Carrier Ethernet

Part II. – EVC Overview

Josef Ungerman
Consulting Systems Engineer, CCIE #6167
Cisco Central & Eastern Europe

Agenda

EVC software infrastructure

- what can I do with EVC configuration?
- EVC 1.0
- EVC 2.0

EVC usage in Carrier Ethernet design

- where should I use EVC and ES cards?
- usage for L2 services
- usage for L3 services

EVC 1.0 Infrastructure Overview

Ethernet Virtual Circuit

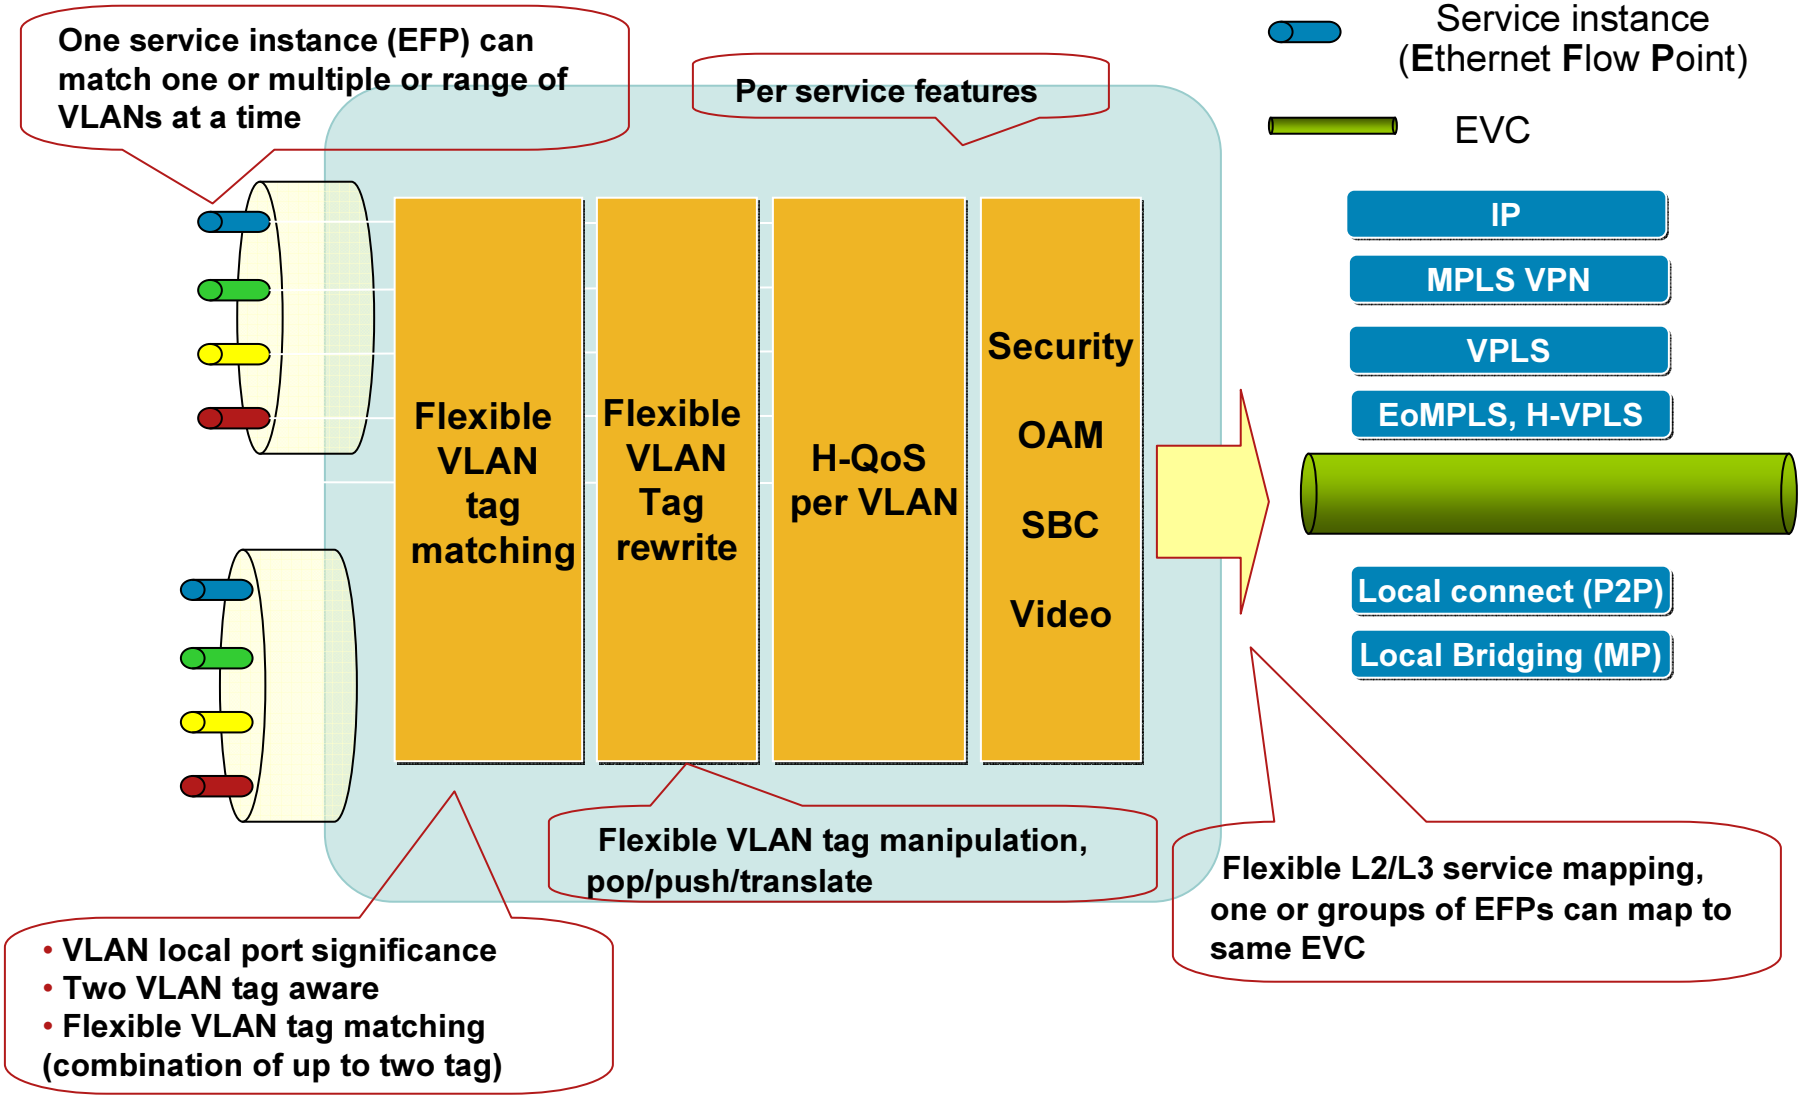

EVC 1.0 – Service-oriented Logic

Traditional router:

interface GigabitEthernet 4/1/0.10

encapsulation dot1q 10 second-dot1q 90

…

Traditional switch:

interface GigabitEthernet 4/1/0

switchport trunk allowed vlan 10

…

EVC infrastructure:

interface GigabitEthernet 4/1/0

service instance 10 ethernet

encapsulation dot1q 10 second-dot1q 10, 90, 155-1600, 1777

rewrite ingress tag pop 1 symmetric

bridge-domain 100 split-horizon

service-policy input POLICE-256

service-policy output SHAPER-QoS-4

service instance 20 ethernet

…

EVC UNI Flexible Frame matching
EFP (Ethernet Flow Points)

- EFP's ...
  - Provide classification of L2 flows on Ethernet interfaces
  - Are also referred to as EVC service-instances
  - Support dot1q and Q-in-Q
  - Support VLAN lists
  - Support VLAN ranges
  - Support VLAN Lists and Ranges combined
  - Coexist with routed subinterfaces
- Longest-match logic
- 8K EFP's (not VLAN's) per SPA
- 32K EFP's per system today
- 120 EFP's per Bridge-Domain

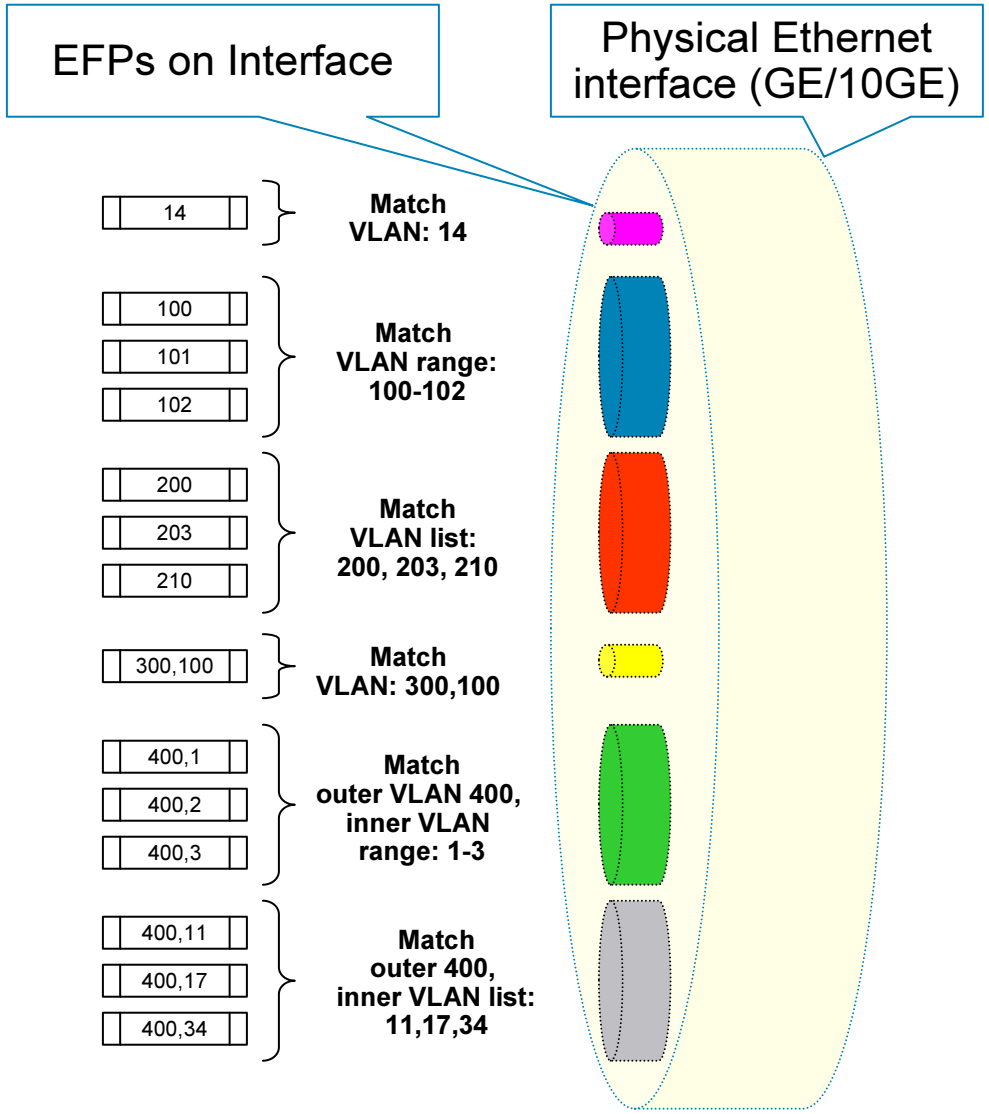

EVC – Encapsulation and Rewrite

NPE1(config-if-srv)#encapsulation ?

default	*catch-all unconfigured encapsulation*
dot1ad	*802.1ad - Provider Bridges*
dot1q	*IEEE 802.1Q Virtual LAN or S-VLAN*
untagged	*Untagged encapsulation*

NPE1(config-if-srv)# encapsulation dot1q 10 second-dot1q {any | 10, 155-160, 2000-4095}

NPE1(config-if-srv)#rewrite ingress tag pop ?

1 Pop the outermost tag ← remove 1 tag

2 Pop two outermost tags ← remove 2 tag

NPE1(config-if-srv)#rewrite ingress tag push dot1q 10 ← add 1 tag

NPE1(config-if-srv)#rewrite ingress tag push dot1q 10 second-dot1q 20 ← add 2 tags

NPE1(config-if-srv)#rewrite ingress tag translate ?

1-to-1 Translate 1-to-1

1-to-2 Translate 1-to-2

2-to-1 Translate 2-to-1

2-to-2 Translate 2-to-2

EVC – Forwarding Actions

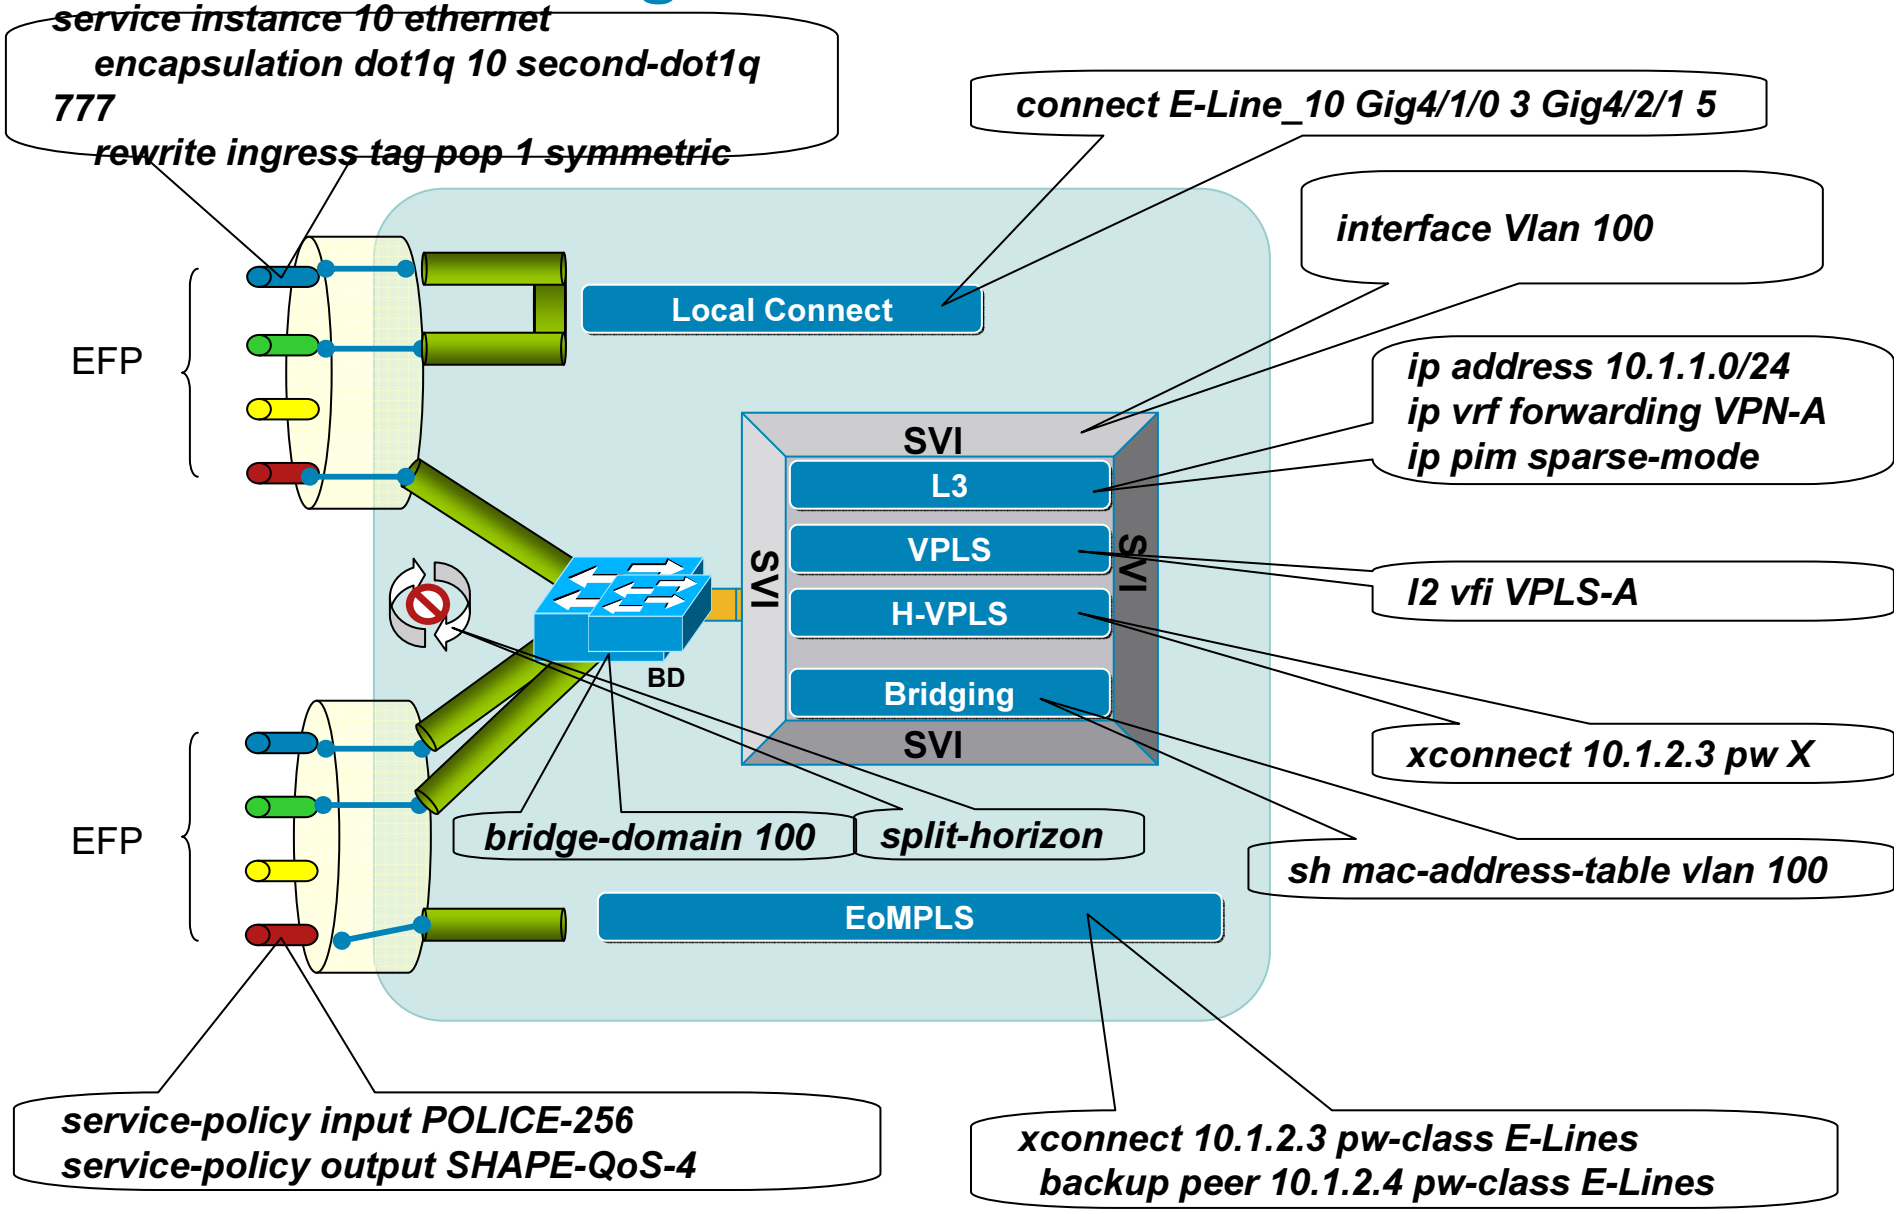

EVC 2.0 – Port/MAC Security

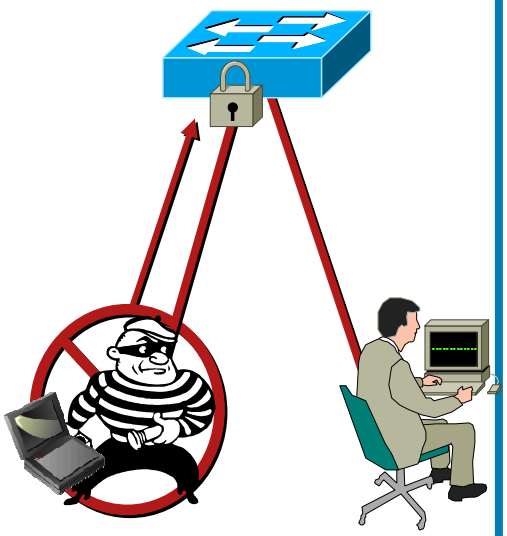

```
 service instance 415 ethernet 415
  encapsulation dot1q 415
  rewrite ingress tag pop 1 symmetric
  bridge-domain 415 split-horizon
  mac security maximum addresses 3
  mac security address permit 0000.0415.0301
  mac security sticky
  mac security violation restrict
  mac security
AGG1-rossi(config-if-srv)# mac security aging ?
  static  Apply aging controls to statically configured addrs also
  sticky  Apply aging controls to persistent ("sticky") addrs also
  time    Configure aging time
```

- Port security works with dynamically learned and static MAC to restrict a port's ingress traffic by limiting the MAC addresses that are allowed to send traffic into the port.
- A security violation occurs in either of these situations:
  - maximum number of secure MAC addresses is reached
  - source MAC address is different from identified secure MAC
  - traffic with a secure MAC address that is configured or learned on one secure port attempts to access another secure port in the same VLAN

EVC 2.0 – DHCP snooping /w Option 82

- Traditional “port + VLAN” information is not enough to identify the subscriber uniquely. It may require access encapsulation VLAN information

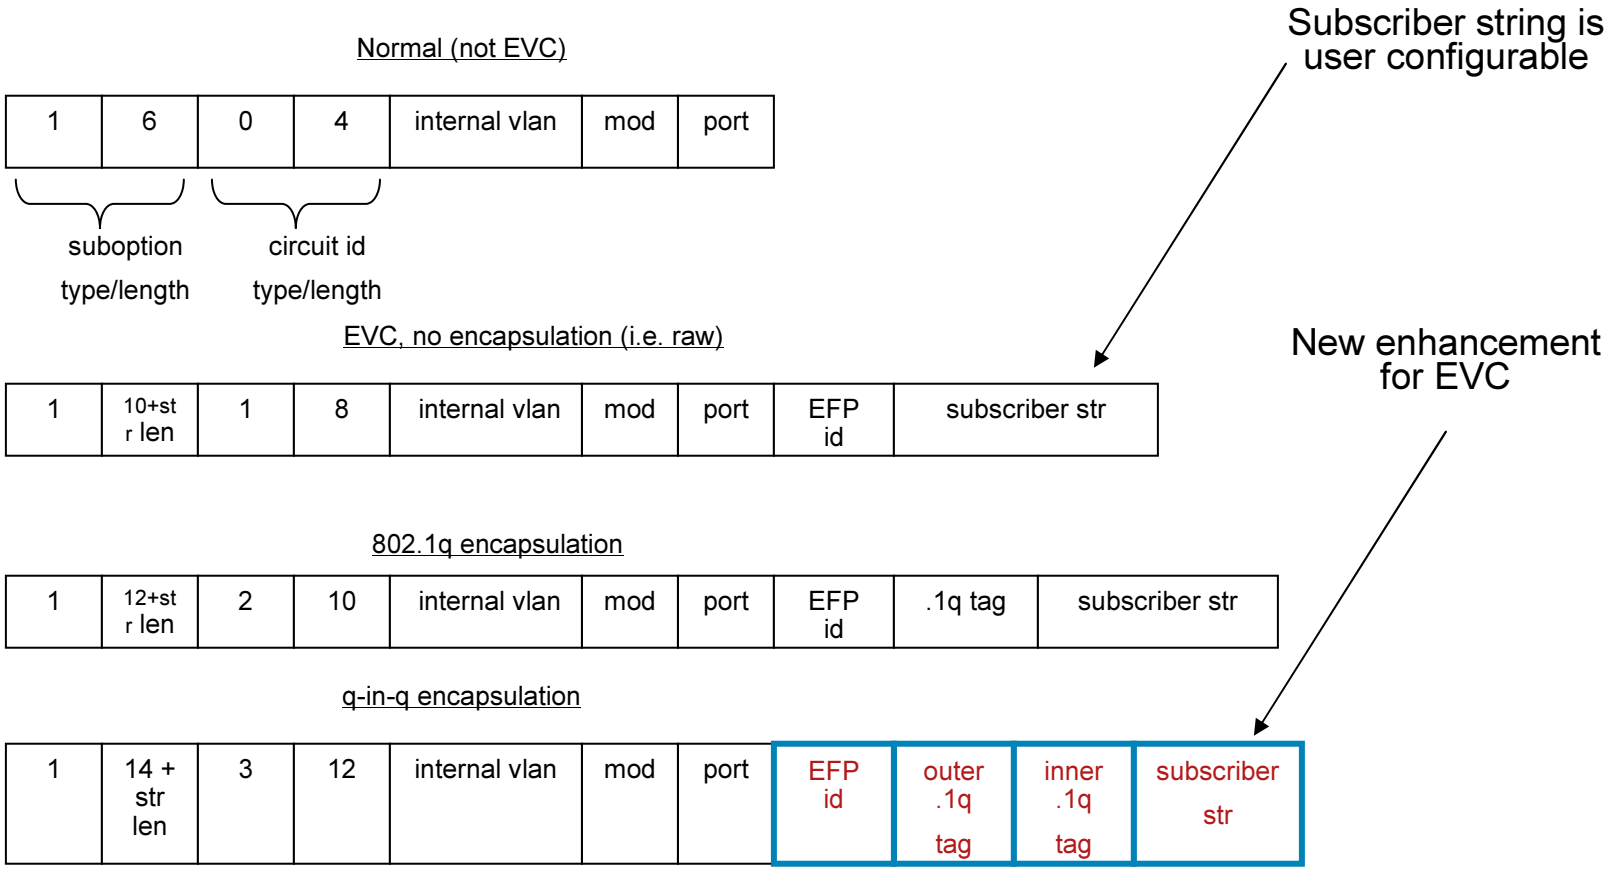

EVC 2.0 – IP Source Guard

- Uses the DHCP Snooping Binding Table Information
- IP Source Guard
  - Operates just like Dynamic ARP Inspection, but looks at every packet, not just ARP Packet

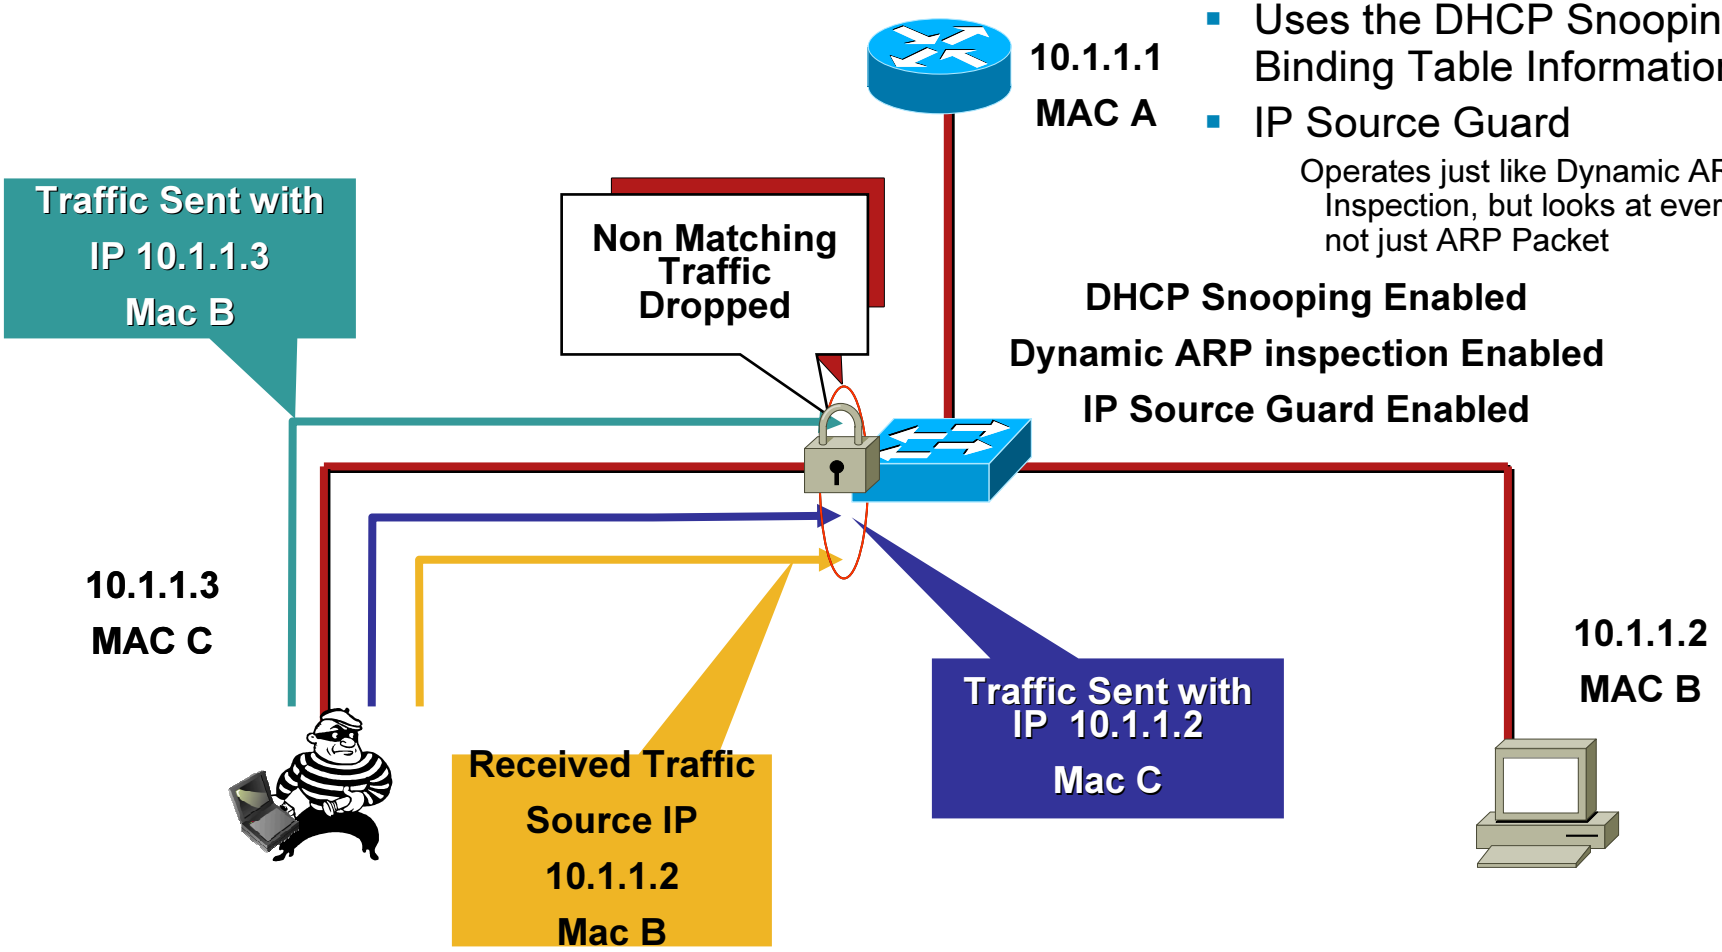


```
AGG1-rossi(config-if-srv)#ip verify source vlan dhcp-snooping
← validate IP only
AGG1-rossi(config-if-srv)#ip verify source vlan dhcp-snooping port-security
← validate both IP and MAC
```

EVC 2.0 – 7600 EVC L2 ACL

- EVC L2 MAC ACL only works for src and/or dst MAC address, it doesn't work for ethertype, VLAN ID, etc
- The ACL counters is per ACL, not per ACE
- Like other IOS ACL, it has implicit "deny any any" at the end of the ACL

```
mac access-list extended mac-415
 permit host 0000.0415.0401 any
 permit host 0000.0415.0401 host 0000.0415.0302

service instance 415 ethernet
  mac access-group mac-415 in



AGG2-duhan#show ethernet service instance id 415 interface gig 2/0/16
detail | inc ACL
L2 ACL (inbound): mac-415
L2 ACL permit count: 189418
L2 ACL deny count: 367339
```

EVC Scalability

- 32k EVCs per system (independent of service – local connect, eompls xconnect, bridge-domain/vpls)
- 16K EVCs per line card (ES20, ES40, SIP-400)
- 8K EVCs per port (4K with QoS)

- Maximum EVCs (service instances) per Bridge-domain
  - –120 EVCs per bridge-domain per SIP-400
  - –220 EVCs per bridge-domain for ES20, 440 EVCs for ES40

- EVC 2.0
  - –E-OAM (CFM including Y.1731 and IP SLA): 9K local + 18K remote MEP's
  - –DHCP snooping and IPSG: 64K
  - –Secure MAC: 32K per system, 10K per BD
  - –64 MST instances
  - –2K ACL's

EVC enabled platforms

TODAY:

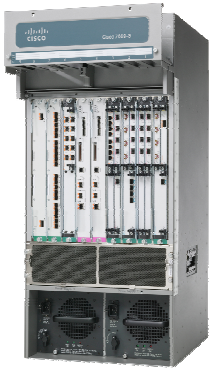

Cisco 7600
- hardware: ES20, SIP-400, ES20+, ES40
- software: IOS 12.2SR

Q1 CY09:

Catalyst ME4500
- hardware: all modules w/ Metro Sup
- software: IOS 12.2XO

Cisco ASR 9000
- hardware: all modules
- software: IOS XR 3.7

Cisco 7600: EVC enabled linecards

TODAY:

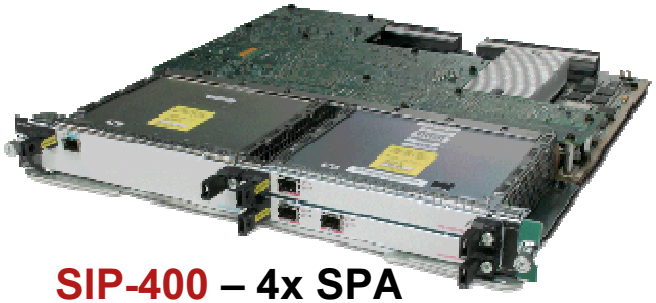

SIP-400 – 4x SPA

ES20 – 2-port 10GE

ES20 – 20-port GE

Q1 CY09:

ES20+

H1/H2 CY09:

IPoDWDM w/ G.709: ES40-based – 2x10GE and 4x10GE

Combo Cards: 10GE + 20xGE, 2x10GE + 20xGE

7600: Carrier Ethernet Module Comparison

	7600-ES+xx	7600-ES20-xx	7600-SIP-400	WS-X67xx
DFC	DFC-3C or 3CXL	DFC-3C or 3CXL	no	Optional 3B/C or 3BXL/CXL
Fabric connection	2 * 20G	2 * 20G	1 * 20G (max. 4 Gbps)	1 or 2 * 20G
Forwarding Rates	Up to 48Mpps	Up to 30Mpps	Up to 6Mpps	Up to 48Mpps with DFC
Interface Density	4*10G/40*1G/2*10G/20*1G	2*10G / 20*1G	1*10G / 10*1G	24*1G/48*1G/4*10G
Interface Support	Ethernet or IPoDWDM	Ethernet	Ethernet, POS, ATM, Serial	Ethernet
Scheduler	CBWFQ	CBWFQ	CBWFQ	DWRR
Ingress Policing	Supported	Supported	Supported	Supported
Egress Policing	Supported	Not Supported	Supported	Supported
Number of Queues	256k for ES40 128k for ES20+	16K	32K	4 or 8 per port
Shaping	Supported	Supported	Supported	Not Supported
HQoS	Supported (4 levels)	Supported (3 levels)	Supported (3 levels)	Not Supported
LLQ	Supported	Supported	Supported	Not Supported (Strict Priority instead)
VPLS uplink	Supported	Supported	Supported	Not Supported
VPLS with MPLS edge	Supported	Supported	Supported	Not Supported
Local vlan significance	Supported	Supported (EVC only)	Supported	
SyncE capability	Yes	No	Yes (SPA)	No

EVC Infrastructure Features
7600 Linecard Comparison

Feature	ES40/ES20+	ES20	SIP-400
DHCP snooping w/ Option 82 insertion and DAI (Dynamic ARP Inspection) on EVC	12.2SRD	12.2SRD	12.2SRD
Port/MAC security on EVC	12.2SRD	12.2SRD	Future
UDLD on all ports w EVC	12.2SRD	12.2SRD	Future
MST on EVC for BD	12.2SRD	12.2SRD	Future
L2 ACLs on EVC	12.2SRD	12.2SRD	Future
Ingress 2R3C Policer with EVC	12.2SRD	12.2SRD	
IP Source Guard on EVC	12.2SRD	Future	Future
REP integration with EVC	Next Release	Future	Future
L3/L4 ACL on EVC	Next Release	Next Release	Future
Broadcast storm control on port with EVC	Next Release	Next Release	Future
Port-channel w/ LACP	Next Release	12.2SRD	Future
LAG core facing	Next Release	Future	Future
802.1ah with EVC	Next Release	Future	Future

E-OAM Features
7600 Linecard Comparison

Feature	ES40/ES20+	ES20	SIP-400
CFM1.0	12.2SRD	12.2SRA	12.2SRA
Link OAM (link and remote loopback)	12.2SRD	12.2SRA	12.2SRA
E-LMI	12.2SRD	12.2SRA	12.2SRA
Link OAM to CFM Interworking	12.2SRD	12.2SRB	12.2SRB
CFM to E-LMInterworking	12.2SRD	12.2SRB	12.2SRB
MPLS PW OAM to E-LMI Interworking	12.2SRD	12.2SRB	12.2SRB
IP SLA for Metro Ethernet	12.2SRD	12.2SRB	12.2SRB
Y.1731 (AIS, RDI)	12.2SRD	12.2SRD	12.2SRD
E-OAM HA (E-LMI, 802.1ag, 802.3ah)	12.2SRD	12.2SRD	12.2SRD
802.1ag CFM MIB	12.2SRD	12.2SRD	12.2SRD
802.3ah MIB	12.2SRD	12.2SRD	12.2SRD
EVC 2.0 MIB	12.2SRD	12.2SRD	12.2SRD
CFM1.0 in EVC mode (BD)	12.2SRD	12.2SRD	Future

7600: 12.2(33)SRD features

Infra / Common	Carrier Ethernet	Carrier Ethernet	Video / Mobile/BB
RSP720-10GE service module support CWDM/DWDM on ES BX Optics (GE SPAand ES) FE (4x, 8xFE) SPA on SIP400 SPA OBFL Q1 CY09: ES40/20+ (2x10, 20x1, 4x10, 40x1), optics supported XFP, SFP-GE-S,Z,L,T	2R3C policers (ingress) on EVC Y.1731 Outward Facing MEP on switchport IP SLAs Metro-Ethernet EVC-1.0 Broadcast storm control on EVC port DHCP snooping w/ Option 82 insertion and DAI on EVC CFM in EVC mode (BD) MST on EVC BD EVC 2.0 MIB BRR (with port channel) on ES UDLD on all ports with EVC QoS parity on service instance encap default L2TPv3 P2P on Excalibur	IP Source Guard on EVC (Excalibur is priority) BRE on SIP-400 E-OAM HA (E-LMI, 802.1ag, 802.3ah) 802.1ag CFM MIB 802.3ah MIB Routed Mode AToM interworking with SIP-400, SIP-600, ES uplink Flexible Service Mapping - Custom Ethertype Private Host - SVI E-OAM on ES20/67xx in switch port mode with PVRSTP Port/MAC security on ES EVC L2 ACLs on ES EVC LACP for Portchannel	Multicast Fast Convergence enhancements Netflow MIB BGP Fast Convergence ph1 IMA Core facing support AToM/cell relay Port Mode support on CEoP and ATM SPA DHCP option 60 + 82 SIP-400 HQF PQ enhancements RSVP-ledge Ph0 SAMI Driver Update End to End VCCV on MS-PW* PW Status TLV * VP QoS on Warlord * MS PW SPE * PW Redundancy w/PW Status TLV *

Agenda

EVC software infrastructure

- what can I do with EVC configuration?
- EVC 1.0
- EVC 2.0

EVC usage in Carrier Ethernet design

- where should I use EVC and ES cards?
- usage for L2 services
- usage for L3 services

Base Design: Carrier Ethernet Network

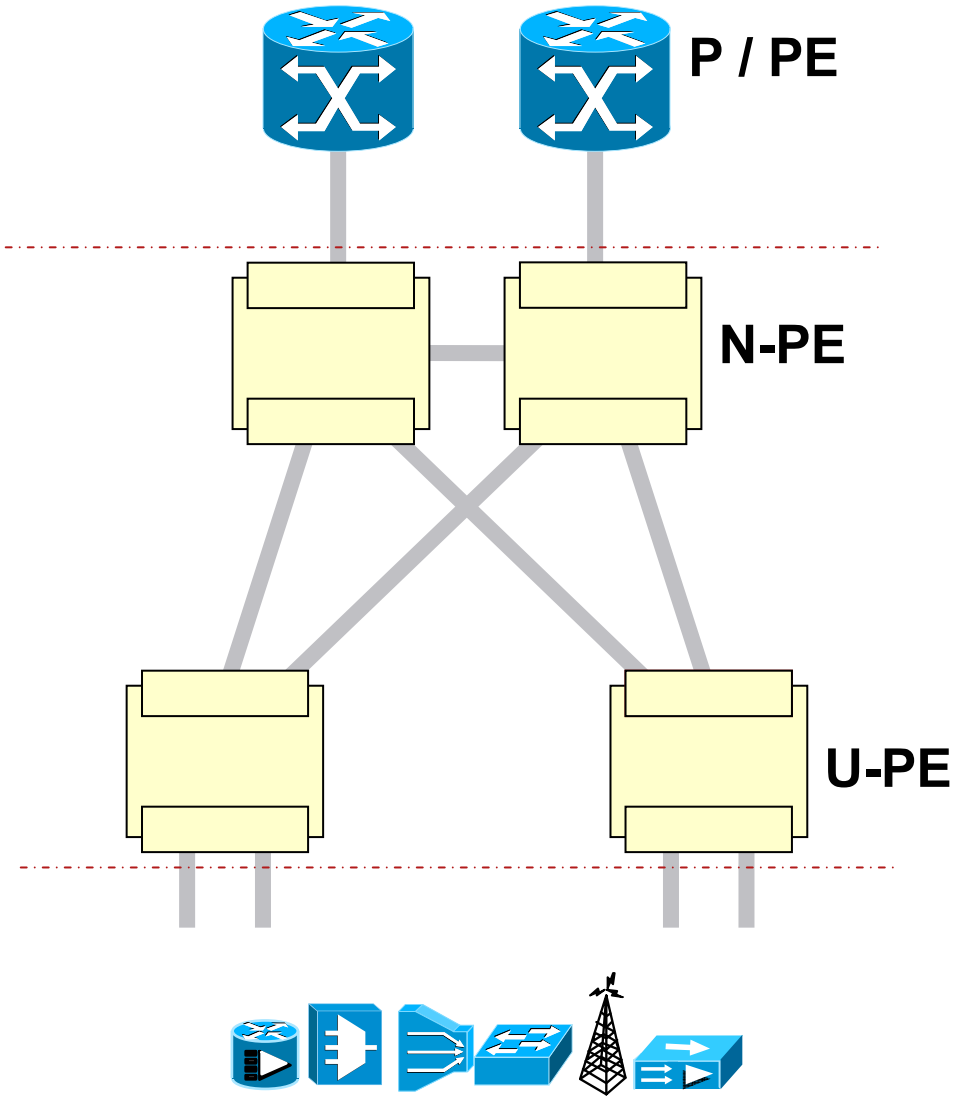

LAN Card

ES20 or SIP

Puzzle #1: Edge-facing

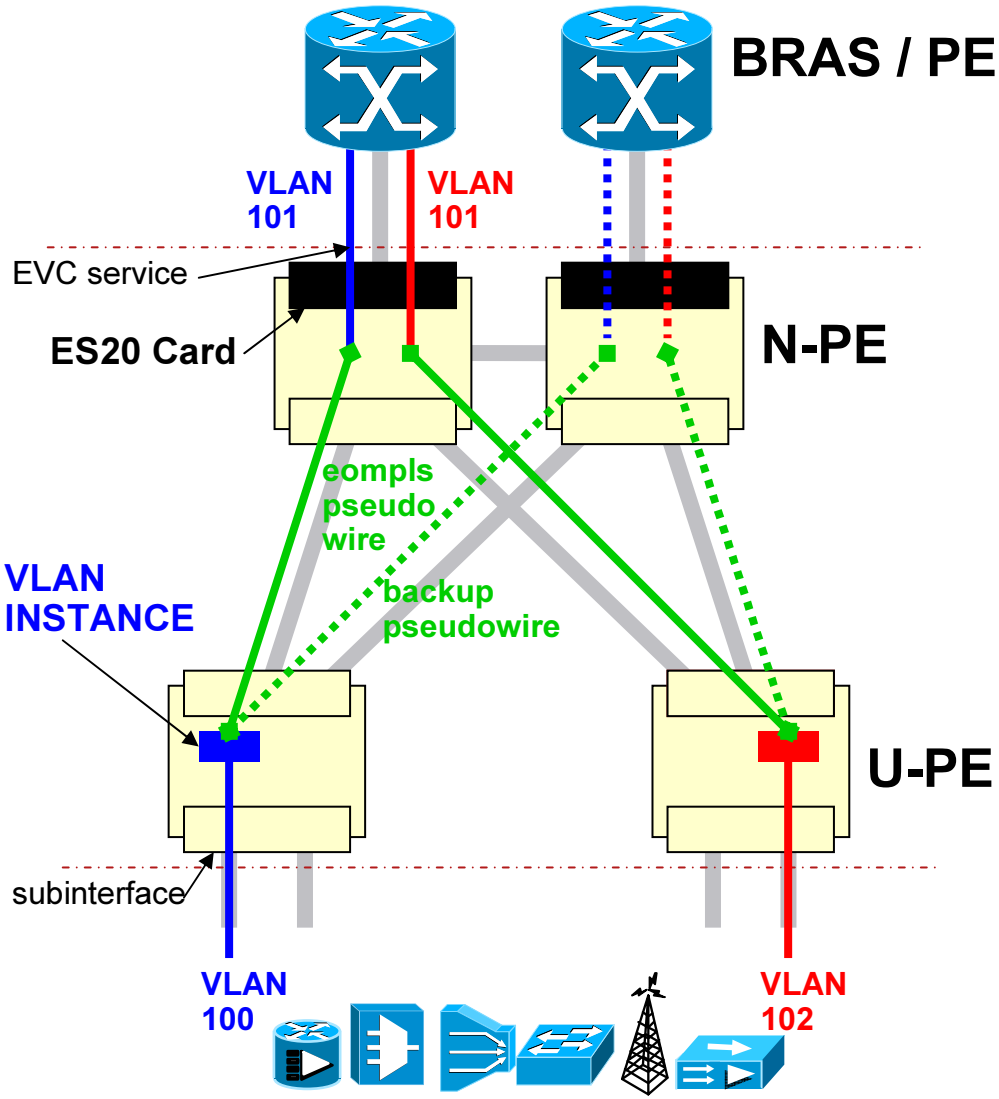

N-PE Scalability
- # of VLANs: 32K
- # of MACs: unlimited

U-PE Scalability
- # of VLANs: 4K
- # of MACs: unlimited

Advantages at N-PE
- PW Scalability (no global VLAN instance is consumed)
- EVC and Flexible QinQ
- Per-Port VLAN significance
- Flexible VLAN translations
- Flexible QinQ CoS mapping

Puzzle #2: Customer-facing

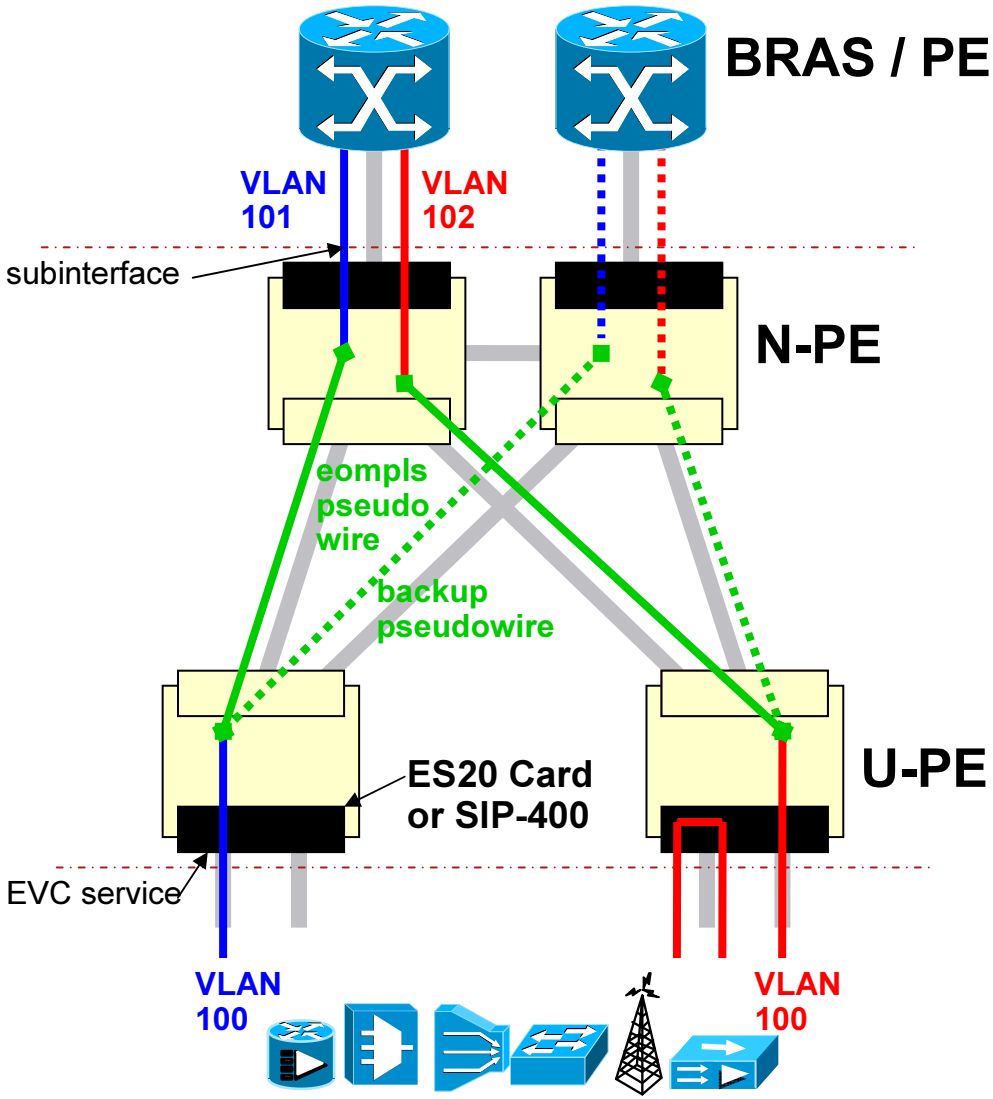

N-PE Scalability
- # of VLANs: 32K
- # of MACs: unlimited

U-PE Scalability
- # of VLANs: 32K
- # of MACs: unlimited

Services Edge:
- E-Line, E-LAN

Advantages at U-PE
- PW Scalability
- EVC and Flexible QinQ
- Per-port VLAN significance
- Flexible VLAN translations
- Flexible QinQ CoS mapping
- HQoS for L2 VPN Customers
- 2 priority Queues (voice/video)
- ELMI for L2 VPN Customers

Puzzle #3: U-PE uplink

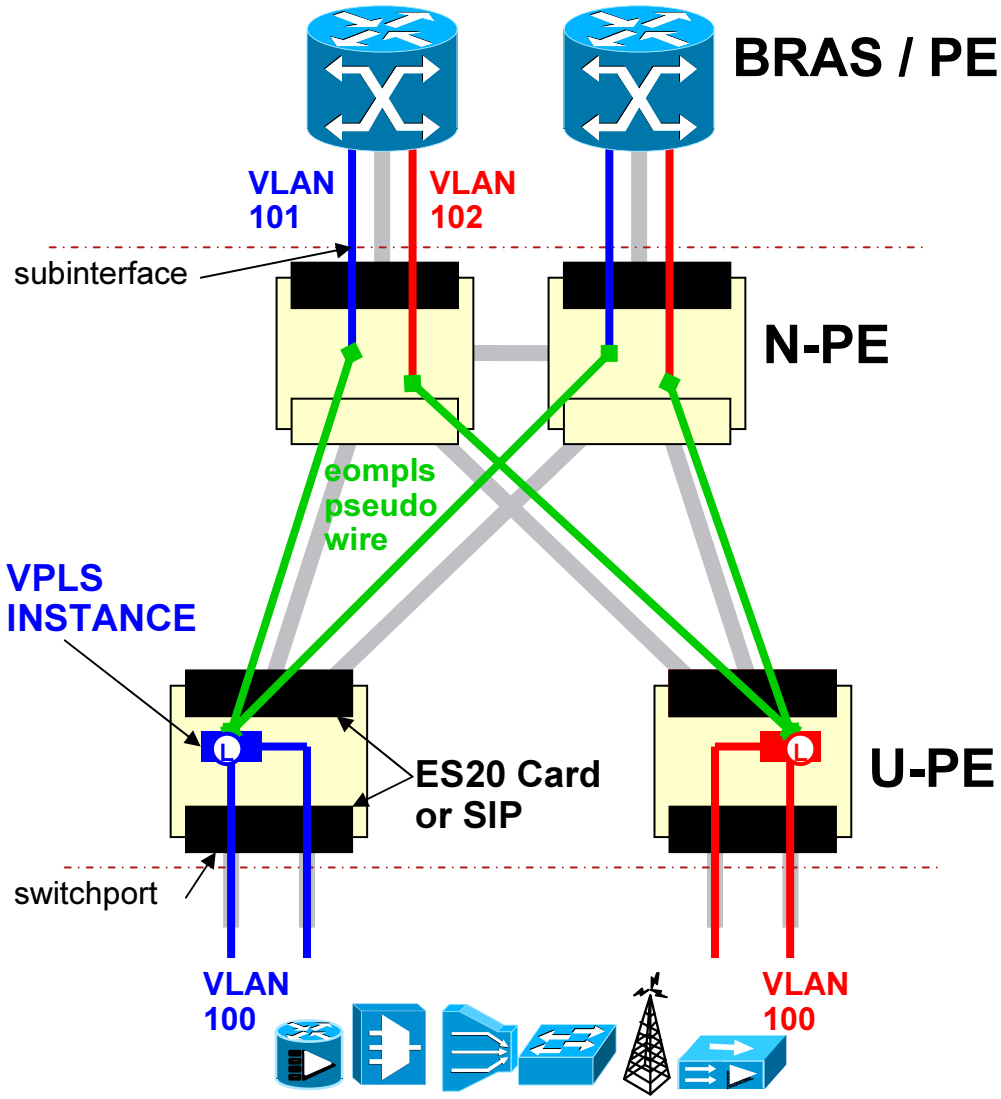

N-PE Scalability
- # of VLANs: 32K
- # of MACs: unlimited

U-PE Scalability
- # of VLANs: 32K
- # of MACs: Nx 80K

Other Limits
- 4K VPLS instances

Advantages at U-PE
- Local Switching with Learning
- Concurrent STP and MPLS
- Active-Active Pseudowires

Puzzle #4: N-PE downlink

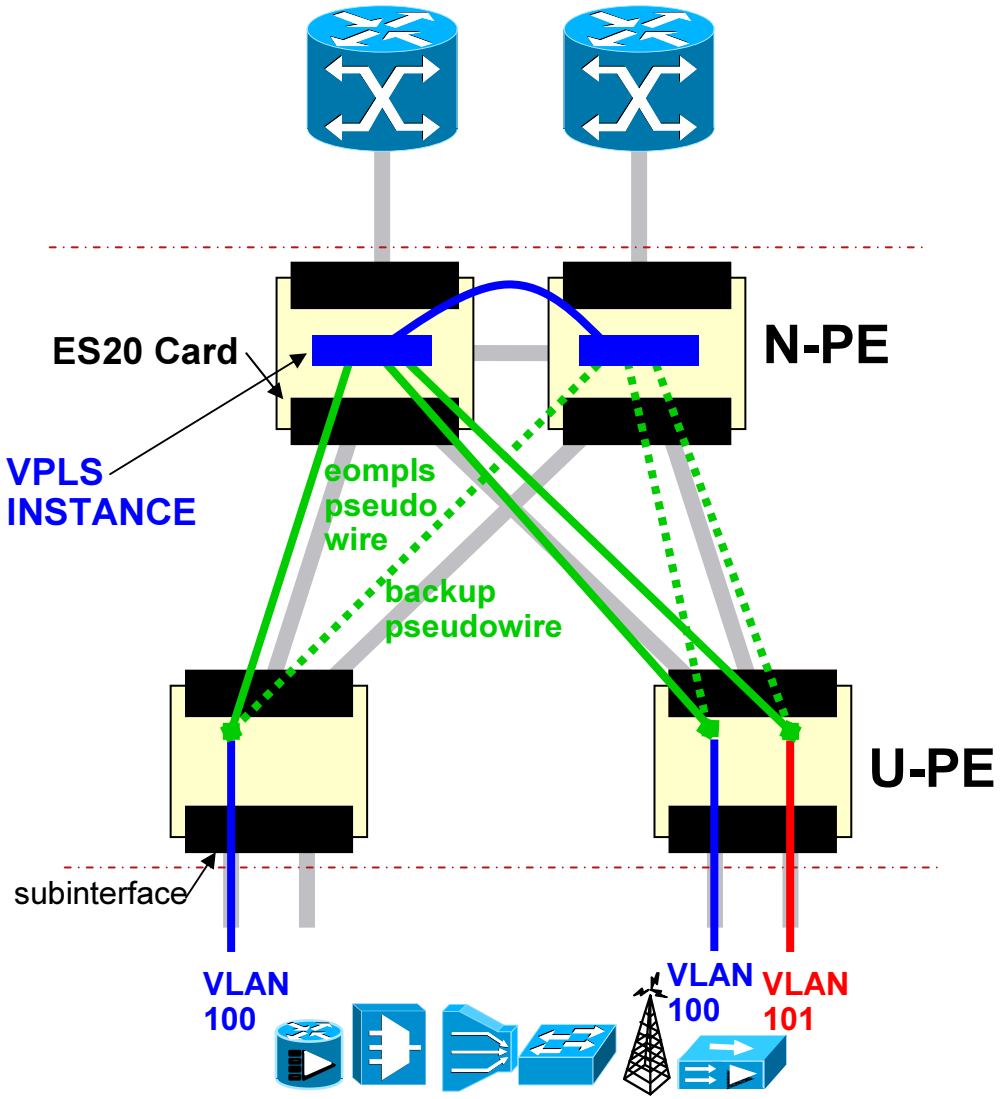

N-PE Scalability
- # of VLANs: 4K
- # of MACs: Nx 80K

U-PE Scalability
- # of VLANs: 32K
- # of MACs: unlimited

Services Edge:
- popular E-LAN services (N-PE is just a H-VPLS hub)

Puzzle #5: Vidmon

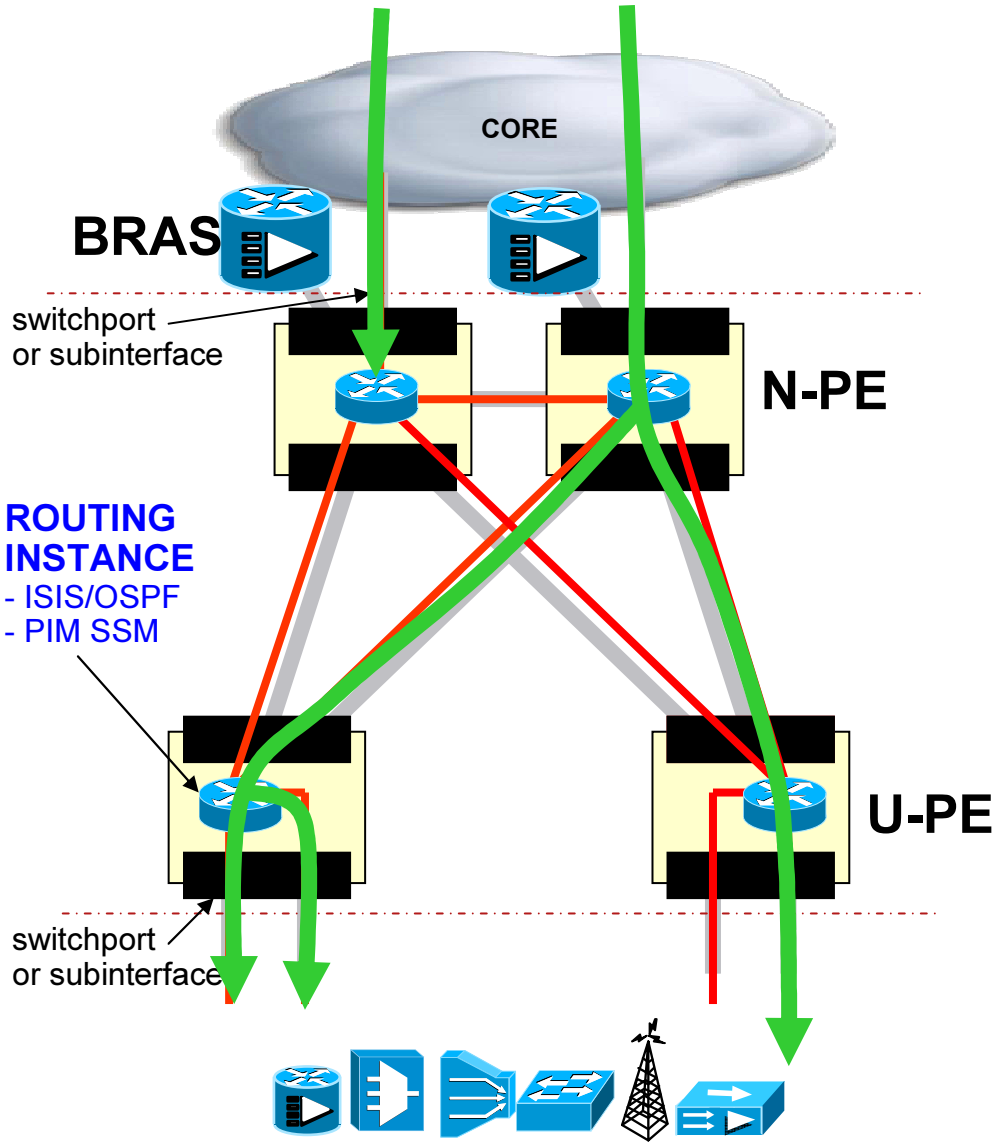

Excellent Replication Capabilities
- **20Gbps of multicast per LC**

ES20+ and ES40+
Next Generation Carrier Ethernet Hardware

- **Video Monitoring Capabilities**
- **MDI (Media Delivery Index)**
- **Delay, Jitter, Loss**
- **Per STREAM statistics**
- **Rapid problem isolation**

Puzzle #6: IP Multicast over Routed PW

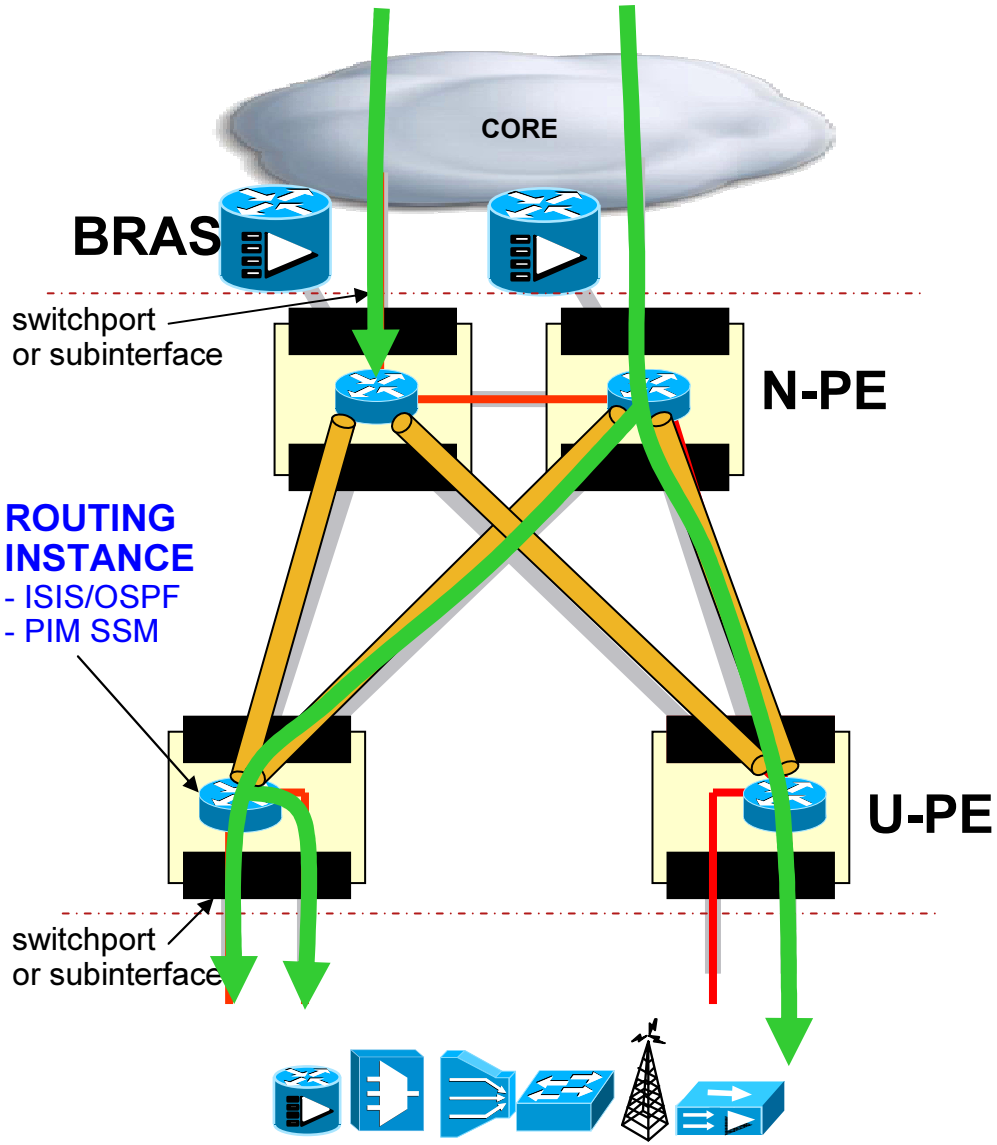

Excellent Replication Capabilities
- **20Gbps of multicast per LC**

IPTV over MPLS Routed PW

- gives you MPLS TE FRR capability
- ES cards needed for any Routed PW
- adds protection (a lot of complexity)

Be warned ;-)

- 50ms for IPTV has no much sense
- Screen is impacted for 1-2s anyway
- FRR doesn't lift TV quality, VQE does!

Puzzle #7: Distributed BRAS and ISG

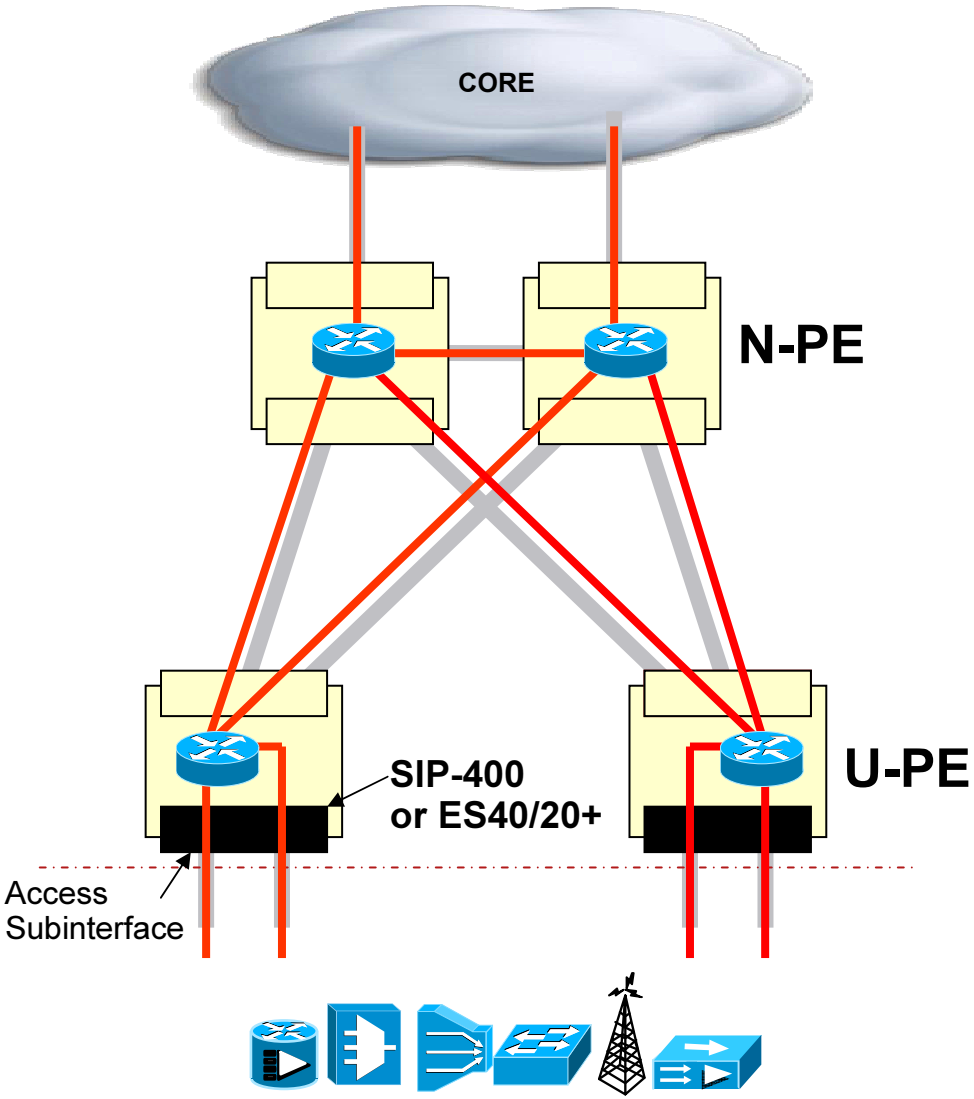

U-PE Scalability
- # of L3 interfaces: 32K
- # of PPPoE or IP sessions: 32K

(per 10GE: 8K VLANs, 8K sessions)

Services Edge:
- integrated residential BRAS
- integrated AAA or SCE
- Video Edge features

Advantages at U-PE
- AAA
- PPPoE sessions
- IP/DHCP sessions
- Change of Authorization
- L4 redirect
- SCE integration bus
- Wholesale models

EVC – Conclusion

1. Flexible service mapping
   - service instance
   - local VLAN significance
   - selective QinQ, push/pop, translations
   - xconnect, bridge-domain, SVI

2. EVC features
   - QoS – similar to what we know from router subinterfaces
   - Security – similar to what we know from switchports
   - MST for EVC
   - E-OAM for EVC

3. EVC roadmap – ES20+ and ES40
   - IPoDWDM
   - ISG
   - Vidmon
   - 802.1ah