

Deploying MPLS L2VPN

PacNOG15

Nurul Islam Roman (nurul@apnic.net)

BUILT FOR
THE HUMAN
NETWORK



1

Abstract

- This session covers the fundamental and advanced topics associated with the deployment of Layer 2 VPNs over an MPLS network.
- The material presents a technology overview with an emphasis on ethernet-based point-to-point and multipoint VPNs. Session content then focuses on deployment considerations including: Signaling/Auto-discovery, OAM, Resiliency and Inter-AS.
- The attendee can expect to see sample configurations (IOS and IOS-XR) associated with the provisioning of L2VPNs.
- This session is intended for service providers and enterprise customers deploying L2VPNs over their MPLS network.

Agenda

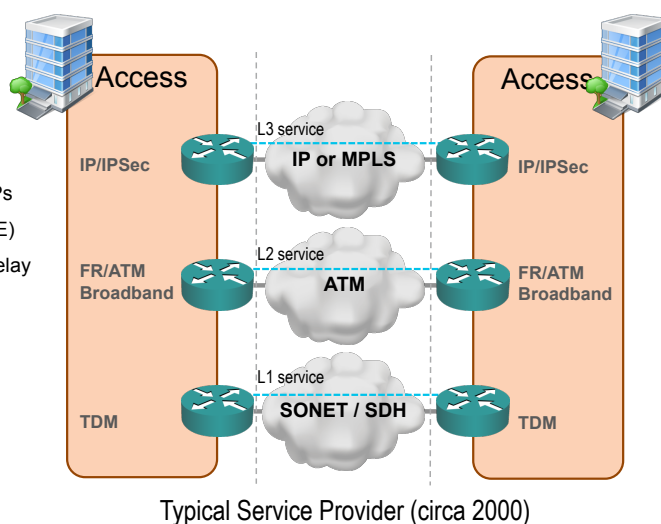
- Layer 2 VPN Motivation and Overview
- VPWS Reference Model
- VPLS Reference Model
- Pseudowire (PW) Signaling and PE Auto-Discovery
- Advanced Topics
- Summary

L2VPN Motivation and Overview

Motivation for L2VPNs

Old and New Drivers

- **Network Consolidation (circa 2000)**
 - Multiple access services (FR, ATM, TDM) required multiple core technologies
- **Enterprise Ethernet WAN Connectivity Services (circa 2005+)**
 - Ethernet well understood by Enterprise / SPs
 - CAPEX (lower cost per bit) / Growth (100GE)
 - Layer 2 VPN replacement to ATM/Frame Relay
 - Internet / Layer 3 VPN access (CE to PE)
- **Data Center Interconnection (DCI)**
- **Mobile Backhaul Evolution**
 - TDM /PDH to Dual/Hybrid to All-packet (IP/ Ethernet)
 - Single (voice + data) IP/Ethernet mobile backhaul universally accepted solution



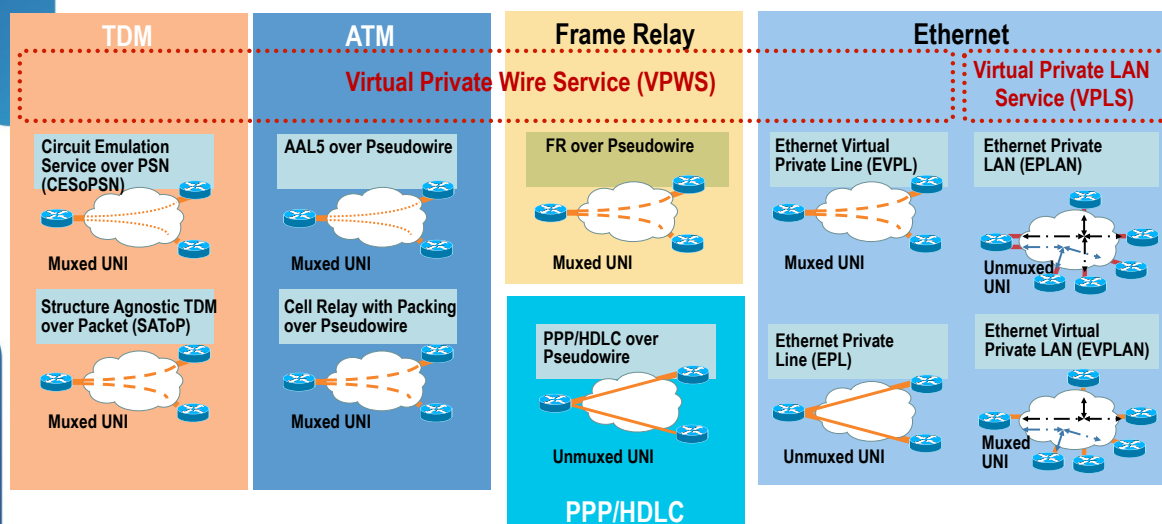
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Service Offerings

L2VPN Transport Services



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Layer 2 VPN Enabler

The Pseudowire

- L2VPNs are built with **Pseudowire (PW)** technology
- PWs provide a common intermediate format to **transport multiple types of network services** over a **Packet Switched Network (PSN)**
- PW technology provides **Like-to-Like** transport and also **Interworking (IW)**



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Virtual Private Wire Service (VPWS) Overview

Pseudowire Reference Model

- Any Transport Over MPLS (AToM) is Cisco's implementation of VPWS for IP/MPLS networks
- An Attachment Circuit (AC) is the physical or virtual circuit attaching a CE to a PE
- Customer Edge (CE) equipment perceives a PW as an **unshared link or circuit**



Ref: RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture, March 2005

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Layer 2 Transport over MPLS

Control Connection

- Targeted LDP session / BGP session / Static
 - Used for VC-label negotiation, withdrawal, error notification

The “emulated circuit” has **three (3) layers of encapsulation**

Tunnelling Component

- Tunnel header (Tunnel Label)
 - To get PDU from ingress to egress PE
 - MPLS LSP derived through static configuration (MPLS-TP) or dynamic (LDP or RSVP-TE)

Demultiplexing Component

- Demultiplexer field (VC Label)
 - To identify individual circuits within a tunnel
 - Could be an MPLS label, L2TPv3 header, GRE key, etc.

Layer 2 Encapsulation


- Emulated VC encapsulation (Control Word)
 - Information on enclosed Layer 2 PDU
 - Implemented as a 32-bit control word

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

VPWS Traffic Encapsulation




PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

VPWS Forwarding Plane Processing



PacNOG15

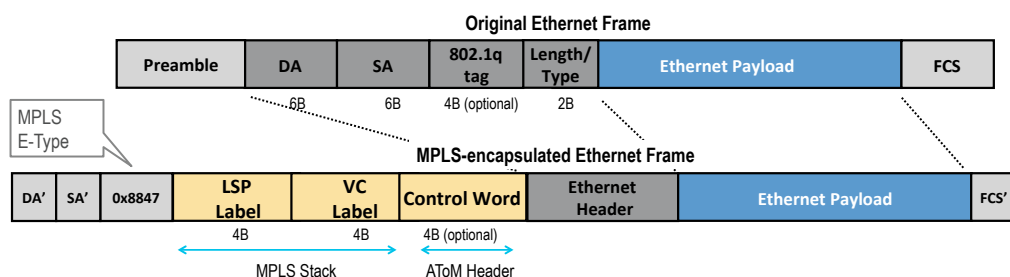
© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Virtual Private Wire Service (VPWS) Ethernet over MPLS (EoMPLS)

How Are Ethernet Frames Transported?

- Ethernet frames transported without Preamble, Start Frame Delimiter (SFD) and FCS
- Two (2) modes of operation supported:
 - **Ethernet VLAN mode** (VC type 0x0004) – created for VLAN over MPLS application
 - **Ethernet Port / Raw mode** (VC type 0x0005) – created for Ethernet port tunneling application



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Ethernet PW VC Type Negotiation

Cisco IOS

- Cisco devices by default will generally attempt to bring up an Ethernet PW using VC type 5
- If rejected by remote PE, then VC type 4 will be used
- Alternatively, Cisco device can be manually configured to use either VC type 4 or 5

```
7604-2(config-pw-class)#interworking ?
 ethernet   Ethernet interworking
 ip         IP interworking
 vlan      VLAN interworking

7604-2#show running-config
pseudowire-class test-pw-class-VC4
 encapsulation mpls
 interworking vlan
!
pseudowire-class test-pw-class-VC5
 encapsulation mpls
 interworking ethernet
```

Ethernet PW VC Type Negotiation

Cisco IOS-XR

- Cisco devices by default will generally attempt to bring up an Ethernet PW using VC type 5
- If rejected by remote PE, then VC type 4 will be used
- Alternatively, Cisco device can be manually configured to use either VC type 4 or 5

```
RP/0/RSP0/CPU0:ASR9000-2(config-l2vpn-pwc-
mpls)#transport-mode ?
 ethernet   Ethernet port mode
 vlan      Vlan tagged mode
RP/0/RSP0/CPU0:ASR9000-2(config-l2vpn-pwc-
mpls)#transport-mode vlan ?
 passthrough passthrough incoming tags

RP/0/RSP0/CPU0:ASR9000-2#show running-config l2vpn
l2vpn
pw-class test-pw-class-VC4
 encapsulation mpls
 transport-mode vlan

pw-class test-pw-class-VC4-passthrough
 encapsulation mpls
 transport-mode vlan passthrough

pw-class test-pw-class-VC5
 encapsulation mpls
 transport-mode ethernet
```


Introducing Cisco EVC Framework


Functional Highlights

Flexible service delimiters

- Single-tagged, Double-tagged
- VLAN Lists, VLAN Ranges
- Header fields (COS, Ethertype)

ANY service – ANY port

- Layer 2 Point-to-Point
- Layer 2 Multipoint
- Layer 3




Ethernet Service Layer

- Ethernet Flow Point (EFP)
- Ethernet Virtual Circuit (EVC)
- Bridge Domain (BD)
- Local VLAN significance

VLAN Header operations - VLAN Rewrites

- POP
- PUSH
- SWAP



PacNOG15

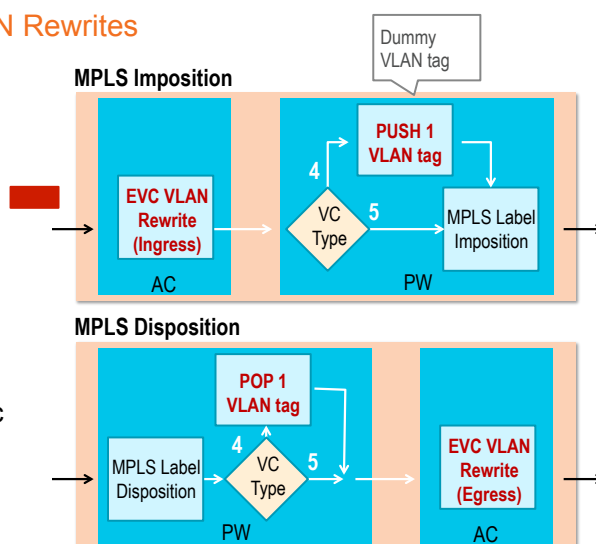
© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Encapsulation Adjustment Considerations

EoMPLS PW VC Type and EVC VLAN Rewrites

- VLAN tags can be added, removed or translated prior to VC label imposition or after disposition
 - Any VLAN tag(s), if retained, will appear as payload to the VC
- VC label imposition and service delimiting tag are independent from EVC VLAN tag operations
 - **Dummy VLAN tag** – RFC 4448 (sec 4.4.1)
- VC service-delimiting VLAN-ID is removed before passing packet to Attachment Circuit processing



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

18

Encapsulation Adjustment Considerations

VC 5 and EVC Rewrites



Single-tagged frame



Double-tagged frame

IOS-XR

```
l2vpn
pw-class class-VC5
encapsulation mpls
transport-mode ethernet

xconnect group Cisco-Live
p2p xc-sample-1
interface GigabitEthernet0/0/0/2.100
neighbor 102.102.102.102 pw-id 111
pw-class class-VC5

interface GigabitEthernet0/0/0/2.100 l2transport
encapsulation dot1q 10
rewrite ingress tag pop 1 symmetric
```

- POP VLAN 10
- No Push of Dummy tag (VC 5)

- No service-delimiting vlan expected (VC 5)
- PUSH VLAN 10

IOS

```
pseudowire-class class-VC5
encapsulation mpls
interworking ethernet

interface GigabitEthernet2/2
service instance 3 ethernet
encapsulation dot1q 10
rewrite ingress tag pop 1 symmetric
xconnect 104.104.104.104 111 encap mpls pw-class class-VC5
```

 MPLS label


PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Encapsulation Adjustment Considerations

VC 4 and EVC Rewrites



Single-tagged frame



Double-tagged frame

IOS-XR

```
l2vpn
pw-class class-VC4
encapsulation mpls
transport-mode vlan

xconnect group Cisco-Live
p2p xc-sample-1
interface GigabitEthernet0/0/0/2.100
neighbor 102.102.102.102 pw-id 111
pw-class class-VC4

interface GigabitEthernet0/0/0/2.100 l2transport
encapsulation dot1q 10
rewrite ingress tag pop 1 symmetric
```

- POP VLAN 10
- Push Dummy tag (VC 4)

- POP service-delimiting vlan (VC 4)
- PUSH VLAN 10

IOS

```
pseudowire-class class-VC4
encapsulation mpls
interworking vlan

interface GigabitEthernet2/2
service instance 3 ethernet
encapsulation dot1q 10
rewrite ingress tag pop 1 symmetric
xconnect 104.104.104.104 111 encap mpls pw-class class-VC4
```

 MPLS label

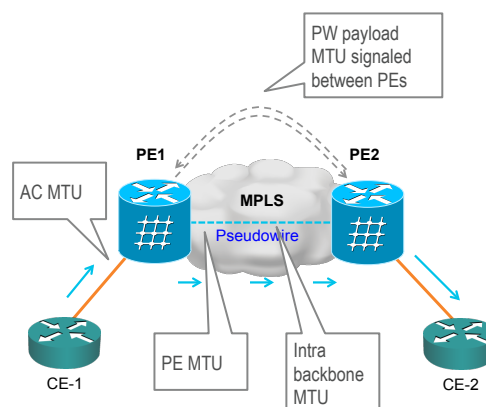
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

MTU Considerations

- No payload fragmentation supported
- Incoming PDU dropped if MTU exceeds AC MTU
- PEs exchange PW payload MTU as part of PW signaling procedures
 - Both ends must agree to use same value for PW to come UP
 - PW MTU derived from AC MTU
- No mechanism to check Backbone MTU
 - MTU in the backbone must be large enough to carry PW payload and MPLS stack



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

21

Ethernet MTU Considerations

Cisco IOS

- Interface MTU configured as largest ethernet payload size
 - 1500B default
 - Sub-interfaces / Service Instances (EFPs) MTU always inherited from main interface
- PW MTU used during PW signaling
 - By default, inherited from attachment circuit MTU
 - Submode configuration CLI allows MTU value to be set per subinterface/EFP in xconnect configuration mode (only for signaling purposes)
 - No MTU adjustments made for EFP rewrite (POP/PUSH) operations

```
interface GigabitEthernet0/0/4
description Main interface
mtu 1600
```

```
ASR1004-1#show int gigabitEthernet 0/0/4.1000 | include MTU
MTU 1600 bytes, BW 100000 Kbit/sec, DLY 100 usec,
```

Sub-interface MTU
inherited from Main
interface

```
interface GigabitEthernet0/0/4.1000
encapsulation dot1Q 1000
xconnect 106.106.106.106 111 encapsulation mpls
mtu 1500
```

PW MTU used during
signaling can be
overwritten

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Ethernet MTU Considerations

Cisco IOS XR

- Interface / sub-interface MTU configured as largest frame size – FCS (4B)

- 1514B default for main interfaces
- 1518B default for single-tagged subinterfaces
- 1522B default for double-tagged subinterfaces

- PW MTU used during PW signaling

- AC MTU – 14B + Rewrite offset
- E.g. POP 1 (- 4B), PUSH 1 (+ 4B)

```
interface GigabitEthernet0/0/0/2
description Main interface
mtu 9000
```

```
interface GigabitEthernet0/0/0/2.100 l2transport
encapsulation dot1q 100
rewrite ingress tag pop 1 symmetric
mtu 1518
```

By default, sub-interface MTU inherited from Main interface

Sub-interface MTU can be overwritten to match remote AC

XC MTU = 1518 – 14 – 4
= 1500B

```
RP/0/RSP0/CPU0:PE1#show l2vpn xconnect neighbor 102.102.102.102 pw-
id 11
Group Cisco-Live, XC xc-sample-1, state is down; Interworking none
AC: GigabitEthernet0/0/0/2.100, state is up
Type VLAN; Num Ranges: 1
VLAN ranges: [100, 100]
MTU 1500; XC ID 0x840014; interworking none
Statistics:
(snip)
```

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

Virtual Private LAN Service (VPLS) Overview

Virtual Private LAN Service

Overview

- Defines Architecture to provide **Ethernet Multipoint** connectivity sites, as if they were connected using a LAN
- VPLS operation **emulates an IEEE Ethernet switch**
- Two (2) signaling methods
 - RFC 4762 (LDP-Based VPLS)
 - RFC 4761 (BGP-Based VPLS)



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Virtual Private LAN Service

Reference Model

- **VFI (Virtual Forwarding Instance)**
 - Also called VSI (Virtual Switching Instance)
 - Emulates L2 broadcast domain among ACs and VCs
 - Unique per service. Multiple VFIs can exist same PE
- **AC (Attachment Circuit)**
 - Connect to CE device, it could be Ethernet physical or logical port
 - One or multiple ACs can belong to same VFI
- **VC (Virtual Circuit)**
 - EoMPLS data encapsulation, tunnel label used to reach remote PE, VC label used to identify VFI
 - One or multiple VCs can belong to same VFI
 - PEs must have a **full-mesh of PWs** in the VPLS core



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Virtual Private LAN Service

Operation

▪ Flooding / Forwarding


- Forwarding based on destination MAC addresses
- Flooding (Broadcast, Multicast, Unknown Unicast)

▪ MAC Learning/Aging/Withdrawal

- Dynamic learning based on Source MAC and VLAN
- Refresh aging timers with incoming packet
- **MAC withdrawal** upon topology changes

▪ Split-Horizon and Full-Mesh of PWs for loop-avoidance in core

- SP does not run STP in the core



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public


Why H-VPLS? Improved Scaling

▪ Flat VPLS

- Potential signaling overhead
- Packet replication at the edge
- Full PW mesh end-end

▪ Hierarchical-VPLS

- Minimizes signaling overhead
- Packet replication at the core only
- Full PW mesh in the core



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

28

VPLS Operation

Loop Prevention

- Core PW – Split Horizon ON
- Spoke PW – Split Horizon OFF (default)
- Split-Horizon Rules
 - Forwarding between Spoke PWs
 - Forwarding between Spoke and Core PWs
 - Forwarding between ACs and Core / Spoke PWs
 - Forwarding between ACs
 - Blocking between Core PWs



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

30/07/14

VPLS Operation

MAC Address Withdrawal

- Remove (flush) dynamic MAC addresses upon Topology Changes
 - Faster convergence – avoids blackholing
 - Uses LDP Address Withdraw Message (RFC 4762)
- H-VPLS dual-home example
 - U-PE detects failure of Primary PW
 - U-PE activates Backup PW
 - U-PE sends LDP MAC address withdrawal request to new N-PE
 - N-PE forwards the message to all PWs in the VPLS core and flush its MAC address table



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

30/07/14

Pseudowire (PW) Signaling and PE Auto-Discovery

VPWS / VPLS

An abstraction

■ Provisioning Model

- What information needs to be configured and in what entities
- Semantic structure of the endpoint identifiers (e.g. VC ID, VPN ID)

■ Discovery

- Provisioning information is distributed by a "discovery process"
- Distribution of endpoint identifiers

■ Signaling

- When the discovery process is complete, a signaling protocol is automatically invoked to set up pseudowires (PWs)

Provisioning Model



Discovery

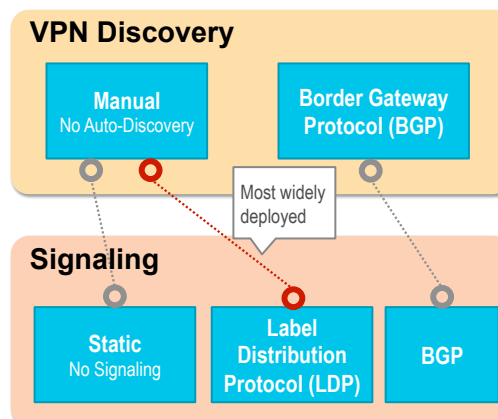


Signaling

VPWS

Discovery and Signaling Alternatives

- VPWS Signaling
 - LDP-based (RFC 4447)
 - BGP-based (informational draft)
draft-kompella-l2vpn-l2vpn
- VPWS with LDP-signaling and No auto-discovery
 - Most widely deployed solution
- Auto-discovery for point-to-point services not as relevant as for multipoint



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

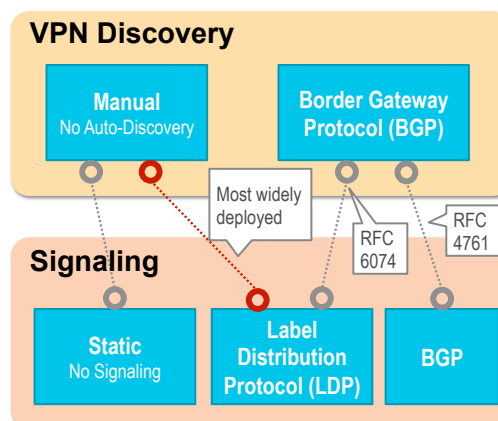
Cisco Public

33

VPLS

Discovery and Signaling Alternatives

- VPLS Signaling
 - LDP-based (RFC 4762)
 - BGP-based (RFC 4761)
- VPLS with LDP-signaling and No auto-discovery
 - Most widely deployed solution
 - Operational complexity for larger deployments
- BGP-based Auto-Discovery (BGP-AD) (RFC 6074)
 - Enables discovery of PE devices in a VPLS instance



PacNOG15

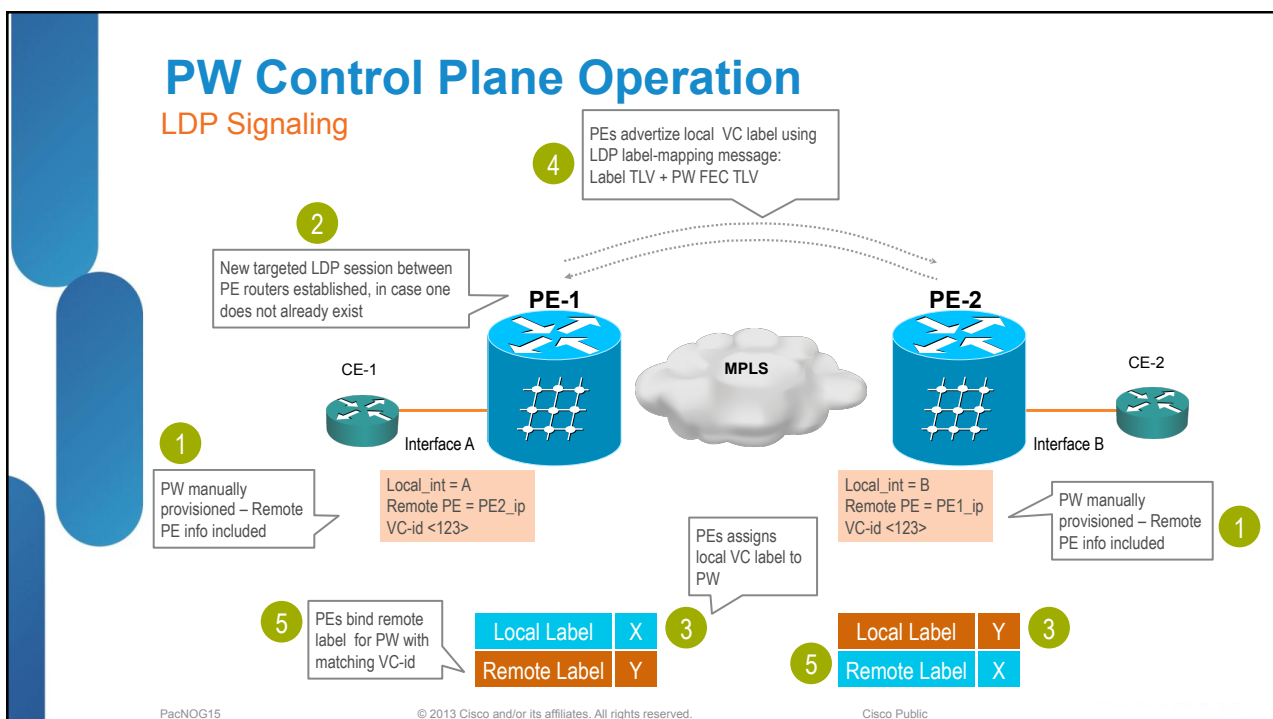
© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

34

Pseudowire (PW) Signaling and PE Auto-Discovery

LDP-based Signaling and Manual Provisioning



VPWS (EoMPLS) LDP Signaling

Cisco IOS (VLAN-based services)

```
hostname PE1
!
interface Loopback0
ip address 106.106.106.106 255.255.255.255
```

Sub-interface
based xconnect

```
interface GigabitEthernet2/4.300
encapsulation dot1q 300
xconnect 102.102.102.102 111 encapsulation mpls
```

OR

```
interface GigabitEthernet2/4
service instance 10 ethernet
encapsulation dot1q 300
rewrite ingress tag pop 1 symmetric
xconnect 102.102.102.102 111 encapsulation mpls
```

Service-Instance
(EFP) based xconnect

OR


```
interface Vlan 300
xconnect 102.102.102.102 111 encapsulation mpls
!
interface GigabitEthernet2/4
switchport mode trunk
switchport trunk allowed vlan 300
```

Interface VLAN (SVI)
based xconnect +
Switchport trunk / access

OR

```
interface Vlan 300
xconnect 102.102.102.102 111 encapsulation mpls
!
interface GigabitEthernet2/4
service instance 10 ethernet
encapsulation dot1q 300
rewrite ingress tag pop 1 symmetric
bridge-domain 300
```

Interface VLAN (SVI)
based xconnect +
Service instance BD



VPWS (EoMPLS) LDP Signaling

Cisco IOS (Port-based services)

```
hostname PE1
!
interface Loopback0
ip address 106.106.106.106 255.255.255.255
```

Main interface
based xconnect

```
interface GigabitEthernet2/5
xconnect 102.102.102.102 222 encapsulation mpls
```

OR

```
interface GigabitEthernet2/5
service instance 1 ethernet
encapsulation default
xconnect 102.102.102.102 111 encapsulation mpls
```

Service-Instance
(EFP) based xconnect
(encap default)

OR


```
interface Vlan 300
xconnect 102.102.102.102 111 encapsulation mpls
!
interface GigabitEthernet2/5
switchport mode dot1q-tunnel
switchport access vlan 300
```

Interface VLAN (SVI)
based xconnect +
Switchport dot1q-tunnel

OR

```
interface Vlan 300
xconnect 102.102.102.102 111 encapsulation mpls
!
interface GigabitEthernet2/5
service instance 1 ethernet
encapsulation default
bridge-domain 300
```

Interface VLAN (SVI)
based xconnect +
Service instance BD



VPWS (EoMPLS) LDP Signaling

Cisco IOS XR

```
hostname PE1
!
interface Loopback0
ipv4 address 106.106.106.106 255.255.255.255
```

```
l2vpn
xconnect group Cisco-Live
p2p xc-sample-1
interface GigabitEthernet0/0/0/2.100
neighbor 102.102.102.102 pw-id 111
```


```
p2p xc-sample-2
interface GigabitEthernet0/0/0/2.200
neighbor 102.102.102.102 pw-id 222
```

```
p2p xc-sample-3
interface GigabitEthernet0/0/0/6
neighbor 102.102.102.102 pw-id 333
```

Single-tagged
VLAN traffic to PW

```
interface GigabitEthernet0/0/0/2.100 l2transport
encapsulation dot1q 100
rewrite ingress tag pop 1 symmetric
```

```
interface GigabitEthernet0/0/0/2.200 l2transport
encapsulation dot1q 999-1010
rewrite ingress tag push dot1q 888 symmetric
```



Single-tagged range
VLAN traffic to PW

OR

Entire port
traffic to PW

```
interface GigabitEthernet0/0/0/6
l2transport
```

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

39

VPLS LDP Signaling / Manual provisioning

Cisco IOS

```
hostname PE1
!
interface Loopback0
ip address 192.0.0.1 255.255.255.255
```

```
l2 vfi sample-vfi manual
vpn id 1111
neighbor 192.0.0.2 encapsulation mpls
neighbor 192.0.0.3 2222 encapsulation mpls
neighbor 192.0.0.4 3333 encapsulation mpls
```

```
!
interface Vlan300
xconnect vfi sample-vfi
```


VFI associated to
VLAN interface (SVI)
via xconnect cmd

Bridge-Domain or
VLAN/switchport
configurations

```
interface GigabitEthernet2/4
service instance 333 ethernet
encapsulation dot1q 333
rewrite ingress tag pop 1 symmetric
bridge-domain 300
```

OR

```
interface GigabitEthernet2/4
switchport mode trunk
switchport trunk allowed vlan 300
```



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public


40

VPLS LDP Signaling / Manual provisioning

Cisco IOS XR

```
hostname PE1
!
interface Loopback0
  ipv4 address 192.0.0.1 255.255.255.255
!
interface GigabitEthernet0/0/0/14.101 l2transport
  encapsulation dot1q 101
  rewrite ingress tag pop 1 symmetric
```

```
l2vpn
  bridge group Cisco-Live
  bridge-domain bd101
  interface GigabitEthernet0/0/0/14.101
  vfi vfi101
    vpn-id 1111
    neighbor 192.0.0.2 pw-id 1111
    neighbor 192.0.0.3 pw-id 2222
    neighbor 192.0.0.4 pw-id 3333
```



GigabitEthernet0/0/0/14.101

Protocol-based CLI:
EFPs, PWs and VFI
as members of
Bridge Domain

VPN ID defined per VFI or
on a per-neighbor basis

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.


Cisco Public

41

H-VPLS LDP Signaling / Manual provisioning

Cisco IOS

```
hostname PE1
!
interface Loopback0
  ip address 192.0.0.1 255.255.255.255
!
l2 vfi sample-vfi manual
  vpn id 1111
  neighbor 192.0.0.2 encapsulation mpls
  neighbor 192.0.0.3 2222 encapsulation mpls
  neighbor 192.0.0.4 3333 encapsulation mpls
  neighbor 192.0.0.5 5555 encapsulation mpls no-split-horizon
  neighbor 192.0.0.6 5555 encapsulation mpls no-split-horizon
!
interface Vlan300
  xconnect vfi sample-vfi
```



Bridge-Domain or
VLAN/switchport
configurations

Spoke
PWs

```
interface GigabitEthernet2/4
  service instance 333 ethernet
  encapsulation dot1q 333
  rewrite ingress tag pop 1 symmetric
  bridge-domain 300
```

```
interface GigabitEthernet2/4
  switchport mode trunk
  switchport trunk allowed vlan 300
```

OR

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

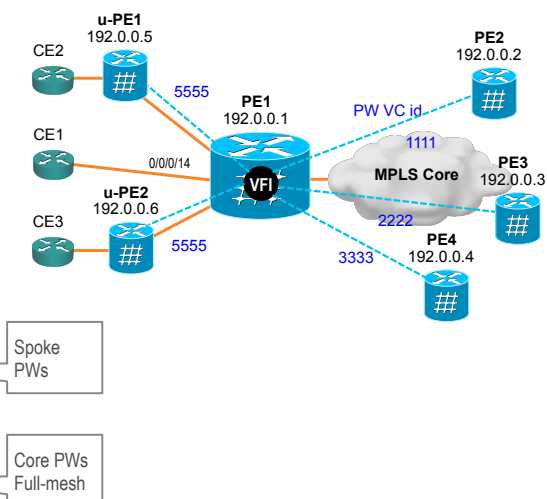
42

H-VPLS LDP Signaling / Manual provisioning

Cisco IOS XR

```
hostname PE1
!
interface Loopback0
  ipv4 address 192.0.0.1 255.255.255.255
!
interface GigabitEthernet0/0/0/14.101 l2transport
  encapsulation dot1q 101
  rewrite ingress tag pop 1 symmetric
```

```
l2vpn
  bridge group Cisco-Live
  bridge-domain bd101
  interface GigabitEthernet0/0/0/14.101
    neighbor 192.0.0.5 pw-id 5555
    neighbor 192.0.0.6 pw-id 5555
  !
  vfi vfi101
    vpn-id 1111
    neighbor 192.0.0.2 pw-id 1111
    neighbor 192.0.0.3 pw-id 2222
    neighbor 192.0.0.4 pw-id 3333
```



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public


43

Pseudowire (PW) Signaling and PE Auto-Discovery

BGP-based AutoDiscovery (BGP-AD) and LDP Signaling

BGP Auto-Discovery (BGP-AD)

- Eliminates need to manually provision VPLS neighbors
- Automatically detects when new PEs are added / removed from the VPLS domain
- Uses BGP Update messages to advertise PE/VFI mapping (VPLS NLRI)
- Typically used in conjunction with BGP Route Reflectors to minimize iBGP full-mesh peering requirements
- Two (2) RFCs define use of BGP for VPLS AD⁽¹⁾
 - RFC 6074 – when LDP used for PW signaling
 - RFC 4761 – when BGP used for PW signaling



(1) VPLS BGP NLRIs from RFC 6074 and 4761 are different in format and thus not compatible, even though they share same AFI / SAFI values

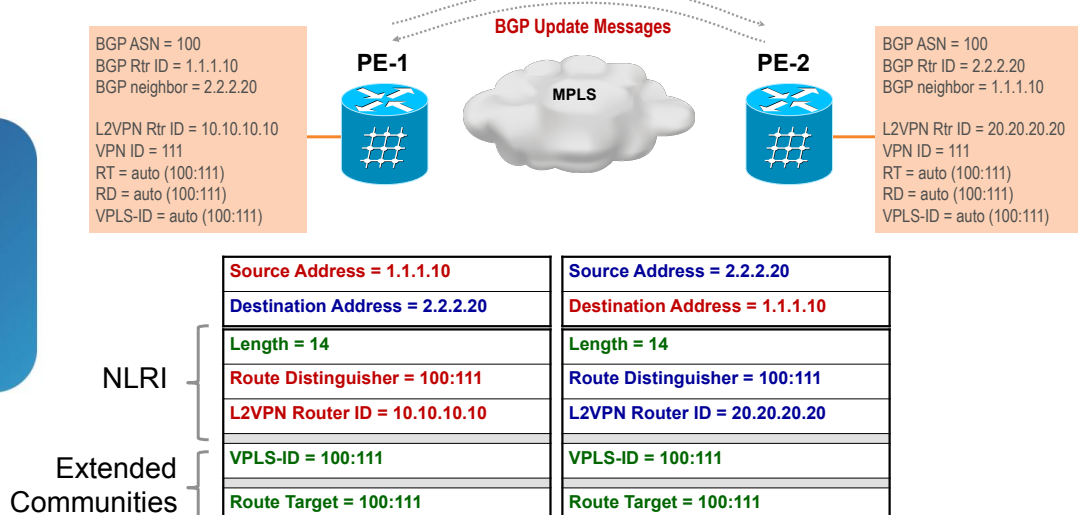
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

45

What is Discovered? NLRI + Extended Communities



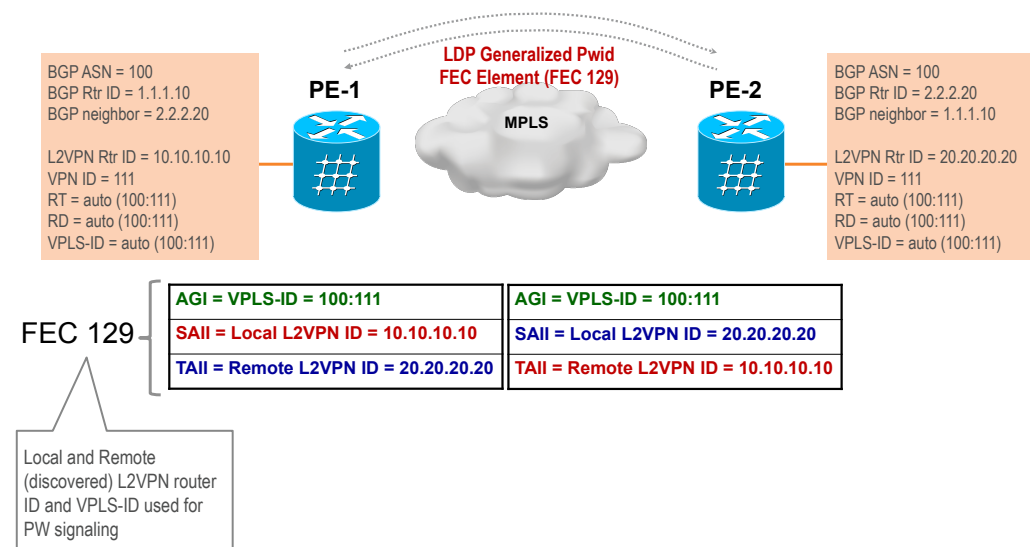
PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

46

What is Signaled?



PacNOG15

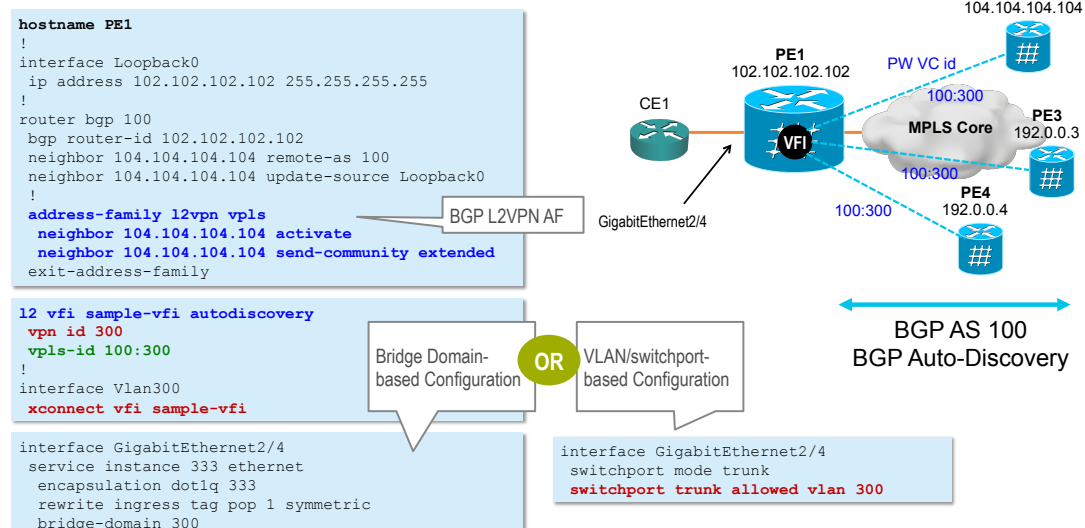
© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

47

VPLS LDP Signaling and BGP-AD

Cisco IOS



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

48

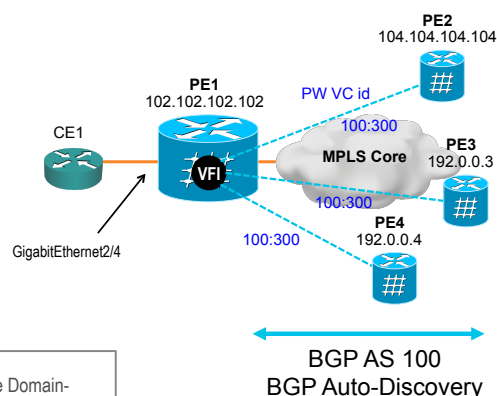
VPLS LDP Signaling and BGP-AD

Cisco IOS (NEW Protocol-based CLI)

```
hostname PE1
!
interface Loopback0
 ip address 102.102.102.102 255.255.255.255
!
router bgp 100
 bgp router-id 102.102.102.102
 neighbor 104.104.104.104 remote-as 100
 neighbor 104.104.104.104 update-source Loopback0
!
 address-family ipv4 vpls
  neighbor 104.104.104.104 activate
  neighbor 104.104.104.104 send-community extended
 exit-address-family

ipv4 vrf context sample-vrf
 vrf id 300
 autodiscovery bgp signaling ldp
 vpls-id 100:300
!
bridge-domain 300
 member vrf sample-vrf
 member GigabitEthernet2/4 service instance 333

interface GigabitEthernet2/4
 service instance 333 ethernet
 encapsulation dot1q 333
 rewrite ingress tag pop 1 symmetric
```



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

49

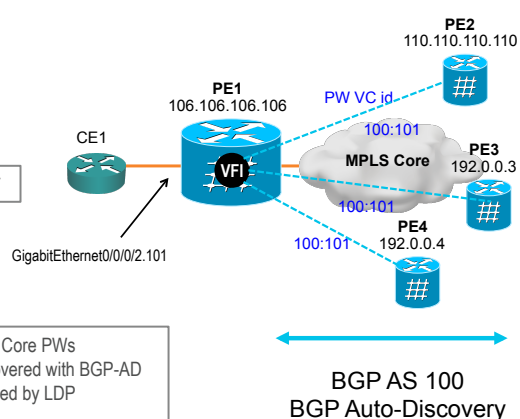
VPLS LDP Signaling and BGP-AD

Cisco IOS XR

```
hostname PE1
!
interface Loopback0
 ip address 106.106.106.106 255.255.255.255
!
interface GigabitEthernet0/0/0/2.101 12transport
 encapsulation dot1q 101
 rewrite ingress tag pop 1 symmetric

router bgp 100
 bgp router-id 106.106.106.106
 address-family 12vpn vpls-vpws
 neighbor 110.110.110.110
 remote-as 100
 update-source Loopback0
 address-family 12vpn vpls-vpws

12vpn
 bridge group Cisco-Live
 bridge-domain bd101
 interface GigabitEthernet0/0/0/2.101
 vfi vfi101
 vpn-id 11101
 autodiscovery bgp
 rd auto
 route-target 100:101
 signaling-protocol ldp
 vpls-id 100:101
```



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public


50

H-VPLS LDP Signaling and BGP-AD / Manual provisioning

Cisco IOS

```
hostname PE1
!
interface Loopback0
ip address 102.102.102.102 255.255.255.255
!
l2 vfi sample-vfi autodiscovery
vpn id 300
vpls-id 100:300
neighbor 192.0.0.5 5555 encapsulation mpls no-split-horizon
neighbor 192.0.0.6 5555 encapsulation mpls no-split-horizon
```

Manually
provisioned
Spoke PWs



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

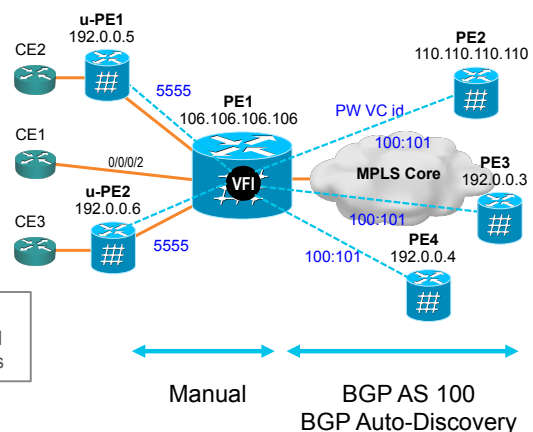
51

H-VPLS LDP Signaling and BGP-AD / Manual provisioning

Cisco IOS XR

```
hostname PE1
!
l2vpn
bridge group Cisco-Live
bridge-domain bd101
interface GigabitEthernet0/0/0/2.101
!
neighbor 192.0.0.5 pw-id 5555
neighbor 192.0.0.6 pw-id 5555
!
vfi vfi101
vpn-id 11101
autodiscovery bgp
rd auto
route-target 100:101
signaling-protocol ldp
vpls-id 100:101
```

Manually
provisioned
Spoke PWs



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

52

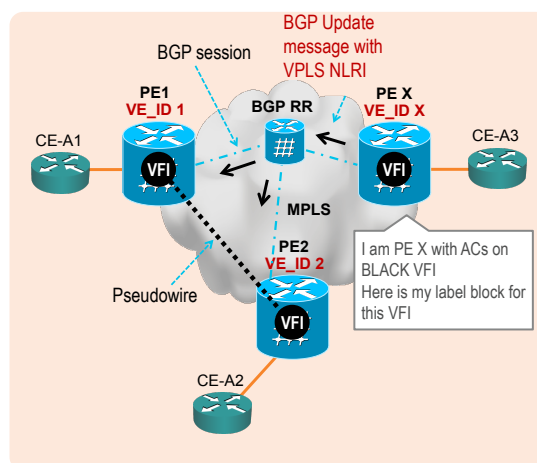
Pseudowire (PW) Signaling and PE Auto-Discovery

BGP-based Signaling and AutoDiscovery and BGP Signaling

BGP Signaling and Auto-Discovery

Overview

- RFC 4761¹ defines use of BGP for VPLS PE Auto-Discovery and Signaling
- All PEs within a given VPLS are assigned a **unique VPLS Edge device ID (VE ID)**
- A PE X wishing to send a VPLS update sends the same **label block** information to all other PEs using **BGP VPLS NLRI**
- Each **receiving PE infers the label** intended for PE X by adding its (unique) VE ID to the label base
 - Each receiving PE gets a unique label for PE X for that VPLS

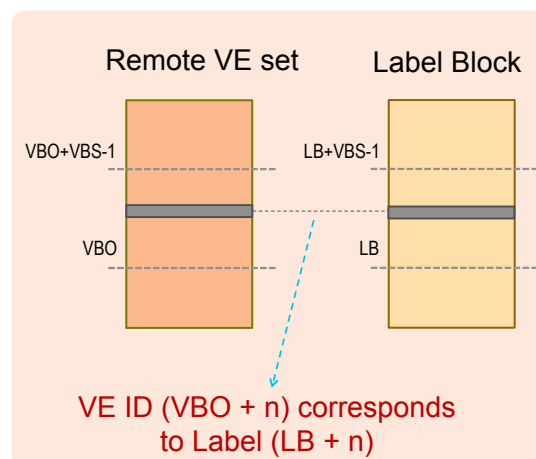


(1) VPLS BGP NLRI from RFC 6074 and 4761 are different in format and thus not compatible, even though they share same AFI / SAFI values

BGP Signaling and Auto-Discovery

Label Blocks

- RFC 4761 is primarily based on the concept of **Label Blocks**
 - Contiguous set of local labels
 - Label Block boundary **advertised using BGP VPLS NLRI**
- **Label Base (LB)** – start of label block
- **VE Block Size (VBS)** – size of label block
- **VE Block Offset (VBO)** – start of remote VE set



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

55


VPLS BGP Signaling and BGP-AD

Cisco IOS XR

```
hostname PE1
!
interface Loopback0
  ipv4 address 106.106.106.106 255.255.255.255
!
router bgp 100
  bgp router-id 106.106.106.106
  address-family l2vpn vpls-vpws
  neighbor 110.110.110.110
  remote-as 100
  update-source Loopback0
  address-family l2vpn vpls-vpws
```

```
l2vpn
  bridge group Cisco-Live
  bridge-domain bd102
  interface GigabitEthernet0/0/0/2.102
  vfi vfi102
    vpn-id 11102
    autodiscovery bgp
    rd auto
    route-target 100:102
    signaling-protocol bgp
    ve-id 5
```

VE-id must be unique in a VPLS instance



BGP Auto-Discovery attributes
VPLS VFI attributes
Signaling attributes

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

56

VPLS BGP Signaling and BGP-AD

Cisco IOS (NEW Protocol-based CLI)

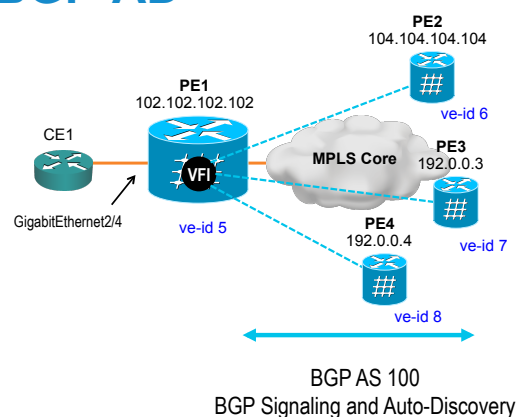
```
hostname PE1
!
interface Loopback0
ip address 102.102.102.102 255.255.255.255
!
router bgp 100
bgp router-id 102.102.102.102
neighbor 104.104.104.104 remote-as 100
neighbor 104.104.104.104 update-source Loopback0
!
address-family 12vpn vpls
neighbor 104.104.104.104 activate
neighbor 104.104.104.104 send-community extended
neighbor 104.104.104.104 suppress-signaling-protocol ldp
exit-address-family
```

```
12vpn vfi context sample-vfi
vpn id 3300
autodiscovery bgp signaling bgp
ve id 5
ve range 10
```

VE-id must be
unique in a
VPLS instance

```
bridge-domain 300
member vfi sample-vfi
member GigabitEthernet2/4 service instance 333
!
interface GigabitEthernet2/4
service instance 333 ethernet
encapsulation dot1q 300
rewrite ingress tag pop 1 symmetric
```

Bridge Domain-
based Configuration



Cisco Public

57

Advanced Topics

Operations, Administration and Management
Virtual Circuit Connection Verification (VCCV)

Virtual Circuit Connection Verification (VCCV)

Overview

- Provides end-to-end fault verification and detection for emulated PW service (RFC 5085)
 - MPLS LSP ping monitors PSN tunnel (PE-PE connectivity)
 - VCCV sends control packets (intercepted by remote PE) in-band of PWs
- VCCV capabilities negotiated during PW signaling phase
- Disposition capabilities to identify VCCV packets are:
 - Type 1: uses PW Control Word with 0001_b as first nibble
 - Type 2: uses MPLS Router Alert (RA) label
 - Type 3: uses MPLS PW label with TTL == 1

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Pseudowire Connectivity Verification

Cisco IOS

```

7604-2#ping mpls pseudowire 104.104.104.104 111000 ?
destination  Destination address or address range
exp          EXP bits in mpls header
interval     Send interval between requests in msec
pad          Pad TLV pattern
repeat       Repeat count
reply        Reply mode
revision     Echo Packet TLV versioning
segment      Segment of the MS-PW
size         Packet size
source       Source specified as an IP address
sweep        Sweep range of sizes
timeout      Timeout in seconds
verbose      verbose output mode

7604-2#ping mpls pseudowire 104.104.104.104 111000
Sending 5, 100-byte MPLS Echos to 104.104.104.104,
        timeout is 2 seconds, send interval is 0 msec:
Codes: '.' - success, 'Q' - request not sent, '-' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'l' - Label switched with FEC change, 'd' - see DDMAP for return code,
       'X' - unknown return code, 'x' - return code 0
Type escape sequence to abort.
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

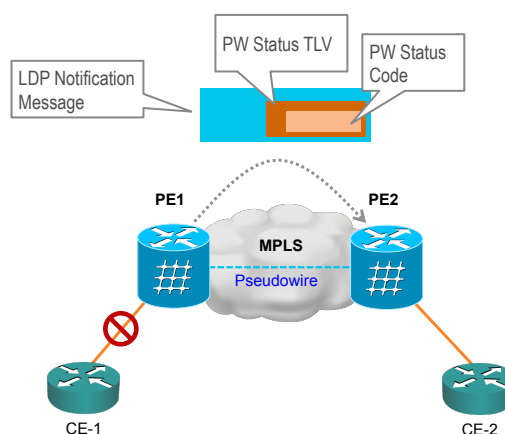
Cisco Public

Advanced Topics

Operations, Administration and Management Pseudowire Status Signaling

MPLS PW Status Signaling

- **PW Status Communication** achieved via
 - Label Withdraw method, or
 - PW Status TLV method
- **Label Withdraw Method**
 - Label Mapping Message sent only when AC in UP state
 - PW status signaled by withdrawing labels
 - Found in earlier implementations
- **PW Status TLV Method**
 - PW status signaled using a PW status TLV (labels not withdrawn)
 - Label Mapping Message sent when PW provisioned (irrespective of AC status)
 - Faster PWs setup as Label allocation processes are independent from each other

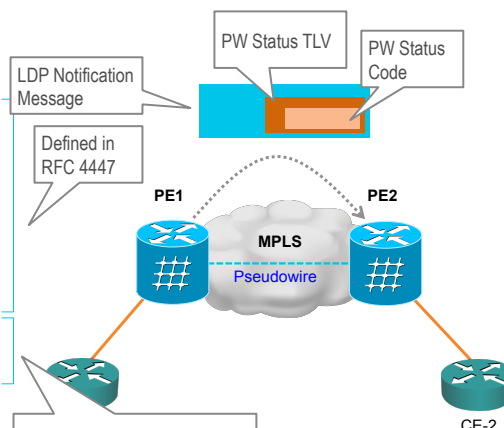


MPLS PW Status Signaling

PW Status Codes

Bit Mask Value	Status Description
0x00000000	Pseudowire forwarding (clear all failures)
0x00000001	Pseudowire Not Forwarding
0x00000002	Local Attachment Circuit (ingress) Receive Fault
0x00000004	Local Attachment Circuit (egress) Transmit Fault
0x00000008	Local PSN-facing PW (ingress) Receive Fault
0x00000010	Local PSN-facing PW (egress) Transmit Fault
0x00000020	PW Preferential Forwarding Status
0x00000040	PW Request Switchover Status

Used in PW redundancy applications. Indicates PW in Active or Standby state



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

MPLS PW Status Signaling

Cisco IOS

```
7604-2#show mpls l2transport vc 2001 detail
Local interface: Gi2/4 up, line protocol up, Eth VLAN 2001 up
Destination address: 101.101.101.101, VC ID: 2001, VC status: up
Output interface: Gi2/2, imposed label stack (41)
Preferred path: not configured
Default path: active
Next hop: 10.10.2.101
Create time: 3d00h, last status change time: 3d00h
Signaling protocol: LDP, peer 101.101.101.101:0 up
Targeted Hello: 102.102.102.102 (LDP Id) -> 101.101.101.101, LDP is UP
Status TLV support (local/remote) : enabled/supported
LDP route watch : enabled
Label/status state machine : established, LruRru
(snip)
Last local LDP TLV status sent: No fault
Last remote LDP TLV status rcvd: No fault
Last remote LDP ADJ status rcvd: No fault
MPLS VC labels: local 33, remote 41
PWID: 70065
(snip)
```

```
7604-1#show mpls l2transport vc 2001 detail | include Status
Status TLV support (local/remote) : enabled/not supported
```

```
7604-2#show mpls l2transport vc 2001 detail | include Status
Status TLV support (local/remote) : disabled/supported
```

pseudowire-class no-status
encapsulation mpls
no status

Status TLV support ON by default
Can be disabled on a per PW class basis

Both Local / Remote PEs support Status TLV

Status TLV not supported by Remote PE

Status TLV disabled on Local PE

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

64

MPLS PW Status Signaling

Cisco IOS-XR

```
RP/0/RSP0/CPU0:ASR9000-2#show l2vpn xconnect neighbor 102.102.102.102 pw-id 111 detail
Group Cisco-Live, XC xc-sample-1, state is up; Interworking none
AC: GigabitEthernet0/0/0/2.100, state is up
(snip)
```

```
PW: neighbor 102.102.102.102, PW ID 111, state is up ( established )
Encapsulation MPLS, protocol LDP
Source address 106.106.106.106
PW type Ethernet VLAN, control word disabled, interworking none
PW backup disable delay 0 sec
Sequencing not set
```

PW Status TLV in use

MPLS	Local	Remote
Label	16010	22
Group ID	0x4000140	0x0
Interface	GigabitEthernet0/0/0/2.100	*** To ME3400-24TS-1 gig0/1 ***
MTU	1500	1500
Control word	disabled	disabled
PW type	Ethernet VLAN	Ethernet VLAN
VCCV CV type	0x2	0x12
	(LSP ping verification)	(LSP ping verification)
VCCV CC type	0x6	0x6
	(router alert label)	(router alert label)
	(TTL expiry)	(TTL expiry)

Incoming Status (PW Status TLV):

Status code: 0x0 (Up) in Notification message

Outgoing Status (PW Status TLV):

Status code: 0x0 (Up) in Notification message

(snip)

```
RP/0/RSP0/CPU0:ASR9000-2(config-l2vpn)#pw-status ?
disable Disable PW status
```

Status TLV support ON by default
Can be disabled globally under l2vpn configuration

Both Local / Remote PEs support Status TLV

65


Advanced Topics

Resiliency

Pseudowire Redundancy

High Availability in L2VPN Networks

Solutions



- IP Fast Re-Route (FRR) / MPLS FRR

1 PSN core failure

- Pseudowire Redundancy:

2 PSN end-to-end routing failure – Redundant PEs

3 PE failure – Redundant PEs

4 Attachment circuit failure – AC Diversity

5 CE failure – Redundant CEs

PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

One-Way Pseudowire Redundancy

Overview

- Allows dual-homing of one local PE to one or two remote PEs
- Two pseudowires - primary & backup provide redundancy for a single AC
- Faults on the primary PW cause failover to backup PW
- Multiple backup PWs (different priorities) can be defined
- Alternate LSPs (TE Tunnels) can be used for additional redundancy



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

Pseudowire Redundancy

Preferential Forwarding Status Bit

- Extensions to PW status codes
 - draft-ietf-pwe3-redundancy-bit
- Allows PEs to signal local forwarding status of the PW (Active or Standby)
- A PW is selected for forwarding when declared as Active by both PEs
- Minimize service downtime during PW failover
 - Backup PWs always signaled before failures and held in Standby mode
- Allows VCCV capability over a backup PW
 - OAM over backup PWs
 - SP monitors backup PWs prior to its usage



PacNOG15


© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

One-Way Pseudowire Redundancy

Failure Protection Points

- Failure 1 - Core failures handled by IGP re-routing / IP/MPLS FRR do not trigger pseudowire switchover
- Failure 2 - Loss of route to remote PE as notified by IGP (PE isolation)
- Failure 3 - Loss of Remote PE
- How to detect PE failures?
 - LDP Fast Failure Detection (FFD) - monitors IGP route availability of LDP peer (2-3 sec or sub-sec with Fast IGP) (a.k.a. Route-Watch)
 - LDP session timeout (default = 3 x 30 sec)
 - BFD timeout (multi-hop PE-to-PE BFD session) (a.k.a. "xconnect client" feature)



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

70

Configuring Pseudowire Redundancy

Cisco IOS

```
hostname PE1
interface Loopback0
ip address 102.102.102.102 255.255.255.255
```

```
interface GigabitEthernet2/4
service instance 170 ethernet
encapsulation dot1q 170
rewrite ingress tag pop 1 symmetric
xconnect 104.104.104.104 170 encapsulation mpls
  backup peer 106.106.106.106 170170
mtu 1500
```

```
7604-2#show xconnect peer 104.104.104.104 vcid 170
Legend:  XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
UP=Up      DN=Down          AD=Admin Down    IA=Inactive
SB=Standby HS=Hot Standby   RV=Recovering NH=No Hardware
```

XC	ST	Segment 1	S1 Segment 2	S2
UP	pri	ac	Gi2/4:170(Eth VLAN)	UP mpls 104.104.104.104:170


```
7604-2#show xconnect peer 106.106.106.106 vcid 170170
Legend:  XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
UP=Up      DN=Down          AD=Admin Down    IA=Inactive
SB=Standby HS=Hot Standby   RV=Recovering NH=No Hardware
```

XC	ST	Segment 1	S1 Segment 2	S2
IA	sec	ac	Gi2/4:170(Eth VLAN)	UP mpls 106.106.106.106:170170

PacNOC15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public



Primary PW in UP state

Redundant PW in Standby state

Configuring Pseudowire Redundancy

Cisco IOS XR

```
hostname PE1
interface Loopback0
  ipv4 address 106.106.106.106 255.255.255.255
!
interface GigabitEthernet0/0/0/2.180 l2transport
  encapsulation dot1q 180
  rewrite ingress tag pop 1 symmetric
```

```
l2vpn
xconnect group Cisco-Live
p2p xc-sample-6
interface GigabitEthernet0/0/0/2.180
  neighbor 104.104.104.104 pw-id 180
  pw-class sample-CW-ON
  backup neighbor 102.102.102.102 pw-id 180180
  pw-class sample-CW-ON
```


```
RP/0/RSP0/CPU0:ASR9000-2#show l2vpn xconnect group Cisco-Live xc xc-sample-6
Sun Apr 15 20:18:50.180 UTC
```

```
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
        SB = Standby, SR = Standby Ready, (PP) = Partially Programmed
```

XConnect Group	Name	ST	Segment 1 Description	ST	Segment 2 Description	ST
Cisco-Live	xc-sample-6	UP	Gi0/0/0/2.180	UP	104.104.104.104 180	UP
					Backup	102.102.102.102 180180

Primary PW in UP state

Redundant PW in Standby state





Advanced Topics

Multi-Segment Pseudowires

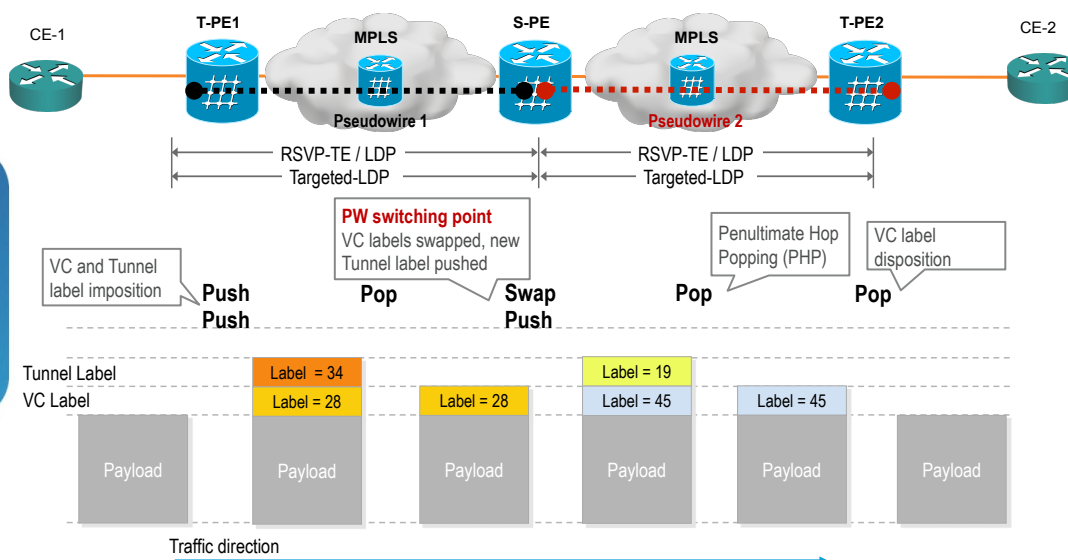
Multi-Segment Pseudowires

Overview

- Separate IGP processes (or areas) for separate MPLS Access networks
- T-PE – Terminating Provider Edge
 - Customer facing PE, hosting the first or last segment of a MS-PW
- S-PE – Switching Provider Edge
 - Switches control / data planes of preceding and succeeding segments
 - Control Word, sequencing, or original packet header not examined
 - VC labels swapped
 - VC Type, MTU should match end-to-end
 - One or more S-PEs can be used depending on number of segments
- MS-PW uses same signaling procedures and TLVs described in RFC 4447



Multi-Segment Pseudowires



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public


Configuring MS-PWs

Cisco IOS

```
hostname S-PE
interface Loopback0
ip address 104.104.104.104 255.255.255.255
```

```
12 vfi sample-ms-pw-1 point-to-point
neighbor 106.106.106.106 222190 encapsulation mpls
neighbor 102.102.102.102 111190 encapsulation mpls
```

MS-PW



```
7604-3#show xconnect peer 102.102.102.102 vcid 111190
```

```
Legend: XC ST=Xconnect State S1=Segment1 State S2=Segment2 State
UP=Up DN=Down AD=Admin Down IA=Inactive
SB=Standby HS=Hot Standby RV=Recovering NH=No Hardware
```

```
XC ST Segment 1 S1 Segment 2 S2
UP mpls 106.106.106.106:222190 UP mpls 102.102.102.102:111190 UP
```

```
7604-3#show xconnect peer 102.102.102.102 vcid 111190 detail
```

```
Legend: XC ST=Xconnect State S1=Segment1 State S2=Segment2 State
UP=Up DN=Down AD=Admin Down IA=Inactive
SB=Standby HS=Hot Standby RV=Recovering NH=No Hardware
```

```
XC ST Segment 1 S1 Segment 2 S2
UP mpls 106.106.106.106:222190 UP mpls 102.102.102.102:111190
Local VC label 65536 Local VC label 65549
Remote VC label 16029 Remote VC label 47
pw-class: pw-class:
```

S-PE labels for each PW segment

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

76


Configuring MS-PWs

Cisco IOS XR

```
hostname S-PE
interface Loopback0
  ipv4 address 106.106.106.106 255.255.255.255
```

```
l2vpn
  xconnect group Cisco-Live
  p2p xc-sample-8
  neighbor 102.102.102.102 pw-id 111200
  !
  neighbor 104.104.104.104 pw-id 222200
```

MS-PW



```
RP/0/RSP0/CPU0:ASR9000-2#show l2vpn xconnect group Cisco-Live xc-name xc-sample-8
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
        SB = Standby, SR = Standby Ready, (PP) = Partially Programmed
```

XConnect			Segment 1		Segment 2	
Group	Name	ST	Description	ST	Description	ST

Cisco-Live	xc-sample-8					
		UP	102.102.102.102	111200	UP	104.104.104.104
						222200
						UP


Advanced Topics

L2VPN and MPLS Traffic Engineering

MPLS Traffic Engineering

AToM Tunnel Selection

- Couples Layer-2 VPNs with MPLS TE
- Static mapping between PW and TE Tunnel on PE
- Implies PE-to-PE TE deployment
- TE tunnel defined as preferred path for pseudowire
- Traffic could fall back to peer LSP if tunnel goes down



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

AToM Tunnel Selection

Cisco IOS

```
7604-2(config-pw-class)#preferred-path interface ?
```

```
Tunnel      Tunnel interface
Tunnel-tp   MPLS Transport Profile interface
```

```
pseudowire-class sample-Tunnel-select
encapsulation mpls
preferred-path interface Tunnel200
```

Tunnel Selection preferred-path interface configured under PW class

```
interface GigabitEthernet2/4
service instance 150 ethernet
encapsulation dot1q 150
rewrite ingress tag pop 1 symmetric
xconnect 106.106.106.106 150 encapsulation mpls pw-class sample-Tunnel-select
mtu 1500
```

```
7604-2#show mpls l2transport vc 150 detail
Local interface: Gi2/4 up, line protocol up, Eth VLAN 150 up
Interworking type is Ethernet
Destination address: 106.106.106.106, VC ID: 150, VC status: up
Output interface: Tu200, imposed label stack {65550 16025}
Preferred path: Tunnel200, active
Default path: ready
Next hop: point2point
Create time: 01:17:27, last status change time: 01:05:29
(snip)
```

PW in UP state and mapped to an egress Tunnel interface

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

80

AToM Tunnel Selection

Cisco IOS XR

```
RP/0/RSP0/CPU0:ASR9000-2(config-l2vpn-pwc-mpls)#preferred-path interface ?
```

```
tunnel-ip Specify IP tunnel interface name for preferred path
tunnel-te Specify TE tunnel interface name for preferred path
tunnel-tp Specify TP tunnel interface name for preferred path
```

```
l2vpn
pw-class sample-Tunnel-select
encapsulation mpls
preferred-path interface tunnel-te 200
```

Tunnel Selection preferred-path interface configured under PW class

```
l2vpn
xconnect group Cisco-Live
p2p xc-sample-4
interface GigabitEthernet0/0/0/2.150
neighbor 102.102.102.102 pw-id 150
pw-class sample-Tunnel-select
```

```
RP/0/RSP0/CPU0:ASR9000-2#show l2vpn xconnect group Cisco-Live xc-name xc-sample-4 detail
Group Cisco-Live, XC xc-sample-4, state is up; Interworking none
AC: GigabitEthernet0/0/0/2.150, state is up
(snip)
```

```
PW: neighbor 102.102.102.102, PW ID 150, state is up ( established )
PW class sample-Tunnel-select, XC ID 0xc0000009
Encapsulation MPLS, protocol LDP
Source address 106.106.106.106
PW type Ethernet, control word enabled, interworking none
PW backup disable delay 0 sec
Sequencing not set
Preferred path tunnel TE 200, fallback enabled
(snip)
```

PW in UP state and mapped to an egress Tunnel interface

PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

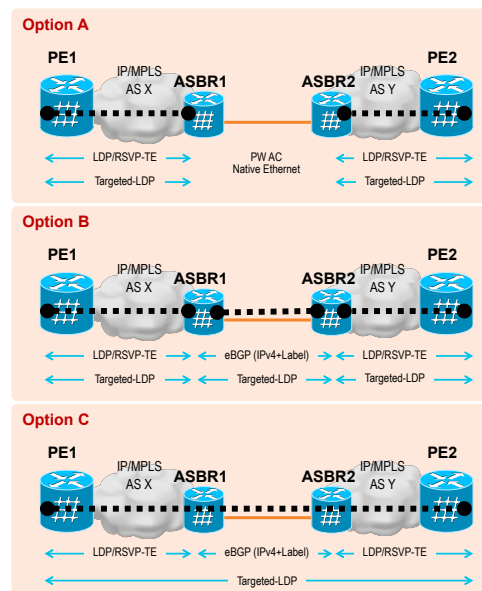
81

Advanced Topics

L2VPN Inter – Autonomous Systems (I-AS)

L2VPN Inter-AS

- Three (3) deployment models
- Option A
 - No reachability information shared between AS
- Option B
 - Minimal reachability information shared between AS
 - ASBR configured as S-PEs (multi-segment PWs)
 - eBGP (IPv4 prefix + label) used to build PSN tunnel between AS
- Option C
 - Significant reachability information shared between AS
 - Single-segment PW signaled across AS boundary



PacNOG15

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

BUILT FOR
THE HUMAN
NETWORK



Cisco *live!*

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public