# Cisco Catalyst 6500 Architecture

This paper will provide an insight into the architecture of the Cisco® Catalyst® 6500 family, including supervisor engines, line cards, chassis, and other components.

Cisco Catalyst 6500 Chassis

The Cisco Catalyst 6500 is a modular chassis family supporting a range of line-card options for all customer network deployments. Its performance scales up to 400Mpps, providing one of the highest performing switch products on the market today. It supports a range of chassis options delivering from 3 to 13 slots. All chassis support redundant supervisor engines as well as redundant power supplies. The chassis options are shown in Figure 1.

Figure 1. Cisco Catalyst 6500 Chassis Family
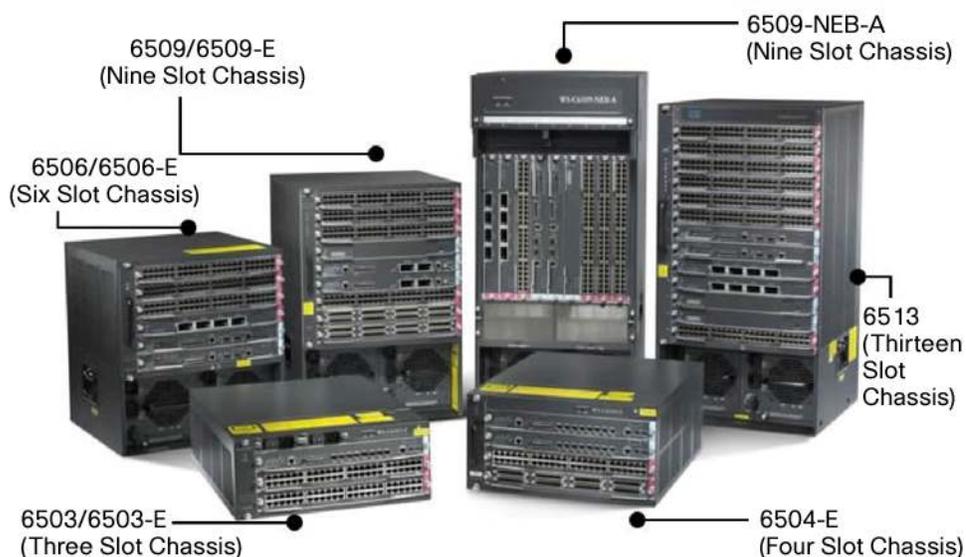
The Cisco Catalyst 6500 chassis family includes the following:

- **Cisco Catalyst 6503 Switch**: Three-slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to two (2) line cards. NEBS L3 compliant.
- **Cisco Catalyst 6506 Switch**: Six slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to five (5) line cards. NEBS L3 compliant.
- **Cisco Catalyst 6509 Switch**: Nine slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to eight (8) line cards. NEBS L3 compliant.
- **Cisco Catalyst 6509-NEB Switch**: NEBS based nine slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to eight (8) line cards. This chassis also supports front to back airflow and is NEBS L3 compliant.

- **Cisco Catalyst 6509-NEB-A Switch**: NEBS based nine slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to eight (8) line cards. This NEBS based chassis differs from the first Cisco Catalyst 6509-NEB in that it can support the new Cisco Catalyst 6500 Series Supervisor Engine 720 without any upgrades to the chassis.

- **Cisco Catalyst 6513 Switch**: Thirteen slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to twelve (12) line cards. This chassis requires the Supervisor Engine 2 as the minimum version of installed supervisor. NEBS L3 compliant.

Cisco has also recently introduced a new "E" series chassis that is designed to support a larger power draw over the backplane and also drive higher power loads to each line-card slot. In all, four new "E" series chassis have been announced:

- **Cisco Catalyst 6503-E Switch**: Three slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to two (2) line cards. NEBS L3 compliant. This chassis increases available power from 50A@42V (~2500W) to 80A@42V (~4000W).

- **Cisco Catalyst 6504-E Switch**: Four slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to three (3) line cards. NEBS L3 compliant. This chassis increases available power from 50A@42V (~2500W) to 80A@42V (~4000W).

- **Cisco Catalyst 6506-E Switch**: Six slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to five (5) line cards. NEBS L3 compliant. This chassis increases available power from 90A@42V (~4000W) to 350A@42V (~12000W).

- **Cisco Catalyst 6509-E Switch**: Nine slot chassis supporting redundant power supplies, redundant supervisor engines, and slots for up to eight (8) line cards. NEBS L3 compliant. This chassis increases available power from 90A@42V (~4000W)to 350A@42V (~12000W).

Cisco Catalyst 6500 Backplane

The Cisco Catalyst 6500 incorporates two backplanes. From its initial release in 1999, the Cisco Catalyst 6500 chassis has supported a 32-Gbps shared switching bus, a proven architecture for interconnecting line cards within the chassis. The Cisco Catalyst 6500 chassis also includes a second backplane that allows line cards to connect over a high-speed switching path into a crossbar switching fabric. The crossbar switching fabric provides a set of discrete and unique paths for each line card to both transmit data into and receive data from the crossbar switching fabric. The first generation switching fabric was delivered by the switch fabric modules (WS-C6500-SFM and WS-C6500-SFM2), each providing a total switching capacity of 256 Gbps. More recently, with the introduction of the Supervisor Engine 720, the crossbar switch fabric has been integrated into the Supervisor Engine 720 baseboard itself, eliminating the need for a standalone switch fabric module. The capacity of the new integrated crossbar switch fabric on the Supervisor Engine 720 has been increased from 256 Gbps to 720 Gbps. The Supervisor Engine 720-3B and Supervisor Engine 720-3BXL also maintain the same fabric capacity size of 720 Gbps.

Depending on the Cisco Catalyst 6500 chassis, the crossbar switching fabric maps out a series of fabric channels (otherwise known as paths into the crossbar) to each line-card slot in a slightly different layout. Each chassis fabric layout is detailed in Table 1.

Table 1.   Chassis Slot Options

Chassis	Supervisor Engine 32/ Supervisor Engine 720 Slots	Classic Line-Card Slots	Single Fabric Connected Line Cards	Dual Fabric Connected Line Cards
Cisco Catalyst 6503	1 and 2	2 and 3	2 and 3	WS-X67XX line cards not supported
Cisco Catalyst 6503-E	1 and 2	2 and 3	2 and 3	2 and 3
Cisco Catalyst 6504-E	1 and 2	2 thru 4	2 thru 4	2 thru 4
Cisco Catalyst 6506	5 and 6	1 through 6	1 through 6	1 through 6
Cisco Catalyst 6506-E	5 and 6	1 through 6	1 through 6	1 through 6
Cisco Catalyst 6509	5 and 6	1 through 9	1 through 9	1 through 9
Cisco Catalyst 6509-E	5 and 6	1 through 9	1 through 9	1 through 9
Cisco Catalyst 6509-NEB	5 and 6	1 through 9	1 through 9	1 through 9
Cisco Catalyst 6509-NEB-A	5 and 6	1 through 9	1 through 9	1 through 9
Cisco Catalyst 6513	7 and 8	1 through 13	1 through 13	9 through 13

In all but the thirteen slot chassis, each line-card slot has two channels in and out of the switching fabric. The thirteen slot chassis has one fabric channel to each slot in slots 1 through 8 and two fabric channels to each slot in slots 9 through 13. The crossbar switching fabric allows each line card to forward and receive data to every other line card over a unique set of transmission paths.

Figure 2 shows the backplane and line-card connectors for a Cisco Catalyst 6509-E chassis.

Figure 2.   Cisco Catalyst 6509-E Backplane
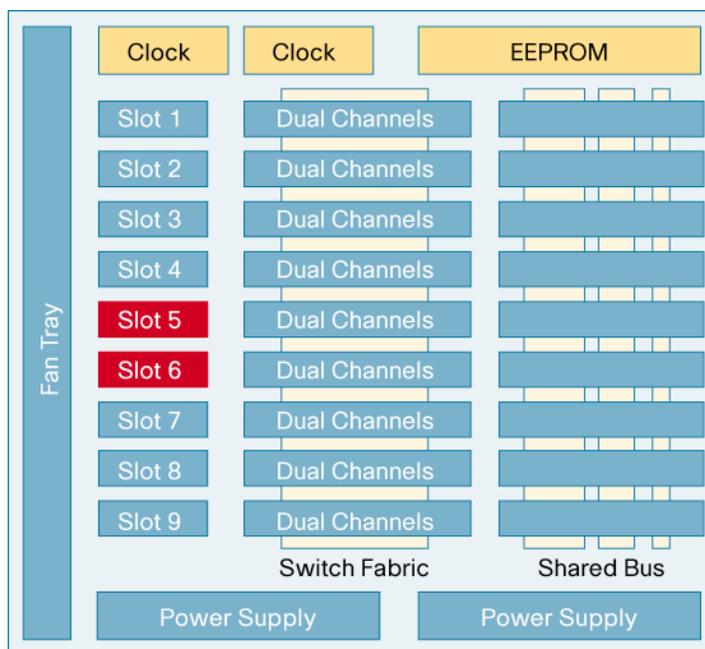
Cisco Catalyst 6500 Fans

The Cisco Catalyst 6500 supports two generations of fans. A new set of high-speed fans was introduced with the announcement of the Supervisor Engine 720. New high-speed fans have been designed for each Cisco Catalyst 6500 chassis, and their primary purpose is to provide additional cooling for new generation line cards that draw more power and generate more heat. If a Supervisor Engine 32, 720, 720-3B, or 720-3BXL is installed in the chassis, then the new high-speed FAN assemblies must be used. The FAN2 assemblies are used with the standard chassis, and the E-Fan is used with the new E series chassis. These high-speed fans can also be used with previous generations of the supervisor engine (1, 1A and 2). (See Figure 3 and Table 2.)

Figure 3.   Cisco Catalyst 6500 Fans

Original Fan   High Speed Fan

Table 2.   Fan Options

Chassis	Original FAN (FAN1)	High-Speed FAN (FAN2)
Cisco Catalyst 6503	FAN-MOD-3	Fan-MOD-3HS
Cisco Catalyst 6503-E	N/A	WS-C6503-E-FAN
Cisco Catalyst 6504-E	N/A	WS-C6504-E-FAN
Cisco Catalyst 6506	WS-C6K-6SLOT-FAN	WS-C6K-6SLOT-FAN2
Cisco Catalyst 6506-E	N/A	WS-C6506-E-FAN
Cisco Catalyst 6509	WS-C6K-9SLOT-FAN	WS-C6K-9SLOT-FAN2
Cisco Catalyst 6509-E	N/A	WS-C6509-E-FAN
Cisco Catalyst 6509-NEB	WS-C6509-NEB-FAN	WS-C6509-NEB-FAN2
Cisco Catalyst 6509-NEB-A	N/A	FAN-MOD-09
Cisco Catalyst 6513	WS-C6K-13SLOT-FAN	WS-C6K-13SLOT-FAN2

Cisco Catalyst 6500 Power Supplies

The Cisco Catalyst 6500 supports a range of AC and DC power supplies to suit a diverse range of customer requirements. The power supplies supported are listed in Tables 3 and 4.

Table 3.   Cisco Catalyst 6500 AC and DC Power Supplies

Power Supply (Part Number)	Cisco Catalyst 6503	Cisco Catalyst 6506	Cisco Catalyst 6509	Cisco Catalyst 6509-NEB	Cisco Catalyst 6509-NEB-A	Cisco Catalyst 6513
950WAC (PWR 950-AC)	X					
950WDC (PWR 950-DC)	X					

Power Supply (Part Number)	Cisco Catalyst 6503	Cisco Catalyst 6506	Cisco Catalyst 6509	Cisco Catalyst 6509-NEB	Cisco Catalyst 6509-NEB-A	Cisco Catalyst 6513
1000WAC (WS CAC-1000W)		X	X	X		
1300WAC (WS CAC-1300W)		X	X	X		
1300WDC (WS CDC-1300W)		X	X	X		
1400WAC (PWR 1400W-AC)	X					
2500WAC (WS CAC-2500W)		X	X	X	X	X
2500WDC (WS CDC-2500W)		X	X	X	X	X
3000WAC (WS CAC-3000W)		X	X	X	X	X
4000WAC (WS CAC-4000W)		X	X	X	X	X
4000WDC (PWR 4000DC)		X	X	X	X	X
6000WAC (WS CAC-6000W)		X	X	X	X	X
8700WAC (WS CAC-8700W-E)		X	X		X	X

Table 4.   Cisco Catalyst 6500 E Series AC and DC Power Supplies

Power Supply (Part Number)	Cisco Catalyst 6503-E	Cisco Catalyst 6504-E	Cisco Catalyst 6506-E	Cisco Catalyst 6509-E
950WAC (PWR-950-AC)	X			
950WDC (PWR-950-DC)	X			
1400WAC (PWR-1400W-AC)	X			
2500WAC (WS-CAC-2500W)			X	X
2500WDC (WS-CDC-2500W)			X	X
2700WAC (PWR-2700W-AC)		X		
3000WAC (WS-CAC-3000W)			X	X
4000WAC (WS-CAC-4000W)			X	X
4000WDC (PWR-4000DC)			X	X
6000WAC (WS-CAC-6000W)			X	X
8700WAC (WS-CAC-8700W-E)			X	X

Each of the power supplies provides a different set of operating characteristics. These are summarized in Table 5.

Table 5.   Cisco Catalyst 6500 Power Supply Characteristics

Power Supply	Part Number	AC Operation	DC Operation	Current Output
950WAC	PWR-950-AC	950W @ 110-220VAC		21.09A
950WDC	PWR-950-DC		950W @ -48 to -60 VDC	21.09A
1000WAC	WS-CAC-1000W	1000W @ 110-220VAC		20.3A
1300WAC	WS-CAC-1300W	1300W @ 110-220VAC		27.46A
1300WDC	WS-CDC-1300W		1300W @ -48 to -60 VDC	27.46A
1400WAC	PWR-1400W-AC	1400W @ 110-220VAC		32.62A
2500WAC	WS-CAC-2500W	2500W @ 220VAC 1300W @ 110VAC		55.5A for 220VAC 27.46A for 110VAC
2500WDC	WS-CDC-2500W		2500W @ -48VDC 2500W @ -60 VDC	55.5A
2700WAC	PWR-2700W-AC	1350W @110V 2700W @220V		63.55A

Power Supply	Part Number	AC Operation	DC Operation	Current Output
3000WAC	WS-CAC-3000W	3000W @ 220VAC 1450W @ 110VAC		65.11 for 220VAC 28.21A for 110VAC
4000WAC	WS-CAC-4000W	4000W @ 220VAC 30A circuit required		91.2A
4000WDC	PWR-4000-DC		1300W @ -48 to -60 VDC	91.2A
6000WAC	WS-CAC-6000W	6000W @ 2 x 220VAC 3000W @ 1 x 220VAC 3000W @ 2 x 110VAC		
8700WAC	WS-CAC-8700W-E	8700W @ 3 x 220VAC 4200W @ 3 x 110VAC 6000W @ 2 x 220VAC 2800W @ 1 x 220VAC 2800W @ 2 x 110VAC		

As noted in the table above, the 2500WAC, 2700WAC, 3000WAC, 6000WAC and 8700WAC power supplies do not deliver full power in an environment running on 110VAC. The 2500W power supply will deliver 1300W at 110VAC, the 2700W will deliver 1350W at 110VAC, the 3000W will deliver 1450W at 110VAC,the 6000W will deliver 3000W with 2 x 110VAC circuits and the 8700W will deliver 4200W with 3 x 110VAC circuits. It is also worth noting that current output is defined as the amount of power made available by the power supplies to the chassis components.

Cisco Catalyst 6500 Supervisor Engines

The Cisco Catalyst 6500 supervisor engine is the primary module where the switch software is loaded and all centralized control and data plane processes are performed. Control plane functions refer to processes that are run in software, whereas data plane functions refer to processes that are run in hardware.

The Supervisor Engine 32 supports a connection to the 32-Gbps shared bus and provides forwarding rates up to 15Mpps. There is no support for a switch fabric in any form with the Supervisor Engine 32. The new Supervisor Engine 720, 720-3B, and 720-3BXL integrate the 720-Gbps crossbar switch fabric onto the supervisor module itself and support a connection into the 32-Gbps bus and a single 20-Gbps connection into the onboard crossbar switch fabric. All Supervisor Engine 720 options support centralized forwarding rates up to 30Mpps and distributed forwarding rates up to 400Mpps.

Supervisor Engine 32

The Supervisor Engine 32 is the most recent addition to the Cisco Catalyst 6500 supervisor engine family. This supervisor engine provides an integrated PFC3B and MSFC2a by default. The PFC3B provides hardware support for security and quality of service (QoS) based access control lists (ACLs). QoS services supported on the PFC3B include ingress traffic policing as well as classification of incoming data allowing the rewrite of class of service (CoS) bits in the Ethernet header and type of service priority bits in the IPV4 header. Performance for these features is supported at a rate of up to 15Mpps.

There are two versions of the Supervisor Engine 32 (Figure 4), one with 8 front Gigabit Ethernet Small Form-Factor Pluggable (SFP) ports and the other with 2 x 10 Gigabit Ethernet front ports. Both of these supervisor engines also have an additional 10/100/1000TX front port. Two USB ports are also present on the front panel, one a type "A" and the other a type "B" USB port.

Figure 4. Supervisor Engine 32-8GE and Supervisor Engine 32-10GE

Supervisor Engine 32-8GE   Supervisor Engine 32-10GE

Cisco Express Forwarding is the forwarding architecture implemented in the Supervisor Engine 32 hardware. As with the Supervisor Engine 720, it also implements hardware registers and control plane policing (CoPP) to limit the effect of denial of service attacks on the control plane. It comes standard with 256MB of Bootflash on the switch processor (SP) side (which can be upgraded to 512MB) and 64MB of Bootflash on the route processor (RP) side. Memory (DRAM) size is 256MB for both the RP and SP and NVRAM is 2MB.

Supervisor Engine 720

The Supervisor Engine 720 was introduced in 2003 and integrates the crossbar switch fabric, Policy Feature Card 3 (PFC3), and Multilayer Switch Feature Card 3 (MSFC3) into the one supervisor module. The PFC3 and MSFC3 are no longer optional.

The crossbar switching fabric on the Supervisor Engine 720 (Figure 5) increases backplane capacity of the crossbar switch fabric from 256 Gbps to 720 Gbps. The crossbar switching fabric supports connections to both the earlier fabric line cards, at 8 Gbps per fabric channel, and the newer fabric line cards, at 20 Gbps per fabric channel. The Supervisor Engine720 also supports classic line cards, thus providing total backward compatibility for all line-card generations. This dual clocking capability preserves customer investment in previous line cards supporting connections into the crossbar switch fabric.

Figure 5. Cisco Catalyst 6500 Supervisor Engine 720

Utilizing the new higher performance line cards with distributed forwarding allows a Supervisor Engine 720 to scale switch performance to 400Mpps.

The Supervisor Engine 720 utilizes the Cisco Express Forwarding architecture to forward packets. Support for up to 30Mpps of Layer 2 and 3 centralized switching of IP traffic is supported. Unlike the earlier PFC1 and PFC2, IPX switching in hardware is not supported on the Supervisor Engine 720 PFC3. The Supervisor Engine 720, however, does still support IPX forwarding in software.

The Supervisor Engine 720 is based on a 600 MHz CPU for the switch processor (SP) and a 600 MHz CPU for the route processor (RP). This supervisor will support up to 1GB of DRAM for the SP and up to 1Gb DRAM for the RP. The default SP bootflash is 512MB, the default RP bootflash is 64 MB and the NVRAM size is 2MB.

Supervisor Engine 720-3B

The Supervisor Engine 720-3B is the most recent addition to the Supervisor Engine 720 family. Architecturally it is the same as the original Supervisor Engine 720 in terms of the switch fabric used and the backplane connections offered. It incorporates a new PFC3B, which increases the functionality of the supervisor engine over its predecessor. It differs from the original Supervisor Engine 720 mainly in functionality. Some of the features that differentiate it from the earlier Supervisor Engine 720 include:

- MPLS support in hardware
- Support for EoMPLS
- Support for Security ACL hit counters
- Multipath URPF check now performed in hardware
- Increased has efficiency (from 50 percent to 90 percent) for storing NetFlow entries in the NetFlow table
- Increased support for ACL labels from 512 to 4096
- QoS policies can now be applied on tunnel interfaces
- Layer 2 ACLs can be applied to IPV4 traffic
- Support for matching on CoS and VLAN in ACLs is now supported
- Support for up to 256K multicast routes in sparse mode

Supervisor Engine 720-3BXL

The Supervisor Engine 720-3BXL was introduced early in calendar year 2004. It is functionally identical to the Supervisor Engine 720-3B, but differs in its capacity for supporting routes and NetFlow entries. Up to 1 million routes can be stored in its forwarding tables and up to 256K NetFlow entries can be stored in the NetFlow tables.

Multilayer Switch Feature Card (MSFC)

The control plane functions in the Cisco Catalyst 6500 are processed by the MSFC and include handling Layer 3 routing protocols, maintaining the routing table, some access control, flow initiation, and other services not found in hardware. Performance of the control plane is dependent on the type and number of processes running on the MSFC. The MSFC3 can support forwarding rates up to 500Kpps.

The MSFC3 is an integrated component on the Supervisor Engine 720, Supervisor Engine 720-3B, and Supervisor Engine 720-3BXL. It is integrated onto all of the Supervisor Engine 720 options along with the PFC3 and crossbar switch fabric. The MSFC3 is shown on a Supervisor Engine 720 in Figure 6.

Figure 6.   MSFC3 on a Supervisor Engine 720

On the MSFC daughter card, the route processor (RP) is located on the MSFC itself. The RP is responsible for a number of processes, including running the Layer 3 routing protocols,

performing address resolution, running ICMP, managing the virtual interfaces (that is, switched virtual interfaces), and the Cisco IOS® Software configuration. The SP is primarily responsible for running the Layer 2 protocols like spanning tree, VLAN Trunking protocol, Cisco Discovery Protocol, and so on as well as programming the FIB tables onto the PFC.

While the multilayer switch feature card maintains the routing tables, it does not actively participate in the forwarding of packets. The MSFC3 still participates in communicating with routing peers to determine the network topology and maintain the routing tables. From the routing tables, the MSFC3 will create a Cisco Express Forwarding table (also known as a Forwarding Information Base—or FIB table) and push this down to the PFC and any DFCs present in the chassis.

The memory configurations across the different generations of supervisor (and MSFC) are detailed in Table 6.

Table 6.   Supervisor Engine and MSFC Memory Facts

Component	Supervisor Engine 32-8GE	Supervisor Engine 32-10GE	Supervisor Engine 720	Supervisor Engine 720-3B	Supervisor Engine 720-3BXL
SDRAM					
Default	SP: 256MB/1GB	SP: 256MB/1GB	SP: 512MB/1GB	SP: 512MB/1GB	SP: 512MB/1GB
Maximum	RP: 256MB/1GB	RP: 256MB/1GB	RP: 512MB/1GB	RP: 512MB/1GB	RP: 512MB/1GB
Bootflash	RP: 64MB SP: 256MB	RP: 64MB SP: 256MB	RP: 64MB SP: 512MB*	RP: 64MB SP: 512MB*	RP: 64MB SP: 512MB*
NVRAM	SP: 2MB RP: 2MB	SP: 2MB RP: 2MB	SP: 2MB RP: 2MB	SP: 2MB RP: 2MB	SP: 2MB RP: 2MB

 * Note that 512MB of DRAM is standard on the Sup720 when ordering IOS 12.2(18)SXE or later

Policy Feature Card

Complementing the MSFC is the policy feature card (PFC). The PFC is a daughter card that sits on the supervisor base board and contains the ASICs that are used to accelerate Layer 2 and Layer 3 switching, store and process QoS and security ACLs, and maintain NetFlow statistics.

The Policy Feature Card 3 (PFC3) is a standard inclusion with the Supervisor Engine 720 and provides centralized forwarding performance up to 30Mpps. It contains a Layer 2 and a Layer 3 forwarding engine. The Layer 2 engine is responsible for:

- Layer 2 MAC address lookups into the Layer 2 CAM table.
- Looking into the packet headers to determine if this switching operation will be a Layer 2 or a Layer 3 operation. If it is going to be a Layer 3 operation, then it will hand off the packet to the Layer 3 engine for further processing.

The Layer 3 Engine is responsible for:

- NetFlow Statistics collection.
- Hardware based forwarding of IPv4, IPv6 and MPLS tagged packets
- QoS mechanism for ACL classification, marking of packets, and policing (rate limiting).
- Security mechanism for validating ACL rules against incoming packets.
- Maintaining Adjacency entries and statistics
- Maintaining Security ACL counters

The PFC3 (Figure 7) supports hardware based Layer 2 and Layer 3 switching, processing security and QoS ACLs in hardware and the collection of NetFlow statistics.

Figure 7.   Policy Feature Card 3 (PFC3)

The PFC3 has also been enhanced and includes support for a set of new features that are now processed in hardware. These features include:

- Network Address Translation (NAT) and Port Address Translation (PAT)
- Multi-Path Unicast Reverse Path Forwarding Check (uRPF)
- Generic Route Encapsulation (GRE)
- IP Version 6 (IPV6) Switching
- IPV6 ACLs
- IPV6 Tunneling
- Multi-Protocol Label Switching (MPLS)
    ○ MPLS Virtual Private Networking (MPLS VPN)
    ○ MPLS Provider and Provider Edge (MPLS P/PE Support)
    ○ MPLS Traffic Engineering (MPLS TE Support)
- Bi Direction Protocol Independent Multicast (BI-DIR PIM)
- User Based Rate Limiting (UBRL)
- Egress Policing

The PFC3 also enhances some capabilities of the PFC2, including:

- Increasing the number of HSRP groups from 16 (in the PFC2) to 255
- Supports a unique MAC address per Layer 3 VLAN interface

The Policy Feature Card 3B (PFC3B) is an enhanced version of the PFC3 that adds a number of new functions in hardware, as well as improving the efficiency levels for storing flow entries in the NetFlow table. The most significant enhancement provided as part of this new daughter card is hardware switching support for MPLS tagged packets. This enables any local Ethernet line card to be able to receive and send MPLS tagged packets. This PFC also now supports Ethernet over MPLS natively.

The PFC3B adds support for a number of new QoS enhancements, including support for applying QoS policies on tunnel interfaces, increasing support for ACL labels to 4096, and allowing ACLs to match on CoS and VLAN values on incoming packets.

The efficiency of the NetFlow algorithm for storing entries in the NetFlow table has been improved in the PFC3B hardware. The PFC3B can now effectively yield up to a 90 percent utilization rate of

the NetFlow table, allowing up to 115K NetFlow entries to be stored in the 128K available table entries.

The PFC3B is an upgradeable option and as such, a Supervisor Engine 720 can utilize this PFC3B to turn the supervisor into a Supervisor Engine 720-3B.

The Policy Feature Card 3BXL (PFC3BXL) is functionally equivalent to the PFC3B, but adds support for up to 1 million routes in its forwarding tables, as well as up to 256K NetFlow entries in its NetFlow tables. Multicast support has also been improved, by providing a way to support up to 256K multicast routes when the switch is running in sparse mode. This is an increase from the effective maximum of 64K entries provided by the original PFC3 and PFC3B. Like the PFC3B, the PFC3BXL is an upgradeable option and as such, a Supervisor Engine 720 can utilize this PFC3BXL to turn the supervisor into a Supervisor Engine 720-3BXL.

Hardware Based Rate Limiters for Control Plane Protection

Another new feature inherent in all the PFC3 hardware options is built-in denial of service (DOS) rate limiters. DOS attacks can slow down or even halt operations of a host, and with the importance of switches in networks, the need to protect them from these types of attacks is becoming more important. The Supervisor Engine 720 incorporates a rate limiting function to limit the impact of a DOS attack. This functionality applies to both unicast and multicast traffic. Table 7 shows summary traffic that the CPU will apply rate limiters to if that traffic goes above a certain threshold.

Table 7. PFC3 Denial of Service Rate Limiters

Traffic Type	Description
Input and Output ACL	Applies to traffic sent to processes that use ACLs—like NAT, ACL Logging, reflexive ACLs, TCP Intercept, etc
Cisco Express Forwarding Receive	Traffic destined for the CPU
Cisco Express Forwarding Glean	Throttling ARPs to unknown hosts
IP Features	Traffic requiring CPU processing like IPSec, Authentication Proxy, Dynamic ARP Inspection, etc
MTU Failures	Packets requiring fragmentation
ICMP Redirect	Packets that require ICMP redirect message
ICMP Unreachable	Limits traffic due to FIB Miss, ACL drop, etc
IP Errors	Limits bad traffic with invalid length or checksums
VACL Logging	When CLI notification is required for packets denied by the VACL process
IP RPF Failure	Packets failing the RPF check
L3 Security Features	Features like Content Based Access Control (CBAC), IPSec traffic, Authentication proxy, etc
IP Options	Packets with IP Options set
TTL Failure	Where the TTL = 1
Capture Packet	Packets punted to CPU when Optimized ACL logging (OAL) enabled
MTU Failure	Packets with MTU larger than interface MTU
L2 PDU	Limits Layer 2 packets like STP, VTP, CDP, UDLD, etc
L2 Tunneling Protocol	Limits tunneled packets destined for CPU like STP, VTP and CDP
Multicast IGMP	Limits IGMP Control Messages sent to CPU
Multicast Partial SC	Limits multicast packets requiring partial CPU processing if special processing is required
Multicast Default Adjacency	Limits Multicast traffic requiring special processing due to FIB miss
Multicast Non-RPF	Multicast traffic failing RPF check

Traffic Type	Description
Multicast IP Options	Multicast traffic with IP Options set

Cisco Express Forwarding vs. Flow Based Switching

Cisco Express Forwarding is now the default forwarding architecture for the Supervisor Engine 32, Supervisor Engine 720, Supervisor Engine 720-3B and Supervisor Engine 720-3BXL. Flow based switching is found in many switching architectures on the market today. Flow based switching has an inherent problem in that it relies on the control plane to forward the first packet of each new flow that enters the switch. In today's application space, many flows are short lived, and combined with increased data load, place a larger burden on the switch control plane. The Cisco Catalyst 6500 has always focused on providing a very high performance control plane. However, control plane performance while in the thousands of packets per second, does not approach the performance provided by hardware based switching (normally in the millions of packets per second). In many customer environments, flow based switching can impose a bottleneck on overall throughput. Cisco devised a new forwarding architecture to greatly enhance the forwarding capabilities of the Cisco Catalyst 6500 switching architecture and eliminate the control place from the forwarding path. (See Figure 8.)

Figure 8. Flow Based vs. Cisco Express Forwarding Switching
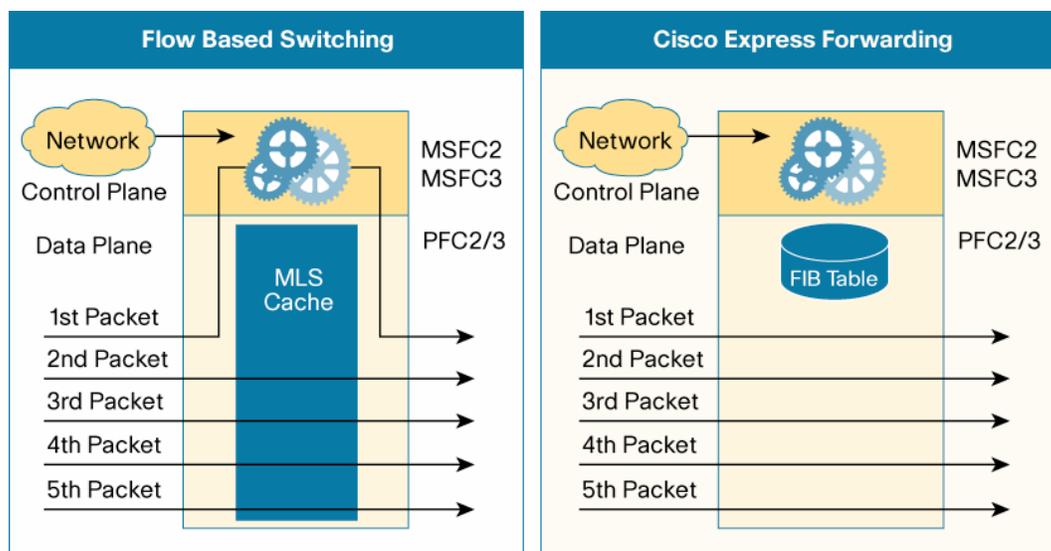
The Cisco Express Forwarding architecture allows the control plane to do what it does best: interact with its routing peers to build a picture of the topology. From this topology, the MSFC builds a Forwarding Information Base (FIB) that is pushed down to the PFC (and any installed DFCs) and programmed into hardware in specialized high performance lookup memory called TCAM. At all times, the PFC has full knowledge of the topology and can make informed decisions on where to forward data. Should the topology of the network change, then the FIB is modified and passed to the PFC and DFCs keeping it current at all times

Cisco Catalyst 6500 Line Cards

The lineup of Cisco Catalyst 6500 line cards provides a full complement of media and speed options to meet the needs of deployment in the access, distribution and core layers of the network. Each line-card slot provides a connection into the 32-Gbps shared bus and the crossbar switch

fabric (if either a Supervisor Engine 720 or switch fabric module is present). Cisco Catalyst 6500 line cards fall into one of four general families of line cards:

- **Classic**: In this mode the line card has a single connection into the 32-Gbps shared bus.
- **CEF256**: The line card in this mode supports a connection into the 32-Gbps shared bus and the switch fabric—these line cards will use the switch fabric for data switching when the Supervisor Engine 720 is present—if a Supervisor Engine 32 is present it will revert back to using the 32-Gbps shared bus.
- **CEF256 and CEF720**: The line card in this mode supports a connection into the 32-Gbps shared bus and the switch fabric: these line cards will use the switch fabric on the Supervisor Engine720 for data switching.
- **dCEF256**: These line cards require the presence of the switch fabric to operate—these line cards do not connect into the shared bus.
- **dCEF720**: Like the dCEF256 linecards, they only require the switch fabric to be present to switch packets. They connect into the switch fabric channels at 20Gbps as opposed to the 8Gbps that the dCEF256 linecards connect.

The architecture of the Cisco Catalyst 6500 is distributed in the sense that not all switching functions are performed on the supervisor. The line cards themselves have intelligence built into the on board ASICs to better scale functions such as:

- Port based buffering
- Port QoS Classification allowing priority bits to be modified
- Input queuing and scheduling
- Congestion management of individual queues using congestion avoidance algorithms like Weighted Random Early Discard (WRED)
- Packet rewrite on transmission of the packet
- Remarking of Type of Service and Class of Service priority bits on egress
- QoS scheduling using Strict Priority and Weighted Round Robin scheduling algorithms

Cisco Catalyst 6500 line cards also use other nomenclature to describe both their backplane connection and forwarding architecture. (See Table 8.)

Table 8.   Cisco Catalyst 6500 Line-Card Types

| Line card	Bus Connection	Fabric Connection	Distributed Forwarding
Classic	Yes	No	No
CEF256	Yes	Yes at 8 Gbps	Upgradeable to DFC, DFC3a, DFC3B, or DFC3BXL
dCEF256	No	Yes at 8 Gbps	Yes, integrated DFC, DFC3a, DFC3B or DFC3BXL,
CEF720	Yes	Yes at 20 Gbps	Upgradeable to DFC3a, DFC3B, or DFC3BXL
dCEF720	No	Yes at 20Gbps	Yes, either a DFC3C or DFC3CXL is present

Most importantly, all of the above line cards can interoperate with one another in the same chassis. The different Cisco Express Forwarding based forwarding architectures that are mentioned in the above table will be explored in a little more detail later in this paper.

Cisco Catalyst 6500 Architecture: Shared Bus

The Cisco Catalyst 6500 shared 32-Gbps bus allows all connected ports to both transmit and receive data. The switching bus is actually comprised of three discrete buses, each serving a

specific function in the switching operation, namely the data bus (DBus), results bus (RBus) and the control bus (CBus). The backplane connectors to the shared bus on the line card are located on the rear right hand side of the card. This can be seen in the diagram below which shows two line cards, the first being a classic line card (shown on the left) and the second being a fabric (CEF256) line card (on the right). The shared bus connector is located in the same position on both cards and is noted in Figure 9 by the red box.

Figure 9.   Shared Bus Connector
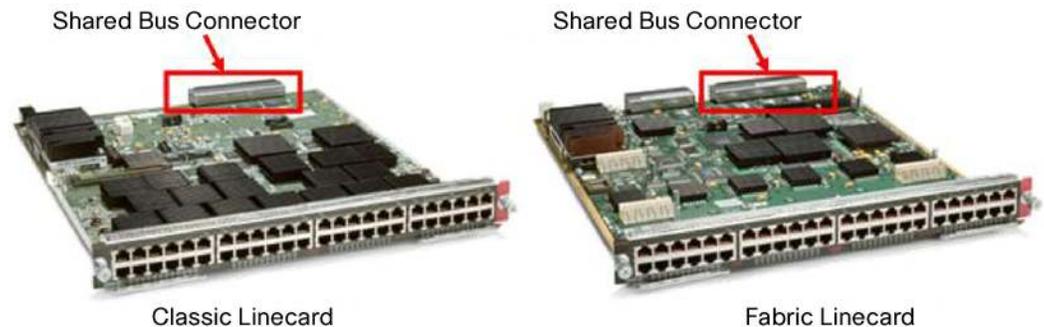
The DBus is the main bus over which all data is transmitted. To transmit the packet, the line card will arbitrate for access to the DBus. It does this by sending a request to a master arbitrator that resides on the primary supervisor. If the bus is not in use, then the master arbiter will grant access allowing the line card to place the frame on the DBus. Once a line card has successfully arbitrated for access to the DBus, it will forward the packet over the DBus to the supervisor. In the same operation, all connected line cards will both see the packet being transmitted and store a copy of the packet. The DBus is the bus that realizes the bandwidth of 32 Gbps. The RBus is the bus that is used by the supervisor engine to forward the result of the forwarding operation to each of the attached line cards. The supervisor will send to each line card either a forward or drop result which each line card will use to either flush the packet from its buffer or forward it onto its destination. The CBus is the control bus that relays control information between the line cards and the Network Management Processor (NMP).

The shared switching bus on the Cisco Catalyst 6500 has evolved from earlier generations of the shared bus architecture found in the Cisco Catalyst 5000 and has improved performance using two features: Pipelining and Burst Mode.

The shared bus used to support a single frame transmission across the bus at any one point in time. After the supervisor received the packet it could initiate a lookup into its local forwarding tables to determine which port the packet should be switched to. Then it would send the result of the lookup to all ports connected to the bus. While the lookup occurred, no subsequent packets could pass over the bus, hence there was some idle time where the use of the bus was not being maximized. Pipelining allows ports to send up to 31 frames across the shared bus before a result is sent via the RBus. Should a 32nd packet be sent, it will be queued locally at the port until a result is received over the RBus. Pipelining reduces idle time on the bus and also improves the overall performance of the shared bus architecture.

The other aspect of using a shared bus is that the bus usage could unfairly be skewed towards ports sending larger frames. If, for example, two ports are requesting access to the bus, and port A is sending 256-byte frames and Port B is sending 512-byte frames, then port B will gain an unfair advantage over a period of time by consuming more bandwidth. This situation is overcome using

the Burst Mode feature of the shared bus. The port ASIC maintains a counter of the byte count it has transmitted and compares this to a local threshold. Provided that the byte count is below the threshold, then the packet can be forwarded. If the threshold is exceeded, then the port ASIC will cease transmitting frames and the arbitration logic removes bus access for this port. The threshold is calculated by the port based on a number of local variables in the system and is calculated automatically to ensure fair distribution.

Cisco Catalyst 6500 Architecture: Crossbar Switching Fabric

The Cisco Catalyst 6500 supports two forms of the crossbar switching fabric. The switching fabric module (SFM) implemented the initial fabric type first available in the Cisco Catalyst 6500. The second implementation was implemented on the Supervisor Engine 720 (and Supervisor Engine 720-3B and Supervisor Engine 720-3BXL). (See Figure 10.)

Figure 10.   Cisco Catalyst 6509 Crossbar Switch Fabric Module Fabric Channel Layout of the Switch Fabric Module
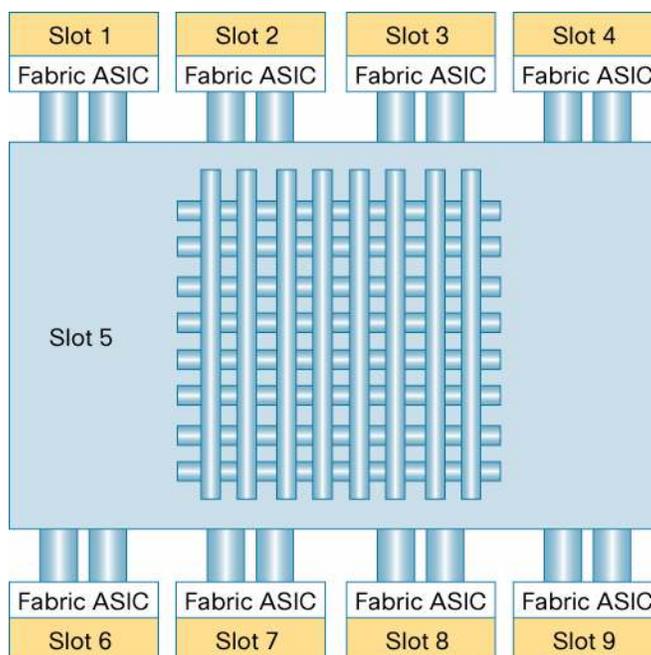
The crossbar switching fabric in both the SFM and in the Supervisor Engine 720, 720-3B, and 720-3BXL is comprised of 18 individual fabric channels that are apportioned across each of the line-card slots in the chassis. For the three, four, six and nine slot chassis, each linecard slot get two fabric channels into the crossbar switch fabric. The nine slot chassis fabric channel layout is shown in the above diagram. In the above diagram, slot 5 is where the Supervisor Engine 720 is located. The Cisco Catalyst 6513 is the exception to the two fabric channels per line-card slot rule. In the Cisco Catalyst 6513, the fabric channels for the crossbar switch fabric are laid out as follows:

- Slots 1 to 8 get a single fabric channel
- Slots 9 to 13 get dual fabric channels

By slots 1 through 8 having a single fabric channel precludes them from housing the dual fabric line cards. Examples of the dual fabric line cards include:

- WS-X6816-GBIC: 16 port GBIC based Gigabit Ethernet (supports 2 x 8Gb fabric channels)

- WS-X6748-GE-TX: 48 port 10/100/1000 RJ45 Gigabit Ethernet (supports 2 x 20Gb fabric channels)
- WS-X6704-10GE: 4 port 10 Gigabit Ethernet (supports 2 x 20Gb fabric channels)
- WS-X6708-10G-3C: 8 port 10 Gigabit Ethernet (supports 2 x 20Gb fabric channels)
- WS-X6708-10G-3CXL: 8 port 10 Gigabit Ethernet (supports 2 x 20Gb fabric channels)
- WS-X6748-SFP: 24 port SFP based Gigabit Ethernet (supports 1 x 20Gb fabric channels)

The fabric trace layout for the Cisco Catalyst 6513 is shown in Figure 11.

Figure 11. Cisco Catalyst 6513 Crossbar Switch Fabric Module Fabric Channel Layout for the Switch Fabric Module
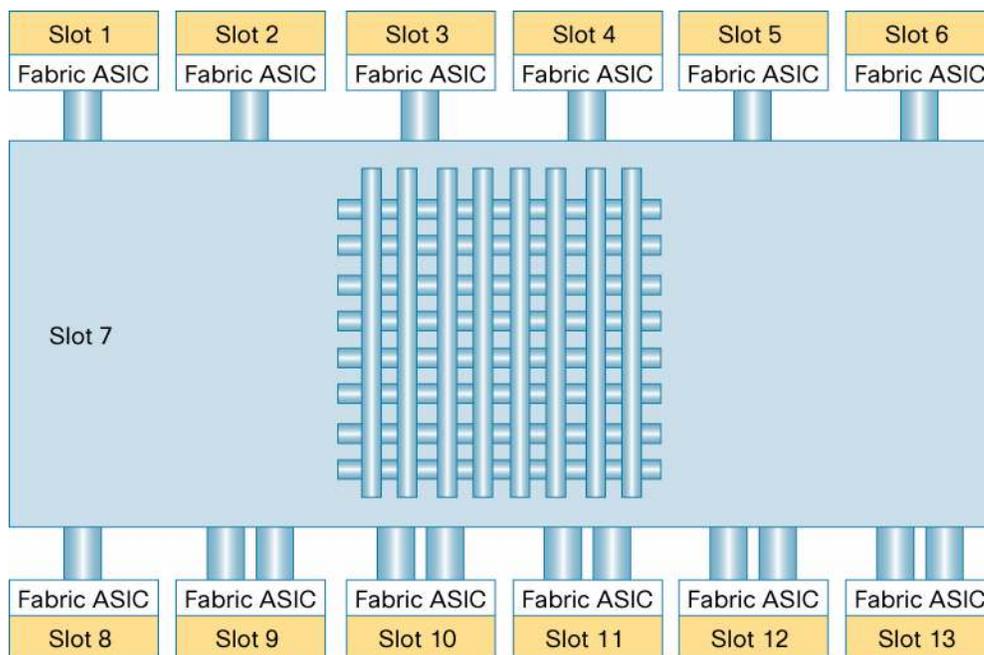
The crossbar switching fabric connector on the line cards is located at the rear left of the card. Figure 12 shows the location of the crossbar connector on a fabric module.

Figure 12. Crossbar Connector on a Fabric Line Card
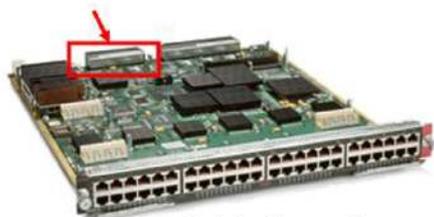
As discussed earlier, the crossbar switch fabric is implemented with the integrated switch fabric on the Supervisor Engine 720, 720-3B, or 720-3BXL. The switch fabric integrated onto the Supervisor Engine 720 clocks each fabric channel at either 8 Gbps or 20 Gbps. The clocking speed is determined by the presence of a specific line card (i.e. CEF256 or CEF720 as an example). The CEF256 and dCEF256 line cards cause the Supervisor Engine 720 switch fabric to clock those fabric channels at 8 Gbps. The CEF720 and dCEF720 line cards will result in the

Supervisor Engine 720 switch fabric clocking those fabric channels at 20 Gbps. The Supervisor Engine 720 switch fabric can simultaneously support fabric channels to different slots at different clock speeds.

This means that the Supervisor Engine 720 could, for instance, support a single fabric channel to a CEF256 line card at 8 Gbps as well as the two fabric channels to a CEF720 line card at 2 x 20 Gbps. This allows all of the Supervisor Engine 720 modules to deliver bandwidth of 40 Gbps to each line-card slot. (See Figure 13.)

Figure 13.   Integrated Switch Fabric on the Supervisor Engine 720

The Cisco crossbar switch fabric architecture uses a combination of buffering and over-speed to overcome any potential congestion and head-of-line blocking conditions. Over-speed is used to clock the paths "internal" to the switch fabric at a speed higher that the fabric channel into the switch fabric. This means that for the Supervisor Engine 720 switch fabric, the internal path is clocked at 60 Gbps for external fabric channels that clock at 20 Gbps. Over-speed is a technique used to accelerate packet switching through the switch fabric to minimize the impact of congestion.

Line rate buffering and queues are also present internally within the switch fabric to overcome any temporary periods of congestion. Buffering is implemented on egress in the switch fabric to assist in eliminating head of line blocking conditions.

Cisco Catalyst 6500 Architecture: Bus Switching Modes

There are three switching modes used by the BUS and fabric ASICs present on CEF256 and CEF720 line cards. These modes determine the header format that is used to transmit the data across the DBus and communicate with other CEF256 and CEF720 line cards. These modes do not apply to line cards that use a DFC. These modes are discussed in more detail below.

Flow-Through Mode

This mode of operation is used by CEF256 modules when there is no crossbar switch fabric present. It enables CEF256 modules to operate as if they were classic line cards. In flow through mode, the entire packet (header plus data) is forwarded by the line card to the supervisor for processing. This mode does not apply to CEF720 modules. When flow through mode is used, performance levels up to 15Mpps can be achieved.

Compact Mode

The compact mode of operation requires a crossbar switch fabric to be present in the system. This can be provided by a switch fabric module, Switch Fabric Module 2 or a Supervisor Engine 720. All modules in the chassis must be fabric enabled (i.e. CEF256, dCEF256, or CEF720) for the switch to operate in Compact Mode. Classic line cards installed in the chassis will negate the ability of the switch to run in Compact Mode. In Compact Mode, only the header is passed over the DBus to the supervisor. The Header is compressed prior to being placed on the DBus which increases the bandwidth available for header transmission. The data portion of the packet is transmitted over the crossbar switch fabric channels. In this mode of operation, the switch can achieve centralized performance of up to 30Mpps independent of packet size.

Truncated Mode

Truncated Mode is used when CEF256 and/or CEF720 line cards are installed in a chassis with a classic line card, but a crossbar switch fabric is present. In this mode, classic line cards will transmit both the header and the data portion of the packet over the DBus. CEF256 and CEF720 line cards will only transmit headers over the DBus and transmit the data portion of the packet over the crossbar switch fabric.

Truncated mode results in centralized forwarding rates up to 15Mpps. In Truncated Mode, since the CEF256 and CEF720 line cards use the crossbar switch fabric to transmit data, overall aggregate bandwidth can actually be higher than the 32-Gbps shared bus capacity.

Performance of DFC enabled line cards are not affected by truncated mode and performance remains the same regardless of the line-card mix in the chassis.

Cisco Catalyst 6500 Architecture: Classic Line Cards

Classic line cards support a connection to the 32-Gbps shared bus but do not have any connections into the crossbar switch fabric. Classic line cards are supported by all generations of the supervisor engines, from the Supervisor Engine 1 through to the Supervisor Engine 720-3BXL. It is worthwhile checking release notes to determine current software support for a given supervisor. When a classic line card is present in a Cisco Catalyst 6500 chassis, it will not allow line cards to operate in compact mode, thus the centralized performance of the PFC will be up to 15Mpps. (See Figure 14.)

Figure 14.   Classic Line Card
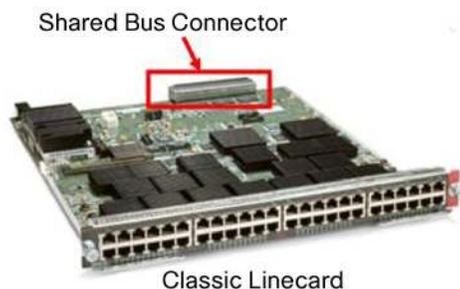
Cisco Catalyst 6500 Architecture: CEF256 Line Cards

CEF256 fabric enabled line cards that connect into the crossbar switch fabric also utilize a local switching bus on the line card itself. The local switching bus is similar in function and operation to the shared 32-Gbps bus that is found connecting to all shared bus capable line cards in a Cisco Catalyst 6500 chassis. The local switched bus resident on the line card is utilized for local switching. By using this bus, a locally switched packet (using a Distributed Forwarding Card (DFC) to determine the forwarding destination) can avoid being transmitted over the chassis shared bus or the crossbar switch fabric. This both reduces overall latency of switching the packet and frees up backplane capacity for those line cards that do not have local switching capabilities.

Fabric enabled line cards that have a single 8-Gbps fabric channel connection to the crossbar switch fabric use a single 16-Gbps shared bus over which local packets are switched. These line cards can also utilize an optional Distributed Forwarding Card (DFC or DFC3a) to provision local switching of packets. The CEF256 line-card architecture is shown in Figure 15.

Figure 15.   CEF256 Line-Card Architecture
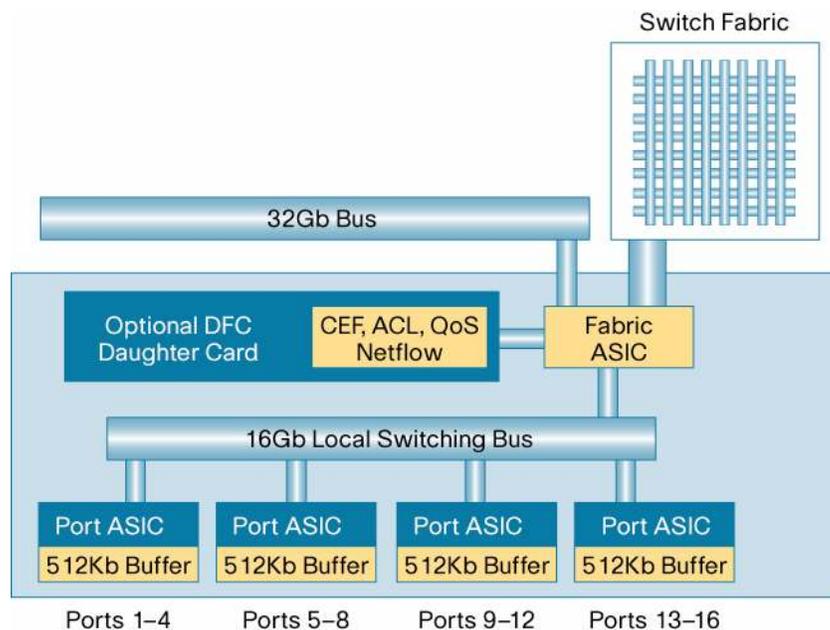
The CEF256 line-card architecture includes a fabric Interface ASIC, which is used as the interface between the ports on the line card and the crossbar switch fabric. This line-card architecture also allows the fabric Interface ASIC to connect into the shared bus. The function of the fabric Interface ASIC is explored in the packet walk section in more detail later in this paper.

Cisco Catalyst 6500 Architecture: dCEF256 Line Cards

The WS-X6816-GBIC line card is an example of a dCEF256 line card that supports dual fabric connections and implements two local 16-Gbps shared buses.

Each of the 16-Gbps shared buses serve one half of the 16 Gigabit Ethernet ports on the line card. Packets switched between each group of 8 ports are locally switched over a 16-Gbps bus and do not venture outside of the line card. Packets that are switched from one block of ports to the other block of ports on the same line card are switched through the crossbar switch fabric. The 6816 line card has an integrated DFC and can also support a DFC3a when used with the Supervisor Engine 720. The architecture of this line card is displayed in Figure 16.

Figure 16.   Dual Fabric Line Card (WS-X6816-GBIC)

In the dCEF256 line-card architecture, the fabric ASIC is used as the interface between the ports on the line card and the crossbar switch fabric. This module architecture has no connection to the 32-Gbps shared bus. (See Figure 17.)

Figure 17.   Line-Card architecture of the WS-X6816-GBIC

Ports 1–4   Ports 5–8   Ports 9–12   Ports 13–16
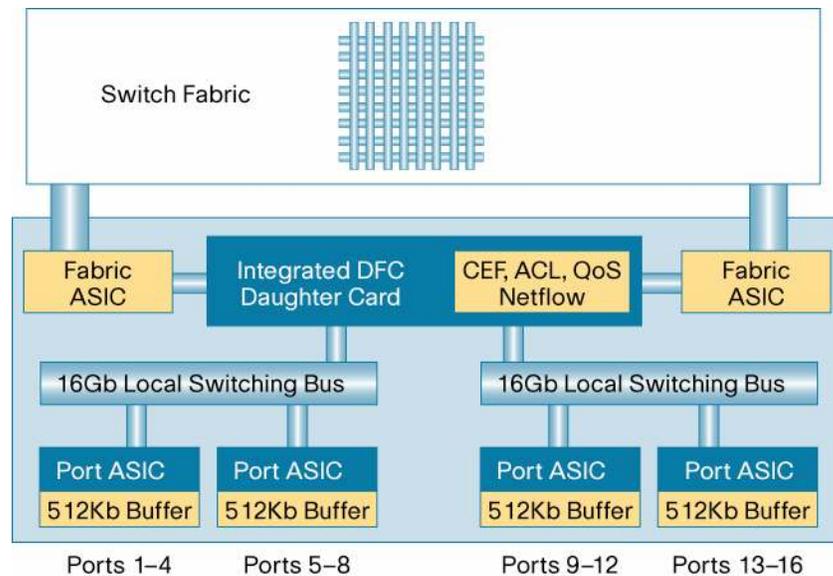
Cisco Catalyst 6500 Architecture: CEF720 Line Cards

In 2003, Cisco announced the next generation Supervisor Engine 720 along with a new series of line cards. They are designed to take advantage of architecture extensions of the Supervisor Engine 720 and the enhanced crossbar switch fabric.

Four Cisco Express Forwarding 720 line cards are available providing higher performance, and higher port densities of high speed interfaces over earlier generations of line cards. The new CEF720 line cards (Figure 18) are:

- **WS-X6724-SFP**: a 24 port Gigabit Ethernet SFP based line card supporting a single 20-Gbps fabric channel to the Supervisor Engine 720 crossbar switch fabric. Also supports an optional Distributed Forwarding Card 3a (DFC3a, DFC3B or DFC3BXL)

- **WS-X6748-GE-TX**: a 48 port 10/100/1000 RJ45 based line card supporting 2 x 20-Gbps fabric channels to the Supervisor Engine 720 crossbar switch fabric. Also supports an optional Distributed Forwarding Card 3a (DFC3a, DFC3B or DFC3BXL)

- **WS-X6748-SFP**: a 48 port 1000Mb SFP based line card supporting 2 x 20-Gbps fabric channels to the Supervisor Engine 720 crossbar switch fabric. Also supports an optional Distributed Forwarding Card 3 (DFC3a, DFC3B or DFC3BXL)

- **WS-X6704-10GE**: a 4 port 10 Gigabit Ethernet Xenpak based line card supporting 2 x 20-Gbps fabric channels to the Supervisor Engine 720 crossbar switch fabric. Also supports an optional Distributed Forwarding Card 3a (DFC3a, DFC3B or DFC3BXL)

- **WS-X6704-10GE**: an 8 port 10 Gigabit Ethernet X2 optic based line card supporting 2 x 20-Gbps fabric channels to the Supervisor Engine 720 crossbar switch fabric. Also supports an integrated Distributed Forwarding Card 3CXL

Figure 18.   CEF720 Line Cards
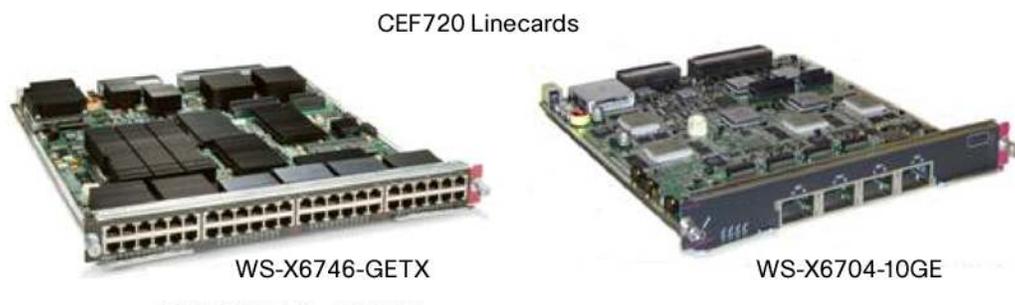
The new line cards utilize a new set of ASICs that provide higher port densities of high speed interfaces (Gigabit Ethernet and 10 Gigabit Ethernet) than previous generations of line cards. The line-card architecture used by CEF720 line cards is shown in Figure 19.

Figure 19.   CEF720 Line-Card Architecture: WS-X6748-GETX Example
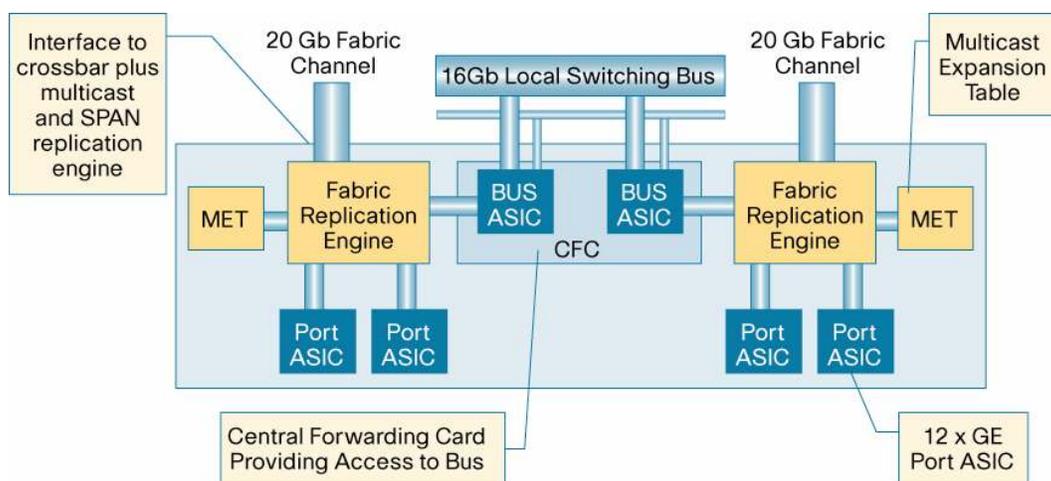
A new fabric ASIC has been developed that replaces the fabric ASIC used in previous fabric line cards. The new fabric ASIC integrates support for Multicast Replication which was previously found in a separate ASIC in earlier line cards. New port ASICs have also been developed to provide for higher port densities of Gigabit Ethernet and 10 Gigabit Ethernet. In the WS-X6724-SFP, WS-X6748-SFP and WS-X6748-GETX a single port ASIC supports 12 local Gigabit Ethernet ports. In the WS-X6704-10GE line card, there are two port ASICs each supporting 2 x 10 Gigabit Ethernet ports. An optional DFC3 can be used with these line cards to provide local switching.

Cisco Catalyst 6500 Architecture: dCEF720 Line Cards

Cisco introduced the WS-X6708-10G-3C and WS-X6708-10G-3CXL line cards in 2006. These linecards provide 8 ports of 10GE using X2 based optics. Unlike the CEF720 linecards, both of these modules ship with an on board DFC. The DFC used on these linecards is either a DFC3C or DFC3CXL. This is an updated DFC daughter card to that used on other current linecards.

Figure 20.   WS-X6708-10G-3C

As with other DFC based linecards, the DFC's used on these linecards provide support for distributed switching up to 48Mpps. This linecard connects into the switch fabric using 2 x 20Gb fabric channels providing a 40Gb connection into the switch backplane. This means that the linecard provides a 2:1 oversubscription which should be factored into any potential designs

Figure 21. dCEF720 Line-Card Architecture: WS-X6708-10G-3C Example
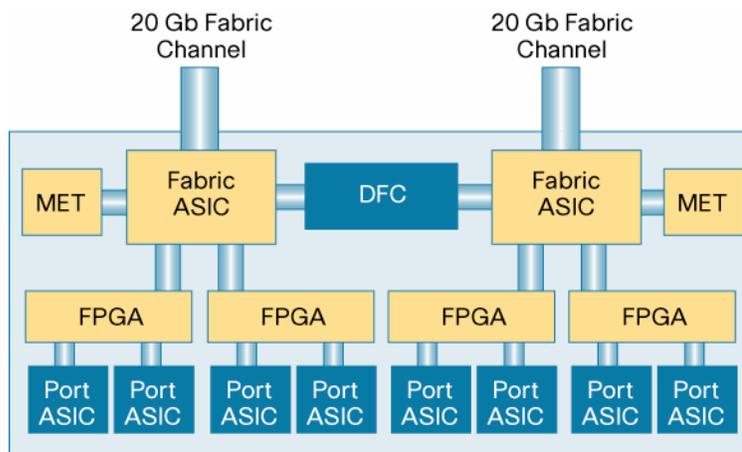
Cisco Catalyst 6500 Architecture: Distributed Forwarding Line Cards

The Distributed Forwarding Card (DFC) can be used on selected line cards to support local switching. The DFC utilizes the same ASICs as those found on the PFC so it supports local Layer 2 and Layer 3 switching, and also holds copies of the access control Lists (ACLs) defined for both QoS and Security. This means that the when switching a packet locally, the DFC can, in parallel, inspect local security and QoS policies defined on the switch and apply those policies to locally switched traffic. (See Figure 22.)

Figure 22. Examples of the Cisco Catalyst 6500 Distributed Forwarding Cards
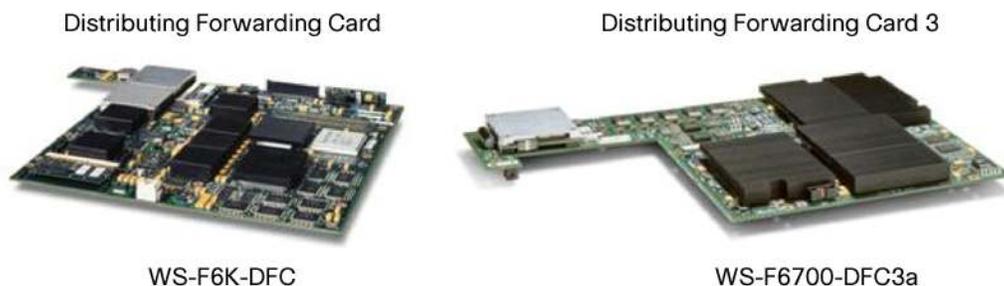
The DFC3a can be used on CEF256 or CEF720 capable line cards when a Supervisor Engine 720 is installed in the chassis. The DFC3a, like the PFC3 can only support hardware based switching of IPV4 and IPV6. The DFC3a does not support the hardware switching of IPX traffic. Two versions of the DFC3a are available, one for CEF256 (part number WS-F6K-DFC3a), one for CEF720

series of line cards (part number WS-F6700-DFC3a). The primary rule for DFC configuration is that different generation forwarding engines cannot be resident in the same chassis at the same time. If a line card with an existing DFC needs to be installed into a Supervisor Engine 720 based chassis, then that DFC will need to be upgraded to a DFC3a. Use of any DFC in the switch chassis will necessitate the use of IOS as the switch OS.

There is also a DFC3B and DFC3BXL available where the PFC in use on the Supervisor Engine 720 is a PFC3B or PFC3BXL. While the mixing of DFC3 and PFC3 is permitted, it is important to understand that when mixed, the PFC/DFC combo will operate at the lowest common denominator. So for instance, if a DFC3a were installed on a system with PFC3BXL, then the PFC3BXL would revert to operating in PFC3a mode. This would have the effect of the PFC3BXL losing its higher capacity for routes and flow entries (back to the levels provided by the PFC3a) and losing any features not found on the DFC3a (i.e. MPLS). The DFC3C and DFC3CXL are integrated with the 8 port 10 Gigabit Ethernet module. These DFC daughter cards are functionally identical to the DFC3B and DFC3BXL respectively.

## Switching Implementations: Making a Switching Decision

This section will explore the process of making a forwarding decision in more detail.

The PFC3 receive the Forwarding Information Base (FIB) from the MSFC3 and install this into a local Cisco Express Forwarding table. This Cisco Express Forwarding table is implemented using Ternary Content Addressable Memory (TCAM) which is designed for very high speed lookups. Even though Cisco Express Forwarding is used to forward packets, a flow entry is still created for each new flow; however, this is used by the NetFlow process to gather statistics.

When a packet enters the PFC2 or PFC3, it performs a lookup into the Cisco Express Forwarding table and performs a longest match lookup based on destination address and the most specific netmask. A packet with a destination address of 192.168.10.1 could result in a lookup for 192.169.10.0/24 rather than the specific destination address. This operation is more efficient as any topology changes are updated immediately by the Cisco Express Forwarding process, making Cisco Express Forwarding more resilient to topology changes.

Both the PFC2 and PFC3 can store up to 256,000 entries in the hardware FIB table. All DFC's can also store the same number of entries in their hardware FIB table. (See Figure 23.)

Figure 23.   Policy Feature Card 3 and Distributed Forwarding Card Functional Elements
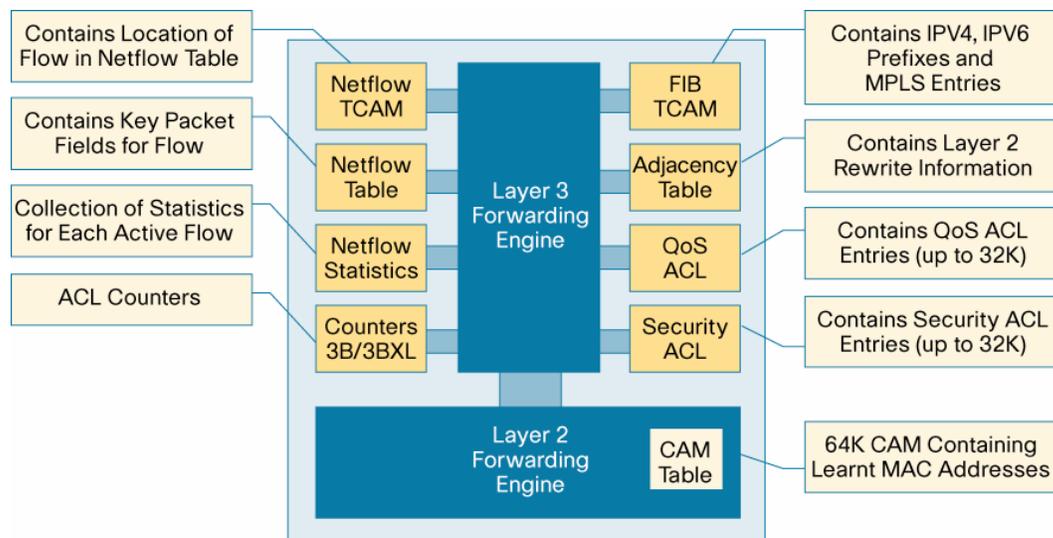
Along with inspecting the FIB table, the PFC3 will also inspect the ACL table to see if any security and/or QoS policies have been defined for this packet flow. This information is retrieved to determine if the packet is to be forwarded or dropped, and if any policing or classification is to be applied to this data. The NetFlow tables are also updated to reflect the statistical counters for this flow, and the adjacency table is inspected to retrieve rewrite information required to modify headers for the packets onward journey.

Quality of Service and Security Access Control Lists

The Cisco Catalyst 6500 implements the processing of QoS and Security ACLs in hardware on all PFCs. An access control List consists of two elements

1. The ACL label which uniquely identifies the ACL: up to 512 unique labels can be supported (up to 4096 ACL labels on the PFC3B and PFC3BXL).

2. The ACL can contain multiple access control entries (ACEs) each defining a permit or deny statement along with some matching characteristics based on Layer 2/3 addressing information and optionally Layer 4 information: up to 32,000 ACEs on the PFC3 can exist across the 512 ACL labels and across the 4096 ACL labels on the PFC3B and PFC3BXL.

To facilitate the fast lookup of ACL entries, Ternary Content Addressable Memory (or TCAM) is used in all PFCs. The TCAM uses a mask to determine where to store the ACL entry. When an ACL is configured, a mask is used to tell the system which of the address bits to check and which bits to ignore. Up to 4K masks on the PFC3 (and PFC3B/PFC3BXL) can be used. The TCAM will use this mask to assist it in storing the ACE in memory. Depending on the number of masks used by the system, actual population of ACEs in the system can vary. If one mask is used, then a maximum of 32K entries can be stored in the TCAM. If more than one mask is used, then the number of TCAM entries is reduced.

Lookups into the TCAM for ACLs is performed in parallel to Layer 2 and Layer 3 lookups. As lookups are performed in hardware, there is no performance hit when ACLs are enabled. Any ACL that is configured in the system is also pushed down the DFC. The DFC can lookup its local TCAM memory for ACLs when performing local switching. The DFC3 supports up to 4K masks and up to 32K ACEs.

Cisco Catalyst 6500: Day in the Life of a Packet

This section will describe the packet walk through the Cisco Catalyst 6500.

Day in the Life of a Packet: Centralized Forwarding

The steps involved in making a centralized forwarding decision are detailed in Figure 22.

Figure 24. Centralized Forwarding Decision
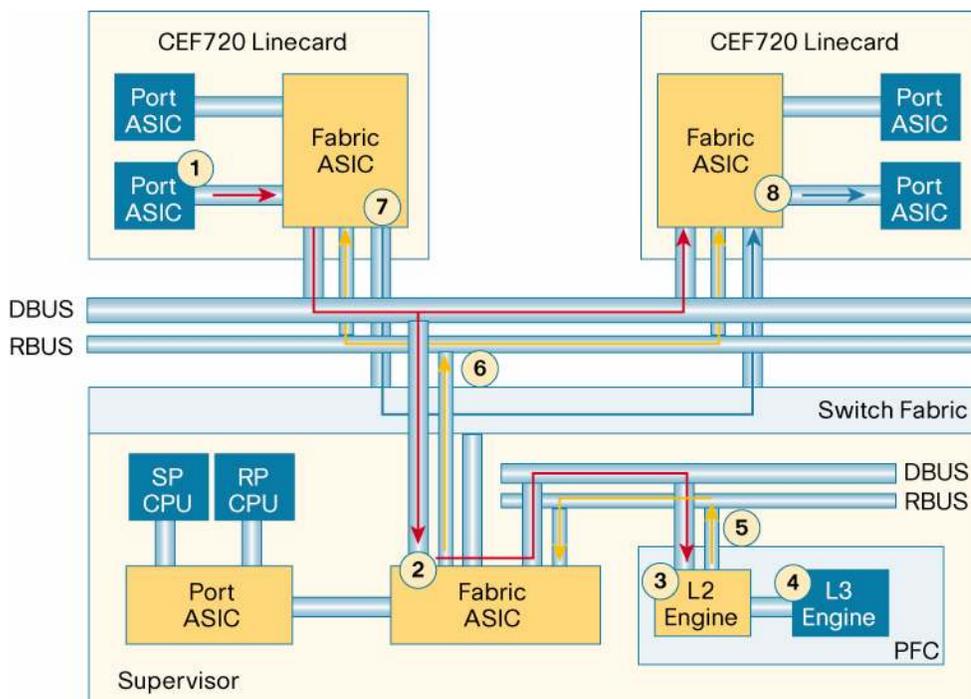
Step 1. The packet arrives at the port and is passed to the fabric ASIC.

Step 2. The fabric ASIC arbitrates for bus access and forwards the header (and not the data payload) over the bus to the supervisor. All line cards connected to the bus will see this header.

Step 3. The supervisor will forward the packet header to the Layer 2 forwarding engine for a Layer 2 lookup.

Step 4. The Layer 2 forwarding engine then forwards the packet to the Layer 3 engine for Layer 3 and 4 processing which includes NetFlow, QoS ACL, Security ACL and Layer 3 lookups.

Step 5. The PFC will combine the results of the multiple lookups and pass the results of the process back to the central supervisor.

Step 6. The supervisor will forward the result of this lookup back over the results bus to all connected line cards.

Step 7. Once the source line card sees the result, it will send the packet data over the switch fabric to the destination line card.

Step 8. The destination line card will receive the packet and forward the data out the destination port.

Day in the Life of a Packet: Distributed Forwarding

The steps involved in making a distributed forwarding decision are detailed in Figure 23.

Figure 25. Distributed Forwarding Decision
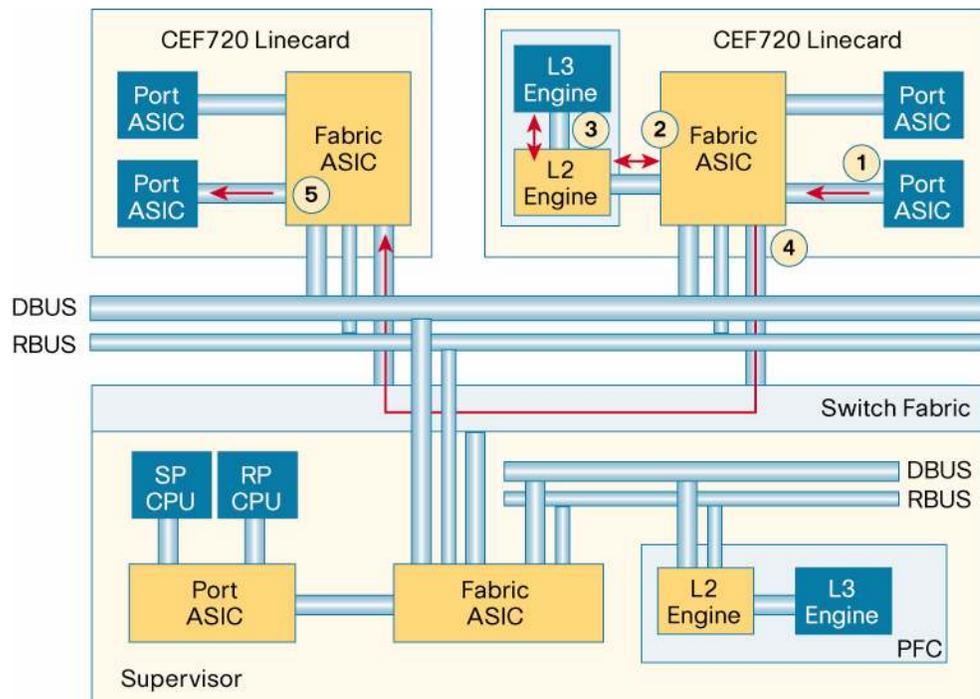
Step 1. The packet arrives at the port and is passed to the fabric ASIC.

Step 2. The fabric ASIC forwards the packet headers are sent to the local DFC.

Step 3. The DFC will perform a forwarding lookup, along with a lookup into the QoS and Security ACLs to determine if any QoS or security policies need to be applied to the packet. The results of the lookup are passed back to the fabric ASIC.

Step 4. The fabric ASIC will forward the packet data over the switch fabric to the destination port.

Step 5. The destination line card will receive the packet and forward the data out the destination port.

Page 27 of 27

CISCO

Printed in USA   C11-395309-01   01/13