Unit 7:
Backbone Networks

Contents

- Introduction; Switched Backbones; Routed Backbones; Virtual LANs (Benefits of VLANs, How VLANs Work); The Best Practice Backbone Design; Improving Backbone Performance (Improving Device Performance, Improving Circuit Capacity, Reducing Network Demand).

Introduction

- Most business organizations realize that information must be stored, retrieved, analyzed, acted on, and shared with others at a moment's notice.
- Without an enterprise-wide network or an Internet connection, moving information from one department LAN to another or to customers is difficult.
- Interconnecting the organization's diverse networks is critical. A **backbone network (BN)** is a high- speed network that connects many networks. BNs typically use higher-speed circuits to interconnect a series of LANs and provide connections to other BNS, MANS, WANs, and the Internet.
- A backbone that connects many BNs spanning several nearby buildings for a single organization is often called a campus network. A BN also may be called an enterprise network if it connects all networks within a company, regardless of whether it crosses state, national, or international boundaries.
- There are two basic components to a BN: the **network cable** and the **hardware devices** that connect other networks to the BN. The cable is essentially the same as that used in LANs, except that it is often fiber optic to provide higher data rates. The hardware devices can be computers or special- purpose devices that just transfer messages from one network to another. These include switches, routers, and VLAN switches.

- ***Switches*** operate at the data link layer. They connect two or more network segments that use the same data link and network protocol. They understand only data link layer protocols and addresses.
- ***Routers*** operate at the network layer. They connect two different TCP/IP subnets. Routers are the "TCP/IP gateways". Routers strip off the data link layer packet, process the network layer packet, and forward only those messages that need to go to other networks on the basis of their network layer address. In general, they perform more processing on each message than switches and therefore operate more slowly.
- ***VLAN switches*** are a special combination of layer 2 switches and routers. They are complex devices intended for use in large networks that have special requirements

Switched Backbones

- Switched backbones are probably the most common type of BN used in the distribution layer. Switched backbone refers to a high-speed, high-capacity network segment that interconnects various network segments, subnets, or LANs within an organization or data center. It serves as a central communication pathway, efficiently managing and directing traffic between different network segments.
- Switched backbones play a crucial role in modern network infrastructures, ensuring that data can flow smoothly and reliably between connected devices. Switched BNs use a star topology with one switch at its center.
- Figure shows a switched backbone connecting a series of LANs. There is a switch serving each LAN (access layer) that is connected to the backbone switch at the bottom of the figure (distribution layer). Most organizations now use switched backbones in which all network devices for one part of the building are physically located in the same room, often in a rack of equipment. This has the advantage of placing all network equipment in one place for easy maintenance and upgrade, but it does require more cable.
- In most cases, the cost of the cable is only a small part of the overall cost to install the network, so the cost is greatly outweighed by the simplicity of maintenance and the flexibility it provides for future upgrades.

With rack-mounted equipment, it becomes simple to move computers from one LAN to another. Usually, all the computers in the same general physical location are connected to the same switch and thus share the capacity of the switch. Although this often works well, it can cause problems if many of the computers on the switch are high-traffic computers. For example, if all the busy computers on the network are located in the upper-left area of the figure, the switch in this area may become a bottleneck.

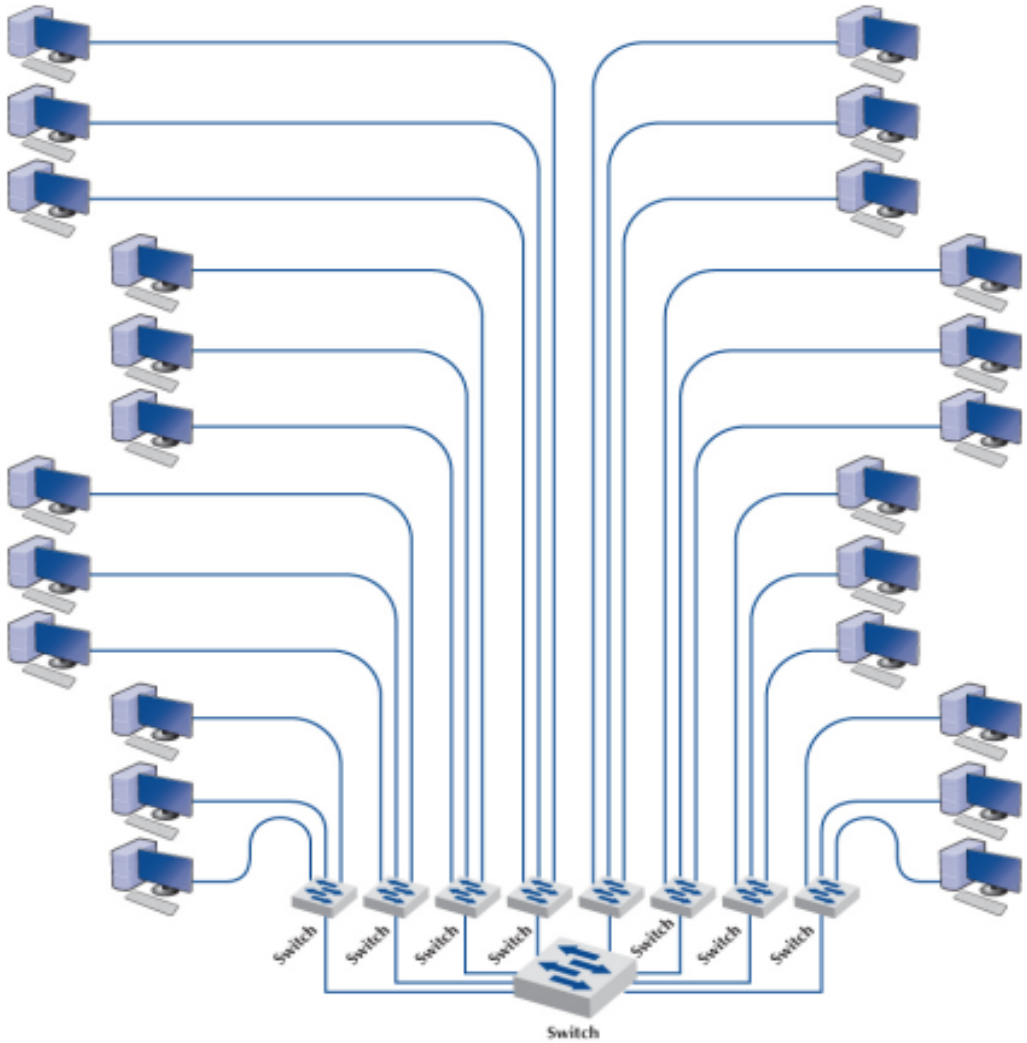

Rack-mounted switched backbone network architecture

- With an Main Distribution Facility (MDF), all cables run into the MDF. If one switch becomes overloaded, it is straightforward to unplug the cables from several high demand computers from the overloaded switch and plug them into one or more less-busy switches.
- This effectively spreads the traffic around the network more efficiently and means that network capacity is no longer tied to the physical location of the computers; computers in the same physical area can be connected into different network segments.

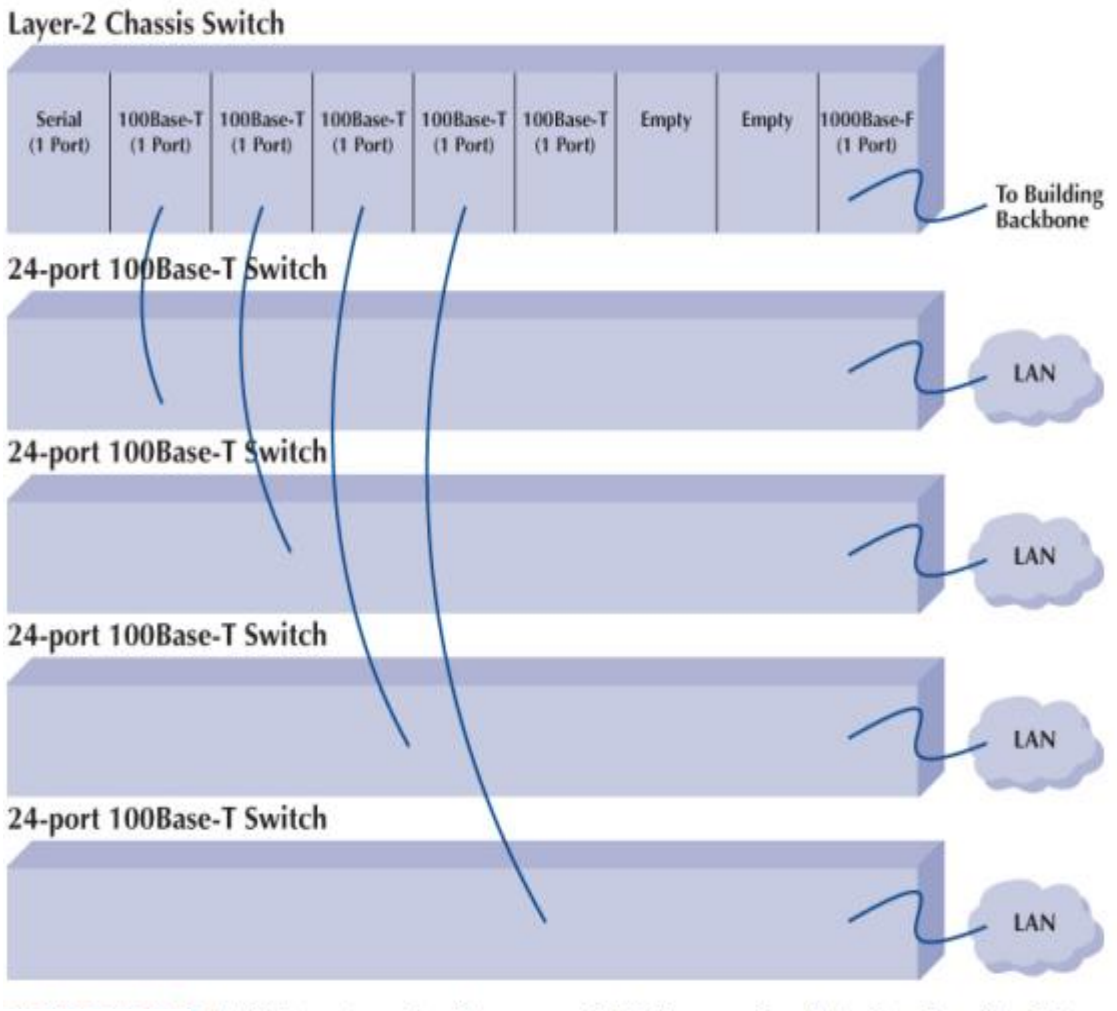

MDF network diagram. MDF = main distribution facility

- Here are some key features and characteristics of switched backbones:
- **Switching Technology:** The term "switched" refers to the use of network switches, which are intelligent devices that operate at Layer 2 (data link layer) or Layer 3 (network layer) of the OSI model. Switches are responsible for forwarding data packets based on MAC addresses (Layer 2) or IP addresses (Layer 3).
- **High-Speed Connectivity:** Switched backbones are designed to provide high-speed connectivity, often using technologies like Gigabit Ethernet or 10 Gigabit Ethernet. This ensures that data can flow quickly and efficiently between network segments.
- **Segmentation:** Switched backbones enable network segmentation, which involves breaking down a large network into smaller, more manageable segments. Each segment can have its own unique characteristics, such as separate IP subnets or VLANs (Virtual LANs).
- **Reduced Broadcast Traffic:** Switches in a switched backbone effectively filter out unnecessary broadcast traffic. Broadcast packets are only sent to devices that need to receive them, reducing network congestion and improving overall performance.

- **Improved Scalability:** A switched backbone can easily accommodate the addition of new devices and network segments. This makes it highly scalable, allowing organizations to expand their network infrastructure as needed.
- **Redundancy:** Redundancy and fault tolerance are often built into switched backbone designs, Redundant switches, links, and failover mechanisms ensure network availability and minimize downtime in case of hardware failures.
- **Quality of Service (QoS):** Switched backbones can implement QoS mechanisms to prioritize certain types of traffic (e.g., voice or video) over others, ensuring a consistent quality of service for critical applications.
- **Security:** Advanced switched backbones can incorporate security features like access control lists, network segmentation protect the network from unauthorized access and threats.
- **Centralized Management:** Network administrators can manage and monitor the entire switched backbone from a central location, allowing for efficient configuration and troubleshooting.
- **Core of Data Centers:** In data center environments, switched backbones are a fundamental component. They connect servers, storage devices, and other critical infrastructure components to ensure rapid and reliable data transfer within the data center.

- Switched backbones are essential for organizations with diverse and extensive networking needs. They facilitate efficient and secure communication between various network segments, ensuring that data can flow seamlessly within the network. As networking technology continues to evolve, switched backbones remain a critical component for handling the ever-increasing volume of data in the digital age.

Routed Backbones

- A routed backbone, is a network infrastructure where data is routed between various subnets or network segments using routers. Unlike switched backbones, which rely on Ethernet switches to forward data at the data link layer, routed backbones operate at the network layer (Layer 3 of the OSI model).
- Routers are the key devices in a routed backbone, and they make forwarding decisions based on IP addresses, allowing data to traverse between different IP subnets or networks. Routed backbones are sometimes called subnetted backbones or hierarchical backbones and are most commonly used to connect different buildings on the same enterprise campus backbone network.

- Figure illustrates a routed backbone used at the core layer. A routed backbone is the basic backbone architecture we used to illustrate how TCP/IP worked. There are a series of LANs (access layer) connected to a switched backbone (distribution layer).
- Each backbone switch is connected to a router. Each router is connected to a core router (core layer). These routers break the network into separate subnets. The LANs in one building are a separate subnet from the LANs in a different building. Message traffic stays within each subnet unless it specifically needs to leave the subnet to travel elsewhere on the network, in which case the network layer address (e.g., TCP/IP) is used to move the packet.
- For example, in a switched backbone, a broadcast message (such as an ARP) would be sent to every single computer in the network. A routed backbone ensures that broadcast messages stay in the one network segment (i.e., subnet) where they belong and are not sent to all computers. This leads to a more efficient network.

- Each set of LANs is usually a separate entity, relatively isolated from the rest of the network. There is no requirement that all LANs share the same technologies. Each set of LANs can contain its own server designed to support the users on that LAN, but users can still easily access servers on other LANs over the backbone, as needed.

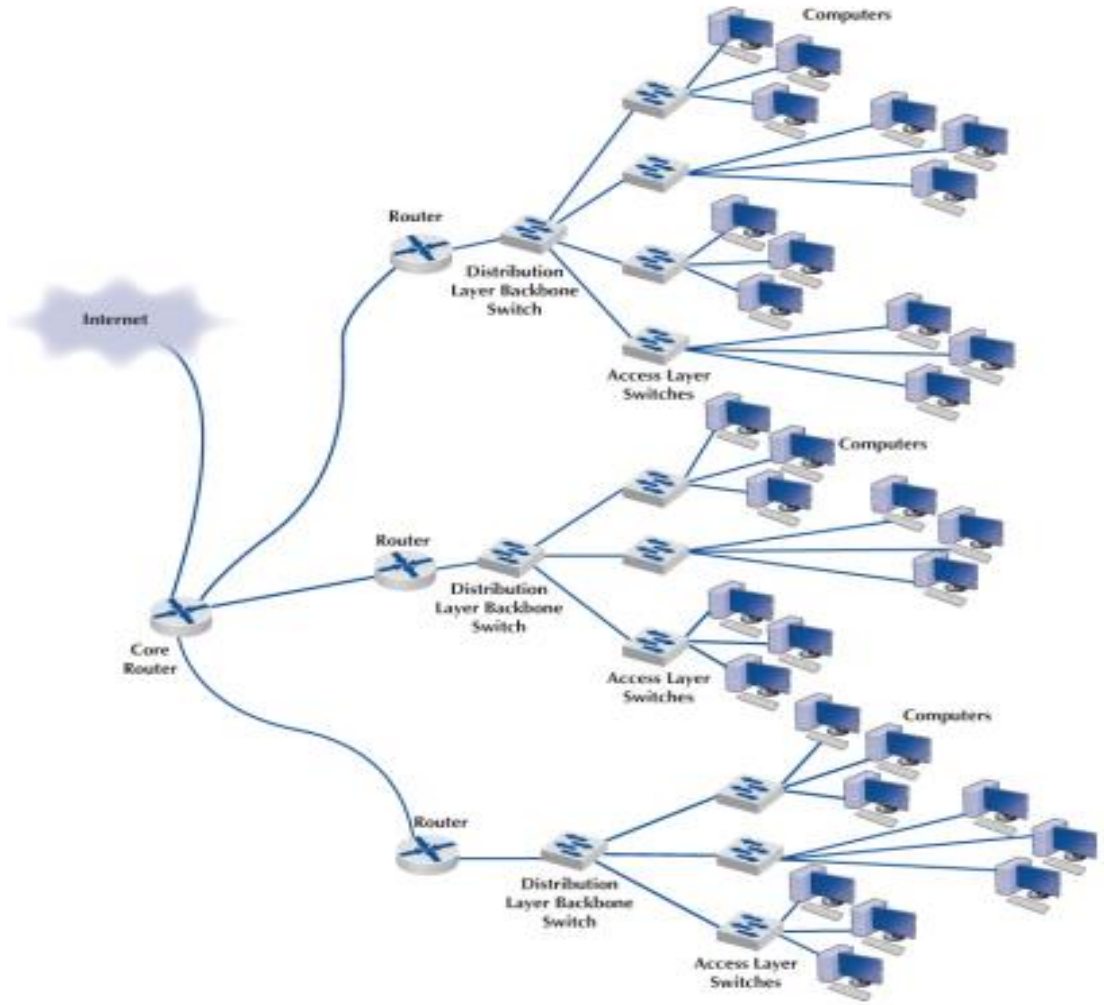

Routed Backbone

- The primary advantage of the routed backbone is that it clearly segments each part of the network connected to the backbone. Each segment (usually a set of LANs or switched backbone) has its own subnet addresses that can be managed by a different network manager. Broadcast messages stay within each subnet and do not move to other parts of the network.
- There are two primary disadvantages to routed backbones. First, the routers in the network impose time delays. Routing takes more time than switching, so routed networks can sometimes be slower. Second, routers are more expensive and require more management than switches.

- Here are some key features and characteristics of routed backbones:
- **IP Routing:** Routed backbones use IP routing to determine the path data should take to reach its destination. Routers maintain routing tables that list the available paths and select the best one based on the destination IP address.
- **Subnet Isolation:** A routed backbone often separates different IP subnets or network segments. This isolation provides network security and reduces the size of broadcast domains, helping to prevent broadcast storms that can occur in flat, switched networks.
- **Interconnectivity:** Routed backbones enable the interconnection of multiple LANs or subnets. This is especially useful in large enterprise networks where different departments or geographical locations need to communicate securely.
- **Traffic Filtering:** Routers in a routed backbone can filter and control the flow of traffic based on various criteria, such as access control lists (ACLs) that specify which types of traffic are allowed or denied.
- **Scalability**: Routed backbones can easily scale as new subnets or network segments are added. This makes them suitable for accommodating network growth and changing requirements.

- **Optimized Routing:** Routers in a routed backbone can implement routing protocols (e.g., OSPF, EIGRP, BGP) to dynamically discover and select the most efficient paths for data to travel within the network.
- **VLANs and Virtual Routing:** Routed backbones can support Virtual LANs (VLANs) to create logical groupings of devices within a routed network. Additionally, routers can provide virtual routing instances to segment the network further.
- **Internet Connectivity:** A routed backbone often serves as the connection point between the internal network and external networks, such as the internet. Routers in the backbone manage the flow of data between internal and external networks.
- **Quality of Service (QoS):** Routed backbones can implement QoS mechanisms to prioritize specific types of traffic for better network performance. This is particularly important for applications that require low latency or high bandwidth.
- **Security and Firewalling:** Routers in a routed backbone can serve as security gateways and , controlling access to the network and protecting it from unauthorized access and cyber threats.

- Routed backbones are commonly used in larger networks, including enterprise networks, where network segmentation, security, and scalability is essential. While they may require more complex configuration and maintenance compared to switched backbones, routed backbones provide the flexibility and control needed to meet the diverse requirements of today's complex networking environments.

1. **Switched Backbone:**
    1. Think of a switched backbone like a super-smart traffic cop. It uses network switches to direct traffic efficiently within the network.
    2. When data (like a file or a webpage) needs to travel from one part of the network to another, the switches help it take the fastest and most direct route to its destination.
    3. Switches make decisions at a lower level based on the MAC (Media Access Control) addresses of devices, allowing for fast and direct communication between devices within the same network.
2. **Routed Backbone:**
    1. Now, imagine a routed backbone as a thoughtful courier navigating through city streets. It uses routers to guide data between different networks.
    2. Routers make decisions based on IP (Internet Protocol) addresses. If data needs to go from one network to another (like from your home network to a server on the internet), routers help it find the best path.
    3. Routers are more aware of the bigger picture, helping data navigate across different networks and ensuring it reaches its destination efficiently.

- In summary:
- **Switched Backbone** focuses on quick and efficient communication within a single network using switches and MAC addresses.
- **Routed Backbone** looks at the broader picture, guiding data between different networks using routers and IP addresses.
- Both switched and routed backbones have their roles in a network, and often, a combination of these technologies is used to create a reliable and efficient overall network infrastructure.

Introduction to VLAN

- Suppose we have two departments in an organization- Sales and Marketing, connected as shown in the figure.
- The sales PC wants to broadcast a message for its department while the message has nothing to do with the marketing department, but what should the switch do in such circumstances?
- Yes, it broadcasts that message to each PC connected to it; hence, the marketing PC will also be reading that particular message.

Does it sound good?

- Apart from unnecessary network congestion occurring due to broadcast, there is a layer two attacks risk.
- One of the solutions is to buy different switches for each department and connect them accordingly.

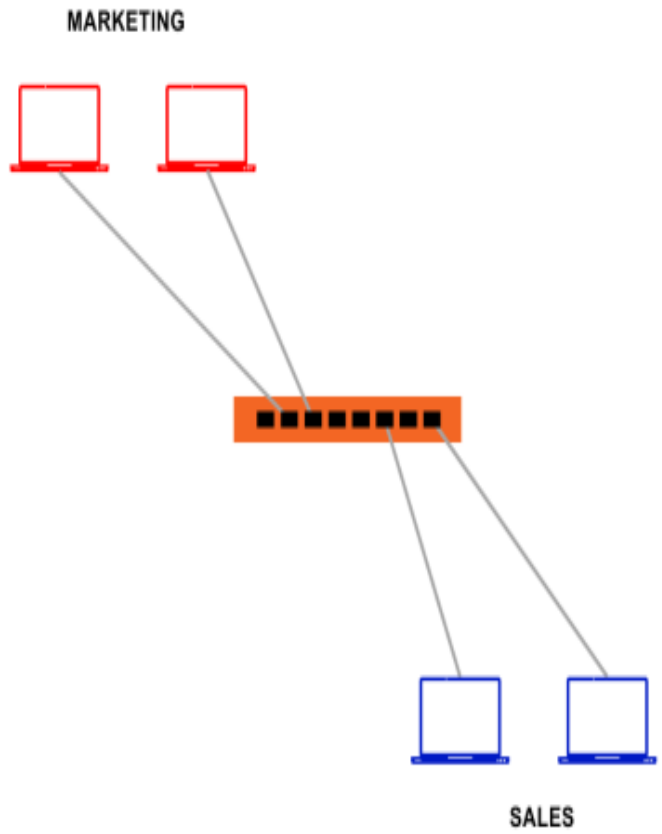

- **But, with this approach, three problems arise-**
- Cost of the infrastructure increases.
- A lot of switch ports might remain vacant.
- What if one department wants to communicate with another department? We need to broadcast that message individually for each department.
- **Thus, VLAN is the concept that can help to eliminate all these problems.**
- VLAN is a logical grouping of network devices connected to a switch. By creating VLAN, we create smaller broadcast domains at layer-2 by assigning different ports to different subnetworks on one switch.
- **In simple words, we are creating a small LAN inside a LAN.**
- With the help of VLAN, frames broadcasted get switched between ports and groups within the same VLAN.

- So now, let us assign VLAN 10 to Sales and VLAN 20 to Marketing, as shown in the figure.
- Now if the Sales PC sends out the broadcast packet, it will also reach another Sales PC or the PCs assigned with VLAN 10.
- Similar is the case with Marketing PCs. Whenever a Marketing PC broadcasts a message, it will reach the PCs with VLAN 20.

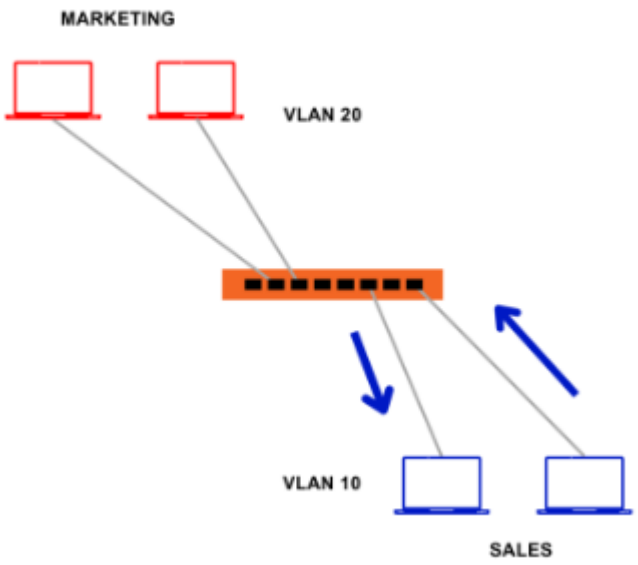

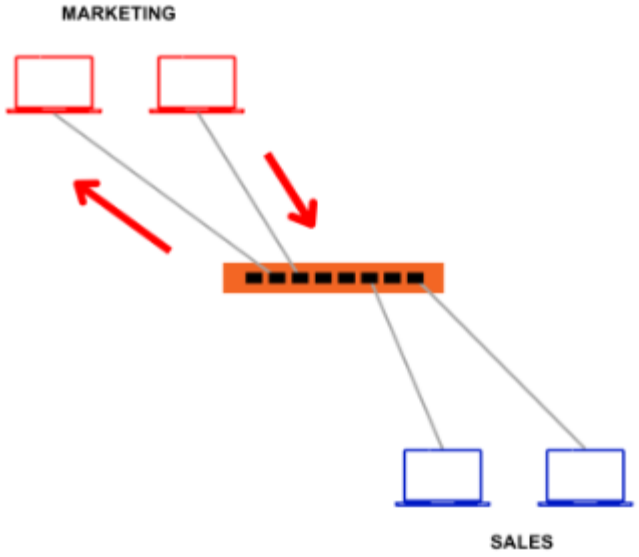

Virtual LANs (VLAN)

- A Virtual Local Area Network (VLAN) is a way of dividing a physical network into multiple logical networks. Instead of relying on the physical layout of devices, VLANs allow network administrators to group devices together based on factors like departments, functions, or security requirements. Devices in the same VLAN can communicate with each other as if they are on the same physical network, even if they are physically located in different areas. VLANs enhance network flexibility, security, and organization by creating isolated groups within a larger network infrastructure.

- Virtual LAN (VLAN) technology enables network architects to segment physical devices into logical subgroups for performance and security reasons.
- A VLAN is a logical subnetwork of devices in a broadcast domain that is partitioned by network switches and/or network management software to act as its own distinct LAN. Switches that support VLANs give network managers the ability to create flexible virtual network segments that are independent of the underlying physical wired or wireless topology.
- VLANs operate at either Layer 2 (data-link layer) or Layer 3 (network layer), depending on the design of the network. Several different network protocols support VLANs, most notably Ethernet and Wi-Fi.
- A VLAN is a virtual LAN that allows you to segment your network without the need for physical segmentation logically. VLANs are very flexible and can be used to provide security, flexibility, and performance benefits. VLANs work by encapsulating Ethernet frames with a VLAN header that contains the VLAN ID. This ID is used to identify which devices are on which VLAN.

- VLANs are created by adding switch ports to a particular VLAN. Devices on the same VLAN can communicate with each other without the need for a router. This makes VLANs very efficient and easy to manage. You can think of a VLAN as a virtual switch that provides isolation between devices.
- VLANs are often used to separate different types of traffic such as voice, video, and data. Voice traffic is typically given priority on a network so it doesn’t experience delays or dropped packets. Video traffic is also given priority on a network so it can stream smoothly without interruptions. Data traffic is typically segregated onto its own VLAN so it doesn’t interfere with real-time traffic such as voice and video.

Benefits of VLAN

- VLANs offer numerous benefits for network management and design, making them a valuable tool in modern networking. Some advantages of using VLANs are:
- **Network Segmentation:** VLANs allow you to logically segment a physical network into multiple virtual networks. This is particularly useful for improving network performance, security, and organization. Different departments or groups can have their own VLANs, keeping their traffic separate from others.
- **Cost and Time Reduction:** VLANs can reduce the migration cost of stations going from one group to another. Physical reconfiguration takes time and is costly. Instead of physically moving one station to another segment or even to another switch, it is much easier and quicker to move it by using software.
- **Improved Security:** VLANs enhance network security by isolating broadcast domains and controlling traffic between VLANs. This segmentation can help prevent unauthorized access to sensitive data or systems, reducing the attack surface for potential security threats.

- **Optimized Traffic Flow:** VLANs enable you to control and prioritize traffic within your network. This is beneficial for ensuring that critical applications or services receive the necessary network resources, leading to better network performance.
- **Simplified Network Management:** Network administrators can manage VLANs more effectively than managing multiple separate physical networks. This simplifies tasks like configuring access controls, monitoring network traffic, and implementing changes across the network.
- **Reduced Broadcast Traffic:** By segregating broadcast domains, VLANs help in reducing broadcast traffic. This minimizes network congestion and prevents broadcast storms that can occur on large, flat networks.
- **Flexibility and Scalability:** VLANs are flexible and can adapt to changing network requirements. Adding new VLANs or reconfiguring existing ones is relatively straightforward, making it easier to accommodate growth and network changes.
- **Geographical Flexibility:** VLANs are not bound by physical locations. They can be used to connect devices and users across different physical sites or locations, making them suitable for geographically distributed organizations.

- **Resource Optimization:** VLANs can help optimize network resources by ensuring that devices that need to communicate with each other are placed within the same VLAN. This reduces unnecessary traffic on the network.
- **Virtual Routing:** Some VLAN setups include virtual routing, allowing traffic between VLANs without the need for a physical router. This simplifies network design and can reduce hardware costs.
- **Disaster Recovery:** VLANs can be a part of disaster recovery and business continuity strategies. Data and services can be replicated across VLANs, ensuring that critical operations continue in case of a network or site failure.
- **Guest Networking:** VLANs are commonly used to create separate guest networks. This allows guest devices to access the internet while keeping them isolated from the organization's internal network for security and compliance reasons.

- **Compliance and Regulation:** VLANs can be a valuable tool for maintaining compliance with regulations such as the Health Insurance Portability and Accountability Act (HIPAA) or the Payment Card Industry Data Security Standard (PCI DSS) by keeping sensitive data separate.
- Overall, VLANs provide a versatile and efficient way to manage and optimize network resources, enhance security, and adapt to changing network requirements. They have become a fundamental element of modern network design and administration.

How VLAN works?

- VLANs work by logically dividing a physical network into multiple isolated virtual networks. These virtual networks enable devices to communicate with each other as if they were on the same physical network, despite being connected to different network switches or segments. Here's how VLANs work
- **Creation and Configuration of VLANs:** Network administrators configure VLANs on network switches or routers that support VLAN functionality. They assign a unique VLAN ID to each VLAN. VLANs can be created based on various criteria, such as department, function, security requirements, or any other logical grouping.
- **Port Assignments:** Administrators assign individual switch ports to specific VLAN based on the connected devices and their locations. Devices connected to a specific port are considered part of the VLAN associated with that port. Some ports can be let unassigned to any VLAN (native VLAN), and any devices connected to these ports belong to that native VLAN.
- **VLAN Tagging:** When a device connected to a VLAN-configured port sends network traffic, the switch adds a VLAN tag (commonly known as an 802.1Q tag) to the Ethernet frames. This tag contains the VLAN ID, which identifies the associated VLAN. The VLAN tag allows switches to distinguish which VLAN the traffic belongs to as it traverses the network.

- **Isolation and Broadcast Control:** Devices within the same VLAN can communicate with each other directly, just as if they were on the same physical network segment. This is accomplished by the switch forwarding traffic only to devices within the same VLAN. Broadcast traffic (e.g., ARP requests) is confined to the VLAN, preventing it from crossing VLAN boundaries. This minimizes network congestion and security risks associated with unnecessary broadcast traffic.
- **Inter-VLAN Routing:** If devices in different VLANs need to communicate, a router or Layer 3 switch is used. These devices act as gateways between the VLANs. The router examines the destination IP addresses of the packets, makes routing decisions, and forwards traffic between VLANs accordingly.
- **Security and Access Control:** VLANs provide network segmentation, which enhances security. Network administrators can control and secure traffic between VLANs using access control lists (ACLs) or firewall rules applied at the router or Layer 3 switch. This isolation helps prevent unauthorized access to sensitive network resources and data.
- **Management and Monitoring:** Network administrators can manage and monitor each VLAN separately. This allows for individual control over configuration settings, security policies, and performance monitoring for each virtual network.

- **Dynamic VLAN Assignment:** In some scenarios, VLANs can be assigned dynamically based on user authentication or specific device characteristics. This approach is commonly used in enterprise and educational environments.
- **Quality of Service (QoS):** VLANs can be used to implement Quality of Service (QoS) policies, ensuring that specific types of traffic receive priority over others, optimizing network performance for critical applications.

VLANs offer scalability and flexibility, allowing network administrators to adapt to changing network requirements and effectively manage complex networks. By separating and controlling network traffic logically, VLANs provide better organization, security, and performance within a single physical network infrastructure.

The Best Practice Backbone Design

- The past few years have witnessed radical changes in the backbone, both in terms of new technologies (e.g., gigabit Ethernet) and in architectures (e.g., VLANs). Fifteen years ago, the most common backbone architecture was the routed backbone, connected to a series of shared 10Base-T hubs in the LAN.
- Today, the most effective architecture for the distribution layer in terms of cost and performance is a switched backbone (either rack-mounted or using a chassis switch) because it provides the best performance at the least cost. For the core layer, most organizations use a routed backbone. Many large organizations are now implementing VLANs, especially those that have departments spread over multiple buildings, but VLANs add considerable cost and complexity to the network.
- Given the trade-offs in costs, there are several best practice recommendations. First, the best practice architecture is a switched backbone or VLAN for the distribution layer and a routed backbone for the core layer. Second, the best practice recommendation for backbone technology is gigabit Ethernet. Considering the LAN and backbone environments together, the ideal network design is likely to be a mix of layer 2 and VLAN Ethernet switches. Figure shows one likely design.

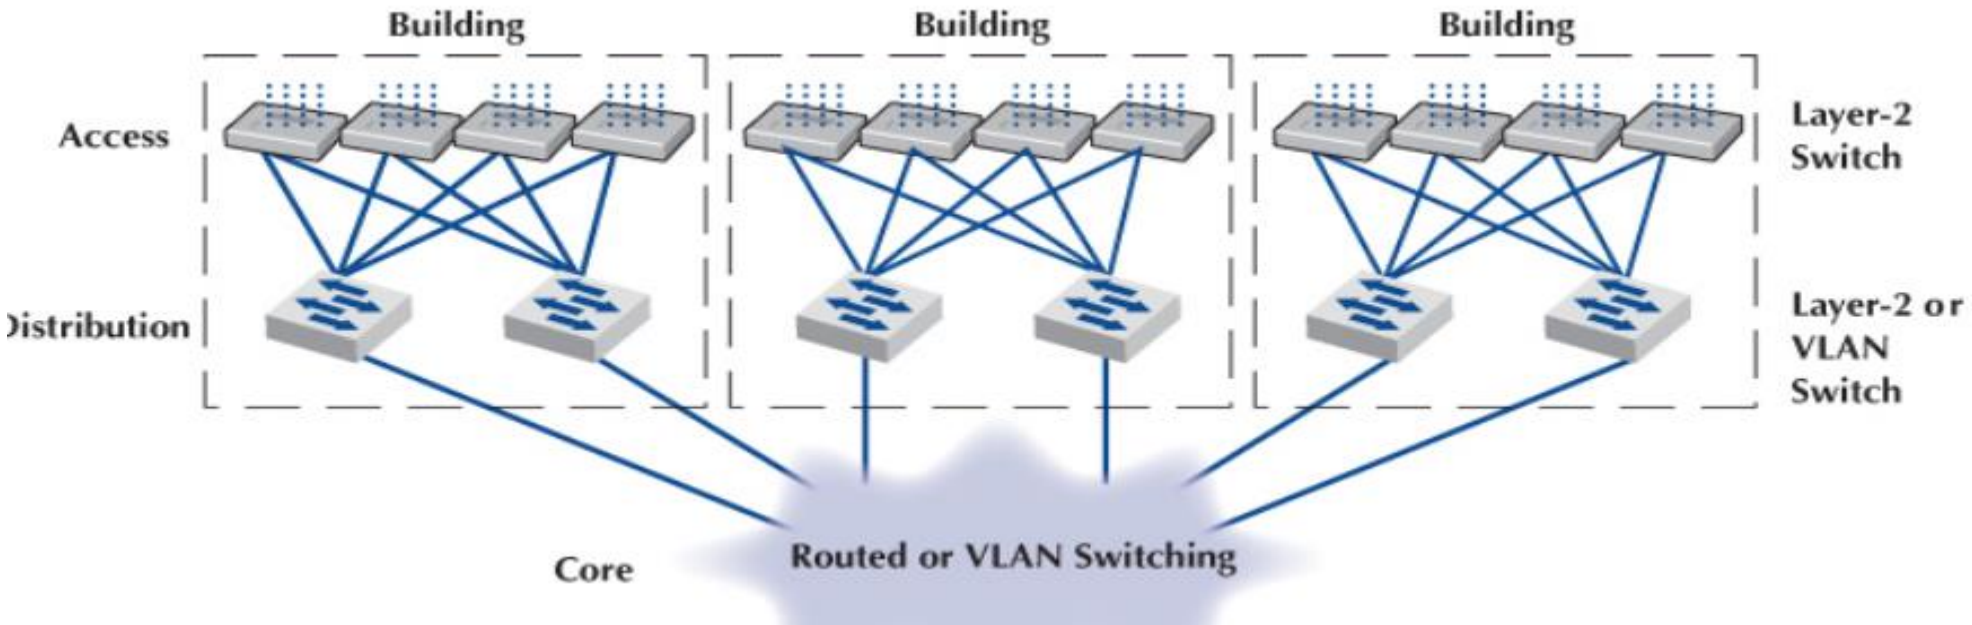

The best practice network design

- The access layer (i.e., the LANs) uses 1000Base-T layer 2 Ethernet switches running on Cat 5e or Cat 6 twisted-pair cables to provide flexibility for 100Base-T or 1000Base-T. The distribution layer uses layer 2 or VLAN switches that use 100Base-T or more commonly 1000Base-T/F (over fiber or Cat 6) to connect to the access layer. To provide good reliability, some organizations may provide redundant switches, so if one fails, the backbone continues to operate. The core layer uses routers or VLAN Ethernet switches running 10 GbE or 40 GbE over fiber.

Improving Backbone Performance

- The method for improving the performance of BNs is similar to that for improving LAN performance. First, find the bottleneck and then remove it (or, more accurately, move the bottleneck somewhere else). You can improve the performance of the network by improving the performance of the devices in the network, by upgrading the circuits between them, and by changing the demand placed on the network.

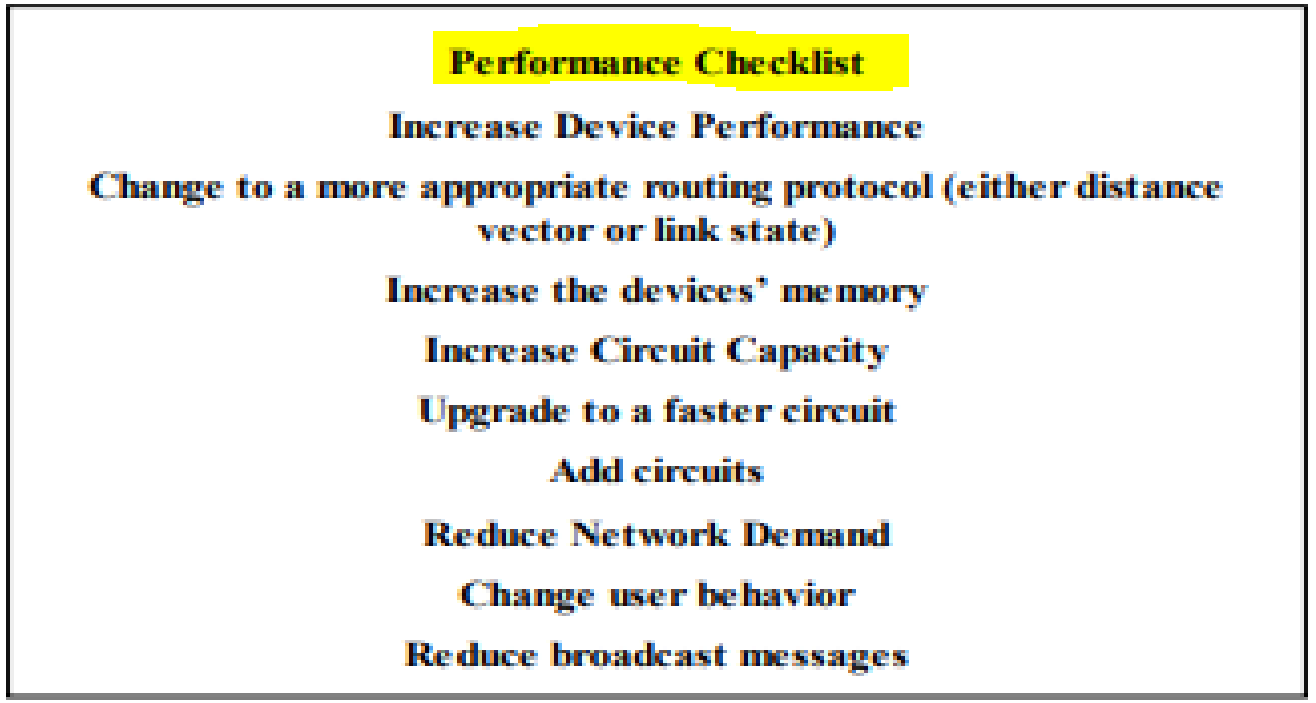

Improving Device Performance

- The primary functions of computers and devices in BNs are forwarding/routing messages and serving up content. If the devices and computers are the bottleneck, routing can be improved with faster devices or a faster routing protocol. Distance vector routing is faster than dynamic routing but obviously can impair circuit performance in high traffic situations. Link state routing is usually used in WANs because there are many possible routes through the network. BNs often have only a few routes through the network, so link state routing may not be too helpful because it will delay processing and increase the network traffic because of the status reports sent through the network. Distance vector routing will often simplify processing and improve performance.
- Most backbone devices are store-and-forward devices. One simple way to improve performance is to ensure that they have sufficient memory. If they don’t, the devices will lose packets, requiring them to be retransmitted.

Improving Circuit Capacity

- If network circuits are the bottlenecks, there are several options. One is to increase circuit capacity (e.g., by going from 100Base-T Ethernet to gigabit Ethernet). Another option is to add additional circuits alongside heavily used ones so that there are several circuits between some devices.
- In many cases, the bottleneck on the circuit is only in one place—the circuit to the server. A switched network that provides 100 Mbps to the client computers but a faster circuit to the server (e.g., 1000Base-T) can improve performance at very little cost.

Reducing Network Demand

- One way to reduce network demand is to restrict applications that use a lot of network capacity, such as desktop videoconferencing, medical imaging, or multimedia. In practice, it is often difficult to restrict users. Nonetheless, finding one application that places a large demand on the network and moving it can have a significant impact.
- Much network demand is caused by broadcast messages, such as those used to find data link layer addresses. Some network operating system and application software packages written for use on LANs also use broadcast messages to send status information to all computers on the LAN. For example, broadcast messages inform users when printers are out of paper or when the server is running low on disk space. When used in a LAN, such messages place little extra demand on the network because every computer on the LAN gets every message.
- This is not the case for routed backbones because messages do not normally flow to all computers, but broadcast messages can consume a fair amount of network capacity in switched backbones. In many cases, broadcast messages have little value outside their individual LAN. Therefore, some switches and routers can be set to filter broadcast messages so that they do not go to other networks. This reduces network traffic and improves performance.