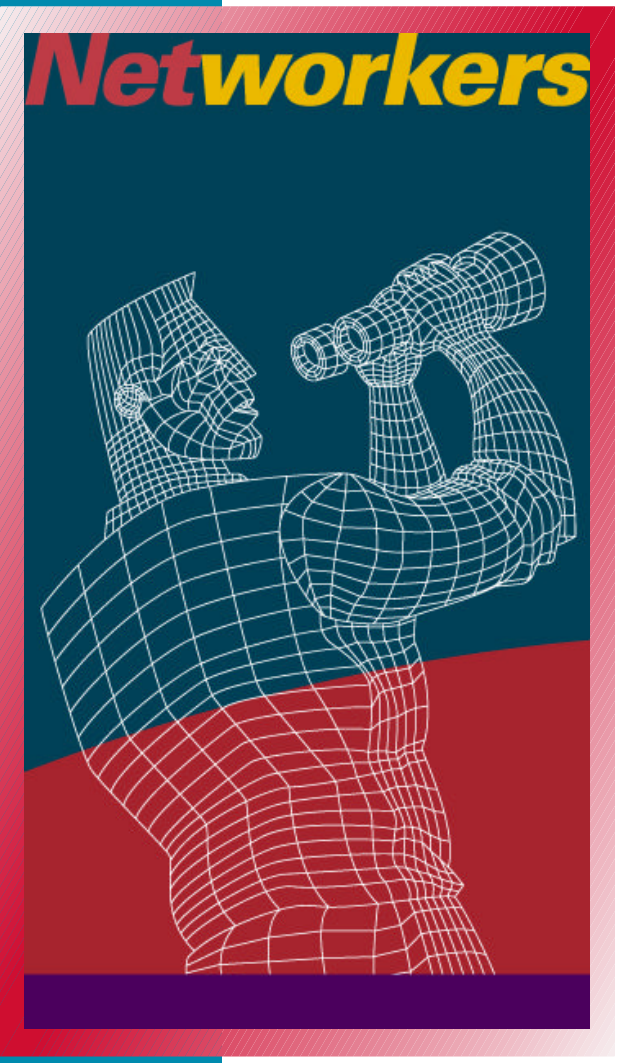

Router Architecture and Performance

Cisco Systems

Agenda

- Perception versus Reality
- Layered Switching
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

Cisco Systems

Perception

Bigger = Better

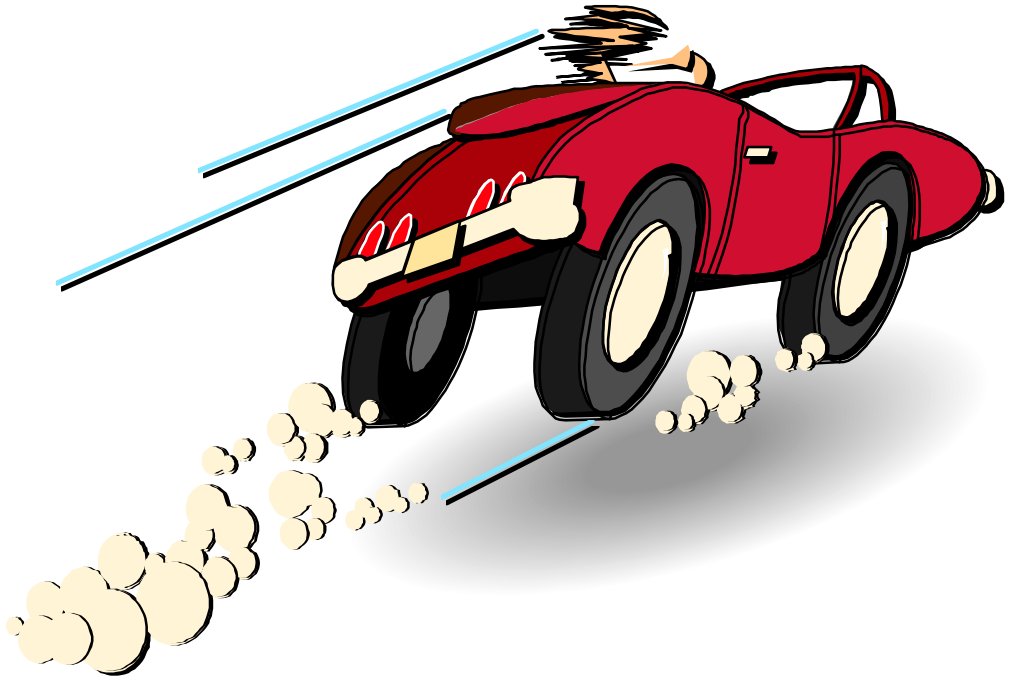

… going 200 MPH

Cisco Systems

Media Characteristics

	IFG	Minimum Valid Frame	Maximum Valid Frame	Bandwidth
Ethernet	96 bits	64 Bytes	1,518 Bytes	10 Mbps
Fast Ethernet	96 bits	64 Bytes	1,518 Bytes	100 Mbps
FDDI	0	34 Bytes	4,500 Bytes	100 Mbps
Token Ring	4 bit	32 Bytes	16K Bytes	16 Mbps
BRI	0	24 Bytes	1500 Bytes	128 Kbps
PRI	0	24 Bytes	1500 Bytes	1.472 Mbps
T1	0	14 Bytes	4500 Bytes	1.5 Mbps
ATM	0	30 Bytes (AAL5)	16K Bytes (AAL5)	155 Mbps

Theoretical Values

Bandwidth ÷ Packet Size = Theoretical Performance

Smaller Packets (Less Efficient, Not Real)

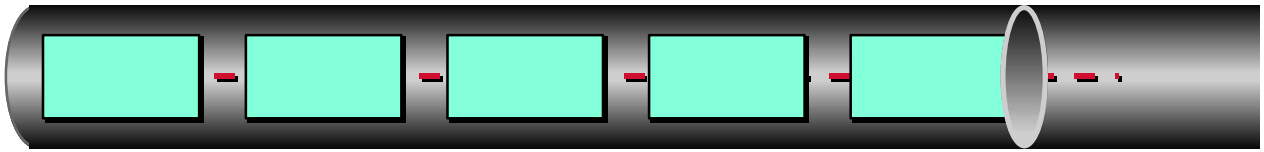

Bigger Packets (Better Utilization)

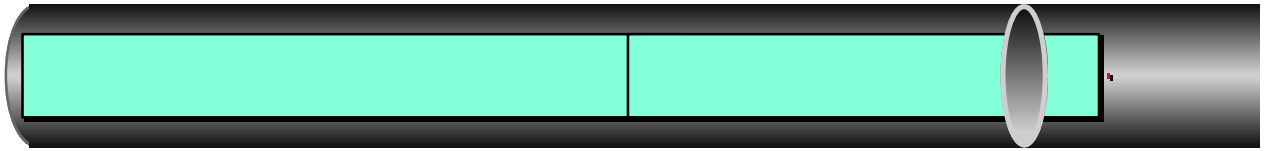

Cisco Systems

Ethernet Maximum Theoretical PPS

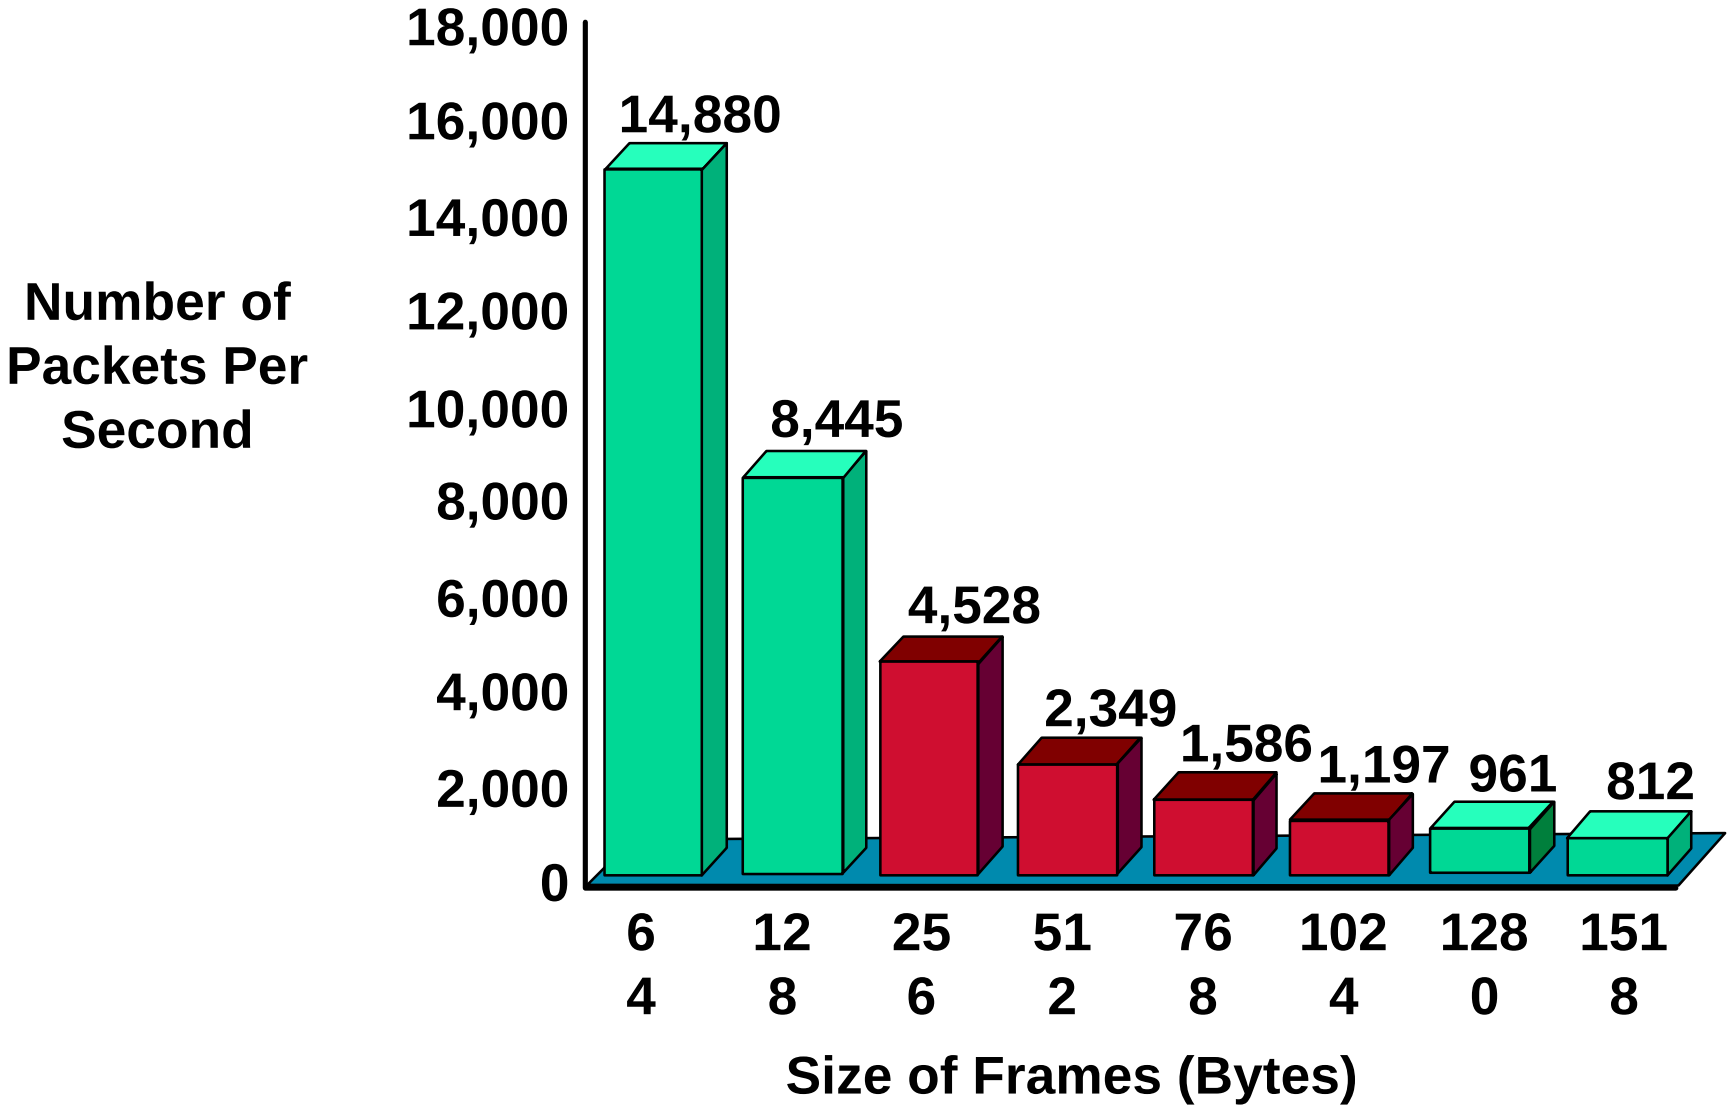

Analysis of Real Network Traffic

Ethernet Load = 40% (Real Heavy)
30% of Traffic =64-Byte Frames
40% of Traffic =512-Byte Frames
30% of Traffic =1,280-Byte Frames

(BW x %) ÷ Frame Size (bits) = PPS
2,343 PPS 390 PPS 117 PPS
= 2,850 PPS

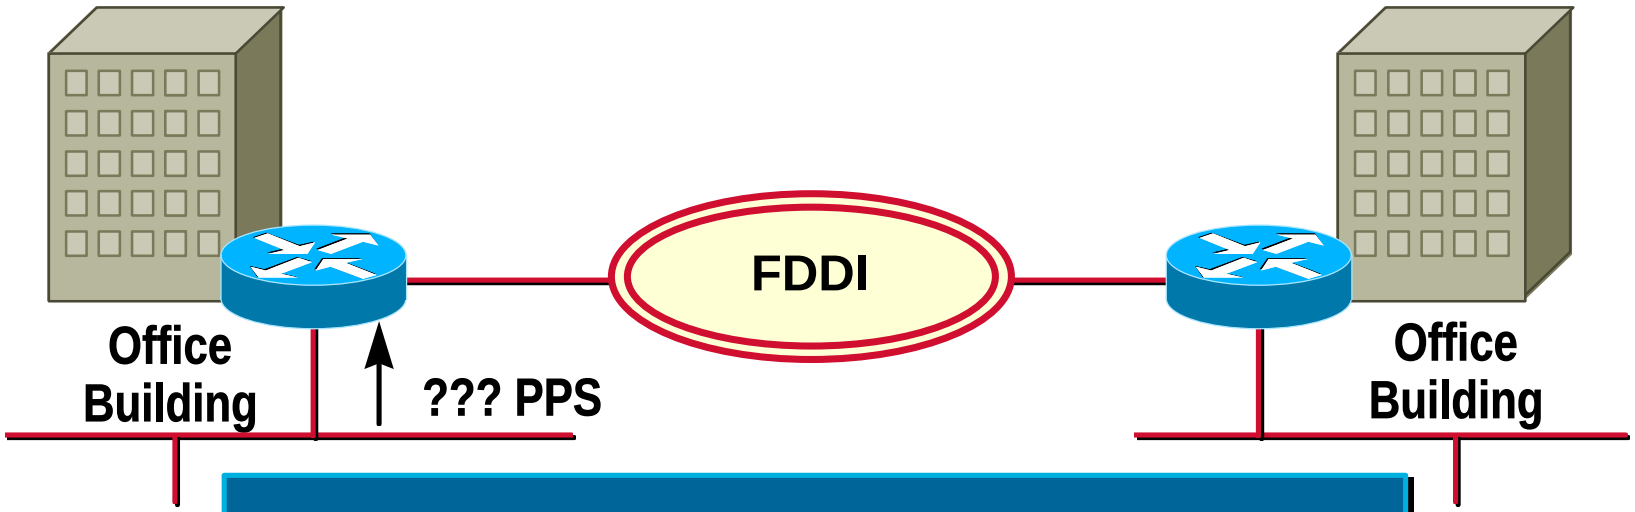

80/20 Rule —> 2,850 X 20% = 570 PPS

30/70 Rule —> 2,850 X 70% = 1,995 PPS

Cisco Systems

Agenda

- Perception versus Reality
- Layered Switching
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

Definition

- **Switching—process of transferring data from an input interface to an output interface**

What Is a Router?

- Routing is an overhead activity
- All performance is derived from the Layer 3 switching mechanism

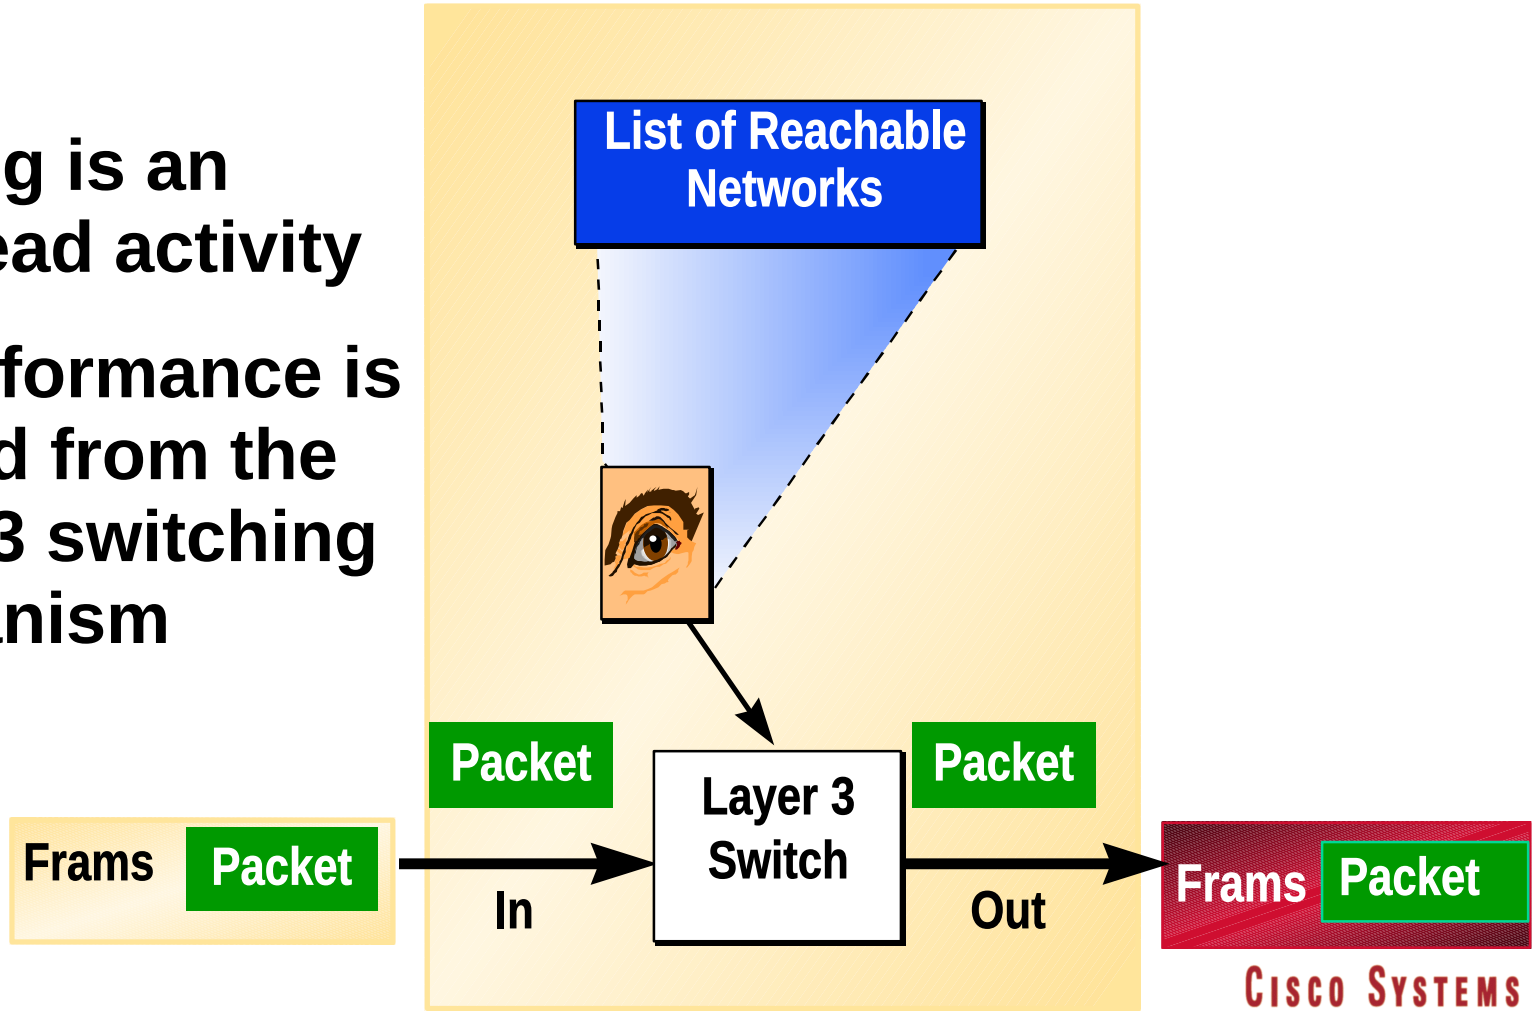

Cisco Systems

Agenda

- Perception versus Reality
- Layered Switching
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

Platform Architectures and Switching Paths

7500 Series

7000 Series

4000

4700

2500 Series

4500

Process

Autonomous

Fast

NetFlow™

Optimum

Silicon

CISCO SYSTEMS

Cisco Low-End/MidRange Architecture

Low-End/Mid-Range Architecture

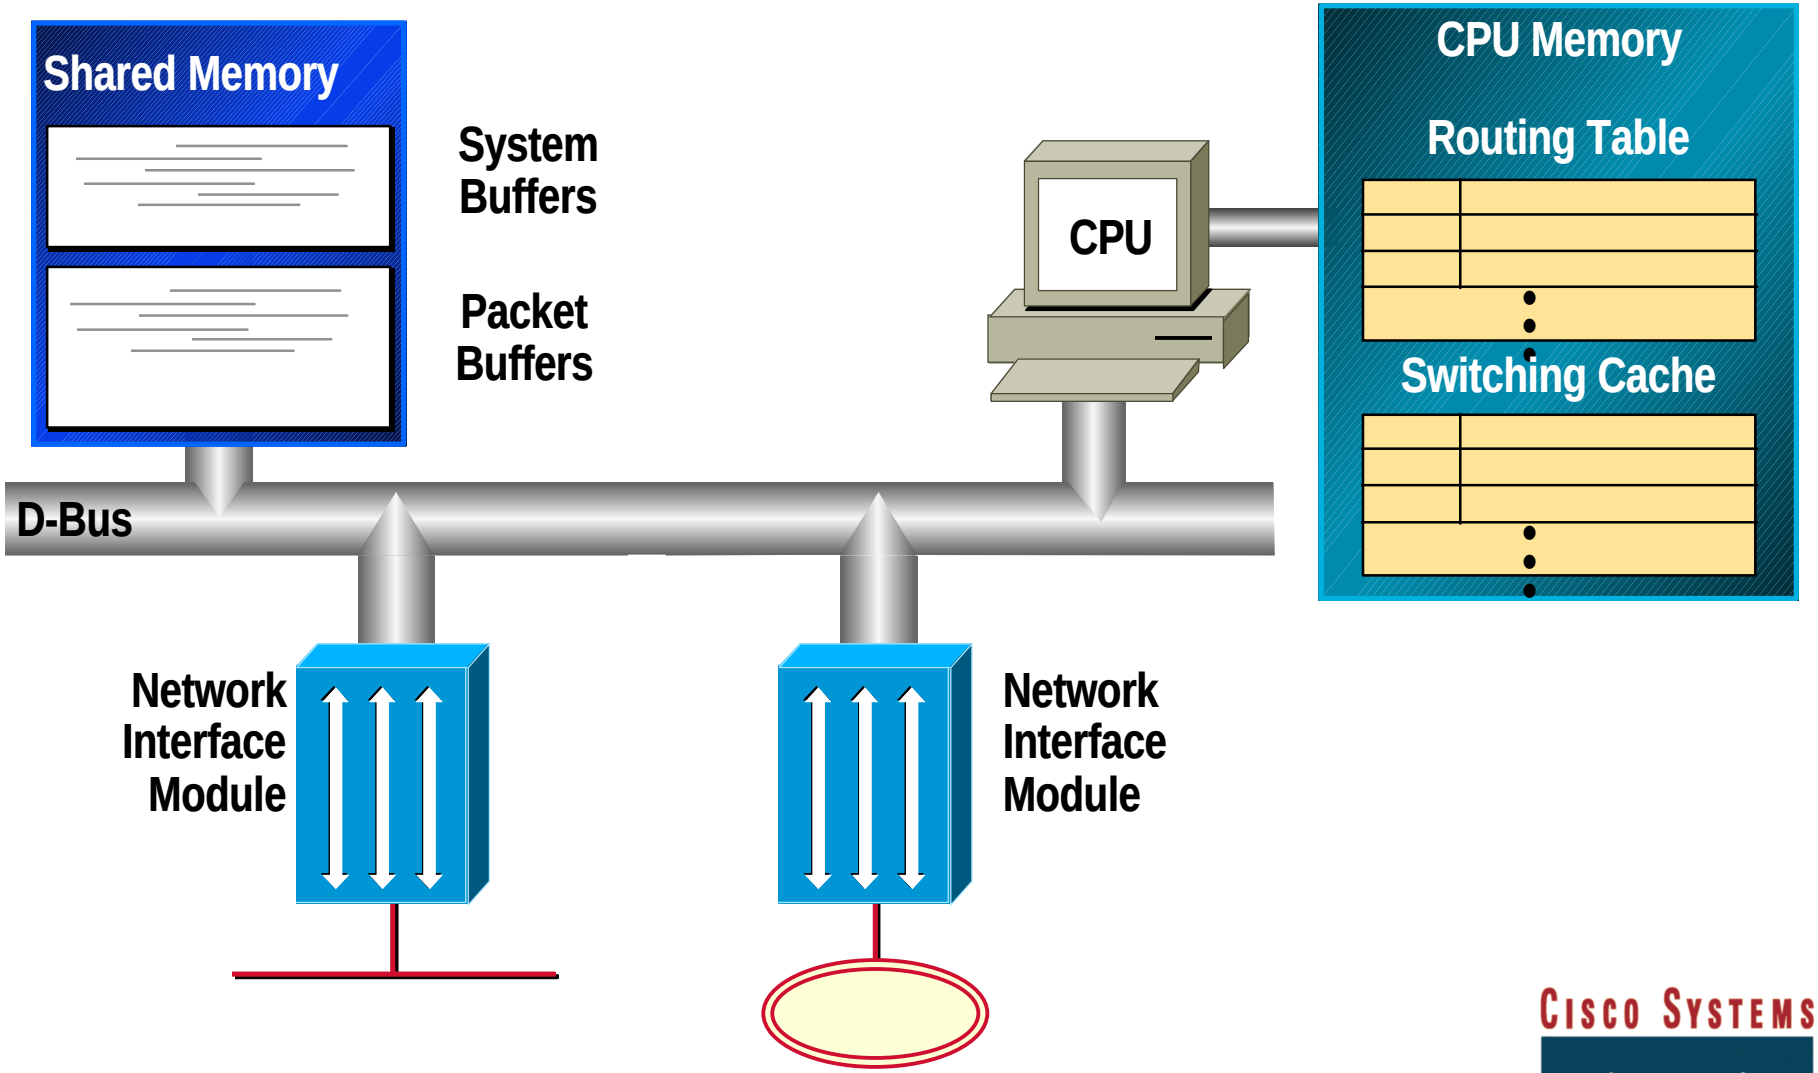

Cisco 2500 Series Hardware Configuration

	Cisco 2500	Used For:
Processor Type	20-MHz 68040	Main System Processor
System Flash Flash EPROM	4 -› 8 MB	System Software Image
Boot Flash	N/A	N/A
Shared DRAM	1 -› 2 MB	Incoming/Outgoing Packets
Main DRAM	1 -› 16 MB	Stores Operating Configuration, Routing Tables and Caches for All Protocols
NVRAM	128 KB	Configuration File

Cisco 3600 Series Hardware Configuration

	Cisco 3620	Cisco 3640	Used For:
Processor Type	80 MHz RISC-IDT4600	133 MHz RISC-IDT4700	Main Processor
System Flash Flash SIMM	4 -› 32 MB	4 –› 32 MB	System Software Image
Boot Flash (PCMCIA)	4 -> 32 MB	4 –› 32MB	System Software Image
DRAM	4 -› 64 MB	4 –› 128 MB	Data and Packet Memory
Boot ROM	1MB	1MB	ROM Monitor
NVRAM	32 KB	128 KB	Configuration

Cisco 4000 Series Hardware Configuration

	Cisco 4000-M	Cisco 4500-M	Cisco 4700
Processor Type	40 MHz Motorola 68030	100 MHz IDT Orion RLSC	133 MHz IDT Orion RLSC
System Flash Flash EPROM	4 -› 16 MB	4 –› 8 or 16 MB	4 -› 8 or 16 MB
Boot Flash	N/A	4 –› 8 or 16 MB	4 -› 8 or 16 MB
Shared DRAM	4 -› 16 MB	4 –› 8 or 16 MB	4 -› 8 or 16 MB
Main DRAM	8 -› 16 or 32 M	8 –› 16 or 32 MB	16 -› 32 MB
Secondary Memory Cache	—	—	512 KB
NVRAM	128 KB	128 KB	128 KB

Cisco Low-End/ Mid-Range Switching Paths

Low-End/Mid-Range Process Switching

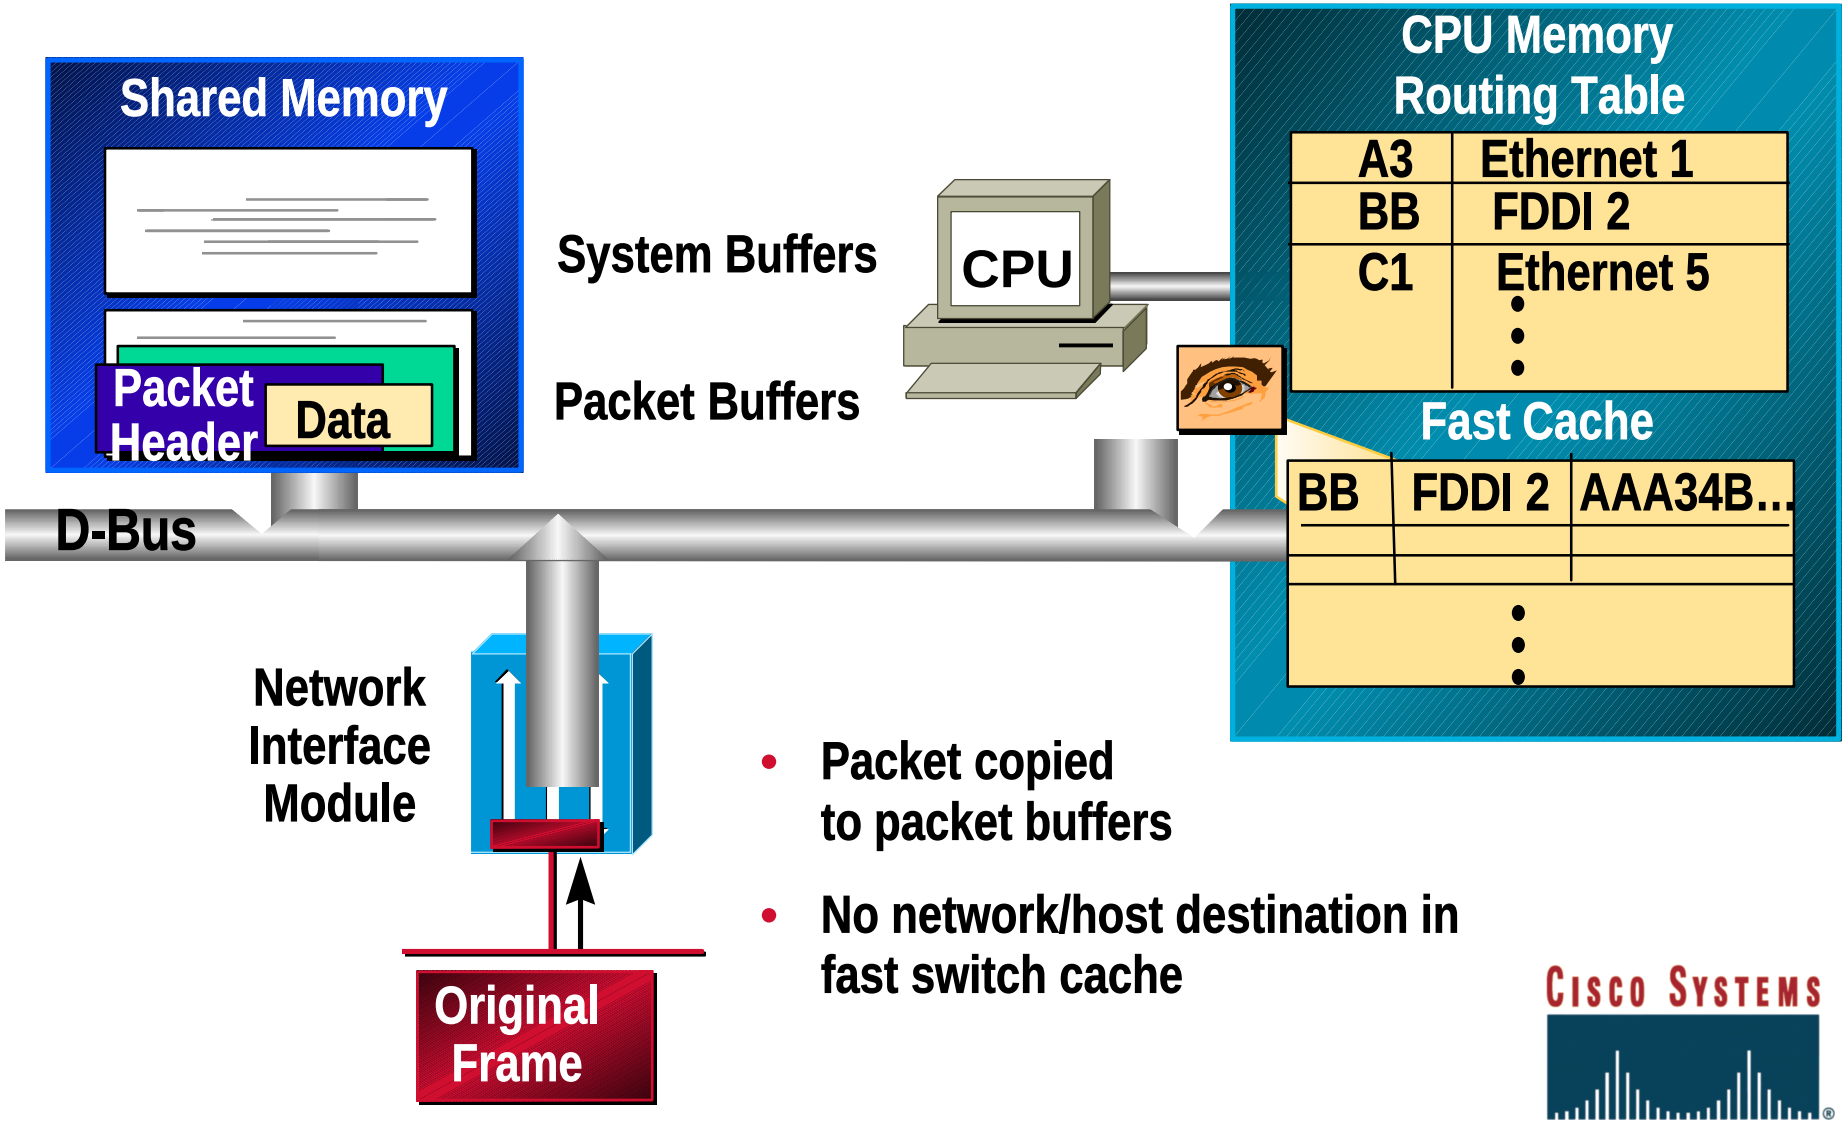

- Packet copied to packet buffers
- No network/host destination in fast switch cache

Cisco Systems

Low-End/Mid-Range Process Switching

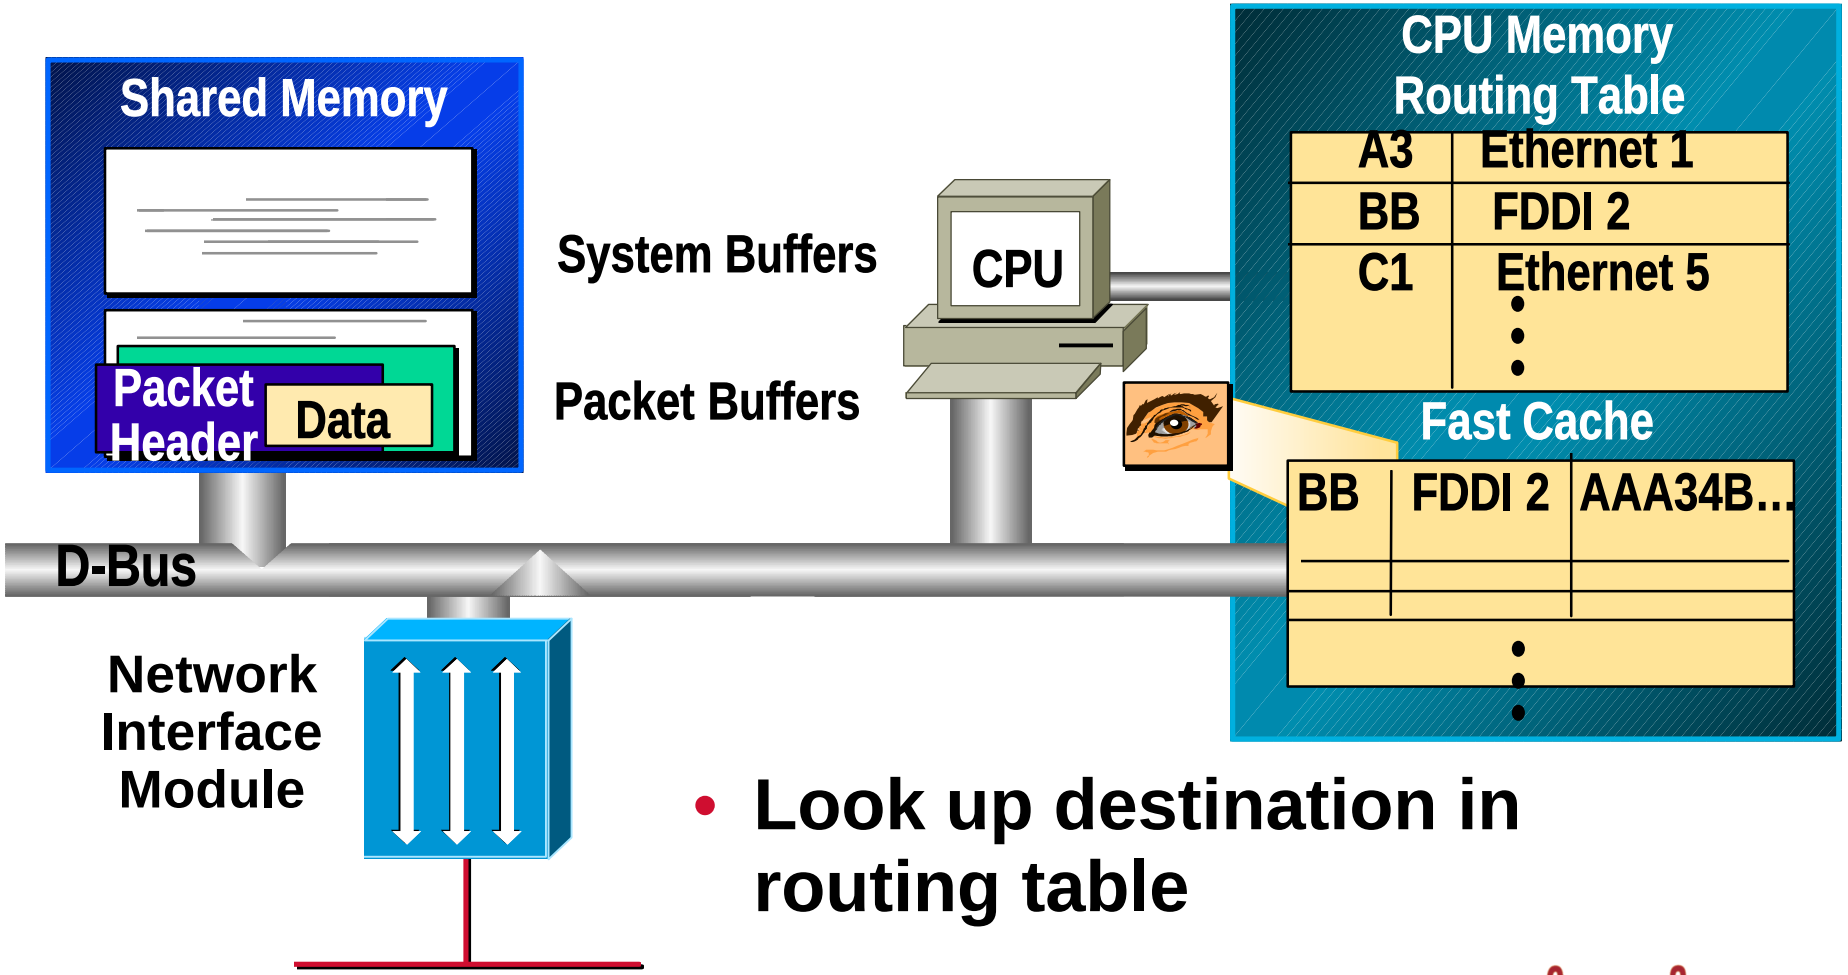

- Look up destination in routing table
- Initialize fast cache

Cisco Systems

Low-End Initializing Fast Switching

```
Oban# show apple cache
AppleTalk Routing Cache, * = active entry,
cache version is 195
Destination      Interface        MAC Header
* 8.100          Ethernet0    00000C35EDED00000C064146
* 7043.42        Fddi0        00000C0DFDD800000C064146
* 7070.28        Fddi0        00000C0DFDD800000C064146
* 7101.105       Fddi0        00000C0DFDD800000C064146
* 7207.1         Fddi0        00000C0DFDD800000C064146
* 7364.2         Ethernet0     00000C37140800000C064146
* 7364.22        Ethernet0     00000C37140800000C064146
* 7369.131       Ethernet11    0800097F6A8900000C064121
```

Low-End/Mid-Range Process Switching

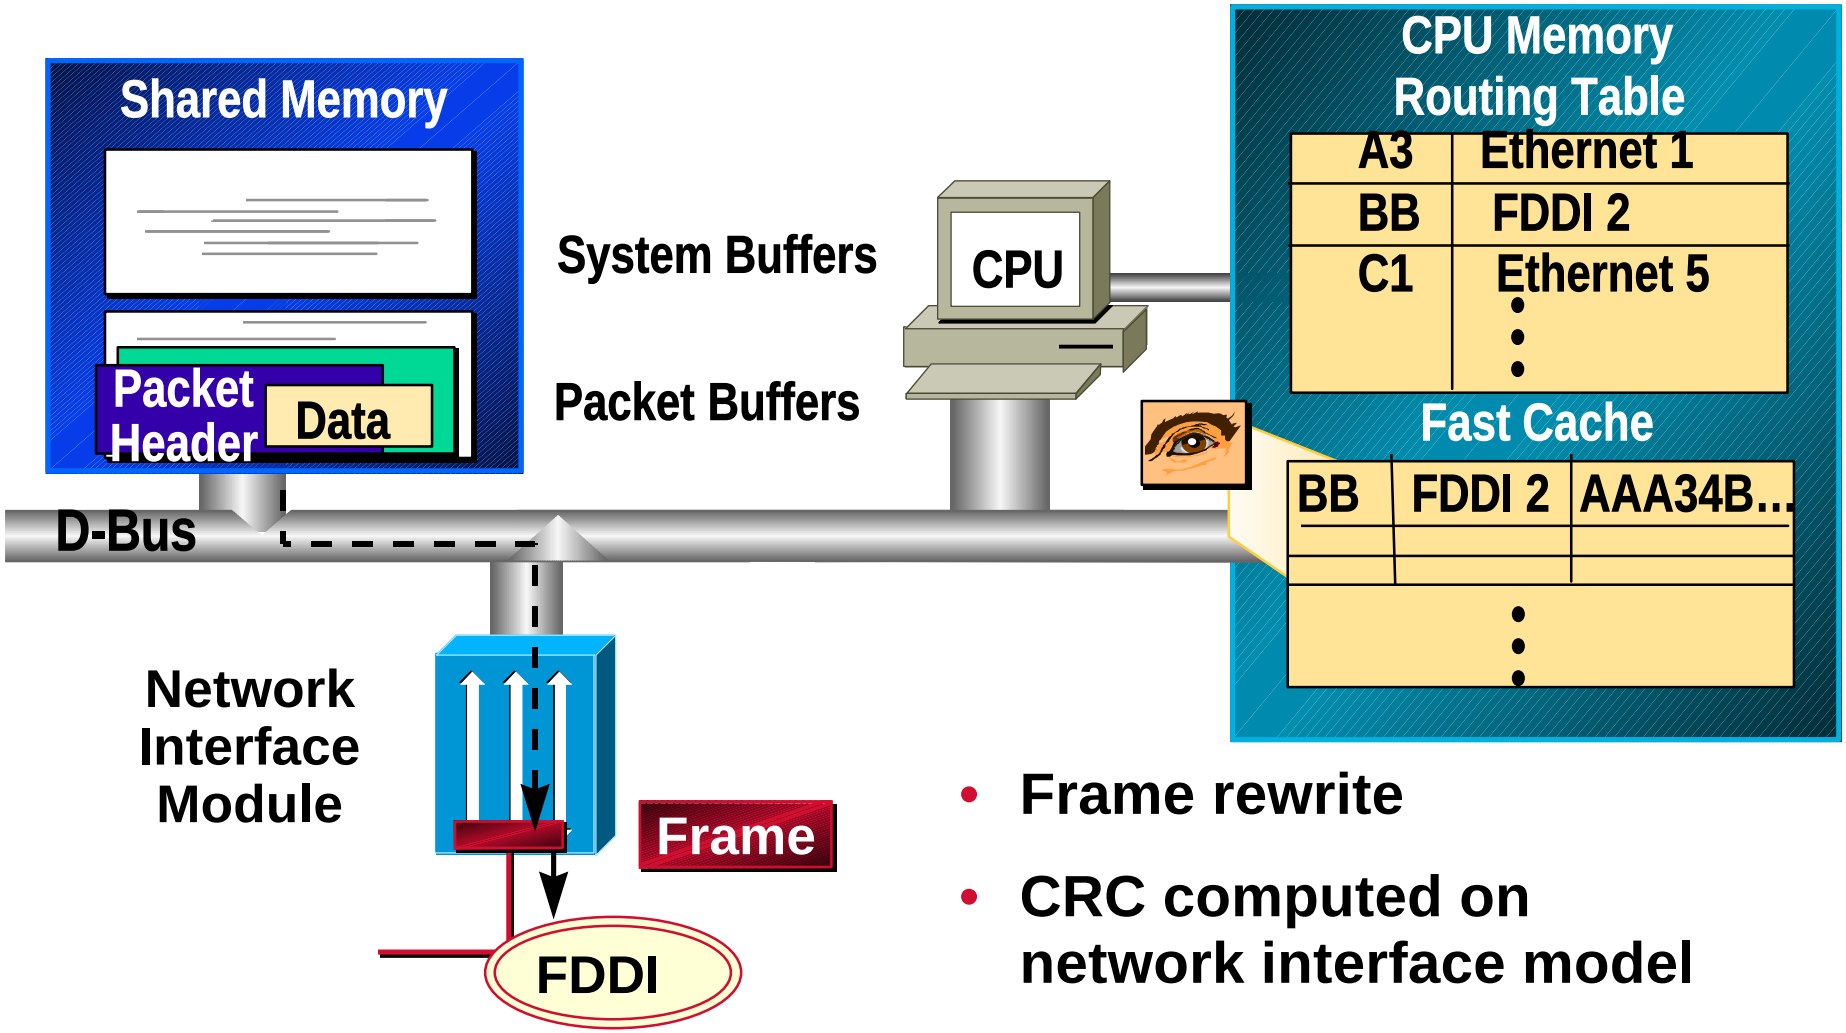

- Frame rewrite
- CRC computed on network interface model

Low-End/Mid-Range Fast Switching

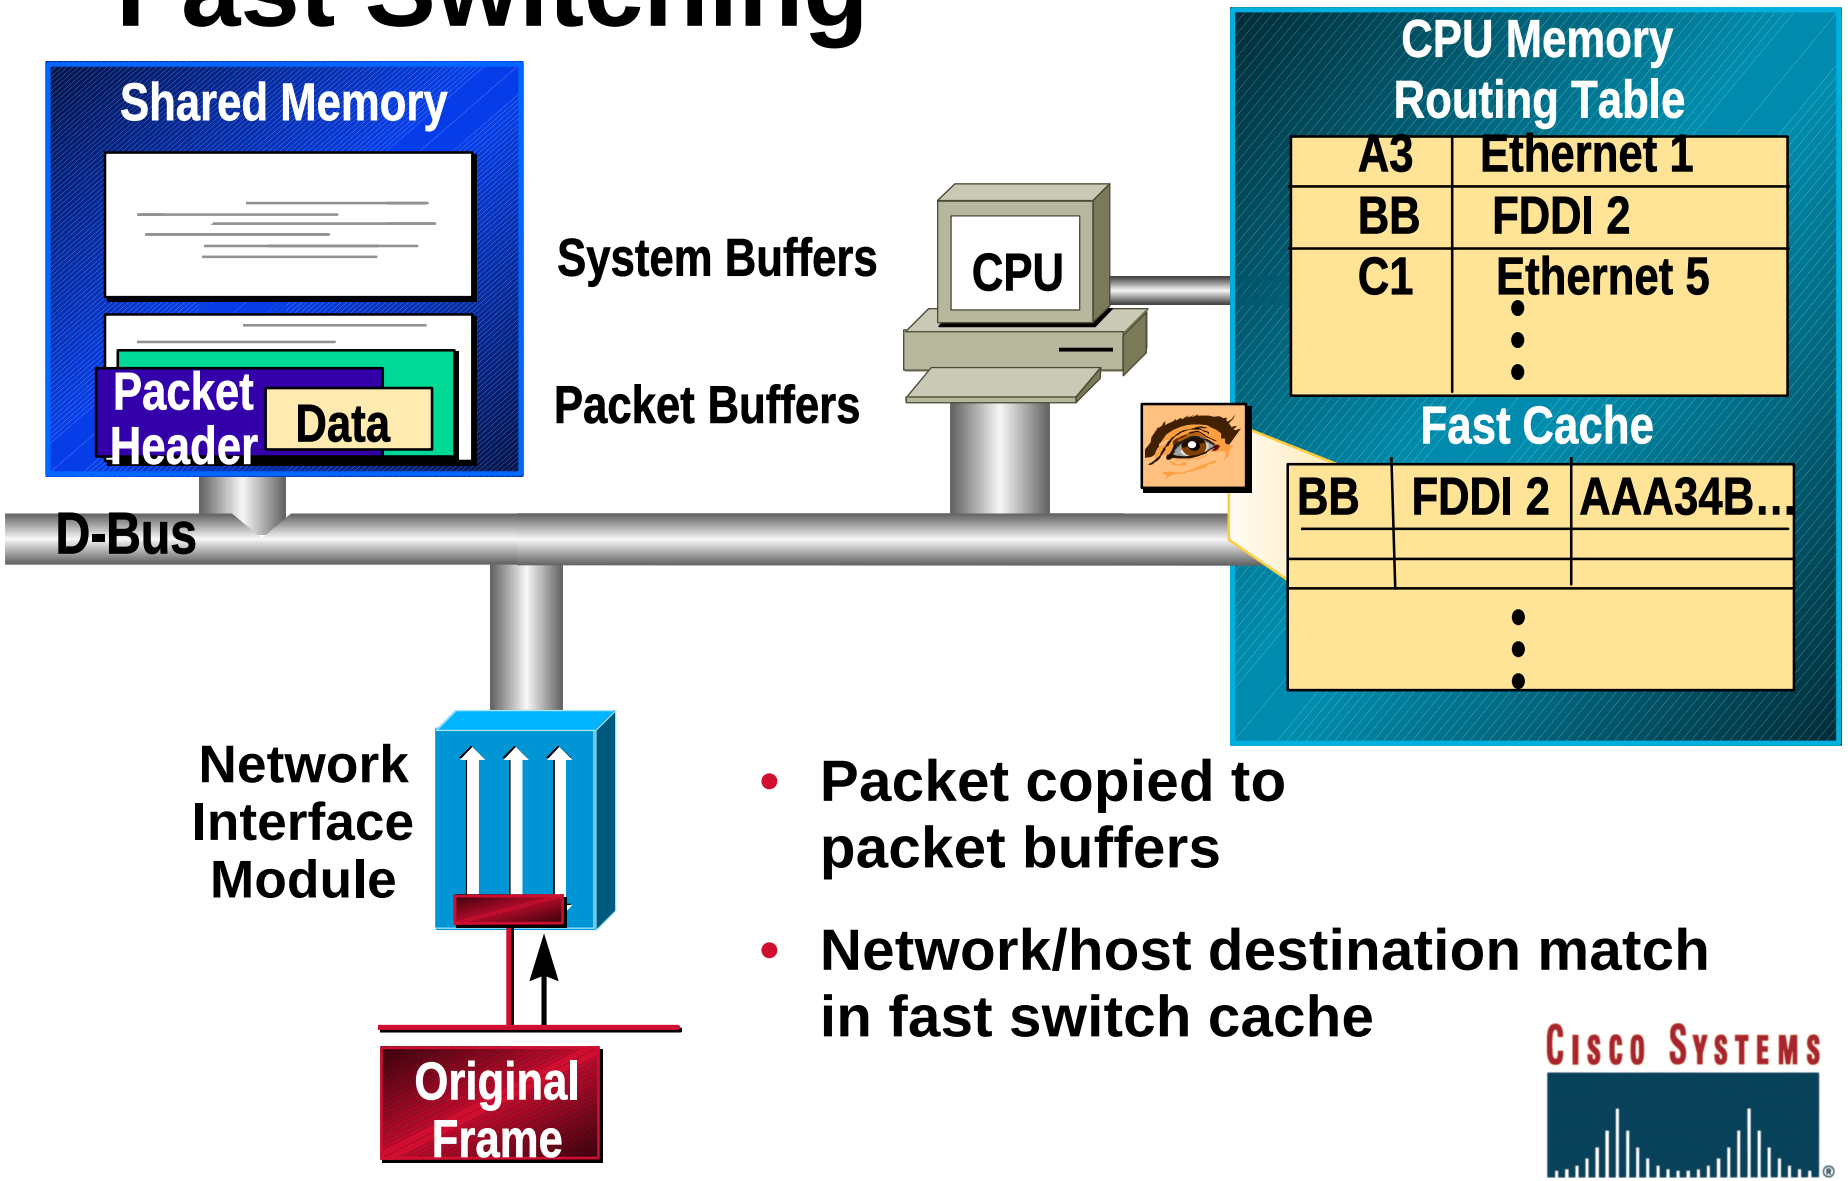

- Packet copied to packet buffers
- Network/host destination match in fast switch cache

Cisco Systems

Low-End/Mid-Range Fast Switching

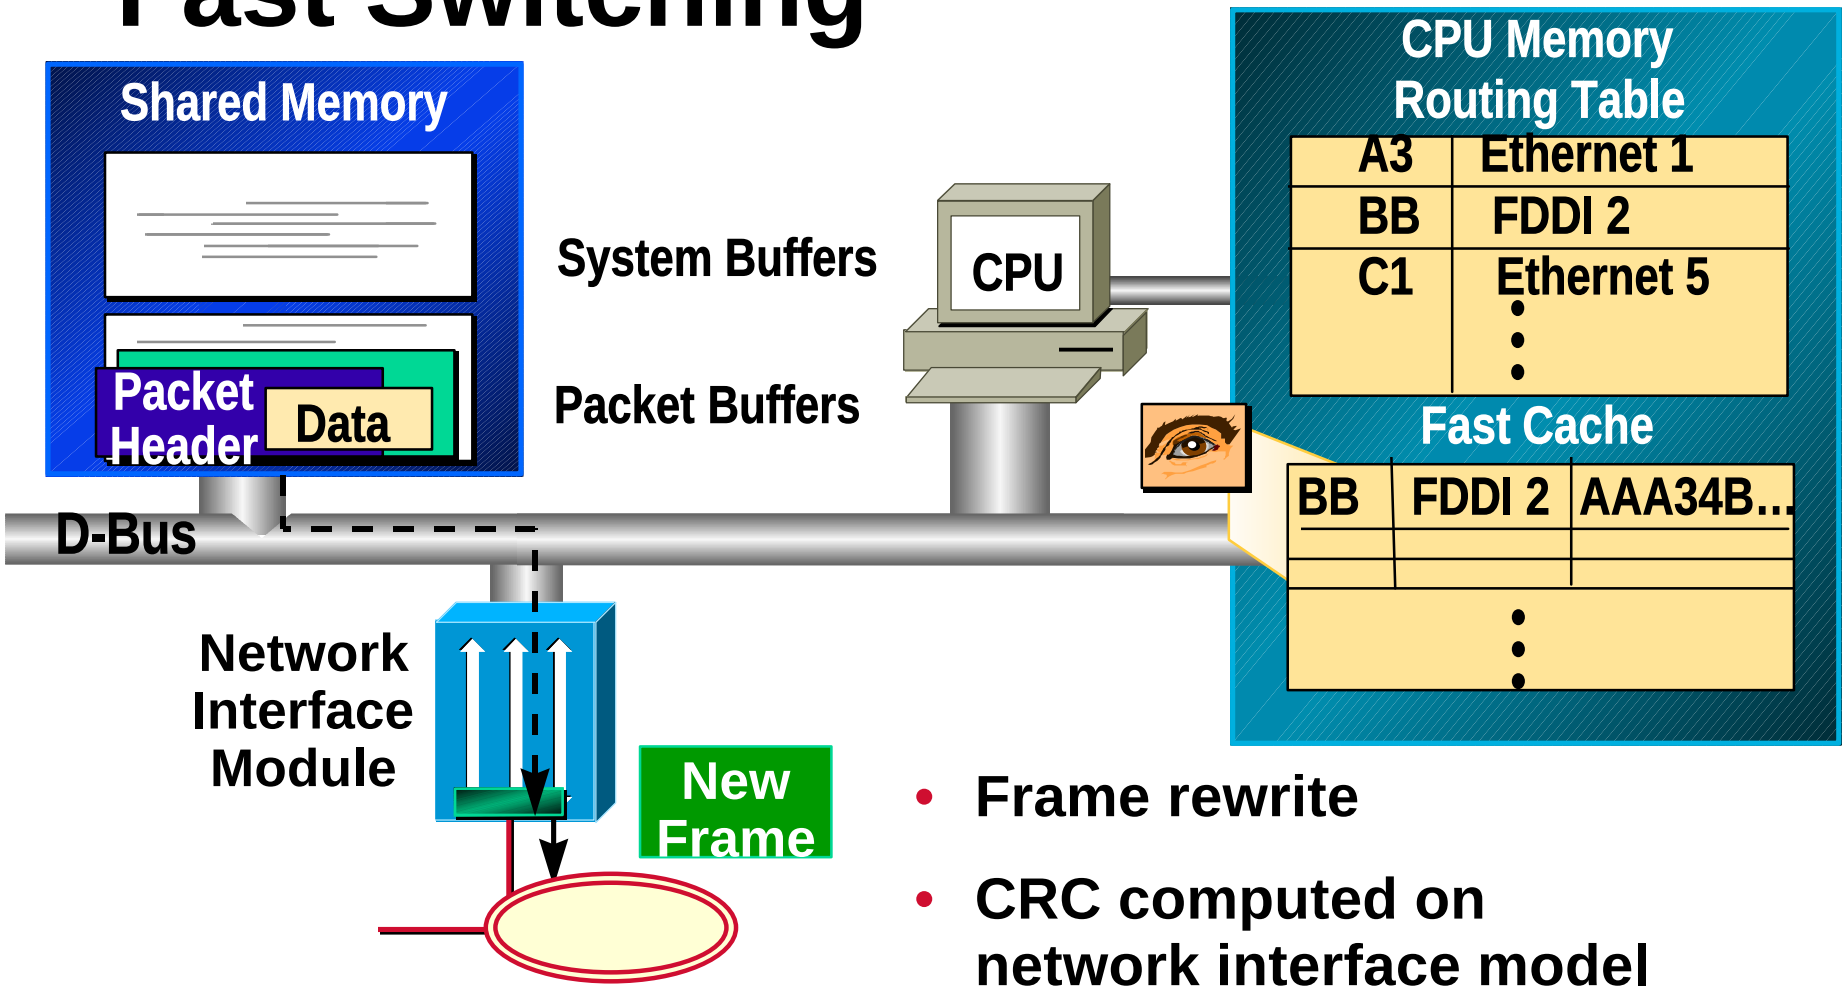

- Frame rewrite
- CRC computed on network interface model

Cisco Systems

Low-End/Mid-Range Aggregate Switching Summary

Platform	Process no <protocol> route cache	Fast <protocol> route cache
Cisco 2500 Series	1,800	6,000
Cisco 3620	2,000	16,000
Cisco 3640	4,000	40,000
Cisco 4000	1,800	14,000
Cisco 4500	10,000	45,000
Cisco 4700	11,000	50,000

Cisco 7000 Series Architecture

Cisco 7000 Architecture

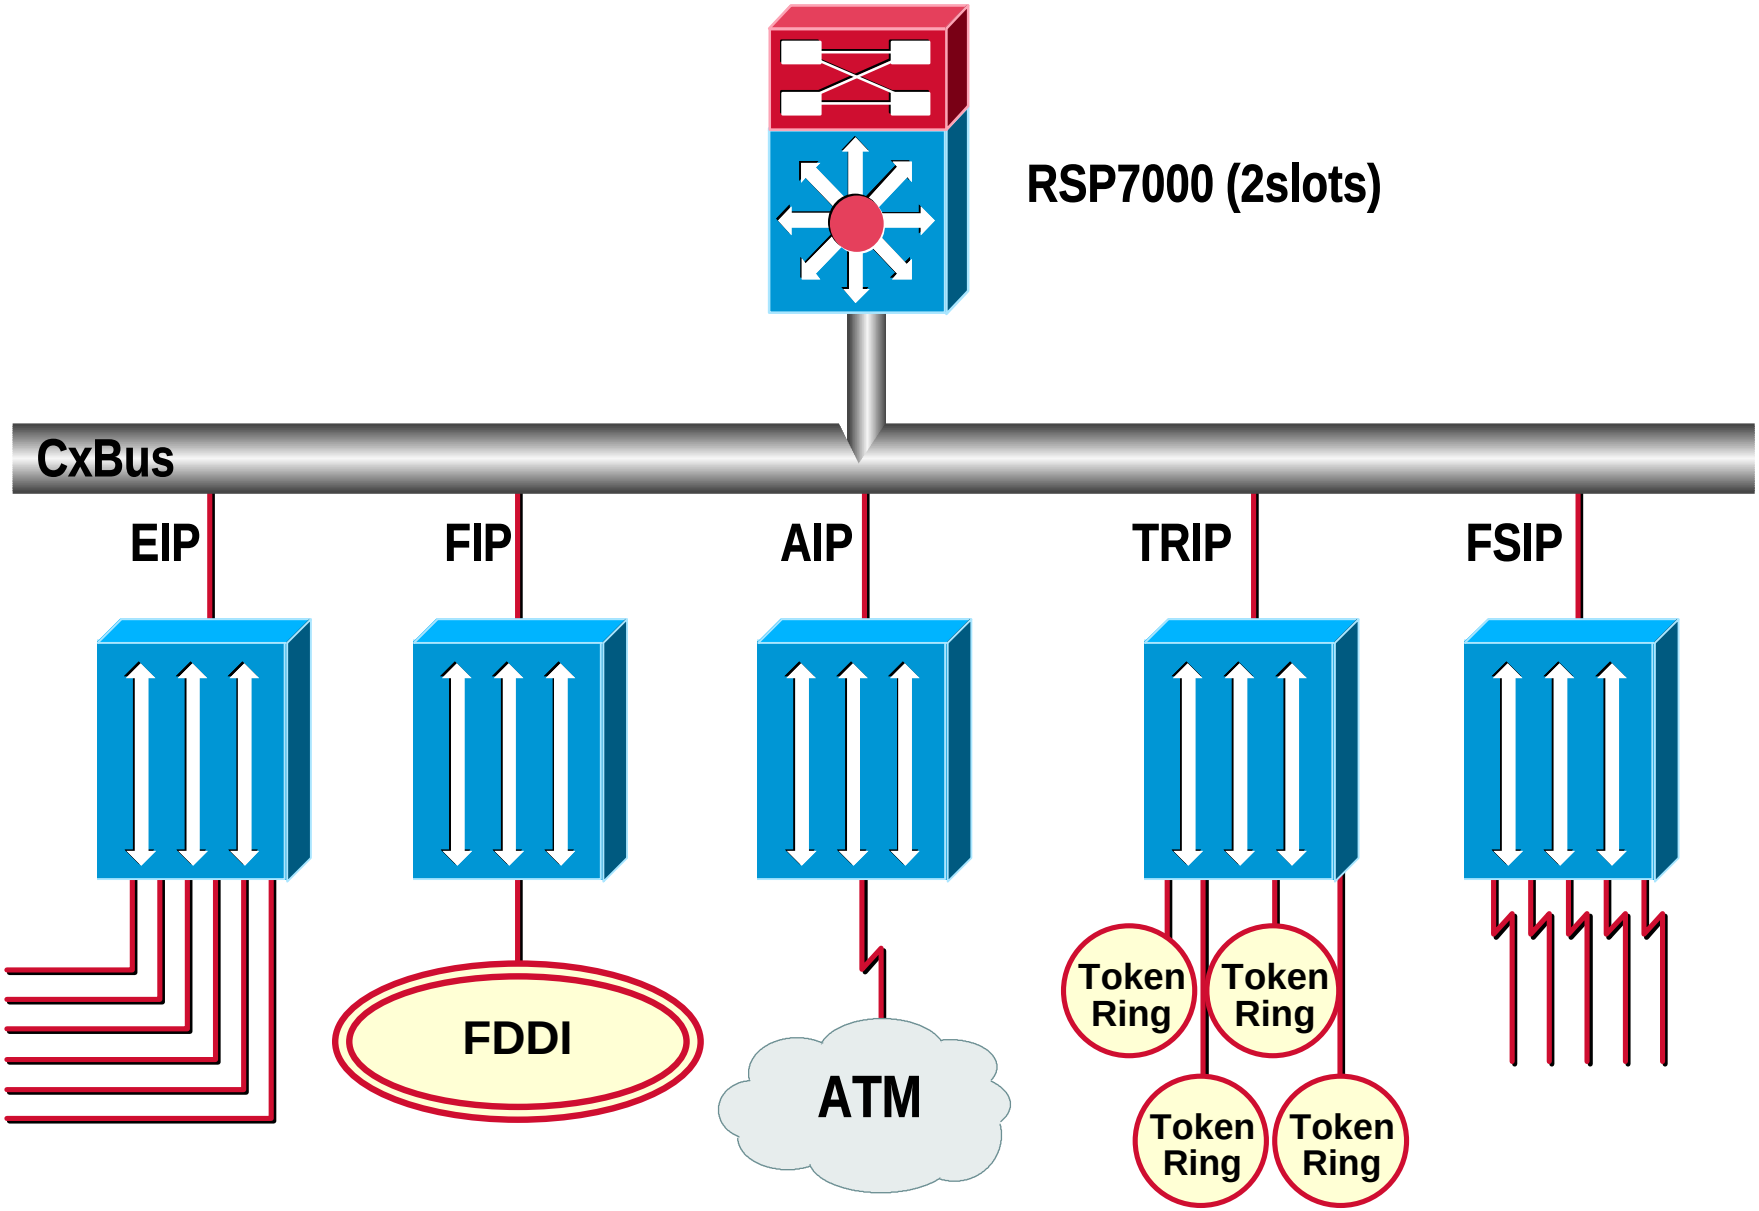

Cisco 7000 Series—Summary of Hardware Features

Characteristics	Cisco 7000	Cisco 7010
Number of Slots	7	5
Processor Type (RP)	25 MHz 68040 CPU	25 MHz 68040 CPU
System Bandwidth	533 Mbps	533 Mbps
Number of Interface Processor Slots	5	3
Power Supplies	2	1

Cisco 7000 Series—RP/SP (SSP) Memory Components

Type	Size	Description
DRAM (RP)	16–64 MB	4- or 8-Mb SIMMs
NVRAM (RP)	128 KB	Non-Volatile EPROM for System Configuration File
Flash EPROM (RP)	4–16 MB	Cisco IOS™ Images and Downloadable Micro-Code
ROM (RP)	8 MB	
Packet Memory (SP)	512 KB	Interface Buffers
Packet Memory (SSP)	512 KB–2 MB	Interface Buffers
Flash Memory Card (PCMCIA)	8–16 MB	Cisco IOS Images and Configuration Files

Cisco Systems

Cisco 7000 Series Switching Paths

Cisco 7000 Series—Switching Paths

Process Switching	→	Initializes Switching Caches
Fast Switching	→	Default for All Protocols
Autonomous Switching	→	Enable per Interface/Protocol
Silicon Switching	→	Need Silicon Switch Processor
RSP7000 Optimum	→	Need RSP7000 Processor
RSP7000 Netflow	→	Need RSP7000 Processor

Cisco 7000 Series—Process Switching

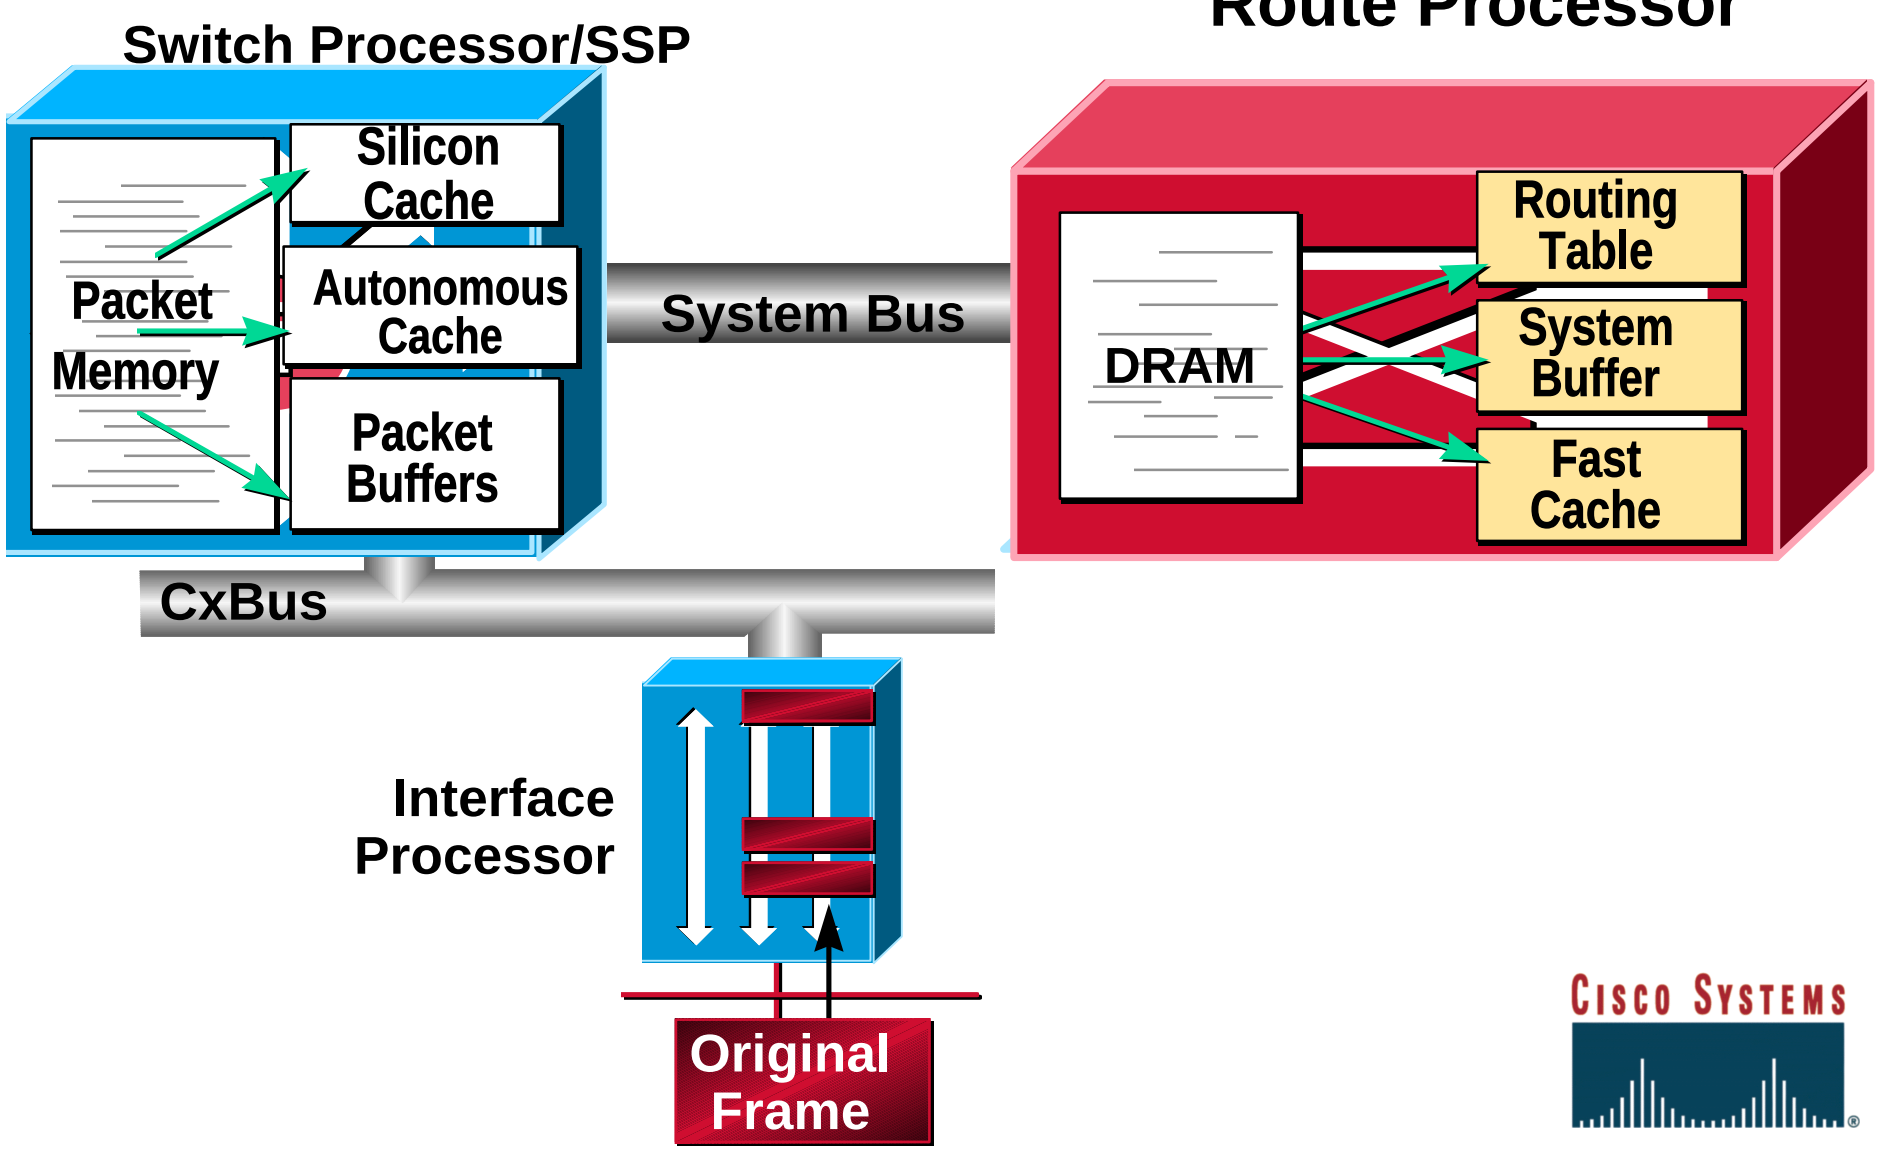

Cisco Systems

Cisco 7000 Series—Process Switching

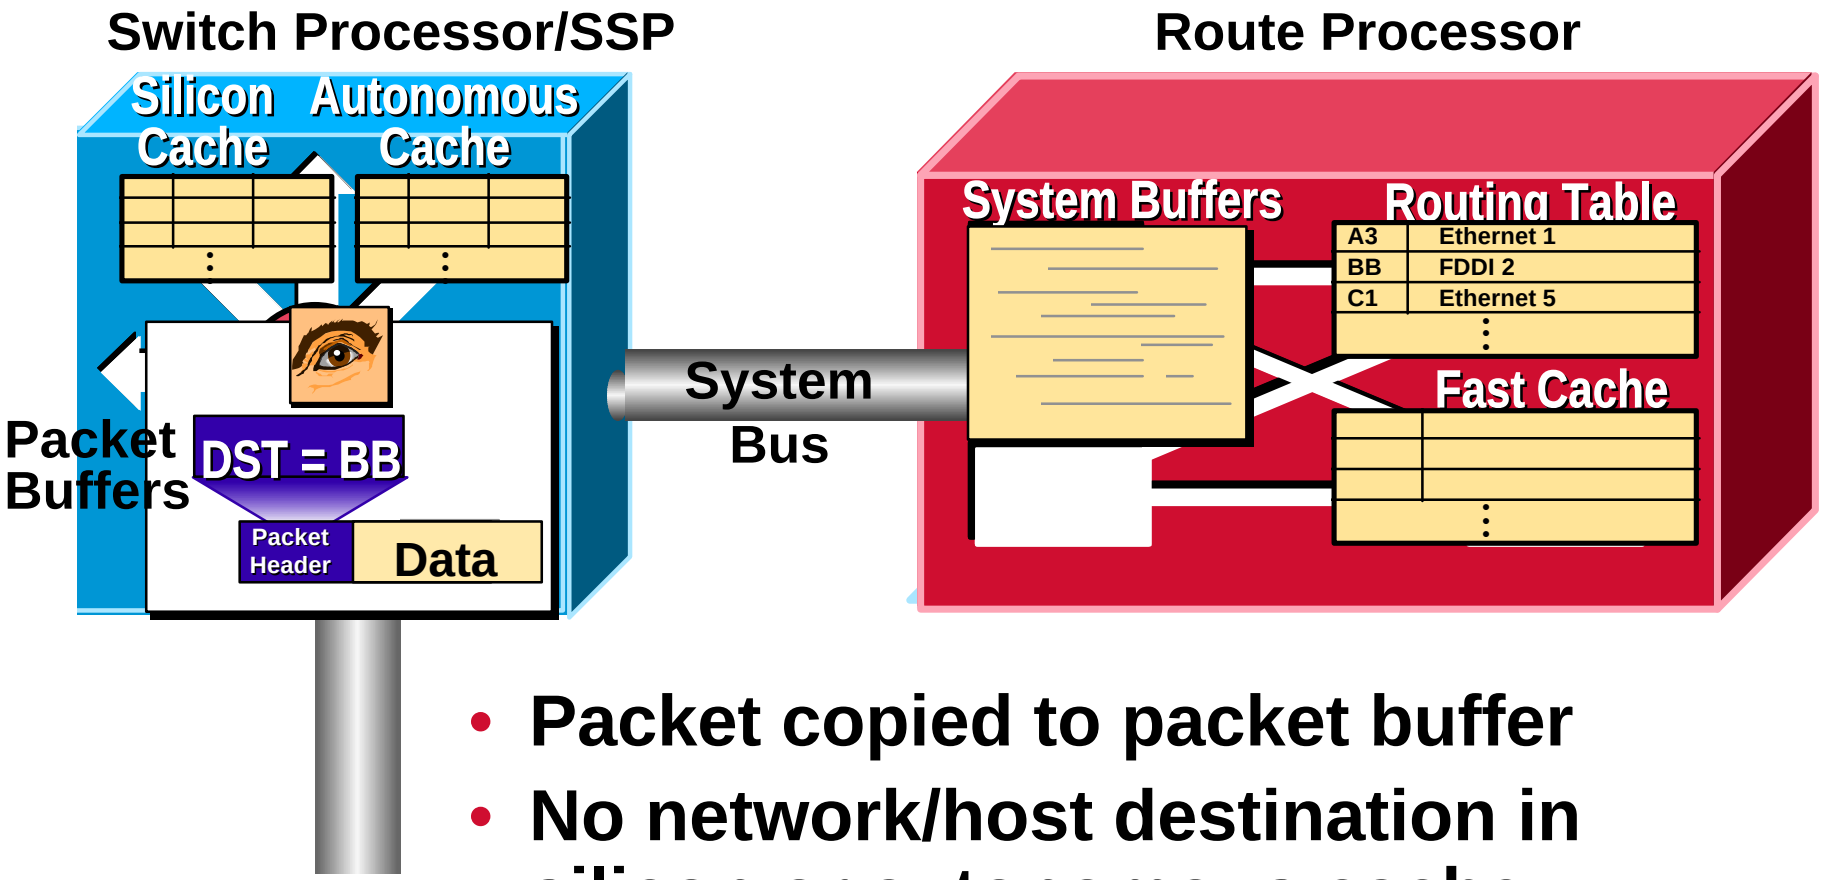

- Packet copied to packet buffer
- No network/host destination in silicon or autonomous cache

Cisco Systems

Cisco 7000 Series—Process Switching

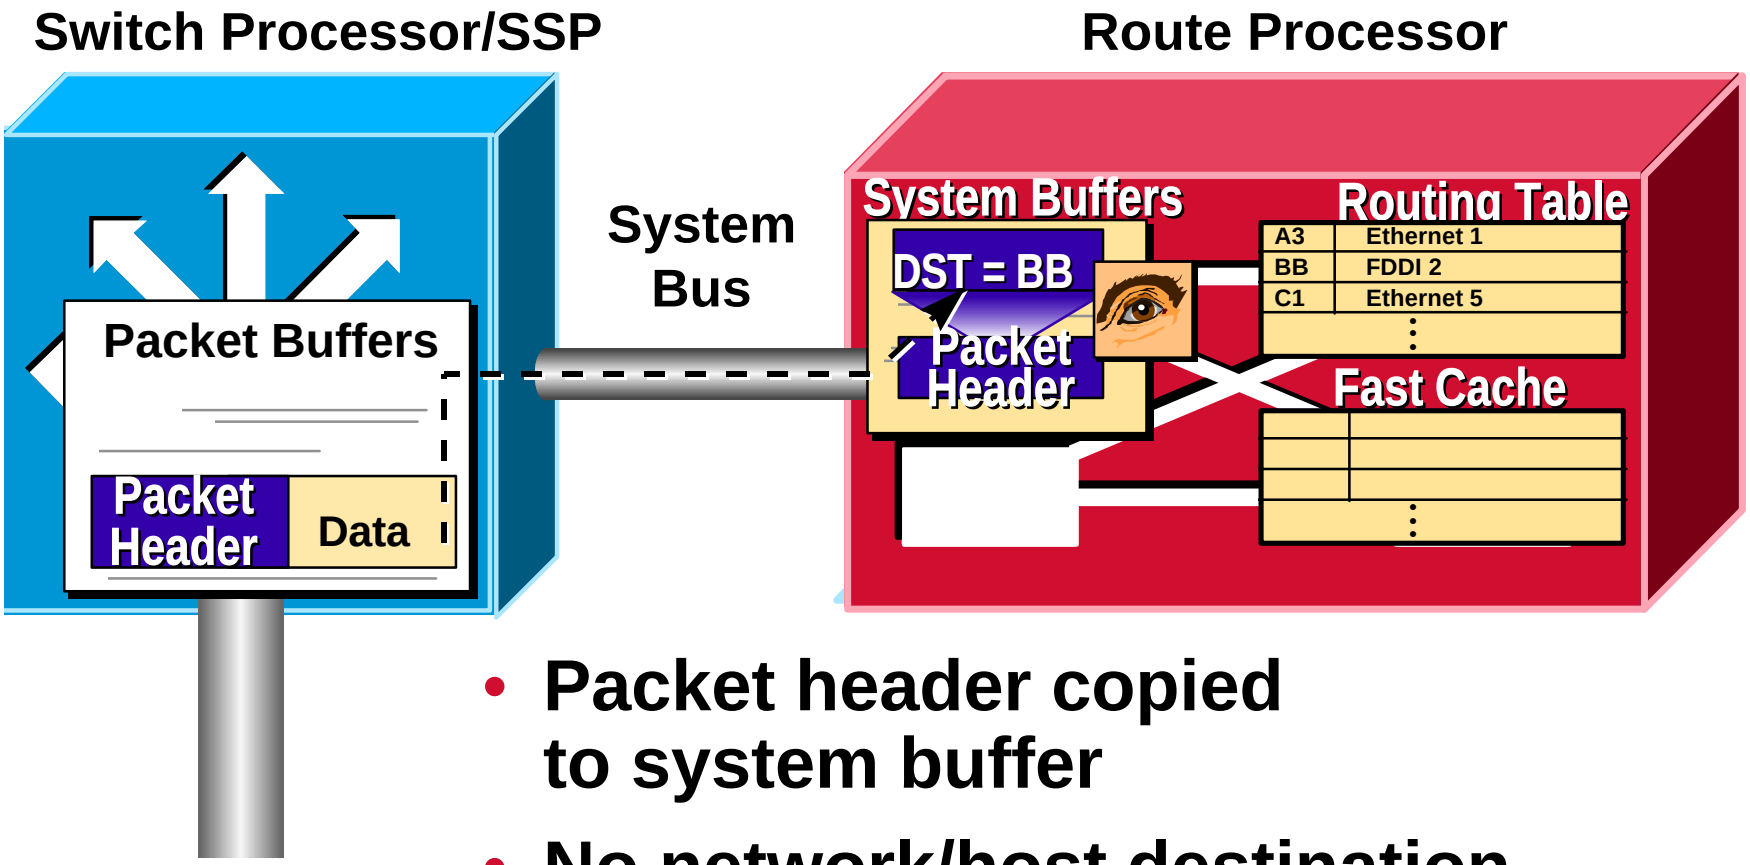

- Packet header copied to system buffer
- No network/host destination in fast switch cache

Cisco Systems

Cisco 7000 Series—Process Switching

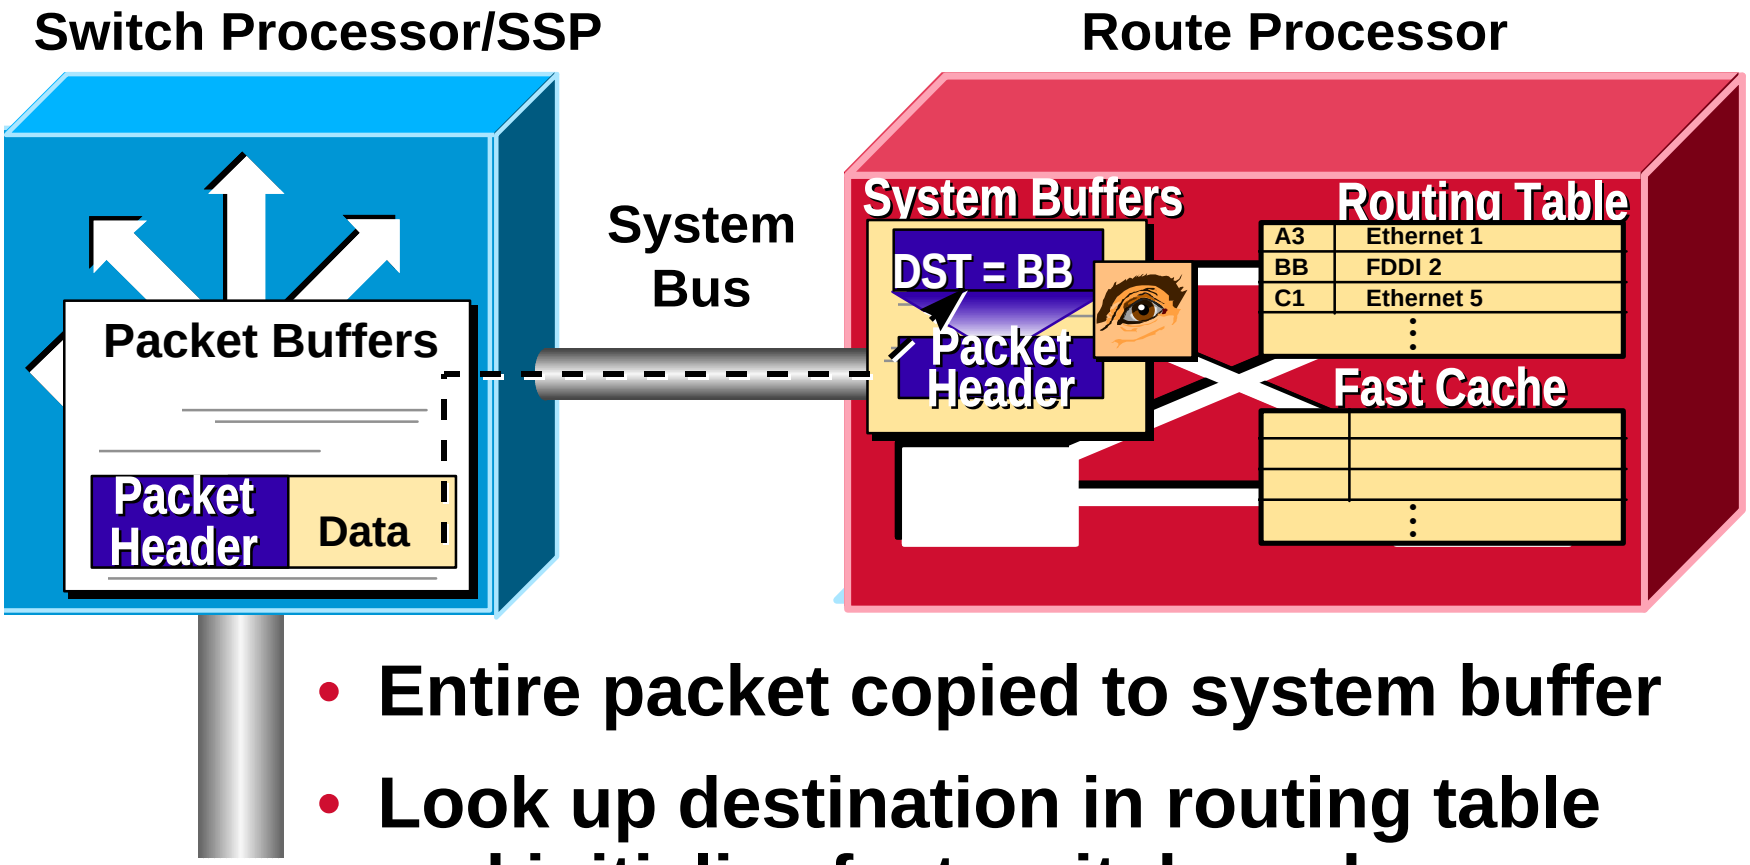

- Entire packet copied to system buffer
- Look up destination in routing table and initialize fast switch cache

Cisco Systems

Cisco 7000 Series—Initializing Autonomous Switching Cache

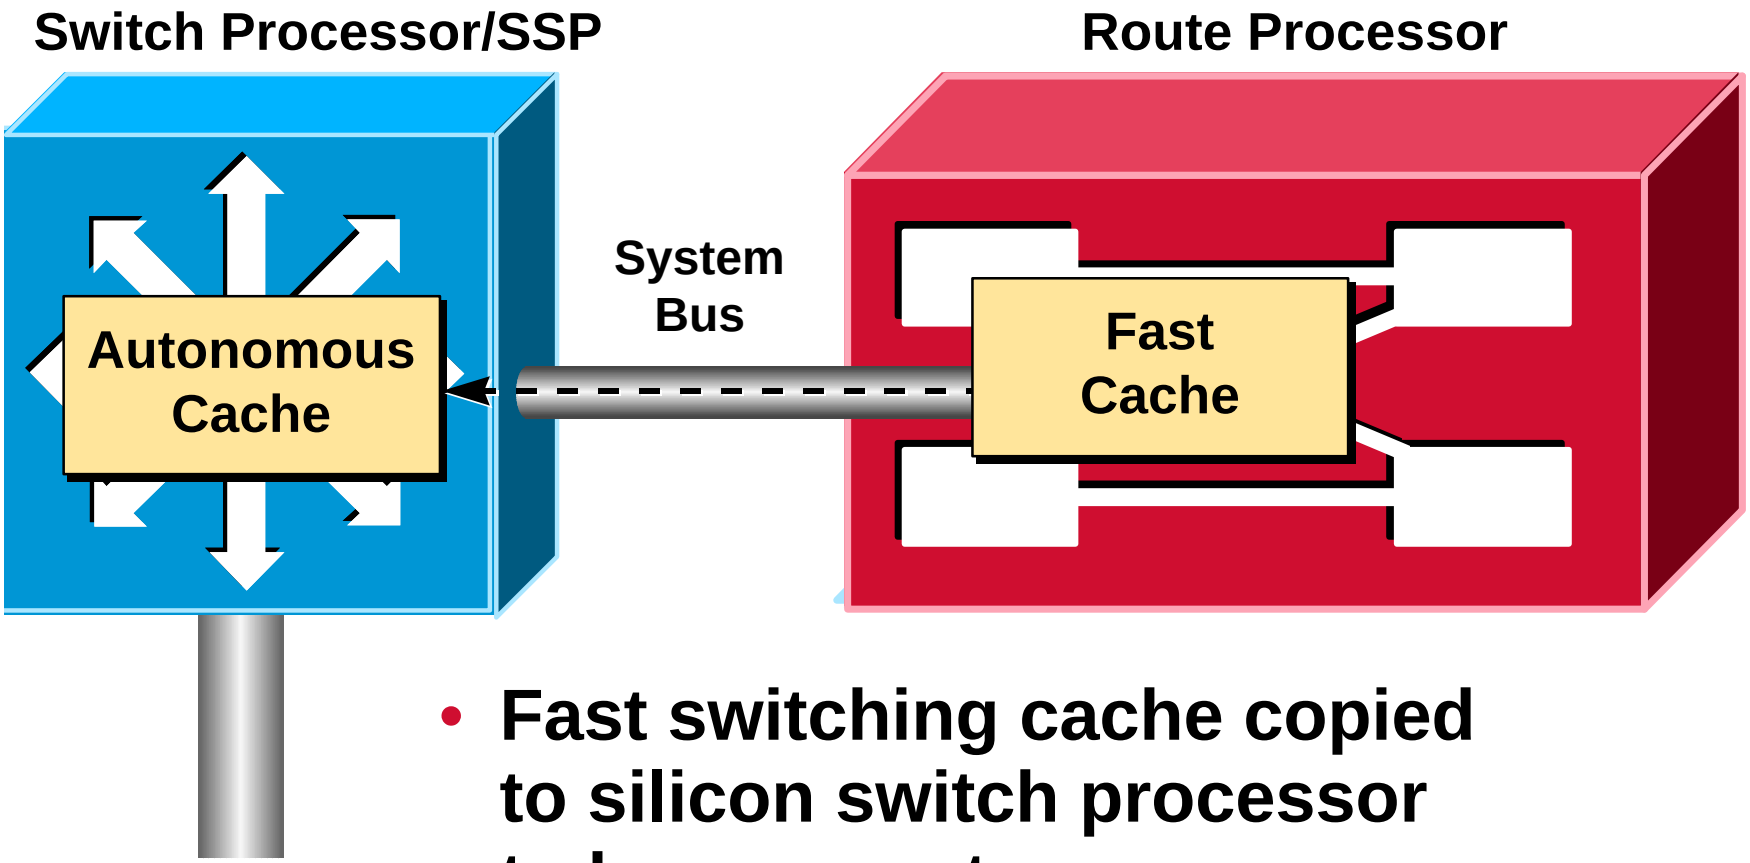

- Fast switching cache copied to silicon switch processor to become autonomous switch cache

Cisco Systems

Cisco 7000 Series—Initializing Silicon Switching Code

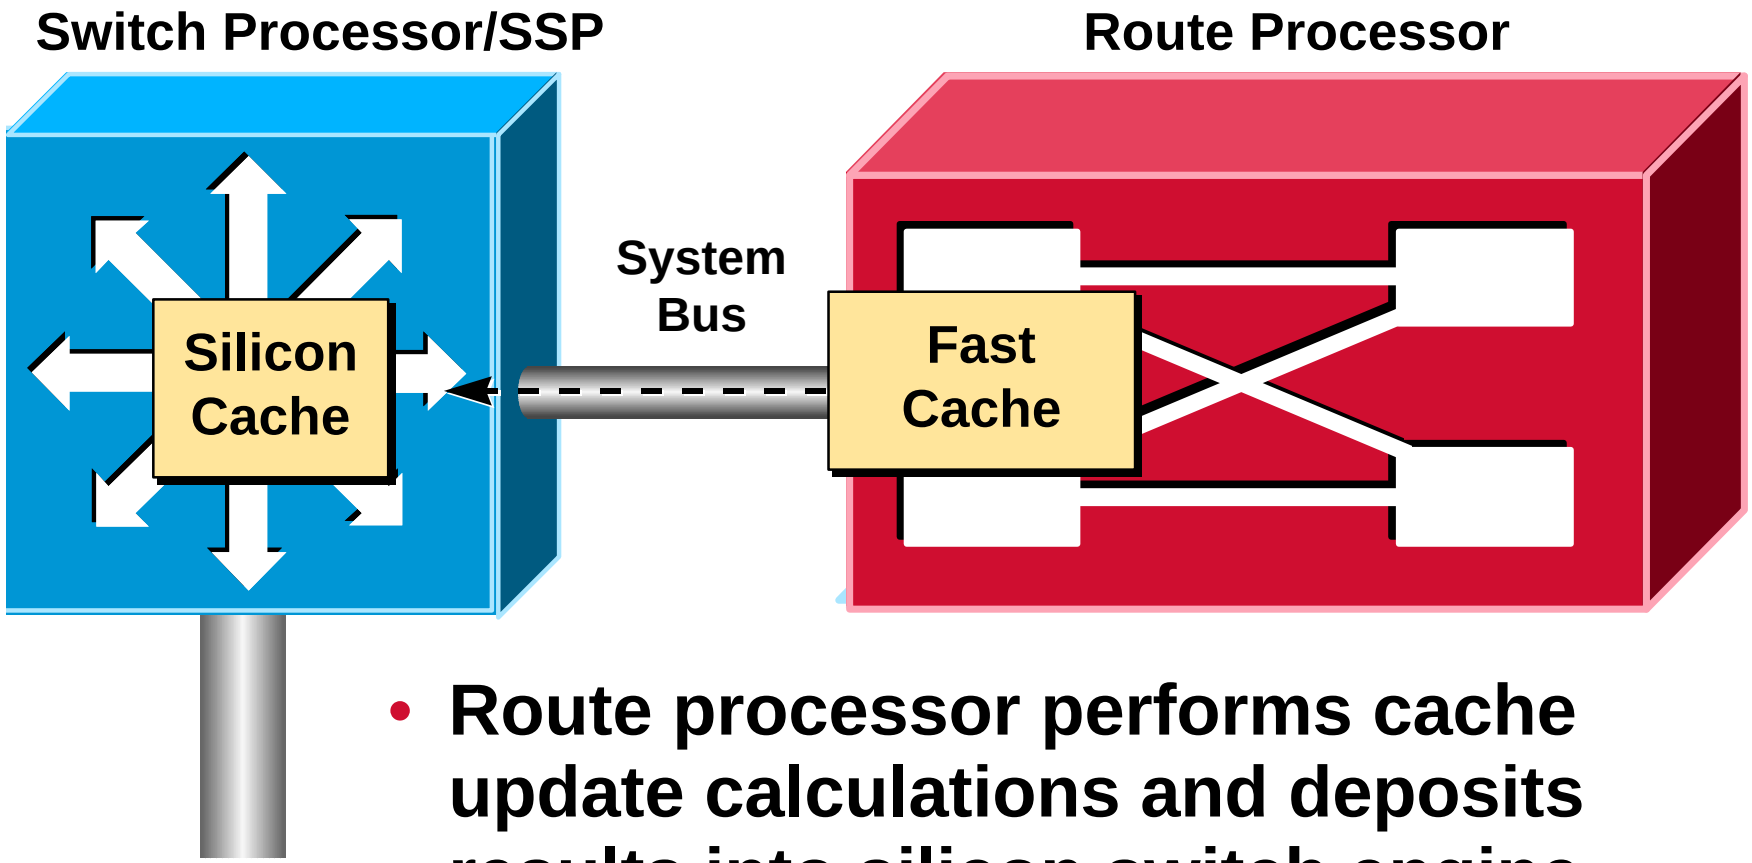

- Route processor performs cache update calculations and deposits results into silicon switch engine

Cisco 7000 Series—Process Switching

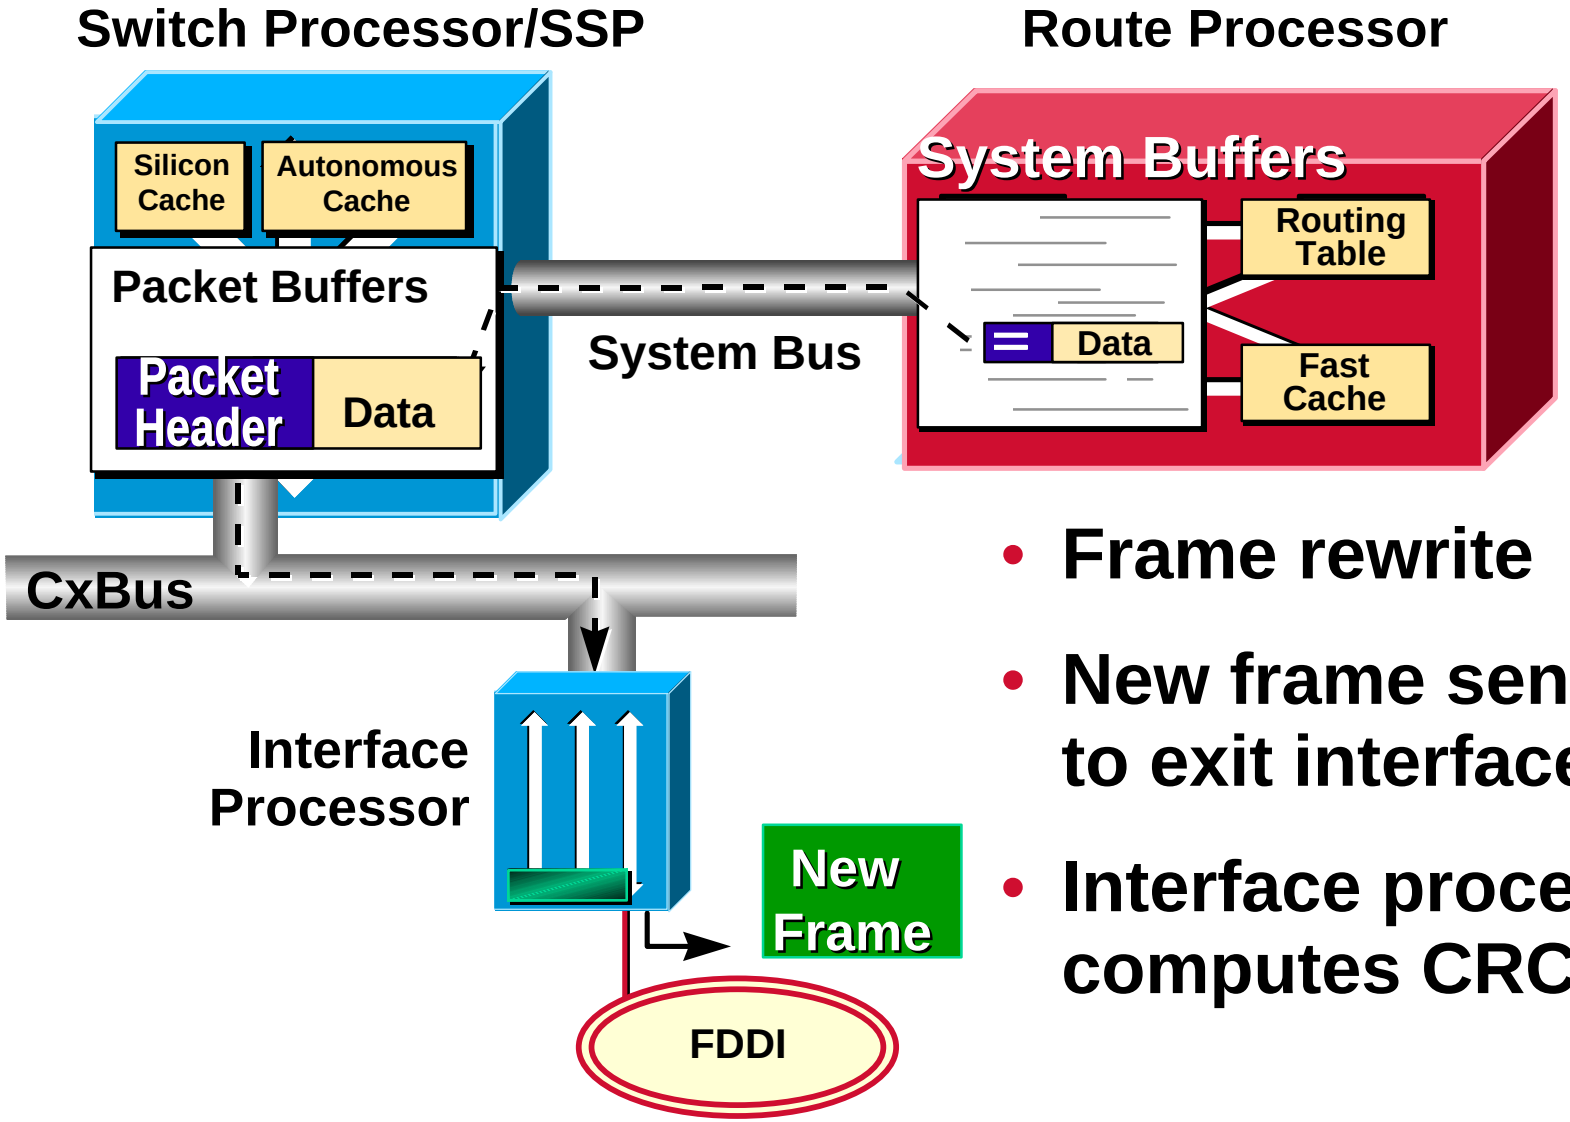

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco 7000 Series—Fast Switching

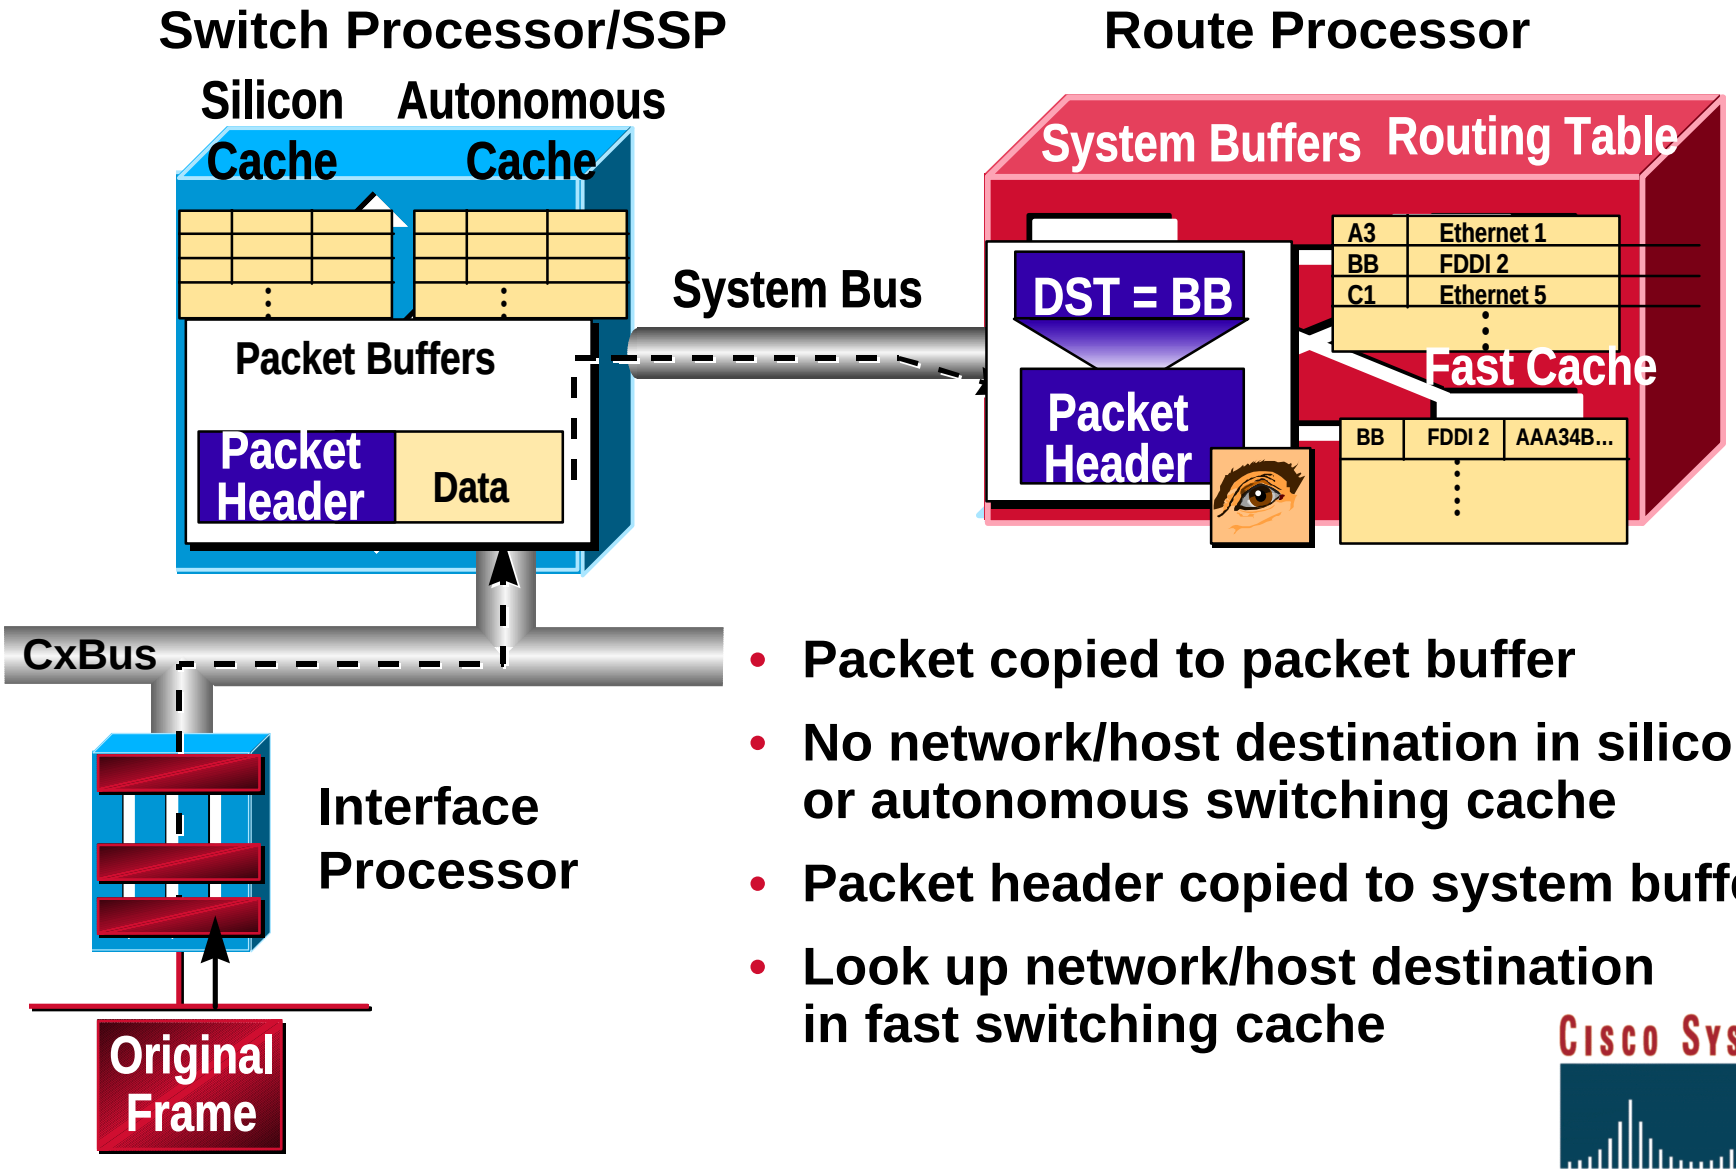

- Packet copied to packet buffer
- No network/host destination in silicon or autonomous switching cache
- Packet header copied to system buffer
- Look up network/host destination in fast switching cache

Cisco Systems

Cisco 7000 Series—Fast Switching

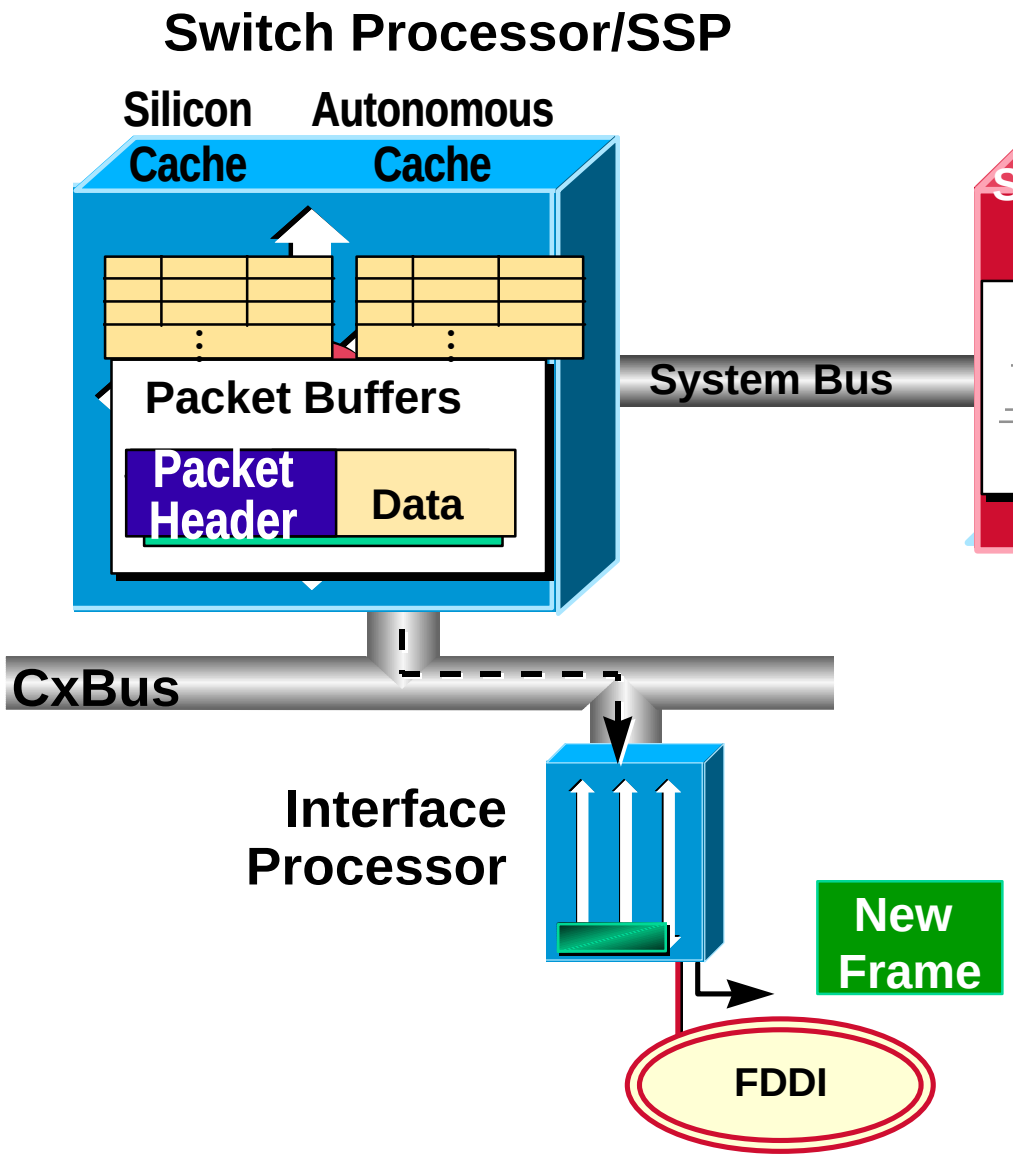

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco 7000 Series Autonomous Switching

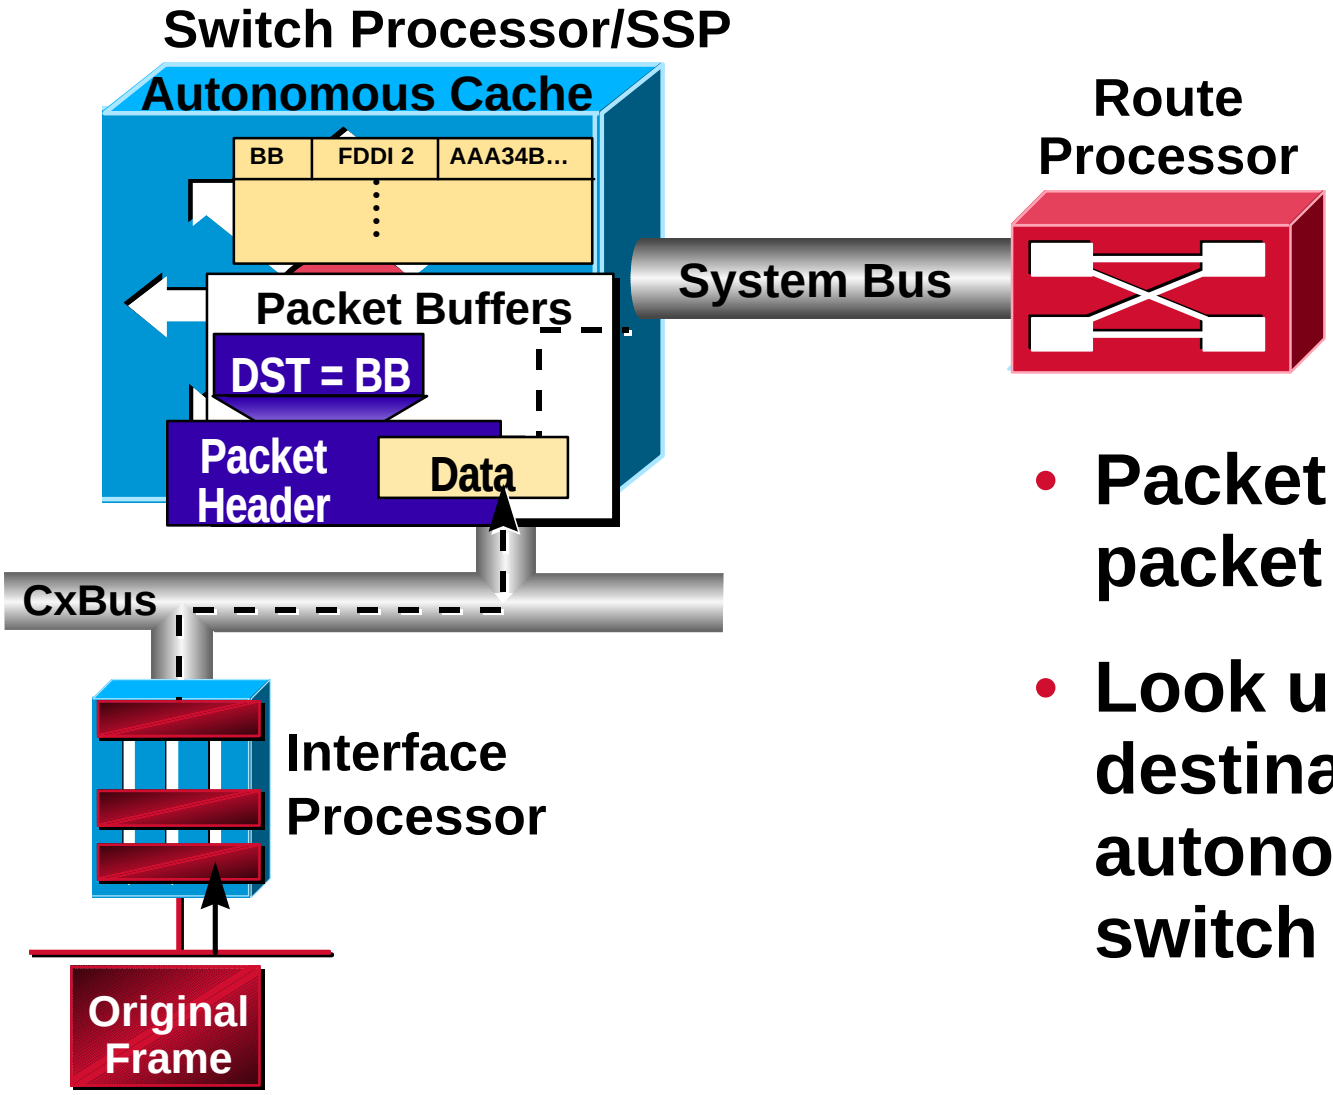

- Packet copied to packet buffer
- Look up Layer 3 destination in autonomous switch cache

Cisco Systems

Cisco 7000 Series Autonomous Switching

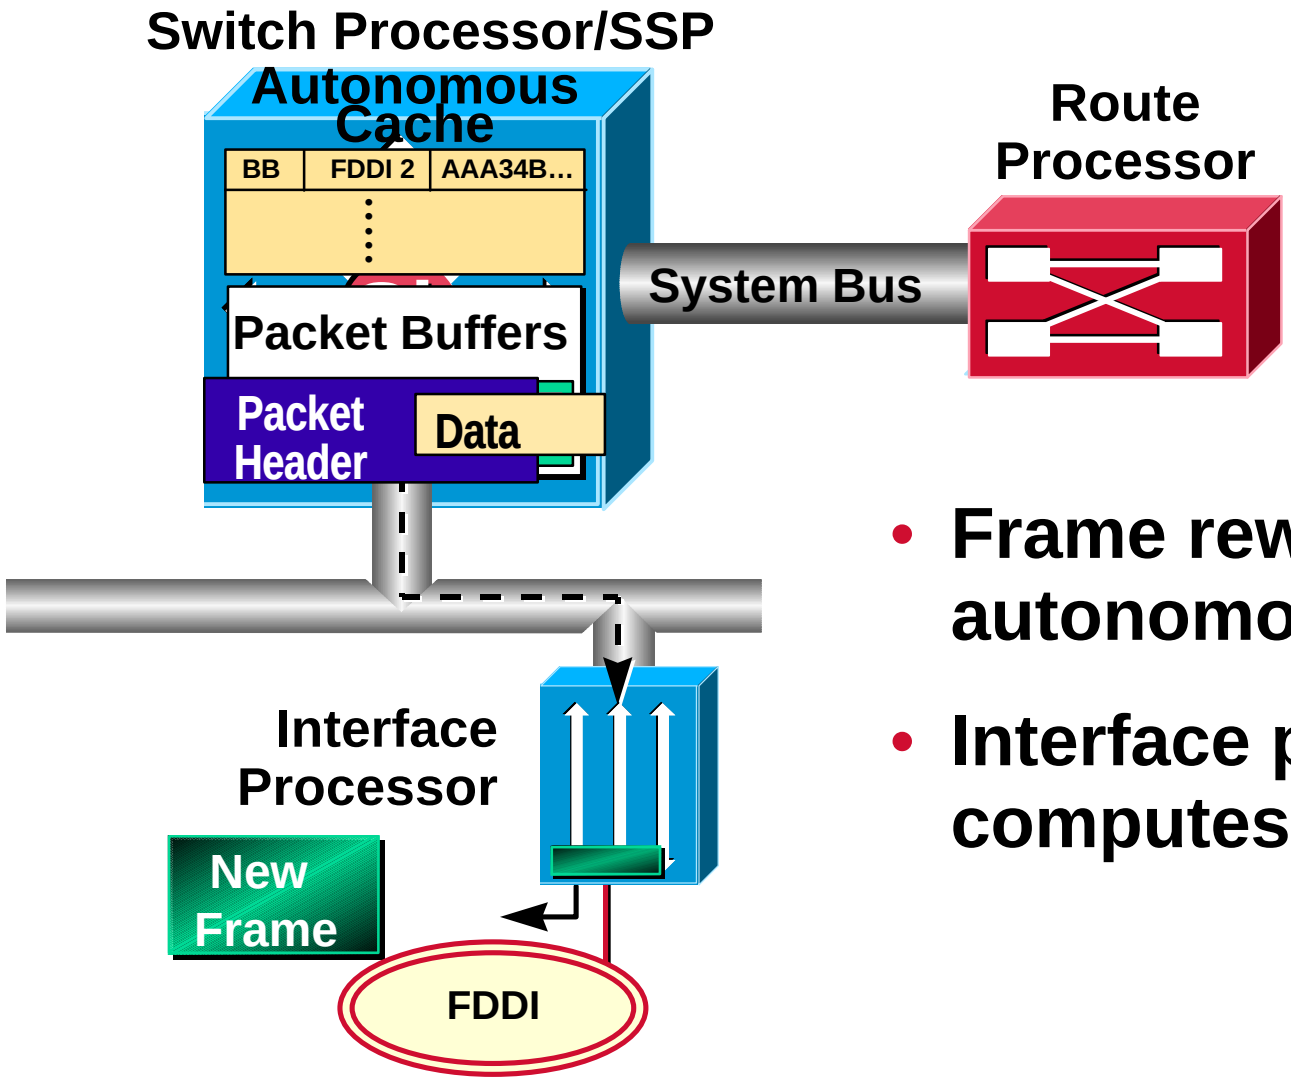

- Frame rewrite using autonomous cache entry
- Interface processor computes CRC

Cisco Systems

Cisco 7000 Series Silicon Switching

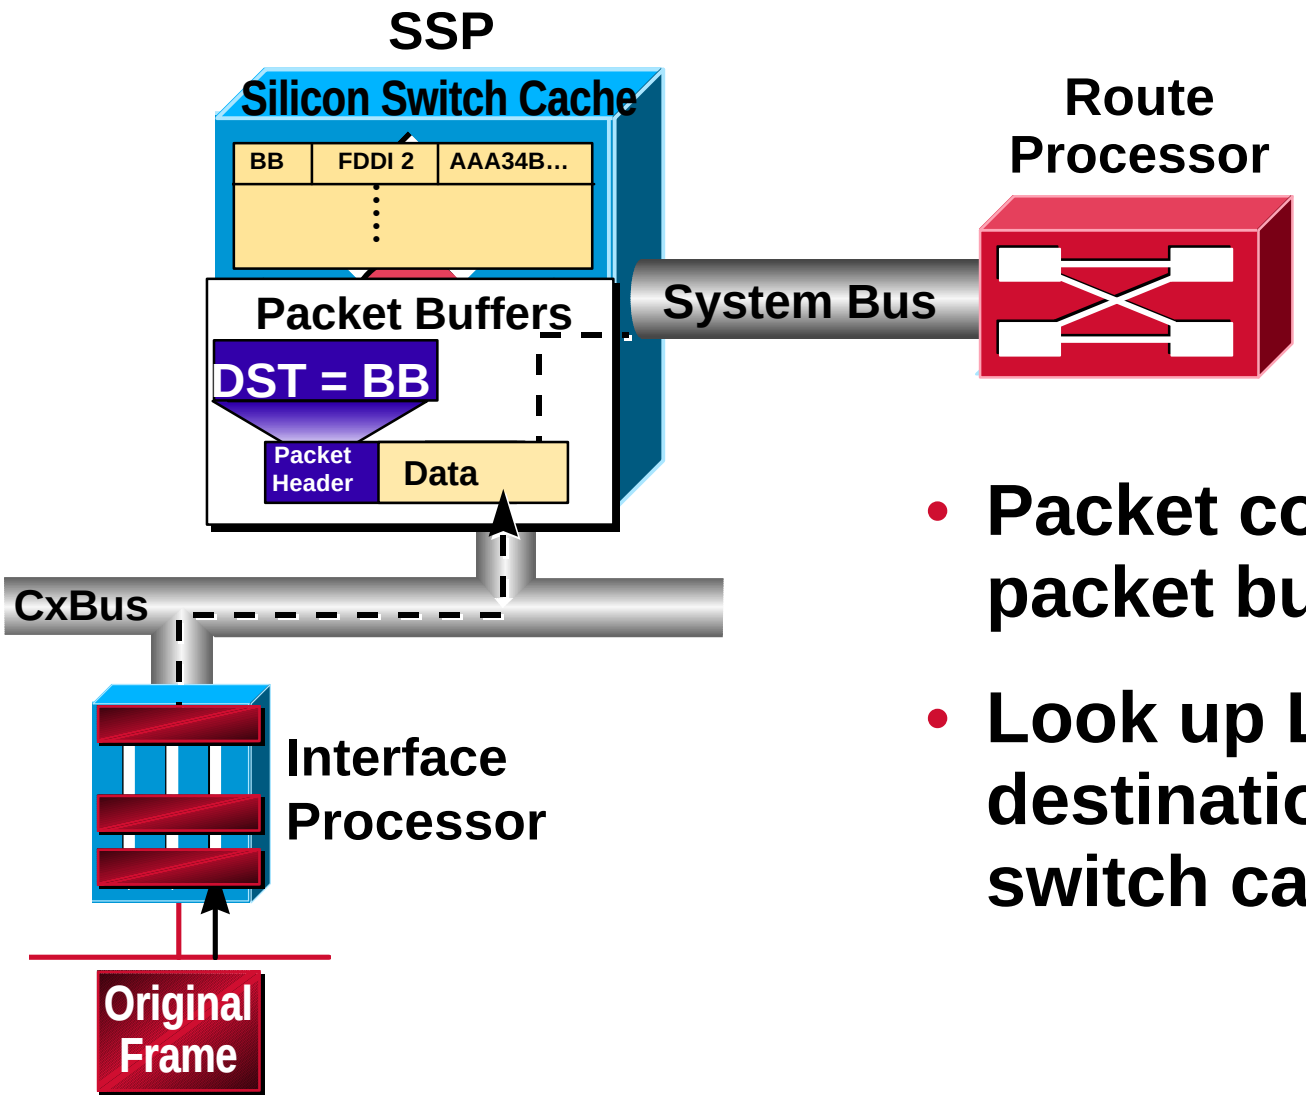

- Packet copied to packet buffer
- Look up Layer 3 destination in silicon switch cache

Cisco Systems

Cisco 7000 Series Silicon Switching

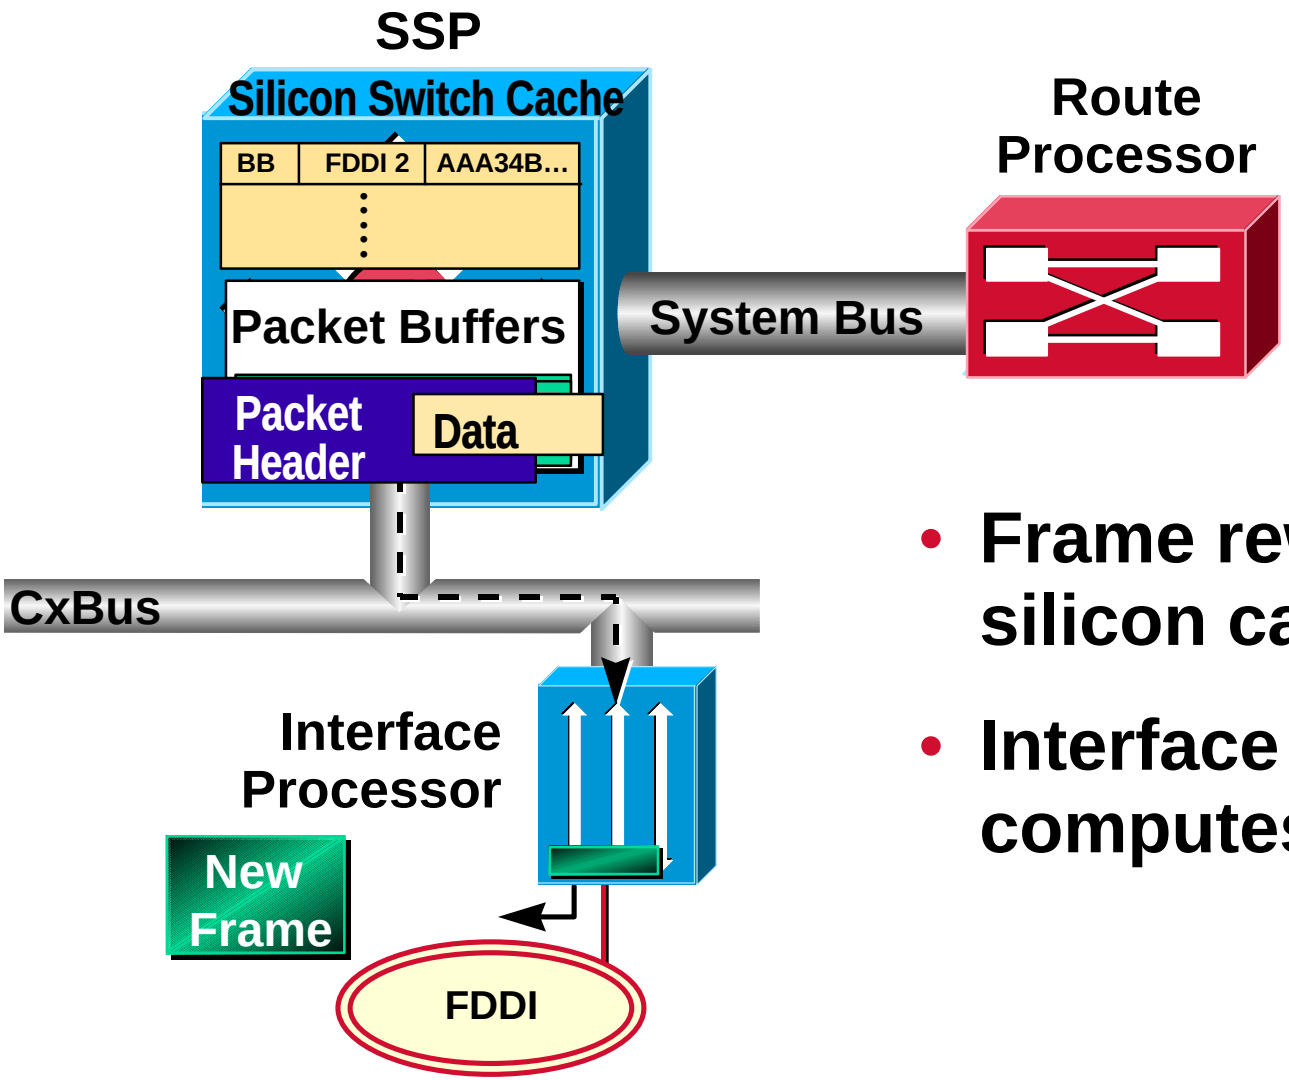

- Frame rewrite using silicon cache entry
- Interface processor computes CRC

Cisco Systems

Cisco 7000 Series—Switching Summary

Process	2,500 pps no <protocol> route-cache
Fast	30,000 pps <protocol> route-cache
Autonomous	200,000 pps <protocol> route-cache Cisco bus
Silicon	271,000 pps <protocol> route-cache SSE

Cisco 7200/7500 Series Architecture

Cisco 7206 Architecture

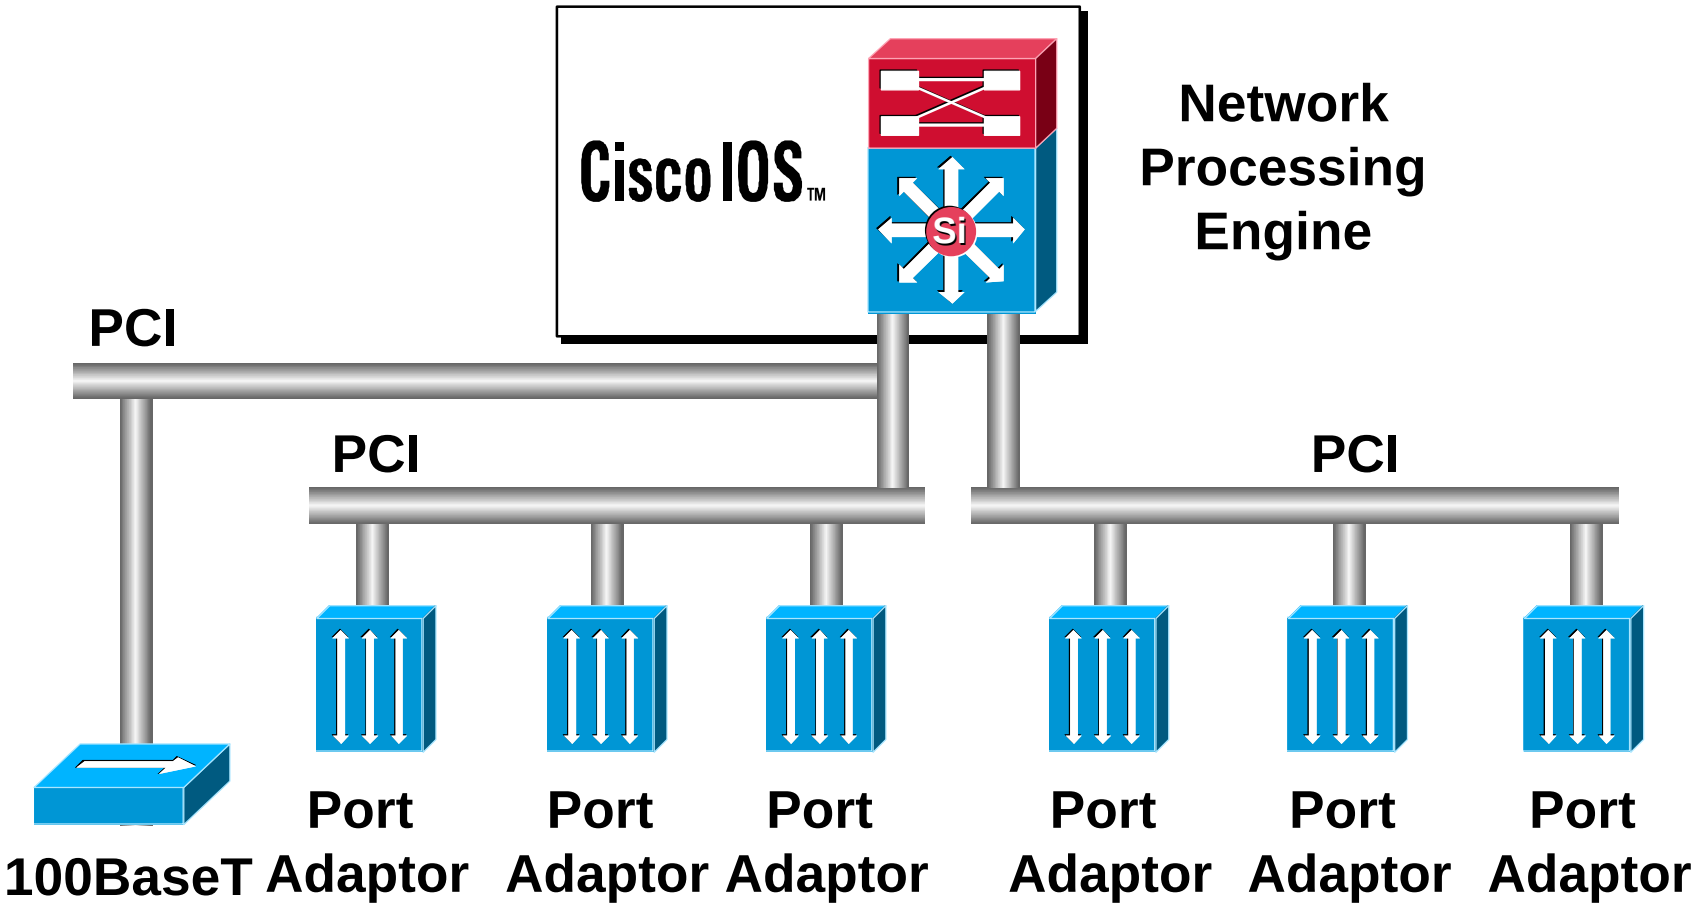

Cisco Systems

Network Processing Engine

	NPE 100	NPE 150	NPE 200
Processor Type	150 MHz Orion R4700 RISC	150 MHz Orion R4700 RISC	200 MHz Orion R5000 RISC
System Controller	Galileo GT-64010	Galileo GT-64010	Galileo GT-64010
Aggregate Bus Bandwidth	200 Mbps	600 Mbps	600 Mbps
DRAM	16 -› 128 MB	16 –› 128 MB	16 –› 128 MB
SRAM	N/A	1 MB	1 MB

CISCO SYSTEMS

Cisco 7505 Architecture

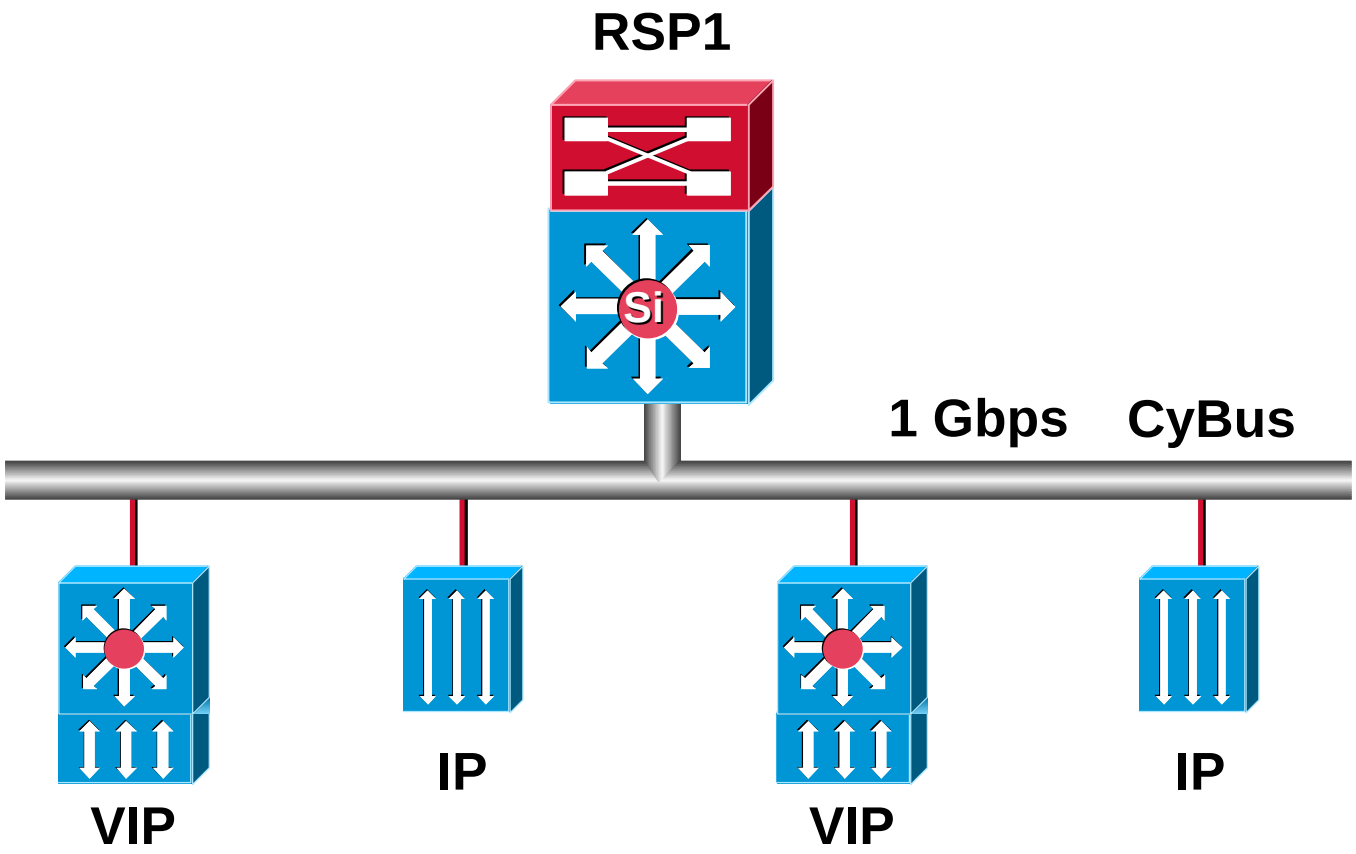

Cisco 7507 Architecture

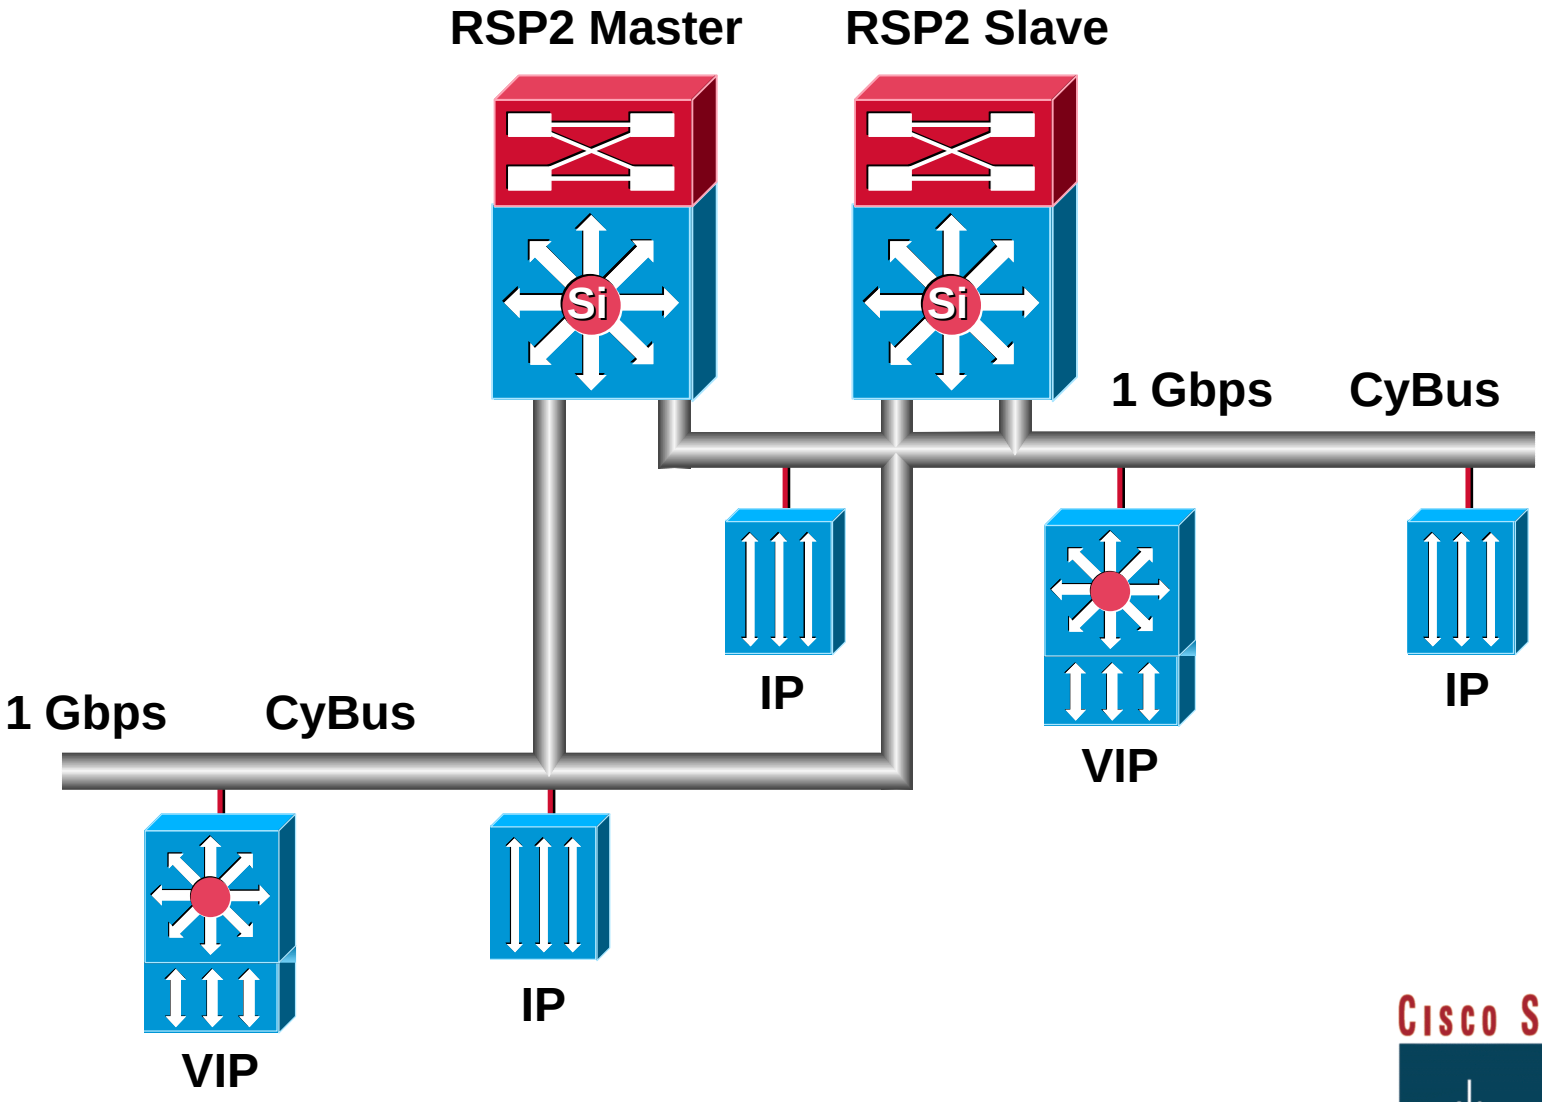

Cisco Systems

Cisco 7513 Architecture

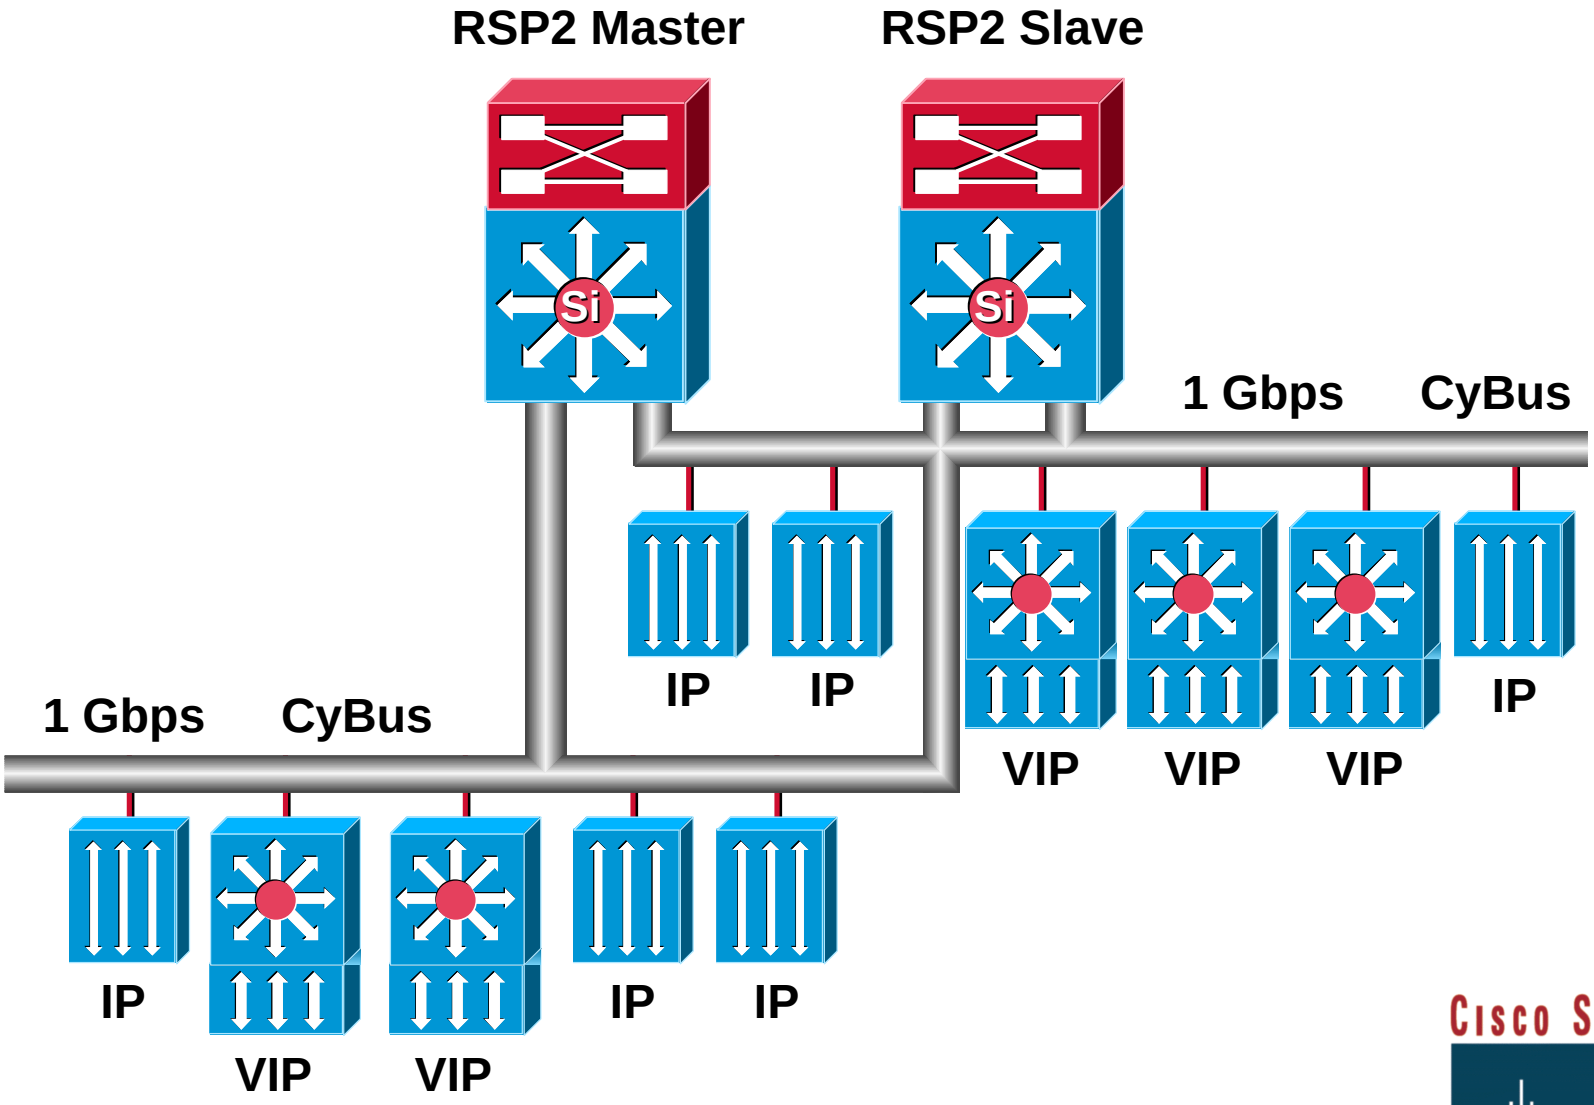

Cisco Systems

Route Switch Module

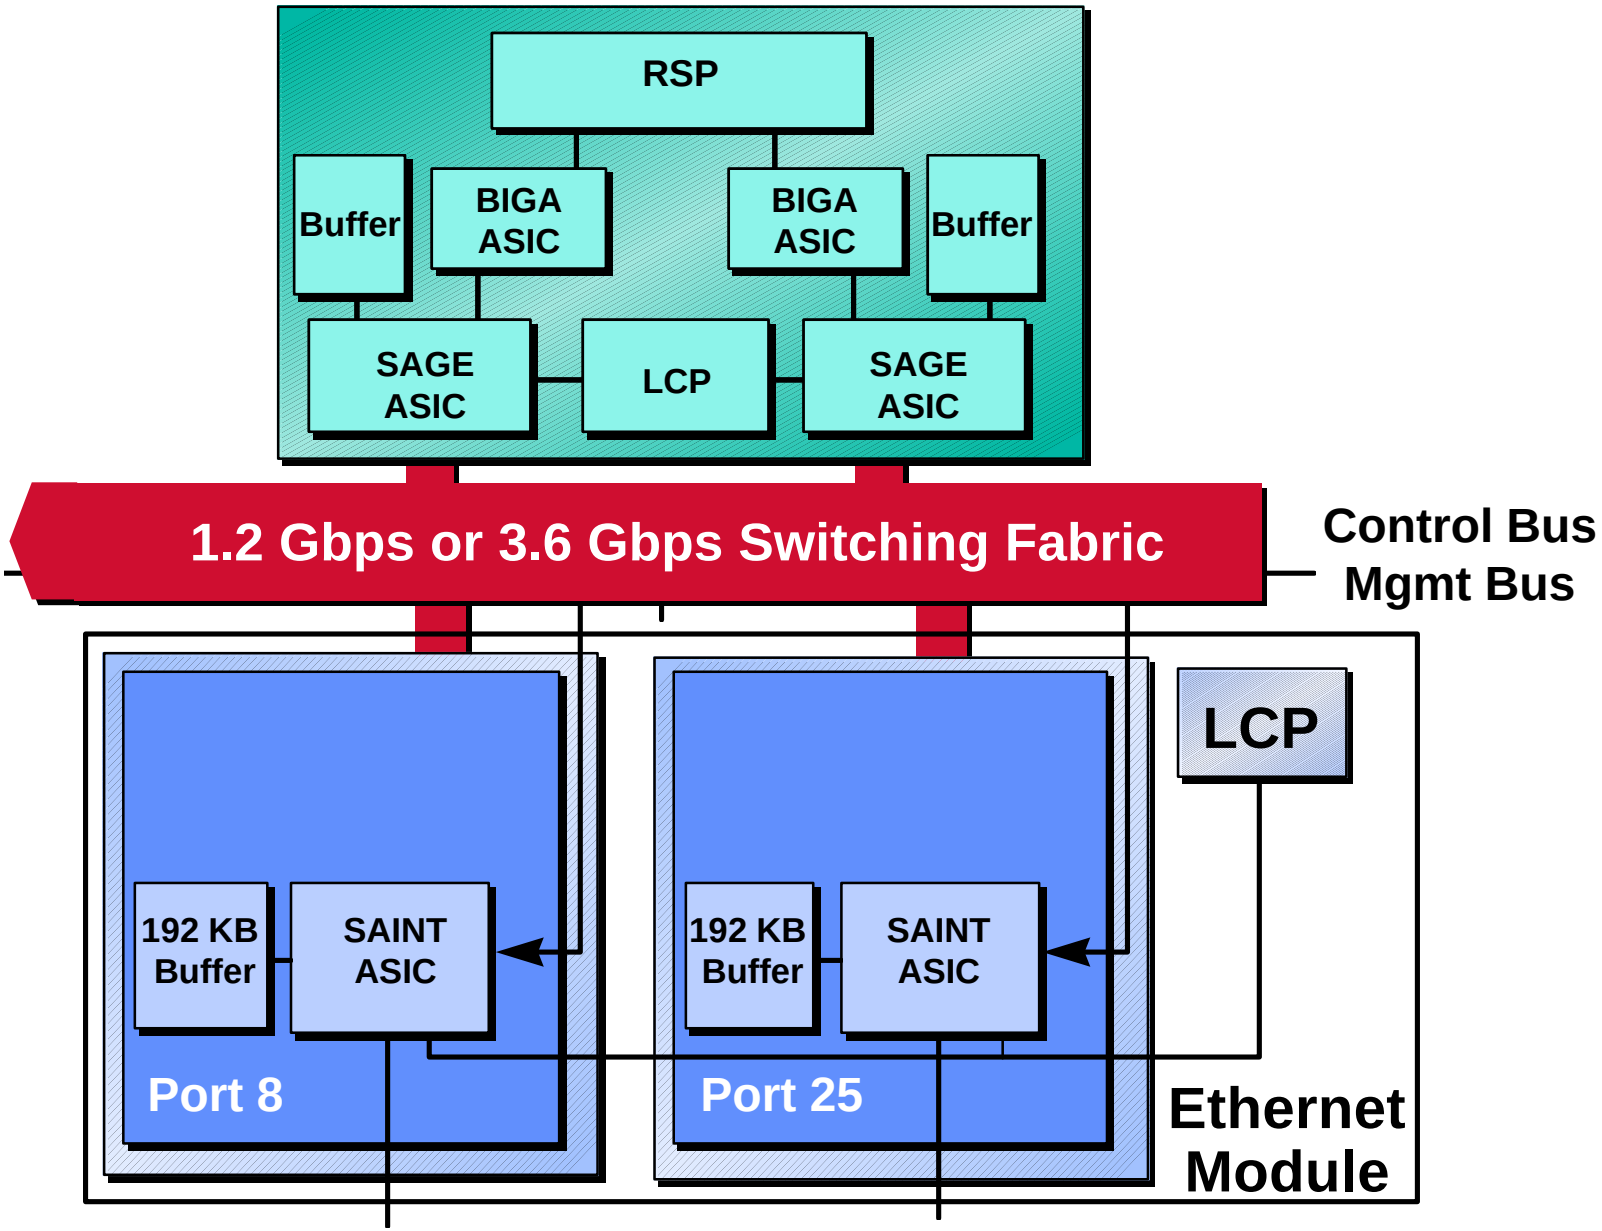

Route Switch Module

Data Flow

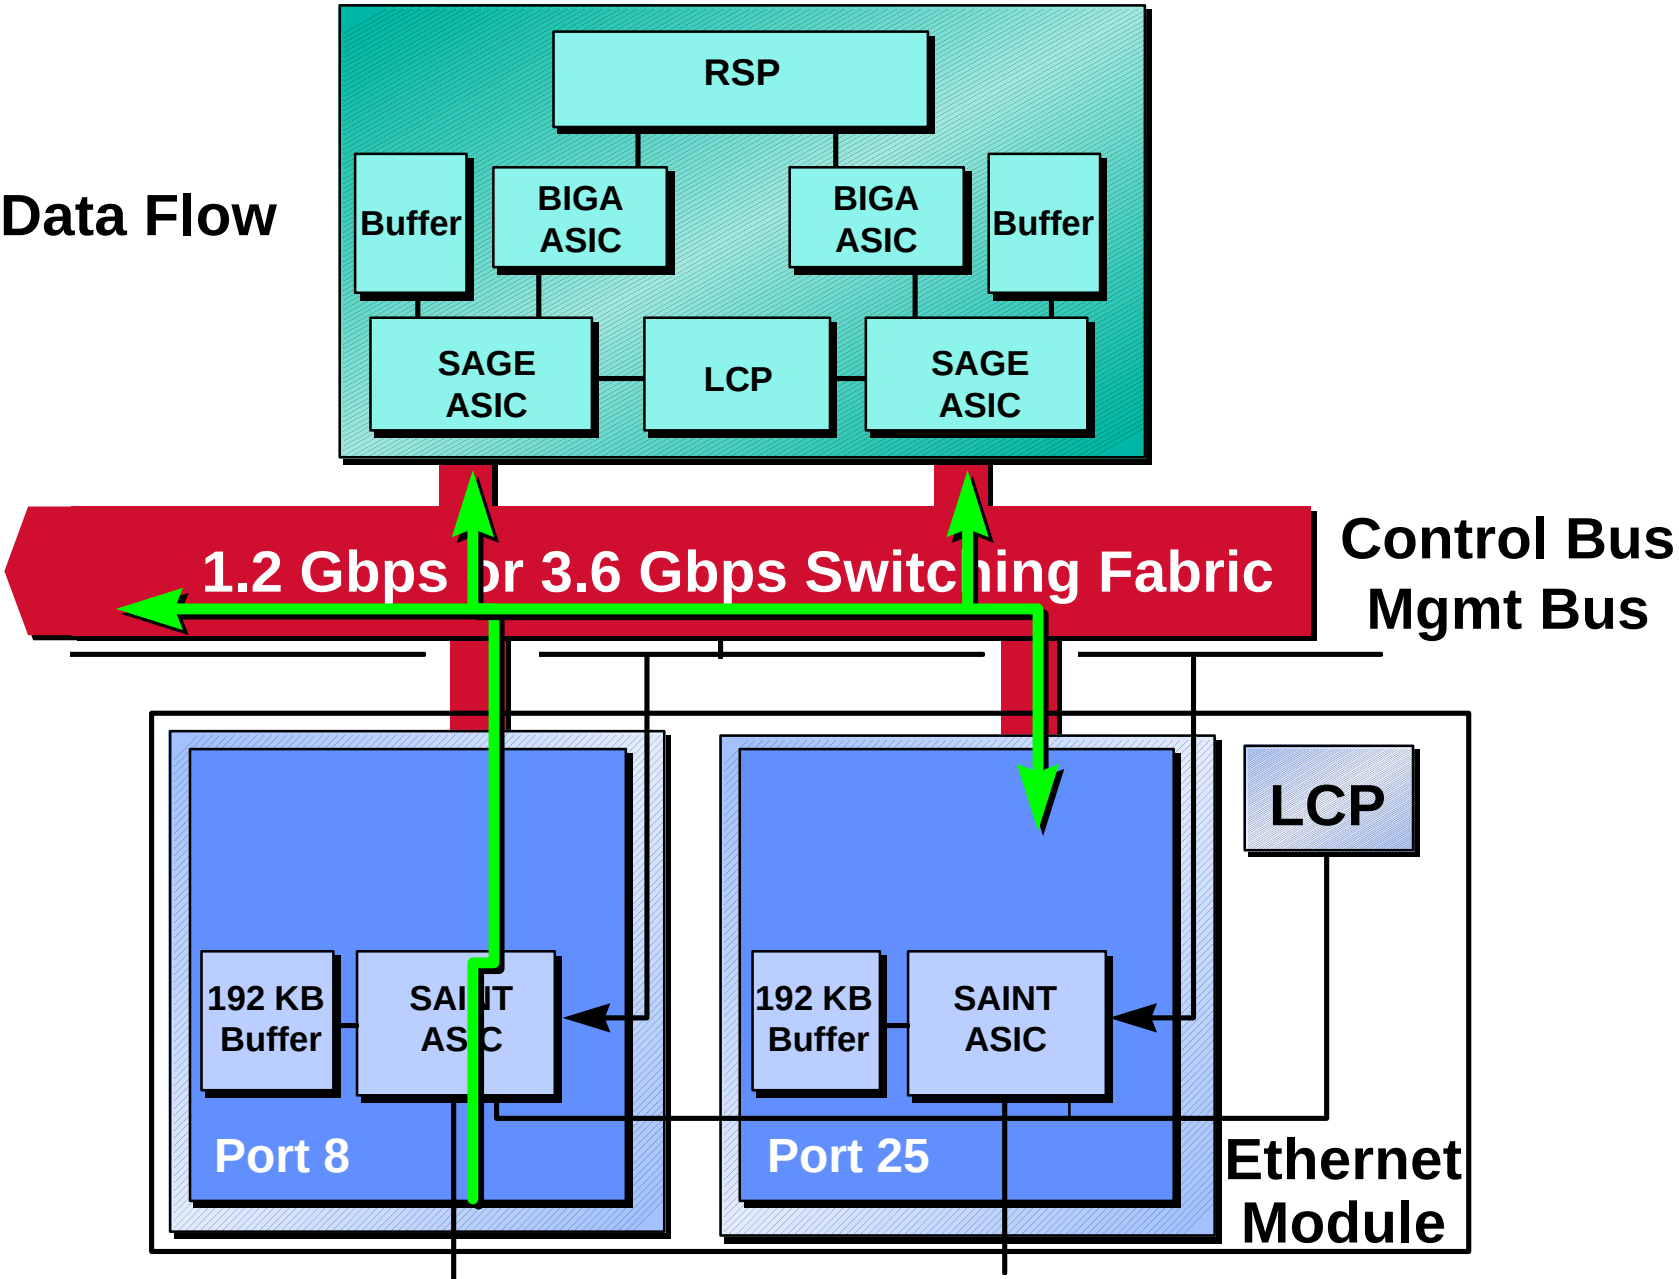

Route Switch Module

Data Flow

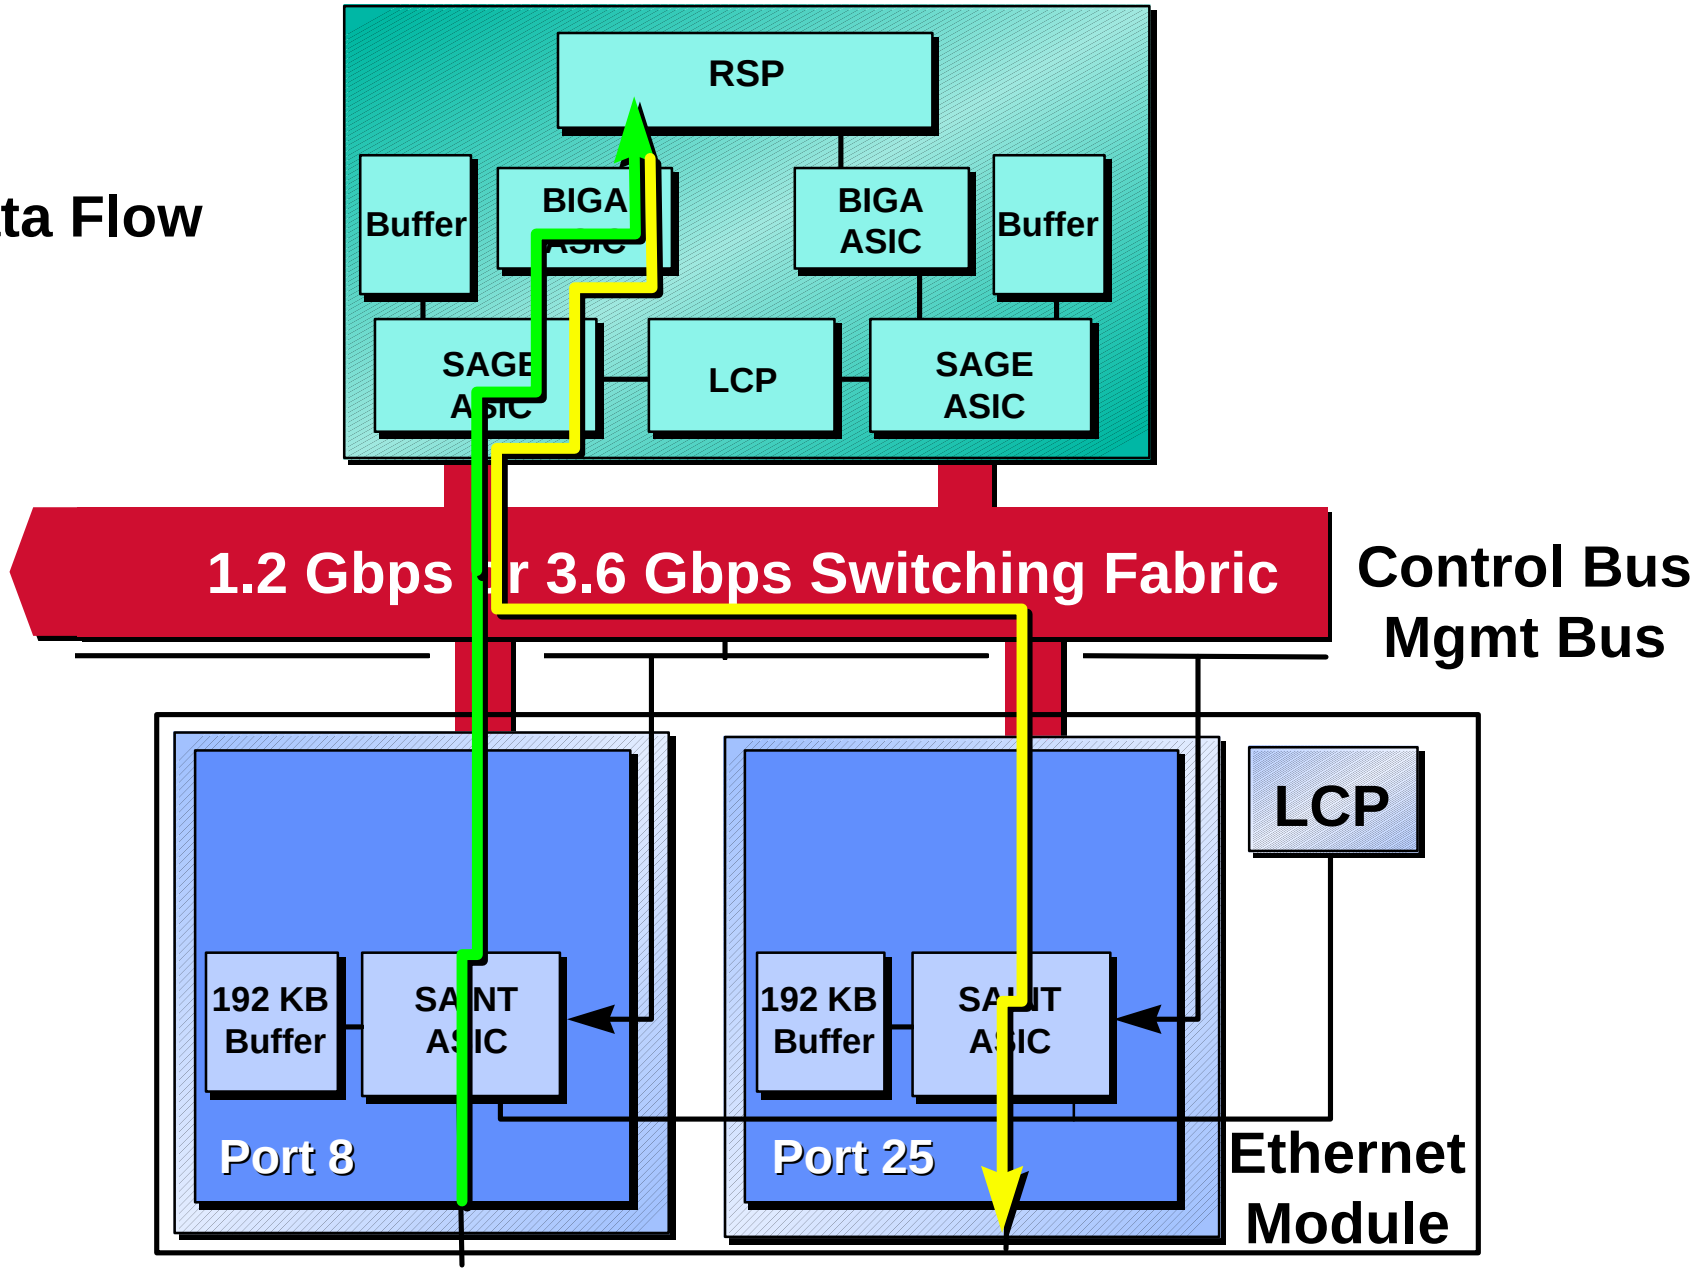

Cisco 7500 Series—Summary of Hardware Features

Characteristics	Cisco 7513	Cisco 7507	Cisco 7505
Number of Slots	13	7	5
Number of Route/Switch Processors	2-RSP2/4s	2-RSP2/4s	1-RSP1
Processor Type	MIPS RISC 100/200 MHz R4600/R5000v	MIPS RISC 100/200 MHz R4600/R5000	MIPS RISC 100 MHz R4600
System Bandwidth	2.132 Gbps	2.132 Gbps	1.066 Gbps
Number of Interface Processor Slots	11	5	4
Power Supplies	2	2	1

Cisco 7500 Series RSP Comparison

Memory Type	Size (RSP1/2)	Size (RSP4)	Description
DRAM	16–128 MB	32-256MB	8-, 16-, or 32-Mb SIMMs
NVRAM	128 Kb	128kb	Non-Volatile EPROM for System Configuration File
Flash SIMM	8MB	16MB	Cisco IOS Boot Images
Flash Memory Card (PCMCIA)	8, 16, 20 MB	8,16,20MB	Cisco IOS Images on Up to Two (2) PCMCIA Cards
Boot ROM	256 Kb	256	EPROM for ROM Monitor Program
Packet Memory	2 MB	2MB	Packet Buffers
Processor Type	100MHz MIPS	200MHz MIPS	Main Processor

Versatile Interface Processor (VIP)

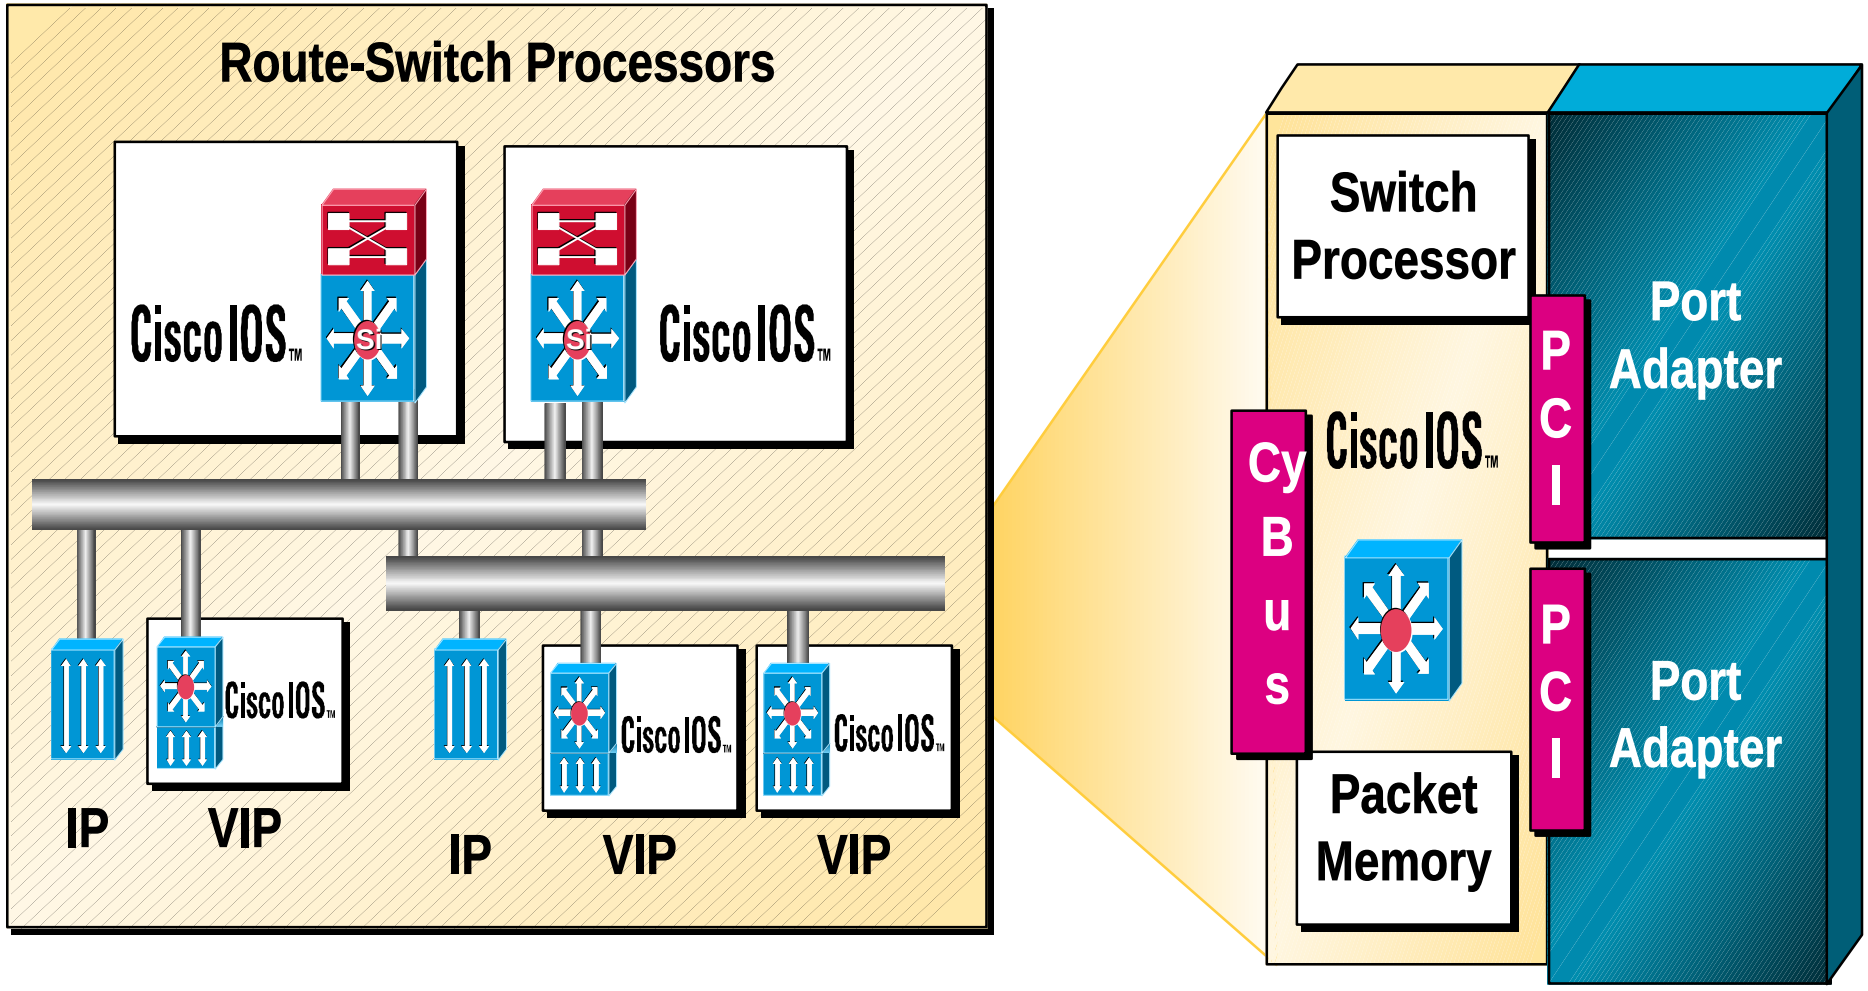

VIP2—Variety of Choices

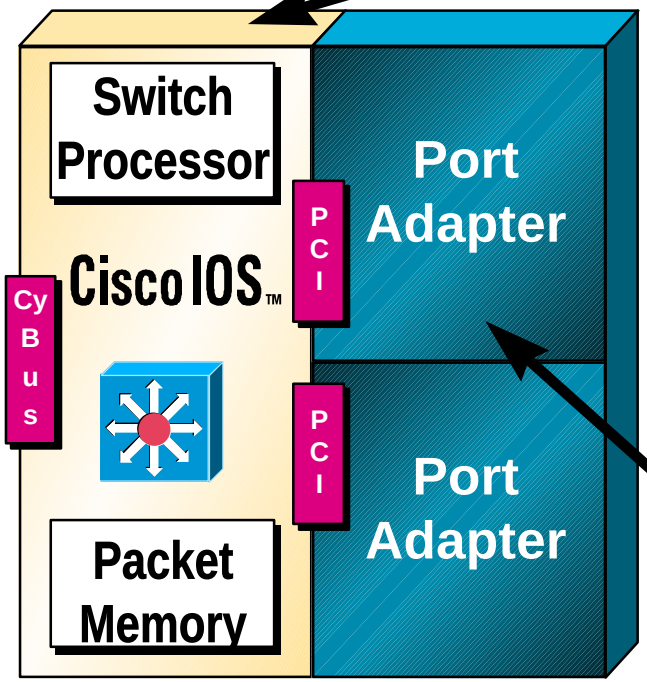

Mother Board Options

PRODUCT	SRAM	DRAM	Dist Sw	Dist Svcs
VIP2-15	1MB	8MB	No	No
VIP2-40	2MB	32MB	Yes	Yes
VIP2-50	8MB	32MB	Yes	Yes

Port Adapter Options

WAN	LAN
4 and 8 Serial	4 and 8 Ethernet
2 PRI—ISDN	5 Ethernet (10BaseFL)
2 Channelized T1/E1	1 Fast Ethernet (TX/FX)
1 and 2 HSSI	4 Token Ring (HD/FD)
	1 FDDI (HD/FD)
	1 100VG
	1 ATM OC-3

Service Adapter Options

Compression

RSP-Based Router Switching Paths

RSP-Based Router Switching Paths

Switching Path	Cisco 7000 w/RSP	Cisco 7200	Cisco 7500	Comments
Process Switching	Yes	Yes	Yes	Initializes Switching Caches
Fast Switching	Yes	Yes	Yes	Default (except IP)
Optimum Switching	Yes	Yes	Yes	Default for IP
NetFlow Switching	Yes	Yes	Yes	Configurable per interface
Distributed Switching	Yes	No	Yes	Using VIP2-20/VIP2-40
CEF	Yes	Yes	Yes	New IP default (IOS 12.0)

Cisco 7500 Series—Memory Allocations

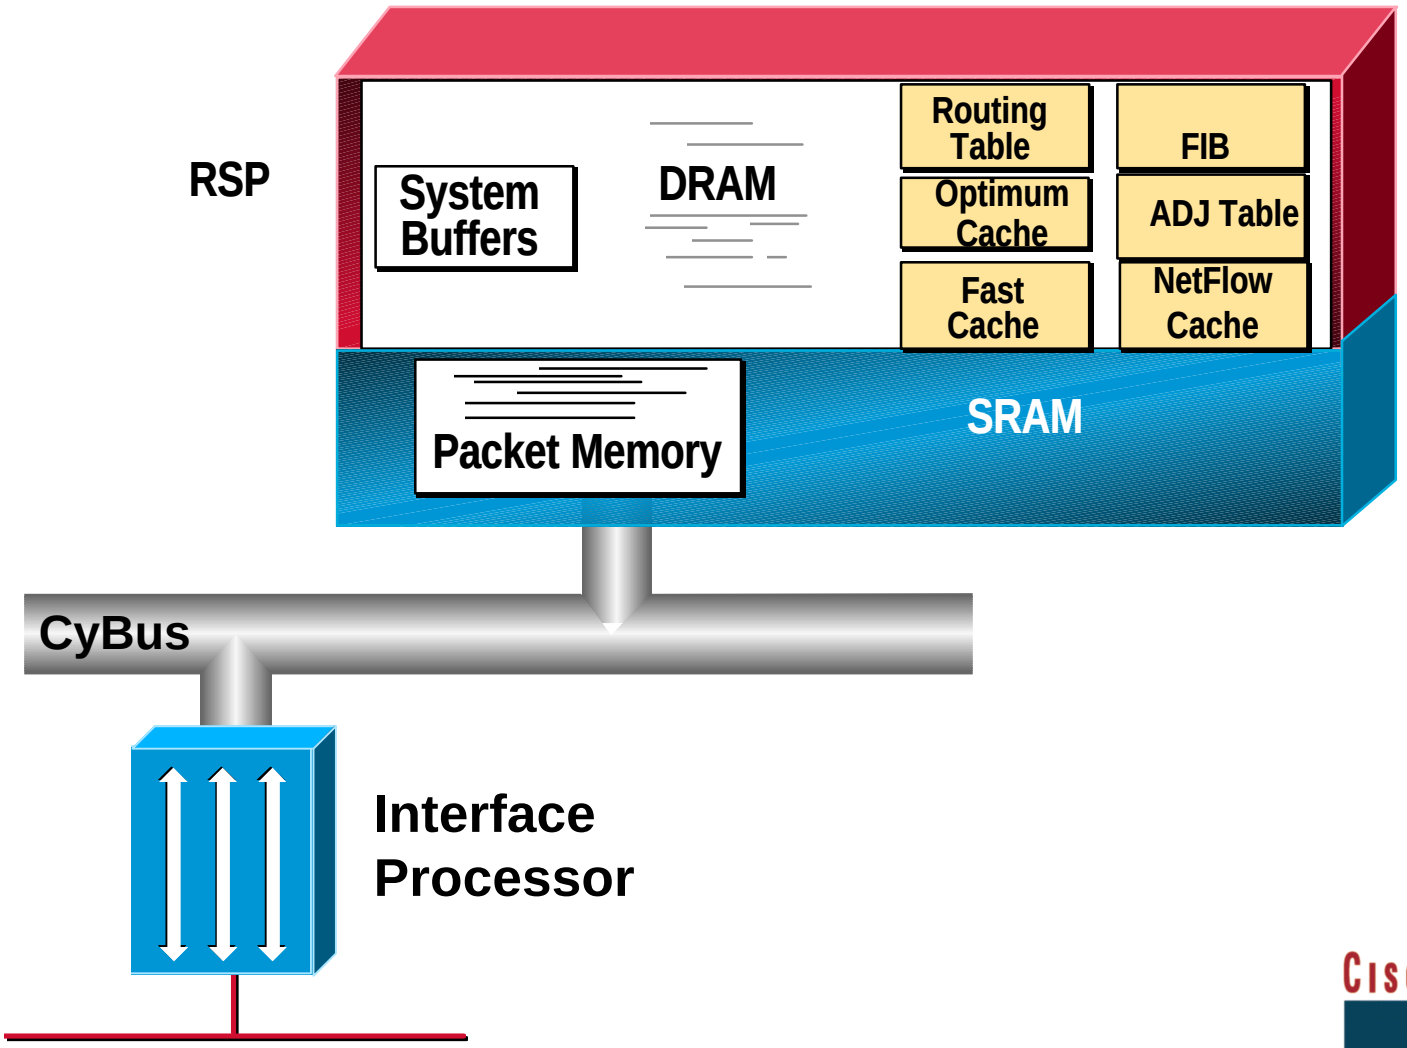

Cisco 7500 Series—Process Switching

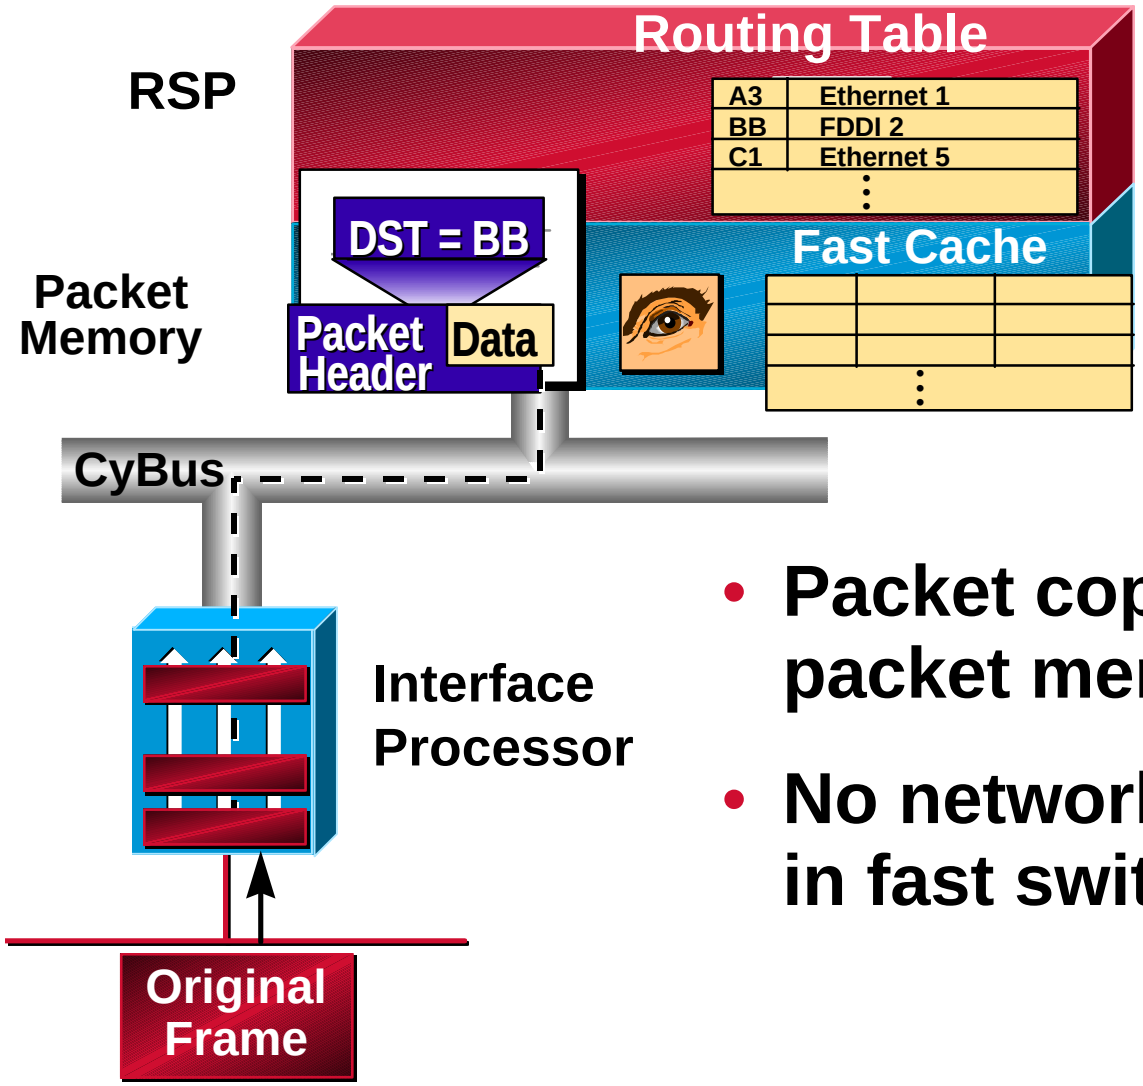

- Packet copied to packet memory
- No network/host destination in fast switching cache

Cisco Systems

Cisco 7500 Series—Process Switching

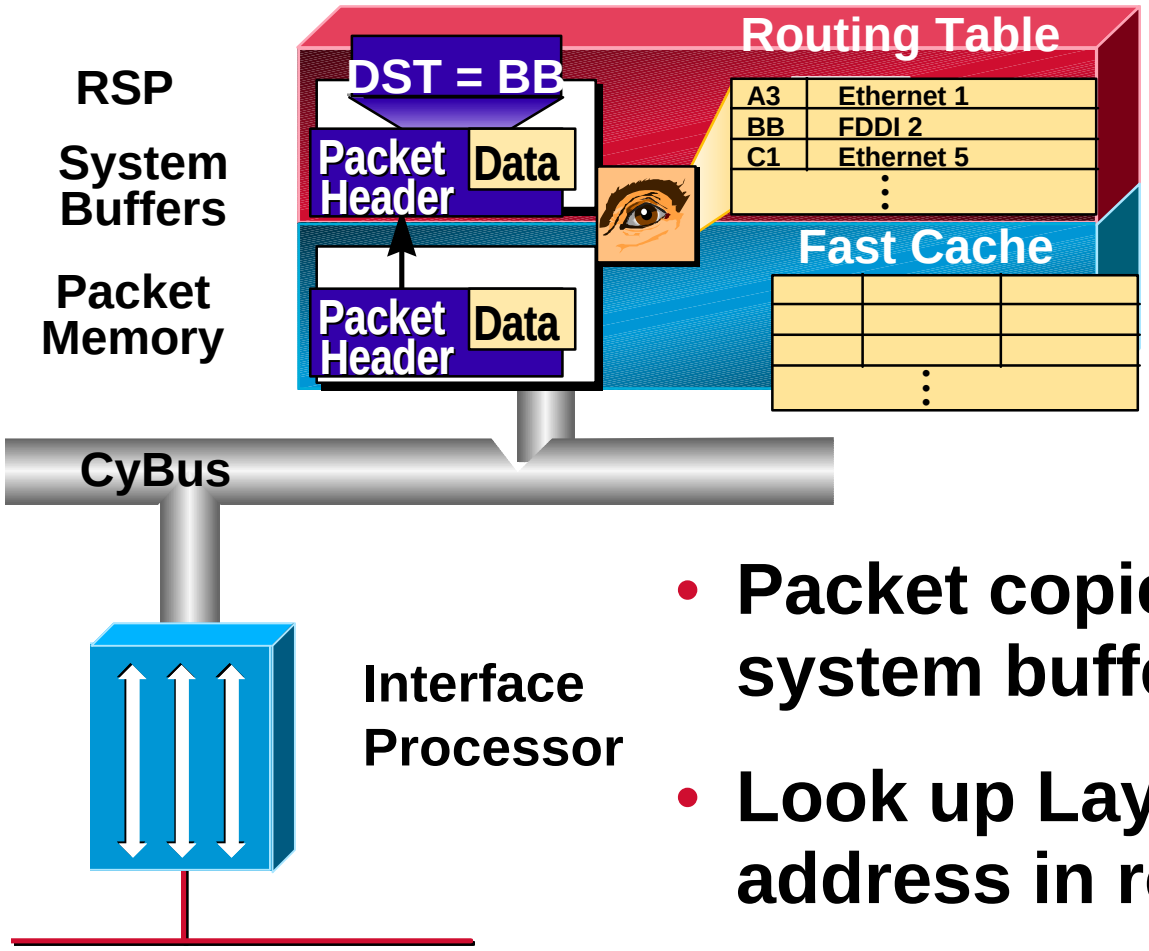

- Packet copied to system buffer
- Look up Layer 3 network address in routing table
- Initialize fast switch cache

Cisco 7500 Series—Process Switching

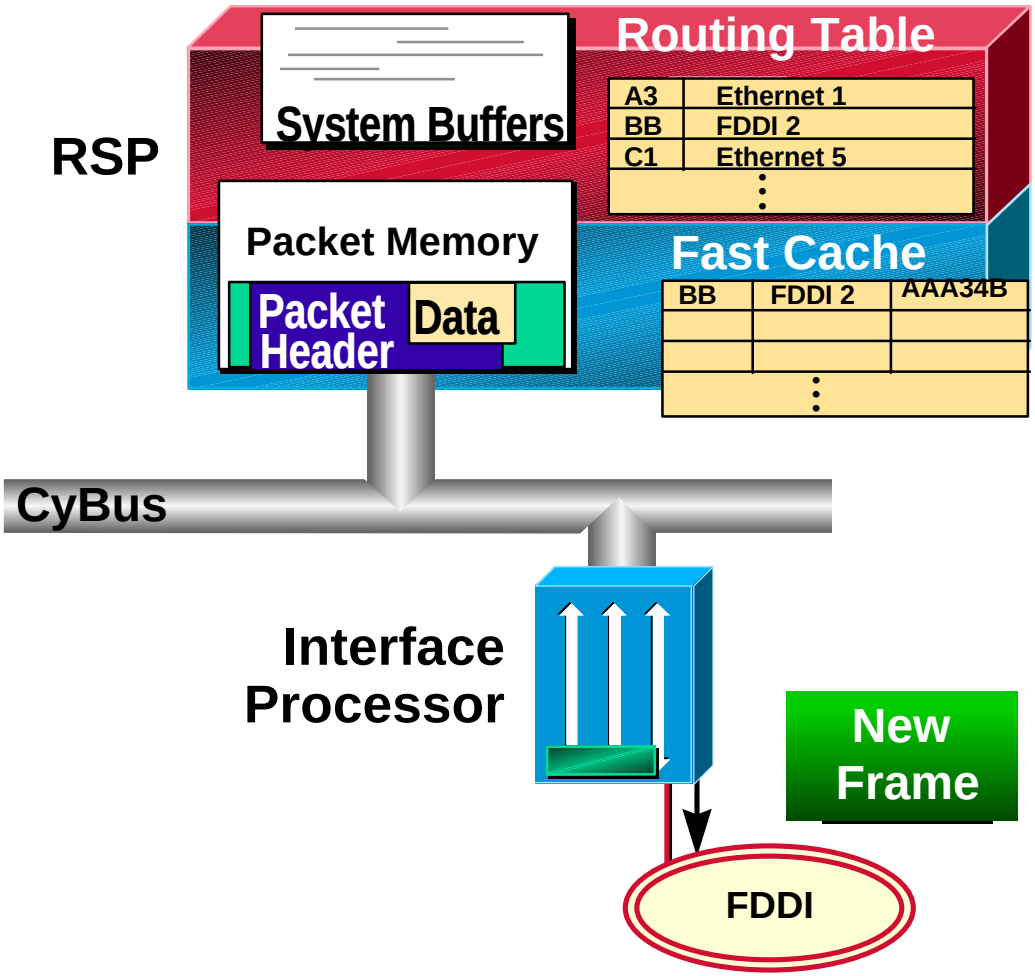

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco Systems

Cisco 7500 Series—Fast Switching

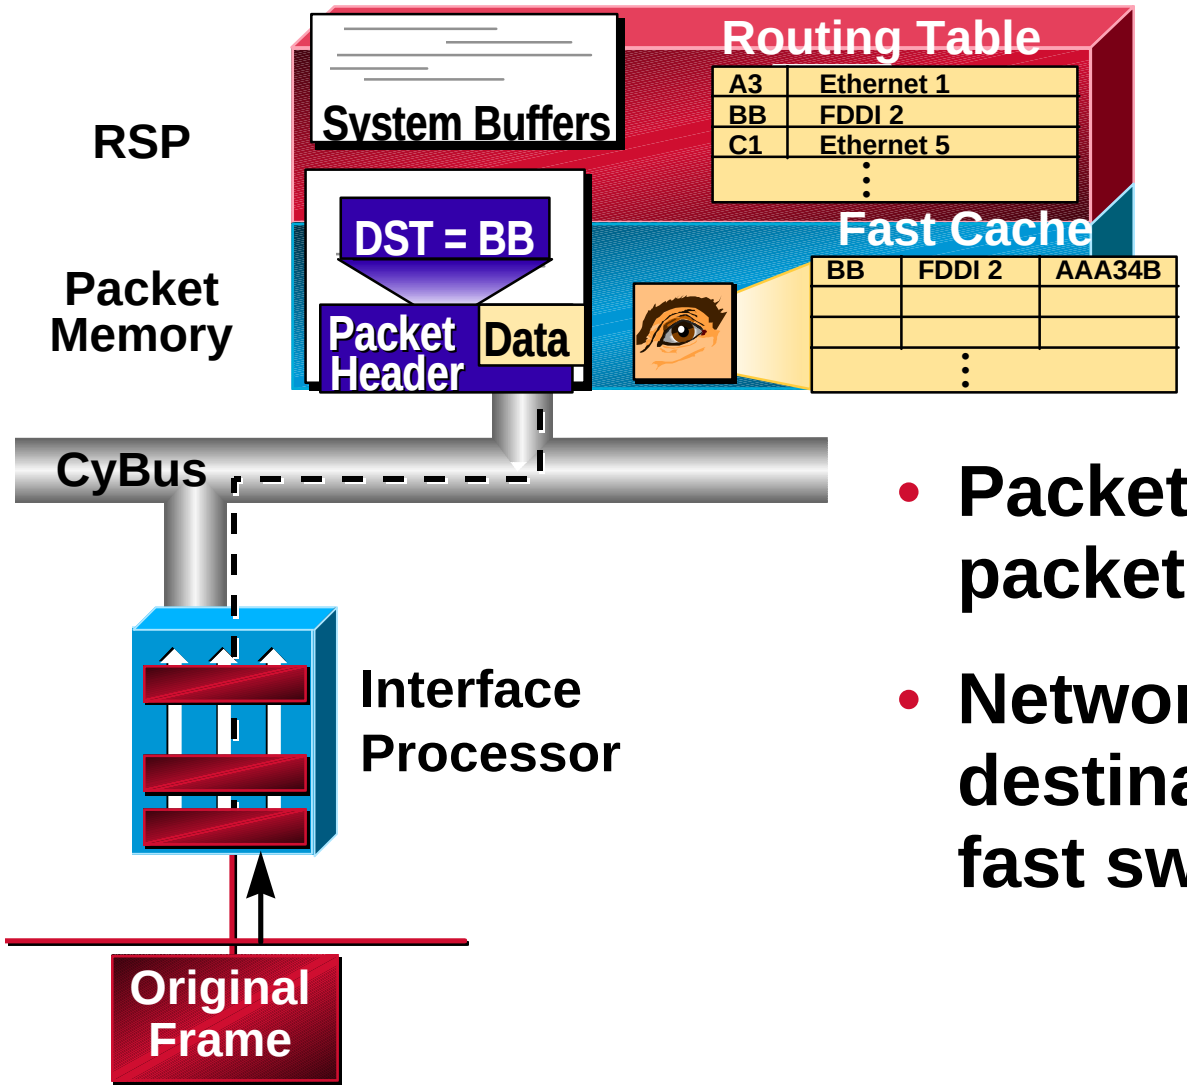

- Packet copied to packet memory
- Network/host destination found in fast switching cache

Cisco Systems

Cisco 7500 Series—Fast Switching

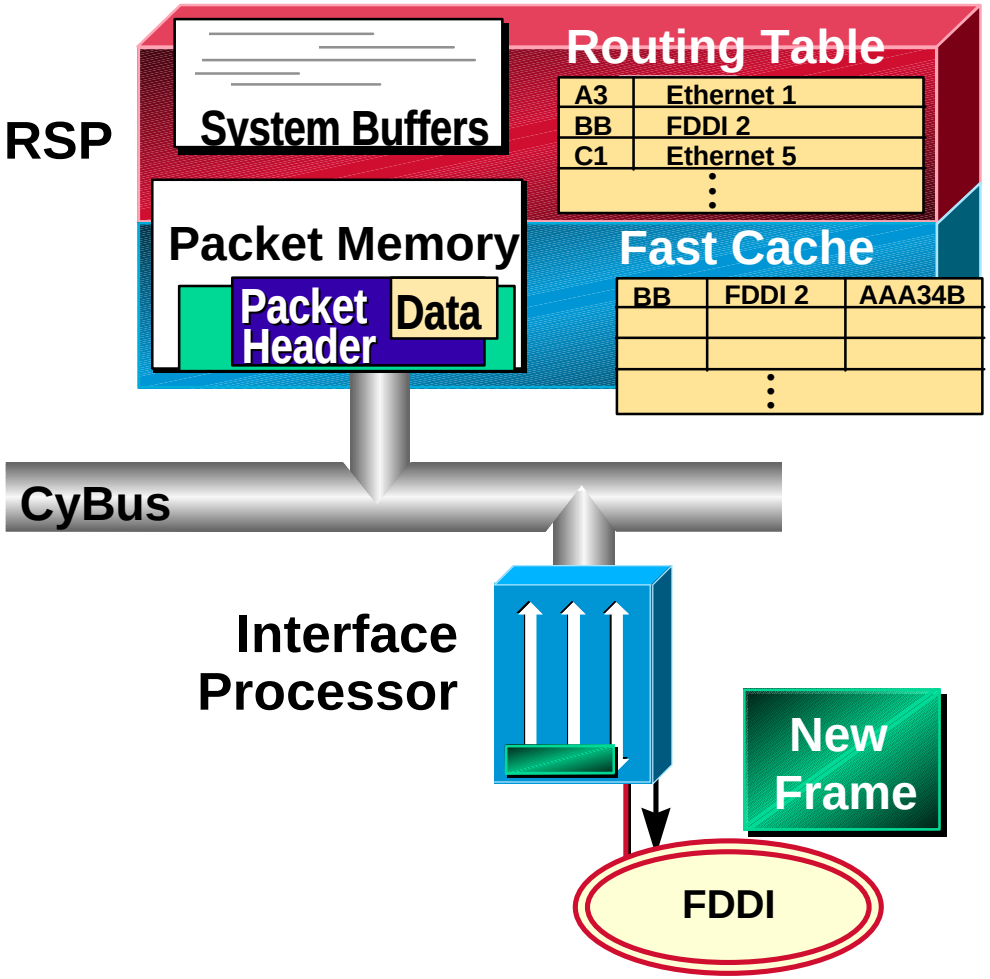

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco Systems

Cisco 7500 Series—Optimum Switching

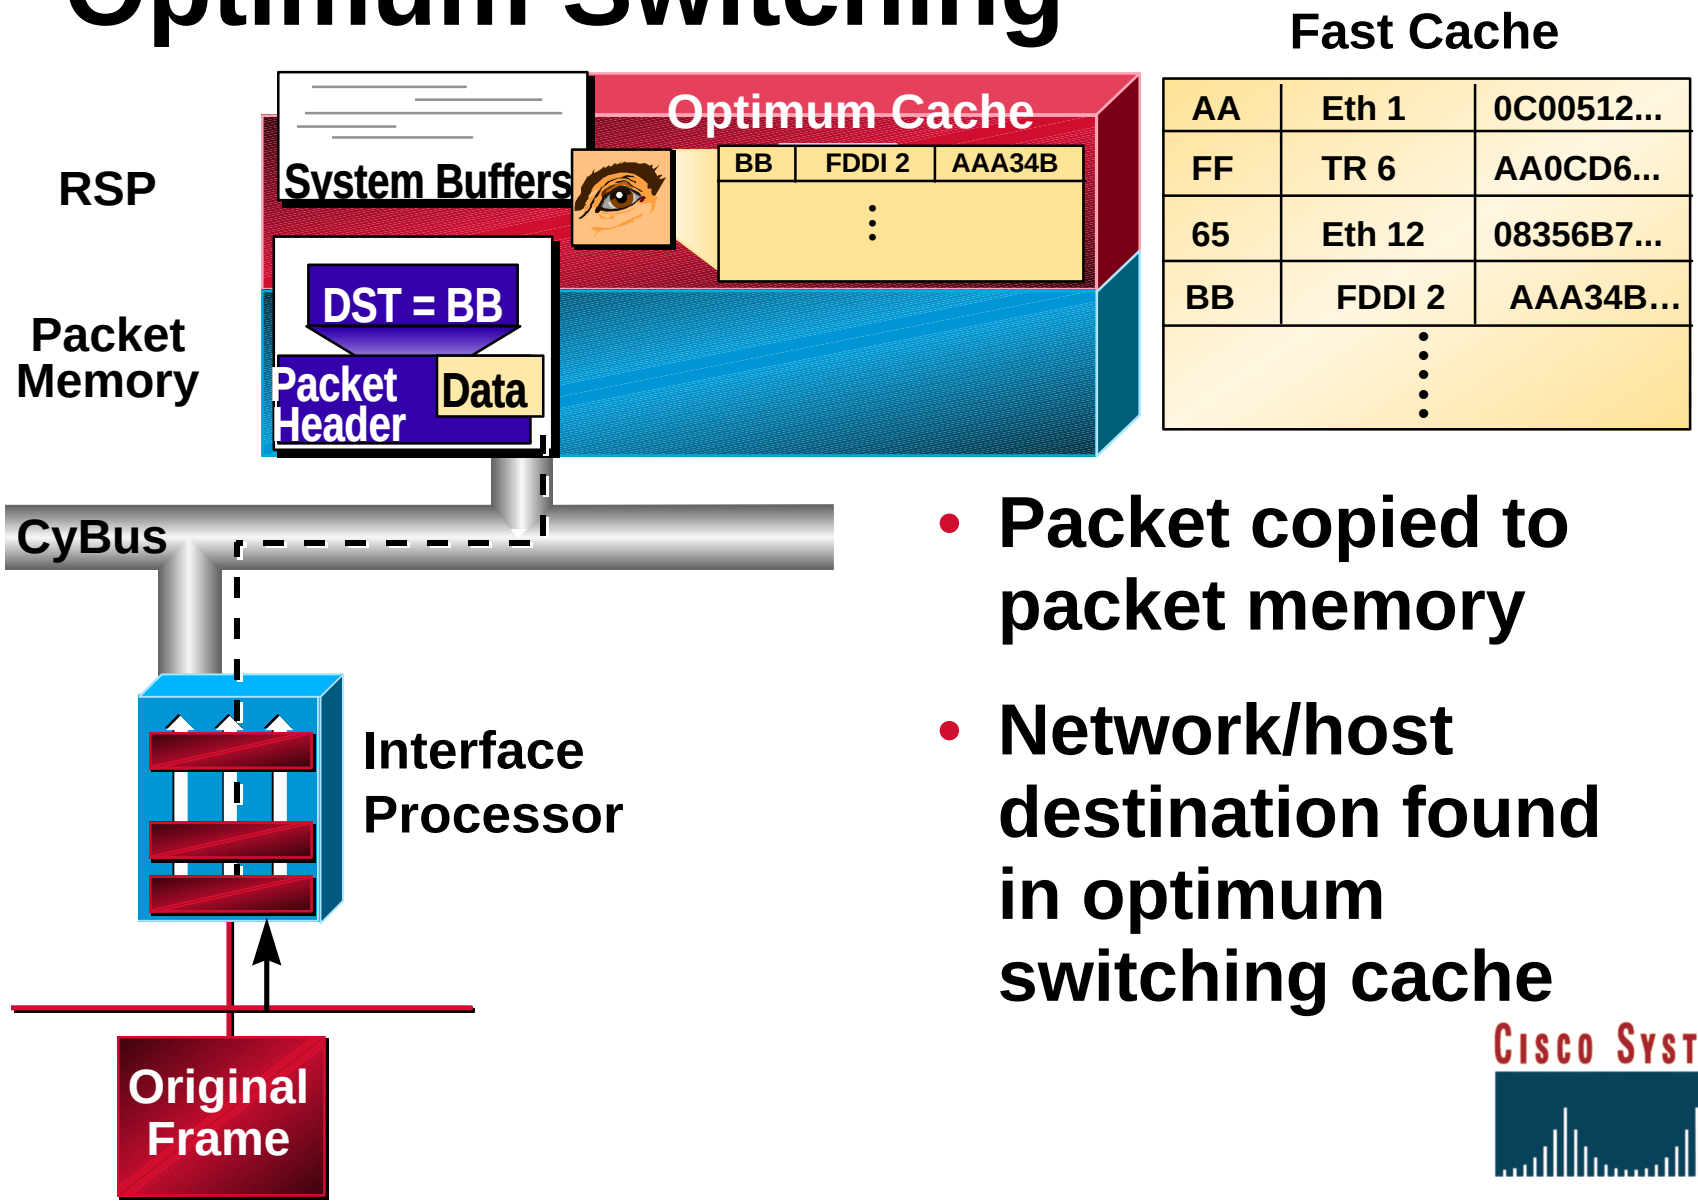

- Packet copied to packet memory
- Network/host destination found in optimum switching cache

Cisco Systems

Cisco 7500 Series—Optimum Switching

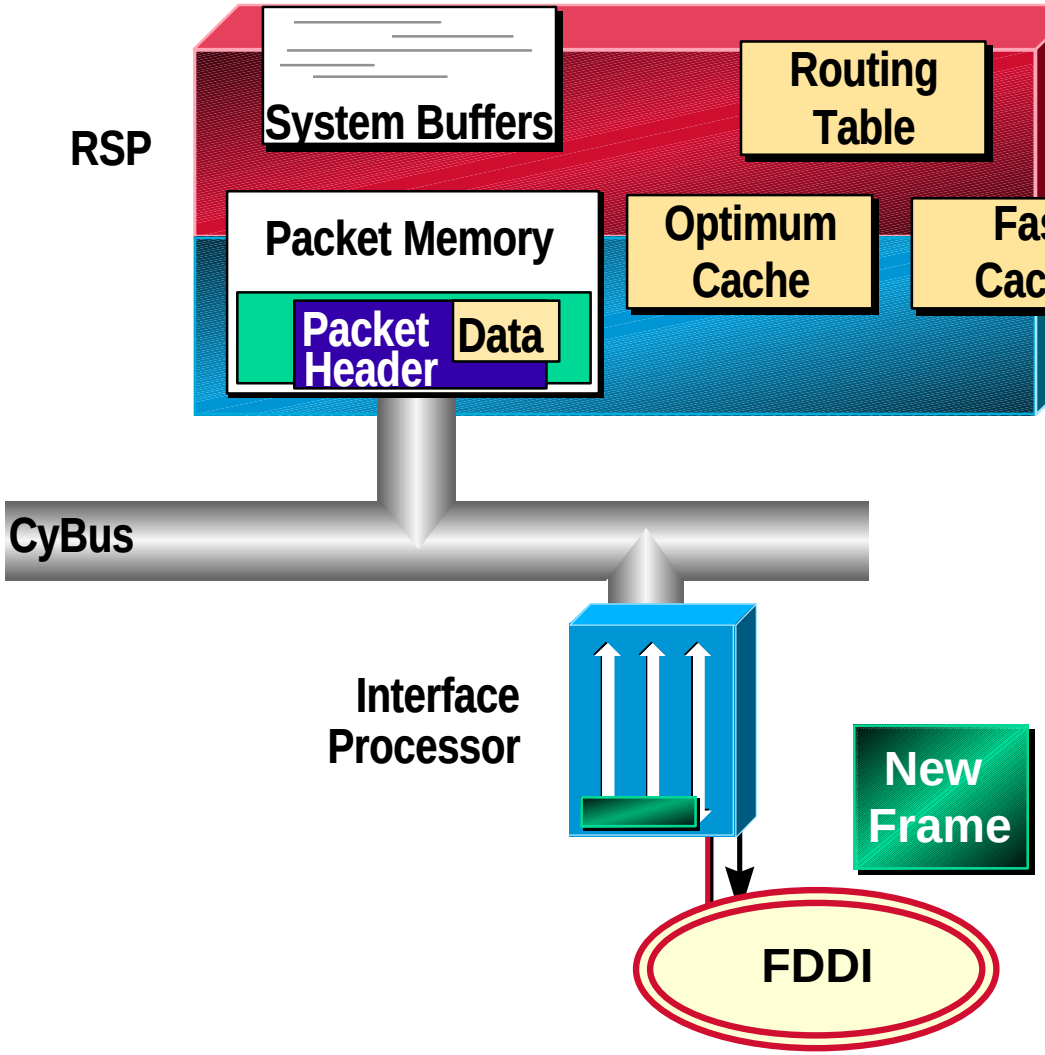

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco Systems

Definition of NetFlow

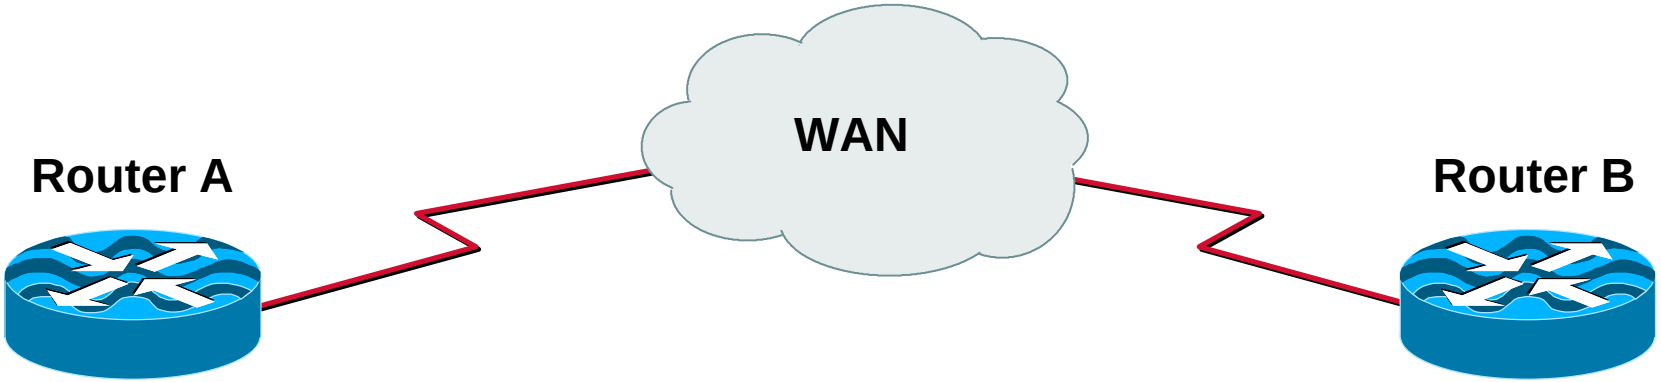

- A unidirectional sequence of packets between a given source and destination

Characterizing NetFlows

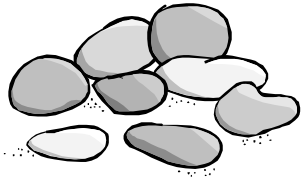

Granularity

- What do I use to define my NetFlow?

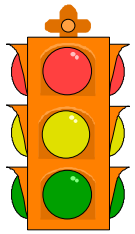

Starting and Stopping

- What determines the start of my NetFlow?
- What determines the end of my NetFlow?

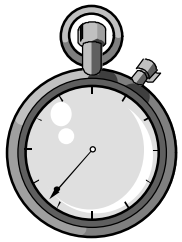

Timeout

- When do I get rid of my NetFlow entry?

CISCO SYSTEMS

NetFlow Granularity
Data Encapsulation

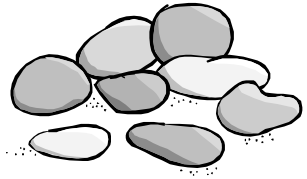

Application Layer
(Telnet, FTP, etc.)

DATA

Transport Layer
(TCP, UDP)

TCP/UDP Header | DATA

Network Layer
(IP)

IP Header | TCP/UDP Header | DATA

Data Link Layer
(Ethernet, TR, etc.)

Ethernet Header | IP Header | TCP/UDP Header | DATA

IP Header

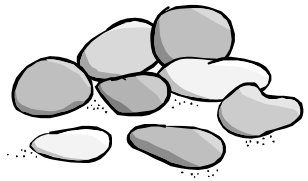

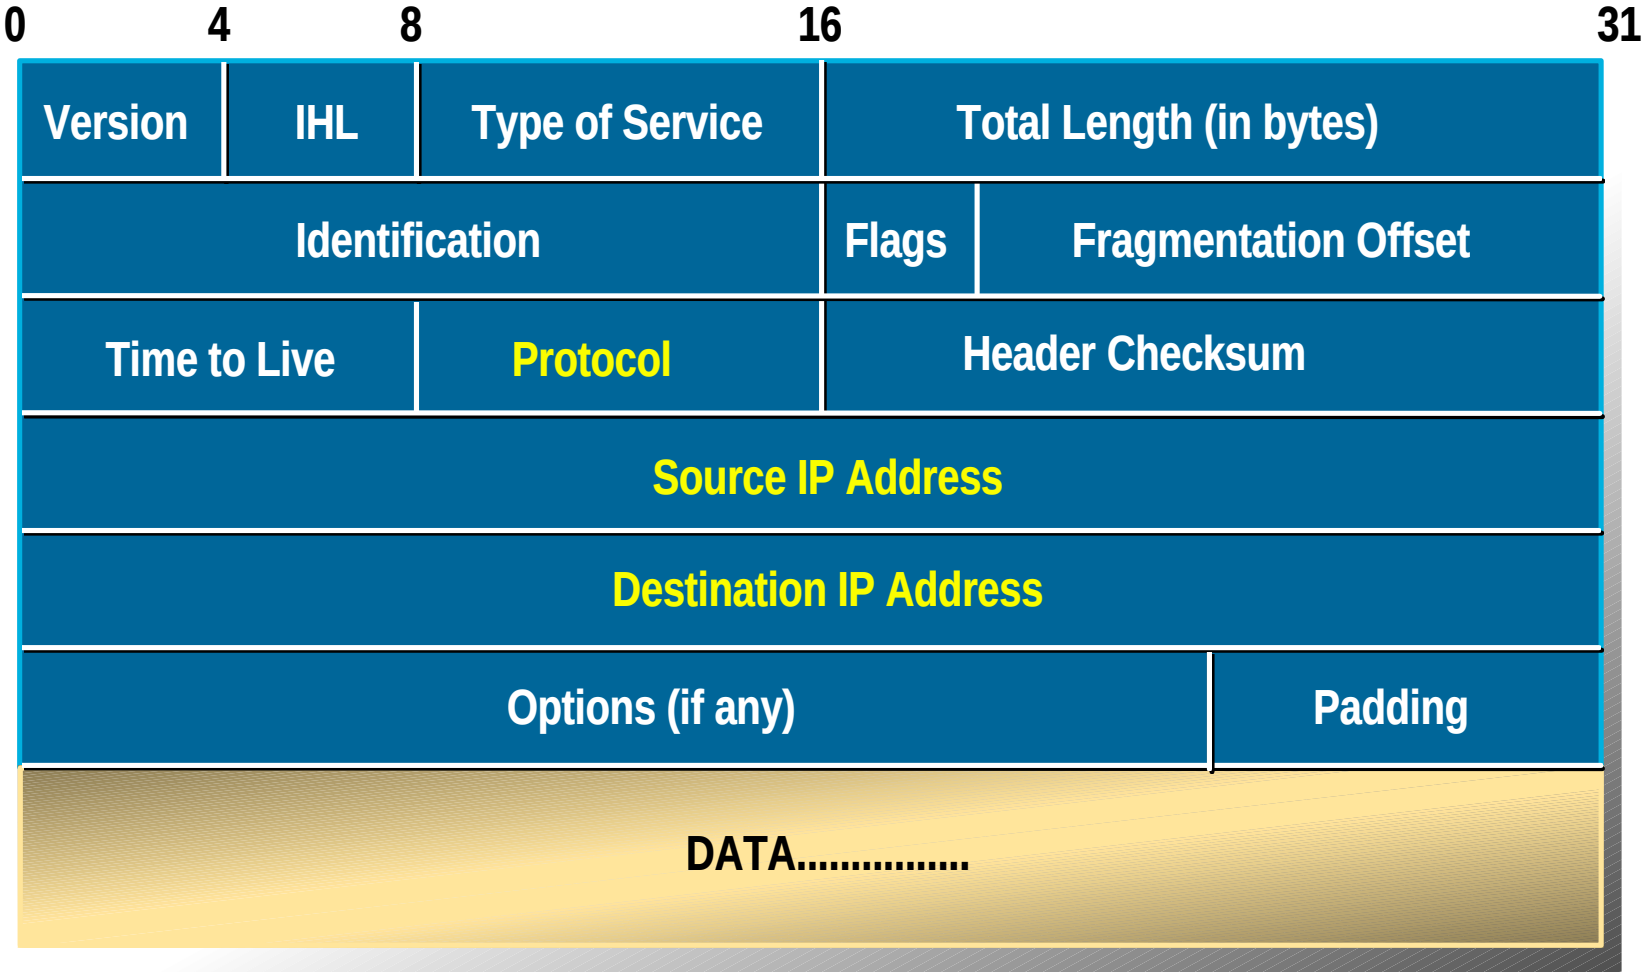
0
4
8
16
31
Version
IHL
Type of Service
Total Length (in bytes)
Identification
Flags
Fragmentation Offset
Time to Live
Protocol
Header Checksum
Source IP Address
Destination IP Address
Options (if any)
Padding
DATA.....

IP Protocol Numbers

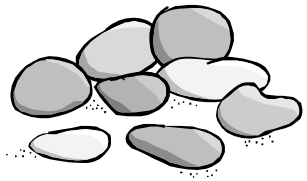

Number	Protocol
1	ICMP—Internet Control Message Protocol
2	IGMP—Internet Group Message Protocol
6	TCP—Transmission Control protocol
7	UDP—User Datagram Protocol

UDP Header

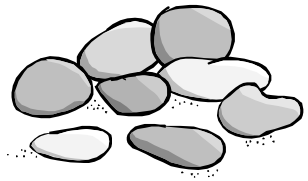

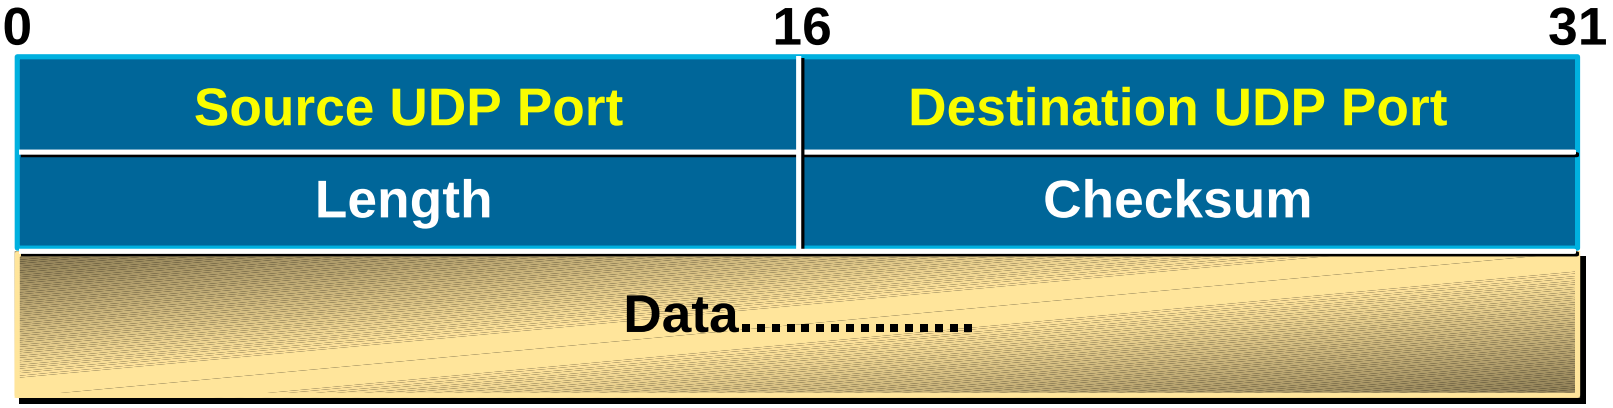

UDP Port Numbers (0-65535)

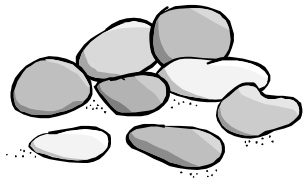

Port Number	Protocol
53	DNS—Domain Name Service
520	RIP—Routing Information Protocol
161	SNMP—Simple Network Management Protocol
69	TFTP—Trivial File Transfer Protocol

TCP Header

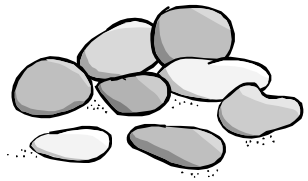

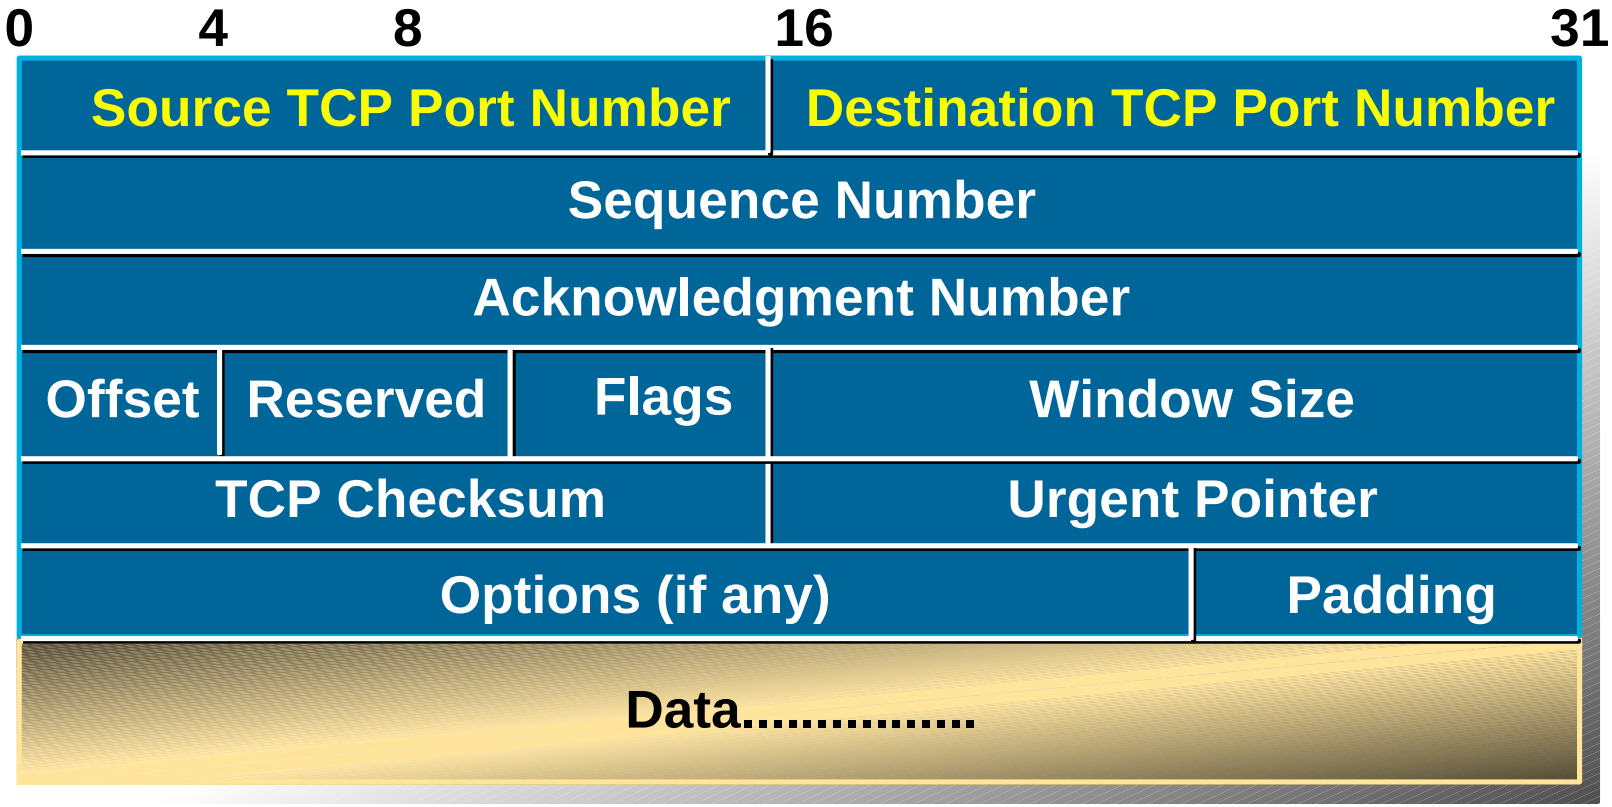

TCP Port Numbers (0-65535)

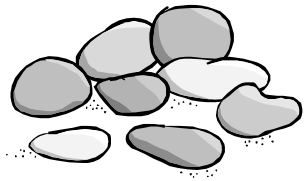

Port Number	Protocol
23	Telnet—Remote Login
21	FTP—File Transfer Protocol
25	SMTP—Simple Mail Transfer Protocol
80	WWW—World Wide Web
179	BGP—Border Gateway Protocol

NetFlow Start/Stop

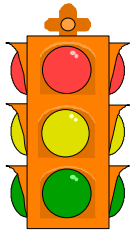

- **Start of a NetFlow**
  - **New NetFlow entry**
- **Stop of a NetFlow**
  - **Protocol flags (TCP FIN)**
  - **Cache entry age**

TCP Flags: Initialization

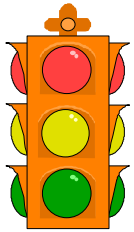

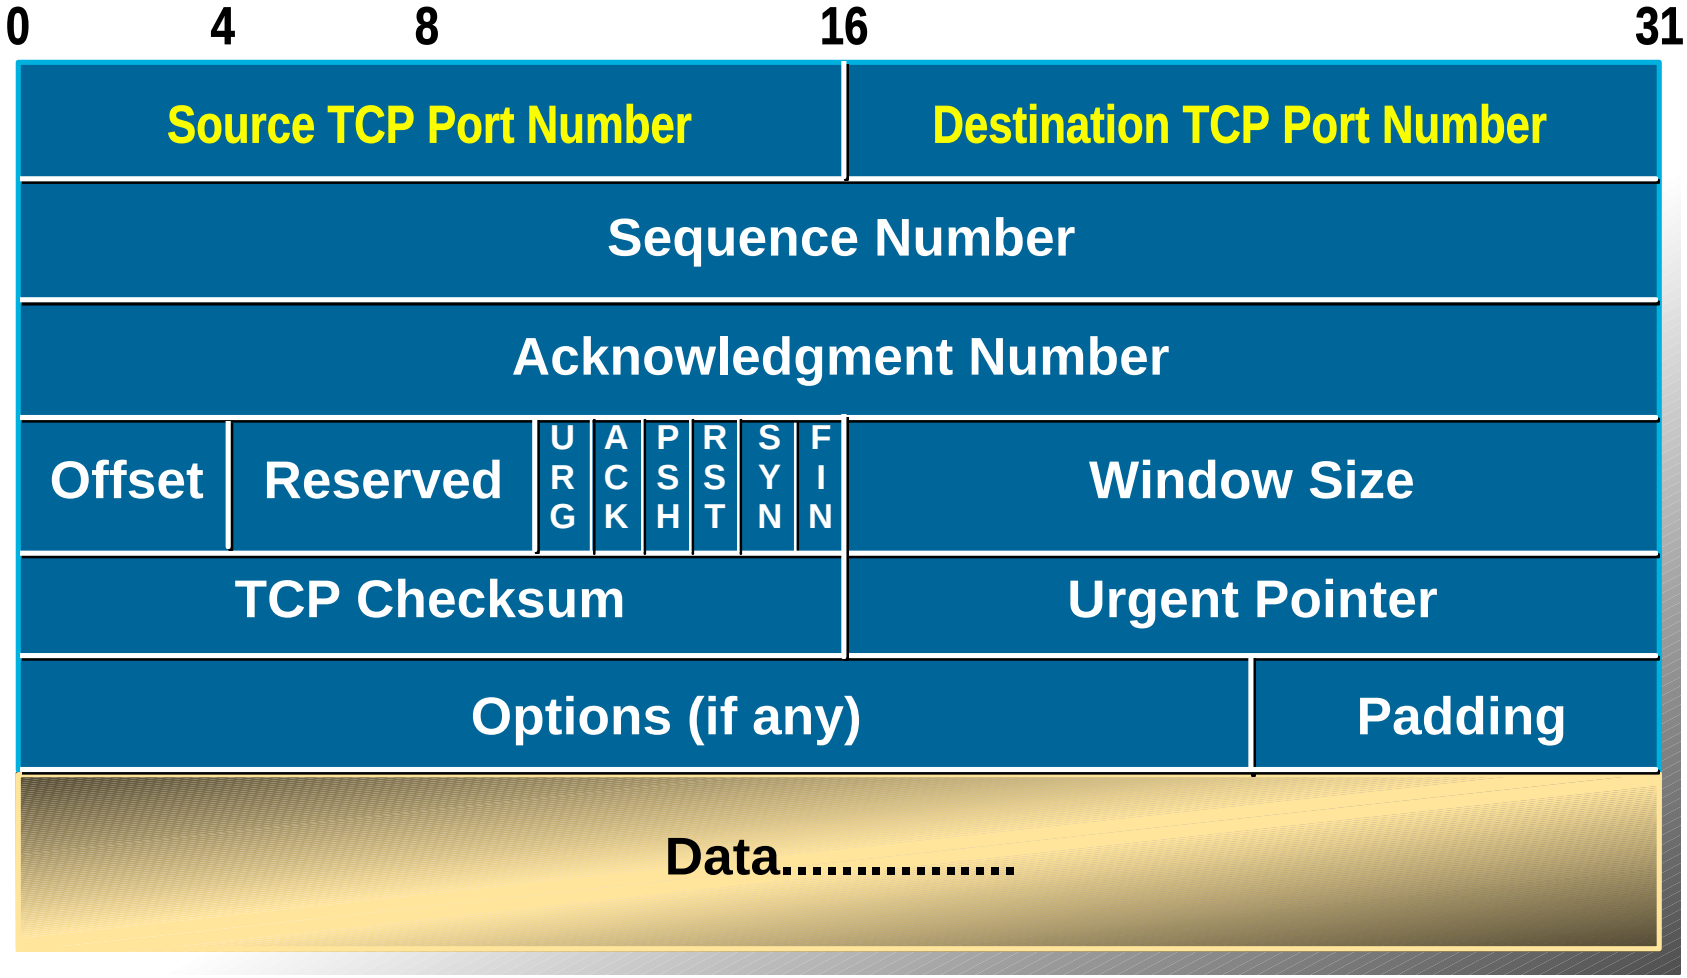

TCP Flags: Initialization

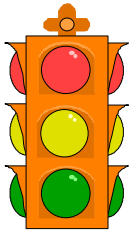

Bob

Jane

SYN
(Can I talk to you?)

New Cache Entry (Bob-Jane)

SYN, ACK
(Sure, what's up?)

New Cache Entry (Jane-Bob)

ACK, Data
(Cool...have you heard blah blah..)

Cisco Systems

New NetFlow

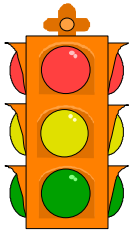

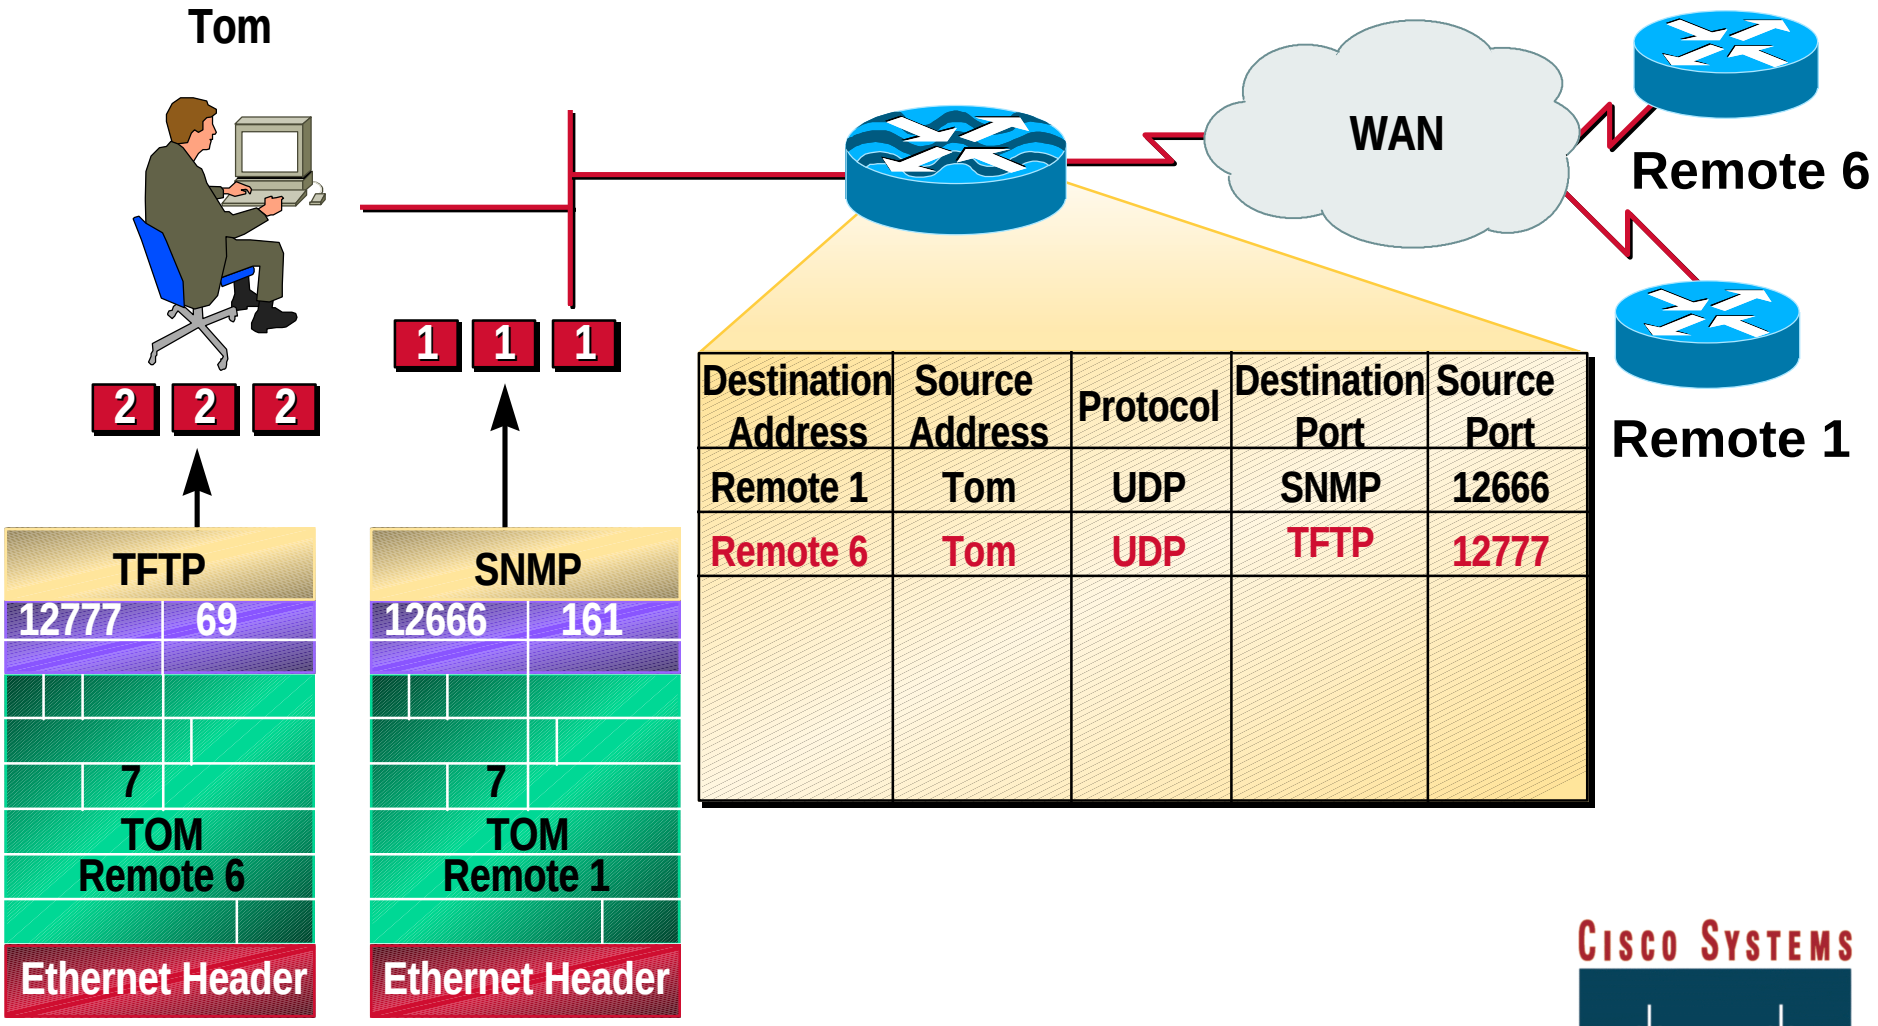

Destination Address	Source Address	Protocol	Destination Port	Source Port
Remote 1	Tom	UDP	SNMP	12666
Remote 6	Tom	UDP	TFTP	12777

CISCO SYSTEMS

TCP Flags: Termination

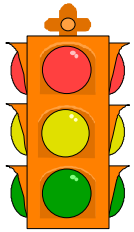

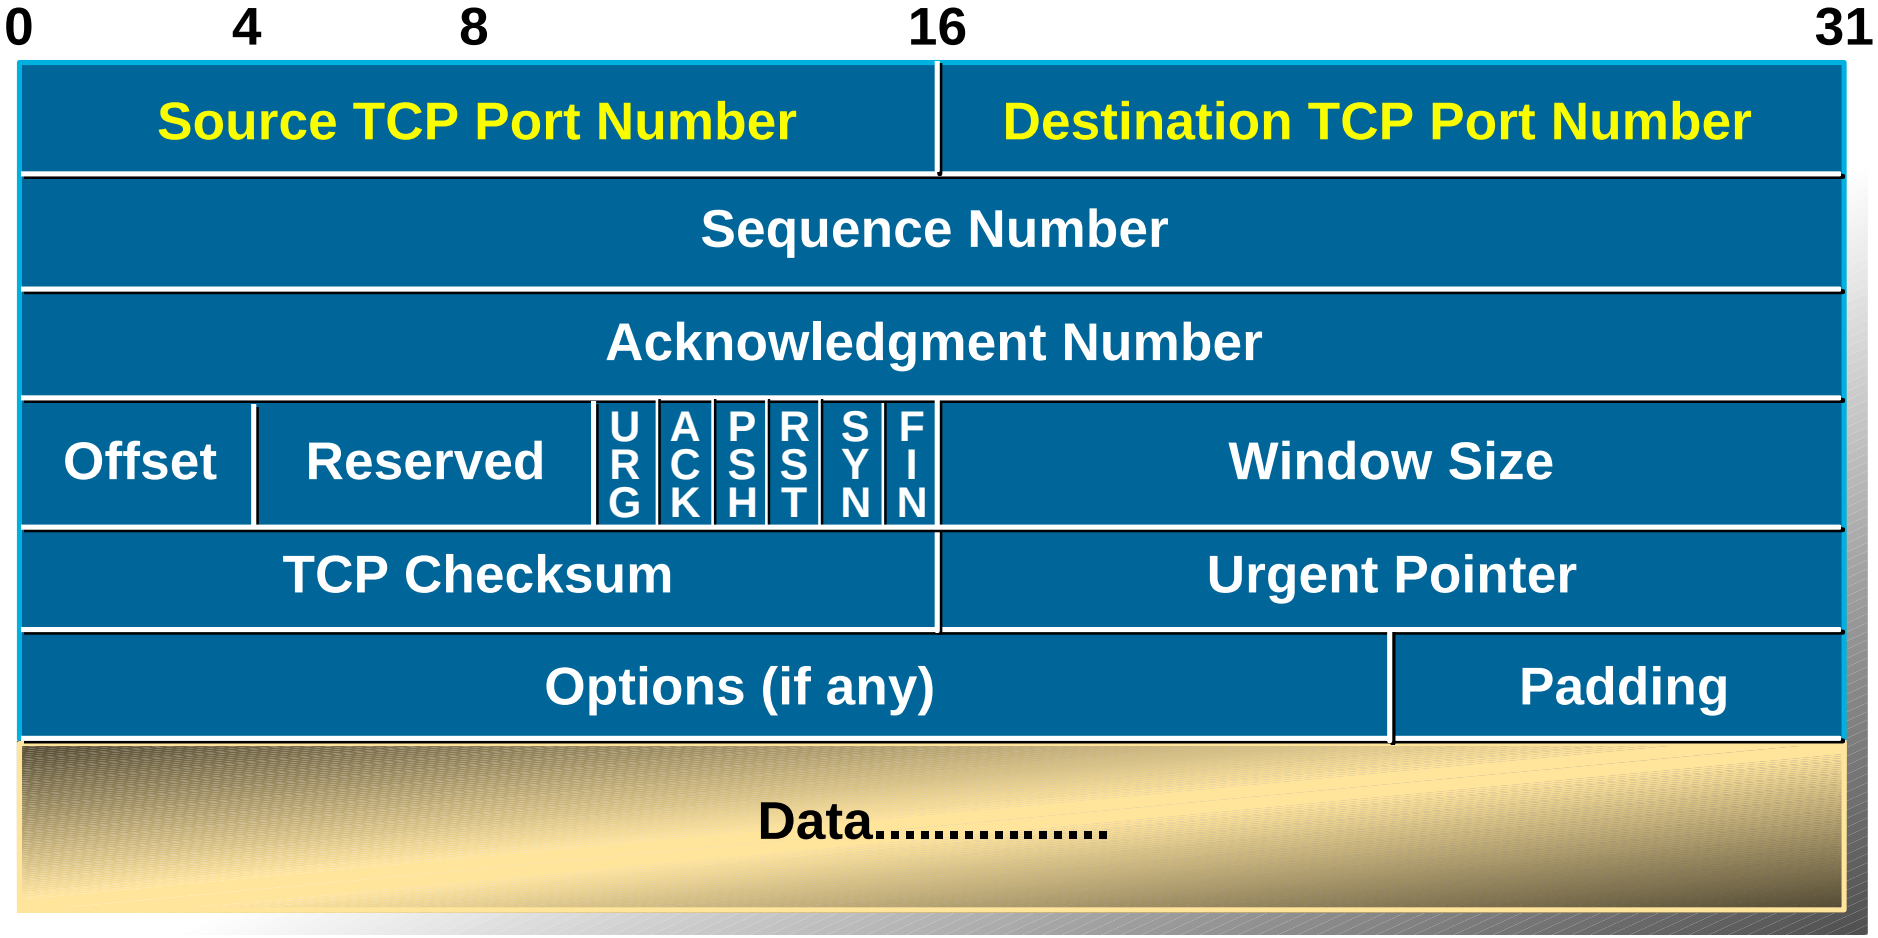

TCP Flags: Termination

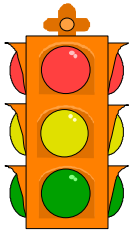

Bob

Jane

FIN
(That's all for now, Bye)

Clear Cache Entry (Bob-Jane)

FIN, ACK
(OK, Bye)

Clear Cache Entry (Jane-Bob)

(FIN)
(Connection Closed)

Cisco Systems

Inter-Frame Gap

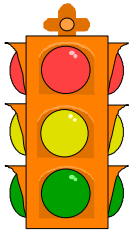

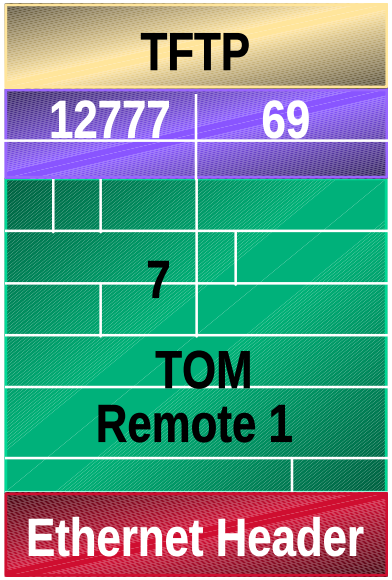

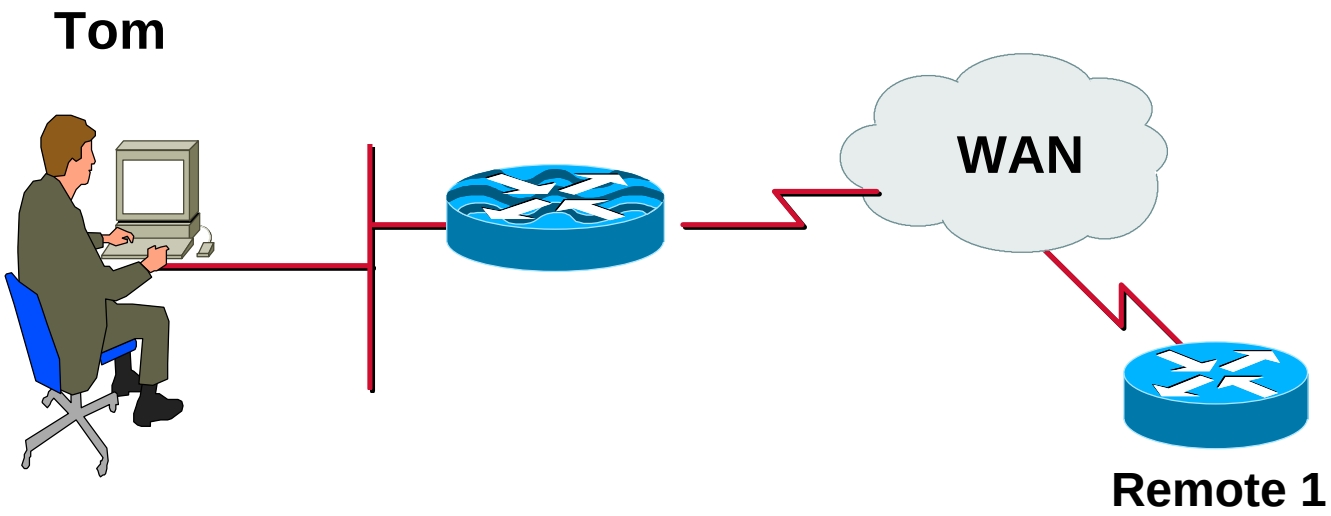

Too Long

OK

NEW NEW NEW 1 1 1 1 1

CISCO SYSTEMS

Cisco 7500 Series—NetFlow Switching

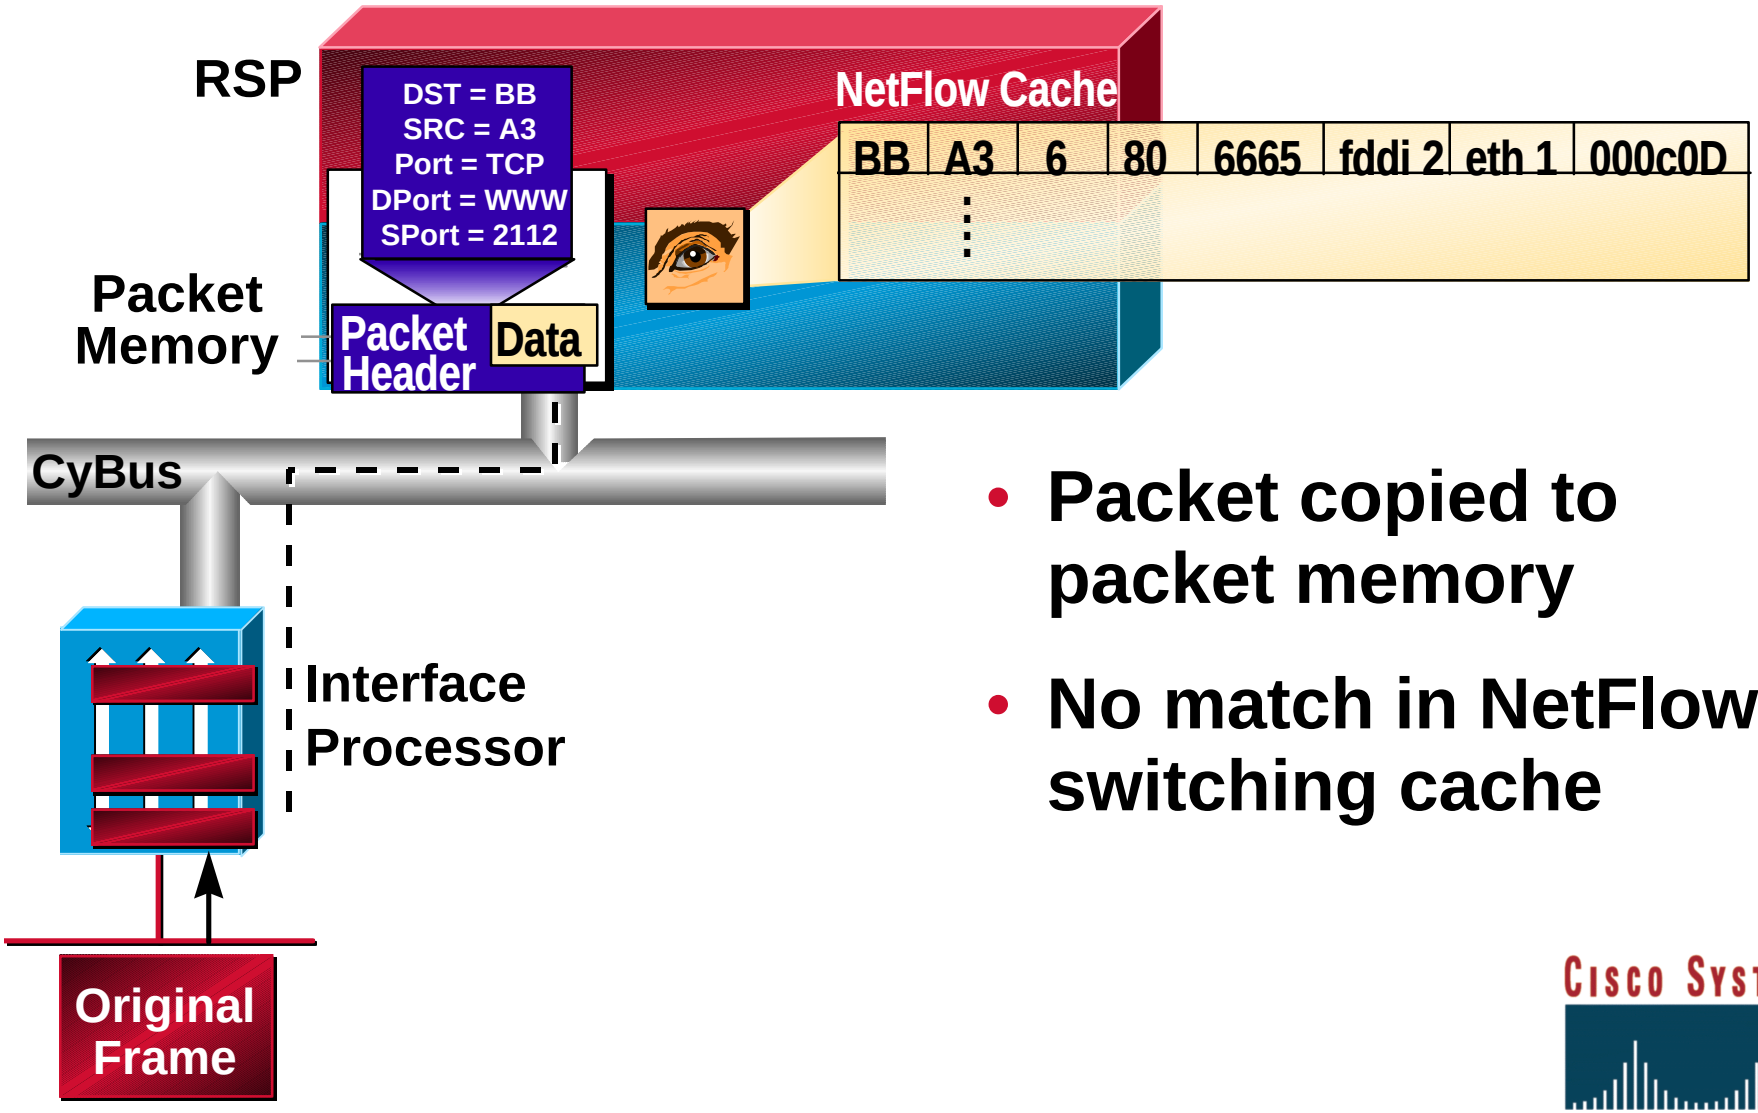

- Packet copied to packet memory
- No match in NetFlow switching cache

Cisco Systems

Cisco 7500 Series—NetFlow Switching

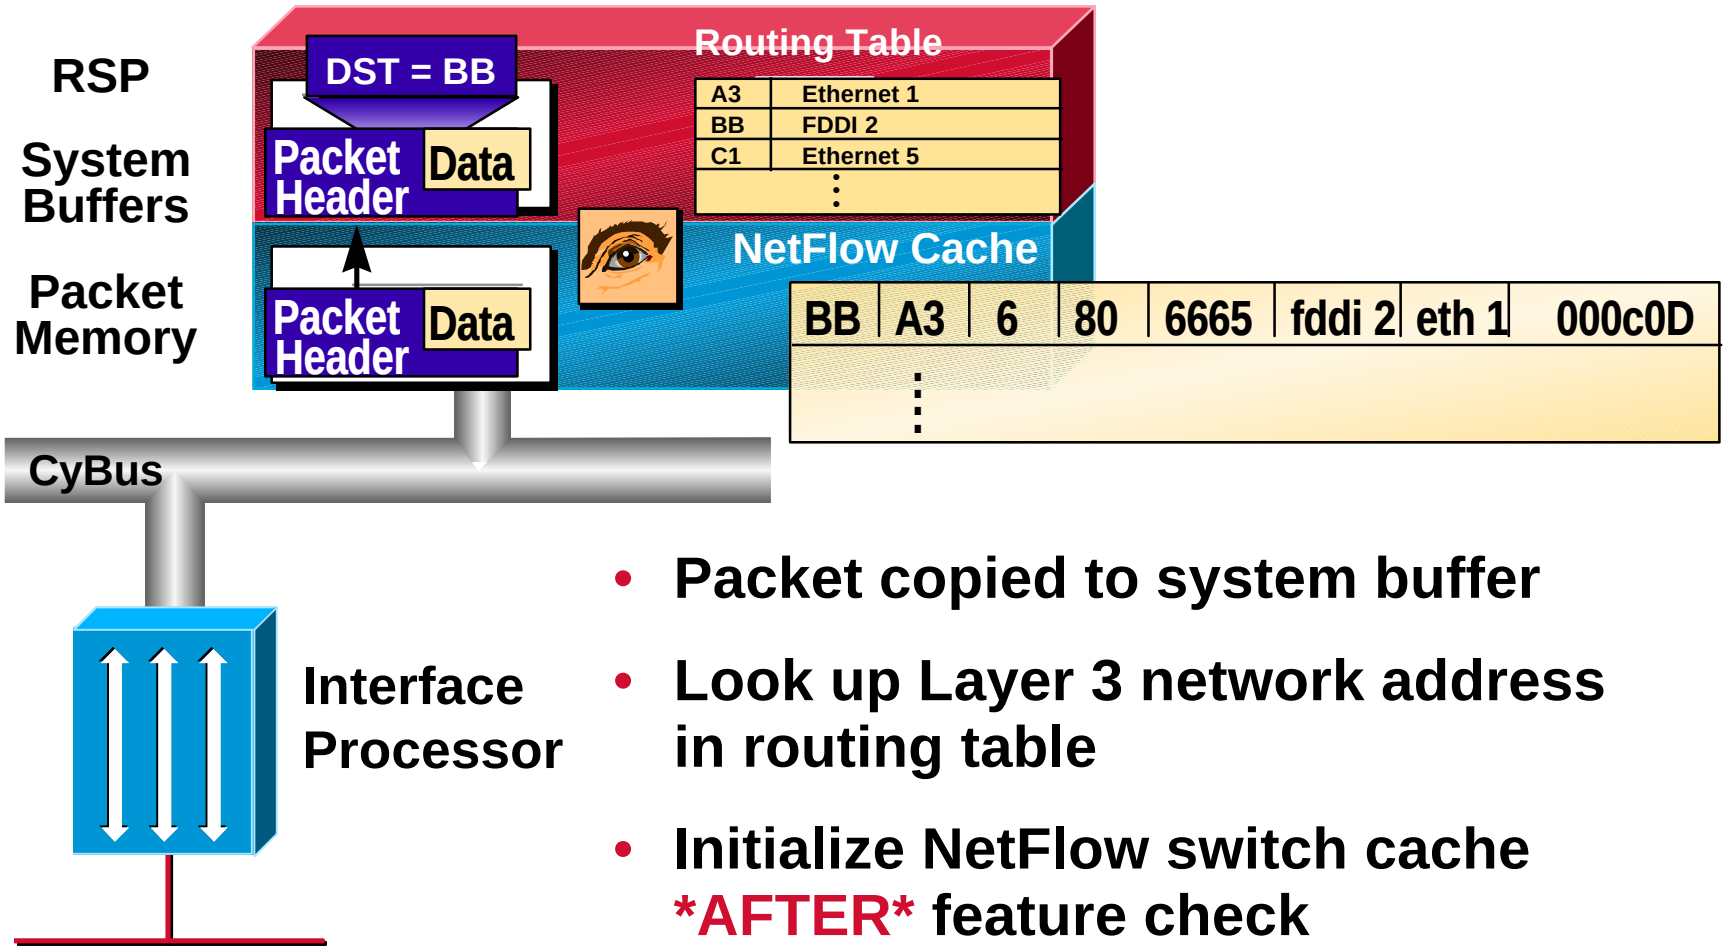

- Packet copied to system buffer
- Look up Layer 3 network address in routing table
- Initialize NetFlow switch cache *AFTER* feature check

Cisco Systems

NetFlow Switching

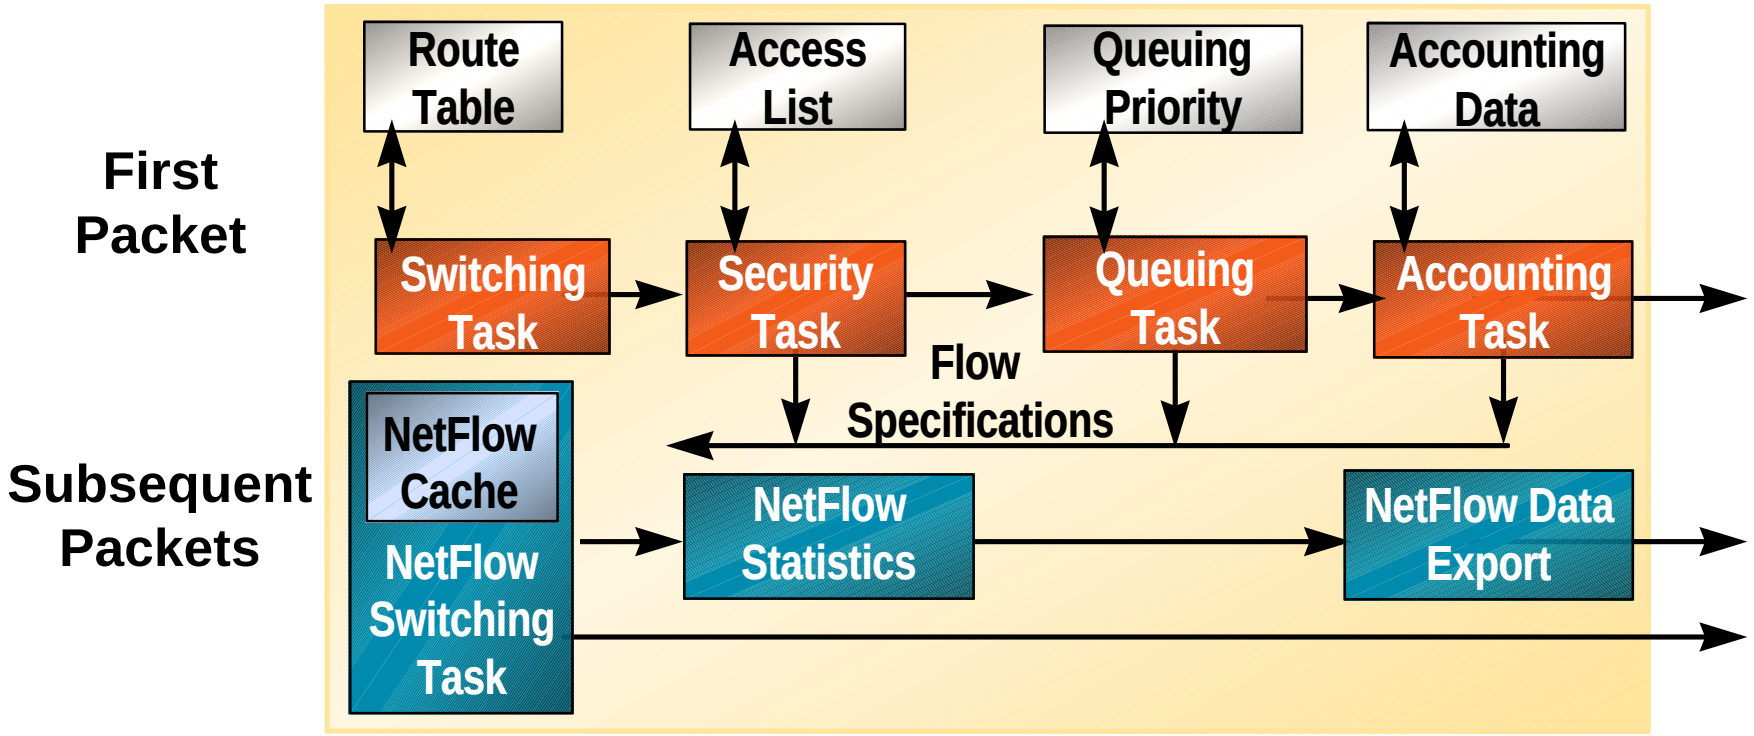

- Only first packet is processed by multiple tasks
- Connection-oriented NetFlow is defined with specific service requirements based on source/destination network address and transport layer port numbers
- Single switching task applies network services and collects traffic statistics

Cisco Systems

Cisco 7500 Series—NetFlow Switching

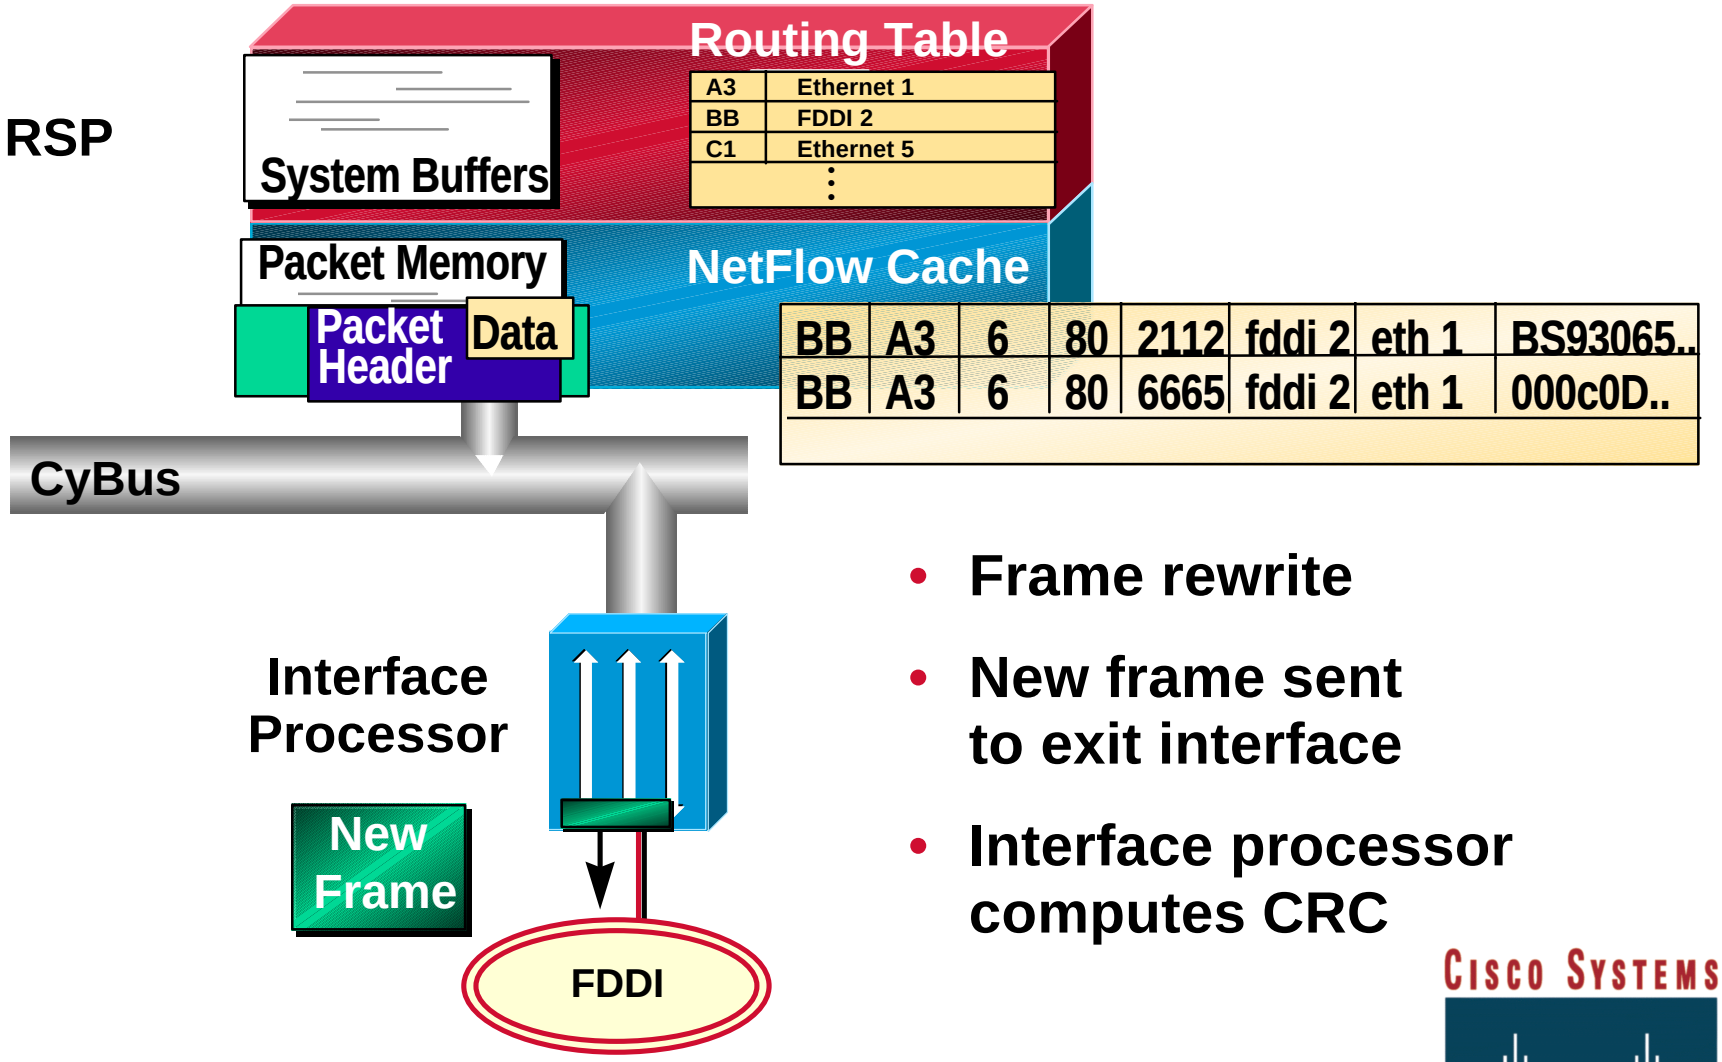

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

CISCO SYSTEMS

Cisco 7500 Series—NetFlow Switching

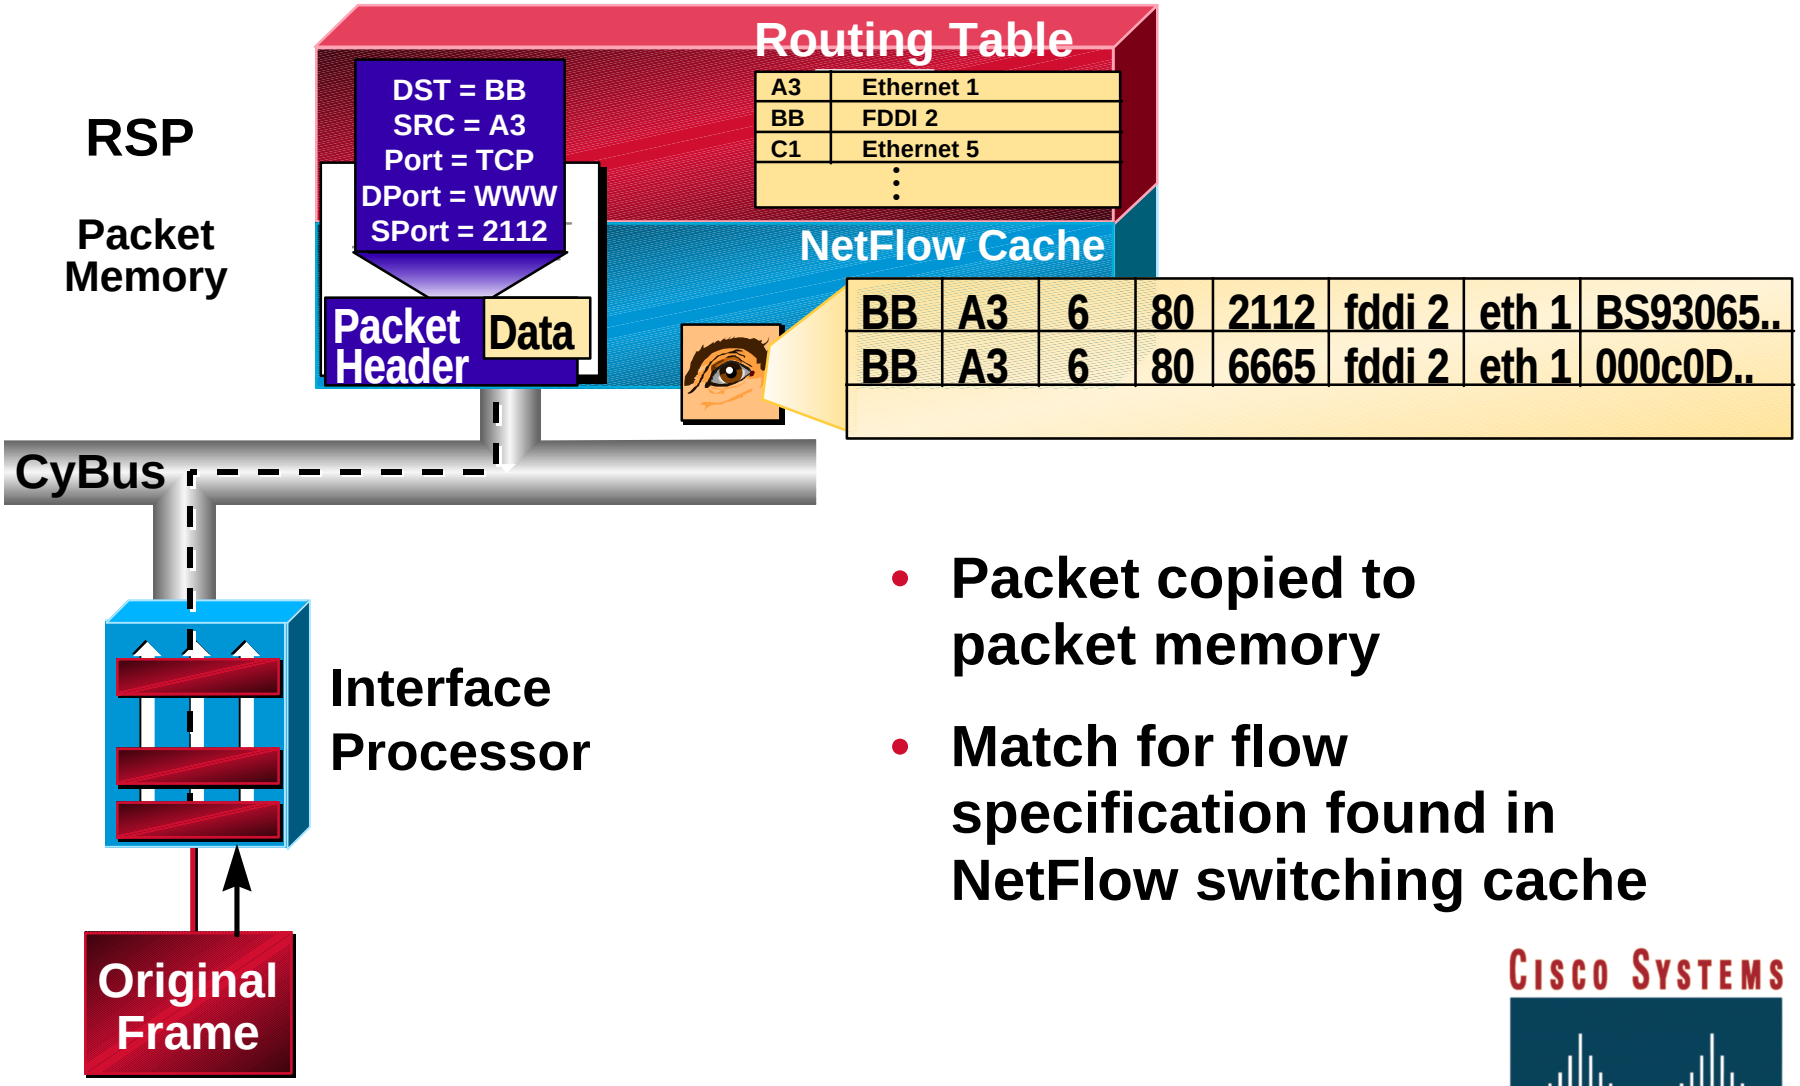

- Packet copied to packet memory
- Match for flow specification found in NetFlow switching cache

Cisco Systems

Cisco 7500 Series—NetFlow Switching

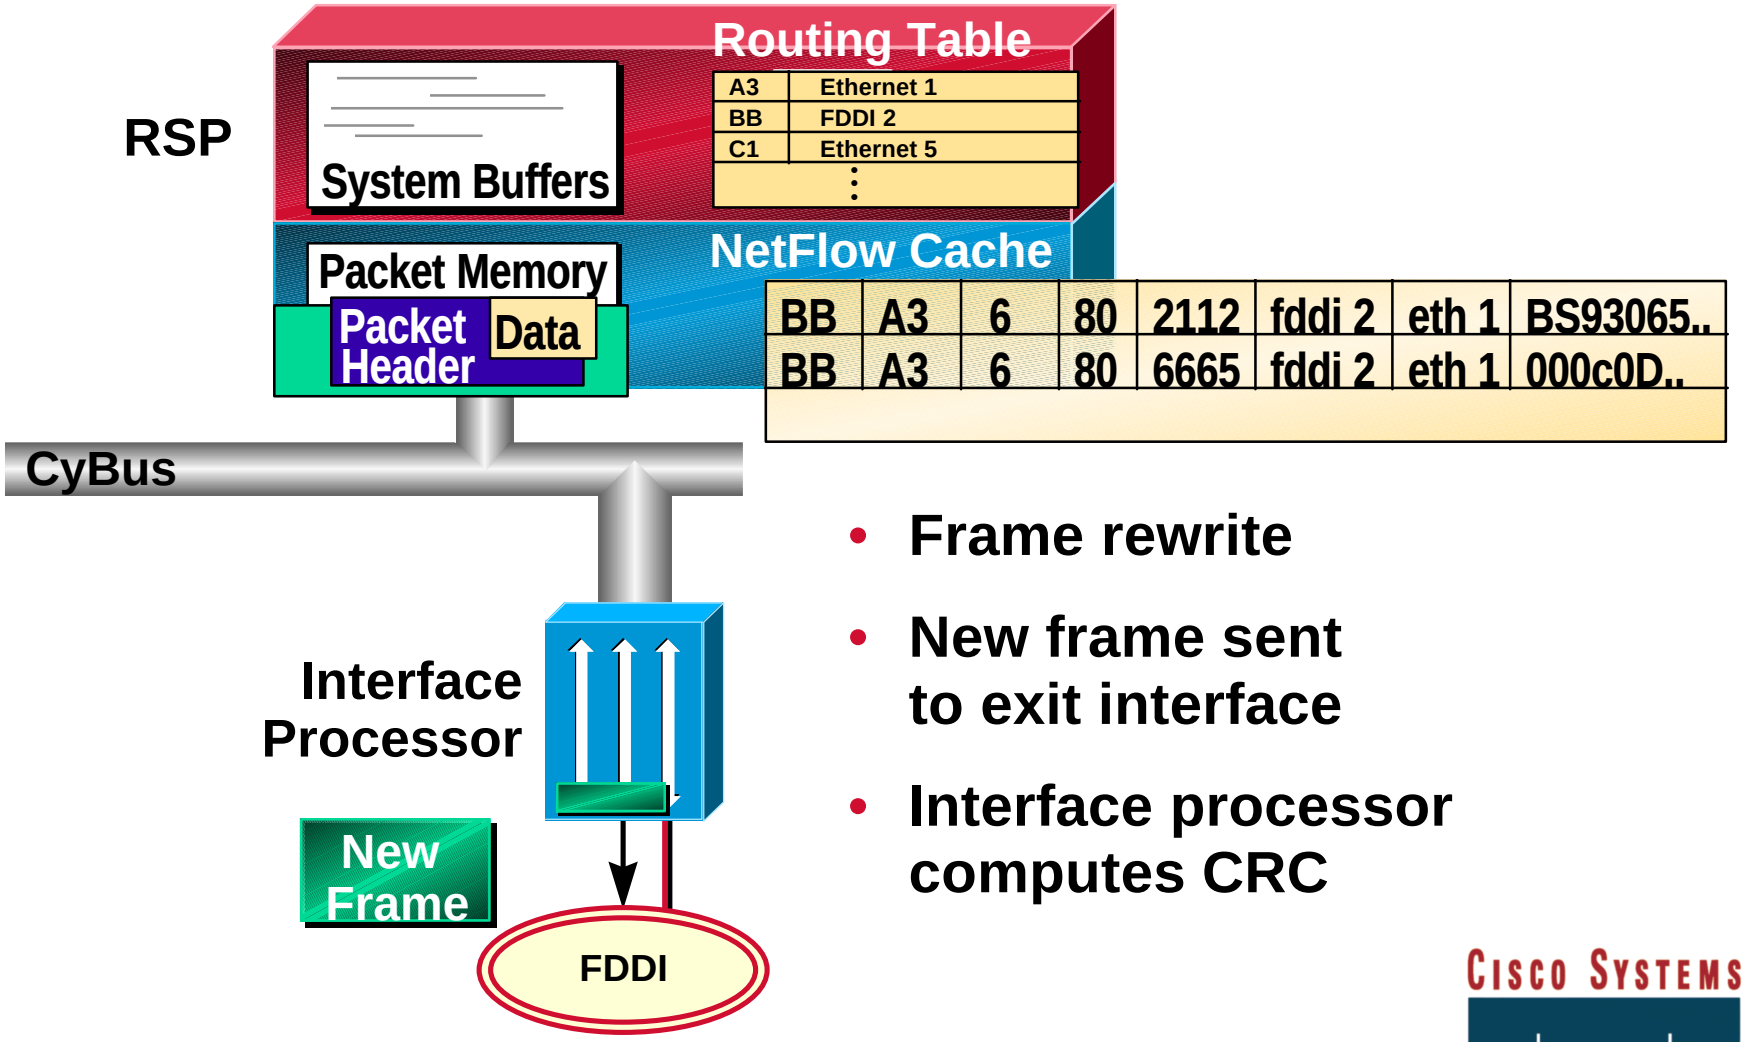

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco Systems

NetFlow Switching—Statistics

- **Detailed traffic statistics on each flow:**
  - **Input/output interface port on router**
  - **Source/destination address**
  - **Source/destination transport port number**
    - **Linked to applications**
  - **Packets and bytes/packet**
  - **Active sec./flow**
  - **CoS precedence per flow**
    - **Packets per CoS precedence**
  - **Encryption (crypto map)**

NetFlow Statistics
IP NetFlow Switching Cache, 29999 active, 2769 inactive, 58411388 added
statistics cleared 141949 seconds ago

	Total Flows	Flows / Sec.	Packets	Bytes	Packets / Sec.		
TCP -	267,034	1.8	233		439.3	182.6	36.5
FTP	1,030,837	7.2	10	78	76.6	22.6	43.7
FTPD	554,967	3.9	164	345	641.3	52.7	15.7
WWW		226.2		247			
	3,526,231	24.8	13	159	323.1		
X	9,600	0.0	121	129	8.2	148.2	
BGP	111,096	0.7			11.5	229.2	61.1
other	5,729,172	40.3	70	220	2858.1		
UDP -	2,398	0.0	3	62	0.0	13.4	69.5
DNS			2	110	195.4	5.4	
other	1,489,072	10.4	30	293	321.8	28.5	68.7
ICMP	665,771	4.6	13	259	62.8	75.7	66.8
IGMP	5,144	0.0	18	278	0.6	82.4	64.3
IPINIP	4,450	0.0	933	377	29.2	166.7	61.0
IP - other	2,693	0.0	11	136	0.2	80.8	65.7
TOTAL		411.3		227		0.0	0.0

Src Intf	Source IP Address	Dest Intf	Dest IP Address	Prt	Src		Bytes / Pkt	Act Sec / Flow	Idle Sec / Flow
Hs3/0	204.119.134.49	Fd0/0	142.35.4.36	6	0050	0610	1	0.0	0.6
Fd0/0		Hs3/0	206.52.126.29	6			105	9.1	1.0
Hs3/0	125.160.1.24	Fd0/0	200.246.225.8	6			745	542	323.0	0.0
...	...	...	...	...	...	...	...	...	...	...

CISCO SYSTEMS

NetFlow Data Collection

- **RMON TrafficProbes:**
  - **Collect and store data from local devices**
  - **For analysis by RMON applications**
- **Cisco data collectors:**
  - **Convert individual flow records to aggregated flow data**
  - **By traffic source/destination:**
    - **Host/subnet, router or interface**
  - **By traffic type:**
    - **By protocol, QoS**
  - **By configurable time period**

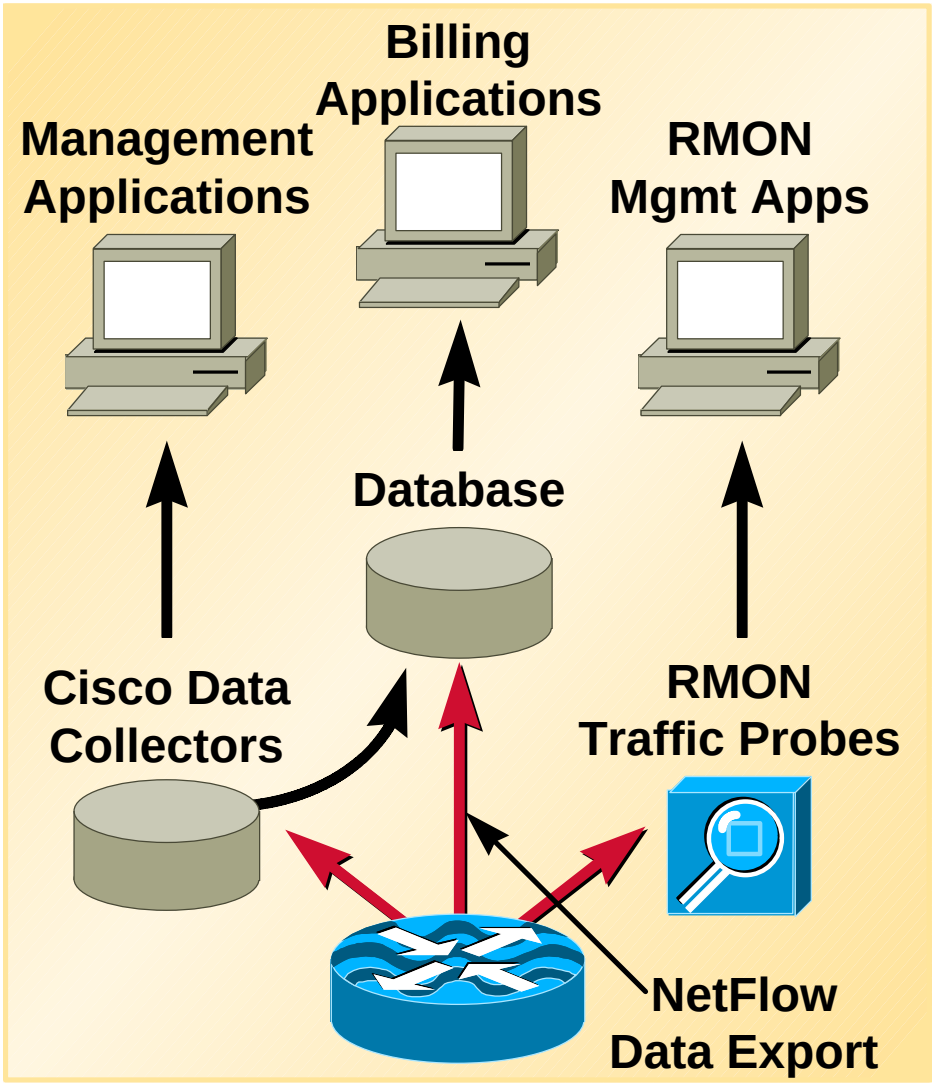

CEF Drivers

The Current Demand based Cache systems have the following characteristics:

- They are traffic driven—i.e. wait for first packet to get full information needed to forward
- The caches can get larger than routing table and consume lots of memory
- Periodic aging consumes significant cpu if cache is large
- Cache invalidation after a route flap causes process switching to repopulate cache
- Time is required to distribute cache to VIPs, packets must be RSP switched until caches-VIP synchronisation
- The inability to do per packet load sharing from interrupt level

CEF Concepts

The following describe the functions involved in CEF

- **Build adjacencies with neighbors**
- **Pre-calculate all information needed to forward packets to VIP's**
- **Create Forwarding Information Base (FIB)**
- **Distribute full forwarding information to line cards**
- **Preclude RSP from switching packets**
- **Reduce aggregate system memory utilization**

CEF Structures and Function

Adjacencies

Nodes are adjacent if they can be reached by one layer 2 hop

The adjacency table is populated by routing protocols and ARP since we hear updates from our neighbors

Types of adjacency

Normal—connected routers
Null—for packets that we want to drop
Glean—for connected interfaces and the hosts on them
Punt—packets that we must send to RSP—Punt order is DFIB —> FIB —> Fast —> Process Drop—drop packet but do accounting
Incomplete—mac rewrite not available

CEF Structures and Function

Adjacency table

A table indexed by layer 3 address (populated when entries are added to ARP table containing):

- **New MAC rewrite string**
- **Connecting Interface**
- **MTU**
- **Counters**

FIB Table

The following describes the contents and structure of the FIB

- mtrie indexed by IP address
- Indexed with 8 bits at each level for 4 level max depth with 32 bit address
- Entries point to adjacency table
- Allows overlapping entries (ie 171.68.118.0 /24 and 171.68.0.0 /16)

Cisco Express Forwarding: Routing Table

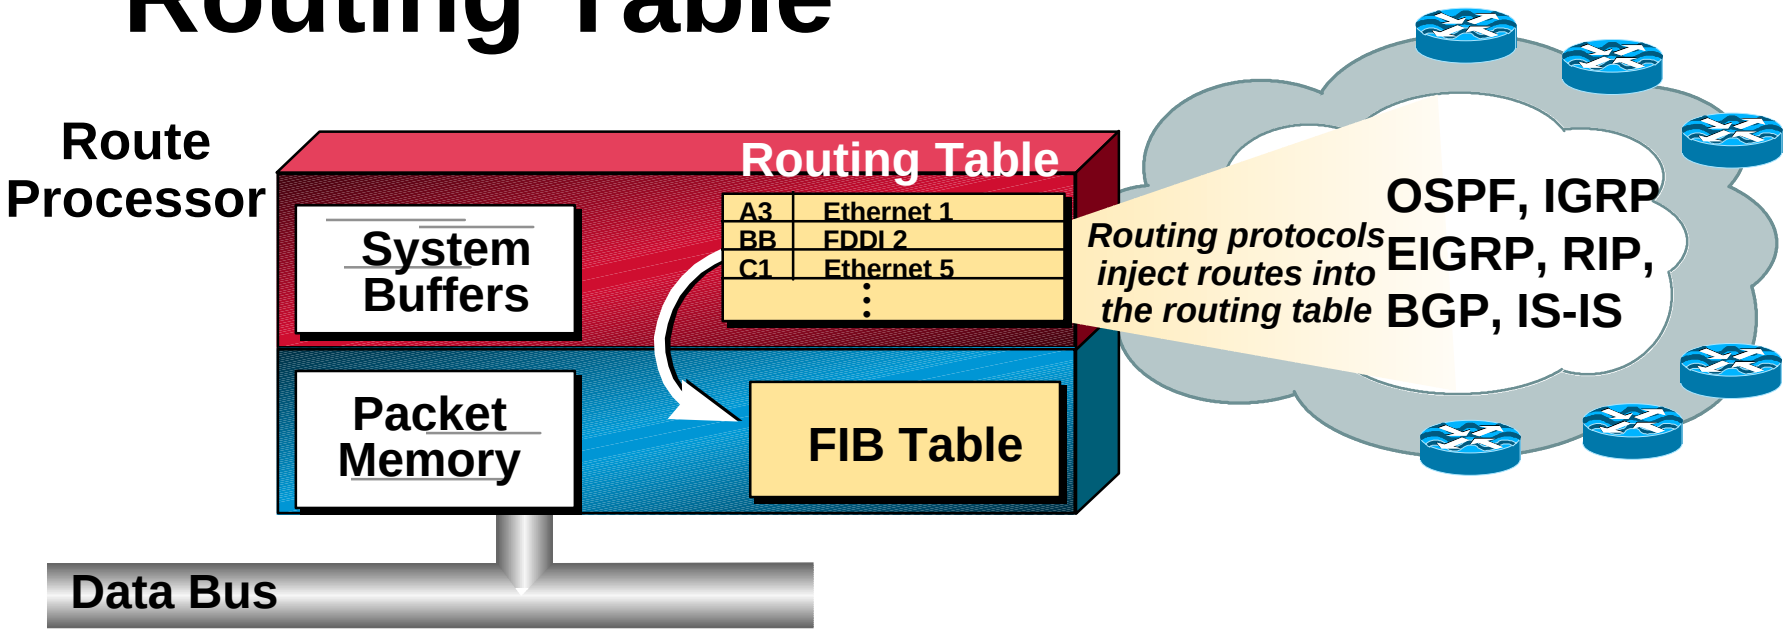

- The routing table contains a list of reachable networks
- The routing information is passed down to the Forwarding Information Base (FIB) manipulated, into an optimized software look-up tree (mtrie)

Cisco Systems

Cisco Express Forwarding: Forwarding Table

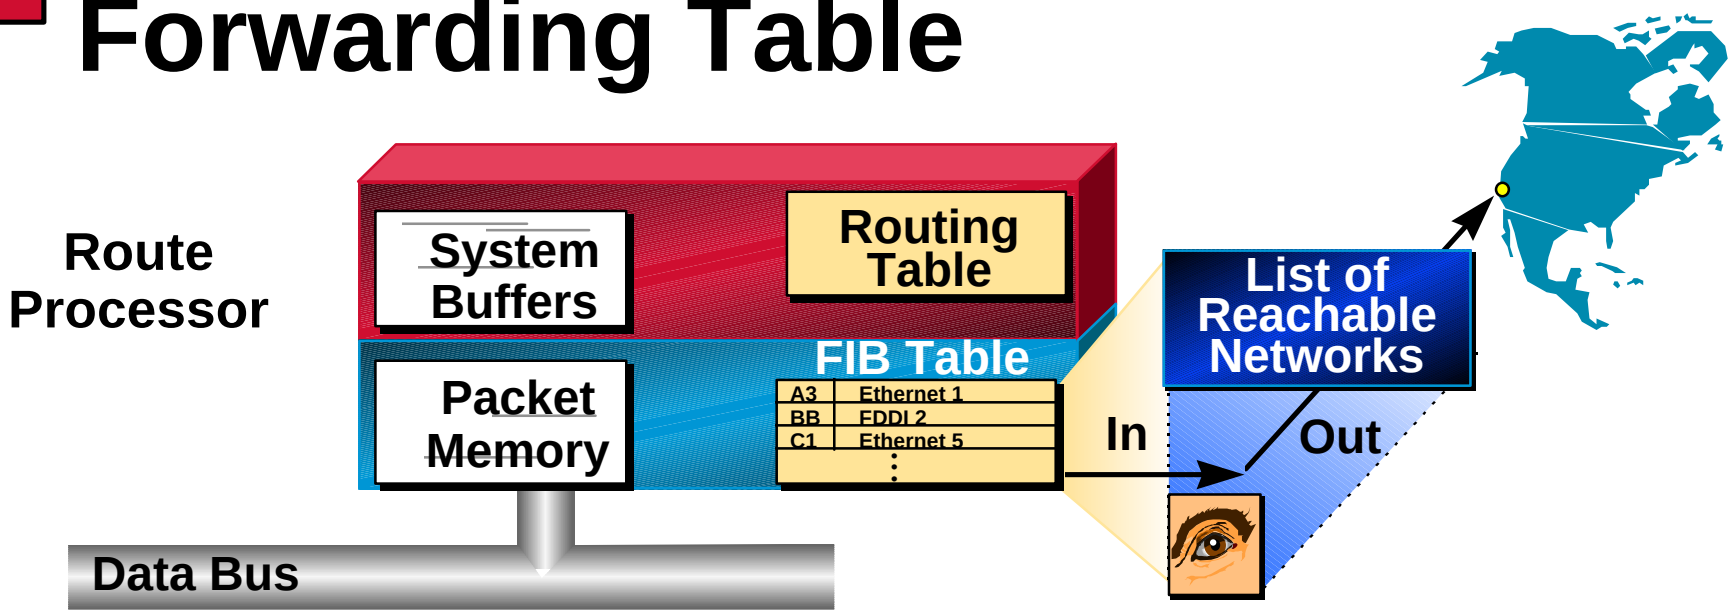

- **The FIB is a database of information which is used when making a packet-forwarding decision**
  - **Updated when a route is added, removed, or changed in the routing table**
  - **The FIB s/w tree contains a list of destination IP prefixes**
  - **Pointer to the adjacency database**

Cisco Systems

Cisco Express Forwarding: Adjacency Table

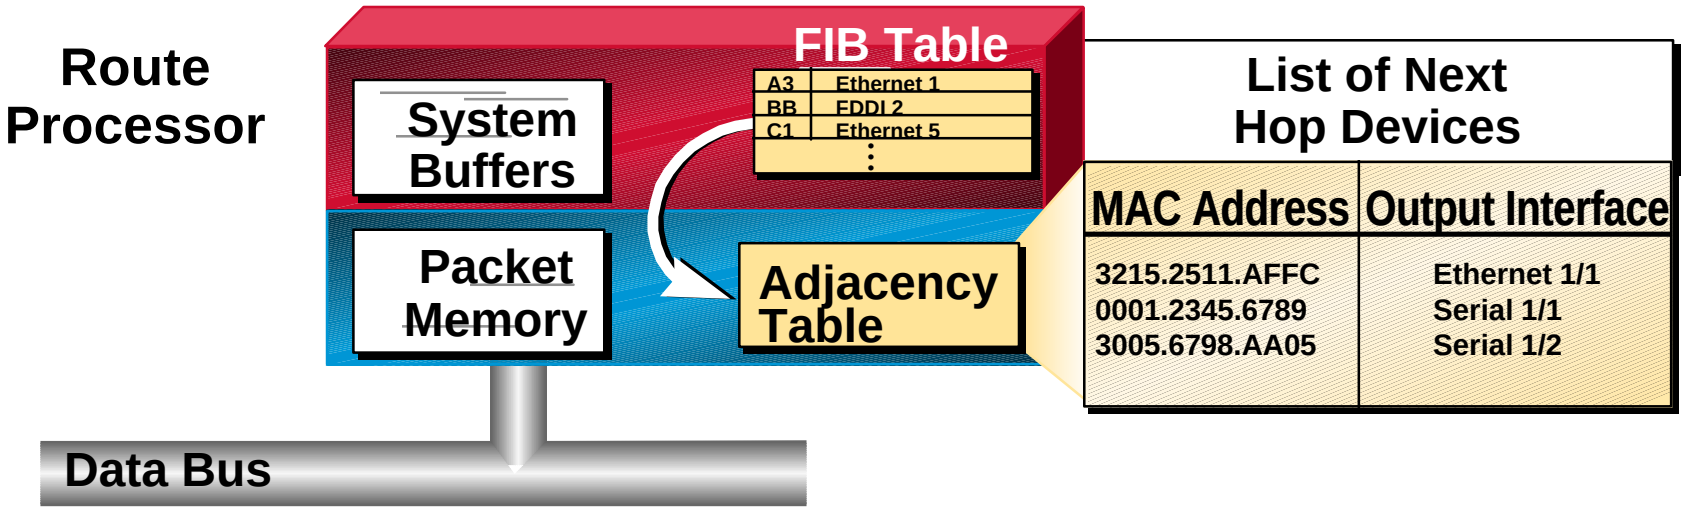

- A list of adjacent nodes and their status up/down
- Adjacency table is populated as it discovers adjacent nodes

  On the creation of an adjacency, a link-layer header to reach the node is computed and stored in the database

Cisco 7500 Series Distributed Switching VIP Distributed/Local Switching

- Distributed ‘fast’ and NetFlow switching for IP traffic
- Locally switched traffic does not cross system bus
- System performance is optimized with both xIPs and VIPs

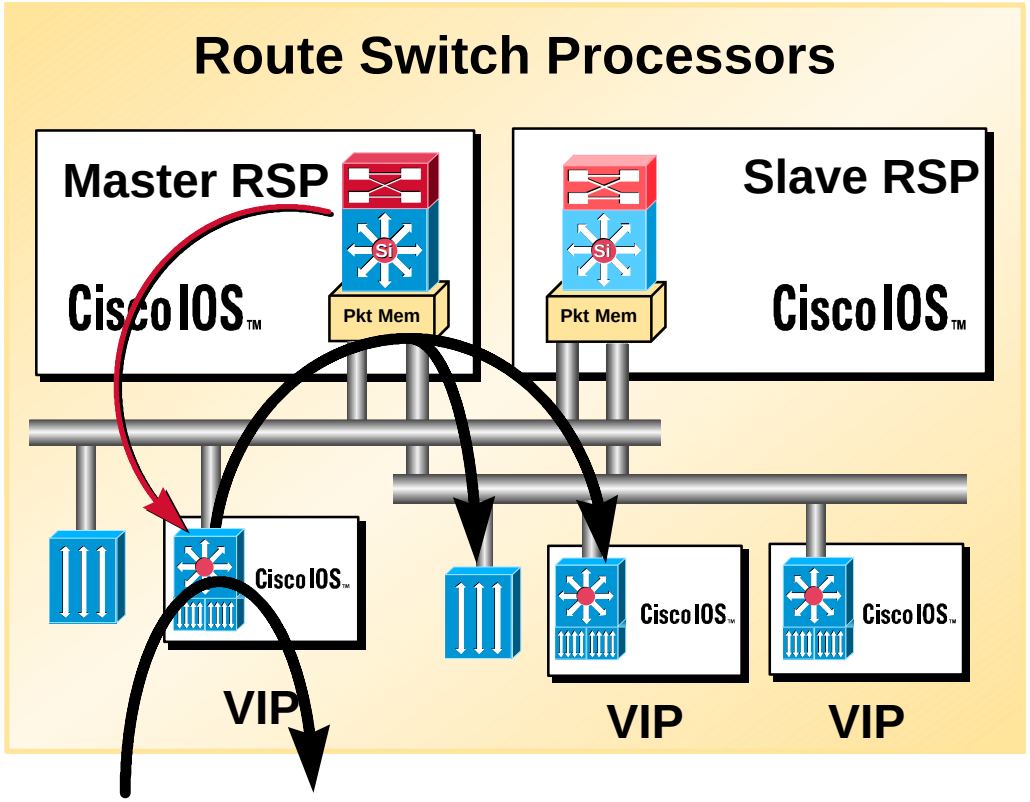

Cisco Systems

Cisco RSP-Based Router Aggregate Switching Summary

Switching Path	Cisco 7000RSP	Cisco 7200	Cisco 7500	Configuration Command
Process Switching	8,000	10,000	10,000	no <protocol> route-cache
Fast Switching	70,000	120,000	150,000	<protocol> route-cache
Optimum Switching	120,000	150,000	275,000	<protocol> route-cache optimum
NetFlow Switching	120,000	150,000	250,000	<protocol> route-cache flow
Distributed Switching	300,000	N/A	> 1 million	<protocol> route-cache distributed

Agenda

- Perception versus Reality
- Layered Switching
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

Features Affecting Performance

- Queuing
- Compression
- Filtering
- Encryption
- Accounting

Queuing

- **The Cisco IOS implements four different queuing algorithms today:**
  - **First in, First Out (FIFO) Queuing**
  - **Priority Queuing**
  - **Custom Queuing**
  - **Weighted Fair Queuing**
- **Queuing occurs when network congestion occurs (i.e., the queue depth => 1), else all packets are sent as they arrive at the interface**

Priority Queuing

- Four queues: high, medium, normal and low
- Priority-list global and priority-group interface command defines rules for assignment of packets to priority queues

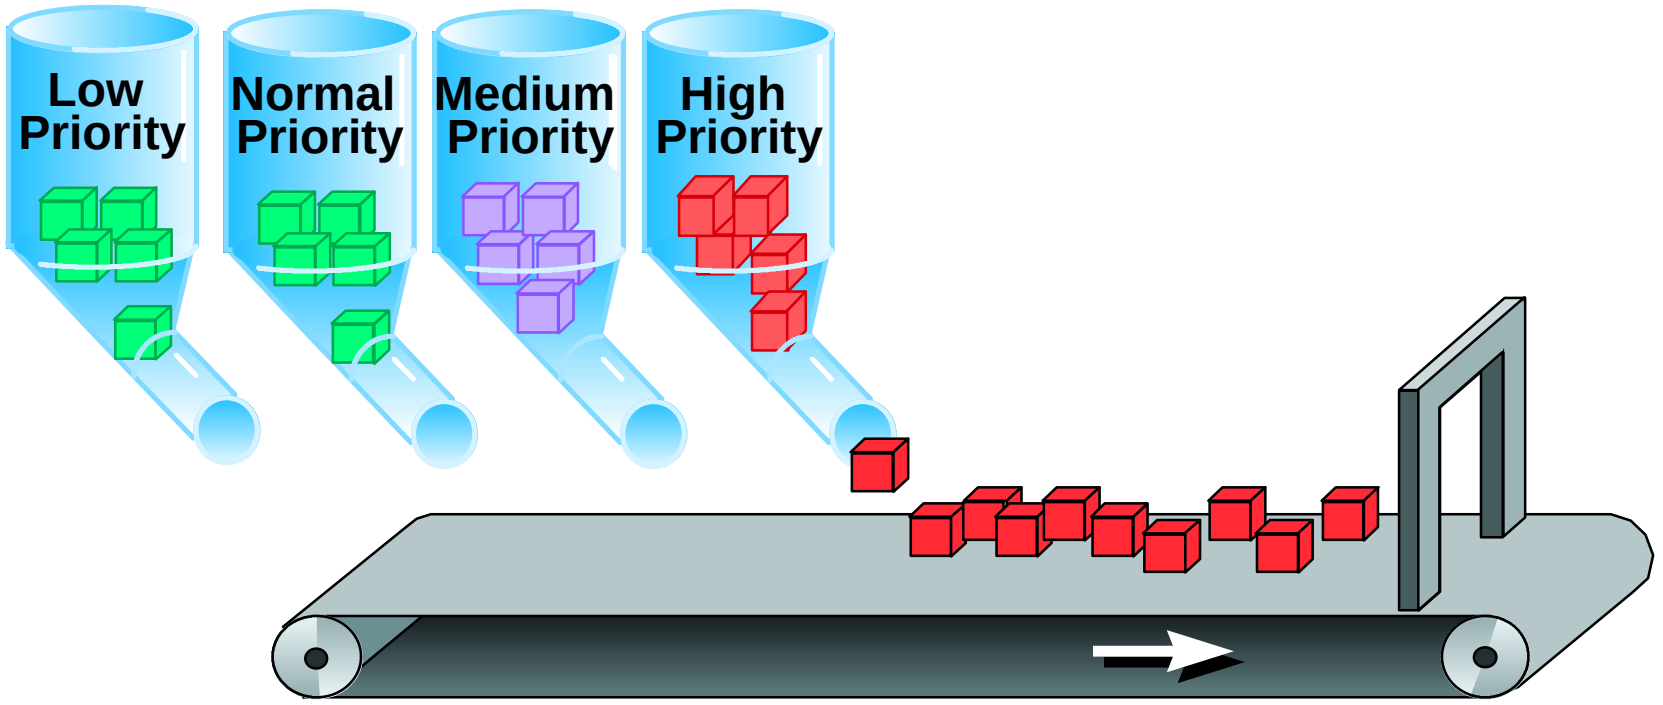

Custom Queuing

- Control % of interface bandwidth for specified traffic
- 17 output queues for each interface [16 configurable]
- Queue-list global and custom-queue-list interface commands

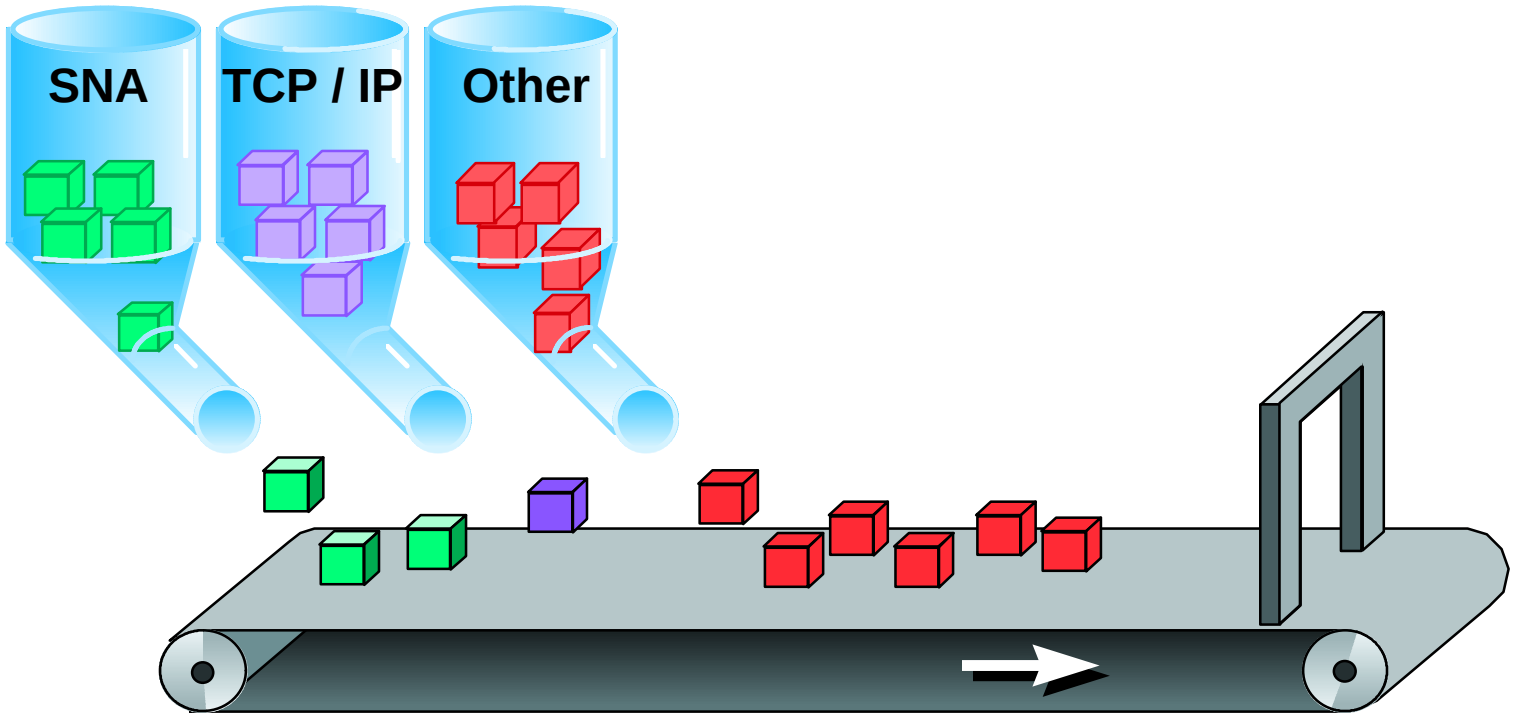

Weighted Fair Queuing

- Automatic traffic priority management
- Low-bandwidth sessions have priority over high-bandwidth sessions
- High-bandwidth sessions assigned weights
- Default for <2.048 Mbps interfaces
- Fair-queue interface command

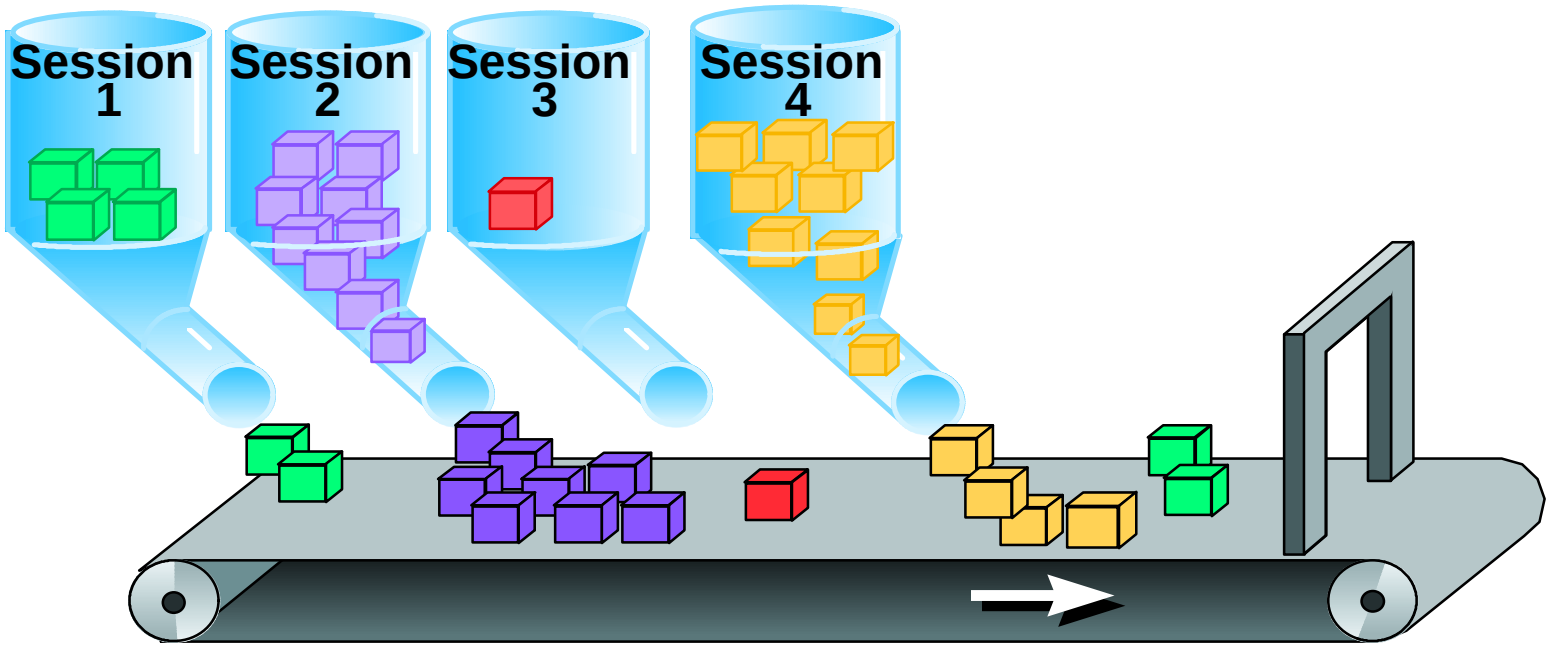

Random Early Detection

- Designed for congestion avoidance
- Prioritize traffic based on TOS (precedence)
- T3, OC-3, ATM
- Random-detect interface command

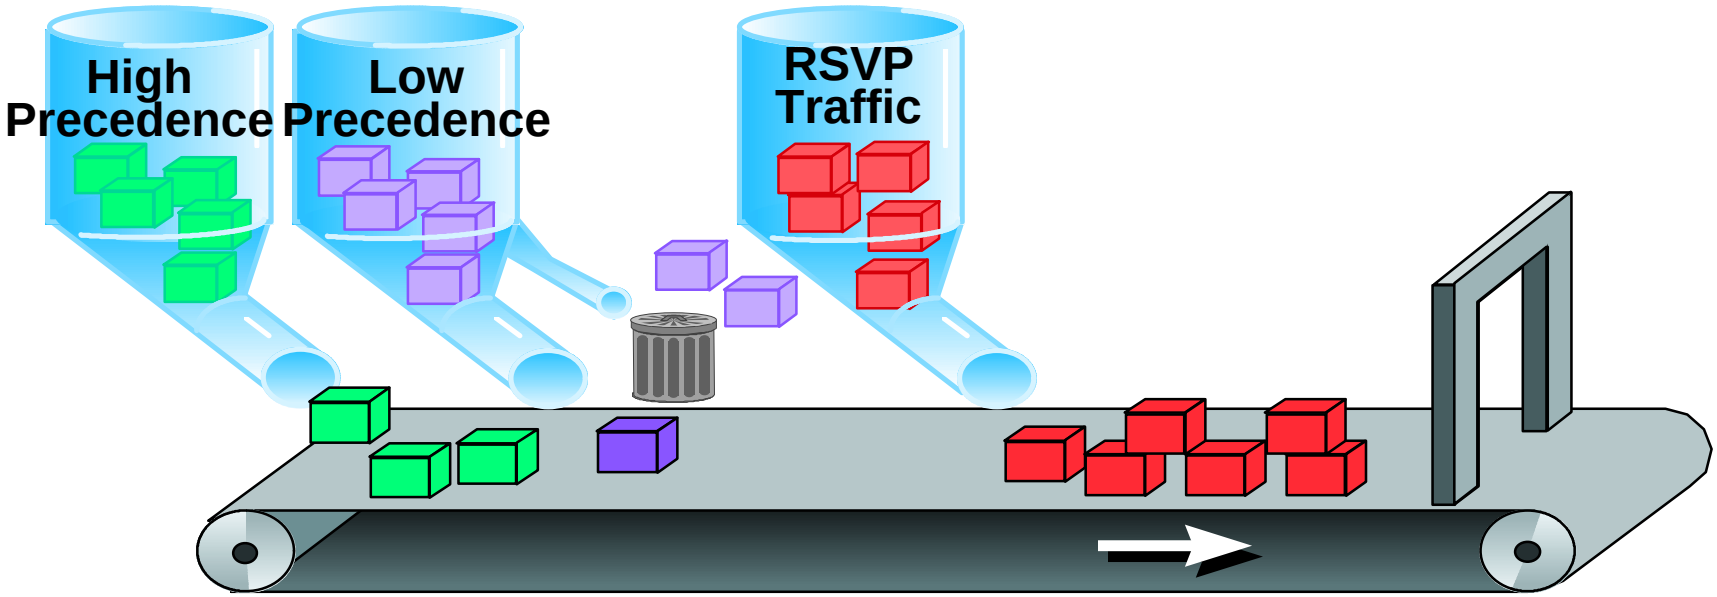

Cisco Systems

Cisco IOS Compression Solutions

Data Compression Method	Supported WAN Encapsulation	Characteristics
TCP/IP Header • Uses Van Jacobson algorithm (rfc 1144)	Frame Relay PPP X.25	• Up to 50% throughput improvement • Saves retransmission of redundant information in header • Delivers bandwidth savings for small-sized packets • Uses standard TCP/IP frames
Per Interface • STAC or Predictor • PPP/CCP supports STAC and Predictor	PPP HDLC LAPB	• Point-to-point lines or ISDN connections • Allows a mix of packet sizes and types • Provides good compression ratio • Protocol independent
Per-Virtual Circuit • STAC or Predictor	Frame Relay X.25	• Individual control of compression on each virtual circuit saves system resources • Protocol independent

STAC vs Predictor

- **STAC**

  Stac performs a little better than predictor at low speeds. (i.e., over ISDN B channels)

  Stac is very CPU intensive but uses little memory

- **Predictor**

  Predictor uses more memory but is not as CPU intensive

Filtering (Access Lists)

```
Taliskar# config terminal

access-list 101 permit igrp 144.254.1.0 0.0.0.255 144.254.1.0 0.0.0.255
access-list 101 permit  tcp  host 144.254.4.6 host 144.254.2.6  eq 80
access-list 101 permit  tcp  host 144.254.5.9 host 144.254.2.6  eq 80
access-list 101 deny tcp any any eq 80
access-list 101 permit tcp any 144.254.3.0 0.0.0.255
access-list 101 permit udp any 144.254.3.0 0.0.0.255

access-list 121 permit  tcp 144.254.2.6 0.0.0.0  any eq 80
access-list 121 deny tcp any any eq 80
access-list 121 permit  tcp  any any
access-list 121 permit udp  any any

interface FDDI 0/0
ip address 144.254.1.1 255.255.255.0
ip rout-cache sse
ip access-group 101 in
ip access-group 121 out
```

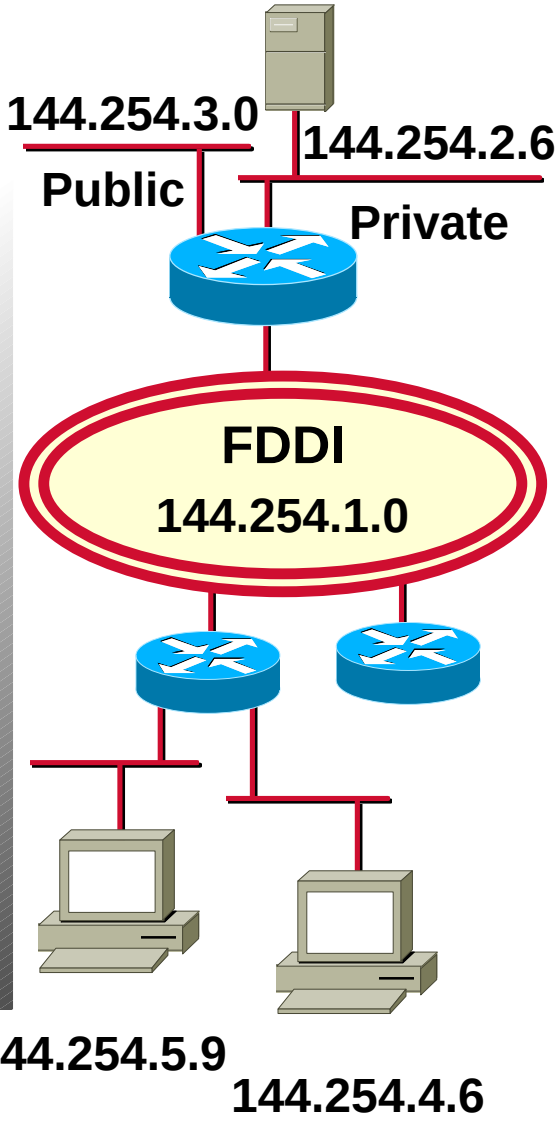

Enhancements to Access Lists

Cisco IOS Release	Performance Enhancement
9.21	Inbound and Outbound Access Lists Can Be Fast Switched
10.0	Standard Outbound Access Lists Can Be SSE Switched on Cisco 7000-Series Routers
10.3	Extended Outbound Access Lists Can Be SSE Switched on Cisco 7000-Series Routers
11.0 (3)	Inbound and Outbound, Standard and Extended Lists Can Be SSE Switched on Cisco 7000-Series Routers
11.1	Access Lists Can Use NetFlow Switching on Cisco 7500 and Cisco 7000-Series Routers with an RSP
11.1(5)	Access Lists Can Use NetFlow Switching on Cisco 7200-Series Routers

Encryption

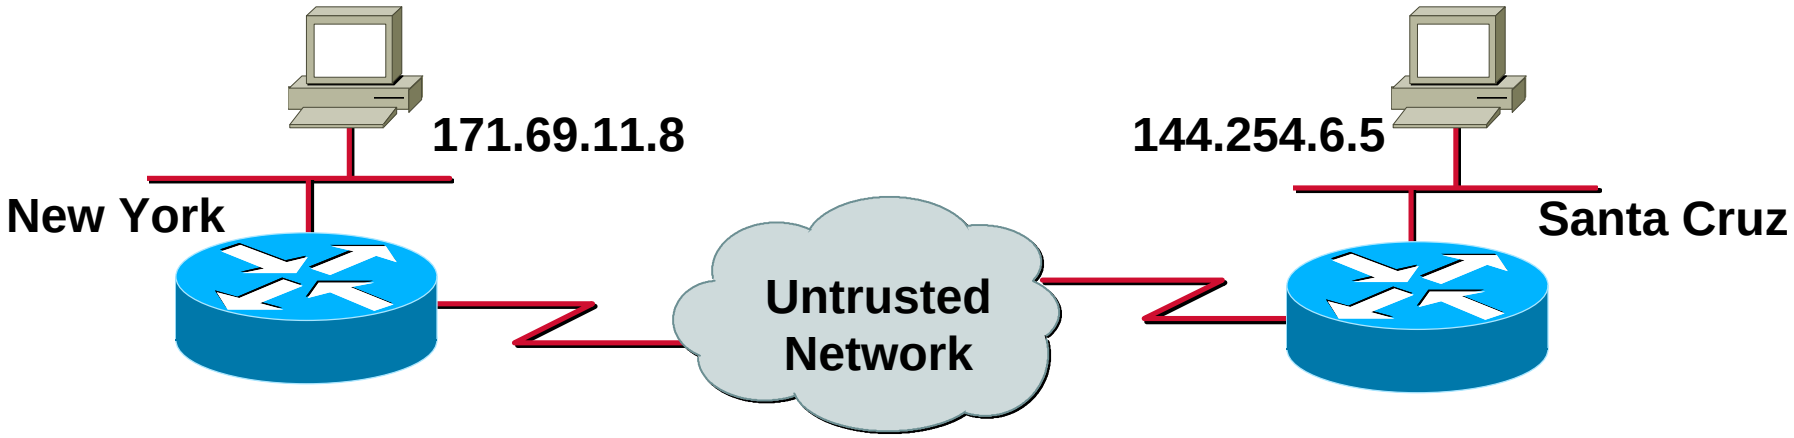


```
access-list 161 permit tcp
host
   171.69.11.8  host
144.254.6.5

crypto-map NY2SC
set algotithm des cfb-8
set peer SantaCruz
match address 161

interface serial0
ip  unnumbered e0
crypto map NY2SC
```

```
access-list 161 permit tcp
host
   144.254.6.5 host
171.69.11.8

crypto-map  SC2NY
set algotithm des cfb-8
set peer NewYork
match address 161

interface serial2/0
ip address 144.254.118.17
crypto map SC2NY
```

Accounting (IP and IPX)

Fast Switch Path…

```
Ardbeg(config-if)#   ip accounting ?
  access-violations  Account for IP packets violating access lists on this
                       interface
  output-packets     Account for IP packets output on this interface


Ardbeg#  show ip accounting output-packets
    Source          Destination      Packets     Bytes
 171.68.156.4    171.68.158.12            1        76
 171.68.158.9    171.68.156.4             1        76
 171.68.158.189   171.68.235.230         11       686
 192.150.42.74   171.68.158.11            2       152
 171.68.158.11   192.150.42.74            4       304
```

Agenda

- Perception versus Reality
- Layered Switching
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

Distributed vs. Centralized Servers

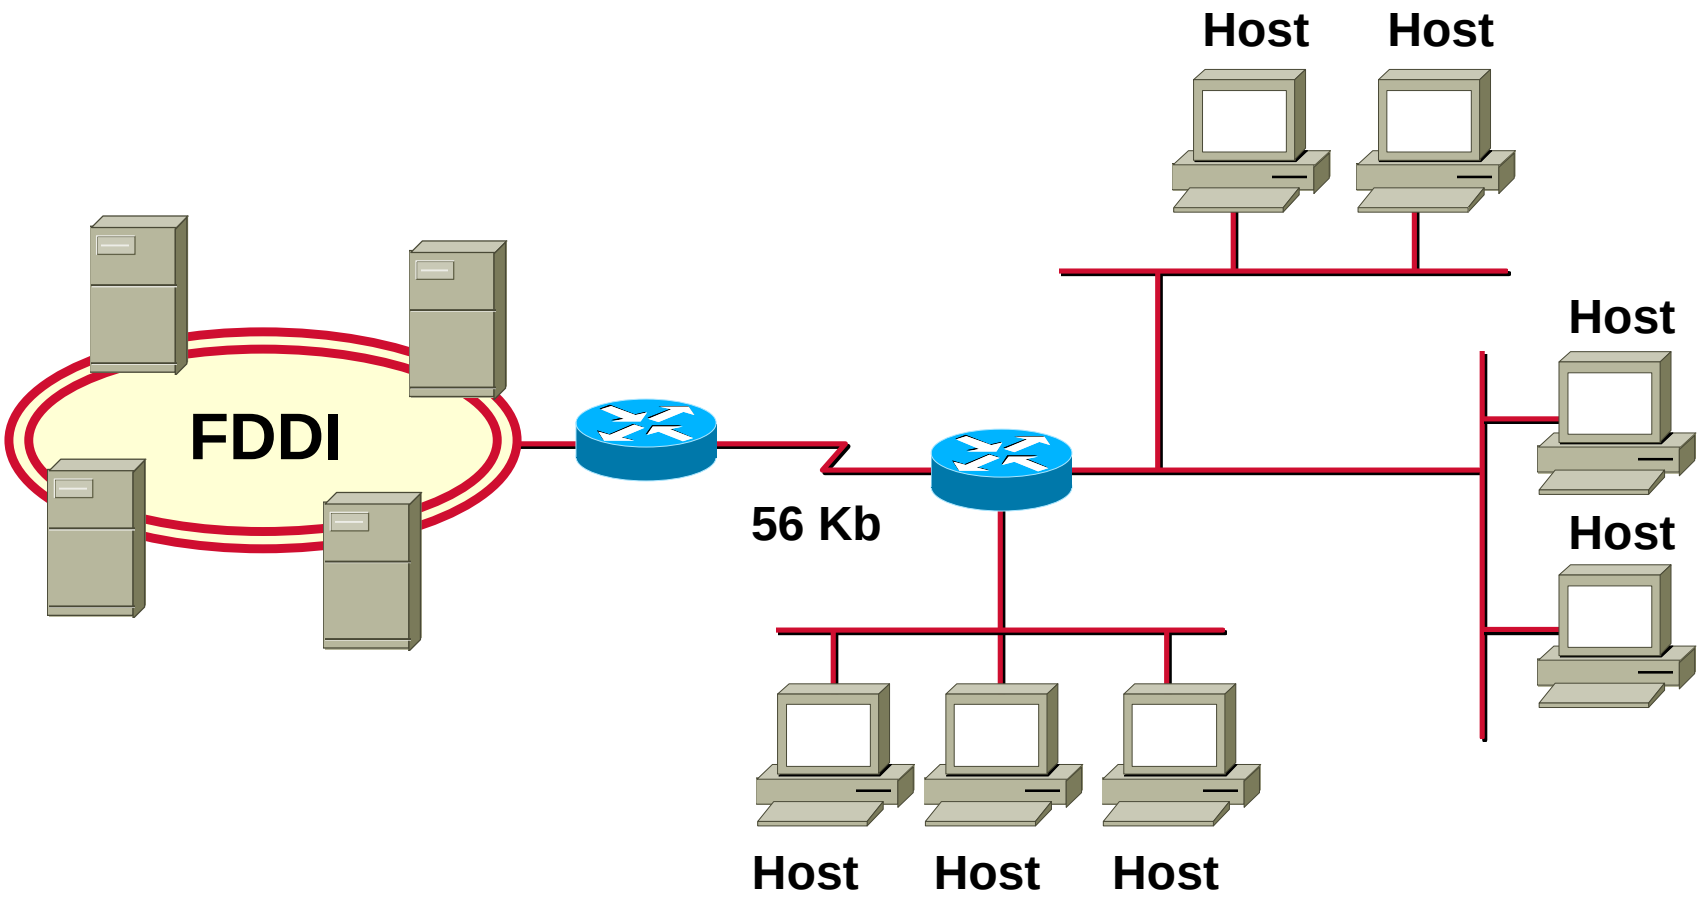

Cisco Systems

Serial-Line, Load-Balancing Process Switching

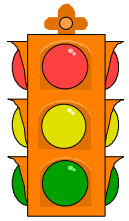

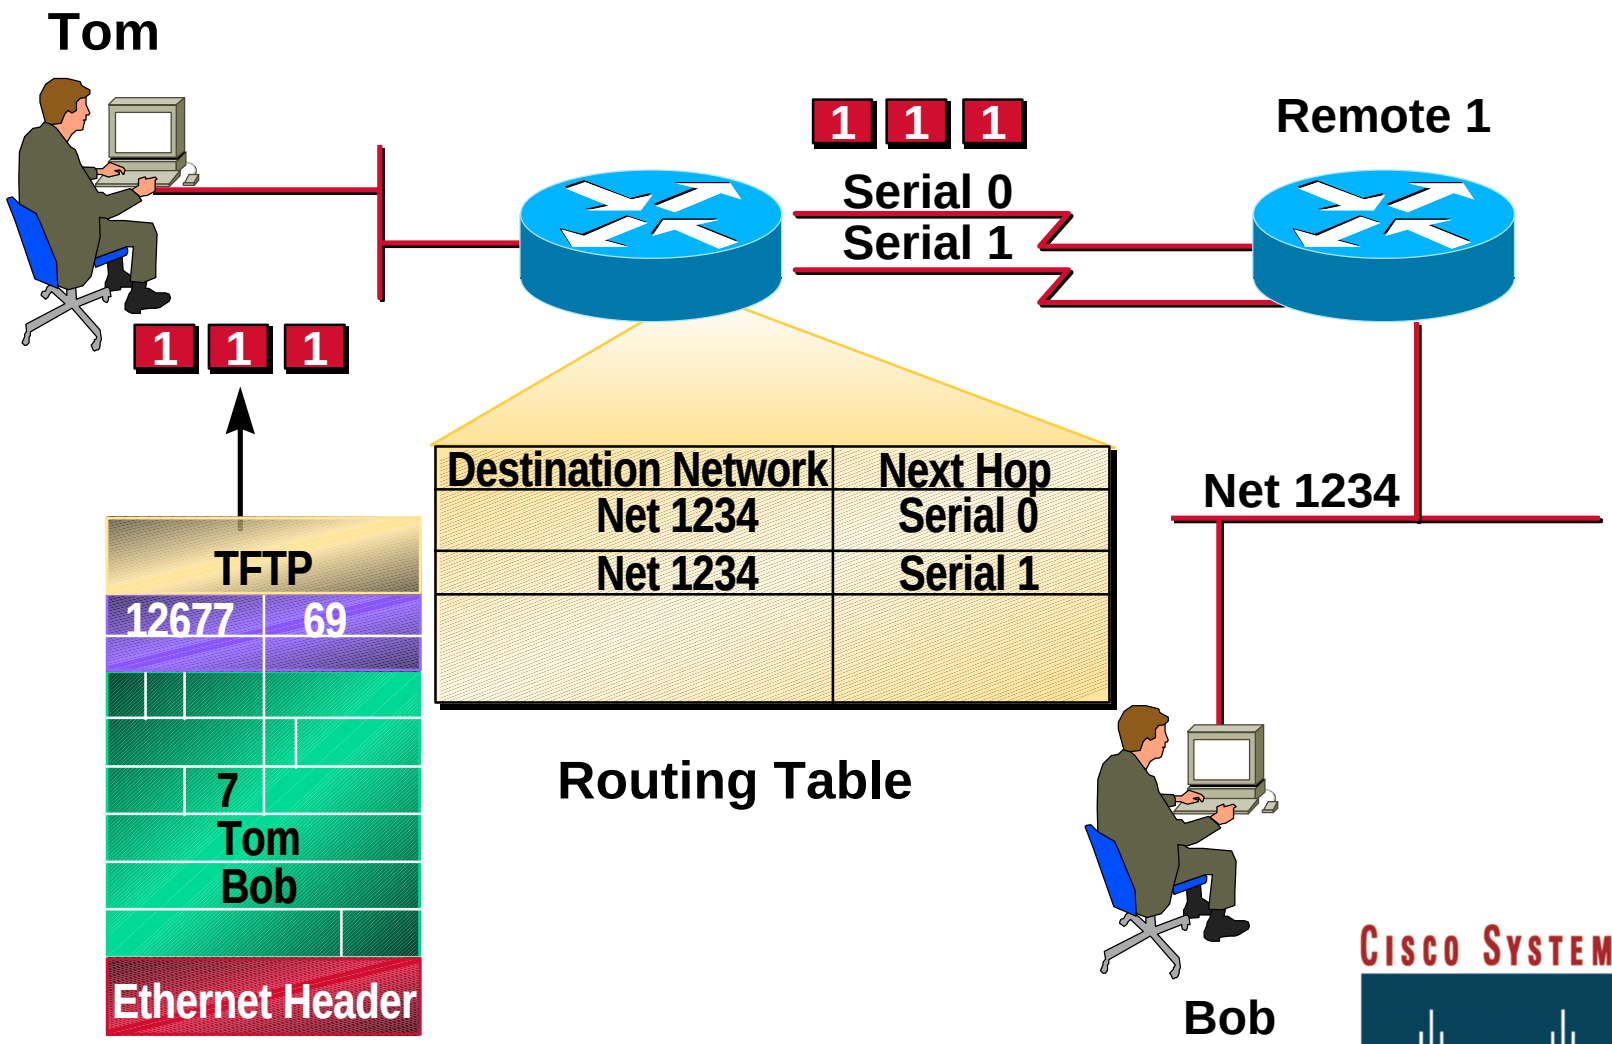

CISCO SYSTEMS

Serial-Line, Load-Balancing Fast/Optimum Switching

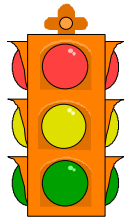

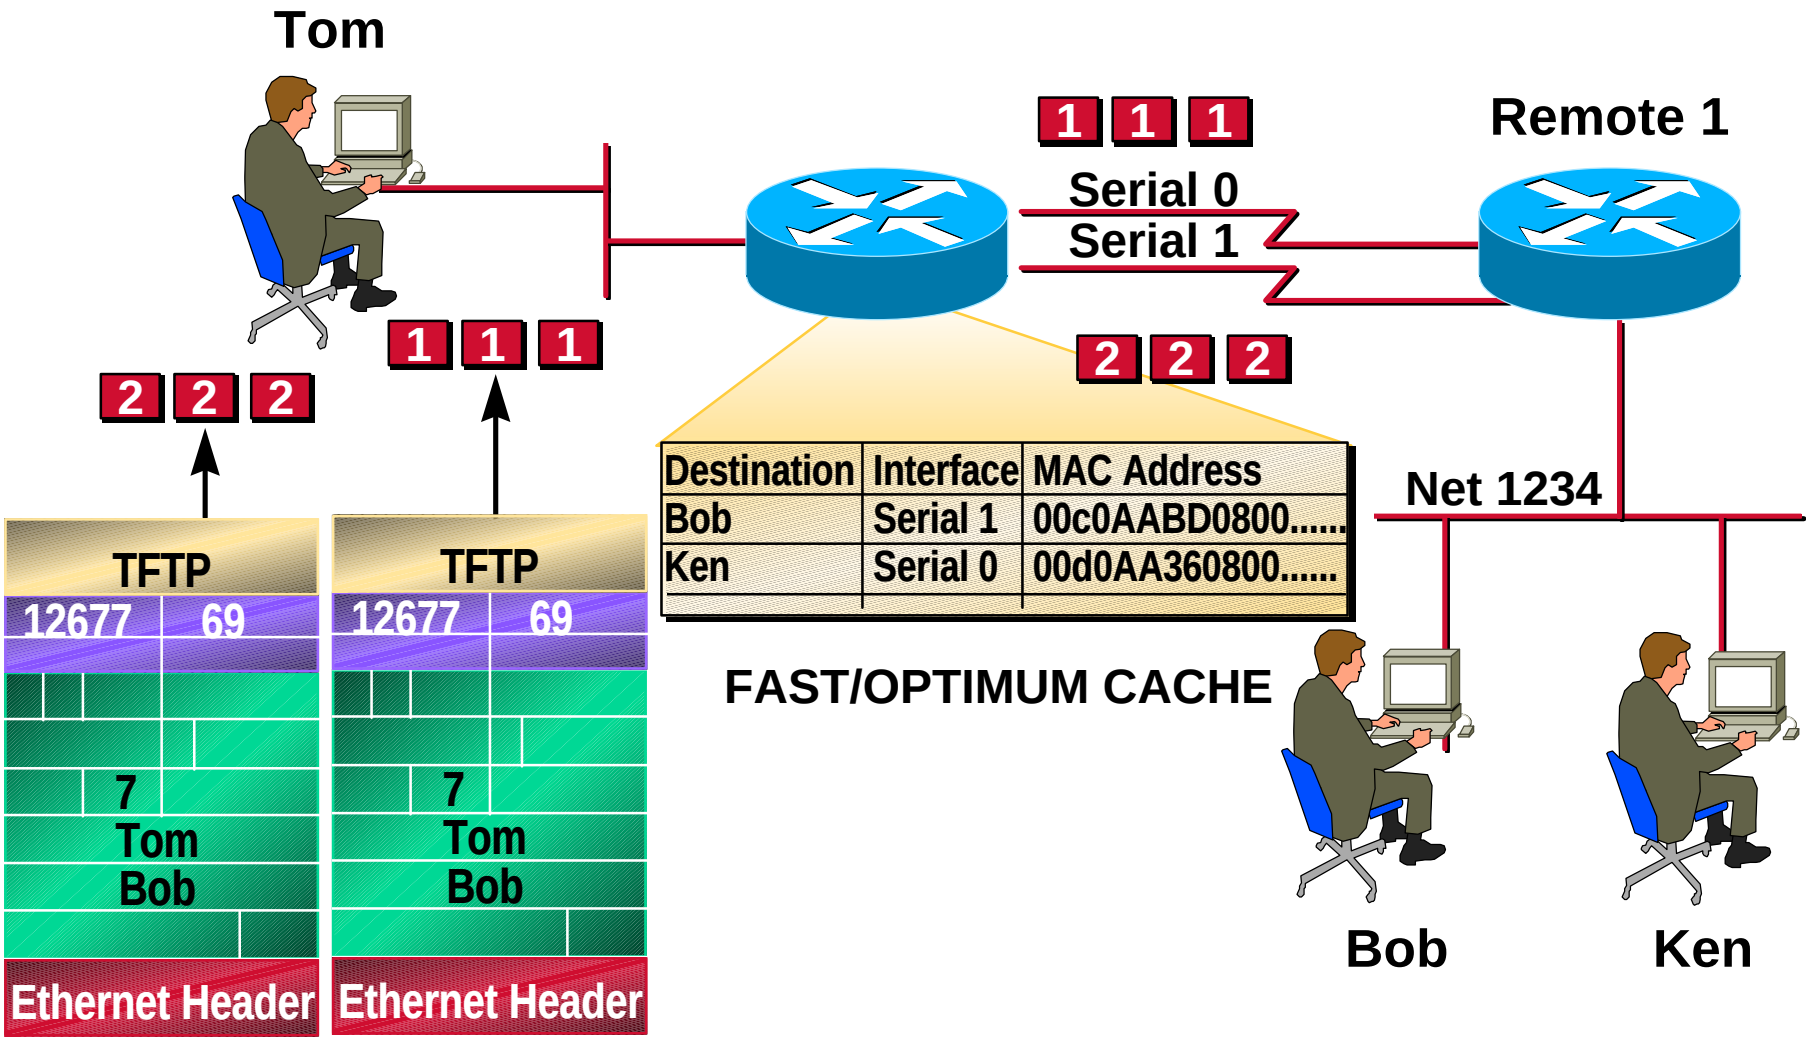

Serial-Line, Load-Balancing NetFlow Switching

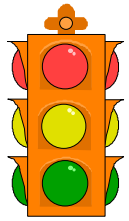

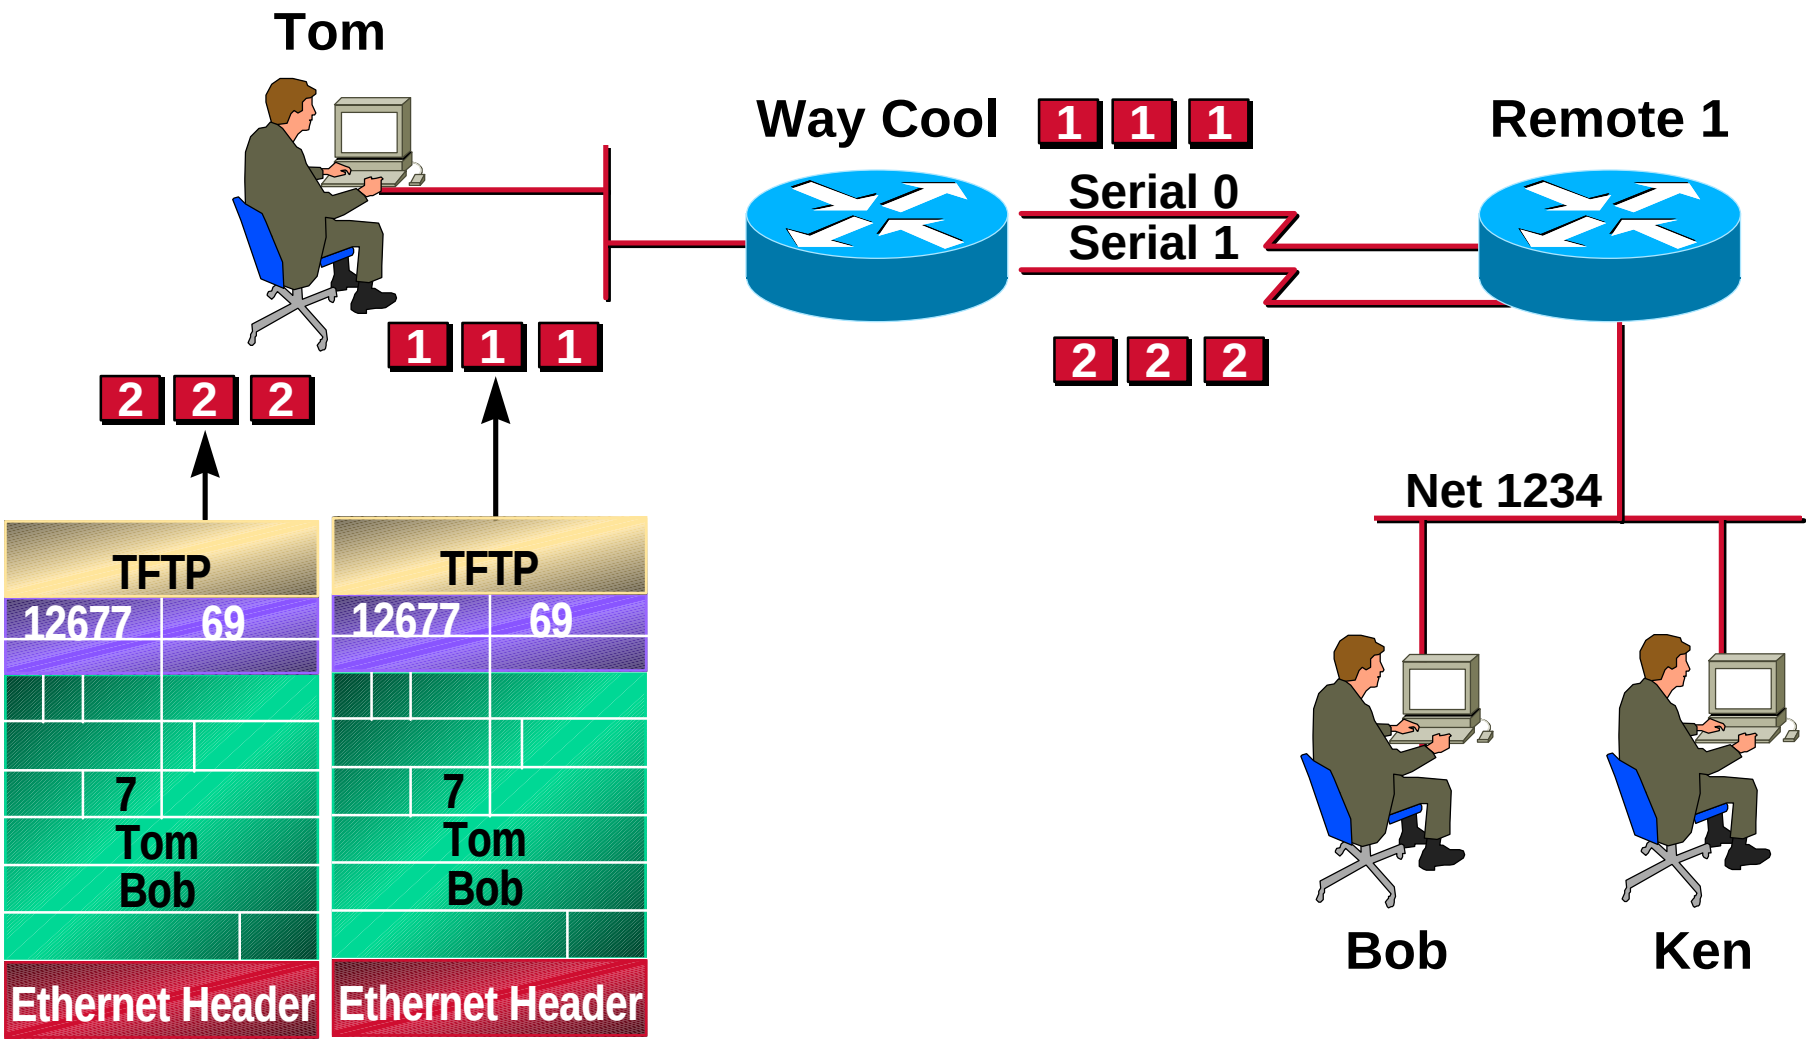

ISDN Aggregate Channels

Dialer Load—Threshold (Fast Switched)

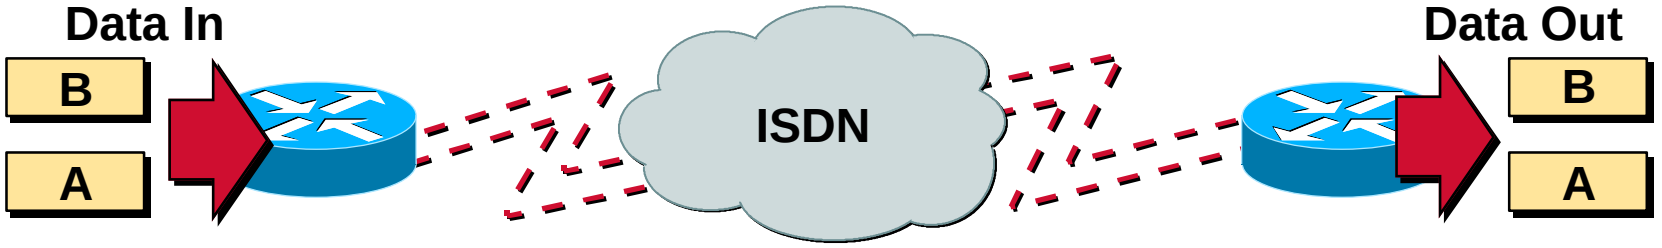

PPP Multilink—RFC1990 (Fast Switched)

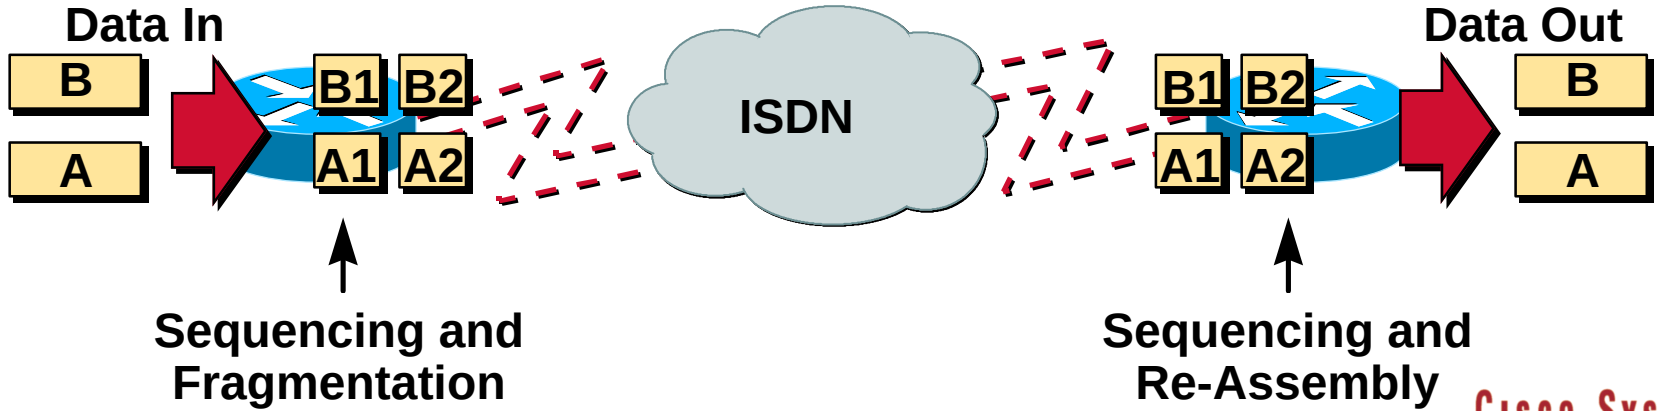

CISCO SYSTEMS

Agenda

- Perception versus Reality
- Layered Switching
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

Buffers and Queues (What Is an Ignore?)

- Ignore incremented when no buffer for incoming frame
- Queue limit can be exceeded to handle bursts

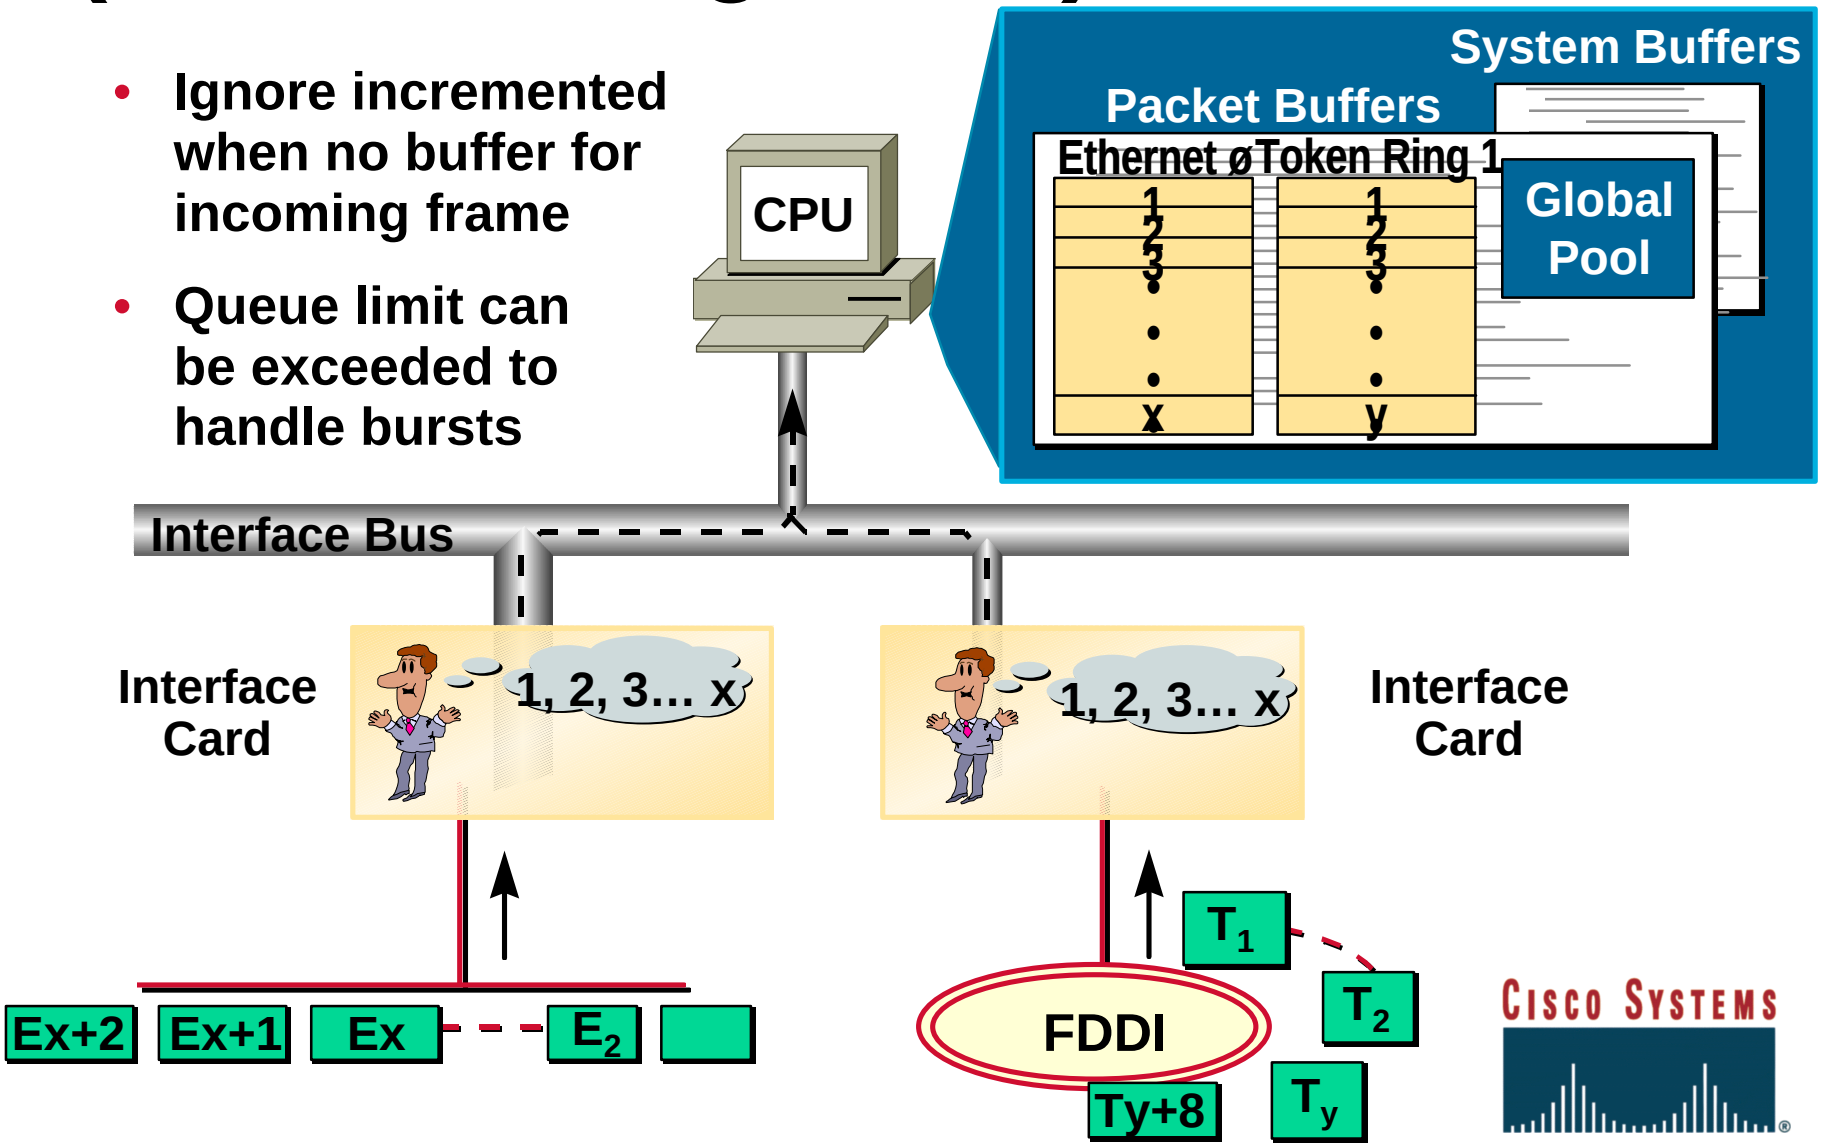

Buffers and Queues (Input Drops)

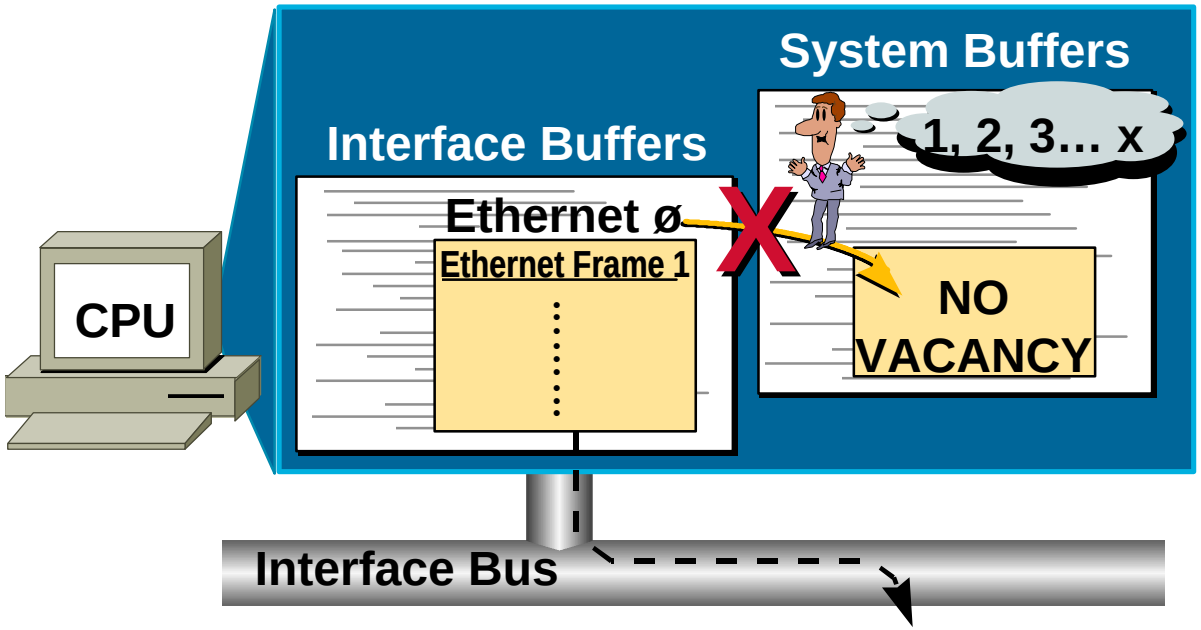

- Input drop increment when no system buffer available
- System buffers used for all process switched packets and packets router generates

Cisco Systems

Buffers and Queues (Output Drops)

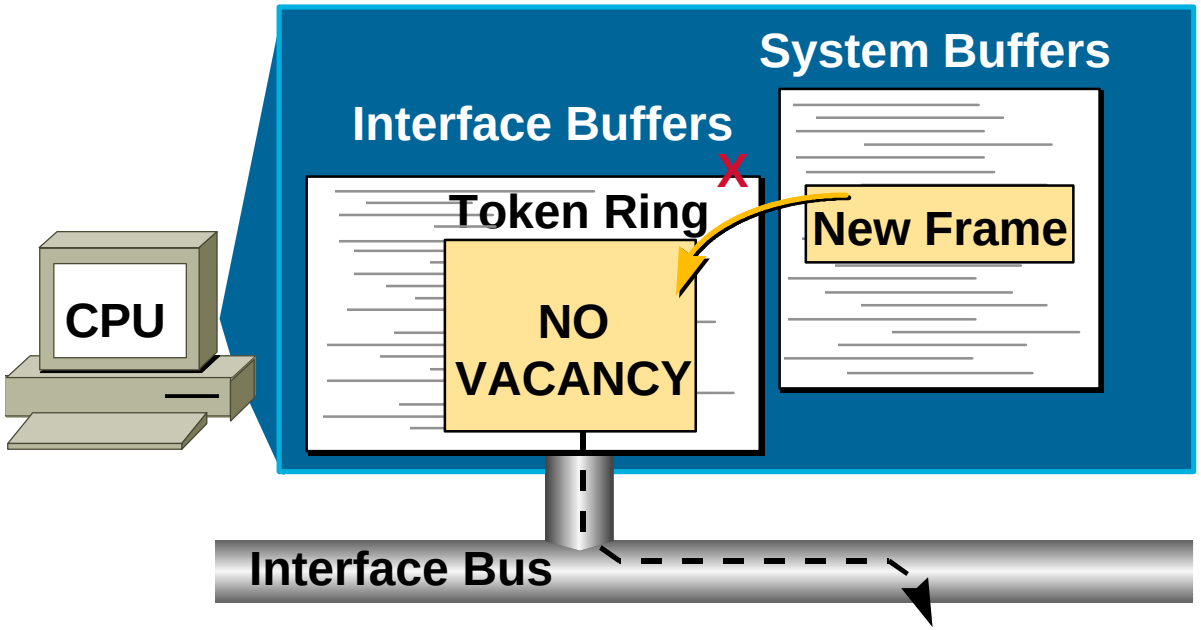

- Output drop incremented when no interface buffer available for outbound frame

Buffers and Queues (Output Drops)

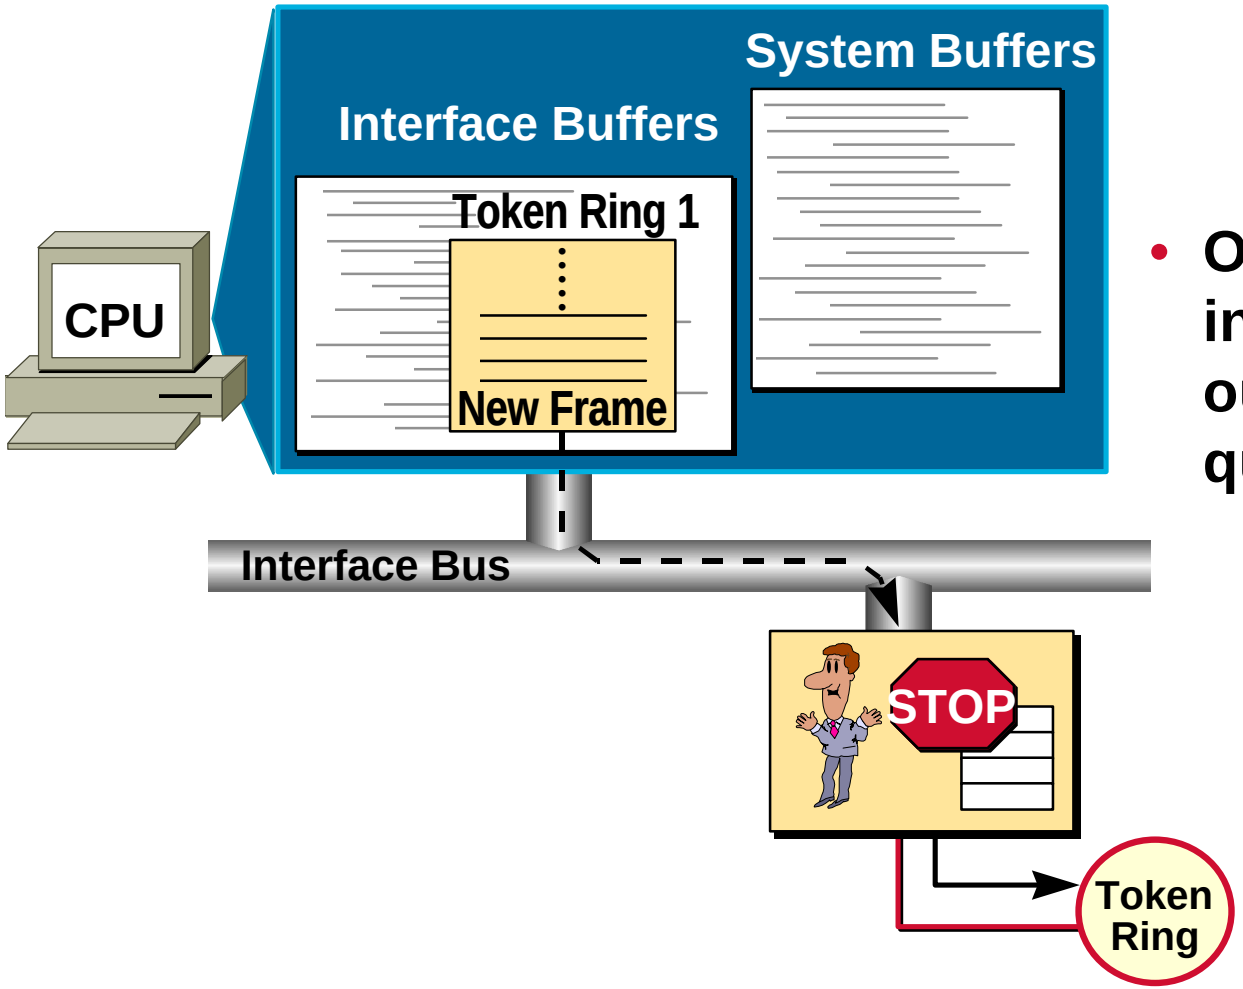

- Output drop incremented when output interface queue limit exceeded

Cisco Systems

Interface Buffers

```
HoStage> show controller cbus

cBus 0, controller type 6.0, microcode version 10.0
  512 Kbytes of main memory, 128 Kbytes cache memory
  134 1520 byte buffers, 65 4496 byte buffers
  Restarts: 0 line down, 0 hung output, 0 controller error

MEC 0, controller type 5.1, microcode version 10.0
  Interface 0 - Ethernet0, station address 0000.0c06.4ae0 (bia 0000.0c06.4ae0)
    11 buffer RX queue  threshold, 18 buffer TX queue limit, buffer size 1520
    ift 0000, rql 11, tq 0000 0000, tql 18
    Transmitter delay is 0 microseconds

CTR 1, controller type 9.0, microcode version 10.1
  Interface 8 - TokenRing0, station address 0000.3060.3219 (bia 0000.3060.3219)
    13 buffer RX queue  threshold, 31 buffer TX queue limit, buffer size 4496
    ift 0005, rql 13, tq 0000 0000, tql 31
    Transmitter delay is 0 microseconds

FDDI-T 3, controller type 7.2, microcode version 10.1
  Interface 24 - Fddi0, station address 0000.0c06.36d7 (bia 0000.0c06.36d7)
    13 buffer RX queue  threshold, 32 buffer TX queue limit, buffer size 4496
    ift 0006, rql 9, tq 0000 0000, tql 32
```

System Buffers

```
HoStage# show  buffers

Buffer elements:
     500 in free list (500 max allowed)
     51640224 hits, 0 misses, 0 created

Small buffers, 104 bytes (total 121, permanent 121):
     119 in free list (20 min, 250 max allowed)
     19229201 hits, 0 misses, 0 trims, 0 created

Middle buffers, 600 bytes (total 90, permanent 90):
     89 in free list (10 min, 200 max allowed)
     20513359 hits, 91 misses, 115 trims, 115 created

Big buffers, 1524 bytes (total 90, permanent 90):
     90 in free list (5 min, 300 max allowed)
     7160285 hits, 0 misses, 0 trims, 0 created

Large buffers, 5024 bytes (total 5, permanent 5):
     5 in free list (0 min, 30 max allowed)
     233295 hits, 0 misses, 0 trims, 0 created

Huge buffers, 18024 bytes (total 0, permanent 0):
     0 in free list (0 min, 4 max allowed)
     0 hits, 0 misses, 0 trims, 0 created
```

Interface Statistics

```
HoStage# show  interface  ethernet 5

Ethernet 5 is up, line protocol is up
  Hardware is cBus Ethernet, address is 0000.0c06.4ae5 (bia 0000.0c06.4ae5)
  Description: Harbrinder Network Connection
  Internet address is 192.150.42.126, subnet mask is 255.255.255.248
  MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 4:00:00
  Last input 0:00:08, output 0:00:04, output hang never
  Last clearing of "show interface" counters 0:00:00
  Output queue 0/40, 0 drops; input queue 0/75, 5 drops
  Five minute input rate 0 bits/sec, 0 packets/sec
  Five minute output rate 2000 bits/sec, 0 packets/sec
     4703028 packets input, 605811959 bytes,     1 no buffer
     Received 184433 broadcasts, 5 runts, 0 giants
     35 input errors, 15 CRC, 15 frame, 0 overrun,           , 0 abort  470 ignored
     16903029 packets output, 2249122236 bytes, 0 underruns
     0 output errors, 6577 collisions, 2 interface resets, 0 restarts
```

Switching Paths Used

CadenHead# **show interface stat**

Ethernet0/0

Switching path	Pkts In	Chars In	Pkts Out	Chars Out
Processor	2778308	269468540	1549954	137370016
Route cache	91	6983	9	756
Autonomous/SSE	106360	18478151	4286	382009
Total	2884759	287953674	1554249	137752781

Ethernet0/1

Switching path	Pkts In	Chars In	Pkts Out	Chars Out
Processor	1958203	193347523	1338734	116950789
Route cache	0	0	0	0
Autonomous/SSE	74478	7844330	1	336
Total	2032681	201191853	1338735	116951125

Interface Configuration

```
CadenHead#  show ip interface

Ethernet0/0 is up, line protocol is up
  Internet address is 171.68.156.3 255.255.255.248
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is enabled
  Multicast reserved groups joined: 224.0.0.10
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Security level is default
```

Interface Configuration (Cont.)

Split horizon is enabled
ICMP redirects are always sent
ICMP unreachables are always sent
ICMP mask replies are never sent
IP autonomous switching is enabled
IP multicast fast switching is enabled
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
Probe proxy name replies are disabled
Gateway Discovery is disabled
Policy routing is disabled

Performance Summary

- Understand performance requirement
- Use fastest supported switching path
- Choose appropriate router platform
- Carefully implement performance-affecting features

Q&A


CISCO SYSTEMS
EMPOWERING THE
INTERNET GENERATION℠