2.1. VLAN Definitions

VLANs provide segmentation and flexibility within switched internetworks by logically grouping devices. They allow devices within a VLAN to communicate as if they are on the same network segment, based on logical connections rather than physical ones.

Administrators use VLANs to segment networks based on functions, project teams, or applications, regardless of user or device location. Devices in a VLAN function as an independent network, even though they share infrastructure with other VLANs. Any switch port can be assigned to a VLAN, forwarding unicast, broadcast, and multicast packets only to end stations within the same VLAN as the source. Each VLAN is a distinct logical network; packets destined for other VLANs require routing through a compatible device.

VLANs create logical broadcast domains that span multiple physical LAN segments, improving network performance by dividing large broadcast domains into smaller, more manageable units. Broadcast Ethernet frames from one VLAN are received exclusively by devices within the same VLAN.

VLANs also facilitate the enforcement of access and security policies based on user groupings. Typically, each switch port is assigned to a single VLAN, with exceptions for ports connected to IP phones or other switches.

In many organizations, frequent organizational changes occur:

- System administrators often spend a significant amount of time unplugging and reconnecting cables when making changes.
- In some cases, changes cannot be implemented due to the physical distance between a user's machine and the correct switch.
- To address these challenges, network vendors started developing a way to reconfigure network structures entirely through software.
- The resulting concept is known as a Virtual LAN (VLAN), which has been standardized by the IEEE 802 committee and is now widely deployed in many organizations.

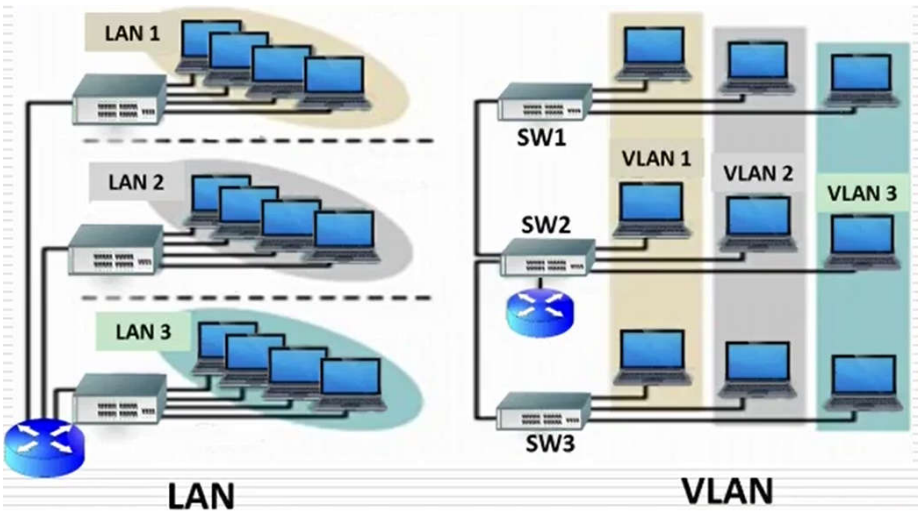

Figure 2.1. LAN vs. VLAN

2.2. VLANs principle

A switch doesn't fully understand broadcast messages. When it receives a broadcast message on one of its ports, it forwards that message out to all other ports (***Un Switch présente un seul domaine de diffusion, des diffusions qui consomment la bande passante***). To illustrate this, consider an 8-port switch. If it receives a broadcast message on port 1, it will relay that message to ports 2 through 8. The following image provides a visual representation of this process.

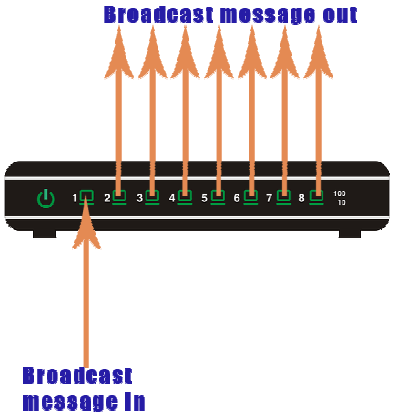

A VLAN is a feature specific to switches, allowing us to group ports that share broadcast messages. If two switch ports belong to different VLANs, they do not share broadcast messages. Conversely, if two ports belong to the same VLAN, they share broadcast messages.

Let's consider the previous example. We create two VLANs: VLAN-10 and VLAN-20 on the switch. We assign ports 1 to 4 to VLAN-10 and ports 5 to 8 to VLAN-20. As a result, ports 1, 2, 3, and 4 will share broadcast messages in VLAN-10, while ports 5, 6, 7, and 8 will share broadcast messages in VLAN-20. The following image illustrates this concept.

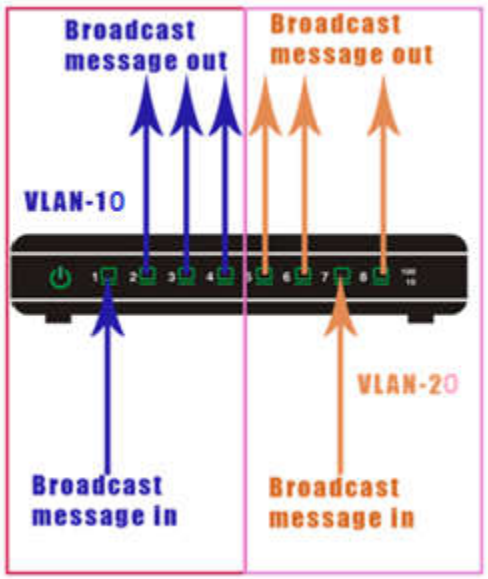

VLANs are not confined to a single switch; you can create and utilize them across multiple switches. This capability allows you to logically organize your network. (***Un VLAN peut exister sur plusieurs commutateurs. Un commutateur peut contenir plusieurs VLANs***)

2.3. Benefits of VLANs

User productivity and network adaptability are crucial for business growth. VLANs facilitate tailored network design to align with organizational objectives. The key advantages of VLAN usage include:

Enhanced Security: VLANs isolate sensitive data groups, minimizing the risk of data breaches. For instance, in the illustration, faculty computers on VLAN 10 are entirely separate from student and guest data traffic.

Cost Efficiency: VLANs reduce the need for costly network upgrades, optimizing existing bandwidth and uplinks for cost savings.

Optimized Performance: Segmenting flat Layer 2 networks into logical workgroups curtails unnecessary network traffic, bolstering overall performance.

Broadcast Domain Reduction: VLAN segmentation diminishes the device count within a broadcast domain. In the example, six computers exist on the network, grouped into three broadcast domains: Faculty, Student, and Guest.

Streamlined IT Management: VLANs simplify network administration by grouping users with similar needs. Configurations and policies are automatically applied when ports are assigned during provisioning of a new switch. IT staff can easily identify a VLAN's purpose through appropriately labeled names, as seen in the figure.

Efficient Project and Application Handling: VLANs aggregate users and devices as per business or geographic requirements. Separate functions facilitate streamlined project management and specialized application usage, such as an e-learning platform for faculty.

Each VLAN in a switched network corresponds to an IP network, necessitating a hierarchical network-addressing approach. Hierarchical network addressing involves methodical assignment of IP network numbers to segments or VLANs, considering the entire network. Contiguous address blocks are designated for specific network sections, as depicted in the figure.

2.4 Types of VLANs

There are a number of distinct types of VLANs used in modern networks. Some VLAN types are defined by traffic classes. Other types of VLANs are defined by the specific function that they serve.

Data VLAN

A data VLAN is a VLAN that is configured to carry user-generated traffic. A VLAN carrying voice or management traffic would not be part of a data VLAN. It is common practice to separate voice and management traffic from data traffic. A data VLAN, is sometimes referred to as a user VLAN. Data VLANs are used to separate the network into groups of users or devices.

Default VLAN

All switch ports become a part of the default VLAN after the initial boot up of a switch loading the default configuration. Switch ports that participate in the default VLAN are part of the same broadcast domain. This allows any device connected to any switch port to communicate with other devices on other switch ports. The default VLAN for Cisco switches is VLAN 1. In the figure, the show vlan

brief command was issued on a switch running the default configuration. Notice that all ports are assigned to VLAN 1 by default.

VLAN 1 has all the features of any VLAN, except it cannot be renamed or deleted. By default, all Layer 2 control traffic is associated with VLAN 1.

Native VLAN

A native VLAN is assigned to an 802.1Q trunk port. Trunk ports are the links between switches that support the transmission of traffic associated with more than one VLAN. An 802.1Q trunk port supports traffic coming from many VLANs (tagged traffic), as well as traffic that does not come from a VLAN (untagged traffic). The 802.1Q trunk port places untagged traffic on the native VLAN, which by default is VLAN 1.

Native VLANs are defined in the IEEE 802.1Q specification to maintain backward compatibility with untagged traffic common to legacy LAN scenarios. A native VLAN serves as a common identifier on opposite ends of a trunk link.

It is a best practice to configure the native VLAN as an unused VLAN, distinct from VLAN 1 and other VLANs. In fact, it is not unusual to dedicate a fixed VLAN to serve the role of the native VLAN for all trunk ports in the switched domain.

Management VLAN

A management VLAN is any VLAN configured to access the management capabilities of a switch. VLAN 1 is the management VLAN by default. To create the management VLAN, the switch virtual interface (SVI) of that VLAN is assigned an IP address and subnet mask, allowing the switch to be managed via HTTP, Telnet, SSH, or SNMP. Because the out-ofthe-box configuration of a Cisco switch has VLAN 1 as the default VLAN, VLAN 1 would be a bad choice for the management VLAN.

In the past, the management VLAN for a 2960 switch was the only active SVI. On 15.x versions of the Cisco IOS for Catalyst 2960 Series switches, it is possible to have more than one active SVI. With Cisco IOS 15.x, the particular active SVI assigned for remote management must be documented. While theoretically a switch can have more than one management VLAN, having more than one increases exposure to network attacks.

Voice VLANs

A separate VLAN is needed to support Voice over IP (VoIP). VoIP traffic requires:

- Assured bandwidth to ensure voice quality
- Transmission priority over other types of network traffic
- Ability to be routed around congested areas on the network
- Delay of less than 150 ms across the network

2.4. IEEE 802.1Q Frame Format (VLAN Frame Format)

A switch identifies packets from different VLANs according to the information contained in its VLAN tags.

VLAN Tagging, also known as Frame Tagging, is a method developed to help identify packets travelling through trunk links. When an Ethernet frame traverses a trunk link, a special VLAN tag is added to the frame and sent across the trunk link. As it arrives at the end of the trunk link the tag is removed and the frame is sent to the correct access link port according to the switch's table, so that the receiving end is unaware of any VLAN information.

A trunk is a point-to-point link between the switch and another networking device. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across an entire network.

Two trunking encapsulations are available on all Ethernet ports:

- ***Inter-Switch Link (ISL)***: ISL is a Cisco-proprietary trunking encapsulation.
- ***802.1Q: 802.1Q*** is an industry-standard trunking encapsulation.

IEEE 802.1Q adds a 4-byte VLAN tag between the Source/Destination MAC address and Length/Type fields of an Ethernet frame to identify the VLAN to which the frame belongs. The following figure shows the position of a VLAN tag in a VLAN data frame.

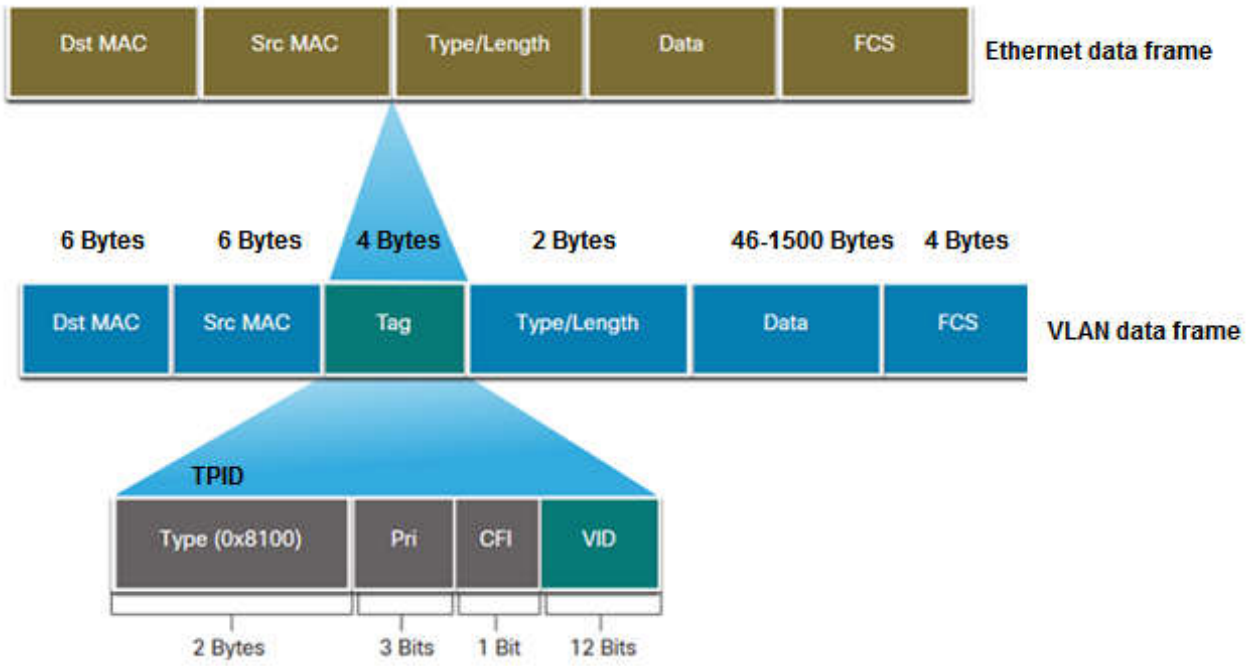

The following are the two types of Ethernet frames in a VLAN:

- Tagged frame: frame with a 4-byte VLAN tag
- Untagged frame: frame without a 4-byte VLAN tag

There are two main types of Ethernet links: access links (transmit untagged frames) and trunk links (transmit tagged frames). The two link types differ in the number of VLANs they can carry traffic for: an access link can carry traffic for only one VLAN, and therefore usually connects a switch to a user terminal, such as a host, a server, or an unmanaged switch; a trunk link can carry traffic for multiple VLANs, and as such usually connects a switch to another switch or a router. The frames transmitted on an access link are untagged frames, and the frames transmitted on a trunk link are tagged frames.

All frames processed on a switch carry VLAN tags. After receiving an untagged frame from a user terminal, the switch adds a VLAN tag to the frame, recalculates the frame check sequence (FCS), and then transmits the frame through a trunk link. Before sending the frame to a user terminal, the switch removes the VLAN tag, and sends the untagged frame to the terminal through an access link.

A VLAN tag contains four fields. Table 1 describes the fields.

Table 2-1 Fields in a VLAN tag

Field	Length	Description	Value
TPID	2 bytes	Tag Protocol Identifier (TPID), indicating the frame type.	The value 0x8100 indicates an IEEE 802.1Q frame. An 802.1Q-incapable device discards 802.1Q frames. Device vendors can define their own TPID values, and users can then change the value to realize interconnection of devices from different vendors.
PRI	3 bits	Priority (PRI), indicating the 802.1p priority of a frame.	The value is in the range from 0 to 7. A larger value indicates a higher priority. If congestion occurs, the switch sends packets with the highest priority first.
CFI	1 bit	Canonical Format Indicator (CFI), indicating whether a MAC address is encapsulated in canonical format over different transmission media. CFI is used to ensure compatibility between Ethernet and token ring networks.	The value 0 indicates that the MAC address is encapsulated in canonical format, and the value 1 indicates that the MAC address is encapsulated in non-canonical format. The CFI field has a fixed value of 0 on Ethernet networks.

Field	Length	Description	Value
VID	12 bits	VLAN ID (VID), indicating the VLAN to which a frame belongs.	The VLAN ID is in the range from 0 to 4095. The values 0 and 4095 are reserved, and therefore available VLAN IDs are in the range from 1 to 4094.

2.5. DTP Dynamic Trunking Protocol

Trunk ports are ports that are used to carry the traffic of more than one VLAN. The port which connects two different switches and the switches have more than one VLAN configured then that port should be made trunk. If all the VLANs are allowed then trunk ports will carry traffic of all the VLANs including native VLANs for which the traffic goes untagged otherwise only the allowed VLANs traffic will be carried by the trunk link. The trunk link traffic would be encapsulated or tagged either by ISL or 802.1Q.

By default, all switch ports are access ports therefore to make a port trunk, the user should manually make it trunk by using DTP.

Dynamic Trunking Protocol –

Dynamic Trunking Protocol is CISCO proprietary protocol used for negotiating a trunk link between two switches as well as the encapsulation type of either 802.1q or ISL (Generally, 802.1q is used because ISL has more overhead than 802.1q). Of course, It is a layer 2 (data link) protocol and is enabled by default.

The Dynamic Trunking Protocol (DTP) manages trunk autonegotiation on LAN ports. DTP supports autonegotiation of both ISL and 802.1Q trunks.

To autonegotiate trunking, the LAN ports must be in the same VTP domain. Use the **trunk** or **nonegotiate** keywords to force LAN ports in different domains to trunk.

Here are the different options available while configuring a switch interface:

switchport mode access (DTP mode OFF) –

This mode puts the switch interface into permanent non-trunking mode regardless of whether the neighbouring interface is a trunk port or trying to become a trunk port that is why it is known as DTP mode OFF. The port is a dedicated layer 2 access port.

switchport mode trunk (DTP mode ON) –

It puts the interface into trunking mode. The interface will become a trunk interface even if the neighbouring ports are trunk or not that is why it is called DTP mode ON.

switchport mode dynamic auto –

This is a default mode on the older CISCO switches. This mode makes the interface able to convert to a trunk link. The interface will become a trunk link if the neighbouring interface

is set to trunk or desirable mode. If both switches interface mode is auto, then the trunk will not be formed.

switchport mode dynamic desirable –

By this mode, the interface will actively attempt to convert the link into a trunk link. The interface will become a trunk port if the neighbouring interface is set to trunk, desirable or auto.

switchport nonegotiate

This mode prevents the interface from generating DTP frames. This command is used only when the switch port mode is access or trunk. You must manually configure the neighbouring interface as a trunk interface to establish a trunk link.

Now, let's see the scenarios in which the switch interface will either become a trunk or access interface.

	Dynamic Auto	**Dynamic desirable**	**trunk**	**access**
Dynamic Auto	**access**	**trunk**	**trunk**	**access**
Dynamic desirable	**trunk**	**trunk**	**trunk**	**access**
trunk	**trunk**	**trunk**	**trunk**	**limited connectivity**
access	**access**	**access**	**limited connectivity**	**access**

By observing this, it is clear that whenever you receive a DTP packet that requests to form a trunk, your interface will be in trunk mode.

Here are some of its features:

- **Automatic negotiation:** DTP allows switches to automatically negotiate the formation of a trunk link without requiring manual configuration.
- **Four modes:** DTP has four modes: "dynamic auto", "dynamic desirable", "trunk", and "access". The mode selected on each switch determines the behavior of the negotiation process.
- **Proprietary protocol:** DTP is a proprietary protocol developed by Cisco and is only supported on Cisco devices.
- **Can pose security risks:** DTP can pose security risks if not configured properly, as it can allow unauthorized devices to form a trunk link with a switch.

- **Can improve network performance:** DTP can improve network performance by allowing switches to form trunk links automatically and efficiently manage network traffic.

- **Can cause issues in mixed environments:** DTP can cause issues in mixed network environments where non-Cisco devices do not support the protocol or do not behave as expected.

Advantages of Dynamic Trunking Protocol (DTP)

- DTP simplifies the process of configuring and managing VLANs on a network. It enables network administrators to dynamically negotiate trunk links between switches without manual configuration.

- DTP reduces the risk of misconfiguration errors and simplifies the task of managing VLANs in large networks with multiple switches.

- DTP allows for automatic creation and deletion of VLANs on switches, which can save time and reduce the risk of configuration errors.

- DTP provides a quick and easy way to connect switches together and establish VLAN communication.

Disadvantages of Dynamic Trunking Protocol (DTP)

- DTP can create security vulnerabilities in the network. It enables automatic negotiation of trunk links, which can lead to unauthorized switches being connected to the network, potentially compromising network security.

- DTP can result in increased network traffic due to the constant negotiation of trunk links between switches.

- DTP can lead to misconfiguration of VLANs in situations where a switch negotiates a trunk link with a different VLAN configuration than intended.

- DTP is a Cisco proprietary protocol, which means it may not be compatible with non-Cisco switches or other network devices. This can limit its usefulness in heterogeneous networks.

2.5. Inter-VLAN Routing (*How do you permit devices on separate VLANs to communicate?*)

VLANs are used to segment switched Layer 2 networks for a variety of reasons, such as improving security, performance, and manageability. However, hosts in one VLAN cannot communicate with hosts in another VLAN unless there is a router or Layer 3 switch to provide routing services.

Inter-VLAN routing is the process of forwarding network traffic from one VLAN to another VLAN. There are three main inter-VLAN routing options:

- Legacy inter-VLAN routing: This is a legacy solution that does not scale well.
- Router-on-a-stick: This is an acceptable solution for small to medium-sized networks.
- Layer 3 switch using switched virtual interfaces (SVIs): This is the most scalable solution for medium to large organizations.

Legacy Inter-VLAN Routing

Legacy inter-VLAN routing is a simple but inefficient method of routing traffic between VLANs. It requires a router with multiple Ethernet interfaces, one for each VLAN. The router interfaces then serve as the default gateways for the local hosts on each VLAN subnet. Legacy inter-VLAN routing is not very scalable and is no longer widely implemented.

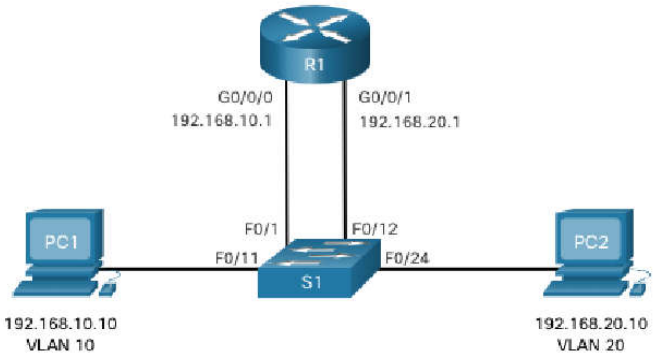

Router-on-a-stick inter-VLAN routing

Router-on-a-stick inter-VLAN routing uses a single physical Ethernet interface to route traffic between multiple VLANs. To do this, the router interface is configured as an 802.1Q trunk and connected to a trunk port on a Layer 2 switch. The router then uses subinterfaces to identify routable VLANs and assign them IP addresses and VLAN tags. When VLAN-tagged traffic enters the router interface, it is forwarded to the appropriate subinterface. The router then makes a routing decision based on the destination IP network address and determines the exit interface for the traffic.

If the exit interface is configured as an 802.1Q subinterface, the data frames are VLAN-tagged with the new VLAN and sent back out the physical interface.

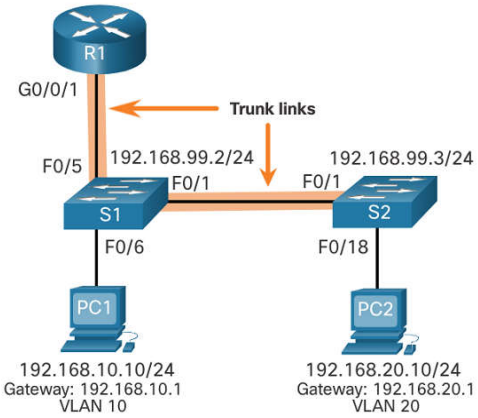

Inter-VLAN Routing on a Layer 3 Switch

The modern method of performing inter-VLAN routing is to use Layer 3 switches and switched virtual interfaces (SVIs). An SVI is a virtual interface that is configured on a Layer 3 switch and assigned to a specific VLAN. SVIs perform the same functions as router interfaces, providing Layer 3 processing for packets that are sent to or from all switch ports associated with that VLAN.

Advantages of using Layer 3 switches for inter-VLAN routing:

- Faster than router-on-a-stick because everything is hardware switched and routed.
- No need for external links from the switch to the router for routing.
- Not limited to one link because Layer 2 EtherChannels can be used as trunk links between the switches to increase bandwidth.
- Lower latency because data does not need to leave the switch in order to be routed to a different network.
- More commonly deployed in a campus LAN than routers.

Disadvantage:

- Layer 3 switches are more expensive.