


Link Layer

1. What is a link?
2. How do we share a network medium?
3. How do we identify link adapters?
4. ~~What is Ethernet?~~
5. How do we interconnect segments at the link layer?

Ethernet Segments Used to be Connected Using Hubs



Hubs

Hubs repeat bits from one port to all other ports

Pros:

- Simple, cheap

Cons:

- inefficient, each bit is sent everywhere
 - limits the aggregates throughput
- limited to one LAN technology
 - can't interconnect different rates/formats
- Lack of traffic isolation
 - Security vulnerabilities

Hubs

Nobody uses hubs. We've moved to **switches**

Switches


Switches connect two or more LANs together at the Link Layer and act as L2 gateways

They are “store and forward” devices:


- extract the destination MAC from the frame
- look up the MAC in a table (using exact match)
- forward the frame on the appropriate interface

Switches are similar to IP routers, except that they operate one layer below

Switches Allow Each LAN to Carry Its Own Traffic



Switches Support Concurrent Communications



Advantages of Switches

- only forward frames where needed
 - avoids unnecessary load on segments
- join segment using different technologies
- improved privacy
 - host can only snoop traffic traversing their segment
- wider-geographic span
 - separates segments allow longer distance
- Plug and play
 - The build their forwarding table on their own

Switches

Switches connect two or more LANs together at the Link Layer and act as L2 gateways

They are “store and forward” devices:


- extract the destination MAC from the frame
- look up the MAC in a table (using exact match)
- forward the frame on the appropriate interface

Switches are similar to IP routers, except that they operate one layer below

Switch Store and Forward

When a frame arrives:

- inspect the source MAC address
- associate the address with the port
- store the mapping in the switch table
- launch a timer to eventually forget the mapping




switch learns how to reach A

If Destination is Not in the Table: Flood

When a frame arrives with an unknown destination

- forward the frame out of all interfaces
 - except for the one where the frame arrived


Hopefully, this is an unlikely event



when in doubt, **shout!**

Flooding Can Create Problems if Network Has a Loop

exponential increase, with no TTL to remove looping frames...



Loop Mitigation: Spanning Trees


Answer:

- Reduce the network to one logical spanning tree
- Upon failure, automatically rebuild a spanning tree

Spanning-trees have only one path between any two nodes

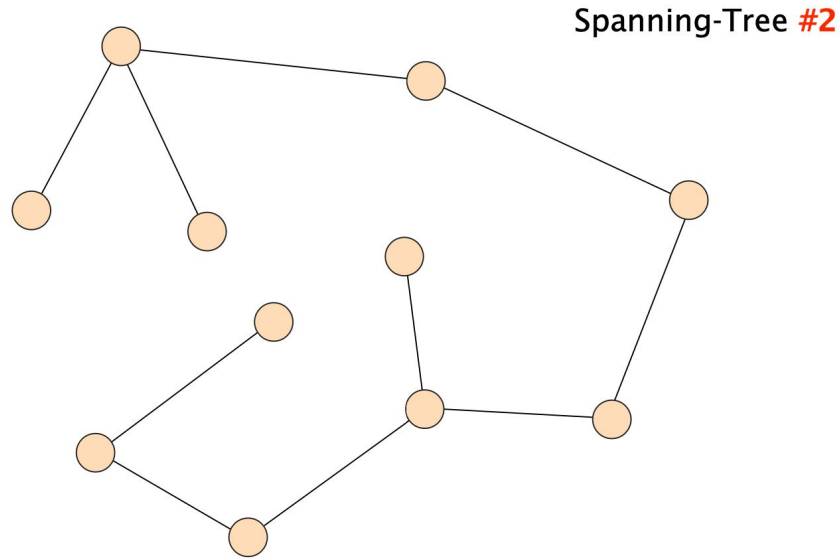
Spanning Trees

There can be many spanning trees for a given topology



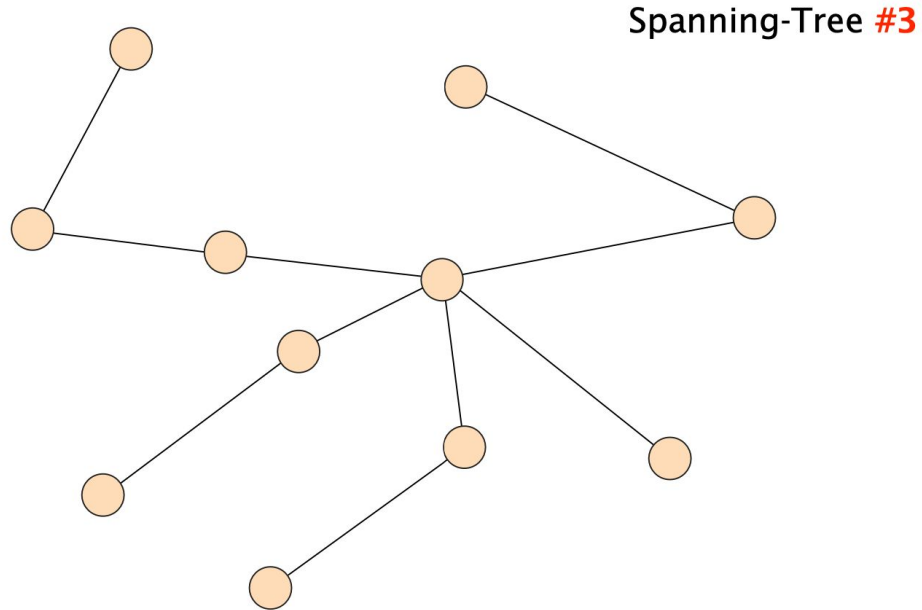
Spanning Trees

There can be many spanning trees for a given topology




Spanning Trees

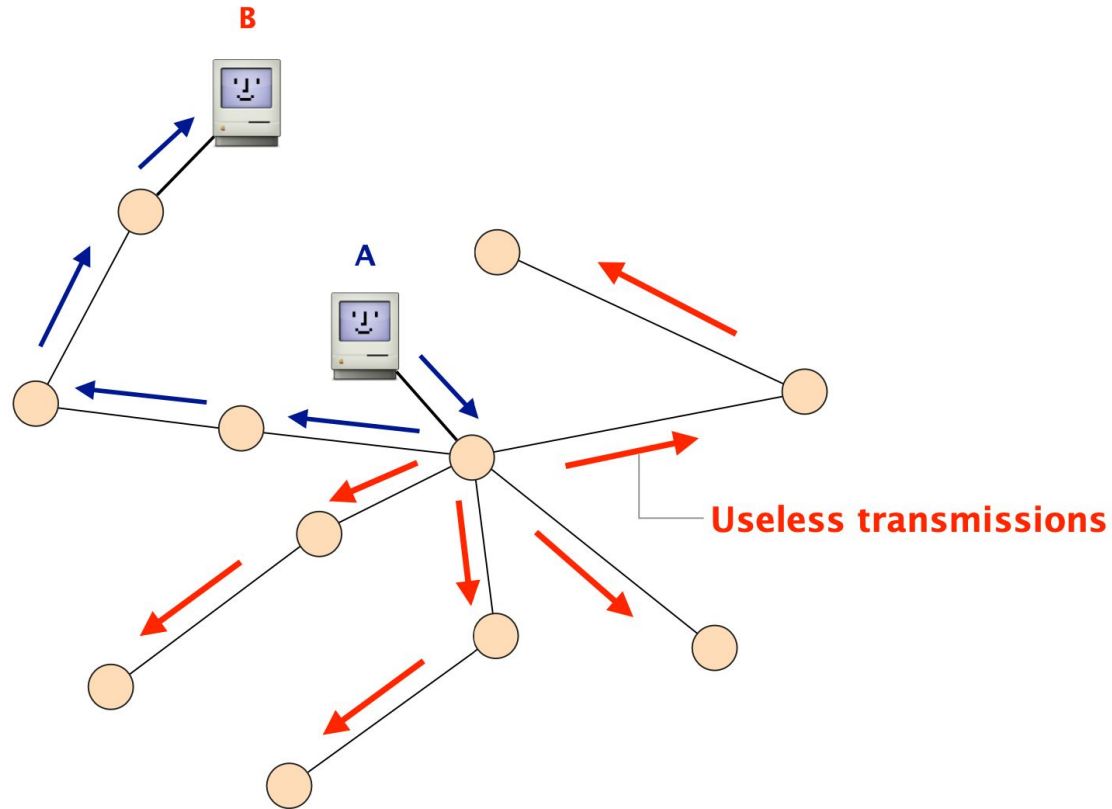
There can be many spanning trees for a given topology



Once You Have a Spanning Tree, Flooding is Safe / Easy



Flooding Remains Wasteful



Constructing a Spanning Tree


Switches...

- elect a root switch
 - the one with the smallest identifier
- determine if each interface is on the shortest-path from the root
 - disable it if not


Constructing a Spanning Tree

- Select a root bridge (typically lowest ID, can be configured)
- All ports on bridge become “designated”
- All ports on other switches facing the root become “root” ports
- Lower ID switches block loop ports

Constructing a Spanning Tree



Constructing a Spanning Tree




STP Must React to Failure

- Any switch, link or port can fail
 - including the root switch
- Root switch continuously sends messages
 - announcing itself as the root, others forward it
- Failures detected through timeout (soft state)
 - if no word from root in X , times out and claims to be the root

STP

Consider the switched network depicted in the figure. It is composed of 5 Ethernet switches, two hosts (connected to switch 3 and 4, respectively) and one IP router acting as default gateway for the hosts. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have equal cost. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

- In the figure, indicate all the links that end up being deactivated in the final state, once all the switches have converged towards the final spanning tree.
- Unsurprisingly, a lot of traffic is exchanged between Host 1 and Internet destinations. Briefly explain two distinct reasons why this configuration is not optimal in terms of network utilization/throughput.




An Ethernet network running the spanning tree protocol.

STP

Consider the switched network depicted in the figure. It is composed of 5 Ethernet switches, two hosts (connected to switch 3 and 4, respectively) and one IP router acting as default gateway for the hosts. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have equal cost. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

- In the figure, indicate all the links that end up being deactivated in the final state, once all the switches have converged towards the final spanning tree.
Solution: Links (4,5), (3,5) and (2,3) end up disabled.
- Unsurprisingly, a lot of traffic is exchanged between Host 1 and Internet destinations. Briefly explain two distinct reasons why this configuration is not optimal in terms of network utilization/throughput.
Solution: Any communication between Host 1 and IP router goes over 4 links. Plus, these links are shared meaning Host 1 and Host 2 will be competing for throughput.




An Ethernet network running the spanning tree protocol.

STP

Consider the switched network depicted in the figure. It is composed of 5 Ethernet switches, two hosts (connected to switch 3 and 4, respectively) and one IP router acting as default gateway for the hosts. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have equal cost. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

- Realizing that there is a problem with their configuration, the network operators ask you to help them improve their network performance. Briefly explain how you would adapt the configuration of the spanning tree protocol (i.e., the switches identifier and/or the link costs) so as to maximize the throughput between Host 1 and Internet destinations.




An Ethernet network running the spanning tree protocol.

STP

Consider the switched network depicted in the figure. It is composed of 5 Ethernet switches, two hosts (connected to switch 3 and 4, respectively) and one IP router acting as default gateway for the hosts. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have equal cost. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

- Realizing that there is a problem with their configuration, the network operators ask you to help them improve their network performance. Briefly explain how you would adapt the configuration of the spanning tree protocol (i.e., the switches identifier and/or the link costs) so as to maximize the throughput between Host 1 and Internet destinations.

Solution: Flipping the switch IDs so that the now-switch 5 becomes the root (e.g. making it switch 1 and the current switch 1, switch 5).




An Ethernet network running the spanning tree protocol.

STP

Consider the switched network depicted in the figure. It is composed of 5 Ethernet switches, two hosts (connected to switch 3 and 4, respectively) and one IP router acting as default gateway for the hosts. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have equal cost. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

- The network operators are happy with your changes. But they now realize that Host 1 and Host 2, in addition to exchanging a lot of Internet traffic, also exchange a lot of traffic between themselves. The network operators ask for your help again. They ask you to find a spanning tree configuration such that: (i) the number of hops between any of these three hosts (Host 1 and 2, and the router) is equivalent; and at the same time (ii) the number of hops is minimal considering the given topology.

Briefly explain how you would configure the spanning tree protocol to achieve these requirements, or why these requirements are impossible to achieve.



An Ethernet network running the spanning tree protocol.


STP

Consider the switched network depicted in the figure. It is composed of 5 Ethernet switches, two hosts (connected to switch 3 and 4, respectively) and one IP router acting as default gateway for the hosts. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have equal cost. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

- The network operators are happy with your changes. But they now realize that Host 1 and Host 2, in addition to exchanging a lot of Internet traffic, also exchange a lot of traffic between themselves. The network operators ask for your help again. They ask you to find a spanning tree configuration such that: (i) the number of hops between any of these three hosts (Host 1 and 2, and the router) is equivalent; and at the same time (ii) the number of hops is minimal considering the given topology.

Briefly explain how you would configure the spanning tree protocol to achieve these requirements, or why these requirements are impossible to achieve.

Solution: Requirements are impossible to achieve: Either the hosts are using their direct link with each other, or with the router. But they cannot all use the direct link between themselves as otherwise that would cause a loop which would be prevented by the spanning tree protocol.



An Ethernet network running the spanning tree protocol.

Broadcast Domains

We've mostly been discussing single broadcast domains


- If one user broadcast a frame, every other user receives it

Larger networks are segmented further


- Improves security
 - smaller attack surface (visibility & injection)
- Improves performance
 - limit the overhead of broadcast traffic (e.g. ARP)
- Improves logistics
 - separates traffic by role (e.g. staff, students, visitors)

Enter VLANs (Virtual Local Area Networks)

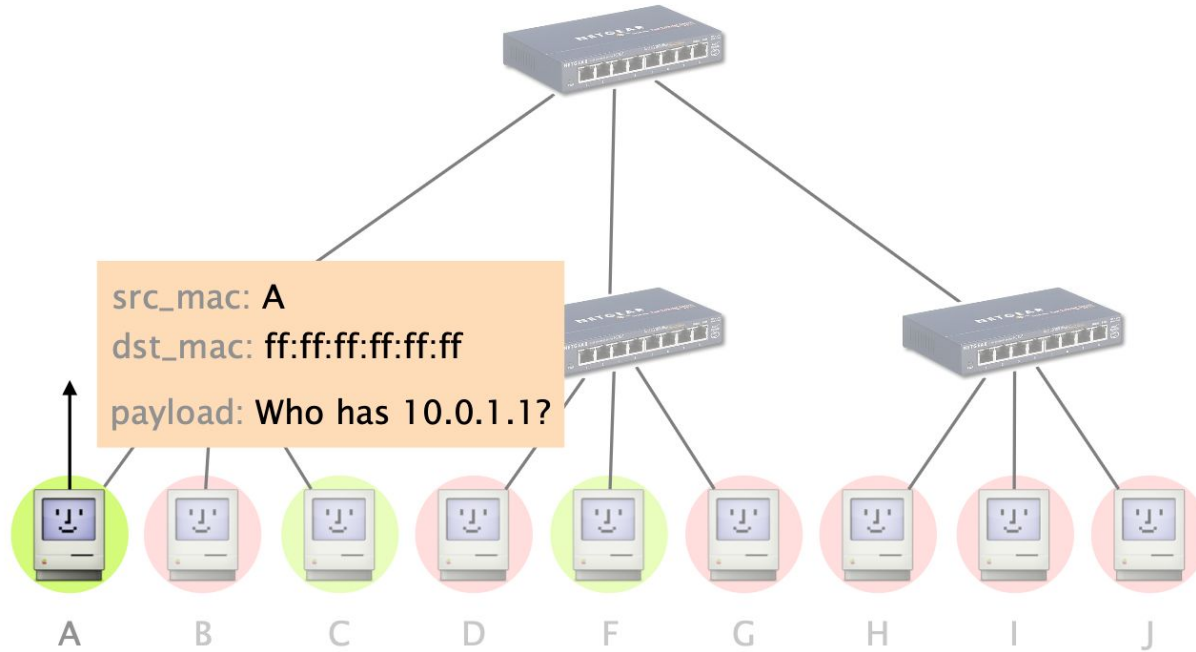
A VLAN logically identifies a set of ports attached to one (or more) Ethernet switches, forming one broadcast domain



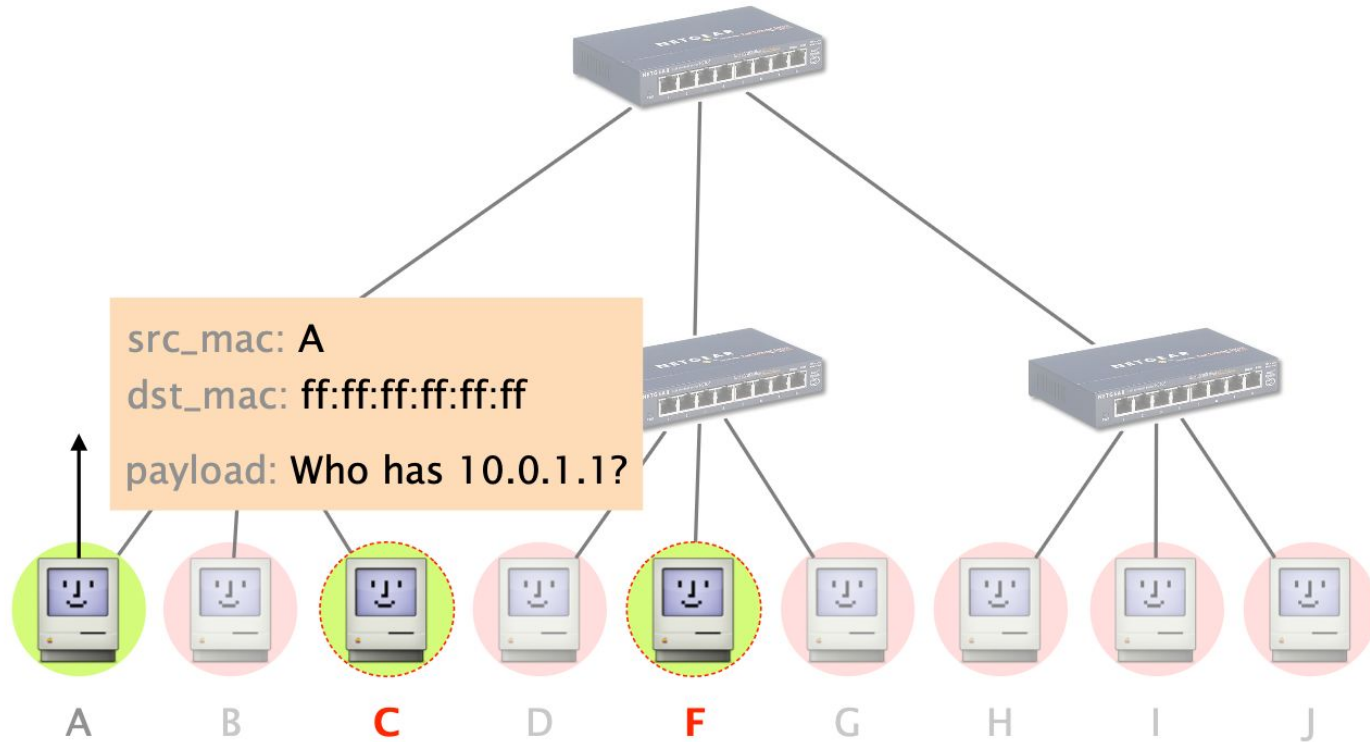
Switches Hold Config Tables Telling them VLAN/Port Info



Broadcasts with VLANs - A sends an ARP request



Only Staff VLAN Members (C and F) Should Receive It



Only Staff VLAN Members (C and F) Should Receive It


src_mac: A

dst_mac: ff:ff:ff:ff:ff:ff

payload: Who has 10.0.1.1?

How does this switch know
where to flood it next?

e.g. pointless to send it right

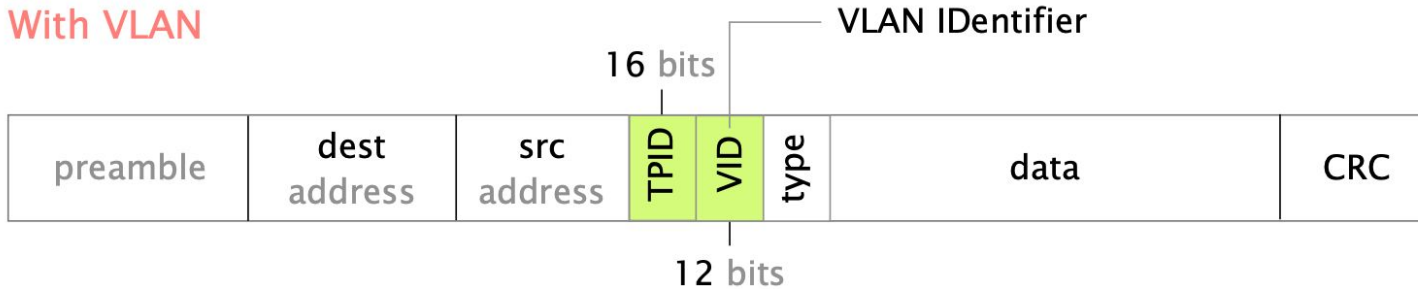


VLAN Tags are Added to the Ethernet Frame

Without VLAN




With VLAN




802.1q Header (4 bytes)
(4 bits missing)

With VLANs, Links are Either “Access” or “Trunk” Links

Access Links Only Belong to One VLAN and Do Not Carry 802.1q Headers



Trunk Links Carry Traffic for More Than One VLAN and Use 802.1q Headers



VLAN Learning

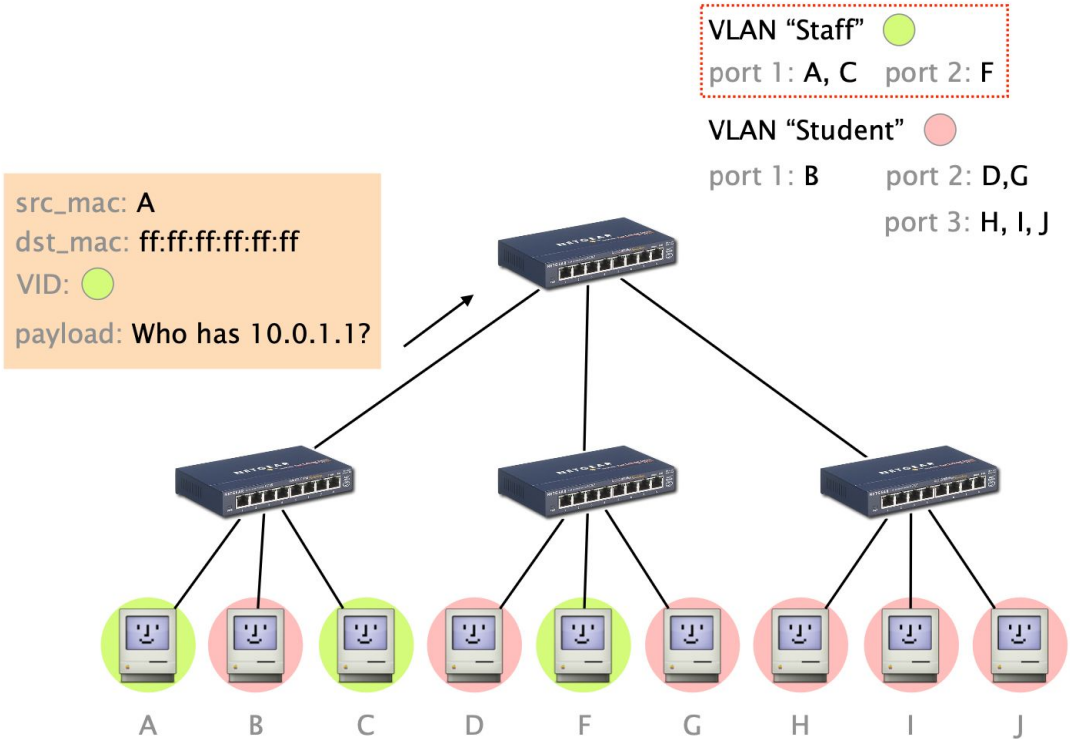
When a switch receives a frame with an unknown or a broadcast destination,

- it forwards it over all the ports that belong to the same VLAN

When a switch learns a source address on a port

- it associates it to the VLAN of this port and only uses it when forwarding frames on this VLAN


VLAN Learning



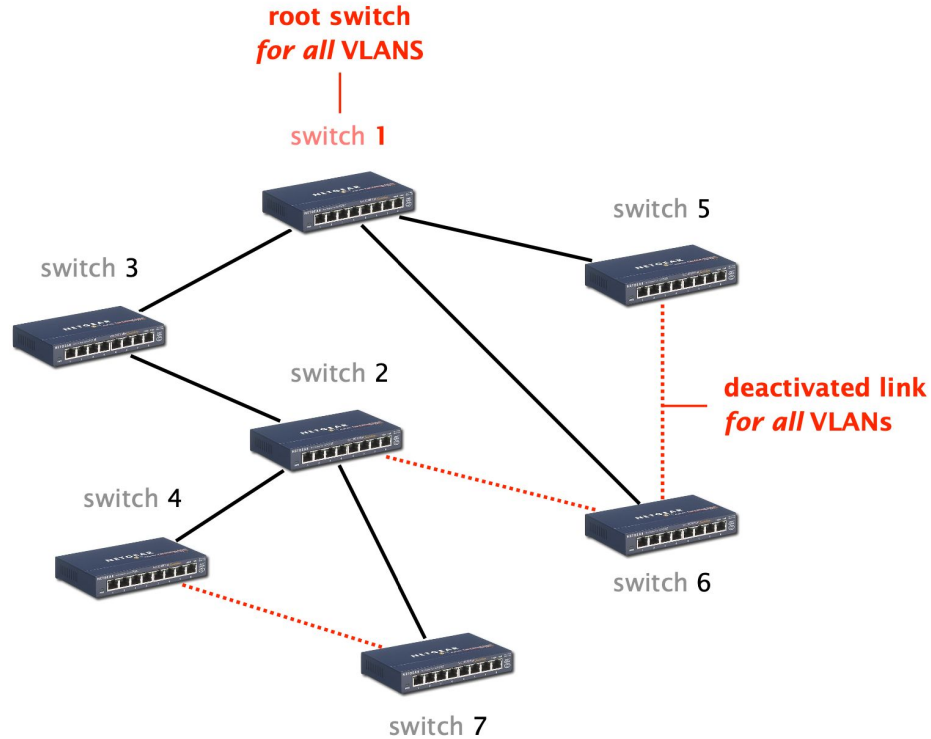
VLAN STP

Switches can also create per-VLAN spanning trees, allowing operators to use more of their links


VLAN STP




VLAN STP - Non Per-VLAN Setup




VLAN STP - Non Per-VLAN Setup




Any Communication Between Red Hosts on 5 and 6 Need to Go Through 1




VLAN STP - Per-VLAN STP Setup



VLAN STP - Per-VLAN STP Setup




VLAN STP - Now Hosts on 5 and 6 Can Use Direct Link



VLAN STP

A network below consists of 9 switches and hosts in two different VLANs (blue and red).


Compute a spanning tree in the network using switch 1 as root. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better). Clearly indicate the type of each link (trunk, access or deactivated).



VLAN STP

A network below consists of 9 switches and hosts in two different VLANs (blue and red).


Compute a spanning tree in the network using switch 1 as root. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better). Clearly indicate the type of each link (trunk, access or deactivated).



VLAN STP

Using the previously computed spanning tree, which path will the red host connected to switch 7 use to communicate with the red host connected to switch 1?

Using the previously computed spanning tree, which path will the red host connected to switch 7 use to communicate with the blue host connected to switch 8?



VLAN STP

Using the previously computed spanning tree, which path will the red host connected to switch 7 use to communicate with the red host connected to switch 1?

Solution: 7-4-1

Using the previously computed spanning tree, which path will the red host connected to switch 7 use to communicate with the blue host connected to switch 8?

Solution: Not possible. A host in the blue VLAN cannot directly communicate with a host in the red VLAN. Traffic would have to go over a layer 3 router to reach the other destination.

