

Cloud Computing

Prepared by Dr. Sherif Mostafa

IT381

Lecture 7



Virtual Network for Virtualization

Cloud Computing (Virtualization – Services – Models)

Virtualization layers

Link virtualization

Virtual local area network

Virtual private network

Data center network virtualization

Benefits of network virtualization

- Virtualization technologies create a virtual version of a server, a storage device, an operating system, or other network and computing resources.
- Virtualization technologies make it possible to run multiple operating systems and multiple applications on the same hardware at the same time, improving the utilization and flexibility of hardware.
- Devices, applications, and human users are able to interact with the virtual resource as if it were a real single logical resource.

- In computing, virtualization is a broad concept that refers to the abstraction of computer resources.
- In networking, virtualization is a method of combining the available resources in a network by splitting up the available bandwidth into channels, each of which is independent from the others and each of which can be assigned (or reassigned) to a particular server or device in real time

Cloud computing - Services

Cloud computing is a method of delivering computing services from a large, highly virtualized data center to many independent end-users, using shared applications and pooled resources.

The National Institute of Standards and Technology (NIST) defined three models for provisioning cloud computing services:

- platform as a service (PaaS)
- Software as a service (SaaS)
- Infrastructure as a service (IaaS)

Cloud computing - Services

Cloud computing services are provisioned according to a resource usage-based pricing model. In each service model, users have different degrees of control of the infrastructure management.

PaaS gives control over the deployed applications and application hosting configurations.

SaaS control is typically limited to user-specific application configuration settings.

IaaS provides control over operating systems, storage deployed applications, and possibly limited control on selecting network components

Cloud computing - Models

Cloud computing resources and services can be deployed in different ways, depending on the organizational structure and the provisioning location.

A cloud deployment model defines **where** the physical servers are deployed and **who** manages them.

The NIST defined four main cloud computing deployment models:

- Public cloud
- Private cloud
- Community cloud
- Hybrid cloud.

Cloud computing - Virtualization

Cloud computing would not be possible without virtualization, not for arcane technical reasons but for one obvious business requirement the need for multi-tenancy.

In order to benefit from economies of scale, cloud computing is predicated upon the sharing of a common infrastructure by multiple groups of users, often referred to as tenants.

Multi-tenancy can only be achieved through some kind of virtualization, at the database level, application server level, kernel level, or CPU level.

Some cloud computing applications require multiple distributed sites to work together. Cloud computing requires a network architecture, which ensures flexibility, built-in cross-domain network with higher bandwidth, lower latency, and higher fault tolerance. There are three principal areas in which the network architecture is of utmost importance to cloud computing:

- Data center network (DCN) that interconnects the infrastructure resources servers like and storage devices within a data center.
- Data center interconnect network (DCIN) that connects multiple data centers in a private, public, or hybrid cloud.
- Internet network that connects end-users to the different public data centers.


Cloud computing - Virtualization


Virtualization technologies provide a virtualized view of resources used to instantiate virtual machines.

A hypervisor manages and multiplexes access to the physical resources, maintaining isolation between VMs.

The pluralist philosophy of network virtualization stipulates that networks are fully virtualized, and services are provisioned in an independent manner from the underlying communication infrastructure.

This has resulted in the two-layer virtualization model, shown in the following Figure (a), where infrastructure providers are responsible for allocating network resources, including physical network nodes' CPU and memory, and physical network links to accommodate services provisioned by service providers.






Virtual networks Vs Physical networks

This results in creating virtual networks operating on top of the physical networks:

A virtual network is composed of a set of virtual nodes, each hosted on a physical network node, and a set of virtual links, each established over a physical network path or a set of physical links. Obviously, a virtual network can be embedded in a single or multiple physical networks.

This approach allows infrastructure providers to manage their physical infrastructure separately while service providers focus on enhancing e2e services provisioned to end-users

- Four-layer virtualization business model, as shown in Figure (b), to describe, request, establish, and manage virtualized networks.
- In this model, the virtualization process is carried out among **infrastructure providers**, **virtual network provider (VNP)**, **virtual network operator (VNO)**, and **service provider (SP)**.
- While service providers SPs focus on providing services to users by requesting resources from VNP (via VNO), VNP is responsible for creating virtual networks based on requests received from SPs, whereas VNO is responsible for managing the successful established virtual networks. VNP and VNO remove the burden of creating and managing tasks from SPs, enabling these to focus only on providing services.



Virtualization layers

Network virtualization technology has evolved from virtual local area networks (VLANs), VPN, active programmable networks, and overlay networks to a pluralist philosophy [10], where many heterogeneous architectures co-exist to offer flexible and scalable cloud computing services.

Based on the virtualization technology and network virtualization concept, traffic load and communication operations could be easily controlled and managed to better support big data applications and parallel processing systems, online gaming, High Definition Television (HDTV), and bandwidth on-demand.

Virtualization layers

different physical nodes, although two virtual nodes from different virtual networks can be mapped to the same physical node. The main resources used by virtual nodes and virtual links are CPU and bandwidth, respectively.

Under the term link virtualization, multiple separate traffic flows are transferred over a shared link (physical or emulated), in such a way that each traffic flow appears to be using a dedicated link referred to as virtual link.

Traffic profiling technique allows to limit the packet sending rate that the deployment of traffic profiling at the end hosts makes bandwidth reservation in switches not necessary if the cross traffic is not exceeding the capacity of a given path. Load balancing mechanisms utilize multiple paths leading to a given destination by spreading the traffic over separate routes and providing an illusion of a single link with increased capacity.

Virtual network device

A virtual network device is any networking device that does not exist in a pure physical form but acts like an analogous physical equivalent. A virtual network device can be broadly classified as follows:

- Device aggregation that enables multiple networking devices to act as a single logical entity;
- Device emulation represented by software emulating a physical device, where emulation can be applied to selected parts or the whole device.
- Software-defined networking (SDN) that separates the control plane from the data plane of a networking device such as a switch or router.

Virtual network device - Device aggregation

VMware offers a virtual network-distributed switch that combines all hypervisors virtual switches into one logical centrally managed unit.

Another example of device aggregation is the Juniper virtual chassis technology, which allows up to 10 switches to be interconnected and managed as a single virtual switch.

Similar approaches are usually applied to provide a single point of management over the aggregated devices, although the authors of observed difficulties in determining the source of failures and errors while using virtual chassis

Virtual network device - Device emulation

Multiple VMs running on a single physical machine are normally communicating using software switch (or bridge) provided by a hypervisor. In this case, the functionality of networking device is emulated by the virtualization software.

Virtual network device – Software defined networking

The concept of SDN is based on the idea of separating the control plane from the data plane of network switches and routers.

So, the control and data planes are decoupled, network intelligence and state are logically centralized, and the underlying network infrastructure is abstracted from the applications.

The control plane is implemented in software that runs on separate servers while the data plane might be realized in software or networking hardware.

Virtual network device – Software defined networking

virtual switch for hypervisors providing network connectivity to VMs and supports migration of VMs by migrating the network state associated with a migrated VM.

These virtual switches maintain a flow table that defines what to do with each flow. It classifies traffic flows and assigns them to a dedicated queue for processing through employing scheduling algorithms along with the traffic profiling.

Network topology virtualization

Network topology virtualization supports multiple virtual networks to co-exist on a same physical network substrate. Such a dynamic virtualization environment can enable the deployment of multiple co-existing heterogeneous network architectures without the inherent limitations found in the existing Internet.

Virtualization Approaches


There are different virtualization approaches that can support network topology virtualization, such as the following:

- Virtual LANs
- Virtual Private Networks

Virtual local area network

A VLAN is a group of hosts with a common interest that are logically brought together under a single broadcast domain regardless of their physical connectivity. Because VLANs are logical entities, that is, configured in software, they are flexible in terms of network administration, management, and reconfiguration.

Virtual local area network




Virtual local area network

All frames in a VLAN bear a common VLAN ID in their Medium Access Controller (MAC) headers, and VLAN-enabled switches use both the destination MAC address and the VLAN ID to forward frames.

This process is known as frame coloring. Multiple VLANs on multiple switches can be connected together using trunking, which allows information from multiple VLANs to be carried over a single link between switches.


Virtual local area network



Virtual local area network




1- separate broadcast domains



2- Enhance network performance



3- Scalable



4-security



Virtual private network

A VPN is a dedicated communications network of one or more enterprises that are distributed over multiple sites and connected through tunnels over public communication networks like the Internet.

Each VPN site contains one or more customer edge (CE) devices like hosts or routers, which are attached to one or more provider edge (PE) routers.


A VPN is managed and provisioned by a VPN service provider, known as provider-provisioned VPN.

Virtual private network


While VPN implementations exist in several layers of the network stack are layer 3 and 2 VPN.

The following figure provides an example of VPN application in extending enterprises data centers to a private cloud.

Virtual private network



A private cloud as an extension of the enterprise 'C' data center using VPN



Layer 3 VPN (L3VPN) is distinguished by its use of layer 3 protocols like IP or multi-protocol label switching (MPLS) in the VPN backbone to carry data between the distributed customer edge (CE) devices like hosts or routers.

L3VPN can be classified into two categories: customer edge CE-based and provider edge PE-based VPNs. In the CE-based VPN approach, CE devices create, manage, and tear up the tunnels without the knowledge of the service provider SP network

Tunneling requires three different protocols:

1. Carrier protocol as the IP, which is used by the SP network to carry the VPN packets.
2. Encapsulating protocol, which is used to wrap the original data. It can range from very simple wrapper protocols, point-to-point tunneling protocol, and layer 2 tunneling protocol to secure protocols.
3. Passenger protocol, which is the original data the IP data encapsulated in the user networks

Sender customer edge CE devices encapsulate the passenger packets and route them into carrier networks.

When the encapsulated packets reach the receiver customer edge CE devices at the end of the tunnels, they are extracted, and actual packets are injected into receiver networks. In provider edge PE-based L3VPNs, the service provider SP knows that certain traffic is a VPN traffic and process them accordingly.

The VPN states are stored in PE devices, and a connected CE device behaves as if it was connected to a private network.

Layer 2 VPN (L2VPN) provides an edge to edge (e2e) layer 2 connection between distributed sites by transporting layer 2 (Ethernet and frame relay frames) between participating sites.

L2VPN support heterogeneous higher-level protocols, but its lack of a control plane takes away its capability of managing reachability across the VPN.

There are two different types of layer 2 VPN services SP could offer to a user:

- point-to-point virtual private wire service (VPWS)
- point-to-multi-point virtual private LAN service (VPLS)


Data center network virtualization

A DCN connects hundreds of thousands of servers to deliver scalable cloud services to cloud end-users. Most public cloud service providers have adopted a two-tier data center architecture

- front-end tier and
- back-end tier

to optimize hardware, operations, and service delivery cost.

Where the front end of a networked system is the side that a user, client, or customer sees while the back end is everything that happens behind the scenes that is not necessarily seen by the user.



A typical Data center network with server virtualization

A DCIN is designed with an objective to support the requirements of seamless user experience in using cloud services, which can guarantee service availability, dynamic service migration, and applications mobility.

So, What is network virtualization?

Network virtualization is the process of decoupling network services and functions from the underlying hardware. These network functions include internet protocol (IP) routing, packet switching and filtering, load balancing, and virtual private networks, and are combined into a single pool of resources in virtual networks.

What is network virtualization?

Where traditional IT infrastructure served up networking resources on physical servers, routers, and switches, virtualized networks create a logical overlay that allows you to treat most network functions separately from the hardware. This means you can create, deploy, provision, and manage virtual networks in software, only relying on the physical infrastructure to handle IP packet forwarding.

What is network virtualization?

With network virtualization, you can consolidate a group of multiple physical networks into a single virtual network, separate a single network into multiple virtual segments, or connect virtual machines (VMs) across distinct domains. Virtual networks can be isolated from each other, even though they all share the same underlying physical hardware.

What is network virtualization?

Similar to how server virtualization enables the creation of multiple virtual machines that reside on a single physical server, network virtualization allows you to have multiple virtual networks that run atop a single physical network infrastructure.

What is network virtualization?

Network virtualization represents a significant departure from traditional IT infrastructure because it frees you from physical hardware limitations and the challenges of physically networking VMs. It provides a great degree of flexibility because you can group or separate virtual networks as needed—setting specific networks aside for designated workloads, for instance—or organize and connect VMs however you want.

Because there's no need to configure the physical hardware to run or create virtual networks, scaling up or down can be fast and easy.

How does network virtualization works?

In a virtual network, specific network functions and resources are sent out to the virtual layer, which is managed by a hypervisor that serves as a virtual network adapter. The virtual layer includes all the nodes and links that comprise the virtual networks, and the hypervisor assigns resources to workloads as needed. Each virtual network is isolated from the others and can have policies that are different from the other networks.

How does network virtualization works?

One of the things that organizations like most about network virtualization is that once you've configured the appropriate network and security policies for a virtually networked application, you can move that application to another host (server) and those network and security configurations go with it. Even better, if you need to create new workloads to scale that application, those new workloads automatically get the same policies you've already configured. This helps keep policies consistent and saves IT time and effort.

How does network virtualization works?

Hosts and workloads such as VMs can communicate with other hosts and workloads on separate virtual networks, using specific host protocols and virtual switches and routers. One advantage of this workload-to-workload communication is that it doesn't travel over the physical network, so latency is reduced.

Types of virtual networks

There are two main types of network virtualization:

External network virtualization

This type allows you to virtualize physical networks that are on the same local area network (LAN), which could span multiple physical servers. Using virtual switches and adapters, external virtual networks can also virtualize multiple physical LANs into a single virtual LAN (VLAN). Organizations use this type of network virtualization to make management more efficient—because IT can manage and configure multiple systems connected to the VLAN instead of working on each machine individually.

Internal network virtualization

This type is confined to one network server. It can help improve performance because it simulates a physical network through the use of distinct software containers, relevant code, and all the various dependencies needed by an application to perform consistently across the entire ecosystem. All the VMs that reside on that server can communicate with each other without having to use an external network.

Benefits of network virtualization

- Simplify network management
- Increase network flexibility
- Boost security across networks
- Reduce costs of hardware

The concept of virtual networks is used to integrate diverse user groups with specific needs, and enables benefits such as increased productivity, improved security, network consolidation, and integration of subsidiaries.



Thank You