Advanced Computer Networks
Professor Doctor Sameer Kulkarni
Department of Computer Science Engineering
Indian Institute of Technology Gandhinagar
Lecture 39
Network Function Virtualization - Concepts, Framework and Architecture - I

(Refer Slide Time: 00:19)

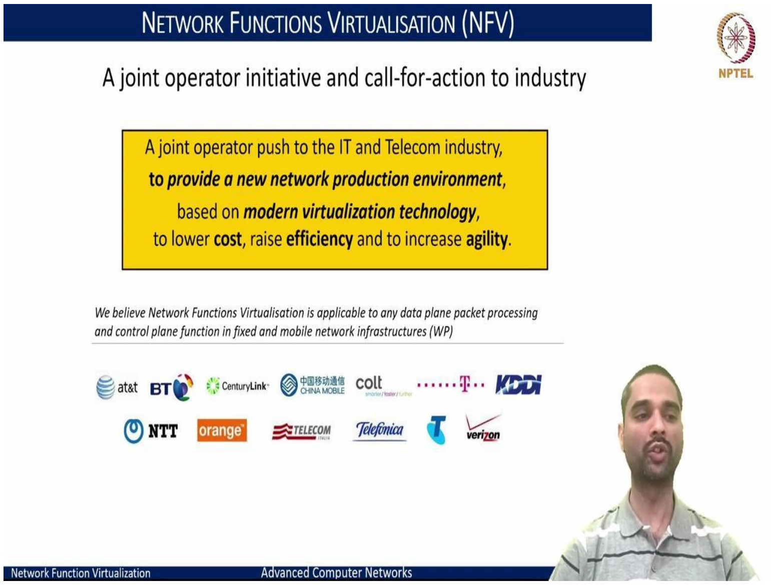

NETWORK FUNCTIONS VIRTUALISATION (NFV)

A joint operator initiative and call-for-action to industry

A joint operator push to the IT and Telecom industry, ***to provide a new network production environment***, based on ***modern virtualization technology***, to lower **cost**, raise **efficiency** and to increase **agility**.

We believe Network Functions Virtualisation is applicable to any data plane packet processing and control plane function in fixed and mobile network infrastructures (WP)

It all started with the joint network and telecommunication operator initiative and call-for-action for the industry in order to provide a new network production environment that was based on modern virtualization technology, which will enable to lower expenses, that is lower reduced the cost, and increased efficiency and agility for managing and operationalizing the networks. And we need to look at what reasons compelled to do this.

(Refer Slide Time: 00:57)

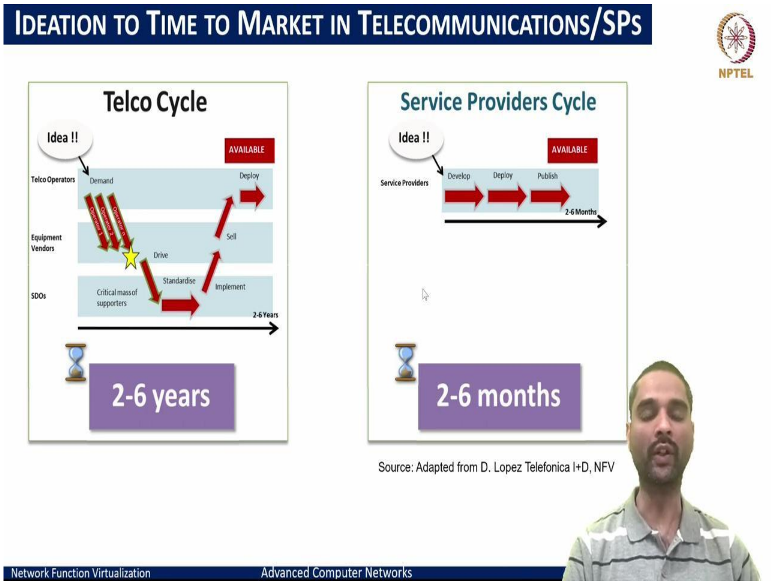

IDEATION TO TIME TO MARKET IN TELECOMMUNICATIONS/SPs

Source: Adapted from D. Lopez Telefonica I+D, NFV

And the most crucial thing being the deployment cycle itself. When you think of, we have the network services that we want to deploy or, we have an idea and to bring that idea to life, what it required for the telecommunication industry to go through? And the diagram over here on the left shows what exactly it means. So, you come up with an idea as a telecommunication operator or a researcher that you want to deploy, and then you need to build this telecom idea into the cycle.

And to do that, we need to ensure that whatever the hardware infrastructures that the telecommunication operators deploy need to be upgraded with the support that is necessary to ensure that kind of service or an idea translates to is deployed. In order to do that, what it meant for the telecommunication operators is to push their ideas onto the equipment vendors, who would then enable to build the required functionalities within those equipments. By equipments, we refer to the routers, the gateways, the switches, and most of the middleboxes that we look and primarily the middleboxes of various kinds.

And even if you think of having a new kind of middlebox, you need the equipment to go to the equipment vendor and get things done. And this would mean that there has to be multiple telecommunication operators who would want to invest on the equipment for long, and the equipments that would be then created, need to be procured in a large scale. That was often a barrier when you want to prototype and test.

Further, when the equipment vendors would take these things, they would then have to ensure that whatever the device is that they are going to deploy meet specific standards. So, they need to get clearance from the standards-developing organizations to ensure that they can readily build such devices, and then they can be operational at different infrastructures.

And this is where most of the time would go in bringing the idea onto the market. And once the standardization would be done, the next part would be to really implement adhering to those standards, and then equipment vendors need to build or produce those kinds of devices and then send them back to the telecommunication operators.

And once the cycle of selling these newer equipments back to the telecommunication operators will be done is when the real deployment of those initial ideas could continue. This would mean that the overall cycle had to go through multiple stages. And there would be multiple players involved in each of the phases. This would roughly account even if there were like a single flag change that you want to experiment with and try out or have to deploy, it will take as minimum as 2 years, and in many cases, 6 to 10 years on average would be the time that your small idea could get into practice.

And this bothered a lot for the telecommunication operators. But then when we contrast the same with the internet service providers, or the service providers on the internet lifecycle, like for example, consider Google, Facebook, Yahoo, who have certain ideas to deploy, they would eventually develop the ideas in-house, deploy them on their servers, and publish them to make them accessible to users. And all of this would need less than a year's worth of time or roughly around 2 to 6 months.

And these were made available for experimentation for users to try in a much quicker fashion. And if we contrast the two aspects here, the whole difference comes in terms of one being driven by the hardware where you are relying on the hardware, the purpose-built hardware to be made ready from the equipments and that is where a major lag in the telecom cycle comes from. While if you look at the service providers, it is all software where you will develop the new software, deploy it on the traditional server machines and run them out.

And this is where the exploitation of an idea and bringing it to the market is feasible. So, if you rethink of this, now, the telecommunication operators also started to see why not have a similar cycle for the telecommunication operator equipments. And what that meant is to get rid of this equipment that were purpose built hardware and move them to the software. And that is what we saw the NFV is all about. And these pin points actually propelled for having NFV the way what we see today.

(Refer Slide Time: 06:05)

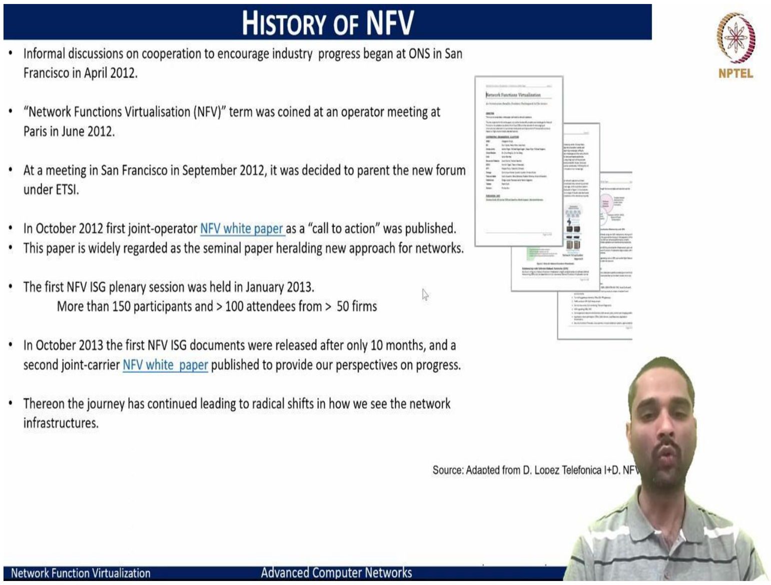

History of NFV

- Informal discussions on cooperation to encourage industry progress began at ONS in San Francisco in April 2012.
- "Network Functions Virtualisation (NFV)" term was coined at an operator meeting at Paris in June 2012.
- At a meeting in San Francisco in September 2012, it was decided to parent the new forum under ETSI.
- In October 2012 first joint-operator NFV white paper as a "call to action" was published.
- This paper is widely regarded as the seminal paper heralding new approach for networks.
- The first NFV ISG plenary session was held in January 2013.
  More than 150 participants and > 100 attendees from > 50 firms
- In October 2013 the first NFV ISG documents were released after only 10 months, and a second joint-carrier NFV white paper published to provide our perspectives on progress.
- Thereon the journey has continued leading to radical shifts in how we see the network infrastructures.

Source: Adapted from D. Lopez Telefonica I+D, NFV

NPTEL
Network Functions Virtualisation
Network Function Virtualization
Advanced Computer Networks

In this, if we look at a brief history of how this NFV evolved and all that happened, it was just about a decade ago, in early April 2012, in an open networking summit in San Francisco, the discussions started about how the network operators can leverage what we saw on the IT side as clouds, which were all virtualized infrastructures that were built to bring in the same aspects on to the telecommunication operator's framework.

Although it all started as a means for telecommunication operator's to get their operations done in a quicker time and make them manage their resources quite handily, but it also is applicable nonetheless to the cloud, to the data centers, to the internet service providers, to the campus networks, enterprise networks and all.

So, the first time this network function virtualization term got coined in the operator's meeting at Paris in around June 2012. And that led to what we now refer to as NFV in common. And then after this, continuous discussions amongst the different network operators from different countries, and telecommunication operators from across continent, continued to keep the momentum going and led to initiation that was basically put under the European Telecommunications Facilitate what we now see as a steering, as an organization to bring together what needs to be done to actually get solve these pressing knee problems for the telecommunication industry.

And these continued interactions under the umbrella of ETSI led to what we call as the first NFV white paper. And this was exactly the call to action in terms of what needs to be done or what were the aspects that we could look forward to basically deossify the telecommunication infrastructures and ensure that the innovations can be done very readily. And in many ways, this NFV white paper that is a call-to-action, published in October 2012, is seen as a very seminal paper that heralded the new approaches that we look forward to have in the networks.

And subsequently, the discussions continued and the NFV ISG sessions were set up and held in around January 2013, the first of the plenary sessions, where more than 150 participants participated, including many of the telecommunication and IT network operators, providing more than around 100 attendees with more than 50 firms. And this culminated into what we started to see as the specifications for how to realize such NFV infrastructure or NFV framework. All of this started to take shape.

And there on, we had a series of publications that started to come, a series of actions that led to the realization of the NFV, and the journey has continued so far and has been radical in defining what we see as what should be the scope for 5G, 6G networks going forward and what kind of network virtualizations can be applied there off for the edge infrastructures, all of these are taking shape as we speak today.

(Refer Slide Time: 09:40)

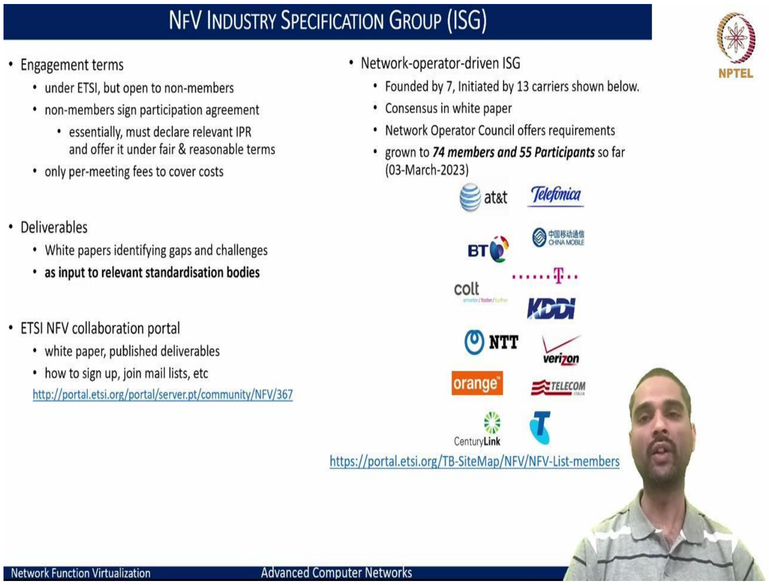

NFV Industry Specification Group (ISG)

- Engagement terms
  - under ETSI, but open to non-members
  - non-members sign participation agreement
    - essentially, must declare relevant IPR and offer it under fair & reasonable terms
  - only per-meeting fees to cover costs
- Deliverables
  - White papers identifying gaps and challenges
  - **as input to relevant standardisation bodies**
- ETSI NFV collaboration portal
  - white paper, published deliverables
  - how to sign up, join mail lists, etc
  - http://portal.etsi.org/portal/server.pt/community/NFV/367
- Network-operator-driven ISG
  - Founded by 7, Initiated by 13 carriers shown below.
  - Consensus in white paper
  - Network Operator Council offers requirements
  - grown to ***74 members and 55 Participants*** so far (03-March-2023)

https://portal.etsi.org/TB-SiteMap/NFV/NFV-List-members

Network Function Virtualization   Advanced Computer Networks

So if we think about this NFV industry specifications group, it was headed under ETSI, but it was open to any of the non-members to sign and participate for developing or coming up with the specifications, and the key deliverables that were considered were to come up with the core requirements and specifications of how to realize an NFV. And this collaborative portal of ETSI hosts all the publications so far, starting from early 2012 all the way up today, in terms of what specifications would shape the future of the networks.

And this network operator driven ISG or the industry specification group, initially, it was founded by just 7 of the telecommunication operators wherein 13 of the operators initiated the cause, the ones that are listed here, that is basically AT&T, Telefonica, British Telecom, China Mobile, KDDI, NTT and so on. So, they were the primary players, but then this group has grown leaps and bounds now with around 74 members across the continent and around 55 participants, who are also there from India; what we see is only Tata Elxsi being one of the participants. And none of the telecom operators so far seem to be part of this, but they are all trying to follow the frameworks and specifications that have been set up by this group.

(Refer Slide Time: 11:20)

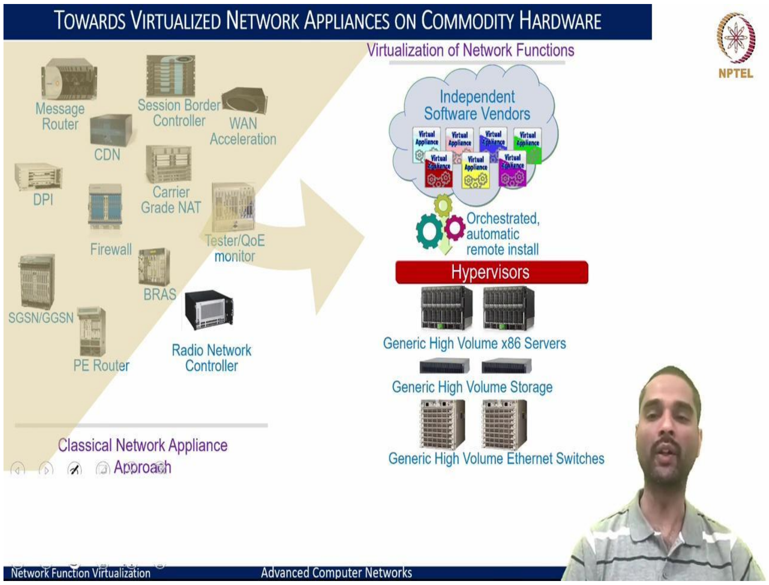

TOWARDS VIRTUALIZED NETWORK APPLIANCES ON COMMODITY HARDWARE

Virtualization of Network Functions

Message Router
Session Border Controller
WAN Acceleration
CDN
DPI
Carrier Grade NAT
Firewall
Tester/QoE monitor
SGSN/GGSN
BRAS
PE Router
Radio Network Controller

Classical Network Appliance Approach

Independent Software Vendors
Virtual Appliance
Orchestrated, automatic remote install
Hypervisors
Generic High Volume x86 Servers
Generic High Volume Storage
Generic High Volume Ethernet Switches

NPTEL

Network Function Virtualization
Advanced Computer Networks

So let us try to understand what this really meant in a nutshell, is that the telecommunication operators had to invariably deal with lots of these classical network appliances. And what this classic network appliance approach meant, is the use of proprietary chassis, that is, the hardware that are woven and provided for the telecommunication operators to manage, all of these may be using similar or dissimilar chips, that is, the compute ASICs that they would run on. But they would be in terms of operation and management, they would be completely different.

So the skills that you would need to manage a message router would be completely different than what you would need to manage a CDN or a SBC and likewise. And this created a lot of complexity through the entire lifecycle of managing these devices also the message router, you purchase it today, and then there is a newer one, there were no standards to say that it would be the same and it would depend even on the equipment manufacturer or the OEM vendor who would give.

So the infrastructure on the management aspects when we take device from Cisco would be completely different from what you would take from some other manufacturer, Broadcom or Arista or someone. So, all of these complexities in terms of how they would be designed, the procurement, the testing that you would need to do, the deployment, the configurations that you have to do post the deployments, or even if the devices go down, the means to repair, maintain, or replace these devices, all called for specialized expertise and high reliance on the equipment manufacturers' expertise. And this made it infeasible in terms of economies of scale for many of the telecommunication operators to manage and build these devices.

And hence, there was a need for a new mode of developing where we could get rid of these operator equipment devices and have a common mechanism where all of these could be virtualized and run on top of the generic compute storage and network devices, that is the x86 servers with generic storage and Ethernet switches, and then enable these instances of these classical network appliances to be run on top of these generic hardware's through the use of virtualization.

And once we have this kind of virtualization approach, we could then deploy all of these as the virtualized network functions. And now, once we have these virtualized network functions, these network functions could be developed and be brought to life by any of the software vendors. And that meant any of the independent software vendors, or ISPs could create and build the message router, the DPIs, and the firewalls, and it suddenly opens up the market for a lot more people to bring in and also for innovation in terms of researchers to try their own instances.

And this, once we build this kind of a framework, we would need to then further see how can this be rationalized because it is all good to create and enable people to try out but when we have multiple people trying out, should there be specific standards? Should there be specific needs or abstractions that we spoke about that can help build in a common, generalized way?

And that was the whole thought of this ETSI framework to come up, to facilitate this virtualized network function environment to create such a ecosystem in open and standardized, and to do this, what we started to see is how this NFV, if we have to build this, what are the key requirements that we need to come up with? How could we facilitate such an open environment? Those were the thoughts that were evolving.

(Refer Slide Time: 15:37)

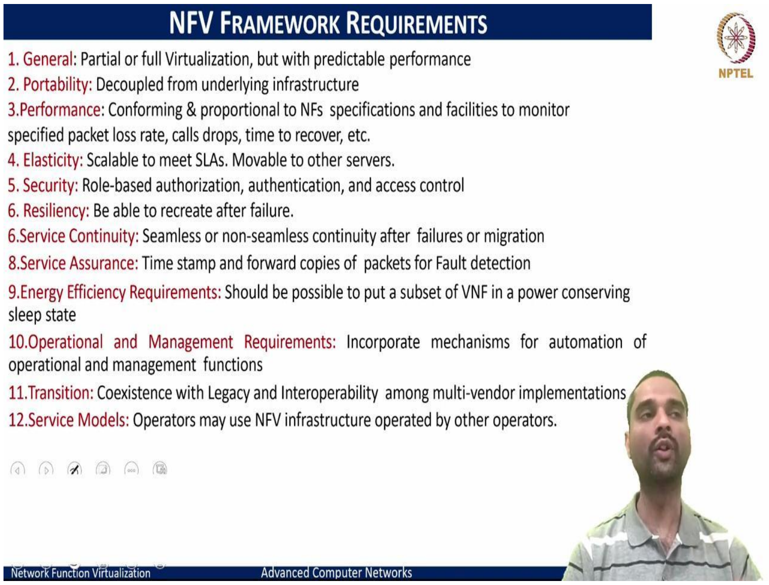

NFV FRAMEWORK REQUIREMENTS

1. General: Partial or full Virtualization, but with predictable performance
2. Portability: Decoupled from underlying infrastructure
3.Performance: Conforming & proportional to NFs specifications and facilities to monitor specified packet loss rate, calls drops, time to recover, etc.
4. Elasticity: Scalable to meet SLAs. Movable to other servers.
5. Security: Role-based authorization, authentication, and access control
6. Resiliency: Be able to recreate after failure.
6.Service Continuity: Seamless or non-seamless continuity after failures or migration
8.Service Assurance: Time stamp and forward copies of packets for Fault detection
9.Energy Efficiency Requirements: Should be possible to put a subset of VNF in a power conserving sleep state
10.Operational and Management Requirements: Incorporate mechanisms for automation of operational and management functions
11.Transition: Coexistence with Legacy and Interoperability among multi-vendor implementations
12.Service Models: Operators may use NFV infrastructure operated by other operators.

Network Function Virtualization Advanced Computer Networks

And as a continued discussion on this, they came up with 12 key NFV framework requirements. One being, that it has to be general, in the sense that having a generic infrastructure that is not coupled with the design of any of the network appliances, we do not want to stick to designing the APIs or abstractions based on a subset of the network appliances. It could be a security appliance, it could be an in-network, custom network appliance. We want to ensure that whatever the abstractions that we want to build, whatever the interfaces that we want to build are generic enough to accommodate all kinds of network appliances that we saw earlier. And it could also employ different kinds of virtualization like we have seen in virtual machines. When we deploy, there is paravirtualization, there is full virtualization and host-based virtualization. So, we should be able to facilitate and support all of those modes of virtualization. And based on the kinds of equipment or appliances that you want to build, they may suit to choose either of these virtualization models. But with a caution that whenever we choose any of these virtualization models or build the virtual appliances, we need to also guarantee predictable performance in terms of what kinds of throughput, what kinds of latency guarantees, or service level agreements these NF instances need to support, need to be still provided.

Second, is about portability. That is, if we have the infrastructure and want to build these network appliances, these network appliances that we want to build should be decoupled from the underlying infrastructure. That means we read the abstraction or APIs at the virtualization layer to say how the NF instances can be used to build these portable instances that can be run on any of the infrastructures. So, we should not be tied to just one kind of compute or one kind of a storage we should be able to leverage any kind of compute, storage and network elements and build our NFV instances.

Third, is the performance like when we want to build these custom middleboxes, or custom network functions, these functions have to meet the requirements of whatever the performance or SLA guarantees that they need to do and facilitate the means to monitor their service aspects in terms of what is the packet loss rate that they are having. Like, when we are talking about telecommunications, what is the call drop rate that they are having, or if it fails, what is the time to recover? All of these aspects with respect to the performance of NFV instances need to be supported.

Next is elasticity. That is, we know that in the real world, these telecommunications infrastructures have to cater to varying demands in terms of the traffic load. So, if we see that there is a variation in the traffic load, we would want to scale and adapt to the current traffic load. And that means they have to be scalable. And if they are virtualized, we will also want them to be movable from one server to the other because one of the physical servers at a time may be overloaded, and other under provision, we may want to move dynamically these services. In a way, we want to facilitate elastic services for the deployment of the network functions.

Next and fundamentally more important is also the security where any access to these network functions or any access to critical data be safeguarded. That means you would want to facilitate role-based authorization and authentication and also enable access control to specific data that you would want to enable, to this data or functions that would run on behalf of users. Think of a billing or a charging system that is running in telecommunication operators that needs to ensure that it is updated only by those who are authorized to make the changes and runs properly.

And next up is once we have the devices, especially the virtual machines, they are prone to failures. So, failure is a common scenario. But failure of an instance, or failure of a device does not have to translate to the failure of a service itself. So, we want to build a resiliency framework where if things fail, the services can be resilient to the failures of the virtual or the physical infrastructure in a way that we should not be able to recreate upon failures and run the services seamlessly so that we ensure the service continuity even in the events of failures. In a way, like when you think of VMs, we can migrate a VM from one machine to the other; when we know that we want to have a downtime because of some hardware updates or patches, we could still run the service on some other device for a time. So, likewise, if there is a failure, we would want to ensure that the services can still continue to operate through the means of running these virtual instances on any of the active devices.

And doing so should also be effective in the sense that the assurance criteria in terms of when the services we want to migrate, you should not be that okay, it is down tonight, and we will bring it up only after 24 hours; we should have control over what kind of assurances we give in terms of when we have recovery. And for that, localized fault detection, quick fault detection of any of the components in the NFV framework should be facilitated.

And once we detect such faults, the recovery needs to be time-bound, in the sense that we can have the infrastructure to be able to recuperate and ensure that the services can be run on top in a seamless fashion.

Also, when we think of these NFV infrastructures that we want to deploy, and it is also imperative that we need to honor to the energy requirements. And we know that these traffic loads, the demands keep on varying, so if we have brought up certain network functions, and they are doing nothing, we should be able to save power by ensuring that some of these network functions can go to sleep state, or they can be hibernated. And these means to ensure that we can harvest on energy would also be important.

Next, when we have these kinds of networks functions that are virtualized functions that we want to deploy, and potentially on a large scale over an infrastructure, we would want to have rich control over how to operate, bring those into operation, manage their lifecycle of bringing up taking services down and also have the means to automate all of these activities. Thinking of doing this manually would be almost infeasible when we think of 1000s of machines that we want to run on a switch.

And then this is where automation becomes an important aspect for managing these virtual network functions. And thus, the framework should facilitate for such automation requirements. Also, these frameworks, when we develop the deployments, could be done by different many different telecommunication operators. But it does not confine that an NFV solution would run only on one kind of infrastructure while not on the other. So, that means we need to have the coexistence of the telecommunication infrastructures for these NFV.

And also, whatever the existing hardware infrastructure they had, it has to coexist with the existing infrastructure. It is not that it is either the hardware infrastructure that you have versus the NFV. And what, in essence, it means is that we want this NFV framework to be incrementally deployable over the existing infrastructure and also facilitate for having the solutions that are going to operate from multiple vendors on the same infrastructure.

Thus these 12 key NFV framework requirements were charted out lower a series of discussions starting in early 2012, which led to a series of specifications that eventually were pulled out from the ETSI, NFV, ISG.

(Refer Slide Time: 24:20)

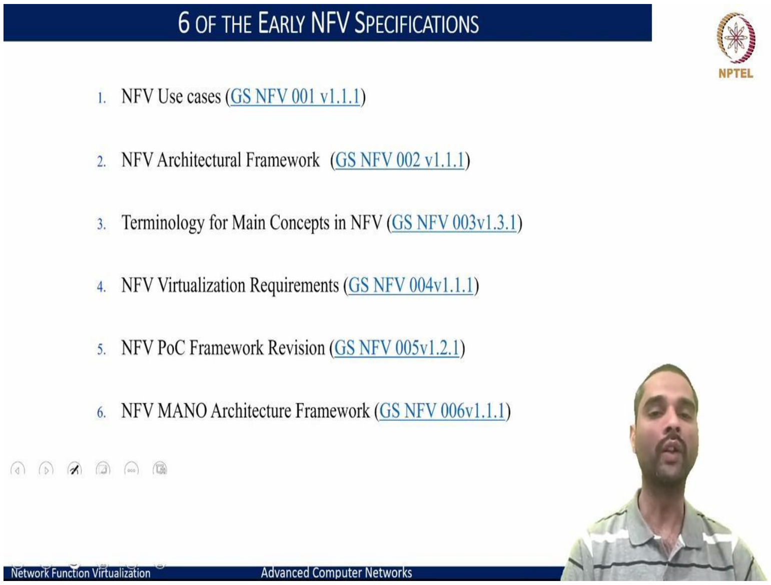

6 OF THE EARLY NFV SPECIFICATIONS

1. NFV Use cases (GS NFV 001 v1.1.1)
2. NFV Architectural Framework (GS NFV 002 v1.1.1)
3. Terminology for Main Concepts in NFV (GS NFV 003v1.3.1)
4. NFV Virtualization Requirements (GS NFV 004v1.1.1)
5. NFV PoC Framework Revision (GS NFV 005v1.2.1)
6. NFV MANO Architecture Framework (GS NFV 006v1.1.1)

And the core six early NFV specifications that came out, were, one, in terms of what are the key use cases that we are trying to think of? And the use cases, not just confined to the telecommunication operator world, but also those were applicable for the IT and Cloud infrastructures, the enterprise networks, campus networks, and even the internet service provider networks or home networks where ISPs could deploy the gateways.

So, most of the use cases were charted out so that it became an incentive for many people to start thinking about NFV, bring in new dimensions, and also innovate in each of these areas. Second, when we want to construct such a framework, we need proper architectural hindsight, and to do that, the NFV group also came up with the NFV architectural framework that we will discuss in a bit. And then once we have to come up with the architecture and make things evolve, there is a lot of jargon, and especially when we think of networks, it is full of jargon, and telecommunications is even more. And we are now trying to bring in the NFV, there were a lot more newer terminologies. But the sense of having the same meaning for a term was much more important, and that is where this terminology for main concepts in the NFV was also charted out so that everyone is on the same page, and thinking of the same terms, and speaking the same language. And then the main core aspect of virtualization, what were the key requirements like the ones that we just discussed, were all charted out in terms of in a much broader sense, in a much more depth manner to say what exactly each of these aspects need to look and how these need to be realized.

And those became one of the other specifications. And then the proof of concept because having everything on terms or on words is of no good until we have a proof of concept framework that can be developed and showcased. And this led to the NFV proof of concept frameworks that were being done and revised over time to see how things can be demonstrated; not necessarily that we have to build the entire framework bottom up, but key pieces that we can bring one at a time, one piece at a time and showcase the benefits that they bring along was the path that was chosen, and these frameworks have evolved thereof.

And overall, like when we think of this NFV infrastructure, where we will have a lot of independent software vendors who are running their virtual appliances, we need a framework to manage and orchestrate these functions, a common framework, which we call NFV MANO or management and network orchestration part were also the specifications of how to do , what are the aspects to do, were also charted out in this early of the 6 NFV specifications.

And ever since the documentation, the outcomes have been tremendous, and there have been a lot more in terms of what should be the performance criteria, what should be the resilience criteria, security criteria, all of these specifications were charted out for different aspects.

(Refer Slide Time: 27:44)

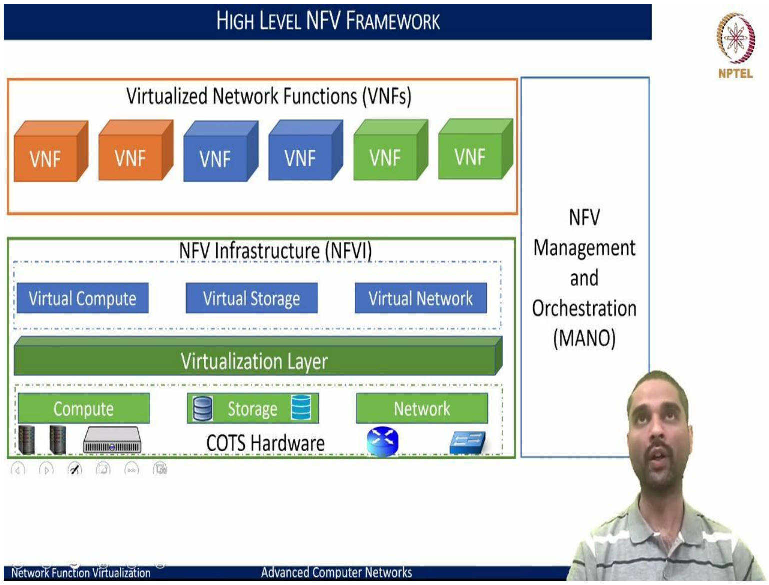

HIGH LEVEL NFV FRAMEWORK

Network Function Virtualization Advanced Computer Networks

And through this, let us now look into the NFV framework, try to understand what it really means, and then look at how this architecture is eventually realized. So, this NFV infrastructure envisages the implementation of the network functions that are software-only entities, which would be run on a common infrastructure. And when we think of this commodity infrastructure, it includes compute, storage, and network elements, which are the commercial off-the-shelf available hardware.

And on top of this hardware, we would want to bring in the virtualization layer, also known as the hypervisor, which basically abstracts the underlying hardware resources, and decouples the VNFs that we want to build from these hardware resources themselves. And what that means is this virtualization layer is then responsible to provide and present the virtualized compute, storage, and network elements in a way that we have logically partitioned the physical resources and abstracted them out and presented the virtual compute, storage and network elements, just as we see with the virtual machines that we work on.

Second, this also has to enable the software implementation of the VNFs to use this underlying virtual infrastructure. That means, if I have to develop a network function, it would use the constructs of the virtualization layer to access the virtual compute, virtual storage and virtual network elements. And this infrastructure is what we call as the NFV infrastructure which encompasses both, the virtual compute storage and network elements built on top of the hardware. And on top of this infrastructure is where any of the virtualized network functions or VNFs can be realized.

And these virtual network functions are, in essence, the software implementation of a network function that is going to be run on top of this NFVI. And now, you can think of multiple vendors trying to come up with different network functions, and each is going to use the common abstractions of the virtualization layer presented through the NFVI to build these virtualized network functions. Now, this part allows us to bring in innovation and deploy new network functions.

But thinking from a telecommunication operator's perspective, if I have an infrastructure where I have to deploy this, I need to manage all of these virtualized network functions, I need to manage my hardware resources, I need to manage my virtualized network infrastructure, all of it. So, the other important piece that we would now need is this management and orchestration framework.

And what this management and orchestration framework is supposed to do is ensure that you are able to have control of how to manage your NFV infrastructure, and how to manage the lifecycle of these virtualized network functions, and ensure that things can be run seamlessly and smoothly and be operated in an automated fashion for running all of these VNFs.

(Refer Slide Time: 31:17)

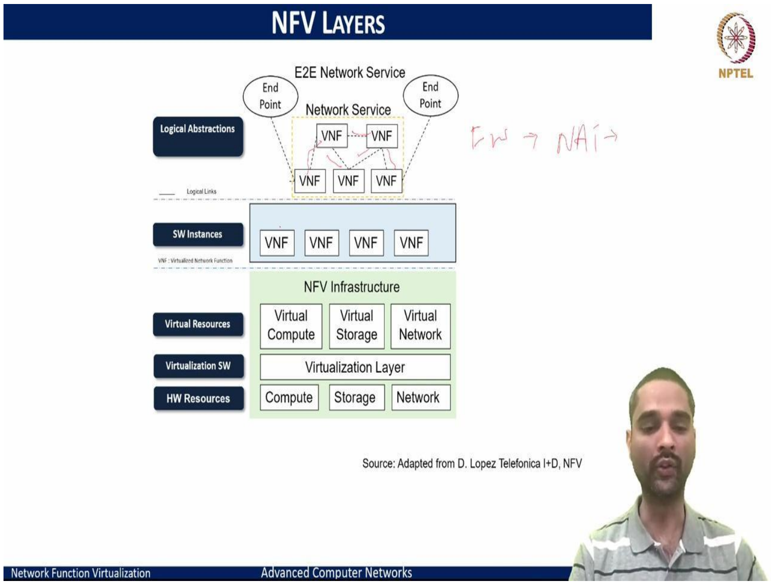

NFV LAYERS

Source: Adapted from D. Lopez Telefonica I+D, NFV

Network Function Virtualization Advanced Computer Networks

And this framework, now, if we rethink of how it has been layered, at the bottom, we have the hardware resources, and on top of this, we have built the virtualization layer. So, it is a virtualization software that virtualizes all of these hardware resources and presents what we can see as virtual resources. And that encompasses the underlying NFV infrastructure with these three layers. And on top of this, we are building these software instances of the virtualized network functions.

So these virtualized network functions could be developed by any of the vendors, any third party can create and deploy the software instances. Now, we can see that the development of VNFs resembles what we saw earlier about how the service providers would develop and deploy the functionalities locally. So, this would give the same freedom now for any telecommunication operators to develop and deploy any of the services that they want to try out and experiment out in a very small fashion and bring it to the market in a much more faster way.

And that is where NFV is thought as an accelerator, and innovator, in bringing things to the world. And the other aspect that we have to see once we have built these software instances, what we really need is to provide the network services, and the network services need not necessarily be just one VNF at a time and like we saw earlier when we think of enterprise networks, data center networks, there would be a lot of network service functions that could be put together or stitched together to provide a common service. For ensuring the network-wide policies that we want to honor in a campus or an enterprise network, we have to stitch together multiple of these VNFs to communicate.

So, see, observe now that these VNFs are running on top of the NFV infrastructure, and now we have to create the logical links to ensure that these VNFs can interact and communicate amongst themselves and provide the notion of the network service that we want to provide.

So, this means that there is a need to provide these logical abstractions of how to stitch these network functions and how to facilitate a chain of services to provide an aspect of what we want to present as a network policy or realize as a network policy or present different services to the end users.

And this is what we call as the virtual network function forwarding graph. That means another implicit requirement would be that the NFV framework should facilitate to have or realize any of the forwarding graphs or a service chain for connecting or interconnecting these virtualized network functions. For example, VNF 1 could be a firewall and VNF 2 could be a NAT or a load balancer, and so on. So, we should be able to facilitate connecting these and route or steer the traffic amongst these network functions.

Also, when we speak of these virtualized network functions, these network functions may span a single instance that would run and could be run on replicated at different locations. So, there is also a location or Point of Presence where these would work and operate when we want to build primary and secondary instances that would want to communicate and then build failure-resilient infrastructure for ensuring that you have fault tolerance in the system; then we can think of many of these VNFs that are also going to talk to the peers, either in active-active mode or active-standby mode to ensure that the services can exchange information amongst themselves and keep the states up to date.

Thus, this NFV as a framework and as a layer can be realized to provide lots of logical abstractions to build the network services. But now we need to understand and look into more detail of what are all the abstraction points or what are the interfaces that we need to build for these cases.