Cisco IOS™

High Availability Overview

Steve Koretsky (skoretsk@cisco.com)
High Availability Product Manager
March, 2011

Agenda

- Cisco IOS High Availability strategy
- Key Cisco IOS High Availability technologies
- Summary

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

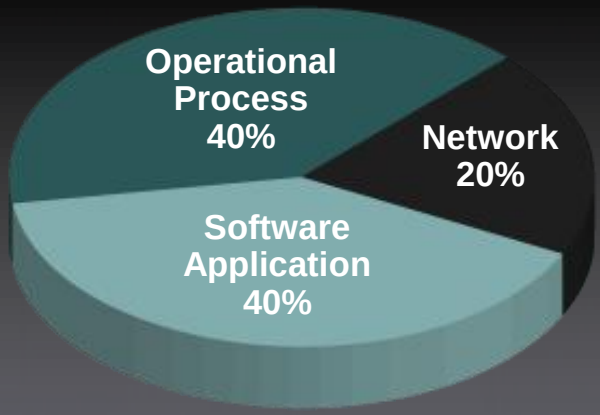

Sources of Network Downtime *

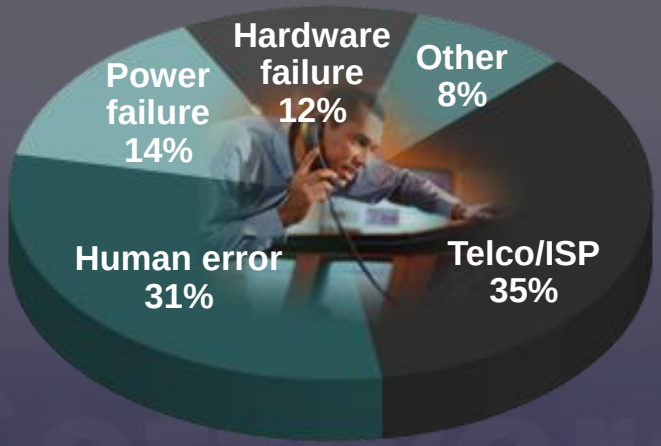

Common causes of Enterprise Network Downtime **

* Source: Gartner Group

** Source: Yankee Group The Road to a Five-Nines Network 2/2004

Cisco IOS High Availability Strategy: Based on Customer Needs

Provide continuous access to applications, data, and content anywhere, anytime

System Level Resiliency

- Reliable, robust hardware
- Cisco IOS software that mitigates fault impact

Network Level Resiliency

Cisco IOS features for faster Convergence, protection, and restoration

Embedded Management

Embedded IOS intelligence for proactive fault/events, configuration, and availability tracking

Address every potential cause of downtime with functionality, design, or best practice

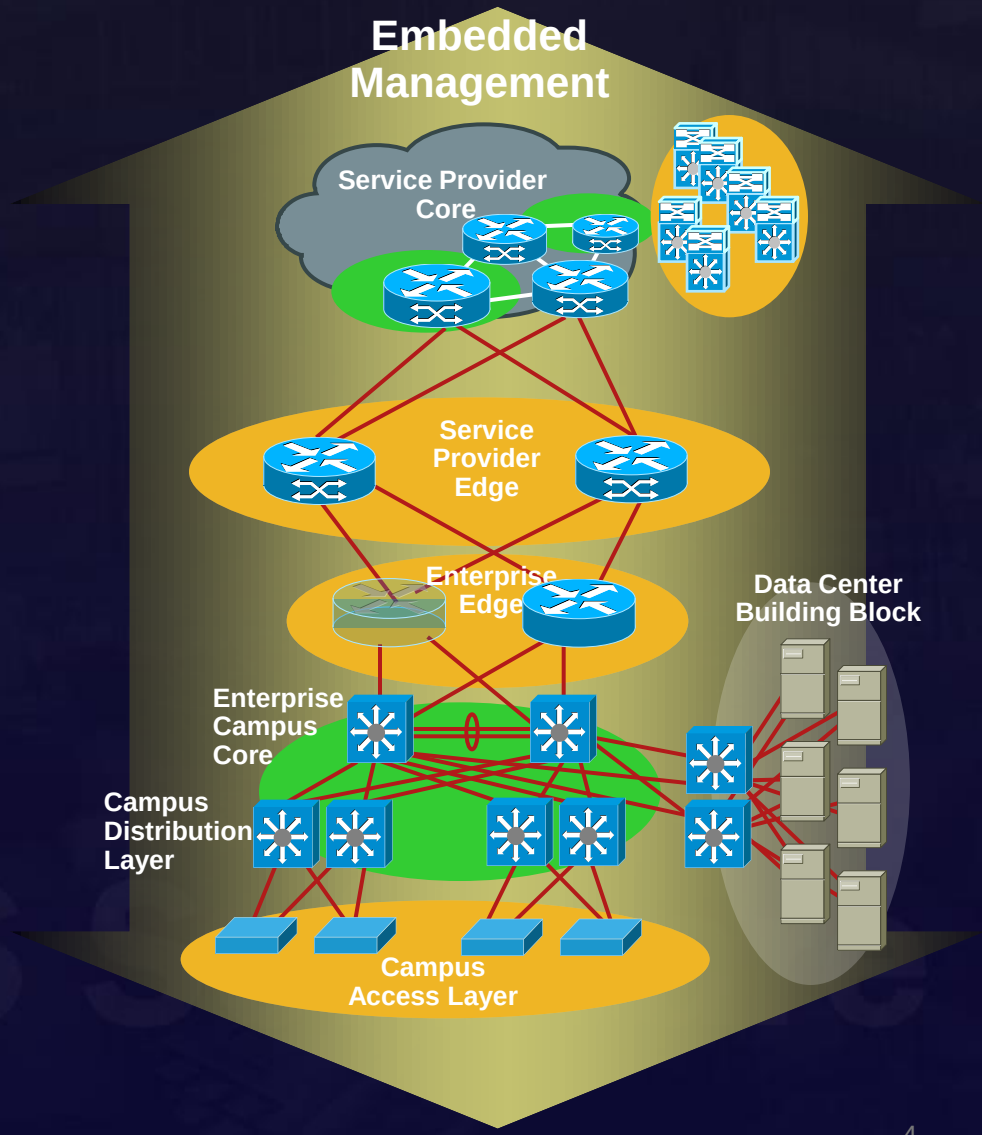

Mitigating the Exposure: Targeting Downtime

Most common causes of downtime

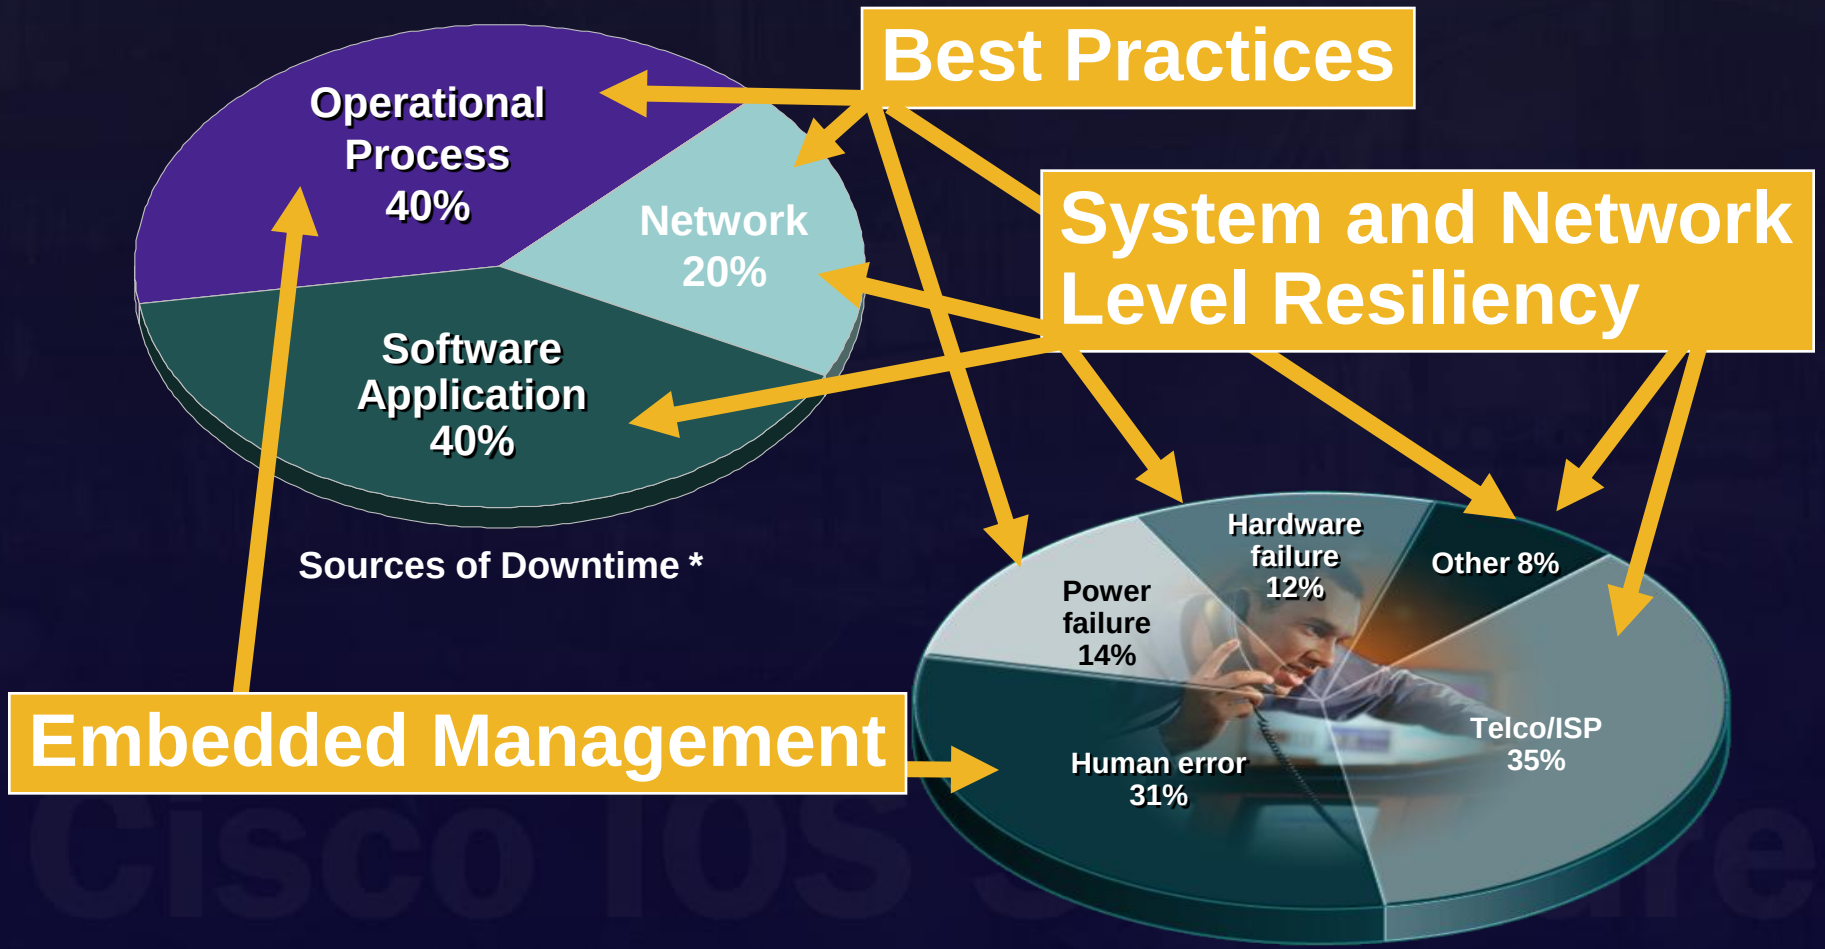

CISCO IOS HIGH AVAILABILITY TECHNOLOGIES

Cisco IOS Software

Industry's Broadest Portfolio of End-to-End High Availability Technologies

Requirements	Technologies
System-Level Resiliency	▪ In-Service Software Upgrade (ISSU) ▪ IP NSF/SSO ▪ MPLS NSF/SSO—LDP, VPNs (including Inter-AS and CsC) ▪ IOS Software Modularity ▪ BGP Nonstop Routing ▪ Line Card Redundancy with Y-Cables ▪ Gateway Load Balancing Protocol ▪ Stateful NAT ▪ Stateful IPSec ▪ Warm Reload ▪ Warm Upgrade ▪ Control Plane Policing
Network- Level Resiliency	▪ NSF/GR Awareness (BGP, OSPF, IS-IS, EIGRP, LDP, RSVP) ▪ Routing Convergence Enhancements BGP Optimization Incremental SPF optimization IP Event Dampening ▪ Multicast Sub-second Convergence ▪ MPLS Fast Reroute ▪ Fast Convergence (OSPF, IS-IS) ▪ Bi-Directional Forwarding Detection (BFD)
Embedded Management	▪ Embedded Event Manager ▪ Component Outage On-Line (COOL) ▪ Embedded Resource Manager (ERM) ▪ Configuration Rollback

System Level Resiliency Overview

Eliminate single points of failure for hardware and software components

Control/data plane resiliency

- Separation of control and forwarding plane
- Fault isolation and containment
- 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

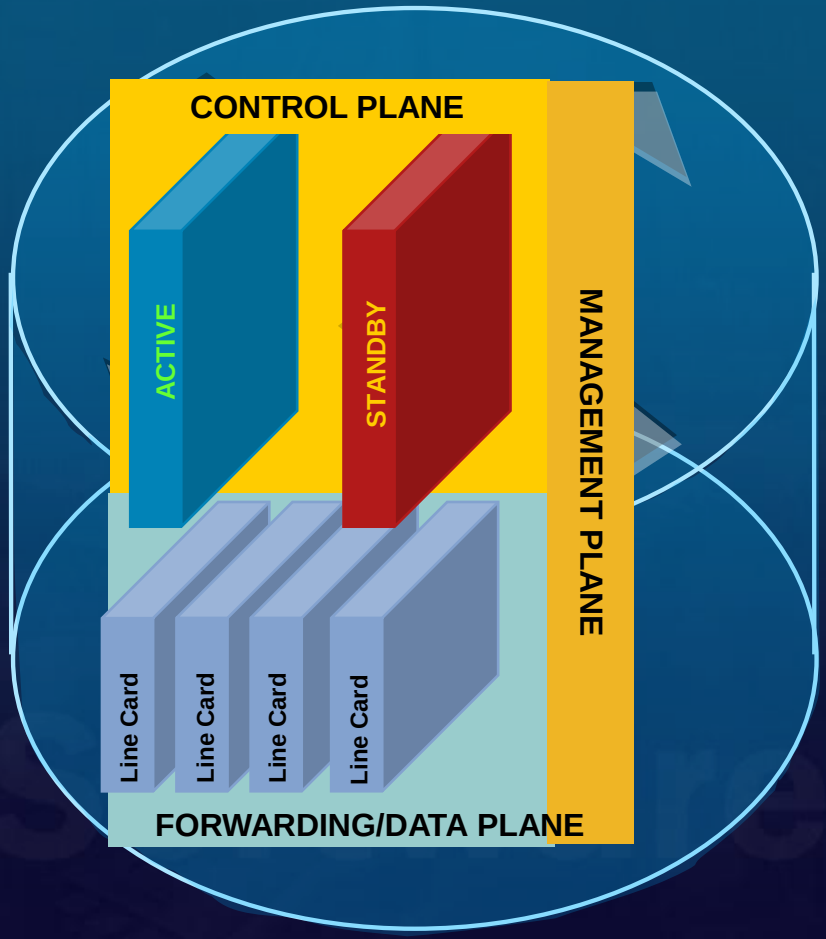

Cisco Nonstop Forwarding with Stateful Switchover

- Network edge is critical
  - Service Provider edge, Enterprise edge
  - Often a single point of failure
- Cisco NSF/SSO increases availability at key edge points
  - Service Provider / Enterprise edge
  - Data center and special distribution blocks
  - Enterprise Campus access

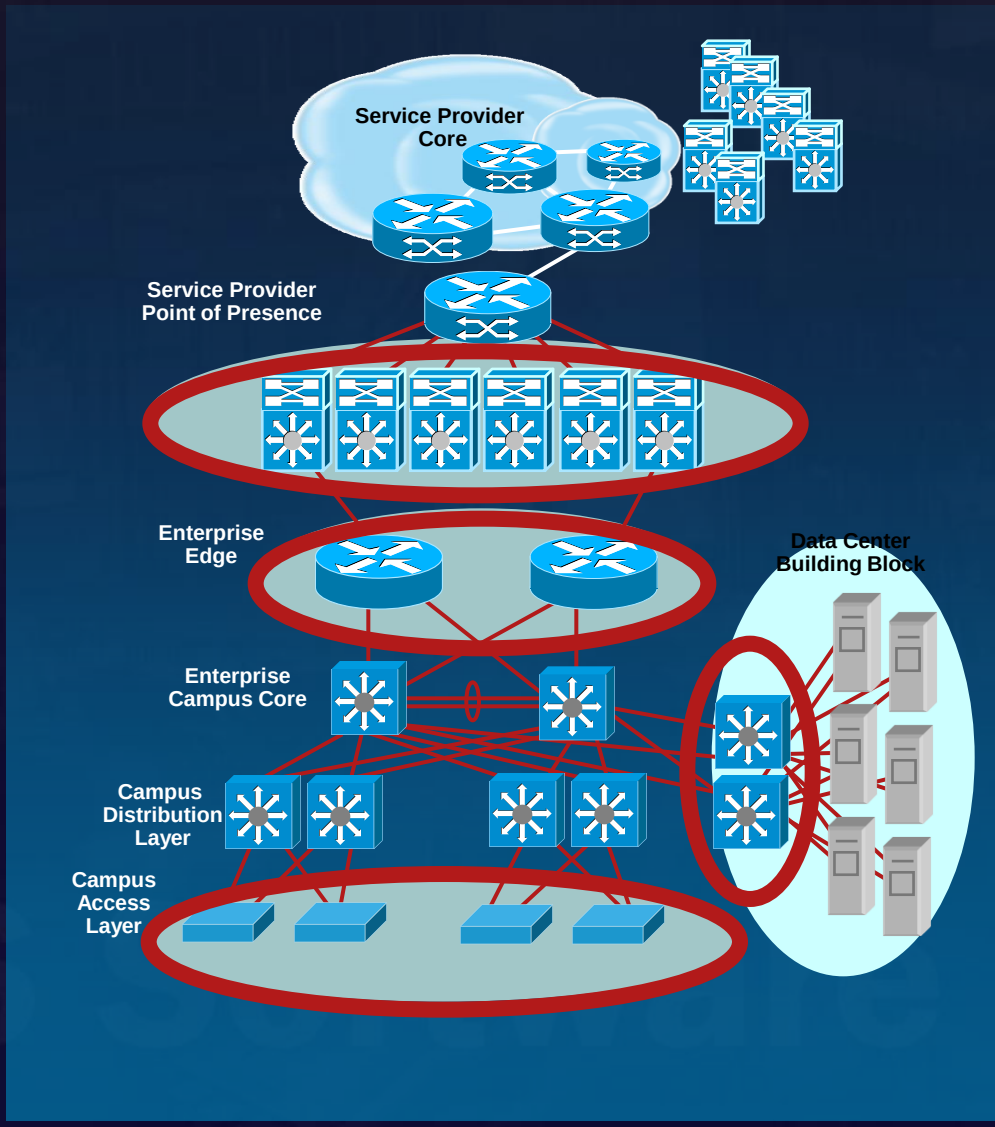

Cisco Nonstop Forwarding with Stateful Switchover

- Network edge is critical
  - Service Provider and Enterprise edge
  - Often a single point of failure
- NSF/SSO increases availability at key edge points
- Employs dual Route Processors
- Cisco SSO maintains connectivity for Layer 2 protocols
- Nonstop forwarding of packets while control plane is reestablished and routing information is validated
  - Packet forwarding continues using current forwarding information base (FIB)
  - Layer 3 (BGP, OSPF, IS-IS) recovers routing information from neighbors, rebuilds routing information base (RIB) and updates FIB

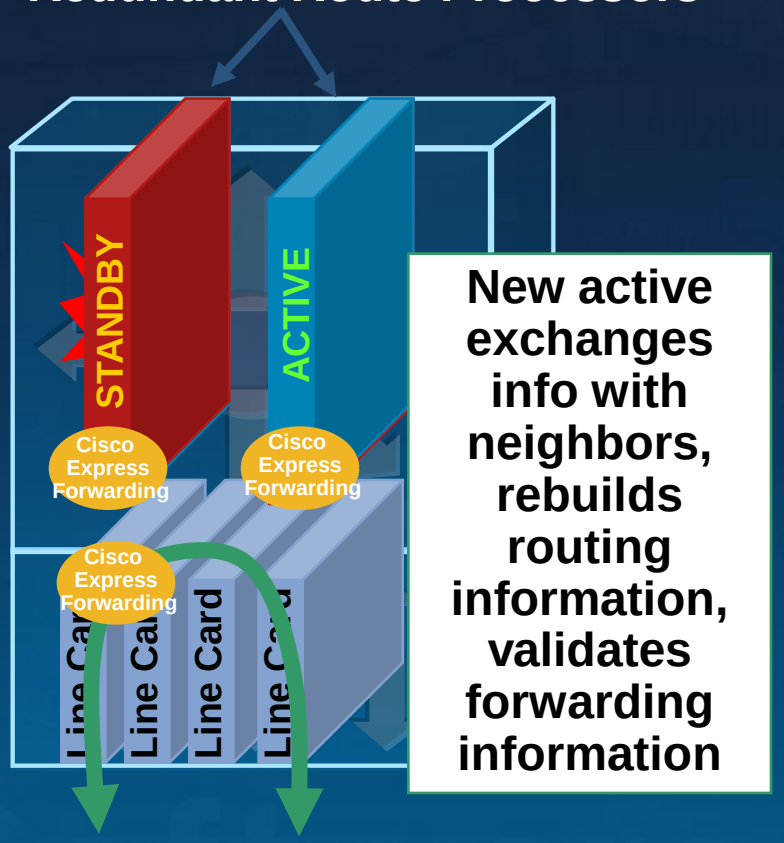

Cisco IOS In-Service Software Upgrade

Industry's First, Comprehensive, In-Service Software Upgrade solution for the IP/MPLS edge router

- Ability to upgrade/downgrade complete Cisco IOS software image
  - With no impact to control plane and minimal impact to packet forwarding
- Leverages dual route processor Cisco NSF with SSO architecture
- Comprehensive upgrade solution covering maintenance-fixes as well as new features
  - Rapid, non-disruptive deployment of new features and services
- Eliminates planned downtime and reduces operational expenses
  - Ability to streamline and minimize planned downtime windows
- Available on:
  - C4500, ASR1000, C10000

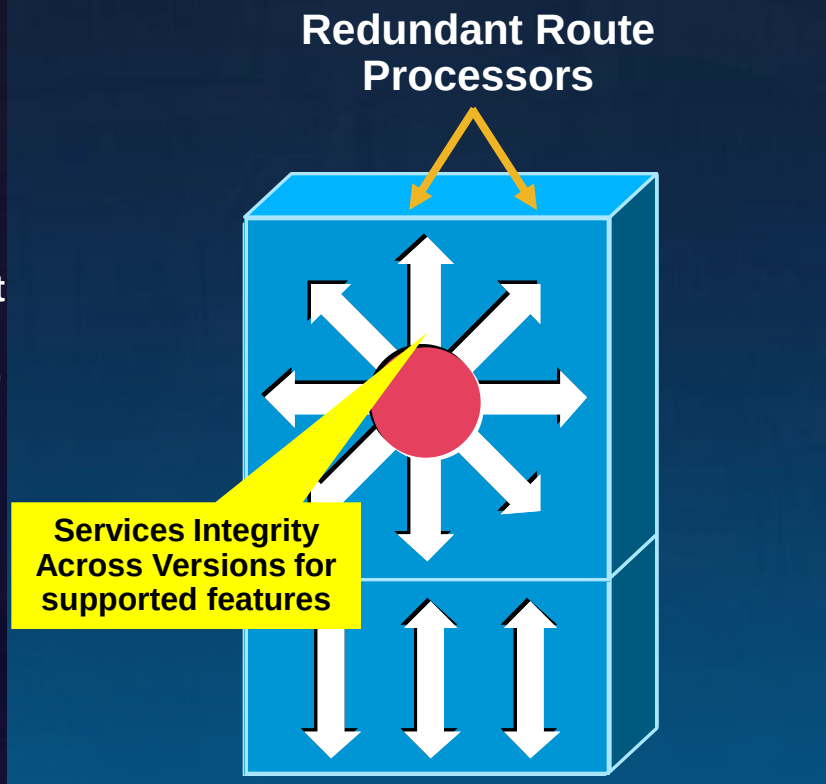

Minimum Disruption Restart (MDR) for Line Card Upgrade with Minimal Impact

Other types of Software Upgrade

Fast Software Upgrade – FSU

- Relies on Dual Processors
- Built on RPR (Route Processor Redundancy)
- Preloads the new software image on Standby RP
- Does not pre-load configuration or state
- Resets & reloads line cards causing link and route flaps

Enhanced Fast Software Upgrade – eFSU

- Relies on Dual Processors
- Built on SSO
- Preloads the new software image on Standby RP
- Does pre-load select line cards
- Reloads line cards using warm upgrade causing route flaps

ASR 1000 & IOS-XE High Availability Highlights

- ASR 1000 leverages Cisco IOS HA infrastructure – NSF/SSO, ISSU
- 1+1 redundancy option for RP and ESP
  - Active and standby
  - No load balancing
- RP's are separate from ESP's
  - Switchover of ESP does not result in switchover of RP
  - Switchover of RP/IOS does not result in switchover of ESP
- Single RP may be configured with dual IOS for SW redundancy (single RP only)
- ISSU of individual SW packages

System Architecture – Distributed Control Plane

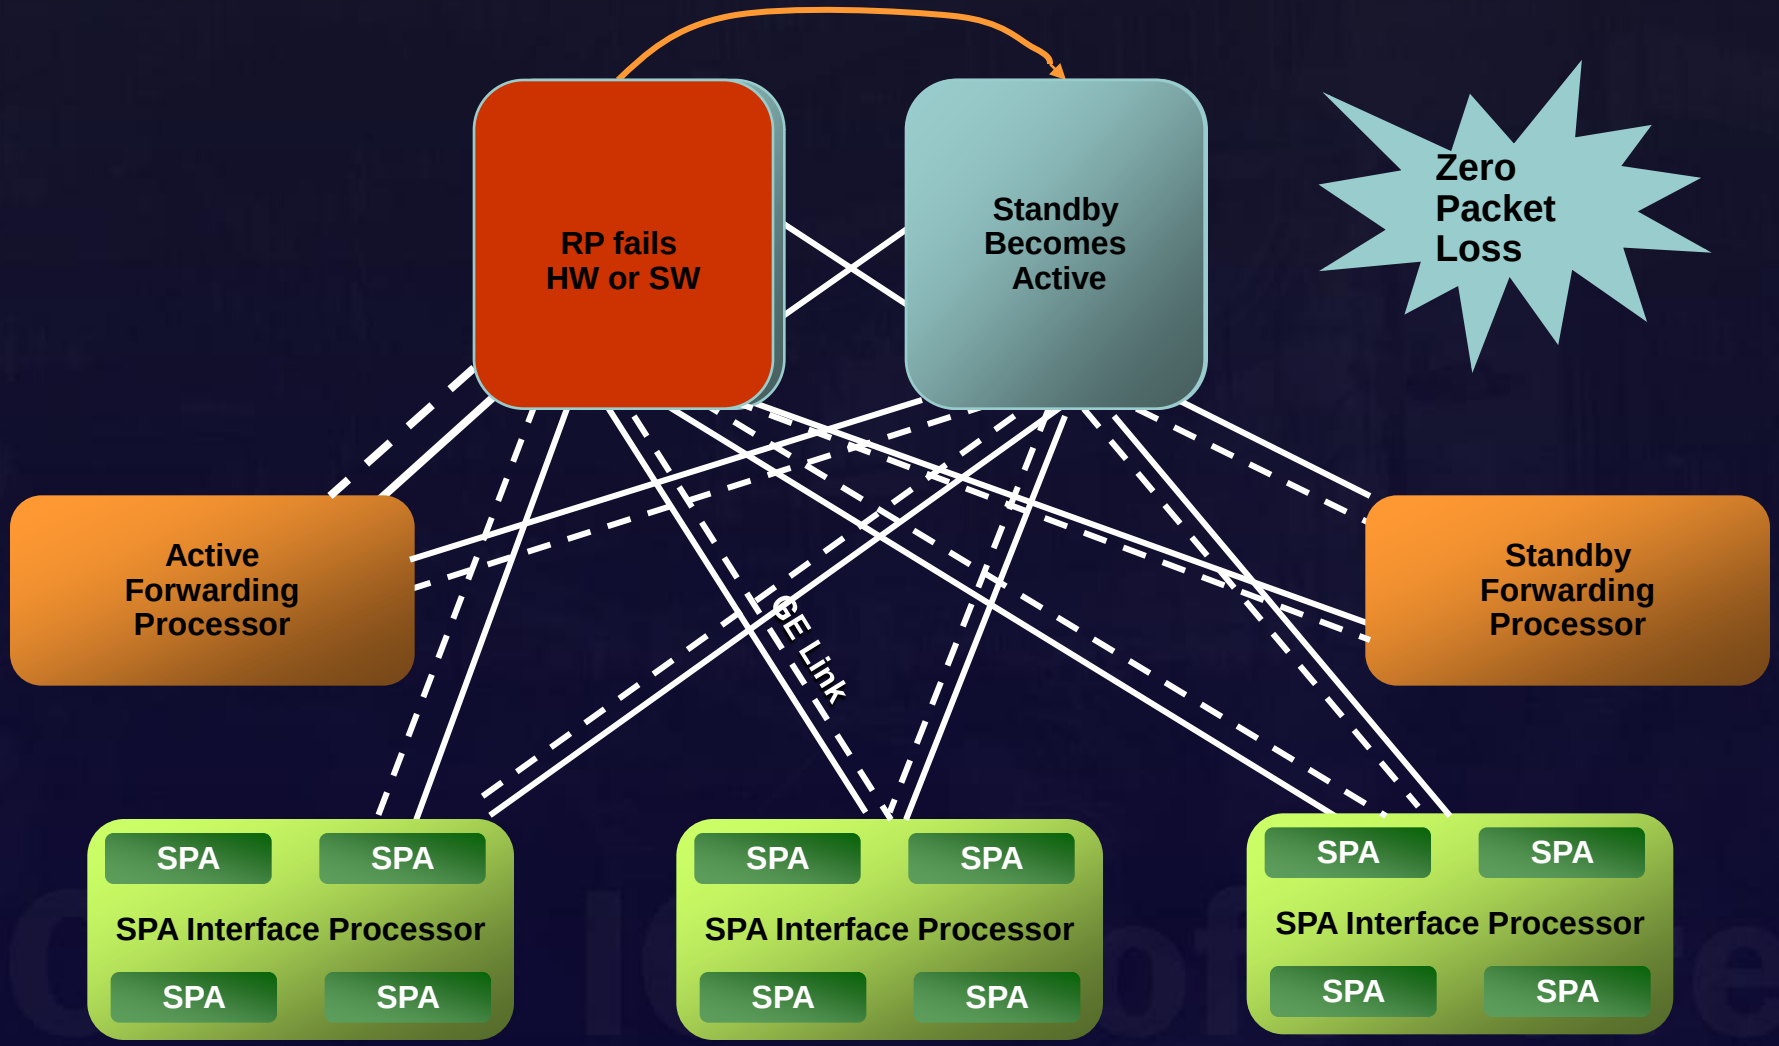

Separate and independent internal communication link for control plane (GE)

System Architecture – Centralized Data Plane

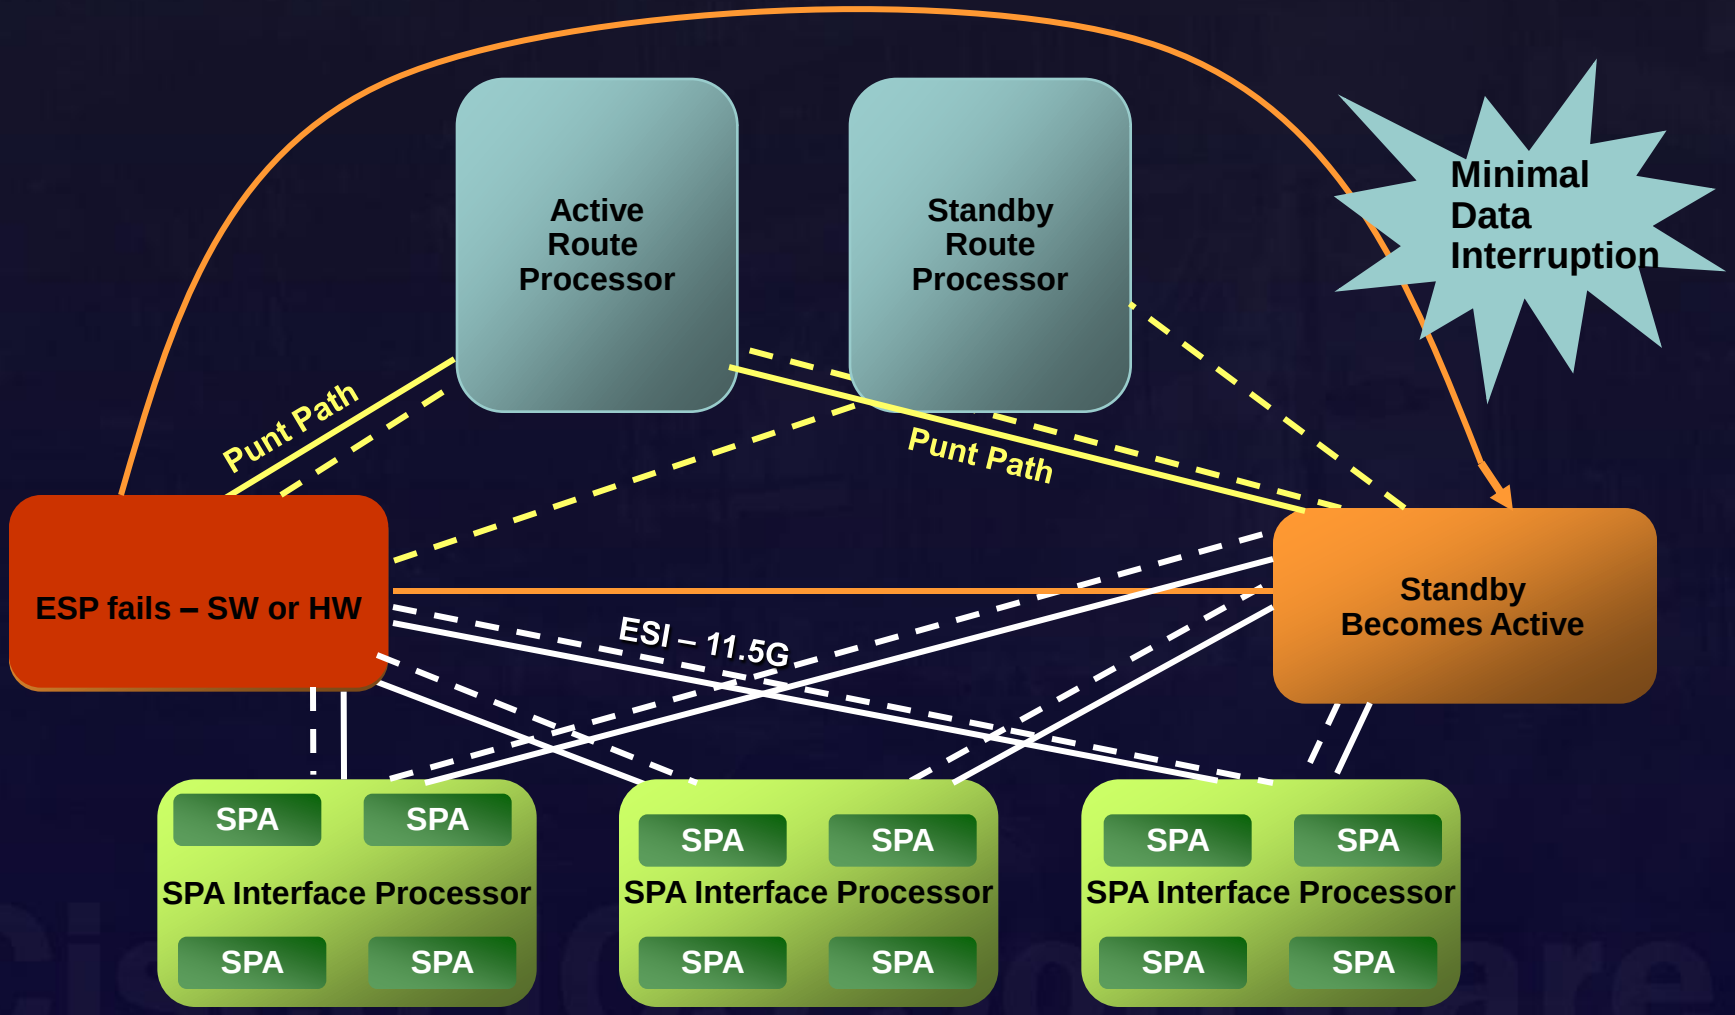

- All packets processed by QFP for forwarding
- Separate and Independent links for Data Plane communication (ESI 11.5G)

Software Architecture – IOS-XE

- IOS-XE = IOS + IOS-XE Middleware + Platform Software
- Operational Consistency - same look and feel as IOS Router
- IOS runs as its own Linux process for control plane (Routing, SNMP, CLI etc). 32bit and 64bit options.
- Linux kernel with multiple processes running in protected memory for
  - Fault containment
  - Re-startability
  - ISSU of individual SW packages
- ASR 1000 HA Innovations
  - Zero-packet-loss RP Failover
  - <50ms ESP Failover
  - "Software Redundancy"

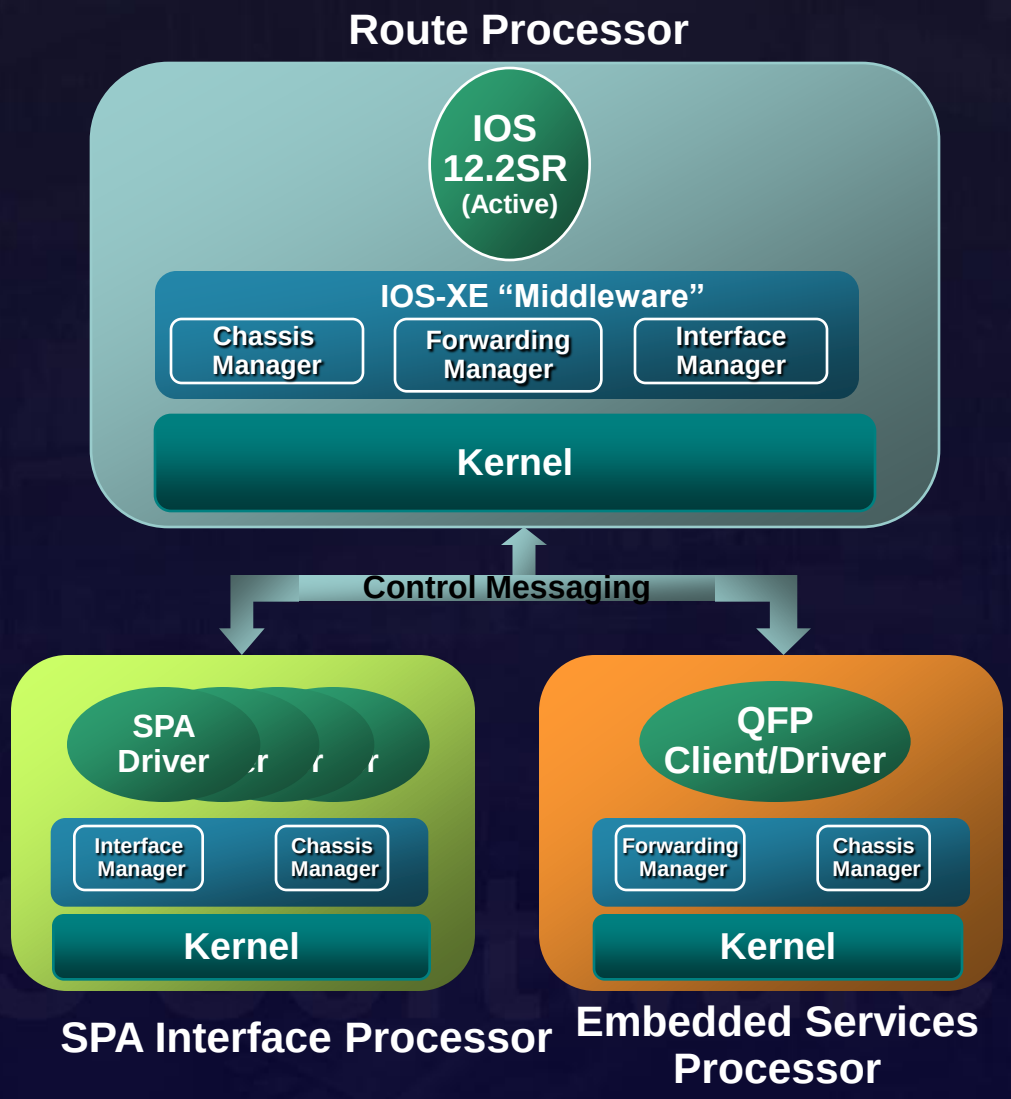

IOS Software Redundancy on 4RU/2RU Single RP/ESP

- Stand-by IOS process in RP in the single-engine 4RU/2RU system
- Two IOS process in a single RP function similar to different processes on separate RP
- Supports all NSF/SSO features supported by dual-RP systems
- Requires additional RP memory – 4G

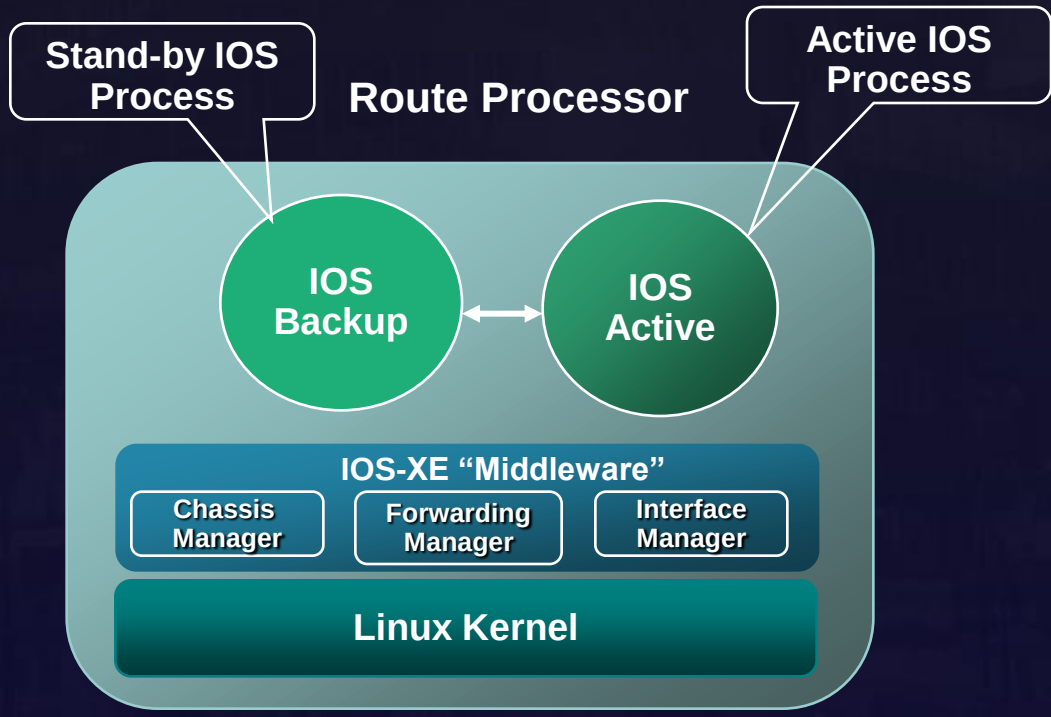

SIP SW Module Upgrade

- SIP redundancy is not supported on ASR 1000
- SIPBase cannot be upgraded ‘in service’
- However, ASR 1000 supports the ability to upgrade selective SPA slots with new SIPSPA software without affecting other SPA slots in the SIP
- This is possible since ASR 1000 supports up to 4 different versions of SIPSPA modules for a single SIP

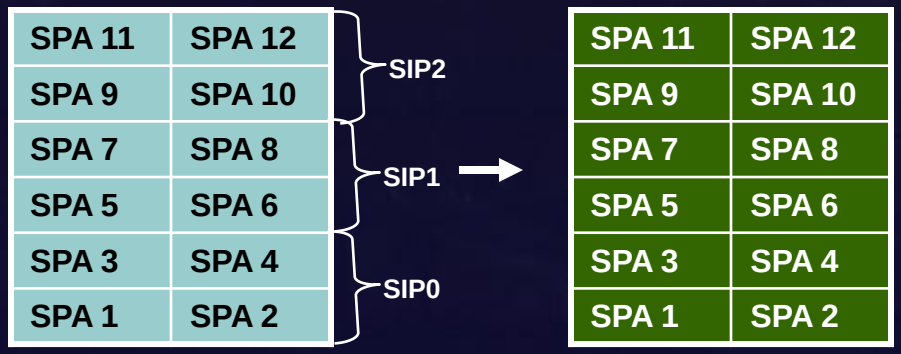

Case 1:

- All SPAs are using default SIPSPA ‘blue’
- New SIPSPA ‘green’ installed and applied to all SPA slots
- Result: All SPAs are reset

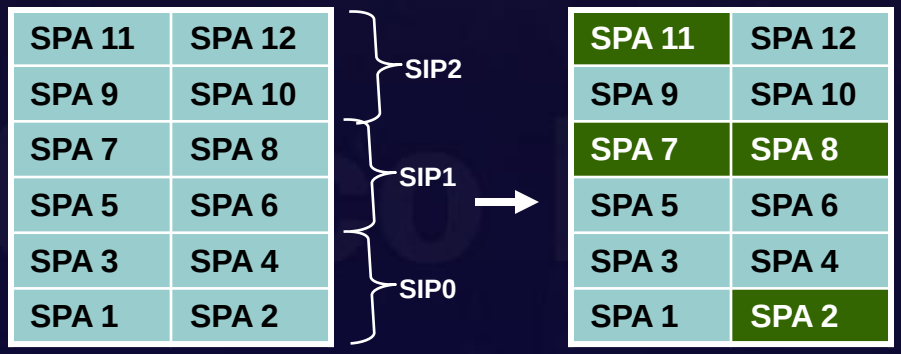

Case 2:

- All SPAs are using default SIPSPA ‘blue’
- New SIPSPA ‘green’ installed and applied to selective SPA slots only
- Result: All SPAs that have the ‘green’ SIPSPA applied to, are reset

BGP Non-Stop Routing

Unique, Self-Contained Edge Routing HA Solution

- Simplifies NSF/SSO deployment by synchronizing edge routes automatically
  - NSF-aware Customer Edge devices not needed
  - Addresses additional network scenarios—e.g. unmanaged CPEs
- Delivers persistent routing for the entire customer edge
- Retains scalability and safety of NSF/GR with benefits of NSR

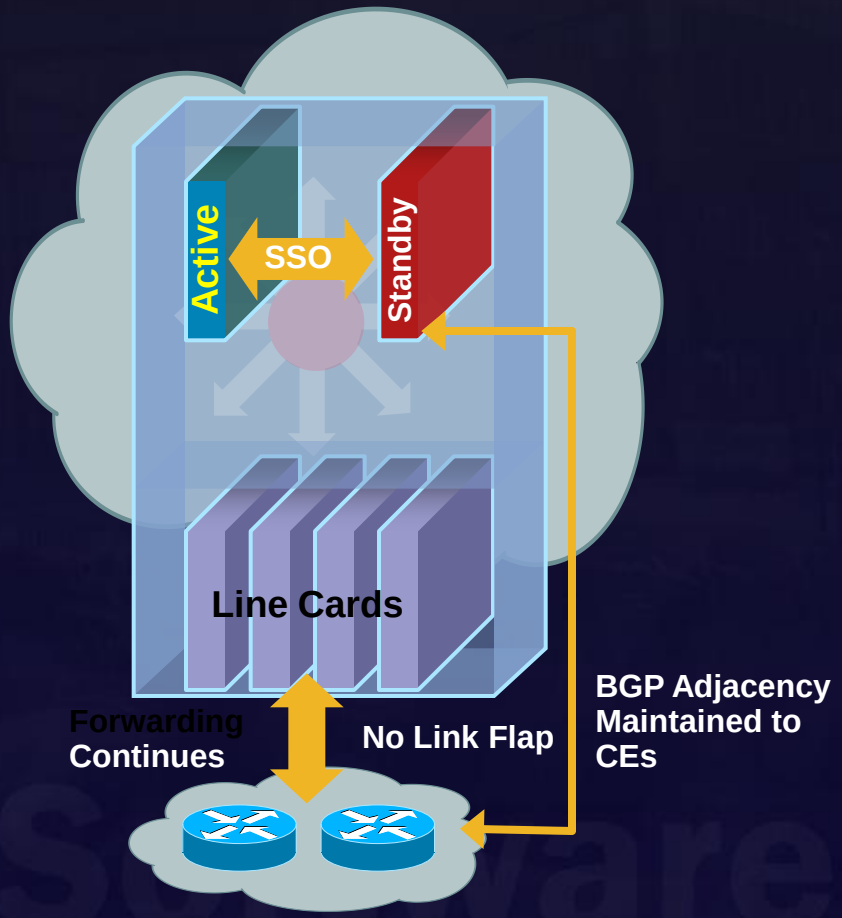

MPLS Nonstop Forwarding and Stateful Switchover

Proven Cisco NSF/SSO Technology for MPLS LDP and VPNs

- NSF/SSO capabilities extended to MPLS forwarding, MPLS LDP and MPLS VPNs (including InterAS and CsC)
- Eliminates service disruptions at IP/MPLS edge
  - Preserves sessions and mitigates outage impact
- Reduces costs
  - Network downtime, SLA penalties, operational costs
- Increases operational efficiency
  - Reduced network administration, troubleshooting, and maintenance

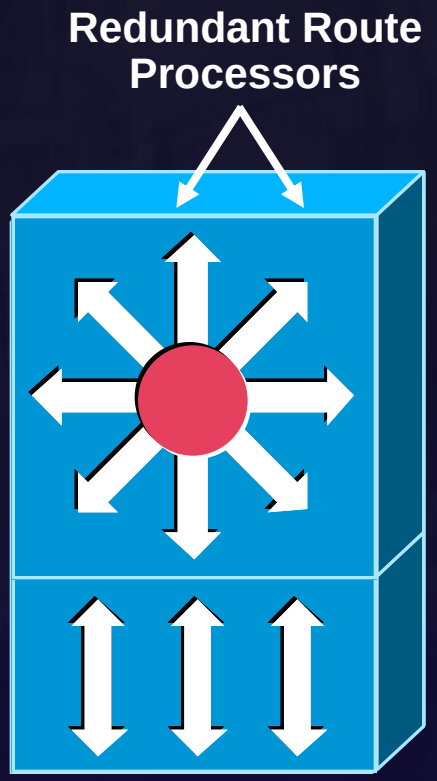

Stateful Switchover (SSO)—Zero Interruption in Layer 2 Connectivity

Nonstop Forwarding (NSF)—Continuous Packet Forwarding with Minimal Packet Loss

Line Card Redundancy

- Provides ability to mitigate unplanned outages due to Line Card hardware or software failures
  - 1+1 APS with No Layer-3 Re-Convergence
  - 1:1 Hot Line Card Redundancy (Y-Cable)

Benefits

- Reduces customer outage due to LC failures from a few hours to a few minutes or seconds
- Enables incremental revenue opportunities
  - Ability to offer better SLAs
- Reduces operational costs
  - Helps avoid SLA penalties, emergency dispatches, repair operations, and additional staffing

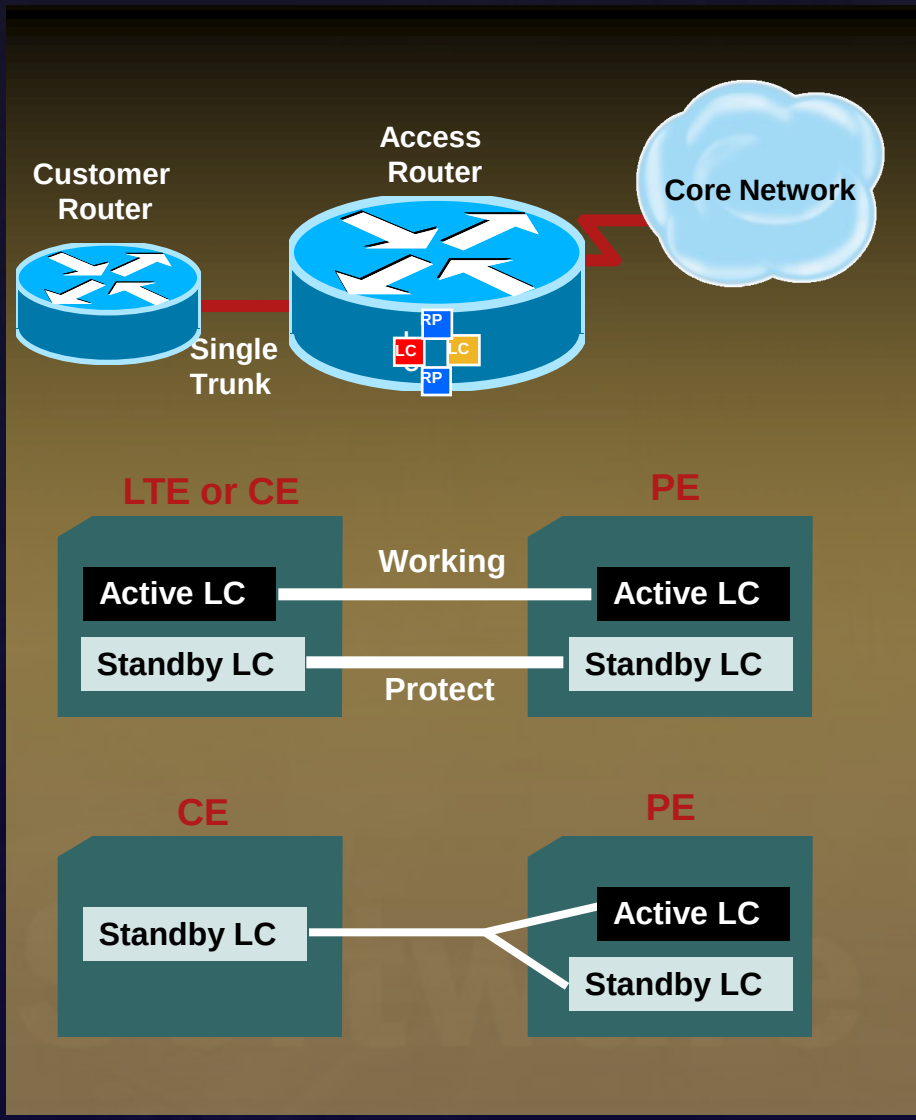

Line Card Redundancy— APS with No Layer 3 Reconvergence

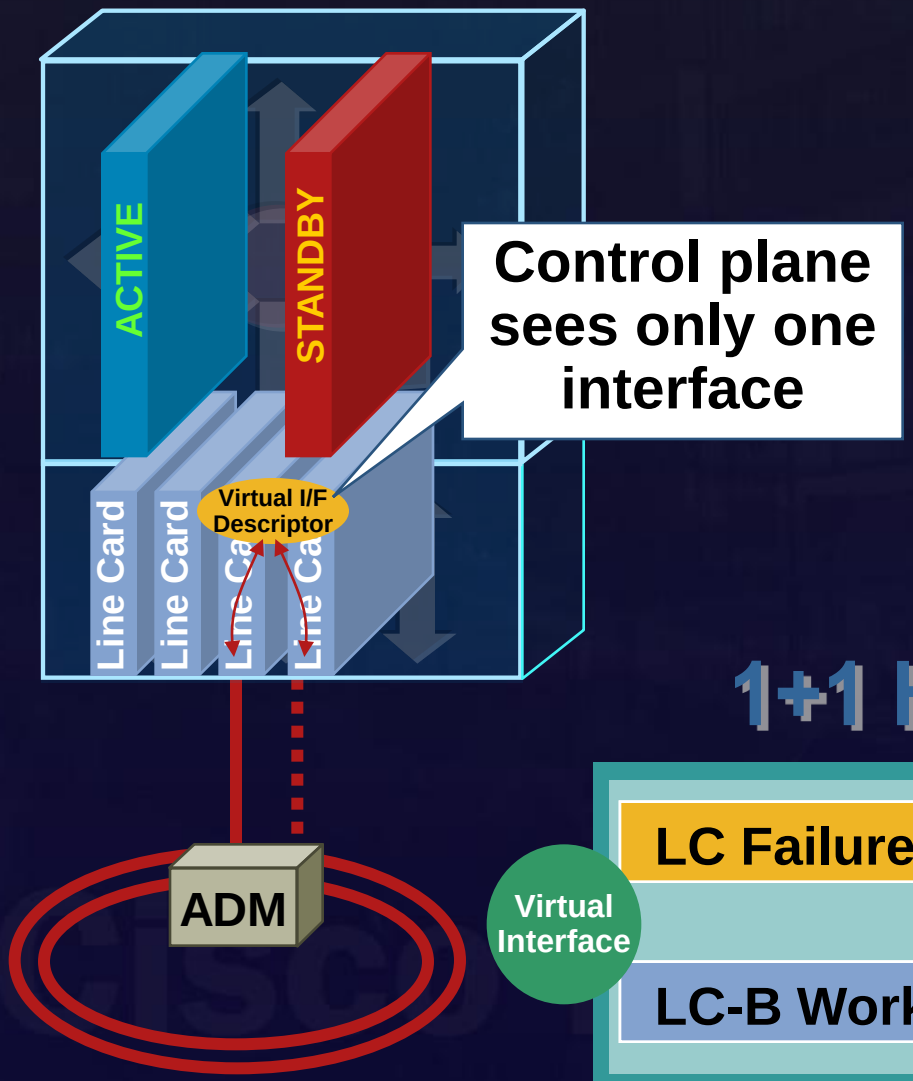

- APS interfaces coupled but separated
- Switch from working to protect results in link “down”, then “up” from routing protocol perspective
- New feature establishes virtual APS interface

1+1 Hot Redundancy

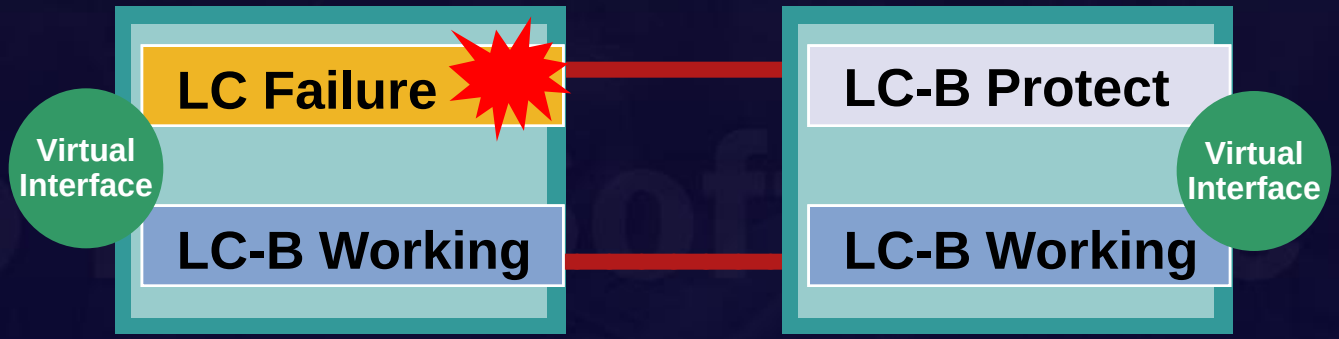

Line Card Redundancy— 1:1 Hot Redundancy with Y-cable

- 1:1 Line Card Redundancy with "Y-cable"
- Optics "ON / OFF" control
- Electrical relays

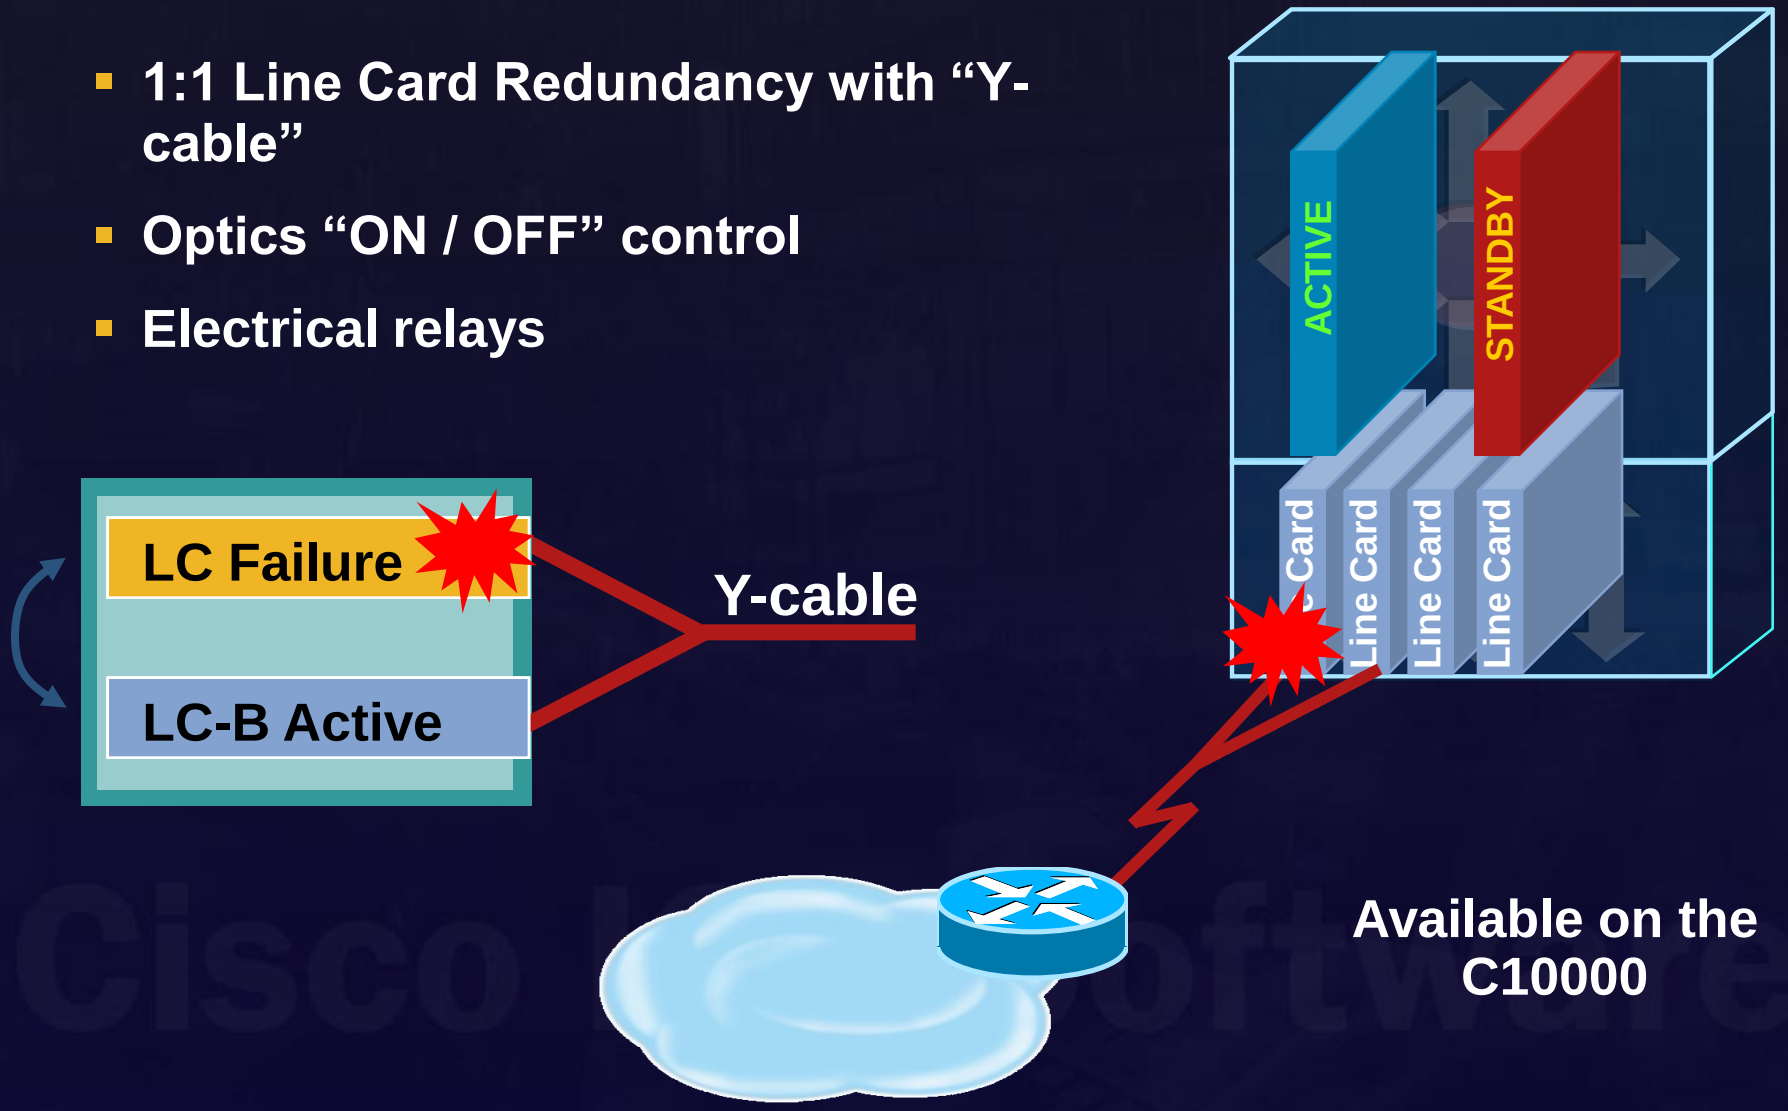

Available on the C10000

Warm Reload and Warm Upgrade

- Warm Reload enables significant reduction in device reboot time, lowering the mean time to repair (MTTR) for software failures
  - Accelerates execution using previously saved, pre-initialized variables
- Warm Upgrades extends the Warm Reload capabilities for planned upgrades or downgrades
  - Router reads and decompresses the new Cisco IOS Software image then transfers the control to it, while packet forwarding is continued

Warm Reload Process

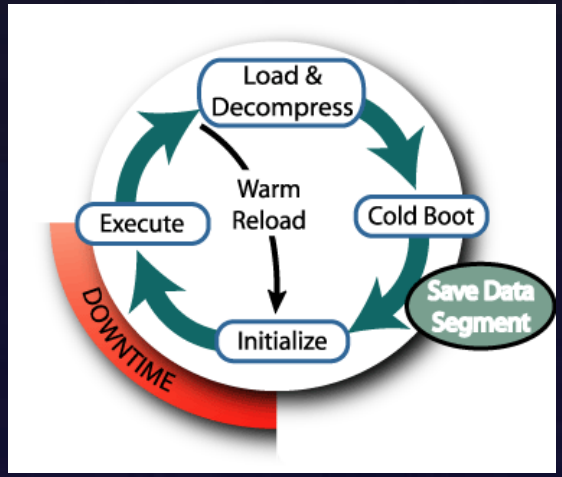

Warm Upgrade process

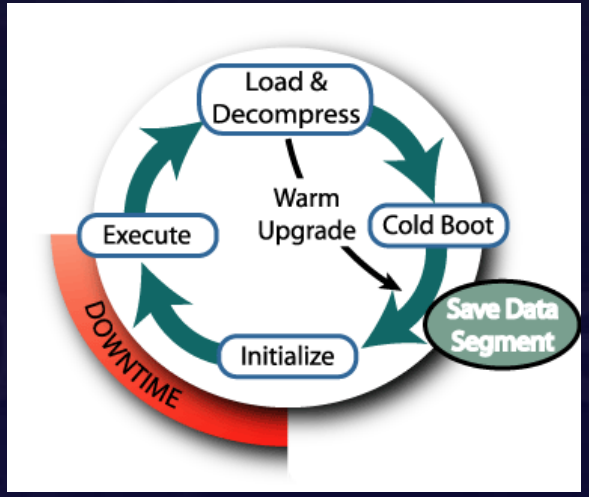

Gateway Load Balancing Protocol

First hop redundancy prior to GLBP

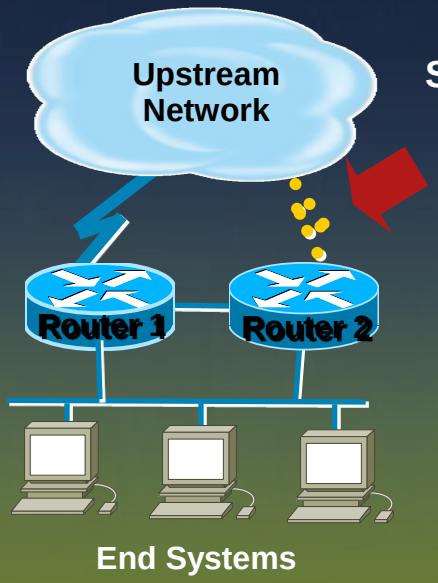

Second uplink not used

- Unused standby capacity
- Seconds of packet loss

First hop redundancy with GLBP

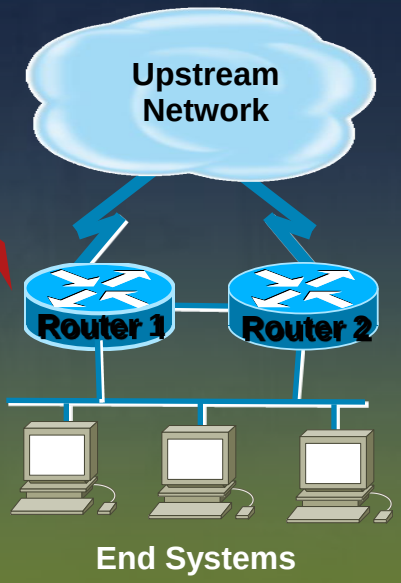

Optimized load sharing across all uplink interfaces

- Cost efficient
- Limited packet loss
- Easier to deploy

*Cisco innovation: Patent pending

Benefits:

- Efficient resource utilization
- Easier deployment and maintainability

Stateful IP Services

- Stateful IPsec maintains connectivity to users and applications

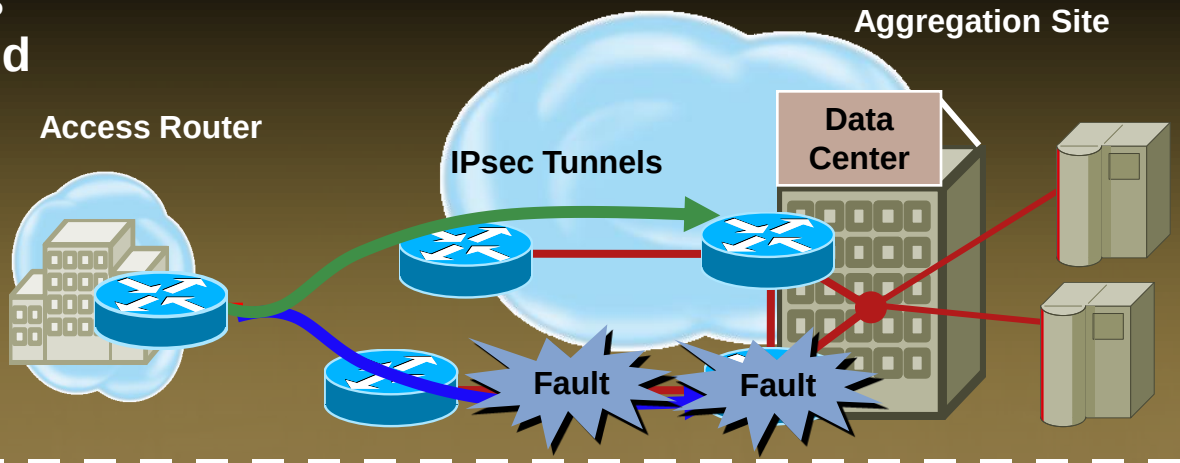

- Stateful NAT maintains application session
- Stateful Cisco IOS Firewall

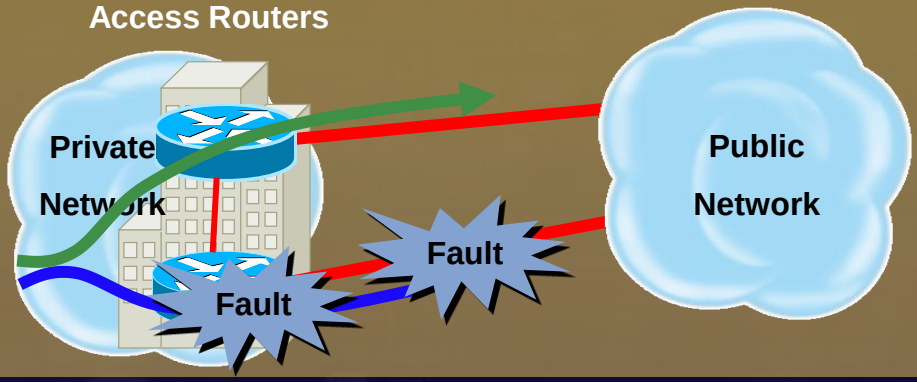

Benefits:

- No Service Disruption
- Reduced Costs

Control Plane Policing

- Limits incoming traffic to the control plane via Quality of Service (QoS)
- Provides better hardware reliability and availability
- Protects against Denial of Service (DoS) attacks
- More control for packets destined to the control plane
- Simplifies configuration for control plane policies

Network Level Resiliency Overview

Deliver industry-leading advances in network convergence times, protection, and restoration

Convergence and self-healing

- Reduce convergence times for major network protocols – IS-IS, OSPF, BGP, EIGRP
- Leverage in network core or where redundant paths exist

Intelligent protocol fabric

- Embed NSF intelligence network-wide in Service Provider Core, Edge Networks, and Enterprise Networks

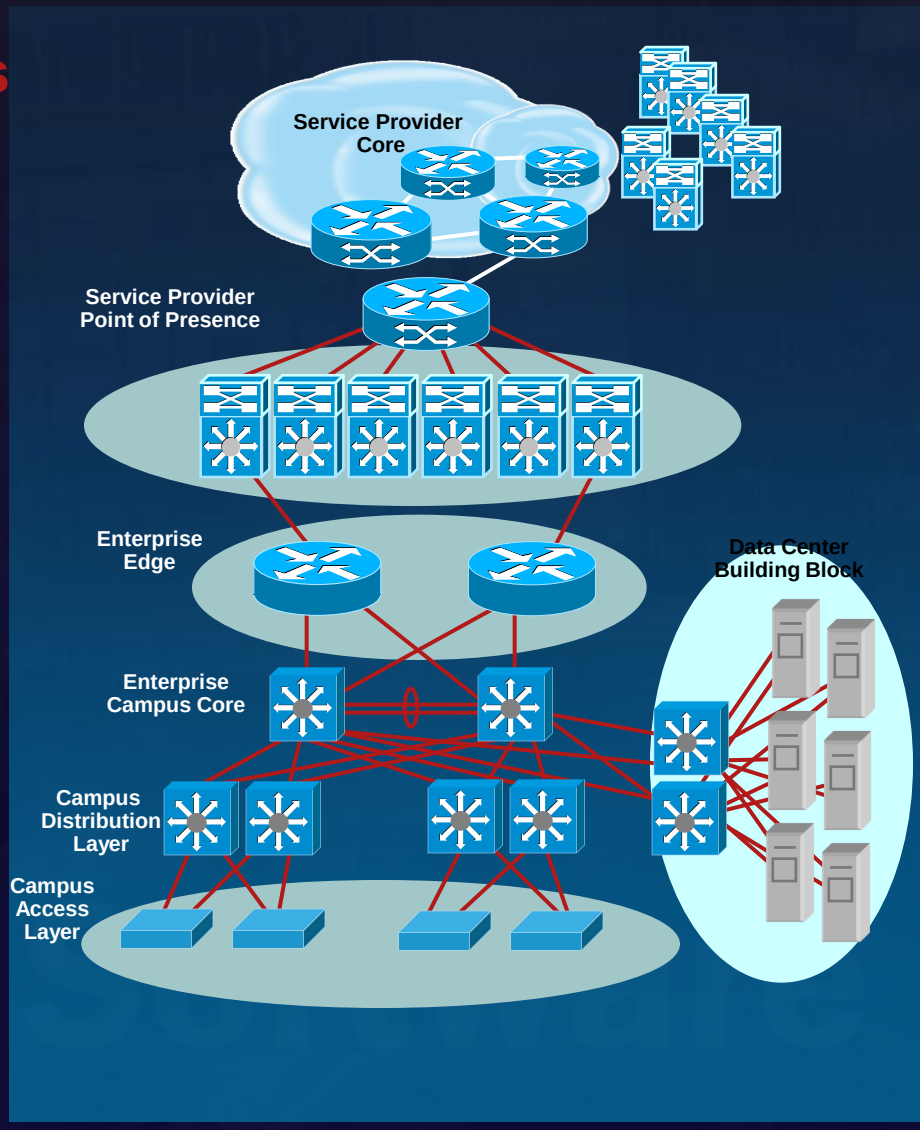

Cisco NSF Awareness

- An NSF-capable router continuously forwards packets during a switchover
- An NSF-aware router assists NSF-capable routers by:
  - Not resetting adjacency
  - Supplying routing information for verification after switchover
- NSF-capable and NSF-aware peers cooperate using Graceful Restart extensions to BGP, OSPF, IS-IS, EIGRP and LDP protocols
- NSF capability at key edge nodes that are single points of failure
- NSF awareness for deployments with redundant topologies having alternate paths
- Cisco advantage:
  - End-to-End NSF Awareness across product portfolio (Cisco 1700 Series Router to CRS-1) maximizes the benefits of NSF

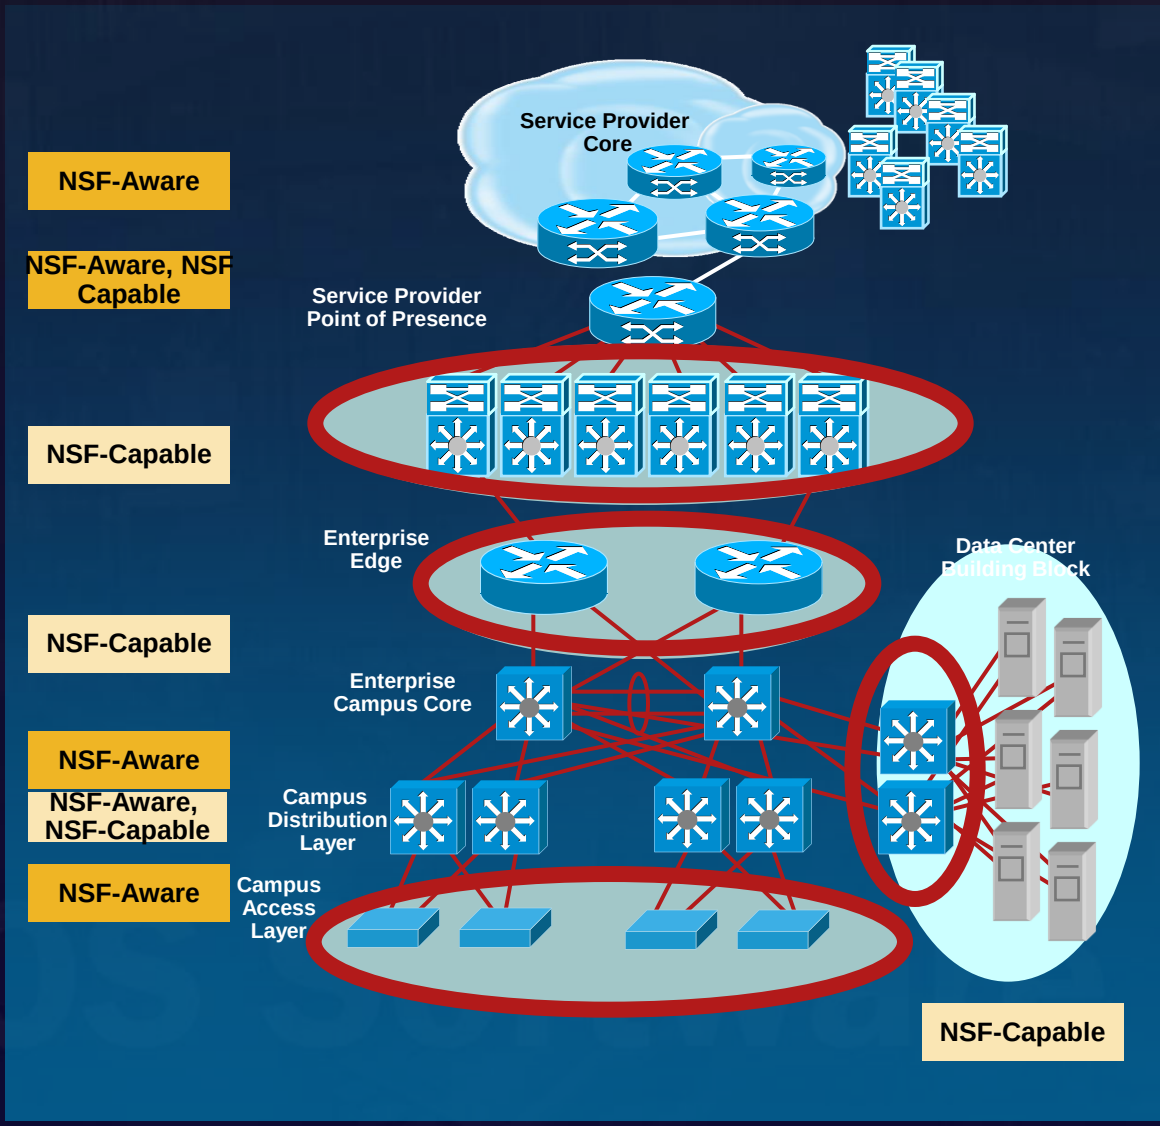

Routing Convergence Enhancements

- BGP Convergence Optimization
  - Improves BGP convergence, router boot time, and transient memory usage
  - 40%–50% convergence time improvement
- Incremental SPF Optimization (OSPF, IS-IS)
  - Allows the system to recalculate only the affected part of the shortest path tree (rather than the entire tree)
  - 10%-90% improvement in convergence time based on distance from network event
- IP Event Dampening
  - Self policing of link flaps reduces packet loss & increases network stability
- OSPF Support for Fast Hello Packets
  - Provides ability to send hello packets in intervals less than one second, enabling faster OSPF convergence
- OSPF Support for Link-State Advertisement (LSA) Throttling
  - Provides a dynamic mechanism for slowing down LSA updates in OSPF during times of network instability
  - Allows faster OSPF convergence by providing LSA rate limiting in milliseconds

Benefits:

- Faster convergence leading to higher availability
- Increased network stability

Bidirectional Forwarding Detection (BFD)

Lightweight Hello Protocol providing sub-hundred milliseconds forwarding plane failure detection using single, common, standardized mechanism, independent of media and routing protocols.

- Sub second link failure detection
- Low overhead
- Media independent
- Protocol independent
- BFD on edge allows applications to do fast failure detection and fast failure recover if there's an alternate path

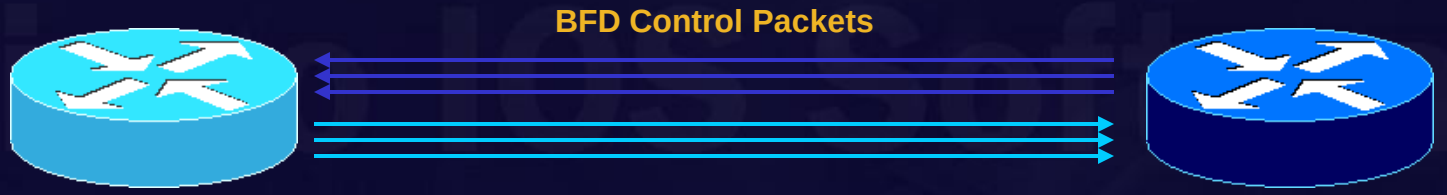

BGP Convergence Optimization

A series of enhancements:

- **BGP Update Packing**
  - Sends BGP NLRI information in a sorted order
  - Greatly speeds best-path calculation
- **Memory backoff algorithm**
  - Reduces memory churn during the transient period of initial convergence
- **TCP MTU Path Discovery**
  - Increases amount of updates that can be sent in one packet
  - Reduces I/O

Further optimization delivers 40% improvement in BGP convergence time for the world's largest BGP production environments – BGP tables with over 200,000 prefixes

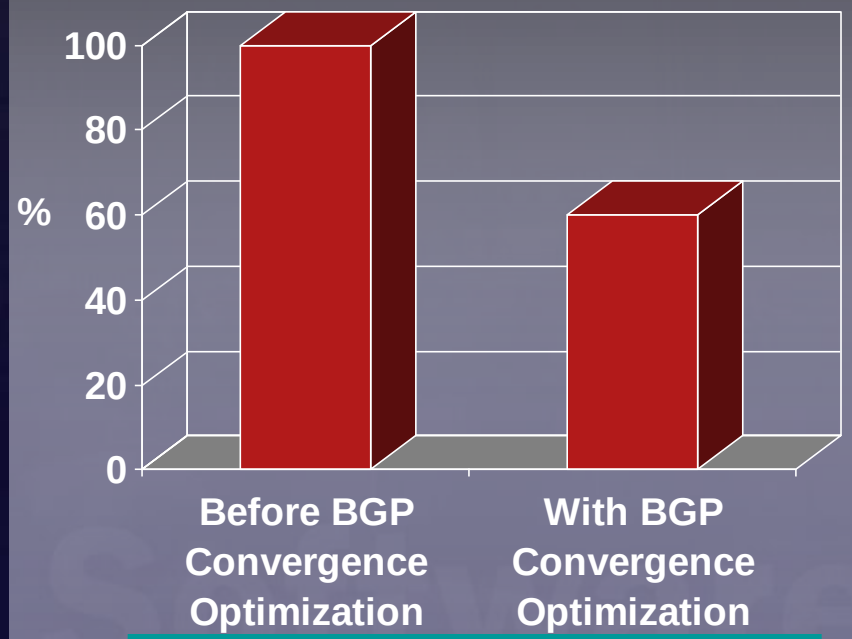

- **Higher availability**
- **Greater stability**

Incremental SPF Optimization

- When there is a state change anywhere within an OSPF area/IS-IS routing domain, Dijkstra's algorithm is run by all routers in the area/domain
- Wasteful use of CPU and memory
  - The change has less impact as it moves further away
- Incremental SPF allows OSPF or IS-IS to shortcut Dijkstra by calculating only the changed portion of the tree
- Saves CPU/Memory

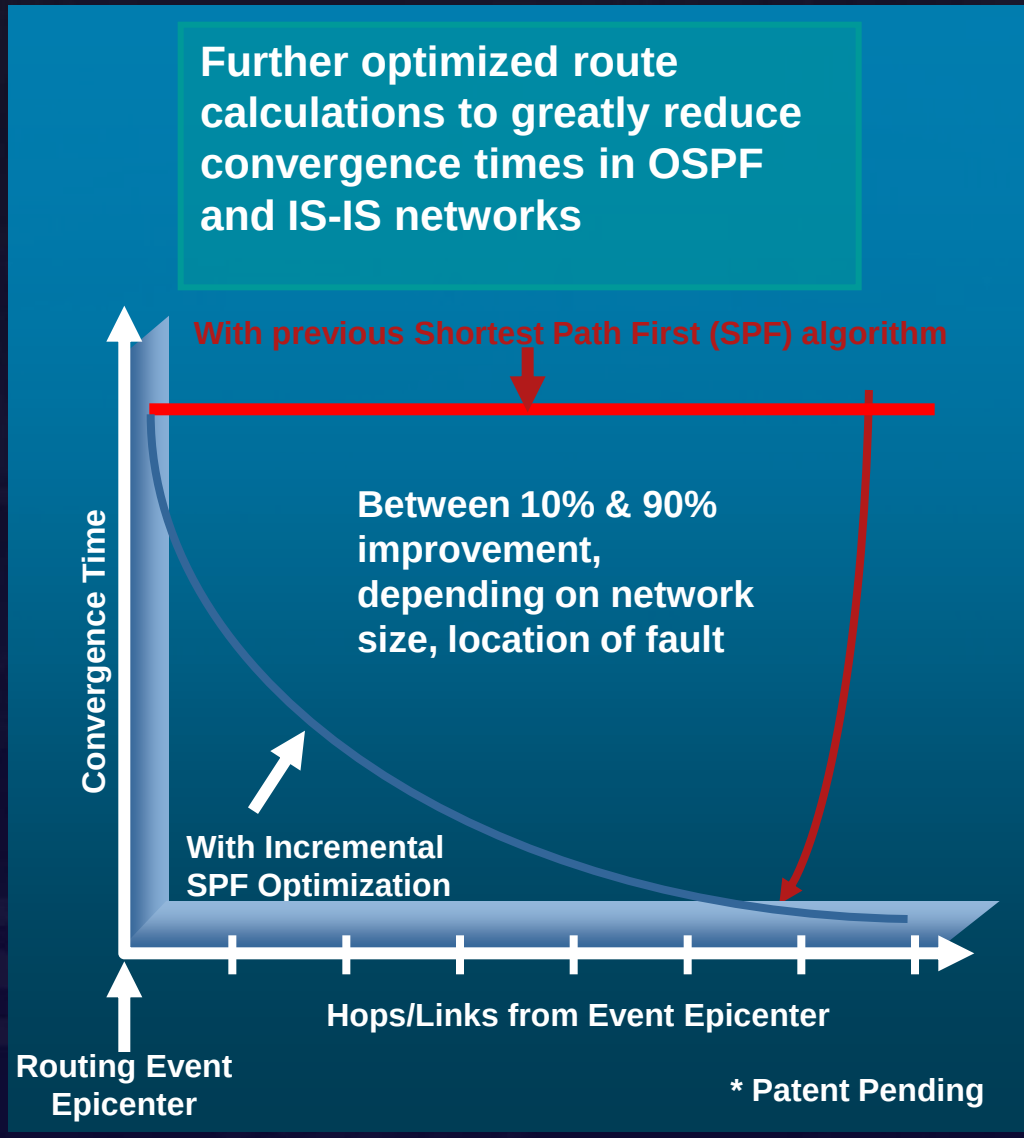

IP Event Dampening

IP Event Dampening logically isolates unstable links by reducing packet loss & routing CPU overhead

Packet loss duration without IP Event Dampening

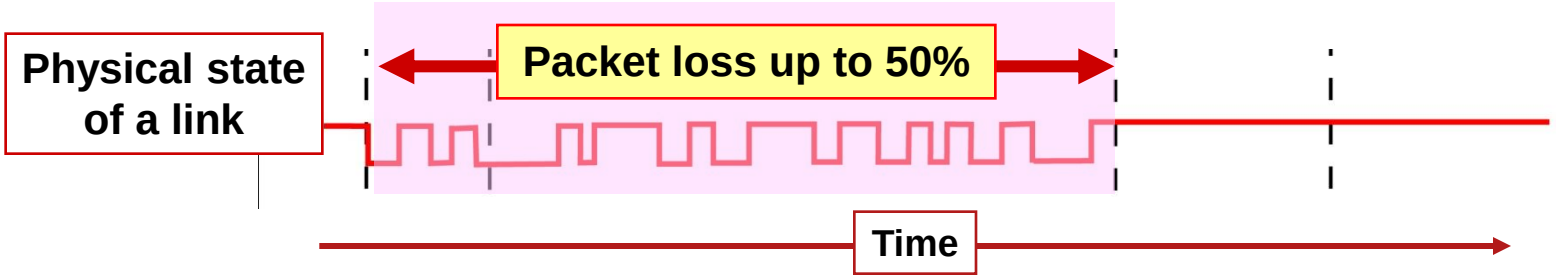

Packet loss window becomes significantly reduced with IP Event Dampening

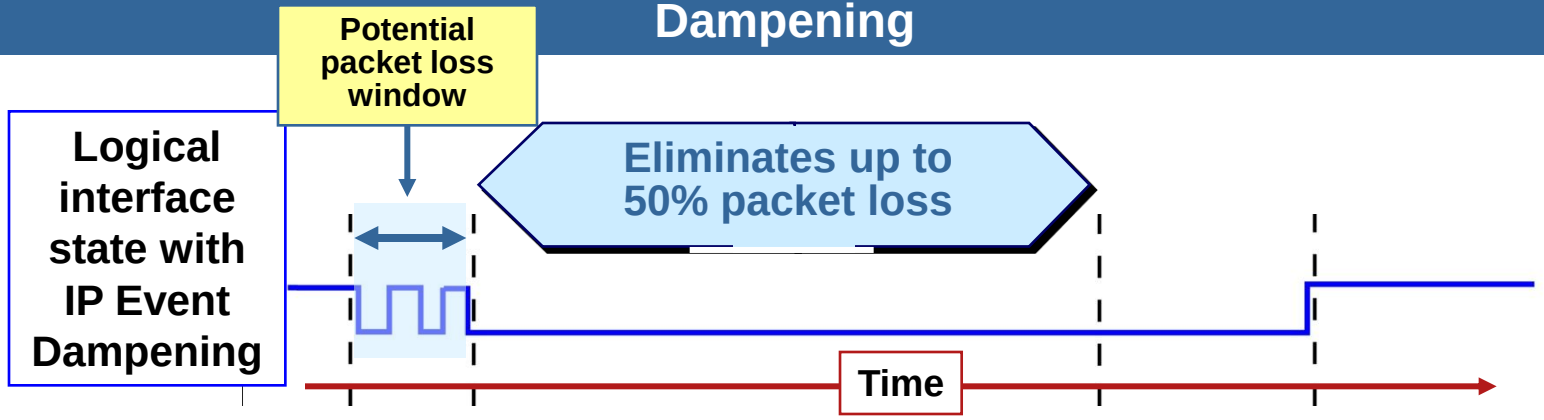

Multicast Subsecond Convergence

Series of a real enhancements:

- **Join/Prune aggregation**
  - Cisco used to send one PIM packet per (S,G) or (*,G) entry after a Rendezvous Point failover
  - Cisco now aggregates these into only a few PIM packets with multiple entries
- **New PIM HELLO option**
  - Added option
  - Advertises the HOLDTIME in milliseconds
  - Allows sub second failover of Designated Router
- **Triggered RPF checks following unicast convergence**
  - Now, as soon as unicast is converged, it causes an instantaneous start of RPF checks (old default was 5 seconds)
- **Important for Multicast applications**
  - Financial trading information
  - Manufacturing process control

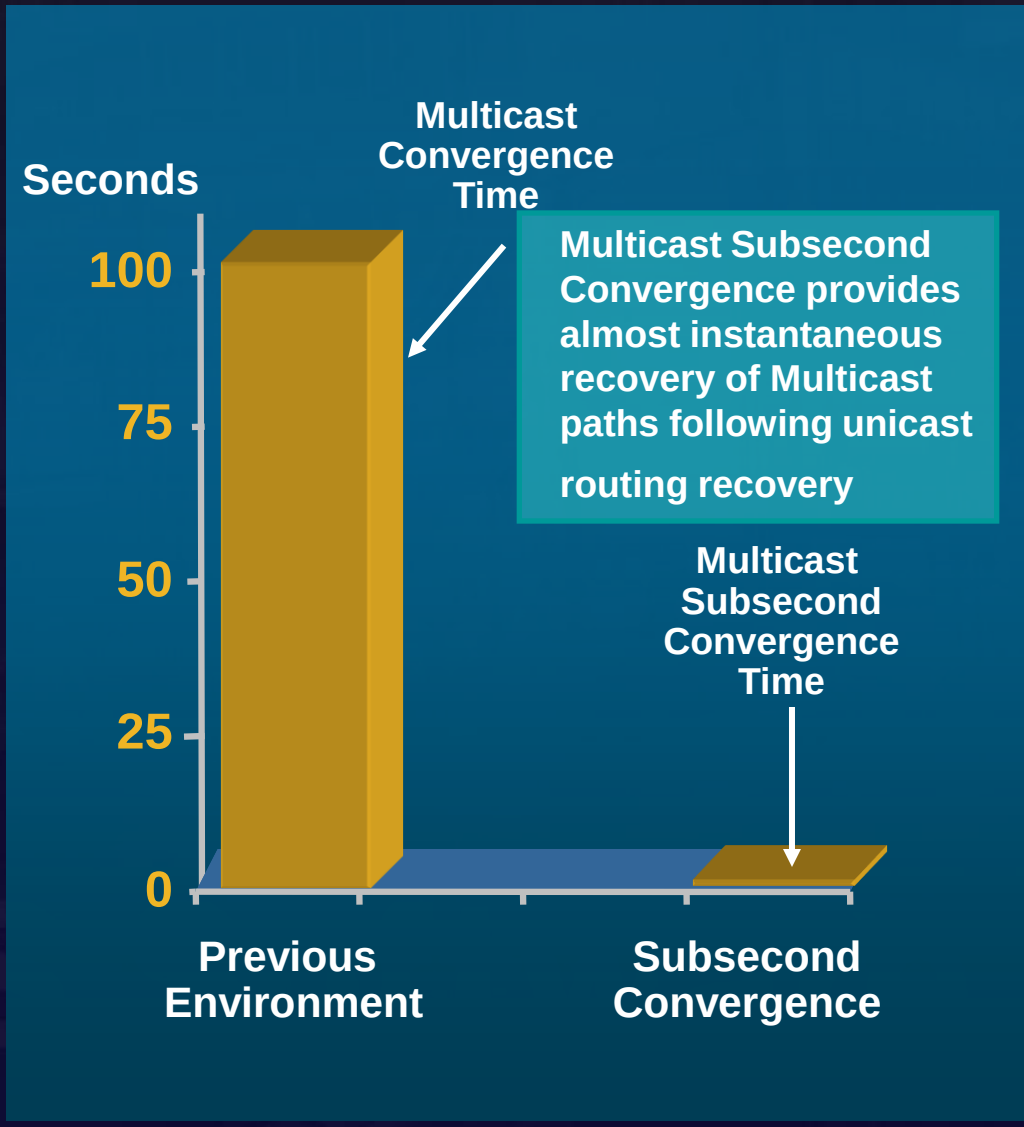

Embedded Management Overview

Deliver embedded intelligence for proactive fault/event management, configuration management and availability measurement

- Embedded fault/event management for proactive maintenance and local recovery
- Automation and configuration management to reduce human errors
- Accurate, consistent, embedded availability measurement

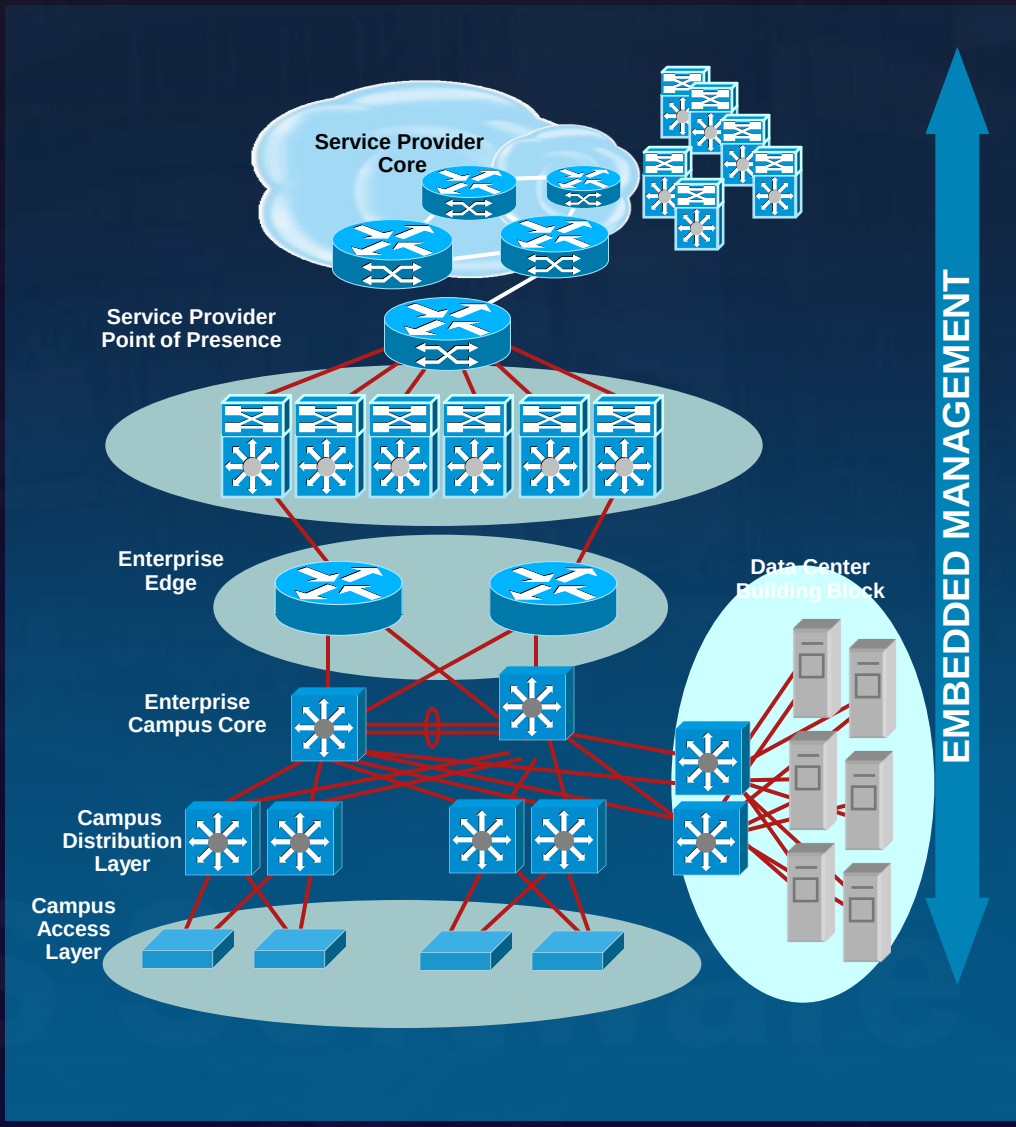

Cisco IOS Embedded Management Infrastructure
Embedded Event Manager -&- Embedded Resource Manager

for Increased Availability, Lower MTTR

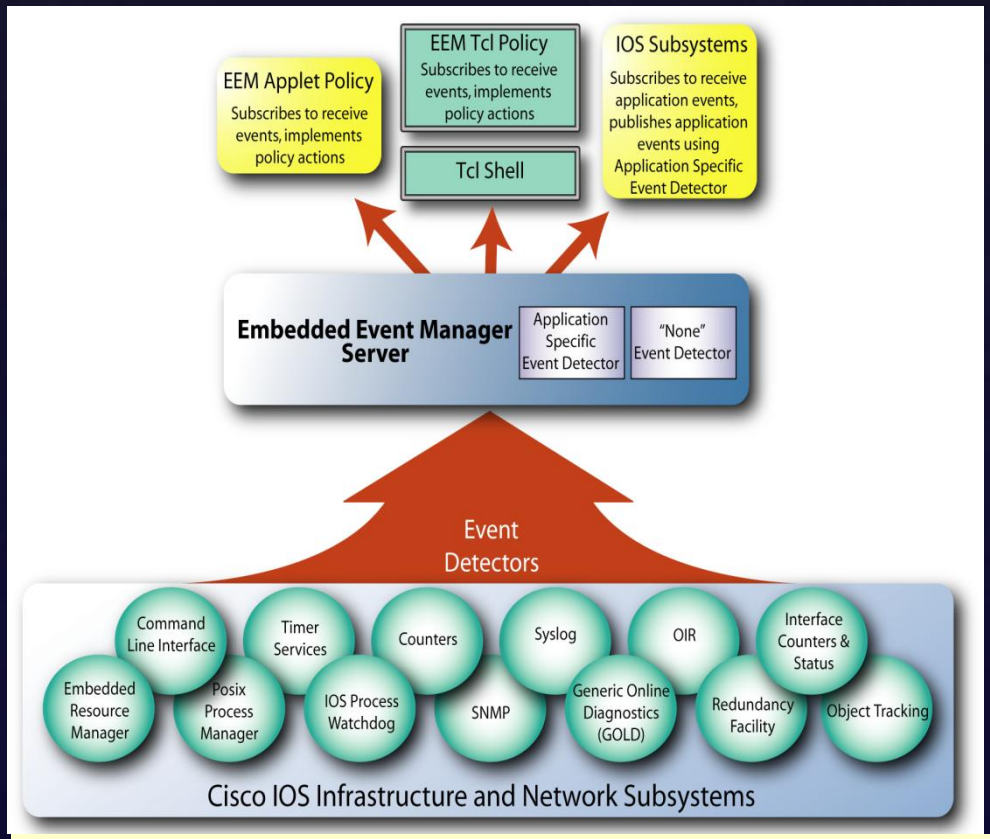

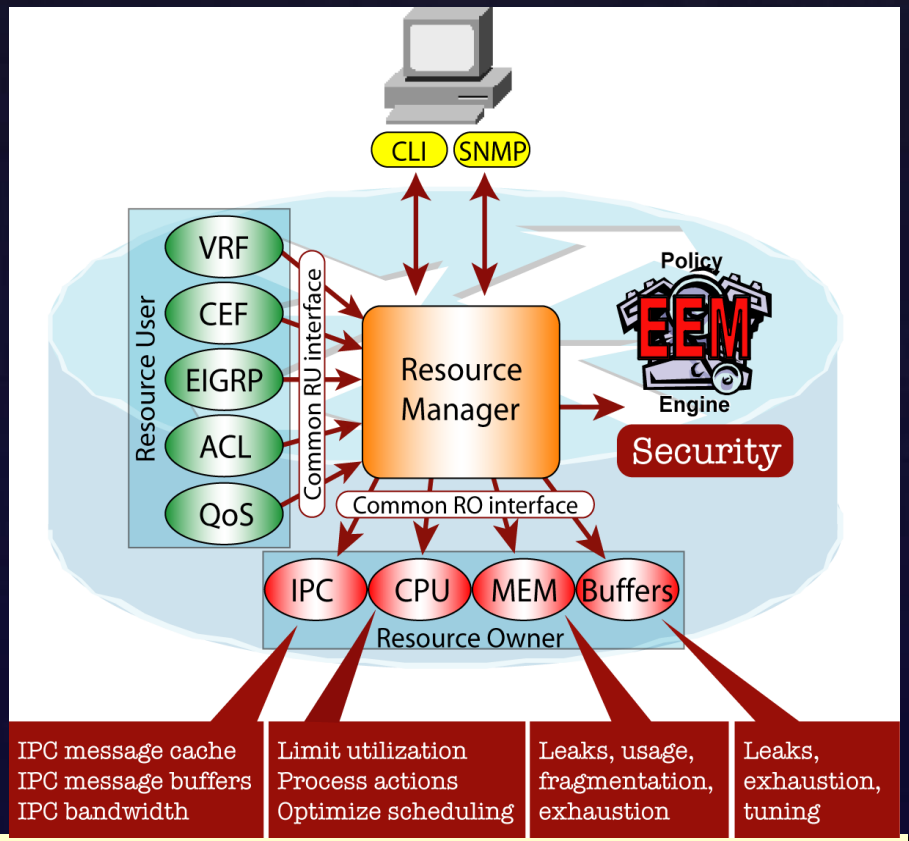

- **New infrastructure capabilities make Cisco routers and switches active participants in network management**
- **Self diagnosis, self awareness, autonomy and automation**

Embedded Event Manager

- Predictive, consistent, scalable fault/event management embedded within Cisco IOS Software devices
- Program local actions based on system state
  - Define policies using TCL scripts
  - Events generated by system state, policies executed
- Flexible and customizable serviceability tool
  - Identify and correct anomalies within user networks in real time

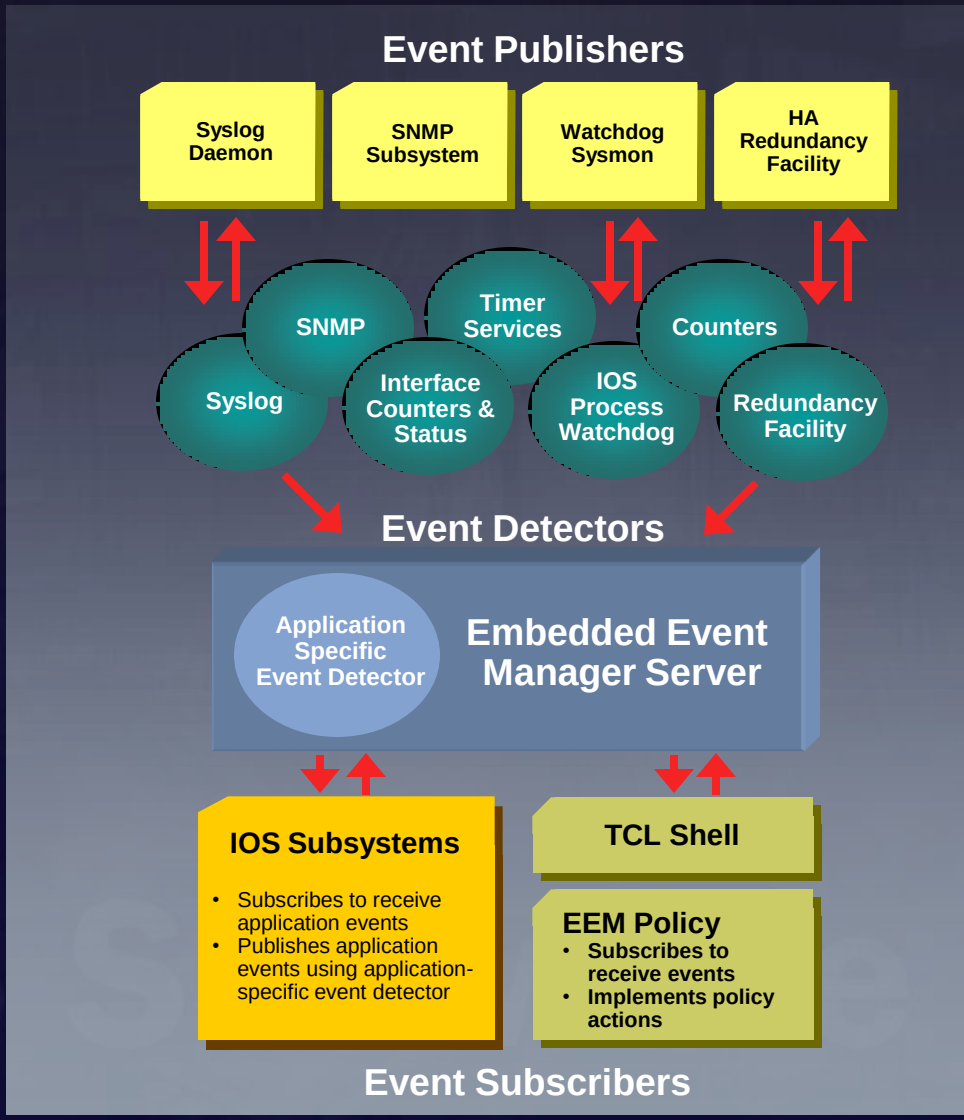

Embedded Resource Manager (ERM)

- Cisco IOS, like all operating systems, must internally manage finite resources
- ERM provides a consolidated, consistent facility to monitor, manage, and react to dynamic changes in resource capacity and availability
- Actionable intelligence - actions to improve the performance and availability of the router
- Arms customers with information…
  - To better understand scalability and resource usage
  - To anticipate future needs and perform capacity planning

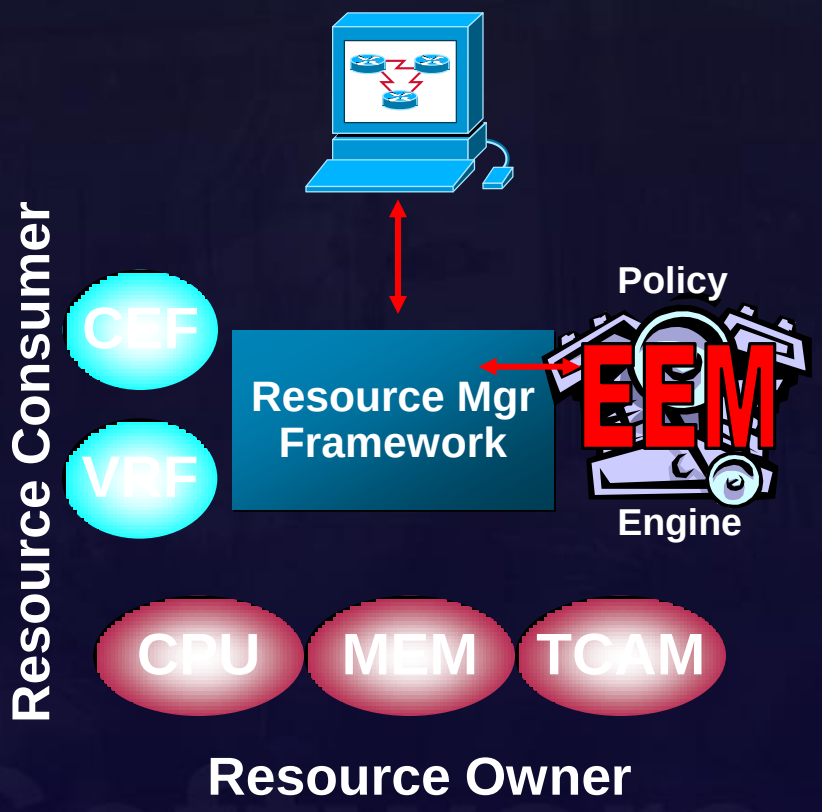

Configuration Rollback

Fault Containment: Reducing Human Error

- Human error represents significant downtime
  - Fat-finger can sometimes bite even seasoned and knowledgeable professionals
  - Mistakes can occur in the heat of troubleshooting
  - Deal with reality while striving to improve
- Cisco IOS Software configuration subsystem and command parser has evolved:
  - New capability offers situational benefit
  - Eases impact of configuration mistakes
  - Provides recovery and prevents escalation and exacerbation
  - Provides interface for extended automation

Configuration Rollback Components

- Contextual Configuration Difference
- Configuration Archiving
- Configuration Replace
- Configuration Logging

Generic Online Diagnostics (GOLD)

Proactive fault detection

- Enables network to proactively detect hardware and software malfunctions before they adversely impact network traffic
- Runtime execution of diagnostic routines
  - Bootup diagnostics (during initial boot & OIR)
  - Health monitoring diagnostics while system is in operation
  - Health monitoring tests are non-disruptive
- On-demand or Scheduled diagnostics
- Integrated with “Call Home” applications

Generic Online Diagnostics

Configuration/Reporting

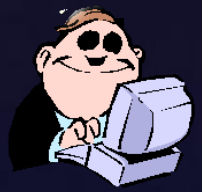

Configure Online Diagnostics and check diagnostics results

Detects bad Hardware

GOLD

Boot-Up Diagnostics

Runtime Diagnostics

- On-Demand
- Scheduled
- Health-Monitoring

Provides Generic Diagnostics and Health-Monitoring Framework

Action

Syslog
EEM
Call Home
SNMP Trap

Automated action based on diagnostics results

Platform Specific Boot up and Runtime Diagnostic

Detect Problems and Identify them before they result in Network Downtime!!

References

Cisco IOS Software

http://www.cisco.com/go/ios

Cisco IOS High Availability Website

Cisco.com: http://www.cisco.com/en/US/products/ps6550/products_ios_technology_home.html

Cisco IOS High Availability White Papers

http://www.cisco.com/en/US/tech/tk869/tk769/tech_white_papers_list.html

http://www.cisco.com/en/US/products/ps6550/prod_white_papers_list.html

Campus Network for High Availability Design Guide

http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/HA_campus_DG/hacampusdg.html

Cisco IOS Software

System-Level Resiliency: Feature Overview (Available Today)

Requirement	Feature	Benefits
Control/ Data Plane Resiliency	Cisco Nonstop Forwarding (NSF) with Stateful Switchover (SSO)	▪ Protection against unplanned control/data plane failure End-to-end protection, reduced downtime, operational costs, SLA penalties; no end-user disruption and loss of productivity
	NSF/SSO—MPLS LDP, L3VPN (Including InterAS and CsC), MPLS QoS	▪ Protection against failures in MPLS VPN networks Greater availability and uptime for MPLS services
	BGP Nonstop Routing	▪ Self-contained stateful switchover solution for routing protocols Seamless routing protocol recovery without requiring graceful restart protocol extension
	Gateway Load Balancing Protocol (GLBP)	▪ Remote site redundancy with load balancing More efficient resource utilization; easier deployment/maintainability
	Stateful IP Services—NAT, IPSec	▪ IP services application (NAT and IPSec) sessions are maintained even when a router fails
	Warm Reload	▪ Quick restoration, lower MTTR for single route processor hardware

System-Level Resiliency: Feature Overview (Available Today)

Requirement	Feature	Benefits
Planned Outages	Cisco IOS In-Service Software Upgrade	▪ Upgrade or downgrade of complete Cisco IOS software images with minimal impact on the data or control plane Minimizes and streamlines planned downtime, reduces operational cost significantly
	Warm Upgrade	▪ Faster upgrade or downgrade of complete Cisco IOS software images on single RP platforms
	Fast Software Upgrade	▪ Faster upgrade or downgrade of complete Cisco IOS software images on dual RP platforms
Link Resiliency	Line Card Redundancy with Y-Cables	▪ Reduce outage due to line card failures from a few hours to minutes/seconds Reduced downtime, lower operational costs, improved SLAs
	Link Bundling—EtherChannels, POS Channels	▪ Load shared protection from link failures Reduced downtime, improved SLAs

Network-Level Resiliency: Feature Overview (Available Today)

Delivers Industry-Leading Advances in Network Convergence Times, Protection, and Restoration

Feature	Benefits
Routing Convergence Enhancements IP Event Dampening, BGP Convergence optimizations, iSPF Optimizations	▪ Increased network stability for Enterprise and Service Provider IP networks Rapid L3 unicast convergence for real-time applications
NSF Awareness—BGP, OSP, IS-IS, EIGRP	▪ End-to-end high availability across network of networks Support for seamless recovery for routing protocols built into device
Multicast Sub-Second Convergence	▪ Recovers multicast path nearly instantly following unicast route recovery ▪ Increases network stability
MPLS Fast ReRoute (FRR)—Link and Node Protection	▪ Link and node protection for MPLS networks SONET equivalent resiliency (under 50ms recovery time)
NSF Awareness—LDP, RSVP	▪ Helper mode graceful restart capabilities for LDP, RSVP-TE to support MPLS high availability functionality
Fast IS-IS Convergence	▪ Sub-second convergence for IS-IS protocol Rapid network convergence for real-time applications
Fast OSPF Convergence	▪ Sub-second convergence for OSPF protocol Rapid network convergence for real-time applications
Bi-Directional Forwarding Detection	▪ Media and protocol independent mechanism to detect data plane connectivity Enables faster network convergence especially over shared media

Embedded Management: Feature Overview (Available Today)

- **Embedded intelligent event management for proactive maintenance**
- **Automation and configuration management for redaction of human errors**

	Feature	Benefits
Available Today	Embedded Event Manager (EEM)	▪ Predictive self-health monitoring capabilities built into Cisco IOS software ▪ Flexible policy support via TCL, delivering enhanced high availability, serviceability, and manageability support
	Component Outage On-Line (COOL) Management Information Base (MIB)	▪ Supports accurate and efficient monitoring of device component outages
	Embedded Resource Manager (ERM)	▪ Intelligent embedded resource management in Cisco IOS Software for enhanced resiliency and performance Graceful handling of out of resource (CPU, memory, IPC, etc.) error conditions Enhanced troubleshooting and debugging
	Configuration Rollback	▪ Enables return to known configuration states and reduces MTTR following configuration errors