Advanced Computer Networks
Dr Sameer Kulkarni
Department of Computer Science Engineering
Indian Institute of Technology, Gandhinagar
Lecture 37
Introduction to Network Function Virtualization - I

Over the last two weeks, we looked at how network virtualization, overlay networks, and software defined networks have enabled to sanitize the networks and empower innovations by facilitating programmability and ease for deployments of prototypes and evaluation.

(Refer Slide Time: 00:43)

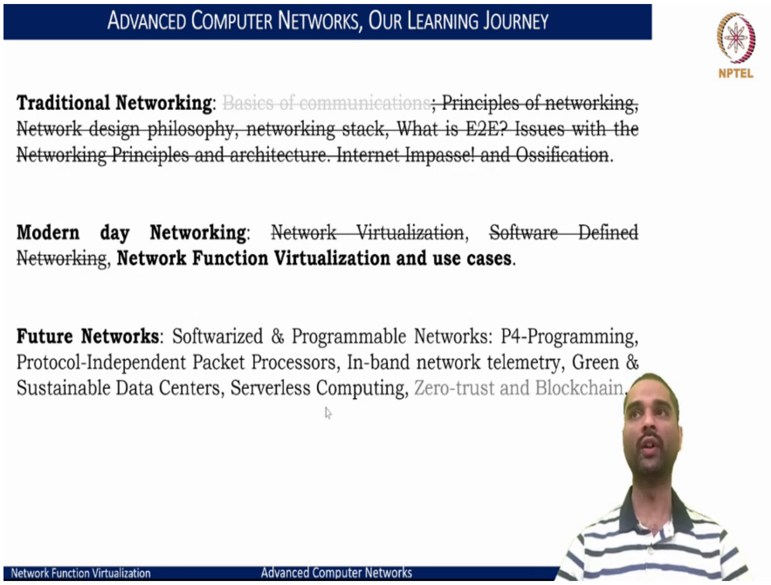

ADVANCED COMPUTER NETWORKS, OUR LEARNING JOURNEY

Traditional Networking: ~~Basics of communications; Principles of networking, Network design philosophy, networking stack, What is E2E? Issues with the Networking Principles and architecture. Internet Impasse! and Ossification~~.

Modern day Networking: ~~Network Virtualization~~, ~~Software Defined Networking~~, **Network Function Virtualization and use cases**.

Future Networks: Softwarized & Programmable Networks: P4-Programming, Protocol-Independent Packet Processors, In-band network telemetry, Green & Sustainable Data Centers, Serverless Computing, Zero-trust and Blockchain.

Network Function Virtualization Advanced Computer Networks

This week, we will look into another major change that is, in fact, radicalizing the way we have the networking infrastructures and also helping to facilitate better management and orchestration of network elements. So our primary focus for this week would be Network Function Virtualization, or more popularly known as NFV, and its associated use cases.

(Refer Slide Time: 01:08)

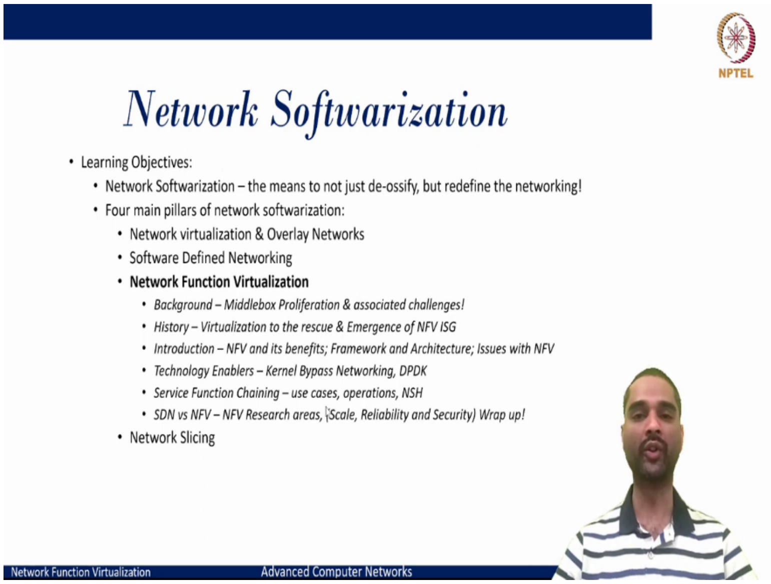

Network Softwarization

- Learning Objectives:
  - Network Softwarization – the means to not just de-ossify, but redefine the networking!
  - Four main pillars of network softwarization:
    - Network virtualization & Overlay Networks
    - Software Defined Networking
    - **Network Function Virtualization**
      - *Background – Middlebox Proliferation & associated challenges!*
      - *History – Virtualization to the rescue & Emergence of NFV ISG*
      - *Introduction – NFV and its benefits; Framework and Architecture; Issues with NFV*
      - *Technology Enablers – Kernel Bypass Networking, DPDK*
      - *Service Function Chaining – use cases, operations, NSH*
      - *SDN vs NFV – NFV Research areas, (Scale, Reliability and Security) Wrap up!*
    - Network Slicing

Network Function Virtualization

Advanced Computer Networks

So in this context, we will try to learn the basics or the background about how the NFV started, or why NFV was needed. And in this aspect, we will learn about the middleboxes, and why the middlebox proliferation happened. And what kind of newer challenges besides the ones that we discussed towards ossification were brought in by these middleboxes. Next, we will try to see, what were the mechanisms that were tried to address the complexities, and the challenges that came along with the middleboxes.

And again, we will see that virtualization came to the rescue, and that led to the emergence of NFV. And in this journey, we will learn about what is NFV, and what are the key benefits that it brings to the table. And when we speak of virtualization, again, we are talking about the right abstractions. So we would be talking about what is the framework and architecture for NFV. And then, we will try to cover the key issues that also the NFV brings along.

So whenever we have the benefits, we will also have to look at the other side of the coin. And that is where the drawbacks or the challenges that vest within NFV need to be understood. And also, when we have the challenges, we also come up with the solutions to address them, and few of the very important and addressing aspects with respect to NFV have been some of the technology enablers that we learn.

And we learn in this context, what we call as Kernel Bypass Networking, and try to understand a brief about DPDK, what we mean by DPDK, and how this helps enable overcome some of the associated challenges in NFV. And then, when we look into the NFV, and the primary use cases, we will learn about the other associated aspects that is Service Function Chaining. And we will learn about, what are the key use cases that would use it,what are the kinds of operations that we do and learn about one specific protocol called Network Service Headers. And finally, we will wrap up with trying to understand what is NFV, distinguish that with the Software Defined Networks, and see how they are actually tangential, and in fact, complement one another. We will also likewise try to learn about the key NFV research areas, in terms of ensuring the scale, reliability, and security aspects, and wrap up our study on Network Function Virtualization.

(Refer Slide Time: 03:51)

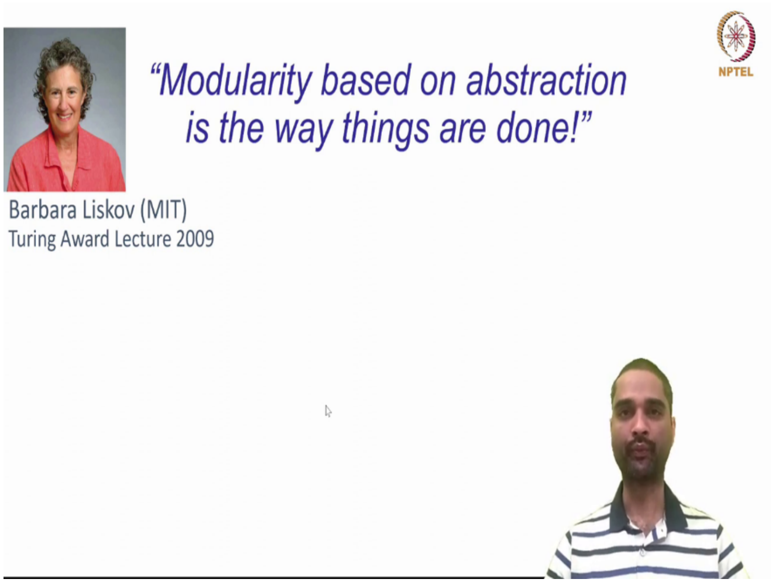

“Modularity based on abstraction is the way things are done!”

Barbara Liskov (MIT)
Turing Award Lecture 2009

And again, I point to this great saying, by Barbara Liskov at MIT. And she is one of the three Turing Award winner females in our generation. And what she pointed out, was that modularity based on abstraction is the way things are done. And in fact, in a way what abstractions really provide are the foundations upon which the computers are built when we look at the computer architecture stack.

And likewise, when we associate that with the development of networks, we are able to fit exactly the propositions that the abstractions brought for the computers, the same that they are bringing for the networks. And in essence, entirely for the network infrastructure, and also the data centers etc, that we see the way they are built. They rely heavily on this abstractions. And if we think over the past 10 to 15 years, the networking has been all about presenting and bringing in the right set of abstractions for better network control, better management of the data.

And if you see even the operations side, they needed the right set of abstractions. And in essence, the network, control, and data plane, they have all been addressed with the right sets of abstractions. And in this context, NFV is also about virtualizing, and bringing the right set of abstractions, in terms of looking at the network appliances.

(Refer Slide Time: 05:36)

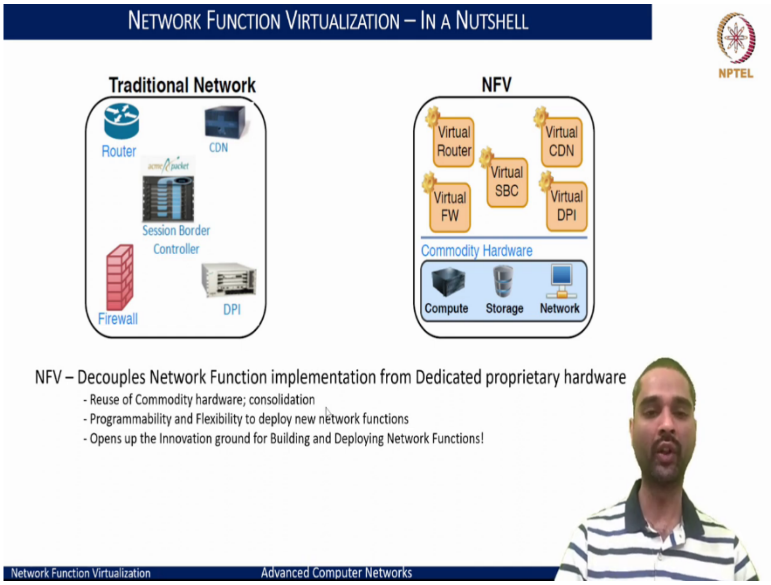

NETWORK FUNCTION VIRTUALIZATION – IN A NUTSHELL

NFV – Decouples Network Function implementation from Dedicated proprietary hardware

- Reuse of Commodity hardware; consolidation
- Programmability and Flexibility to deploy new network functions
- Opens up the Innovation ground for Building and Deploying Network Functions!

Network Function Virtualization Advanced Computer Networks

So let us put this in context, and try to understand what NFV is in a nutshell, before we dig further deeply into any of this aspects. I introduced you earlier about the networking elements to a very brief extent though, but we have known that we have been operating with switches and routers for long. And then we also have a notion, what kind of a firewall or DPIs are being used for Content Delivery Networks (CDNs), and Session Border Controllers, all of these kinds of traditional networking appliances.

And what NFV tries to do is, if we see these devices, we are seeing them they are all purpose built dedicated hardwares. And that means if I have a network with the router, with a CDN, with a firewall, I am talking about at least three different devices that I have to manage and setup. And this in a way makes it a lot more harder for not just in terms of trying to buy them and set them up on a network, it is a one-time capital expenses that we will have to spend.

The other is also that each of these devices might come from different vendors or different OEMs. That means the way I have to interact, I have to set them up would also vary based on who the vendor is, what is the manual in terms of what is the mode of configuring a router, what is the mode of configuring a firewall, all of these have their own set of APIs to deal with. And that makes it even more hard to even manage these devices, over the years, we need experts to handle them. And that is where, and they would also incur heavily on the operational side of the expenses. And this is where many of the network operators were facing a lot of concerns. And as a result of how to overcome the solution came the NFV. And NFV, what it tries to do is, I do not want to be looking at each of these distinct devices as individual devices and have to acquire expertise of managing all these devices. Instead, I can think of these as software programs, and ensure that these software programs could be run on a common platform or a common hardware. That means I am trying to decouple these purpose built devices which facilitate specific network function on specific hardware to make them run on a generalized hardware, but as specific functions.

So we are in essence, decoupling the hardware and software. And we are trying to get rid of these dedicated proprietary hardwares, and run all of these network appliances as software functions. So you consider a router, which is both a hardware and a software component that is purpose built. We are now trying to break that and make just the software virtual router, which would be run on commodity hardware. And by commodity hardware we mean off-the-shelf available devices like the server machines, x86 server machines with compute, storage and network.

And likewise, if we take a firewall, now we would run a firewall as a software that is running again, on top of the same commodity hardware. And this way, for any of the network appliances that otherwise we would have from a proprietary vendor, buying a proprietary hardware, would now be transformed into running on a single hardware. And we are going to reuse this commodity hardware by ensuring that these virtual routers, virtual firewalls, virtual DPI, all of these could be run as software instances on top of this commodity hardware.

Think of these as processes that are running on your PC, running on your laptops. That is the flexibility that NFV would want to provide and for very simple reasons now that if we have this kind of a model, then we can have multiple of these functionalities being run on a single machine. So we do not need to buy specific hardwares from different places, or have to buy n distinct hardwares for running n distinct functions.

We can consolidate them and run just one or two of these hardwares, which would facilitate to run multiple of these network functions. And that is also going to provide us another benefit in terms of when I think of virtual router, firewalls all being run on same hardware. This is equivalent to how we saw the SDN help us with respect to the centralized control plane. So now we have one point or one device where we can control likewise, because we are able to run these on a same commodity hardware.

The other benefits that we leveraged in SDN were also the programmability. And now because these are softwarized instances, we would again get the programmability aspects, which brings in a lot of flexibility for us in terms of how we want to deploy these network functions. In fact, when we decouple the hardware from the software, we can run it anywhere on any machine, any node. And this is the other flexibility that it is going to provide when we think of going towards NFV. And the benefits are likewise again, similar with respect to what we learned in SDN, but for different planes.

And again, what SDN enabled to break the ossification in terms of the developments that we wanted to do at the network, or the link layer protocols that could be enabled or the network functions that we spoke about as routing and others, which we could readily deploy to newer algorithms. And likewise, what NFV tries to do is when we make these different network appliances or softwares, it is now no more tied to a particular vendor to deploy and provide us the solutions. If we want to innovate in the networking stack, starting thereof from the transport all the way to the Application layer. Think of these as applications, then we could innovate in trying to bring up newer applications and deploy them as just softwares, not necessarily rely on any dedicated hardwares, or VMs, to provide us the solutions.

So this way, it opens up the ground for innovation, and enables for accelerated deployments of the network functions that we can build as softwares, and deploy on the testbeds as we need. So these are in essence, the key aspects with respect to NFV. And our goal now is to understand these aspects in detail about how NFV is shipping all of these that we just discussed.

And to do that, we need to understand the left hand side of the part in a bit more detail that is especially trying to understand what these network middleboxes that we are trying to say,

(Refer Slide Time: 13:19)

NPTEL

Let us look at the typical Networks, Middleboxes and address some of the fundamental challenges with the adoption of these networking appliances!

Network Function Virtualization

Advanced Computer Networks

what are the aspects that they bring, what are the fundamental challenges associated with them, all need to be understood to some greater degree. So first, let us try to do that, and then jump on to the lane of NFV.

(Refer Slide Time: 13:34)

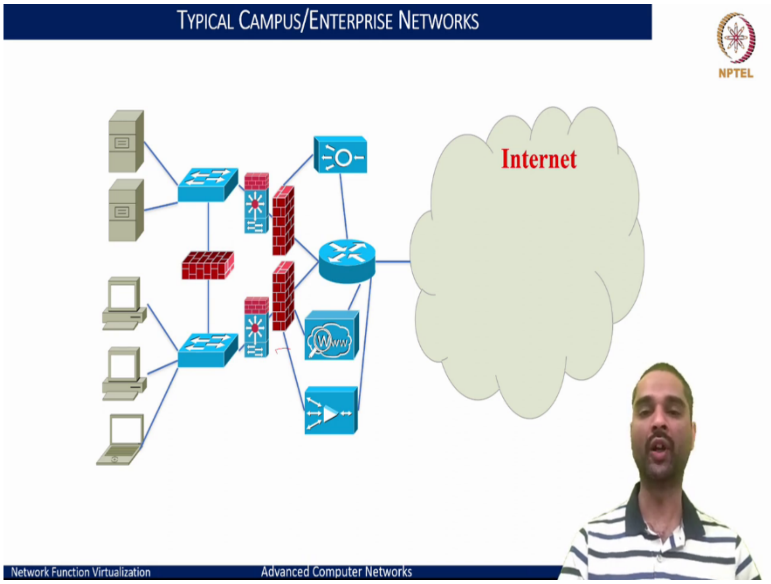
TYPICAL CAMPUS/ENTERPRISE NETWORKS

Network Function Virtualization

Advanced Computer Networks

So let us begin our journey trying to understand, what a typical campus network would look like or campus or enterprise networks that we really either work on or live in. And even this is similar for our home networks, only that there will be a vendor provided CPE element that is going to be put and the way like if we were the ISP manages that CPE for us.

But if we look at our campus, or the enterprise networks, if you are working in any of the companies, what it typically looks like is, on one end we have a bunch of local area networks where we users often try to connect to the internet.

And on the other side, we would have a DMZ zone where you typically host some of the server, storage, and databases that would be running and facilitating the connection and services to the internet. And the way these are going to be connected are often using the typical Layer 2 devices like switches, and Layer 3 devices like routers and edge router eventually for the means to connect us with the Internet. That this is how we visioned how the networks to be. But what we would see in reality is that the campus or enterprise networks, or including the data center, or any other networks are interspersed with a lot of the middleboxes. Like, but on the ends, either the people within the network in a LAN, you are not aware where or how they, like firewall in this case has been deployed.

And likewise, we would see from outside of the Internet neither the other end would know where the NAT device would reside or what the firewalls that are residing within this network. So these kind of firewalls that we just added, are added transparently into the campus or enterprise networks without the notion of awareness for end users on either ends. And it could be a lot more devices than just the firewalls.

Like we could have the packet gateways that are going to be adding some filtering capabilities that can also be added. We can have several other optimizers, like one optimizer is to be able to review the caches, all of these being inserted within our campus network. And, in fact, in reality, what we really have is not just the routers and switches, but so many of the different middleboxes that are being set up on the network.

And any traffic that now goes through this for us to facilitate connection could be within a LAN where we are trying to access the specific services inside our network. Or it could be our connections to the internet, based on where we are trying to connect, we may see that our traffic goes through different middleboxes. And this is very common scenario on most of the networks that we have today.

And the number of middleboxes could keep on growing in terms of like, if you are trying to build so many of the services, we want to add the load balancers, we want to add the proxy devices that is reverse proxy, forward proxy based on where and how the services are going to be exported and provided. So all of these devices add up making our networks start to look lot more complex.

And in terms of network operators, they have to manage not just the switches and routers, but so many of these variety of different network appliances. And this is where the challenges start to begin.

(Refer Slide Time: 17:33)

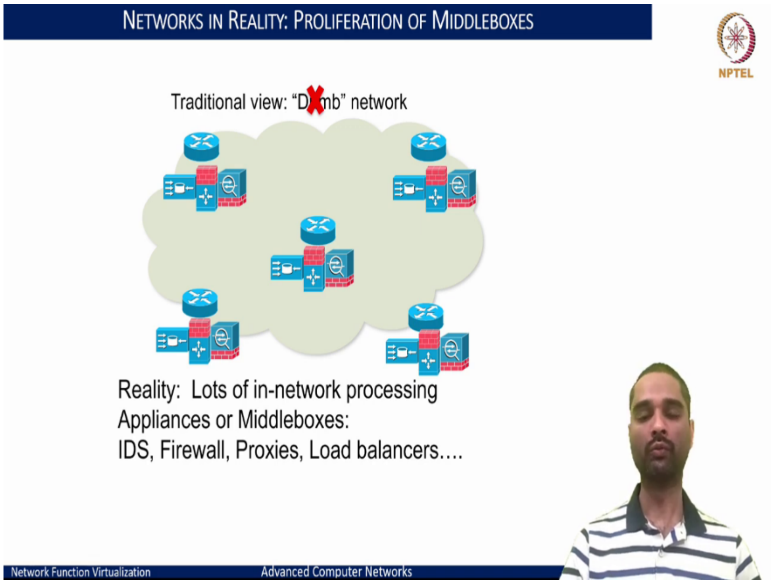

NETWORKS IN REALITY: PROLIFERATION OF MIDDLEBOXES

Traditional view: “D~~u~~mb” network

Reality: Lots of in-network processing
Appliances or Middleboxes:
IDS, Firewall, Proxies, Load balancers….

And if we start to now look at how this proliferation has happened at across different networks, including the internet service providers, telecommunication networks or the data centers, what we can see in reality is that the networks which we earlier considered as dumb in terms of being just the pipes that would facilitate routing the end users content is no more true. And what we have is not a dumb network, but an intelligent network that is operating on the data that is being passed through such a network.

And there is a lot of network processing, or what we call in-network processing that is happening because of these variety of middleboxes that we have to deal with. And these appliances come in variety of forms and provide variety of functionalities, like IDS, firewall, proxies, load balancers, and so many. And our network that we can envision is being flood with lots of these kinds of middleboxes.

So you can see if we have just these five routers, and each of these connecting us to one of the WAN or a LAN networks and connecting us over an internet, each have their own number of middleboxes that are going to be deployed. And you can see these number of middleboxes keeps growing. And as the in-network service or in-network functions increase, the deployment of the middleboxes would also increase.

(Refer Slide Time: 19:22)

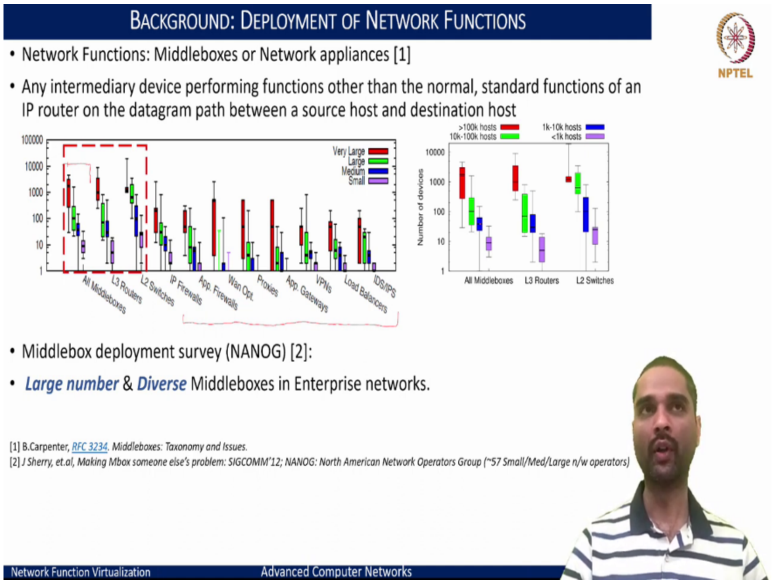

Background: Deployment of Network Functions

- Network Functions: Middleboxes or Network appliances [1]
- Any intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host

- Middlebox deployment survey (NANOG) [2]:
- ***Large number*** & ***Diverse*** Middleboxes in Enterprise networks.

[1] B.Carpenter, RFC 3234. *Middleboxes: Taxonomy and Issues.*
[2] *J Sherry, et.al, Making Mbox someone else's problem: SIGCOMM'12; NANOG: North American Network Operators Group (~57 Small/Med/Large n/w operators)*

Network Function Virtualization   Advanced Computer Networks

So we referred in brief about the deployment of network functions earlier. And we said the network functions are basically the network appliances or middleboxes, as the intermediary devices that are placed on the datagram path between a source and a destination host that perform specific functions which we had called them as not the normal functions that a standard IP router would do that is in essence, forwarding the packets. And they are doing more than just forwarding function.

And these network functions were getting more and more heavily deployed on many of the network functions in the early 90s, where we said the internet community started to sense a lot of these kinds of emergence of these devices. And especially one of the main reasons we said was it started with the exhaustion of the IPv4 and transition to the model of IPv6, or making sure that IPv4 would last with a transition to private and public IP address model with the NATs.

All of these led to the proliferation of middleboxes. And what I am showing here in this box plot, which we referred in very brief context, but let us try to dig in a lot more detail and understand what we are seeing here. If we focus on the x-axis, we are saying that there are a variety of middleboxes, the first part here that is being shown. And the middleboxes may be of variegated type. So we are seeing the list of middleboxes as examples here, starting as IP firewalls, application firewalls, WAN optimizers, proxies, application gateways, VPNs, load balancers, IDS IPS, etc. And they have been deployed of variety of kinds of middleboxes. That is the variety that we are trying to show here. And it is a small fraction of numbers that are significant in a given network operators domain.

On the right, what we are trying to show is that the number of middleboxes when you sum them up in a network and try to compare them to the numbers of the L2 switches, and L3 routers, we are seeing the proportions to be almost similar. And like I said, this was a survey that was carried out on the North American Network Operators Group, which consisted of 57, of small, medium, large and very large network operators. And the distinctions of small, were based on the number of hosts that they were serving, like, when we said small, it has hosted less than 1,000 hosts. And when he said medium, it was hosting around 1,000 to 10,000 hosts. And large 10,000 to 100,000. And very large means more than 100,000 hosts facilitated in a given network.

And with these distributions of the way the networks were defined, and if we see across all of these networks, we see that the number of middleboxes, often in very small ones, you would see that it is much more than the L3 routers. But as the internet and the connections grow as the network scales up, we are seeing still the number of middleboxes that are deployed, that is the red box plot or a green or a blue one. And compare them to the L3 and L2 devices, their numbers seem to be on the same lines.

Note that this y-axis is on a log scale. So we are talking about the number of devices in the orders of 100s to 1000s of middleboxes that are being deployed in the medium to the large, very large networks. And here these middleboxes are also coming in a variety of types. And when we look at what exactly or why exactly do we need, they serve different purposes and different functionalities.

And if we zoom in now into just the left hand side and try to see how the middleboxes have been, we can clearly see that the number of middleboxes in the order of 1,000s to 10,000s in a very large networks is seemingly very common. And when we look at the very small networks, even in small networks where we have the routers and switches in the orders of 10s or 100s, the number of middleboxes is also in the same scale and around 10 to 15 of these middleboxes for every router.

So more than the routers, you may end up seeing the middleboxes that are being set. Hence, we need to understand, why this proliferation has happened and how are these middleboxes being used? And I referred earlier as well they are built for various purposes.