Cisco.com

Gigabit to the Desktop

Gigabit to the Desktop

What we are seeing:

- Today’s driver for Gigabit Ethernet to the Desktop is not a single application but the simultaneous use of multiple applications
- Product availability from Cisco, Dell, Foundry, Extreme, 3Com with more coming soon…
- Gig enabled PCs/Workstations (LOM) – Dell, HP/Compaq, Apple, Sun, many Linux hardware manufacturers
- Cisco 10/100/1000 port sales +328% CY01 vs. CY02

Gigabit Solutions

- **Improve the end-user experience**
- **Increased throughput with 10/100/1000 Ethernet and 10 GE interfaces**
- **Reduce wire time, buffer congestion & relieve flow control mechanisms**
- **Convergence of applications at the desktop requires increased throughput end-to-end ✍ Gigabit to the Desktop (GTTD)**
- **End-to-end Intelligent Network Services: QoS, Security, High Availability, Manageability**

Real Time Traffic Modeling

Cisco.com

- Look at day to day functions – compute profile based off of review of real time application use
- Multiple applications running simultaneously
- Transactional, bulk transfer and streaming applications within the compute profile
- Same hosts and same scripts used in each test at 10, 100 and 1000

Compute Profile Applications

- Outlook, Ariba, Connected TLM, Reflection X, Clarify, Oracle, Siebel, IPTV, IP Telephony, WBT and other web based tools.
- Gains are seen in each application but it is the overall use of network resources where the true gains lie.

Network Time by Connect Speed

- **3.1 hours of network time at 10mbps**
- **47 minutes of network time at 100mbps**
- **27 minutes of network time at 1000mbps**
- **At gig speed we spend 88% less time on the network than 10mbps and 44% less than 100mbps**

Network Response Improvements 10mbps v. 1000mbps

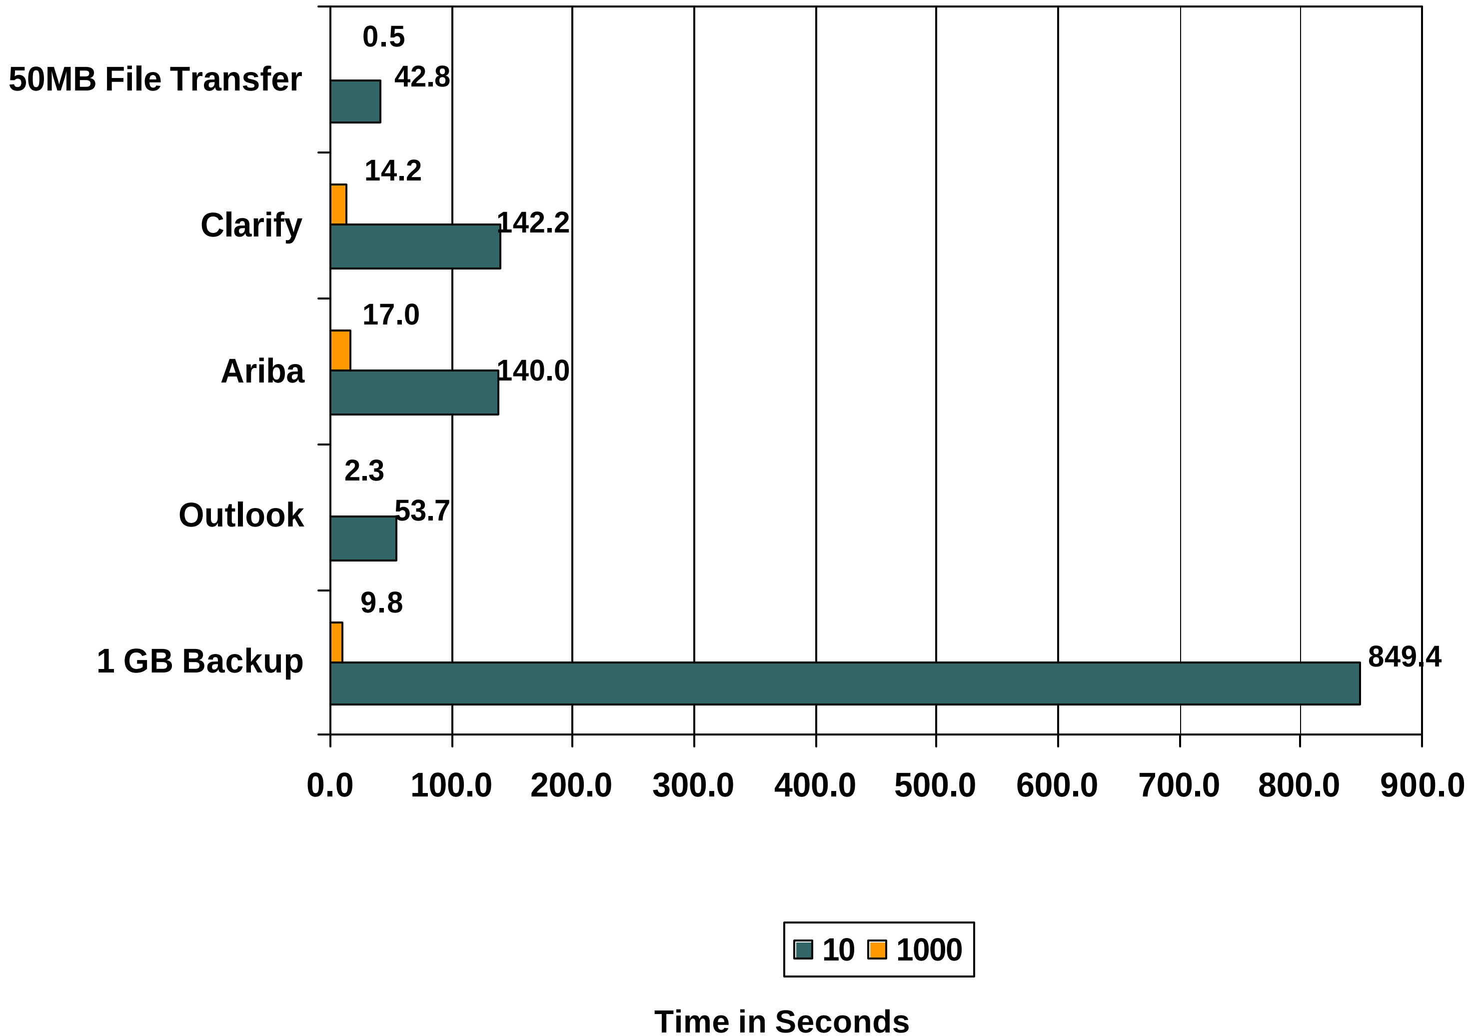

Network Response Improvements 100mbps v. 1000mbps

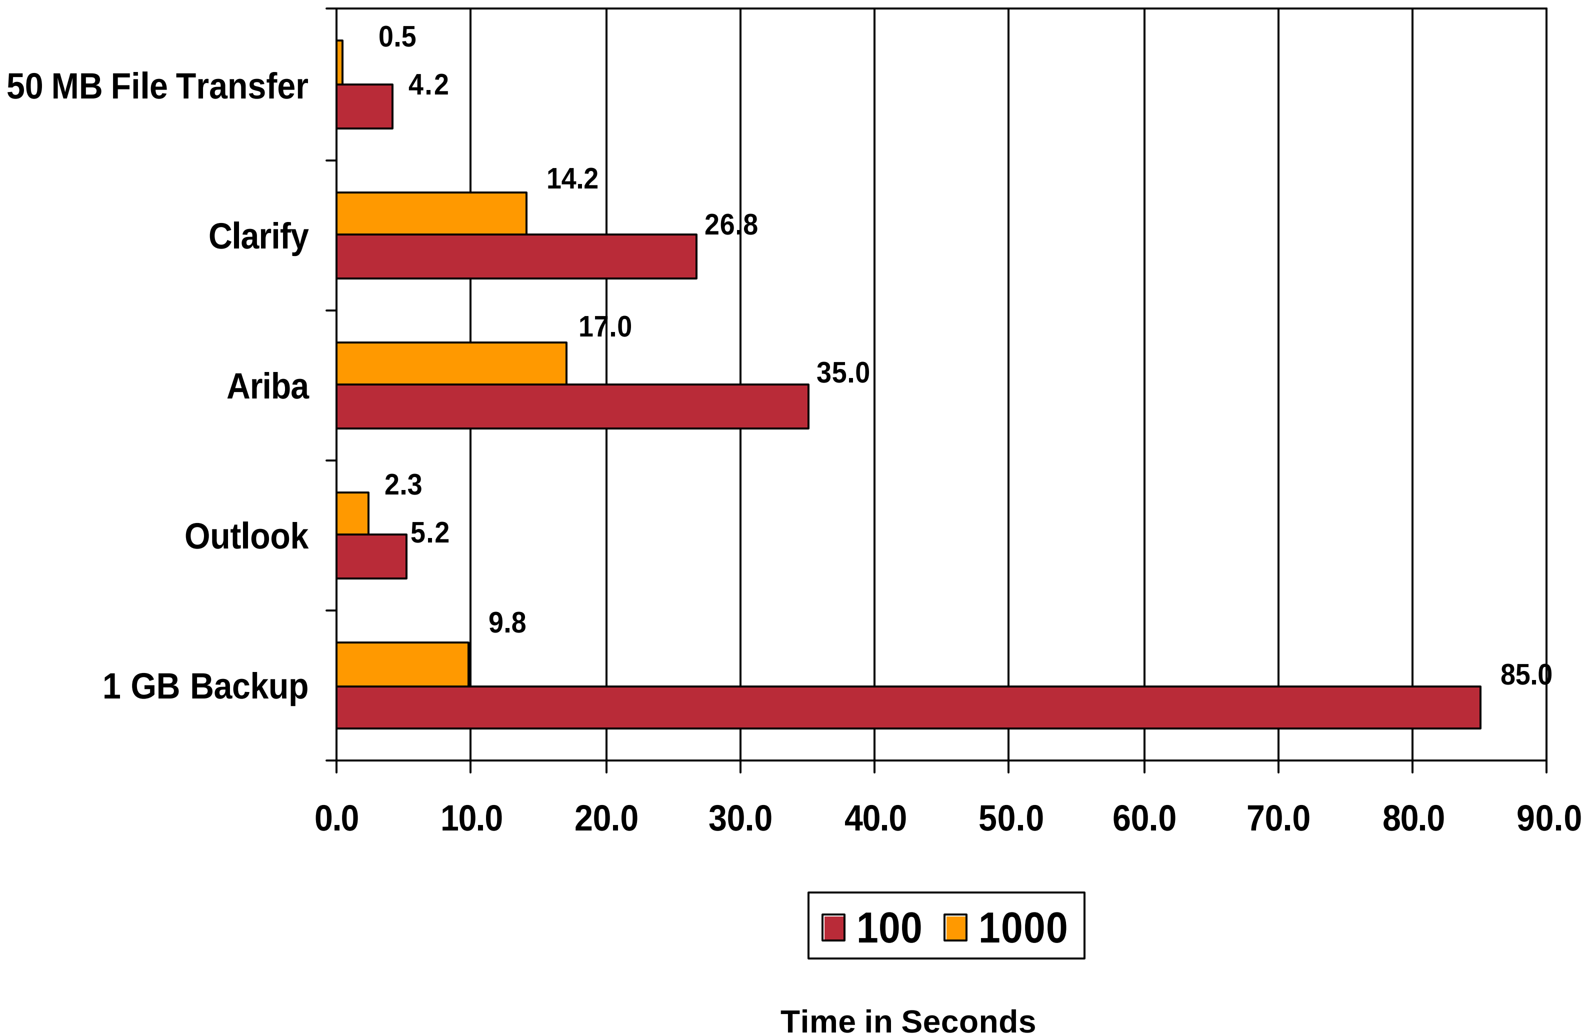

Network Response Improvements 10mbps v. 1000mbps

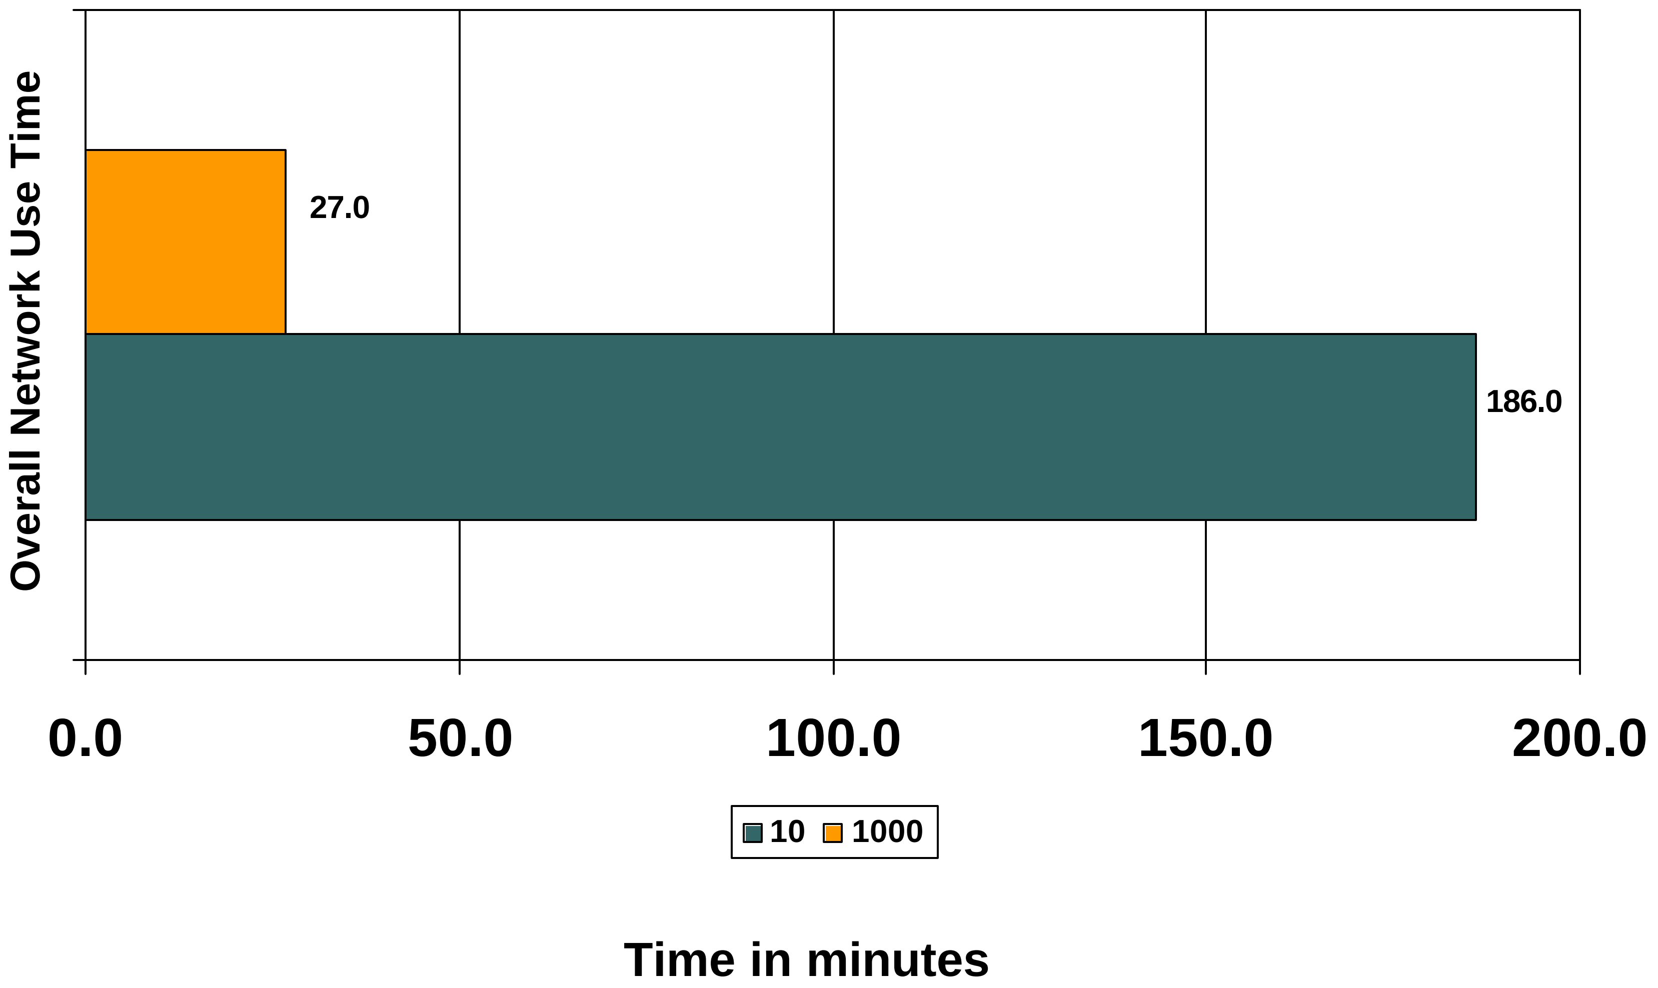

Network Response Improvements 100 vs. 1000

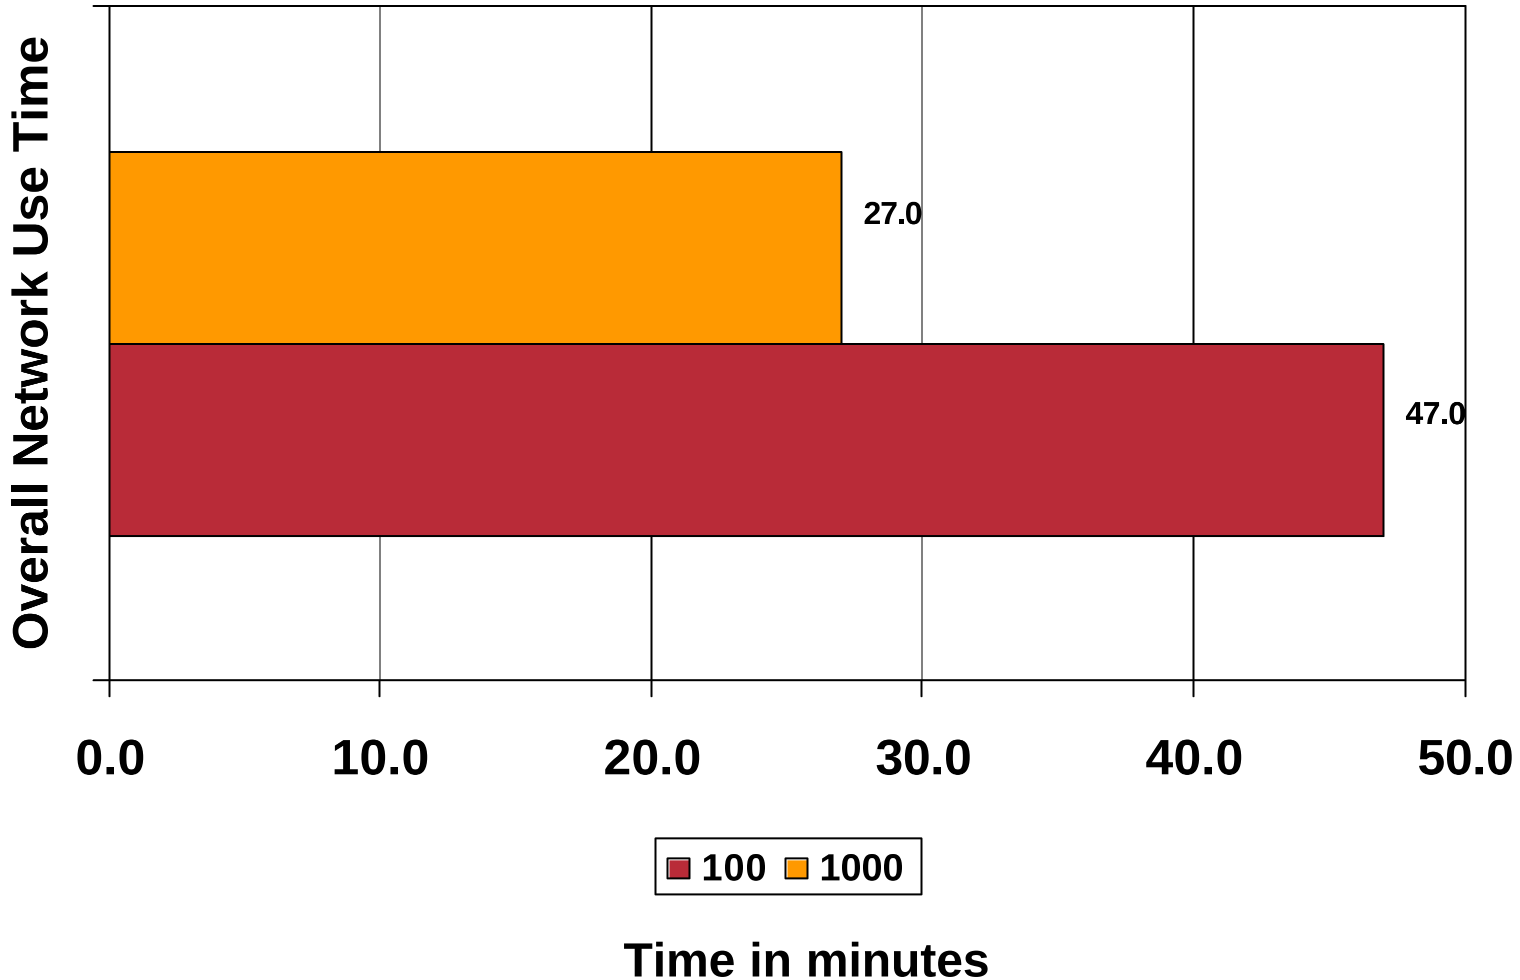

Productivity Enablement

- Late 90's – 2001 60% of productivity improvements are directly related to IT*
- The network is an entity to further enable productivity
- Lower costs of PC's, silicon for memory and processors.
- Increase in the performance of the end stations
- Application development and delivery of information to the desktop that help enterprises reduce hard dollar costs

*** Source: Federal Reserve http://www.federalreserve.gov/pubs/feds/2002/200229/200229pap.pdf**

Productivity Enablement

- **Given the contribution to productivity gains with the existing and emerging desktop technologies, improving software applications, increase in technology awareness and lastly the existing and emerging technologies in the underlying network infrastructure, further gains will be made to productivity levels with the increase in speeds to the desktop.**

Gigabit Costs

Cisco.com

- Manufacturing cost of 10/100/1000 dropping to and below that of 10/100 for both NICs and switch ports
- Gigabit Ethernet NICs as low as $80
- Cisco chassis port costs are nearing the $200 per port level
- Cisco stackable (3750) per port costs of $250

Gigabit Switch Port Sales

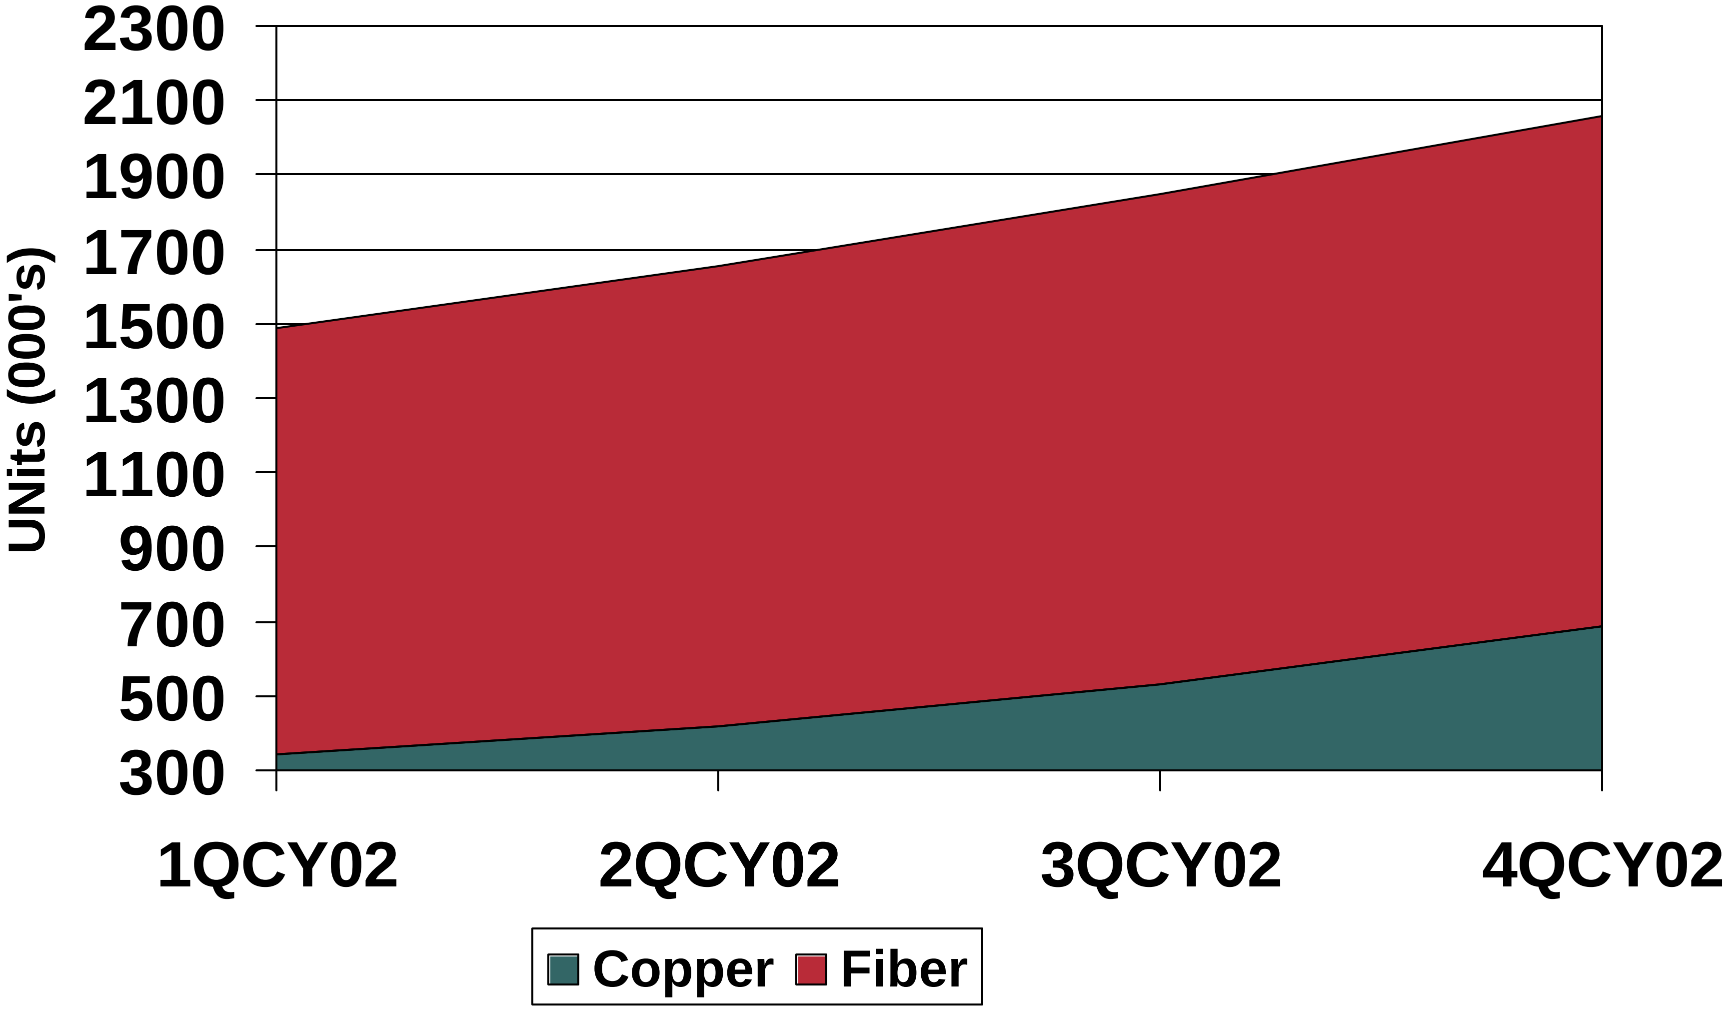

Market Share Gigabit Copper Ports

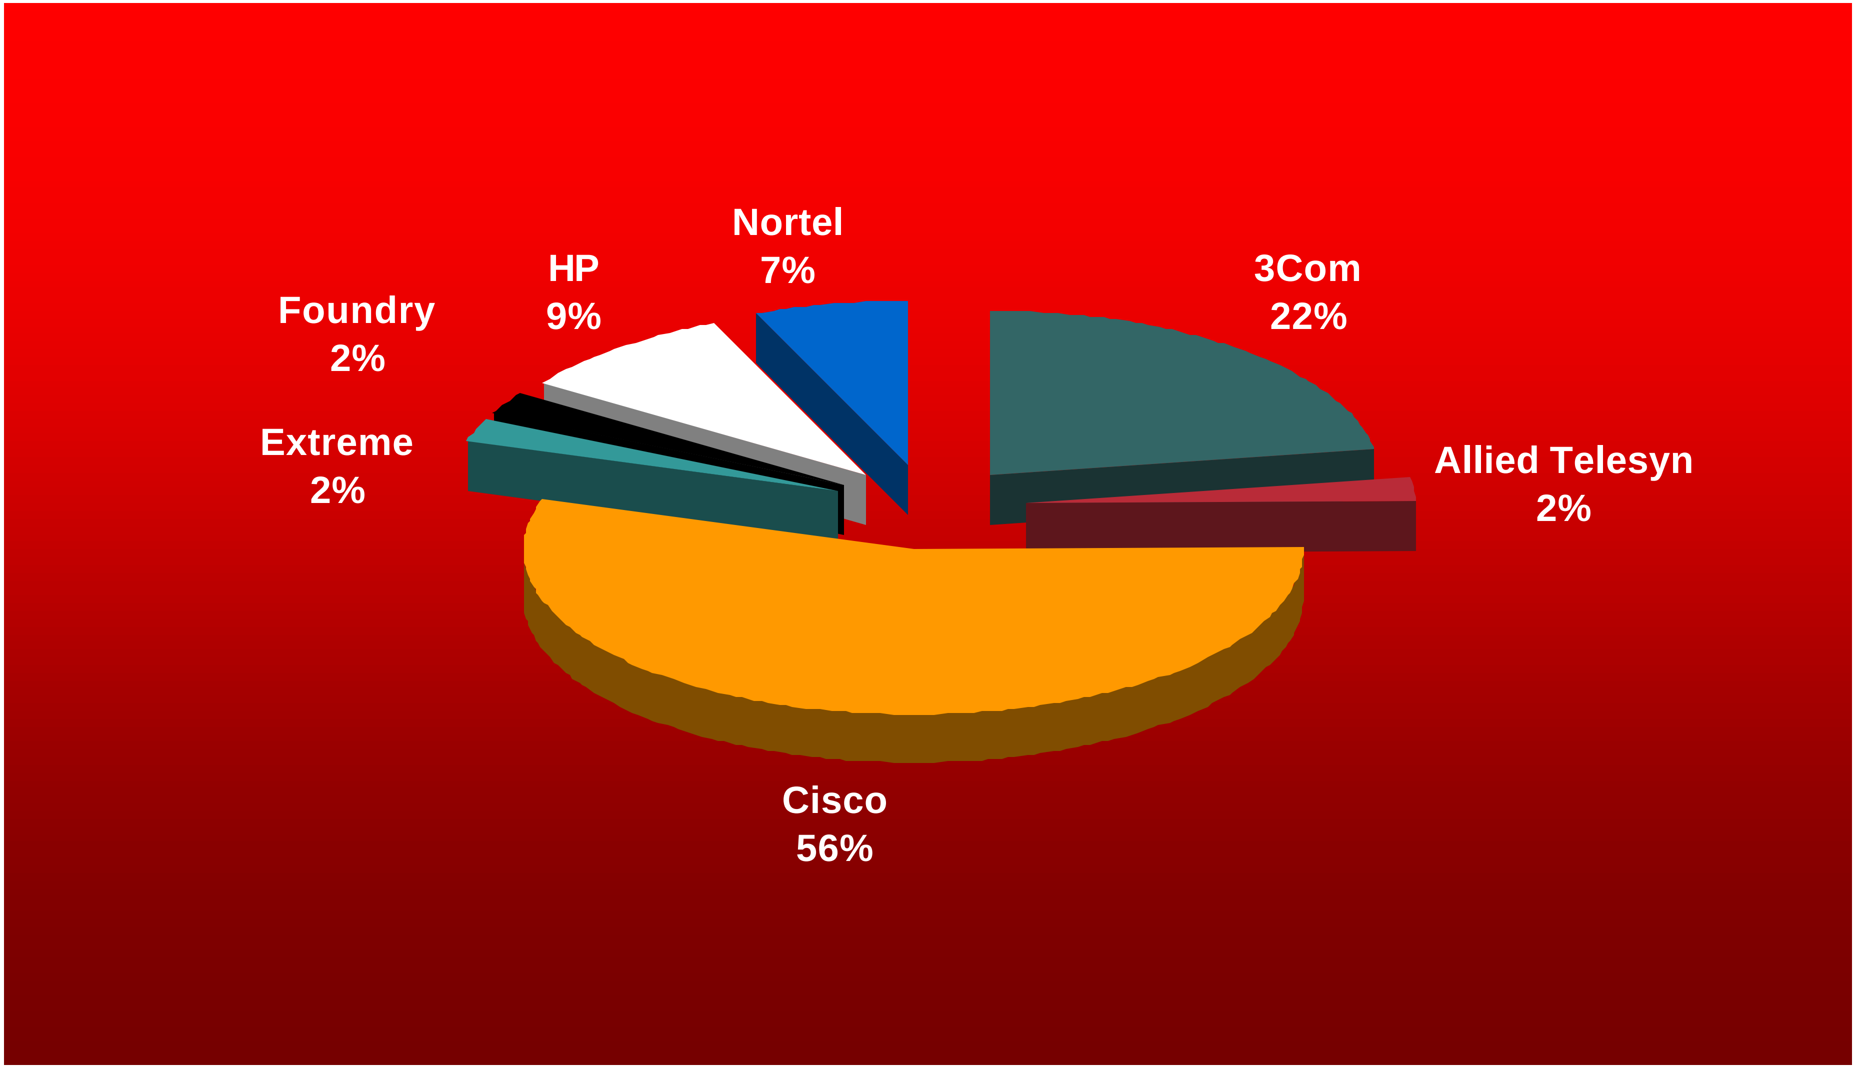

Gigabit Adapter Sales

Cisco.com

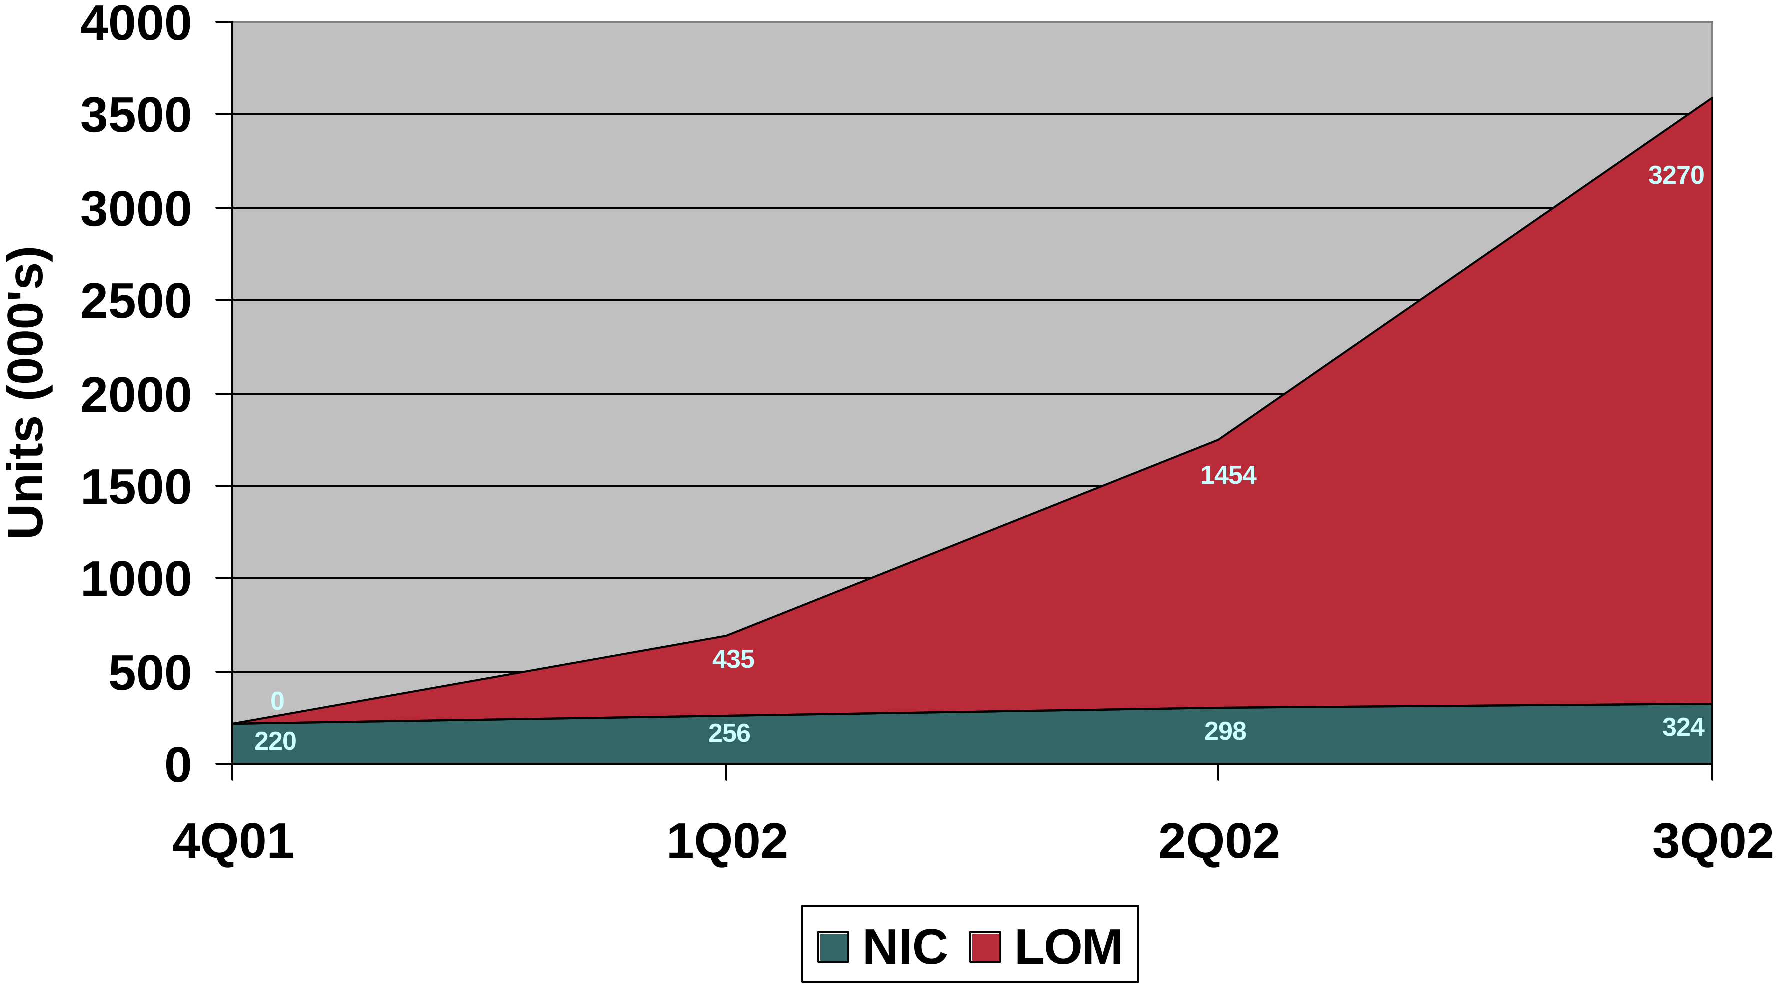

Gigabit Intelligent Campus Network Design

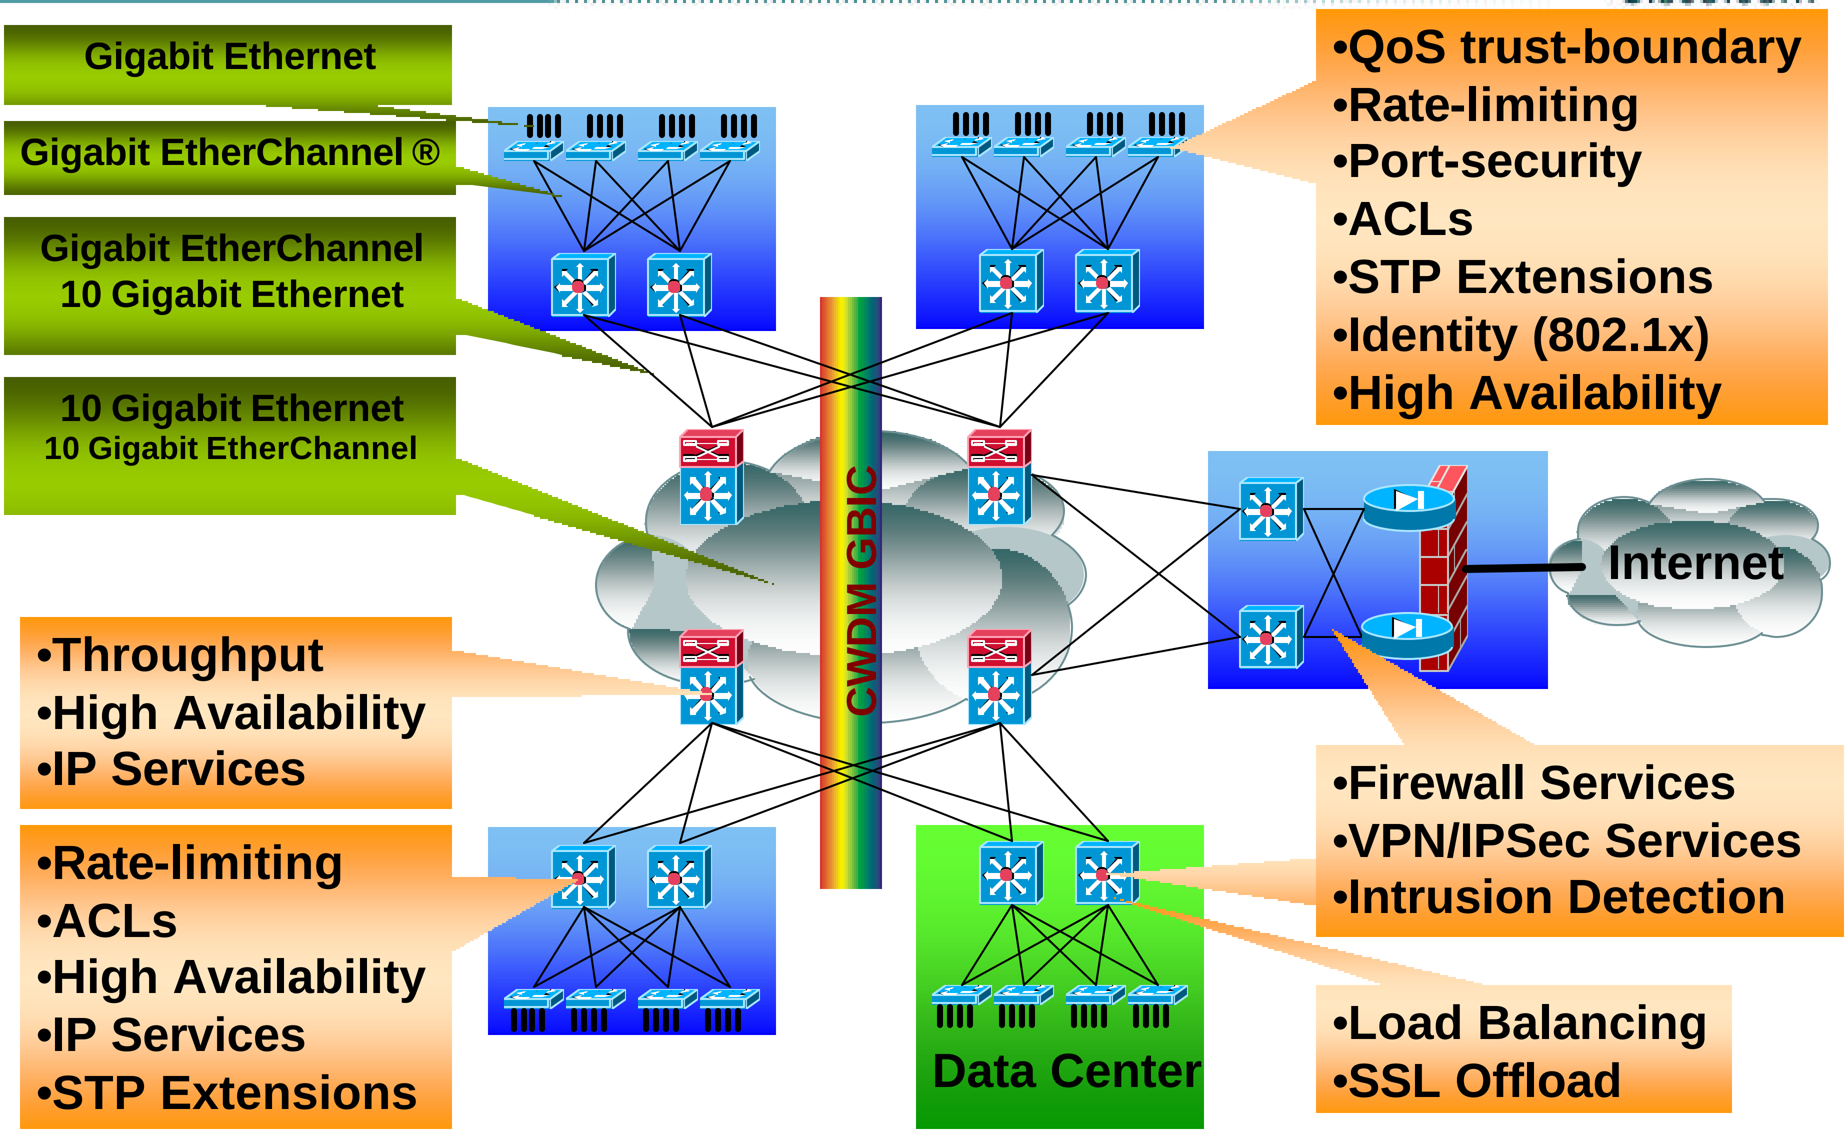

Gigabit Campus Network Design

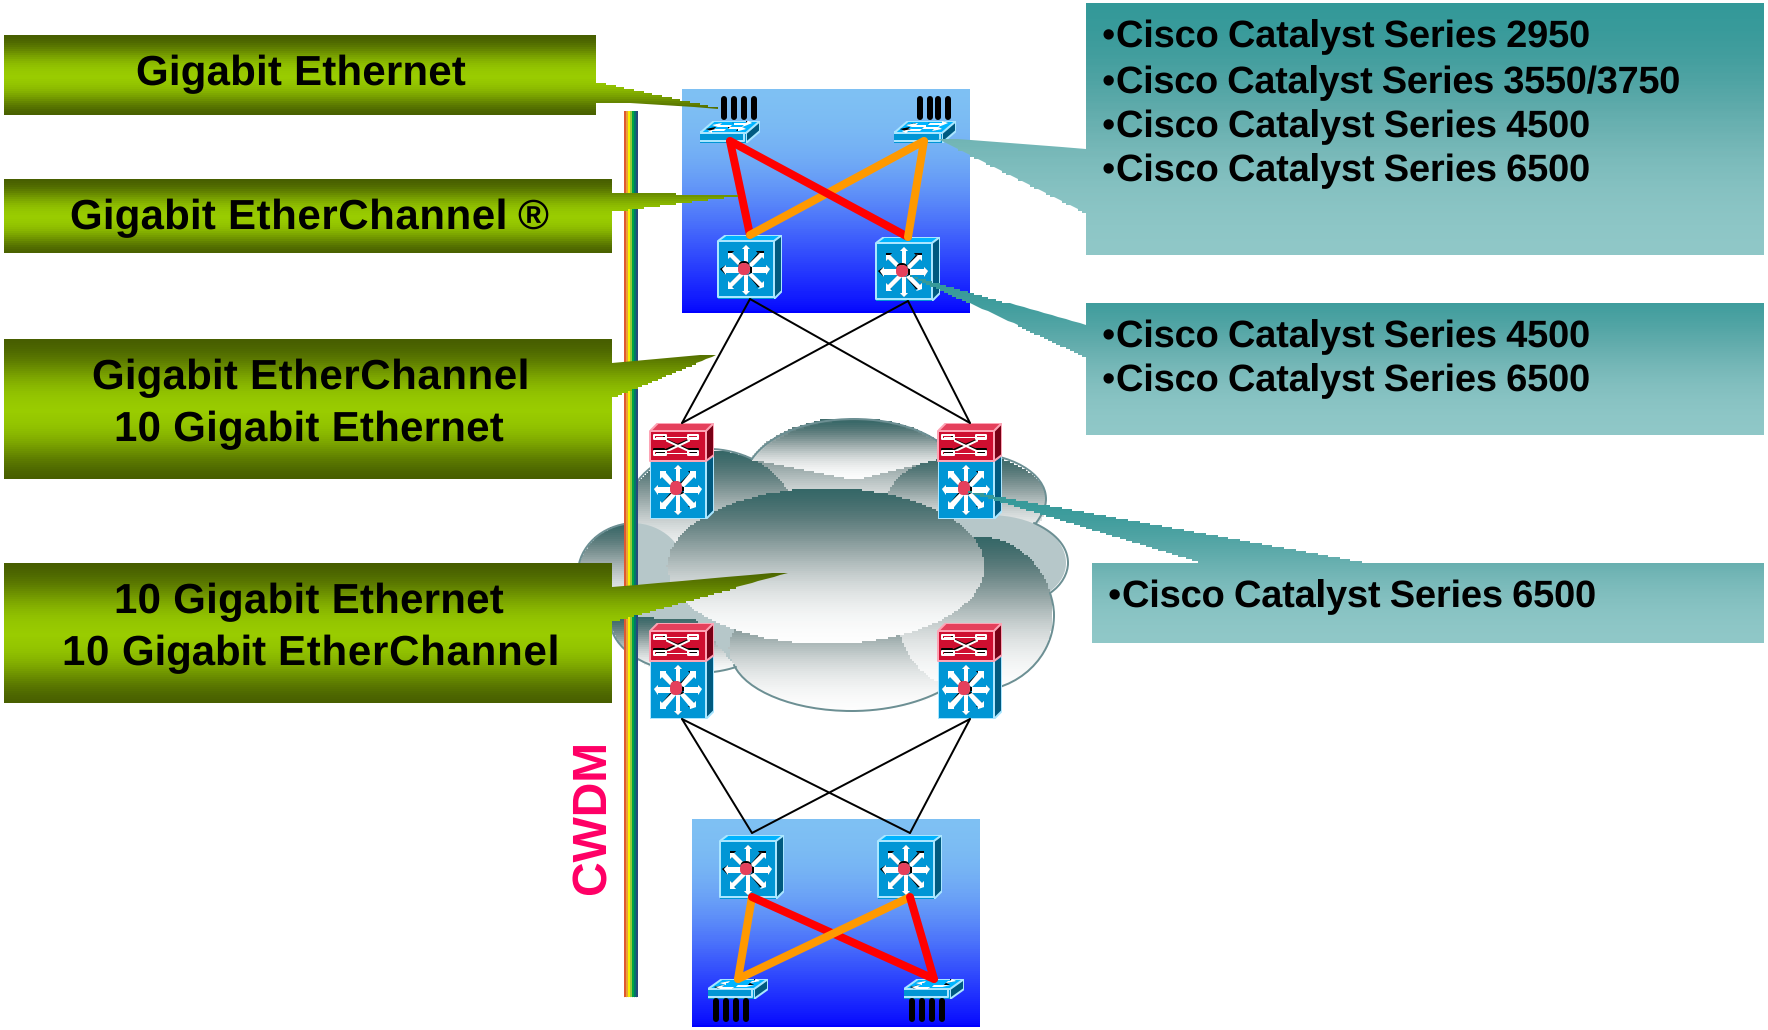

Gigabit in the Cisco Catalyst Series line of Switches

Cisco Catalyst Series	1 GE Uplinks	10/100/1000 Ports	10 GE Ports	1GE fiber (GBIC & SFP)
2950	2 ✓	2/module ✓	✗	2/module (GBIC) ✓
3550	2 ✓	10/module 50/stack ✓	✗	10/module (GBIC) 50/stack ✓
3750	4 ✓	24/module 216/stack ✓	✗	4 SFP uplinks ✓
4500	2/sup ✓	48/blade 240/chasis ✓	✗	48/blade (SFP) 18/blade (GBIC) 240/chasis ✓
6500	2/sup ✓	48/blade 384/chasis ✓	4/blade (Xenpak) 32/chasis ✓	48/blade (SFP) 16/blade (GBIC) 384/chasis ✓

Over Subscription

- **1:1 v. 4:1 v. 8:1 designs**
- **True traffic patterns are random and bursty in nature**
- **Large data transfers happen much faster minimizing traffic overlap and congestion – eg. 9 seconds vs 85 seconds for 1 GB**
- **Most mission-critical business applications and most Web transactions use TCP**
- **TCP is adaptive, rate based, and connection-oriented; it is a well behaved protocol especially when oversubscribed.**
- **TCP applications run as fast as they can, but gracefully back down when faced with congestion.**

Over Subscription

- Multimedia and IP Telephony traffic use the UDP/RTP protocol
- UDP applications don't have feedback mechanisms so when faced with congestion, these protocols don't back off – frames are just dropped, degrading the quality of what is received.
- QoS is designed to deal with the problems caused by network oversubscription
- To further tame TCP for the benefit of UDP/RTP, WRED can be used
- Test data at : http://172.19.191.11/gttd.htm

Gigabit and QoS

- **With higher bandwidth availability and significantly faster delivery of information on and off the wire the paradigm for QoS changes**
- **QoS will continue to be used to guarantee and protect the service for latency and jitter sensitive applications such as voice and video**

Gigabit and QoS

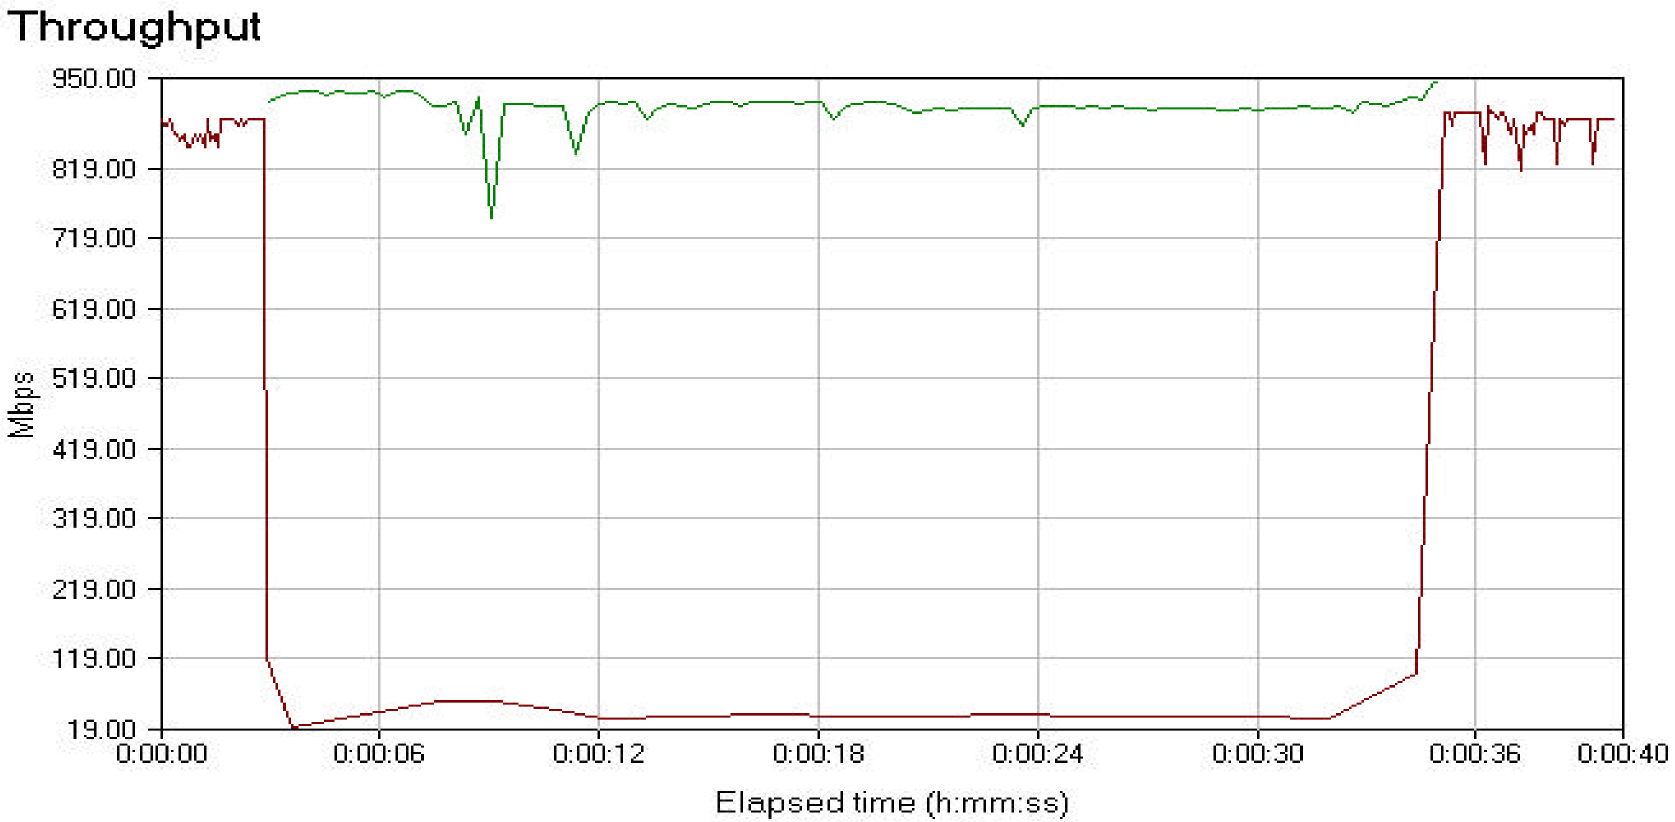

- **TCP (red) and UDP (green) Streams with Gig attached hosts**
- **1GB data using TCP with 0MB Loss**
- **3.7GB data using UDP with 23MB lost**
- **15K of 22.5M datagrams lost - 154 max consecutive loss**

Jumbo Frames Support

- Cisco Catalyst Series 6500 and 4500 on non-blocking modules only
- 6748-GE-TX with Sup720 will support but this is a future product
- 3750 supports on the 10/100/1000 configurations
- TCP Offload Engines (TOE) alleviate the need for jumbo frames unless the application requires
- MTU discovery
  - MTU Falls to lowest MTU in path

Tuning TCP

Improving Throughput – TCP Windows

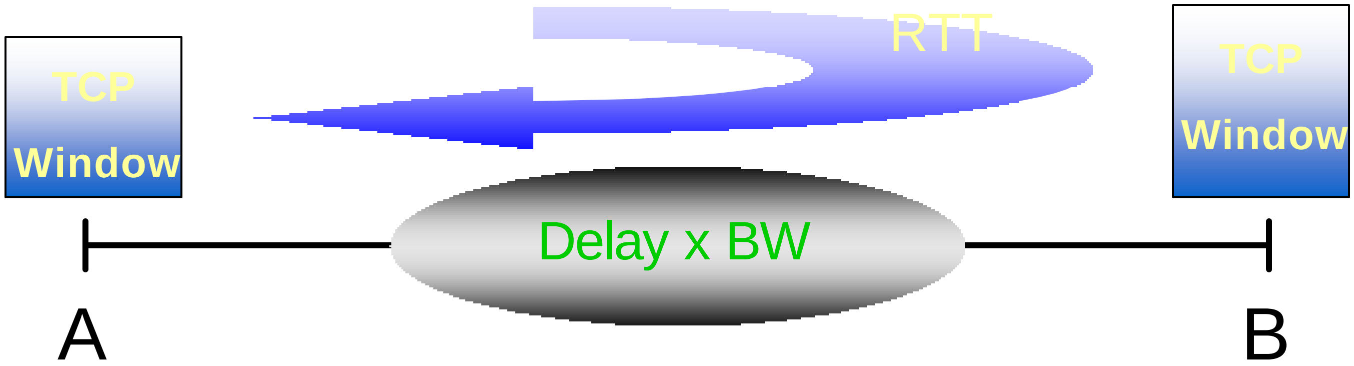

Window size below RTT x BW / 8bits
Results in low throughput

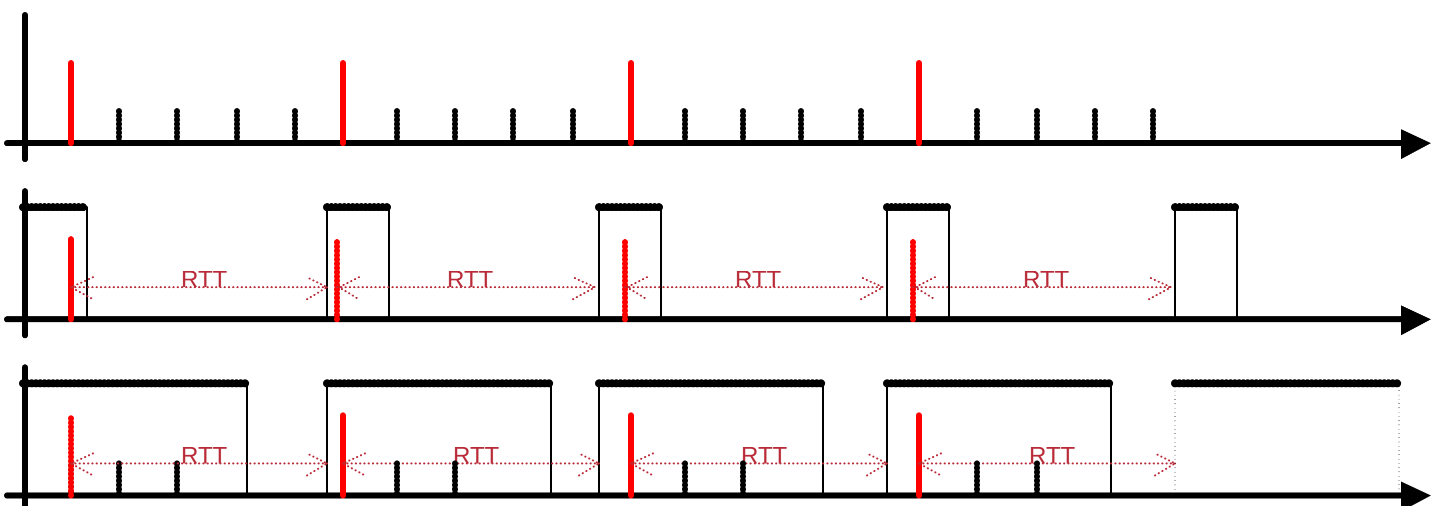

Improving Throughput – TCP Windows

Cisco.com

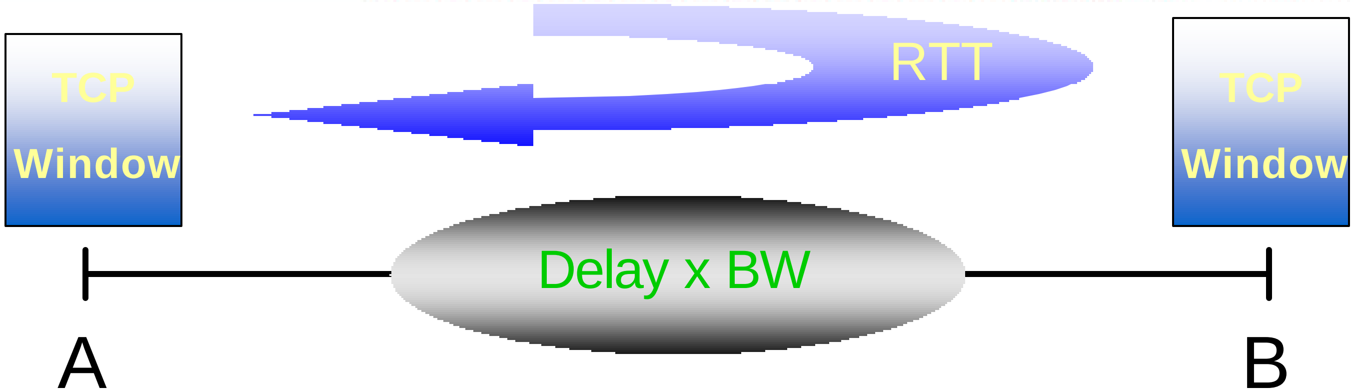

Window size > RTT x BW / 8bits
(for max throughput)

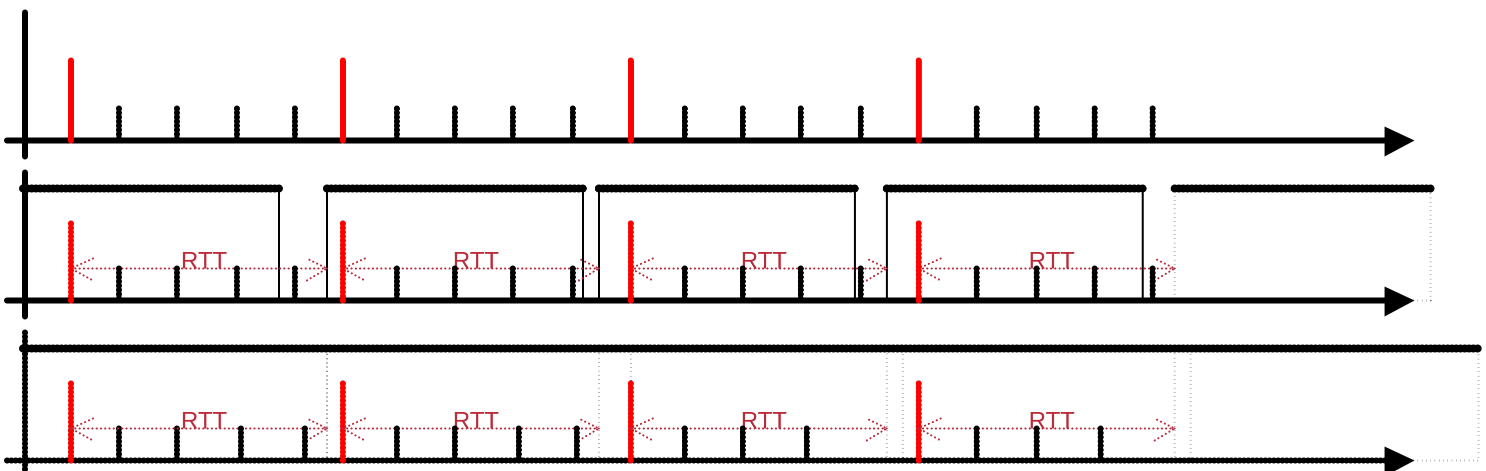

TCP Windows

Cisco.com

- **Windows control the amount of data that is allowed to be “in flight” in the network**
- **Maximum throughput is one window full per round trip time**
- **The sender, receiver, and the network each determine a different window size**

TCP Throughput (window/rtt)

Cisco.com

- **The smallest of three windows determines throughput**
- **sbuf, or sender side socket buffers rwin, the receive window size cwin, TCP congestion window**
- **Receive window (rwin) and/or sbuf are still the most common performance limiters**

  E.g. 8kB window, 87 msec ping time = 753 kbps

  E.g. 64kB window, 14 msec rtt = 37 Mbps

Improving Throughput – TCP Windows

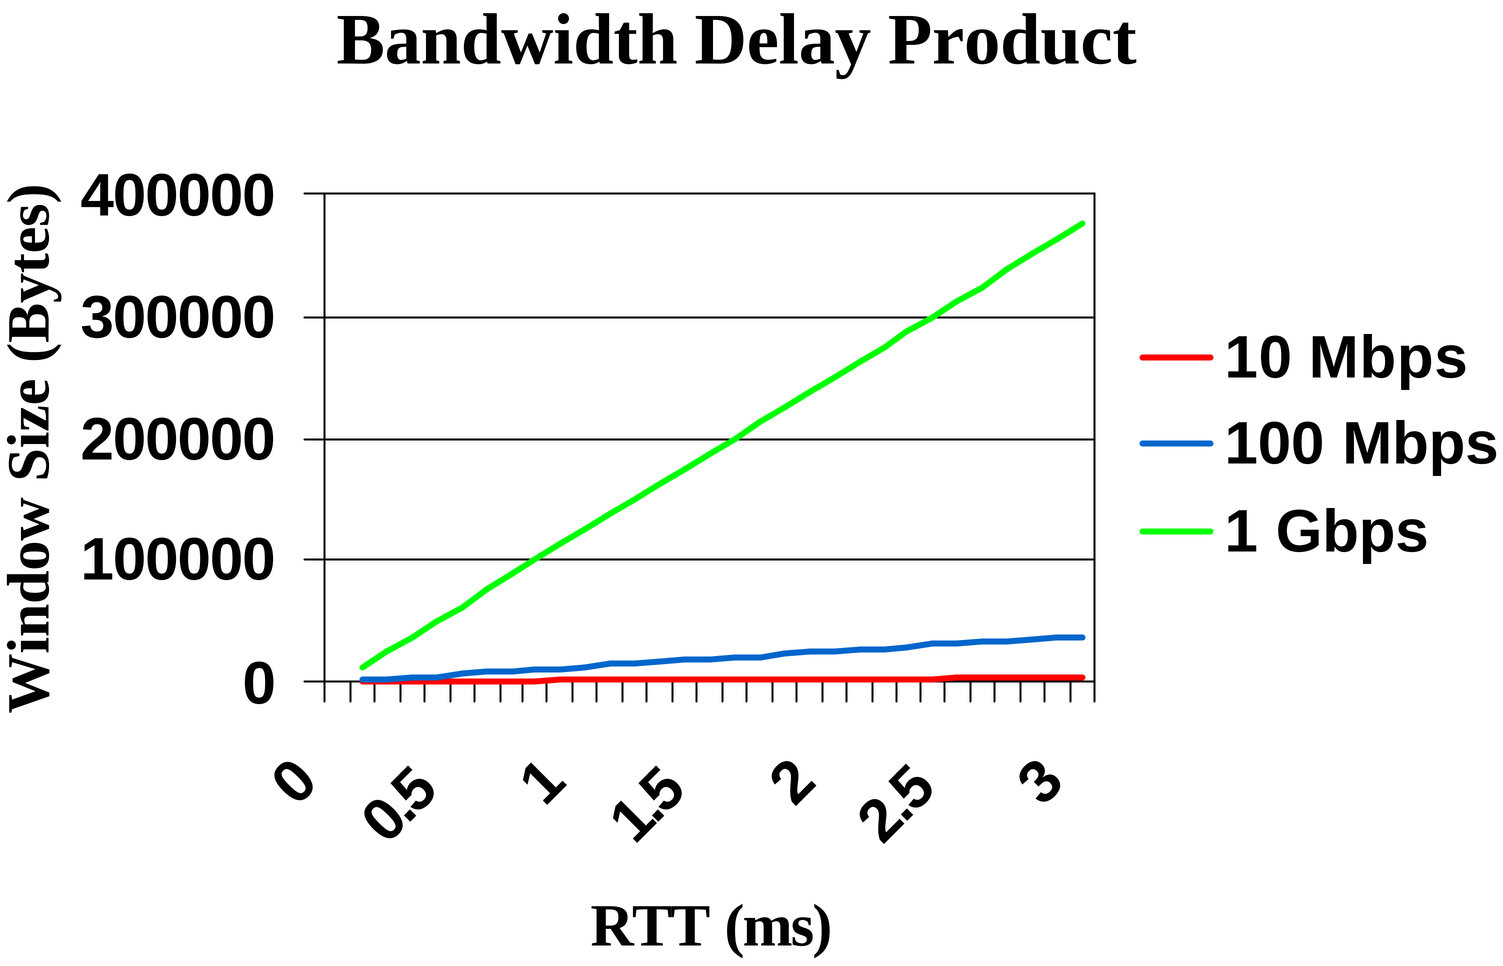

Bandwidth*Delay Product and TCP

- TCP needs a receive window (rwin) equal to or greater than the BW*Delay product to achieve maximum throughput
- TCP needs sender side socket buffers of 2*BW*Delay to recover from errors
- You need to send about 3*BW*Delay bytes for TCP to reach maximum speed

Bandwidth*Delay Product

- The number of bytes in flight to fill the entire path
- Includes data in queues if they contributed to the delay
- Example:
    - 100 Mbps path
    - ping shows a 75 ms rtt
    - BDP = 100 * 0.075 = 7.5 million bits (916 KB)

What to Do

- **Make sure your HPC apps offer sufficient receive windows and use sufficient send buffers**
- **But don’t run your system out of memory**
- **Find out the rtt with ping, compute BDP**
- **Can tune system wide, by application, or automatically**
- **Check your TCP for high performance features**
- **Look for sources of loss**
- **Watch out for duplex problems (late collisions?)**

TCP Window Size and Throughput

- **In our tests we initially ran a 32k windows size, then increased it to 64K, 128K, 256K, 300K and finally 1M. The performance we saw is as follows:**

32K	**720mbps average**
64K	**886mbps average**
128K	**903mbps average**
300K	**936mbps average**
1M	**941mbps average**

Host Side Changes

Cisco.com

- **Endstation – PCI Bus -> PCI-X -> PCI-X+**

  32 bit (33MHz)-> 64 bit (66MHz) -> 64 bit (133MHz)

- **Processor Speed**

  P3 -> P4/Athlon/Xeon

- **Memory**

  Double Data Rate (DDR) and Intel RAMBUS memory

- **New LOM technology**

  Full Duplex Gigabit speed

- **TCP Offload Engines**

Gigabit to the Desktop Resources

Cisco and Intel Gigabit Website

http://www.gigabitsolution.com

Windows Stack Tuner

http://www.dslreports.com/front/drtcp.html

CISCO SYSTEMS

EMPOWERING THE INTERNET GENERATION