


Objective

To understand the how to measure a device's network ability to protect application networking resource performance **and** scalability.

Agenda


- Definition of Terms
- Performance Test Tools
- Testing Methods and Metrics
- Example—Buffering and Architecture
- Example—Flow Lookup and Architecture
- Example—WRR and Architecture
- Example—SLB and Architecture

Definition of Terms




- **Device under test (DUT)**—A single stand-alone or modular unit which receives frames on one or more of its interfaces and then forwards them to one or more interfaces according to the addressing information contained in the frame

Definition of Terms (Cont.)




- **System Under Test (SUT)**—The collective set of network devices to which stimulus is offered as a single entity and response measured

Definition of Terms (Cont.)



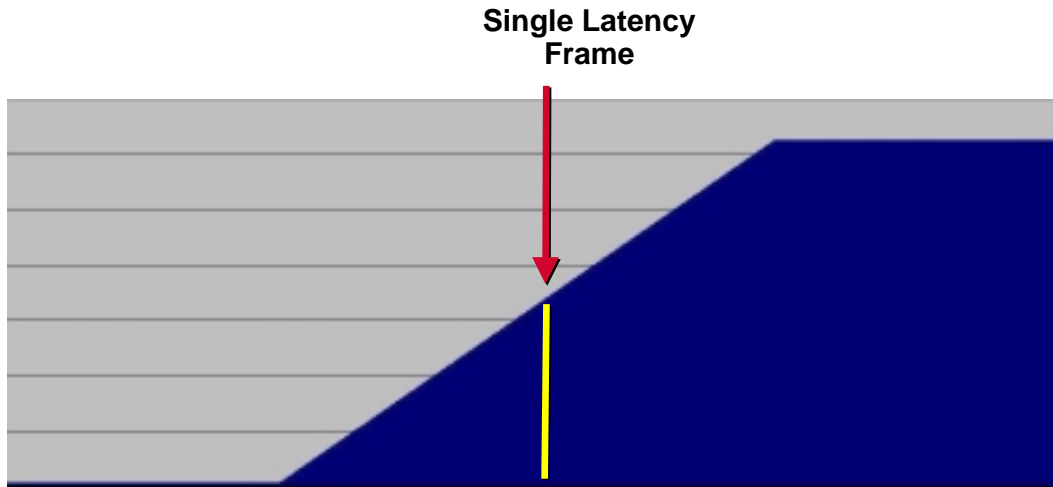
- **Queuing**—buffer mechanisms used to control congestion by temporarily storing packets in memory

Definition of Terms (Cont.)




- **Latency (RFC 1944)**—Time in microseconds which a DUT takes to transmit one packet with same size packet at X rate

Definition of Terms (Cont.)



- **Latency (RFC 1944)**—Time in microseconds which a DUT takes to transmit one packet with same size packet at X rate

Definition of Terms (Cont.)




All Frames tagged
with latency and
sequence

- **Latency (RFC 2285)**—Time in microseconds which a DUT takes to transmit all packets with sequence and latency tags



Definition of Terms (Cont.)

All Frames tagged
with latency and
sequence




- **Latency (RFC 2285)**—Time in microseconds which a DUT takes to transmit all packets with sequence and latency tags

Definition of Terms (Cont.)




- **Throughput**—Highest rate in packets-per-second through a DUT which frame loss does not occur

Definition of Terms (Cont.)



- **Frame loss**—Percentage of lost frames when running at said frame size often confused with throughput

Definition of Terms (Cont.)



- **Back-to-back burst** —The minimum interval between bursts which the DUT can process with no frame loss

Definition of Terms (Cont.)



128-->1518+ Layer 7

48 bytes Layer 4

32 bytes Layer 3

12 bytes Layer 2


- **Depth of inspection—**
Amount of packet information processed to make a switching decision

Agenda

- Definition of Terms
- Performance Test Tools
- Testing Methods and Metrics
- Example—Buffering and Architecture
- Example—Flow Lookup and Architecture
- Example—WRR and Architecture
- Example—SLB and Architecture

Test Tools


Real PC User



PC Based Testing

No Industry Standard

Application (Packet) Scripts



Application to Application Measurement

A Few Scripts

Metrics Unknown
Metrics Unknown
Metrics Unknown
Metrics Unknown

Not Wire Speed
No Accuracy

No Network Measurement

Only Round Trip Times

Testing with Appliance

Device to Device Measurement

Scales to Many Users

Setup Timing etc
Latency, Jitter etc.
No OS Delay
Throughput, Loss, etc.


Wire Speed
Nanosec Accuracy

Full Visibility of Network Elements

One Way Times, Accurate GPS timing

ML-7710

- TX INIT
- TRIG CLS
- RX ERR
- LINK DPLX
- 100M PAUSE



2802

1272_05_2000_c2


© 2000, Cisco Systems, Inc.

cisco.com

17

Test Tools


Real PC User



PC Based Testing

No Industry Standard

Fake Application (Packet) Scripts



Application to Application Measurement

A Few Scripts

Metrics Unknown
Metrics Unknown
Metrics Unknown
Metrics Unknown

Not Wire Speed
No Accuracy

No Network Measurement

Only Round Trip Times

Testing with Appliance

Device to Device Measurement

Scales to Many Users

Setup Timing etc
Latency, Jitter etc.
No OS Delay
Throughput, Loss, etc.


Wire Speed
Nanosec Accuracy

Full Visibility of Network Elements

One Way Times, Accurate GPS timing

ML-7710

- TX INIT
- TRIG CLS
- RX ERR
- LINK DPLX
- 100M PAUSE



2802

1272_05_2000_c2


© 2000, Cisco Systems, Inc.

cisco.com

18


Application Testing

- Only true application test
- User “feel” valuable but intangible
- No standard metrics
- Not wire speed



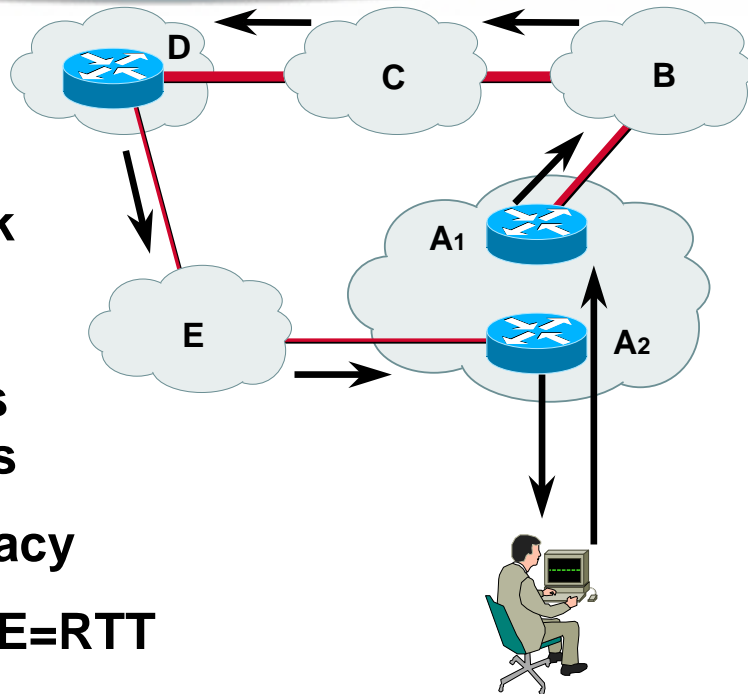
Appliance Testing

- End-to-End Latency
- Multi-user
- Multi-network GPS SYNC
- Great for L3 & L2
- OK for L4



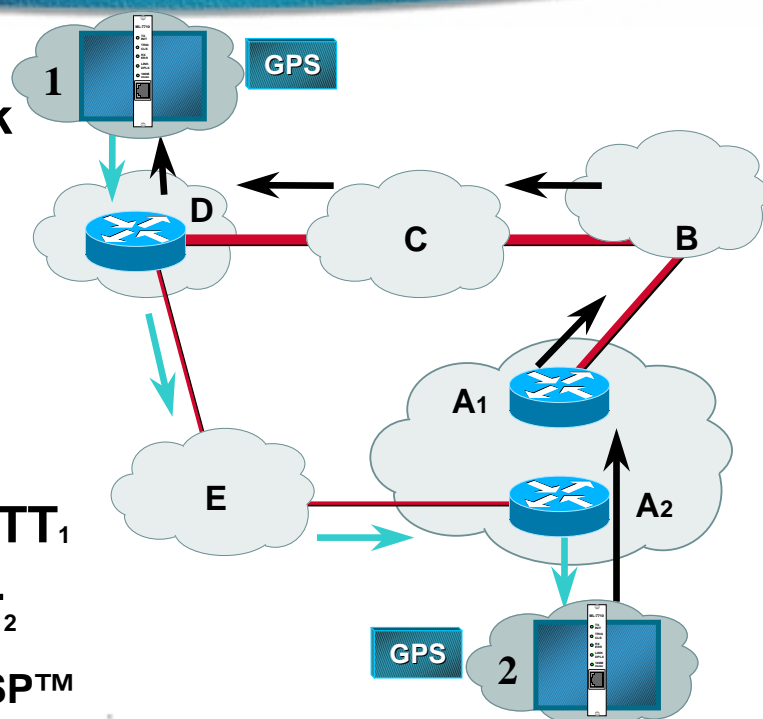
Why Ping-Pong is all Wrong

- Single user
- Single network
- Single Sample
- Asynchronous Routing Issues
- Minimal accuracy
- $A_1 + A_2 + B + C + D + E = RTT$



Why Ping-Pong is all Wrong

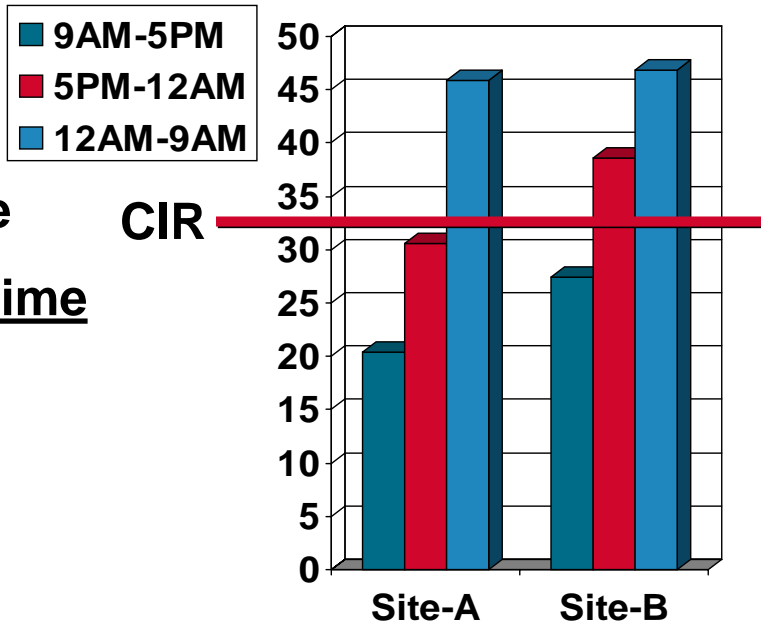
- Multi-user
- Multi-network
- Multi-Sample
- End-to-End Latency
- GPS timing
- $A_1 + B + C + D = RTT_1$
- $D + E + A_2 = RTT_2$



*Netcom's SmartISP™

Verify your Service

- Service Level Agreements
- CIR-Committed Information rate
- With Respect to time
- Latency & Deviation
- Throughput & Deviation



Performance Test Tools

PC Test Tool

- Application to Application Measurement
- A Few Scripts
- Metrics Unknown
- Metrics Unknown
- Metrics Unknown
- Metrics Unknown
- Not Wire Speed
- No Accuracy
- No Network Measurement
- Only Round Trip Times

Appliance

- Device to Device Measurement
- Scales to Many Users
- Setup Timing Etc
- Latency, Jitter Etc.
- No OS Delay
- Thru'Put, Loss, Etc.
- Wire Speed
- Nanosec Accuracy
- Full Visibility of Network Elements
- One Way Times, Accurate GPS Timing

Performance Test Tools Links

PC Test Tool

<http://www.ganamede.com/>

- LanMark Pro™

<http://www.zdnet.com/etestinglabs>

[/filters/benchmarks/](#)

- WebBench™

- NetBench™

Appliance

<http://www.netcomsystems.com/>

- SmartWindows™

- SmartFlow™

- SmartTCP™

<http://www.ixiacom.com/>

- TCLGui™



cisco.com

Application vs. Performance Test Tool

Application Tool

Not Scalable
Real World

Performance Tool

Scalable
Not Real World

Combine
and
Prosper

Types of Tests

Type of Test Typical Tests

1.Performance	Tests for delays in sending data or setting up calls Tests for variation in delay and effect on applications Tests for user application responsiveness Tests use of theoretical bandwidth
2.Capacity	Tests number of addresses/flows supported
3.Correct Operation	Tests data integrity and data loss/ RFC conformance Tests of device/network functionality under load
4.Resilience	Network ability to recover from errors or changes Map to Service Level Agreement
5.Security	Tests loss of security under load Tests loss of security based on port scanners Tests loss of security based on new attacks

Agenda


- Definition of Terms
- Performance Test Tools
- **Testing Methods and Metrics**
- **Example—Buffering and Architecture**
- **Example—Flow Lookup and Architecture**
- **Example—WRR and Architecture**
- **Example—SLB and Architecture**

Resource Constraints and Bottlenecks

Application Characteristics

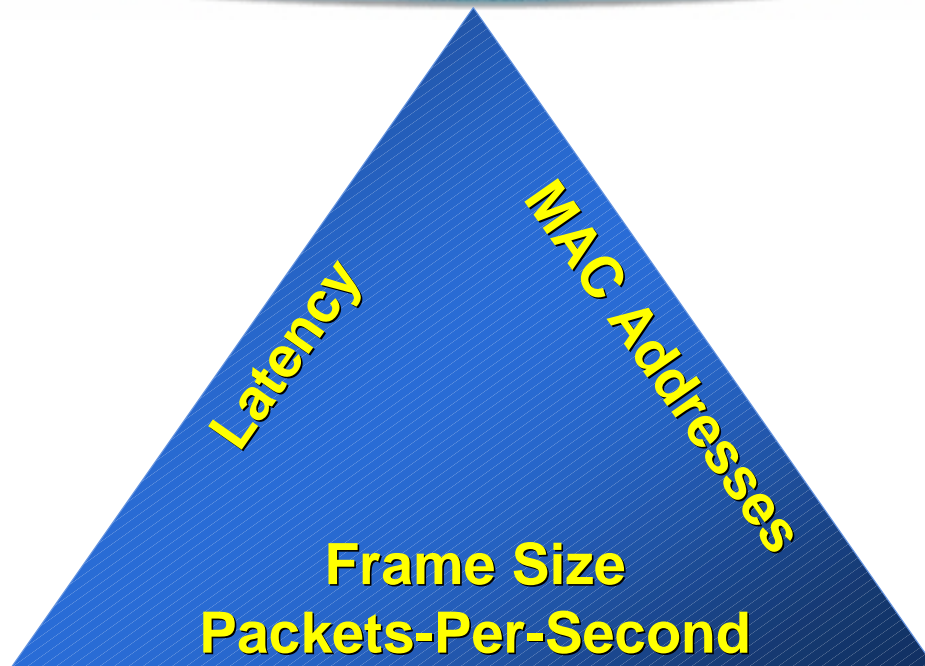
- Buffering, queuing, and latency
- Interface and pipe sizes, speed
- Control plane, IP flow speed (if flow-based)
- Memory, IP flow capacities, buffering capacities, control plane and data plane

Relational Perflow Metrics




**Any One Change
Affects the Other**

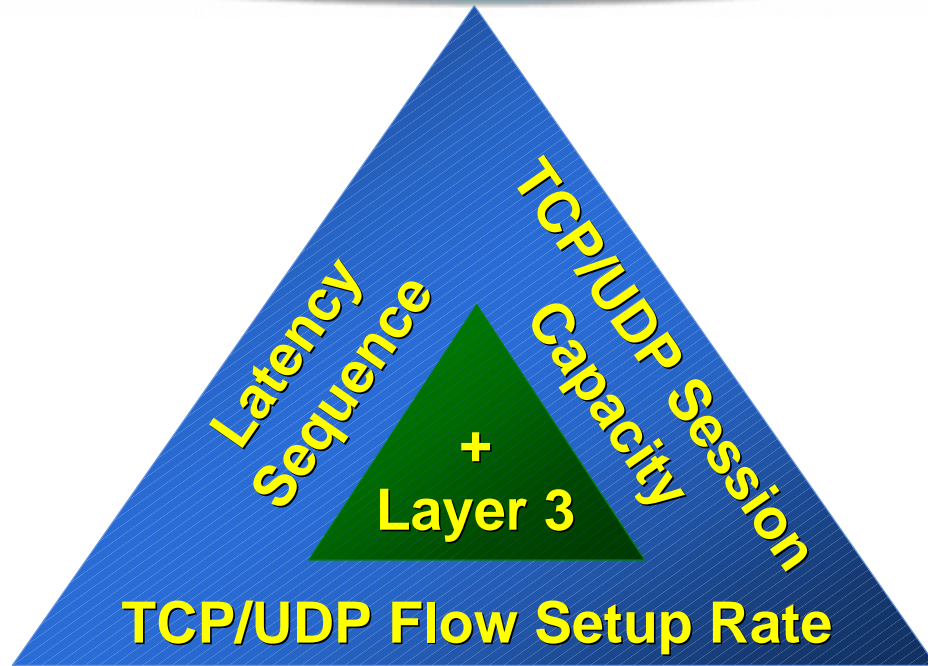
Relational Layer 2 Metrics



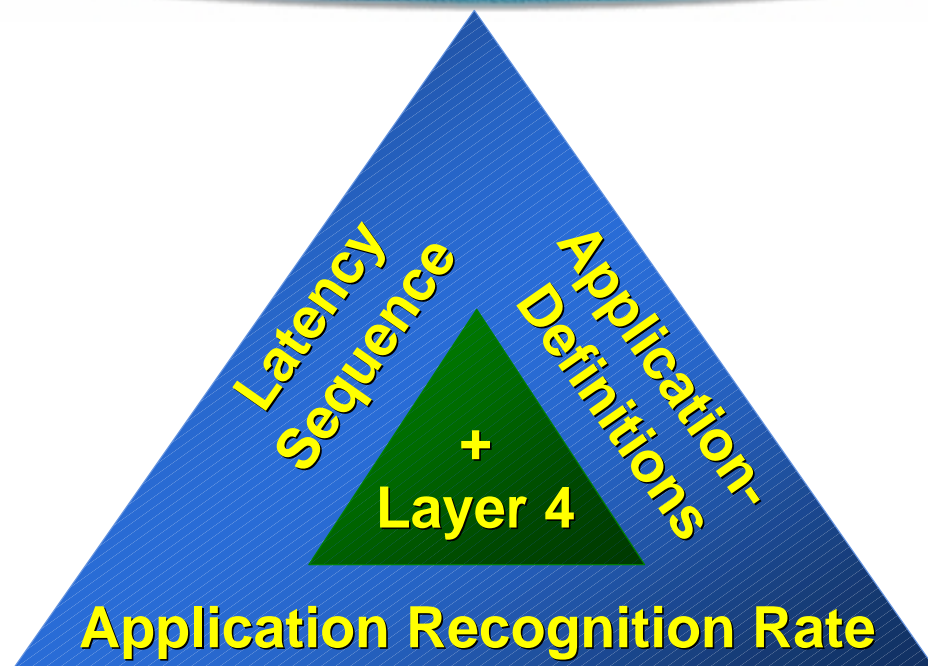
Relational Layer 3 Metrics



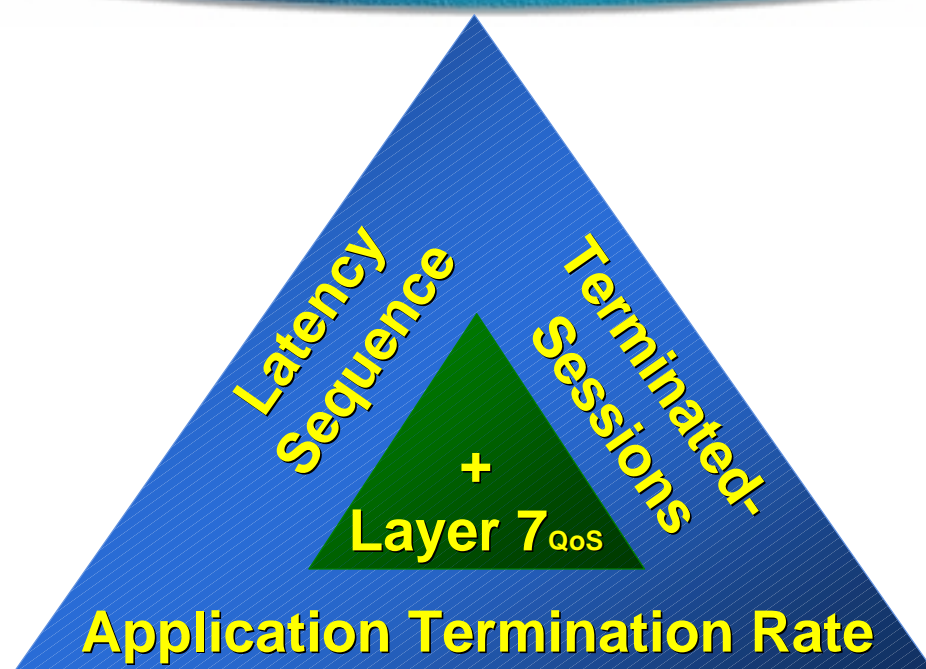
Relational Layer 4 Metrics



Relational Layer 7 QoS



Relational Layer 7 Termination




Agenda

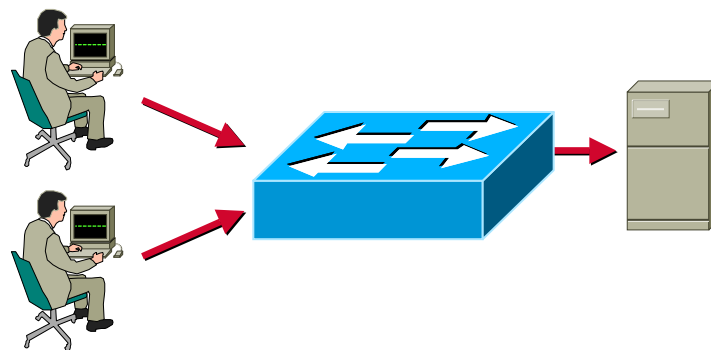
- Definition of Terms
- Performance Test Tools
- Testing Methods and Metrics
- **Example—Buffering and Architecture**
- **Example—Flow Lookup and Architecture**
- **Example—WRR and Architecture**
- **Example—SLB and Architecture**

Buffering and Switching Architectures

- **Congestion management**
- **Switching decision**

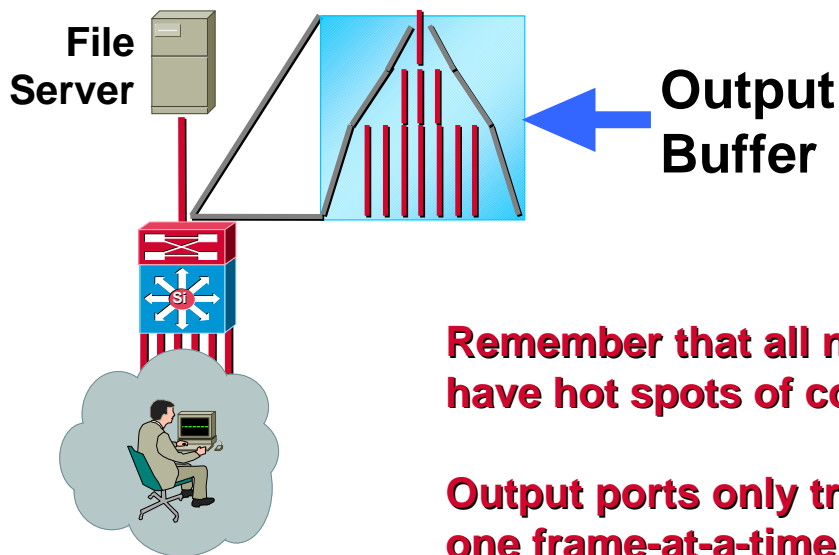


Congestion Management



- **Required when multiple ports are contending for the same port**
- **Congestion occurs during short bursts**


Closet/Server Buffering



Remember that all networks have hot spots of congestion

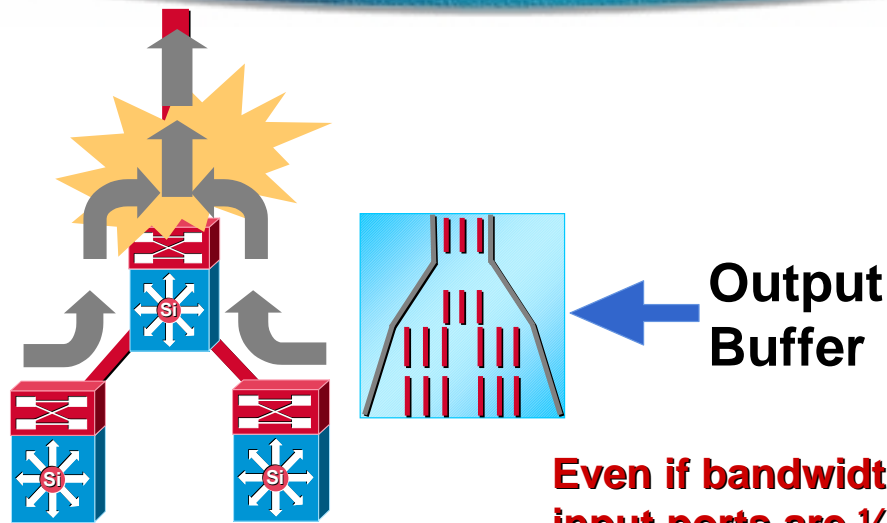
Output ports only transmit one frame-at-a-time

High-Speed Buffering




**1 Gigabit Ethernet port at 100% fills a 500 KB buffer in:
4.1 ms with 1500 Byte MTU
5.5 ms with 64 Byte MTU**

Instantaneous Over Subscription



Even if bandwidth of two input ports are ½ the speed of the output port congestion will occur!

Dynamic Buffer Queuing



- Each buffer element fixed in small increments (for example, 64 bytes each)
- Allows for efficient use of buffers
- **More expensive**

Fixed Buffer Queuing



- Buffer element is a fixed size
- Inefficient use of buffers
- Less controller overhead by 31 times—easier to make a faster switch!


Dynamic vs. Fixed Buffering

Dynamic
Ethernet—Frame Size
Voice
Data Throughput

Large—Small
Buffers Available

Fixed
ATM—Cell Size
Voice
Data Requires SAR


Only Small
Buffers Available



Buffering Tests 2:1

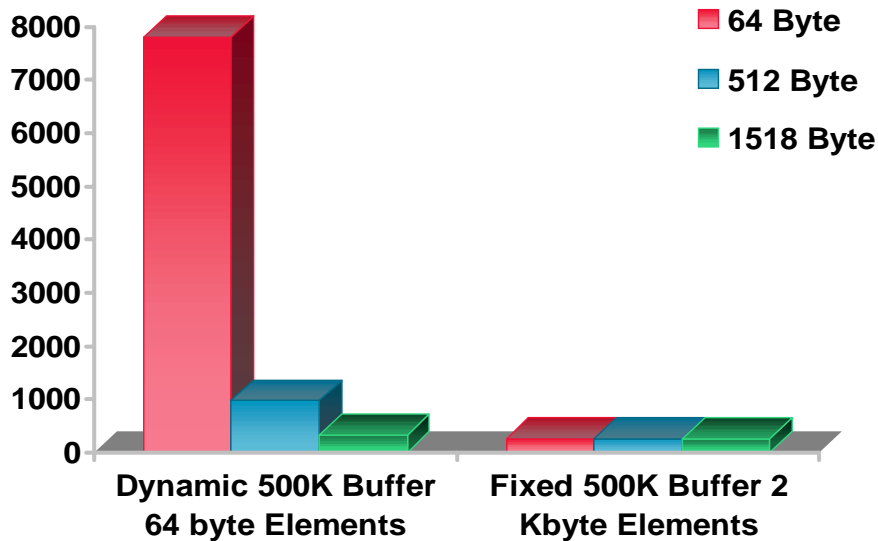
Backbone Test

- Two 1000 Mbps—
one 1000 Mbps
- Burst 100–3000 frames for
2:1 gigabit Ethernet ports
at 5–50% utilization

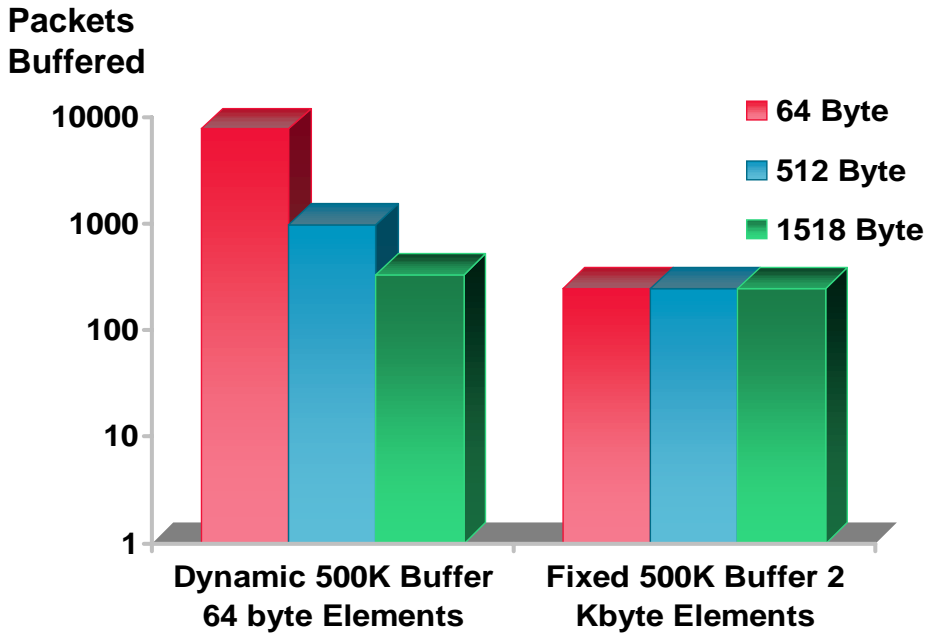


Buffering Capacity


Packets
Buffered




Buffering Capacity



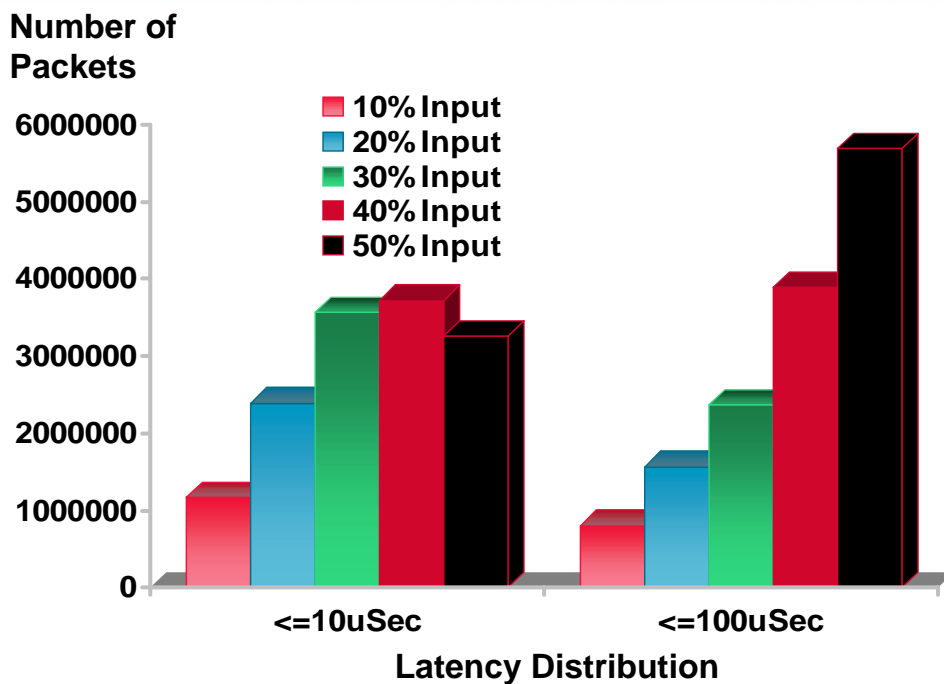
Buffering Test 2:1



Buffering Test 2:1

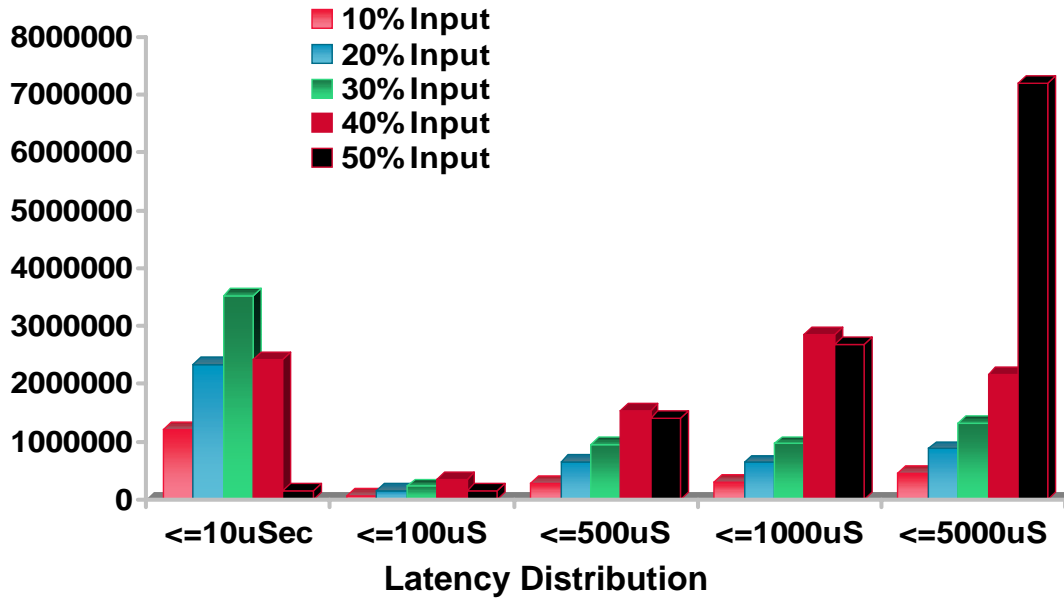


Fixed Queuing Latency Distribution 2:1



Dynamic Latency Distribution 2:1

Number of Packets



2802
1272_05_2000_c2

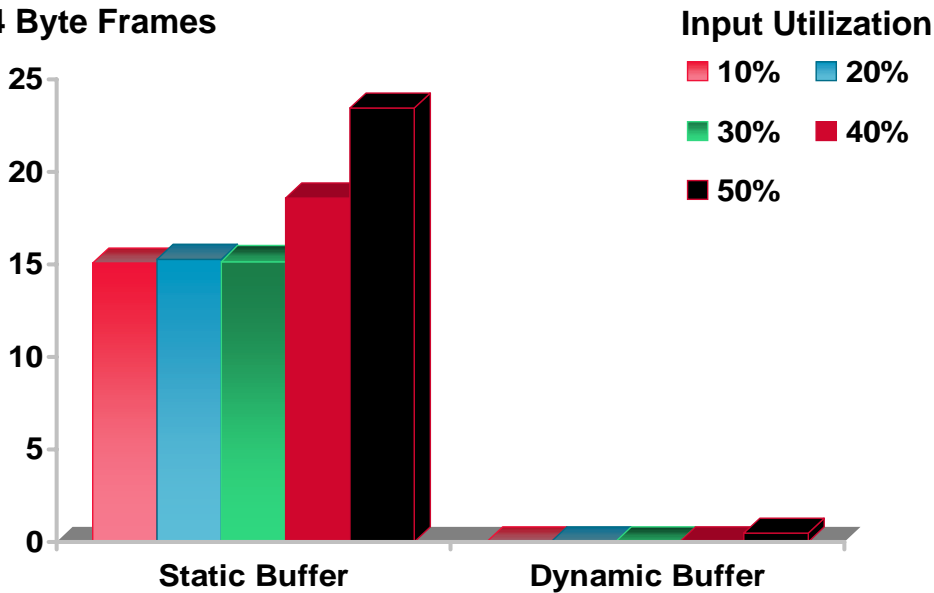
© 2000, Cisco Systems, Inc.

cisco.com

51

Buffering Results 2:1

Percent Lost with 64 Byte Frames




2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

52

Buffering Tests 40:1



Closet Test

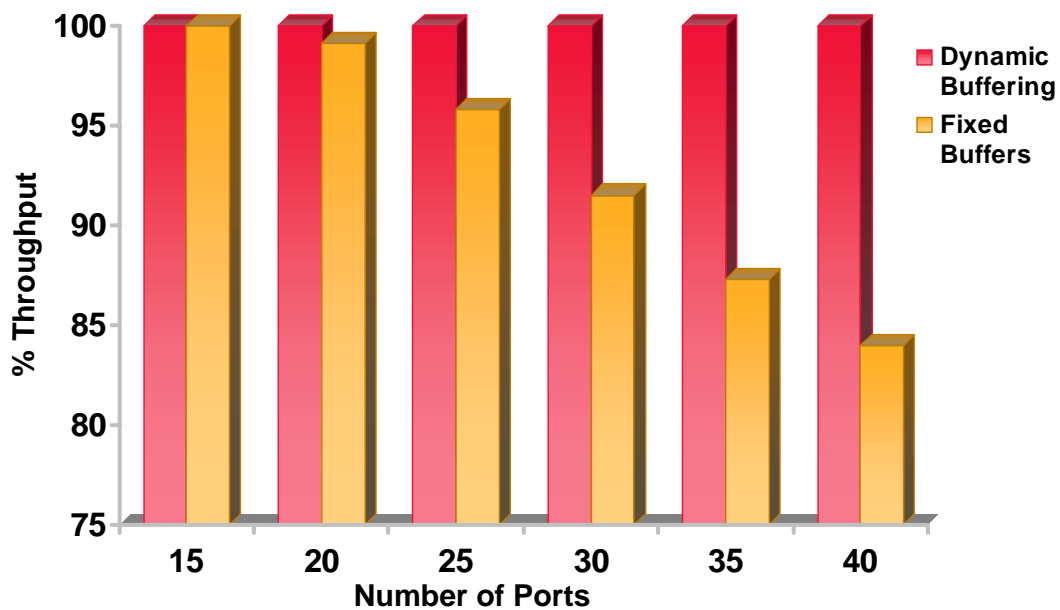
- 40 100 Mbps—
One 1000 Mbps
- Burst 10–30 frames for 20:1
100 Mbps Ethernet ports at
1–5% utilization

2802
1272_05_2000_c2 © 2000, Cisco Systems, Inc.

cisco.com

53

Test Results— 64 bytes 10 ppb




2802
1272_05_2000_c2 © 2000, Cisco Systems, Inc.


cisco.com

54


Test Results— 500 bytes 10 ppb




Test Results— 1518 bytes 10 ppb



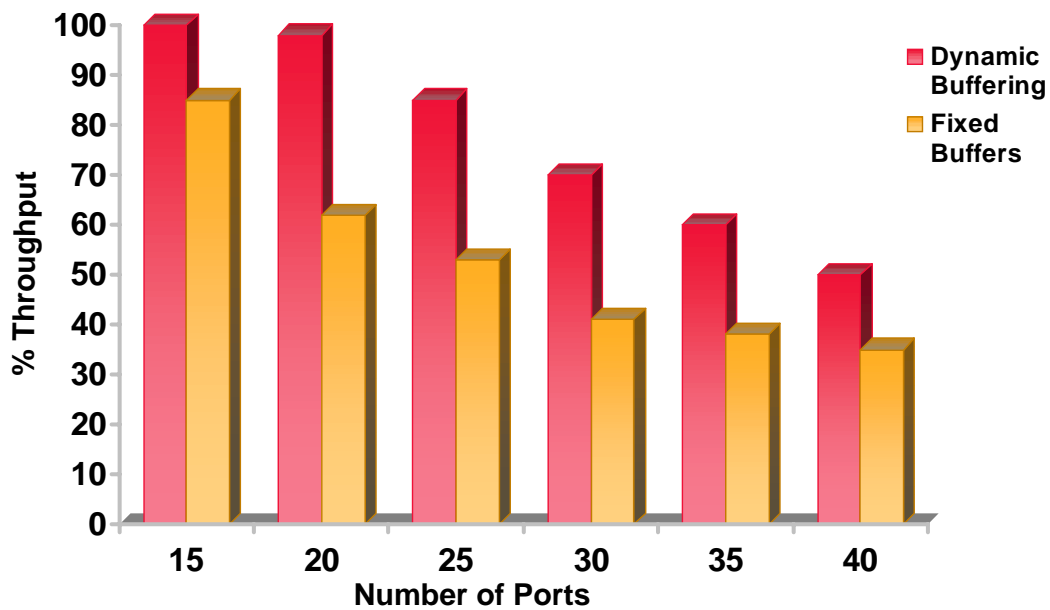
Test Results— 64 bytes 24 ppb



Test Results— 500 bytes 24 ppb



Test Results— 1518 bytes 24 ppb



2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

59

Agenda

- Definition of Terms
- Performance Test Tools
- Testing Methods and Metrics
- Example—Buffering and Switch Architecture
- **Example—Flow Lookup and Architecture**
- **Example—WRR and Architecture**
- **Example—SLB and Architecture**

2802
1272_05_2000_c2


© 2000, Cisco Systems, Inc.

cisco.com

60

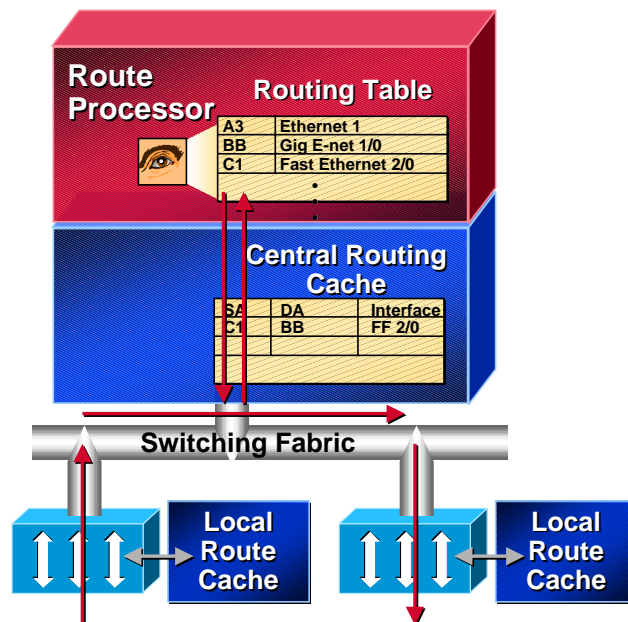
Flow Lookup and Switch Architectures

- Congestion management
- **Switching decision**



Route Caching (Flow-Based Switching)

- First packet processed switched by route processor
- Flow cache enabled (centrally and/or locally)




CPU-Control Plane

CPU-Control Plane

- Run routing protocols
- Maintain routing tables
- Check for CLI commands
- Accounting counters
- ICMP
- Queue packets, etc...

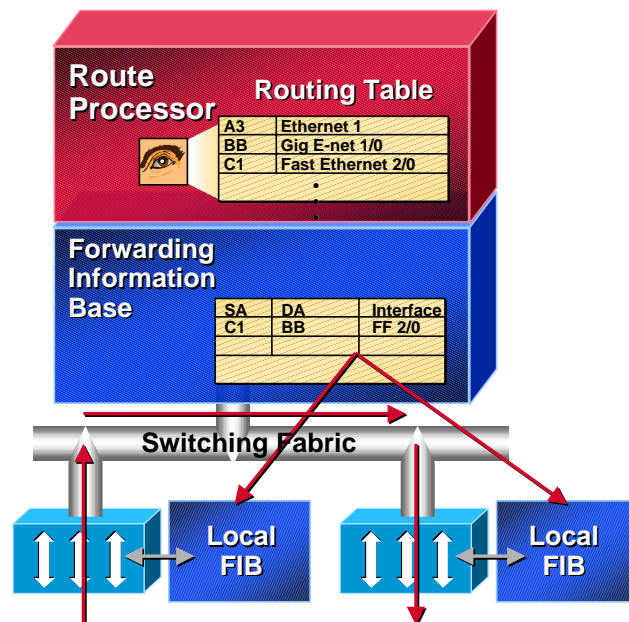
CPU-Control Plane Issues

- Control Plane must manage flow setup and system management
- What happens to Spanning Tree, VRRP and routing updates if the CPU is busy processing flows?

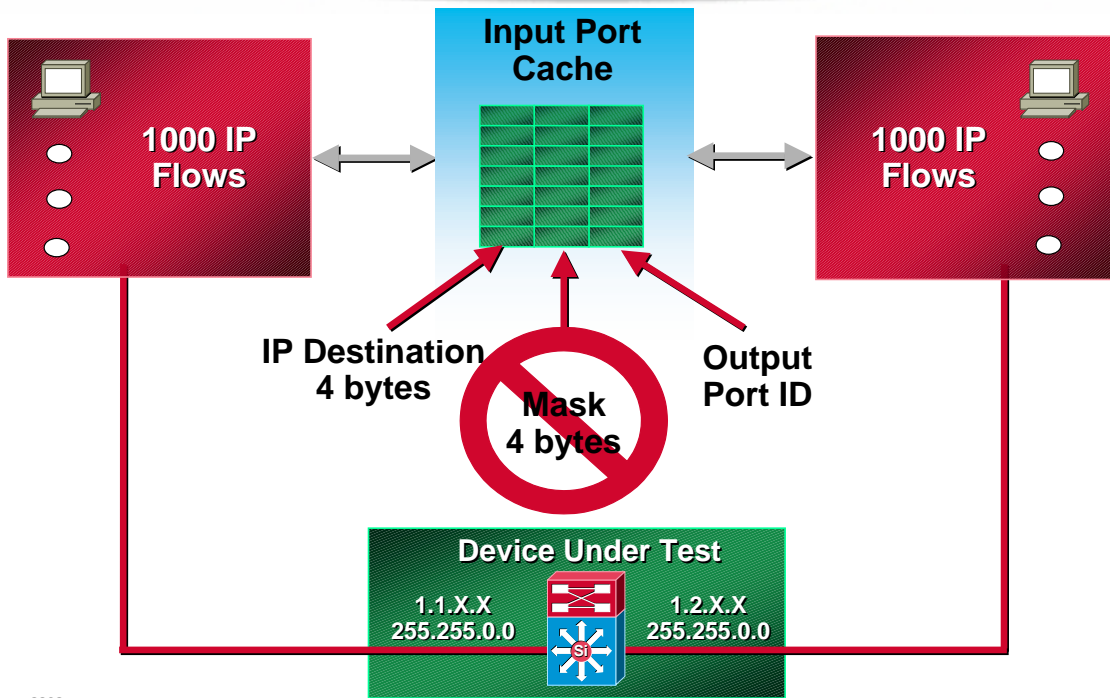


Cisco Express Forwarding (FIB-Based Switching)

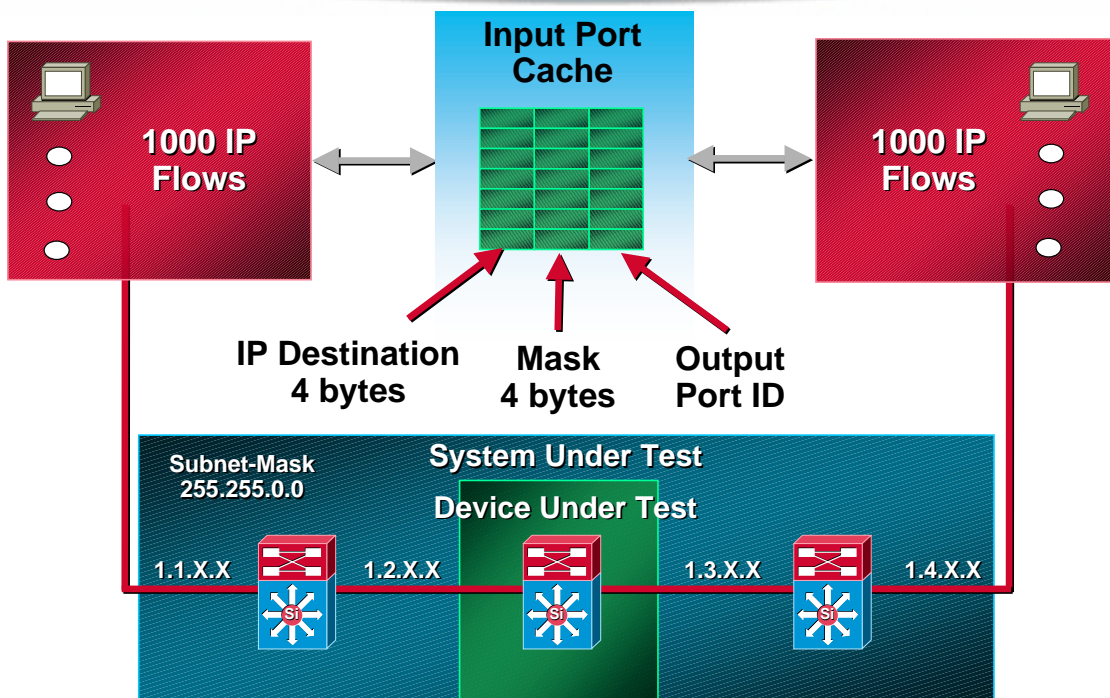
- No process switching necessary
- FIB based on routing table entries, not traffic flows
- FIB can be kept central or distributed



IP-Host Flow Test




IP-Recursive Flow Test



Flow Test Results—Loss

Packets Transmitted



2802
1272_05_2000_c2

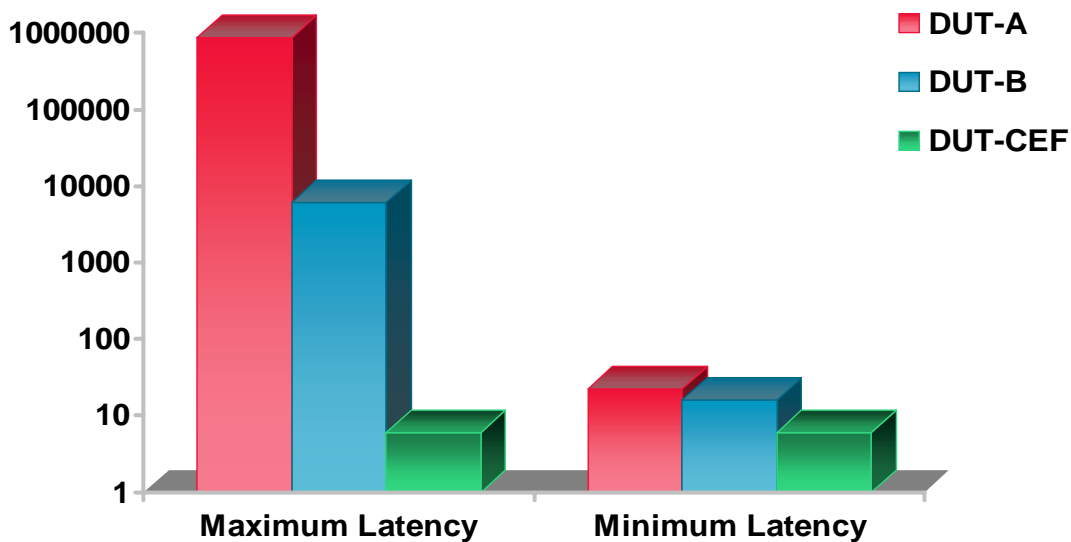
© 2000, Cisco Systems, Inc.

cisco.com

67

Flow Test Results—Jitter

Latency in Microseconds



2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.


cisco.com

68

Flow Test Results—Raw using Smart Windows

L3Tests 1	Port	VTE	Sequence	TxTime	RxTime	DUT-A
1	1	1	0	632079	640709	863.0
2	1	2	0	632149	642458	1030.9
3	1	3	0	632220	644057	1183.7
4	1	4	0	632290	645658	1336.8
5	1	5	0	632361	647272	1491.1
6	1	6	0	632431	648887	1645.6
7	1	7	0	632501	650607	1810.6
8	1	8	0	632572	652236	1966.4
9	1	9	0	632642	653978	2133.6
10	1	10	0	632713	655607	2289.4
11	1	11	0	632783	657360	2457.7
12	1	12	0	632853	658982	2612.9
13	1	13	0	632924	660781	2785.7
14	1	14	0	632994	662465	2947.1

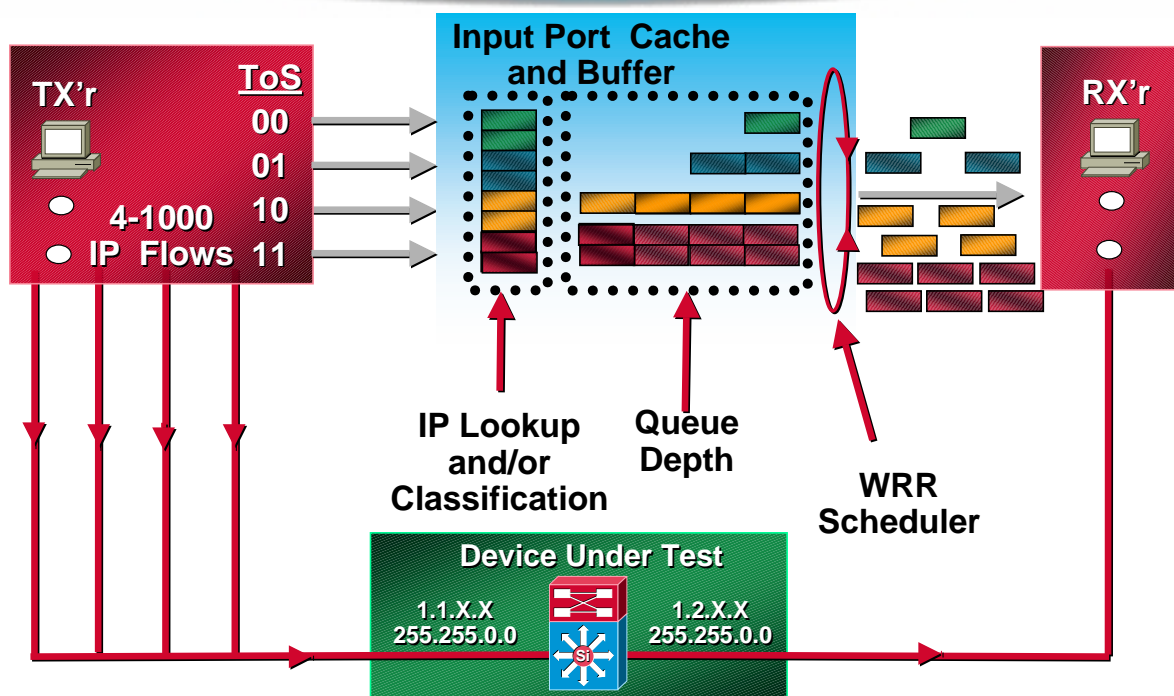
Flow Test Results-Setup Time




Agenda

- Definition of Terms
- Performance Test Tools
- Testing Methods and Metrics
- Example—Buffering and Switch Architecture
- Example—Flow Lookup and Architecture
- **Example—WRR and Architecture**
- **Example—SLB and Architecture**

IP Weighted-Round Robin Queues Test



IP WRR Baseline No Congestion



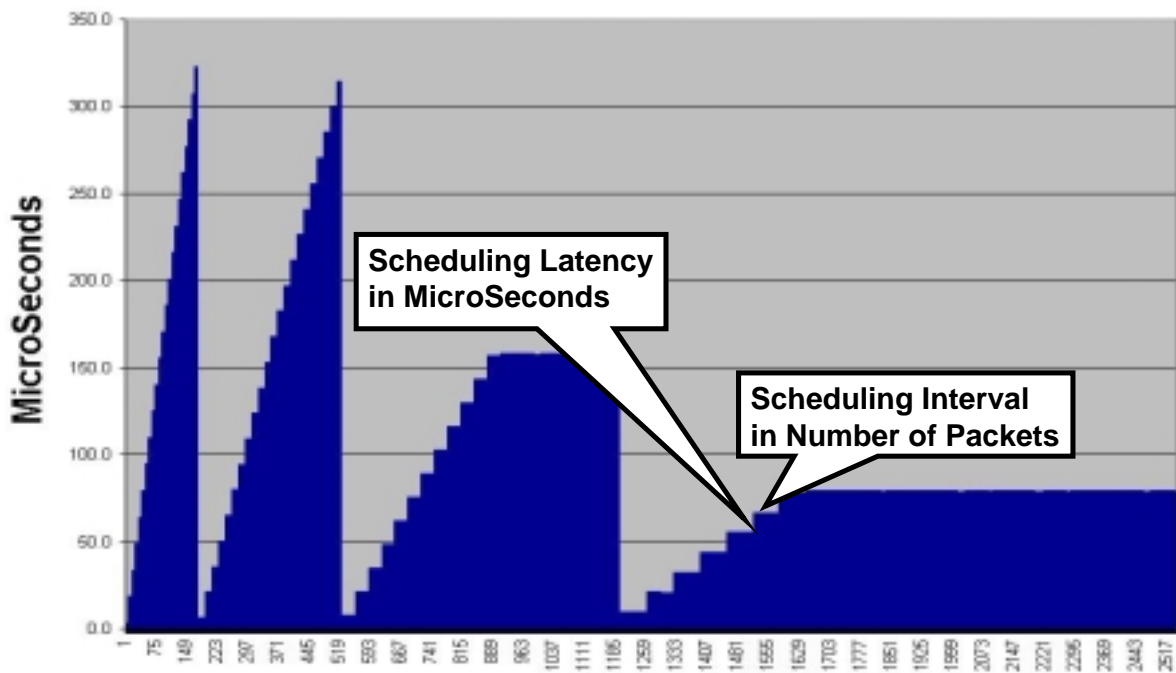
2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

73

IP WRR With Congestion



2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com


74

Agenda

- Definition of Terms
- Performance Test Tools
- Testing Methods and Metrics
- Example—Buffering and Switch Architecture
- Example—Flow Lookup and Architecture
- Example—WRR and Architecture
- **Example—SLB and Architecture**

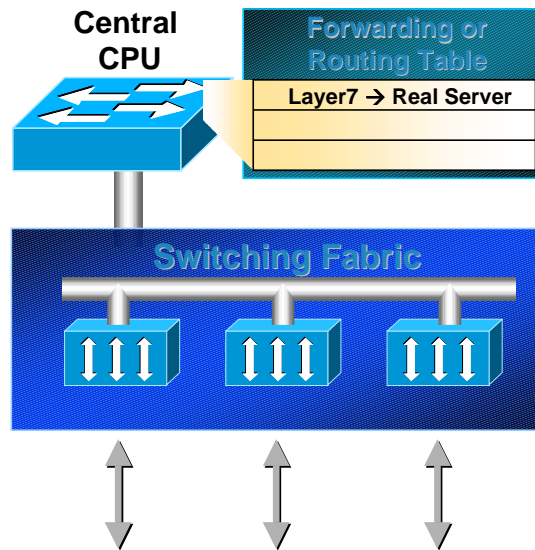
Centralized L4 Switching

- **More memory = More connections**
- **Centralized State**
- **Centralized Persistence**
- **Physical port independent**

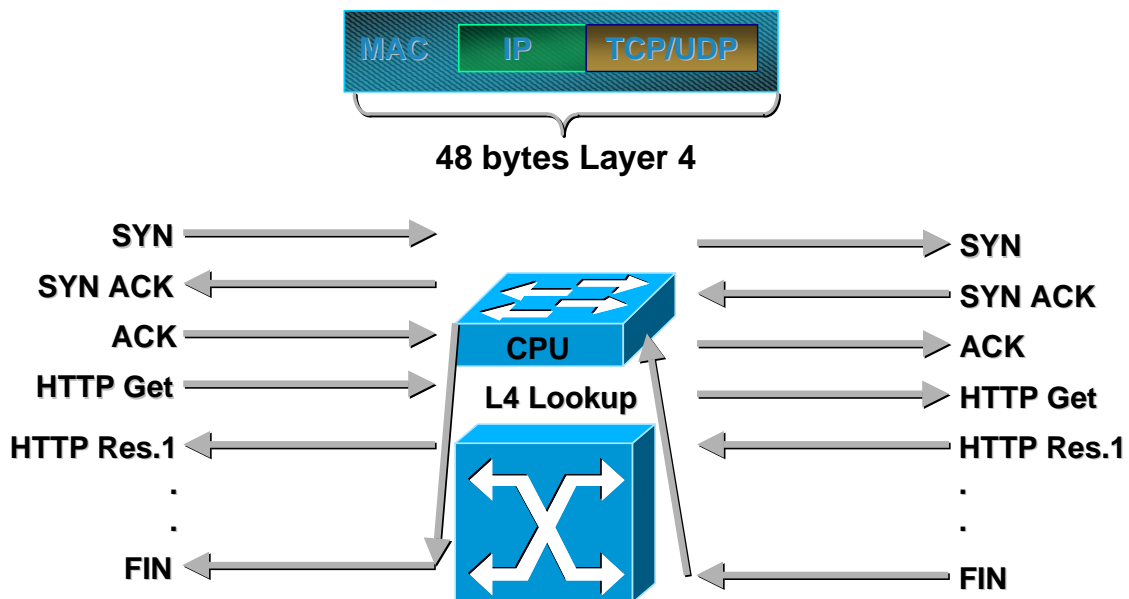


Centralized L7 Switching

- One place to keep persistent information
 - Cookies
 - URL
 - SSL
- Works with and without interface caching
- Interface “punts” to CPU
- Cache examples- Vito,C6k(Dispatch) IOSSLB & ASLB, Arrowpoint

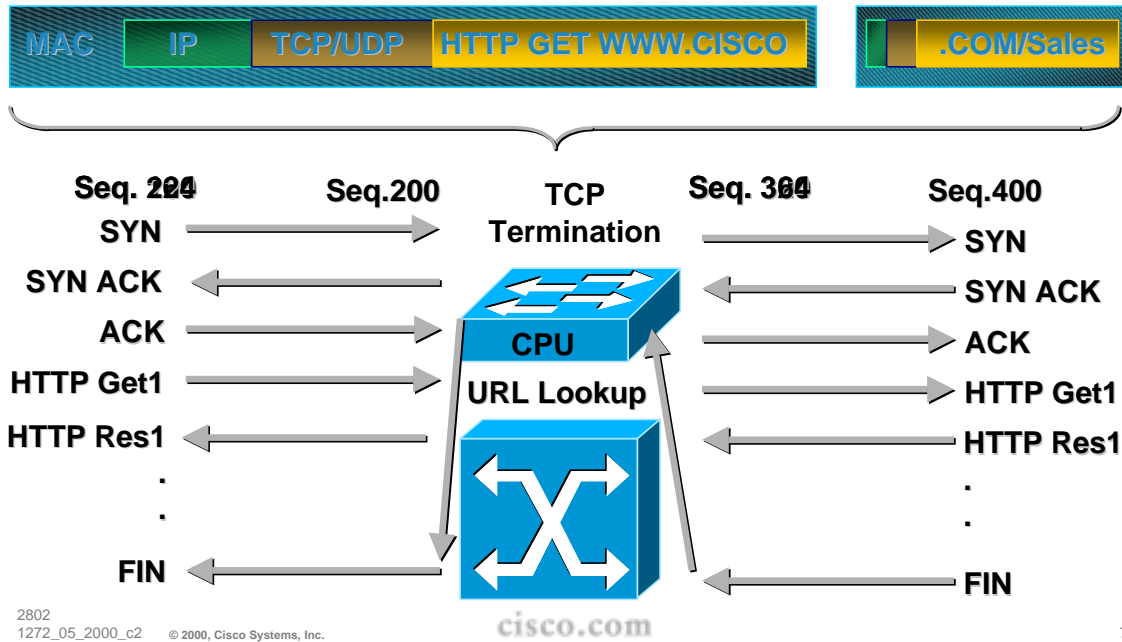


Layer 4 Load Balancing



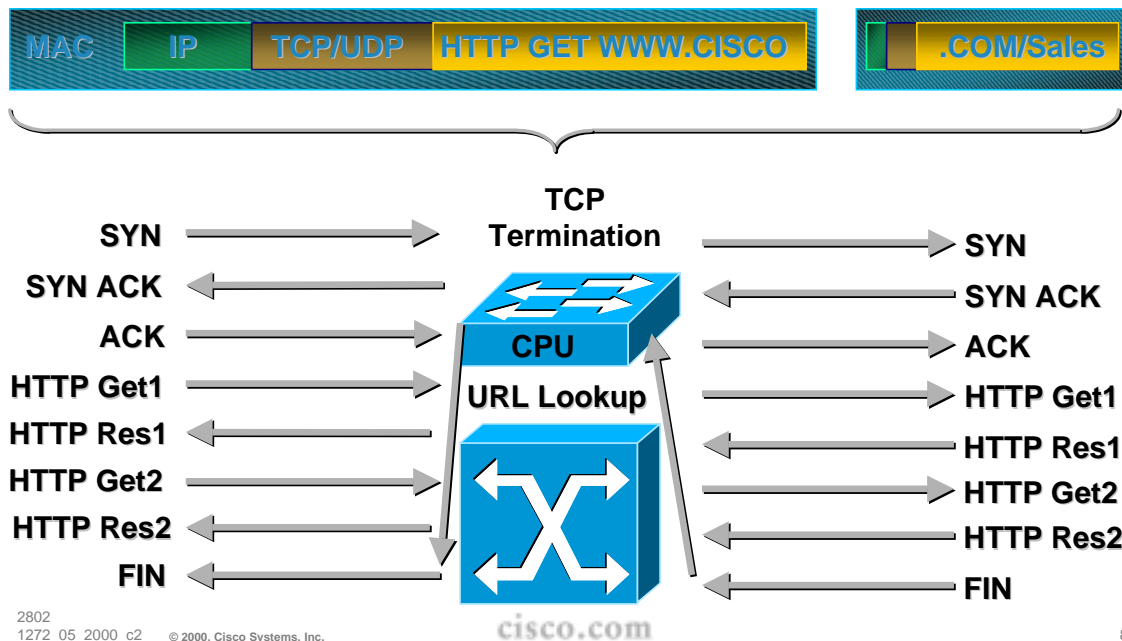
Layer 7 Load Balancing

HTTP 1.0 Single HTTP Get Per-Session

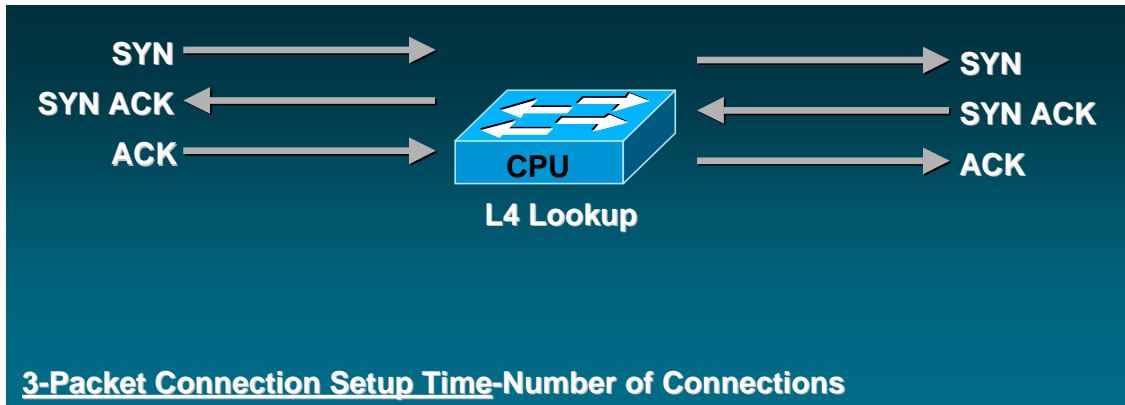



Layer 7 Load Balancing

HTTP 1.1 Multi HTTP Get Per-Session



Layer 4 Testing-SmartTCP™




2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

81

Layer 4 Testing-SmartTCP™



Test Completes as a BURST of connections NOT continuous

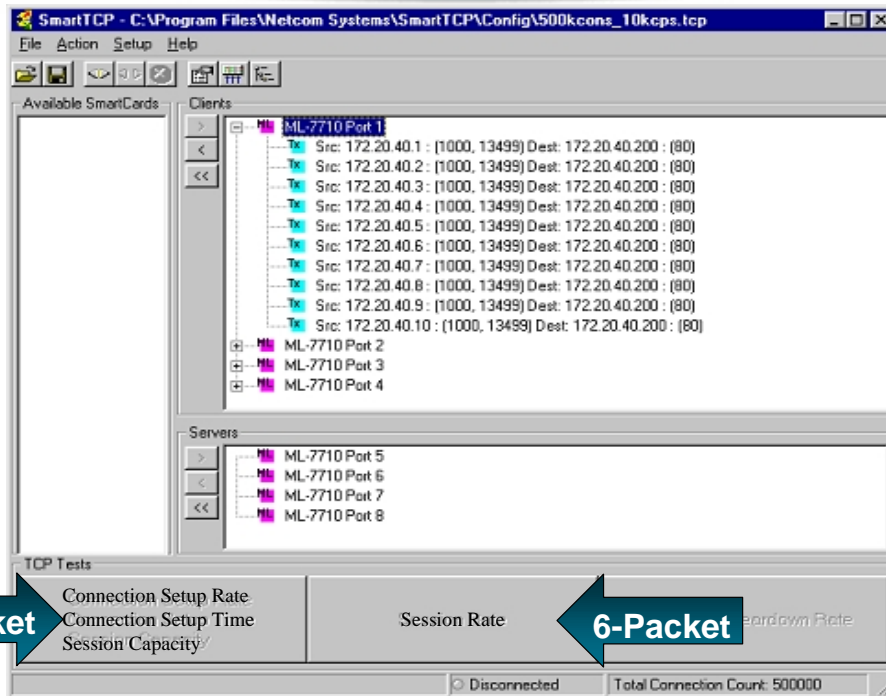
2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

82

Layer 4 Testing-SmartTCP™



2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

83

Layer 4 Testing-SmartTCP™

Card ID	Speed	Duplex	MAC Address	IP Address	Gateway	Response Delay (µsec)	Auto-Negotiation
Client Port 1	100M	Full	00 00 00 00 00 01	172.020.040.001	000.000.000.000	0.96	Off
Client Port 2	100M	Full	00 00 00 00 00 02	172.020.040.011	000.000.000.000	0.96	Off
Client Port 3	100M	Full	00 00 00 00 00 03	172.020.040.021	000.000.000.000	0.96	Off
Client Port 4	100M	Full	00 00 00 00 00 04	172.020.040.031	000.000.000.000	0.96	Off
Server Port 5	100M	Full	00 00 00 00 00 05	172.020.040.201	000.000.000.000	0.96	Off
Server Port 6	100M	Full	00 00 00 00 00 06	172.020.040.202	000.000.000.000	0.96	Off
Server Port 7	100M	Full	00 00 00 00 00 07	172.020.040.203	000.000.000.000	0.96	Off
Server Port 8	100M	Full	00 00 00 00 00 08	172.020.040.204	000.000.000.000	0.96	Off

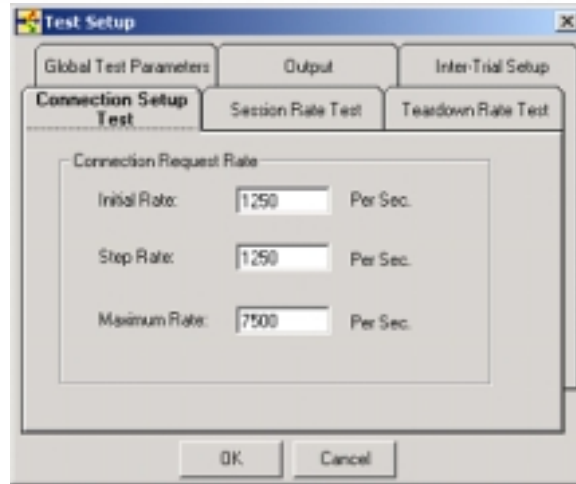
2802
1272_05_2000_c2

© 2000, Cisco Systems, Inc.

cisco.com

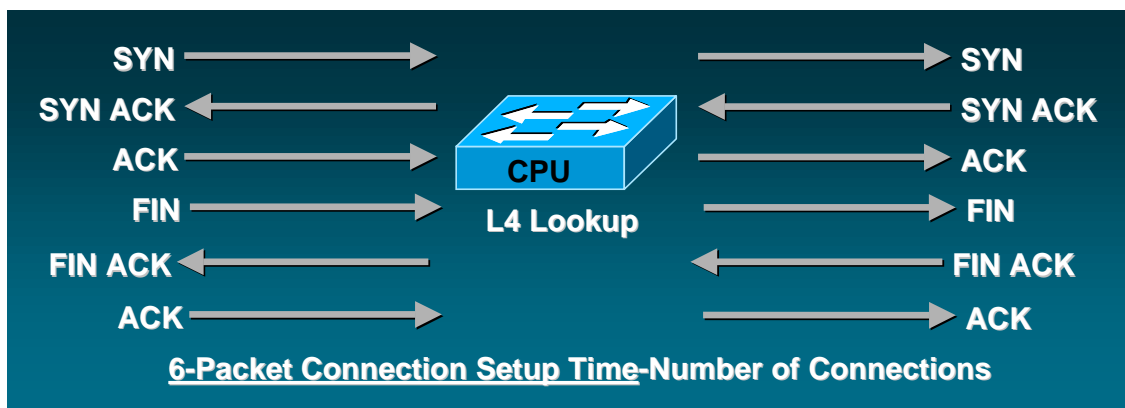

84

Layer 4 Testing-SmartTCP™

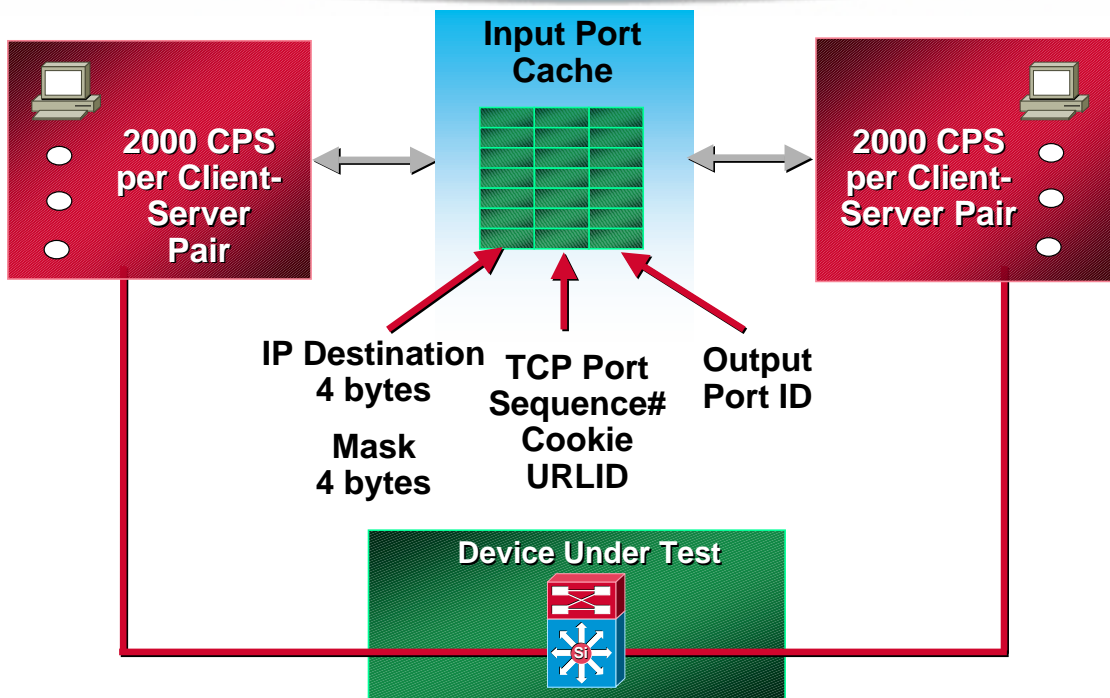


$$\text{RATE} \times (\text{NUMBER OF CLIENT CARDS}) = \text{TOTAL NUMBER OF CONNECTIONS PER SECOND}$$

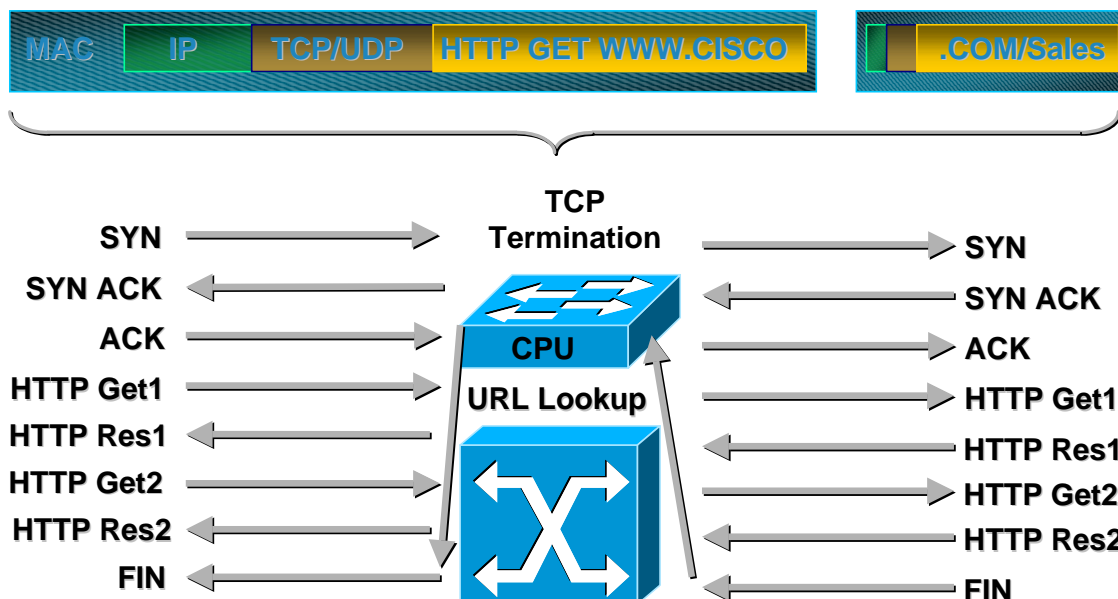
Layer 4 Testing (no data)



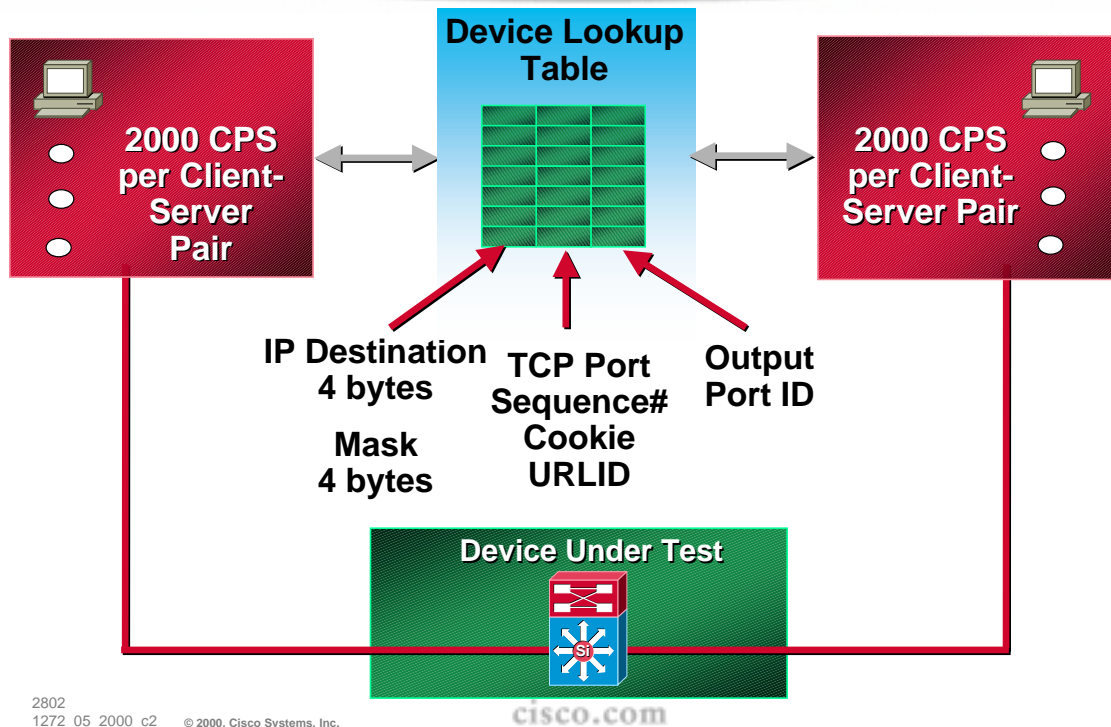
Layer 4 Testing



Layer 7 Testing 1.0 vs. 1.1




Layer 7 Testing 1.0 vs. 1.1



Layer 7 Testing 1.0 vs. 1.1

- WebStone
- WebBench-Ziffdavis (400cps)
Best functionality tool
- Tarantula-Cisco/Arrowpoint
(4000cps) High Performance
- Socrates-Very Easy to setup
Good functionality tool



Performance Metrics for Multilayer Switches

Session 2802



Please Complete Your Evaluation Form

Session 2802

CISCO SYSTEMS



EMPOWERING THE
INTERNET GENERATIONSM