



6PE over CsC – The SEEREN case

Thanassis Liakopoulos
Greek Research & Technology Network (GRNET)

17th TF-NGN Meeting
Zurich, 14th April 2005

Presentation Outline

■ Introduction

- What is 6PE? - Basic characteristics
- Applicability
- Alternative solutions

■ 6PE Implementation Details

- MPLS VPN - 6PE innovation
- Configuration Overview - Label Distribution
- Packet Forwarding

■ The SEEREN Case

- The network
- What is CsC? – 6PE over CsS
- Configuration - Troubleshooting

■ Conclusions



Presentation Outline

■ Introduction

- What is 6PE? - Basic characteristics
- Applicability
- Alternative solutions

■ 6PE Implementation Details

- MPLS VPN - 6PE innovation
- Configuration Overview - Label Distribution
- Packet Forwarding

■ The SEEREN Case

- The network
- What is CsC? – 6PE over CsS
- Configuration - Troubleshooting

■ Conclusions

What is 6PE?

- A method that allows an MPLS-based ISP to offer IPv6 interconnection services to his/her customers without upgrading the entire network to support IPv6 protocols.

Basic characteristics

- **6PE is similar to MPLS VPNs model in terms of technical implementation and complexity.**
 - Label encapsulation is used for transporting IPv6 packets.
- **IPv6 functionality is enabled only at the edge routers.**
 - A 6PE router, in simple terms, is an edge router with IPv6 functionality.
 - Core routers are IPv6 unaware, a.k.a. support IPv4 only protocols.

Applicability

- 6PE is deployed by ISPs that have MPLS core networks and (possible) offers MPLS VPN (or other) services.
- IPv6 services are requested by a small number of customers.
 - If the IPv6 customers are limited, an L2 tunnelled solution may be preferable.
 - If the IPv6 customers require most of the (access) edge routers to be upgraded, ISP may consider to upgrade to whole network.
- The ISP wants to avoid either to fully upgrade the core network or to deploy L2/L3 tunnels.
 - Why?

Alternative Solutions

■ Dual stack network

- IPv6 and IPv4 traffic is treated identically by the core routers. 😊
- Requires IPv4/IPv6 routing and label distribution protocols. 😞

■ L2 MPLS VPNs

- Use MPLS to create point-to-point tunnels that terminate at the edge routers.
- IPv6 is needed only at edge routers. No changes are needed to core routers. 😊
- Scalability problems arise in heavily (L2 tunnel) mesh topologies. 😞
- Edge routers has to support the appropriate functionality, e.g. ATOM or CCC/TCC. 😞

Presentation Outline

■ Introduction

- What is 6PE? - Basic characteristics
- Applicability
- Alternative solutions

■ 6PE Implementation Details

- MPLS VPN - 6PE innovation
- Configuration Overview - Label Distribution
- Packet Forwarding

■ The SEEREN Case

- What is CsC? – 6PE over CsS
- Configuration - Troubleshooting

■ Conclusions

MPLS VPNs Simplified

Terminology

- **CE:** Customer Edge Router (C1,C2)
- **PE:** Provider Edge Router (Ra, Rd)
- **P:** Provider Core Router (Rb, Rc)

Packet Forwarding based on two labels imposed at the ingress point of the MPLS network

Label Distribution for IGP routes

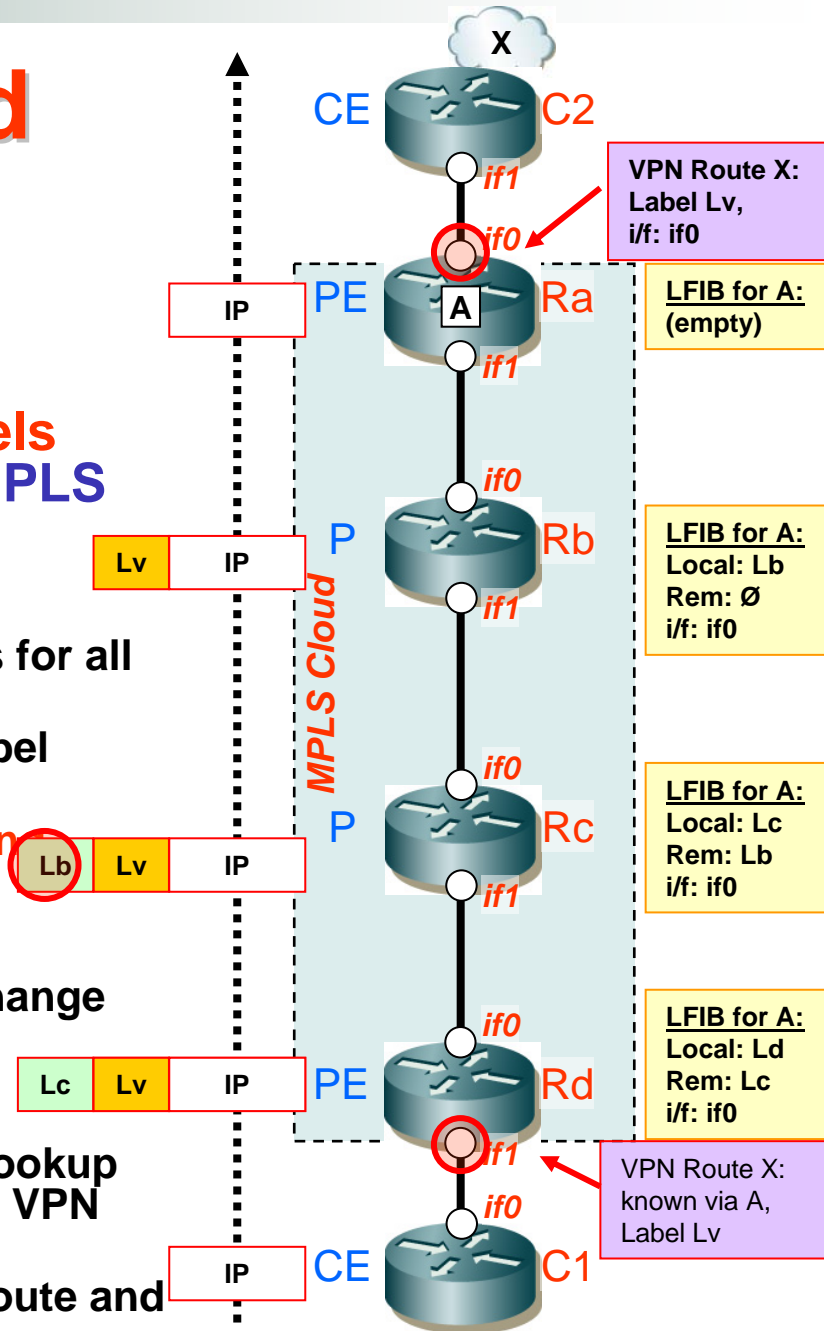
- Routers in the MPLS domain create labels for all the known **IPv4 IGP** routes.
- Label information is exchanged with a Label Distribution Protocol, e.g. LDP or BGP.
- Labels + Routing Table → **Label Forwarding Information Base (LFIB)**

Label Distribution for VPNs

- MP-BGP is used for VPN information exchange
- VPN Routes + VPN Label for each PE

Packet Forwarding

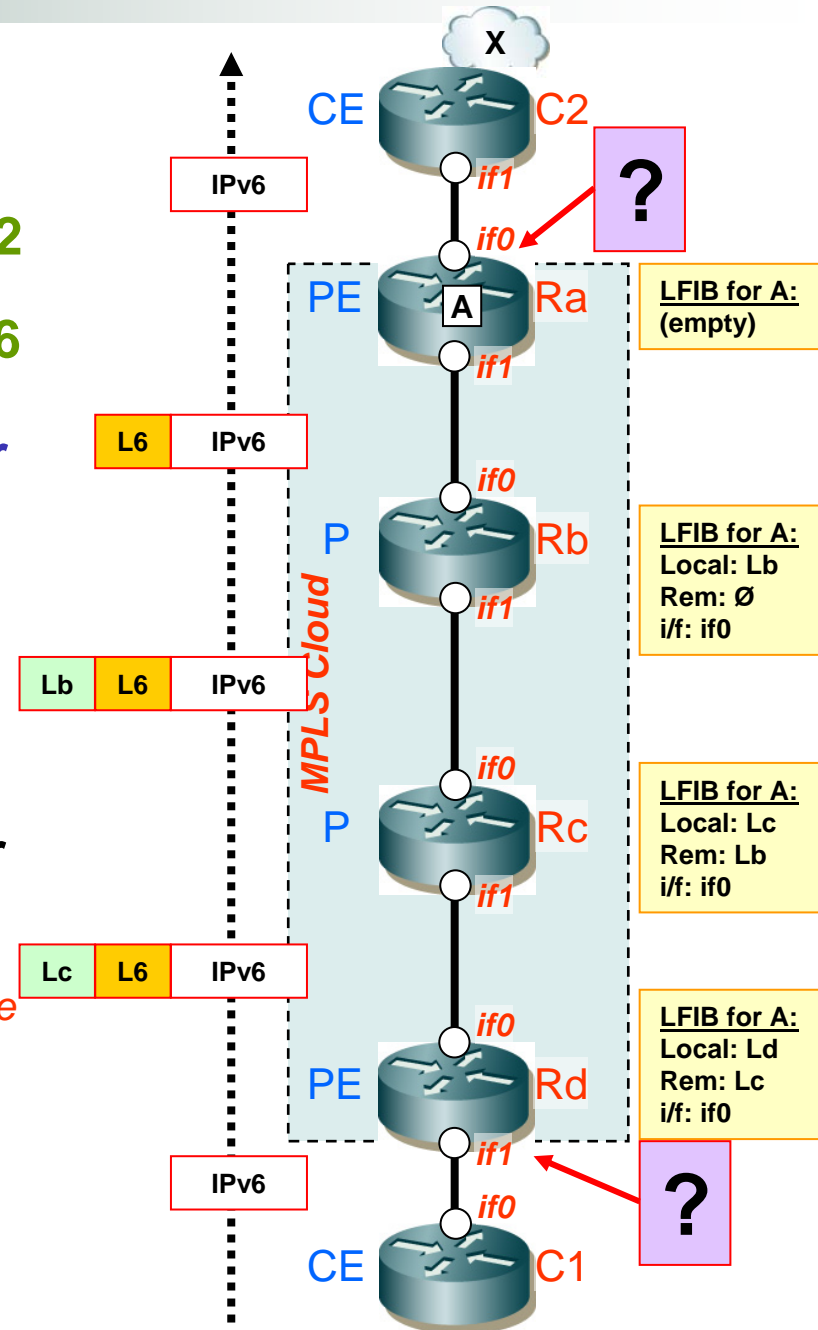
- **Ingress VPN interface:** On packet arrival lookup the VPN table and find (a) the label for the VPN route and (b) the egress router
- **Push two labels; Inner** label for the VPN route and **Outer** label towards the egress router



6PE «innovation»

Provided that routers Ra, Rd, C1, and C2 are dual stack could the MPLS VPN technique be used for forwarding IPv6 traffic from C1 to C2?

- Why not? One label could be used for frame forwarding through the MPLS (IPv4) network while the other label could be used for identifying the IPv6 route.
- There are some problems, though!
 - How Rd will find the egress router Ra?
 - Not with IPv4 IGPs (cannot carry IPv6)
 - Not with IPv6 IGPs (are not supported by the core)
 - Solution → MP-BGP
 - How will Rd determine the inner label?
 - Solution → MP-BGP again



Configuration Overview

- **Step 1: Upgrade PEs to become Dual Stack**

- **Step2: Populate the PE's IPv6 Routing Table**
 - Establish IPv6 route exchange with the CE


- **Step3: Establish 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: Built 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 (e.g. ISIS)
 - Static or default routes
-  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

■ Exchange IPv6 Routes

- 6PE routers establish MP-BGP sessions running over IPv4
- BGP Next Hop field is the IPv4-mapped IPv6 address of the 6PE router

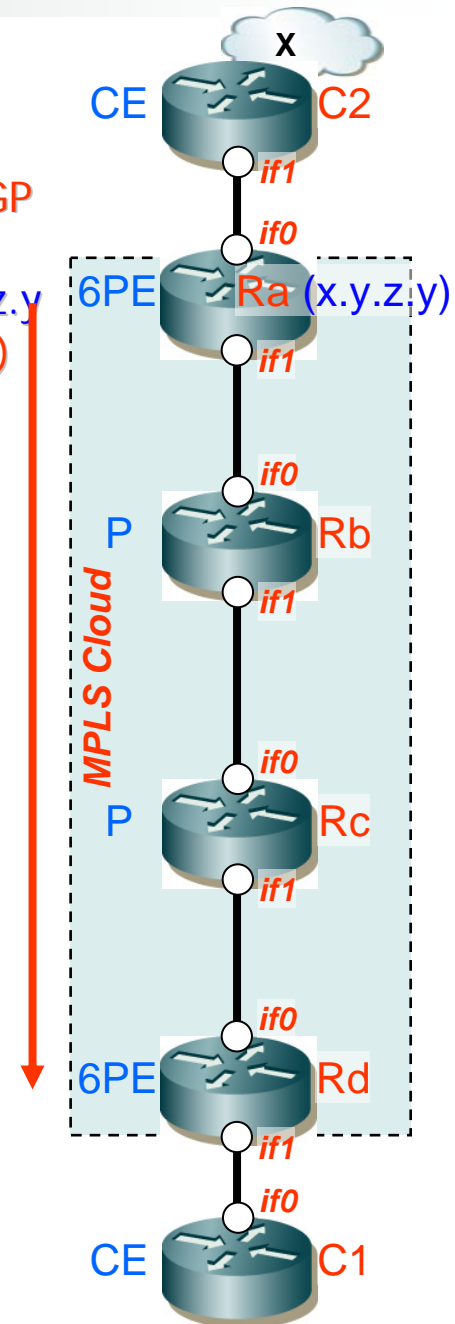
■ *Mapping: X.Y.Z.W → ::FFFF:X.Y.Z.W*

- IPv6 prefixes are exchanged

■ Exchange Labels

- An “Aggregated IPv6 Label” is sent for IPv6 routes. This will be used as the “inner” label

Advertise X over MP-iBGP
(X is reachable via BGP
Next Hop = ::FFFF:x.y.z.y
+ bind a BGP label to X)



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#

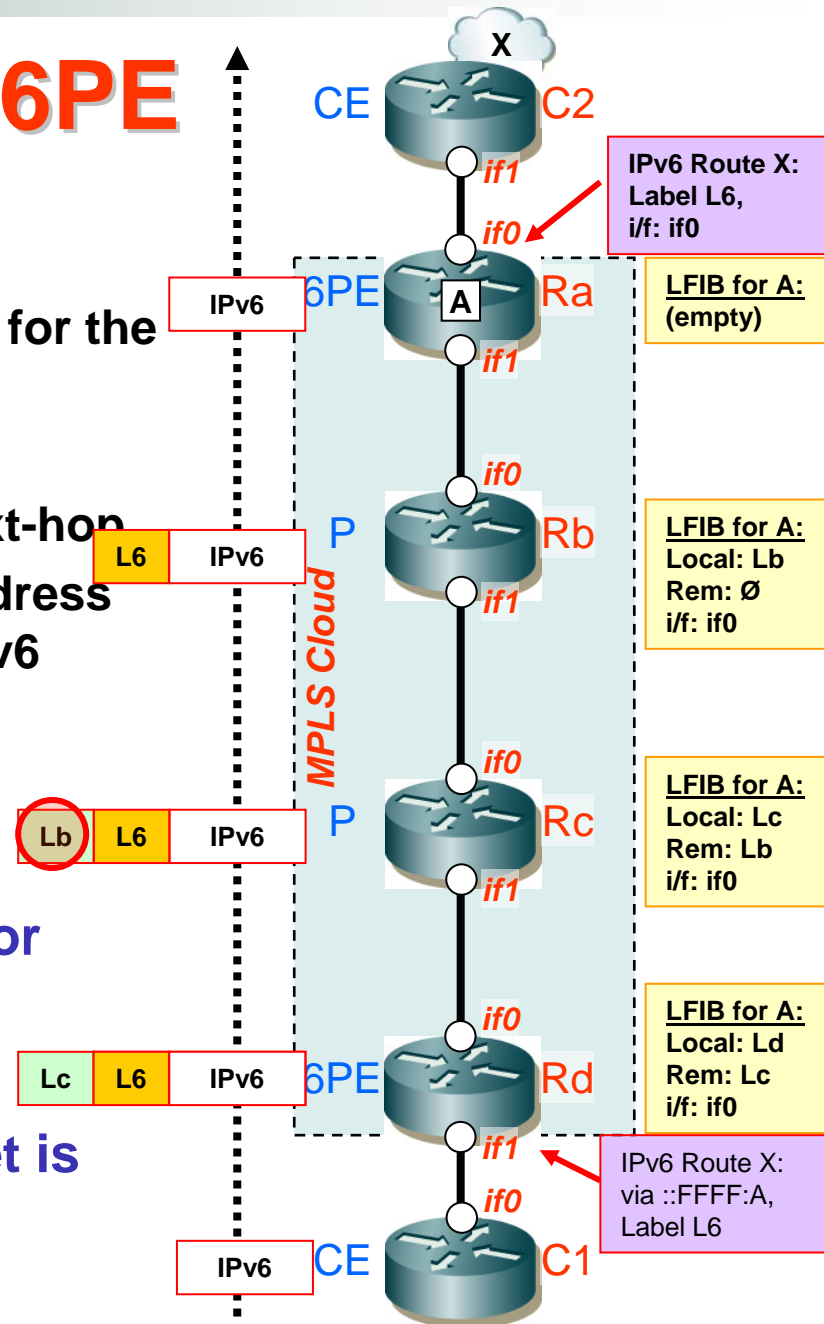
```
router bgp 5408
 neighbor 194.22.15.3 remote-as 5408
 neighbor 194.22.15.3 update-source
 loopback0
!
address-family ipv6
 neighbor 194.22.15.3 activate
 neighbor 194.22.15.3 send-label
 redistribute rip RIPv6
 redistribute connected
```

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

- On packet arrival at the ingress 6PE:
 - The **IPv4-mapped IPv6** BGP next-hop for the destination address is found
 - Two labels are pushed:
 - Inner label:** The label for the BGP next-hop
 - Outer label:** The label for the IPv4 address corresponding to the IPv4-mapped IPv6 address
- **Forwarding:** Swap exterior label
- **Penultimate Hop Popping (PHP)**
- **Last Hop:** Packet received with interior label only
- Inner label → VPN & egress interface
- The last label is popped, the IP packet is forwarded to the CE router



Presentation Outline

■ Introduction

- What is 6PE? - Basic characteristics
- Applicability
- Alternative solutions

■ 6PE Implementation Details

- MPLS VPN - 6PE innovation
- Configuration Overview - Label Distribution
- Packet Forwarding


■ The SEEREN Case

- The network
- What is CsC? – 6PE over CsS
- Configuration - Troubleshooting

■ Conclusions

Case Study: The network


- **SEEREN: South Eastern European Research & Educational Network**
 - **Upstream:** GRNET (Greek Research Network)
 - **Backup Upstream:** RoEduNet (via GRNET)
 - **Peers:** Serbia & Montenegrina, Bulgarian, Albanian, FYROM, Bosnian & Herzegovinan NRENs
 - **Connectivity:** Oteglobe, via a **CsC MPLS VPN**




What is CsC?

Carrier-Supporting-Carrier

- Designed for ISPs that are VPN customers of other (larger MPLS) ISP.
 - A VPN with very small virtual routing table (VRF)
- MPLS between CsC-CE and CsC-PE
 - CsC-CEs exchange limited labels with CsC-PE




CsC: Control & Data planes



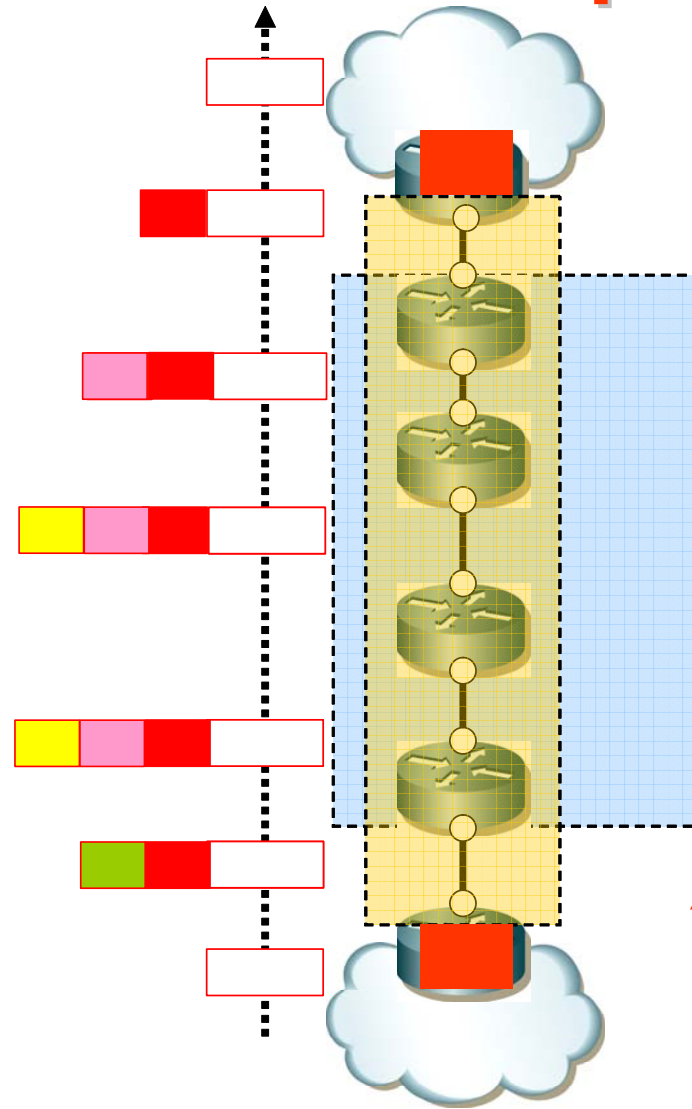
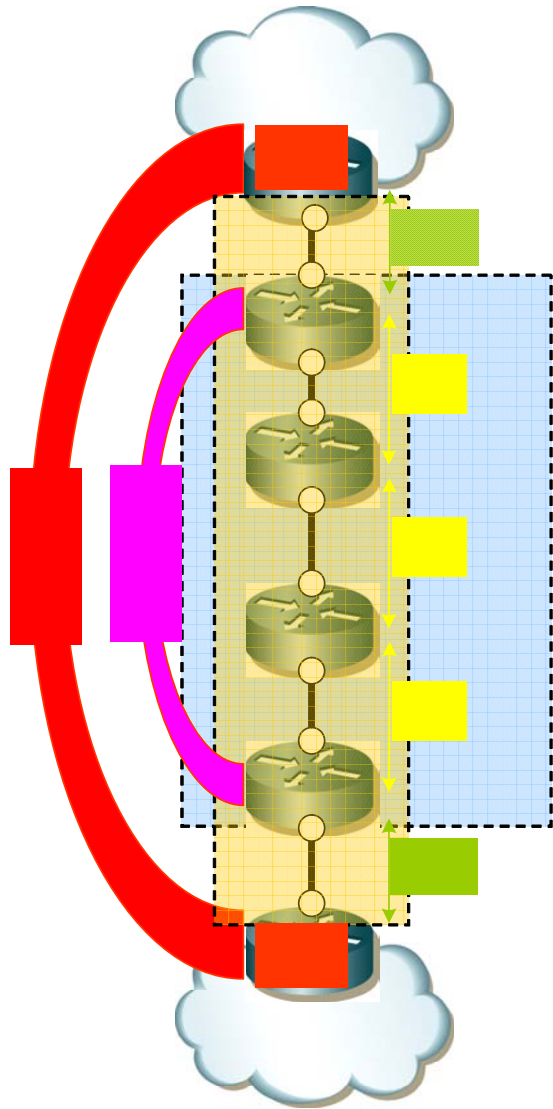
AMN

6PE over CsC

- **Integration of two techniques**
 - 6PE functionality is installed on the CEs instead of the PEs!
 - 6PE peers belong to different administrative domains.
 - **Three Label Stack !!!**
 - Feasible, because CE-PE connection uses MPLS.



6PE over CsC: Control & Data planes



AMREJ

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.
(OteGlobe)

Send Labels for
IPv6 routes
(AMREJ)

```
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 nei 147.91.0.112
BGP neighbor is 147.91.0.112, remote AS 13092, external link
Description: AMREJ-YUGOSLAVIA
BGP version 4, remote router ID 147.91.0.112
BGP state = Established, up for 2d00h
Last read 00:00:16, 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
```

```
athens-2#sh ipv6 cef 2001:4170::/32 internal
2001:4170::/32
  path list pointer 4355B5E0
    1 path -
      Nexthop path_pointer 43558B80 traffic share 1 path_list
      pointer 4355B5E0
        nexthop ::FFFF:147.91.0.112
        next_hop_len 0 adjacency pointer 4351A6B8
        refcount 2
        no loadinfo
        fast tag rewrite with AT1/0/0.1, point2point, tags
        imposed: {17 30}
```

6PE verification & troubleshooting

```
athens-2#sh ipv6 route 2001:4170::  
IPv6 Routing Table - 474 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   2001:4170::    via ::FFFF:147.91.0.112, IPv6-mpls
```

```
athens-2#show bgp ipv6 unicast 2001:4170::BGP routing table entry for 2001:4170::Paths: (1 available, best #1, table Global-IPv6-Table)  
Advertised to update-groups:  
    2  
13092, (received & used)  
    ::FFFF:147.91.0.112 from 147.91.0.112 (147.91.0.112)  
    Origin IGP, metric 0, localpref 100, valid, external, best
```

6PE verification & troubleshooting

```
athens-2#sh bgp ipv6 u labels | b 2001:4170::  
2001:4170::/32    ::FFFF:147.91.0.112  
                    nolabel/30
```

```
athens-2#sh ip cef 147.91.0.112  
147.91.0.112/32, version 48051, epoch 0, cached adjacency to  
ATM1/0/0.1  
0 packets, 0 bytes  
Flow: AS 12713, mask 32  
tag information set, shared, all rewrites owned  
local tag: BGP route head  
fast tag rewrite with ATM1/0/0.1, point2point, tags  
imposed: {18}  
via 62.75.33.245, 4 dependencies, recursive  
next hop 62.75.33.245, ATM1/0/0.1 via 62.75.33.245/32  
valid cached adjacency  
tag rewrite with ATM1/0/0.1, point2point, tags imposed:  
{17}
```



Presentation Outline

■ Introduction

- What is 6PE? - Basic characteristics
- Applicability
- Alternative solutions

■ 6PE Implementation Details

- MPLS VPN - 6PE innovation
- Configuration Overview - Label Distribution
- Packet Forwarding

■ The SEEREN Case

- The network
- What is CsC? – 6PE over CsS
- Configuration - Troubleshooting

■ Conclusions

Conclusions

- **6PE is applicable only if there is an MPLS core network and limited number of customers asks for IPv6 services. Otherwise, consider other solutions.**
 - Avoid changing the (IPv4) MPLS core network, a.k.a. minimal operational cost and risk

- **6PE is similar to MPLS VPNs technique in terms of technical implementation and complexity.**
 - Scalability properties

- **Requirements:**
 - CE is dual stack and exchange IPv6 routing information with local PE router.
 - P and PE routers support MPLS and IPv4 IGP!
 - PE routers support also MP-BGP and 6PE functionality

Acknowledgements

The oscar for animations goes to ...

Andreas Polyrakis (A.Polyrakis@noc.ntua.gr)

Dimitrios Kalogeras (D.Kalogeras@noc.ntua.gr)



Questions?

Thank You!

Thanassis Liakopoulos
aliako@grnet.gr