MPLS HIGH AVAILABILITY OVERVIEW & LABEL DISTRIBUTION PROTOCOL HIGH AVAILABILITY

MARCH 2004

Agenda

- Introduction to Cisco IOS High Availability (HA)
- IP Nonstop Forwarding with Stateful Switchover (NSF with SSO)
- MPLS Co-existence with IP NSF with SSO
- MPLS Forwarding Infrastructure (MFI)
- MPLS HA Overview
- LSD NSF with SSO
- Summary

Customers Demand More from IP/MPLS Networks

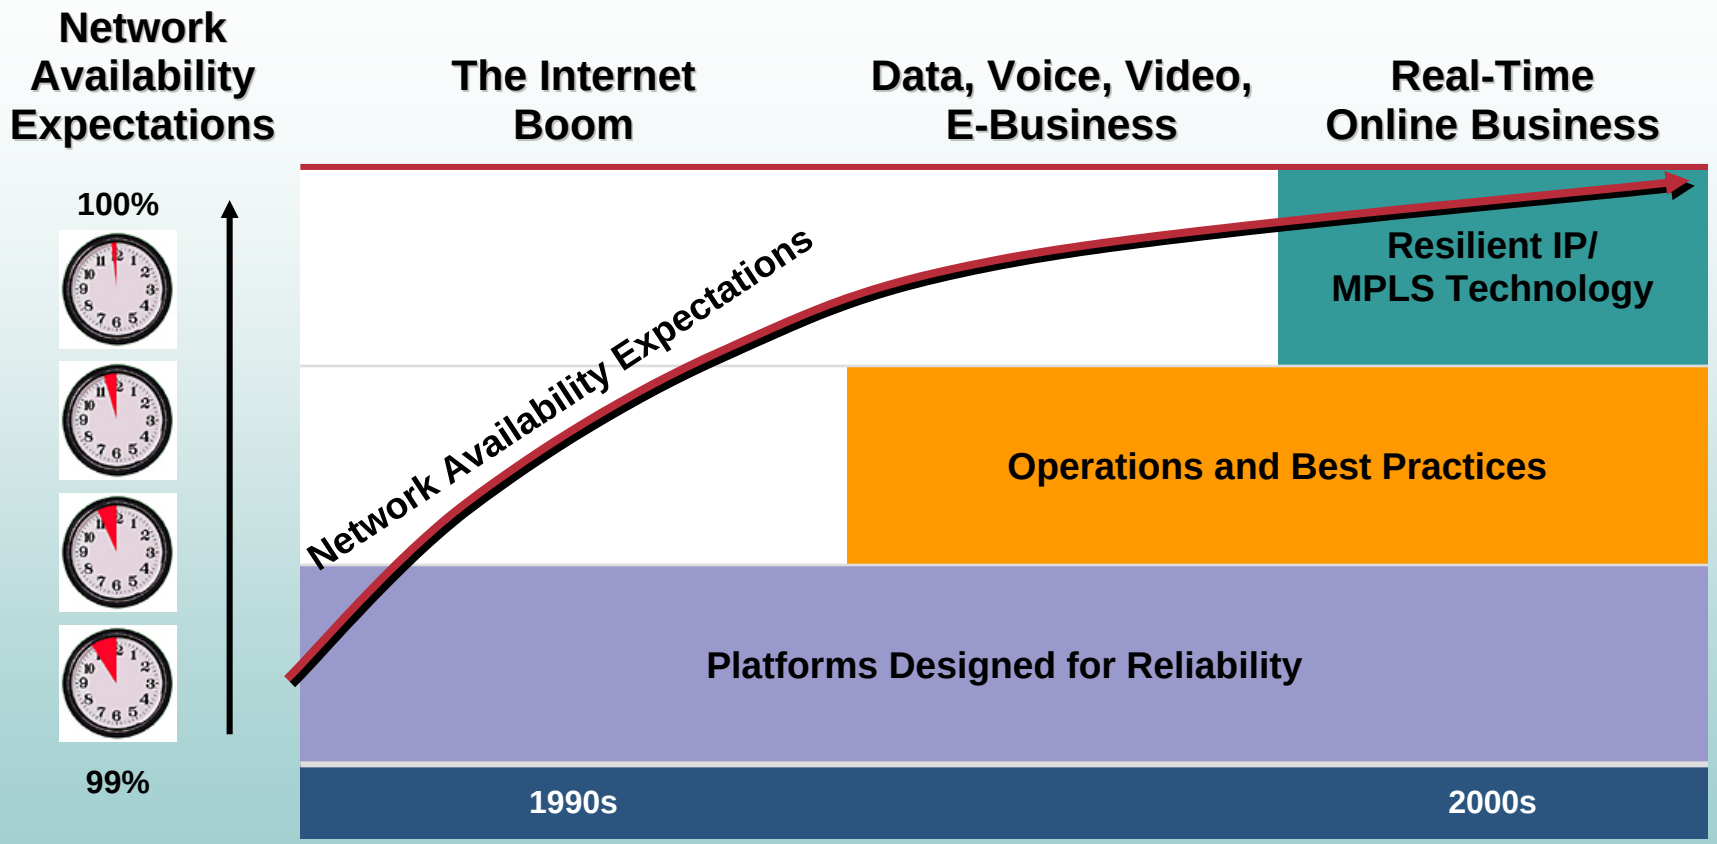

IP Convergence Drives The Need For Network-wide Availability

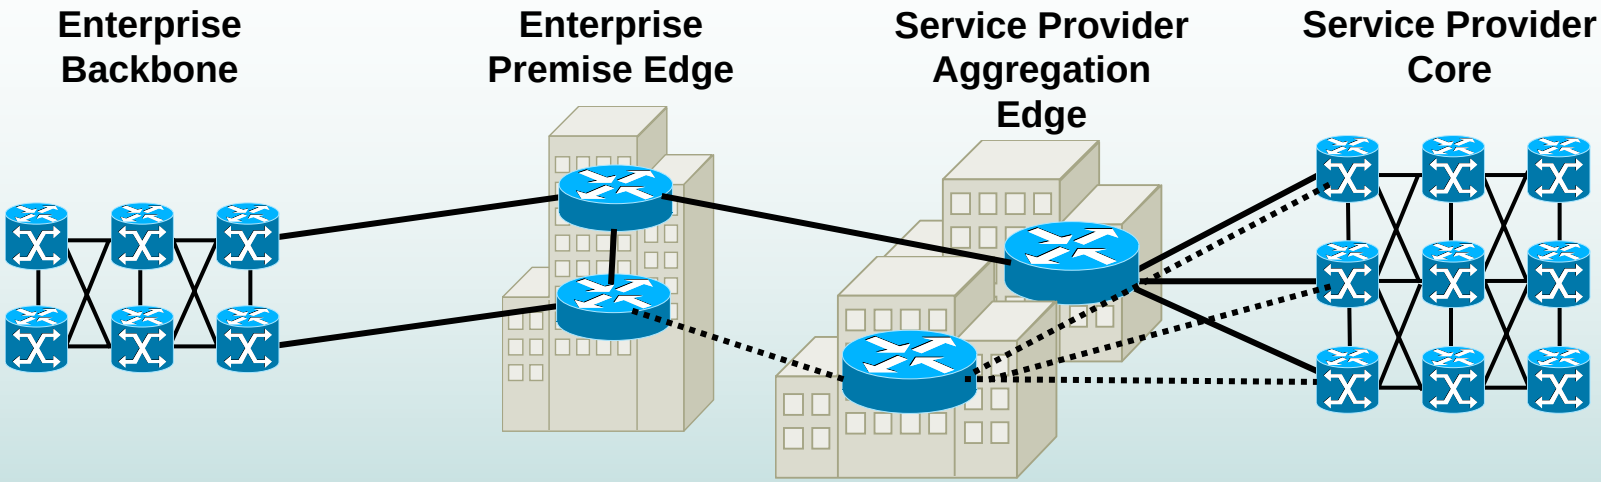

A network is only as available as its weakest link

Different network architectures and equipment create a need for appropriate software solutions in order to recover from network failures

Cisco IOS Software High Availability: Protecting the *Entire Network*

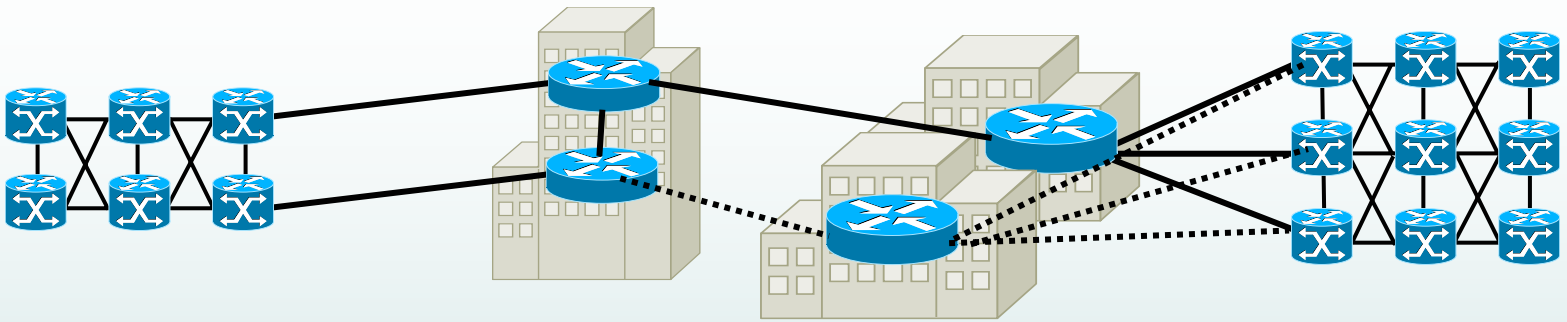

- Multicast Sub-Second Convergence
- IGP Fast Convergence

- Gateway Load Balancing Protocol (GLBP)
- Stateful NAT
- Stateful IPsec

- Cisco NSF with SSO
- MPLS High Availability
- ISSU
- LC Redundancy

- MPLS Fast Reroute - Node Protection
- IGP Fast Convergence

Routing Protocol Convergence Enhancements

Delivering Network-Wide Resilience

- Industry's broadest portfolio of end-to-end high availability features
- Industry first routing resiliency: zero packet loss (Cisco 12000 Series Internet Router)
- No fork-lift upgrades: high availability functionality on existing Cisco hardware
- Industry-leading advances in convergence times and recovery
- Load sharing between routers creates superior thruput and redundancy

Cisco IOS Software High Availability Dimensions

Component and Device Level Resiliency

- Component Resiliency
- Device resiliency and failover
- Resilient IP Services
- Denial of Service, self protection

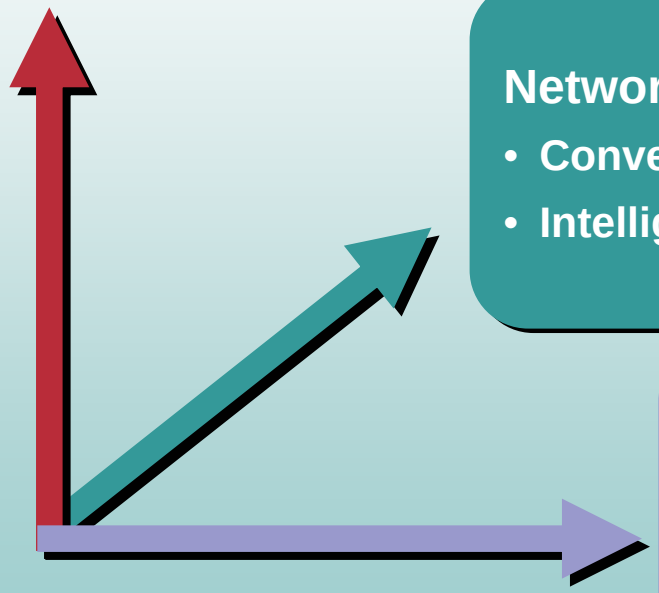

Network Level Resiliency

- Convergence and self-healing
- Intelligent protocol fabric

Operations and Management

- Hot software upgrades
- Automated, low-error configuration
- Fault / event management
- Availability measurement

Cisco IOS Software High Availability Overview

Protect networks from failures *end-to-end*

- **Component and Device Level Resiliency**
  - **Across the board component resiliency—line card, route processor, software**
  - **Radical reduction in unplanned outages for IP/MPLS network**
  - **More efficient and cost-effective redundancy for remote devices**
  - **Self-protection from Denial of Service attacks built into the device**
- **Network Level Resiliency**
  - **Optimize convergence algorithms, speeding up network recovery**
  - **Intelligent protocol fabric with network-wide NSF awareness**
- **Operations and management**
  - **Embed intelligent event management for proactive maintenance**
  - **Automation and configuration rollback to reduce human errors**
  - **Embedded, lightweight measurement of availability metrics**

Cisco IOS Software High Availability Overview (Cont.)

Requirements	Features	
	Available Today	Available CY04/1H05
Component and Device Level Resiliency	• Cisco NSF with SSO • Gateway Load Balancing Protocol • Stateful NAT • Warm Reboot • Control Plane Policing	• Stateful IPsec • MPLS High Availability • Line Card Redundancy (1+1 APS with No Layer3 Re-Convergence)
Network Level Resiliency	• Routing Convergence Enhancements: BGP Optimization Incremental SPF Optimization IP Event Dampening • Multicast Sub-Second Convergence • MPLS Fast ReRoute • NSF Awareness	• Graceful Restart Enhancements • IGP Fast Convergence
Operations and Management	• Embedded Event Manager (EEM) • Component Outage On-Line (COOL) MIB • Cisco AutoQoS • Cisco AutoSecure	• In-Service Software Upgrades (ISSU) • Configuration Rollback

Agenda

- Introduction to Cisco IOS High Availability (HA)
- IP Nonstop Forwarding with Stateful Switchover (NSF with SSO)
- MPLS Co-existence with IP NSF with SSO
- MPLS Forwarding Infrastructure (MFI)
- MPLS HA Overview
- LSD NSF with SSO
- Summary

Protecting the Service Aggregation Edge: Cisco NSF with SSO

- Cisco NSF: minimal or no packet loss for the Edge and Core
  - Packet forwarding continues while peering relationships are reestablished
  - No route flaps between participating neighbor routers
  - Transparent route convergence: Layer 3 routing information recovered from peers (BGP, OSPF, IS-IS, EIGRP)
- SSO: zero interruption to Layer 2 sessions
  - Protection from hardware *or* software faults
  - Active route processor synchronizes information with standby route processor
  - Maintains session state for high availability-aware protocols (Frame Relay, ATM, PPP, MLPPP, cHDLC, APS) on the standby route processor
  - Standby route processor immediately takes control when active route processor is compromised

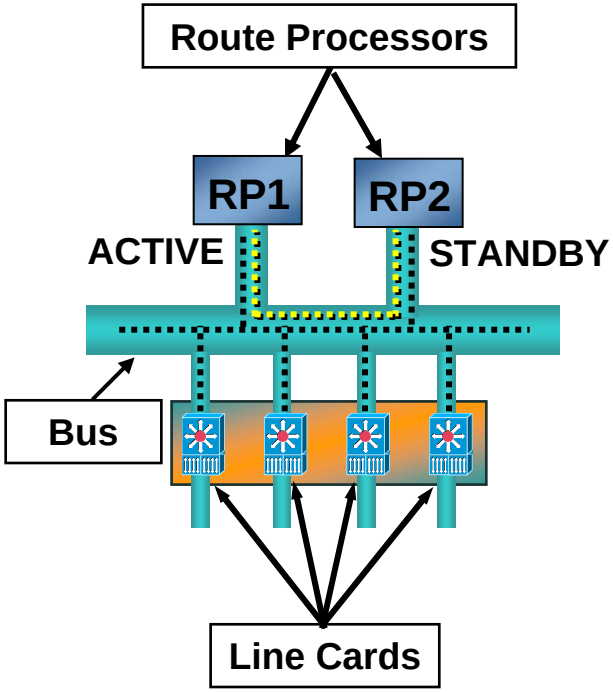

HOT Redundancy Solution

SSO Architecture: Switchover

1. RPs initialize, assign roles (active or standby), negotiate mode (SSO), begin checkpointing state

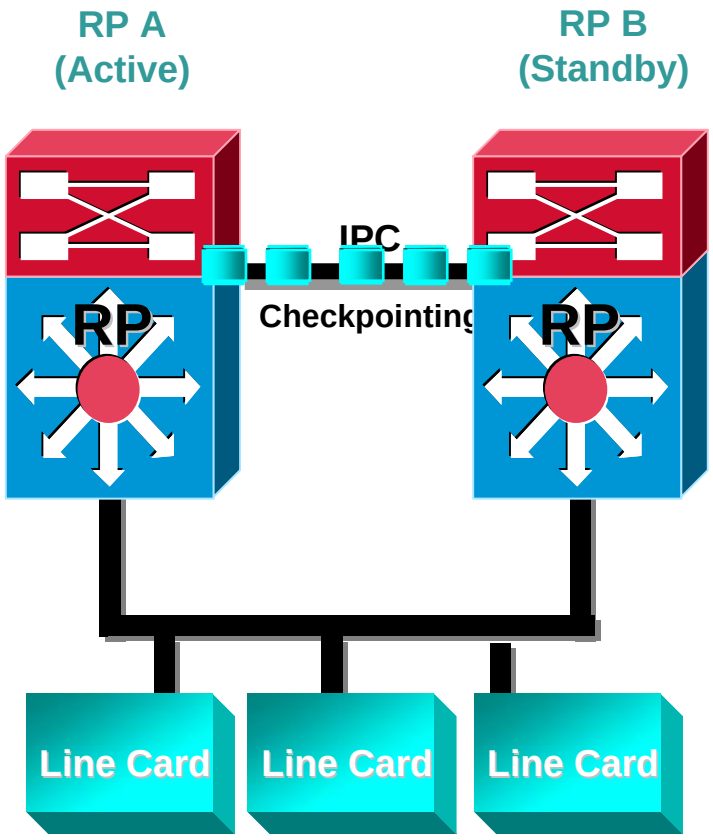

SSO Architecture: Switchover (Cont.)

1. RPs initialize, assign roles (active or standby), negotiate mode (SSO), begin checkpointing state

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

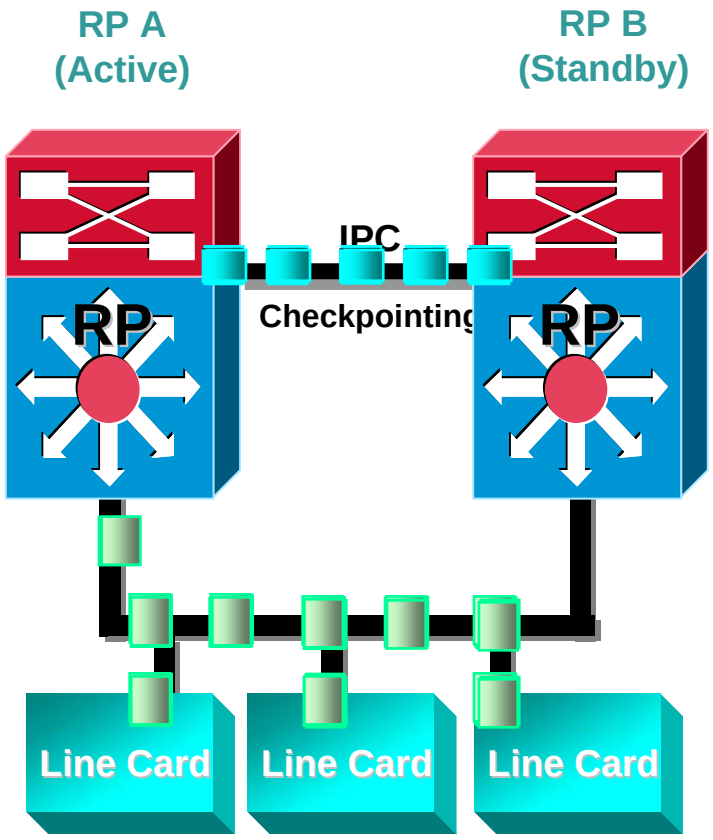

SSO Architecture: Switchover (Cont.)

1. RPs initialize, assign roles (active or 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 and checkpointing stops**
- **Forwarding continues on LCs/FP**
- **RP B assumes Active role and begins providing Layer 2 and 3 services**
- **Layer 2 continues where it left previously stopped**
- **Layer 3 reconverges, updates RIB then FIB**

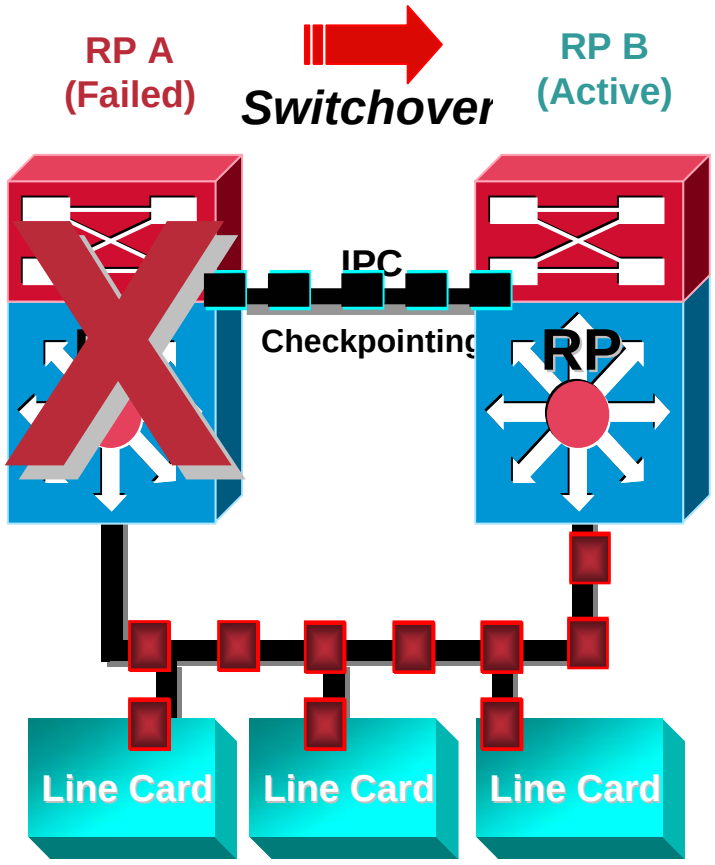

SSO Architecture: Switchover (Cont.)

1. RPs initialize, assign roles (active or 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 and checkpointing stops
   - Forwarding continues on LCs/FP
   - RP B assumes Active role and begins providing Layer 2 and 3 services
   - Layer 2 continues where it left previously stopped
   - Layer 3 reconverges, updates RIB then FIB

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

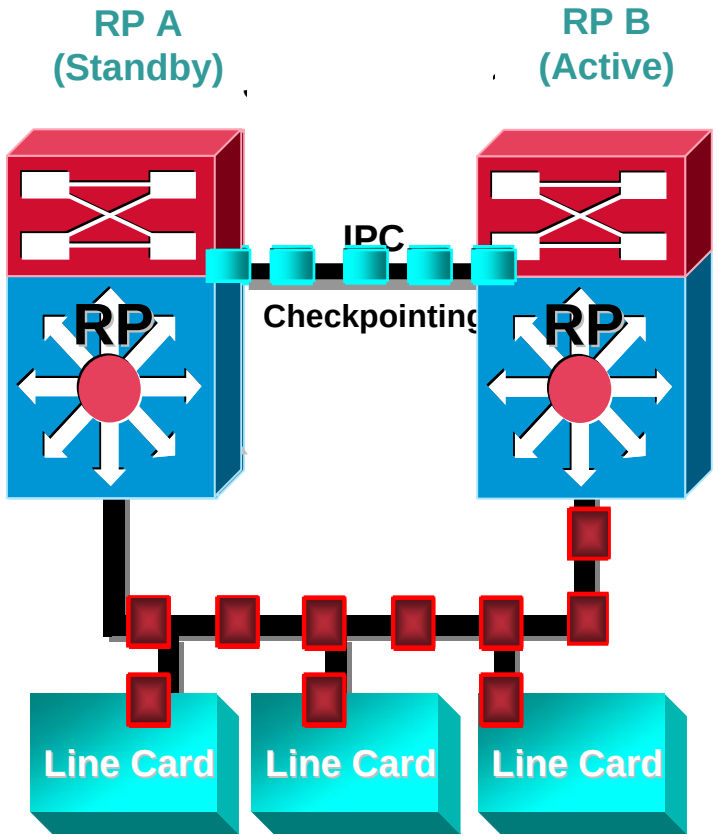

Benefits of Cisco NSF with SSO

- **No service disruptions**

  Preserves user sessions and mitigates impacts of service outage on network users

- **Increased operational efficiencies**

  Reduces network administration, troubleshooting and maintenance due to minimal downtime

- **Reduced costs**

  Decreases downtime costs (ie: Service Level Agreement (SLA) penalties, lost revenue opportunities, and productivity costs for users and administration)

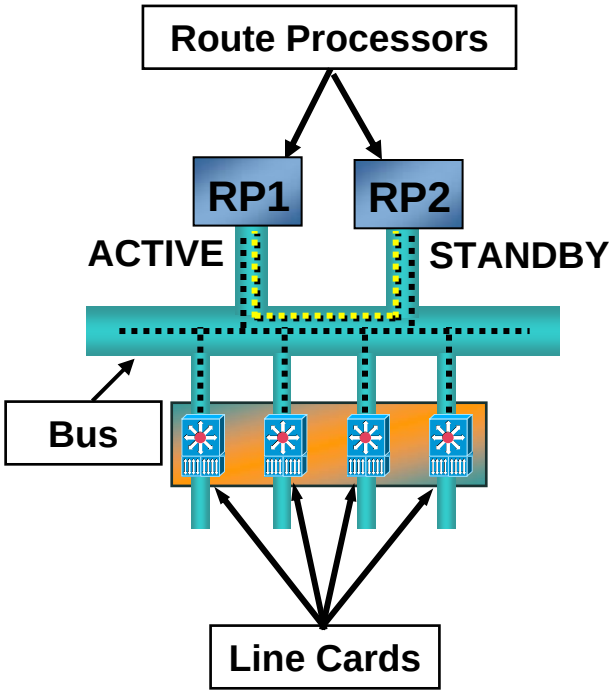

Agenda

- Introduction to Cisco IOS High Availability (HA)
- IP Nonstop Forwarding with Stateful Switchover (NSF with SSO)
- MPLS Co-existence with IP NSF with SSO
- MPLS Forwarding Infrastructure (MFI)
- MPLS HA Overview
- LSD NSF with SSO
- Summary

SSO Coexistence

- Each feature needs SSO implementation
- Service Providers may want to use some SSO capable features along with other features that do not support SSO
- In order to use the mix of SSO and non-SSO features, the non-SSO features may need to be modified to handle the switchover
- SSO coexistence allows the mix of SSO and non-SSO features at the same time
- MPLS SSO coexistence

  During the switchover, MPLS forwarding entries are removed from the linecards and MPLS forwarding is stopped

Cisco NSF with SSO for IP Improves MPLS VPN Recovery Time

- **Initial Cisco NSF with SSO support provides more than coexistence with MPLS**
  - **SSO with MPLS will reduce MTTR by up to several minutes versus a single Route Processor**
- **MPLS tunnels with SSO begin rebuilding more quickly after a switchover to the standby route processor**
  - **No waiting for the route processor or LCs to reload**
  - **No loss of Layer 2 connectivity, so it does not need to be re-established**

Various MPLS SSO-Coexistence Scenarios

- **All peers are SSO capable/aware:**

  IP over MPLS (Label Distribution Protocol (LDP))

  MPLS VPN

- **All PEs are SSO capable/aware:**

  IP over MPLS (LDP)

  MPLS VPN

Agenda

- Introduction to Cisco IOS High Availability (HA)
- IP Nonstop Forwarding with Stateful Switchover (NSF with SSO)
- MPLS Co-existence with IP NSF with SSO
- MPLS Forwarding Infrastructure (MFI)
- MPLS HA Overview
- MPLS HA – LDP NSF with SSO
- Summary

MPLS Forwarding Infrastructure TFIB/LFIB Shortcomings

- **Existing MPLS forwarding is based on TFIB design**
  - **TFIB was designed as a part of TDP/LDP in 1996**
  - **All FEC bound to IPv4 prefixes**
  - **Always expects to have an incoming label**
- **TFIB design has certain deficiencies in handling newer MPLS applications such as TE and AToM**
  - **For a TE tunnel head, the MPLS rewrite is associated with an interface (not IPv4 prefix) and there is no incoming label**
  - **AToM has similar requirements**
- **MFI is a prerequisite for MPLS HA**
- **MFI solves these design issues**

MFI Solution

- Modular Design with clearly defined APIs
- Clean separation between applications and label management
- MFI Control Plane Services
  - Label allocation and management
  - Rewrite setup
  - Interaction with the forwarding paths
  - Interface management
- MFI Data Plane Services
  - Imposition, disposition, swapping
  - Forwarding table management

MPLS Forwarding Infrastructure

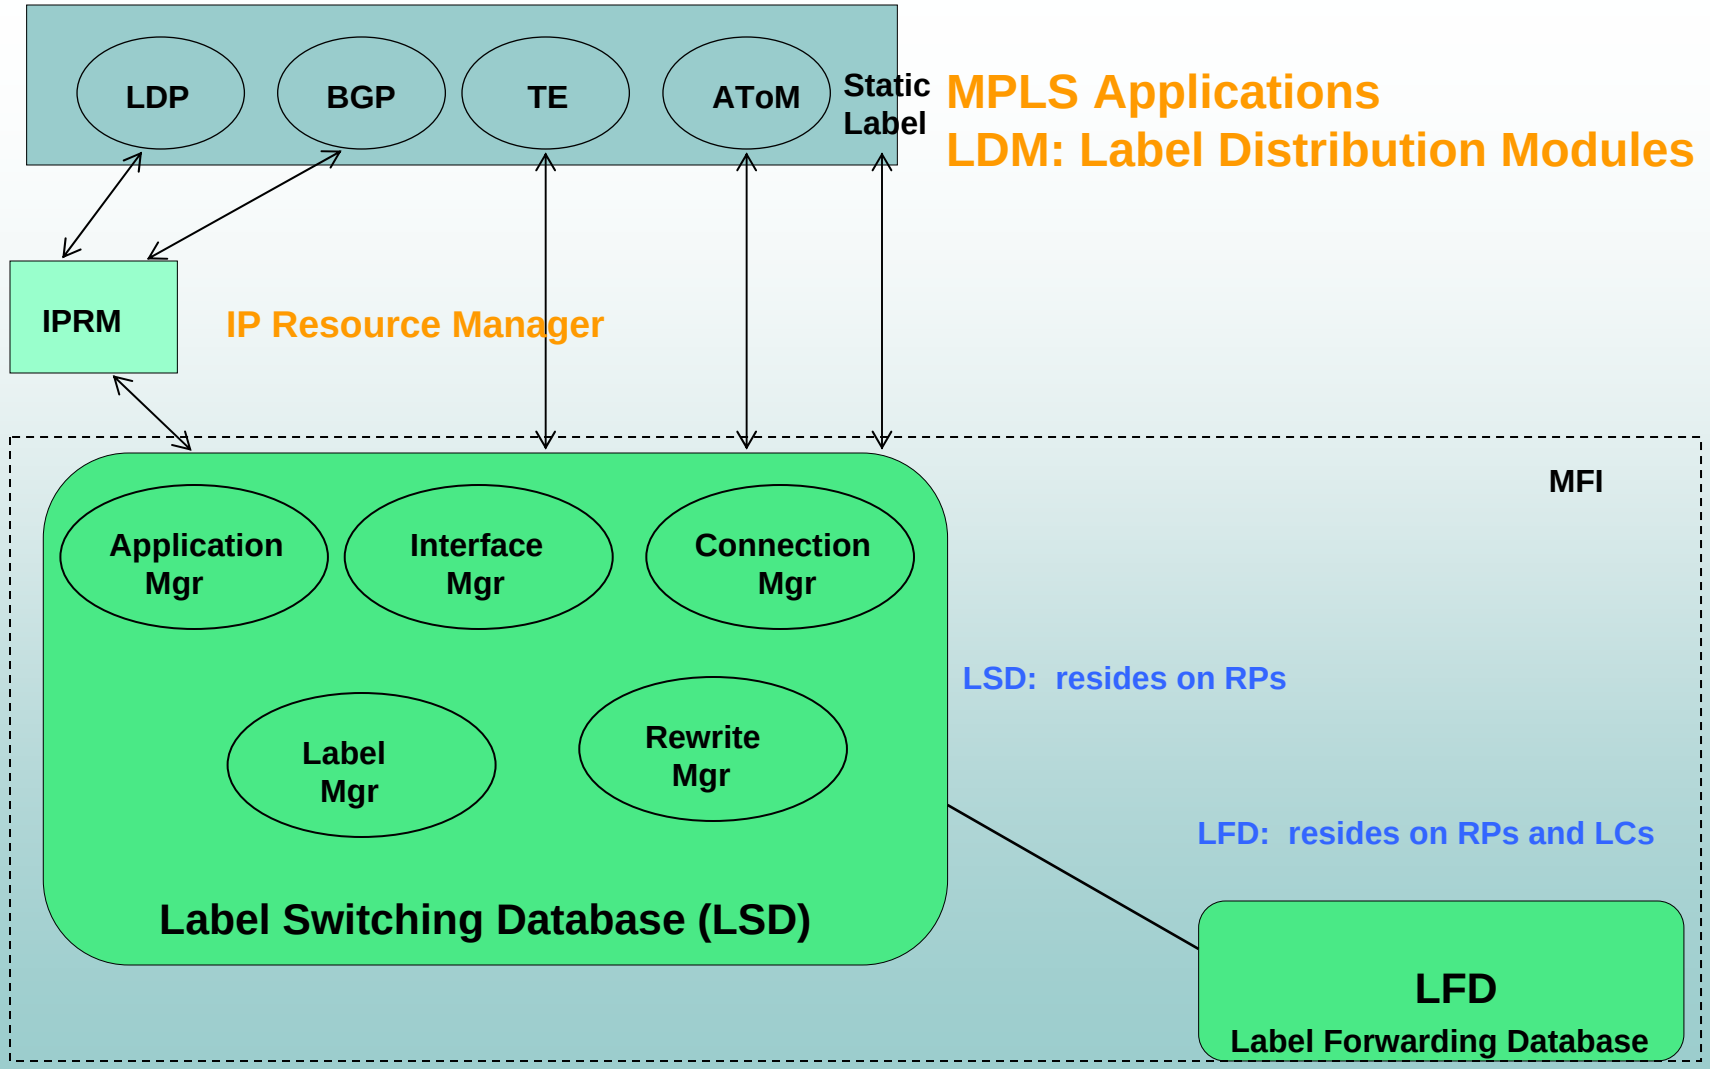

MFI System Architecture

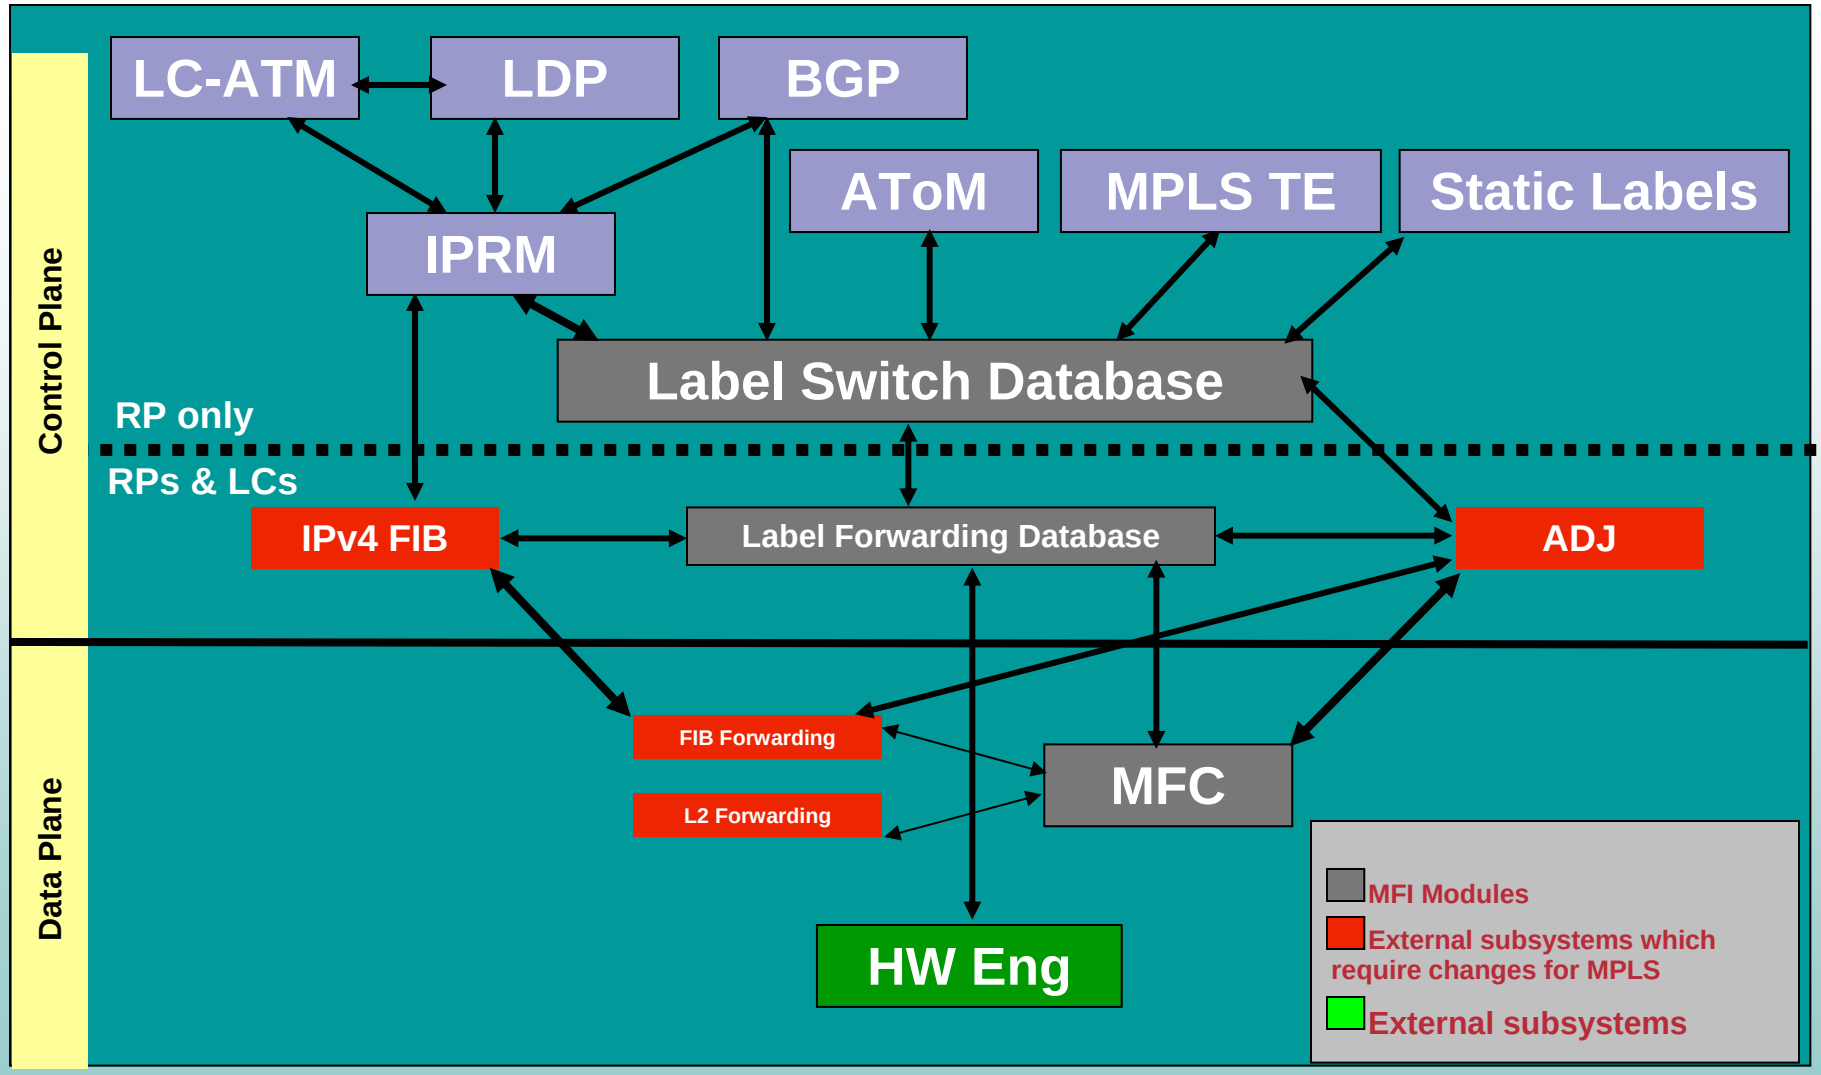

MFI NSF

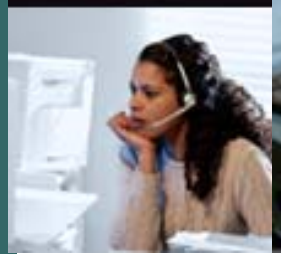

MFI HA NSF support: Before Switchover

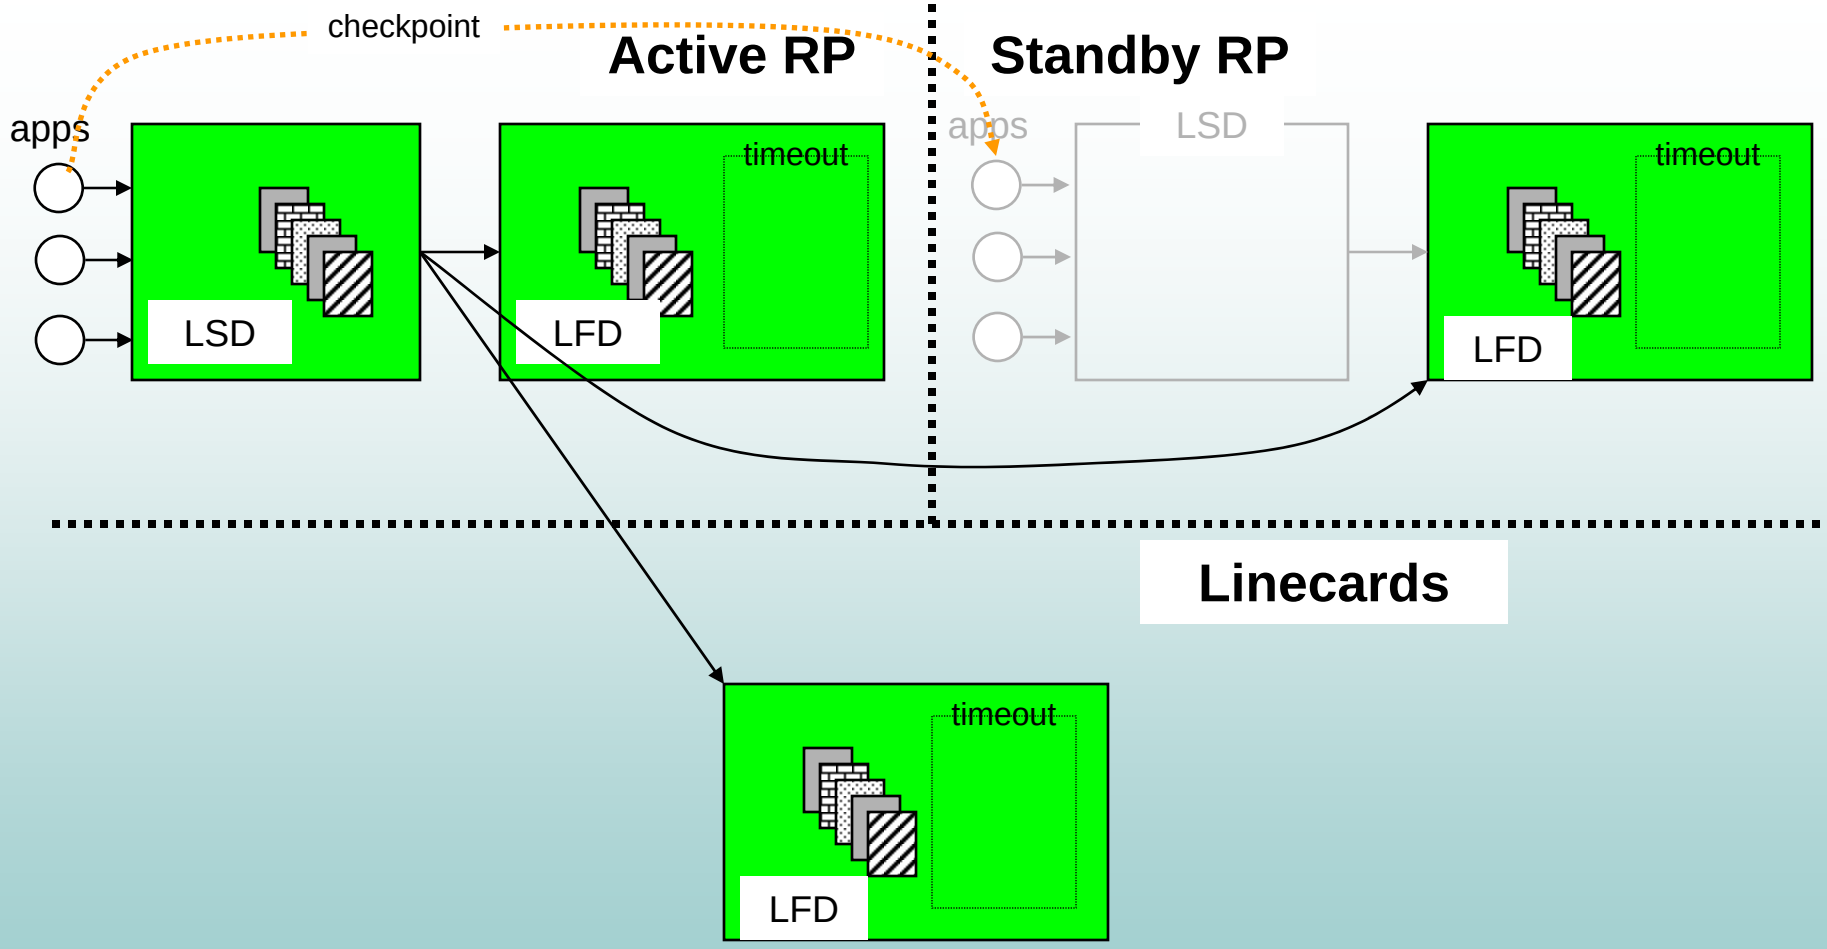

Normal Operation

MFI HA NSF support: Active RP Failed

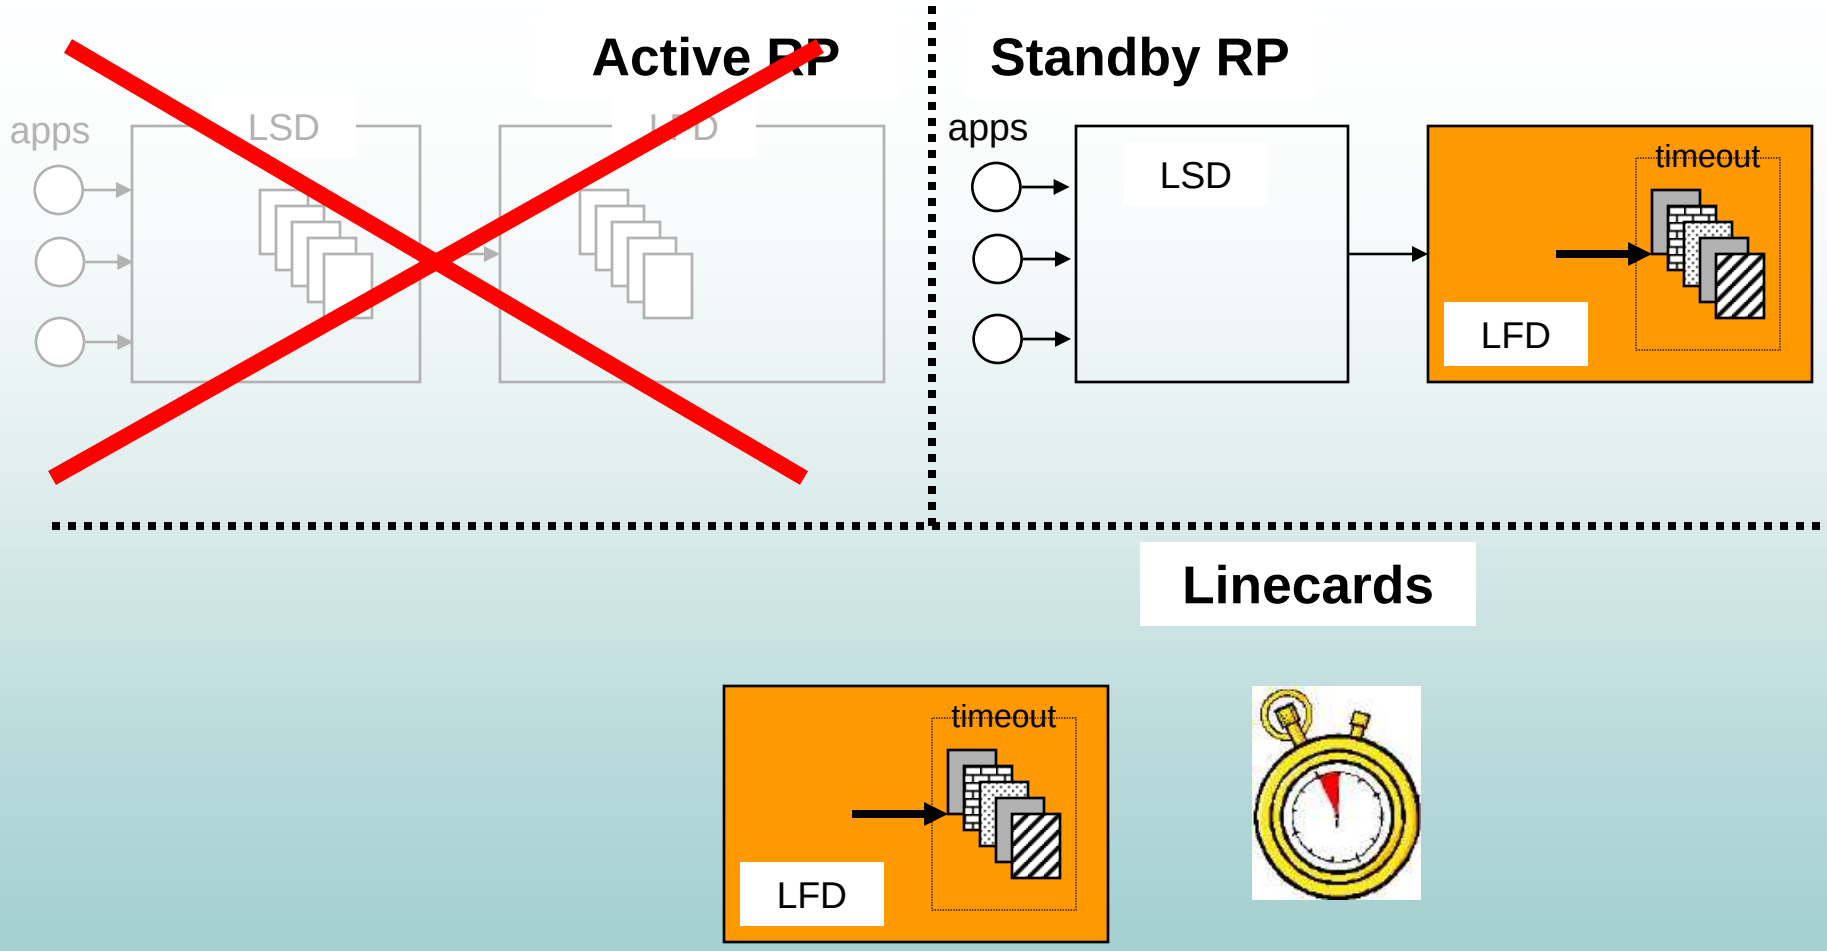

Move the rewrites to timeout queue. timeout wait begins

MFI HA NSF support: timeout-wait

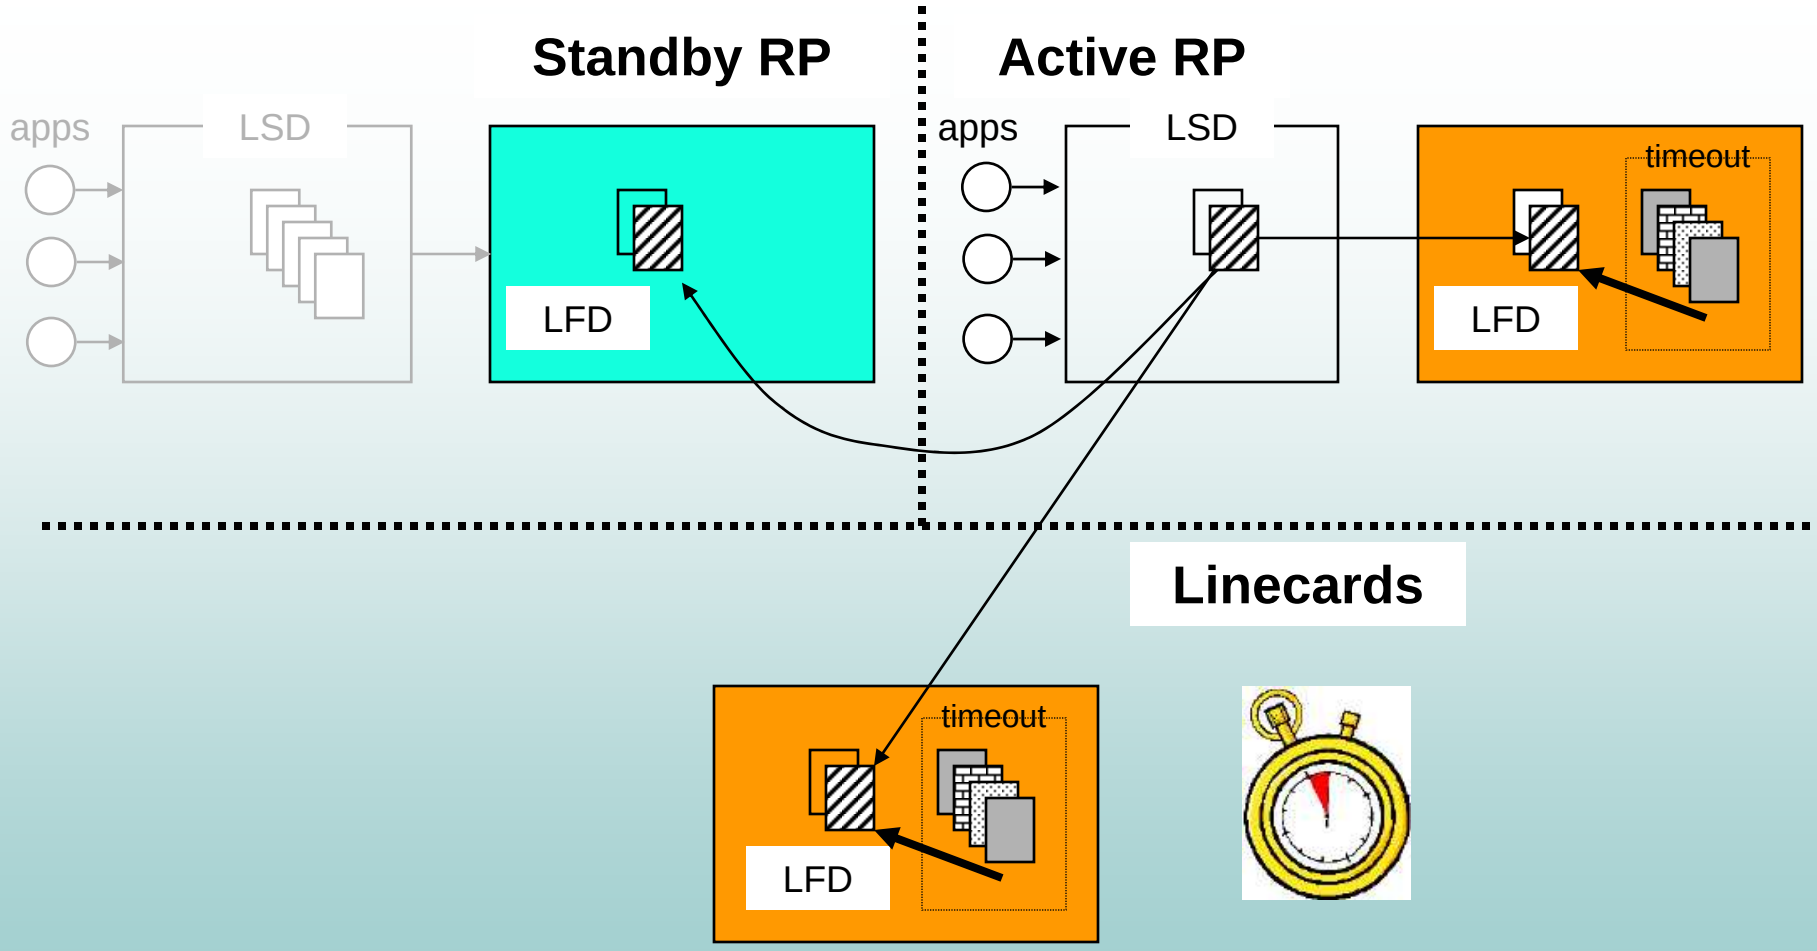

Newly created rewrite rescues the existing matching rewrite from the timeout queue

MFI HA NSF support: timeout ends Normal Operation

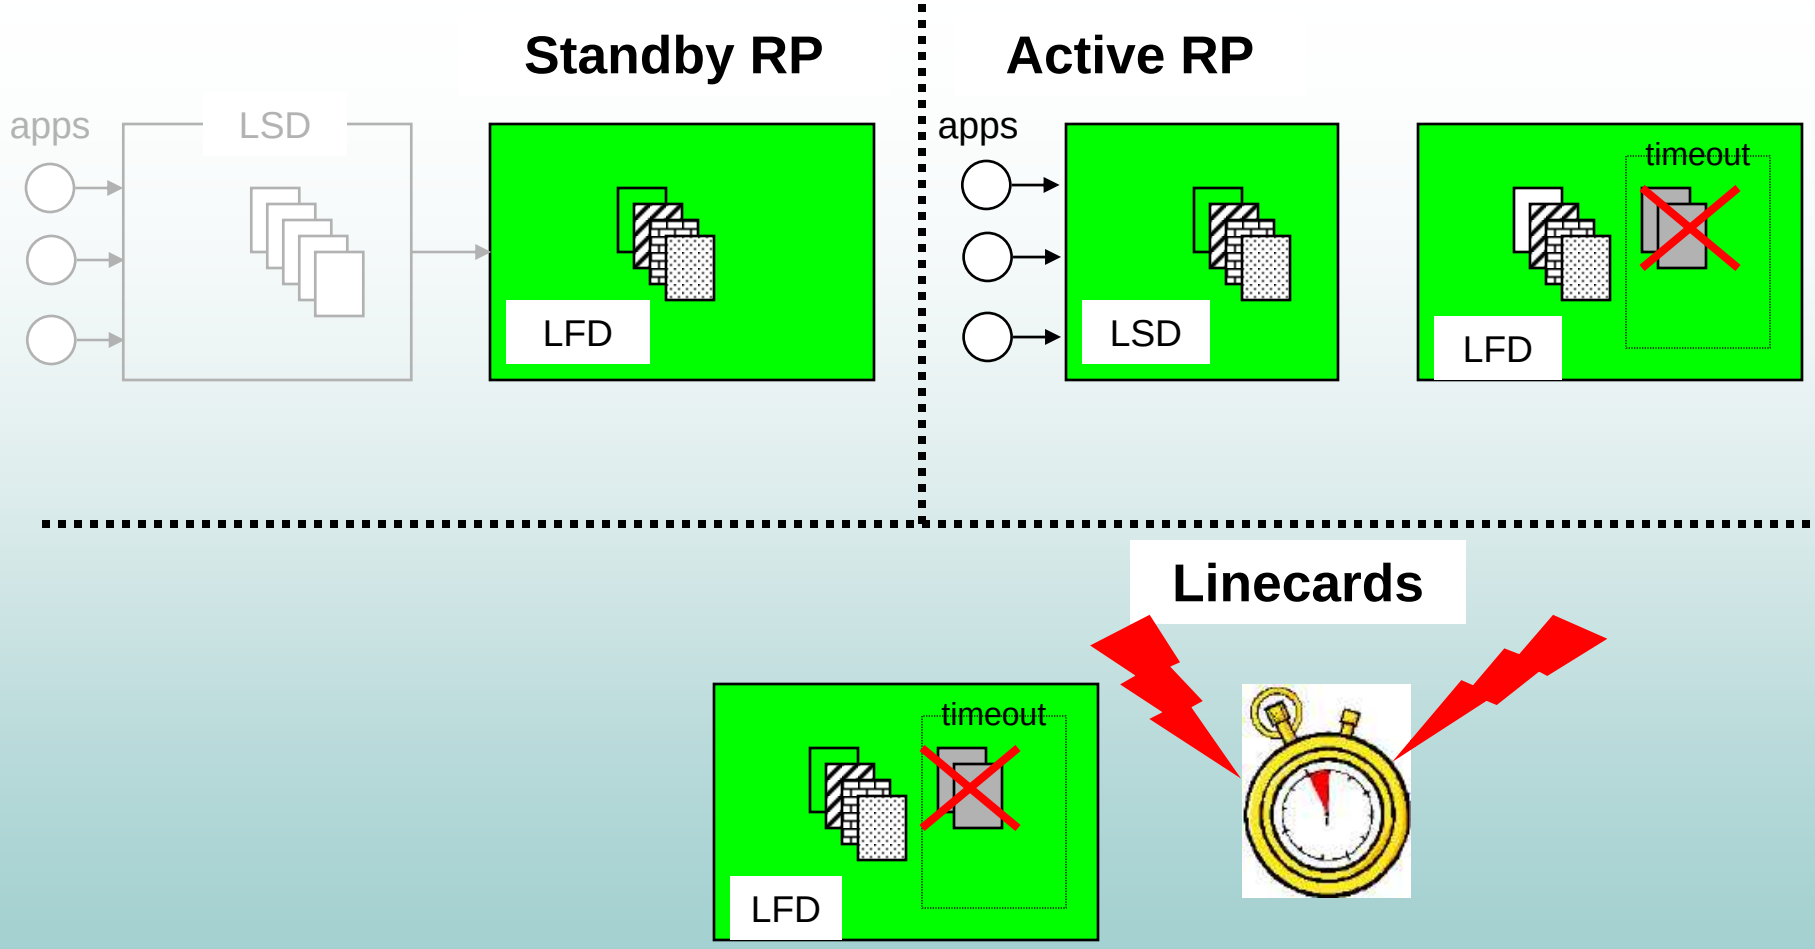

timeout period ends, delete any rewrites remaining on timeout queue

LSD HA Behavior

- **LSD resides on RPs**
  - **Application makes a label request from LSD**
  - **Label manager on active LSD allocates a label, checkpoints label and/or ownership info to the standby LSD, and returns the label to the Application**
  - **Application on the standby LSD can make request for this specific local label, build the rewrite, and hands it to the LSD**
  - **Application builds the rewrite and hands it to the LSD. The rewrite is pushed to LFD which in turn installs it in appropriate forwarding table**
- **LSD behavior after RP switchover**
  - **LSD on new active RP starts a per Application timer.**
  - **Since the Applications on the standby LSD are allowed to connect and install info, they don't need to re-update LSD with local label/rewrite info. This enables a faster recovery after for such applications after switchover .**
  - **Applications normally checkpoints the local label and recovers the remote label from the peers. It needs to relearn remote labels, build rewrites, give it to LSD so that it can be pushed to LFD before LFD timer expires.**
  - **All unclaimed local labels will be timedout by LSD after Apps. timers have expired.**

LSD HA Behavior (Cont.)

- LFD resides on both RPs as well as on line cards.
- LFD behavior after RP switchover:
  - LFD on LC detects communication loss with the LSD
  - Marks rewrites stale and starts a LFD stale timer.
  - Rewrites with NSF clear flag will be deleted after switchover
  - Continue to use the stale rewrites for forwarding (NSF)
  - LFD reconnects to LSD on the new active RP, the LSD will push all rewrites to LFD before the LFD timer expires
  - On LFD timer expiry, any rewrite entry still marked as stale is deleted

MFI HA Summary

If RP crashes, LFD on LC loses contact with LSD on RP:

1. **LFD retains the existing rewrites**
2. **Flags the rewrite as stale and starts the stale timer**
3. **Applications may recover the existing rewrites**
4. **If the stale timer goes off, withdraw the rewrites**
5. **Label Distribution Protocol handles the label changes**
6. **LSD must not reallocate any previously allocated local labels to new requests during switchover**
7. **Allocated local labels are checkpointed and saved across switchover**

Agenda

- Introduction to Cisco IOS High Availability (HA)
- IP Nonstop Forwarding with Stateful Switchover (NSF with SSO)
- MPLS Co-existence with IP NSF with SSO
- MPLS Forwarding Infrastructure (MFI)
- MPLS HA Overview
- MPLS HA – LDP NSF with SSO
- Summary

MPLS High Availability

- **MPLS High Availability features extend Cisco NSF with SSO capabilities for:**
  - **Label distribution protocol (LDP)**
  - **MPLS Forwarding**
  - **Virtual Private Networks (VPNs)**
  - **Traffic engineering and Frame Relay**
  - **AToM**
- **Minimal disruption to MPLS forwarding plane due to route processor control plane failures**
  - **Includes MPLS control plane failures (LDP,BGP, RSVP)**

BENEFITS

- **Greater uptime and lower network outages for MPLS L3VPNS and L2VPNs**

MPLS Control and Forwarding Planes

- **MPLS architecture can be decomposed into two planes:**
  - **Control plane**
  - **Forwarding plane**
- **MPLS control plane**
  - **Distributes labels and establishes label switched paths**
- **MPLS allows multiple control planes:**
  - **LDP**
  - **BGP**
  - **RSVP-TE**
- **MPLS forwarding plane**
  - **Used for actual data packet forwarding**

MPLS Control and Forwarding Planes Summary

- **In many Cisco routers, control/forwarding planes are separate:**
  - **Control plane resides on route processors (RPs)**
  - **Forwarding plane resides on line cards (LCs)**
- **Certain router failures are confined to control plane**
  - **Hardware failure on active RP: switchover to standby**
  - **Software failure on active RP: switchover to standby**

Highly Available Network Requirements

- For system and network high availability, disruption in data plane must be kept to an absolute minimum
- Separation of control and forwarding plane should allow forwarding to continue while control plane recovers (NSF)

Why MPLS HA: Problem Description

- TCP session between two LDP/BGP peers may fail for various reasons, such as RP switchover due to HW/SW failures
- LDP and BGP use TCP as a reliable transport mechanism for its protocol messages
- On detection of TCP session failure (hardware or software failure), existing LDP and BGP control plane components would disrupt their forwarding state

Continuous Forwarding of Traffic During Control Plane Service Disruption

Before Cisco NSF/SSO Support and Graceful Restart

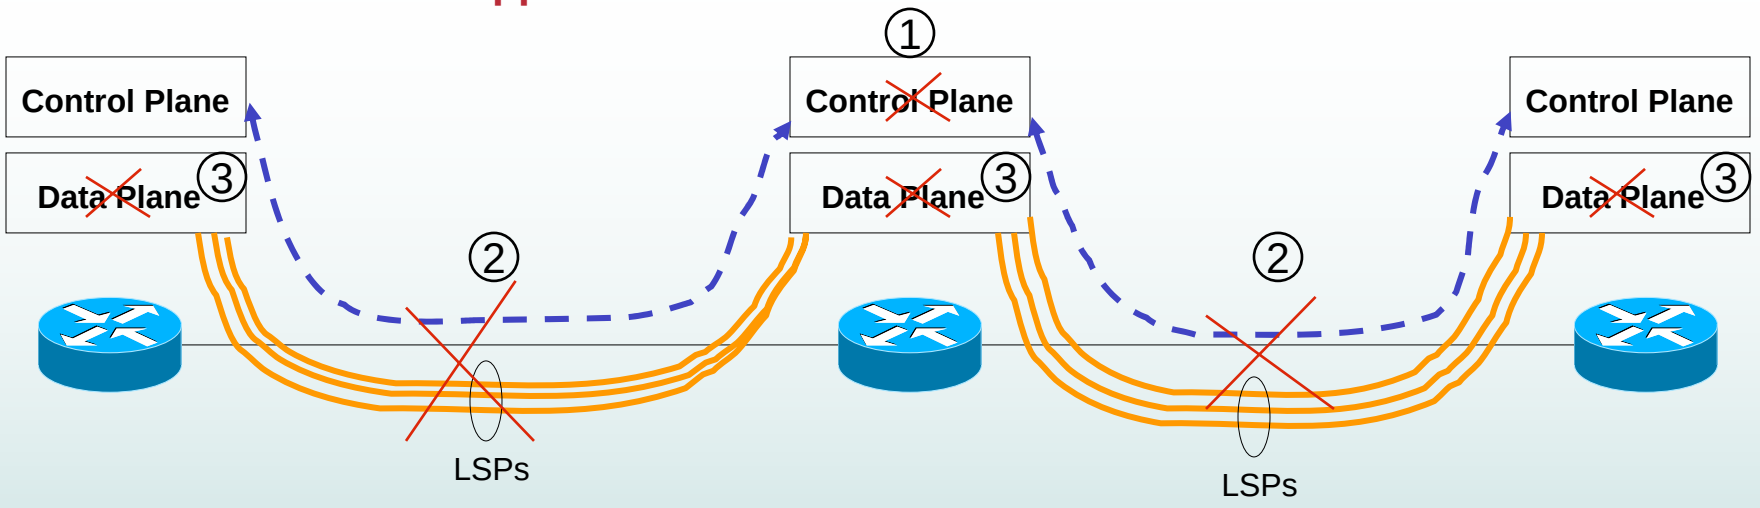

After Cisco NSF/SSO Support and Graceful Restart

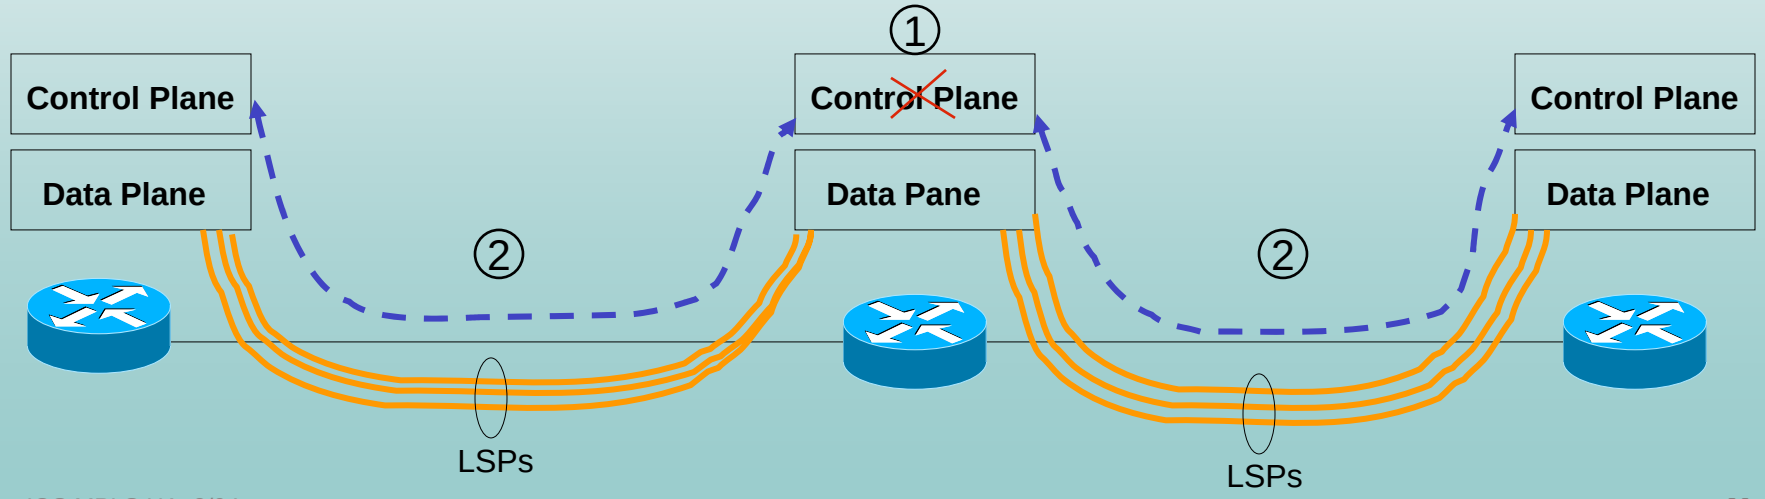

Why MPLS HA: Solution

- **Enhance the key protocols used in MPLS Control plane to minimize disruption in MPLS forwarding plane due to MPLS control plane restart**
- **Protocol enhanced:**

  LDP

  MP-BGP

  RSVP-TE

High Availability Umbrella

Cisco NSF with SSO for High Availability

IP HA

Local Checkpoint + Graceful Restart

MPLS HA

Local Checkpoint +Graceful Restart

High Availability Umbrella

Cisco NSF with SSO for High Availability

IP HA

	EGP
	eBGP
IGP	OSPF
OSPF	EIGRP
IS-IS	RIP

MPLS HA

LDP
BGP
RSVP-TE
AToM
LC-ATM

Deploying MPLS HA

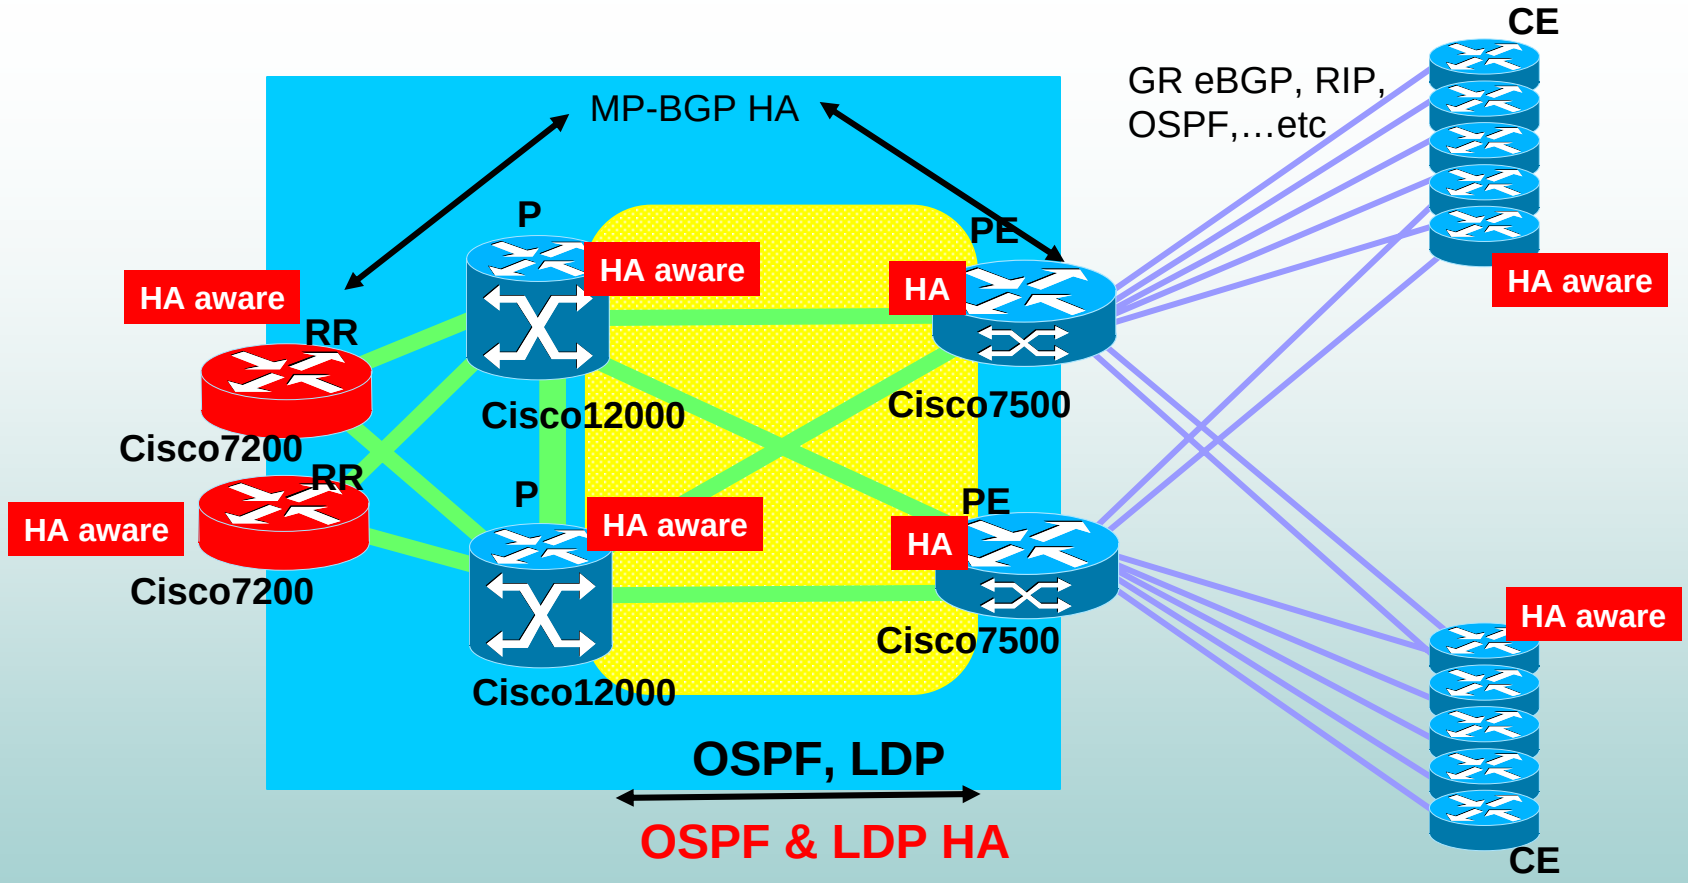

HA aware devices must support Graceful Restart for deployed Routing Protocol.

MPLS HA Summary

1. Active RP checkpoints related information after the MPLS forwarding is updated to back up RP
2. Graceful Restart is executed
3. Active RP again checkpoints related information after the MPLS forwarding is updated

Agenda

- Introduction to Cisco IOS High Availability (HA)
- IP Nonstop Forwarding with Stateful Switchover (NSF with SSO)
- MPLS Co-existence with IP NSF with SSO
- MPLS Forwarding Infrastructure (MFI)
- MPLS HA Overview
- MPLS HA – LDP NSF with SSO
- Summary

MPLS HA – LDP NSF with SSO

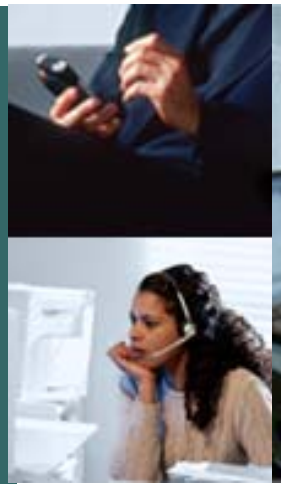

Why LDP Graceful Restart?

- **Allows a router to recover from disruption in control plane service without losing its LDP bindings and MPLS forwarding state**

LDP HA Key Elements

1. **Checkpoint local label bindings to backup RP**
2. **Execute LDP Graceful Restart**

   Allows a router to recover from disruption in control plane service without losing its LDP bindings and MPLS forwarding state

3. **Checkpoint the refreshed/new local label bindings**

Why LDP Checkpointing?

- **LSRs that support LDP require that the restarting control plane be able to query the forwarding plane for local label binding information**
  - **MPLS forwarding plane will not support such queries**
- **LDP checkpointing is the mechanism that provides this functionality**
  - **Makes local label bindings available to the restarting (standby) LDP control plane**

MPLS LDP Local CheckPointing Key Points

- Enabled by default via a separate process
- Copies active RP's LDP local label bindings to the backup RP
  - All labels are copied from active to backup RP in the first round
- Periodic incremental updates
  - Reflect new routes that have been learned
  - Indicate the removal of routes
  - Allocation of labels
- Checkpointing stops during control plane disruptions, GR, and recovery
- Label bindings on backup RP are marked as checkpointed
- This marking is removed when it becomes active RP

LDP HA Approaches

- **Three approaches have been proposed to minimize disruption of MPLS forwarding due to LDP restart:**
  - **Fault Tolerance for the LDP: draft-ietf-mpls-ldp-ft-06.txt**
  - **Checkpointing Procedures for LDP: a slight variation of FT-LDP**
  - **Graceful Restart Mechanism for LDP RFC3478**

LDP HA Approaches Comparison

- **Similarities**
  - **Fault Tolerant Session Type Length Value in the LDP Initialization message**
  - **Allow data forwarding to continue while LDP recovers**
- **Differences**
  - **Fault Tolerance-LDP**
    - **Requires preservation of both LDP and MPLS forwarding state across LDP restart**
    - **Requires exchange of only unacknowledged LDP state after LDP restart**
    - **Covers both downstream unsolicited (DU) & downstream on demand (DoD) modes**
  - **Graceful Restart-LDP**
    - **Requires preservation of only MPLS forwarding state across LDP restart.**
    - **Requires exchange of all LDP-related state information after LDP restart**
    - **Covers only the DU mode (Cisco to extend it for DoD)**

LDP HA Approaches Solution

- FT-LDP is obsolete
- FT-LDP required too many checkpoints including remote checkpoints
- Cisco supports Graceful Restart mechanism

LDP Graceful Restart Mechanism

- LSR sends the LDP initialization message to a neighbor to establish an LDP session
- Fault Tolerant (FT) session type length value (TLV) is included in the LDP initialization message

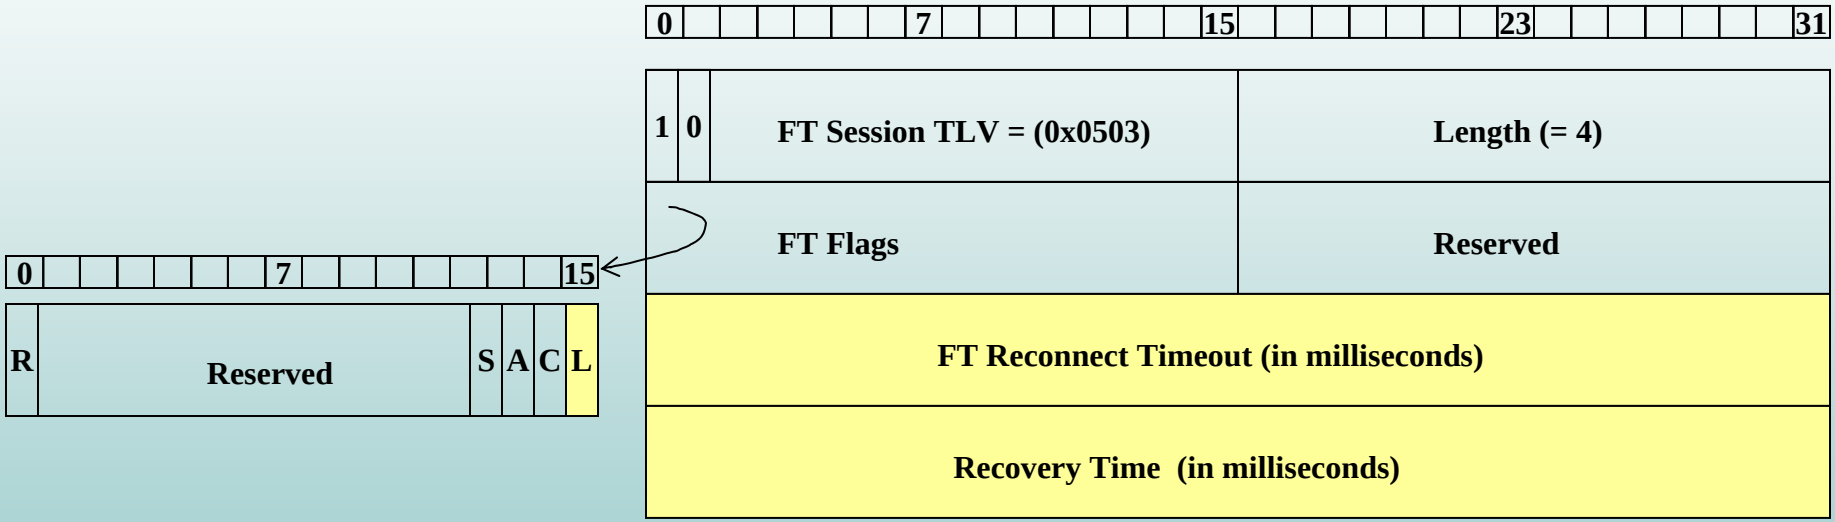

L=1 means GR-LDP is selected
FT Reconnect = 0 means LSR is not NSF capable
FT Recovery Time = 0 means LSR was unable to preserve MPLS forwarding state across restart

FT Session TLV Fields

- **Learn from Network (L):**

  Flag is set to 1

  LSR is configured to perform MPLS LDP GR

  Remainder of the FT fields are set to 0 and ignored by the receiving LSRs

- **FT Reconnect Timeout**

  Time that the sender of the TLV would like the receiver of that TLV to wait after the receiver detects the failure of LDP communication with the sender

  Setting the FT Reconnect Timeout to 0 indicates that the sender of the TLV will not preserve its forwarding state across the restart, yet the sender supports LDP GR

- **FT Recovery time**

  Field shows the duration (the LSR should retain the MPLS forwarding state that it preserved across the restart during a reconnection

  If an LSR did not preserve the MPLS forwarding state before the restart of the control plane, the LSR sets the recovery time to 0

LDP Graceful Restart Key Points

- **MPLS LDP GR must be enabled before an LDP session is established**
- **Works with LDP sessions between directly connected and non-directly connected peers (targeted sessions)**
- **LDP GR supports both failure cases:**
  - **LDP Restarts**
    - **Stateful Switchover event interrupts LDP communication with all LDP neighbors**
    - **Backup router retains MPLS forwarding state and reestablishes communication with the LDP neighbors**
    - **Backup router ensures that the MPLS forwarding state is recovered.**
  - **LDP Session Resets**
    - **Individual LDP session has been interrupted, but not due to a Stateful Switchover event; rather, it may have been due to a TCP or UDP communication problem**
    - **LSR associates a new session with the previously interrupted session**
    - **The LDP and MPLS forwarding states are recovered when the new session is established**

LDP Graceful Restart Key Points (Cont.)

- **Timers can be set to limit the LDP session reestablishment time to restart the router**
- **If an LSR doesn’t have LDP GR and it attempts to establish and LDP session with an LSR that is configured with MPLS LDP Graceful Restart, the following actions occur:**
  1. **LSR that is configured with MPLS LDP Graceful Restart sends an initialization message that includes the FT TLV value to the LSR, which is not configured with MPLS LDP Graceful Restart**
  2. **LSR that is not configured for MPLS LDP Graceful Restart receives the LDP initialization message and discards the FT TLV value**
  3. **Two LSRs create a normal LDP session but do not have the ability to perform MPLS LDP Graceful Restart**

LDP Graceful Restart Process Summary

1. Is LDP GR supported on LSRs - negotiate Restart capabilities?
2. Update bulk of current label binding info from active to backup RP
3. Retain old forwarding plane info – LDP bindings - during (FT Reconnect) timeout
4. Restarting/Recover

LDP GR on SSO-LSR and SSO-aware Peer

- **SSO capable (Restarting peer) LSR:**
  - **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 (Receiving peer) LSR:**
  - **LDP failure detected.**
  - **Mark the forwarding state as stale, and retain it**
  - **Reestablished LDP session**

LDP Graceful Restart Operation Overview

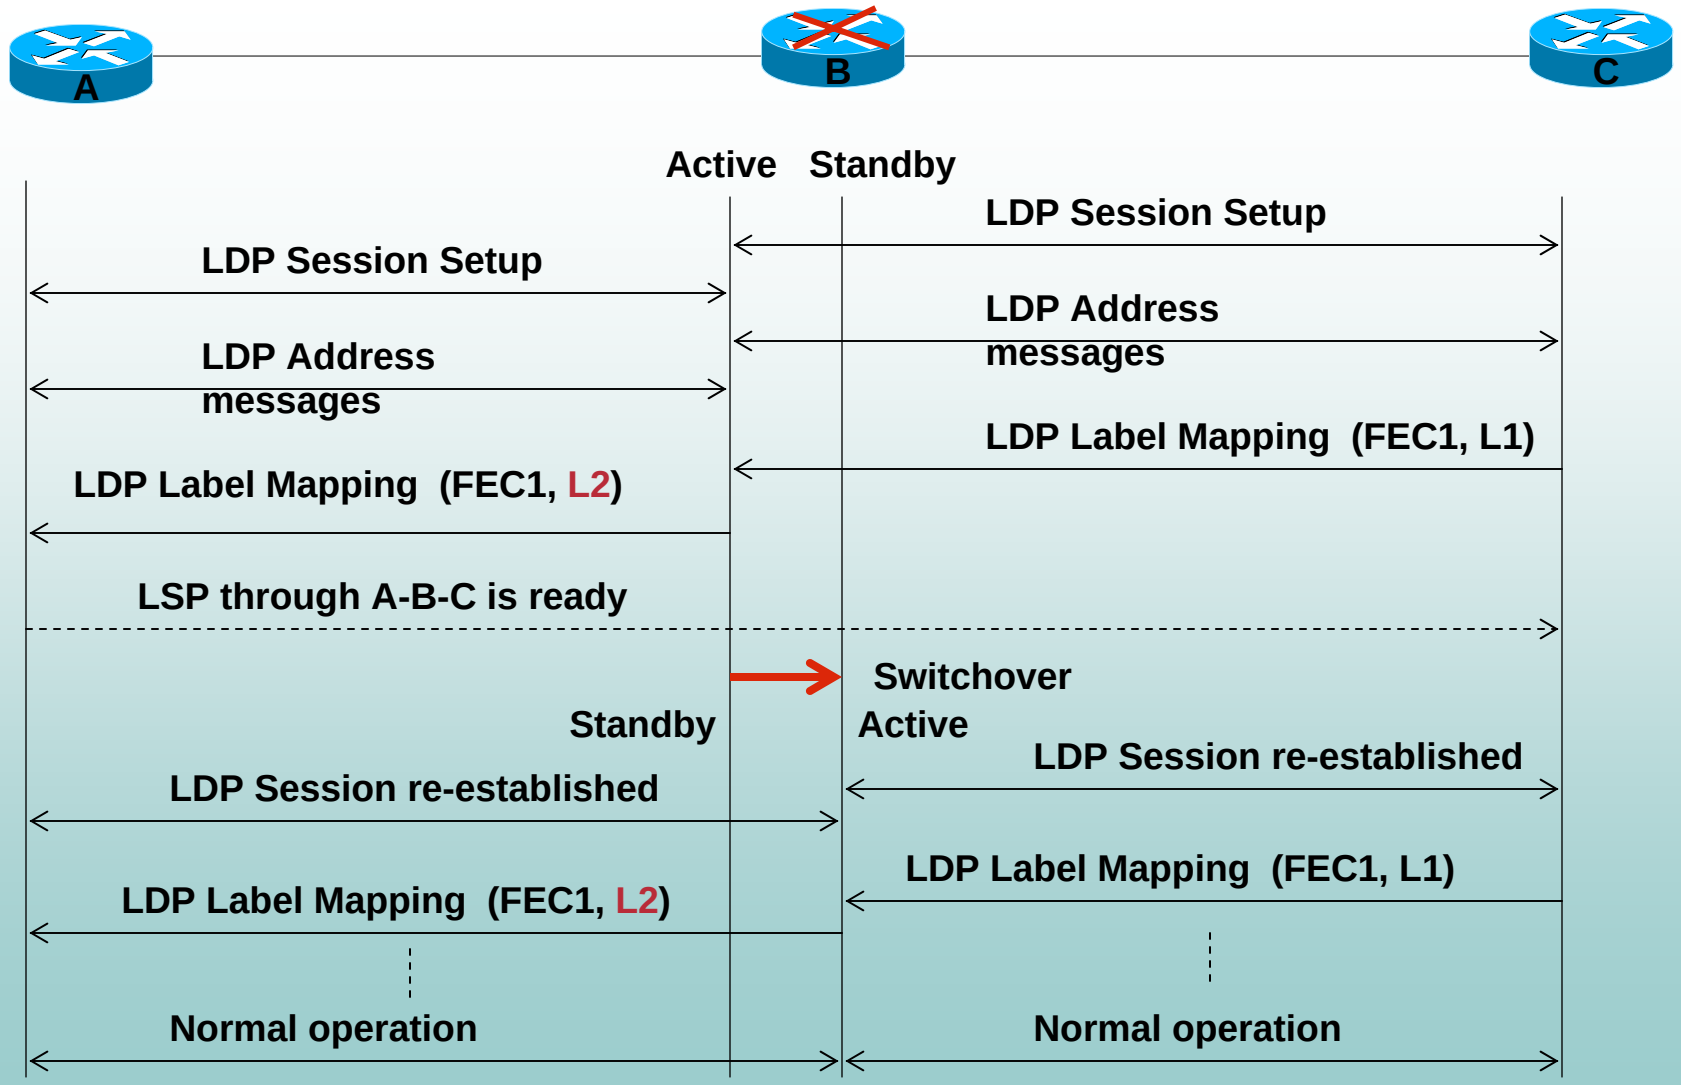

LDP Graceful Restart Operation

- 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 fits forwarding entries as stale. As 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 GR Security Considerations

- **LSR with LDP GR capabilities is subject to additional DoS attacks:**

  An intruder may impersonate an LDP peer in order to force a failure and reconnection of the TCP connection, where an intruder sets the Recovery Time to 0 on reconnection. This will force all labels received from the peer to be released.

  An intruder could intercept the traffic between LDP peers and set Recovery Time to 0. This will force all labels received from the peer to be released.

- **Use MD5 authentication between LDP peers**

LDP Graceful Restart Restrictions

- **Tag Distribution Protocol (TDP) sessions are not supported.**
- **MPLS LDP Graceful Restart cannot be configured on label controlled ATM (LC-ATM) interfaces**

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 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 Troubleshooting

Cisco.com

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

Debug mpls ldp graceful-restart

- **Shows events and errors related to MPLS LDP Graceful Restart**

```
Router# debug mpls ldp graceful-restart
```

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 & remote bindings…**

```
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 & Info. 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
```

Debug mpls ldp checkpoint

- **Shows events and errors related to LDP checkpointing**

show mpls ldp checkpoint

- Displays current check pointing data on the active RP
- Status of local label bindings: this list will explain the different states of the local label bindings

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
- One the standby, it deletes all LIB entries learned via checkpointing

Agenda

- Introduction to Cisco IOS High Availability (HA)
- IP Nonstop Forwarding with Stateful Switchover (NSF with SSO)
- MPLS Co-existence with IP NSF with SSO
- MPLS Forwarding Infrastructure (MFI)
- MPLS HA – LDP NSF with SSO
- Summary

Summary

- Cisco is enhancing its portfolio to add features for improved full HA solution
- MPLS HA features provide stateful switchover and NSF capability for VPN, LDP, TE, etc
- MPLS VPN HA requires MFI HA, LDP HA, BGP HA
- Need IP HA enabled to support MPLS HA

  GR must be enabled on all participating RPs (OSPF, BGP, IS-IS) on P, PE, and CE routers
- HA Capable system is enabled with full Cisco NSF with SSO

  Peers only need to support Graceful Restart

Summary

- AToM Cisco NSF with SSO is exactly the same as directed LDP
- AToM application will only check for local labels
- TE Cisco NSF with SSO is defined in GMPLS-RSVP(TE) RFC 3473
- FRR Cisco NSF with SSO is same as NSF/SSO for TE
- LC-ATM Cisco NSF with SSO is DoD LDP GR

References

- www.cisco.com/go/mpls