

IT Infrastructure Architecture


Infrastructure Building Blocks
and Concepts

Networking

Network virtualization

Virtual LAN (VLAN)

- VLANs enable logical grouping of network nodes on the same LAN
 - Configured on network switches
 - Operate at the Ethernet level



Virtual LAN (VLAN)

- VLANs:
 - Allow segmenting a network at the data link layer
 - Allow end stations to be grouped together even if they are not physically connected to the same switch
 - Can adapt to changes in network requirements and allow simplified administration
 - Enhance security by preventing traffic in one VLAN from being seen by hosts in a different VLAN
- For VLANs to communicate with each other a router is needed

VXLAN

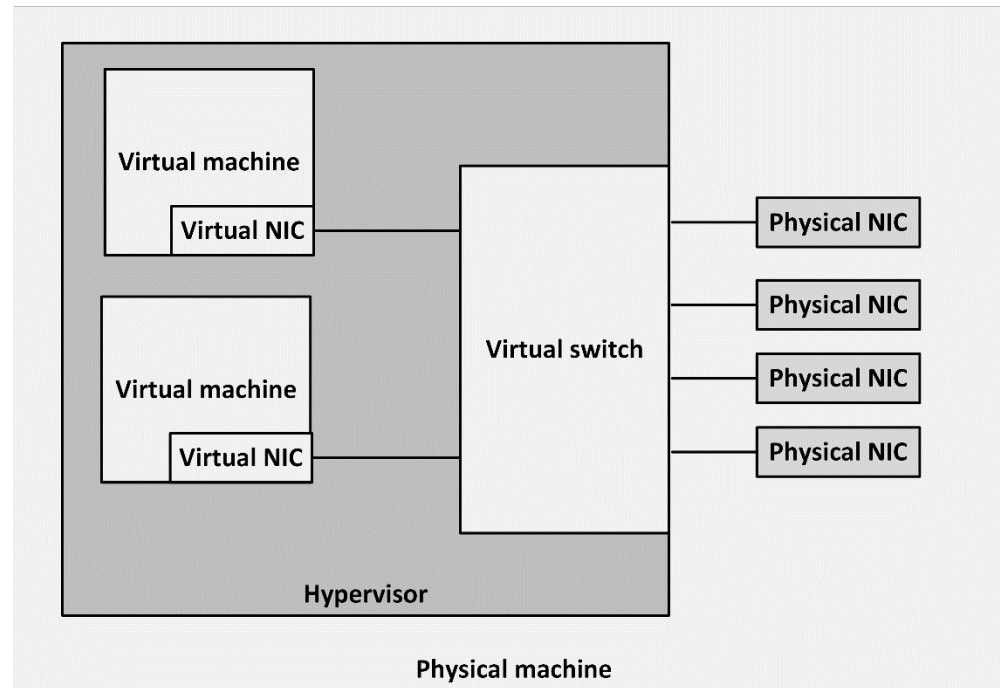
- Virtual Extensible LAN (VXLAN) is an encapsulation protocol
- Can be used to create a logical switched layer 2 network across routed layer 3 networks
- Only servers within the same logical network can communicate with each other
- VXLANs are heavily used in multi-tenant cloud environments

Virtual NICs

- Virtual machines are only aware of virtual Network Interface Controllers (NICs) provided to them
- Virtual machines running on physical machines share physical NICs
- Communications between virtual machines on the same physical machine are routed directly in memory space by the hypervisor, without using the physical NIC
- The hypervisor routes Ethernet packages from the virtual NIC on the virtual machine to the physical NIC on the physical machine

Virtual switch

- Virtual NICs are connected to virtual switches
- A virtual switch is an application running in the hypervisor, with most of the capabilities of a physical network switch
- A virtual switch is dynamically configured
 - Ports in the virtual switch are configured at runtime
 - The number of ports on the switch is in theory unlimited



Virtual switch


- Availability:
 - No cable disconnects
 - No need for auto-detecting network speed
 - No network hubs, routers, adapters, or cables that could physically fail
- Security:
 - No easy way to intercept network communications between virtual machines from outside of the physical machine

Software Defined Networking

- Software Defined Networking (SDN) allows networks to be defined and controlled using software external to the physical networking devices
- A set of physical network switches can be programmed as a virtual network:
 - Hierarchical
 - Complex
 - Secured
- A virtual network can easily be changed without touching the physical network components

Software Defined Networking

- Control plane resides centrally
- Data plane (the physical switches) remain distributed



Software Defined Networking

- SDN can be controlled from a single management console
- Provides open APIs that can be used to manage the network using third party software
- In an SDN, the distributed data plane devices are only forwarding network packets based on ARP or routing rules that are preloaded into the devices by the SDN controller in the control plane
 - This allows the physical devices to be much simpler and more cost effective

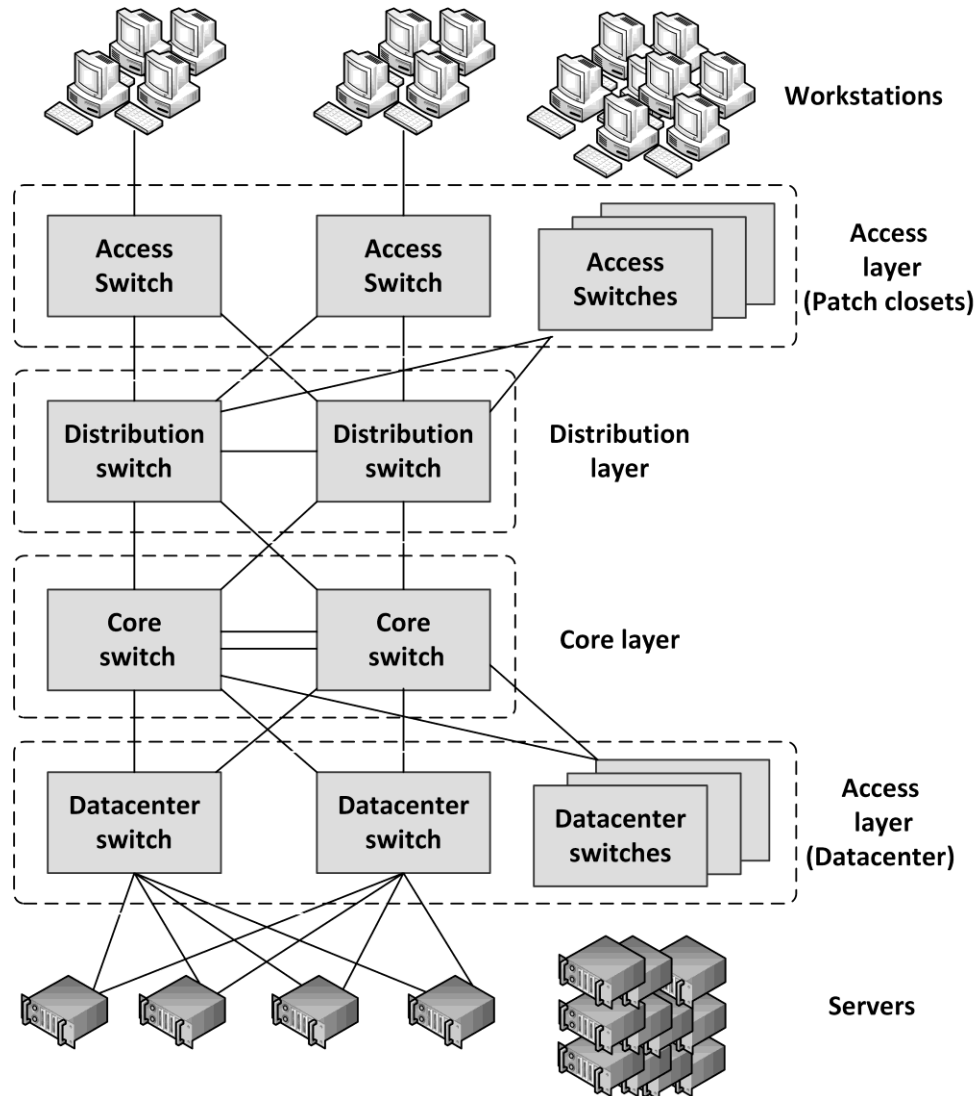
Network Function Virtualization

- Network Function Virtualization (NFV) is a way to virtualize networking devices
 - Firewalls
 - VPN gateways
 - Load balancers
- NFV appliances are implemented as virtual machines running applications that perform the network functions
- NFV virtual appliances can be created and configured dynamically and on-demand using APIs
- Example:
 - Deploy a new firewall as part of a script that creates a number of connected virtual machines in a cloud environment

Network availability

Layered network topology

- A network infrastructure should be built up in layers
 - Improve availability and performance
 - Provides scalability
 - Provides deterministic routing
 - Avoids unmanaged ad-hoc data streams
- Provides high availability
 - Because the layering provides multiple paths to any piece of equipment

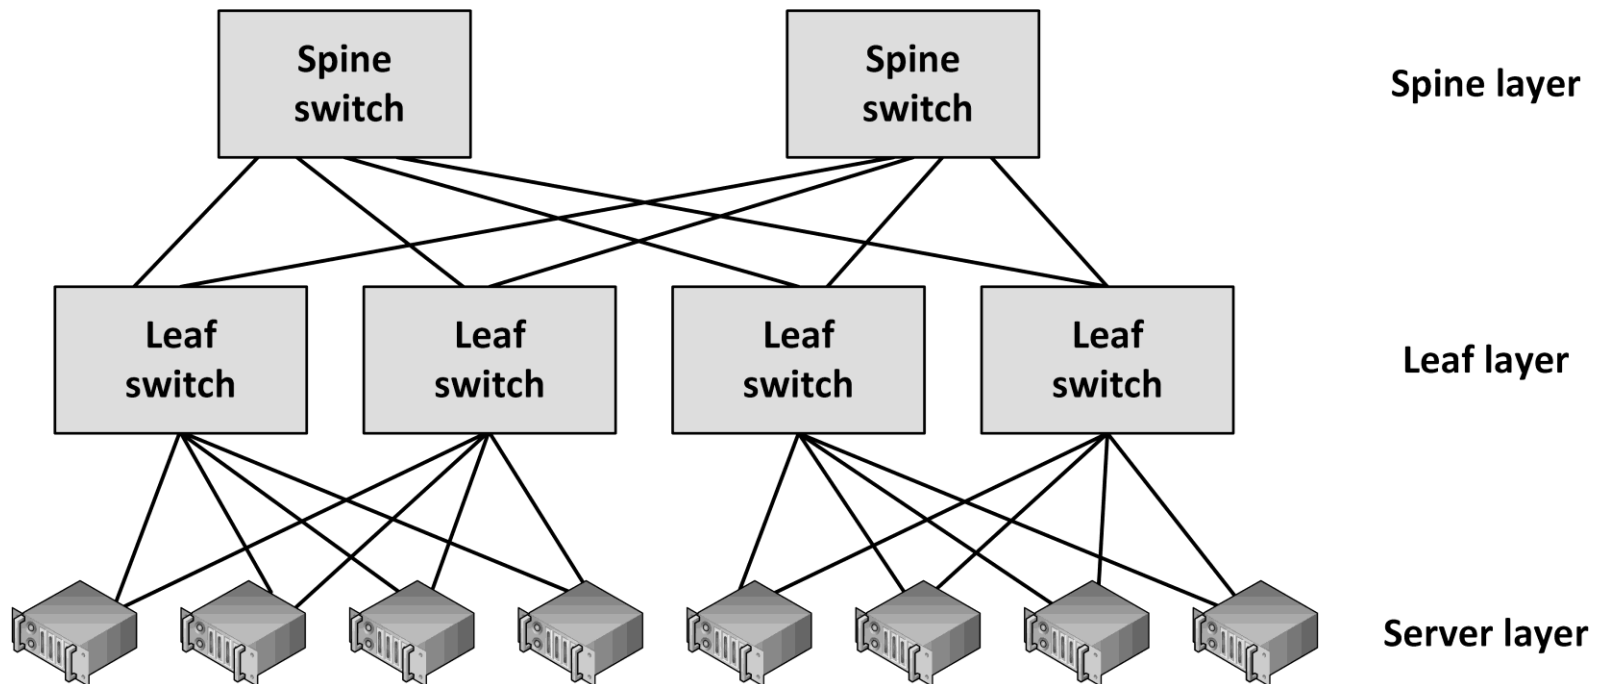


Layered network topology

- Core layer
 - This is the center of the network
- Distribution layer
 - An intermediate layer between the core layer in the datacenter and the access switches in the patch closets
 - Combines the access layer data and sends its combined data to one or two ports on the core switches
- Access layer
 - Connect workstations and servers to the distribution layer
 - For servers, located at the top of the individual server racks or in blade enclosures
 - For workstations, placed in patch closets in various parts of the building

Spine and Leaf topology

- In a SDN, a simple physical network is used that can be programmed to act as a complex virtual network
- Such a network can be organized in a spine and leaf topology



Spine and Leaf topology


- Characteristics:
 - The spine switches are not interconnected
 - Each leaf switch is connected to all spine switches
 - Each server is connected to two leaf switches
 - The connections between spine and leaf switches typically have ten times the bandwidth of the connectivity between the leaf switches and the servers

Spine and Leaf topology

- Benefits:
 - Highly scalable
 - There are no interconnects between the spine switches
 - Simple to scale
 - Just add spine or leaf servers
 - With today's high density switches, many physical servers can be connected using relatively few switches
 - Each server is always exactly four hops away from every other server
 - Leads to a very predictable latency

Network teaming

- Network teaming is also known as:
 - Link aggregation
 - Port trunking
 - Network bonding
- Provides a virtual network connection using multiple physical cables for high availability and increased bandwidth





Network teaming

- Network teaming bonds physical NICs together to form a logical network team
 - Sends traffic to the team's destination to all NICs in the team
 - Allows a single NIC, cable, or switch to be unavailable without interrupting traffic

Spanning Tree Protocol (STP)

- STP is an Ethernet level protocol that runs on switches
- Guarantees that only one path is active between two network endpoints at any given time
- Redundant paths are automatically activated when the active path experiences problems
- Ensures no loops are created when redundant paths are available in the network
- A disadvantage of using the spanning tree protocol is that it is not using half of the network links in a network, since it blocks redundant paths
- Rapid Spanning Tree Protocol (RSTP) provides for fast spanning tree convergence after a topology change (6 s instead of 30-60 s)

Spanning Tree Protocol




Multihoming

- Connecting a network to two different Internet Service Providers (ISPs) is called multihoming
- Four options for multihoming:
 - Single router with **dual links** to a single ISP
 - Single router with dual links to **two separate** ISPs
 - **Dual routers**, each with its own link to a single ISP
 - Dual routers, each with its own link to a **separate ISP**
- It is not always guaranteed that multiple network paths actually run on a different set of cables
 - WAN cables are typically installed alongside highways and railway tracks
 - Cables are used by multiple carrier providers

Network performance

Nielsen's law

- Network connection speeds for high-end home users increase 50% per year, they double every 21 months
- Bandwidths should be 15 Gbit/s in 2025, for about \$50 per month



Please note that the vertical scale is logarithmic instead of linear

Throughput and bandwidth

- Throughput is the amount of data that is transferred through the network during a specific time interval
- Throughput is limited by the available bandwidth
- When an application requires more throughput than a network connection can deliver:
 - Queues in the network components temporarily buffer data
 - Buffered data is sent as soon as the network connection is free again
 - When more data arrives than the queues can store in the buffer, packet loss occurs

Latency

- Latency is defined as the time from the start of packet transmission to the start of packet reception
- Latency is dependent on:
 - The physical distance a packet has to travel
 - The number of switches and routers the packet has to pass
- Rules of thumb:
 - 6 ms latency per 100 km
 - WANs: Each switch in the path adds 10 ms to the one-way delay
 - LANs: add 1 ms for each switch

Latency

- One-way latency: the time from the source sending a packet to the destination receiving it
- Round-trip latency: the one-way latency from source to destination plus the one-way latency from the destination back to the source
- “ping” can be used to measure round-trip latency

Quality of Service (QoS)

- Quality of service (QoS) is the ability to provide different data flow priority to different applications, users, or types of data
- QoS allows better service to certain important data flows compared to less important data flows
- QoS is mainly used for real-time applications like video and audio streams and VoIP telephony

Quality of Service (QoS)

- Four basic ways to implement QoS:
 - Congestion management
 - Defines what must be done if the amount of data to be sent exceeds the bandwidth of the network link
 - Packets can either be dropped or queued
 - Queue management
 - When queues are full, packets will be dropped
 - Queue management defines criteria for dropping packets that are of lower priority before dropping higher priority packets

Quality of Service (QoS)

– Link efficiency

- Ensures the link is used in an optimized way
- For instance by fragmenting large packets with a low QoS, allowing packets with a high QoS to be sent between the fragments of low QoS packets

– Traffic shaping

- Limiting the full bandwidth of streams with a low QoS to benefit streams with a high QoS
- High QoS streams have a reserved amount of bandwidth

WAN link compression

- Data compression reduces the size of data before it is transmitted over a WAN connection
- WAN acceleration appliances:
 - Provide compression
 - Perform some caching of regularly used data at remote sites