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)

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

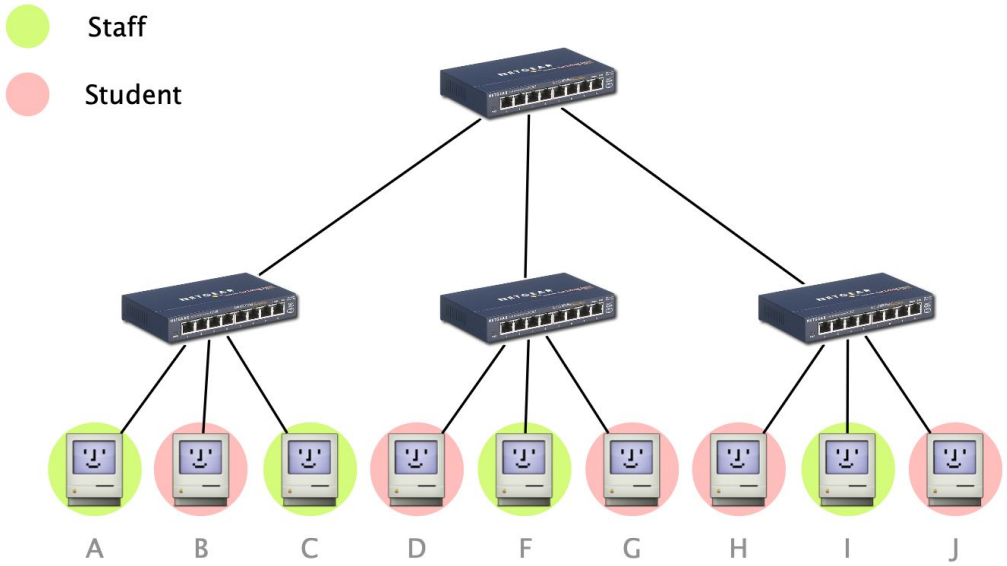

Switches Hold Config Tables Telling them VLAN/Port Info

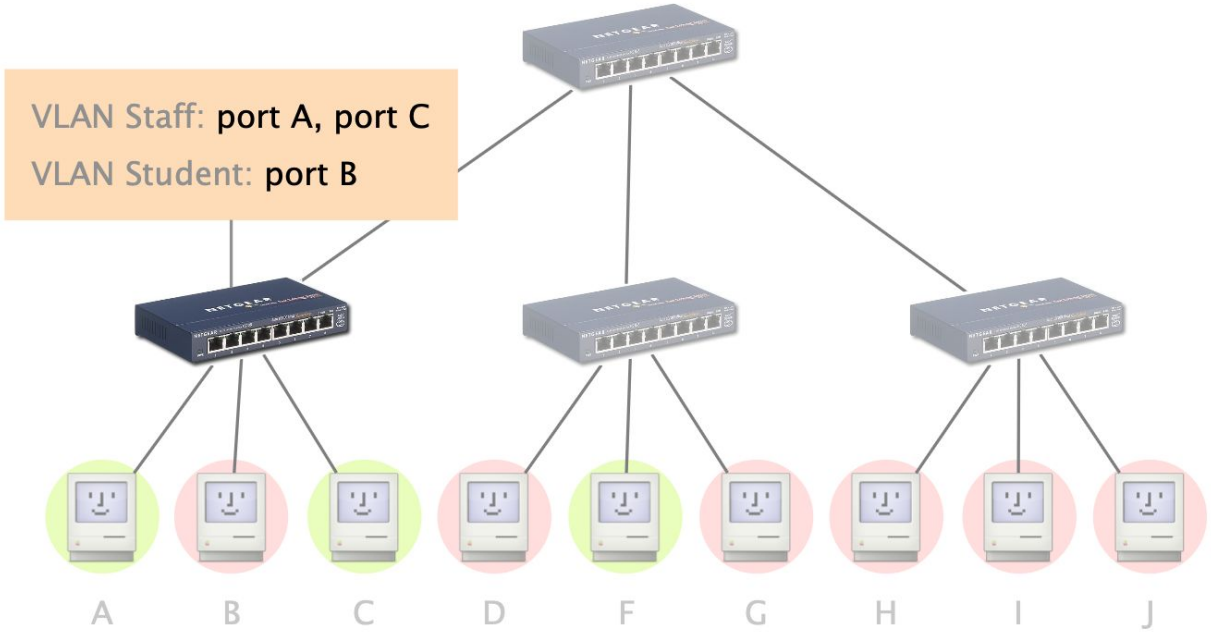

Broadcasts with VLANs - A sends an ARP request

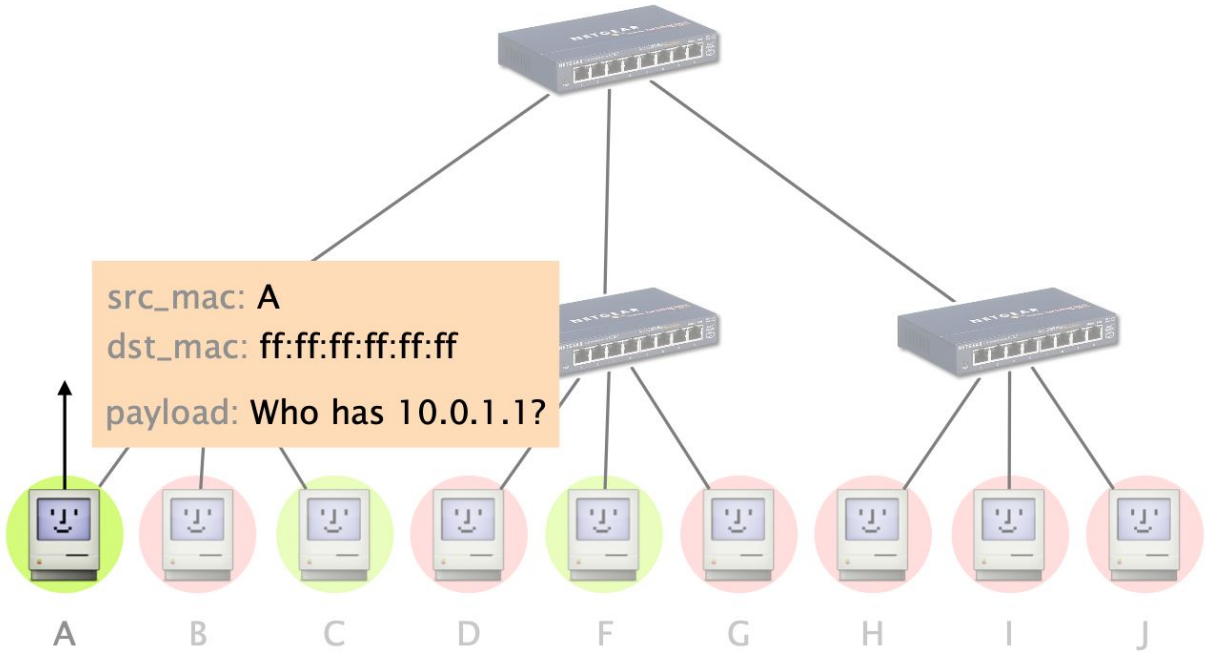

Only Staff VLAN Members (C and F) Should Receive It

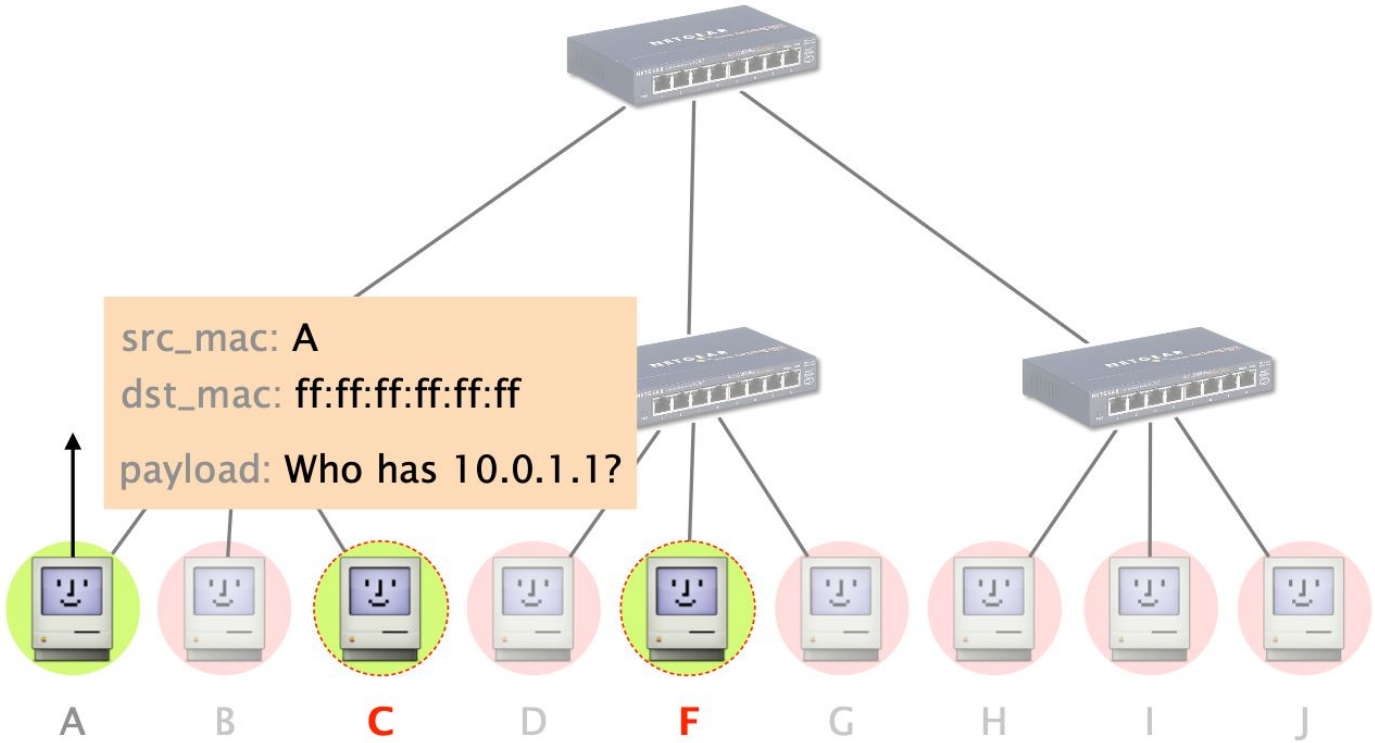

Only Staff VLAN Members (C and F) Should Receive It

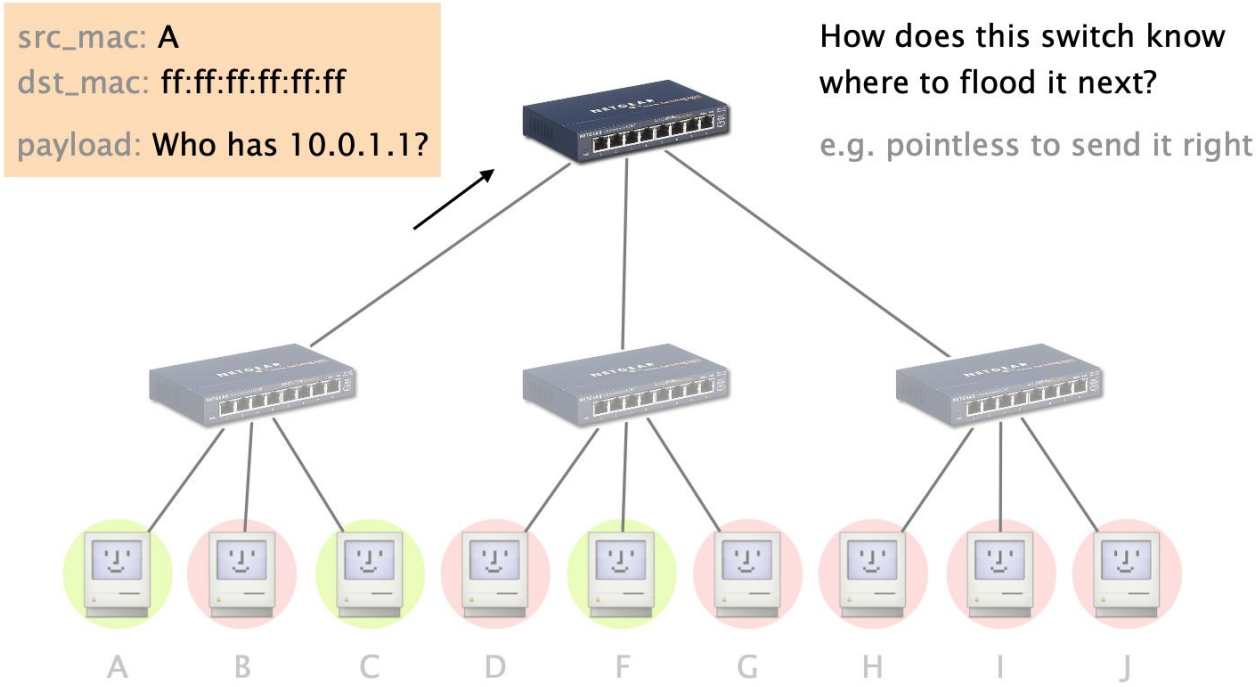

VLAN Tags are Added to the Ethernet Frame

Without VLAN

With VLAN

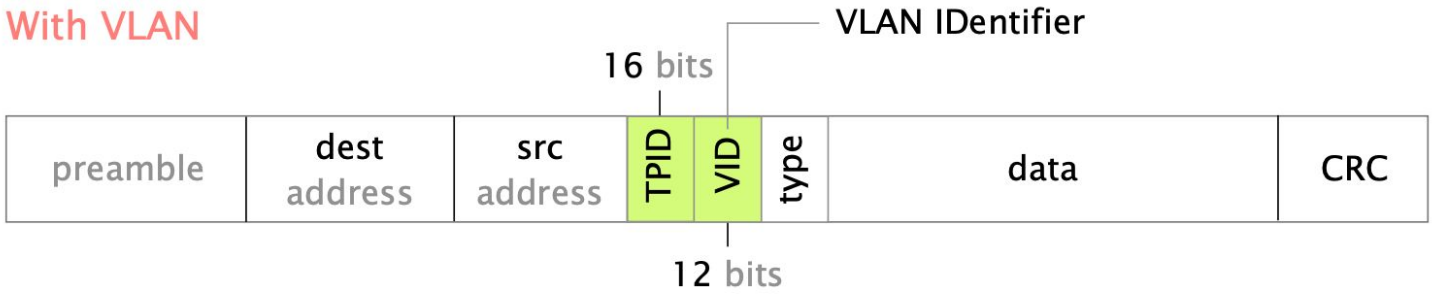

802.1q Header (4 bytes)
(4 bits missing)

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

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

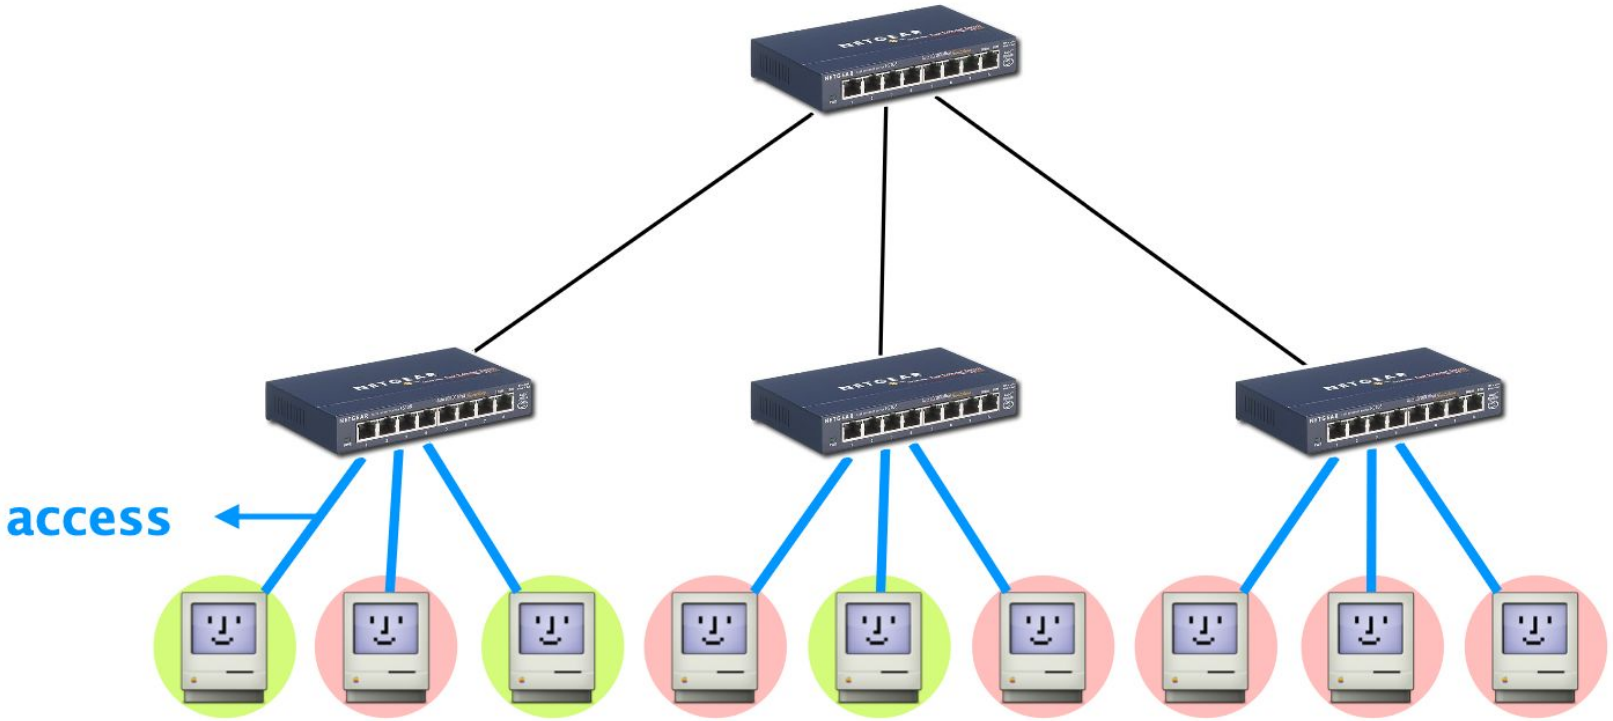

Trunk Links Carry Traffic for More Than One VLAN and Use 802.1q Headers

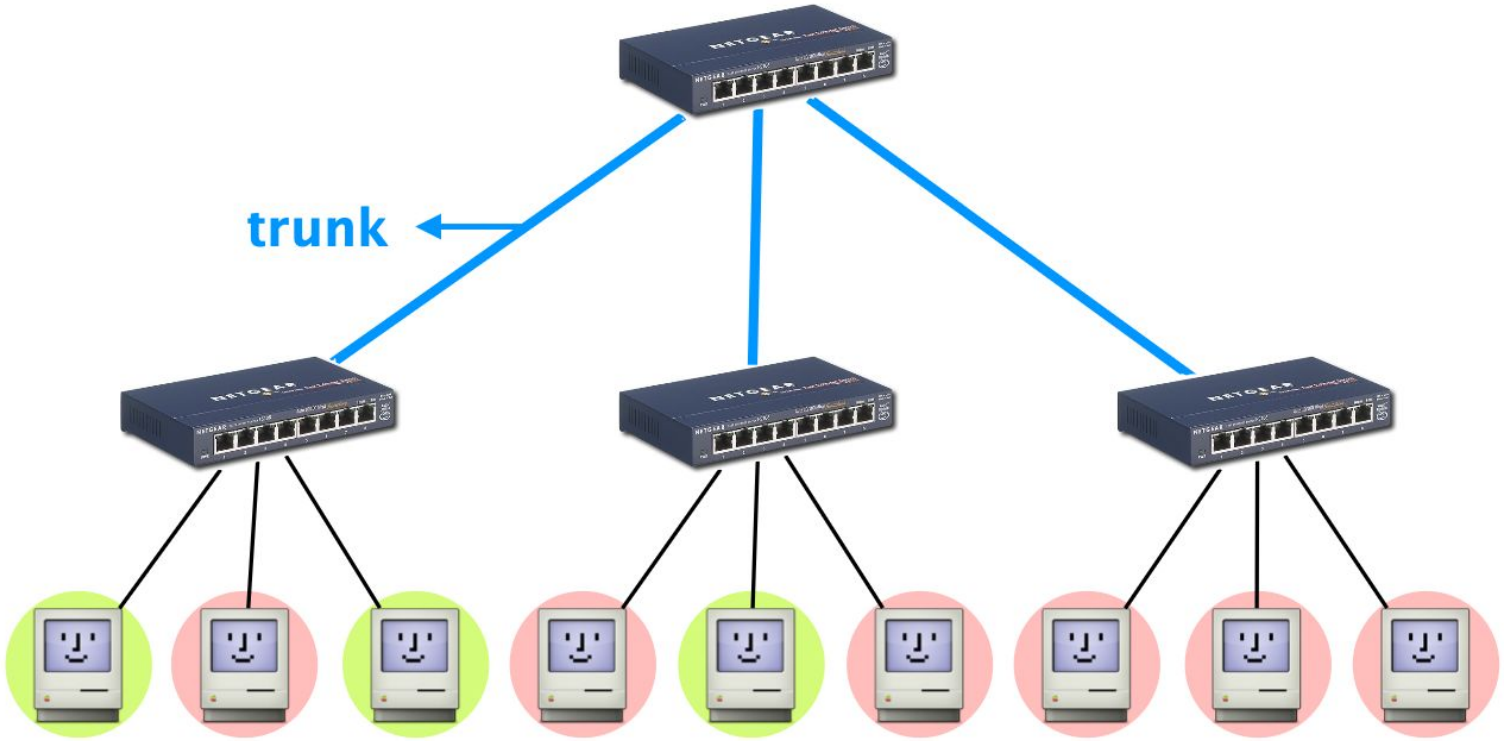

VLAN Learning

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

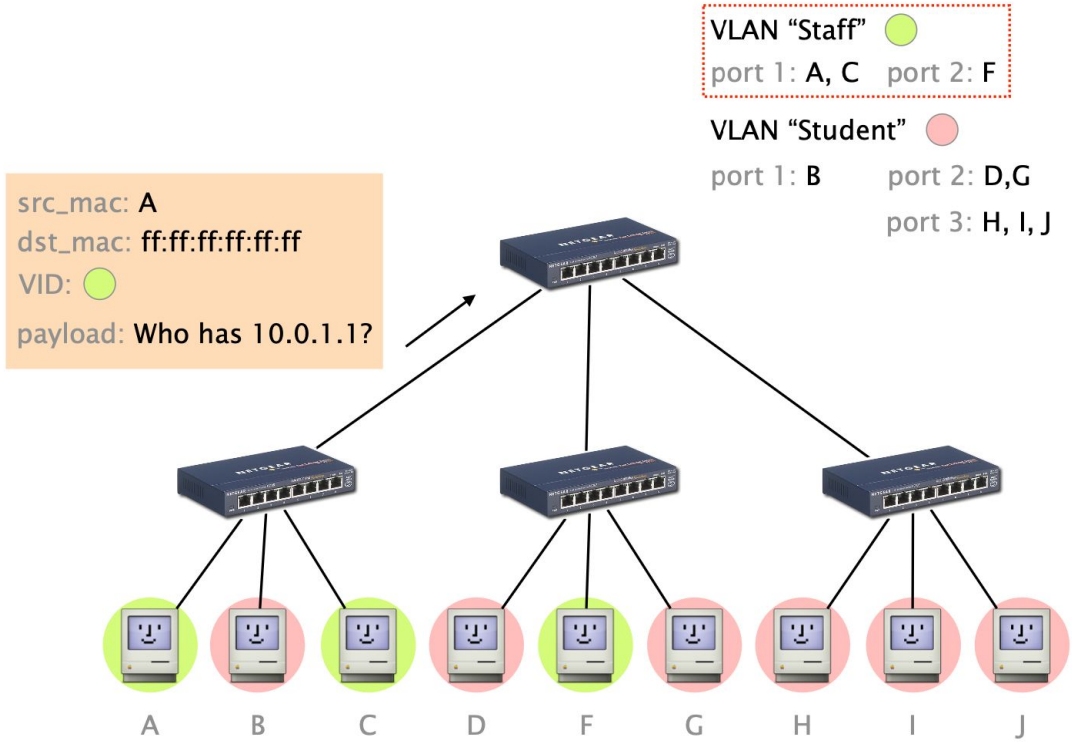

VLAN STP

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

VLAN STP

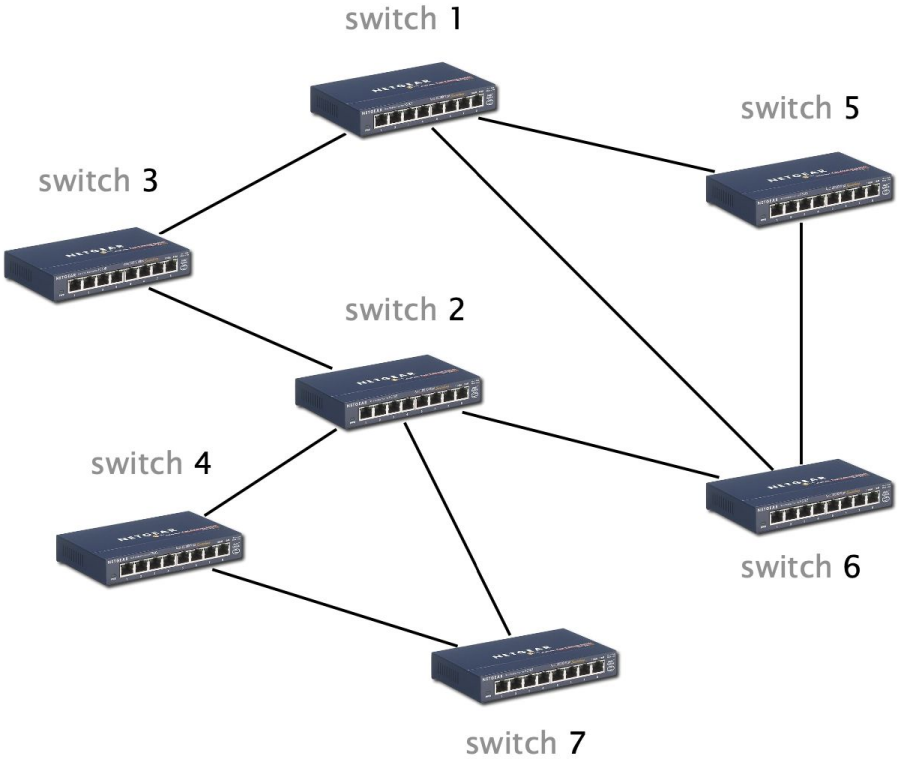

VLAN STP - Non Per-VLAN Setup

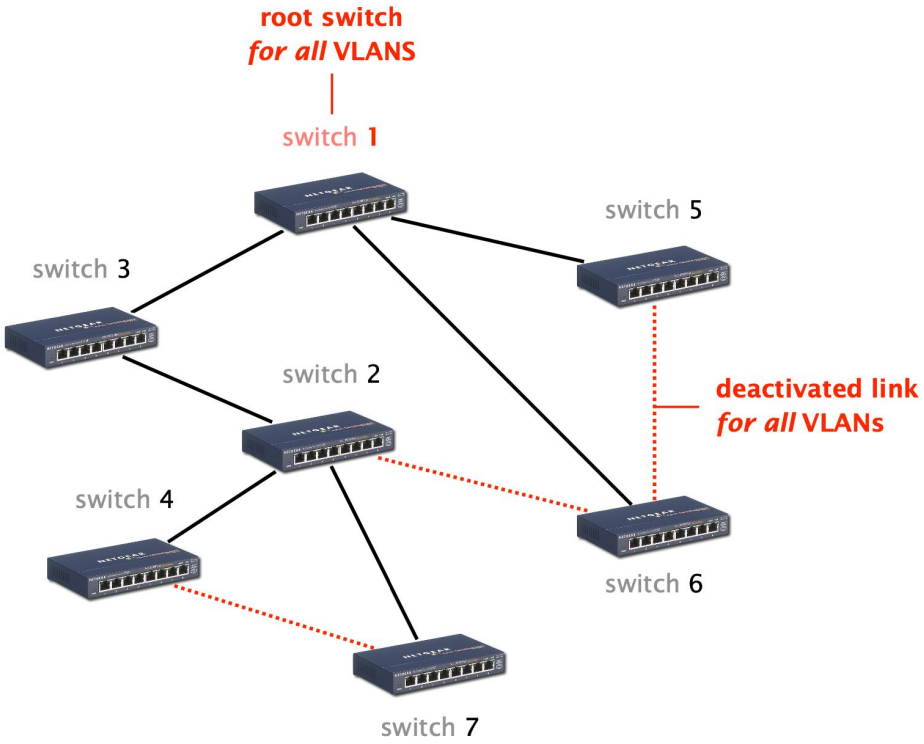

VLAN STP - Non Per-VLAN Setup

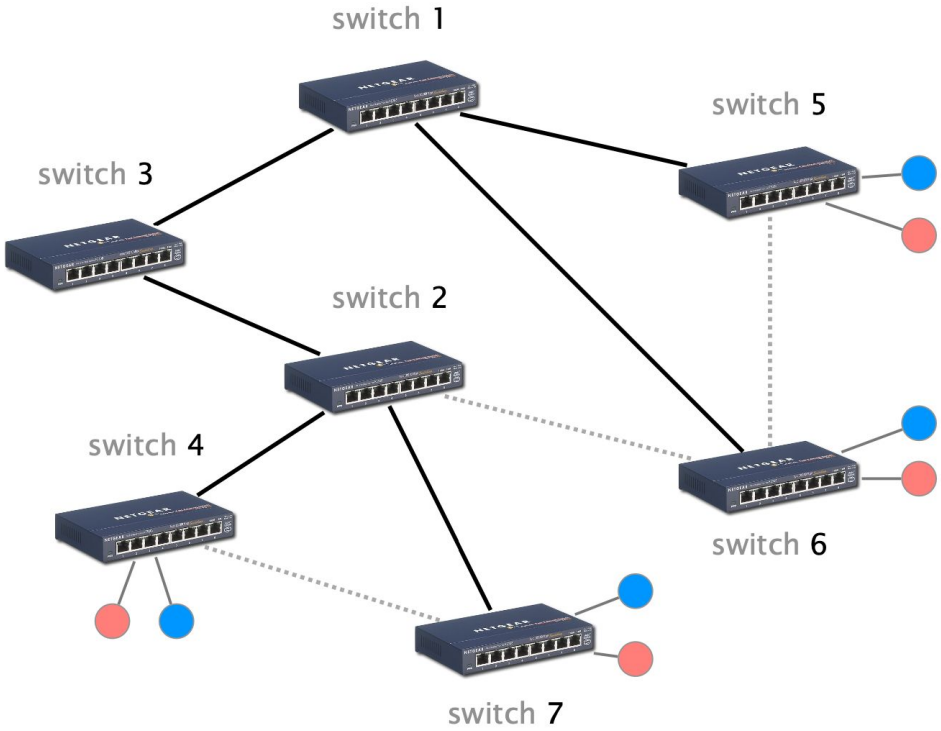

Any Communication Between Red Hosts on 5 and 6 Need to Go Through 1

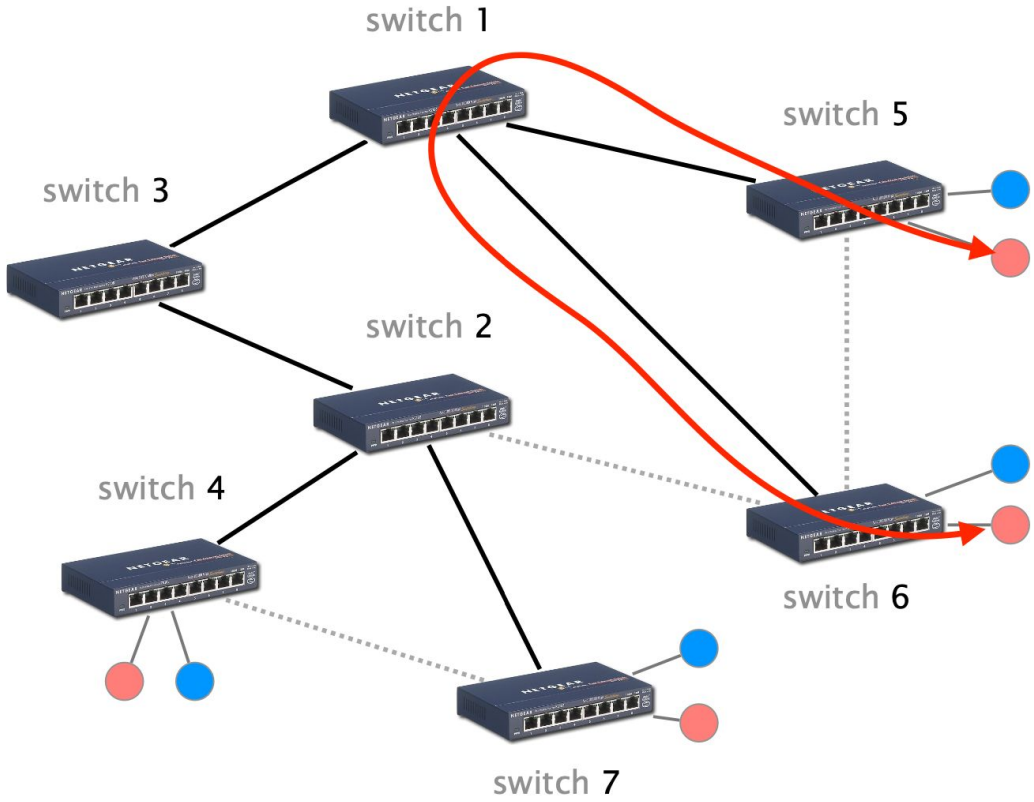

VLAN STP - Per-VLAN STP Setup

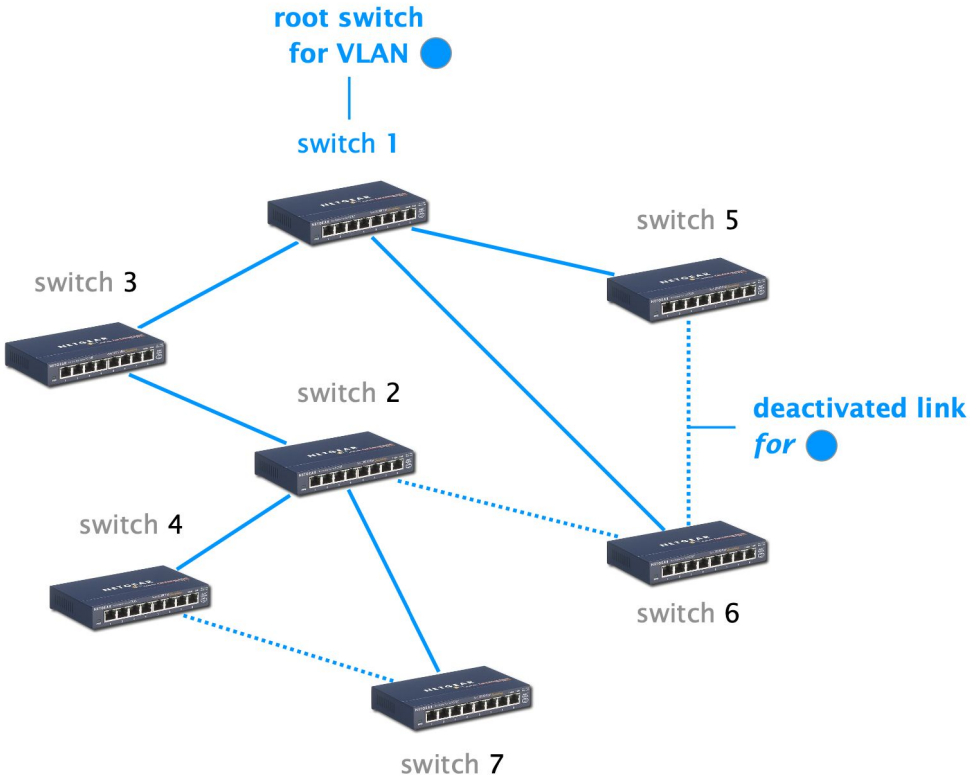

VLAN STP - Per-VLAN STP Setup

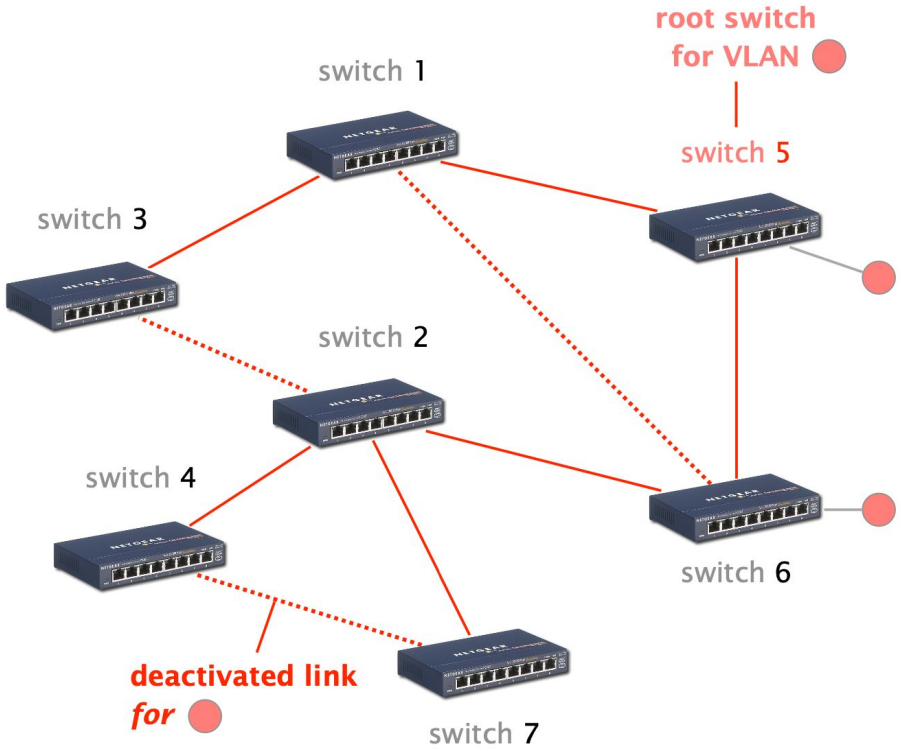

VLAN STP - Now Hosts on 5 and 6 Can Use Direct Link

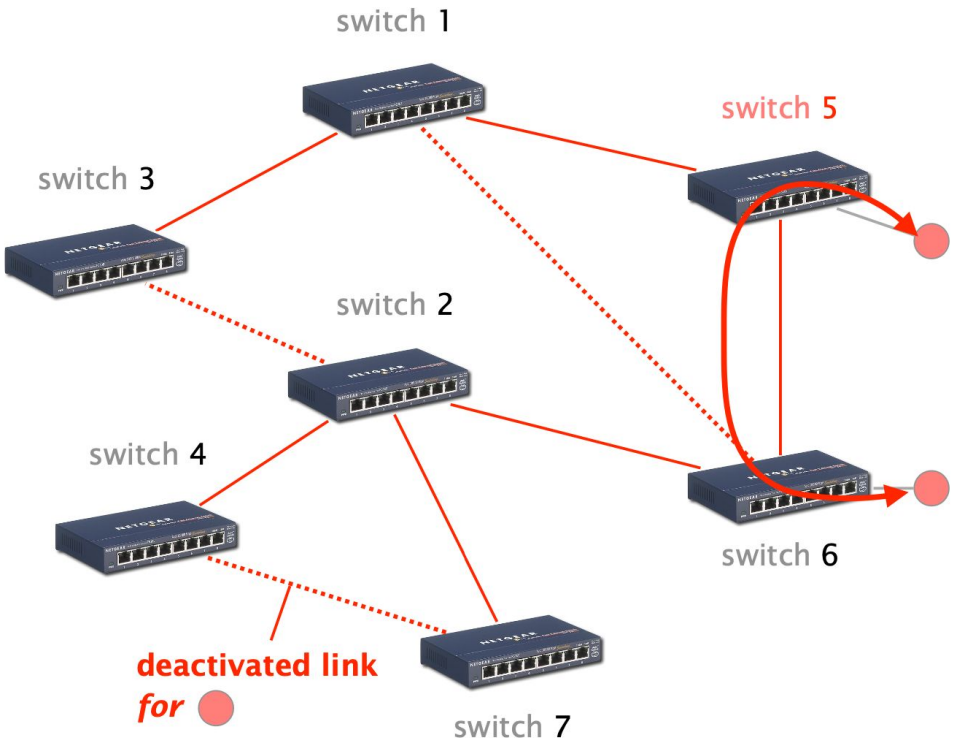

VLAN STP

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).

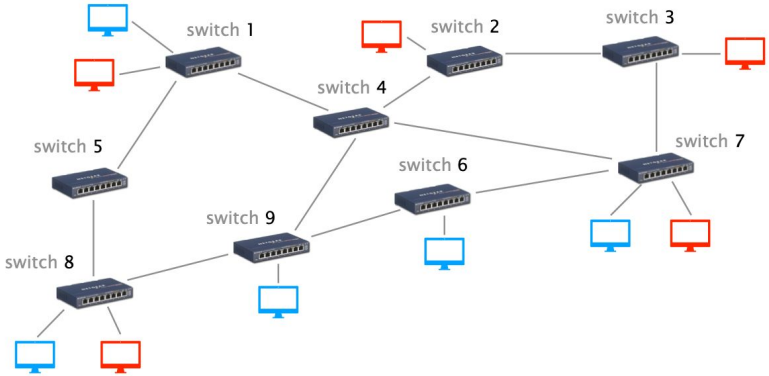

VLAN STP

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).

Solution:

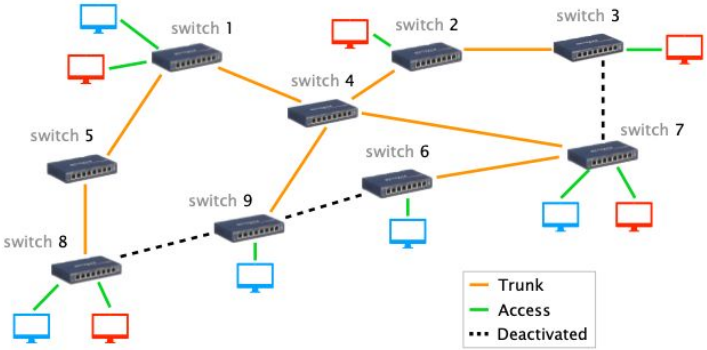

VLAN STP

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?

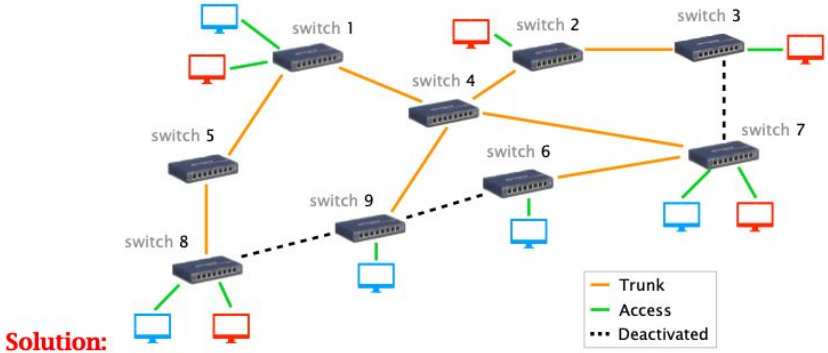

Solution:

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.

Solution:

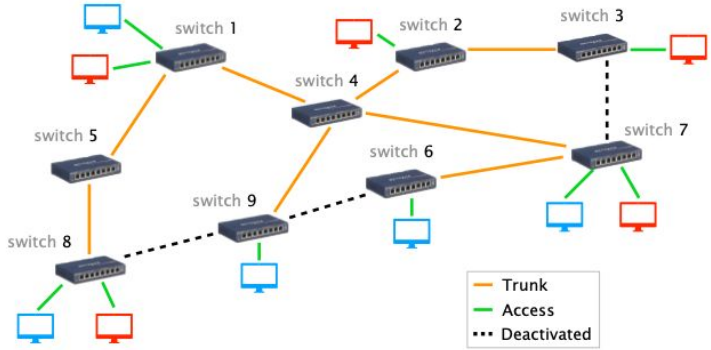

Network Layer

Network Layer - Moving up the Stack

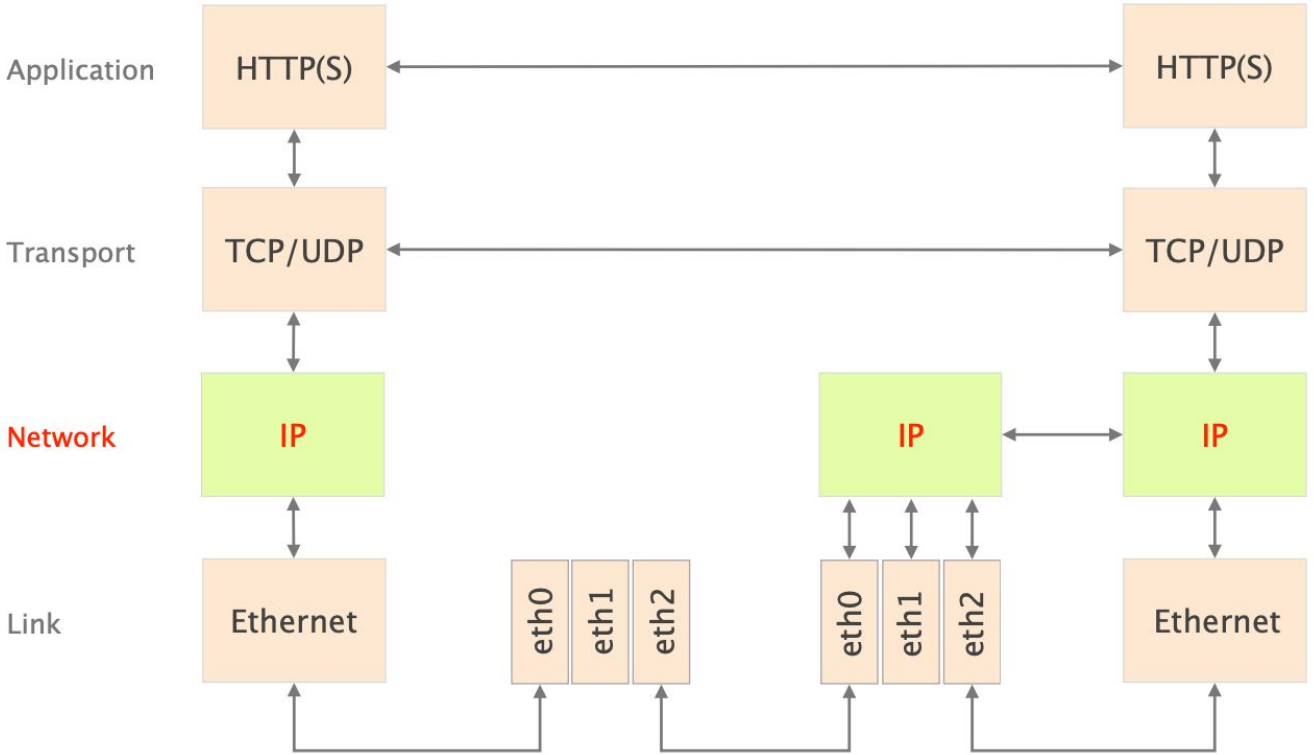

Network (IP) Layer

1. IP addresses
    - use, structure, allocation
2. IP forwarding
    - longest prefix match rule
3. IP header
    - IPv4 and IPv6, wire format

Network (IP) Layer

1. IP addresses
    - use, structure, allocation
2. IP forwarding
    - longest prefix match rule
3. IP header
    - IPv4 and IPv6, wire format

IP(v4) Addresses are Unique 32-bit Numbers Associated with a Host/Router

IP addresses are usually written

using dotted-quad notation

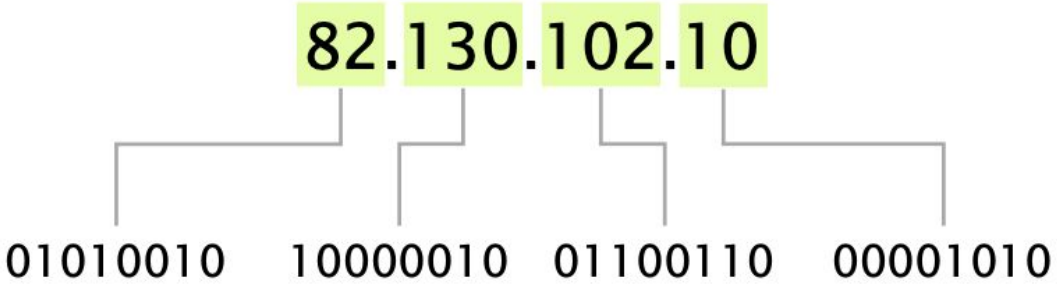

IP(v6) Addresses are Unique 128-bit Numbers Associated with a Host/Router

Notation	8 groups of 16 bits each separated by colons (:) Each group is written as four hexadecimal digits
Simplification	Leading zeros in any group are removed **One** section of zeros is replaced by a double colon (::) Normally the longest section
Examples	1080:0:0:0:8:800:200C:417A → 1080::8:800:200C:417A FF01:0:0:0:0:0:0:0101 → FF01::101 0:0:0:0:0:0:0:1 → ::1

Routers Forward Based on IP Destination

If IPs Were Individually Handed Out, Routers Would Require Massive Forwarding Tables

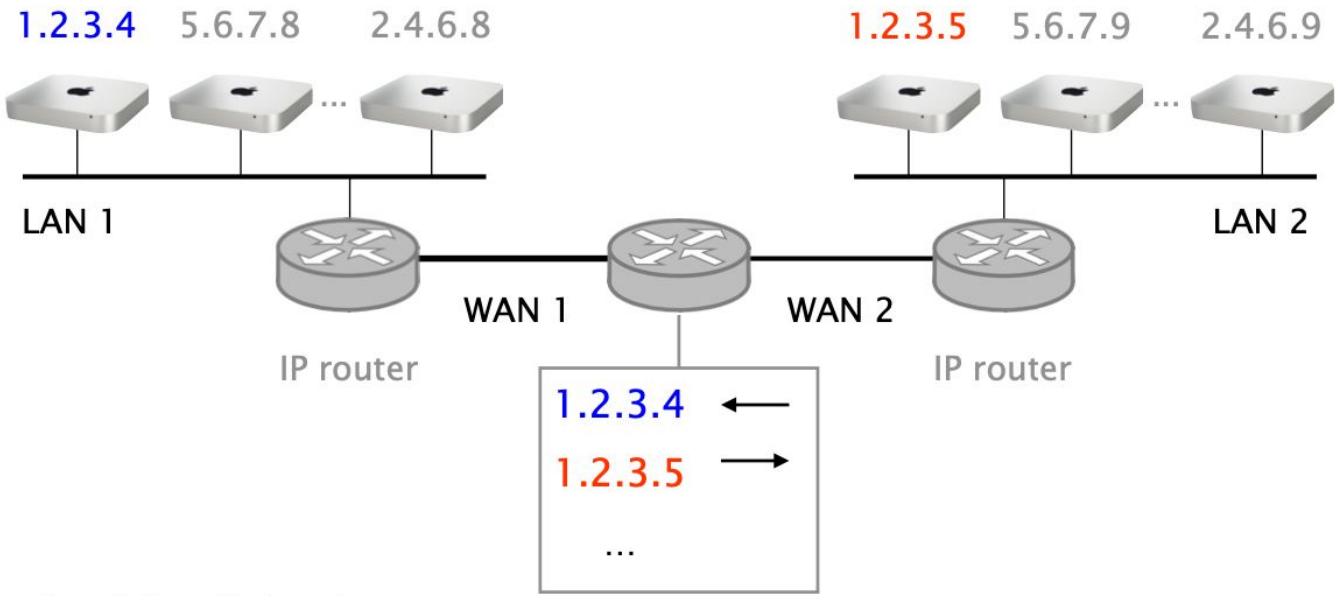

LAN Local Area Network

WAN Wide Area Network

Universal Answers to Such Problems

When you need… more flexibility,

you add… a layer of indirection

When you need… more scalability,

you add… a hierarchical structure

IPs are Hierarchically Allocated

Example: Mail service

- Deliver the letter to the post office responsible for the zip code
- Assign letter to the mail person covering the street
- Drop letter into the mailbox attached to the building
- Hand in the letter to the appropriate person

IPs are Hierarchically, Composed of Prefix (network address) and Suffix (host address)

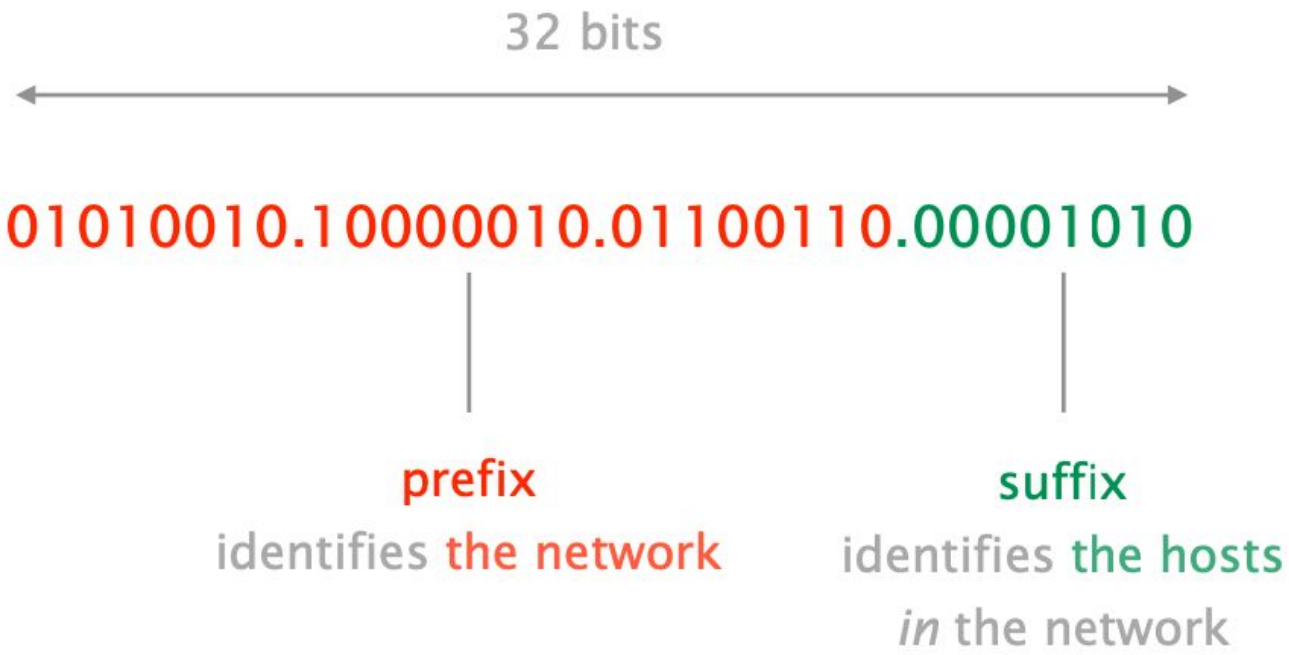

Prefixes Have Varying Lengths, Usually Written Using “slash notation”

/24 Means We Have 8 bits for Host Addresses, enough for 256 Hosts

82.130.102.0 /24

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0
01010010.10000010.01100110.	00000001	82.130.102.1
01010010.10000010.01100110.	00000010	82.130.102.2
01010010.10000010.01100110.	11111110	82.130.102.254
01010010.10000010.01100110.	11111111	82.130.102.255

In Practice the First and Last IP Addresses of a Prefix are not Usable

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0
01010010.10000010.01100110.	11111111	82.130.102.255

In Practice the First and Last IP Addresses of a Prefix are not Usable

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0

All 0s Identifies the Network Itself

01010010.10000010.01100110.	11111111	82.130.102.255

In Practice the First and Last IP Addresses of a Prefix are not Usable

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0

All 1s Identifies the Broadcast Address

01010010.10000010.01100110.	11111111	82.130.102.255