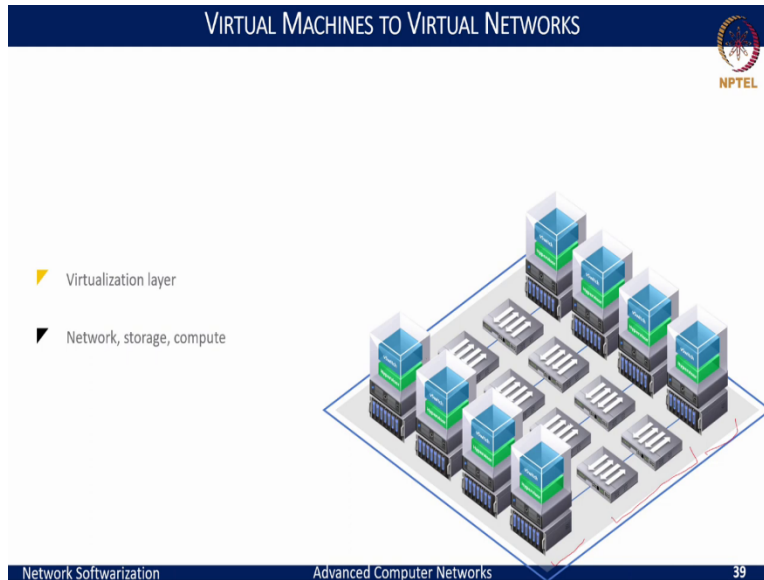


Advanced Computer Networks
Professor Dr. Sameer Kulkarni
Department of Computer Science Engineering
Indian Institute of Technology, Gandhinagar
Lecture 27
Network Virtualization: Part 2

(Refer Slide Time: 0:19)

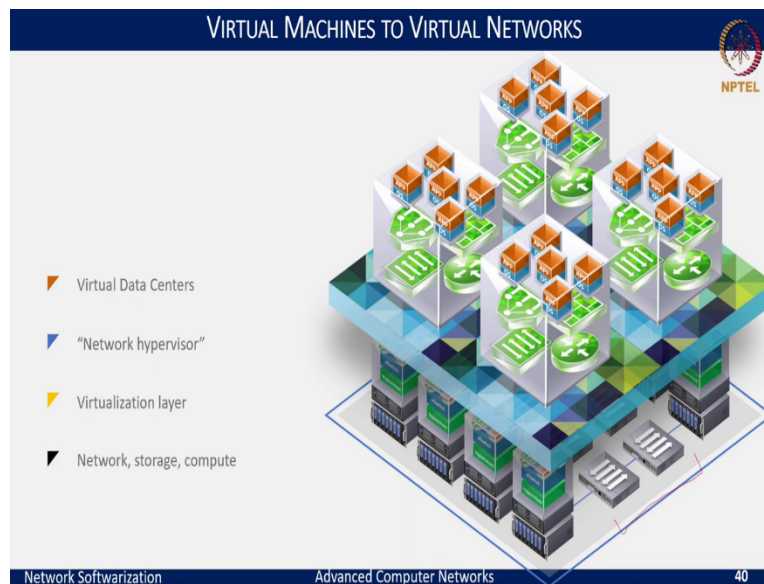


Let us try to understand about Network Virtualization in a bit more detailed fashion. The slide here with the picture shows exactly the composition of a typical enterprise network or a campus network, or even it could be mapped to a small-scale data center where we have the physical compute devices and storage devices and also the network elements like the switches and routers that connect these compute and storage elements.

And when we look at this kind of infrastructure and when we see how the virtual machines were envisioned, we think of having the hypervisors that are built on top of the compute elements which abstract out the underlying hardware details and enable us to build the virtual machines on top of each of these compute and storage elements.

And likewise, the v-switch component here enables us to build the networking for these virtual machines that are built on top of the hypervisor. And the v-switch component is acting as a virtualization layer within the hypervisor within the perimeter of each of the physical resources or a physical compute device.

(Refer Slide Time: 1:40)



And thus, through this abstraction, we would be able to implement and have what we can call or think of as virtual routers, virtual switches, virtual firewalls, and other networking elements that would want to implement. But this would be within the scope of a single device or a single physical resource. What we can achieve in this is to have a bunch of networking elements that can be constituted and be able to make the virtual machines communicate and interact amongst these devices. Note, however, that we have not covered the actual networking elements or the physical switches and the routers in our vision of this virtualization layer.


Hence, if we want to build virtual routers and virtual switches that can enable us to communicate amongst multiple of these compute and storage devices and still present the same set of right abstractions with which we are able to build the virtual machines on a single physical resource, we want to build this networking capability on top of these resources. Then we need what we would call as a network hypervisor that would encompass all of these physical resources as well as the virtual machines that we would have built and enable the interfaces so that we can rightly build what we can think of as the real virtual networks and even to build what we can see as virtual data centers.

So, here now, you can see that we are able to truly build the virtual routers that can connect anywhere across the physical infrastructure and even likewise the virtual switches and routers

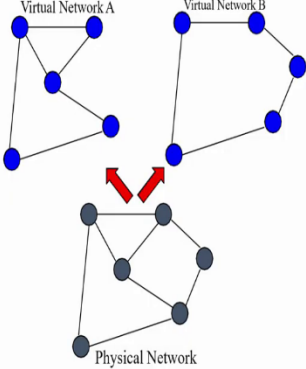
and run the specific workloads, the networking applications on top of these well-interconnected network hypervisors.

(Refer Slide Time: 3:36)

NETWORK VIRTUALIZATION – A MEANS AND ENDS TO BREAK THE IMPASSE


NPTEL

- Network virtualization offers **all-sizes-fit-into-one** solution
- Open and expandable model
 - multiple heterogeneous architectures on shared physical substrate
 - promotes innovation and customized services/applications
- Testbed for deployment/evaluation of new network architectures/protocols



Network Softwarization Advanced Computer Networks 41

And to achieve this is the means for using network virtualization, and hence this network virtualization would then allow the administrators to create multiple virtual networks wherein we can logically segment or group the physical networks and make them operate as single or multiple independent networks. And these virtual networks may span across both the virtual machines as well as the physical networks and also share the physical and virtual switches and network resources.

To sum it, what network virtualization offers is as all sizes fit into one solution. In the sense that now we are able to break and disentangle the physical infrastructure and build the virtual infrastructure the way we would want to suit the needs of our elements and also expand the model that we would want based on the characteristics that is independent of the physical network.

And this, in a way, provides us also to mix and match or have multiple heterogeneous architectures that would be running on the same physical network. And most importantly, this also enables us, to have the innovations of the networking technologies done independently without having to affect or disrupt any of the running ongoing networking services, and

instantaneously such a method would promote for innovations and help build customized services and applications much more freely.

And such a thing would give us a testbed for the deployment and evaluation of new network architectures, which was the need of the hour. And overall, this network virtualization enables us to enforce the routing for communications between the virtual networks the way we would want them to do.

We could even restrict or manage the traffic and also enable the functional grouping of the nodes within the network like we have seen a virtual network A and virtual network B grouped on some principles of logical services that we would want to build. Nonetheless, all of this appears as a physical network that are independent on their own. One cannot distinguish between a virtual network and a physical network while operating on just the desired virtual networks.

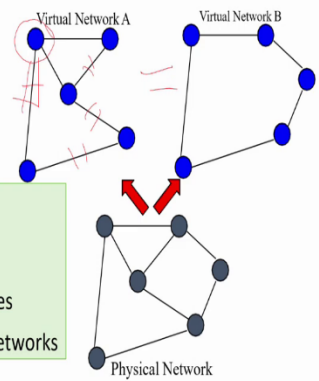
(Refer Slide Time: 6:08)

NETWORK VIRTUALIZATION – SLICE THE NETWORK!

- Decouples the services provided by a network from the physical infrastructure
- Virtual network is a “container” of network services, provisioned by software
- Allows the creation of multiple virtual networks on same physical substrate
- Allows for faithful reproduction of services provided by a physical network

Each virtual network:

- is a collection of virtual nodes and virtual links
- is a subset of underlying physical network resources
- co-exists with, but is isolated from, other virtual networks



http://www.opennetsummit.org/pdf/2013/presentations/bruce_davie.pdf

Network Softwarization Advanced Computer Networks 42

So, to sum up, what network virtualization enables us is to have a slice of a network and enables us to use it the way a user wants to use it without worrying about any of the other implications that it may have or disruptions that it may cause to anywhere else. And in a way, this is possible only because it is able to decouple the services that are provided by the physical network from the services that you would want to extract or choose out from the virtual networks.

And in essence, this virtual network is like a container of network services that would be provisioned by the software and instantiated on demand. While the physical structure or the

physical network is set in stone and built, which cannot be changed anymore. So, this allows for the creation of multiple virtual networks on the same physical substrate, and more importantly, it is not just about creating the virtual networks, but also it enables the faithful reproduction of the services that would be provided by the underlying physical network. Since these virtual networks that we have the ability to build share the same underlying physical network.

Hence, we can consider each virtual network as a collection of virtual nodes, that is, the virtual routers, switches, and the workloads that would operate on these virtual elements, including the virtual links that enable to connect these virtual elements. And thus, it is a subset of the underlying physical network resources and coexists with other such virtual networks that are isolated from each other.

(Refer Slide Time: 7:57)

NETWORK VIRTUALIZATION – SLICE THE NETWORK!

- Sharing the network and Topology Abstraction
 - Arbitrary/custom network topologies can be realized very easily.
 - Different controllers for different users/traffic; Isolation (bandwidth, flow space)


- Experiments vs. operational network
 - Support research without breaking real services
 - Failure resiliency

- Multiple administrative groups & multiple customers
 - Researchers on a shared infrastructure
 - Different departments on a campus

- Multiple services or applications in one domain
 - Tenants in a shared data center

Facets of Network Virtualization

- Virtual LAN (VLAN)
- Virtual Extensible LAN (VXLAN)
- Virtual private network (VPN)
- Active & programmable networks
- Overlay networks



Network SoftwarizationAdvanced Computer Networks43

And to this means we are now able to share the network and build the desired topology abstractions, and by this, what I mean is that we can now have the ability to build arbitrary or custom network topologies which may not even exist. Meaning we would be able to alter the topology of the network link characteristics and test or experiment out without even worrying to actually deploy such a physical infrastructure.

Second, when we share, it also means that we would want this sharing to be done in a way that the usage of the shared resources does not intervene or intervene with each other, and hence the

isolation guarantees through virtualization by means of having different controllers for different users and traffic is also crucial in this mode of network virtualization.

With all of this, it greatly enables the researchers to do carry out their operational experiments, whether in bringing the innovative networking technologies or trying out some protocols which otherwise may not be supported by many of the physical devices much more readily and easily.

And also, very importantly if any of the services on a network virtualization pattern fail they are confined to a specific virtual group, and in essence, they would also try to meet video resiliency in a way that when any of the experiments are run simultaneously, the breaking of a one would not affect the other and more importantly we would also see that there is a need for multiple administrative groups and to facilitate the support for multiple customers. As researchers are going to be operating on a shared infrastructure, everyone would want to control the way their network would want to appear, the way their network topology, the link characteristics, the loss characteristics, delay or bandwidth characteristics, all these aspects. That means we need to enable the management and administration of these virtual networks to various people.

As an analogy, if you think of different departments that stay on the same campus, every department would want to enforce or have some control over the parameters that you would want to define and likewise with respect to the networks. And overall, with this kind of network virtualization, where we are able to slice the physical network into multiples of virtual networks, this also lends us to build multiple services on the same physical infrastructure and enable the support for multiple applications in one single domain. And in fact, this as in model helped build what we call as the shared data centers where multiple tenants can use and reuse the same underlying physical infrastructure. And this is the aspect brings multi-tenancy in data centers.

And overall, when we want to build the abstractions for topology and provide the customization capabilities for the user to build specific topologies, quick network virtualization it becomes very easy. And if we look at the key facets that have evolved with network virtualization, we can see and think of the virtual LAN which allows to partition the group of devices into communication domains to provide better security, monitoring, and management capabilities.

And an extension of the virtual LAN is the VX LAN or virtual extensible LAN. This, on the other hand, enables to extend the layer 2 subnets across the layer 3 networks, and they provide

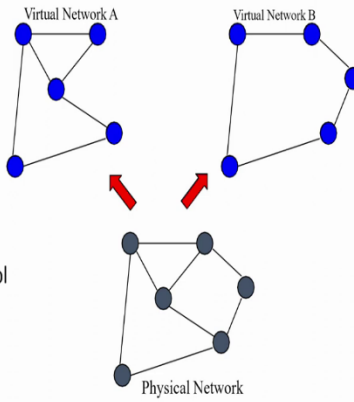
similar benefits as that of a VLAN across wide area networks, that is, to provide better security, monitoring, and management capabilities amongst the WAN networks.

And the other one that we typically hear is the virtual private network or the VPN's which allows users from anywhere on the internet to connect to one or more resources on a campus or enterprise network. It allows to securely connect and access the resources that reside within a confined domain, and also the other facets include the active and programmable networks that we will speak about in the coming days and also we looked at the overlay networks, which really enabled us to build at an application level the isolated networking workloads that can run and coexist and run together over the same physical substrate.

(Refer Slide Time: 12:53)

WHY ABSTRACT THE TOPOLOGY?

- Single Physical network appears as multiple logical networks.



- **Simplicity**
 - Hide inessential details, churn, migration, ...
- **Privacy**
 - Hide internal details of the network
- **Scalability**
 - Present a smaller topology and fewer events
- **Partial deployment**
 - Tunnel through components you don't control
- **Experimentation**
 - Try topologies that don't really exist

Network Softwarization Advanced Computer Networks 44

So, the question may come like why we would want to abstract the topology or have this characteristic to build a different kind of network links than what we have in the underlying physical network. And this stems from various aspects; one is the simplicity, that is, you would want to create the network the way that you would want without even trying to have the burden of trying to replicate the same physical infrastructure.

Let us say we want to test out a protocol and see how it behaves when there are a lot of packet losses. Then we cannot tweak the parameters on the physical links, but if we have the virtual links and virtual nodes, we can control these explicit characteristics of what would be the delay in packet processing, what would be the loss rate that it would encounter, and so on.

So, this provides us with a very simple means to configure and set these characteristics. Also, our physical infrastructure may be limited in its way, in its size, and because of economic constraints, we may not be able to scale it out. But now, with virtual networks, we would be able to scale the topologies at will without any of the cost overheads, and this allows us to experiment and see what would happen for various protocols and various technologies when we try to deploy and see how they would work with changing network parameters.


And also, in a sense, these virtual networks are isolated and provide us the means to run specific aspects in a much more contained fashion, or they promote for privacy while trying to hide the internal details of the network.

And lastly, it is also about experimentation, where we would want to try out specific things on the topologies that in the real world may not even exist. So, this gives us the full freedom and flexibility to try out many things, experiment the way the users would want, and try to test out the things before any things can be deployed on the scale of the internet.

(Refer Slide Time: 15:00)

OVERLAY NETWORKS: PLANETLAB – A SUCCESSFUL STORY!

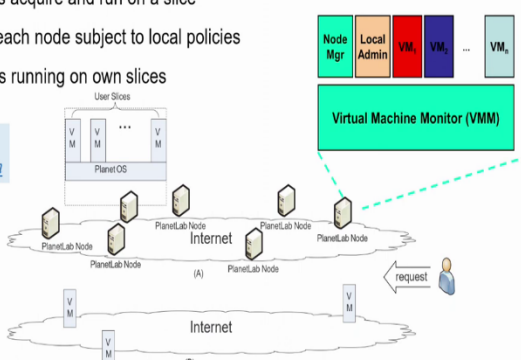
- Overlay testbed for geographically distributed network services.
- At its peak, PlanetLab consisted of 1353 nodes at 717 sites spanning 48 countries.
- **Service-oriented** network architecture:
 - sliceability → applications acquire and run on a slice
 - decentralized control → each node subject to local policies
 - management sub-services running on own slices



Decommissioned in May 2020!
Farewell Note by Larry Peterson

Node Mgr	Local Admin	VM ₁	VM ₂	VM _n
----------	-------------	-----------------	-----------------	-----------------

Virtual Machine Monitor (VMM)



Network Softwarization
Advanced Computer Networks
45

So, let us consider one example infrastructure that was built using the overlay networks that we just discussed about and the most successful and history is the PlanetLab story, and this PlanetLab is a planet-wide testbed for the R&D of network applications and distributed computing, and in its prime, it spanned over 1300 nodes including around 717 geographically

distributed sites, across 48 different countries primarily in the U.S and Europe and even in India we had the PlanetLab presence.

What this PlanetLab led was to build the service-oriented network architecture on which the researchers could deploy and run and experiment out their networking aspects or even the distributed computing or distributed protocols that they would want to try out and see how they would really run over this geographically distributed network.

And this was the first service-oriented network architecture presented by the works that we saw and when we discussed earlier about Larry Peterson's work and by David Culler and others, led to this building of the PlanetLab with the virtualized networks and virtualized testbed. This provided the means to the resources in a sliceable fashion where applications can acquire the set of slices and run them on the desired set of slices, and it also presented the way for a decentralized control wherein each node could be subject to different local policies and still be able to run all the network services.

And if we look at this, the underlying substrate was the physical internet on top of which the PlanetLab's overlay network was built, which included what we call as various of these PlanetLab nodes ranging in the number of thousand 1353 nodes across the world and what each of the PlanetLab nodes gave was to have a virtualized environment on which you would have the hypervisor, the Planet OS on top and you could instantiate various of the virtual machines that constituted the users slice on that PlanetLab node.

And multiple of PlanetLab nodes spread across the geographical region could communicate and could have slices in different regions, and if we look into what this PlanetLab node really constitutes was to provide a virtual machine monitor, the Planet OS and you have various user-defined virtual machines that would be the workloads that would run for the user. And also the administration and management aspects through the node manager and local admin support. And the users would request for specific PlanetLab nodes and deploy their use cases, change or enable them to apply whatever the protocols, whatever the networking technologies that you would want to deploy and experiment out. As if they are experimenting on the real internet.

And this service ran for around 18 years, and eventually, it got decommissioned very recently in May 2020. And now we have a lot more other kinds of services, including the CloudLab that is

been set up at the U.S site. But this success story really enabled for the users to come out of the impasse that the internet had posed and readily try out and experiment on such a large-scale geographically distributed PlanetLab network.

(Refer Slide Time: 18:46)

MORE FOOD FOR THOUGHT - RELEVANT PUBLICATIONS FOR THE INTERESTED



[[rfc3724](#)] “The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture” - Kempf et al.

[[Blumenthal](#)] “Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world”, *ACM Transactions on Internet Technology*, Vol. 1, No. 1, August 2001, pp 70-109. [~700 citations]

This is an Optional Reading material for the interested!

There are some other pertinent and most relevant works that I consider are worth reading, and one of them is an RFC this RFC 3724, the rise of the middle, and the future of end-to-end reflections on the evolution of the internet architecture by Kempf et al. It essentially re-examines the interpretation of end-to-end arguments starting from the early internet days up until 2004.

In fact, it bats for the end-to-end principle stating that the end user choice and empowerment, the integrity of the networking services, support for trust, and also, if you consider the good network citizen behavior. These, in fact, have stemmed and developed as a consequence of the end-to-end principle.

Nonetheless, with the increasing pressure to incorporate the services within the network as we saw with the rise of the middle boxes and various other states that are being now plumped into the network itself, it emphasizes and questions that any such proposal where we would want to incorporate these services which would add on to the state in the network, we need to weigh against the merits in a very thorough manner in terms of how it affects the integrity of the service, how it affects the support for trust and how it is in fact affecting the end user's choices.

And all these cases need to be thought through before even such proposals we met, and that was the whole crux of this RFC 3724, but it really gives a good picture of how this end-to-end argument which started in the early days of an argument, like I said, propelled for different design principles and in itself being the fundamental principle how its interpretation changed over time in terms of the middle and the future that was considered in the early 2000s.

And the other work by Blumenthal and David Clark rethinking the design of the internet, the end-to-end argument versus the brave new world. It is a very interesting paper, and of all the changes that are transforming the internet, it batted to reconsider the loss of trust as early in the early 2000s to see what are the aspects that we need to think when we are trying to think of tomorrow's internet. And it also emphasized that the simple model of an early internet that is a group of mutually trusting users attached to a network is gone forever and said that we should really be thinking for a tomorrow as global communications which are eroded of the global trust but only based on the local trust.

And it also operates two specific pictures of the constraints that the technology imposed on the future internet. One of them being the technological solutions that being fixed and rigid in a way and what we discussed at lines about the ossification, and the other, a continuing tussle between those who would impose the controls and those who would evade them. And it calls to rethink the trust and security aspects in the internet.

I have given the links for which you can access the resource materials but note that this is an optional reading material, and for the interested, I would really recommend that you go read these works.

(Refer Slide Time: 22:22)



- Evolution of the Internet
 - A Historical perspective and meteoric rise of the Internet

- Internet Impasse and Ossification
 - Different aspects that created the barriers for networking innovation

- Network Virtualization as a means and Ends to break the Impasse!
 - How Network virtualization and Service oriented architecture came to the rescue.

- Next up in Network Softwarization:
 - Road to Software-Defined Networking

And to summarize, we have looked at the evolution of the internet starting from the early days of 1960s, starting as a packet-switched network, to its meteoric rise of the internet, where almost half the world's population is hooked to it as of today.

And in doing so, we looked at the inherent design principles and learned that as with any highly successful technology, the ossification of the internet was also a natural evolutionary stage. And this problem in fact, was very acute in the context of networking technologies because it has started to shield effective computation and also became obstacle to deploy any of the innovative networking technologies, which is also compounded by the high cost of the infrastructure and the need for agreement amongst the large number of stakeholders and organizations who often were seen to have competing interests and these aspects created the barriers for network Innovations. This impasse, network virtualization was considered as the means to rescue from this phase, and virtualization support in network devices helped overcome internet ossification and also to address specialized requirements to deploy and experiment innovative networking technologies and protocols on the virtual testbeds without disrupting the established internet and over the next few lectures, we will then start looking into SDN or the software-defined networking.

