Layer 3 Switches

Campus Network Design & Operations Workshop

UNIVERSITY OF OREGON

Last updated 25th October 2022

NSRC
Network Startup Resource Center

What's a Layer 3 switch?

- It's an Ethernet switch!
  - Can look at Ethernet headers
  - Builds MAC address table
- And it's a router!
  - Can look at IP headers
  - Has IP forwarding table and ARP table
- Which function it performs depend on how you configure it
- Out-of-the-box it will default to a simple L2 Ethernet switch

Factory Default

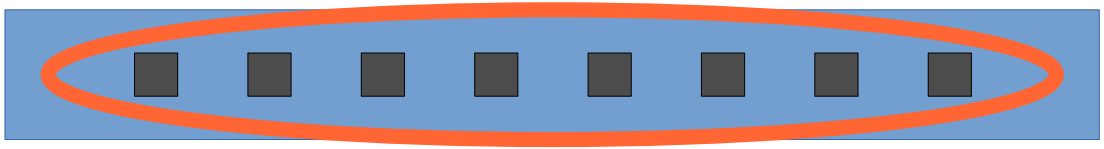

```
vlan 1
!
interface range Gi 1 - 8
  no shutdown
  switchport
  switchport mode access
  switchport access vlan 1
!
```

VLANs

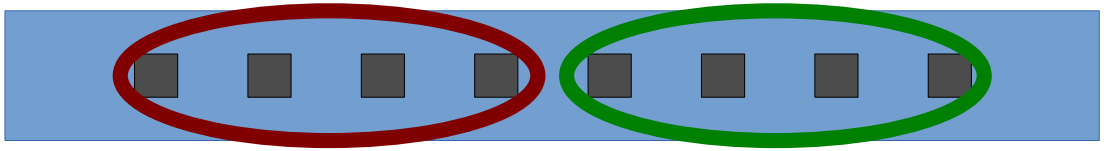

```
vlan 10,20

interface range Gi 1 - 4
  switchport mode access
  switchport access vlan 10

interface range Gi 5 - 8
  switchport mode access
  switchport access vlan 20
```

Question: how does the device behave differently after this config change?

VLAN trunking

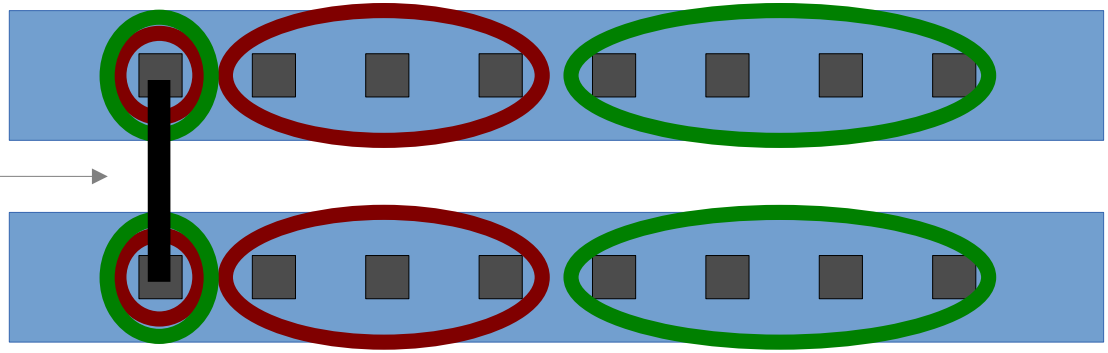

```
interface Gi1
  switchport mode trunk
  switchport trunk allowed vlan 10,20
```

Question: what is different about the frames on this wire?

Management IP address

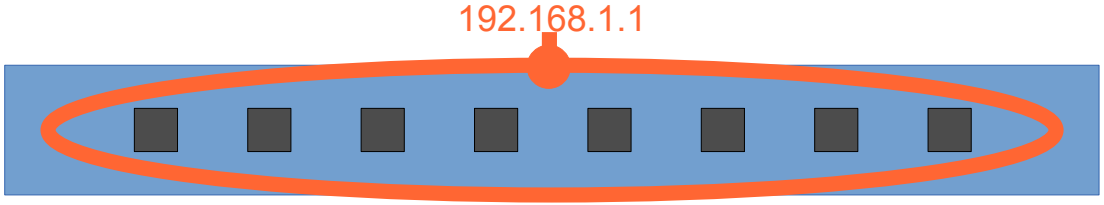


```
vlan 1

interface range Gi 1 - 8
  switchport access vlan 1

interface Vlan1
  ip address 192.168.1.1 255.255.255.0

ip default-gateway 192.168.1.254
! or: ip route 0.0.0.0 0.0.0.0 192.168.1.254
```

The Management Interface

- The switch has its own IP interface on VLAN 1, with its own IP address
- Imagine the switch's CPU is plugged into VLAN 1 (but without using up a physical port)
- You use this to manage the switch (ssh, snmp)
- Like any other IP device, it needs a default gateway to be able to send packets to a destination address on a different subnet

IP routing

- Extend this by giving the switch an IP address on multiple VLANs
  - Each address is of course within the IP subnet for that particular VLAN
- Enable the internal router within the switch
- It can receive datagrams on one VLAN, and resend them on another
- You have a layer 3 switch!

IP routing

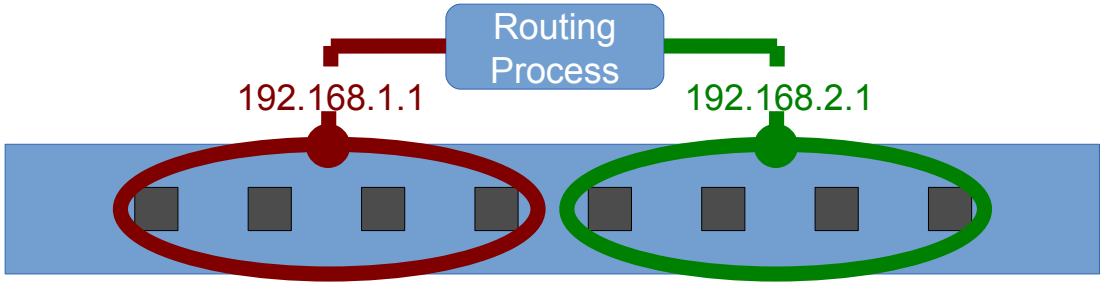


```
vlan 10,20
ip routing

interface Vlan10
  ip address 192.168.1.1 255.255.255.0
interface Vlan20
  ip address 192.168.2.1 255.255.255.0

ip route 0.0.0.0 0.0.0.0 192.168.1.254
```

Routed VLAN interfaces

- It's really that simple!
- We have an IP address on each VLAN
- Other devices can point their default gateway at us
- We will forward datagrams on their behalf
  - Based on our IP forwarding table
  - Connected routes, static routes etc.

Acting as a gateway

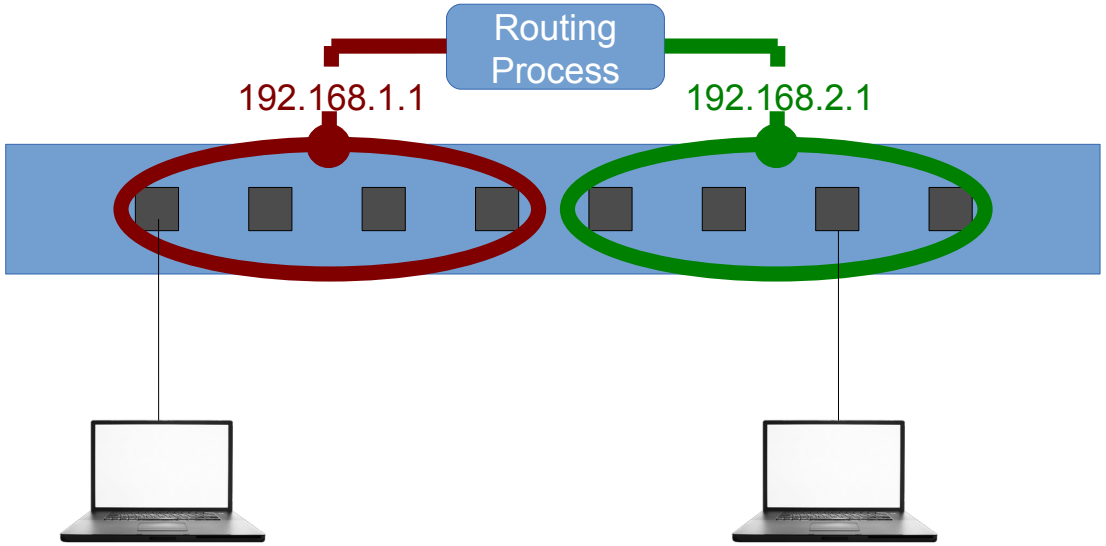

IP addr: 192.168.1.50
Gateway: **192.168.1.1**

IP addr: 192.168.2.99
Gateway: **192.168.2.1**

IPv6 is the same

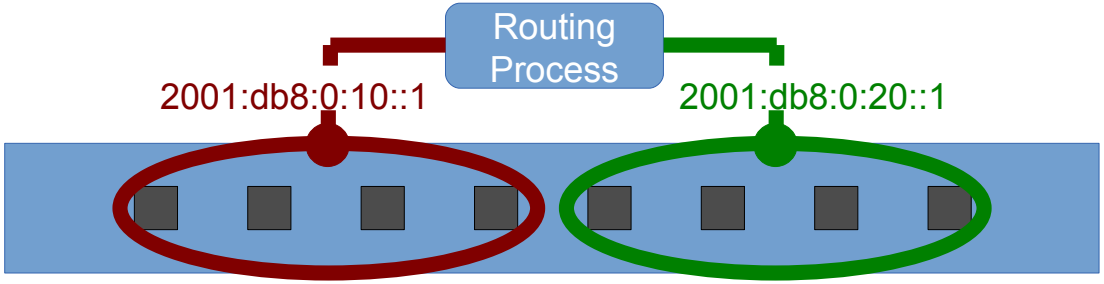


```
ipv6 unicast-routing
!
interface Vlan10
  ipv6 address 2001:db8:0:10::1/64
interface Vlan20
  ipv6 address 2001:db8:0:20::1/64
!
ipv6 route ::/0 2001:db8:0:10::ff
```

Simple campus: 1 subnet/building

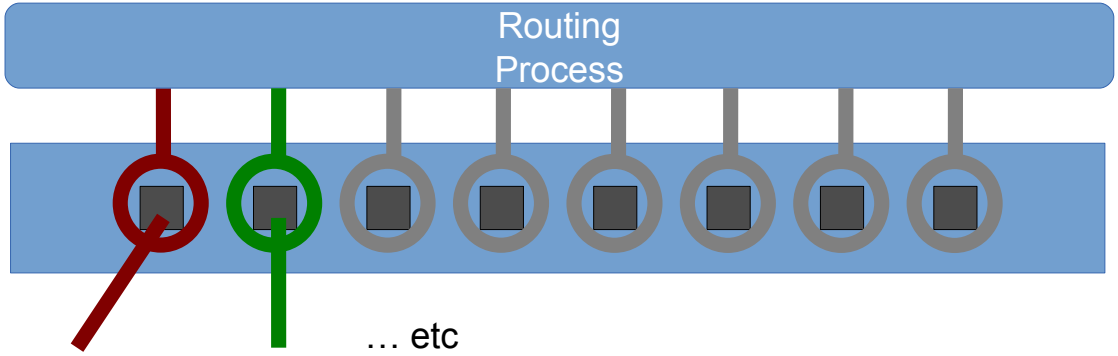


```
interface Gi1
  switchport mode access
  switchport access vlan 10

interface Vlan10
  ip address 192.168.1.1 255.255.255.0
```

```
interface Gi2
  switchport mode access
  switchport access vlan 20

interface Vlan20
  ip address 192.168.2.1 255.255.255.0
```

Multiple subnets per building

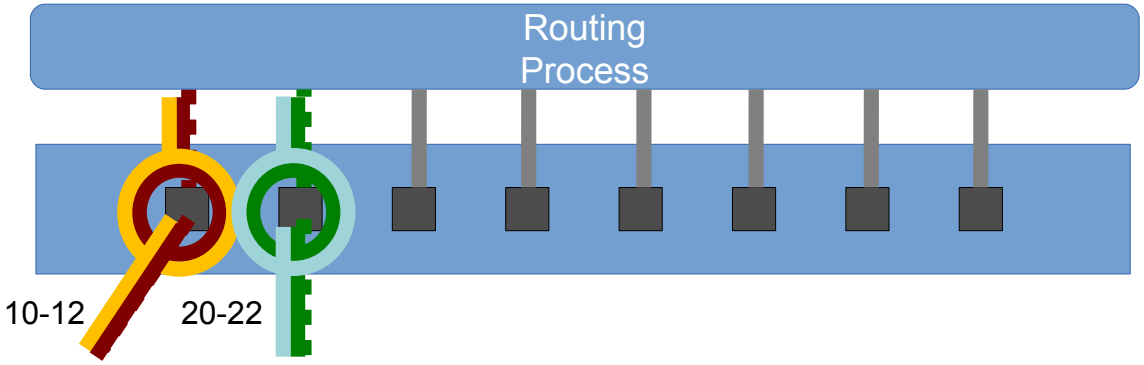


```
interface Gi1
  switchport mode trunk
  switchport trunk allowed vlan 10-12

interface Vlan10
  ip address 10.1.0.1 255.255.255.0
interface Vlan11
  ip address 10.1.1.1 255.255.255.0
interface Vlan12
  ip address 10.1.2.1 255.255.255.0
```

```
interface Gi2
  switchport mode trunk
  switchport trunk allowed vlan 20-22

interface Vlan20
  ip address 10.2.0.1 255.255.255.0
interface Vlan21
  ip address 10.2.1.1 255.255.255.0
interface Vlan22
  ip address 10.2.2.1 255.255.255.0
```


Question: what has to be different at the building aggregation switch?

Hints and tips

- Remember, one subnet = one VLAN
- Don’t use VLAN 1
  - It’s the “default VLAN” and often has special default behavior
  - It may appear by default on all ports
  - It’s often hard to use with tagging
  - Better to ignore it or remove it completely (if possible)
  - In general, VLANs 2 to 4094 are usable
    - some vendors reserve some for special purposes: check the documentation

Hints and tips

- Don’t enable the same VLAN on links to different buildings!
  - A layer 3 switch lets you do this but that doesn’t mean it’s a good idea. “VLAN spaghetti”
- Implies: a management VLAN per building, a wired VLAN per building, a WiFi VLAN per building etc
- Choose a consistent scheme
  - e.g. VLAN 2-9 for core, VLAN 10-19 for building 1, VLAN 20-29 for building 2 etc.

Routed interfaces / subinterfaces

- Some layer 3 switches let you configure “routed ports”
  - This makes it work exactly like a router instead of a switch
- Some also support routed **sub-interfaces** with VLAN tags
  - This means you can route multiple subnets to each building, without having to create any actual VLANs
  - You can re-use the *same* VLAN tags for *different subnets in different buildings!*
  - Makes the distribution/edge switch configs almost identical everywhere
- Routed ports don't participate in STP, so each building becomes its own independent spanning tree domain

Fully routed interfaces

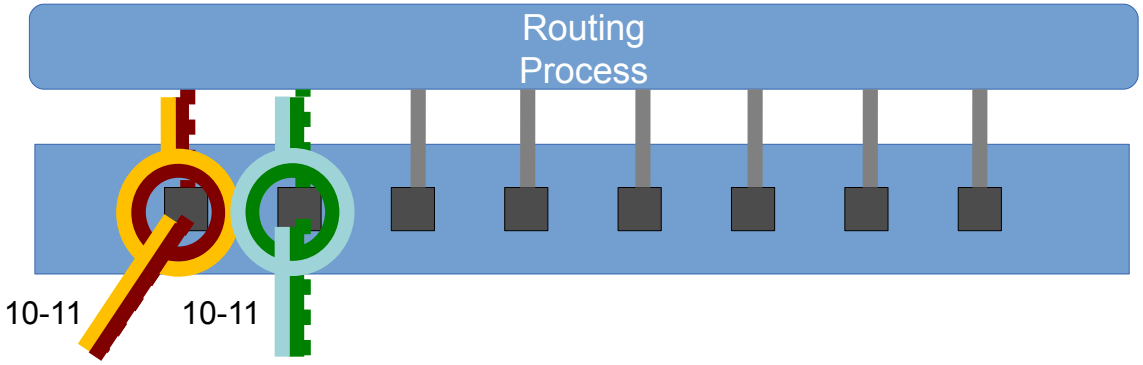


```
interface Gi1
  no switchport

interface Gi1.10
  encapsulation dot1q 10
  ip address 10.1.0.1 255.255.255.0

interface Gi1.11
  encapsulation dot1q 11
  ip address 10.1.1.1 255.255.255.0
```

```
interface Gi2
  no switchport

interface Gi2.10
  encapsulation dot1q 10
  ip address 10.2.0.1 255.255.255.0

interface Gi2.11
  encapsulation dot1q 11
  ip address 10.2.1.1 255.255.255.0
```

Both buildings use VLAN tags 10-11 but these are different, isolated subnets

Questions?