Understanding Cat65k/Cisco76k

Kinczli Zoltán – kz@cisco.com

Cisco Systems, Advanced Services

Agenda

- C65k/C76k távoltról
- Történelem
- C65k/C76k közelebbről
- Architektúra
- Packet-flow
- C65k/C76k egészen közelről
- CEF és MPLS ‘troubleshooting’ - understanding

C65k/C76k – távolról

- Nagy – gyors – komplex
- Sok module, sok port
- L2/L3, switch/router - ?

  Ezek kombinációja – de nem L2 router ☺
- *Fancy* modulok

  OSM, FWSM, CSM, VPN, NAM, IDS, FlexWan, …
- Régi emlékeink szerint: valaha LAN switch volt
- Ma: (I)SP core/EDGE, DC, Enterprise…

Történelem – Honnan ‘indultunk’ ?

- CiscoFusion in-a-box – emlékszik még valaki? ☺
- Koncepció
  - ‘lassú’ router + ‘buta’ switch = ‘Mpps’ ‘mega’-switch/router
  - MLS: MultiLayer Switching
    - CLI: mls rp, mls sp
  - HW: router + L2 switch w/ NFFC
- in-a-box
  - C5/6K: RSFC, MSM, MSFC
    - SUP: MSFCx+PFCy
      - 1;1, 1;2, 2;2, 3;3(a/b/bxl), 3;2a
    - PFC1: data-driven, flow-based
    - PFC2,3: topology driven, CEF/FIB
      - örökség: CLI ‘sh mls…’

Layer 3 Packet Processing
Multilayer Switching w/ PFC2/PFC3

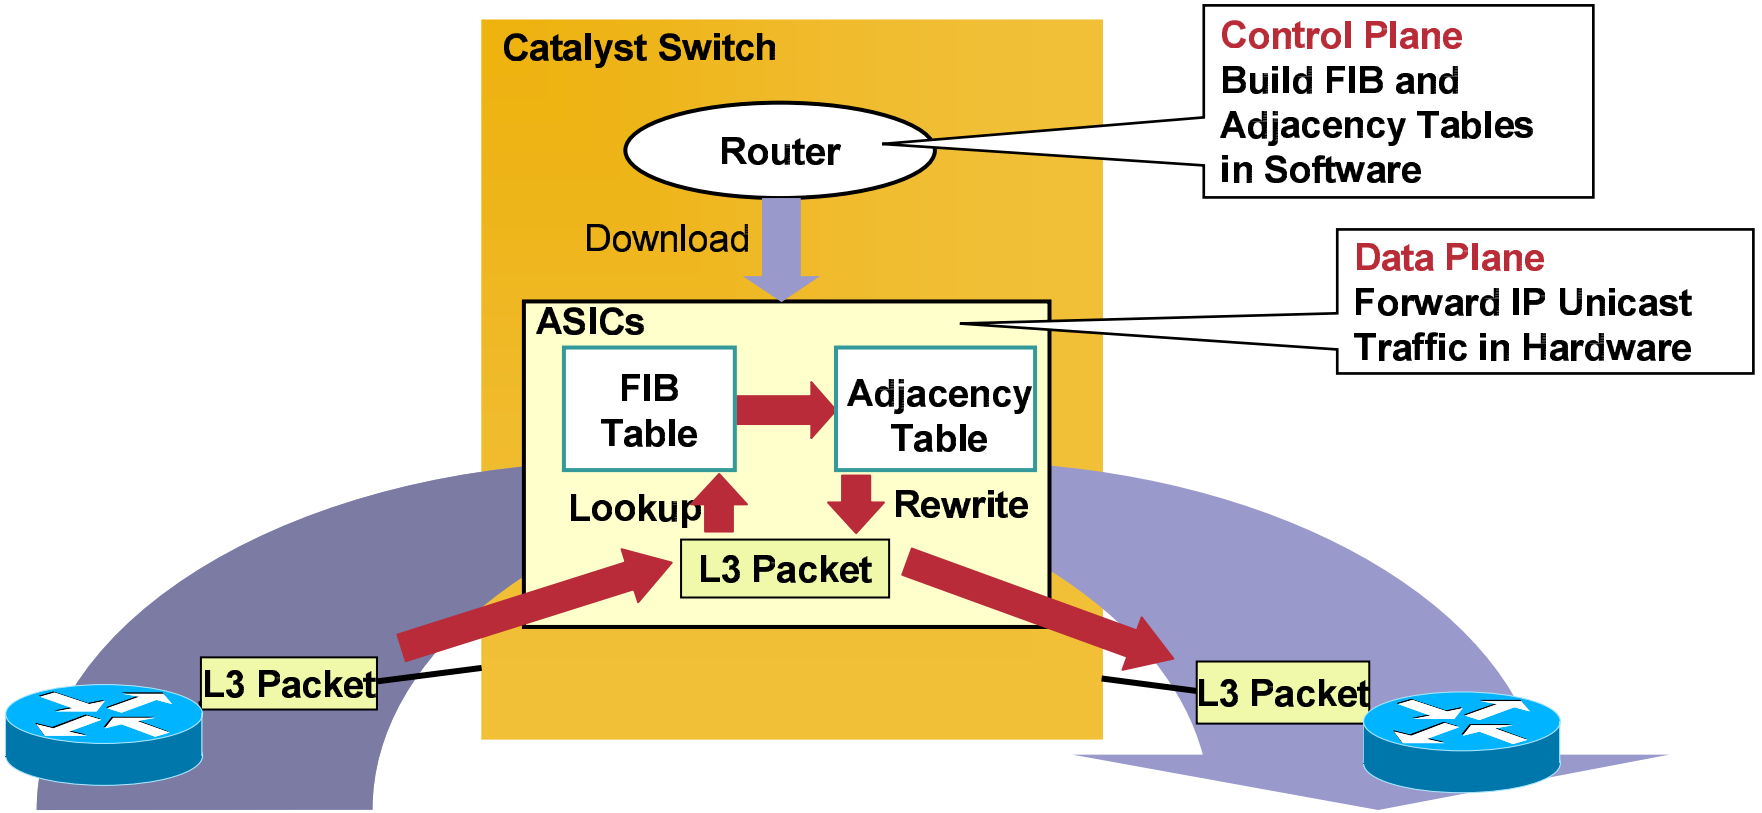

C65k/C76k – közelebbről

- RP, SP – CPU

  Console owner – RP (de boot), tshoot, crashinfo…
- superman, tycho; pinnacle, hyperion… – ASIC
- PXF – OSM, FlexWan
- DBUS (16G), RBUS, EOBC
- fabric/X-bar (256G@fdx/720G@fdx)
- LC: classic, cef256, cef720, osm, sip, service modulok
- CatOS, IOS – hibrid (2 külön img), natív (2 img ‘egyben’)
- v4/6, MPLS, MPLS-TE, MCast, VPNv4/6, MVPN, VPLS, ATOM + L2 ‘dolgok’ + service modulok + …

Catalyst 6503/6503E and 6504E

- Slots 1 and 2—Supervisor engine, or switching module
- Other slots—Any switching module
- 2 fabric channels per slot
- Power supplies in rear
- 950W AC/DC and 1400W AC power supplies for 6503/6503E
- 2700W AC/DC power supplies for 6504E

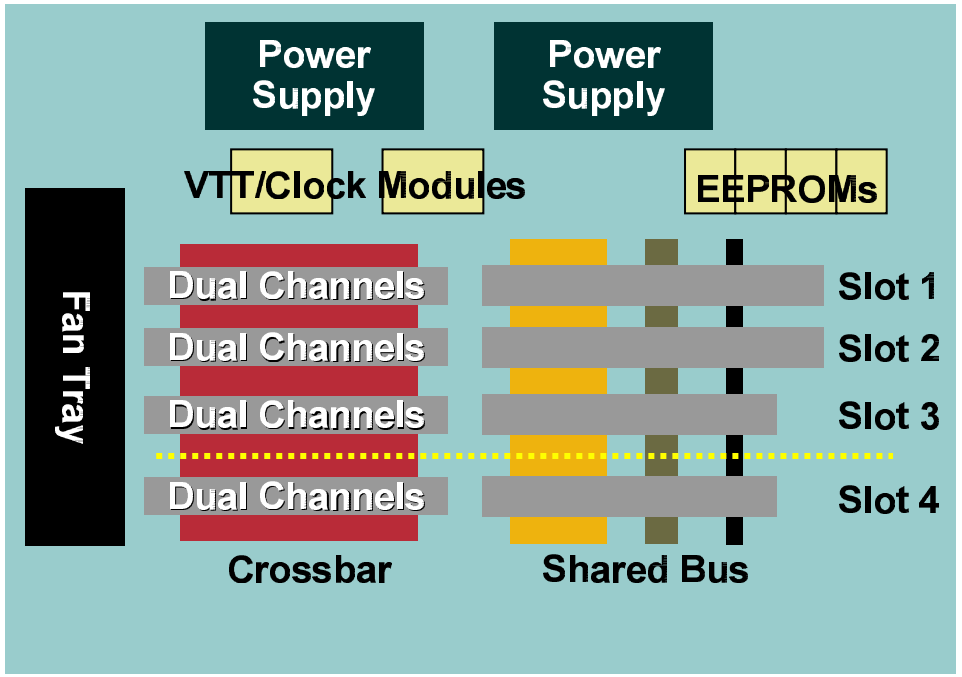

Note: CEF720 modules not supported in Catalyst 6503 (non-E) chassis

Catalyst 6506/6509 and 6506E/6509E

- Slots 1 and 2—Supervisor Engine 2, or switching module
- Slots 5 and 6—Supervisor Engine 32/720, or switching module
- Other slots—Any switching module
- 2 fabric channels per slot
- Wide variety of power supplies, from legacy 1000W to new 6000W—E chassis requires at least 2500W PS
- NEB-A chassis has vertical slot alignment, dual fan trays, front-to-back air flow, air filtration system

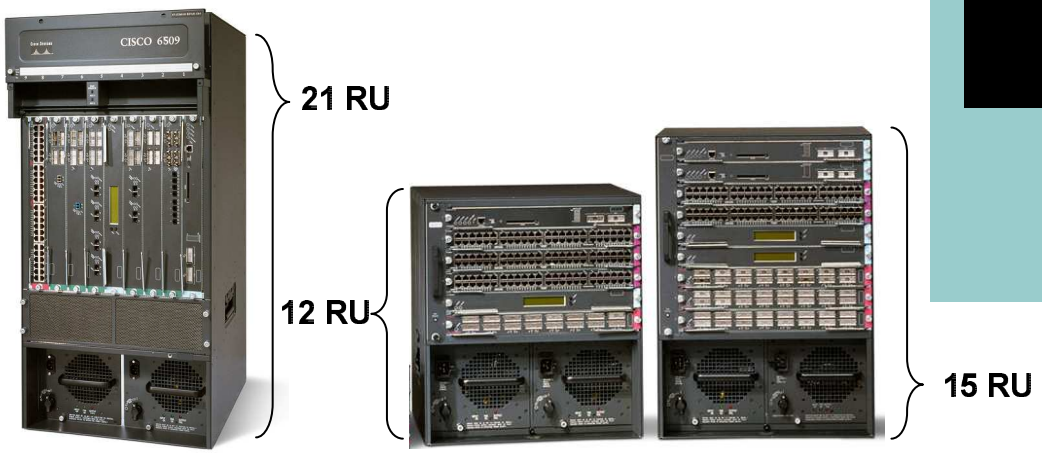

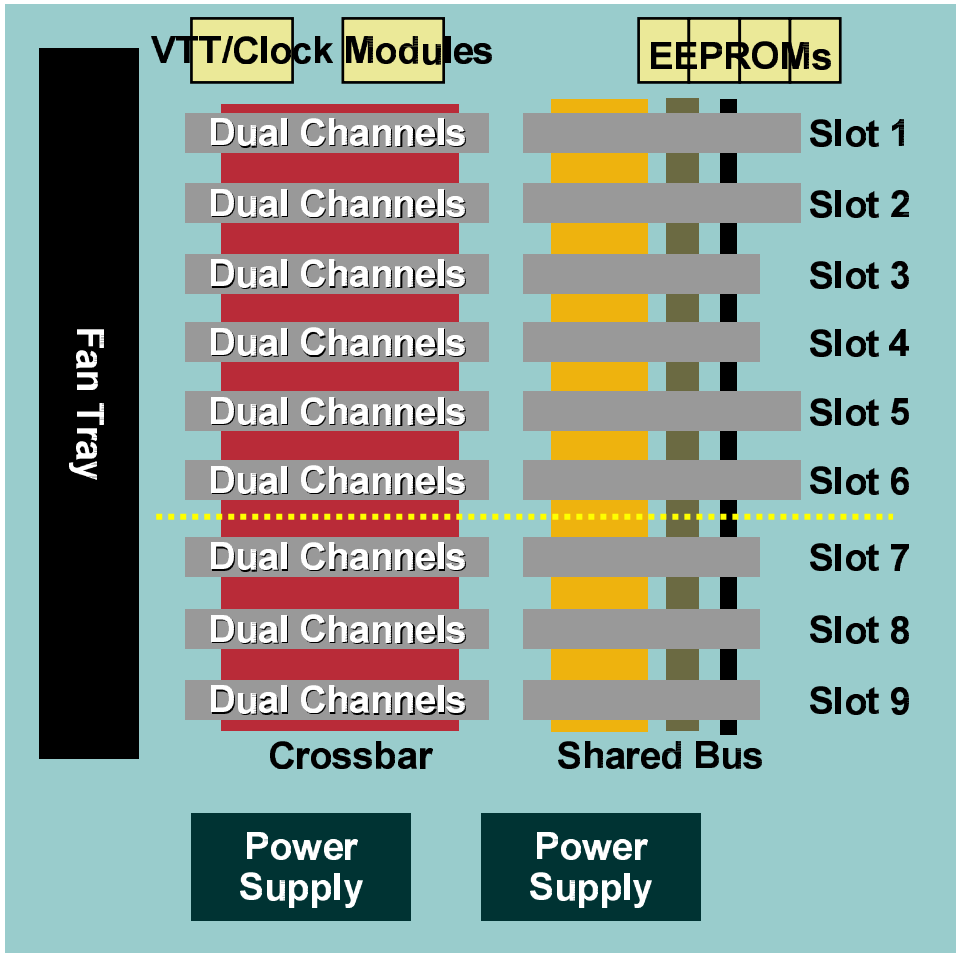

Catalyst 6513

- Slots 1 and 2—Supervisor Engine 2, or switching module
- Slots 7 and 8—Supervisor Engine 32/720, or switching module
- Wide variety of power supplies, from 2500W to new 6000W
- 1 fabric channel slots 1–8

  Dual-fabric modules not supported in slots 1–8!
- 2 fabric channels slots 9–13

  Any switching module

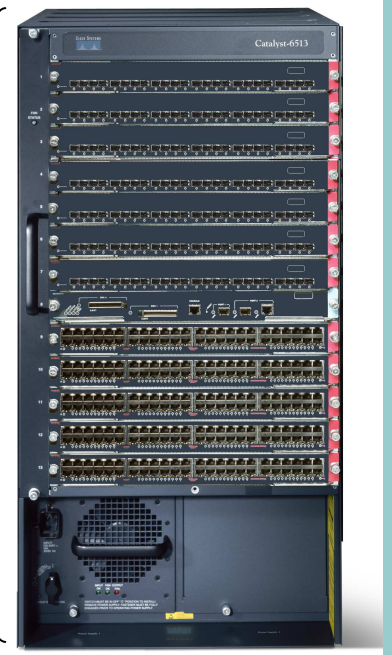

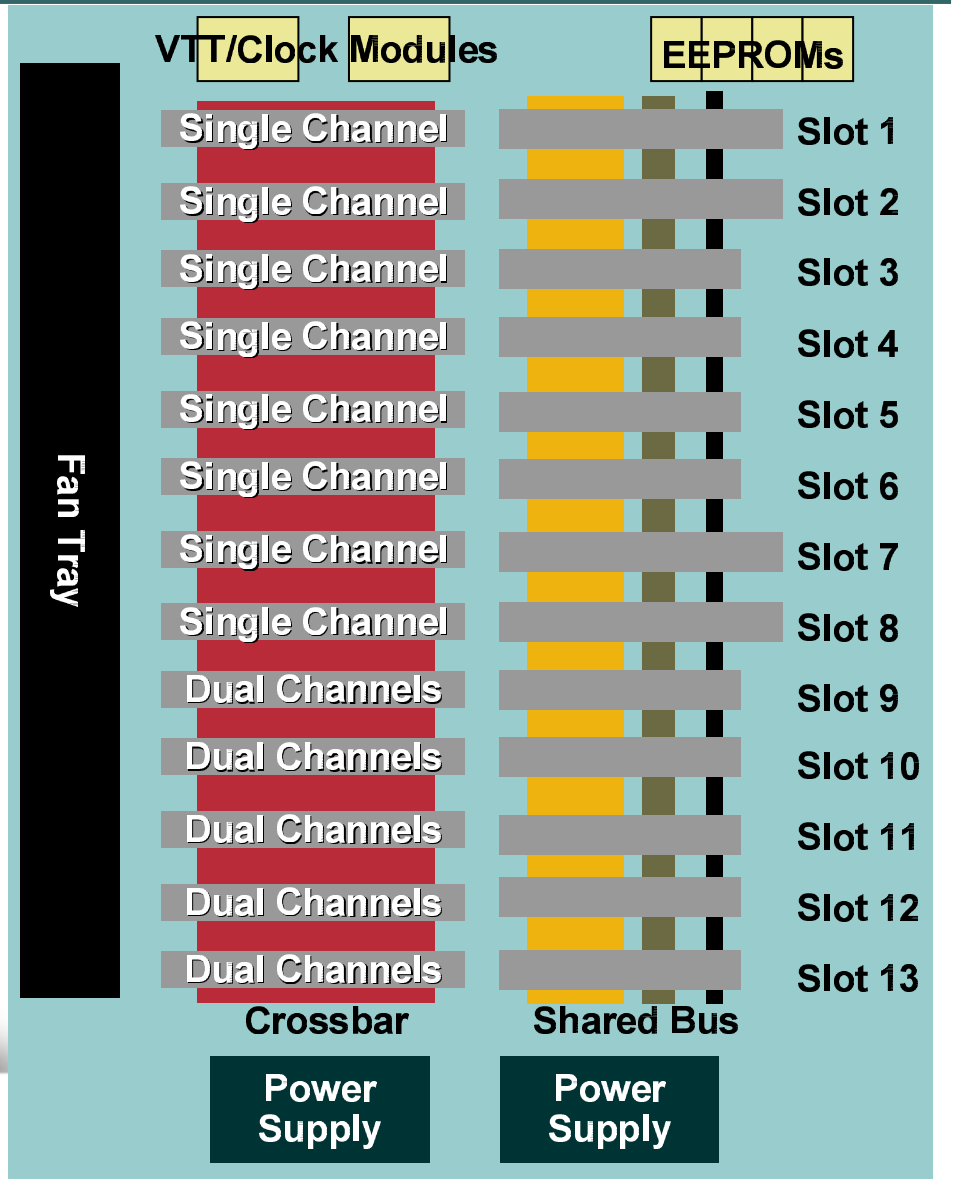

Supervisor Engine 2/PFC2 Architecture

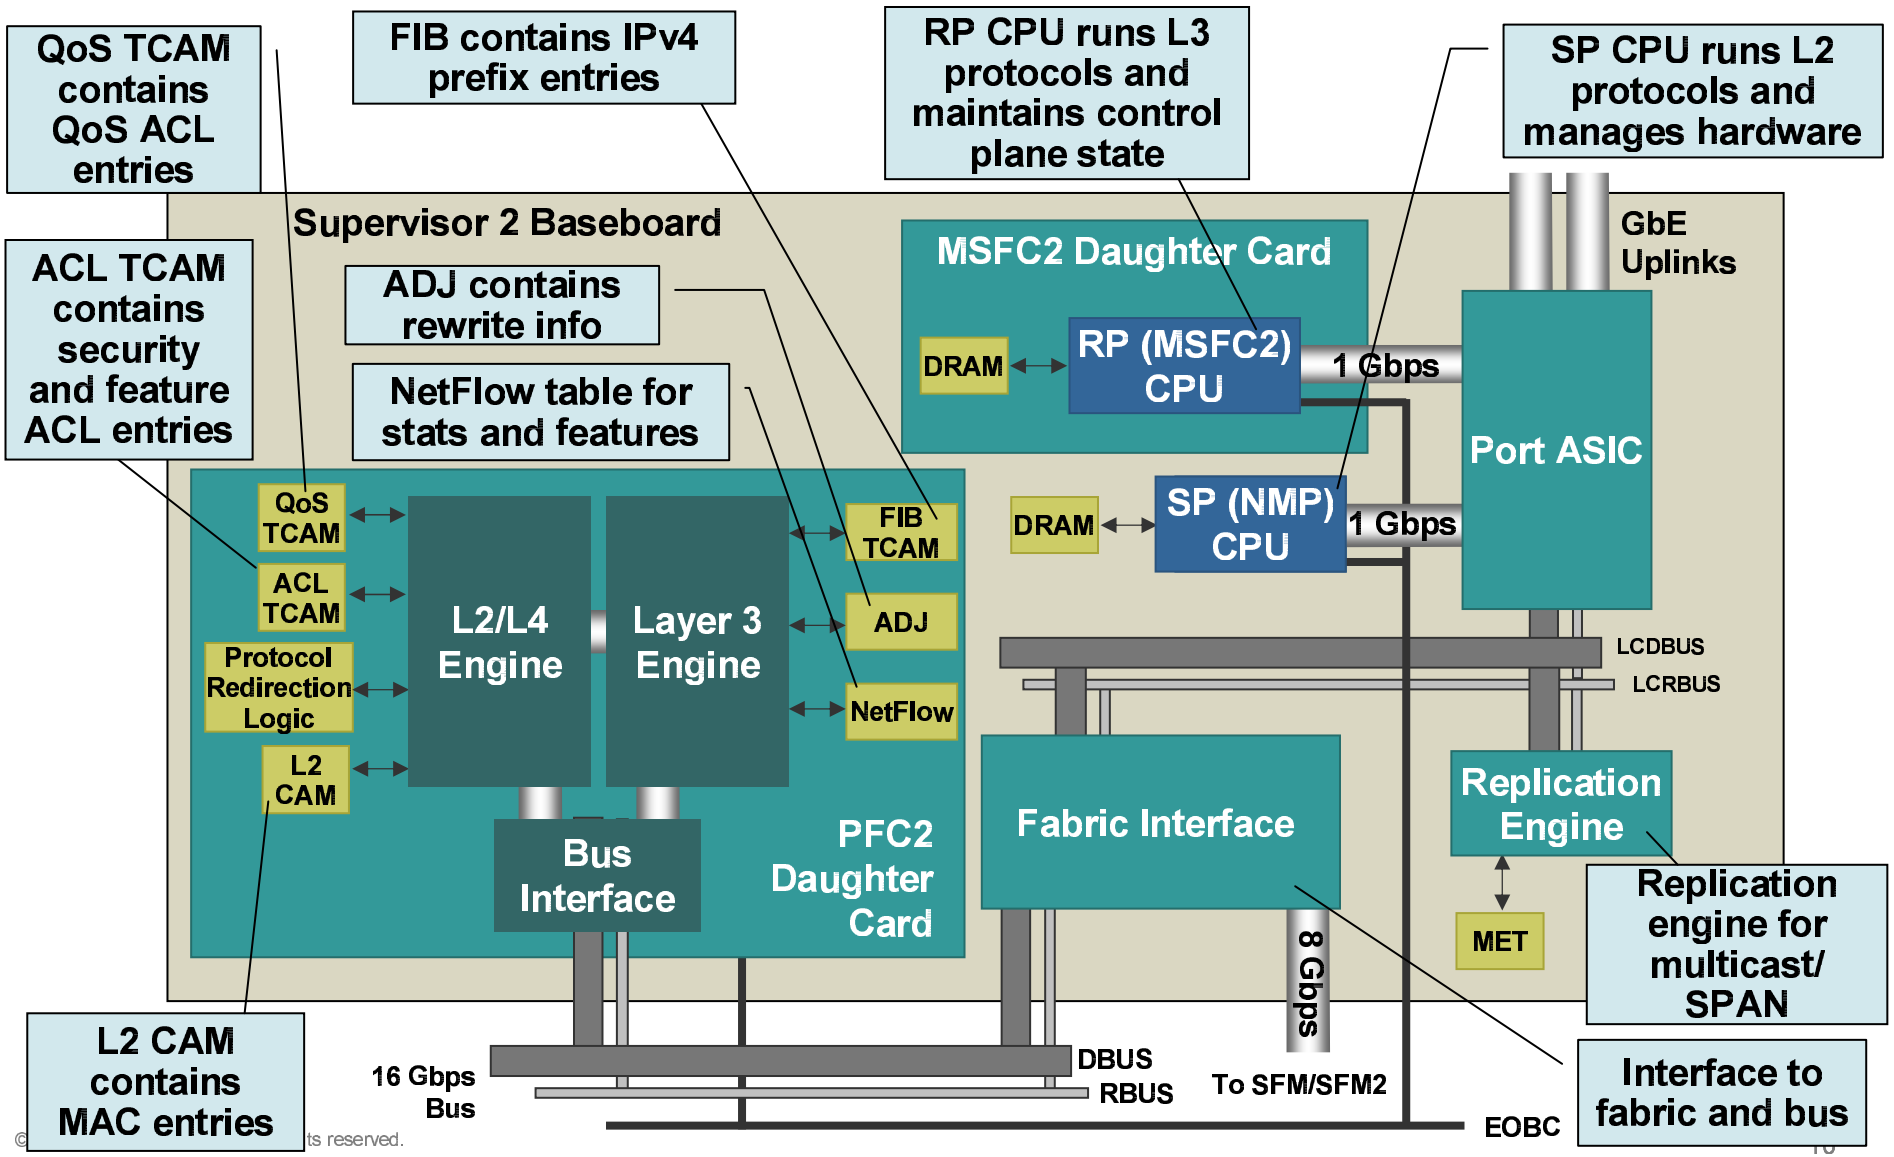

Supervisor Engine 720/PFC3 Architecture

Cisco.com

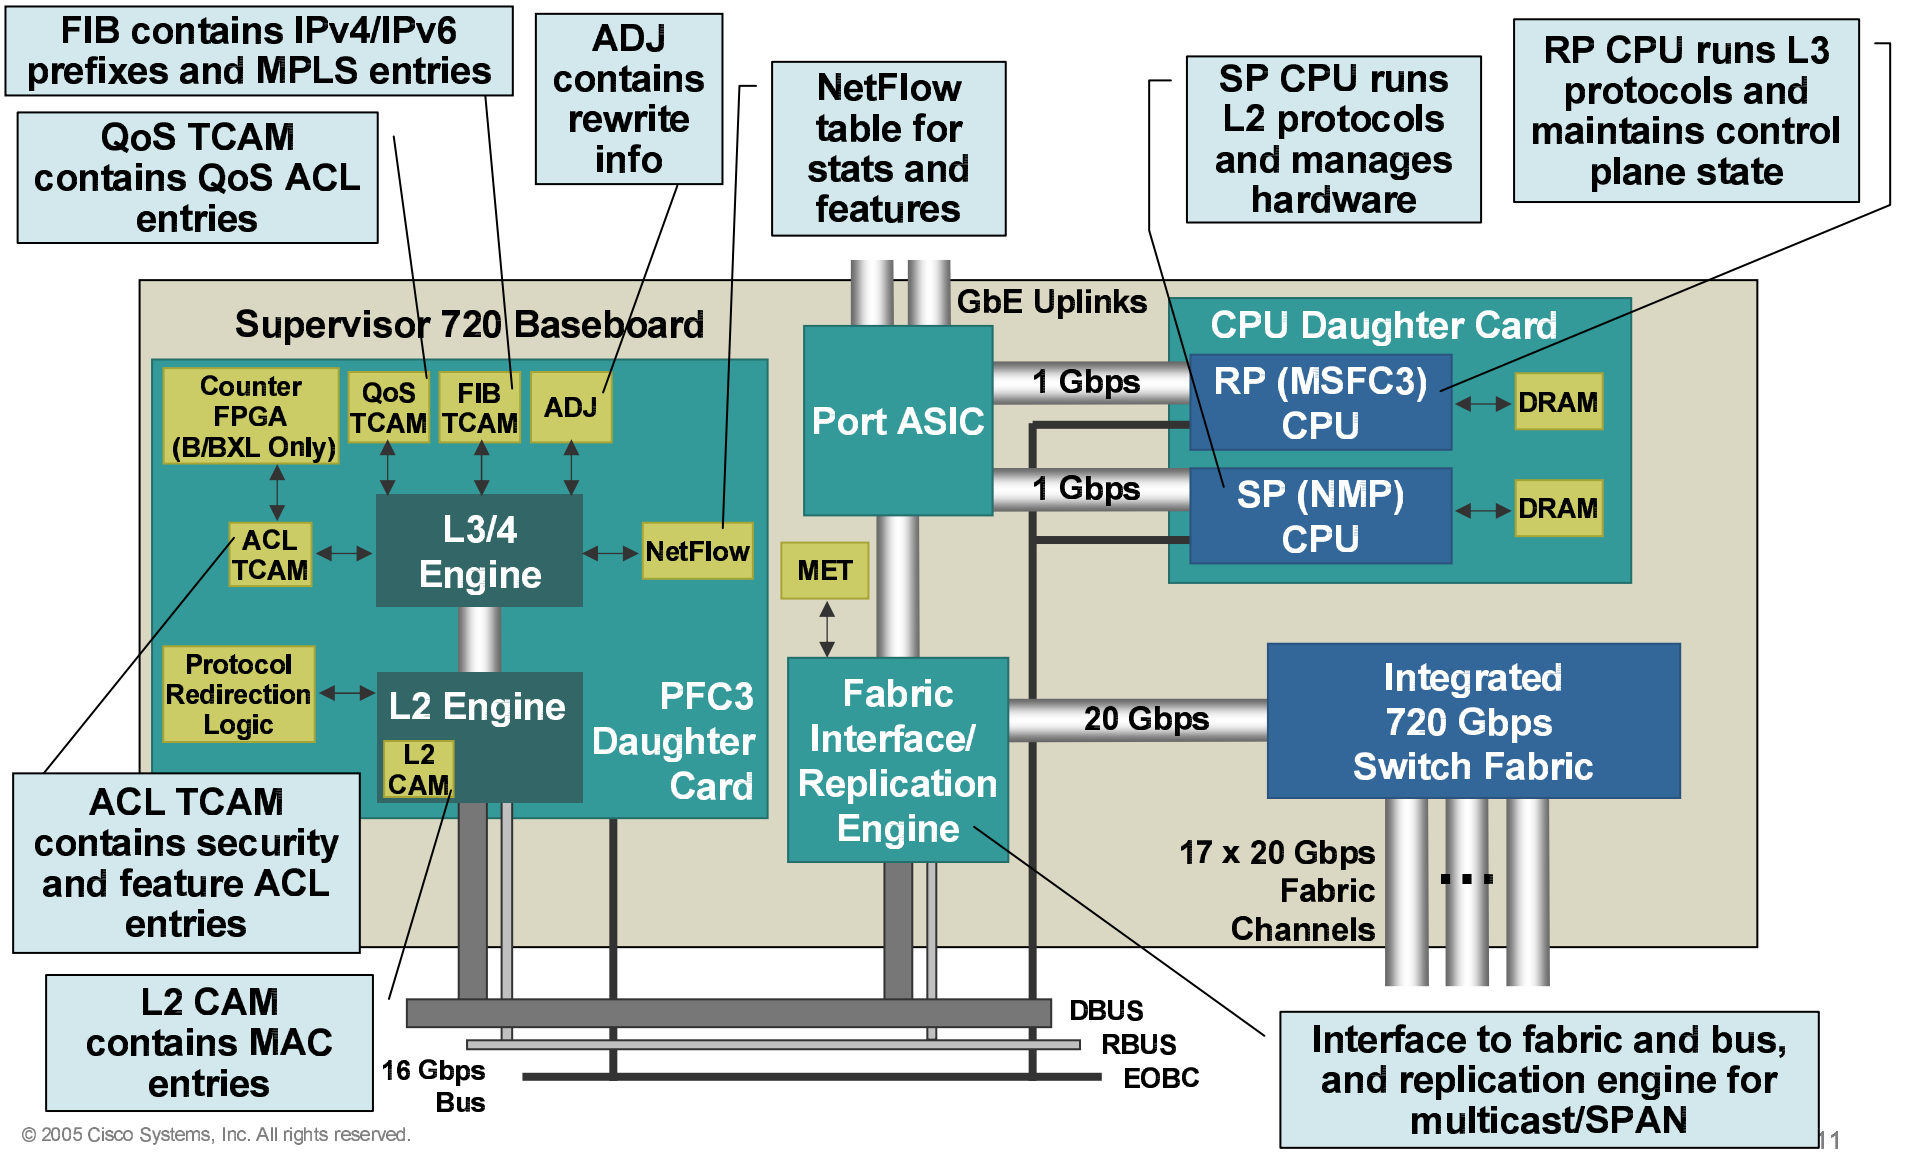

Supervisor Engine 32/PFC3 Architecture

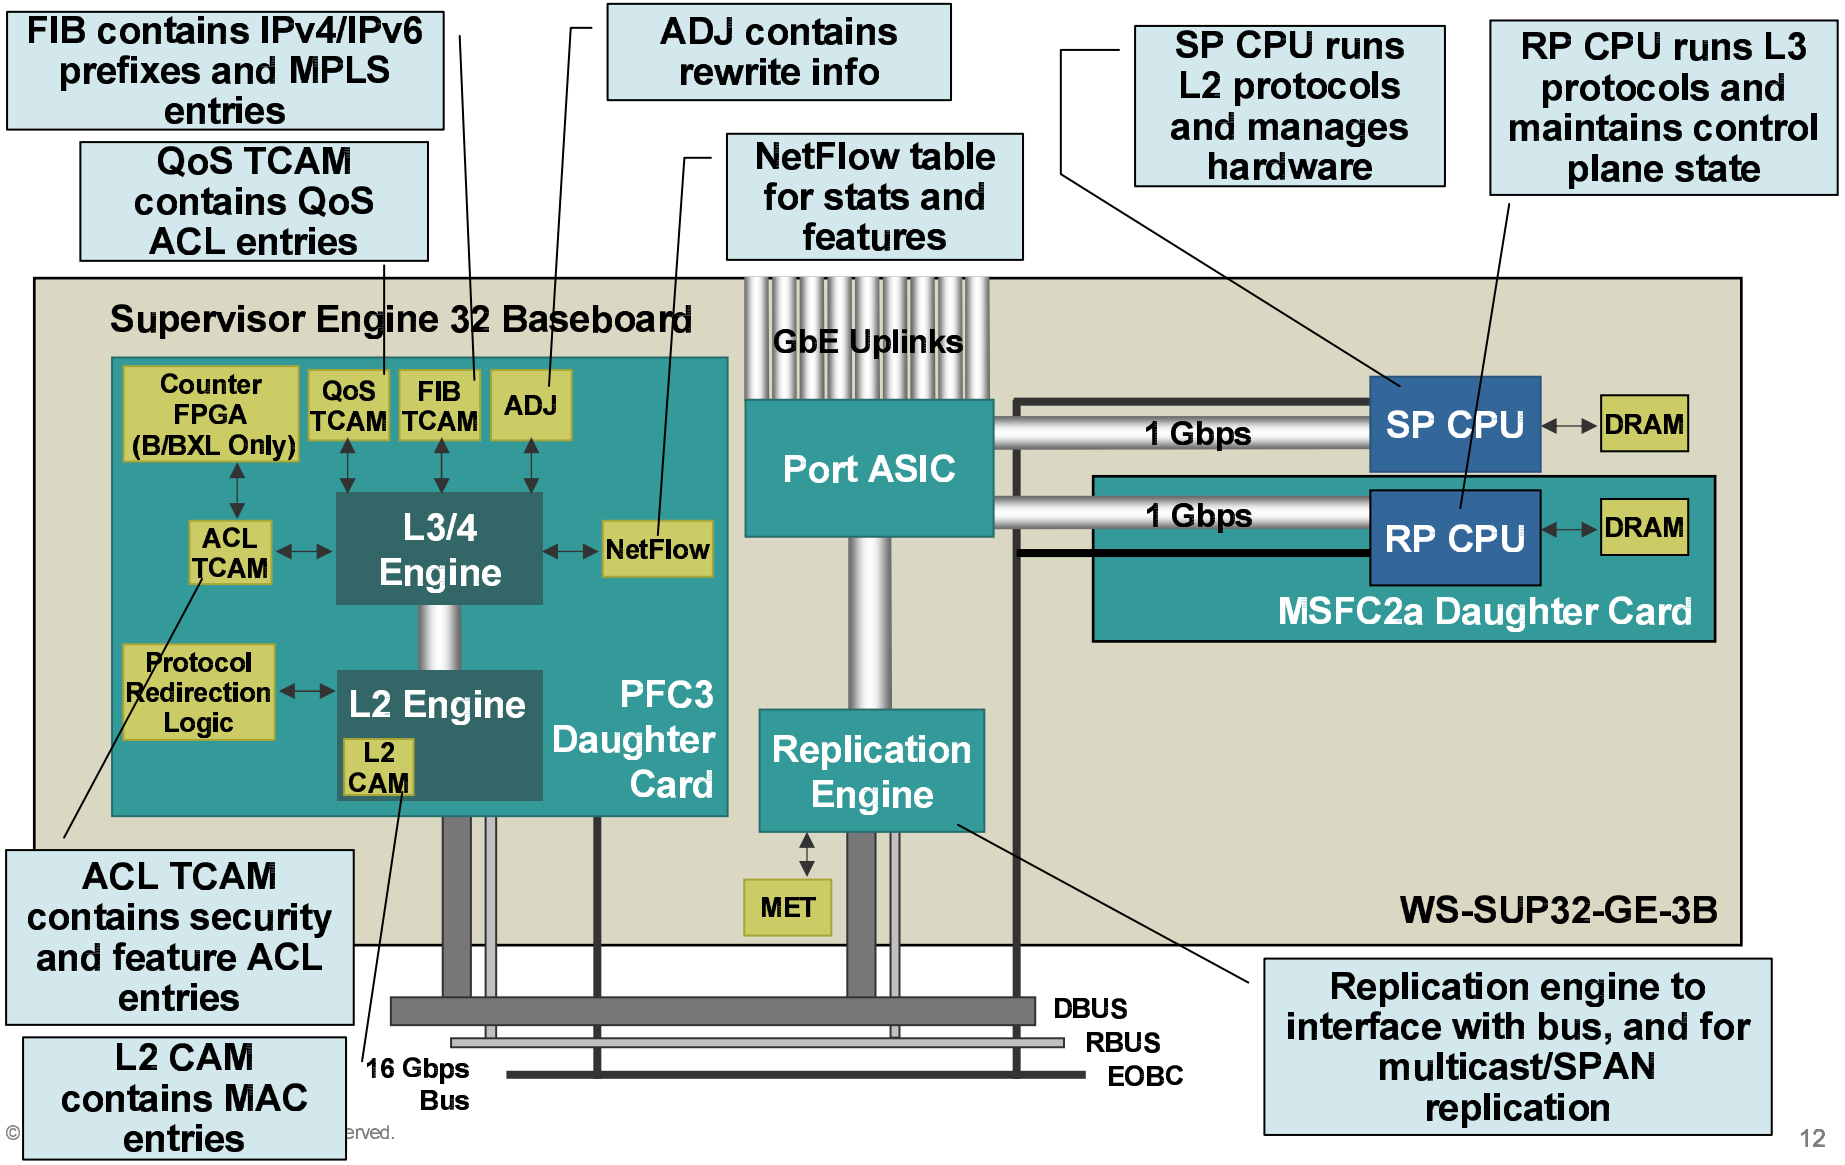

Catalyst 6500 Architecture
with Supervisor 720

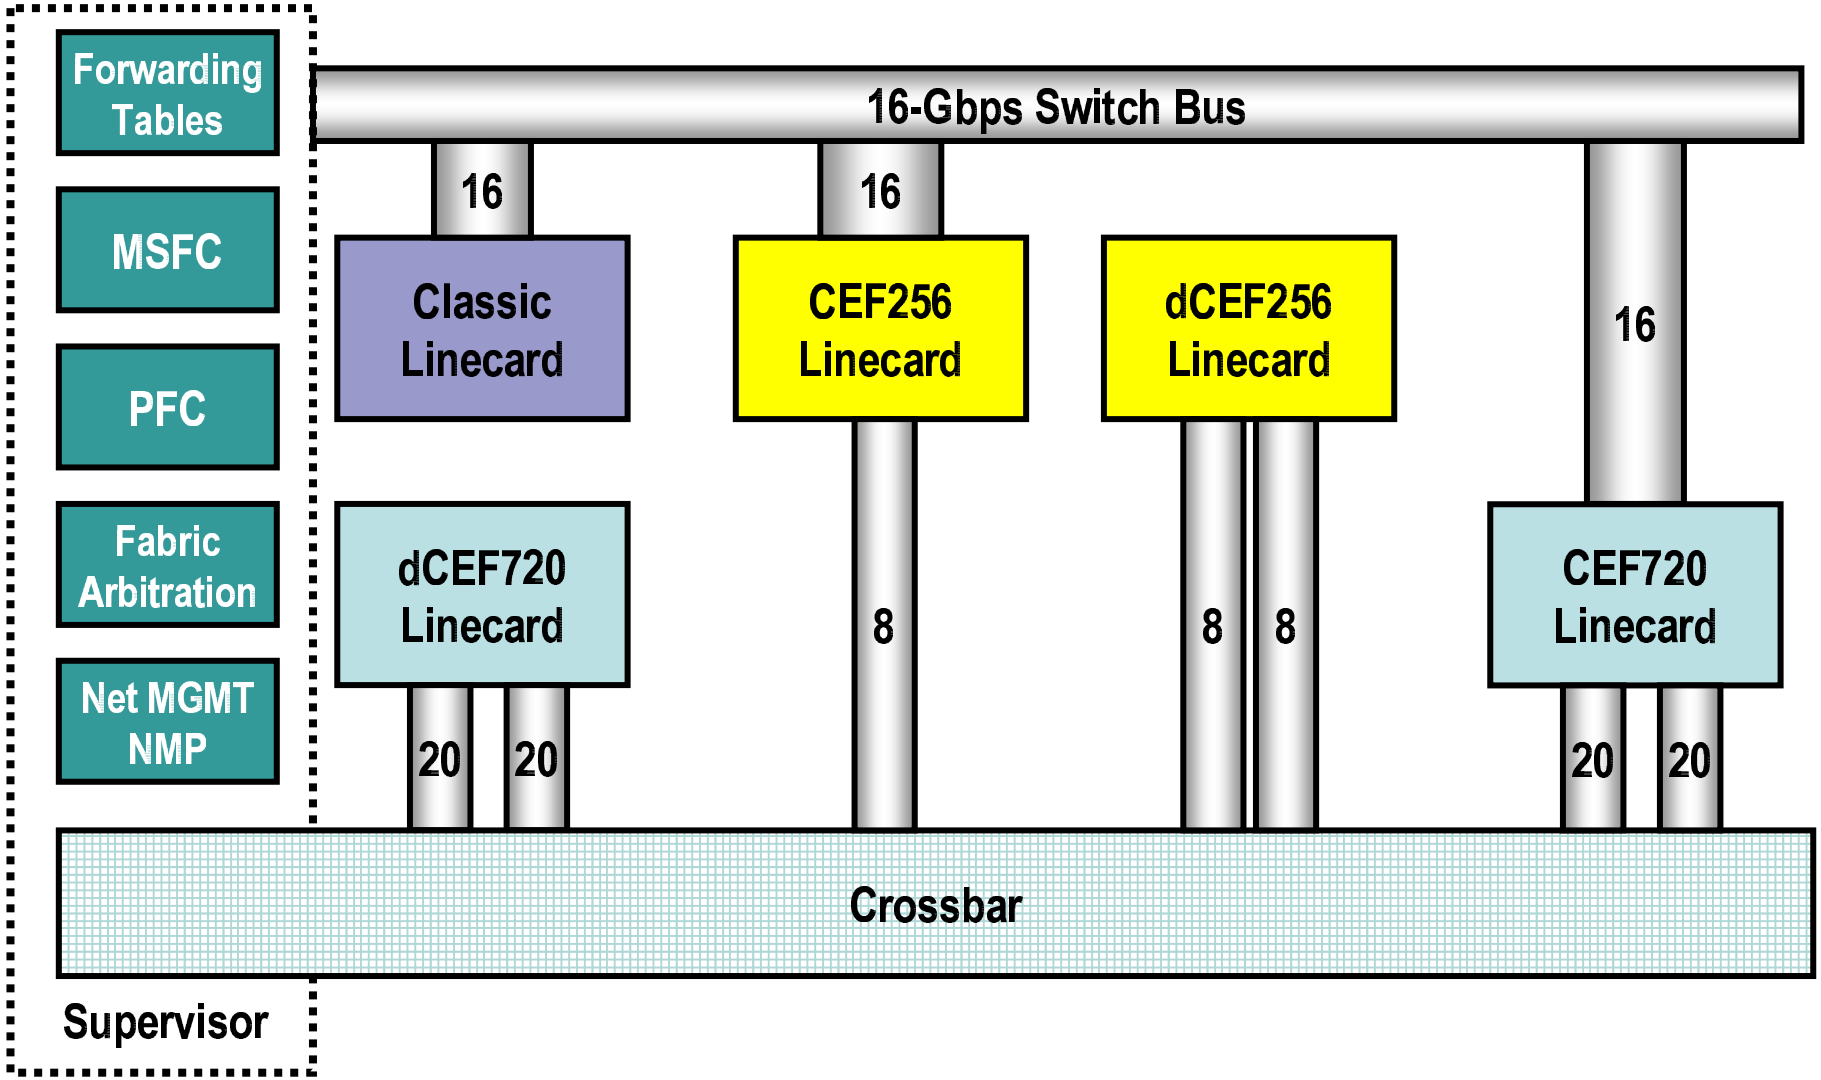

Catalyst 6500 Architecture
Linecard Architecture with Supervisor 32

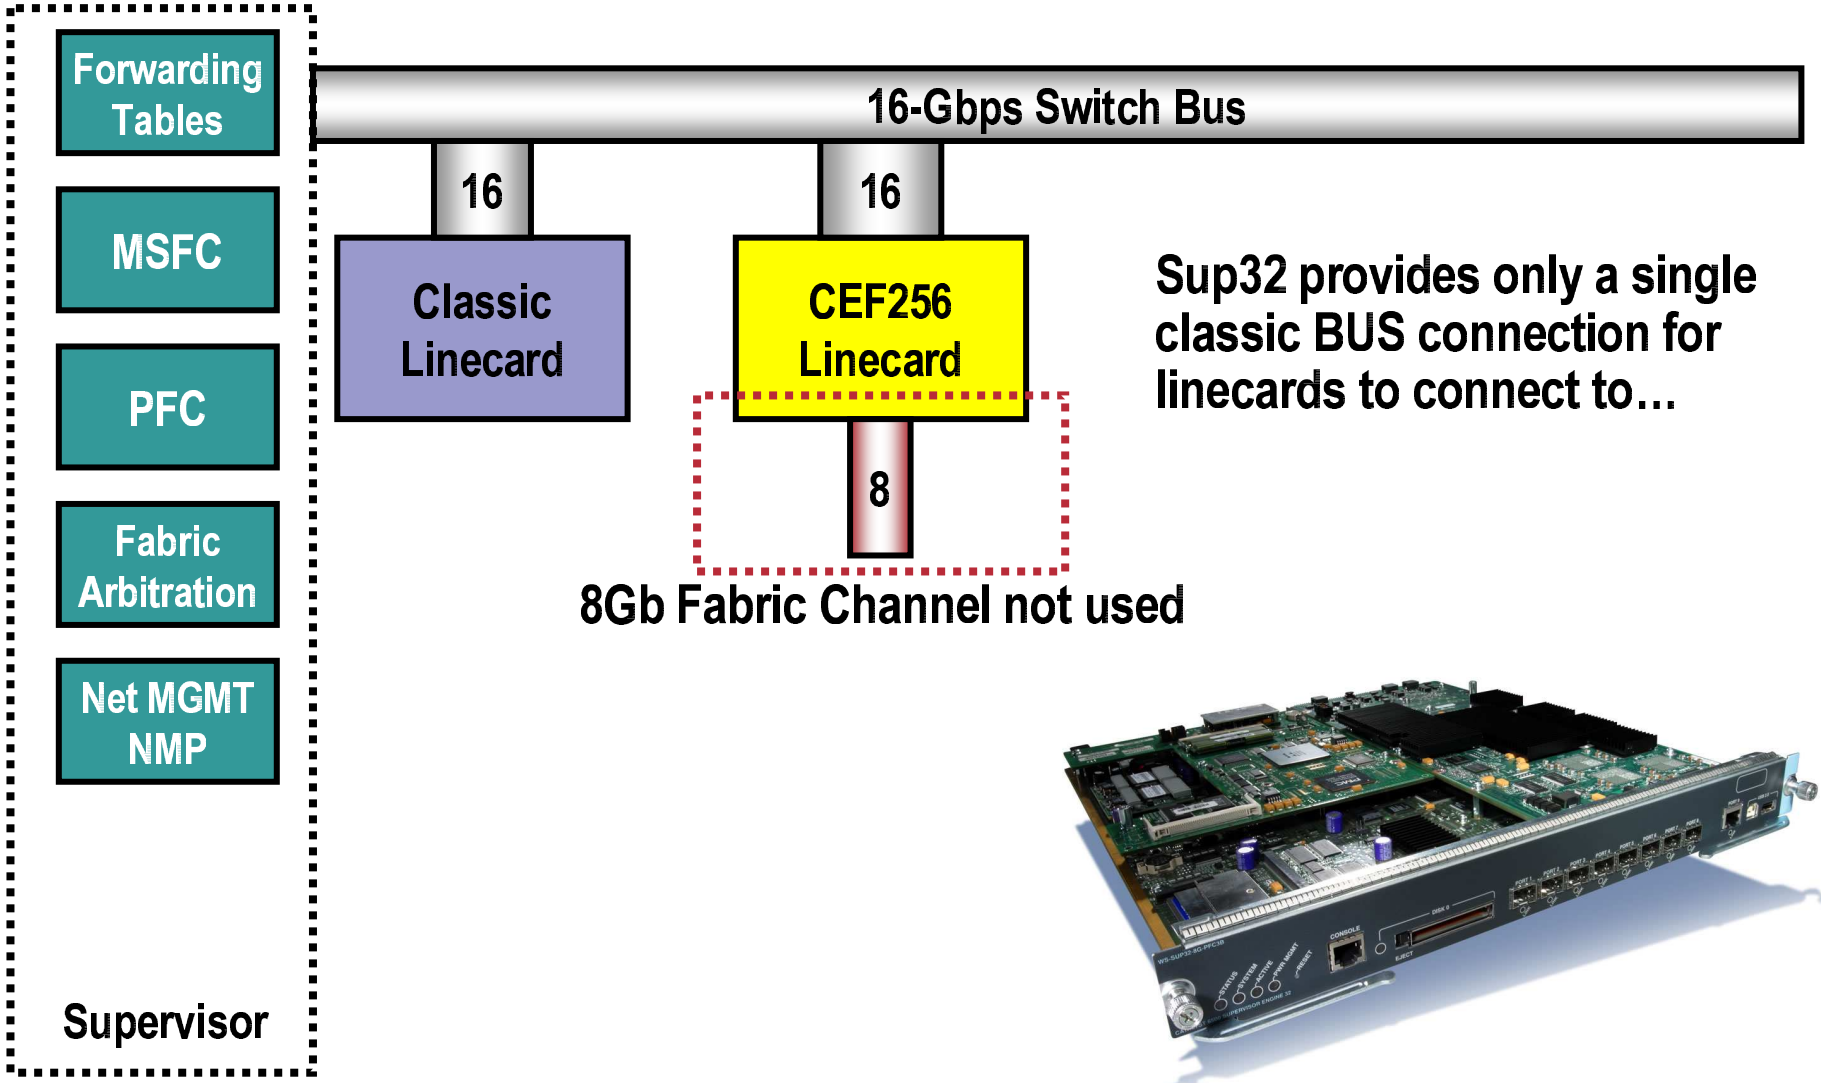

Classic to Classic Centralized Forwarding

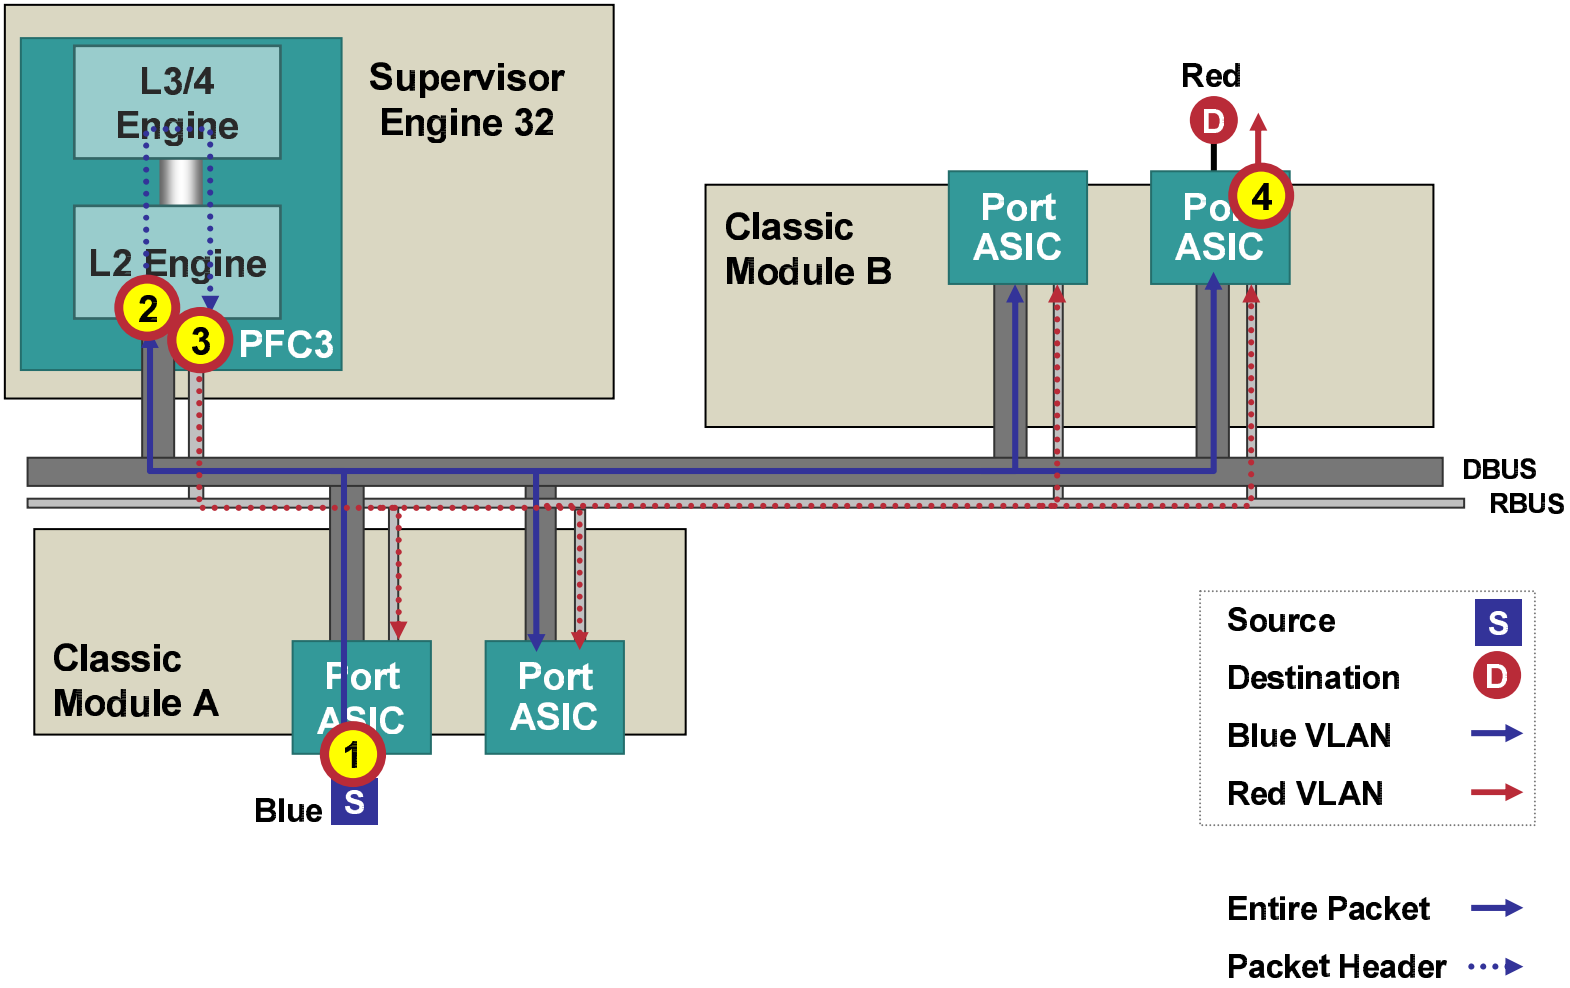

Reference: Classic to Classic Centralized Forwarding

1. Unicast IPv4 packet received on Classic Module A; entire packet is flooded on DBUS and all devices, including the PFC on the supervisor engine, receive it
2. PFC makes a forwarding decision for the packet
3. PFC floods forwarding decision result on RBUS
4. Egress port ASIC on Classic Module B is selected to transmit the packet—all other devices on the bus discard the packet

CEF256 to CEF256 Centralized Forwarding

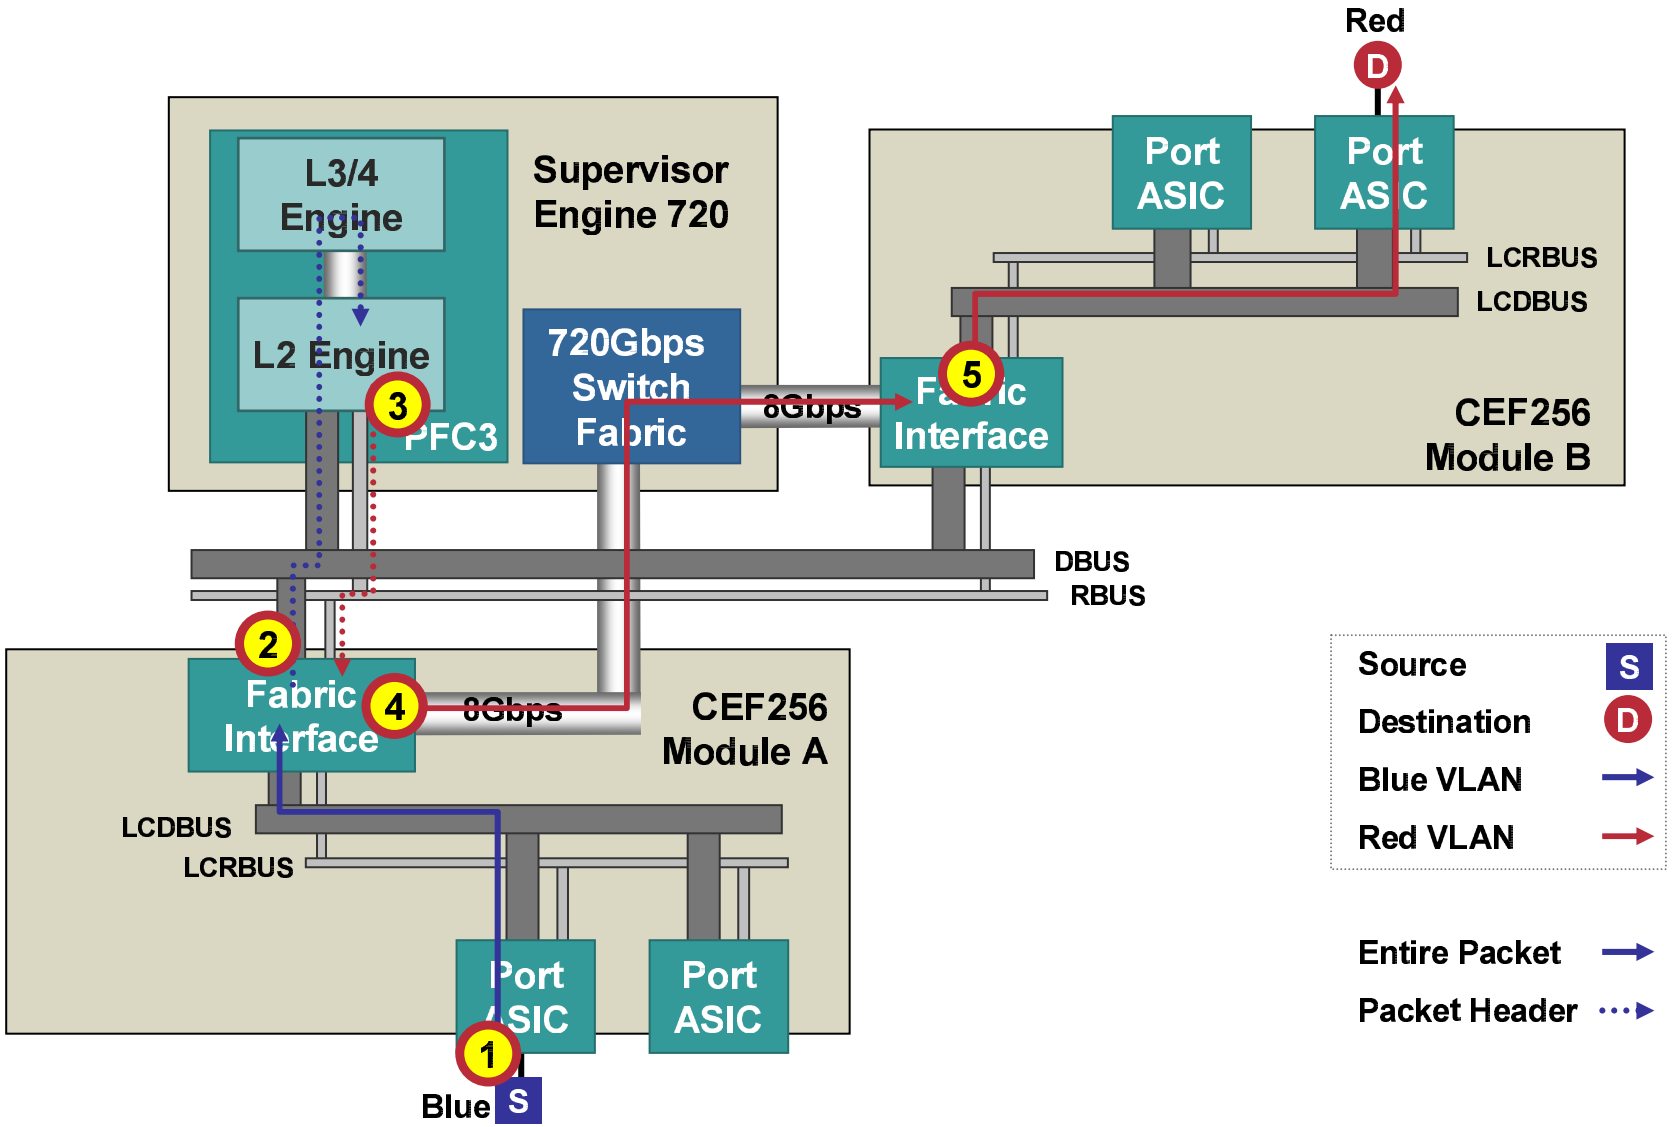

Reference: CEF256 to CEF256 Centralized Forwarding

1. Unicast IPv4 packet received on CEF256 Module A; entire packet is flooded on LCDBUS and fabric interface receives it
2. Fabric interface floods just the packet header on the DBUS; PFC receives packet header and makes a forwarding decision for the packet
3. PFC floods forwarding decision result on RBUS
4. Fabric interface transmits packet across the fabric
5. CEF256 Module B receives the packet and transmits the packet, and the result, on its LCDBUS; the egress port ASIC is selected to transmit the packet

CEF720/DFC3 to CEF720/DFC3 Distributed Forwarding

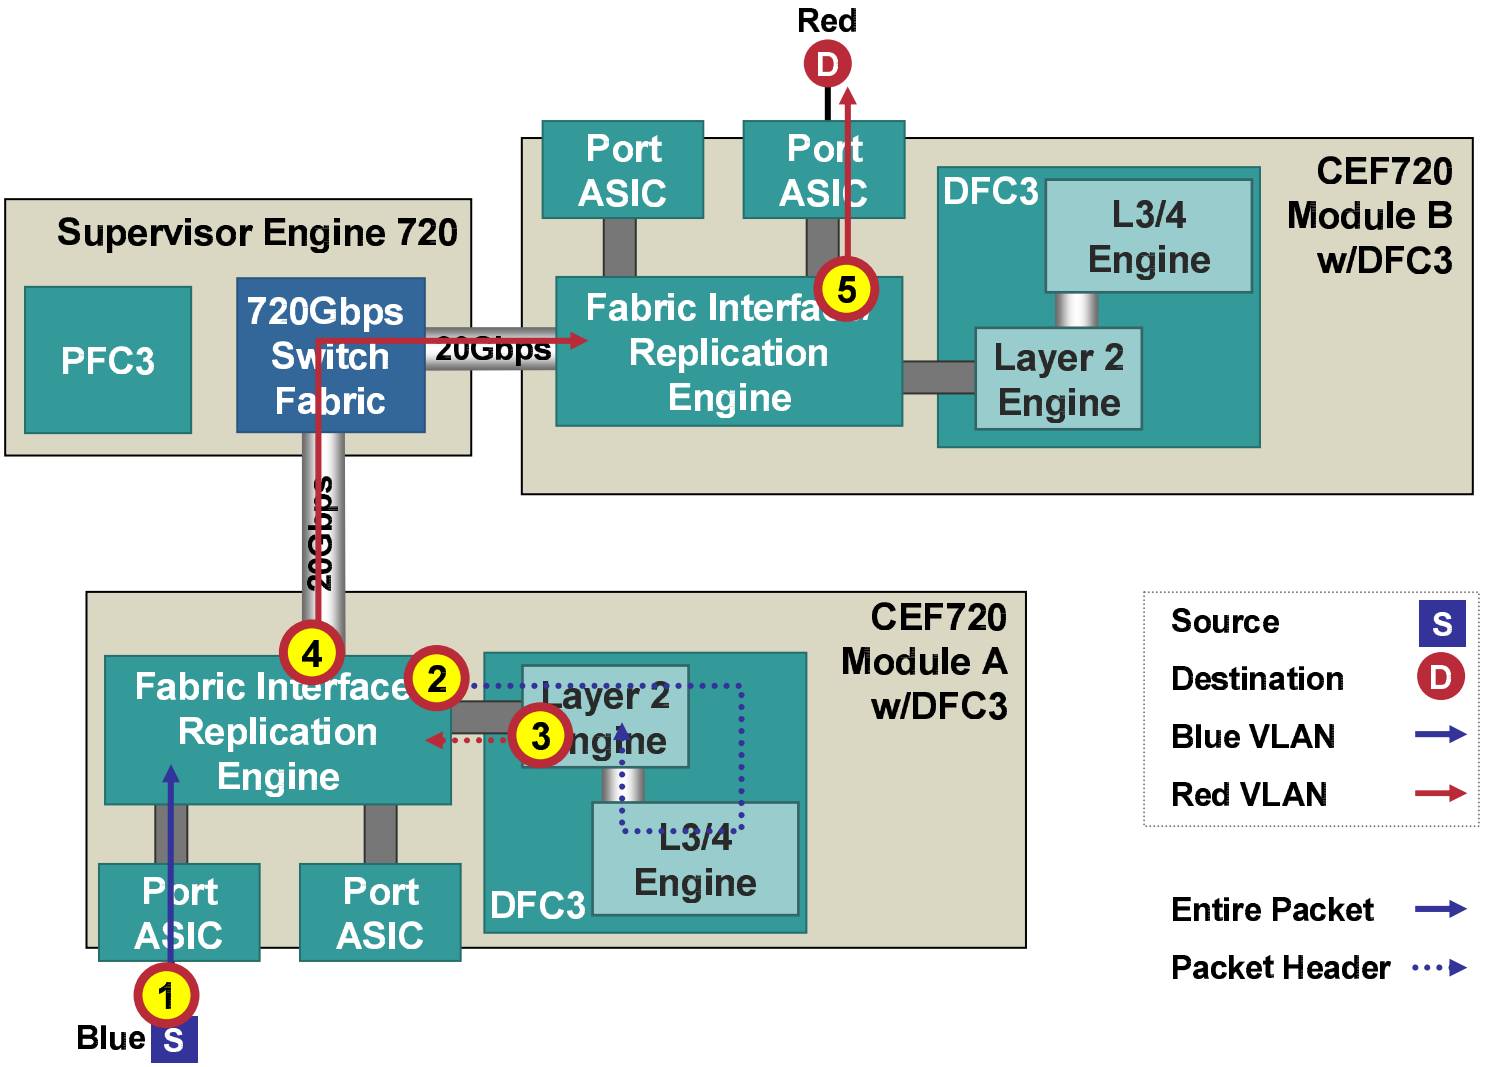

Reference: CEF720/DFC3 to CEF720/DFC3 Distributed Forwarding

1. Unicast IPv4 packet received on CEF720 Module A; entire packet is forwarded to the fabric interface
2. Fabric interface sends just the packet header to the DFC; DFC makes a forwarding decision for the packet
3. DFC returns the forwarding decision result to the fabric interface
4. Fabric interface transmits packet across the fabric
5. CEF720 Module B receives the packet and transmits the packet to the egress port ASIC

CEF – Programming ASIC's

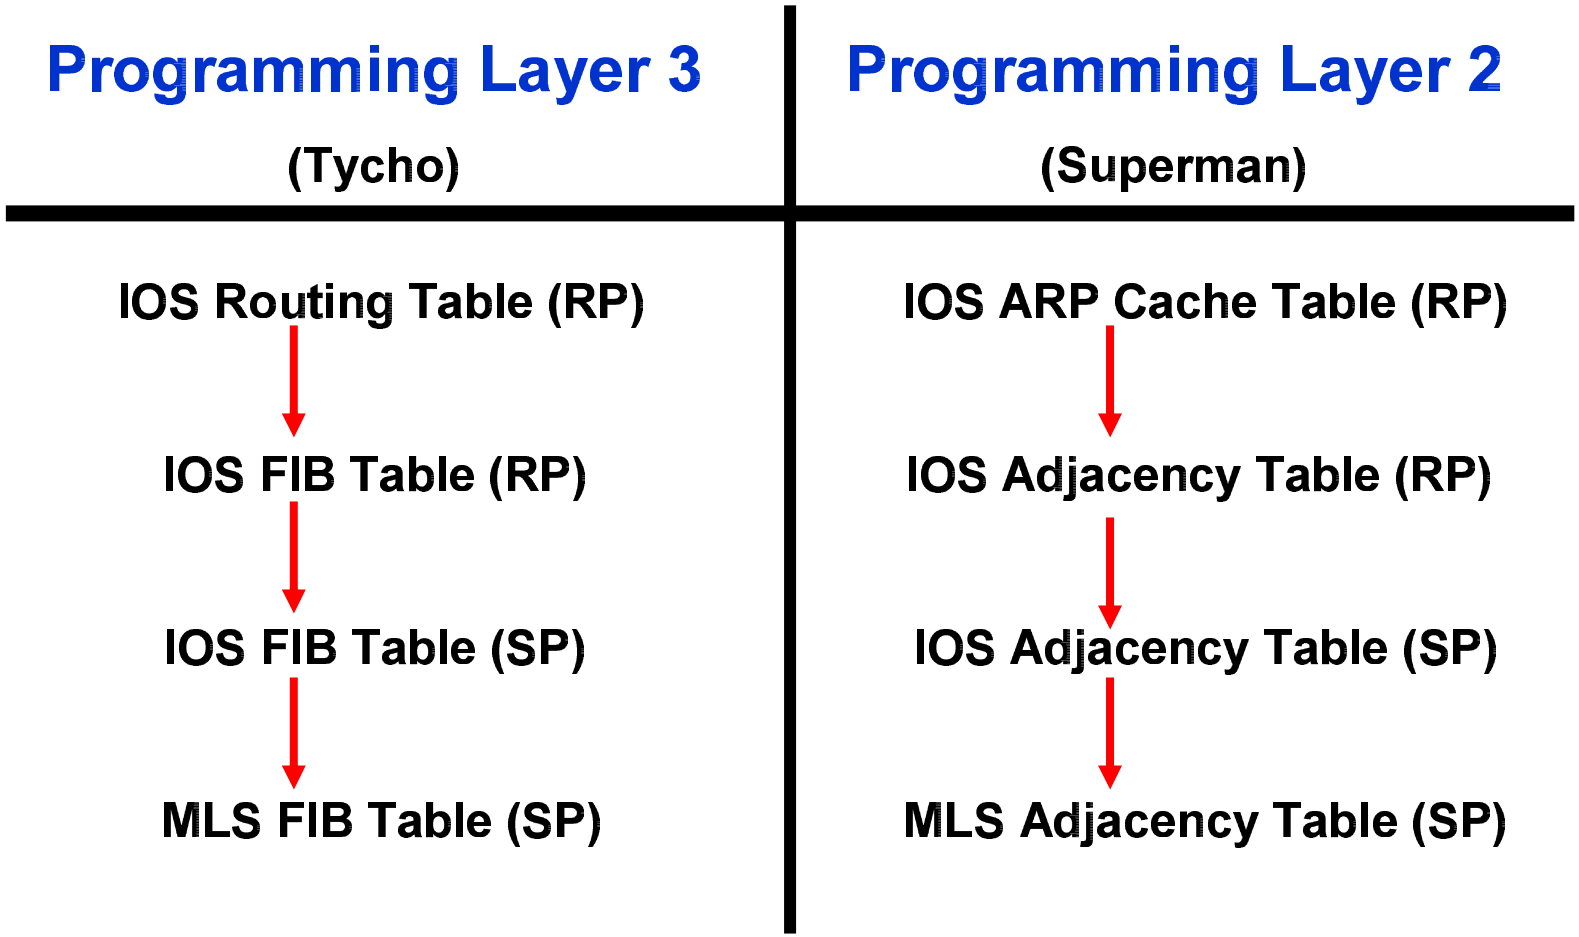

Basic MPLS Path Flows

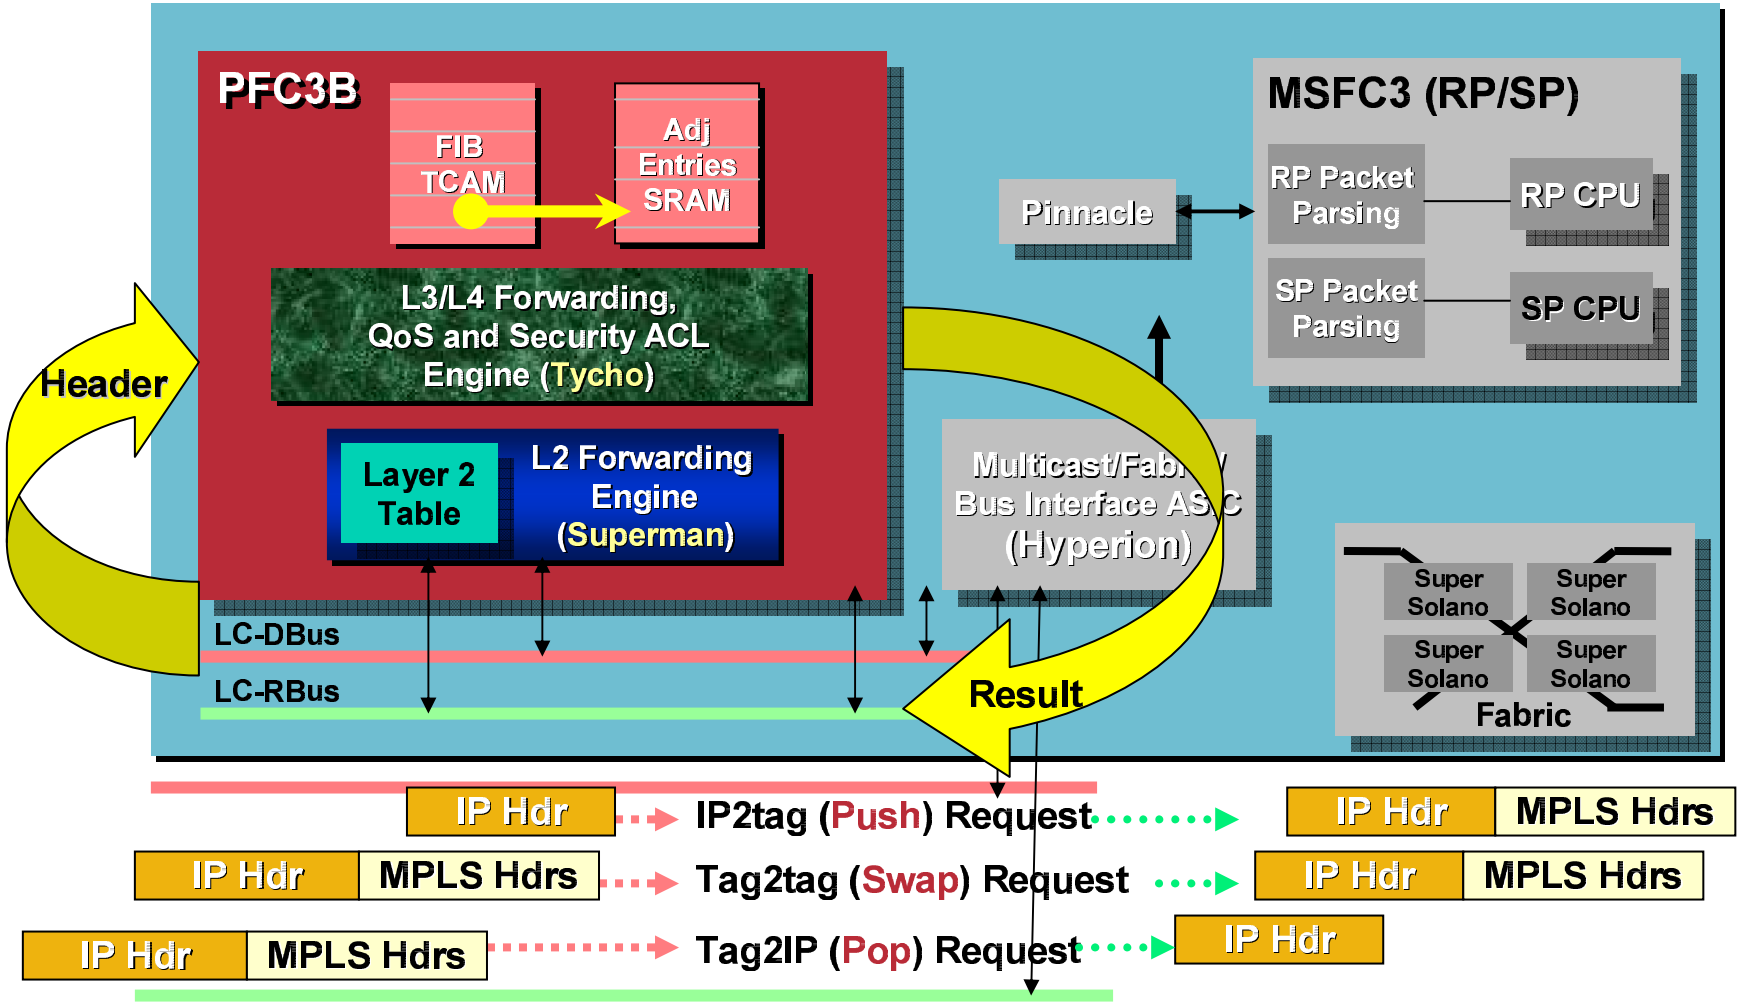

Distribution of IP/MPLS Routing Information

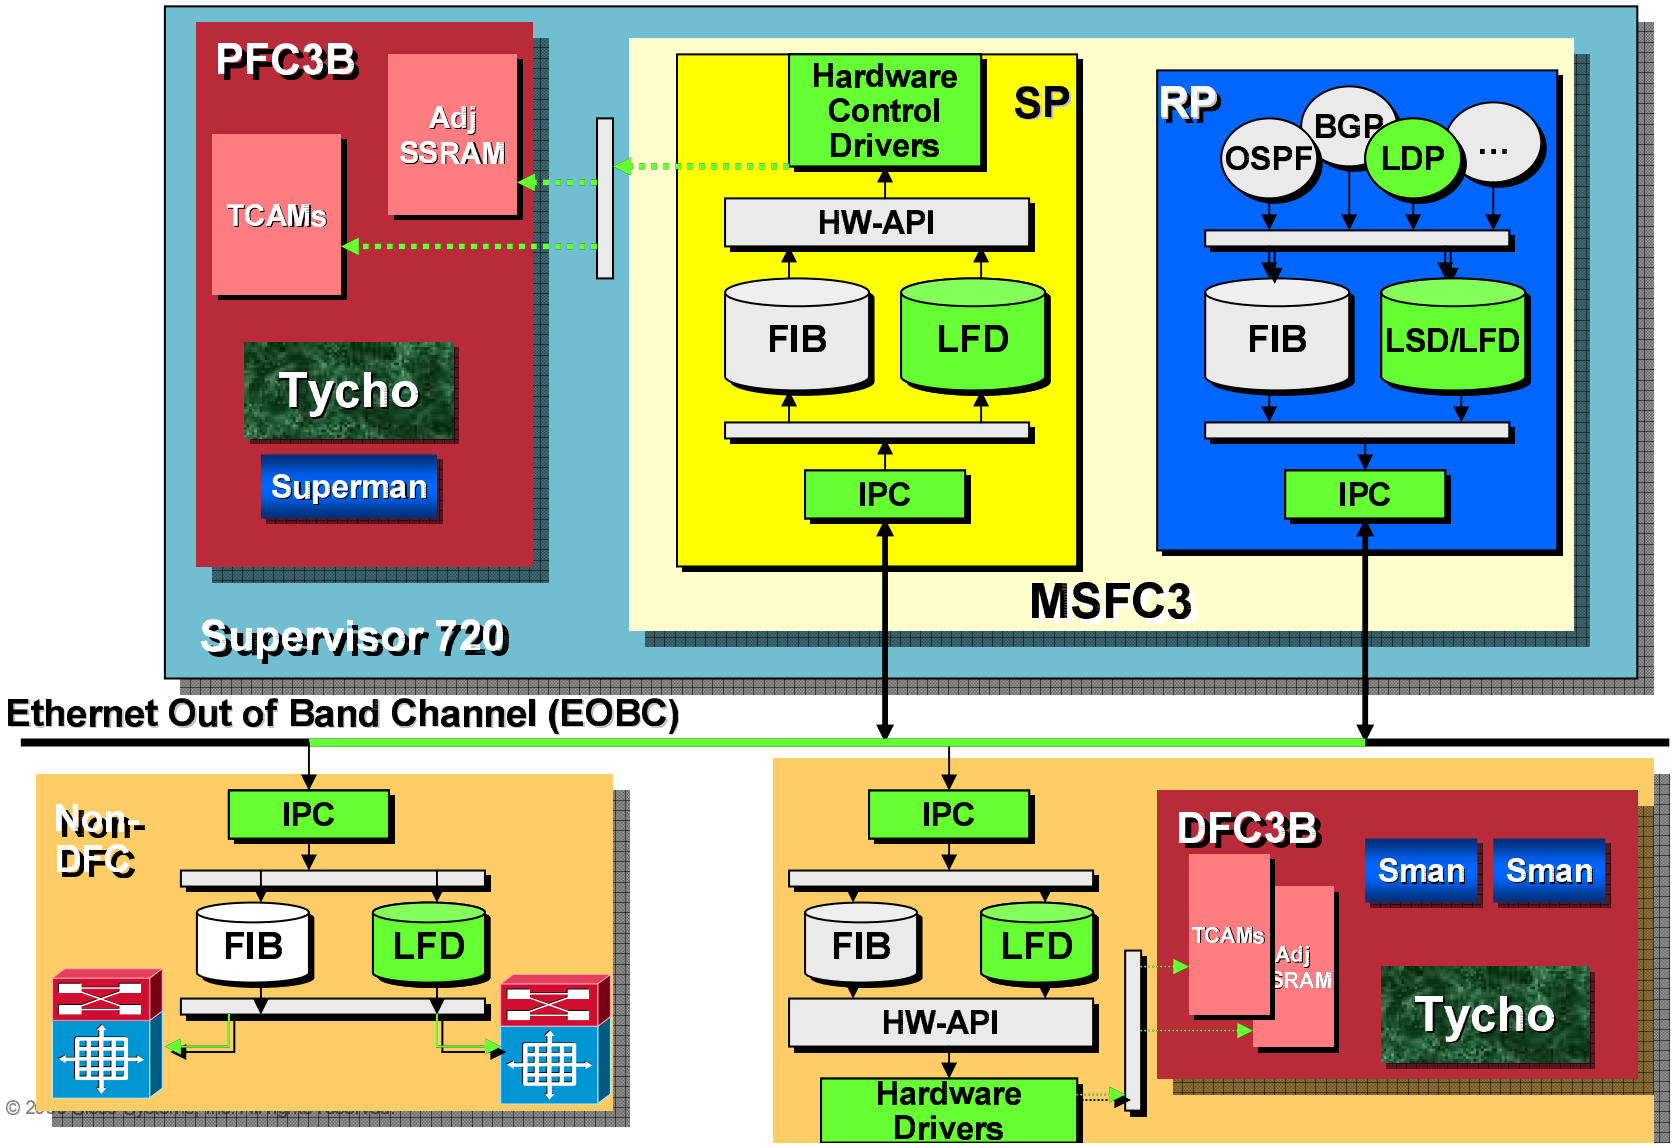

MPLS Control / Data Plane

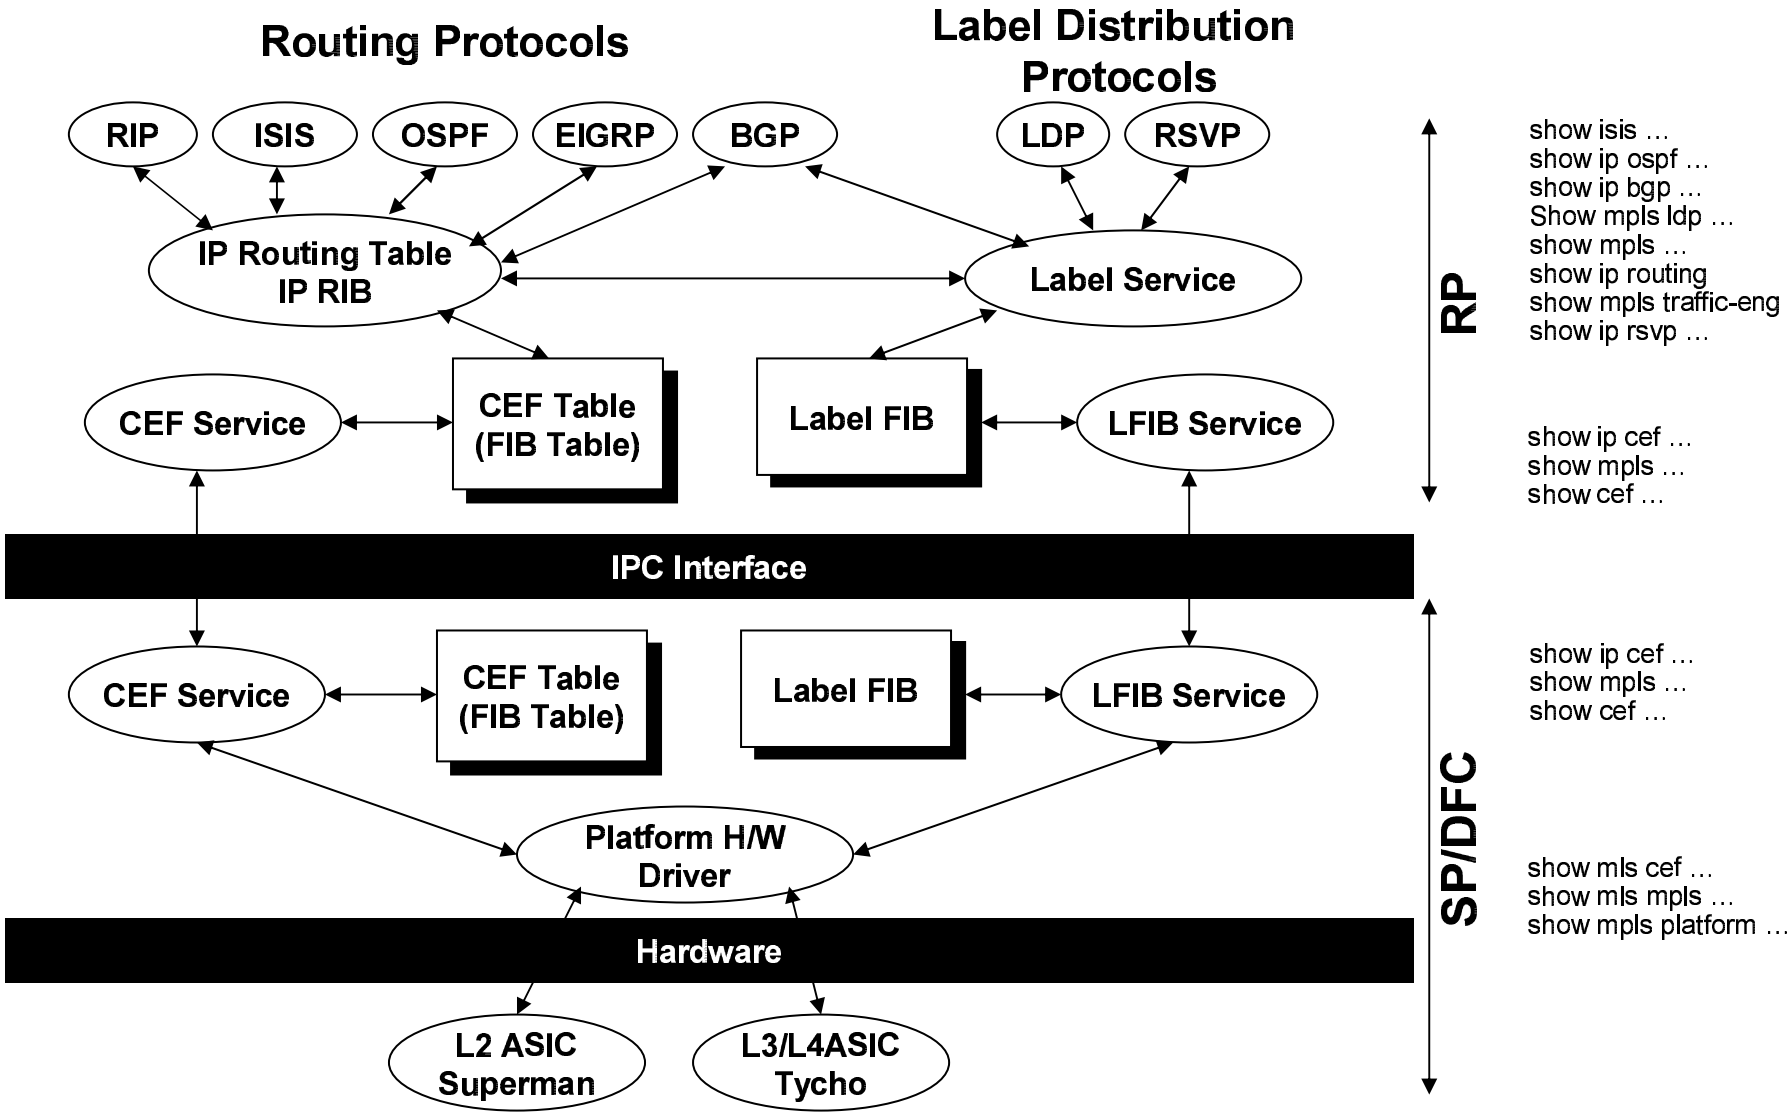

The PFC3B & PFC3BXL MPLS Architecture Components

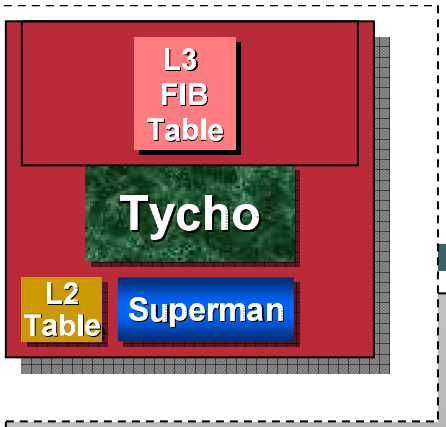

FIB/MPLS/VRF TCAM

Used for fast lookup prefix/label entries

VLAN/VPN Table

MPLS VPN Table

The VLAN Table is indexed with the vlan of the packet stored in the dbus header of packet. The VLAN is used to retrieve the VPN of the packet used for FIB Lookup.

The MPLS VPN Table is indexed when a VPN CAM hit occurs. This can potentially modify the TOS of a packet based on E-LSP or L-LSP QOS*.

Tycho

Internal DSCP

Superman

VPN CAM

The VPN CAM is used for VPN Aggregate Label matching up to 512 entries.

MPLS: IP to TAG Packets

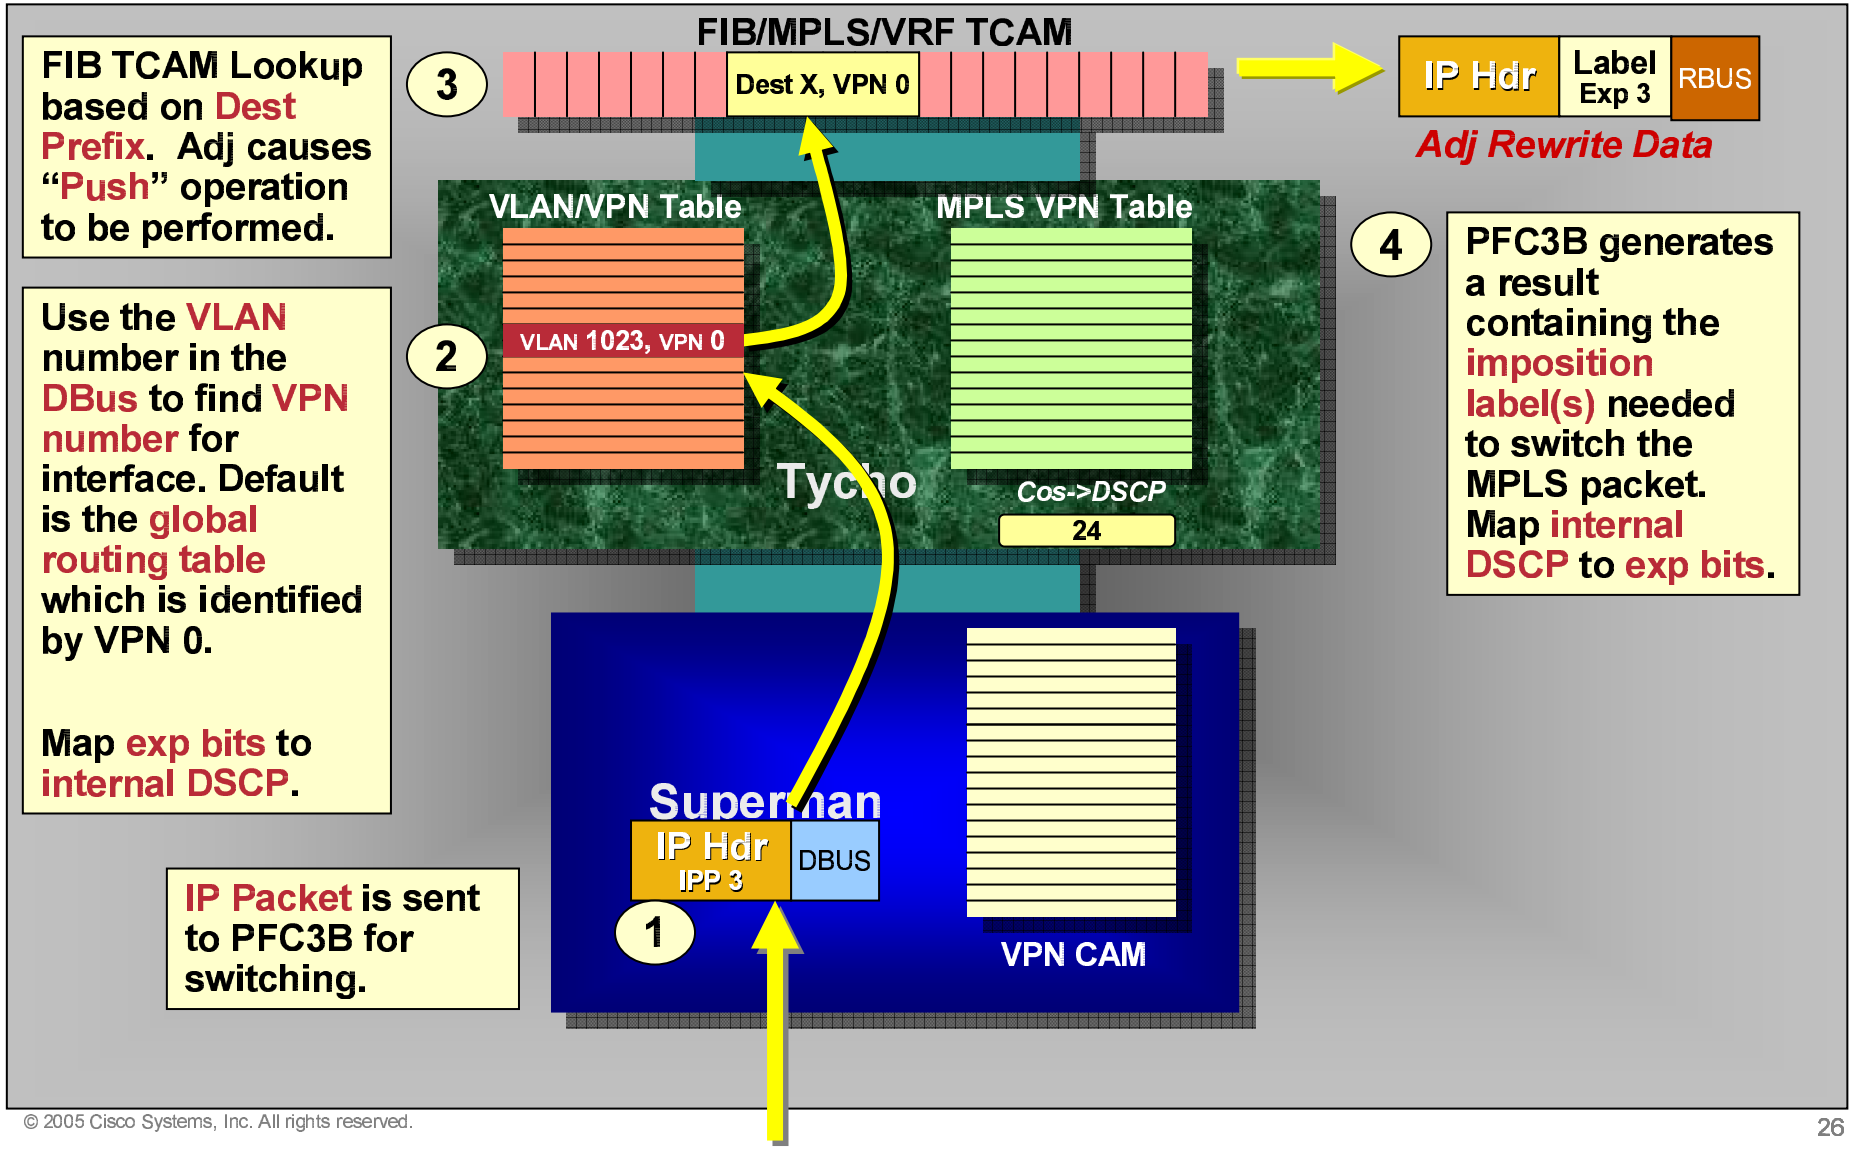

How to look at the VLAN/VPN Table

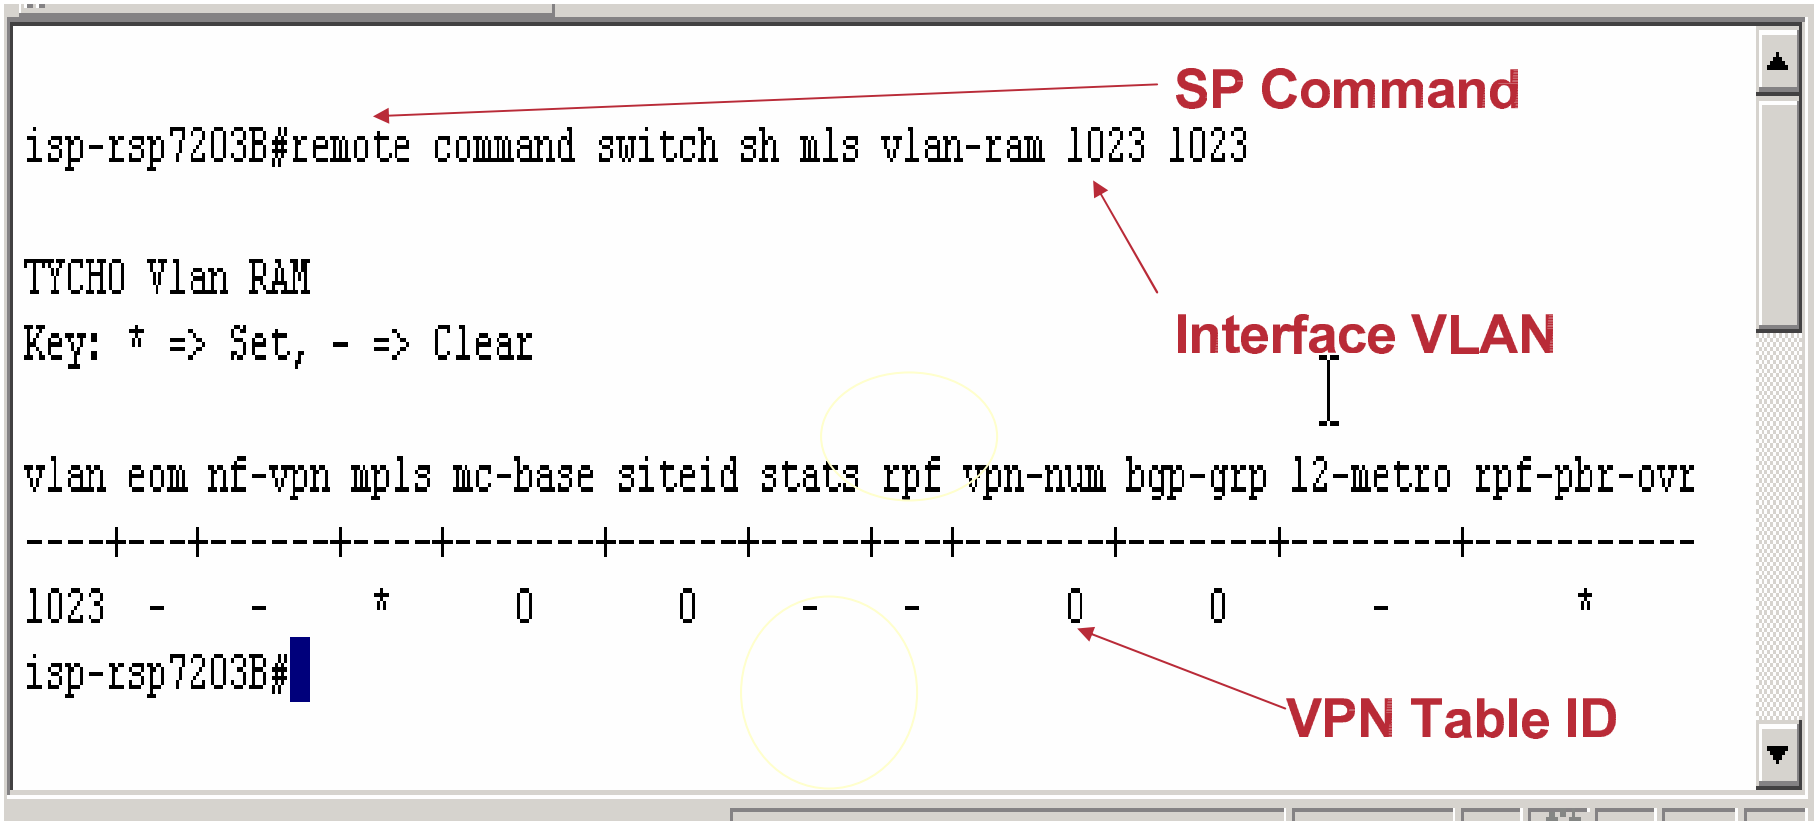


```
isp-rsp7203B#remote command switch sh mls vlan-ram 1023 1023

TYCHO Vlan RAM
Key: * => Set, - => Clear

vlan eom nf-vpn mpls mc-base siteid stats rpf vpn-num bgp-grp l2-metro rpf-pbr-ovr
----+---+------+----+------+------+-----+---+------+------+-------+----------
1023  -    -     *      0       0      -     -       0      0       -         *
isp-rsp7203B#
```

MPLS: TAG to IP Packets

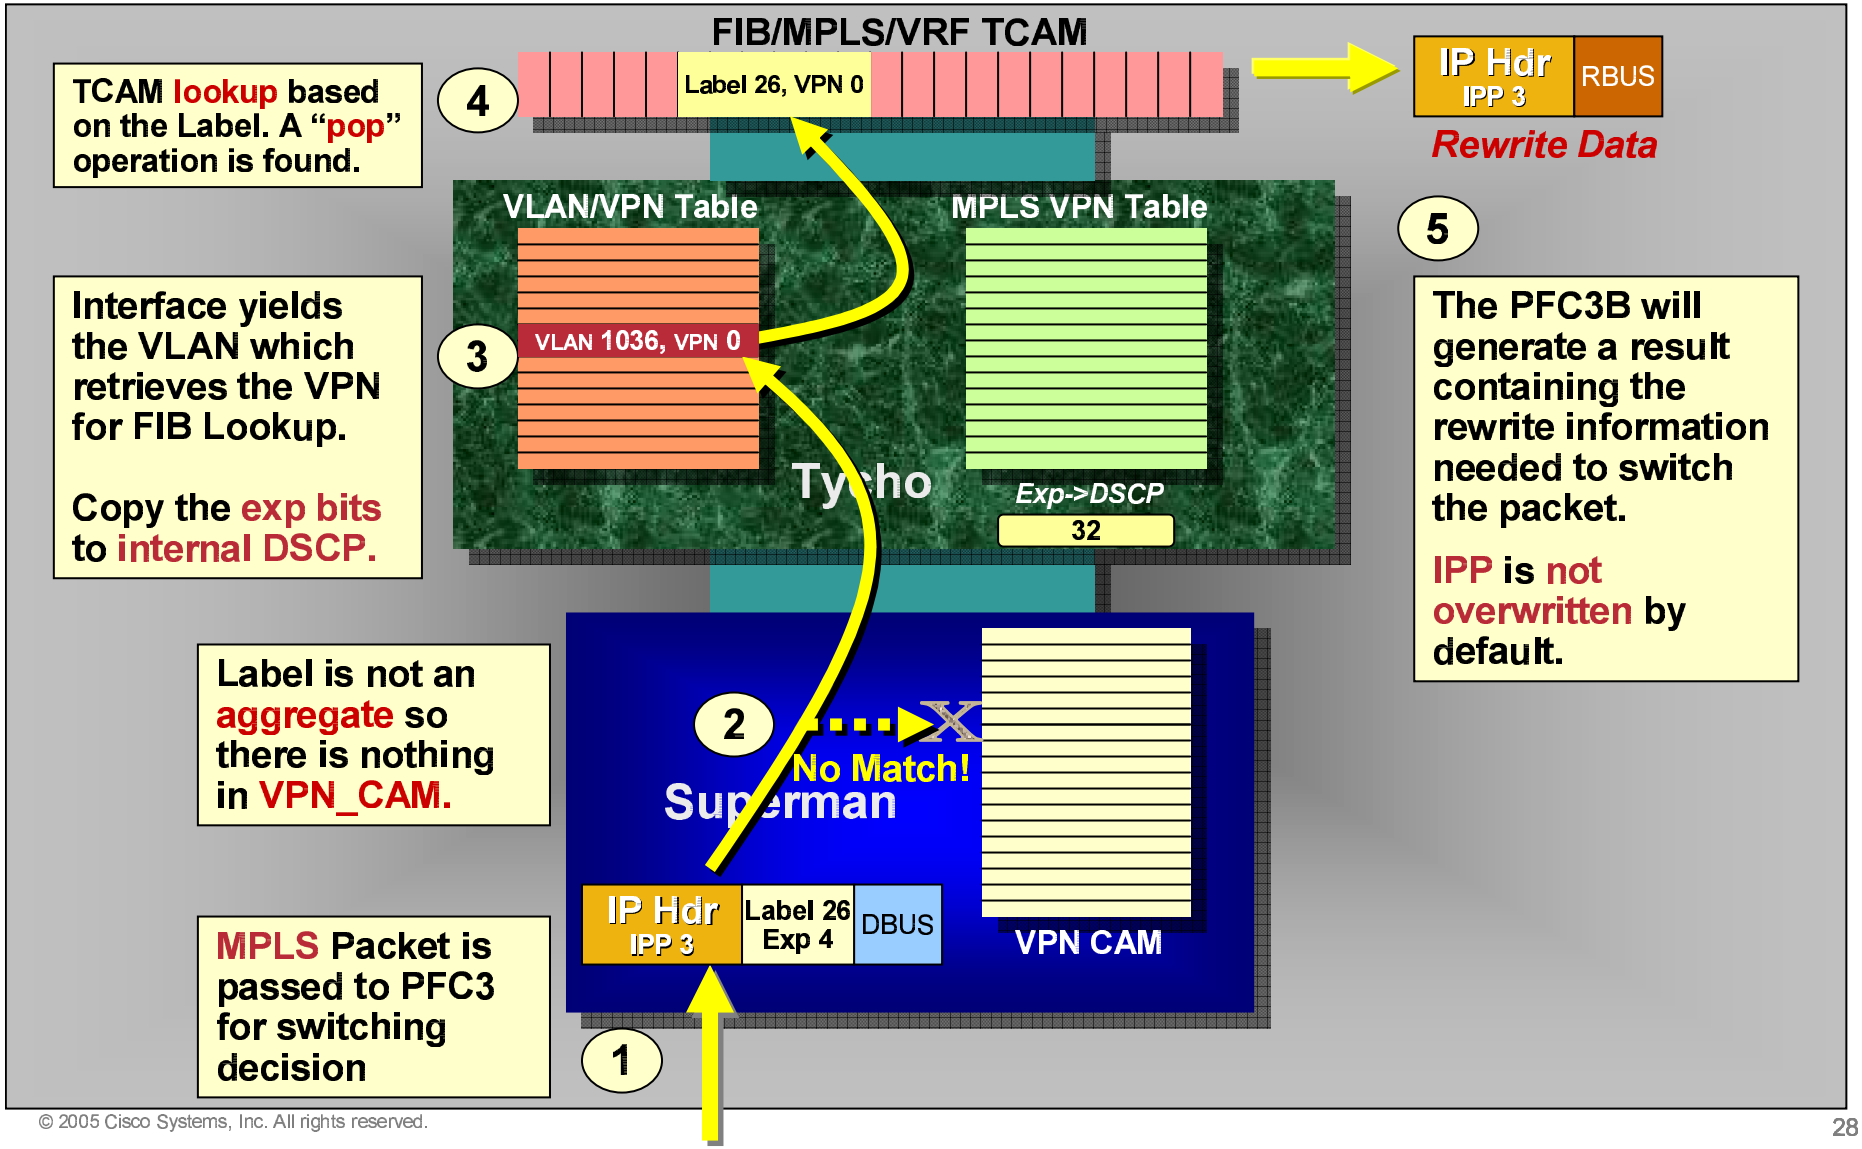

Looking at the MPLS Disposition Entry


```
isp-rsp7203B#sh mls cef mpls label 26

Codes: + - Push label, - - Pop Label          * - Swap Label
Index  Local     Label                  Out i/f
       Label      Op
103    26        (-)                    Fa3/1          , 00e0.f9aa.a054

isp-rsp7203B#
```

Incoming Label

Pop Operation

Rewrite Data

Network Topology for Examples

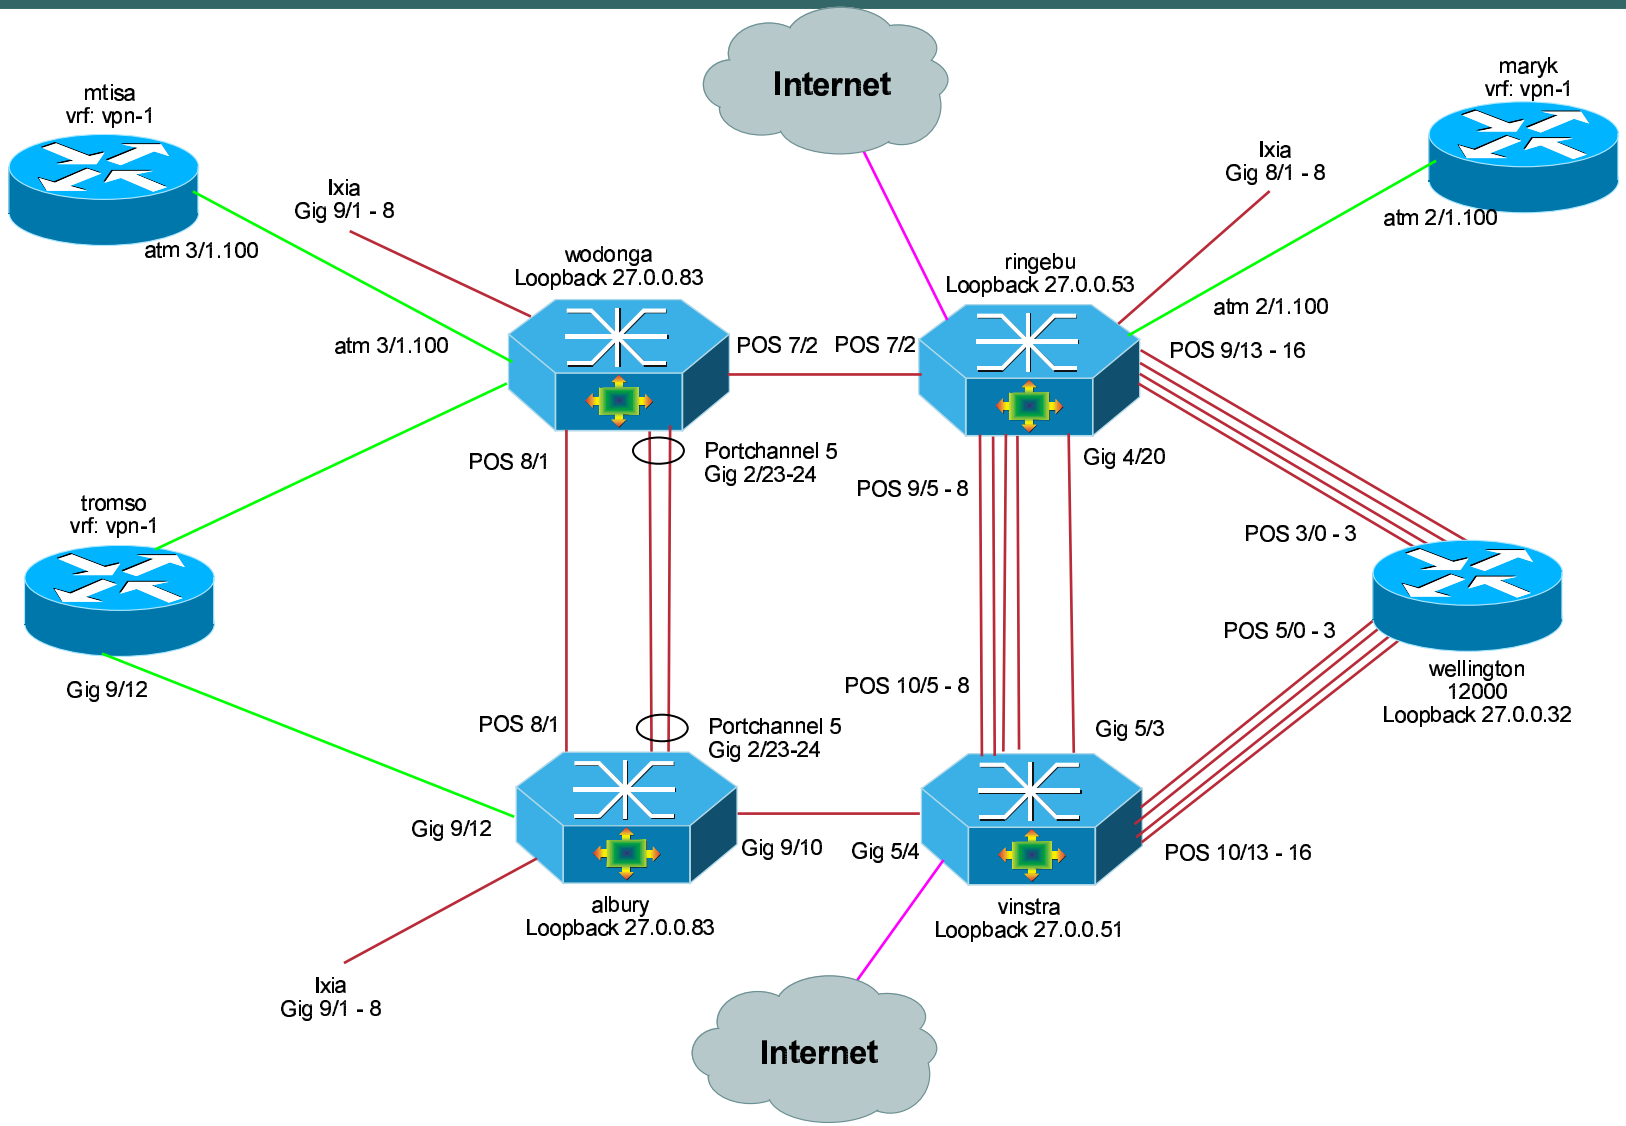

LER - Push

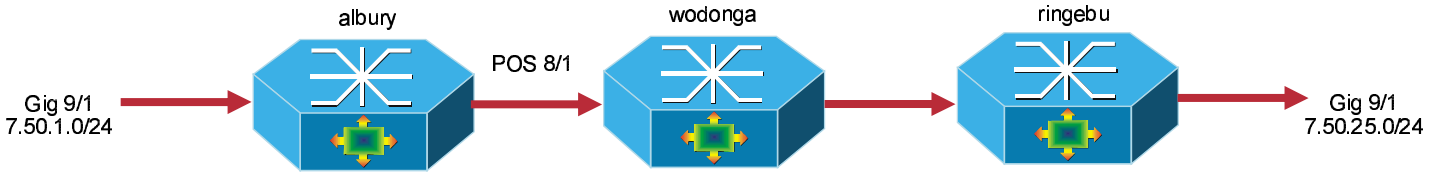

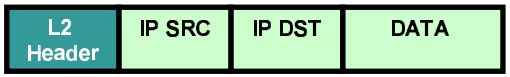

Label 34

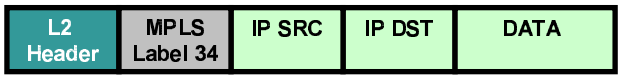


```
albury#sh ip ro 7.50.25.0
Routing entry for 7.50.25.0/24
  Known via "ospf 100", distance 110, metric 18, type intra area
  Last update from 42.52.1.1 on POS8/1, 00:00:13 ago
  Routing Descriptor Blocks:
  * 42.52.1.1, from 27.0.0.53, 00:00:13 ago, via POS8/1
      Route metric is 18, traffic share count is 1


albury#sh mpls forwarding-table 7.50.25.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
61     34          7.50.25.0/24      0          PO8/1       point2point
       MAC/Encaps=4/8, MRU=9216, Tag Stack{34}
       FF030281 00022000
       No output feature configured
    Per-packet load-sharing

albury#sh mls cef 7.50.25.0

Codes: decap - Decapsulation, + - Push Label
Index  Prefix              Adjacency
61019  7.50.25.0/24        PO8/1          34(+)
albury#sh mls cef 7.50.25.0 detail

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(61019  ): E | 1 FFF  0 0 0 0   255.255.255.0
V(61019  ): 8 | 1 0    0 0 0 0   7.50.25.0          (A:278541 ,P:1,D:0,m:0 ,B:0 )
```

H/W CEF Table

VPN ID	IP FIB
0	DST 7.50.25.0/24

Adjacency Table

Interface	Out Label
POS 8/1	34

```
albury#sh mls cef adjacency entry 278541 detail

Index: 278541  smac: 00d0.7995.8400, dmac: 0000.0800.0000
               mtu: 9234, vlan: 1168, dindex: 0x0, l3rw_vld: 1
               format: MPLS, flags: 0x8418
               label0: 0, exp: 0, ovr: 0
               label1: 0, exp: 0, ovr: 0
               label2: 34, exp: 0, ovr: 0
               op: PUSH_LABEL2
               packets: 50899, bytes: 4745124


albury#sh vlan internal usage

VLAN Usage
---- -----------------
...
1168 POS8/1
...
```

LSR – Swap

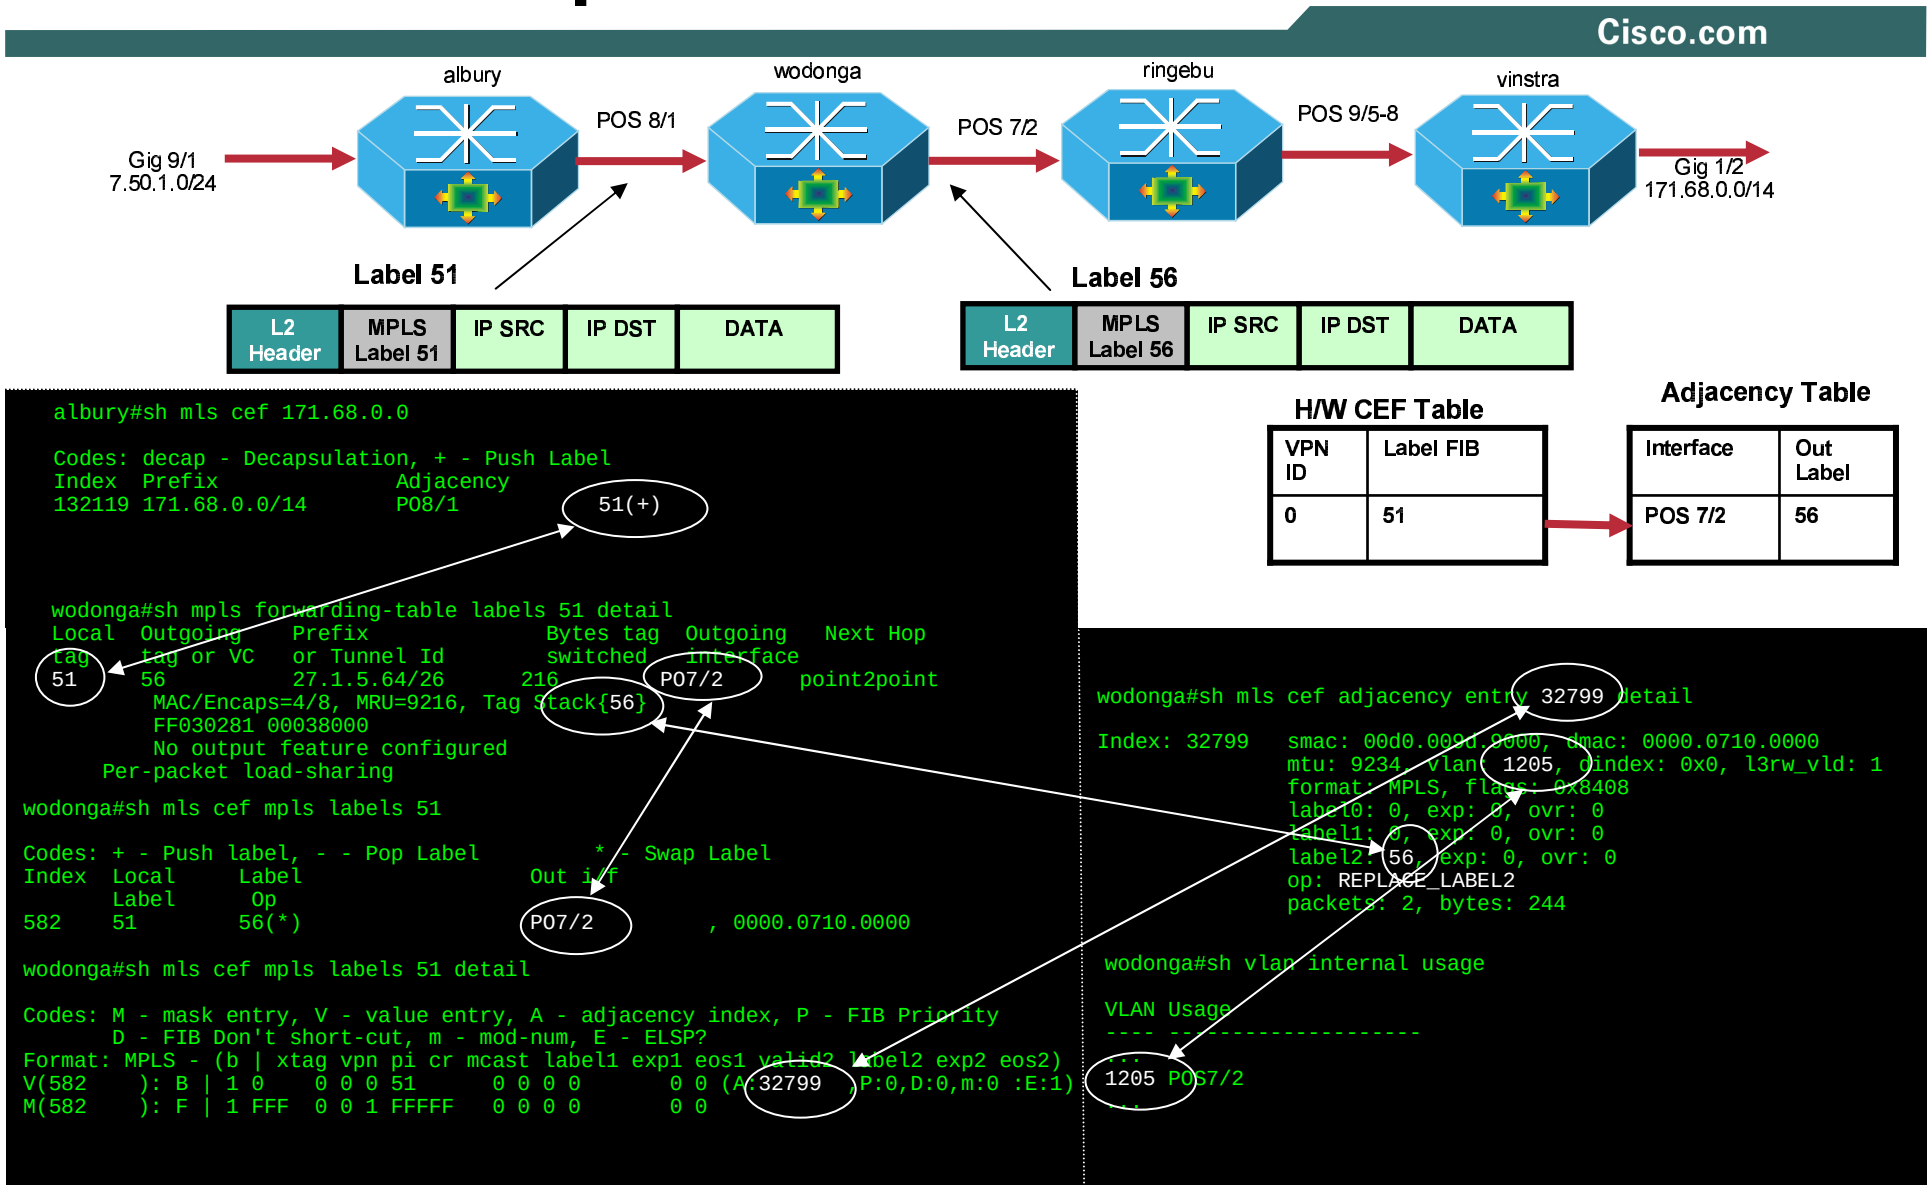

LSR/LER - Pop

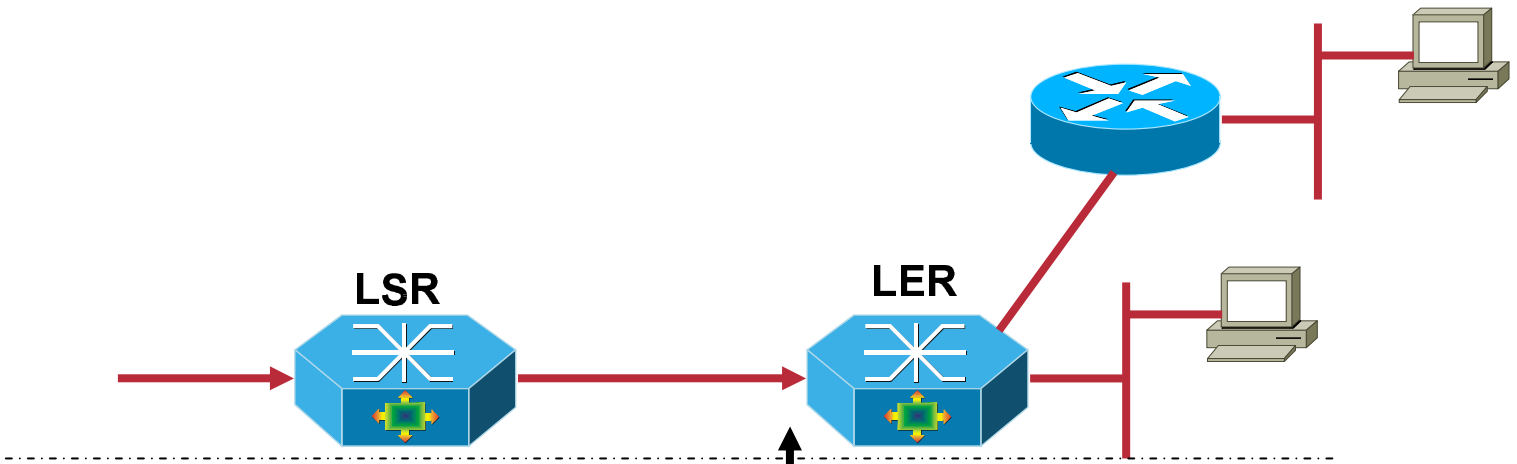

Penultimate Hop Popping

Label 20

L2 Header	MPLS Label20	IP SRC	IP DST	DATA

L2 Header	IP SRC	IP DST	DATA

Explicit Null

L2 Header	MPLS Label20	IP SRC	IP DST	DATA

Label 20

L2 Header	MPLS Null Label	IP SRC	IP DST	DATA

Null Label

LSR – PHP Pop (Implicit Null)

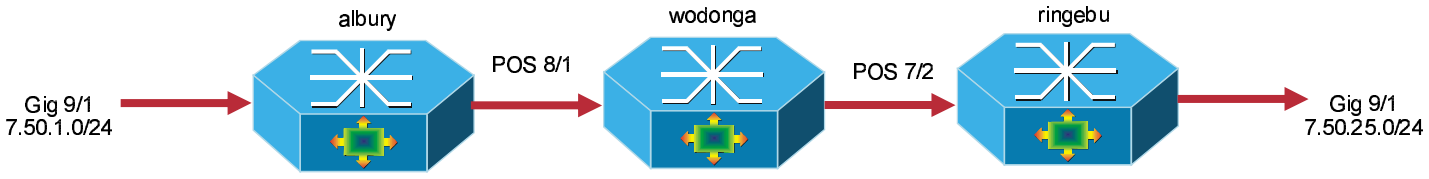

Label 34

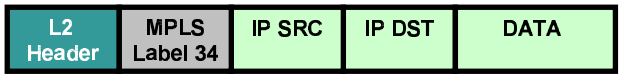

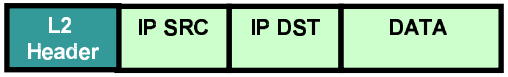

H/W CEF Table

VPN ID	IP FIB
0	34

Adjacency Table

Interface	Out Label
POS 7/2	-

```
wodonga#sh mpls forwarding-table labels 34 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
34     Pop tag     7.50.25.0/24      0          P07/2       point2point
       MAC/Encaps=4/4, MRU=9220, Tag Stack{}
       FF030281
       No output feature configured
    Per-packet load-sharing


wodonga#sh mls cef mpls labels 34

Codes: + - Push label, - - Pop Label          * - Swap Label
Index  Local     Label                 Out I/f
       Label      Op
565    34        (-)                   P07/2          , 0000.0710.0000
wodonga#sh mls cef mpls labels 34 detail

Codes: M - mask entry, V - value entry, A - adjacency index, P - FIB Priority
       D - FIB Don't short-cut, m - mod-num, E - ELSP?
Format: MPLS - (b | xtag vpn pi cr mcast label1 exp1 eos1 valid2 label2 exp2 eos2)
V(565    ): B | 1 0    0 0 0 34      0 0 0 0       0 0  A:81921  ,P:0,D:0,m:0 :E:1)
M(565    ): F | 1 FFF  0 0 1 FFFFF   0 0 0 0       0 0
```

```
wodonga#sh mls cef adjacency entry 81921 detail

Index: 81921    smac: 00d0.009d.9000, dmac: 0000.0710.0000
                mtu: 9234, vlan: 1205, dindex: 0x0, l3rw_vld: 1
                format: MPLS, flags: 0x1000008408
                label0: 0, exp: 0, ovr: 0
                label1: 0, exp: 0, ovr: 0
                label2: 0, exp: 0, ovr: 0
                op: POP
                packets: 10, bytes: 1220


wodonga#sh vlan internal usage

VLAN Usage
----------------------
...
1205 POS7/2
...
```

LER – Explicit Null

mpls ldp explicit-null

Cisco.com

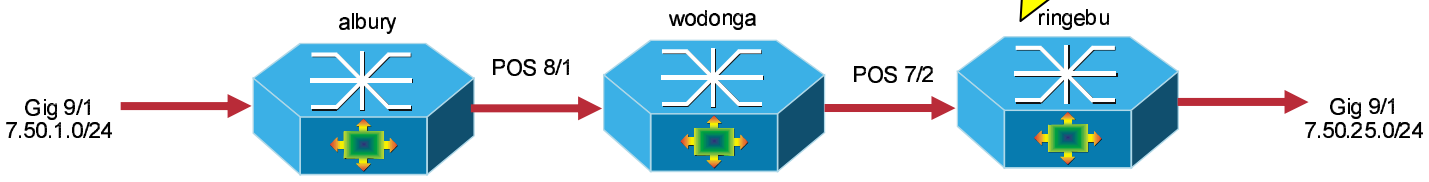

Label 34

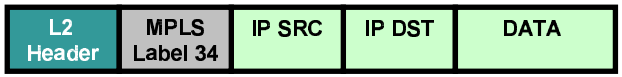

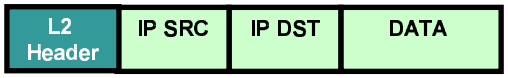

H/W CEF Table

VPN ID	IP FIB
0	34

Adjacency Table

Interface	Out Label
POS 7/2	0

```
albury#sh mpls forwarding-table 7.50.25.0
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
61     34          7.50.25.0/24      0          PO8/1       point2point
albury#
```

```
wodonga#sh mls cef mpls labels 34

Codes: + - Push label, - - Pop Label         * - Swap Label
Index  Local     Label                 Out i/f
       Label      Op
556    34        0(*)                  PO7/2          , 0000.0710.0000

wodonga#sh mls cef mpls labels 34 detail

Codes: M - mask entry, V - value entry, A - adjacency index, P - FIB Priority
       D - FIB Don't short-cut, m - mod-num, E - ELSP?
Format: MPLS - (b | xtag vpn pi cr mcast label1 exp1 eos1 valid2 label2 exp2 eos2)
V(556    ): B | 1 0    0 0 0 34      0 0 0 0       0 0  A:32804  ,P:0,D:0,m:0 :E:1)
M(556    ): F | 1 FFF  0 0 1 FFFFF   0 0 0 0       0 0
```

```
wodonga#sh mls cef adjacency entry 32804 detail

Index: 32804    smac: 00d0.000d.9000, dmac: 0000.0710.0000
                mtu: 9234, vlan: 1205, dindex: 0x0, l3rw_vld: 1
                format: MPLS, flags: 0x8408
                label0: 0, exp: 0, ovr: 0
                label1: 0, exp: 0, ovr: 0
                label2: 0, exp: 0, ovr: 0
                op: REPLACE_LABEL2
                packets: 20, bytes: 2440

wodonga#sh vlan internal usage

VLAN Usage
----------------------
...
1205 POS7/2
...
```

LER – Route behind non-MPLS router

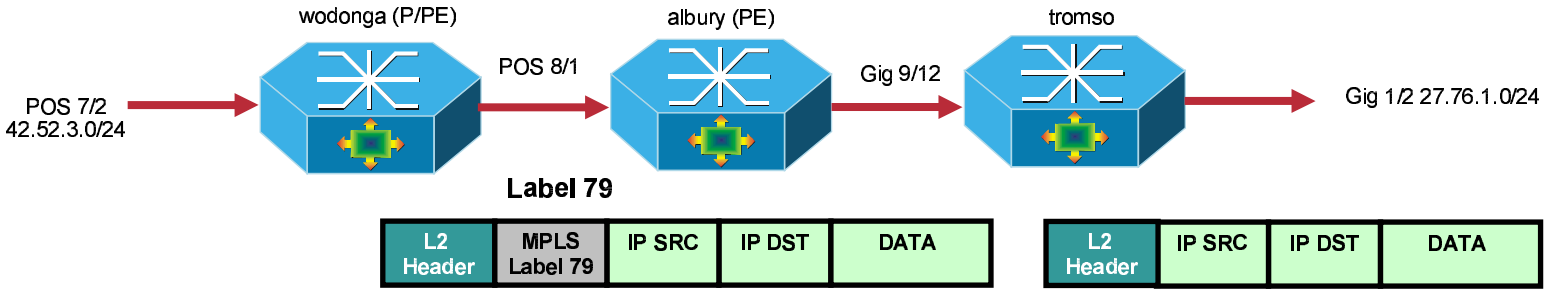


```
wodonga#sh mpls forwarding-table 27.76.1.85 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
185    79          27.76.1.0/24      0          PO8/1      point2point
        MAC/Encaps=4/8, MRU=9216, Tag Stack{79}
        FF030281 0004F000
        No output feature configured
    Per-packet load-sharing

albury#sh mpls forwarding-table labels 79 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
79     Untagged    27.76.1.0/24      0          Gi9/12     42.220.254.1
        MAC/Encaps=0/0, MRU=9220, Tag Stack{}
        No output feature configured
    Per-packet load-sharing

albury#sh mls cef mpls labels 79

Codes: + - Push label, - - Pop Label         * - Swap Label
Index  Local     Label                 Out i/f
       Label      Op
928    79        (-)                   Gi9/12        , 0030.7b4a.940a

albury#sh mls cef mpls labels 79 detail

Codes: M - mask entry, V - value entry, A - adjacency index, P - FIB Priority
       D - FIB Don't short-cut, m - mod-num, E - ELSP?
Format: MPLS - (b | xtag vpn pi cr mcast label1 exp1 eos1 valid2 label2 exp2 eos2)
V(928    ): B | 1 0    0 0 0 79      0 0 0 0       0 0 (A:311333  P:0,D:0,m:0 :E:1)
M(928    ): F | 1 FFF  0 0 1 FFFFF   0 0 0 0       0 0
```

H/W CEF Table

VPN ID	Label FIB
0	79

Adjacency Table

Interface	Out Label
POS 7/2	0

```
albury#sh mls cef adjacency entry 311333 detail

Index: 311333  smac: 00d0.7995.9400, dmac: 0030.7b4a.940a
               mtu: 9234, vlan: 1157, dindex: 0x0, l3rw_vld: 1
               format: MPLS, flags: 0x1000008408
               label0: 0, exp: 0, ovr: 0
               label1: 0, exp: 0, ovr: 0
               label2: 0, exp: 0, ovr: 0
               op: POP
               packets: 0, bytes: 0


wodonga#sh vlan internal usage

VLAN Usage
----------------------
...
1157 Gi9/12
...
```

Load Balancing and Etherchannel

L3 Load Balancing

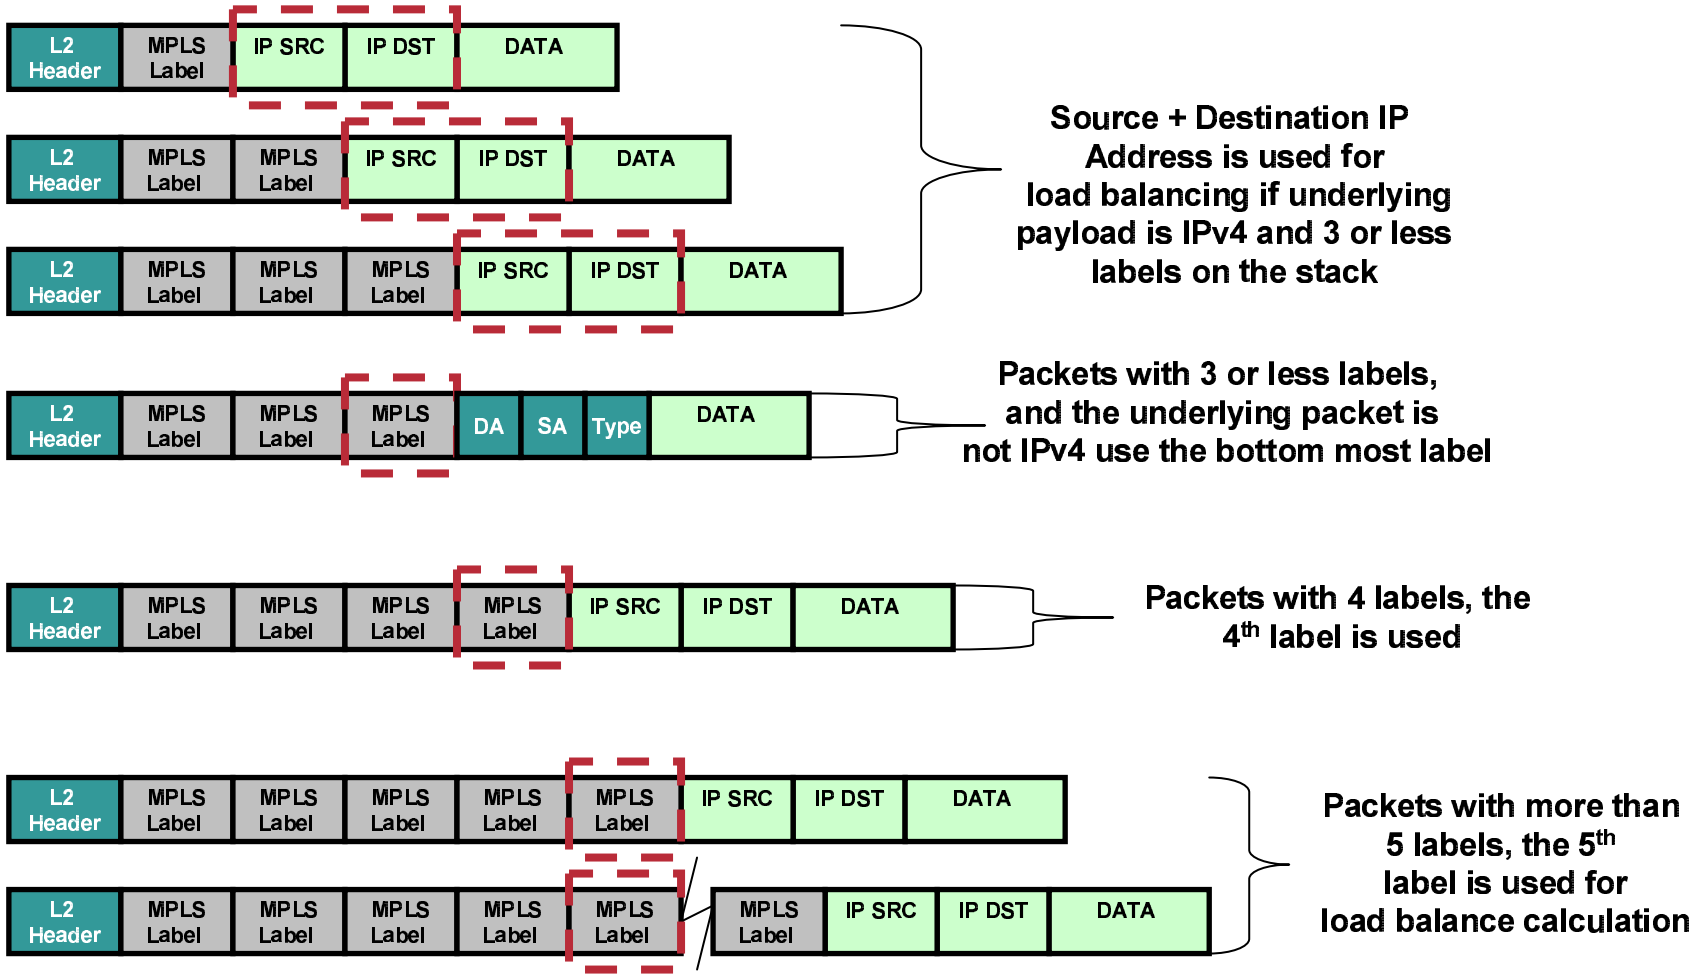

Etherchannel Load Balancing

Cisco.com

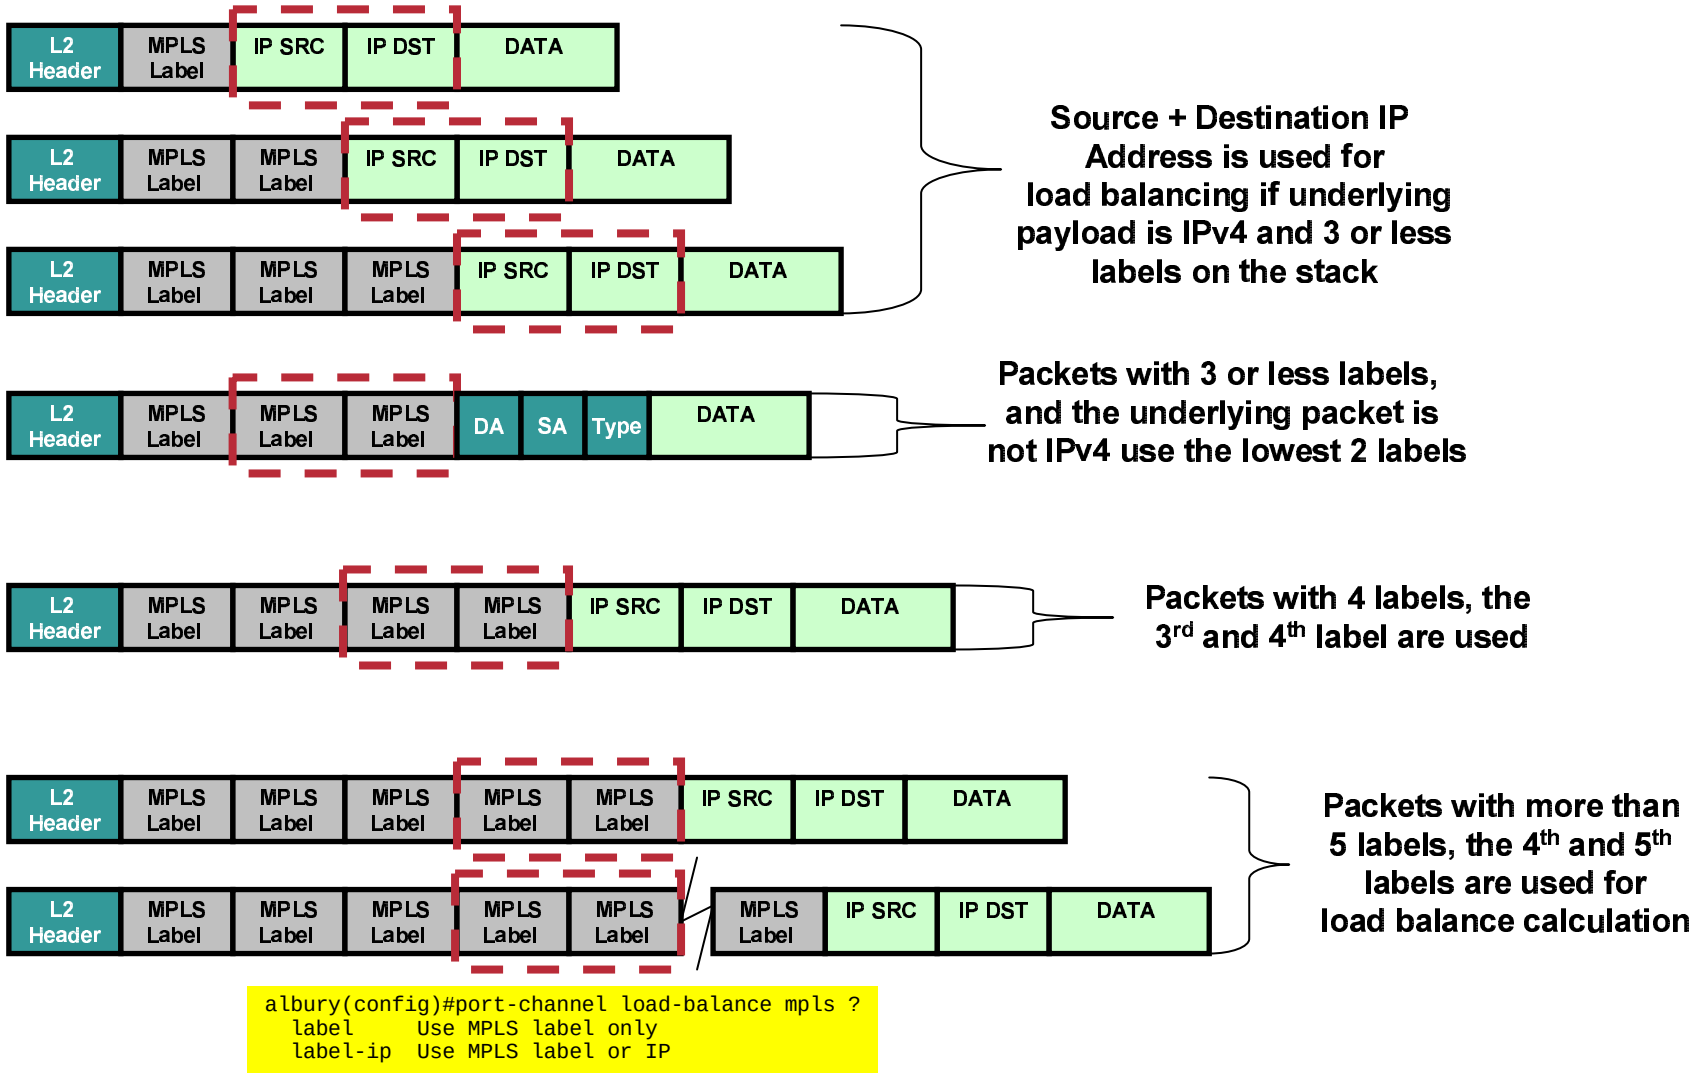


```
albury(config)#port-channel load-balance mpls ?
  label     Use MPLS label only
  label-ip  Use MPLS label or IP
```

IP to Label L3 Load Balance

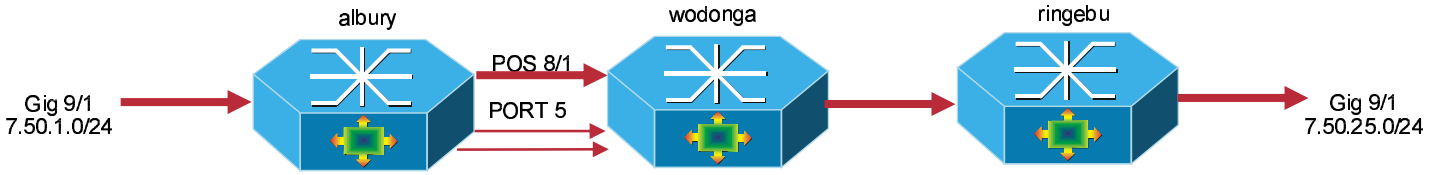


```
albury#sh ip ro 7.50.26.0
Routing entry for 7.50.26.0/24
  Known via "ospf 100", distance 110, metric 18, type intra area
  Last update from 42.52.1.1 on POS8/1, 00:08:40 ago
  Routing Descriptor Blocks:
  * 42.52.1.1, from 27.0.0.53, 00:08:40 ago, via POS8/1
      Route metric is 18, traffic share count is 1
    42.54.1.2, from 27.0.0.53, 00:08:40 ago, via Port-channel5
      Route metric is 18, traffic share count is 1


albury#sh mpls forwarding-table 7.50.26.0
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
62     35          7.50.26.0/24      0          PO8/1       point2point
       35          7.50.26.0/24      0          Po5         42.54.1.2
albury#sh mls cef 7.50.26.0 detail

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(85817  ): E | 1 FFF  0 0 0 0   255.255.255.0
V(85817  ): 8 | 1 0    0 0 0 0   7.50.26.0          (A:213062 ,P:1,D:0,m:1 ,B:0 )
```

H/W CEF Table

VPN ID	IP FIB
0	DST 7.50.25.0/24

Adjacency Table

Interface	Out Label
POS 8/1 PortChannel 5	35

```
albury#sh mls cef adjacency entry 213062 detail

Index: 213062  smac: 00d0.7995.9400, dmac: 00d0.009d.9000
               mtu: 9234, vlan: 1013, dindex: 0x0, l3rw_vld: 1
               format: MPLS, flags: 0x8418
               label0: 0, exp: 0, ovr: 0
               label1: 0, exp: 0, ovr: 0
               label2: 35, exp: 0, ovr: 0
               op: PUSH_LABEL2
               packets: 0, bytes: 0

albury#sh mls cef adjacency entry 213063 detail

Index: 213063  smac: 00d0.7995.9400, dmac: 0000.0800.0000
               mtu: 9234, vlan: 1168, dindex: 0x0, l3rw_vld: 1
               format: MPLS, flags: 0x8418
               label0: 0, exp: 0, ovr: 0
               label1: 0, exp: 0, ovr: 0
               label2: 35, exp: 0, ovr: 0
               op: PUSH_LABEL2
               packets: 0, bytes: 0
```

Label to Label Load Balance

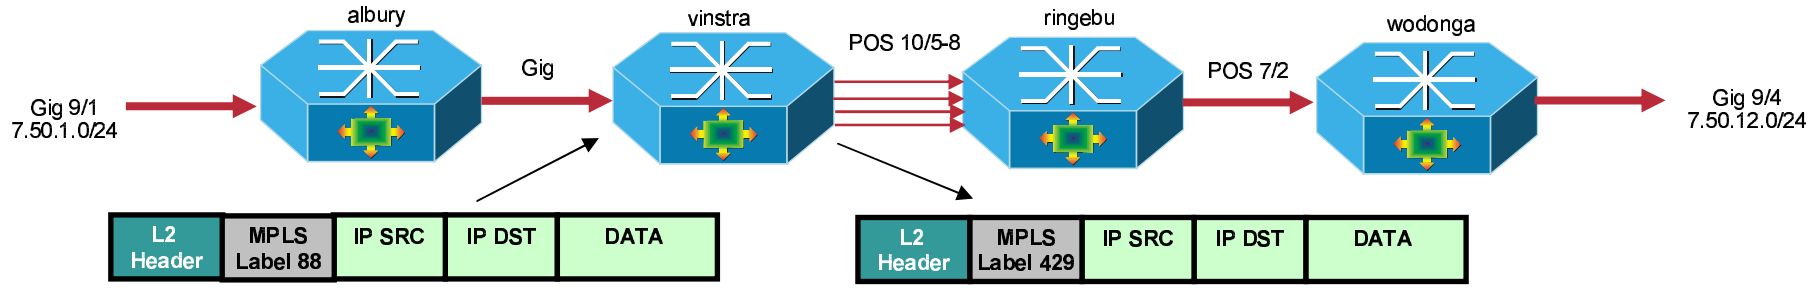


```
albury#sh ip ro 7.50.12.0
Routing entry for 7.50.12.0/24
  Known via "ospf 100", distance 110, metric 88, type intra area
  Last update from 42.55.200.2 on GigabitEthernet9/10, 00:08:42 ago
      ...
albury#sh mpls forwarding-table 7.50.12.0
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
89      88          7.50.12.0/24     0          Gi9/10     42.55.200.2
vinstra#sh mpls forwarding-table labels 88
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
88     429         7.50.12.0/24      1620       PO10/5     point2point
       429         7.50.12.0/24      0          PO10/6     point2point
       429         7.50.12.0/24      0          PO10/7     point2point
       429         7.50.12.0/24      1620       PO10/8     point2point

       vinstra#sh mls cef mpls labels 88
Codes: + - Push label, - - Pop Label         * - Swap Label
Index  Local     Label                 Out i/f
       Label      Op
118    88        429(*)                PO10/5        , 0000.0a40.0000
                 429(*)                PO10/6        , 0000.0a50.0000
                 429(*)                PO10/7        , 0000.0a60.0000
                 429(*)                PO10/8        , 0000.0a70.0000
vinstra#sh mls cef mpls labels 88 detail
Codes: M - mask entry, V - value entry, A - adjacency index, P - FIB Priority
       D - FIB Don't short-cut, m - mod-num, E - ELSP?
Format: MPLS - (b | xtag vpn pi cr mcast label1 exp1 eos1 valid2 label2 exp2 eos2)
V(118    ): B | 1 0    0 0 0 88      0 0 0 0       0 0 (A:32803  ,P:0,D:0,m:4 :E:1)
M(118    ): F | 1 FFF  0 0 1 FFFFF   0 0 0 0       0 0
```

H/W CEF Table

VPN ID	IP FIB
0	DST 7.50.25.0/24

Adjacency Table

Interface	Out Label
POS 8/1 PortChan 5	35

```
vinstra#sh mls cef adjacency entry 32803 detail
Index: 32803   smac: 0007.4fa6.bd40, dmac: 0000.0a40.0000
               ...
               label2: 429, exp: 0, ovr: 0
               op: REPLACE_LABEL2
vinstra#sh mls cef adjacency entry 32804 detail
Index: 32804   smac: 0007.4fa6.bd40, dmac: 0000.0a40.0000
               ...
               label2: 429, exp: 0, ovr: 0
               op: REPLACE_LABEL2
vinstra#sh mls cef adjacency entry 32805 detail
Index: 32805   smac: 0007.4fa6.bd40, dmac: 0000.0a50.0000
               ...
               label2: 429, exp: 0, ovr: 0
               op: REPLACE_LABEL2
vinstra#sh mls cef adjacency entry 32806 detail
Index: 32806   smac: 0007.4fa6.bd40, dmac: 0000.0a60.0000
               ...
              label2: 429, exp: 0, ovr: 0
               op: REPLACE_LABEL2
```

RFC2547 VPN Forwarding

MPLS LER VPN

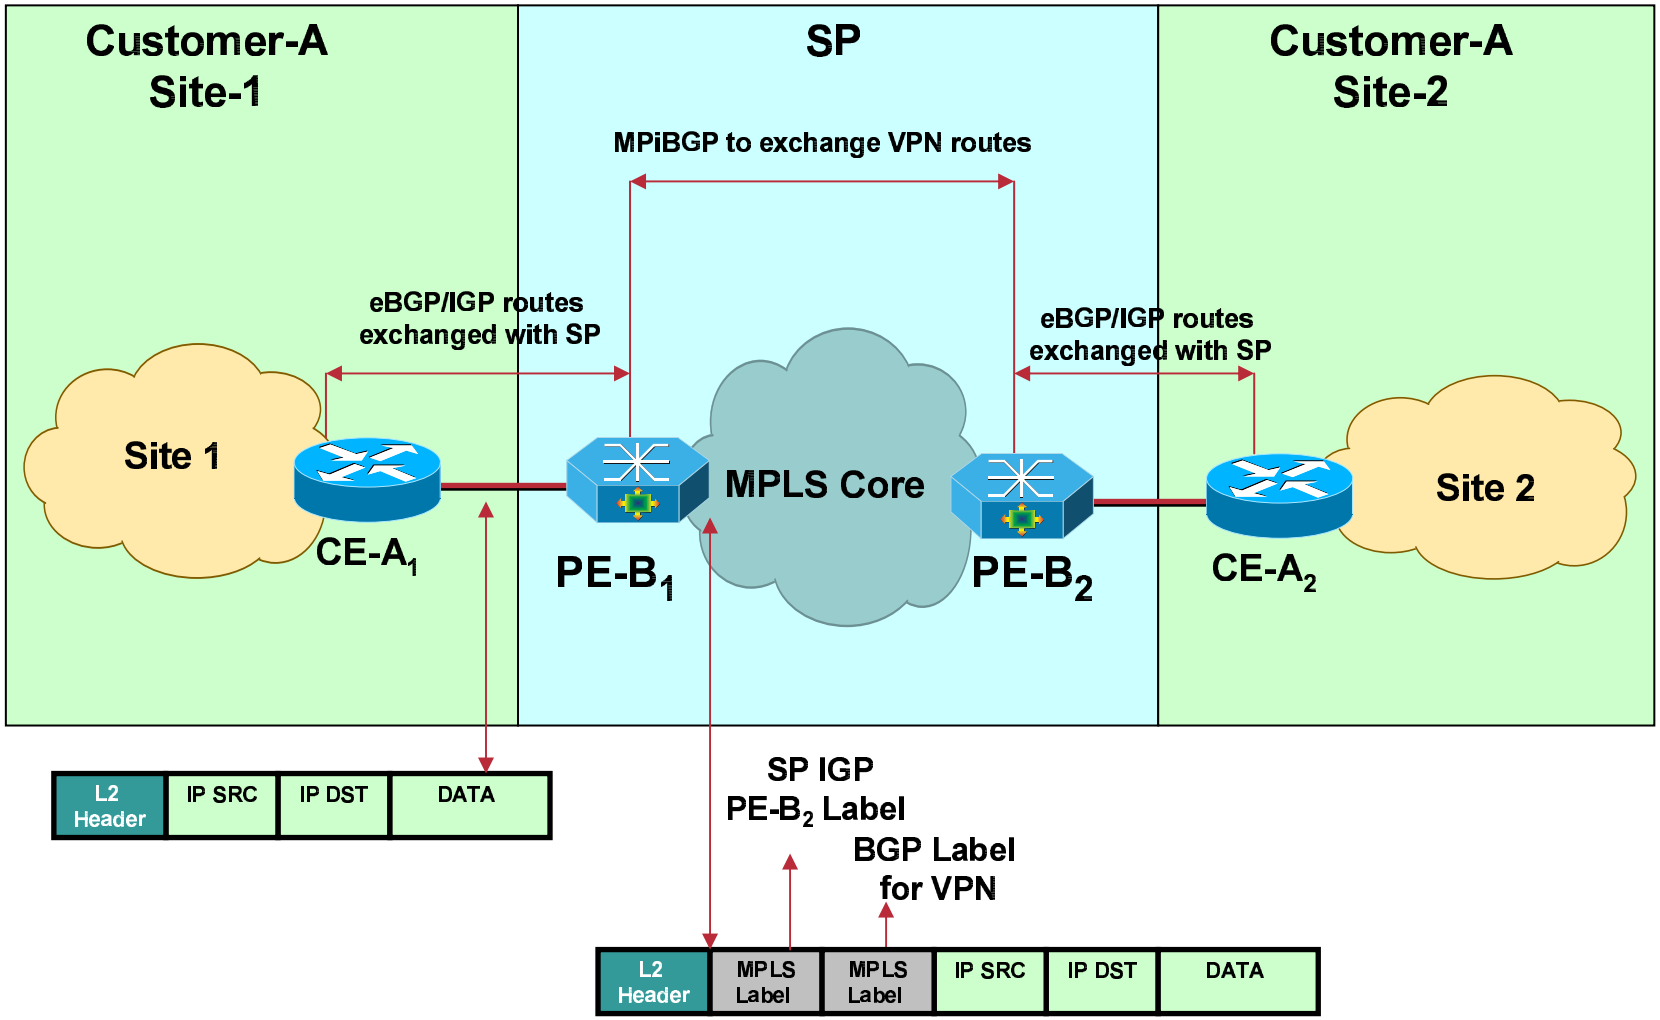

MPLS/VPN: IP to TAG Packets

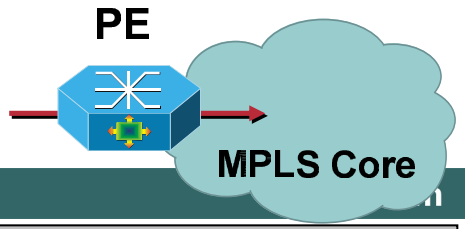

Lookup Dest Address in correct VPN

4

FIB/MPLS/VRF TCAM

Dest n, VPN 256

IP Hdr IPP 4 | Label 4118 exp 4 | Label 16 exp 4 | RBUS

5

Rewrite Data

VLAN/VPN Table

MPLS VPN Table

The VPN for this interface is retrieved.

3

VLAN 1027, VPN 256

Tycho

COS-> DSCP

32

VLAN is used to index into VLAN/VPN Table

2

PFC3B generates a result containing the two imposition labels needed to switch the MPLS/VPN packet.

The exp bits are mapped from the internal DSCP

VLAN is stored in DBUS Header

Superman

IP Header IPP 4 | DBUS

VPN CAM

MPLS/VPN IP Packet is passed to PFC3 for switching decision

1

Examining Hardware Imposition Labels


```
isp-rsp7203B#sh mls cef vrf test

Codes: decap - Decapsulation, + - Push Label
Index   Prefix               Adjacency
66      0.0.0.0/32           receive
67      255.255.255.255/32   receive
93      55.0.0.1/32          receive
94      55.0.0.0/32          receive
95      55.255.255.255/32    receive
3201    224.0.0.0/24         receive
134403  55.0.0.0/8           PO2/1,           0000.0200.0000
134411  98.0.0.0/8           GE1/4,           4118(+),16(+)
134412  78.0.0.0/8           GE1/4,           4117(+),16(+)
134413  31.0.0.0/8           GE1/4,           4116(+),16(+)
134414  29.0.0.0/8           GE1/4,           4115(+),16(+)
134415  128.0.0.0/8          PO2/1,           0000.0200.0000
134416  126.0.0.0/8          PO2/1,           0000.0200.0000
isp-rsp7203B#
```

Dest Prefix

Output Interface

Imposition Labels

MPLS/VPN: TAG to TAG Packets

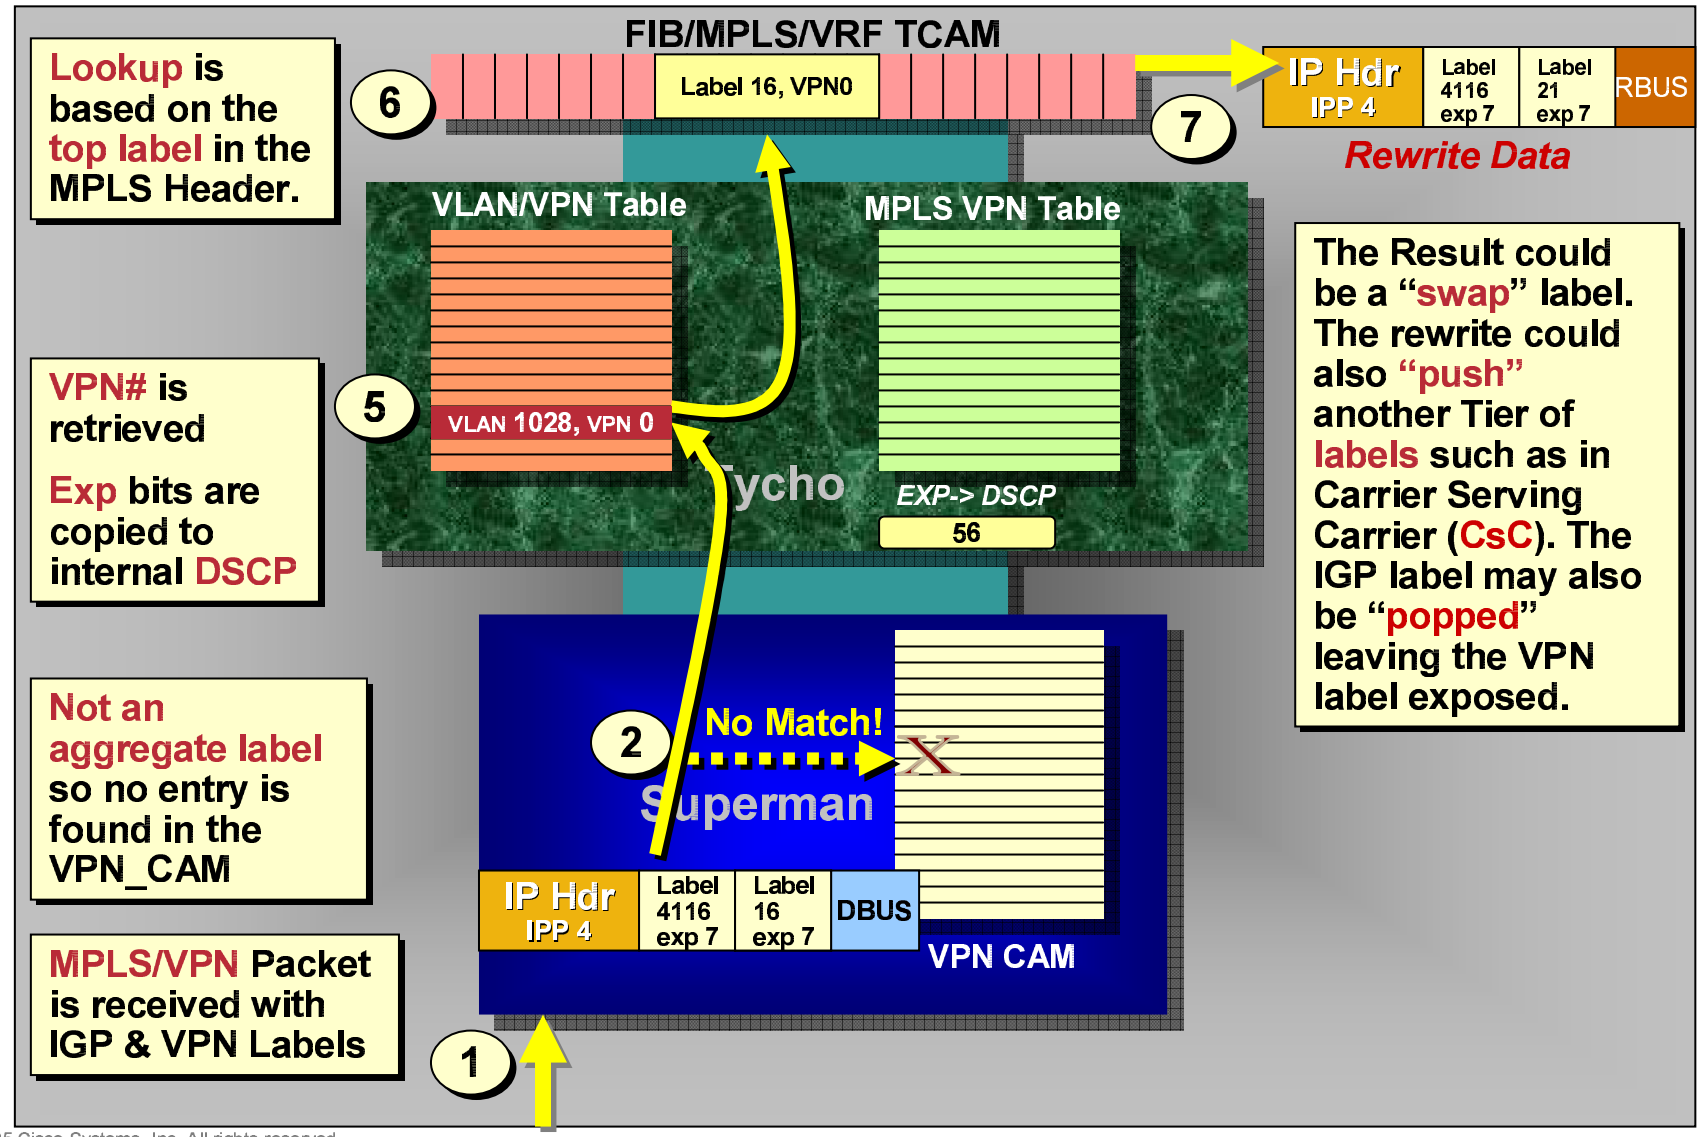

MPLS/VPN: TAG to IP Packets
Case: *Aggregate Label*

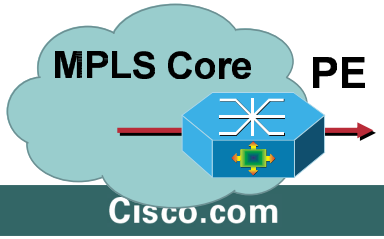

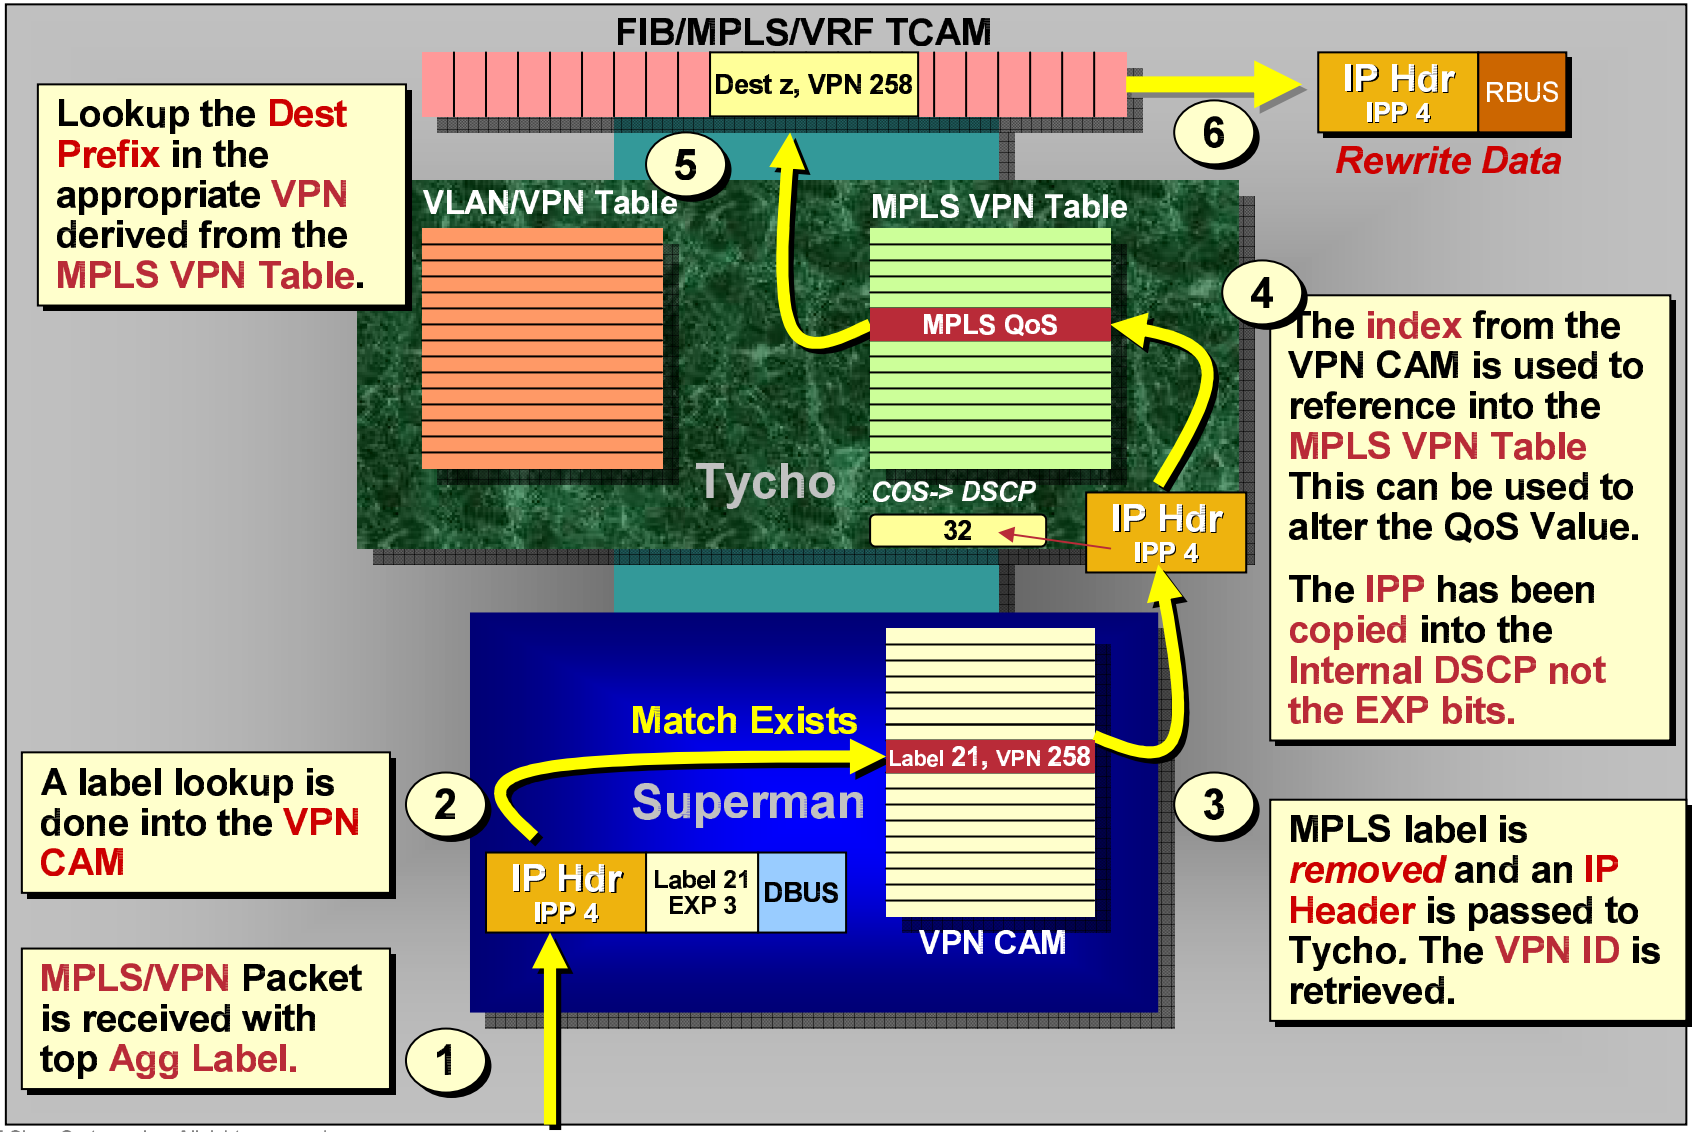

Looking at the Superman VPN CAM

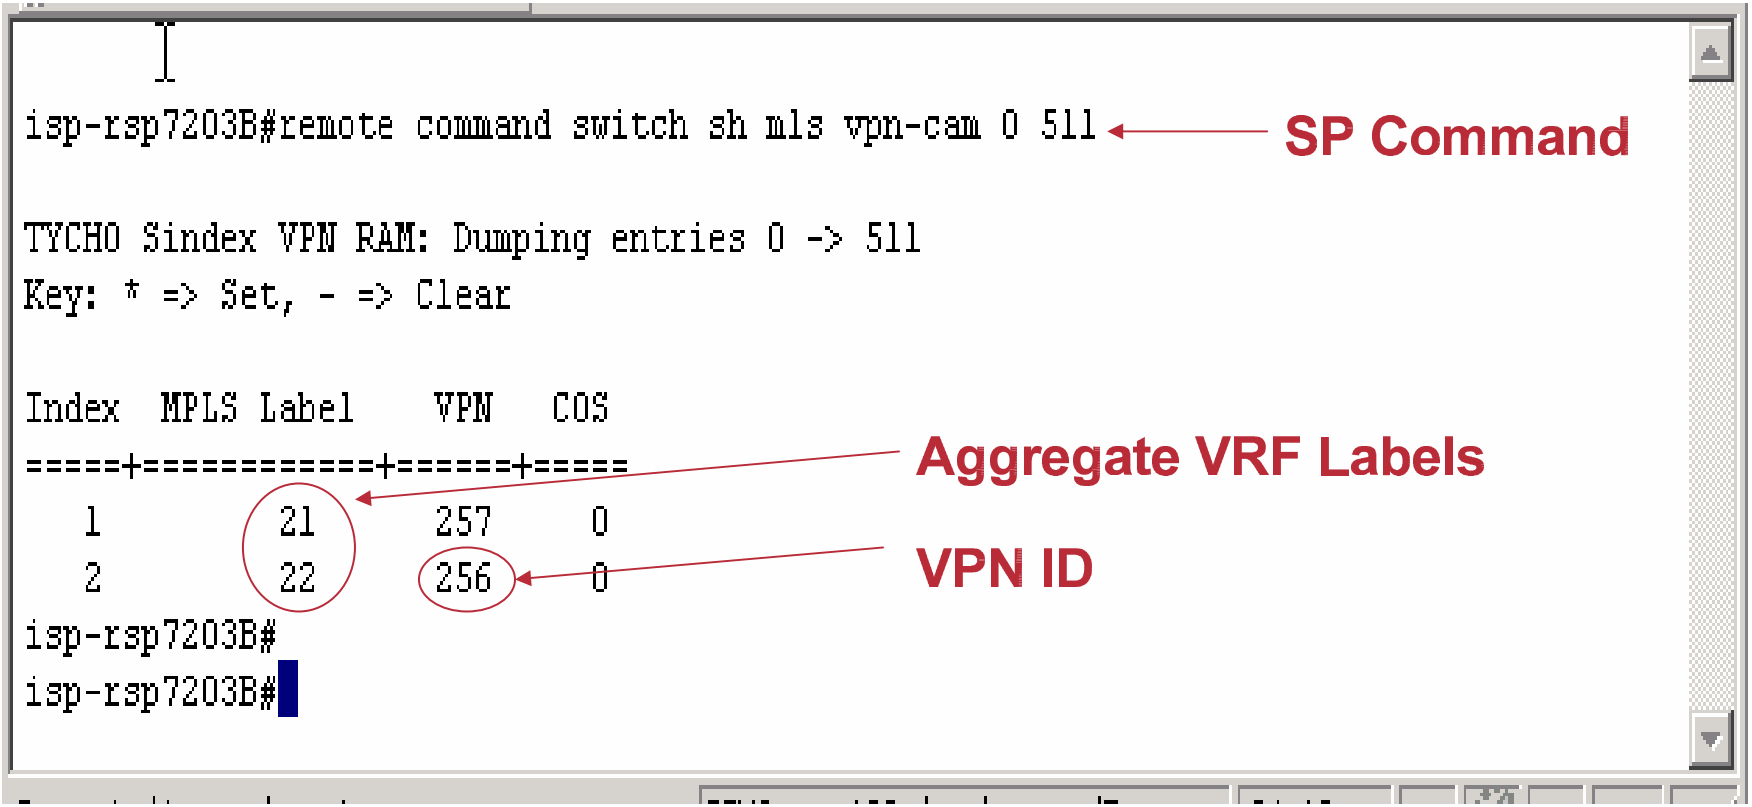

MPLS/VPN TAG to IP Packets
Case: *Aggregate* Labels above 512 VRFs

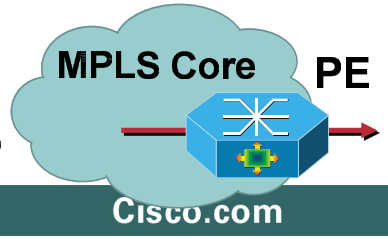

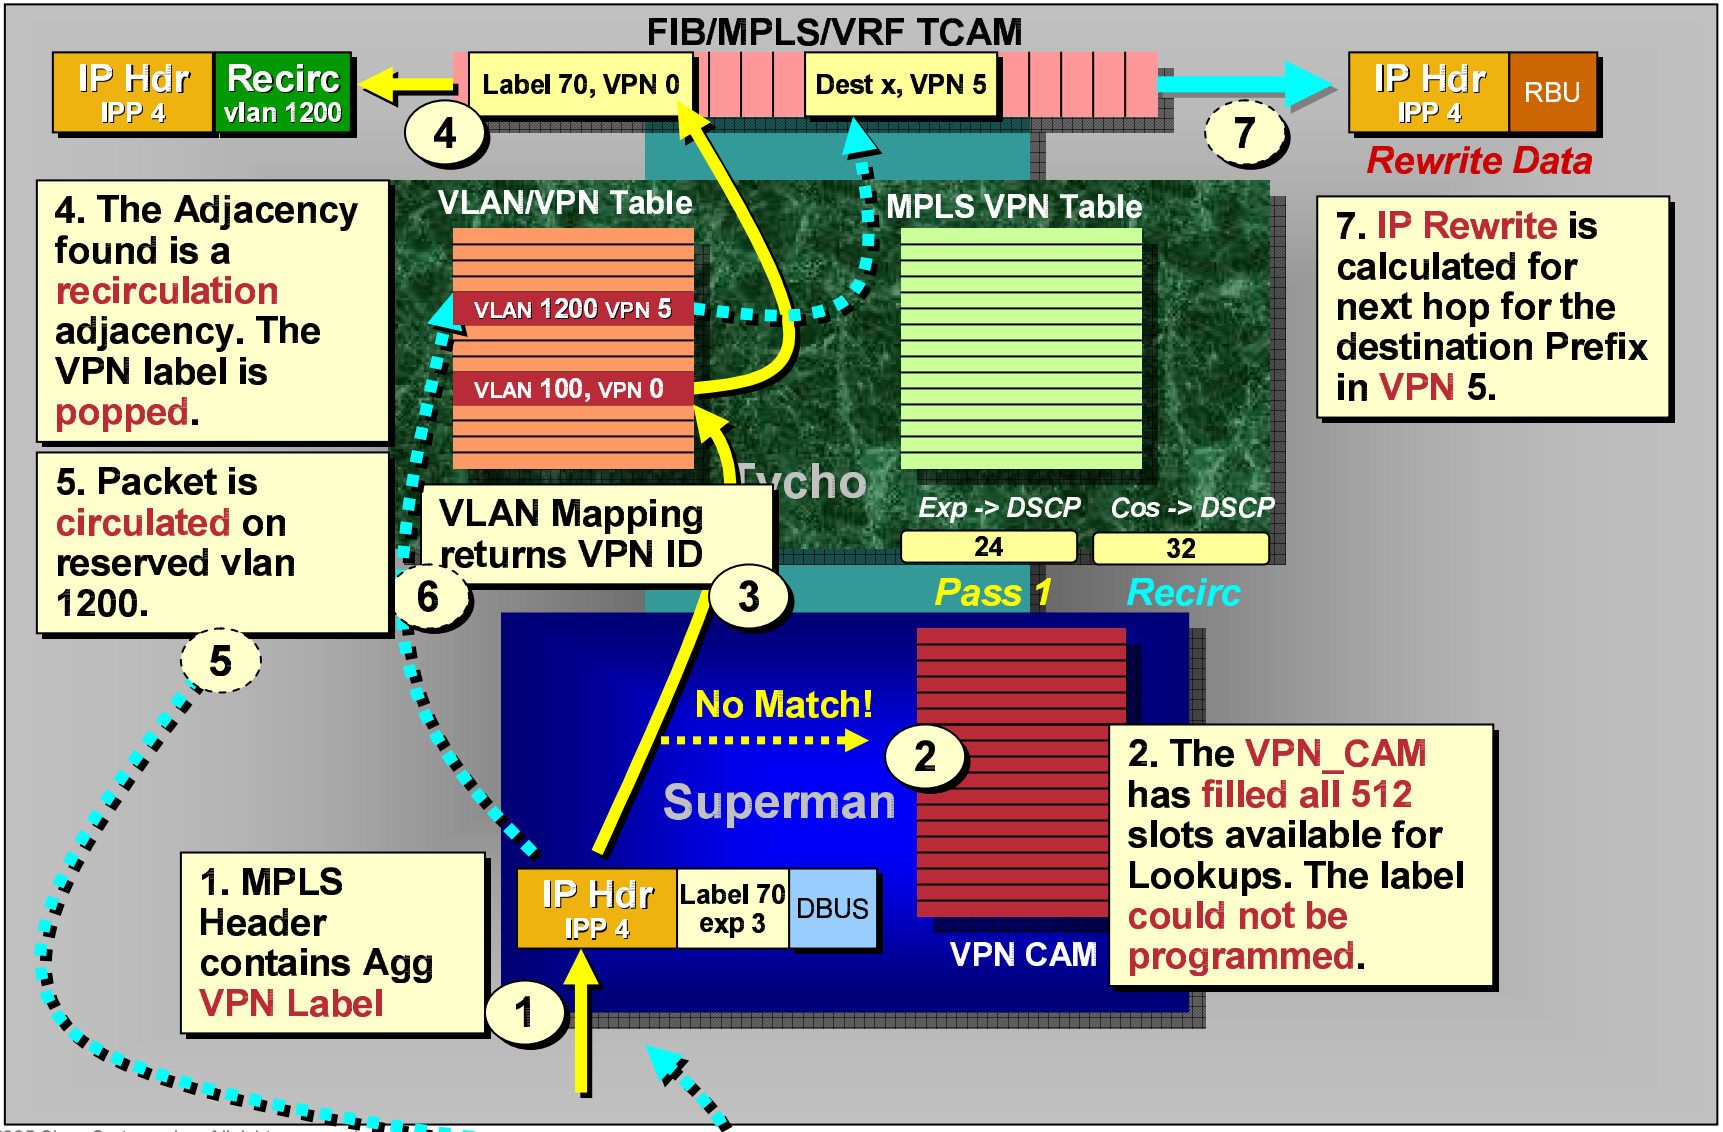

Simplifying the recirculation path (Reference)

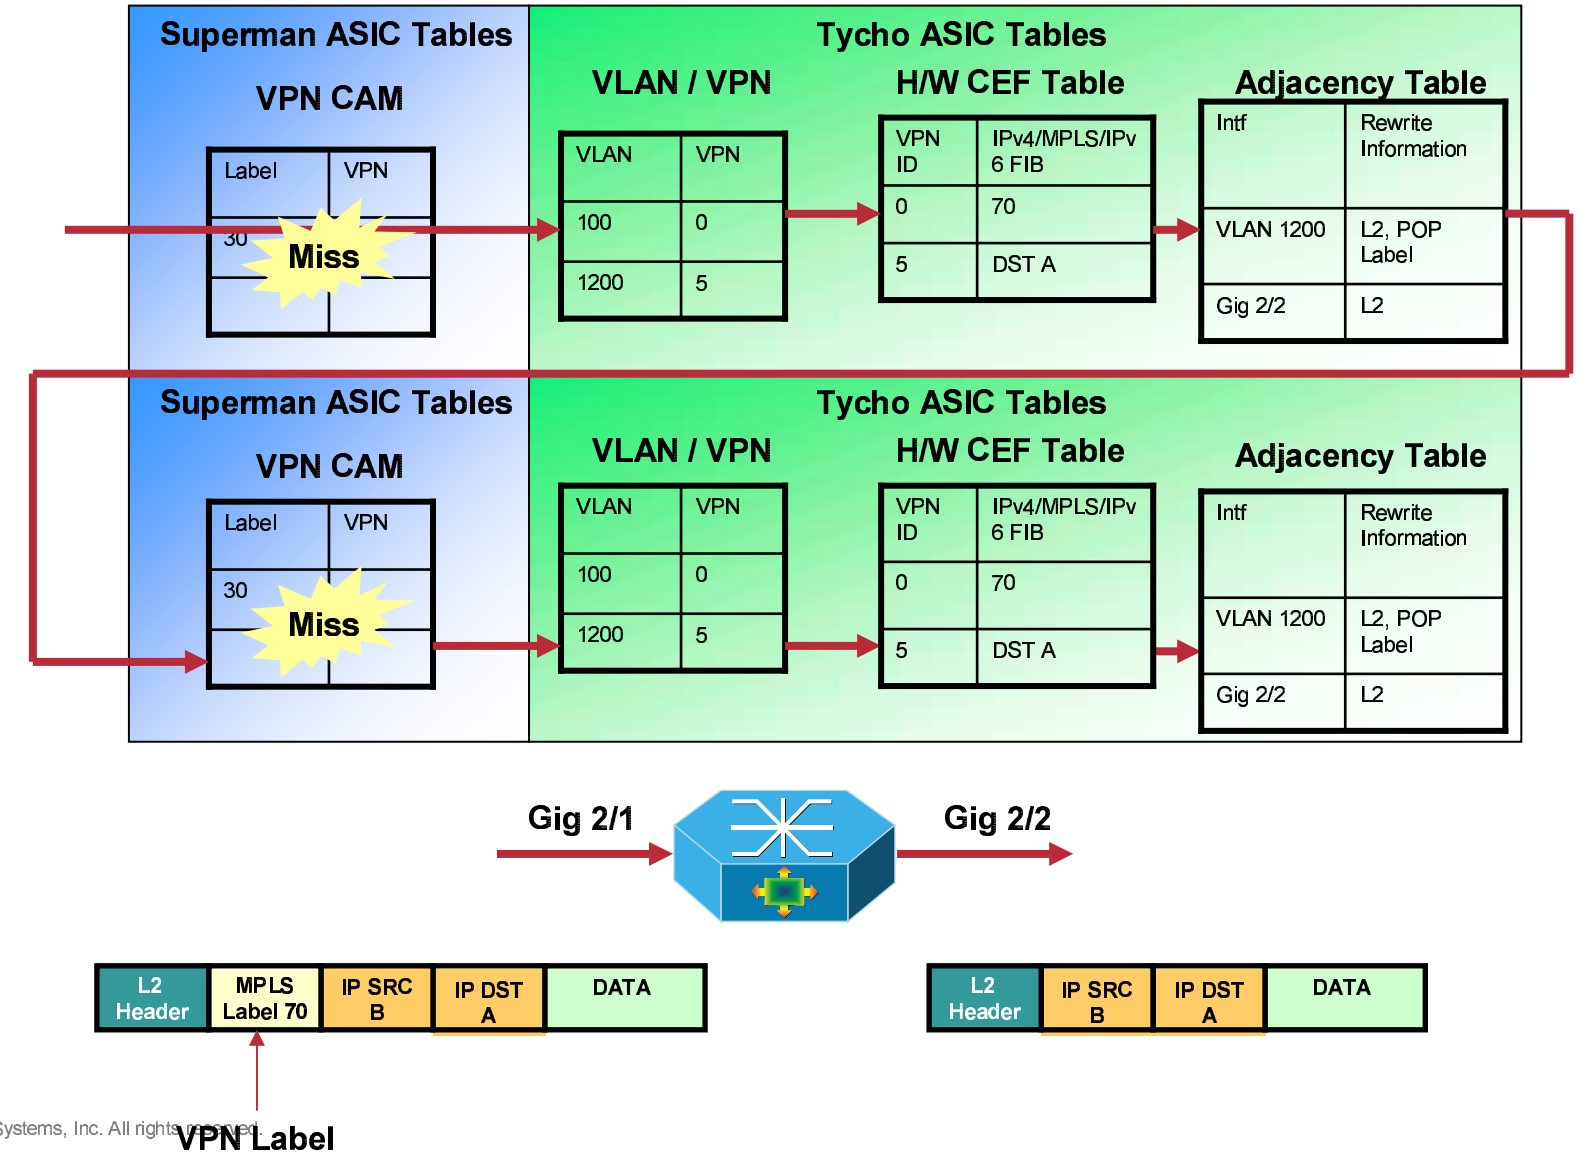

MPLS/VPN: TAG to IP Packets - Recirculation

Example: Exp NULL + VPN Label Double Recirculation

Cisco.com

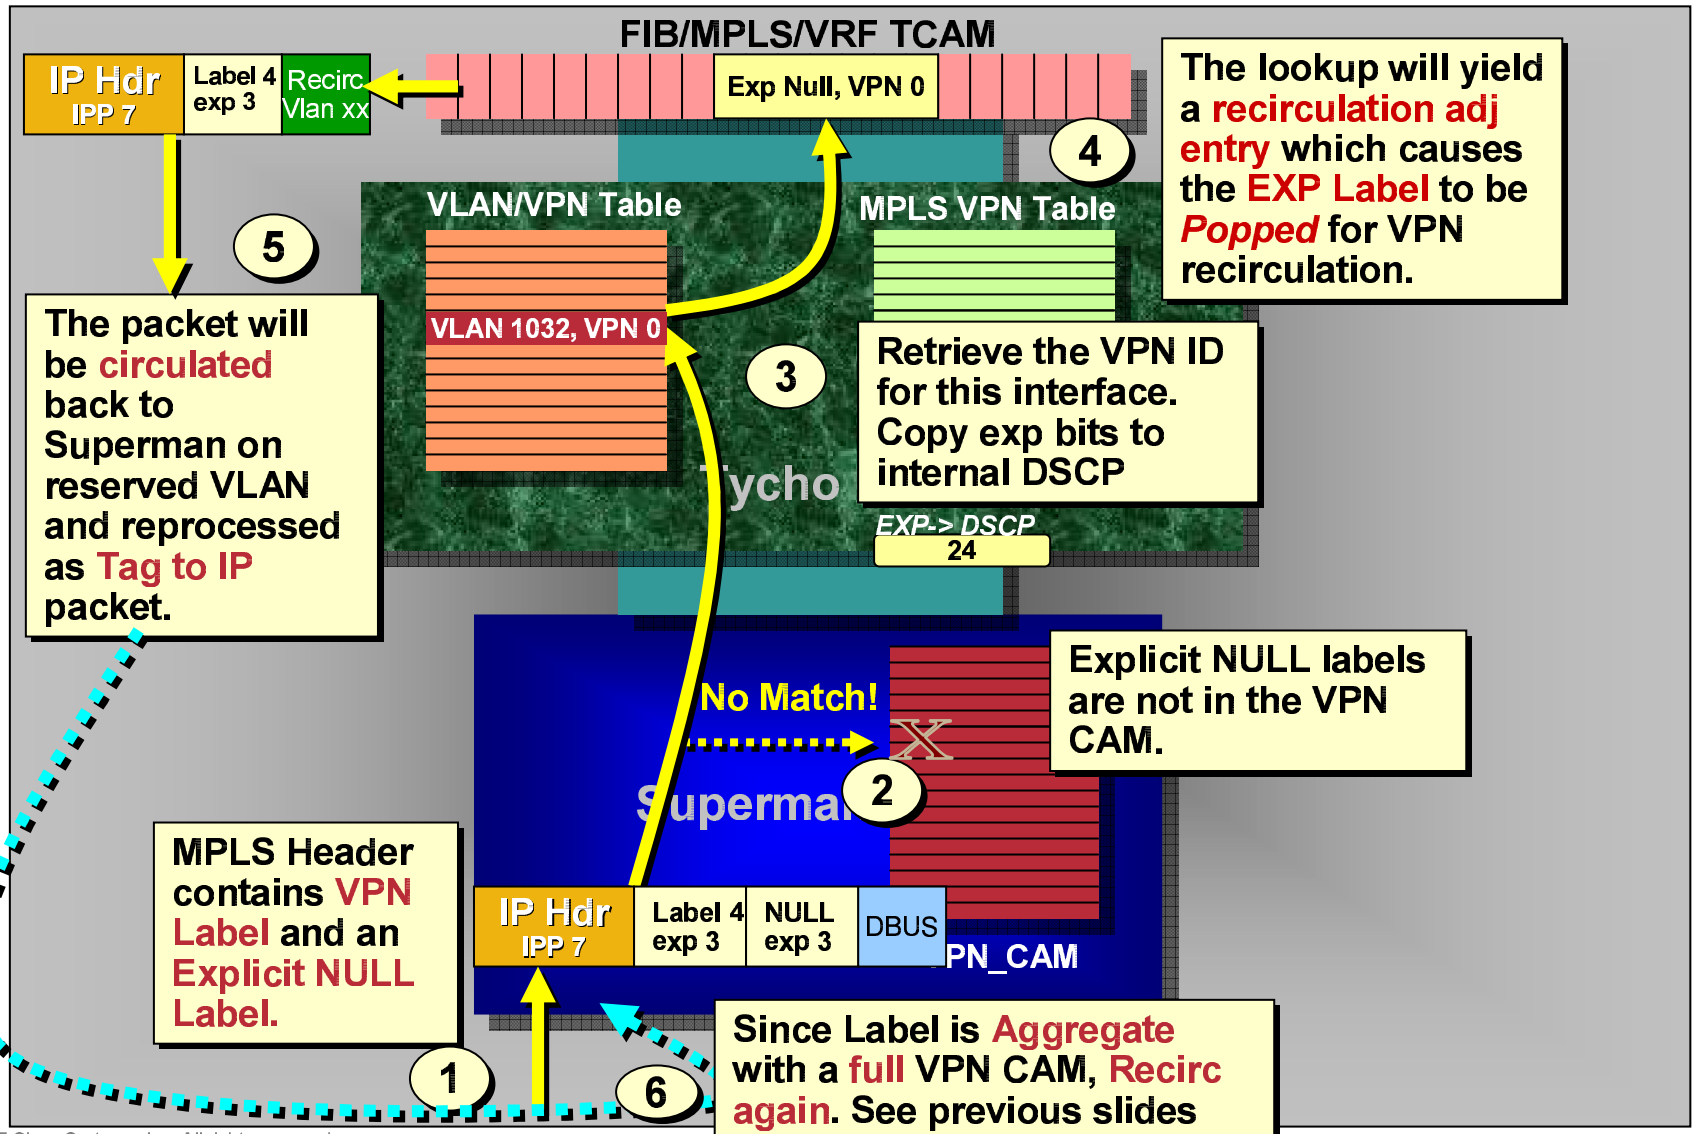

Ingress RFC2457 VPN

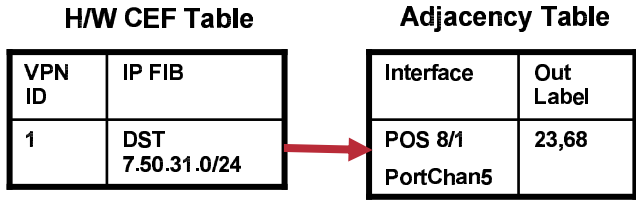

Cisco.com

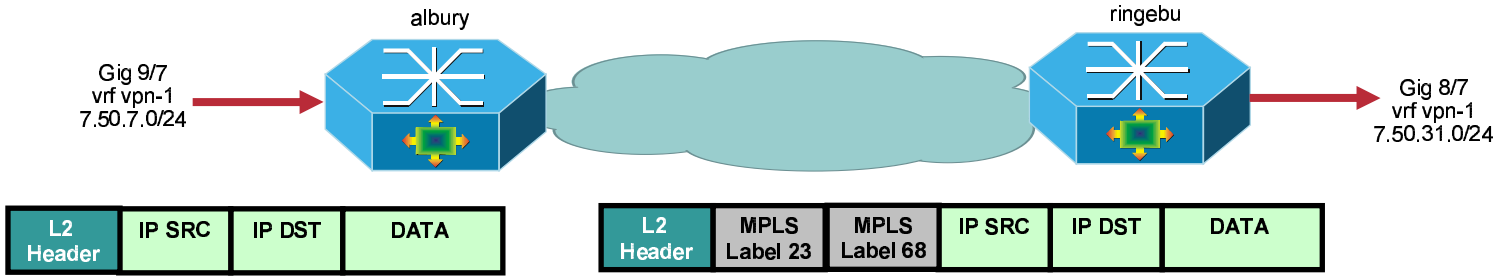


```
albury#sh ip ro vrf vpn-1 7.50.31.0
Routing entry for 7.50.31.0/24
  Known via "bgp 56001", distance 200, metric 0, type internal
  Last update from 27.0.0.53 00:04:28 ago
  Routing Descriptor Blocks:
  * 27.0.0.53 (Default-IP-Routing-Table), from 27.0.0.53, 00:04:28 ago
      Route metric is 0, traffic share count is 1
      AS Hops 0
albury#sh ip bgp vpnv4 vrf vpn-1 labels
   Network          Next Hop      In label/Out label
Route Distinguisher: 56001:1 (vpn-1)

   1.0.81.1/32      27.0.0.53       nolabel/169
   1.80.0.1/32      0.0.0.0         Per VRF Aggregate Tag:66/aggregate(vpn-1)
   7.50.15.0/24     27.0.0.83       nolabel/74
   7.50.31.0/24     27.0.0.53       nolabel/68


albury#sh mls cef vrf vpn-1 7.50.31.0

Codes: decap - Decapsulation, + - Push Label
Index  Prefix              Adjacency
87328  7.50.31.0/24        Po5             68(+),23(+) (Hash: 0001)
                           PO8/1           68(+),23(+) (Hash: 0002)
albury#sh mls cef vrf vpn-1 7.50.31.0 detail

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(87328  ): E | 1 FFF  0 0 0 0   255.255.255.0
V(87328  ): 8 | 1 256  0 0 0 0   7.50.31.0      (A:230008 ,P:1,D:0,m:1 ,B:0 )
```

```
albury#sh mpls forwarding-table 27.0.0.53
Local  Outgoing    Prefix              Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id        switched   interface
40     23          27.0.0.53/32      117005630  PO8/1       point2point
       23          27.0.0.53/32      32133      Po5         42.54.1.2

albury#sh mls cef adjacency entry 230008 detail

Index: 230008  smac: 00d0.7995.9400, dmac: 00d0.009d.9000
               mtu: 9234, vlan: 1014, dindex: 0x0, l3rw_vld: 1
               format: MPLS, flags: 0x8418
               label0: 0, exp: 0, ovr: 0
               label1: 68, exp: 0, ovr: 0
               label2: 23, exp: 0, ovr: 0
               op: PUSH_LABEL2_LABEL1
               packets: 0, bytes: 0

albury#sh mls cef adjacency entry 230009 detail

Index: 230009  smac: 00d0.7995.9400, dmac: 0000.0800.0000
               mtu: 9234, vlan: 1019, dindex: 0x0, l3rw_vld: 1
               format: MPLS, flags: 0x8418
               label0: 0, exp: 0, ovr: 0
               label1: 68, exp: 0, ovr: 0
               label2: 23, exp: 0, ovr: 0
               op: PUSH_LABEL2_LABEL1
               packets: 0, bytes: 0
```

Egress RFC2457 VPN
Non-Aggregate Label Per Prefix Label

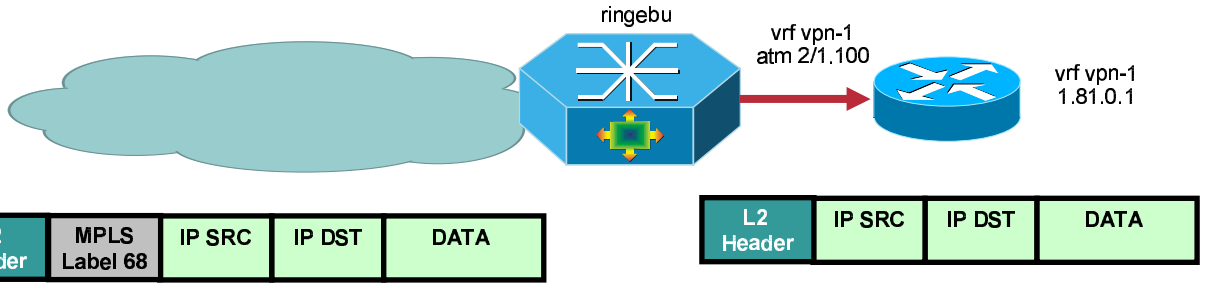


```
ringebu#sh mpls forwarding-table labels 39
Local  Outgoing    Prefix             Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id       switched   interface
39     Untagged    1.81.0.1/32[V]    1948       AT2/1.100  point2point

ringebu#sh mls cef mpls labels 39

Codes: + - Push label, - - Pop Label         * - Swap Label
Index  Local     Label                 Out i/f
       Label      Op
2001   39        (-)                   AT2/1.100     , 0000.0200.0003

ringebu#sh mls cef mpls labels 39 detail

Codes: M - mask entry, V - value entry, A - adjacency index, P - FIB Priority
       D - FIB Don't short-cut, m - mod-num, E - ELSP?
Format: MPLS - (b | xtag vpn pi cr mcast label1 exp1 eos1 valid2 label2 exp2 eos2)
V(2001   ): B | 1 0    0 0 0 39      0 0 0 0       0 0 (A:409607 ,P:0,D:0,m:0 :E:1)
M(2001   ): F | 1 FFF  0 0 1 FFFFF   0 0 0 0       0 0
ringebu#sh mls cef adjacency entry 409607 detail

Index: 409607  smac: 0005.9a3b.7240, dmac: 0000.0200.0003
               mtu: 4488, vlan: 1330, dindex: 0x0, l3rw_vld: 1
               format: MPLS, flags: 0x1000008408
               label0: 0, exp: 0, ovr: 0
               label1: 0, exp: 0, ovr: 0
               label2: 0, exp: 0, ovr: 0
               op: POP
               packets: 18, bytes: 2236
```

If outbound IP features are configured on the PE the packets will need to get re-circulated for the feature to be applied. The same mechanism is used for a miss in the VPN CAM and this is shown in an upcoming slide.

LER (PE)
VPN Egress Aggregate Label in VPN CAM

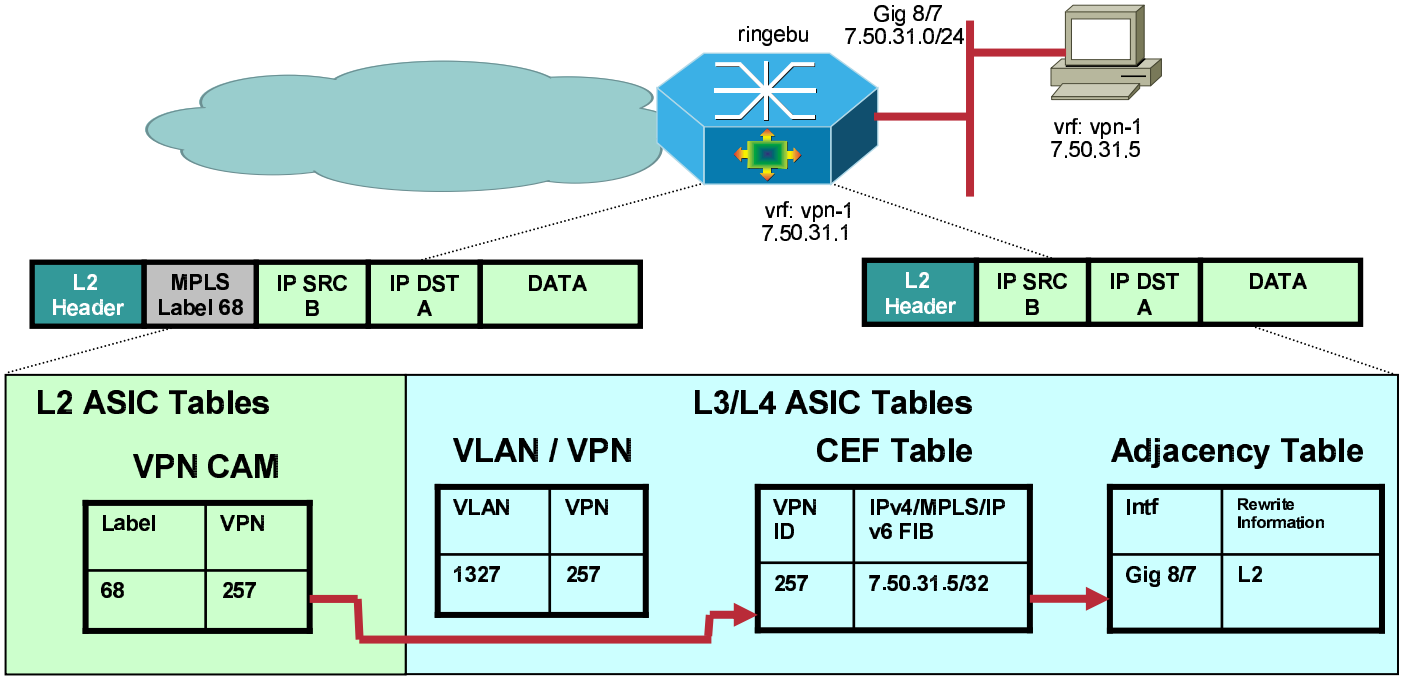

LER (PE)
VPN Egress Aggregate Label in VPN CAM (cont.)


```
ringebu#sh mpls forwarding-table labels 68
Local  Outgoing    Prefix              Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id        switched   interface
68     Aggregate   vrf:vpn-1           4288

ringebu#sh mls cef mpls labels 68

Codes: + - Push label, - - Pop Label         * - Swap Label
Index  Local     Label                  Out i/f
       Label      Op

ringebu# sh ip ro vrf vpn-1 7.50.31.5
Routing entry for 7.50.31.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via bgp 56001
  Advertised by bgp 56001
  Routing Descriptor Blocks:
  * directly connected, via GigabitEthernet8/7
      Route metric is 0, traffic share count is 1

ringebu#sh ip cef vrf vpn-1 7.50.31.5
7.50.31.5/32, version 37, epoch 0, connected, cached adjacency 7.50.31.5
0 packets, 0 bytes
  via 7.50.31.5, GigabitEthernet8/7, 0 dependencies
    next hop 7.50.31.5, GigabitEthernet8/7
    valid cached adjacency

ringebu#sh mls cef vrf vpn-1 7.50.31.5

Codes: decap - Decapsulation, + - Push Label
Index  Prefix              Adjacency
2338   7.50.31.5/32        Gi8/7          , 0000.0732.1f05

ringebu#sh mls cef vrf vpn-1 7.50.31.5 detail

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(2338   ): E | 1 FFF 0 0 0 0   255.255.255.255
V(2338   ): 8 | 1 257 0 0 0 0   7.50.31.5          (A:229421 ,P:1,D:0,m:0 ,B:0
```

```
ringebu#sh mpls platform vpn-vlan-mapping | inc VPN#|257
VPN#  Rsvd Vlan  IDB Created  Feature  Has agg label  In superman    EoM data
1          No          No       Yes           Yes            Yes

ringebu-sp#sh platform software vpn mapping
   Type    |   VRF Name           | Table id  || HW table id   |
-----------+----------------------+-----------++--------------+
 IOS       | Default              | 0         || 0            |
 IOS       | vpn-1                | 1         || 257          |
```

Aggregate label does not get installed into the L3/L4 ASIC FIB when it is installed into the L2 ASIC VPN CAM. The L2 ASIC removes the Aggregate label and passes just the IPv4 packet to the L3/L4 ASIC. The L3/L4 ASIC handles the packet as an IP packet. The L3/L4 ASIC performs all L3/L4 features like ACLs and Policing.

Use 'sh mls vpn-cam 0 511' to look up the label-to-vpn mapping 68 → 257 in this case

Check the mls adj entry, like earlier

LER (PE) VPN Egress Aggregate Label not in VPN CAM, IP Egress Feature on Per Prefix Label

Cisco.com

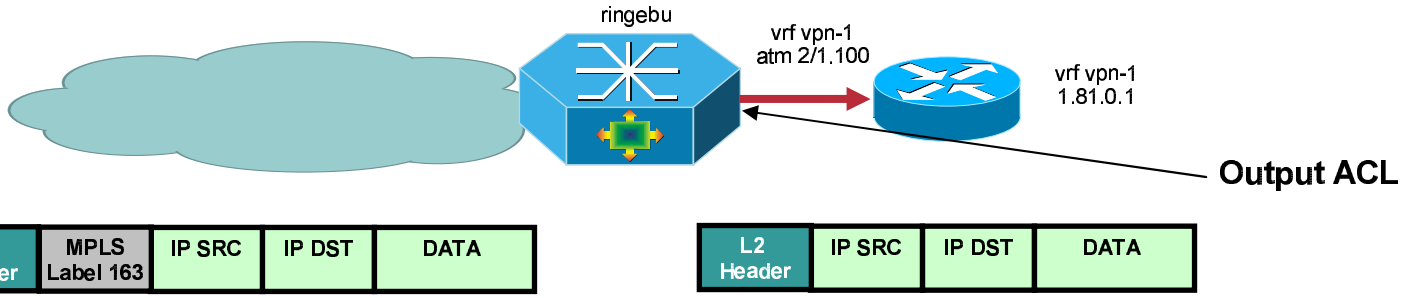

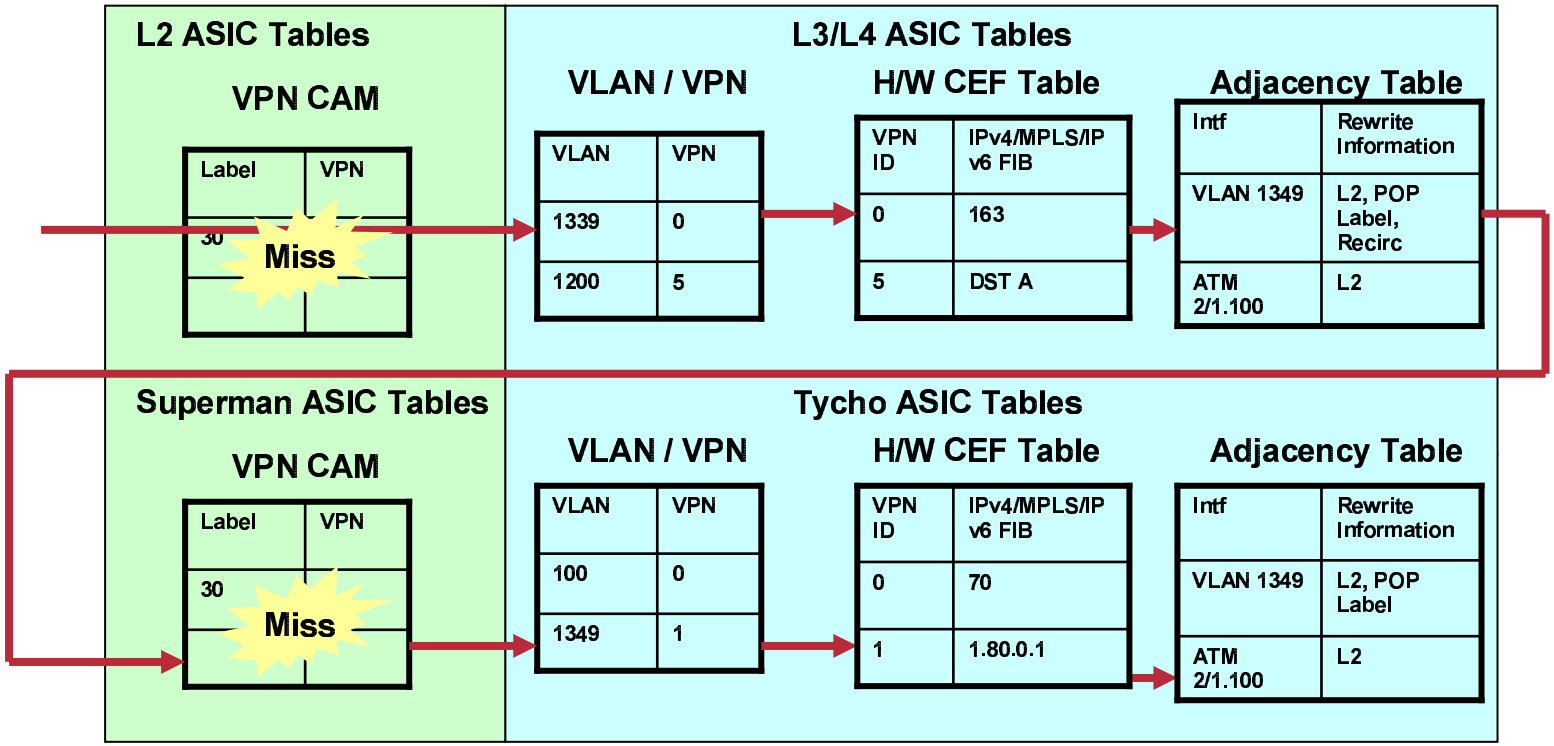

LER (PE)
VPN Egress Aggregate Label not in VPN CAM


```
ringebu#sh mpls forwarding-table labels 163 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
163    Untagged    1.81.0.1/32[V]    540        AT2/1.100  point2point
        MAC/Encaps=0/0, MRU=4474, Tag Stack{}
        VPN route: vpn-1
        No output feature configured
    Per-packet load-sharing

ringebu#sh mls cef mpls labels 163
Codes: + - Push label, - - Pop Label         * - Swap Label
Index  Local     Label                 Out i/f
       Label      Op
133558 163       (-)                   recirc

rngebu#sh mls cef mpls labels 163 detail
Codes: M - mask entry, V - value entry, A - adjacency index, P - FIB Priority
       D - FIB Don't short-cut, m - mod-num, E - ELSP?
Format: MPLS - (b | xtag vpn pi cr mcast label1 exp1 eos1 valid2 label2 exp2 eos2)
V(133558 ): B | 1 0    0 0 0 163     0 0 0 0       0 0 (A:409606 ,P:0,D:0,m:0 :E:1)
M(133558 ): F | 1 FFF  0 0 1 FFFFF   0 0 0 0       0 0

ringebu#sh mls cef vrf vpn-1 1.81.0.1
Codes: decap - Decapsulation, + - Push Label
Index  Prefix             Adjacency
2107   1.81.0.1/32         AT2/1.100       , 0000.0200.0003

ringebu#sh mls cef vrf vpn-1 1.81.0.1 detail
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(2107   ): E | 1 FFF  0 0 0 0   255.255.255.255
V(2107   ): 8 | 1 257  0 0 0 0   1.81.0.1          (A:98304 ,P:1,D:0,m:0 ,B:0 )
```

```
ringebu#sh mls cef adjacency entry 409606 detail

Index: 409606  smac: 0000.0000.0000, dmac: 0005.9a3b.7240
               mtu: 65535, vlan: 1349, dindex: 0x7FFA, l3rw_vld: 1
               format: MPLS, flags: 0x1000008600
               label0: 0, exp: 0, ovr: 0
               label1: 0, exp: 0, ovr: 0
               label2: 0, exp: 0, ovr: 0
               op: POP
               packets: 3, bytes: 582

ringebu#sh platf soft vlan mapp | inc 1349
IOS      | vpn-1

ringebu#sh mls cef adjacency entry 98304 detail

Index: 98304   smac: 0005.9a3b.7240, dmac: 0000.0200.0003
               mtu: 4488, vlan: 1204, dindex: 0x0, l3rw_vld: 1
               format: MAC_TCP, flags: 0x2000208408
               delta_seq: 0, delta_ack: 0
               packets: 0, bytes: 0

ringebu#sh mpls platform reserved-vlans vlan 1349 | inc protocol|L3
VRF_1_vlan1349 is up, line protocol is up
  L3 in Switched: ucast: 3 pkt, 570 bytes - mcast: 0 pkt, ...
  L3 out Switched: ucast: 3 pkt, 582 bytes mcast: 0 pkt, ..

ringebu-sp#sh platform software vpn mapping
  Type   |   VRF Name          | Table id  || HW table id   |
---------+---------------------+-----------++---------------+
IOS      | Default             | 0         || 0            |
IOS      | vpn-1               | 1         | 257            |
```

Köszönöm a figyelmet!