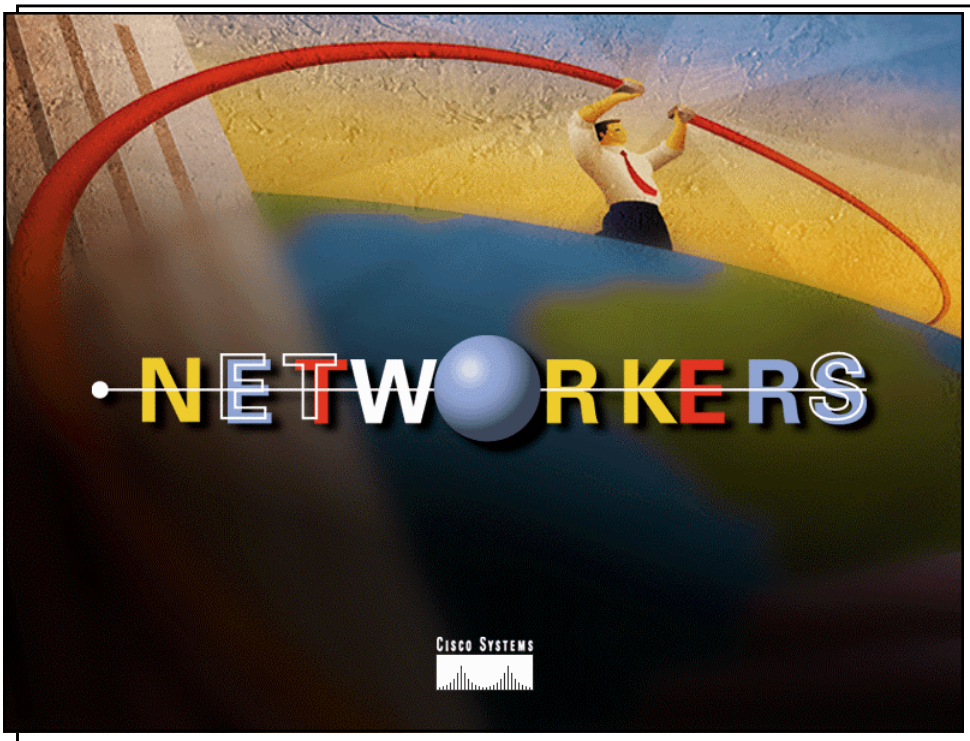

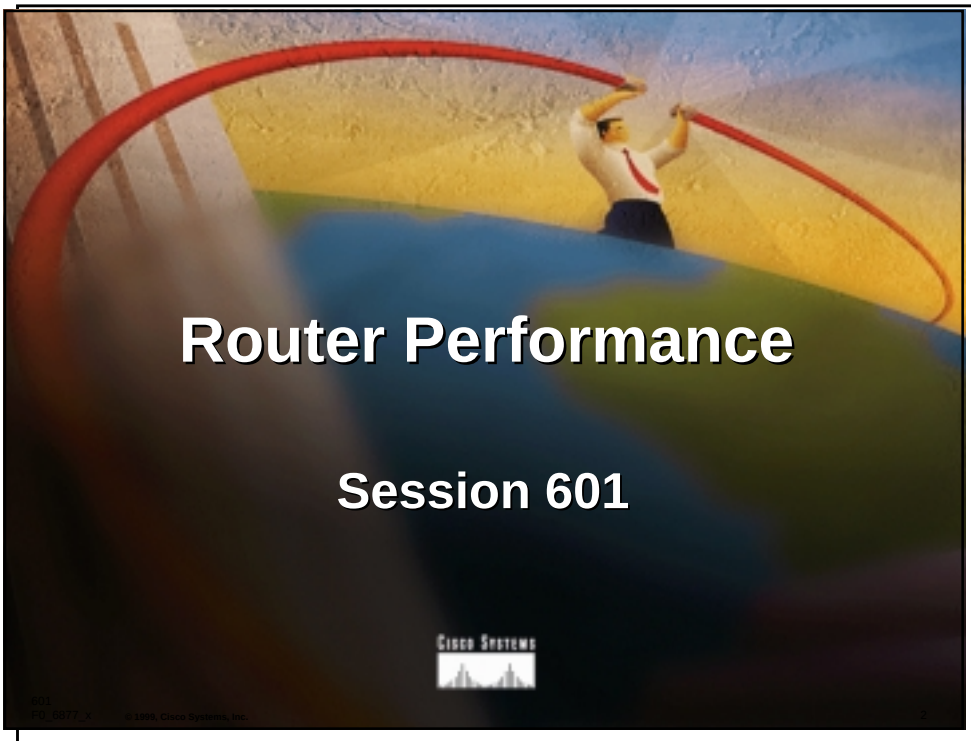

Agenda

- Measuring Performance
- Router Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

www.cisco.com

Performance Reality

"Vendors often engage in 'Specmanship' in an attempt to give their products a better position in the marketplace. This usually involves much 'smoke and mirrors' to confuse the user."

RFC 1242/1944 Scott Bradner

www.cisco.com

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

www.cisco.com

Theoretical Values

Bandwidth ÷ Packet Size = Theoretical Performance

Smaller Packets (Less Efficient, More pps)

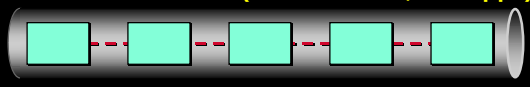

Bigger Packets (Better Utilization, Less pps)

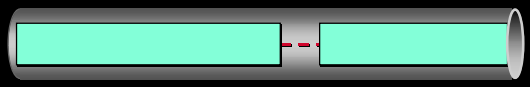

www.cisco.com

Maximum Theoretical PPS

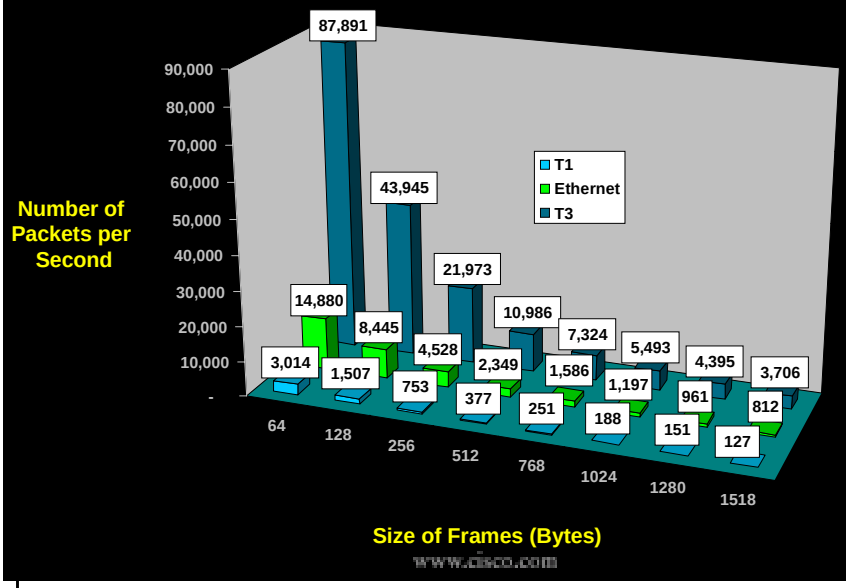

www.cisco.com

Cisco Low-End/ Mid-Range Switching Paths

www.cisco.com

Low-End/Mid-Range Process Switching

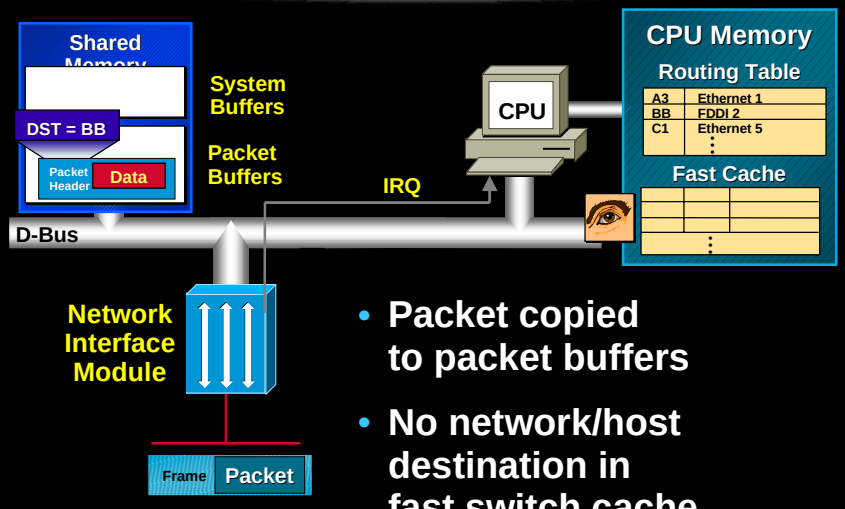

- Packet copied to packet buffers
- No network/host destination in fast switch cache

www.cisco.com

Low-End/Mid-Range Process Switching

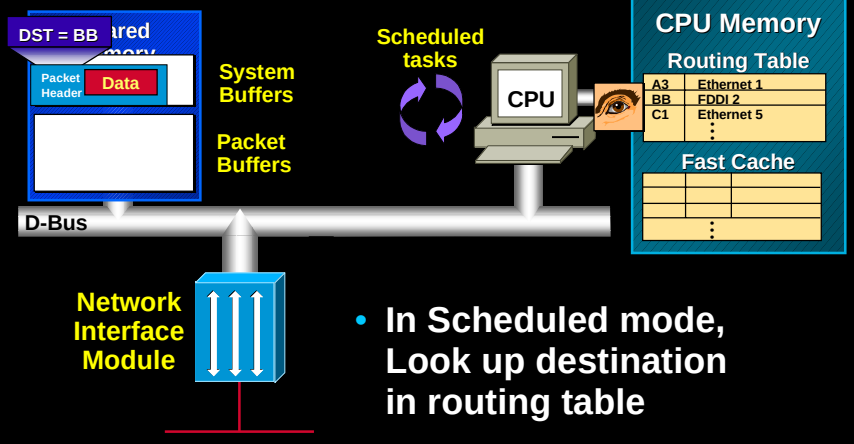

- In Scheduled mode, Look up destination in routing table
- Initialize fast cache

www.cisco.com

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
```

www.cisco.com

Low-End/Mid-Range Process Switching

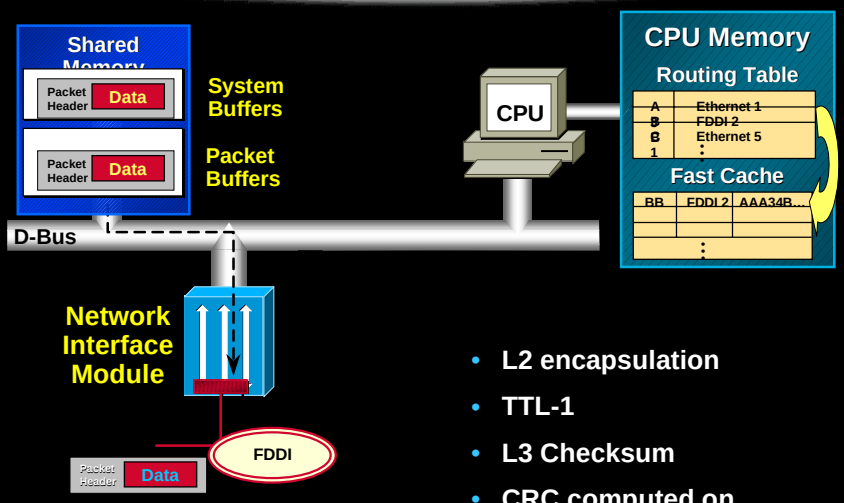

- L2 encapsulation
- TTL-1
- L3 Checksum
- CRC computed on network interface model

www.cisco.com

Fast Switching

- Recursion resolution at process level
- Classfull
- Per-destination load sharing
- Cache entries aged to limit the cache size

www.cisco.com

Fast Switching

- Overhead inherent to cache maintenance
  - Route change
  - Interface state change
  - Configuration change
- Assumes finite number of active flows

www.cisco.com

Deployment Information

- Available in all platforms
- Enabled using the ‘ip route-cache’ interface command
- Deprecated

www.cisco.com

Low-End/Mid-Range Fast Switching

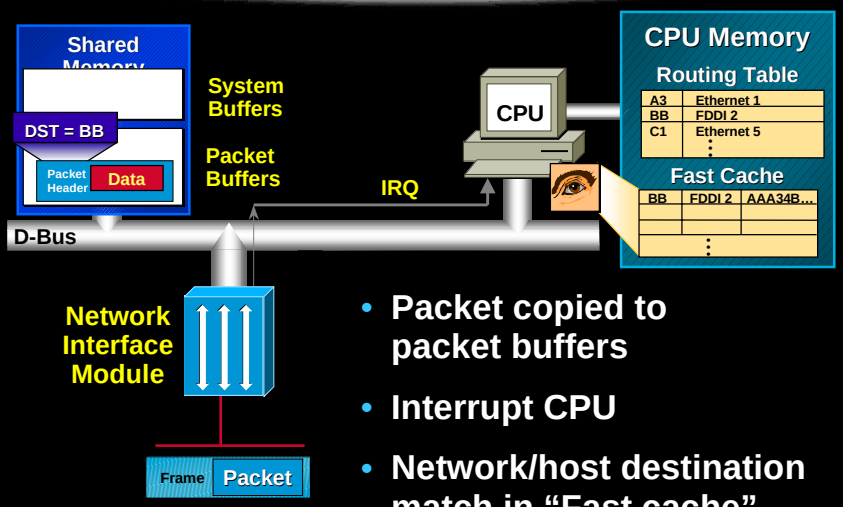

- Packet copied to packet buffers
- Interrupt CPU
- Network/host destination match in “Fast cache”

www.cisco.com

Low-End/Mid-Range Fast Switching

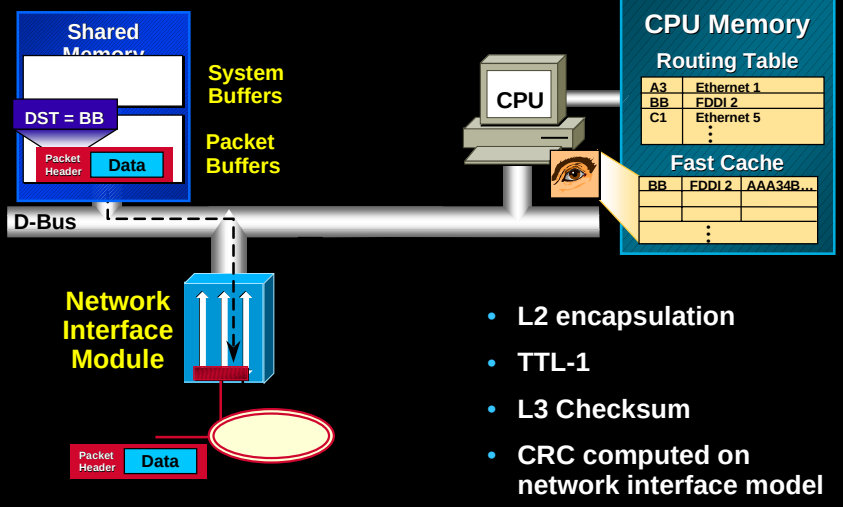

- L2 encapsulation
- TTL-1
- L3 Checksum
- CRC computed on network interface model

www.cisco.com

Low-Mid Range Performance

Platform	Process	Fast
1720	3,000	15,000
25xx	760	6,000
261x	1,500	12,000
262x	2,000	14,000
3620	2,000	16,000
3640	4,000	40,000
4500	10,000	45,000
4700	11,000	50,000
command	no <protocol> route cache	<protocol> route cache

www.cisco.com

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

www.cisco.com

Cisco 7000 Series—Process Switching

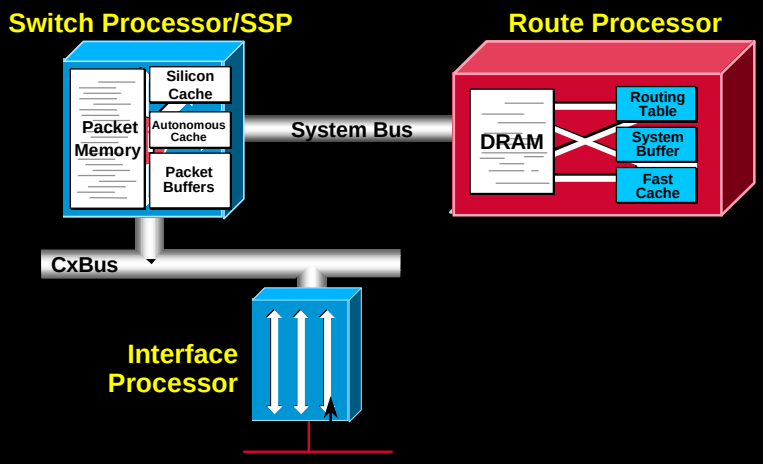

www.cisco.com

Cisco 7000 Series—Process Switching

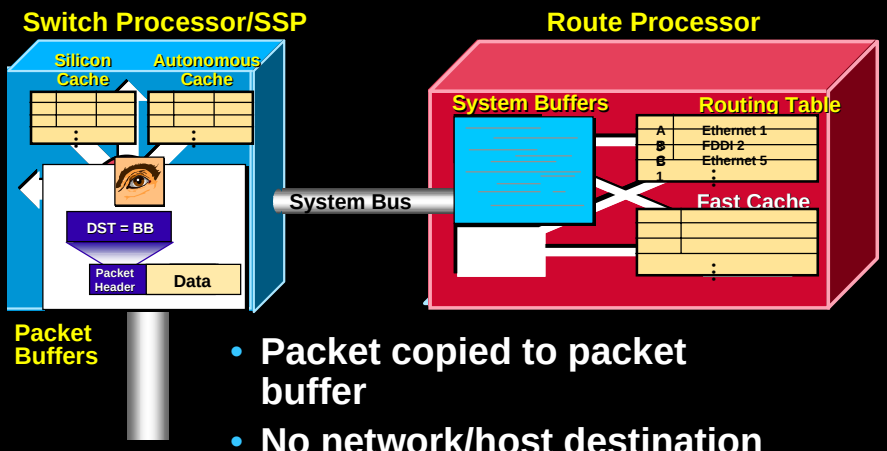

- Packet copied to packet buffer
- No network/host destination in silicon or autonomous cache

www.cisco.com

Cisco 7000 Series—Process Switching

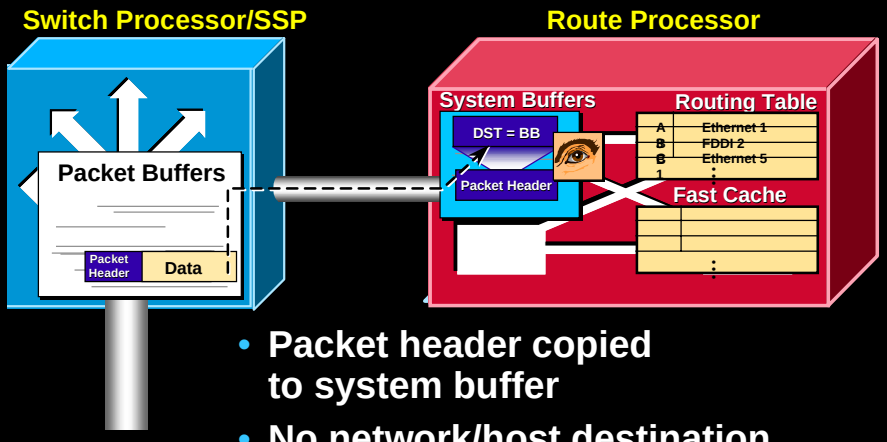

- Packet header copied to system buffer
- No network/host destination in fast switch cache

www.cisco.com

Cisco 7000 Series—Process Switching

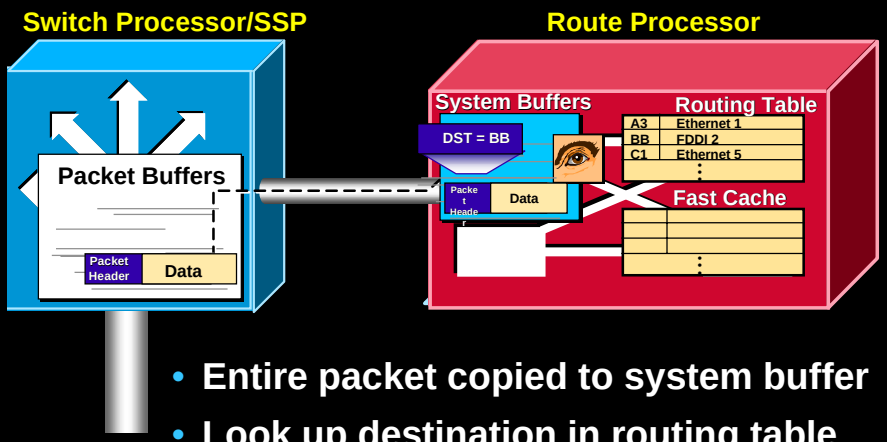

- Entire packet copied to system buffer
- Look up destination in routing table and initialize fast switch cache

www.cisco.com

Autonomous/SSE Switching

- Hardware/microcode assist
- Extension of fast switching cache
- Increased performance/reduce functionality
- Cache misses bubble up the packet
- Same issues as fast switching
- Now deprecated

www.cisco.com

Cisco 7000 Series—Initializing Autonomous Switching Cache

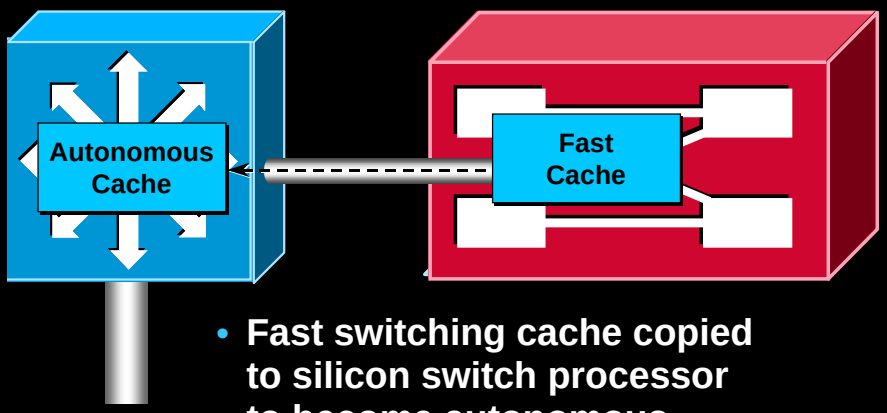

- Fast switching cache copied to silicon switch processor to become autonomous switch cache

www.cisco.com

Cisco 7000 Series—Initializing Silicon Switching Code

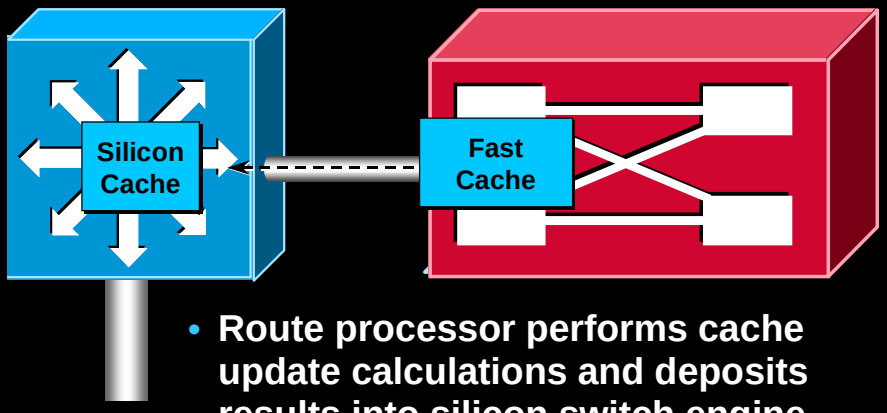

- Route processor performs cache update calculations and deposits results into silicon switch engine

www.cisco.com

Cisco 7000 Series—Process Switching

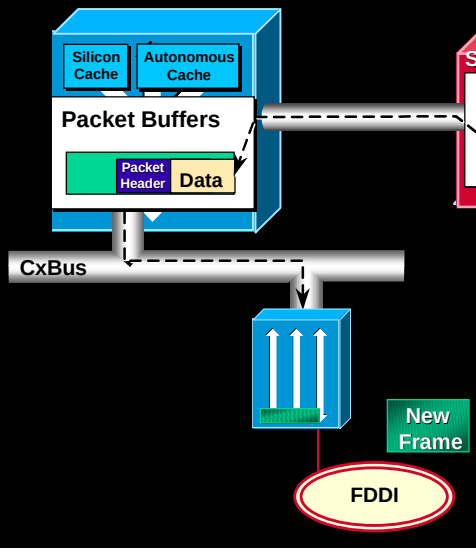

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco 7000 Series—Fast Switching

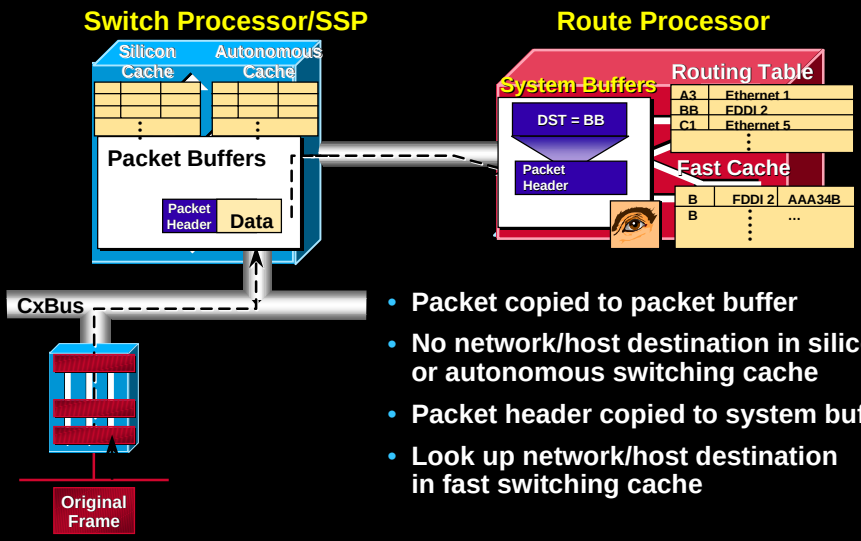

- Packet copied to packet buffer
- No network/host destination in silicon or autonomous switching cache
- Packet header copied to system buffer
- Look up network/host destination in fast switching cache

Cisco 7000 Series—Fast Switching

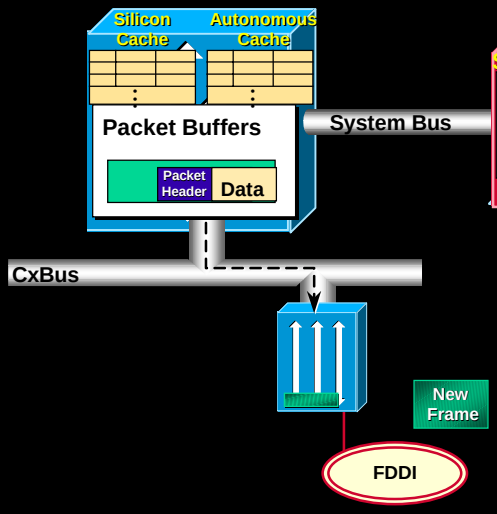

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco 7000 Series Autonomous Switching

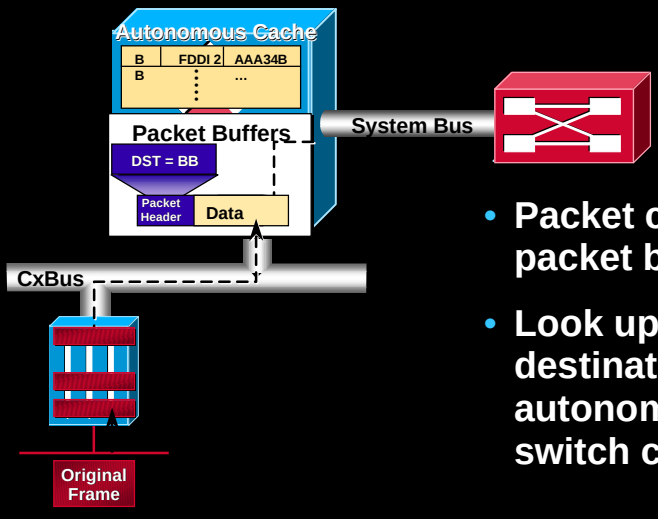

- Packet copied to packet buffer
- Look up Layer 3 destination in autonomous switch cache

Cisco 7000 Series Autonomous Switching

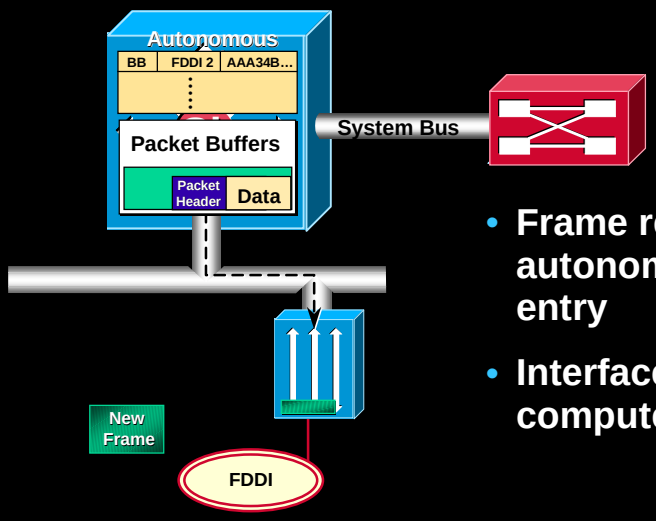

- Frame rewrite using autonomous cache entry
- Interface processor computes CRC

www.cisco.com

Cisco 7000 Series Silicon Switching

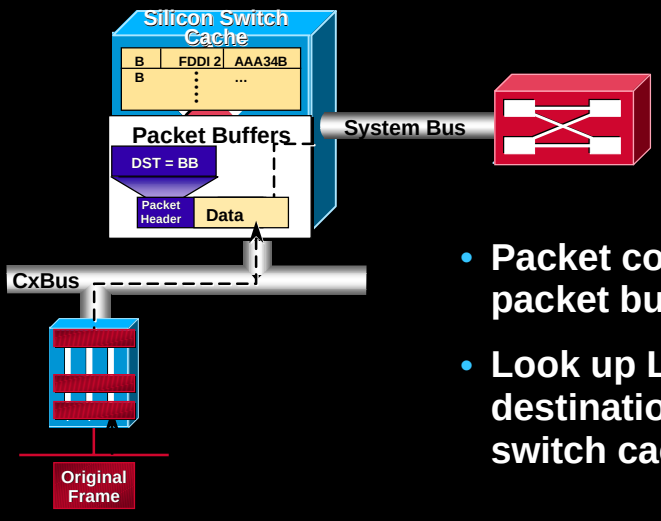

- Packet copied to packet buffer
- Look up Layer 3 destination in silicon switch cache

www.cisco.com

Cisco 7000 Series Silicon Switching

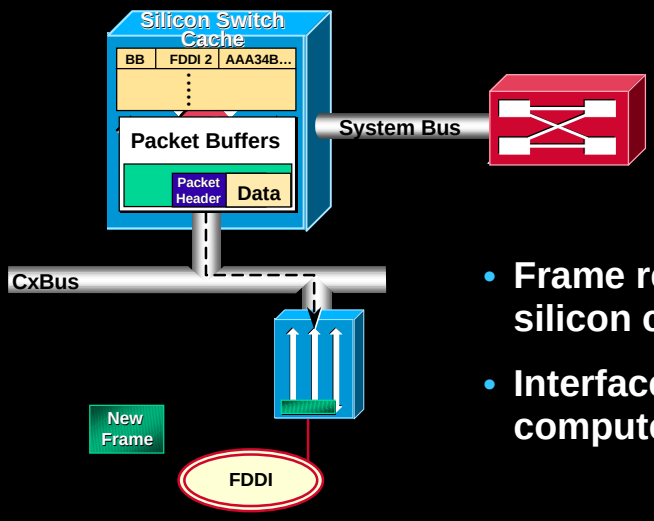

- Frame rewrite using silicon cache entry
- Interface processor computes CRC

www.cisco.com

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

www.cisco.com

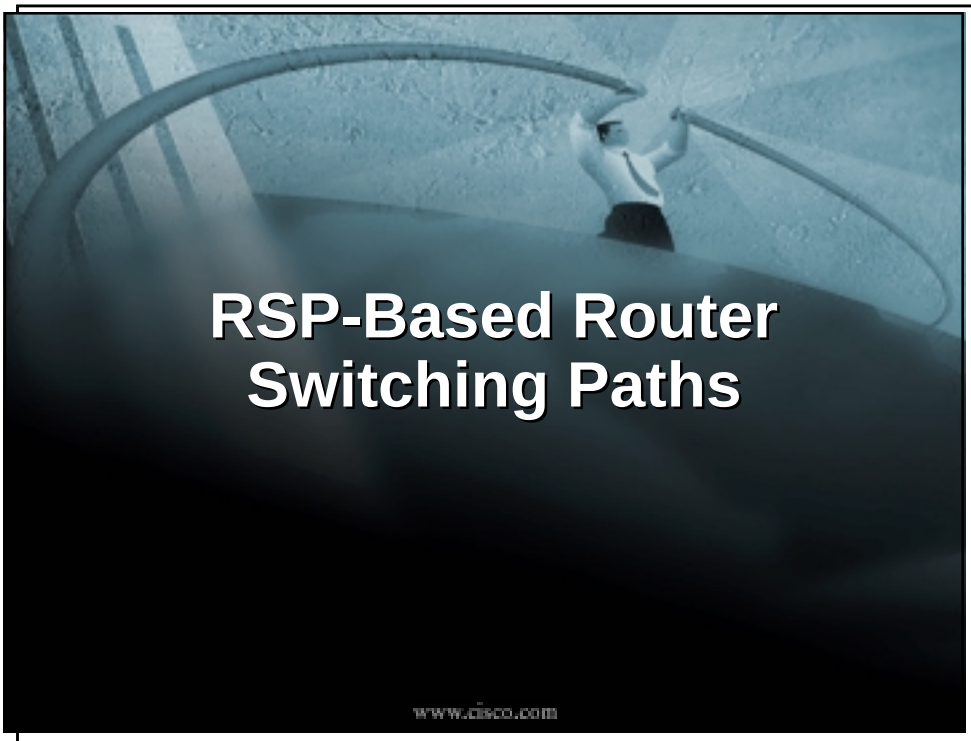

RSP-Based Router Switching Paths

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)

www.cisco.com

Cisco 7500 Series—Memory Allocations

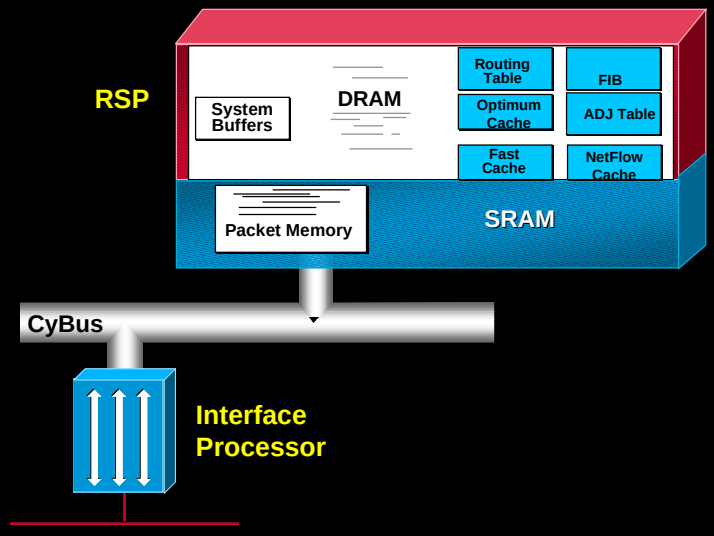

Cisco 7500 Series—Process Switching

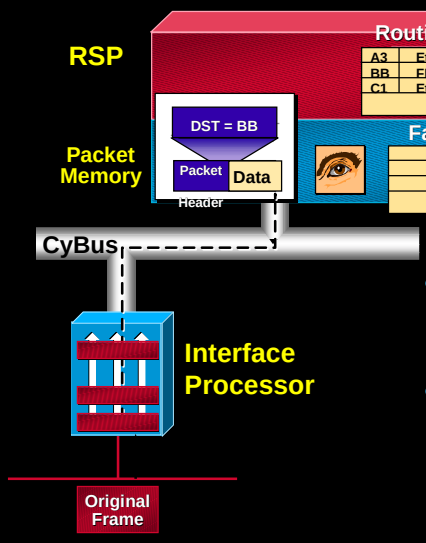

- Packet copied to packet memory
- No network/host destination in fast switching cache

Cisco 7500 Series—Process Switching

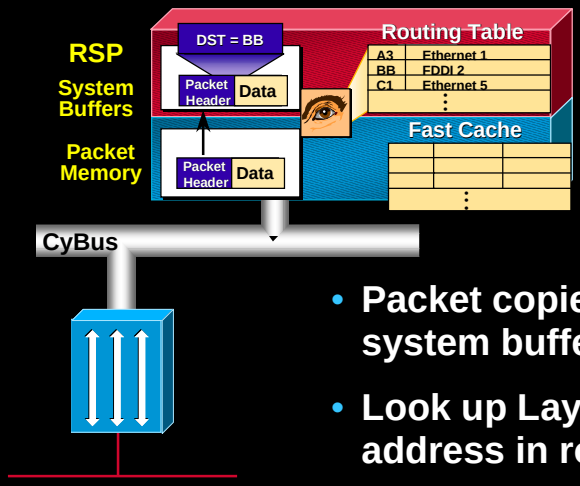

- Packet copied to system buffer
- Look up Layer 3 network address in routing table
- Initialize fast switch cache

Cisco 7500 Series—Process Switching

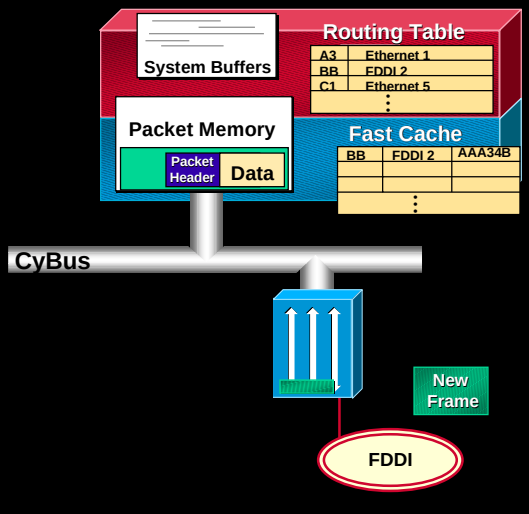

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

Cisco 7500 Series—Fast Switching

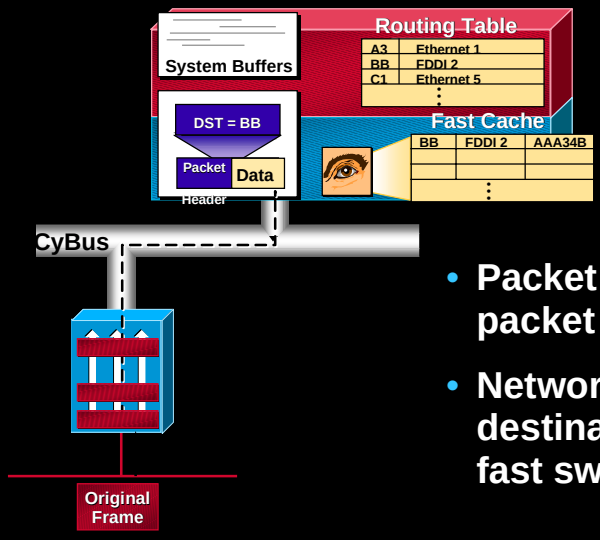

- Packet copied to packet memory
- Network/host destination found in fast switching cache

www.cisco.com

Cisco 7500 Series—Fast Switching

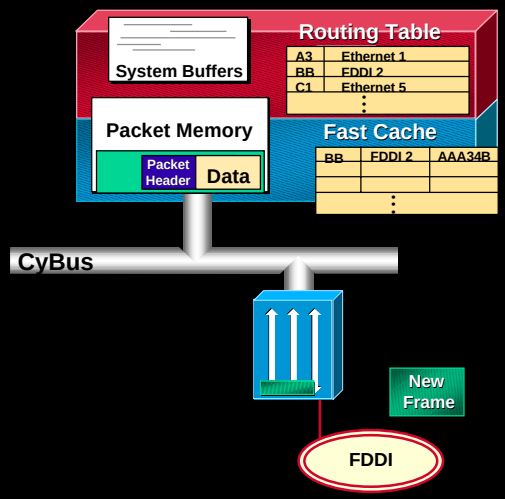

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

www.cisco.com

Optimum Switching

- Extension of fast switching
- Optimum cache
- Optimized for higher performance
- Same issues
- Deprecated in Cisco IOS 12.0

www.cisco.com

Cisco 7500 Series—Optimum Switching

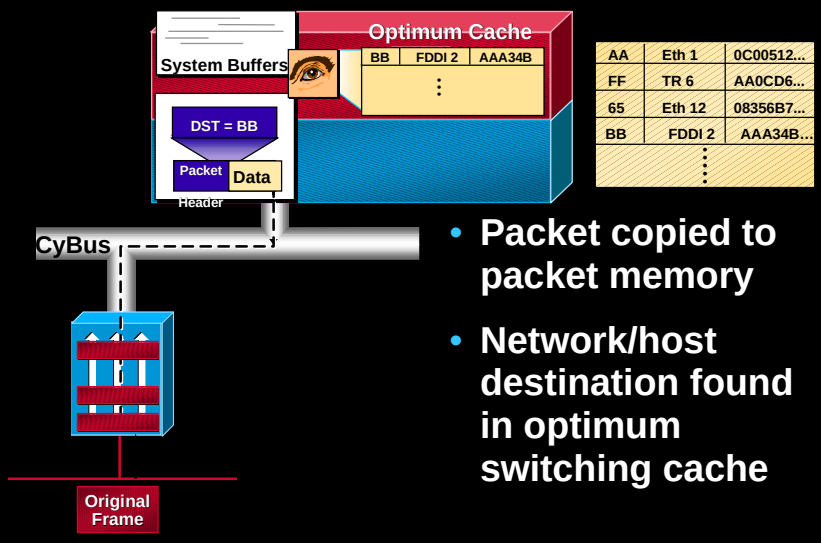

- Packet copied to packet memory
- Network/host destination found in optimum switching cache

www.cisco.com

Cisco 7500 Series—Optimum Switching

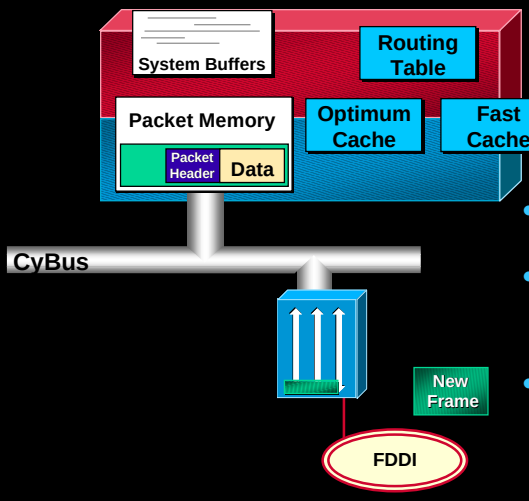

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

www.cisco.com

Cisco 7500 Series—NetFlow Switching

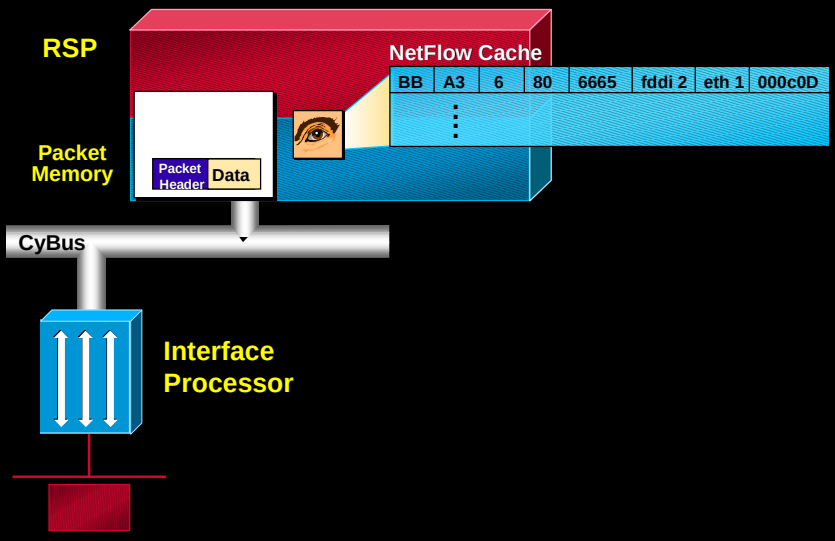

www.cisco.com

Definition of NetFlow

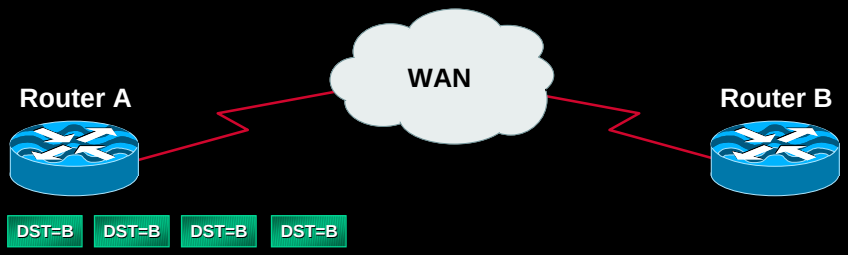

- A unidirectional sequence of packets between a given source and destination

www.cisco.com

NetFlow Granularity Data Encapsulation

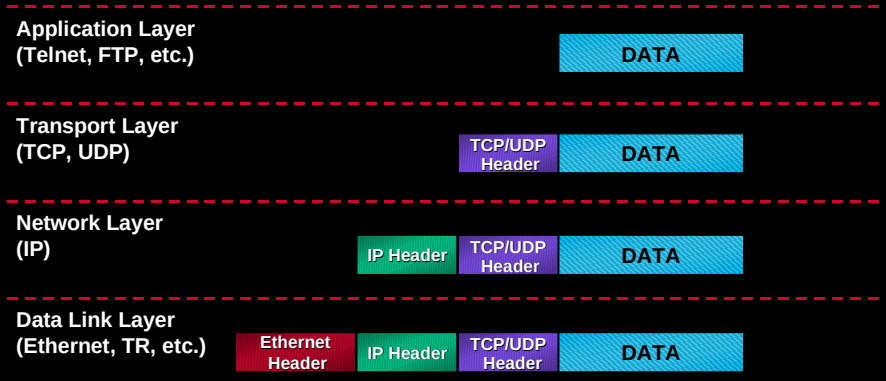

www.cisco.com

Cisco 7500 Series— NetFlow Switching

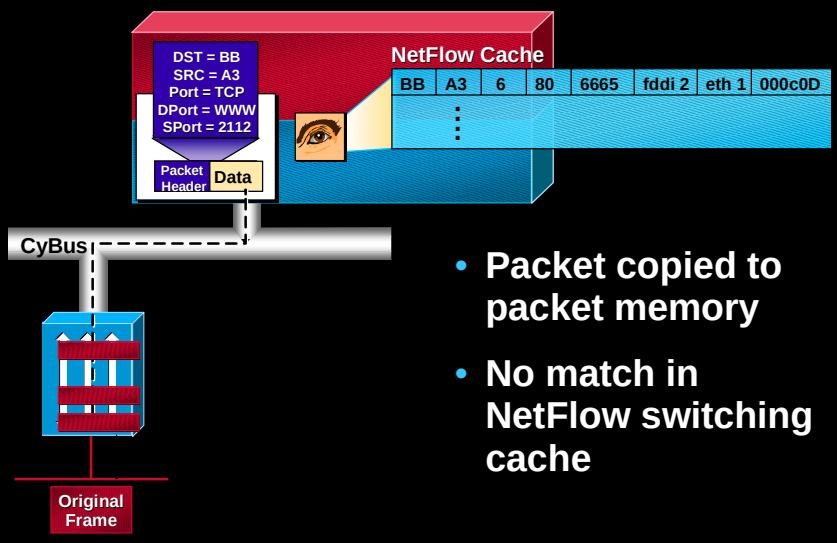

- Packet copied to packet memory
- No match in NetFlow switching cache

www.cisco.com

Cisco 7500 Series— NetFlow Switching

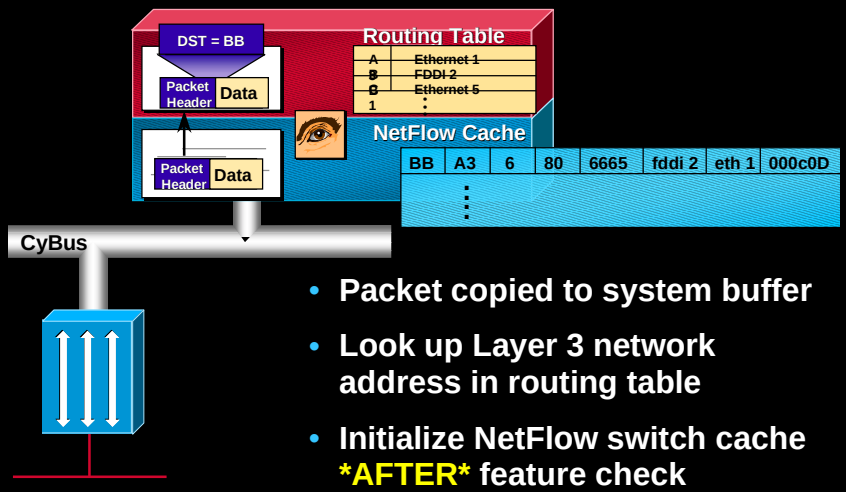

- Packet copied to system buffer
- Look up Layer 3 network address in routing table
- Initialize NetFlow switch cache ***AFTER*** feature check

www.cisco.com

NetFlow Switching

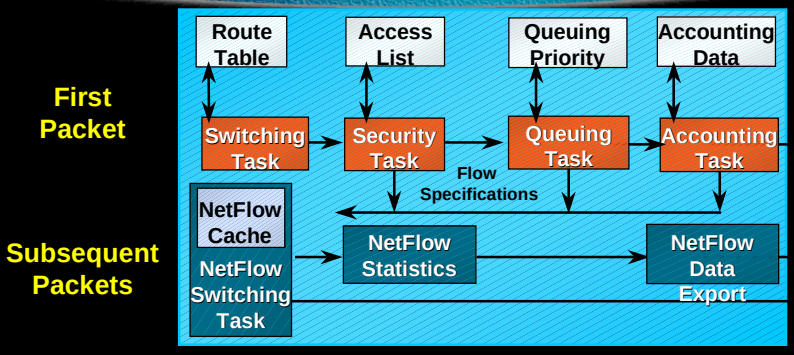

- Only first packet is processed by multiple tasks
- Connection-oriented NetFlow is defined with specific service requirements based on source/destination network address and transport layer port numbers
- Single switching task applies network services and collects traffic statistics

www.cisco.com

Cisco 7500 Series— NetFlow Switching

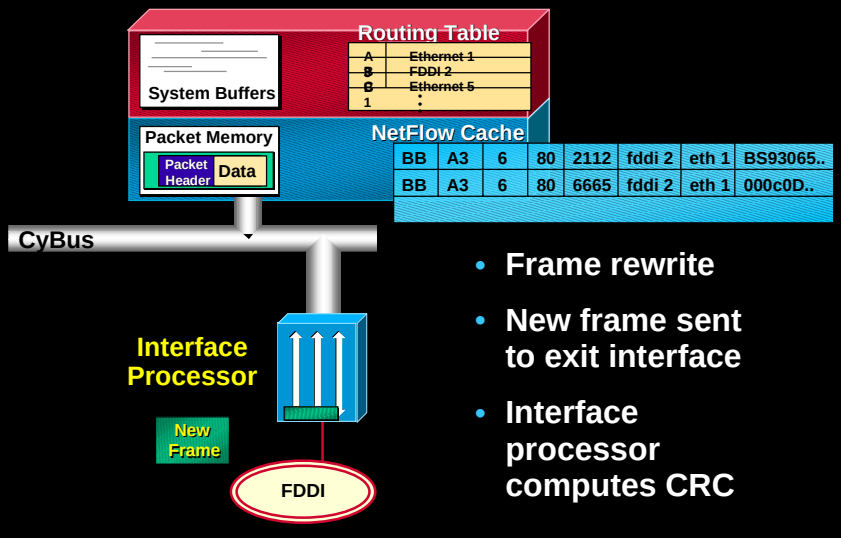

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

www.cisco.com

Cisco 7500 Series— NetFlow Switching

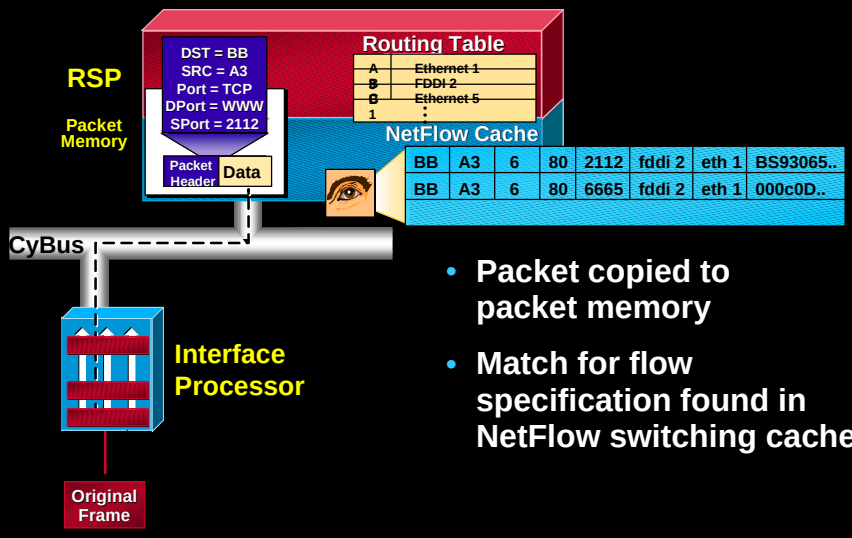

- Packet copied to packet memory
- Match for flow specification found in NetFlow switching cache

www.cisco.com

Cisco 7500 Series— NetFlow Switching

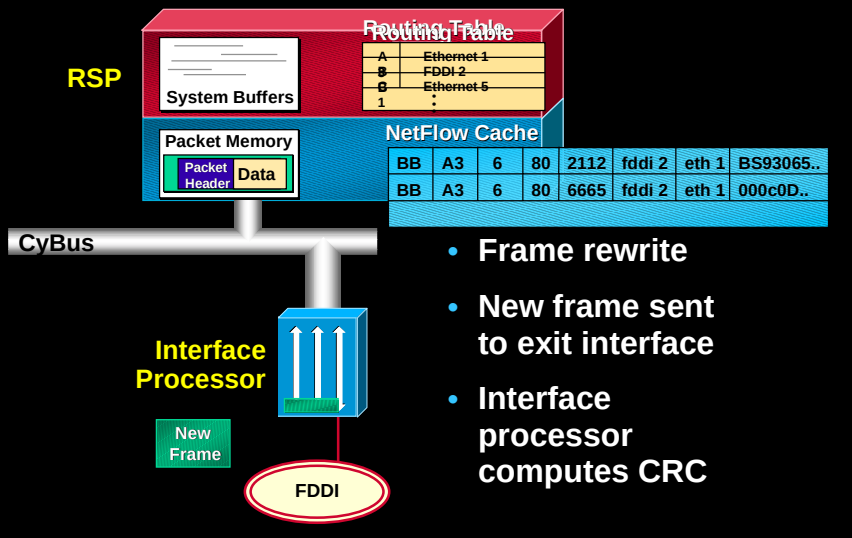

- Frame rewrite
- New frame sent to exit interface
- Interface processor computes CRC

www.cisco.com

NetFlow Services—Manageability/Accounting

IP NetFlow Switching Cache, 29999 Active, 2769 Inactive, 58411388 added Statistics Cleared 141949 Seconds Ago

Protocol	Total Flows	Flows/ Sec.	Packets/ Flow	Bytes/ Pkt	Packets/ Sec.	Active Sec/ Flow	Idle Sec/ Flow
TCP—Telnet	267,034	1.8	233	75	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	32,107,858	226.2	15	247	3610.6	13.5	28.1
SMTP	3,526,231	24.8	13	159	323.1	10.2	23.6
X	9,600	0.0	121	129	8.2	148.2	55.1
BGP	111,096	0.7	14	77	11.5	229.2	61.1
Other	5,729,172	40.3	70	220	2858.1	71.0	41.3
UDP—TFTP	2,398	0.0	3	62	0.0	13.4	69.5
DNS	12,875,077	90.7	2	110	195.4	5.4	43.6
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	58,381,400	411.3	20	227	8579.4	0.0	0.0

- Extensive statistics maintained on L3 device
- Snapshot summary traffic characterization

NetFlow Policy Routing (NPR)

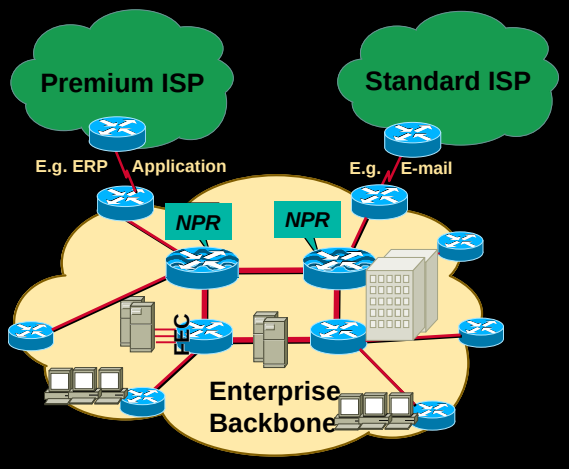

- Powerful traffic engineering tool
- ISP and/or Application selection
- Distributed performance and flow acceleration
- IP Precedence based QoS

NetFlow Policy Routing

- NetFlow calls policy routing at flow creation time
- Policy routing checks if the packet needs to be policy routed
- If not convey that to NetFlow
- Subsequent packets for the flow don’t go through policy routing, hence minimize forwarding overhead

www.cisco.com

NetFlow Policy Routing

- If it requires policy routing, policy route it and attach state to the flow
- Subsequent packets for the flow are fast-policy routed using the attached state
- Distributed version for increased performance

www.cisco.com

NetFlow Feature Acceleration

- NetFlow is leveraged to accelerate specific Cisco IOS features, including:
  - NetFlow Policy Routing (NPR)
  - Router-based Network Data Encryption
  - Network Address Translation (NAT)
  - Committed Access Rate (CAR)
  - Web Cache Control Protocol (WCCPv2)
  - Others

www.cisco.com

RSP2 Optimum vs. NetFlow Switching

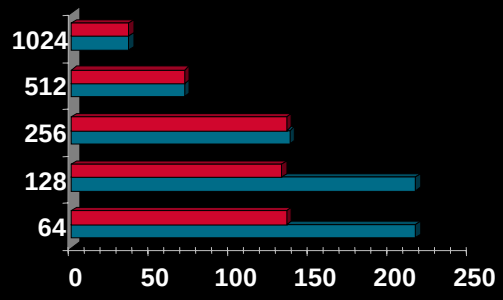

	64	128	256	512	1024
Flow	135	132	135	71	36
Optimum	216	216	137	71	36

Note: NDR stands for No Drop Rate

CPU% stands for RSP CPU Utilization

www.cisco.com

RSP2 Optimum vs. NetFlow with Export

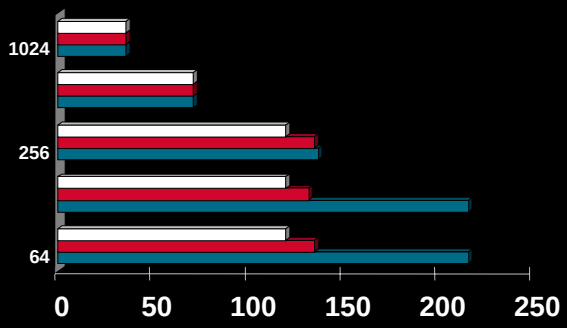

	64	128	256	512	1024
Flow W/ Export	120	120	120	71	36
Flow	135	132	135	71	36
Optimum	216	216	137	71	36

www.cisco.com

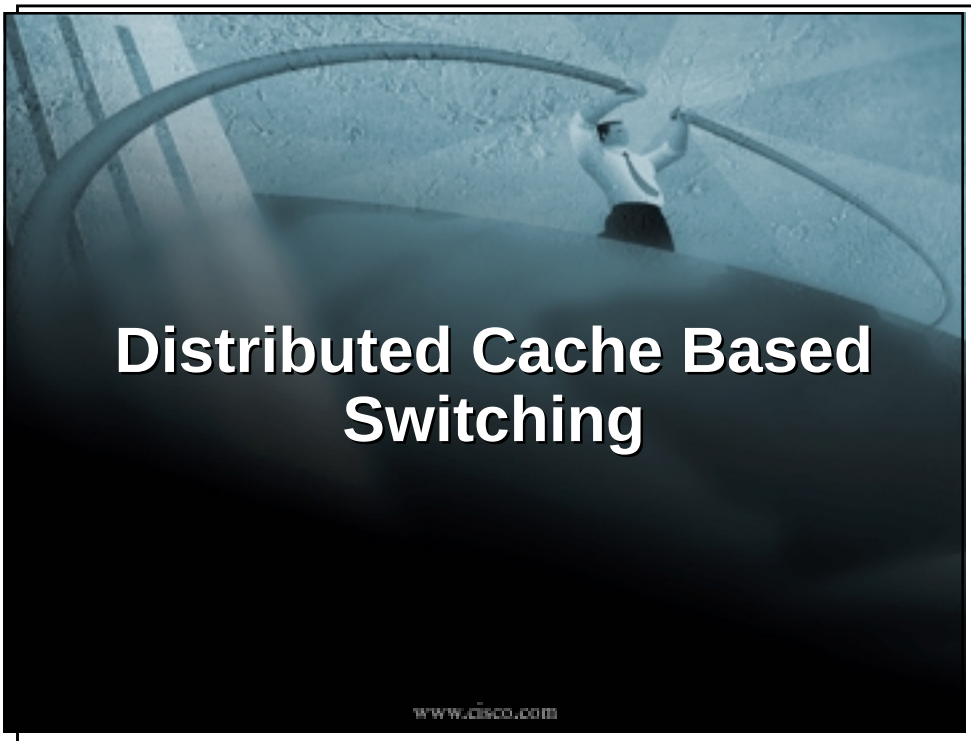

Cisco 7500 Series Distributed Switching VIP Distributed/Local Switching

- 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

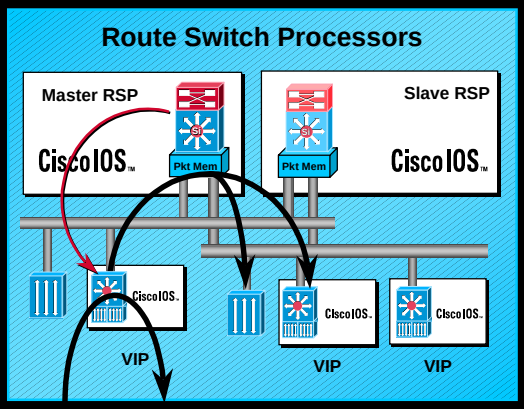

Cisco Express Forwarding

- New topology driven architecture
- Main components
  - Forwarding Information Base (FIB)
  - Adjacency table
- No process switching of packets

www.cisco.com

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

www.cisco.com

FIB Table

- Shadow copy of the IP routing table
- Classless
- Routing protocol independent
- One for each route in IP routing table
- Each entry has one or more path
- Each path has nexthop IP address and nexthop interface
- Each path points to an adjacency

www.cisco.com

Adjacency Table

- Maintains IP address to Mac-rewrite mapping
- Populated by ARP table, Frame Relay map table and ATM map table, etc.
- Mac-rewrite of the nexthop is all that's required to switch packet

www.cisco.com

Cisco Express Forwarding: Routing Table

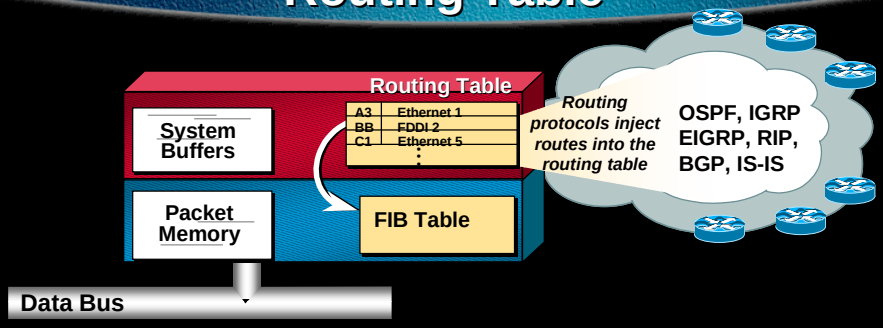

- The routing table contains a list of reachable networks
- The routing information is passed down to the Forwarding Information Base (FIB) manipulated, into an optimized software look-up tree (mtrie)

www.cisco.com

Cisco Express Forwarding: Forwarding Table

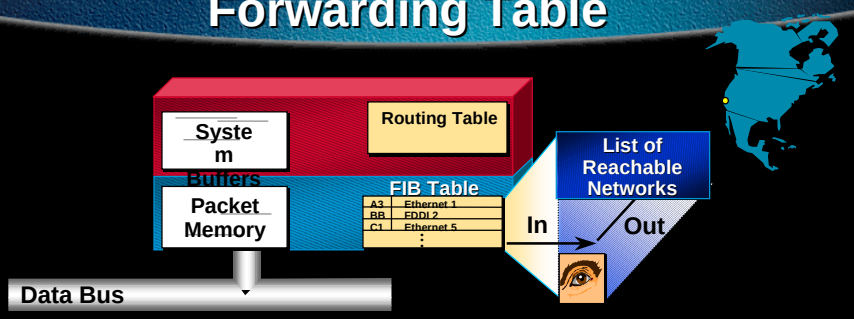

- The FIB is a database of information which is used when making a packet-forwarding decision
  - Updated when a route is added, removed, or changed in the routing table
  - The FIB s/w tree contains a list of destination IP prefixes
  - Pointer to the adjacency database

Cisco Express Forwarding: Adjacency Table

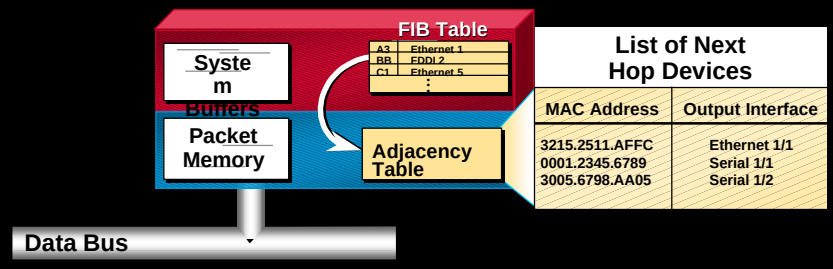

- A list of adjacent nodes and their status up/down
- Adjacency table is populated as it discovers adjacent nodes
  - On the creation of an adjacency, a link-layer header to reach the node is computed and stored in the database

Distributed CEF—For Example, dCEF on the Cisco 7500

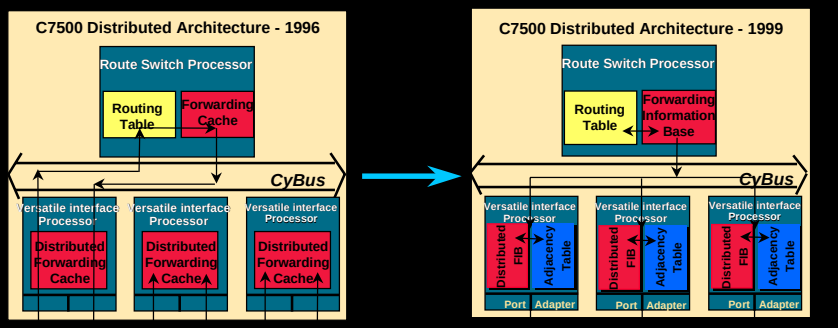

First Packet 'Process Switched'

Subsequent Packets 'Fast Switched'

All packets forwarded via switching line cards

Adjacency Table:

- IP Address L2/MAC mapping
- Populated by ARP
- Maintains counters/statistics
- Adjacencies defines as Normal, Null, Receive, Punt, Incomplete etc.

Forwarding Information Base:

- Constructed by routing process
- Foundation for Tag Information Base (TIB)

www.cisco.com

CEF Load-Sharing

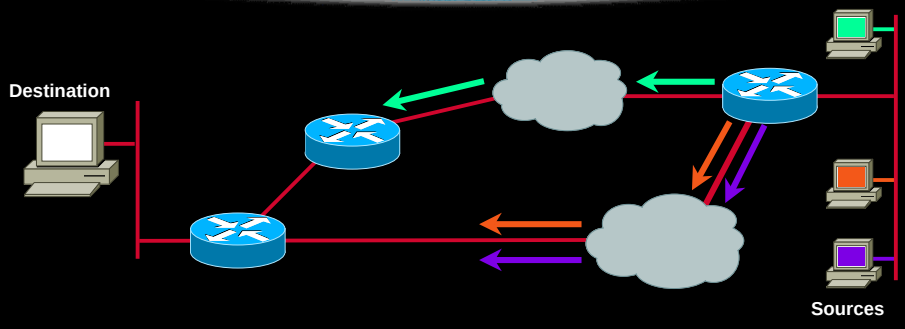

- Per packet and enhanced per destination
- Enhanced per destination is based on source and destination IP addresses
- Each destination flow takes a single, separate path
- Reduces need for per packet load-sharing

www.cisco.com

CEF and QoS

- CAR classifies packets into classes
- Polices within each class as “in” and “out” of profile
- Mark “in” with higher precedence than “out”
- Can be used to drop “out” packets

www.cisco.com

CEF and QoS

- WRED can manage queues that develop in the routers
  - Prioritizes “in” traffic over the rest
- WFQ can be used to allocate bandwidth to each CoS
- Unused bandwidth from a CoS available for others

www.cisco.com

Agenda

- Measuring Performance
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

www.cisco.com

Features Affecting Performance

- Multicast
- Queuing
- Compression
- Filtering
- Encryption

www.cisco.com

Multicast Buffer Cloning

- For Multicast we only need to replicated the Packet Header Particle. All “Data” particles need not be replicated.
- The Buffer Header Pointers associate the cloned header particle with the data particles.

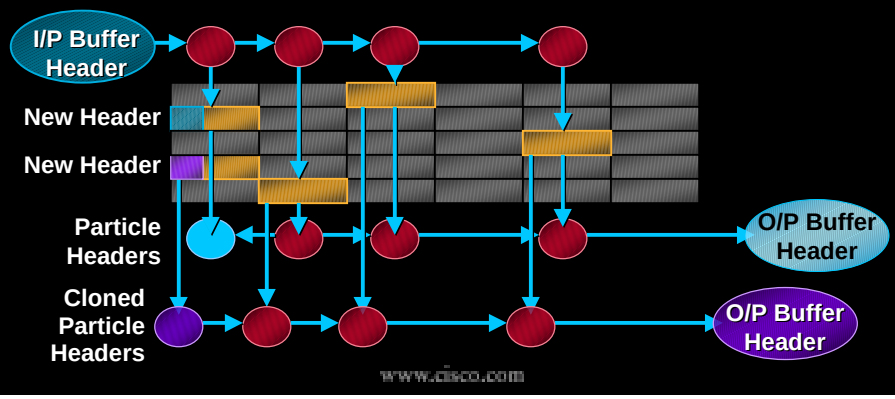

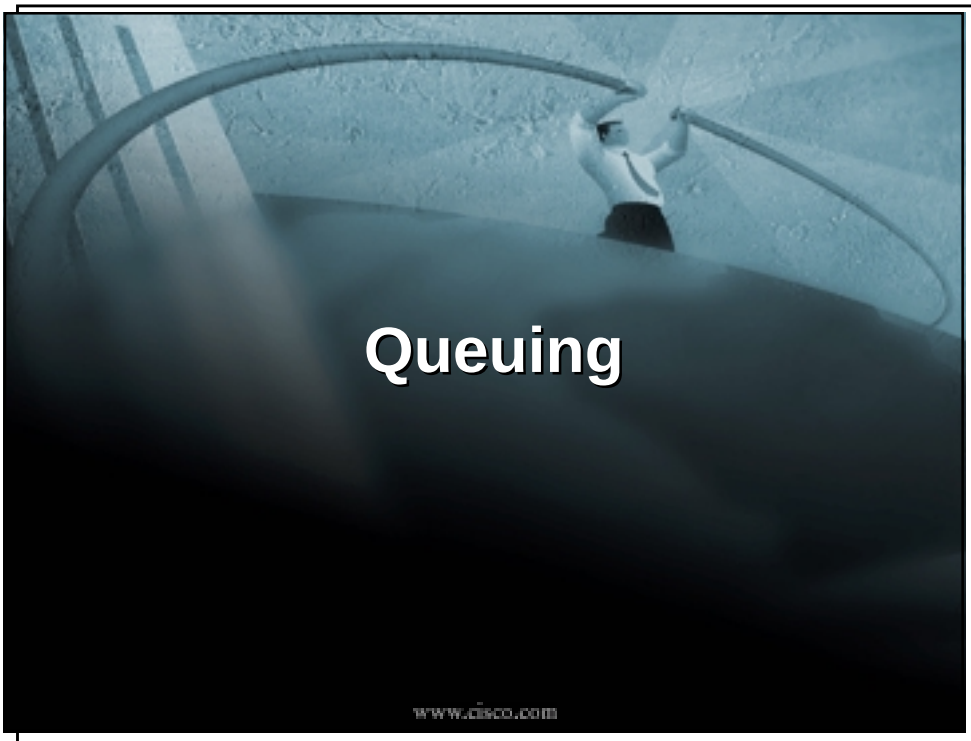

Queuing

- 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

www.cisco.com

Priority Queuing

- Four queues: high, medium, normal and low
- Higher level queue starves the queues below

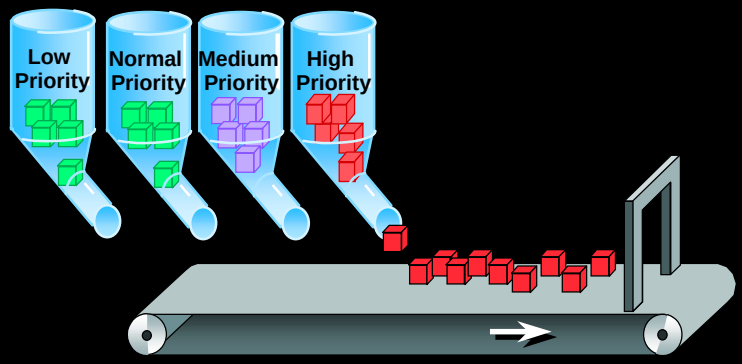

Custom Queuing

- Control % of interface bandwidth for specified traffic
- 17 output queues for each interface [16 configurable]

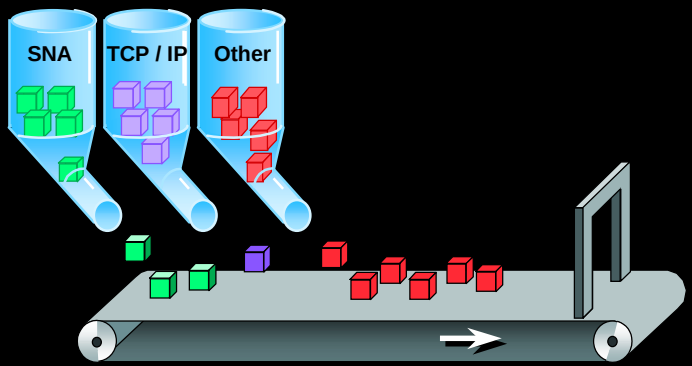

Weighted Fair Queuing

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

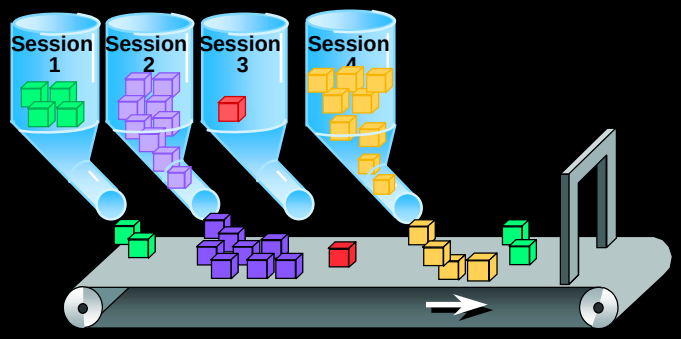

www.cisco.com

Processes and Functions

- Dequeuing is the function used when a packet has been process switched and needs to be passed to the output interface, or when an interface has “fancy” queuing enabled and again packets are stored in system buffers until they are dequeued (based on queuing mechanism)

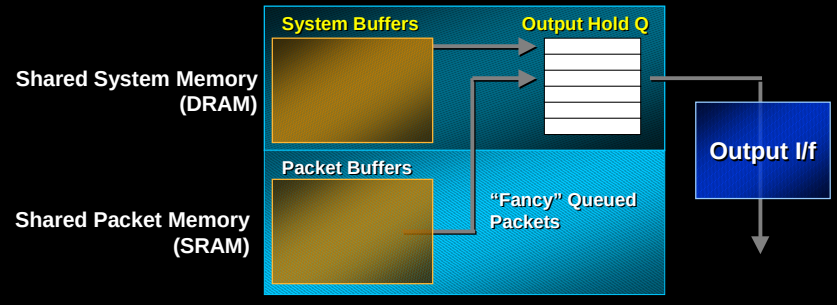

www.cisco.com

Processes and Functions

- In the case of “fancy” queuing being applied to an interface, the Output Hold Queue represents either the (4) Priority, (>16) Custom or (>256) Weighted Fair queues
- In Process switching this is a FIFO queue by default

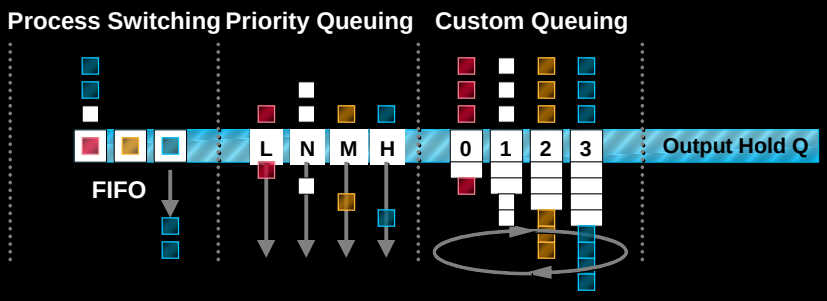

www.cisco.com

Compression

www.cisco.com

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

ACL Development

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

www.cisco.com

RSP2 ACl's with Optimum vs. NetFlow

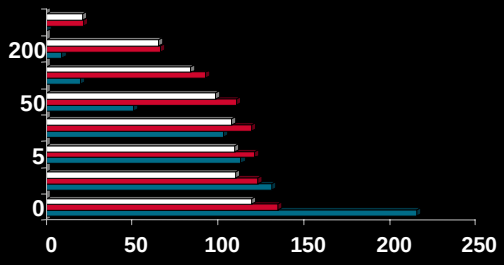

	0	1	5	10	50	100	200	900
Flow w/ Export	120	110.63	109.6	108	99	84	65.25	21.19
Flow	135	123.38	121.53	120	111	93	66.75	21.75
Optimum	216	131.25	113.63	103.5	51	20.25	9.19	0.46

RSP4 ACl's with Optimum vs. NetFlow

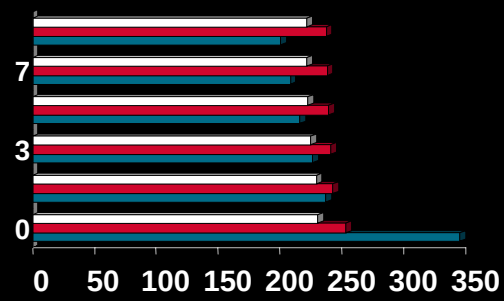

	0	1	3	5	7	9
Flow with Export	231.02	229.5	225	222.6	222	221.85
Flow	253.51	243.01	241.51	239.56	238.5	237.76
Optimum	345.77	237	226.5	216.01	208.5	201.1

(Note: this test is done for 64 byte packets only)

RSP2 ACl's with Optimum vs. CEF vs. NetFlow

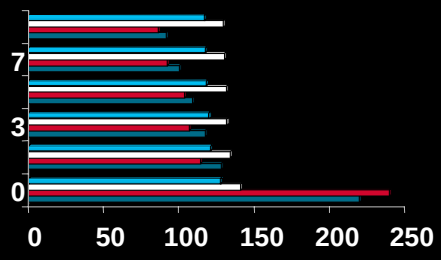

	0	1	3	5	7	9
CEF_Flow w/ Export	128	121	120	118	118	117
CEF_Flow	141	134	132	132	131	129
CEF	240	114	107	104	93	86
Optimum	221	128	118	109	101	92

(Note: this test is done for 64 byte packets only)

www.cisco.com

IPSec VPN Router Performance

- Performance tests done by EWAN BU and MSABU in Feb'99
- Includes Cisco 7200/NPE300, 7200/NPE200, 3640, 3620, 2620, 2610, and 1720

www.cisco.com

Performance Test Setup

- Cisco IOS 12.0
- Two identical routers connected back-to-back
- Via serial link clocked at
  - 4 Mbps for 7200
  - 1.5 Mbps for all others
- Full-duplex traffic from SmartBits

Performance Test Coverage

- Ethernet pkt sizes: 64, 128, 256, 512, 1024 bytes
- IPSec ESP encryption: DES, 3DES
- IPSec ESP authentication: SHA, MD5
- Crypto access-list length: 1, 25

Cisco 1720

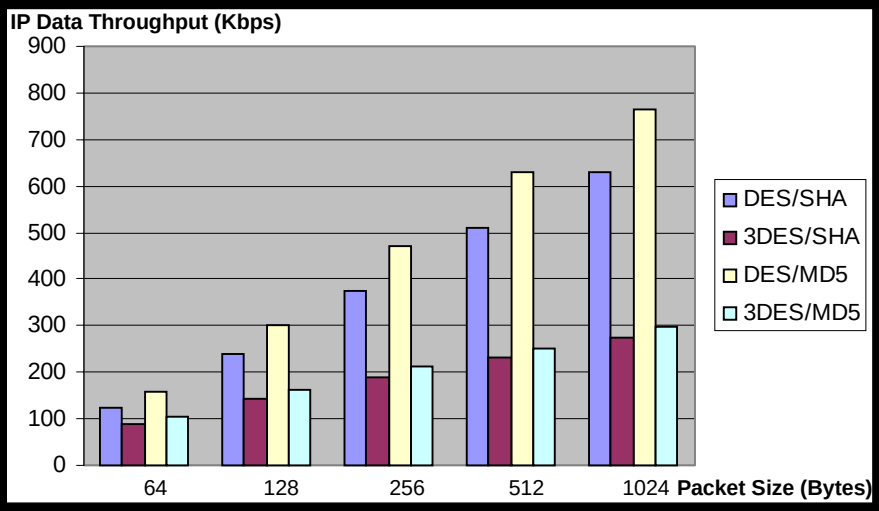

www.cisco.com

Cisco 2610

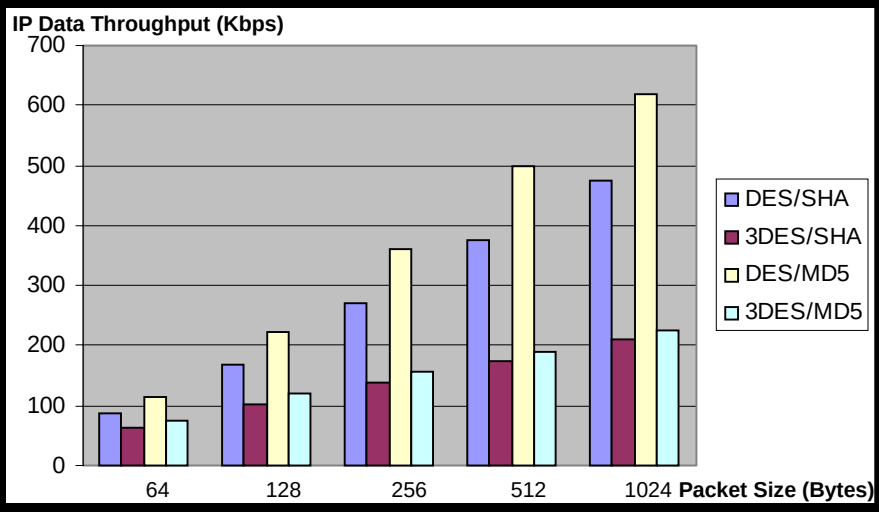

www.cisco.com

Cisco 2620

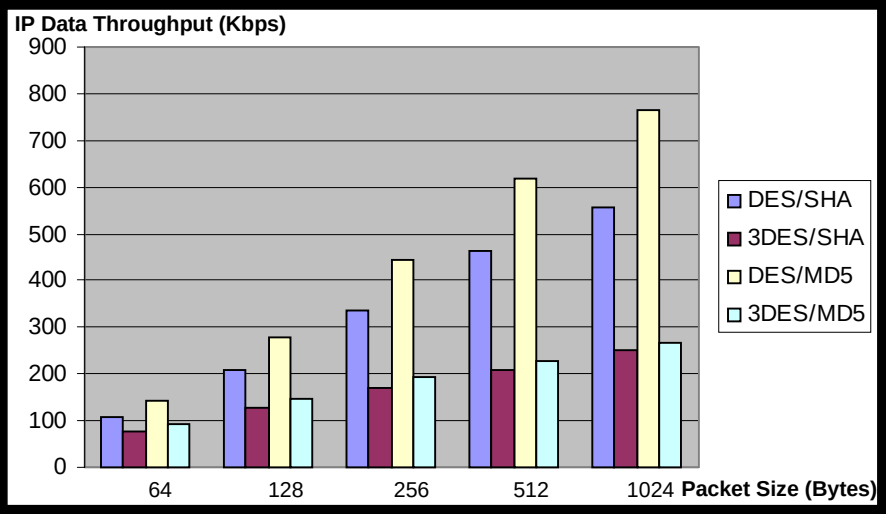

www.cisco.com

Cisco 3620

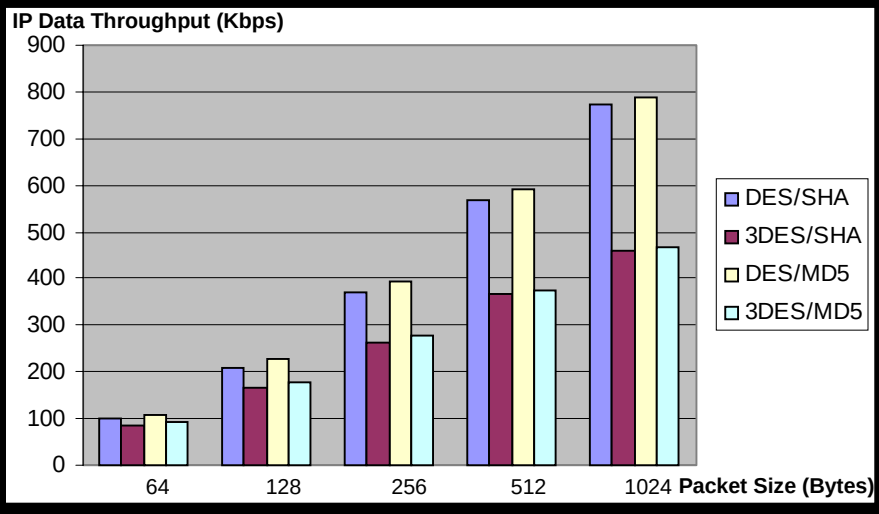

www.cisco.com

Cisco 3640

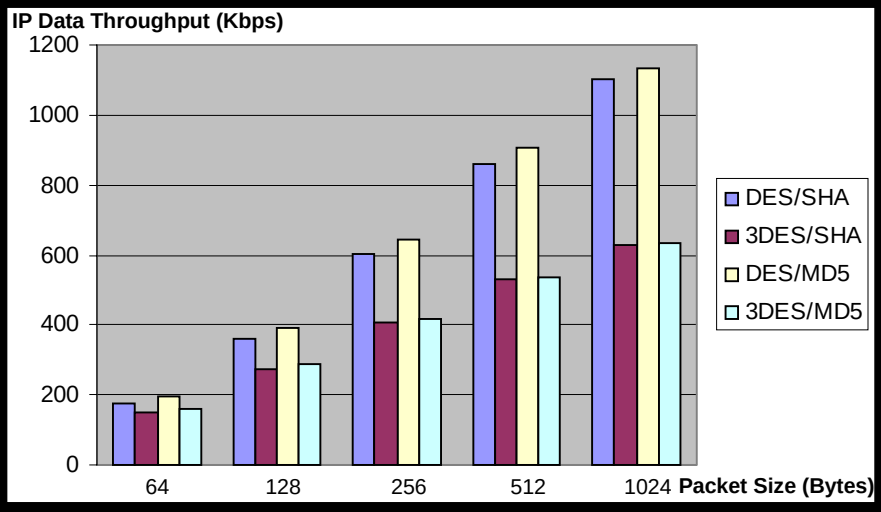

www.cisco.com

Cisco 7200/NPE200

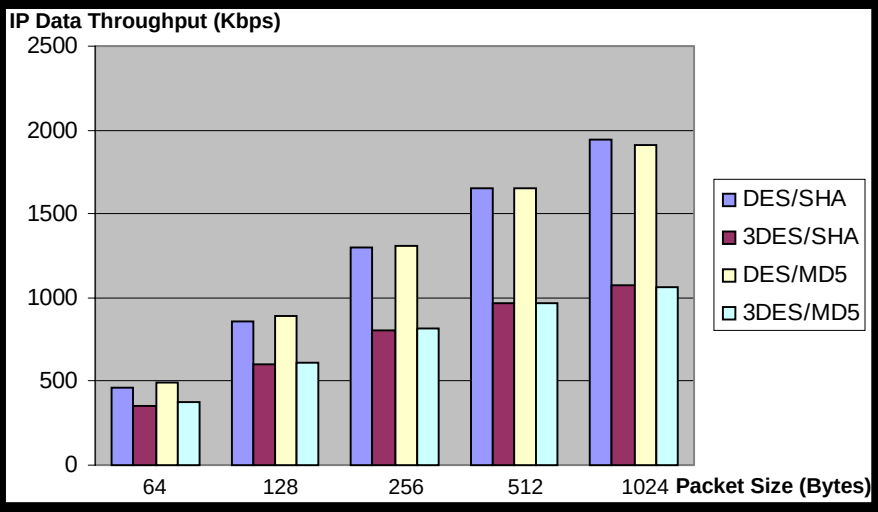

www.cisco.com

Cisco 7200/NPE300

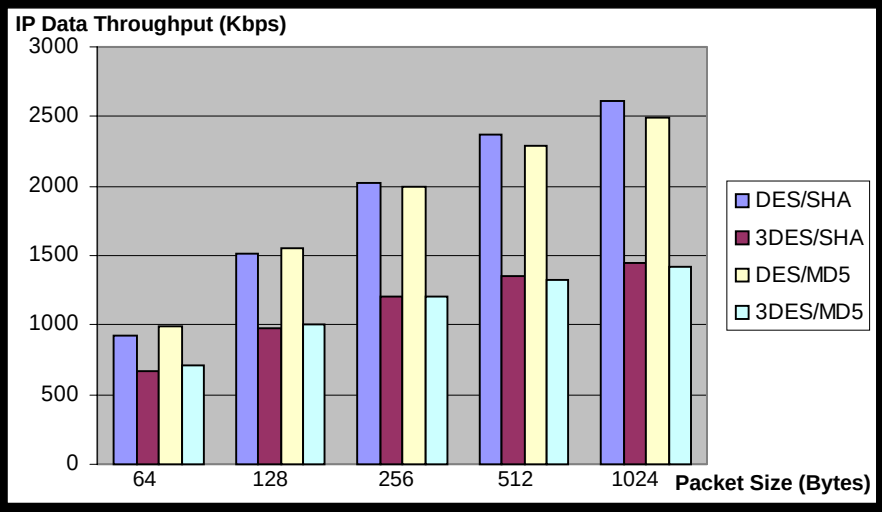

www.cisco.com

Performance Test Results

- Charts show “one-way” throughput (total traffic through router is twice as much)
- Throughput based on end-to-end IP packets (i.e. IPSec and MAC overhead not counted)
- Charts show results for crypto access-list length of 1 (less than 5% degradation for access-list length 25)

www.cisco.com

Performance Observations

- Throughput improves with packet size because of reduced per-packet IPSec overhead
- Results verify 3DES more CPU-intensive than DES, and SHA more CPU-intensive than MD5
- DES vs. 3DES performance difference is not “triple”

Agenda

- Measuring Performance
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

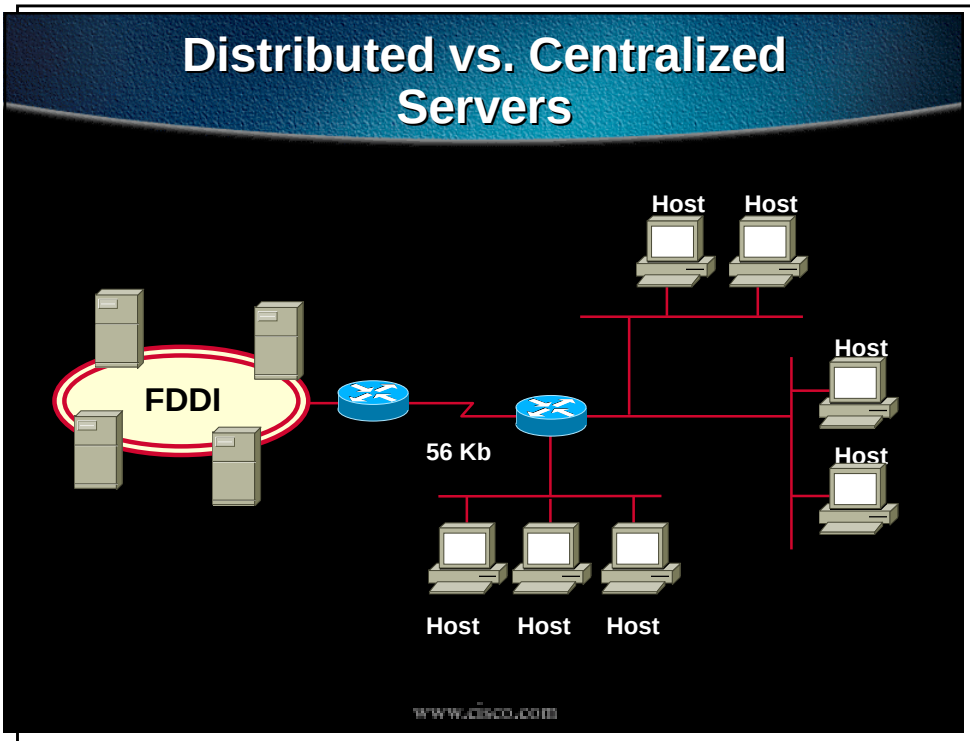

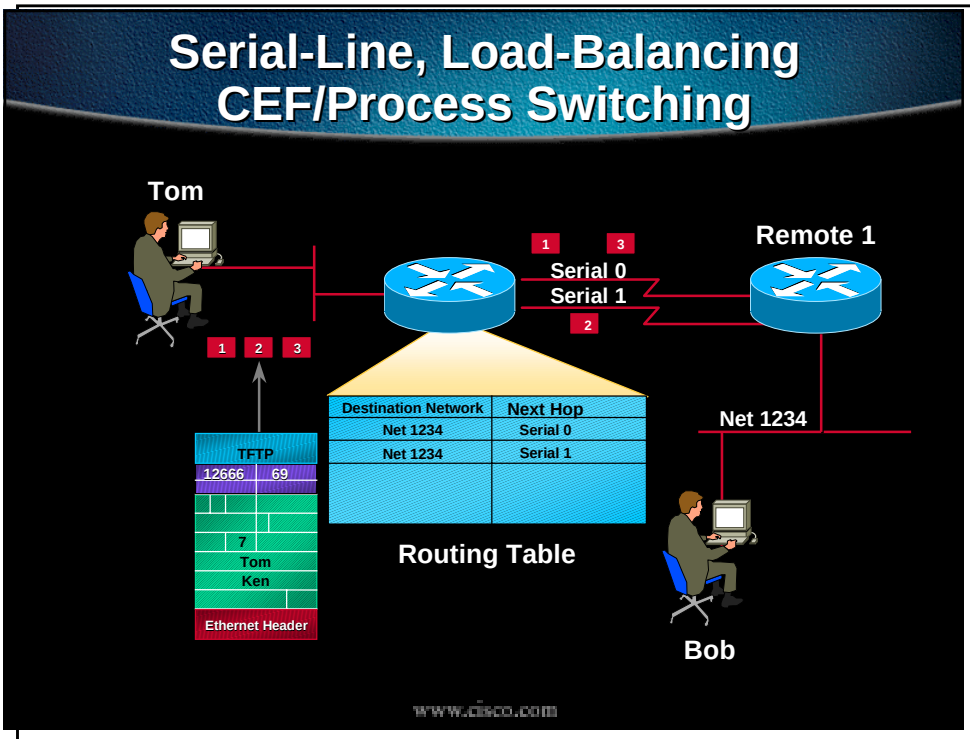

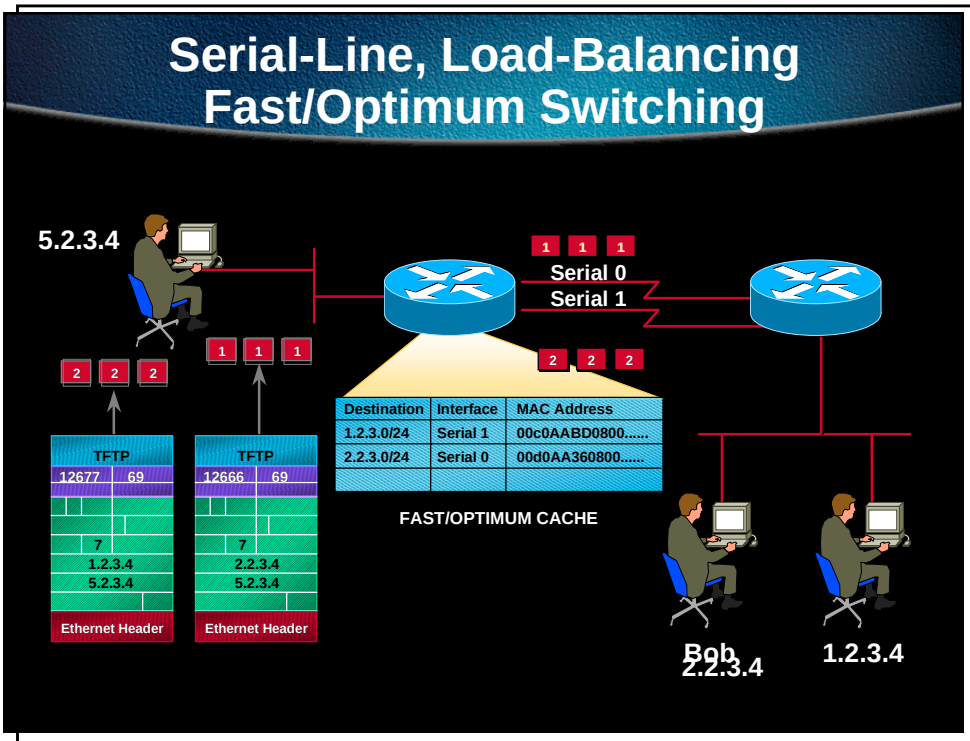

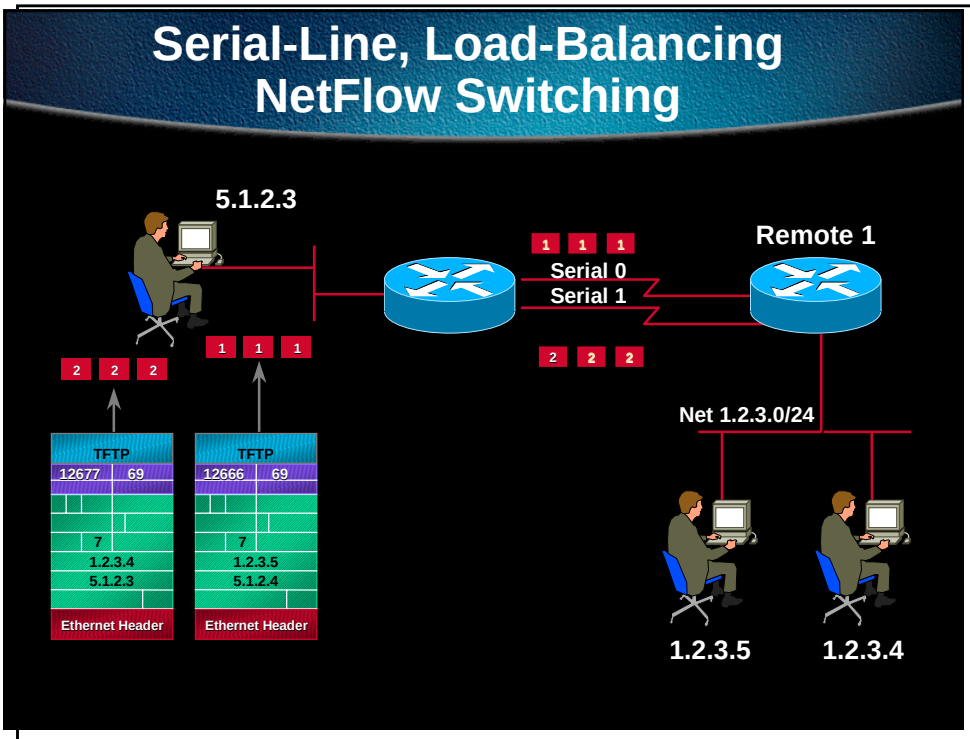

Agenda

- Measuring Performance
- Router Architectures/Switching Paths
- Features Affecting Performance
- Optimized Network Design
- Troubleshooting

www.cisco.com

Receiving Packet

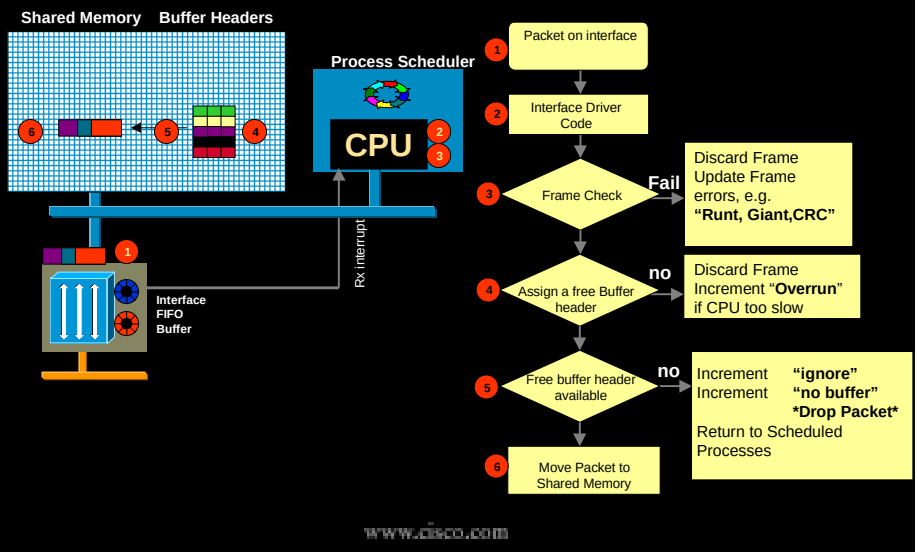

www.cisco.com

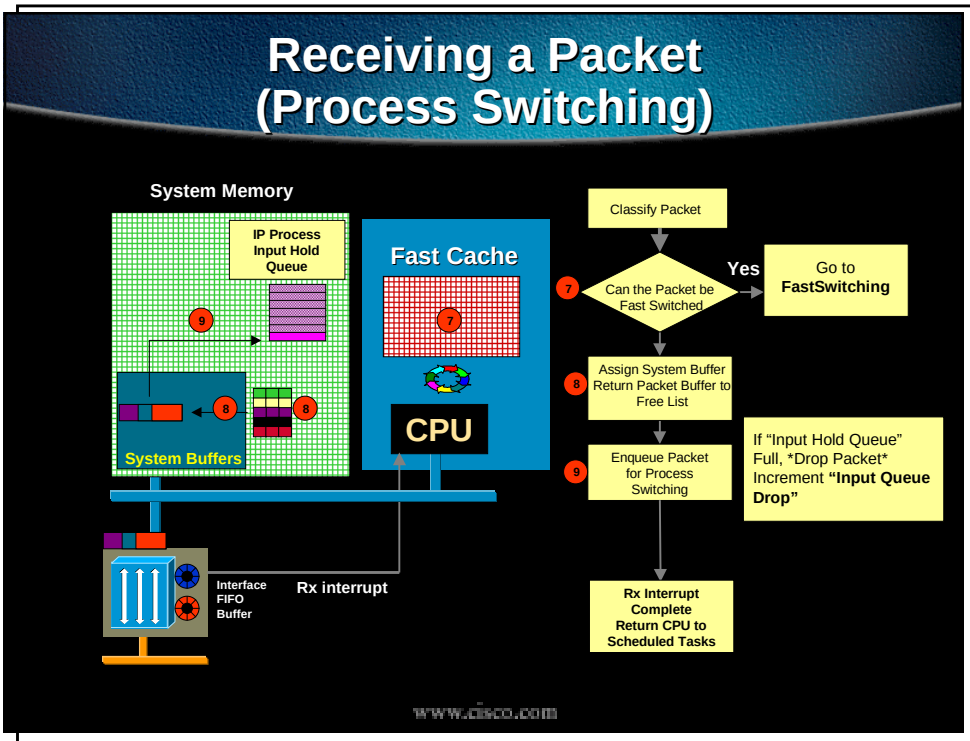

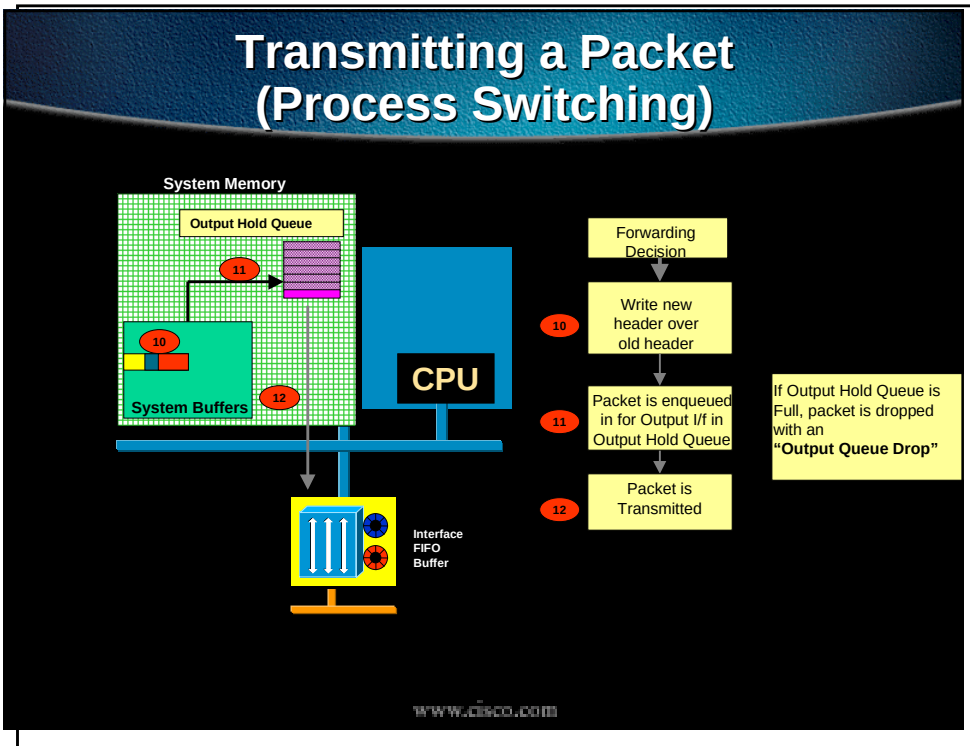

Transmitting a Packet (Fast Switching)

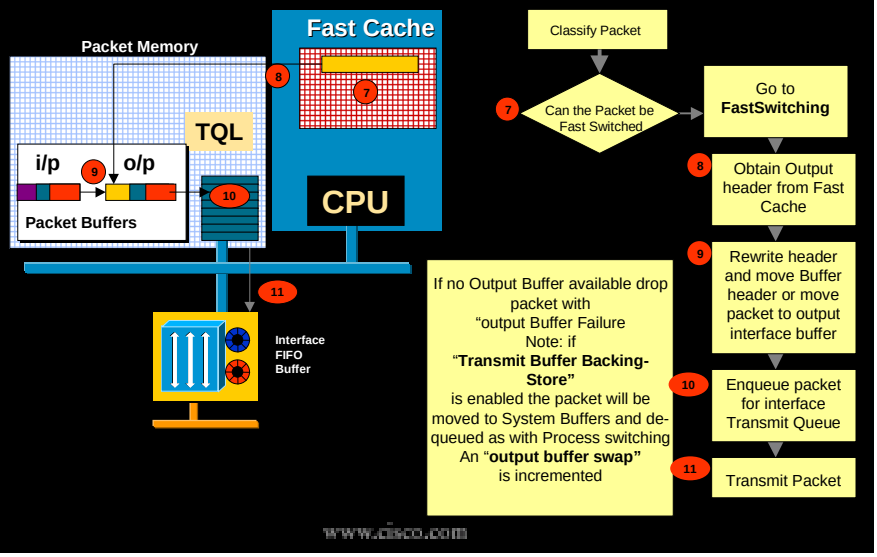

Switching Paths Used

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

www.cisco.com