Catalyst 6800 Series Switch Architectures

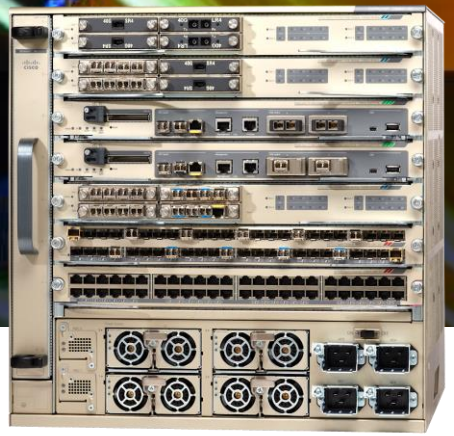

Blake Gao, TAC Engineer

March, 2016

Session Objectives

Session Objectives

What we will cover…

Provide you with a *working understanding* of the new Catalyst C6807-XL and C6880-X Campus Backbone Switches

- Chassis Architectures
- Forwarding Engine Functions
- Hardware Feature Operations
- Basic Packet Flows

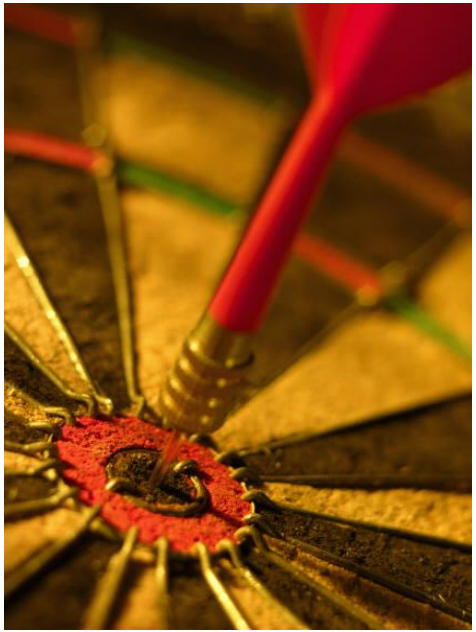

Catalyst 6800 Series

3000+ Catalyst 6500 Features

MPLS, VPLS & EVN

10 + Year MPLS Maturity	✓
L2 VPN	✓
L3 VPN	✓
L2/L3 VPN over mGRE	✓
MPLS TE	✓
VPLS/A-VPLS/H-VPLS	✓
Label Switch Multicast	✓
MACsec over EoMPLS	✓

Next-Gen Solutions

FEX Controller	✓
L2 Campus LISP	✓
L3 Campus LISP	✓
SDN / ONEpk	✓
Smart Call Director	✓
1G/10G/40G VSL	✓
Multicast BGP PIC	✓

VSS & HA

8+ Year VSS Maturity	✓
MACsec on VSL	✓
BFD on SVI & MEC	✓

QoS & Security

MACsec & SGT	✓	DHCP Snooping	✓
SGACL & L3 SGT	✓	Dynamic ARP Inspection	✓
Ingress/Egress ACL	✓	Span with ACL	✓
Time-Based ACL	✓	Identity 2.0	✓
ACL Statistics	✓	ACL-Based QoS Classification	✓
Port Security	✓	Marking (DSCP/COS)	✓

IPV4 Unicast

OSPFv2/v3	✓
OSPFv3 VRF-Lite	✓
VRF-Aware Unicast	✓
IP Tunnel HA	✓
BFD SVI-GRE	✓
IPv6–IPv4 HW Parity	✓
BGP PIC	✓
OSPFv3 VRF PE-CE	✓

IPV4 Multicast

IGMPv3 and MLDv2 Snooping in HW	✓
PIM-SM "Dual-RPF" in HW	✓
PIM-DM, PIM-Bidir	✓
PIM-Bidir, PIM Snooping	✓
PIM Register in HW	✓
mVPN, MSR, mcast BFD	✓
BGPv6, IS-ISv6	✓

IPV6 Features

MPLS at Access with I-Sat	✓
VRF-Lite EVN	✓
VRF-Aware Multicast	✓
VRF-Aware Net-flow	✓
IPv6 in IPv4 Tunnels, VRF-Aware IPv6 Tunnels	✓
MLD, MLD Snooping	✓
6VPE and 6PE	✓

Large Tables & Scalability

IPv4 Routing Capability	**256K-2M**
Multicast Routes (IPv4)	64K-128K
Number of Adjacencies	1M
MAC Addresses	128K
ECMP (v4 and v6)	16
Security & QoS ACL	64K-256K
Flexible Netflow	128K-5M
MPLS Label Push/Pop in 1pass	5/3

Rich Media

IPv6 FHS Security	✓
Flexible NetFlow	✓
Egress NetFow	✓
ACL Atomic Commit/Dry Run	✓
Sampled NetFlow	✓
IPv6 uRPF	✓
NDE (Full & Sampled)	✓
Video Monitoring	✓
Mediatrace	✓
Metadata QoS	✓
Multicast Service Reflection (MSR)	✓

Management & Services

Microflow Policing	✓
WCCPv3	✓
PBR IPv4/IPv6	✓
Advanced CoPP	✓
NAT/PAT	✓
IPv6 VACL	✓
GRE	✓
ERSPAN	✓
GOLD	✓
ISE	✓
Cisco Prime	✓
Mini Protocol Analyzer	✓

Agenda

- **Chassis & Power**
  - **C6807-XL**
  - C6880-X
- Supervisor Architectures
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Catalyst 6807-XL Chassis

High-Level Overview

Up to 880G/Slot capable

7 Slots
10 RU

Catalyst 6000 DNA

Next-Generation
ASIC Ready

Investment Protection:
Compatible with Sup2T, 6700, 6800, 6900 and Latest Service Modules

Low-Power & Noise
High-Efficiency Fans

Backwards compatible backplane connectors

Up to 4 (N+1)
3000W PSU
redundancy

4 x 220VAC
Power Inputs

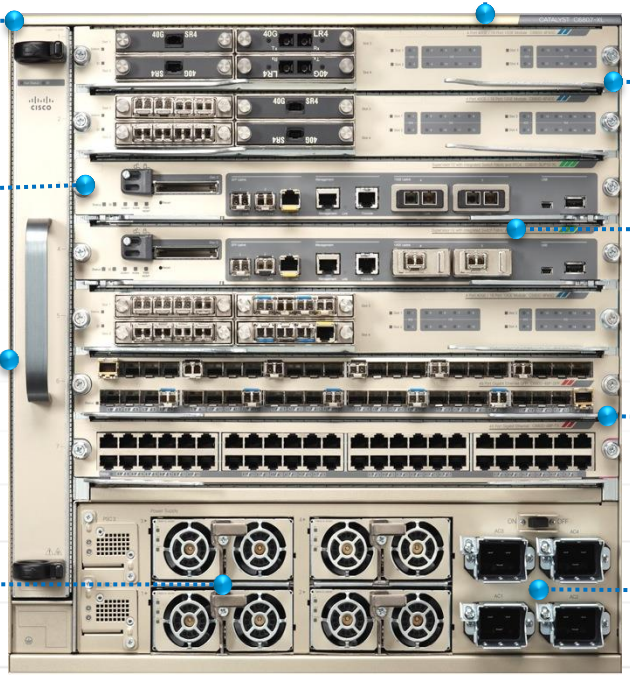

C6807-XL

Catalyst 6807-XL:

Environmental Overview

High Efficiency
4500 RPM
Redundant Fans

Platinum Efficient
3000W AC
Power Supplies

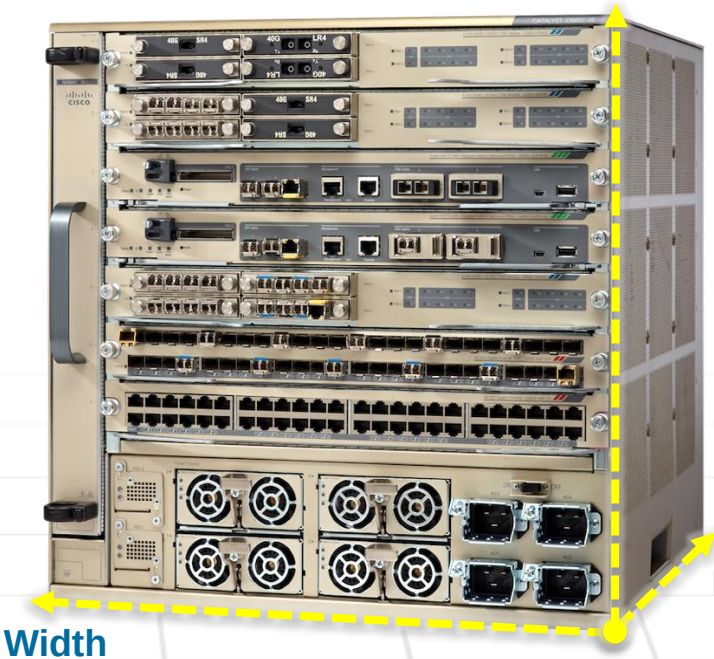

Height
17.5” (10RU)

Depth
18.10”

Width
17.36”

Catalyst 6807-XL:

Mechanical View

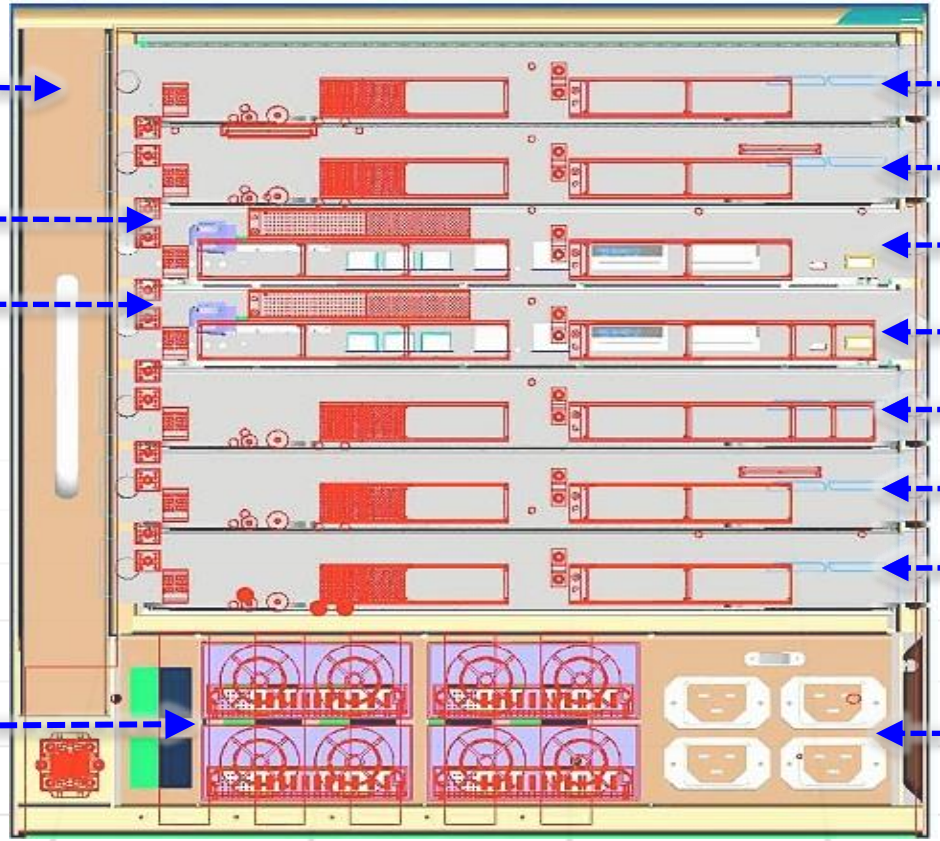

Fan Tray

Extra Fabric Connectors

Power Supplies

Slot 1: Line Card

Slot 2: Line Card

Slot 3: Supervisor *

Slot 4: Supervisor *

Slot 5: Line Card

Slot 6: Line Card

Slot 7: Line Card

Power Inputs

Catalyst 6807-XL

C6807-XL Supports:

- **Supervisor 2T**
- **Current Fabric Cards**
  - 6900, 6800 & 6700 (CFC or DFC4)*
- **Current Service Modules**
  - NAM-3, ASA-SM, WISM2, ACE-30
- **4 x 3000W AC Power Supplies**
- **8 x Fabric Channels Per Slot**
  - **4 Channels to each Supervisor**
  - Up to 220G with Sup2T in Active / Standby
- **Future Supervisors & Cards**
  - **Each Channel can operate @ 110Gbps!**

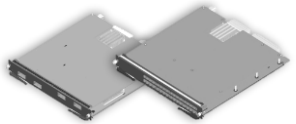

NO support for Sup720!

Supervisor Engine
VS-S2T-10G
VS-S2T-10G-XL
Line Cards
WS-X6904-40G-2T *
WS-X6908-10G-2T *
WS-X6824-SFP-2T *
WS-X6848-SFP-2T *
WS-X6848-TX-2T *
WS-X6816-10T-2T *

Catalyst 6807-XL

Fan Redundancy & Air Flow

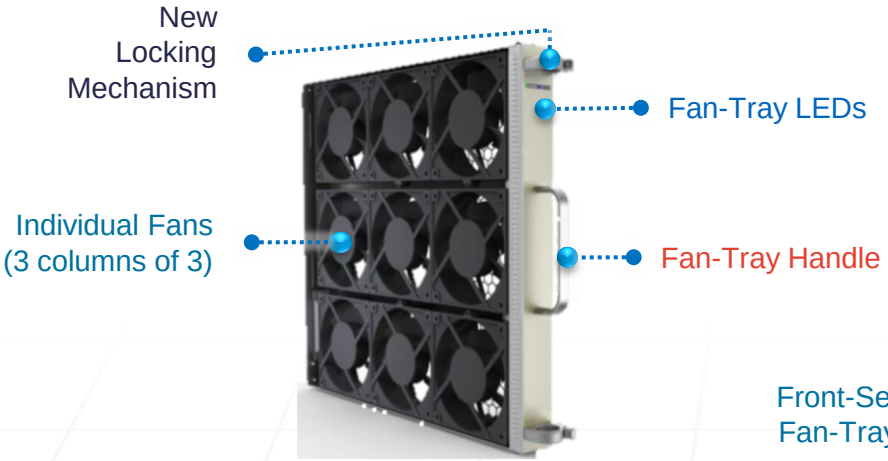

Air Flow:
Side to Side

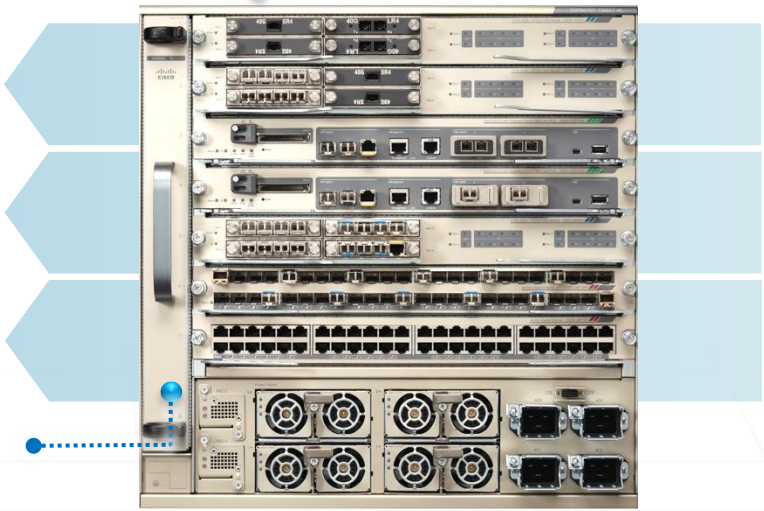

Fan-Tray Highlights:

- Has 9 variable-speed High-Efficiency Fans (850 CFM)
- Supports 4 speeds between 3000 & 4500 RPM per Fan
- Capable of cooling Slots operating up to 800W per Slot
- Can still operate with up to 3 individual fan failures
- Supports Fan-Tray “OIR” for minimum of 120 seconds

LED	Color	Status	Description
FAN		Solid	Fan-Tray OK
FAN		Solid	Fan-Tray Fault
ID		Solid	Identifies Fan-Tray

Catalyst 6807-XL

PSU Redundancy & Inputs

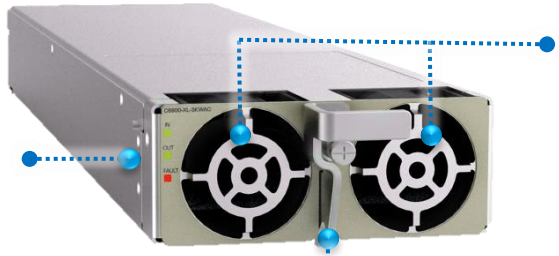

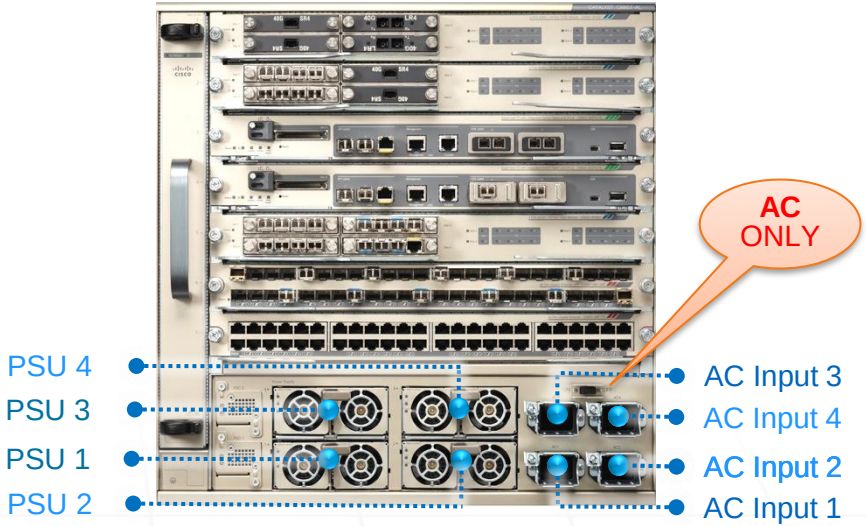

Power Supply Highlights:

- Hold-up time is ~ 20 msecs at 100% load
- Up to 92% Power Efficiency at 50-100% of load
- Max output is 3000W @ 220V (or 1300W @ 110V)
- Dual “Front to Back” Variable-Speed Cooling Fans
- Supports both Combined & Redundant (N+1) mode

LED	Color	Status	Description
IN		Solid	Input OK
IN		Blinking	Under-Current
OUT		Solid	Output OK
OUT		Blinking	Over-Current

Power Supply Redundancy

Catalyst 6807-XL

For Your Reference

Catalyst 6807-XL Can Utilize **Four** Power Supplies in Either **Redundant** or **Combined** Mode

N+1 Redundant Mode

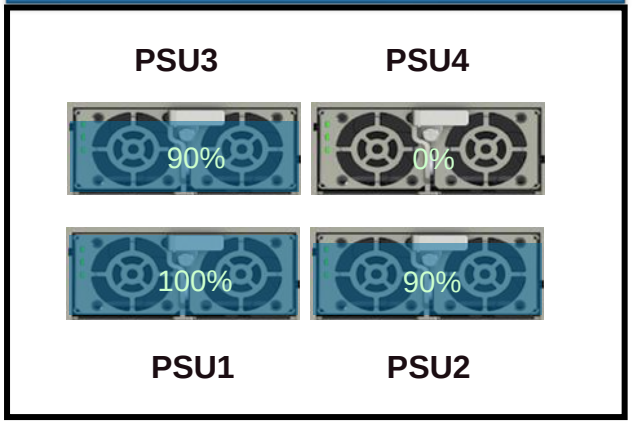

- **Adds +1 to total # of Redundant PSU**
- **First PSU operates @ 100% of capacity**
- Each Additional PSU @ 90% (100+90*N), with the +1 Redundant @ 0%
- With 1+1, 2+1 & 3+1 redundancy, when one PSU fails, the +1 PSU will take over
- **This is the default & recommended mode**

Combined Mode

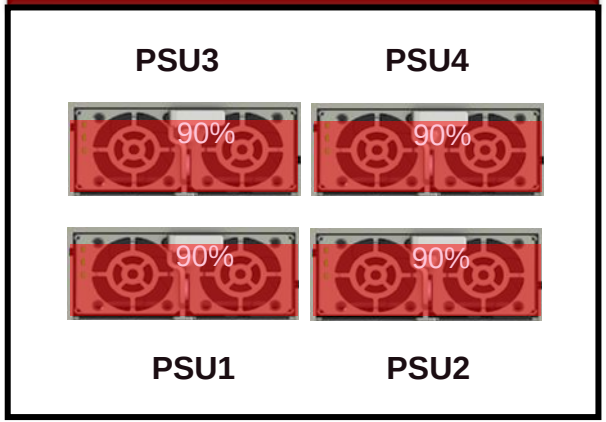

- **Same operational behavior as 6500-E**
- **Each PSU provides ~90% of capacity**
- The total system power is ~366% of the capacity of a single PSU
- Pseudo-redundant behavior, but this is not equivalent to 1:1 or N+1 redundancy.
- **This is not a recommended mode**

Agenda

- **Chassis & Power**
  - C6807-XL
  - **C6880-X**
- Supervisor Architectures
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Catalyst 6880-X Chassis

High-Level Overview

Catalyst 6000
Feature Set

(e.g. MPLS/VPLS, LISP, Multicast, IPv6, CTS, VSS, FEX)

5 RU
(4 Half Slots)

Up to 220G/Slot
Backplane

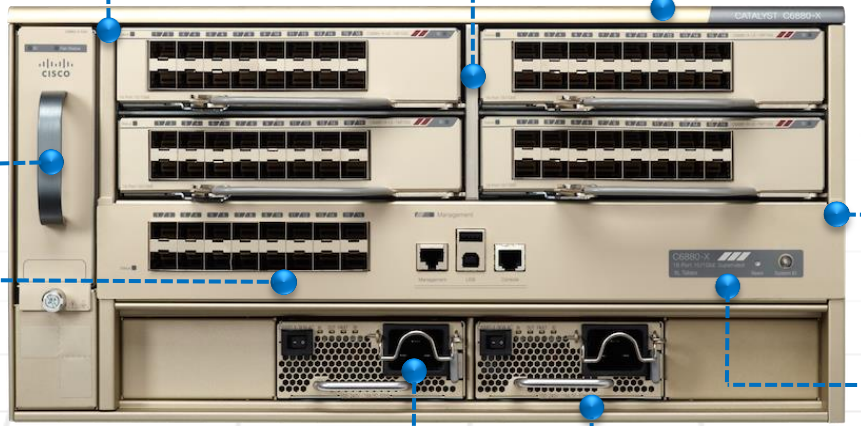

Low-Power & Noise
High-Efficiency Fans

Fixed RP Slot with
Built-in 16P 1/10G

Powerful X86
2.0GHz CPU
& 4GB DRAM

Up to 2M FIB
Up to 5M Netflow
Up to 256K ACL

3000W (AC/DC)
Power Supplies

1+1 Redundant PS
(Front Serviceable)

C6880-X

Catalyst 6880-X

Environmental Overview

High Efficiency
4500 RPM
Redundant Fans

Platinum Efficient
3000W AC
Power Supplies

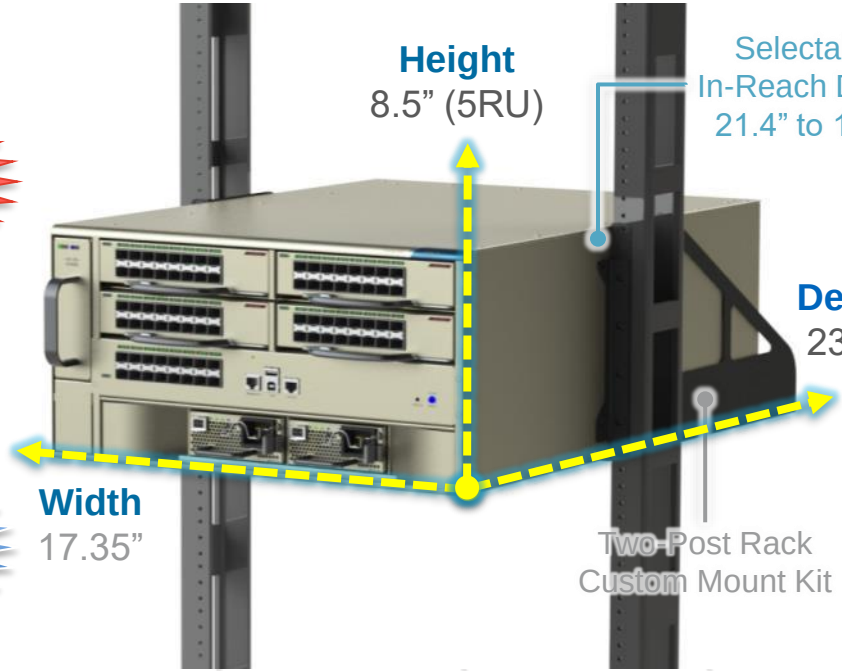

Flexible Mounting Brackets

Catalyst 6880-X

Mechanical View

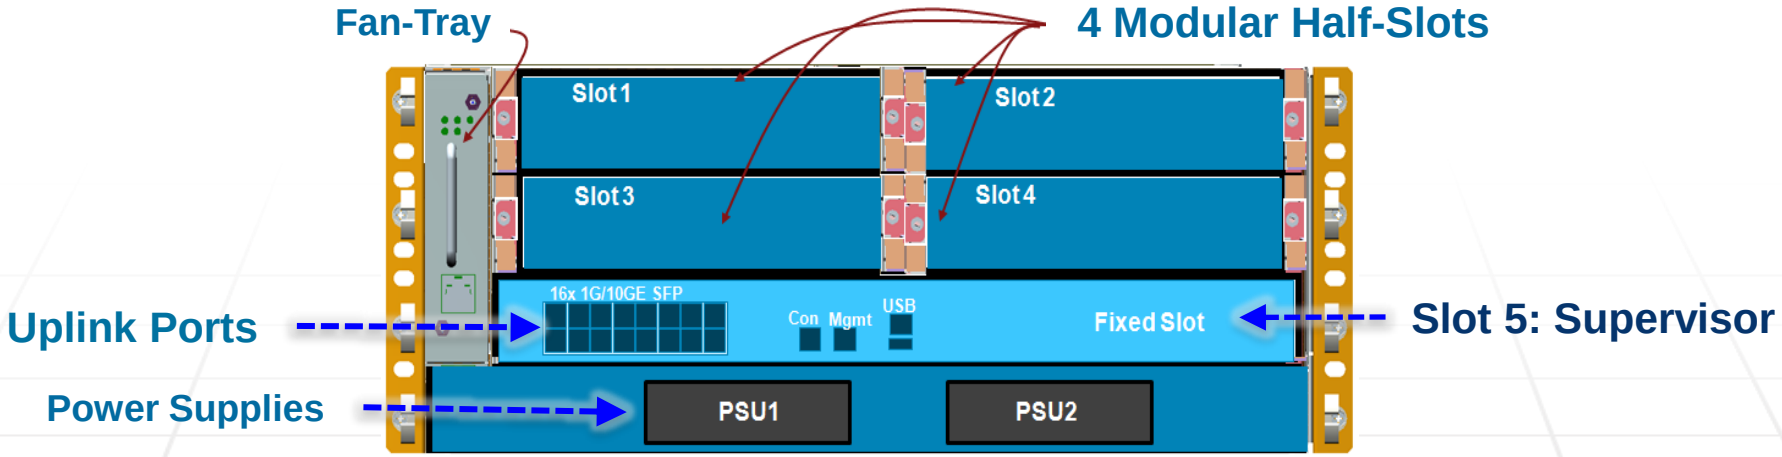

Catalyst 6880-X

Fan Redundancy & Air Flow

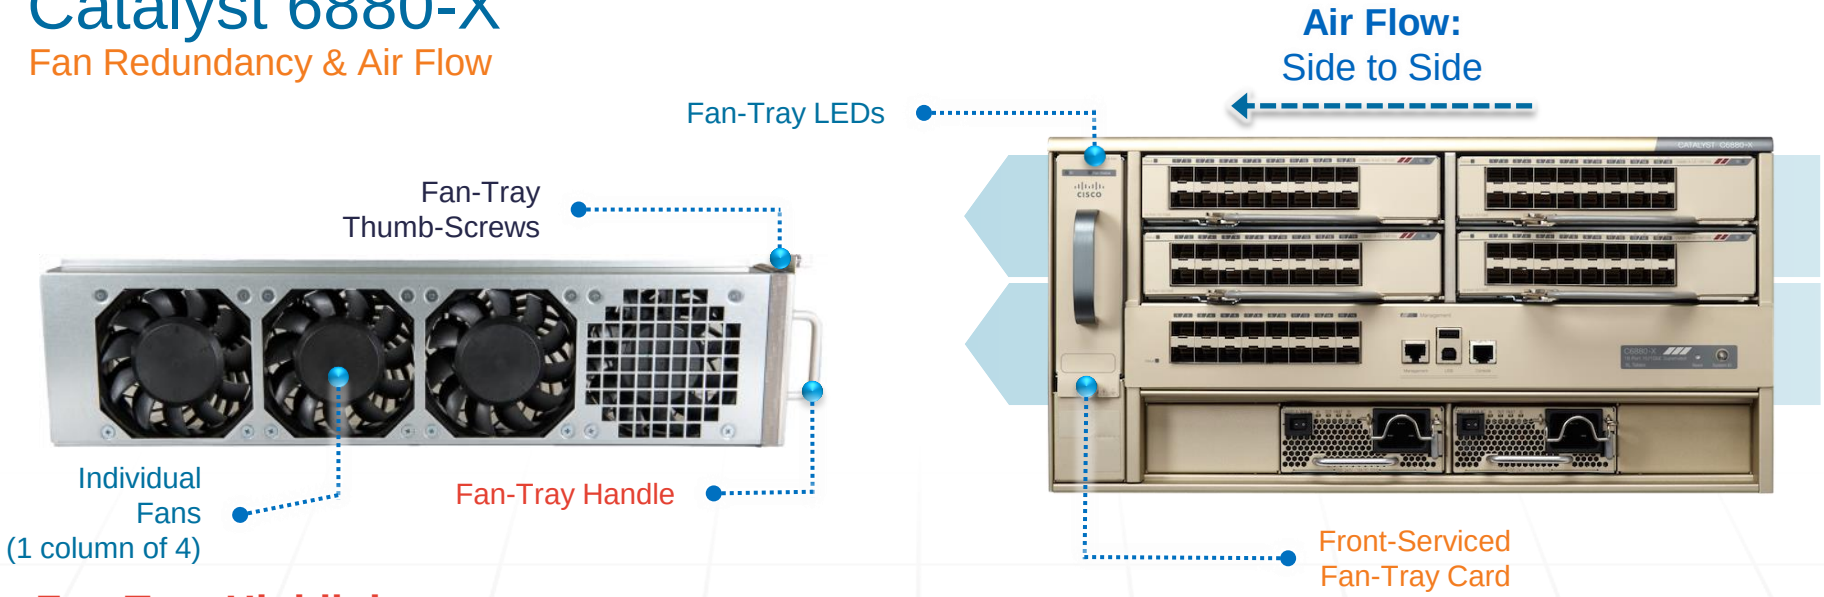

Fan-Tray Highlights:

- Has 4 variable-speed High-Efficiency Fans (250 CFM)
- Supports 4 speeds between 3000 & 4500 RPM per Fan
- Capable of cooling Slots operating up to 800W per Slot
- Can still operate with up to 2 individual fan failures
- Supports Fan-Tray “OIR” for minimum of 120 seconds

LED	Color	Status	Description
FAN		Solid	Fan-Tray OK
FAN		Solid	Fan-Tray Fault
ID		Solid	Identifies Fan-Tray

Catalyst 6880-X

PSU Redundancy & Inputs

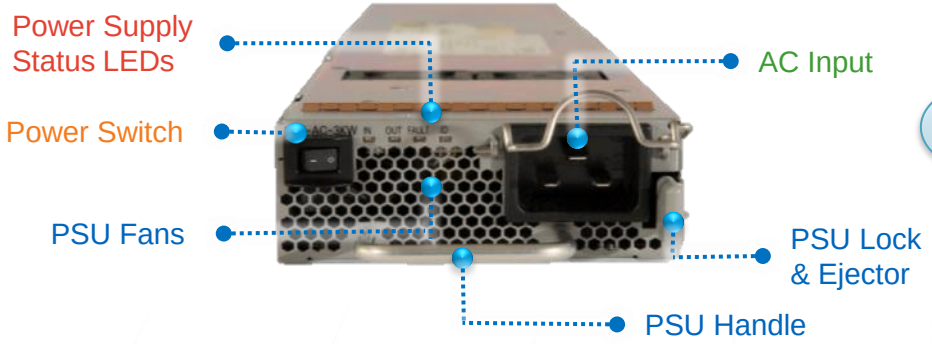

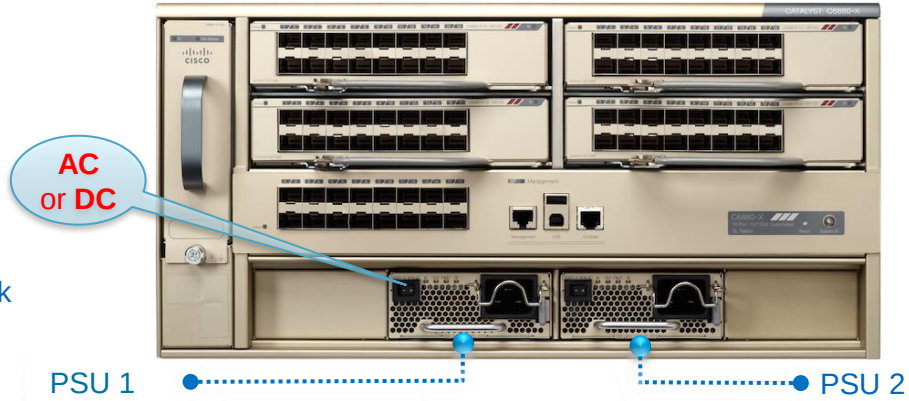

Power Supply Highlights:

- Hold-up time is ~ 20 msecs at 100% load
- Up to 92% Power Efficiency at 100% of load
- Max output is 3000W @ 220V (or 1300W @ 110V)
- Dual “Front to Back” Variable-Speed Cooling Fans
- Supports both Combined & Redundant (1:1) mode

LED	Color	Status	Description
IN		Solid	Input OK
IN		Blinking	Under-Current
OUT		Solid	Output OK
OUT		Blinking	Over-Current

Power Supply Redundancy

Catalyst 6880-X

For Your Reference

The **Catalyst 6880-X** Utilizes **Two** Power Supplies in Either **Redundant** or **Combined** Mode

Redundant Mode

PSU 1 PSU 2

- Each supply provides ~50% of power needs
- Neither supply operates at >60% or <40% capacity
- Either supply can power the system on its own
- This is the **BEST PRACTICE** mode of operation

Combined Mode

PSU 1 PSU 2

- Each supply provides up to 83% of its capacity
- Total power available is 167% o a single supply
- A single supply may not power the whole system
- **NOT** the recommended mode for production

Agenda

- Chassis & Power
- **Supervisor Architectures**
  - **VS-S2T-10G**
  - MSFC, PFC & Fabric
  - C6880-X Baseboard
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Supervisor Engine 2T
Quick Facts

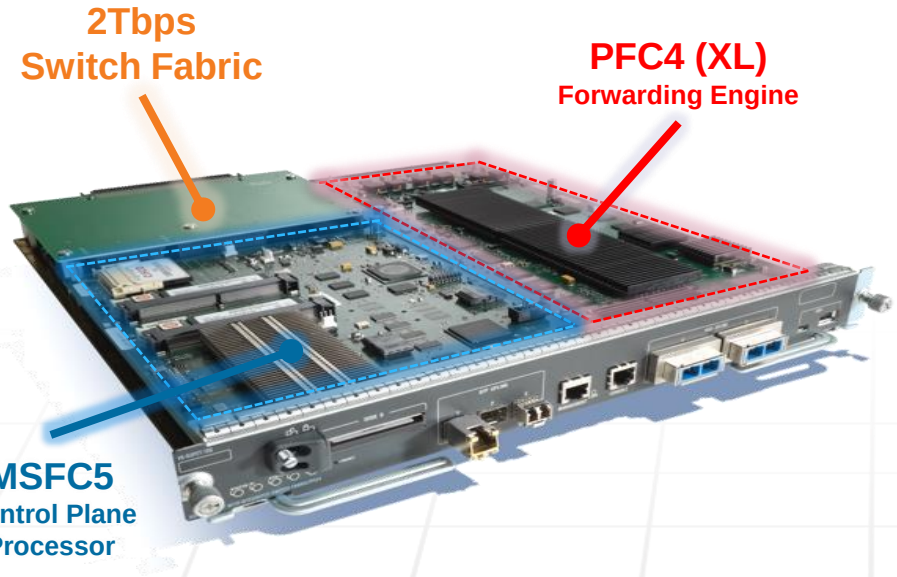

VS-S2T-10G

PFC4 (4XL)

QUICK FACTS

- ✓ **Integrated 2Tbps Switch Fabric**
- ✓ **Integrated Policy Feature Card 4 (PFC4) supporting L2/L3+ hardware acceleration**
- ✓ **Integrated Multilayer Switch Feature Card 5 (MSFC5) with a single CPU for both Layer 2 & Layer 3 functionality**
- ✓ Two 10GE & Three 1GE Uplink Ports
- ✓ Connectivity Management Processor (CMP) for improved switch management
- ✓ IPv4, IPv6, Multicast, MPLS / VPLS, VSS & Instant Access (FEX) support

Supervisor 2T

Block Diagram

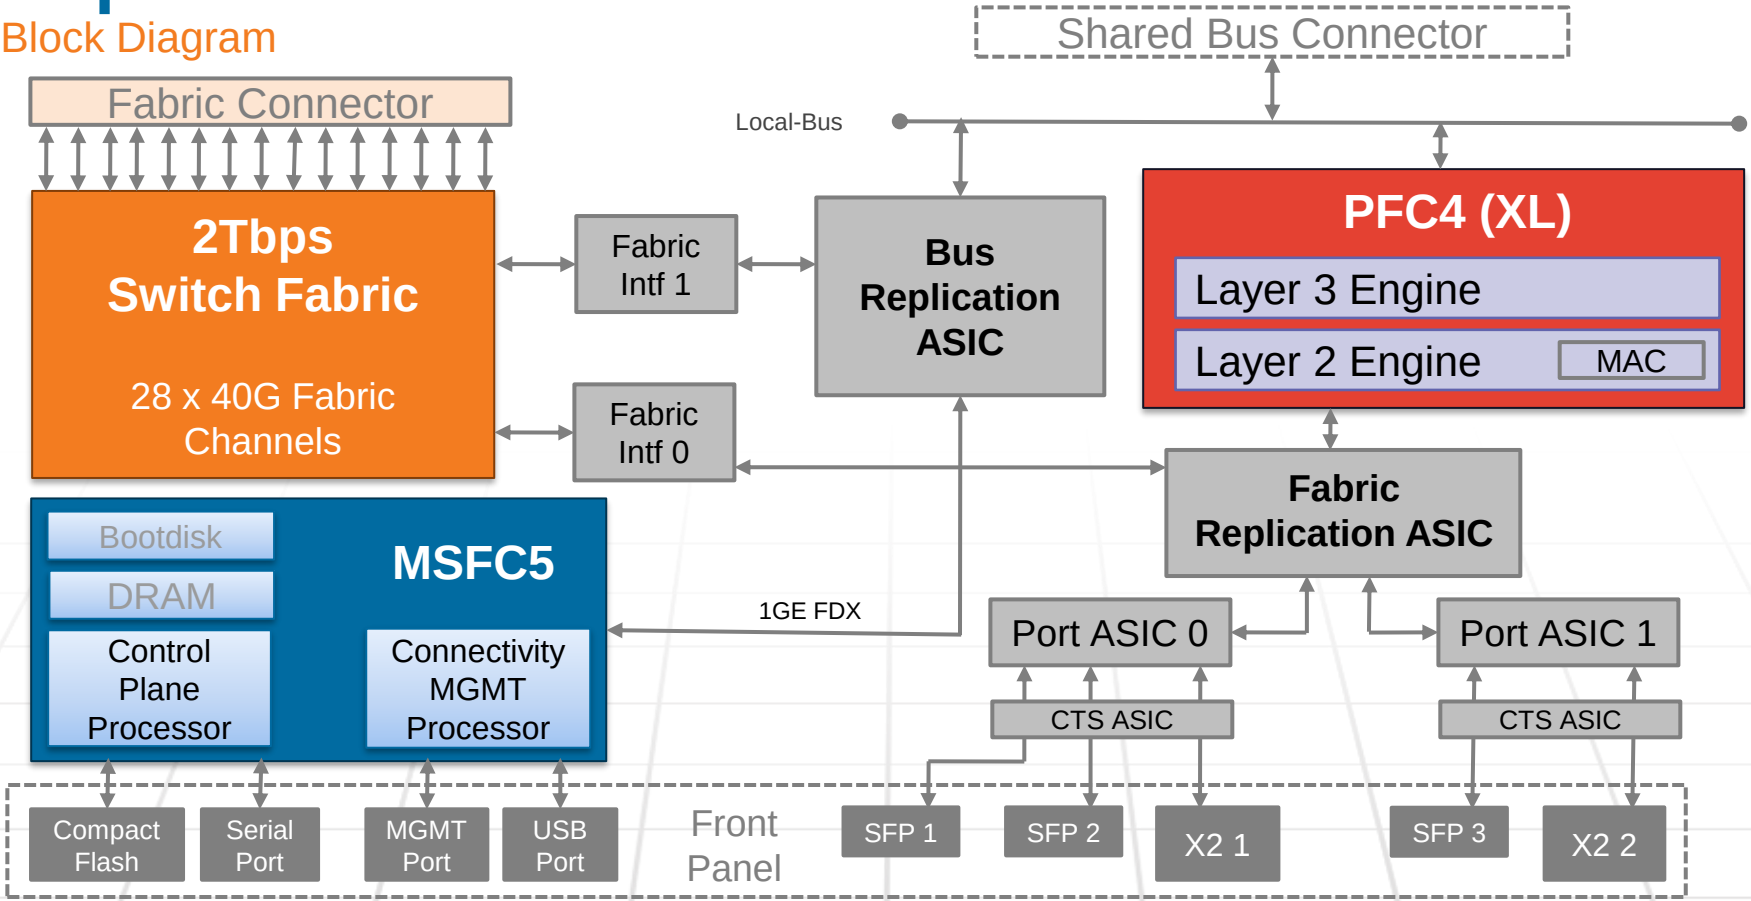

Agenda

- Chassis & Power
- **Supervisor Architectures**
  - VS-S2T-10G
  - **MSFC, PFC & Fabric**
  - C6880-X Baseboard
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

MSFC = Multilayer Switch Feature Card

The “Software” Control Plane for the System, where IOS runs…

Multilayer Switch Feature Card

MSFC5 Introduction

- **Single Dual-Core Processor**
  - Combines functionalities of the Switch Processor (SP) & the Route Processor (RP)
- **1.5Ghz CPU Performance**
- **2GB or 4GB DDR2 DRAM**
- Single Bootdisk file system
- Connectivity Management Processor (CMP)
- On-Board Failure Logging (OBFL)
- Mini Protocol Analyzer (MPA)

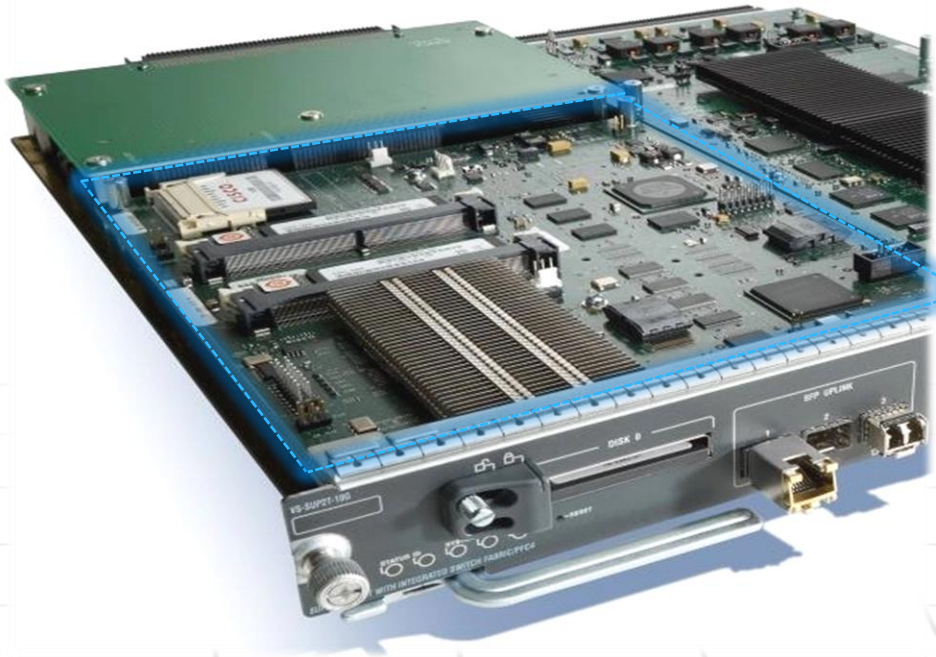

Multilayer Switch Feature Card 5

Block Diagram

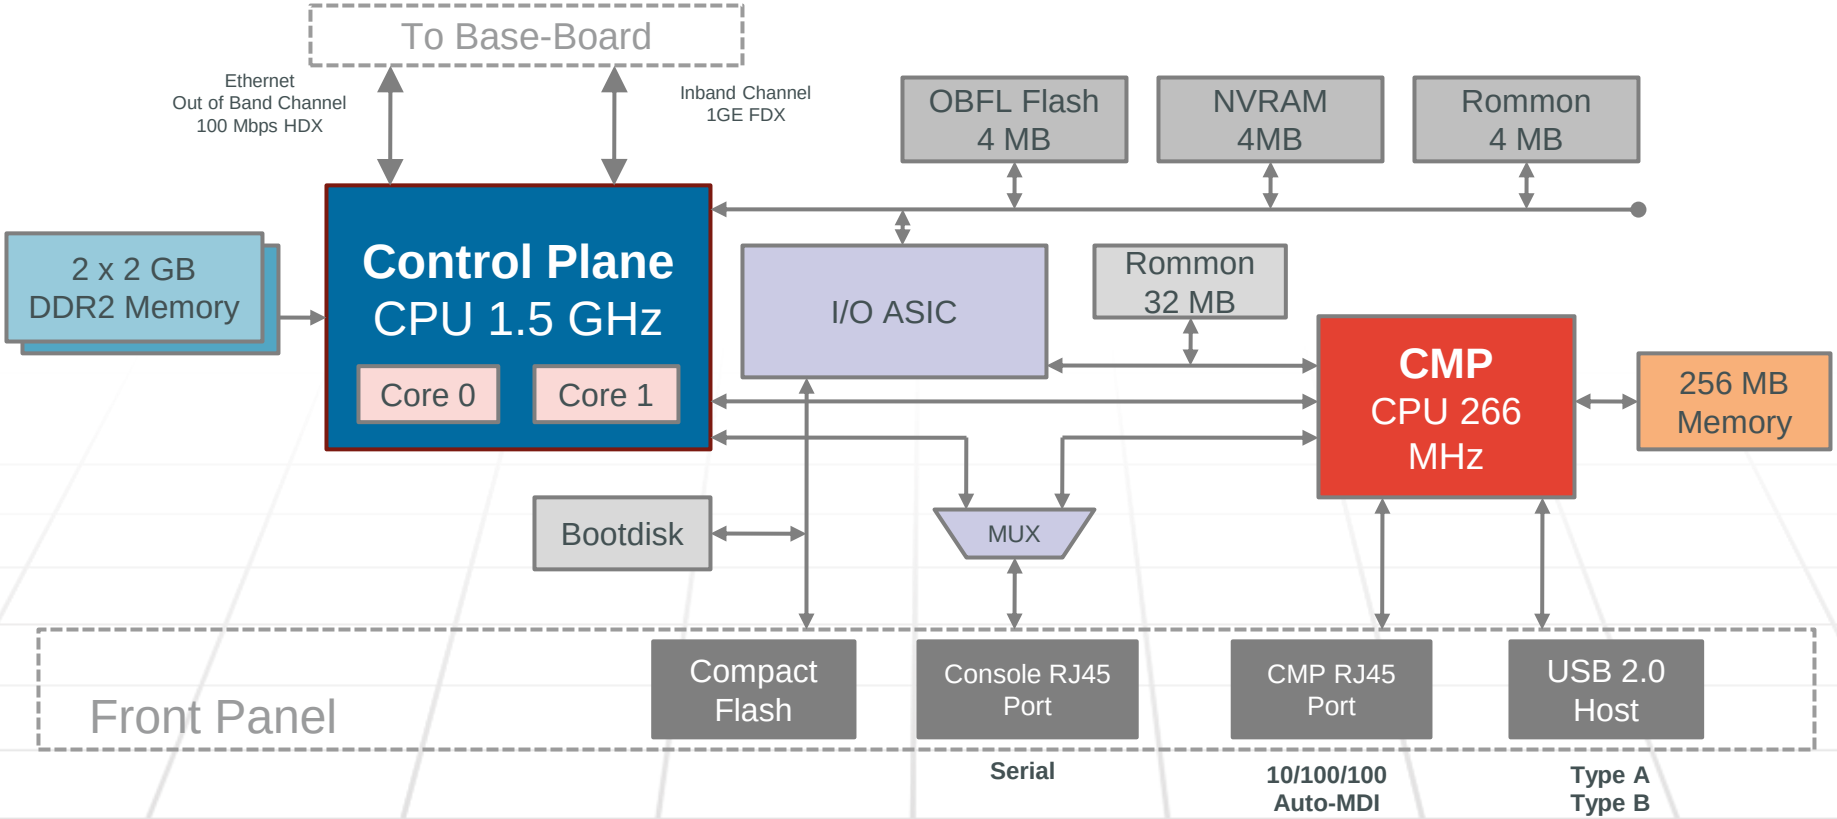

PFC = Policy Feature Card

The “Hardware” Control Plane, based on information learned by MSFC…

Policy Feature Card 4

PFC4 Introduction

Also applies to DFC4

- **Increased Hardware Performance**
  - **Up to 60Mpps L2/L3 Forwarding**
  - New IFE/OFE Lookup Process
- **Increased Hardware Scalability**
  - **256K or 1M FIB TCAM Entries**
  - 128K MAC Address CAM Entries
  - 64K or 256K Security & QoS ACL Entries
  - 512K or 1M Flexible Netflow (FnF) Entries
  - 16K Virtual Routing & Forwarding (VRF) Instances
  - 16K Bridge Domains & 128K Logical Interfaces
- **New & Enhanced Feature Capabilities**
  - **SGT & MACSEC for Cisco Trustsec (CTS)**
  - L2 + L3 + L4 Access Control List (ACL) Support
  - IPv4 & IPv6 RPF check for up to 16 Paths
  - Improved EtherChannel Load-Balancing

Policy Feature Card 4

Block Diagram

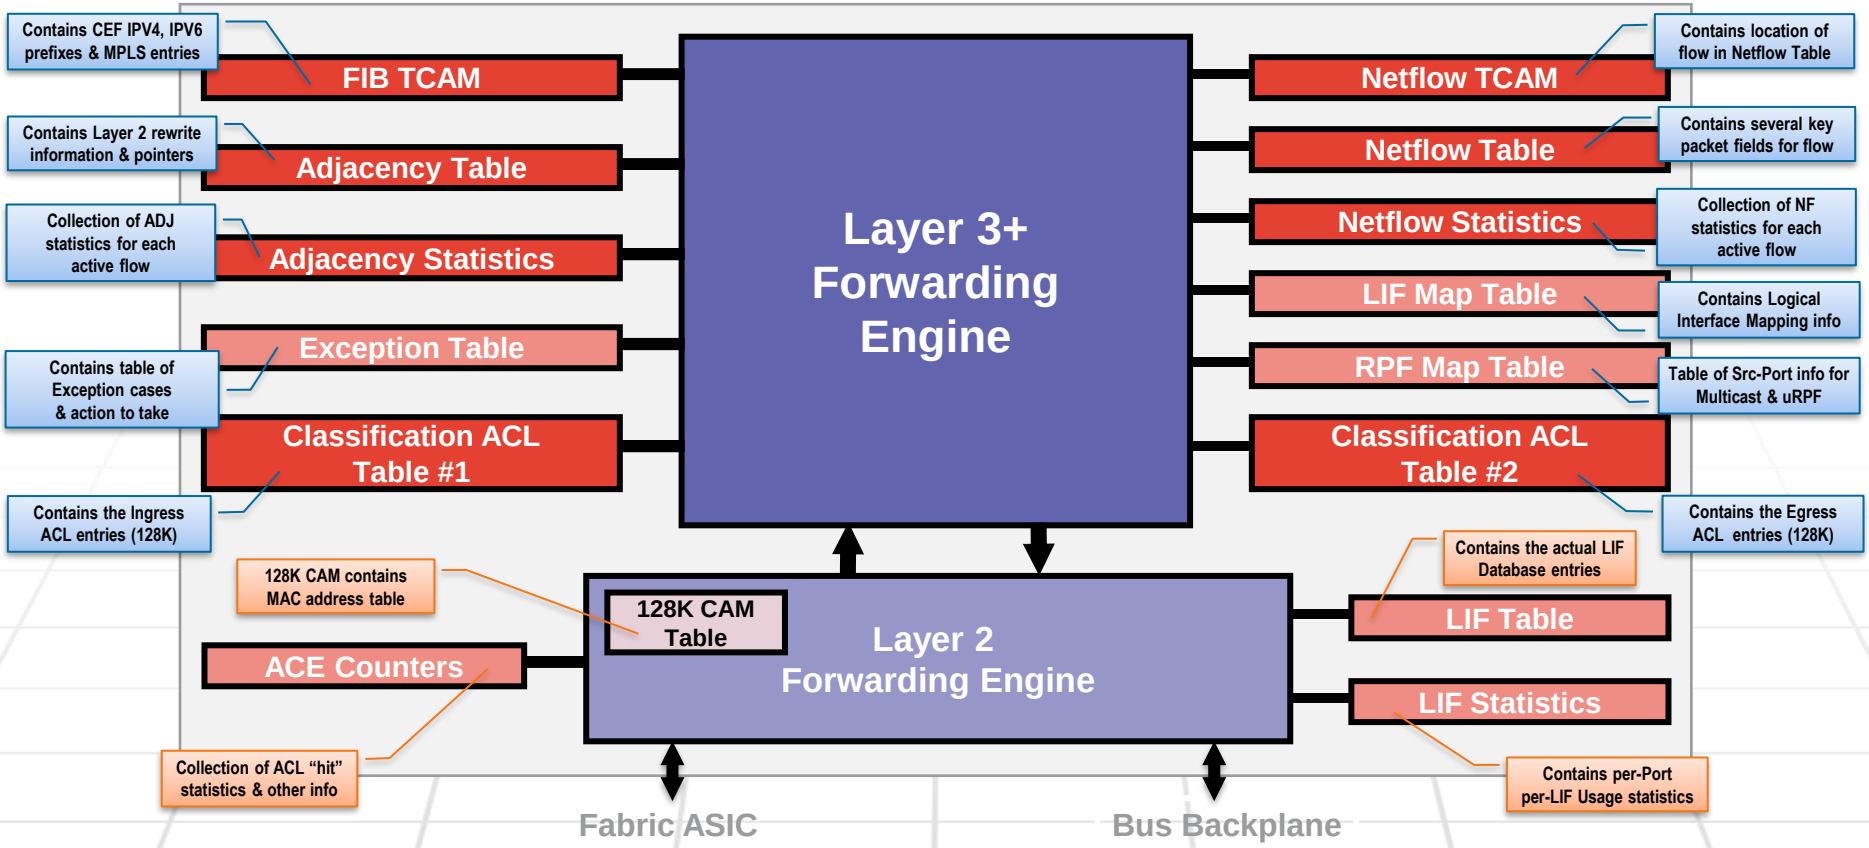

Policy Feature Card 4

EARL8 IFE/OFE Processing

The Forwarding Engine ASIC has 2 processing pipelines @ 60Mpps:

1. **Input Forwarding Engine (IFE)**
2. **Output Forwarding Engine (OFE)**

STEP 1 - As each packet Header enters the L3 ASIC, the “IFE” pipeline will perform an L3 Lookup and *Ingress* Security, QoS & Netflow processing…

STEP 2 - The Header is then merged with the IFE result and passed to the “OFE” pipeline, which does *Egress* Security, QoS & Netflow processing… to generate a final result.

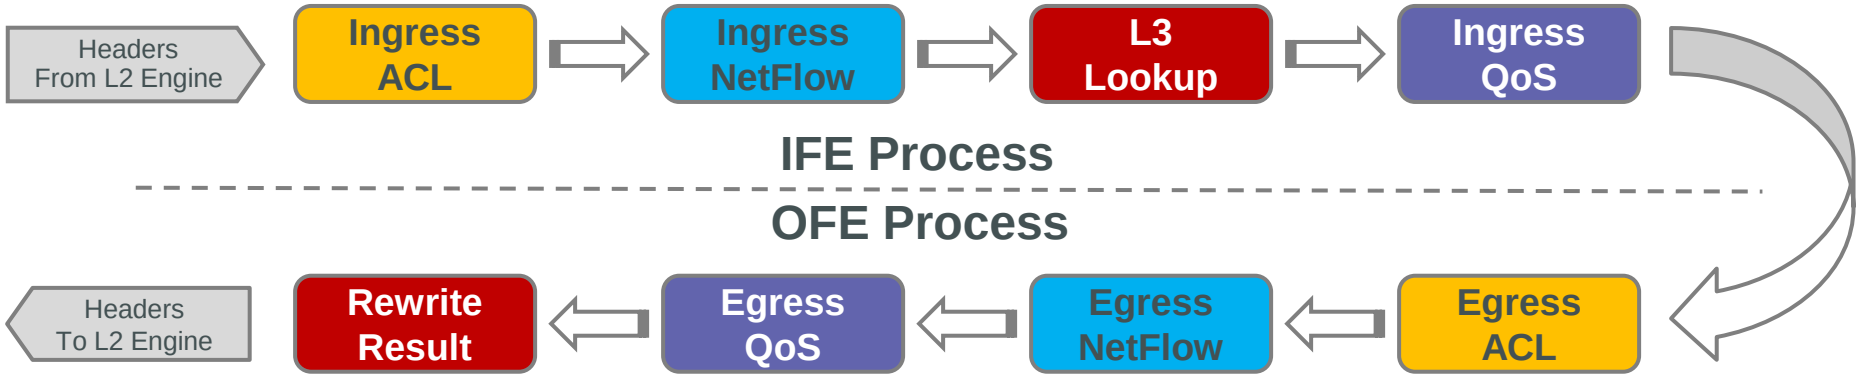

Switch Fabric = Data Plane (Back Plane)

A dedicated set of Crossbar Channels that interconnect All Cards…

2T Switch Fabric

Introduction

- **Integrated 2Tbps Switch Fabric**
  - **28 Channels** to support 6513-E & 6807-XL
  - Dual Queues (Lo & Hi) per Fabric Channel
  - Redundant Channel to Standby Fabric for faster traffic convergence, during an SSO...
- **Provides Backplane Interconnects**
  - Multiple Fabric Channels are distributed to each slot
  - Each Channel can independently operate at either **20Gb/sec** or **40Gb/sec***
  - Mixing 6700 & 6900 (20G & 40G) modules does not affect speeds of other modules

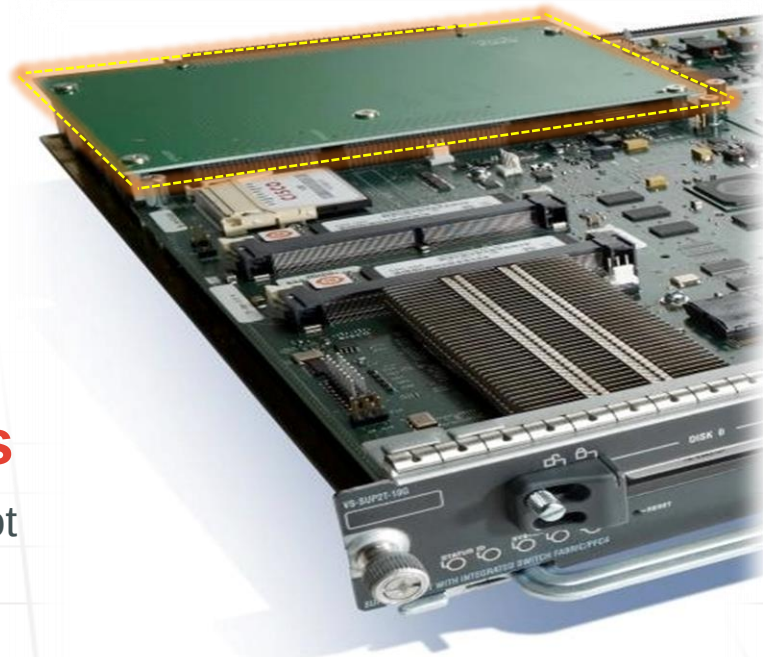

Crossbar Switch Fabric

Introduction

The Catalyst 6500 & 6800 series eliminates the earlier Bus-based limitations by using a "Crossbar" Switch Fabric as its backplane.

The Crossbar architecture is essentially 2N busses (where N is the number of LC's connected to the Switching Fabric), connected by N*N cross-points.

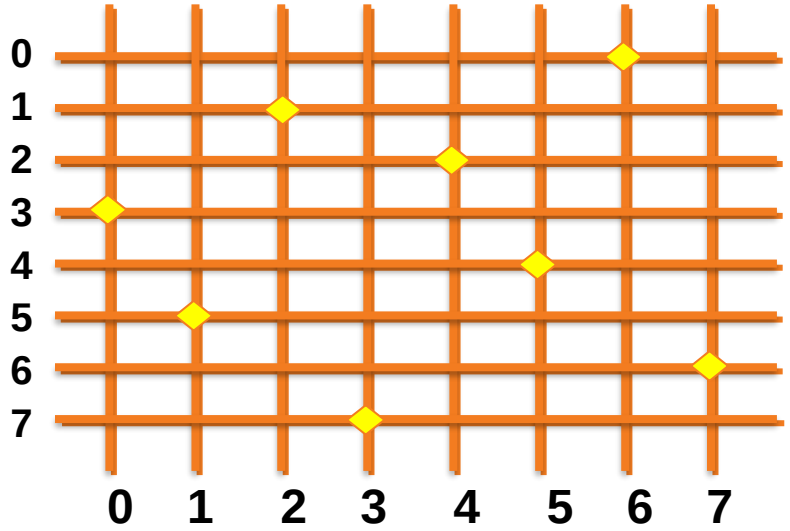

Switch Fabric - Logical Architecture

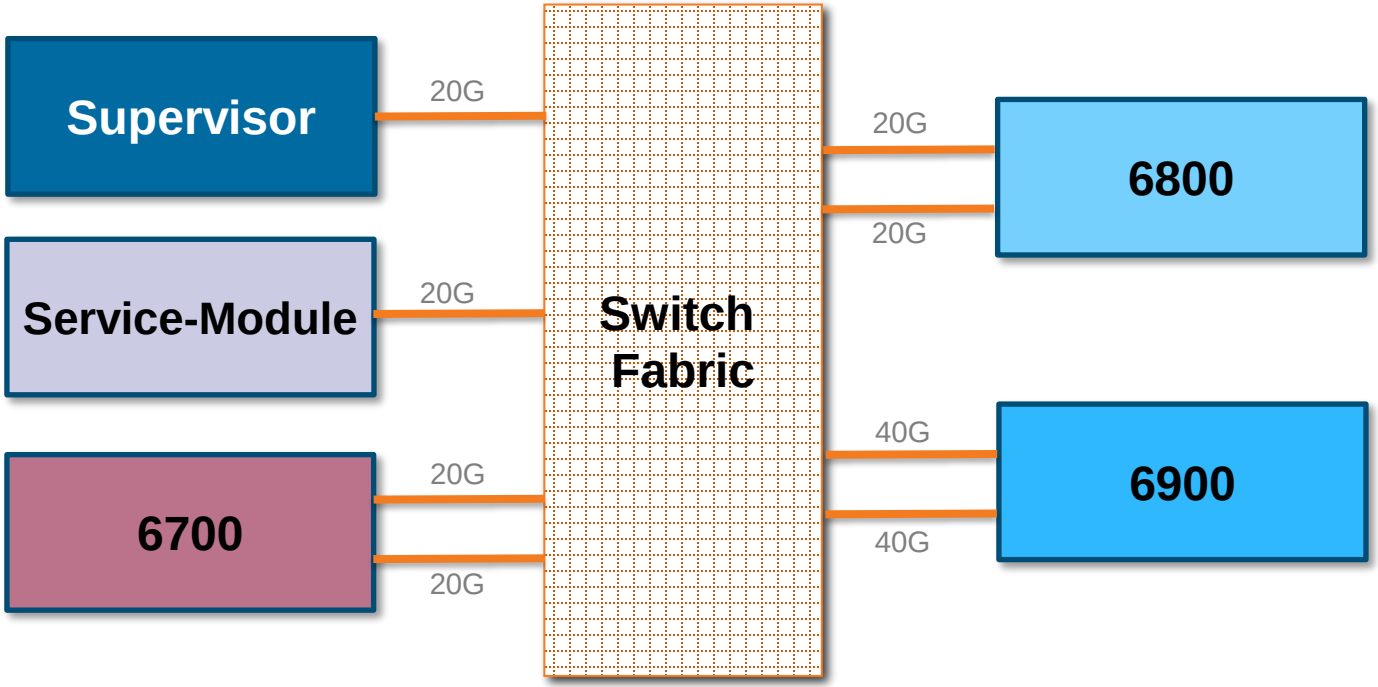

2T Switch Fabric

Block Diagram

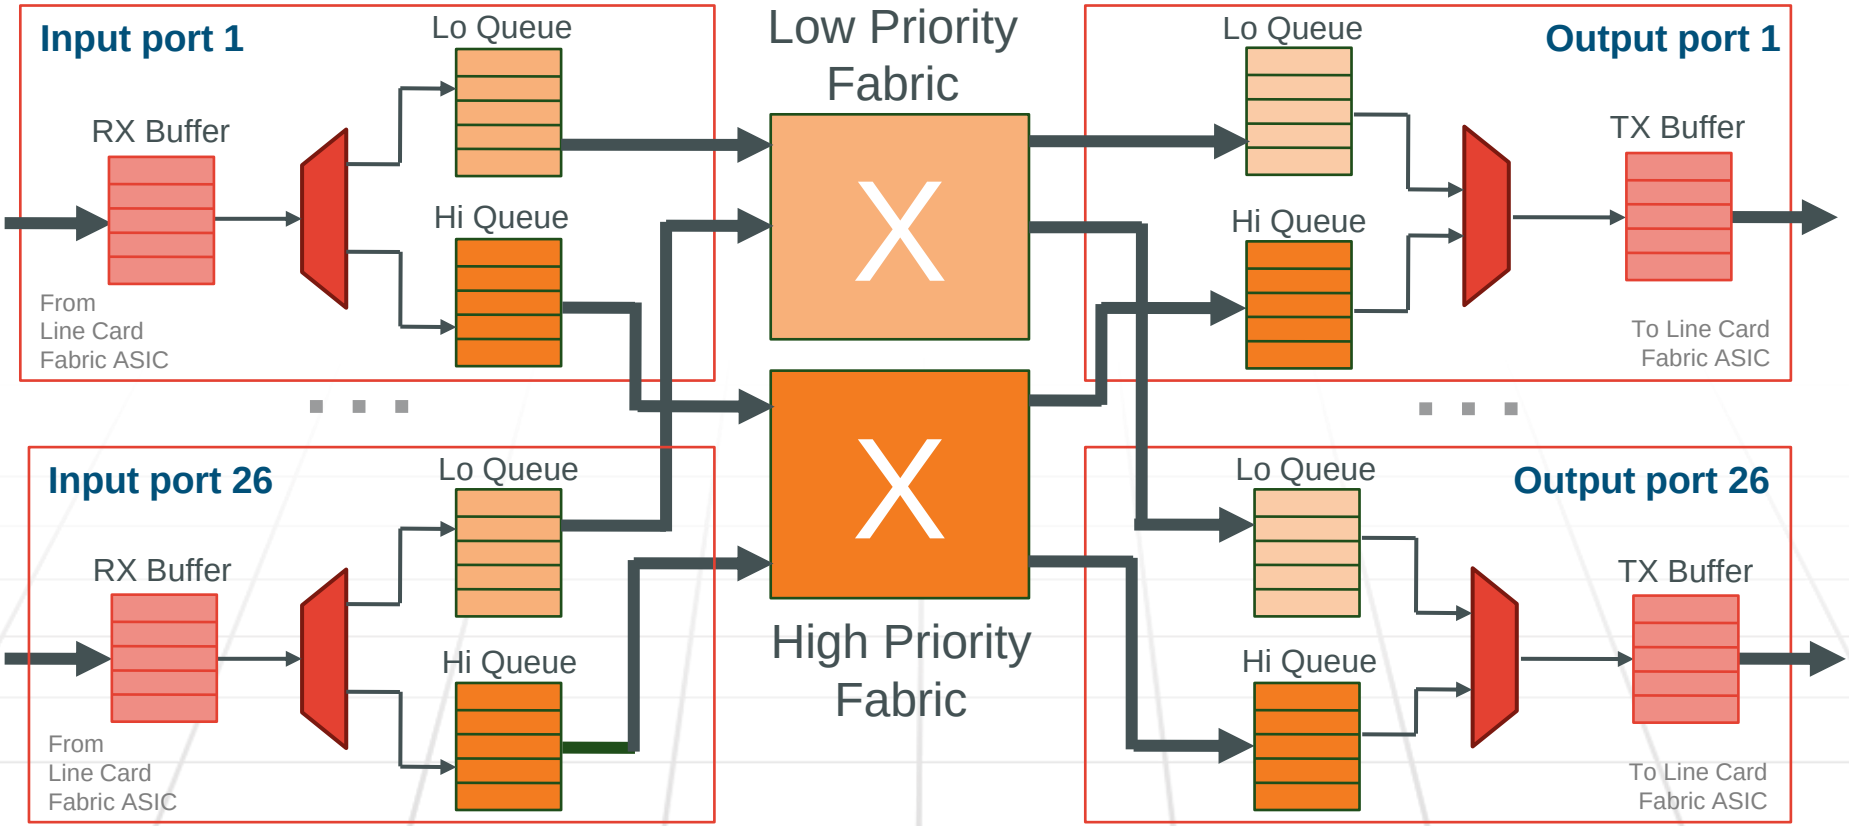

Catalyst 6807-XL

Fabric Channel Distribution

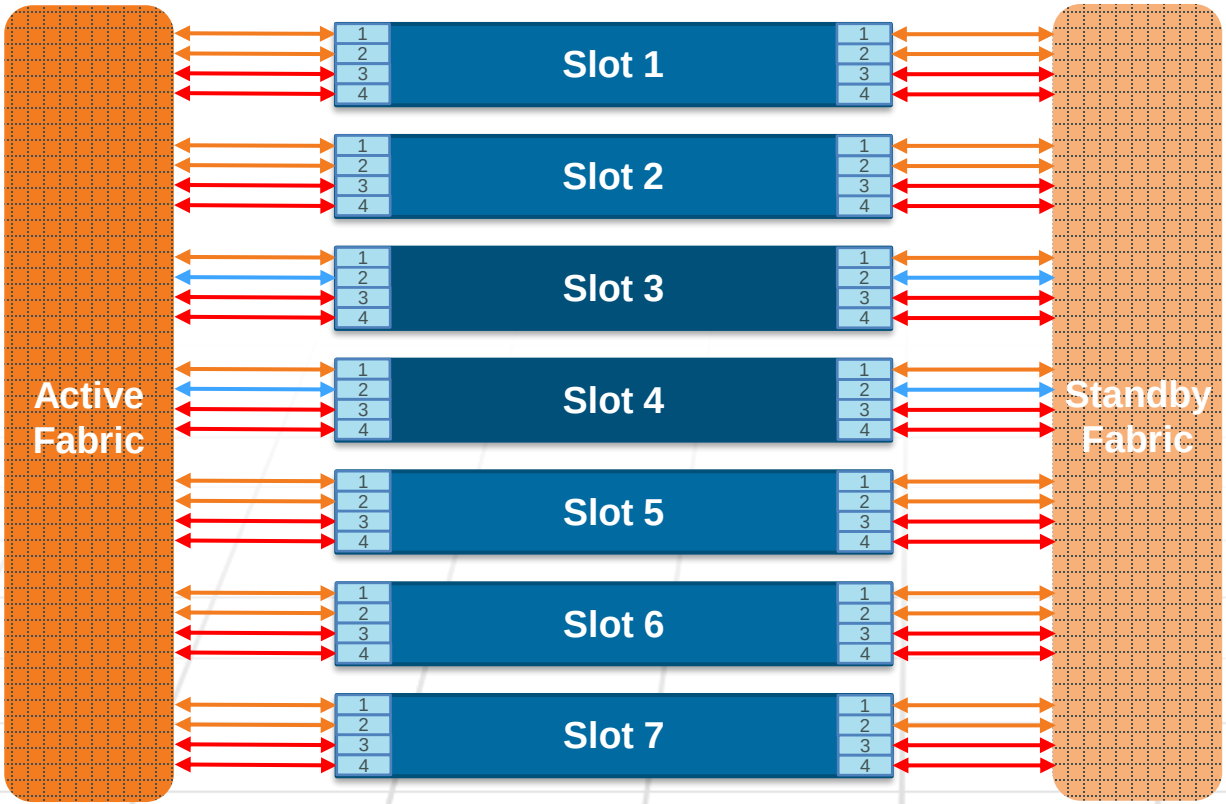

Current Fabric Channels

Sup (HA) Fabric Channels

Extra Fabric Channels (for future use)

Each “Fabric Channel” uses 8 SerDes Lanes

Each SerDes can use any of the following clock frequencies:

- 3.13 GHz for **20 Gbps**
- 6.25 GHz for **40 Gbps**
- 7.50 GHz for **55 Gbps**
- 15.0 GHz for **110 Gbps**

Catalyst 6807-XL

How Supervisor 2T operates

Support for up to 220G per slot

Supervisor Engine 2T:

- Using the same XBAR Fabric ASIC
- Can use 1 to 4 channels (per Card)
- **Increased per Channel bandwidth**
  - New Clock Frequencies (7.5 Ghz / 55G)
  - New Line Encodings (24/26 or 64/66b)
  - **Applicable Only to the Next-Gen Cards**
- No changes to MSFC5 or PFC4
- Local Channel for Uplinks @ 20Gbps

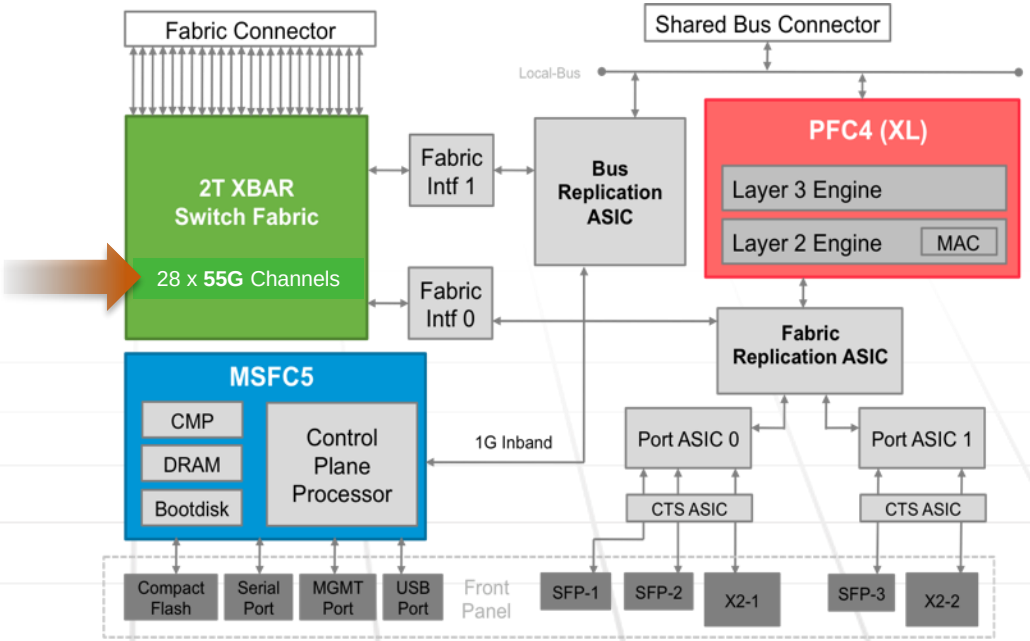

Agenda

- Chassis & Power
- **Supervisor Architectures**
  - VS-S2T-10G
  - MSFC, PFC & Fabric
  - **C6880-X Baseboard**
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Catalyst 6880-X

Base Board & System Controller

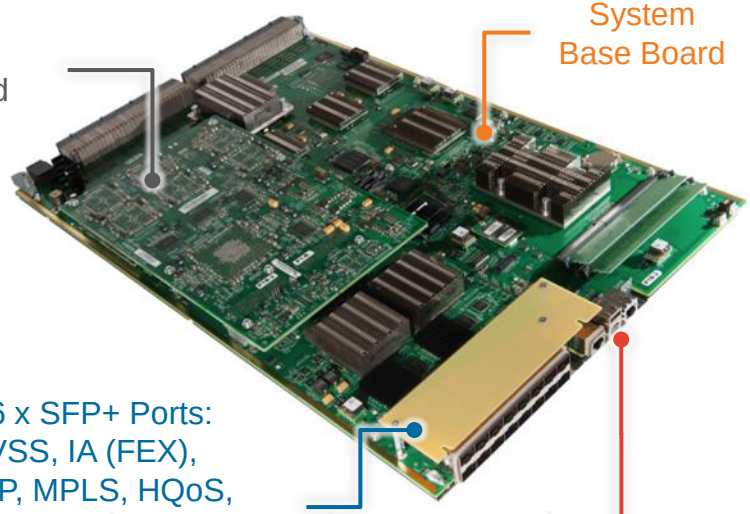

Two HW Options	6880-X-LE	6880-X
IPv4/v6 Routing Capability	256K/128K	2M/1M
Multicast Routes (IPv6)	64K	256K
Number of Adjacencies	1M	1M
MAC Addresses	128K	128K
L3 Interfaces	128K	128K
Security and QoS ACL	64K	256K

* Hardware Capable

Enhanced Control-Plane Scale with new X86 2GHz Dual Core CPU

Catalyst 6880-X

2M FIB TCAM

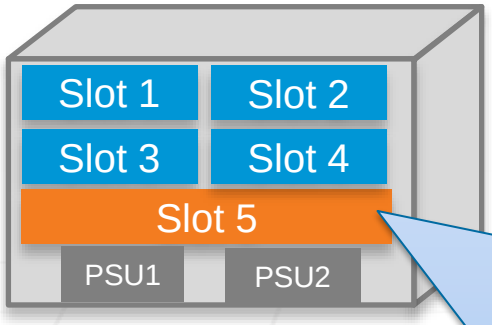


```
C6880-X.VSS# show module
Mod Ports Card Type                                  Model              Serial No.
--- ----- -------------------------------------- ------------------ -----------
  5    20  6880-X 16P SFP+ Multi-Rate (Active)     C6880-X   SAL17152E9G

C6880-X.VSS# show platform hardware capacity system
System Resources
  PFC operating mode: PFC4XXL
  Supervisor redundancy mode: administratively sso, operationally sso
  Switching resources: Sw/Mod    Part number                Series       CEF mode
                       1/5       C6880-X                 supervisor           CEF
                       2/5       C6880-X                 supervisor           CEF

C6880-X.VSS# show platform hardware capacity forwarding
L2 Forwarding Resources
           MAC Table usage:   Sw/Mod  Collisions  Total       Used       %Used
                              1/5              0 131072         8          1%
                              2/5              0 131072         8          1%

L3 Forwarding Resources
             FIB TCAM usage:                        Total       Used       %Used
                  72 bits (IPv4, MPLS, EoM)  2097152          51           1%
                 144 bits (IP mcast, IPv6)   1048576          31           1%
                 288 bits (IPv6 mcast)        524288           1           1%
```

Catalyst 6880-X:

Fixed Supervisor Design

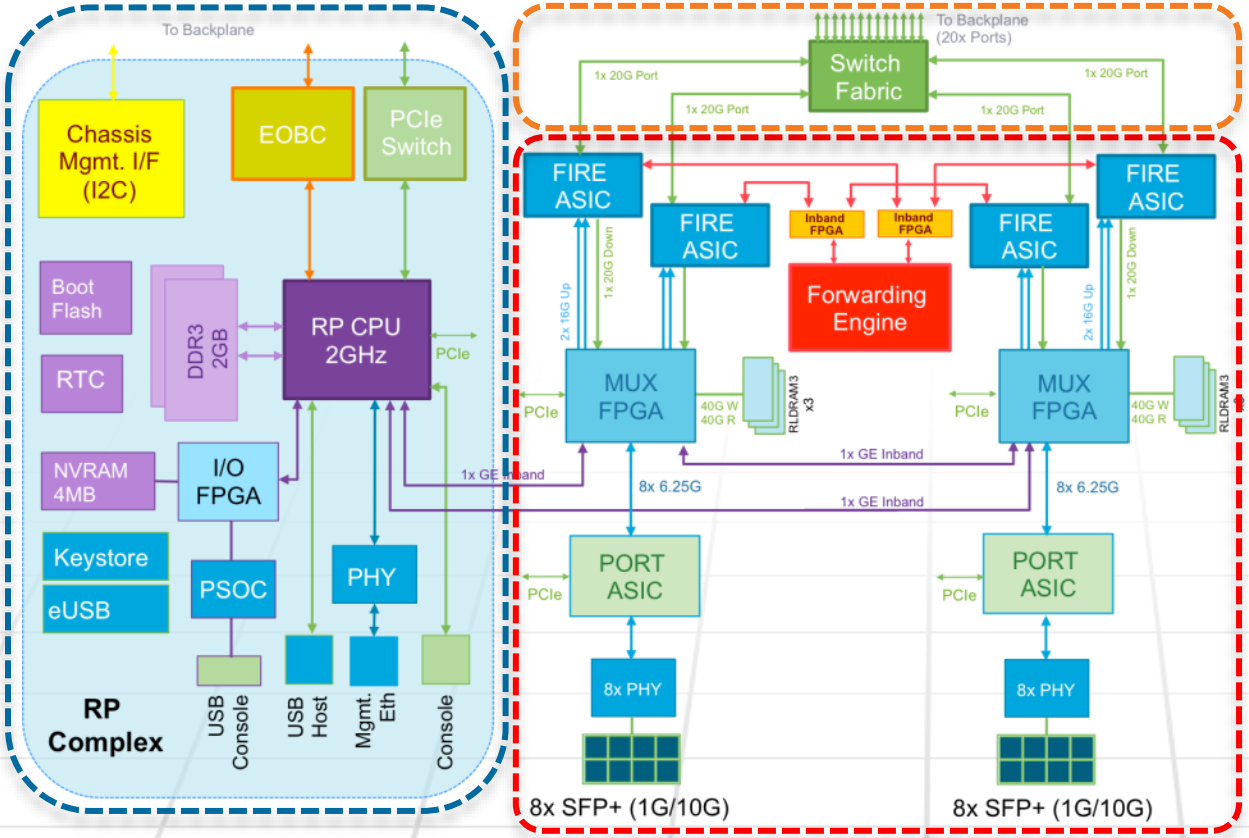

3 Main Components:

- **RP Complex**
- **Baseboard**
- **Switch Fabric**

Catalyst 6880-X

Fixed Supervisor - RP Complex

NEW

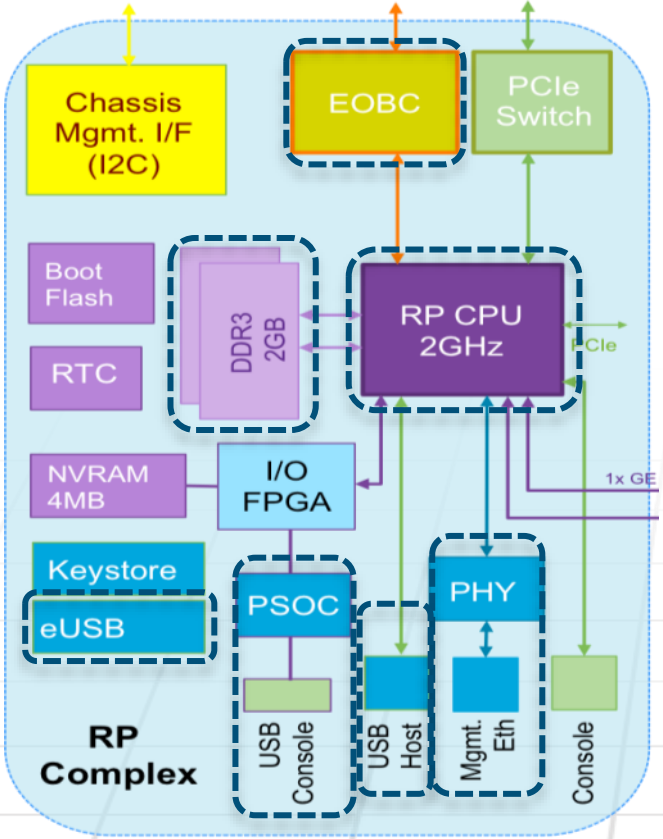

RP Complex Highlights:

- Essentially the same as MSFC5
- New 2.0Ghz X86 Dual Core CPU
- 2 or 4GB of DDR3 ECC SDRAM
- CMP replaced with direct RJ45 (Mgmt0)
- Support for USB Type A File System
- Support for USB Type B Serial Console
- Compact Flash replaced with eUSB
- New Switched EOBC Interface

Catalyst 6880-X:
Fabric Channel Distribution

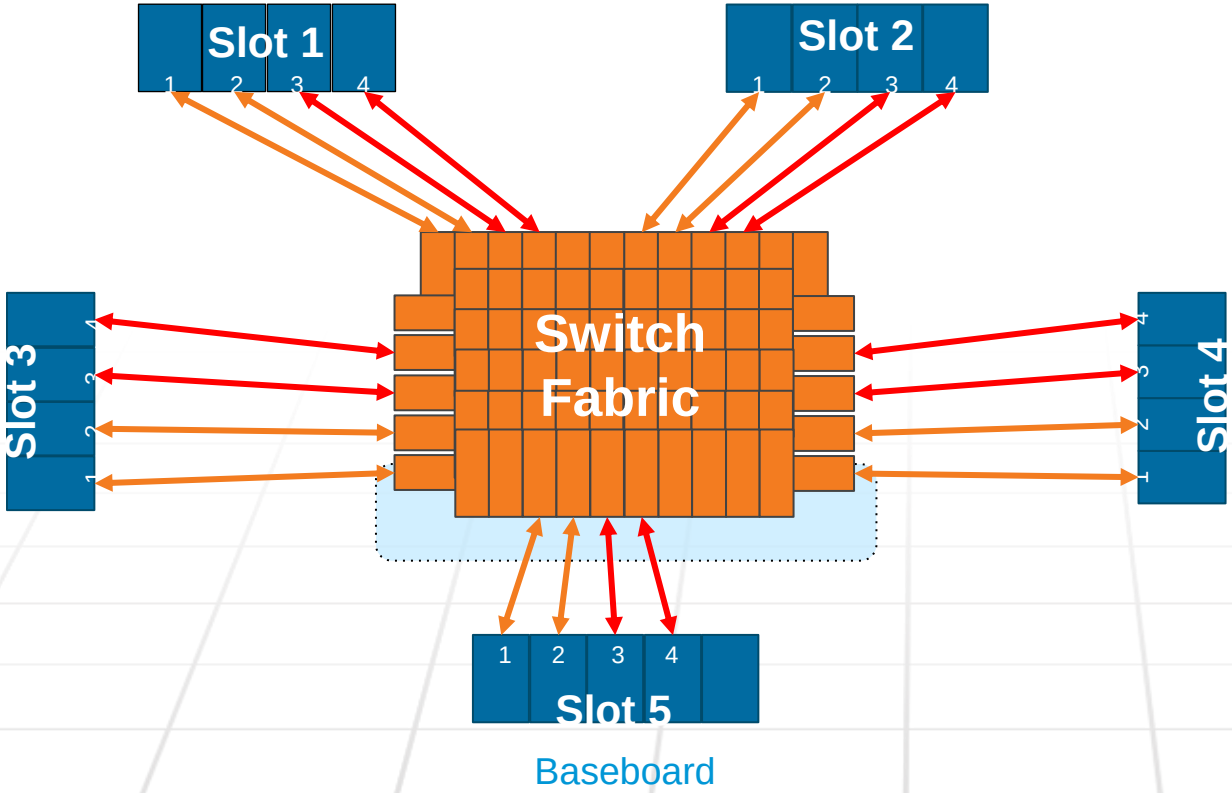

Current Fabric Channels

Extra Fabric Channels (for future use)

Each “Fabric Channel” uses 8 SerDes Lanes

Each SerDes can use any of the following clock frequencies:

- 6.25Ghz for 40Gbps
- 7.50Ghz for 55Gbps

Catalyst 6880-X

Fixed Supervisor - Baseboard & 16 x 1/10GE Uplink

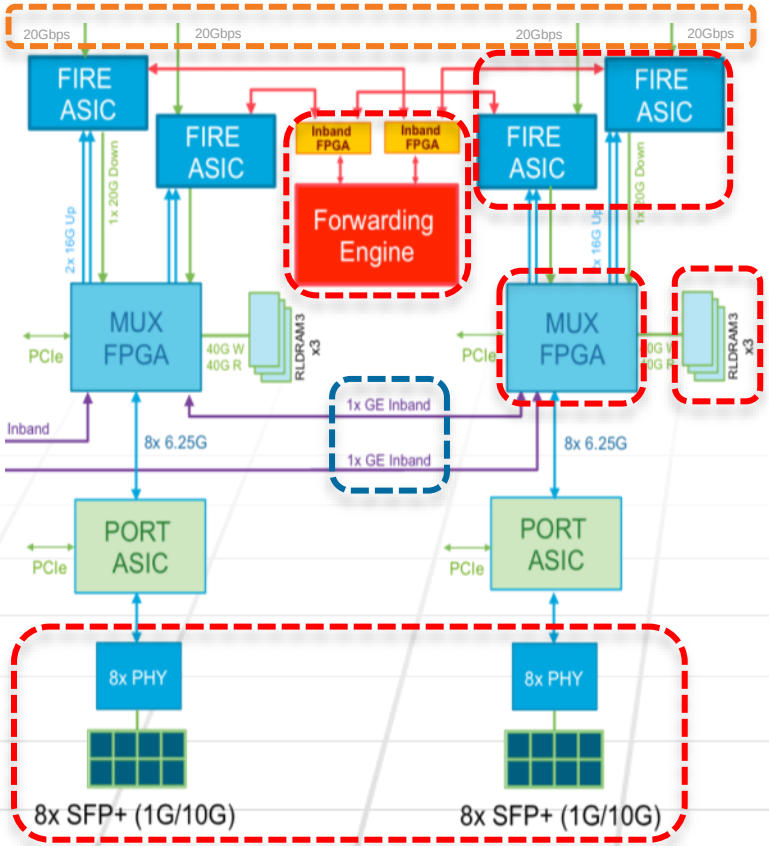

Baseboard Highlights:

- Essentially same as Sup2T + 6904-40G
- 16 x SFP+ (Multi-Rate) Ethernet Ports
- 80Gbps to Switch Fabric (2 Modes)
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- Enhanced DFC4-E Forwarding Engine
- Improved 40Gbps Fabric/Replication ASIC
- New 40+Gbps Port Interface MUX FPGA
- Local 1GE Inband Links to LCP Complex
- RLDRAM3 Packet Buffers on MUX FPGA
- 192MB per MUX FPGA, 24/48MB per Port

Agenda

- Chassis & Power
- Supervisor Architectures
- **Module Architectures**
  - **6700 Series (CEF720)**
  - **6800 Series (dCEF720)**
  - **6900 Series (dCEF2T)**
  - Integrated Service Modules
  - C6880-X Series (dCEF2T)
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Catalyst 6500-E & 6807-XL Line Cards

6700 Series
with CFC

6800 Series
with DFC4

6900 Series
with DFC4

Catalyst 6500 & 6800 - Fabric Line Cards

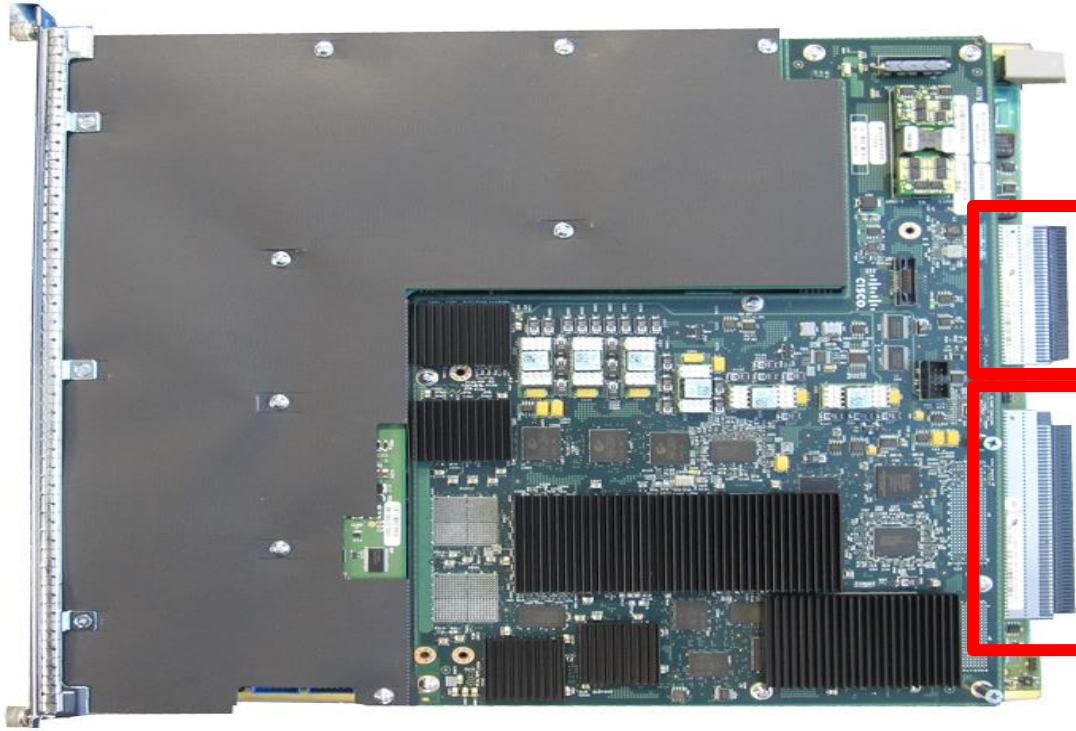

Fabric Connector

Connects 1-2 fabric channels that connect to Switch Fabric

Each Fabric Channel runs @ 20G (Sup720) or 40G (Sup2T)

Used to forward Data portion of packet to other Line Cards

Compact Packet Headers sent to Forwarding Engine

Classic Connector

CEF720 Architecture (6700 Series)

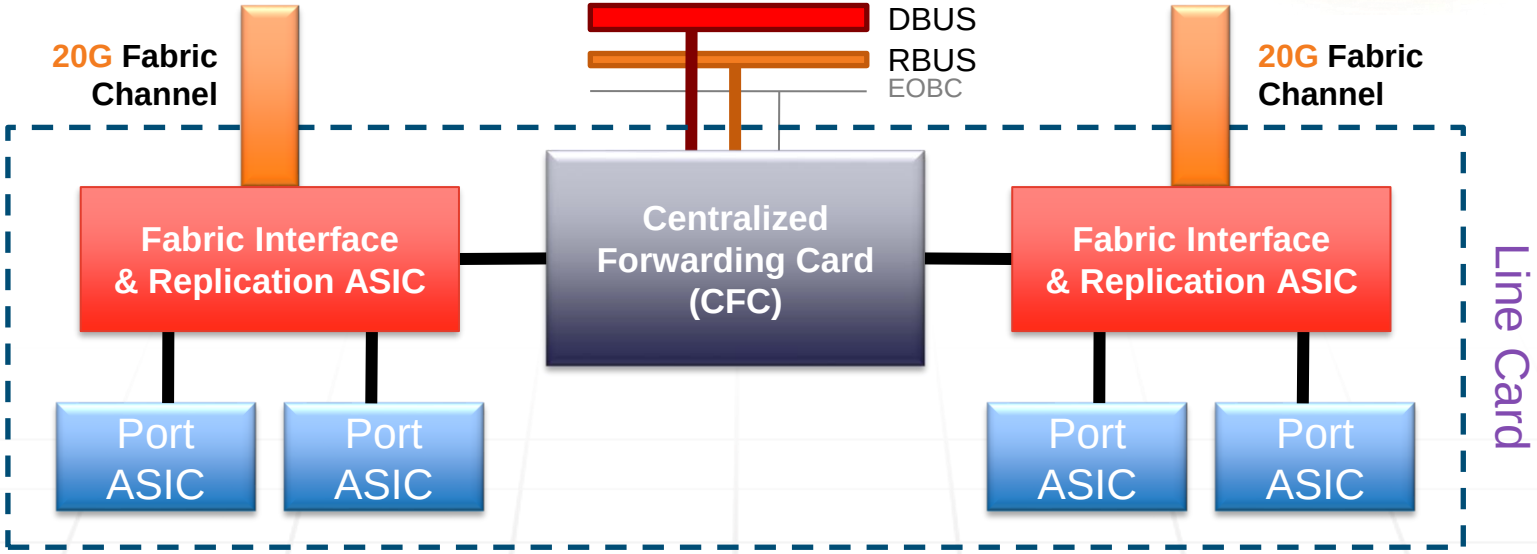

CEF720 has NO local forwarding lookup capability

Uses CFC card to send Packet header to Supervisor (PFC) over Central BUS

Ingress & Egress packet queuing and scheduling is done in the Port ASIC

All Data traffic is sent over Fabric Channels to destination Line Card

dCEF720 Architecture (6700 & 6800 Series)

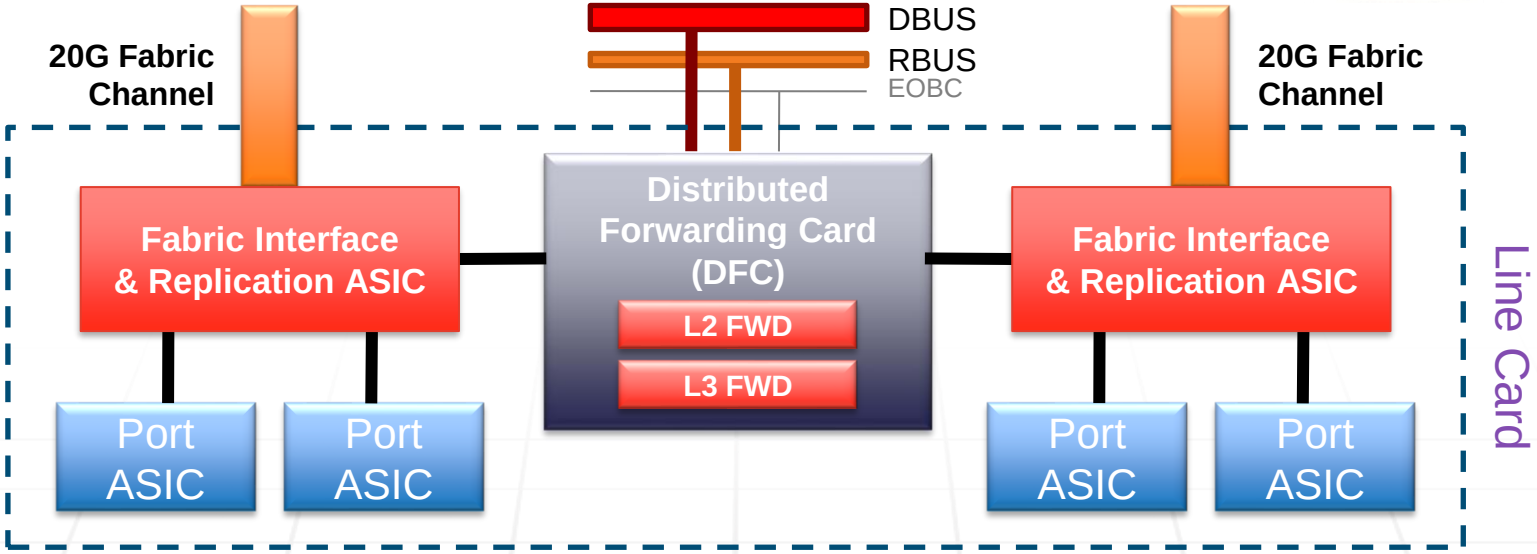

dCEF720 uses DFC3 or DFC4 for local (distributed) forwarding lookup

Module has NO connection to Central Bus

DFC3 / DFC4 contains same Hardware & Logic as PFC3 / PFC4 on Supervisor

Ingress & Egress packet queuing and scheduling is done in the Port ASIC

Centralized Forwarding Card (CFC)

The Centralized Forwarding Card (CFC) provides BUS connectivity for *centralized* (PFC-based) *forwarding lookups* ONLY…

The CFC comes default on 6700 modules and provides a connection to the DBUS & RBUS

All L2 / L3 Forwarding "decisions" are made by the PFC and "results" are returned on the RBUS

Actual DATA Forwarding is via the Switch Fabric...

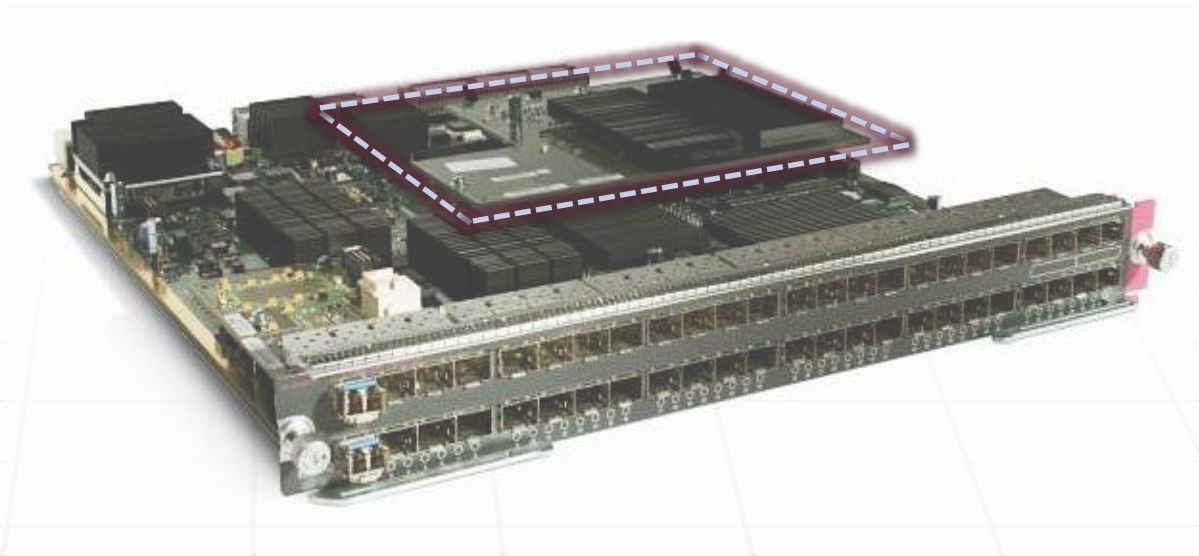

Distributed Forwarding Card 4

The DFC4 is an option for 6700 Cards, and comes pre-installed on 6800 & 6900 Series Cards

The DFC4 stores a local copy of the L2/L3+ forwarding info, as well as Netflow, Security & QoS ACL's

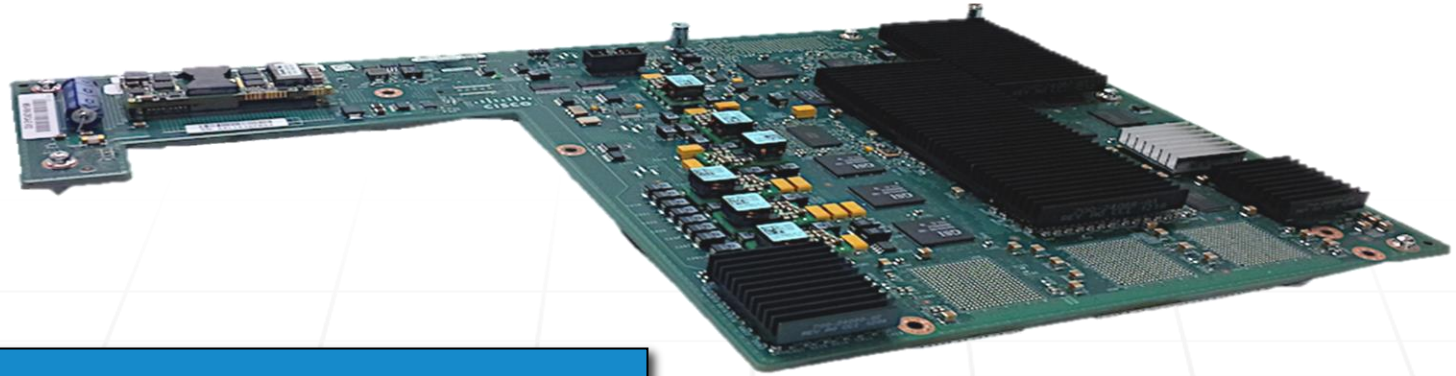

The DFC4 supports forwarding rates up to 60Mpps

The DFC4 includes same IFE/OFE capabilities & increased table sizes

Two different versions of the DFC4 are supported…

- **DFC4-A (AXL)**
- **DFC4-E (EXL)**

Catalyst 6500/6800 Modules

DFC3/4 interoperability with PFC3/4

For Your Reference

	PFC3A	PFC3B	PFC3BXL	PFC3C	PFC3CXL	PFC4	PFC4XL
DFC3A		Operate as PFC3A	Operate as PFC3A	Operate as PFC3A	Operate as PFC3A	X	X
DFC3B	Operate as DFC3A		Operate as PFC3B	Operate as PFC3B	Operate as PFC3B	X	X
DFC3BXL	Operate as	Operate as		Operate as PFC3B and	Operate as	X	X

CFC or DFC = Centralized or Distributed

CFC connects to DBUS/RBUS, so that PFC can perform Forwarding Lookup

DFC enables Local (Distributed) Forwarding Lookup on each Line Card

dCEF2T Architecture (6908-10G)

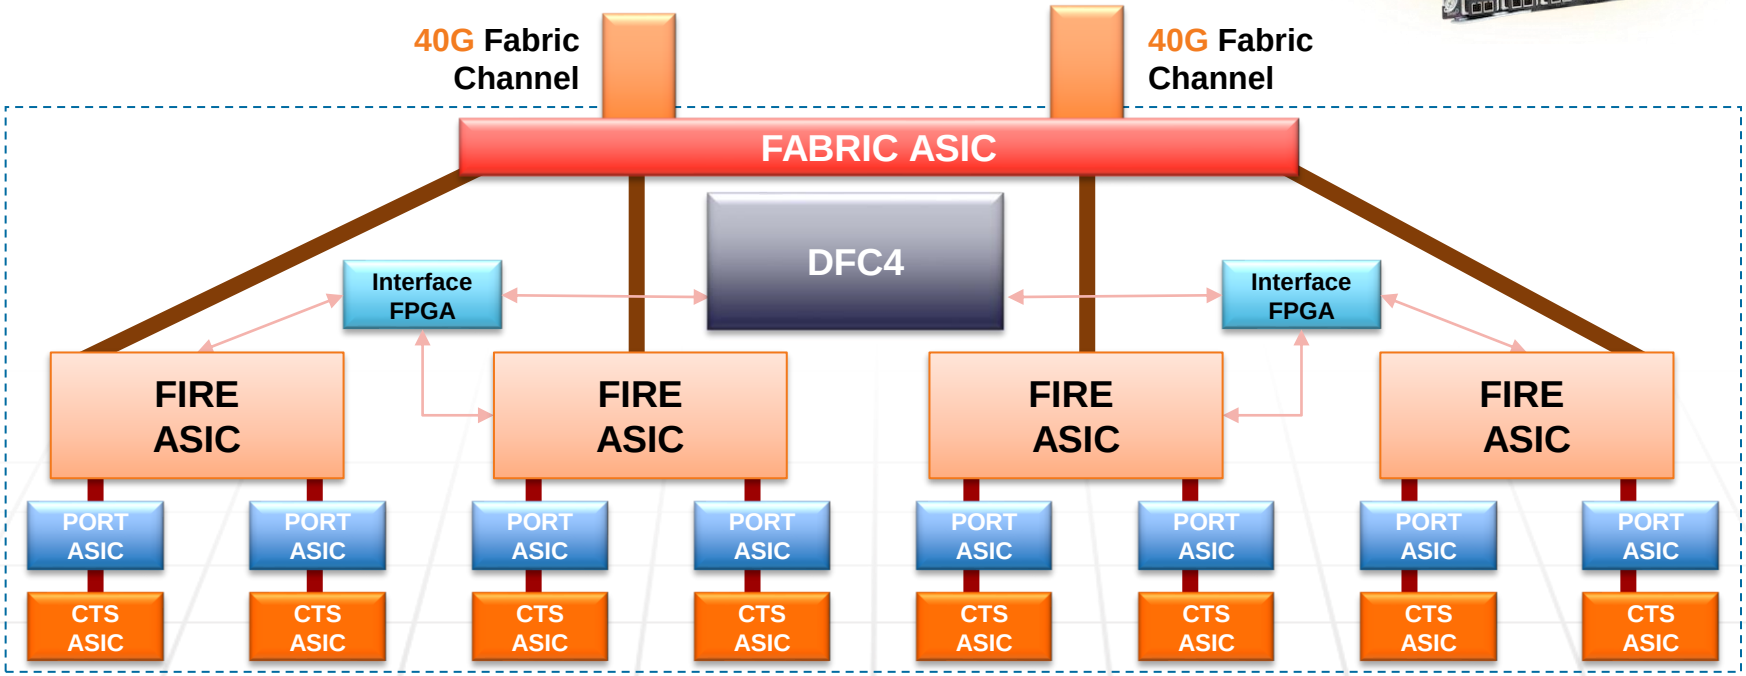

- dCEF2T uses DFC4 for local forwarding
- Module has NO connection to Central Bus
- Double the number of PORT & FIRE ASICs
- CTS ASICs provide wire-rate Encryption & Decryption

dCEF2T Architecture (6904-40G)

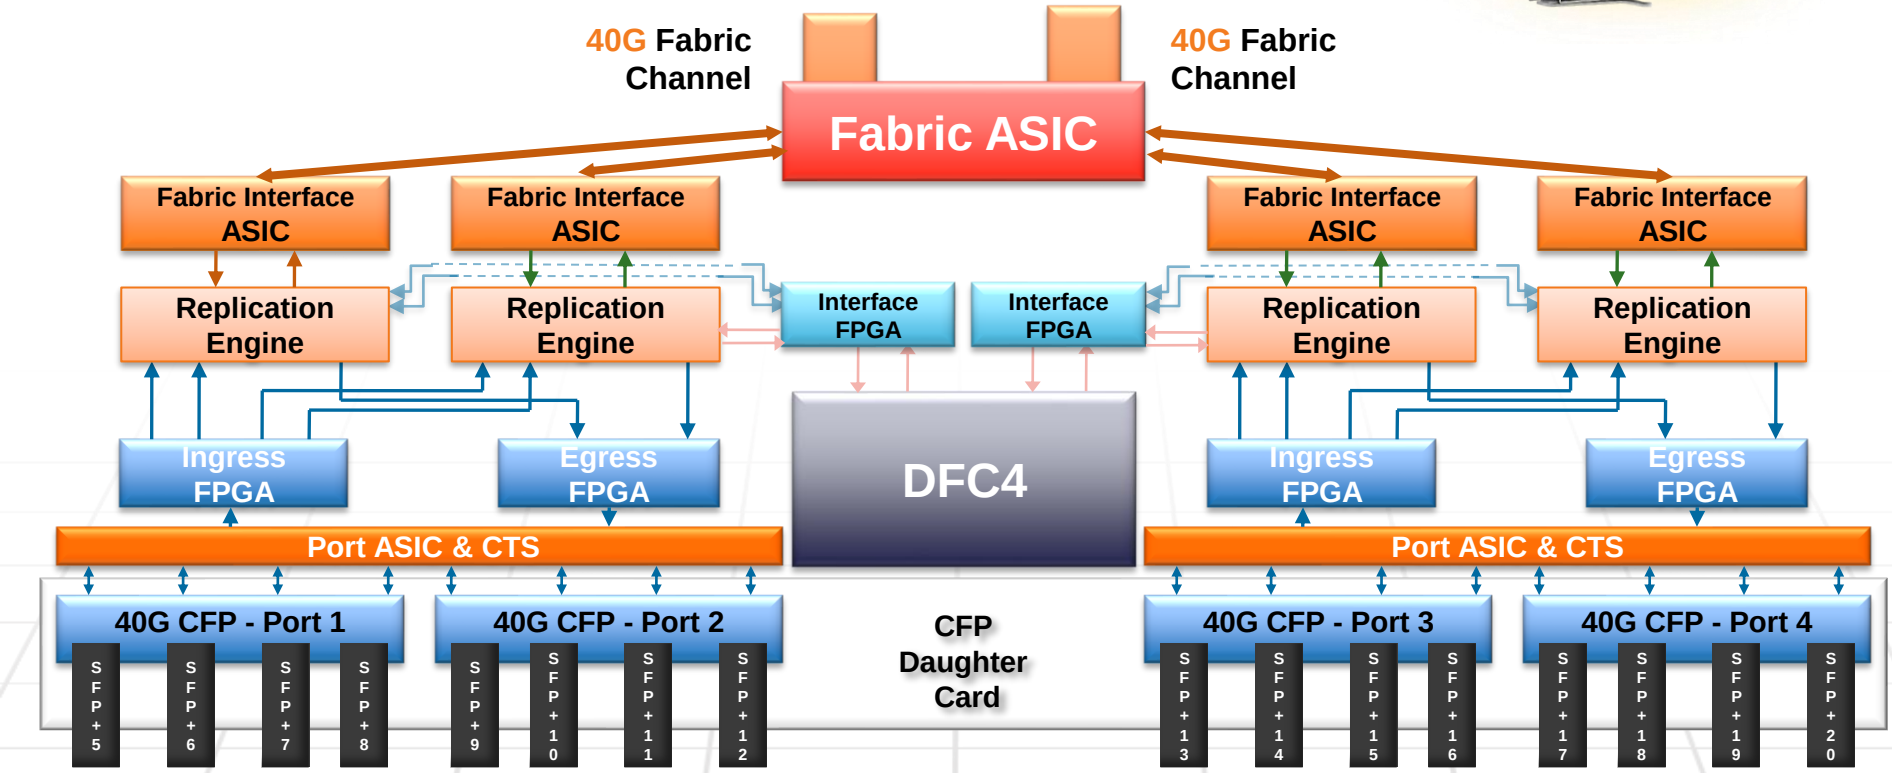

Catalyst 6807-XL

How current Line Cards operate

Support for 40G & 80G per slot

6700/6800 & 6900 Series:

- Single DFC4 (or CFC)
- 2 Fabric Channels (per Sup)
- **Per-Channel bandwidth:**
  - **40G for 6700/6800 series**
  - **80G for 6900 series**
- Same performance and operation as with 6500-E
- No performance issue with mixed speeds of 20G and 40G

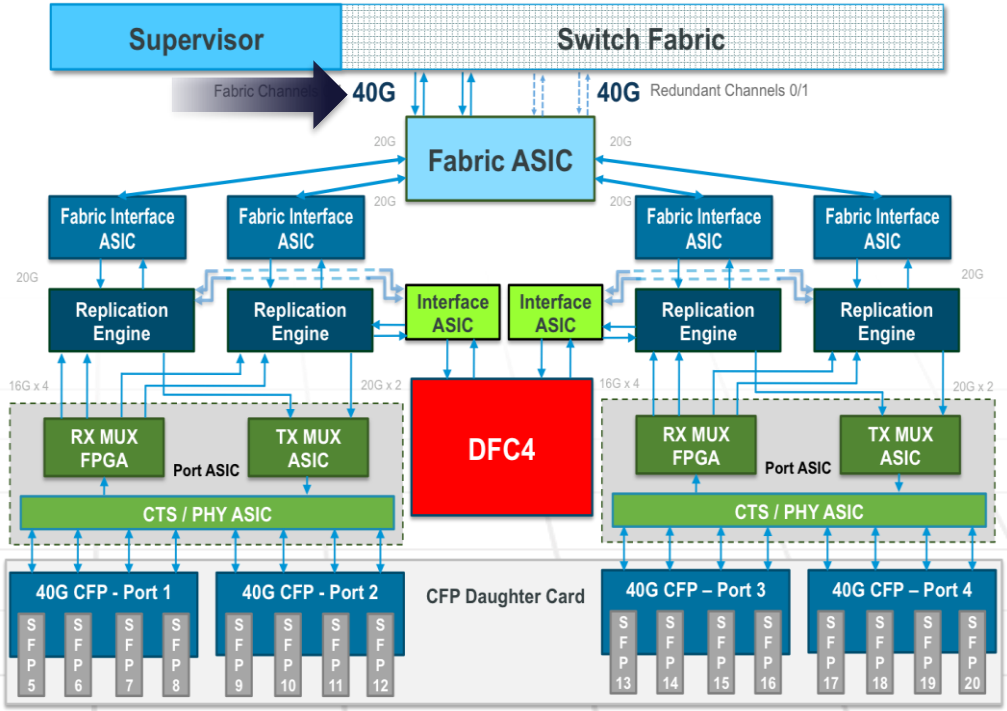

Agenda

- Chassis & Power
- Supervisor Architectures
- **Module Architectures**
  - 6700 Series (CEF720)
  - 6800 Series (dCEF720)
  - 6900 Series (dCEF2T)
  - **Integrated Service Modules**
  - C6880-X Series (dCEF2T)
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Catalyst 6500-E & 6807-XL

High Performance Integrated L4-L7 Service Modules

Integrate Wired & Wireless Management

BYOD

Next Generation Wireless Controller: **WiSM2**

Performance	**20 Gbps**
Access Points	**500 - 1000**
Wireless Clients	**15,000**
Concurrent AP Upgrades	**Up to 500**
Mobility, Domain Size	**Up to 18,000 APs**

Accelerate & Balance Application Performance

SLB

Next Generation Load Balancer: **ACE-30**

16 Gbps	Performance
6 Gbps	Compression
30,000	Transactions per Second
250	Virtual Context
4000	VLANs

Enhance Visibility, Accelerate Troubleshooting

NMS

Next Generation Network Analysis: **NAM3**

Monitoring Performance	**16 Gbps**
Capture to External Disk	**6 Gbps**
Performance Analytics	**1588 Timestamps**

Deliver Robust, Integrated, Streamlined Security

SEC

Next Generation Firewall & DPI: **ASA-SM**

64 Gbps	System Performance
16 Gbps	Performance per SM
10,000,000	Concurrent Sessions
300,000	Connections per Second

69

Agenda

- Chassis & Power
- Supervisor Architectures
- **Module Architectures**
  - 6700 Series (CEF720)
  - 6800 Series (dCEF720)
  - 6900 Series (dCEF2T)
  - Integrated Service Modules
  - **C6880-X Series (dCEF2T)**
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Catalyst 6880-X:

16-port SFP+ Multi-Rate Port Card

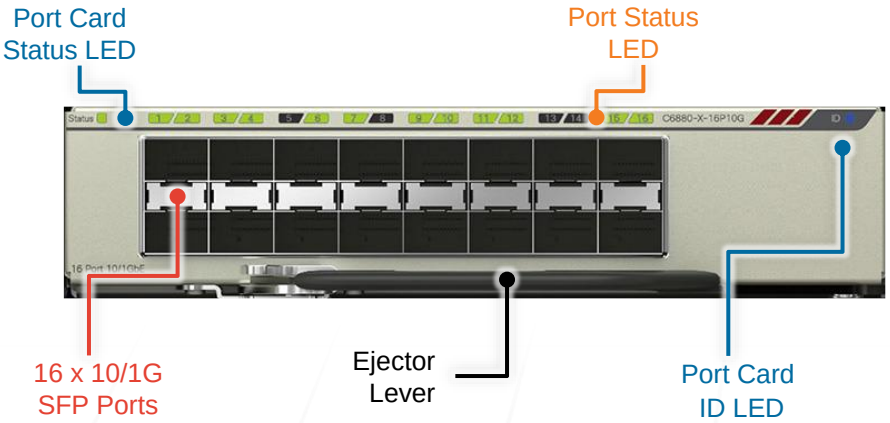

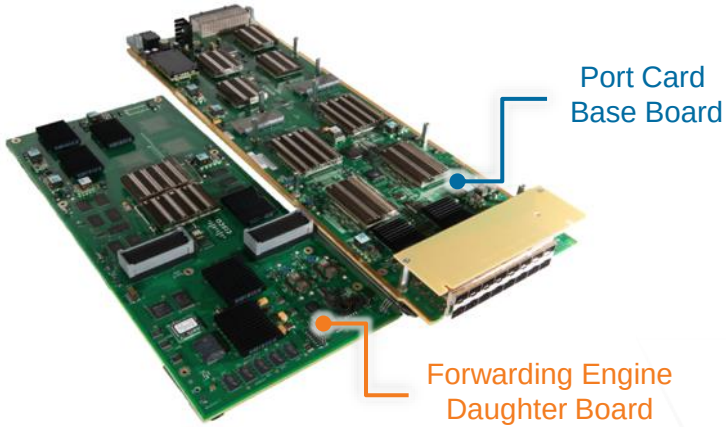

Two Versions	Standard (LE)	Large Tables
FIB Table v4/v6	256K/128K	**2M/1M**
NetFlow Table	512K	**1M**
Security ACL Table	64K	**256K**
Port Buffering	48MB / Port	48MB / Port

Port Speed & Type	Number of Ports
10/100/1000 Mb/s Copper	16 (GLC-T)
1 Gb/s Fiber	16 (SFP)
10 Gb/s Fiber	16 (SFP+)
40 Gb/s Fiber	4 (SFP-QSFP)

MacSec, FEX, VSS, LISP, SGT, 1588 Capable on Every Port

Catalyst 6880-X:

C6880-X-16P10G Module Architecture

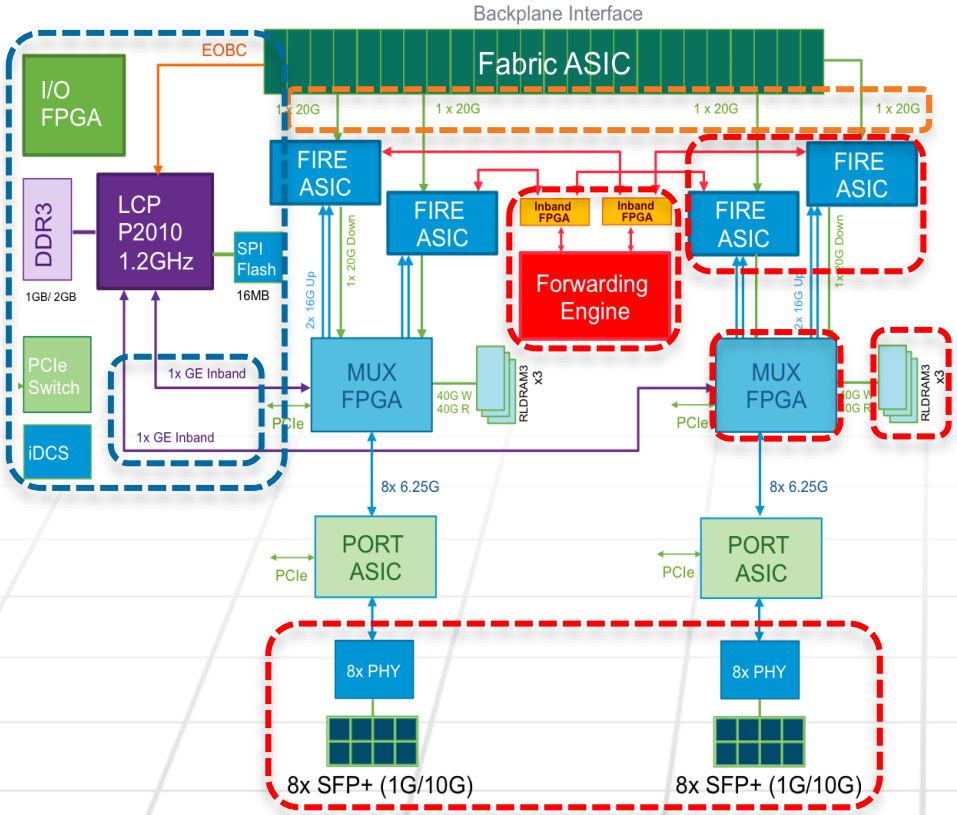

16P-10G Port Card Highlights:

- Same as Baseboard + 1.2Ghz LCP Complex
- 16 x SFP+ (Multi-Rate) Ethernet Ports
- 80Gbps to Switch Fabric (2 Modes)
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- Enhanced DFC4-E Forwarding Engine
- Improved 40Gbps Fabric/Replication ASIC
- New 40+Gbps Port Interface MUX FPGA
- Local 1GE Inband Links to LCP Complex
- RLDRAM3 Packet Buffers on MUX FPGA
- 192MB per MUX FPGA, 24/48MB per Port

Agenda

- Chassis & Power
- Supervisor Architectures
- Module Architectures
- **L2 Packet Forwarding**
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- Packet Walks

Layer 2 Switching

Layer 2 Table Structure

The PFC has an integrated CAM Table with 4096 rows * X pages = MAC address space

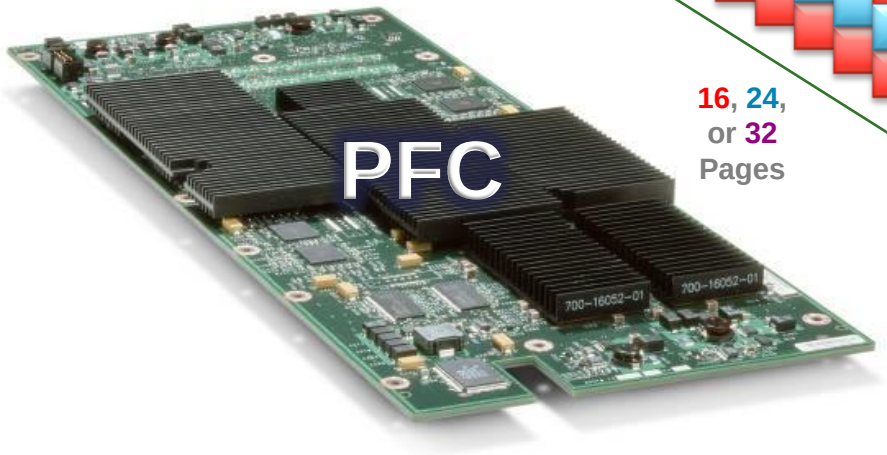

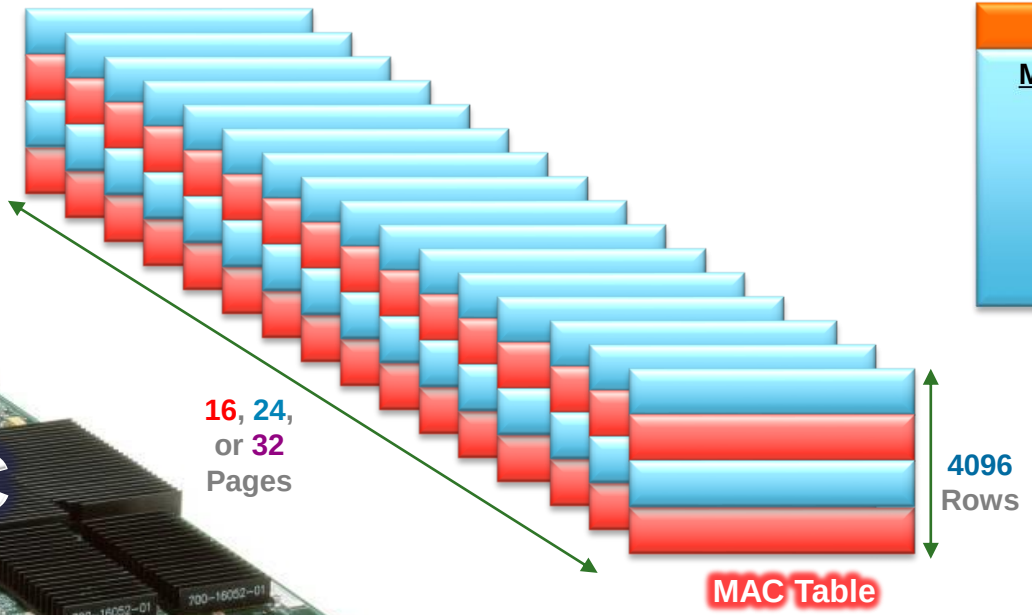

MAC Table	
MAC	Port
A	1
B	2
C	3
D	4
E	5
F	6

PFC3B/BXL = 16 pages (64K entries)
PFC3C/CXL = 24 pages (96K entries)
PFC4/XL = 32 pages (128K entries)

Layer 2 Switching

Layer 2 Forwarding Operation

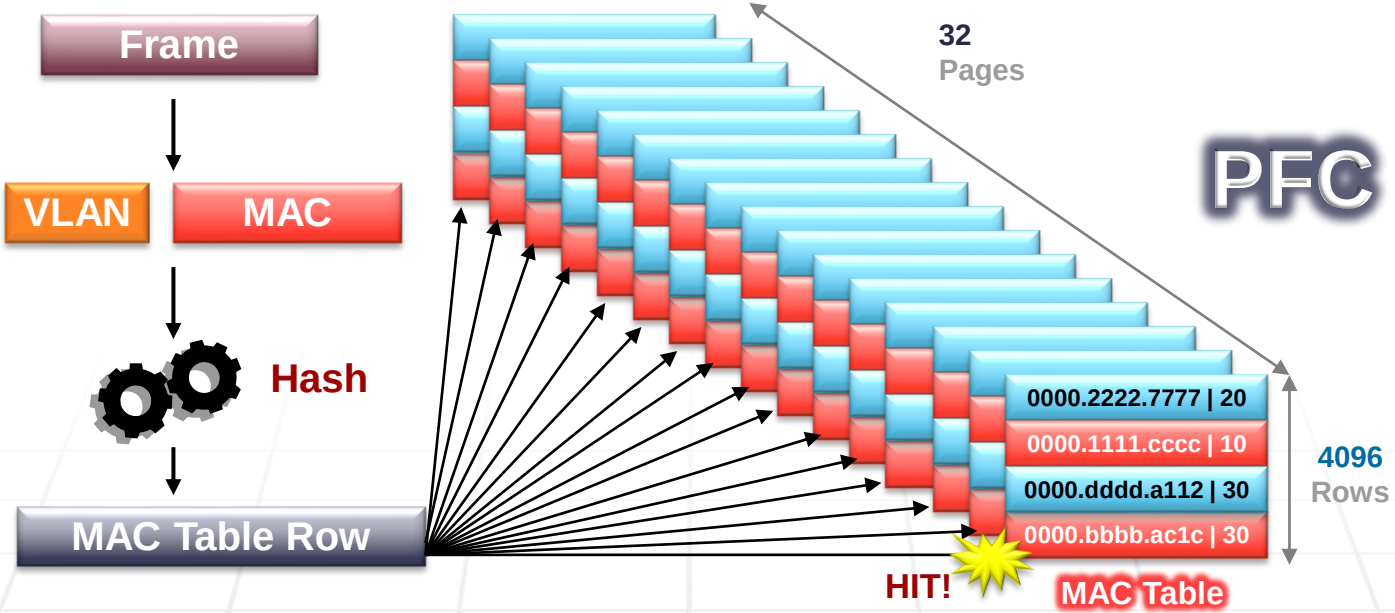

1. Hash result identifies the starting Page and Row in MAC table
2. Lookup key (VLAN + MAC) compared to contents of indexed line on each page, sequentially
3. Destination lookup: Match returns destination interface(s), Miss results in Flood
4. Source lookup: Match updates age of matching entry, Miss installs new entry in table

Catalyst 6500/6800 Internals

EtherChannel

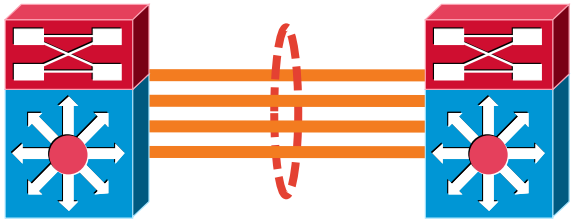

- Combines **Multiple** physical ports into **ONE** logical port
- Deterministic **Hash-based** Channel **Load-Balancing**
- **PFC3** hash algorithm supports **8 results (3 bits)**
- **PFC4** hash algorithm supports **256 results (8 bits)**
- Load Sharing is always **Per Flow** (Not Per Packet)
- EtherChannel can be configured for L2 and L3 interfaces

EtherChannel Load-Balancing

PFC3 Flow Distribution

EtherChannel Hash 3 bit result

E/Chan Bundle	Link1	Link2	Link3	Link4	Link5	Link6	Link7	Link8
2 Links	50%	50%	--	--	--	--	--	--
3 Links	37.5%	37.5%	25%	--	--	--	--	--

Even Distribution for Flows for ONLY those cases highlighted in RED

EtherChannel Load-Balancing

PFC4 Flow Distribution

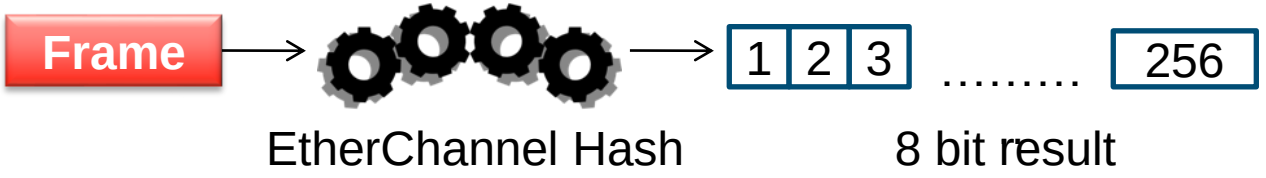

E/Chan Bundle	Link1	Link2	Link3	Link4	Link5	Link6	Link7	Link8
2 Links	50%	50%	--	--	--	--	--	--
3 Links	33.6%	33.2%	33.2%	--	--	--	--	--

Nearly Even Distribution for Flows in ODD & EVEN cases!

Agenda

- Chassis & Power
- Supervisor Architectures
- Module Architectures
- L2 Packet Forwarding
- **L3 Packet Forwarding**
- NetFlow & NDE
- Access Control Lists
- Packet Walks

IP Unicast Forwarding

High-Level L3 Forwarding on PFC

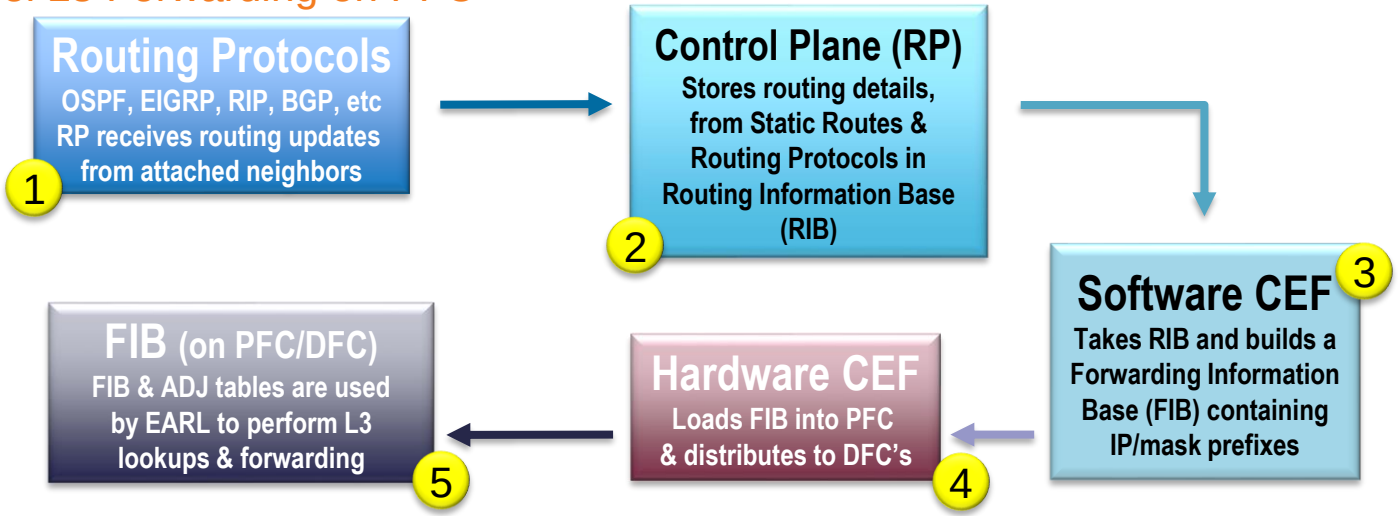

Hardware-based CEF Process

1. FIB lookup based on Destination prefix (longest-match)
2. FIB "Hit" returns an Adjacency pointer
3. Adjacency contains Rewrite (next-hop) information
4. ACL, QoS & NetFlow lookups occur IN PARALLEL (may effect final result)

IP Unicast Forwarding

FIB & Adjacency Overview

Each PFC/DFC stores a copy of the “FIB” & “Adjacency Table”…

The FIB contains:

CEF entries arranged from *MOST* to *LEAST* specific (based on /Mask)

Overall FIB hardware shared by:

- IPv4 Unicast
- IPv4 Multicast
- IPv6 Unicast
- IPv6 Multicast
- MPLS

The Adjacency Table:

- L2 “Re-Write” information and / or pointers for replication
- Hardware Adjacency table also shared among protocols

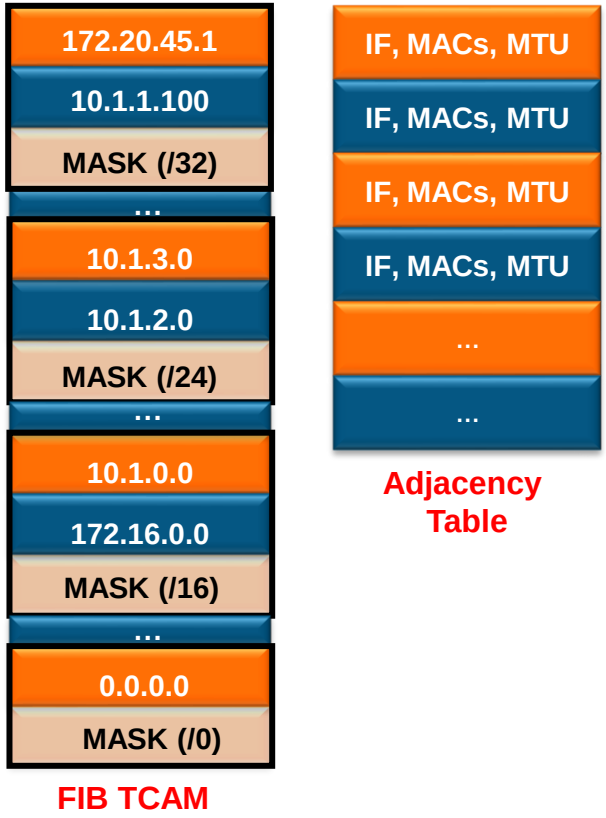

IP Unicast Forwarding

Layer 3 Forwarding on PFC

Lets assume a lookup needs to be performed for a packet with a destination of **10.1.5.2 /24**, then the following would occur…

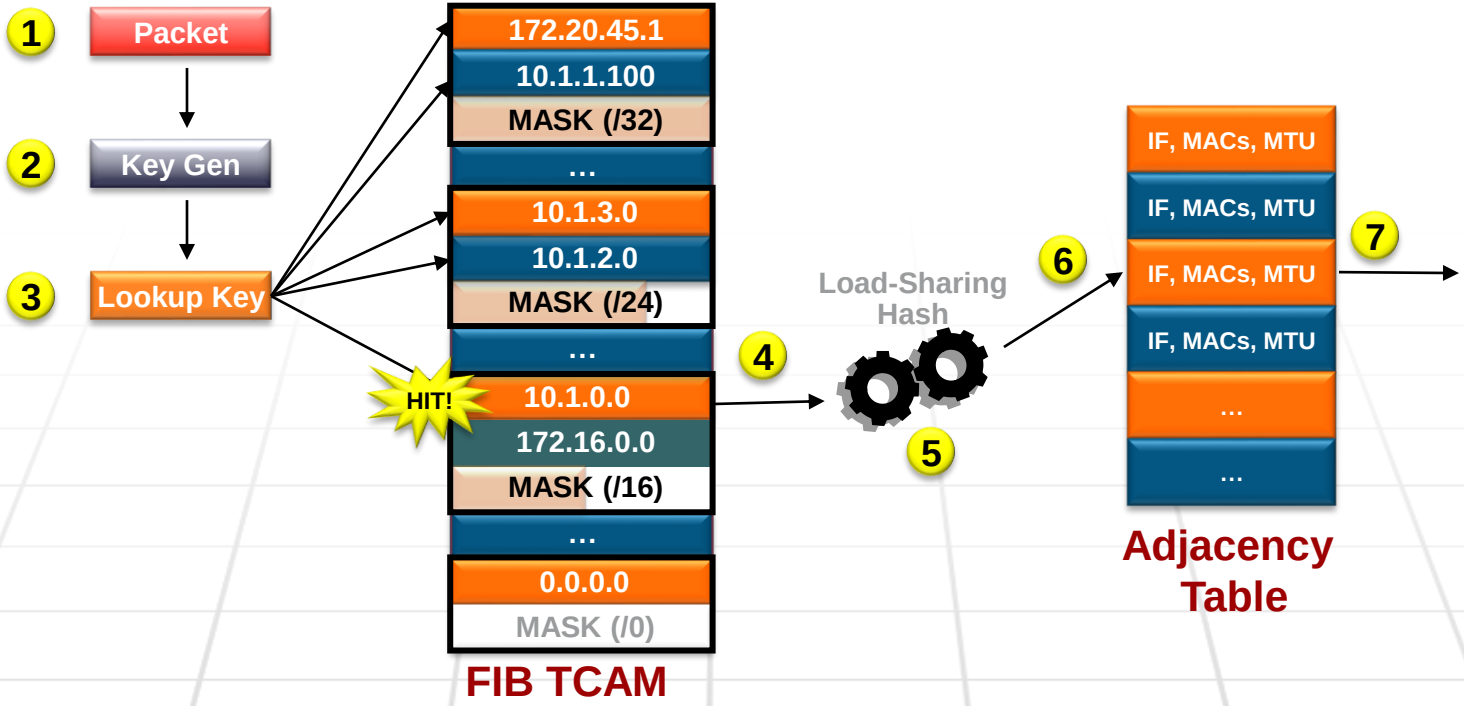

Supervisor FIB TCAM Resources

Defaults and Changes

For Your Reference

IPv6 Unicast & IPv4 Multicast require 2 entries

- MPLS and IPv4 Unicast only 1 entry

XL PFCs = 1M entries

Non-XL PFCs = 256K entries

Default TCAM allocation shown below

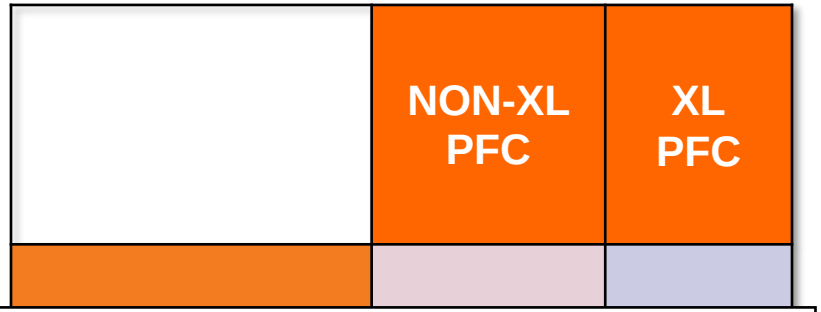

SUP2TXL Example

```
SUP2T#sh platform cef maximum-routes
FIB TCAM maximum routes :
=========================
Current :-
-------
 IPv4 + MPLS      - 512k (default)
 IPv6 + IP Multicast - 256k (default)
```

Changing default (requires Reboot!)

```
SUP2T(config)#platform cef maximum-routes
?
 ip                  number of ip routes
 ip-multicast        number of multicast
routes
 ipv6                number of ipv6 routes
 mpls                number of MPLS labels
```

Displaying IPv4 Forwarding Summary

For Your Reference

```
SUP2T#show platform hardware capacity forwarding
...
L3 Forwarding Resources
              FIB TCAM usage:                      Total          Used       %Used
                   72 bits (IPv4, MPLS, EoM)      196608            28          1%
                  144 bits (IP mcast, IPv6)        32768             7          1%

                  detail:      Protocol                           Used       %Used
                               IPv4                                 28          1%
                               MPLS                                  0          0%
                               EoM                                   0          0%

                               IPv6                                  1          1%
                               IPv4 mcast                            3          1%
                               IPv6 mcast                            3          1%

            Adjacency usage:                       Total          Used       %Used
                                                 1048576           171          1%
```

Displaying Hardware IPv4 Prefix Entries

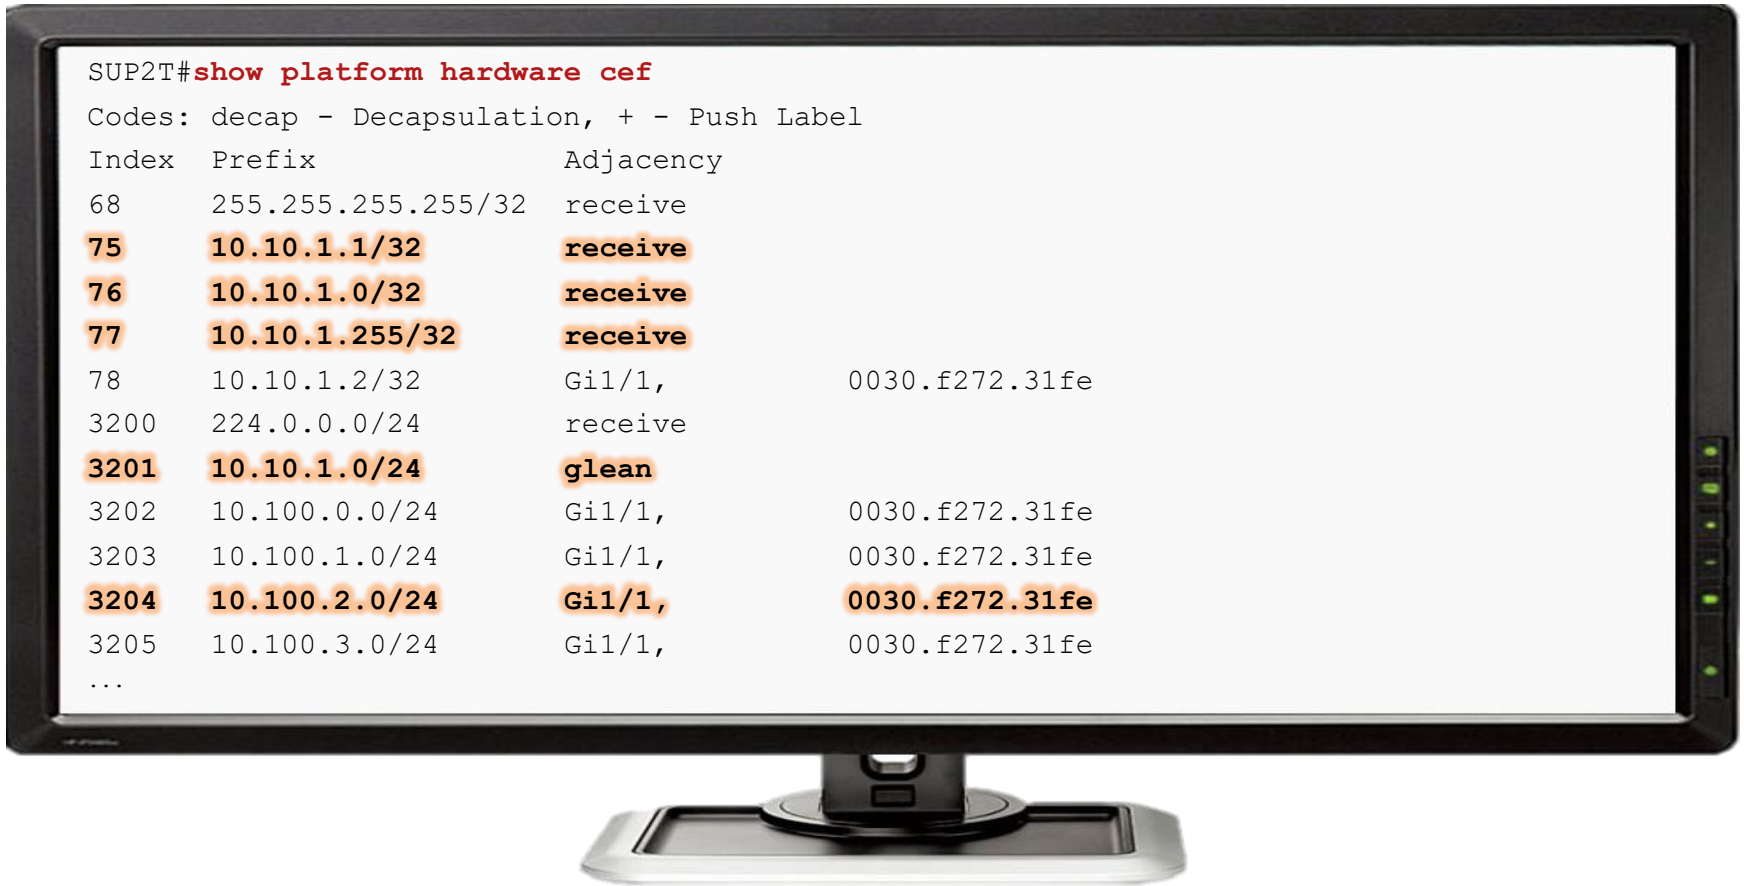

```
SUP2T#show platform hardware cef
Codes: decap - Decapsulation, + - Push Label
Index   Prefix              Adjacency
68      255.255.255.255/32  receive
75      10.10.1.1/32        receive
76      10.10.1.0/32        receive
77      10.10.1.255/32      receive
78      10.10.1.2/32        Gi1/1,          0030.f272.31fe
3200    224.0.0.0/24        receive
3201    10.10.1.0/24        glean
3202    10.100.0.0/24       Gi1/1,          0030.f272.31fe
3203    10.100.1.0/24       Gi1/1,          0030.f272.31fe
3204    10.100.2.0/24       Gi1/1,          0030.f272.31fe
3205    10.100.3.0/24       Gi1/1,          0030.f272.31fe
...
```

Finding the Longest-Match Prefix Entry

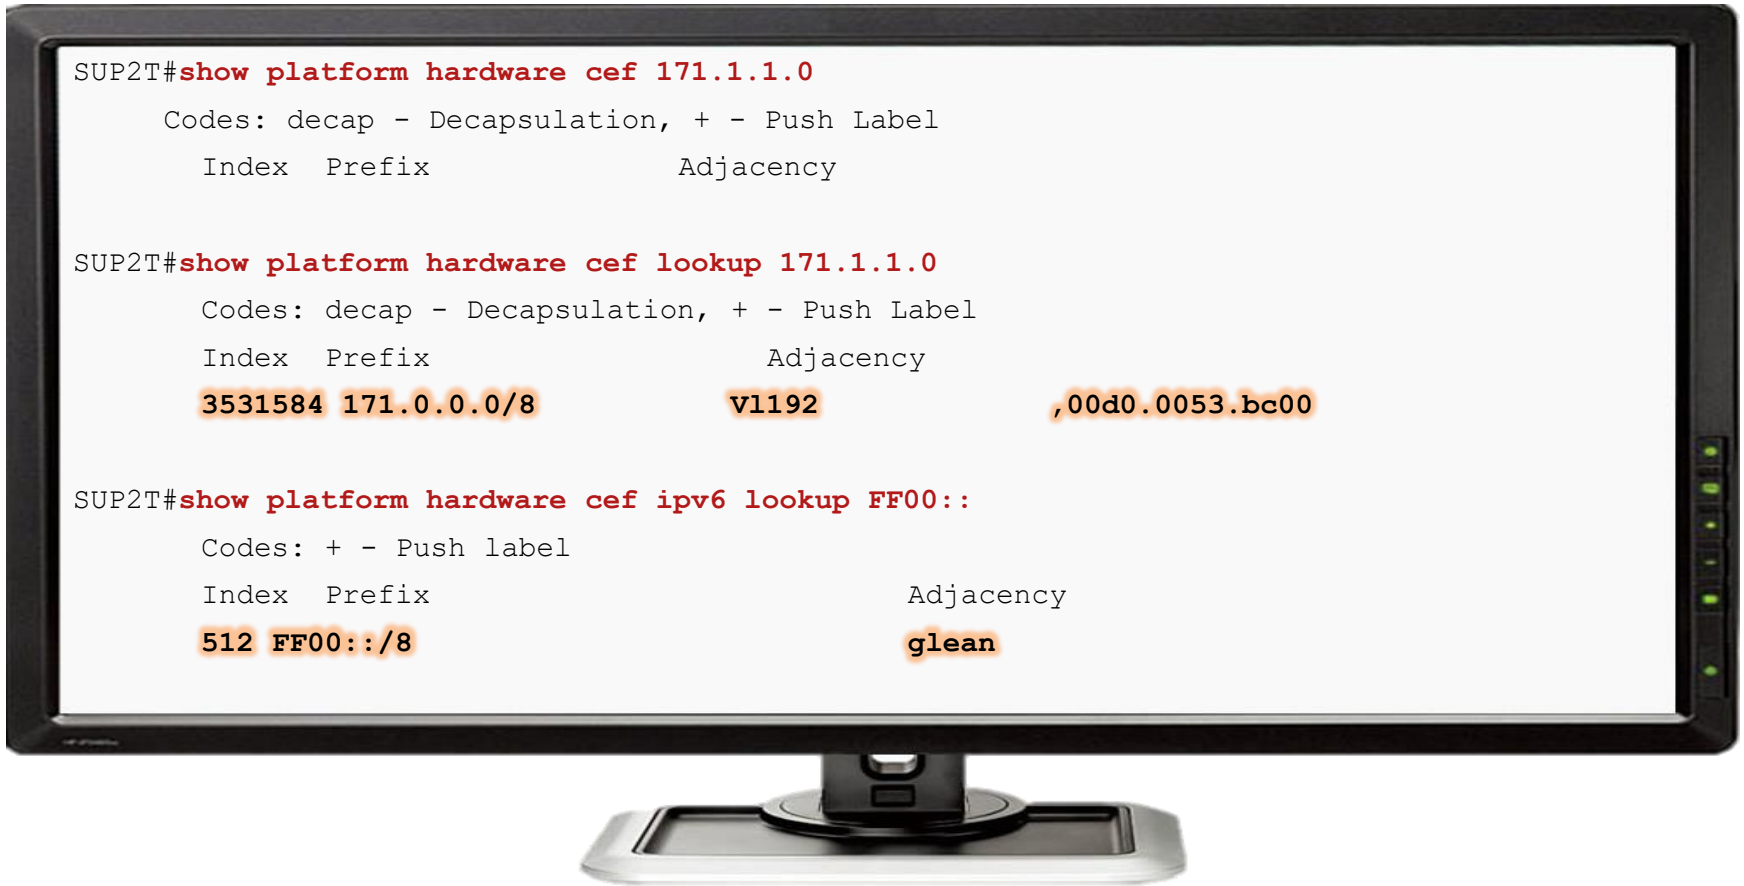

```
SUP2T#show platform hardware cef 171.1.1.0
     Codes: decap - Decapsulation, + - Push Label
       Index  Prefix              Adjacency

SUP2T#show platform hardware cef lookup 171.1.1.0
       Codes: decap - Decapsulation, + - Push Label
       Index  Prefix                   Adjacency
       3531584 171.0.0.0/8             Vl192              ,00d0.0053.bc00

SUP2T#show platform hardware cef ipv6 lookup FF00::
       Codes: + - Push label
       Index  Prefix                                Adjacency
       512 FF00::/8                                 glean
```

Agenda

- Chassis & Power
- Supervisor Architectures
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- **NetFlow & NDE**
- Access Control Lists
- Packet Walks

Hardware NetFlow

NetFlow Export Process

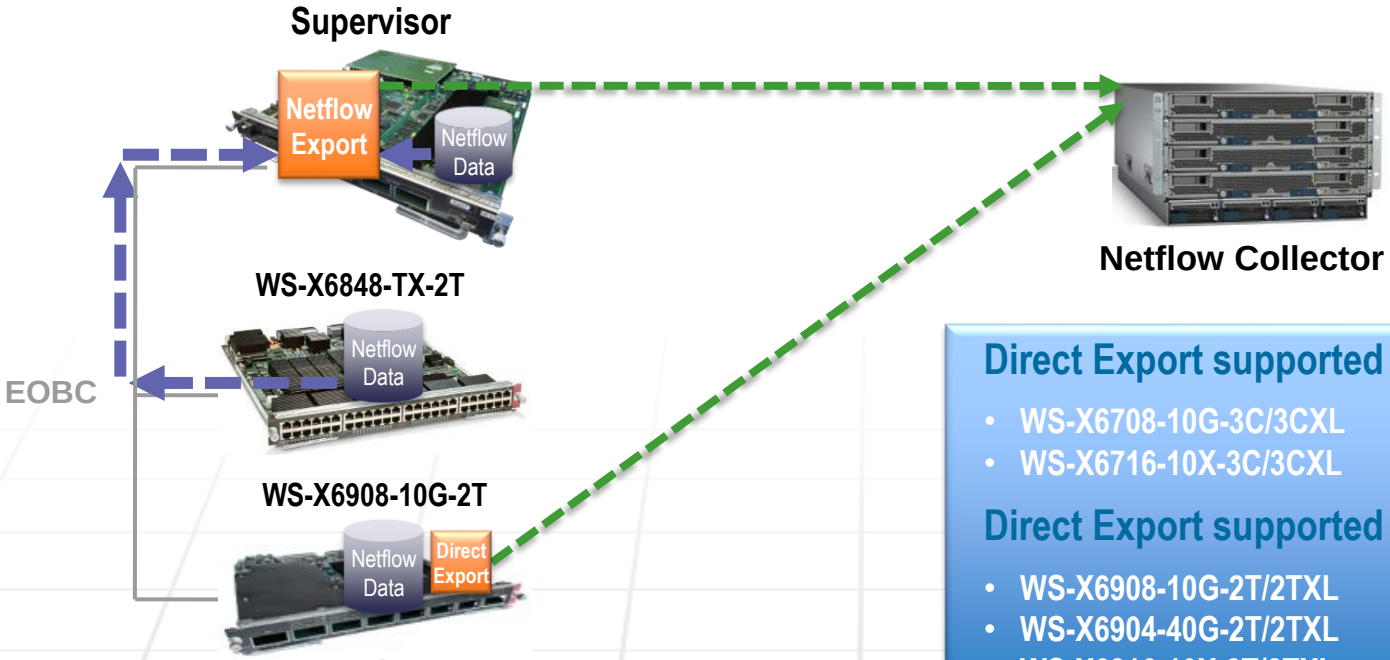

Direct Export supported with Sup720:

- WS-X6708-10G-3C/3CXL
- WS-X6716-10X-3C/3CXL

Direct Export supported with Sup2T:

- WS-X6908-10G-2T/2TXL
- WS-X6904-40G-2T/2TXL
- WS-X6816-10X-2T/2TXL
- WS-X6716-10X with DFC4-E/EXL

Hardware NetFlow

Supervisor 2T Enhancements

Optimal CPU Utilization with Yielding NDE & Direct Export from a Line Card

Increased customization and scale by selecting the fields to Match and Collect for both IPv4 and IPv6

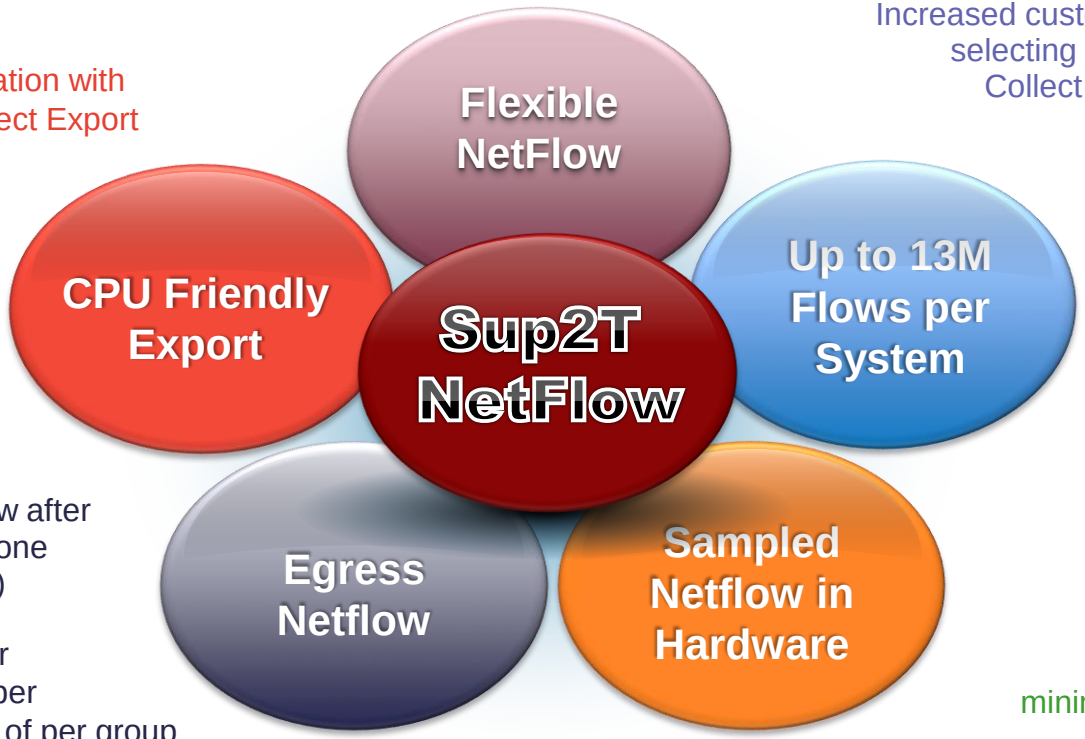

Bigger tables mean MORE entries per system, giving you better visibility into your network (up to 13 million NF entries with a 13 slot chassis)

Allow to use Netflow after ingress lookup is done (NetFlow on CoPP)

Allow to account for IP Multicast traffic per destination instead of per group

Optimize the Netflow Tables utilization and minimize load on Analyzers

Hardware NetFlow

TCAM Lookup on PFC4

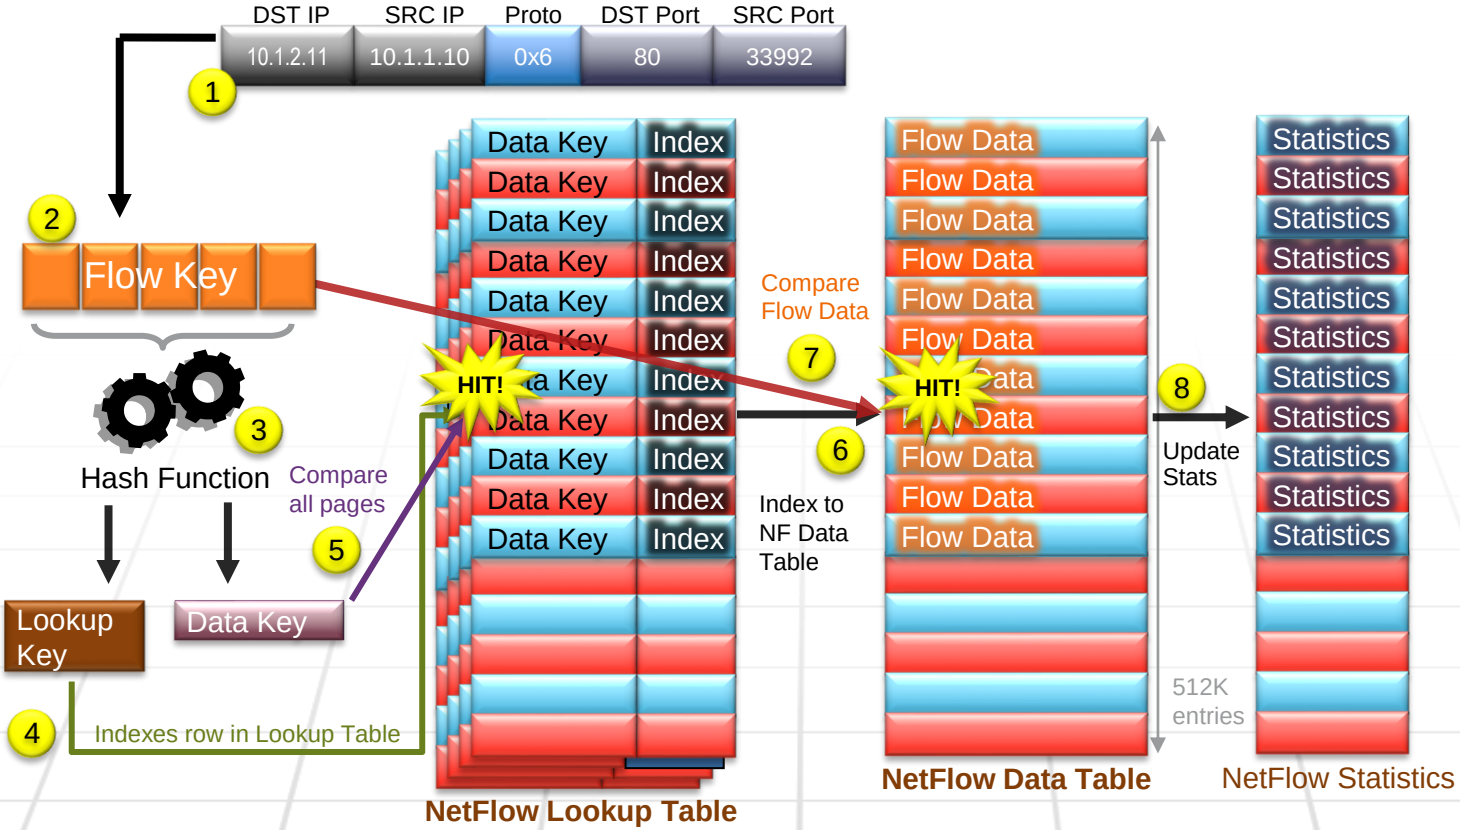

Hardware NetFlow

Sup2T “CPU Friendly” Netflow Export

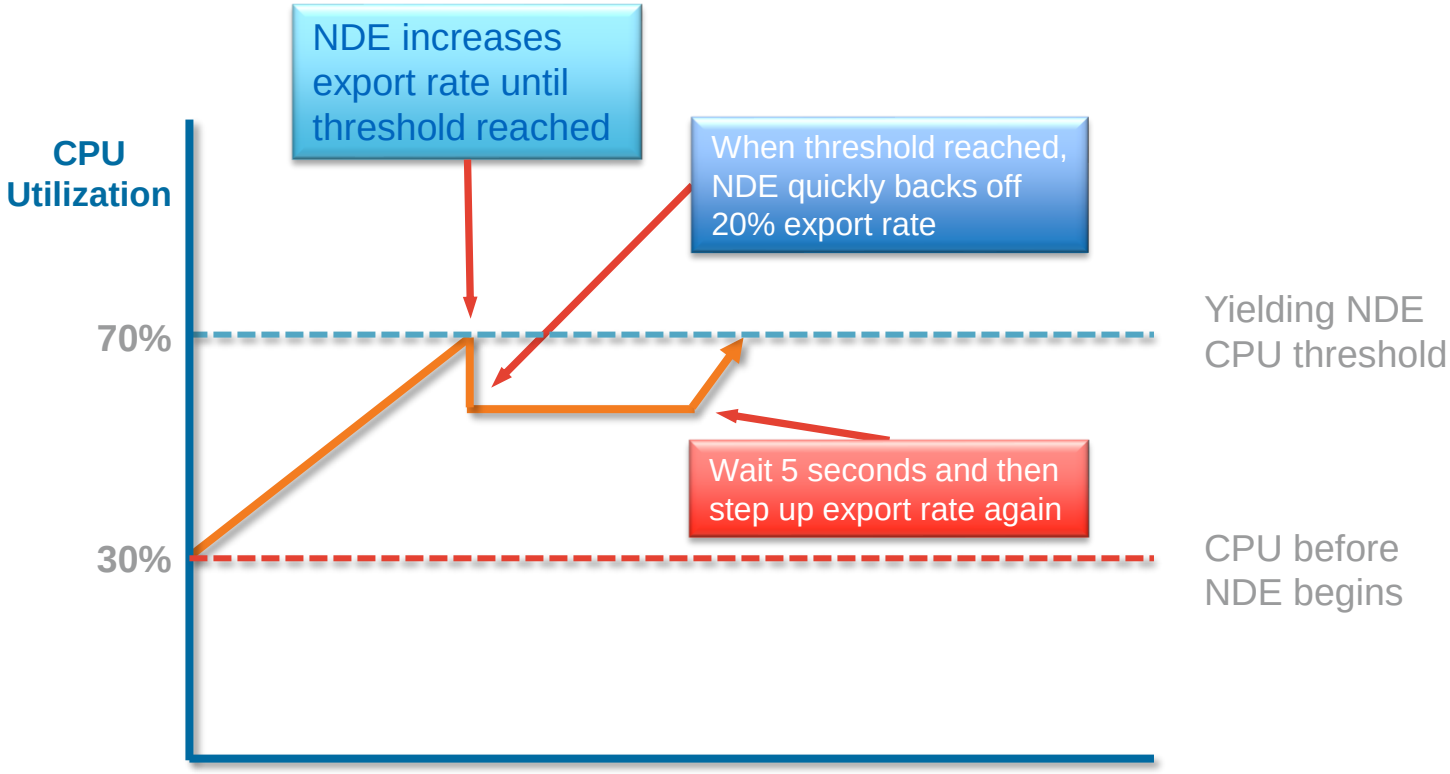

Displaying NetFlow Utilization

For Your Reference

```
SUP2T#show platform hardware capacity netflow
Netflow resources:
  Netflow table size: 515032 entries total

  Netflow table usage: Module/Instance     Input flows     Output flows
                                     3             10%              10%
                                     7             25%              25%
```

Agenda

- Chassis & Power
- Supervisor Architectures
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- **Access Control Lists**
- Packet Walks

Access Control Lists

Hardware Support

1

Create the ACL or classification policy using CLI or Network Management System

```
ip access-list extended Internet
 permit ip any host 10.2.2.4
 permit ip any host 10.5.2.33
 permit ip any host 10.11.0.0
 permit ip any host 10.4.0.0
```

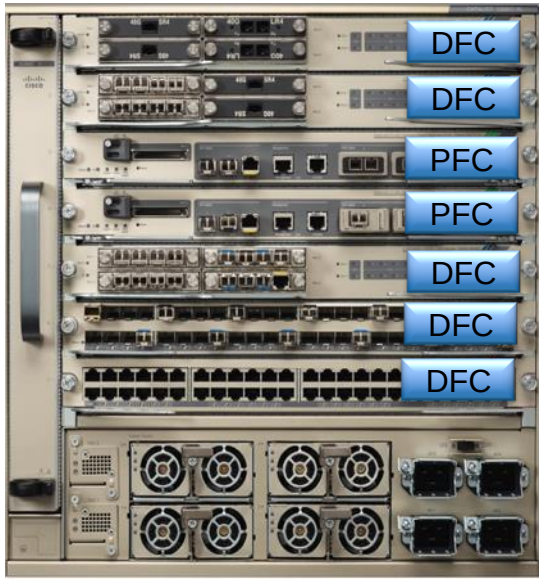

Hardware Support
Policy Feature Card (PFC)
Distributed Forwarding Card (DFC)

Router ACLs
VLAN ACLs
Port Based ACLs
Role Based ACLs

2

Hardware- Assist ACL Features 3

Netflow
WCCP
Reflexive ACLs
NAT & PAT
Cisco Trust Sec

Access Control Lists

Three Forms of Security ACLs

The PFC3/PFC4 supports three forms of Security ACLs: the RACL, VACL and PACL…

Router ACL (RACL)	VLAN ACL (VACL)	Port ACL (PACL)
Used to permit or deny the movement of traffic between Layer 3 Subnets	Used to permit or deny the movement of traffic between Layer 3 Subnets & VLANs or within a VLAN	Used to permit or deny the movement of traffic between Layer 3 Subnets & VLANs or within a VLAN
Applied as an input or output policy to a Layer 3 interface	Applied as a policy to a VLAN - is inherently applied to both inbound and outbound traffic	Applied as a policy to a Layer 2 Switch port interface - is applied for inbound traffic only

Access Control Lists

ACL Order of Processing

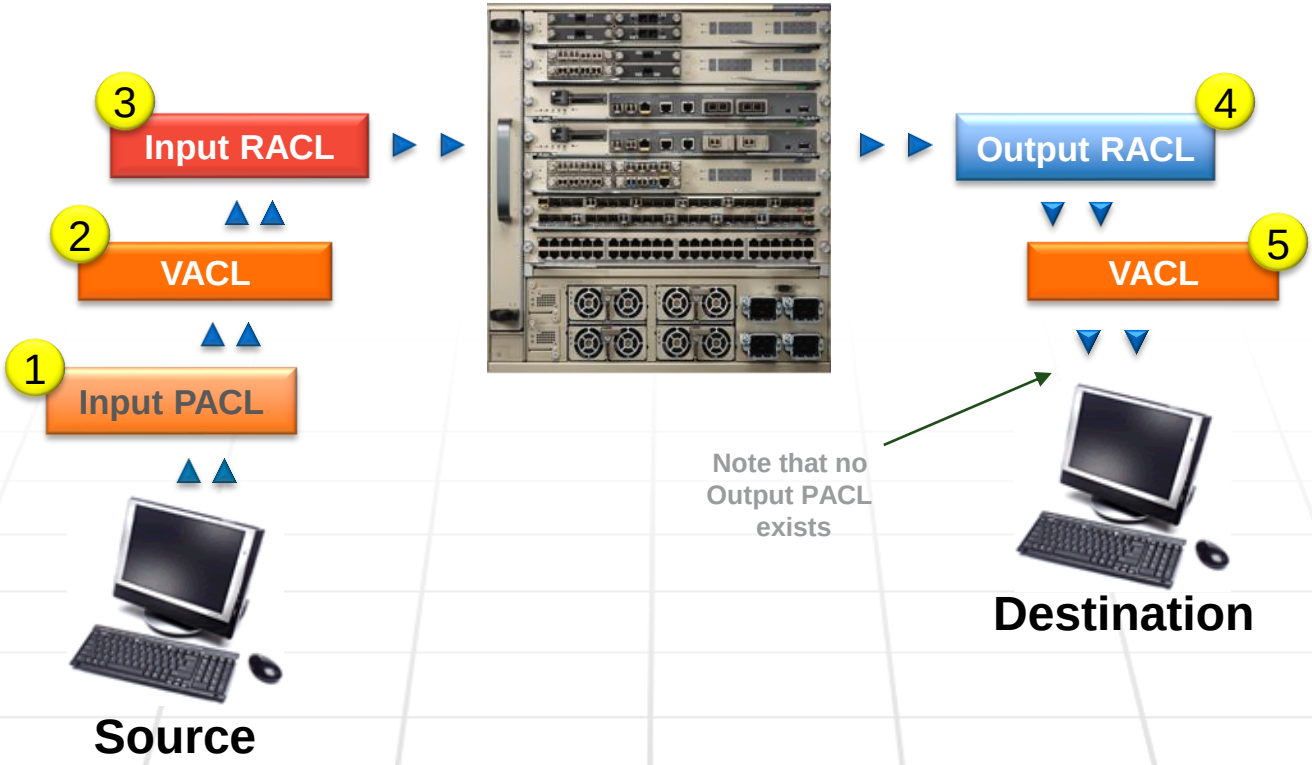

Access Control Lists

PFC4 TCAM Lookup

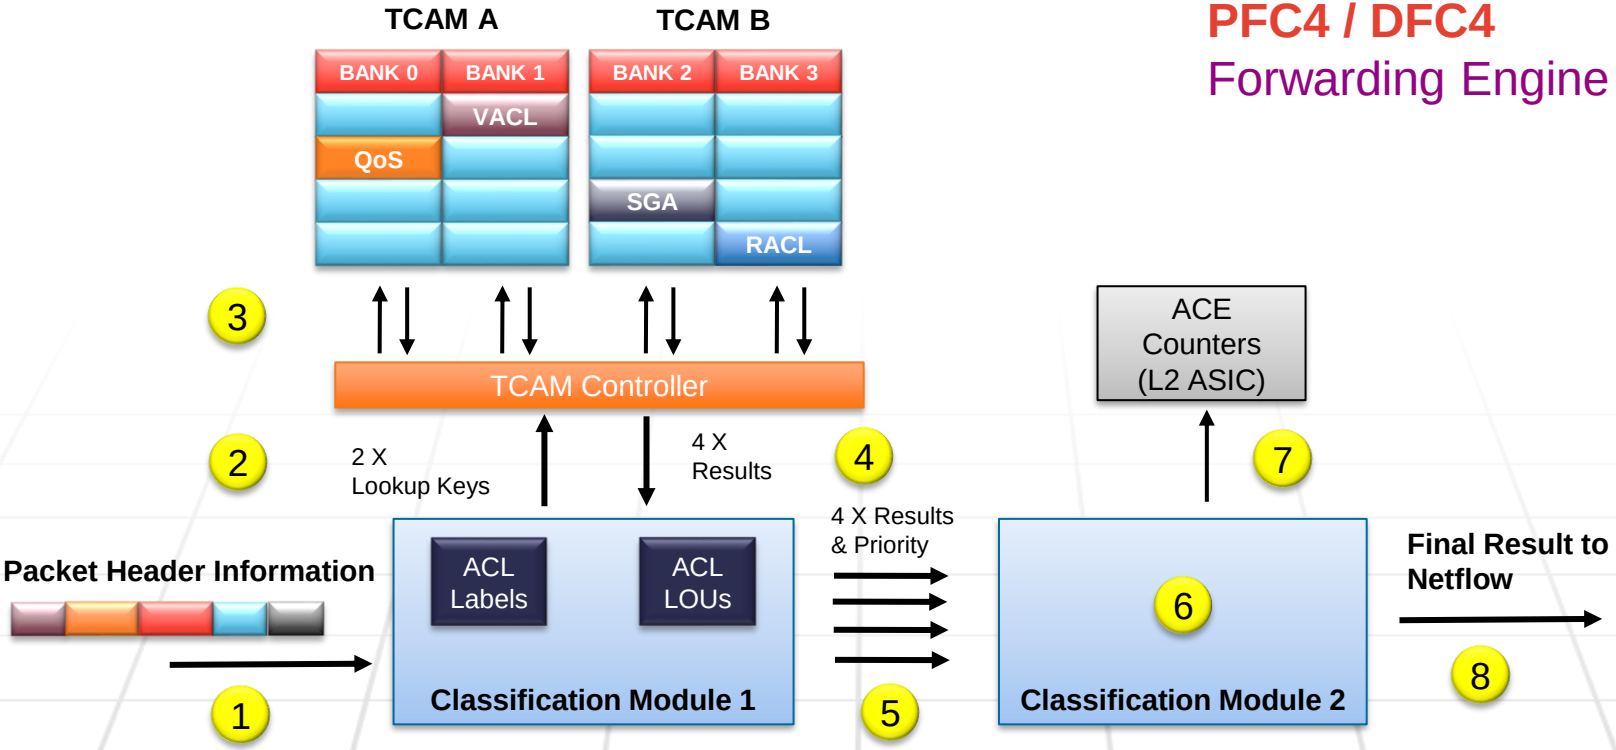

PFC4 / DFC4
Forwarding Engine

Agenda

- Chassis & Power
- Supervisor Architectures
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- **Packet Walks**
  - **6700 Series (CEF720)**
  - **6800 Series (dCEF720)**
  - **6900 Series (dCEF2T)**
  - C6880-X Series (dCEF2T**)**

6700 to 6700 - Centralized Forwarding (CEF720)

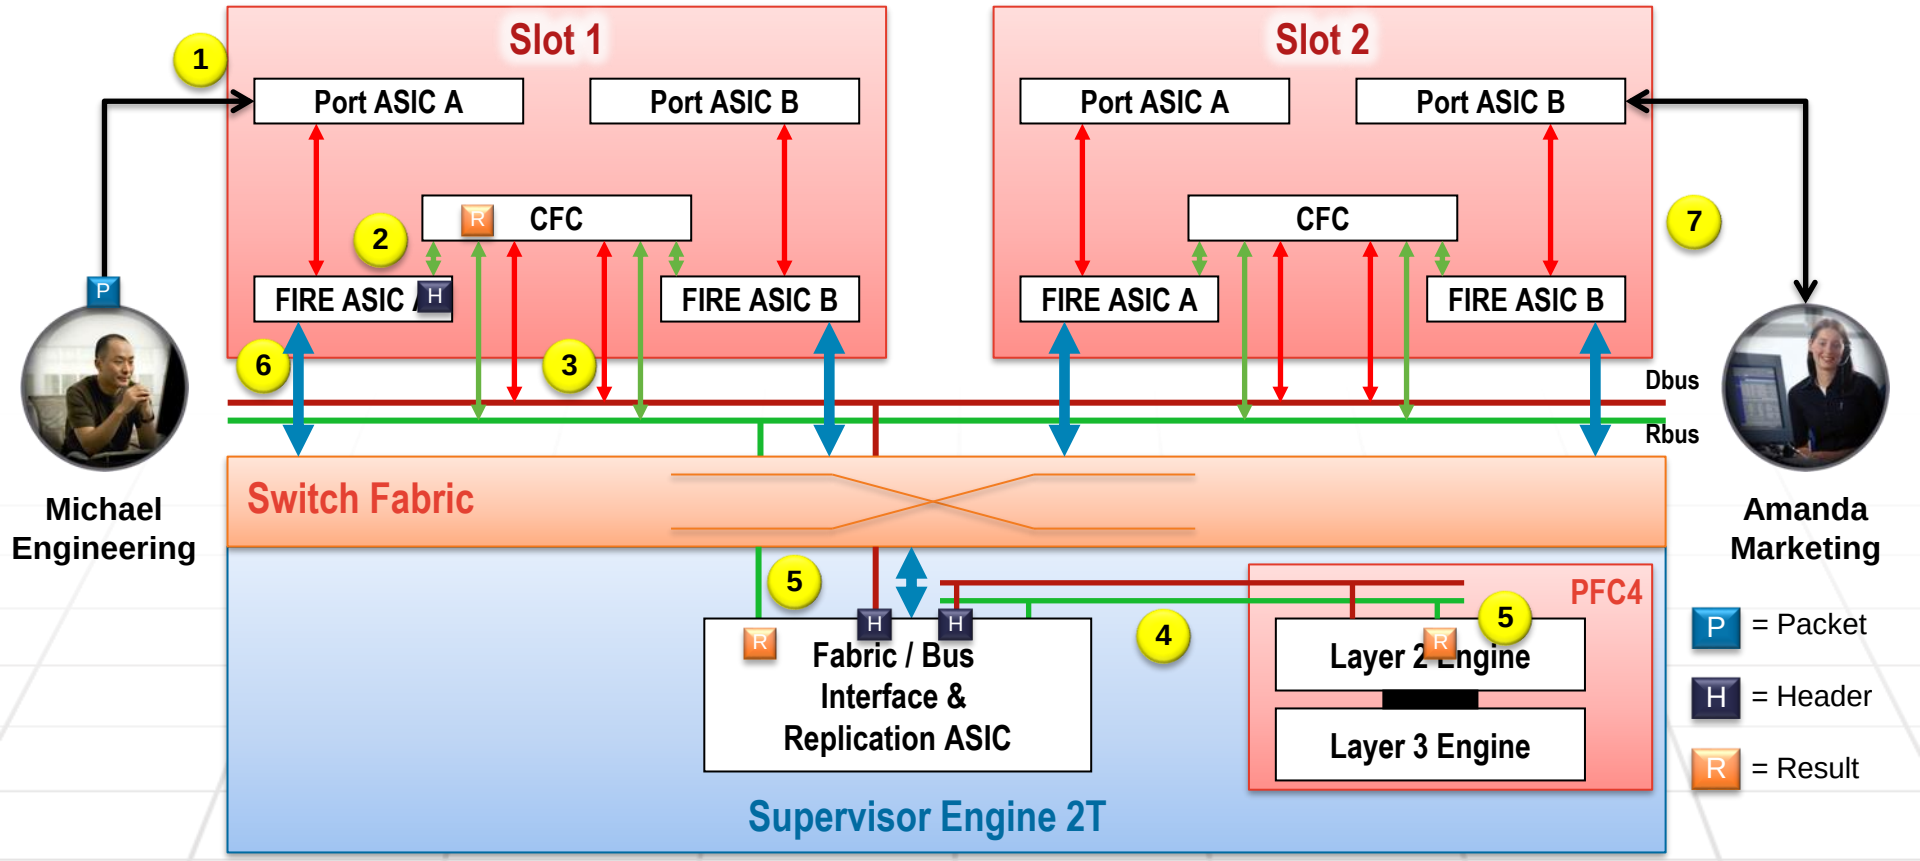

6800 to 6800 - Distributed Forwarding (dCEF20)

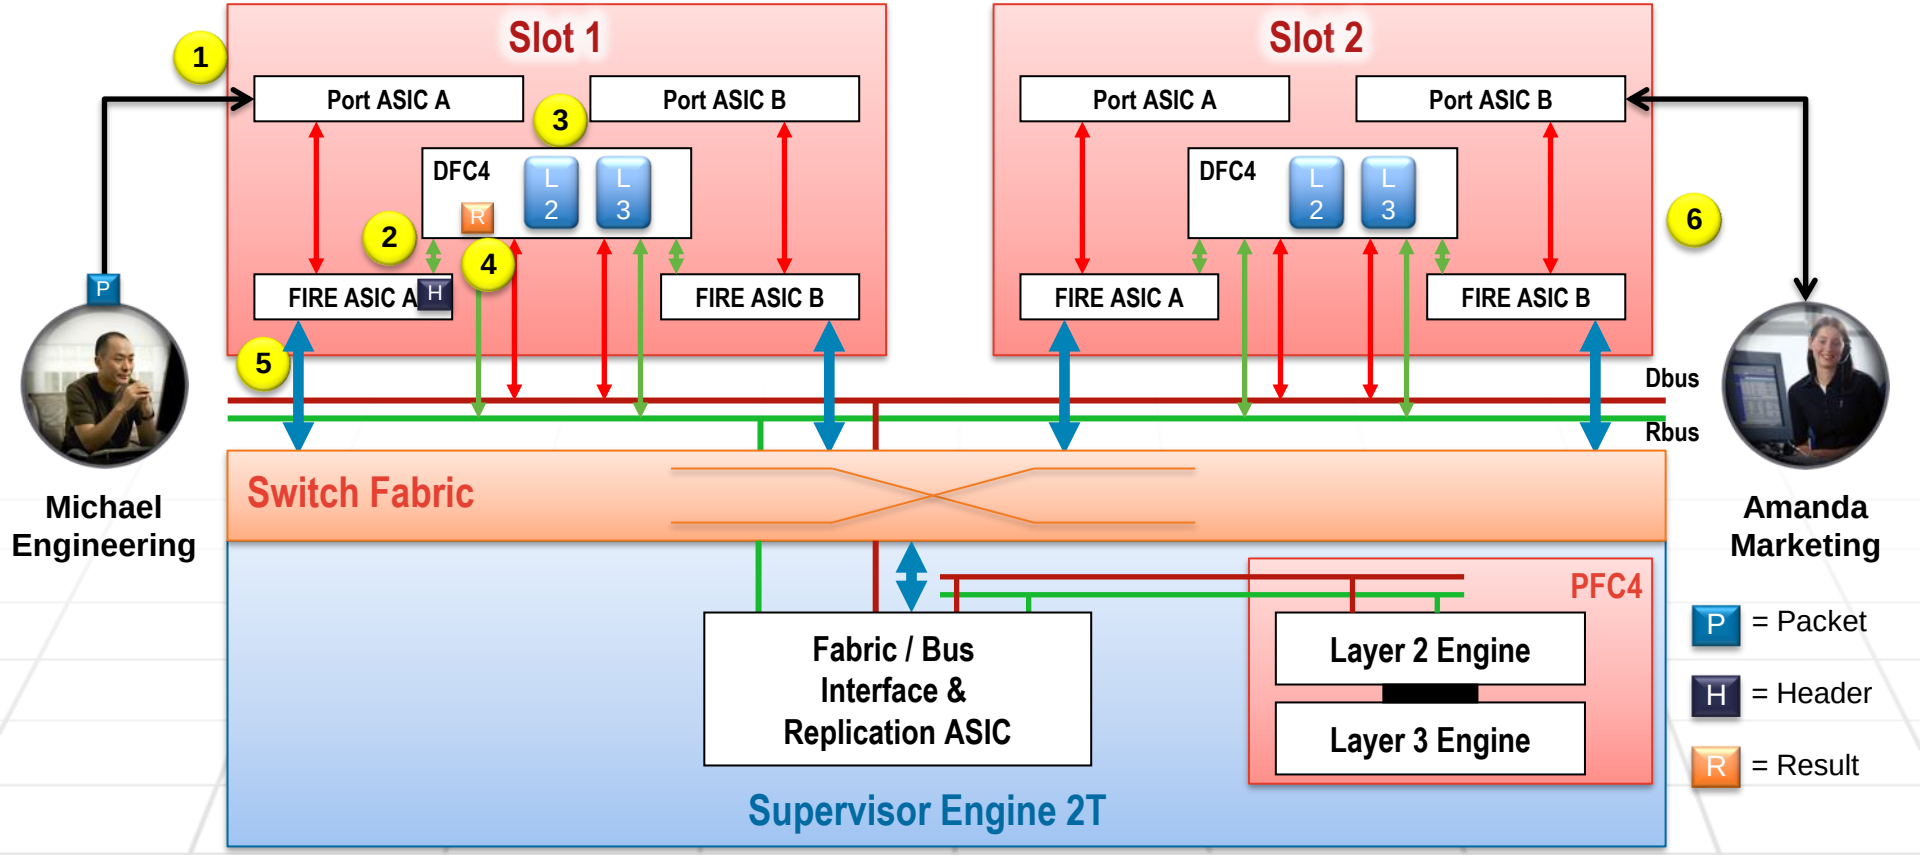

6900 to 6900 - Distributed Forwarding (dCEF2T)

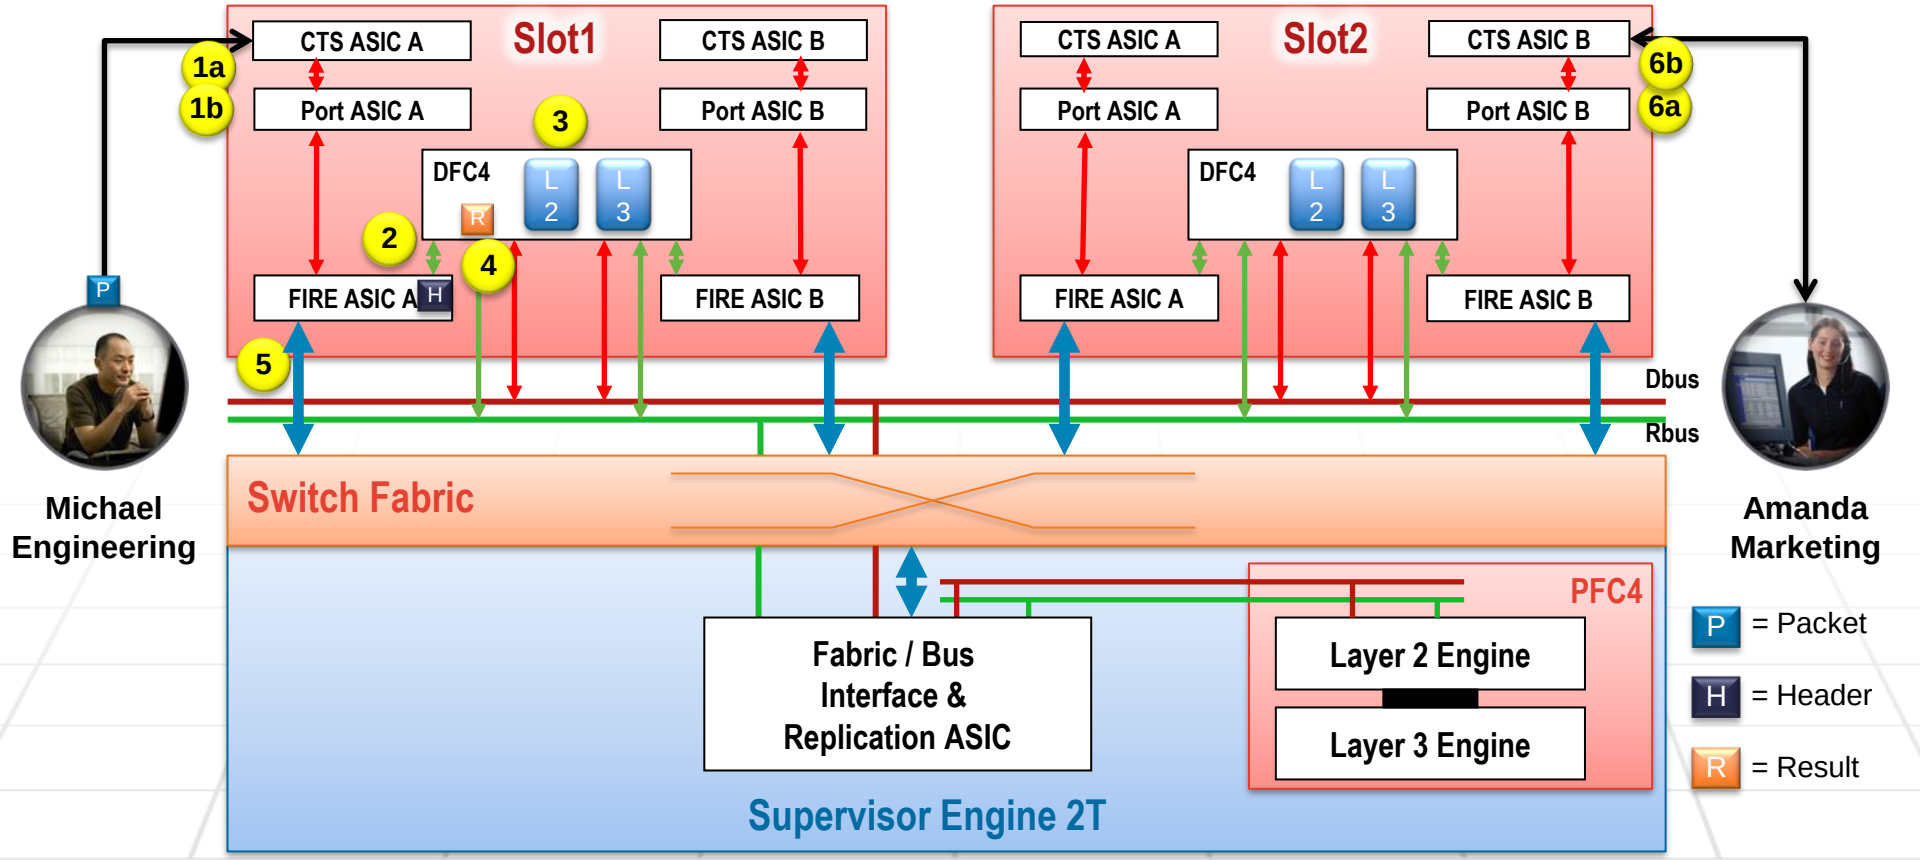

Agenda

- Chassis & Power
- Supervisor Architectures
- Module Architectures
- L2 Packet Forwarding
- L3 Packet Forwarding
- NetFlow & NDE
- Access Control Lists
- **Packet Walks**
  - 6700 Series (CEF720)
  - 6800 Series (dCEF720)
  - 6900 Series (dCEF2T)
  - **C6880-X Series (dCEF2T)**

Catalyst 6880-X: Packet Walks

Remote Forwarding (Ingress)

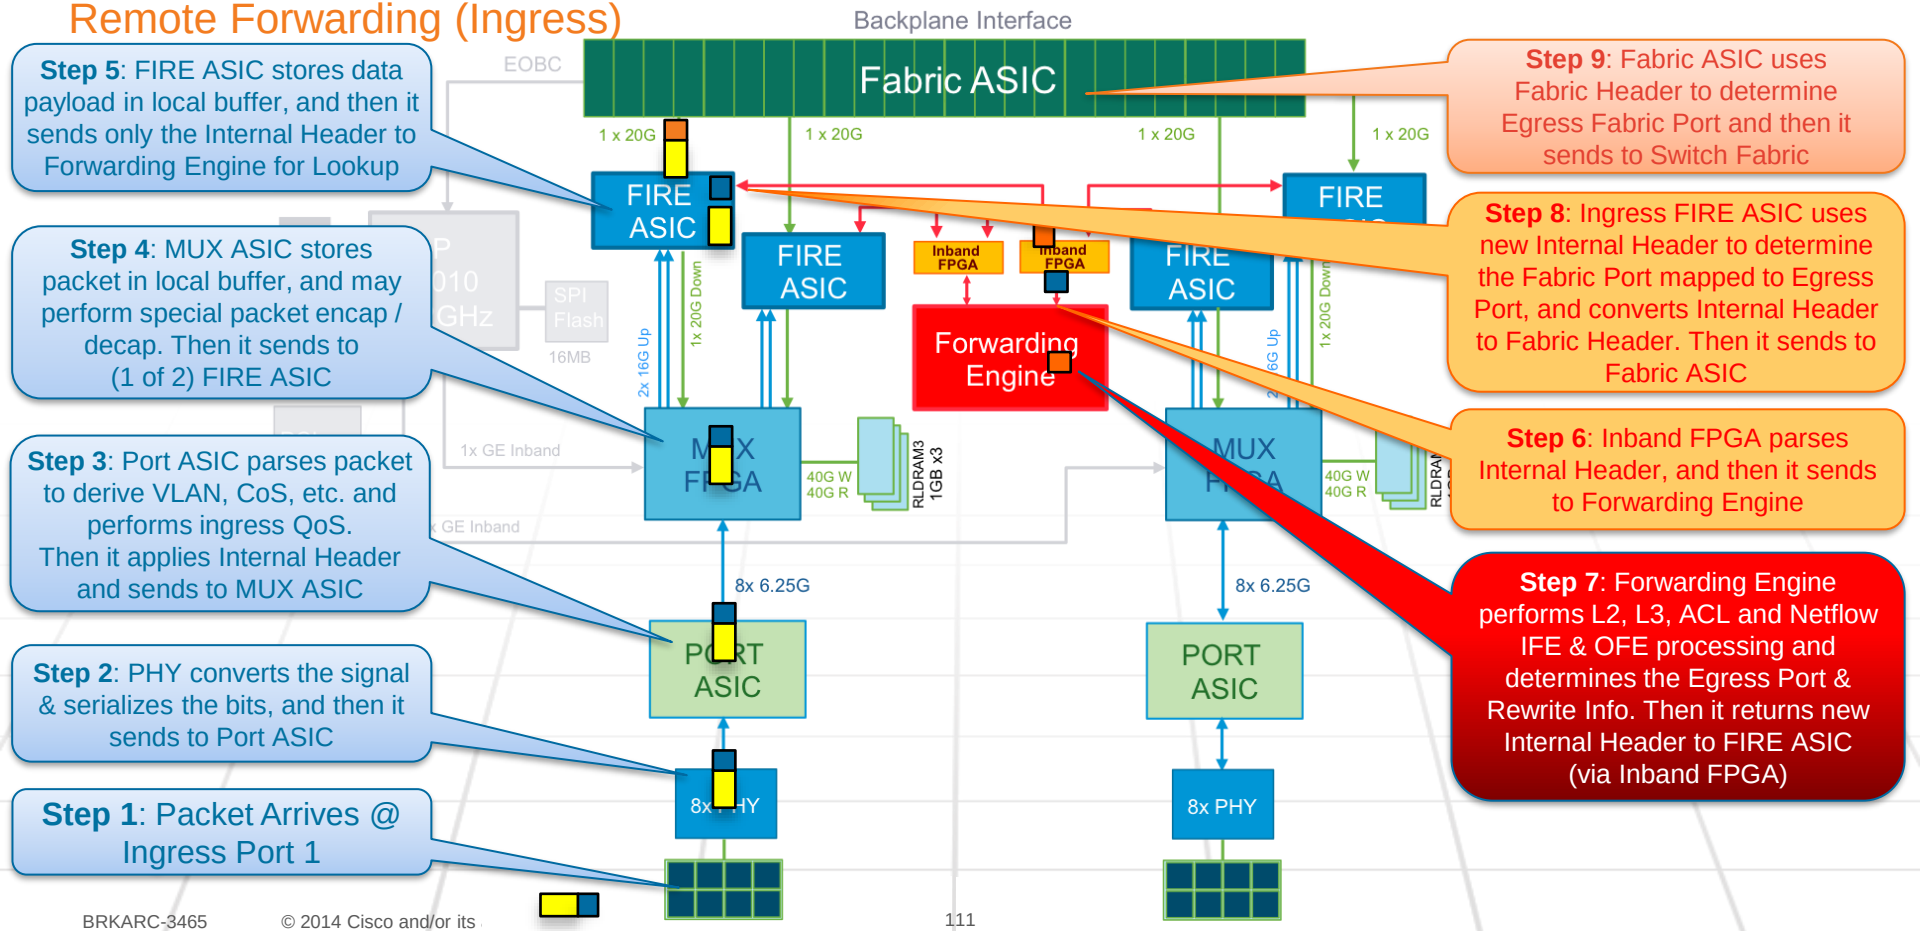

Step 5: FIRE ASIC stores data payload in local buffer, and then it sends only the Internal Header to Forwarding Engine for Lookup

Step 4: MUX ASIC stores packet in local buffer, and may perform special packet encap / decap. Then it sends to (1 of 2) FIRE ASIC

Step 3: Port ASIC parses packet to derive VLAN, CoS, etc. and performs ingress QoS. Then it applies Internal Header and sends to MUX ASIC

Step 2: PHY converts the signal & serializes the bits, and then it sends to Port ASIC

Step 1: Packet Arrives @ Ingress Port 1

Step 9: Fabric ASIC uses Fabric Header to determine Egress Fabric Port and then it sends to Switch Fabric

Step 8: Ingress FIRE ASIC uses new Internal Header to determine the Fabric Port mapped to Egress Port, and converts Internal Header to Fabric Header. Then it sends to Fabric ASIC

Step 6: Inband FPGA parses Internal Header, and then it sends to Forwarding Engine

Step 7: Forwarding Engine performs L2, L3, ACL and Netflow IFE & OFE processing and determines the Egress Port & Rewrite Info. Then it returns new Internal Header to FIRE ASIC (via Inband FPGA)

Catalyst 6880-X: Packet Walks

Remote Forwarding (Egress)

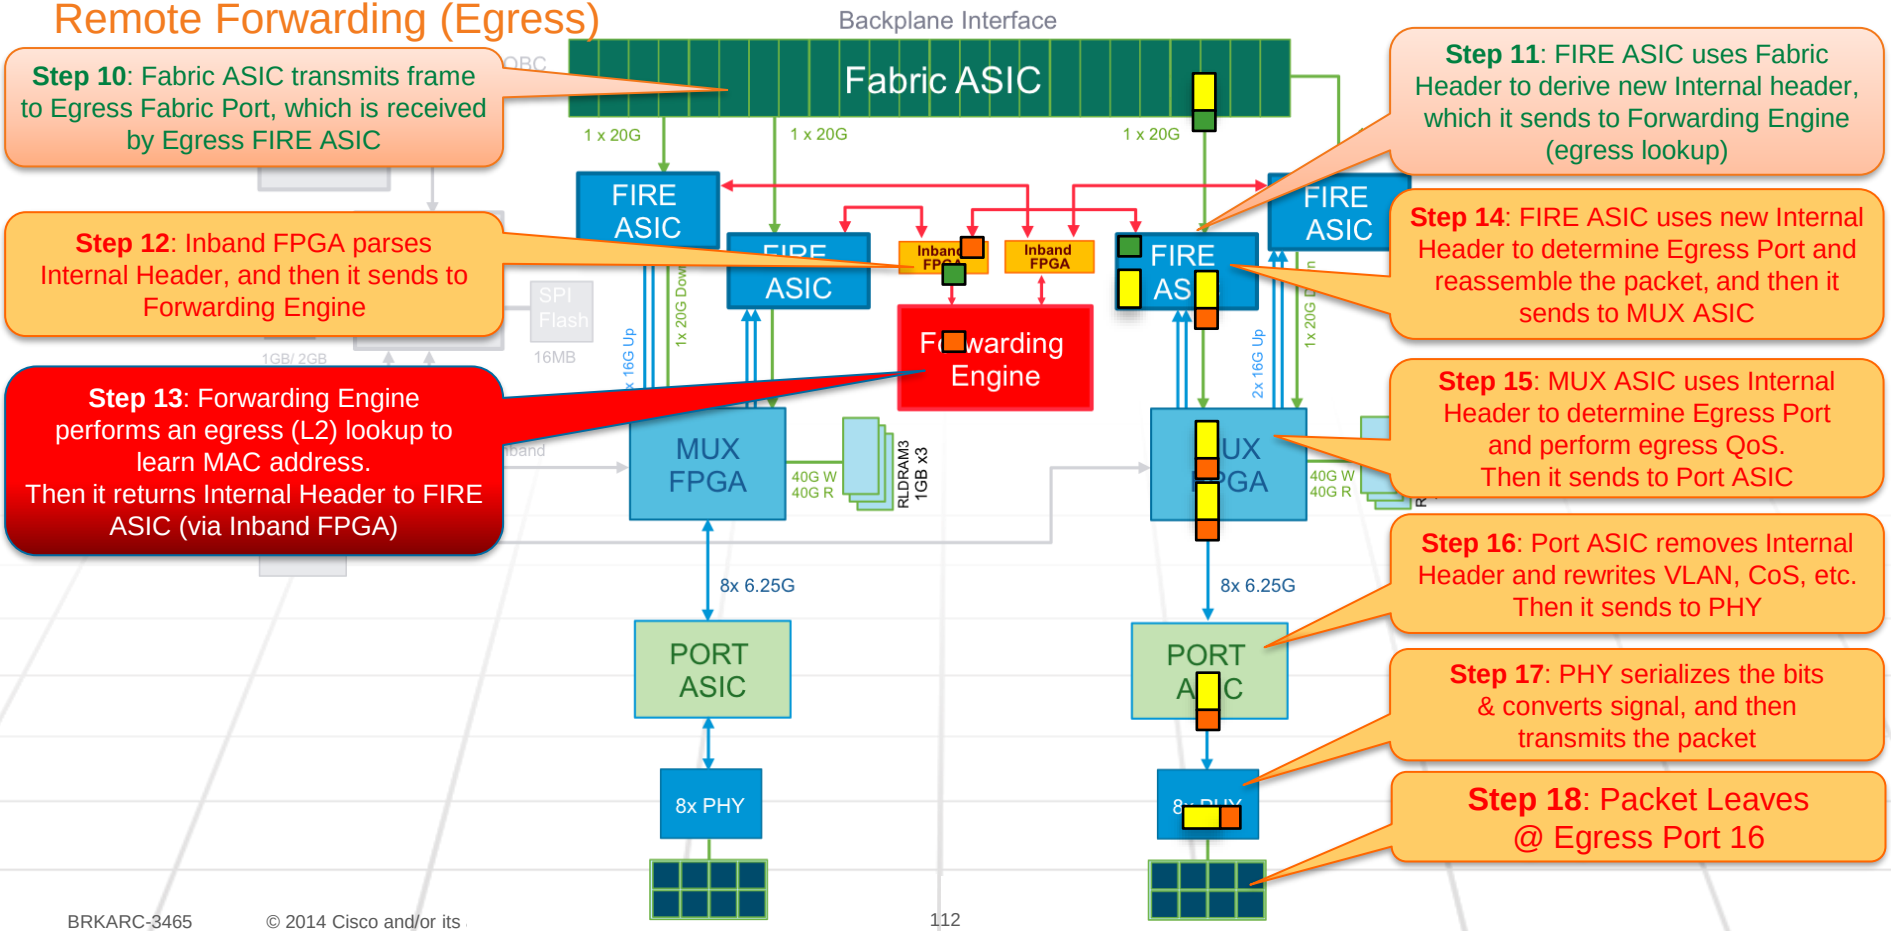

Catalyst 6880-X: Packet Walks

Local Forwarding (No Fabric)

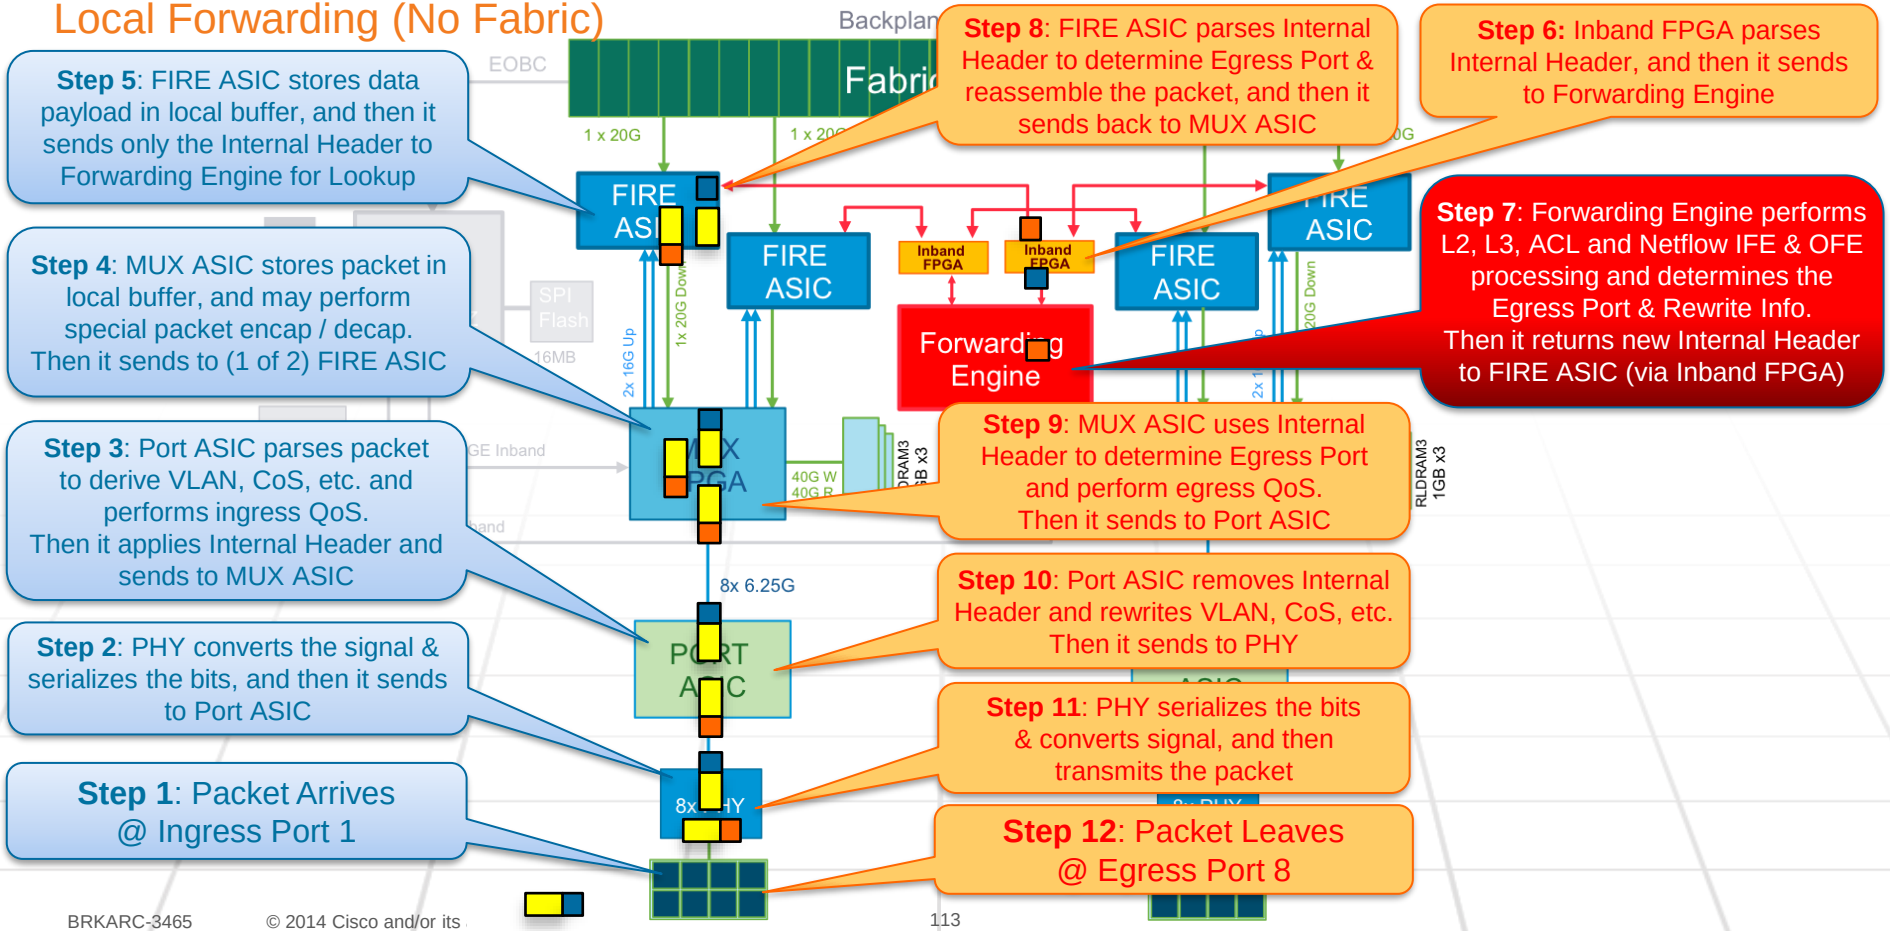

Summary

- **Innovation with Investment Protection**
  The new Catalyst 6800 architectures are based on the “Gold Standard” Catalyst 6500 Campus Switch, leveraging the same ASICs and Software, while providing a foundation for next-generation hardware.
- **Hardware Multi-Layer Switching**
  Get the best of both worlds! L2 & L3 forwarding, policies and statistics collection are performed by the same ASIC hardware, so there is minimal difference in performance and scale.
- **Combined Forwarding & Features**
  Enabling multiple additional hardware features such as Netflow, QoS and Security can be done without impacting performance or scale, because these hardware features are processed in parallel.

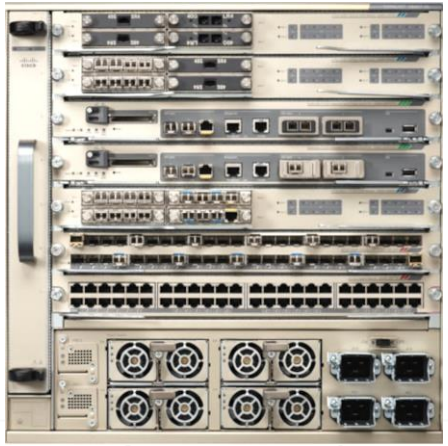

Thank you☺