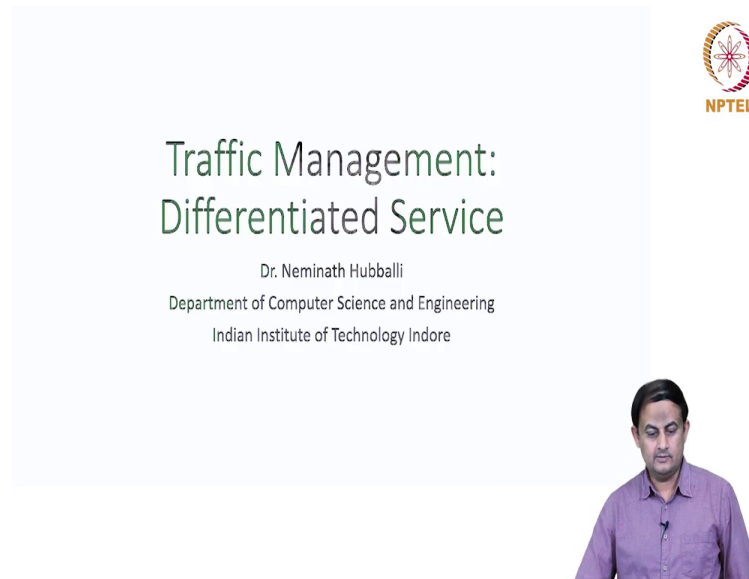


Advanced Computer Networks
Professor Neminath Hubballi
Department of Computer Science Engineering
Indian Institute of Technology, Indore

Lecture 16
Traffic Management - Part 3

(Refer Slide Time: 00:18)



Welcome back in the last lecture, we started our discussion on traffic management. And we looked at one of the service models called as the integrated service model. Today we will discuss about another service model called as the differentiated service model.

(Refer Slide Time: 00:36)

Overview

- ❑ Differentiated service
- ❑ Code points
- ❑ Per hop behavior

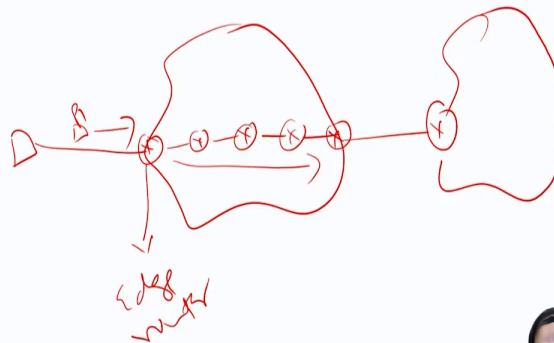


And here is the agenda about what will be discussing in this lecture. So, we will look into the motivation for the differentiated service where exactly the integrated service actually lacks and we will also look into something called as the code points and the per-hop behavior. So, this is the agenda for today's lecture.

(Refer Slide Time: 00:55)

Differentiated Service

- ❑ Per flow based resource reservation puts pressure on the network routers
- ❑ Differentiated service is simpler
- ❑ Much of the needed complexity is placed in the edge routers
- ❑ Operation mode - SLAs
- ❑ Traffic conditioning agreement - a set of rules to realize the SLAs



So, the differentiated service model originated because of the inefficiencies of the integrated service model. What the integrated service model offers you is a kind of service, which is making a reservation of the resources all around the path. So, this means if the source and destination want to engage in conversation, or two endpoints want to engage in the conversation. So the first thing that is there is a signaling protocol called RSVP that runs and makes a reservation on the path at every router. And this reservation actually puts a lot of pressure on the intermediate routers. So they need to remember which endpoints require what kind of bandwidth and other things. And by virtue of that, remembering that itself and then at runtime, and

comparing them would be a challenge. So what the differentiated service model brings you is it simplifies the reservation model, there is nothing called as the reservation.

What it does is it identifies a set of classes of the service model and maps any application requirements into one of these available service models. And then the traffic actually flows following that nomenclature. And what it also does is, by simplifying that, it also puts a lot of the whatever the working needs to be done at the endpoints, so much of the complexity are at the endpoints. So let us take an example.

So let us say this is the network, and the end user is somewhere here. This is the computer that connects to this network via this link. And there is a router here and this router, we call it as the remember, this is not the router at the boundary of the network somewhere which is where the end user's computer is connecting called as the edge router. So edge router or the service router provides services to the end users. And there are a bunch of intermediate routers in this network through which the packet actually flows. And it reaches something called as the border router and then enters another network's territory.

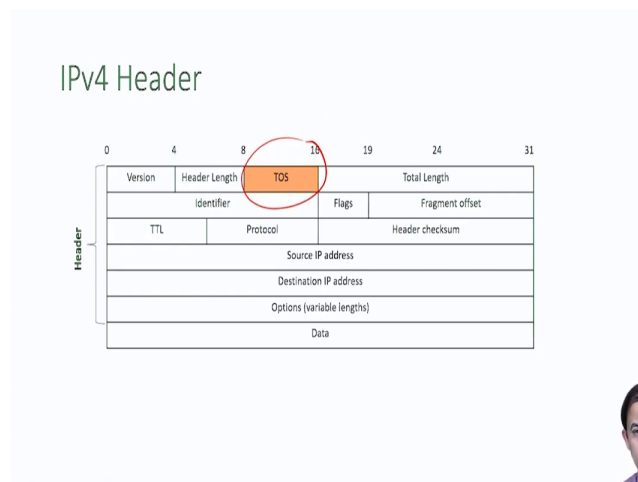
And what the differentiated service tells you is once the packet arrives at the edge router let us call this packet as the P1. When it arrives at the edge router, you mark appropriately this packet. This means depending upon the what type of content is actually going through this packet, or what application actually has sent it, or which customer has sent it based on that, I am going to put something called the code point or some marking is done to this packet using which, the rest of the routers within this network are going to just look at that code information or the marking that is done and then forward the packets. So, forwards the packet with the appropriate priority and when to drop, differentiating when to drop which packet by looking at the code. That is what differentiated service is all about. There is only a handful number of such code points that the differentiated service defines, but using that, the actual routing or the prioritized service or the quality of services actually provides the differentiated services.

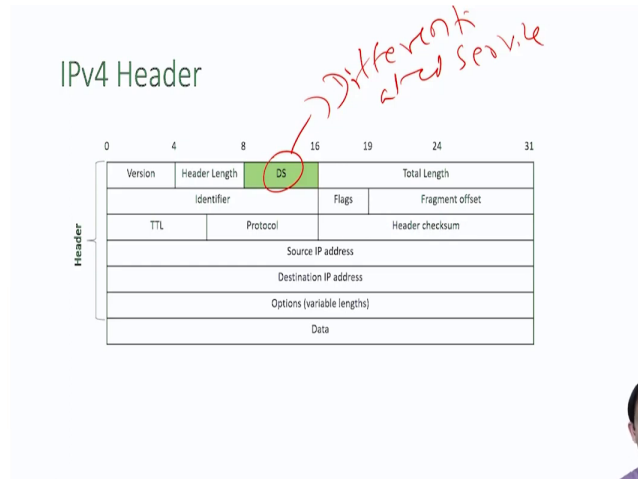
And the way the differentiated service model operates is: it defines a set of service level agreements to the end users, so the service level agreement might say that I am going to provide you, the network service provider might tell you that I am going to forward your packet with a certain priority.

With so much of the fraction of the total traffic that I received, maybe 99 percent of the traffic that you sent to me, I am going to forward it to the correct destination within this timeline. That is what the service level agreement looks like. And in order to realize the service level agreement, there are certain rules defined.

So, the set of rules is called as the traffic conditioning. So, that is an agreement between the end user and the service provider. So, that is how the differentiated service actually operates. So, we will look into the details of how these code points and other differentiation actually look like and how the marking is done.

(Refer Slide Time: 05:50)





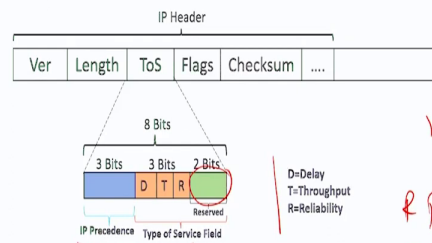
So before that, let us look at the IP version 4 header, and then carefully examine where exactly this marking is actually put. So, here is a picture of the IP version 4 header, and there is one field of interest to us, which is called as the type of service. And often, you might have seen that this field is sometimes marked as the type of service field and sometimes marked as the differentiated service, here it is DS stands for differentiated service.

So this is the same 8-bit one or one octet field from the bit number 8 to 15. So this field is sometimes referred to as the type of service, and sometimes it is referred as the differentiated service. So differentiated service is the one which is in our context, what we are going to be discussing, but differentiated services has a history.

The reason why this is marked as a type of service is that before differentiated service actually came into the picture earlier, the standard was talking about something called as the type of service. So that is why sometimes you see the IP version 4 header have been marked as type of service and sometimes there is a differentiated service, this particular 1 byte data.

(Refer Slide Time: 07:23)

IP Precedence



0 1 2 3 4 5

1981
RFC
791



So this type of service earlier I said that the differentiated service model came later. And prior to that, there was something called a type of service. So the type of service standard was defined to provide something called as the precedence, that precedence is called as the IP precedence, precisely in the year 1981.

RFC number 791 actually defined this standard providing some kind of differentiated service or some kind of quality of service to the different packets. So the type of service field 1 byte was divided into three parts. To be precise, two parts, one first 3 bits from the bit number 0 to 2, 0 1 and 2, define something called the IP precedence that indicates when to drop the packet which packet needs to be dropped if at all, a router has to drop a packet.

And the remaining 5 bits are supposed to be indicating something called as the type of service field. And out of these 5 fields, 2 bits are reserved for future use as RFCs in the 791. So this was not defined. But the other 3 bits, bit numbers precisely 3, 4, and 5, were indicating three different quality of service parameters called as D stands for the delay, T stands for throughput, and R stands for reliability.

You can ask I want so and so delay, I want so and so throughput, I want so and so reliability from the network, and the network will try to provide that kind of service to you. So these 8 bits are divided into 3 parts, the first 3 bits as IP, preference, and the next 3 bits, indicating delay, and throughput, and reliability, and the next 2 bits are actually reserved for future use as for the RFC, 791.

(Refer Slide Time: 09:40)

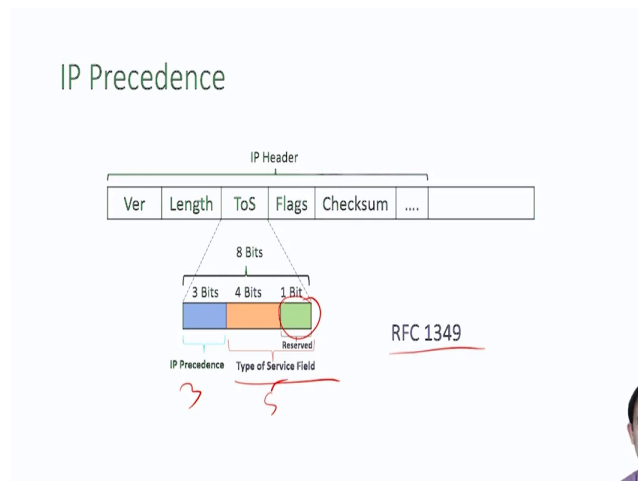
IP Precedence

IP Precedence Decimal	IP Precedence Binary	IP Precedence Name	Delay	Throughput	Reliability
0	000	Routine	0	0	0
1	001	Priority	0	0	1
2	010	Immediate	0	1	0
3	011	Flash	0	1	1
4	100	Flash-Override	1	0	0
5	101	Critical	1	0	1
6	110	Internetwork Control	1	1	0
7	111	Network Control	1	1	1

Bit 3: 0 = Normal Delay, 1 = Low Delay.
Bits 4: 0 = Normal Throughput, 1 = High Throughput.
Bits 5: 0 = Normal Reliability, 1 = High Reliability.
Bit 6-7: Reserved for Future Use.

Default

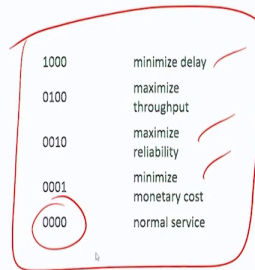
001 791



IP Precedence and Type of Service



IP Precedence Decimal	IP Precedence Binary	IP Precedence Name
0	000	Routine
1	001	Priority
2	010	Immediate
3	011	Flash
4	100	Flash-Override
5	101	Critical
6	110	Internetwork Control
7	111	Network Control



And as we can actually check IP precedence being the 3 bit field. So, there are 8 different values that IP precedence can take from all 0s to all 1s. The first column here indicates the decimal values corresponding to these 8 numbers from 0 to 7, and the name given to this precedence or to all 0s, when all 3 bits set to 0 indicate the routine traffic that is you do not do any differentiation the best effort IP service model using that you forward the packet.

And the next priority level is indicated with the 001. And if that is, if the first 3 bits within that 8 bits are set to 001 then you need to give some level of priority and if the number is 010 or 011, something like that, then the priority level keeps increasing, from this being the best effort service model the routine. So, this is the default service model and from here to here, as you go, the priority level actually increases.

So meaning if at all, you have to drop a packet, you need to drop a packet with the lower priority, preserving the higher priority packets in the queue. So, that is how this IP precedence is actually defined. And this is our story about the first 3 bits and the next 3 bits, as I said, indicate the delay throughput and reliability.

So, these 3 bits are independently defined, meaning when the bit number 3 is set to 0, it indicates that you are not expecting or you are not telling the router or the network to provide any kind of load delay, whatever the default delay with which the packet can be transmitted you just transmit with that if it is set to 1 then you are telling the network or the router that this packet needs to be forwarded with low delay.

And similarly, bits 4 and 5, if bit 4 is 0 then it indicates the normal throughput if it is 1 that you are expecting high throughput. So bit 5 similarly if it is 0 it is normal reliability, and if you want higher reliability, you set it to 1. So as you can see, what is shown in this table is the possible combination of the values that you can specify.

So, by setting appropriately the delay or throughput or reliability component to either 0 or 1. So often although this is possible to ask the RFC 791, so mandates that you do not put or ask more than two of these parameters in one go, this means that if I want the best possible cases, you want a very low delay, very high throughput, and very high reliability, often these requirements are very hard to meet.

So, some of them conflict with each other. So in that sense, so putting all these fields and asking for all these parameters to be met in one go, actually complicates and the network may not be able to meet those requirements. So what it says at max, you set to two of them, either you ask for the low delay, or you ask for the low throughput and not the other one or any two of them.

So in that sense, you are actually not asking too much from the network. So back then, in 1981, and subsequently, the network was not as we see today; the bandwidth was a very scarce resource. So meeting these quality of service requirements or expectations would be very hard in those days. That is why it is actually asked to limit to only two of them.

And that is the story, that is the beginning of or the first case when some kind of the quality of service notion was brought into computer networks, subsequently RFC 1349 or something approximately 10 years later came and modified this type of service field. So the divisions still look the same, you use the 3 bits for the IP precedence and then another 5 bits for the type of service field.

But in the previous case, you were using 3 bits for delay, reliability, and throughput. Now it says that the combination of the 4 bits should be used to indicate different kinds of the type of service, and it reserved only 1 bit for the subsequent or future case. So the changes made in the RFC 1349 look something like this. The IP precedence values still remain the same thing. There are no changes in those 8 values that the first 3 bits define, but what it changes using the next 4 bit combinations you define a different service class.

So, for example, 1000 indicates you need to require a minimum delay from the network that is the expectation. So 0100 indicates you need to maximize the throughput, and so forth. So 0010 indicates maximum reliability, 0001 indicates you consider some of the type cost parameters into account, I want to send this packet with the minimum over header or minimum cost from source to destination. That is the meaning of 0001.

And all 0 set for all the 4 bits indicates the default service, which is the best efforts service model that the IP offers you. So that is how, the changes were made to the IP precedence field. So, although this actually defined the notion of the quality of service in the IP service model, it never really took into the real networks; none of the practical network service providers actually used this IP precedence model or the type of service model in their network and implemented and provided the differentiated service. So, it is more of a theoretical exercise than a real implementation in a real network. So, in 1991, where the RFC 1349 actually came into picture.

(Refer Slide Time: 16:40)

The slide is titled "Service Level Agreements and Traffic Conditioning Agreement". It contains the following bullet points:

- SLA defines what the service provider offers to customers
 - Bandwidth, penalties, contact person, resolution plan, escalation ladder...
- Traffic Conditioning Agreement details how to realize the guarantees described in SLAs
 - Describe traffic profiles, classifier rules applied, policy for discarding and shaping rules

Handwritten in red ink on the slide is "12 bps".

In the top right corner is the NPTEL logo. In the bottom right corner is a video inset of a man in a purple shirt speaking.

So people found that this service model, the IP type of service model, is also inadequate. So it does not provide the kind of differentiated services that I want in the network. So that is where the differentiated service actually came into the picture. So these people started talking about how we actually provide differentiated service over the best-effort IP delivery service model.

So the network service providers started talking about providing the service level agreements to the end users, meaning, as I was referring to in the first slide, the network service providers will come and tell you, I am going to give you so, and so amount of the bandwidth and this bandwidth I am going to commit to you 99 percentage of the time and the maximum downtime that I would experience would be less than 1 percentage or 0.5 percentage of the total time duration.

And I am going to give you end-to-end delay of so and so, something like this. So, the service providers started engaging the customers by formally defining the service level agreements. So meaning, so by specifying all of them, the service level agreements also specified something called as the penalties if I do not meet these requirements, what is specified inside the service level agreement, and what kind of penalties is the service provider going to take?

So, if there is some issue with the service that they provided, then whom to contact? So, if there is an issue, then how exactly do they resolve that issue? And what is the escalation matter? if the first person who is the point of contact is not able to resolve the issue, how or whom to contact, and how the escalation matter actually works? So, all these details were actually put into this service-level agreement.

And so, this service level agreement is more of a nontechnical specification. And in order to realize these nontechnical requirements into reality, there is something called as traffic conditioning that is defined. So traffic conditioning is, so I am going to give you this kind of service model, this kind of service, and in order to realize that service level agreements, so, the end user traffic need to be conforming to something called as the profile.

So for example, if you transmit less than 1 Gbps of the data, then, then only these service level agreements are going to be met as long as the traffic rate or the peak rate is less than 1 Gbps. I am going to give you a bandwidth of so and so, I am going to give you an end-to-end delay of this much, and so forth.

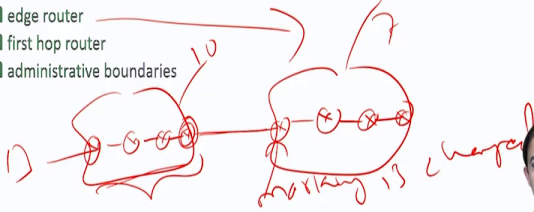
So if you violate that 1 Gbps, that is the profile 1 Gbps, then there are consequences for the end user, the service provider is not going to guarantee you the same kind of reliability, the same kind of end-to-end delay for that kind of traffic. So, that is what the traffic conditioning agreement is. So, you define what is the traffic profile.


So, if the end-user traffic profile is meeting, what is defined in that profile, then the service level agreement is going to be enforced, and then you can actually expect that kind of service from the service provider.

(Refer Slide Time: 20:28)

Differentiated Service

- Its a framework and building block to enable deployment of scalable service discrimination in the Internet.
- Implemented with configuring parameters in the forwarding path.
- Packet marking is done at
 - edge router
 - first hop router
 - administrative boundaries





So, using that model of the service level agreement, the differentiated service actually took up from the base, the real implementations actually came into the service provider networks. So, the differentiated service actually defined a kind of a framework or a kind of building block with which you can actually realize this kind of differentiated service or the quality of service in the internet.

So, it actually provided a kind of configurable parameters using which the traffic from the different applications and different users are treated differently. So, as I said, much of this complexity is placed at the edge router or the first hop router, where the packet marking is done, and then subsequently, in the network, the routers are looking at that particular marking and then doing the forwarding differentiation, meeting the quality of service requirements of that particular packet.

So, what can often happen is although this is defined at the first hop router, what can happen is, this is the first hop router, and then, as I said, there are a bunch of intermediate routers, and then you go to the next level of the network. So, from one network to another network, you, when the

packet transits, when I say this the edge router, the marking is done that marking is actually valid only in this particular network. So, for example, it is possible that the kind of the quality of service the way it interprets the quality of service in this network is different than the way this service provider or the network actually interprets it.

So, for example, hypothetically, there are 10 different classes of the quality of service type of the service, but here, there may not be 10 there, there might be only 7 of them. So, I need to find a mapping of these 10 quality of service parameters or what we call it in the differentiated service model as the code points to the mapping to these 7 of them, which one of this one maps to which one of the different types of the code points that we have in the second model.

So, by virtue of this what I want to say is, so, when the marking is done, that marking is valid in this particular network. When it hits the boundary of this network and enters into another network. The same packet might be remarked at the edge router or the boundary router of the second network. So here, the marking might be changed, marking is changed or not necessarily changed.

If the marking is same in the first network and second network, then no remarking is required if the number of the quality of service parameters is different, then you require the remarking at the boundary of the network. So the marking is valid only within the boundary or the administrative boundary of the one network. So you need to find out a mapping between the first level marking i.e., the marking done in the first level to the second level when it crosses another network.