Switching Architectures: VLANs

Campus Network Design & Operations Workshop

UNIVERSITY OF OREGON

Last updated 1st October 2024

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

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
  - Using sub-interfaces in Cisco routers
  - Using IRB interfaces in Juniper routers

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
  - Using sub-interfaces in Cisco routers
  - Using IRB interfaces in Juniper routers
- VLANs are also useful in servers especially with virtualization
  - Virtual Machines (VMs) for different networks (public vs private or student vs admin) can exist on the same virtualization host

Local VLANs

- Two or more VLANs within a single switch
- The switch behaves as several virtual switches, sending traffic only within VLAN members
- ***Access ports***, where end nodes are connected, are configured as members of a VLAN
- By default, all ports of a switch are members of VLAN 1 or default VLAN (***VLAN ID*** = 1)
- Newly created VLANs must have a VLAN ID other than 1
  - Then add ports by moving them out of VLAN 1 into our new VLAN

Local VLANs

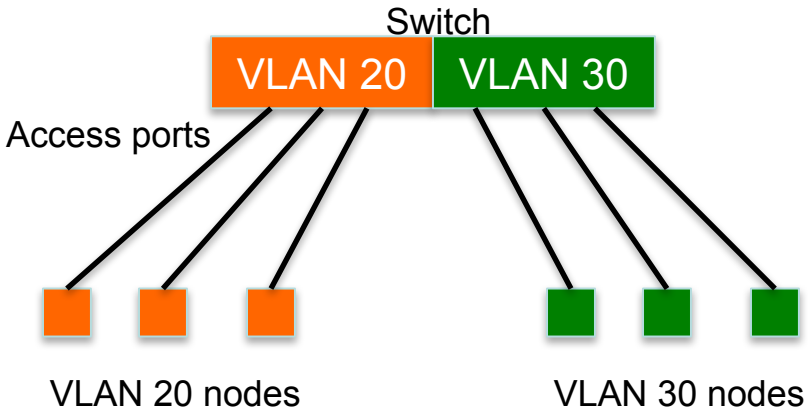
Switch
VLAN 20
VLAN 30
Access ports
VLAN 20 nodes
VLAN 30 nodes

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

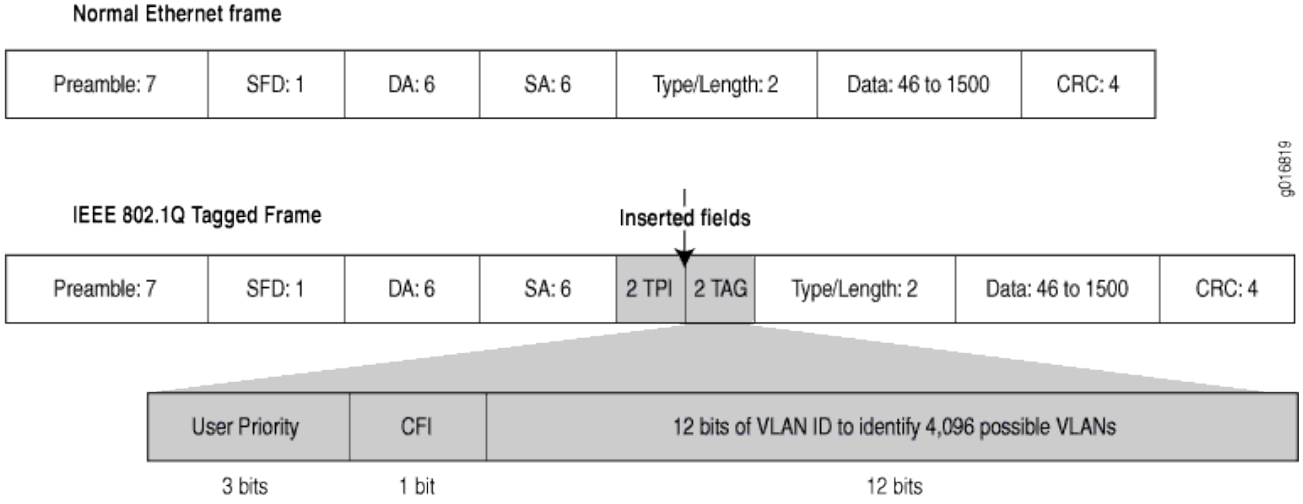

TPI is the Tag Protocol Identifier, set to 0x8100 to indicate 802.1Q

VLANs across switches

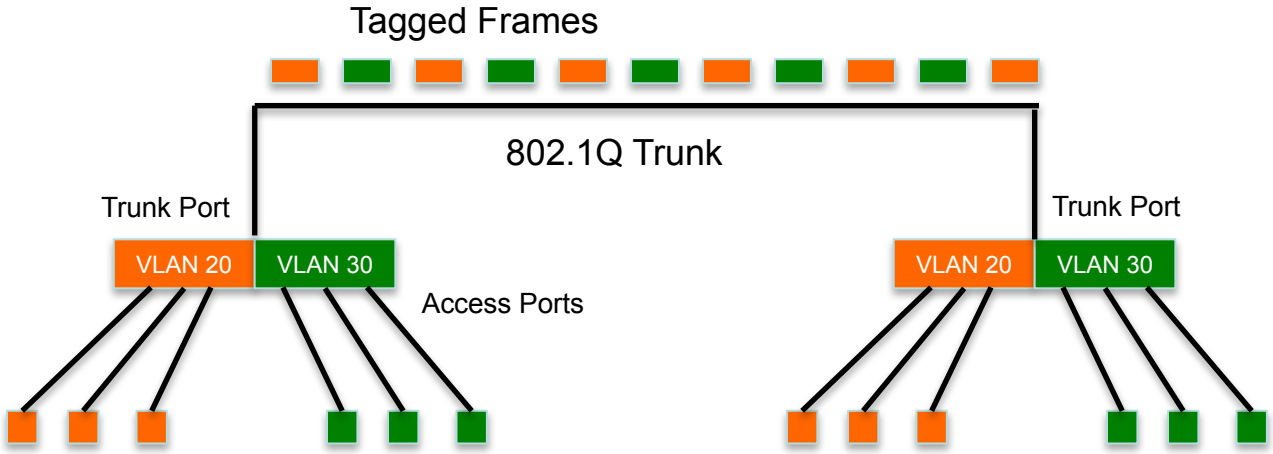

This is called “VLAN Trunking”

Tagged vs. Untagged

- Frames sent and received on access ports are **not tagged**
- You only need to tag frames in switch-to-switch links (trunks), when transporting multiple VLANs
- However, a trunk **can** transport both tagged and untagged frames
  - As long as the two switches agree on how to handle untagged frames
  - Only **one** VLAN can be untagged ("native") on a given link
  - Usually best avoided, but there are some cases where this is useful

Routing Inter-VLAN traffic

Traffic between VLANs must now go through a router.

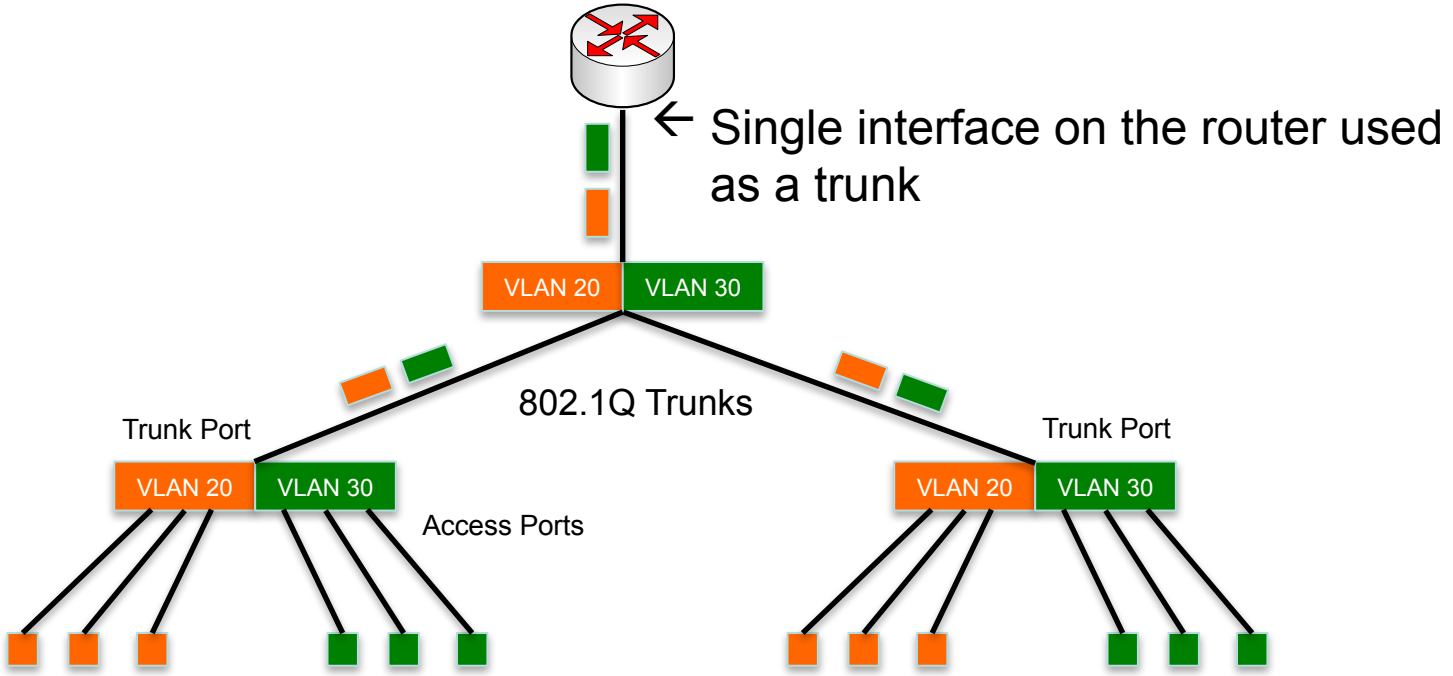

Routing Inter-VLAN traffic (2)

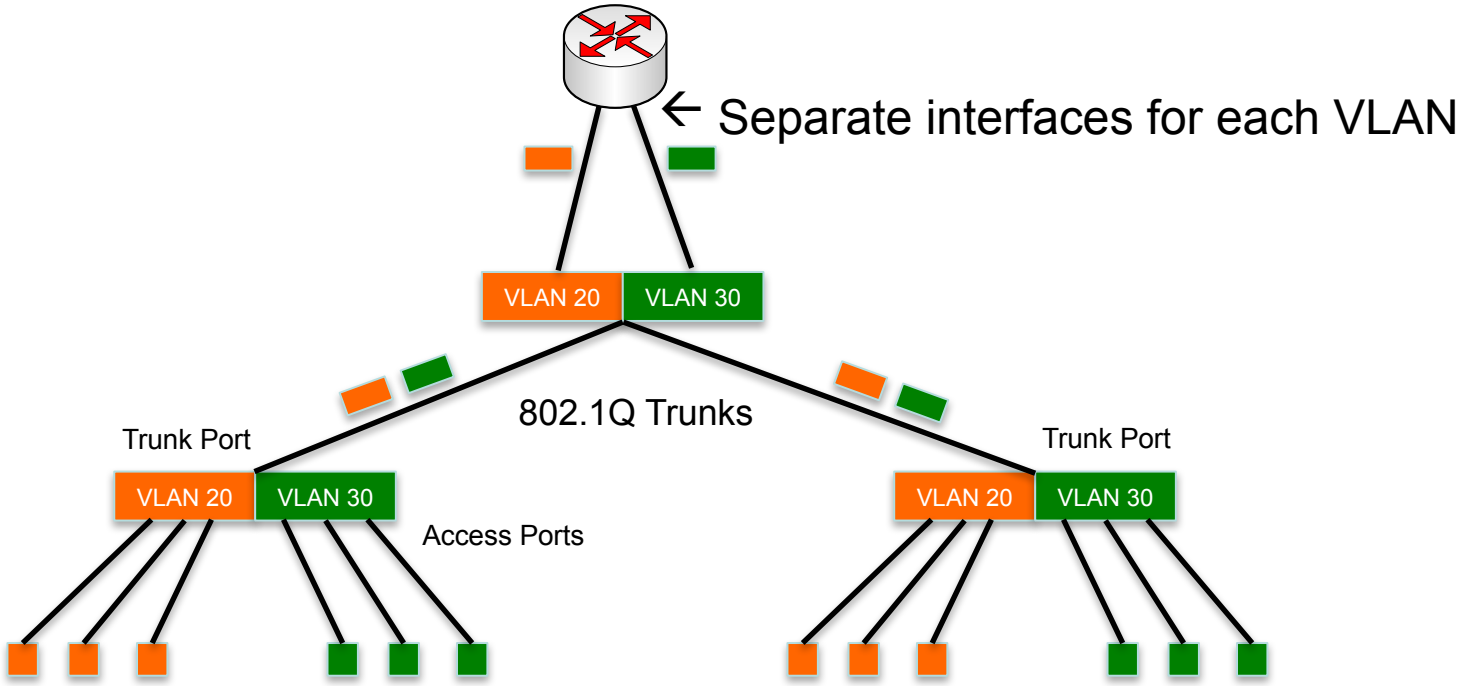

Routing Inter-VLAN traffic (3)

Or we can use an 802.1Q compliant Layer-3 switch to do switching as well routing

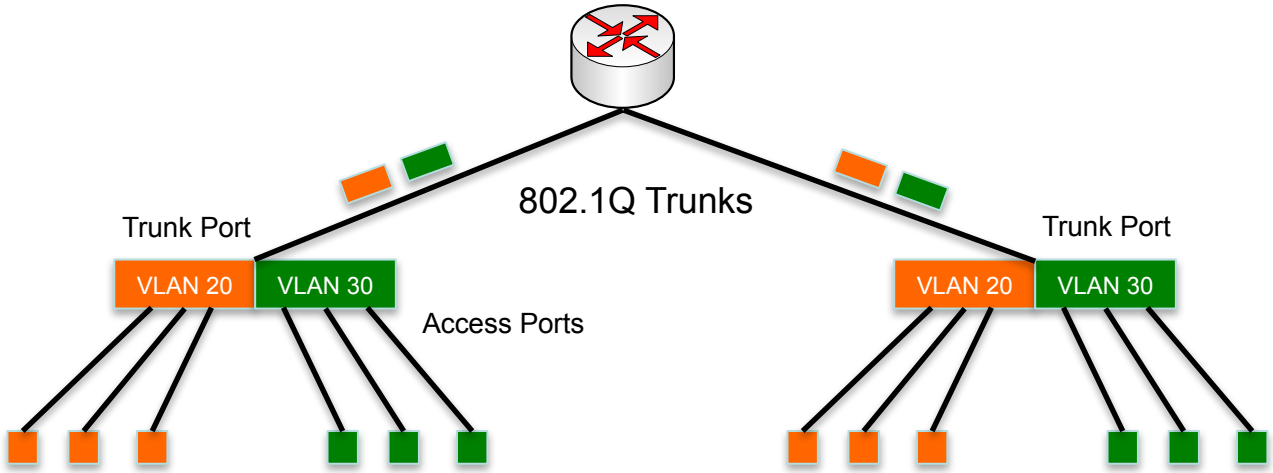

VLANs increase complexity

- 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 configured to carry frames of all the necessary VLANs
  - Need to keep in mind when adding/removing VLANs

Good reasons to use VLANs

1. You want multiple subnets in a building, and carry them over a single fibre to your core router

2. You want to segment your network into multiple subnets, without buying more switches
   - Separate broadcast domains for wired, wireless, phones, device management etc.

3. Separate control traffic from user traffic
   - Restrict who can access your switch management address

Bad reasons to use VLANs

1. Because you can, and you feel cool 😉

2. Because they will completely secure your hosts (or so you think)

3. 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”

- VLAN “spaghetti” means extending a VLAN to multiple buildings across trunk ports

Do not build “VLAN spaghetti”

- VLAN “spaghetti” means 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 IOS VLAN 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 IOS mis-features

- *Disable* VLAN Trunking Protocol (VTP)

```
vtp mode off
```

or

```
vtp mode transparent
```

- *Disable* Dynamic Trunking Protocol (DTP)

```
interface range Gi 1/0/1 - 8
 switchport mode [trunk|access]
 switchport nonegotiate
```

Cisco NX-OS differences

- Cisco's Nexus switches run NX-OS not IOS and are Layer-3 switches which could make a good choice for a core router.
- They can be set so that the configuration assumes most ports will be used as routed Layer-3 ports (our preference) or as switched Layer-2 ports.
  - the default can be set:

```
no system default switchport
```

  - then any ports designated to be switched (Layer-2) are set explicitly before applying access or trunk configuration as shown earlier.

```
interface Ethernet1/10
  switchport
  switchport access vlan 10
```

Cisco NX-OS differences(2)

- Support for VTP commands and the protocol is disabled by default.
  - If your Nexus has vtp commands, you can turn it off altogether

```
no feature vtp
```

- There is fortunately no support for DTP even as an option.
- Only VLAN trunking protocol available is standards based 802.1q.
- The VLANs NX-OS reserves for internal use are different from those reserved in IOS

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
```

Juniper configuration

- Configure access ports

```
set interfaces ge-0/0/1.0 family bridge interface-mode access vlan-id 10
```

- Configure trunk ports

```
set interfaces ge-0/0/2.0 family bridge interface-mode trunk vlan-id-list [ 10 20 30 ]
```

Netgear configuration (weird)

- Configure access ports

```
interface g3
vlan pvid 10
vlan acceptframe admituntaggedonly
vlan participation include 10
```

*Incoming **untagged** frames are assigned to this VLAN*

- Configure trunk ports

```
interface g1
vlan acceptframe vlanonly
vlan participation include 10,20,30
vlan tagging 10,20,30
```

Allow outgoing frames from these VLANs

Add tags when sending these VLANs

Optional: "vlan participation exclude 1" to remove the default vlan
To remove "vlan participation [include|exclude] X" you write "vlan participation auto X"

Linux Server (using netplan)

```
network:
  version: 2
  renderer: networkd
  ethernets:
    enp0s25:
      dhcp4: false
  vlans:
    vlan10:
      id: 10
      link: enp0s25
      dhcp4: false
      addresses: [1.2.3.4/24]
```

Questions?