Advanced Computer Networks
Professor Sameer Kulkarni
Department of Computer Science Engineering
Indian Institute of Technology Gandhinagar
Lecture 38
Introduction to Network Function Virtualization - II

(Refer Slide Time: 0:21)

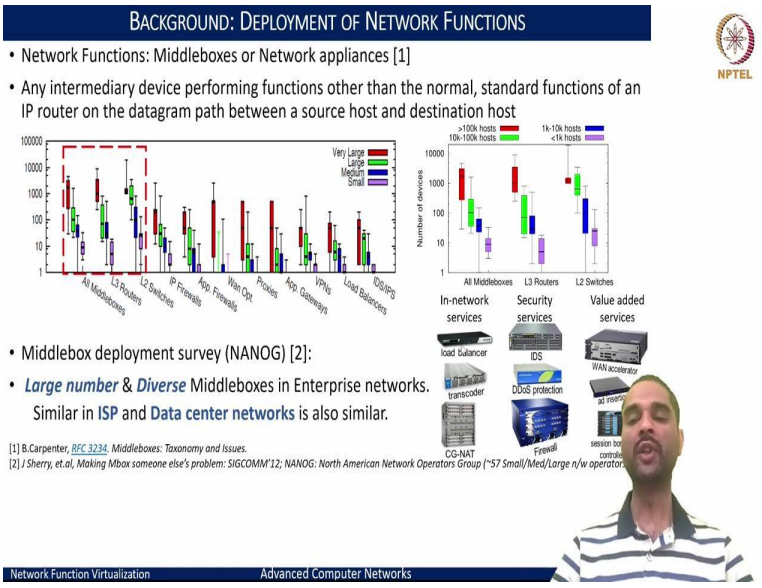

In essence, these middleboxes could be employed to provide various services. For example, the traditional network services like the carrier-grade NAT, wherein we could be using these devices to translate the private network address to a public IP and public to the private vice versa.

So, that we are able to facilitate connecting many of the users within and align to the network services using a single public IP address. Or consider a load balancer, which would be used to ensure that any incoming traffic for the dedicated services that would be running on our network is balanced with respect to the load for each of the different instances of the network servers that we could be running.

And likewise, they could also be used to facilitate and provide security services. For example, firewalls, the most common re-employed middleboxes, the types of firewalls could range from being simple stateless firewalls to the rule-based firewalls that can operate at layer 3, layer 4, or even up to the application layer firewalls or the next-generation firewalls that we think of.

And when we speak of other kinds of security devices, we typically think of intrusion detection or intrusion prevention systems using the network intrusion, prevention, and detection systems as the middleboxes that could be deployed within the network.

And again, we could also have deep packet classifiers, wherever whenever we suspect any malicious content and want to analyze packet-by-packet contents, we could have the DPI engines that could be running within the network. And also, to safeguard from the DDoS attacks that typically happen on the services that are being run on an enterprise campus or data center networks, we would want to build specific purpose-built middle network appliances like DDoS protection and so on.

Besides, they could also be used to provide specific value-added services for the services that are being hosted within a network. For example, in an enterprise network or campus networks, we could have WAN accelerators that could either be providing the cache of the contents that we typically access we could also have them compress and decompress the packets before transmission.

And if we think of the telecommunication or ISP networks, we could even have ad insertion modules within there, including data centers, where when you try to search for Google, the Google data center would have ad insertion devices, which would insert the ads that you would see on your pages.

When we consider the telecommunication networks, we could have these SBCs, or the session border controls that typically manage and handle the wipe-specific services over the telecom networks. So likewise, we can see that these middleboxes, although like when we said a single firewall, they may come in different types and shapes. So, the amount of services kind of service that you want, all would require that these are specific purpose-built hardware for specific network function.

And when we look at all of this, the takeaway here is that there is a significant usage of the middlebox that is it highlights the dominance of the middleboxes. And it is not that these are essential devices that came in; these are also playing a very significant role in different networks, be it the campus or the enterprise networks, or be it the data centers, or Telecommunication or internet service provider networks. These middleboxes play a crucial role in each of them.

(Refer Slide Time: 4:21)

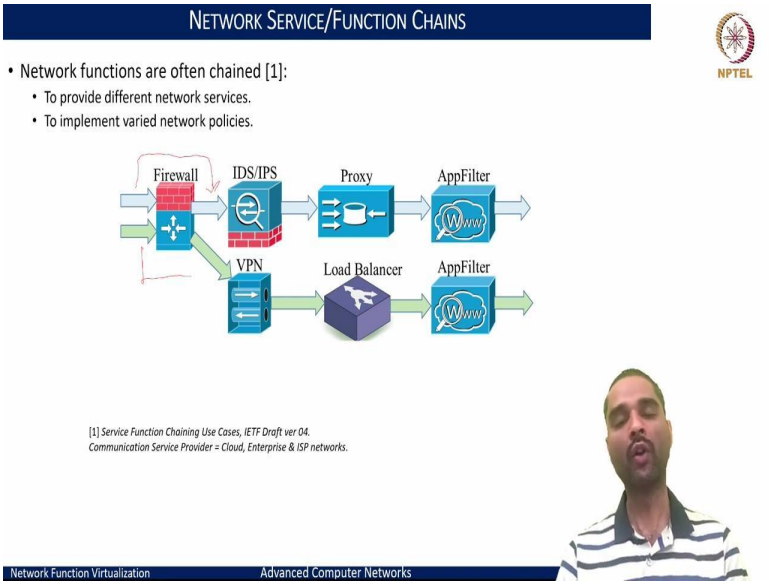

So, another important aspect, when we tie with these middleboxes, is often that it is not just one middlebox that is used in isolation. And what we eventually see is that these middleboxes are often used in conjunction with other kinds of middleboxes. And there is a dependence on how these middleboxes are used, set up, and deployed. And that is what we call as the service-based network functions that we are going to be using.

And the reason is that we use these middleboxes to facilitate a subset of the network services that we want, or when we think of campus or enterprise networks, we want to enforce certain policies for different kinds of traffic. And when we say these policies need to be applied, we are talking about what kind of middleboxes need to be employed for enforcing these kinds of policies.

So, for example, in this diagram, if we look at two different kinds of traffics, very simple traffic that is originating from within the network, or it could be traffic that is coming outside of the network from WAN, or it could be that we have within the campus, we have multiple of the departments that could be using, so, traffic that is originating from students and it could be the traffic that is coming from the faculty and staff.

So, there are different ways we can classify the traffic. And for each of these kinds of traffics, we may have different policies to enforce. And if we consider this example here, what we are showing is a blue kind of traffic and a green kind of traffic. And consider the green is the one

that is coming from external to the campus network, which means maybe it is going to use the services that are being put within the network through the use of the VPN devices, and VPNs and other kinds of most prominently used middleboxes.

So, when the traffic comes in green, what you are trying to show here is any traffic that is coming through the firewall into the campus network, could connect through the VPN, and then go through the load balancers and layer 7 optimizers, like application filters for www HTTP contents that we may be using, and ensure that the services are provided in the best way.

And likewise, for the blue kind of traffic, we may want to inspect for any malicious content through the IDS and IPS services and then connect them through the forwarding proxy and connect them with the application filters, all the way to the necessary services. So, this shows that when we defined these middleboxes, we are also defining how the traffic needs to be routed from one device to the other. And just the firewall itself, maybe trying to classify the blue versus the green kind of traffic and route them to different instances of middleboxes.

Also, as traffic gets forwarded, we see that different middleboxes take the decisions on their own, like IDS or IPS may even drop the packets, the firewalls could drop the packets, the load balancers would choose what would be the end host where you would want to send, which of the end servers that you would really want to send a packet, send the traffic towards.

So, all of these middleboxes are going to take their decisions on their own, and also say where to route or impact the way that traffic would be forwarded from one instance of a network appliance to the next. And this is where we see the complexity in terms of how these have to be set up.

(Refer Slide Time: 8:13)

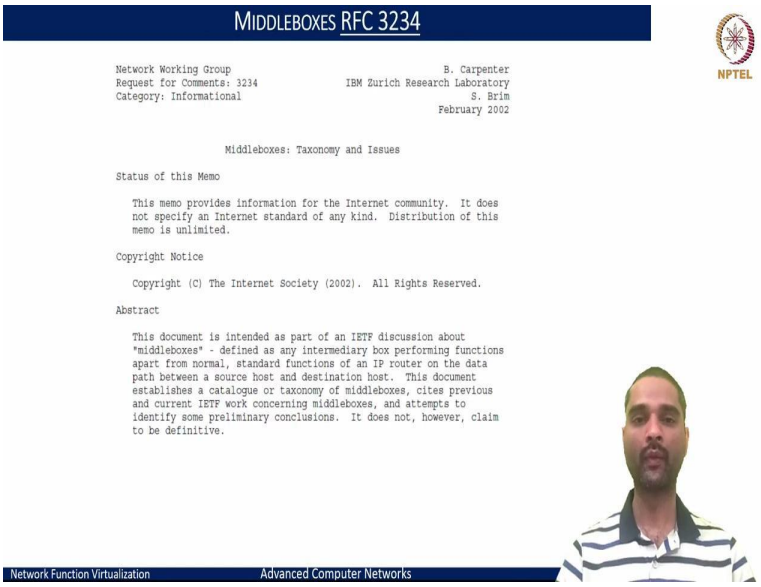

And this was analyzed very thoroughly. And the kinds of middleboxes that were in the early 2000s were studied and put up in RFC 3234, which basically classified the networking elements that existed in terms of the network appliances and provided the means to understand the kinds of middleboxes and what are the challenges that each would bring along. And the core of this document is about the cataloging of the number of types of middleboxes. And then when we speak of these middleboxes, and when we said they have different kinds of services they can provide. But then the question becomes how do we classify them and what is the means to build the classification.

And this document this study put forth that there are multiple facets to which we can classify and it could be a multi-dimensional taxonomy that we may be thinking of, but we need to consider specific characteristics in terms of how the middleboxes can be classified. And this could still be as simple as saying that based on what layer of protocol or layer of the network stack these devices could operate.

For instance, considering a protocol layer at which these middleboxes could operate could be one of the means to classify these devices. That means whether the middlebox acts at the IP layer or does it acts at a TCP layer, or does it act at layer 7, or the HTTP layer, could be one of the means to classify devices. For example, the firewalls we could have a layer 3 firewall to layer 7 firewalls. And likewise, when we have the load balancers, we could have the basic layer 3 load balancers or layer 4 load balancers.

So, these classifications for the same middlebox can even be done for a different protocol level. And another way to consider this classification could be based on like how these services are being provided, whether it is the middlebox function, or an explicit design for a protocol. Like when we consider SMTP replays, which are meant for specific SMTP functionality only, or is it an add-on that is being foreseen for trying to circumvent certain networking aspects, for example, the NAT device, because we are trying to add these NATs as translators to circumvent the issue of the public and private IPs.

So, we can call these as explicit versus implicit kinds of networking devices that we can have. It could also be on the basis of like, the way these devices could carry the state, they could be stateless, to begin with, the most simple where you just have the rules in the stateless firewalls where they would just allow or deny. Or you could have the firewalls themselves split the state with respect to the traffic they are serving. And if you take the load balancers, they are maintaining the state for each of the incoming flows to which server, the map for that flow is being set. So, these kinds of things become like a state.

And why state is very crucial is if the services break, that is, if the middlebox breaks, can it recover on its own and facilitate that the end-to-end connection remains intact? Or does it have a repercussion in breaking the end-to-end connection entirely? Like if we had a NAT, and the NAT tools are being taken off, or a NAT device restarts, and there is no state anymore for the ongoing connection.

Now we need these connections to be rebuilt. And this cannot happen until and unless the information is brought back into the NAT device through the flows that originate again. And that is where the state can make aspect to classify these middleboxes. And we call this as a soft state versus a hard state. That is, if the box loses the information and still the sessions can continue. And this is typically true in the case of value-added services that are being provided, while like the cases where the state is implicit and necessary for the service to happen, like in the case of an ad and others, we call that as a hard state wherein if that fails, the connection failed.

And the classifications can be in many more diverse ways. I think this document talks a whole lot about the way the soft states, and hard states way the issues with respect to how the optimizations are being provided, what kind of processing the devices are doing, whether they are manipulating the traffic at layer 3, layer 4, layer 7, what are the changes that are adding, what is the side effect with respect to any of the devices, all of this has been discussed in this RFC 3234.

And understanding of this is a must to understand why a middlebox is a real benefit or to see the other side of what are the challenges that they bring. So, to summarize this in context, the rise of middleboxes has given both to the positive edge as well as a negative impact. And if we think of end-to-end principle, and especially when I spoke about the hard state, they, in fact, break the end-to-end principle.

And nonetheless, it does not nullify the desirable properties of the middleboxes because we need these NAT devices, without which we could not have been able to successfully support a lot of users to connect to the network.

However, the future application protocols that we would want to define how to recognize how these middleboxes operate, what are the means or constraints that they bring along? And that is where it becomes essential to question how we want to build or what aspects we want to build in middleboxes. And having carried these thoughts said, the number of middleboxes and the middleboxes by themselves have become inevitable in the networks.

(Refer Slide Time: 14:51)

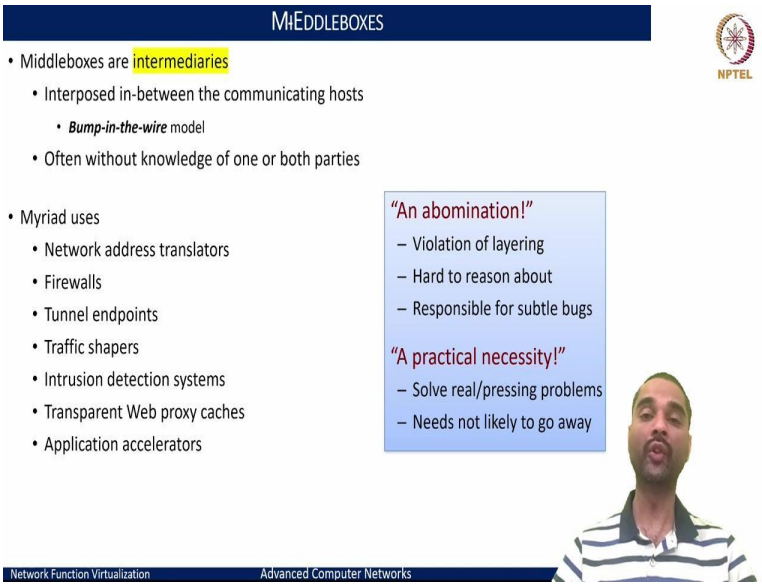

And when we look into these middleboxes we often refer to them as just intermediaries, that is, their existence is almost transparent. In essence, neither of the two communicating parties or the hosts would know of their presence. And this model is often called as the bump-in-the-wire model. When you are driving on a highway and you have a bump, you slow down, and then again pick up your speed, it is the same what happens with respect to the packets that are communicating between the two end hosts.

And when the packets are destined from one host to the other, the routers and switches would only process in terms of where to forward where to send a packet next, but these middleboxes act as the bump, wherein you now have to provide additional services or additional processing on these packets, for example, you want to build a NAT, then you are trying to modify and translate the packet headers. Likewise, with respect to load balancers, you are looking up the state and trying to match which of the servers that you would want to send; you are even changing the steering of the packet to altogether a different server. And this processing comes with a cost. And that is one, what really impacts the packets that are flowing between the two ends is the added delay or additional latency in reaching the two end hosts. And that is why the model of this middlebox that is being interspersed, always add to the latency. And that is why they are referred to as a bump-in-the-wire model.

And again, when we are having this kind of communication that is without the knowledge of either of the parties, you are not sure what kind of processing they are doing or it is tangibly

professed for providing the same data as what you would expect. On the other end, there is no guarantee on this. But we trust in the devices to be providing the key services to enable the kind of functionalities. But it also means that these are, in a way, a threat to this, ensuring that integrity, confidentiality and reliability of the data that we will be transacting. These middleboxes come in variety of uses like NATs, firewalls, tunnel endpoints, traffic shapers, and application accelerators, and so on.

And this is where we see the two faces. One, they are a practical necessity in terms of their deployment is necessary to solve networking problems. And their needs are not likely to go away by just putting the functionalities at the end hosts. Hence they become quite essential devices to be facilitated within a network, and we cannot get rid of them.

Second, by principle, they are breaking the foundations of how we envisioned the internet because they are violating the layering principle; they are violating the end-to-end principle, making the networks carry state besides just carrying the information. And once we have anything to do with the state, we know that failure in such things could add too much harder repercussions. And the devices, when they are more, they may have bugs in themselves. They may be having issues with configuring these kinds of devices as well.

So, in a way, they are also an abomination, in the sense that it's minimum use is better, but the reliance on them could create troubles for us. But because they are also a practical necessity, we have to live with them. That means we need to think of the ways where we can work with these middleboxes in a better and simpler way and make the workings simple.

(Refer Slide Time: 18:54)

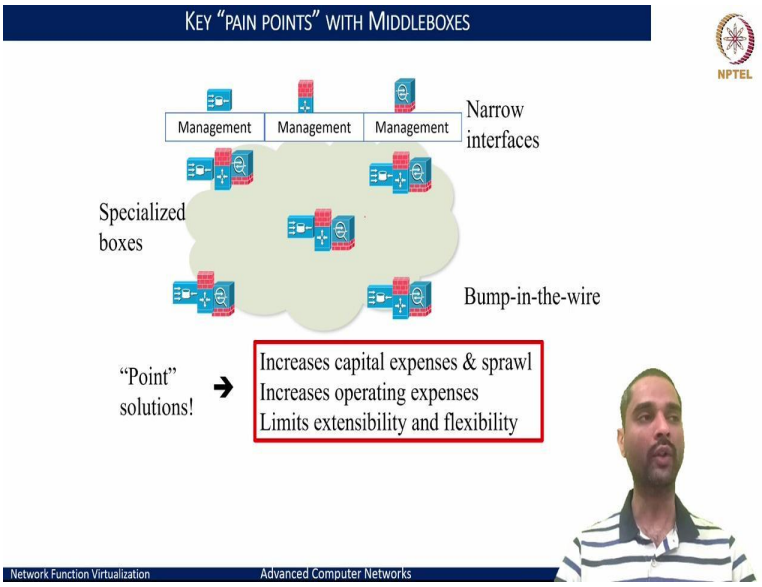

And this is where we need to now look at what are the key pain points that these middleboxes bring along and how we can really try to address them. So, again, if we revisit the same networks that we thought of earlier and now see that we have just one simple middlebox that is being deployed. The first question that these middleboxes bring as a pain point is how do we manage these devices, because when we deploy the routers and the switches, we manage them in a specific way likewise, when we deploy this kind of middleboxes, which are again, like purpose-built, built by a specific vendor specific OEM, and manuals for these devices to handle configure could be altogether completely different. So, we need the skills and expertise to deal with these kinds of specialized devices. And that is where management becomes a pain point.

Next, when we scale these different devices, the management problem is going to scale again, because the firewalls could be provided by a specific OEM, their specs would be completely different than the proxy devices which could be provided by some other vendor, the load balancer network appliance could be from so many other vendors. And managing with each of these requires specific skills of having to build. And that means we cannot leverage the same skill, I am well versed with the interfacing with the firewall, I can use the same knowledge interface with a load balancer that will not work.

So, we need different kinds of skills to deal with these kind of specialized devices. And that is where each of these devices present a narrow set of interfaces, where one needs to get the expertise with this management in terms of how do we configure these devices, how do we set

up these devices? And how do we interface with these devices and ensure that things operate in normal way in a network. So, we would now be asking for specialized skill set to manage these specialized devices.

And second, like I said, these are all like a bump-in-the-wire, the more we add these devices, the more we are going to end up paying the price in terms of the latency. And hence, they could also create operational hazard in terms of when the processing happen. And if they silently discard or drop the packets, we would have to then know where the things went wrong, or what are the consequences thereof. So, all of these aspects make middleboxes the real pain points in deploying.

And first of the issue is like these middleboxes, like we said they are purpose built assets, which are going to be deployed, that means they are going to increase the capital expense necessary to deploy these devices upfront. And you would see many of these middleboxes cost tens of thousands to hundreds of thousands of dollars.

And being able to making them afford and deploy them becomes a challenge. And that is where the increase in capital expense would be enormous. And second, once we deploy these, we need to operate with them continuously. And that means they are going to add also to the operational expenses to deal with these kind of devices. And once we deploy them another major problem is like we said, that routes for these devices, how they are connected, which middlebox would serve the traffic at which endpoint and where is the exact traffic going, to what middlebox it needs to go.

And if we have to scale up these services or add more, we have to reconfigure the entire networks. So, this greatly limits the flexibility of adding or removing the middleboxes once they are deployed. If I want to take out now the firewall, I have to rewire the entire thing to see you know the data that is flowing from one middlebox to the other are setup appropriately. Or if I want to add another new middlebox, which is going to be the same kind of problem, and that is where they greatly limit the extensibility and flexibility, making it a lot more problem in terms of even the physical operations.

(Refer Slide Time: 23:26)

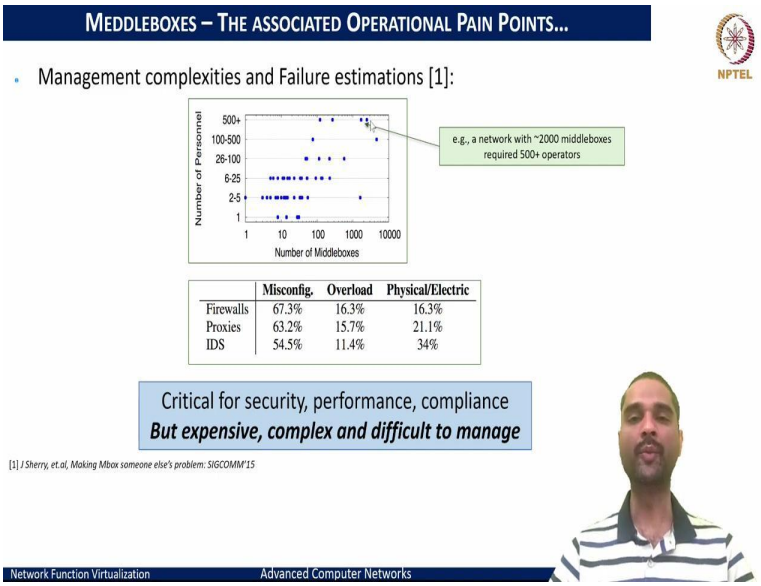

To address more on the management complexities, the survey that was conducted earlier also tried to study on the aspects of how the middleboxes affect the failures and what are the complexities involved in managing of the middleboxes. And the survey of various network operators showed that when we speak of the different middleboxes, firewalls prominently add to lot of miss configuration issues especially that means it is complex to manage or handle with the firewalls and getting the expertise to configure them properly is an important aspect. And that is a challenge.

And likewise, the proxy devices, how are you setting up and configuring these proxies for seen to be a major aspect which was resulting in the failures and like 67 percent for firewalls to almost 63 percent for proxies to 54 percent for the IDS where the network operators voted saying that these were more prone or susceptible for misconfigurations that means the lack of expertise in handling these devices.

And likewise, like how these devices when they become overload or lead to the potential problems in network traffic. And like firewalls again seem to be the prominent cases when they are straightful firewalls resulting in lot of processing overheads, lot of delays, resulting in the failures. So, it becomes interesting to see that these middleboxes, although like it enumerates a few of them, there can be lot of repercussions in terms of the downtime that we may see because of misconfigurations. Or, in fact, adding to lot of adverse effects in the network, creating lot of

processing and overheads, consuming lot of power, and so on, could be the issues that these middleboxes bring around.

So, the survey tried to further look at, like, if I have a network with high number of middleboxes, what is the resource personnel that we may need to ensure that the middleboxes can be configured and set up. And what this shows was that a network with around 2000 of the middleboxes would require around 500 or more operators to manage them.

So, the x axis is showing the number of middleboxes that you would want you have in your network and for the corresponding number of middleboxes, what are the specific resource personnel with specific expertise that you are hiring to manage those network appliances, and you can see that for around 1000 or so, 2000 of the network middleboxes, you are going to spend on 500 or more operators to manage them. That means it is going to be a huge operational expenses that these middleboxes bring along.

And to take out, these middleboxes are critical for security, performance, compliance or policies, aspects that you would want to enforce, but are really expensive, complex and difficult to manage. And this is where we need to see how we can address these specific challenges with the middleboxes.

(Refer Slide Time: 27:02)

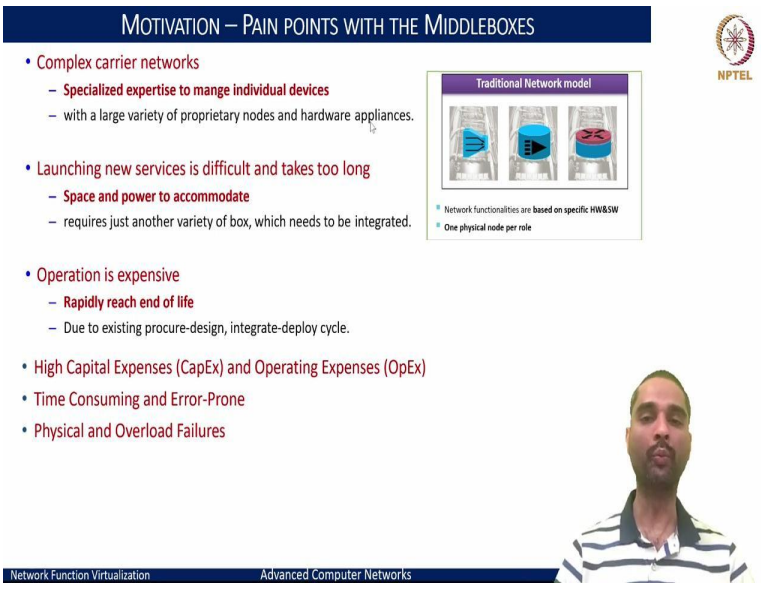

And these were primarily the key motivation aspects in trying to see that we want to operate with the middleboxes, but alleviate these pain points. And to sum up the pain points these middleboxes, in essence they make the carrier networks more complex, and you would need specialized expertise to manage these individual devices because each is a dedicated hardware and software. And the skills for operating with one hardware and software would be drastically different from operating with the other devices.

And moreover, when we deploy them in large numbers they make the networks essentially very complex, making them very difficult to manage, or very difficult to even alter the topology of the network services that we would want to add. So, if we want to bring in new services on table, we have to run and work through the current model of how the new middleboxes are being set up, where to set up a new middle box, what kind of integration plan needs to done, all of this has to be well planned.

And again, these kinds of middleboxes, they also account for a space and power in the network infrastructure. So, the space and capacity planning needs to be done to ensure that we can accommodate these middleboxes. And besides operations being expensive, and we have seen, like when we see the firewalls, they started with the simple stateless to the enhancements and improvements and these middleboxes have seen drastic improvements and enhancements over the kind of devices.

And that means when you spend hundreds to thousands of dollars in deploying these devices, you may want to keep on upgrading and change them as you go that means they may rapidly reach the end of the life and moving out or moving in with the new devices could also add to lot of again the capital expenses coming up.

And even the most other significant aspect, worrying aspect is we have a solution that we need today but we cannot have it until it is being shipped and provided from the specific vendors. That means we have a procure stage and then we design the things and we see, whether they fit or we need additional things. And once we say there are additional requirements, these need to be designed deployed by third parties or the OEMs and then we need to again integrate and deploy them back in our networks.

So, the cycle of procure design and integration deploy could span years, which will also make it more difficult to address the immediate or pressing functions that we would want to serve with the middleboxes. So, overall, these middleboxes end up providing very high capital and operational expenses. And deployment of them and management could be time consuming. And as we saw misconfigurations could be the casual occurrences, and hence leading to lots of errors. And they could also be subject to the physical and overload failures that would otherwise also impact the way the network services would be facilitated.

And hence, having these traditional networks model where we think of these network appliances has dedicated hardware and software, which are being run as one physical entity per node, like one hardware, one functionality, could all mean that there is a lot to pay for. And hence, we need to simplify these aspects.

(Refer Slide Time: 30:55)

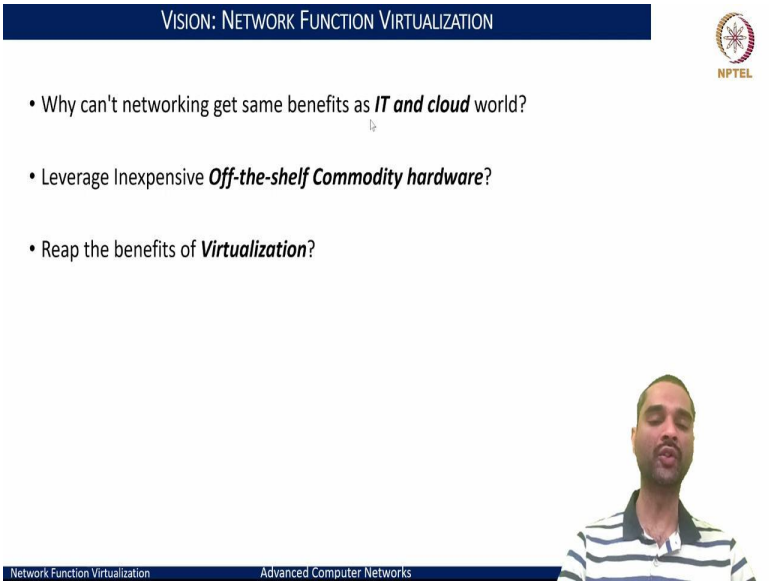

And now if we think of these problems, and if we jump back and see how the things were, when it was with respect to the clusters, or the grid networks, or the cloud infrastructures themselves, in essence, we had to deal with lot more physical devices, physical service, bring up the network storage, compute all of those devices together, and manage them. And how it was done in the IT or the cloud services was essentially with the use of off-the-shelf commodity hardware.

And more prominently, the way we would want to scale or setup the things was through the means of virtualization. So, when we virtualized these off-the-shelf commodity hardwares, we built off-the-shelf commodity software or virtualized instances, wherein I could consolidate many of the servers that would otherwise require a dedicated hardware now to be run on a single hardware. And that was the beauty of the virtual machines that we got. So, now if we retrospect and see whether this would fit for these network appliances that we just discussed, and in what way can this really aid to break the challenges that we just saw.

(Refer Slide Time: 32:15)

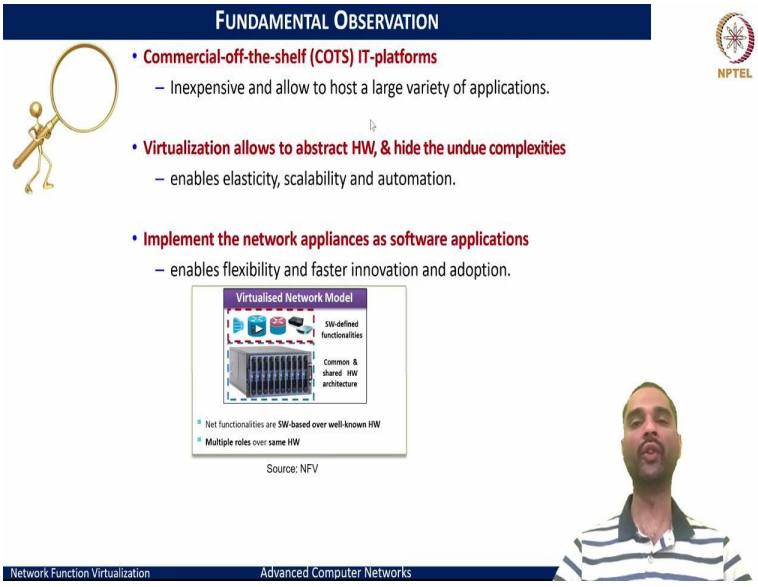

And we can see that with the use of the commercial-off-the-shelf, or the COTS or the commodity IT-platforms, we could basically use a single hardware to facilitate multiple of the network appliances, the functionalities of different network appliances. And this would also mean that now we would not have to spend lot more on for getting these middleboxes, but it would make these middleboxes very, very cheap or inexpensive.

And it would also allow us to host a large number of these devices on a single commodity hardware. So, I can mix and match different kinds of network appliances that can be run on the commodity hardware.

Second, if we rather than just limiting ourselves to the commodity hardware, and add the virtualization argument that with the virtualization, this allows us to completely abstract the

hardware and hide any of the unnecessary complexities that we have to deal with in managing the hardware and the use of virtual machines and virtual networks put together.

This would now enable us to have much desired elasticity in terms of we can bring on multiple instances of the same virtual machine as and how we need and we can scale to the network demands. And virtual machines allow us to automate at one point and start spawning these virtual machines rather than going and setting up the hardware, booting them up, wiring them, all of these hazards could be taken off.

And hence the use of virtualization, we can see that it can really enable us a lot more flexible means to manage and set up these devices. And third, when we make these network appliances as software applications that are going to be run on these virtual machines. Now, we can again bring the abstractions in terms of how the interfaces could be designed and aspects of how to manage these devices, which can again be unified in one way. That way we can eliminate the overheads of having specific resource personnel to manage these devices.

Nonetheless, we may not be able to get rid of all the complexities thereof because we still need the expertise who could understand what this software appliance would do. But we could provide better abstractions like an intent based configurations that we spoke earlier. We do not have to understand the nitty gritties of the software's but provide the intents which can then be translated to the device specific characteristics that they need or software specific parameters that can be broken down and setup.

Thus, these observations that we make here, make us rethink of using the hardwares and try to see whether we can softwarize these dedicated appliances and provide software defined functionalities first, which could be run on the commodity or the shared hardware architecture.

What this would mean is, now a single hardware could be facilitating multiple roles, or multiple network appliances all run on one hardware. So, management wise, it becomes lot more easier infrastructure capacity planning, again, simplified. And when we have the software and build the right abstractions, also the operational costs and management costs could be contained to a greater extent. And this is exactly what the network function virtualization deals with.

(Refer Slide Time: 35:51)

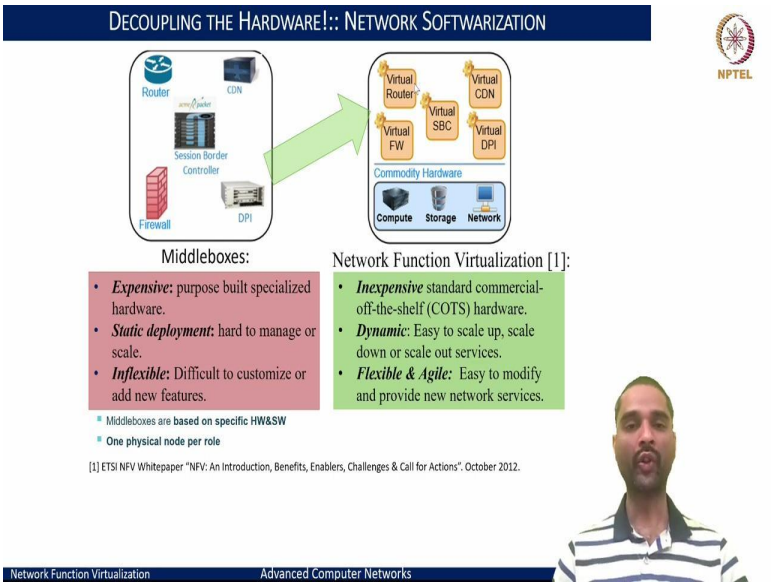

So, instead of these middleboxes, which are specific, or a proprietary hardware software combinations, and each middleboxe being one physical on a node, which was seen to be very expensive, because they are purpose built hardwares, which also have these issues with deployment, because you have to statically, set them up, wire them up, plumb the connections, making them hard to manage or scale.

And having them in hardware makes it difficult to customize or add new features making them highly inflexible, two what we speak of network function virtualization is to run these hardware or software entities where you are going to run the router CDN any of the network appliances as a virtual router, virtual CDN, virtual firewall and likewise, and these are all going to be run on the commodity hardware, which is the shared hardware infrastructure where you have the compute servers, storage servers, network connections being made for these servers, all of these as a commodity devices now, which are going to run these network appliances all in one, so you have made them go inexpensive, because you are running on a standard off-the-shelf hardware, or the server machines, x86 Machines that we speak of. And now because we can bring these appliances at will as like trying to spawn a new process or spawn a new virtual machine that is going to run on this hardware, it makes it very dynamic infrastructure where we can easily scale up or scale down or scale out the services as we need.

And because these are softwarized instances, it becomes a lot more easier to manage and modify these devices, or even to bring innovation readily onto table because we can write the codes, we

have the power to write the code, power to create new kinds of these virtual middleboxes without having to depend on the OEMs or any proprietary vendors to provide us these functionalities. And that is where network function virtualization tried to shape what we see as a means for network softwarization, helping immensely, the network operators in building and managing these networks.

(Refer Slide Time: 38:29)

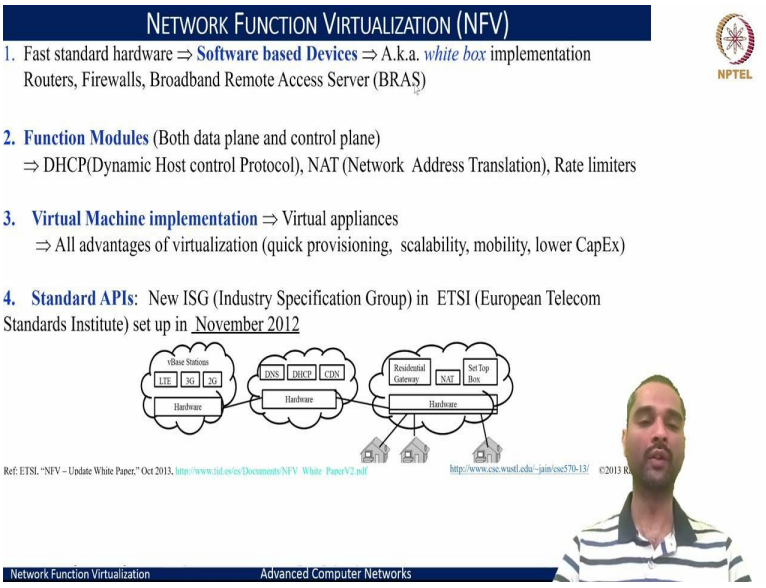

So, to sum up, what we really see is that this NFV brings in a means to run these network functions or software based devices on a fast standard hardwares, the x86 server machines to think of making the entire implementation of the network appliances as a white box model, wherein we have the control over both the software and hardware and enable any of these functionalities like routers, firewalls, broadband accesses, and so on.

And second, because these are now the function modules that we are thinking of, we have full control over what the data plane has to be and the control plane has to be for these models. Like if we think of a DHCP, now we can build DHCP software, run it on other Linux machines, NAT which you could run on our x86 servers, the rate limiters that we want to deploy or employ. All of these can be thought of as very simple functional modules, and we get to have control and data plane operations done quite easily.

And third, because these are virtual machine implementations, or rather, we now call these network appliances as virtual appliances, we leverage all the advantages that come along with the virtualization or the virtual machines that is very easy to provision the resources. And it can be automated.

In fact, sitting at one point we can spawn multiple of these virtual machines. We can scale on demand rather than waiting for a new hardware to be procured, set it up, we can have a very timely scalability, we can even move from one server to the other when we want to manage or if some devices going down, we can migrate the virtual machine from one hardware to the other. This greatly adds to the mobility and enables us to have uptime a lot more.

And overall, because these are virtual machines, we do not have to spend anything on buying the hardware, it significantly lowers the capital expenses. And having softwarized these devices, it gives us a new means to say how we can standardize the API's like what way we can build the abstraction, so that we can deal with the different middleboxes in a simple coherent manner.

And this is where like the formation of the new industry specifications group in ETSI happened around November 2012, which has been laying the foundations of how middlebox abstractions can be brought, what are the frameworks for deploying these middlebox, what should be the architecture for building the network from NFV or virtualized network functions that we will start to understand in the next part.