

Cisco Expo  
2008


# Cisco 7600 Series Update

... introducing the  
12.2SRC Release



**Christian Schmutzer**  
**Technical Marketing Engineer, ERBU**  
**March 2008**


**Enable Your Network**  
**Empower Your Business**



# Agenda

- The role of the Cisco 7600 in NGN Networks
- Overview of new Hardware and Software Features in Cisco IOS 12.2(33)SRC
- Details of Key Software Features introduced in Cisco IOS 12.2(33)SRC
- Summary

# Cisco 7600's Role in NGN Networks



...

# Cisco IP/MPLS Portfolio

Cisco 7600



Triple Play  
Service Edge  
20G=>40G=>80G+

**NG Evolution**

Cisco 10000



Integrated  
Broadband and  
Leased Line  
Edge

**NG Evolution**

Cisco XR 12000



Multiservice  
Edge  
10G => 20G

**Evolution**



Cisco CRS-1



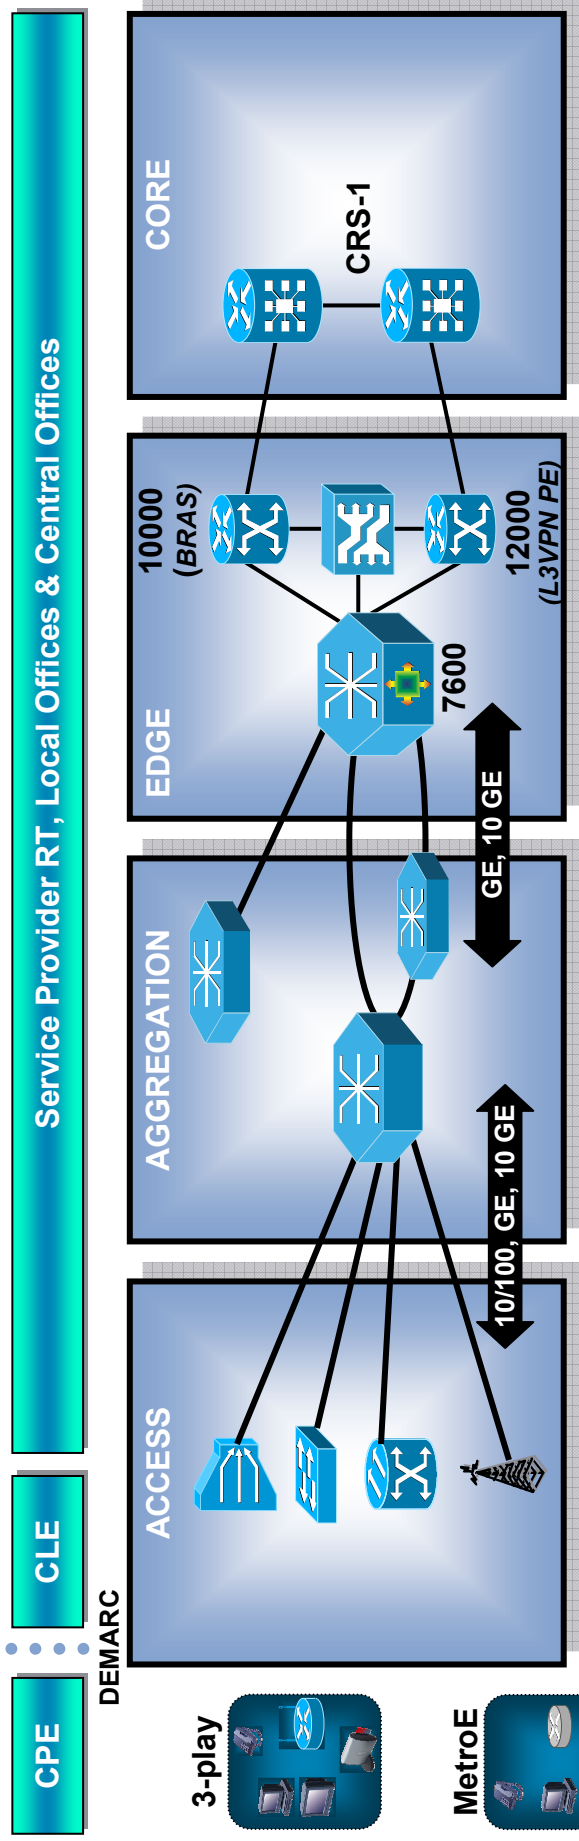
Redefining  
IP  
Communications  
40G => 100G+

**Evolution**

# Service Transformation



# Next Generation Aggregation Edge Areas of Focus



**20-400G**, 32kVPLS, 64kVPWS, VPWS, VPLS, **H-VPLS**, Mac-in-Mac, IP Multicast/Unicast

4-Level **H-QoS**, Highly Flexible Classification, TE  
>1M Queues, 80k 2R3C Policers, Ingress / Egress Policing

**Add/Remove/Rewrite 2-Tags** . 1q/.1ad/VC/MPLS Tunnel  
802.1 Legacy Integration, EoS

802.1ag, **802.3ah**, **E-LMI**, **VCCV**, **BFD**, **FRR**, **EFP**  
IOS-XR : ISSU, NSF, & Process Manageability

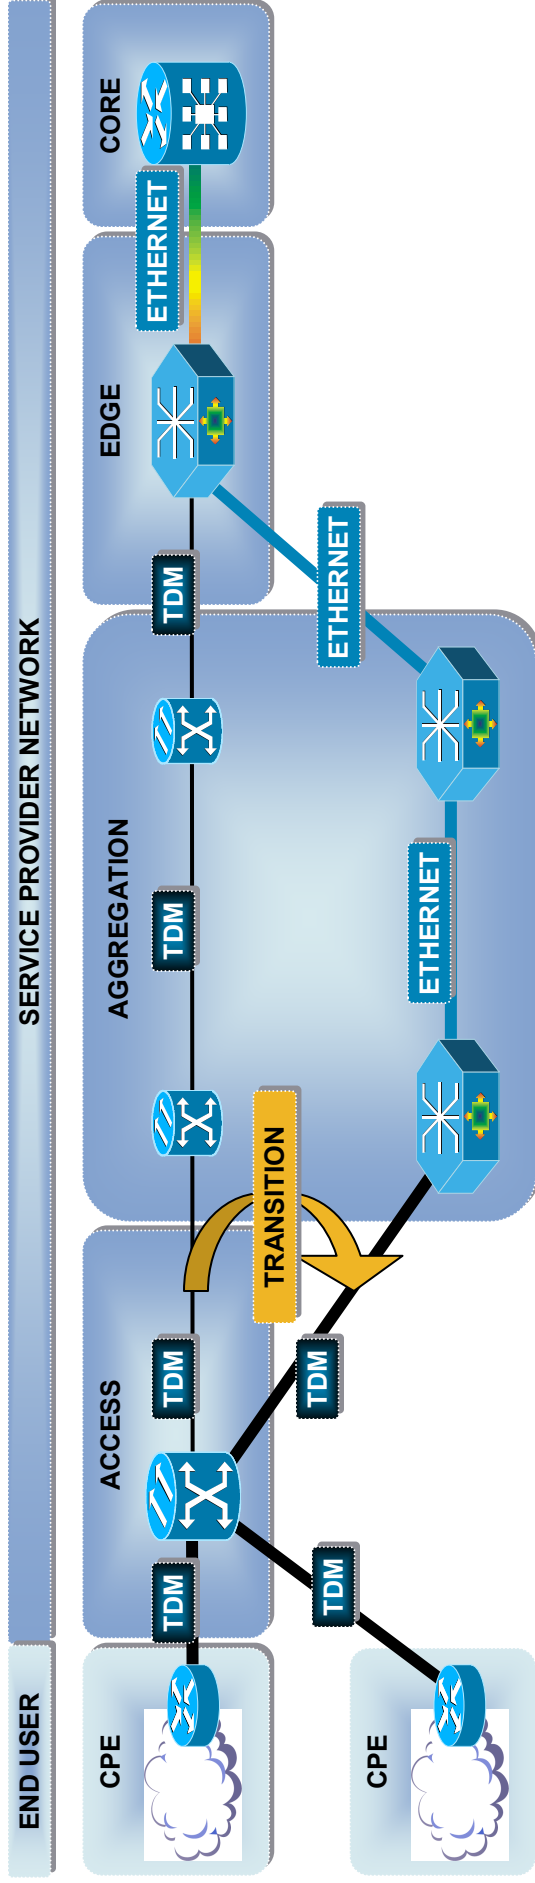
Service / Network Scalability

QoS Management

Interface / Service Flexibility

Carrier Class Manageability

# Legacy – Overlay – NGN (Transition)



## Legacy Interfaces

FR/ATM/IP over E1/STM-x

CEoP + Clocking

## Packet based NGN "Bridge"


Legacy Adaptation  
over Packet  
(PWE3/MPLS)

## NG Edge

Packet => Edge and  
Core

Legacy TDM => Legacy  
TDM (Mobile)

# Overview of new HW and SW Features in 12.2(33)SRC



...




# The Cisco 7600 Family




## Engines

Supervisor 32  
Supervisor 720  
Route Switch Processor 720  
Route Switch Processor 720-10GE




## SPA Interface Processors

Modular Carrier Cards  
for WAN and Metro  
Shared Port Adapters



## High-Density Ethernet Modules

High-Density GE and 10GE  
with Distributed, Line-rate Performance




## Enhanced FlexWAN

7500 Parity and  
PA Investment Protection



## Ethernet Services Modules

GE and 10GE  
with Rich QoS, Distributed,  
Line-rate Performance




## S Chassis

7609-S, 7606-S,  
7603-S, 7604

## Traditional Chassis

7606, 7609,  
7613



## Services Modules


Distributed Security;  
IPSEC, Firewall, IDS,  
DoS Protection

# Introducing the latest Software Release

## Cisco IOS 12.2(33)SRC

- Key Initiatives in this release
  - Improve SIP200/SIP400 SPA parity
  - Many new High Availability & Convergence Features
  - Introduction of ISG

Over 150 New Software Features!



# Cisco 7600 - “S” Class Chassis

## • High Availability






- Enables fast fabric sync (<100 ms switchover)
- Redundant channels for inter-line card control plane communication
- Redundant fan trays in 7609-S

## • 80G/slot ready

## • Power and Cooling

- Power & cooling capacity of over 600W per slot to support new high-density Ethernet linecards

Not EC'd yet

	<b>Cisco 7613-S</b>	Rack Units	18	Rack Density	2
		Slots	13		
	<b>Cisco 7609-S</b>	Rack Units	21	Rack Density	2
		Slots	9		
	<b>Cisco 7606-S</b>	Rack Units	7	Rack Density	6
		Slots	6		
	<b>Cisco 7604</b>	Rack Units	5	Rack Density	9
		Slots	4		
	<b>Cisco 7603-S</b>	Rack Units	4	Rack Density	11
		Slots	3		

S Chassis	Availability
7609S	Shipping
7606S	Shipping
7604	Shipping
7603S	Shipping
7613S	Radar

# 7600 Engines Comparison



	SUP720	RSP720	RSP720-10GE
Control Plane	MSFC3	MSFC4	MSFC4
Ctrl Plane CPU	600Mhz MIPS	1.2GHz PowerPC	1.2GHz PowerPC
DRAM	1GByte (DDR)	Up to 4GByte (DDR2)	Up to 4GByte (DDR2)
NVRAM	2Mbyte	4MByte	4MByte
Bootflash/bootdisk	64MByte	512MByte	512MByte
Forwarding Plane	PFC3B, PFC3BXL	PFC3C, PFC3CXL	PFC3C, PFC3CXL
MAC (CAM) Table Size (pract./theor.)	32k/64k	80k/96k	80k/96k
IP Subscriber Termination	x	32k	32k
IP Forwarding	30Mpps	30Mpps	30Mpps
MPLS Forwarding	20Mpps	20Mpps	20Mpps
On board Ports	2x GE	2x GE	3x GE, 2x 10GE

# Ethernet Services 20G Modules

## Overview

- Two 20 Gbps (full-duplex) linecard options:
  - 2-port 10 Gigabit Ethernet, XFP LAN-PHY optics  
*Future support for XFP WAN-PHY optics*
  - 20-port Gigabit Ethernet, SFP optics
- Up to 30 Mpps distributed performance per linecard
  - Line rate for 64-byte packets, L2 or L3
- 512 MB packet buffering (200 ms) per 10G ports
- Full support for online insertion & removal (OIR)
- Choice of hardware-based scale for L2 & L3 services:

DFC-3C daughter-card


TCAM entries: 256K IP FIB & 128K Netflow

DFC-3CXL daughter-card

TCAM entries: 1M IP FIB & 256K Netflow



7600-ES20-10G  
(aka BALDUR)



7600-ES20-GE  
(aka LOKI)

# ES20 SW Feature Highlights

The most flexible line card in the industry – full L2 & L3 support

- L2 Switchport Feature Parity
- EVC Ethernet Infrastructure
  - VLAN local significance
  - VLAN Scale (16K per ES20)
  - Advanced VLAN translation
  - Flexible Service mapping
- Rich QoS Feature
  - 16K hardware queues
  - 2K hierarchy policies
  - Egress H-QoS (shaping, LLQ/CBWFQ)
  - Ingress Policing
  - Ingress and egress marking and classification
- L2VPN Support
  - VPLS/H-VPLS
  - Fast LoS Detection for TE/FRR
  - PW Redundancy, H-VPLS Redundancy
  - PW Stitching
- Ethernet and MPLS OAM
- L3 VPN, IPv6, L3 Multicast and so on

# Cisco 7600 Flexible Ethernet UNI

## Convergence of Residential Quad Play + Business VPN

- **The Flexible Ethernet UNI** defines a unique, virtual L2 or L3 service instance per customer

- A service instance can be a MAC address, VLAN, Q-in-Q VLAN, L2 VPLS pseudowire, IP address, or L3 MPLS VPN

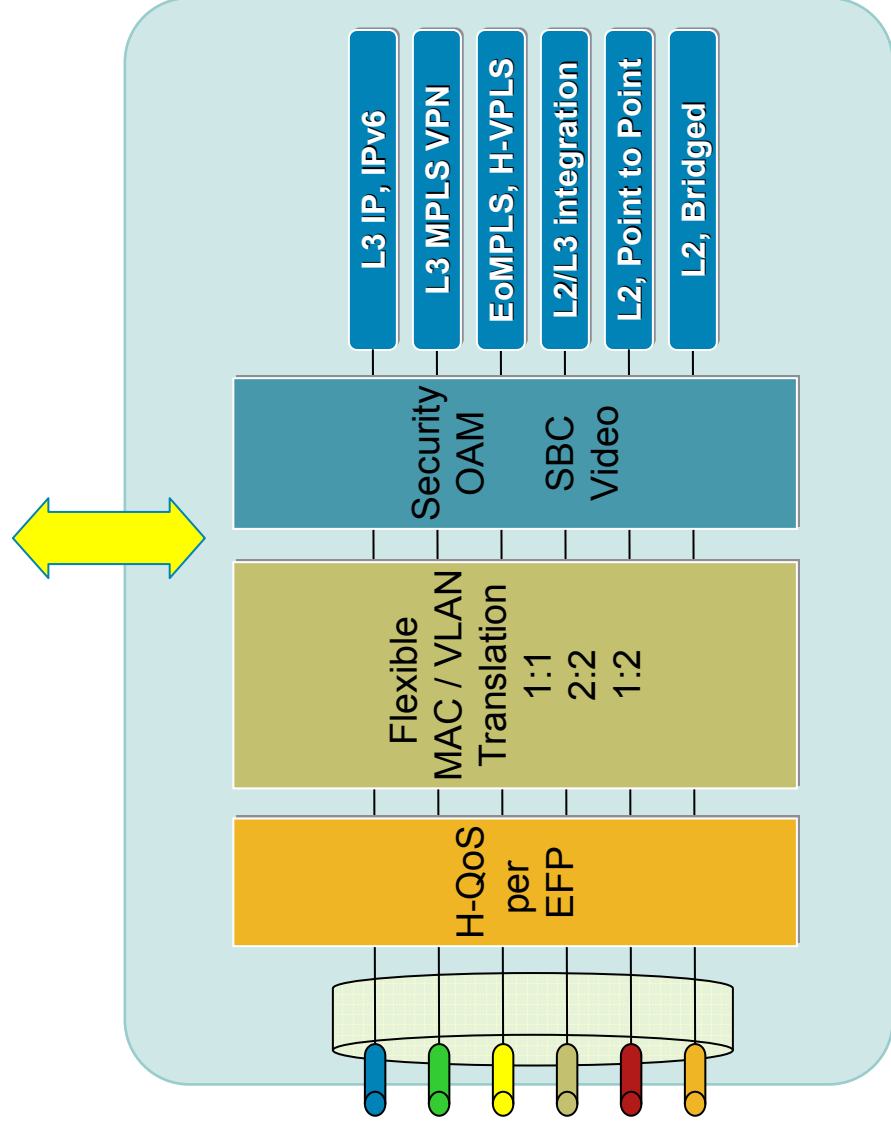
- **For each service instance, Flexible UNI offers:**

- Unique ID with service separation via VLAN or MAC translation
- H-QoS with shaping per VC
- IP+MAC spoofing prevention
- Ethernet and MPLS OAM

- **Each service instance can in turn be flexibly mapped to:**

- L2: Pseudowires, H-VPLS
- L3: IP, IPv6, MPLS VPN

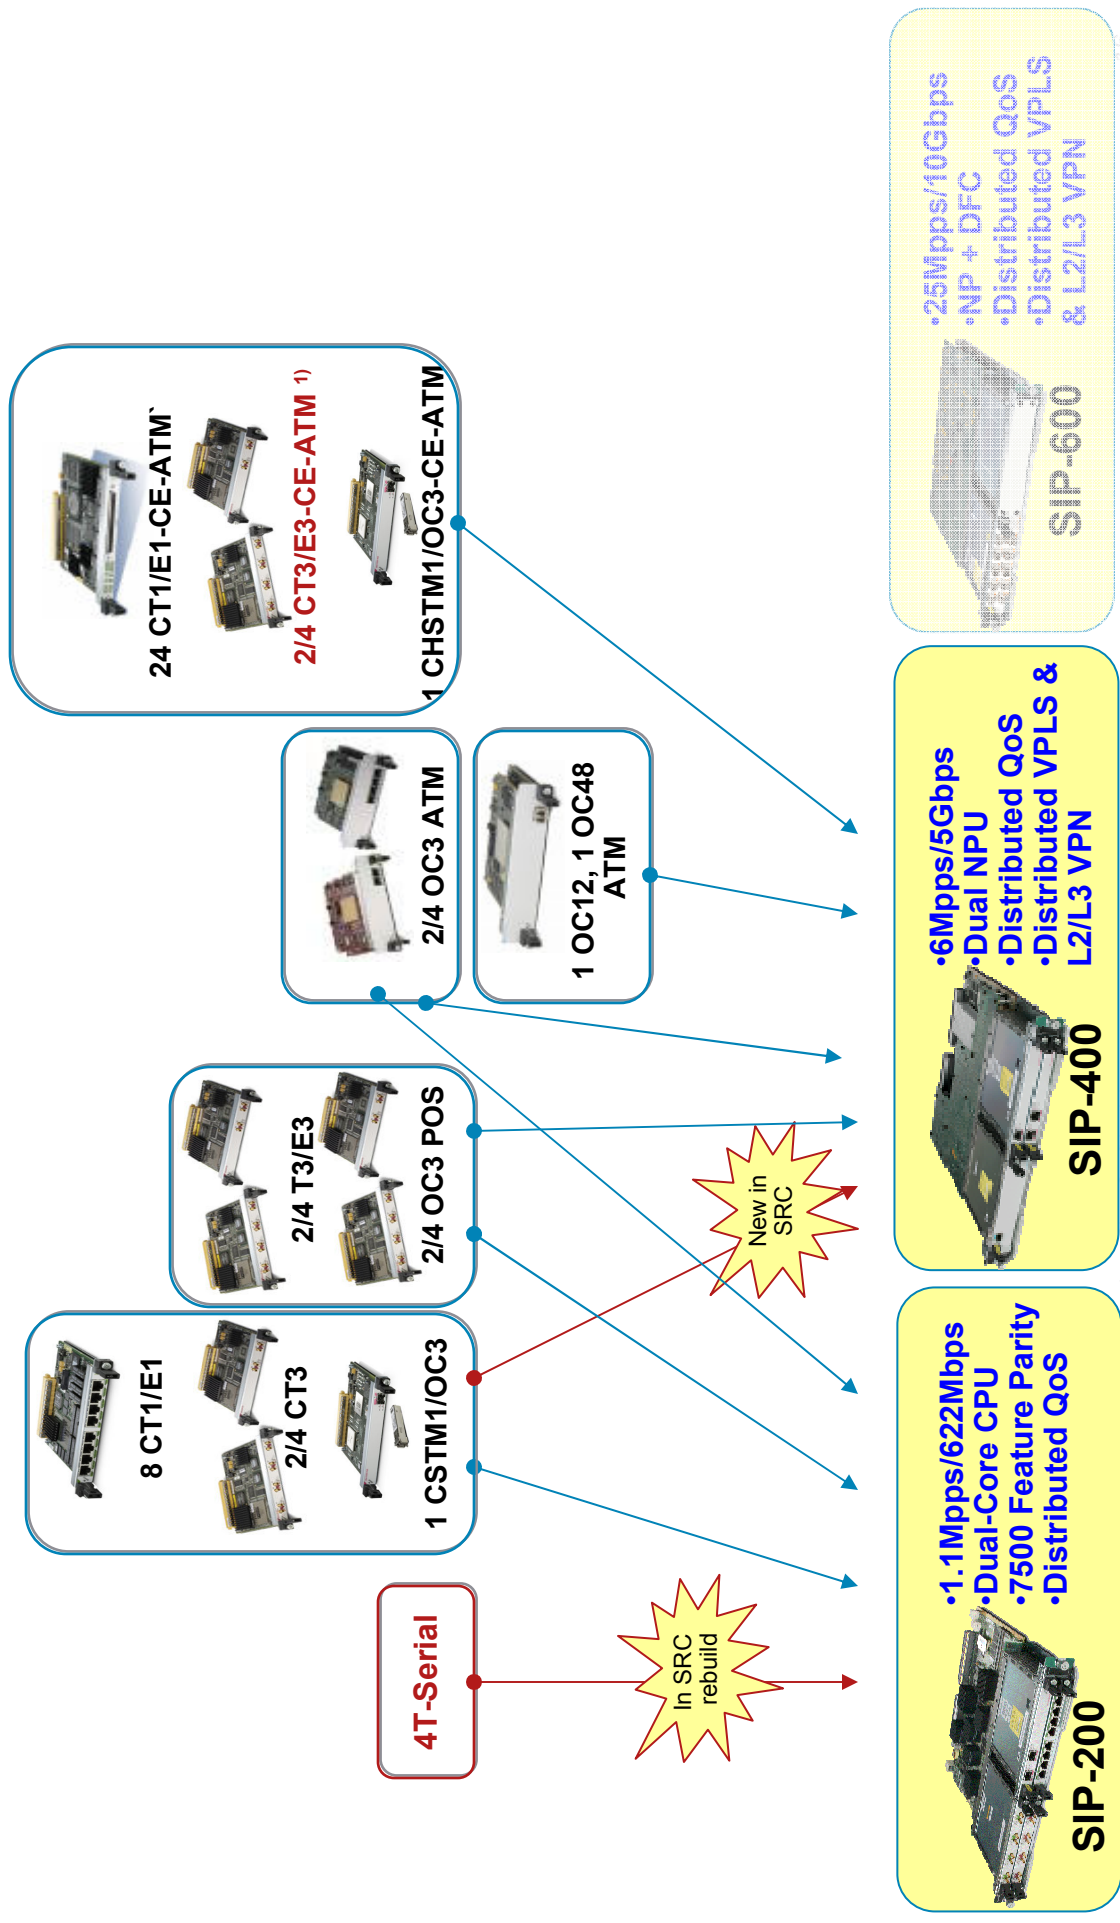
### OSS / Policy Management



# Cisco 7600 SPA/SIP Family

## I-Flex Interface Portfolio below 1Gbits

1) Clear Channel T3 ATM only in SRC







# Doubling the 10GE Density on the 7600

- New card called WS-X6708
- Already supported in SXF on cat6500 and 7600, but now in also in SR train
- Integrated DFC3C or DFC3CXL daughter-card for up to 48Mpps
- 40Gbits Fabric connectivity
- Increased Performance for Multicast as well as complex Feature combinations
- Cisco 7609 can now support to up 7x8=65 10GE ports !
- Card can be configured to only enable 4 ports (performance mode) or all 8 ports (2:1 oversubscribed mode)
- **No VPLS, per VLAN/Service Queuing or EVC support**



# Service Module Support

- SUP720 has the widest Service Module Support
- RSP720 got introduced in 12.2(33)SRB but without Service Module support
- RSP720 picked up SM supported in 12.2(33)SRC

	sup720		RSP720	
	SRB	SRC	SRB	SRC
NAM-1	Yes	Yes	NO	Yes
NAM-2	Yes	Yes	NO	Yes
CSG1	Yes	Yes	NO	NO
MWAM	Yes	Yes	NO	NO
SAMI	Yes	Yes	NO	NO
CSG2	Yes	Yes	NO	NO
PSD	Yes	Yes	NO	Yes
IDSM2	Yes	Yes	NO	Yes
FWSM	Yes	Yes	NO	Yes
ACE	Yes	Yes	NO	Yes
ADM	Yes	Yes	NO	Yes
AGM	Yes	Yes	NO	Yes
7600-SSC-400/IPSec	NO	Yes *SRC rebuild	NO	Yes *SRC rebuild
WiSM1	NO	Yes	NO	NO
SBC	Yes	Yes	NO	Yes

# New System HA Features

- ATOM NSF/SSO/eFSU/ISSU  
ATM, CEM, FR, HDLC, Ethernet, PPP
- VPLS NSF/SSO/eFSU/ISSU
- BFD SSO Admin Down (i.e. Phase 1 of BFD SSO)
- FHRP/VRRP SSO/ISSU
- OSPF Graceful Shutdown
- Hot Fabric Sync

# New Convergence Features

## Layer 2 related

- VPLS N-PE Redundancy with MPLS Access
- VPLS N-PE Redundancy with QinQ Access (MST)
- VPLS MAC Address Withdrawal
- REP
- CEOPs AC Redundancy for Local Switching
- ATOM PW Status Signaling
- LACP Fast Mode

## Layer 3 related


- VRF aware BFD
- BFD for Static/Connected routes
- BFD on WAN interfaces
- OSPFv3 (IPv6) Fast Convergence
- TBAC (Video)

# New Convergence Features

## Core Infrastructure

- LDP Local Label Allocation Filtering
- BFD triggered FRR
- TE FRR over Link Bundles
- TE FRR Path Protection

# Details on new SW Features



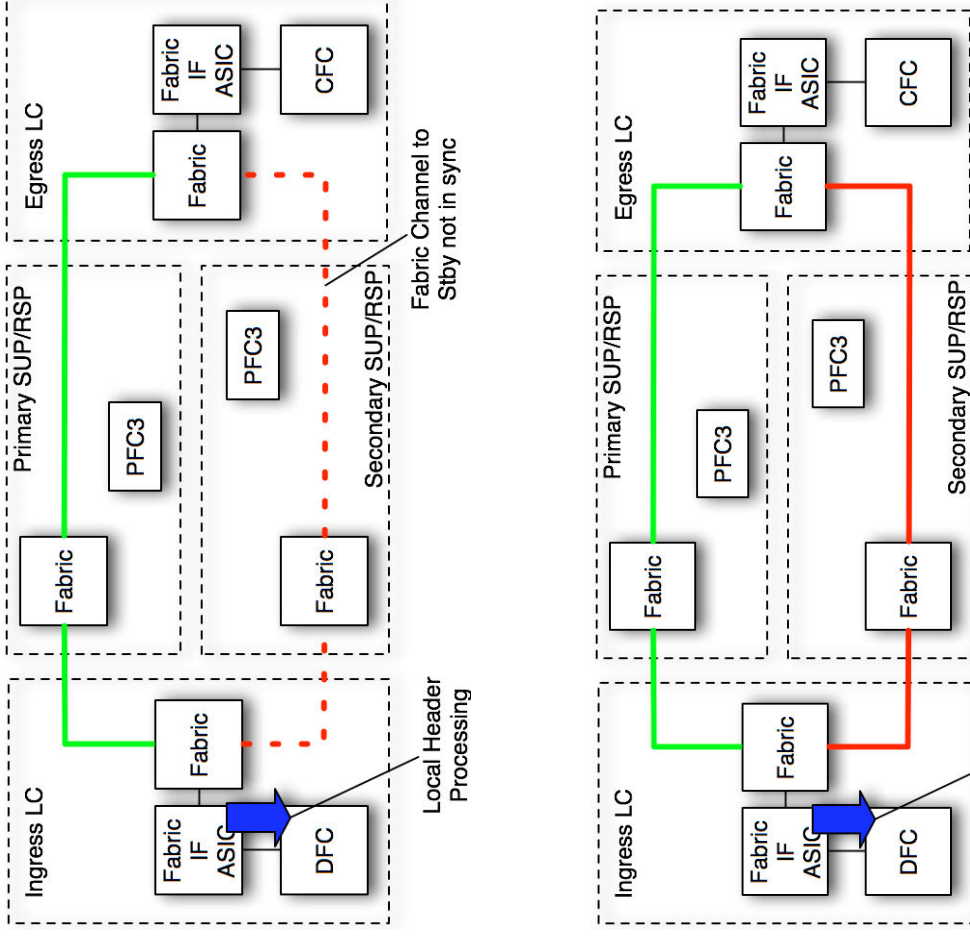
...

# Hot Fast Fabric Sync

## 76xx-S and 7604 Chassis Only

### Before Hot Fabric Sync

- Fabric Chips on LC and SUP/RSP maintain only one channel in sync
- **LCs are informed via SCP message by SUP/RSP to switch over**
- Fabric Chips need to bring up 2nd Channel
- Traffic can pass again



### With Hot Fabric Sync

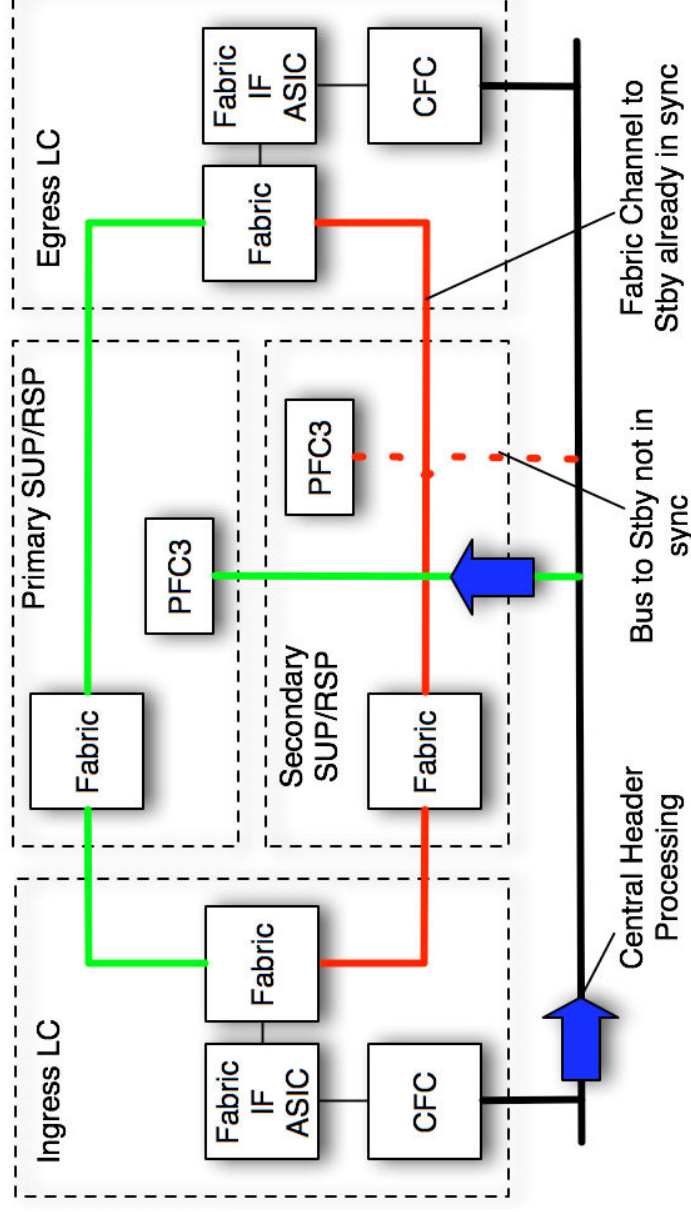
- Fabric Chips on LC and SUP/RSP maintain both fabric channels in sync
- Linecard only needs to be informed to use standby channel
  - a) **SCP Message from SUP/RSP to LC**
  - b) **Interrupt from SUP/RSP to LC**
- Traffic can pass again



# Hot Fabric Sync Performance per Linecard

	Today	With Hot Fabric Sync	Limiting Factor
<b>ES20, SIP600</b>	<b>1-1.2sec</b>	<b>&lt;100msec</b>	<b>Interrupt driven Fabric Sync</b>
SIP400, SIP200, eFlexWAN	0.8-1.5sec	600-800msec	Bus Sync
67xx with DFC	1-1.2sec	200-400msec	SCP Messages for Fabric Sync
67xx with DFC (post Cobra)	1-1.2sec	<100msec	Interrupt driven Fabric Sync
67xx with CFC	0.8-1.5sec	600-800msec	Bus sync
6xxx (classic)	0.6-0.8sec	600-800msec	Bus Sync

# Hot Fabric Sync with non-DFC cards



- **Bus Sync** is the limiting Factor = 600-800msec
- Header processing of ingress LC is done on central PFC3
- Hot Fabric Sync happens fast, but does not matter as we can not perform lookups until Bus in operational again

# Monitoring Hot Fabric Sync

- You can display Fabric channel status with the following command
- In addition “debug fabric ...” commands are available on SP console

```

tuonno#sh fabric status
slot channel speed module fabric support
1 0 20G OK OK
2 0 20G OK OK
2 1 20G OK OK
3 0 20G OK OK
3 1 20G OK OK
4 0 20G OK OK
4 1 20G OK OK
5 0 20G OK OK
6 0 20G OK OK
7 0 20G OK OK

tuonno#sh mod
Mod Ports Card Type Model
-----
1 24 CEF720 24 port 10000mb SFP WS-X6724-SFP
2 20 ESM20G 7600-ES20-GE3C
3 4 CEF720 4 port 10-Gigabit Ethernet WS-X6704-10GE
4 8 CEF720 8 port 10GE with DFC WS-X6708-10GE
5 2 Route Switch Processor 720 (Hot) RSP720-3CXL-GE
6 2 Route Switch Processor 720 (Active) RSP720-3CXL-GE
7 0 4-subslot SPA Interface Processor-400 7600-SIP-400
9 0 4-subslot SPA Interface Processor-200 7600-SIP-200

```

# L2VPN SSO/ISSU

- No extra commands introduced
- Supported for targeted LDP and local switching configurations
- ATOM related commands are sync'd as part of the config between active and standby SUP/RSP
- If there is a version command mismatch, the router will revert back to RPR+
  - Commands with version mismatch are reported in ISSU show/debug outputs
- Features supported:
  - AToM P2P: Eth (all flavors), ATM, FR, HDLC, PPP, CEM
  - VPLS
  - ATM/FR Local Switching, TDM
  - Tunnel Select
  - Interworking

# BFD SSO Admin-down

- Provide a mechanism by which BFD does not inform clients (any RP like OSPF or other functions like TE FRR, ...) of a failure in the forwarding plane when a planned switchover occurs
  - BFD “poll message” is sent to BFD peer in case of a switchover
  - BFD session is put in “Admin-down” state by peer and confirmation is sent back via “Final message”
  - Upon receipt of Final message, local BFD session state is sent to Admin-down as well
- Planned switchover is initiated via following command
  - `Router#redundancy force-switchover`
- Implementation is based on draft-ietf-bfd-base-05
- Only supported with SSO, no plans for RPR or RPR+

# OSPF Graceful Shutdown

- Used if a network operator wants to temporarily disable a protocol without losing the configuration
  - e.g. operator wants to upgrade the router but wants to minimize network wide downtime due to ISSU not available on the router
- Router tells neighbors to not send traffic anymore as he will go down
  - This is complete opposite to the NSF concept!
- Neighbors reroute around the router being shut down
- Traffic loss is minimized
  - Neighbors passive detection (time depends on OSPF tuning) is avoided by active information

# OSPF Graceful Shutdown

- New CLI at router/protocol and interface level
- What happens?

Flooding of self-generated MaxAge Router-LSA

OSPF process is killed internally

MaxAge LSA might not be received by neighbor, hence 1x Empty Hello is sent that will bring adjacency down

If also the empty hello is not received, dead timer will bring adjacency down

```
tuonno(config-if)#int gi 1/1
tuonno(config-if)#ip ospf shutdown ?
<CR>
```

```
tuonno#sh ip ospf int gi 1/1
GigabitEthernet1/1 is down, line protocol is down
(notconnect)
Internet Address 3.3.3.1/24, Area 0
OSPF Routing Process is shutdown
...
```

```
tuonno(config)#router ospf 4711
tuonno(config-router)#shutdown ?
<CR>
```

```
tuonno#sh ip ospf 4711
Routing Process "ospf 4711" with ID 10.1.6.48
Start time: 1d04h, Time elapsed: 00:02:26.380
Routing Process is shutdown
...
```

# LDP Local Label Allocation Filtering

- Per default LDP generates label bindings for all IGP or statically derived prefixes
- To optimize MPLS L3VPN end2end VPN convergence, people want to limit label bindings to PE loopbacks in order to make IGP converge faster
- This feature allows use of **prefix-lists to control label bindings being generated**

```
tuonno(config)#mpls ldp la
tuonno(config-ldp-lbl)#allocate ?
global Specify global Routing/Forwarding instance

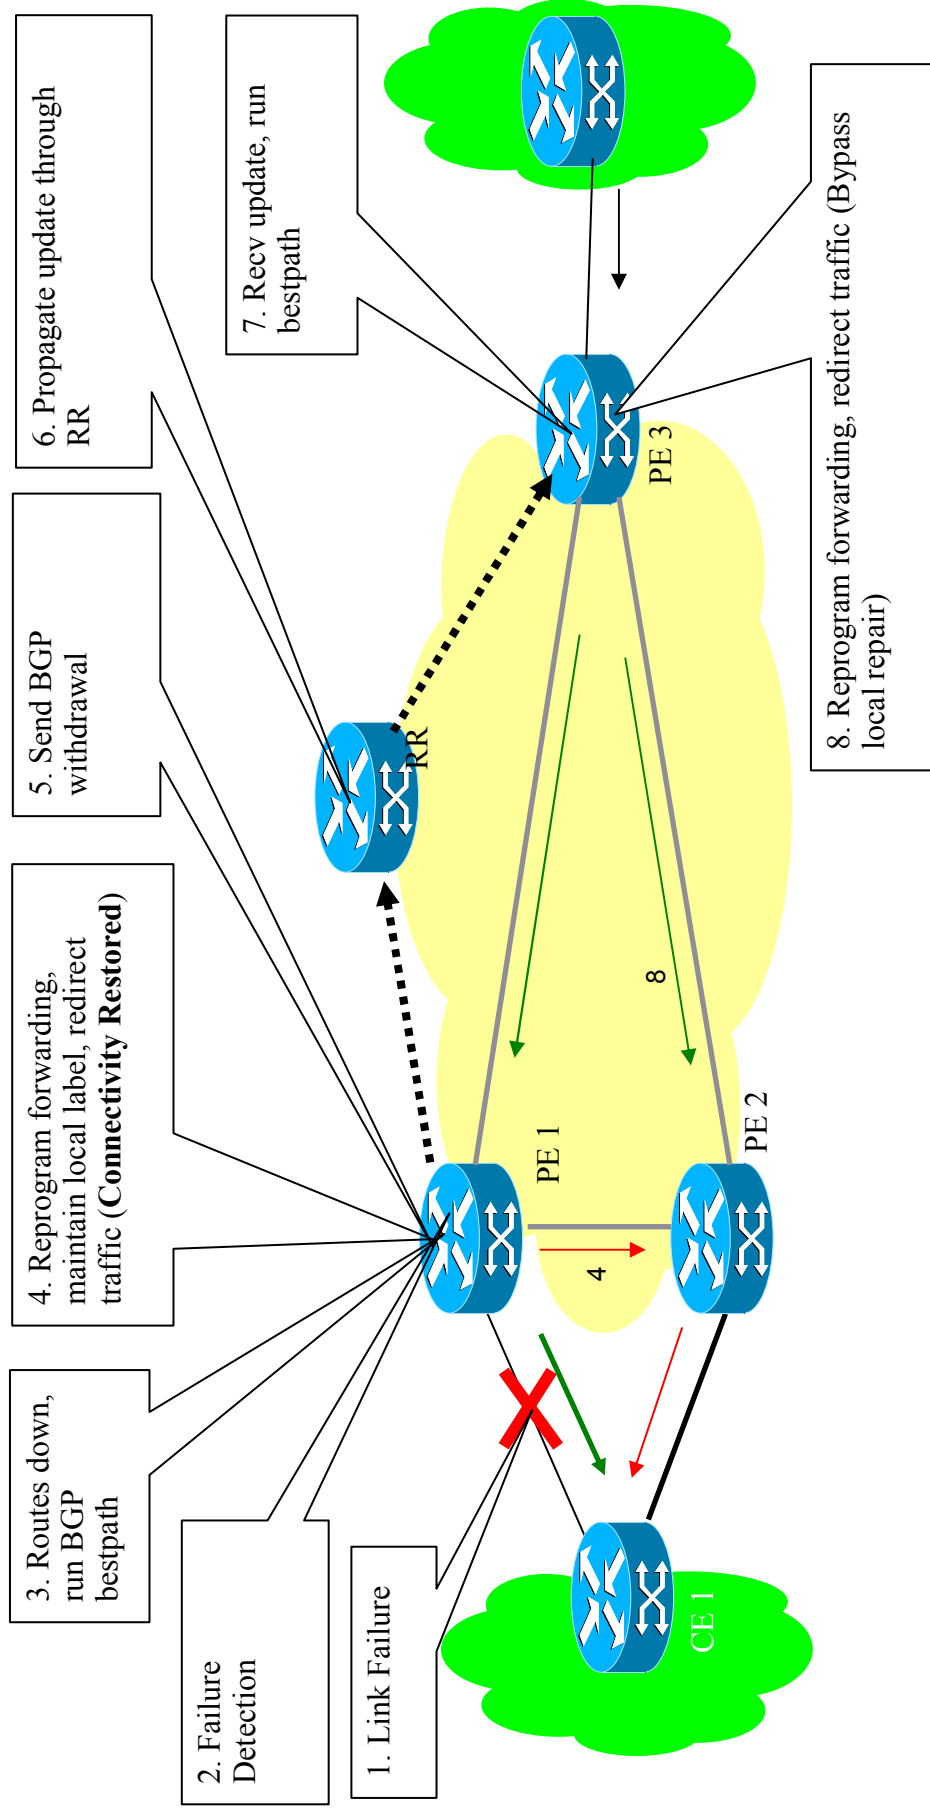
tuonno(config-ldp-lbl)#allocate global ?
host-routes allocate local label for host routes only
prefix-list Specify a prefix list for local label filtering
<cr>
tuonno(config-ldp-lbl)#allocate global prefix-list ?
WORD IP prefix-list for destination prefixes;
name or number (1-99)
```

```
tuonno(config)#ip prefix-list test seq 4 permit ?
A.B.C.D IP prefix <network>/<length>,
e.g., 35.0.0.0/8

tuonno#sh ip prefix-list
ip prefix-list test: 2 entries
seq 5 permit 1.1.0.0/23
seq 34 deny 32.0.0.0/3
```

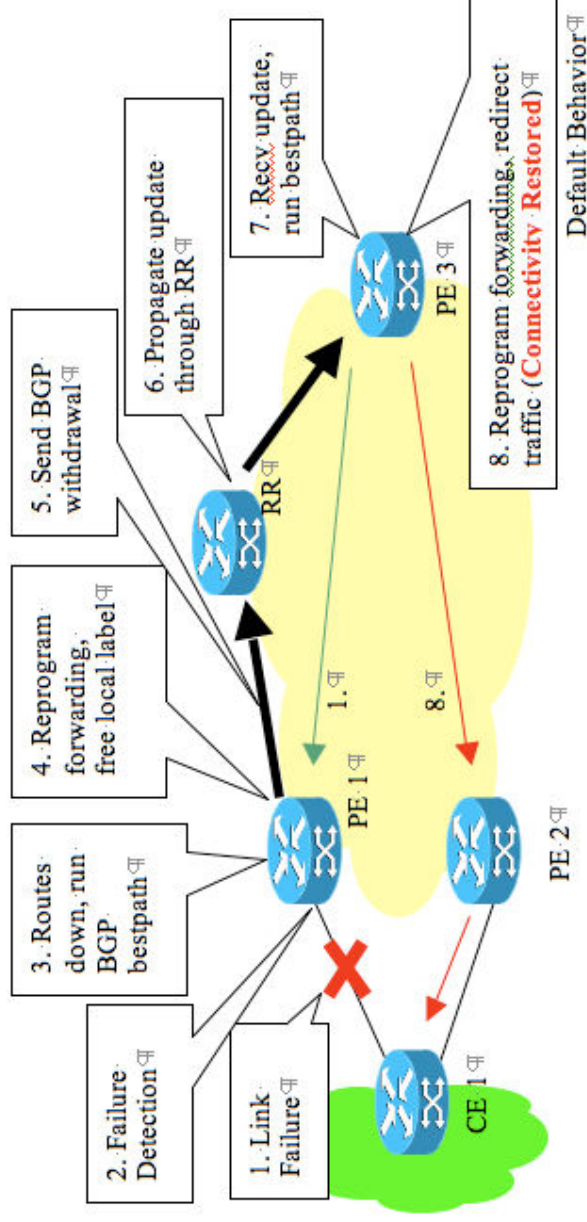


# PE-CE Local Link Protection



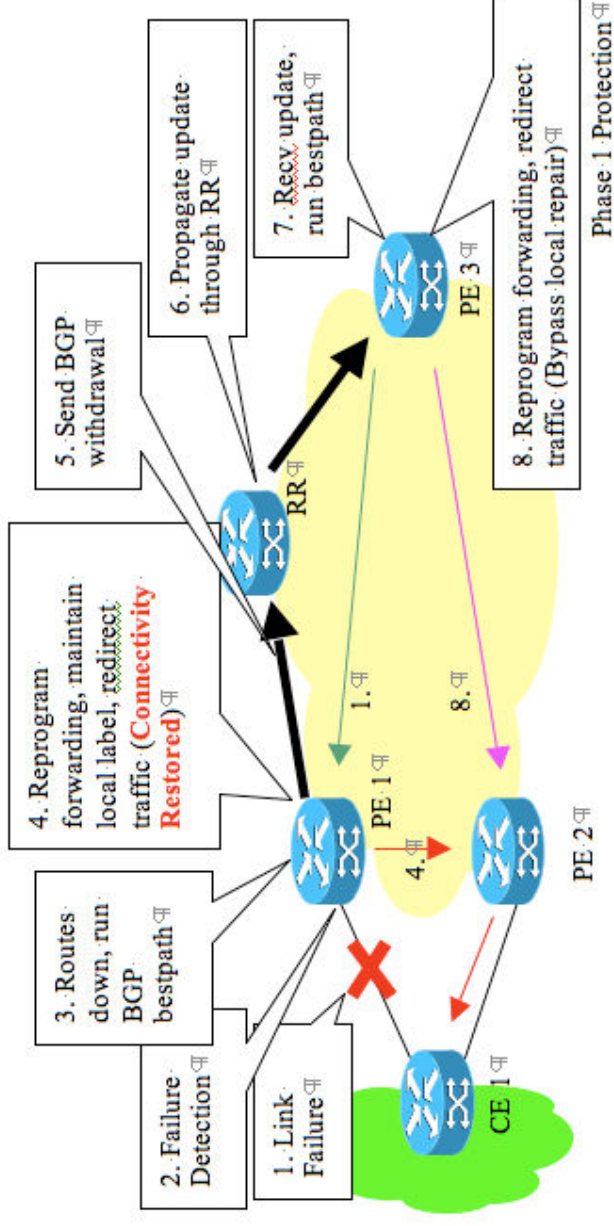
Source: Ted Qian (NSSTG PM)

# PE-CE Local Link Protection



- Edge Link Failures trigger BGP convergence
- Local Labels are removed
- Traffic loss happens until remote PE's BGP LFIB has converged
- BGP convergence is directly proportional to number of VRF routes

# PE-CE Local Link Protection



- BGP convergence is still triggered
- Local Labels are kept and traffic is restored locally
- Cleanup (removal of Local labels) will happen after 5 minutes

# Configuring PE-CE Local Link Protection

- Can be enabled/disabled on a per VRF basis
- Cleanup Timer of 5 minutes is not configurable

```
tuonno(config-vrf)#protection ?  
local-prefixes Enable protection for local prefixes
```

```
tuonno#sh ip vrf detail vrf1  
VRF vrf1 (VRF Id = 1): default RD 4711:1; default VPNID <not set>  
No interfaces  
VRF Table ID = 1  
Export VPN route-target communities  
  RT:4711:1  
Import VPN route-target communities  
  RT:4711:1  
No import route-map  
No export route-map  
VRF label distribution protocol: not configured  
VRF label allocation mode: per-prefix  
vrf-conn-aggr for connected and BGP aggregates (No Label)  
Local prefix protection enabled
```


# PE-CE Link Protection

## Test Setup Details

- 100k routes on UUT
  - 10 VRFs
  - 10k routes per VRF
  - Remote : Local ration = 99:1
- Number or protected VRFs
  - 1, 2, 5, 10 (for each VRF 100 local routes are affected)
- Failure Detection
  - VRF aware BFD

# PE-CE Link Protection


## Test Results



Source: Ted Qian (NSSTG PM)

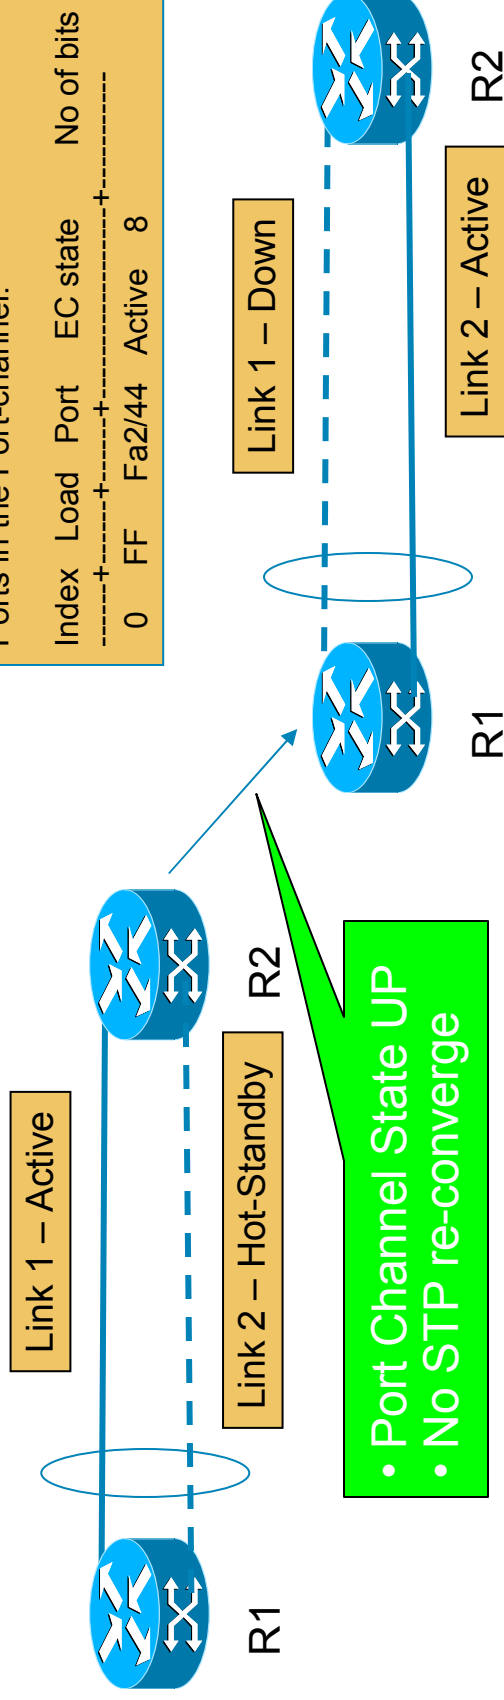
# LACP 1:1 Link Redundancy Enhancement

- **Drivers**  
QoS  
Attaching a QoS policy onto a port-channel will create an independent policy per member link  
**Simulate SONET/SDH APS/MSP**  
People used to TDM like the active/standby approach
- **Problem Details:**
  - Current LACP hot-standby can be configured for max-bundle 1-8
  - By setting, LACP max-bundle = 1, we achieve LACP 1:1 link redundancy
  - 30 second traffic loss during switchover or reversion mainly due to STP re-convergence



# LACP 1:1 Link Redundancy Enhancement

- **Solution Details:**
  - Supports LACP 1:1 link redundancy with:
    - Fast switchover
    - Automatic reversion
  - Active link has higher system and port priority and lower one is hot-standby
  - Port-channel remains up UP and no STP re-convergence during:
    - fast switchover
    - automatic reversion
  - traffic switchover loss < 1 second



```

R(config)#int po6
R(config-if)#lACP fast-switchover
R(config-if)#lACP max-bundle 1

R#sh etherc 6 port-channel

-----
Port-channels in the group:
-----
Port-channel: Po6 (Primary Aggregator)
-----
Age of the Port-channel = 0d:01h:54m:42s
Logical slot/port = 14/3    Number of ports = 1
Port state = Port-channel Ag-Inuse
Protocol = LACP
Fast-switchover = enabled

Ports in the Port-channel:

Index Load Port EC state No of bits
-----+-----+-----+-----+-----
0 FF Fa2/44 Active 8
  
```




# TE FRR over Link Bundles

- TE supported on  
L3 EtherChannel,  
SVI across L2 EtherChannel  
MLPPP
- FRR supported, triggered only by bundle link going down due to member links dropping **below configured min-links**.
- FRR cutover time is depending on link down detection time of above option and configuration. (for sure not <50msec)
- Bandwidth reservation/preemption is supported.  
This means we are NOT limiting ourselves to zero bandwidth tunnels
- Support for setting percentage of global pool bandwidth on bundled interface.

# TE FRR over Link Bundles – Examples


## Layer 3 Port-Channel



```
Int gi 1/1
 channel-group 1
Int gi 1/2
 channel-group 1

Int port-channel 1
 ip address ...
 mpls ip
 mpls traffic-eng tunnels
 mpls traffic-eng backup-path tun 1
 port-channel min-links 2
 ip rsvp signalling hello
```

## Layer 2 Port-Channel & SVI



```
Int gi 1/1
 switchport
 switchport mode trunk
Int gi 12/
 switchport
 switchport mode trunk

Int vlan 100
 ip address ...
 mpls ip
 mpls traffic-eng tunnels
 mpls traffic-eng backup-path tun 1
 port-channel min-links 2
 ip rsvp signalling hello
```

# TE FRR Path Protection

- Provides a backup tunnel between Headend and Tailend PE router
- Failure of the primary LSP can be caused by
  - RSVP refresh timeout
  - Receipt of PathErr or ResvTear
  - IGP removal of a link in the primary LSP path
  - ...
- Backup LSP is pre-signalized to reduce downtime during switchover
- Configuration

```
tunnel mpls traffic-eng path-option 10 explicit name pri-path  
tunnel mpls traffic-eng path-option protect 10 explicit name alt-path
```

# VRF Enhancements

- Per VRF AAA
- VRF aware LI
- VRF aware BFD
- Half Duplex VRF
- Several Broadband VRF Features (i.e. VRF transfer)

# Various BFD Enhancements

- Session Scale increased to 128
- VRF aware BFD
  - Very important for 67xx based FastConvergence Designs to make PE-CE link failures predictable
- BFD over WAN interfaces
  - ATM port or Sub-Interface with various AAL5 encaps
  - POS with HDLC or PPP
  - Serial with HDLC or PPP
  - POS Sub-Interface with FR
  - Serial Sub-Interface with FR
- BFD for Static Routes

```
ip route static bfd GigabitEthernet5/2 10.1.6.254  
ip route 3.3.3.0 255.255.255.0 10.1.6.254
```

```
tuonno#sh ip static route bfd  
Codes in []: R - Reachable, U - Unreachable, L - Loop  
GigabitEthernet5/2 10.1.6.254 [R]  
...
```

# Access Circuit Redundancy for ATM/TDM Local Switching

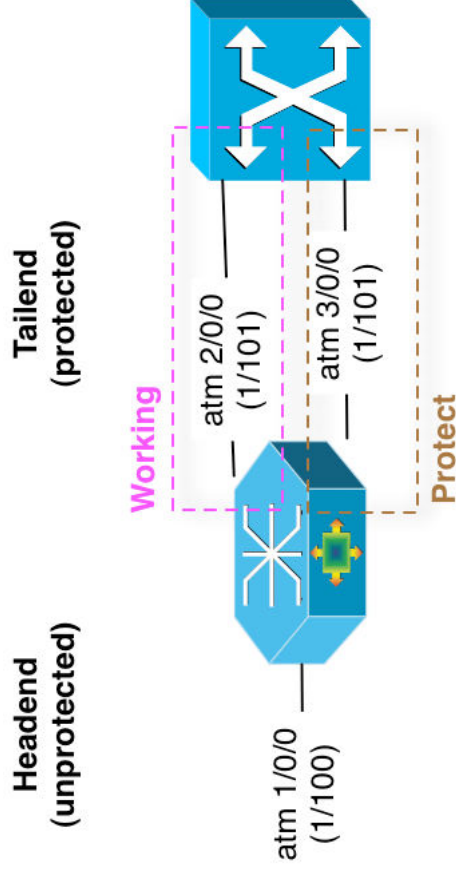
- Supported Media Types  
ATM PVC, ATM PVP, CEM

- Restrictions

Only Tailend can be protected

Port Mode is not supported

Primary and Backup circuit must be of same type




```
connect testvci A TM1/0/0 1/100 A TM2/0/0 1/101  
backup interface A TM3/0/0 1/101
```

# New Metro Redundancy Options

- Redundancy in Layer 2 networks often means Loops  
Spanning Tree is historical the protocol to prevent loops
- Popular Scenarios
  - Layer 2 rings
  - Layer 2 Access rings with VPLS core
  - Layer 2 CPE dual homed to VPLS core (i.e. datacenter interconnect)
  - Layer 3 Access rings with PW redundancy into VPLS core
- List of options that get added in 12.2(33)SRC
  - REP
  - MAC withdraw
  - MST on N-PE over special PW
  - MPLS PW Redundancy into N-PW with MAC-Withdraw

# H-VPLS with PWE3 Access

## Problem: Possible packet black hole upon PW switchover




- Assume N-PE11 is the primary N-PE. All of its MPLS links to access network go down, thus PWs to u-PE11 go down as well. Since its MPLS uplink to core network is still up, the PW to other n-PEs are still up
- For the traffic going from right to left side, N-PE21 will continue forwarding to N-PE11, and N-PE11 will drop the packet. This packet black hole will continue until N-PE21's MAC table get updated
- For bi-directional traffic, as soon as N-PE21 receive traffic from N-PE12, it will update it's MAC table accordingly. This packet black hole will stop
- For uni-directional traffic, N-PE21 will continue forwarding packet to wrong place until MAC aging out



# H-VPLS with PWE3 Access

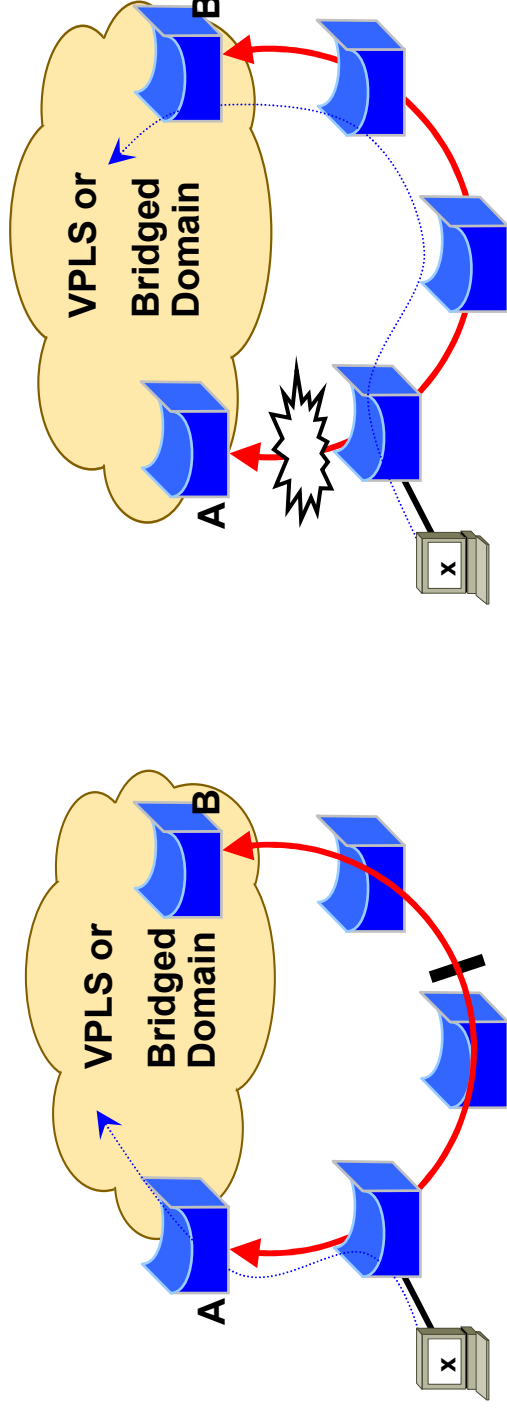
## Solution: VPLS MAC withdrawal



- U-PE switchover to backup PW if primary PW goes down
- As soon as backup PW come active, U-PE will generate MAC withdrawal message via D-LDP to N-PE12. N-PE12 will flush it's MAC table and forward this message to all its remote PEs
- After receive the MAC withdrawal message, remote PEs will flush it's MAC address table
- Packet from N-PE21 will flood to all N-PEs including N-PE12 which will forward it to U-PE. N-PE11 and N-PE31 will drop the packet
- Flooding will stop if PE receive packet from the reverse direction. No packet black hole

# Resilient Ethernet Protocol (REP)

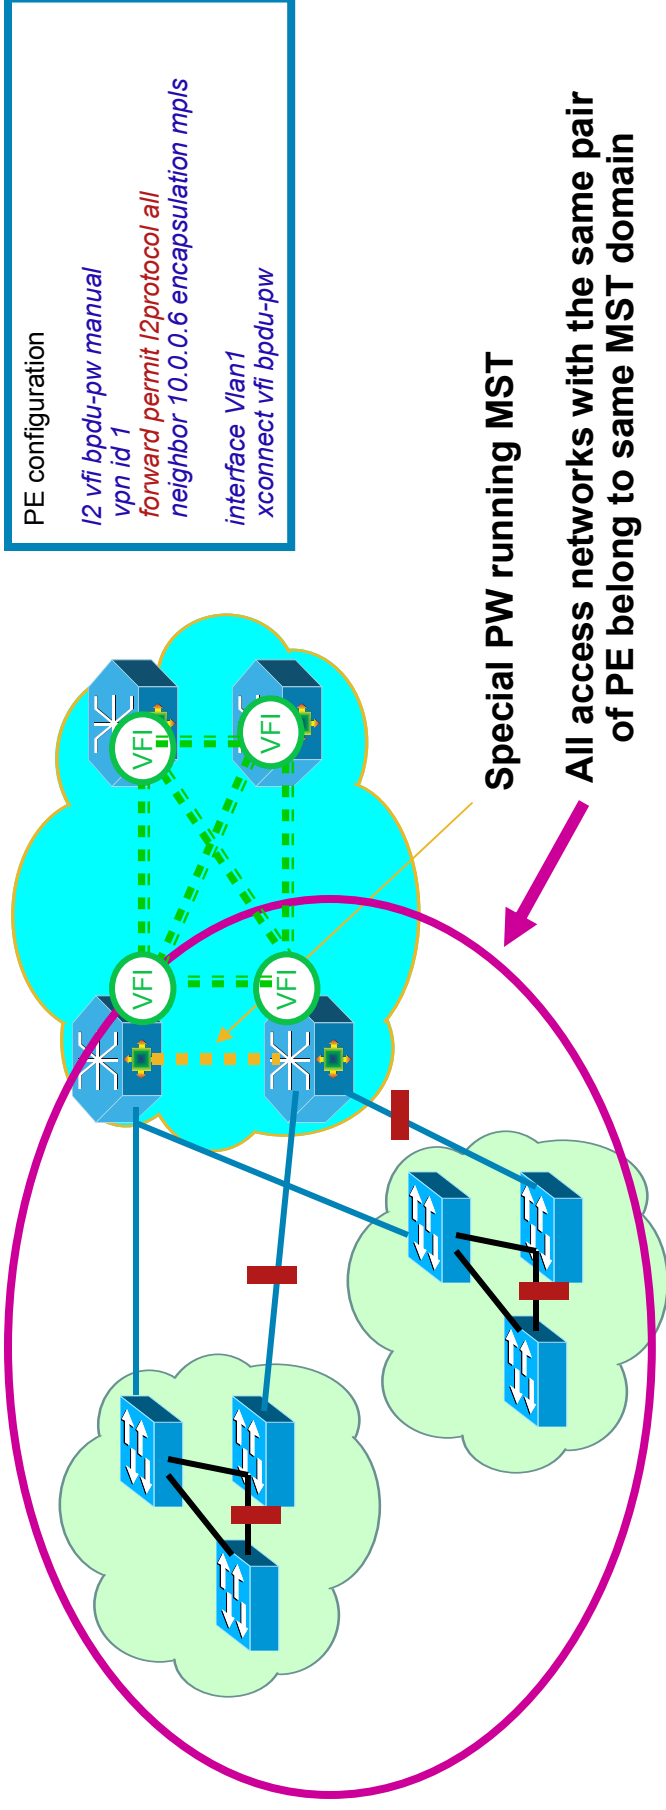
## Its a Ring/Segment Protocol



- Sub-200msec convergence on 7600
- Supported on Switchport in SRC. **EVC only supported in future release**
- REP is loop free L2 control protocol
- REP segment comprises from Port A to Port B
- Ports A and Port B are called as REP Edge Ports
- Ports are explicitly configured to be part of a segment.
- When all the links in the segment are operational, a blocked port is determined so that there is no connectivity between the edges A,B through the segment
- If a failure occurs within the segment, the blocked port goes forwarding

# H-VPLS with Ethernet Access N-PE Redundancy

## MST over PW on NPE



- All CE networks with the same pair of PEs belong to same MST domain.
- Special PW between two PEs run MST. Special PW is associated to native VLAN, **thus not pass CE data packet. Data packet goes through the regular VPLS VC**
- STP is configured in such a way to make sure
  - **Special PW is never blocked, by assigning very low STP port cost internally**
- STP TCN can trig VPLS MAC withdrawal in PE network

# Increased Metro / L2VPN Scale

- Increased Scale for EVCs and PWs
  - Pre SRC: 16k PWs for Customer Services
    - 32k PWs for VPLS mesh
  - With SRC:
    - 32k PWs for Customer Services
    - 32k PWs for VPLS mesh
- Total number of PWs per system has to be <32k
- Per Linecard EVC scale
  - ES20 - 16k
  - SIP400 - up to 16k (8k per SPA installed)
- Reminder on VPLS scale (unchanged for SRC)
  - Up to 60 peers
  - Up to 32k PWs for Mesh
  - Up to 4k VFIs

# Port-channel Support with EVC

Support is introduced on ES20

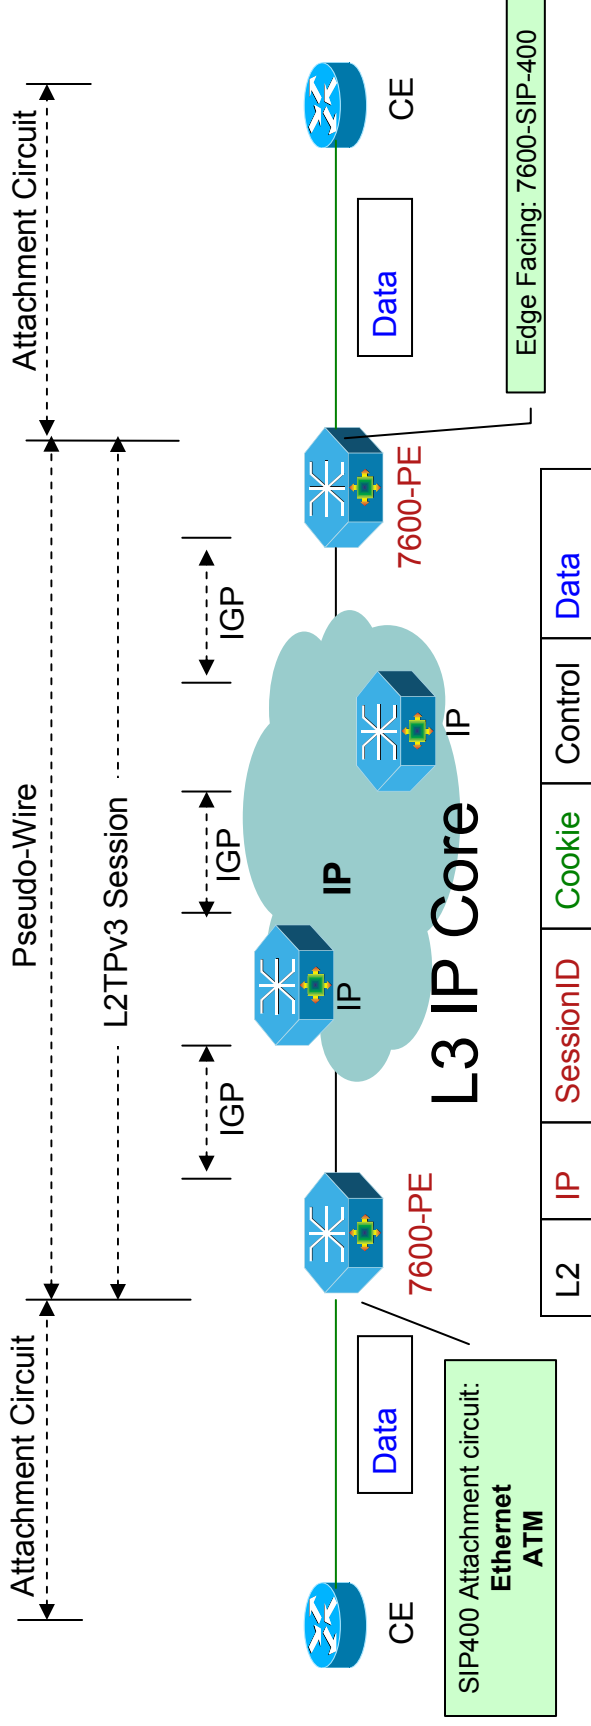
- Up to 8 member ports per channel group
- Member ports can be on the same or different ES20 card
- Service instance configuration is under port-channel interface
- LACP or PAGP is not supported in SRC
  - Only static configuration.
  - LACP plan to be supported in SRD release and SRC1
- Load balancing is not per flow, it's per service instance
- SIP-400 doesn't support port-channel

## But please be aware of the following QoS restrictions

- Ingress policing is not supported on port-channel in SRC release, plan to support in SRD
- Due to lack of BRR (bandwidth remaining ratio) support on ES20, the sum of the bandwidth configured for a particular EVC on a port-channel interface can't exceed 1Gbps (for gig port channel).
  - For example, even it has 4 member ports, the aggregated bandwidth configured under port-channel will still be limited to 1 Gbps. This restricts will be removed in SRD release along with BRR


# L2VPN over IP/L2TPv3

## Data Network Elements



- Requires SIP400 CE facing
  - L2TPv3 imposition is done on SIP400, EARL only sees IP packet
- Scale
  - Up to 512 L2TPv3 Tunnels (PE peers) per SIP
  - Up to 8k Pseudo-wires per SIP

# L2VPN over IP/L2TPv3 Configuration



```

pseudowire-class L2TPv3-class10
encapsulation l2tpv3
protocol l2tpv3
ip local interface loopback 10
  
```

```

pseudowire-class L2TPv3-class20
encapsulation l2tpv3
protocol l2tpv3
ip local interface loopback 20
  
```

```

interface Loopback10
ip address 48.10.0.14 /32
mls reserve l2tpv3 slot 7
  
```

```

interface Loopback20
ip address 48.20.0.14/32
mls reserve l2tpv3 slot 8
  
```

```

interface ATM7/0/0.1 point-to-point
pvc 0/200 l2transport
  
```

```

xconnect 48.10.0.26 200 pw-class L2TPv3-class10
  
```

```

interface GigabitEthernet 8/1/0.1
encapsulation dot1q 100
xconnect 48.20.0.26 100 pw-class L2TPv3-class20
  
```

```

pseudowire-class L2TPv3-class10
encapsulation l2tpv3
protocol l2tpv3
ip local interface loopback 10
  
```

```

pseudowire-class L2TPv3-class20
encapsulation l2tpv3
protocol l2tpv3
ip local interface loopback 20
  
```

```

interface Loopback10
ip address 48.10.0.26 /32
mls reserve l2tpv3 slot 1
  
```

```

interface Loopback20
ip address 48.20.0.26 /32
mls reserve l2tpv3 slot 2
  
```

```

interface ATM1/0/ima0.1 point
pvc 0/200 l2transport
  
```

```

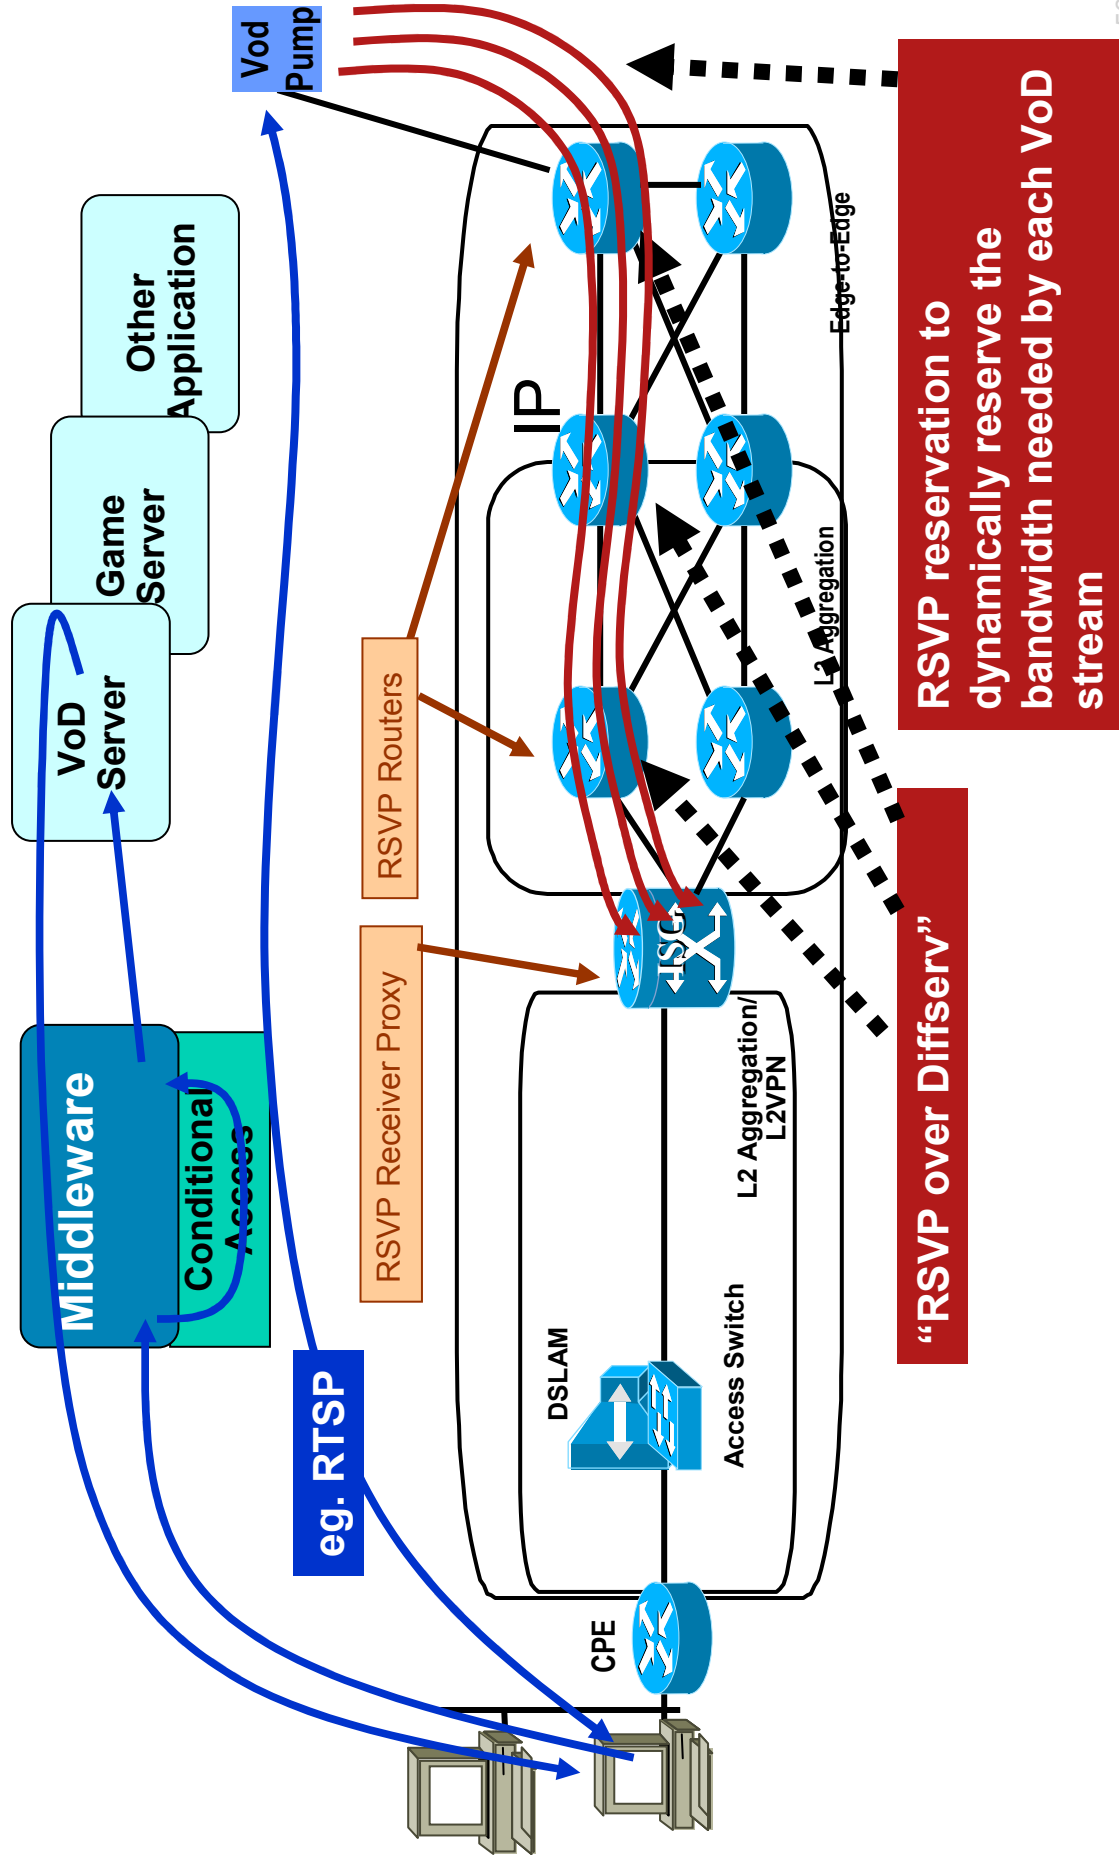
xconnect 48.10.0.14 200 pw-class L2TPv3-class10
  
```

```

interface GigabitEthernet2/1/0
encapsulation dot1q 100
xconnect 48.20.0.14 100 pw-class L2TPv3-class20
  
```

# On-Path CAC with Cisco's RSVP (Big Picture)

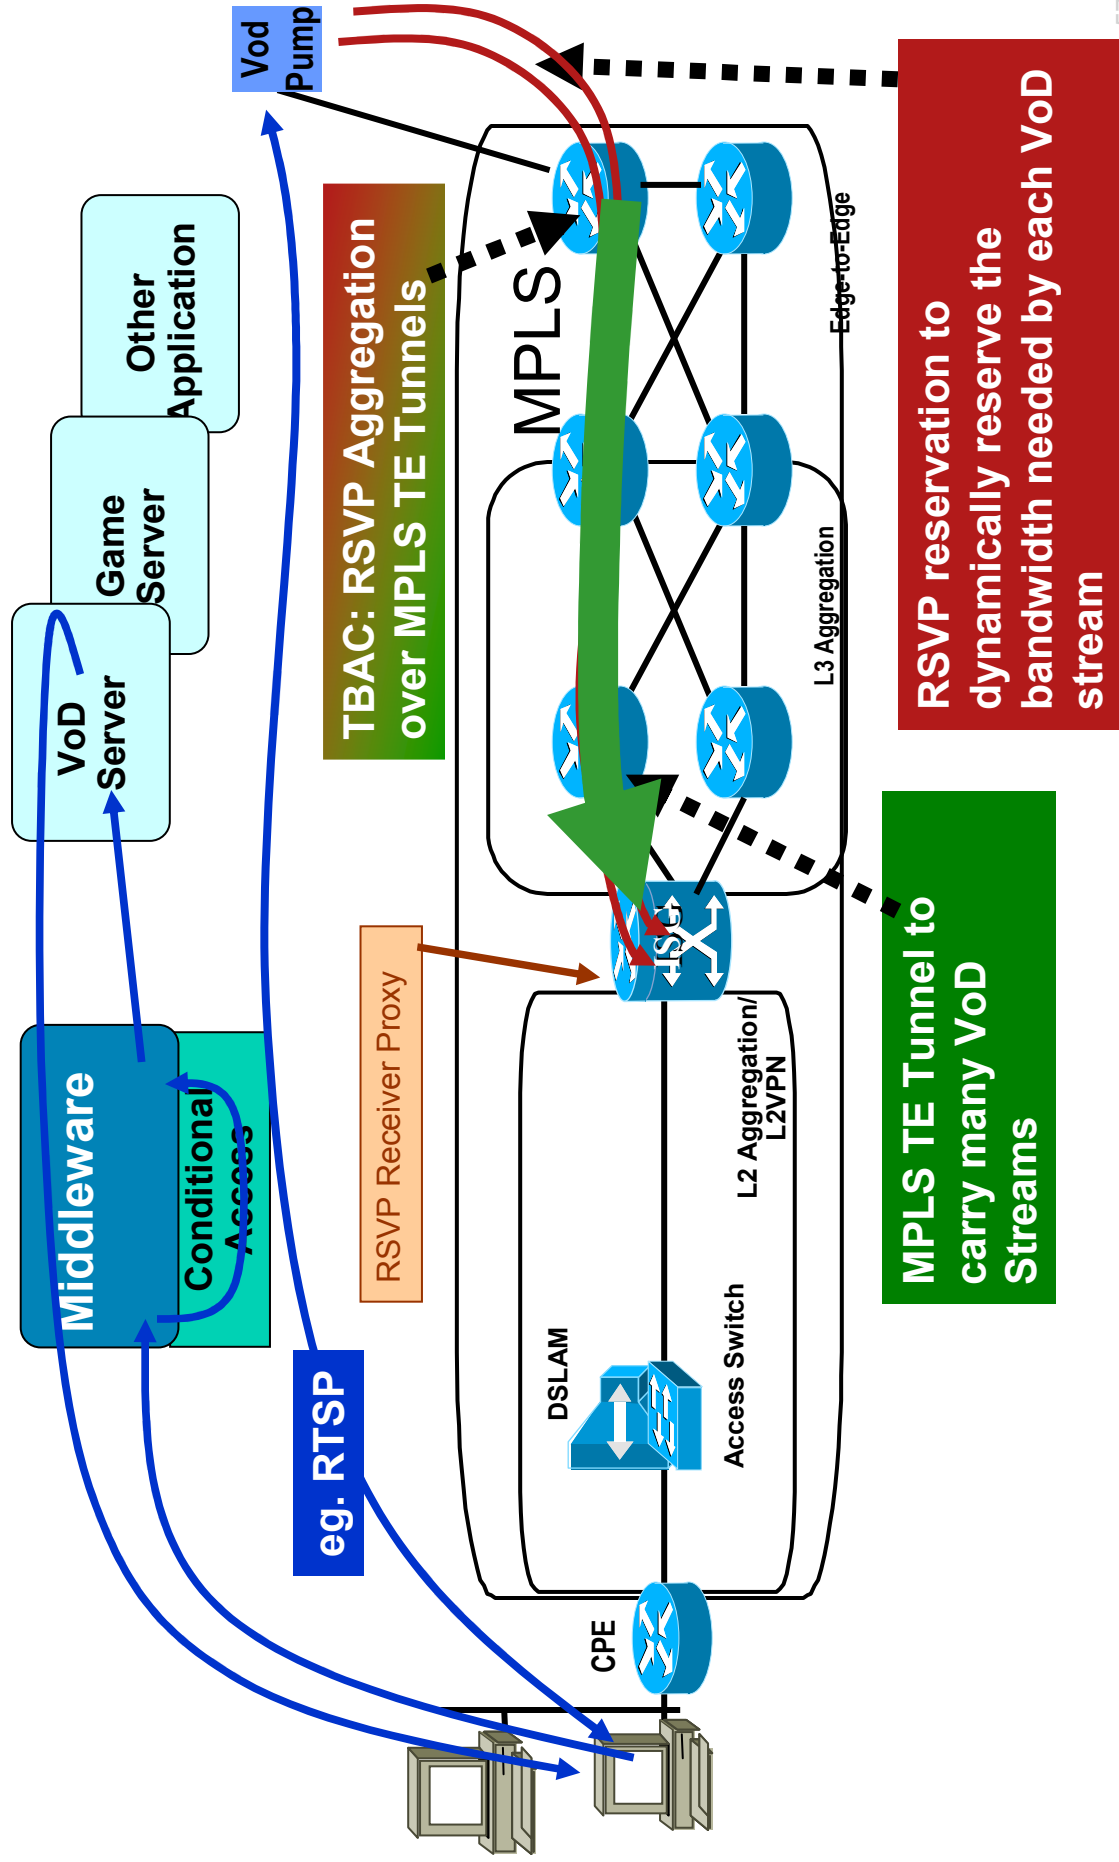
## Case Study: Admission Control for Triple-Play VoD



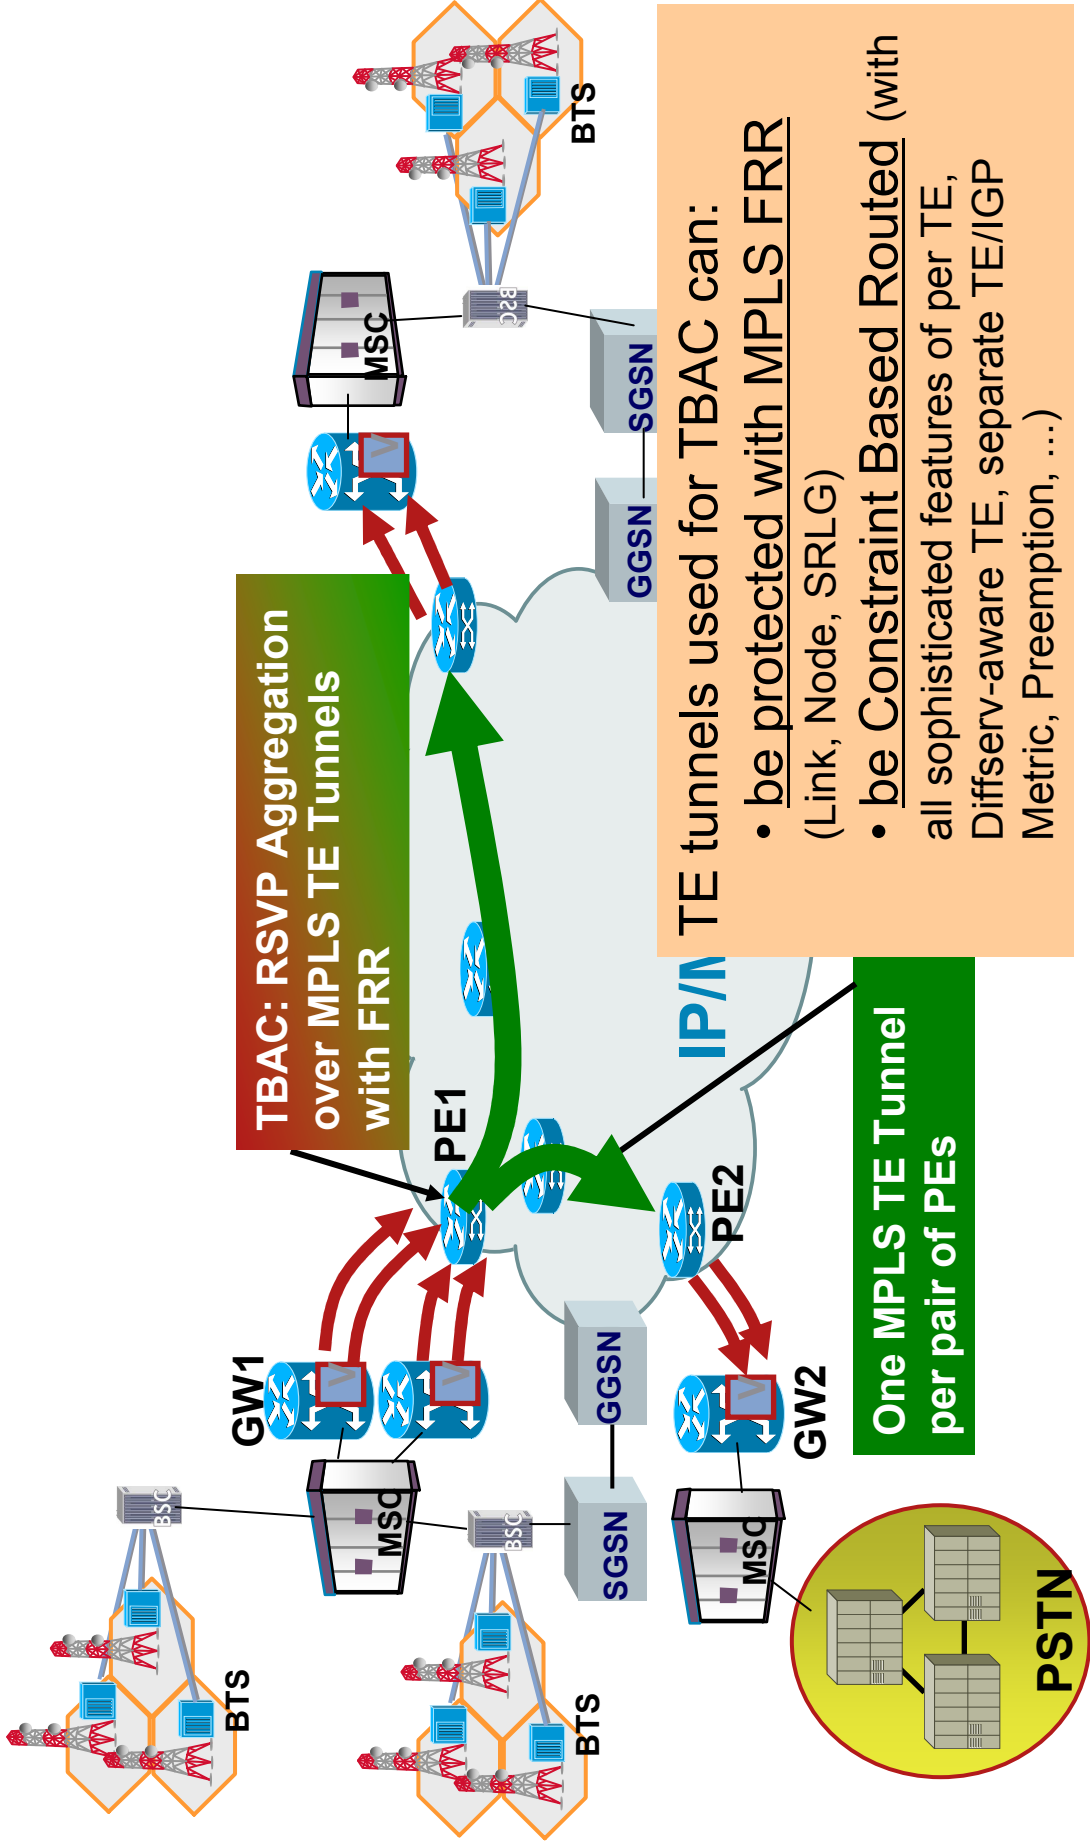


# On-Path CAC with Cisco's TBAC (Big Picture)

## Case Study: Admission Control for Triple-Play VoD



# TBAC can be used in conjunction with all MPLS TE and MPLS FRR features




# Summary

- Some new Hardware
- Lots of new System HA Features
  - L2VPN NSF/SSO
  - N-PE redundancy
- Lots of new Convergence Features
  - PE-CE Local Link Protection
  - LDP Local Label Allocation Filtering
  - OSPFv3 Fast Convergence
- **BROADBAND on SIP400 (ISG)!**
- **L2TPv3 on SIP400**

# Q and A





**cisco**