Chapter 4 – Maintaining & Troubleshooting Campus Switched LANs Objectives

- Describe VLAN and inter-VLAN troubleshooting techniques.
- Describe STP troubleshooting techniques.
- Describe FHRP troubleshooting techniques

Possible LAN Issues

- Physical problems
- Bad, missing, or mis-wired cables
- Bad ports
- Power failure
- Device problems
- Software bugs
- Performance problems
- Misconfiguration
- Missing or wrong VLANs
- Misconfigured VLAN Trunking Protocol (VTP) settings
- Wrong VLAN setting on access ports
- Missing or mis-configured trunks
- Native VLAN mismatch
- VLANs not allowed on trunk

VLAN Types

- Data – user data, with the switching block
- Voice – VoIP telephony
- Management – device management for administrators
- Native – supports untagged traffic (802.1q only)

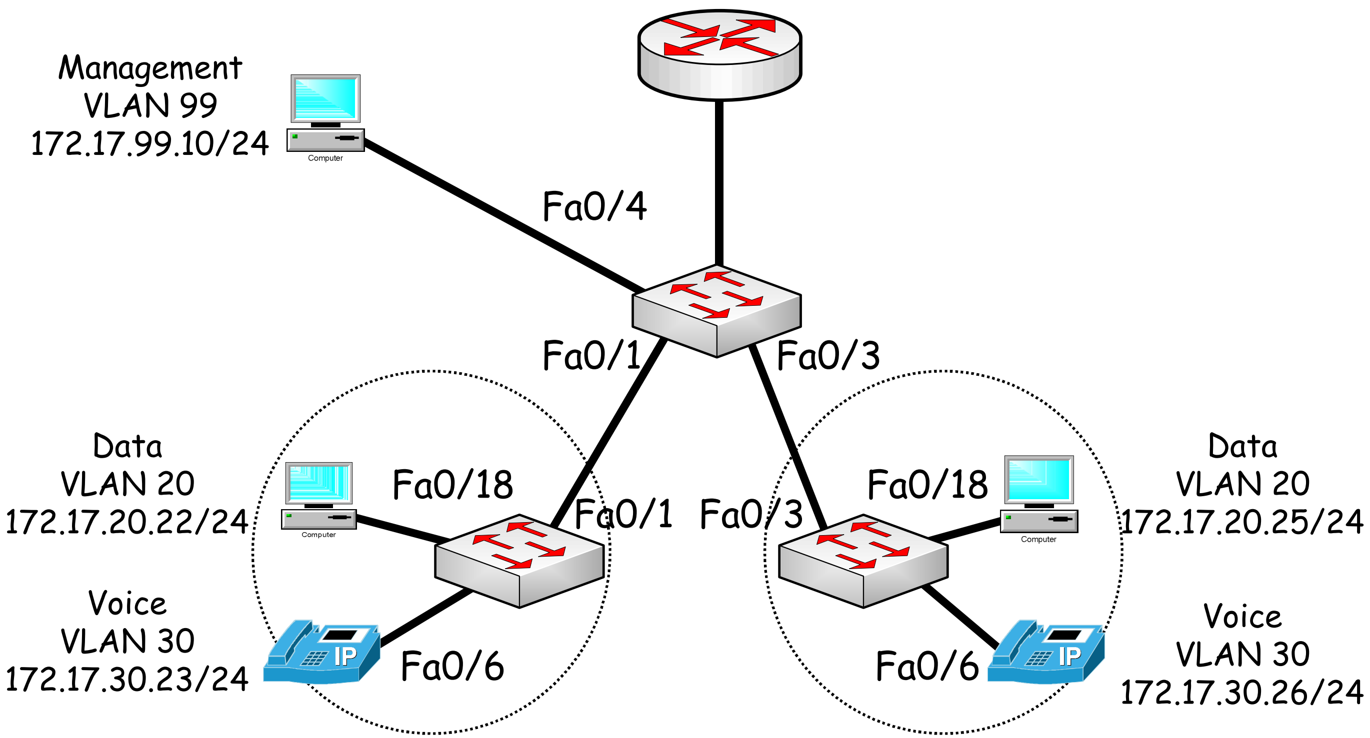

VLAN Trunks

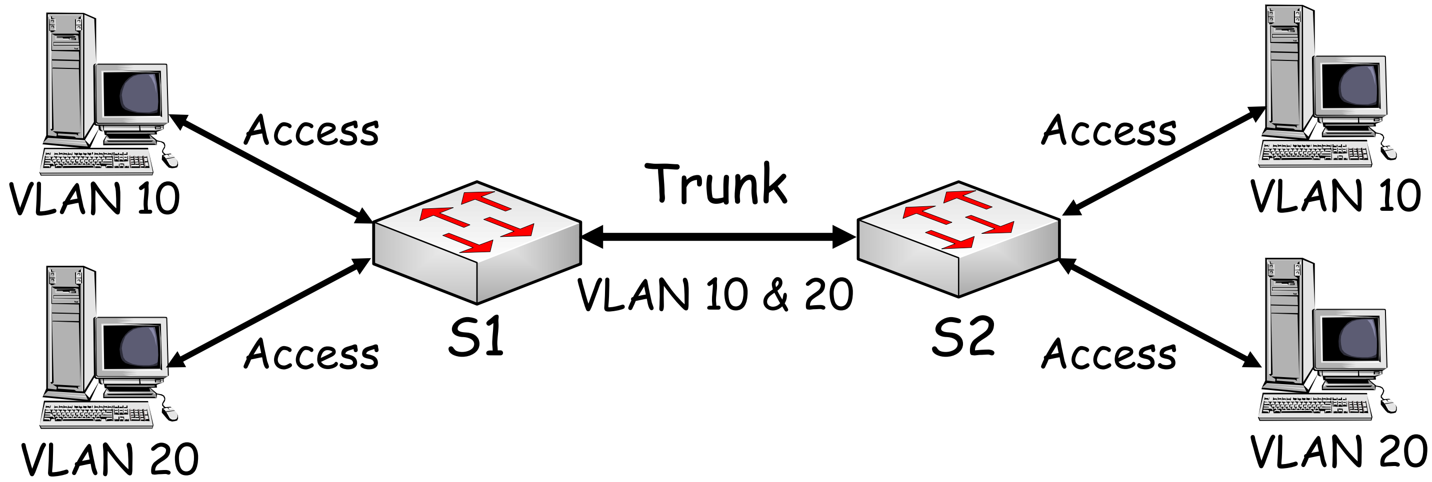

- A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device, such as a router or a switch. Ethernet trunks carry the traffic of *multiple* VLANs over a *single* link.

- A VLAN trunk allows extension of VLANs across an entire network. Cisco supports *IEEE 802.1Q* for coordinating trunks on Fast Ethernet and Gigabit Ethernet interfaces.

- A VLAN trunk does not belong to a specific VLAN, rather it is a conduit for VLANs between switches and routers.

802.1Q Tagging

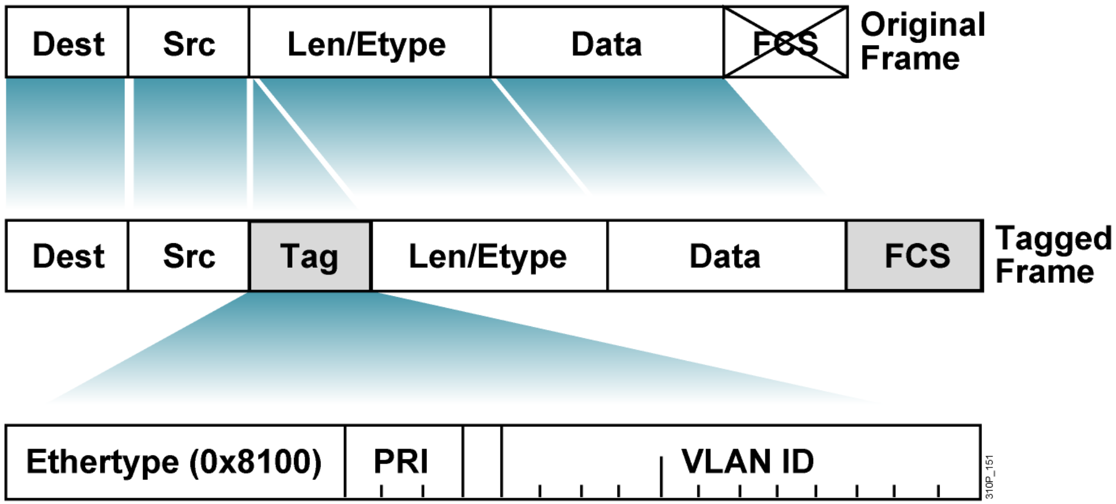

- 802.1Q does not encapsulate the original frame, but modifies the Ethernet type field by adding a Tag Control Information (TCI) field.
- A TCI contains a 12-bit VLAN identifier (VID), uniquely identifying the VLAN to which the frame belongs (4,096 VLANs max, with 0 and 4095 reserved).
- Because inserting this header changes the frame, 802.1Q encapsulation forces a recalculation of the original FCS field in the Ethernet trailer.

Native VLANs

- A *native VLAN* is assigned to all 802.1Q trunk ports.
- An 802.1Q trunk port must supports traffic coming from many VLANs (tagged traffic) as well as traffic that does not come from a VLAN (untagged traffic).
- Untagged traffic is carried over the native VLAN.

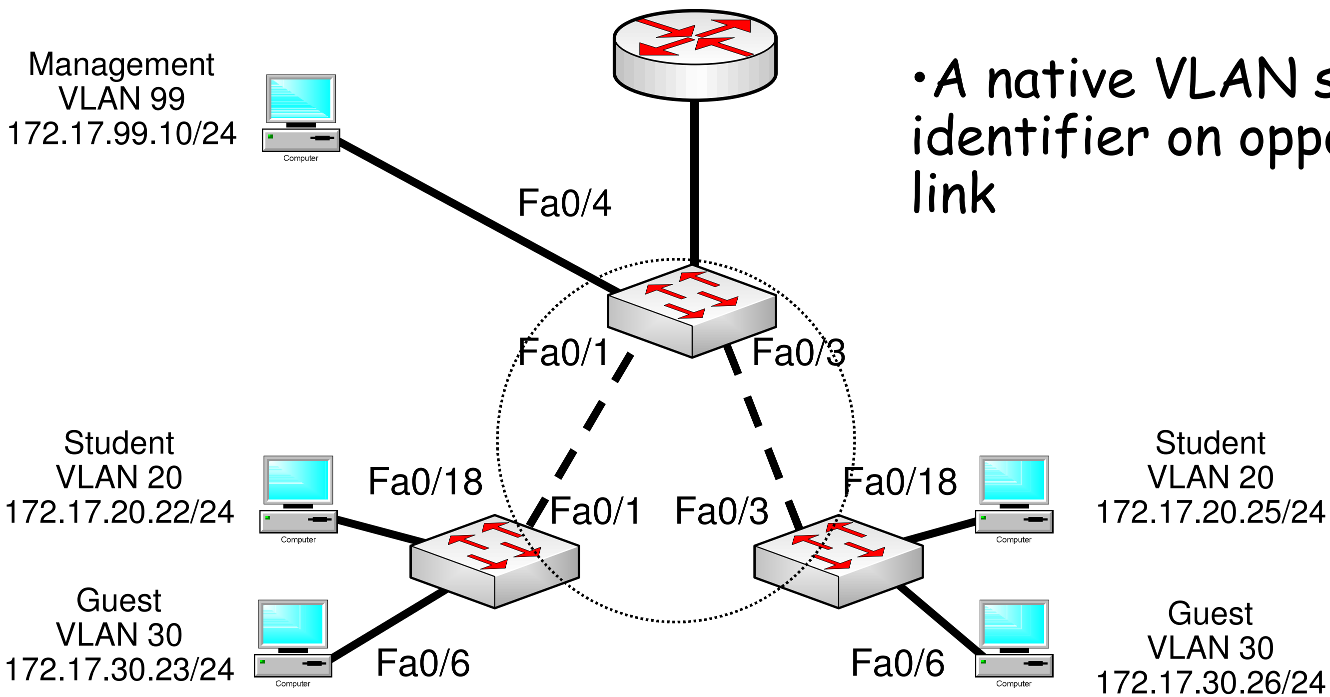

- Trunks are used to allow the same VLAN to span different switches

- A native VLAN serves as a common identifier on opposing ends of a trunk link

Verifying Layer 2

Cat3550# **show vlan**

VLAN	Name	Status	Ports
1	default	active	Gi0/7, Gi0/8, Gi0/10, Gi0/11, Gi0/12
10	VLAN0010	active	
20	SPAN	active	
261	VLAN0261	active	Gi0/1, Gi0/2, Gi0/3, Gi0/4 Gi0/5, Gi0/6
262	VLAN0262	active	Gi0/1, Gi0/2, Gi0/3, Gi0/4

Cat3550# **show interfaces trunk**

Port	Mode	Encapsulation	Status	Native vlan
Gi0/9	desirable	n-isl	trunking	1

Port	Vlans allowed on trunk
Gi0/9	1-4094

Port	Vlans allowed and active in management domain
Gi0/9	1,10,20,261-262

Port	Vlans in spanning tree forwarding state and not pruned
Gi0/9	1,10,20,261-262

Verifying Layer 2

Cat3550# show mac address-table

```
Mac Address Table
-------------------------------------------
Vlan          Mac Address          Type       Ports
---- ------------ ----- -----
1        0000.865c.7fc2          DYNAMIC    Gi0/7
1        0009.1260.0aee          DYNAMIC    Gi0/9
1        0009.b7fa.d1e1          DYNAMIC    Gi0/9
10       0009.1260.0aee          DYNAMIC    Gi0/9
261      0004.27d4.0b21          DYNAMIC    Gi0/3
261      0008.a3b8.945e          DYNAMIC    Gi0/1
Total Mac Addresses for this criterion: 48
...OUTPUT OMITTED...
```

clear mac address-table dynamic
Clears dynamically learned MAC addresses from a switch's MAC address table; this can help determine if a previously learned MAC address is relearned

Switching Loops

The addition of redundant paths creates switching loops. These can lead to the following problems:

- Broadcast Storms
- Multiple Frame Transmission
- MAC Database Instability

Spanning Tree Protocol 802.1d (STP)

- The solution is to allow physical loops, but create a *loop free logical* topology called a *tree*.

- It is a spanning-tree because all devices in the network are reachable or *spanned*.

- The algorithm used to create this loop free logical topology is the *spanning-tree algorithm*.

- STP exchanges information called *Bridge Protocol Data Units* (BPDUs).

- A new algorithm called the *rapid spanning-tree algorithm* was developed to reduce the time for a network to compute a loop free logical topology .

STP/RSTP Port State Comparison

Operational Port State	STP Port State	RSTP Port State
Enabled	Blocking	Discarding
Enabled	Listening	Discarding
Enabled	Learning	Learning
Enabled	Forwarding	Forwarding
Disabled	Disabled	Discarding

RSTP Port roles

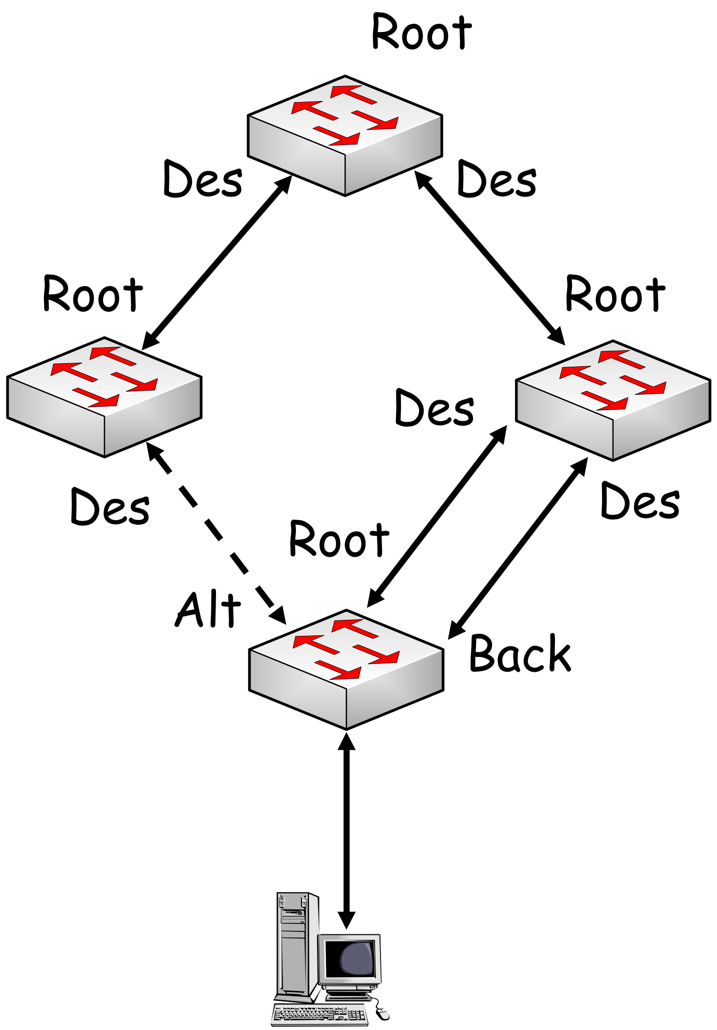

- *Alternative port:* switch port that offers an alternative path toward the root bridge.
- The alternative port assumes a *discarding* state in a stable, active topology.
- Backup port: additional switch port on the designated switch with a *redundant* link to the segment for which the switch is designated.
- A backup port has a *higher* port ID than the designated port on the designated switch.
- The backup port assumes the *discarding* state in a stable, active topology.

Step 1 - Root Bridge Election Process

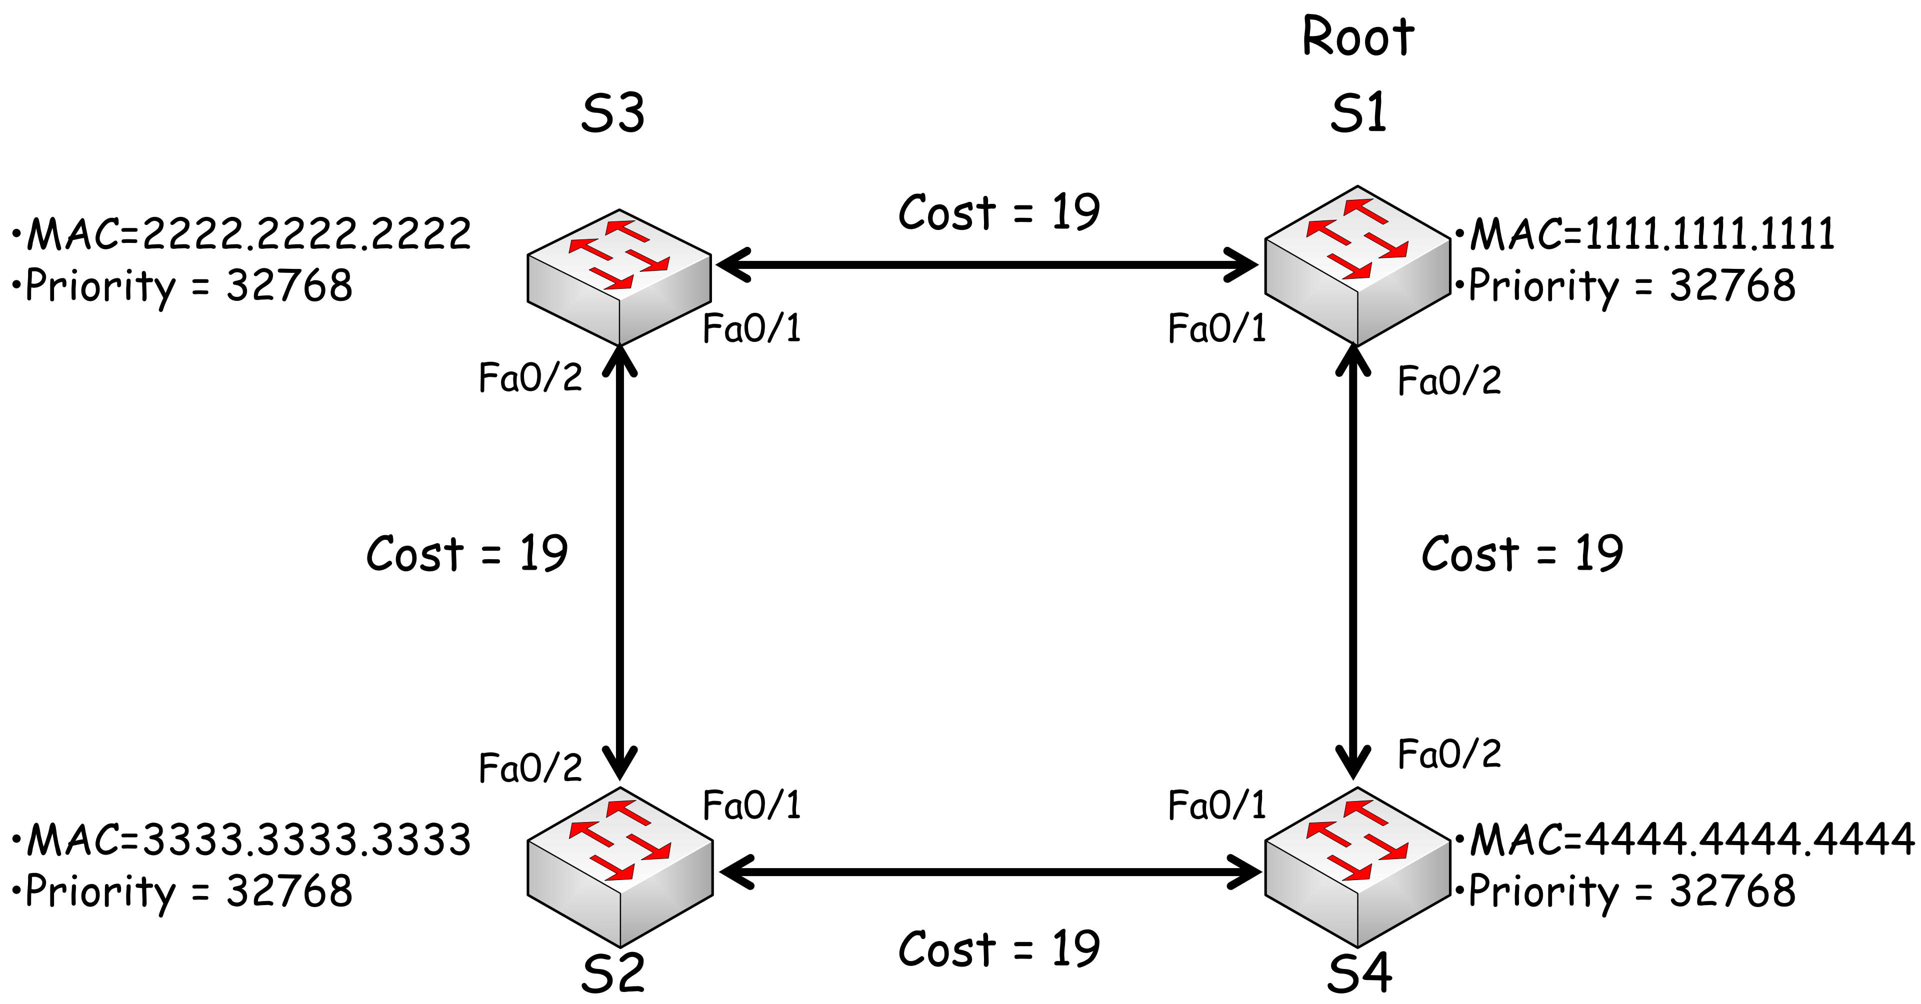

Step 2 - Root Port Election Process

Link Speed	Cost (Revised IEEE Specification)
10 Gb/s	2
1 Gb/s	4
100 Mb/s	19
10 Mb/s	100

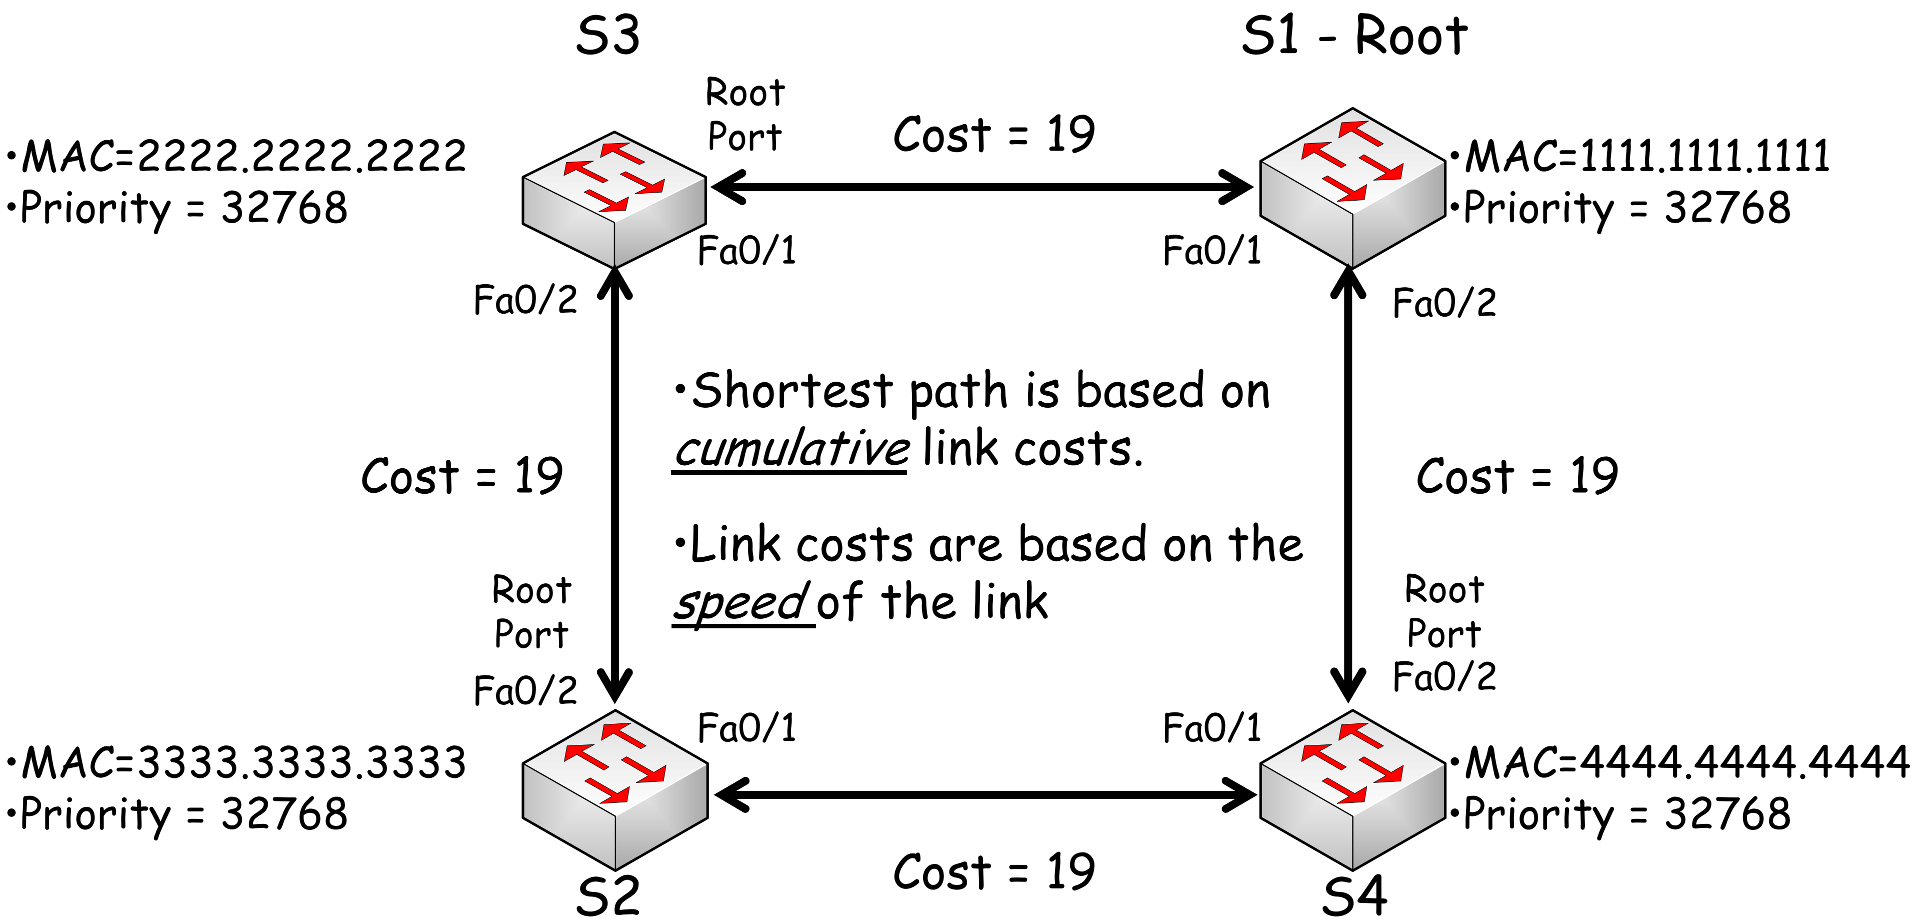

- Shortest path is based on *cumulative* link costs.
- Link costs are based on the *speed* of the link

1. Lowest root Bridge ID (BID)
2. Lowest path cost to root bridge
3. Lowest sender bridge ID
4. Lowest port ID

Step 3 - Designated Port Election Process

Link Speed	Cost (Revised IEEE Specification)
10 Gb/s	2
1 Gb/s	4
100 Mb/s	19
10 Mb/s	100

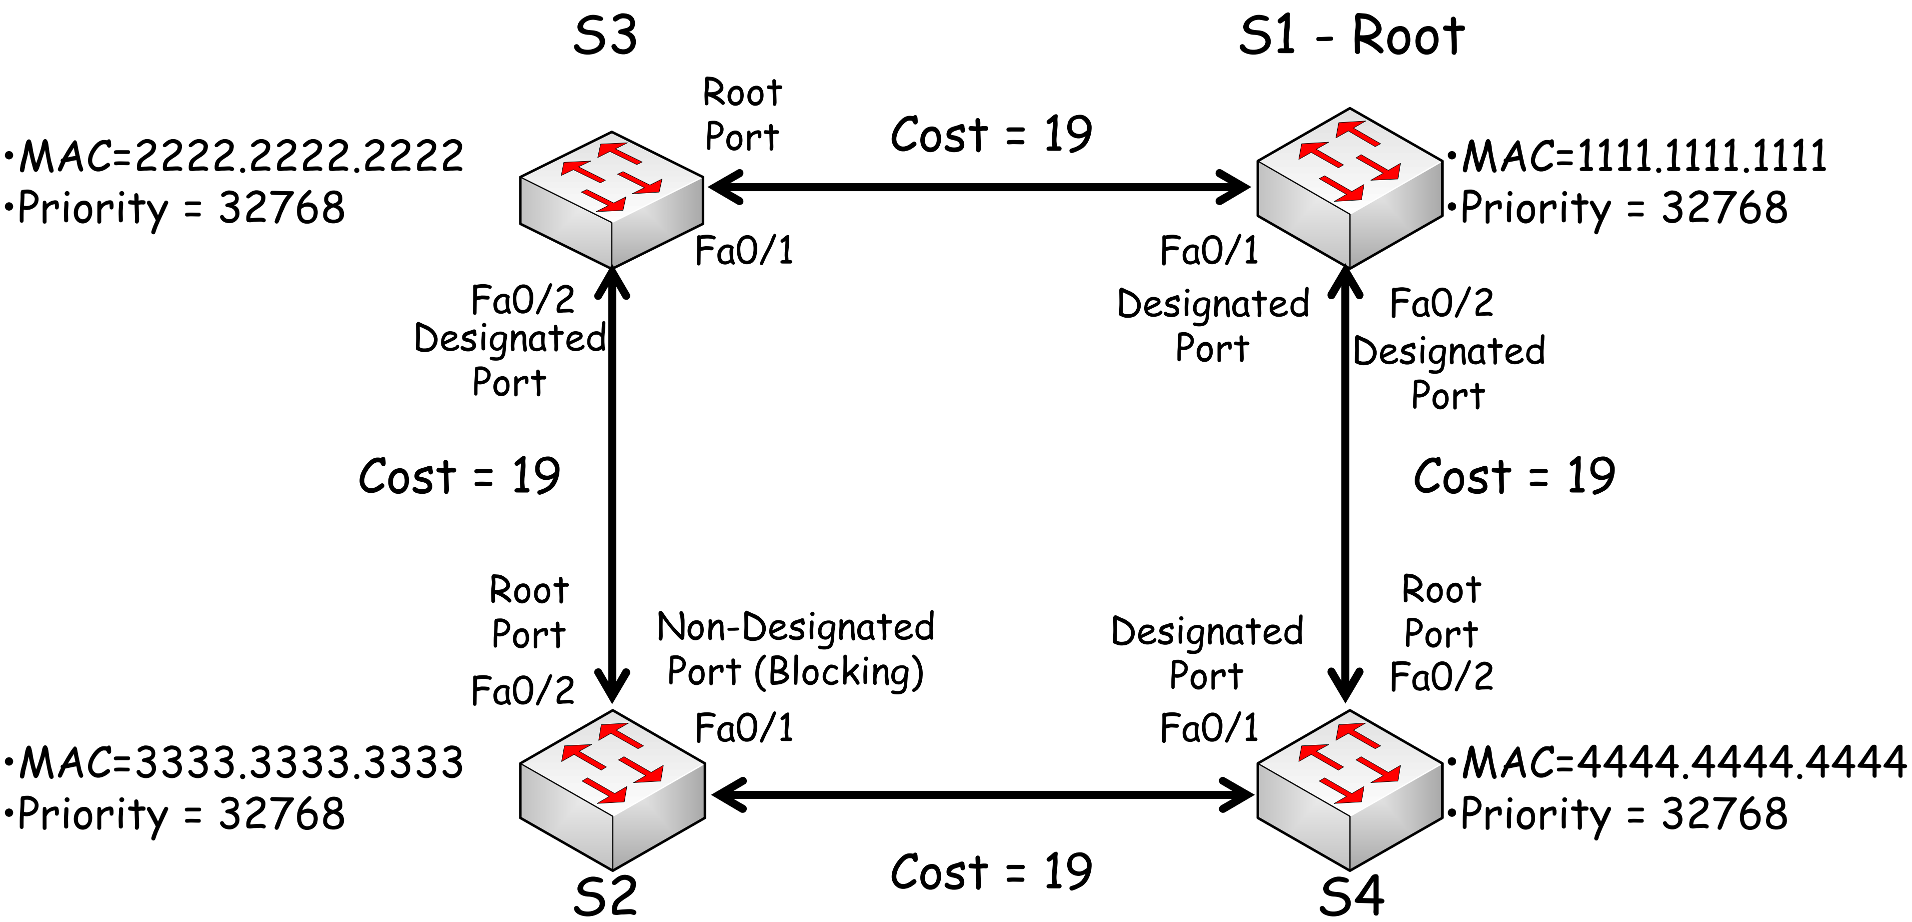

1. Lowest root Bridge ID (BID)
2. Lowest path cost to root bridge
3. Lowest sender bridge ID
4. Lowest port ID

Verifying STP

SW2# show spanning-tree vlan 1

```
VLAN0001
Spanning tree enabled protocol ieee

Root ID       Priority    32768
              Address     0009.122e.4181
              Cost 19
              Port 9 (GigabitEthernet0/9)
              Hello Time 2 sec     Max Age 20 sec     Forward Delay 15 sec

Bridge ID     Priority    32769 (priority 32768 sys-id-ext 1)
              Address     000d.28e4.7c80
              Hello Time 2 sec     Max Age 20 sec     Forward Delay 15 sec
              Aging Time 300

Interface  Role   Sts    Cost    Prio.Nbr   Type
----------------------------------------------------------------------------
Gi0/9      Root   FWD    19      128.9      P2p
Gi0/10     Altn   BLK    100     128.10     Shr
```

SW2# show spanning-tree interface gig 0/9 detail

```
Port 9 (GigabitEthernet0/9) of VLAN0001 is root forwarding
Port path cost 19, Port priority 128, Port Identifier 128.9.
Designated root has priority 32768, address 0009.122e.4181
Designated bridge has priority 32768, address 0009.122e.4181
Designated port id is 128.303, designated path cost 0
Timers: message age 2, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDU: sent 1, received 1245
```

Spanning-Tree Failure

- STP failure/misconfiguration may result in a broadcast storm, exhibiting the following symptoms:
  1. The load on all links in the switched LAN will quickly start increasing as traffic increases exponentially
  2. Layer 3 protocol's cpu utilisation increases.
  3. Devices become unreachable via SSH/Telnet.
- One intrusive but effective way to start tackling severe spanning-tree problems is eliminating topological loops by either physically disconnecting links or by shutting down interfaces if that is still possible.
- Once the loops are broken, the traffic and CPU loads should quickly drop to normal levels

Link Aggregation With EtherChannel

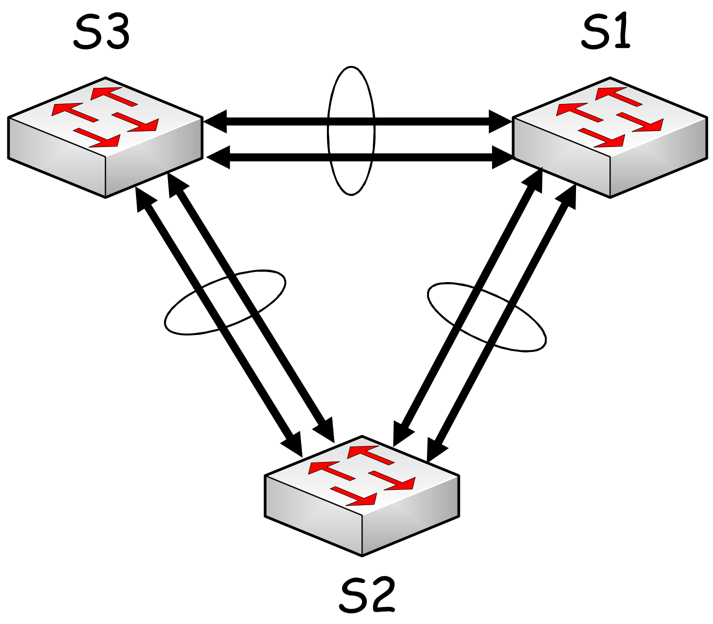

- Allows for the creation of a very-high-*bandwidth* logical link
- *Load balances* amongst the physical links involved
- Provides *automatic* failover
- *Simplifies* subsequent logical configuration (configuration is per logical link instead of per physical link)

- EtherChannel bundles individual Ethernet links into a single logical link that provides bandwidth up to 1600 Mbps (Fast EtherChannel, full duplex) or 16 Gbps (Gigabit EtherChannel) between two Cisco Catalyst switches.
- All interfaces in each EtherChannel must be the same speed,duplex, trunk mode & native VLAN.

Link Aggregation With PAgP

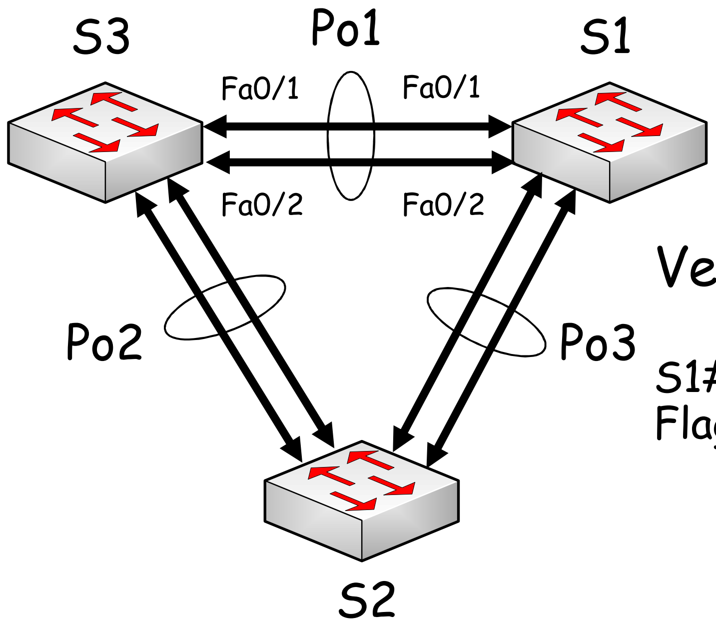

S1(config-if-range)#interface range fa0/1 – 2
S1(config-if-range)#channel-protocol pagp
S1(config-if-range)#channel-group 1 mode on

Verify:

```
S1#sh etherchannel summary
Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port


Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------
-
1       Po1(SU)         PAgP   Fa0/1(P) Fa0/2(P)
```

EtherChannel Failure

•Control protocols like STP or routing protocols will only interact with this single port-channel interface and not with the associated physical interfaces.

•There are three common EtherChannel problems:

1. Inconsistencies between the physical ports that are members of the channel.
2. Inconsistencies between the ports on the opposite sides of the Ether-Channel link.
3. Uneven distribution of traffic between EtherChannel bundle members.

EtherChannel Failure

DLS1# **show etherchannel summary**

```
Flags: D - down P - bundled in port-channel
I - stand-alone            s - suspended
H - Hot-standby (LACP only)
R - Layer3        S - Layer2
U - in use          f - failed to allocate aggregator
M - not in use, minimum links not met
u - unsuitable for bundling
w - waiting to be aggregated
d - default port

Number of channel-groups in use: 2
Number of aggregators: 2
Group  Port-channel    Protocol         Ports
------+--------------------+------------+-----------------------
1        Po1(SD)          -          Fa0/5(s) Fa0/6(s)
2        Po2(SU)          -          Fa0/3(P) Fa0/4(P)
```

EtherChannel Failure

DLS1# show etherchannel 1 detail

```
Group state = L2
Ports: 2 Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol: -
Minimum Links: 0
Ports in the group:
-------------------
Port: Fa0/5
------------
Port state = Up Cnt-bndl Suspend Not-in-Bndl
Channel group = 1          Mode = On          Gcchange = -
Port-channel = null        GC = - Pseudo      port-channel = Po1
Port index = 0             Load = 0x00        Protocol = -

Age of the port in the current state: 0d:00h:25m:13s

Probable reason: vlan mask is different

Mar 20 08:12:39 PDT: %EC-5-CANNOT_BUNDLE2: Fa0/5 is not compatible with Po1 and
will be suspended (vlan mask is different)
Mar 20 08:12:39 PDT: %EC-5-CANNOT_BUNDLE2: Fa0/6 is not compatible with Po1 and
will be suspended (vlan mask is different)
```

Layer-3 Switch

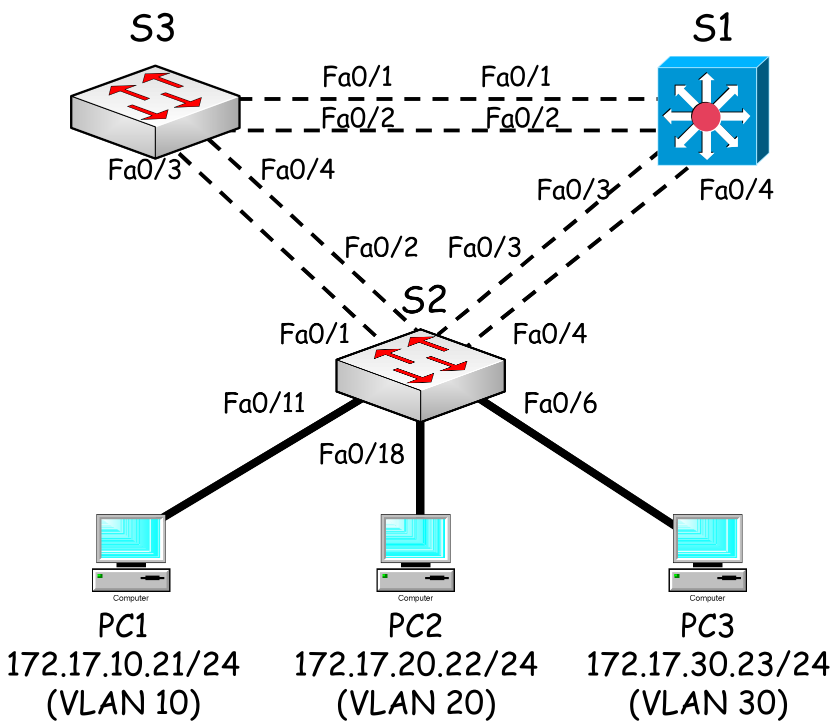

<u>S1 VLAN Interfaces</u>

172.17.10.1 – Default Gateway to VLAN 10

172.17.20.1 – Default Gateway to VLAN 20

172.17.30.1 – Default Gateway to VLAN 30

- Multilayer switches can perform Layer 3 functions, replacing the need for dedicated routers to perform basic routing on a network.

- Multilayer switches are capable of performing inter-VLAN routing.

- To enable a multilayer switch to perform routing functions, VLAN interfaces on the switch need to be configured with the appropriate *IP addresses* that match the subnet that the VLAN is associated with on the network. The multilayer switch also must have *IP routing* enabled.

Layer-3 Switch SVI Configuration

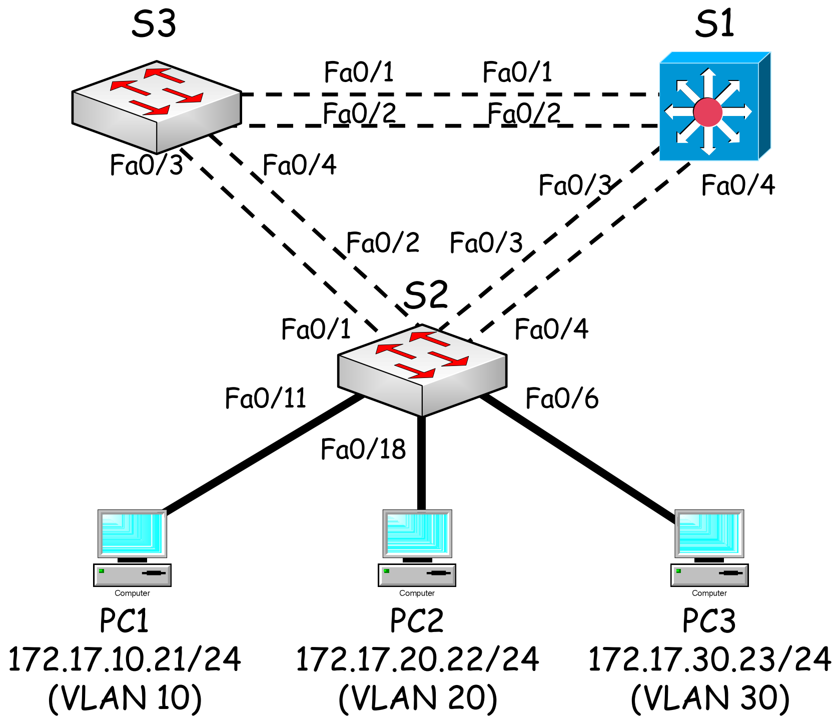

Configure SVI Addresses:

```
S1(config)#int vlan 10
S1(config-if)#ip add 172.17.10.1 255.255.255.0
S1(config-if)#int vlan 20
S1(config-if)#ip add 172.17.20.1 255.255.255.0
S1(config-if)#int vlan 30
S1(config-if)#ip add 172.17.30.1 255.255.255.0
```

Configure Routing:

```
S1(config)#ip routing
S1(config)#exit
S1#sh ip route

     172.17.0.0/24 is subnetted, 3 subnets
C       172.17.10.0 is directly connected, Vlan10
C       172.17.20.0 is directly connected, Vlan20
C       172.17.30.0 is directly connected, Vlan30
```

- An SVI is not a physical interface, and so it generally does not fail, but its status directly depends on the status of the _VLAN_ with which it is _associated_.
- An SVI stays up as long as there is at least _one port associated_ to the corresponding VLAN; that port has to be up and in the spanning-tree forwarding state.

Multilayer Switch Packet Forwarding Process

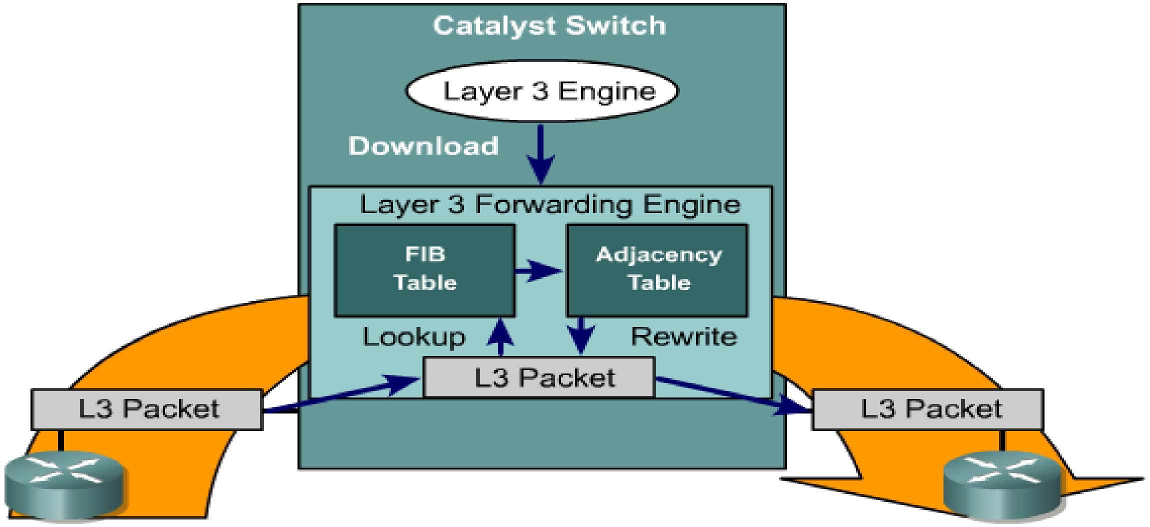

•CEF expediently switches data packets to their destination. It *caches* information generated by the Layer 3 routing engine.

•CEF caches routing information in the *Forwarding Information Base* (FIB), and caches Layer 2 next-hop addresses for all FIB entries in an adjacency table.

•Because CEF maintains multiple tables for forwarding information, parallel paths can exist and enable CEF to load balance per packet.

CEF Tables

IP Routing Table

Protocol	Address	Prefix	Next-Hop	Exit-Int
D	10.1.1.0	/24	10.1.1.2	Fa0/1
D	10.1.2.0	/24	10.1.2.2	Fa0/2

ARP Table

IP Address	MAC Address
10.1.1.2	0c.00.11.22.33.44
10.1.2.2	0c.00.11.55.66.77

Control Plane

FIB Table

IP Address	Prefix	Adjacency Pointer
10.1.1.0	/24	10.1.1.2
10.1.2.0	/24	10.1.2.2

Adjacency Table

IP Address	Layer 2 Header	
10.1.1.2	DA	SA
10.1.2.2	DA	SA

Data Plane

CEF Operation in Switches

- In contrast to routers, multilayer switches do not just use the FIB & Adjacency tables for packet forwarding, they compile and download the information they contain into the *ternary content-addressable memory* (TCAM).

- Using specialised hardware, multilayer switches forward packets at high speeds based on the information contained in the TCAM.

- The hardware uses *Application-Specific Integrated Circuits* (ASIC).

CEF Adjacencies

- **Null adjacency:** Packets destined for a null0 interface are dropped. This can be used as an effective form of access filtering.
- **Glean adjacency:** When a router is connected directly to several hosts, the FIB table on the router maintains a prefix for the subnet rather than for the individual host prefixes. The subnet prefix points to a glean adjacency. When packets need to be forwarded to a specific host, the adjacency database is gleaned for the specific prefix.
- **Punt adjacency:** Features that require special handling, or features that are not yet supported in conjunction with CEF switching paths, are forwarded to the next switching layer for handling. For example, the packet may require CPU processing. Features that are not supported are forwarded to the next-higher switching level.
- **Discard adjacency:** Packets are discarded.
- **Drop adjacency:** Packets are dropped, but the prefix is checked.

Verifying CEF

ALS1# show ip cef

Prefix	Next Hop	Interface
0.0.0.0/32	receive	
192.168.199.0/24	attached	Vlan1
192.168.199.0/32	receive	
192.168.199.1/32	receive	
192.168.199.2/32	192.168.199.2	Vlan1
192.168.199.255/32	receive	

Configure & Verify CEF

Configure CEF:

S1 (conf)#ip cef
S1 (conf-if)#ip route-cache cef

Verify CEF:

S1#sh ip cef
S1#sh ip cef fa0/1 detail
S1#sh adjacency fa0/1 detail
S1#show ip cef summary
S1#show ip cef vlan 10

- If CEF is enabled globally, it is automatically enabled on all interfaces as long as IP routing is enabled on the device.

- CEF can be enabled/disabled on a per interface basis.

- Cisco recommends that CEF be enabled on all Layer 3 interfaces.

Layer-3 Switch Routed Port Configuration

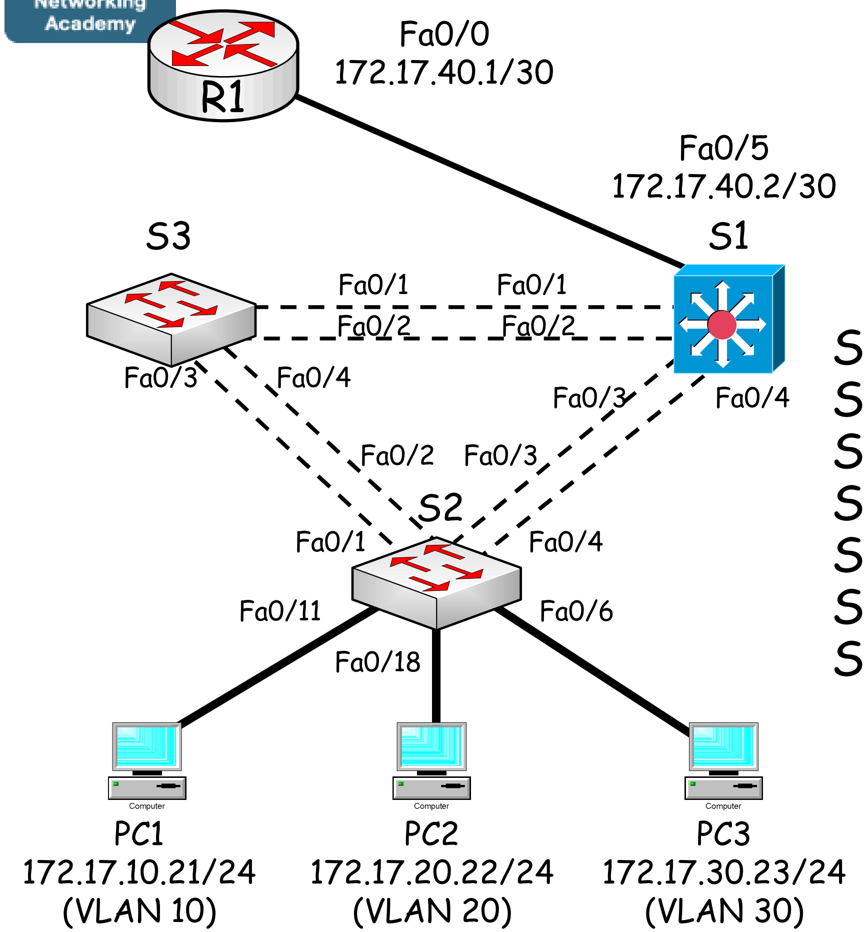

Configure Routed Port:

```
S1(config)#int fa0/5
S1(config-if)#no switchport
S1(config-if)#ip add 172.17.40.2 255.255.255.0
S1(config-if)#no sh
S1(config-if)#exit
S1(config)#router eigrp 1
S1(config-router)#network 172.17.40.0 0.0.0.3
```

A routed port has the following characteristics and functions:

- Physical switch port with Layer 3 capability
- Not associated with any VLAN
- Serves as the default gateway for devices out that switch port
- Layer 2 port functionality must be removed before it can be configured

First Hop Redundancy Protocols (FHRP)

- FHRPs such as HSRP, VRRP, and GLBP all serve the same purpose: They provide a *redundant* default gateway on a subnet and do this in such a way that actions such as *failover* and *load balancing* remain entirely transparent to the hosts.
- These protocols provide a *virtual* IP address (and the corresponding virtual MAC address) that can be used as the default gateway by the hosts on the subnet.
- This virtual IP address is not bound to any particular router, but can be controlled by a router within a group of routers participating in the protocol's scheme.
- Under normal circumstances, at any given moment, only one router, the *active* router, has control over the virtual IP address.

HSRP Operation

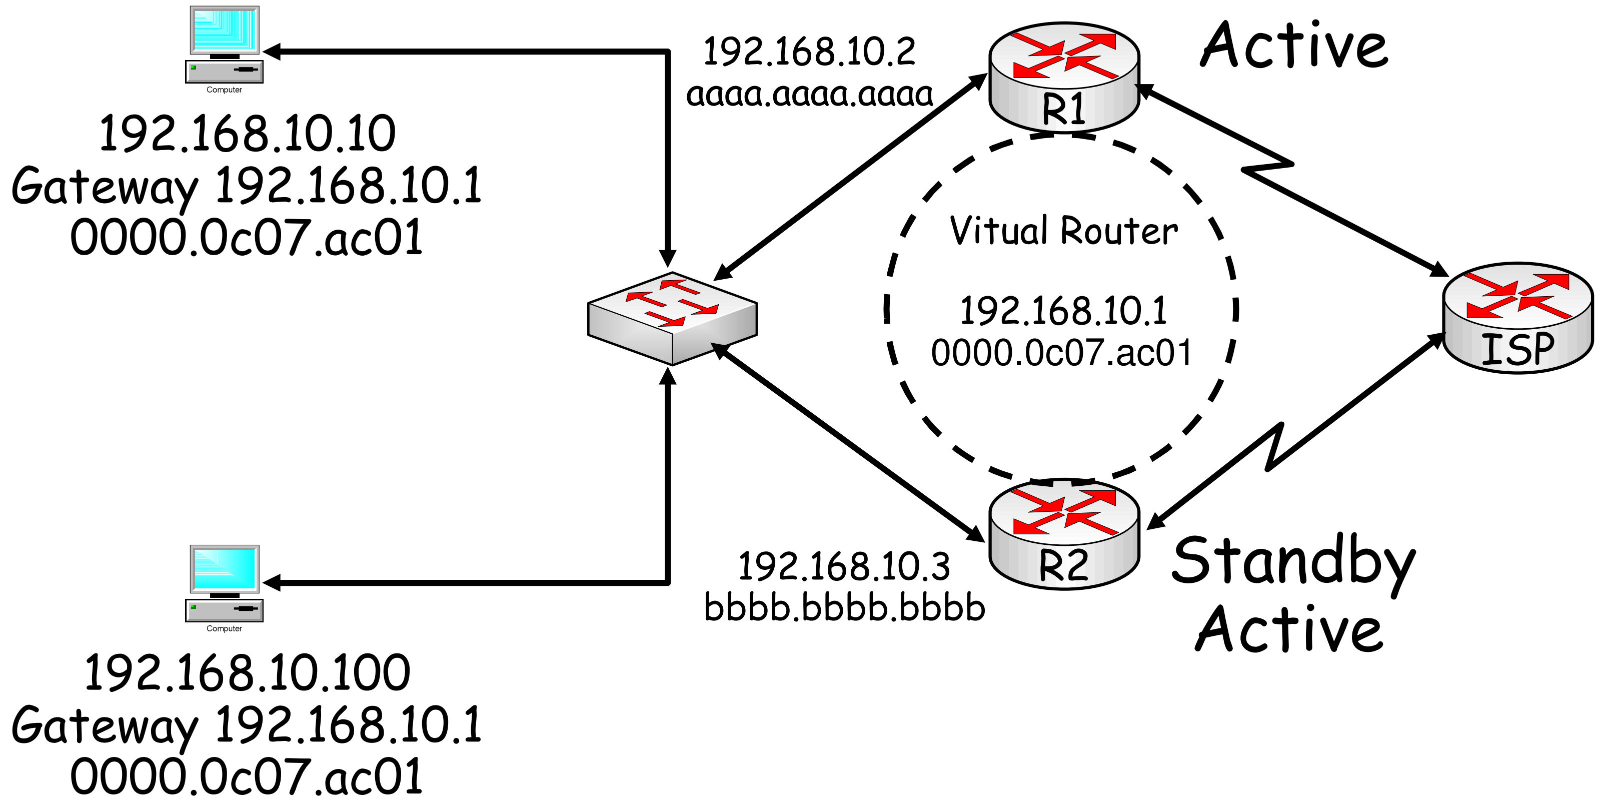

- The *virtual router* is simply an IP and MAC address pair that end devices have configured as their default gateway.
- The *active* router processes all packets and frames sent to the virtual router address.
- The HSRP standby router monitors the operational status of the HSRP group and quickly assumes packet-forwarding responsibility if the active router becomes inoperable.

HSRP Group Members

- ***Active router*** - Does the forwarding of data packets and transmits hello messages to other routers informing them of its status.
- ***Standby router*** - Monitors the status of the active router and quickly begins forwarding packets in the event of an active router failure.
- ***Virtual router*** - Does not exist! Represents a consistently available router with an IP address and a MAC address to the hosts on a network.
- ***Other routers*** - Monitor HSRP hello messages but do not respond. function as normal routers that forward packets sent to them but do not forward packets addressed to the virtual router.

HSRP Configuration

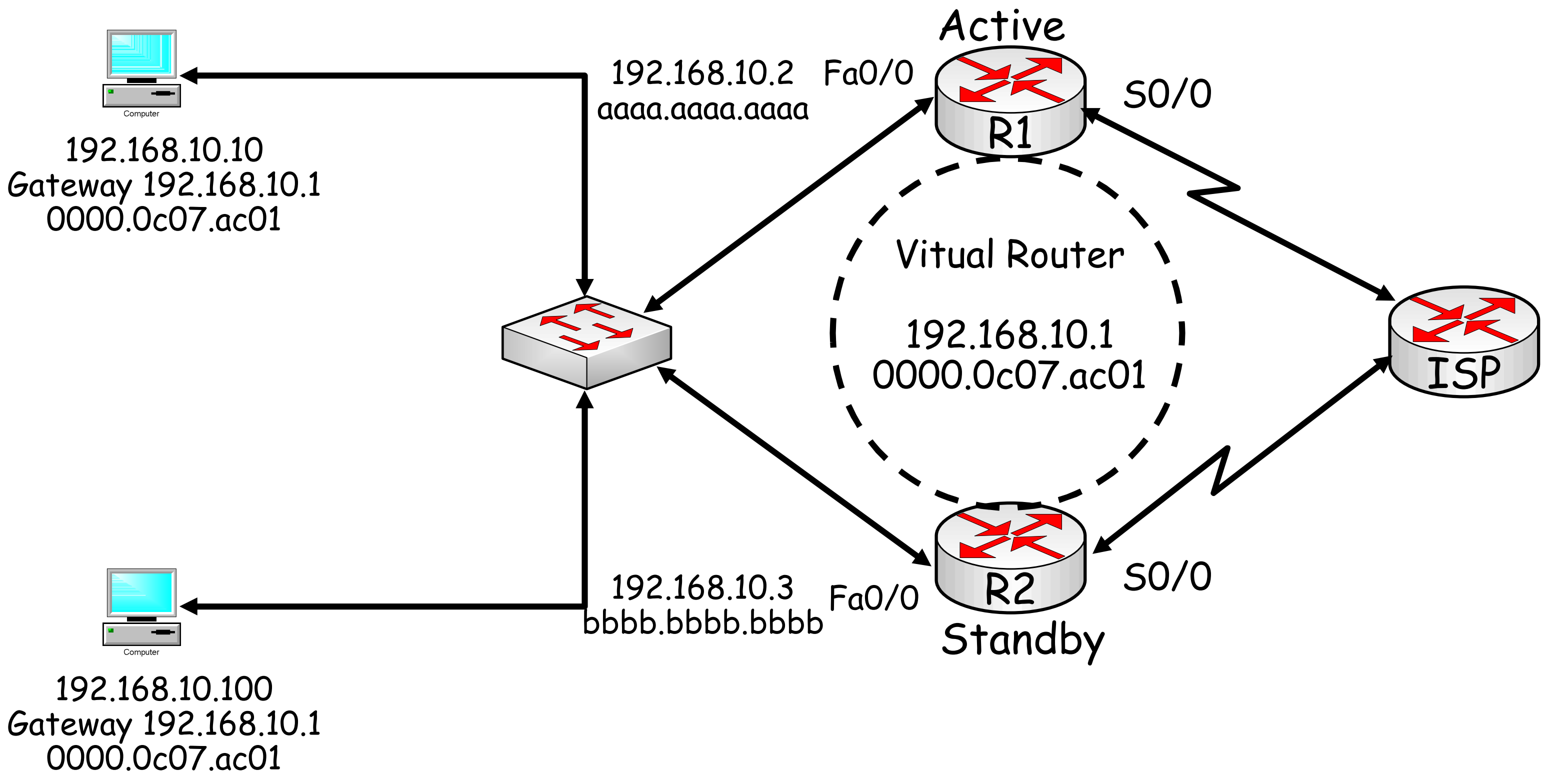

Active	Standby
R1(conf)#int fa0/0	R2(conf)#int fa0/0
R1(conf-if)#standby 1 ip 192.168.10.1	R2(conf-if)#standby 1 ip 192.168.10.1
R1(conf-if)#standby 1 priority 150	R2(conf-if)#standby 1 priority 100
R1(conf-if)#standby 1 preempt	R2(conf-if)#standby 1 preempt
R1 (conf-if)#standby 1 track s0/0 55	R2 (conf-if)#standby 1 track s0/0 55

HSRP Configuration

- The *group* number can be from 0 – 255; Default is 0. HSRP supports a maximum of 16 groups.
- The *priority* value can be from 0 – 255; Default is 100. If no priority, router with *highest* IP address on HSRP interface becomes active.
- The *hellotime* can be from 1 – 255; default is 3.
- The *holdtime* can be from 1 – 255; default is 10.
- The track command default decrement of the priority is 10

Verifying HSRP

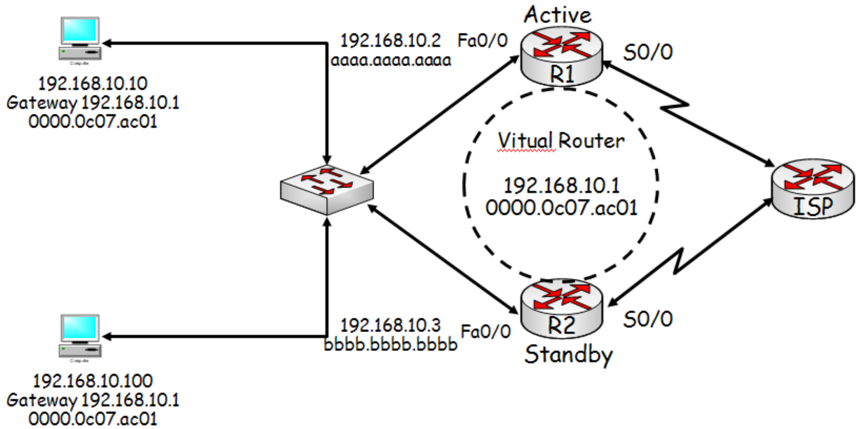


```
C:\>arp -a

Interface: 192.168.10.1          --- 0x4
Internet Address    Physical Address    Type
192.168.10.1        00-00-0c-07-ac-01   dynamic
```

R2# show standby brief

P indicates configured to preempt.

Interface	Grp	Prio	P State	Active	Standby	Virtual IP
Fa0/0	1	100	Standby	192.16.10.2	local	192.168.10.1

R1#show standby

FastEthernet0/0 - Group 1
Local state is Active, priority 150, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:01.028
Hot standby IP address is 192.168.10.1 configured
Active router is local
Standby router is 192.168.10.3 expires in 00:00:08
Tracking interface states for 1 interface, 1 up: Up Serial0/0

R2#show standby

FastEthernet0 - Group 1
Local state is Standby, priority 100, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:00.772
Hot standby IP address is 192.168.10.1
Active router is 192.168.10.2 expires in 00:00:09
Standby router is local
Standby virtual mac address is 0000.0c07.ac01
Tracking interface states for 1 interface, 1 up: Up Serial0/0

HSRP Debug

- R2#debug standby terse
- *Mar 1 00:16:23.555: HSRP: Fa0/0 Grp 1 Coup in 192.168.10.2 Listen pri 110 vIP 192.168.10.1
- *Mar 1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active: j/Coup rcvd from higher pri router
- (110/192.168.10.2)
- *Mar 1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active router is 192.168.10.2, was local
- *Mar 1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active -> Speak
- *Mar 1 00:16:23.555: %HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Active -> Speak
- *Mar 1 00:16:23.555: HSRP: Fa0/0 Grp 1 Redundancy "hsrp-Fa0/0-1" state Active -> Speak
- *Mar 1 00:16:33.555: HSRP: Fa0/0 Grp 1 Speak: d/Standby timer expired (unknown)
- *Mar 1 00:16:33.555: HSRP: Fa0/0 Grp 1 Standby router is local
- *Mar 1 00:16:33.555: HSRP: Fa0/0 Grp 1 Speak -> Standby
- *Mar 1 00:16:33.555: %HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Speak -> Standby
- *Mar 1 00:16:33.559: HSRP: Fa0/0 Grp 1 Redundancy "hsrp-Fa0/0-1" state Speak -> Standby

Chapter 4 – Maintaining & Troubleshooting Campus Switched LANs Objectives

- Describe VLAN and inter-VLAN troubleshooting techniques.
- Describe STP troubleshooting techniques.
- Describe FHRP troubleshooting techniques

Any Questions?