



*LET'S  
BUILD  
TOMORROW  
TODAY*

# Catalyst 6800 Series Switch Architectures

BRKARC-3465

**Shawn Wargo**

Technical Marketing Engineer




# *Session Objectives*

# Session Objectives

What we will cover...

Help you become familiar with:

- **Catalyst 6807-XL**
- **Catalyst 6880-X**
- **Catalyst 6840-X** 
  - Chassis Architectures
  - Supported Modules
  - Forwarding Engine Functions
  - Hardware Feature Operations
  - Basic Packet Flows





# Session Objectives

What we will not cover...

**We will NOT cover...**

- **Virtual Switching System (VSS)**

- Refer to **BRKARC-3035**

- **Instant Access (IA)**

- Refer to **BRKARC-3502**

Visit our **C6800 VSS & Instant Access Demo Booth** in the **World of Solutions!**

Try our **VSS & IA Hands-On Lab**

- Visit **LTRCRS-2004**





***Why is this session  
important?***

# Why is this session important?


## Catalyst 6800 is an Architectural Foundation for Next Generation Campus

- **Increased Port Speeds** – To support the advent of next-generation integrated Wireless (e.g. 802.11ac) and integrated HD Video Conferencing.
- **Increased Port Density** – To support the increased number of directly-connected hosts (including Wireless) and adjacent network devices.
- **Increased Fabric Speeds** – To support the increased number of Speeds & Density
- **Improved Control-Plane** – To support the increased number of directly-connected hosts (e.g. ARP/ND) and adjacent network devices.
- **New Hardware Capabilities** – To off-load traditional software (CPU) capabilities (e.g. BFD & NDE), as well as to enable new features (e.g. FEX & AVB).
- **Innovation with Investment Protection** – To provide a support infrastructure that continues to allow older technologies (software & hardware), while enabling new ones.

# Agenda


## ❖ Chassis & Power

- ❖ C6807-XL
- ❖ C6880-X
- ❖ C6840-X
- ❖ Supervisor Architectures
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks



# Catalyst 6807-XL

## Modular Chassis Overview



# Catalyst 6807-XL

- **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
- **NEW 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 *
WS-X6816-10G-2T *
WS-X6716-10G *
WS-X6716-10T *
WS-X6704-10GE *
WS-X6724-SFP *
WS-X6748-SFP *
WS-X6748-GE-TX *
Service Modules
WS-SVC-NAM-3
WS-SVC-ASA-SM1
WS-SVC-WISM2-K9
ACE30-MOD-K9



# Catalyst 6807-XL:

## Environmental Overview




High Efficiency  
**4500 RPM**  
Redundant Fans



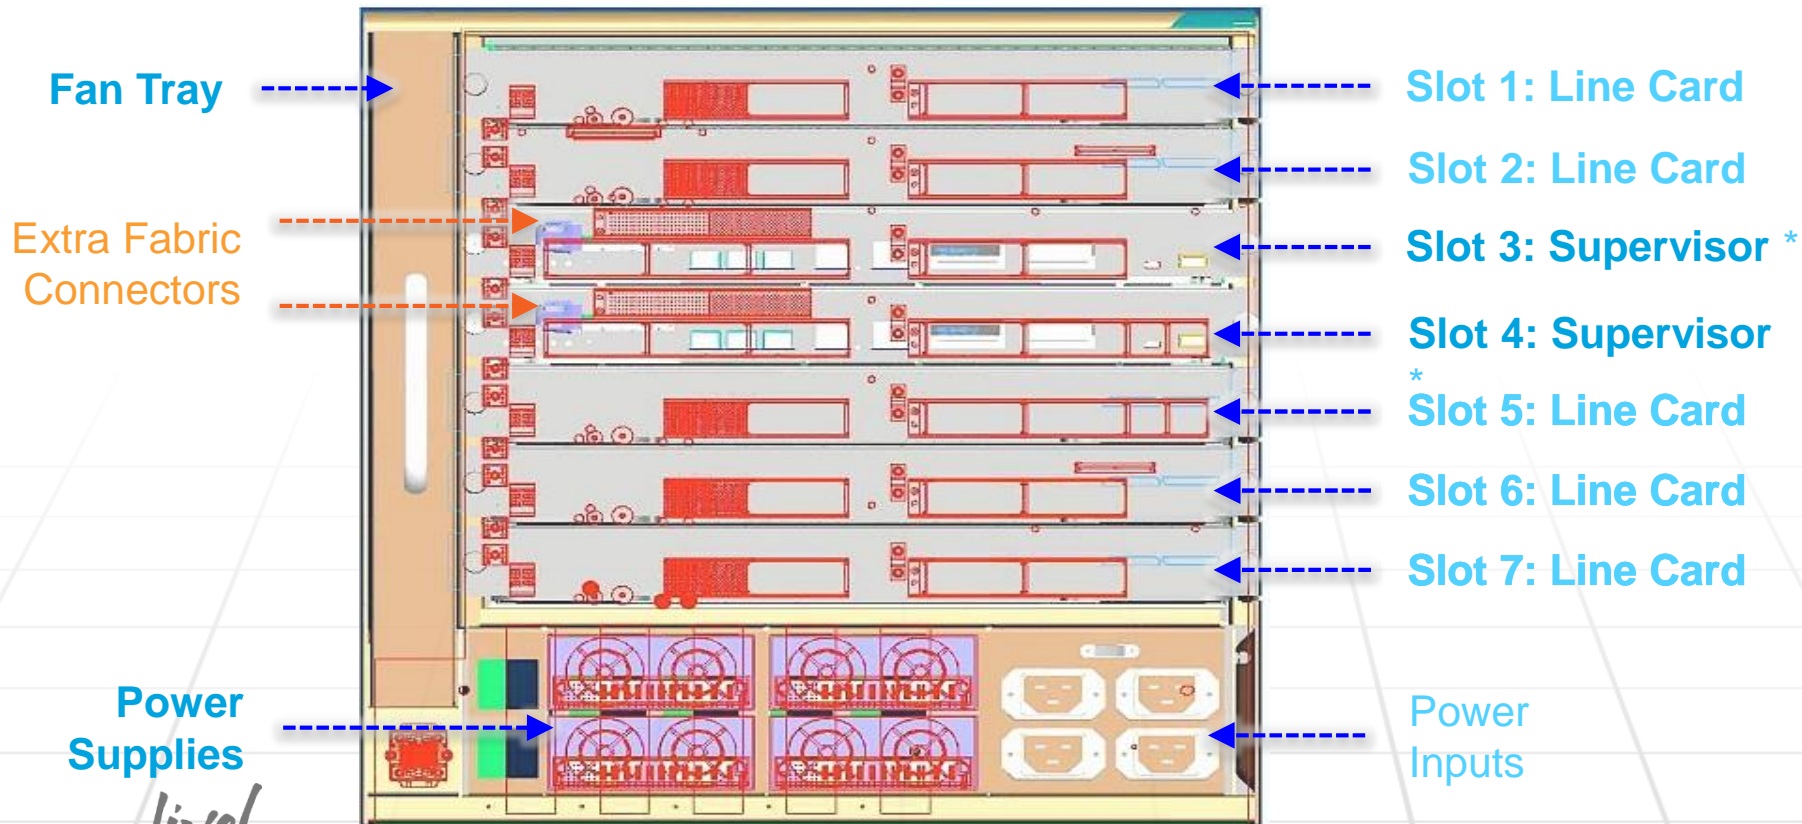
Platinum Efficient  
**3000W**  
Power Supplies

Cisco *live!*



# Catalyst 6807-XL

## Mechanical View



Cisco *live!*




# Catalyst 6807-XL

## Fan Redundancy & Air Flow




For Your  
Information

**Air Flow:**  
Side to Side



Front-Service  
Fan-Tray



### 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




For Your Information



### Power Supply Highlights:

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




LED	Color	Status	Description
IN	Green	Solid	Input OK
IN	Green	Blinking	Under-Current
OUT	Green	Solid	Output OK
OUT	Green	Blinking	Over-Current
Fault	Red	Solid	Malfunction

# Power Supply Redundancy

## Catalyst 6807-XL


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: if one PSU fails, then the +1 PSU will take over
- This mode is **Recommended (Default)**

### Combined Mode



- Same operational behavior as 6500-E
- Each PSU provides ~90% of capacity
- The total system power is ~360% of the capacity of a single PSU
- Pseudo-redundant behavior, but this is not equivalent to 1:1 or N+1 redundancy.
- This mode is **NOT Recommended**

# Agenda


## ❖ Chassis & Power

- ❖ C6807-XL
- ❖ **C6880-X**
- ❖ C6840-X
- ❖ Supervisor Architectures
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks



# Catalyst 6880-X

## Semi-Fixed Chassis Overview



Cisco *live!*



# Catalyst 6880-X

## Environmental Overview




High Efficiency  
**4500 RPM**  
Redundant Fans



Platinum Efficient  
**3000W**  
Power Supplies

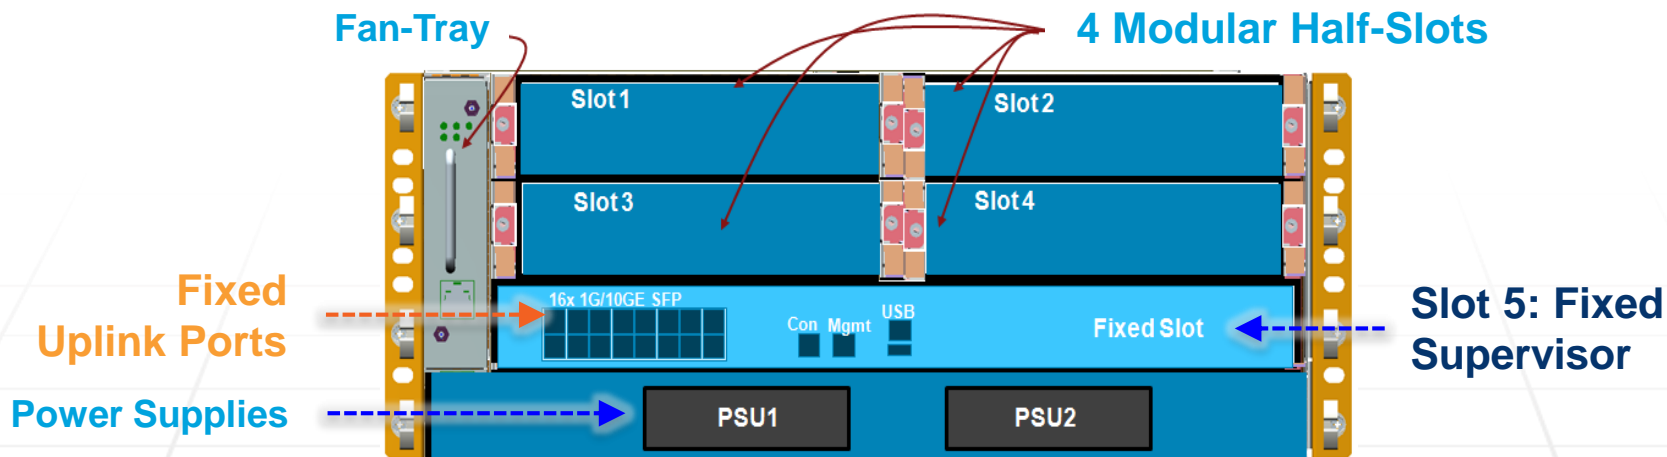
*Cisco*live!



Flexible Mounting Brackets

# Catalyst 6880-X



## Mechanical View



# Catalyst 6880-X

## Fan Redundancy & Air Flow

**Air Flow:**  
Side to Side



### 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









C6880-X-3KW-AC (DC)



### Power Supply Highlights:

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


LED	Color	Status	Description
IN		Solid	Input OK
IN		Blinking	Under-Current
OUT		Solid	Output OK
OUT		Blinking	Over-Current
Fault		Solid	Malfunction
ID		Solid	Identifies PSU

# Power Supply Redundancy

## Catalyst 6880-X

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

### Redundant Mode



PSU 1



PSU 2

- Each PSU provides ~50% of power needs
- Neither PSU operates at >60% or <40% capacity
- Either PSU can power the system on its own
- This mode is **Recommended (Default)**

### Combined Mode



PSU 1



PSU 2

- Each PSU provides up to 90% of its capacity
- Total power available is 180% of a single supply
- A single PSU may not have power for the system
- **This mode is NOT Recommended**

# Agenda

NEW

## ❖ Chassis & Power

- ❖ C6807-XL
- ❖ C6880-X
- ❖ **C6840-X**
- ❖ Supervisor Architectures
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks



# Catalyst 6840-X Chassis

## High-Level Overview



For Your  
Information


### Catalyst 6000 Feature Set

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

**2 RU**  
(Fixed Chassis)

Low-Power & Noise  
High-Efficiency Fans

**4 Models with  
up to 48 x 10G  
and 12 x 40G**



**Up to 256K FIB  
Up to 1.5M NetFlow  
Up to 64K ACL**

Powerful X86  
2.0GHz CPU  
& 4GB DRAM

**750/1100W (AC/DC)  
Power Supplies**


**1+1 Redundant PS  
(Front Serviceable)**

*Cisco live!*

# Catalyst 6840-X

## Environmental Overview


**NEW**



High Efficiency  
**4500 RPM**  
Redundant Fans




Platinum Efficient  
**750 / 1100W**  
Power Supplies



**Cisco**live!

# Catalyst 6840-X

## Mechanical View






# Catalyst 6840-X

## Fan Redundancy & Air Flow

**Air Flow:**  
Front to Back



Airflow  
Intake Vents


### Fan-Tray Highlights:

- Has 4 variable-speed High-Efficiency Fans (200 CFM)
- Supports 4 speeds between 3000 & 4500 RPM per Fan
- Capable of cooling of ALL 4 C6804-X chassis models
- 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

# Catalyst 6840-X

## PSU Redundancy & Inputs








C6840-X-750W-AC (DC)  
C6840-X-1110W-AC (DC)



## Power Supply Highlights:

- Two AC or DC models @ 110V: **750W** or **1100W**
- Up to 92% Power Efficiency @ 100% of load
- Power Hold-Up Time is ~20 msec @ 100% load
- “Front to Back” Variable-Speed Cooling Fans
- Supports **Combined** or **Redundant (1:1)** mode

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



# Power Supply Redundancy

## Catalyst 6840-X

The **Catalyst 6840-X** Utilizes **Two** Power Supplies in Either **Redundant** or **Combined** Mode

### Redundant Mode



PSU 1

PSU 2

- Each PSU provides ~50% of power needs
- Neither PSU operates at >60% or <40% capacity
- Either PSU can power the system on its own
- This mode is **Recommended (Default)**

### Combined Mode




PSU 1

PSU 2

- Each PSU provides up to 90% of its capacity
- Total power available is 180% of a single supply
- A single PSU may not have power for the system
- This mode is **NOT Recommended**


# Agenda

- ❖ Chassis & Power
- ❖ **Supervisor Architectures**
  - ❖ **VS-S2T-10G**
  - ❖ MSFC, PFC & Fabric
  - ❖ C6880-X
  - ❖ C6840-X
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks



# Supervisor Engine 2T

## Quick Facts



# VS-S2T-10G

## PFC4 (XL)

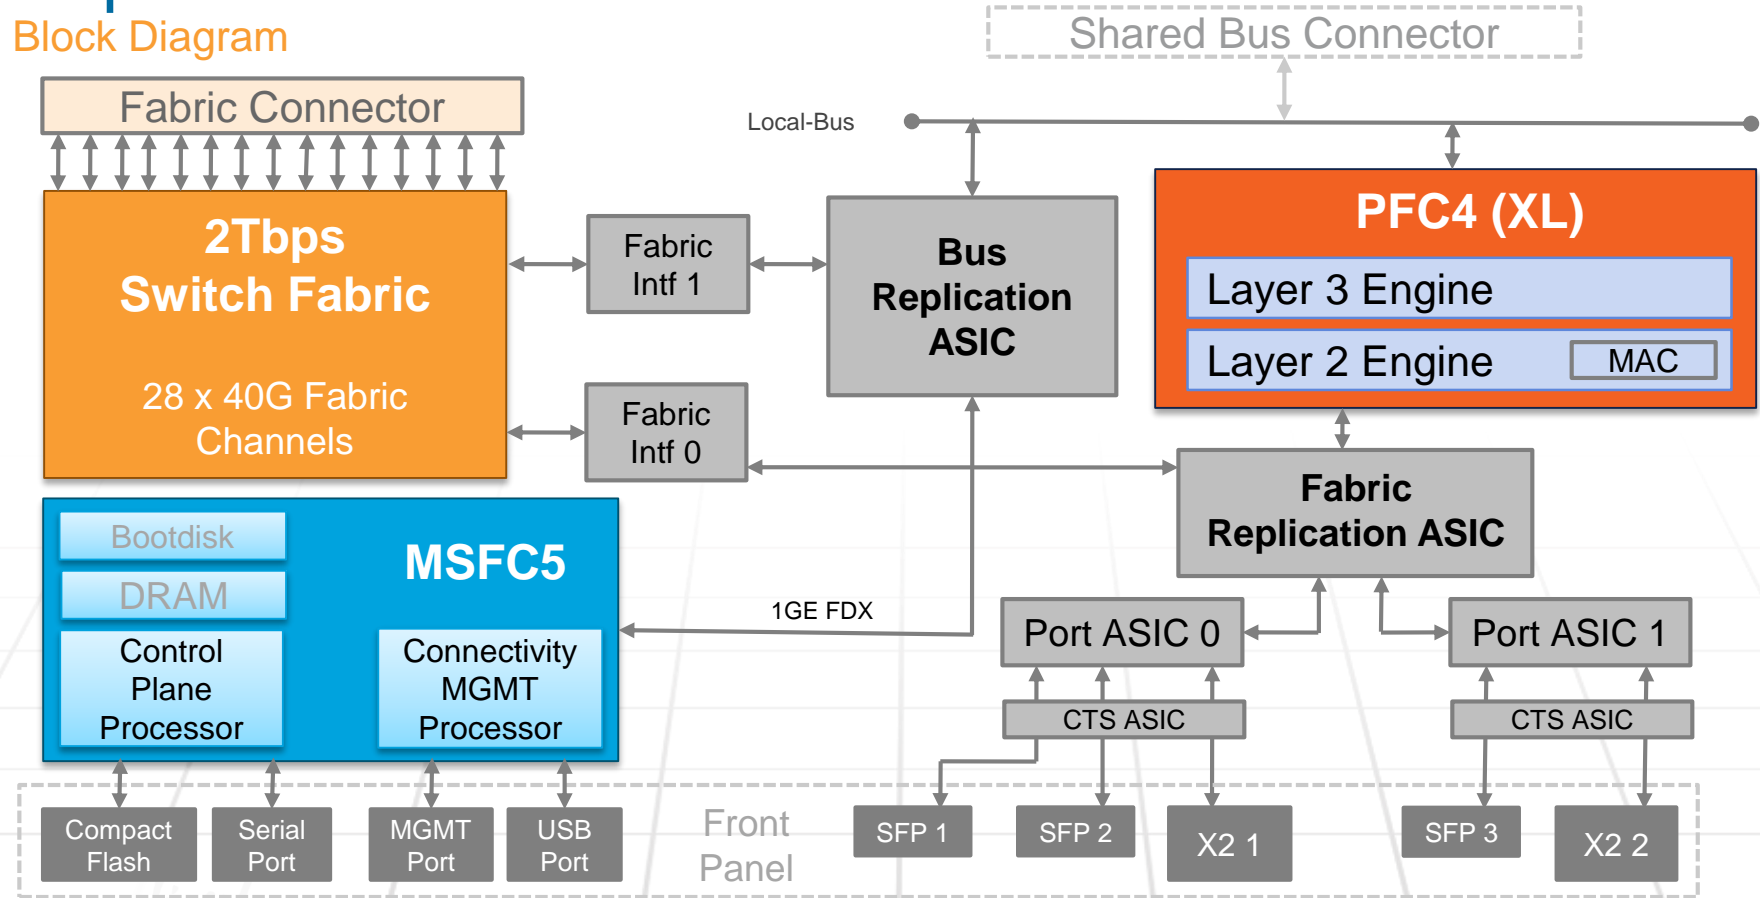
Cisco *live!*

## QUICK FACTS

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

# Supervisor 2T

## Block Diagram



# Supervisor Chassis Requirements



	Supervisor 720-10G	Supervisor 2T
Chassis	6500 (Non-E) <b>6500-E</b>	6500-E <b>6807-XL</b>
Fan Trays	E-Fan for E-Series Fan2 for Non-E Series	E-Fan for E-Series 6807-XL Fan-Tray
Power Supplies	2500W AC / DC or greater	
Supervisor Slots	6503-E: 1 and 2 6504-E: 1 and 2 6506-E : 5 and 6 6509-E : 5 and 6 6513-E : 7 and 8 <b>6807-XL : 3 and 4</b>	

E-Fan cannot be used in non-E Series Chassis


Fan2 cannot be used in E-Series

With Sup720 and 6513-E, the fabric channel distribution is the same as with Supervisor 720 and 6513.

With Sup2T and 6513-E or 6807-XL, only Supervisors (or Service Modules) are allowed in the Supervisor Slots

# Agenda

- ❖ Chassis & Power
- ❖ **Supervisor Architectures**
  - ❖ VS-S2T-10G
  - ❖ **MSFC, PFC & Fabric**
- ❖ C6880-X
- ❖ C6840-X
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks





# Multilayer Switch Feature Card

## MSFC5 Introduction


- **Single Dual-Core Processor**

- Combines functionalities of the Switch Processor (SP) and the Route Processor (RP)

- **1.5Ghz CPU Performance**

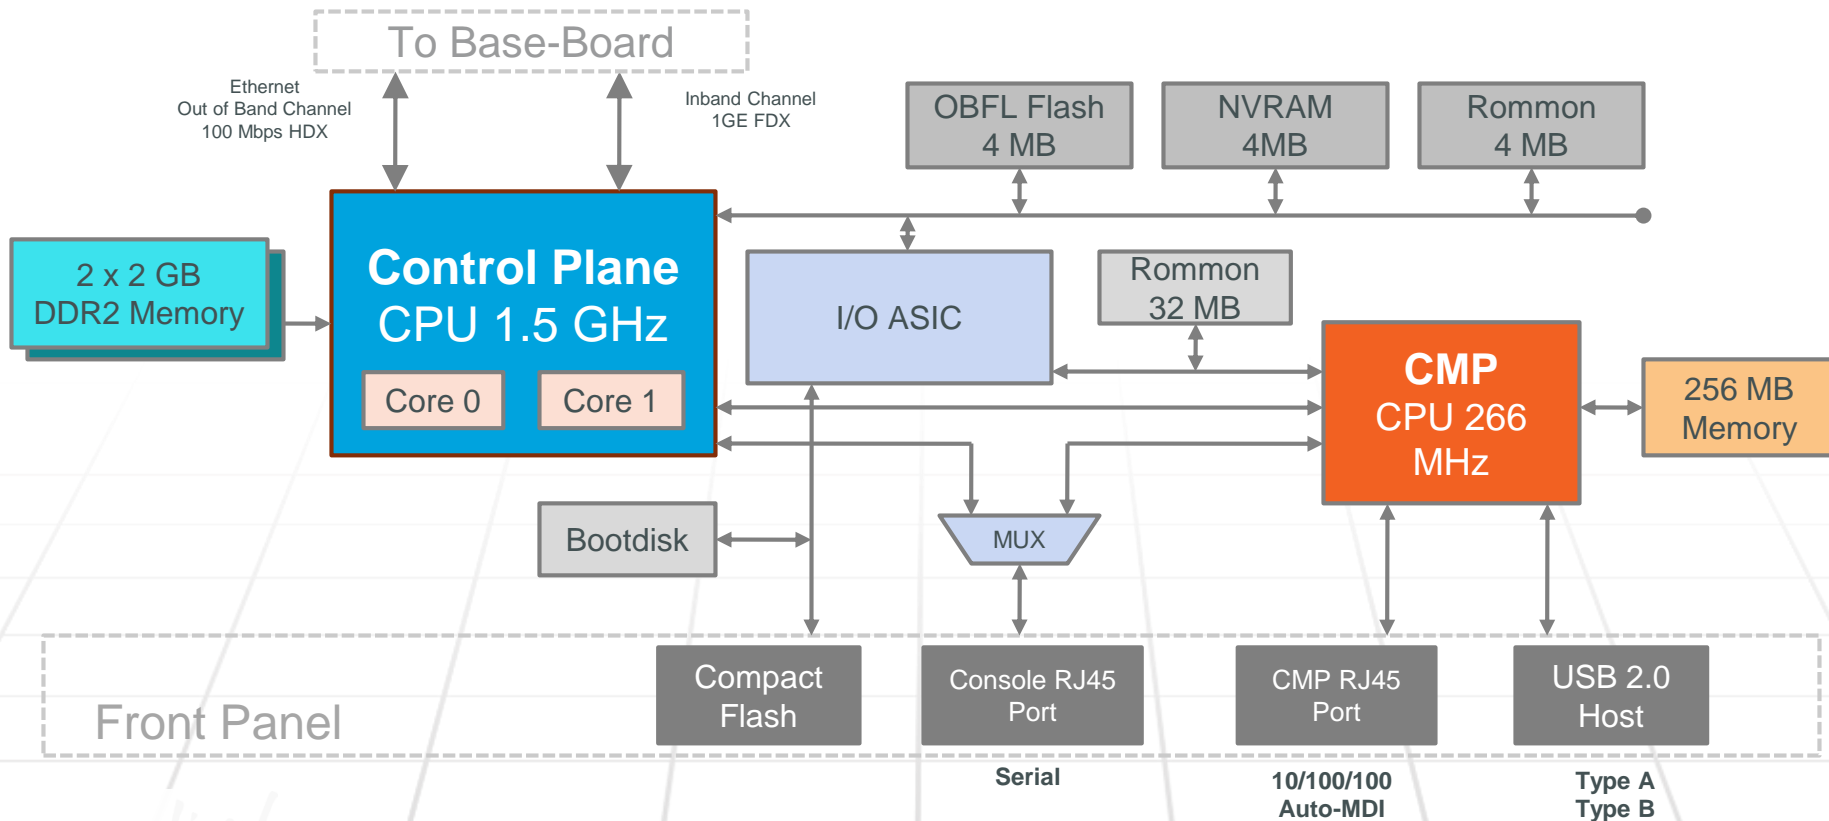
- **2GB or 4GB DDR2 DRAM**

- Single Bootdisk: File System
- Connectivity Management Processor (CMP)
- On-Board Failure Logging (OBFL)
- Mini Protocol Analyzer (MPA)



# Multilayer Switch Feature Card 5

## Block Diagram






# Multilayer Switch Feature Card 5

## Connectivity Management Processor



The new **Connectivity Management Processor (CMP)** supports out-of-band control of the system, completely independent of the RP CPU status...



### RP Image Recovery

- TFTP boot of the system

### RP File Transfer

- Image copy via TFTP

### Remote RP Reset

- Hard or Soft reset

### RP Console Logging

### USB Support

- Booting via approved USB flash
- USB Serial Console Port access

**Removes the need for a separate Telnet Server for console access**

# MSFC3 vs. MSFC5

Feature	MSFC3 (Supervisor 720-10G)	MSFC5 (Supervisor 2T)
CPU Speed	SP CPU – 600Mhz RP CPU – 600Mhz	Dual Core @ 1.5Ghz
Number of CPU cores	1	2
DRAM	SP CPU – Up to 1GB RP CPU – Up to 1GB	2GB (default) (4GB upgrade)
Connectivity Management Processor (CMP)	No	Single core @ 266Mhz 32MB Boot Flash 256MB System Memory
NVRAM	2MB	4MB
OBFL Flash	No	4MB
Bootflash / Bootdisk	SP CPU – 1GB (CF) RP CPU – 64MB (flash)	1GB (CF)

# MSFC = Multilayer Switch Feature Card

---

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

---



# Policy Feature Card 4

## PFC4 Introduction

Also applies to DFC4

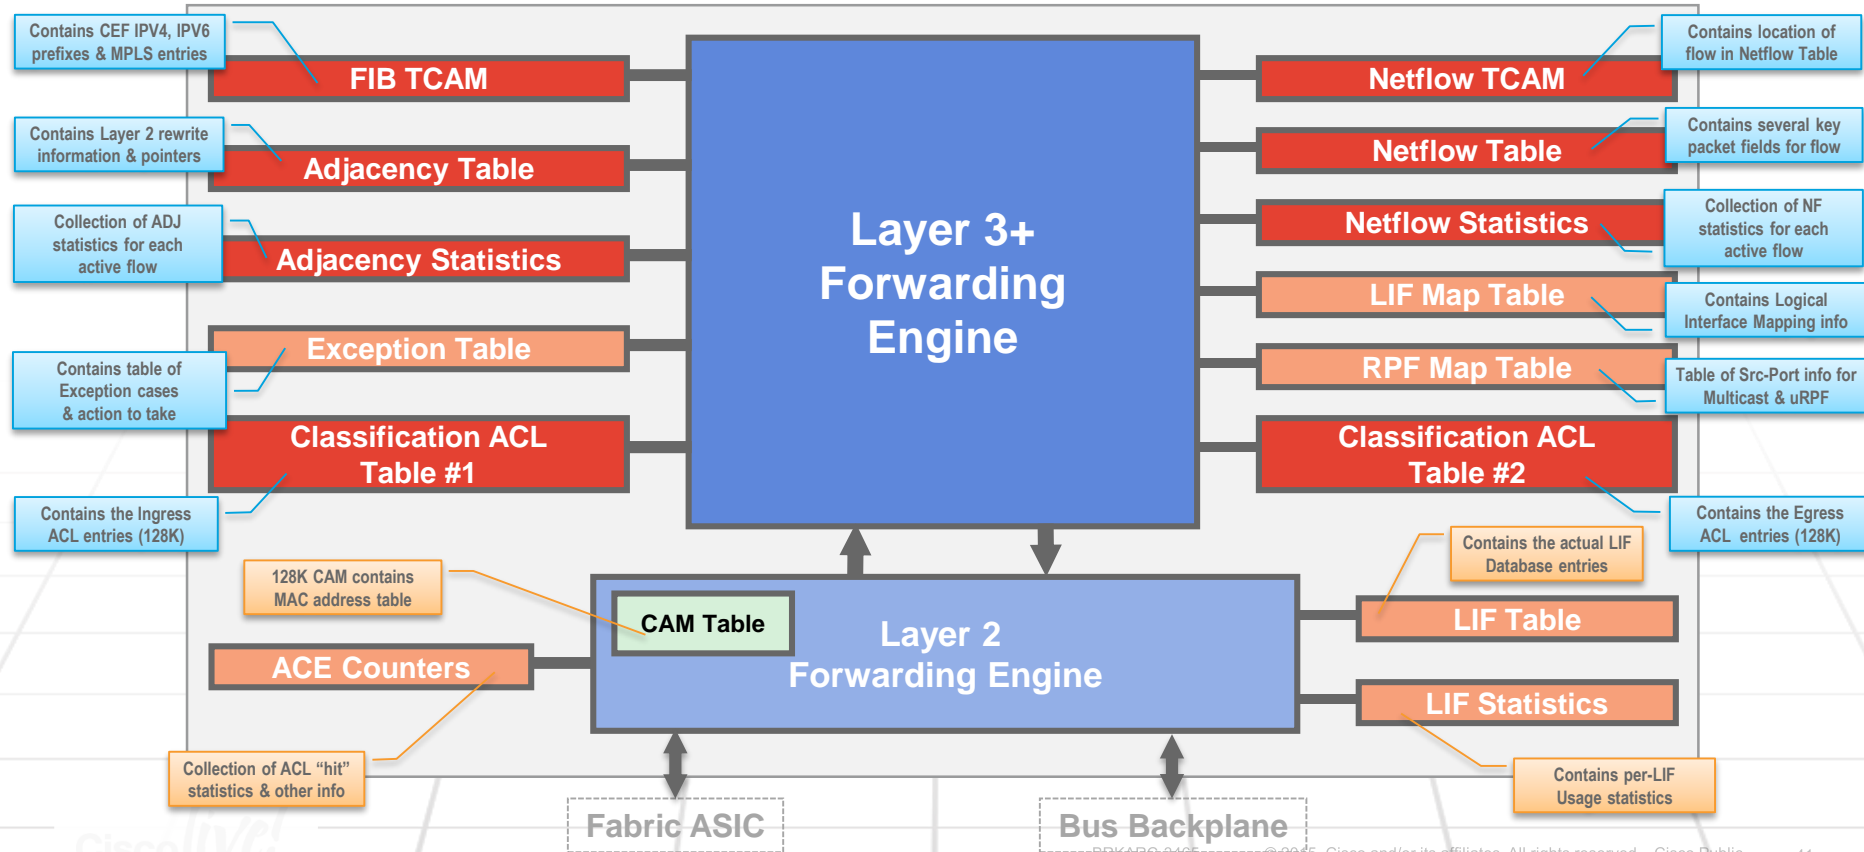


**EARL8**

- **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 8-bit EtherChannel Load-Balancing

# Policy Feature Card 4

## Block Diagram



# Policy Feature Card 4

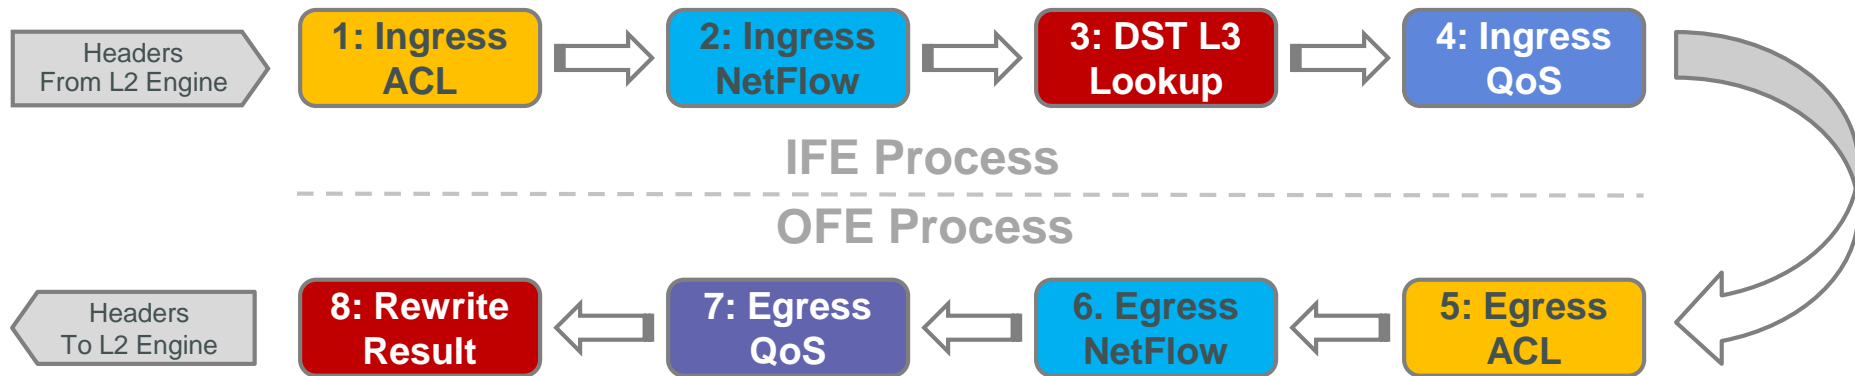
## EARL8 IFE / OFE Processing

The L3 Forwarding Engine ASIC has 2 processing pipelines @ 60Mpps:

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

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

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





# PFC4/DFC4 Lookup Process

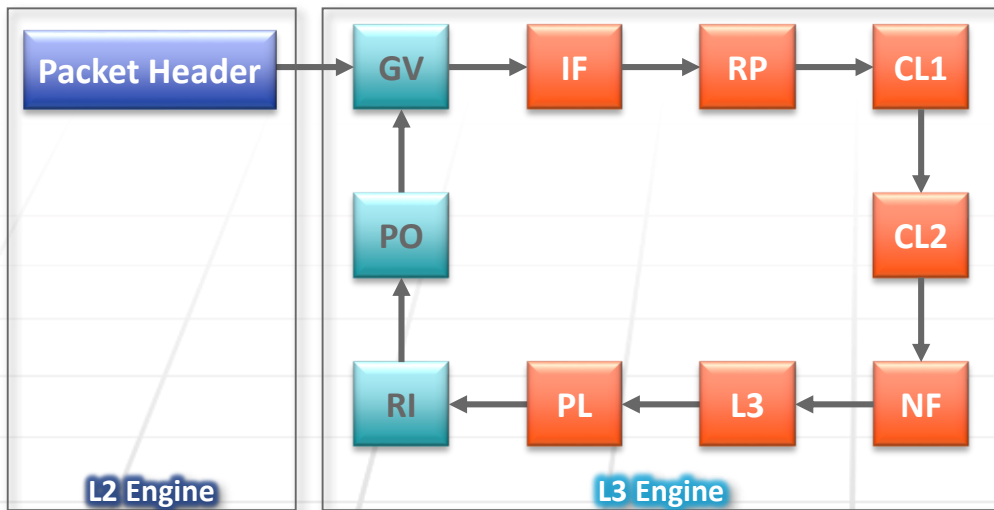
## IFE Processing



Architecturally, the PFC/DFC4 is very similar to the PFC/DFC3

What's new is the *dual-cycle* **Input (IFE)** and **Output (OFE)** Processing

Here we perform the **Input Forwarding Engine (IFE)** pass...



### IFE process:

1. **IF**: Get Port & Ingress LIF QoS info
2. **RP**: Src FIB and Src QoS Lookup
3. **CL1**: Ingress ACL Lookup
4. **CL2**: Select Ingress Class & Policy
5. **NF**: Ingress NetFlow lookup
6. **L3**: Dst FIB and Dst QoS Lookup
7. **PL**: Apply Ingress Policing & Marking

# PFC4/DFC4 Lookup Process

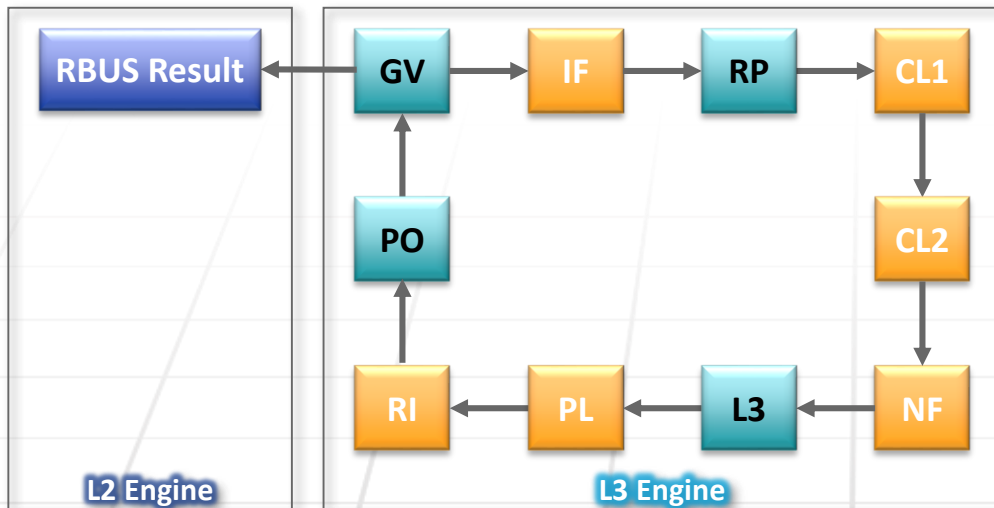
## OFE Processing



Architecturally, the PFC/DFC4 is very similar to the PFC/DFC3

What's new is the *dual-cycle* **Input (IFE)** and **Output (OFE)** Processing

Here we perform the **Output Forwarding Engine (OFE)** pass...



### OFE process:

1. **IF**: Get Egress LIF QoS info
2. **CL1**: Egress ACL TCAM lookup
3. **CL2**: Select Egress Policy & Class
4. **NF**: Select NF Egress Policy & Class
5. **PL**: Apply Egress Policing & Marking
6. **RI**: Generate final RBUS result

# PFC3 vs. PFC4

Feature	PFC3B/BXL	PFC3C/CXL	PFC4/XL
IPv4 Forwarding	Up to 30Mpps	Up to 48Mpps	Up to 60Mpps
IPv6 Forwarding	Up to 15Mpps	Up to 24Mpps	Up to 30Mpps
FIB TCAM IPv4	256K / 1M	256K / 1M	256K / 1M
FIB TCAM IPv6	128K / 500K	128K / 500K	128K / 500K
Adjacency Table	1M	1M	1M
Netflow Table	Up to 256K (XL)	Up to 256K (XL)	Up to 1M (XL) (Ingress 512K : Egress 512K)
MAC Table	64K (32K)	96K (80K)	128K
Egress Netflow	No	No	Yes
Flexible Netflow	No	No	Yes
MPLSoGRE	No	No	Yes
IPv6 uRPF	No	No	Yes

# PFC3 vs. PFC4

Feature	PFC3B/BXL	PFC3C/CXL	PFC4/XL
ACL Labels	4K	4K	16K
Security ACEs	Up to 32K	Up to 32K	Up to 192K (XL Default)
QoS ACEs	Up to 32K	Up to 32K	Up to 64K (XL Default)
Port ACLs	2K	2K	8K
Aggregate Policers	1023	1023	6K
Shared Microflow Policers	63	63	512
Egress Microflow Policing	No	No	Yes
Distributed Policers	No	No	Yes
Packet or Byte Based Policing	No	No	Yes
RPF Interfaces	2	2	16
Native VPLS	No	No	Yes
VSS	No	Yes	Yes

# PFC = Policy Feature Card

---

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

---



# 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 Fabric Channel can independently operate at **20Gb/sec** or **40Gb/sec\***
- Mixing of 6800 & 6900 (20G & 40G) modules does not affect speeds of other modules






# Crossbar Switch Fabric

## More Details

The **Catalyst 6500 & 6800** series **eliminated** the earlier **Bus-based** limits by creating a new “**Crossbar**” **Switch Fabric** for its backplane.

A **Crossbar Architecture** is essentially **2N busses** (where N is the number of ASICs connected to the Switch Fabric) **connected by N\*N cross-points**.



# Crossbar Switch Fabric


More Details



**This allows multiple LC's to transmit and receive data simultaneously**

**A scheduler is responsible for selecting which LCs transmit and which LCs receive data, during any given fabric cycle.**

**This can be viewed as (1 or more) dedicated Fabric channels to each LC**



**Per Channel bandwidth is defined by 2 factors:**

- Clock Speed (in Hz) defines the maximum BPS
- Line Encoding (e.g. 8/10b) defines usable bits

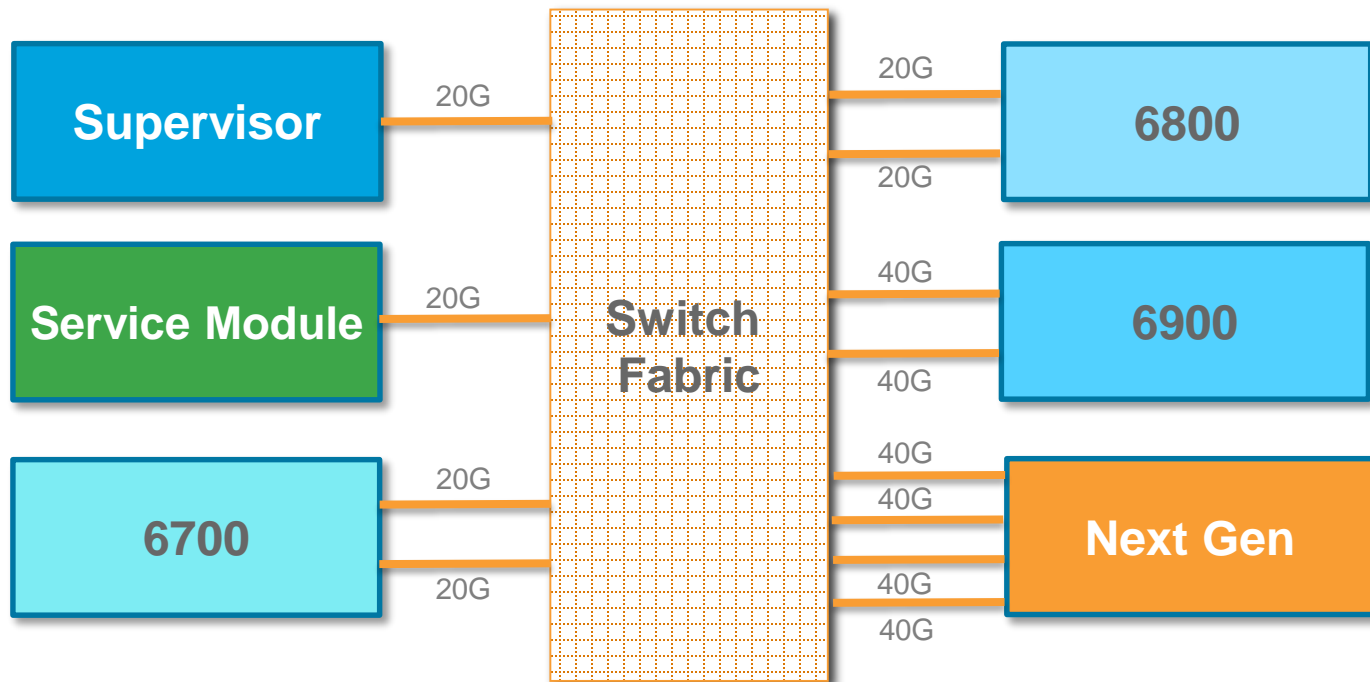
**Per Slot bandwidth can increase with multiple Fabric Channels:**

- $X \text{ bps} * Y \text{ channels} = Z \text{ bandwidth}$

Cisco *live!*

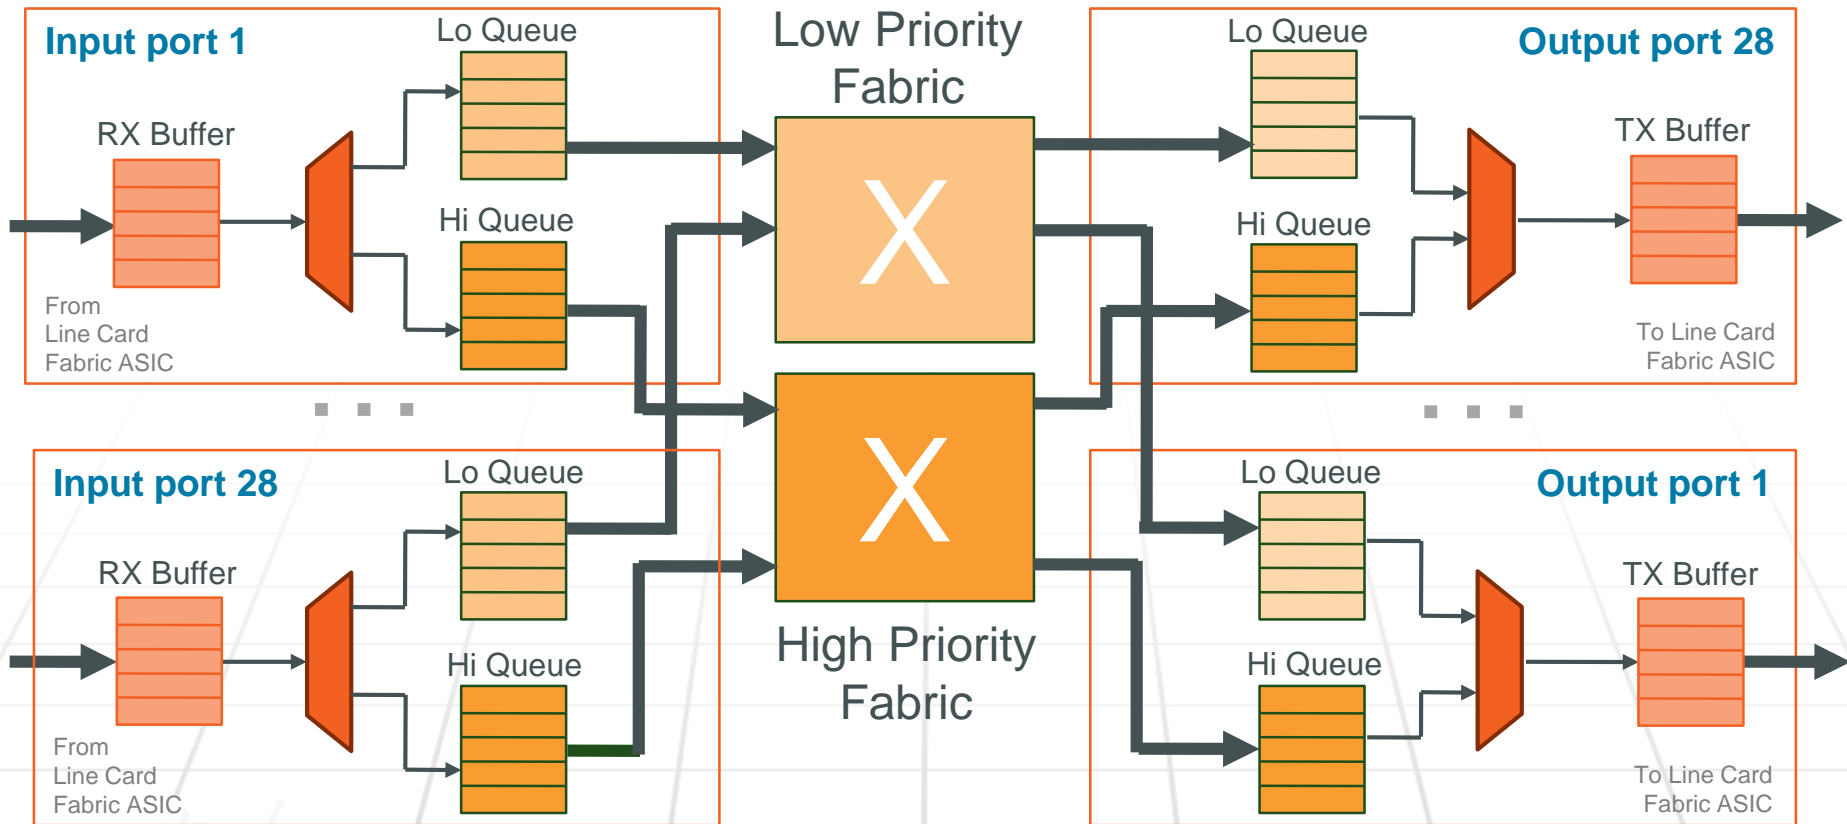
# Crossbar Switch Fabric

## Logical Architecture



# 2T Switch Fabric

## Block Diagram




# 2T Switch Fabric

## High Availability



- When using a **VS-S2T-10G** with a redundant Supervisor, two (2) fabric channels \* are connected "back to back"
  - **Standby Supervisor uplink connection to Active Supervisor**
  - **Active Supervisor bus connection to Standby Supervisor**
- The Standby Supervisor is in DFC Mode, with its Bus connection disabled
- The redundant Standby Supervisor enables its redundant channels to dCEF2T (e.g. **WS-X6908-10GE** and **WS-X6904-40GE**) modules for ~50ms failover...



\* "Back-to-Back"  
Fabric Channels, for  
sub-second (~50ms)  
SSO switchover

# 720G vs 2T Fabric



Feature	Sup720	Sup720-10G	Sup2T
Number of Channels	18	20*	28
Aggregate Bandwidth	720 Gbps	720 Gbps	2 Tbps
Channel Speeds (bps)	8G / 20G	8G / 20G	20G / 40G
Fabric Redundancy	Yes	Yes	Yes
SSO Fabric Hot Synch	No	Yes*	Yes
Redundant Channels	No	No	Yes
Fabric Priority (QoS)	Single Fabric Hi / Lo Priority	Single Fabric Hi / Lo Priority	Hi Priority Fabric Lo Priority Fabric
Clear Block Support	Yes	Yes	Yes
Switching Modes (DBUS Header Size)	Bus, Truncated, Compact	Bus, Truncated, Compact	Truncated, Compact
Requires E-Series	No	No*	Yes



# Switch Fabric = Data Plane

---

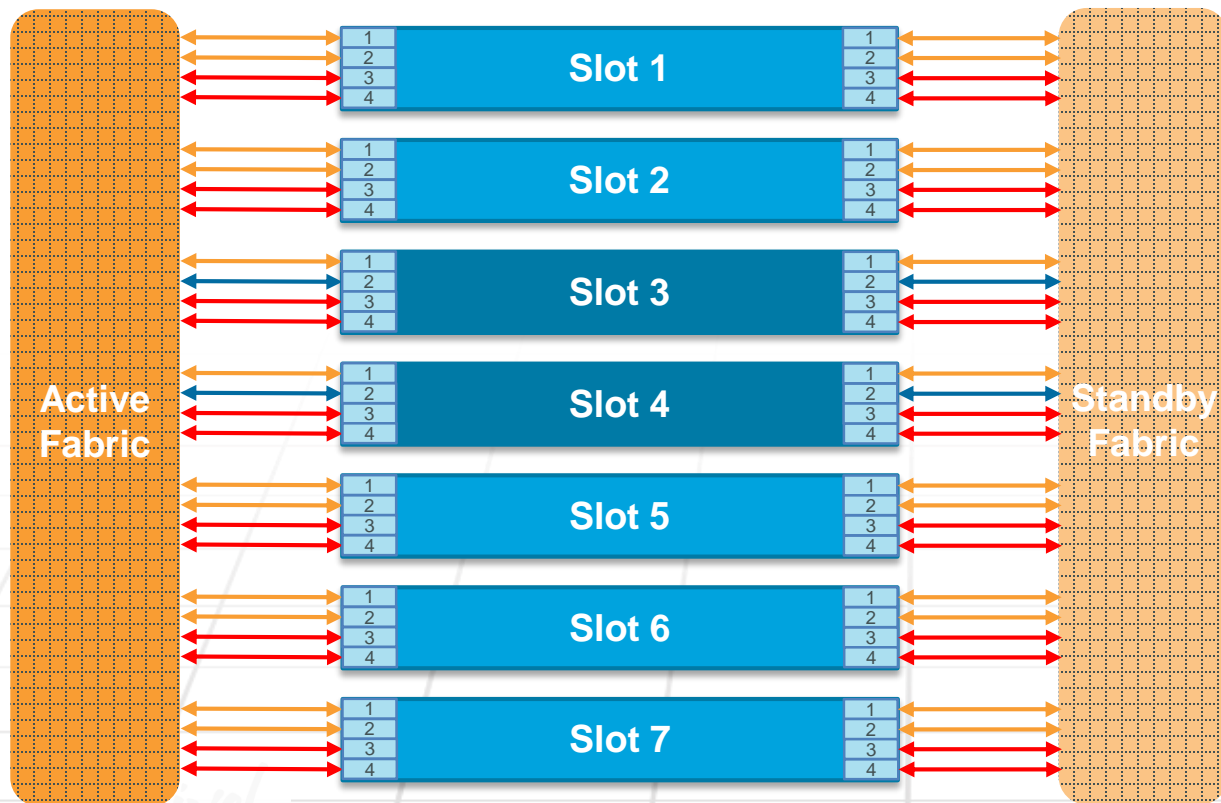
A dedicated set of Crossbar Channels that interconnect All Cards...

---



# 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 “Channel” 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**

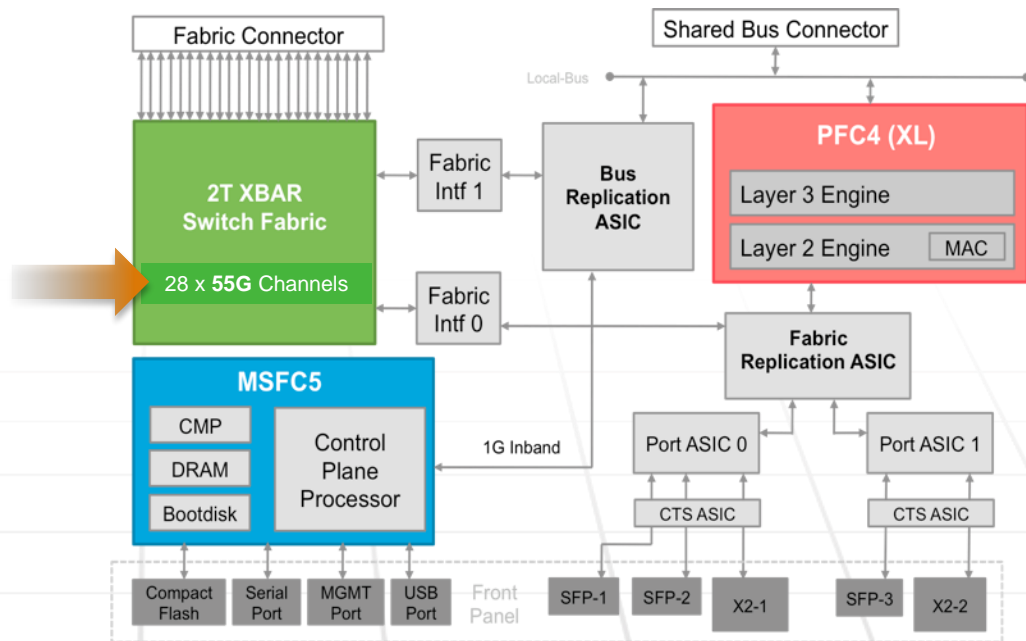
# Catalyst 6807-XL

## How Supervisor 2T operates



## Supervisor Engine 2T:

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




# Agenda

- ❖ Chassis & Power
- ❖ **Supervisor Architectures**
  - ❖ VS-S2T-10G
  - ❖ MSFC, PFC & Fabric
  - ❖ **C6880-X**
  - ❖ C6840-X
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks




# Catalyst 6880-X

## Base Board & System Controller



Forwarding  
Daughter Board



16 x SFP+ Ports:  
VSS, IA (FEX),  
LISP, MPLS, HQoS,  
MACSEC, SGT,  
1588 PTP & AVB\*  
available on Every Port

USB Host (Type A)  
USB Console (Type B)  
RJ-45 Console and  
Management Ports

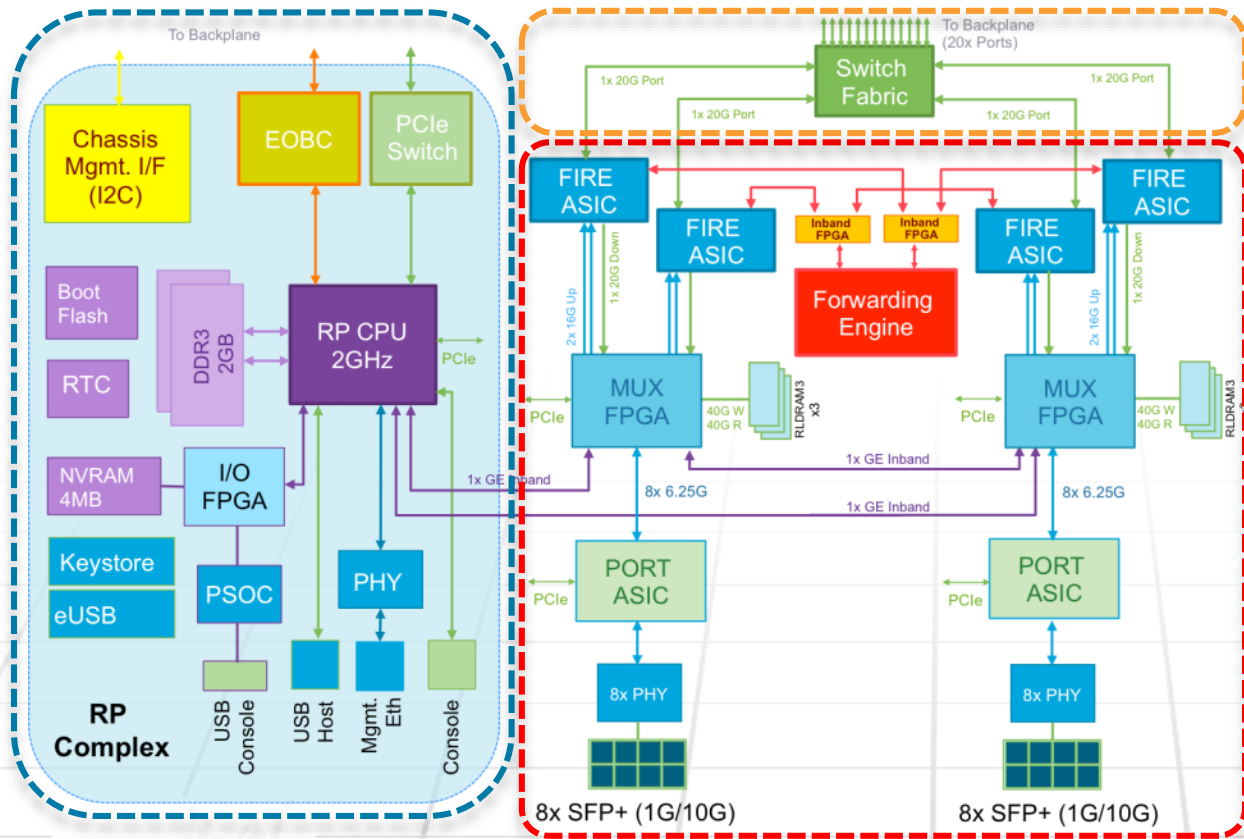
Two HW Options	6880-X-LE	6880-X
IPv4 Routes	256K	2M
IPv6 Routes	128K	1M
Multicast Routes	64K	128K
MAC Addresses	128K	128K
Security ACL Entries	128K	128K
QoS ACL Entries	64K	256K
Flexible NetFlow Entries	512K	1M

\* Hardware Capable

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

# Catalyst 6880-X:

## Fixed Supervisor Design



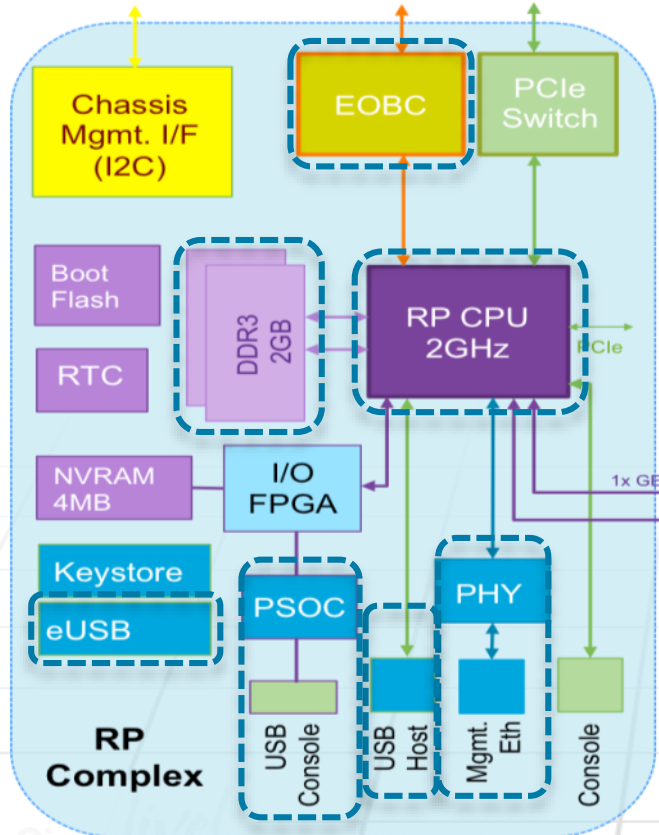
### 3 Main Components:

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



# Catalyst 6880-X

## Fixed Supervisor - RP Complex

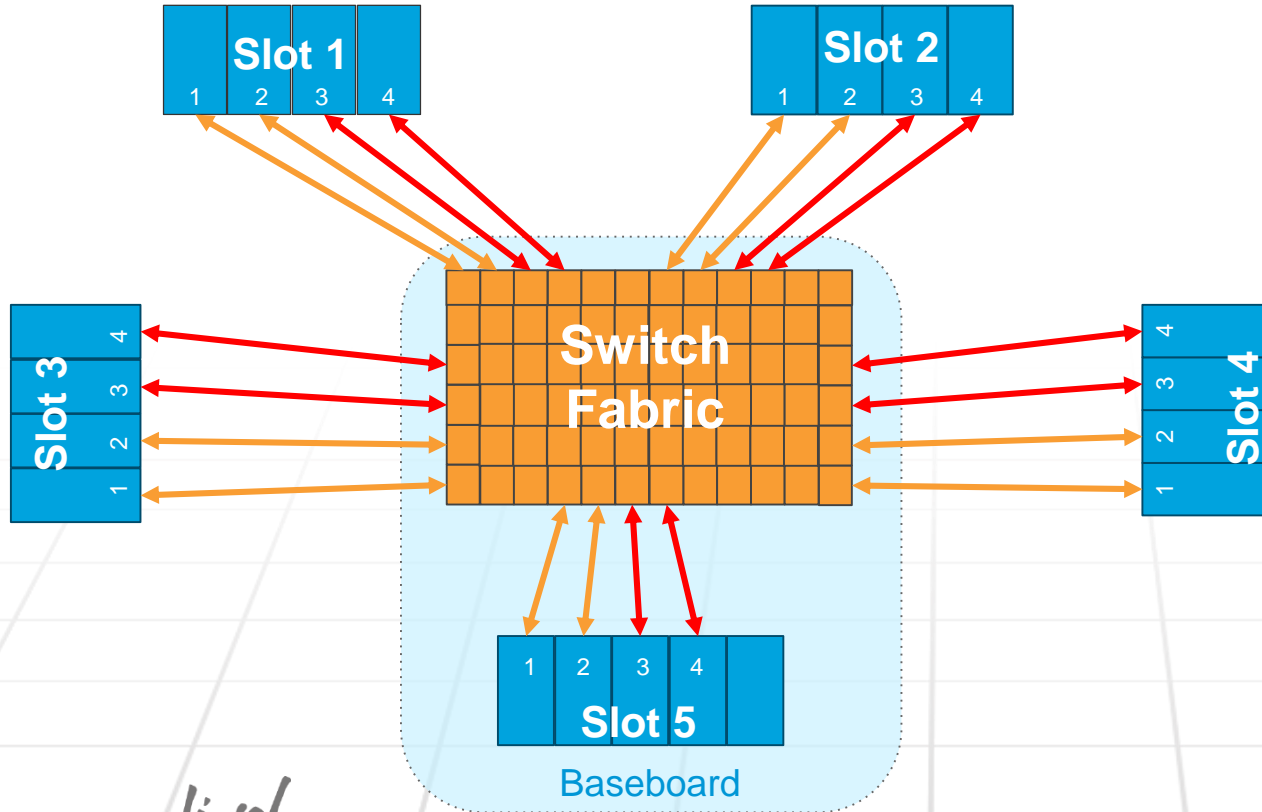


## Essentially the same as MSFC5!

- 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
- Bootdisk 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 "Channel" can use any of the following Clock Frequencies:

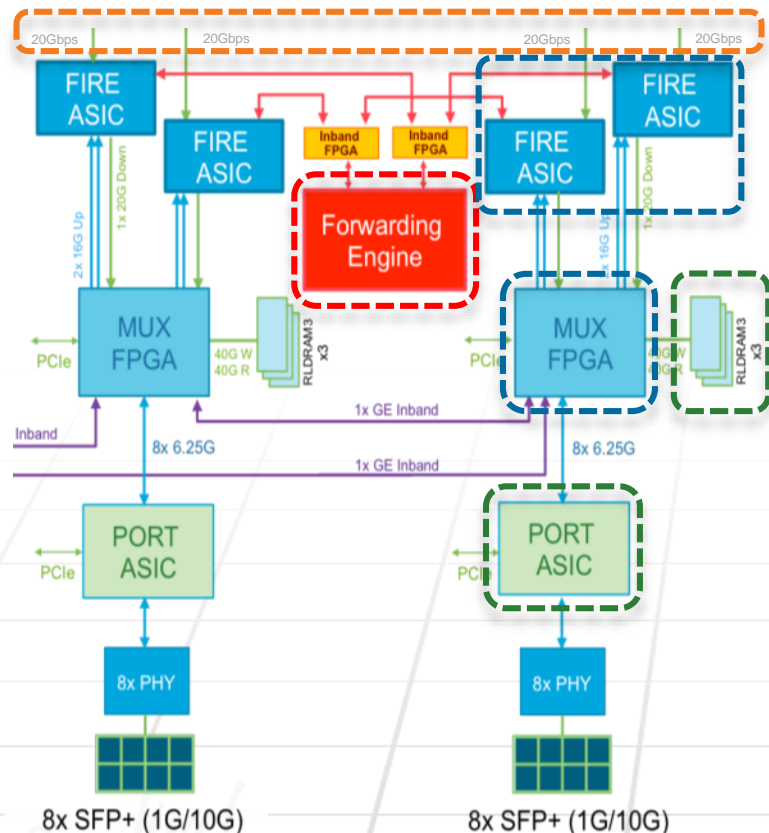
- 6.25 GHz for **40 Gbps**
- 7.50 GHz for **55 Gbps**

# Catalyst 6880-X

Fixed Supervisor – Baseboard 16 x SFP+ Ports



FIRE = Fabric Interface & Replication Engine



**Essentially same as Sup2T + 6904-40G**

- **16 ports of 1/10G SFP+**
  - 2 x 8 SFP+ Port-Groups
- **Enhanced DFC4-E Forwarding Engine**
- **80Gbps to Switch Fabric (2 Modes)**
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- **Improved 40Gbps Fabric/Replication ASIC**
- **New 40Gbps Port Interface MUX FPGA**
- **RLDRAM Packet Buffers on MUX & Port ASIC**
  - 1.25 or 2.5MB RX per Port (10MB per Port ASIC)
  - 24 or 48MB TX per Port (192MB per MUX FPGA)

# Agenda

NEW

- ❖ Chassis & Power
- ❖ **Supervisor Architectures**
  - ❖ VS-S2T-10G
  - ❖ MSFC, PFC & Fabric
  - ❖ C6880-X
  - ❖ **C6840-X**
- ❖ Module Architectures
  - ❖ L2 Packet Forwarding
  - ❖ L3 Packet Forwarding
  - ❖ NetFlow & NDE
  - ❖ Access Control Lists
  - ❖ Packet Walks



# Catalyst 6840-X Series

## Fixed Chassis Portfolio

**NEW**

### C6816-X-LE



### C6832-X-LE



### C6824-X-LE-40G



### C6840-X-LE-40G



Native Optics

SFP/SFP+

SFP/SFP+

SFP/SFP+ and QSFP

SFP/SFP+ and QSFP

# of 10G Ports

**16**

**32**

**24 +**  
8 using breakout cable

**40 +**  
8 using breakout cable

# of 40G Ports

4 using reverse adapter

8 using reverse adapter

**2 +**  
6 using reverse adapter

**2 +**  
10 using reverse adapter

IPv4 / v6 Routes

**256K / 128K**

**256K / 128K**

**256K / 128K**

**256K / 128K**

Multicast Routes

128K / 64K

128K / 64K

128K / 64K

128K / 64K

MPLS Labels

256K

256K

256K

256K

MAC Addresses

128K

128K

128K

128K

Security ACL

64K (Shared)

64K (Shared)

64K (Shared)

64K (Shared)

QoS ACL

64K (Shared)

64K (Shared)

64K (Shared)

64K (Shared)

Flexible NetFlow

**512K**

**1M**



**1M**

**1.5M**

# Catalyst 6840-X:

## Fixed Chassis Design

NEW



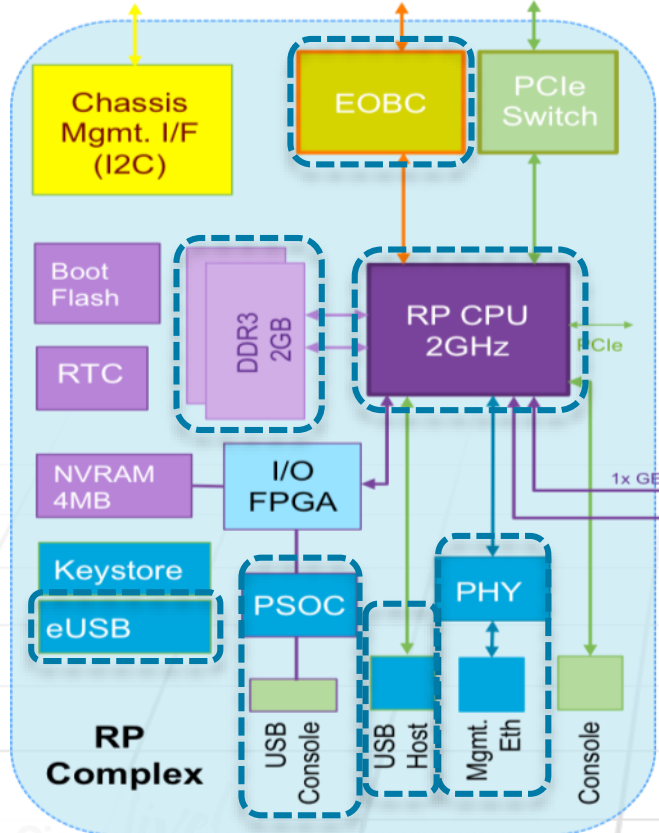
### 3 Main Components:

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

# Catalyst 6840-X

## Fixed Chassis - RP Complex

NEW



## Essentially the same as C6880-X RP


- 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
- Bootdisk Flash replaced with eUSB
- New Switched EOBC Interface



# Catalyst 6840-X:

## Fabric Channel Distribution

**NEW**



↔ Fabric Channels

Each "Fabric Channel" uses 8 SerDes Lanes

Each "Channel" uses the following Clock Frequency:

- 6.25 GHz for **40 Gbps**

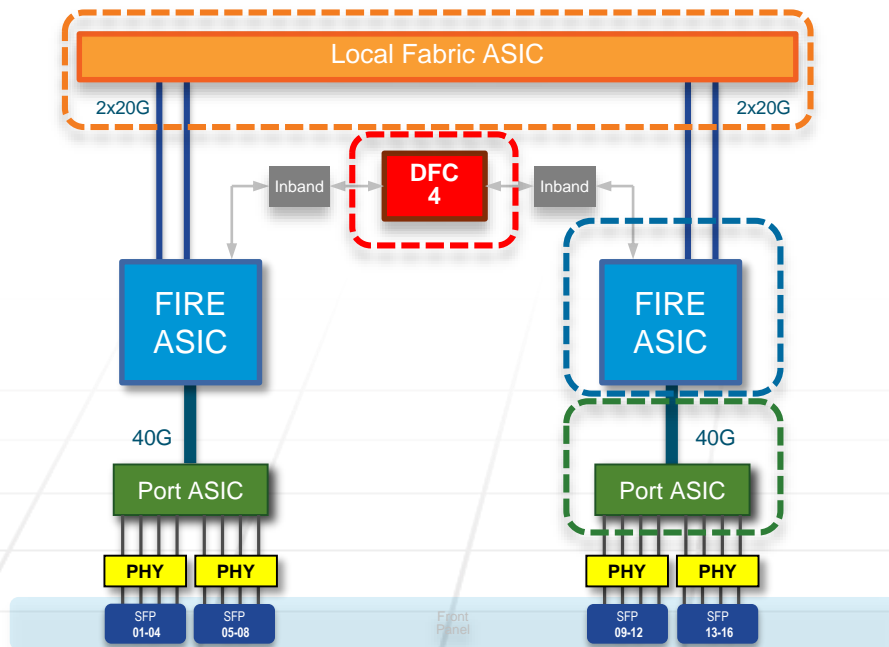
# Catalyst 6840-X

## C6816-X-LE – 16 x SFP+ Ports



**NEW**

FIRE = Fabric Interface & Replication Engine



## Essentially same as 6880-X Baseboard

- **16 ports of 1/10G SFP+**
  - 2 x 8 SFP+ Port-Groups
- **Enhanced DFC4-E Forwarding Engine**
- **80Gbps to Switch Fabric (2 Modes)**
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- **New 40Gbps Fabric & Replication ASIC**
- **Combines FIRE ASIC & MUX FPGA of 6880-X**
- **DDR3 Packet Buffers on FIRE & Port ASIC**
  - 1.25 or 2.5MB RX per Port (10MB per Port ASIC)
  - 250 or 500MB TX per Port (2GB per FIRE ASIC)

CiscoLive!

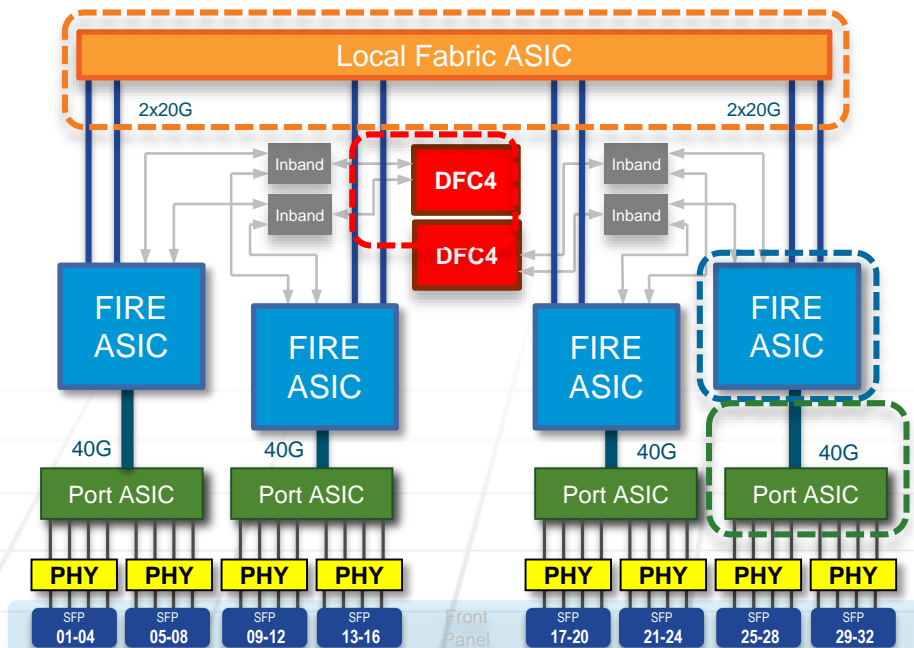
# Catalyst 6840-X

## C6832-X-LE – 32 x SFP+ Ports



**NEW**

FIRE = Fabric Interface & Replication Engine



## Essentially same as 6880-X Baseboard

- **16 ports of 1/10G SFP+**
  - 2 x 8 SFP+ Port-Groups
- **2 x DFC4-E Forwarding Engines**
- **80Gbps to Switch Fabric (2 Modes)**
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- **New 40Gbps Fabric & Replication ASIC**
- **Combines FIRE ASIC & MUX FPGA of 6880-X**
- **DDR3 Packet Buffers on FIRE ASIC**
  - 250 or 500MB TX per Port (2GB per FIRE ASIC)

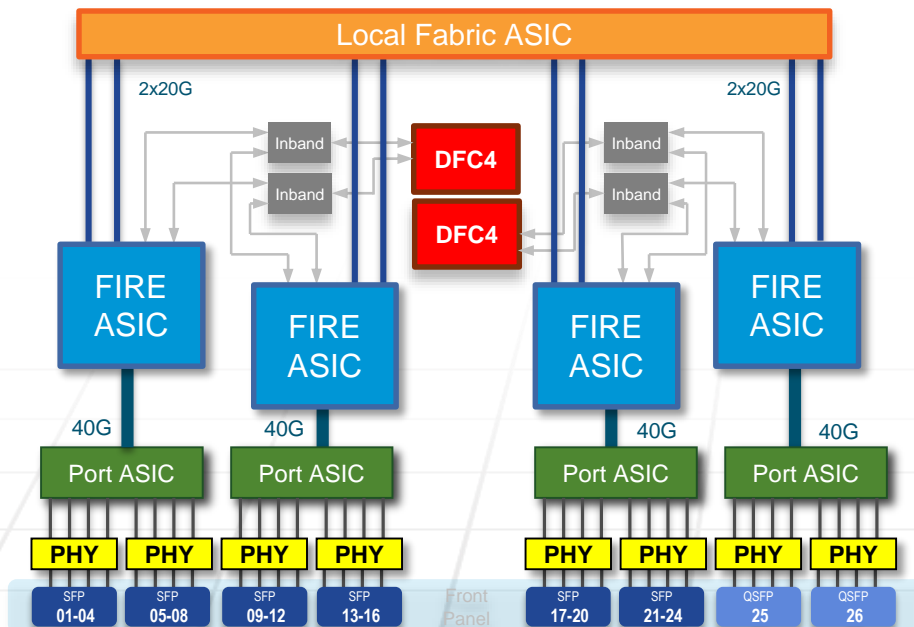
# Catalyst 6840-X

C6824-X-LE-40G – 24 x SFP+ & 2 x QSFP



**NEW**

FIRE = Fabric Interface & Replication Engine



## Essentially same as 6880-X Baseboard

- **24 ports of 1/10G SFP+, 2 ports of 40G QSFP**
  - 3 x 8 SFP+ Port-Groups, 1 x 2 QSFP Port-Group
- **2 x DFC4-E Forwarding Engines**
- **80Gbps to Switch Fabric (2 Modes)**
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- **New 40Gbps Fabric/Replication ASIC**
- **Combines FIRE ASIC & MUX FPGA of 6880-X**
- **DDR3 Packet Buffers on FIRE ASIC**
  - 250 or 500MB TX per Port (2GB per FIRE ASIC)

*Cisco*live!

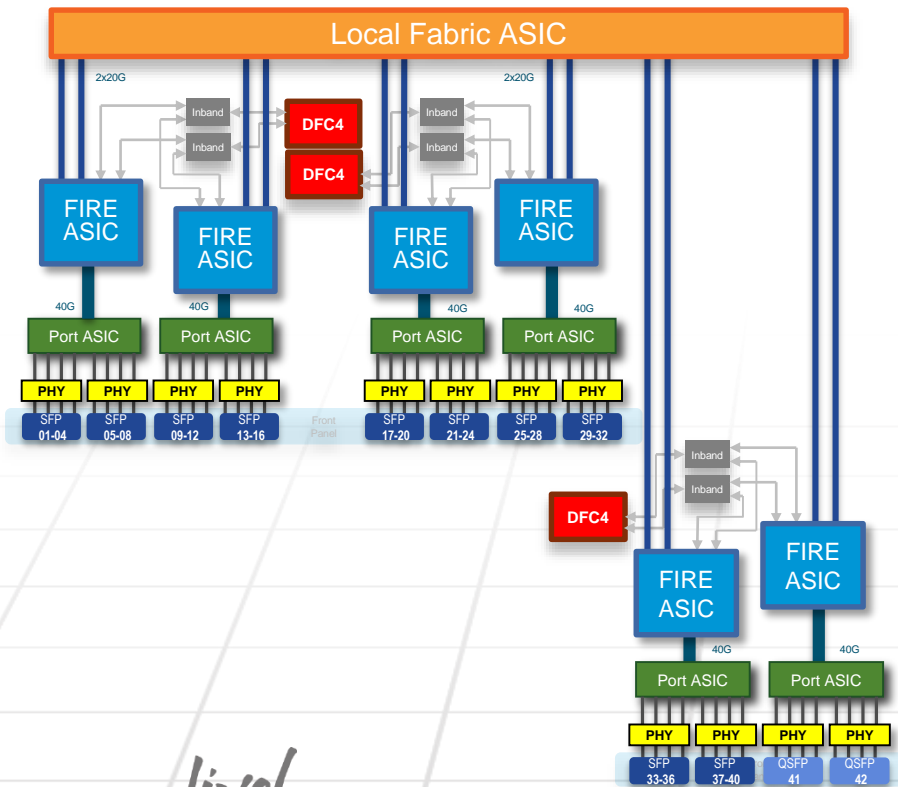
# Catalyst 6840-X

C6840-X-LE-40G – 40 x SFP+ & 2 x QSFP



**NEW**

FIRE = Fabric Interface & Replication Engine




## Essentially same as 6880-X Baseboard

- **40 ports of 1/10G SFP+, 2 ports of 40G QSFP**
  - 4 x 8 SFP+ Port-Groups, 1 x 2 QSFP Port-Group
- **3 x DFC4-E Forwarding Engines**
- **80Gbps to Switch Fabric (2 Modes)**
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- **New 40Gbps Fabric/Replication ASIC**
- **Combines FIRE ASIC & MUX FPGA of 6880-X**
- **DDR3 Packet Buffers on FIRE ASIC**
  - 250 or 500MB TX per Port (2GB per FIRE ASIC)

CiscoLive!

# Agenda

- ❖ Chassis & Power
- ❖ Supervisor Architectures
- ❖ **Module Architectures**
  - ❖ 6700 & 6800 Series
  - ❖ 6900 Series
  - ❖ Integrated Service Modules
  - ❖ C6800 1G & 10G Series
  - ❖ C6880-X Series
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks



# Catalyst 6500-E & 6807-XL Line Cards



## 6700 & 6800 Series

with CFC or DFC4



## 6900 Series

with DFC4



## New C6800 1G Series

with DFC4



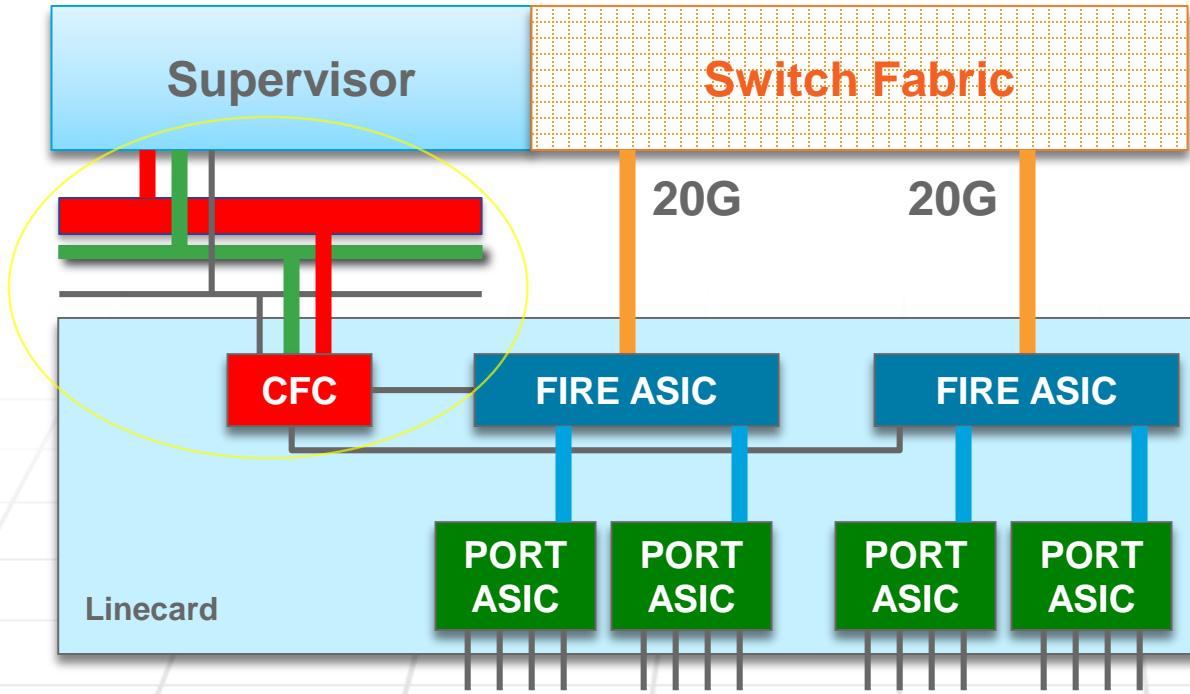
## New C6800 10G Series

with DFC4

Cisco *live!*



# 6700 Series (with CFC)



The 6700 modules connect to the BUS using a Centralized Forwarding Card (CFC)

They connect to the Switch Fabric via Dual **20G** Fabric channels

**NOTE:** The DBUS is NOT used. Its ONLY for control traffic


# Centralized Forwarding Card (CFC)

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

The CFC comes on legacy 6700 modules to provide 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...



# CFC Centralized Forwarding Card

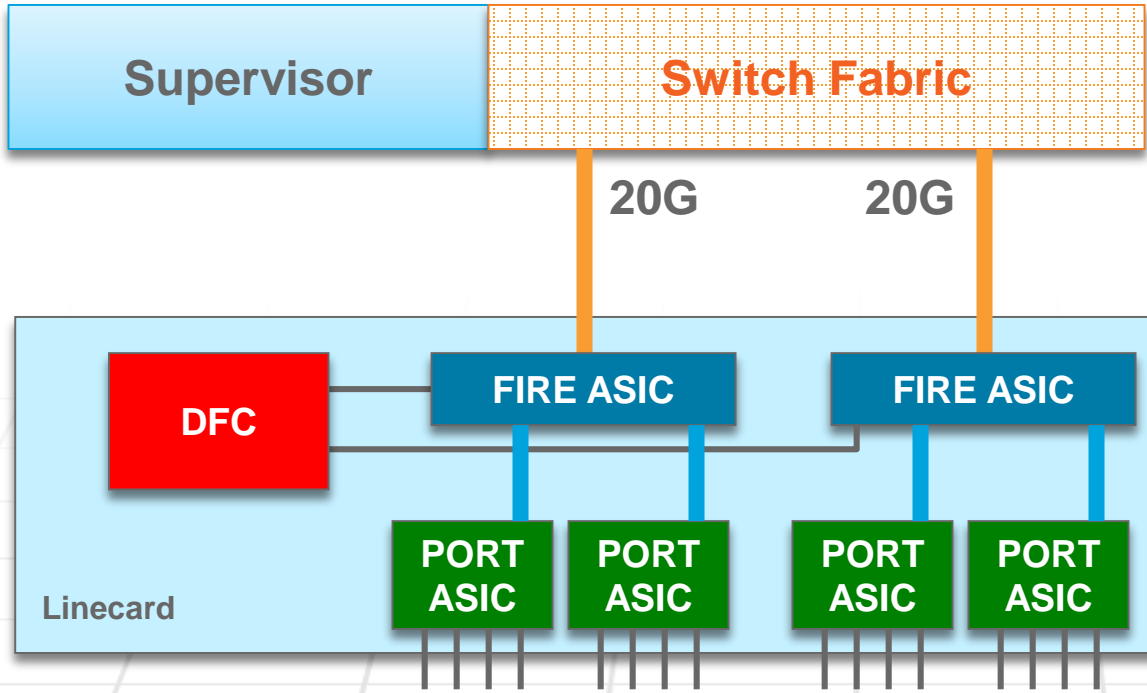
---

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

---



# 6700 & 6800 Series (with DFC)



The **6700** Series supports a **Distributed Forwarding Card (DFC)**

The **6800** Series use preinstalled **Distributed Forwarding Card 4**

They connect to the Switch Fabric via **Dual 20G Channels**

# Distributed Forwarding Card 3

The DFC3 is an option for 6500 & 6700 Series, and pre-installed on 6708 & 6716 Cards

The DFC3 stores a local copy of the L2/L3+ forwarding info, as well as Netflow, Security & QoS ACL's



The DFC3 supports forwarding rates up to **48Mpps**

The DFC3 offloads forward lookup work from the PFC, increasing performance & reducing latency


**Three versions** of the DFC3 are supported...

- DFC3A
- DFC3B (3BXL)
- DFC3C (3CXL)

# Distributed Forwarding Card 4

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

Each DFC4 stores a Local Copy of the L2/L3+ forwarding info, as well as NetFlow, Security & QoS ACL's



The **DFC4** supports local forwarding rates up to **60Mpps**

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

**Two models** of the DFC4 are supported:

- DFC4-A (XL) for 1G
- DFC4-E (XL) for 10G+

# DFC = Distributed Forwarding Card

---

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

---





# Catalyst 6500/6800 Modules

DFC3/4 interoperability with PFC3/4



For Your  
Information


	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 DFC3A	Operate as DFC3B		Operate as PFC3B and DFC3B	Operate as PFC3BXL	X	X
DFC3C	Operate as DFC3A	Operate as DFC3B	Operate as PFC3B and DFC3B		Operate as PFC3C	X	X
DFC3CXL	Operate as DFC3A	Operate as DFC3B	Operate as DFC3BXL	Operate as DFC3C		X	X
DFC4	X	X	X	X	X		Operates as PFC4
DFC4XL	X	X	X	X	X	Operates as DFC4	

# WS-X6704-10G

## 6700 Series Line Cards



- **4 ports 10G Xenpak**
- **Supports CFC or DFC4-A**
- **Up to 60Mpps with DFC installed**
- **2 x 20G Channels to Switch Fabric**
- **Connection to the Shared Bus**
- 16MB Packet Buffers per port
- Egress Multicast Replication
- Supports Strict Priority queue on TX
- Supports 2 receive queues per port
- Supports 8 transmit queues per port
- Supports Weighted Round Robin



Cisco *live!*

# WS-X6748-GETX and 6848-TX

6700 and 6800 Series Line Cards



- **48 ports 10/100/1000M RJ45**
- **Supports CFC or DFC4-A**
- **Up to 60Mpps with DFC installed**
- **2 x 20G Channels to Switch Fabric**
- **Connection to the Shared Bus**

- 1.2Mb Packet Buffers per port
- Egress Multicast Replication
- Supports Strict Priority queue on TX
- Supports 2 receive queues per port
- Supports 4 transmit queues per port
- Supports Weighted Round Robin




Cisco *live!*

# WS-X6748-SFP and 6848-SFP

6700 and 6800 Series Line Cards



- **48 ports 1G SFP**
- **Supports CFC or DFC4-A**
- **Up to 60Mpps with DFC installed**
- **2 x 20G Channels to Switch Fabric**
- **Connection to the Shared Bus**
- 1.2Mb Packet Buffers per port
- Egress Multicast Replication
- Supports Strict Priority queue on TX
- Supports 2 receive queues per port
- Supports 4 transmit queues per port
- Supports Weighted Round Robin




Cisco *live!*

# WS-X6716-10G and 6816-10G

## 6700 and 6800 Series Line Cards



- 16 ports 10GE (X2) in MUX mode
- 4 Programmable “port-groups” in 2 modes
- **Requires integrated DFC4-E**
- **Up to 60Mpps with DFC installed**
- **2 x 20G Channels to the Switch Fabric**
- VSL support on 4 ports in Transparent mode
- Up to 256MB Packet Buffers per-port
- Egress Multicast Replication
- Supports Strict Priority Queue on TX
- Supports 2 receive queues per port
- Supports 8 transmit queues per port
- Supports Weighted Round Robin & Shaped Round Robin






# WS-X6716-10T and 6816-10T

6700 and 6800 Series Line Cards



- 16 ports 10GE (RJ45) in MUX mode
- 4 Programmable “port-groups” in 2 modes
- **Requires integrated DFC4-E**
- **Up to 60Mpps with DFC installed**
- **2 x 20G Channels to the Switch Fabric**
- VSL support on 4 ports in Transparent mode
- Up to 256MB Packet Buffers per-port
- Egress Multicast Replication
- Supports Strict Priority Queue on TX
- Supports 2 receive queues per port
- Supports 8 transmit queues per port
- Supports Weighted Round Robin & Shaped Round Robin



Cisco *live!*

# WS-X6908-10G

## 6900 Series Line Cards



- **8 ports 10GE (X2 based)**
- **Pre-Installed with DFC4-E**
- **Up to 60Mpps with DFC installed**
- **2 x 40G Channels to Switch Fabric**
- Supports Cisco TrustSec (CTS)
- L2/L3 Encryption on all ports
- Supports VSL on all ports

- Up to 256MB Packet Buffers per port
- Egress Multicast Replication
- Supports Strict Priority Queue on TX
- Supports 2 receive queues per port
- Supports 8 transmit queues per port
- Supports Deficit Weighted Round Robin & Shaped Round Robin



X2-10G-SR/LR



CVR-SFP10G  
(OneX)

Cisco *live!*

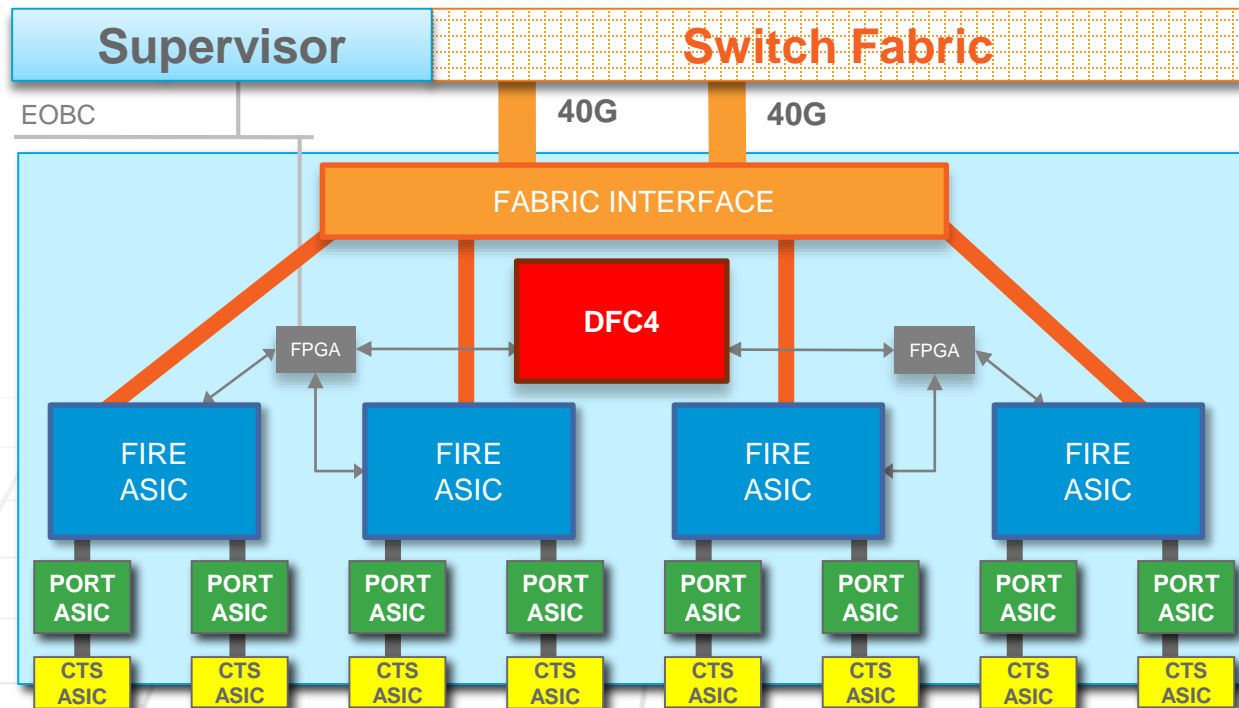


# 6900 Series Cards

WS-X6908-10G



FIRE = Fabric Interface & Replication Engine



The **6900** Series use preinstalled **Distributed Forwarding Card 4**

They connect to the Switch Fabric via **Dual 40G Channels**

They come with a **Cisco Trust Security (CTS) ASIC** built in

*CiscoLive!*

# WS-X6904-40G

## 6900 Series Line Cards




- **4 ports 40GE (CFP)**
- **16 ports 10GE (FourX & SFP+)**
- **Pre-Installed with DFC4-E**
- **Up to 60Mpps with DFC installed**
- **2 x 40G Connections to Switch Fabric**
- Supports Cisco TrustSec (CTS)
- L2/L3 Encryption on all ports
- Supports VSL & IA on all ports

- Up to 10MB Packet Buffers per 10G port
- Up to 40MB Packet Buffers per 40G port
- Egress Multicast Replication
- Supports Dual Strict Priority Queue on TX
- Supports 2 receive queues per port
- Supports 8 transmit queues per port
- Supports Deficit Weighted Round Robin & Shaped Round Robin



CFP-40G-LR4



CFP-40G-SR4

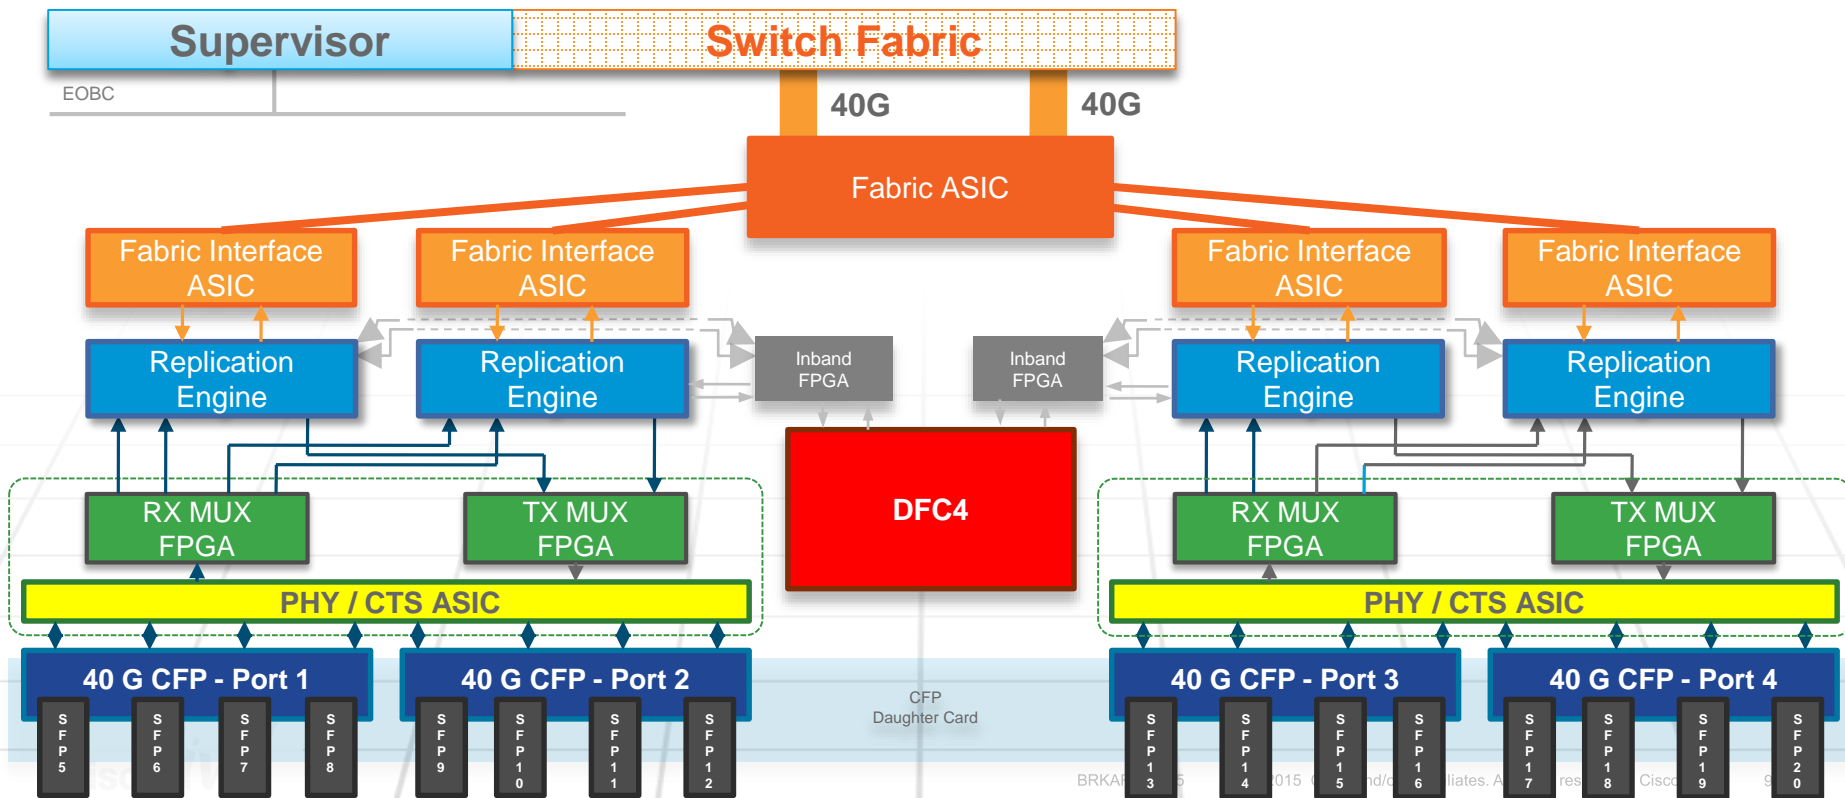


CVR-4SFP10G  
(FourX)

Cisco *live!*


# 6900 Series Cards

The 6900 Series supports pre-installed DFC4, Dual 40G Fabric channels and Built-In CTS ASIC...



# Agenda

- ❖ Chassis & Power
- ❖ Supervisor Architectures
- ❖ **Module Architectures**
  - ❖ 6700 & 6800 Series
  - ❖ 6900 Series
  - ❖ **Integrated Service Modules**
  - ❖ C6800 1G & 10G Series
  - ❖ C6880-X Series
  - ❖ C6840-X Series
- ❖ 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 per WISM2
Concurrent AP Upgrades	Up to 500
Mobility, Domain Size	Up to 18,000 APs

Enhance Visibility,  
Accelerate Troubleshooting

NMS



### Next Generation Network Analysis: **NAM3**

Monitoring Performance	16 Gbps
Capture to External Disk	6 Gbps
Performance Analytics	1588 Timestamps
Hardware Filters & Packet Captures	SPAN, FnF, SNMP Port Monitoring

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

Deliver Robust,  
Integrated, Streamlined Security

SEC

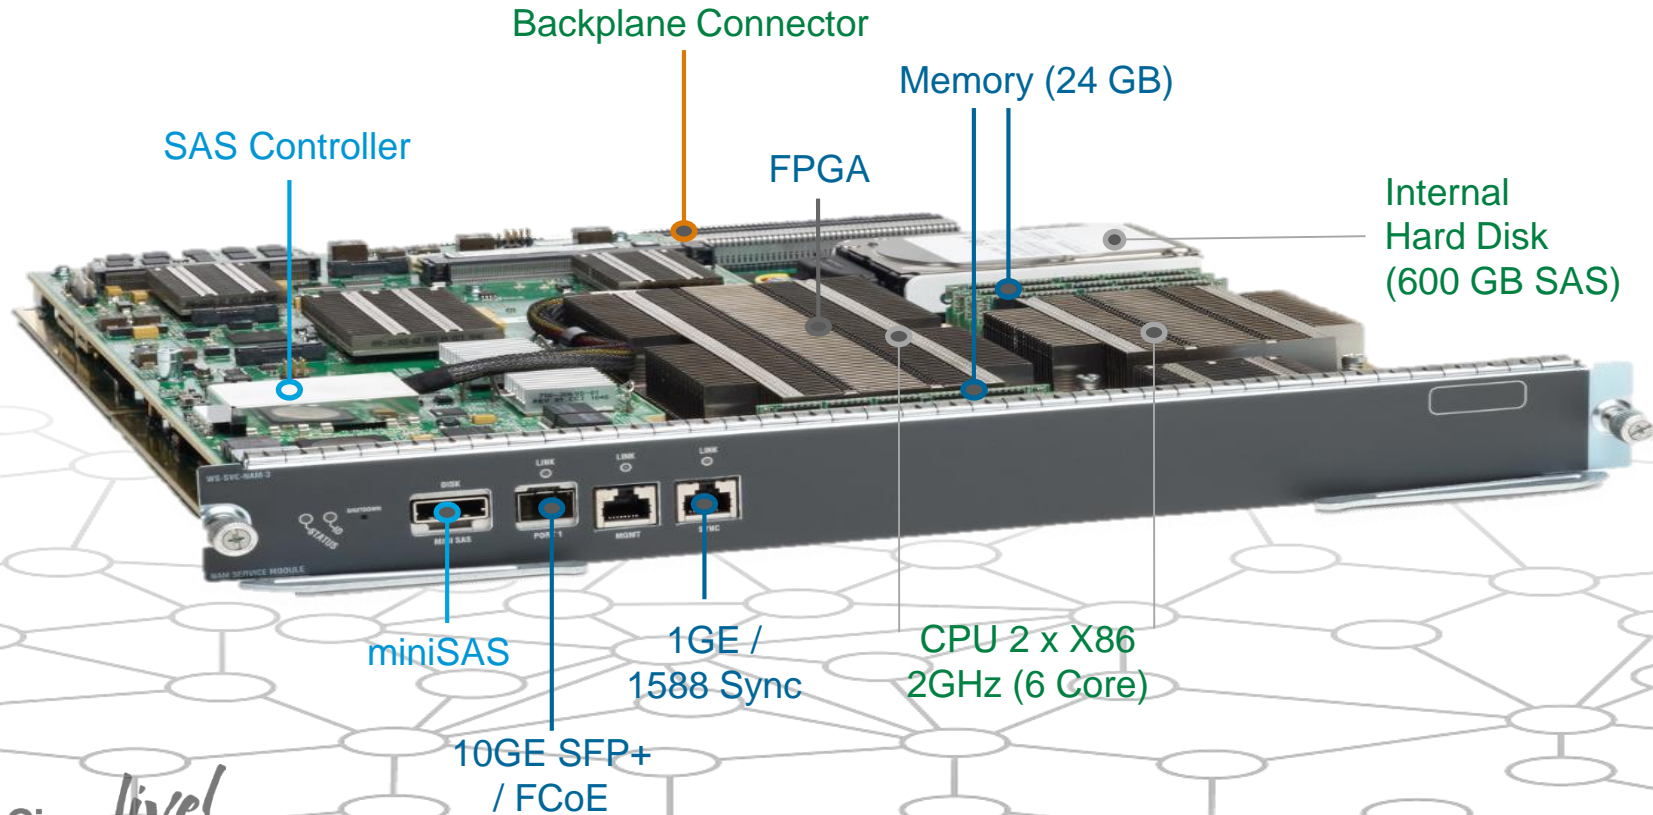


### Next Generation Firewall & DPI: **ASA-SM**

64 Gbps	System Performance
16 Gbps	Performance per ASA-SM
10,000,000	Concurrent Sessions
300,000	Connections per Second
250	Security Contexts
1,000	VLANs

# Network Analysis with NAM-3

WS-SVC-NAM3-6G-K9



Cisco *live!*




# Wireless Services with WISM-2

## WS-SVC-WISM2-K9

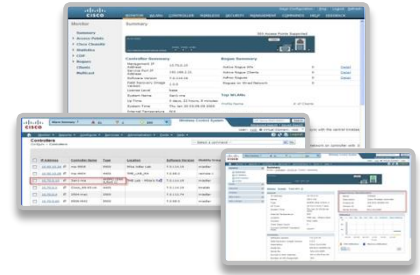
One device for converged Wireless and Wired Services supporting next-generation wiring closet infrastructures...

### Reduced Operational Costs

- **Scalability**
  - 1000 Access Points
  - 15,000 Clients
- **Central Maintenance**
  - Troubleshooting
  - Simultaneous AP Upgrade
- **Wireless Mobility**
  - 36,000 AP in Domain
  - Fast Roaming
- **Performance**
  - 16+ Gbps Throughput
- **New Features**
  - Wireless AVC
  - Flexible NetFlow
  - Bonjour Gateway
  - NMSP Location Services
  - Stateful AP Failover with VSS



### WLC GUI TOOLS



PRIME



ISE

Cisco *live!*



# Firewall & VPN Services with ASA-SM

WS-SVC-ASA-SM1-K9

**NAT64, IPSec,  
VPN Site-to-Site  
Services**

## Multi-Gigabit Fabric

- 20G Connection
- Virtualized interfaces
- Module-to-module communications

## Multiple Contexts (250)

- High Capacity
- Memory for high session counts
- 24 GB of memory

## Dual-Crypto Accelerators

- Hardware Processing
- Accelerated Virtual Private Networking
- Unified Communications Encryption


## Security Service Processors

- Multi-Services Capable
- Dedicated 64-bit multi-core Processors
- Future-proof Hardware

**Cisco** live!

# Agenda

- ❖ Chassis & Power
- ❖ Supervisor Architectures
- ❖ **Module Architectures**
  - ❖ 6700 & 6800 Series
  - ❖ 6900 Series
  - ❖ Integrated Service Modules
  - ❖ **C6800 1G & 10G Series**
  - ❖ C6880-X Series
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ Packet Walks




# C6800-48P-TX

Updated 1GE Line Cards

**40G**  
per slot

Supported With  
**15.2(1)SY**

- **48 ports of RJ45**
- **40Gbps Capacity per Slot**
- **Single DFC4 for 60Mpps**
- **Integrated DFC-4A (XL)**
- **Updated Port ASICs**
- **1.5MB per Port Buffers**
- **New Hardware RFID**

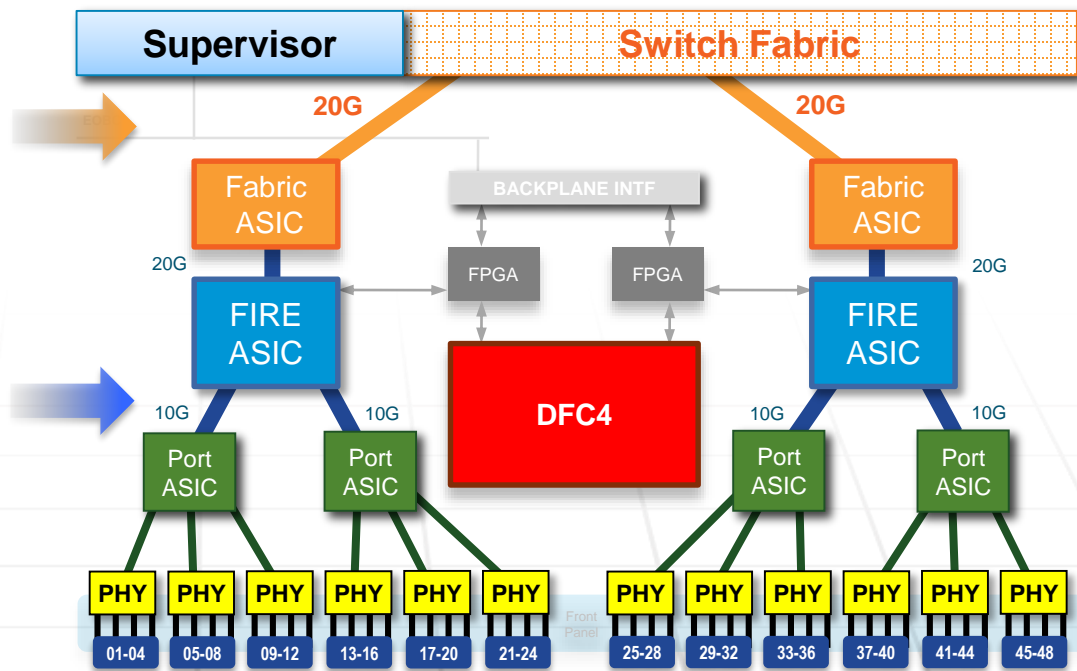


Based on  
**WS-X6848-TX**

*Cisco*live!

## Block Diagram

- **12 x 1G Ports per Port ASIC**
- **DFC4-A supports 48 x 1G Ports**
- **Improved FIRE & Fabric ASICs**
  - **20Gbps per Fabric Channel**
  - **20Gbps per FIRE ASIC**
  - **1.2:1 Oversubscribed @ Port ASIC**
- **New Port ASIC with more Buffers**
  - **1.5MB RX and 1.5MB TX**



# C6800-48P-SFP


Updated 1GE Line Cards

40G

per slot

Supported With  
15.2(1)SY

- 48 ports of SFP
- 40Gbps Capacity per Slot
- Single DFC4 for 60Mpps
- Integrated DFC-4A (XL)
- Updated Port ASICs
- 1.5MB per Port Buffers
- New Hardware RFID



Based on  
WS-X6848-SFP

Cisco *live!*

## Block Diagram

**FIRE – Fabric Interface & Replication Engine**

- 
- The diagram illustrates the architecture of a 40Gbps Ethernet switch. At the top, the **Supervisor** (blue box) and **Switch Fabric** (orange dotted box) are connected. The Supervisor is linked to two **Fabric ASIC** (orange boxes) via 20G connections. The Switch Fabric is also connected to these Fabric ASICs via 20G connections. Each Fabric ASIC is connected to a **FIRE ASIC** (blue boxes) via 20G connections. The FIRE ASICs are connected to **Port ASIC** (green boxes) via 10G connections. The Port ASICs are connected to **PHY** (yellow boxes) via 10G connections. The PHYs are organized into two groups: **ODD** (01-07, 09-15, 17-23, 25-31, 33-39, 41-47) and **EVEN** (02-08, 10-16, 18-24, 26-32, 34-40, 42-48). A central **DFC4** (red box) is connected to the FIRE ASICs via **BACKPLANE INTF** (grey box) and **FPGA** (grey boxes). A **Front Panel** (light blue box) is located between the two PHY groups.



# C6800-32P10G

New High Density Multi-Rate Line Cards

**160G**  
per slot

Supported With  
**15.2(1)SY**

- **32 ports of SFP / SFP+**
- **Up to 8 Ports of QSFP\***
- **160Gbps Capacity per Slot**
- **Dual DFC4 for 120Mpps**
- **Integrated DFC-4E (XL) x 2**
- Support for 1 / 10GE and 40GE\*
- 250 / 500MB Per Port TX Buffers
- New Integrated Hardware RFID
- VSL and IA capable on all ports



Based on  
**WS-X6904-40G**

**Cisco**live!

\* with new CVR-4SFP-QSFP adapter cable

BRKARC-3465

© 2015 Cisco and/or its affiliates. All rights reserved. Cisco Public

102

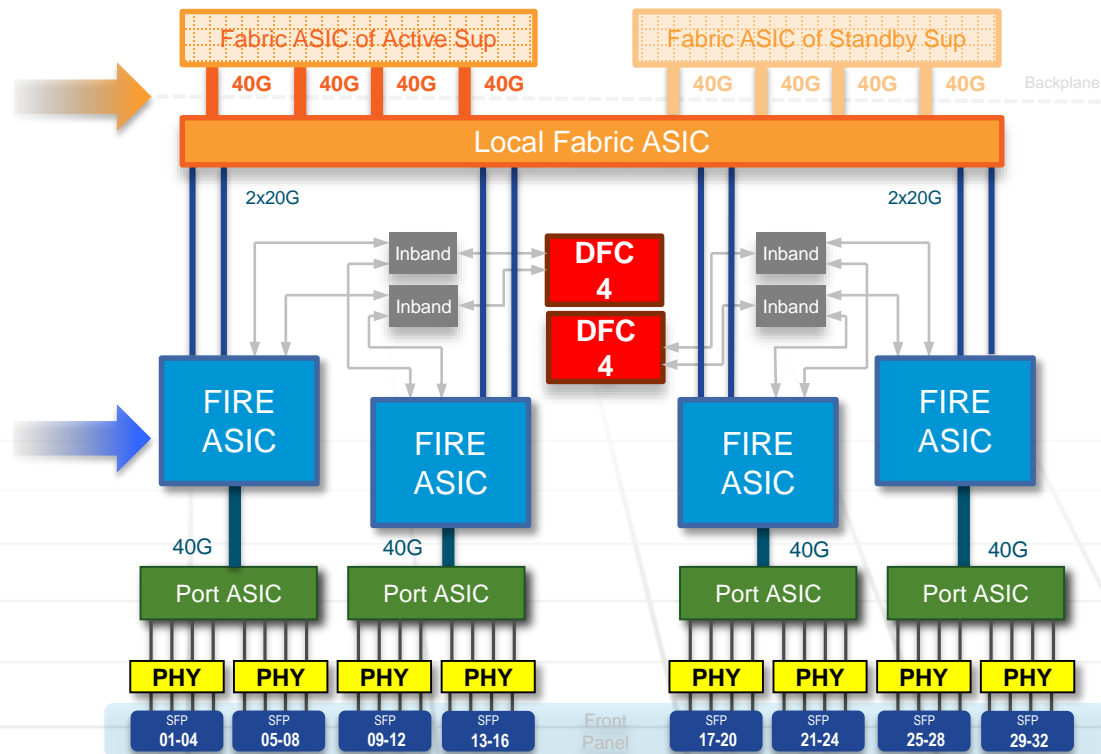


# C6800-32P10G

## Block Diagram

FIRE = Fabric Interface & Replication Engine

- **4 x 8 SFP+ Port-Groups**
- **DFC4-E supports 16 x SFP Ports**
- **New FIRE & Fabric ASICs**
  - **40Gbps per Fabric Channel**
  - **40Gbps per FIRE ASIC**
  - **2:1 Oversubscribed @ PORT ASIC**
- **Performance Mode per Port-Group**
- **1/10 or 40GE Mode per Port-Group\***
- **1.25 / 2.5MB RX and 250 / 500MB TX**




# C6800-16P10G

New High Density Multi-Rate Line Cards

- 16 ports of SFP / SFP+
- Up to 4 Ports of QSFP\*
- 80Gbps Capacity per Slot
- Single DFC4 for 60Mpps
- Integrated DFC-4E (XL)
- Support for 1 / 10GE and 40GE\*
- 250 / 500MB Per Port TX Buffers
- New Integrated Hardware RFID
- VSL and IA capable on all ports

80G  
per slot

Supported With  
15.2(1)SY



Based on  
WS-X6904-40G

Cisco *live!*

\* with new CVR-4SFP-QSFP adapter cable

BRKARC-3465

© 2015 Cisco and/or its affiliates. All rights reserved. Cisco Public

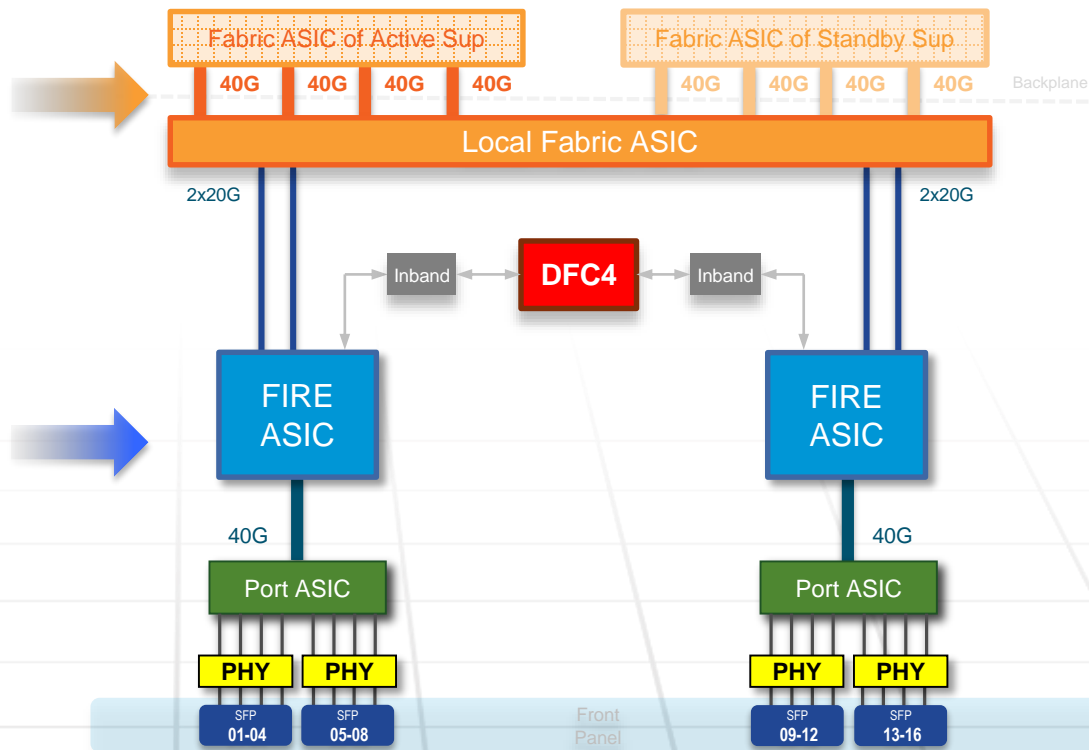
104

# C6800-16P10G

## Block Diagram

FIRE = Fabric Interface & Replication Engine

- **2 x 8 SFP+ Port-Groups**
- **DFC4-E supports 16 x SFP Ports**
- **New FIRE & Fabric ASICs**
  - 40Gbps per Fabric Channel
  - 40Gbps per FIRE ASIC
  - 2:1 Oversubscribed @ Port ASIC
- Performance Mode per Port-Group
- 1/10 or 40GE Mode per Port-Group\*
- 1.25 / 2.5MB RX and 250 / 500MB TX



# C6800-8P10G

New High Density Multi-Rate Line Cards

80G  
per slot

Supported With  
15.2(1)SY

- 8 ports of SFP / SFP+
- Up to 2 Ports of QSFP\*
- 80Gbps Capacity per Slot
- Single DFC4 for 60Mpps
- Integrated DFC-4E (XL)
- 500MB Per Port TX Buffers
- Support for 1 / 10GE and 40GE\*
- New Integrated Hardware RFID
- VSL and IA capable on all ports



Based on  
WS-X6904-40G

Cisco *live!*

\* with new CVR-4SFP-QSFP adapter cable

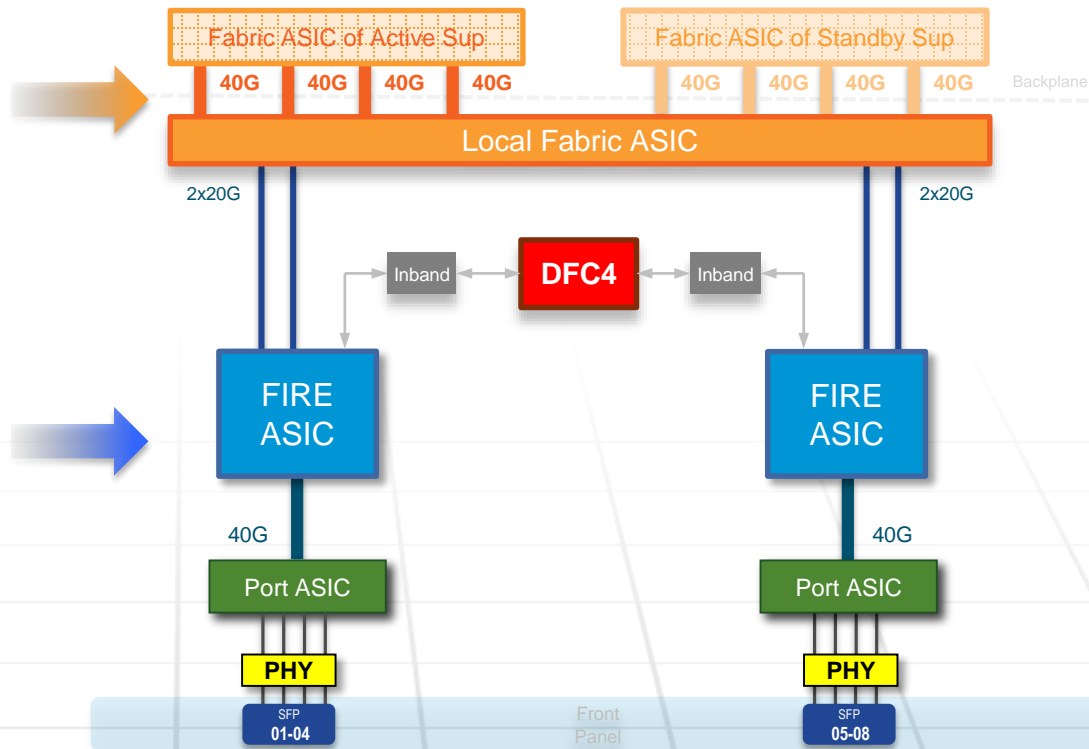
BRKARC-3465

© 2015 Cisco and/or its affiliates. All rights reserved. Cisco Public 106

# C6800-8P10G

## Block Diagram

- **2 x 4 SFP+ Port-Groups**
- **DFC4-E supports 8 x SFP Ports**
- **New FIRE & Fabric ASICs**
  - 40Gbps per Fabric Channel
  - 40Gbps per FIRE ASIC
  - Non Oversubscribed (1:1)
- No Need for Performance Mode
- 1/10 or 40GE Mode per Port-Group\*
- 2.5MB RX and 500MB TX





# Agenda

- ❖ Chassis & Power
- ❖ Supervisor Architectures
- ❖ **Module Architectures**
  - ❖ 6700 & 6800 Series
  - ❖ 6900 Series
  - ❖ Integrated Service Modules
  - ❖ C6800 1G & 10G Series
  - ❖ **C6880-X Series**
- ❖ 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

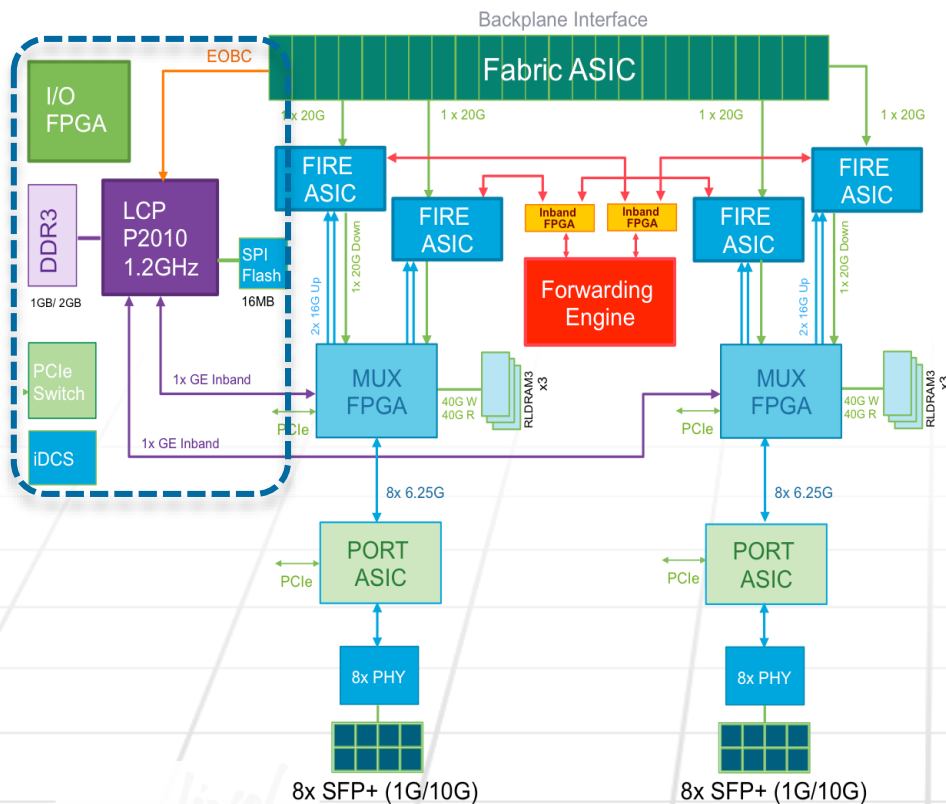
Port Speed & Type	Number of Ports
10/100/1000 Mb/s Copper	16 (GLC-T)
1 Gb/s Fiber	16 (SFP)
10 Gb/s Fiber	16 (SFP+)
40 Gb/s Fiber	4 (SFP-QSFP)

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



# Catalyst 6880-X

## C6880-X-16P10G Architecture




### Same as Baseboard + 1.2Ghz LCP

- **16 ports of 1/10G SFP+**
  - 2 x 8 SFP+ Port-Groups
- **Enhanced DFC4-E Forwarding Engine**
- **80Gbps to Switch Fabric (2 Modes)**
  - Performance (8P @ 1:1)
  - Oversubscribe (16P @ 2:1)
- **Improved 40Gbps Fabric/Replication ASIC**
- **New 40Gbps Port Interface MUX FPGA**
- **RLDRAM Packet Buffers on MUX & Port ASIC**
  - 1.25 or 2.5MB RX per Port (10MB per Port ASIC)
  - 24 or 48MB TX per Port (192MB per MUX FPGA)

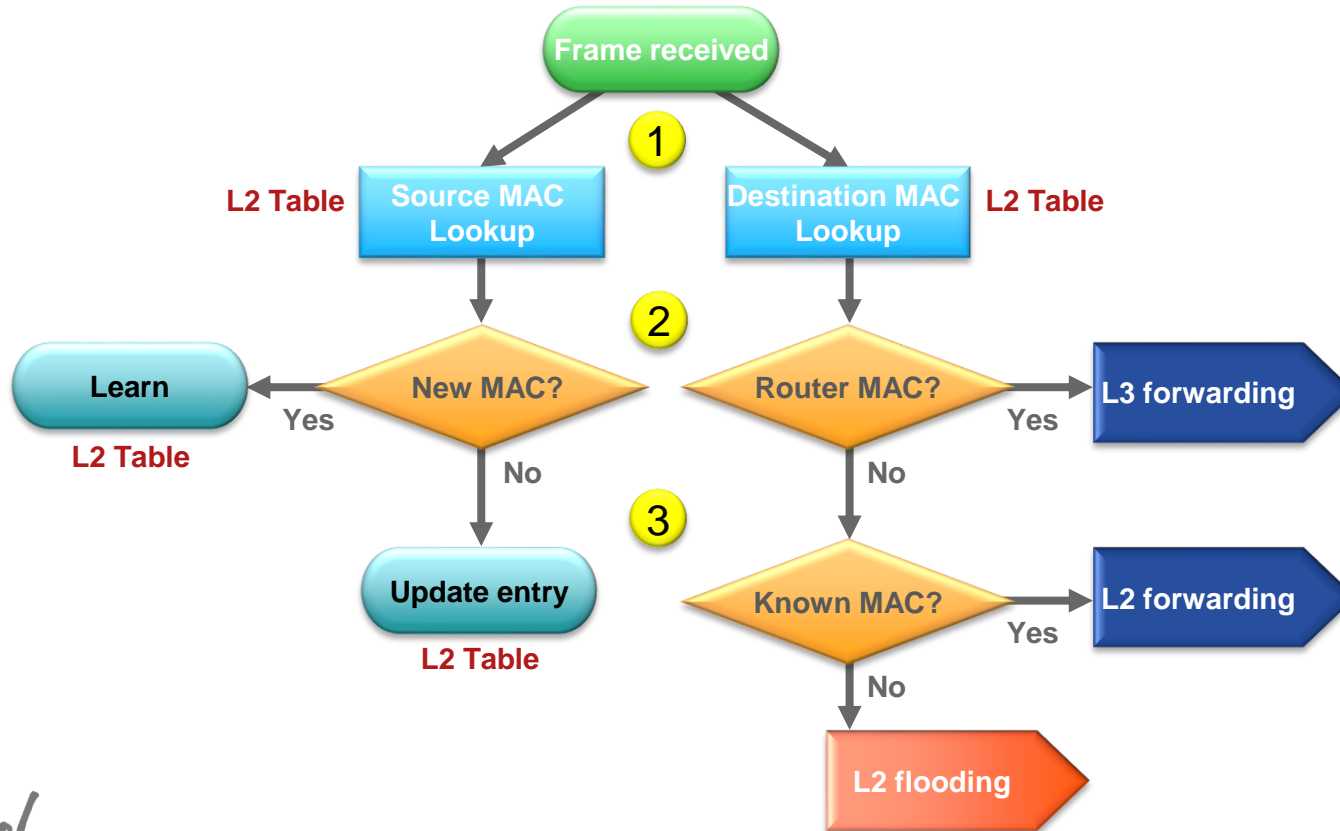
# Agenda

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



# Layer 2 Switching



## L2 Forwarding Steps



# Layer 2 Switching

## Layer 2 MAC Table

PFC / DFC has an integrated CAM with  
**N pages x 4096 rows =**  
**MAC address space**




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

**PFC4 = 128K Entries**  
**(32 x 4096)**

# Layer 2 Switching

## Layer 2 Forwarding Operation



1. **Hash result identifies the starting Page and Row in the MAC table**
2. **Lookup Key (VLAN + MAC) compared to contents of entries on each page (sequentially)**
  - **DST MAC Lookup:** Match returns Destination interface(s) OR Miss results in Flood
  - **SRC MAC Lookup:** Match updates age of current entry OR Miss installs new entry

# Displaying the Layer 2 Table

```
SUP2T#show mac address-table
```

Legend: \* - primary entry

age - seconds since last seen; n/a - not available; S - secure entry;


R - router's gateway mac address entry; D - Duplicate mac address entry

Displaying entries from active supervisor:

	vlan	mac address	type	learn	age	ports
*	192	00d0.0053.bc00	dynamic	Yes	5	Gi7/3
R	205	0024.c4dc.d740	static	No	-	Router
R	20	0024.c4dc.d740	static	No	-	Router
*	192	0014.5e31.4220	dynamic	Yes	65	Gi7/3
*	60	00d0.2bfc.23f5	dynamic	Yes	30	Gi5/14
*	192	00e0.1e5d.e9ff	dynamic	Yes	30	Gi7/3
...						

# Catalyst 6500/6800

## Understanding EtherChannels




- Combines **Multiple** physical ports into **One** logical port
- Deterministic **Hash-based** Channel **Load-Balancing**
- **Configurable** Hash uses **SRC, DST, L2** and / or **L3**
- **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 Load-Balancing


## PFC3 Flow Distribution



Channel Bundle	Link1	Link2	Link3	Link4	Link5	Link6	Link7	Link8
2 Links	50%	50%	--	--	--	--	--	--
3 Links	37.5%	37.5%	25%	--	--	--	--	--
4 Links	25%	25%	25%	25%	--	--	--	--
5 Links	25%	25%	25%	12.5%	12.5%	--	--	--
6 Links	25%	25%	12.5%	12.5%	12.5%	12.5%	--	--
7 Links	25%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	--
8 Links	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%

# EtherChannel Load-Balancing

## PFC4 Flow Distribution



Channel Bundle	Link1	Link2	Link3	Link4	Link5	Link6	Link7	Link8
2 Links	50%	50%	--	--	--	--	--	--
3 Links	33.6%	33.2%	33.2%	--	--	--	--	--
4 Links	25%	25%	25%	25%	--	--	--	--
5 Links	20.4%	19.9%	19.9%	19.9%	19.9%	--	--	--
6 Links	16.8%	16.8%	16.8%	16.8%	16.4%	16.4%	--	--
7 Links	14.5%	14.5%	14.5%	14.5%	14%	14%	14%	--
8 Links	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%	12.5%


# PFC4 EtherChannel Inputs

EtherChannel Uses a Hash Algorithm to Determine which Link in the Bundle to Use -  
The Inputs to the Algorithm Are a Combination of L2, L3 and / or L4 addresses

<code>dst-ip</code>	Dst IP Addr
<code>dst-mac</code>	Dst Mac Addr
<b><code>dst-mixed-ip-port</code></b>	<b>Dst IP Addr and TCP/UDP Port</b>
<code>dst-port</code>	Dst TCP/UDP Port
<code>mpls</code>	Load Balancing for MPLS packets
<code>src-dst-ip</code>	Src XOR Dst IP Addr
<code>src-dst-mac</code>	Src XOR Dst Mac Addr
<b><code>src-dst-mixed-ip-port</code></b>	<b>Src XOR Dst IP Addr and TCP/UDP Port</b>
<code>src-dst-port</code>	Src XOR Dst TCP/UDP Port
<code>src-ip</code>	Src IP Addr
<code>src-mac</code>	Src Mac Addr
<b><code>src-mixed-ip-port</code></b>	<b>Src IP Addr and TCP/UDP Port</b>
<code>src-port</code>	Src TCP/UDP Port
<code>vlan-dst-ip</code>	Vlan, Dst IP Addr
<b><code>vlan-dst-mixed-ip-port</code></b>	<b>Vlan, Dst IP Addr and TCP/UDP Port</b>
<code>vlan-src-dst-ip</code>	Vlan, Src XOR Dst IP Addr
<b><code>vlan-src-dst-mixed-ip-port</code></b>	<b>Vlan, Src XOR Dst IP Addr and TCP/UDP Port</b>
<code>vlan-src-ip</code>	Vlan, Src IP Addr
<b><code>vlan-src-mixed-ip-port</code></b>	<b>Vlan Src IP Addr and TCP/UDP Port</b>

# Agenda

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



# Interface Management

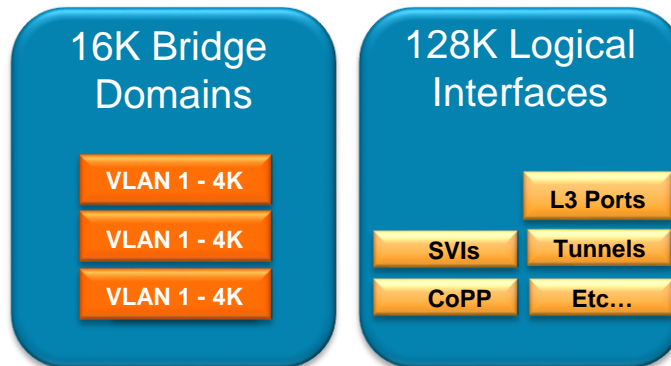
## VLANs & Interfaces

### Supervisor 720



- VLANs used for both L2 Bridging and L3 Routing
- Each L3 Interfaces consumes an internal VLANs from total 4K VLAN pool

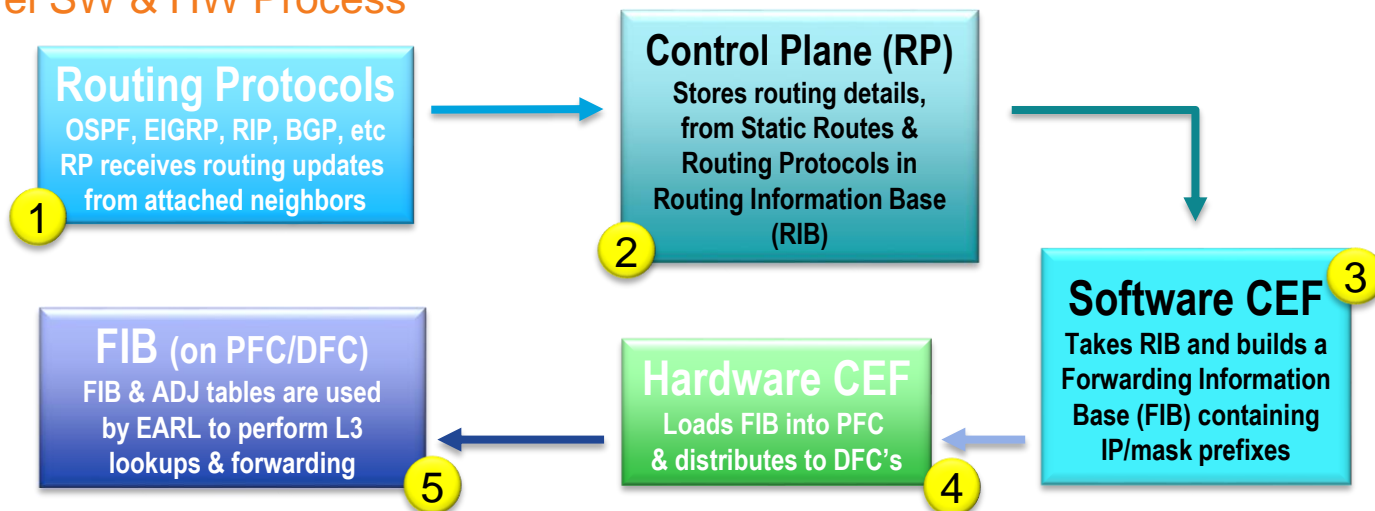
### Supervisor 2T



- Separate L2 Bridging and L3 Routing
- Breaks the 4K VLAN barrier
- Allows VLAN reuse on Per Port basis
- Massive scale for L3 interfaces

# L3 Forwarding

## High-Level SW & HW Process



### 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)

# L3 Forwarding

## FIB & Adjacency Overview

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

### FIB contains:

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

#### FIB TCAM shared by:

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

### Adjacency Table:

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

172.20.45.1
10.1.1.100
MASK (/32)
...
10.1.3.0
10.1.2.0
MASK (/24)
...
10.1.0.0
172.16.0.0
MASK (/16)
...
0.0.0.0
MASK (/0)

**FIB TCAM**

IF1, MAC, MTU
IF2, MAC, MTU
IF3, MAC, MTU
IF4, MAC, MTU
...
...


**Adjacency Table**



# L3 Forwarding

## FIB & ADJ Lookup in PFC/DFC

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

## Defaults and Changes

IPv4 Unicast & MPLS require only 1 entry

IPv6 Unicast & IPv4 Multicast require 2 entries

**Standard PFC/DFC = 256K entries**

**“XL” PFC/DFC = 1M entries**

Default TCAM allocation shown below

	NON-XL PFC	XL PFC
IPv4, MPLS	192k	512k
IPv6, Multicast	32k	256k

### 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

```
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
<b>75</b>	<b>10.10.1.1/32</b>	<b>receive</b>
<b>76</b>	<b>10.10.1.0/32</b>	<b>receive</b>
<b>77</b>	<b>10.10.1.255/32</b>	<b>receive</b>
78	10.10.1.2/32	Gi1/1, 0030.f272.31fe
3200	224.0.0.0/24	receive
<b>3201</b>	<b>10.10.1.0/24</b>	<b>glean</b>
3202	10.100.0.0/24	Gi1/1, 0030.f272.31fe
3203	10.100.1.0/24	Gi1/1, 0030.f272.31fe
<b>3204</b>	<b>10.100.2.0/24</b>	<b>Gi1/1, 0030.f272.31fe</b>
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	V1192, 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




# Catalyst Hardware NetFlow

Netflow is a process designed to collect information about traffic flows that pass through a switch

● **Data Flow (PFC)**

● **Exported Netflow Record (MSFC)**





# Catalyst 6500 NetFlow

## NetFlow Flow Masks

### Full Interface Mask

INTERFACE	SRC IP	DEST IP	IP PROTOCOL	SRC PORT	DEST PORT
-----------	--------	---------	-------------	----------	-----------

### Full Mask

INTERFACE	SRC IP	DEST IP	IP PROTOCOL	SRC PORT	DEST PORT
-----------	--------	---------	-------------	----------	-----------

### Destination-Source Interface Mask

INTERFACE	SRC IP	DEST IP	IP PROTOCOL	SRC PORT	DEST PORT
-----------	--------	---------	-------------	----------	-----------

### Source-Only Mask

INTERFACE	SRC IP	DEST IP	IP PROTOCOL	SRC PORT	DEST PORT
-----------	--------	---------	-------------	----------	-----------

### Destination-Only Mask

INTERFACE	SRC IP	DEST IP	IP PROTOCOL	SRC PORT	DEST PORT
-----------	--------	---------	-------------	----------	-----------

### Destination-Source Mask


INTERFACE	SRC IP	DEST IP	IP PROTOCOL	SRC PORT	DEST PORT
-----------	--------	---------	-------------	----------	-----------

# 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

## TCAM Lookup on PFC4




# Hardware NetFlow

Sup2T Flexible NetFlow



## Flow Record



**Key Fields** trigger the creation of a new Flow entry every time their value change

**Non-Key Fields** are data that is indexed by the Key Fields.

Key Fields are defined using the “match” statement  
Non-Key-Fields are defined using the “collect” statement

## Flow Export



Multiple Exporters can be associated with a single FNF monitor

## Flow Profile

## Flow Monitor



## Interfaces

Same Flow Monitor can be associated with multiple Interfaces

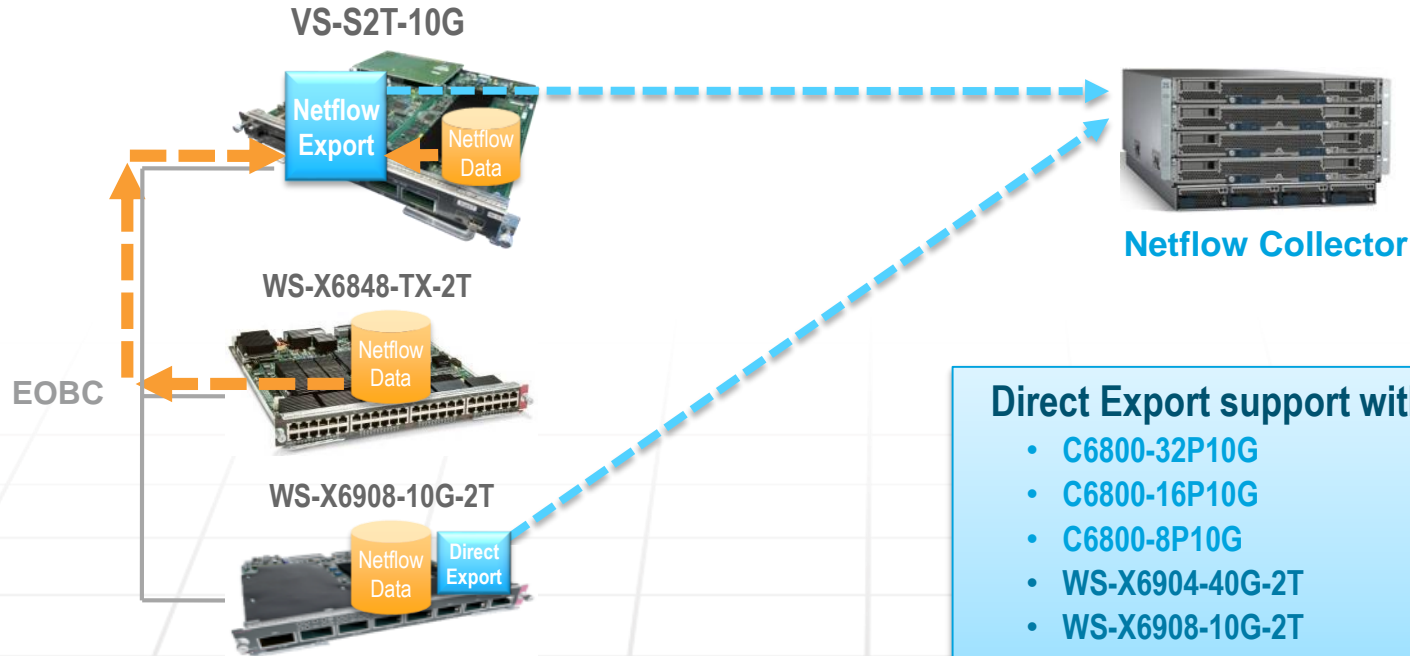


# PFC4 Key & Non-Key Fields

Field Name	Key	Field Name	Key	Field Name	Key
datalink vlan input	Yes	flow direction	Yes	ipv6 source address	Yes
ipv4 tos	Yes	flow cts source group-tag	Yes	ipv6 source prefix	No
ipv4 precedence	Yes	flow cts destination group-tag	Yes	ipv6 source mask	No
ipv4 protocol	Yes	interface input (snmp)	No	ipv6 destination address	Yes
ipv4 source address	Yes	interface output (snmp)	No	ipv6 destination prefix	No
ipv4 source prefix	No	routing destination as	No	ipv6 destination mask	No
ipv4 source mask	No	routing forwarding-status	Yes	transport source-port	Yes
ipv4 destination address	Yes	routing next-hop address ipv4 (bgp)	Yes	transport tcp flags	No
ipv4 destination prefix	No	routing next-hop address ipv6 (bgp)	Yes	counter Byte	No
ipv4 destination mask	No	routing source as	No	counter packet	No
ipv6 traffic-class	Yes	routing vrf input	Yes	timestamp sys-uptime first	No
ipv6 protocol	Yes	transport destination-port	Yes	timestamp sys-uptime last	No
input Interface	Yes	output Interface	Yes	input physical interface	Yes

# Hardware NetFlow

## NetFlow Data Export Process

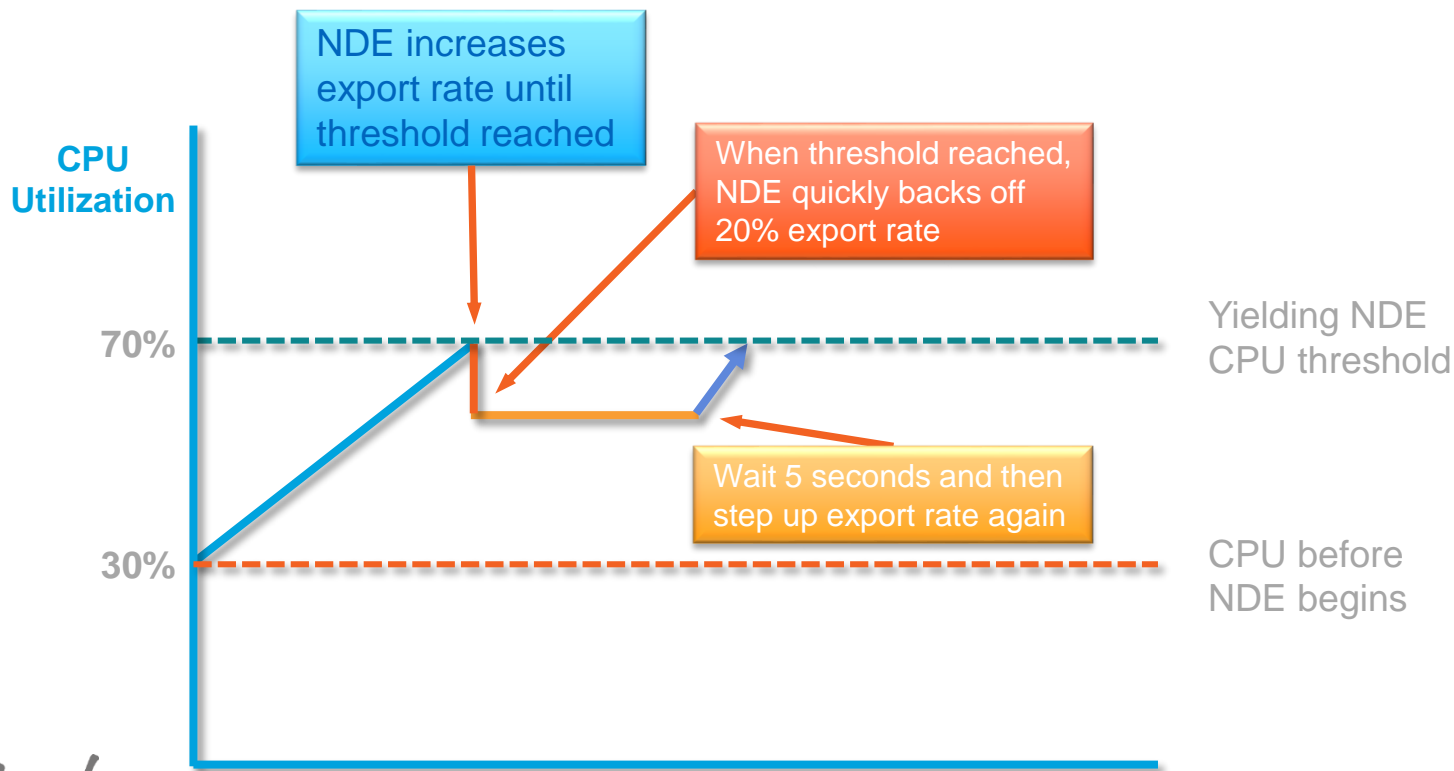


### Direct Export support with Sup2T:

- C6800-32P10G
- C6800-16P10G
- C6800-8P10G
- WS-X6904-40G-2T
- WS-X6908-10G-2T
- WS-X6816-10X-2T
- WS-X6716-10X (DFC4-E)

# Hardware NetFlow

## Sup2T "CPU Friendly" Netflow Export





# Displaying NetFlow Utilization

```
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%

# Sup720 vs. Sup2T- Flexible Netflow

Feature	Sup720	Sup2T
Netflow Entries	256 K (Ingress Only)	512 K Ingress – 512 K Egress
Shared Netflow Policers	N/A	512
Netflow Samplers	N/A	1K
Class Maps per Policy Map	1K	4K
Aggregate Policers	1K	6K
Distributed Policers	N/A	4K
Microflow Policers	63	128
Flexible Netflow	N/A	Yes
Egress Netflow	N/A	Yes
VRF-aware Netflow	N/A	Yes*
Netflow Export Enhancements	Direct Export	Direct Export, CPU Yield, EEM Event Manager
Packet or Byte Policing	Byte-based	Packet or Byte-based

# 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

Create an ACL or classification policy, using the CLI or Network Management System (NMS)

1

```
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
```



### Full Hardware Support Policy Feature Card Distributed Forwarding Card

- Router ACLs
- VLAN ACLs
- Port Based ACLs
- Role Based ACLs

2



### Hardware- Assisted ACL Features

- NetFlow
- NAT & PAT
- PBR
- WCCP
- Reflexive ACLs
- Cisco Trust Sec

3

# 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)

Used to permit or deny the movement of traffic between Layer 3 Subnets

Applied as an input or output policy to a Layer 3 interface

### VLAN ACL (VACL)

Used to permit or deny the movement of traffic between Layer 3 Subnets & VLANs or within a VLAN

Applied as a policy to a VLAN - is inherently applied to both inbound and outbound traffic

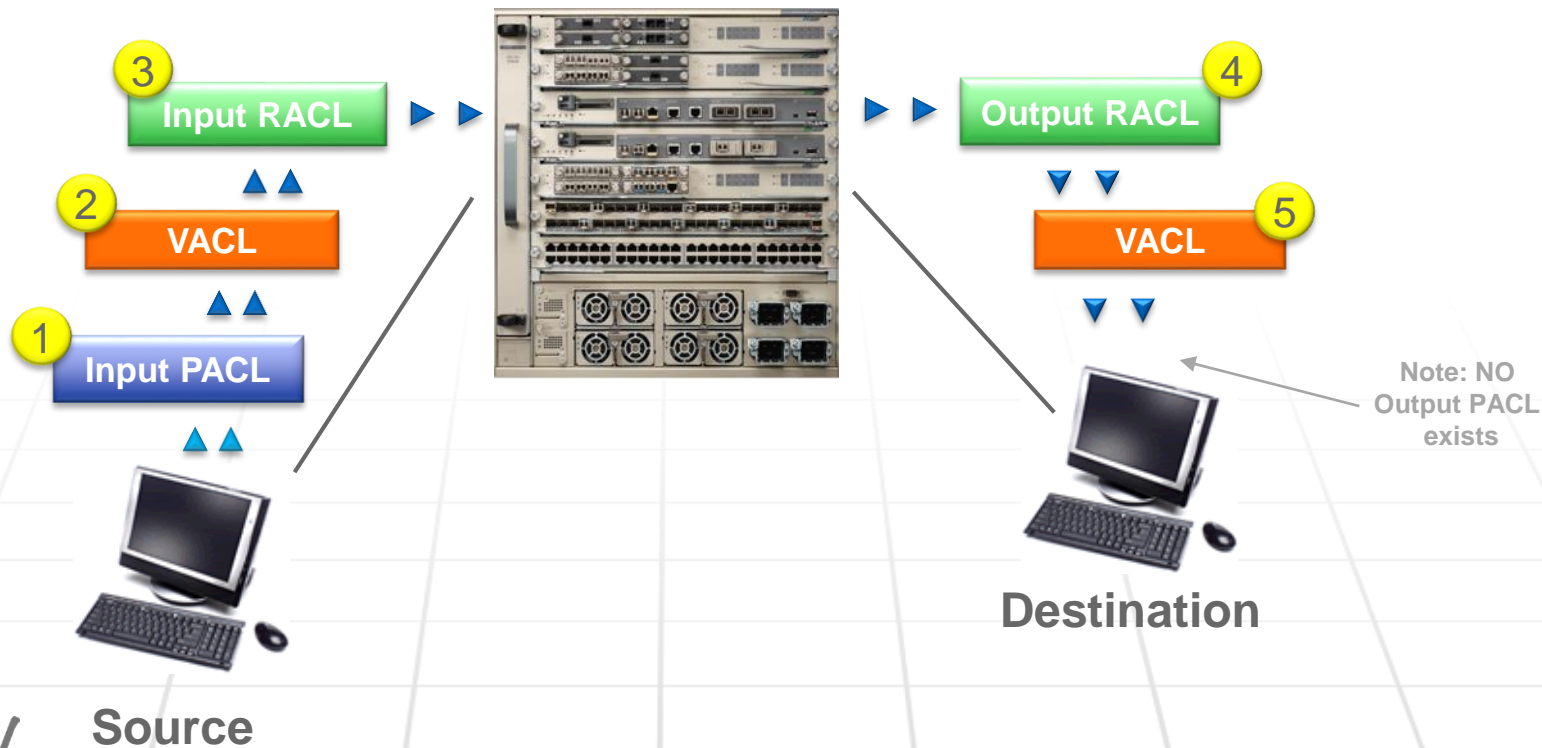
### Port ACL (PACL)

Used to permit or deny the movement of traffic between Layer 3 Subnets & VLANs or within a VLAN

Applied as a policy to a Layer 2 Switch port interface - is applied for inbound traffic only

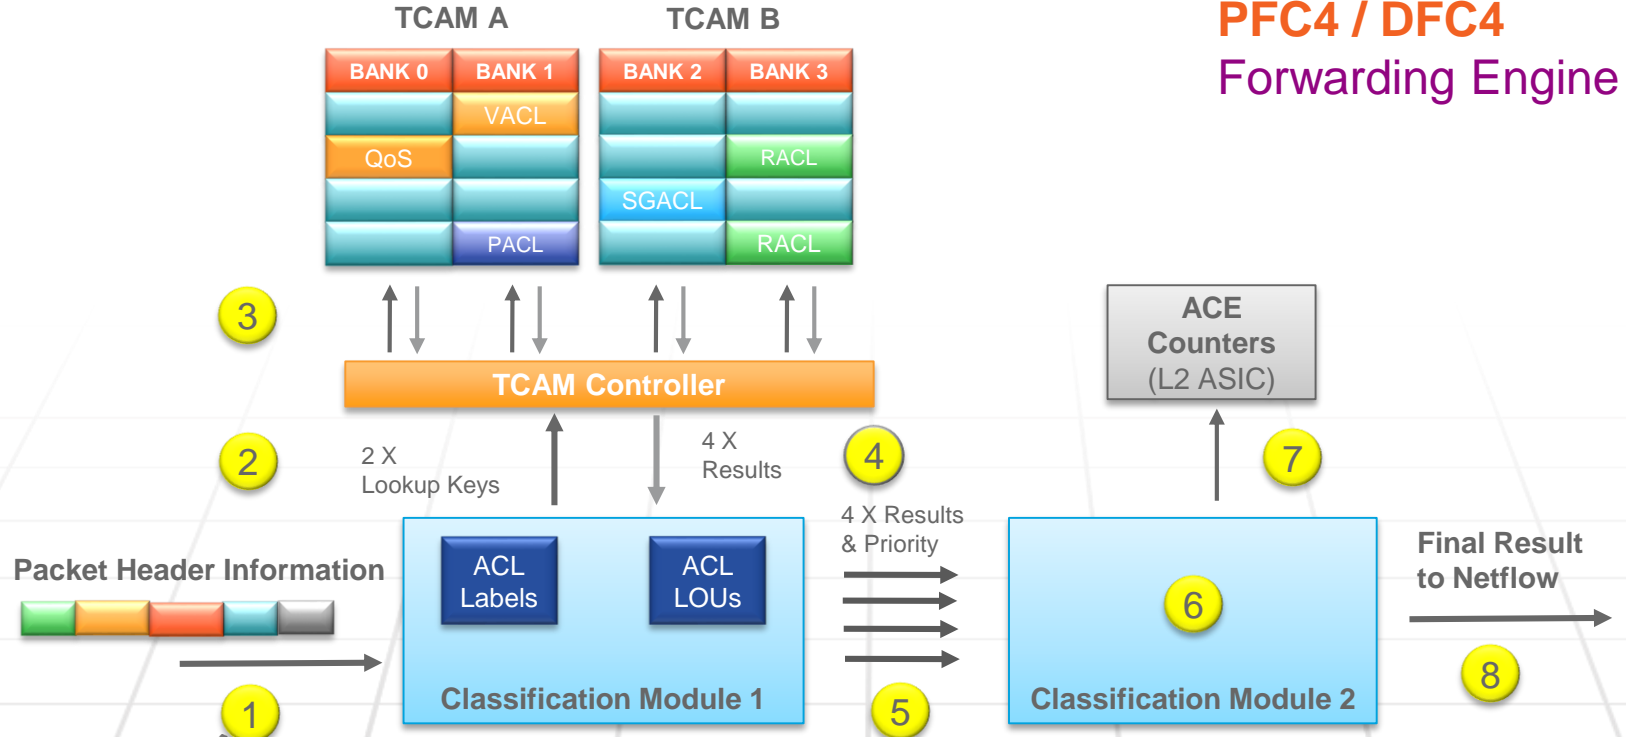
# Access Control Lists

## ACL Order of Processing



# Access Control Lists

## PFC4 TCAM Lookup



## PFC4 / DFC4 Forwarding Engine



# Access Control Lists

## Mask Utilization: PFC3 vs PFC4

```
permit ip 10.1.1.0 0.0.0.255 any
permit ip 10.2.1.0 0.0.0.255 any
permit ip 10.3.0.0 0.0.255.255 any
```

### PFC3 ACL TCAM

MASK  0.0.0.255	10.1.1.0	permit
	10.2.1.0	permit
	-	-
	-	-
	-	-
	-	-
MASK  0.0.255.255	10.3.0.0	permit
	-	-
	-	-
	-	-
	-	-
	-	-

16 ACEs  
used

### PFC4 ACL TCAM

Mask 0.0.0.255	10.1.1.0	permit
Mask 0.0.0.255	10.2.1.0	permit
Mask 0.0.255.255	10.3.0.0	permit
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-
-	-	-

3 ACEs  
used

# Access Control Lists

## PFC4 “ACL Dry Run” Feature

- **Make sure an ACL will fit in the TCAM before you apply the ACL**
  - ACLs that do not fit can cause software forwarding and possible high CPU utilization
- **Special configuration session**
  - Create and Edit the ACL
  - Verifies if the changes will fit within the hardware resources
- The actual changes are not programmed into the hardware during the configuration session
- Configuration changes can be verified step by step

SUP2T-E# **show configuration session test status**


=====

Status of last config validation:

Timestamp: 2013-09-20@17:27:06

=====

SLOT = [1] **Result = Configuration will fit in TCAM**



# PFC4 ACL Dry Run

## Configuration Example

### 1- Start the configuration session

```
SUP2T-E#configure session dryrun
SUP2T-E(dry-run-config)#
```

### 2- Enter the ACL to be modified/created

```
SUP2T-E(dry-run-config)#ip access-list extended dryrun3
SUP2T-E(dr-config-ext-nacl)#
```

### 3- Make the changes to the ACL

```
SUP2T-E(dr-config-ext-nacl)#permit ip any host 6.1.1.1
SUP2T-E(dr-config-ext-nacl)#permit ip any host 6.1.1.2
SUP2T-E(dr-config-ext-nacl)#permit ip any host 6.1.1.3
SUP2T-E(dr-config-ext-nacl)#permit ip any host 6.1.1.4
SUP2T-E(dr-config-ext-nacl)#permit ip any host 6.1.1.5
SUP2T-E(dr-config-ext-nacl)#permit ip any host 6.1.1.6
```

### 4- Exit out of the ACL config mode

```
SUP2T-E(dr-config-ext-nacl)#exit
SUP2T-E(dry-run-config)#
```

### 5- Execute the Validate command

```
SUP2T-E(dry-run-config)#validate
SUP2T-E(dry-run-config)#
*Feb 20 17:27:08.021:
%FM-6-SESSION_VALIDATION_RESULT_INFO:
Session Validation Result :
"Validation Completed Successfully."
. Please use 'show config session dryrun status' to get the more details of the config validation
status
SUP2T-E(dry-run-config)#
```

### 6- Exit out of the configuration session and verify if the validated session changes will fit in the hardware resources


```
SUP2T-E(dry-run-config)#end
SUP2T-E#
SUP2T-E#show configuration session dryrun status
=====
Status of last config validation:
Timestamp: 2010-02-20@17:27:06
=====
SLOT = [1] Result = Configuration will fit in TCAM
```

# Access Control Lists


## PFC4 “ACL Hitless Update” Feature



- **Allows updates to an ACL without interrupting traffic**
- Multiple features updated at once
  - IPv4, IPv6, MAC...
  - RACL, VACL, PBR...
- Global configuration option (default is on)
- Feature does consume double the number of TCAM entries



↓  
**ACL  
Updates**



# Access Control Lists

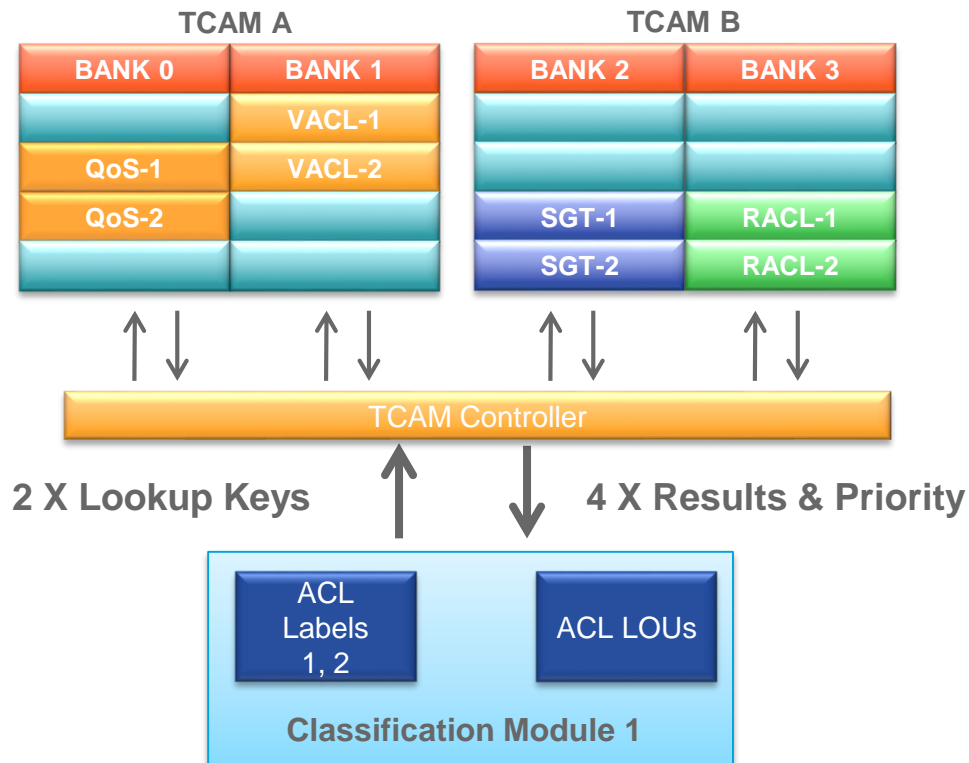
## PFC4 ACL Hitless Update

- Each ACL feature is initially programmed into two different spaces into the TCAM

Primary Space (Label-1)

Shadow Space (Label-2)

- While an ACL is being updated the PFC4 will use a temporary label that points to the shadow TCAM space
- Once the ACL changes have been completed the then PFC4 will then use the original label again




# Sup720 vs. Sup2T- Security

Feature	Sup720	Sup2T
ACL TCAM Size	32K	64K or 256K (Programmable)
ACL Labels	4K	16K
ACE Mask Ratio	8:1	1:1
ACL LOU's	64	208
Port ACL's	2K	8K
MAC ACL Support	Software	Hardware
Per-Port Per-VLAN ACL's	N/A	Yes
Security Group ACL's	N/A	32K
Security Group Tagging (SGT)	CTS 1.5 (SXP)	CTS 2.0
802.1ae Encryption Support	N/A	Yes (Line-Rate)
Unicast RPF	IPv4 (2 paths)	IPv4 & IPv6 (16 paths)
Hardware Rate Limiters	Layer3: 8 Layer2: 4	Layer3: 31 Layer2: 26

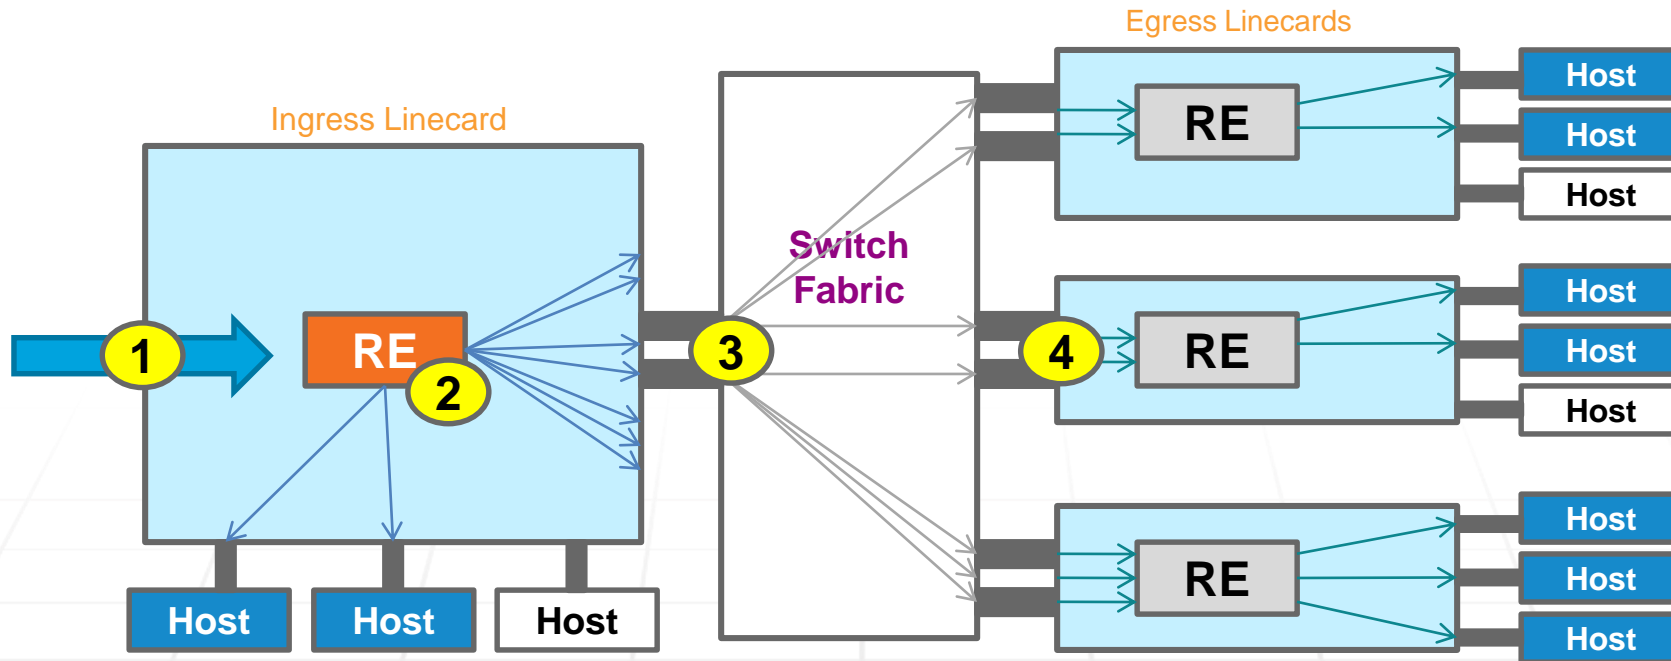
# Agenda

- ❖ Chassis & Power
- ❖ Supervisor Architectures
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ **Packet Walks**
  - ❖ **6700/6800 & 6900 Series**
  - ❖ C6800 1G & 10G Series
  - ❖ C6880-X & C6840-X Series





# Multicast Ingress Replication Model

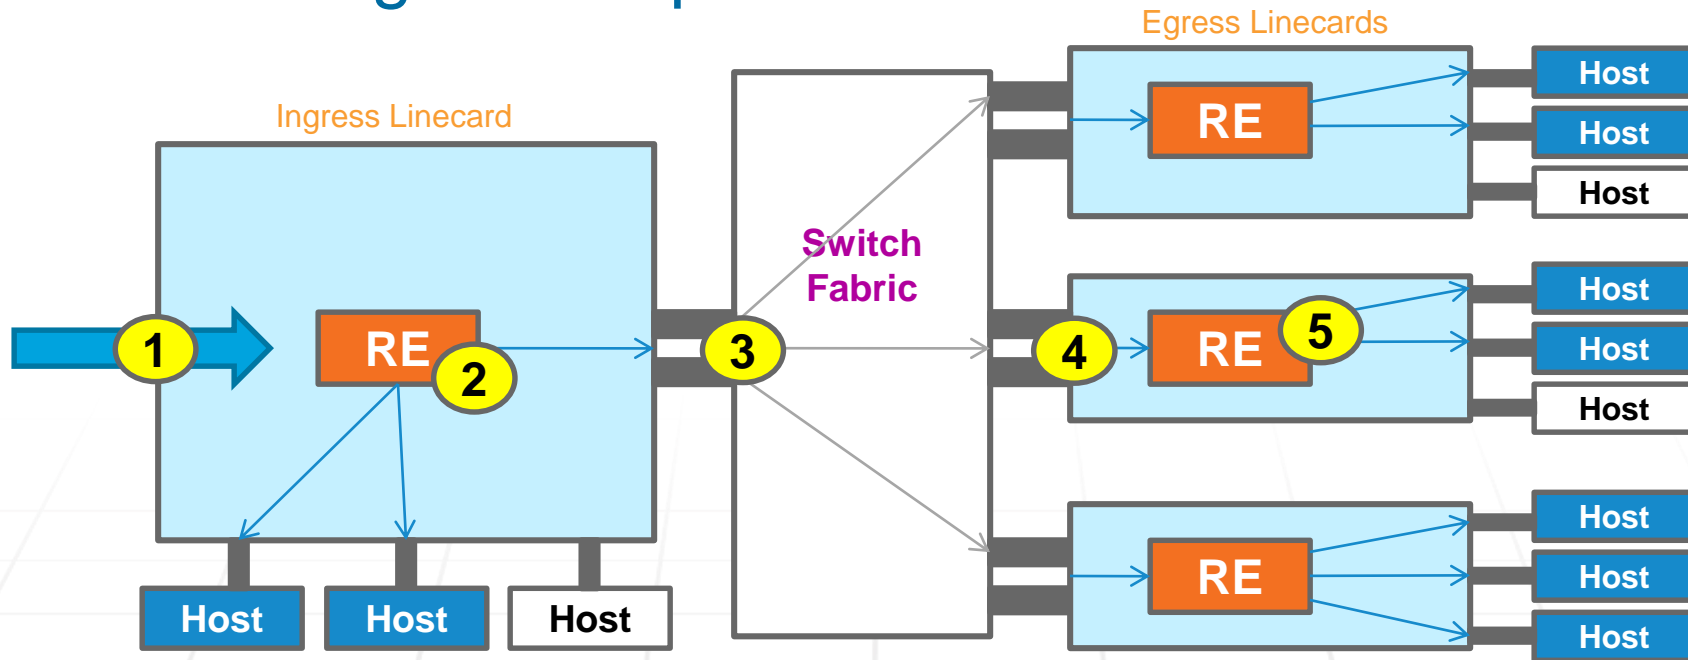


Ingress Replication Engine (RE) responsible for replication to ALL OIF's

Replicates for each egress linecard

Switch Fabric sends copies for each OIF


# Multicast Egress Replication Model



Ingress Replication Engine responsible  
for replication to local OIF's  
Replicates a SINGLE copy to Fabric

Switch Fabric replicates for each OIF  
Local Linecard RE replicates for local OIF's

# 6700 to 6700 - Centralized Forwarding




Michael Engineering



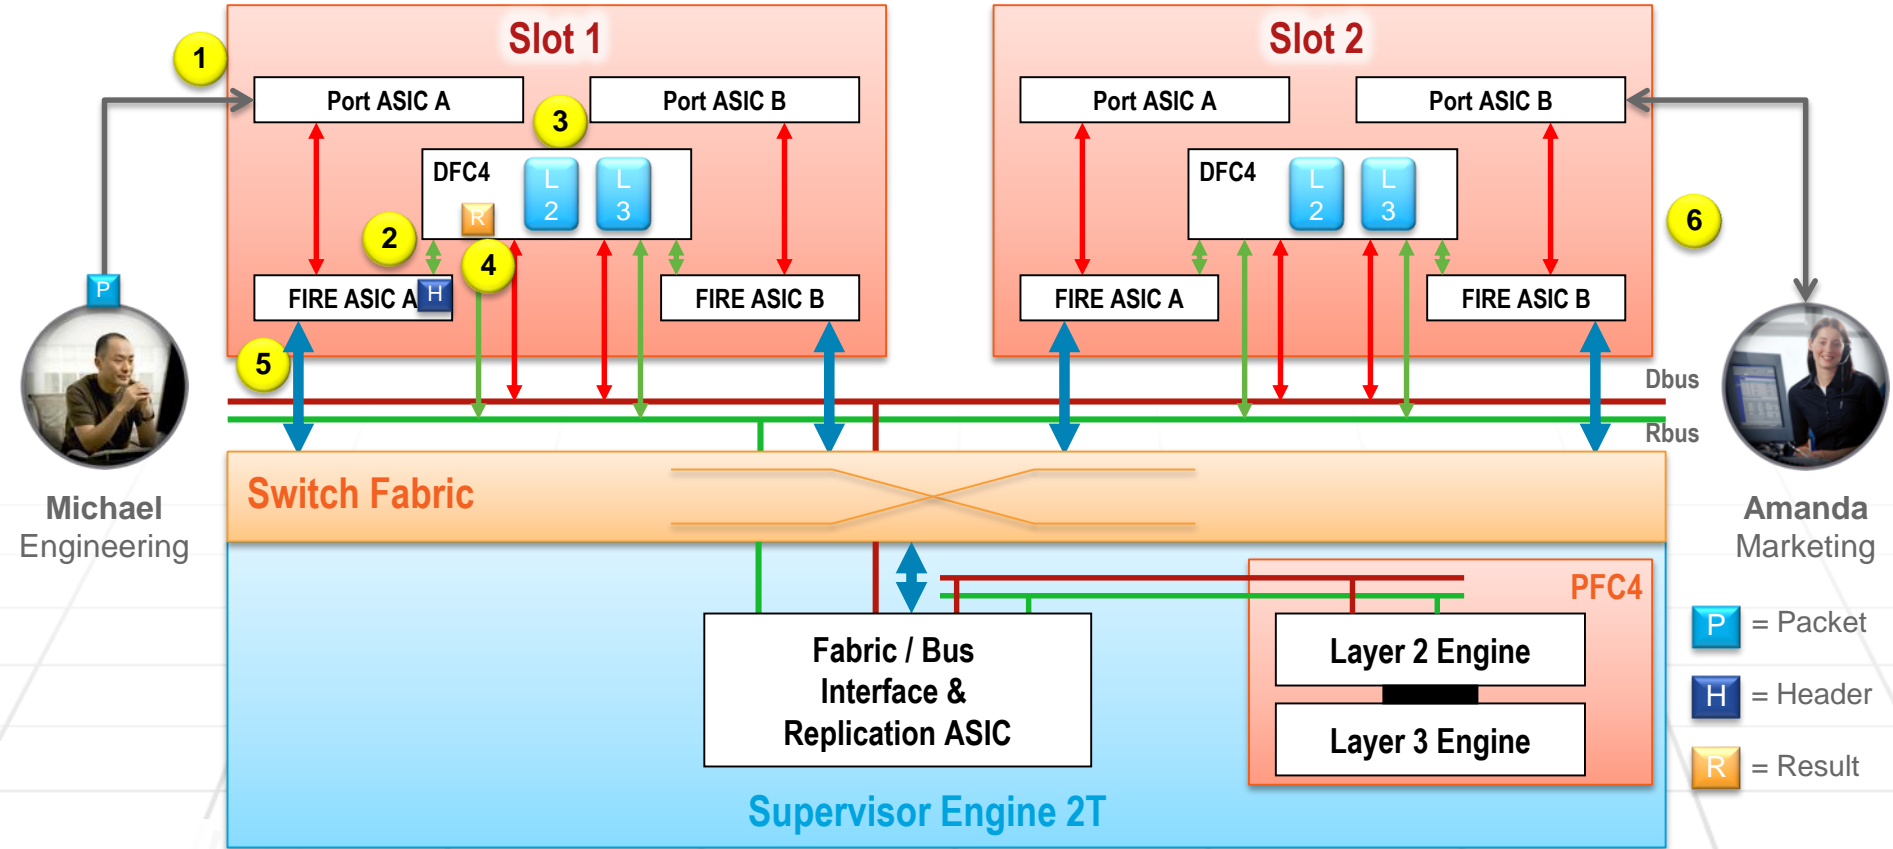
Amanda Marketing

# 6700 (CFC)

## Local Forwarding



# 6800 to 6800 - Distributed Forwarding (DFC)




Michael Engineering



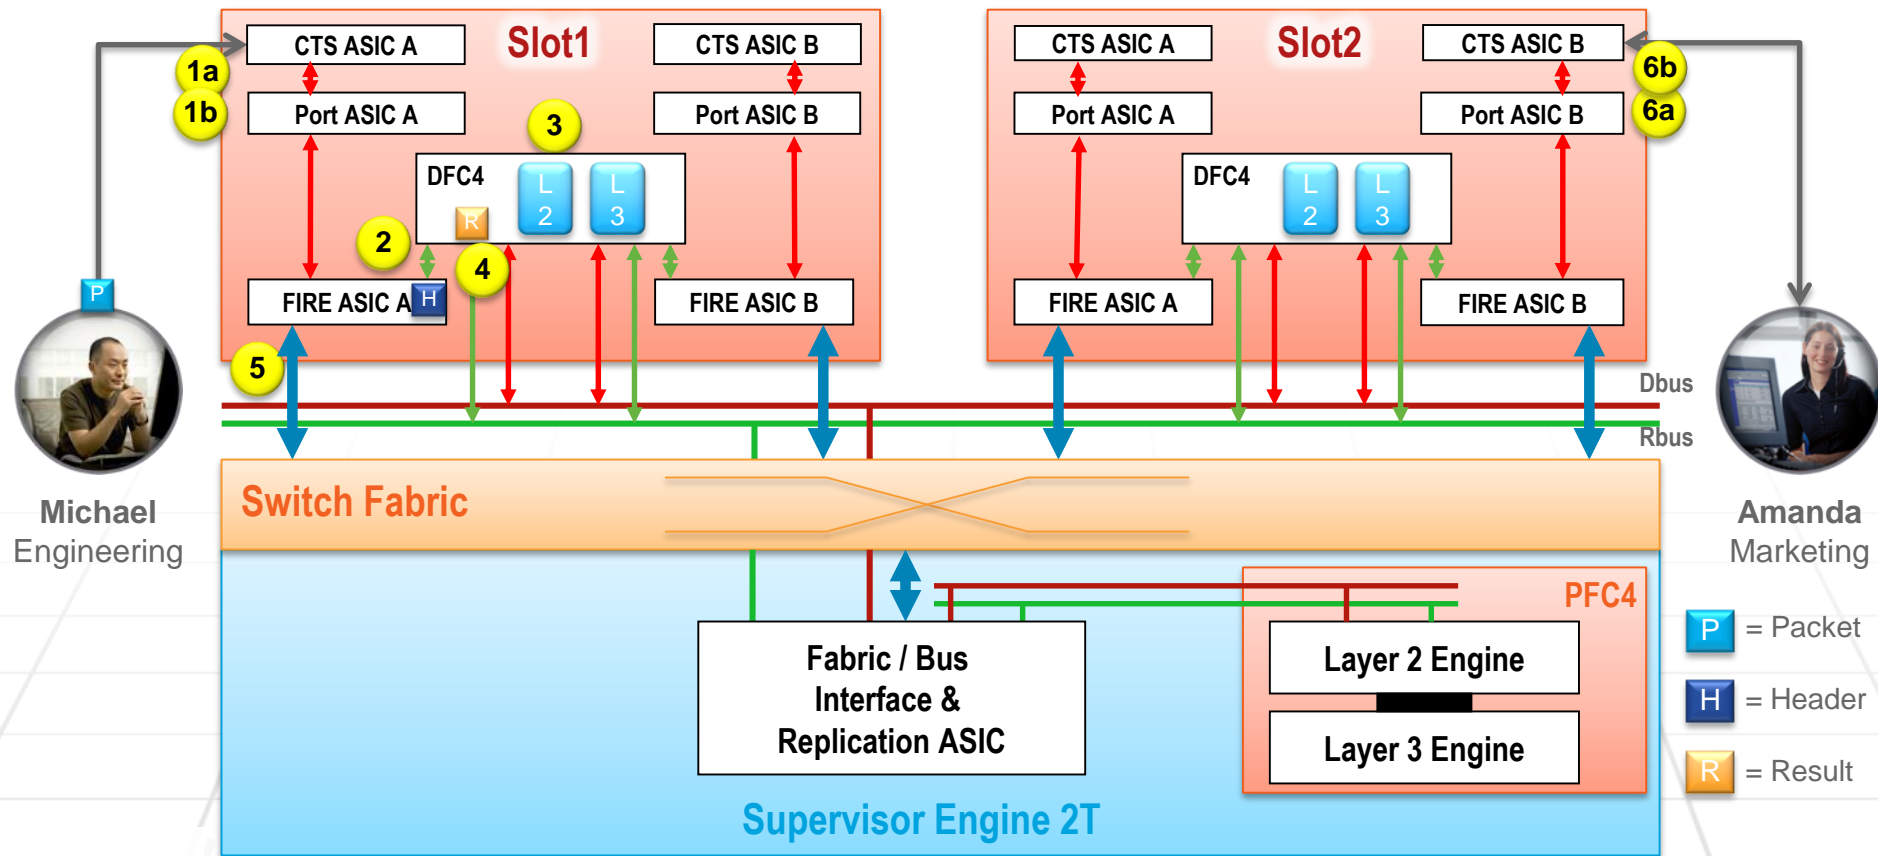
Amanda Marketing

# 6800 (DFC)

## Local Forwarding



# 6900 to 6900 - Distributed Forwarding (DFC)




Michael  
Engineering

Amanda  
Marketing




# 6900 (DFC)

## Local Forwarding



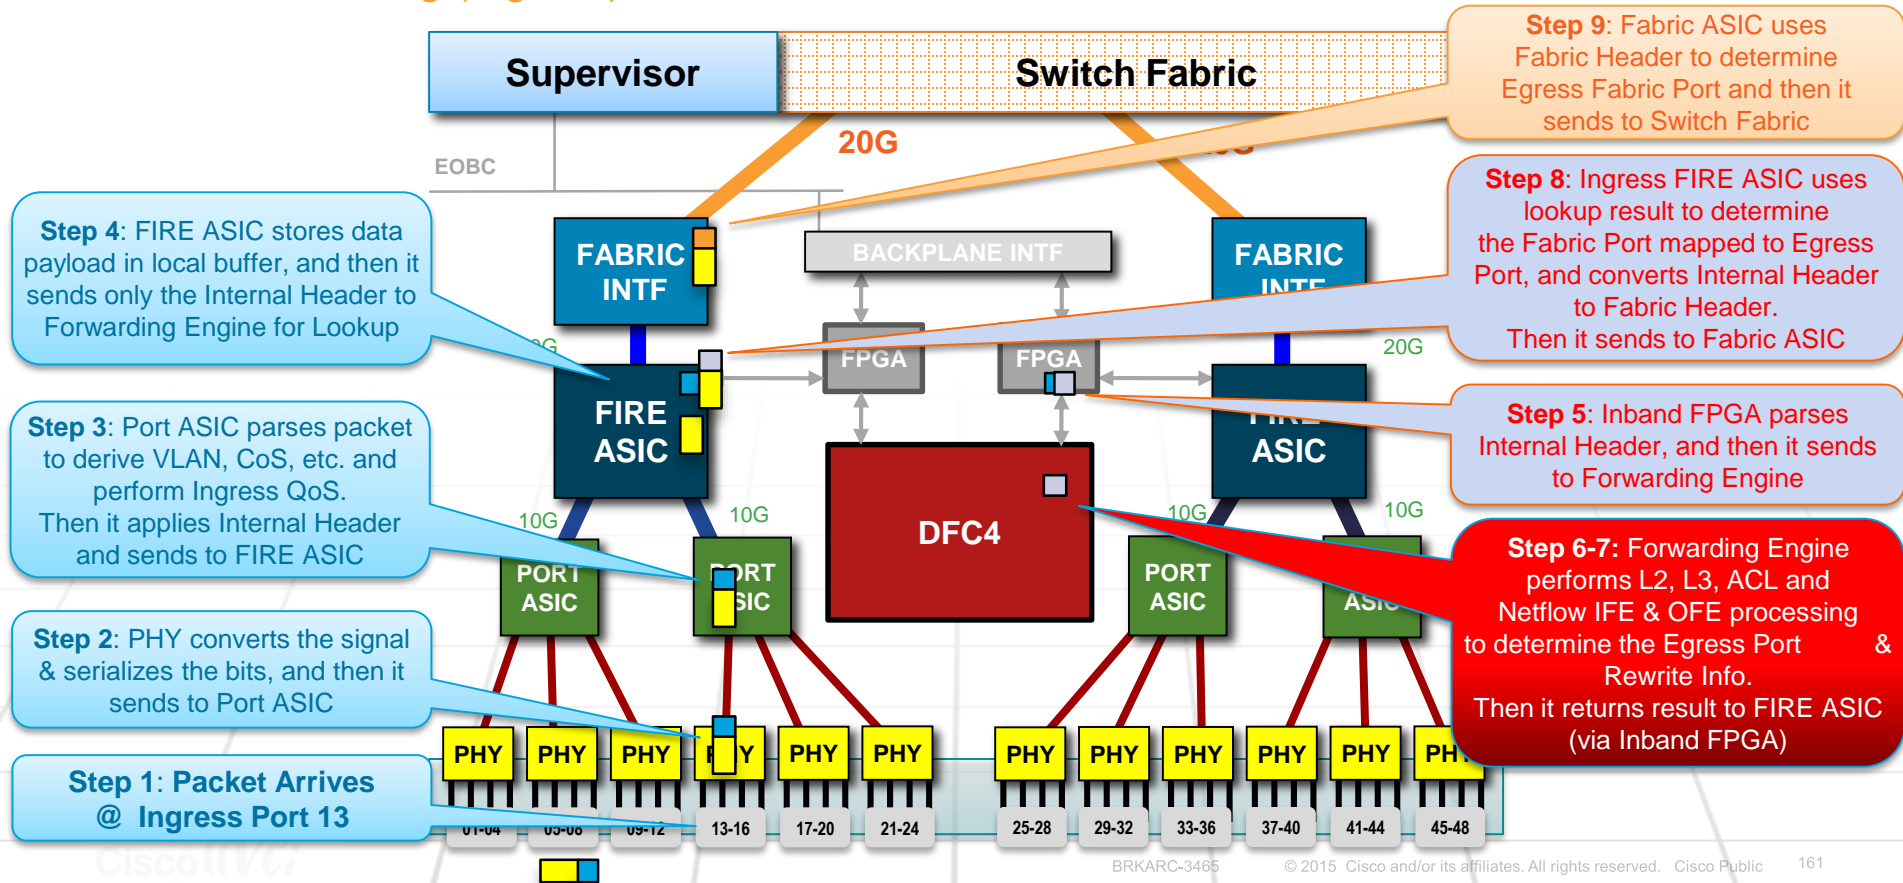
# Agenda

- ❖ Chassis & Power
- ❖ Supervisor Architectures
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ **Packet Walks**
  - ❖ 6700/6800 & 6900 Series
  - ❖ **C6800 1G & 10G Series**
  - ❖ C6880-X & C6840-X Series



# C6800 1G - Packet Walks

## Remote Forwarding (Ingress)



# C6800 1G Packet Walks

## Remote Forwarding (Egress)

**Step 10:** Switch Fabric transmits frame to Egress Fabric ASIC, which sends to Egress FIRE ASIC

**Step 11:** FIRE ASIC uses Fabric Header to derive new Internal header, which it sends to Forwarding Engine (Egress Lookup)

**Step 12:** Inband FPGA parses Internal Header, and then it sends to Forwarding Engine

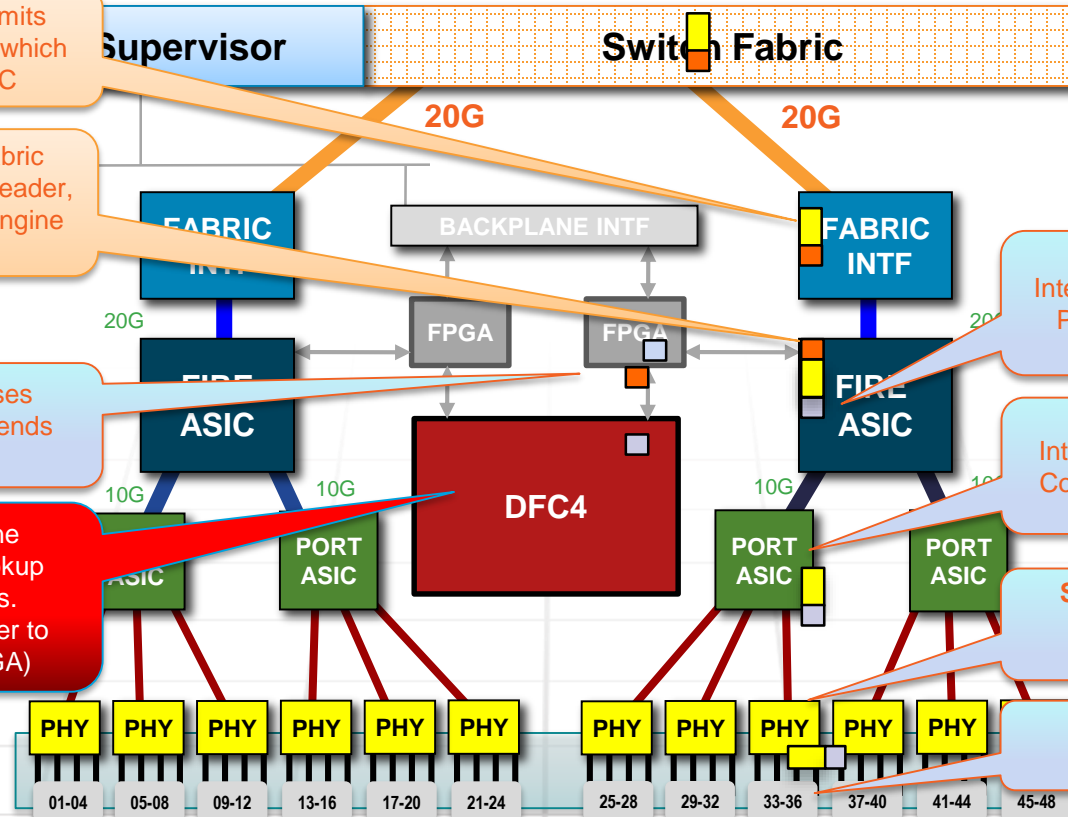
**Step 13:** Forwarding Engine performs an Egress (L2) Lookup to learn SRC MAC address. Then it returns Internal Header to FIRE ASIC (via Inband FPGA)

**Step 14:** FIRE ASIC uses Internal Header to determine Egress Port and reassemble the packet. Then it sends to Port ASIC

**Step 15:** Port ASIC removes Internal Header and rewrites VLAN, CoS, etc. and perform Egress QoS. Then it sends to PHY

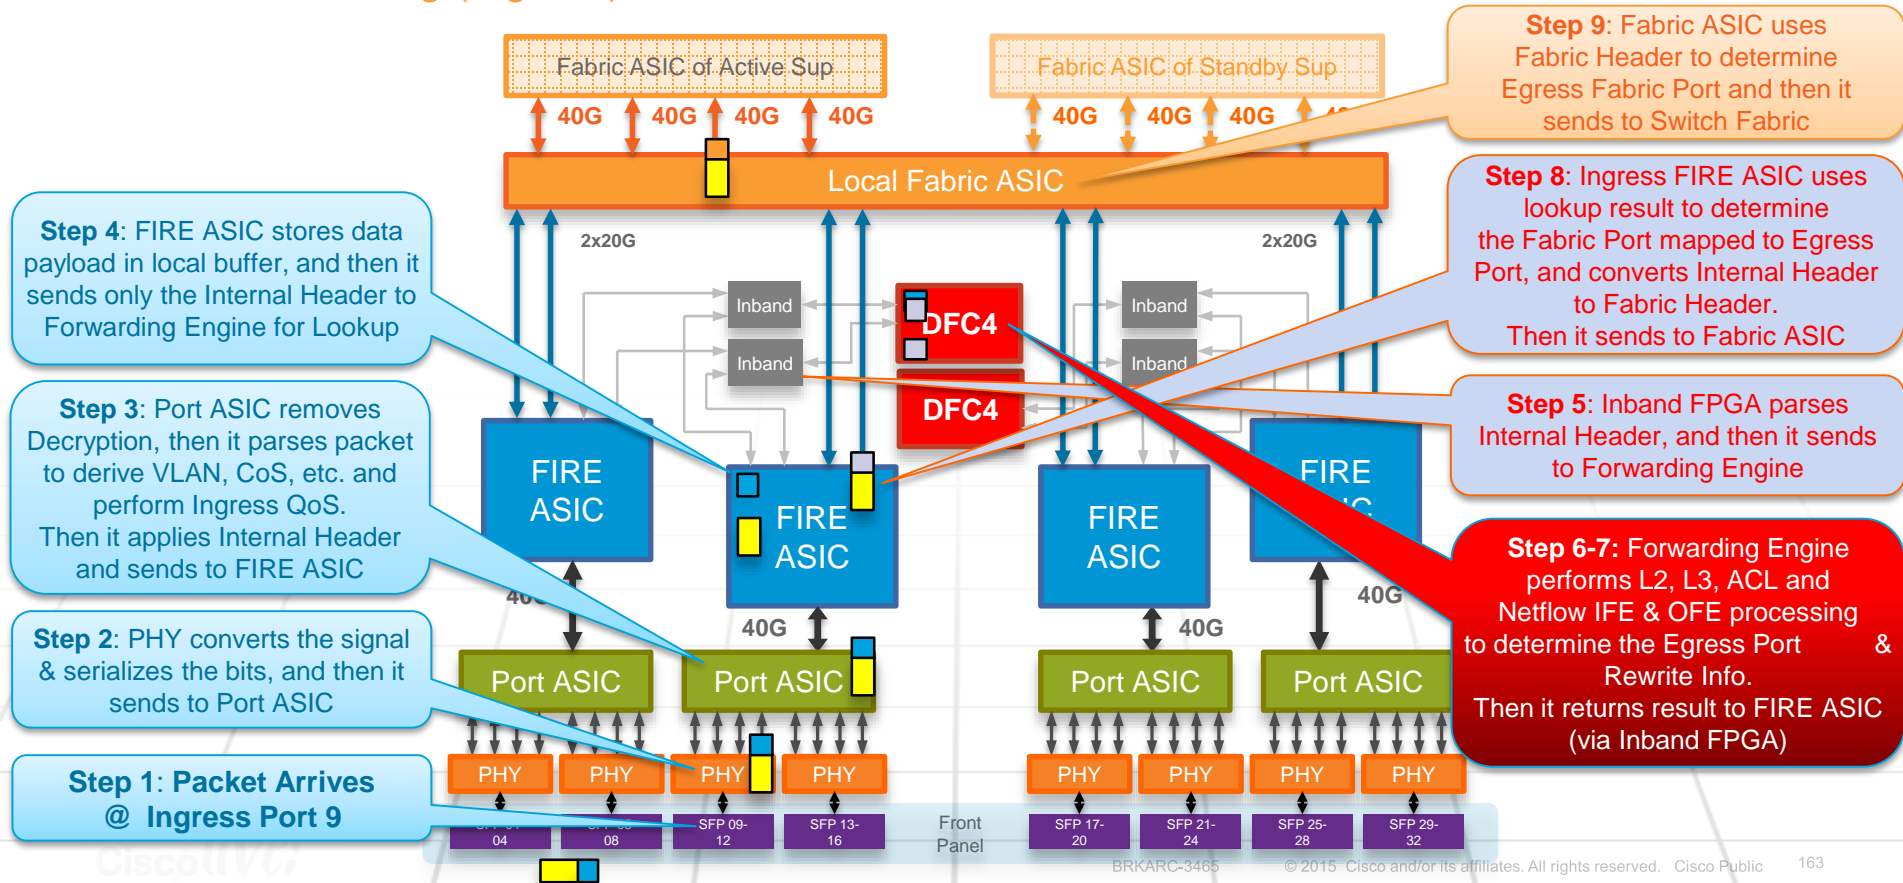
**Step 16:** PHY serializes the bits & converts signal, and then transmits the packet

**Step 17:** Packet Leaves @ Egress Port 36



# C6800 10G - Packet Walks

## Remote Forwarding (Ingress)



# C6800 10G Packet Walks

## Remote Forwarding (Egress)

**Step 10:** Switch Fabric transmits frame to Egress Fabric ASIC, which sends to Egress FIRE ASIC

**Step 11:** FIRE ASIC uses Fabric Header to derive new Internal header, which it sends to Forwarding Engine (Egress Lookup)

**Step 12:** Inband FPGA parses Internal Header, and then it sends to Forwarding Engine

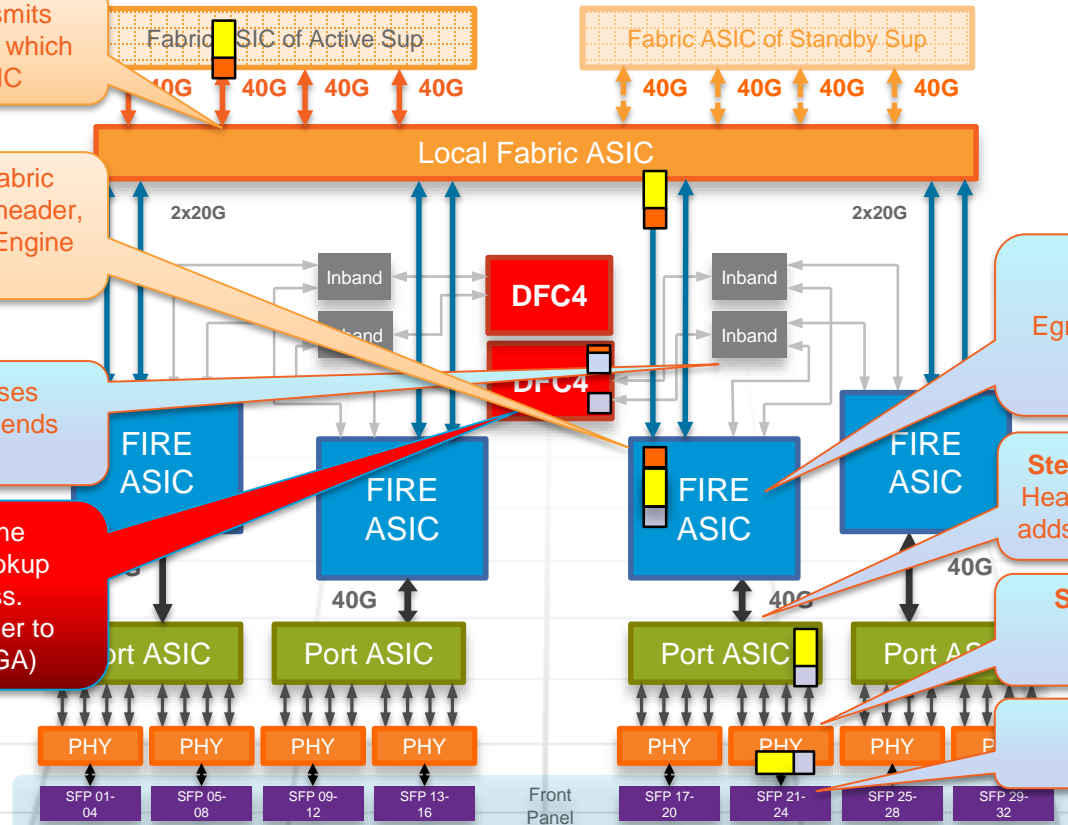
**Step 13:** Forwarding Engine performs an Egress (L2) Lookup to learn SRC MAC address. Then it returns Internal Header to FIRE ASIC (via Inband FPGA)

**Step 14:** FIRE ASIC uses Internal Header to determine Egress Port, reassemble the packet and perform Egress QoS. Then it sends to Port ASIC

**Step 15:** Port ASIC removes Internal Header, rewrites VLAN, CoS, etc. and adds Encryption. Then it sends to PHY


**Step 16:** PHY serializes the bits & converts signal, and then transmits the packet

**Step 17: Packet Leaves @ Egress Port 24**



# Agenda

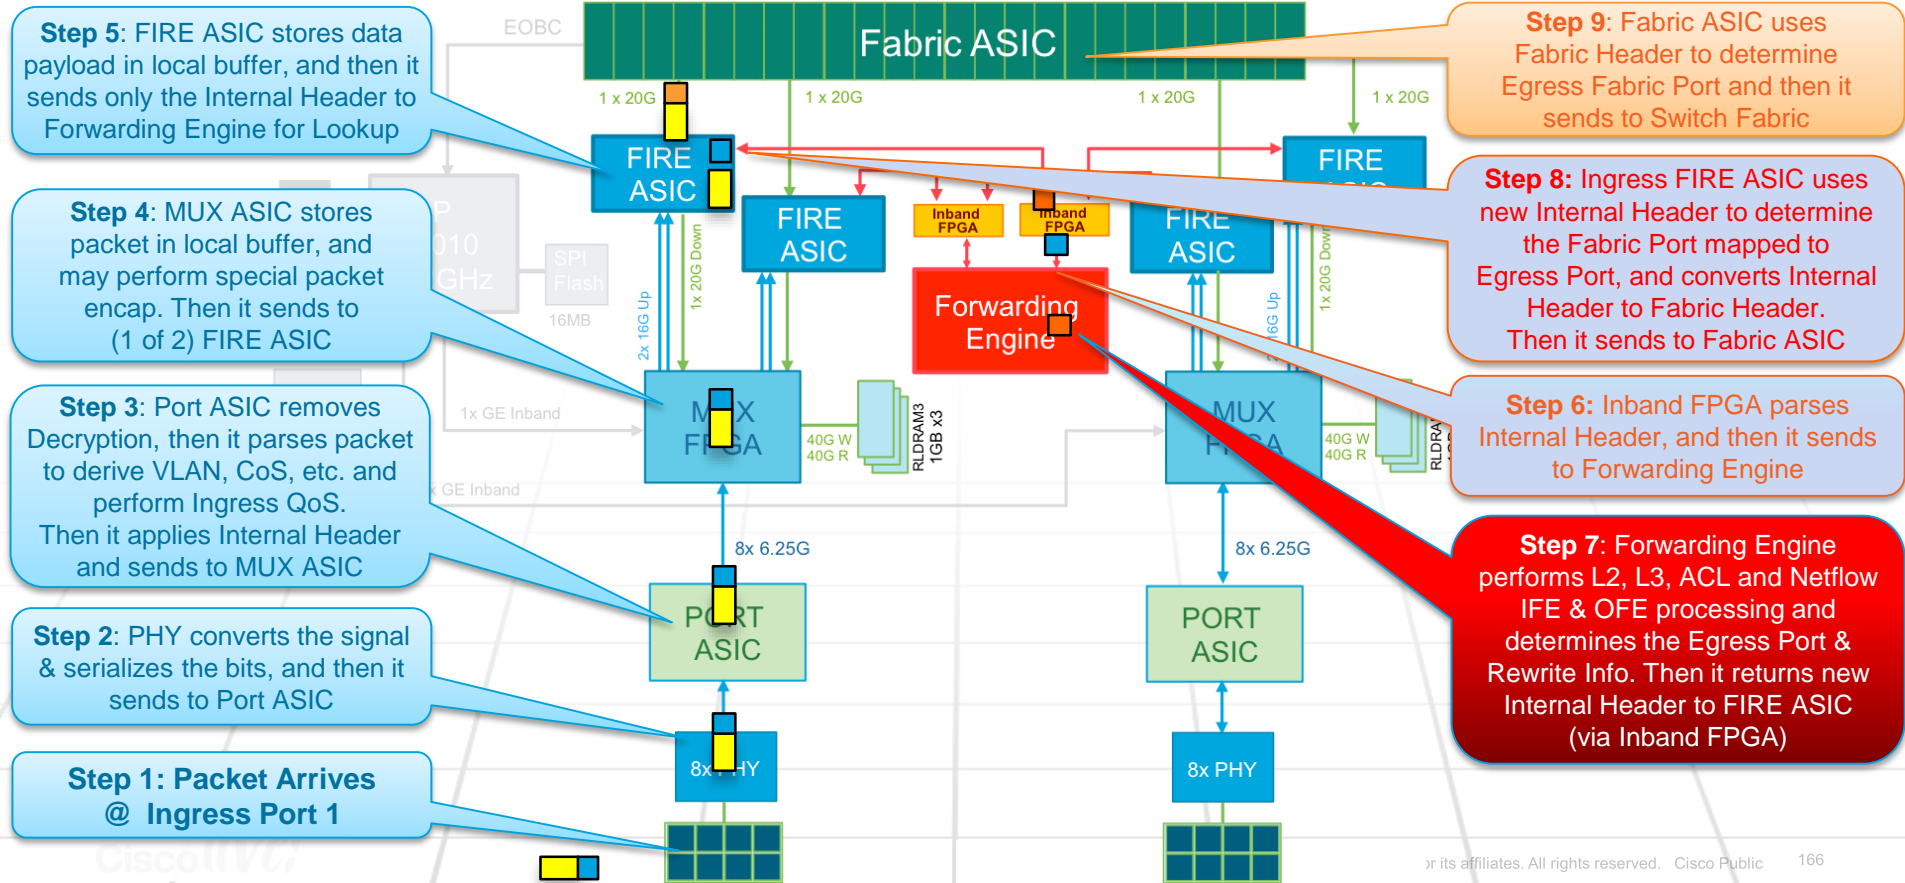
- ❖ Chassis & Power
- ❖ Supervisor Architectures
- ❖ Module Architectures
- ❖ L2 Packet Forwarding
- ❖ L3 Packet Forwarding
- ❖ NetFlow & NDE
- ❖ Access Control Lists
- ❖ **Packet Walks**
  - ❖ 6700/6800 & 6900 Series
  - ❖ C6800 1G & 10G Series
  - ❖ **C6880-X & C6840-X Series**





# C6880-X & C6840-X Packet Walk

## Remote Forwarding (Egress)





# C6880-X & C6840-X Packet Walk

## Remote Forwarding (Egress)

**Step 10:** Fabric ASIC transmits frame to Egress Fabric Port, which is received by Egress FIRE ASIC

**Step 12:** Inband FPGA parses Internal Header, and then it sends to Forwarding Engine

**Step 13:** Forwarding Engine performs an egress (L2) lookup to learn MAC address. Then it returns Internal Header to FIRE ASIC (via Inband FPGA)

**Step 11:** FIRE ASIC uses Fabric Header to derive new Internal header, which it sends to Forwarding Engine (egress lookup)

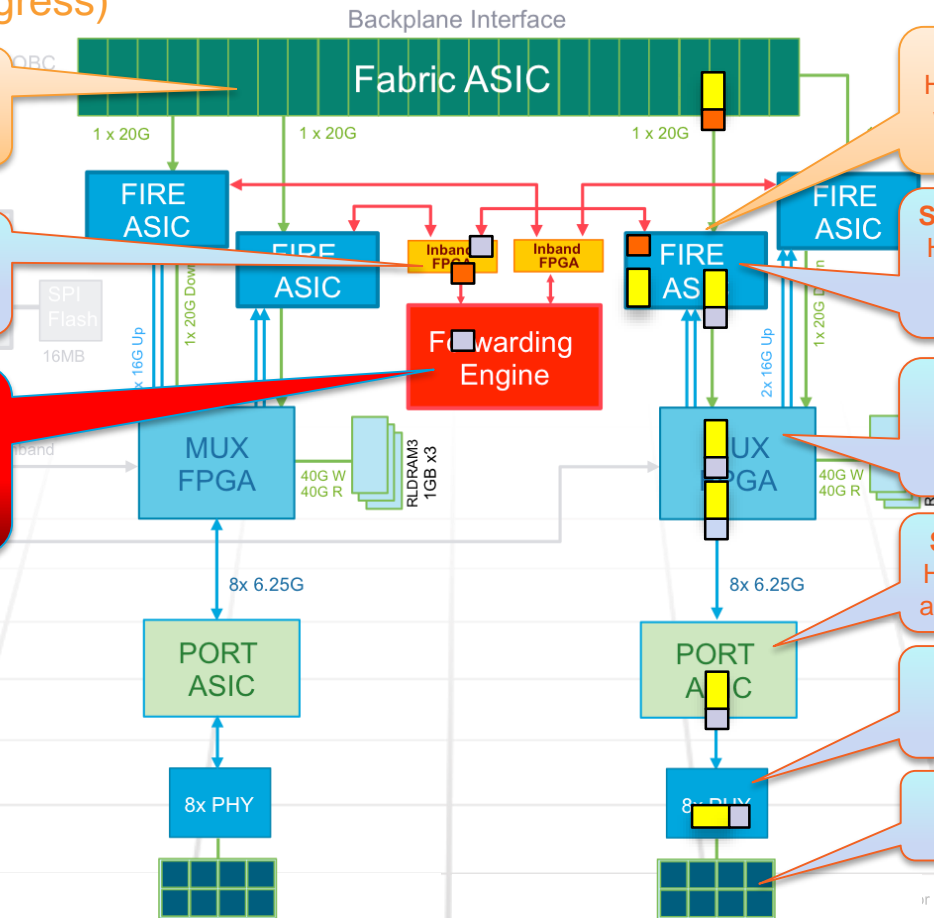
**Step 14:** FIRE ASIC uses new Internal Header to determine Egress Port and reassemble the packet, and then it sends to MUX ASIC

**Step 15:** MUX ASIC uses Internal Header to determine Egress Port and perform Egress QoS. Then it sends to Port ASIC

**Step 16:** Port ASIC removes Internal Header, rewrites VLAN, CoS, etc. and adds Encryption. Then it sends to PHY

**Step 17:** PHY serializes the bits & converts signal, and then transmits the packet

**Step 18:** Packet Leaves @ Egress Port 16



# Summary

- **Innovation & 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 hardware features such as Netflow, QoS and Security can be done without impacting performance or scale, because these hardware features are processed in parallel.



# Catalyst 6800 Series

3500+ Catalyst 6500 Features

## MPLS, VPLS & EVN

15 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	✓
MPLS at Access with I-Sat	✓
VRF-Lite	✓
EVN	✓
VRF-Aware Multicast	✓
VRF-Aware NetFlow	✓

## Next-Gen Solutions

Instant Access Controller	✓
L2 Campus LISP	✓
L3 Campus LISP	✓
SDN - OpenFlow/APIC	✓
Smart Install Director	✓

## VSS & HA

10 Year VSS Maturity	✓
MACsec on VSL	✓
1G/10G/40G VSL	✓
SSO / NSF & NSR	✓
BGP PIC & IP FRR	✓
BFD on SVI & MEC	✓



## Security & QoS

MacSec & SGT	✓	DHCP Snooping	✓
SGACL & L3 SGT	✓	Dynamic ARP Inspection	✓
Ingress/Egress ACL	✓	SPAN with ACL	✓
Time-Based ACL	✓	Identity 2.0	✓
ACL Statistics	✓	ACLQoS Classification	✓
Port Security	✓	Marking (DSCP/COS)	✓
IPv6 FH Security	✓	Microflow Policing	✓
ACL Atomic Commit/Dry Run	✓	Advanced CoPP	✓
IPv6 uRPF	✓	IPv6 VLAN ACL	✓

## IPv4 Unicast

OSPFv2/v3	✓
OSPFv3 VRF-Lite	✓
VRF-Aware Unicast	✓
IP Tunnel HA	✓
BFD SVI-GRE	✓
BGP PIC	✓

## IPv4 Multicast

IGMPv3 & MLDv2 Snooping in HW	✓
PIM-SM "Dual-RPF" in HW	✓
PIM-DM, PIM-Bidir	✓
L3 PIM Snooping	✓
PIM Register in HW	✓
mVPN, MSR, mcast BFD	✓

## IPv6 Features

IPv6 - IPv4 HW Parity	✓	IPv6 in IPv4 Tunnels, VRF-Aware IPv6 Tunnels	✓
OSPFv3 VRF PE-CE	✓	BGPv6, IS-ISv6	✓
MLD, MLD Snooping	✓	6VPE and 6PE	✓
IPv6 Mcast HA	✓	PIM Sparse Mode (PIM-SM)	✓
IPv6 SLA, TCL, LLDP	✓	IPv6 QoS, PACL, RACL	✓

## 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
Aggregate Policers	8K
Number of VPNs	16K

## Rich Media


Flexible NetFlow	✓
Egress NetFlow	✓
Sampled NetFlow	✓
NDE (Full & Sampled)	✓
Video Monitoring	✓
Mediatrace	✓
Metadata QoS	✓
Multicast Service Reflect (MSR)	✓
Multicast VPN (MVPN)	✓

## Management & Services

WCCPv3	✓
PBR IPv4/IPv6	✓
NAT/PAT	✓
GRE/mGRE	✓
ERSPAN	✓
GOLD	✓
Cisco ISE	✓
Cisco Prime	✓
Mini Protocol Analyzer	✓

# Related Sessions, Demos & Labs

- **World of Solutions – Enterprise Networks**
  - Lazy Susan Static Displays
  - Equipment Rack Displays
  - VSS & Instant Access Demo
- **TECCRS-2932** - Cisco Campus LAN Switching Techtorial
- **BRKCRS-3035** - Advanced Enterprise Campus Design: VSS
- **BRKCRS-3502** - Advanced Enterprise Campus Design: Instant Access
- **LTRCRS-2004** - VSS & Instant Access Hands On Lab



# Complete Your Online Session Evaluation

- Give us your feedback to be entered into a Daily Survey Drawing.
- **A daily winner will receive a \$750 Amazon gift card.**
- Complete your session surveys though the Cisco Live mobile app or your computer on Cisco Live Connect.



**Don't Forget:** All Cisco Live sessions will be available for viewing on-demand after the event at [CiscoLive.com/Online](https://ciscolive.com/online)



# *Thank You*



*TOMORROW starts here.*