



*TOMORROW
starts here.*

Cisco *live!*



Abstract

As Service Providers are deploying value-added triple-play or quadruple play services to maintain or generate a higher average revenue per user, overall Service Availability becomes increasingly important. High Availability techniques such as Fast Convergence or MPLS TE FRR have focused on raising the availability of the network core in the past. These techniques are being increasingly deployed in Ethernet Aggregation networks, for example by introducing MPLS TE FRR in the aggregation. Also, additional high-availability mechanism are being developed to enhance the resilience of the IP Edge against failures. Examples of new developments include ASR 9000 nV, BGP Prefix Independent Convergence for both the Core and Edge, or even stateful application inter-chassis redundancy mechanisms to overcome single-system outages. This Session aims to provide the audience with best current practices to increase service availability by deploying Cisco High-Availability mechanisms in both the Aggregation and the IP Edge. Traditional HA techniques such as NSF/SSO, BFD, Fast convergence or NSR are reviewed. The details of new technologies such as BGP PIC are discussed in depth. Furthermore, advanced topics such as achieving HA for Layer 4-7 services or stateful inter-chassis redundancy solutions are introduced. The Session provides the best current practices of deploying the tools offered by the Cisco High-availability toolset, in particular the deployment of ASR 9000 nV technologies clustering, which SPs may use to increase the availability of their IP Edge architecture.

Glossary



NHAT	next hop address tracking	EOBC	Ethernet out of band management
ACL	Access Control List	ESP	Embedded Services Processor
ACT	Active	EVC	Ethernet Virtual Circuit
APS	Automatic Protection Switching	EVDO	Evolution Data Only
ARP	Address Resolution Protocol	FECP	Forwarding Engine Control Processor
AS	autonomous System	FIB	Forwarding Information Base
ATM	Asynchronous Transfer Mode	FM	Forwarding Manager
bfd	Bi Directional Forwarding Detection	FR	Frame Relay
BNG	Broadband Network Gateway	FRR	Fast Re Route
BW	Bandwidth	FSOL	First Sign of Life
CC	Continuity Check	FWLB	Firewall Loadbalancing
CC	control connection	GEC	Gigabit Ether Channel
CDR	call detail record	GLBP	Global Load Balancing Protocol
CE	Customer Edge	GR	Graceful Restart
CE	Customer Edge	GRE	Generic Route Encapsulation
CF	checkpoint facility	GW	Gateway
CFM	Configuration and Fault Management	HA	High Availability
CLI	Command Line Interface	HSRP	Hot Standby Routing Protocol
CM	Chassis Manager	HW	Hardware
CP	Control Plane	IETF	Internet Engineering Task Force
CPLD	Complex Programmable Logic Device ?	IF	Interface
CSC	Carrier's Carrier	IGP	Internal Gateway Protocol
DHCP	Dynamic Host Configuration Protocol	IOCP	Input Output control Processor
DP	Data Plane	IOS	Internet Operating System
DPM	Defects per Million	IP	Internet Protocol
DSLAM	DSL Access Multiplexer	IPC	Inter process communication
E2E	End to end	ISG	Intelligent Services Gateway
ECMP	equal cost multipath	iSPF	incremental Shortest Path First
EEM	Embedded Event Manager	ISSU	in service software upgrade
EOAM	Ethernet OAM	IWF	Interworking function

Glossary (Cont.)



L2TP	Layer 2 transport protocol	NIC	network interface card
LAC	L2TP access concentrator	Nr	receive sequence number
LACP	Link aggregation control Protocol	Ns	send sequence number
LAN	Local Area Network	NSF	non stop forwarding
LC	Linecard	NSR	non stop routing
LDP	label Distribution Protocol	NVRAM	non volatile random access memory
LFA	loop free alternate	OAM	operations, administration and maintenance
LI	Lawful Intercept	OCE	Object Chain Element
LMI	Local management interface	OIR	online insertion and removal
LNS	L2TP network Server	OS	operating system
LOS	Loss of signal	PADR	PPP active discovery
LSDB	link state database	PE	provider edge
LSP	label switched path	PIC	prefix independent convergence
LTE	long term evolution	PIM	protocol independent multicast
MC LAG	multi chassis link aggregation	PPP	Point to point protocol
mcast	multicast	PS	power supply
MD5	message Digest algorithm 5	PSN	Packet Switched Network
MFIB	multicast forwarding information base	PTA	PPP termination and aggregation
MLD	multicast listener discovery	PVRSTP	Per VLAN rapid spanning tree
MME	mobile management entity	PW	pseudowire
MoFRR	Multicast Only fast reroute	QFP	Quantum flow Processor
MPLS	Multiprotocol label switching	RADIUS	remote authentication dial in user service
MRIB	multicast routing information base	RF	redundancy facility
MSC	mobile switching center	RMA	Return material authorization
MSPP	Multi-service provisioning platform	RNC	radio network controller
MST	Minimum spanning tree	RP	route processor
MTBF	mean time between failures	RPR	route processor redundancy
MTSO	mobile telephone switching office	RSP	route switch processor
MTTR	mean time to repair	RSVP	resource reservation protocol
NAT	network address translation	SAA	service assurance agent

Glossary (Cont.)



SBC	session border controller
SBY	standby
SGW	SAE gateway
SIP	Session initiation protocol
SLA	service level assurance
SLB	server loadbalancing
SP	service Provider
SPA	Shared port adapter
SPF	shortest path first
SRLG	shared risk link group
SSH	secure shell
SSO	stateful switchover
STP	spanning tree protocol
SW	software
T&C	terms & conditions
TCAM	ternary content addressable memory
TE	traffic engineering
TR	Traceroute
UC	unified communications
uRPF	unicast reverse path forwarding
VAI	virtual access interface
VC	Virtual Circuit
VCCV	VC connection verification
VIP	virtual IP
VLAN	virtual LAN
VMAC	virtual MAC
VPN	virtual private network
VRF	virtual routing and forwarding table
VRRP	virtual router redundancy protocol
WAN	wide area network



Best Practices to Deploy High-Availability in Service Provider Edge and Aggregation Architectures

BRKSPG-2402

Matthias Falkner, Distinguished Engineer, Technical Marketing

Cisco *live!*

Agenda BRKSPG-2402

- Motivation for High Availability in SP Aggregation Networks
- Network Level High Availability
- System High Availability
- Stateful Inter-chassis Redundancy
- Service High Availability
- Case Studies
- Summary and Conclusions





Motivation for High-Availability in the IP Edge and Aggregation

High Availability and Service Level Agreements

- Many SPs specify their SLAs in the T&Cs
- Important characteristic of both business and residential services
- Historically given for Core network, but expanding to end-end SLAs
- Metrics
 - Service Availability (averaged over time)
 - Mean time to repair (MTTR)
 - Packet Loss / Delay / Jitter
 - Increasingly also offering MOS

Examples

[AT&T](#)

[Verizon Business](#)

[Time Warner Telecom](#)

tw telecom IP Network Performance



Monthly IP Network Performance Averages

Month	Packet Delivery%	Latency (ms)	Jitter (ms)
Mar	100	37	0
Feb	100	37	0
Jan	100	37	0

 *National IP Network*
 *Regional Interconnection Sites*
 *Performance metrics available*
 *National Operations Center*

[Click here for historical data](#)

[Click here for Honolulu metrics](#)

[Click here for Methodology](#)

Metrics for individual city pairs are not the basis for SLA comparison. For SLA metrics, please see the Monthly IP Network Performance Averages.

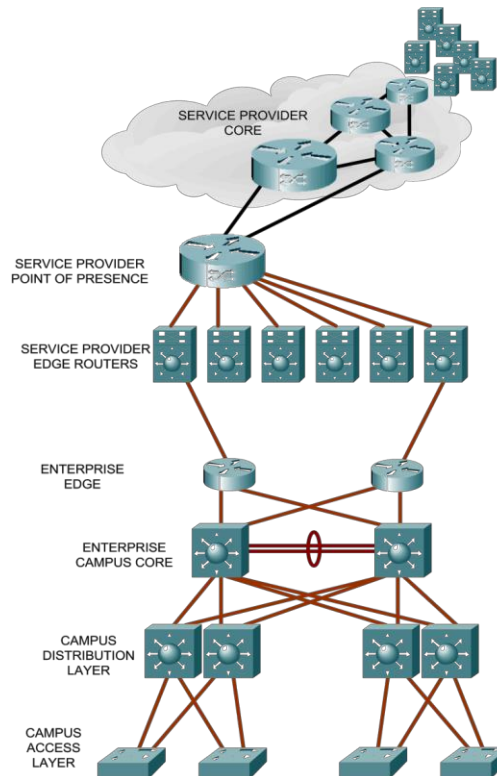




Availability Definitions

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

- ‘Uptime divided by the total time’ to create the percentage time your network is operational
- **MTBF** is **M**ean **T**ime **B**etween **F**ailure
When does it fail?
- **MTTR** is **M**ean **T**ime **T**o **R**epair
How long does it take to fix?





Calculated vs. Measured Availability

- **Calculated Availability** based on:

- Network design

- Component MTBF and MTTR

- different underlying models, simulations

- Cisco uses Industry standards to compute Hardware MTBF

- Basic Availability Calculation Formulae:

$$A_{\text{Series}} = \prod_{k=1}^N A_k = A_1 A_2 \dots A_N$$
$$A_{\text{Parallel}} = 1 - \prod_{K=1}^N (1 - A_k) = 1 - (1 - A_1) \dots (1 - A_N)$$

- **Measured Availability** based on:

- ICMP Reachability (E2E, Device)

- Cisco Service Assurance Agent (SAA)

- Trouble Ticket / Outage Log Analysis

- Observed Method: Shipping/RMA Method




Reduction of MTTR

- Stateful inter-chassis redundancy allows for additional resilience against System Failures
Interface Failures
- Issue is not really MTBF of hardware modules, but rather Line failures / optical path failures
Interface failures
Power outages
- Goal of stateful inter-chassis redundancy is sub-second failover with state preservation for applications

Product ID	Predicted MTBF (Hours)	Product ID	Predicted MTBF (Hours)
ASR1000-RP2	185857	A9K-2X100GE-SE	101500
ASR1000-SIP40	283225	A9K-36X10GE-SE	80520
ASR1000-ESP20	178658	A9K-36X10GE-TR	82560
ASR1000-ESP40	118790	A9K-40GE-B	141680
ASR1006	1986649	A9K-4T-B	150000
ASR1013	1528905	ASR-9006-DC	712690
SPA-1x10GE-L-ITUC	437278	ASR-9010-DC	392470
SPA-10x1GE	590295	SPA-1X10GE-L-ITUC	437278
SFP-GE-L	2294776	XFP-10GER-OC192IR	2294776
SPA-10X1GE	590295	XFP-10GZR-OC192LR	2294776



Device Availability Calculation



$$IOS = \frac{30.000}{30.000 + 0.1} = 0.999997$$

$$IF1 = \frac{1.120.000}{1.120.000 + 4} = 0.999996$$

$$IF2 = \frac{600.000}{600.000 + 4} = 0.999993$$

$$CPU = \frac{490.000}{490.000 + 4} = 0.999992$$

$$BB = \frac{460.000}{460.000 + 8} = 0.999983$$


$$P/S = \frac{750.000}{750.000 + 4} = 0.999995$$

$$\text{System Availability} = 0.999997 * \dots * 0.999983 * (1 - (1 - 0.999995)^2) = 0.999961 = 99.9961\%$$

Calculated MTBF Values from Cisco Database



Network Availability Calculation



1

Availability of R1 and R2 in Series
 $= (0.999961 * 0.999961) = 0.99992175$

2

Availability of Parallel Network Path (R1-R4)
 $= 1 - ((1 - 0.999921)(1 - 0.999921)) = 0.999999994$

3

Network Availability = 99.9999%
Only Based on Device Availability Values

but **not considered**:

- Links (WAN, LAN)
- Computer NICs
- Computer OS
- Computer Applications

Cisco High-Availability Focus

System Level Resilience

- increase MTBF using resilient Hardware & software
- minimize MTTR for system failures (RP, Linecards and Software)
- Mitigate planned outages by providing hitless software upgrades

Network Level Resilience


- in the core and where redundant paths exist
- Deliver features for fast network convergence, protection & restoration

Embedded Management and Automation

- Embed intelligent event management for proactive maintenance
- Automation and configuration management to reduce human errors

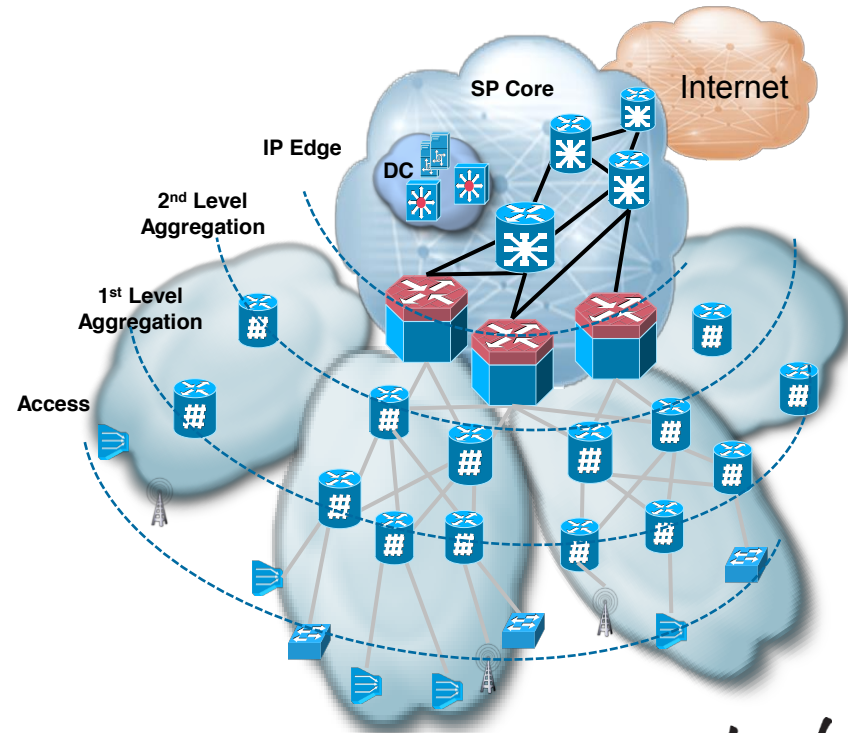
Examples for System Level HA Mechanisms

- RPR, SSO
- NSF, **NSR**
- SSO Multirouter APS
- Stateful NAT/IPSec/Firewall/SLB stateful failover within single chassis
- MPLS HA (L3VPN, L2VPN, InterAS, CSC, TE, FRR)
- IOS / IOS XR / IOS XE ISSU, dual IOS XE
- **ASR 9000 nV**




Examples for Network Level HA Mechanisms

- Network Design Resiliency: Dual-homing, APS, GEC, MC-LAG
- Event Dampening
- Fast Convergence: iSPF Optimization (OSPF, IS-IS), BGP Optimization, Fast BGP Convergence, **BGP PIC**
- Graceful Restart (MBGP, OSPF, RSVP, LDP)
- EMCP, Anycast, dual RR
- VRRP/HSRP/GLBP/SLB/FWLB
- MPLS High Availability: LDP Graceful Restart, MPLS/VPN NSF
- **BFD**
- **MPLS FRR Path Protection**
- **MoFRR**
- IP FRR
- **Pseudowire Redundancy**
- Spanning Tree (MST, PVRSTP...)



Cost of High Availability

- Designing a network for higher Service Availability comes at a cost
 - Redundant Network Elements
 - Redundant Links
 - Redundant System Components (route processors, forwarding processors, power supplies, etc.)
- Operational costs
 - Lower steady-state Utilization levels
 - Increased configuration and management
 - Tighter maintenance windows



Cost of High-Availability: Example

- Large SP Network

Residential Services (3-Play)

10M Subscribers


1.25 Mbps / subscriber

- Up to 96% increased CAPEX for full redundancy!

Opex increased due to higher number of network elements

Redundancy Scheme	Total Cost \$M	Chassis Costs \$M	Interface Costs (SPA, SFPs), \$M	Number of nodes
No redundancy	\$1,232	\$529	\$704	4658
Access NW uplink redundancy (Agg1, Agg2, Agg3)	\$1,250	\$529	\$721	4658
AN uplink redundancy	\$1,563	\$531	\$1,032	4680
Access Network node redundancy (Agg1, Agg2, Agg3)	\$2,423	\$1,044	\$1,379	9222
Edge link redundancy	\$2,425	\$1,044	\$1,381	9222
Edge Node redundancy	\$2,437	\$1,056	\$1,381	9296

Values for AN, Agg1, Agg2, Agg3 and Edge nodes only (No Pp-routers). Cumulative redundancy Schemes, GPL




Subscribers	AN	Agg1	Agg2	Agg3	BNG	P
Locations	200,000	4000	500	74	74	74
System Type	Generic	ASR 9000	ASR 9000	ASR 9000	ASR 1013	CRS-3

A nighttime photograph of a city street. In the background, there are modern buildings with lit windows and a pedestrian bridge with blue lighting. The foreground is dominated by long, colorful light trails from moving vehicles, creating a sense of motion and energy. The text is overlaid on a dark horizontal band across the middle of the image.


Network High Availability Best Practices Failure Detection and Network Recovery

HA Network Map



L0/1	Failure Detection	Interrupts	Loss of Signal	Interrupts	Loss of Signal	Interrupts	
	Recovery	Module Redundancy	Path diversity / dual homing	Module Redundancy	Path diversity / dual homing	Module Redundancy	

HA Network Map



		AN	<=>	AGG	<=>	Edge	<=>	Core
L4-7 App	Failure Detection					Keepalives		
	Recovery					Stateful App Redundancy		
L3	Failure Detection	BFD				Keepalives	BFD, Keepalives	
	Recovery	MPLS TE FRR, IP Event Dampening, Fast Convergence, IP FRR,				nV NSR / NSF HSRP / VRRP/GLBP/SLB/FWLB Multicast HA	ECMP, iSPF, BGP PIC Core / Edge, IP / MPLS FRR, LNS Load sharing / Anycast / Dual RR, Fast Hello	
L2	Failure Detection	EOAM, (VCCV)		Keepalives	VCCV, EOAM, MPLS Ping / TR	Keepalives	EOAM, MPLS Ping / TR	
	Recovery	GEC / APS / MC-LAG		PW redundancy Bridge Domains	GEC / APS / MC-LAG	PPP / FR / ATM / HDLC / GE SSO	GEC, APS, MC-LAG	
L0/1	Failure Detection	Interrupts	Loss of Signal	Interrupts	Loss of Signal	Interrupts		
	Recovery	Module Redundancy	Path diversity / dual homing	Module Redundancy	Path diversity / dual homing	Module Redundancy		



Failure Detection Best Practices

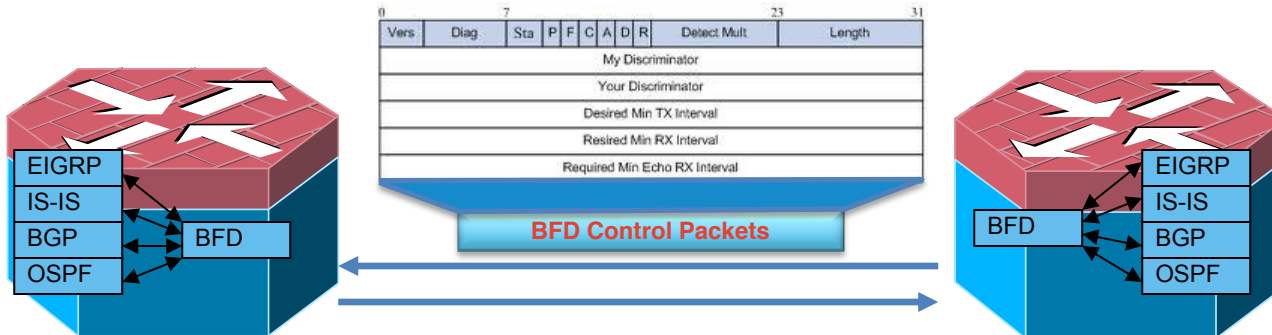
Best Practice (BP): Link Failure Detection Mechanisms

- Link failure detection based on multiple mechanisms
1. Loss of Signal (LOS) key to link failure detection
O(ms) with interrupt driven LoS detection on ASR 9000
Carrier Delay may be used to become resilient to link flaps
 2. BFD for end-end (e.g. PE-CE)
 3. Ethernet OAM or MPLS OAM for Aggregation network


Media type	CC CP	CC DP	Loopback	Performance	Traceroute
Ethernet Last Mile	IEEE 802.1ah		-	-	-
Ethernet Provider Bridge	IEEE 802.1ag (MAC: Broadcast Domain)				
MPLS LDP	LDP Hello	BFD, Y.1713, Y.1711	LSP Ping	-	LSP TR
MPLS TE	RSVP Hello			-	
MPLS PW	LDP Hello	BFD, Y.1711	VCCV Ping	-	-
IPv4	IGP/BGP Hello	BFD	IP Ping	-	IP TR

BP: Run BFD for L3 Failure Detection

- Accelerates convergence by running fast keepalives in a consistent, standardized mechanism across routing protocols
- Lightweight hello protocol
- Neighbors exchange hello packets at negotiated regular intervals
- Configurable transmit and receive time intervals
- Unicast packets, even on shared media
- No discovery mechanism
- BFD sessions are established by the clients e.g. OSPF, IS-IS, EIGRP, BGP, ...
- Client hello packets transmitted independently




BP: Ethernet OAM Failure Detection



- **Ethernet LMI:** Automated configuration of CE based on EVCs and bandwidth profiles
L2 connectivity management
- **IEEE 802.3ah:** When applicable, physical connectivity management between devices.
- **IEEE 802.1ag: Connectivity Fault Management (CFM)**
Uses Domains to contain OAM flows and bound OAM responsibilities
Provides per EVC connectivity management and fault isolation
Three types of packets: Continuity Check, L2 Ping, L2 Traceroute

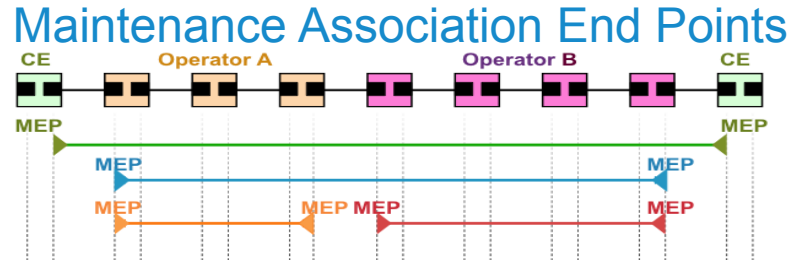
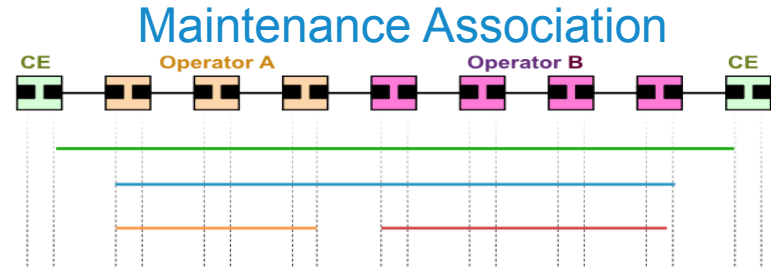
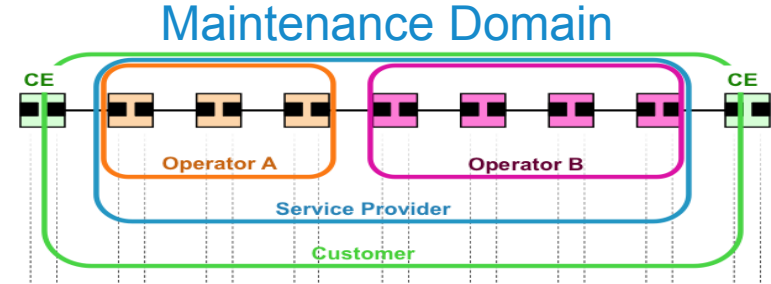
Ethernet OAM Overview

- E-LMI - Provides protocol and mechanisms used for:
 - Notification of EVC addition, deletion or status to CE
 - Communication of UNI and EVC attributes to CE
 - CE auto-configuration
 - Notification of Remote UNI name and status to CE
- IEEE 802.3ah
 - OAM Discovery
 - Link Monitoring
 - Fault Signaling
 - Remote MIB Variable Retrieval
 - Remote Loopback
- IEEE 801.3ag (CFM)
 - Family of protocols that provides capabilities to detect, verify, isolate and report end-to-end ethernet connectivity faults
 - Protocols (Continuity Check, Loopback and Linktrace) used for Fault Management activities



IEEE 802.1ag CFM Concepts

- **Nested Maintenance Domains (MDs)**
 - break up the responsibilities for network administration of a given end-to-end service
 - Defined by operational boundaries
 - Nest & touch, but do not intersect
- **Maintenance Associations (MAs)**
 - monitor service instances under a given MD
 - Defined by set of MEPs at the edge of a domain
 - Identified by {MA Name + MD ID}
- **Maintenance Association End Points (MEPs)**
 - generate and respond to CFM PDUs
 - Define boundaries of MD
 - Initiate & respond to CFM PDUs
- **Per-Maintenance Association multicast “heart-beat” messages**
 - Carries status of port on which MEP is configured
 - Uni-directional (no response required)
 - Transmitted at a configurable periodic interval by MEPs
- **Catalogued by MIPs at the same MD-Level and service, Terminated by remote MEPs in the same MA**



ITU-T Y.1731 Overview

- OAM Functions for **Fault Management**

- Ethernet Continuity Check (ETH-CC) (Y.1731 adds unicast CCM)

- Ethernet Loopback (ETH-LB) (Y.1731 adds multicast LBM)

- Ethernet Linktrace (ETH-LT)

- Ethernet Remote Defect Indication (ETH-RDI)

- Ethernet Alarm Indication Signal (ETH-AIS)

- Ethernet Locked Signal (ETH-LCK)

- In addition: ETH-TEST, ETH-APS, ETH-MCC, ETH-EXP, ETH-VSP

- OAM Functions for **Performance Management**

- Frame Loss Measurement (ETH-LM)

- Frame Delay Measurement (ETH-DM)

BP: Deploy VCCP For MPLS PW Domains in the Aggregation

- Checks connectivity between egress and ingress PEs
- VCCV allows sending control packets in band of pseudowires (PW)
 - Signaling: communicate VCCV capabilities as part of VC label
 - Switching: cause the PW payload to be treated as a control packet
- VCCV capability is negotiated when the AToM tunnel is brought up
 - depends on the LDP peer and the VC type
 - Both endpoints must have the same capabilities
- Marks the payload as control packet for switching purpose; packet follows the PW data path
- Control packets sent over the AToM tunnels are intercepted by the egress PE

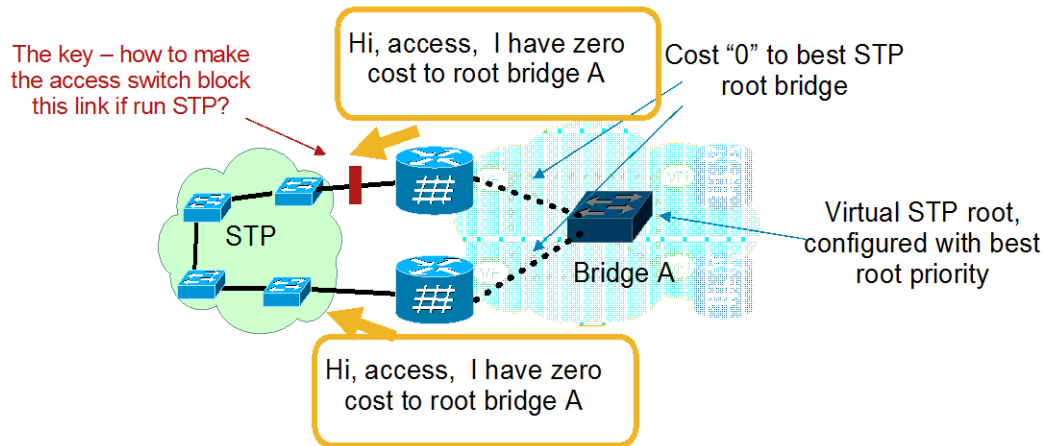
Type 1 (in-band vccv)	To signal in-band VCCV [RFC4385] using PW ID from PW Control Word
Type 2 (out-of-band VCCV)	Signal out-of-band VCCV inserting MPLS router alert label between tunnel and PW Labels
Type 3 (TTL expiry)	Manipulate and Signal TTL exhaust (TTL == 1) for multiple switching point PEs



Recovery and Resilience Best Practices

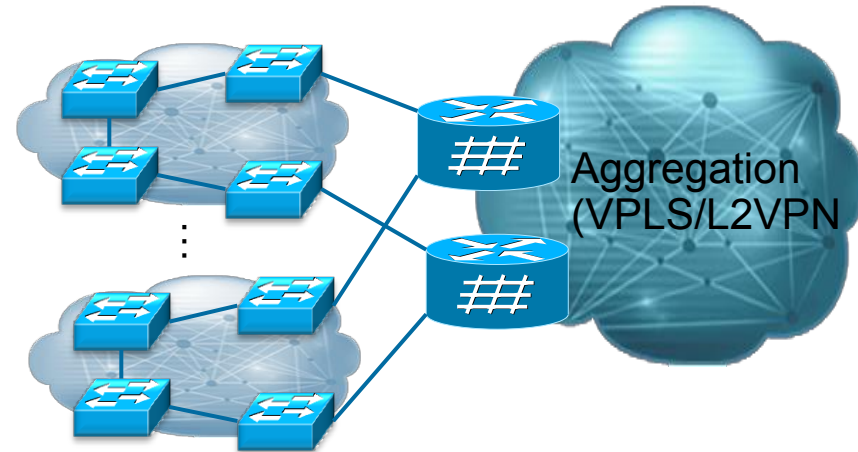
BP: Run MST Access Gateway towards Access

- MST access Gateway **sends pre-canned BPDUs** into the access network at the hello timer BPDUs have zero cost to the best STP root bridge (the latter is statically configured)
Forces access network to block one of the access links.
- Also **snoops the MSTP Topology Change Guard** from the access network, flushes the local MAC and triggers VPLS MAC withdrawal (in case of VPLS)



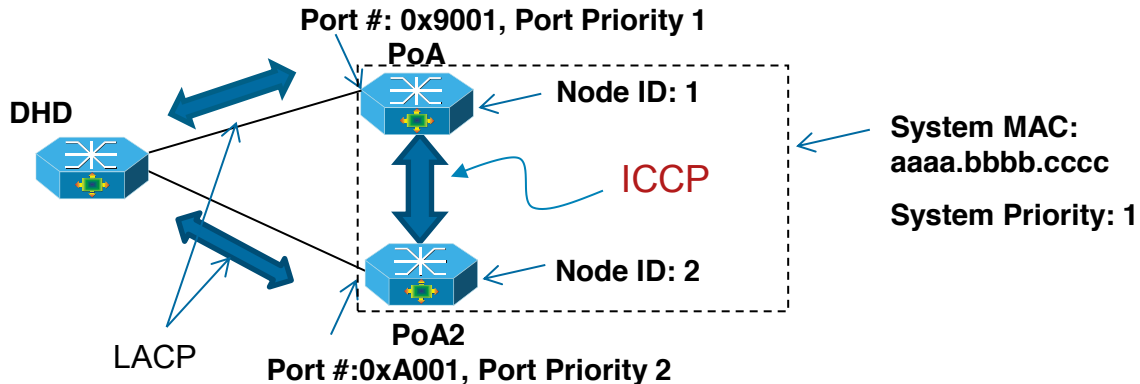
MST Access Gateway - Motivation

- Access Gateway aggregates many L2 Access networks into L2/L3-based aggregation networks
- Problem: consistent L2 Topology view?
- Could run 1 STP per access network
Scalability on access Gateway
- Could maintain single access topology across all access networks
Maintenance difficult
- Could tunnel BPDUs between legs of access network
Convergence issues upon topology changes O(sec)



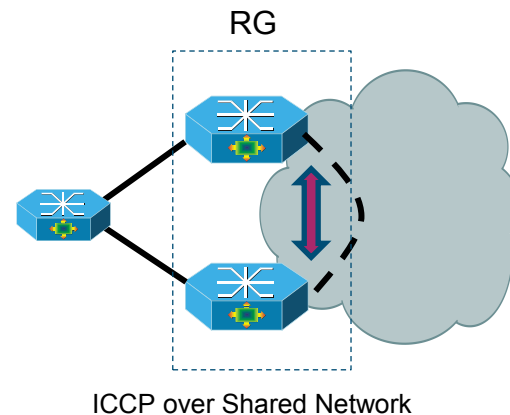

BP: Access Link Redundancy using Multi-chassis LAG

- mLACP uses **ICCP** to **synchronize LACP configuration & operational state** between PoAs, to provide DHD the perception of being connected to a single switch
- All PoAs use the **same System MAC Address & System Priority** when communicating with DHD
 - Configurable or automatically synchronized via ICCP
- Every PoA in the RG is configured with a **unique Node ID** (value 0 to 7). Node ID + 8 forms the most significant nibble of the Port Number
- For a given bundle, all links on the same PoA must have the same Port Priority



Inter-chassis Communication Protocol

- ICCP allows two or more devices to form a ‘Redundancy Group’
- ICCP provides a control channel for synchronizing state between devices
- ICCP uses TCP/IP as the underlying transport
 - ICCP rides on targeted LDP session, but MPLS need not be enabled
- Various **redundancy applications** can use ICCP:
 - mLACP
 - Pseudowire redundancy



Network Convergence — Why It Takes So Long

Detection of Link layer failure


Report failure to Route Controller

Generate and flood an LSP

Trigger and Compute an SPF

Communicate new FIB entries to linecards

Install new FIB entries into linecard HW path



Network Convergence — Why It Takes So Long

Detection of Link layer failure

Optimize IGP
Convergence

Report failure to Route Controller

Optimize LDP & BGP
Convergence


Generate and flood an LSP

BGP PIC

Trigger and Compute an SPF

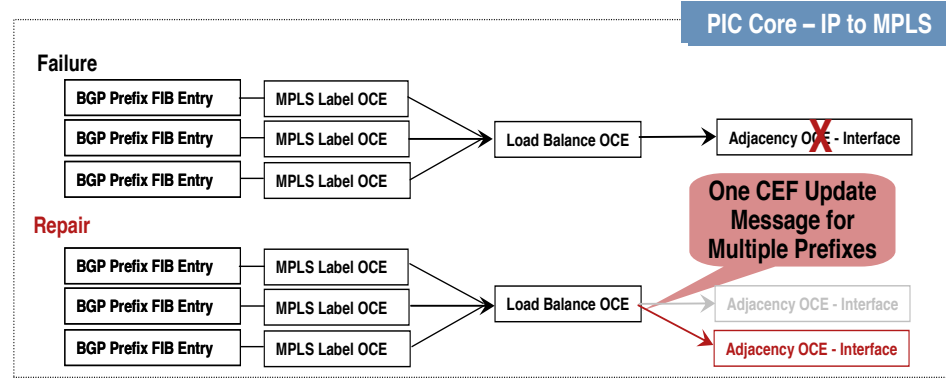
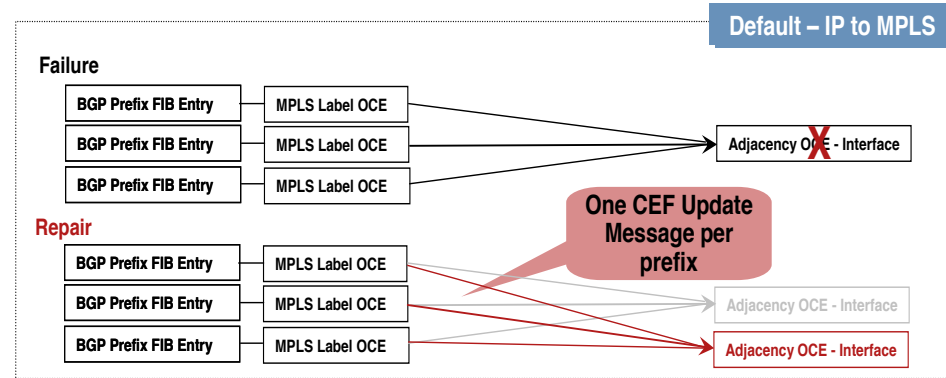
Communicate new FIB entries to linecards

Install new FIB entries into linecard HW path



Hierarchical CEF

- Optimizes the data plane for sub-second convergence
- CEF Data Structure Enhancements
 - Solves the FIB Download Convergence Bottleneck
 - LSP and Prefix Independent
 - Optimizes FIB
- Hierarchical CEF Technologies
 - MPLS-FRR
 - IP-FRR
 - BGP PIC Core
 - BGP PIC Edge
- Non-Hierarchical CEF Technologies
 - MPLS Path Protection



BP: Deploy MPLS FRR Node Protection in the Aggregation

- Key Features

 - Fast Convergence for Link and Node Failures

 - Supported Across all Network Topologies

 - MPLS-TE Traffic Management

 - SRLG

 - BW Reservation

 - Per Tunnel Traffic Statistics

- Caveats

 - Requires MPLS and MPLS-TE


 - No Protection for Ingress or Egress Tunnel Failures

 - Requires Pre-Computed Backup Paths

 - Requires “(n-1)!” Tunnels for Full Protection

- Applicability

 - Protecting Links in the aggregation network



Cisco *live!*

BP: Deploy MPLS Path Protection in the Aggregation

- Key Features

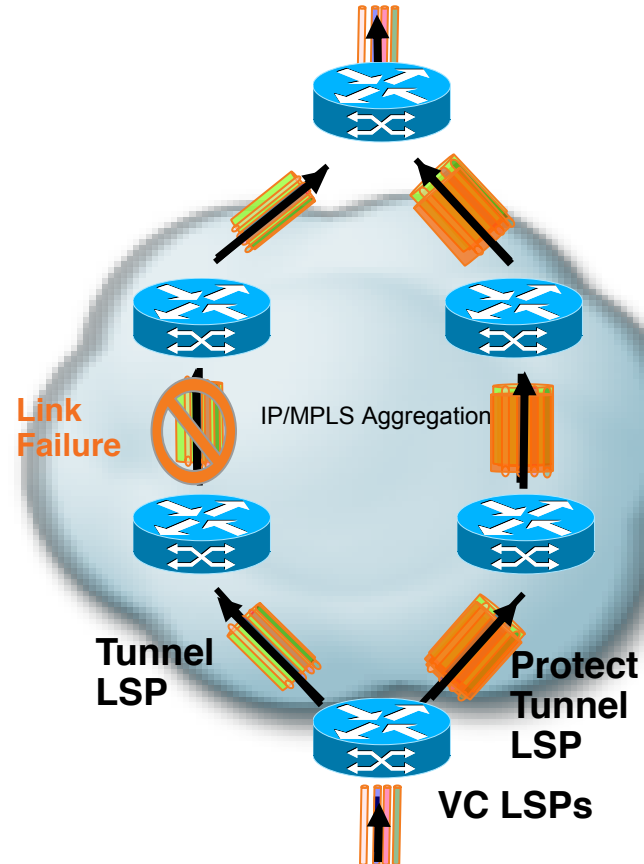
- Optimized for Ring Topologies
- Utilizes Pre-Signaled Backup Tunnel
- MPLS-TE Traffic Management
- SRLG
- BW Reservation
- Per Tunnel Traffic Statistics

- Caveats

- Requires MPLS and MPLS-TE
- No Protection for Ingress or Egress Tunnel Failures
- Convergence Dependant on IGP Prefixes and L2VPN LSPs Under Protection


- Applicability

- Protecting Ring Topologies

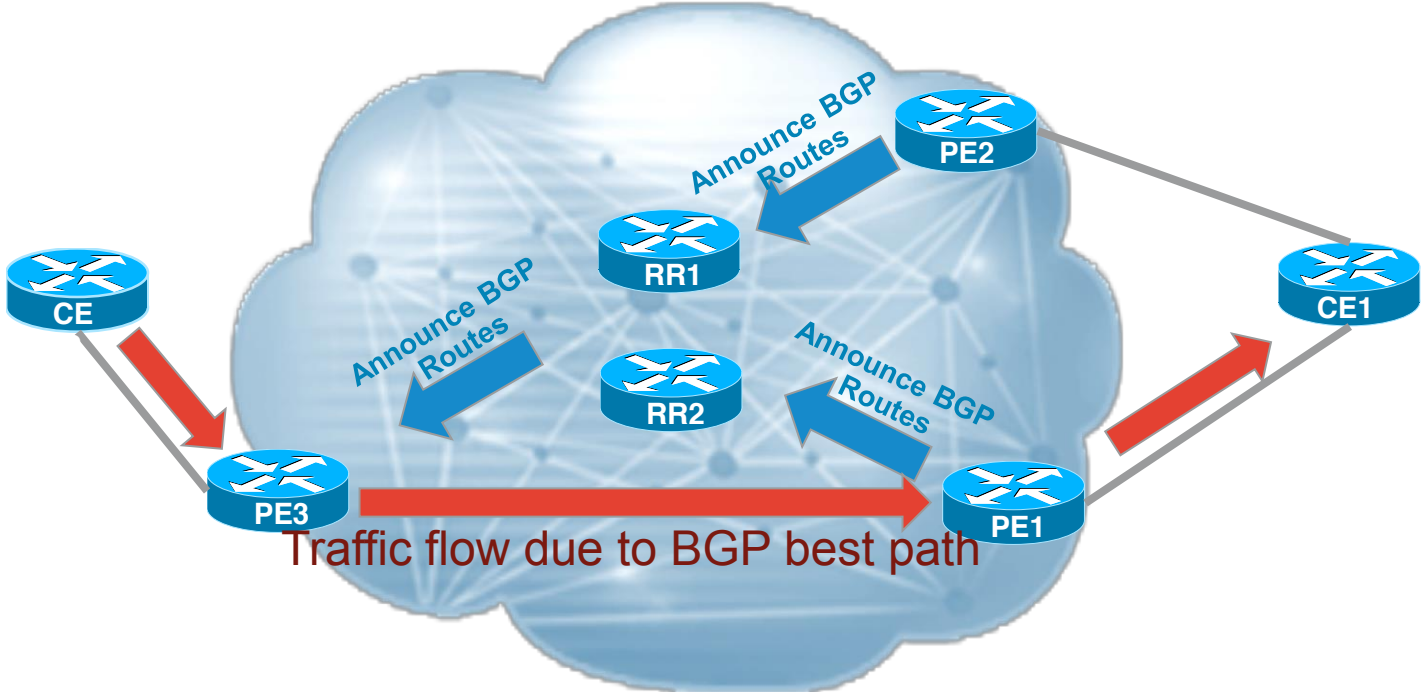


BP: Configure BGP PIC at the Edge

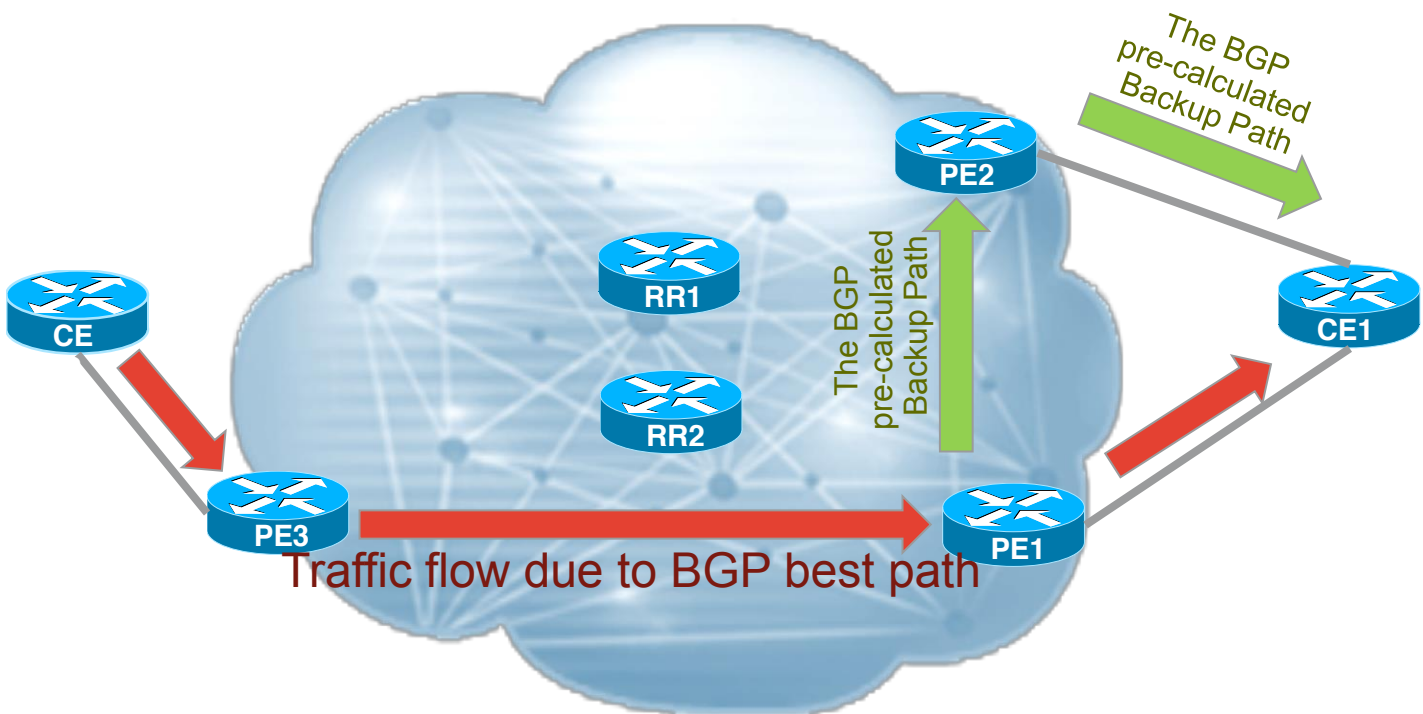
- Optimizes BGP Convergence for BGP Next-Hop Change
 - PE to CE Link Failures
 - PE Node Failures
 - CE Node Failures
- Applicability
 - PE Routers
- Requires “bgp advertise-best-external” to enable



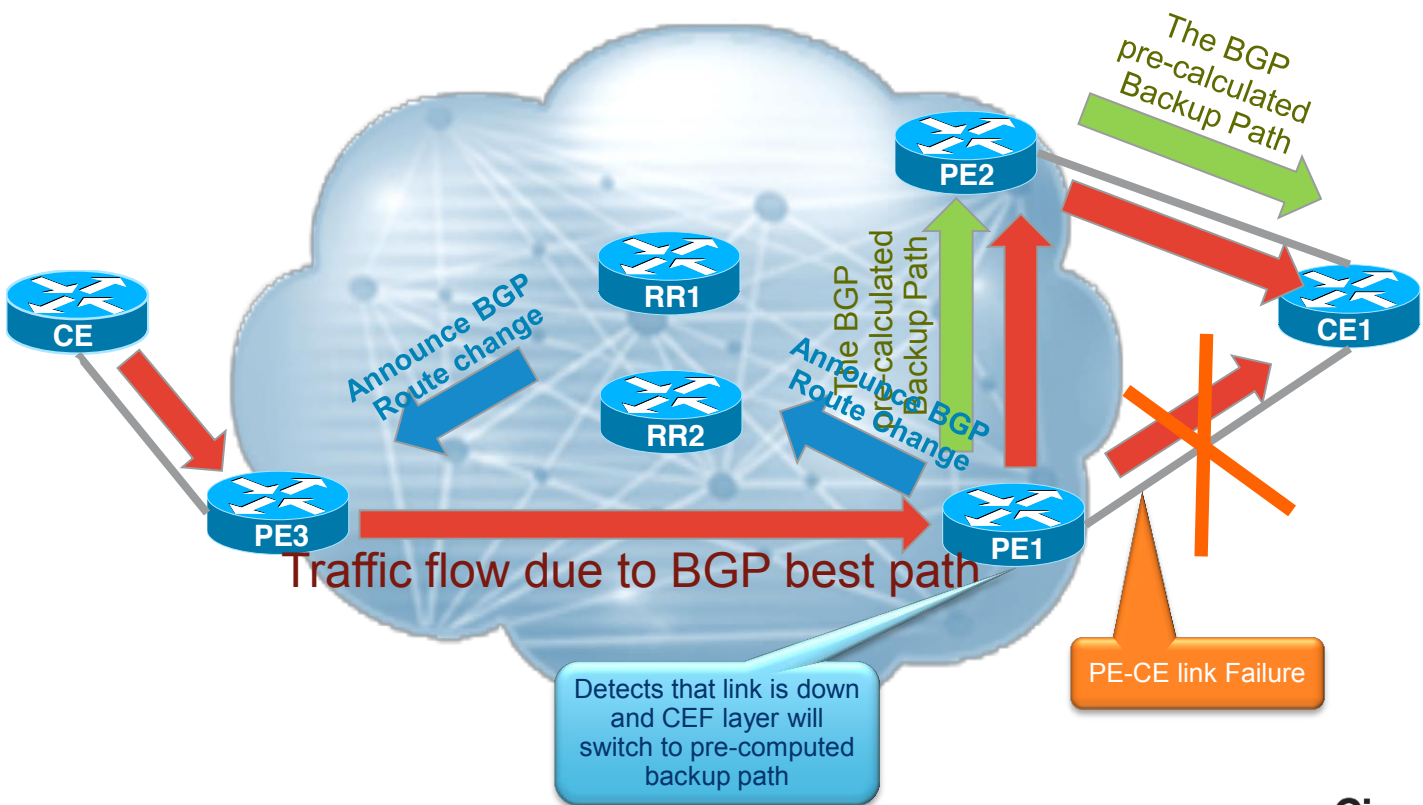
BGP PIC Edge PE-CE Link Protection



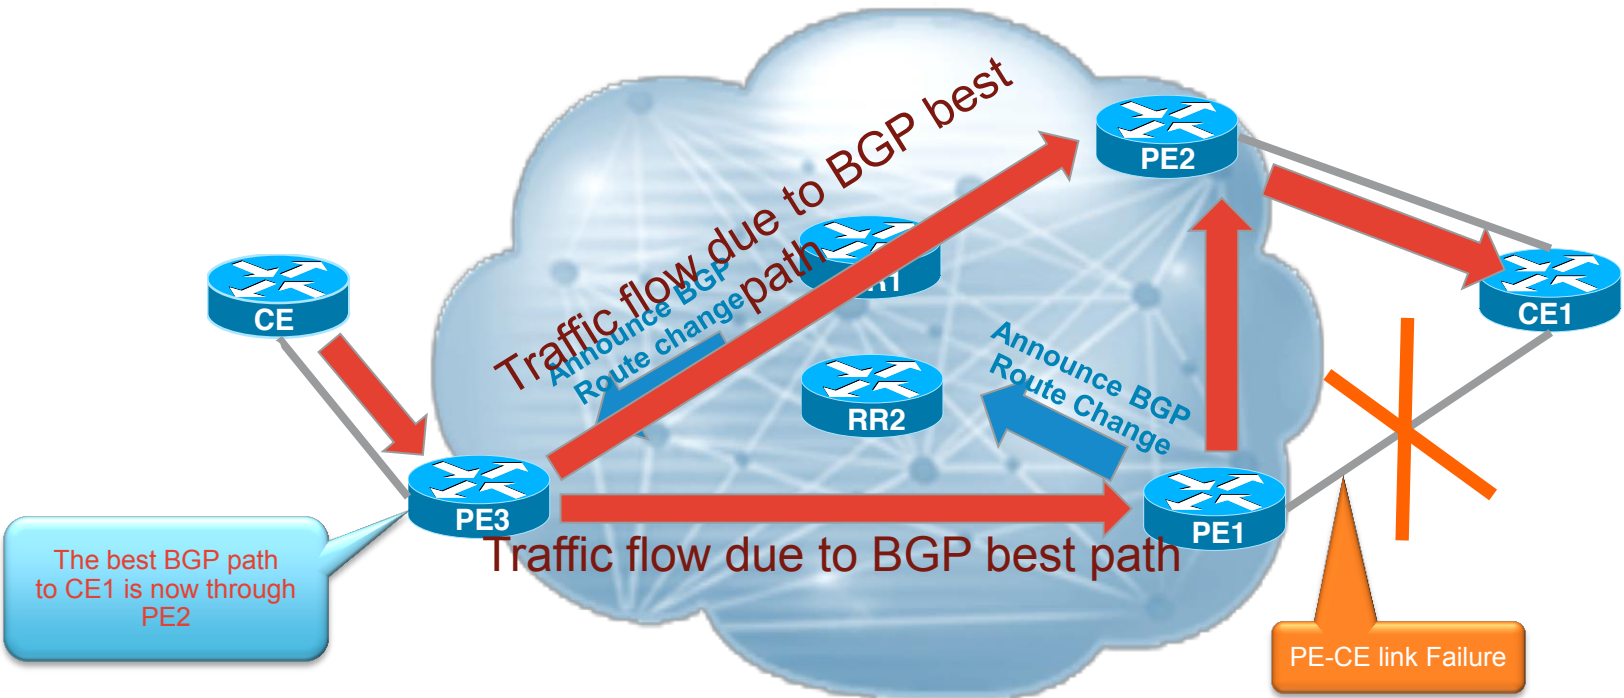
BGP PIC Edge PE-CE Link Protection



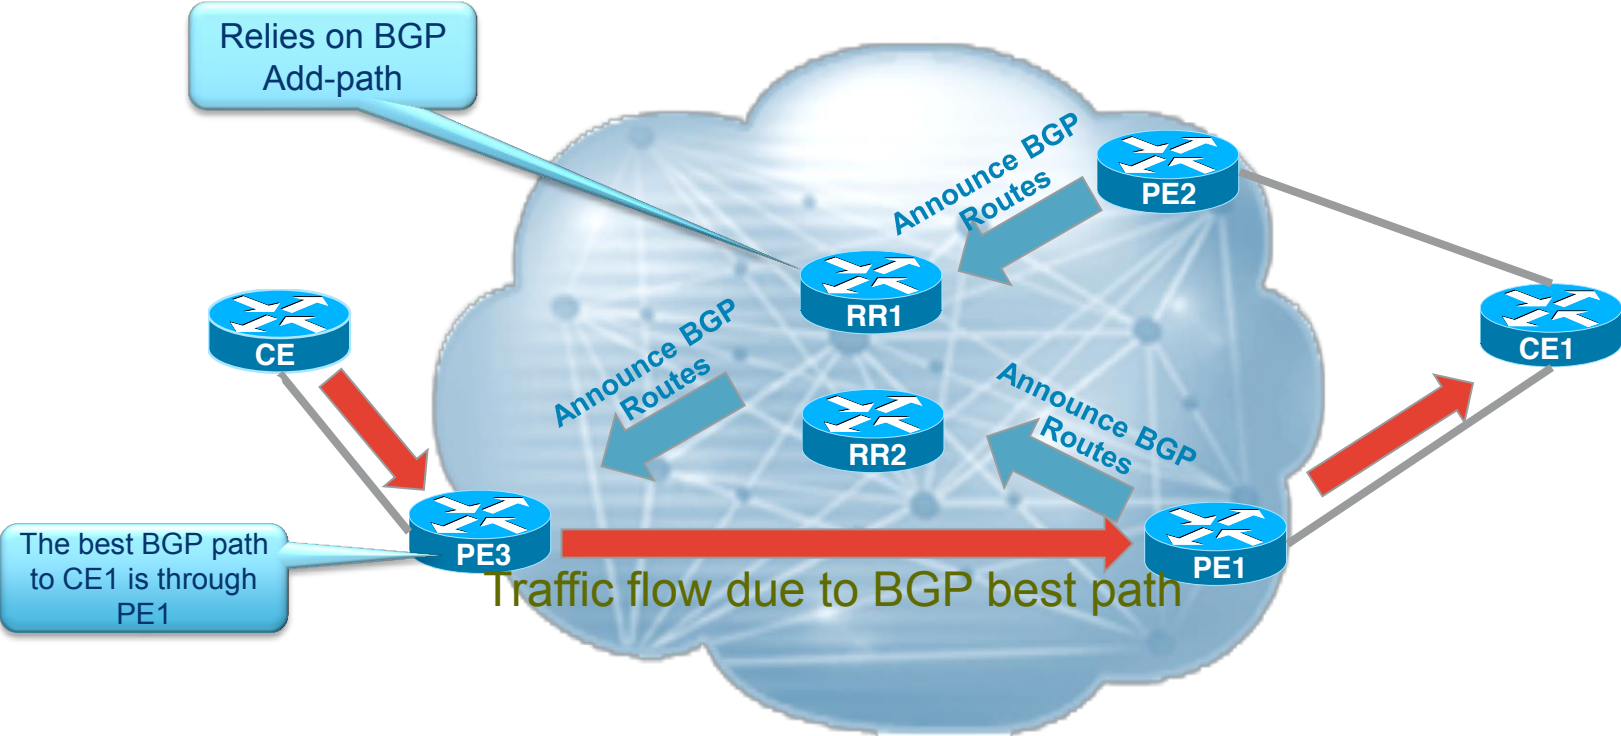
BGP PIC Edge PE-CE Link Protection



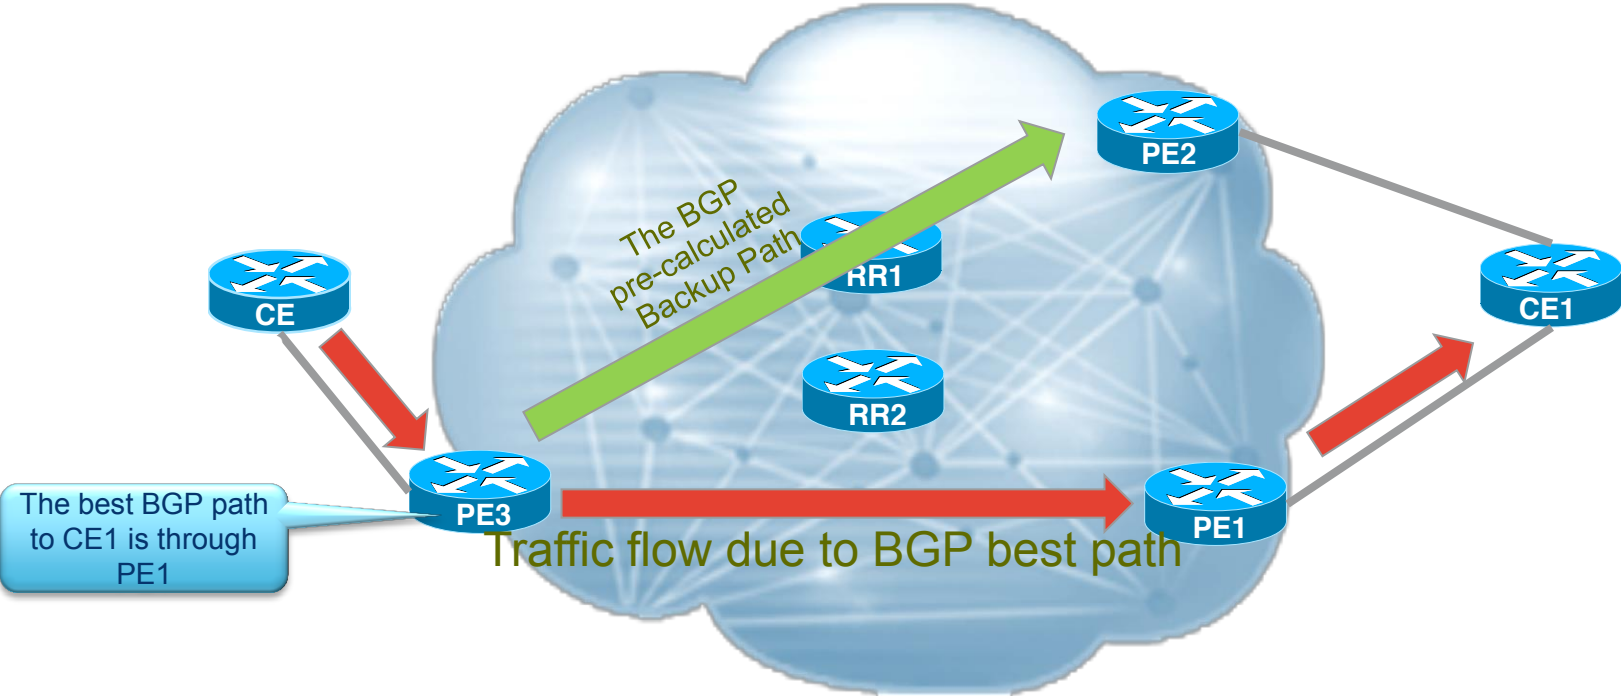
BGP PIC Edge PE-CE Link Protection



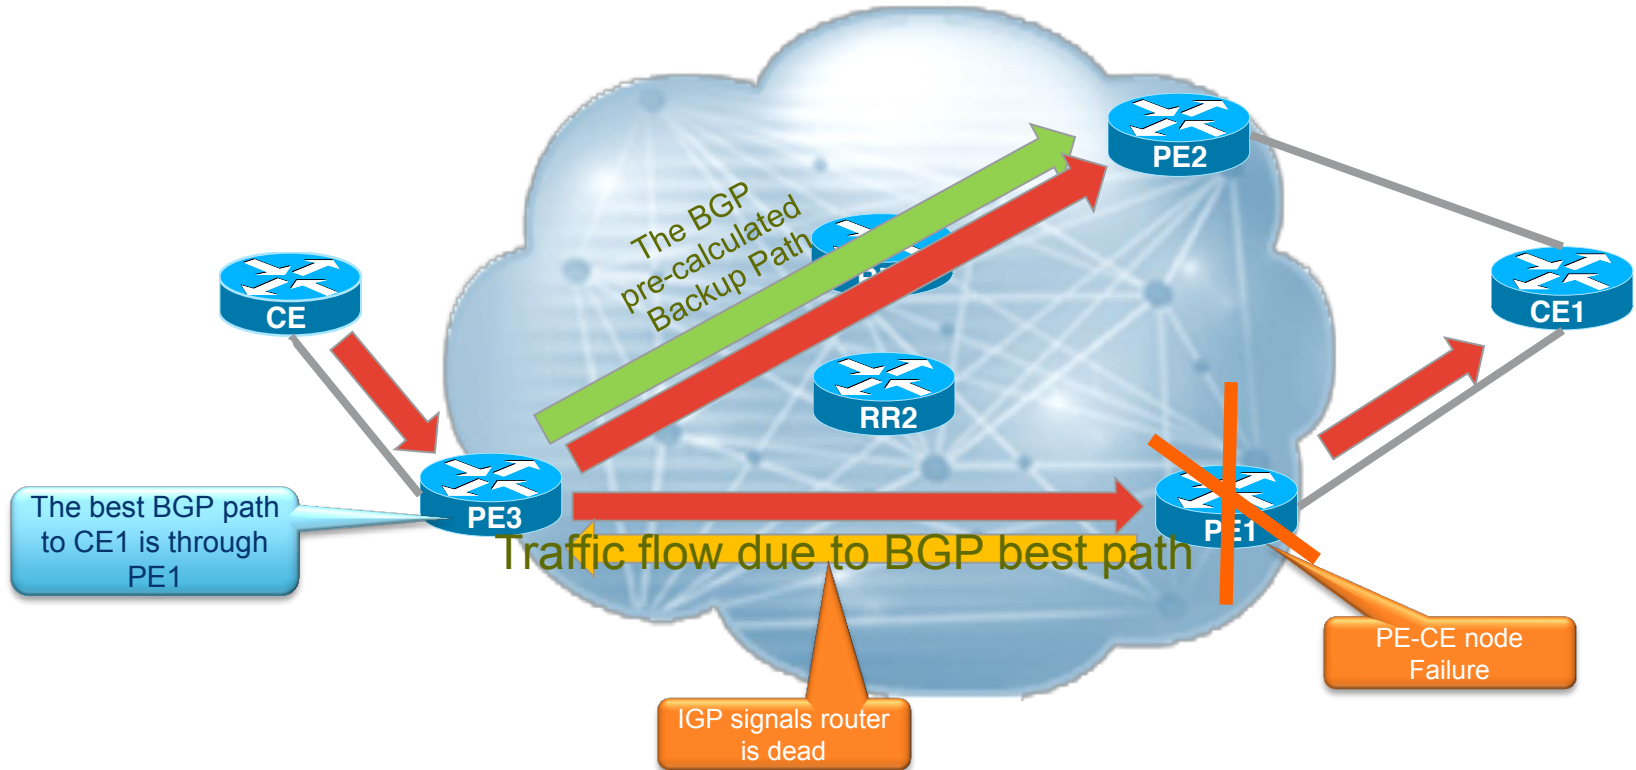
BGP PIC Edge PE-Node Protection



BGP PIC Edge PE-Node Protection



BGP PIC Edge PE-Node Protection

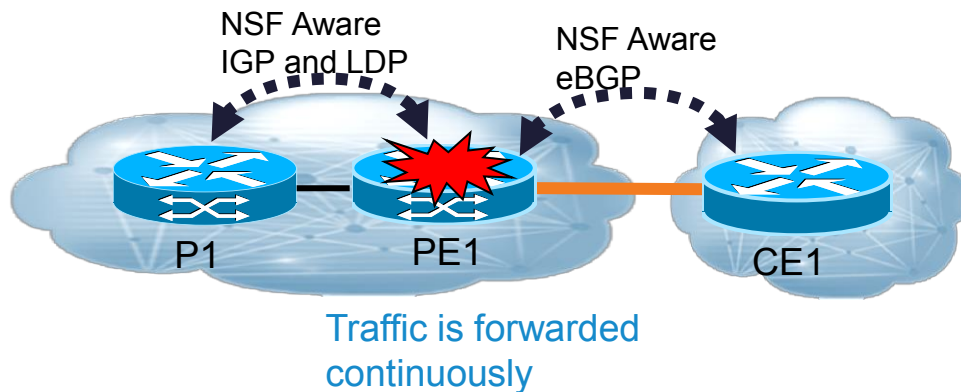




BGP PIC Edge Notes

- Supported for IPv4/v6 and VPNv4/v6 in IOS XE
 - Not supported for L2VPN and mVPN address families
 - Not supported for MPLS VPN Inter-AS Option B
 - Not supported in combination with Best External
- Supported for IPv4/v6 and VPNv4/v6 in IOS XR
 - Requires IOX XR 4.2.1
 - Including L2VPN, L3VPN, VPLS, 6PE, 6VPE
- Failures detected using BFD or IGP

Non Stop Forwarding (NSF)



- Routers to maintain forwarding state when communication between them is lost
- Routing sessions are established with NSF aware peers. Upon HA event, neighboring peers maintain forwarding until routing sessions are reestablished.
- Copy of FIB maintained on secondary and used on failure for continuously traffic flow.
- Requires neighboring routers to be NSF aware.



System High Availability Best Practices

BP: Deploy Hardware Redundancy!

- Redundant hardware components
 - Power Supplies
 - Route Processors
 - Forwarding Processors
 - Switching Matrix
 - SPA Interface Cards
- Interface Redundancy typically achieved using IEEE 802.3ad / LACP or APS
- Hardware Redundancy needs to be complemented by Software redundancy Features
- Cisco Platforms supporting hardware redundancy



CRS-3



ASR 9000



ASR 5000



ASR 1000




Cisco 12000



Cisco 7600

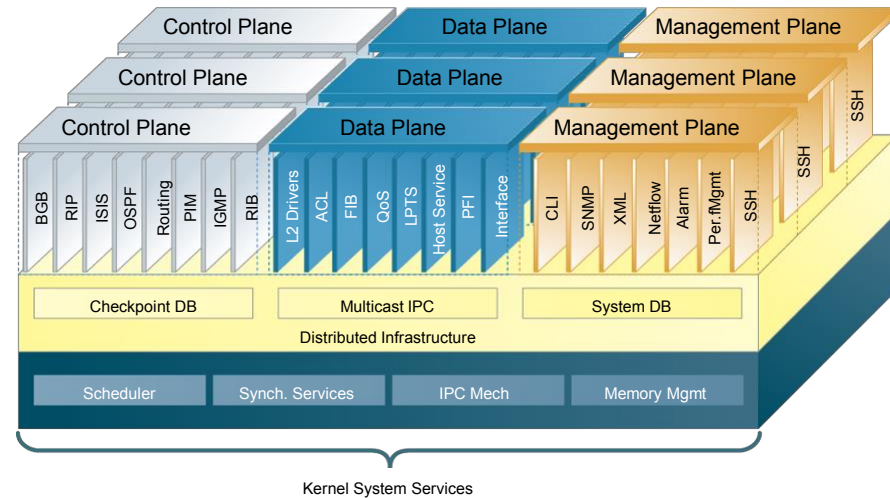
ASR 9000 System Architecture



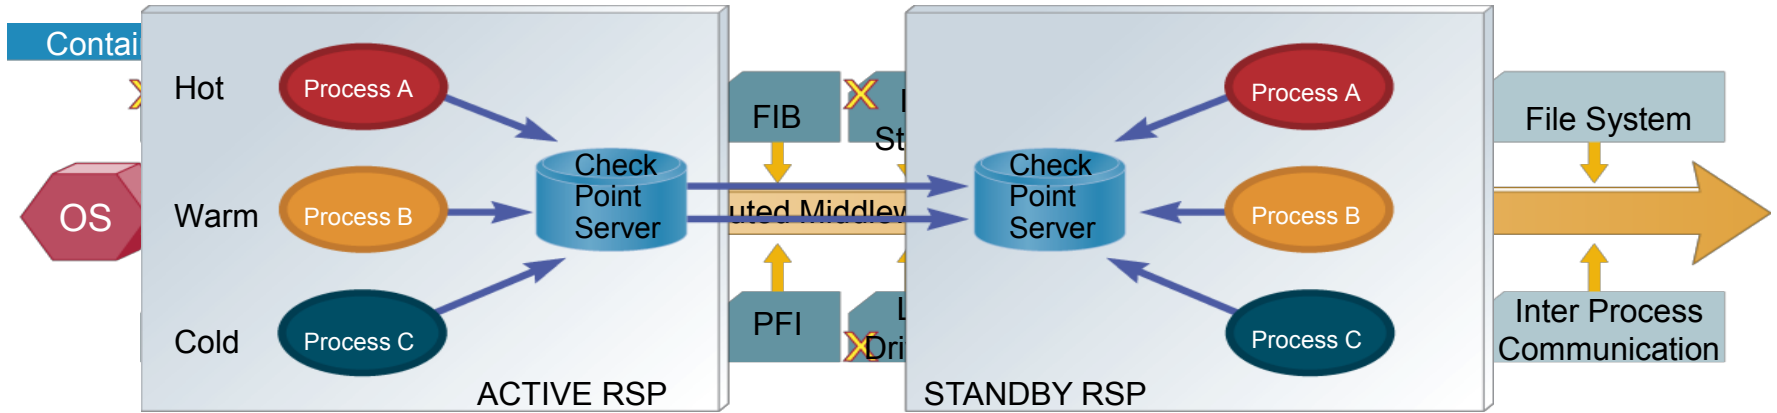
- **Distributed Forwarding Plane for Performance**
 - Up to Eight Linecards
 - (Autonomous Forwarding)
- **Distributed IOS®XR based Control Plane for Scale**
 - Dual Route Switch Processors (RSPs)
 - Dual-Core CPU on Each Linecard
- **Active/Active Switch Fabric for HA**
 - Non-blocking Memory-less Fabric
 - Service Intelligence with Hi / Lo Priorities,
 - Unicast & Multicast Recognition, and VoQ's
 - Redundant EOBC, Fan Trays, Power supplies (not shown)

ASR 9000 Software Architecture

- IOS XR offers distributed software architecture
 - Distributed control plane (e.g. routing)
 - Linecards have independent FIB/IDB for only local interfaces
 - Two stage forwarding
- Control plane information pushed from RSP to linecards using reliable multicast channel (IPC)
 - Improves scale and performance
- RSPs maintain single consolidated system view

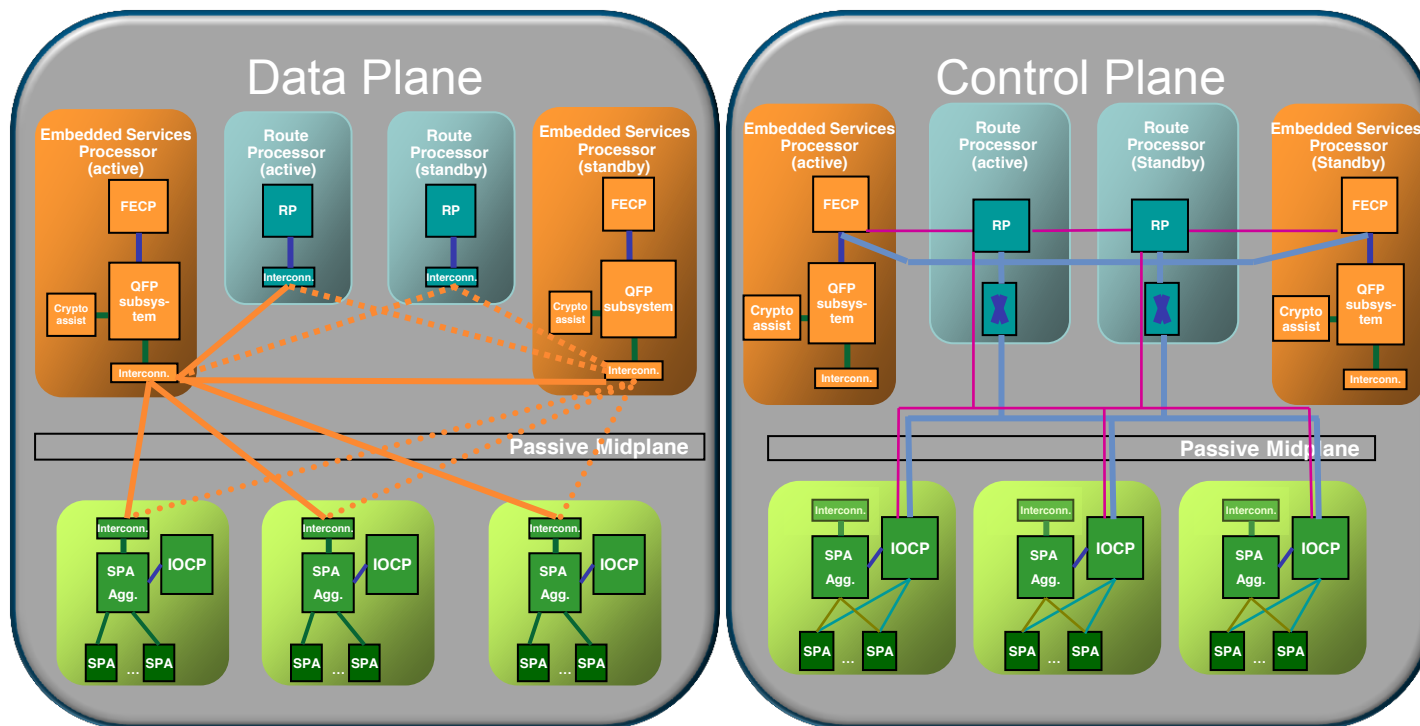


ASR 9000 High availability Infrastructure



- **Distribution improves fault tolerance** and recovery time by localizing the database and system management functionality to each node
- **Granular process restart** allows for fast recovery from failures
Hitless process re-starts!
- IOS XR is designed to **optimize the switch over between redundant hardware** elements
IOS XR is designed to route around fabric failure
Line cards are protected by link bundling, APS, IPS, ECMP etc.
RSP failover has no impact on the control protocols or packet forwarding running on the linecards

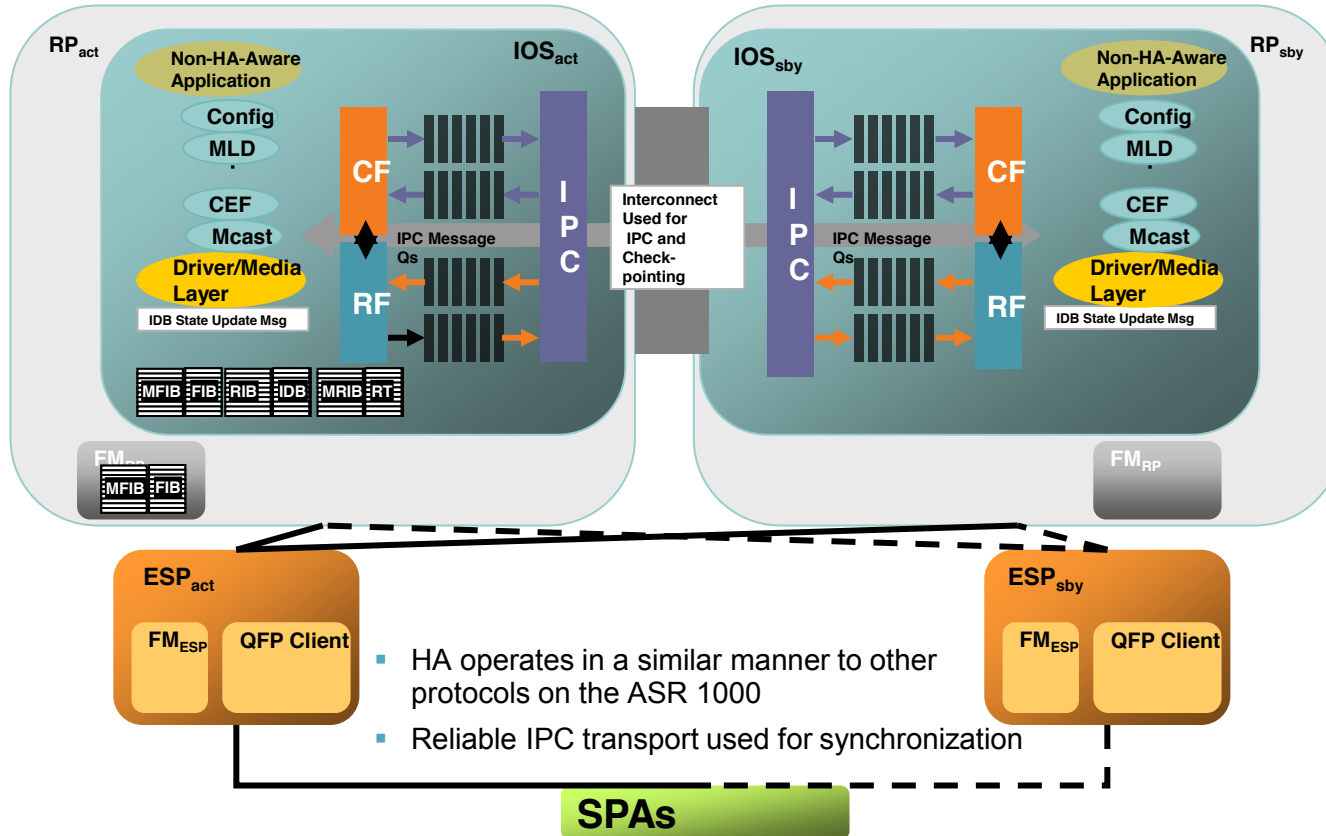
Example: ASR 1000 System Redundancy



— ESI, (Enhanced Serdes) 11.5Gbps
— SPA-SPI, 11.2Gbps
— Hypertransport, 10Gbps

— GE, 1Gbps
— IFC
— SPA Control
— SPA Bus

ASR 1006 / 1013 System High Availability

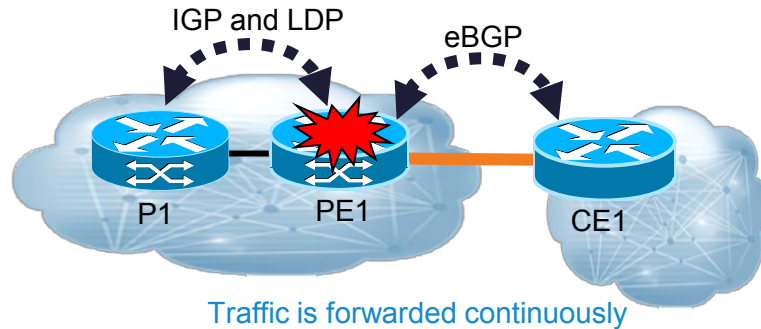


Cisco Software High-Availability Support

- Stateful Switchover (SSO) support for features provides the synchronization of dynamic feature state between hardware modules
- Configuration synchronization ensures that the running config is synchronized on the route processors

Dynamic State Preservation	ASR 1000	ASR 9000
Connectivity Protocols & Infrastructure	FR, ATM, PPP, MLPPP, HDLC, 802.1Q, BFD (BGP, IS-IS, OSPF), CEF, SNMP,	BFD (OSPF, BGP, IS-IS, Static)
Routing & IP Services	IS-IS, ARP, HSRP/GLBP, QoS, OSPF, DHCP, VRRP NSR: BGP, OSPF, MP-iBGP (roadmap for IS-ISv6) NSF/GR: IPv6, OSPFv3, IS-IS, RIPng, PIMv6, MLD, (roadmap for IS-ISv6)	NSF (ISIS, OSPF, BGP), NSR (ISIS, OSPFv2, OSPFv3, BGP)
Multicast	MFIB, IGMP, PIM-SM, PIM-SSM, PIM-BiDir, MLD, MSDP, Auto RP, BSR, mVPN, mLDP GR, NSR for mLDP	NSF Multicast, BFD for PIM, MoFRR
MPLS Protocols	L3VPN, LDP, VRF-aware BFD, QoS, TE, EoMPLS, L2VPN, L2TPv3, LDP, t-MDP, NSR MPLS TE (with OSPF, IS-IS, RSVP)	NSF (LDP, T-LDP, RSVP-TE) NSR (LDP, T-LDP), BFD for MPLS FRR, VRF-aware BFD
Broadband	PPPoE, L2TP (LAC, LNS), DHCPv4/v6, AAA, session state (virtual templates), ISG/IP Sessions, ANCP, LI	PPPoE (including nV)
Security	SSO, Stateful Inter-chassis redundancy for FW / NAT	Roadmap
SBC	SSO	n.a.

BP: Use Non Stop Routing (NSR) for PE-CE Resilience



- Routing sessions are maintained between processors on a failure, allowing routing sessions to stay up with Peer
No peer session flaps
- Copy of FIB maintained on secondary and used on failure for continuously traffic flow
- No need for neighboring routers to be NSF aware or capable. Can give high reliability without upgrading CE.
- Configuration

```
router bgp <asn>  
neighbor x.x.x.x ha-mode sso
```


```
router ospf <process id> [vrf <vrf name>]  
Nsr
```



BP: Stateful Inter-chassis Redundancy

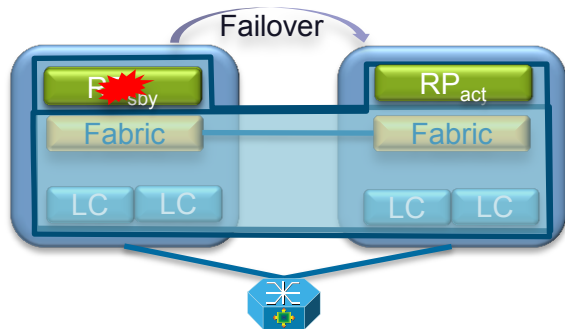
Stateful Inter-chassis Redundancy - Motivation

- Current Intra-chassis HA typically **protects against**
 - Control Plane (RP) Failures
 - Forwarding Plane (ESP) failuresInterface failures can be mitigated using link bundling (e.g. GEC)
- Any other failures may result in recovery times O(hours)
- Inter-chassis redundancy provides **additional resilience against**
 - Interface Failures
 - System failures
 - Site failures (allowing for geographic redundancy)



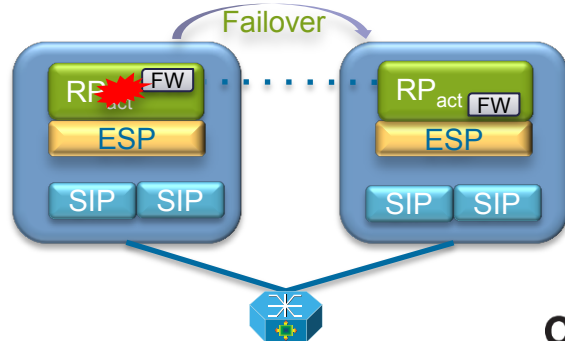
System Level Redundancy

- Example: VSS, nV
- Failover Granularity at the **System Level**
- **Control-plane active-standby**
Active RP considers 'remote' linecards under its control
- **Forwarding-plane active-active**
- No application granularity for failover
Need to ensure all features are SSO capable



Application Level Redundancy

- Example: RG Infra
- Failover Granularity at the **Application Level** (NAT, Firewall, SBC etc)
- **Control plane active-active**
Each RP only considers its own linecards, but synchronizes application state
- **Forwarding-plane active-active**
- E.g. can have one set of firewall services resilient, and other set of firewall services non-resilient





Stateful System Redundancy Models

- Different deployment models

 - 1+1 – one system is **actively** processing and passing traffic, the other in **standby** mode.

 - 1:1 – **two systems** are **actively** processing and passing traffic, and backing each other up

 - N+1 – **N systems** are **actively** processing and passing traffic, and share a single standby

- System vs. Application

Is the inter-chassis resilience applicable to ALL of the features / functions configured on the system, or only for a particular application?

System-level: provide **resilience for ALL applications** and traffic configured on a system

Application-Level: provide **resilience for a particular application** and its traffic

- Hot-standby vs. Cold-standby

Cold-standby: **FIB / adjacency updates are NOT synchronized** between active and standby system

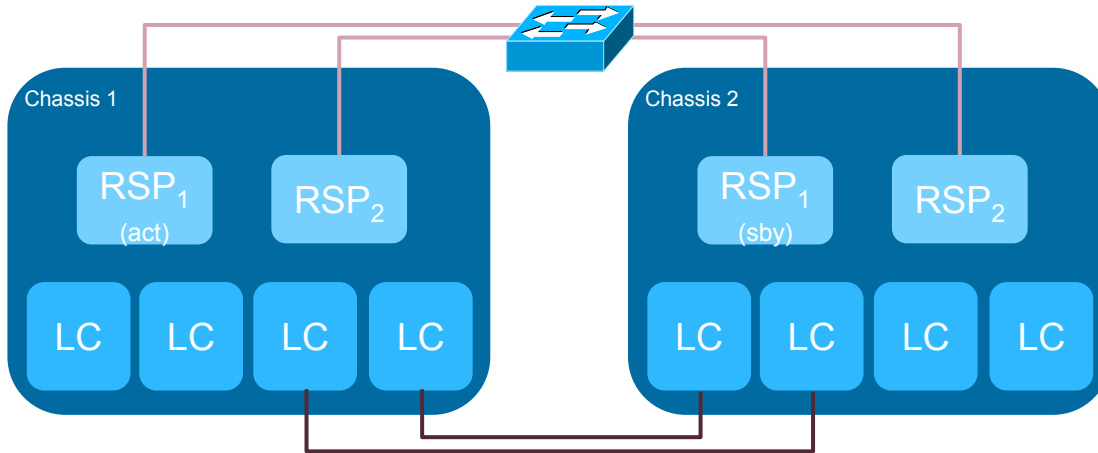
Hot-standby: **forwarding/state information is synchronized** between active and standby system

- Different Approaches can also be categorized into

 - Control plane active-standby / active-active

 - Forwarding plane active-standby / active-active


BP: Deploy ASR 9000 nV Edge




- **Chassis Members are linked via L1/L2 control plane links** (10ms latency assumed)
Dedicated (special) ports on RSP
Extend internal EOBC between the chassis members
1GE or 10GE using Ethernet SNAP encapsulation on internal MAC addresses
- **Data plane fabric is extended between the chassis via regular Ethernet VLAN line-card ports**
Can mix ports on a linecard for data plane extension and downstream connectivity
No requirement for external fabric chassis
10G or 100GE up to 32 ports
- **Both Control plane and Data plane extensions are assumed to be redundant**

Achieving Redundancy via Clustering

Redundancies required in current architecture: New protocols, extra routing peers



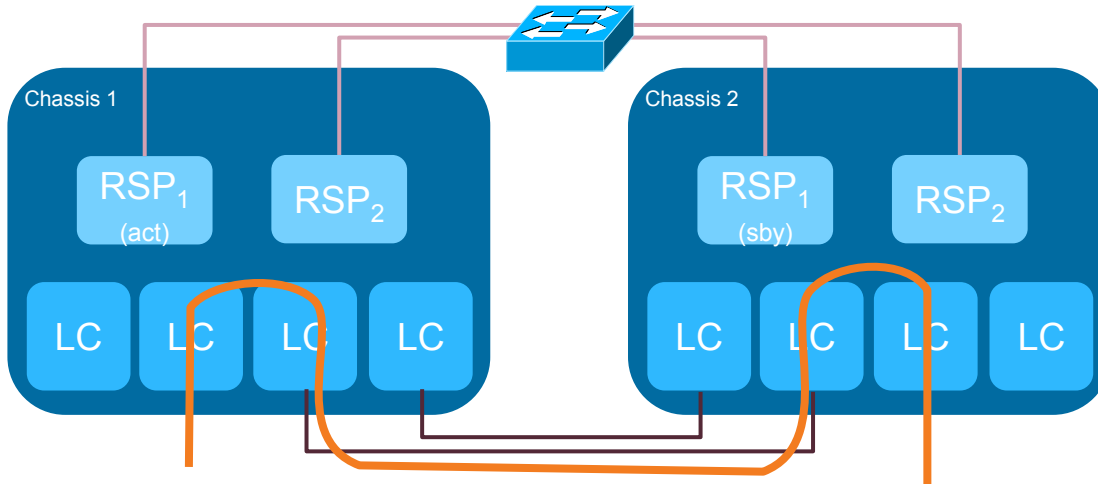
Clustering simplifies UNI & NNI-facing links, end-to-end network connectivity



Key ASR 9000 nV Technology Advantages

- ✓ Plain straightforward network protocol configuration and design greatly simplified network protocol redundancy design
no more mandatory complex HSRP/VRRP, BGP edge PIC, PWE-RED, MST/REP-AG, RG/MC-LAG etc
more efficient and full network link bandwidth resource usage for all active active operation, no more active-standby links
- ✓ One logical system internal infrastructure convergence and service level state synchronization by control plane
Maintain service level state persistence and eliminate any need for service rebuilding up upon network layer convergence post any link/node failure
Enable service scale independent convergence
Such as BNG subscriber sessions (PPPoE, RADIUS, DHCP binding table)
IGMP group membership, ARP table, doesn't even need to run MoFRR

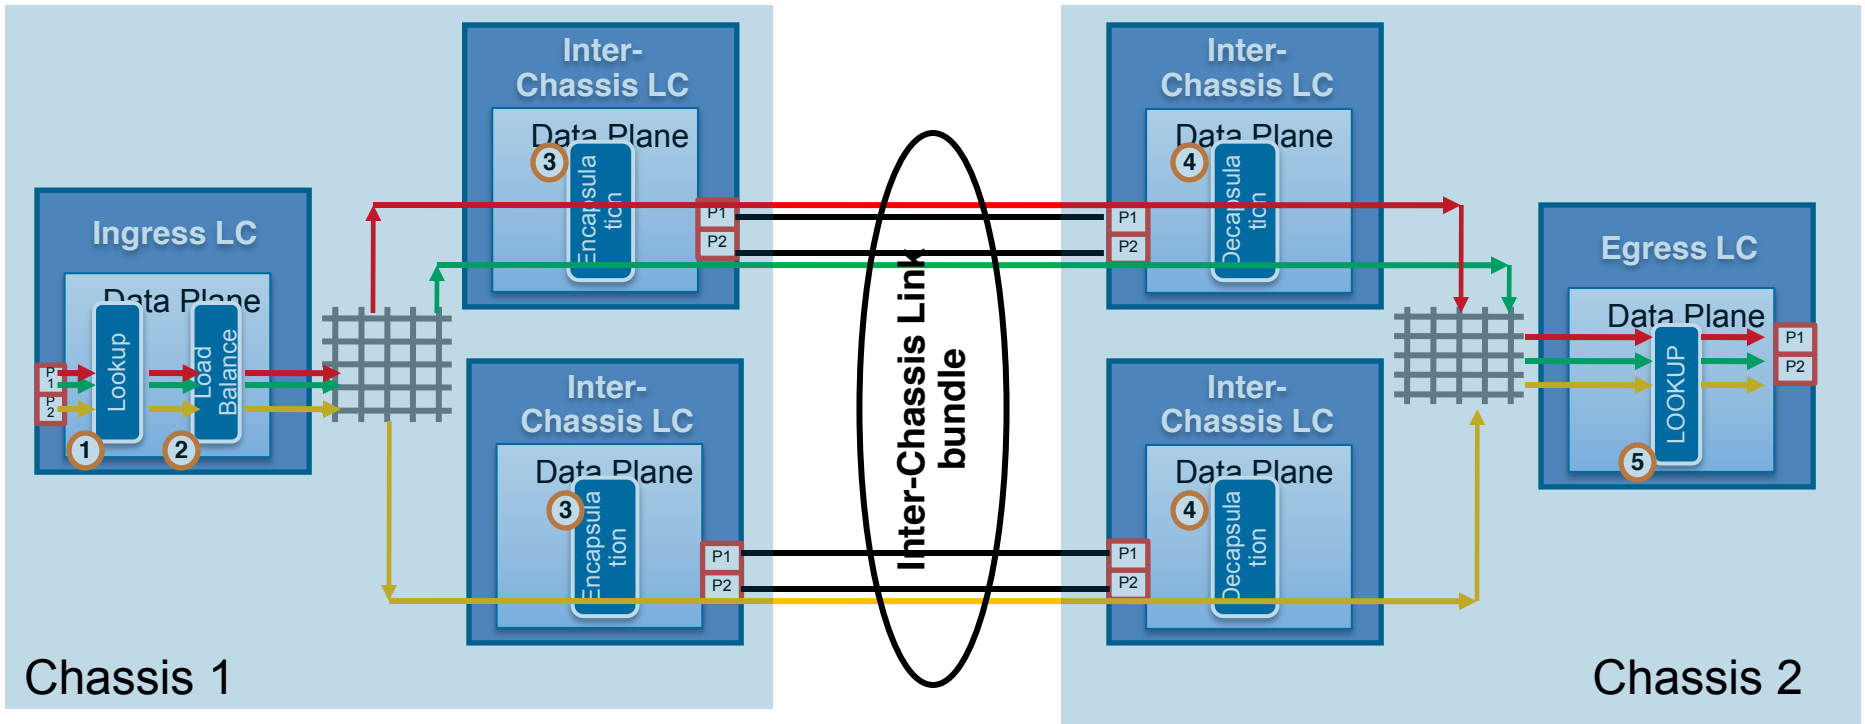
ASR 9000 nV Edge Forwarding



- Keep the existing IOS-XR two-stage forwarding model
- Inter-chassis data links simulate the switch fabric
 - has similar features as switch fabric, for example, fabric qos.
 - Packet load balancing over inter-chassis links is same as regular link bundle: per-flow based
- In case of ECMP or link bundle paths cross two chassis, **local port preferred over load balancing packet** to the other chassis.
 - Can be turn off by user CLI
- Only **single Multicast copy is sent** over inter-chassis link. Multicast replication is done on egress line cards and fabric on the local chassis



nV Edge Data Forwarding



Chassis 1

Chassis 2

1 Ingress Forwarding Lookup → L2/L3/Mcast regular lookup

2 Inter-Chassis Load Balance → Load balance across multiple inter-chassis links

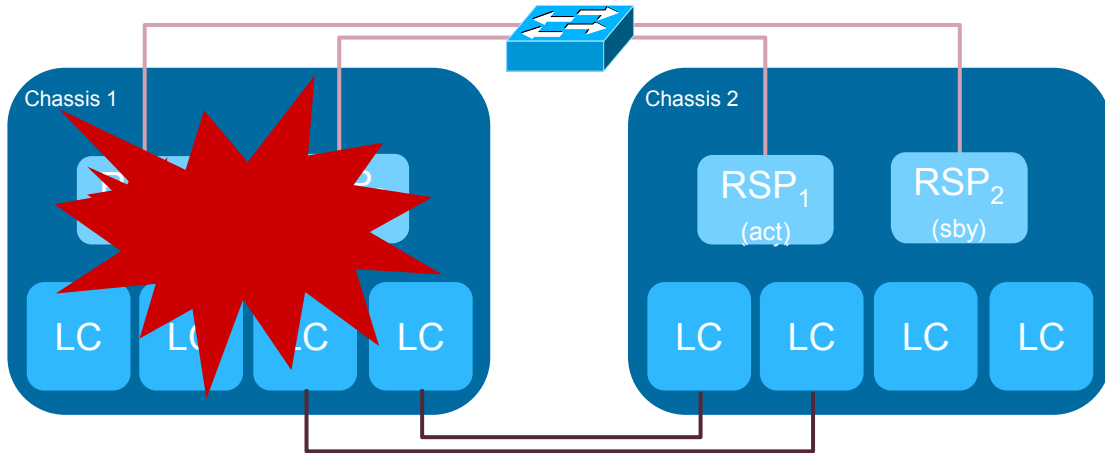
3 Inter-Chassis Encapsulation

4 Inter-Chassis Decapsulation

5 Egress Forwarding Lookup → L2/L3/Mcast regular lookup



ASR 9000 nV Edge Failures



- Only one Active RSP, Only one standby RSP at a given time, which are located on two different chassis

SSO/NSF/NSR works exactly the same way as two RSPs on the same chassis

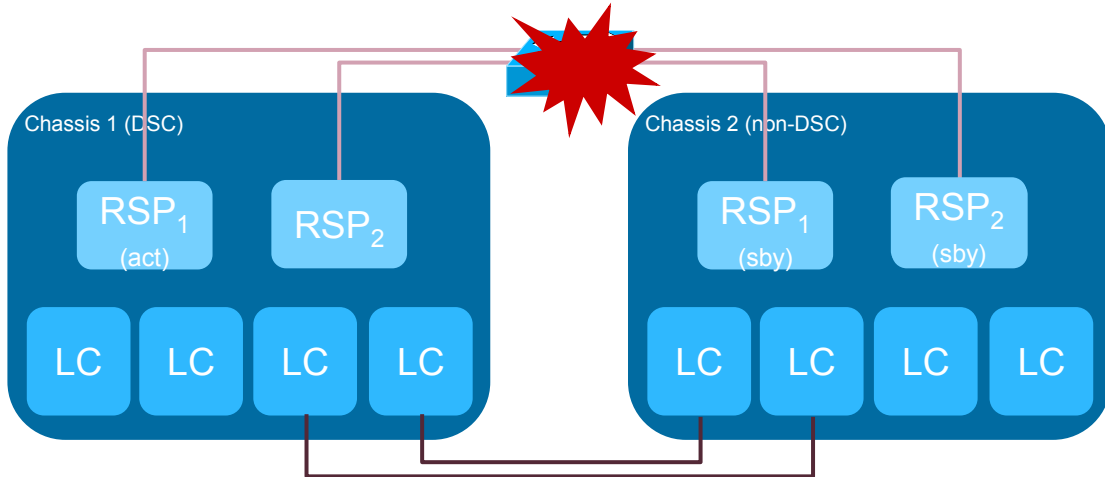
Reliable out of band control channel between two chassis

IOS-XR control plan can tolerate hundreds of msec latency*, although the latency can impact overall service convergence time

- Virtual Chassis is always on as long as there is one chassis and one RSP alive

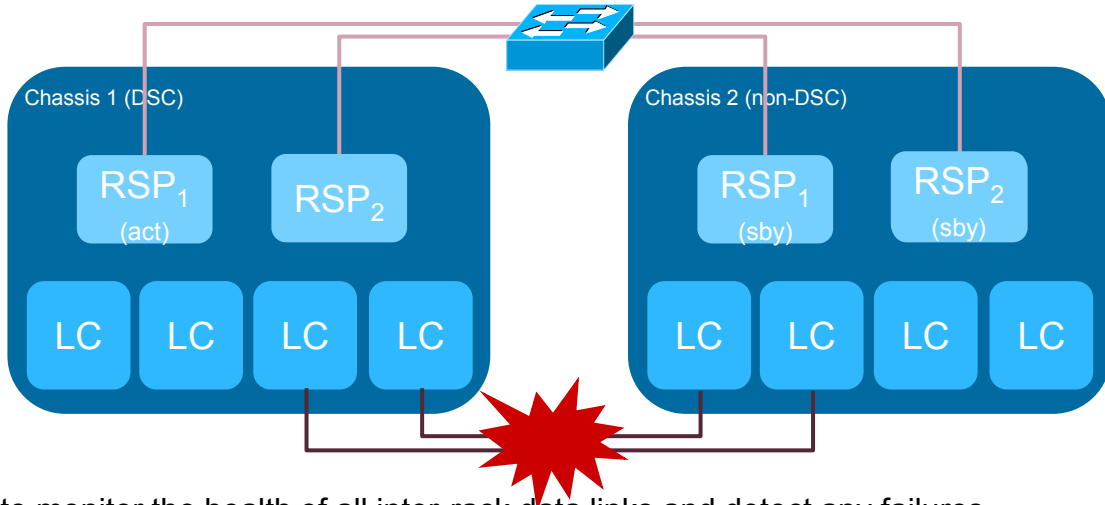
CiscoLive!

ASR 9000 nV Control Plane Failures



- RSPs in the cluster multicast **fast heartbeat hello messages** to other RSPs and EOBC control plane loss can be detected in a matter of ~1s
- In case of all **EOBC connection loss**, DSC rack continues to work as usual
- **Backup DSC queries DSC** rack via data plane extension on detection of complete CP connection loss
 - If dsc rack is alive by data plane query, non DSC will bring its own rack down to avoid any potential network mis-or duplicate identity and routing peers or service disruption
 - If unattended, the downed non-DSC rack1 will periodically boot up to try to rejoin the cluster
 - if CP is restored, rack1 will boot up to rejoin the cluster and service will be restored to use rack1 data path

ASR 9000 nV Data Plane Failures



- **UDLD** is used to monitor the health of all inter-rack data links and detect any failures
Inter-rack data link failure can be detected in ~200ms
- In case all inter-rack data links fail, by default, **nV edge system** will go into a **self-protected** operation mode:
DSC rack continues to work as usual, no change in its operation
Non-dsc rack brings down all its data ports to preempt possible unpredictable traffic blackholing
Non-dsc rack keeps probing its inter-rack link. Once they become available, bring up its service data ports
- **Config options** to allow both nv edge racks to continue **work as-is** without inter-rack data link or only bring down specific data ports in the cluster

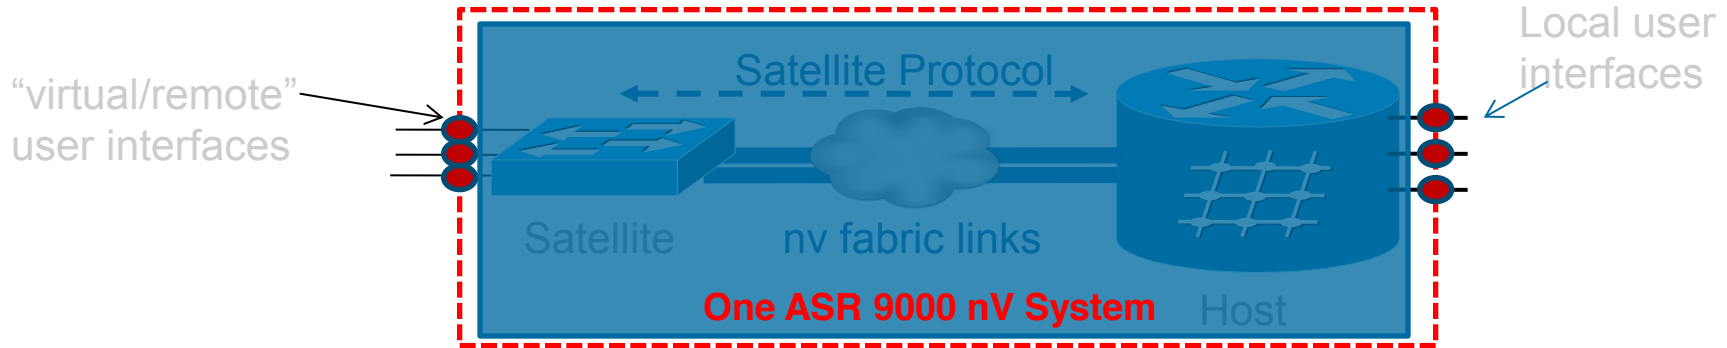
nV Edge System Rack Locality

- Bundle **local member links or ECMP paths** on a local rack are **preferred** over any cross rack member links or ECMP paths
- Link bundle rack locality is supported in 4.2.1
locality is per link bundle, applies to only flow based loadbalance
configurable under bundle interface

```
Interface bundle-ether 100
    bundle load-balancing localize threshold links < n >
```

If the number of active member links on local rack > n, enable locality for the bundle
- Rack locality **supports unicast paths only**
L2 multicast/flood traffic continue to use all member links in both racks
The same is for L3 multicast
The same applies to ICL bundle to/from ASR9000v
- Rack locality is **off by default**

BP: Complement nV Edge with nV Satellite



- Satellite and ASR 9000 host run satellite auto-discovery and control protocol (SDAC) for host based provisioning and management
- Satellite and Host could co-locate or in different location. There is **no distance limit** between satellite and Host
- Satellite uplink to host forms “nv fabric link”, which could be L1 or over L2 virtual circuit
- ASR9K Host and its associated satellites are one virtual Router system, running one OS: IOS XR
- From end user point of view, Satellite functions like a “remote or virtual” line card to host ASR9k. The interfaces on satellite look/feel/work the same as any interfaces on any local ASR9K line cards


Satellite is plug-n-play, zero touch configuration/management

Cisco *live!*


nV Satellite Network Topology




Single home, static pinning



Single home, fabric link bundle



Dual home to cluster, static pinning




Dual home to cluster with fabric link bundle

BP: Use RG-Infra on ASR 1000 for NAT/FW/SBC

- RG Infra: IOS Redundancy Group Infrastructure
- Enables **synchronization of application state** data between different physical systems
 - Does the job of RF/CF between chassis
- Infrastructure provides the functions to
 - Pair two instances of RG** configured on different chassis for application redundancy purposes
 - Determine active/standby** state of each RG instance
 - Exchange application state data** (e.g. for NAT/Firewall)
 - Detect failures** in the local system
 - Initiate & **manage failover** (based on RG priorities, allows for pre-emption)
- Assumptions
 - Application state has to be supported** by RG infra (ASR 1000 currently supports NAT, Firewall, SBC)
 - Connectivity redundancy solved at the architectural level
 - need to 'externalize' the redundant ESI links of the intra-chassis redundancy solution

Redundancy Groups Functions

- Registers applications as clients
- Registers (sub)interfaces / {SA/DA}-tuplets in case of firewall
- Communicates control information between RGs using a redundancy group protocol
 - Advertisement of RGs and RG state
 - Determination of peer IP address
 - Determination of presence of active RG
- Synchronizes application state data using a transport protocol
- Manages Failovers!





Redundancy Groups Functions — Details

- **Configuration of stateful system redundancy**
 - Priority (similar to HSRP priority for RG state determination)
 - Preemption, Name
- **RG State control**
 - Init, Active, Standby, disabled
 - Communicating state changes to other software entities in the system (e.g. QFP software)
- **Synchronization management**
 - Synchronization state tracking (standby has to request bulk-updates from active)
 - Determines when synchronization is started (e.g. ensures transport is available)
- **Peer Management**
 - Maintain information about peers
- **Fault Handling**
 - Changing priorities of RG (may affect RG state)
 - Fault event dampening
 - Logging
 - Integration with Enhanced Object tracking / BFD
- **Transport Connectivity**
 - Knows via which interface application state is synchronized
 - Can be different for application state data and RG control messages



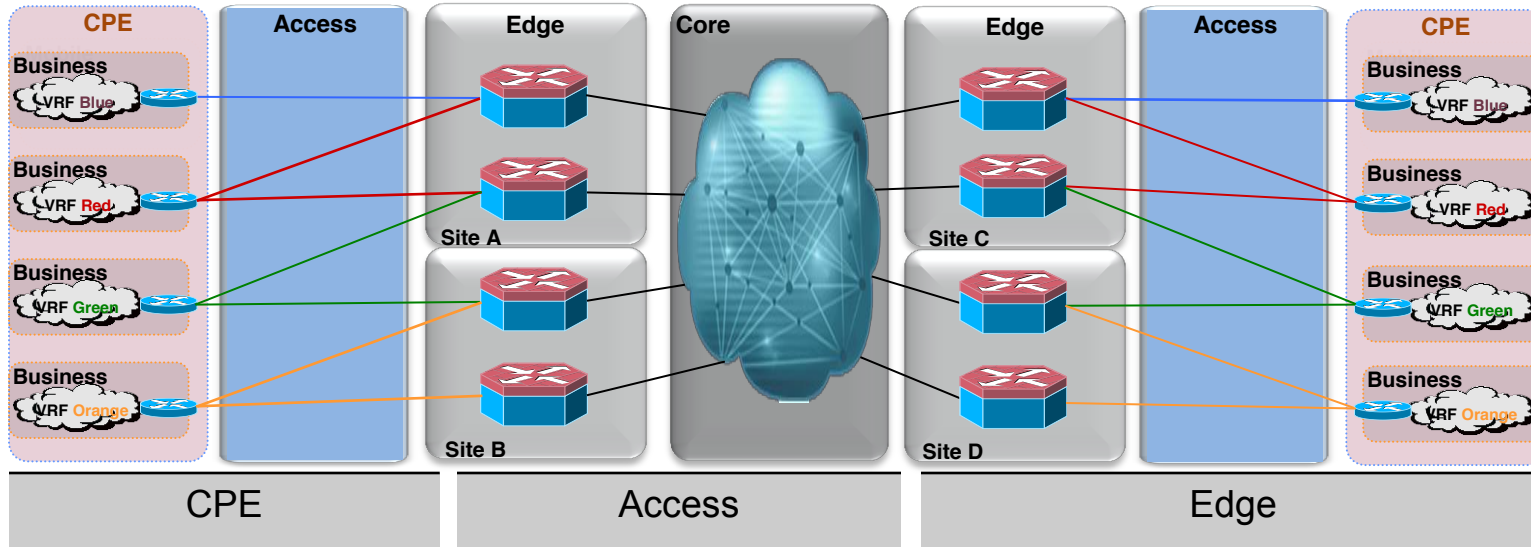
Service High Availability

Cisco *live!*

High Availability for Advanced Service Models

- Many SP Services already go beyond standard L3VPN / L2VPN / transport services
- Increasing subscriber management capabilities and L4-L7 services
- Examples:
 - Subscriber Management
 - Multicast
 - Session Border Controller
 - Firewall
 - IPSec
 - LI
- Some Services can be made highly-available using Intra-chassis redundancy (e.g. IPSec, Firewall, NAT, PPPoX, L2TP)
- Stateless inter-chassis redundancy available for BNG
- Stateful Inter-chassis redundancy available for NAT, Firewall and SBC on the Cisco ASR 1000

L3VPN Key HA Technologies



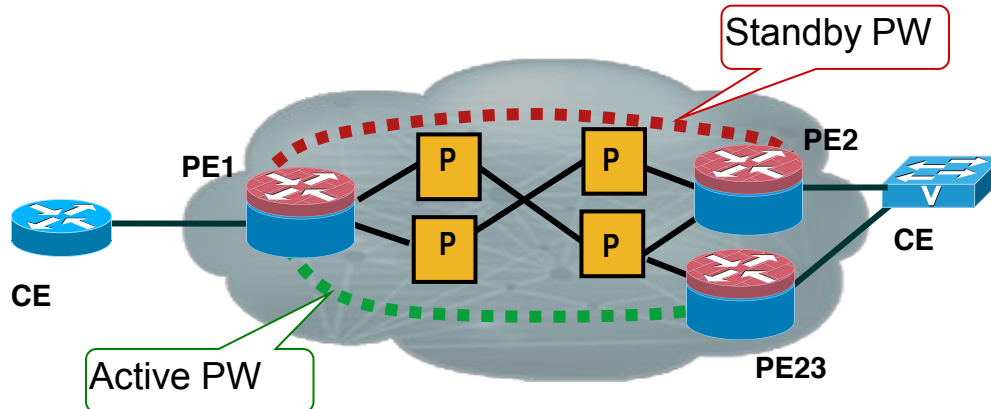
- BFD for PE-CE Link Detection
- NSF/NSR for Chassis HA
- PE Multihoming
 - Intra-Site PE for PE Diversity
 - Inter-Site for SP Facility Diversity

- Circuit Diversity - Physical Diversity for Multihomed CPE
 - Physical Circuit Diversity is Not the Default
 - Must be Requested from the SP

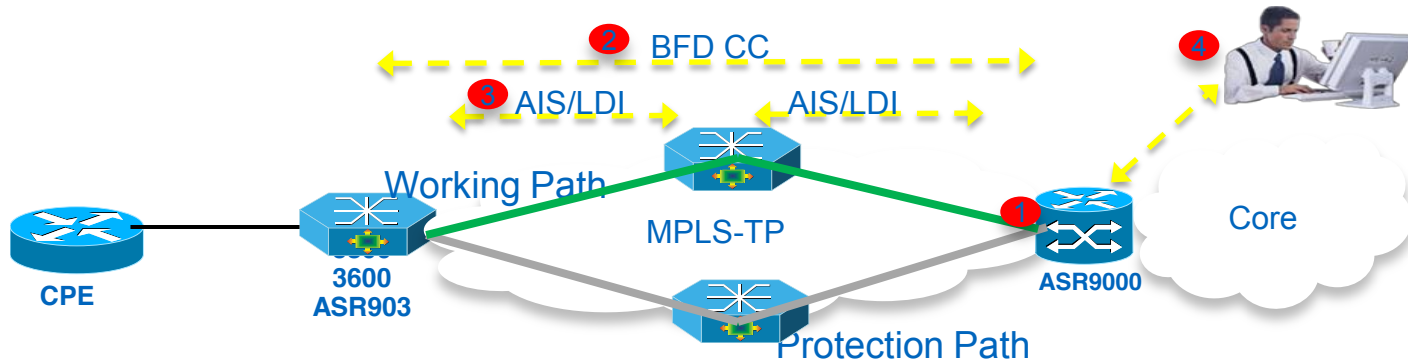
- BFD for PE-CPE / PE-P Link Detection
- NSF/NSR for Chassis HA
- IP Event Dampening for PE-CPE
- IP-FRR for PE-P
 - For Cost Effective PE-P Bandwidth
- BGP PIC Core
- BGP PIC Edge for Multi-Homed CPE

L2VPN — Pseudowire Redundancy

- Active-Standby PW Access Circuit Redundancy
L2TPv3 and MPLS Support
- Detection Mechanisms
 - IGP Convergence for Remote PE Failure
 - LDP Signaling for PE-CE Failure
 - LDP Timeout for Remote PE Software Failure

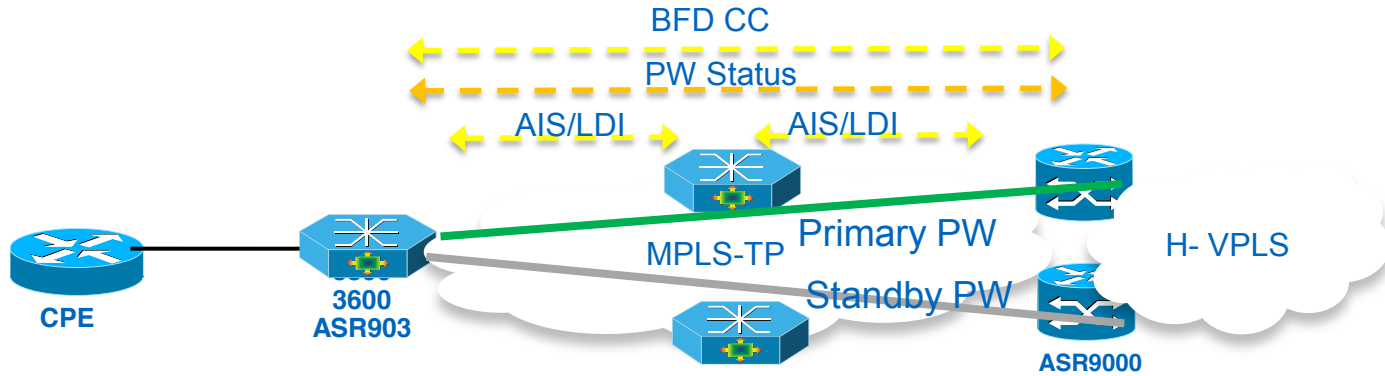


MPLS-TP Infrastructure Resiliency



- Statically setup Working and Protection LSP, 1:1 protection
- Protection switching can be triggered by
 1. LOS, Physical failure
 2. BFD timeout
 3. Detected defect condition (LDI/AIS, LKR)
 4. Administrative action (lockout), Far end request (lockout)
- Revertive mode, after wait to restore the traffic is restored over working path

MPLS-TP Service Resiliency



- Service Resiliency is based on PW redundancy
- Works for P2P and H-VPLS
 - For static PW, MAC-WD will be triggered based on PW status going down on primary PW
- Required for node protection or AC protection only. Link protection provided by MPLS-TP LSP

Multicast High-Availability Behavior

- Before failure

Multicast state is synchronized from RP_{act} to RP_{sby}

Configuration

MLDv1/v2 state information

PIM or MRIB state are NOT synchronized

MFIB also synched to ESP_{act} and ESP_{sby}

- After failure

RP_{sby} sends out PIM hellos to all neighbors

PIM neighbors re-send PIM state


Newly active RP re-builds the PIM state

IGP reconverges to assure uRPF check

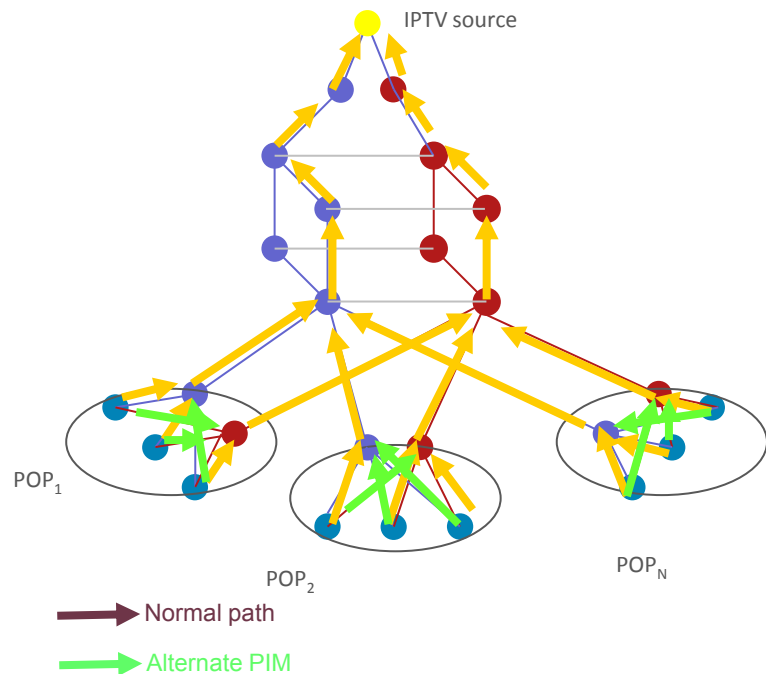
MFIB and ESP updates proceed to incorporate refreshed PIM state

- ESP_{act} continues to forward multicast traffic based on its version of the MFIB

Forwarding of multicast packets is NOT disrupted



Multicast only Fast Re-Route (MoFRR)



- Receiver
 - Multicast join on primary path
 - Multicast join on backup path
- Data packets are received from the primary and secondary paths
- The redundant packets are discarded at topology merge points due to RPF checks
- Failure:
 - Interface change on where packets are accepted
 - Backup path interfaces become 'active'


Configuration and Restrictions

Dependency on ECMP and will not work without it
Disabled by default and enabled through a cli
Applicable to IPv4 multicast only and not IPv6 multicast
Works only for SM S,G and SSM routes
Works where the rpf lookups are done in a single vrf
Extranet routes are not supported
Both primary and secondary paths should exist in the same multicast topology.

Stateful Application Switchover: PPP

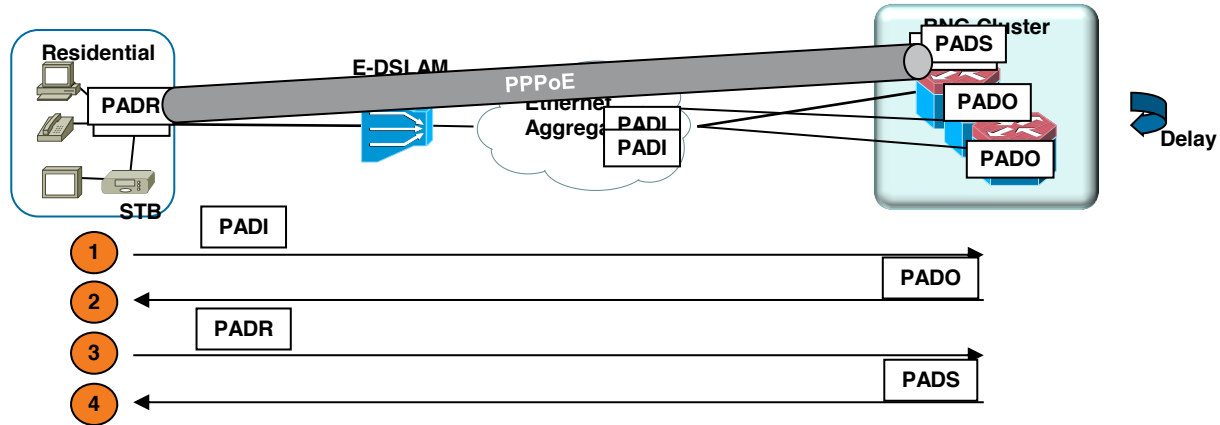
- Copies state information for PPP, PPPoE, and PPPoEoVLAN Sessions
- Switch-over is transparent to peers
 - Sessions are not torn-down / re-established
- PPP, PPPoE, and PPPoEoVLAN Session States:
 - Configuration (through config synch), including QoS configuration, ACLs
 - Session identifiers
 - PADR frame (cached)
 - RADIUS session attributes
 - Physical interface
 - VAI identifier
 - MD5 signature
- Statistics are synchronized on ASR 1000!

Stateful Application Switchover: L2TP



- RP_{act} synchronizes state with RP_{sby}
State includes configuration, PPP session IDs, L2TP CC sequence numbers etc.
Sequence numbers (N_s , N_r) for L2TP Control Connections (CC) are only synched once for a packet window of X (i.e. once every X L2TP control packets)

BNG Service Edge High Availability

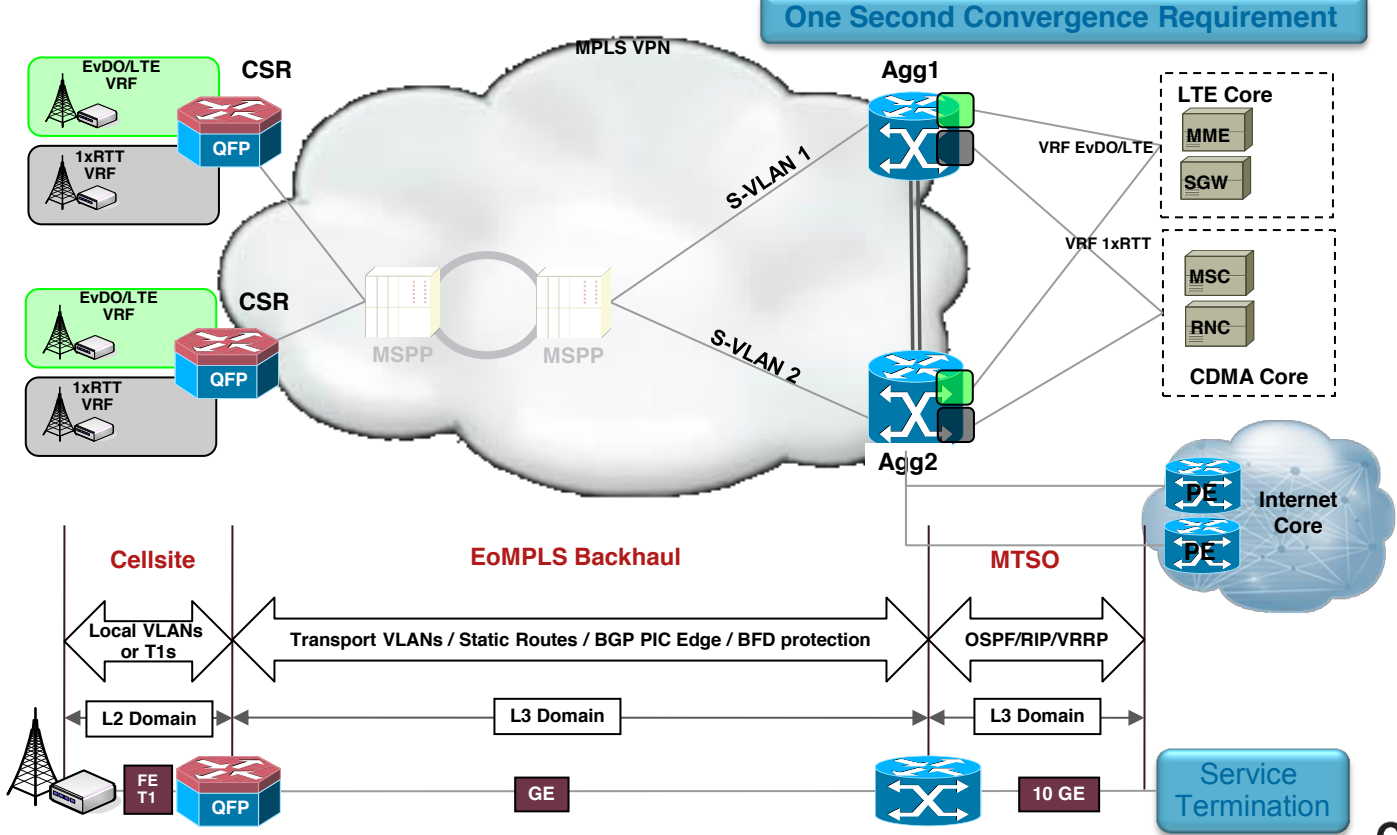


- PPP Smart Server Selection allows user to configure specific PADO delay for a received PADI packet
 - Can be configured per bba-group or based on circuit-id/remote-id
- In case of an outage of a BNG in the cluster, other BNG stand ready to accept subscriber sessions
 - Detection of failure possible at both ends of PPPoE session because of missing keepalives
 - Subscriber sessions have to be re-established
- Allows BNG redundancy with predictable behavior



Case Studies

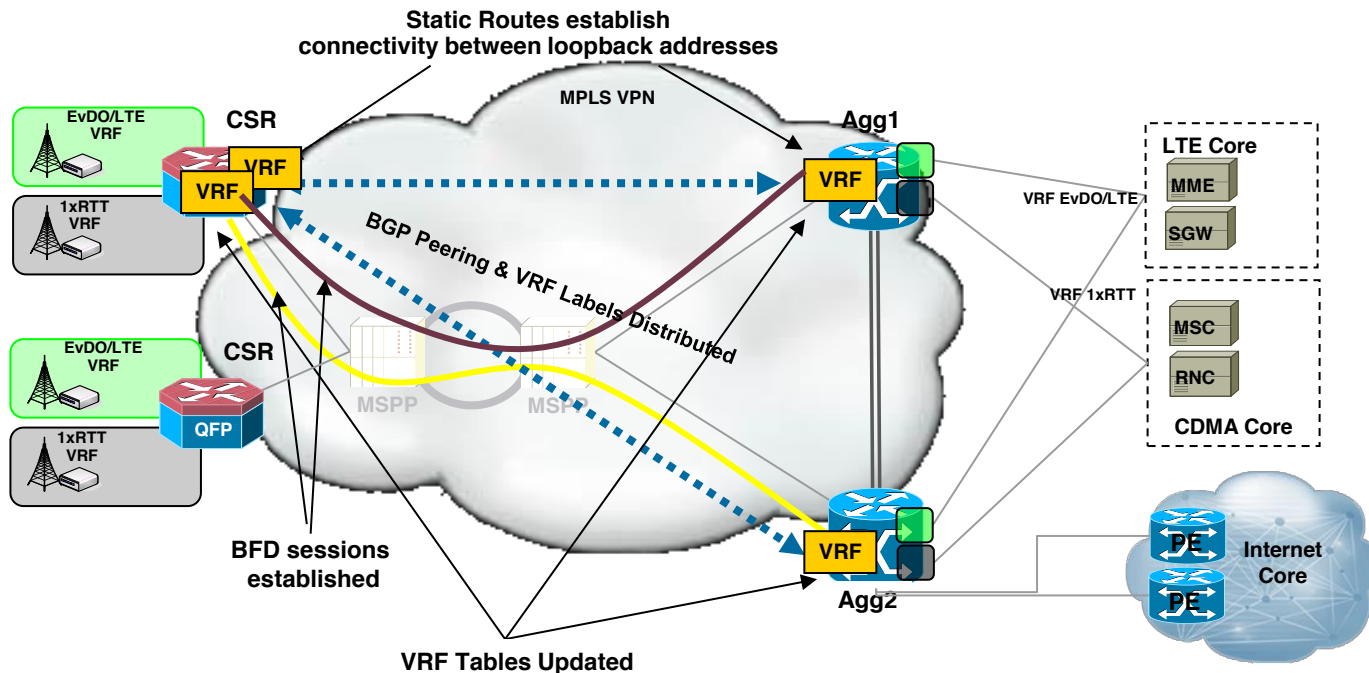
Case Study: Highly Available IP Architecture for Mobile



One Second Convergence Requirement




Case Study: Highly Available IP Architecture for Mobile — Transport



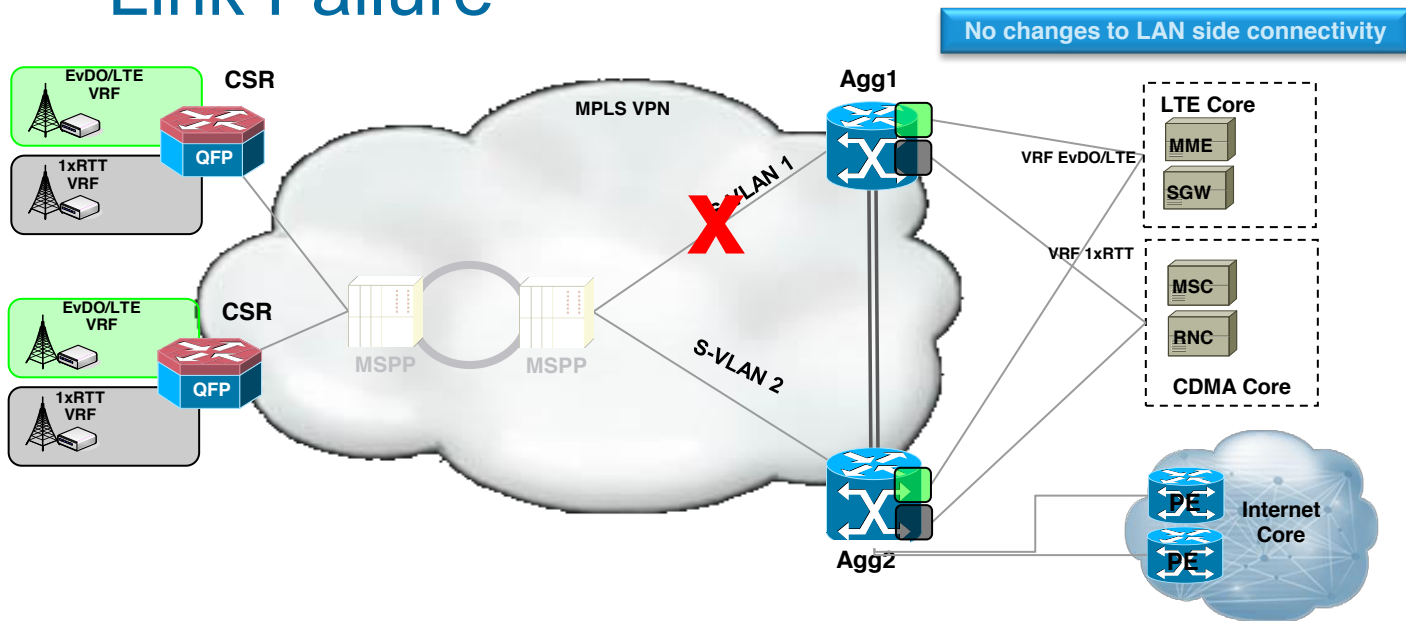
- Static routes for cellsite reachability
- BGP PIC Edge for Layer-3 convergence
- VRRP for MTSO

Case Study: Highly Available IP Architecture for Mobile — Steady-State Traffic Flows



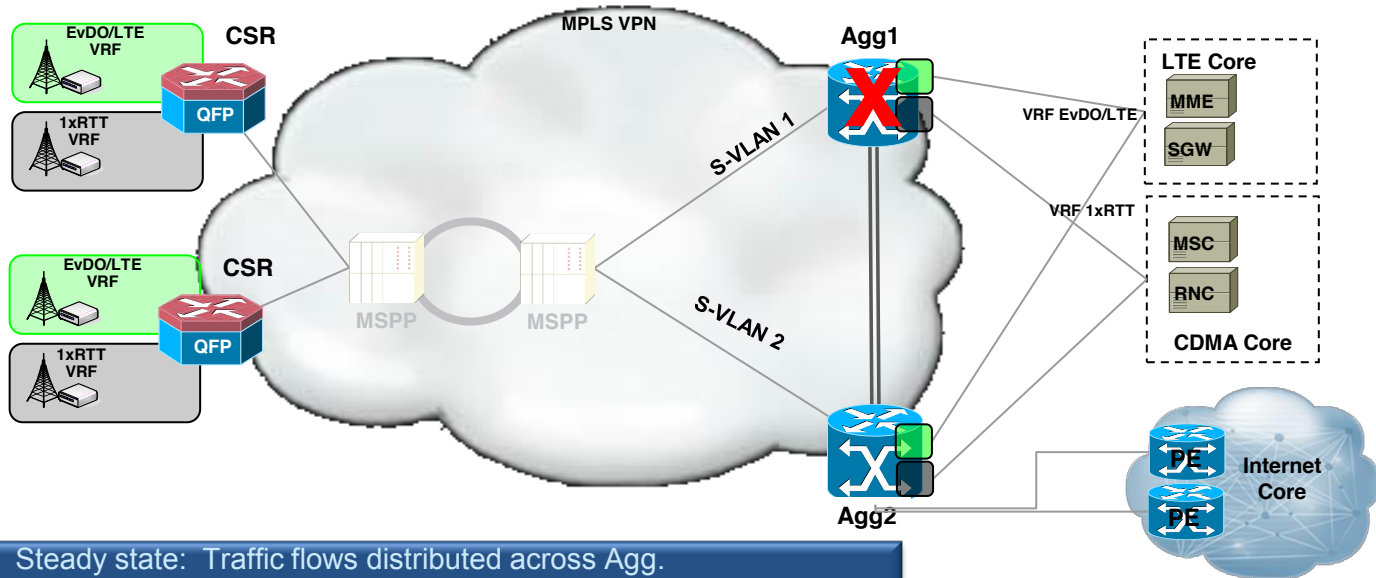
- Steady state:
 - CSR distributes flows across both Agg's using ECMP.
 - Traffic could flow across Agg inter switch links.
 - Each Agg handles traffic related to all services from the cell-site.

Case Study: Highly Available IP Architecture for Mobile — Link Failure




- Steady state: Traffic flows distributed across both Agg.
- Failure: GE link from MSPP to Agg1 fails.
- Action:
 - BFD session to Agg1 times out at CSR.
 - Agg1 next hop removed from forwarding table.
 - Traffic flows resume across existing path to Agg2.
- Results: Traffic flows to Agg1 via Agg2.

Case Study: Highly Available IP Architecture for Mobile — Aggregation Switch Failure



- Steady state: Traffic flows distributed across Agg.
- Failure: Agg1 power outage.
- Action:
 - BFD and VRRP sessions time out
 - BGP and OSPF neighbors drop due to BFD
 - BGP PIC Edge ensures sub-second convergence
 - Traffic flows resume across existing path thru Agg2.
- Results: Traffic flows via Agg2 to end hosts.

Case Study: Highly Available IP Architecture for Mobile — CSR Failure



- Steady state: Traffic flows distributed across CSR.
- Failure: CSR power outage.
- Action:
 - BFD sessions time out
 - BGP neighbors drop due to BFD
 - Mobile handsets resync to neighboring cell site
- Results: Mobile handset voice connectivity is maintained.



Summary

Summary

- Motivation for High Availability in SP Aggregation Networks
- Network Level High Availability
- System High Availability
- Stateful Inter-chassis Redundancy
- Service High Availability
- Case Studies
- Summary and Conclusions

Key Takeaways

- **High-Availability** becoming **increasingly deployed** in Aggregation Networks
Motivated by experiences with MPLS Core Networks
- Many high-availability techniques deployed in the core are now applied in MPLS aggregation networks
MPLS TE FRR, BFD, EOAM, Pseudowire Redundancy ...
- Service High Availability **requires comprehensive approach** including the deployment of
 - Network level resiliency
 - System Level resiliency
 - L4-7 service resiliency
- **Stateful Inter-chassis redundancy** increasingly being considered to provide geographic redundancy for applications
- **ASR 9000 nV Edge and Satellite greatly reduce complexity**
- High Availability **comes at a cost** (CAPEX & OPEX)!

Complete Your Online Session Evaluation

- Give us your feedback and you could win fabulous prizes. Winners announced daily.
- Complete your session evaluation through the Cisco Live mobile app or visit one of the interactive kiosks located throughout the convention center.



Don't forget: Cisco Live sessions will be available for viewing on-demand after the event at [CiscoLive.com/Online](https://www.ciscolive.com/online)

Continue Your Education

- Demos in the Cisco Campus
- Walk-in Self-Paced Labs
- Table Topics
- Meet the Engineer 1:1 meetings

Recommended Reading



- **N. Stringfield et. Al, “Cisco Express Forwarding”,**

ISBN-13: 978-1-58705-236-1

- **D. C. Lee, “Enhanced IP Services for Cisco Networks”,**

ISBN-13: 978-1-58770-106-3

- **A. Sayeed, M. Morrow, “MPLS and Next-Generation Networks”,**

ISBN-13: 978-1-58720-120-2

- **J. Davidson et. al, “Voice over IP Fundamentals”, 2nd Edition,**

ISBN-13: 978-1-58705-257-6

- **V. Bollapragada et. Al, “Inside Cisco IOS Software Architecture “,**

ISBN-13: 978-1-58770-181-0.

- **R. Wood, “Next-generation Network Services”,**

ISBN-13: 978-1-58705-159-3.

- **K. Lee, F. Lim, B. Ong, “Building Resilient IP Networks”,**

ISBN-13: 978-1-58705-215-6

- **T. Szigeti, C. Hattingh, “End-to-End QoS Network Design: Quality of Service in LANs, WANs, and VPNs;”,**


ISBN-13: 978-1-58705-176-0

- **B. J. Carroll, “Cisco Access Control Security”,**

ISBN-13: 978-1-58705-124-1.

- **A. Khan, “Building Service-Aware Networks: The Next-Generation WAN/MAN”,**

ISBN-13: 978-1-58705-788-5



Whitepapers on CCO

- Cisco IOS High Availability
 - 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 Validated Designs
 - http://www.cisco.com/en/US/netsol/ns742/networking_solutions_program_category_home.html
- ASR 9000
 - [Cisco ASR 9000 Series High Availability: Continuous Network Operations](#)
 - [Introduction to Cisco ASR 9000 Series Network Virtualization Technology](#)
 - [Distributed Virtual Data Center for Enterprise and Service Provider Cloud](#)
- ASR 1000
 - [Cisco ASR 1000 Series Aggregation Services Routers](#)
 - [Cisco ASR 1000 Series: ISSU Deployment Guide and Case Study](#)
 - [Cisco Unified Border Element \(SP Edition\) on Cisco ASR 1000 Series](#)
 - [Cisco Unified WAN Services: Services, Security, Resiliency, and Intelligence](#)



CISCO