Campus Network Design Workshop

Layer 2 Engineering – VLANs

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

Virtual LANs (VLANs)

- Allow us to split switches into separate (virtual) switches
- Only members of a VLAN can see that VLAN's traffic
  - Inter-vlan traffic must go through a router
- Allow us to reuse router interfaces to carry traffic for separate subnets
  - E.g. sub-interfaces in Cisco routers

Local VLANs

- 2 VLANs or more within a single switch
- ***Edge ports***, where end nodes are connected, are configured as members of a VLAN
- The switch behaves as several virtual switches, sending traffic only within VLAN members

Local VLANs

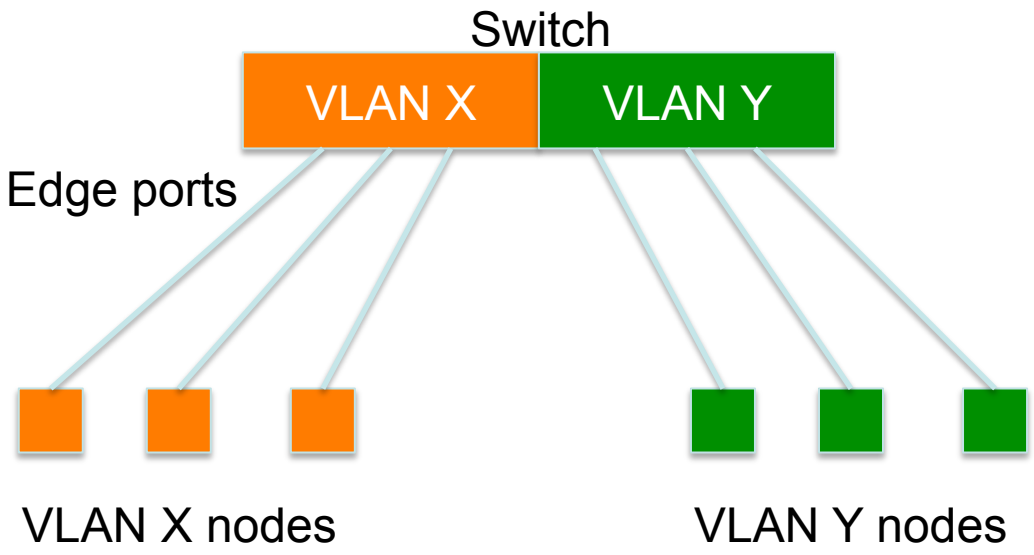

VLANs across switches

- Two switches can exchange traffic from one or more VLANs
- Inter-switch links are configured as ***trunks***, carrying frames from all or a subset of a switch's VLANs
- Each frame carries a ***tag*** that identifies which VLAN it belongs to

802.1Q

- The IEEE standard that defines how ethernet frames should be ***tagged*** when moving across switch trunks
- This means that switches from *different vendors* are able to exchange VLAN traffic.

802.1Q tagged frame

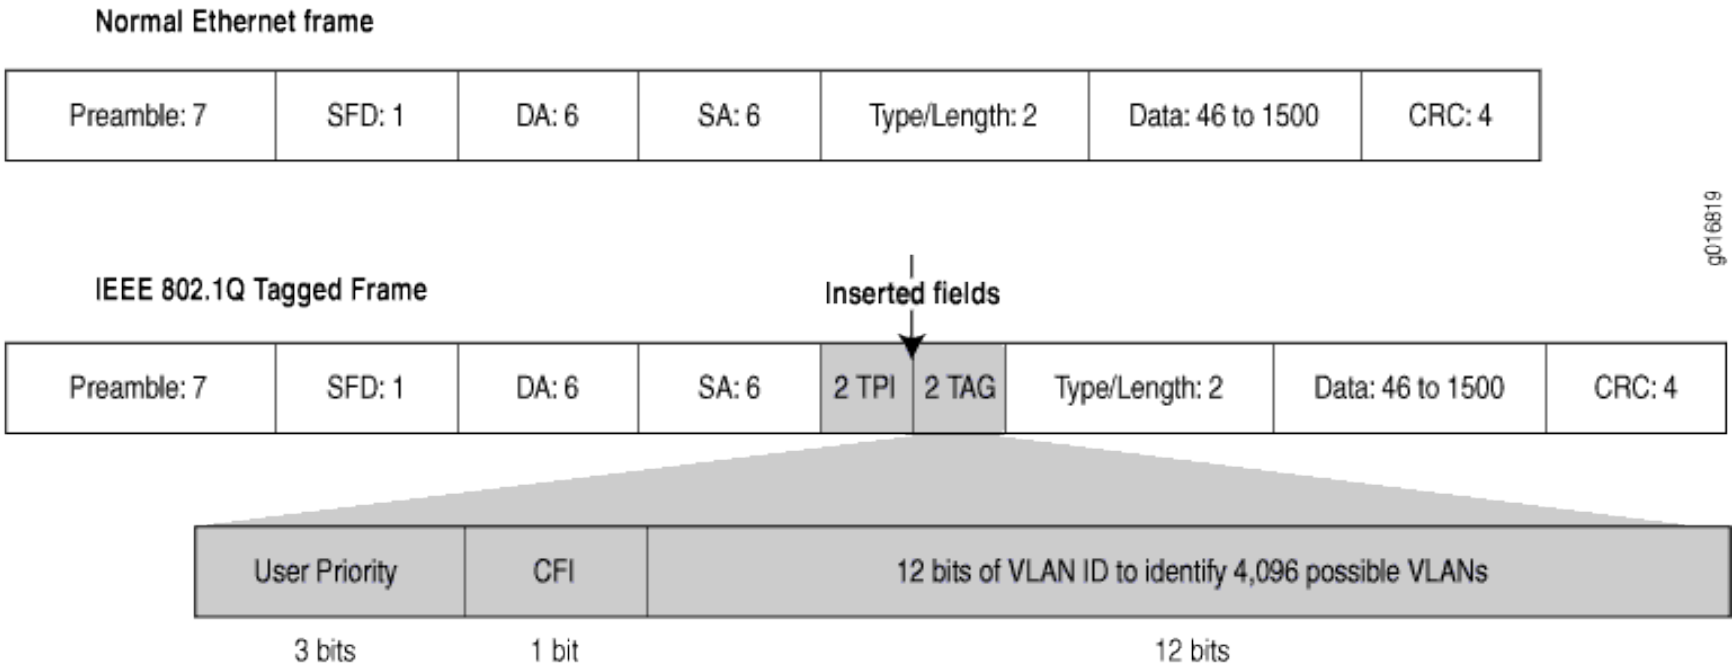

VLANs across switches

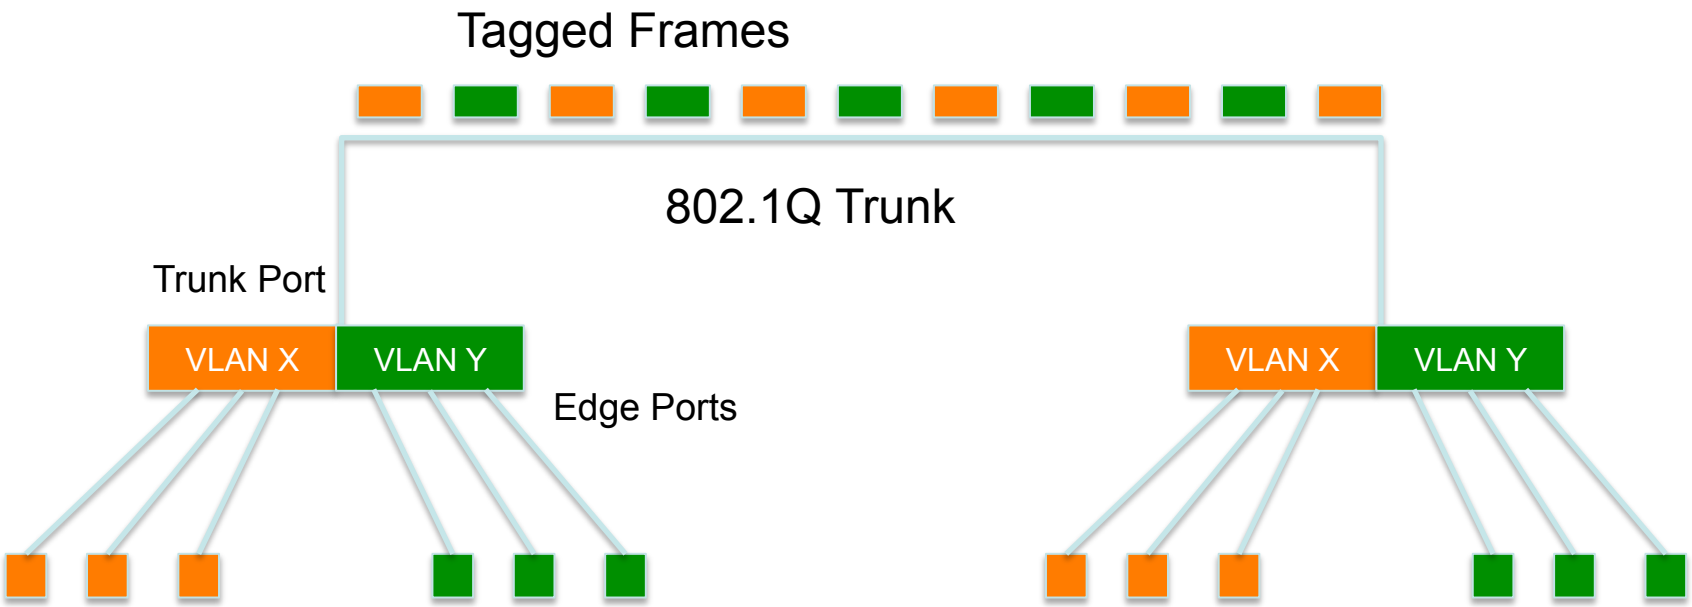

This is called “VLAN Trunking”

Tagged vs. Untagged

- Edge ports are not tagged, they are just "members" of a VLAN
- You only need to tag frames in switch-to-switch links (trunks), when transporting multiple VLANs
- A trunk can transport both tagged and untagged VLANs
  - As long as the two switches agree on how to handle those

VLANs increase complexity

- You can no longer “just replace” a switch
  - Now you have VLAN configuration to maintain
  - Field technicians need more skills
- You have to make sure that all the switch-to-switch trunks are carrying all the necessary VLANs
  - Need to keep in mind when adding/removing VLANs

Good reasons to use VLANs

- You want multiple subnets in a building, and carry them over a single fibre to your core router
- You want to segment your network into multiple subnets, without buying more switches
  - Separate broadcast domains for wired, wireless, phones, device management etc.
- Separate control traffic from user traffic
  - Restrict who can access your switch management address

Bad reasons to use VLANs

- Because you can, and you feel cool ☺
- Because they will completely secure your hosts (or so you think)
- Because they allow you to extend the same IP network over multiple separate buildings
  - This is actually very common, but a bad idea

Do not build “VLAN spaghetti”

- Extending a VLAN to multiple buildings across trunk ports
- Bad idea because:
  - Broadcast traffic is carried across all trunks from one end of the network to another
  - Broadcast storm can spread across the extent of the VLAN, and affect all VLANS!
  - Maintenance and troubleshooting nightmare

Cisco configuration

- Configure access port
  - `interface GigabitEthernet1/0/3`
    `switchport mode access`
    `switchport access vlan 10`
- Configure trunk port
  - `interface GigabitEthernet1/0/1`
    `switchport mode trunk`
    `switchport trunk allowed vlan 10,20,30`

Cisco mis-features

- *Disable* VLAN Trunking Protocol (VTP)
  - `vtp mode transparent`
- *Disable* Dynamic Trunking Protocol (DTP)
  - `interface range Gi 1 - 8`
    `switchport mode [trunk|access]`
    `switchport nonegotiate`

HP configuration

- Configure access ports
  - `vlan 10`
    `untagged 3,5-7,12`
- Configure trunk ports
  - `vlan 10`
    `tagged 1-2`
    `vlan 20`
    `tagged 1-2`
    `vlan 30`
    `tagged 1-2`

Questions?

Link Aggregation

- Also known as *port bundling, link bundling*
- You can use multiple links in parallel as a single, logical link
  - For increased capacity
  - For redundancy (fault tolerance)
- LACP (Link Aggregation Control Protocol) is a standardized method of negotiating these bundled links between switches
- Proprietary methods exist too (PAgP, EtherChannel)

LACP Operation

- Two switches connected via multiple links will send LACPDU packets, identifying themselves and the port capabilities
- They will then automatically build the logical aggregated links, and then pass traffic.
- Switch ports can be configured as active or passive

LACP Operation

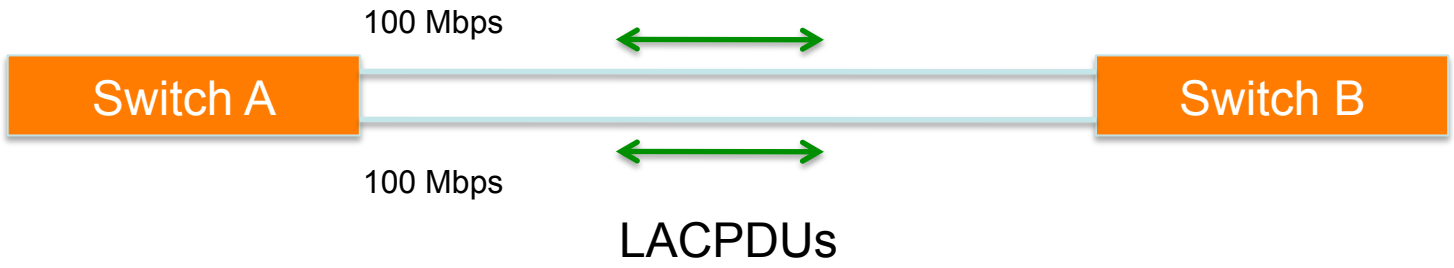

- Switches Switches A and B are connected to each other using two sets of Fast Ethernet ports
- LACP is enabled and the ports are turned on
- Switches start sending LACPDUs, then negotiate how to set up the aggregation

LACP Operation

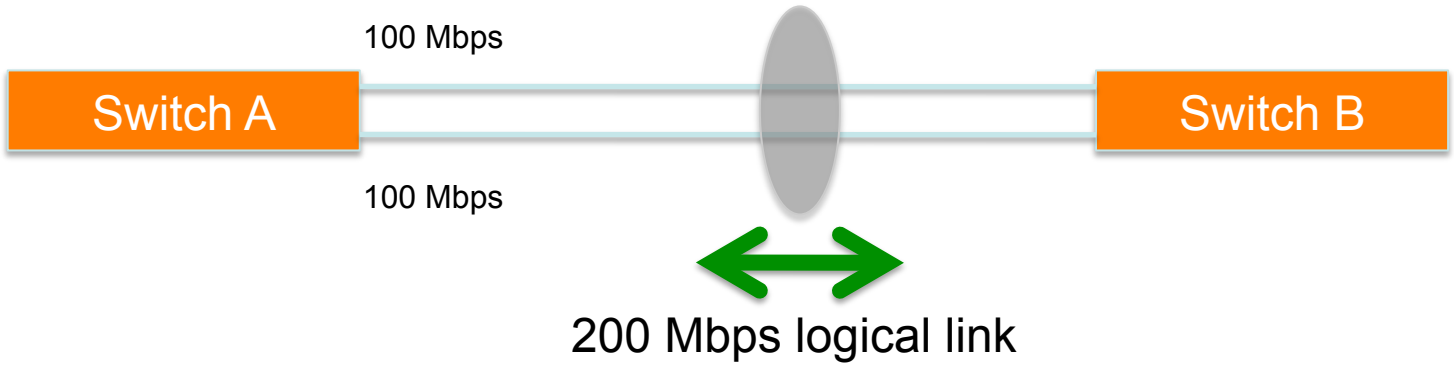

- The result is an aggregated 200 Mbps logical link
- The link is also fault tolerant: If one of the member links fail, LACP will automatically take that link off the bundle, and keep sending traffic over the remaining link

Distributing Traffic in Bundled Links

- Bundled links distribute frames using a hashing algorithm, based on:
  - Source and/or Destination MAC address
  - Source and/or Destination IP address
  - Source and/or Destination Port numbers
- This can lead to unbalanced use of the links, depending on the nature of the traffic
- Always choose the load-balancing method that provides the most distribution

Questions?

Selecting Switches

- Minimum features:
  - Standards compliance
  - Encrypted management (SSH/HTTPS)
  - VLAN trunking
  - Spanning Tree (RSTP at least)
  - SNMP
    - At least v2 (v3 has better security)
    - Traps
  - Remote management and config backup
    - CLI preferred

Selecting Switches

- Other recommended features:
  - DHCP Snooping
    - Prevent end-users from running a rogue DHCP server
      - Happens a lot with little wireless routers (Netgear, Linksys, etc) plugged in backwards
    - Uplink ports towards the legitimate DHCP server are defined as “trusted”. If DHCPOFFERs are seen coming from any untrusted port, they are dropped.

Selecting Switches

- Other recommended features:
  - Dynamic ARP inspection
    - A malicious host can perform a man-in-the-middle attack by sending gratuitous ARP responses, or responding to requests with bogus information
    - Switches can look inside ARP packets and discard gratuitous and invalid ARP packets.

Selecting Switches

- Other recommended features:
  - IGMP Snooping:
    - Switches normally flood multicast frames out every port
    - Snooping on IGMP traffic, the switch can learn which stations are members of a multicast group, thus forwarding multicast frames only out necessary ports
    - Very important when users run Norton Ghost, for example.

Network Management

- Enable SNMP traps and/or syslog
  - Collect and process in centralized log server
    - Spanning Tree Changes
    - Duplex mismatches
    - Wiring problems
- Monitor configurations
  - Use RANCID to report any changes in the switch configuration

Network Management

- Collect forwarding tables with SNMP
  - Allows you to find a MAC address in your network quickly
  - You can use simple text files + grep, or a web tool with DB backend
- Enable LLDP (or CDP or similar)
  - Shows how switches are connected to each other and to other network devices

Documentation

- Document where your switches are located
  - Name switch after building name
    - E.g. building1-sw1
  - Keep files with physical location
    - Floor, closet number, etc.
- Document your edge port connections
  - Room number, jack number, server name

Questions?

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center