AOS-CX 10.10 Update

aruba
a Hewlett Packard Enterprise company

LDP (Label Distribution Protocol)

GR (Graceful Restart)

Presenters
- Daryl Wan
- Yash NN

Agenda

1. Overview
2. Use Cases
3. Details and Caveats
4. Configuration
5. Best Practices
6. Troubleshooting
7. Demo
8. Additional Resources

Overview

LDP GR Overview

- 10.10 adds LDP GR support for MPLS deployments
- Preserves data plane forwarding during control plane restart
- Supports both GR restarting LSR (LSR that performs GR) or GR helper (LSR that helps GR restarter complete GR)

- Supported platform:
  - 8360

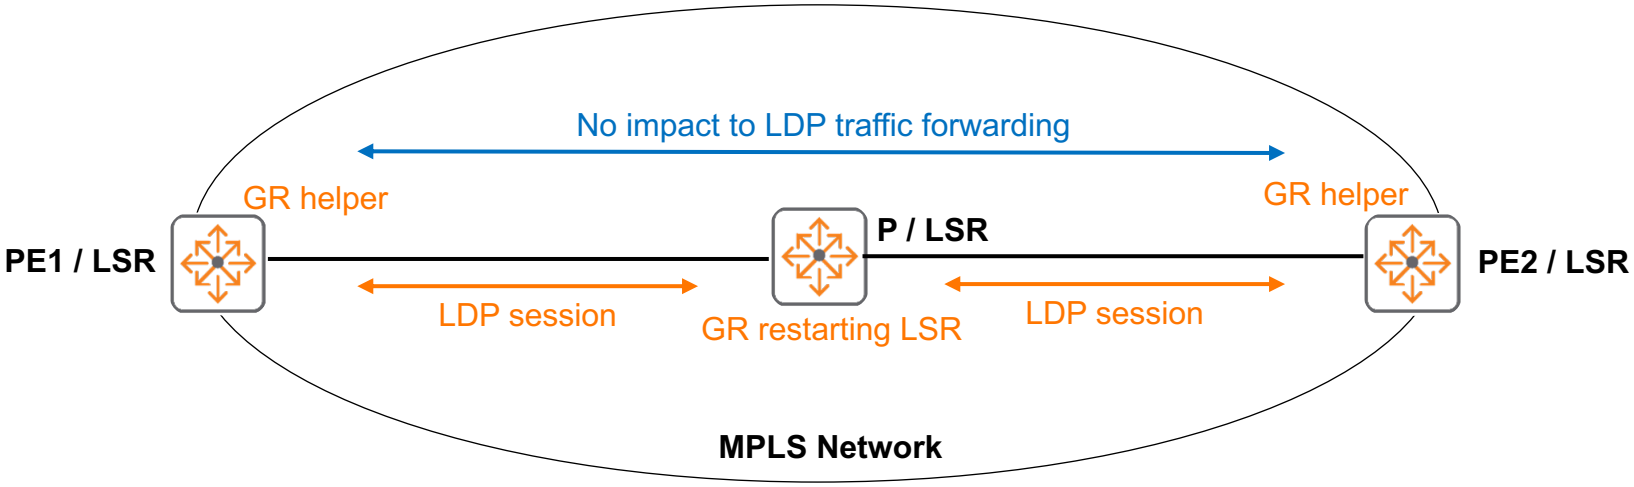

Example of LDP control plane restart on P

LDP GR Overview

- 10.10 adds LDP GR support for MPLS deployments
- Preserves data plane forwarding during control plane restart
- Supports both GR restarting LSR (LSR that performs GR) or GR helper (LSR that helps GR restarter complete GR)

- Supported platform:
  - 8360

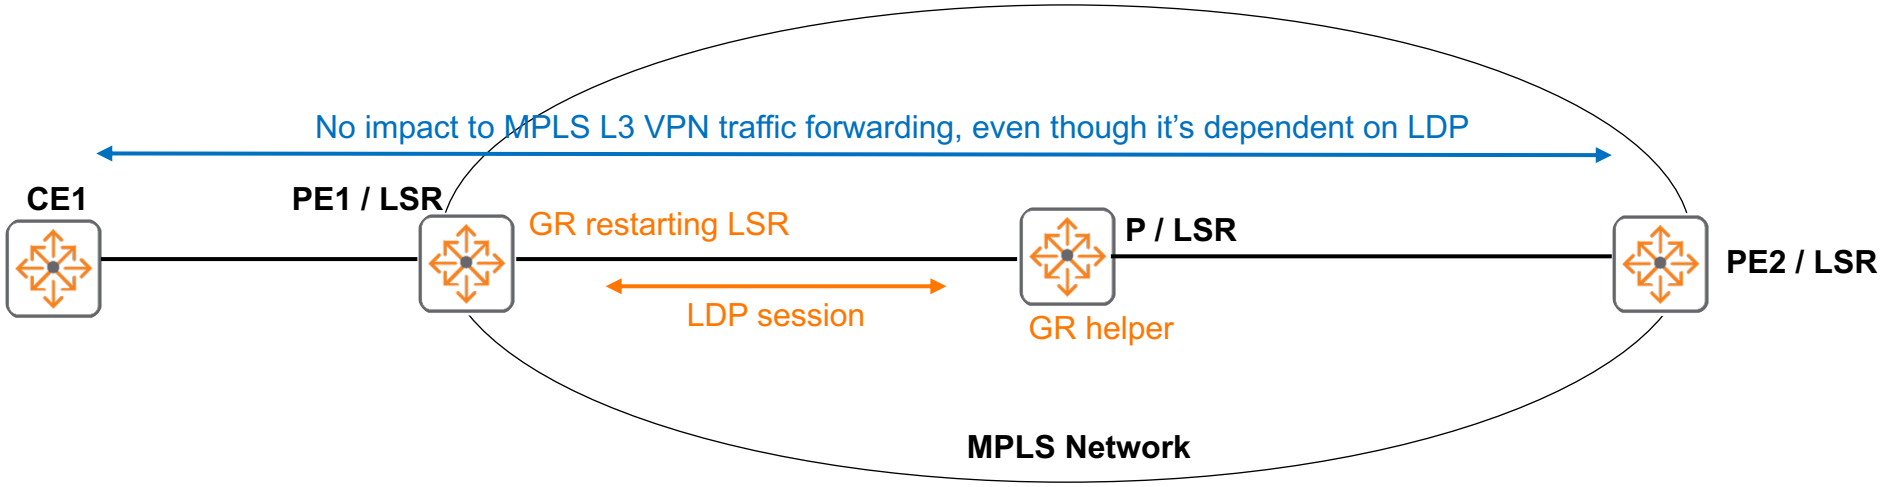

Example of LDP control plane restart on PE

Use Cases

Control plane restart - Due to routing daemon

- LDP traffic would normally be impacted by routing daemon restart
- LDP GR allows LDP traffic forwarding during routing daemon restart

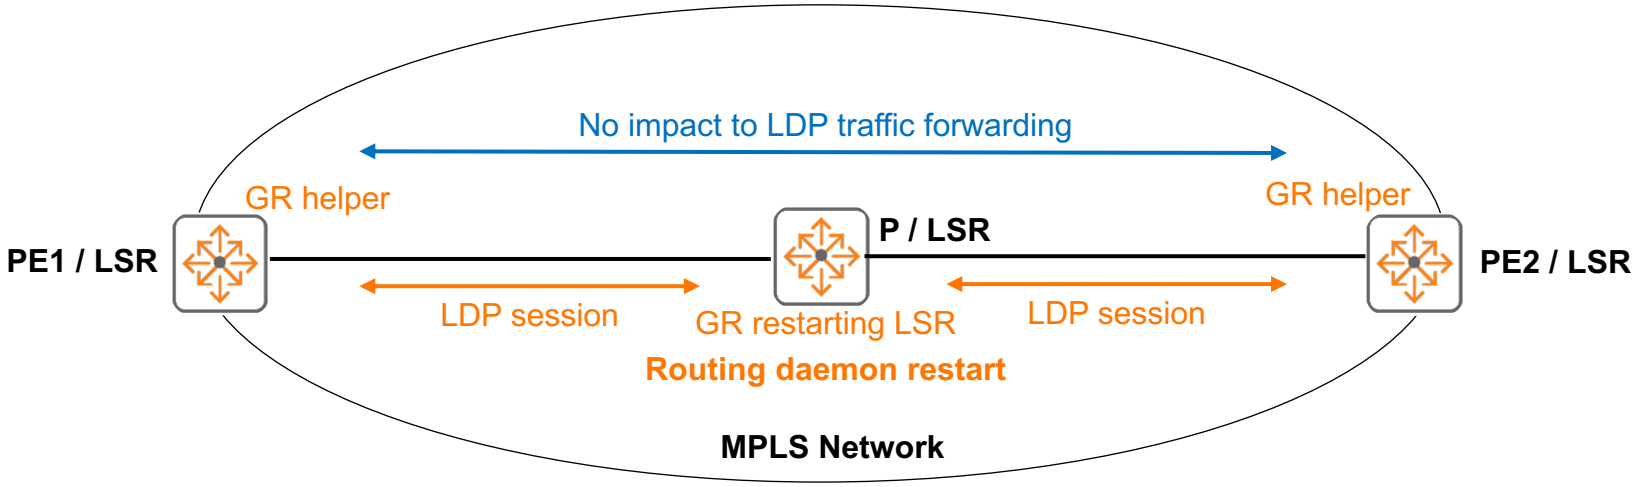

Example of routing daemon restart on P

Control plane restart - Due to management module failover

- Not applicable for 10.10 since 8360 is a fixed switch, will be applicable to modular switches in future
- LDP traffic would normally be impacted by redundant management module failover
- LDP GR allows LDP traffic forwarding during redundant management module failover

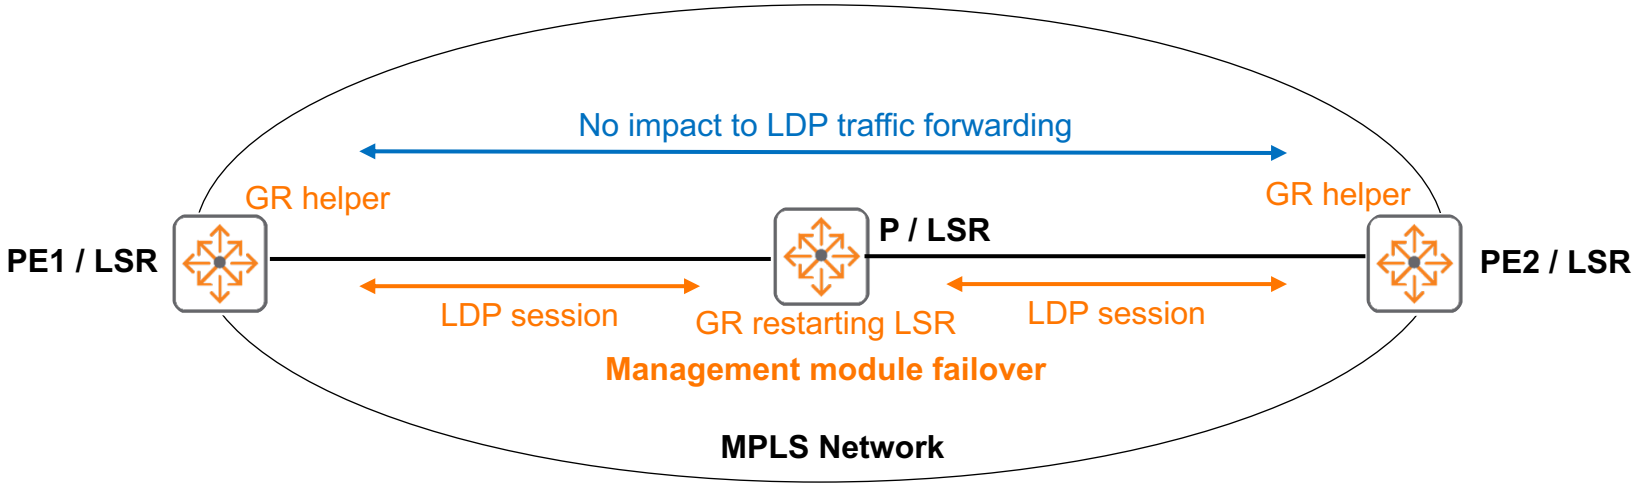

Example of management module failover on P

Details

LDP GR Details

- Feature is enabled by default
  - Cannot be disabled
- Based on RFC 3478, Graceful Restart Mechanism for LDP
- Requires LDP neighbor to also support LDP GR
- Example of LDP neighbor session without LDP GR

```
PE2# sh mpls ldp neighbor
Local LDP Identifier: 2.2.2.2:0, Peer LDP Identifier: 3.3.3.3:0
    Graceful Restart: No
    Peer Reconnect Time: 0 sec
    Peer Recovery Time: 0 sec
    Session Holdtime: 40 sec
    Up time: 00:06:03
    State: operational
    LDP Discovery Sources: 1/1/15
    Addresses bound to this peer:
         3.3.3.3 30.1.0.1 60.1.0.1
```

- Example of LDP neighbor session with LDP GR

```
PE2# sh mpls ldp neighbor
Local LDP Identifier: 2.2.2.2:0, Peer LDP Identifier: 3.3.3.3:0
    Graceful Restart: Yes
    Peer Reconnect Time: 120 sec
    Peer Recovery Time: 450 sec
    Session Holdtime: 40 sec
    Up time: 01:16:12
    State: operational
    LDP Discovery Sources: 1/1/15
    Addresses bound to this peer:
         3.3.3.3 30.1.0.1 60.1.0.1
```

LDP GR Details

– Requires GR flag to be set to 1 during LDP session establishment

```
PE2# debug mpls ldp

PE2# sh debug buffer reverse
! snip
2022-03-10:18:40:54.741100|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|LDP SC 1 Session has become operational,
local LDP ID 02020202 0000, LDP Entity index 0X00000007, peer LDP ID 03030303 0000.

2022-03-10:18:40:54.739908|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|P2MP MLDP Present           = 0
2022-03-10:18:40:54.739892|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Recovery Time              = 0
2022-03-10:18:40:54.739876|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Reconnect Timeout          = 120000
2022-03-10:18:40:54.739860|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Graceful Restart Flags     = 0X0001
2022-03-10:18:40:54.739844|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Graceful Restart(FT)Present = 1
2022-03-10:18:40:54.739828|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Max PDU length              = 1440
2022-03-10:18:40:54.739812|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|PV limit                   = 0
2022-03-10:18:40:54.739795|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Loop detection            = 0
2022-03-10:18:40:54.739779|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Advertisement discipline  = 1
2022-03-10:18:40:54.739763|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Keepalive time            = 40
2022-03-10:18:40:54.739746|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Session takes active role = False

2022-03-10:18:40:54.739729|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|LDP SC 1 sent a Session Initialization
msg, local LDP ID 02020202 0000, LDP Entity index 7, peer LDP ID 03030303 0000.
```

– FT refers to Fault Tolerant TLV

LDP GR Details

– Packet capture of LDP initialization message

```
    80 62.492971400  3.3.3.3              2.2.2.2               LDP       118 Initialization Message

> Frame 80: 118 bytes on wire (944 bits), 118 bytes captured (944 bits) on interface MirrorRxNet, id 0
> Ethernet II, Src: ArubaaHe_42:c7:00 (8c:85:c1:42:c7:00), Dst: HewlettP_54:4c:60 (b4:99:ba:54:4c:60)
> Internet Protocol Version 4, Src: 3.3.3.3, Dst: 2.2.2.2
> Transmission Control Protocol, Src Port: 54739, Dst Port: 646, Seq: 1, Ack: 1, Len: 52
v Label Distribution Protocol
    Version: 1
    PDU Length: 48
    LSR ID: 3.3.3.3
    Label Space ID: 0
  v Initialization Message
      0... .... = U bit: Unknown bit not set
      Message Type: Initialization Message (0x200)
      Message Length: 38
      Message ID: 0x00000001
    > Common Session Parameters
    v FT Session TLV
        10.. .... = TLV Unknown bits: Unknown TLV, do not Forward (0x2)
        TLV Type: FT Session TLV (0x503)
        TLV Length: 12
      v FT Session Parameters
        v Flags: 0x0001 (Using LDP Graceful Restart)
            0... .... .... .... = R bit: LSR has not preserved state and resources for all FT-Labels
            .000 0000 0000 .... = Reserved: 0x000
            .... .... .... 0... = S bit: FT Protection TLV not supported on other than KeepAlive
            .... .... .... .0.. = A bit: May treat some labels as FT and others as non-FT
            .... .... .... ..0. = C bit: Check-Pointing procedures not in use
            .... .... .... ...1 = L bit: Re-learn the state from the network
          Reserved: 0x0000
          Reconnect Timeout: 200000
          Recovery Time: 300000
```

LDR GR Process

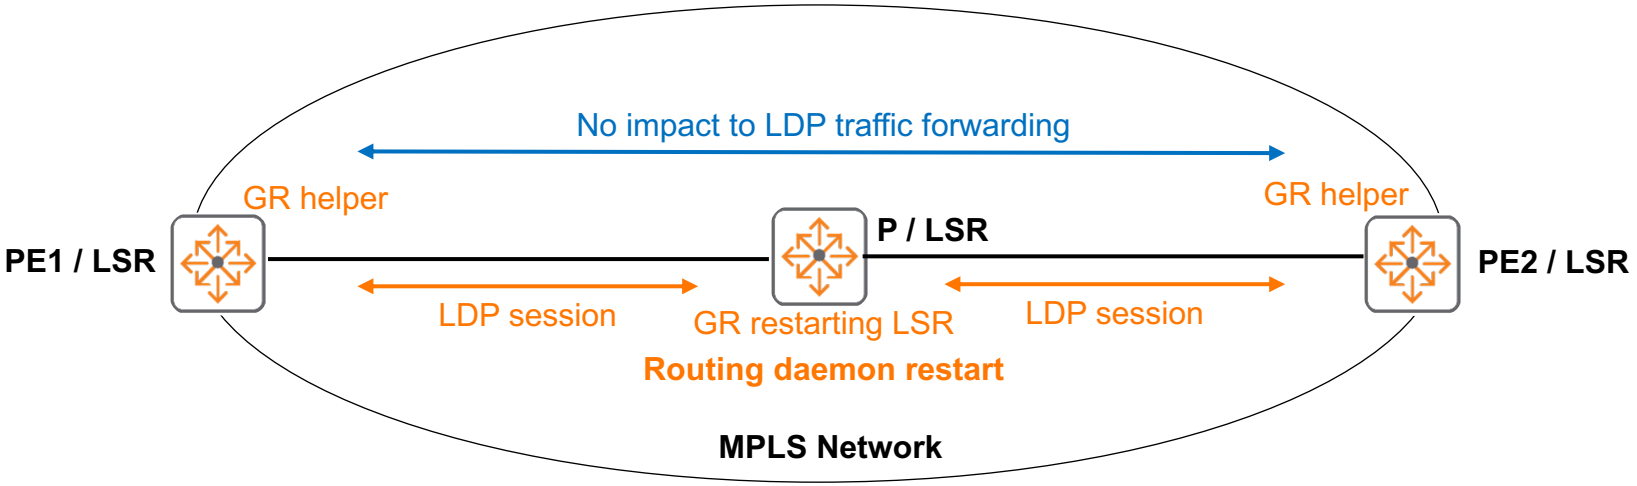

– PE1 and PE2 detects service interruption with P

```
2022-03-10:22:53:13.083572|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|LDP SC 1 Session has ceased to be
operational, local LDP ID 02020202 0000, LDP Entity index 0X0000000A, peer LDP ID 03030303 0000.
```

– PE1 and PE2 marks all the label bindings from P as stale, continues to use these bindings for MPLS forwarding, no impact to LDP traffic forwarding

– PE1 and PE2 tries to re-establish LDP session with P

– PE1 and PE2 will delete all the stale LDP bindings associated with P when “neighbor-liveness/peer reconnect” countdown ends

LDP GR Caveats

- GR timers across the protocols need to meet the condition that OSPF < LDP < BGP
- Current 10.10 defaults
  - OSPF GR (120s) < LDP GR (forward hold interval 150s) < BGP GR (120s)
- In 10.10, if BGP GR is enabled, the BGP GR timer below will need to be modified to 180s in order for LDP GR to function correctly

```
PE1(config-bgp)# bgp graceful-restart
  restart-time     Set the max time that the BGP Speaker requires to restart and
                   establish BGP session with peers
```

- Plan is for above BGP GR default timer to be modified to 180s in future software release

Configuration

Configuration Info

- Feature is enabled by default, no configuration required
  - Cannot be disabled

- Timers can be modified if desired

```
PE2(config)# mpls
PE2(config-mpls)# label-protocol ldp
PE2(config-mpls-ldp)# graceful-restart-timers
  forwarding-holding  Configure the forwarding holding time
  max-recovery        Configure the max recovery time
  neighbor-liveness   Configure the neighbor liveness time
```

- Forwarding-holding
  - The amount of time (seconds) that the MPLS forwarding state should be preserved after the control plane restarts (Default 150)

- Max-recovery
  - The amount of time (seconds) that the stale label bindings should be kept on the router after the LDP session has been re-established (Default 120)

- Neighbor-liveness
  - The amount of time (seconds) that the router will wait for the LDP session to be re-established. If the router cannot re-establish the LDP session within that time, the router deletes all the stale LDP bindings received from that LDP neighbor (Default 120).

LDP GR Timers

- Timers do not need to match on LDP neighbors
- Timers are exchanged between neighbors and they each pick the lower value and use them for their sessions during GR
- The GR status and timers shown for the neighbor are the received values

```
PE2# sh mpls ldp neighbor
Local LDP Identifier: 2.2.2.2:0, Peer LDP Identifier: 3.3.3.3:0
    Graceful Restart: Yes
    Peer Reconnect Time: 200 sec
    Peer Recovery Time: 300 sec
    Session Holdtime: 40 sec
    Up time: 00:10:47
    State: operational
    LDP Discovery Sources: 1/1/15
    Addresses bound to this peer:
         3.3.3.3 30.1.0.1 60.1.0.1

PE2# sh mpls ldp graceful-restart
Max recovery time                : 120 sec
Neighbor liveness time           : 120 sec
Forwarding holding time          : 600 sec
Number of graceful restart events: 1
Graceful restart in progress     : false
Forwarding holding time remaining: 0 sec
Current graceful restart status  : complete
```

```
P3# sh mpls ldp neighbor
Local LDP Identifier: 3.3.3.3:0, Peer LDP Identifier: 2.2.2.2:0
    Graceful Restart: Yes
    Peer Reconnect Time: 120 sec
    Peer Recovery Time: 600 sec
    Session Holdtime: 40 sec
    Up time: 00:10:29
    State: operational
    LDP Discovery Sources: 1/1/47
    Addresses bound to this peer:
         2.2.2.2 30.1.0.2

P3# sh mpls ldp graceful-restart
Max recovery time                : 200 sec
Neighbor liveness time           : 200 sec
Forwarding holding time          : 300 sec
Number of graceful restart events: 2
Graceful restart in progress     : false
Forwarding holding time remaining: 0 sec
Current graceful restart status  : complete
```

- Peer reconnect time = Neighbor liveness timer configured in the peer device.
- Peer recovery time = Forwarding holding time configured in the peer minus the amount of time the control plane has been up

Best Practices

Best Practices

- Leave timers to default unless there is a reason to modify
- Changing a timer will restart LDP sessions and may impact traffic forwarding
- Since LDP sessions are restarted, this is a GR event

```
P3# sh mpls ldp graceful-restart
Max recovery time                : 120 sec
Neighbor liveness time           : 120 sec
Forwarding holding time          : 600 sec
Number of graceful restart events: 1
Graceful restart in progress     : false
Forwarding holding time remaining: 0 sec
Current graceful restart status  : complete

P3# conf
P3(config)# mpls
P3(config-mpls)# label-protocol ldp
P3(config-mpls-ldp)# graceful-restart-timers max-recovery 200
Changing the timer value will restart any LDP sessions.
This may result in traffic loss.

Continue (y/n)? y
```

Best Practices

– As LDP sessions are restarted, GR status “in-progress” will be shown

```
P3(config)# sh mpls ldp graceful-restart
Max recovery time                : 200 sec
Neighbor liveness time           : 200 sec
Forwarding holding time          : 300 sec
Number of graceful restart events: 1
Graceful restart in progress     : true
Forwarding holding time remaining: 88 sec
Current graceful restart status  : in-progress
```

– Once “forwarding-holding-time” is over, GR status “complete” will be shown

```
P3(config)# sh mpls ldp graceful-restart
Max recovery time                : 200 sec
Neighbor liveness time           : 200 sec
Forwarding holding time          : 300 sec
Number of graceful restart events: 2
Graceful restart in progress     : false
Forwarding holding time remaining: 0 sec
Current graceful restart status  : complete
```

Troubleshooting

LDP GR Troubleshooting

- Have a topology diagram ready
- Ensure IPs, interface details are included
- Check physical cabling and generate “show tech” when opening a TAC case
- Check network: show LLDP neighbor, ensure directly connected neighbors have connectivity by using ping and traceroute between loopbacks and interfaces, fix any issues found

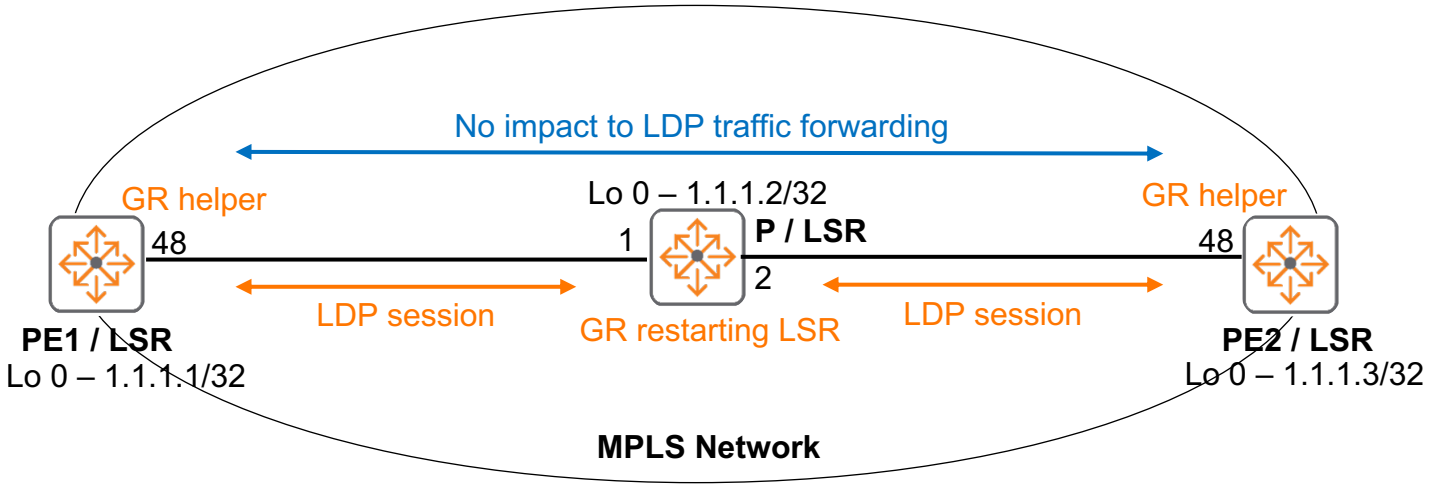

- Recommended troubleshooting flow

1. Check LDP neighbors support LDP GR capability
2. Verify LDP neighbors peer with LDP GR enabled
3. Enable MPLS LDP debugs
4. Verify data plane forwarding during control plane restart

1. Check LDP neighbors support LDP GR capability

– Check required software is loaded, e.g. 10.10

2. Verify LDP neighbors peer with LDP GR enabled

– Ensure peer shows GR: Yes

```
PE2# sh mpls ldp neighbor
Local LDP Identifier: 2.2.2.2:0, Peer LDP Identifier: 3.3.3.3:0
    Graceful Restart: Yes
    Peer Reconnect Time: 200 sec
    Peer Recovery Time: 300 sec
    Session Holdtime: 40 sec
    Up time: 00:10:47
    State: operational
    LDP Discovery Sources: 1/1/15
    Addresses bound to this peer:
         3.3.3.3 30.1.0.1 60.1.0.1
```

3. Enable MPLS LDP debugs

– If required, enable MPLS LDP debugs to check further

```
PE2# debug mpls ldp

PE2# sh debug buffer reverse
! snip
2022-03-10:18:40:54.741100|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|LDP SC 1 Session has become operational,
local LDP ID 02020202 0000, LDP Entity index 0X00000007, peer LDP ID 03030303 0000.

2022-03-10:18:40:54.739908|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|P2MP MLDP Present           = 0
2022-03-10:18:40:54.739892|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Recovery Time              = 0
2022-03-10:18:40:54.739876|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Reconnect Timeout          = 120000
2022-03-10:18:40:54.739860|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Graceful Restart Flags     = 0X0001
2022-03-10:18:40:54.739844|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Graceful Restart(FT)Present = 1
2022-03-10:18:40:54.739828|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Max PDU length              = 1440
2022-03-10:18:40:54.739812|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|PV limit                   = 0
2022-03-10:18:40:54.739795|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Loop detection            = 0
2022-03-10:18:40:54.739779|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Advertisement discipline  = 1
2022-03-10:18:40:54.739763|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Keepalive time            = 40
2022-03-10:18:40:54.739746|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|Session takes active role = False

2022-03-10:18:40:54.739729|hpe-routing|LOG_INFO|AMM|-|MPLS|MPLS_LDP|LDP SC 1 sent a Session Initialization
msg, local LDP ID 02020202 0000, LDP Entity index 7, peer LDP ID 03030303 0000.


PE2# no debug all
```

– Ensure FT TLV is exchanged

4. Verify data plane forwarding during control plane restart

- Generate data plane traffic and verify traffic is not impacted due to control plane restart
  - From routing daemon restart
  - From management module failover (future)

Demo

MPLS LDP GR

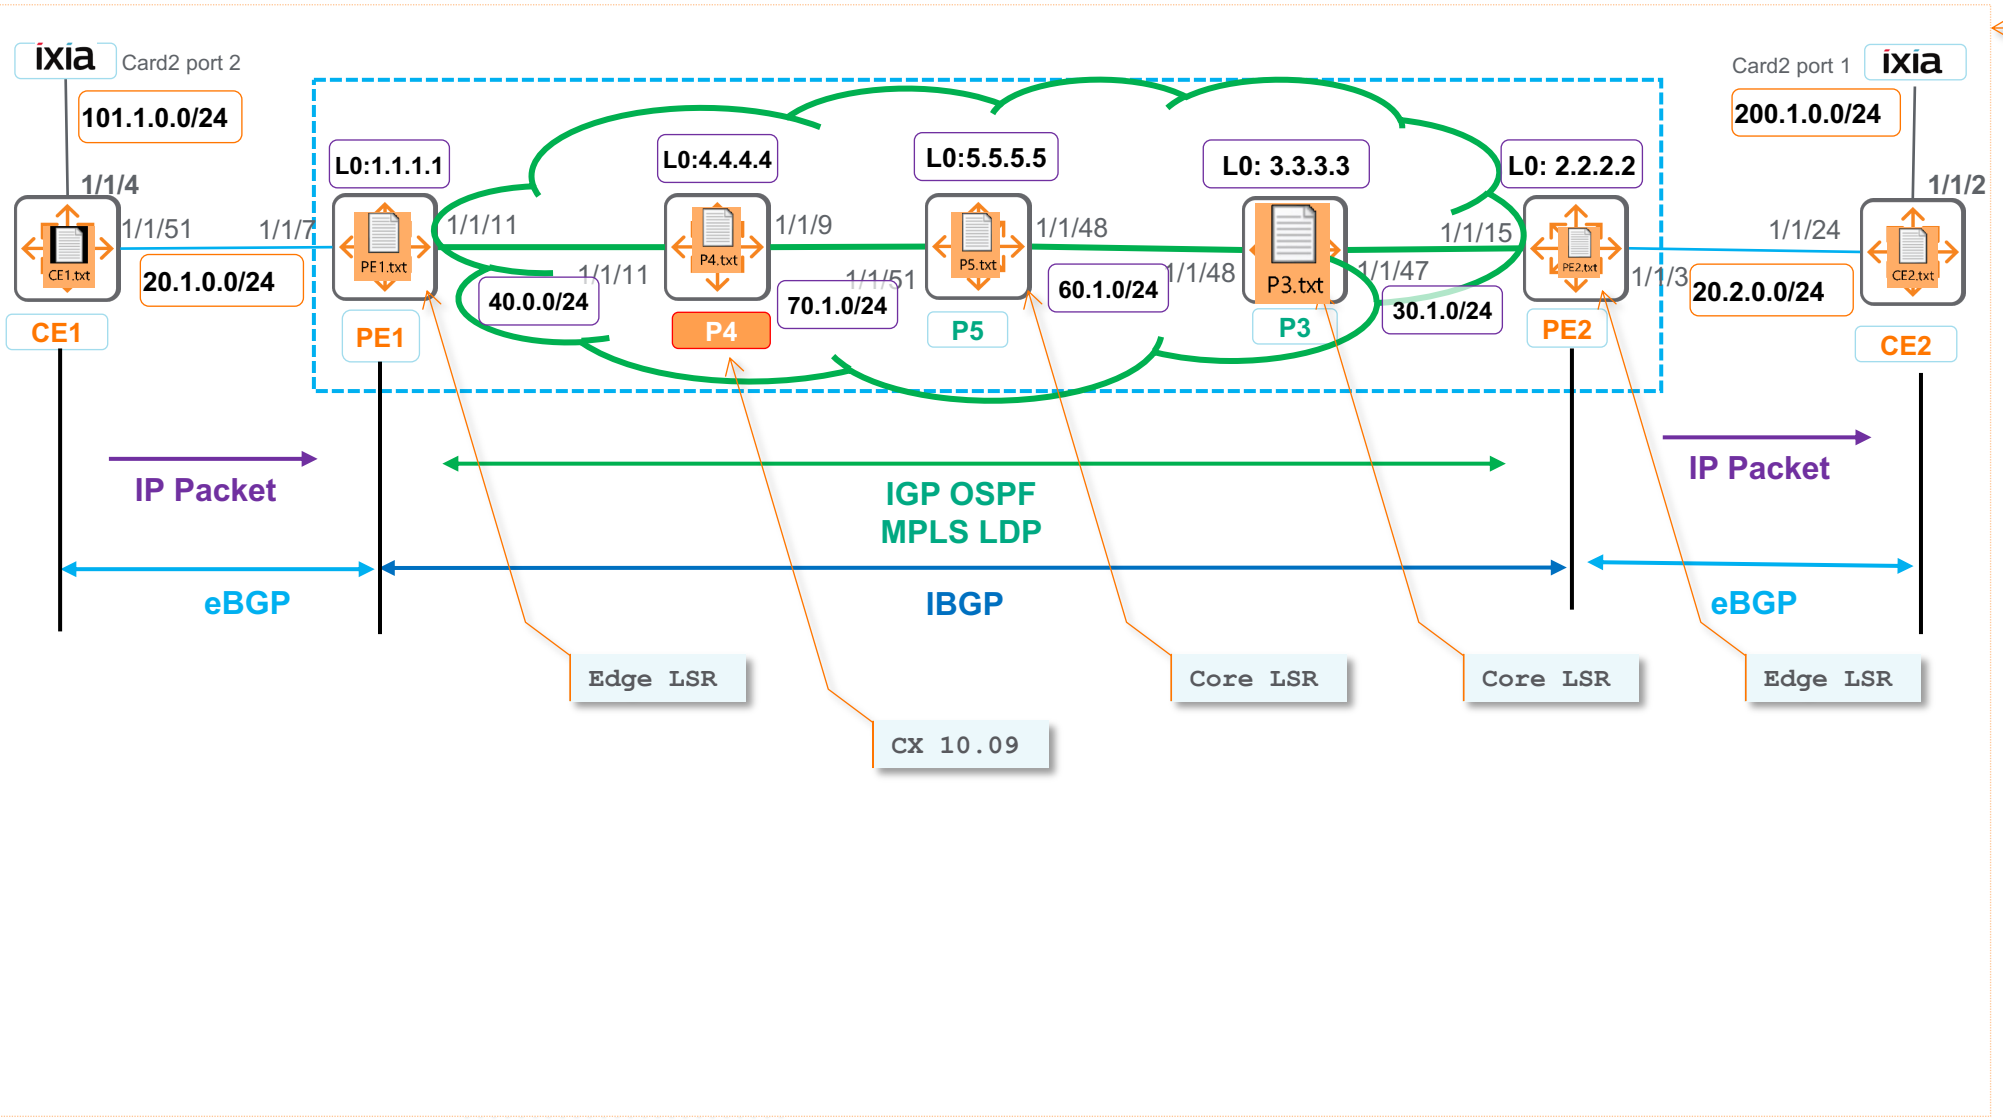

Step1:
Configuration
Step2:
OSPF / unicast routing
Step3:
LDP neighbor
Step4:
Labels for all LSRs & labels for remote VPNv4
Step5:
MP-BGP VPNv4 peering
Step6:
VRF routing table
Step7:
PE and CE Connectivity
Step8:
Ping CE's

Resources

Feature/Solution References

– https://datatracker.ietf.org/doc/html/rfc3478

aruba
a Hewlett Packard Enterprise company

Thank you

daryl.wan@hpe.com
yashavantha.n.n@hpe.com