Introduction

- Welcome to Jeremy's IT Lab, a free, complete course for the CCNA.
- In part 1, the first half covered **voice VLANs** and **Power over Ethernet**; the second half briefly introduced **QoS**.
- **Purpose of QoS:** to prioritize certain traffic such as voice and video traffic to reduce the **delay, jitter, and loss**, and also to ensure it gets the required bandwidth.
- This video goes more in depth about the terms in **exam topic 4.7**: **classification, marking, queuing, congestion, policing, and shaping**.
- **Note:** QoS design and configuration are **not needed for the CCNA**. As long as the topics introduced in this video are understood, QoS questions on the CCNA exam should be answerable.

What Will Be Covered in Today's Video

- **Classification and marking** — how routers and switches identify which traffic should be given higher priority.
- **Queuing and congestion management** — already introduced in the previous video, now covered more in depth.
- **Shaping and policing** — two ways of controlling the rate at which traffic enters or exits an interface.
- Watch until the end of the video for a **bonus practice question** from Boson Software's ExSim for CCNA, the best practice exams for the CCNA.

Classification

- The **purpose of QoS** is to give certain kinds of network traffic **priority** over others during times of **network congestion**.

- **Classification** organizes network traffic, meaning packets, into **traffic classes** or **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.

Methods of Classifying Traffic

- There are **many methods** of classifying traffic.

Example: ACL

- Traffic **permitted** by the ACL will be given certain treatment, for example it could be treated as **high-priority**, and other traffic which is **denied** by the ACL will not be given that treatment.
- This is like how ACLs were used for **dynamic NAT** — the ACL isn't being used to actually permit or deny traffic, it's just being used to **identify** certain traffic.

Example: NBAR (Network Based Application Recognition)

- Sometimes just looking at the **lower-layer information**, such as the **TCP or UDP port number** and **IP address**, isn't enough to identify exactly what kind of traffic it is.
- NBAR performs what's called a **"deep packet inspection"** to look beyond the Layer 3 and 4 information all the way up to **Layer 7** to identify the specific kind of traffic.
- **Note:** We won't focus on methods like these in this video.

Fields Used for Classification in Layer 2 and Layer 3 Headers

- In the Layer 2 and Layer 3 headers there are **specific fields** used for the purpose of classifying traffic.
- The **PCP (Priority Code Point)** field of the **802.1Q tag** can be used to identify high or low priority traffic.
    - **Keep in mind:** this field can only be used when there is a **dot1q tag**, only when a VLAN tag is added to the Ethernet header.
- The **DSCP (Differentiated Services Code Point)** field of the **IP header** can also be used to identify high or lower priority traffic.

PCP (Priority Code Point)

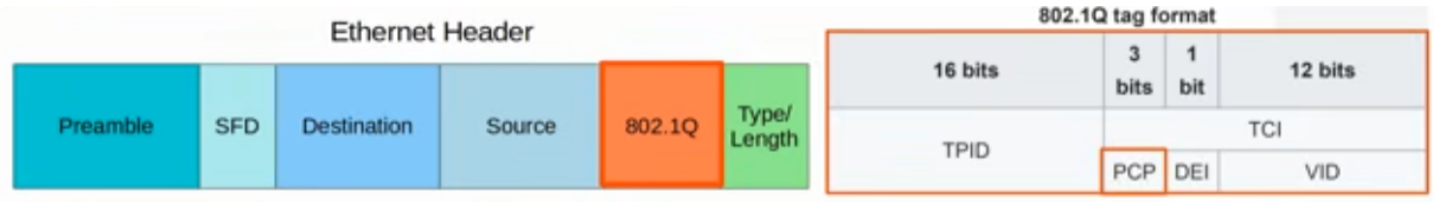

- The **PCP field** is located in the **dot1q tag** (VLAN tag) of the Ethernet header.
- It is a **3-bit field**, providing **8 possible values** from 0 to 7.
- **Note:** This field may also be referred to as **CoS (Class of Service)**. Do not confuse it with the entire concept of QoS — **CoS** just refers to this part of the dot1q tag.
- Its use is defined by **IEEE 802.1p**.

Recommended PCP Values

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

- The professor recommends remembering the following values:
  - **0 — Best Effort**
  - **3 — Critical Applications**
  - **4 — Video**
  - **5 — Voice**

- **Best Effort delivery** means there is no guarantee that data is delivered or that it meets any QoS standard. It is regular traffic, not high-priority.
- **Note:** Best Effort is the **default** — regular traffic will have a value of **0** in the PCP field.

Example: IP Phone Markings

- IP phones mark their **call signaling traffic** as **PCP3**. Call signaling traffic is used to establish phone calls.
- When the call is established, the actual voice traffic itself (the audio packets) is marked as **PCP5**.

Marking

- **Mark** is another important QoS term.
- To **mark** traffic is to **set the value** in the **PCP or DSCP fields**.
- Network devices look at those markings and use them to **classify the traffic** as high-priority, low-priority, etc.
- When an IP phone marks its voice traffic as **PCP5**, it wants the routers and switches to classify those packets as **high-priority**.

PCP (Priority Code Point) — Network Example

Example: Simple Network with PCs and IP Phones

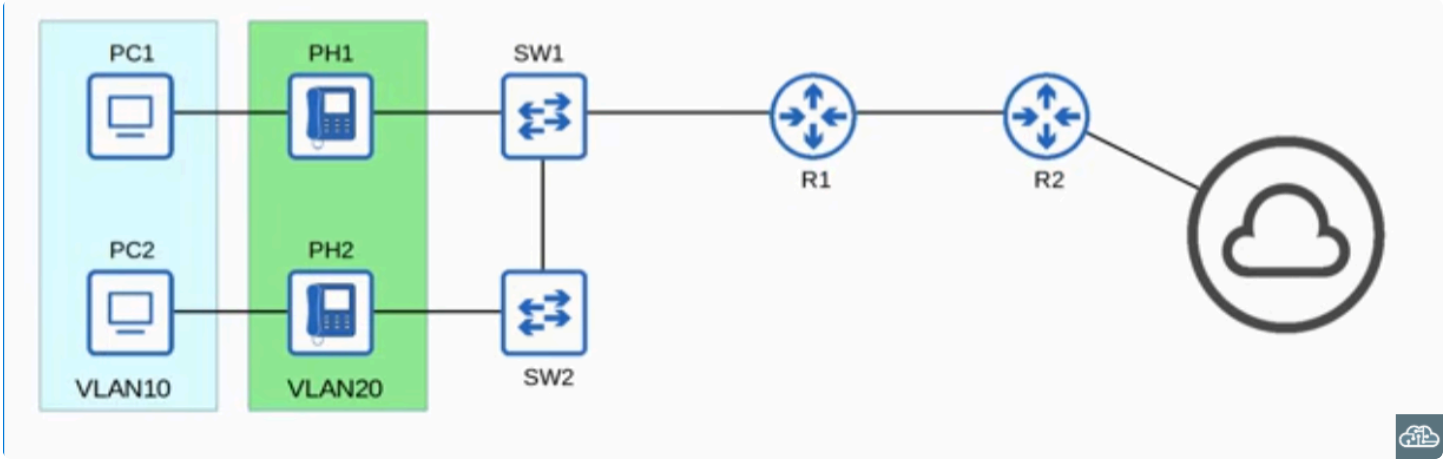

- The network consists of a couple of **PCs in VLAN10** and **IP phones in VLAN20**.

- This example demonstrates a very important point about using **PCP** to classify traffic.

PCP Usage Limitations

- Because the **PCP field** is found in the **dot1q header**, it can only be used over the following connection types:
  - **Trunk links** — traffic sent over a trunk link is tagged with dot1q, unless the traffic is in the **native VLAN**, but the professor assumes the native VLAN isn't being used.
  - **Access links with a voice VLAN** — voice traffic is tagged as well even over access links.

Identifying Links Where PCP Can Be Used

- Assuming the network is configured so that all PCs and phones can communicate with each other as well as over the external network, the following links are either **trunk links** or **access links with a voice VLAN**:
  - The connections between the **phones and the switches** are **access ports with a voice VLAN**. Traffic to and from the phones will be **tagged**.
  - The two connections between the switches are **trunk links**, tagging traffic in **VLAN 10 or 20**.
- Over these connections, devices can **mark traffic** with a specific **PCP value** to tell other devices to treat the traffic with a certain level of priority.
- Devices receiving that marked traffic can **classify the traffic** as high, medium, low, or whatever priority based on the **PCP marking**.

Links Where PCP Cannot Be Used

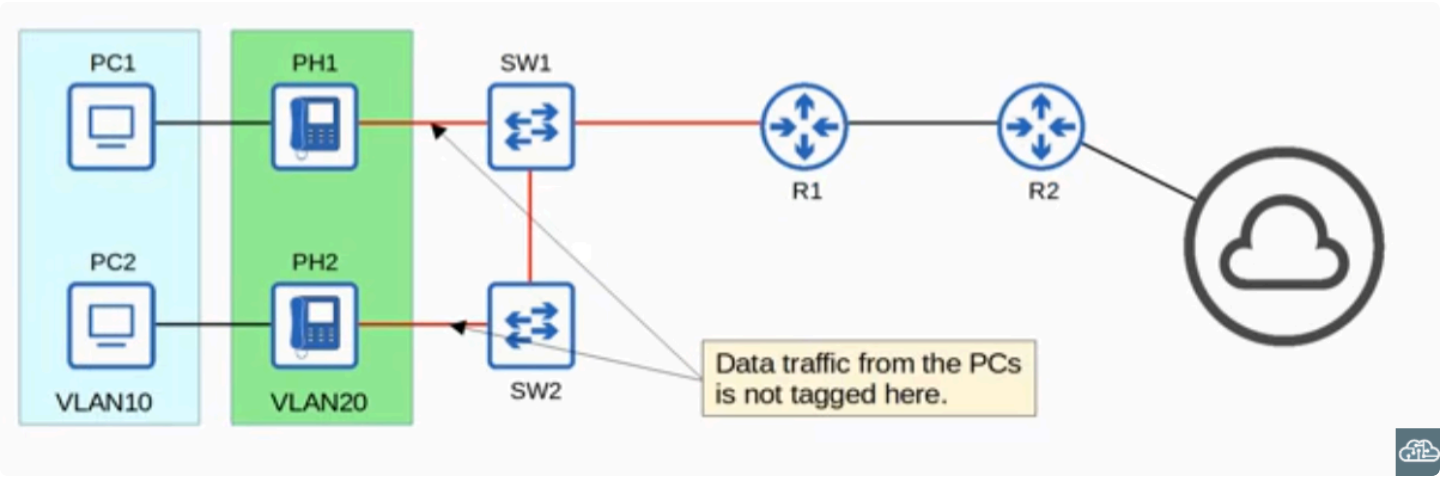

- In this type of network design, **traffic from the PCs is not tagged with dot1q**, so it can't have a **PCP marking** over this link.
- In addition, all traffic between **R1 and R2**, or between **R2 and external destinations**, will not have a **dot1q tag**.
- Traffic over those links cannot be marked with a **PCP value — PCP cannot be used to classify traffic**.

Note: This is a **major limitation** of using **PCP**, a limitation which the **Layer 3 classification methods** don't have.

Layer 3 Marking and Classification

- Now let's look at how **marking and classification** is done at **Layer 3**.
- In the **IPv4 header**, there is a byte referred to as the **ToS byte** (Type of Service).
- **IPv6** also has a byte called the **traffic class byte** used for QoS, but we will focus on the IPv4 header.
- The ToS byte is located in the IPv4 header, specifically the **second byte** after the **version and IHL fields**.

Modern Use of the ToS Byte

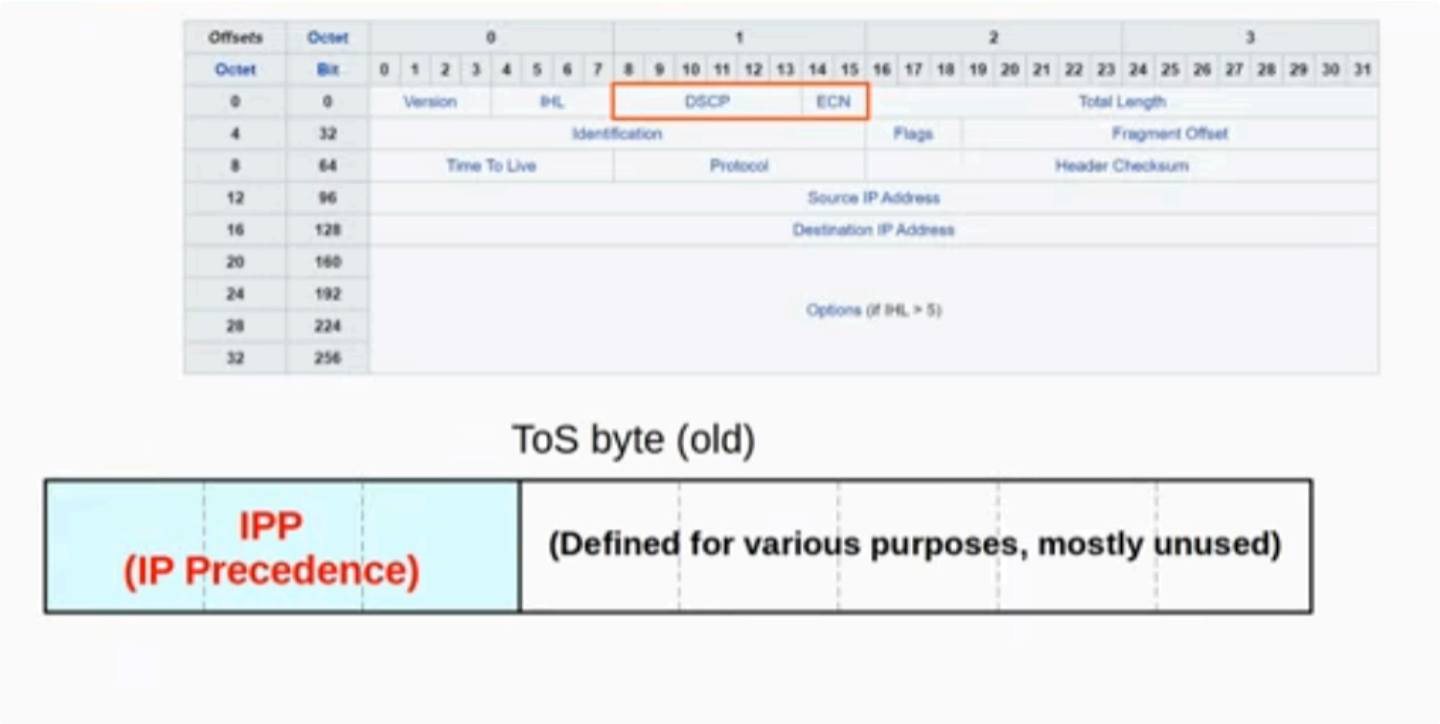

- The modern use of the ToS byte consists of **two fields**:
  - **DSCP** (Differentiated Services Code Point)
  - **ECN** (Explicit Congestion Notification)
- The current ToS byte is organized as **6 bits for DSCP** and **2 bits for ECN**.
- 6 bits in DSCP allows for a total of **64 values**, which gives a lot of flexibility regarding how traffic can be marked and classified in the network.

Old Use of the ToS Byte

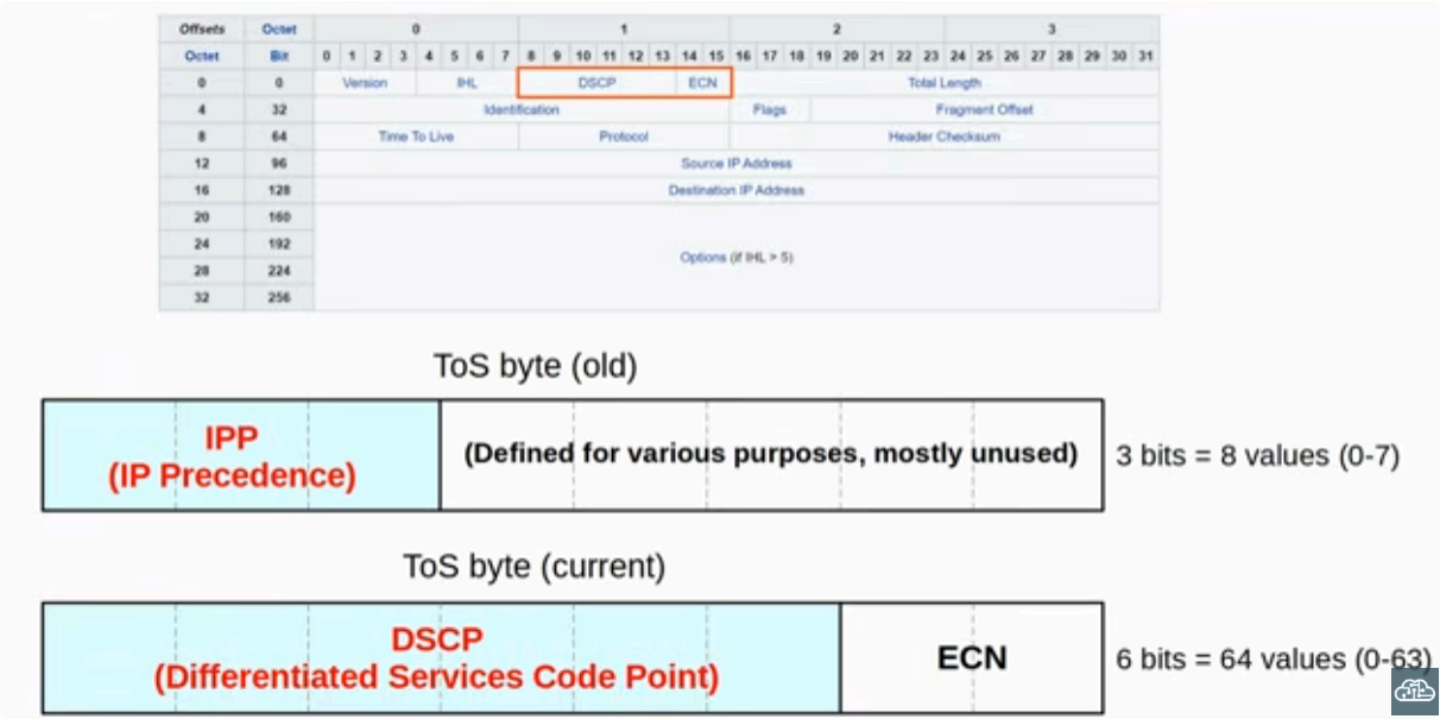

- Previously, the ToS byte was organized **differently**.
- **3 bits** were used for the **IPP** (IP Precedence) field.
  - IPP was used to **mark packets according to their priority**, like the PCP field in the dot1q header.
- The remaining **5 bits** were mostly defined for some other purpose, but the professor's understanding is that they **weren't really used**.
- 3 bits gives **8 possible values**, just like in the PCP field of the dot1q header.

IPP (IP Precedence)

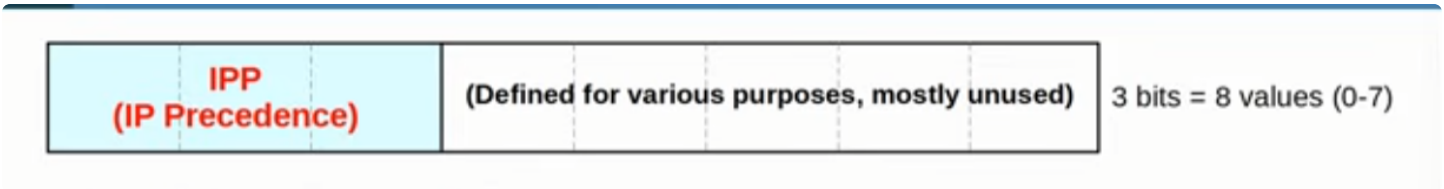

- Before going in depth about **DSCP**, let's take a quick look at **IPP**.
- The standard **IPP markings** are similar to **PCP**.

- **6** and **7** are reserved for **network control** traffic — that refers to traffic like **OSPF messages** sent between routers.
- **Interactive voice** traffic is marked as **IPP 5**.
- **Interactive video** traffic is marked as **IPP 4**.
- **IPP 3** is used for **voice signaling traffic** — for setting up and tearing down phone calls.
- **IPP 0** is used for **best effort traffic** — regular data traffic without any special requirements.
- With **6** and **7** reserved, only **6 possible values** remain.
- Although 6 is sufficient for many networks, the **QoS requirements** of some more complex networks demand more flexibility.
- Because IPP only used **3 bits** of the ToS byte and the other 5 largely weren't used, it was decided that an **extra 3 bits** would be added to IPP to make **DSCP** (Differentiated Services Code Point).

DSCP (Differentiated Services Code Point)

- **RFC 2474**, published in **1998**, defines the **DSCP field**, and then other **DiffServ** (differentiated services) RFCs after it elaborate on the use of the field.
- With **IPP** updated to **DSCP**, new standard markings had to be decided upon.
- Because of generally agreed upon standard markings for different kinds of traffic:
  - **QoS design and implementation** is simplified.
  - **QoS works better** between ISPs and enterprises because they agree upon the markings that will be used.
  - There are other benefits too.
- A few sets of **industry-standard markings** were developed.

- **Note:** It might feel overwhelming trying to remember all of these. Flashcards are included to help memorize them, but the only ones that really need to be remembered are some of the more common ones. Those will be pointed out after all of the markings are introduced.

Standard DSCP Markings

- The following standard markings should be known:
  - **DF (Default Forwarding)** — the marking for **best effort traffic**, which doesn't have any particular QoS requirements.
  - **EF (Expedited Forwarding)** — used for traffic that requires **low loss, latency, and jitter**, which is usually **voice traffic**.
  - **AF (Assured Forwarding)** — not one marking, but a set of **12 standard values**.
  - **CS (Class Selector)** — a set of **8 standard values**. **CS** provides **backward compatibility with IPP**.

QoS Configuration Example (Class-Map)

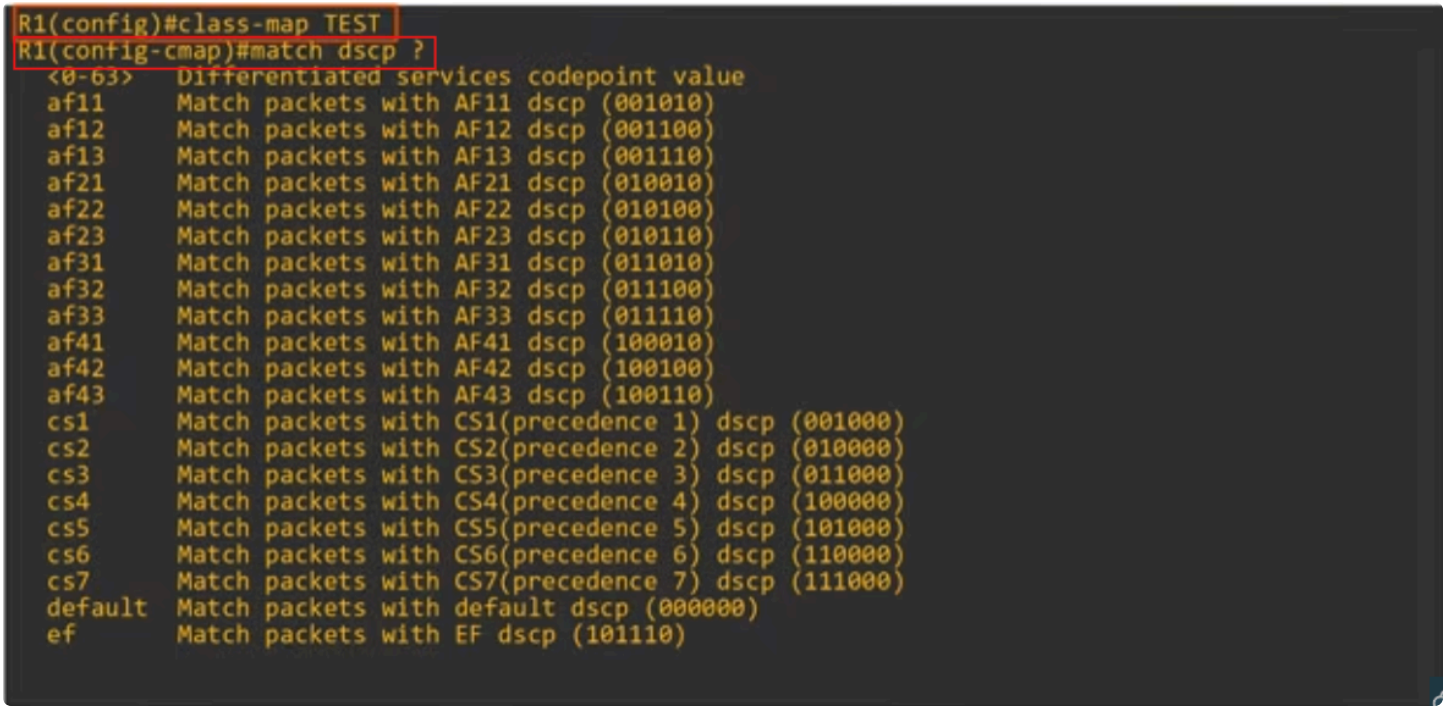

- **Note:** QoS configuration is **not covered** because it is **not a CCNA exam topic**, but the professor shows a configuration example.

- He configured a **class-map** called **TEST**.
- When configuring QoS, **class maps** are used to identify which traffic you want to match.
- He configured a **MATCH statement**, specifying that he wants to match traffic based on the **DSCP value** in the IP header.
- The options for matching DSCP include:

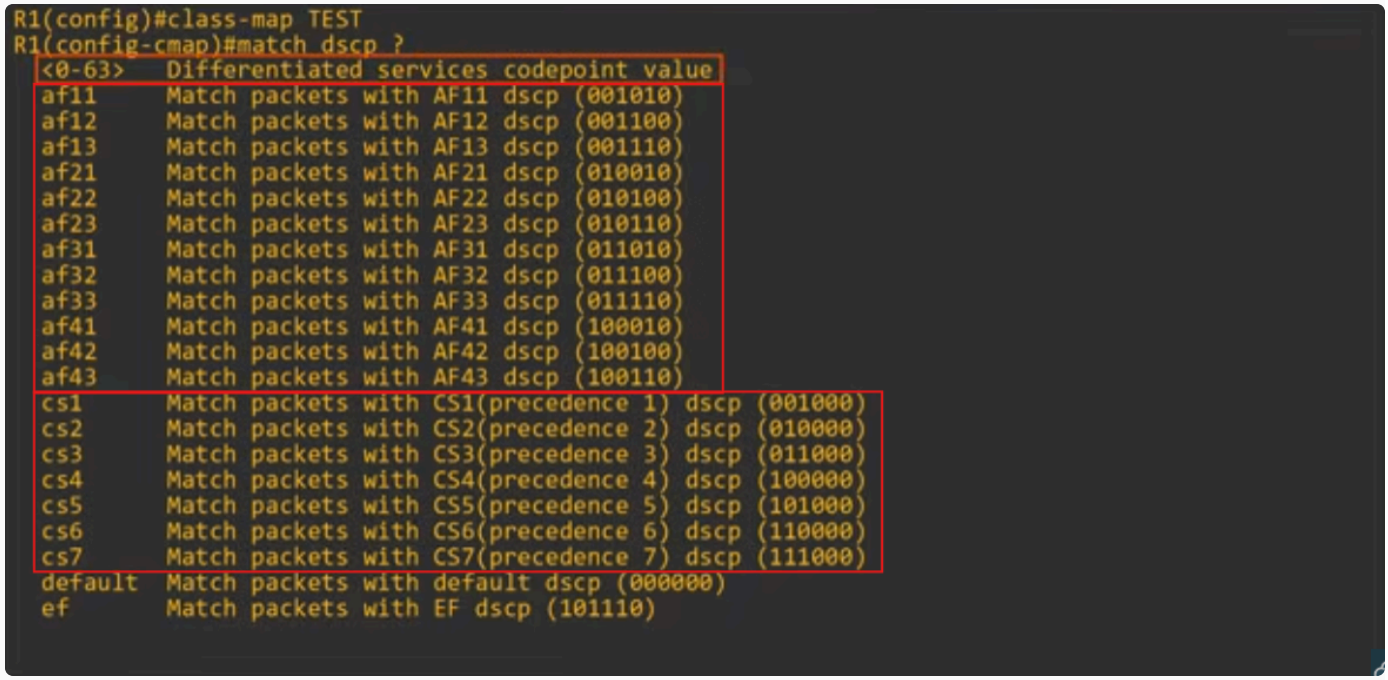

  - The choice to simply configure the **DSCP value in decimal**, from **0 to 63**.
  - The **12 AF values**.
    - Example: **AF11** has a binary value of **001 010**.
  - The **CS values**.
    - There are **8 CS values**, but only 7 are shown because the other one, **CS0**, is the same as **DF** (default forwarding), which is **000 000**.
  - **EF** (Expedited Forwarding) at the bottom.
- **Note:** The binary values are displayed on the right for AF and CS.

DF and EF

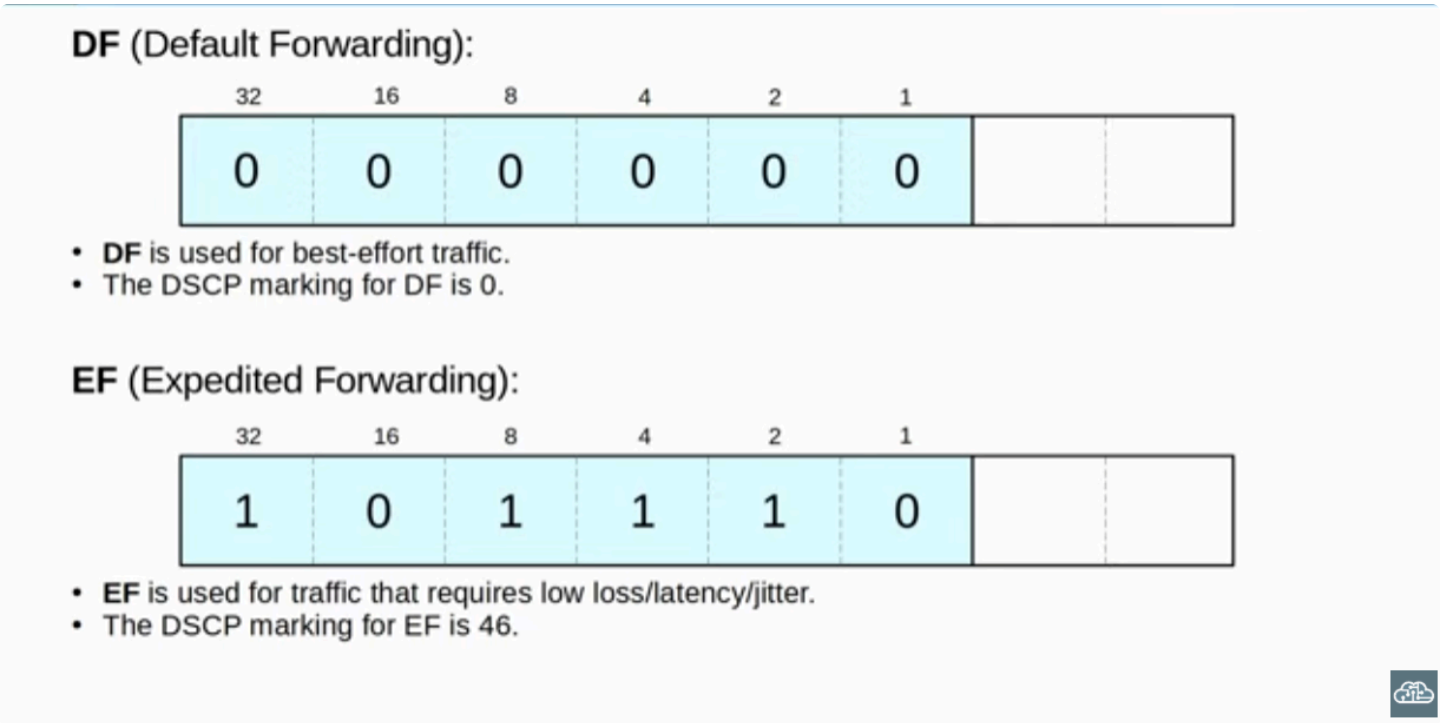

- **DF** is used for **best-effort traffic**. The recommended **DSCP marking** is **0**. So, all 6 bits will be set to **0**.
- To make things easier to understand, the professor writes the **decimal value of each binary bit** above the bits.
- **Example:** Standard network traffic like sending an email or downloading a file will probably have **0 0 0 0 0 0** in the DSCP field of the IP header, indicating that it doesn't require any special treatment.
- **EF (Expedited Forwarding)** is used for traffic that requires **low loss, latency, and jitter**. Typically, **voice traffic** is marked as **EF**.
- The **DSCP marking for EF** is **46**. In binary, this is **1 0 1 1 1 0**.
- Summary:
  - **DF** is used for **best-effort traffic**, and the **DSCP value** is **0**.
  - **EF** is used for traffic requiring **low-loss, low-latency, and low-jitter**, usually **voice traffic**, and the **DSCP value** is **46**.
- **Note:** The professor indicates that things will get a bit more complicated next.

AF (Assured Forwarding)

- **AF** can be a little tricky to understand, but the professor explains it clearly.
- **AF** defines **four traffic classes**. All packets in a class have the **same priority**.
- A **higher class number** means **higher priority** — their packets will be forwarded with **better service** than lower-priority packets.
- Within each class, there are **three levels of drop precedence**.
- A **higher drop precedence** means the packet is **more likely to drop** during congestion due to **WRED**.

AF Bit Structure

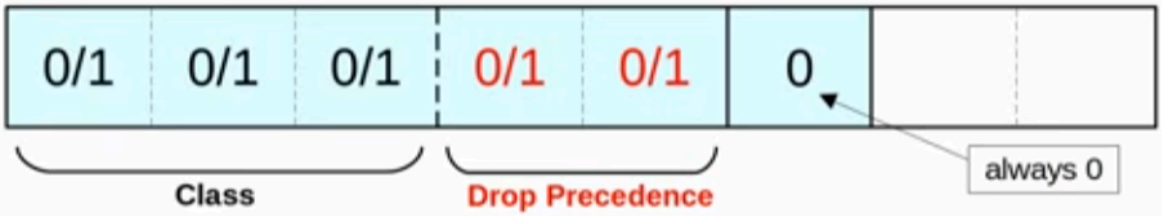

- The 6 DSCP bits are set up a little differently for AF:
  - In AF, the **bit on the end is always set to 0**. The professor is not exactly sure of the reason — it might just be because the designers of AF decided they didn't need so many values.
  - **Two bits** represent the **drop precedence**.
  - **Three bits** represent the **class**.
- When writing an AF value, it is written as **AF X Y**, with **X** being the decimal number of the **class** and **Y** being the decimal number of the **drop precedence**.

Example: Binary DSCP Value 001 010

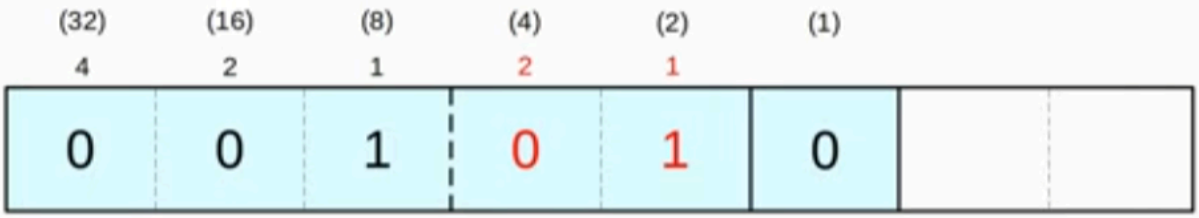

= AF11

(DSCP 10)

- Split it up into two parts: three bits for the class, two bits for the drop precedence.
- The decimal of the class is **1**, and the decimal of the drop precedence is also **1**.
- So this is **AF11** (or AF1 1).
- This is really just a 6-bit DSCP number. Using the normal 6-bit decimal values (1 2 4 8 16 32), this is **DSCP 10**.
- **AF11 is the same value as DSCP10** — this is because both represent the same 6-bit binary pattern.
- **AF is just a set of standard DSCP markings**, which is **easier to work with** than simply having **64 DSCP values** with no system of standard values.

Example: Binary DSCP Value 001 100

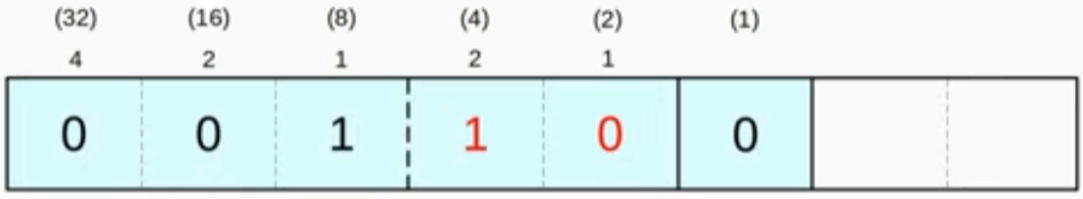

= AF12

(DSCP 12)

- This time it is class **1**, drop precedence **2**.
- So it is **AF12** (or AF1 2).

- As a normal decimal DSCP value, **AF12 is the same as DSCP 12**.

Example: Binary DSCP Value 010 011 (Practice)

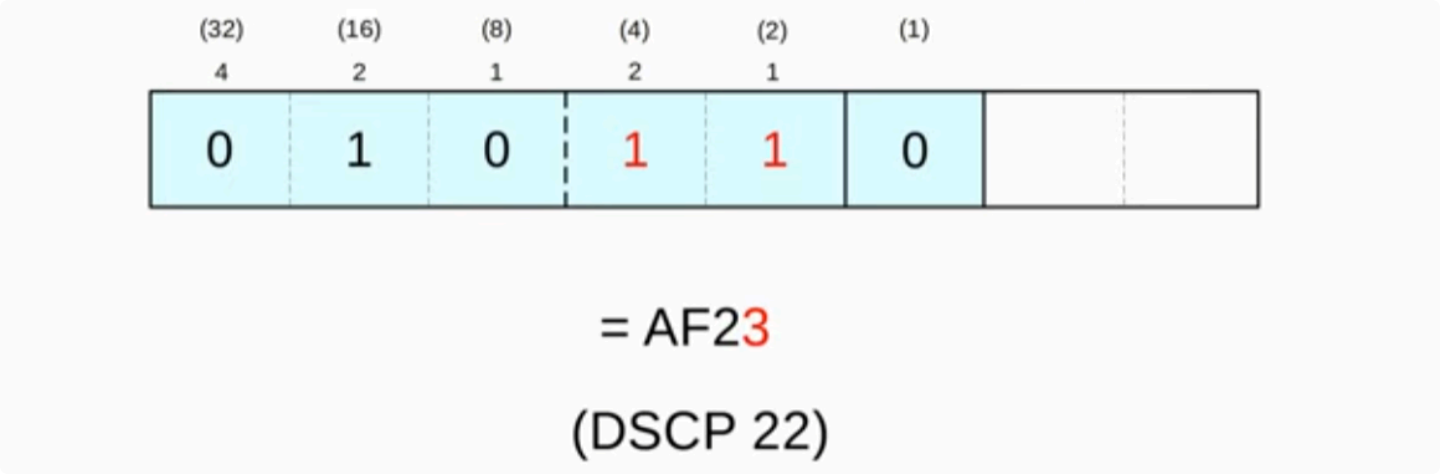

- This is **AF23** (or AF2 3).
- Written as a normal DSCP value, it is **DSCP 22**.

Example: Binary DSCP Value 011 100

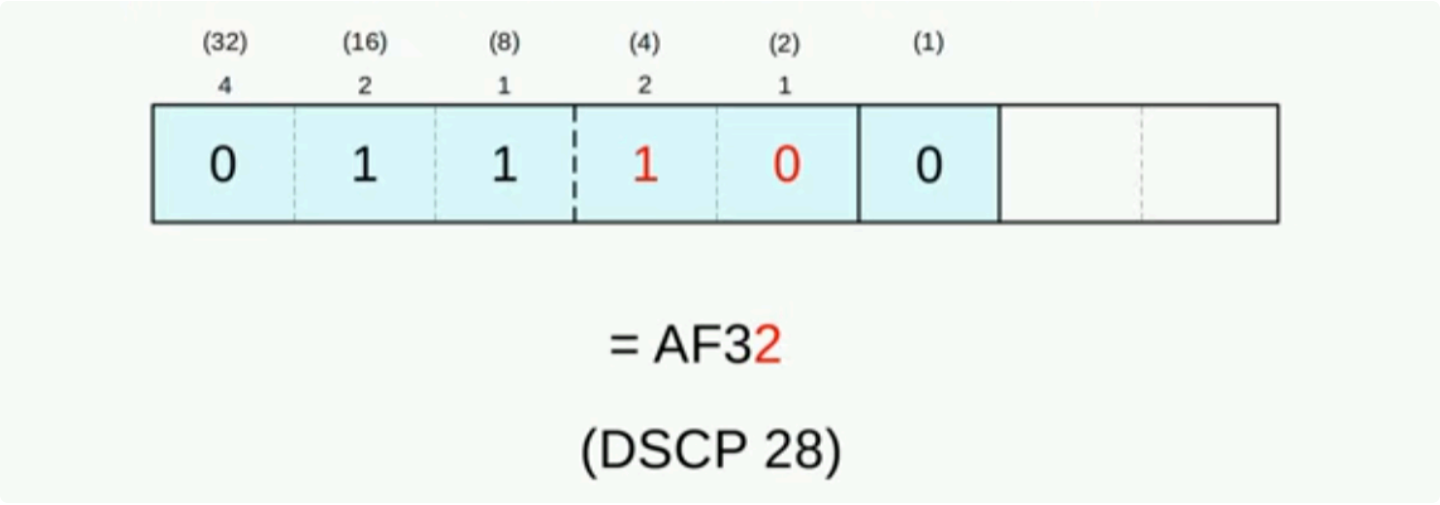

- This is **AF32** (or AF3 2).
- Written as a normal DSCP value, it is **DSCP 28**.
- So, **AF32 is the same as DSCP 28**.

Example: Binary DSCP Value 100 011

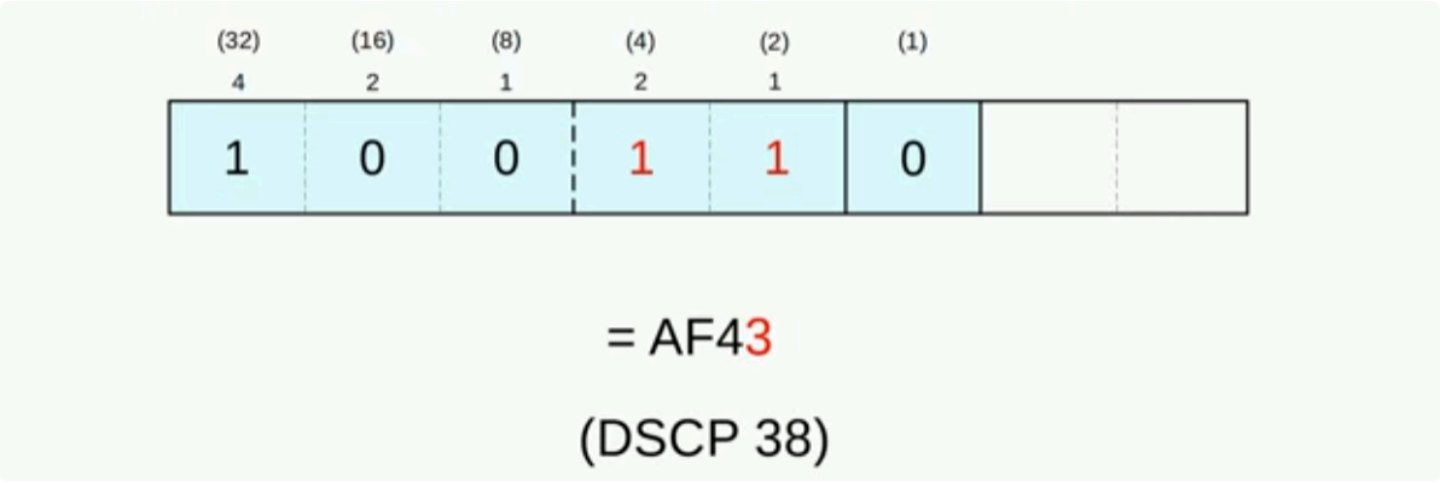

- This is **AF43** (or AF4 3).
- Written as a normal DSCP value, it is **DSCP 38**.
- So, **AF43 is equivalent to DSCP 38**.
- **Note:** Within AF, **43 is the highest value**. There is no class 5, 6, or 7.

Converting AF to DSCP — Shortcut Formula

- To convert between AF numbers and DSCP values without writing out the binary every time, use the formula:

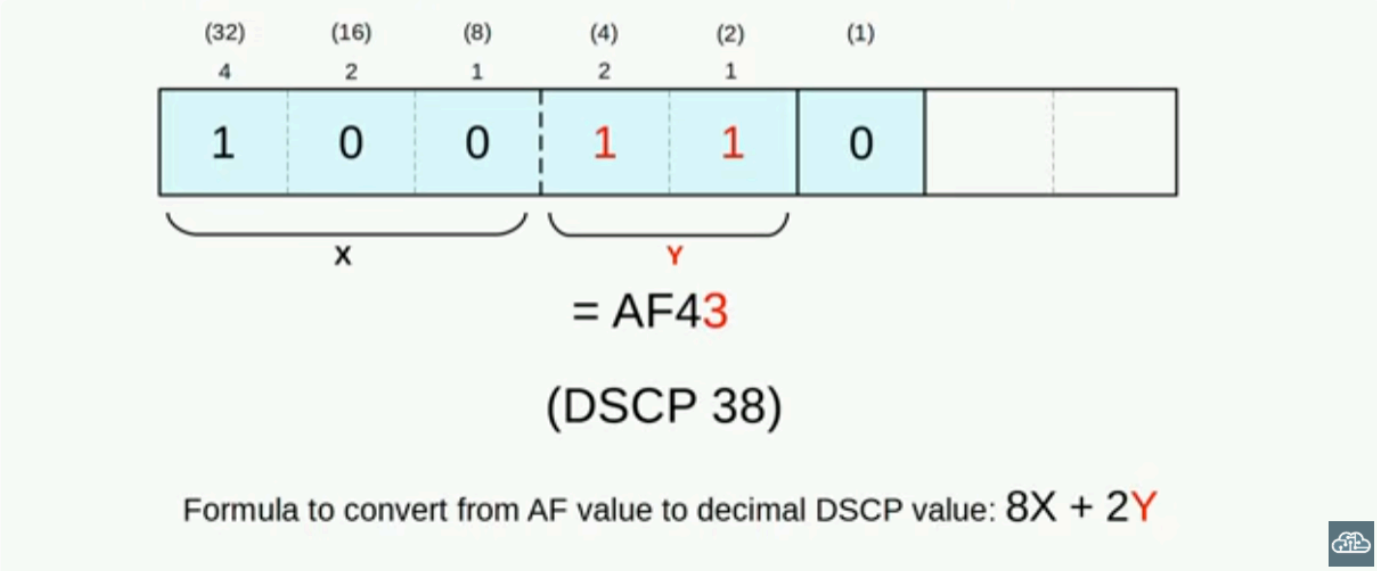

  - **DSCP = 8X + 2Y**, where **X** is the class and **Y** is the drop precedence.
- The reason for this is that **8** is the lowest value of the 'X' portion, and **2** is the lowest value of the 'Y' portion.

- **Note:** It is always best to understand the binary underneath it all, whether learning IPv4 addresses, subnetting, IPv6, matching with ACLs, QoS, etc. But after understanding the binary, it is also nice to have shortcuts like this.

AF (Assured Forwarding) — Summary of AF Values

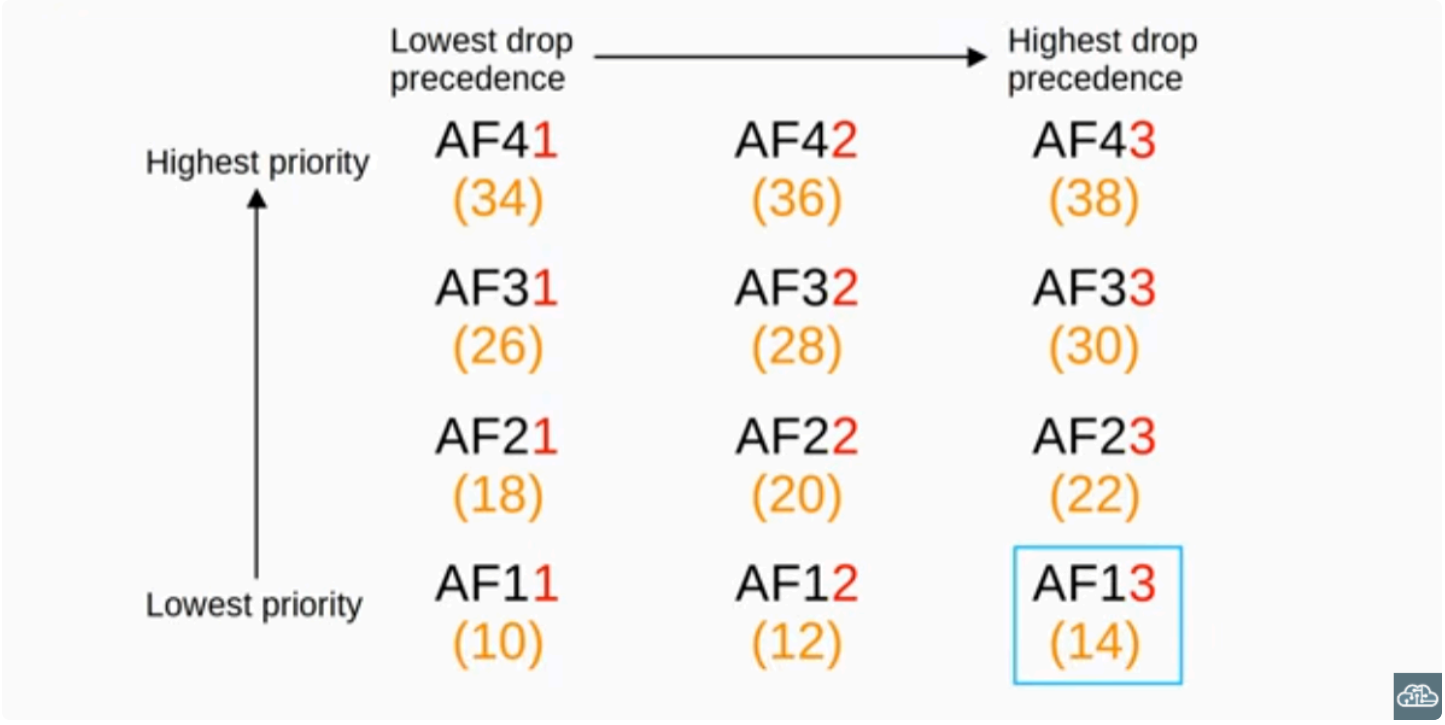

- Here is a **summary of all of the AF values**: **4 classes** with **3 drop precedence values** each.
- Within these AF values:
  - Traffic marked as **AF41** gets the **best treatment** — it is in the **highest priority class** and has the **lowest drop precedence**.
  - Traffic marked as **AF13** gets the **worst treatment** — it is in the **lowest priority class** and has the **highest drop precedence**.
- The professor will include **flashcards** for all of these to test the ability to convert them from the **AF values** to the **regular DSCP values**.
- **Note:** Students are **not expected to memorize** them, but should be able to **calculate the DSCP value** of a given AF value.

CS (Class Selector)

- Fortunately, **CS** is actually **quite simple**.
- **CS (Class Selector)** defines **eight DSCP values** for **backward compatibility with IPP**.

How CS Backward Compatibility Works

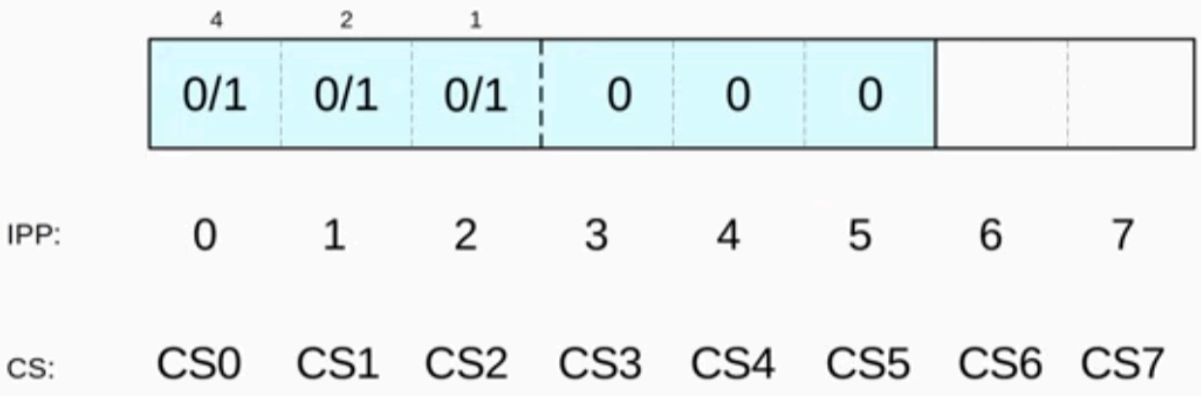

- The **three bits** that were added for DSCP are set to **0**, and the **original IPP bits** are used to make **8 possible values**.
- The **three bits on the right** are set to **0**, and the **three bits on the left** (the original IP precedence bits) can be either **0 or 1**.
- The decimal values of those three bits are **1, 2, and 4**.
- So, **IPP** gives us values from **0 through 7**.
- The equivalent values in **CS** are: **CS0, CS1, CS2, CS3, CS4, CS5, CS6, and CS7**.

Converting CS to Decimal DSCP Values

- Although CS is a way of organizing the DSCP field into a set of 8 values, they are still **DSCP values**.

- The conversion is quite simple: it is **8 multiplied by the CS number**.

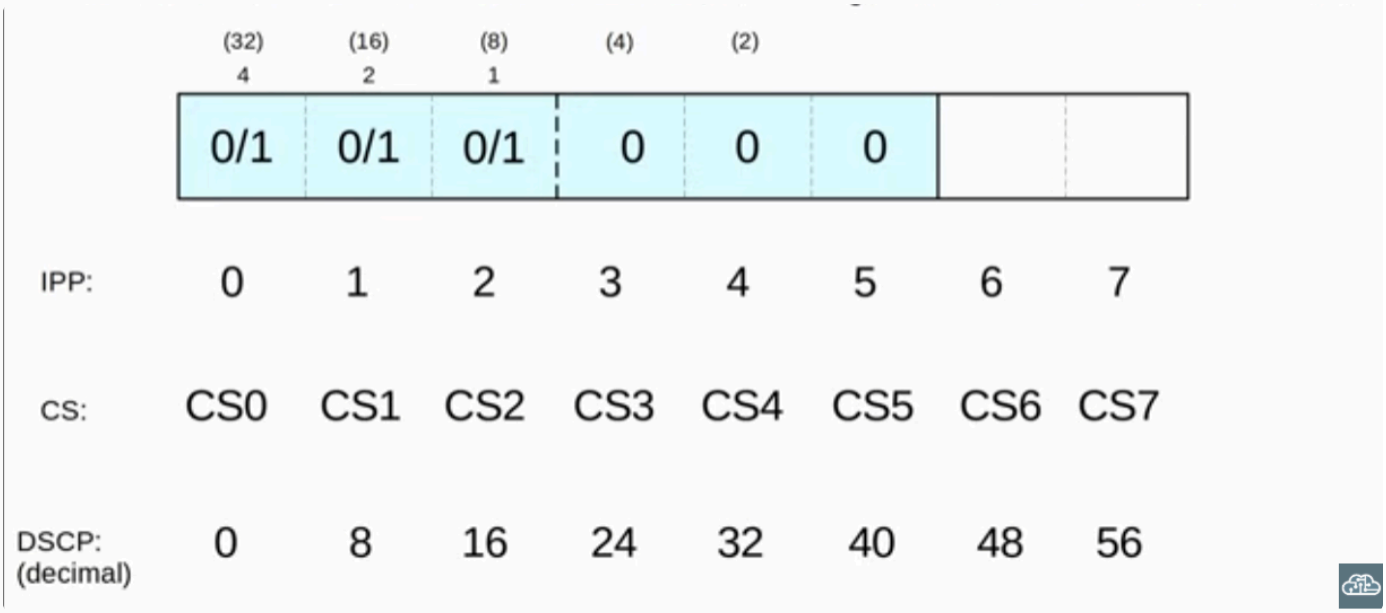

  - **CS0** is **0**
  - **CS1** is **8**
  - **CS2** is **16**
  - **CS3** is **24**
  - **CS4** is **32**
  - **CS5** is **40**
  - **CS6** is **48**
  - **CS7** is **56**

Using DSCP Values — Standard Recommendations

- Now that we've got **DF, EF, AF, and CS**, how are we supposed to take those values and actually use them?
- **RFC 4954** was developed with the help of **Cisco** to bring all of these values together and **standardize their use**.

- Many specific recommendations are given in the RFC, but here are a few key ones:
  - **Voice traffic** is marked as **EF**, because it requires **very low delay, jitter, and loss**.
  - **Interactive video** should be marked as **AF4x**, meaning AF41, 2, or 3.
  - **Streaming video** is marked as **AF3x**.
  - **High priority data** is **AF2x**.
  - **Best effort data** is marked as **DF**, a DSCP value of **0**.
- There are many more recommendations given within the RFC — a Google search can be done if interested.
- **Note:** In the end, it's up to the **engineer designing the QoS policy** of the network which traffic will get which markings. These are just **standard recommendations**.

Trust Boundaries

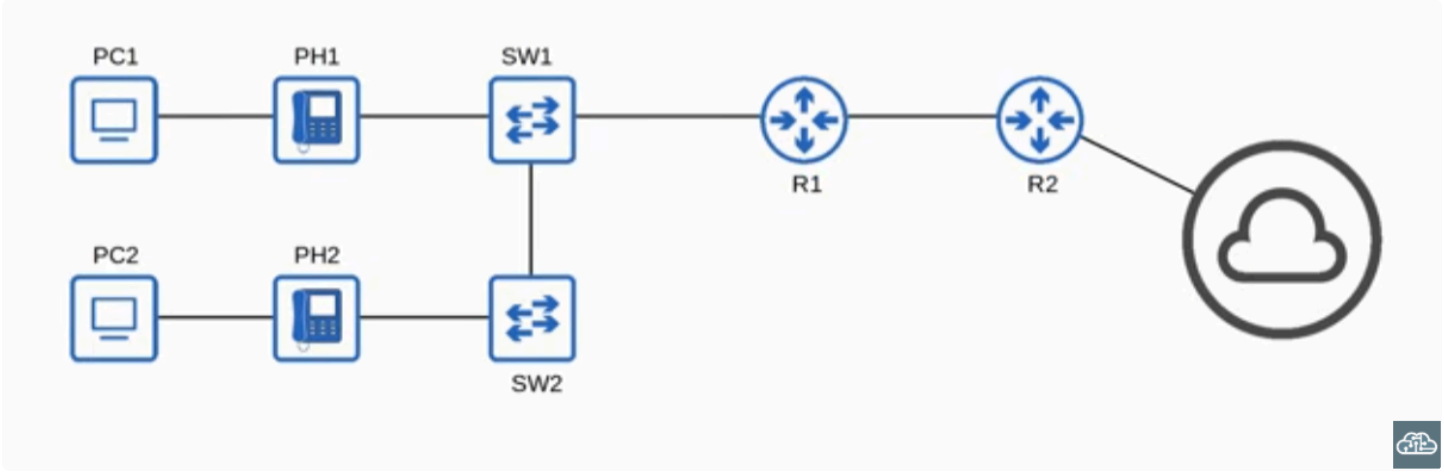

- The **trust boundary** of a network defines where devices **trust** and **don't trust** the **QoS markings** of received messages.
- If the markings are **trusted**, that means the device will **forward the message without changing** the markings.
- But if the markings **aren't trusted**, the device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

Example: Untrusted Markings at SW1

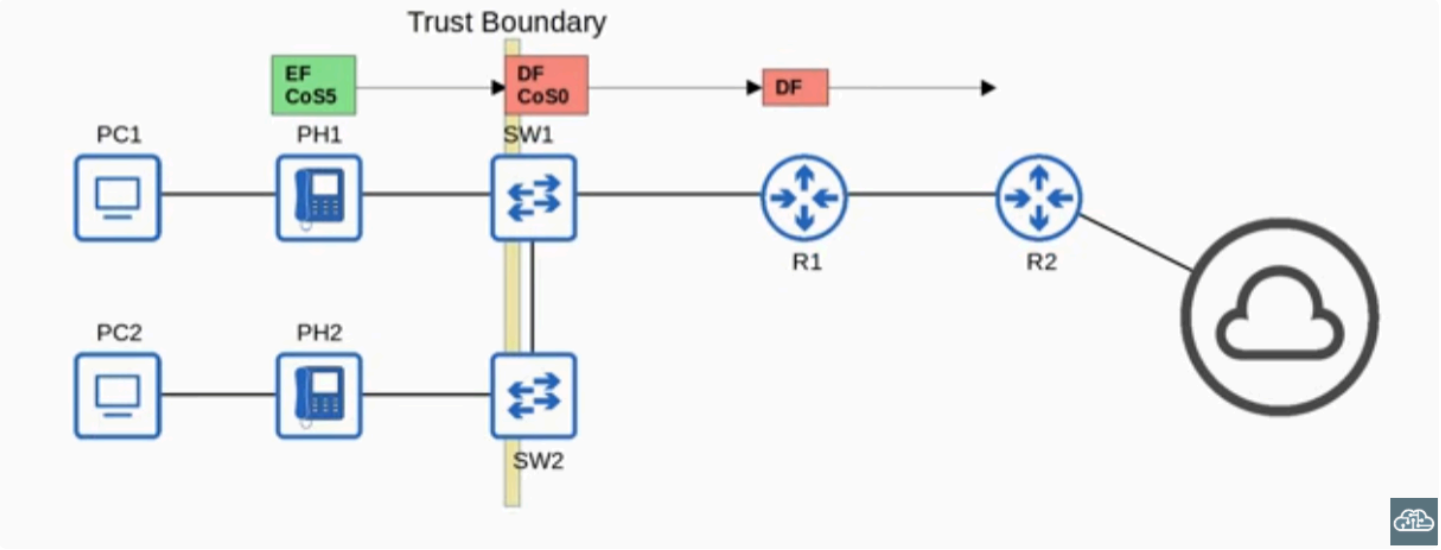

- Let's say the trust boundary is at **SW1**.
- **Phone1** sends a message marked as **EF** and **CoS5** to SW1.
    - **Note:** CoS is referring to the **PCP field** in the dot1q header. You may hear it called **CoS5** or **PCP5**, you should be familiar with both terms.
- SW1 **doesn't trust** the markings from Phone1 because it's from **outside of the trust boundary**.
- So perhaps it changes the **DSCP marking to DF** and the **CoS marking to 0**, before forwarding it to R1.
- R1 forwards it on to R2 with just the **DF marking** because there is **no dot1q header**.
- **Note:** This configuration isn't ideal. Usually it's **best to trust the markings from an IP phone** because we want its traffic to be high priority anyway.

Moving the Trust Boundary

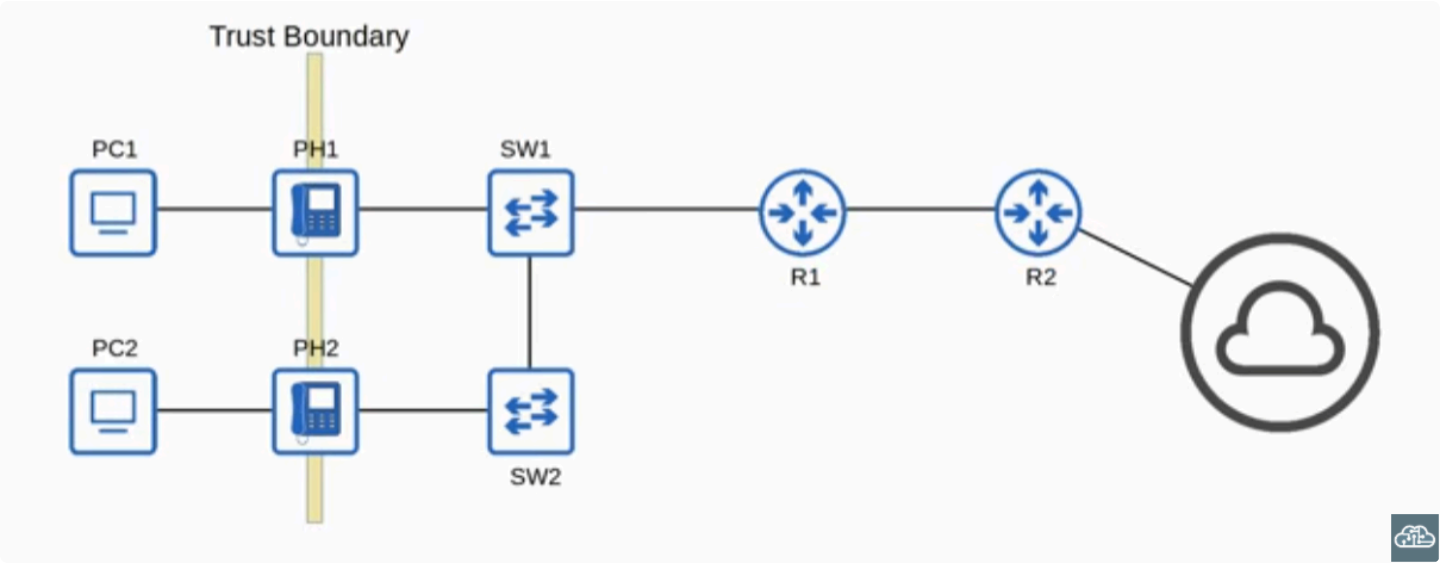

- If an IP phone is connected to the switchport, it is **recommended to move the trust boundary to the IP phone**.
- This is done via configuration on the **switch port connected to the IP phone**, by the way, not directly on the phone itself.
- So, if a tech-savvy user is able to mark their PC's traffic with a **high-priority marking** to get faster service, the marking will be **changed** according to the configured policy.
- But traffic sent from the **phone itself** will be **trusted** by the switch.

Example: Trusted Markings from IP Phone vs. PC

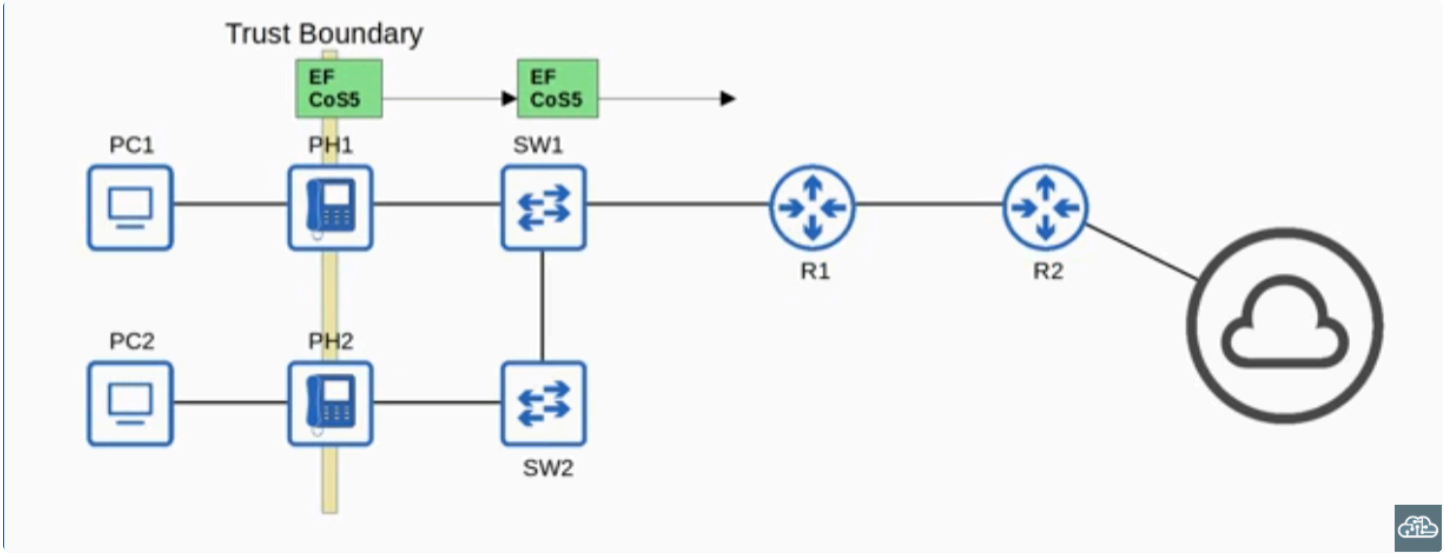

- If **Phone1** sends an **EF and CoS5**-marked message, SW1 will **trust** those markings and **not change** them.

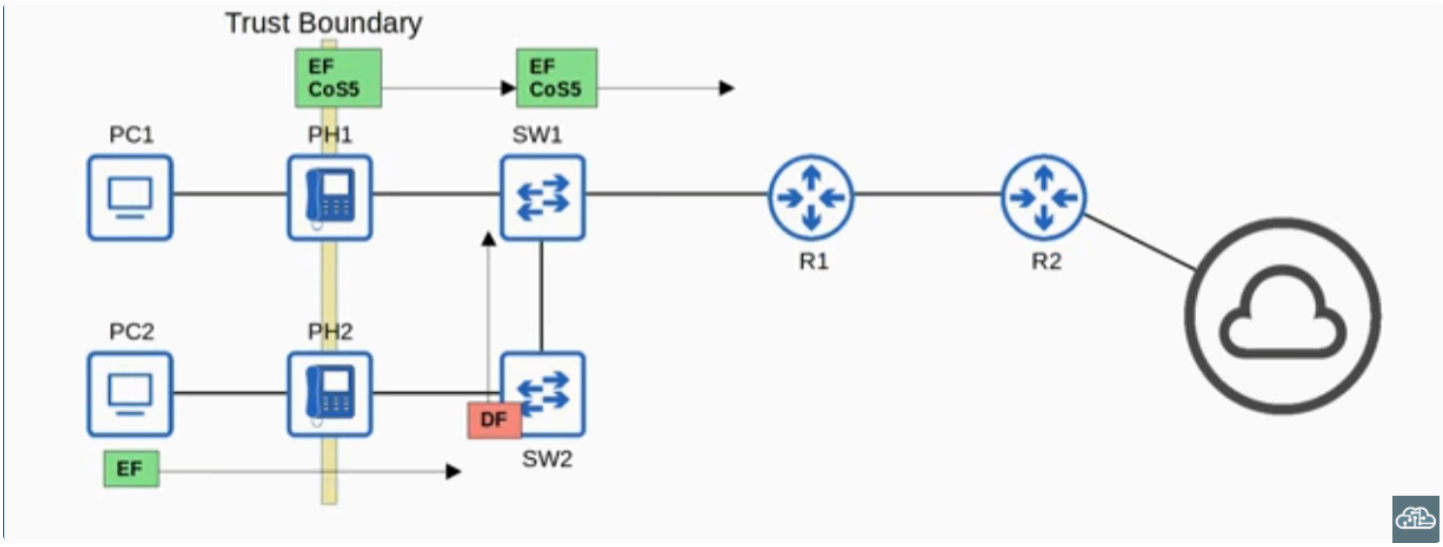

- On the other hand, if **PC2** sends an **EF**-marked packet, the switch should **not trust** that, and for example it might change the **EF marking (DSCP 46)** to **DF (DSCP 0)**.
- **Reason:** We don't want data applications on the PC being treated with the same priority as the voice traffic from the phones.
- **Note:** You don't need to know how to configure trust boundaries for the exam, but just be aware of the **concept of trust boundaries** and how they work.
- This concludes **classification and marking**.

Queuing and Congestion Management

- This topic was already introduced in the previous video, but there's a bit more to be covered for the CCNA.

Review: Queues and Tail Drop

- 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**, and this is called **tail drop**.
- **RED** and **WRED**, which were already introduced, **drop packets early** to avoid tail drop.

Example: FIFO Queue Filling Up

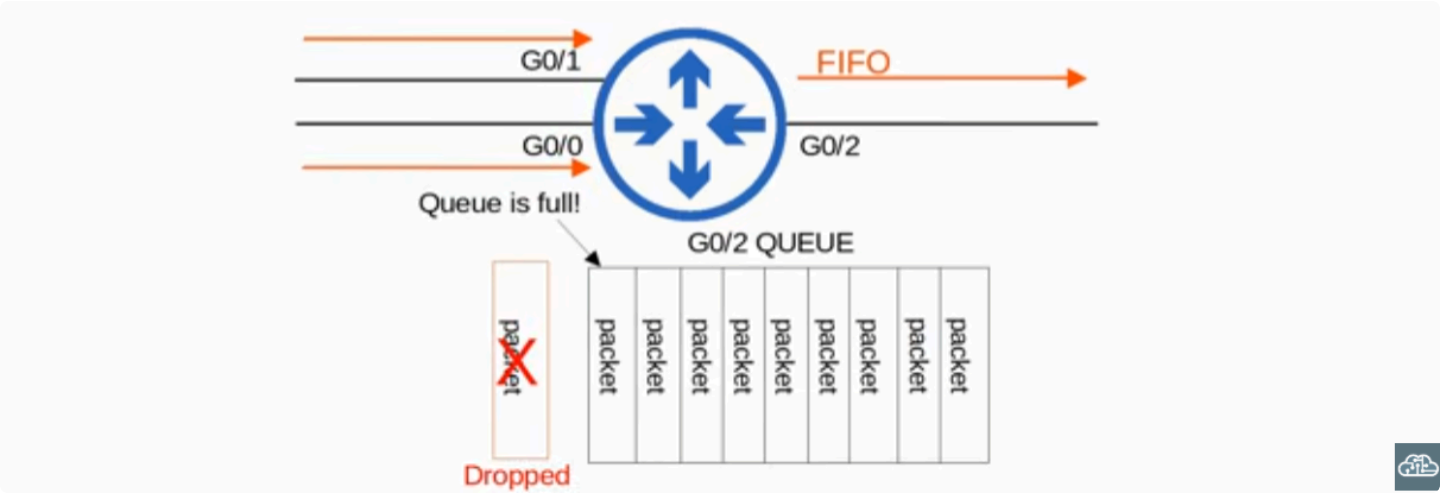

- The router is forwarding packets in a **FIFO manner**, but the queue gets full and packets start getting dropped.
- **RED** or **WRED** could help avoid this by dropping packets earlier.

Multiple Queues and Classification

- An essential part of QoS is the use of **multiple queues**, not just a single queue.
- This is where **classification** really plays a role.
- The device can match traffic based on various factors, for example the **DSCP marking** in the IP header, but also many other things, and then place the traffic 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.

Example: Router with Multiple Queues

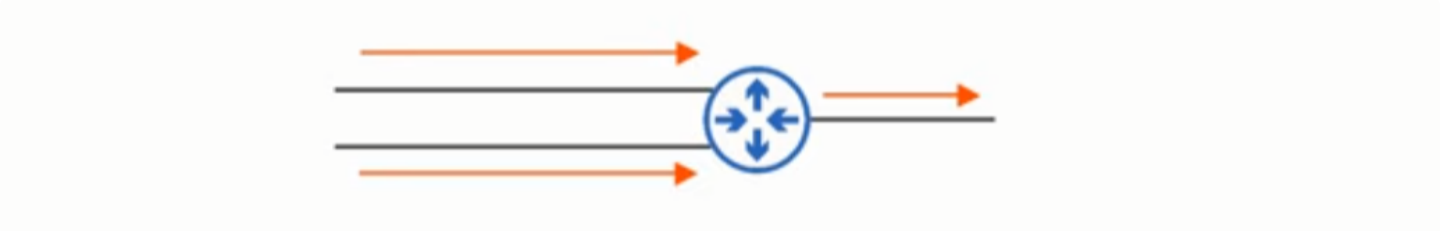

- The router's interface forwarding the traffic has been configured with **multiple queues**.

- Here's how it works:

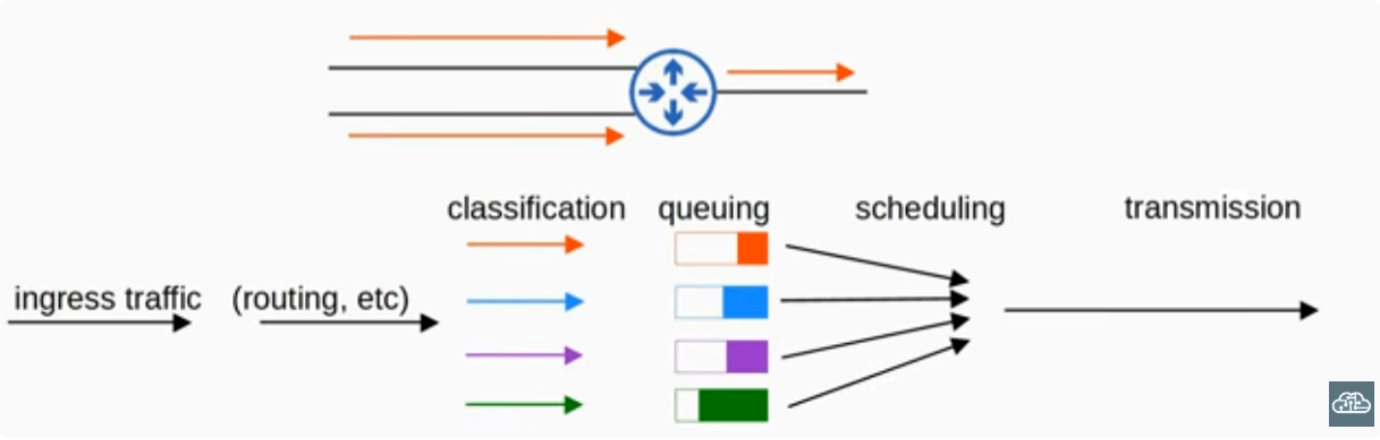

  - **Ingress traffic** is received by the router. Ingress just means **incoming traffic**, traffic entering the router.
  - Then it performs **routing**, meaning it decides which interface to send it out of, as well as other things like **NAT** if necessary.
  - Then it **classifies** the traffic and places it into the **appropriate queue**.
  - In this case there are **four queues** and traffic is classified and placed in a queue depending on, for example, the **DSCP marking**.
  - Then the **scheduler** decides how much traffic to send from each queue, in which order, and the router forwards the traffic, **one packet at a time**.
- The professor notes this is an **oversimplification**, but it's basically how it works.
- After the router decides which interface to forward the packet out of, it is **classified, queued, scheduled, and then transmitted**.

Scheduling Methods

Weighted Round-Robin

- A **common scheduling method** is **weighted round-robin**.
- **Round-robin** means that packets are taken from each queue **in order, cyclically**.
- **Weighted** means that **more data is taken from high priority queues** each time the scheduler reaches that queue, so those queues get **higher output**.

CBWFQ (Class-Based Weighted Fair Queuing)

- **CBWFQ** is a term the professor says you definitely should know.
- **CBWFQ** 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 times of congestion.
- Here's the process again:

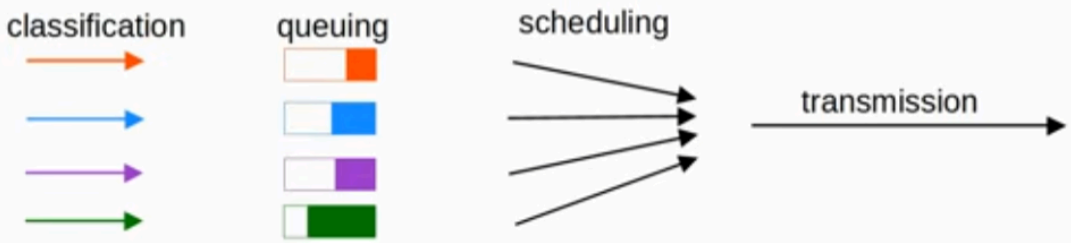

  - Classify the traffic.
  - Place it in queues.
  - Schedule it.
  - Transmit.

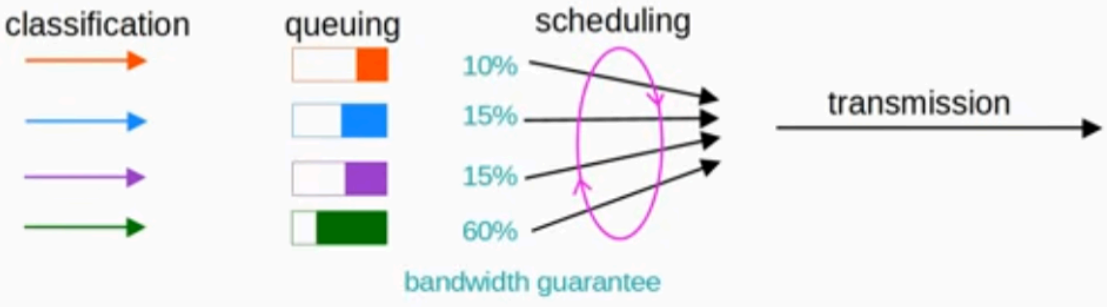

- The device is using a **weighted round-robin scheduler**, sending a **certain amount of traffic from each queue in cycles**.
- On top of that, each queue gets a **guaranteed minimum amount of bandwidth**, even when the queues are congested.

- **Note:** This is a lot more advanced than just a **single output queue**, but it's still **not ideal**.
  - Specifically, it's **not ideal for voice and video traffic**.
  - Even if the traffic receives a **guaranteed minimum amount of bandwidth**, round-robin can add **delay and jitter** because even the **high-priority voice and video queues** have to **wait their turn** in the scheduler.

LLQ (Low Latency Queuing)

- To solve the problem of delay and jitter with round-robin scheduling, we can configure **LLQ (Low Latency Queuing)**.
- **LLQ** designates **one, or more, queues** as **strict priority queues**.
- **Strict priority** 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 and video traffic.
- As soon as traffic enters the priority queue, the scheduler will forward that traffic.

Example: LLQ with a Strict Priority Queue

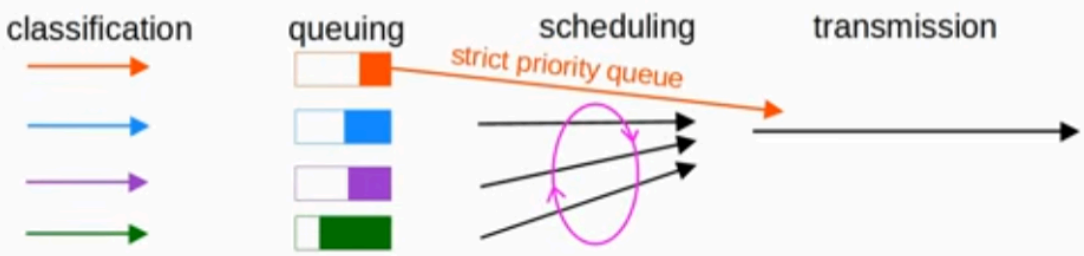

- The top queue is a **strict priority queue**.
- Currently there is traffic in the queue, so the scheduler will **send all of that traffic** before continuing the **weighted round-robin scheduling** of the other queues.
- **Potential problem:** LLQ has the downside of potentially **starving other queues** if there is always traffic in the designated strict priority queue.
  - The other queues might **never get a turn** to send traffic.

- **Note: Policing**, which will be covered in the 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**.

Queuing and Congestion Management — Summary

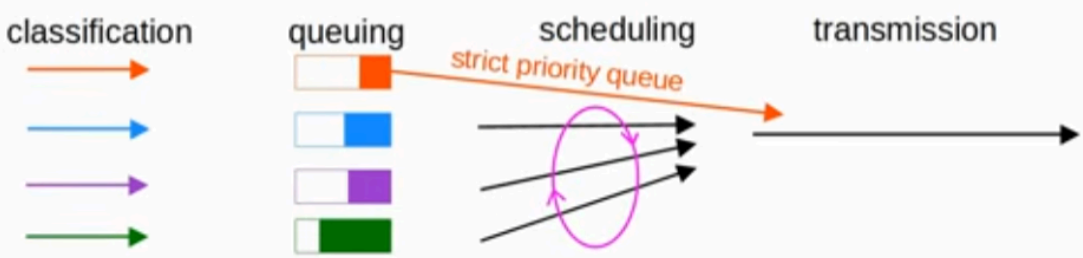

- This section expanded on the idea of queuing introduced in the previous video and examined the use of **multiple queues**.
- A router **classifies traffic**, for example by looking at the **DSCP value**, then places it in the **appropriate queue**.
- The scheduler, for example using **weighted round-robin logic**, forwards packets from each queue in a **cycle**.
- With the addition of **LLQ**, a **strict priority queue** can be used to **immediately forward high priority packets**.
- **Note:** Within each of these queues, congestion prevention tools like **RED** or **WRED** can be used to **avoid tail drop** if that queue becomes full.

Shaping and Policing

- Traffic **shaping** and **policing** are both used to **control the rate of traffic**.
- In the previous examples of queuing and scheduling, the professor assumed that the interfaces are operating at **full capacity**, or **beyond full capacity**, which is why packets need to be queued.
- However, there are situations in which it is **desirable to limit the rate of traffic** to **below the actual maximum capacity** of the link.

Shaping

- **Shaping buffers traffic in a queue** if the traffic rate goes over the **configured rate**.
- This is the same concept demonstrated earlier, however instead of the **actual capacity of the link** being the limiting factor, it's a **maximum traffic rate configured on the link**.

Policing

- **Policing** isn't as polite as shaping. It **drops traffic** if the traffic rate goes over the **configured rate**.
- However, there is some **flexibility**.
  - **Burst traffic** over the configured rate is **allowed for a short period of time**.
  - This accommodates data applications which are **"bursty"** in nature — instead of a constant stream of data, they tend to send data in **bursts**.
- Just like the **policed data rate**, the amount of **burst traffic allowed** is also **configurable**.
- In both cases, **classification** can be used to allow for **different rates for different kinds of traffic**.

Example: Customer and ISP Network

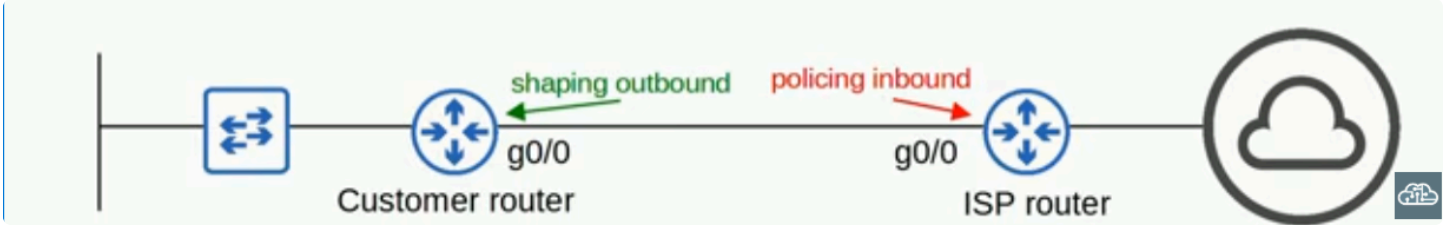

- A **customer router** is connected to an **ISP (Internet Service Provider) router**.
- The customer configures **shaping outbound** on the **G0/0 interface** of the router.
- The ISP configures **policing inbound** on the **G0/0 interface** of the ISP router.

- Reason why this might be done:

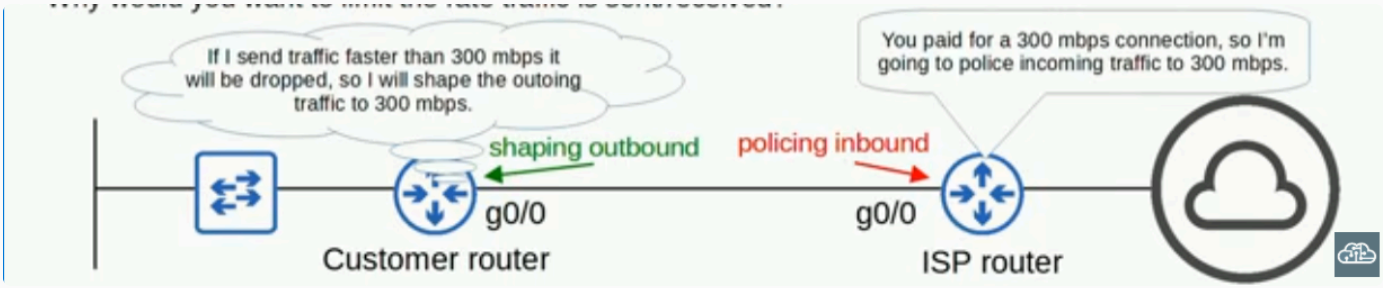

  - Although the physical interfaces are **gigabit ethernet, 1000 megabits per second**, perhaps this customer is paying for a **300 megabit per second** connection.
  - The ISP says: "You paid for a 300 megabit per second connection," so the ISP will **police incoming traffic to 300 megabits per second**.
  - The customer then decides: "If I send traffic faster than 300 megabits per second it will be **dropped by the ISP**," so the customer will **shape the outgoing traffic to 300 megabits per second**.
- There are various possible uses for shaping and policing, but this is a **common use** of these tools.

Review

- That was a lot of material to cover; the professor reviews what was covered before moving on to the quiz.

Classification and Marking

- **Classification** means identifying different kinds of traffic so that you can then treat it with an **appropriate level of priority**.
- **Marking** refers to **setting the values of certain fields** in the **Layer 2 and Layer 3 headers** for use in classification.

- Topics covered:
  - The **CoS field** in the **dot1q tag**.
  - The **ToS byte** in the **IP header**, including **IP Precedence** and **DSCP** (differentiated services code point).
  - The concept of **trust boundaries**.

Queuing and Congestion Management

- Expanded on queuing and congestion management, which was introduced in the previous video.
- Topics covered:
  - The concept of **multiple queues**.
  - **Weighted round-robin scheduling**.
  - **CBWFQ**.
  - **LLQ**.
- **Note:** Marking a packet as **EF**, for example, doesn't actually do anything on its own. You have to configure tools like **CBWFQ** and **LLQ** to make the devices treat those packets as **high priority**.

Shaping and Policing

- **Shaping** and **policing** are both tools to **control the rate** traffic is sent or received.
- Watch until the end of the quiz for a **bonus practice question** from Boson Software's ExSim for CCNA.
