Advanced Computer Networks
Professor Doctor Sameer Kulkarni
Department of Computer Science Engineering
Indian Institute of Technology, Gandhinagar
Lecture 30
Data and Control Plane Separation

(Refer Slide Time: 00:17)

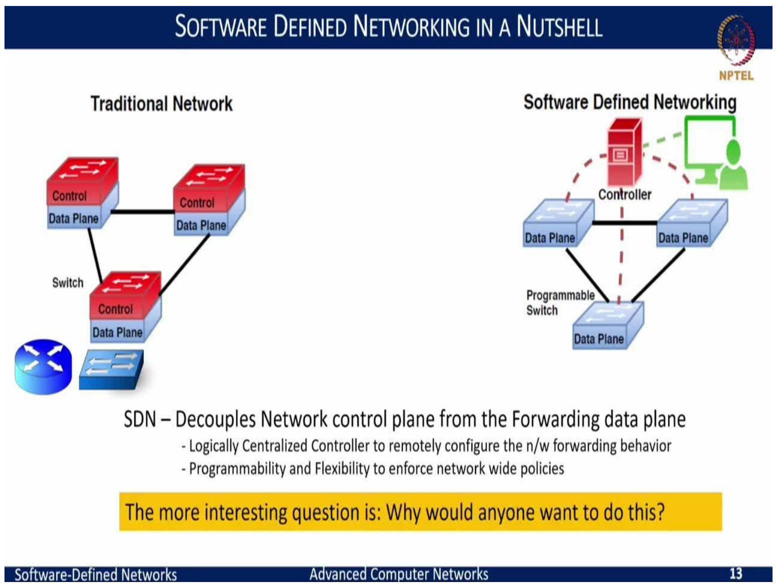

So, the most interesting question would be to say, why would we even want to do this decoupling of control from the data and centralize this control outside? And to understand this, we need to see how the traditional networks operated with the control and data being coupled. What were the challenges, and where this approach of decoupling helps?

(Refer Slide Time: 00:47)

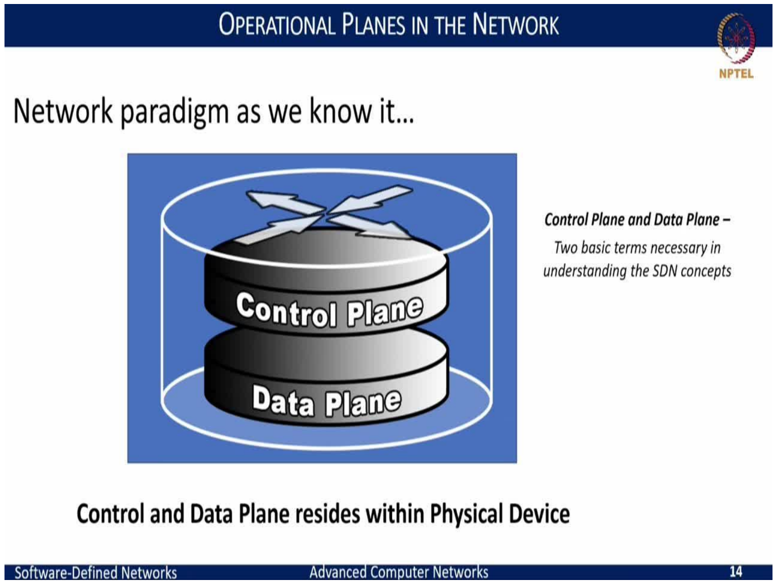

So, let us try to look at the traditional networking devices in a much more detailed manner. For most, like what I am showing you here, is the traditional networking router, where it has the two plates embedded in it one is the data plane and the other the control plane, and the understanding of these two necessary planes or the terms is more vital in understanding their operations and what impact they have over each other. And whether this coupling is necessary or not. But the takeaway here is that the control and data plane, they reside within the same physical device, be it the switches or routers, as I have shown them here.

(Refer Slide Time: 01:35)

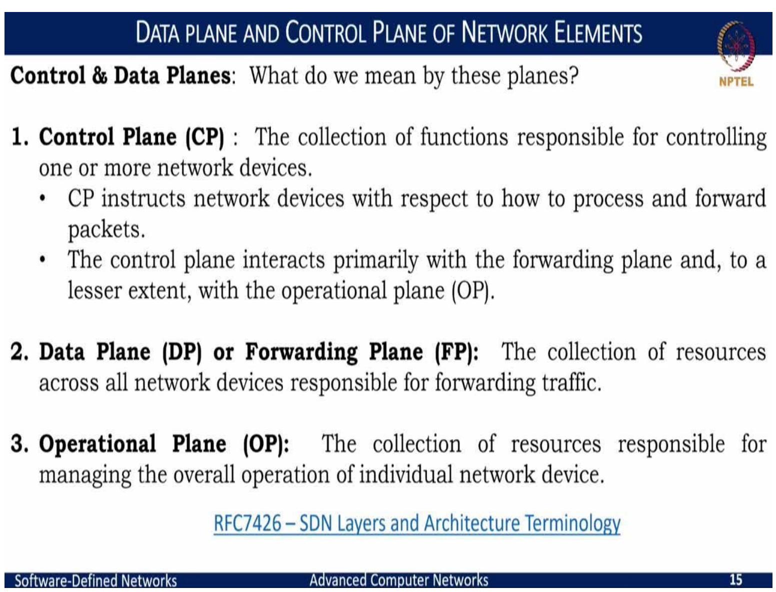

So, what do we really mean by these control and data plane? So, for this, if we refer to RFC 7426, which put forth the SDN layers and architecture terminology, although there are many variants of definitions that are possible, I feel this would give us the right perspective of defining and demarcating the control and data plane.

So, by the control plane, what we really mean is the collection of functions that are responsible for controlling one or more network devices. And when we say one or more network devices, we are referring traditionally to the switches and routers. In fact, this control plane, when we say that we want the functions that are responsible for controlling it means that we are having this plane that will instruct the network devices, that is, the switches and routers, in terms of how to do specific things or how to do a certain action like processing and forwarding of packets as an instruction to these devices. And that is what we mean by control planes.

And primarily, when we say this as a control plane, it is trying to set certain aspects, which would then affect the way the device would process the packets and forward them along. And it would also have an impact in terms of how the device would operate. And this is where this separation of what the device would do with respect to the incoming packets, how it would process being split as a data plane or also termed as a forwarding plane functionality. And by this, what we mean is the collection of resources across all network devices that are responsible for forwarding the traffic that is strictly tied with respect to saying, given a packet that arrives, where should I send it?

In the sense of as a packet comes, what should I do with it and that actions that are taken on a per packet basis is what defines the forwarding plane or a data plane of a device. While the operational plane, although at the level of each device, is responsible for managing the overall operations of the individual network devices that is, you want to enable specific monitoring, filtering of the packets data, etc. that you want to send in terms of how the device would operate, I want to operate my router in OSPF, or I want to operate my router with IS-IS, all of these become basically the modes of operation for a given device.

And once that is set, the control plane is going to dictate how the actual routing as functionality needs to be carried out. And then that data plane would look into this and say what are the means for it to forward a packet given that incoming packet with a certain signature receive where it would send and you have looked in terms of all how the routers work.

(Refer Slide Time: 04:55)

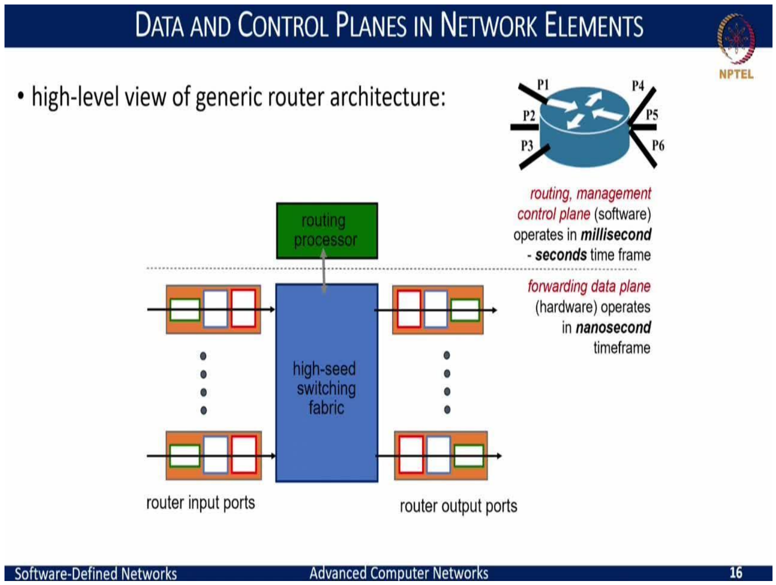

So, let us start from there and see how these data and control planes map. We can see in the figure here that we have a router with 6 ports, and what this really reflects back is basically you have 6 of the ports, which are both inputs as well as outputs. And they are connected through a high-speed switching fabric, and you have a router processor that is also augmenting to say the decisions that need to be taken. And each of these switching fabrics that we are showing is also termed as a line card for each of the ports like P1; we have the port here with a line card and P2, P3 likewise, and for P4, P5, and P6, we will have these line cards.

And here, when we see the functionality of as the packets arrive on each of these line cards distinctly, first it is to terminate the line that is the physical bits that are being accumulated, and the green box here emphasizes on the line termination. And once you terminate, you prepare a frame, the datalink layer 2 frames that is done by the blue box, and then you see the packet, queue it up and see what route you need to apply so that you decide which port the packet needs to go out on or what is the action that needs to be taken on a given packet. And this is carried out through the aspect of high-speed switching fabric to move data from one port to the other. And this constitutes exactly what we call as a forwarding plane or a data plane functionality. And remember, these operations happen for every incoming packet. And hence like, if we are talking about gigabit lines or multi-hundred several hundred-gigabit lines, we are talking about a timeframe of every nanosecond you are processing a specific packet on these data links.

While on the top, what we are seeing is a routing processor that is involved in setting up the key aspects that are needed for the forwarding plane to operate, that is, when I want to run this route, should I be running the link state routing algorithm, should I be running OSPF mode or distance vector mode of routing algorithm, all of those have to be decided, and those are the aspects that play how a route will be decided for a given incoming packet and then which port needs to be taken at the data plane.

And hence, all the computation of the routing, the management of the device in terms of what routing algorithms to do, what ports that we need to be enabling disabling all of these basically constitute what we call as the routing and management part of a router, which we can club as a control plane.

In fact, you can think of the control plane as software that dictates how the underlying hardware, the line cards, and the switching fabric to operate in forwarding the packets. And note here the timeline at which these routing processes need to operate is much more coarse-grained, that is, a few milliseconds to a few seconds at a time and only when there is a state change of a link or when there is a state update from a neighbor, you would want to update your routes, recompute your routes and update. And hence, there is a distinct requirement of the timescale of operation that they have and how they would work. And this is where we see this as a separation of data and control planes.

(Refer Slide Time: 08:36)

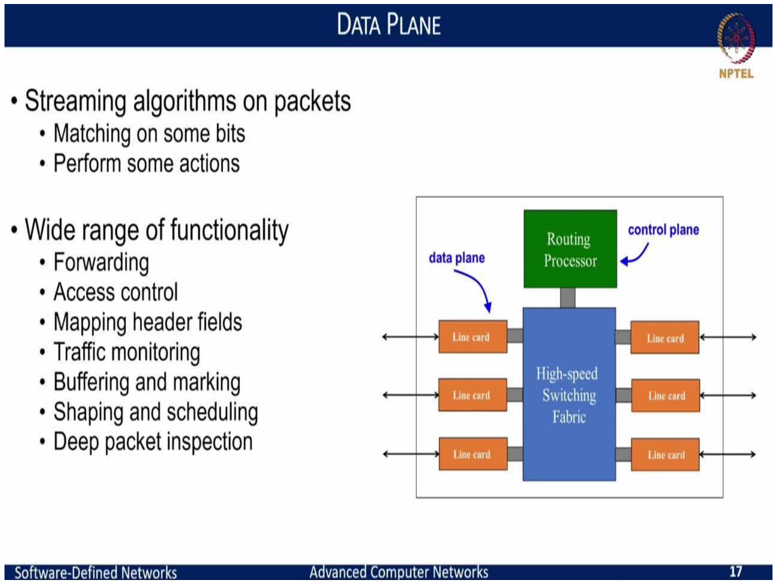

So, to look at this in a much more elaborate way on a data plane, data planes are actually the ones essential for routing the packets within a fabric from one port to the other. And the means to do this is what we call as a forwarding functionality. The packets you receive on one port will be forwarded out onto the other port. And in essence, they would also be doing several of the aspects with respect to the incoming traffic or data, be it the access control, be it the counters of in terms of how many packets were received for a particular source or how many packets were destined to a particular source. And when the packets arrive, you want to buffer and mark like you will see in the congestion control of ECN or in scenarios where you want to either take a decision to drop the packet when to schedule a particular packet. And when it comes to the security features, you would want to inspect the content of the packet either at the header level or at the header and both the data, so the deep Packet inspections could be the functionalities. So, these are all the different functions when we speak of what a data plane would do. So, in a nutshell, they would match on specific bits, perform certain actions, and forward the packets out onto a specific port.

(Refer Slide Time: 10:03)

DATA PLANE VS CONTROL PLANE

Processing Plane	Data Plane	Control Plane
Operation Time-scale	Packet (nanosec) Billions/Millions of packets per second.	Event (millisec to sec) Few Thousands of packets per second.
Tasks/operations performed	**Forwarding**, buffering, filtering, scheduling, QoS Marking, Policing, Classification etc.	**Routing** (OSPF, BGP, IS-IS), AAA, SYSLOG, SNMP circuit set-up, etc.
Location – where do they run?	Line-card hardware (dedicated HW ASICs)	Router software (switch CPU/router processor)

Software-Defined Networks Advanced Computer Networks 18

While the control plane is trying to operate in a different fashion. So, if we see any burst of packets that arrive, the data plane has to process them. So, the data plane could operate on a nanosecond timescale with expectations of billions or millions of packets per second are served through each of these line cards, where the control planes are more or less event triggered, where any update from the neighboring device or updates or configuration changes on the device to trigger for certain activities on the routers or switches.

And this might mean that we may have a few thousands of packets per second to process at best. The key operations in the data plane is primarily forwarding, but also includes buffering of the packets when you have to queue the burst of packets that come; you may have to filter specific packets if access controls are being set in terms of what needs to be allowed to pass or what packets signatures are not allowed. And then, when you have the data on a multiple of the line cards, how do you schedule these packets, and what should be the quality of service parameters that you need to honor while trying to schedule the incoming packets, marking of the packets in terms of whether there is congestion perceived and any information that you want to relate back to the end hosts. All of these become the functionalities at the data plane.

While at the control plane, it is on a coarse grain operation where you would want to see how you would want to go to routing; it could be OSPF that you would want to apply or BGP or IS-IS, which of the routing algorithms need to be run to get the network, understand the network topology and ensure that you are able to forward the packets towards the right destination. This also includes like if you want to check the management aspects of the device in terms of the status to the behavior of the current device through SNMP or extract the operational logs in terms of the Syslogs or facilitate the authentication, authorization, and access control aspects for a particular device. All of these can be done on the control plane. And what this also means is the data plane typically operates on the line card hardware that is typically with dedicated hardware ASICs. While the control plane is primarily router software that is ingrained into the switch and would run on a switch CPU or a router processor.

And notably, this data plane is an aspect that is unique to a particular device in isolation. But on the other end the control plane, when we speak of routing, and you have put multiple routers in a network, they have to work together I cannot have one router running OSPF while the other running IS-IS or link state in some other fashion they all have to be coherent. And this is where again, what it means in a network or a topology of a network, although the data plane and control plane are embedded in a single device, the data plane is primarily distinct for operations in terms of each device, but some aspects not all, but some aspects of the control plane have to be common across different network elements in a given topology, in a given domain.

(Refer Slide Time: 13:31)

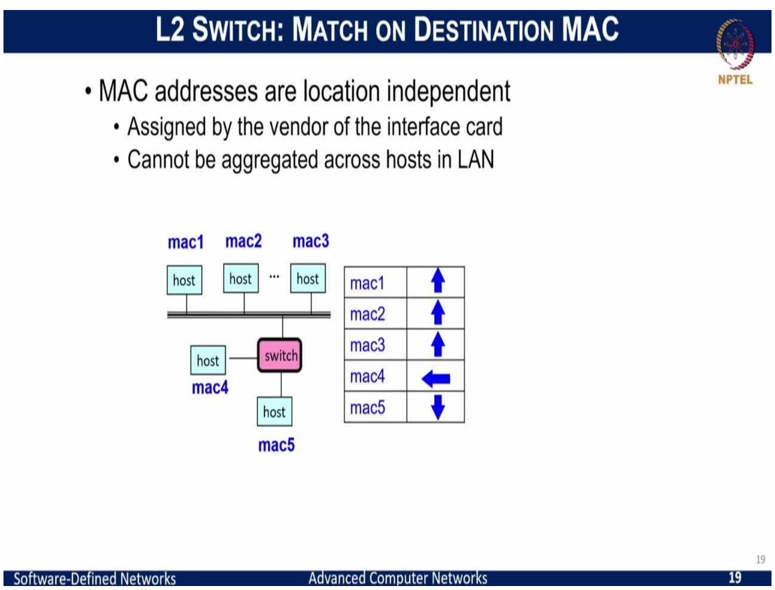

So, let us look at what we mean by the operation of a switch in a given network and typically, these switches operate in a LAN, and what they allow is for multiple hosts to communicate within a given LAN. So, if we look at the diagram here, if we have 5 different hosts, a switch would facilitate forwarding the packets and enable communication between each of these hosts.

In doing so, it will learn and remember what or where to route a particular packet given the destination Mac matches one of these Mac 1 to Mac 5. So, as a very simple illustration here, the switch would basically forward upwards for any of the destination Macs that are set as Mac 1, 2, or 3, and it would send it to its towards down for the Mac 5 and to its left for the mac 4, this way the switch has a control plane which tells where to route a particular packet when you see a particular destination mac address and it has fundamentally a data plane which would look this table and make a decision that, okay I got a packet with a matching destination MAC address of mac 1 then I would send it upwards.

Alright, this is how basically the Macs operate. But note that the MAC address is somewhat vendor-defined, although you could change them not that it is fixed and you cannot change, but they do not specify anything about the location, but within a LAN, a switch would make this location as a part of which port of the switch that it is connected to and accordingly maintained this information to forward the packets. So, at the data plane, it is MAC address to the port address mapping is what is more relevant.

(Refer Slide Time: 15:34)

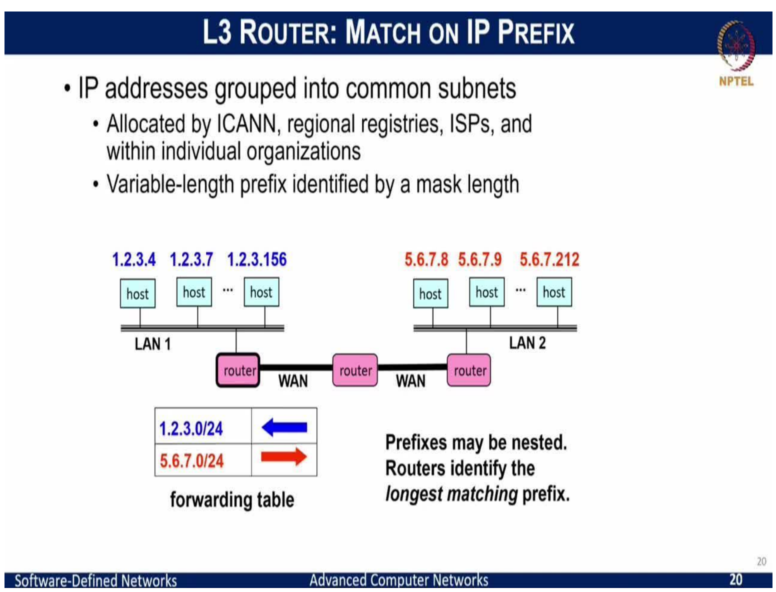

Likewise, when we look at the routers, which are layer 3 devices, you have seen this earlier in earlier lectures as well; they match on the IP prefix, that is, the destination IP address. And once they look up the parts of destination IP addresses, you can aggregate specific bits of IP addresses and see which way to route, and here, the routing table would say that LAN 1 corresponds to anything with the IP address 1.2.3.0 with 1.2.3 being fixed, in any of the addresses from 0 to 255 on the last bit would be allowed to be routed towards LAN 1 and likewise for 5, 6, 7, 8 towards the right where you would see that any address with 5.6.7, a decision at a router would be made to send them towards the right.

And here, you can see that the prefix is in the way in the state, and routers identify where to forward the packets based on these longest prefix matches which you already covered in the earlier parts of the course. So, in terms of when we look at the router, the forwarding table dictates the data plane functionality, and the means to populate that forwarding table is being dictated by the control plane functionality of what routers 1, 2, and 3 are together building routing algorithm based on what routing algorithm they are trying to use would eventually dictate what forwarding table would be set up for each of these routers.

(Refer Slide Time: 17:23)

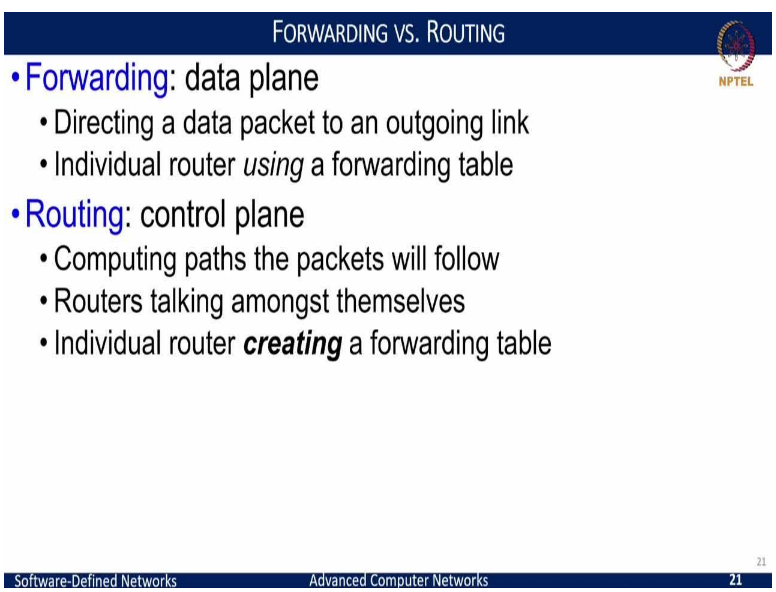

So, in a nutshell, when we look at the data plane and the control plane elements, think of forwarding as a means or a main function of a data plane that tells where to direct a given data packet, which outgoing link to choose, be it a switch or a router while one does it on an L2 destination MAC while the other does on L3 IP address. Whereas the control plane is primarily about routing, that is, how to compute these parts so that the forwarding plane can decide where to send the packet. And in doing so, the routers do talk amongst themselves and enable to create the route while the switches learn based on the packets that they receive and understand although they may not coordinate with each other, they would still be able to learn based on the traffic that is moving in and out. Hence the data plane and control plane are clearly distinct entities.

(Refer Slide Time: 18:29)

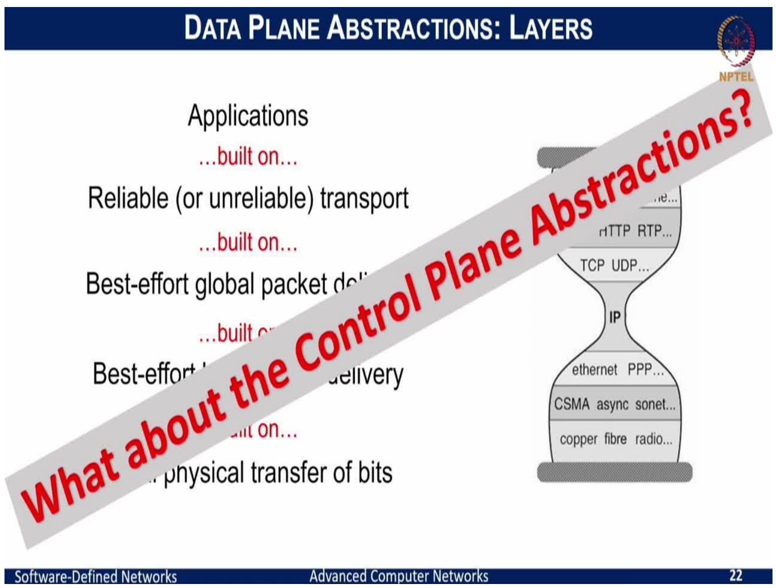

And when we look at these operations in the data plane that is basically forwarding the packet from one end to the other, we have seen the layering principle that puts forward a nice set of abstractions wherein we set at the end host it would operate all the way from layer 1 to layer 7, while the switches operate on layer 1 and layer 2, the routers operate on layer 1, layer 2 and layer 3 and so on. And this enabled us to build the right set of abstractions where applications could to be built on a reliable or unreliable transport ruling on top of the IP, which is our best effort, global packet delivery, and so on.

So, the abstractions were neatly defined for the data plane, which confined the roles of the networking elements to either layer 2 or layer 3 operations. But then, when we speak about the control plane operation that is doing the routing. Now, do we see any such abstractions in here that is the main question how would, if any, be it needed to do the control plane operations? So, first is do we need the abstractions? Second, if we need what are they, and how should they be brought up? Let us try to look at it in detail in the next lecture.