Multiprotocol Label Switching Nonstop Forwarding with Stateful Switchover Technical Presentation

ITD Product Management

Agenda

- Introduction to Cisco IOS High Availability
- Multiprotocol Label Switching (MPLS) Non-Stop Forwarding (NSF)/Stateful Switchover (SSO) Overview
- MPLS NSF/SSO Components
  - MPLS NSF/SSO —LDP NSF/SSO
  - MPLS NSF/SSO—BGP VPNv4 NSF/SSO
  - MPLS NSF/SSO—Management
- Summary

Cisco IOS High Availability Strategy Based on Customer Needs

- **Overarching requirement is to provide continuous access to applications, data, and content from anywhere and anytime**
- **Nonstop application delivery**
  - **End-to-end**
  - **Systems approach**
  - **Target every potential cause of downtime with functionality, design, or best practice to mitigate the impact**

Cisco Approach

Understand

- Research and development
- IP network expertise, experience
- Listening and responding to customer requests
- Standards development / participation
- Field interaction
- Customer partnership

Identify

- Identify service
  - Affecting conditions
    - Hardware
    - Software
    - Traffic and protocol
    - Performance
    - Management
    - Security
- Leverage broad experience, deployment history, and analysis
- Industry experiences shared

Design & Engineer

- Focused solutions
  - Targeted to specific, known causes of downtime
- Engineered solutions
  - Architected for long term, infrastructure reuse
- Service related, rather than platform related
- Considers network topology and design - Eliminate single points of failure
- Automation and embedded intelligence

Best Practice

- Test solution for compliance with objectives
- Institute best practices and designs using high availability solutions
- Continual process
- Evolutionary

Where is the Exposure?
Most Common Causes of Network Outages

- **Network and software applications**
  - Hardware failures
  - Software failures
  - Link failures
  - Power/environment failures
  - Resource utilization issues
- **Operational processes**
  - Network design issues
  - Lack of standards (hardware, software, configurations)
  - No fault management or capacity planning
  - Inadequate change of management, documentation, and staff training
  - Lack of ongoing event correlation
  - Lack of timely access to experts/knowledge

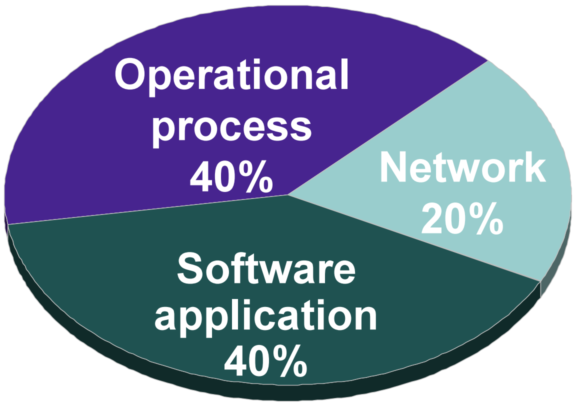

Sources of network downtime *

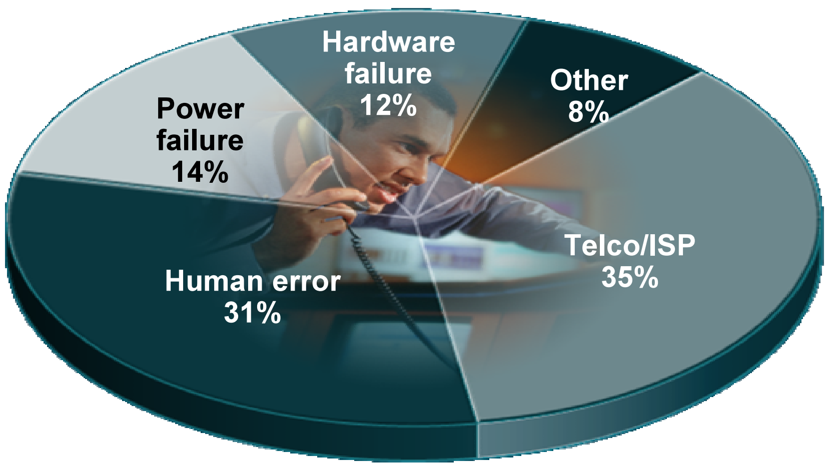

Common causes of Enterprise network downtime **

Pie graph source*: * Gartner Group ** Yankee Group The Road to a Five-Nines Network* 2/2004

Systematic, End to End Approach Targeting Downtime

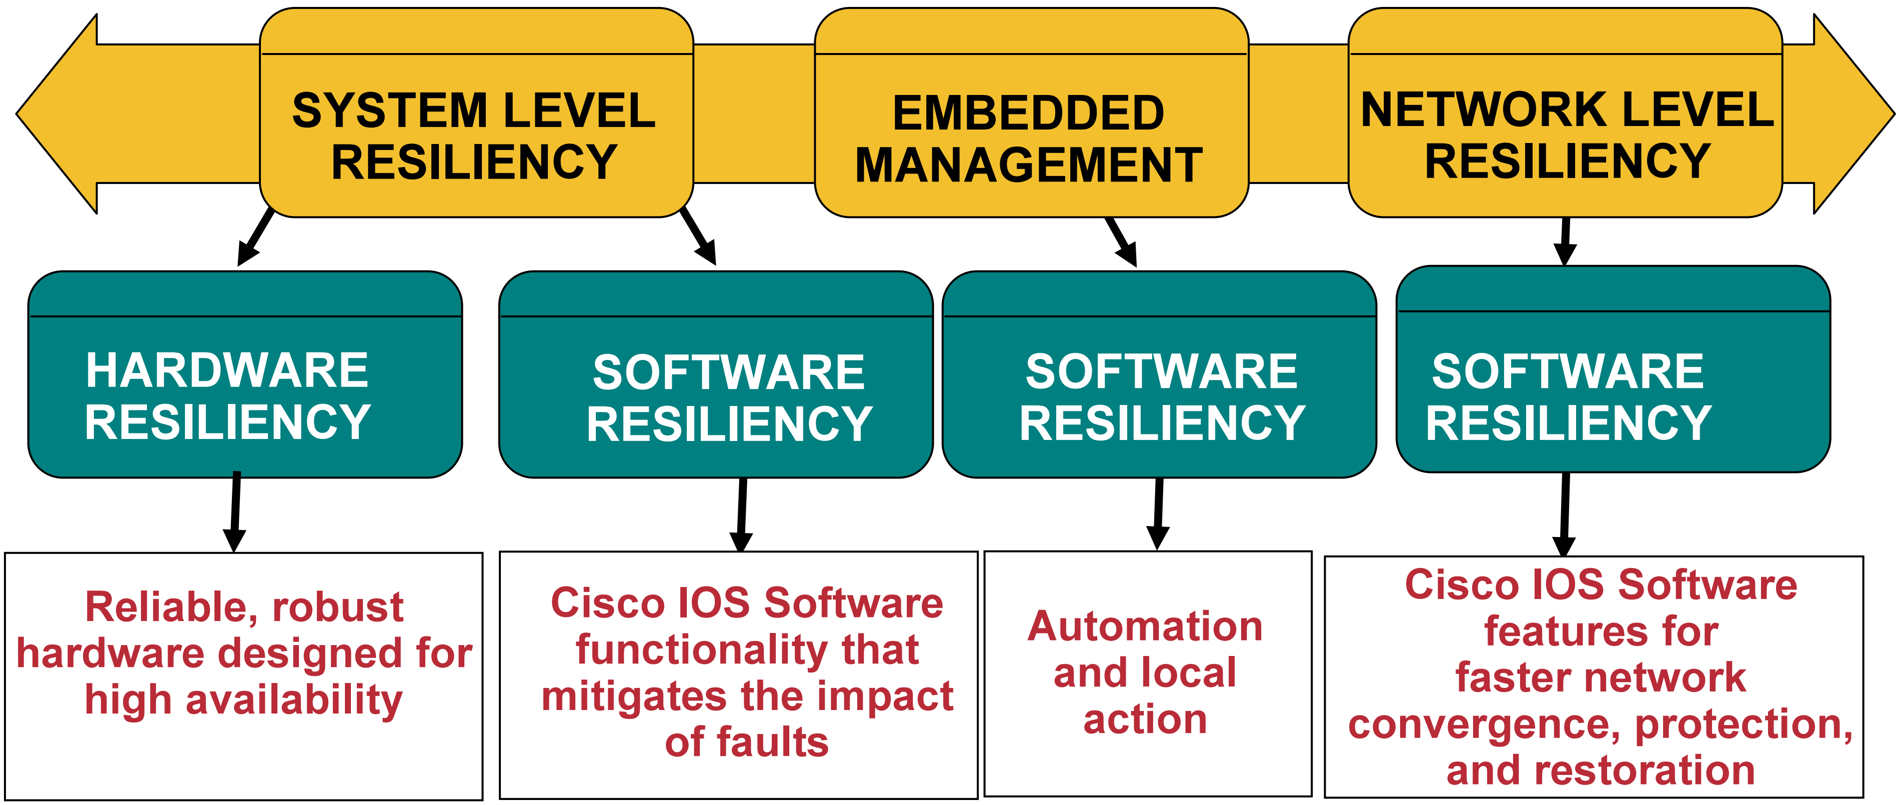

INVESTMENT PROTECTION IS A KEY COMPONENT

Mitigating the Exposure
Targeting Downtime

Most common causes of downtime

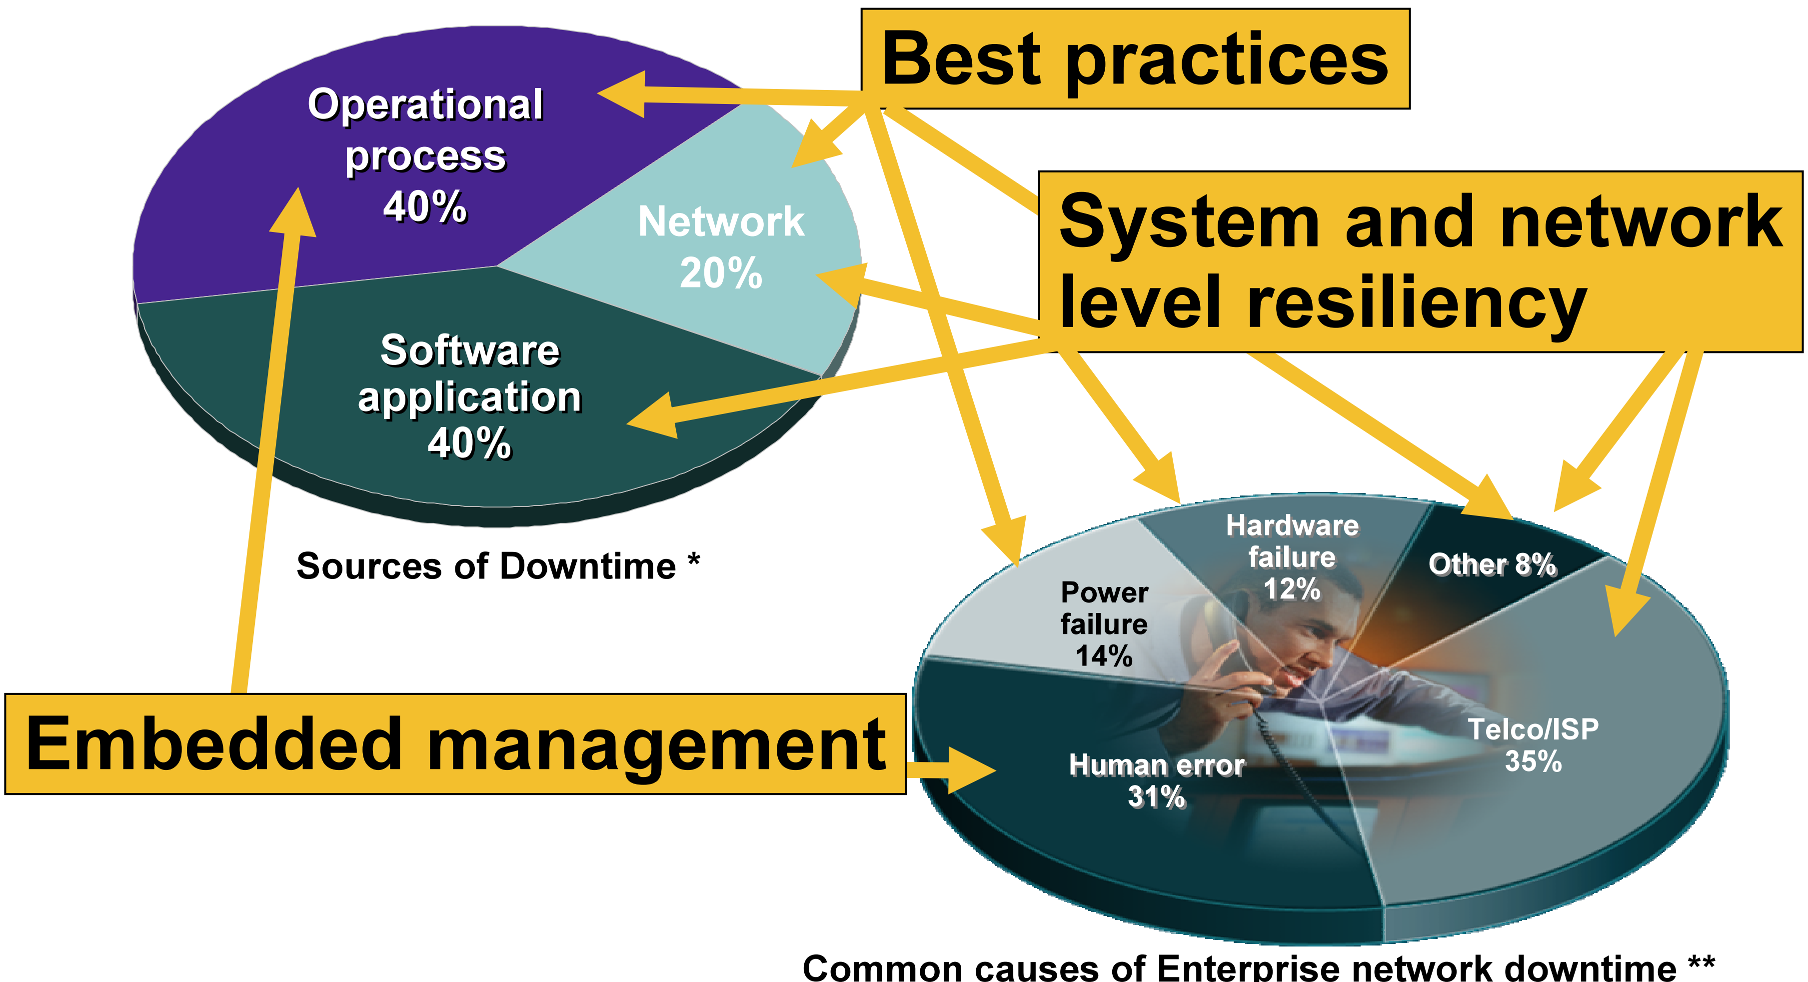

Cisco IOS High Availability Focused Approach

- Start with resilient hardware foundation
  - ↑ MTBF
  - Eliminate single points of failure
- **System level resiliency** at critical network edges
  - ↓ MTTR for system failures including route processors, line cards and software component failures
  - Reduce outage for both dual and single route processor systems
  - Mitigate planned outages by providing in-service software upgrades
  - Protect remote devices while creating more efficient and cost-effective resiliency
- **Network level resiliency** in the core and where redundant paths exist
  - Deliver industry-leading features for fast network convergence, protection & restoration
- Account for services and protocols
- **Embedded management** & automation
  - Embed intelligent event management for proactive maintenance
  - Automation and configuration management to reduce human errors

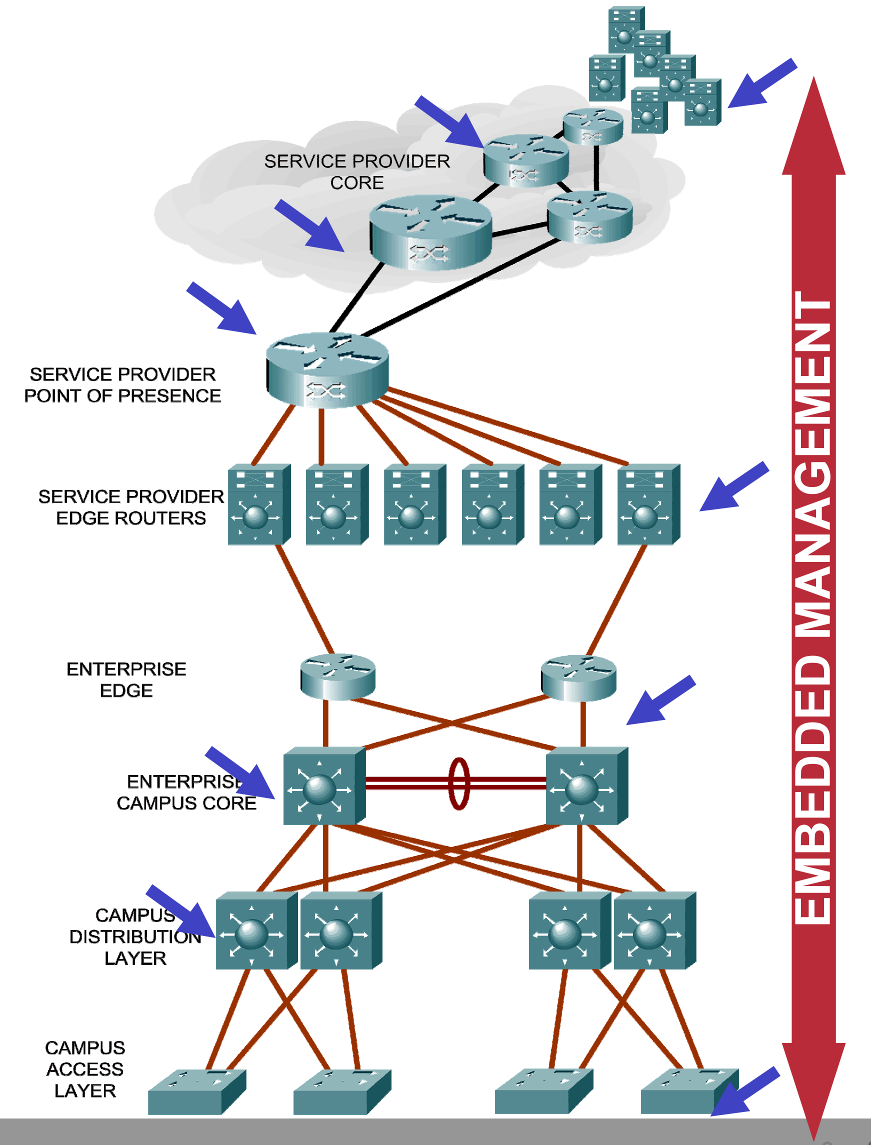

System Level Resiliency Overview
Reduce Mean Time to Repair

Eliminate single points of failure for hardware and software components

- Control/data plane resiliency
  - Separation of control and forwarding plane
  - Control plane handles signaling and network awareness
  - Forwarding/data plane rapidly switches packets
  - Seamless restoration of route processor control and data plane failures
- Link resiliency
  - Reduced impact of line card hardware and software failures
- Planned outages
  - Seamless software and hardware upgrades
- Fault isolation and containment
  - Process independence
  - Target single processor systems

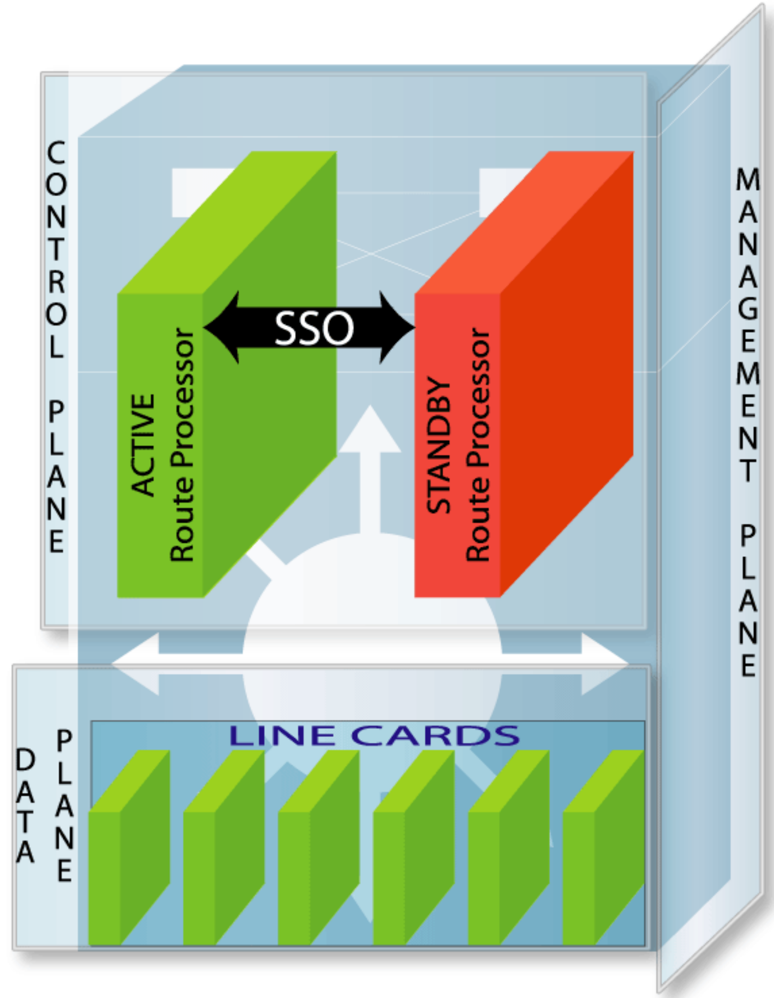

Agenda

- Introduction to Cisco IOS High Availability
- MPLS NSF/SSO Overview
- MPLS NSF/SSO Components

  MPLS NSF/SSO —LDP NSF/SSO

  MPLS NSF/SSO—BGP VPNv4 NSF/SSO

  MPLS NSF/SSO—Management
- Summary

System Level Resiliency
MPLS NSF/SSO

MPLS NSF/SSO = Dual RPs + Local Checkpoint + Graceful Restart + IP HA

Enhance the key protocols used in MPLS control plane to minimize the disruption in MPLS forwarding plane

IP HA

- BGP
- OSPF
- EIGRP
- IS-IS

MPLS HA

- LDP
- MP-BGP

NSF/SSO Architecture

- 1. RPs initialize, assign roles
- (Active/Standby),
  negotiate mode (SSO),
  begin
  checkpointing state

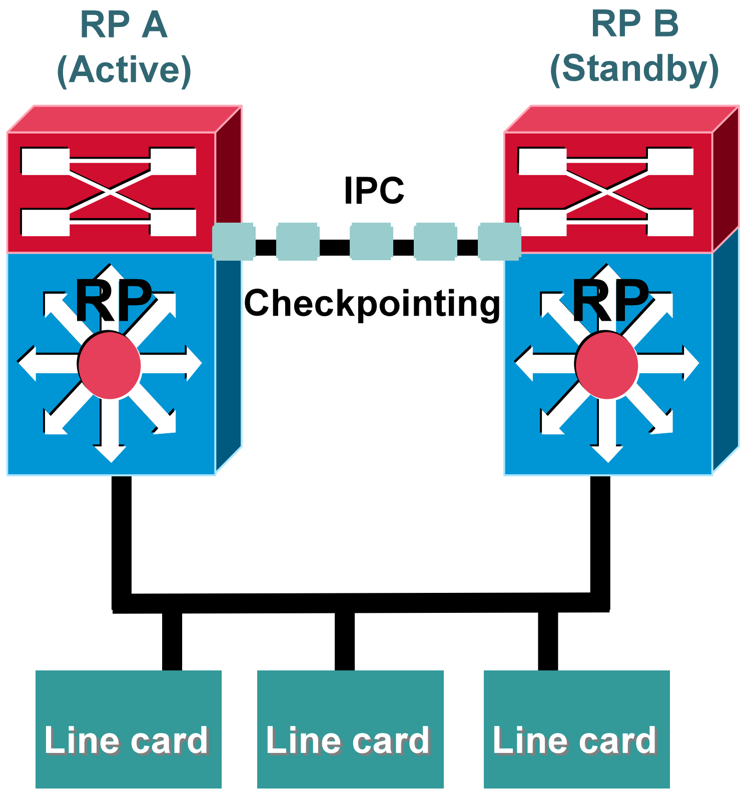

NSF/SSO Architecture

1. RPs initialize, assign roles (Active/Standby), negotiate mode (SSO), begin checkpointing state

2. L2/L3 services provided by Active Forwarding done directly via line cards/ forwarding ASICs

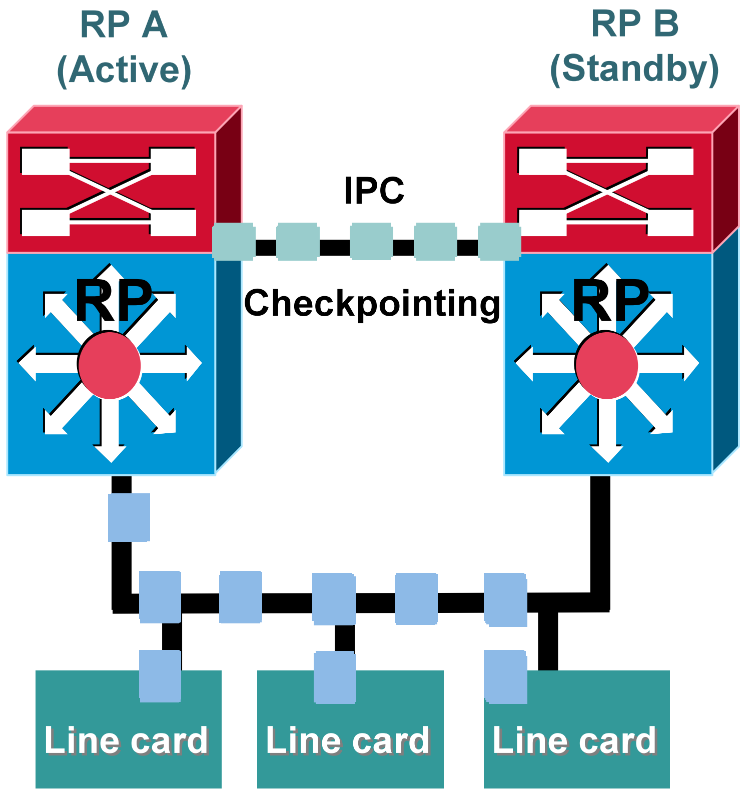

NSF/SSO Architecture

1. RPs initialize, assign roles (Active/Standby),
   negotiate mode (SSO), begin checkpointing state
2. L2/L3 services provided by Active
   Forwarding done directly via line cards/ forwarding ASICs
3. Active RP fails
   Switchover starts, checkpointing stops
   Forwarding continues on LCs/FP;
   RP B assumes Active role and begins providing L2, L3 services;
   L2 continues where it left off;
   L3 reconverges, updates RIB then FIB

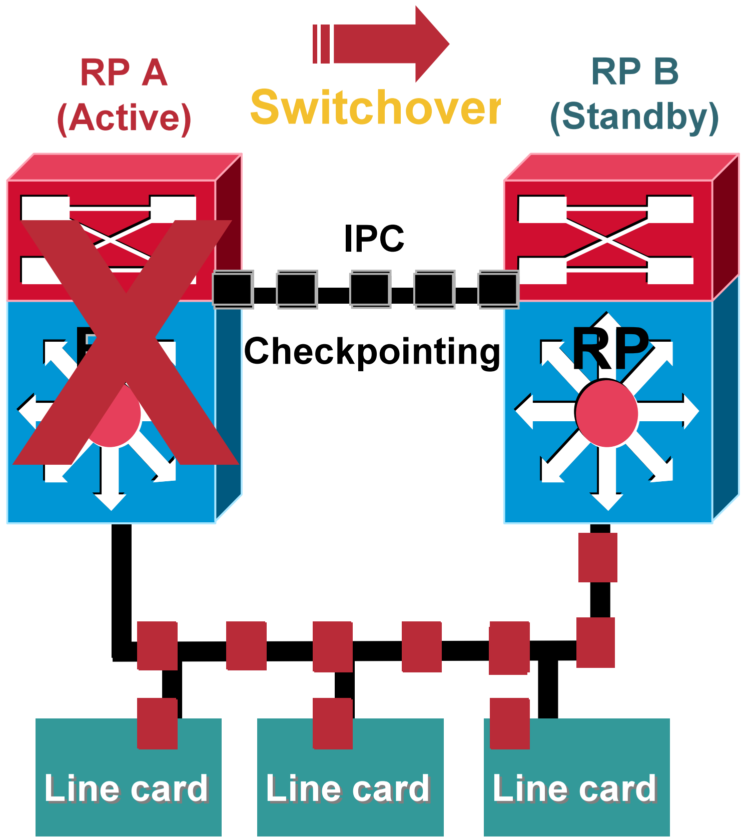

NSF/SSO Architecture

1. RPs initialize, assign roles (Active/Standby), negotiate mode (SSO), begin checkpointing state

2. L2/L3 services provided by Active Forwarding done directly via line cards/ forwarding ASICs.

3. Active RP fails
   Switchover starts, checkpointing stops
   Forwarding continues on LCs/FP;
   RP B assumes Active role and begins providing L2, L3 services;
   L2 continues where it left off;
   L3 reconverges, updates RIB then FIB

4. **RP A reloads, reboots, reinitializes and rejoins as Standby;**
   Checkpointing resumes from Active to Standby

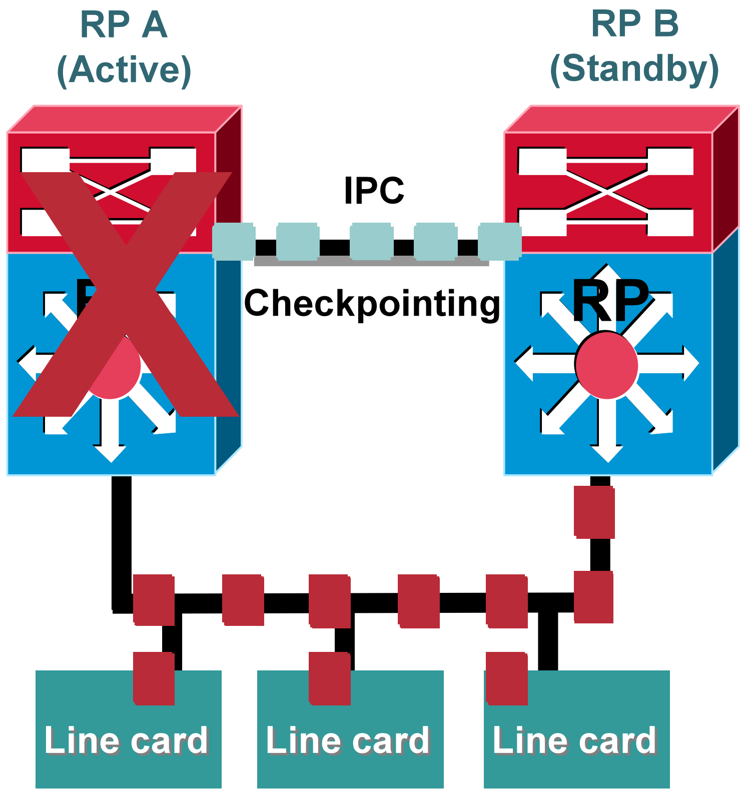

Agenda

- Introduction to Cisco IOS High Availability
- MPLS NSF/SSO Overview
- MPLS NSF/SSO Components

  MPLS NSF/SSO —LDP NSF/SSO

  MPLS NSF/SSO—BGP VPNv4 NSF/SSO

  MPLS NSF/SSO—Management
- Summary

LDP HA Key Elements

- **Checkpointing local label bindings**
  - **On devices with route processor redundancy**
- **LDP graceful restart capability**
  - **On participating PEs, RRs, and P routers**

Why LDP Checkpointing?

- The LSRs that support LDP GR, require that the restarting control plane be able to query the forwarding plane for local label binding information; the MPLS forwarding plane will not support such queries
- LDP checkpointing makes local label bindings available to the restarting (standby) LDP control plane

MPLS LDP Local Checkpointing Key Points

- The checkpointing function is enabled by default and done using a separate process
- The checkpointing function copies active RP's LDP local label bindings to the backup RP; for the first round, all the labels are copied from active to back up RP
- Periodic incremental updates are done to reflect new routes that have been learned or routes that have been removed and/or when labels are allocated or freed

MPLS LDP Local Checkpointing Key Points (Cont.)

- Checkpointing stops during control plane disruptions, GR, and recovery
- Label bindings on backup RP are marked checkpointed
- This marking is removed when it becomes active RP

LDP Graceful Restart Mechanism

- Described in RFC3478
- The LSR sends the LDP initialization message to a neighbor to establish an LDP session
- The Fault Tolerant (FT) session Type Length Value (TLV) is included in the LDP initialization message

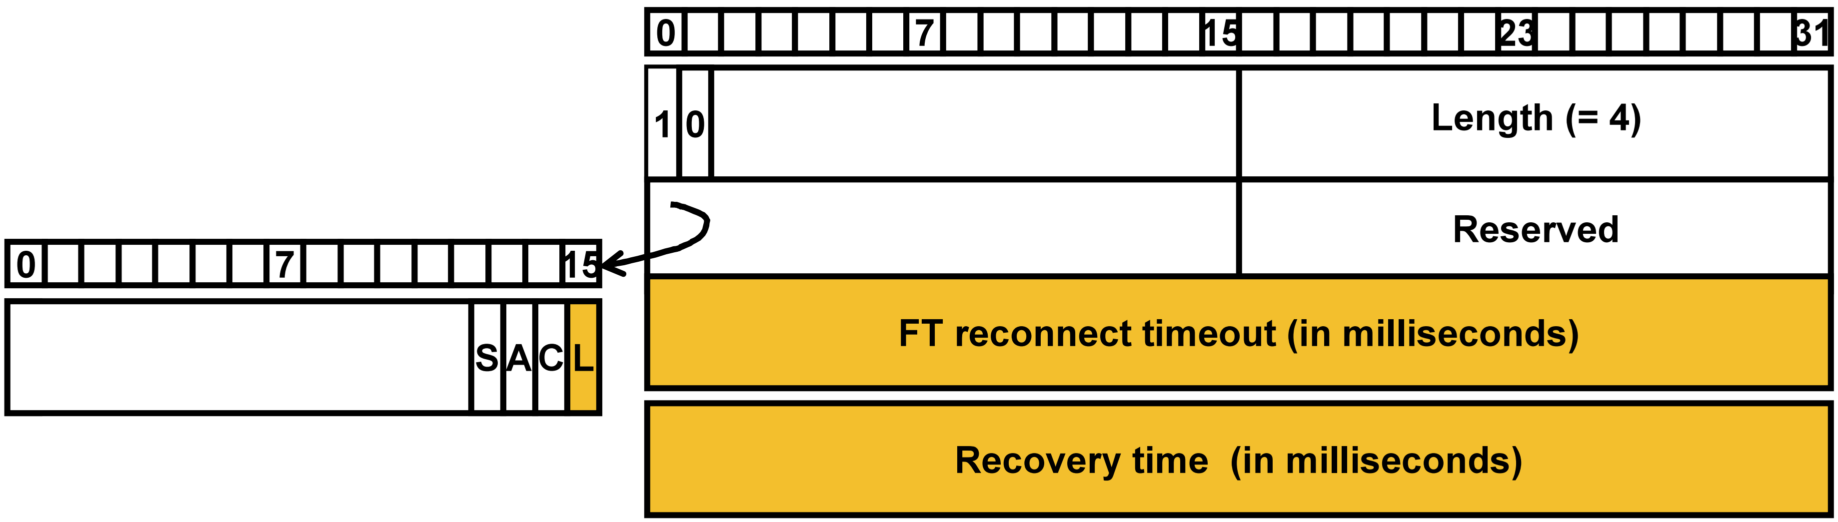

L=1 means GR-LDP is selected
FT Reconnect = 0 means LSR is not NSF capable
FT Recovery Time = 0 means LSR was unable to preserve MPLS forwarding state across restart

LDP Graceful Restart Key Points

- MPLS LDP GR must be enabled before an LDP session is established on all the LSRs
- Both directly connected and non-directly connected peers (targeted sessions) are supported
- LDP GR supports both failure cases
    1. LDP restarts
    2. LDP session resets
- Restart timers can be adjusted to limit the session re-establishment time at restart
- If an LSR proposes LDP GR capability to a non-LDP GR capable LSR

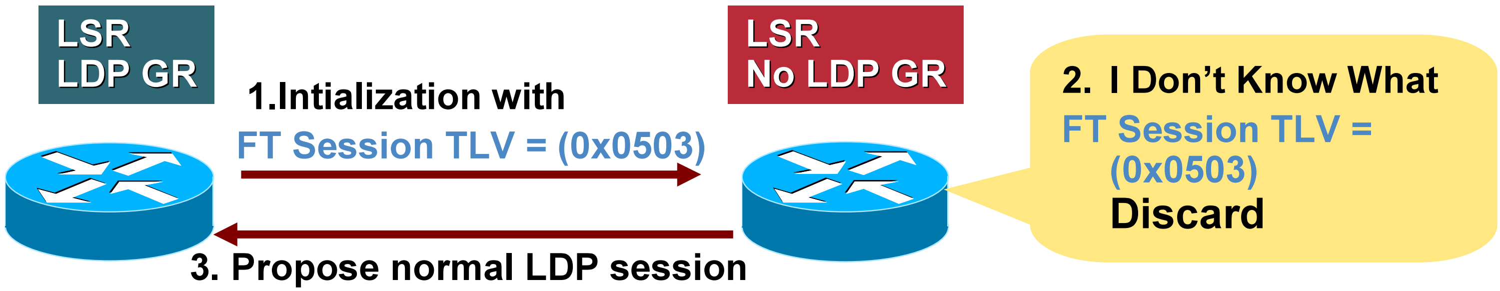

LDP Graceful Restart Process

- **Is LDP GR supported on LSRs—negotiate restart capabilities**
- **Retain old forwarding plane info—LDP bindings—during (fault tolerant reconnect) timeout**
- **Restart/recover**

Deploying MPLS HA Example

HA aware devices:
Graceful restart functionality for all the related protocols CEs, Ps, and RRs

HA capable devices
Full NSF/SSO functionality (PE1 and PE2)

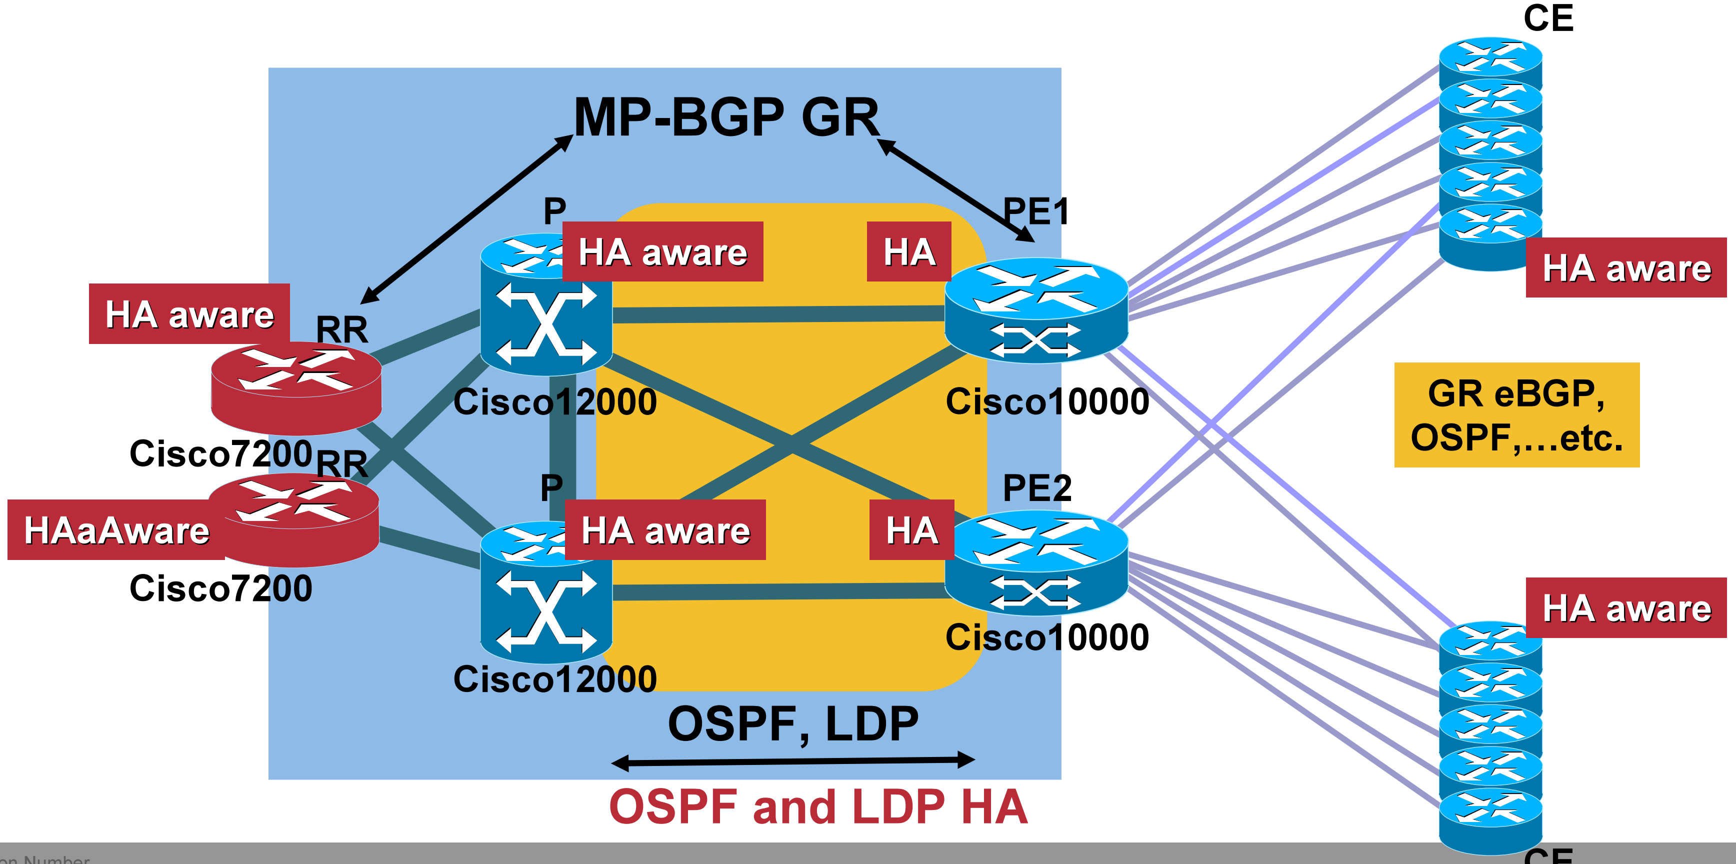

LDP GR on SSO-LSR and SSO-Aware Peer

- **SSO capable LSR failed**
  - Active RP failed
  - Continue forwarding using the stale state
  - Standby RP becomes Active
  - Mark the forwarding state as stale, and retain it
  - Reestablish LDP session
- **SSO aware neighbor LSR**
  - LDP failure detected
  - Mark the forwarding state as stale, and retain it
  - Reestablish LDP session

LDP Graceful Restart Operation

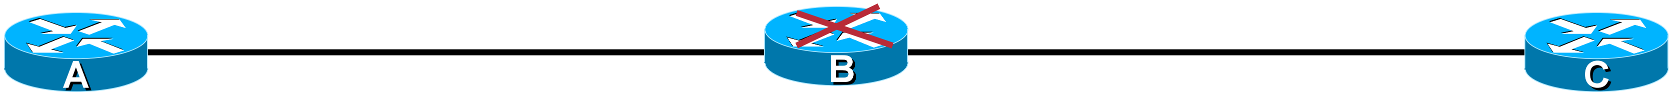

- LDP paths established, LDP GR negotiated
- When RP fails on LSRb, communication between peers is lost; LSRb encounters a LDP restart, while LSRa and LSRc encounter an LDP session reset
- LSRa and LSRc mark all the label bindings from LSRb as stale, but continue to use the same bindings for MPLS forwarding
- LSRa and LSRc attempt to reestablish an LDP session with Rb
- LSRb restarts and marks all of its forwarding entries as stale; at this point, LDP state is in restarting mode
- LSRa and LSRc reestablish LDP sessions with Rb, but keep their stale label bindings; at this point, the LDP state is in recovery mode
- All routers re-advertise their label binding info; stale flags are removed if a label has been relearned; new LFIB table is ready with new local label, outgoing label or VC, prefix or tunnel ID, label-switched bytes, outgoing interface and Next Hop

LDP Graceful Restart Configuration

If LDP session doesn't reestablish or if LSR doesn't restart

```
mpls ldp graceful-restart
mpls ldp graceful-restart timers forwarding-holding
mpls ldp graceful-restart timers neighbor-liveness
mpls ldp graceful-restart timers max-recovery
```

Configuration options:

- **If enabled globally, don't need to enable per interface**
- **Don't need to enable globally, can enable on selected interfaces**

LDP Graceful Restart Configuration (Cont.)

- If LDP sessions are not reestablished or if the LSR does not restart, various timers can be adjusted at global config level

  `mpls ldp graceful-restart timers forwarding-holding`

  - Amount of time the MPLS forwarding state should be preserved after the control plane restarts (in seconds)
  - If the timer expires, all the entries marked stale are deleted

  `mpls ldp graceful-restart timers neighbor-liveness`

  - Amount of time an LSR should wait for an LDP session to reconnect (in seconds)
  - If the LSR cannot recreate an LDP session with the neighbor in the time allotted, the LSR deletes the stale label-FEC bindings received from that neighbor (def: 5sec)

  `mpls ldp graceful-restart timers max-recovery`

  - Amount of time an LSR should hold stale label-FEC bindings after an LDP session has been re-established
  - After the timer expires, all prefixes are removed from the binding and forwarding table
  - Set this timer to a value that allows the neighboring LSRs to re-sync the LSPs without creating congestion in the LDP control plane (def: 120sec)

LDP Graceful Restart Configuration

LDP GR must be enabled on all the LSRs to take full advantage of LDP GR in a network

1. Enable ip cef

    ip cef <distributed> (if distributed mode)

2. Enable GR globally (must do before enabling LDP)

    mpld ldp graceful-restart

3. Enable MPLS (global)

    mpls ip

4. Enable LDP on appropriate interfaces

    mpls label protocol ldp

LDP Graceful Restart Troubleshooting

Show mpls ldp neighbor	Display LDP/TDP neighbor info
Show mpls ldp discovery	Display status of the LDP discovery process
show mpls ldp neighbor graceful-restart	Verify all LDP sessions are configured for LDP GR
Show mpls bindings	Look at LIB table, make sure active and backup processor has identical copies of the local label bindings
Show mpls ldp checkpoint	To display local chekpoint info on the active RP
Clear mpls ldp checkpoint	Clear checkpoint info in LIB on the Active RP & delete all LIB entries learned by checkpointing on the standby RP
Debub mpls ldp checkpoint Debug mpls ldp grace-restart	For debugging sessions

show mpls ldp graceful-restart

- **Shows a summary of the global LDP restart states**

```
Router# show mpls ldp graceful restart
LDP Graceful Restart is enabled
Neighbor Liveness Timer: 5 seconds
Max Recovery Time: 200 seconds
Down Neighbor Database (0 records):
Graceful Restart-enabled Sessions:
Graceful Restart-enabled Sessions:
VRF default:
Peer LDP Ident: 18.18.18.18:0, State: estab
Peer LDP Ident: 17.17.17.17:0, State: estab
```

show mpls ldp bindings

- **Shows contents of the Label Information Base (LIB), local and remote binding**

```
Router# show mpls ldp bindings
34.0.0.0/8, rev 9
local binding: label: imp-null
remote binding: lsr: 155.0.0.55:0, label: 17
remote binding: lsr: 66.66.0.66:0, label: 18
remote binding: lsr: 144.0.0.44:0, label: imp-null
45.0.0.0/8, rev 17 local binding: label: 19
remote binding: lsr: 155.0.0.55:0, label: imp-null
remote binding: lsr: 66.66.0.66:0, label: 16
remote binding: lsr: 144.0.0.44:0, label: imp-null
66.66.0.66/32, rev 19
```

show mpls ldp neighbor

- **Displays the status of LDP sessions and information about LDP neighbors for the default routing domain**

```
Router# show mpls ldp neighbor
Peer LDP Ident: 203.0.7.7:2; Local LDP Ident 8.1.1.1:1
TCP connection: 203.0.7.7.11032 - 8.1.1.1.646
State: Oper; Msgs sent/rcvd: 5855/6371; Downstream on demand
Up time: 13:15:09
LDP discovery sources:
ATM3/0.1
Peer LDP Ident: 7.1.1.1:0; Local LDP Ident 8.1.1.1:0
TCP connection: 7.1.1.1.646 - 8.1.1.1.11006
State: Oper; Msgs sent/rcvd: 4/411; Downstream
Up time: 00:00:52
LDP discovery sources:
Ethernet1/0/0
Addresses bound to peer LDP Ident:
2.0.0.29 7.1.1.1 59.0.0.199 212.10.1.1 10.205.0.9
```

Clear ldp checkpoint

- **On the active RP, this command clears the checkpoint-related state from the specified LIB entries**
- **Triggers an immediate checkpointing attempt for those entries**
- **On the standby, it deletes all LIB entries learned via checkpointing**

Agenda

- Introduction to Cisco IOS High Availability
- MPLS NSF/SSO Overview
- MPLS NSF/SSO Components
  - MPLS NSF/SSO —LDP NSF/SSO
  - MPLS NSF/SSO—BGP VPNv4 NSF/SSO
  - MPLS NSF/SSO—Management
- Summary

MPLS VPNv4 HA Elements

- **MPLS VPN checkpointing capability**
  - –Active RP checkpoints the following information to the back up RP after the MPLS forwarding is updated

    {<VRFID>, <prefix>, <mask>, <local label>}

- **BGP graceful restart capability**

MPLS BGP VPNv4 GR Mechanism

- **MPLS BGP GR mechanism defines preservation of forwarding state across BGP restart**

  draft-ietf-mpls-bgp-mpls-restart

- **A new graceful restart capability is carried in BGP open message**

  A BGP update message with no reachable NLRI and empty withdrawn NLRI is specified as an end-of-RIB marker

  AFI/SAFI (Subsequent Address Family Identifier) pairs in the graceful restart cap is used by an LSR for indicating its ability to preserve MPLS forwarding state across BGP restart

  SAFI in the advertise capability indicates that NLRI field carries addressing information as well as labels

MPLS VPN: BGP Graceful Restart Procedure

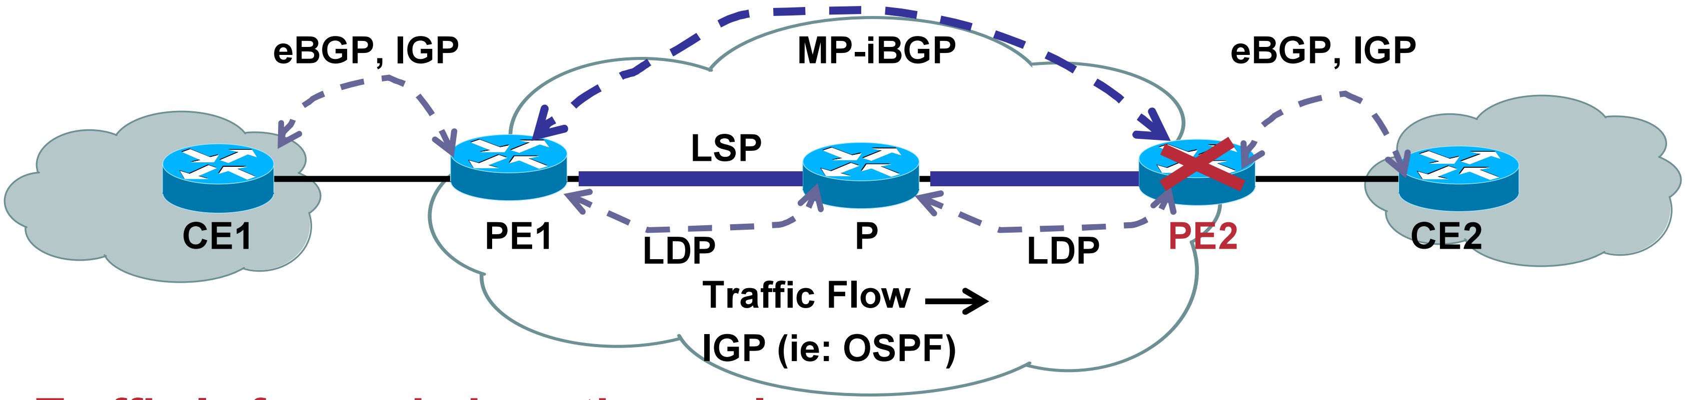

Traffic Is forwarded continuously

- PE1-PE2 exchange graceful restart cap (restart time, AFI/SAFI, etc.)
- PE1 and PE2 exchange routing information via BGP update messages
- Assume LSP has been established from PE1 to PE2
- PE2 restarts (active RP fails, switchover occurs)

MPLS VPN: BGP Graceful Restart Procedure (Cont.)

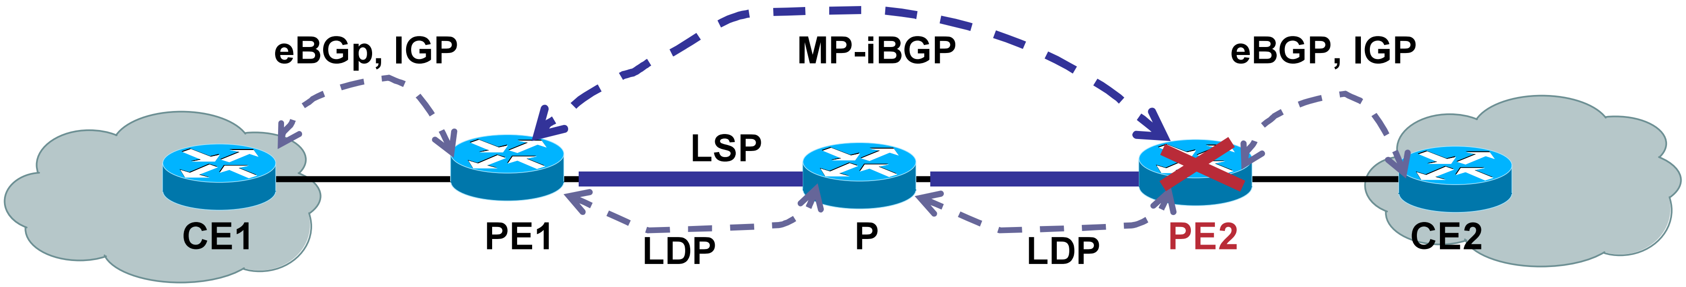

- PE1 detects PE2's failure, retains its last Adj-RIB-In and forwarding state learned from PE2; PE1 will delete this state if session is not re-established within restart time
- The BGP session between PE1 and PE2 goes down
- PE1 marks the entries in its BGP table which it learned from PE2 as stale but does not delete the routes from its VRF tables; hence it continues to forward traffic to PE2
- PE2 switches over to back up RP; continues to forward traffic using the backed up info

MPLS VPN: BGP Graceful Restart Procedure (Cont.)

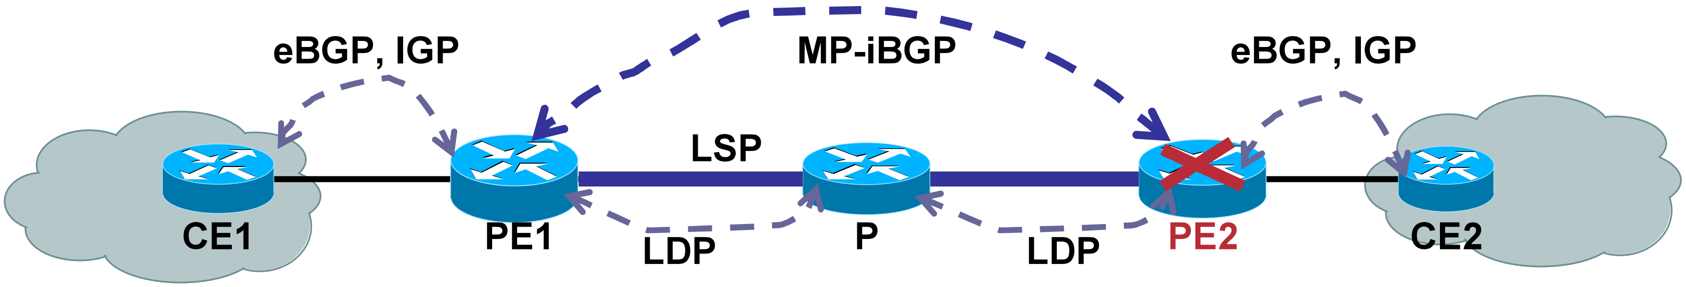

- PE2 re-establishes TCP session with PE1 (hopefully within restart time-180seconds default)
- PE1 sends BGP updates from Adj-RIBs-Out to PE2 along with the label mapping; on completion, sends End-of-RIB marker
- PE2 runs decision process after receiving End-of-RIB markers from all BGP peers, updates its Loc-RIB, FIB, Adj_RIBs-Out and advertise its routes to PE1; on completion sends End-of-RIB marker

MPLS VPN: BGP Graceful Restart Procedure (Cont.)

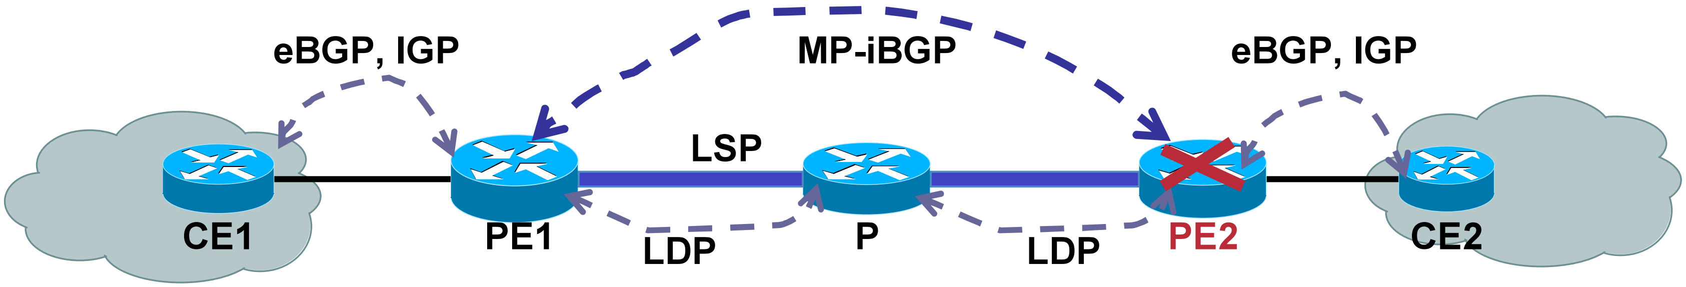

- **Suppose PE2 had bound an in label L1 to a FEC, it picks the same label (if checkpointed) or allocates a new one and advertises it to PE1 for this route**
- **PE2 updates its Adj-RIBs-In, on receipt of End-of-RIB marker, deletes stale entries, runs its decision process, updates its Loc-RIB, FIB, and Adj-RIBs-Out; removes stale state if the labels are the same (routing info and Out_Label are recovered from peer; In_Label is either checkpointed or a new one allocated)**
- **Back to normal operation**

How to Enable BGP VPNv4 SSO/NSF

The following components must be enabled to support BGP VPNv4 NSF

- Enabling SSO
  - http://www.cisco.com/univercd/cc/td/doc/product/software/ios122s/122snwft/release/122s20/fssso20s.htm
- Enabling LDP graceful restart
  - http://www.cisco.com/univercd/cc/td/doc/product/software/ios122s/122snwft/release/122s25/fsldpgr.htm
- Enabling nonstop forwarding for routing protocols
  - http://www.cisco.com/univercd/cc/td/doc/product/software/ios122s/122snwft/release/122s20/fsnsf20s.htm
- Enabling NSF support for basic VPNs
- Configuring NSF support for VPN interfaces that use BGP for label distribution
- Verifying VPN NSF (optional)

MPLS L3 VPN Features

Supported features

- **MPLS VPN BGP GR for VPNv4**
- **MPLS VPN BGP checkpointing**
- **MPLS VPN SSO/NSF support for VRF**
- **MPLS VPN SSO/NSF support for I-AS**
- **MPLS VPN SSO/NSF support for CSC**

  Additional features will be supported in subsequent releases; please refer to Cisco Feature Navigator for the latest information

  www.cisco.com/go/fn

BGP VPNv4 GR Configuration

- **Enable MPLS nonstop forwarding on an interface that uses BGP as the label distribution protocol**

```
mpls forwarding bgp
```

BGP VPNv4 GR Restrictions

- LDP sessions are supported, but Tag Distribution Protocol (TDP) sessions are not supported
- VPN NSF cannot be configured on label-controlled ATM (LC-ATM) interfaces
- VPN NSF requires that neighbor networking devices be NSF-aware

  Peer routers must support the graceful restart of the protocol used to communicate with the VPN NSF-capable router

Cisco 10000 Series Routers MPLS SSO Sample Test Topology

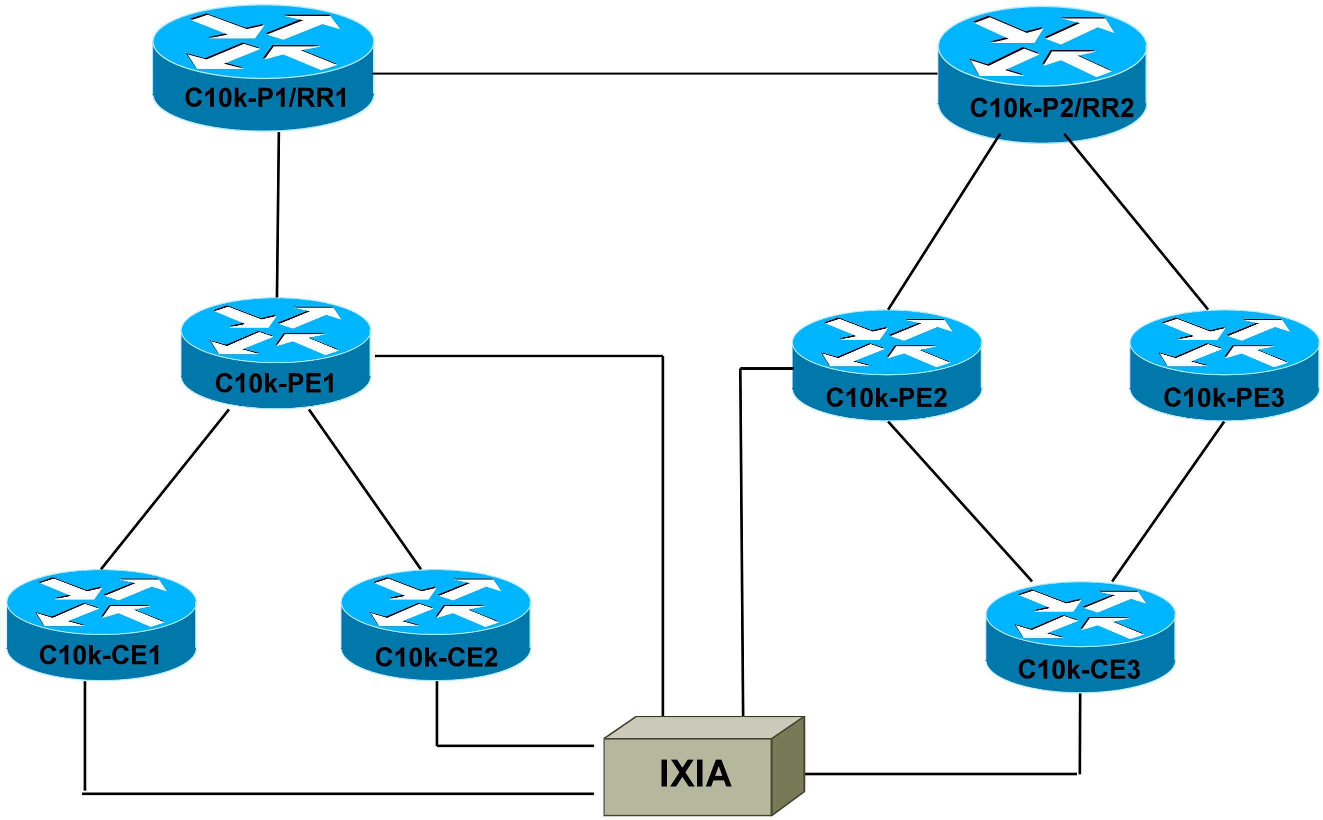

- C10k-P1/RR1, C10k-P2/RR2: MPLS SSO/NSF aware
- C10k-PE1, C10k-PE2, C10k-PE3 are SSO/NSF capable with dual RP.
- C10K-CEs are enabled with RP Graceful restart capability

```
Core Routing:  OSPF/ISIS
Edge Routing:  eBGP/OSPF/eIGRP/RIP/Static
PE-P LCs:      HHGIGE/OC3POS/OC12ATM
PE-CE LCs:     HHGIGE/6CT3/OC12ATM
```

BGP VPNv4 Graceful Restart Troubleshooting

Show ip bgp labels	Displays information about MPLS labels from eBGP route table
Show ip bgp vpnv4 all labels	Displays info on the active and standby RPs when they are configured for MPLS VPN NSF
Debug ip bgp vpnv4 checkpoint	Displays the events for the VRF checkpointing system between the active and standby RP
Debug ip bgp vpnv4 nsf	Displays the NSF events for the VRF tableid synch subsystem between the active and standby RPs
Debug mpls checkpoint label-binding	Displays the events for the checkpoint label bindings of MPLS applications running on the router

Show ip bgp labels

Displays EBGP labels associated with an ASBR

```
Router#show ip bgp labels
Network           Next Hop         In Label/Out Label
3.3.0.0/16        0.0.0.0          imp-null/exp-null
15.15.15.15/32    15.15.15.15      18/exp-null
16.16.16.16/32    0.0.0.0          imp-null/exp-null
17.17.17.17/32    34.0.0.1         20/exp-null
18.18.18.18/32    43.0.0.1         24/31
18.18.18.18/32    38.0.0.1         24/33
19.19.19.19/32    43.0.0.1         25/32
19.19.19.19/32    38.0.0.1         25/34
20.20.20.20/32    43.0.0.1         21/30
20.20.20.20/32    38.0.0.1         21/32
33.0.0.0          15.15.15.15      19/exp-null
34.0.0.0          0.0.0.0          imp-null/exp-null
35.0.0.0          43.0.0.1         22/29
```

Show ip bgp vpnv4 all labels

- **NH and network prefixes are from BGP table**
- **From the same LSR, Info on Standby RP**
- **InLabels are local labels, which can be verified by comparing the log to SHOW IP BGP LABELS. OutLabels are assigned by BGP NH router. They should be the same, because BGP is used as a label distribution protocol**
- **From the same LSR, Info on Active RP**

```
Router#show ip bgp vpnv4 all labels
Network          Next Hop    In label/Out label
Route Distinguisher: 100:1 (vpn1)
12.12.12.12/32   0.0.0.0     16/aggregate(vpn1)
135.0.0.0/8      0.0.0.0     17/aggregate(vpn1)
Route Distinguisher: 609:1 (vpn0)
13.13.13.13/32   0.0.0.0     18/aggregate(vpn0)
```

- **From the same LSR, Info on Standby RP**

```
Router#show ip bgp vpnv4 all labels
Network        Masklen    In label
Route Distinguisher: (dec) 0001000001
12.12.12.12    /32         16
135.0.0.0      /8          17
Route Distinguisher: (dec) 002970001
13.13.13.13    /32         18
```

Debug ip bgp vpnv4 checkpoint

Displays the events for the VRF checkpointing system between the active and standby RP

```
Router#debug ip bgp vpnv4 checkpoint
3d18h: %HA-5-SYNC_NOTICE: Config sync started.
3d18h: vrf-nsf: vrf vpn2 tableid 1 send OK
3d18h: vrf-nsf: vrf tableid bulk sync
complete msg send OK
3d18h: vrf-nsf: CF send ok
3d18h: vrf-nsf: CF send ok
3d18h: %HA-5-SYNC_NOTICE: Config sync completed.
3d18h: %HA-5-SYNC_NOTICE: Standby has restarted.
3d18h: %HA-5-
MODE: Operating mode is sso,config mode is sso
```

Debug ip bgp vpnv4 nsf on the Active RP

Router#debug ip bgp vpnv4 nsf

```
MPLS VPN NSF Processing debugging is on
Router(config)#ip vrf vpn3

3d18h: vrf-nsf: vrf vpn3 tableid 2 send rpc OK

Router(config-vrf)#no ip vrf vpn3

%IP addresses from all interfaces in VRF vpn3 have been rem
oved

3d18h: vrf-nsf:
rx vrf tableid delete complete msg, tid = 2, name = vpn3
```

Debug ip bgp vpnv4 nsf on the Standby RP

Router#debug ip bgp vpnv4 nsf

```
MPLS VPN NSF Processing debugging is on
00:05:21: vrf-
nsf: rx vrf tableid rpc msg, tid = 2, name = vpn3
%IP addresses from all interfaces in VRF vpn3 have been rem
oved 00:06:22: vrf-

nsf: vrf vpn3 tableid 2 , delete complete, send OK
```

Debug mpls checkpoint label-binding

Output log when the command is issued on the standby route processor

Router#debug mpls checkpoint label-binding

```
MPLS Label Binding Checkpoint debugging is on
3d17h: mpls_lbl_bind_chkpt: client ID 13 up, total client 1
3d17h: mpls_lbl_bind_chkpt: msg rx for 1D, vers 0, type 1
action 56, len 0, state 4, peer 13
3d17h: mpls_lbl_bind_chkpt: post msg type 1
3d17h: mpls_lbl_bind_chkpt: msg rx for 1D, vers 0, type 1
action 56, len 0, state 4, peer 13
3d17h: mpls_lbl_bind_chkpt: post msg type 1
3d17h: mpls_lbl_bind_chkpt: msg rx for 1D, vers 0, type 1
action 56, len 0, state 4, peer 13
3d17h: mpls_lbl_bind_chkpt: post msg type 1
3d17h: mpls_lbl_bind_chkpt: appl_id 13, KEY 000C800018888200
3d17h: mpls_chkpt_db: AVL insert successful, Key 000C800018888200
action Add, label 19
3d17h: mpls_lbl_bind_chkpt: appl_id 13, KEY 000C800013200080
3d17h: mpls_chkpt_db: AVL insert successful, Key 000C800013200080
action Add, label 20
3d17h: mpls_lbl_bind_chkpt: appl_id 13, KEY 000C8000138383838200
3d17h: mpls_chkpt_db: AVL insert successful, Key 000C8000138383838200
action Add, label 21 3d17h: Stby RP OR CF peer not ready, don't send msg
3d17h: mpls_lbl_bind_chkpt: client ID 13 down, total client 0
3d17h: mpls_lbl_bind_chkpt: msg rx for 1D, vers 0, type 1
action 56, len 2, state 4, peer 13
3d17h: mpls_lbl_bind_chkpt: post msg type 1
3d17h: mpls_lbl_bind_chkpt: appl_id 13, KEY  action NSF unconfig, appl id 13
```

Requirements to support VPNv4 GR

- **All peering LSRs need to support restart to take full advantage of VPNv4 GR**
- **Default to nonGR BGP if one of the neighbors doesn't support it, it will default back to original BGP**
- **If label binding on an LSR is created/maintained by multiple protocols (LDP, RSVP-TE) individually, LSR needs to preserve across the restart the info about which protocol has assigned which labels**

Agenda

- Introduction to Cisco IOS High Availability
- MPLS NSF/SSO Overview
- MPLS NSF/SSO Components
    - MPLS NSF/SSO —LDP NSF/SSO
    - MPLS NSF/SSO—BGP VPNv4 NSF/SSO
    - MPLS NSF/SSO—Management
- Summary

MPLS HA: Management MIBs

- **MPLS VPN: SSO/NSF aware VPN MIB traps**
- **MPLS: SSO/NSF aware LDP MIB traps**

Agenda

- Introduction to Cisco IOS High Availability
- MPLS NSF/SSO Overview
- MPLS NSF/SSO Components
    - MPLS NSF/SSO —LDP NSF/SSO
    - MPLS NSF/SSO—BGP VPNv4 NSF/SSO
    - MPLS NSF/SSO—Management
- Summary

Summary

- Cisco is enhancing its portfolio to add features for improved full HA NSF/SSO solution

    MPLS NSF/SSO features are part of that solution and provide stateful switchover and NSF capability for MPLS forwarding, LDP and VPN

- Need IP NSF/SSO enabled to support MPLS NSF/SSO; GR must be enabled for all participating routing protocols (OSPF, BGP, IS-IS) on P, PE, and CE routers
- For MPLS VPN NSF/SSO: MFI HA, LDP HA, BGP HA is required
- Some CLI changes