

What You Make Possible



Troubleshooting Catalyst 4K and 6K

BRKRST-3067

Agenda

Troubleshooting Catalyst 4500

- Architecture
- Interface Packet Drops
- High CPU
- IOS and IOS-XE Crashes
- IOS-XE Licenses



Catalyst 4500-E Series Switches

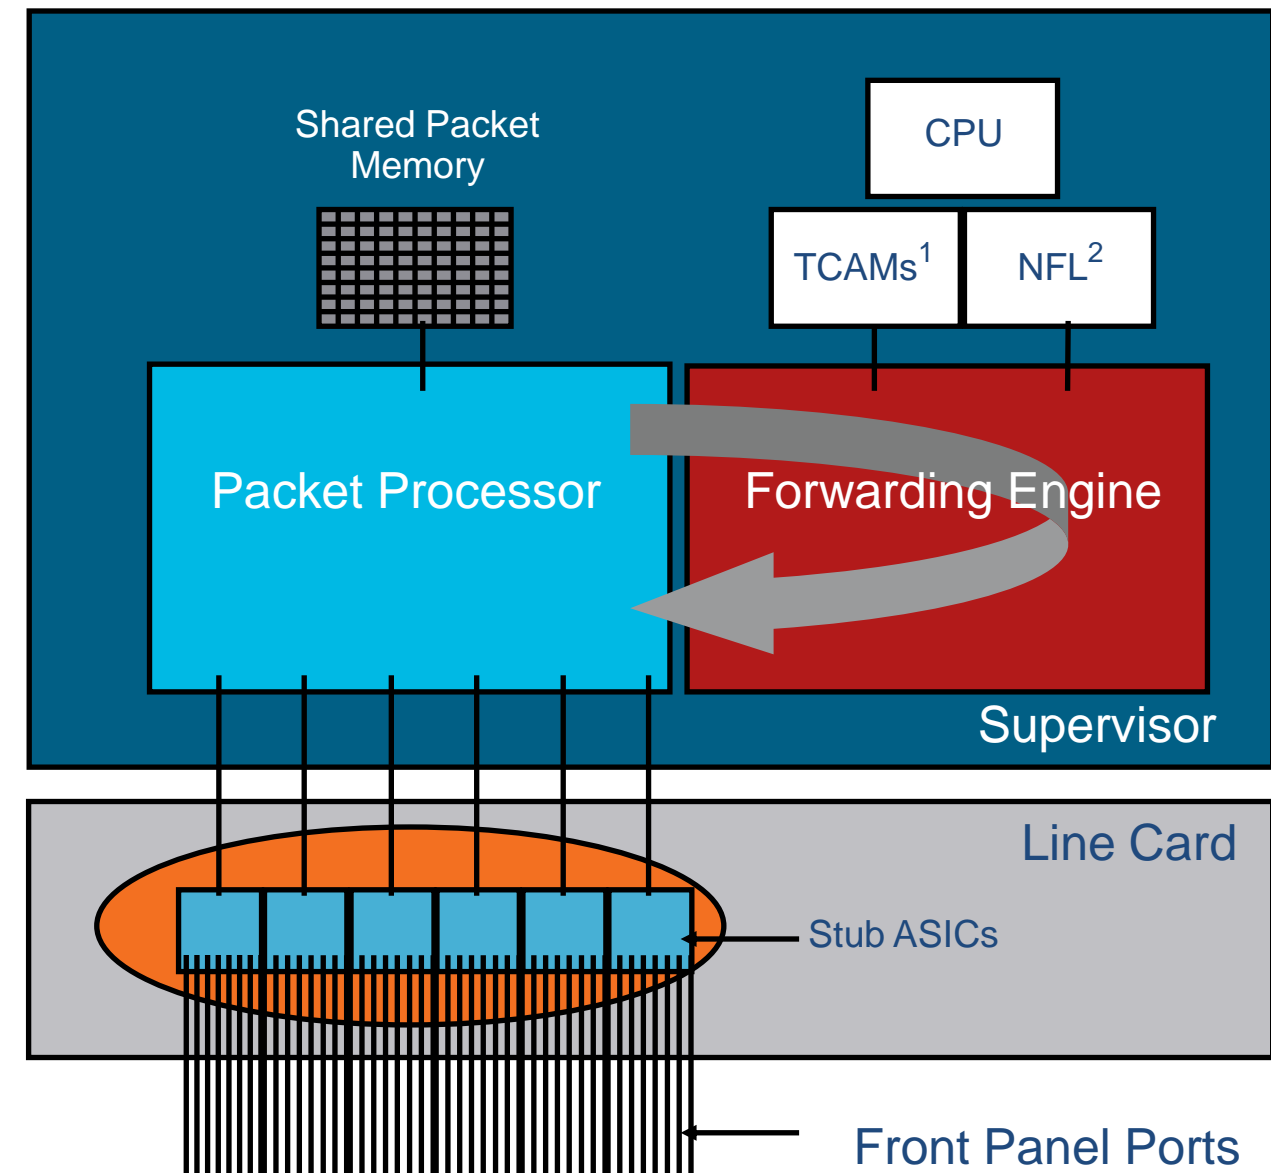


Catalyst 4500 Series Switches



Architecture Overview


- **Intelligent Supervisors**
 - Supervisor Engine 7-E, 6-E, 6L-E, V-10GE, V, IV,
 - II-Plus-10GE, II-Plus, II-Plus
- **Switching ASICs**
 - Packet Processor
 - Forwarding Engine
- **Specialised Hardware**
 - TCAMs for ACLs, QoS, L3 forwarding
 - NetFlow (NFL) for statistics gathering
- **Transparent Line Cards**
 - Wire-rate, oversubscribed, PoE
 - 10/100, 10/100/1000, GE, 10GE



Supervisor and Blocking Line Card Block Diagram



Next Generation Supervisors

- Supervisor 6L-E
- Supervisor 6-E
- Supervisor 7-E



Troubleshooting Interface Drops

- Packet drops in transmit direction
 - TX queue is full
 - Oversubscription
 - Pause frames received
- Packet drops in receive direction
 - Minimal buffer on receiving queue



Troubleshooting Interface Drops

TX Queue is Full

```
SUP6-E# show interfaces g2/47 counters detail | begin Drops
Port      Tx-Drops-Queue-1  Tx-Drops-Queue-2  Tx-Drops-Queue-3  Tx-Drops-Queue-4
Gi2/47    0                 0                 0                 0
Port      Tx-Drops-Queue-5  Tx-Drops-Queue-6  Tx-Drops-Queue-7  Tx-Drops-Queue-8
Gi2/47    0                 0                 0                 37748571

SUP6-E# show interfaces g2/47 counters detail | begin Drops
Port      Tx-Drops-Queue-1  Tx-Drops-Queue-2  Tx-Drops-Queue-3  Tx-Drops-Queue-4
Gi2/47
Port      Tx-Drops-Queue-5  Tx-Drops-Queue-6  Tx-Drops-Queue-7  Tx-Drops-Queue-8
Gi2/47    0                 0                 0                 37874327
```

Default queue 8 when QoS is disabled

■ Next steps

- Consider implementing a QoS policy to share traffic amongst queues
- Increment default output queue-limit of 40 packets to allow for more buffer space

Troubleshooting Interface Drops

Peer is Sending Pause Frames

```
SUP6-E# show interfaces g2/47 counters detail | begin RxPause
Port      Rx-No-Pkt-Buff  RxPauseFrames  TxPauseFrames  PauseFramesDrop
Gi2/47    0                130            0              0

SUP6-E# show platform software interface g2/47 all | inc Busy
Switch Phyport Gi2/7 BusyStatus : ON, PauseStatus : OFF
BusyBitActivityCount      : 474
Ignore Subport Busy       : False
```

BusyBit is activated when received pause frames exceed threshold.

```
%C4K_HWPORTRMAN-4-BLOCKEDTXQUEUE: Blocked transmit queue HwTxQId7 on Switch Phyport Gi2/47,
count=116
```

Error log displayed when BusyBit is set

Troubleshooting Interface Drops

Peer is Sending Pause Frames Continued

- Next steps
 - Disable flow control receive on the local interface
 - Disable flow control send on the peer

Troubleshooting Interface Drops

Minimal buffer on receiving stub ASIC (applicable to linecards with oversubscription ratio)

■ Next steps

- Enable flow control send on the local interface
- Enable flow control receive on the peer

```
Sup6-E# show interface gi1/13 | include overrun
0 input errors, 0 CRC, 0 frame, 86432 overrun, 0 ignored
Sup6-E# show interface gi1/13 | include overrun
0 input errors, 0 CRC, 0 frame, 206658 overrun, 0 ignored
```

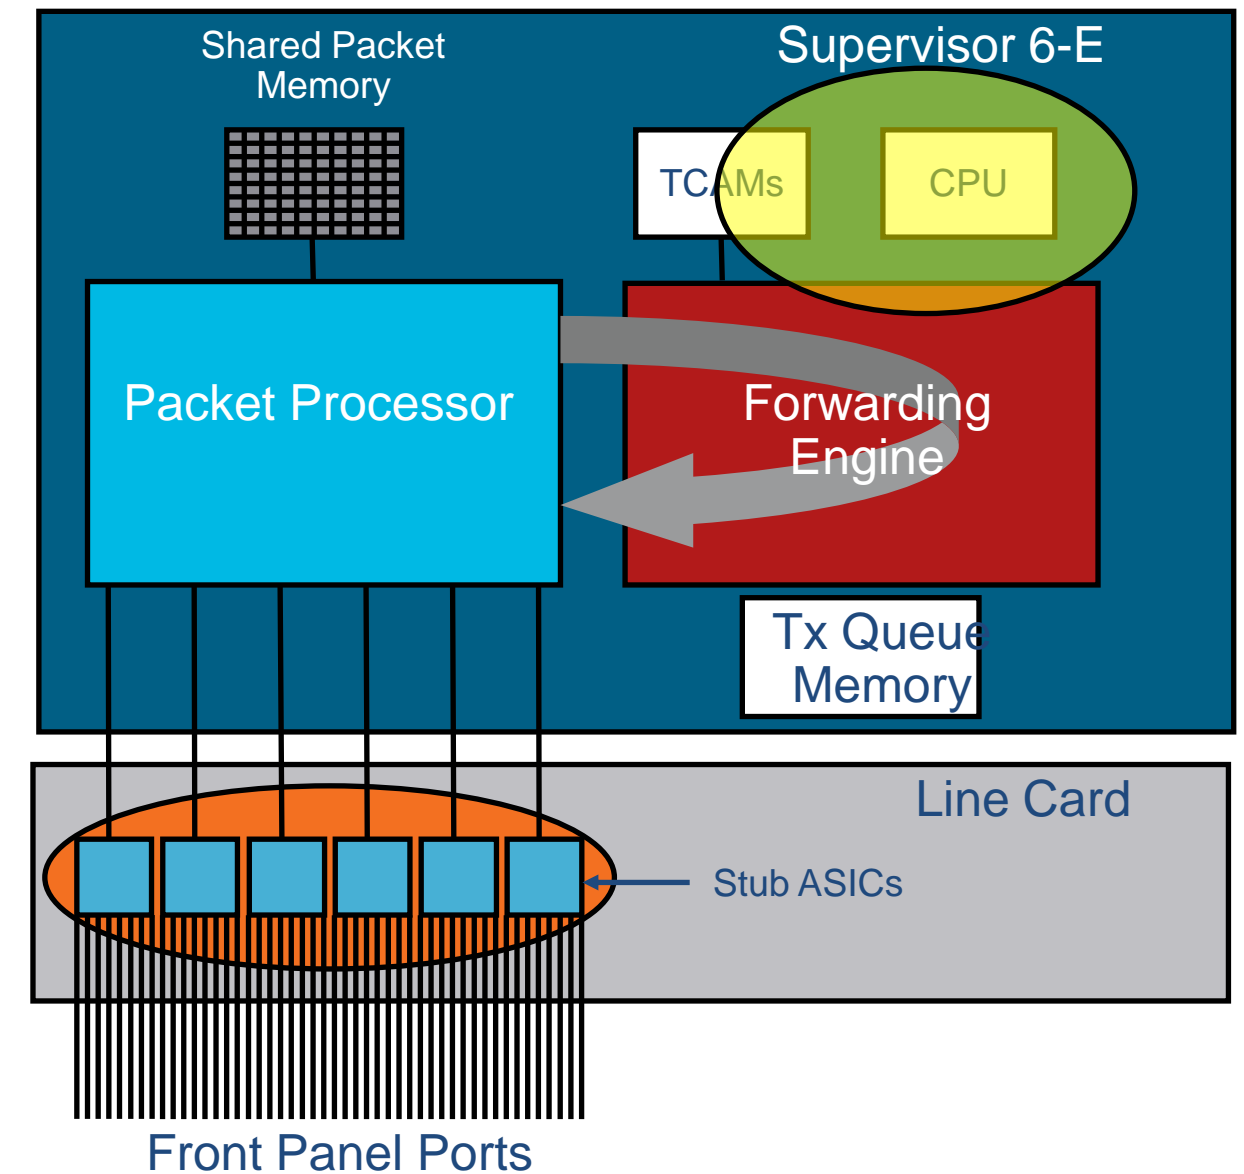
```
Sup6-E# show interface gi1/13 counter all | begin Rx-No
Port      Rx-No-Pkt-Buff  RxPauseFrames  TxPauseFrames  PauseFramesDrop
Gi1/13    206658          0               0               0
```

```
Sup6-E# show platform software interface g1/13 stub stat | in Overrun
OverrunPackets      : 206658    (look for Rx Stats)
```

Troubleshooting High CPU

Functions of the 4500 CPU

- Sends and receives control plane traffic
 - STP, CDP, PAgP, VTP, DTP, routing protocols
- Program dynamic entries into hardware
 - ACLs, CEF entries
- Manage access to the switch
 - Telnet, SSH
- Manage system components
 - Fan tray, power supply, PoE



Troubleshooting High CPU

Software-forwarded Data Traffic

- Packets copied to CPU but originally switched in hardware
 - Host mac address learning
- Packets punted to CPU for processing
 - Routing updates, BPDUS, flood of traffic
- Packets sent to the CPU for forwarding
 - AppleTalk, IPX

Troubleshooting High CPU

Usage on IOS Threads

Dual core

Process
breakdown

```
SUP7-E# show processes cpu sorted detail
```

```
Core 0: CPU utilization for five seconds: 4%; one minute: 2%; five minutes: 2%
```

```
Core 1: CPU utilization for five seconds: 6%; one minute: 3%; five minutes: 2%
```

PID	T	C	TID	Runtime(ms)	Invoked	uSecs	5Sec (%)	1Min (%)	5Min (%)	TTY	Process
9433	L			2946931	1131416	0	6.70605	A 4.39062	4.12207	0	iosd
9433	L	1	11383	984896	4669930	0	7.11	A 3.22	3.00	0	iosd
9433	L	1	9433	1961205	6644042	0	6.22	A 5.44	5.11	0	iosd
9433	L	1	11386	829	18630	0	0.00	A 0.00	0.00	0	iosd
71	I			55575	8787502	0	1.11	R 1.00	1.00	0	Cat4k Mgmt HiPri
52	I			4221576	2152734	0	0.33	R 0.33	0.33	0	IDB Work
72	I			1033445	1988579	0	0.33	R 0.33	0.33	0	Cat4k Mgmt LoPri
89	I			4	132	0	0.00	R 0.00	0.00	0	Exec

```
SUP6-E#sh proc cpu sorted
```

```
CPU utilization for five seconds: 6%/0%; one minute: 6%; five minutes: 6%
```

PID	Runtime(ms)	Invoked	uSecs	5Sec	1Min	5Min	TTY	Process
51	12294972	8770348	1401	3.75%	3.78%	3.79%	0	Cat4k Mgmt LoPri
50	2556152	16464011	155	2.47%	2.49%	2.47%	0	Cat4k Mgmt HiPri
92	0	38032	0	0.07%	0.00%	0.00%	0	Ethchnl
111	20	142	140	0.07%	0.00%	0.00%	0	Exec
38	2044	380316	5	0.07%	0.06%	0.07%	0	IDB Work
99	44600	106342	419	0.07%	0.04%	0.05%	0	CDP Protocol

Troubleshooting High CPU

Usage on platform dependent threads

- These processes use CPU under the **Cat4K Mgmt HiPri** and **Cat4K Mgmt LoPri**
- **HiPri:** within CPU target **LoPri:** exceeds CPU target

```
SUP6-E#show platform health
```

	%CPU Target	%CPU Actual	RunTimeMax Target	RunTimeMax Actual	Priority Fg	Priority Bg	Average 5Sec	Average Min	Average Hour	Total CPU
RkiosObflMan	0.50	0.00	4	0	100	500	0	0	0	0:15
GalChassisVp-review	3.00	0.09	10	33	100	500	0	0	0	5:34
S2w-JobEventSchedule	10.00	0.00	10	0	100	500	0	0	0	0:00
Stub-JobEventSchedul	10.00	0.50	10	5	100	500	0	0	0	13:58
Lj-poll	1.00	0.01	2	0	100	500	0	0	0	1:18
StatValueMan Update	1.00	0.01	1	0	100	500	0	0	0	2:18
Pim-review	0.10	0.00	1	0	100	500	0	0	0	0:18
Ebm-host-review	1.00	0.00	8	0	100	500	0	0	0	0:05
Ebm-host-util-review	1.00	0.00	10	0	100	500	0	0	0	0:00
Ebm-port-review	0.10	0.00	1	0	100	500	0	0	0	0:01
Protocol-aging-revie	0.20	0.00	2	0	100	500	0	0	0	0:00
EbmHostRedundancyMan	2.00	0.00	20	0	100	500	0	0	0	0:00
Acl-Flattener	1.00	0.00	10	0	100	500	0	0	0	0:00
IrmFibThrottler Thro	2.00	0.00	7	0	100	500	0	0	0	0:26



Troubleshooting High CPU

Traffic in the CPU Queue

- Events that trigger packets to be sent to the CPU queue

```
SUP7-E# show platform cpu packet statistics
Packets Dropped In Processing Overall

Total                5 sec avg 1 min avg 5 min avg 1 hour avg
-----
                36802848                0                0                0                0

Packets Dropped In Processing by CPU event

Event                Total                5 sec avg 1 min avg 5 min avg 1 hour avg
-----
Sa Miss                36778332                0                0                0                0
Input Acl Fwd                19                0                0                0                0
Input ACl Copy                24497                0                0                0                0
```

- Next steps

Mac flap?

– Identify the event that is dropping packets so that we know what type of packet to look for in further CPU debugging

Troubleshooting High CPU - Tools

In-Built CPU Sniffer

- Packets that are punted to CPU are sent to a circular buffer
- Buffer retains 1024 packets at a time and does not occupy CPU cycles

```
SUP6-E# debug platform packet all buffer
platform packet debugging is on
SUP6-E# show platform cpu packet buffered
Total Received Packets Buffered: 1024
-----
Index 0:
3 days 23:23:18:54927 - RxVlan: 1006, RxPort: Gi1/1
Priority: Normal, Tag: No Tag, Event: 11, Flags: 0x40, Size: 64
Eth: Src 00:00:0B:00:00:00 Dst 00:22:90:E0:D6:FF Type/Len 0x0800
Ip: ver:IpVersion4 len:24 tos:0 totLen:46 id:0 fragOffset:0 ttl:64 proto:tcp
    src: 10.10.10.100 dst: 172.16.100.100 hasIpOptions firstFragment lastFragment
Remaining data:
 0: 0x0  0x64 0x0  0x64 0x0  0x0  0x0  0x0  0x0  0x0
10: 0x0  0x0  0x50 0x0  0x0  0x0  0x8A 0x37 0x0  0x0
20: 0x0  0x1  0xB5 0x77 0x6A 0x7E
```

Ingress interface of packet

Source/destination address

Troubleshooting High CPU - Tools

SPAN Capture for CPU

- Monitor CPU bound traffic both in send and receive direction

```
SUP6-E# show running-config | include monitor  
  
monitor session 1 source cpu rx  
monitor session 1 destination interface Gi1/48
```

```
SUP6-E# show monitor session 1
```

```
Session 1
```

```
-----
```

```
Type                : Local Session  
Source Ports        :  
  RX Only           : CPU  
Destination Ports   : Gi1/48  
Encapsulation       : Native  
  Ingress           : Disabled  
  Learning          : Disabled
```

Connect laptop with Wireshark

- Next steps

– Is there a pattern and is this legitimate traffic?

IOS Process Crash

- Reasons for a crash
 - Bus error, forced by software, address errors, watchdog timeouts
- All crashes generate a crashinfo
 - Present in crashinfo or platform crashdump

```
----- show platform crashdump -----  
Last powerfail: 03/09/2006 02:10:24  
Current time: 02/19/2007 16:09:59  
Last reload status: 00008800 038D0000  
Last crash: 02/19/2007 15:44:40  
Build: 12.2(31)SG ENTSERVICES  
buildversion addr: 12288034  
pc=112088A4 lr=11208824 msr=20029030 vector=00000300  
cr=20000042 ctr=11207418 xer=00000000  
r0=8000FBCE r1=13FFE680 r2=0000C000 r3=13FFE6C8 r4=13FFE620 r5=00000002 r6=00000000 r7=000000FF  
r8=FFFFFFFF r9=00000000 r10=00000002 r11=00000008
```

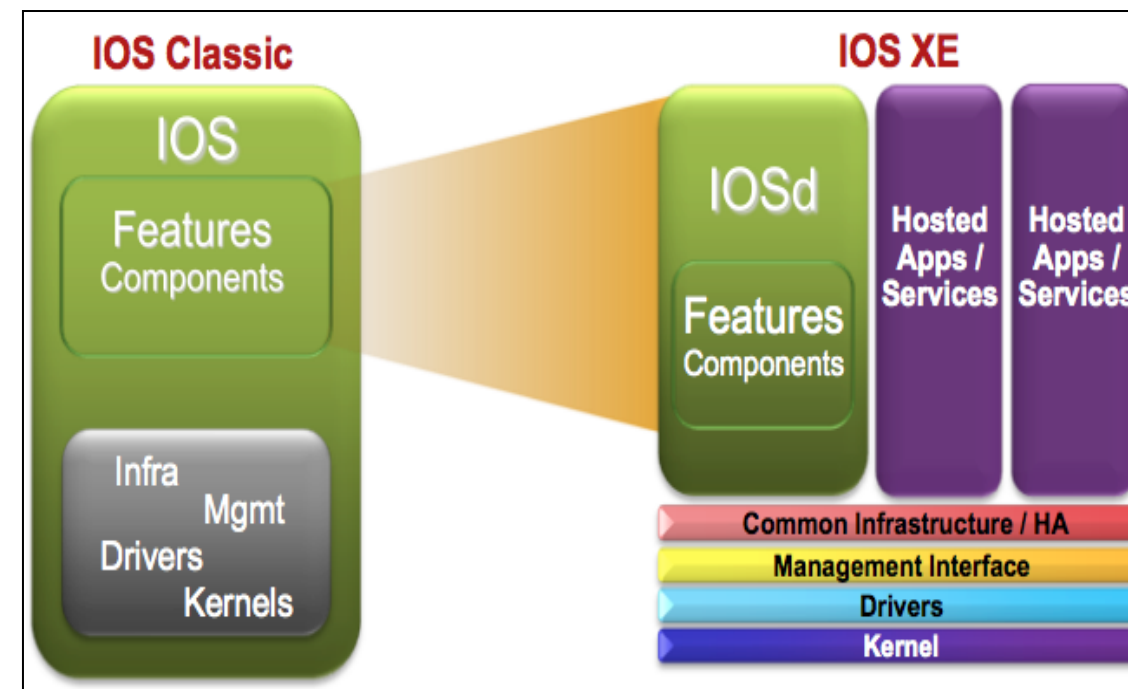
Verify time of outage

Tracebacks that need to be decoded

IOS-XE Process Crash

IOS-XE vs. IOS

- IOS-XE kernel is Linux
- Runs several different processes
 - IOSd, FFM, HA, licensing
- IOS runs as one process in a single unprotected memory space
- IOS-XE crash infrastructure collects crashinfo files for any failing process



IOS-XE Process Crash

- Files generated upon a crash: system kernel info, crashinfo file and core file
- By default, crashinfo files are stored in the crashinfo directory (135MB partition of bootflash)

Enables generation of process core dump file

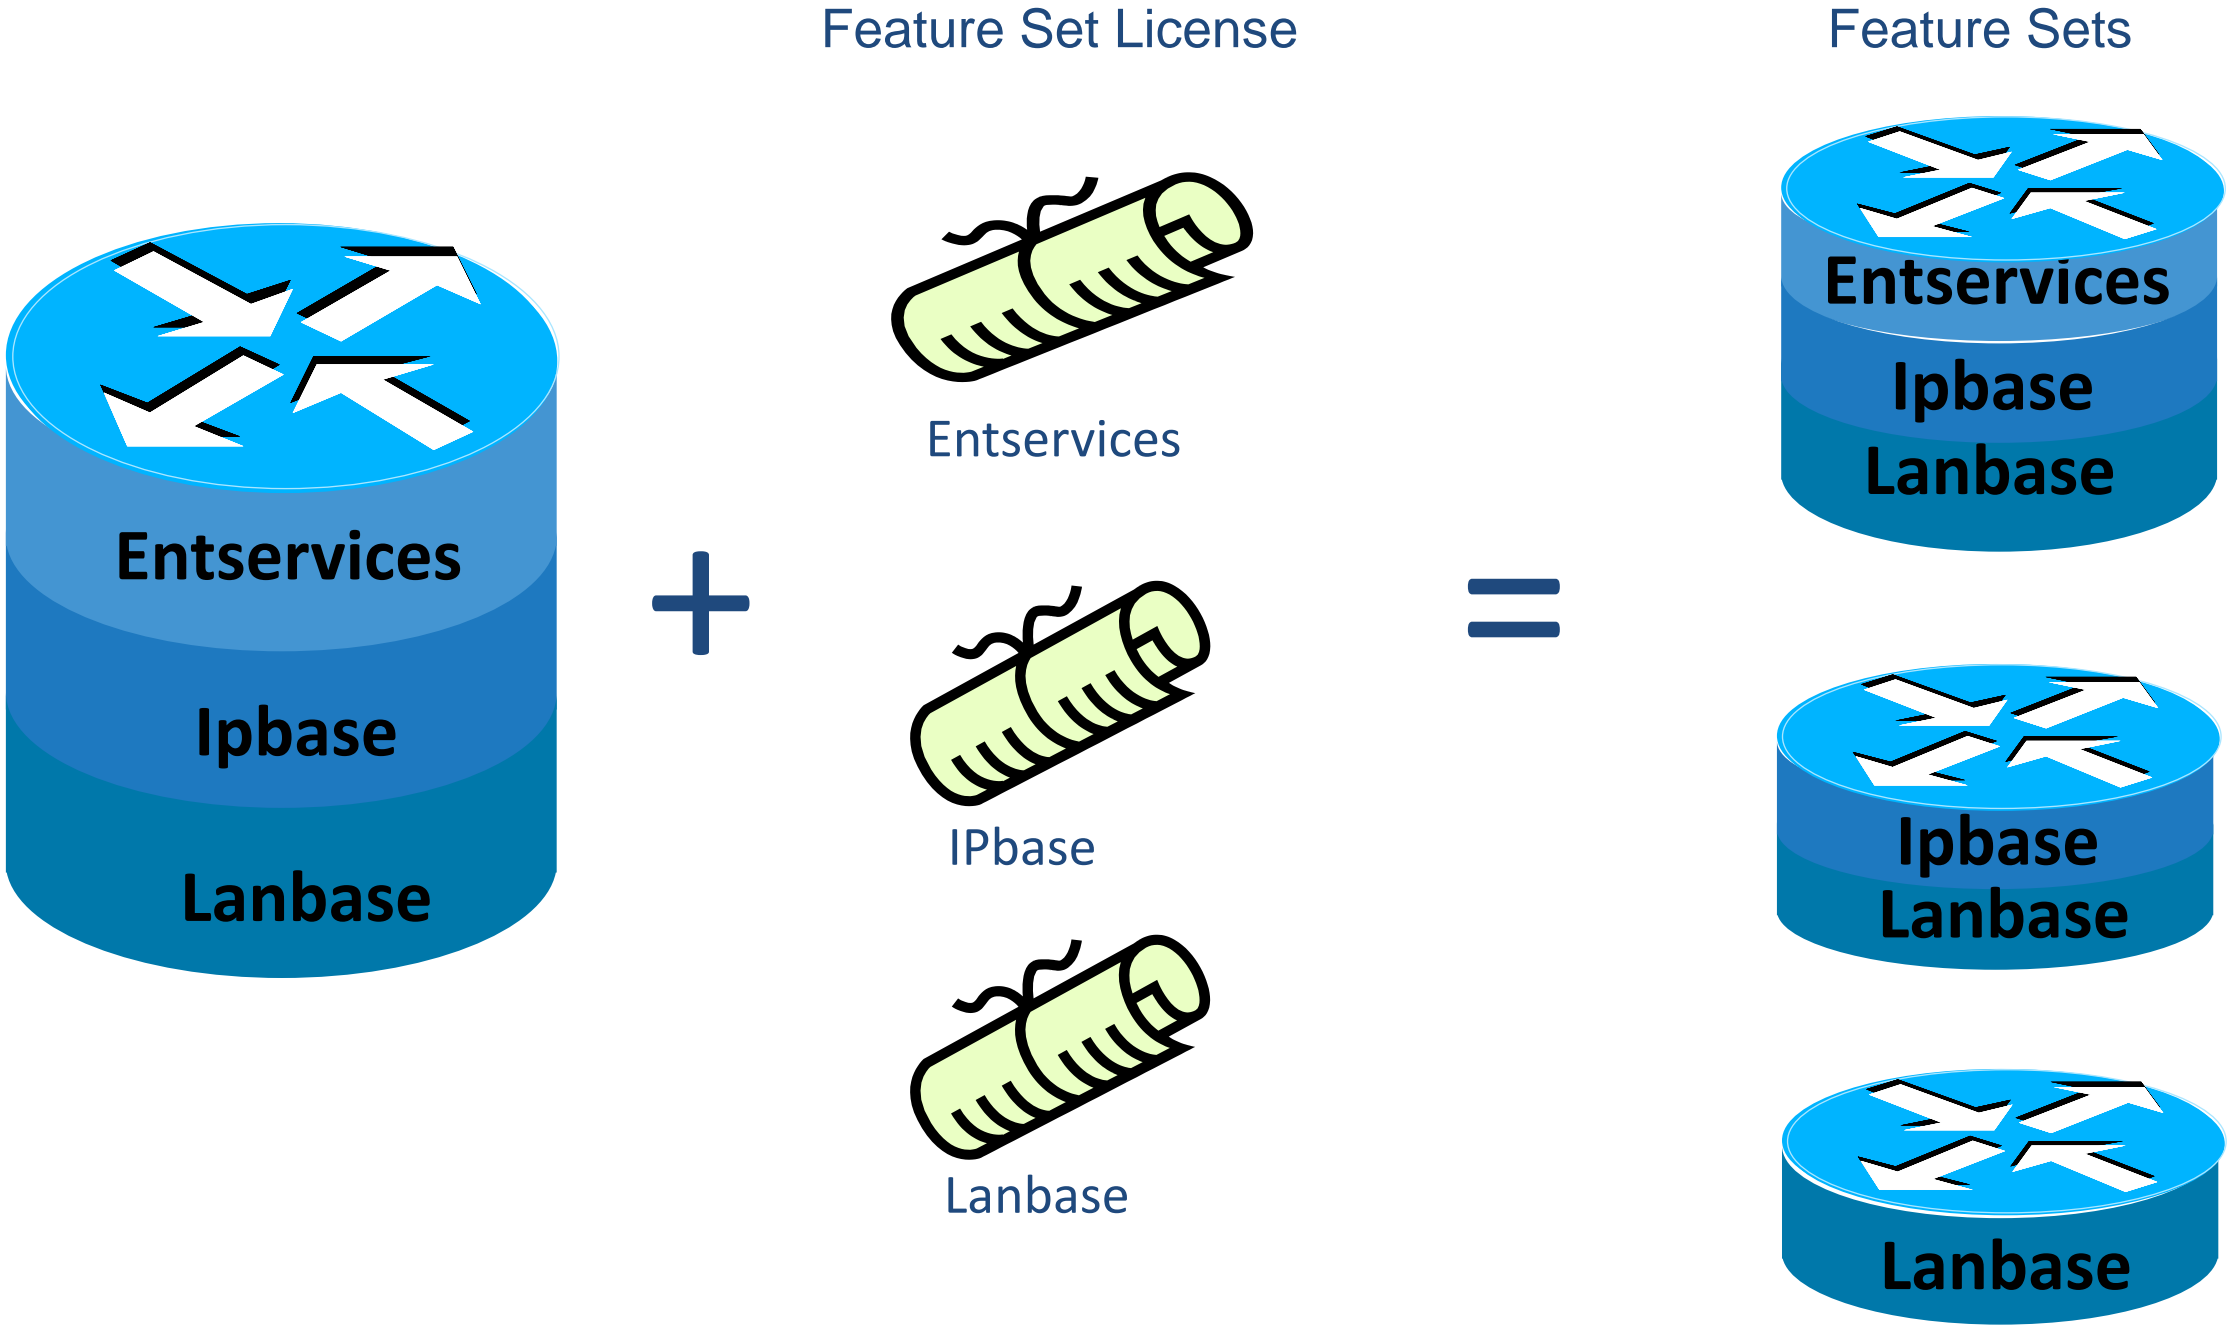
```
Sup7-E(config)# exception coredump
```

- Generation of core files requires configuration

```
SUP7-E#show exception information  
Exception configuration information  
Coredump file - disabled,compressed  
Maximum number of files  
Core - 10 file(s)  
Process crashinfo - 10 file(s)  
Configured storage devices  
1 - crashinfo:  
2 - not assigned  
3 - not assigned  
Dump protocol - not configured (not supported)
```

Default – 10 crashinfo files and 10 core files saved

IOS-XE Licensing (CSL)



IOS-XE Licensing (CSL)

- Permanent license is node-locked
- License is for a chassis UDI (Universal Device Identifier), but stored on Supervisor bootflash
- License synced to hot standby supervisor
- No Product Activation Key (PAK) is generated for customer

License Show Commands

```
Sup7e# show license all
License Store: Primary License Storage1
StoreIndex: 0 Feature: internal_service Version: 1.0 License Type: Evaluation
Evaluation period left: 23 hours 59 minutes License State: Active, Not in Use, EULA
accepted2 ...
StoreIndex: 2 Feature: entservices Version: 1.0 License Type: Permanent
License State: Active, In Use
...
License Store: Dynamic License Storage
StoreIndex: 0 Feature: entservices Version: 1.0 License Type: Evaluation
Evaluation total period: 8 weeks 4 days Evaluation period left: 4 weeks 3 days
License State: Inactive
...
```

Permanent node-locked license

Temporary License

Troubleshooting License Installation

License must be installed and operational

- The system will boot up with default Lanbase license level if no licenses are installed

```
Sup7e# dir bootflash:  
44268 ... Jan 4 2011 21:46:41 ...7slot_ent_FOX1418GEW0_20110103155106655.lic
```

```
Sup7e#license install bootflash:7slot_ent_FOX1418GEW0_20110103155106655.lic  
Installing licenses from "bootflash:7slot_ent_FOX1418GEW0_20110103155106655.lic  
Installing...Feature:entservices...Successful:Supported  
1/1 licenses were successfully installed  
0/1 licenses were existing licenses  
0/1 licenses were failed to install
```


Troubleshooting License Installation

- Is the license operational?

```
Sup7e#show license all
License Store: Primary License Storage
StoreIndex: 2 Feature: entservices Version: 1.0
License Type: Permanent
License State: Active, Not in Use
License Count: Non-Counted
License Priority: Medium
```

- Reboot required to make the license operational

```
Sup7e# show version
Cisco IOS Software, IOS-XE Software, Catalyst 4500 L3 Switch Software (cat4500e-UNIVERSAL-M), Version
03.01.00.SG RELEASE SOFTWARE (fc4)
License Information for 'WS-X45-SUP7-E'
License Level: lanbase Type: Default. No valid license found
Next reboot license Level: entservices
```

Troubleshooting License Installation

License Bootlevel: HA Upgrade

- Standby supervisor will always boot to the same license level as that of Active supervisor
- If running SSO and Active and Standby are running different license levels, need to manually set level.

```
Sup7e(config)#license boot level entservices
```

Troubleshooting Catalyst 6500 Series Switches

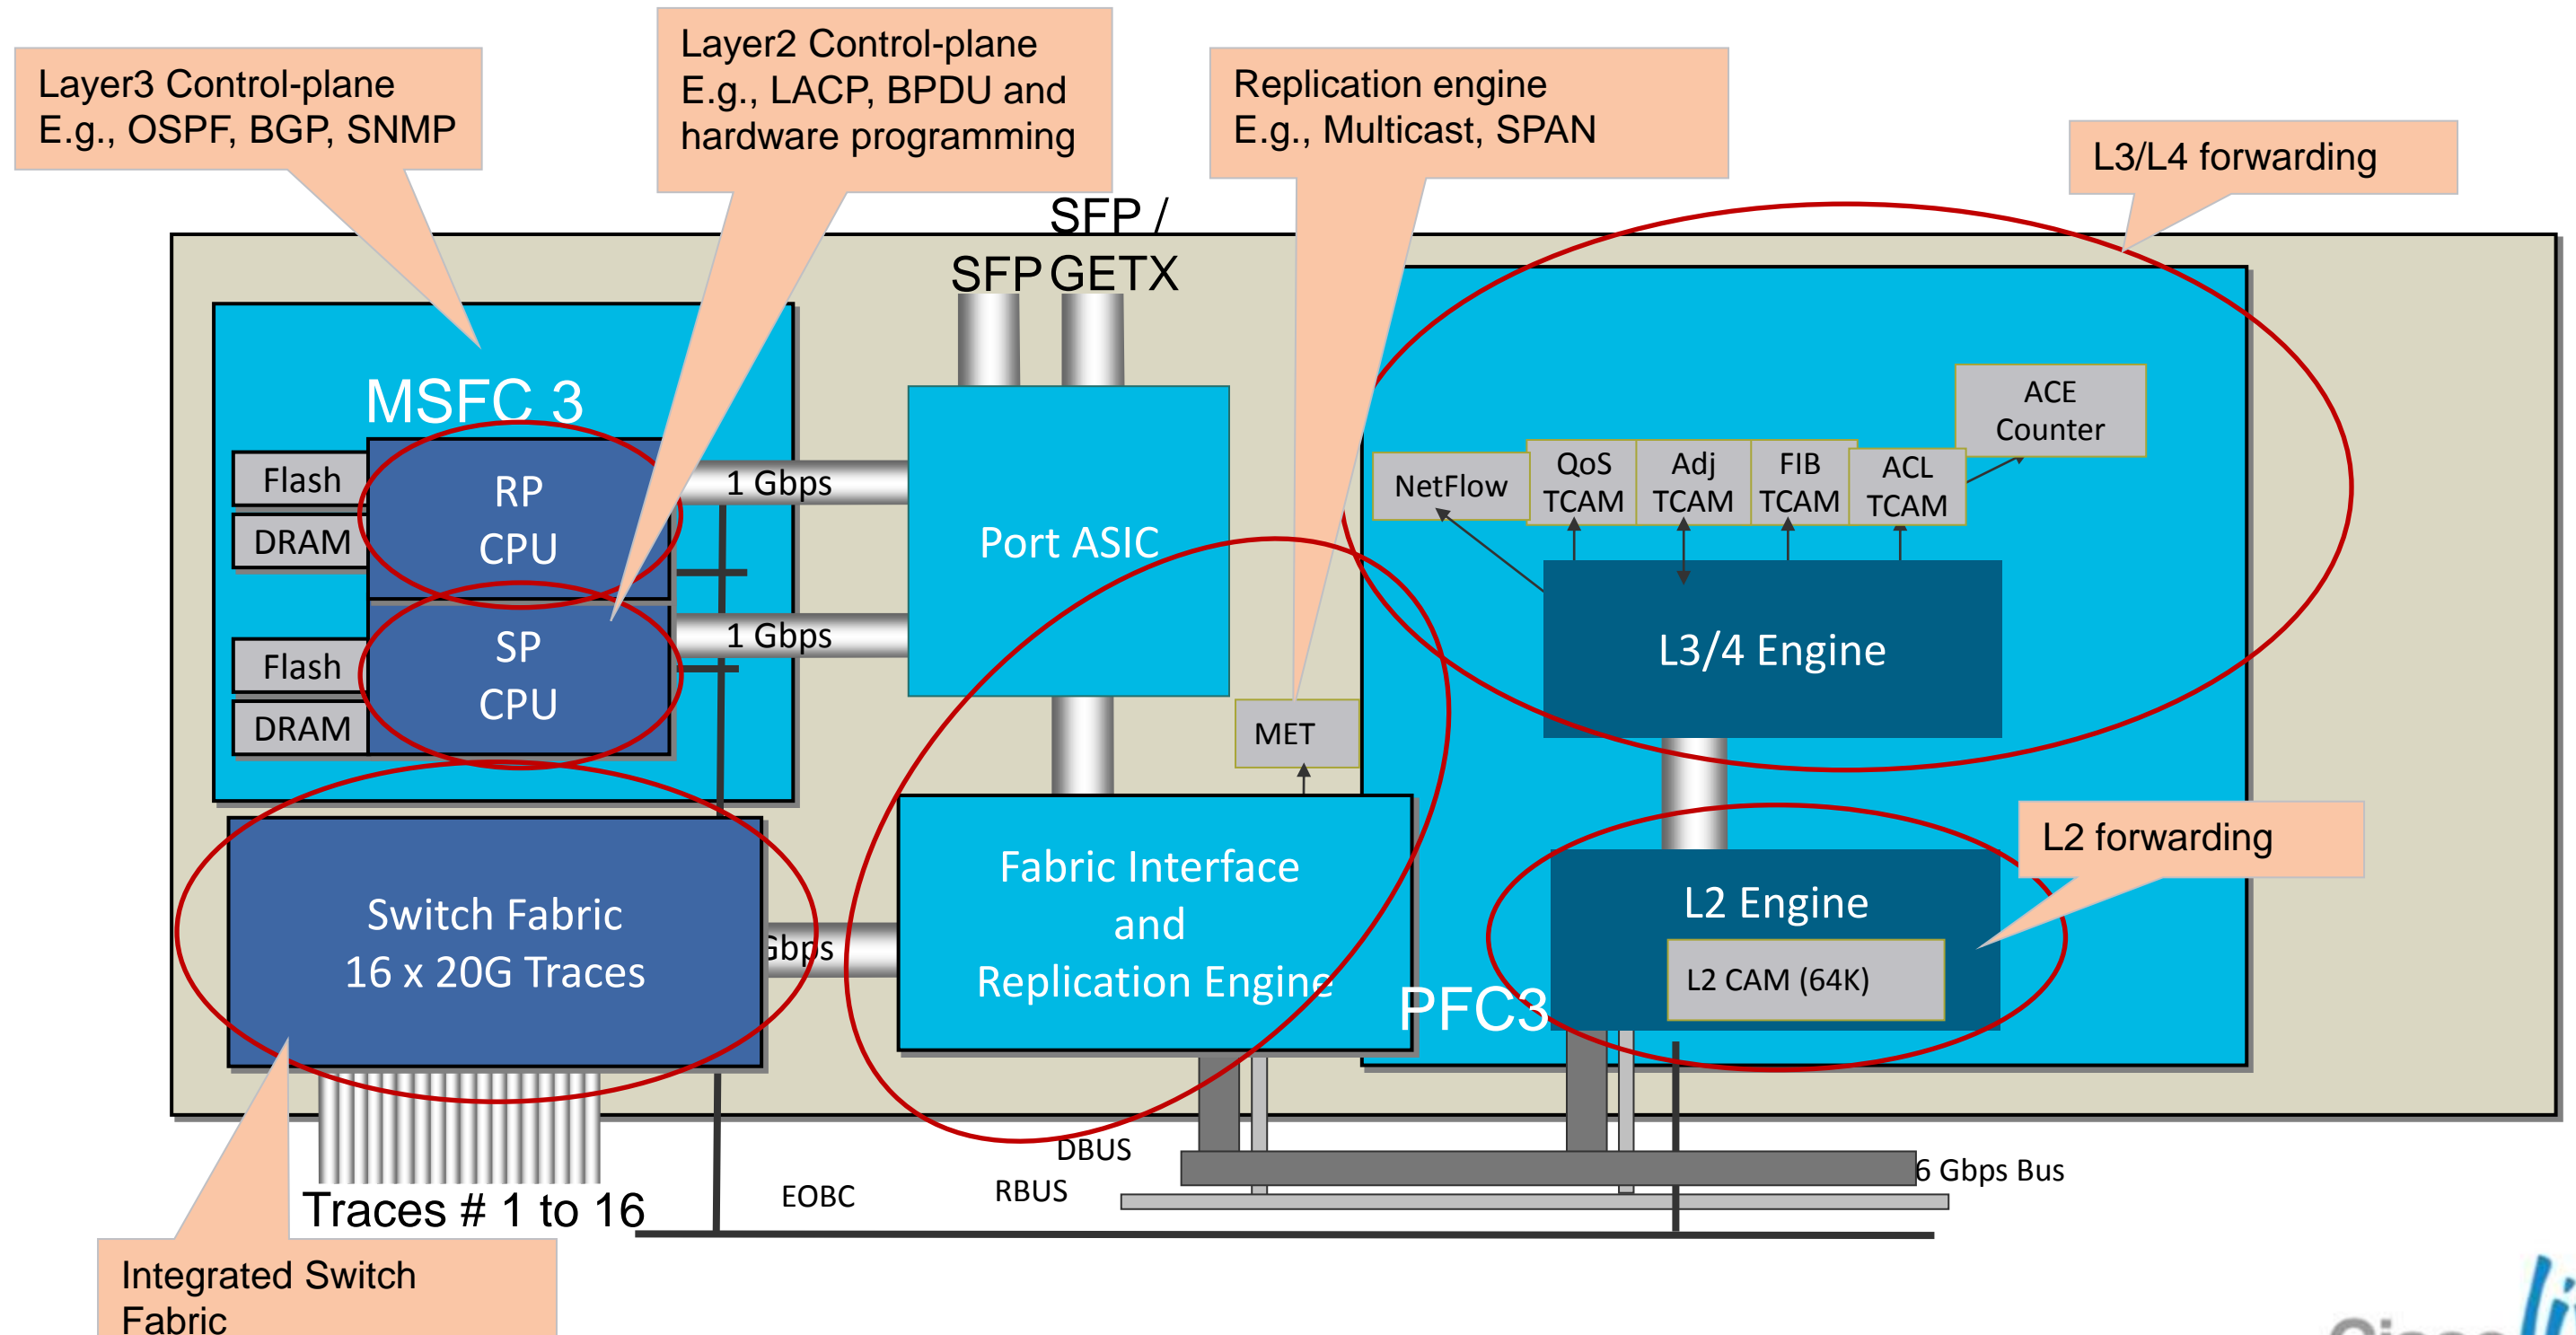


Agenda

Troubleshooting Catalyst 6500

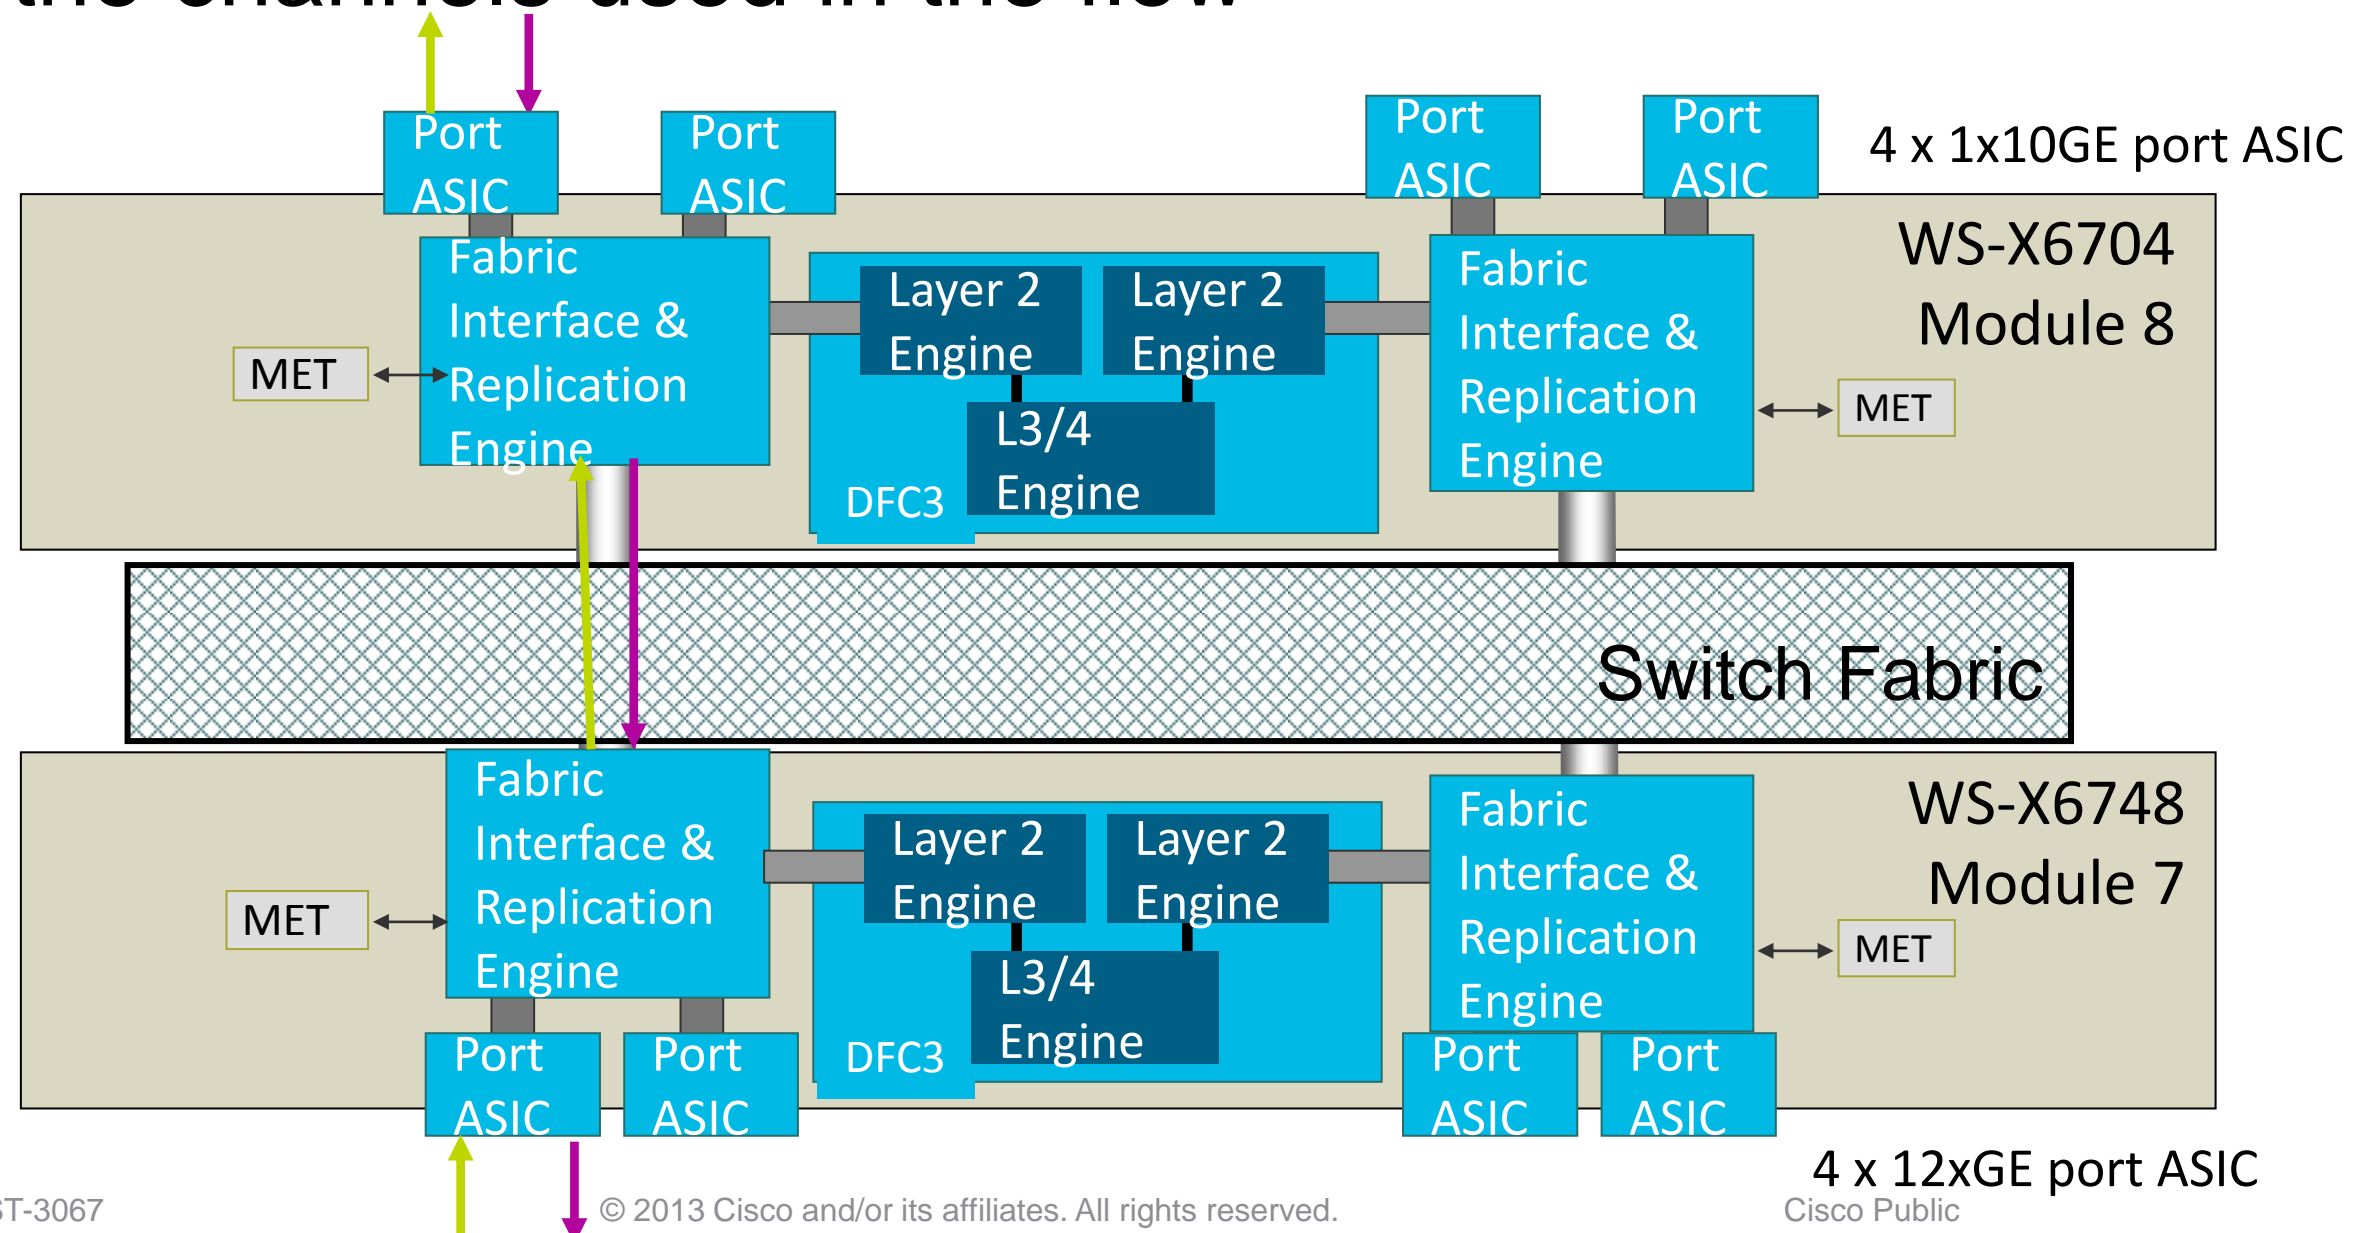
- Sup720 Architecture
- Sup720 Layer 2 and Layer 3 Unicast Troubleshooting
- High CPU utilisation
- Virtual Switch System (VSS) Troubleshooting
- Sup2T Architecture
- Sup2T Layer 3 Unicast Troubleshooting

Supervisor 720/PFC3 Architecture



L2 Packet Flow Troubleshooting

- Check the L2 forwarding engine counters
- Verify the channels used in the flow



Layer 2 Learning and Forwarding


- Layer 2 forwarding is based on {VLAN, MAC} pairs
 - Entries are stored based on result of hash done on MAC and VLAN
- MAC learning is done per PFC or DFC
 - Each PFC/DFC maintains separate L2 CAM table
- PFC and DFCs age entries independently
 - Refreshing of entries based on “seeing” traffic from specific host
 - New learns on one forwarding engine communicated to other engines via MAC-Sync process
- MAC table size:
 - 64K entries on PFC3A / 3B / 3BXL and DFC3A / 3B / 3BXL
 - 96K entries on PFC3C / 3CXL and DFC3C / 3CXL

L2 Unicast Traffic

- Verify that ARP entry is present for both hosts

```
Cat6K#show ip arp 7.0.1.1
Protocol Address      Age (min)  Hardware Addr  Type   Interface
Internet 7.0.1.1         -          000b.fca2.fe0a ARPA   Vlan700

Cat6K#show ip arp 7.0.1.2
Protocol Address      Age (min)  Hardware Addr  Type   Interface
Internet 7.0.1.2         -          0011.bced.e400 ARPA   Vlan700
```



L2 Unicast Traffic

Verify the interface that both mac-addresses are learnt on

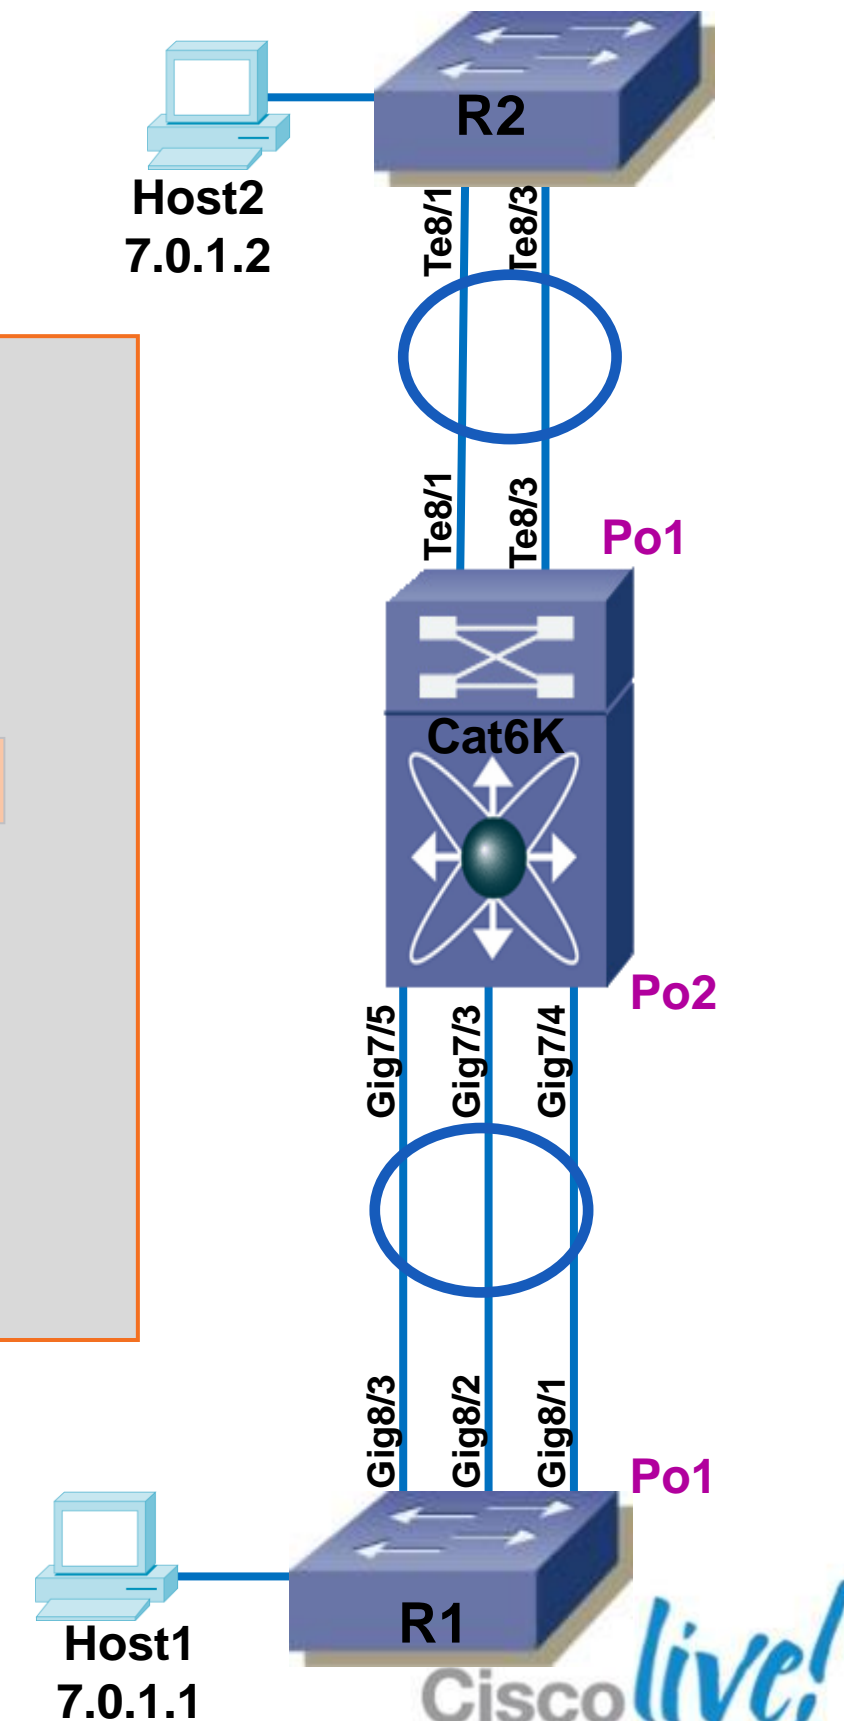
```
Cat6K#show mac-address-table address 000b.fca2.fe0a vlan 700
Legend: * - primary entry; age - seconds since last seen
```

vlan	mac address	type	learn	age	ports
Module 7[FE 1]:					
* 700	000b.fca2.fe0a	dynamic	Yes	50	Po2
Module 7[FE 2]:					
* 700	000b.fca2.fe0a	dynamic	Yes	170	Po2

Two Forwarding Engines

```
Cat6K#show mac-address-table address 0011.bced.e400 vlan 700
Legend: * - primary entry; age - seconds since last seen
```

vlan	mac address	type	learn	age	ports
Module 8[FE 1]:					
* 700	0011.bced.e400	dynamic	Yes	170	Po1
Module 8[FE 2]:					
* 700	0011.bced.e400	dynamic	Yes	170	Po1



L2 Unicast Traffic


Verify that there is an entry for all forwarding engines (FE)

```
Cat6K#show mac-address-table address 000b.fca2.fe0a vlan 700 all
```

```
Legend: * - primary entry; age - seconds since last seen
```

vlan	mac address	type	learn	age	ports
Module 1:					
700	000b.fca2.fe0a	dynamic	Yes	170	Po2
Active Supervisor:					
700	000b.fca2.fe0a	dynamic	Yes	170	Po2
Standby Supervisor:					
700	000b.fca2.fe0a	dynamic	Yes	170	Po2
Module 7[FE 1]:					
* 700	000b.fca2.fe0a	dynamic	Yes	50	Po2
Module 7[FE 2]:					
* 700	000b.fca2.fe0a	dynamic	Yes	170	Po2
Module 8[FE 1]:					
700	000b.fca2.fe0a	dynamic	Yes	170	Po2
Module 8[FE 2]:					
700	000b.fca2.fe0a	dynamic	Yes	170	Po2

Primary Entry



L2 Unicast Traffic

Check load balancing configuration. Use ingress Module number in command in case per-module load-balancing is configured (SXH images and later)

```
R1#show etherchannel load-balance module 1
```

```
EtherChannel Load-Balancing Configuration:
```

```
dst-ip
```

```
mpls label-ip
```

```
EtherChannel Load-Balancing Addresses Used Per-Protocol:
```

```
Non-IP: Destination MAC address
```

```
IPv4: Destination IP address
```

```
IPv6: Destination IP address
```

```
MPLS: Label or IP
```

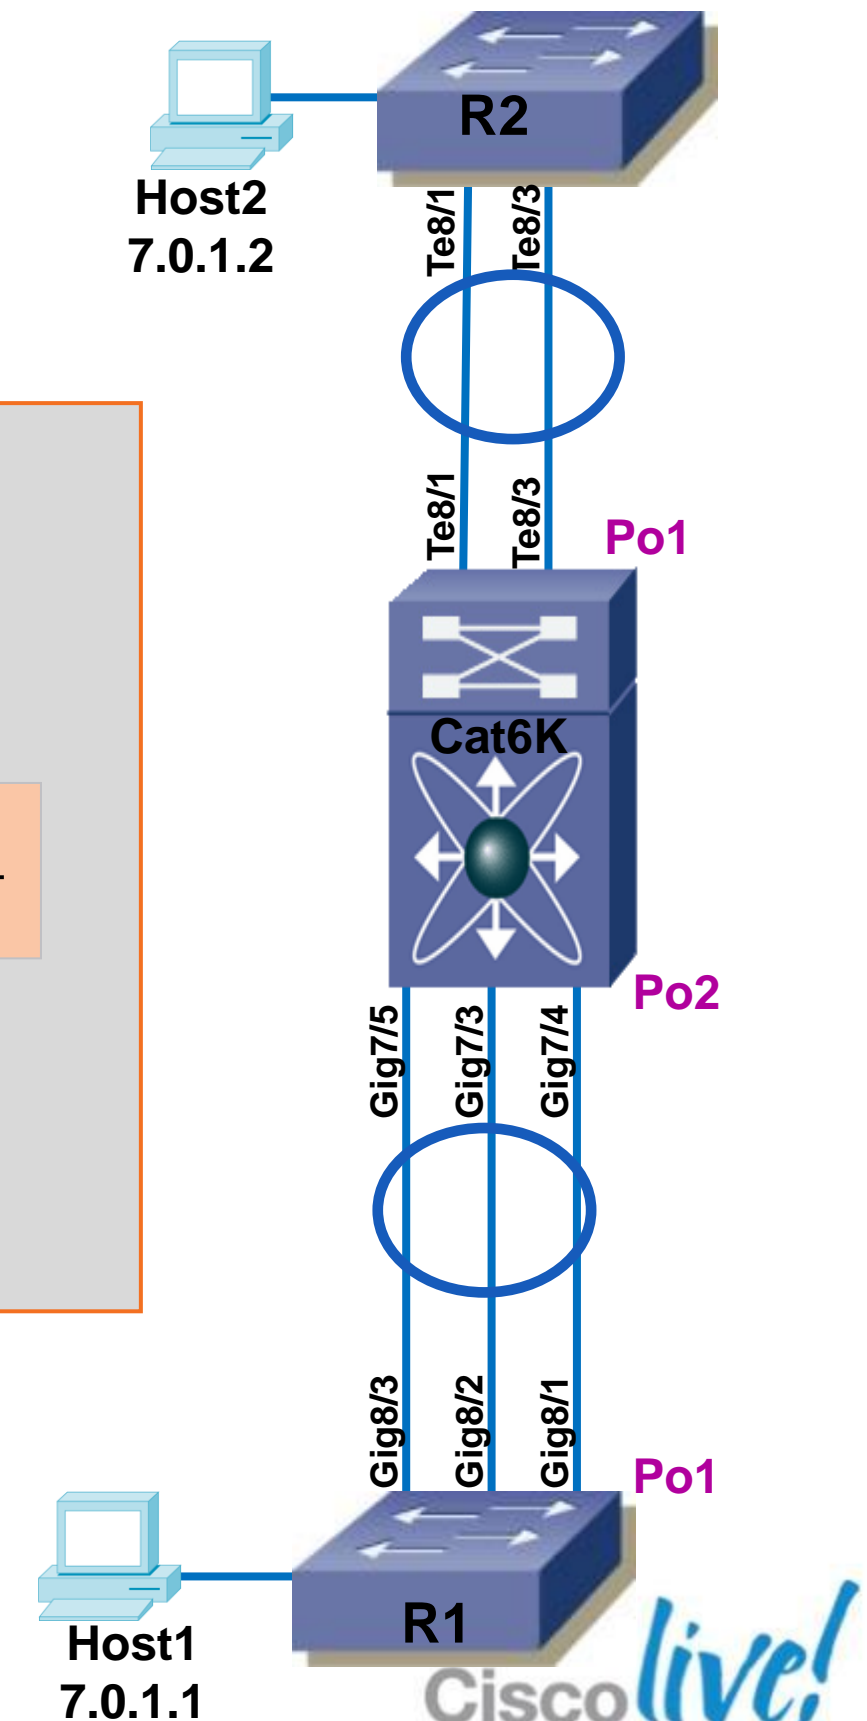
```
R1#show etherchannel load-balance hash-result interface po1 ip 7.0.1.2
```

```
Computed RBH: 0x1
```

```
Would select Gi8/1 of Po1
```

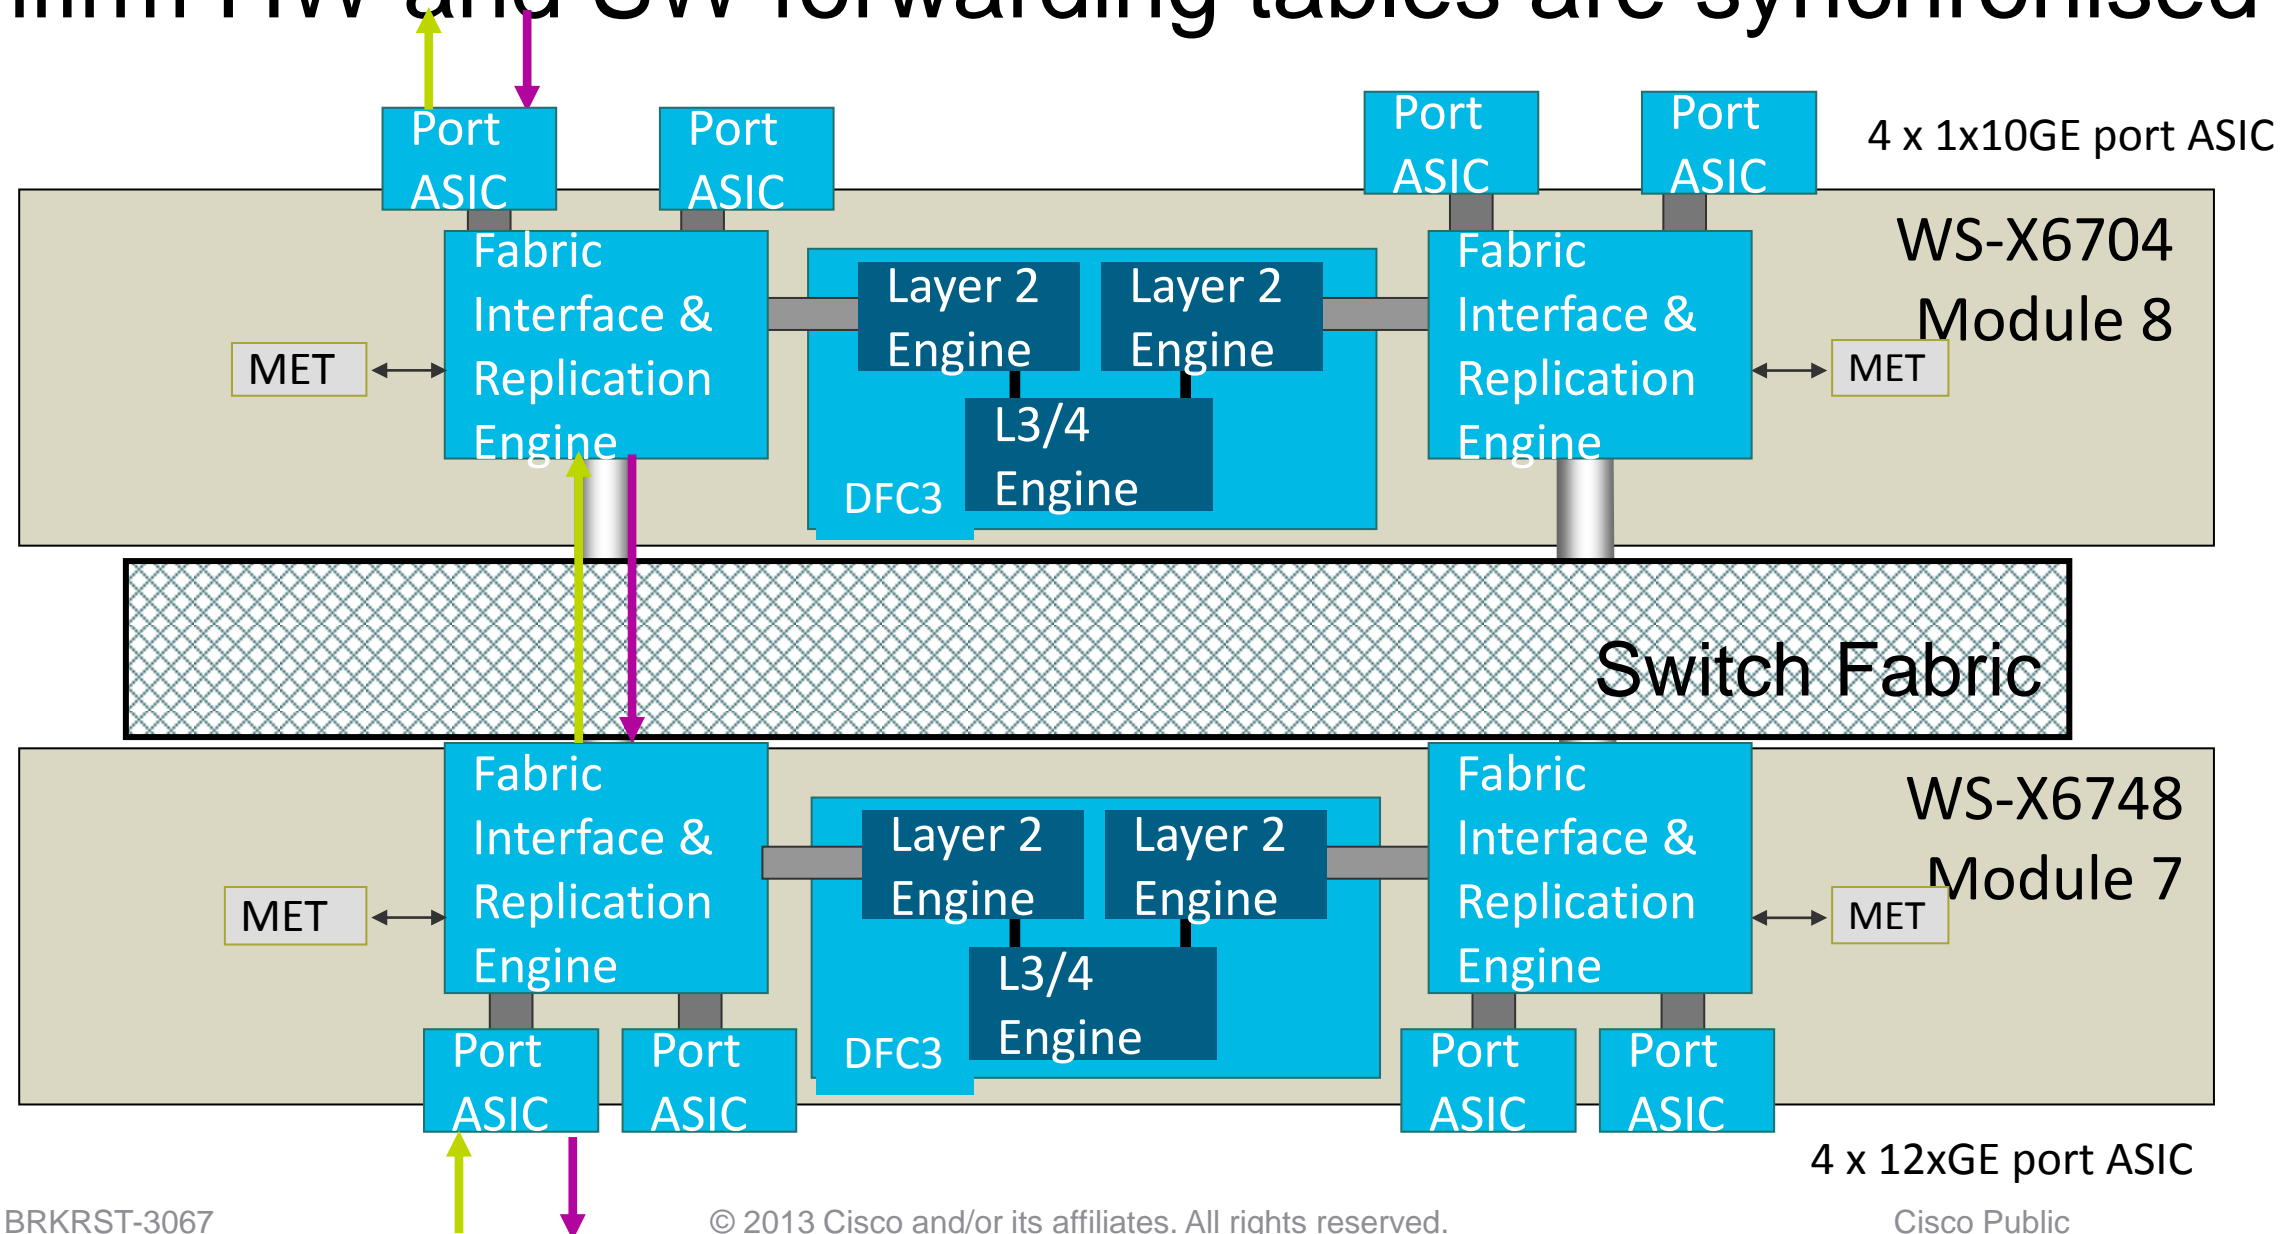
Mode is "dst-ip". Only use dest IP as argument. Prior to 12.2(33)SXH, use test etherchannel load-balance ...(same arguments) on the SP

Link selected is Gi8/1 in Po1 of R1 for traffic to 7.0.1.2 leaving R1




L3 Packet Flow Troubleshooting

- Check L3/L4 tables
- Confirm HW and SW forwarding tables are synchronised



L3 Unicast Traffic



```

R1#show ip route 9.0.1.0 | include via
Known via "eigrp 700", distance 90, metric 3328, type internal
Redistributing via eigrp 700
* 7.2.1.2, from 7.2.1.2, 00:21:58 ago, via Vlan702
  7.5.1.2, from 7.5.1.2, 00:21:58 ago, via Vlan705
  7.4.1.2, from 7.4.1.2, 00:21:58 ago, via Vlan704
  7.3.1.2, from 7.3.1.2, 00:21:58 ago, via Vlan703
  7.1.1.2, from 7.1.1.2, 00:21:58 ago, via Vlan701

R1#show ip cef exact-route 8.0.1.1 9.0.1.2
8.0.1.1      -> 9.0.1.2      : Vlan701 (next hop 7.1.1.2)

R1#show mls cef exact-route 8.0.1.1 0 9.0.1.2 0
Interface: v1705, Next Hop: 7.5.1.2, Vlan: 705, Destination Mac:
0050.f0f8.7400
    
```

SW

Equal Cost Routes to the destination prefix

Next hop used for SW based CEF (SW forwarding data path)

HW

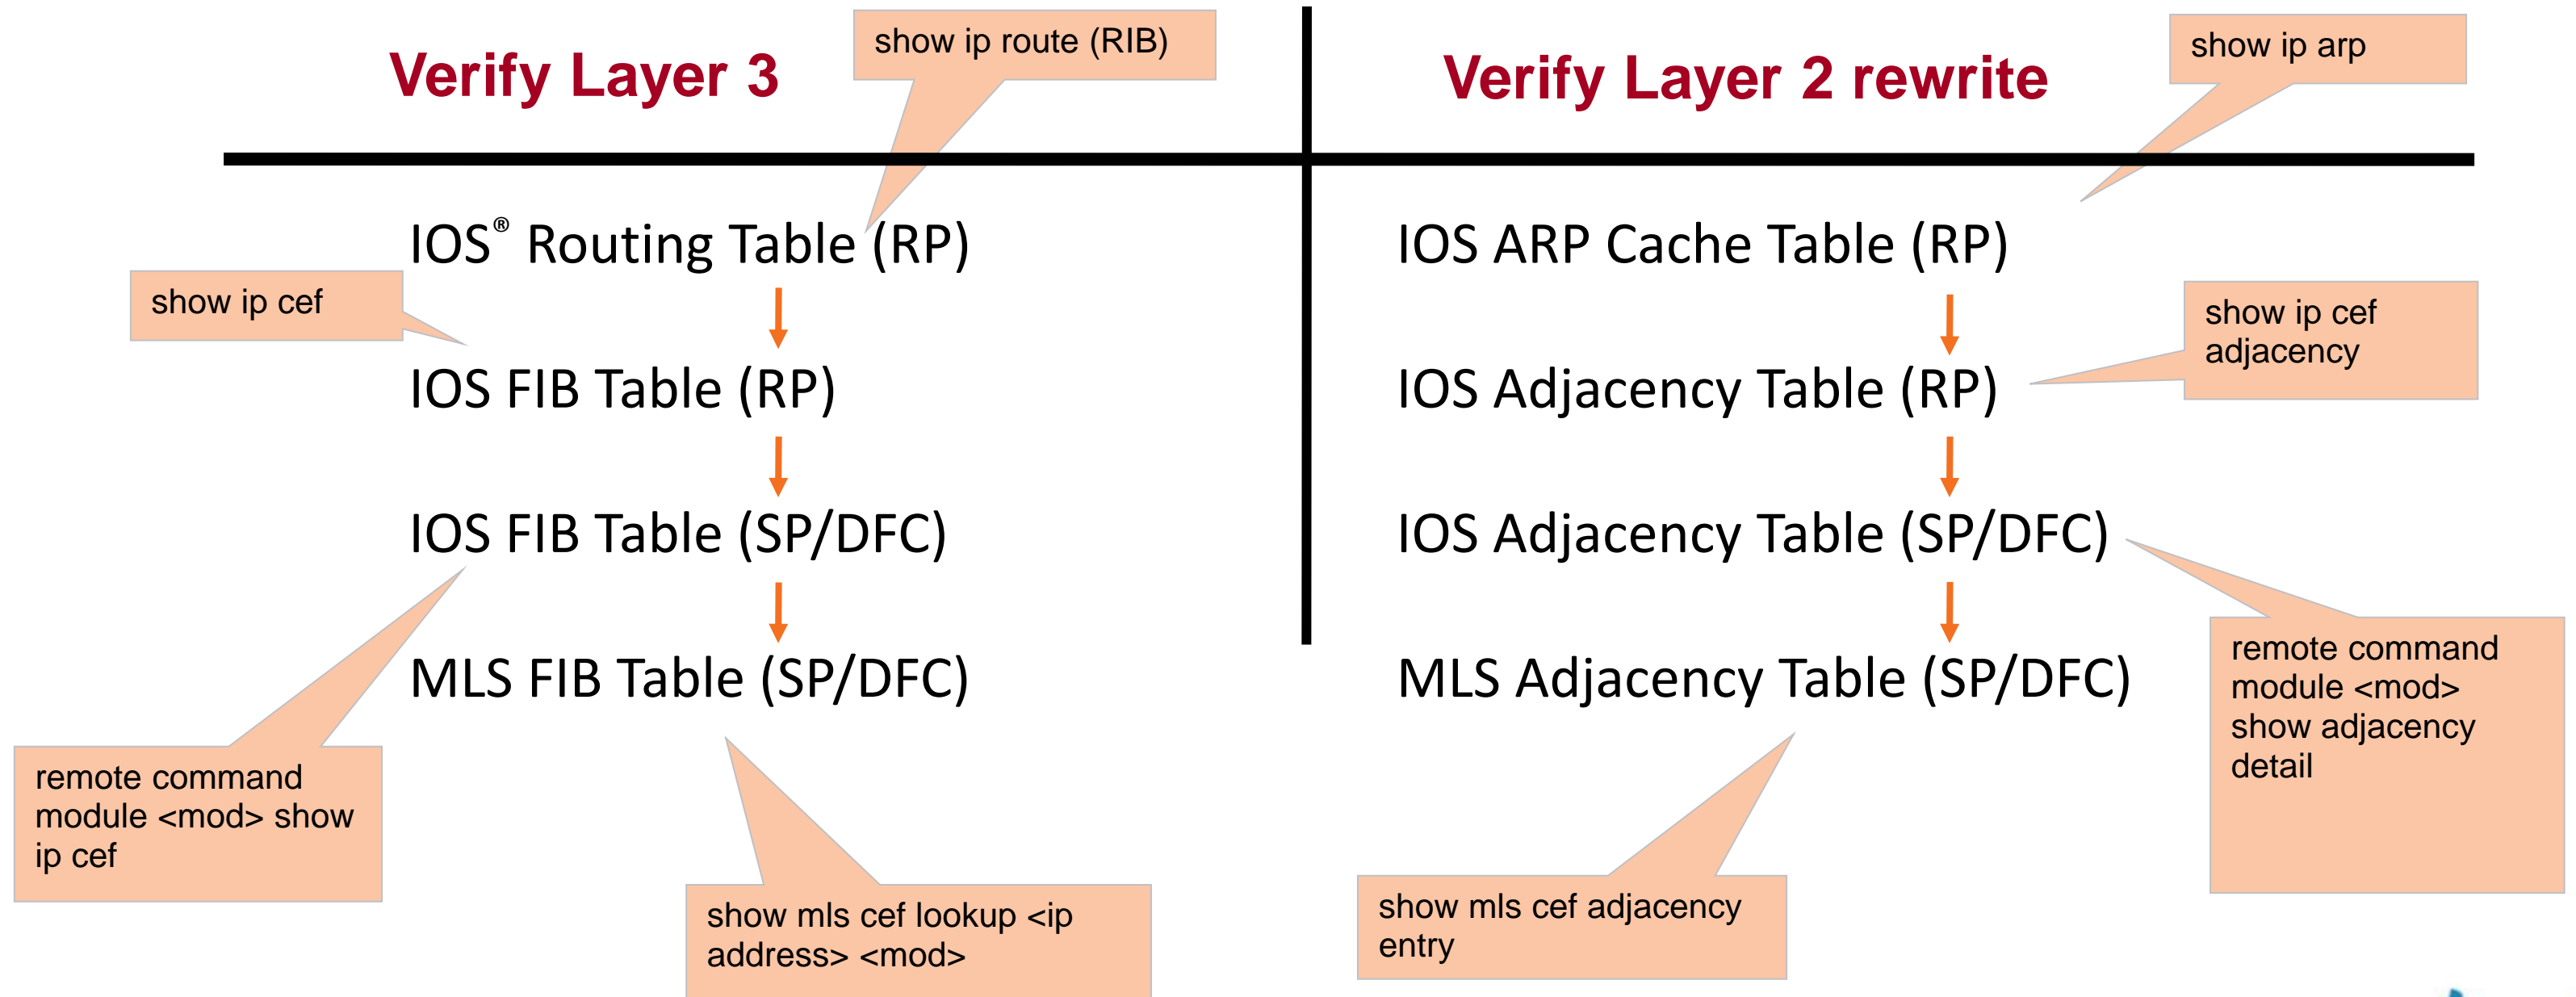
Check which link between R1 and Cat6K is chosen.

Next hop used for HW based CEF (HW forwarding path). Note: "0" is used for both src and dest L4 port numbers as test flow was ICMP echo



FIB/Adjacency Tables

L3 FIB Table Programming Flow




L3 Unicast Traffic

```
Cat6K#show ip route 9.0.1.0 | i via
Known via "eigrp 700", distance 90, metric 3072, type internal
Redistributing via eigrp 700
* 7.7.1.2, from 7.7.1.2, 00:07:33 ago, via TenGigabitEthernet8/3
  7.6.1.2, from 7.6.1.2, 00:07:33 ago, via TenGigabitEthernet8/1
```

```
Cat6K#show vlan internal usage
[snip]
1090 TenGigabitEthernet8/3
1091 TenGigabitEthernet8/1
[snip]
```

L3 Interface map internally to a "1-port" VLAN



L3 Unicast Traffic

```
Cat6K#show ip cef 9.0.1.2
```

SW

```
9.0.0.0/8
```

```
  nexthop 7.6.1.2 TenGigabitEthernet8/1
```

```
  nexthop 7.7.1.2 TenGigabitEthernet8/3
```

```
Cat6K#show ip cef exact-route 8.0.1.1 9.0.1.2
```

```
8.0.1.1 -> 9.0.1.2 => IP adj out of TenGigabitEthernet8/1, addr 7.6.1.2
```

```
Cat6K#show ip cef adjacency TenGigabitEthernet 8/1 7.6.1.2
```

```
7.6.1.2/32
```

```
  attached to TenGigabitEthernet8/1
```

```
9.0.0.0/8
```

```
  nexthop 7.6.1.2 TenGigabitEthernet8/1
```

```
Cat6K#show mls cef lookup 9.0.1.2 mod 7
```

HW

```
Codes: decap - Decapsulation, + - Push Label
```

```
Index Prefix
```


```
Adjacency
```

```
108749 9.0.0.0/8 Te8/1 , 000f.f8e4.d000 (Hash: 007F)
```

```
Te8/3 , 000f.f8e4.d000 (Hash: 7F80)
```

```
Cat6K#show mls cef exact-route 8.0.1.1 0 9.0.1.2 0 module 7
```

```
Interface: Te8/3, Next Hop: 7.7.1.2, Vlan: 1090, Destination Mac: 000f.f8e4.d000
```



L3 Unicast Traffic

```
Cat6K#show adjacency ten 8/3 7.7.1.2 detail
```

```
Protocol Interface Address
IP          TenGigabitEthernet8/3 7.7.1.2 (17)
           2001 packets, 228114 bytes
           epoch 0
           sourced in sev-epoch 774
           Encap length 14
           000FF8E4D0000050F0F874000800
           ARP
```

Aggregate HW adjacency statistics
(SW collects it from all DFC/PFC's for
all prefixes linked to this adjacency)

Rewrite information
(Dmac|Smac|0800): verify it is conform
with next hop rewrite info

To get HW adjacency statistic
for this prefix on this module

```
Cat6K#show mls cef lookup 9.0.1.2 detail mod 7
```

```
Codes: M - mask entry, V - value entry, A - adjacency index, P - priority bit
       D - full don't switch, m - load balancing modnumber, B - BGP Bucket sel
       V0 - Vlan 0, C0 - don't comp bit 0, V1 - Vlan 1, C1 - don't comp bit 1
       RVTEN - RPF Vlan table enable, RVTSEL - RPF Vlan table select
```

```
Format: IPV4_DA - (8 | xtag vpn pi cr recirc tos prefix)
```

```
Format: IPV4_SA - (9 | xtag vpn pi cr recirc prefix)
```

```
M(108749 ): E | 1 FFF 0 0 0 0 255.0.0.0
```

```
V(108749 ): 8 | 1 0 0 0 0 0 9.0.0.0 (A:294933 , P:1, D:0, m:14, B:0)
```

Start adjacency pointer is 294933, 14 + 1 = 15
adjacencies linked to the prefix

L3 Unicast Traffic

```
Cat6K#show mls cef adjacency entry 294933 to 294947 mod 7
```

```
Index: 294933 smac: 0050.f0f8.7400, dmac: 000f.f8e4.d000
      mtu: 9234, vlan: 1091, dindex: 0x0, l3rw_vld: 1
      packets: 0, bytes: 0
Index: 294947 smac: 0050.f0f8.7400, dmac: 000f.f8e4.d000
      mtu: 9234, vlan: 1090, dindex: 0x0, l3rw_vld: 1
      packets: 0, bytes: 0
```

15 HW adjacencies linked to this prefix: which one is really used ?

```
Cat6K#show mls cef adjacency entry 294933 to 294947 mod 7 | i packets
```

```
294933 → packets: 0, bytes: 0
294934 → packets: 0, bytes: 0
294935 → packets: 0, bytes: 0
294936 → packets: 2001, bytes: 236118
294937 → packets: 0, bytes: 0
294938 → packets: 0, bytes: 0
294939 → packets: 0, bytes: 0
```

Based on the packet counts, we see that the 4th adjacency (entry 294936) is being used.

Verify that the rewrite information is correct for the adjacency.

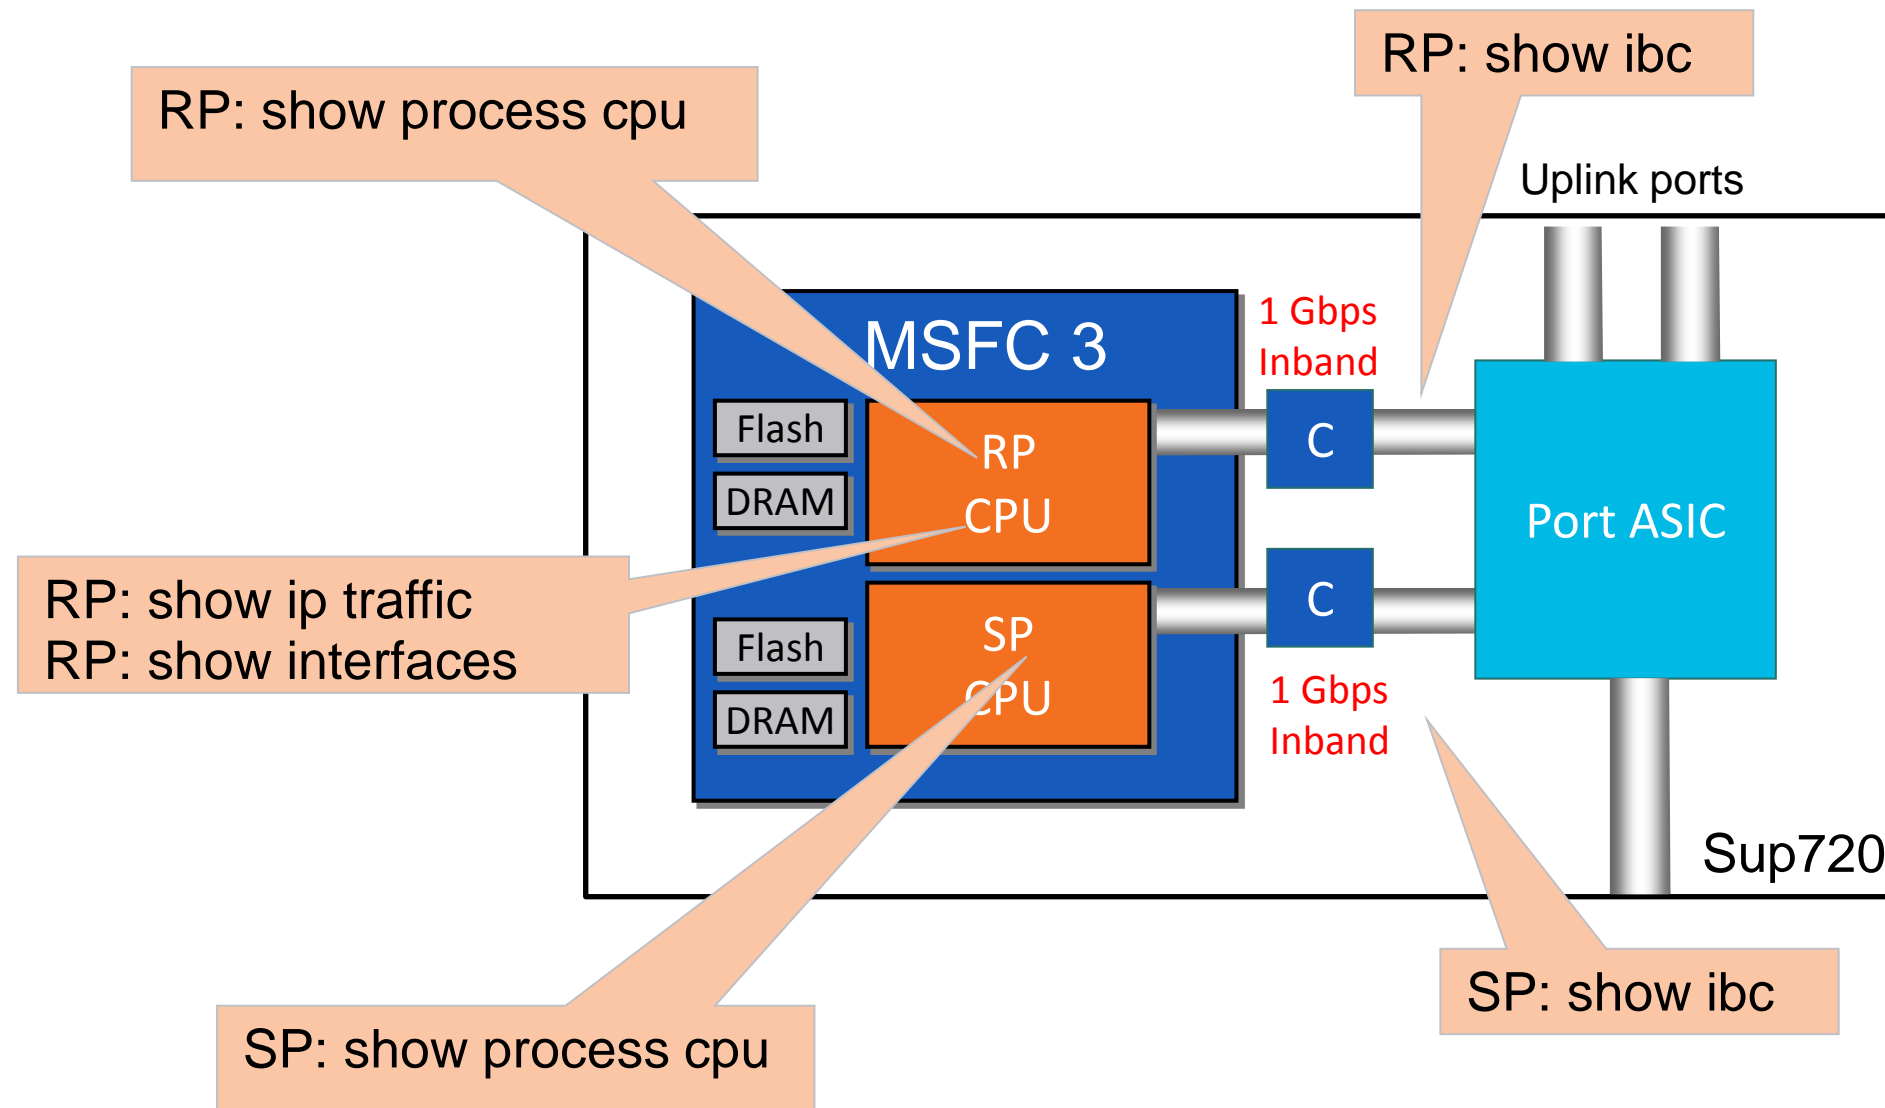
```
Cat6K#show mls cef adjacency entry 294936 detail mod 7
```

```
Index: 294936 smac: 0050.f0f8.7400, dmac: 000f.f8e4.d000
...
packets: 0, bytes: 0
```

Note counter is cleared here; counter is cleared when adjacency is read.

Troubleshooting High CPU

Components Involved



High CPU Utilisation - Process

Process: ARP Input

- Caused by ARP flooding
- Static route configured with next-hop interface instead of IP

Incrementing at very high rate

```
Cat6K#show ip traffic | begin ARP
ARP statistics:
  Rcvd: 6512 requests, 2092 replies, 0 reverse, 0 other
  Sent: 258 requests, 707 replies (0 proxy), 0 reverse
  Drop due to input queue full: 20
<snip>
Cat6K#show interfaces | include line protocol|rate
Vlan501 is up, line protocol is up
  5 minute input rate 23013521 bits/sec, 2535 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
```

Look for abnormal input rate

VLAN SVI is Virtual and counter gives only the amount/rate of process-switched or control-pane traffic

High CPU Utilisation

- Identify if it is process driven or interrupts

Total CPU usage (Process + Interrupt)

CPU usage due to Interrupt

```
DUT# show process cpu
CPU utilization for five seconds: 99%/90%; one minute: 9%; five minutes: 8%
PID  Runtime(ms)  Invoked  uSecs   5Sec   1Min   5Min  TTY  Process
  2      720          88    8181   9.12%  1.11%  0.23%  18  Virtual Exec
```

- Next Steps

- Process: recurring events, control plane process etc.
- Interrupts: incorrect switching path, system exceeding hardware resources

High CPU Utilisation - Process

Process: IP Input

- Caused by traffic that needs to be process switched or sent to the CPU
- Common Causes and Next Steps
 - Broadcast Storm
 - Traffic with IP-Options enabled
 - Traffic to which ICMP redirect or Unreachable required e.g., TTL=1, ACL Deny
 - Traffic that needs further CPU processing e.g., ACL Logging

Find the source host and isolate. Consider per-port broadcast storm-control

Find the source host and isolate. Consider HW Rate-limiters

Consider Optimised ACL Logging (OAL)

Under VLAN SVIs configure:
no ip unreachable
no ip redirects
no ip proxy-arp
Consider HW Rate-limiters

High CPU Utilisation - Process

Process: Exec and Virtual Exec

```
Cat6K#show process cpu | include CPU|Virtual |Exec
CPU utilization for five seconds: 30%/0%; one minute: 8%; five minutes: 5%
PID Runtime(ms)   Invoked    uSecs   5Sec   1Min   5Min TTY Process
  3      272         194      1402  29.00% 2.12% 1.89% 0 Exec
 54      180        1443      124   0.00% 0.00% 0.00% 1 Virtual Exec
```

Responsible for tty lines
(console, auxiliary)

High CPU when too many messages
sent to console/vty

Responsible for vty lines (telnet,
SSH)

■ Next steps

- Check if any debug is enabled via “show debug”. Issue “undebug all” if it is not needed
- Disable logging by “no logging console” or “no logging terminal”

High CPU Utilisation - Traffic

Analyse Input Buffer

- Use when an input queue is oversubscribed

```
Cat6K#show buffers assigned
Header  DataArea  Pool Rcnt  Size Link  Enc  Flags  Input  Output
46FDBC14 8029784 Small  1    77  36   1    200  V1100  None
46FE0010 802CBC4 Small  1    77  36   1    200  V1100  None
. . .

Cat6K#show buffers input-interface vlan 100 dump
Buffer information for RxQ3 buffer at 0x378B3BC
 data_area 0x7C05EF0, refcount 1, next 0x0, flags 0x200
 linktype 7 (IP), enctype 1 (ARPA), encsize 14, rxtyp 1
 if_input 0x46C7C68 (Vlan100), if_output 0x0 (None)
 inputtime 2d03h (elapsed 00:00:01.024)
 outputtime 00:00:00.000 (elapsed never), oqnumber 65535
 datagramstart 0x7C05F36, datagramsize 62, maximum size 2196
 mac_start 0x7C05F36, addr_start 0x7C05F36, info_start 0x0
 network_start 0x7C05F44, transport_start 0x7C05F58, caller_pc 0x6C1564

 source: 137.34.219.3, destination: 224.0.0.2, id: 0x0000, ttl: 1,
 TOS: 192 prot: 17, source port 1985, destination port 1985
```

Find the interface that's holding most of the buffers

Packet details

High CPU Utilisation - Traffic

Debug Netdr Capture

- Capture packets being received and sent by RP to buffer space
- Does not introduce extra CPU processing

```
Cat6K#debug netdr cap ?
acl                (11) Capture packets matching an acl
and-filter         (3) Apply filters in an and function: all must match
continuous        (1) Capture packets continuously: cyclic overwrite
destination-ip-address (10) Capture all packets matching ip dst address
dstindex          (7) Capture all packets matching destination index
ethertype         (8) Capture all packets matching ethertype
interface         (4) Capture packets related to this interface
or-filter         (3) Apply filters in an or function: only one must
match
rx                (2) Capture incoming packets only
source-ip-address (9) Capture all packets matching ip src address
srcindex          (6) Capture all packets matching source index
tx                (2) Capture outgoing packets only
vlan             (5) Capture packets matching this vlan number
<cr>
```

High CPU Utilisation - Traffic

Netdr Capture Output

```
Cat6K#show netdr cap
```

```
A total of 10 packets have been captured  
The capture buffer wrapped 0 times  
Total capture capacity: 4096 packets
```

Vlan

```
----- dump of incoming inband packet -----  
interface Vl10, routine draco2_process_rx_packet_inline  
dbus info: src_vlan 0xA(10), src_indx 0xC0(192), len 0x76(118)  
  bpdv 0, index_dir 0, flood 0, dont_lrn 0, dest_indx 0x380(896)  
  08020401 000A0000 00C00000 76080000 00010428 0E000040 00000000 03800000  
mistral_hdr: req_token 0x0(0), src_index 0xC0(192), rx_offset 0x76(118)  
  requeue 0, obl_pkt 0, vlan 0xA(10)  
destmac 00.23.04.18.F8.80, srcmac 00.13.7F.8B.84.C1, protocol 0800  
protocol ip: version 0x04, hlen 0x05, tos 0x00, totlen 100, identifier 13202  
  df 0, mf 0, fo 0, ttl 255, src 10.10.10.1, dst 10.10.10.3  
  icmp type 8, code 0
```

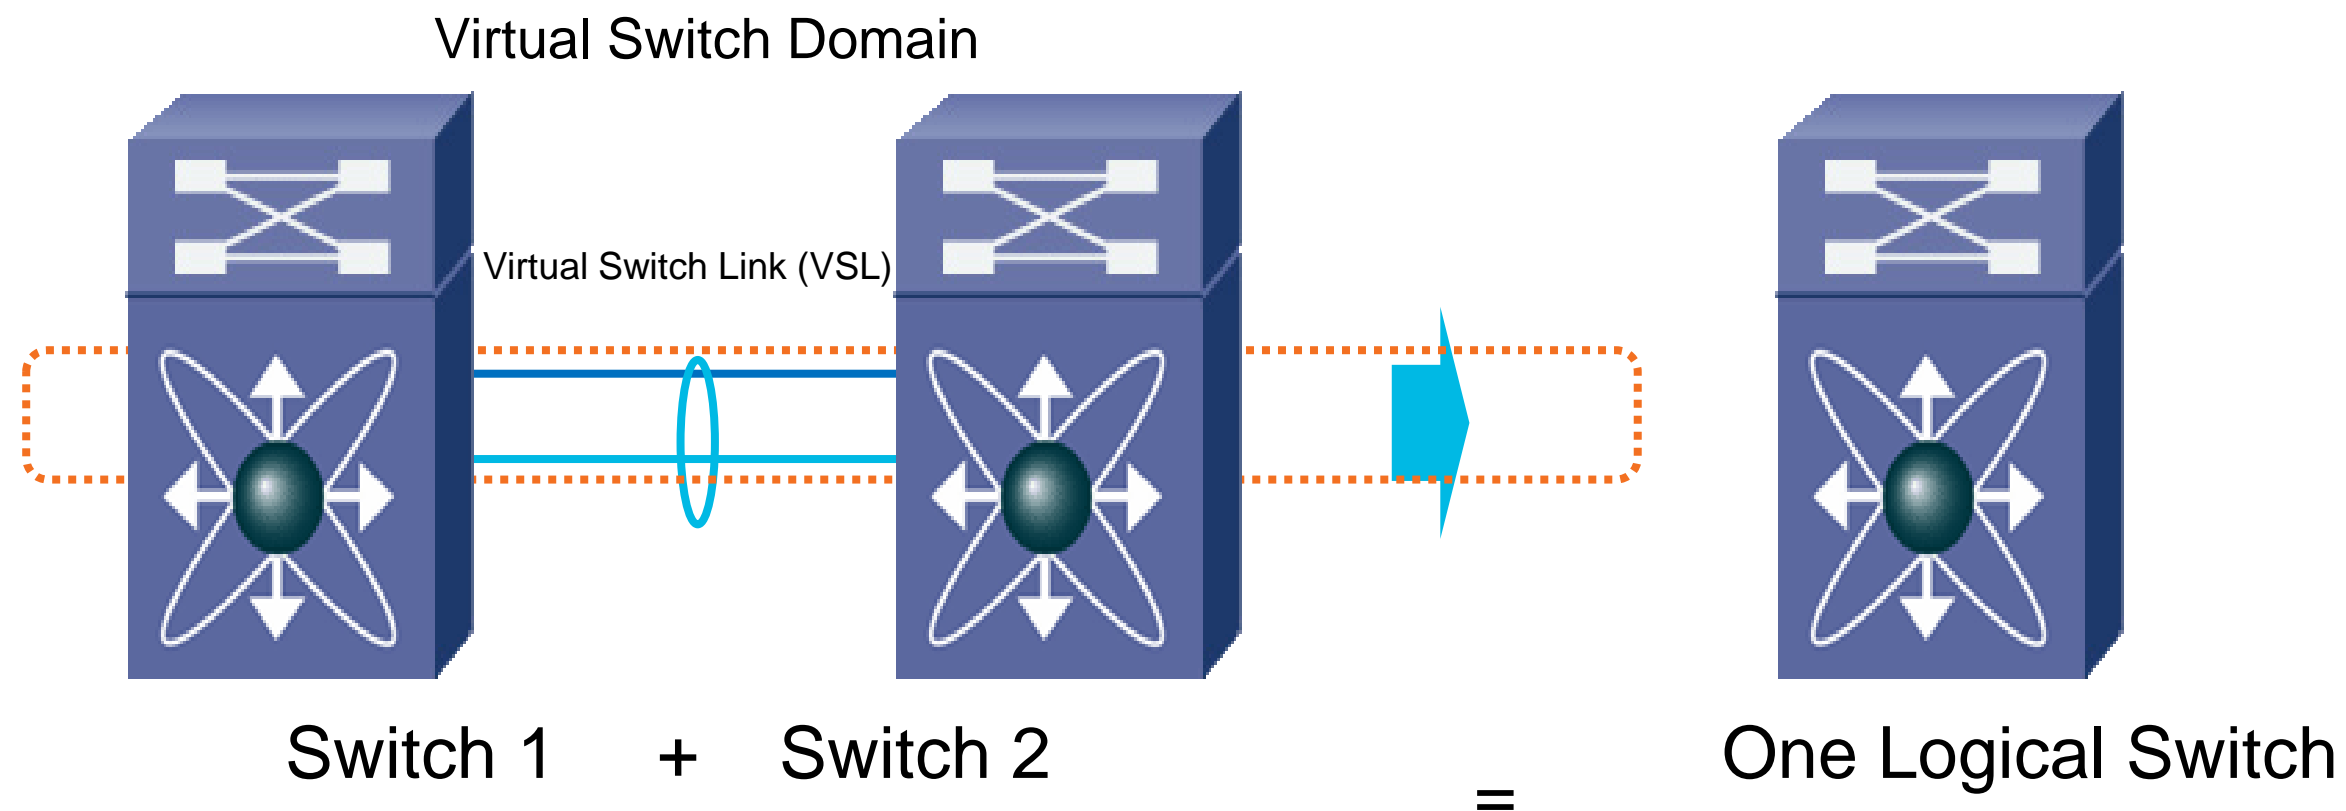
Ethertype

Source/Destination

```
----- dump of incoming inband packet -----  
interface Vl10, routine draco2_process_rx_packet_inline  
dbus info: src_vlan 0xA(10), src_indx 0xC0(192), len 0x76(118)  
  bpdv 0, index_dir 0, flood 0, dont_lrn 0, dest_indx 0x380(896)
```

Virtual Switch System (VSS)

- Virtual Switching System consists of two Cisco Catalyst 6500 Series switches connected through a special etherchannel called a Virtual Switch Link (VSL) to become one logical entity.



VSS Control Plane

VSS Specific Protocols

- VSL Protocol (VSLP) runs between active and standby switch over the VSL, and has two components:
 - LMP and RRP
- Link Management Protocol (LMP): Runs over each individual link in VSL
 - Bi-directionality; keepalives; connectivity (peer-to-peer)
- Role Resolution Protocol (RRP): Runs on each side (each peer) of the VSL port channel
 - HW/SW compatibility; role negotiation; control-link selection

VSS Control Plane

VSS Configuration Check

```
Cat6K#show switch virtual role
```

Switch	Switch Number	Status	Preempt Oper (Conf)	Priority Oper (Conf)	Role	Session ID Local	Session ID Remote
LOCAL	1	UP	TRUE (Y*)	200 (200)	ACTIVE	0	0
REMOTE	2	UP	TRUE (Y*)	100 (100)	STANDBY	2977	3643

```
Standby configured preempt timer(switch 2): 5 minutes
```

```
Active configured preempt timer(switch 1): 5 minutes
```

```
In dual-active recovery mode: No
```

Switch 1 is active, switch 2 is standby, both are up

```
Cat6K#show switch virtual link port-channel
```

```
Flags: D - down P - bundled in port-channel
```

Group	Port-channel	Protocol	Ports
256	Po256 (RU)	-	Te1/3/3 (P) Te1/5/4 (P)
255	Po255 (RU)	-	Te2/2/3 (P) Te2/5/4 (P)

Interfaces are identified by [switch#]/[mod#]/[port#]

Switch 1 side of the VSL

Switch 2 side of the VSL

VSS Control Plane

VSS LMP Check

```
Cat6K#show switch virtual link
```

```
VSL Status : UP
```

```
VSL Uptime : 18 hours, 12 minutes
```

```
VSL SCP Ping : Pass
```

```
VSL ICC Ping : Pass
```

```
VSL Control Link : Te1/5/4
```

```
Cat6K#show switch virtual link port
```

```
LMP summary
```

```
Link info:
```

```
Configured: 4
```

```
Operational: 4
```

```
Peer Peer Peer Peer Timer(s) running  
Flag MAC Switch Interface (Time remaining)
```

```
-----  
Te1/5/4 vfs operational vfs 0011.bc75.4400 2 Te2/5/4 T4 (220ms)  
T5 (175s)  
Te1/3/3 vfs operational vfs 0011.bc75.4400 2 Te2/2/6 T4 (220ms)  
T5 (175s)  
Te1/3/4 vfs operational vfs 0011.bc75.4400 2 Te2/2/8 T4 (220ms)  
T5 (175s)  
Te1/3/6 vfs operational vfs 0011.bc75.4400 2 Te2/2/3 T4 (768ms)  
T5 (175s)
```

```
Flags: v - Valid flag set f - Bi-directional flag set s - Negotiation flag set
```

Status and uptime of the VSL

Control link carries EOBC and IBC control messages (SCP and ICC/IPC)

Check LMP state and Flags (vf) of the links in the VSL bundle

VSS Control Plane

VSS LMP Check

```

Cat6k#show vslp lmp counters
Instance #1:
  LMP counters

```

Interface	Tx		Bidir	Rx		
	OK	Fail		Uni	Fail	Bad
Te1/5/4	12649	0	12675	1	0	0
Te1/3/3	12000	0	12024	0	0	0
Te1/3/4	11999	0	12024	0	0	0
Te1/3/6	12001	0	12025	0	0	0


```

Rx error details

```

Interface	My info mismatch	My info absent	Bad MAC Address	Bad switch id	Domain id mismatch	Peer info mismatch
Te1/5/4	0	1	0	0	0	0
Te1/3/3	0	0	0	0	0	0
Te1/3/4	0	0	0	0	0	0
Te1/3/6	0	0	0	0	0	0

Indicates problem sending or receiving LMP packets from peer

Indicates problem sending LMP packet to the VSL peer

Packets received from VSL peer without our info, proving the link is unidirectional at that moment

These errors usually indicate a misconfiguration on one of the peers

1 unidirectional packet when first link in VSL came up

VSS Control Plane

VSS RRP Check

```
Cat6K#show switch virtual role detail
```

Switch	Switch Number	Status	Preempt Oper (Conf)	Priority Oper (Conf)	Role	Session ID Local	Session ID Remote
--------	---------------	--------	---------------------	----------------------	------	------------------	-------------------

Check 1 is active, 2 is standby

LOCAL	1	UP	TRUE (Y*)	200 (200)	ACTIVE	0	0
REMOTE	2	UP	TRUE (Y*)	100 (100)	STANDBY	2977	3643

Standby configured preempt timer (switch 2): 5 minutes

Active configured preempt timer (switch 1): 5 minutes

RRP Counters:

Inst.	Peer	Direction	Req	Acc	Est	Rsugg	Racc
1	1	Tx	0	2	0	2	6
1	1	Rx	2	0	2	0	6

State machine info on RRP protocol; current state is "role resolved"

RRP FSM info

```
sm(vslp_rrp RRP SM information for Instance 1, Peer 1), running yes, state role_res
```

```
Last transition recorded: (req)-> hold (srt_exp)-> hold (est)-> role_neg (srt_exp)-> role_neg (est)-> role_neg (racc)-> role_res (srt
```

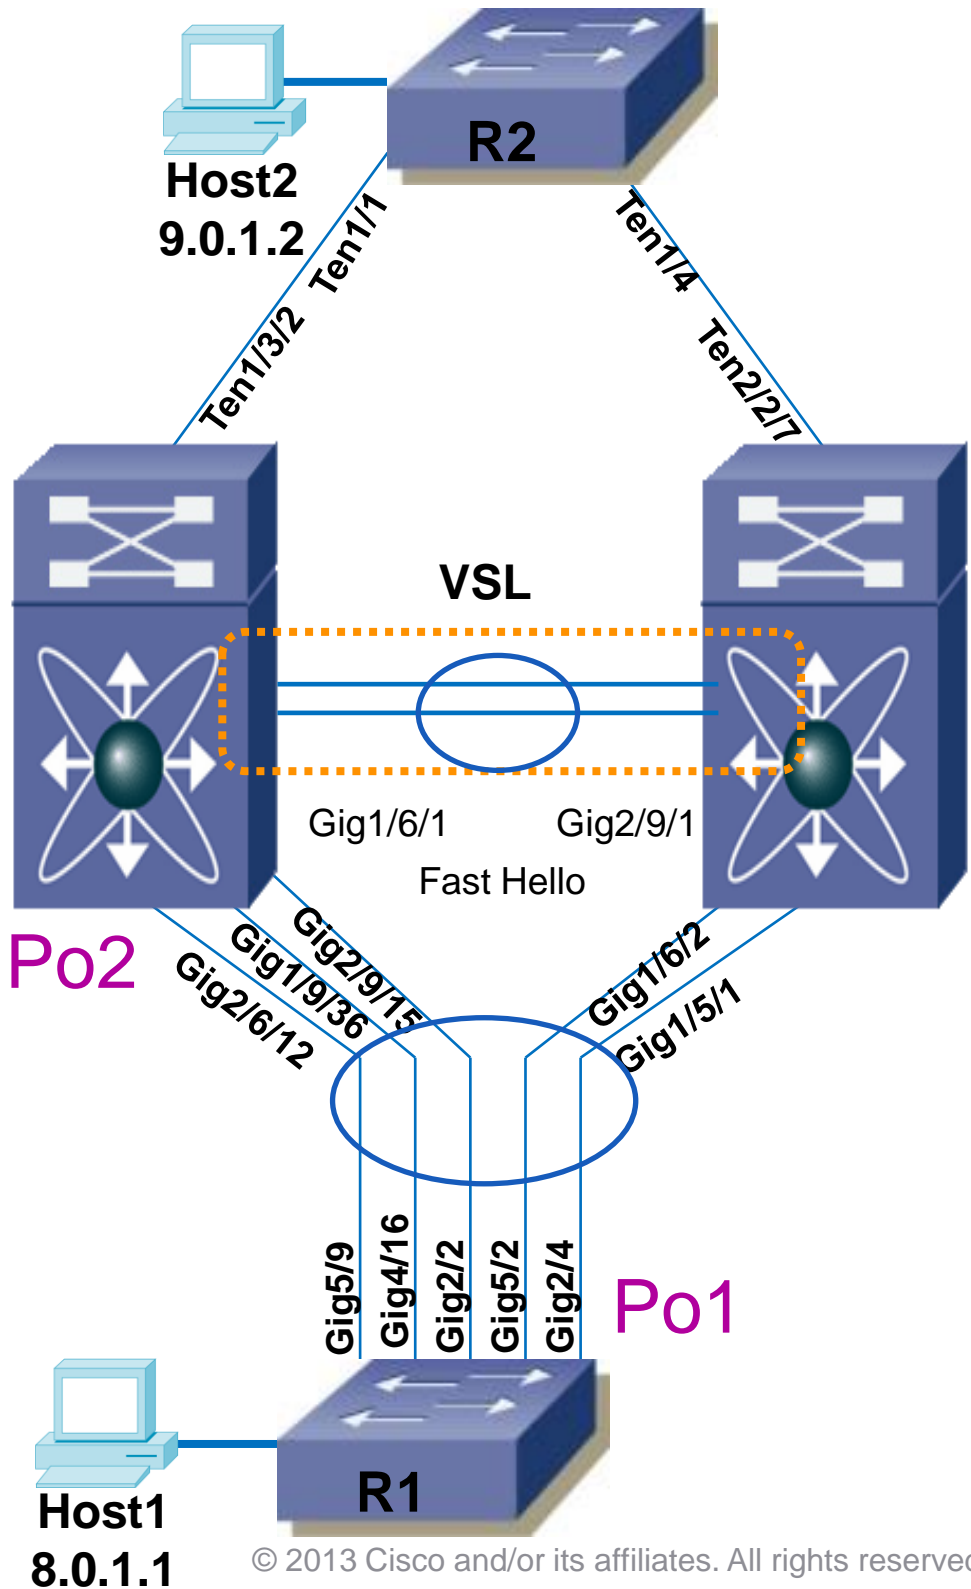
```
_. . .
```

In dual-active recovery mode: No

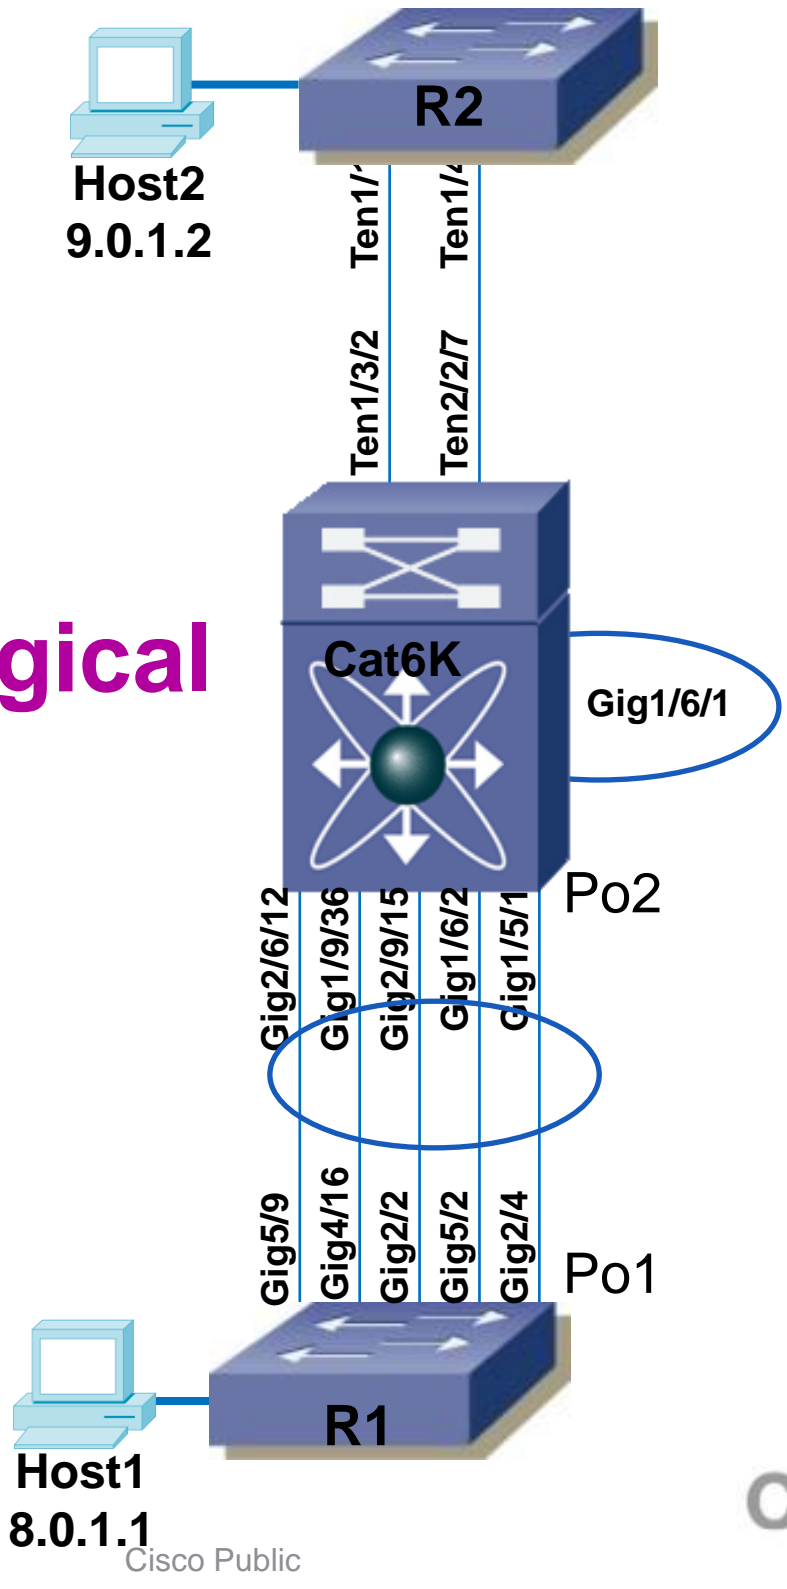
Switch is not in dual active recovery mode

VSS L2/L3 Forwarding Path

Physical




Logical



VSS Data Plane Forwarding Path

VSS Data Plane Design: Minimise Load on VSL

- **Multi-chassis Ether Channel (MEC):**
 - Channel hash is modified on VSS so that local links in the MEC are preferred over links on remote switch
- **Equal Cost Multi Path (ECMP):**
 - Adjacency table is modified on VSS to prefer next hops attached to local switch



VSS L2/L3 Forwarding

VSS Data Plane Troubleshooting L2 MEC

- Verify the load-balancing algorithm used

VSS specific commands augmented with switch id

```
Cat6K#show etherchannel load-balance switch 2 module 2
EtherChannel Load-Balancing Configuration:
    src-dst-ip enhanced
    mpls label-ip
EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
    IPv4: Source XOR Destination IP address
    IPv6: Source XOR Destination IP address
MPLS: Label or IP
```

- Identify the physical path for flow from host 2 → host 1

```
Cat6K#show etherchannel load-balance hash-result interface Port-channel 2 switch 1 ip 9.0.1.2 vlan 705
8.0.1.1
Computed RBH: 0x6
Would select Gi1/6/2 of Po2
Cat6K#show etherchannel load-balance hash-result interface Port-channel 2 switch 2 ip 9.0.1.2 vlan 705
8.0.1.1
Computed RBH: 0x6
Would select Gi2/9/15 of Po2
```

Important: Only use parameters consistent with the configured load-balancing algorithm. Command uses all the specified arguments to calculate the hash.

VSS L2/L3 Forwarding

- Routing table shows two equal cost paths to 9.0.0.0/8

```
Cat6K#show ip route 9.0.0.0 | i via
Known via "eigrp 101", distance 90, metric 3072, type internal
Redistributing via eigrp 101
 7.7.1.2, from 7.7.1.2, 1d00h ago, via TenGigabitEthernet2/2/7
* 7.6.1.2, from 7.6.1.2, 1d00h ago, via TenGigabitEthernet1/3/2
```

- Looking at the HW table shows next hop directly attached to local switch is preferred

```
Cat6K#show mls cef lookup 9.0.1.0 switch 1 mod 3
Codes: decap - Decapsulation, + - Push Label
Index Prefix Adjacency
108775 9.0.0.0/8 Te1/3/2 , 000f.35ed.7c00
```

Packet coming in on switch 1 module 3, for 9.0.0.0/8 prefers next hop attached to local switch id 1

```
Cat6K#show mls cef lookup 9.0.1.0 switch 2 mod 2
Codes: decap - Decapsulation, + - Push Label
Index Prefix Adjacency
108775 9.0.0.0/8 Te2/2/7 , 000f.35ed.7c00
```

Packet coming in on switch 2 module 2, for 9.0.0.0/8 prefers next hop attached to local switch id 2

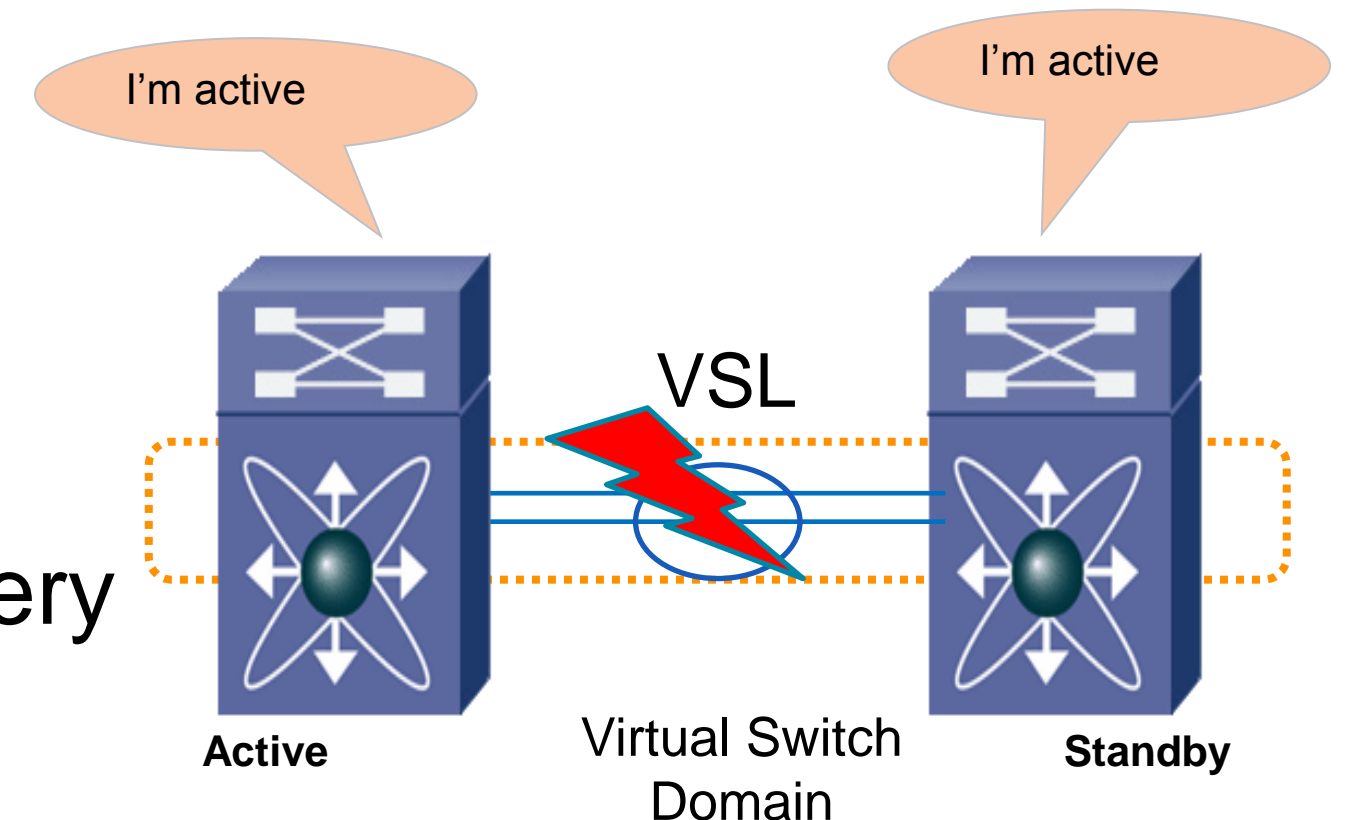
```
Cat6K#show mls cef exact-route 8.0.1.1 0 9.0.1.2 0 switch 1 mod 3
Interface: Te1/3/2, Next Hop: 7.6.1.2, Vlan: 4064, Destination Mac: 000f.35ed.7c00
```

```
Cat6K#show mls cef exact-route 8.0.1.1 0 9.0.1.2 0 switch 2 mod 2
Interface: Te2/2/7, Next Hop: 7.7.1.2, Vlan: 4056, Destination Mac: 000f.35ed.7c00
```

VSS Dual-Active Detection

VSL Failure

- Both switches assume the active role
- Leads to two independent routers with identical configurations
- Three methods for detection and recovery
 - VSLP Fast Hello
 - Enhanced Port Aggregation Protocol (ePAgP)
 - Bi-directional Forwarding Detection (BFD)



VSS Dual-Active Detection

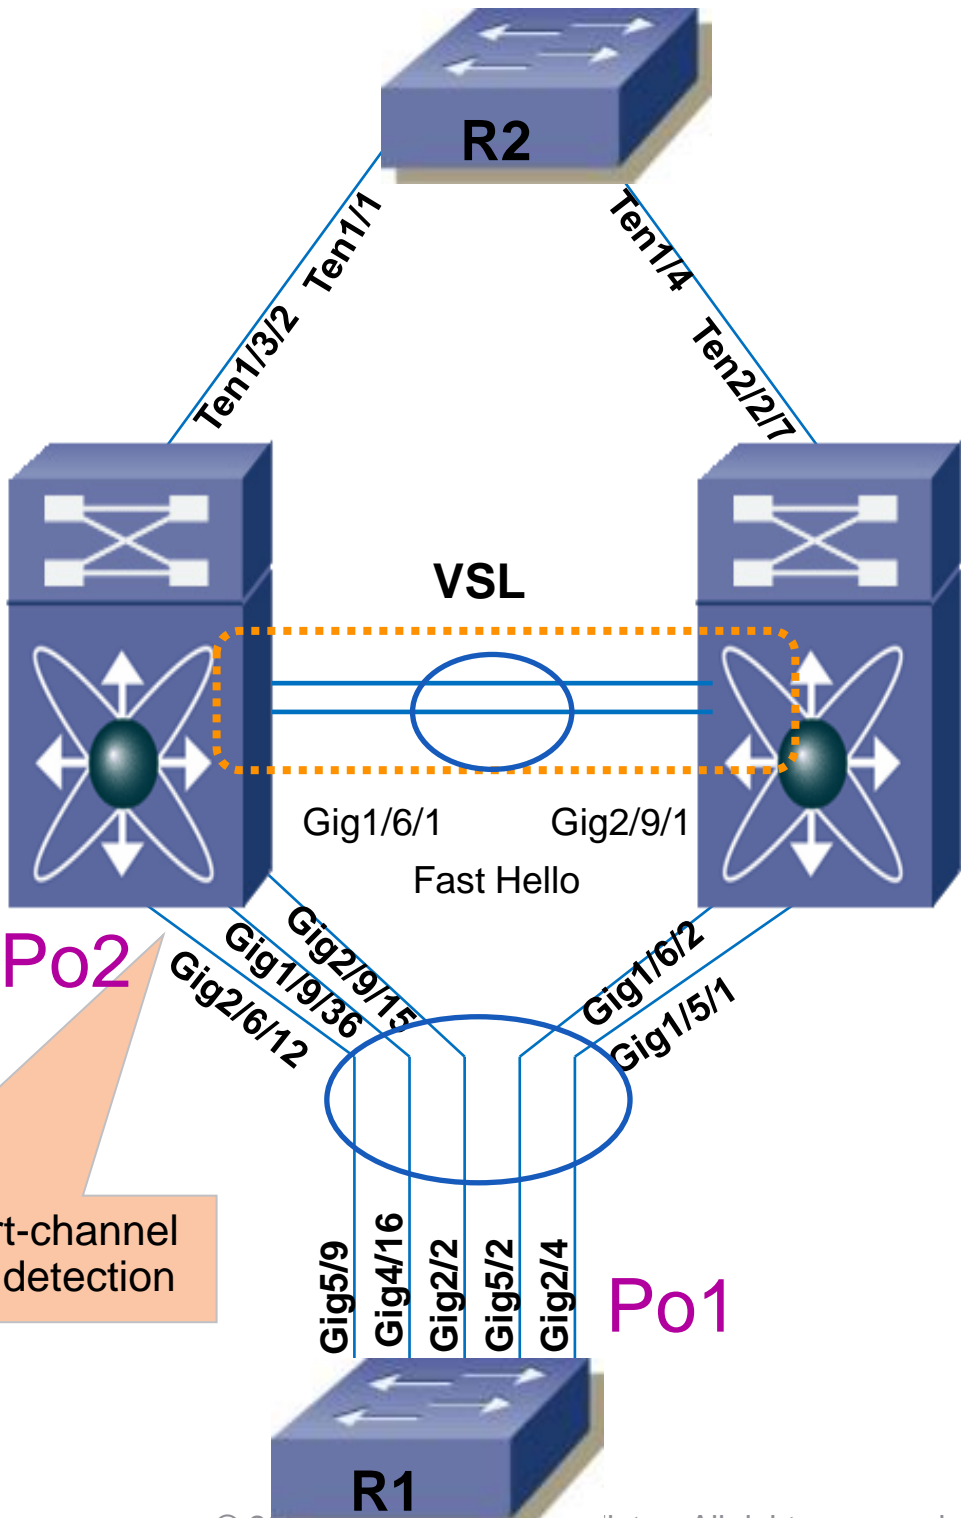
VSS Dual Active Detection Methods

- VSLP Fast Hello: (available in 12.2(33)SXI onwards)
 - Uses a dedicated heartbeat link between the VSL peers to send periodic hellos
 - Received hellos contain remote VSL peer's current state and role
- Enhanced PAgP
 - Uses new Type-Length-Value (TLV) field in PAgP packets. contains the ID of the active switch
 - Requires Multi-chassis EtherChannel (MEC) switch support ePAgP
- BFD
 - Uses a dedicated connection between VSL peers
 - Requires IP addresses in different subnet on the interfaces in the BFD link

VSS Dual-Active Detection

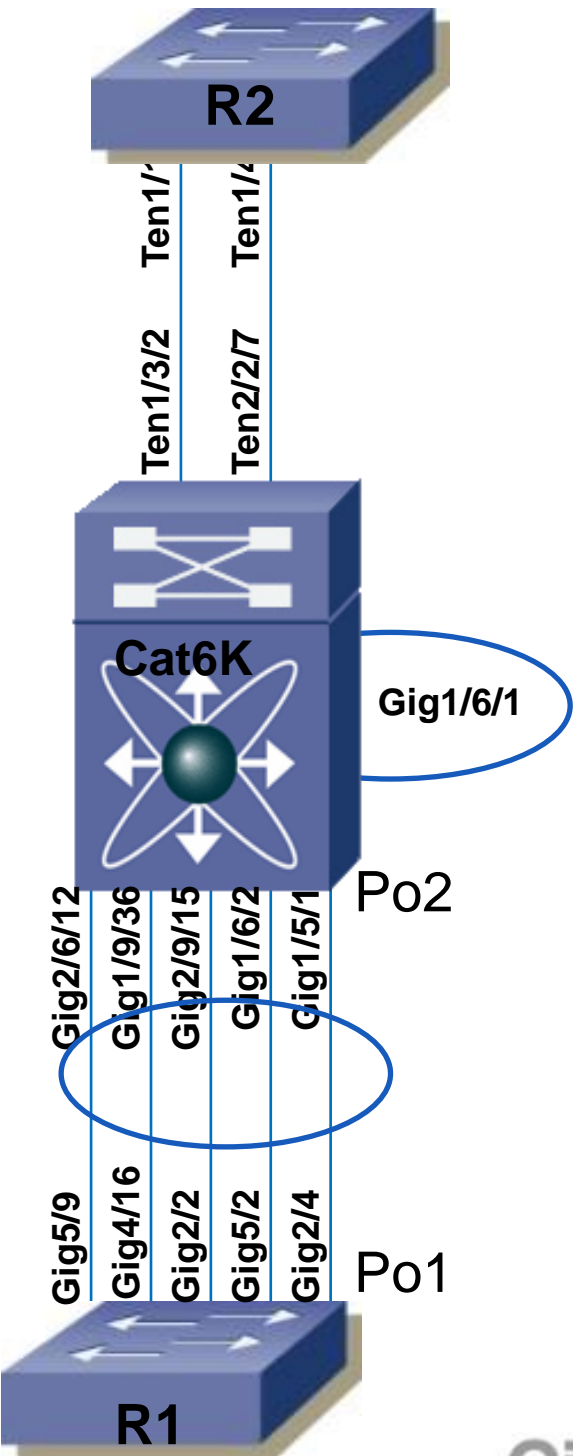
VSS Dual Active Detection Setup

Physical



Port-Channel 2 is the only port-channel trusted for ePAGP dual-active detection

Logical



VSS Dual-Active Detection

VSS Dual Active Detection Troubleshooting

```
Cat6K#show startup-config | begin switch virtual
switch virtual domain 1
  switch mode virtual
  ...
  no dual-active detection bfd
  dual-active detection pagp trust channel-group 2
  dual-active exclude interface GigabitEthernet1/5/3
  dual-active exclude interface GigabitEthernet2/5/3
  !
  ...
interface GigabitEthernet1/6/1
  no switchport
  no ip address
  dual-active fast-hello
  ...
interface GigabitEthernet2/9/1
  no switchport
  no ip address
  dual-active fast-hello
```

BFD Dual Active detection is explicitly disabled

Enhanced PAgP messages are only sent on channels configured in trust mode

Interfaces can be excluded from recovery mode. They will not go down when the switch goes into recovery mode

VSS Dual-Active Detection

VSS Dual Active Detection Troubleshooting

```
Cat6K#show switch virtual dual-active summary
```

```
Pagp dual-active detection enabled: Yes
```

```
Fast-hello dual-active detection enabled: Yes
```

```
Interfaces excluded from shutdown in recovery mode:
```

```
Gi1/5/3
```

```
Gi2/5/3
```

```
In dual-active recovery mode: No
```

```
Cat6K#show switch virtual dual-active pagp
```

```
PAGP dual-active detection enabled: Yes
```

```
PAGP dual-active version: 1.1
```

```
Channel group 2 dual-active detect capability w/nbrs
```

```
Dual-Active trusted group: Yes
```

Port	Dual-Active Detect	Capable	Partner Name	Partner Port	Partner Version
Gi1/5/1	Yes		R1	Gi2/4	1.1
Gi1/6/2	Yes		R1	Gi5/2	1.1

```
.....
```

Shows that both enhanced PAGP and fast hello methods of dual active detection are enabled

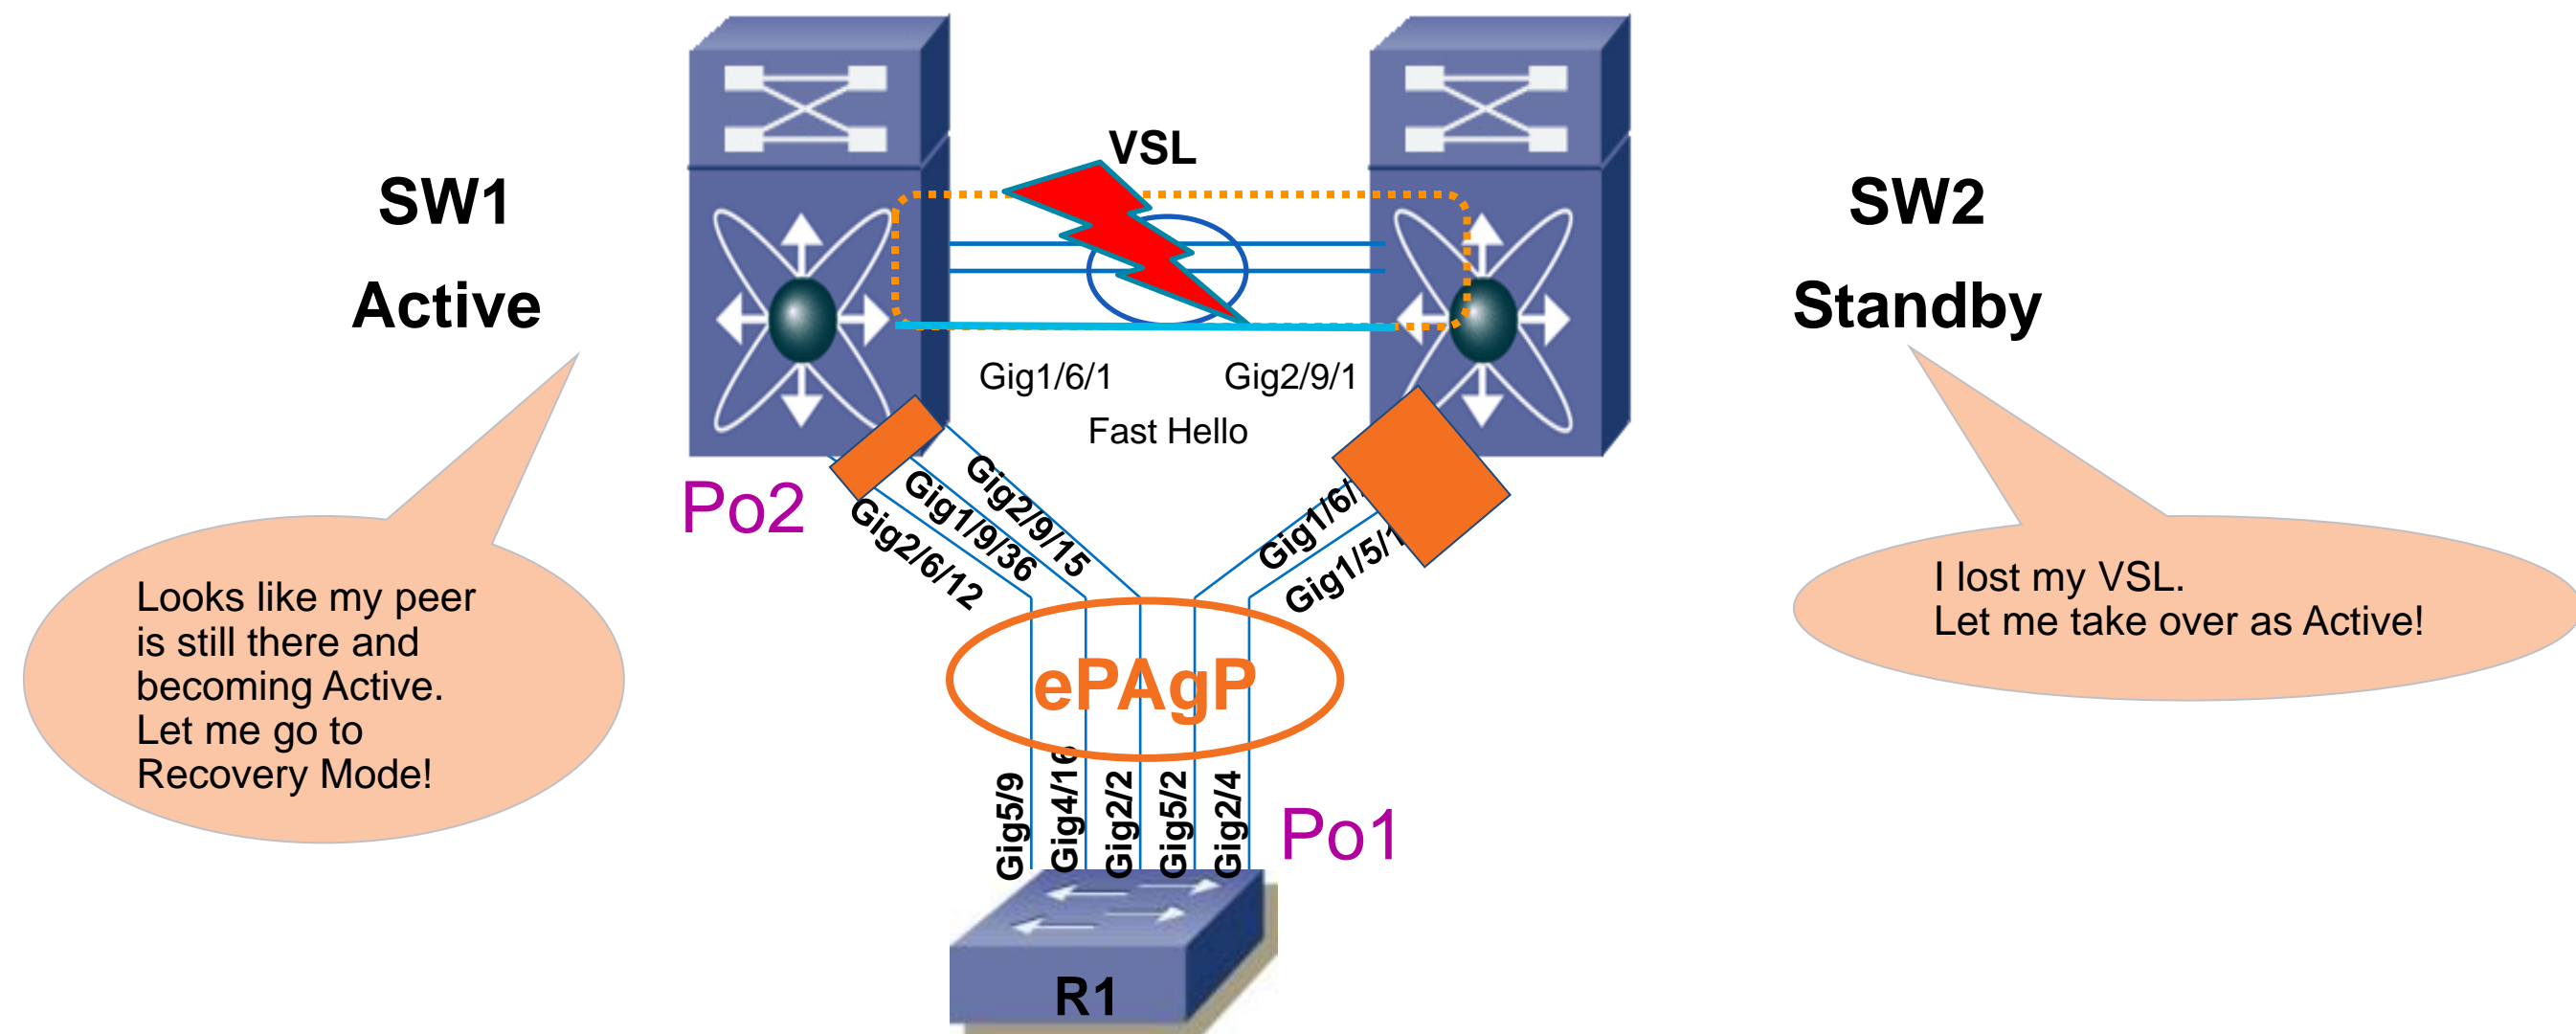
Port channel 2 is trusted for Enhanced PAGP dual active detection ... at least 1 trusted port channel needed !!

Check that the neighbour runs a SW version that supports Enhanced PAGP ... if not, no dual active detection !!

VSS Dual-Active Detection

VSS Dual Active Detection Troubleshooting

- Recovery Mode: all non-VSL and all non-excluded ports will be internally shutdown



VSS Dual-Active Detection

VSS Dual Active Detection Troubleshooting

- Trigger dual active situation by bringing down the VSL

Enhanced PAGP detected both switch ids were active at the same time

```
*Apr 1 12:40:22.885 CET: %PAGP_DUAL_ACTIVE-SW1_SP-1-RECOVERY: PAgP running on Gi1/5/1 triggered dual-active recovery: active id 0011.bc75.4400 received, expected 0011.5d54.6800
*Apr 1 12:40:22.945 CET: %DUAL_ACTIVE-SW1_SP-1-DETECTION: Dual-active condition detected: all non-VSL and non-excluded interfaces have been shut down
```

Original active switch (switch 1 in example) goes into recovery mode

- On switch in recovery mode, do “show ip interface brief | inc up” to check the interfaces that are up

```
Apr 1 12:40:20.096 CET: %VSLP-SW2_SPSTBY-2-VSL_DOWN: All VSL links went down while switch is in Standby role
*Apr 1 12:40:20.096 CET: %DUAL_ACTIVE-SW2_SPSTBY-1-VSL_DOWN: VSL is down switchover, or possible dual-active situation has occurred
*Apr 1 12:40:20.100 CET: %PFREDUN-SW2_SPSTBY-6-ACTIVE: Initializing as Virtual Switch ACTIVE processor
```

Original Standby switch goes into Active mode

VSS Dual-Active Detection

VSS Dual Active Detection Troubleshooting

```
Cat6K#show switch virtual role
```

Switch	Switch	Status	Preempt	Priority	Role	Session ID	
	Number		Oper (Conf)	Oper (Conf)		Local	Remote

```
-----  
LOCAL 1 UP TRUE (Y*) 200(200) ACTIVE 0 0
```

```
Active configured preempt timer(switch 1): 5 minutes
```

```
In dual-active recovery mode: Yes
```

```
Triggered by: PAgP detection
```

```
Triggered on interface: Gi1/5/1
```

```
<snip>
```

On switch id 1: originally active, now in recovery mode

Does not see switch 2 (as VSL is still down)

Mechanism that detected dual active was Enhanced PAgP, via link 1/5/1

```
Cat6K#show switch virtual role
```

Switch	Switch	Status	Preempt	Priority	Role	Session ID	
	Number		Oper (Conf)	Oper (Conf)		Local	Remote

```
-----  
LOCAL 2 UP TRUE (Y*) 100(100) ACTIVE 0 0
```

```
Active configured preempt timer(switch 2): 5 minutes
```

```
In dual-active recovery mode: No
```

Switch id 2 is now the "real" active switch, and doesn't see switch id 1 as long as the VSL is down

VSS Dual-Active Detection

Bringing up the VSL

```
*Apr  1 12:49:29.513 CET: %DUAL_ACTIVE-1-VSL_RECOVERED: VSL has recovered during dual-active situation:
Reloading switch 1

*Apr  1 12:49:29.513 CET: %VS_GENERIC-5-VS_CONFIG_DIRTY: Configuration has changed. Ignored reload request
until configuration is saved

*Apr  1 12:49:32.781 CET: %LINK-SW2_SP-3-UPDOWN: Interface TenGigabitEthernet2/5/4 changed state to up

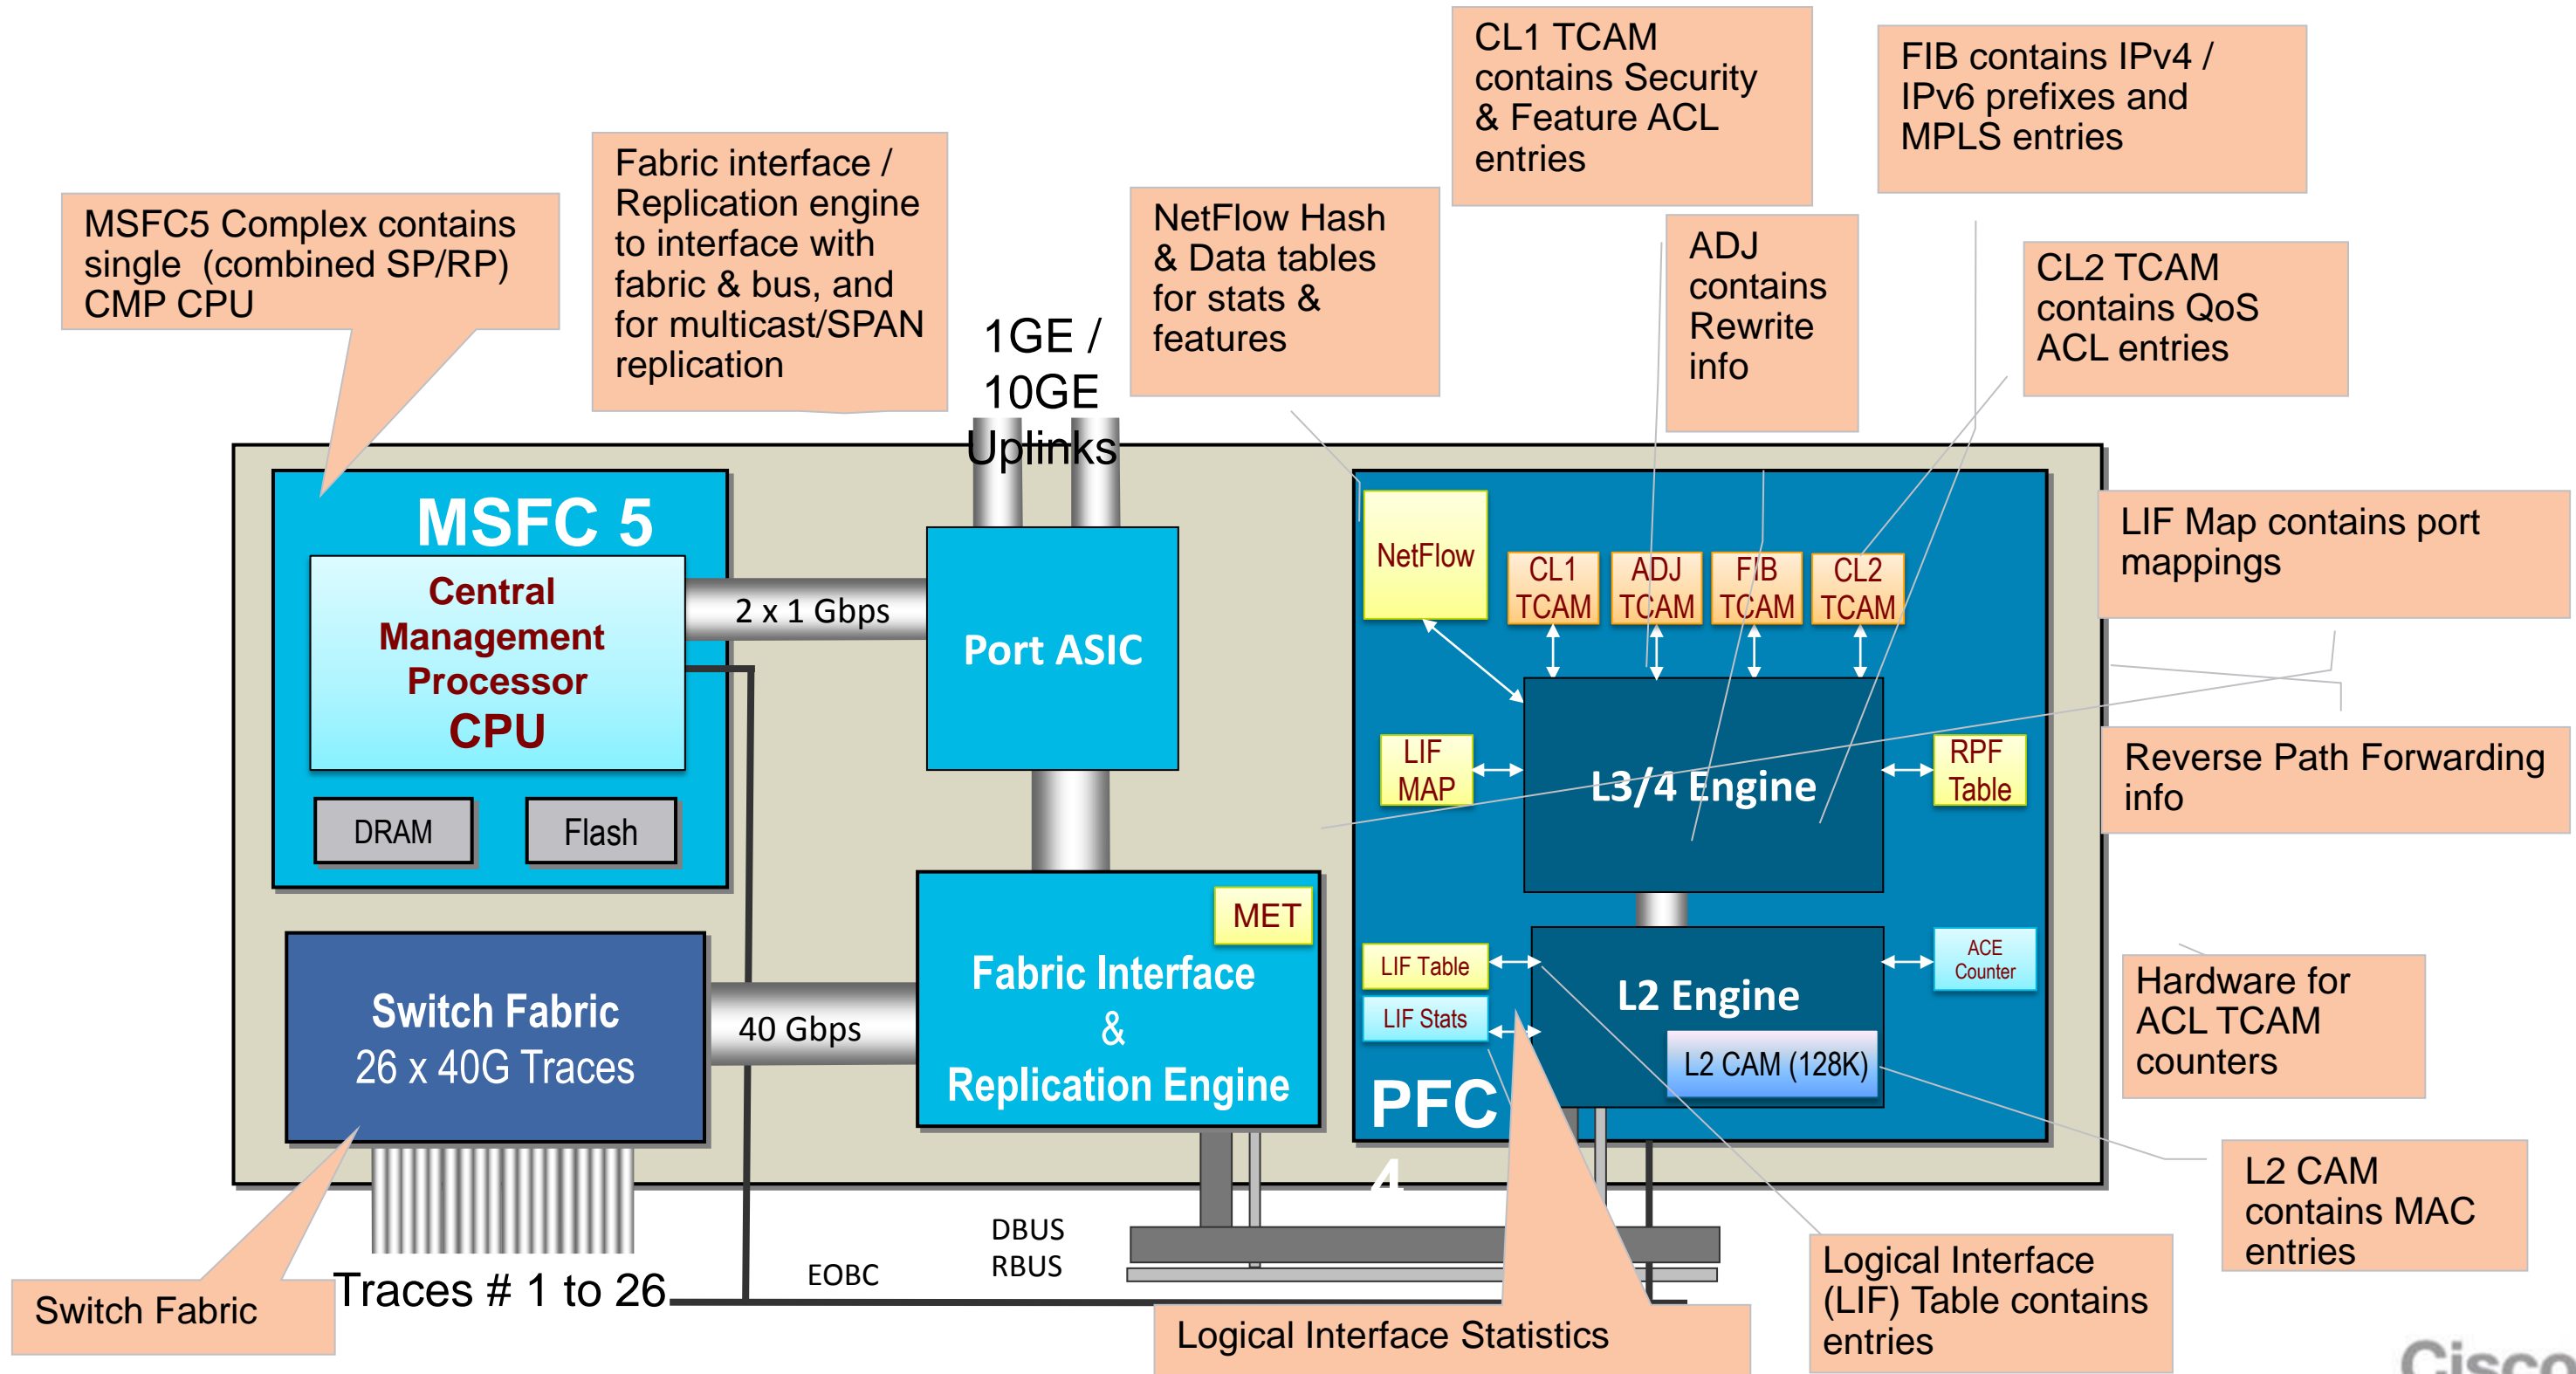
*Apr  1 12:49:49.128 CET: %VSLP-SW2_SP-5-VSL_UP: Ready for Role Resolution with Switch=1,
MAC=0011.5d54.6800 over Te2/2/6

*Apr  1 12:49:50.320 CET: Initializing as Virtual Switch ACTIVE processor

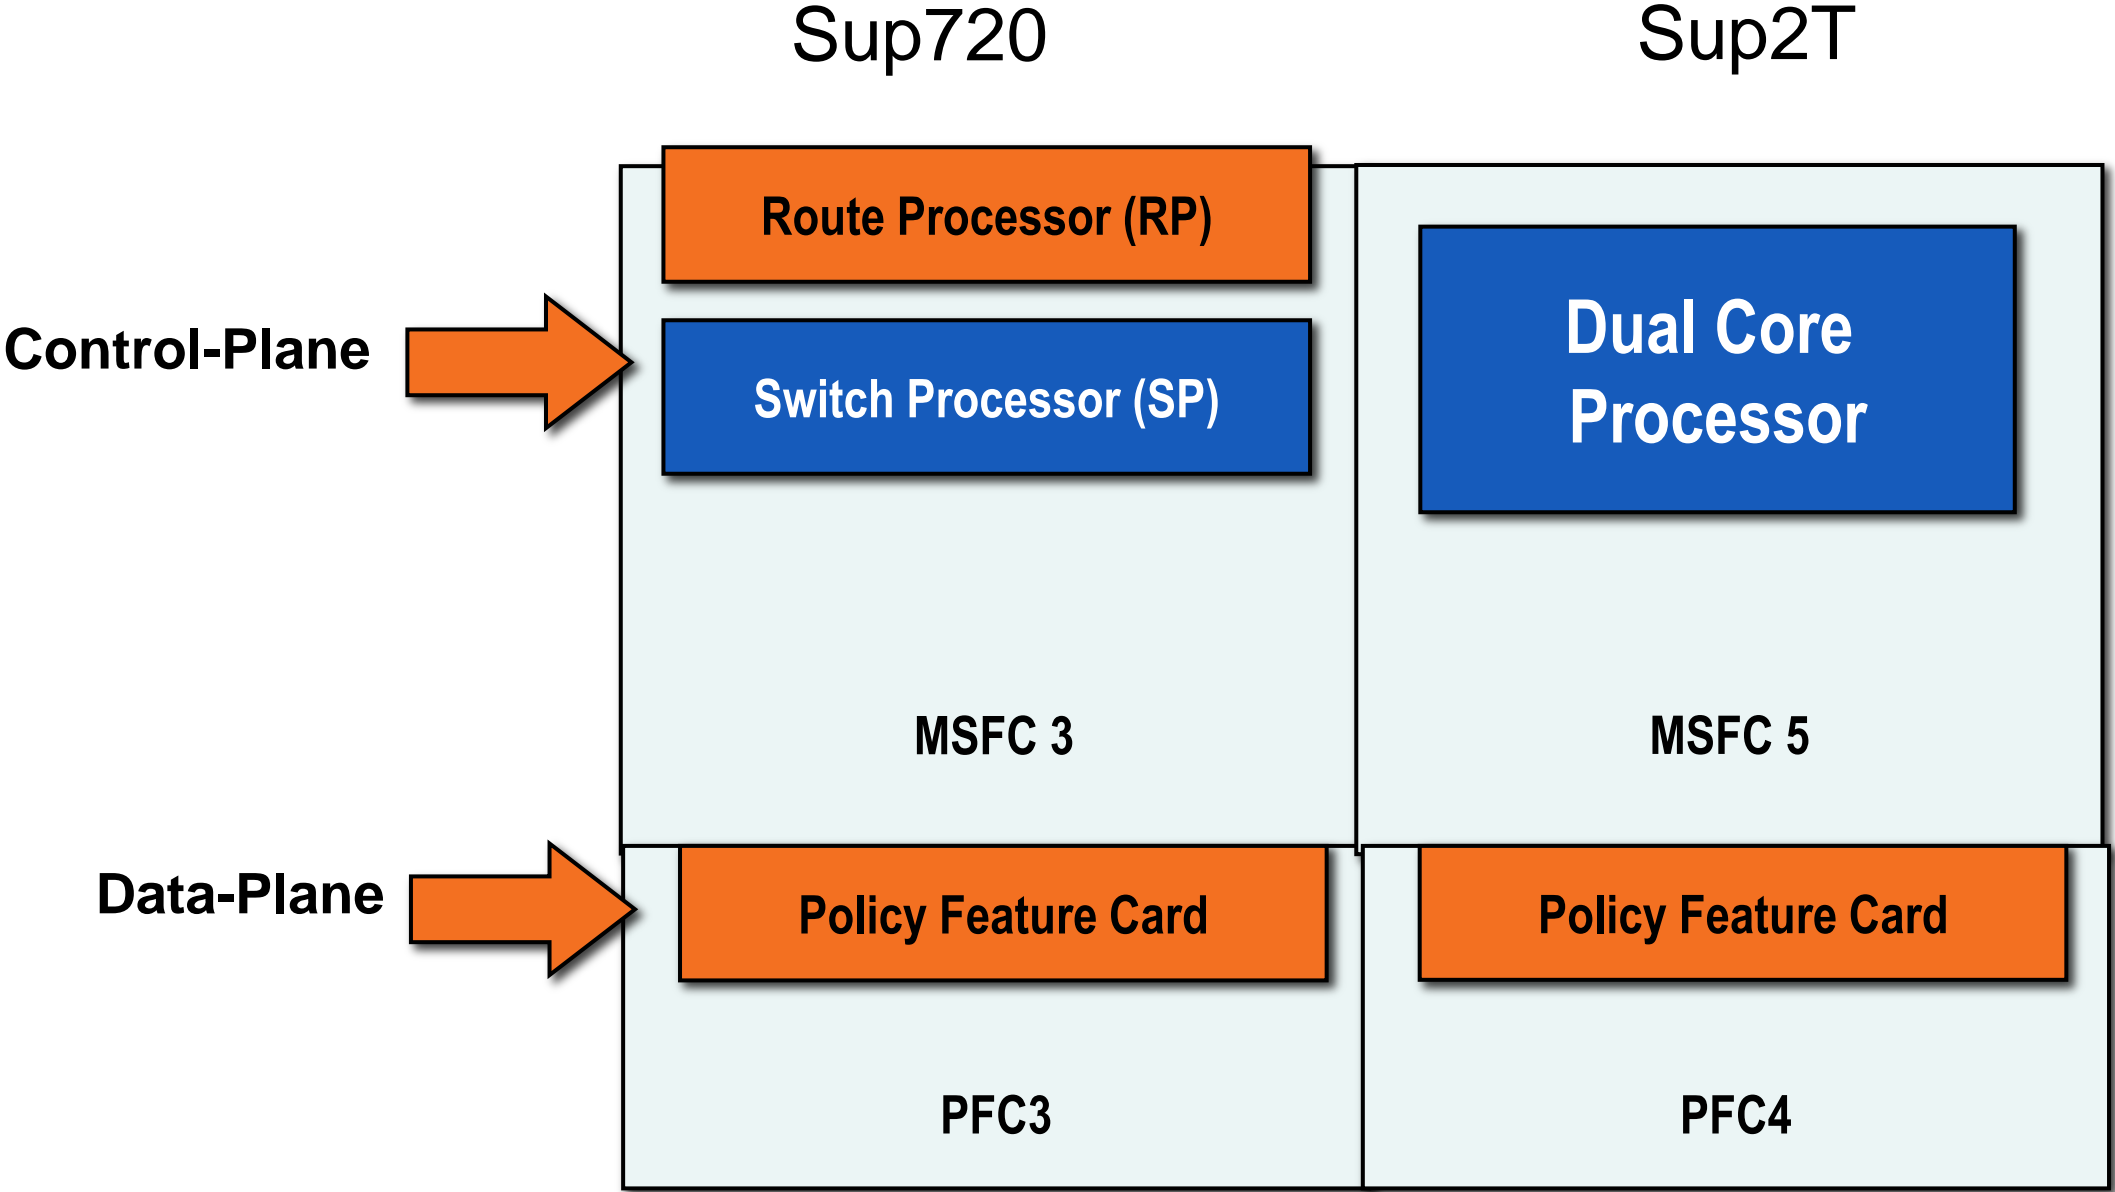
*Apr  1 12:49:52.140 CET: %VSLP-SW2_SP-5-RRP_MSG: Peer Switch with unsaved configurations needs to be
reloaded. Please save relevant configurations on the peer switch and reload it.
```

- Configuration deemed “DIRTY” since configuration mode was entered and configuration was not written. Do NOT change configuration when the switch is in Recovery mode.
- Configuration must be saved MANUALLY and reloaded to bring the switch back up into the VSS

Supervisor 2T/PFC4 Architecture



Sup720 vs. Sup2T Overview



Sup2T Overview

- Integrated with Control Management Processor (CMP)
- USB-based console and flash support
- Key improvement: Single-Cycle Input & Output Processing
- Has superior Control-Plane protection
- Supports all Sup720 features, including VSS

L3 Packet Flow Troubleshooting

Verify IP Route

```
Sup2T#show ip route 192.168.200.5
```

```
Routing entry for 192.168.200.4/30
```

```
Known via "static", distance 1, metric 0
```

```
Routing Descriptor Blocks:
```

```
  192.168.25.5
```

```
    Route metric is 0, traffic share count is 1
```

```
* 192.168.25.1
```

```
    Route metric is 0, traffic share count is 1
```

```
Sup2T#show ip route 192.168.200.9
```

```
Routing entry for 192.168.200.8/30
```

```
Known via "static", distance 1, metric 0
```

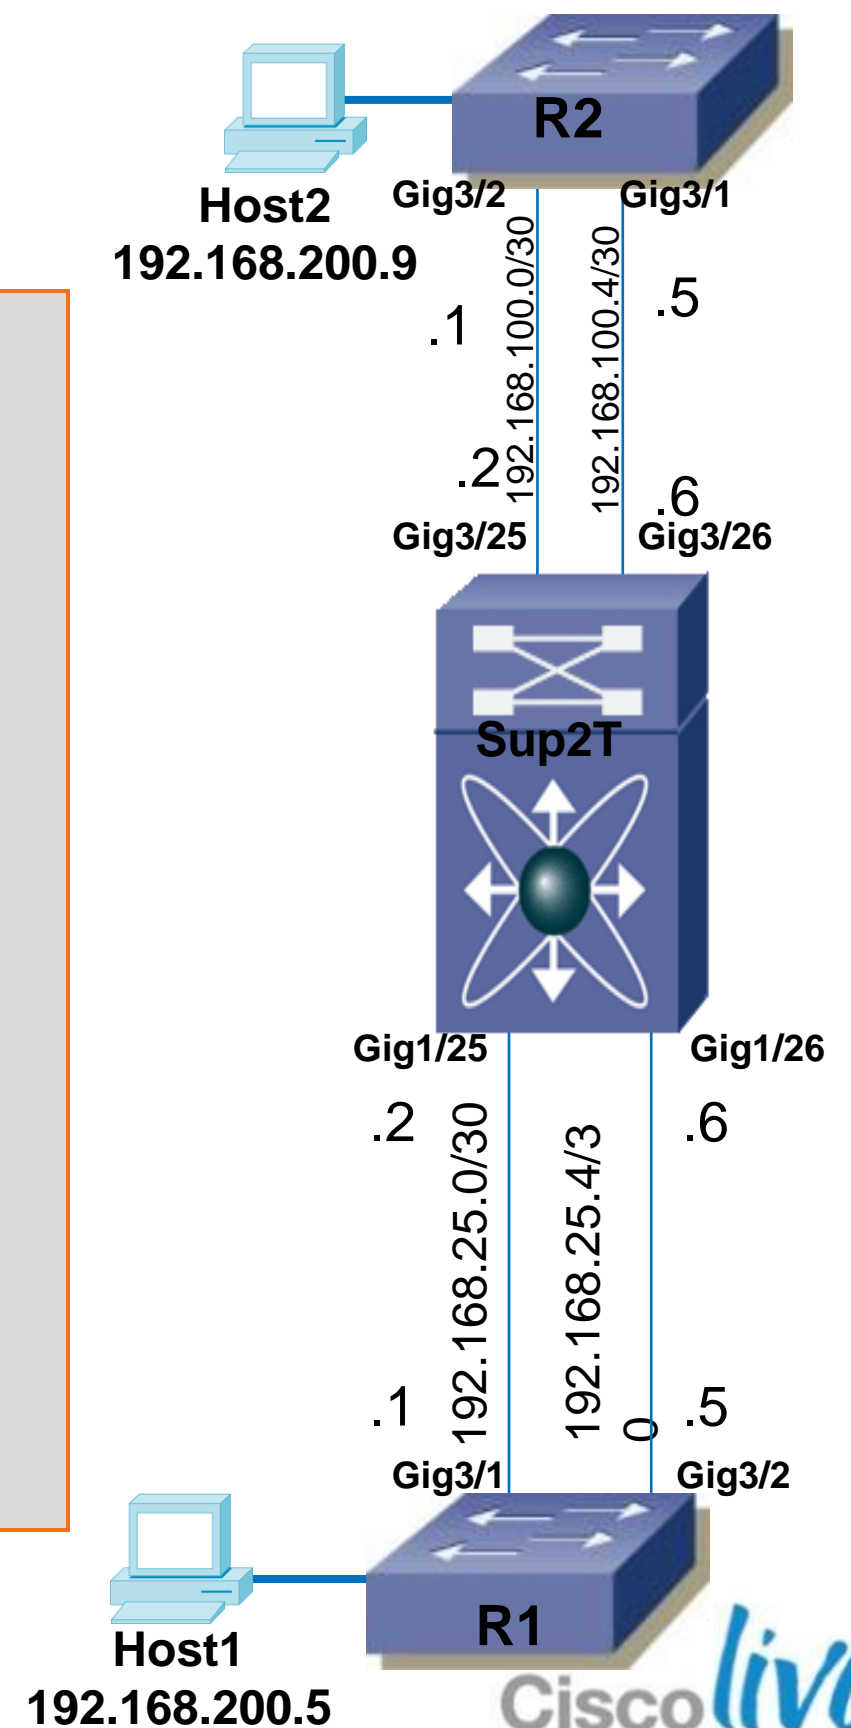
```
Routing Descriptor Blocks:
```

```
  192.168.100.5
```

```
    Route metric is 0, traffic share count is 1
```

```
* 192.168.100.1
```

```
    Route metric is 0, traffic share count is 1
```



L3 Packet Flow Troubleshooting

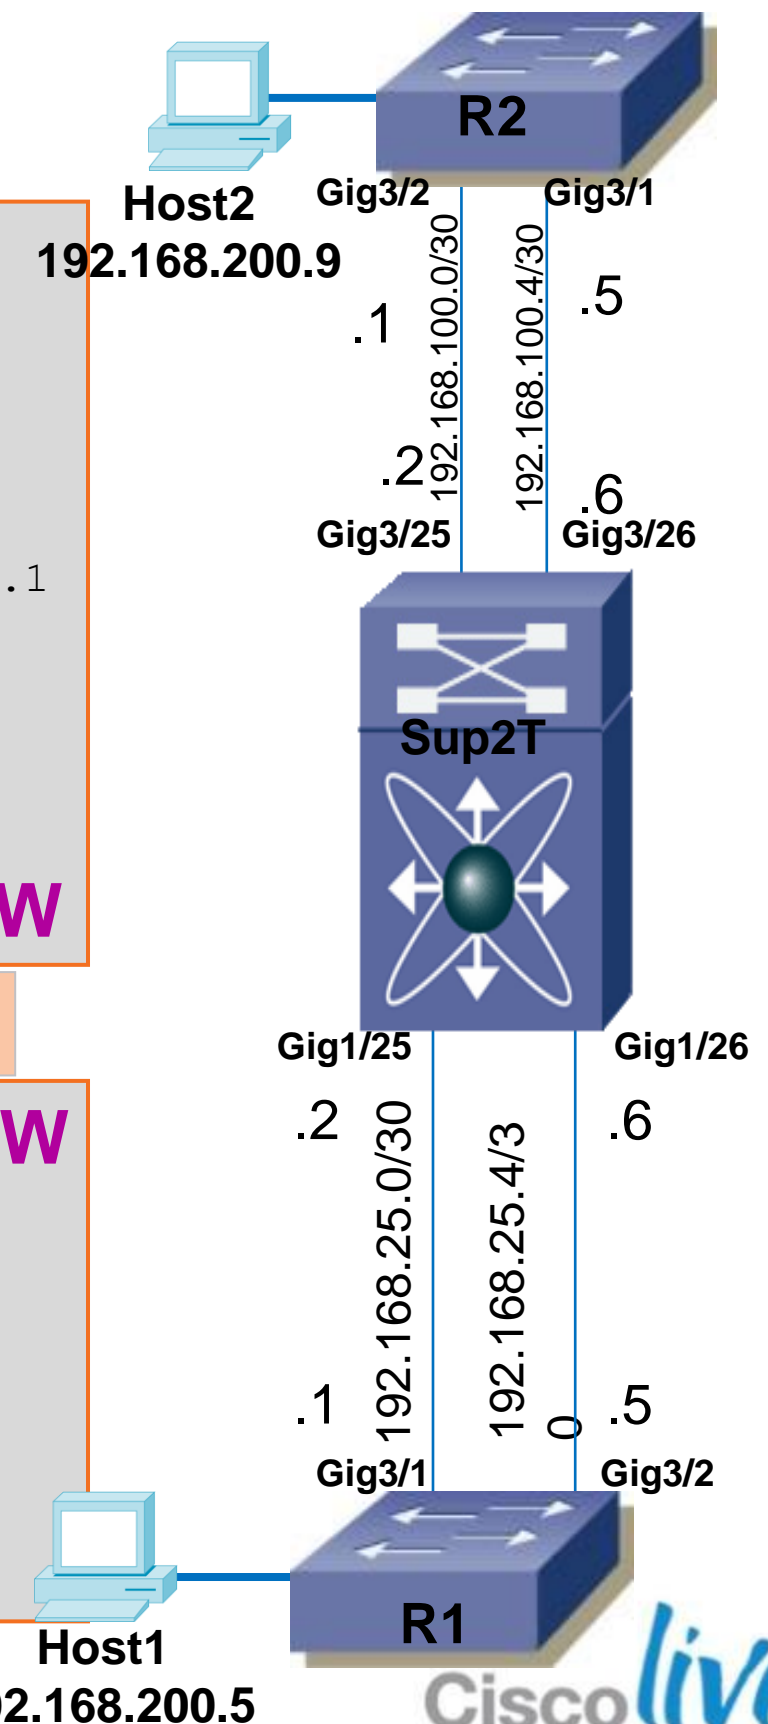
```
Sup2T#show ip cef 192.168.200.9
192.168.200.8/30
  nexthop 192.168.100.1 GigabitEthernet3/25
  nexthop 192.168.100.5 GigabitEthernet3/26
Sup2T#show ip cef exact-route 192.168.200.5 192.168.200.9
192.168.200.5 -> 192.168.200.9 => IP adj out of GigabitEthernet3/25, addr 192.168.100.1
Sup2T#show ip cef adjacency GigabitEthernet 3/25 192.168.100.1
192.168.100.1/32
  attached to GigabitEthernet3/25
192.168.200.8/30
  nexthop 192.168.100.1 GigabitEthernet3/25
```

SW

No more MLS. Use "platform hardware".

```
Sup2T#show platform hardware cef lookup 192.168.200.9
Codes: decap - Decapsulation, + - Push Label
Index Prefix Adjacency
219265 192.168.200.8/30 Gi3/25 ,001d.e656.cc00 (Hash: 007F)
      Gi3/26 ,001d.e656.cc00 (Hash: 7F80)
Sup2T#show platform hardware cef exact-route 192.169.200.5 0 192.168.200.9 0
Interface: Gi3/26, Next Hop: 192.168.100.5, ifnum: 0x7E, Destination Mac:
001d.e656.cc00 LIF: 0x2000411B
```

HW



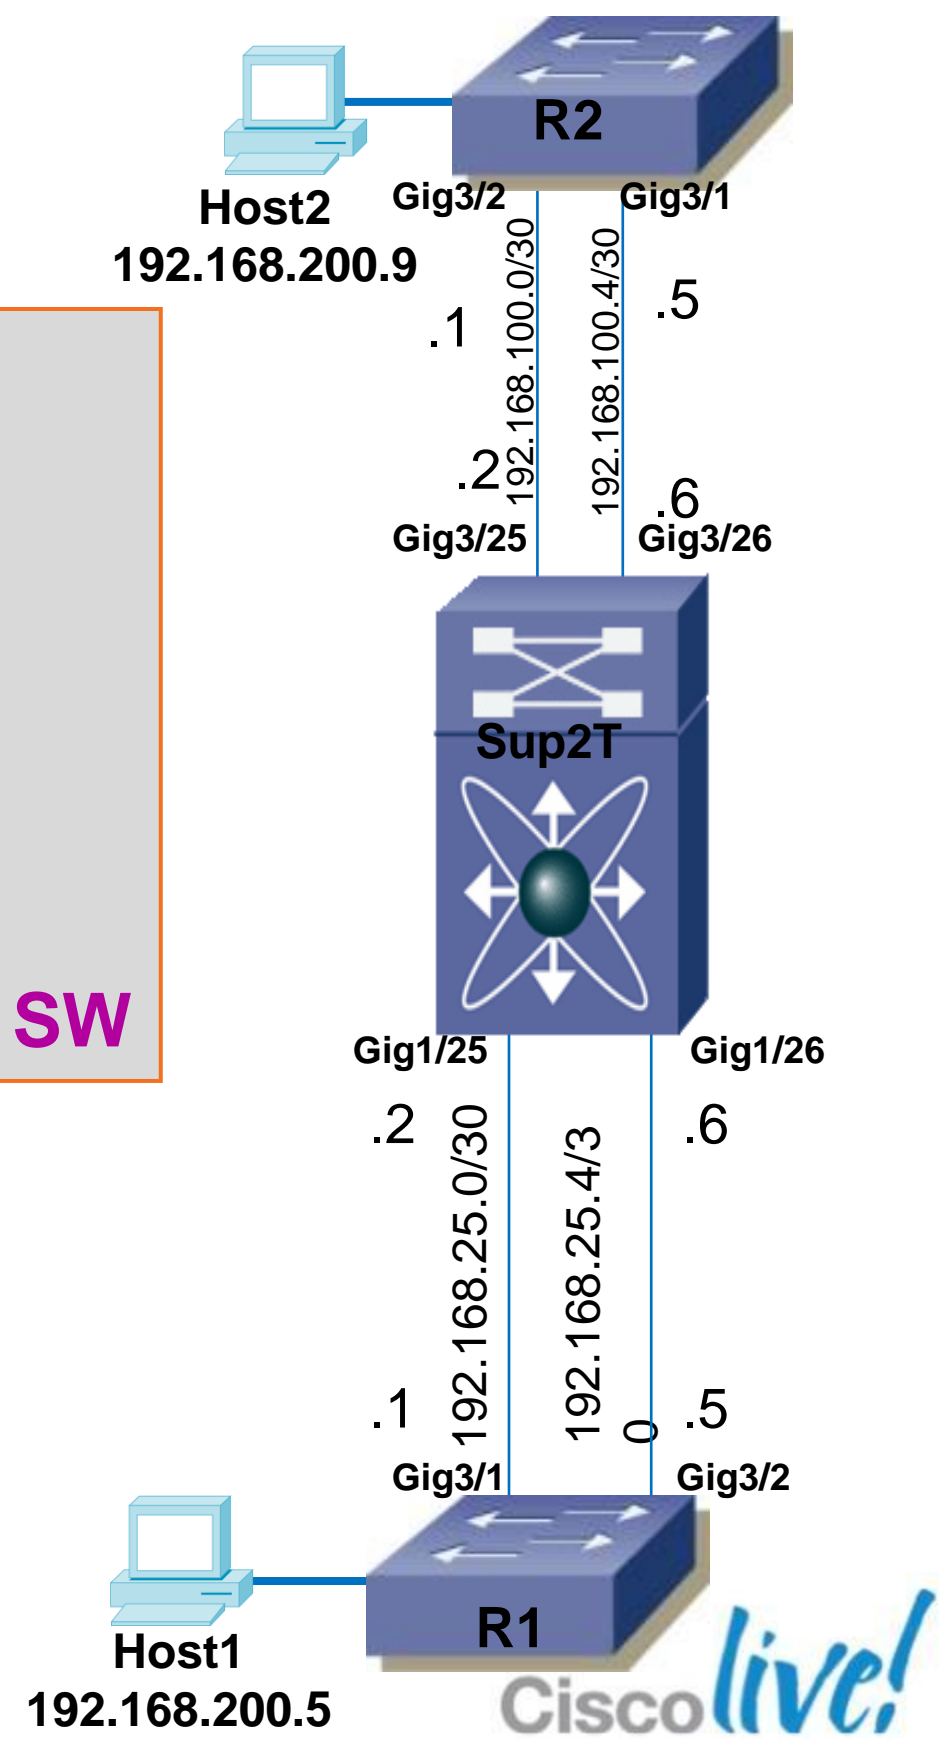
L3 Packet Flow Troubleshooting

L3 FIB Counters and Tables

```
Sup2T#show adjacency Gig3/25 192.168.100.1 detail
Protocol Interface          Address
IP      GigabitEthernet3/25  192.168.100.1 (13)
                2001 packets, 11682 bytes
                epoch 2
                sourced in sev-epoch 29
                Encap length 14
                001DE656CC0001DE65828000800
                L2 destination address byte offset 0
                L2 destination address byte length 6
                Link-type after encap: ip
                ARP
```

SW

Rewrite information (Dmac|Smac|0800): verify it is conform with next hop rewrite info



L3 Packet Flow Troubleshooting

L3 FIB Counters and Tables

```
Sup2T#show platform hardware cef ip 192.168.100.5 detail
```

HW

Codes: M - mask entry, V - value entry, A - adjacency index, NR- no_route bit
LS - load sharing count, RI - router_ip bit, DF: default bit
CP - copy_to_cpu bit, AS: dest_AS_number, DGTv - dgt_valid bit
DGT: dgt/others value

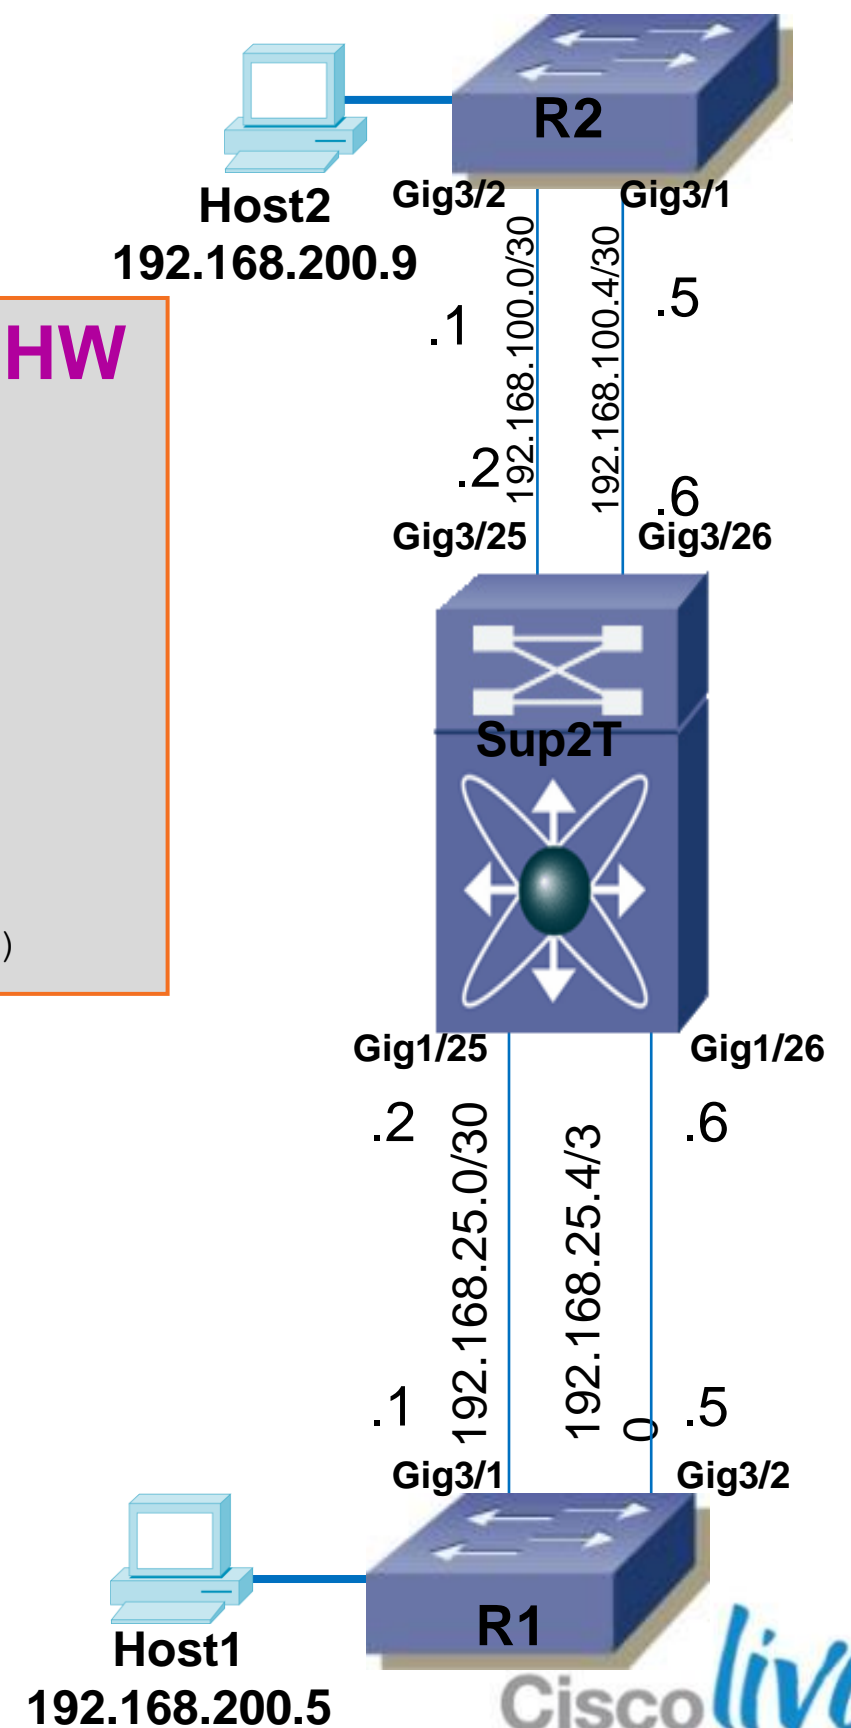
Format:IPV4 (valid class vpn prefix)

M(54794): 1 F 3FFF 255.255.255.255

V(54794): 1 0 0 192.168.100.5

(A:376833, LS:0, NR:0, RI:0, DF:0 CP:0 DGTv:1, DGT:0)

Start adjacency pointer is 376833



L3 Packet Flow Troubleshooting

L3 FIB Counters and Tables

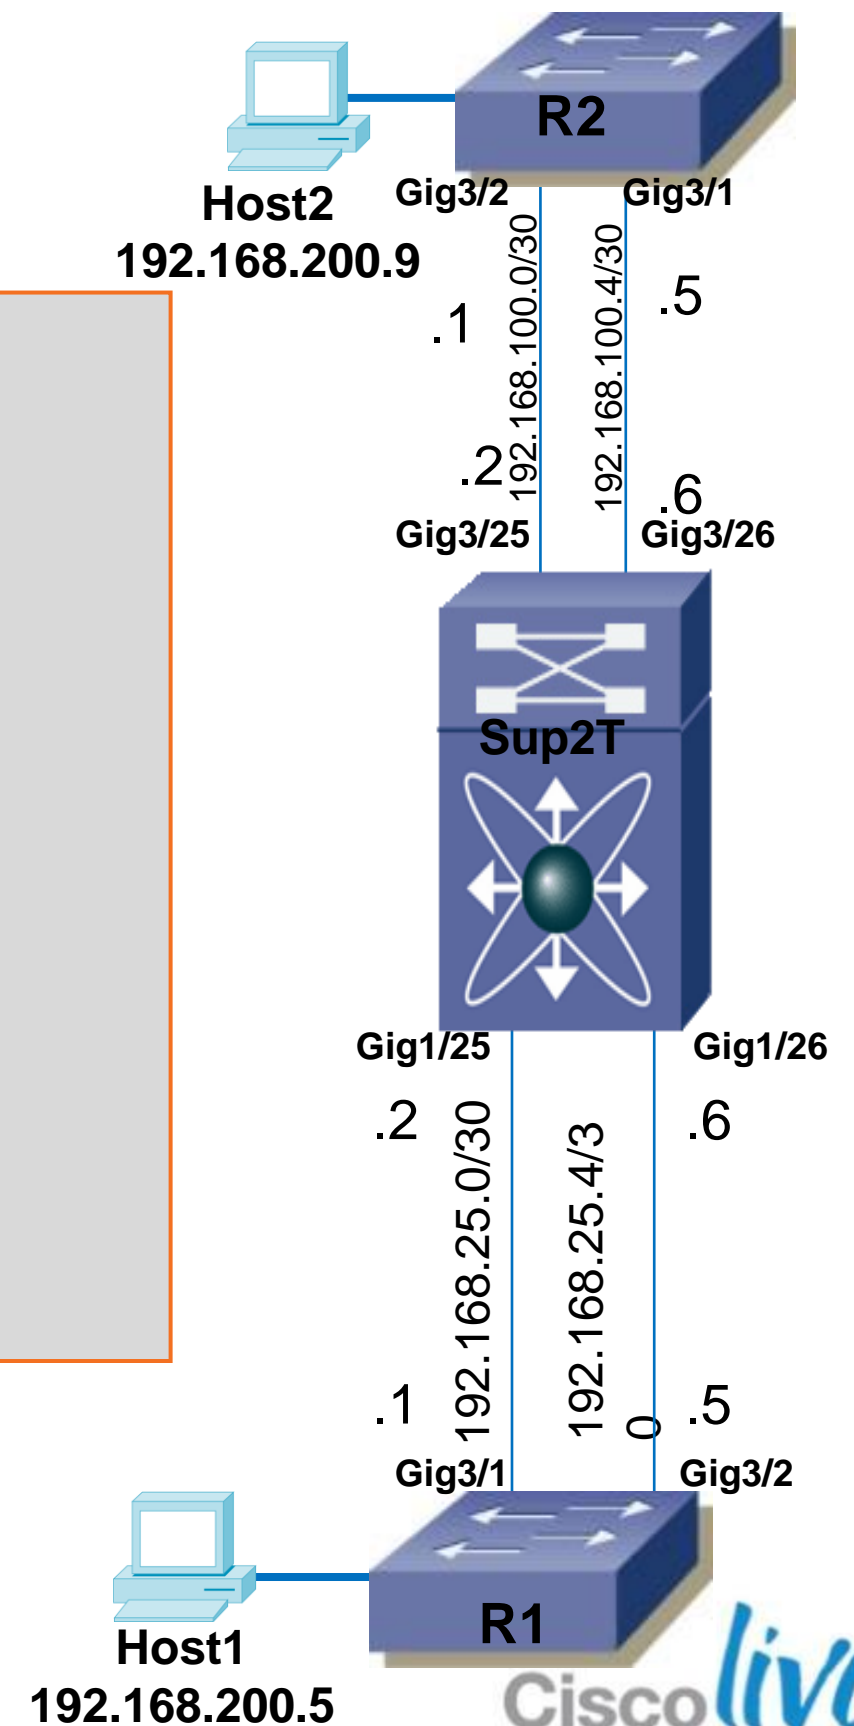
```
Sup2T#show platform hardware cef adjacencies entry 376833 detail
```

```
Index: 376833 -- Valid entry (valid = 1) -
```

```
Adjacency fields:
```

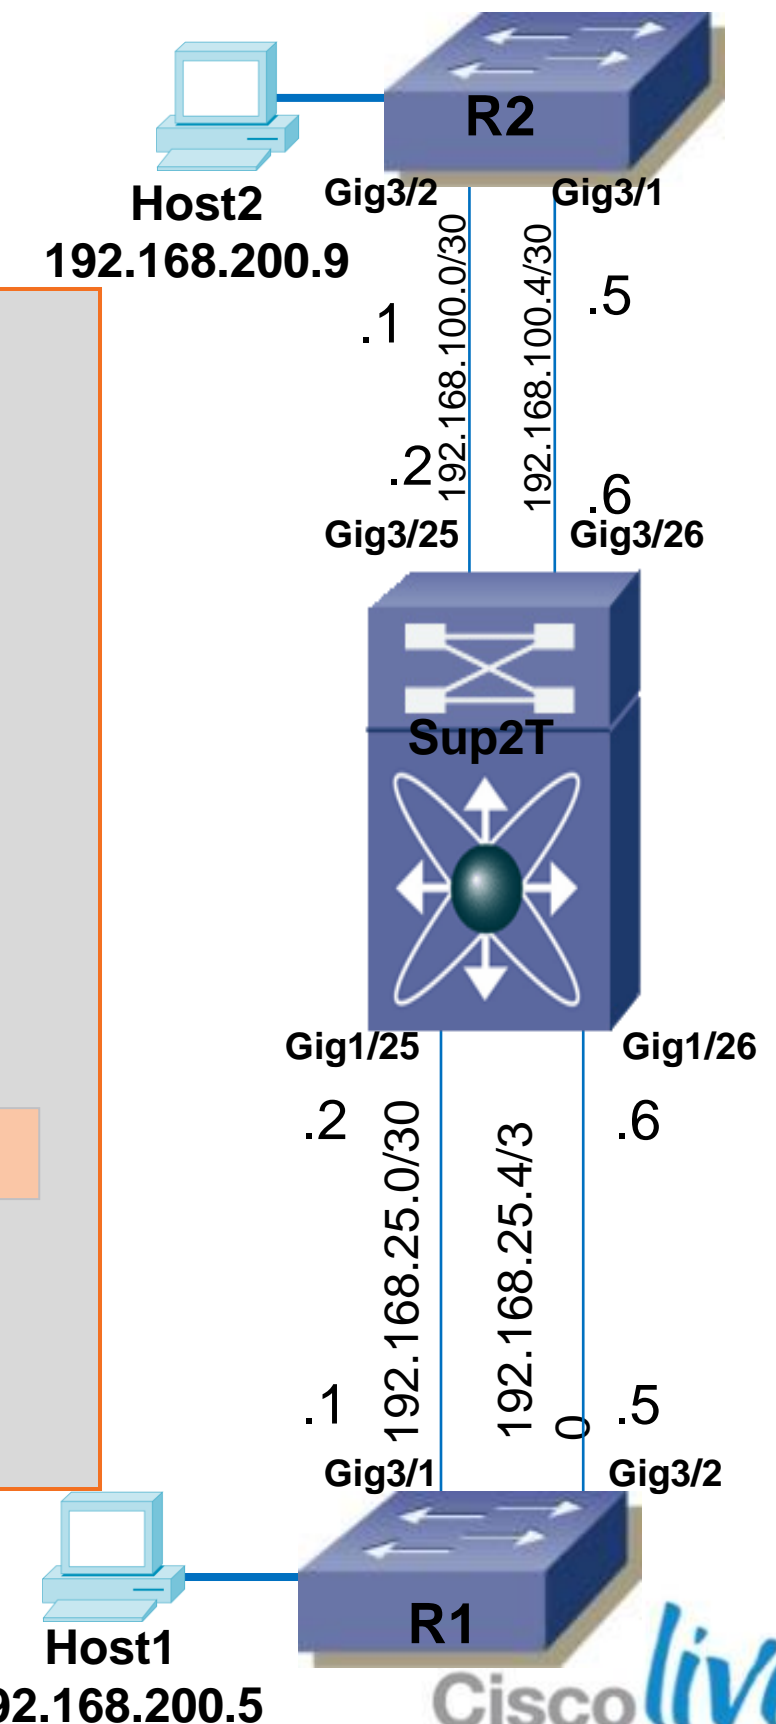
adj_stats = EN	fwd_stats = EN	trig = 0
l3_enable = ON (classify as Layer3)	age = 0	
format = IP	rdt = ON	ignr_emut = 0
vpn = 0x3FFF	elif = 0x411B	ri = 3
top_sel = 0	zone_enf = OFF	fltr_en = OFF
frr_te = OFF	idx_sel = 0	tnl_encap = 0
rw_hint = 0	ttl_control = 4	

```
. . . Continued . . .
```



L3 Packet Flow Troubleshooting

L3 FIB Counters and Tables



```

. . . Continued . . .

RIT fields: The entry has a Layer2 Format

|decr_ttl = YES      | pipe_ttl = 0      | utos = 0
|_____|_____|_____
|l2_fwd = 0         | rmac = 0         | ccc = L3_REWRITE
|_____|_____|_____
|rm_null_lbl = YES | rm_last_lbl = YES| pv = 0
|_____|_____|_____
|add_shim_hdr= NO  | rec_findex = N/A | rec_shim_op = N/A
|_____|_____|_____
|rec_dti_type = N/A|                   | rec_data = N/A
|_____|_____|_____
|modify_smac = YES| modify_dmac = YES| egress_mcast = NO
|_____|_____|_____
|ip_to_mac = NO
|_____|_____|_____
|dest_mac = 001d.e656.cc00 | src_mac = 001d.e658.2800
|_____|_____|_____
|
Statistics: Packets = 1138
           Bytes   = 569016
    
```

Rewrite MAC info

Statistics

Summary

- Troubleshooting Catalyst 4500
- Troubleshooting Catalyst 6500
- Sup-720 and Sup-2T
- Troubleshooting High CPU
- Troubleshooting VSS

Q & A



Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm



Cisco *live!* 365

Don't forget to activate your Cisco Live 365 account for access to all session material,

communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.

www.ciscoliveaustralia.com/portal/login.wv

Cisco *live!*

