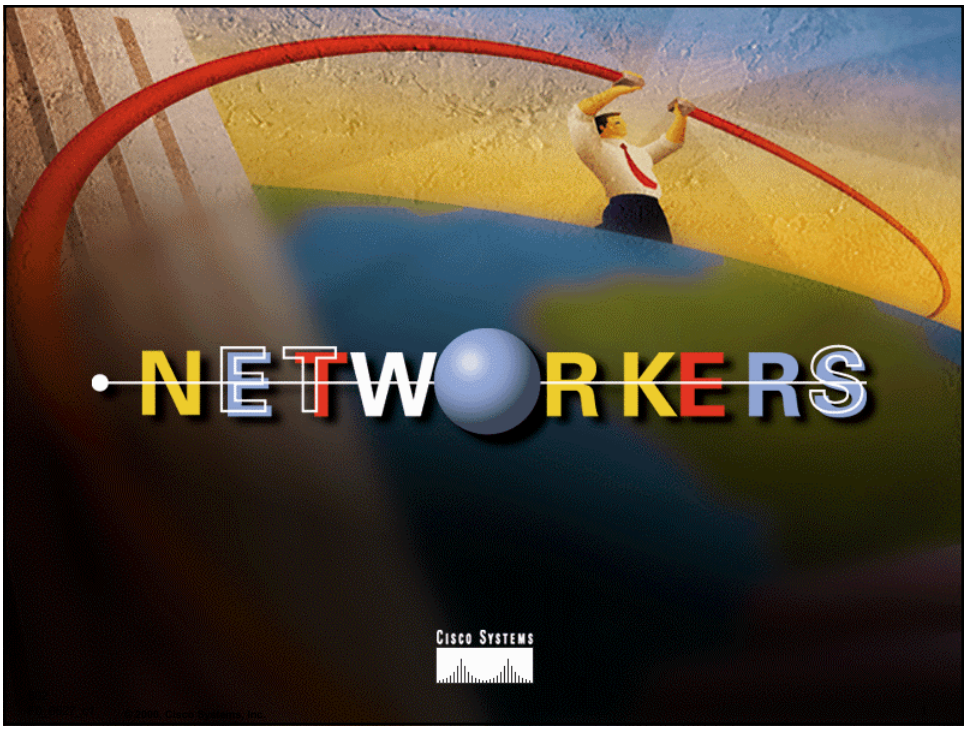

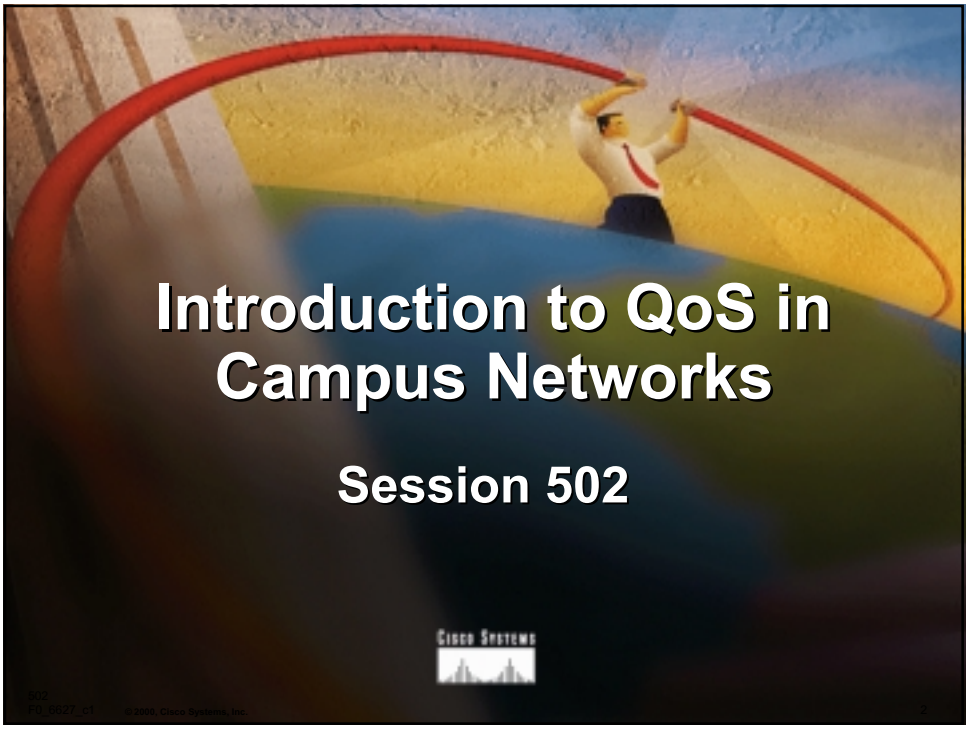

Agenda

- Networking Trends
- QoS Demystified
- QoS in Campus: Justifications
- QoS Model and Techniques
- Implementing QoS
- Summary
- Q & A

www.cisco.com

Networking Trends

Evolution

Bridges to first generation routers — Early 1980s

Multiprotocol router
Dynamic routing protocols — Mid 1980s

Multicast, switching, scaling, DHCP — 1990s

Single network for delay sensitive (voice), high bandwidth (video), standard data (e-mail) and mission-critical data (ERP) — Mid 1990s–2000

www.cisco.com

Voice, Data, Video Communications

Voice | E-Mail | Internal Web | Video | Mission Critical

www.cisco.com

What Is Mission Critical?

- ERP applications
- SNA applications
- Selected physical ports
- Selected hosts/clients

www.cisco.com

Enterprise Resource Planning (ERP) Networked Applications

- Order entry
- Finance
- Manufacturing
- Human Resources
- Supply chain management
- Sales force automation

ORACLE

SAP

PEOPLESoft

BaaN

www.cisco.com

Trend—Integrated Multiservice Network: Data, Voice, Video

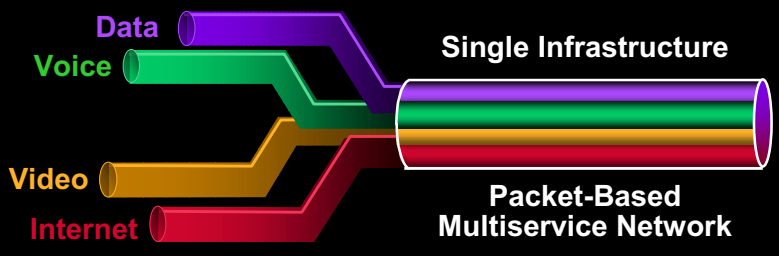

- Integration of data, voice, and video services into a single packet-based infrastructure using IP
- Both in enterprise and public service provider networks infrastructure

www.cisco.com

Data, Voice, Video Benefits of Integration

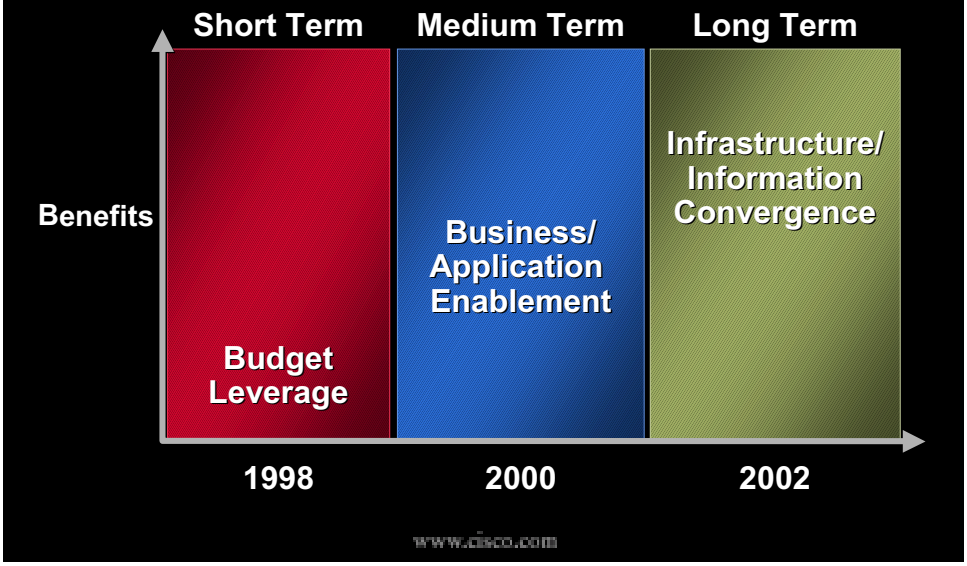

Agenda

- Networking Trends
- QoS Demystified
- QoS in Campus: Justifications
- QoS Model and Techniques
- Implementing QoS
- Summary
- Q & A

www.cisco.com

QoS—What Is It ?

- A traditional network is best-effort
- All traffic get the same service, i.e., the forwarding behavior by a network device is FIFO
- QoS prioritizes traffic into different service levels and provides preferential, forwarding treatment to some data traffic at the expense of lower-priority traffic
- QoS = preferential treatment

www.cisco.com

QoS—Why Is It Needed? Motivation for QoS Is Manifold

- Integrated networks carry different traffic types from a variety of business-enabling applications
- Business drivers and policies dictate preferential treatment for some type of traffic over other(s)
- Convergence of voice and data networks force us to consider servicing two different types of traffic on a single wire

www.cisco.com

Different Requirements

Data	Voice
• Bursty	• Smooth
• Greedy	• Benign
• Loss Sensitive	• Loss Insensitive
• Delay Insensitive	• Delay Sensitive
• TCP Retransmits	• UDP Best Effort

www.cisco.com

QoS Components

Bandwidth

Loss

Delay

Delay Variation (Jitter)

www.cisco.com

Not All Traffic Is Created Equal

	Voice	Video	Data (Best-Effort)	Mission-Critical Data
Bandwidth	Low to Moderate	Moderate to High	Moderate to High	Low to Moderate
Random Drop Sensitivity	Low	Low	High	Moderate to High
Delay Sensitivity	High	High	Low	Moderate to High
Jitter Sensitivity	High	High	Low	Low to Moderate

www.cisco.com

What Is Driving the Need for Preferential Treatment in Networks?

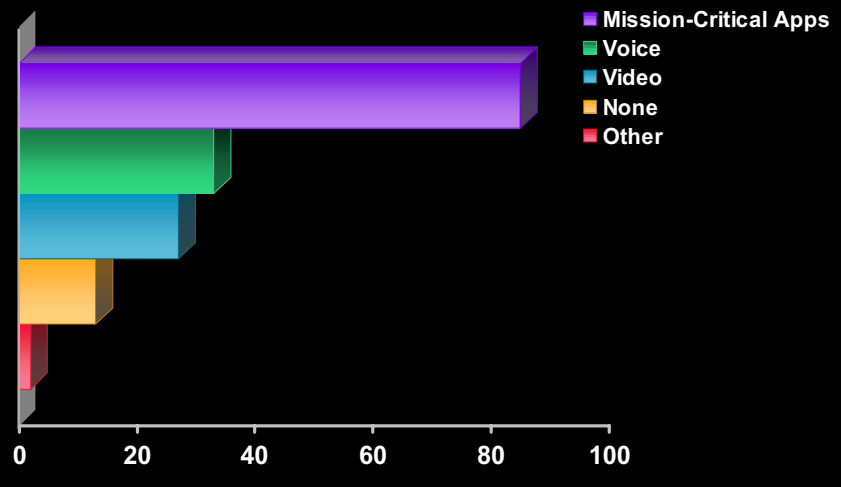

www.cisco.com

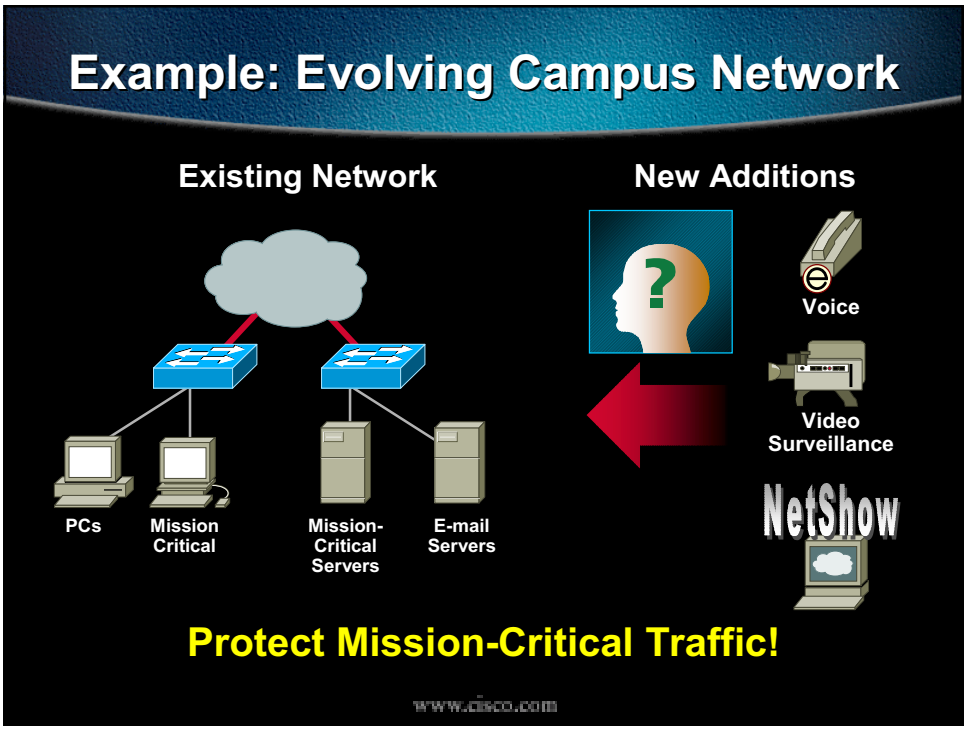

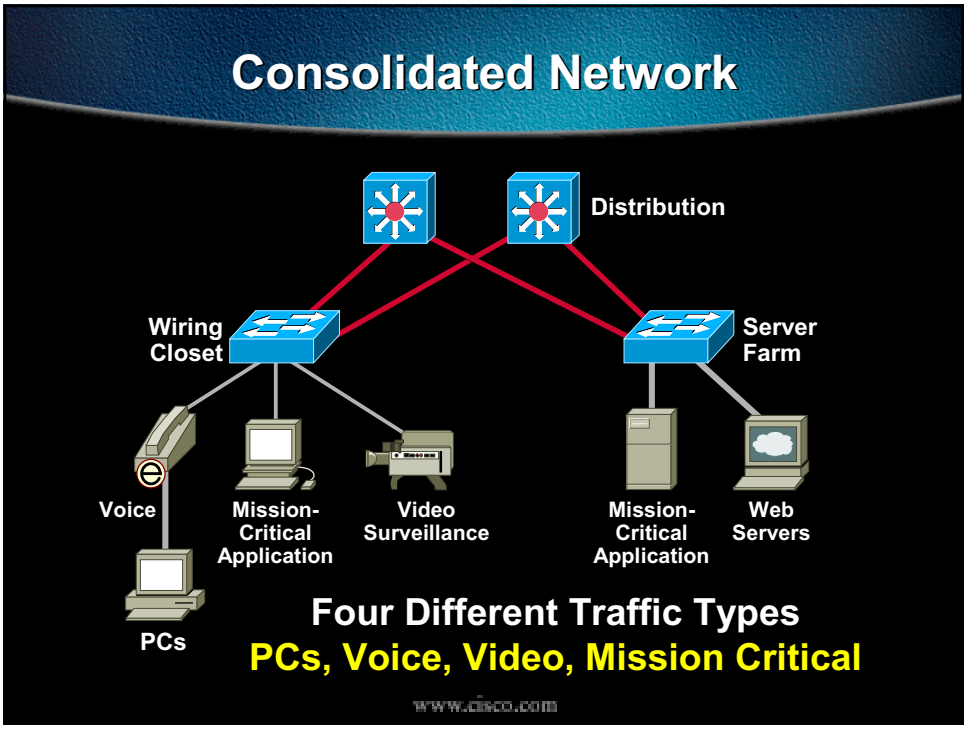

Agenda

- Networking Trends
- QoS Demystified
- QoS in Campus: Justifications
- QoS Model and Techniques
- Implementing QoS
- Summary
- Q & A

www.cisco.com

QoS in Campus Networks?

"Why Not Just Increase the Link Speed? Throw More Bandwidth at the Problem!"

www.cisco.com

Maybe You Can, Maybe You Can't

- QoS in the campus is primarily about buffer management
- Buffer management is required to control delay or drops
- TCP will eventually retransmit dropped packets
- Do you care about delay or drops?

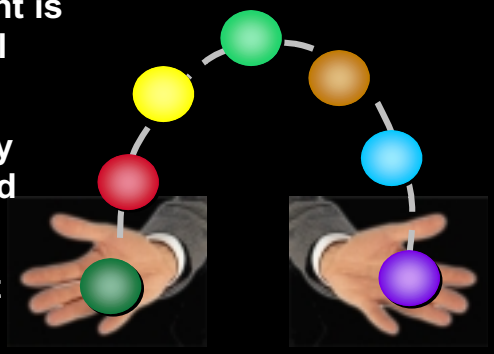

Buffers Fill for Various Reasons

- Rate mismatch (1000M to 10M Ethernet)
- Many to one (multiple interfaces talking to the same interface all with the same rate)
- Aggregation points

TCP Traffic Behavior

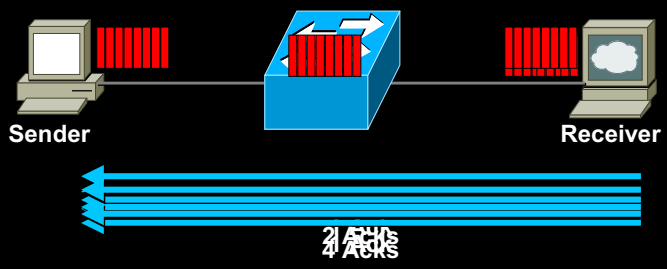

Buffers are available at sender and receiver.
Media is a buffer of sorts as well because of RTT

Because of Almost Zero RTT in LAN,
Switch Has to Buffer the TCP Window

www.cisco.com

TCP Traffic Behavior

Example: Transmission of a 13 Segment Window

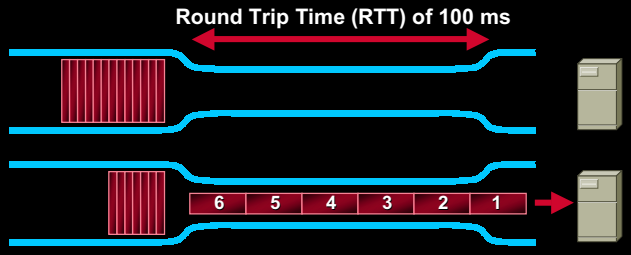

www.cisco.com

TCP Traffic Behavior

Window Perfectly Matched with RTT

Example: Link holds 12 IP segments worth of data which is equal to TCP window

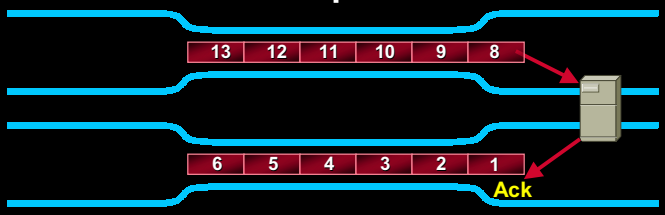

Next window is transmitted upon the receipt of Segment 1 acknowledgement

www.cisco.com

TCP Traffic Behavior

Window Not Perfectly Matched

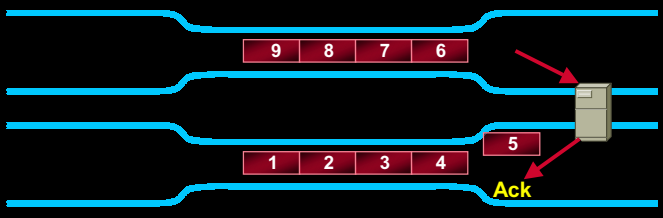

Link Buffers, Less Than the TCP Window Which Causes Transmit Buffer to Fill

www.cisco.com

Buffer Congestion Speed Mismatch

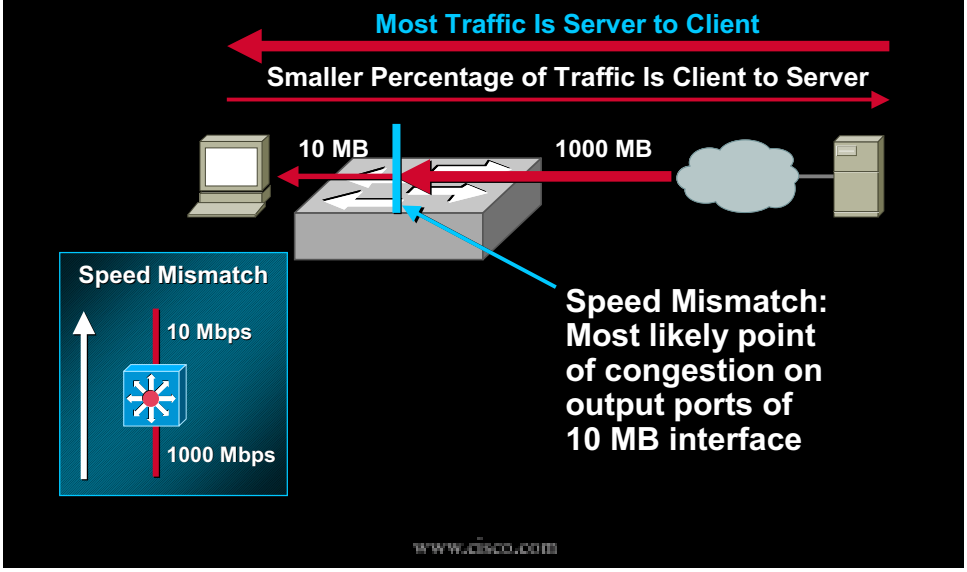

Buffer Congestion Many to One

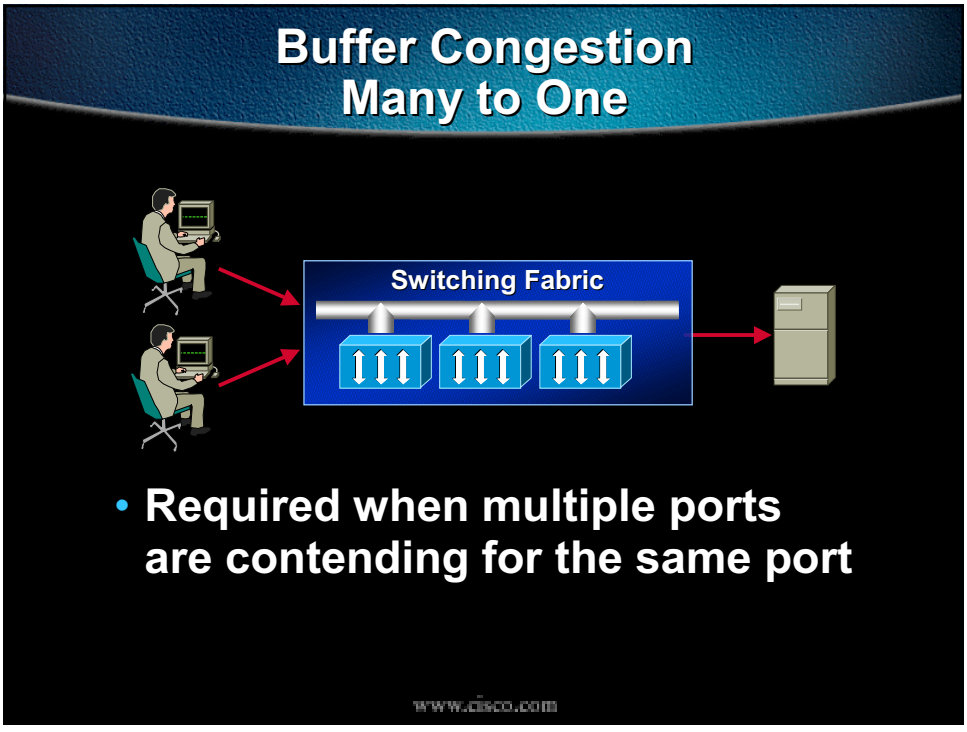

Buffer Congestion Aggregation Point

- Links oversubscribed
- Buffers full (clue: packets dropped)
- Buffering reduces loss but delay-sensitive application could be negatively impacted

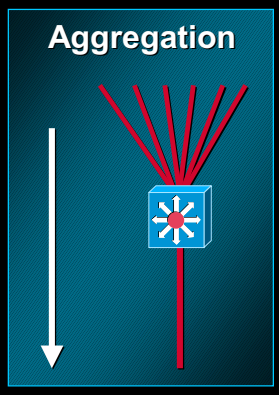

www.cisco.com

Result

Even though the average link utilization is below 100%, buffers can still fill up

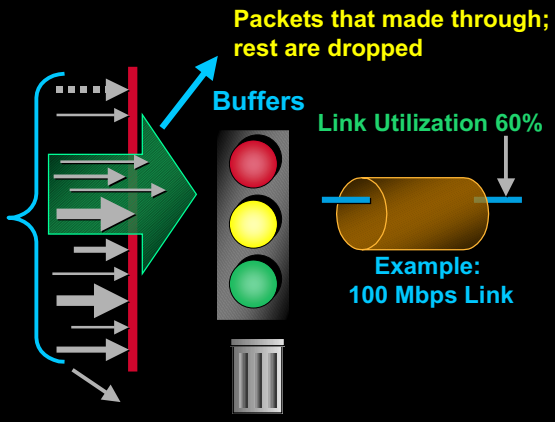

www.cisco.com

What Happens When Multiple Packets Are Dropped Within the Same Window?

- Window size is reduced exponentially since it is halved for each failed ACK
- TCP window may be reduced to the minimum
- Transmitter of the sender tends to lock-up

www.cisco.com

Voice Delay Budget: 150 ms

Cumulative Transmission Path Delay

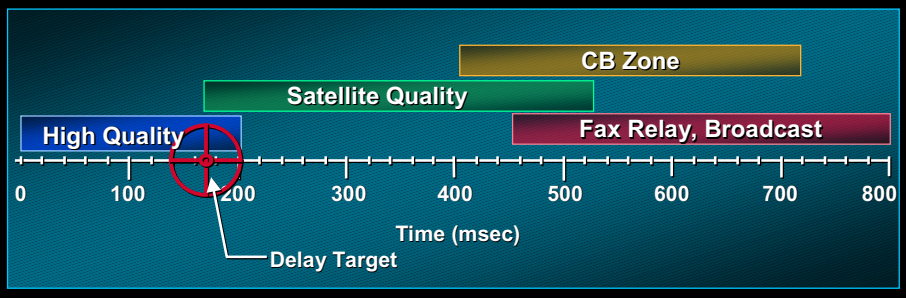

ITU G.114 Recommendation <= 150 msec 1-Way Delay

www.cisco.com

Buffer Management

Delay in Campus Network

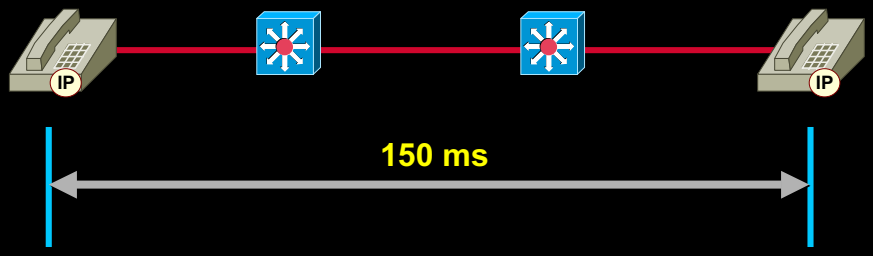

www.cisco.com

Conclusions

- Buffers can congest in LANs
- QoS required when there is congestion in buffers
- Buffer Management can help reduce loss
- Buffering reduces loss but delay sensitive application could be negatively impacted

www.cisco.com

Agenda

- Networking Trends
- QoS Demystified
- QoS in Campus: Justifications
- QoS Model and Techniques
- Implementing QoS
- Summary
- Q & A

www.cisco.com

QoS Operational Model

Admission Control	Allow or disallow frames/packets based on application type or policy (security, time of day, etc.) E.G., Application recognition
Input Classification	Assigning a label based on examining fields in the L2 (802.1p, ISL) or L3 Header (IP precedence, DSCP codepoint from Diffserv)
Input Q Scheduling	Assigning the frame/packet to an input queue and/or drop threshold based on the label assigned
Policing	Ensuring traffic rate is within specified contract if policing is turned on. Excess is dropped or marked down to a lower priority
Re-classification	Assign a new label to the frame/packet if necessary, based on configuration (policy, security, congestion, etc.) or policing action
Output Q Scheduling	Assigning the frame/packet to an output queue and/or drop threshold based on original or new label assigned

www.cisco.com

Admission Control

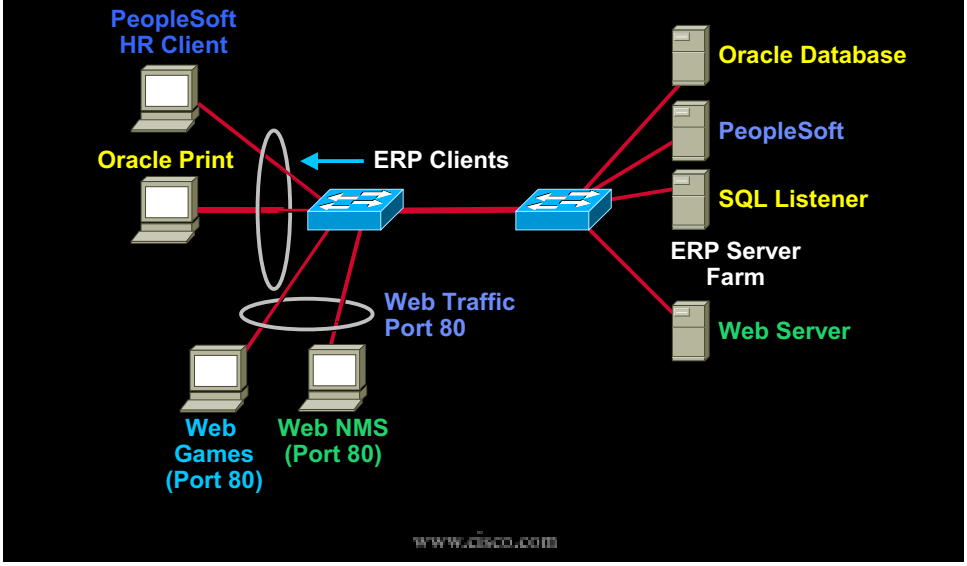

Classify at Layer 3 and Layer 2

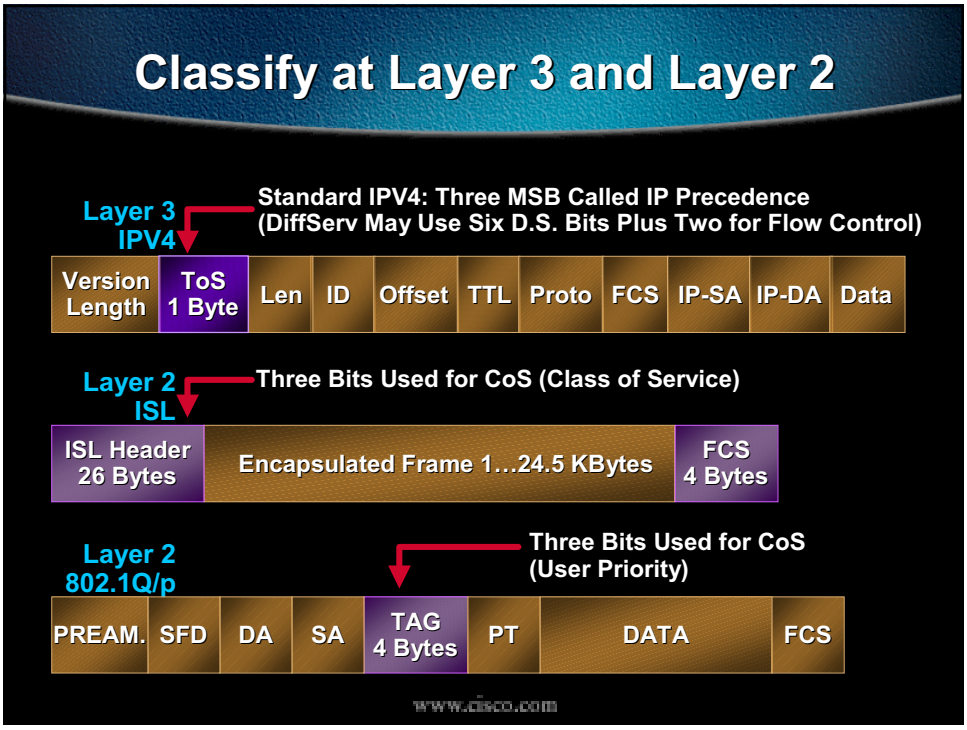

Input Classification

- Priority can be based on ToS in a packet
- Priority could also be determined by the 3-bit user-priority field in the ISL/.1Q header
  - This field is commonly called Class of Service (CoS)
- Value of zero is for best effort and seven is for the highest priority traffic

www.cisco.com

CoS Assignment

- Untagged packets, i.e., packets without an ISL/.1Q tag such as those on an access port are assigned an initial label by the ingress port
- Tagged packets can carry a CoS assigned by say an end-station

www.cisco.com

Trusted vs. Untrusted

- Cisco's QoS model assumes that the CoS carried in a packet may or may not be trusted by the network device
- For example, some servers maybe reliable enough to correctly tag the packets
- End stations like user PCs can mostly not be trusted to tag a packet's priority correctly

www.cisco.com

CoS for Untrusted Ports

- "Incorrect Tag" means tag in a way different from what the administrative policy specifies
- Network Administrators can designate which ports can be "trusted" to tag packets correctly
- On untrusted ports, the ingress port assigns a default CoS to all packets—tagged or not

www.cisco.com

CoS Assignment (Cont.)

- On trusted ports, ingress port assigns CoS only to untagged packets
- CoS assigned by an ingress port is independent of the VLAN, i.e., on an untrusted trunking port, packets of all VLANs get tagged by the same CoS

www.cisco.com

Random Early Detection (RED)

- RED reduces long term average queue
- Packet drops are randomized throughout queue depth
- Drop rate is increased as queue depth is increased

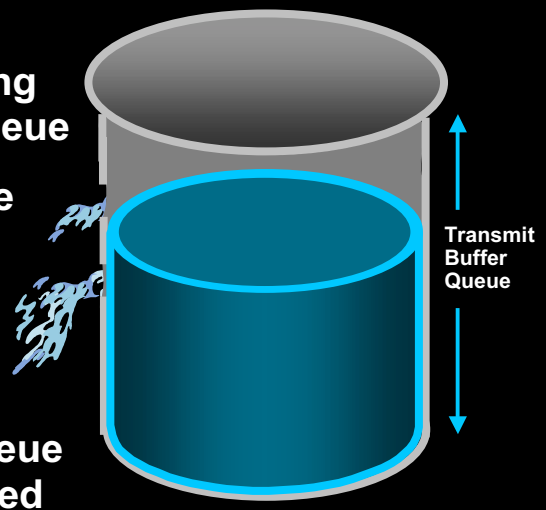

www.cisco.com

Random Early Detection Benefits

- Average delay reduced
- Reduced TCP slow start conditions
- Reduced global synchronization
- Reduces negative bias towards light users

www.cisco.com

Weighted RED

- WRED addresses:

  In the event packets need to be dropped, what class of packets should be dropped?

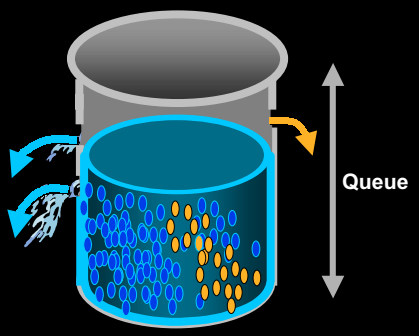

Packets classified as Blue start dropping at a 50% queue depth. drop rate is increased as queue depth is increased

Packets classified As Gold are dropped At 90% queue depth

www.cisco.com

RED for Congestion Avoidance

Uncontrolled Congestion

Managed Congestion

- Maximize goodput
- Accommodate burstiness
- Minimize delay

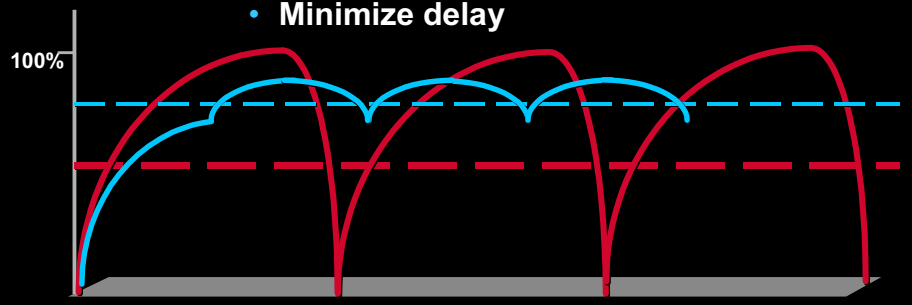

www.cisco.com

WRED Congestion Avoidance—Two Drop Thresholds

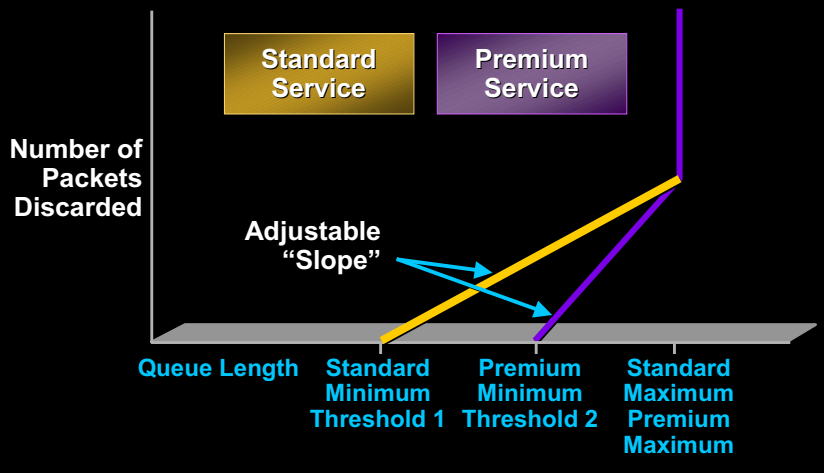

www.cisco.com

Weighted Random Early Detection (WRED)

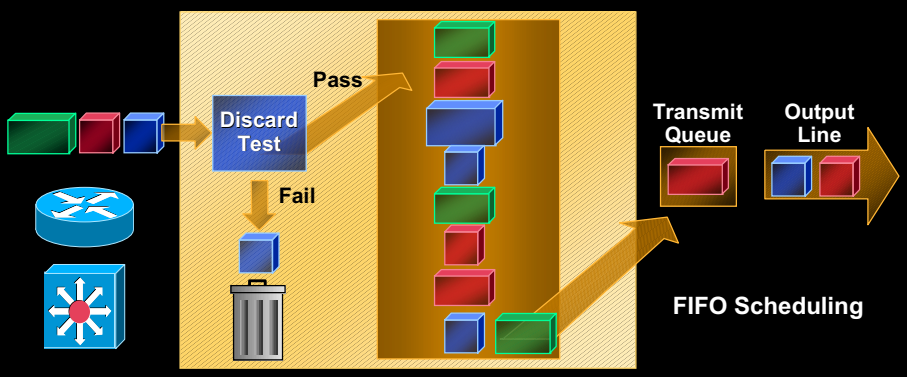

www.cisco.com

Link Management: Increased Link Utilization

Data from a Burst E1 (2.0 Mbps) Courtesy of Sean Doran

RED was turned on Friday at 10:00 am; Link utilization goes up to near 100%

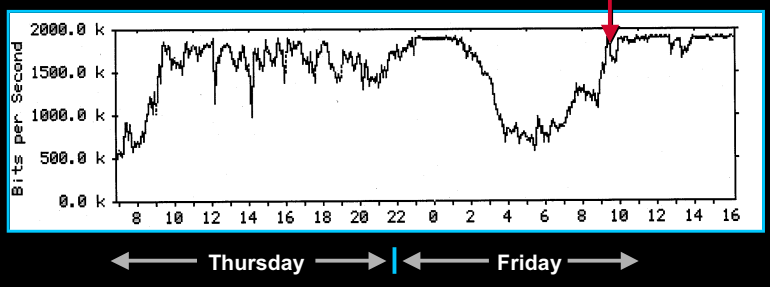

www.cisco.com

Link Management (Cont.)

Link Utilization Went from about 1700 kbps to 2.0 Mbps

Average TCP window size is higher;
Multiple packets from same window are not dropped

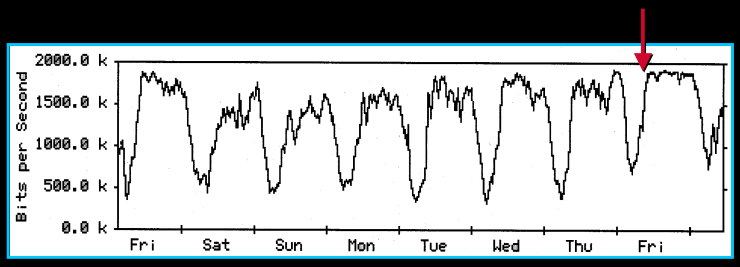

www.cisco.com

Other Benefits of RED

- **Beyond the higher link utilization other benefits that are not graphically shown include:**
  - **Small users are not negatively biased**
  - **Fewer retransmissions because there are fewer restarts**

www.cisco.com

Policing Action

- Device (switch/router) has a contract with the Network, e.g., only allow 3M of HTTP traffic
- Rest is discarded (e.g., Token Bucket) or marked down from the current value
- A new label is assigned based on the configuration

www.cisco.com

Token Bucket for Policing

- Start with a bucket full of tokens
- Tokens can be removed at a bursty rate
- Tokens are replaced at a specified constant rate

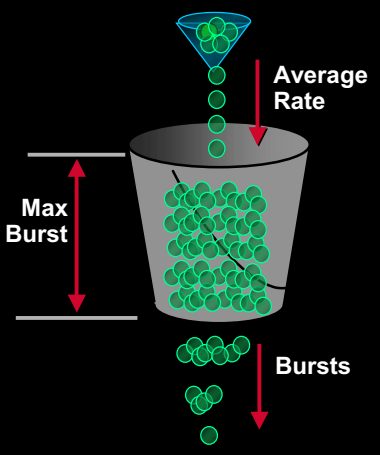

www.cisco.com

Re-Classification

- CoS can also be assigned by the Layer 2 forwarding engine
  - The user can configure mac addresses to be associated with a CoS
  - When packets arrive destined to such a mac address, their CoS is re-assigned to the value specified
- Assignment of CoS via dest-mac overrides trust and port-based CoS assignment

www.cisco.com

Output Scheduling

- On output, there are multiple queues, some high priority and some low priority
- Strict priority servicing of the queue would cause starvation of low priority traffic in the presence of high priority traffic
- However, delay requirements of high priority traffic must also be respected

www.cisco.com

Output Scheduling—WRR

Weighted Round Robin Used to Schedule Between Queues

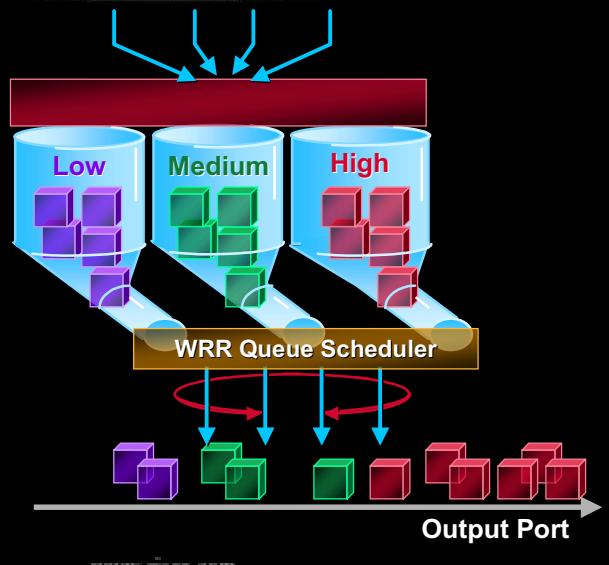

www.cisco.com

Weighted Round Robin

- The **Network Administrator specifies weights** associated with each of the queues
  - The weight ranges from 1–255
- Link is **shared in the proportion** specified by the weights

www.cisco.com

Example with Dual Queues

Two Queues with Four Drop Thresholds

Service	Delay Insensitive	Delay Sensitive
Mission Critical	7	6
Premium Best Effort	5	4
Regular Best Effort	3	2
Discard Eligible	1	0

Drop Priority

Delay Priority

www.cisco.com

QoS-Enabled Switching

Wiring Closet (Access Layer)	Backbone (Distribution Layer)	WAN Edge (Core Layer)
QoS Ingress	QoS Core	QoS WAN Edge
Admission Control Classification Policing	Congestion Avoidance Congestion Management Scheduling	Congestion Avoidance Congestion Management Policing, Shaping, Fragmentation, Header Compression

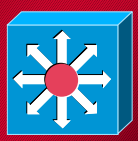

www.cisco.com

Agenda

- Networking Trends
- QoS Demystified
- QoS in Campus: Justifications
- QoS Model and Techniques
- Implementing QoS
- Summary
- Q & A

www.cisco.com

Example 1: Prioritization of IP Telephony

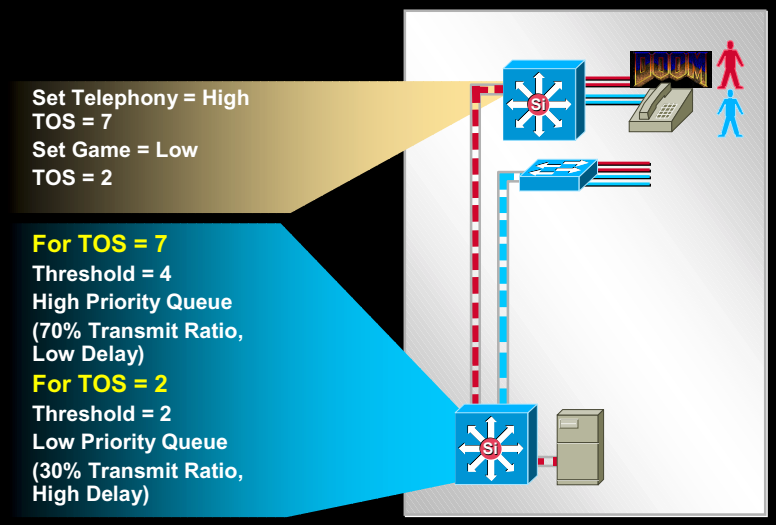

www.cisco.com

Example 2: ERP Application

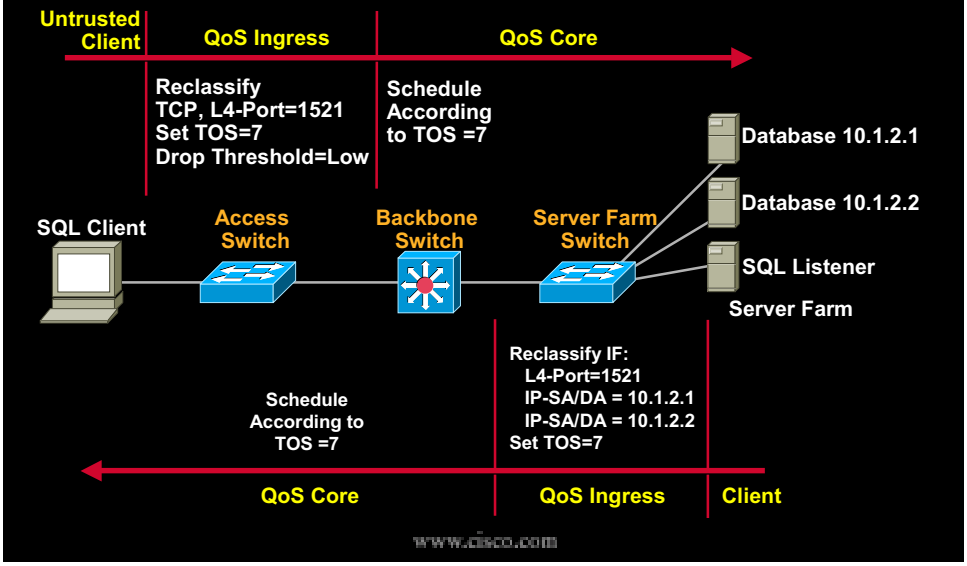

Example 3: SNA Traffic

L2 ISL or 802.1p Class of Service Reclassification

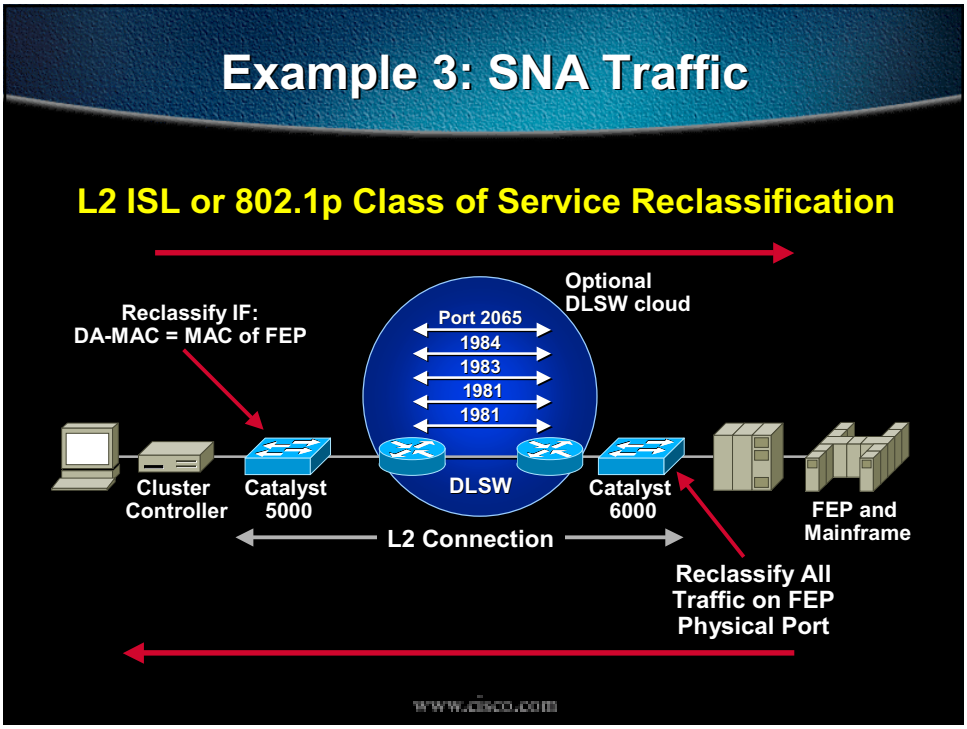

Agenda

- Networking Trends
- QoS Demystified
- QoS in Campus: Justifications
- QoS Model and Techniques
- Implementing QoS
- Summary
- Q & A

www.cisco.com

Conclusions

- QoS in campus is:
  - Meaningful
  - Important
  - Needed
- Campus (LAN) QoS requirements are emerging
- Prerequisite to D/V/V deployment, is the need to prioritize mission-critical applications

www.cisco.com

Agenda

- Networking Trends
- QoS Demystified
- QoS in Campus: Justifications
- QoS Model and Techniques
- Implementing QoS
- Summary
- Q & A

www.cisco.com