE1#

router bgp 5408
 neighbor 194.22.15.3 remote-as 5408
 neighbor 194.22.15.3 update-source
    loopback0
!
address-family ipv6
neighbor 194.22.15.3 activate
neighbor 194.22.15.3 send-label
redistribute rip RIPv6
redistribute connected
```

```
PE2#

router bgp 5408
 neighbor 194.22.15.2 remote-as 5408
 neighbor 194.22.15.2 update-source
    loopback0
!
address-family ipv6
neighbor 194.22.15.2 activate
neighbor 194.22.15.2 send-label
redistribute rip RIPv6
redistribute connected
```

Packet Forwarding in 6PE

- MP-BGP exchanges IPv6 Routes & Labels
- On packet arrival at the ingress 6PE:
  - The IPv4-mapped IPv6 BGP next-hop for the destination address is found
  - These two labels are pushed:

    Inner label: The label for the BGP next-hop

    Outer label: The label for the IPv4 address correspinding to the IPv4-mapped IPv6 address
- Forwarding: Swap exterior label
- Penultimate Hop Popping (PHP)
- Last Hop: Packet received with interior label only
- Inner label à VPN & egress interface
- The last label is popped, the IP packet is forwarded to the CE router

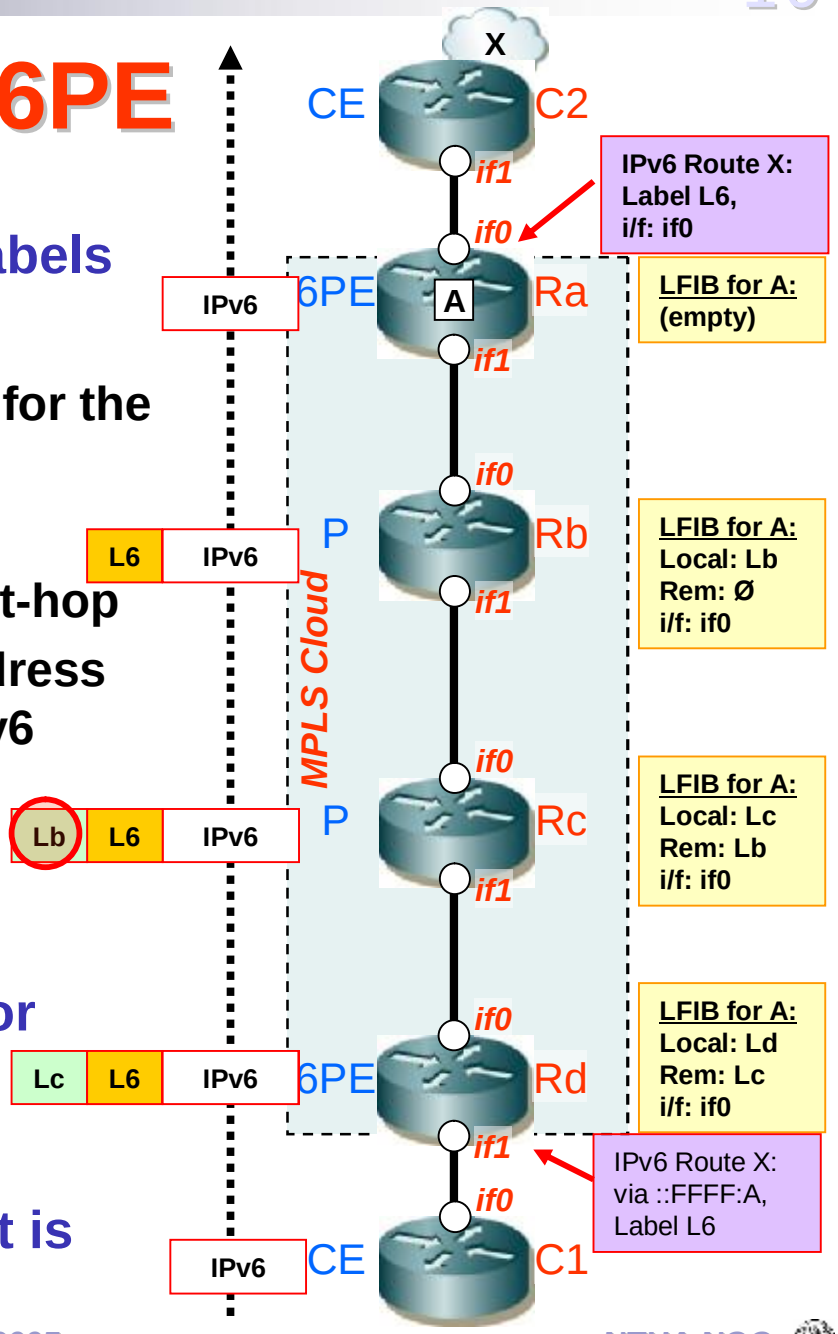

Routing Interactions in 6PE

- **Client Network**
  - **Forwarding:** IPv6
  - **Routing:** IGPv6 (IS-IS, OSPF), static
- **PE-CE**
  - **Forwarding:** IPv6
  - **Routing:** eBGP, IGPv6 (IS-IS, OSPF), static
- **PE-PE**
  - **Forwarding:** MPLS
  - **Routing:** MP-BGP, IGP
  - **Label Distribution:** MP-BGP (V6), LDP (V4)

Advanced Topics: 6PE over CSC

Carrier-Supporting-Carrier (CSC)

- For ISPs that are VPN customers of MPLS Backbones
- A VPN with very small virtual routing table (VRF)
- CSC-CEs exchange limited labels with CSC-PE
- MPLS between CSC-CE and CSC-PE

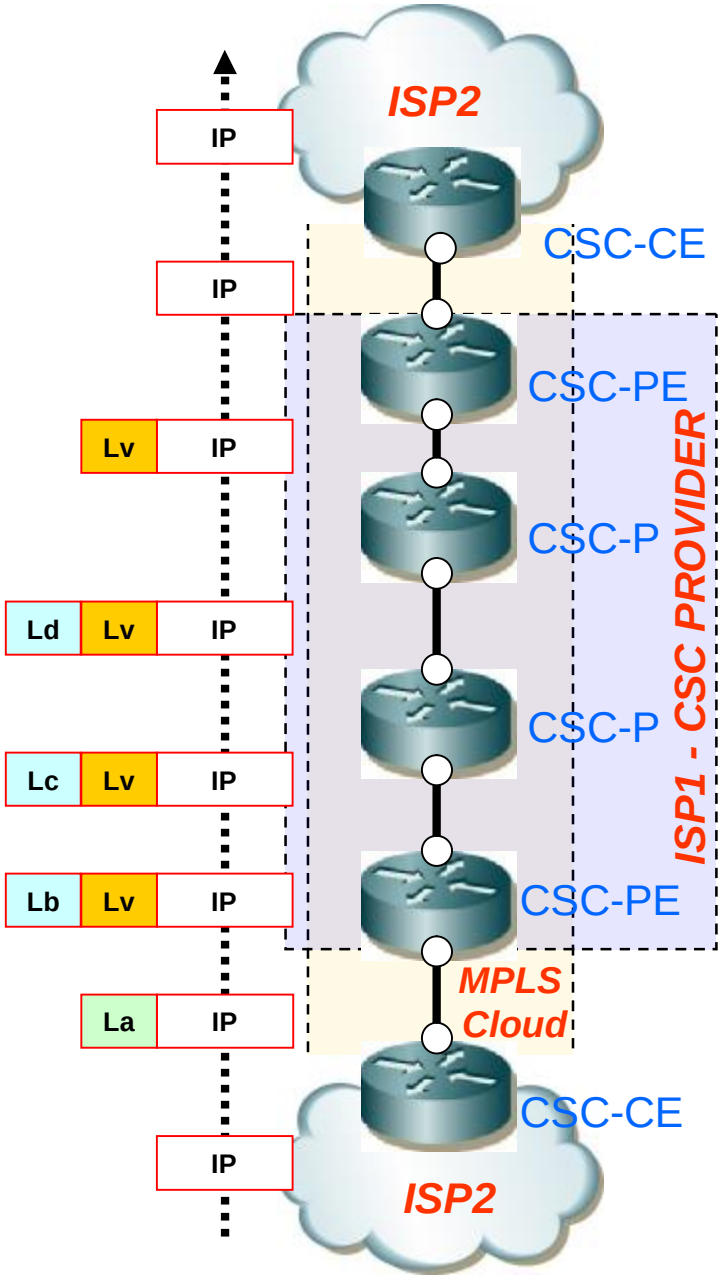

Case Study: The SEEREN network

- **SEEREN: South Eastern European Research & Educational Network**
  - **Upstream:** GRNET (Greek Research Network)
  - **Backup Upstream:** Roedunet (via GRNET)
  - **Peers:** Serbia & Montoragevo, Bulgaria, Albania, FYROM, Bosnia & Herzecovinova
  - **Connectivity:** Oteglobe, via a **CSC MPLS VPN**

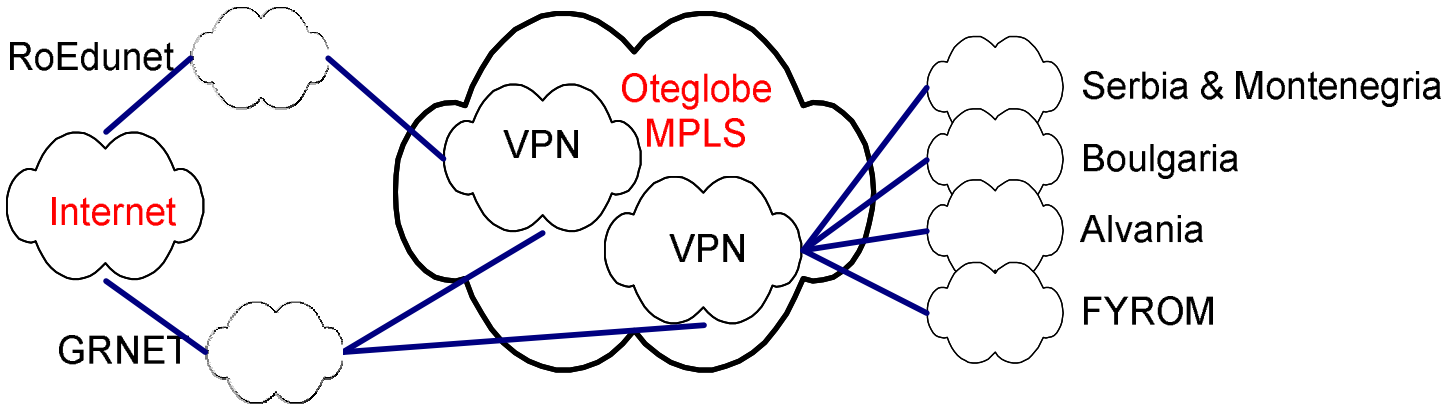

6PE in SEEREN

Not a typical 6PE network:

- "Typical" PE routers belong to Oteglobe
- 6PE is installed on the CEs instead of the PEs!!!
- 6PE peers belong to different ASes !!!
- Feasible, because CE-PE connection uses MPLS
- **Three Label Stack !!!**

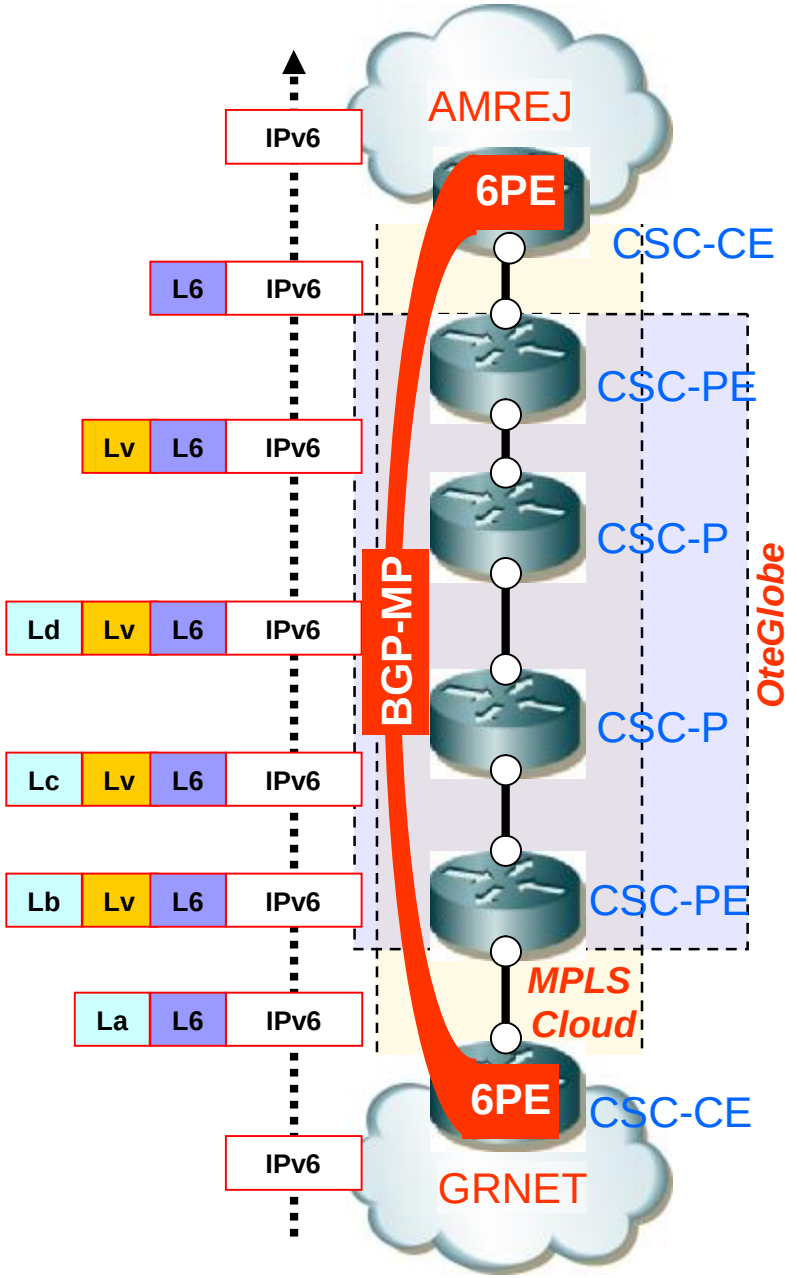

SEEREN 6PE Configuration #1

```
hostname athens-2
!
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 ip address 194.177.210.40
    255.255.255.255
!
interface ATM1/0/0.1 point-to-point
 description SEEREN via OteGlobe VPN
 ip address 62.75.33.246 255.255.255.252
!
interface GigabitEthernet3/0/0
 description Athens2 - Athens3
 ipv6 address 2001:648:2FFF:106::2/126
 ipv6 router isis
!
router isis
...
...
```

Enable IPv6

IPv4 Loopback: Necessary for BGP

IPv4 Link with OteGlobe

IPv6 GRNET Network

IPv6 Routing Protocol. Populates the IPv6 Routing Table

SEEREN 6PE Configuration #2

MP-BGP

Send Labels for IPv4 routes. (BGP is used for Label Distribution with OteGlobe)

Send Labels for IPv6 routes

```
router bgp 5408
 neighbor 62.75.33.245 remote-as 12713
 neighbor 62.75.33.245 description OTEGLOBE PE
 neighbor 147.91.0.112 remote-as 13092
 neighbor 147.91.0.112 description AMREJ-YUGOSLAVIA
 neighbor 147.91.0.112 ebgp-multihop 5
 neighbor 147.91.0.112 update-source Loopback0
 !
 address-family ipv4
 neighbor 62.75.33.245 activate
 neighbor 62.75.33.245 send-community
 neighbor 62.75.33.245 remove-private-as
 neighbor 62.75.33.245 soft-reconfiguration inbound
 neighbor 62.75.33.245 send-label
 neighbor 147.91.0.112 activate
 neighbor 147.91.0.112 send-community
 neighbor 147.91.0.112 remove-private-as
 neighbor 147.91.0.112 soft-reconfiguration inbound
 !
 address-family ipv6
 neighbor 147.91.0.112 activate
 neighbor 147.91.0.112 send-community
 neighbor 147.91.0.112 remove-private-as
 neighbor 147.91.0.112 soft-reconfiguration inbound
 neighbor 147.91.0.112 send-label
 exit-address-family
```

6PE verification & troubleshooting

```
athens-2# sh ip bgp  neighbor 194.141.252.13
BGP neighbor is 194.141.252.13,  remote AS 6802, external link
 Description: UNICOM-BOULGARIA
  BGP version 4, remote router ID 194.141.252.13
  BGP state = Established, up for 2d08h
  Last read 00:00:25, hold time is 90, keepalive interval is 30 seconds
  Configured hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
    ipv6 MPLS Label capability: advertised and received
  ipv6 MPLS Label capability: advertised and receive
...
```

```
 athens-2# sh bgp ipv6 unicast
BGP table version is 26881, local router ID is 194.177.210.40
Status codes: s suppressed, d damped, h history, * valid, > best, i -
     internal,  S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2001:648:F080::/42
                    ::FFFF:194.141.252.13
                                             0             0 6802 i
```

6PE verification & troubleshooting

```
athens-2# sh bgp ipv6 unicast labels
  Network          Next Hop      In label/Out label
…
  2001:648:F080::/42
                   ::FFFF:194.141.252.13
                                 216/47
…
```

```
athens-2# sh ipv6 route 2001:648:F080::/47
IPv6 Routing Table - 487 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS
    summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2
    - OSPF ext 2
B   2001:648:F080::/42 [20/0]
     via ::FFFF:194.141.252.13, IPv6-mpls
```

Conclusions

- **Purpose: Provide native IPv6 services to customers without changing the IPv4 MPLS core network**
  - **minimal operational cost and risk**
  - **scalability**
  - **fits very well into the general MPLS philosophy è Flexibility**
- **Requirements:**
  - **P routers : no modification!**
  - **PE routers : (a) must be dual stack and (b) must support MP-BGP**
  - **CE routers : just need to be dual stack**
    - ***But also recommended to run an IPv6-capable routing protocol with PE***
- **But do not:**
  - **Use 6PE when most customers require IPv6. Consider upgrading the entire network to dual stack**
  - **Upgrade your entire network to MPLS only just to support 6PE. Use tunnels instead.**

Questions?

Thank You!

Andreas Polyrakis
A.Polyrakis@noc.ntua.gr

Dimitrios Kalogeras
D.Kalogeras@noc.ntua.gr