Cisco live!

Networkers

June 22 - 26, 2008 · Orlando, FL

The Power of Collaboration

CISCO

Enterprise Quality of Service

TECRST-2500

QoS Techtorial Agenda

- Introduction to QoS
  - Michael Lin/Sid Nag
- Campus QoS Design
  - Kevin Turek
- WAN and Branch QoS Design
  - Steve Short
- Network Management
  - Sid Nag
- Summary & Closure

What Is Quality of Service?

- To the end user

  User's perception that their applications are performing properly

  Voice - No drop calls, no static

  Video - High quality, smooth video

  Data - Rapid response time

- To The Network Manager

  Maximize network bandwidth utilization while meeting performance expectations

  Control Delay - The finite amount of time it takes a packet to reach the receiving endpoint

  Jitter -The difference in the end-to-end delay between packets.

  Packet Loss - relative measure of the number of packets that were not received compared to the total number of packets transmitted.

Why Enable QoS?

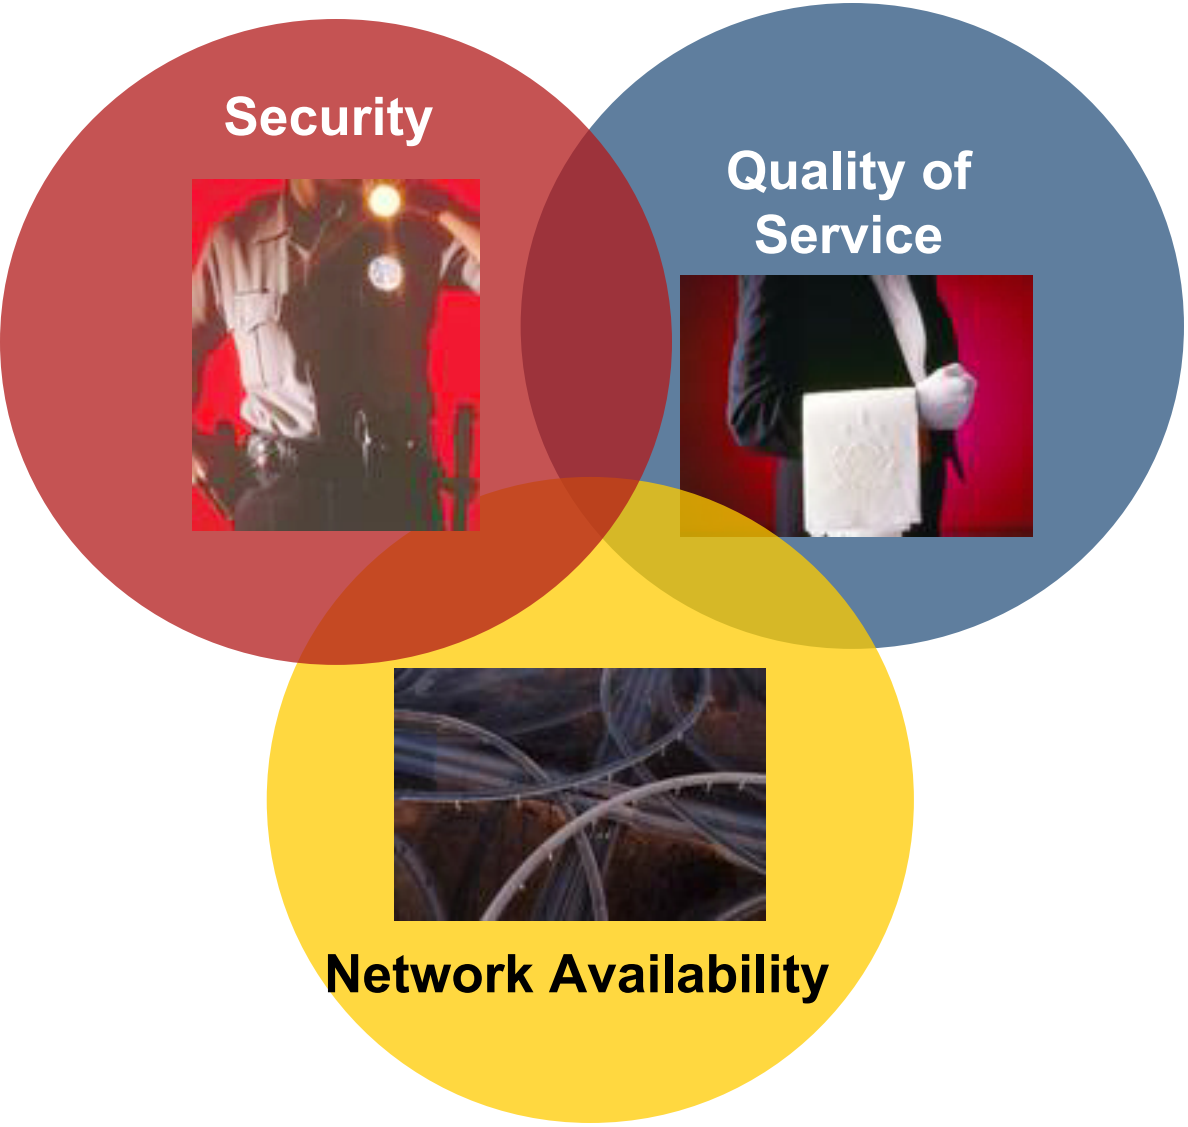

- Optimize bandwidth utilization for Video, Voice & Data apps.
- Drives productivity by enhancing service-levels to mission-critical applications
- Helps maintain network availability in the event of DoS/worm attacks

Quality of Service Operations
How does it work & essential elements

- **Classification & Marking:**

  The first element to a QoS policy is to classify/identify the traffic that is to be treated differently. Following classification, marking tools can set an attribute of a frame or packet to a specific value.

- **Policing:**

  Determine whether packets are conforming to administratively-defined traffic rates and take action accordingly. Such action could include marking, remarking or dropping a packet.

- **Scheduling (including Queuing & Dropping):**

  Scheduling tools determine how a frame/packet exits a device. Queueing algorithms are activated only when a device is experiencing congestion and are deactivated when the congestion clears.

- **Link Specific Mechanisms (Shaping, Fragmentation, Compression, Tx Ring)**

  Offers network administrators tools to optimize link utilization

QoS
Deployment Principles

How Is QoS Optimally Deployed?

1. Strategically define the business objectives to be achieved via QoS
2. Analyze the service-level requirements of the various traffic classes to be provisioned for
3. Design and test the QoS policies prior to production-network rollout
4. Roll-out the tested QoS designs to the production-network in phases, during scheduled downtime
5. Monitor service levels to ensure that the QoS objectives are being met

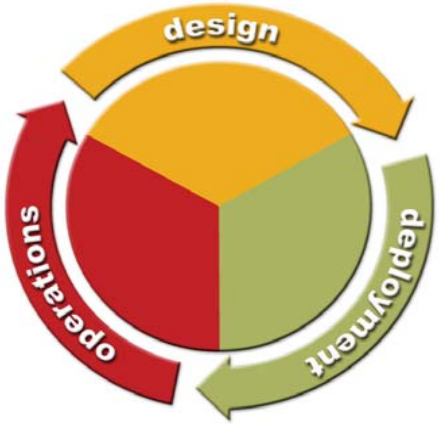

General QoS Design Principles

- Clearly define the organizational objectives

  Protect voice? Video? Data?
  DoS/worm mitigation?

- Assign as few applications as possible to be treated as "mission-critical"
- Seek executive endorsement of the QoS objectives prior to design and deployment
- Determine how many classes of traffic are required to meet the organizational objectives

  More classes = More granular service-guarantees

How Many Classes of Service Do I Need?

Example Strategy for Expanding the Number of Classes of Service over Time

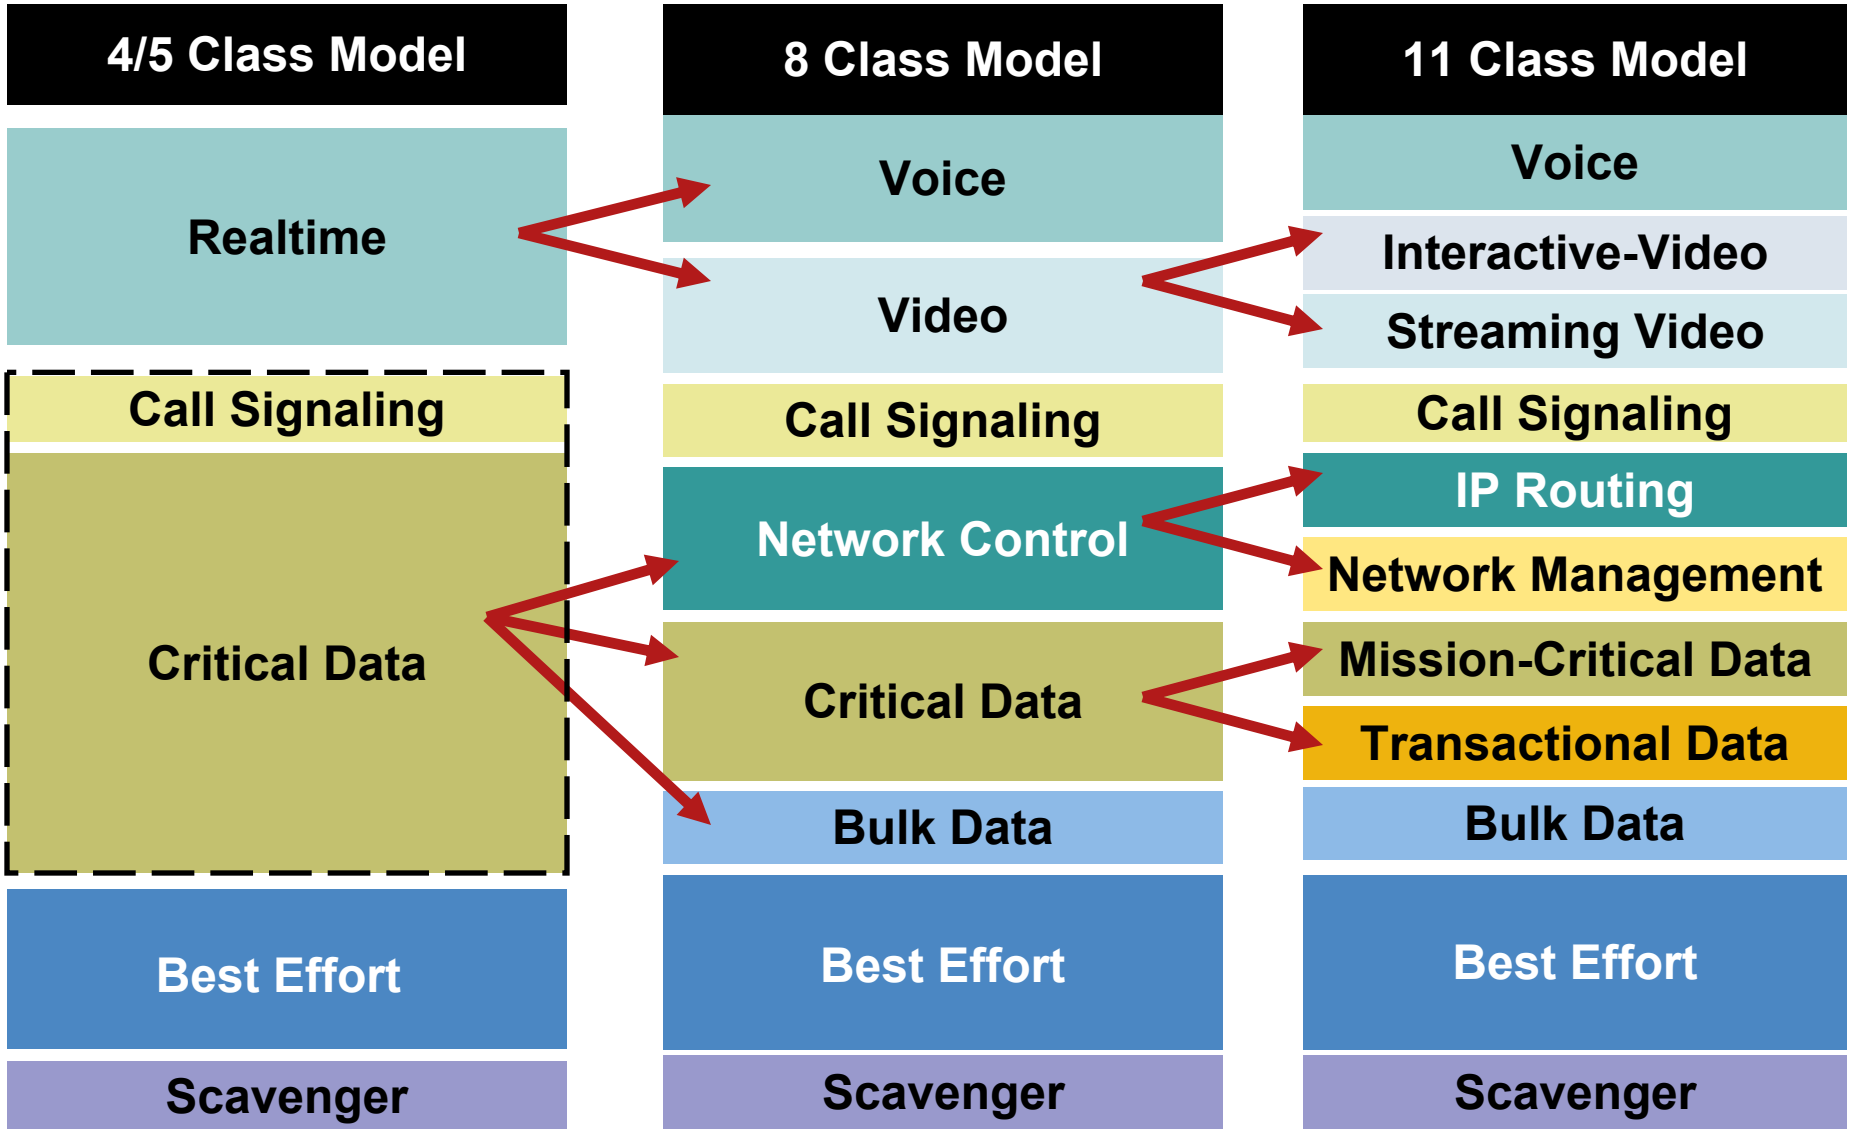

QoS Technologies Review

- QoS Overview
- Classification Tools
- Policing
- Scheduling Tools
- Shaping Tools
- Link-Specific Tools
- Signalling Tools (RSVP)
- AutoQoS Tools
- Management Tools

Classification Tools – Layer 2 Ethernet 802.1Q Class of Service

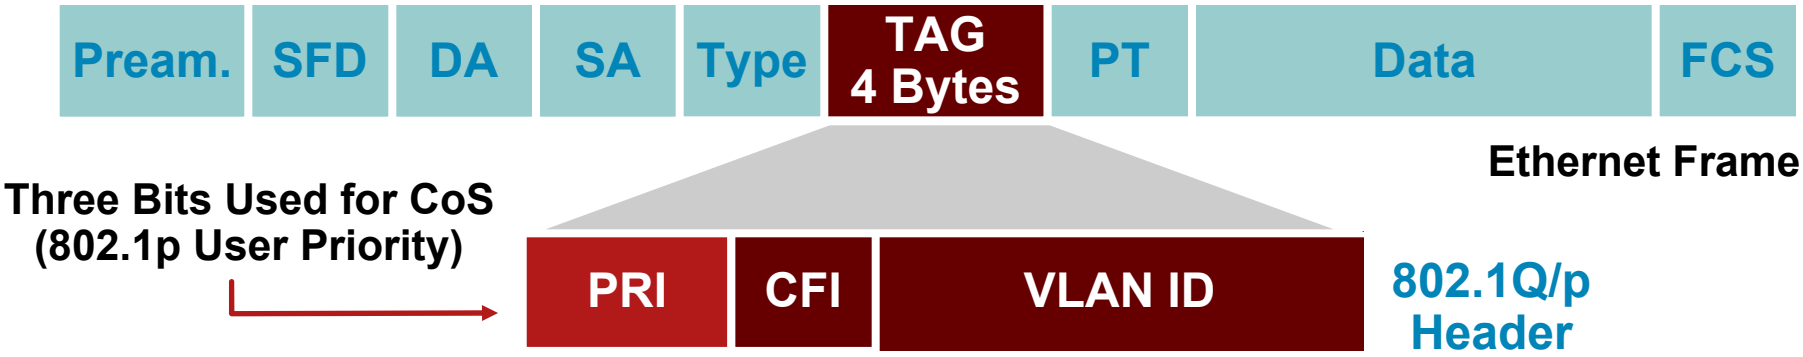

- 802.1p user priority field also called Class of Service (CoS)
- Different types of traffic are assigned different CoS values
- CoS 6 and 7 are reserved for network use

CoS	Application
7	Reserved
6	Routing
5	Voice
4	Video
3	Call Signaling
2	Critical Data
1	Bulk Data
0	Best Effort Data

Classification Tools – Layer 3
IP Precedence and DiffServ Code Points

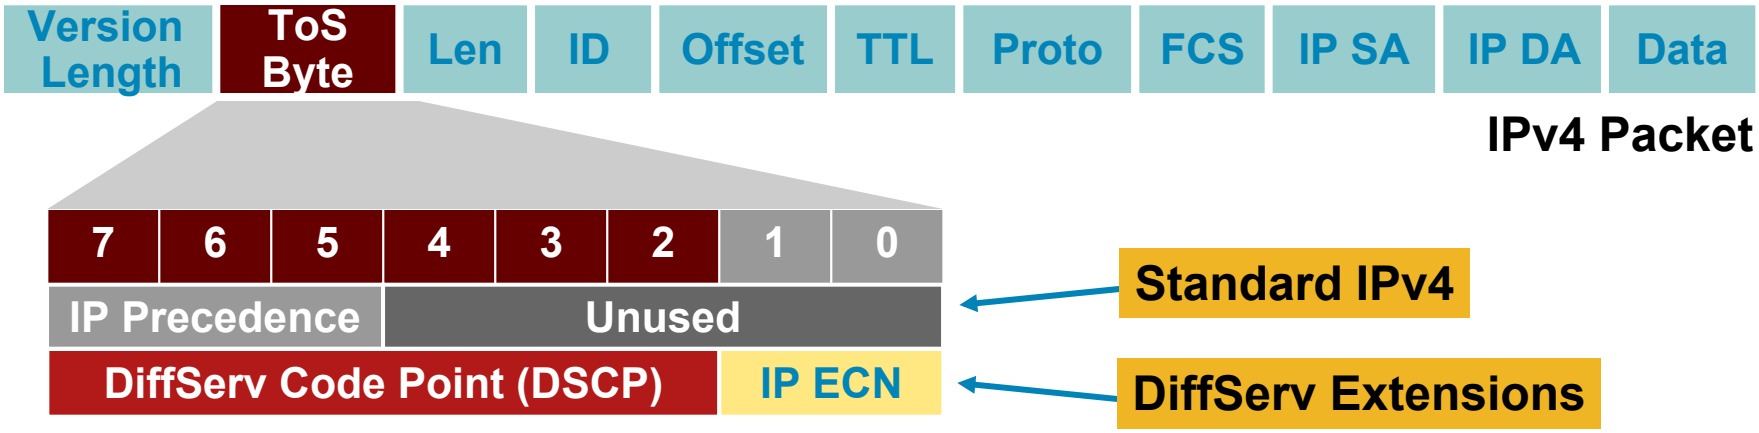

- IPv4: Three most significant bits of ToS byte are called IP Precedence (IPP)—other bits unused
- DiffServ: Six most significant bits of ToS byte are called DiffServ Code Point (DSCP)—remaining two bits used for flow control
- DSCP is backward-compatible with IP precedence

Classification Tools
MPLS EXP Bits

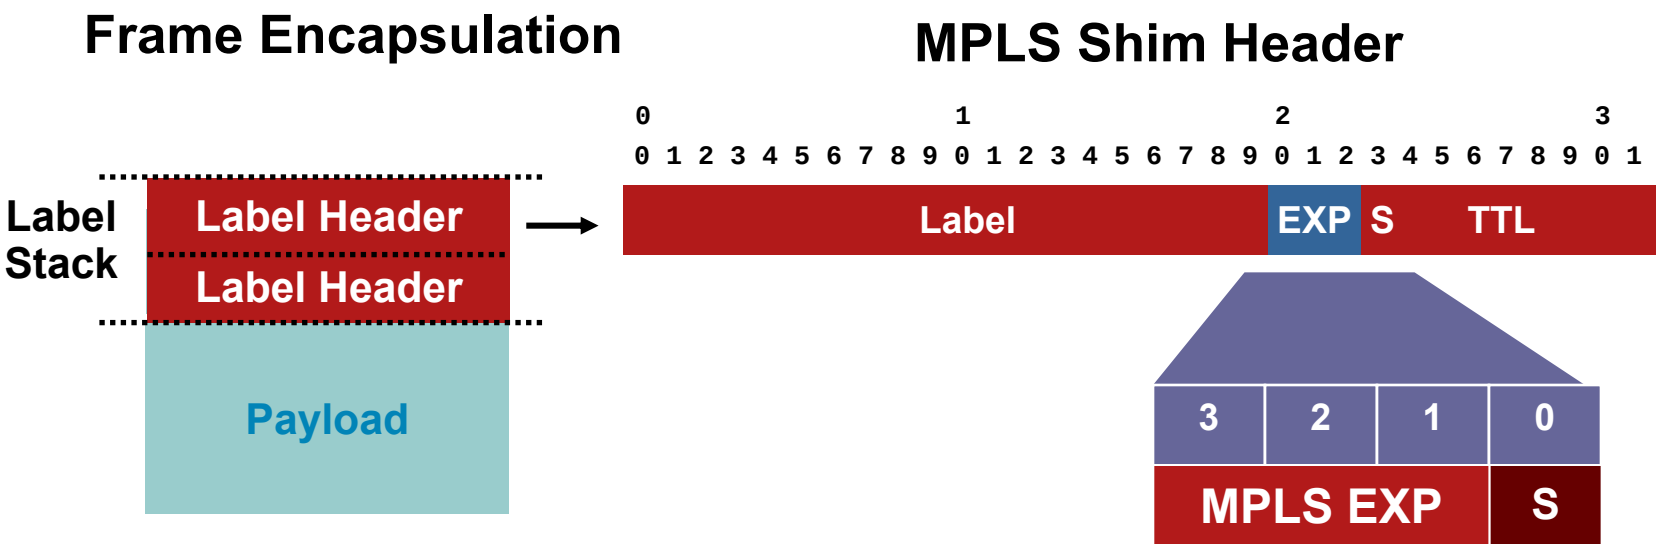

- Packet Class and drop precedence inferred from EXP (three-bit) field
- RFC3270 does not recommend specific EXP values for DiffServ PHB (EF/AF/DF)
- Used for frame-based MPLS

Classification Tools
DSCP Per-Hop Behaviors

- IETF RFCs have defined special keywords, called Per-Hop Behaviors, for specific DSCP markings
- EF: Expedited Forwarding (RFC3246)

  (DSCP 46)
- CSx: Class Selector (RFC2474)

  Where x corresponds to the IP Precedence value (1–7)

  (DSCP 8, 16, 24, 32, 40, 48, 56)
- AFxy: Assured Forwarding (RFC2597)

  Where x corresponds to the IP Precedence value
  (only 1–4 are used for AF Classes)

  And y corresponds to the Drop Preference value (either 1 or 2 or 3)

  With the higher values denoting higher likelihood of dropping

  (DSCP 10/12/14, 18/20/22, 26/28/30, 34/36/38)
- BE: Best Effort or Default Marking Value (RFC2474)

  (DSCP 0)

Classification Tools
Network-Based Application Recognition

Stateful and dynamic inspection

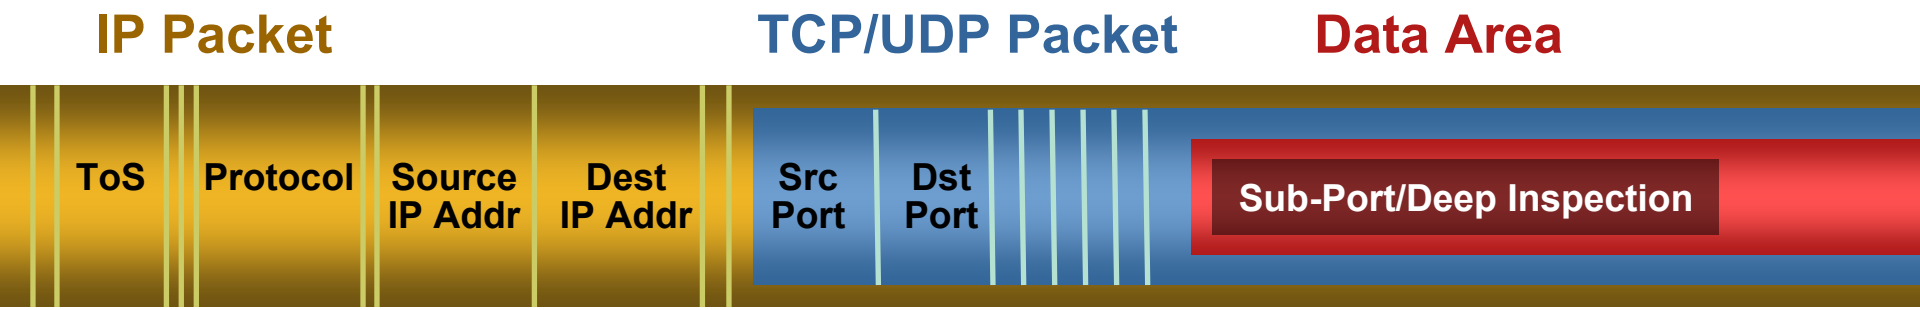

- Identifies over 90 applications and protocols TCP and UDP port numbers
  - Statically assigned
  - Dynamically assigned during connection establishment
- Non-TCP and non-UDP IP protocols
- Data packet inspection for matching values

Policing Tools

Policing Tools
RFC 2697 Single Rate Three Color Policer

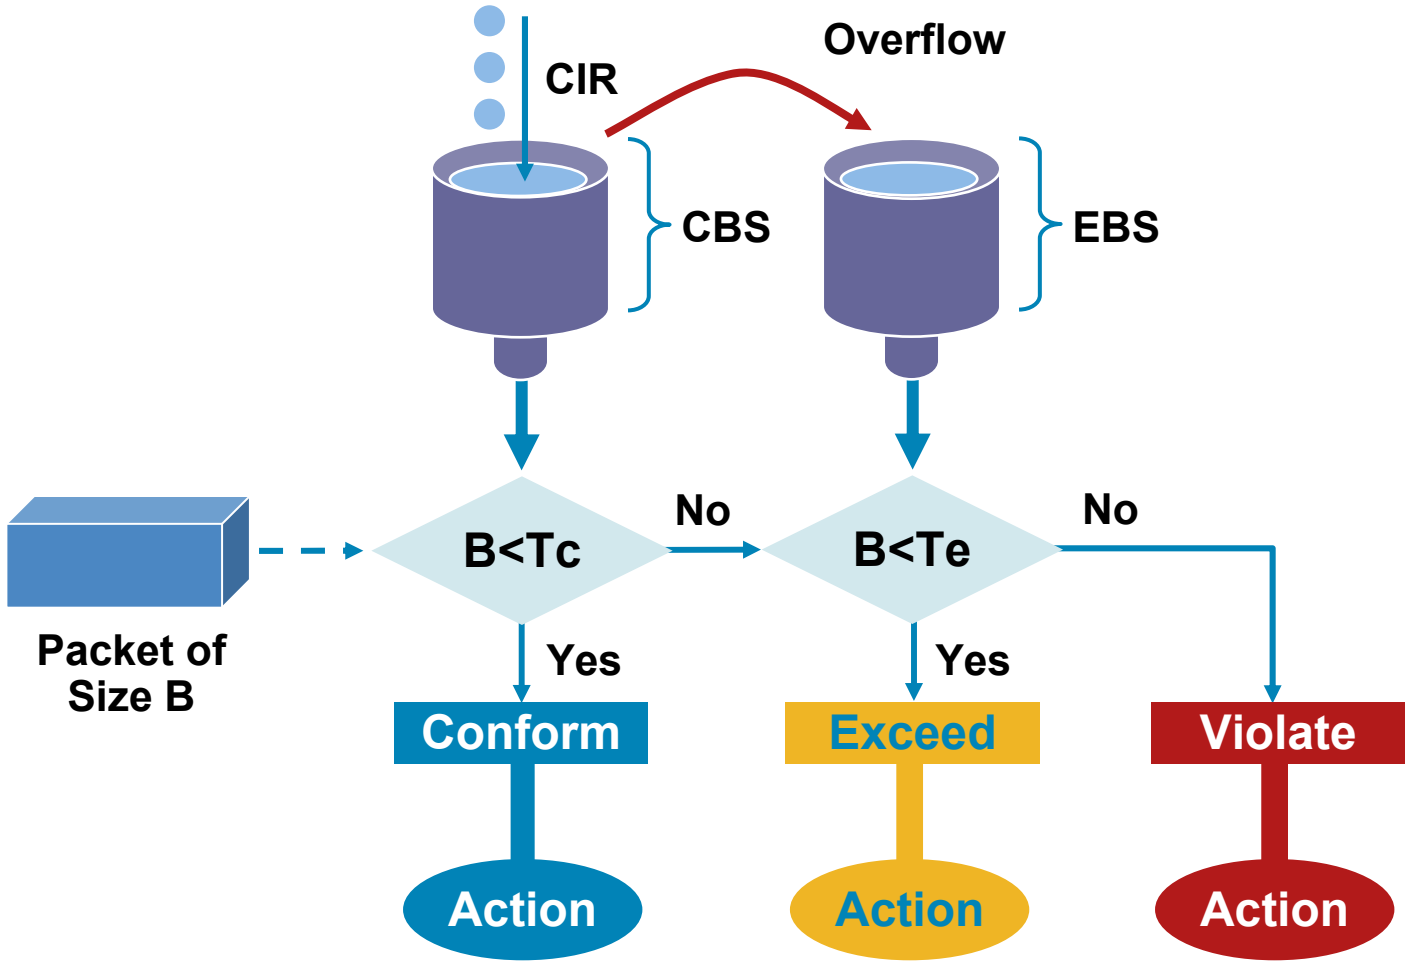

EBS=Excess BW Scheduler; CBS Committed Burst size

Policing Tools
RFC 2698 Two Rate Three Color Policer

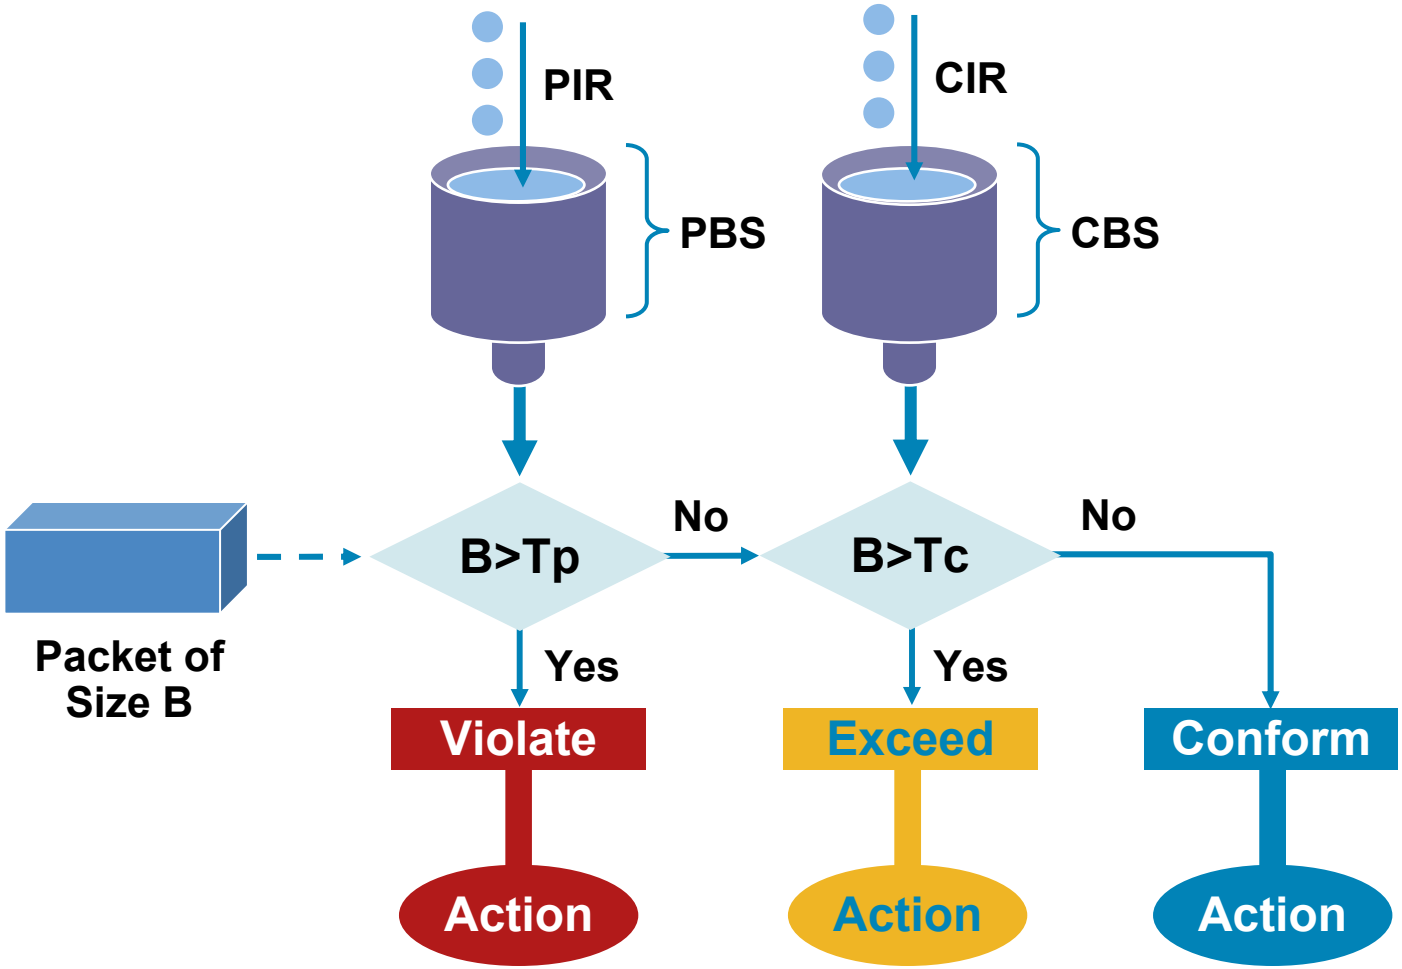

PIR=Peak Information Rate; PBS – Priority based Schedulling

Scheduling Tools
Queuing Algorithms

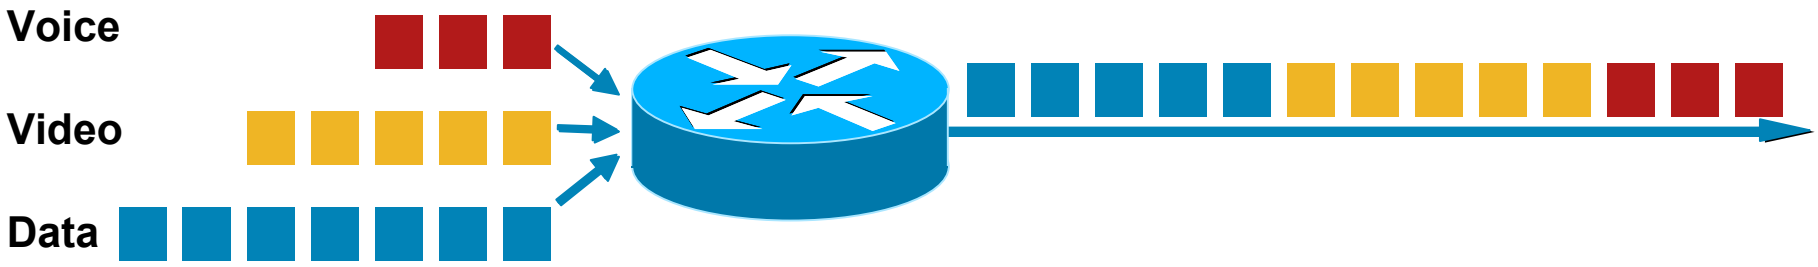

- Congestion can occur at any point in the network where there are speed mismatches
- Routers use Cisco IOS-based software queuing
  - Low-Latency Queuing (LLQ) used for highest-priority traffic (voice/video)
  - Class-Based Weighted-Fair Queuing (CBWFQ) used for guaranteeing bandwidth to data applications
- Cisco Catalyst switches use hardware queuing

Scheduling Tools

TCP Global Synchronization: The Need for Congestion Avoidance

- All TCP Flows Synchronize in Waves
- Synchronization Wastes Available Bandwidth

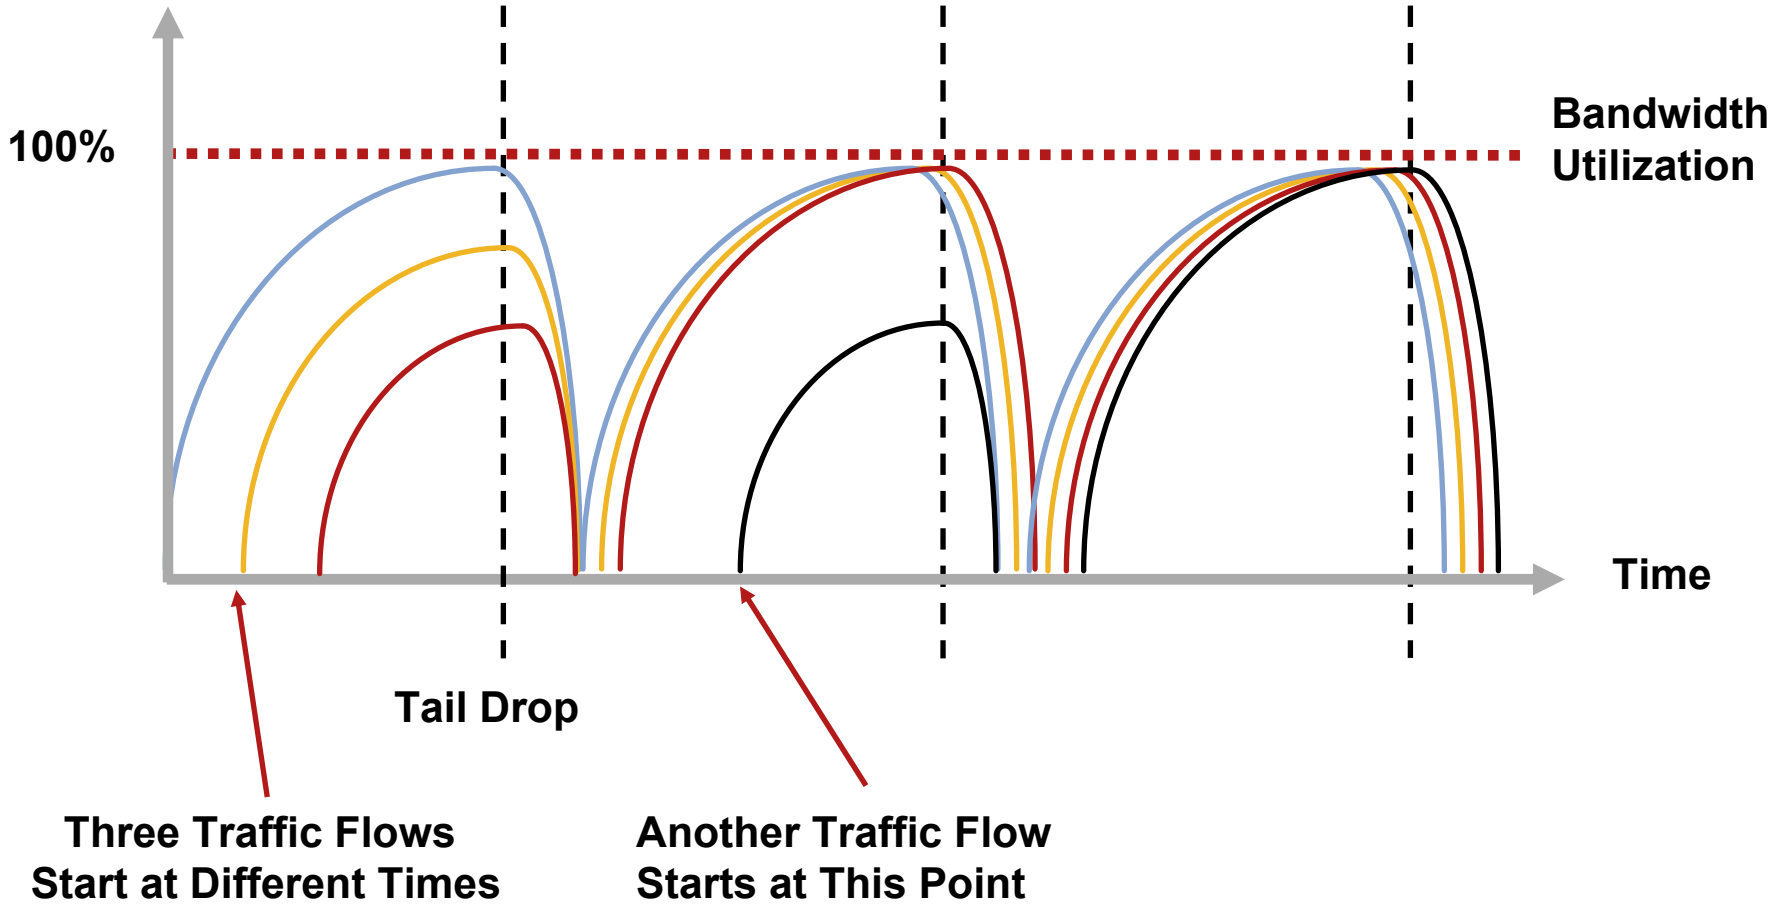

Scheduling Tools
Congestion Avoidance Algorithms

WRED

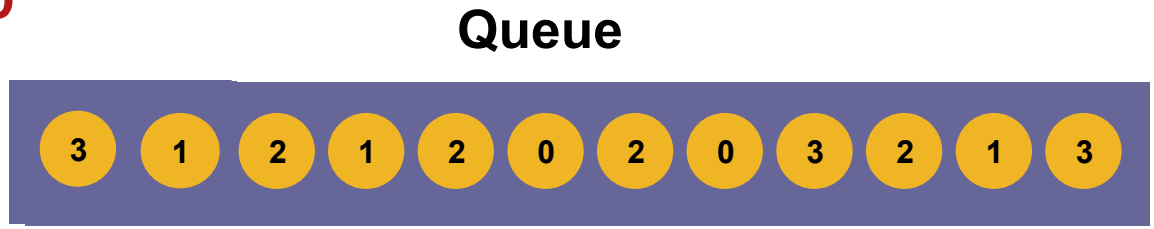

- Queueing algorithms manage the front of the queue
  - → which packets get transmitted first
- Congestion avoidance algorithms manage the tail of the queue
  - → which packets get dropped first when queuing buffers fill
- Weighted Random Early Detection (WRED)
  - WRED can operate in a DiffServ-compliant mode
  - → Drops packets according to their DSCP markings
  - WRED works best with TCP-based applications, like data

Scheduling Tools
DSCP-Based WRED Operation

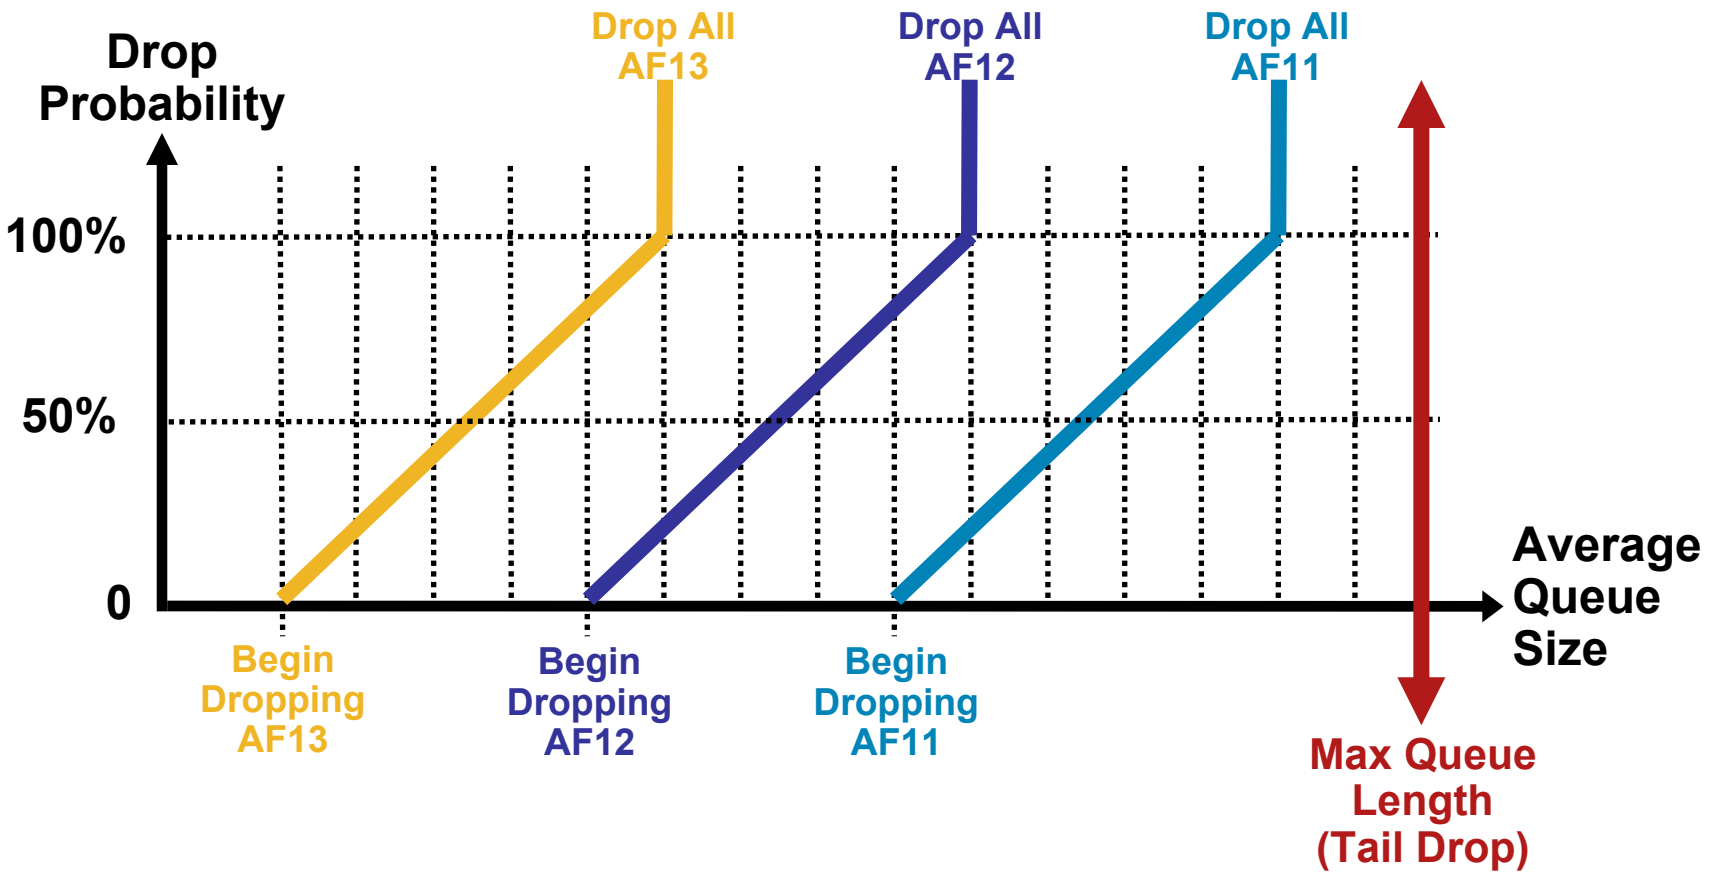

AF = (RFC 2597) Assured Forwarding

Congestion Avoidance

- IP Header Type of Service (ToS) Byte
- Explicit Congestion Notification (ECN) Bits

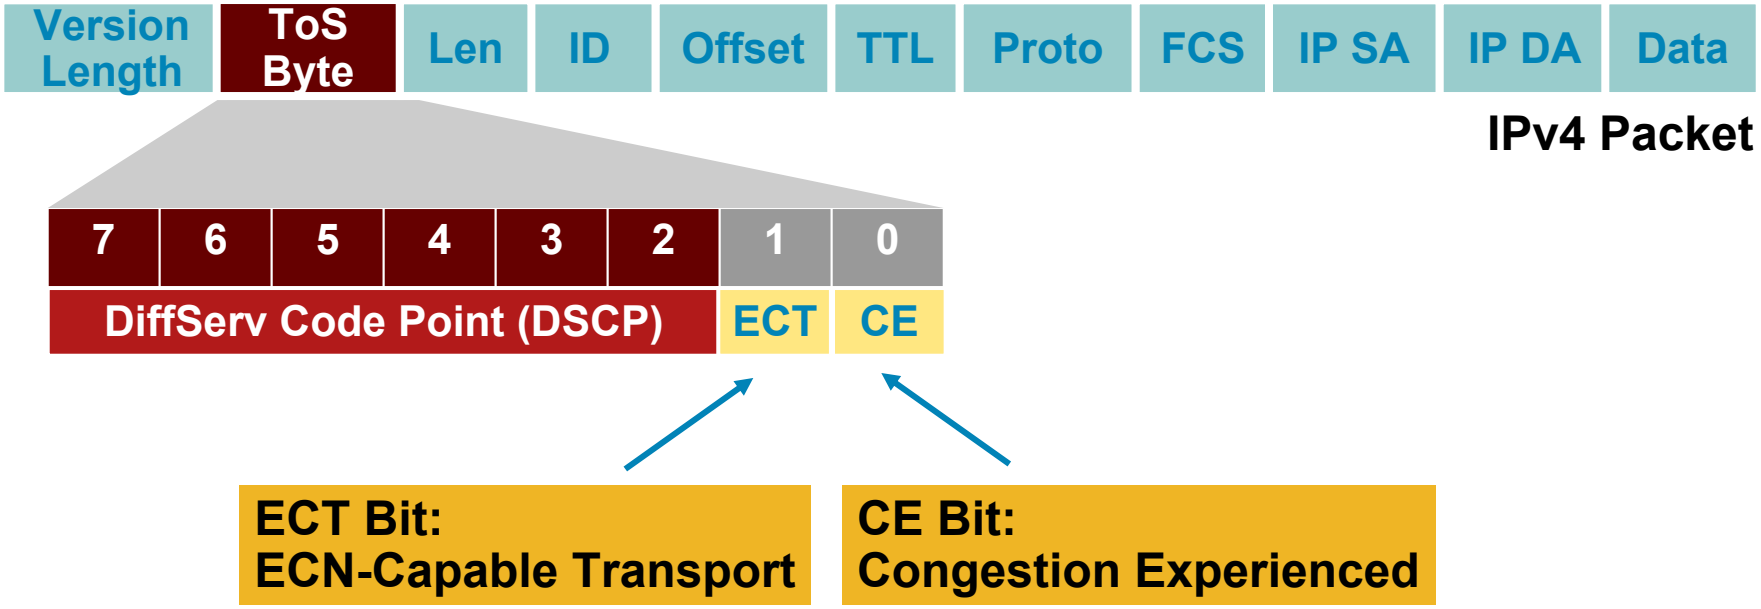

RFC3168: IP Explicit Congestion Notification

Traffic Shaping

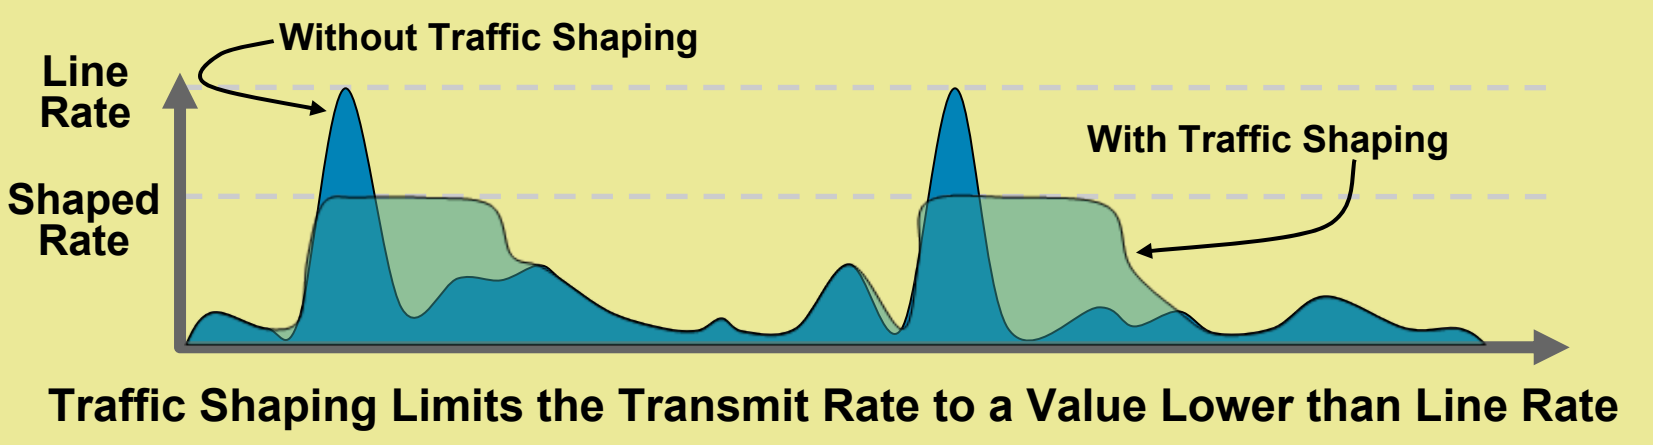

- Policers typically drop traffic
- Shapers typically delay excess traffic, smoothing bursts and preventing unnecessary drops
- Very common on Non-Broadcast Multiple-Access (NBMA) network topologies such as Frame-Relay and ATM

Shaping Tools/Link Specific Tools

Link-Specific Tools
Link-Fragmentation and Interleaving

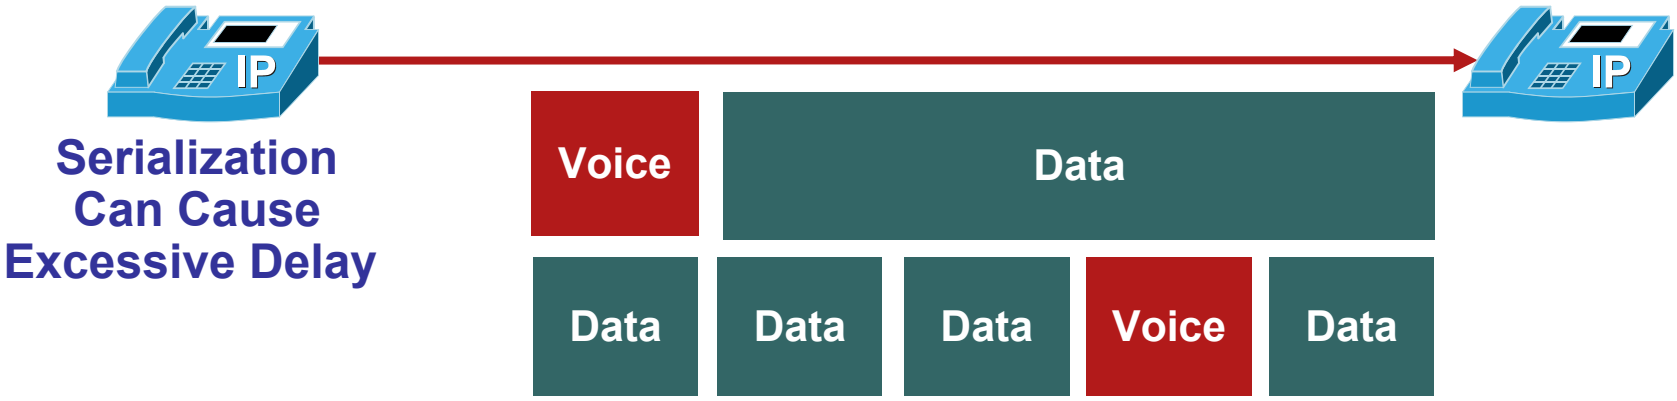

- Serialization delay is the finite amount of time required to put frames on a wire
- For links ≤ 768 kbps serialization delay is a major factor affecting latency and jitter
- For such slow links, large data packets need to be fragmented and interleaved with smaller, more urgent voice packets

Link-Specific Tools
IP RTP Header Compression

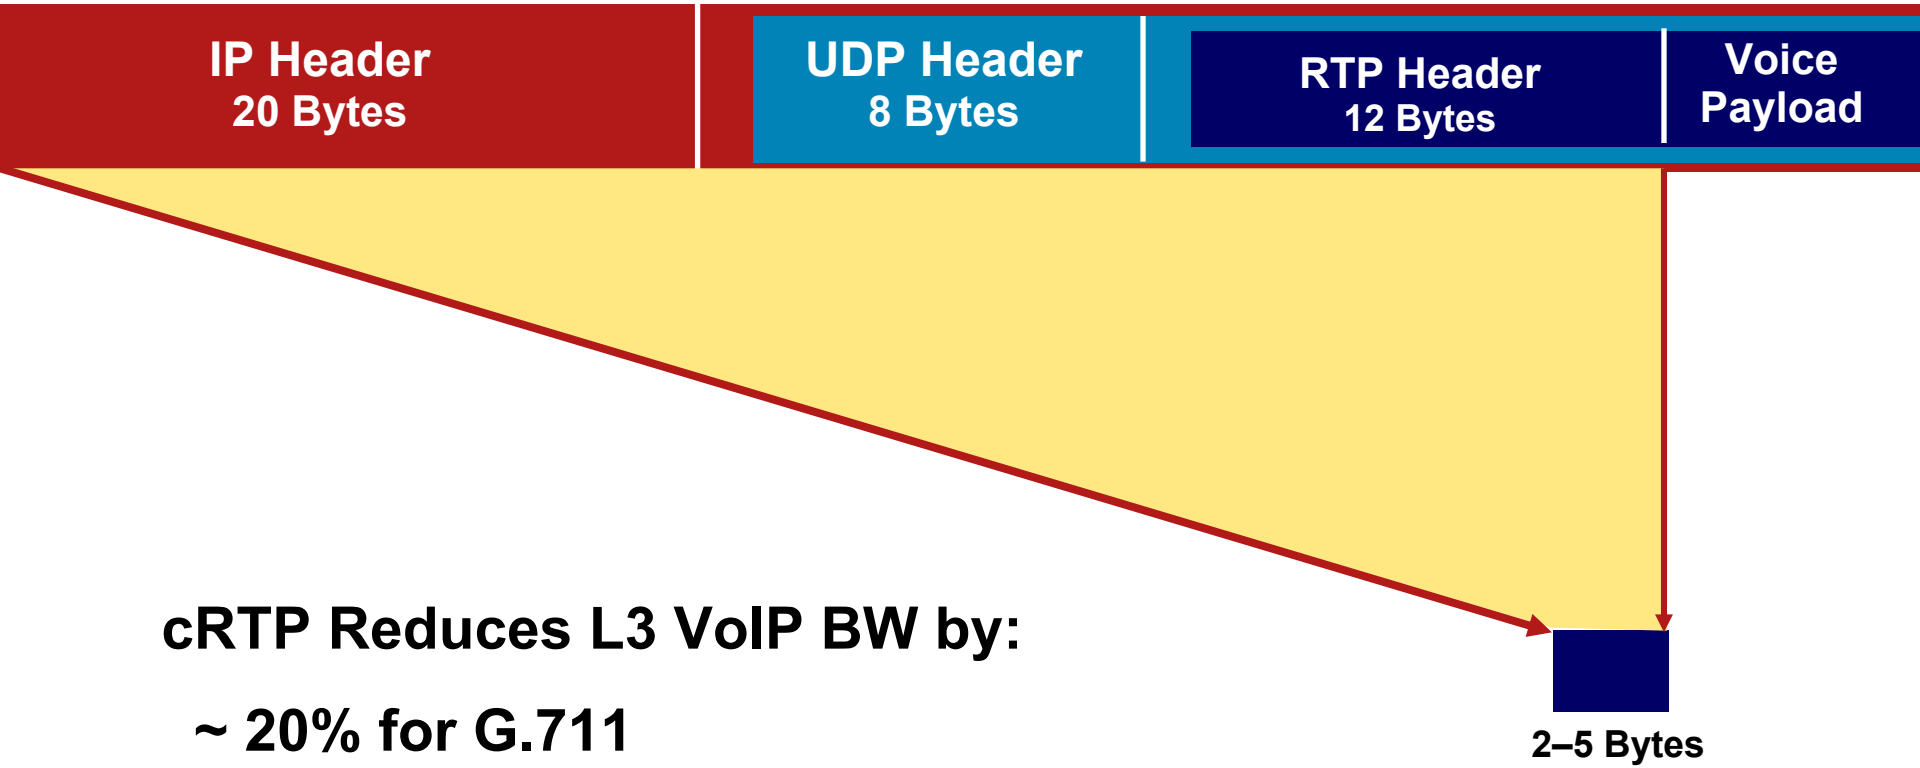

cRTP Reduces L3 VoIP BW by:

- **~ 20% for G.711**
- **~ 60% for G.729**

Signaling Tools
Resource Reservation Protocol (RSVP)

- RSVP QoS services
  - Guaranteed service
    - Mathematically provable bounds on end-to-end datagram queuing delay/bandwidth
  - Controlled service
    - Approximate QoS from an unloaded network for delay/bandwidth
- RSVP provides the policy to WFQ and LLQ

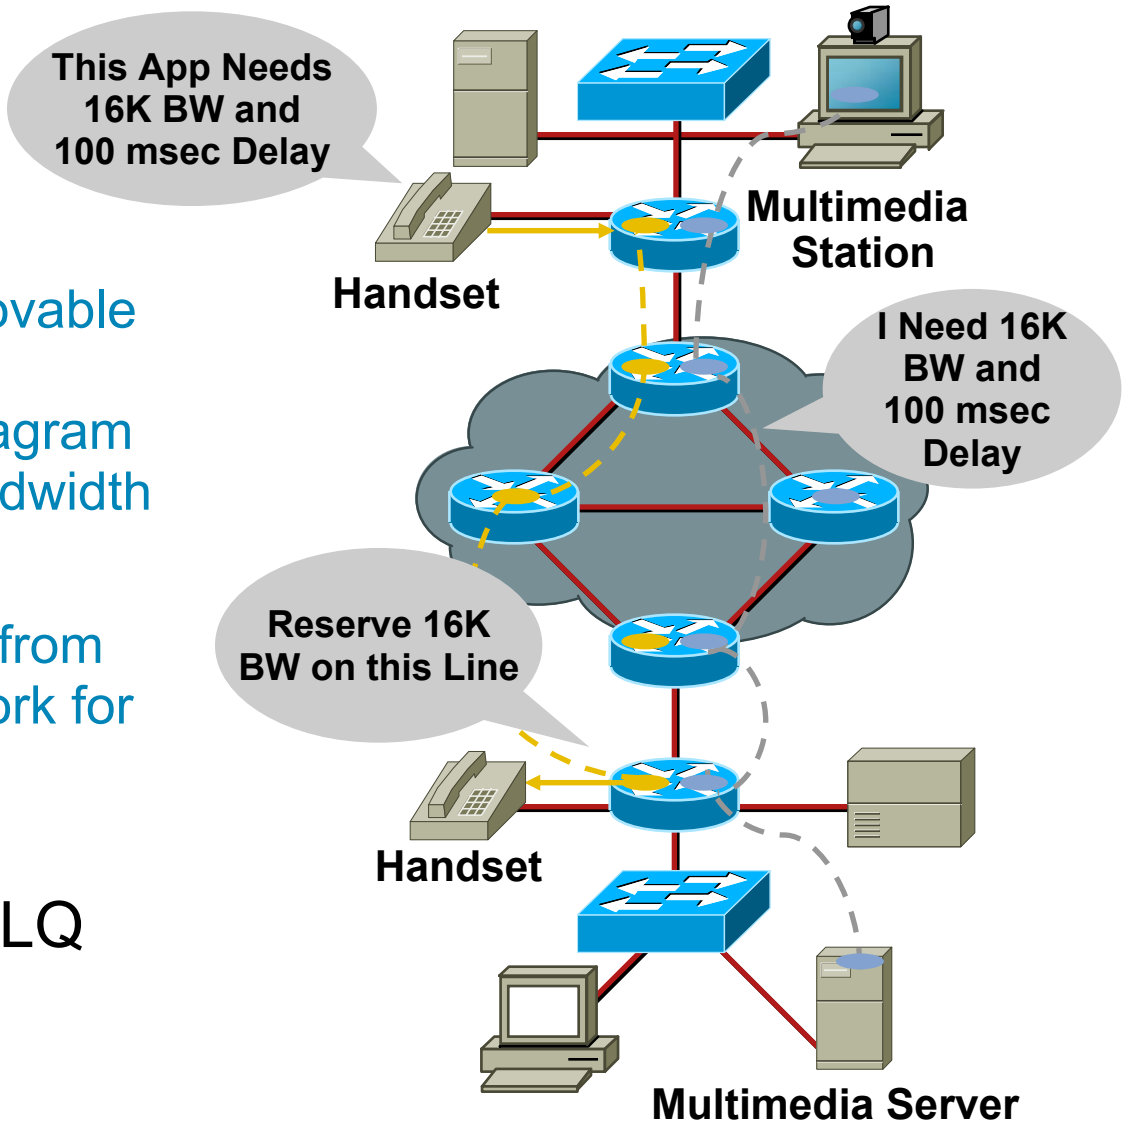

Cisco AutoQoS Phase 1 – 'Automatic QoS for VoIP Traffic' (AutoQoS - VoIP)

Configures each switch or router

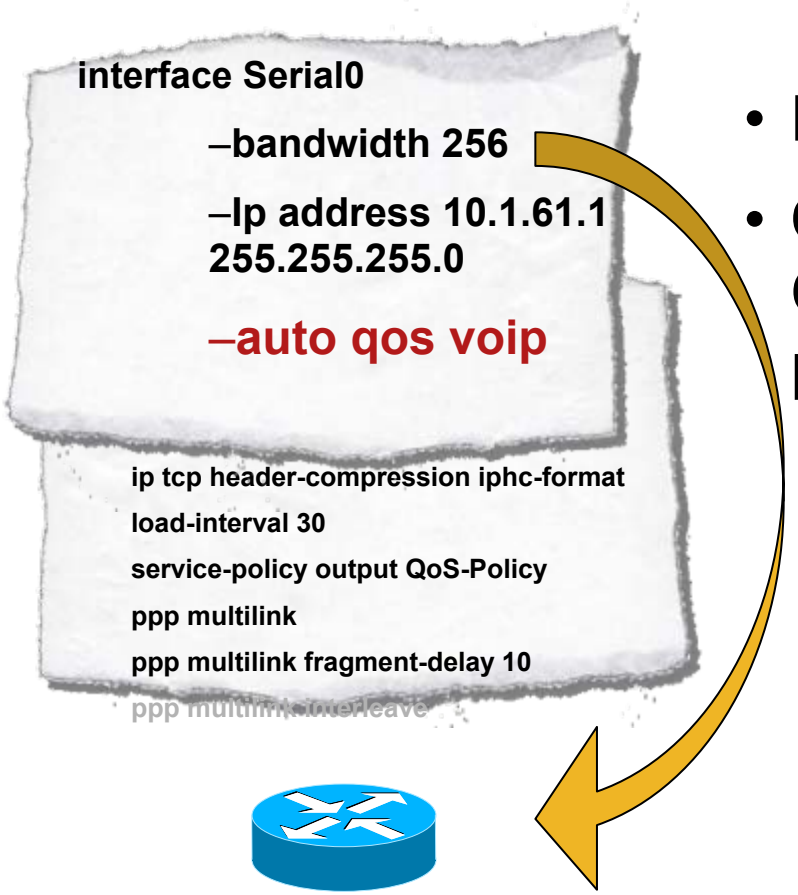

- **LAN and WAN - Routers & switches**
- **One single command enables Cisco QoS for VoIP on a port/interface/PVC!**

Cisco AutoQoS – Automating the Key Elements of QoS Deployment

1. Application classification
   - Example: automatically discovering applications and providing appropriate QoS treatment
2. Policy generation
   - Example: auto-generation of initial and ongoing QoS policies
3. Configuration
   - Example: providing high level business knobs, and multi-device / domain automation for QoS
4. Monitoring and reporting
   - Example: generating intelligent, automatic alerts and summary reports
5. Consistency
   - Example: enabling automatic, seamless interoperability among all QoS features and parameters across a network topology – LAN, MAN, and WAN

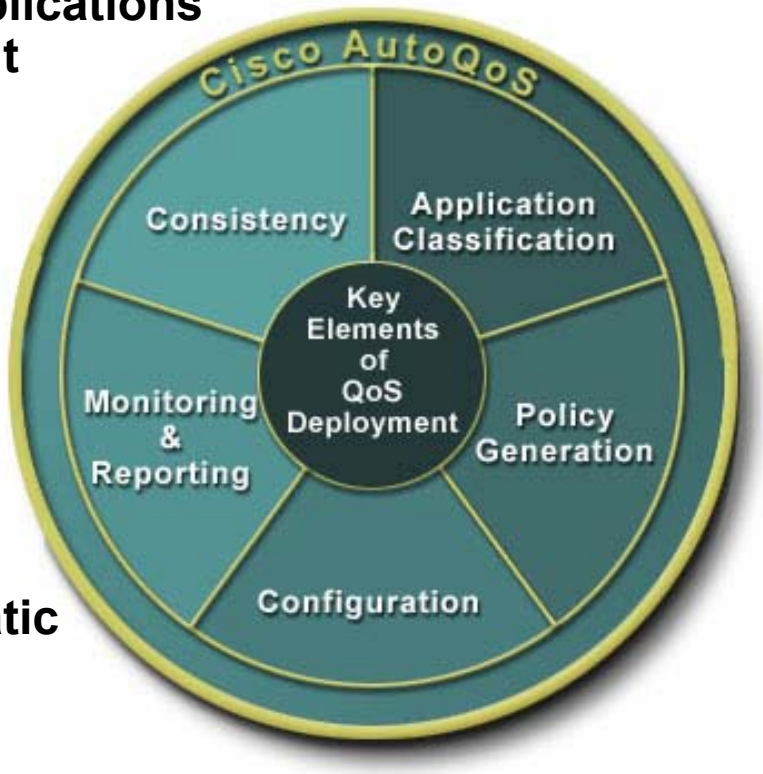

QoS for Convergence

Voice QoS Requirements
End-to-End Latency

Avoid the "Human Ethernet"

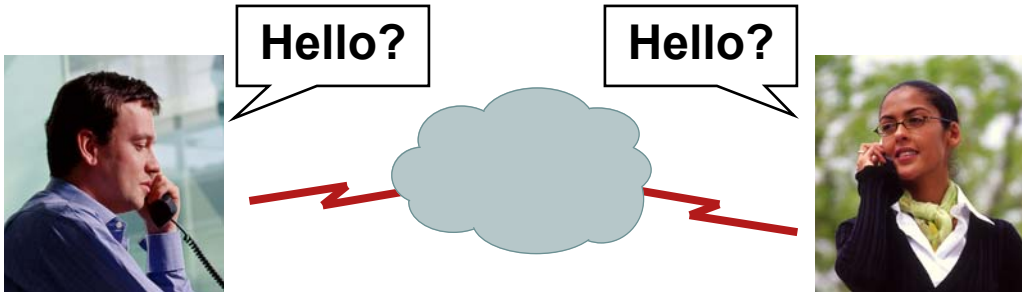

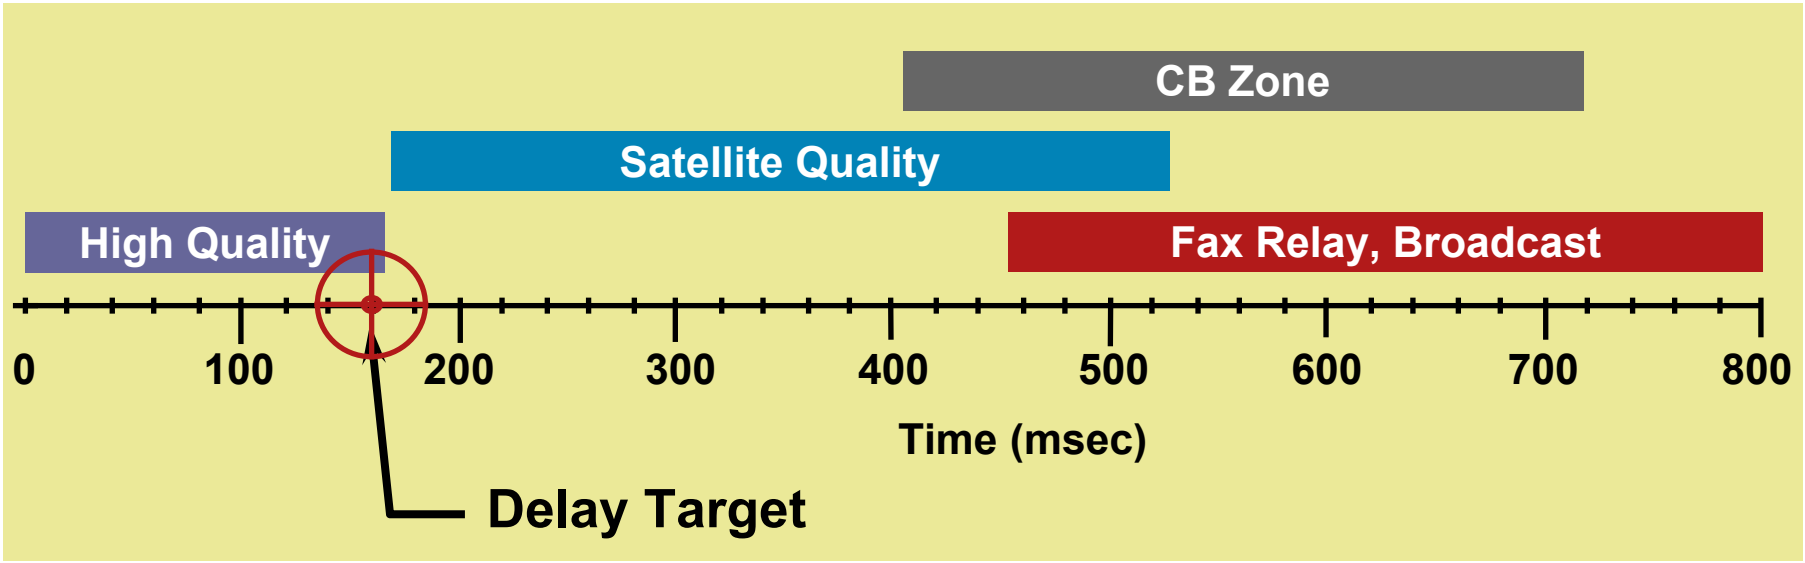

ITU's G.114 Recommendation: ≤ 150msec One-Way Delay

Voice QoS Requirements
Elements That Affect Latency and Jitter

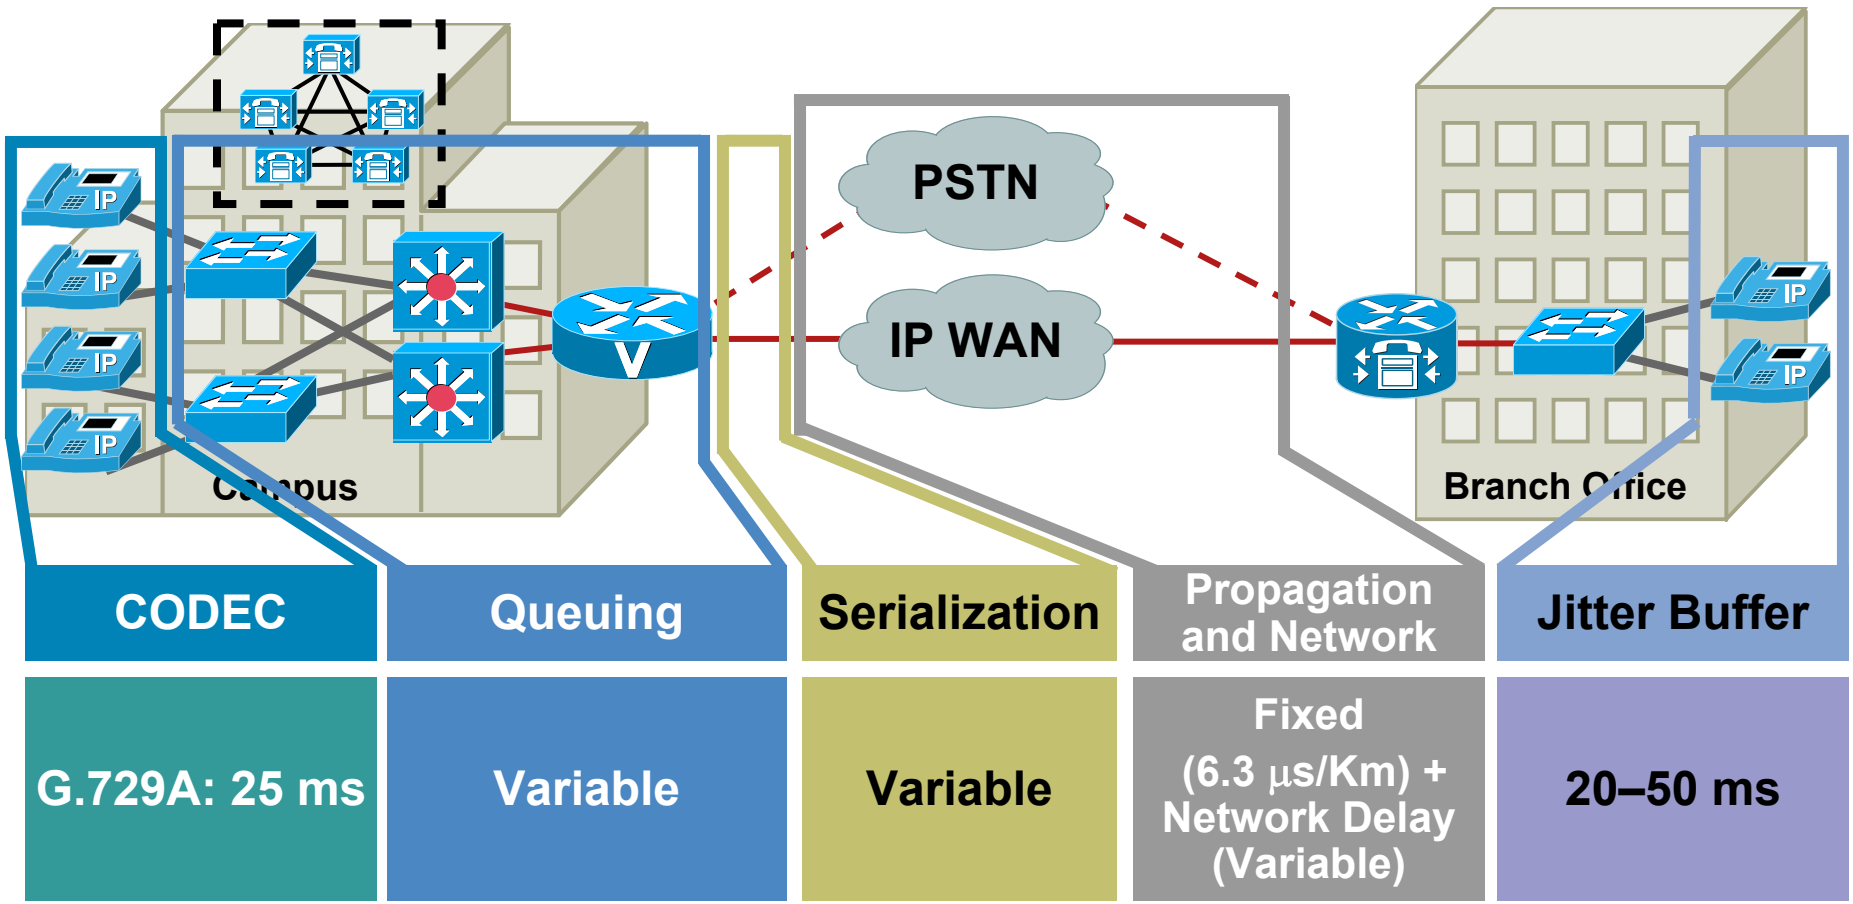

Voice QoS Requirements
Packet Loss Limitations

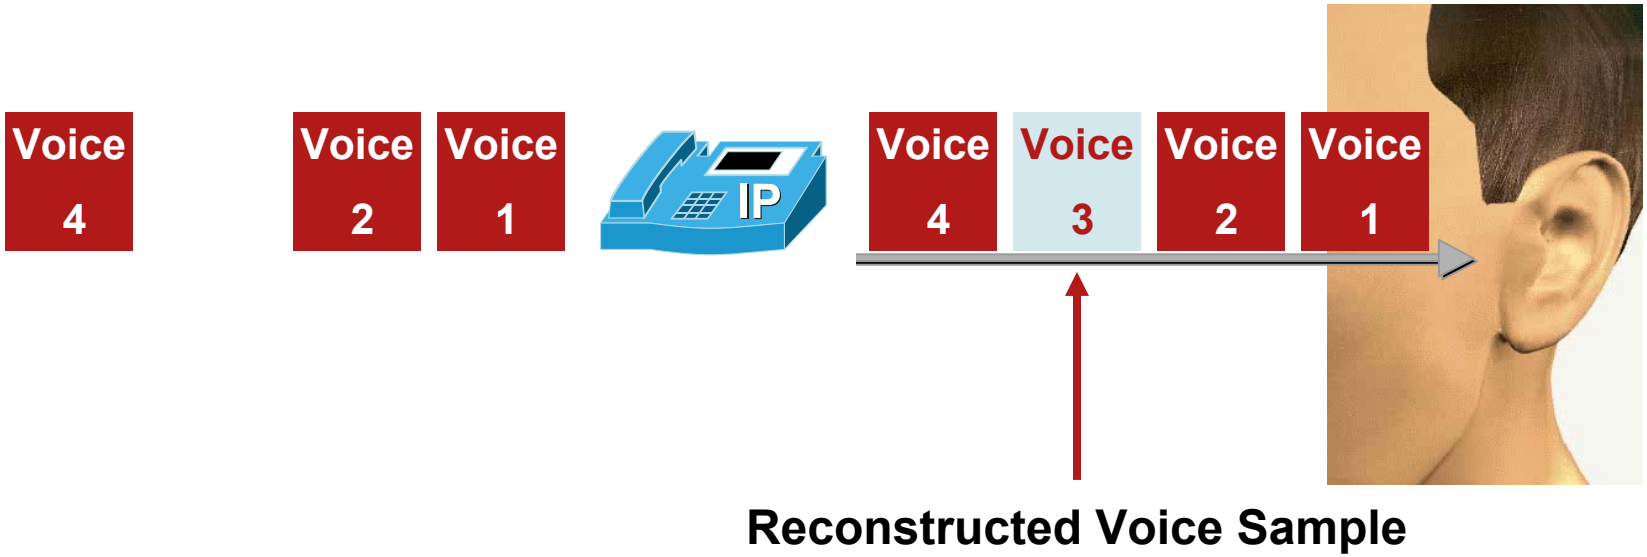

- Cisco DSP codecs can use predictor algorithms to compensate for a single lost packet in a row
- Two lost packets in a row will cause an audible clip in the conversation

Voice QoS Requirements
Provisioning for Voice

- Latency ≤ 150 ms
- Jitter ≤ 30 ms
- Loss ≤ 1%

One-Way Requirements

- 17–106 kbps guaranteed priority bandwidth per call
- 150 bps (+ Layer 2 overhead) guaranteed bandwidth for Voice-Control traffic per call
- CAC must be enabled

Voice

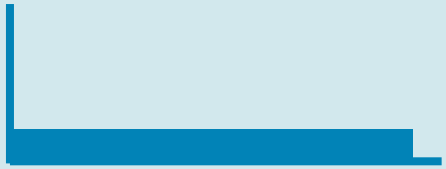

- Smooth
- Benign
- Drop sensitive
- Delay sensitive
- UDP priority

Video QoS Requirements
Video Conferencing Traffic Example (384kbps)

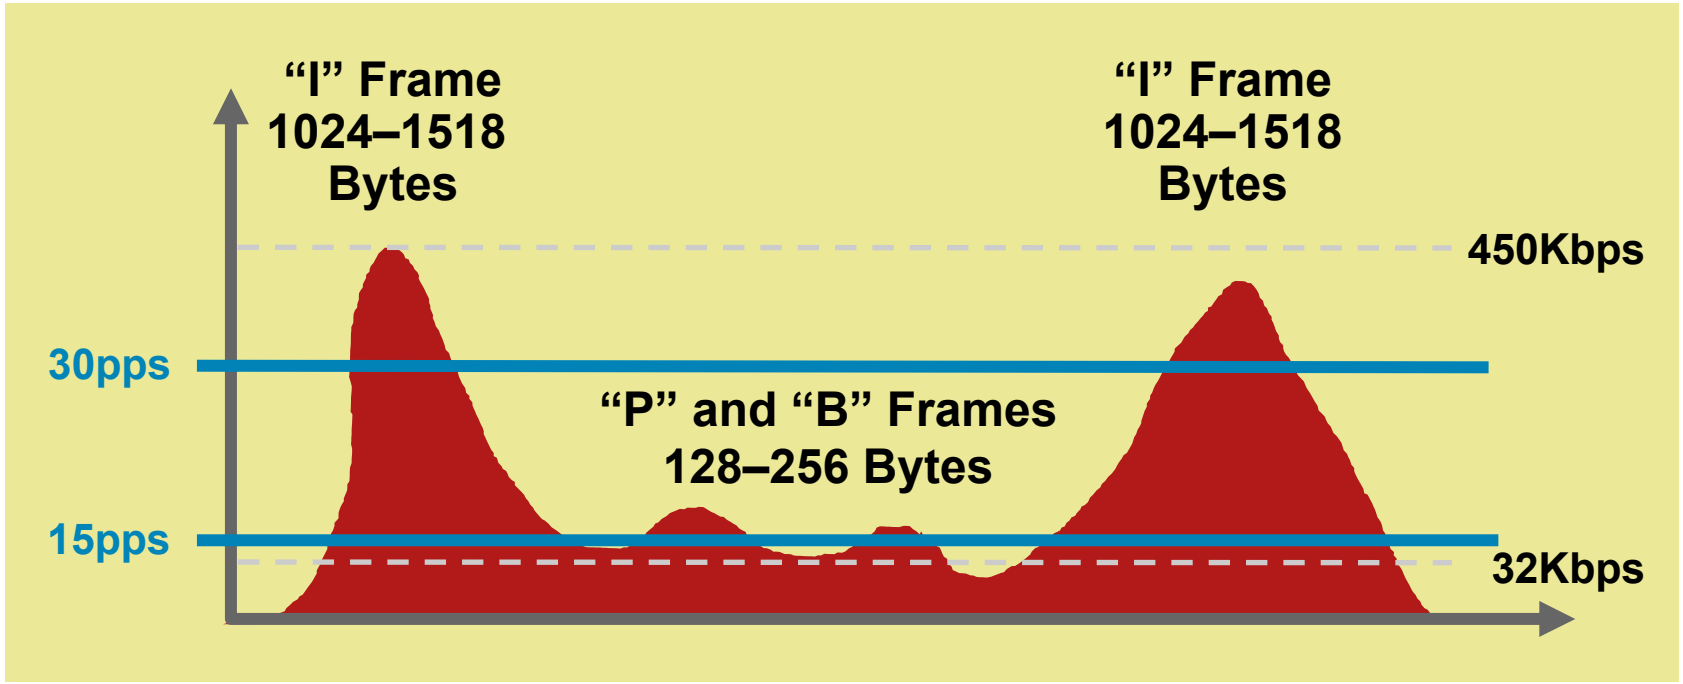

- "I" (infra) frame is a full sample of the video
- "P" (predictive) & "B" (Bi-dir)frames use quantization via motion vectors and prediction algorithms
- Key point is that dealing with large bursty I frames

Video QoS Requirements
Video Conferencing Traffic Packet Size Breakdown

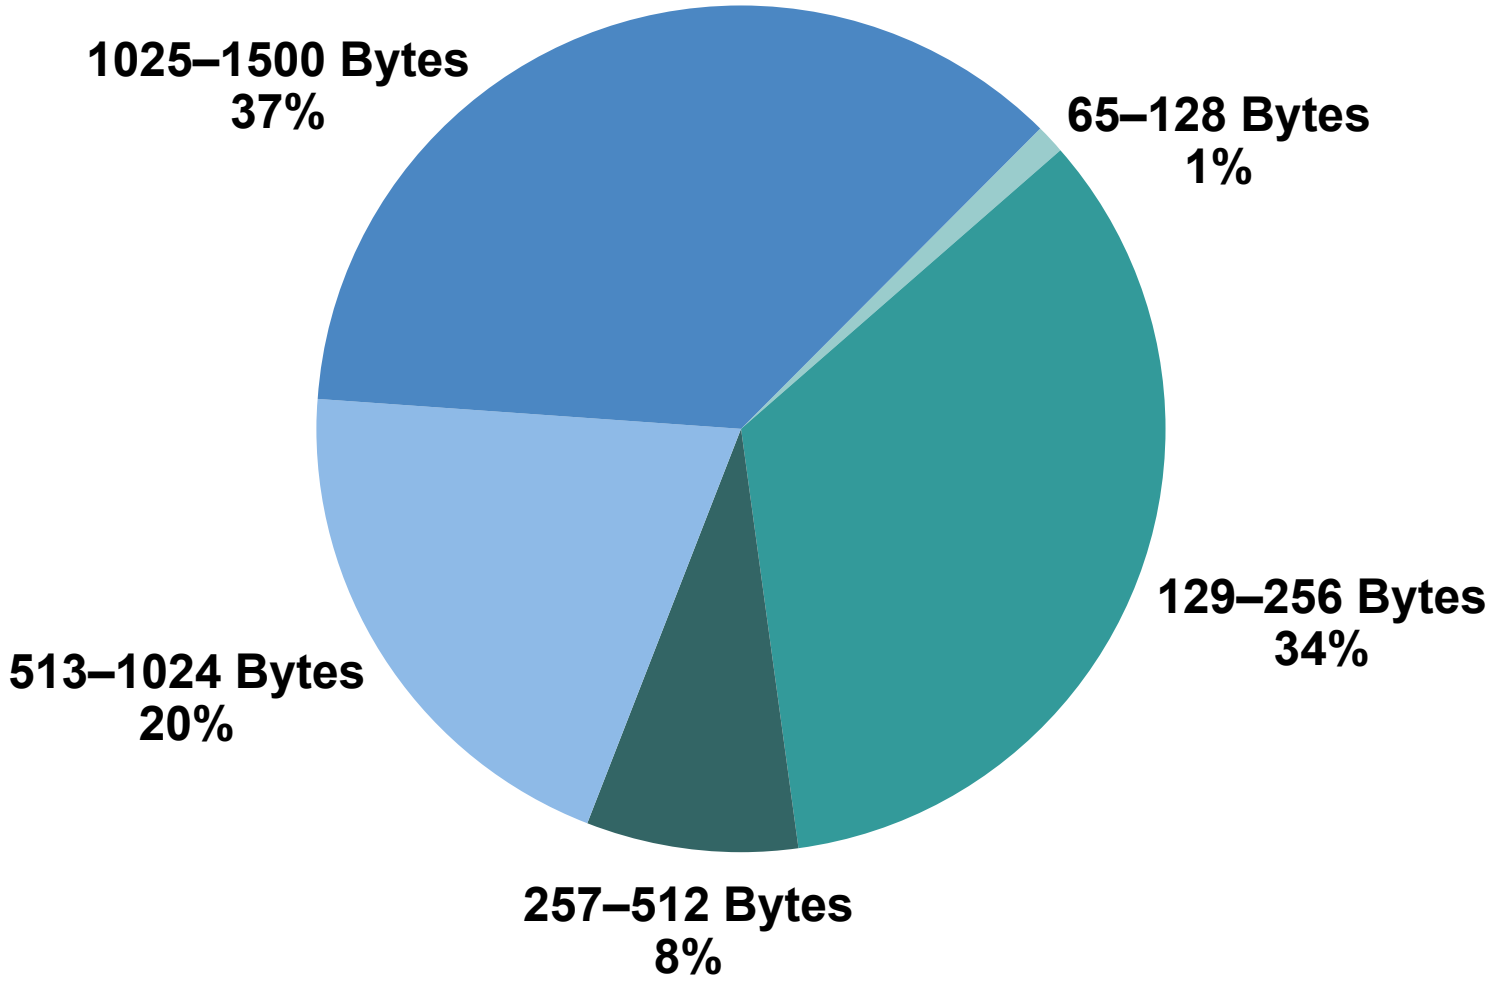

Video QoS Requirements Provisioning for Interactive Video

- Latency ≤ 150 ms
- Jitter ≤ 30 ms
- Loss ≤ 1%

One-Way Requirements

- Minimum priority bandwidth guarantee required is:

  Video-stream + 10–20%

  e.g., a 384 kbps stream could require up to 460 kbps of priority bandwidth

- CAC must be enabled

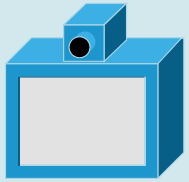

Video

- Bursty
- Drop sensitive
- Delay sensitive
- UDP priority

Data QoS Requirements Application Differences

Oracle

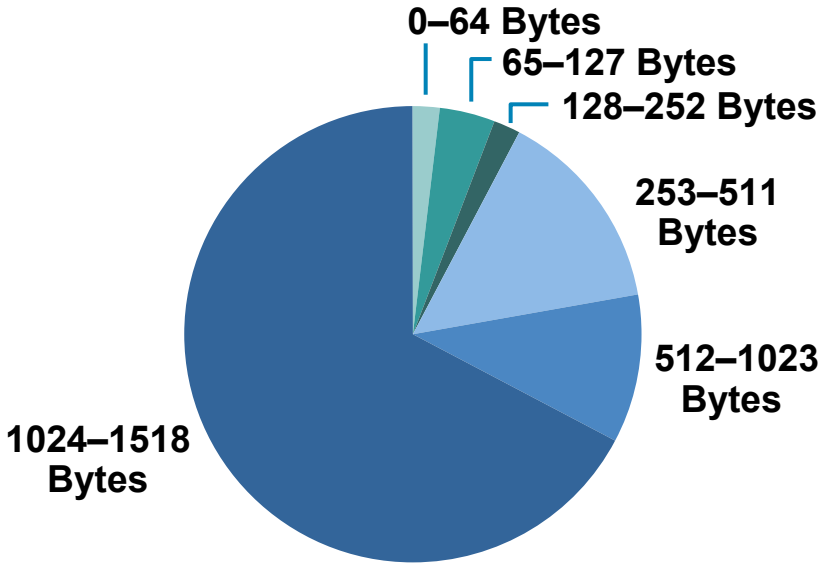

SAP R/3

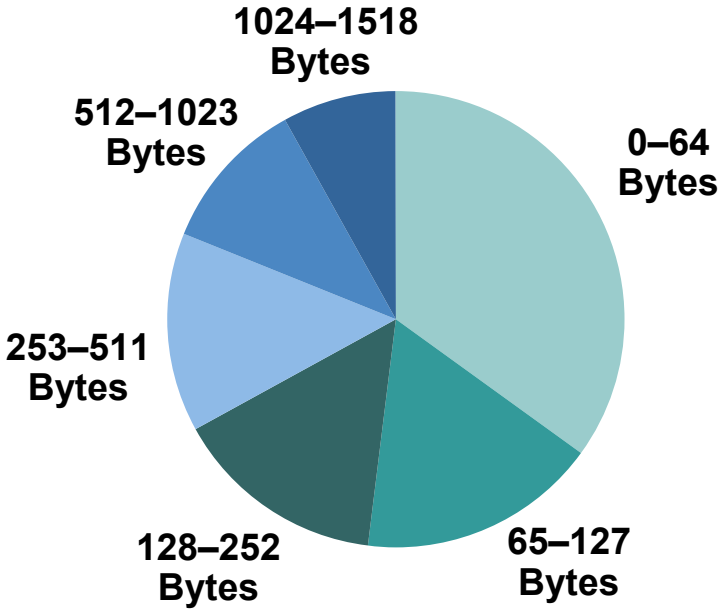

Data QoS Requirements Version Differences

Same Transaction Takes Over 35 Times More Traffic from One Version of an Application to Another

SAP Sales Order Entry Transaction

Client Version	VA01 # of Bytes
SAP GUI Release 3.0 F	14,000
SAP GUI Release 4.6C, No Cache	57,000
SAP GUI Release 4.6C, with Cache	33,000
SAP GUI for HTML, Release 4.6C	490,000

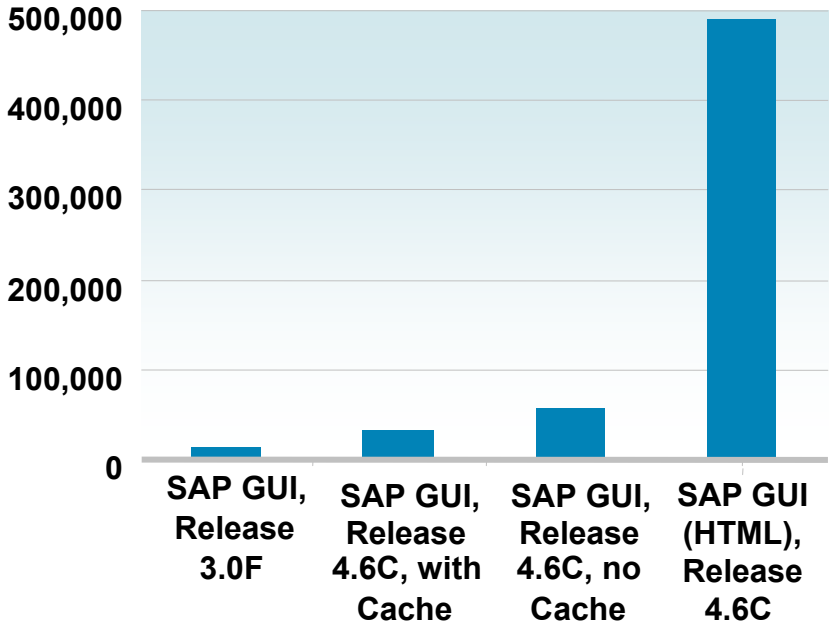

Data QoS Requirements Provisioning for Data (Cont.)

- Use four/five main traffic classes:

  Mission-critical apps—business-critical client-server applications

  Transactional/interactive apps—foreground apps: client-server apps or interactive applications

  Bulk data apps—background apps: FTP, e-mail, backups, content distribution

  Best effort apps—(default class)

  Optional: Scavenger apps—peer-to-peer apps, gaming traffic

- Additional optional data classes include internetwork-control (routing) and network-management
- Most apps fall under best-effort, make sure that adequate bandwidth is provisioned for this default class

Data QoS Requirements Provisioning for Data

- Different applications have different traffic characteristics
- Different versions of the same application can have different traffic characteristics
- Classify data into four/five data classes model:
  - Mission-critical apps
  - Transactional/interactive apps
  - Bulk data apps
  - Best effort apps
  - Optional: Scavenger apps

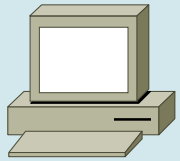

Data

- Smooth/bursty
- Benign/greedy
- Drop insensitive
- Delay insensitive
- TCP retransmits

Scavenger-Class
What Is the Scavenger Class?

- The Scavenger class is an Internet 2 Draft Specification for a "less than best effort" service
- There is an implied "good faith" commitment for the "best effort" traffic class
  - It is generally assumed that at least some network resources will be available for the default class
- Scavenger class markings can be used to distinguish out-of-profile/abnormal traffic flows from in-profile/ normal flows
  - The Scavenger class marking is CS1, DSCP 8
- Scavenger traffic is assigned a "less-than-best effort" queuing treatment whenever congestion occurs

QoS for Security

Business Security Threat Evolution
Expanding Scope of Theft and Disruption

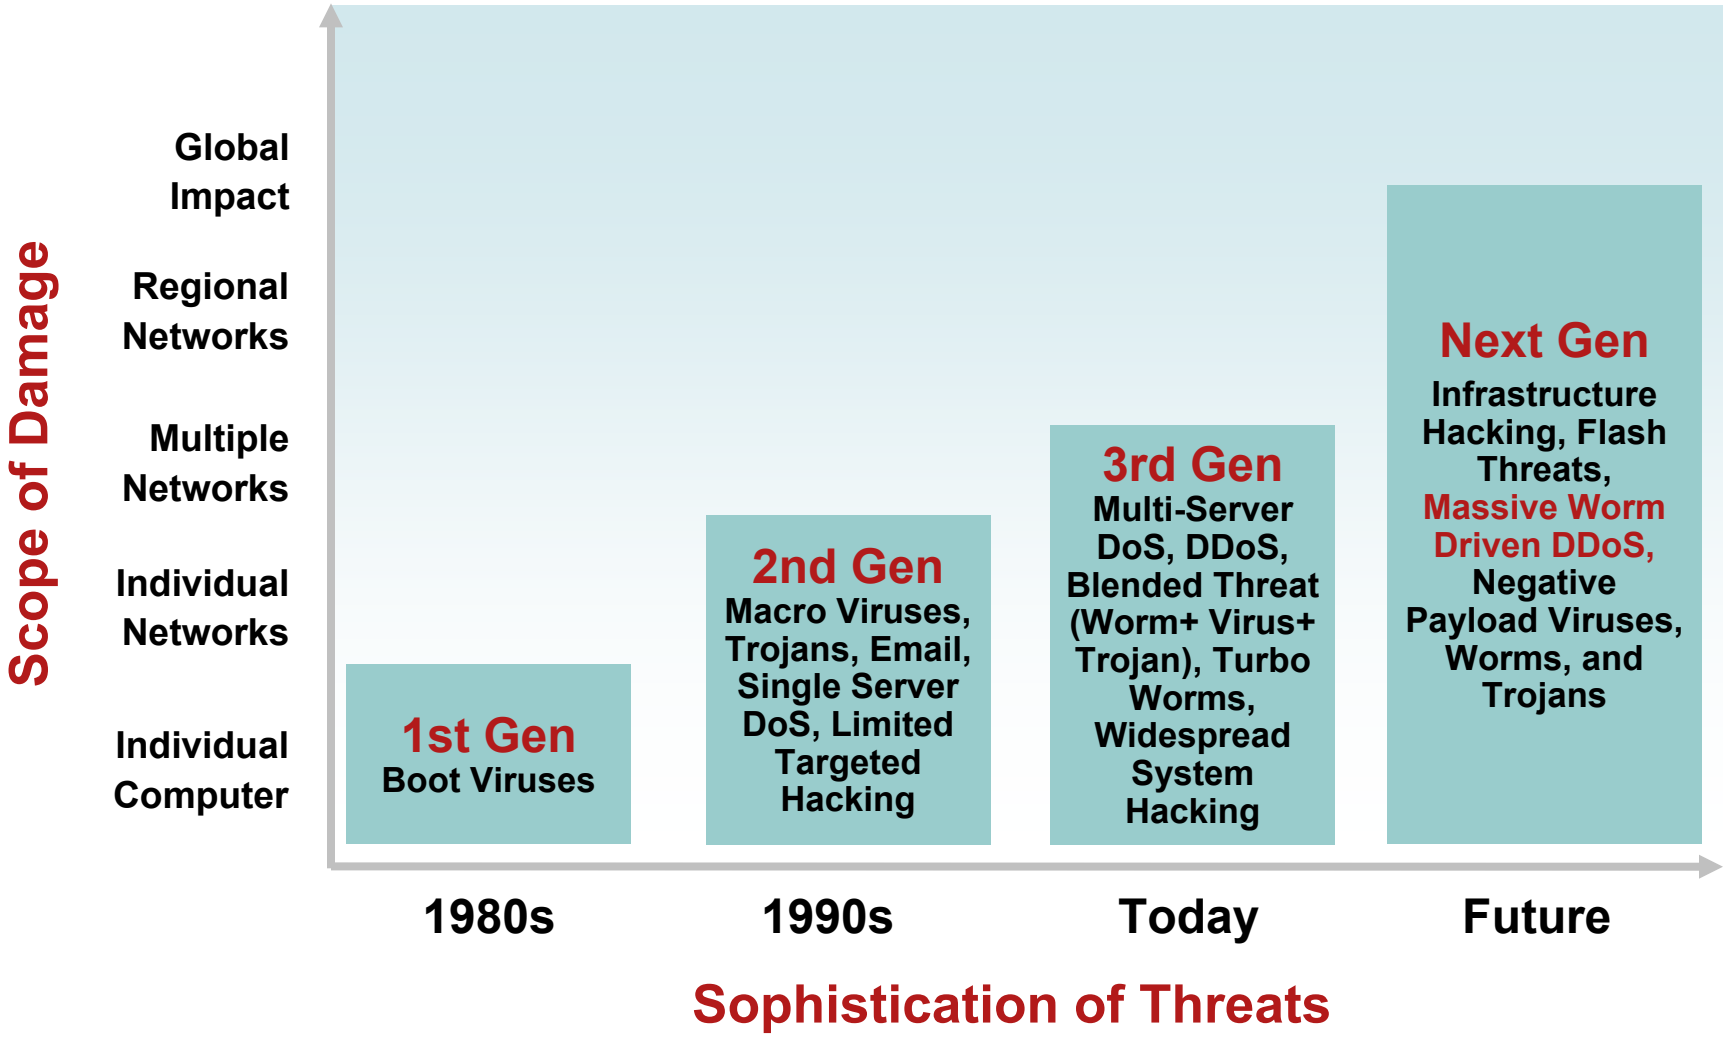

Impact of an Internet Worm
Anatomy of a Worm: Why It Hurts

1—The Enabling Vulnerability

2—Propagation Mechanism

3—Payload

Impact of an Internet Worm
Direct and Collateral Damage

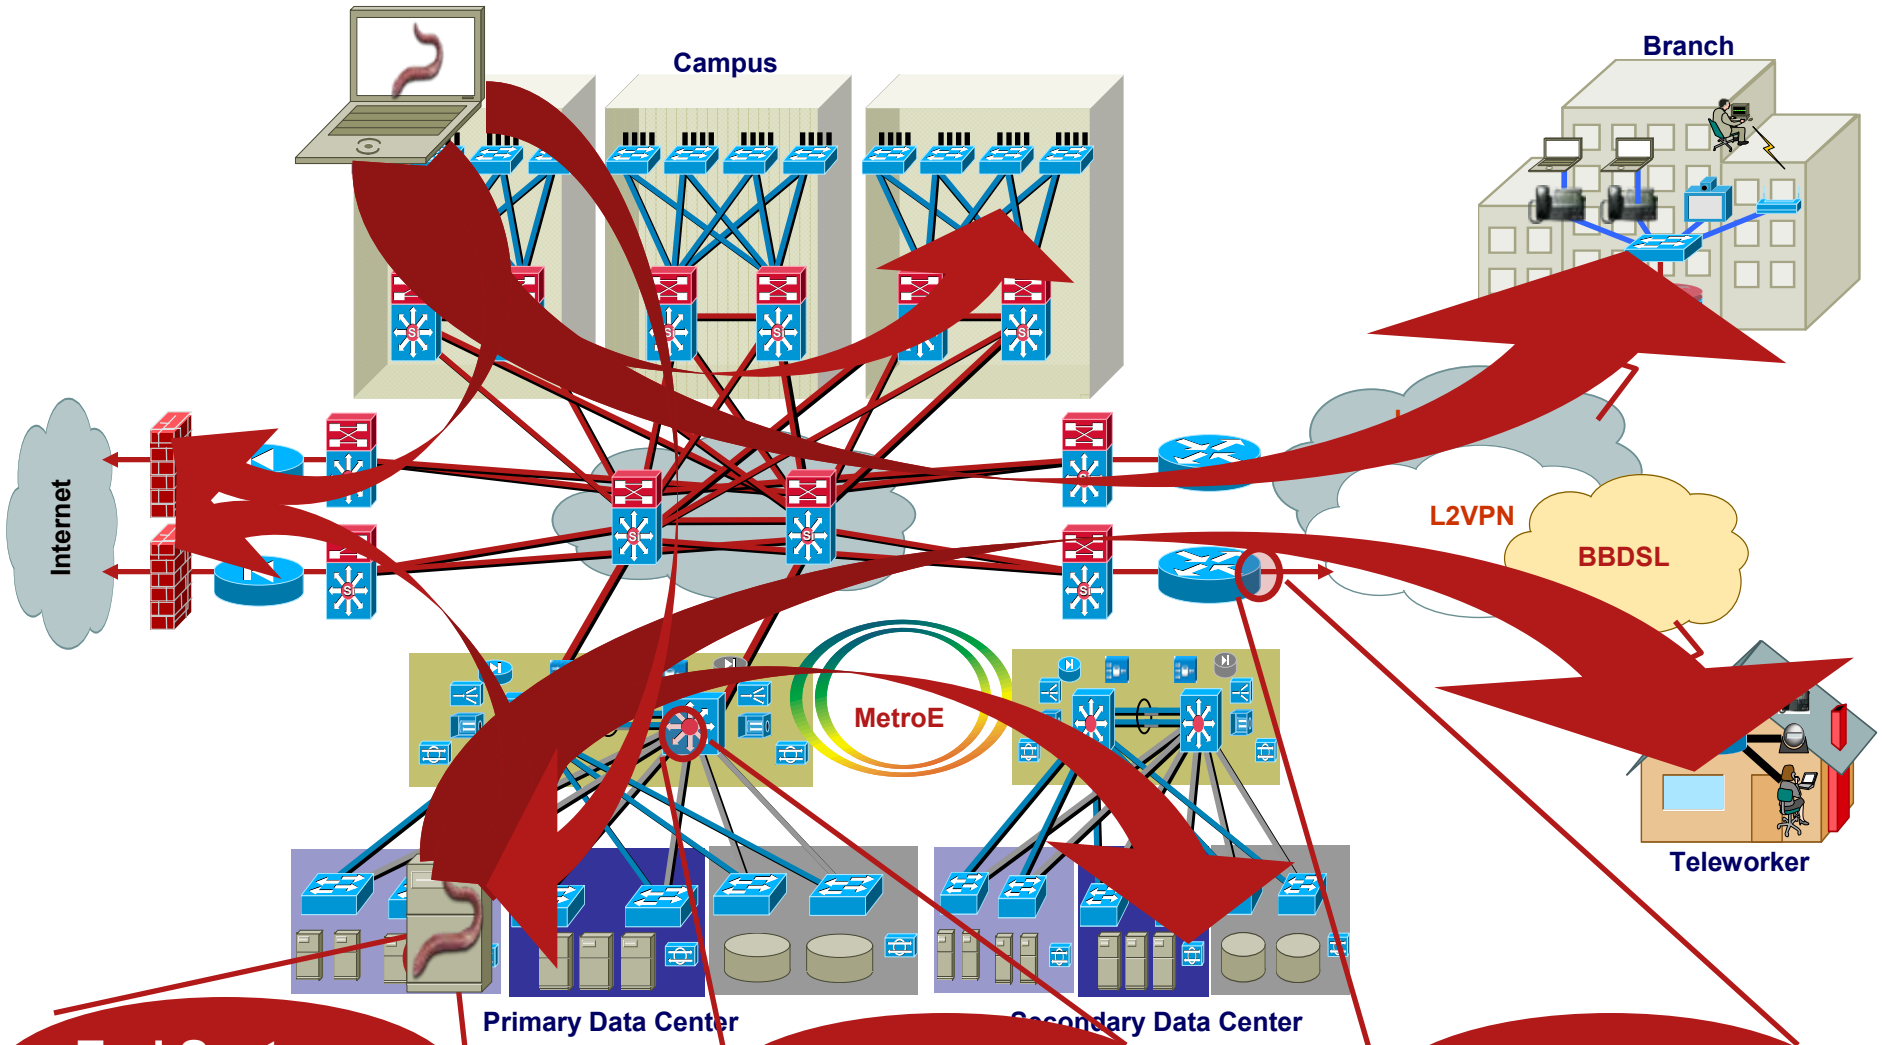

QoS Tools and Tactics for Security
QoS for Self-Defending Networks

- Control plane policing
- Data plane policing (Scavenger-Class QoS)
- NBAR for known-worm policing

Control Plane Policing Overview

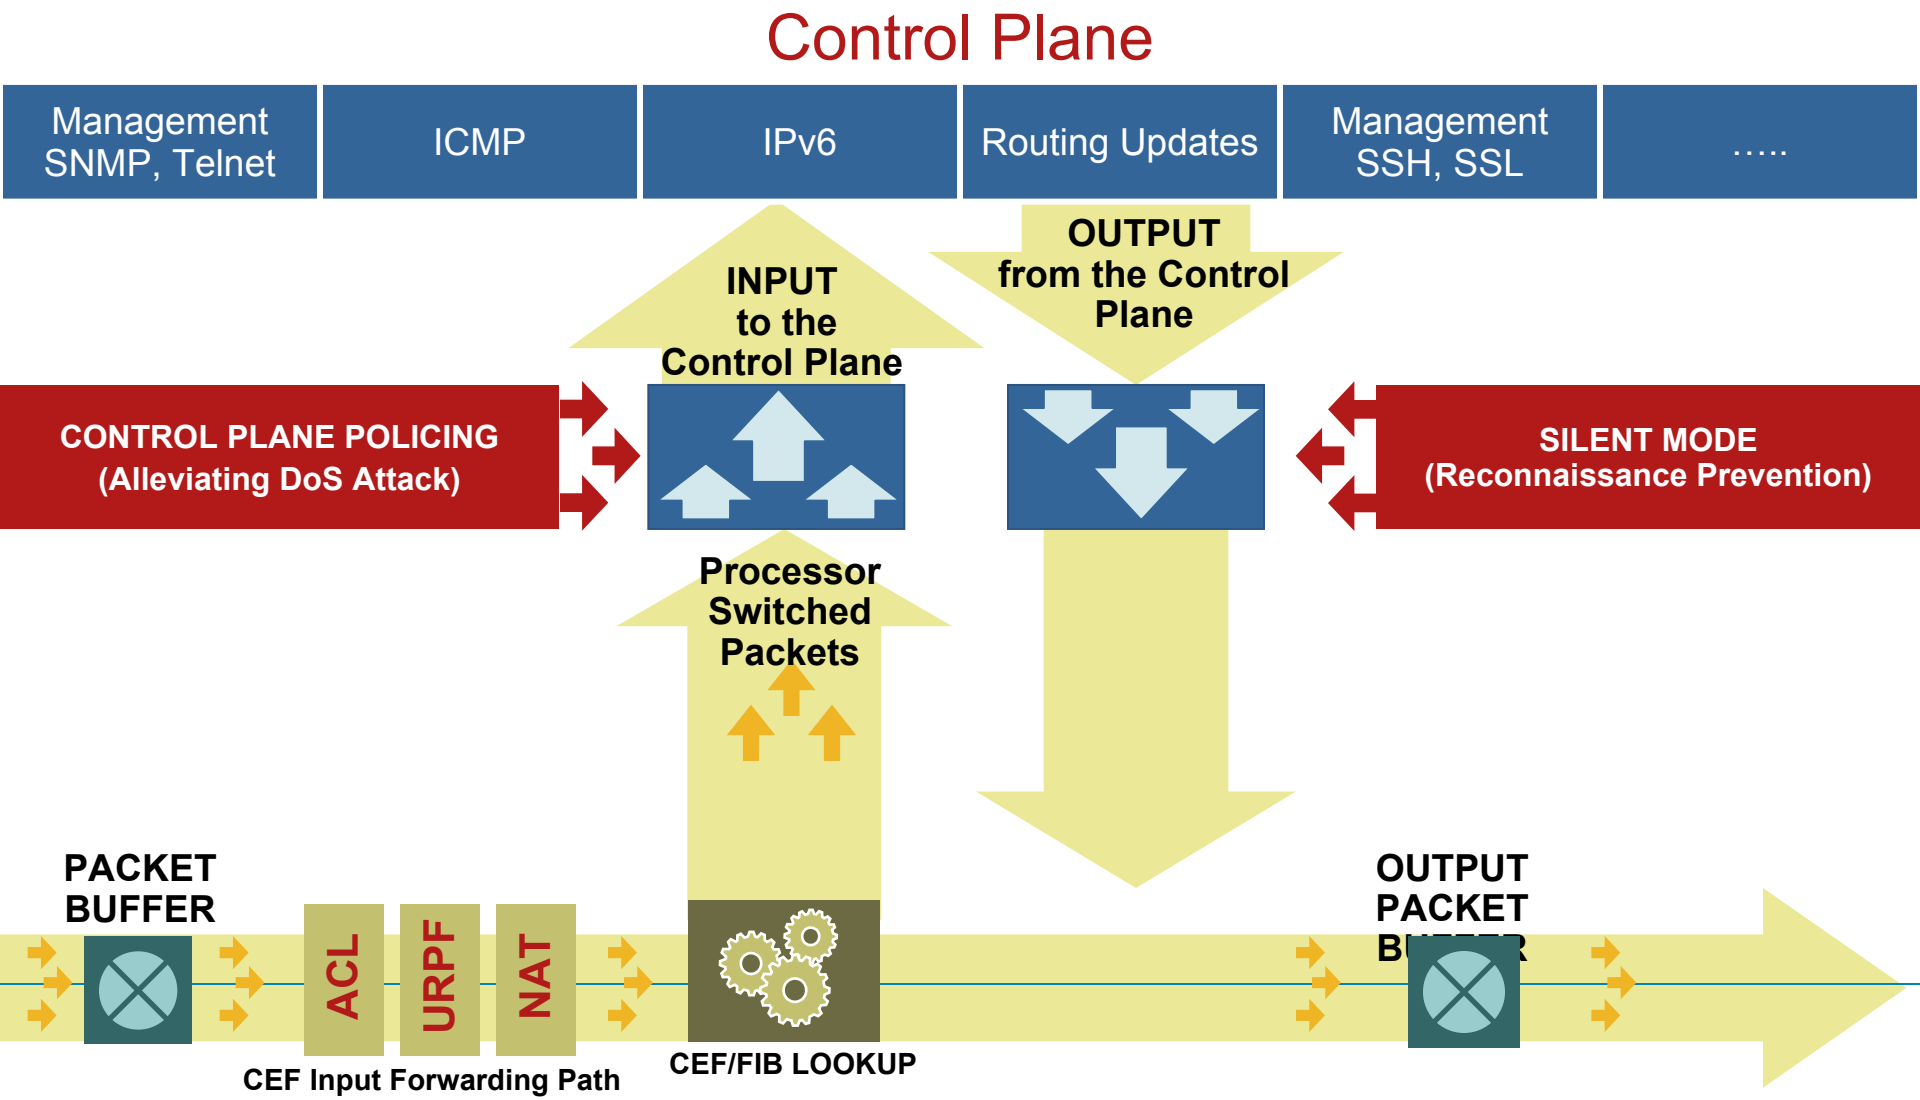

Data Plane Policing (Scavenger-Class QoS)
Part 1: First Order Anomaly Detection

- All end systems generate traffic spikes, but worms create sustained spikes
- Normal/abnormal threshold set at approx 95% confidence
- No dropping at campus access-edge! Only remarking

Policing and Remarking (if necessary)

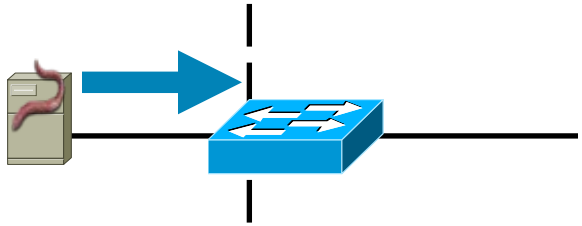

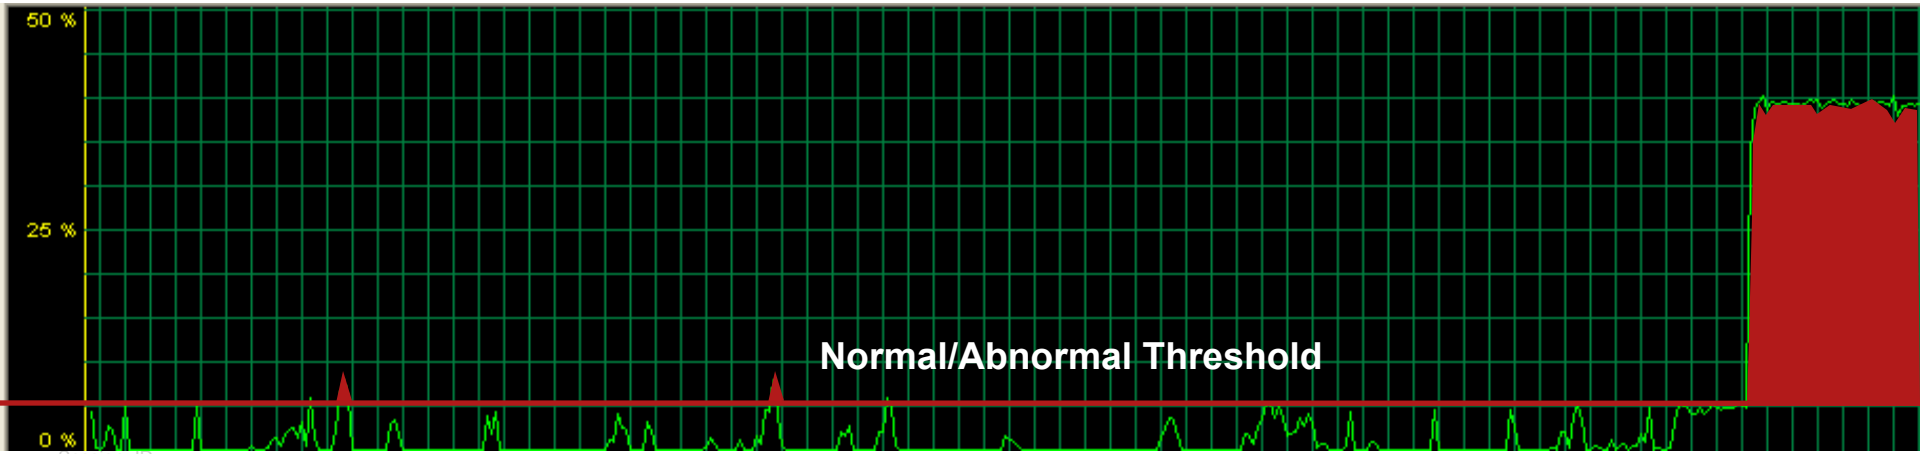

Data Plane Policing (Scavenger-Class QoS)
Part 2: Second Order Anomaly Reaction

- Queuing only engages if links become congested
  - When congestion occurs, drops will also occur
- Scavenger-class QoS allows for increased intelligence in the dropping decision
  - "Abnormal" traffic flows will be dropped aggressively
  - "Normal" traffic flows will continue to receive network service

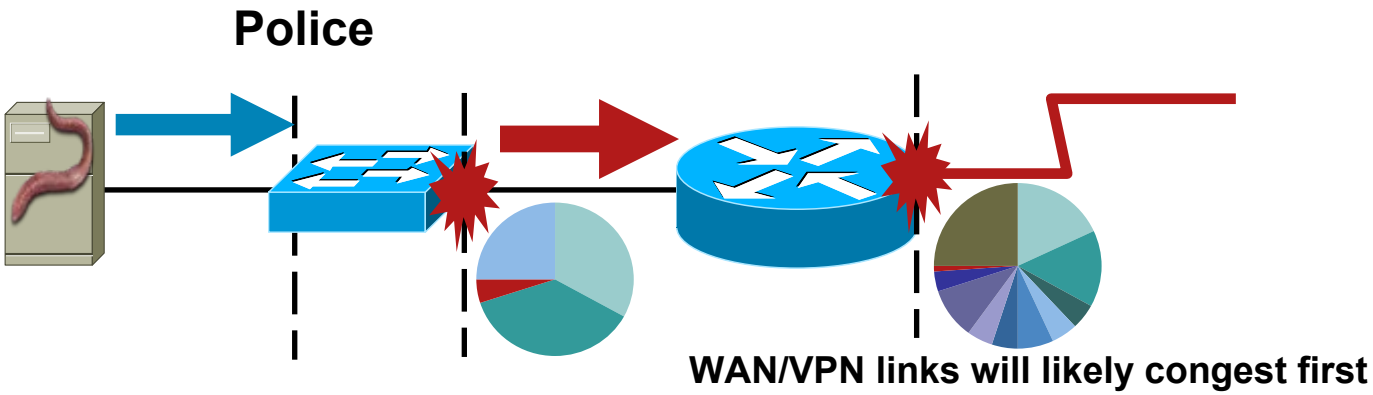

Queuing Will Engage When Links Become Congested and Traffic Previously Marked as Scavenger Is Dropped Aggressively

NBAR Known-Worm Policing
NBAR vs. Code Red Example

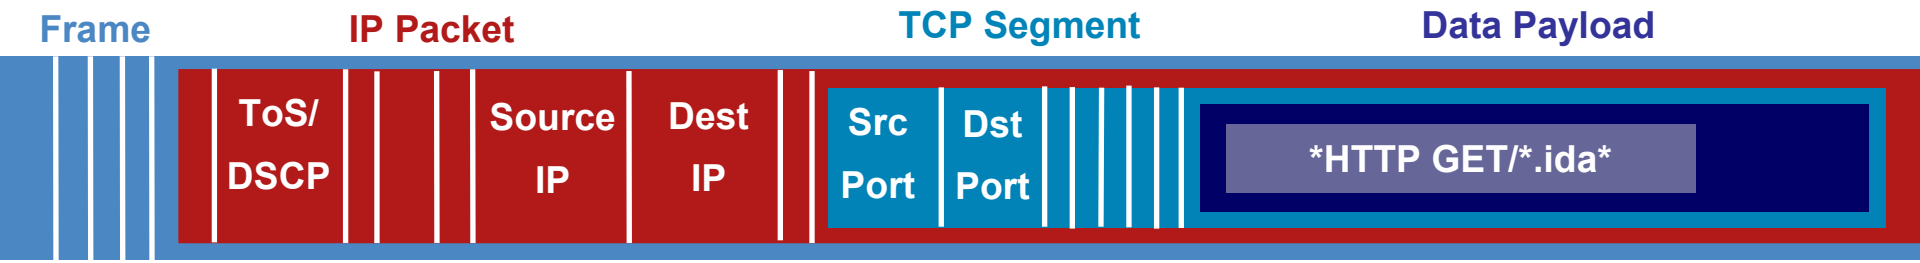

- First released in May 2001
- Exploited a vulnerability in Microsoft IIS and infected 360,000 hosts in 14 hours
- Several strains (CodeRed, CodeRedv2, CodeRed II, Code,Redv3, CodeRed.C.)
- Newer strains replaced home page of Web servers and caused DoS flooding-attacks
- Attempts to access a file with ".ida" extension

```
class-map match-any CODE-RED
   match protocol http url "*.ida*"
   match protocol http url "*cmd.exe*"
   match protocol http url "*root.exe*"
```

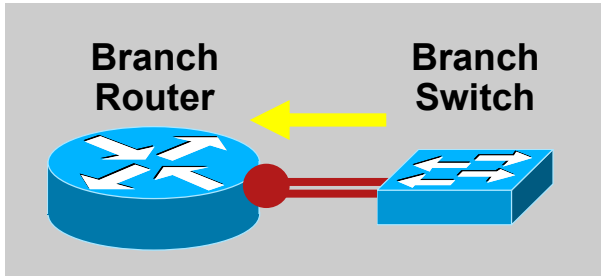

Network Management Tools

SNMP MIB
Cisco-Class-Based-QoS-MIB

- Primary accounting mechanism for QoS:
  - Policing, Classification, Shaping, Queuing, Congestion avoidance
- Long-term QoS monitoring
  - Cisco QoS Policy Manager
- Only accounts on configured QoS behaviour
  - Does not inspect packets for TOS/DSCP
- Provides equivalent statistics to “Show policy-map interface”
  - Counters can not be reset
- Navigation is complex
- Supported platforms:
  - All routers (only policing and marking on 7600)
  - Cat6K running IOS native (only policing and marking)
  - not supported on Cat2xxx/3xxx/4xxx

NetFlow

- NetFlow answers the who, what, when, where, and how network traffic is flowing
  - Provides flow information Per Class of Service (TOS)
  - Provides pre and post policy QoS classification
- Flows are defined by seven keys:
  - Source IP address
  - Destination IP address
  - Source port
  - Destination port
  - Layer 3 protocol
  - **TOS byte (DSCP)**
  - Input interface (ifIndex)

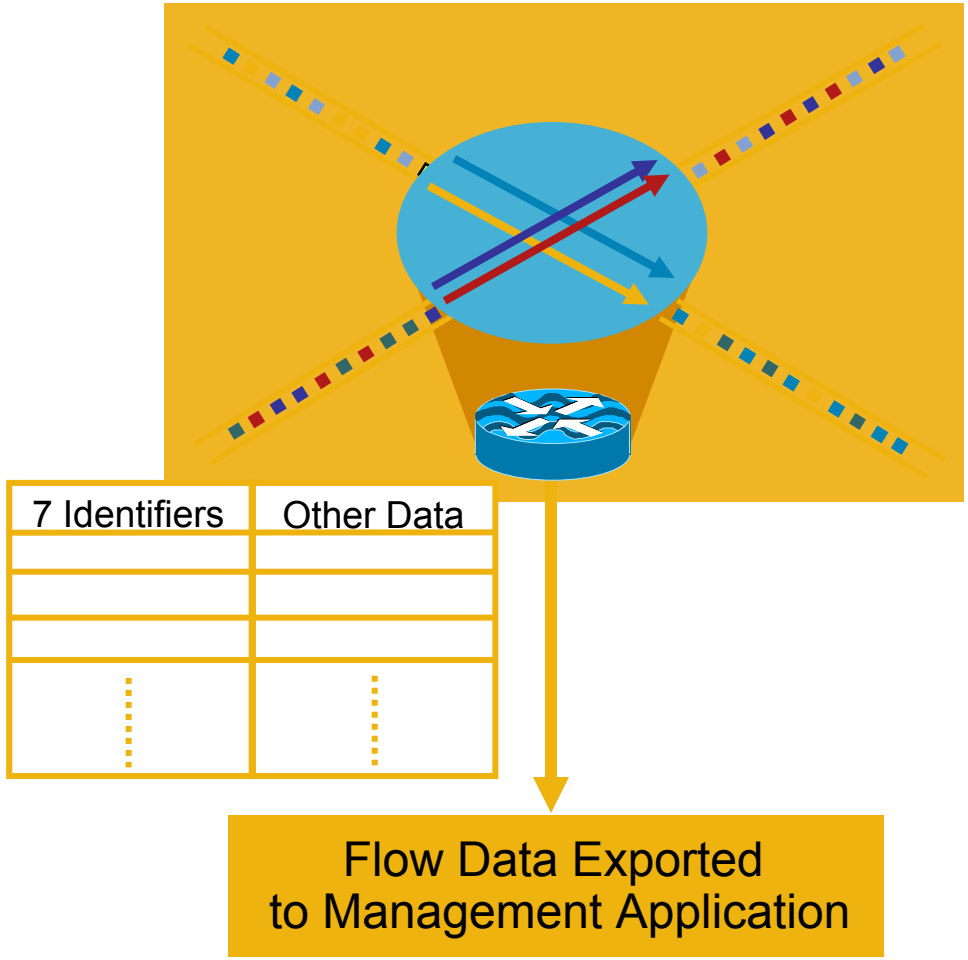

What Is a Traditional IP Flow

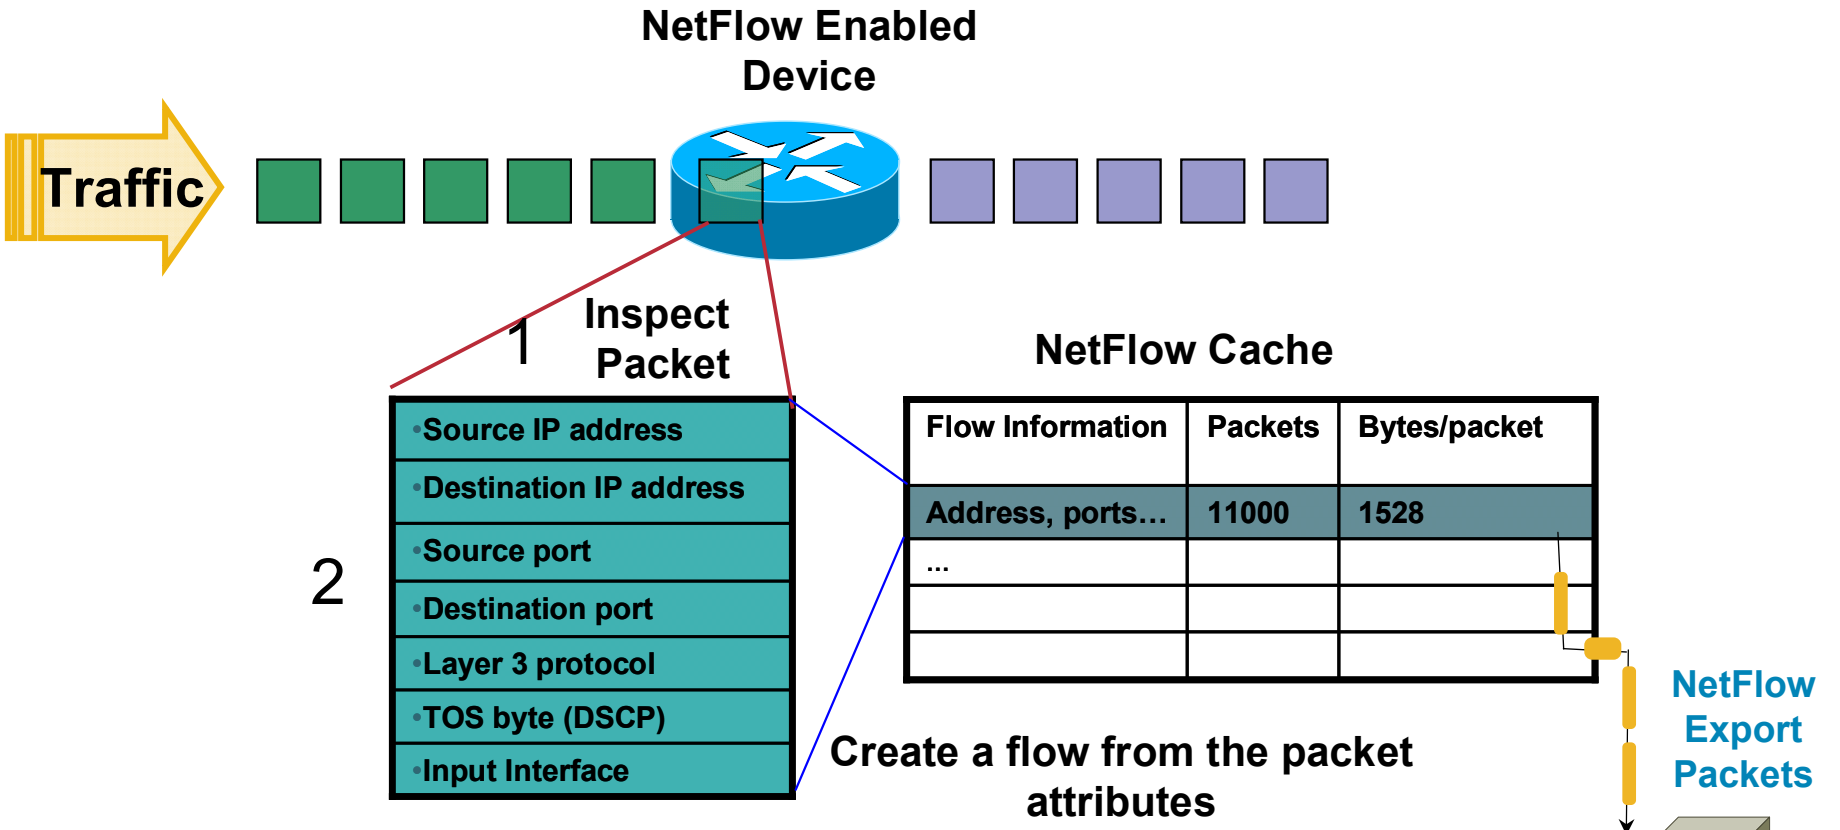

1. Inspect a packet's seven key fields and identify the values
2. If the set of key field values is unique create a flow record or cache entry
3. When the flow terminates export the flow to the collector

Network Management Tools
Network-Based Application Recognition

Stateful and dynamic inspection

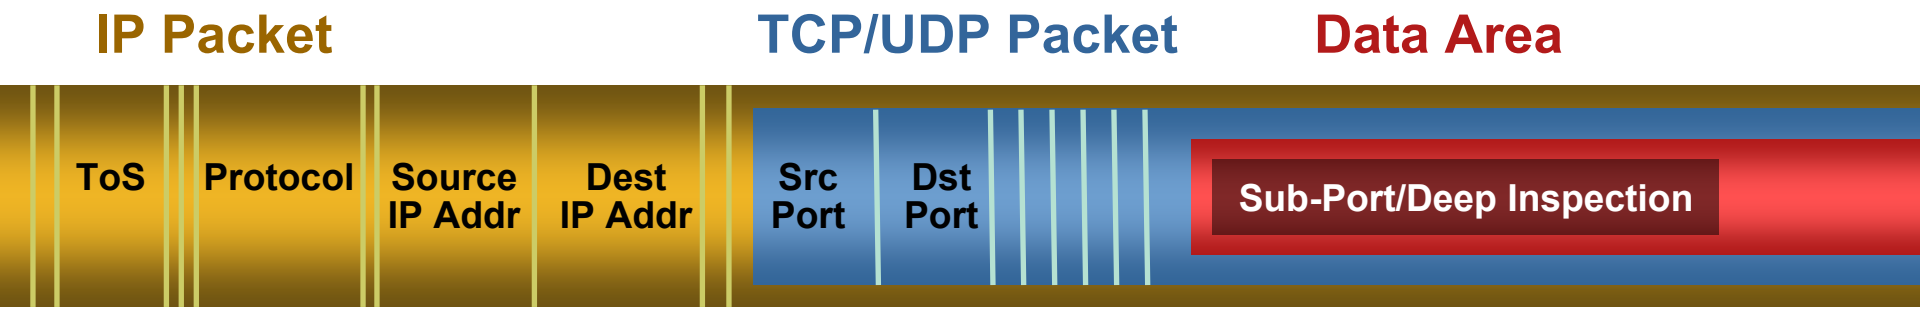

- Identifies over 90 applications and protocols TCP and UDP port numbers
  - Statically assigned
  - Dynamically assigned during connection establishment
- Non-TCP and non-UDP IP protocols
- Data packet inspection for matching values

Measurement Technology: IP SLAs

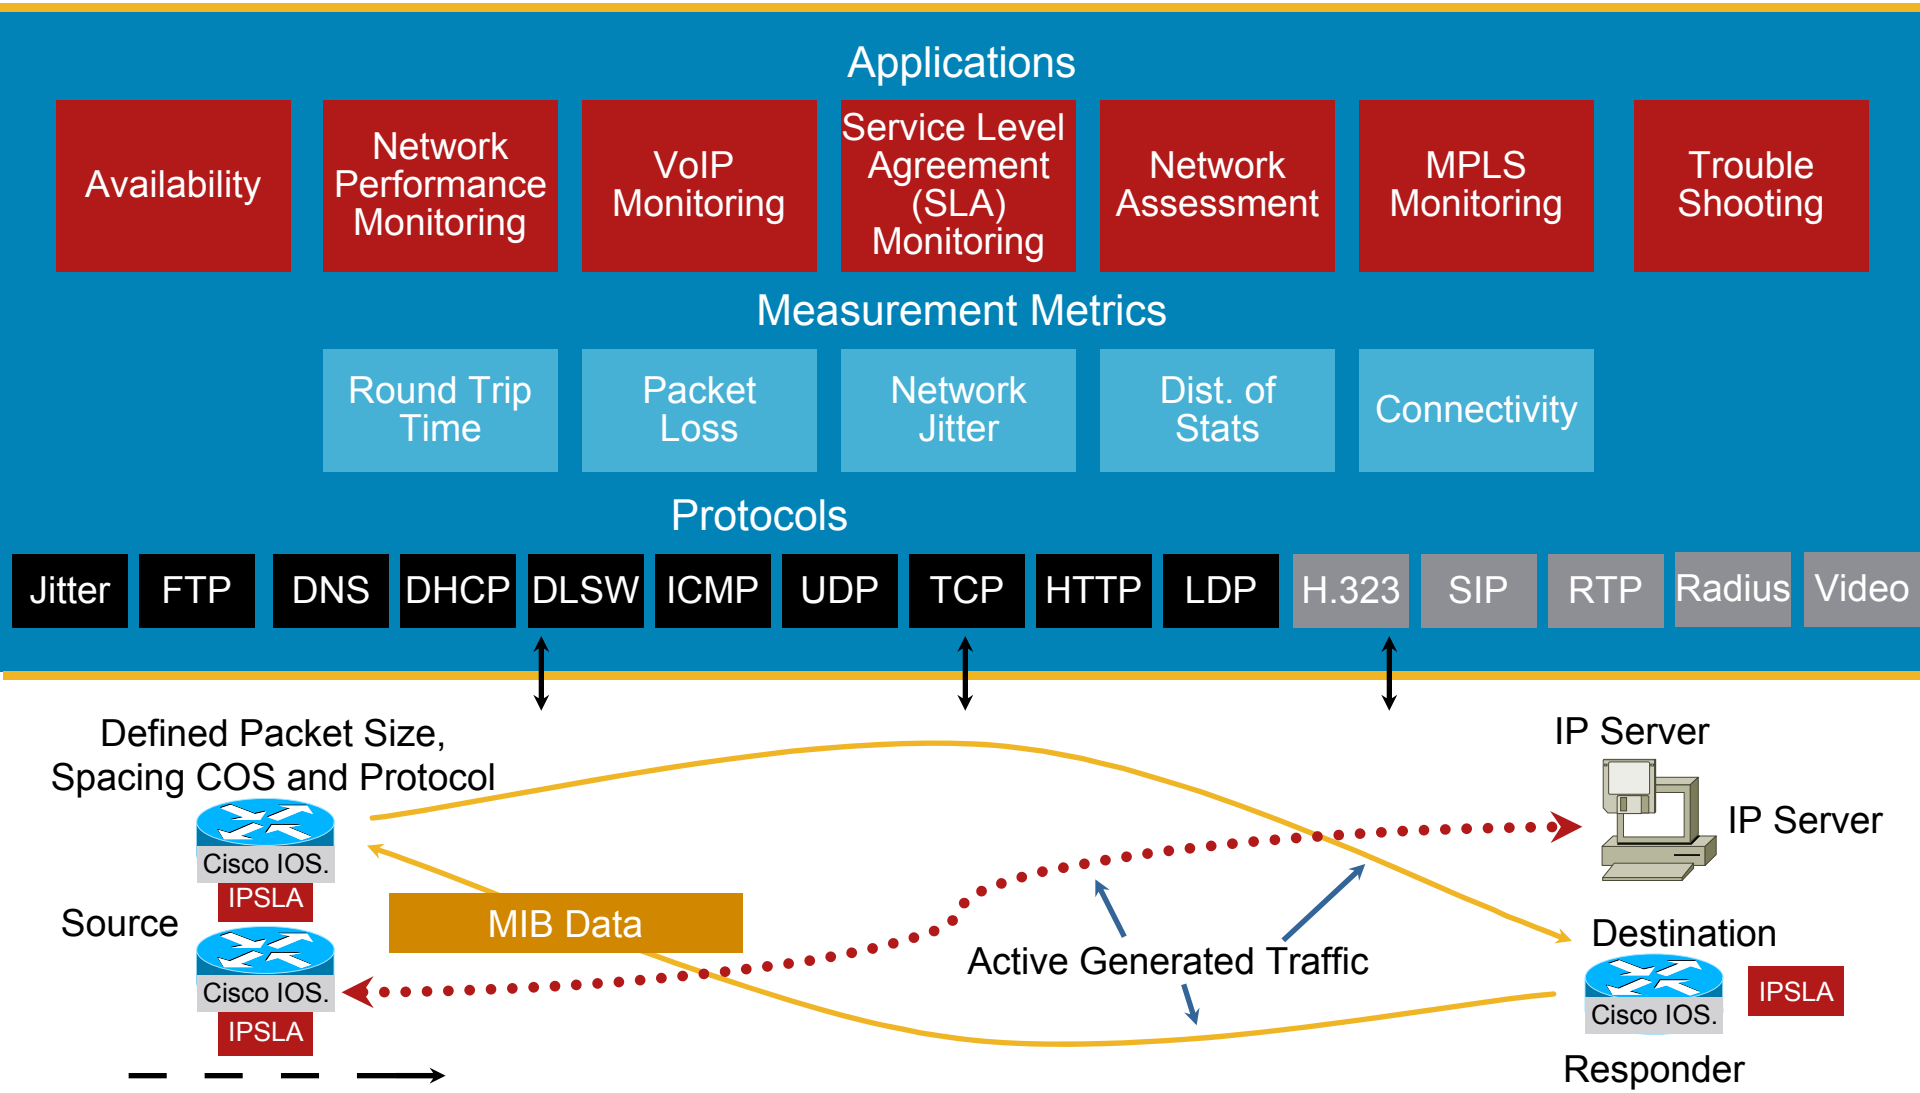

What are IP SLAs?

- IP SLA is an active probing and monitoring feature in Cisco IOS
- Wide protocol and applications coverage: UDP, TCP, ICMP, HTTP, DNS, DHCP, FTP,…
- Microsecond granularity
- Use it through SNMP or CLI
- Already in Cisco IOS (available on most platforms and interfaces type)

Q and A

AT-A-GLANCE SUMMARIES

QoS TOOLS AT–A–GLANCE

Quality of Service (QoS) is the measure of transmission quality and service availability of a network (or internetworks). The transmission quality of the network is determined by the following factors: Latency, Jitter, and Loss.

QoS technologies refer to the set of tools and techniques to manage network resources and are considered the key enabling technologies for the transparent convergence of voice, video, and data networks. Additionally, QoS tools can play a strategic role in significantly mitigating DoS/worm attacks.

Cisco QoS toolset consists of the following:

- Classification and Marking tools
- Policing and Markdown tools
- Scheduling tools
- Link-specific tools
- AutoQoS tools

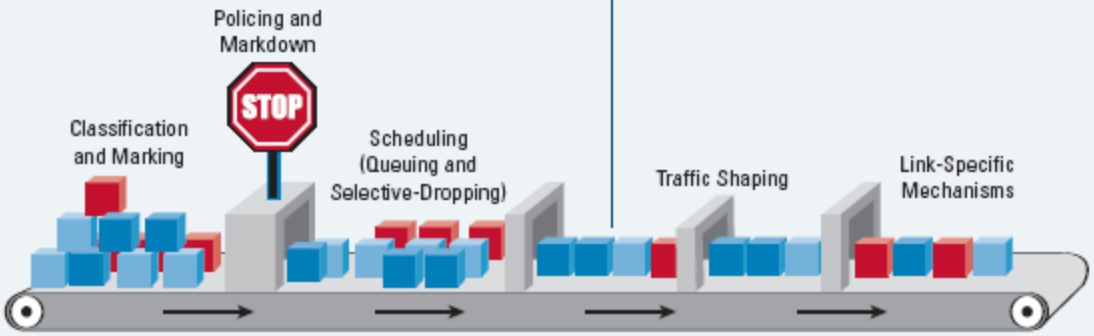

Classification can be Done at Layers 2–7

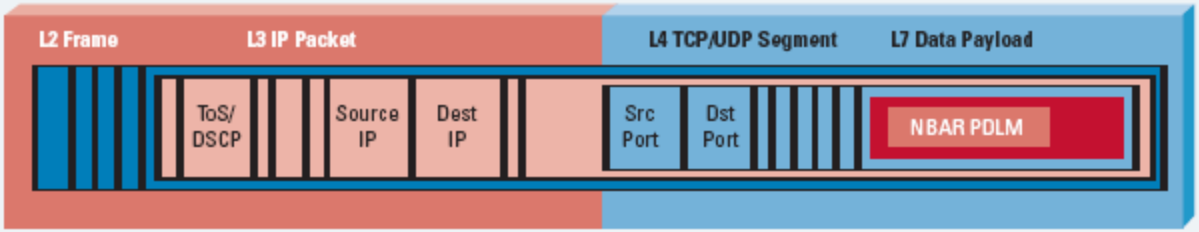

Marking can be done at Layers 2 or Layer 3:

- Layer 2: 802.1Q/p CoS, MPLS EXP
- Layer 3: IP Precedence, DSCP and/or IP ECN

Layer 3 (IP ToS Byte) Marking Options

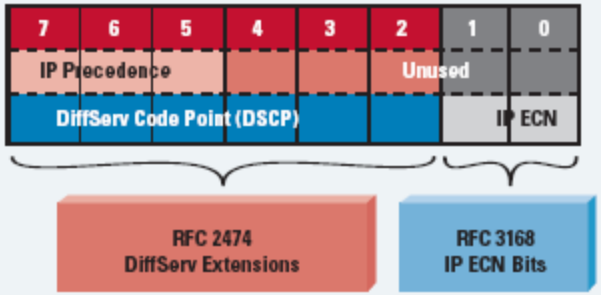

Cisco recommends end-to-end marking at Layer 3 with standards-based DSCP values.

Policing tools can complement marking tools by marking metering flows and marking-down out-of-contract traffic.

Policers Meter Traffic Into Three Categories:

- Violate: No More Traffic is Allowed Beyond This Upper-Limit (Red Light)
- Exceed: Moderate Bursting is Allowed (Yellow Light)
- Conform: Traffic is Within the Defined Rate (Green Light)

Scheduling tools re-order and selectively-drop packets whenever congestion occurs.

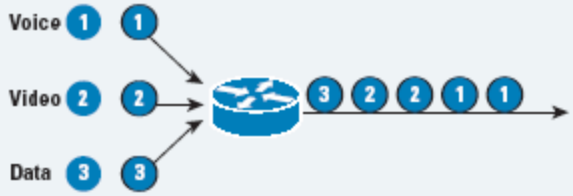

Link-Specific tools are useful on slow-speed WAN/VPN links and include shaping, compression, fragmentation, and interleaving.

AutoQoS features automatically configure Cisco recommended QoS on Cisco Catalyst® switches and Cisco IOS® Software routers with just one or two commands.

204170.k_ETMG_AE_4.05

THE QoS BASELINE
AT-A-GLANCE

The QoS Baseline is a strategic document designed to unify QoS within Cisco. The QoS Baseline provides uniform, standards-based recommendations to help ensure that QoS products, designs, and deployments are unified and consistent.

The QoS Baseline defines up to 11 classes of traffic that may be viewed as critical to a given enterprise. A summary of these classes and their respective standards-based markings and recommended QoS configurations are shown below.

Interactive-Video refers to IP Video-Conferencing; Streaming Video is either unicast or multicast uni-directional video.

The (Locally-Defined) Mission-Critical class is intended for a subset of Transactional Data applications that contribute most significantly to the business objectives (this is a non-technical assessment).

The Transactional Data class is intended for foreground, user-interactive applications such as database access, transaction services, interactive messaging, and preferred data services.

The Bulk Data class is intended for background, non-interactive traffic flows, such as large file transfers, content distribution, database synchronization, backup operations, and email.

The IP Routing class is intended for IP Routing protocols, such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and etc.

The Call-Signaling class is intended for voice and/or video signaling traffic, such as Skinny, SIP, H.323, etc.

The Network Management class is intended for network management protocols, such as SNMP, Syslog, DNS, etc.

Standards-based marking recommendations allow for better integration with service-provider offerings as well as other internetworking scenarios.

In Cisco IOS Software , rate-based queuing translates to CBWFQ; priority queuing is LLQ.

Application	L3 Classification PHB	DSCP	Referencing Standard	Recommended Configuration
IP Routing	CS6	48	RFC 2474-4.2.2	Rate-Based Queuing + RED
Voice	EF	46	RFC 3246	RSVP Admission Control + Priority Queuing
Interactive-Video	AF41	34	RFC 2597	RSVP + Rate-Based Queuing + DSCP-WRED
Streaming Video	CS4	32	RFC 2474-4.2.2	RSVP + Rate-Based Queuing + RED
Mission-Critical	AF31	26	RFC 2597	Rate-Based Queuing + DSCP-WRED
Call-Signaling	CS3	24	RFC 2474-4.2.2	Rate-Based Queuing + RED
Transactional Data	AF21	18	RFC 2597	Rate-Based Queuing + DSCP-WRED
Network Mgmt	CS2	16	RFC 2474-4.2.2	Rate-Based Queuing + RED
Bulk Data	AF11	10	RFC 2597	Rate-Based Queuing + DSCP-WRED
Scavenger	CS1	8	Internet 2	No BW Guarantee + RED
Best Effort	0	0	RFC 2474-4.1	BW Guarantee Rate-Based Queuing + RED

DSCP-Based WRED (based on RFC 2597) drops AFx3 before AFx2, and in turn drops AFx2 before AFx1. RSVP is recommended (whenever supported) for Voice and/or Interactive-Video admission control

Cisco products that support QoS features will use these QoS Baseline recommendations for marking, scheduling, and admission control.

The Scavenger class is based on an Internet 2 draft that defines a "less-than-Best Effort" service. In the event of link congestion, this class will be dropped the most aggressively.

The Best Effort class is also the default class. Unless an application has been assigned for preferential/deferential service, it will remain in this default class. Most enterprises have hundreds—if not thousands—of applications on their networks; the majority of which will remain in the Best Effort service class.

The QoS Baseline recommendations are intended as a standards-based guideline for customers—not as a mandate.

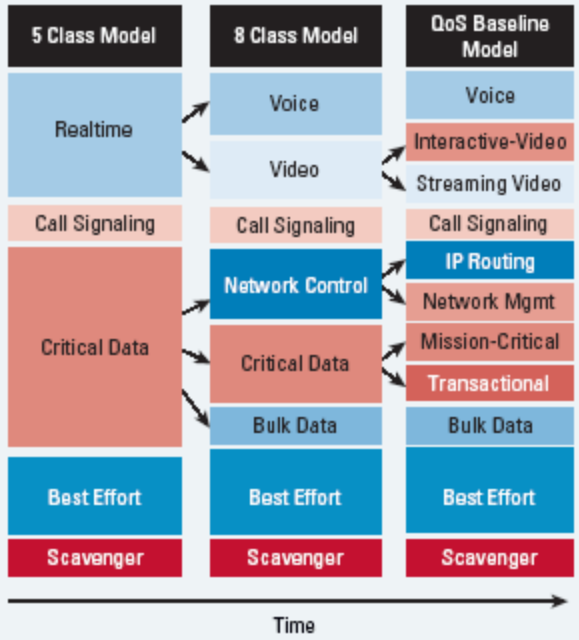

204170.1_ETMG_AE_4.05

QoS BEST-PRACTICES
AT–A–GLANCE

A successful QoS deployment includes three key phases:

1) Strategically defining the business objectives to be achieved via QoS
2) Analyzing the service-level requirements of the traffic classes
3) Designing and testing QoS policies

1) STRATEGICALLY DEFINING THE BUSINESS OBJECTIVES TO BE ACHIEVED BY QOS

Business QoS objectives need to be defined:

- Is the objective to enable VoIP only or is video also required?
- If so, is video-conferencing or streaming video required? Or both?
- Are there applications that are considered mission-critical? If so, what are they?
- Does the organization wish to squelch certain types of traffic? If so, what are they?
- Does the business want to use QoS tools to mitigate DoS/worm attacks?
- How many classes of service are needed to meet the business objectives?

Because QoS introduces a system of managed unfairness, most QoS deployments inevitably entail political repercussions when implemented. To minimize the effects of non-technical obstacles to deployment, address political/organizational issues as early as possible, garnishing executive endorsement whenever possible.

2) ANALYZE THE APPLICATION SERVICE-LEVEL REQUIREMENTS

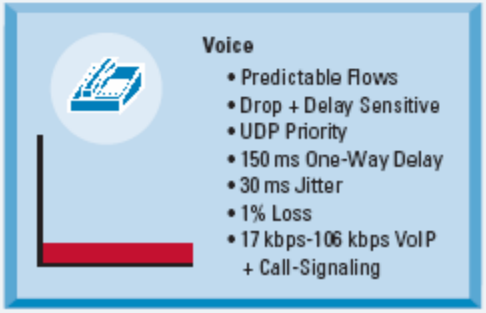

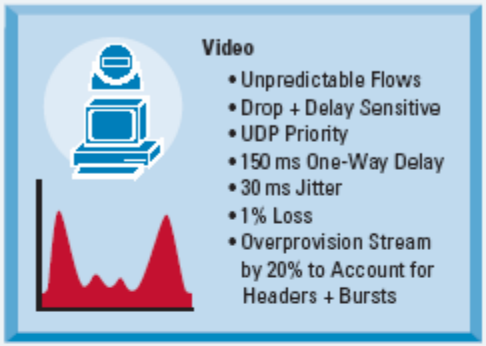

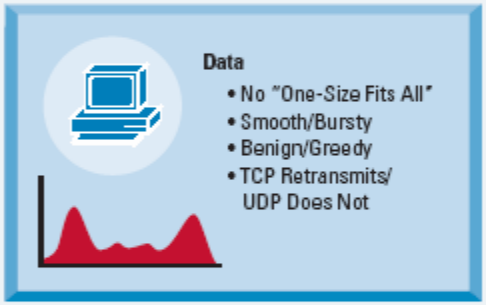

3) DESIGN AND TEST THE QoS POLICIES

Application	L3 Classification PHB	DSCP
Routing	CS6	48
Voice	EF	46
Interactive-Video	AF41	34
Streaming Video	CS4	32
Mission-Critical	AF31	26
Call-Signaling	CS3	24
Transactional Data	AF21	18
Network Mgmt	CS2	16
Bulk Data	AF11	10
Scavenger	CS1	8
Best Effort	0	0

Classify, mark, and police as close to the traffic-sources as possible; following Differentiated-Services standards, such as RFC 2474, 2475, 2597, 2698 and 3246.

Provision queuing in a consistent manner (according to hardware capabilities).

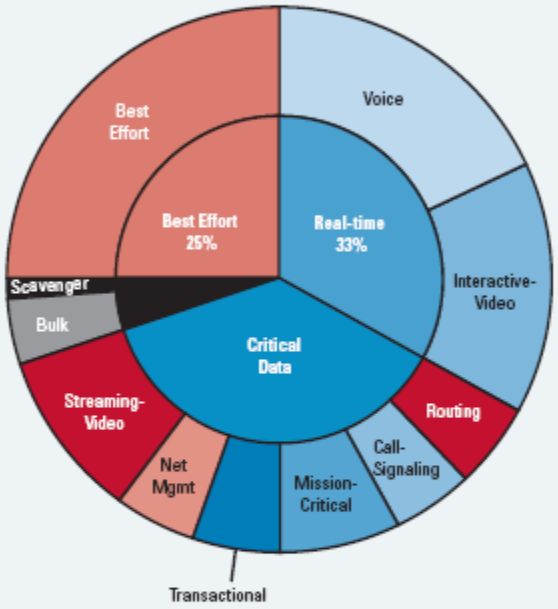

Thoroughly test QoS policies prior to production-network deployment.

A successful QoS policy rollout is followed by ongoing monitoring of service levels and periodic adjustments and tuning of QoS policies.

As business conditions change, the organization will need to adapt to these changes and may be required to begin the QoS deployment cycle anew, by redefining their objectives, tuning and testing corresponding designs, rolling these new designs out and monitoring them to see if they match the redefined objectives.

204170.m_ETMG_AE_4.05

SCAVENGER-CLASS QoS STRATEGY FOR DOS/WORM ATTACK MITIGATION
AT–A–GLANCE

DoS and worm attacks are exponentially increasing in frequency, complexity, and scope of damage.

QoS tools and strategic designs can mitigate the effects of worms and keep critical applications available during DoS attacks.

One such strategy, referred to as Scavenger-class QoS, uses a two-step tactical approach to provide first- and second-order anomaly detection and reaction to DoS/worm attack-generated traffic.

The first step in deploying Scavenger-class QoS is to profile applications to determine what constitutes a normal vs. abnormal flow (within a 95% confidence interval).

Application traffic exceeding this normal rate will be subject to first-order anomaly detection at the Campus Access-Edge, specifically: excess traffic will be marked down to Scavenger (DSCP CS1/8).

Note that anomalous traffic is not dropped or penalized at the edge; it is simply remarked.

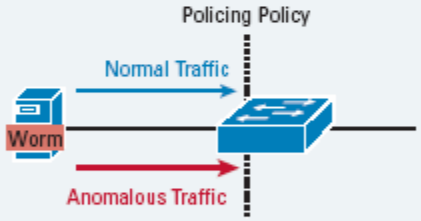

Only traffic in excess of the normal/abnormal threshold is remarked to Scavenger.

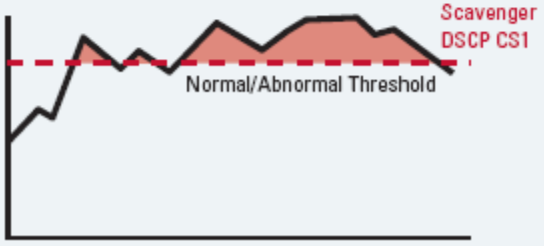

Campus Access-Edge policing policies are coupled with Scavenger-class queuing policies on the uplinks to the Campus Distribution Layer.

Queuing policies only engage when links are congested. Therefore, only if uplinks become congested, traffic begin to be dropped.

Anomalous traffic—previously marked to Scavenger—is dropped the most aggressively (only after all other traffic types have been fully-serviced).

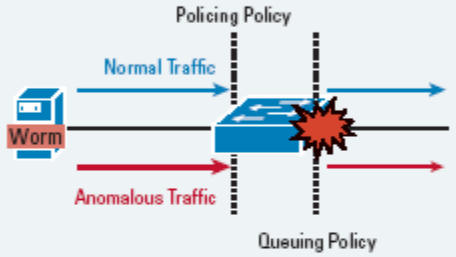

A key point of this strategy is that legitimate traffic flows that temporarily exceed thresholds are not penalized by Scavenger-class QoS.

Only sustained, abnormal streams generated simultaneously by multiple hosts (highly-indicative of DoS/worm attacks) are subject to aggressive dropping—and such dropping only occurs *after* legitimate traffic has beean fully-serviced.

The Campus uplinks are not the only points in the network infrastructure where congestion could occur. Typically WAN and VPN links are the first to congest.

Therefore, Scavenger-class "less-than-Best-Effort" queuing should be provisioned on all network devices in a consistent manner (according to hardware capabilities).

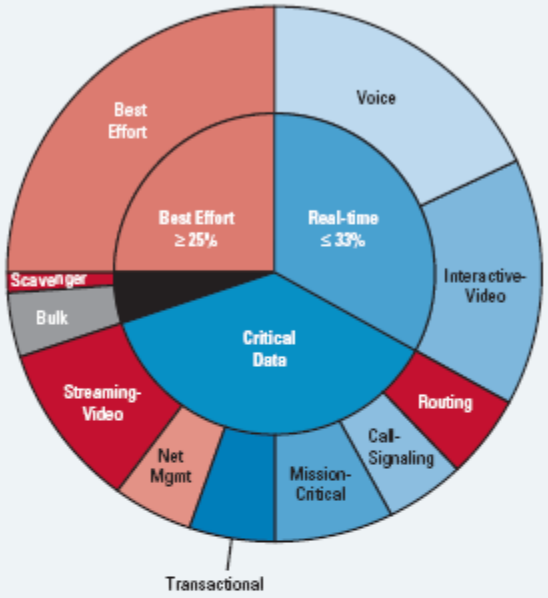

Thoroughly test QoS policies prior to production-network deployment.

It is critically important to recognize, that even when Scavenger-class QoS has been deployed end-to-end, this tactic only mitigates the effects of certain types of DoS/worm attacks, and does not prevent them or remove them entirely. Scavenger-class QoS is just one element of a comprehensive Cisco Self-Defending Networks (SDN) strategy.

204170.n_ETMG_AE_4.05

CAMPUS QoS DESIGN
AT-A-GLANCE

QoS policies should always be enabled in Cisco Catalyst® switches—rather than router software—whenever a choice exists.

Three main types of QoS policies are required within the Campus:

1) Classification and Marking

2) Policing and Markdown

3) Queuing

Classification, marking, and policing should be performed as close to the traffic-sources as possible, specifically at the Campus Access-Edge. Queuing, on the other hand, needs to be provisioned at all Campus Layers (Access, Distribution, Core) due to oversubscription ratios.

Classify and mark as close to the traffic-sources as possible following Cisco QoS Baseline marking recommendations, which are based on Differentiated-Services standards, such as: RFC 2474, 2597 & 3246.

Application	L3 Classification PHB	DSCP
Routing	CS6	48
Voice	EF	46
Interactive-Video	AF41	34
Streaming Video	CS4	32
Mission-Critical	AF31	26
Call-Signaling	CS3	24
Transactional Data	AF21	18
Network Mgmt	CS2	16
Bulk Data	AF11	10
Scavenger	CS1	8
Best Effort	0	0

Access-Edge policers, such as this one, detect anomalous flows and remark these to Scavenger (DSCP CS1).

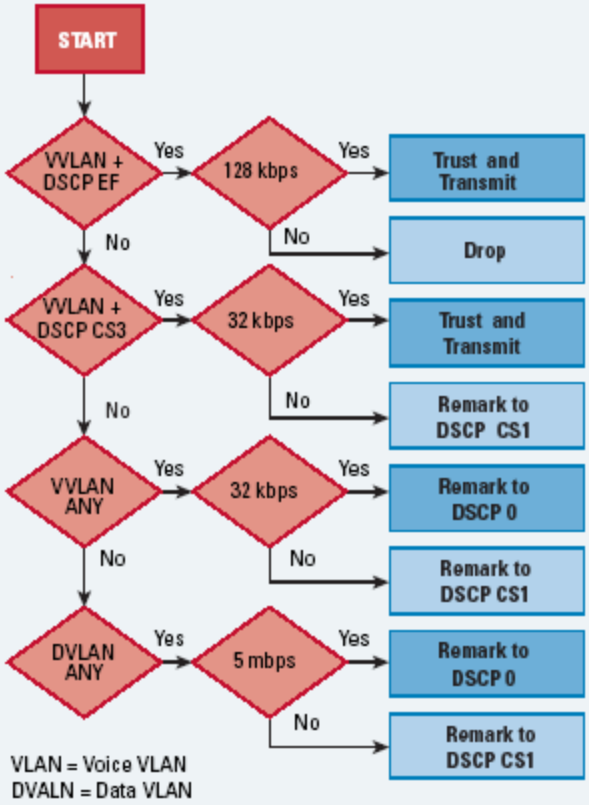

VLAN = Voice VLAN
DVALN = Data VLAN

Queuing policies will vary by platform:

E.g. 1P3Q1T P = Priority Queue
Q = Non-Priority Queue
T = WRED Threshold

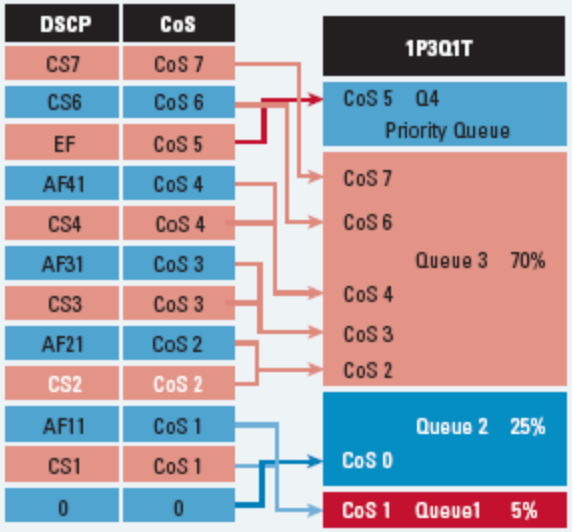

Campus Access switches require the following QoS policies:

- Appropriate (endpoint-dependant) trust policies, and/or classification and marking policies
- Policing and markdown policies
- Queuing policies.

Campus Distribution and Core switches require the following QoS policies:

- DSCP trust policies
- Queuing policies
- Optional per-user microflow policing policies (only on distribution layer Catalyst 6500s with Sup720s.)

204170.o_ETMG_AE_4.05

WAN QoS DESIGN
AT-A-GLANCE

In an enterprise network infrastructure, bandwidth is scarcest—and thus most expensive—over the WAN. Therefore, the business case for efficient bandwidth optimization via QoS technologies is strongest over the WAN.

WAN QoS policies need to be configured on the WAN edges of WAN Aggregator (WAG) routers and Branch routers. WAN edge QoS policies include queuing, shaping, selective-dropping, and link-specific policies.

The number of WAN classes of traffic is determined by the business objectives and may be expanded over time.

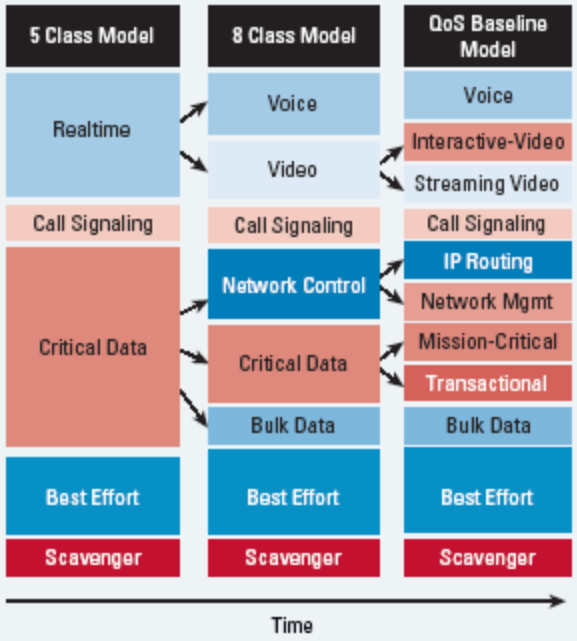

WAN links can be categorized into three main speed groups:

- Slow-Speed (≤ 768 kbps)
- Medium-Speed (>768 kbps & ≤ T1/E1)
- High-Speed (≥ T1/E1)

Queuing Models for 5/8/11 Classes of Service

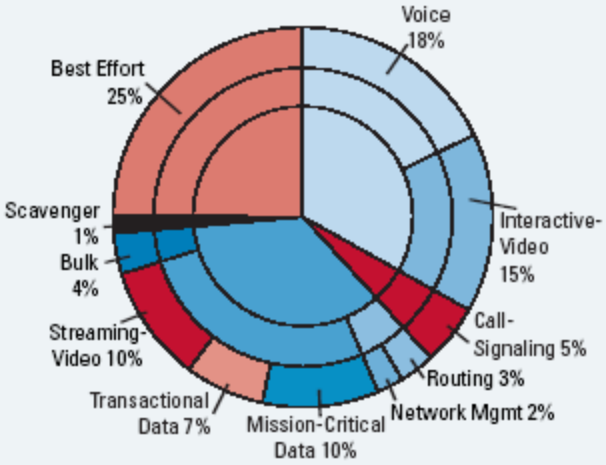

WAN QoS Tools: RTP Header Compression (cRTP)

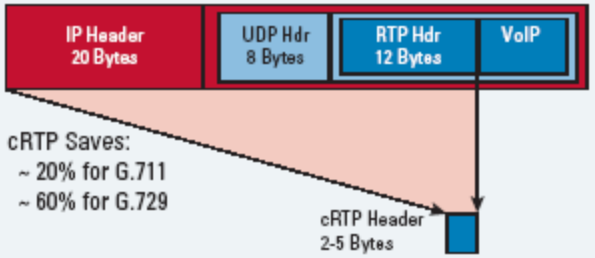

WAN QoS Tools: Link Fragmentation and Interleaving

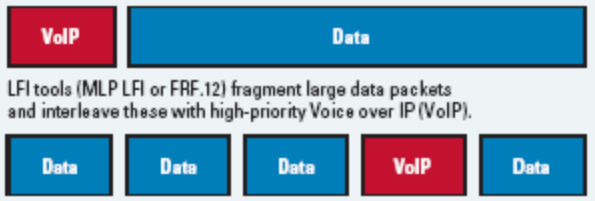

LFI tools (MLP LFI or FRF.12) fragment large data packets and interleave these with high-priority Voice over IP (VoIP).

LINK-SPECIFIC DESIGN RECOMMENDATIONS

Leased-Line (MLP) Link

- Use MLP link fragmentation and interleaving (LFI) and cRTP on Slow-Speed links

Frame Relay Link

- Use Frame-Relay traffic shaping
  - Set CIR to 95% of guaranteed rate
  - Set Committed Burst to CIR/100
  - Set Excess Burst to 0
- Use FRF.12 and and cRTP on Slow-Speed links

ATM Link

- Use MLP LFI (via MLPoATM) and cRTP on Slow-Speed links
- Set the ATM PVC Tx-Ring to 3 for Slow-Speed links

204170.p_ETMG_AE_4.05

BRANCH QoS DESIGN
AT-A-GLANCE

Branch routers are connected to central sites via private-WAN or VPN links which often prove to be the bottlenecks for traffic flows. QoS policies at these bottlenecks align expensive WAN/VPN bandwidth utilization with business objectives.

QoS designs for Branch routers are—for the most part—identical to WAN Aggregator QoS designs. However, Branch routers require three unique QoS considerations:

1) Unidirectional applications
2) Ingress classification requirements
3) Network Based Application Recognition (NBAR) policies for worm policing

Each of these Branch router QoS design considerations will be overviewed.

1) UNIDIRECTIONAL APPLICATIONS

Some applications (like Streaming Video) usually only traverse the WAN/VPN in the Campus-to-Branch direction; and therefore, do not require provisioning in the Branch-to-Campus direction on the Branch router's WAN edge.

Bandwidth for such unidirectional application classes can be reassigned to other critical classes, as shown in the following diagram. Notice that no Streaming Video class is provisioned and the bandwidth allocated to it (on the Campus side of the WAN link) is reallocated to the Mission-Critical and Transactional Data classes.

An Example 10-Class QoS Baseline Branch Router WAN Edge Queuing Model

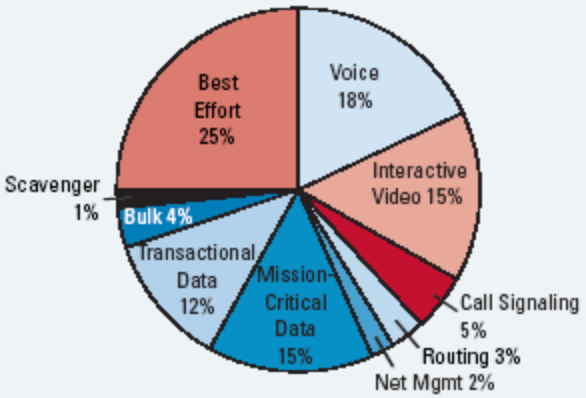

2) INGRESS CLASSIFICATION

Branch-to-Campus traffic may not be correctly marked on the Branch Access Layer switch.

These switches—which are usually lower-end switches—may or may not have the capabilities to classify and mark application traffic. Therefore, classification and marking may need to be performed on the Branch router's LAN edge (in the ingress direction).

Furthermore, Branch routers offer the ability to use NBAR to classify and mark traffic flows that require stateful packet inspection.

3) NBAR FOR KNOWN WORM POLICING

Worms are nothing new, but they have increased exponentially in frequency, complexity, and scope of damage in recent years.

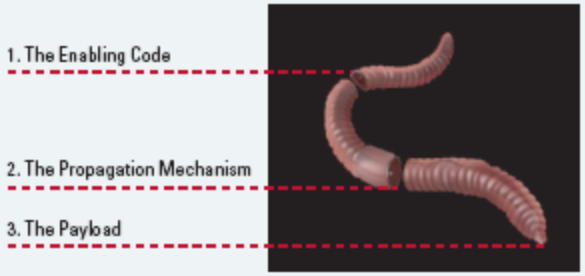

The Branch router's ingress LAN edge is a strategic place to use NBAR to identify and drop worms, such as CodeRed, NIMDA, SQL Slammer, MS-Blaster, and Sasser.

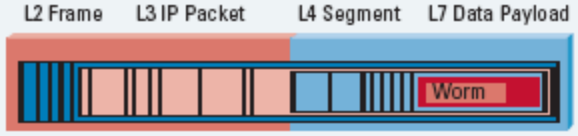

NBAR extensions allow for custom Packet Data Language Modules (PDLMs) to be defined for future worms.

Where is QoS Required on Branch Routers?

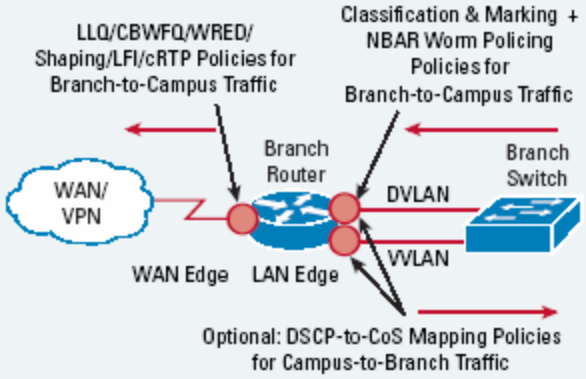

204170.q_ETMG_AE_4.05

QoS DESIGN FOR MPLS VPN SUBSCRIBERS
AT-A-GLANCE

QoS design for an enterprise subscribing to a Multiprotocol Label Switching (MPLS) VPN requires a major paradigm shift from private-WAN QoS design.

This happens because with private-WAN design, the enterprise principally controlled QoS. The WAN Aggregator (WAG) provisioned QoS for not only Campus-to-Branch traffic, but also for Branch-to-Branch traffic (which was homed through the WAG).

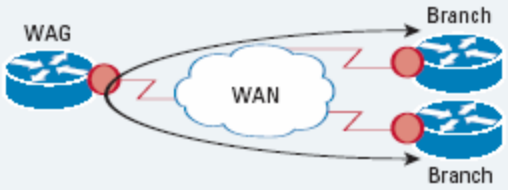

However, due to the any-to-any/full-mesh nature of MPLS VPNs, Branch-to-Branch traffic is no longer homed through the WAG. While Branch-to-MPLS VPN QoS is controlled by the enterprise (on their Customer-Edge—CE—routers), MPLS VPN-to-Branch QoS is controlled by the service provider (on their Provider Edge—PE—routers).

Therefore, to guarantee end-to-end QoS, enterprises must co-manage QoS with their MPLS VPN service providers; their policies must be both consistent and complementary.

MPLS VPN service providers offer classes of service to enterprise subscribers.

Admission criteria for these classes is the DSCP markings of enterprise traffic. Thus, enterprises may have to remark application traffic to gain admission into the required service provider class.

Some best practices to consider when assigning enterprise traffic to service provider classes of service include:

- Do not put Voice and Interactive-Video into the Realtime class on slow-speed (≤ 768 kbps) CE-to-PE links
- Do not put Call-Signaling into the Realtime class on slow-speed CE-to-PE links
- Do not mix TCP applications with UDP applications within a single service provider class (whenever possible); UDP applications may dominate the class when congested

Example—enterprise subscriber DSCP Remarking Diagram and CE Edge Bandwidth Allocation Diagram.

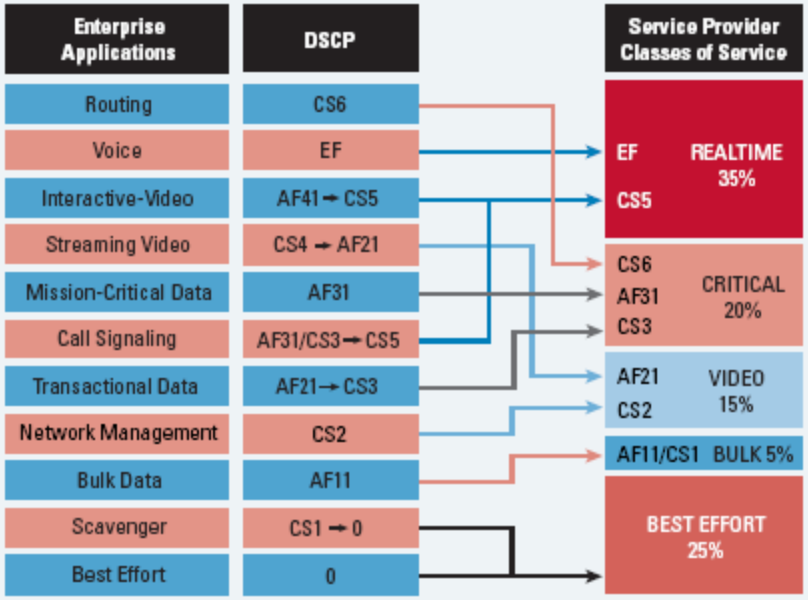

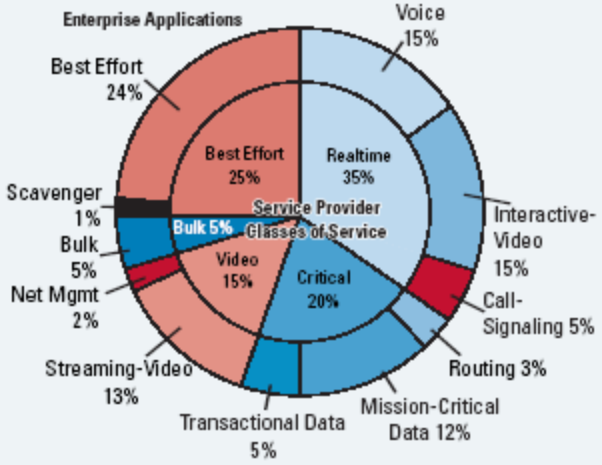

A general DiffServ principle is to mark or trust traffic as close to the source as administratively and technically possible. However, certain traffic types might need to be re-marked before handoff to the service provider to gain admission to the correct class. If such re-marking is required, it is recommended that the re-marking be performed at the CE's egress edge, not within the campus. This is because service-provider service offerings likely will evolve or expand over time, and adjusting to such changes will be easier to manage if re-marking is performed only at CE egress edges.

204170.r_ETMG_AE_4.05

QoS DESIGN FOR MPLS VPN SERVICE PROVIDERS
AT–A–GLANCE

In order to support enterprise-subscriber voice, video, and data networks, service providers must include QoS provisioning within their Multiprotocol Label Switching (MPLS) VPN service offerings.

This is due to the any-to-any/full-mesh nature of MPLS VPNs, where enterprise subscribers depend on their service providers to provision Provider-Edge (PE) to Customer-Edge (CE) QoS policies consistent with their CE-to-PE policies.

In addition to these PE-to-CE policies, service providers will likely implement ingress policers on their PEs to identify whether traffic flows are in- or out-of-contract. Optionally, service providers may also provision QoS policies within their core networks, using Differentiated Services and/or MPLS Traffic Engineering (TE).

In order to guarantee end-to-end QoS, enterprises must co-manage QoS with their MPLS VPN service providers; their policies must be both consistent and complementary.

Service providers can mark at Layer 2 (MPLS EXP) or at Layer 3 (DSCP).

RFC 3270 presents three modes of MPLS/DiffServ marking for service providers:

1) Uniform Mode: SP can remark customer DSCP values

2) Pipe Mode: SP does not remark customer DSCP values (SP uses independent MPLS EXP markings); final PE-to-CE policies are based on *service provider's* markings

3) Short Pipe Mode (shown below): SP does not remark customer DSCP values (SP uses independent MPLS EXP markings); final PE-to-CE policies are based on *customer's* markings

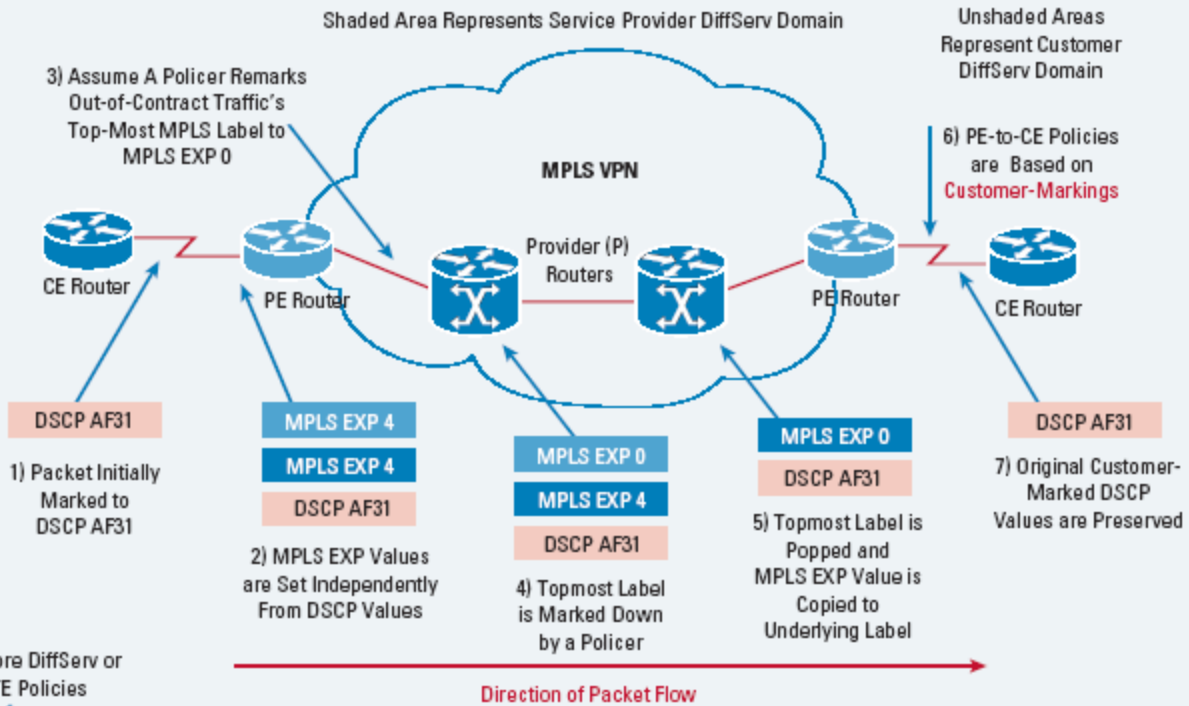

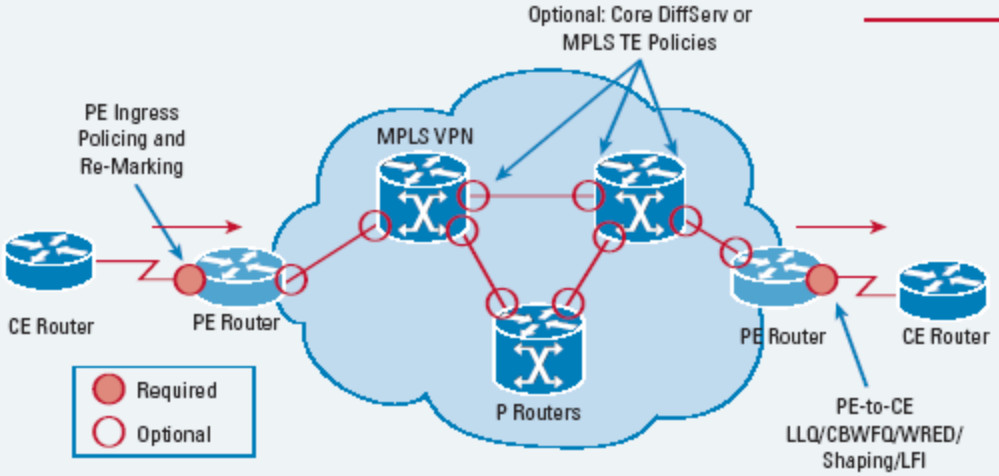

Service providers can guarantee service levels within their core by:

1) Aggregate Bandwidth Overprovisioning: adding redundant links when utilization hits 50% (simple to implement, but expensive and inefficient)

2) Core DiffServ Policies: simplified DiffServ policies for core links

3) MPLS TE: TE provides granular policy-based control over traffic flows within the core

204170.s_ETMG_AE_4.05

QoS DESIGN FOR IPsec VPNs
AT–A–GLANCE

IPsec VPNs achieve network segregation and privacy via encryption. IPsec VPNs are built by overlaying a point-to-point mesh over the Internet using Layer 3-encrypted tunnels. Encryption/decryption is performed at these tunnel end-points, and the protected traffic is carried across the shared network.

Three main QoS considerations specific to IPsec VPNs are:

1) Additional bandwidth required by IPsec encryption and authentication

2) Marginal time element required at each point where encryption/decryption takes place

3) Anti-Replay interactions

1) IPsec BANDWIDTH OVERHEAD

The additional bandwidth required to encrypt and authenticate a packet needs to be factored into account when provisioning QoS policies.

This is especially important for Voice over IP (VoIP), where IPsec could more than double the size of a G.729 voice packet, as shown below.

The Layer 3 data rate for a G.729 call (at 50 pps) is 24 kbps (60 Bytes * 8 bits * 50 pps). IP GRE tunnel overhead adds 24 bytes per packet. IPsec ESP adds another 52 bytes. The combined additional overhead increases the rate from 24 kbps (clear voice) to just less than 56 kbps (IPsec ESP tunnelmode encrypted voice).

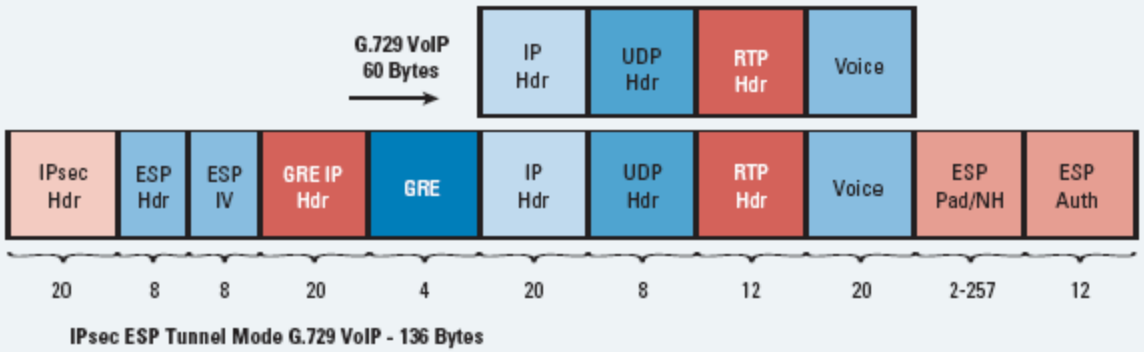

IPsec ESP Tunnel Mode G.729 VoIP - 136 Bytes

2) ENCRYPTION/DECRYPTION DELAYS

A marginal time element for encryption and decryption should be factored into the end-to-end delay budget for realtime applications, such as VoIP. Typically these processes require 2–10 ms per hop, but may be doubled in the case of spoke-to-spoke VoIP calls that are homed through a central VPN headend hub.

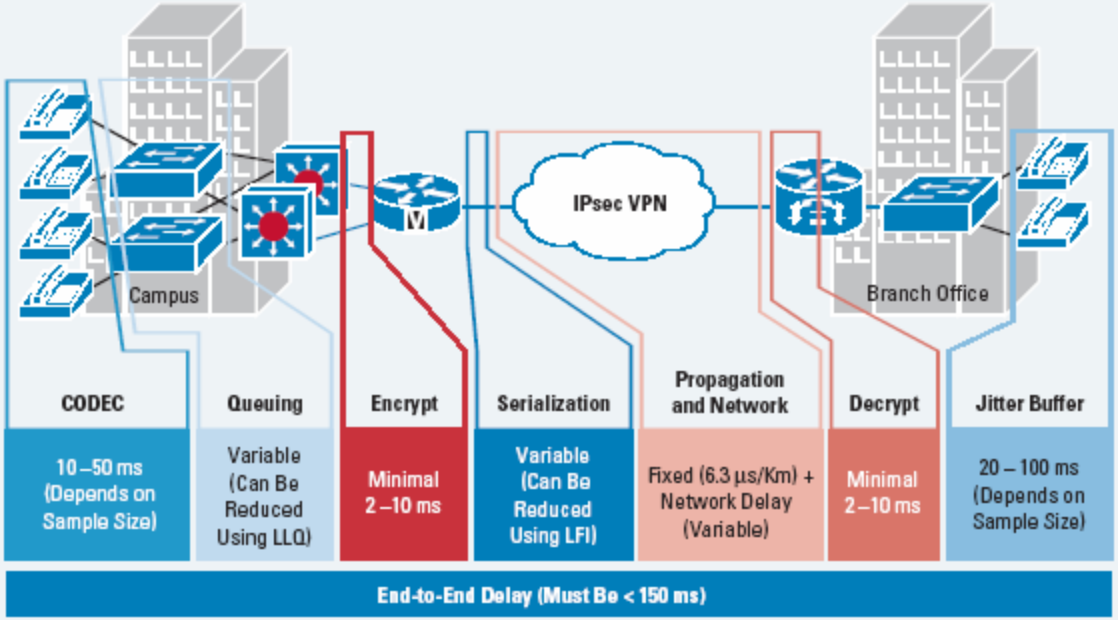

3) ANTI-REPLAY INTERACTIONS

Anti-Relay is a standards-defined mechanism to protect IPsec VPNs from hackers. If packets arrive outside of a 64-byte window, then they are considered hacked and are dropped prior to decryption. QoS queuing policies may re-order packets such that they fall outside of the Anti-Replay window. Therefore, IPsec VPN QoS policies need to be properly tuned to minimize Anti-Replay drops.

204170.r_ETMG_AE_4.05

References

Reference Materials
DiffServ Standards

- RFC 2474 “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers” http://www.apps.ietf.org/rfc/rfc2474.html
- RFC 2475 “An Architecture for Differentiated Services” http://www.ietf.org/rfc/rfc2475.txt
- RFC 2597 “Assured Forwarding PHB Group” http://www.ietf.org/rfc/rfc2597.txt
- RFC 2697 “A Single Rate Three Color Marker” http://www.ietf.org/rfc/rfc2697.txt
- RFC 2698 “A Two Rate Three Color Marker” http://www.ietf.org/rfc/rfc2698.txt
- RFC 3246 “An Expedited Forwarding PHB (Per-Hop Behavior)” http://www.ietf.org/rfc/rfc3246.txt
- Configuration Guidelines for DiffServ Service Classes http://www.ietf.org/rfc/rfc4594.txt

Reference Materials
Cisco AutoQoS Documentation

- AutoQoS VoIP for the Cisco Catalyst 2950
  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2950/12120ea2/2950scg/swqos.htm#wp1125412
- AutoQoS VoIP for the Cisco Catalyst 2970
  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2970/12220se/2970scg/swqos.htm#wp1231112
- AutoQoS VoIP for the Cisco Catalyst 3550
  http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/12120ea2/3550scg/swqos.htm#wp1185065
- AutoQoS VoIP for the Cisco Catalyst 3750
  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/12220se/3750scg/swqos.htm#wp1231112
- AutoQoS VoIP for the Cisco Catalyst 4550
  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_2_18/config/qos.htm#1281380
- AutoQoS VoIP for Cisco IOS Routers (Cisco IOS 12.2(15)T)
  http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t15/ftautoq1.htm
- AutoQoS Enterprise for Cisco IOS Routers (Cisco IOS 12.3(7)T)
  http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_7/ftautoq2.htm

Recommended Reading

Enterprise QoS Solution Reference Network Design Guide

Version 3.3
November 2005

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

- http://www.cisco.com/go/srnd
- QoS Design Overview
- Campus QoS Design
- WAN QoS Design
- Branch QoS Design
- MPLS VPN (CE) QoS Design

Reference Materials
Cisco Press Book: End-to-End QoS Design

http://www.ciscopress.com/title/1587051761

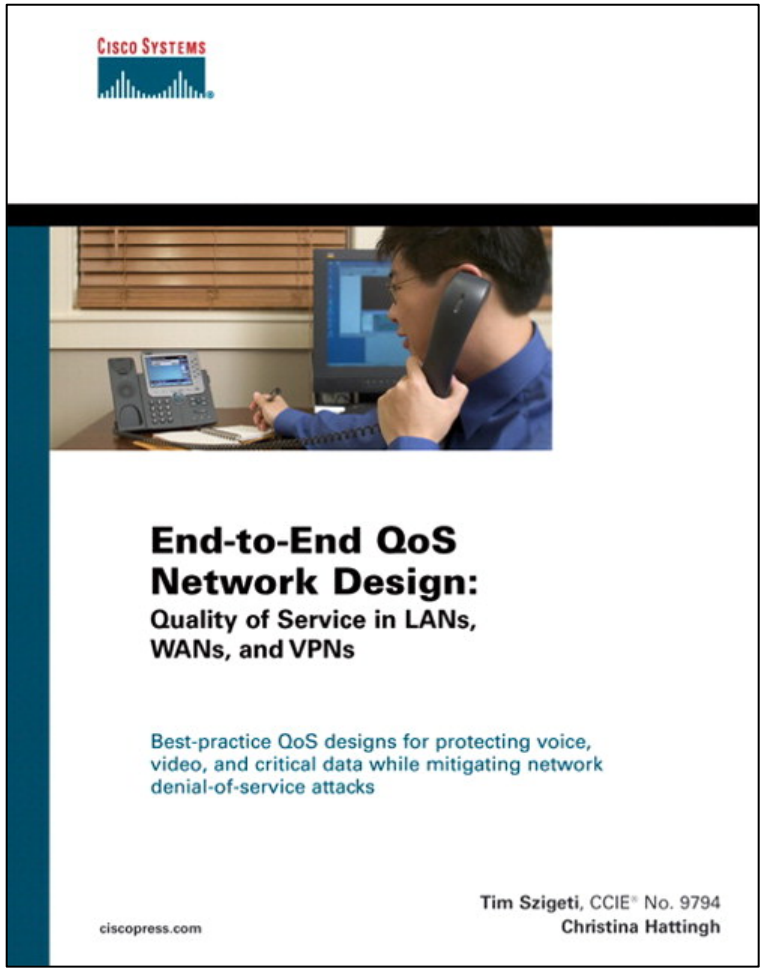

ISBN: 1587051761

Publish Date: Nov 2004

- LAN
  - Catalyst 2950
  - Catalyst 3550
  - Catalyst 2970/3560/3750
  - Catalyst 4500
  - Catalyst 6500
- WAN/Branch
  - Leased Lines
  - Frame Relay
  - ATM
  - ATM-to-FR SIW
  - ISDN
  - NBAR for Worm Policing
- VPN
  - MPLS (for Enterprise Subscribers)
  - MPLS (for Service Providers)
  - IPSec (Site-to-Site)
  - IPSec (Teleworker)

Recommended Reading

- Continue your Networkers at Cisco Live learning experience with further reading from Cisco Press
- Check the Recommended Reading flyer for suggested books

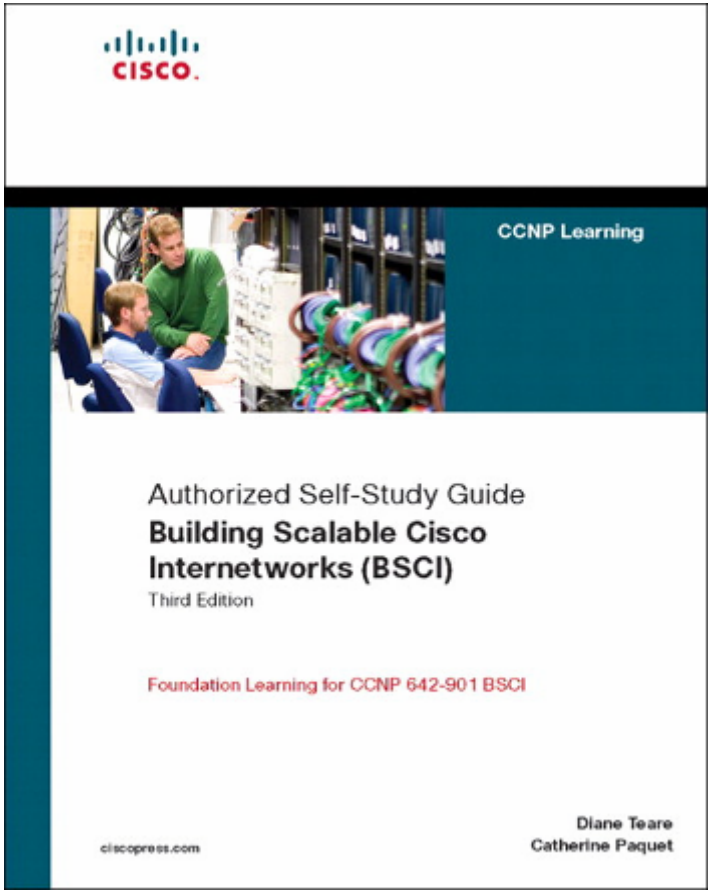

Available Onsite at the Cisco Company Store

Recommended Reading

- Continue your Cisco Live learning experience with further reading from Cisco Press
- Check the Recommended Reading flyer for suggested books

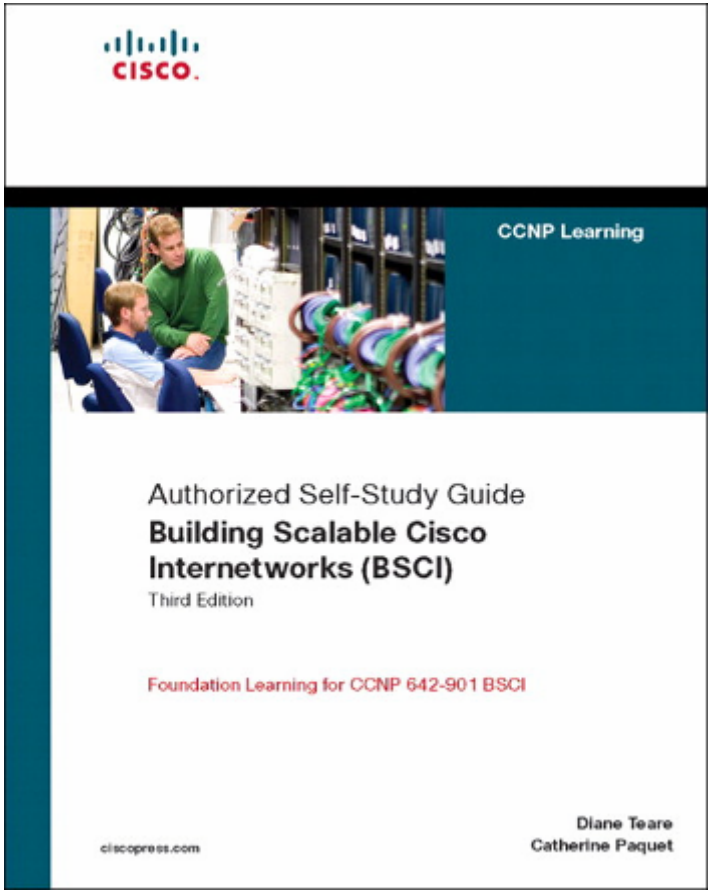

Available Onsite at the Cisco Company Store

Complete Your Online Session Evaluation

- Cisco values your input
- Give us your feedback—we read and carefully consider your scores and comments, and incorporate them into the content program year after year
- Go to the Internet stations located throughout the Convention Center to complete your session evaluations
- Thank you!