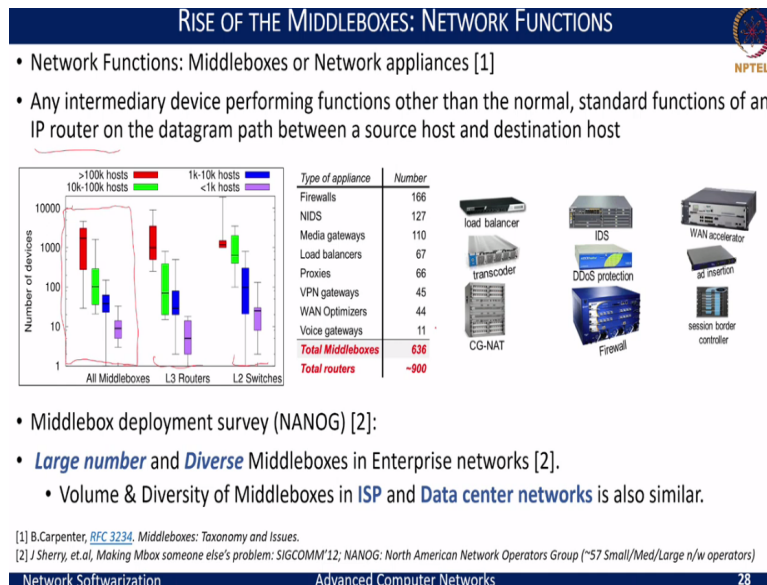


**Advanced Computer Networks**  
**Professor Dr. Sameer Kulkarni**  
**Department of Computer Science Engineering**  
**Indian Institute of Technology, Gandhinagar**  
**Lecture 25**  
**Network Ossification**

(Refer Slide Time: 00:16)



First, let us look at what additional changes or network transitions that happened with the widespread adoption of the internet. Foremost is the rise of the middleboxes. The phrase middlebox was coined by Lixia Zhang, a computer science professor from UCLA. It is termed as the graphic description of a phenomenon that stormed the internet in the early 90s. By definition, a middlebox is an intermediary device that is placed on the datagram path between the source host and the destination host.

And by functionality, they performed functions other than the normal or the standard function of an IP router, and these are also known as the network appliances when they are purpose-built hardware or by the term network functions or, more recently, the virtualized network functions and traditionally these middleboxes were deployed in communication service provider networks for various purposes.

For example, in the early 90s, when the internet community sensed the end of IPv4 address space and needed newer alternatives to keep up or meet the scale and growth of the internet, many

interim solutions were solved, and as an interim solution, the networking community came up with a gateway based address translation services encompassing public and private addresses. And this resulted in what we started to see as the rise of the network address translators for the IPv4 case, and as things progressed and IPv6 was proposed as a full fledged solution, still the usage of the network address and protocol translators or the NAT devices started to emerge.

Now when you look at these devices, these devices are transparent to the host on either end, that is, both the source host and the destination host. But they would translate the IP address of an incoming packet from a private to a public address and route it out. So, these kinds of functionalities now that were built into the networks are what we refer as the middleboxes, and typically, the NAT or the NAT devices were built they were part of the gateway devices, but they could also be custom proprietary built devices as well.

So, in that sense, these middleboxes are employed to provide basic in-network services like NAT, and we can even consider the load balancers, the proxies, and the transcoders as specific in-network services that these middleboxes provide and second with the need to enhance the security several of the other kinds of middleboxes started to emerge like the firewalls that are deployed at the periphery of a typical enterprise network or the campus networks. And the DDoS protection services or the IDS or the intrusion detection systems or intrusion prevention systems that are built for protecting the network from malicious packets, and these kinds of security-oriented services also started to emerge, and we start to see the second phase of the middleboxes that were deployed.

And more so, there is also another set of middleboxes, like a performance enrichment or the value-added service category of these middleboxes. For example, the WAN accelerators enabled you to perform the TCP acceleration or the HTTP accelerations that you would need when you want to communicate across WAN and also like ad insertions that are typically used to introduce the ads that you typically see when you grow specific sites. So, all of these kinds of different devices started to emerge and which led to the rise of the middleboxes.

And these middleboxes, be it the campus networks or the enterprise network or even in the internet service provider or the ISP networks and telecommunication networks and likewise in the data centers, all began to flood with the use of these various kinds of middleboxes and a

survey was conducted by a team at UC Berkeley around 2012 on the deployment of middleboxes in the North American network operators group that constituted of around 57 network operators.

And the graph that I am showing here presents exactly the results of what they observed and to quickly summarize the plot here what they tried to do is compare different network operators and categorize them in terms of the number of hosts that they had into large, small and medium-sized network sectors and based on the number of hosts that they host the red here corresponds to a large scale while the purple here corresponds to a small scale, and the two green and the blue here are the intermediaries of the mid-scale network operator centers and what they saw in each of these cases was in terms of the number of middleboxes that they employed versus the number of the layer 2 switches and the layer 3 routers that those networks had and invariably we can see that right here the portion depicting the number of middleboxes that were present in each of the network operators group in each case you can see they correspond almost alike to the number of L3 routers and the L2 switches that you see on each of these networks. So, this indicates the widespread usage of the middleboxes that had already spanned several network operators

And if we take just one of the network operator's cases of medium-sized network operators case, what they saw is the type of appliances or these middleboxes that were used for various purposes. Those included the firewalls, the network intrusion detections, the media gateways, the load balancers, the proxies, and so on, and the numbers you could see is almost on the orders of the total number of routers or the typical network elements that we expect to operate at layer 3 or the layer 3 routers the numbers were almost the same, and this rise of the middleboxes was also studied in other ISP and data center networks, and the numbers were found to be in similar orders. So, the survey, in a way, highlights the dominance and importance of middleboxes in different enterprise networks and how they started to shape the internet.

(Refer Slide Time: 07:05)

## RISE OF THE MIDDLEBOXES: IMPACT ON THE ARCHITECTURE/DESIGN PRINCIPLES



- The End2End Principle:
  - Network pipe is no more stateless, middleboxes maintain a load of state in the network.
  - Loss of end-2-end address transparency – A major deviation from the CATENET concept.
- The Hourglass model:
  - Middleboxes dilute the significance of IP layer as the single necessary feature of all communications sessions.
- Fate-sharing principle:
  - Failure in the middleboxes may disrupt the end to end communication, without the failure of either of end hosts.
- Layering principle: No more an Invariant (Horizontal communication)
  - End hosts may receive completely different data at each layer than the one intended by the sender at the origin.

So, in a way, the rise of the middleboxes and their emergence had a significant impact on the internet architecture and the design principles that we discussed earlier. Unlike the traditional end-to-end principle now with the middleboxes, instead of concentrating the diversity and functionalities at just the end host, these middleboxes tried to take and spread these functions within the network, and this was seen throughout the network as well.

And besides the routing and forwarding tables that used to be maintained as the state necessary at the layer 2 switches and layer 3 routers. Now these middleboxes also maintain a large number of state for different purposes. For example, if we consider the NAT device again, a NAT needs to maintain the address and port translation for each of the outgoing and incoming connections. So, this, in a sense is a new state which is now becoming part of the network, and likewise, the load balancers, the reverse proxies also maintain the state mapping, the incoming connection to the mapped back-end server, and so on. Hence, what we argued earlier was the networks were just dumb pipes are no more true, and the networks are no more stateless. We cannot consider the network anymore to be as dumb and could impact the way the networks operate.

Further, now with a NAT again, we can see that the address transparency is completely broken, and as we discussed in the early internet principles that laid foundations by Vinton Cerf and Robert Kahn's work or what is also termed as the CATENET concept, which is a concatenation of the networks to build the internet which was based on a clear assumption that a single logical address space would cover the whole internet. But now, with the NAT, this is no more true, and

likewise with the emergence of private network addresses, which basically is what the concept that the NATs use and hence what this meant is that packets essentially cannot anymore flow unaltered throughout the network that is given a source and destination address these are no more to be considered as unique labels to route in the internet. And what you would need is basically the translations of these addresses to say how to route, and once the routing has a certain endpoint which is not exactly the endpoint as the end hosts, but thereof, you would change the addressing to route it in a private network. So, we can clearly see that the end-to-end principle was almost violated with the rise of these middleboxes.

Further, in the traditional internet architecture, we discussed earlier about the hourglass model where the only box in the neck of the hourglass was the IP protocol, and hence the IP routers belonged to this neck and their only function was to determine the routes and forward the packets. While also trying to update some necessary fields for the purpose of forwarding, but now this is no more true with the middleboxes because now the middleboxes can alter the entire packets data structure and also occupied the same neck space as the IP routers.

And this is where we start to observe the deviations in the hourglass model as these middleboxes would perform functions other than just the IP forwarding. For example, the layer 3 load balancer or the layer 3 NAT that translates the addresses, and further, they are not just confined to layer 3 anymore; although they reside in the network, they may operate at even higher layers like layer 4, where they may even terminate a TCP connection or even at layer 7 terminating the HTTP connection for the proxies typically do and modify even the HTTP options and security parameters.

So, in one way, when we look at the middleboxes, they are a challenge to the hourglass model as the middleboxes dilute the significance of IP as a single necessary feature of all communication sessions. Further, if we look at it in another way, these middleboxes are a challenge to the transparency of the network layer and may break the network functionality independent of the end-to-end systems.

Hence, when we consider the fate-sharing principle earlier, we said that failure or success is collective information that is because of the ends, but now the systems may fail to operate not because of the ends but because of the failure in the network for one of these middlebox devices.

So, now it is no more okay to lose if any unrelated entity goes down while the end connections are still up.


So, having the stateful middleboxes, we end them violating this fate-sharing principle. Again, for example, if we consider the NAT, if the NAT fails or the NAT device gets broken, then the connection between the two end hosts also gets broken because there is no other translator that can facilitate the routing within each of the private networks or provide the network as a proxy to the other side of the public network.

And likewise, it would be with load balancers and proxies. Hence, the fate-sharing principle also is a deviation when it comes with the middleboxes, and moreover, as these middleboxes may transparently modify the state in the network packets, the end host may end up receiving completely different data at each layer than the one that was intended by the center at the origin.

What this means in another way is that we have violated the invariant principle that was the crux of the layering principle. So, overall now, this horizontal communication can no more guarantee this invariance and thus, we start to see the deviations and the violations of a majority of the internet architecture and design principles that we discussed before because of the widespread proliferation of these middleboxes.

(Refer Slide Time: 13:36)

EVOLUTION OF INTERNET - THE COMPROMISING CHANGES



- **Success of TCP and Socket APIs**
  - Success of TCP/UDP – Monopoly of TCP and the slow transport evolution!
  - CAIDA traces in the early 2000, showed that 95% of WAN bytes and 85% of packets were TCP alone.
  - Socket API's directly bind the application to the transport protocol service type!
- **Success of the Web**
  - The protocol stack started to get hardened! HTTP, TLS, TCP and IP.
  - Internet increasingly shaped by commercial interests and less influenced by research.
  - Unwillingness of the Internet Infrastructure players to develop/deploy disruptive technologies.
- **Scale of the network and proprietary hardware and protocols**
  - *monolithic* router contains switching hardware, runs proprietary implementation of Internet standard protocols (IP, RIP, IS-IS, OSPF, BGP) in proprietary router OS (e.g., Cisco IOS)
  - *different "middleboxes"* for different network layer functions: firewalls, load balancers, NAT, etc.
  - Very hard to debug or make changes to the network layer functionalities.

Besides, we also need to consider as the networks evolved, what other comprising changes occurred and we can see that there were several compromising changes that made us drift from the initial principles of the networking and, for example, consider the success of TCP and the socket APIs. So, at one end, the success of the web, the TCP/IP greatly helped scale to a large audience bringing diverse networking services.

But on the other end, the services started to ossify around this protocol stack. So, what I meant by this is when you see the widespread use of a specific protocol amongst the set of protocols, most of the services try to adopt towards the same protocol, and this in a way, stifled the adoption of newer protocols and rather hardened the network stack often resulting in just the adoption of the TCP over IP as well preferred protocol stack. The successful and widely adopted technologies started to pave assumptions about the underlying traffic; for example, now with the middleboxes on one end and the middleboxes need to support the different protocols and started to bring different aspects of how the monopoly over specific protocols could occur at the network level.

And there was a study that was done in the early 90s and that showed that was the CAIDA traces that showed that around 95% of the WAN bytes, around 85% of the packets were just TCP alone, in the sense that the various other protocols that corresponded to the transport layer were more or less getting neglected or were not able to keep up with the traffic that we see because of the services hardening around the TCP IP stack.

Also, the success of the socket APIs that provide the mapping of an obligation to the transport protocols, you could see that in most cases, the type of a socket that you would create would invariably map to TCP as its default transport protocol and thus we started to see this form of a hardening and also called the ossification in this case. Moreover, this progression again also broke the end-to-end argument, like I said because the network is no more dumb, but an intelligent stateful pipe.

And all these aspects, along with the success of the web, also started to shape what looks like the most preferred network stack, and that looks like it would be HTTP at the application layer, TLS for the security and TCP and IP, and most of the application that built around would only part of

the application stack would change, but most of the services the other TLS, TCP and IP would remain intact.

And this was again another form of ossification that happened within the networks. So, internet now could no more be flexed, modified, or even experimented with and especially if the data has to go through the network, being no more just a stateless pipe but lot of a dispersed middleboxes, whatever the protocols that middleboxes would support are the ones that would make it to the end host.

And if there were protocols and the middleboxes in the path of the internet were not supporting them, then it is likely that the packets would be dropped at the middleboxes, and your protocol would never make it to the other end. So, all of these aspects tied with the commercial interest started to affect the way internet research would progress and typically, you would have lot of newer protocols that you would want to experiment with, but the internet became resilient to all of these protocols that you would want to develop and deploy and try.

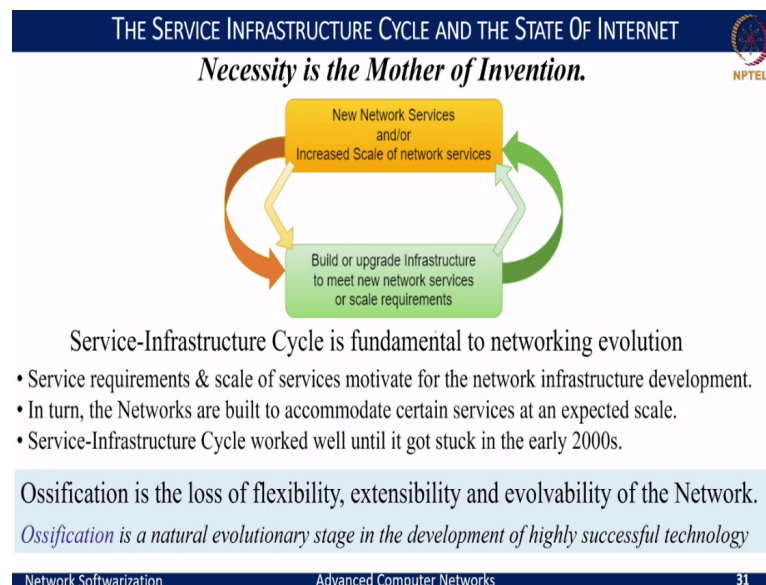
And it was also the unwillingness of the internet infrastructure players to accept or bring in the changes that would allow or adopt to facilitate supporting these networking technologies or the newer disruptive technologies and on the other hand, when the scale of the network and the proprietary hardware and protocol started to build on it made it even worse to intervene and do any disruptive or progressive technological advancements and try it out on the internet testbed and especially even we consider these monolithic routers which were basically developed and built by proprietary implementations from various manufacturers. They built around the standardized protocols that they would support, and any other protocol if you would want to bring and experiment, was not open to experiment or even open to deploy and try.

And this led to another form of ossification because of the market influence and also, like we said, because of the different middleboxes, which were proprietary-built devices again a new scale of the manufacturers who try to control what kind of protocols would work and what would not. So, ossification to think of those who will not familiar with the term is a kind of hardening like in our bones when we are young. They are somewhat agile, but as we age out, they become hard, and the only way that they can be bent is break. They would no more bent and take, for example, the same with the wood as the wood ages you will call that petrified wood, and this

when you try to break it would be like also fragile in the sense that it would completely break and you hit it, but not have dents or bent around when you try to use it. So, with a state of this is exactly what happened with the internet.

In a sense, the protocols that they used to support, the functionalities that you would want to apply and exercise were all getting confined, and more and more they were getting aged and robust over, and it was not easy to tweak any of the aspects over the internet, and this is what we precisely termed as an internet impasse because we would no more allow the researchers to try and experiment something new in the internet and also not allow any new aspects to be brought in a easy fashion even though there is a need to bring. To put an example, when IPv6 started in the late 90s as of today, we have nearly passed two and half decades; still, we have reached around less than around 60% perforation of IPv6, and that is again due to a kind of ossification that we see around the IPv4 address space and unwillingness of many of the players to change to IPv6.

(Refer Slide Time: 20:55)



And to look at it in another perspective, whenever there is a technology that grows and whenever there is a technology that is taken in terms of the services that it provides, we see there is a nice service infrastructure cycle which is fundamental for the evolution of any of the technology, and

it was the same with the networking. On one end, we would have the kind of network services that would grow up, and these kinds of network services as they grow and or a given network service, the demand for kind of network service increases, you would need to build the infrastructure to cater to those demands of newer network services and of the increased scale of network services that you would want to support, and that is exactly what the green box shows where you are trying to build or upgrade your infrastructure to beat the service requirements and the scale requirements of the services that are provided over the internet.

And this chain worked well up till the early 90s or the late 90s, and then we started to see that the networks that were built to accommodate certain services, the certain expected scale started to see a huge upsurge in the scale, and the services were no more able to be accommodated and second it also what the points that we discussed earlier because the infrastructure grew to such a wide scale that making the changes became almost difficult. And it was not easy to make a change and adopt it at the scale of the network, and thus, this cycle got stuck in the early 2000s making many of the network researchers and network community think of the alternatives on how this impasse can be broken and this in a sense led to what we call as the ossification or the loss of flexibility, extensibility and evolvability of the network.

And it was argued by many that this ossification is a natural evolutionary stage in the development of highly successful technology, and the internet just became a victim to it in the early 2000s. These new amplifications or significant increase in the scale of network usage often stresses the network infrastructure in a way that we are required to rethink of the current mechanisms or significant step increase in the network resources that we would need to accommodate if we have to break over this internet impasse.

(Refer Slide Time: 23:34)



- The Internet evolved as an experimental packet-switched network.
  - But, grew beyond leaps and bounds! Diverse Protocols and Services.
- As networks grow they become resistant to change. → i.e. *Become Ossified!*
  - many aspects appear to be "set in stone".
  - large investment in physical infrastructure
  - many stake-holders with often competing interests
  - even modest changes (e.g. IPv6) are difficult to deploy.
- Ossification is a serious concern. → *also creates an impasse for innovation!*
  - the internet, while useful, is far from perfect.
  - if no prospect for innovations to effect networking practice, innovators will go elsewhere
- Problem not unique to networking, but more severe.
  - dominant players in other domains can be slow to change
  - but, competitive environment allows new entrants to innovate and forces incumbents to respond, preventing ossification
- Networking much less open to effective competition.

So, to summarize, the internet evolved as an experimental packet-switched network but grew beyond leaps and bounds with diverse protocols and services, and a variety of middleboxes started to take shape and the network, as they grew, they also become resistant to the changes in terms of the protocols that you would want to put or adapt or change any aspects in the networking stack and thus the network became ossified.

And in the sense many aspects appeared to be now set in stone, and typically you would hinge around the same set of network protocols, and it became very difficult to make any changes, and more so, this ossification is a serious concern, especially when we have a network that requires a lot of innovations and especially when we have the network that is far from perfect.

So, whenever we have to devise and design solutions to fix those imperfections, we also need to bring those into effect, experiment out and see how these can be brought out, but unfortunately, this ossification created an impasse for innovation, and interestingly, this problem is not unique to networking, but it is just that it is more severe when it comes to the networking because the experimentation realizes on having an open internet aspect where you would deploy and test. But unfortunately, there is no means now that you can make these changes and experiment out on an internet scale, and also the market plays its own role in terms of the dominant player who dominates and lay the terms of what protocols need to be there and what cannot be supported, and it becomes very hard to change and also makes the progression very slow, and these competitive environments often allow new entries in many ways, and that would be the barrier

breakers, but when the scale of internet it becomes really difficult to even respond to those newer competitive emergent aspects. Hence, the network, unlike the other domains, was less open to effective competition, and this made the problem even worse.

(Refer Slide Time: 25:50)

TO SUM UP



## Network needs the Change!

*-- new means to break the Impasse & De-ossify the networks*

So, overall what we would see is that there is a need to rethink about the networks, especially in terms of trying to have the researchers try out different aspects over these networks in a way that they could experiment, and try out different innovations much more easily. So, this demand to break the impasse and de-ossify the networks.