Chapter 6: Quality of Service

Connecting Networks

Cisco | Networking Academy®
Mind Wide Open™

6.1 QoS Overview

Cisco | Networking Academy®
Mind Wide Open™

Network Transmission Quality

Prioritizing Traffic

- Queuing packets causes delay -> new packets cannot be transmitted
- Number of packets increases -> memory fills up, packets are dropped.
- Congestions occurs when multiple links aggregate into a single device.

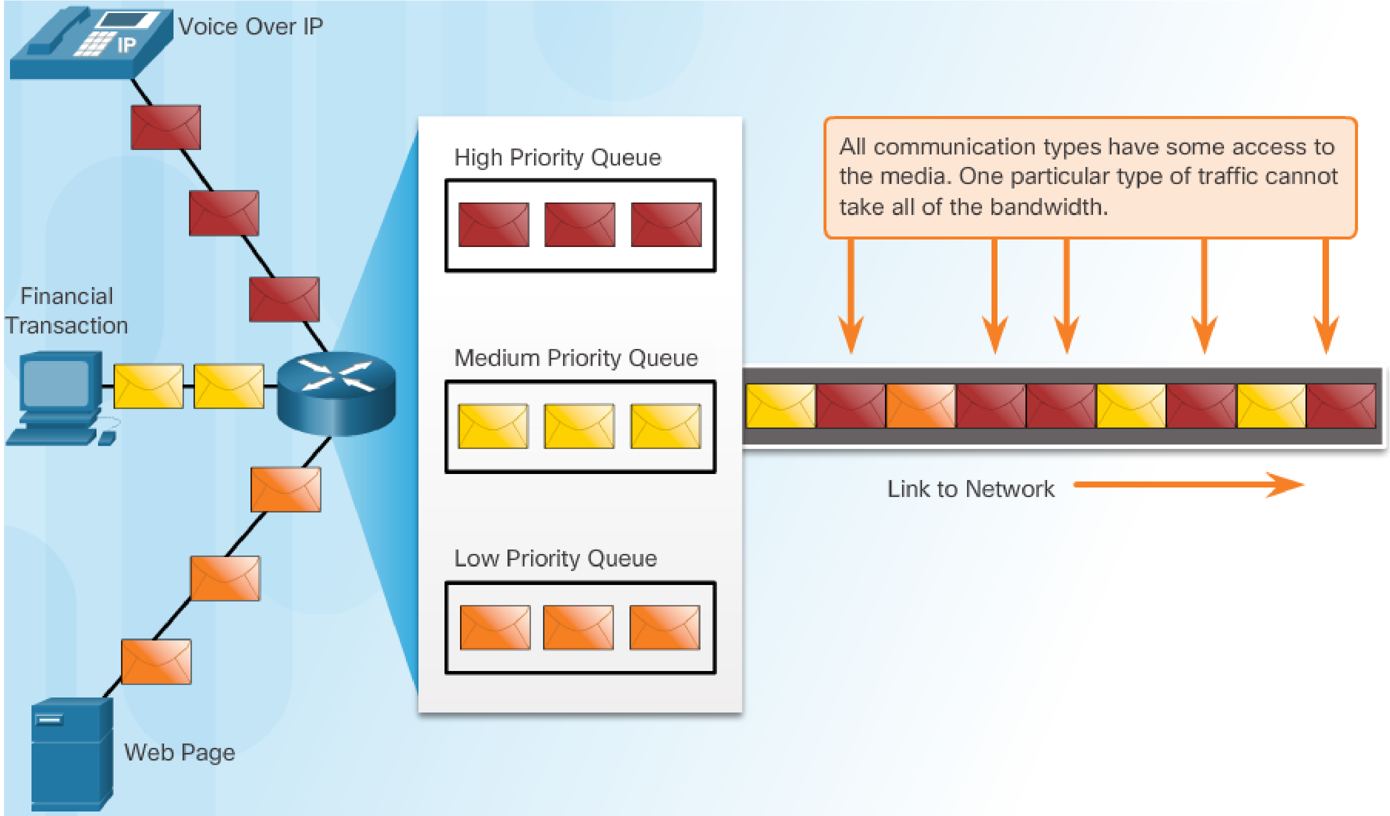

Typical Congestion Points

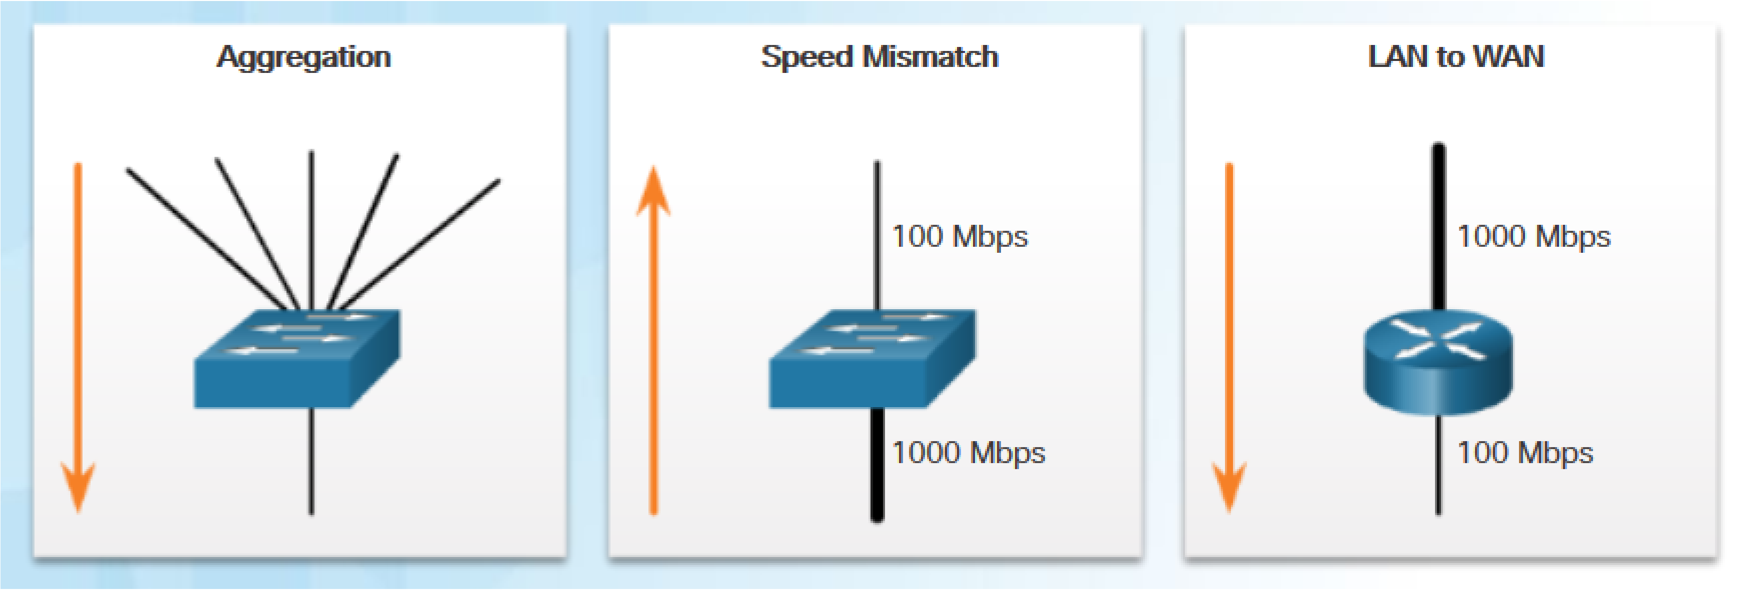

Delays

- Network congestion causes delay.

Delay (latency)

- Delay is the time it takes for a packet to travel from the source to the destination.
- Fixed delay (transmission, packetization, coding)
- Variable delay (queuing time, propagation)

Jitter is the variation in the delay of received packets.

Delay	Description
Code delay	The fixed amount of time it takes to compress data at the source before transmitting to the first internetworking device, usually a switch.
Packetization delay	The fixed time it takes to encapsulate a packet with all the necessary header information.
Queuing delay	The variable amount of time a frame or packet waits to be transmitted on the link.
Serialization delay	The fixed amount of time it takes to transmit a frame onto the wire.
Propagation delay	The variable amount of time it takes for the frame to travel between the source and destination.
De-jitter delay	The fixed amount of time it takes to buffer a flow of packets and then send them out in evenly spaced intervals.

Packet Loss

- When congestion occurs, network devices can drop packets.
- Packet loss is a common cause of voice quality problems on an IP network.
- In a properly designed network, packet loss should be near zero.
- QoS mechanisms is used to classify voice packets for zero packet loss.

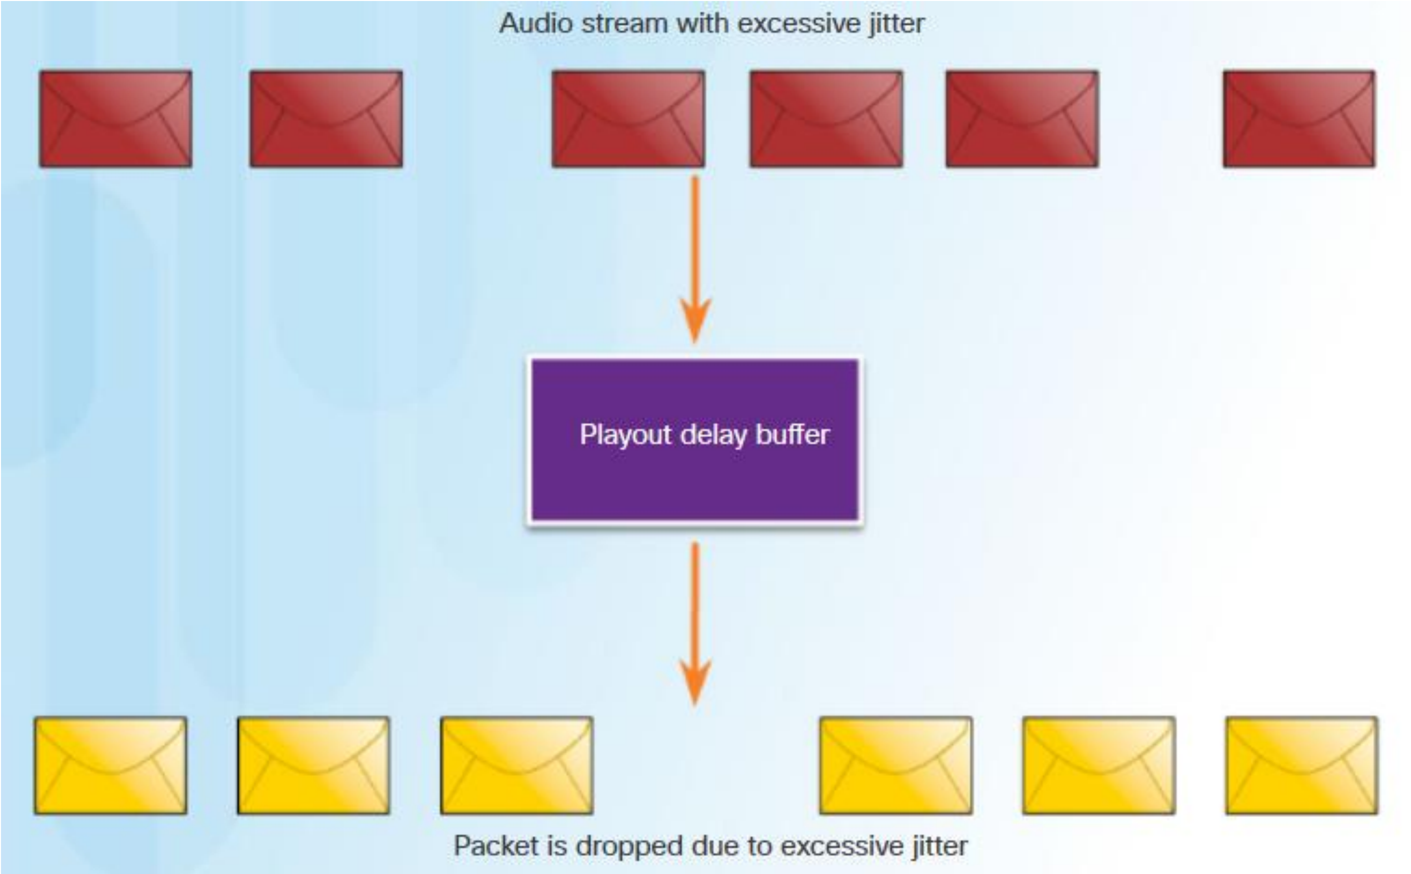

Network Traffic Trends

Voice

- Demands on voice, video, and data traffic are very different.
- Voice is very sensitive to delays and dropped packets:
  - Packets are not retransmitted if they are lost.
  - Must receive a higher priority than other types of traffic.
  - Voice can tolerate a certain amount of latency, jitter, and loss without any noticeable effects.

Voice

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

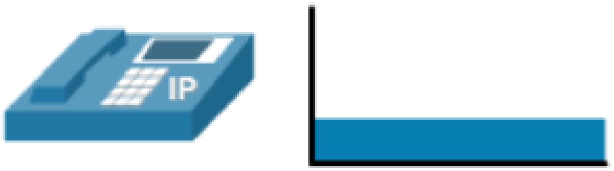

One-Way Requirements

- Latency ≤ 150 ms
- Jitter ≤ 30 ms
- Loss ≤ 1%
- Bandwidth (30 - 128 Kb/s)

Network Traffic Trends

Video

- Compared to voice, video is less resilient to loss and has a higher volume of data per packet.
- Video can tolerate a certain amount of latency, jitter, and loss without any noticeable affects.

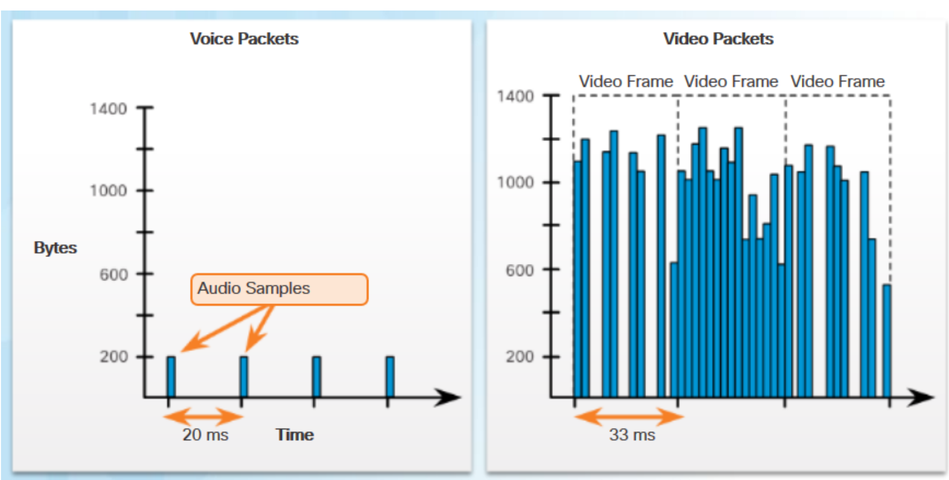

Video

- Bursty
- Greedy
- Drop sensitive
- Delay sensitive
- UDP priority

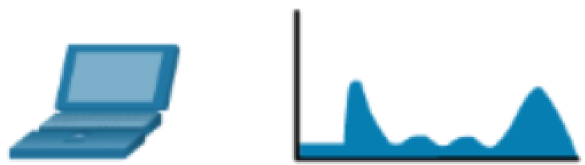

One-Way Requirements

- Latency ≤ 200-400 ms
- Jitter ≤ 30-50 ms
- Loss ≤ 0.1-1%
- Bandwidth (384 Kb/s - 20+ Mb/s)

Network Traffic Trends

Data

- Data applications that have no tolerance for data loss, such as email and web pages, use TCP that ensures that lost packets are resent.
- Data traffic is relatively insensitive to drops and delays compared to voice and video.

Data

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

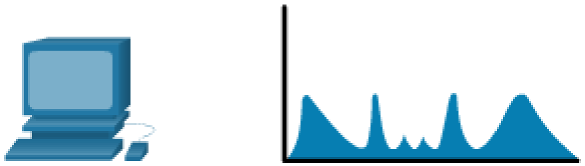

Factor	Mission Critical	Not Mission Critical
Interactive	Prioritize for the lowest delay of all data traffic and strive for a 1 to 2 seconds response time.	Applications could benefit from lower delay.
Not interactive	Delay can vary greatly as long as the necessary minimum bandwidth is supplied.	Gets any leftover bandwidth after all voice, video, and other data application needs are met.

Queueing Algorithms

First In First Out (FIFO)

- No concept of priority or classes of traffic.
- The fastest method of queuing.
- Effective for large links that have little delay and minimal congestion.
- Default method on network interfaces (except serial)

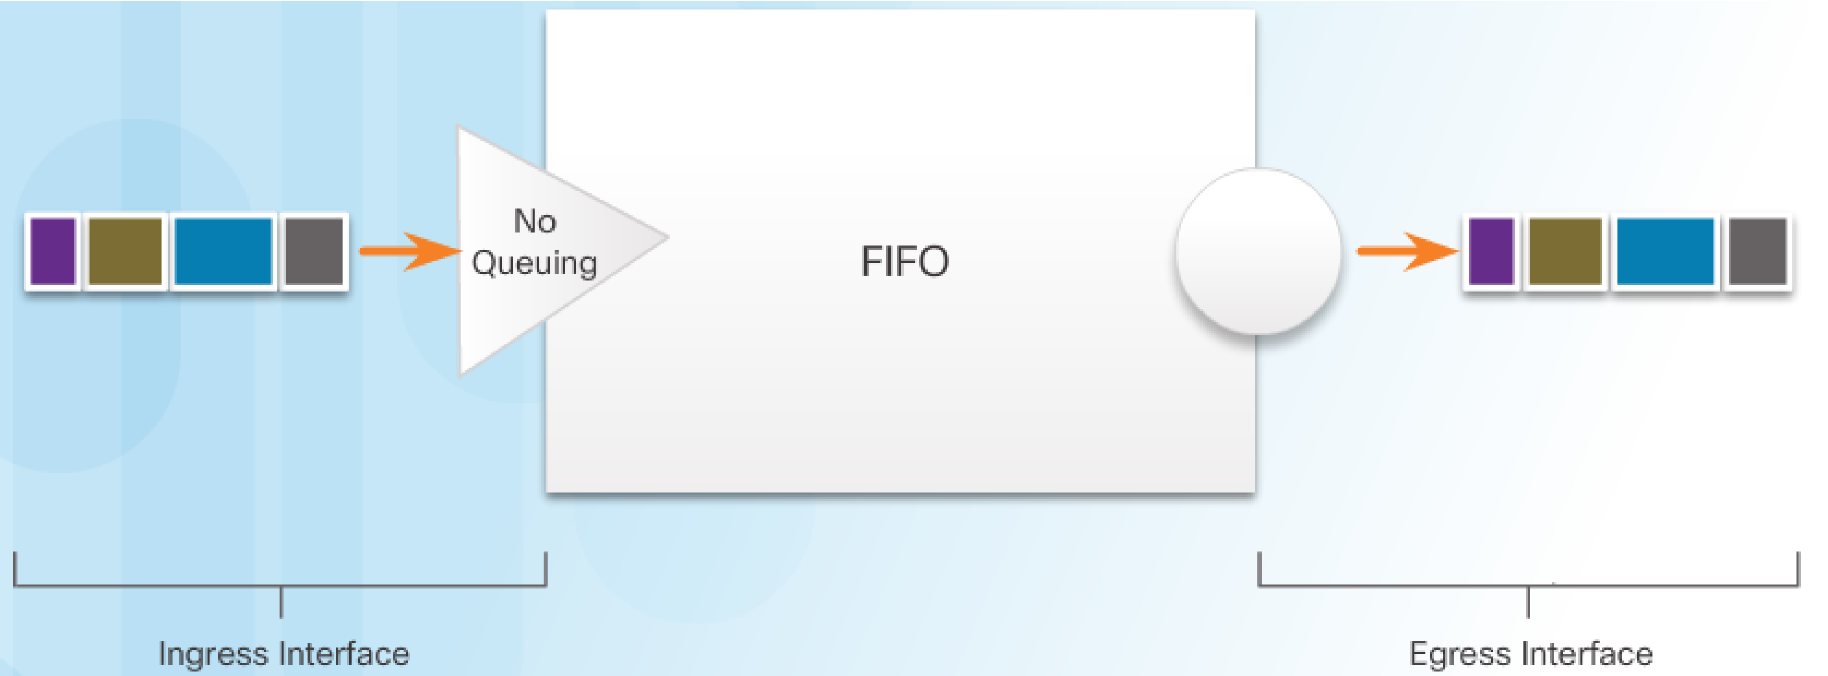

Queueing Algorithms

Weighted Fair Queuing (WFQ)

- An automated scheduling method that provides fair bandwidth allocation.
- Applies priority/weights to identified traffic and classifies it into flows.
  - Identification based on IP/MAC addresses, protocol, port numbers
  - Classification sets the ToS value
- Not supported with tunneling and encryption -> they modify the packet.

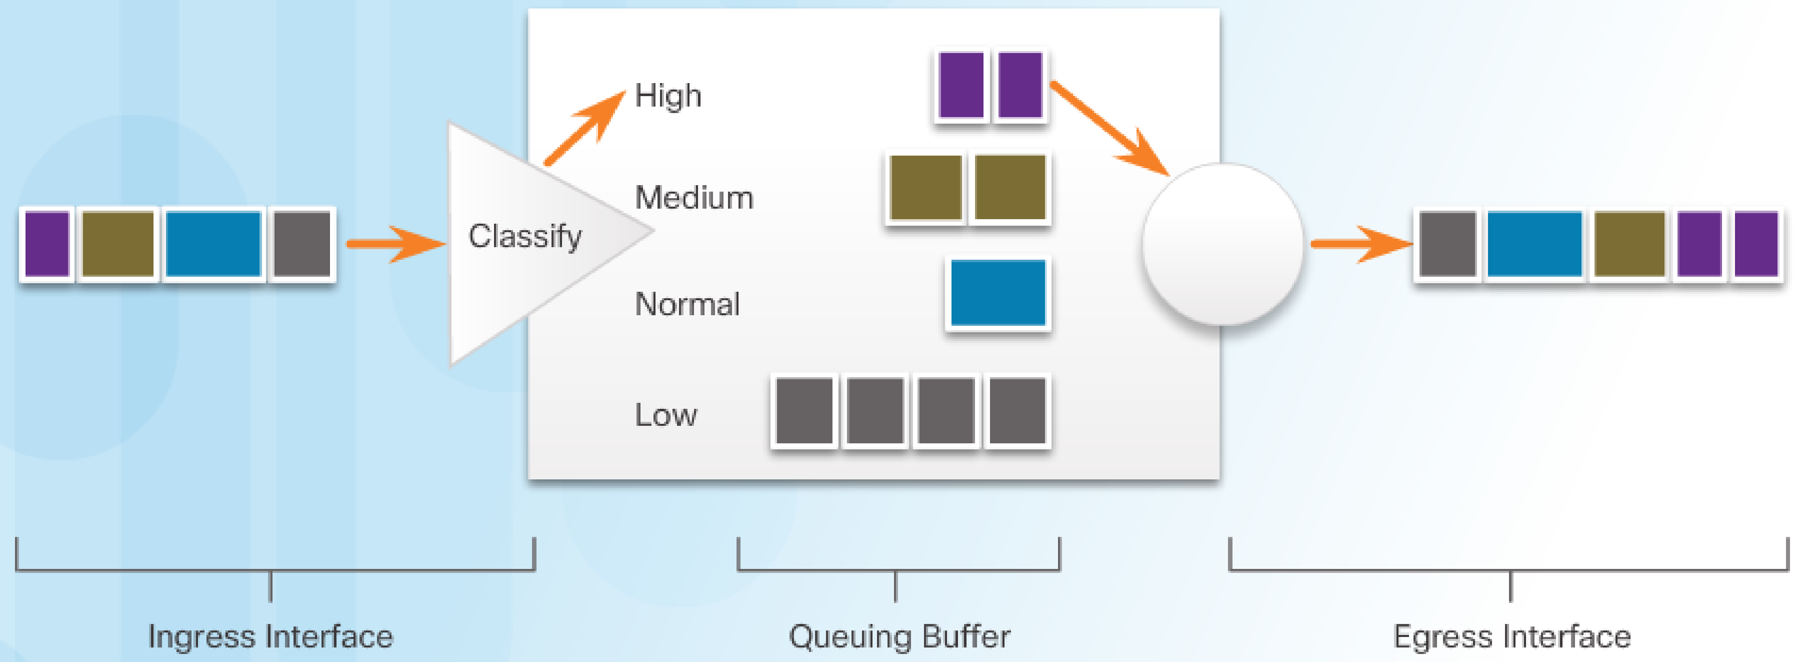

Queueing Algorithms

Class-Based Weighted Fair Queuing (CBWFQ)

- Extends the standard WFQ functionality -> user-defined traffic classes.
- Classes based on match criteria: protocols, ACLs, interfaces.
- A class is assigned bandwidth, weight, and maximum packet limit.
- The queue limit sets the maximum no. of packets allowed in the queue.

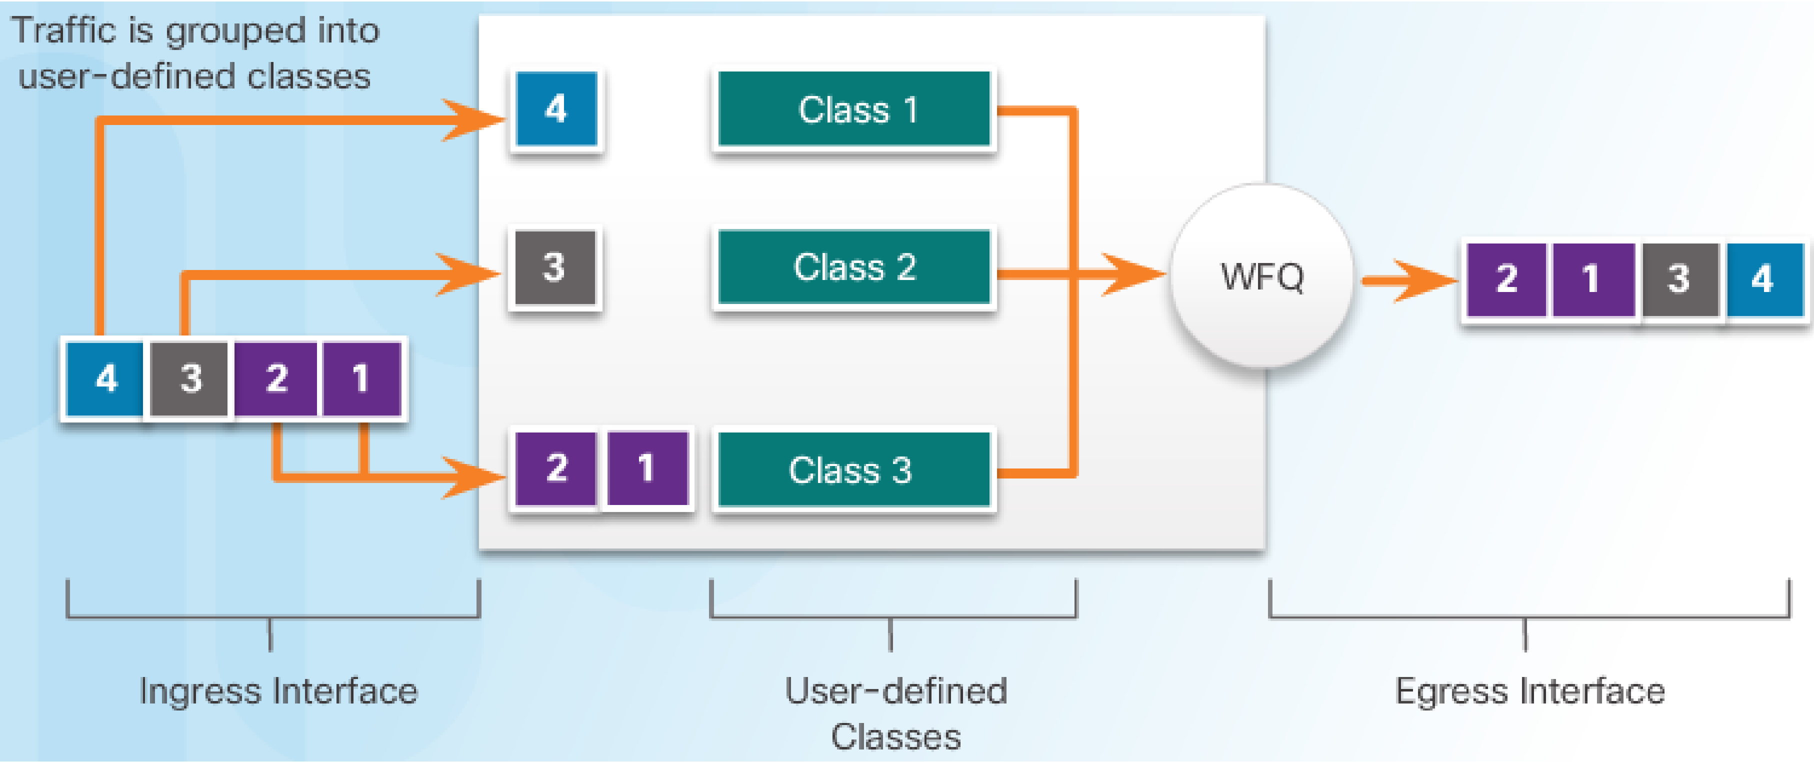

Queueing Algorithms

Low Latency Queuing (LLQ)

- Adds strict priority queuing (PQ) for CBWFQ: delay-sensitive data first.
- Without LLQ, all packets are serviced fairly based on weight only.
- With LLQ, delay-sensitive data such as voice is sent first.

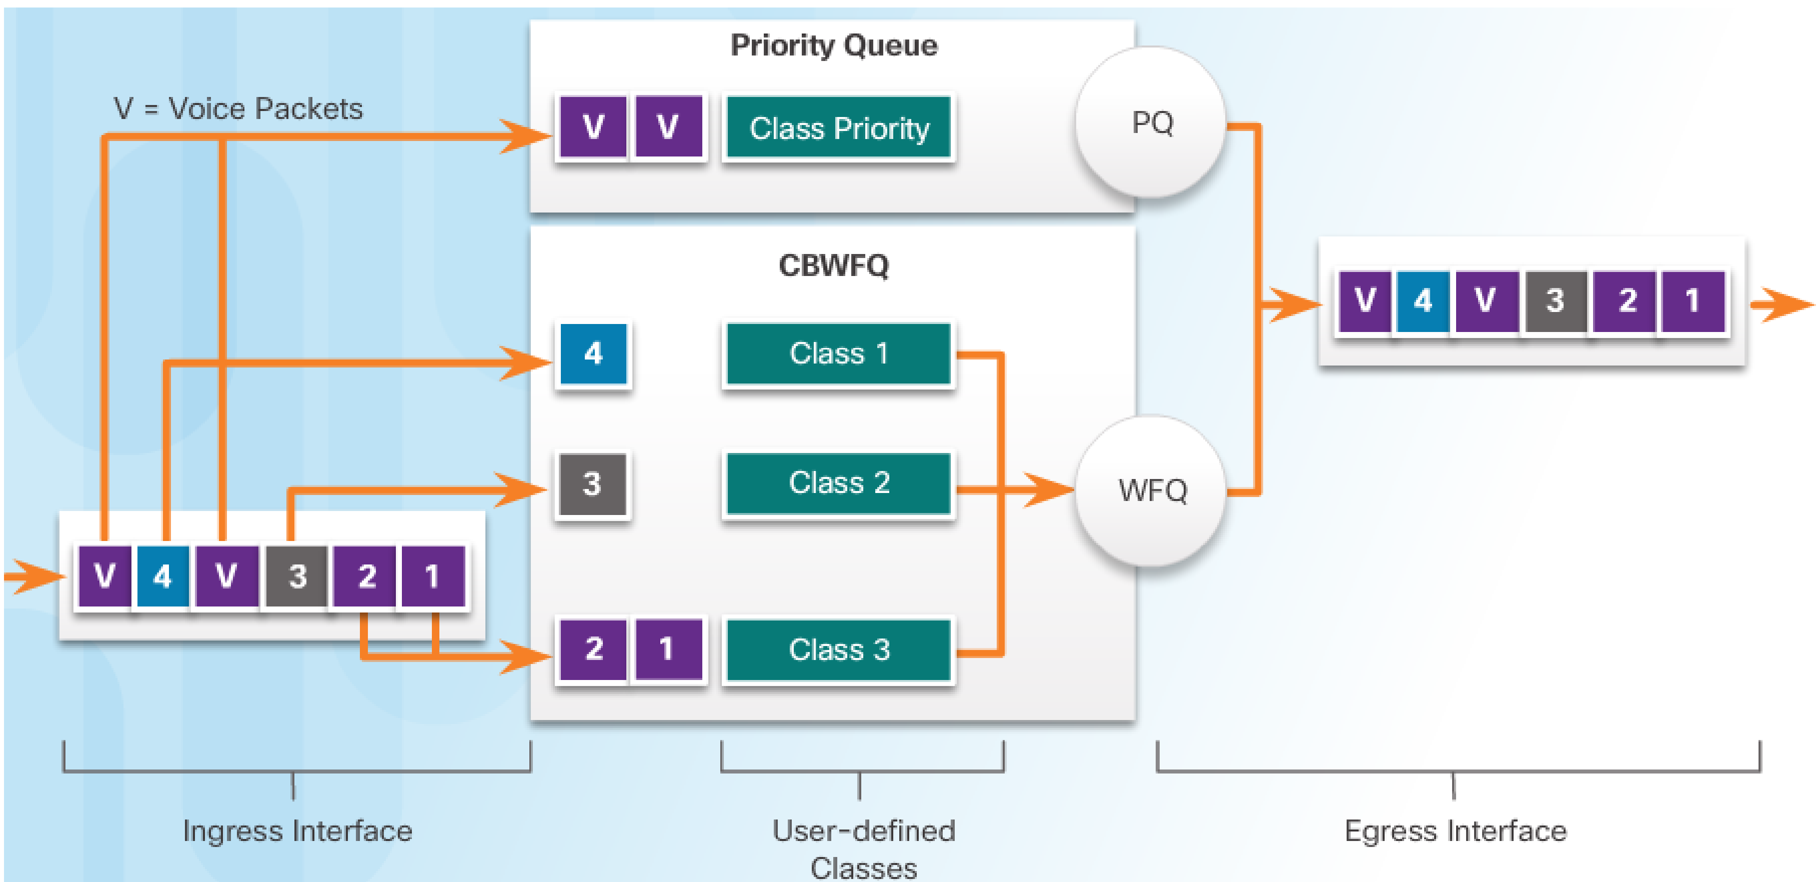

6.2 QoS Mechanisms

Cisco | Networking Academy®
Mind Wide Open™

Models for Implementing QoS

Three models for implementing QoS policy

Model	Description
Best-effort model	• Not really an implementation as QoS is not explicitly configured. • Use when QoS is not required.
Integrated services (IntServ)	• Provides very high QoS to IP packets with guaranteed delivery. • It defines a signaling process for applications to signal to the network that they require special QoS for a period and that bandwidth should be reserved. • However, IntServ can severely limit the scalability of a network.
Differentiated services (DiffServ)	• Provides high scalability and flexibility in implementing QoS. • Network devices recognize traffic classes and provide different levels of QoS to different traffic classes.

Best-Effort Model

- No guarantee for packet delivery.
- Default for IP networks.
- All packets treated in the same way.

Benefits	Drawbacks
The model is the most scalable.	There are no guarantees of delivery.
Scalability is only limited by bandwidth limits, in which case all traffic is equally affected.	Packets will arrive whenever they can and in any order possible, if they arrive at all.
No special QoS mechanisms are required.	No packets have preferential treatment.
It is the easiest and quickest model to deploy.	Critical data is treated the same as casual email is treated.

Integrated Services (IntServ)

- Provides resource reservation and admission-control to establish and maintain QoS.
- The edge router performs admission control to ensure that available resources are sufficient in the network.
- Hard QoS approach: guarantees traffic characteristics such as bandwidth, delay, packet-loss rates between end points.

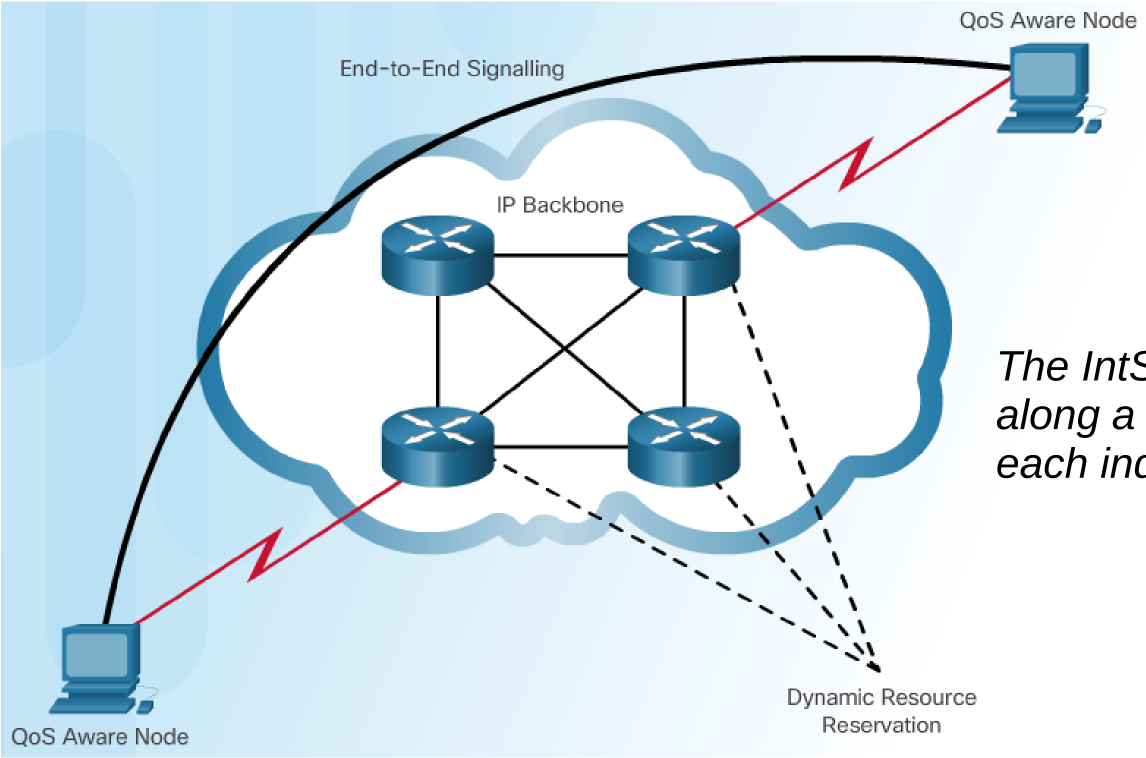

The IntServ standard assumes that routers along a path set and maintain the state for each individual communication

Differentiated Services (DiffServ)

- DiffServ divides network traffic into classes based on business requirements.
- Each of the classes can then be assigned a different level of service.
- Network elements sets multiple classes of traffic for QoS requirements.
- Soft QoS approach: not an end-to-end QoS strategy.

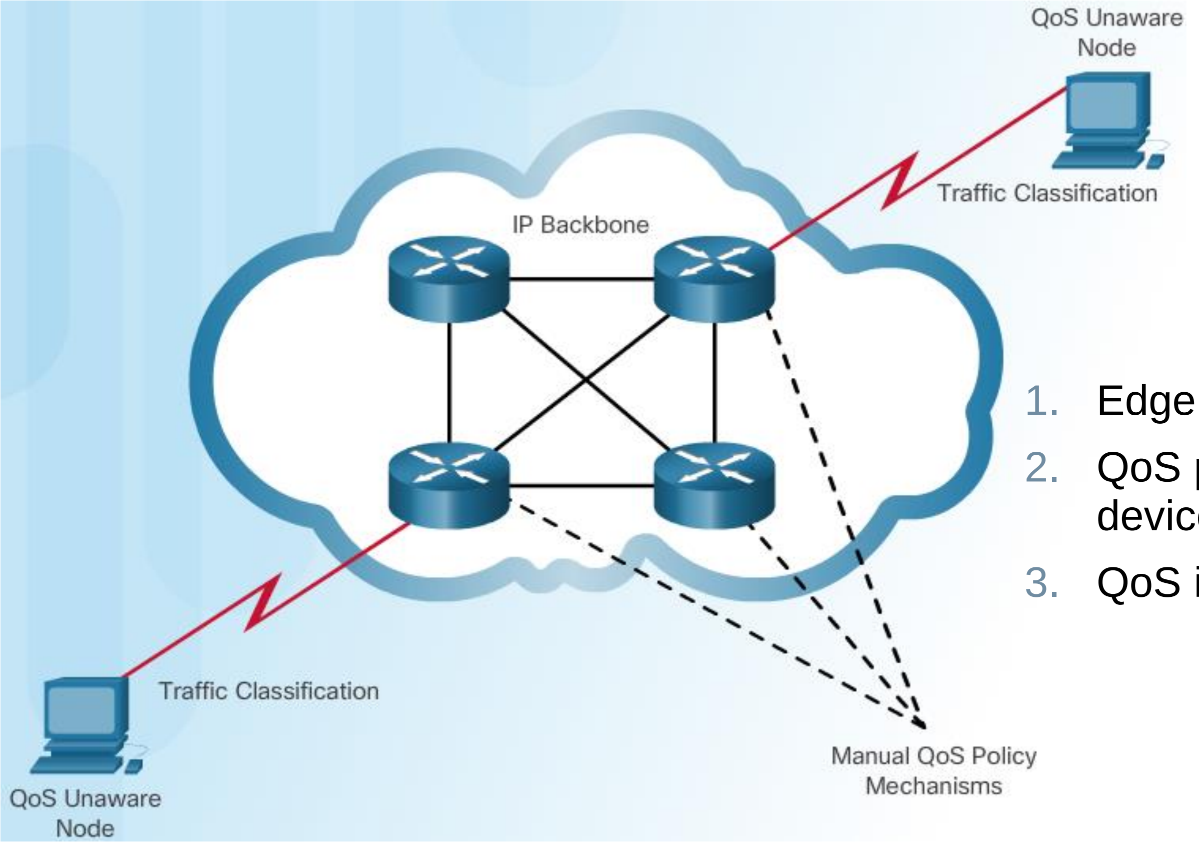

1. Edge router classifies a packet.
2. QoS policy is configured on all network devices.
3. QoS is enforced on a hop-by-hop basis.

Comparison

Integrated Services

Benefits	Drawbacks
• Explicit end-to-end resource admission control • Per-request policy admission control • Signaling of dynamic port numbers	• Resource intensive due to the stateful architecture requirement for continuous signaling. • Flow-based approach not scalable to large implementations such as the Internet.

Differentiated Services

Benefits	Drawbacks
• Highly scalable • Provides many different levels of quality	• No absolute guarantee of service quality • Requires a set of complex mechanisms to work in concert throughout the network

Avoiding Packet Loss

- Packet loss is a result of congestion on a interface.
- Dropped TCP segments cause TCP sessions to reduce their window sizes.
- Some applications do not use TCP and cannot handle drops.
- QoS Sequence

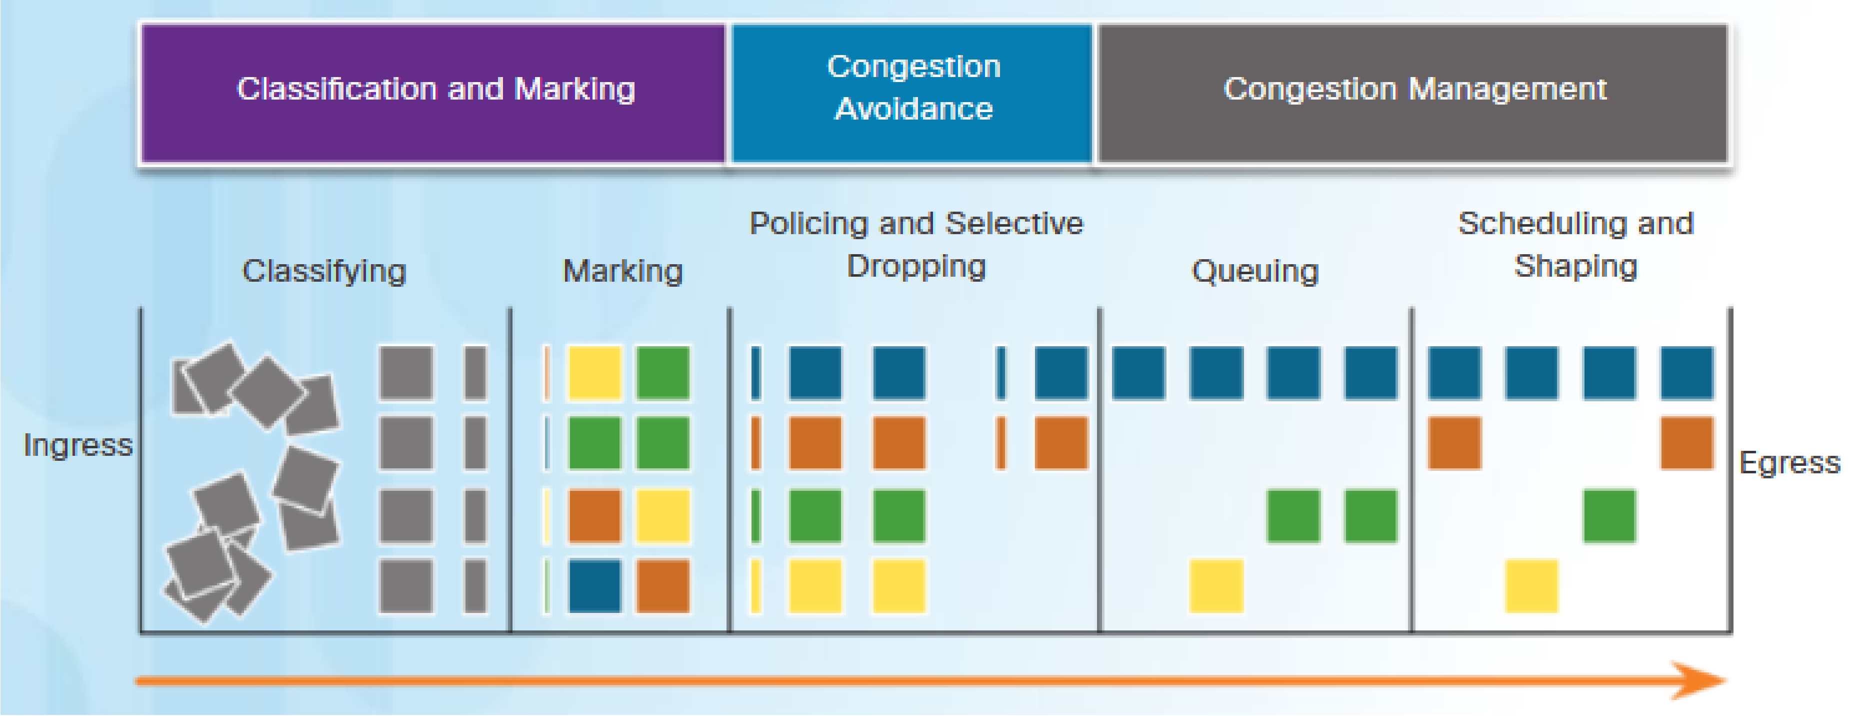

Classification and Marking

Before QoS policy is applied, packet has to be classified.

- ACLs, class maps, Network Based Application Recognition (NBAR)

Marking = adding a value to the packet header

- Marking on L2 or L3
- Depends on transmission technology

QoS Tools	Layer	Marking Field	Width in Bits
Ethernet (802.1Q, 802.1p)	2	Class of Service (CoS)	3
802.11 (Wi-Fi)	2	Wi-Fi Traffic Identifier (TID)	3
MPLS	2	Experimental (EXP)	3
IPv4 and IPv6	3	IP Precedence (IPP)	3
IPv4 and IPv6	3	Differentiated Services Code Point (DSCP)	6

Classification and Marking: Layer 2

Standards IEEE 802.1Q and 802.1P

- Three bits for the Class of Service (CoS)

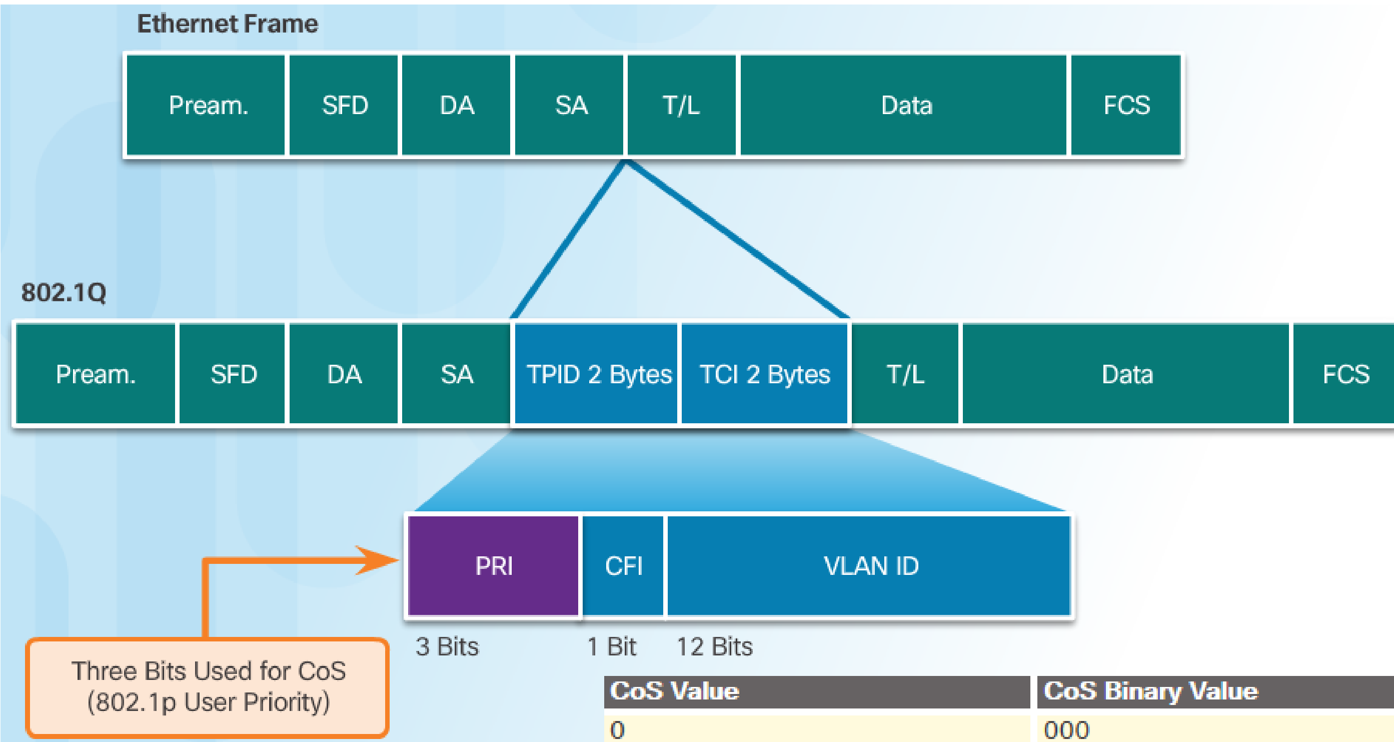

CoS Value	CoS Binary Value	Description
0	000	Best-Effort Data
1	001	Medium-Priority Data
2	010	High-Priority Data
3	011	Call Signaling
4	100	Videoconferencing
5	101	Voice bearer (voice traffic)
6	110	Reserved
7	111	Reserved

Classification and Marking: Layer 3

IPv4: Type of Service (ToS)

- 3-bits of IP precedence or 6 bits for DSCP code

IPv6: Traffic Class

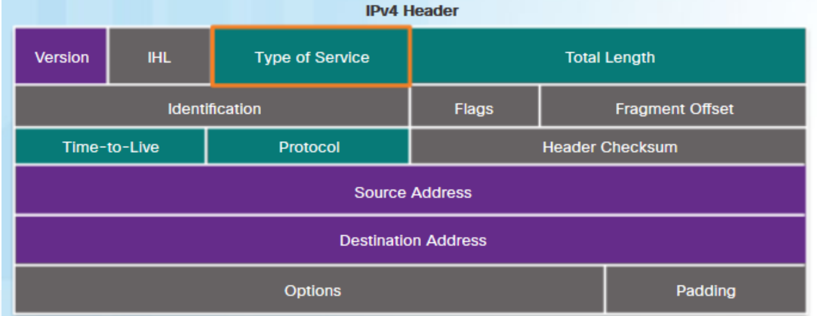

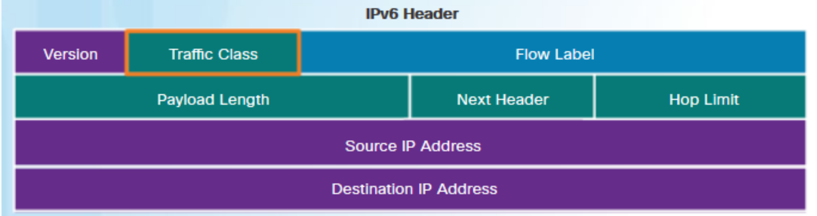

Best Queue

Worst Queue

	Low Drop	Medium Drop	High Drop
Class 4	AF41 (34)	AF42 (36)	AF43 (38)
Class 3	AF31 (26)	AF32 (28)	AF33 (30)
Class 2	AF21 (18)	AF22 (20)	AF23 (22)
Class 1	AF11 (10)	AF12 (12)	AF13 (14)

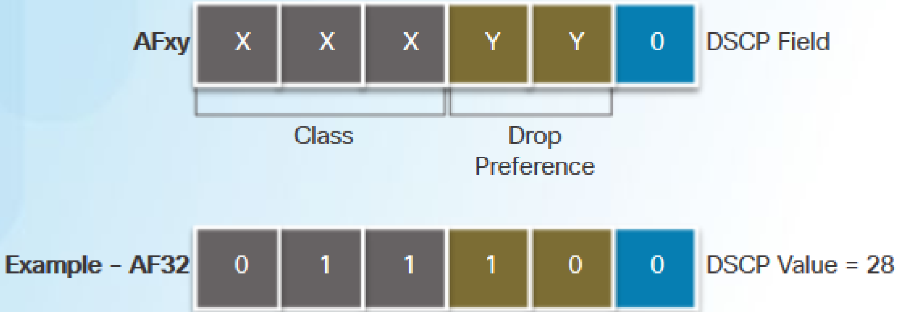

Mapping L2 Marking to L3 Marking

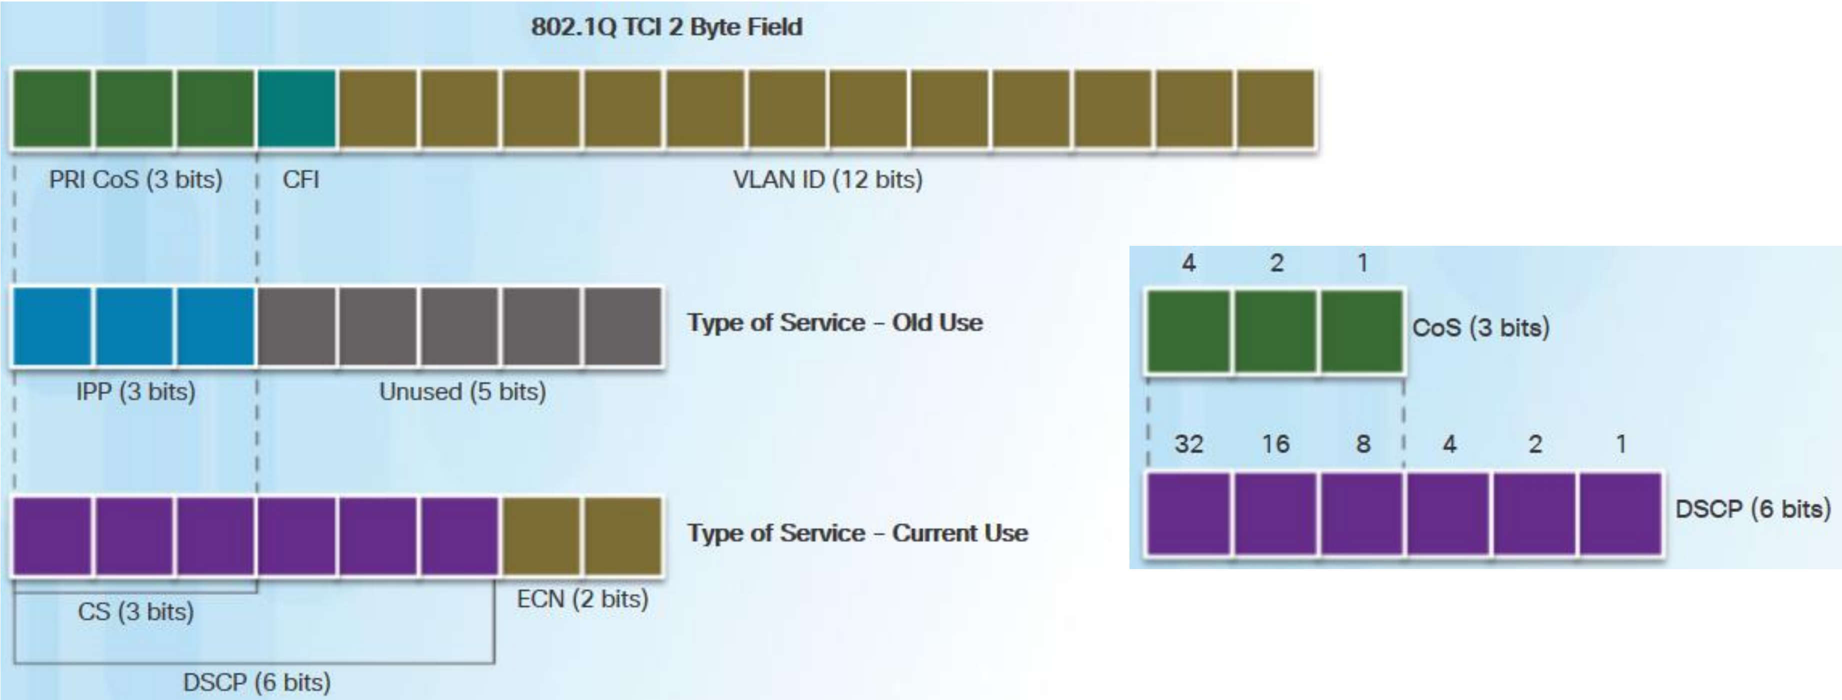

CoS Value	CoS Binary Value	Class Selector (CS)	CS Binary	DSCP Decimal Value
0	000	CS0*/DF	000 000	0
1	001	CS1	001 000	8
2	010	CS2	010 000	16
3	011	CS3	011 000	24
4	100	CS4	100 000	32
5	101	CS5	101 000	40
6	110	CS6	110 000	48
7	111	CS7	111 000	56

Trust Boundaries

- Traffic should be classified and marked as close to its source as technically and administratively feasible.
  - Trusted endpoints mark application traffic using L2 and/or L3 values.

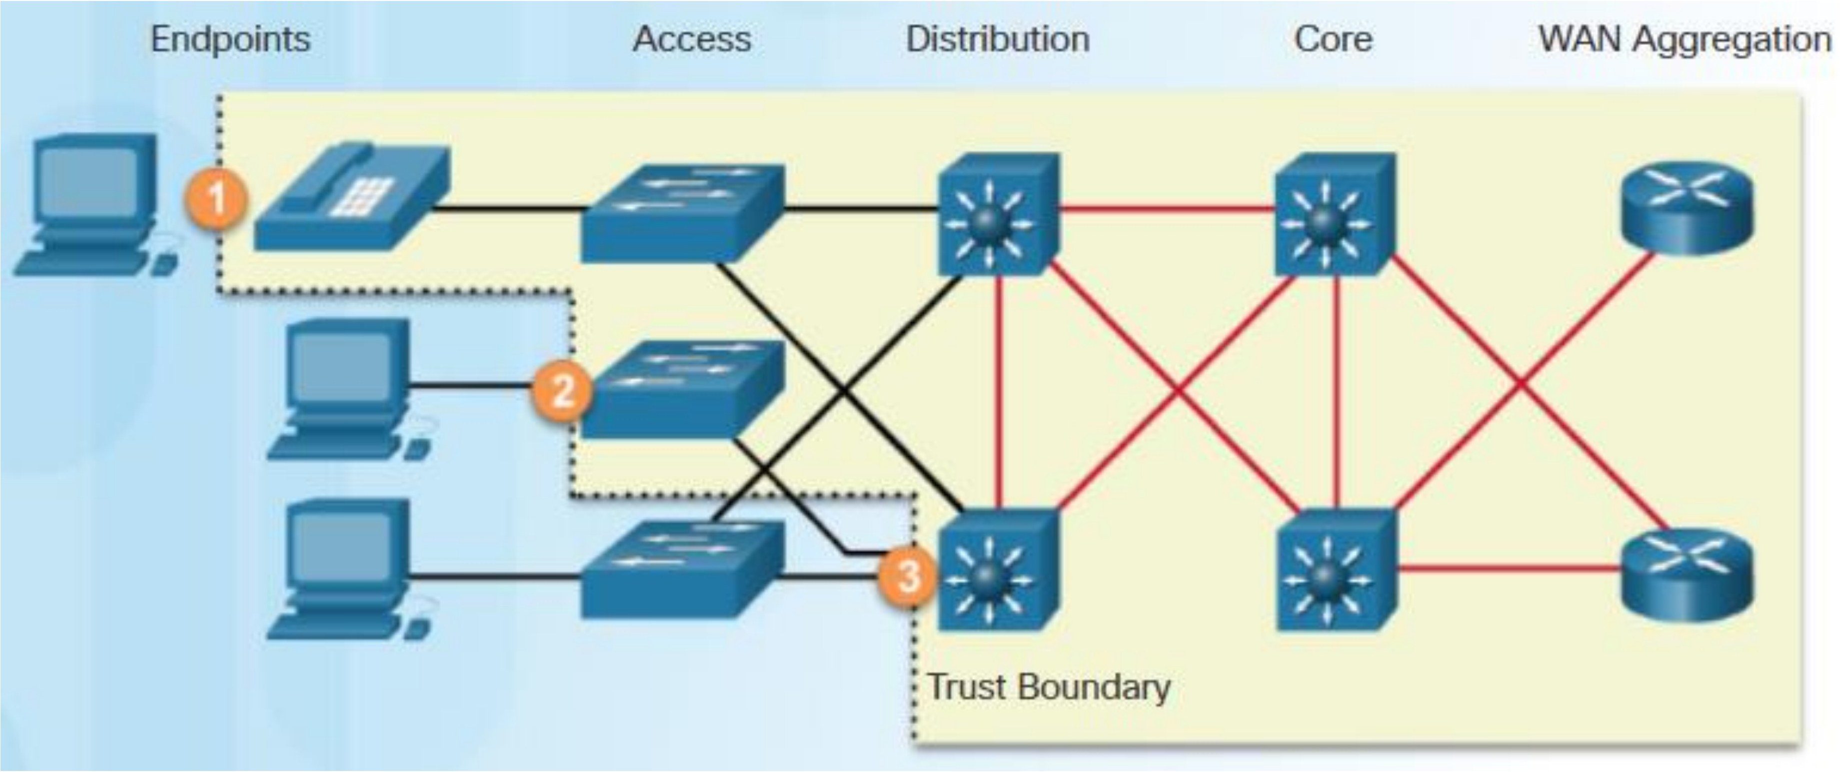

Congestion Avoidance

- Congestion management includes queuing and scheduling.
- Congestion avoidance tools:
  - Monitor network traffic load
  - Depending on the current queue size, packet starts to be dropped.
  - Weighted Random Early Detection (WRED) algorithm is applied.

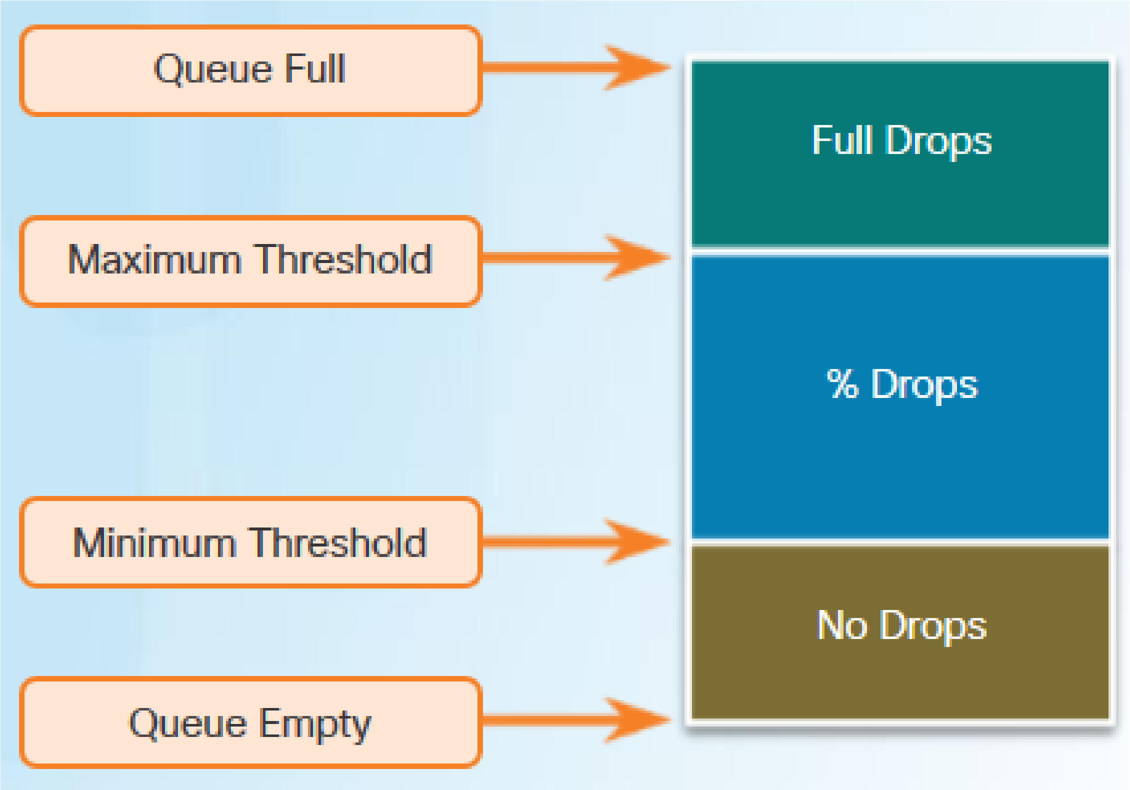

Shaping and Policing

- Traffic shaping retains excess packets in a queue and then schedules the excess for later transmission over increments of time.

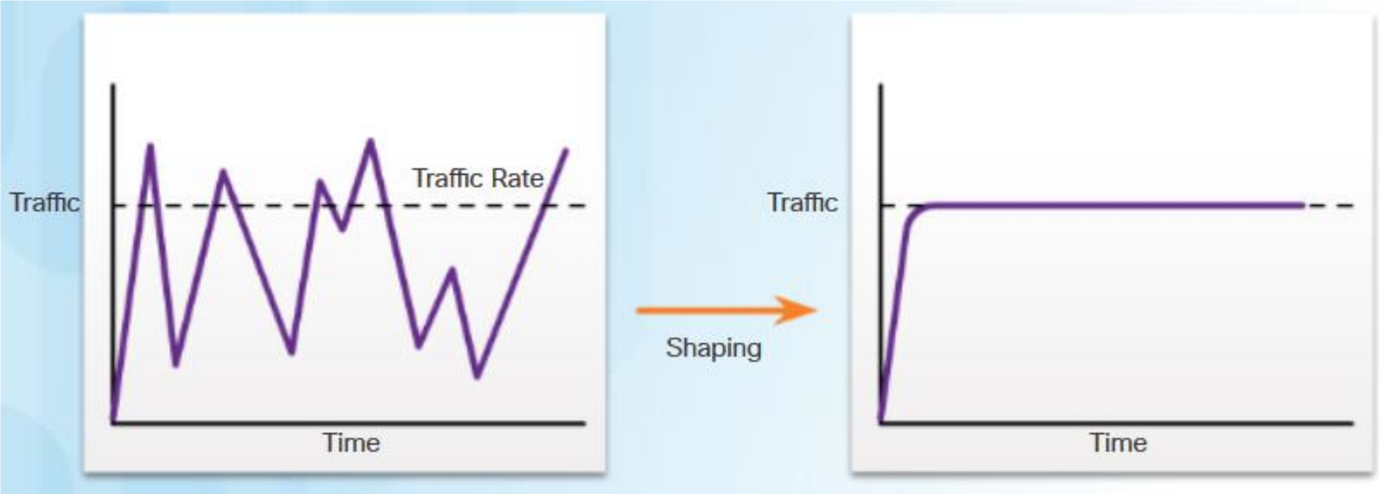

- Policing is applied to inbound traffic on an interface. When the traffic rate reaches the configured maximum rate, excess traffic is dropped (or remarked).