

# CCNA Day 47

## Quality of Service (Part 2)

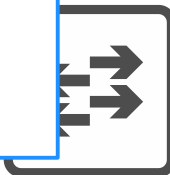


### 4.0 IP Services

10%



- 4.1 Configure and verify inside source NAT using static and pools
- 4.2 Configure and verify NTP operating in a client and server mode
- 4.3 Explain the role of DHCP and DNS within the network
- 4.4 Explain the function of SNMP in network operations
- 4.5 Describe the use of syslog features including facilities and levels
- 4.6 Configure and verify DHCP client and relay
- 4.7 Explain the forwarding per-hop behavior (PHB) for QoS such as classification, marking, queuing, congestion, policing, shaping
- 4.8 Configure network devices for remote access using SSH
- 4.9 Describe the capabilities and function of TFTP/FTP in the network



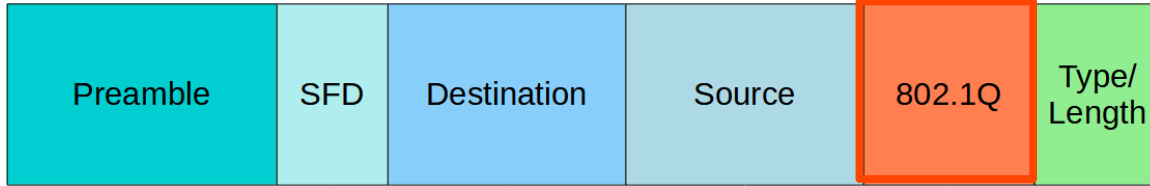
# Things we'll cover

- Classification/Marking
- Queuing/Congestion Management
- Shaping/Policing

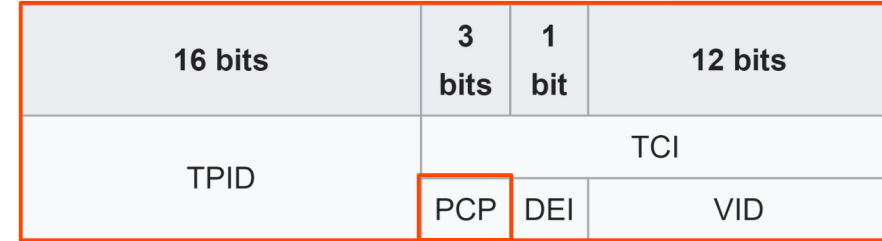
# Classification

- The purpose of QoS is to give certain kinds of network traffic priority over others during congestion.
- **Classification** organizes network traffic (packets) into traffic classes (categories).
- Classification is fundamental to QoS. To give priority to certain types of traffic, you have to identify which types of traffic to give priority to.
- There are many methods of classifying traffic. Some examples:
  - An ACL. Traffic which is permitted by the ACL will be given certain treatment, other traffic will not.
  - **NBAR** (Network Based Application Recognition) performs a *deep packet inspection*, looking beyond the Layer 3 and Layer 4 information up to Layer 7 to identify the specific kind of traffic.
  - In the Layer 2 and Layer 3 headers there are specific fields used for this purpose.
- The **PCP** (Priority Code Point) field of the 802.1Q tag (in the Ethernet header) can be used to identify high/low priority traffic.
  - Only when there is a dot1q tag!
- The **DSCP** (Differentiated Services Code Point) field of the IP header can also be used to identify high/low priority traffic.

## Ethernet Header



## 802.1Q tag format



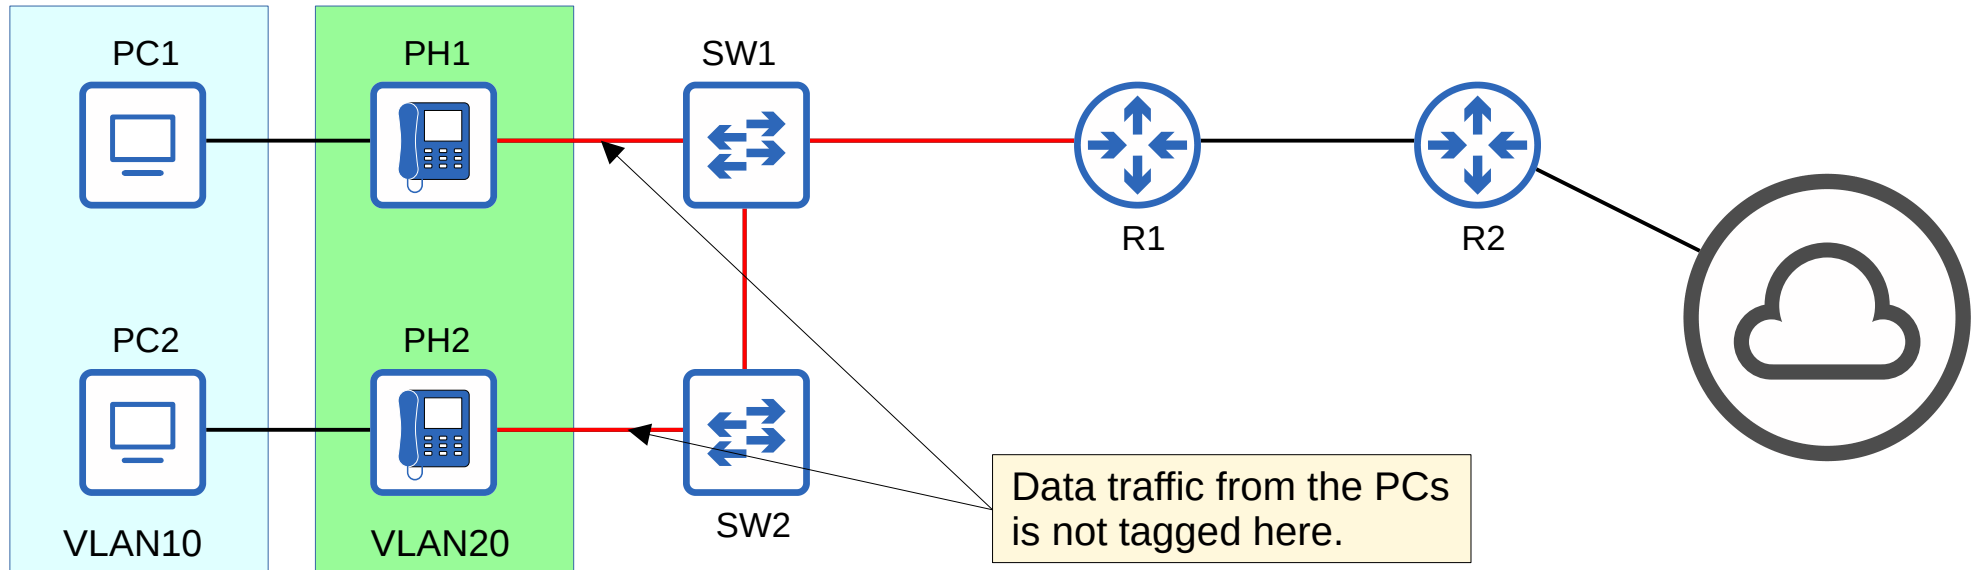
- PCP is also known as CoS (Class of Service). Its use is defined by IEEE 802.1p.
- 3 bits = 8 possible values ( $2^3 = 8$ ).

'Best effort' delivery means there is no guarantee that data is delivered or that it meets any QoS standard. This is regular traffic, not high-priority.

IP phones **mark** call signaling traffic (used to establish calls) as PCP3. They **mark** the actual voice traffic as PCP5.

PCP value	Traffic types
0	Best effort (default)
1	Background
2	Excellent effort
3	Critical applications
4	Video
5	Voice
6	Internetwork control
7	Network control

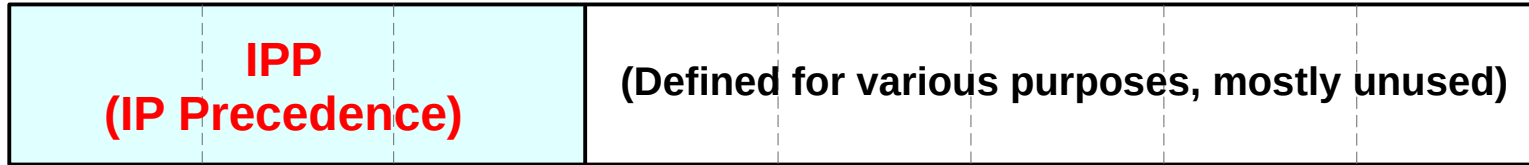
- Because PCP is found in the dot1q header, it can only be used over the following connections:
  - trunk links
  - access links with a voice VLAN
- In the diagram below, traffic between R1 and R2, or between R2 and external destinations will not have a dot1q tag. So, traffic over those links PCP cannot be marked with a PCP value.



# The IP ToS Byte

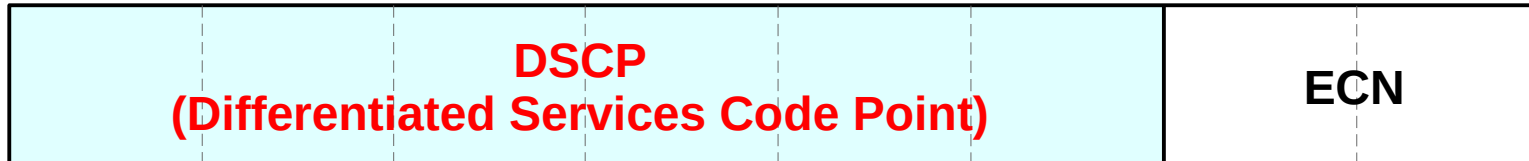
Offsets	Octet	0								1								2								3											
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	Version				IHL				DSCP						ECN		Total Length																			
4	32	Identification																Flags				Fragment Offset															
8	64	Time To Live								Protocol								Header Checksum																			
12	96	Source IP Address																																			
16	128	Destination IP Address																																			
20	160	Options (if IHL > 5)																																			
24	192																																				
28	224																																				
32	256																																				

ToS byte (old)



3 bits = 8 values (0-7)

ToS byte (current)



6 bits = 64 values (0-63)

# IP Precedence

IPP (IP Precedence)	(Defined for various purposes, mostly unused)	3 bits = 8 values (0-7)
------------------------	---	-------------------------

- Standard IPP markings are similar to PCP:
  - 6 and 7 are reserved for 'network control' traffic (ie. OSPF messages between routers)
  - 5 = voice
  - 4 = video
  - 3 = voice signaling
  - 0 = best effort
- With 6 and 7 reserved, 6 possible values remain.
- Although 6 values is sufficient for many networks, the QoS requirements of some networks demand more flexibility.

**DSCP**  
**(Differentiated Services Code Point)**

**ECN**

6 bits = 64 values (0-63)

- RFC 2474 (1998) defines the DSCP field, and other 'DiffServ' RFCs elaborate on its use.
- With IPP updated to DSCP, new standard markings had to be decided upon.
  - By having generally agreed upon standard markings for different kinds of traffic, QoS design & implementation is simplified, QoS works better between ISPs and enterprises, among other benefits.
- You should be aware of the following standard markings:
  - Default Forwarding (DF) – best effort traffic
  - Expedited Forwarding (EF) – low loss/latency/jitter traffic (usually voice)
  - Assured Forwarding (AF) – A set of 12 standard values
  - Class Selector (CS) – A set of 8 standard values, provides backward compatibility with IPP

```
R1(config)#class-map TEST
```

```
R1(config-cmap)#match dscp ?
```

```
<0-63> Differentiated services codepoint value
```

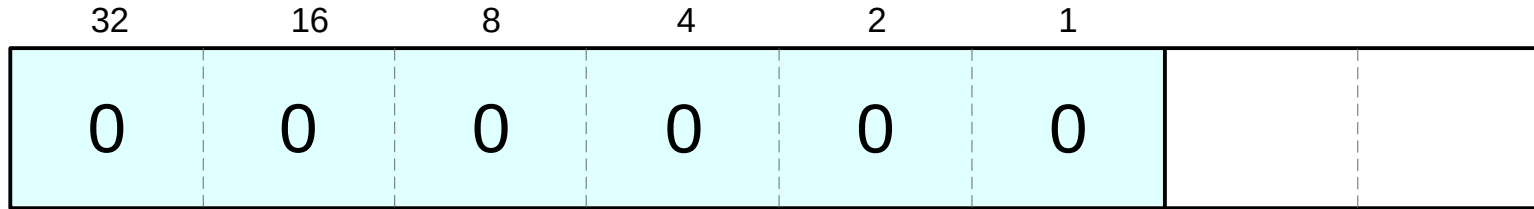
```
af11    Match packets with AF11 dscp (001010)
af12    Match packets with AF12 dscp (001100)
af13    Match packets with AF13 dscp (001110)
af21    Match packets with AF21 dscp (010010)
af22    Match packets with AF22 dscp (010100)
af23    Match packets with AF23 dscp (010110)
af31    Match packets with AF31 dscp (011010)
af32    Match packets with AF32 dscp (011100)
af33    Match packets with AF33 dscp (011110)
af41    Match packets with AF41 dscp (100010)
af42    Match packets with AF42 dscp (100100)
af43    Match packets with AF43 dscp (100110)
```

```
cs1     Match packets with CS1(precedence 1) dscp (001000)
cs2     Match packets with CS2(precedence 2) dscp (010000)
cs3     Match packets with CS3(precedence 3) dscp (011000)
cs4     Match packets with CS4(precedence 4) dscp (100000)
cs5     Match packets with CS5(precedence 5) dscp (101000)
cs6     Match packets with CS6(precedence 6) dscp (110000)
cs7     Match packets with CS7(precedence 7) dscp (111000)
```

```
default Match packets with default dscp (000000)
```

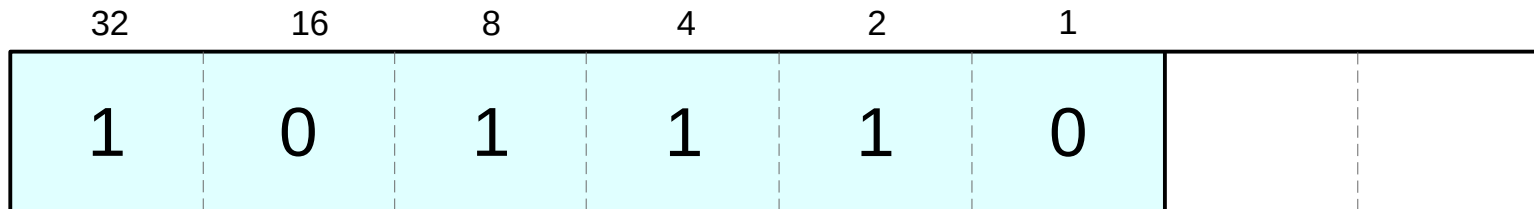
```
ef      Match packets with EF dscp (101110)
```

## DF (Default Forwarding):



- **DF** is used for best-effort traffic.
- The DSCP marking for DF is 0.

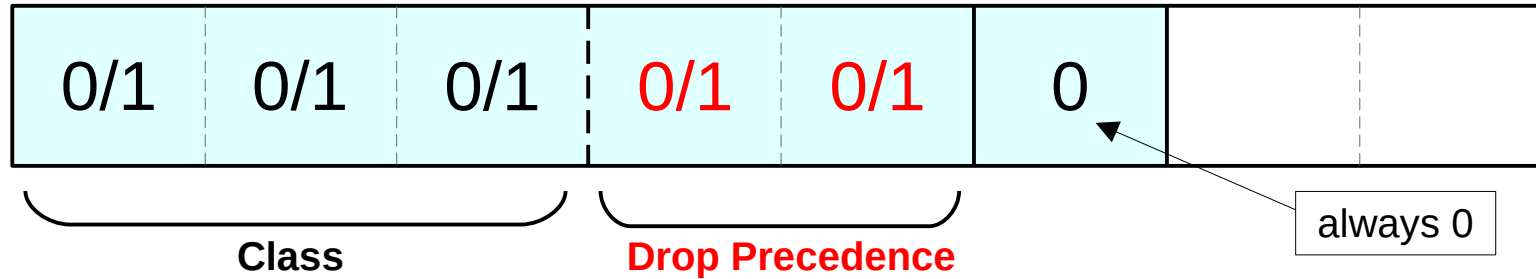
## EF (Expedited Forwarding):



- **EF** is used for traffic that requires low loss/latency/jitter.
- The DSCP marking for EF is 46.

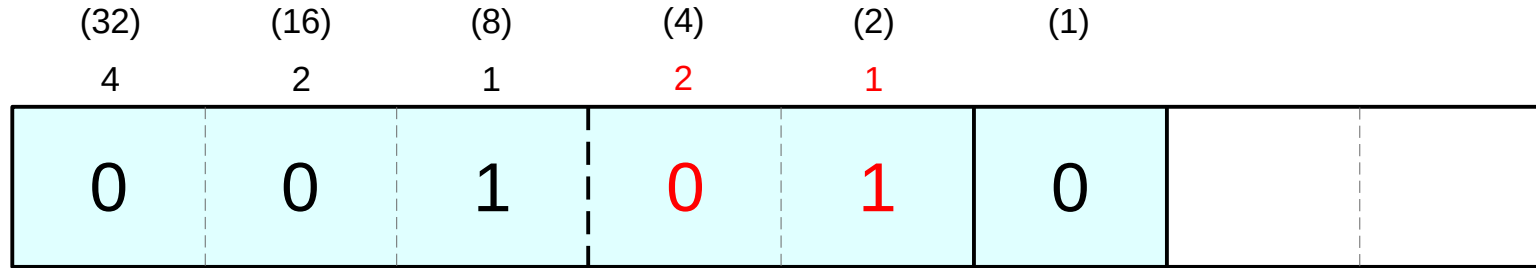
```
R1(config)#class-map TEST
R1(config-cmap)#match dscp ?
<0-63>    Differentiated services codepoint value
af11      Match packets with AF11 dscp (001010)
af12      Match packets with AF12 dscp (001100)
af13      Match packets with AF13 dscp (001110)
af21      Match packets with AF21 dscp (010010)
af22      Match packets with AF22 dscp (010100)
af23      Match packets with AF23 dscp (010110)
af31      Match packets with AF31 dscp (011010)
af32      Match packets with AF32 dscp (011100)
af33      Match packets with AF33 dscp (011110)
af41      Match packets with AF41 dscp (100010)
af42      Match packets with AF42 dscp (100100)
af43      Match packets with AF43 dscp (100110)
cs1       Match packets with CS1(precedence 1) dscp (001000)
cs2       Match packets with CS2(precedence 2) dscp (010000)
cs3       Match packets with CS3(precedence 3) dscp (011000)
cs4       Match packets with CS4(precedence 4) dscp (100000)
cs5       Match packets with CS5(precedence 5) dscp (101000)
cs6       Match packets with CS6(precedence 6) dscp (110000)
cs7       Match packets with CS7(precedence 7) dscp (111000)
default   Match packets with default dscp (000000)
ef        Match packets with EF dscp (101110)
```

- **AF** (Assured Forwarding) defines four traffic classes. All packets in a class have the same priority.
- Within each class, there are three levels of *drop precedence*.
  - Higher drop precedence = more likely to drop the packet during congestion



= AFXY

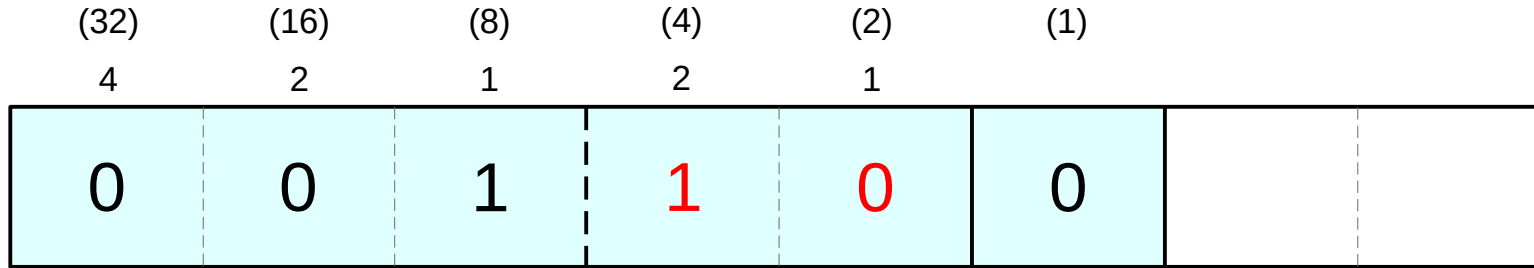
- **AF** (Assured Forwarding) defines four traffic classes. All packets in a class have the same priority.
- Within each class, there are three levels of *drop precedence*.  
→ Higher drop precedence = more likely to drop the packet during congestion



= AF11

(DSCP 10)

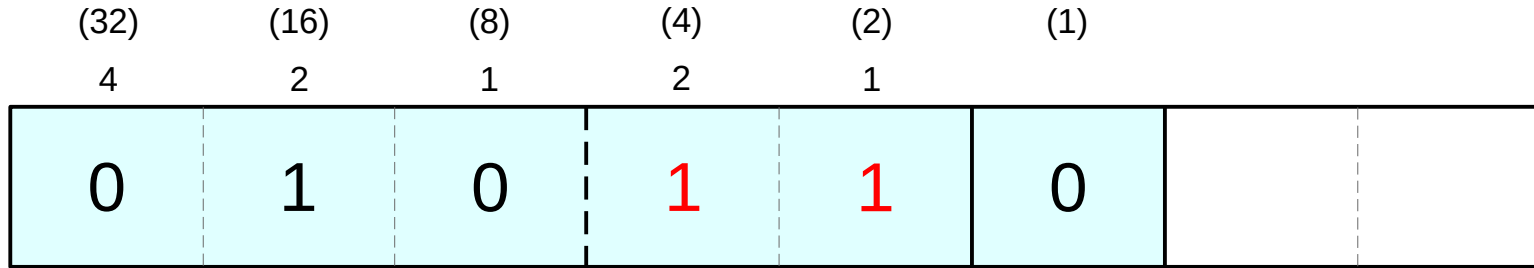
- **AF** (Assured Forwarding) defines four traffic classes. All packets in a class have the same priority.
- Within each class, there are three levels of *drop precedence*.  
→ Higher drop precedence = more likely to drop the packet during congestion



= AF12

(DSCP 12)

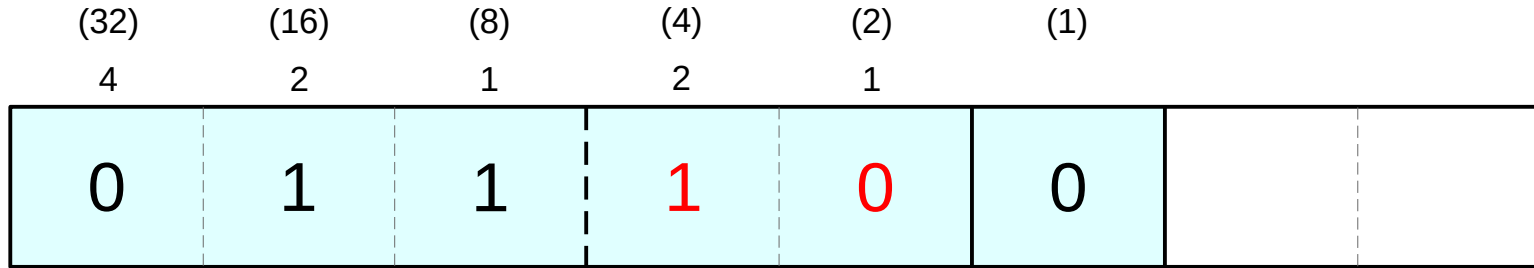
- **AF** (Assured Forwarding) defines four traffic classes. All packets in a class have the same priority.
- Within each class, there are three levels of *drop precedence*.  
→ Higher drop precedence = more likely to drop the packet during congestion



= AF23

(DSCP 22)

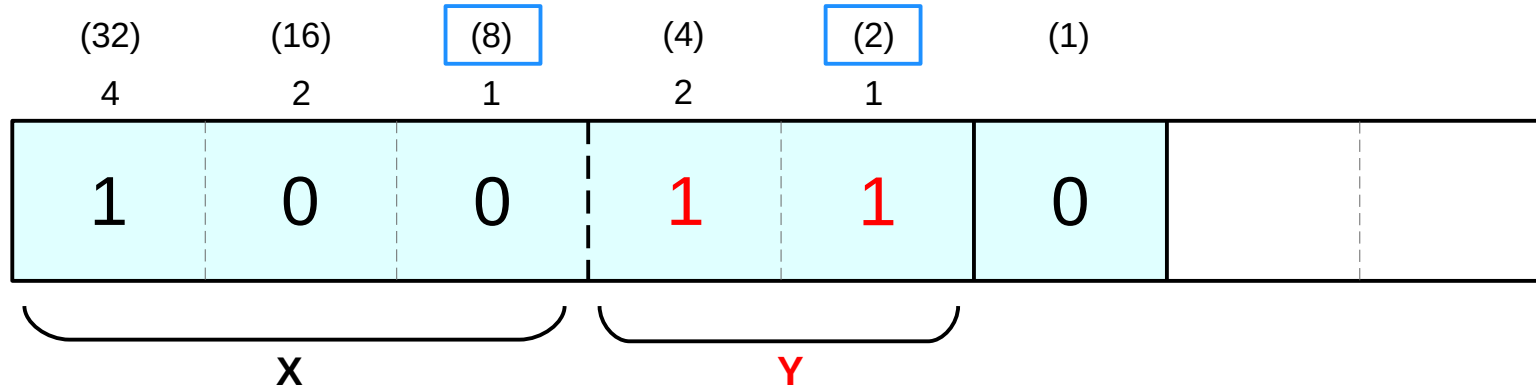
- **AF** (Assured Forwarding) defines four traffic classes. All packets in a class have the same priority.
- Within each class, there are three levels of *drop precedence*.  
→ Higher drop precedence = more likely to drop the packet during congestion



= AF32

(DSCP 28)

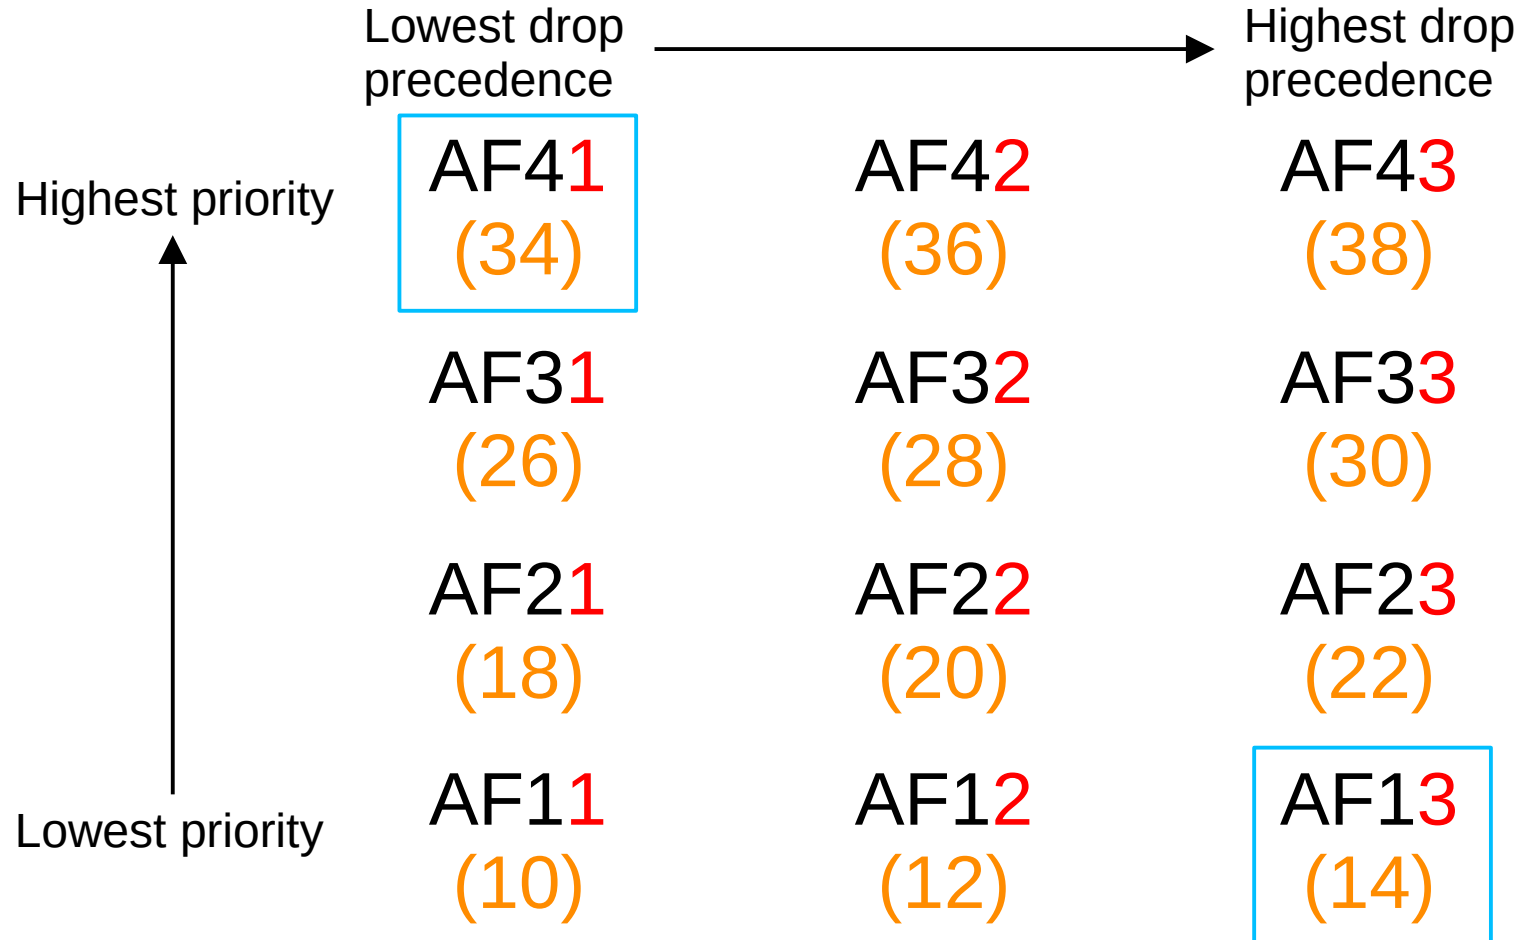
- **AF** (Assured Forwarding) defines four traffic classes. All packets in a class have the same priority.
- Within each class, there are three levels of *drop precedence*.  
→ Higher drop precedence = more likely to drop the packet during congestion



= AF43

(DSCP 38)

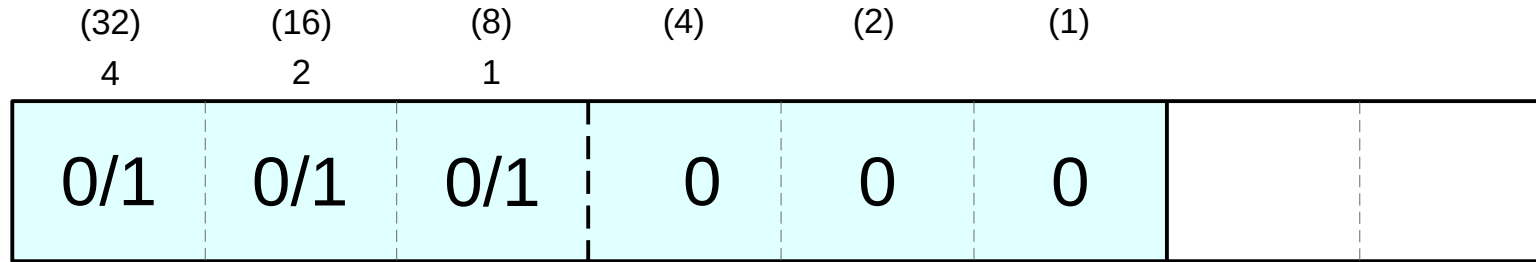
Formula to convert from AF value to decimal DSCP value:  $8X + 2Y$



```

R1(config)#class-map TEST
R1(config-cmap)#match dscp ?
<0-63>    Differentiated services codepoint value
af11      Match packets with AF11 dscp (001010)
af12      Match packets with AF12 dscp (001100)
af13      Match packets with AF13 dscp (001110)
af21      Match packets with AF21 dscp (010010)
af22      Match packets with AF22 dscp (010100)
af23      Match packets with AF23 dscp (010110)
af31      Match packets with AF31 dscp (011010)
af32      Match packets with AF32 dscp (011100)
af33      Match packets with AF33 dscp (011110)
af41      Match packets with AF41 dscp (100010)
af42      Match packets with AF42 dscp (100100)
af43      Match packets with AF43 dscp (100110)
cs1       Match packets with CS1(precedence 1) dscp (001000)
cs2       Match packets with CS2(precedence 2) dscp (010000)
cs3       Match packets with CS3(precedence 3) dscp (011000)
cs4       Match packets with CS4(precedence 4) dscp (100000)
cs5       Match packets with CS5(precedence 5) dscp (101000)
cs6       Match packets with CS6(precedence 6) dscp (110000)
cs7       Match packets with CS7(precedence 7) dscp (111000)
default   Match packets with default dscp (000000)
ef        Match packets with EF dscp (101110)
  
```

- CS** (Class Selector) defines eight DSCP values for backward compatibility with IPP.
- The three bits that were added for DSCP are set to 0, and the original IPP bits are used to make 8 values.

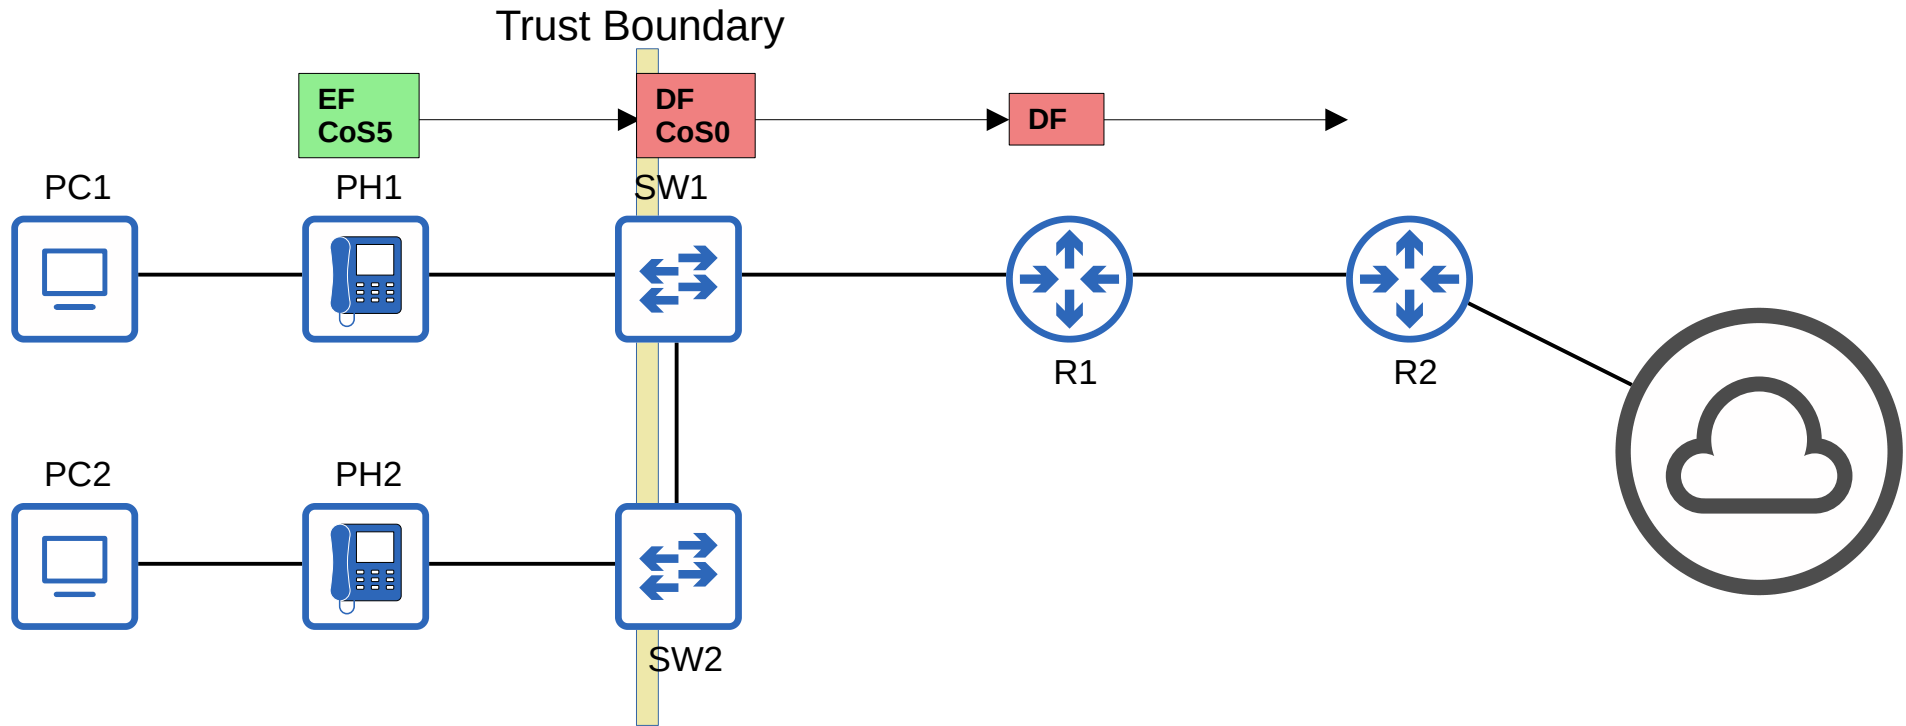


IPP:	0	1	2	3	4	5	6	7
CS:	CS0	CS1	CS2	CS3	CS4	CS5	CS6	CS7
DSCP: (decimal)	0	8	16	24	32	40	48	56

- RFC 4594 was developed with the help of Cisco to bring all of these values together and standardize their use.
- The RFC offers many specific recommendations, but here are a few key ones:
  - Voice traffic: **EF**
  - Interactive video: **AF4x**
  - Streaming video: **AF3x**
  - High priority data: **AF2x**
  - Best effort: **DF**

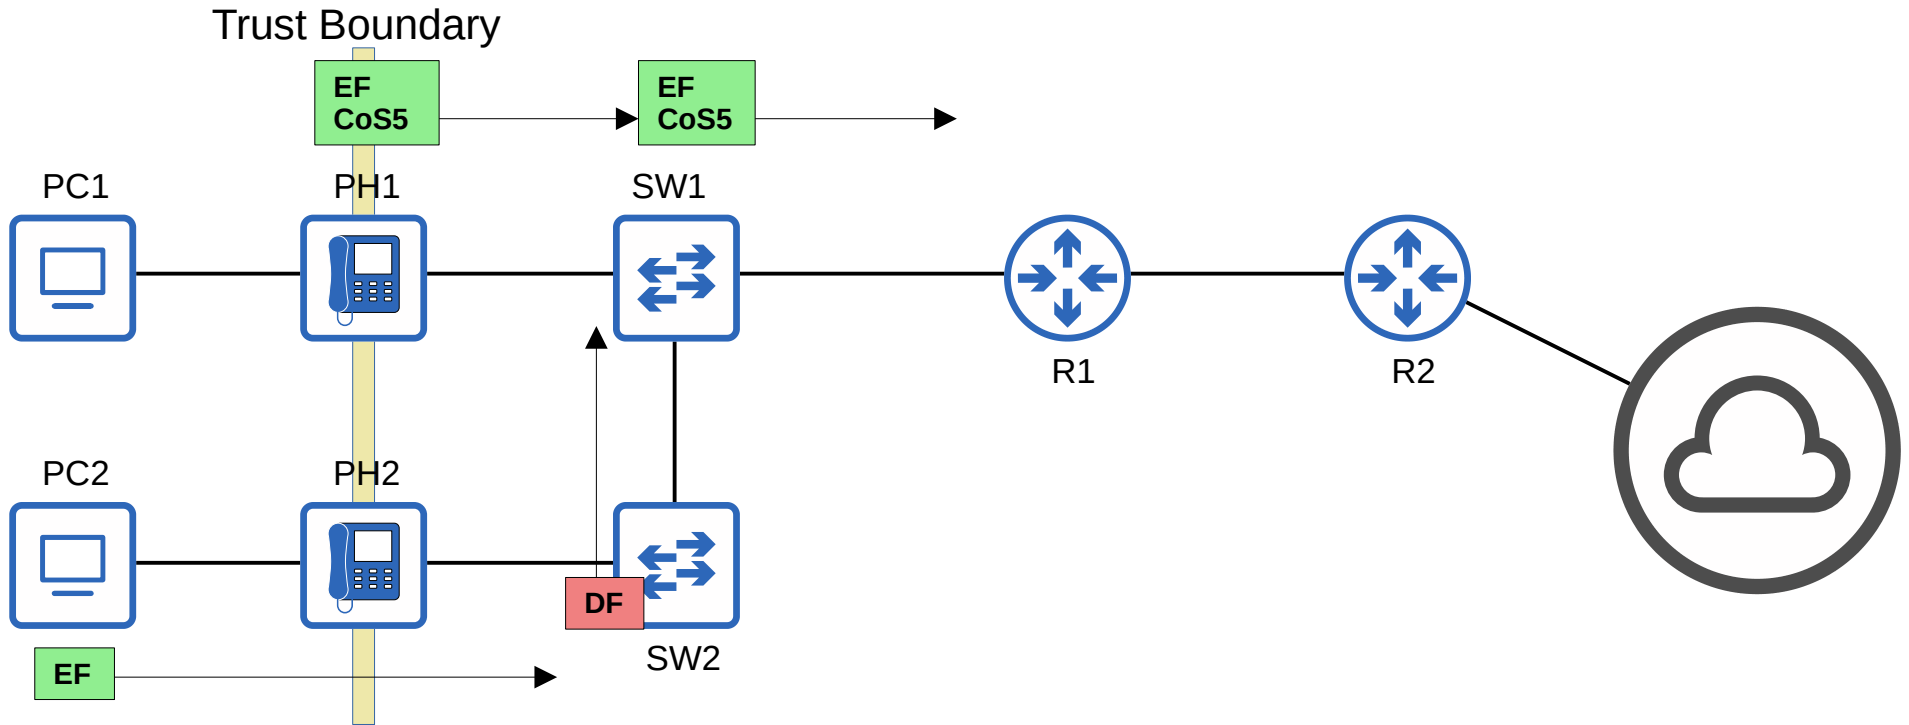
# Trust Boundaries

- The *trust boundary* of a network defines where devices trust/don't trust the QoS markings of received messages.
- If the markings are trusted, the device will forward the message without changing the markings.
- If the markings aren't trusted, the device will change the markings according to the configured policy.



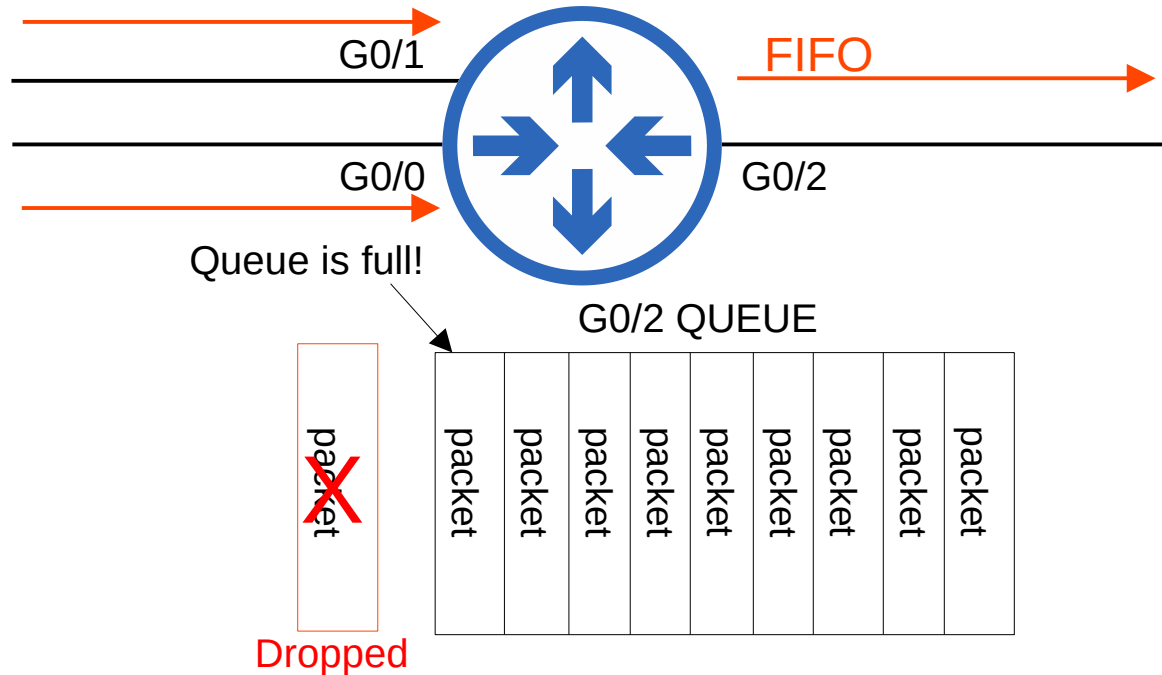
# Trust Boundaries

- If an IP phone is connected to the switch port, it is recommended to move the trust boundary to the IP phones.
- This is done via configuration on the switch port connected to the IP phone.
- If a user marks their PC's traffic with a high priority, the marking will be changed (not trusted)



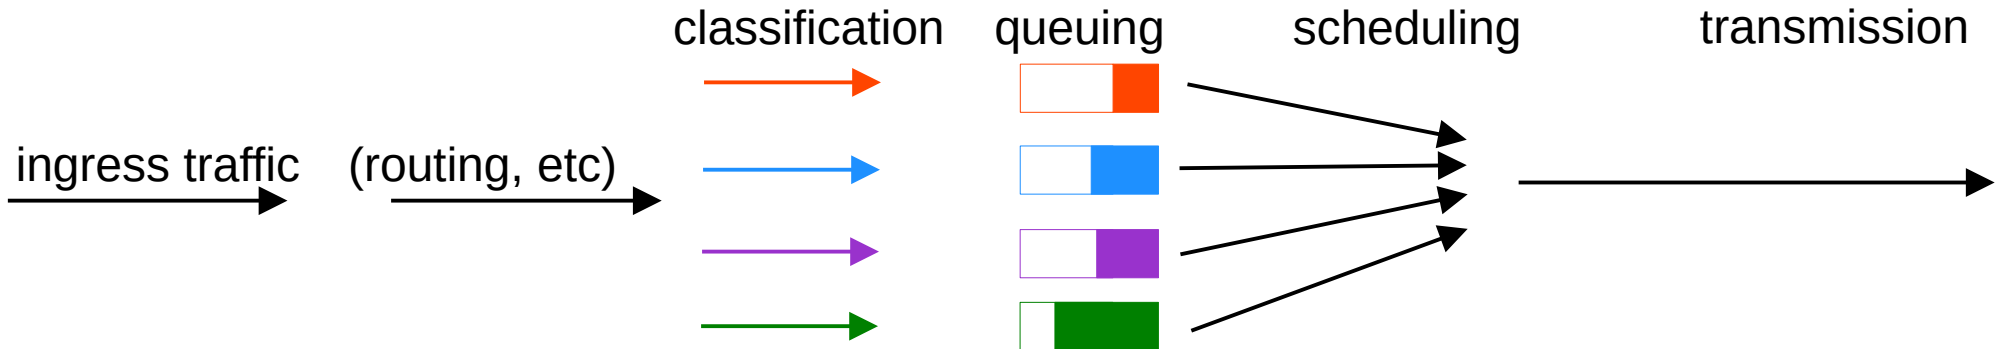
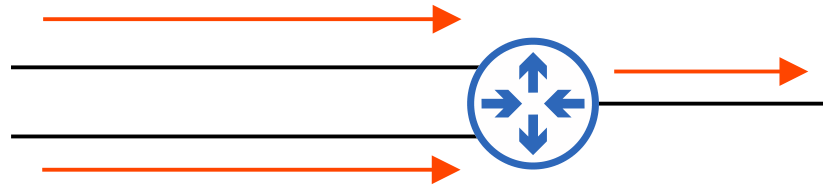
# Queuing/Congestion Management

- When a network device receives traffic at a faster rate than it can forward the traffic out of the appropriate interface, packets are placed in that interface's queue as they wait to be forwarded.
- When the queue becomes full, packets that don't fit in the queue are dropped (tail drop).
- RED and WRED drop packets early to avoid tail drop.



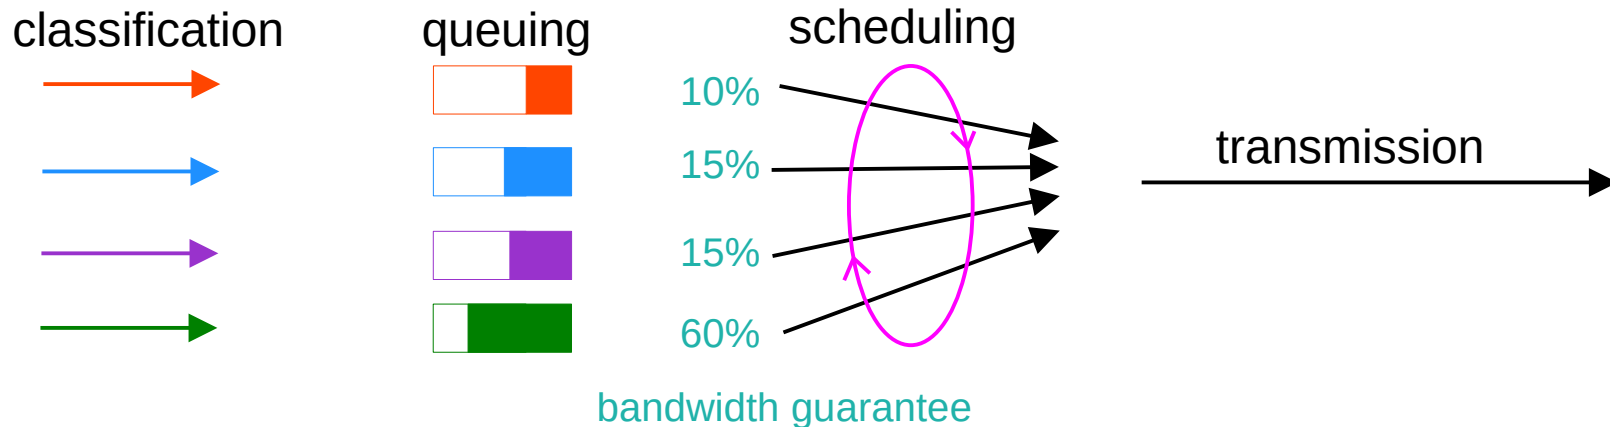
# Queuing/Congestion Management

- An essential part of QoS is the use of multiple queues.
  - This is where classification plays a role. The device can match traffic based on various factors (for example the DSCP marking in the IP header) and then place it in the appropriate queue.
- However, the device is only able to forward one frame out of an interface at once, so a *scheduler* is used to decide which queue traffic is forwarded from next.
  - *Prioritization* allows the scheduler to give certain queues more priority than others.



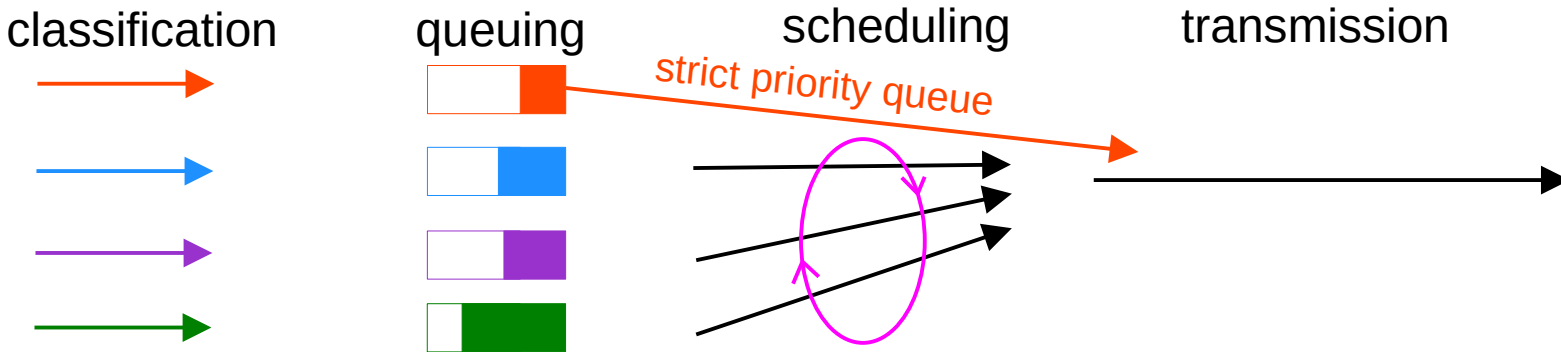
# Queuing/Congestion Management

- A common scheduling method is *weighted round-robin*.
  - **round-robin** = packets are taken from each queue in order, cyclically
  - **weighted** = more data is taken from high priority queues each time the scheduler reaches that queue
- CBWFQ** (Class-Based Weighted Fair Queuing) is a popular method of scheduling, using a weighted round-robin scheduler while guaranteeing each queue a certain percentage of the interface's bandwidth during congestion.
- Round-robin scheduling is not ideal for voice/video traffic. Even if the voice/video traffic receives a guaranteed minimum amount of bandwidth, round-robin can add delay and jitter because even the high priority queues have to wait their turn in the scheduler.



# Queuing/Congestion Management

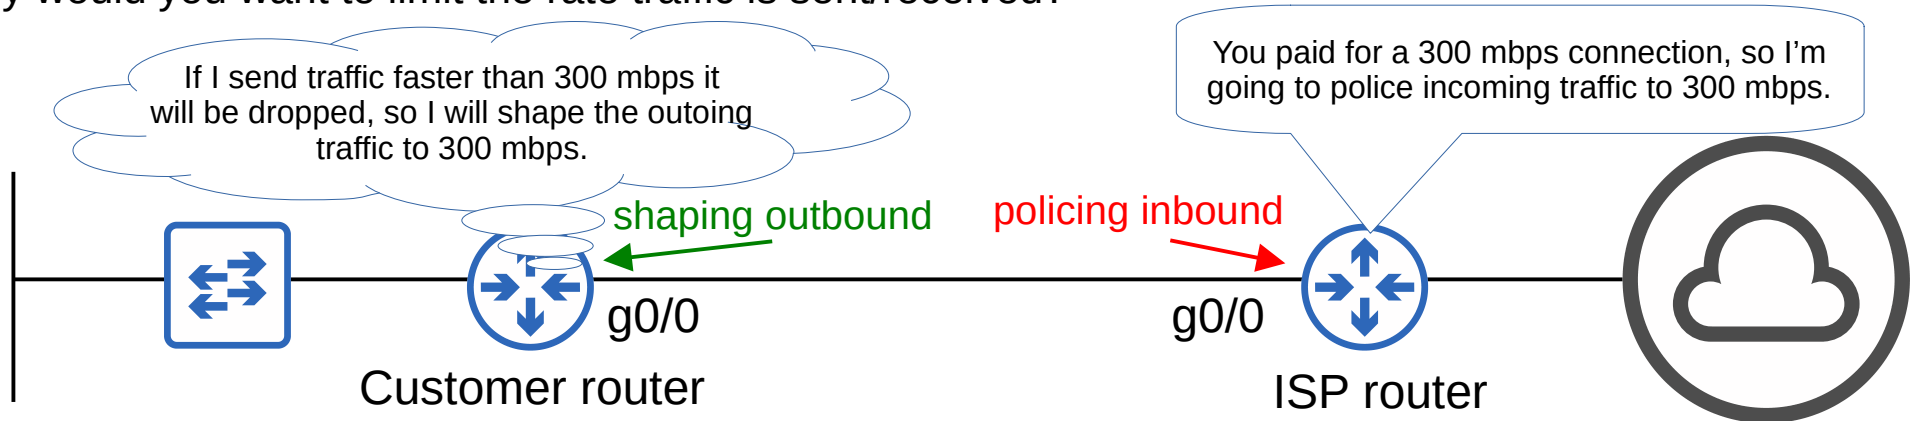
- **LLQ** (Low Latency Queuing) designates one (or more) queues as *strict priority queues*.  
→ This means that if there is traffic in the queue, the scheduler will always take the next packet from that queue until it is empty.
- This is very effective for reducing the delay and jitter of voice/video traffic.
- However, it has the downside of potentially starving other queues if there is always traffic in the designated strict priority queue.  
→ *Policing* (next slide) can control the amount of traffic allowed in the strict priority queue so that it can't take all of the link's bandwidth.



# Shaping and Policing

- Traffic **shaping** and **policing** are both used to control the rate of traffic.
- Shaping** buffers traffic in a queue if the traffic rate goes over the configured rate.
- Policing** drops traffic if the traffic rate goes over the configured rate.
  - 'Burst' traffic over the configured rate is allowed for a short period of time.
  - This accommodates data applications which typically are 'bursty' in nature. Instead of a constant stream of data, they send data in bursts.
  - The amount of burst traffic allowed is configurable.
- In both cases, classification can be used to allow for different rates for different kinds of traffic.
- Why would you want to limit the rate traffic is sent/received?

\*policing has the option of re-marking the packets instead of dropping them



# Things we covered

- Classification/Marking
- Queuing/Congestion Management
- Shaping/Policing

Which of the following CoS markings are consistent with standard practice?  
(select three)

a) Best effort = CoS 7

b) Best effort = CoS 0

c) Voice = CoS 4

d) Voice = CoS 5

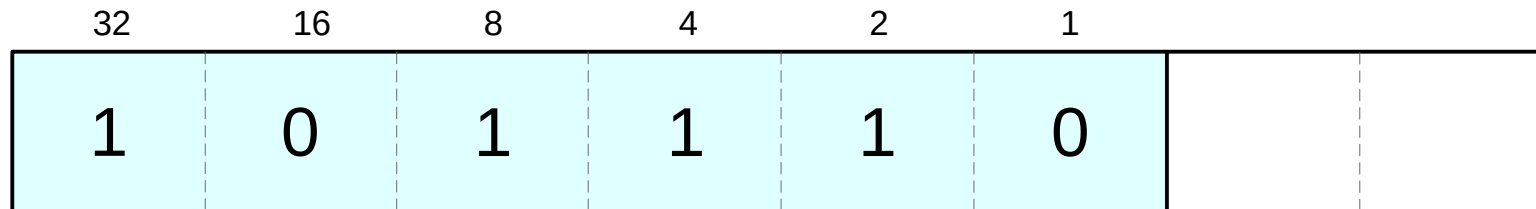
e) Video = CoS 4

f) Video = CoS 5

PCP value	Traffic types
0	Best effort (default)
1	Background
2	Excellent effort
3	Critical applications
4	Video
5	Voice
6	Internetwork control
7	Network control

What bit pattern would you find in the DSCP field of a packet marked as EF?

- a) 101 011
- b) 111 111
- c) 111 110
- d) 101 110



# Quiz 3

Which of the following AF markings provides the best service?

	Lowest drop precedence		Highest drop precedence
a) AF43	AF4 <b>1</b> (34)	AF4 <b>2</b> (36)	AF4 <b>3</b> (38)
b) AF41			
c) AF51	AF3 <b>1</b> (26)	AF3 <b>2</b> (28)	AF3 <b>3</b> (30)
d) AF11	AF2 <b>1</b> (18)	AF2 <b>2</b> (20)	AF2 <b>3</b> (22)
e) AF13			
f) AF61	AF1 <b>1</b> (10)	AF1 <b>2</b> (12)	AF1 <b>3</b> (14)

Diagram illustrating the relationship between AF markings, drop precedence, and priority:

- A horizontal arrow points from "Lowest drop precedence" to "Highest drop precedence".
- A vertical arrow points from "Lowest priority" to "Highest priority".

Which of the following statements represents general best practice regarding QoS?

- a) Trust markings from IP phones. Don't trust markings from PCs.
- b) Trust all QoS markings.
- c) Don't trust any QoS markings.
- d) Trust markings from PCs. Don't trust markings from IP phones.

Which of the following creates a strict priority queue for data that requires low delay/jitter/loss?

- a) CBWFQ
- b) LLQ
- c) Weighted round-robin
- d) WFQ