Advanced Computer Networks
Professor Neminath Hubballi
Department of Computer Science Engineering
Indian Institute of Technology, Indore

Lecture 17
Traffic Management - Part 4

(Refer Slide Time: 00:20)

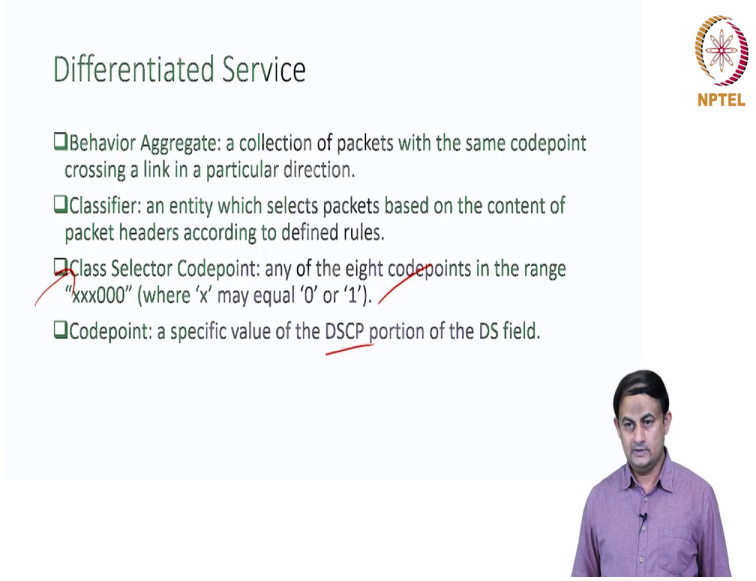

So, in order to understand this differentiated service, we need to define certain terminology, something called the behavior aggregate. So, this is a collection of the packet which receives the common treatment that is identified by code point. So, all these packets, which are sharing a common code point and receiving a common kind of treatment, it is called the behavior aggregate.

And how do you assign this code point to the packets coming from a particular user or the application, there is a classifier sitting at the entry point which actually identifies which application is transmitting using a bunch of parameters and then selects what kind of the code point to be assigned to this, that is called as the classifier, it might use a set of rules similar to that the one we used in the packet classifier.

And then, there is something called the Class Selector. Once you classify the packet, then you assign a code point to that particular packet; differentiated service uses a set of classes out of

those classes; you map the incoming packet to one of these classes. That is the job of the class selector.

And the code point is simply the number that is put into that differentiated service field, and it is also called as the DSCP field. In this particular field, you write a number, that number is called as the code point.

(Refer Slide Time: 01:58)

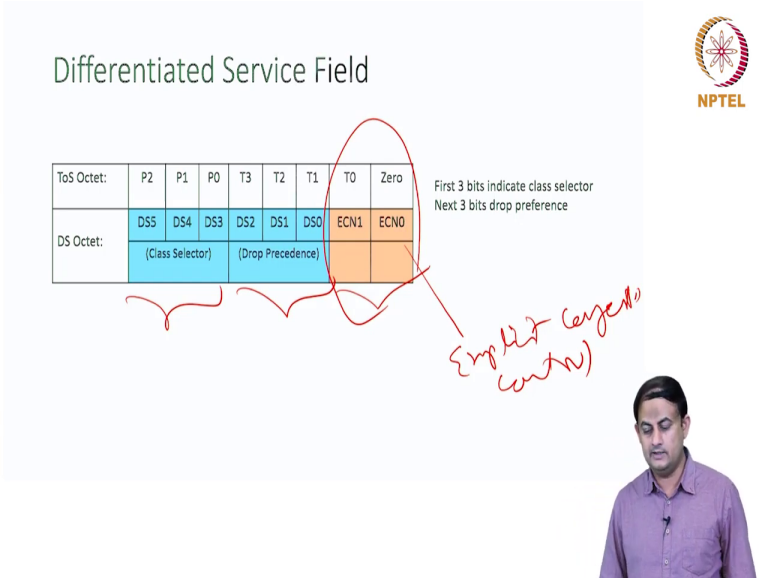

And in the differentiated service model, the 8 bits of the type of service field in the earlier version would take the following meaning: the first 3 bits out of the 8 bits indicate something called the class selector, they identify a particular class of the traffic. And then, the remaining 3 bits actually indicate the precedence or the priority with which the packet is actually dropped.

A combination of these two first selects what class the traffic is; within that class, there are multiple dropping preferences. So, let us say hypothetically, there are two classes; in class one you have got might be 1 2 3 4, kind of the dropping preferences; in class two as well, we have 1 2 3 4 dropping preferences.

So, that is how a combination of these two, first you select the class, and then you select what is the dropping preference. And remaining 2 bits in this out of these 8 bits are actually not used, or if at all they are used, they are used for something called as Explicit Congestion Control. So, ECN stands for Explicit Congestion Notification.

There are two kinds of congestion control. One is the implicit congestion control that the TCP, by default uses, which actually interprets the congestion based on the delay variations and the packet drops that are happening in the network. The second kind of congestion control is something called network congestion control, where the network tells the end host that I am experiencing so and so kind of congestion, you better reduce the transmission rate.

So, by setting the values in these two bits, the network actually or the router precisely tells the end host that you need to slow down the transmission; for that purpose, these two bits are actually used. So, out of the 8 bits, 6 of them are used for assigning the code point, and the remaining two bits are used for explicit congestion control.

(Refer Slide Time: 04:22)

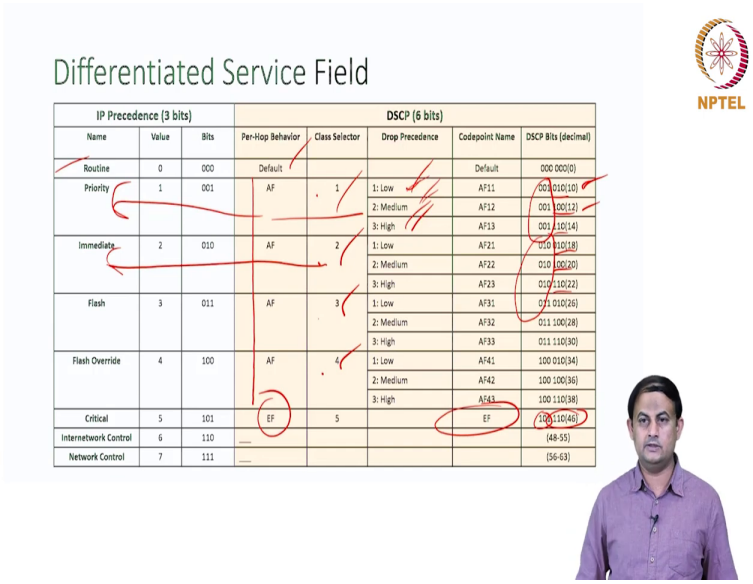

And the way DSCP actually defined this service model or the code points is it actually maps the different kinds of classes to what the IP preference model actually previously defined. So, for example, there is something called the default one, the best-effort service delivery model, there is no change in that one. And that maps to the routine service model or the routine preference we had in the IP precedence model or type of service model. And it defined 4-different types of the classes called as the Assured Forwarding model; AF stands for assured forwarding they are identified by numbers 1, 2, 3, and 4.

So, within that, for the assured forwarding number 1, the class selector roughly maps assured forwarding number 1 to what is defined as the priority in the IP precedence model. And similarly, assured forwarding with number 2 roughly maps to what is defined in the IP precedence as the immediate; you need to transmit that packet immediately. And so forth, the 3 defined compares to the flash, and 4 compares to the flash override. And 5 compares to something called as critical 5 here it is; this EF stands for Expedited Forwarding. So, the default one class, assured forwarding: 4 different classes were identified by numbers 1,2,3, and 4, and the fifth one is something called the Expedited Forwarding model.

So, that is how the classes are actually selected. So, in some sense, this is actually giving you some kind of backward compatibility, the same 3 bits that are used to define the IP precedence model or in the type of service field are mapped to the class selector in this differentiated service model.

And within that, subsequent are the next 3 bits to define the priority of the dropping. So if we were to drop, then the priority of that dropping is low and medium and then high. So, in every class, these three types of dropping priority are defined. So, you assign a binary number to this low, you assign a binary number to medium, and you assign a number to this high.

So, is, what is the point here? So, low is indicated with the number 010, the medium is indicated with 100, and high is indicated with 110. And the class selector here, assured forwarding class number 1, is identified by writing 1 in the first three bits 001. That is the decimal 1, I am talking about. So, you combine these two when you see it in binary. So, 001 followed by 010 indicate that this is the assured forwarding class, this packet actually belongs to the assured forwarding class number 1.

And the priority for dropping this packet is low, so whenever you want to make a choice, you drop that with the low priority, meaning you can actually drop that; if you go to the medium, then with a certain higher priority, you need to forward the packet and you need to drop that with the even lower priority.

When you go to the higher onto drop precedence, then you need to forward this packet with the higher priority and then drop it with the very low priority that is the meaning of that. So, you

treat these two, 3 3 bits independently. So, you might be wondering, what is this 10, 12 why is there no 11 and 13, and so forth? That is because of the way these have been numbered.

So, the service classes always in all the cases is 001. And these three are taking the binary value 010, 100, and 110, which is the drop precedence. So, using which you combine these two independently, that is why it is actually taking this kind of number. So, the same is repeated for the assured forwarding class numbers 2, 3, and 4.

So, low is taking the value 010, the medium is taking the value 100, and high is taking the value 110. So, all three are actually repeated. And the assured forwarding class number 2 is actually identified by binary number 010. And similarly, 011 defines class 3, and 100 indicates the assured forwarding class number 4.

So, that is how it is actually mapped to the IP precedence model. The Expedited Forwarding model is actually the highest priority model that you can think of, so that is actually identified by the number 5, 101. And then, the highest priority within the drop precedence is 110.

Together, these two numbers indicate the decimal equivalent of 46. So, that is what is actually shown here. So, in a nutshell, these values are called as code points, and these many code points are defined. And any user traffic is actually mapped to only these many numbers of the code points.

So, thereby, it actually simplifies the notion at each router. So, the operations at the routers are now simplified; what it needs to do is to first look into this field and extract the 8 bits from the different set of service fields, look at the first 3 bits, and identify the class and the next 3 bits and identify what is the precedence for dropping, using which you make the forwarding decisions.

Router will understand if I have to drop, I need to drop this packet with this particular priority if I have to forward this packet, then I need to forward it, this is expecting so and so, and this belongs to the class EF, if I need to forward it to a very high, priority without much of the delay, and if it is of class 4 then, slightly less priority is given to that and forwarding something like that.

So, that is how you actually simplify the operations at each of the routers, now, routers do not have to remember which source is sending and what it is sending, what application it is, and

what resources it requires. None of these are actually required; it only needs to look into this particular field and then make the appropriate forwarding decision. So, that is what the code point and the differentiated service are all about.

(Refer Slide Time: 11:41)

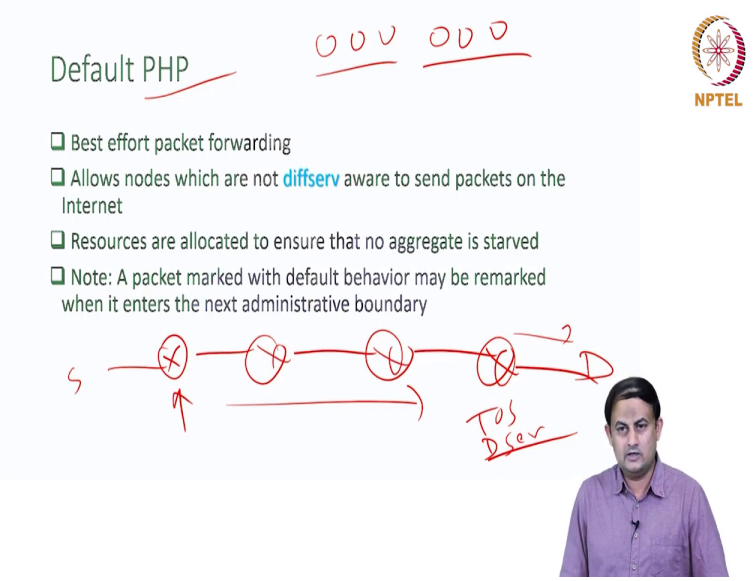

So, as I said, there are three classes: one is the default, which is the default best-effort service delivery model of the IP, and when you set all the fields, the first 3 bits as 0, then that indicates the default. So, default is called as the per-hop behavior; per-hop behavior is how an individual router will behave by looking at the values in this particular field, which is the DSCP field. So, PHB per hop, PHB stands for per-hop behavior, and per-hop behavior is how exactly the router is going to behave.

So, if it is a default per-hop behavior, then if all the 6 bits are set to 0, the class and the next 3 bits within the drop precedence are set to 0, that is the default one. And what it allows you is to see if there is some router along the path, so if my source is here, the packet is going to D, and there are a bunch of intermediate routers. And if let us say the marking is done right here, and it crosses two of them, and then if one of the routers here probably does not understand what is the type of service field or the differentiated service field, then it just can still forward that particular packet because you do not need any differentiation right here when all the bits are actually set to 0. So, that is how it actually maps. So, if there is a router, which also does not understand the type of service field, then also forwarding can be done. And the default strategy just, although it

says that it is the best-effort delivery model, it tries to ensure that none of the packets coming from a particular source are continuously dropped at a particular router.

(Refer Slide Time: 13:53)

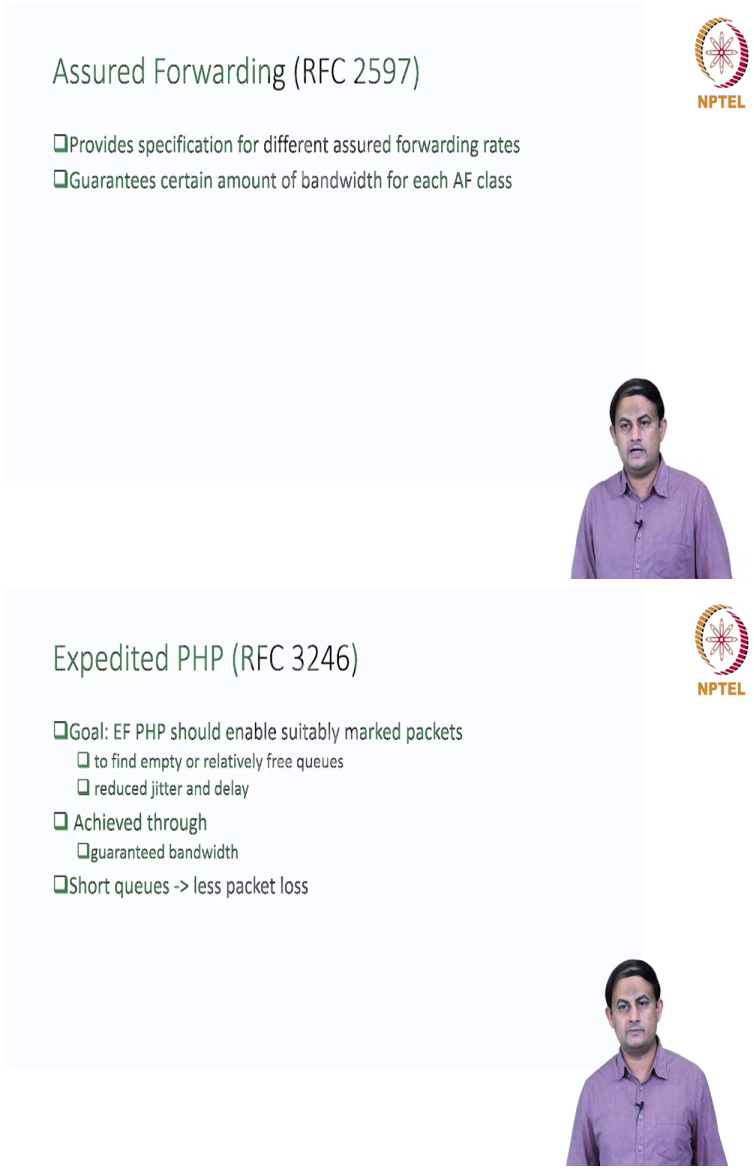

The second class is called as I said is Assured Forwarding. In short, that is denoted as the AF. There are four different classes, and these four classes actually guarantee some amount of the bandwidth and priority for different classes that are defined. And the RFC 2597 is the one that actually standardizes or defines these classes from 1 to 4. Similarly, RFC 3246, the one which defines the expectations or what the network or the router are supposed to do, when a packet with this class is, this code point is actually set, and it actually arrives at the router.

So, basically, the goal is to expedite service actually, as the name suggests, is to forward the packet with a very high priority. And the routers need to point out always when a packet with the expedited forwarding set arrives at the router, you need to forward it with the very minimum jitter and delay and also you need to ensure that the packet always finds an empty queue so that it can be scheduled for transmission immediately as and when it arrives.

And it also expects some amount of the bandwidth to be guaranteed to this class of traffic. So, in short, you want to minimize the queue length. So, shorter the queue length then the chances of the packet getting lost or dropped at this particular router is actually minimal. So, that is what the expedited service code point is actually telling you.

(Refer Slide Time: 15:46)

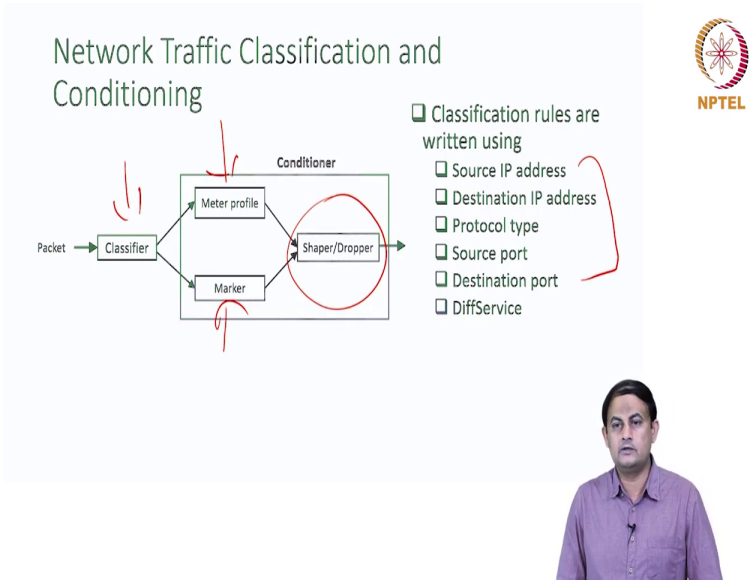

And that is how you have a set of code points, and those code points are put inside the IP packet header by the edge router or first hop router, and then there are subsequent routers within the administrative domain of that network as that of the edge router, which is putting that label or going to make the forwarding decisions or giving the priorities to the packet using that code point. And now, the question is how exactly this is done, how does the edge router actually place this appropriate code point inside the packet. So, there is something called as the classifier, which I was referring to in the earlier discussion. So, the classifier actually identifies using a bunch of parameters, what are those parameters, the source and destination IP address, the protocol type field, and the source and destination port number.

And possibly if it is not the router that is connected to the end-user or it is receiving a packet from a previous network, a network which is under the administrative control of the different unit, you receive a packet from them and then you look at the differentiated service field that is set inside the packet by the previous network.

And then, you decide what code point you want to assign to this particular packet at this level. And then, the marker is the one which actually places that code point, and the meter or the profiler is actually going to measure whether the traffic sent by this particular end host or this particular network is actually meeting the criteria of what is called as traffic conditioning.

So, is it falling under the profile, or in other words, in simple terms, what it means is if you are an end user, so whether the end user is transmitting it at the rate that he or she is supposed to transmit, if it is violating then what to do, so keeping the profile.

Nowadays also, you might see that service providers also often ask that the peak bandwidth detour. I am going to give you 1 Gbps of the transmission rate till you exhaust cumulatively for maybe, let us say, up to 20 GB, and once you exhaust that or consume that 20 GB, then the rate of transmission would be reduced. So, in order to do that kind of differentiation, whether to transmit this particular traffic or packet with that priority or not, the metering is also done at the edge router.

Taking inputs from this one, the next component within this one is called as the shaper or the dropper that decides to either drop the packet or, if it is fluctuating if, let us say, the traffic is not to confined to the profile, then the shaper would probably put them in the queue and make them exit at a constant rate to form to this particular network. So, that is what the whole system actually looks like.

So, I repeat the classifier using a bunch of parameters within the header fields of the IP packet, and the transport layer packet actually identifies what kind of the packet it is, and the metering and the profiler actually look into the historical data transmission and keep track of how much data has been transmitted, the marker is the one which actually places the code point inside the packet, and the shaper and the dropper make the decisions at what rate and when the packet needs to be transmitted? if at all, it is dropped? and when it is to be dropped.

(Refer Slide Time: 19:49)

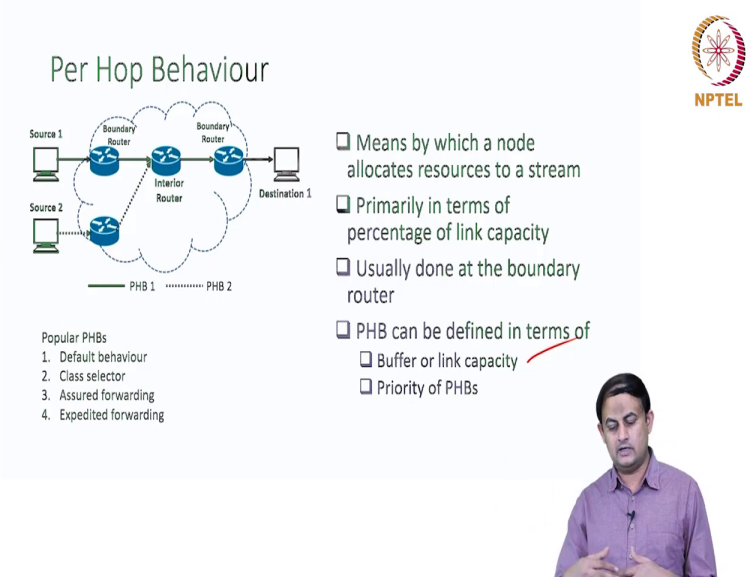

So, this all as I was refering in the differentiated service model, every router at least within the administrative domain of a single administrative entity is referring to the code point written inside the packet and making the forwarding decisions. And this behavior is called as I said that per hop behavior. Per hop behavior is mainly defined, for example, if you have these many classes and then within that the dropping preferences, and for every class of the traffic and every drop preference, router might reserve a certain amount of the bandwidth.

That is how it actually brings the notion of Expedited Forwarding or differentiated forwarding to the network. So, either you define the link capacity, or you commit a certain amount of the buffer space to that particular class of the traffic so that the packets from that particular profile or the class are not actually dropped. So, if you have four or five classes, for each class, you commit a certain amount of the bandwidth and the buffer capacity.

(Refer Slide Time: 21:16)

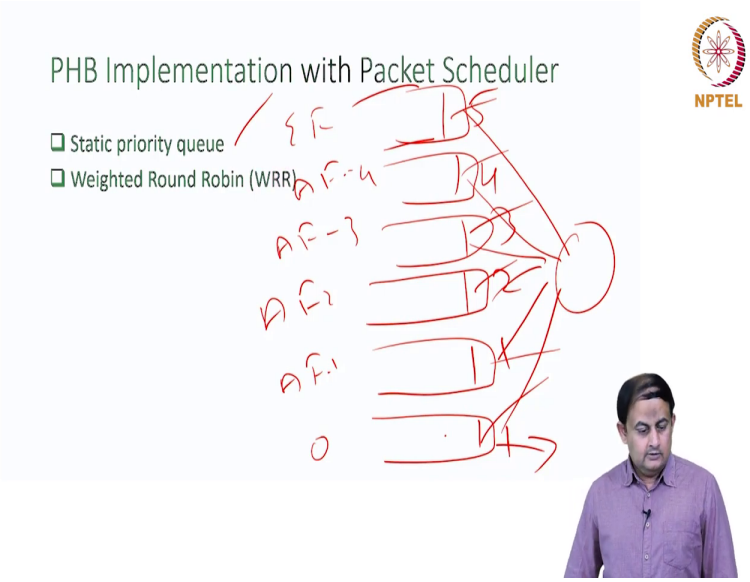

The differentiated service model actually assigns a priority or a code point to every packet that is incoming and typically, the edge router or the first hop router does that and the subsequent routers, at least within the administrative control of the same service provider, those routers are actually looking into the code point written inside the packet and based on that they are making the prioritized forwarding decisions.

So, now the question is, how do we actually bring the prioritized service to these packets? So, one way to think of this is I have got a set of queues, and each queue is holding the packets corresponding to a different class of priority, maybe the default service, which is class 0, those packets are coming and sitting in this particular queue.

The Assured Forwarding class 1 is actually coming and sitting in this queue. Assured Forwarding class 2 is actually coming and sitting in this queue. And assured forwarding class 3 is coming and sitting in this queue. And Assured Forwarding class 4 is coming and sitting in this queue, and then the last class, Expedited Forwarding packets, are coming in this queue. And the scheduler might pick up the packets always from the higher priority queue.

And then, if there are no packets available in the higher priority queue, then it can go and select the packets from the next higher priority queue and so forth. So, the best-effort service packets will get the chance of being forwarded only when there are no packets that are sitting in the higher priority queues.

And it is obvious to note that if there is a continuous arrival of the packets with the higher priority set, for example, if there are enough packets coming with the Expedited Forwarding code point at this queue, then the subsequent other code point or other packets with other priorities will not get their chance to be scheduled for transmission.

So, that is the drawback with static priority, when you assign the priority, you always pick up from the highest priority queue then this is going to be an issue. So, packets in this queue might wait indefinitely they may not find the chance of being forwarded on to the next hop.

So, in order to mitigate this issue, people also talked about something called the weighted round-robin scheduling model where you assign a weight maybe for every five packets that I transmit from the Expedited Forwarding queue, I pick up 4 packets from here, 3 from here, 2 from here, 1 from here and probably 1 from here or with fraction 0.4, a fraction of the packets I am going to forward from this queue, another 0.2 from this one, 0.15 from this queue, 0.1 from this one and 0.5 from this one and whatever is left over is from the other region.

So, something like this, by telling your packet scheduler at what order and what rate they need to pick up the packets from different queues you can bring this notion of the differentiated service. And again, where to put that particular packet inside the queue is defined by that code point that is the inside the packet.

With this, I will stop here. We will come back and look at different kinds of scheduling operations and what kind of schedulers I have in the next class. Thank you.