Chapter 4: Maintaining and Troubleshooting Campus Switched Solutions

CCNP TSHOOT: Maintaining and Troubleshooting IP Networks

Cisco | Networking Academy®
Mind Wide Open™

Troubleshooting VLANs

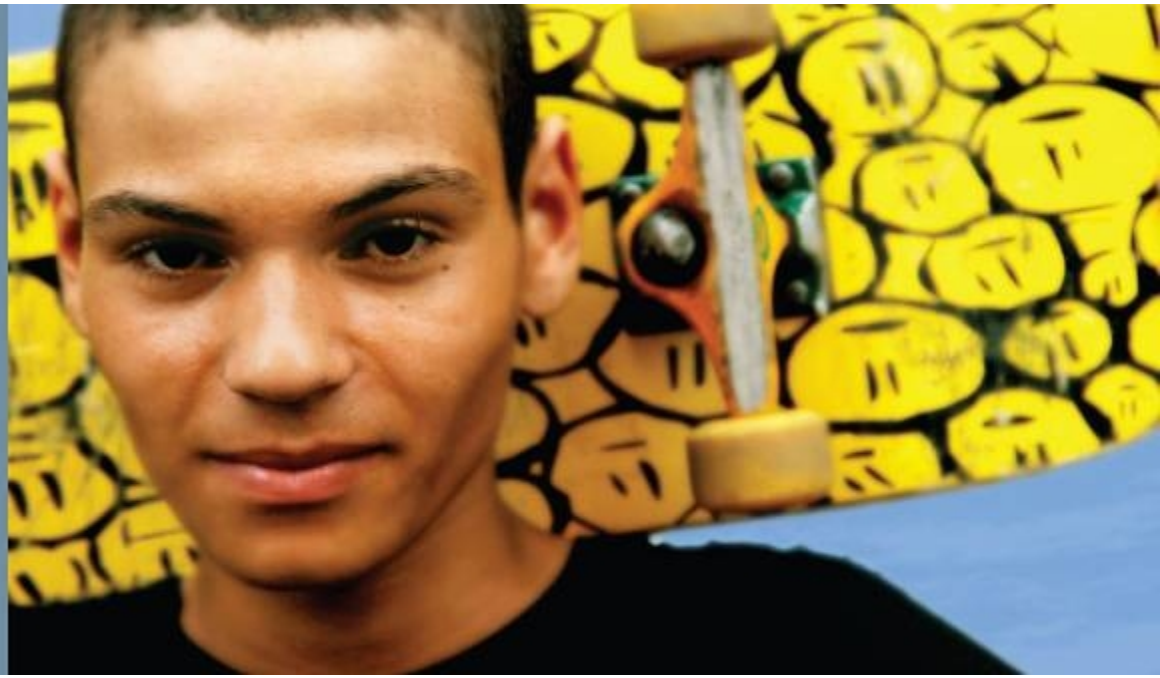

LAN Switch Operation

- Host A pings Host B on the same VLAN (subnet).
- Host A determines that destination (Host B) IP is on the same subnet.
- Host A consults its ARP cache, encapsulates the IP packet in an Ethernet frame and transmits the frame to Host B.

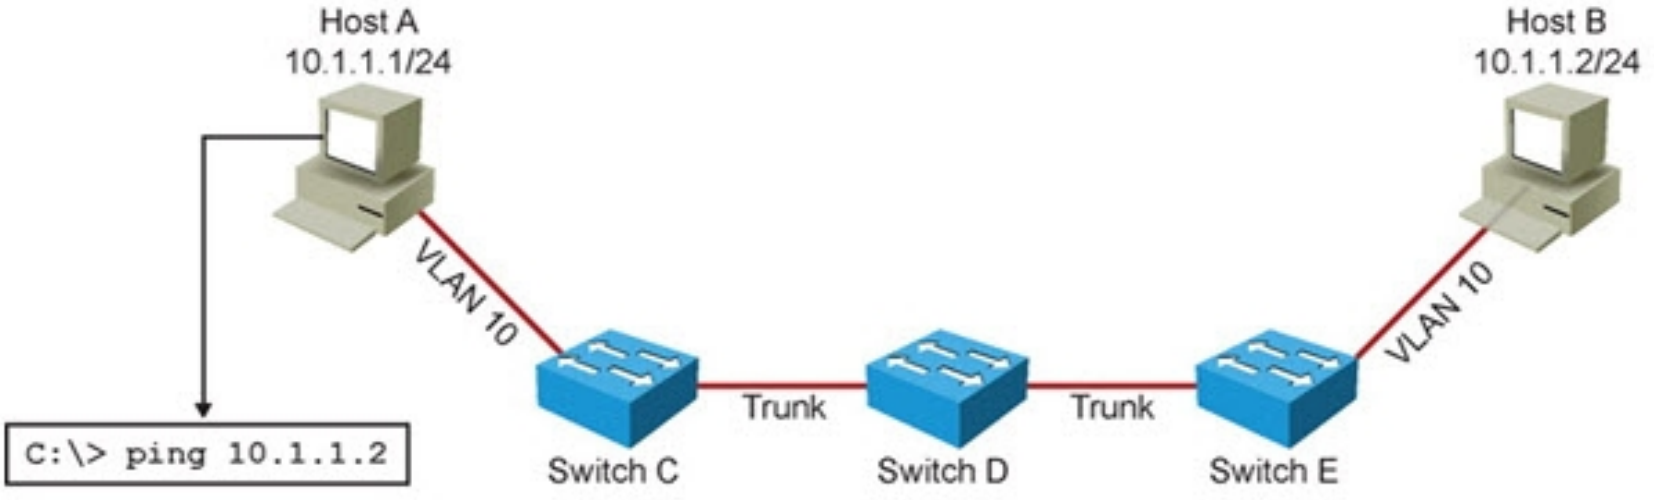

LAN Switch Operation – Cont.

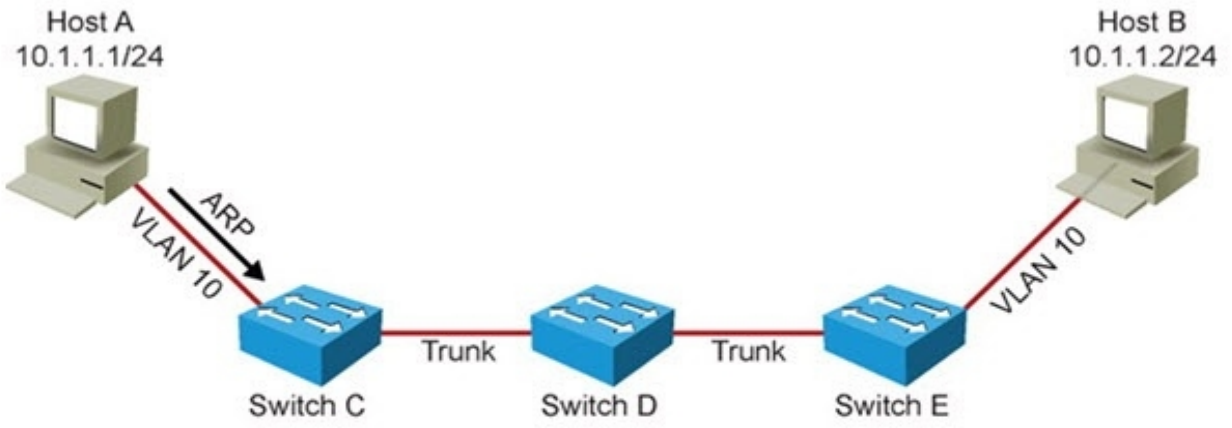

DMAC	SMAC	Type	Data	FCS
BCAST	MAC A	0x0806	ARP Request	CRC

- If Host A does not have an entry for Host B in its ARP cache, it will ARP for the Host B MAC address.

LAN Switch Operation – Cont.

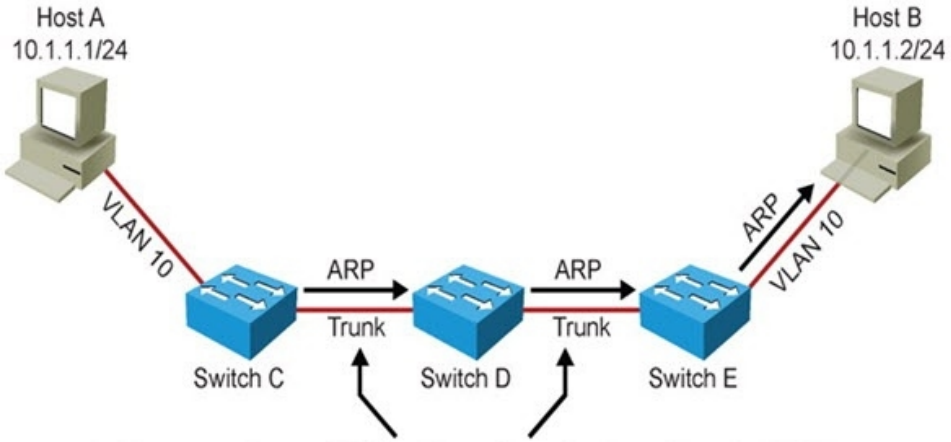

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
BCAST	MAC A	0x8100	VLAN 10	0x0806	ARP Request	CRC

- The intermediate switches flood the ARP request over the 802.1Q trunk links.

LAN Switch Operation – Cont.

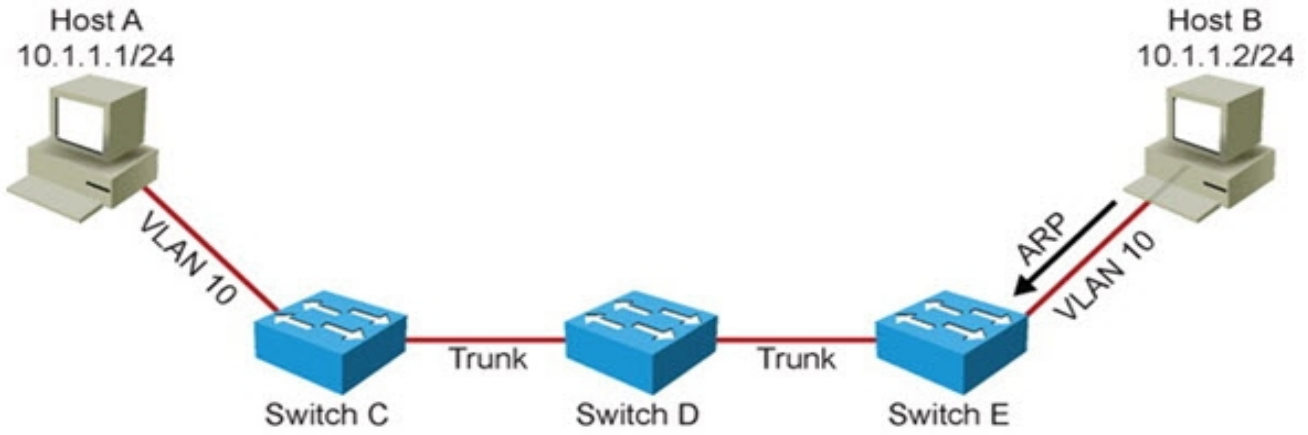

DMAC	SMAC	Type	Data	FCS
MAC A	MAC B	0x0806	ARP Reply	CRC

- Host B sends a unicast ARP reply back to Host A.

LAN Switch Operation – Cont.

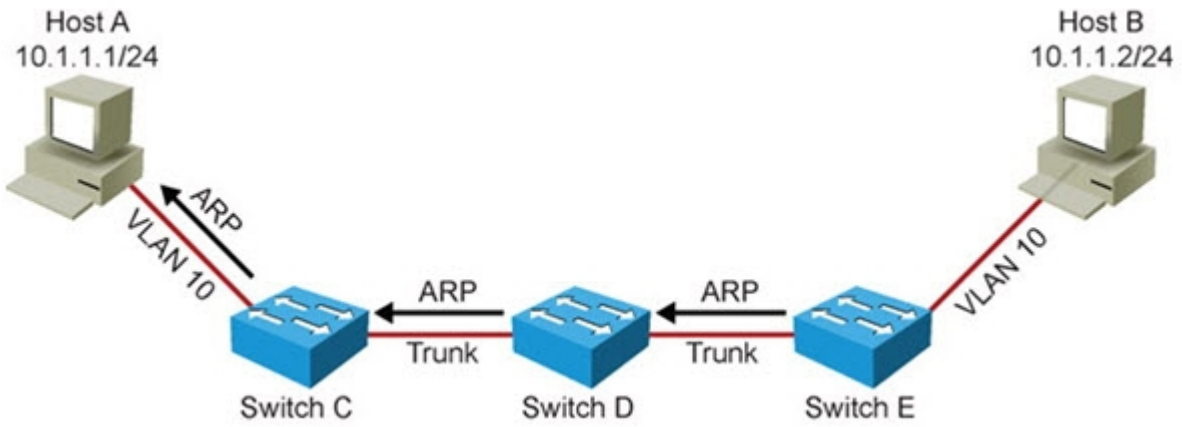

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
MAC A	MAC B	0x8100	VLAN 10	0x0806	ARP Reply	CRC

- Switches forward the ARP reply unicast frame toward Host A.

LAN Switch Operation – Cont.

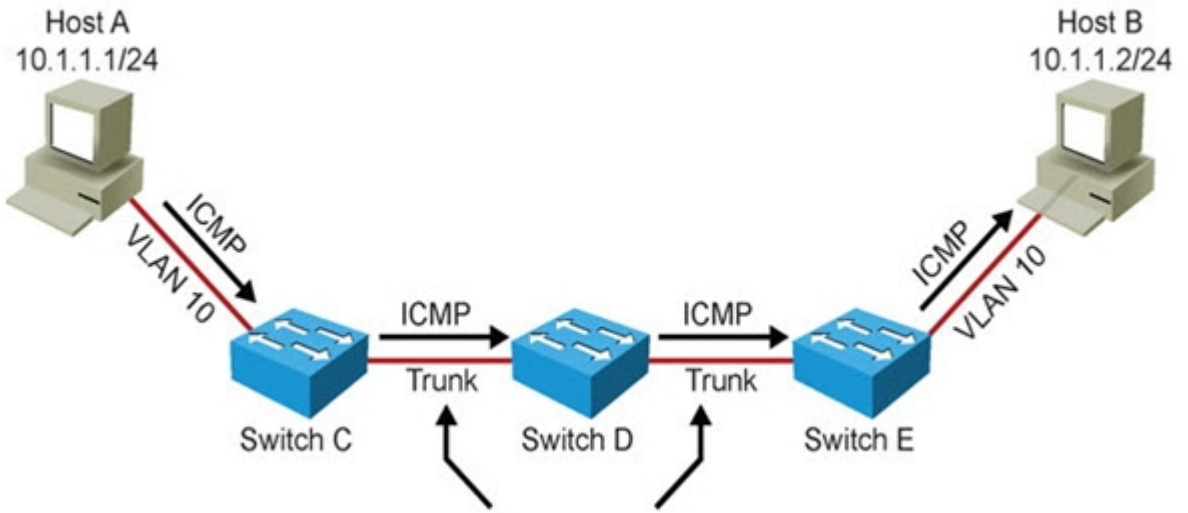

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
MAC B	MAC A	0x8100	VLAN 10	0x0800	ICMP Echo Request	CRC

- Host A encapsulates the IP packet (ICMP Echo Request) in a unicast frame and sends it to Host B.
- Switches forward ICMP Echo Request unicast frame toward Host B.

LAN Switch Operation – Cont.

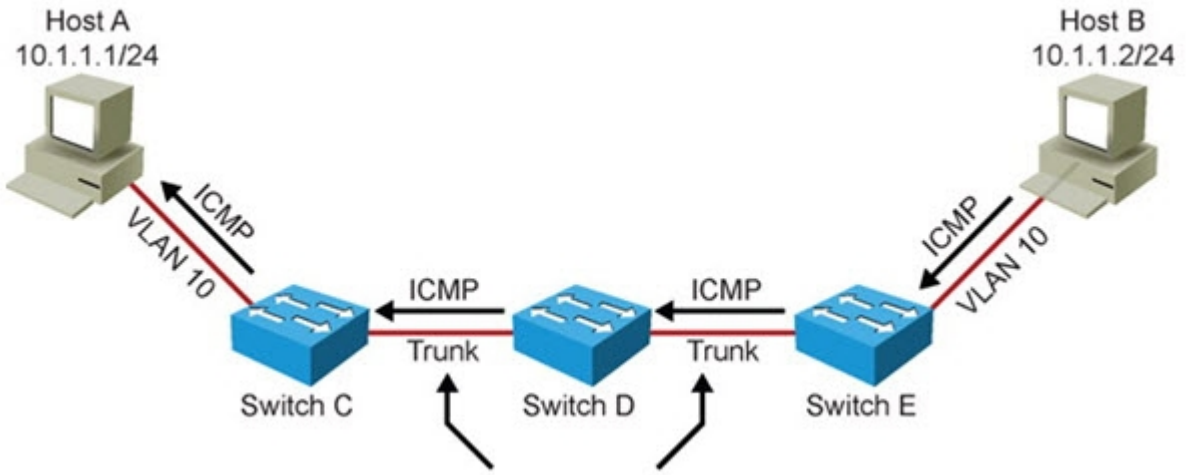

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
MAC A	MAC B	0x8100	VLAN 10	0x0800	ICMP Echo Reply	CRC

- Switches forward ICMP Echo Reply unicast frame toward Host A.

LAN Switch Operation – Cont.

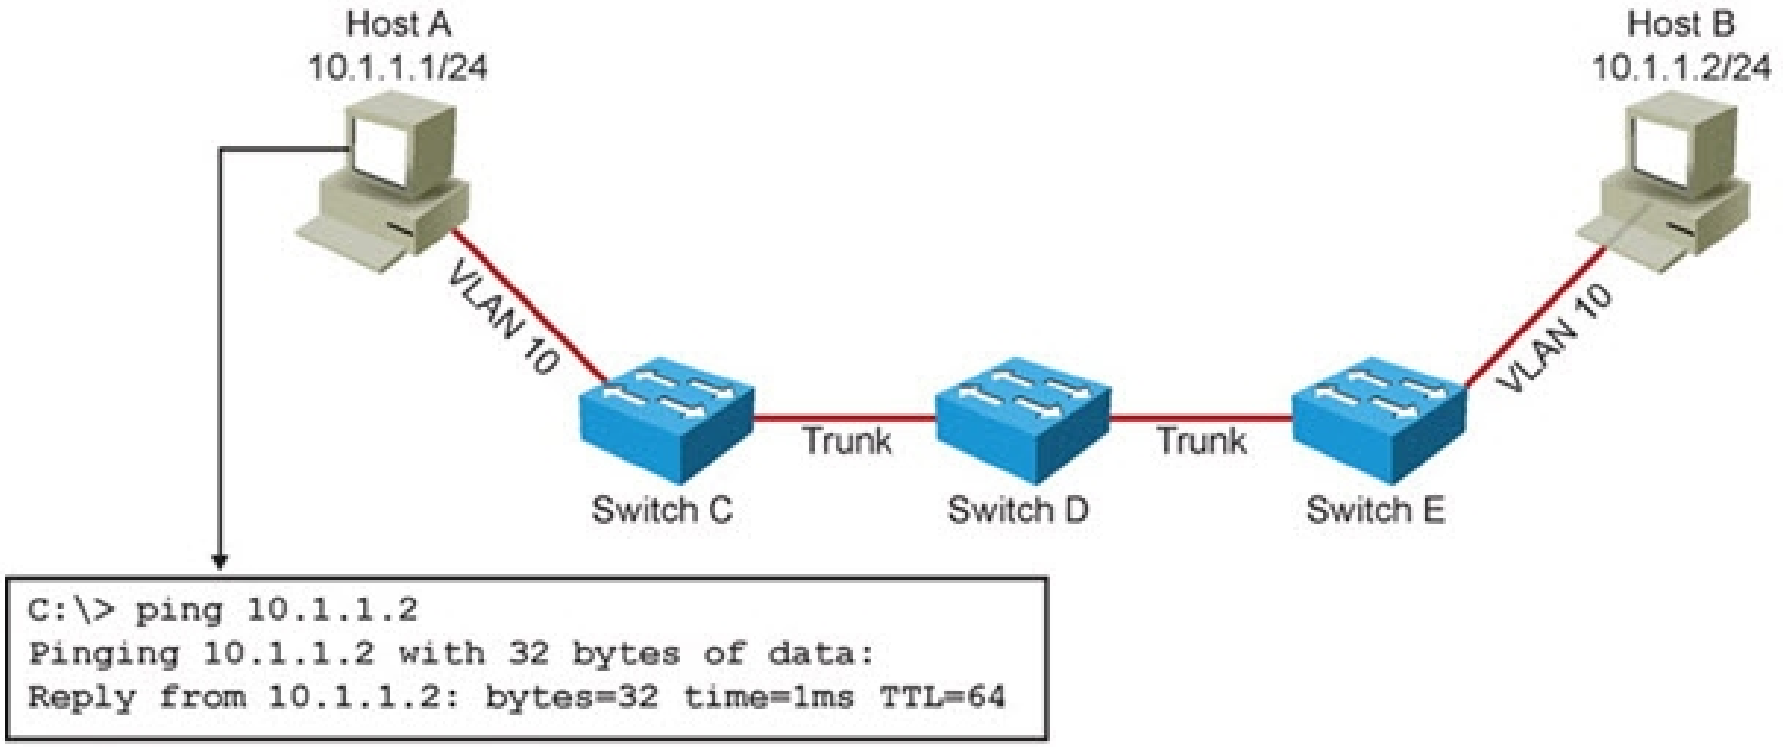

- Host A Receives ICMP Echo Reply Back from Host B.

LAN Switch Operation – Cont.

Issues that could cause the communication to fail:

- Physical problems
- Bad, missing, or miswired cables
- Bad ports
- Power failure
- Device problems
- Software bugs
- Performance problems
- Misconfiguration
- Missing or wrong VLANs
- Misconfigured VTP settings
- Wrong VLAN setting on access ports
- Missing or misconfigured trunks
- Native VLAN mismatch
- VLANs not allowed on trunk

Verifying Layer 2 Forwarding

Common findings when following the path of the frames through the switches:

- **Frames are not received on the correct VLAN**: This could point to VLAN or trunk misconfiguration as the cause of the problem.
- **Frames are received on a different port than you expected**: This could point to a physical problem, spanning tree issues, a native VLAN mismatch or duplicate MAC addresses.
- **The MAC address is not registered in the MAC address table:** This tells you that the problem is most likely upstream from this switch. Investigate between the last point where you know that frames were received and this switch.

Verifying Layer 2 Forwarding – Cont.

Useful Layer 2 diagnostic commands:

- **`show mac-address-table:`** Shows learned MAC addresses and corresponding port and VLAN associations.
- **`show vlan`**: Verifies VLAN existence and port-to-VLAN associations.
- **`show interfaces trunk`**: Displays all interfaces configured as trunks, VLANs allowed and what the native VLAN is.
- **`show interfaces switchport`**: Provides a summary of all VLAN related information for interfaces.
- **`show platform forward`** *`interface`*: Used to determine how the hardware would forward a frame.
- **`traceroute mac`**: Provides a list of switch hops (layer 2 path) that a frame from a specified source MAC address to a destination MAC address passes through. CDP must be enabled on all switches in the network for this command to work.
- **`traceroute mac ip`:** Displays Layer 2 path taken between two IP hosts.

Troubleshooting Spanning Tree

Spanning Tree Operation

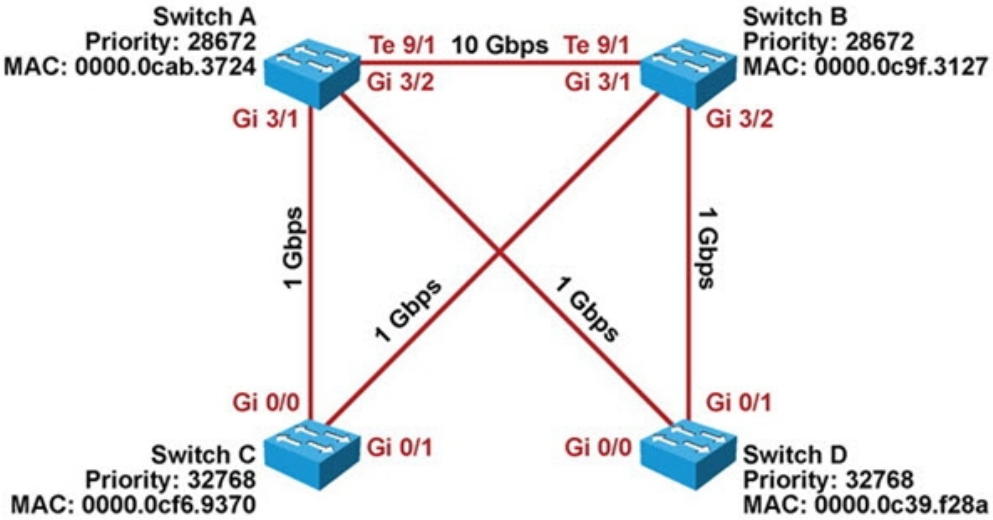

- Elect a Root Bridge/Switch.
- Select a Root Port on each Bridge/Switch (except on the Root bridge/switch).
- Elect a Designated device/port on each network segment.
- Ports that are neither Root Port nor a Designated Port go into Blocking state.

Spanning Tree Operation – Cont.

1. Elect a Root Bridge/Switch.

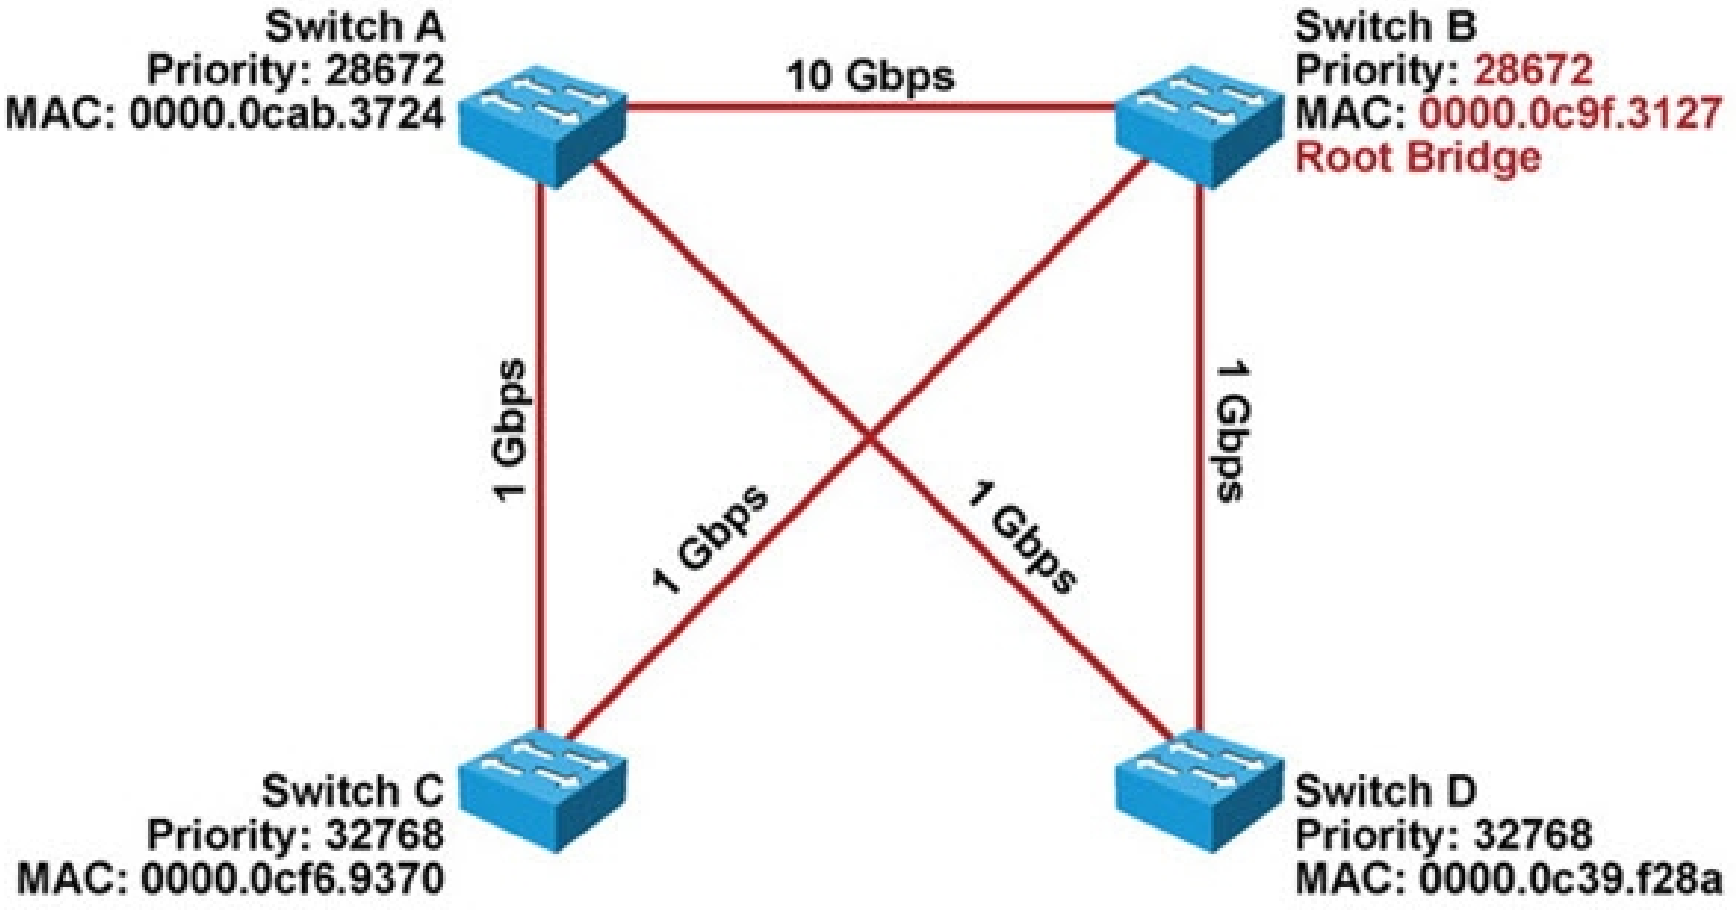

Spanning Tree Operation – Cont.

2. Select a Root Port on each bridge/switch.

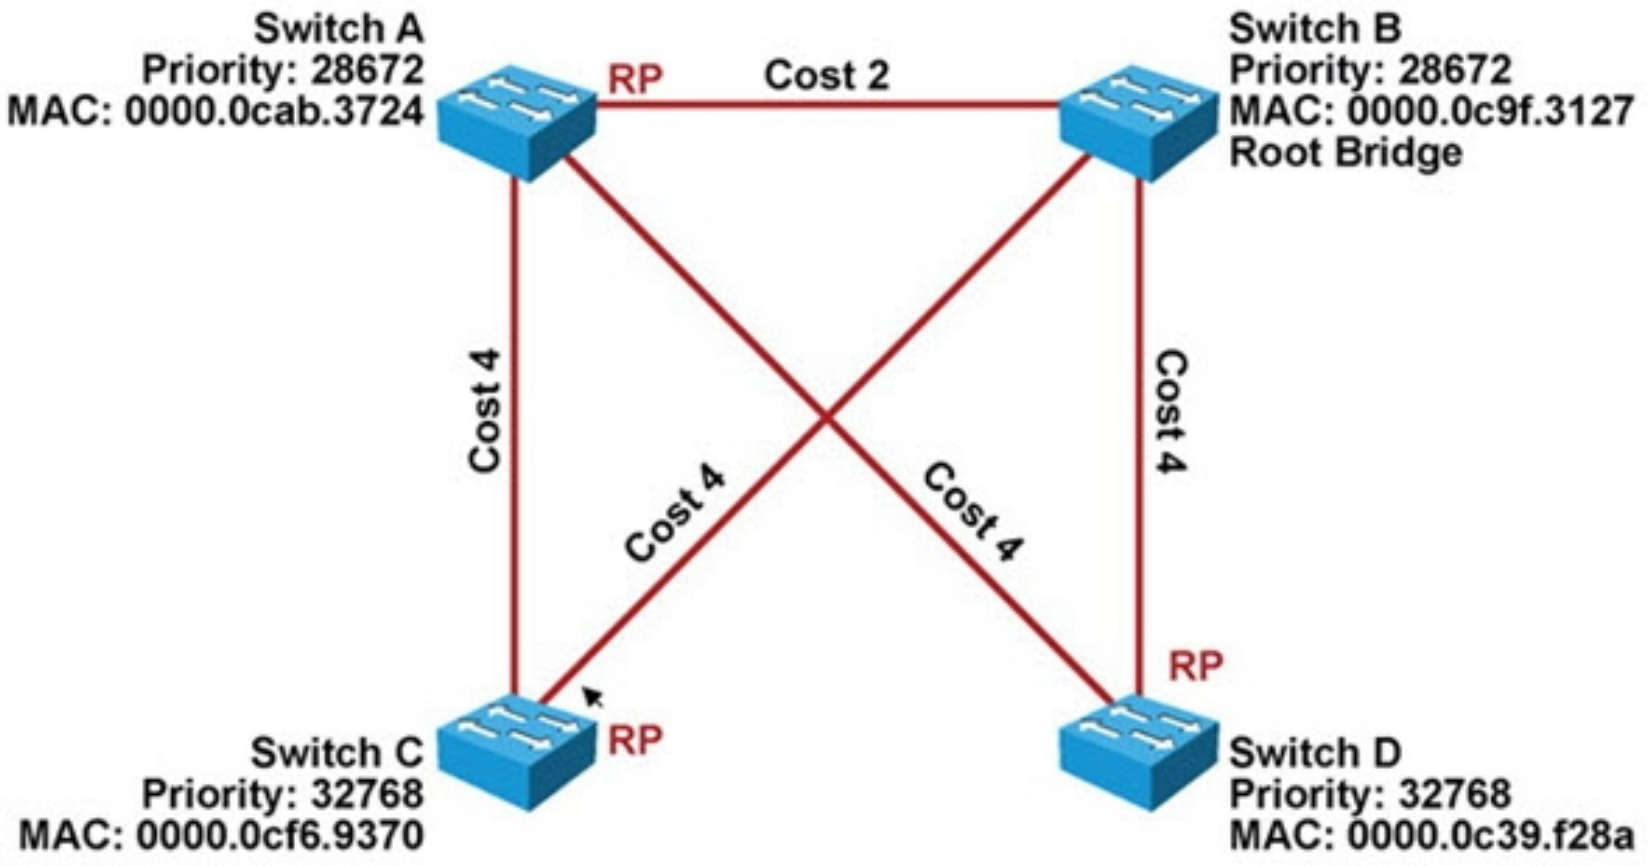

Spanning Tree Operation – Cont.

3. Elect a Designated device/port on each network segment.

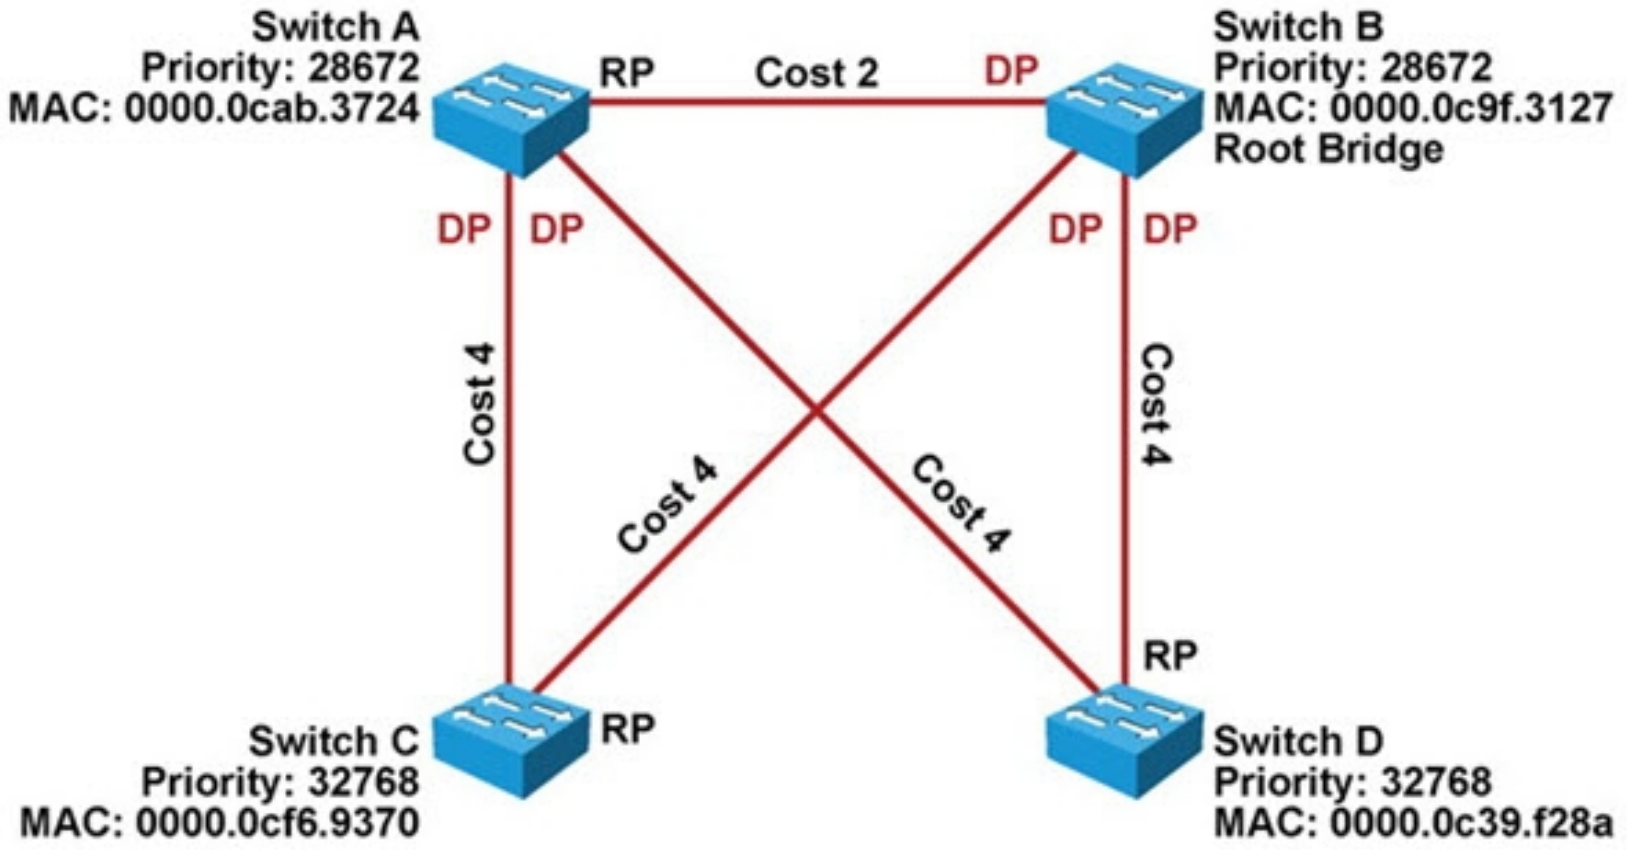

Spanning Tree Operation – Cont.

4. Place ports in Blocking state.

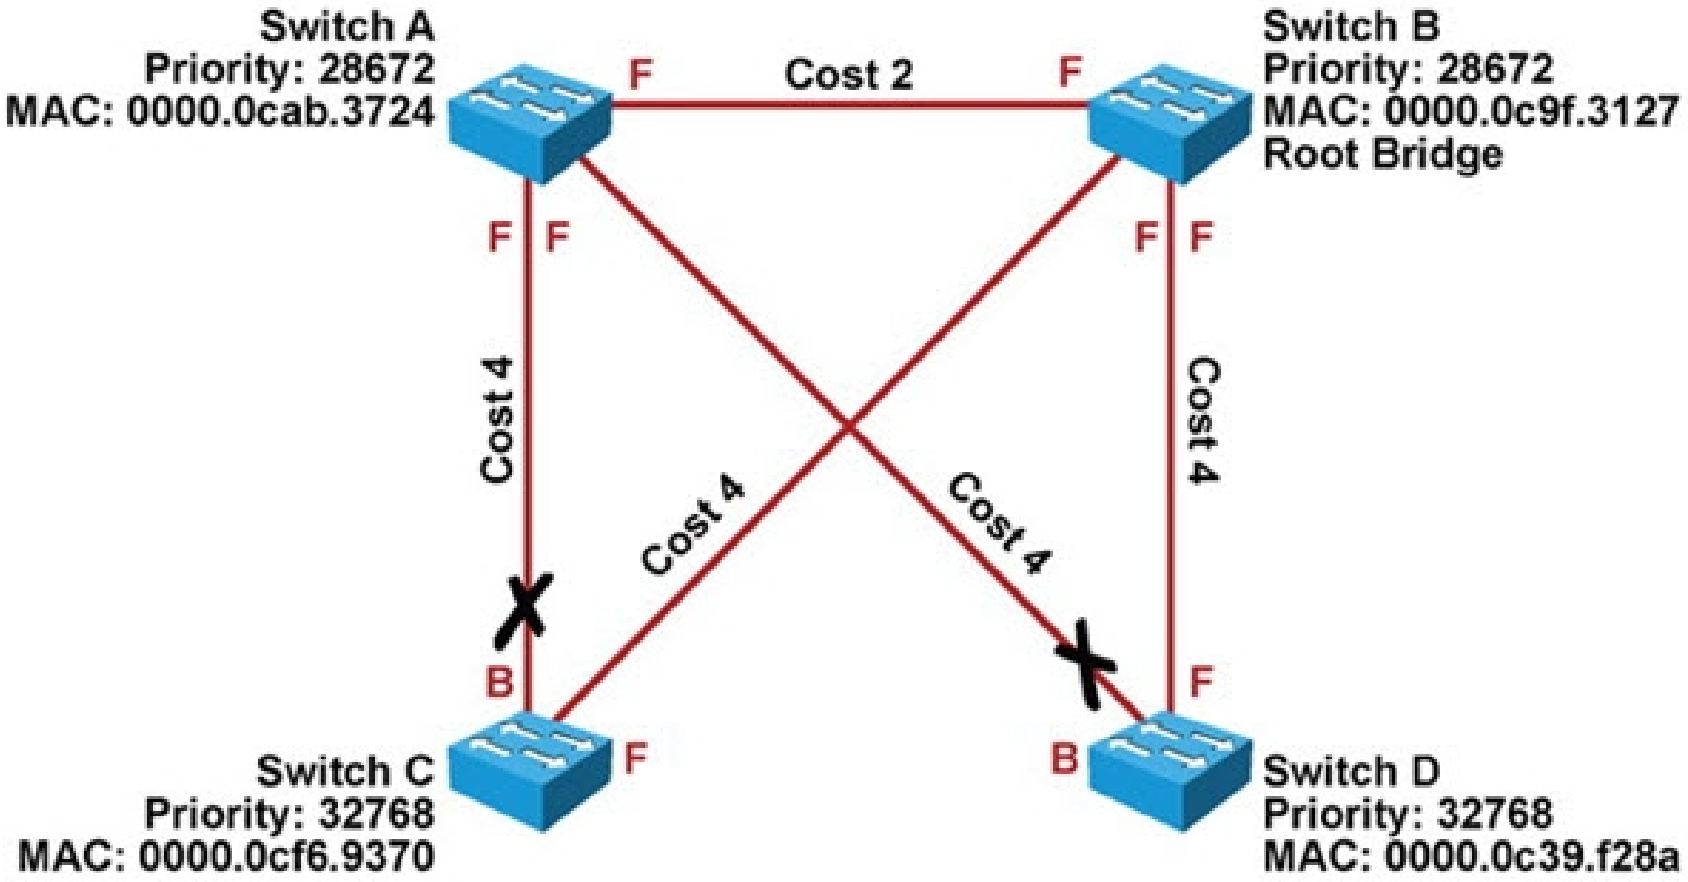

Spanning Tree Operation – Cont.

- Sample output from the **`show spanning-tree vlan`** command.

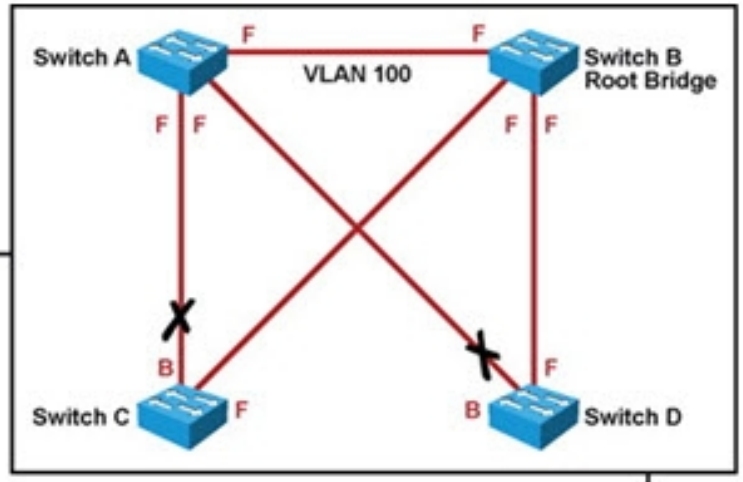


```
SwitchA#show spanning-tree vlan 100

VLAN0100
  Spanning tree enabled protocol rstp
  Root ID     Priority    28772
              Address     0000.0c9f.3127
              Cost        2
              Port        88 (TenGigabit9/1)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    28772  (priority 28672 sys-id-ext 100)
              Address     0000.0cab.3724
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Gi3/1            Desg FWD 4         128.72   P2p
Gi3/2            Desg FWD 4         128.80   P2p
Te9/1            Root FWD 2         128.88   P2p
```

Spanning Tree Operation – Cont.

- Sample output from the **`show spanning-tree interface`** command.

```
SwitchA#show spanning-tree interface Ten 9/1 detail
 Port 88 (TenGigabitEthernet9/1) of VLAN0100 is root forwarding
   Port path cost 2, Port priority 128, Port Identifier 128.88.
   Designated root has priority 28772, address 0000.0c9f.3127
   Designated bridge has priority 28772, address 0000.0c9f.3127
   Designated port id is 128.88, designated path cost 0
   Timers: message age 15, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 10, received 670
```

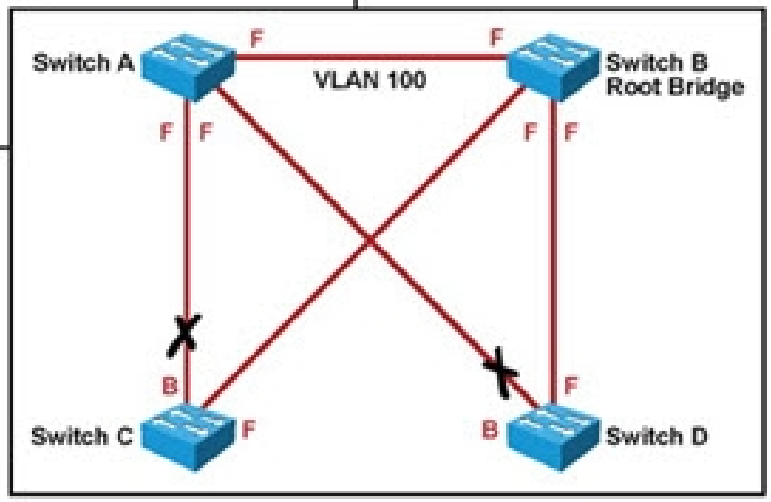

Spanning Tree Failures

- STP is a reliable but not an absolutely failproof protocol.
- If STP fails there are usually major negative consequences.
- With Spanning Tree, there are two different types of failures.
  - Type 1 - STP may erroneously block certain ports that should have gone to the forwarding state. You may lose connectivity to certain parts of the network, but the rest of the network is unaffected.
  - Type 2 - STP erroneously moves one or more ports to the Forwarding state. The failure is more disruptive as bridging loops and broadcast storms can occur.

Spanning Tree Failures – Cont.

- Type 2 failures can cause these symptoms.
  - The load on all links in the switched LAN will quickly start increasing.
  - Layer 3 switches and routers report control plane failures such as continual HSRP, OSPF and EIGRP state changes or that they are running at a very high CPU utilization load.
  - Switches will experience very frequent MAC address table changes.
  - With high link loads and CPU utilization devices typically become unreachable, making it difficult to diagnose the problem while it is in progress.
- Eliminate topological loops and troubleshoot issues.
  - Physically disconnect links or shut down interfaces.
  - Diagnose potential problems.
  - A unidirectional link can cause STP problems. You may be able to identify and remove a faulty cable to correct the problem.

EtherChannel Operation

- EtherChannel bundles multiple physical Ethernet links (100 Mbps,1 Gbps, 10 Gbps) into a single logical link.
- Traffic is distributed across multiple physical links as one logical link.
- This logical link is represented in Cisco IOS syntax as a “Port-channel” (Po) interface.
- STP and routing protocols interact with this single port-channel interface.
- Packets and frames are routed or switched to the port-channel interface.
- A hashing mechanism determines which physical link will be used to transmit them.

EtherChannel Problems

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 EtherChannel link
3. Uneven distribution of traffic between EtherChannel bundle members

Troubleshooting Example: Switch Replacement Gone Bad

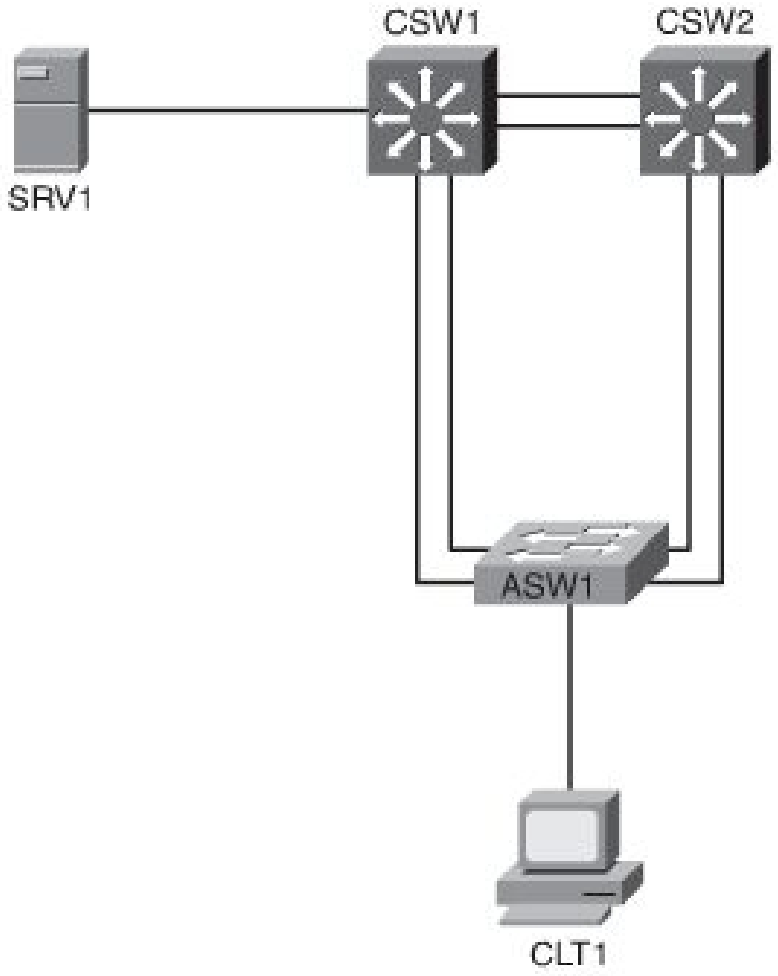

- A broken access switch ASW1 has been replaced but has 3 problems.
- On CSW2, port channel 1, which connects to ASW1, is down.
- On ASW1, the console log indicates a STP problem on Po2: %SPANTREE-2-PVSTSIM_FAIL: Blocking designated port Po2: Inconsistent superior PVST BPDU received on VLAN 17,claiming root 24593:001f.2721.8400.
- On ASW1, interface VLAN 128 is down.

EtherChannel Diagnostic Commands

Using the **`show etherchannel summary`** command

```
DSW2# 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 Diagnostic Commands

Using the **`show etherchannel 1 detail`** command

```
DSW2# 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

<output omitted>
```

EtherChannel Diagnostic Commands

```
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)
```

- You could also find the problem indication from the log, which contains the messages shown in this graphic.
- With this information, you would compare the port-channel interface to the physical interfaces to find out that the VLAN allowed list is missing on physical interfaces Fa0/5 and Fa0/6of ASW1.
- The junior staff member should change the configuration by adding allowed VLAN list to physical interfaces Fa0/5 and Fa0/6.

Spanning Tree Diagnostics

Using the **`show spanning-tree`** command to examine STP

```
ASW1# show spanning-tree vlan 17

MST0
  Spanning tree enabled protocol mstp
  Root ID     Priority    32768
              Address     001e.79a9.b580
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32768  (priority 32768 sys-id-ext 0)
              Address     001e.79a9.b580
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Fa0/7               Desg FWD 200000    128.9    P2p Edge
Po1                 Desg BLK 100000    128.56   P2p
Po2                 Desg BKN*100000    128.64   P2p Bound(PVST) *PVST_Inc
```

Spanning Tree Diagnostics

```
ASW1# show ip interfaces brief ¦ exclude unassigned
Interface                  IP-Address      OK? Method Status                Protocol
Vlan128                    10.1.156.1      YES NVRAM  up                    down
```

- For the third problem, checking the status of VLAN128, the graphic shows that the VLAN interface 128 is indeed down (and not administratively).
- A VLAN interface is up as long as the VLAN exists and there is an active port in that VLAN that is in spanning-tree forwarding state.

Spanning Tree Diagnostics

```
ASW1# show spanning-tree vlan 128
Spanning tree instance(s) for vlan 128 does not exist.



ASW1# show vlan id 128
VLAN id 128 not found in current VLAN database
```

- When a VLAN interface is down, it is a good idea to first check the spanning-tree status for that VLAN.
- Using the **show spanning tree vlan 128** and the **show vlan id 128** commands, you will discover that spanning tree is not running for VLAN 128., as shown in the graphic.
- That leads to the theory that VLAN 128 probably does not exist on ASW1, which is confirmed. VLAN 128 is add aand problem solved

Troubleshooting Switched Virtual Interfaces and Inter-VLAN Routing

Inter-VLAN Routing and Multilayer Switching

- Similarities between multilayer switches and routers
- Both routers and multilayer switches use routing protocols or static routes to maintain information about the reachability and direction to network destinations (prefixes), and record this information in a routing table.
- Both routers and multilayer switches perform the same functional packet switching actions:
  - They receive a frame and strip off the Layer 2 header.
  - They perform a Layer 3 lookup to determine the outbound interface and next hop.
  - They encapsulate the packet in a new Layer 2 frame and transmit the frame.

Troubleshooting Routers and Multi-Layer Switches

Sample Data Plane and Control Plane commands for routers and multi-layer switches

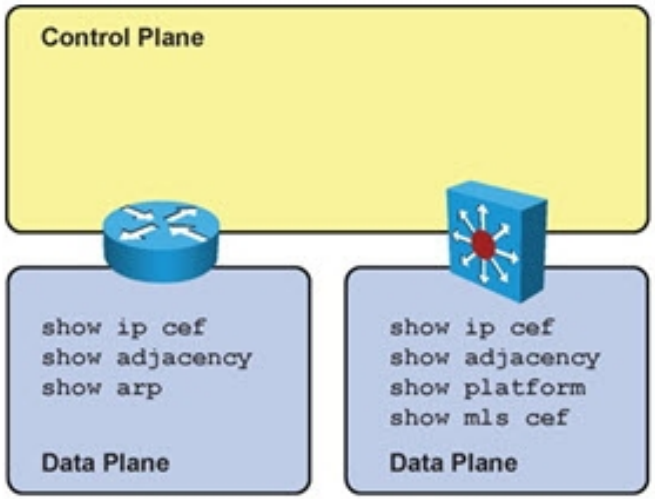

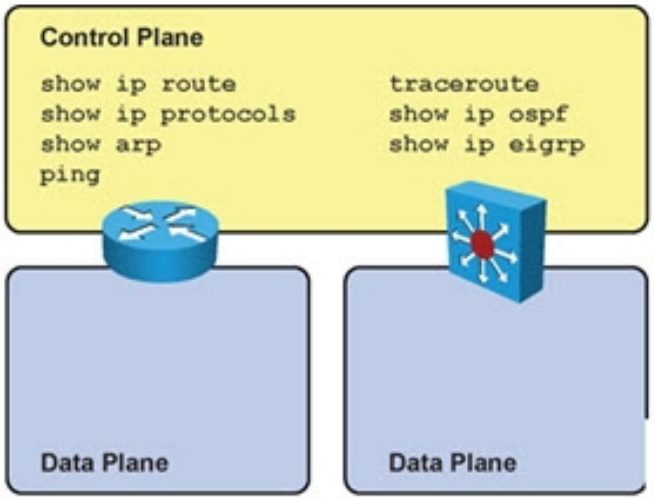

Cisco 7206

Catalyst 6504

Troubleshooting Routers and Multi-Layer Switches – Cont.

Commands to check the CEF data structures for routers and multi-layer switches.

`show ip cef`

- Displays the content of the CEF FIB.
  - The FIB reflects the content of the routing table with all the recursive lookups resolved already and the output interface determined for each destination prefix.
  - The FIB also holds additional entries for directly connected hosts, the router's own IP addresses, and multicast and broadcast addresses.

`show adjacency`

- Displays the content of the CEF adjacency table.
  - This table contains preconstructed Layer 2 frame headers with all necessary fields already filled in. These frame headers are used to encapsulate the egress CEF-switched packets and deliver them to appropriate next hop devices..

Troubleshooting Multi-layer Switches

Commands to check forwarding behavior of switches from the content of TCAM on Catalyst switches:

`show platform`

- On the Catalyst 3560, 3750 and 4500 platforms, the show platform family of commands can be used to obtain detailed information about the forwarding behavior of the hardware.

`show mls cef`

- On the Catalyst 6500 platform, the show mls cef family of commands can be used to obtain detailed information about the forwarding behavior of the hardware.

Switched Virtual Interfaces (SVIs) and Routed Ports

A Logical Demonstration of a Multilayer Switch:

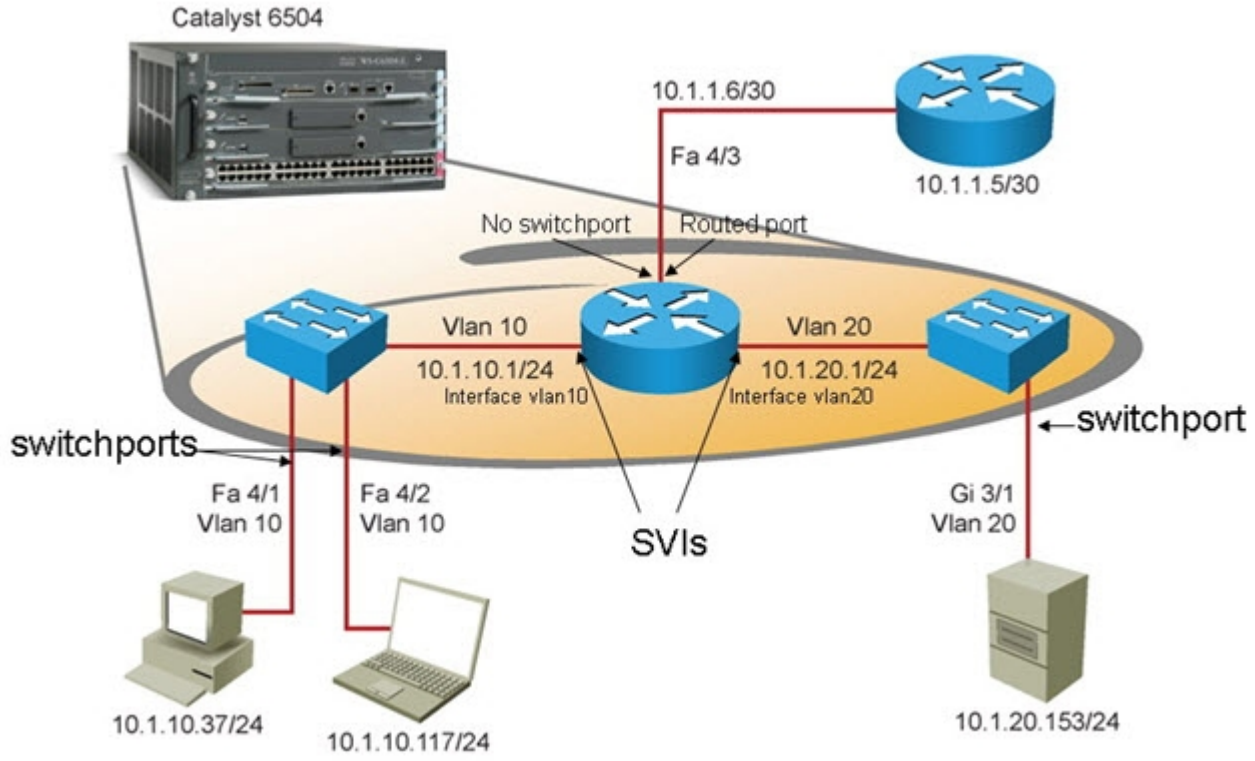

Multilayer Switches, VLANs and Routing

A multilayer switch provides three different core functions in a single device:

1. Layer 2 switching within each VLAN:
   - The traffic switched between ports that belong to the same VLAN
   - The MAC address tables for different VLANS are logically separated.
   - No IP or Layer 3 configuration is necessary.
1. Routing and multilayer switching between the local VLANs:
   - Layer 3 switching between VLANs requires SVIs
   - Each SVI requires an appropriate IP address and subnet mask.
   - Hosts on the can use the SVI's IP address as default gateway.
   - IP routing must be enabled.

Multilayer Switches, VLANs and Routing – Cont.

3. Routing and multilayer switching between the local VLANs and one or more routed interfaces:

- A regular physical switched port can be made a routed port.
- A routed interface does not belong to any user-created or default VLAN and has no dependency on VLAN status (unlike an SVI).
- Traffic on this port is not bridged (switched) to any other port
- There is no MAC address table associated to it.
- The port acts like a regular router interface and needs its own IP address and subnet mask.

Routed Interfaces vs. SVIs

The main differences between SVIs and routed interfaces are:

- A routed interface is not a Layer 2 port – Layer 2 protocols, such STP and DTP are not active.
- The status of a routed interface is directly related to the availability of the corresponding directly-connected subnet.
- If a routed interface goes down, the corresponding connected route will immediately be removed from the routing table.
- An SVI is not a physical interface so it generally doesn't fail.
- An SVIs status is directly dependent on the status of the VLAN with which it is associated. The VLAN must be defined in the VLAN database.
- 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.
- An SVI can only go down when the last active port in the VLAN goes down or loses its Spanning Tree forwarding status (and the corresponding connected subnet will be removed from the routing table).

Troubleshooing First Hop Redundancy Protocols

First Hop Redundancy Protocols (FHRPs)

- FHRP is an important element in building highly available networks.
- Clients and servers normally point to a single default gateway and lose connectivity to other subnets if their gateway fails.
- FHRPs provide redundant default gateway functionality that is transparent to the end hosts.
- These protocols provide a virtual IP address and the corresponding virtual MAC address.
- Examples of FHRPs include:
  - Hot Standby Router Protocol (HSRP) – Cisco
  - Virtual Router Redundancy Protocol (VRRP) – IETF standard
  - Gateway Load Balancing Protocol (GLBP) – Cisco
- The mechanisms of these protocols revolve around these functions:
  - Electing a single router that controls the virtual IP address
  - Tracking availability of the active router
  - Determining if control of the virtual IP and MAC addresses should be handed over to another router

Using First Hop Redundancy

Sample HSRP Configuration

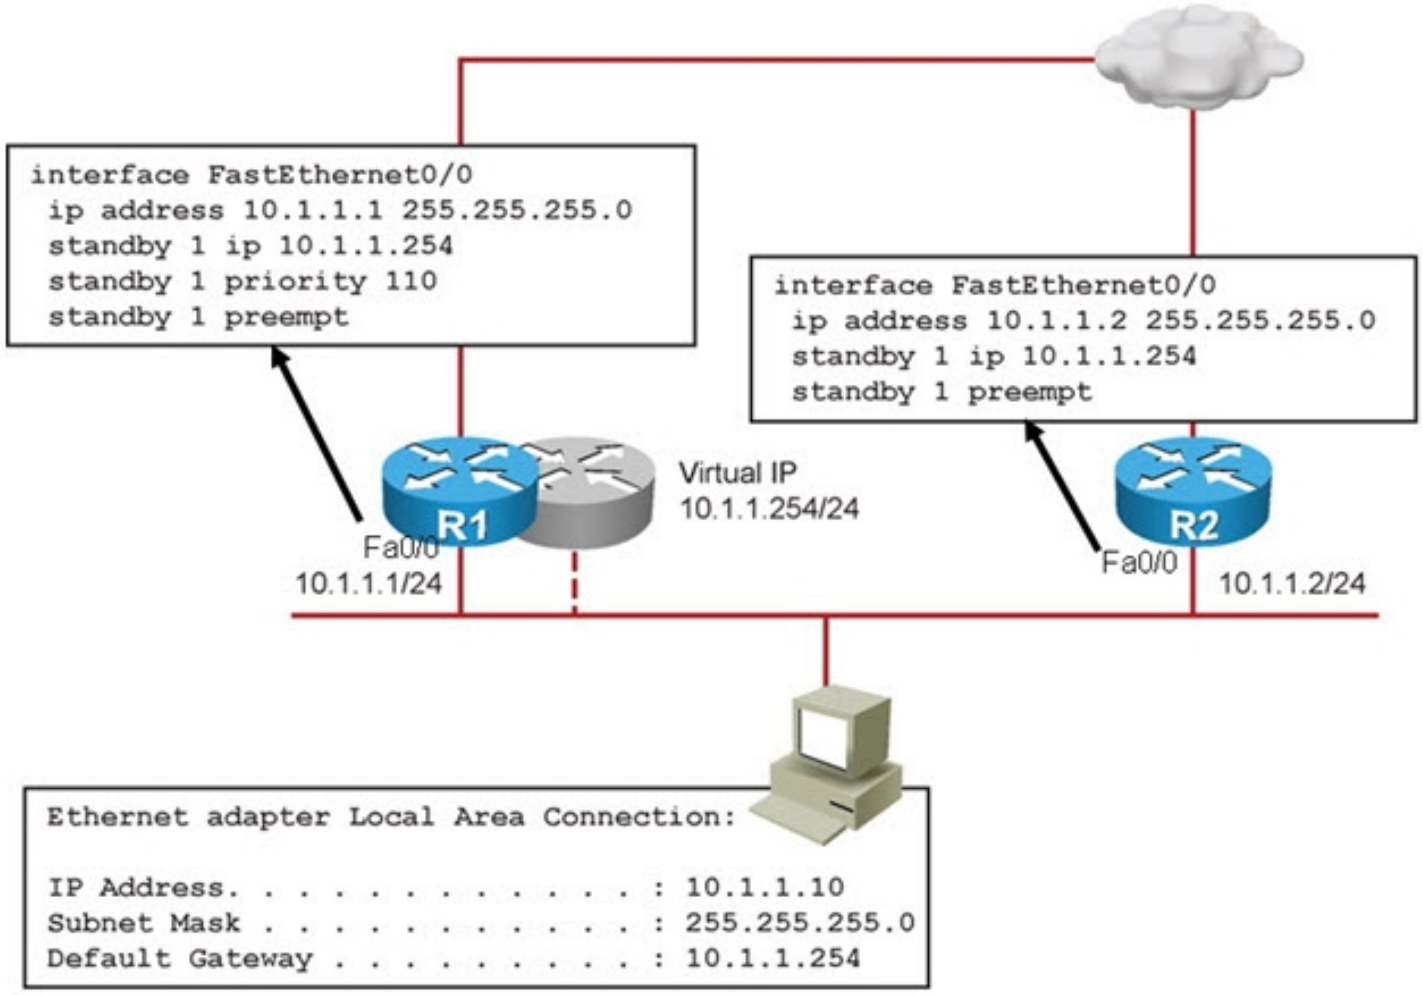

Verifying HSRP Operation

Sample output from the **`show standby brief`** command

```
R1#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0       1   110  P Active   local           10.1.1.2        10.1.1.254
```

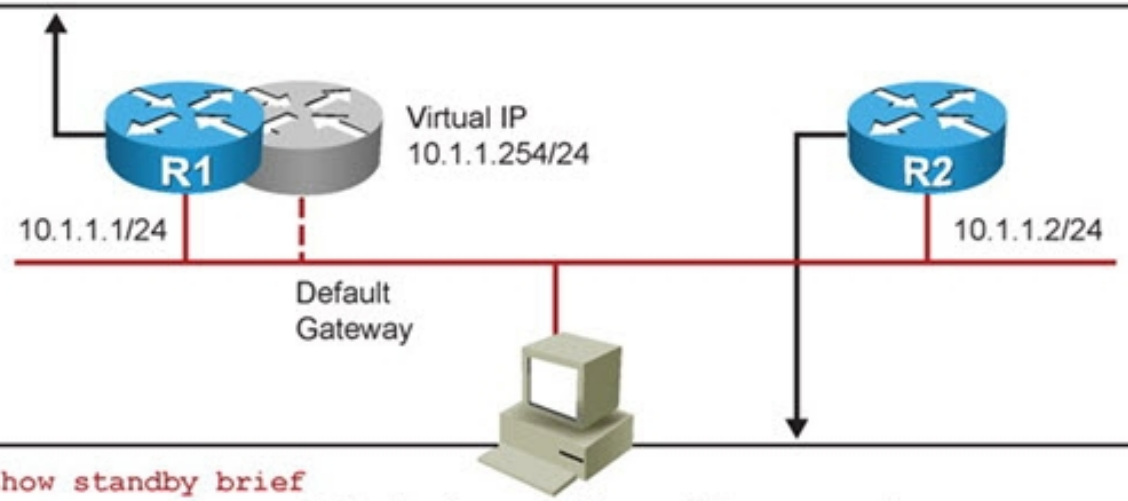


```
R2#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0       1   100  P Standby  10.1.1.1        local           10.1.1.254
```

Verifying HSRP Operation – Cont.

Sample output from the **`show standby`** *`interface-id`* command

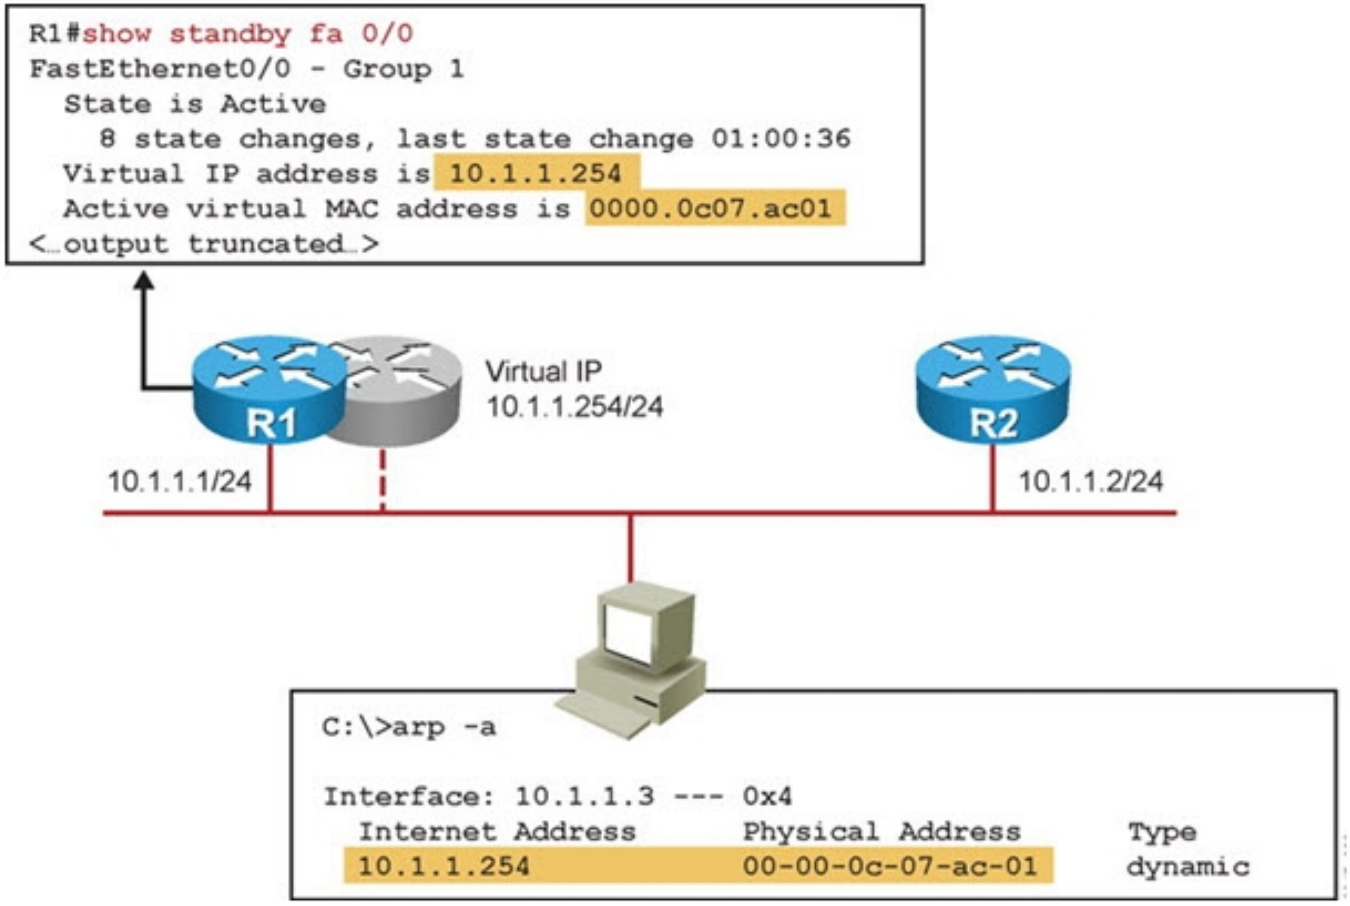

Verifying HSRP Operation – Cont.

The interface of a router participating in HSRP is shutdown.

```
interface FastEthernet0/0
 ip address 10.1.1.1 255.255.255.0
 standby 1 ip 10.1.1.254
 standby 1 priority 110
 standby 1 preempt
 shutdown
```

```
interface FastEthernet0/0
 ip address 10.1.1.2 255.255.255.0
 standby 1 ip 10.1.1.254
 standby 1 preempt
```

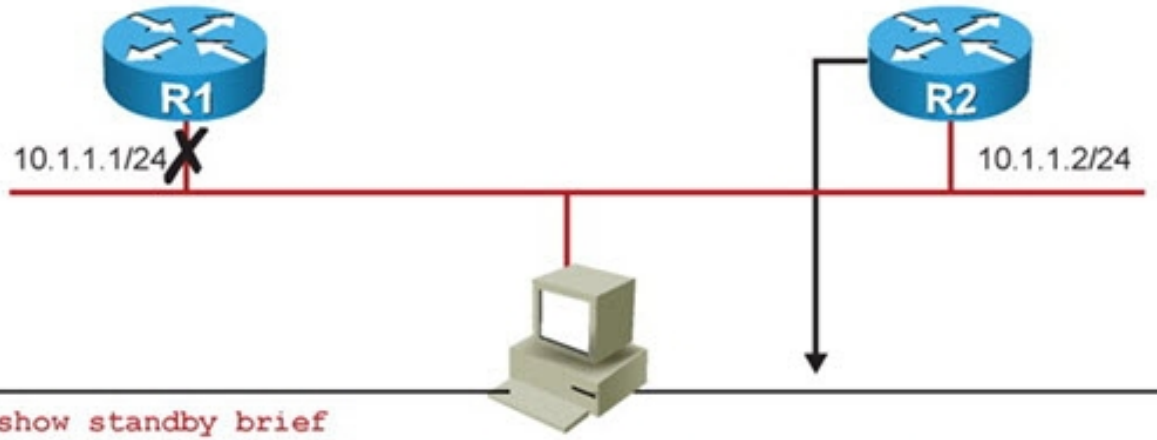


```
R2#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0       1   100  P Active   local           unknown         10.1.1.254
```

Verifying HSRP Operation – Cont.

While **`debug standby terse`** is enabled on R2, R1's interface is enabled.

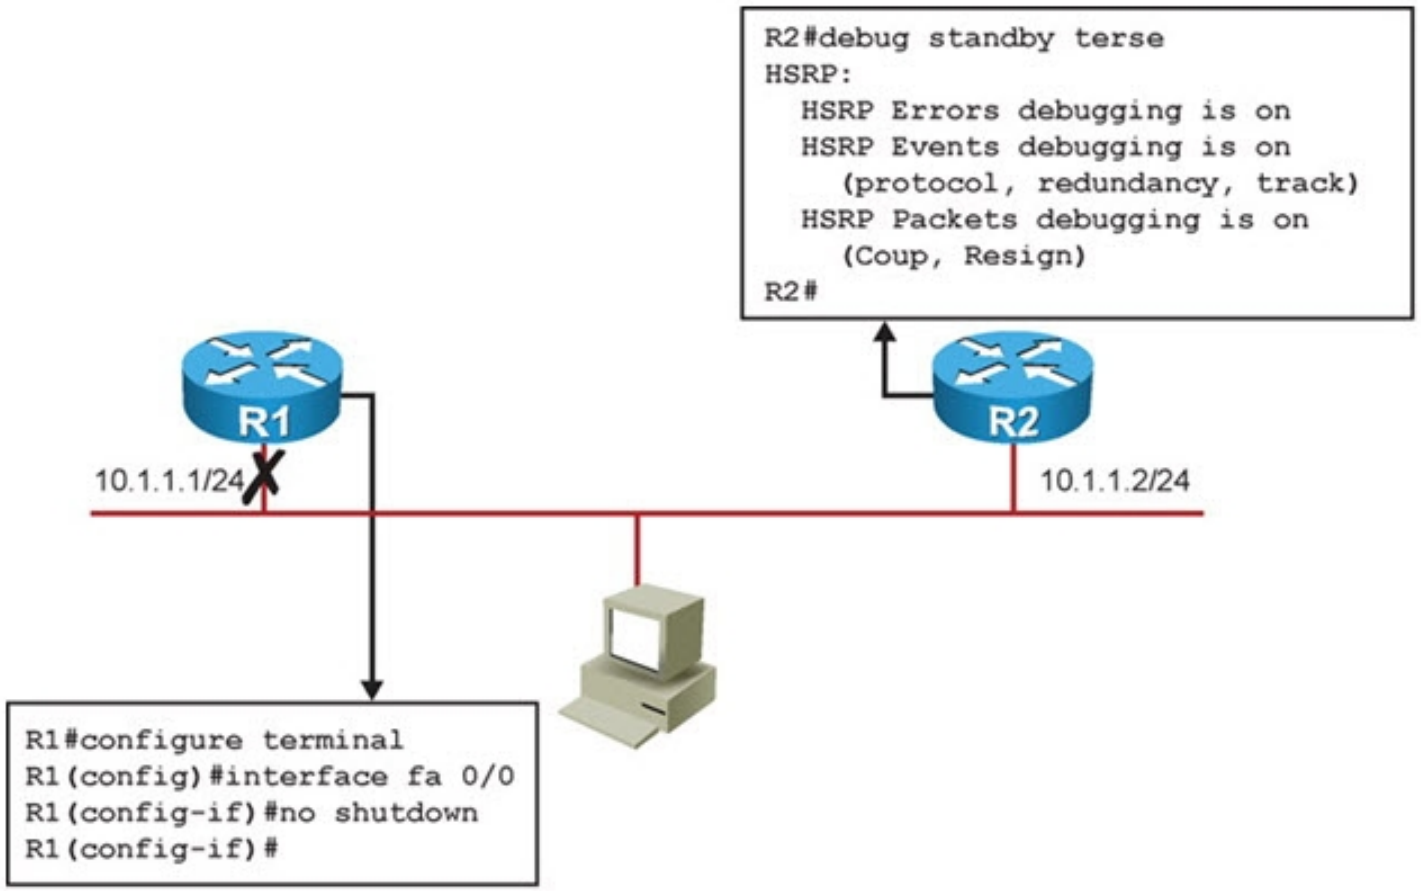

Verifying HSRP Operation – Cont.

Output of **`debug standby terse`** on R2 as R1's interface is enabled

```
R2#
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Coup   in  10.1.1.1 Listen  pri 110
vIP 10.1.1.254
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active: j/Coup rcvd from higher pri
router (110/10.1.1.1)
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active router is 10.1.1.1, 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
R2#
```

Alternatives to HSRP

Operational differences between HSRP, VRRP, and GLBP

Feature	HSRP	VRRP	GLBP
Transparent default gateway redundancy	Yes	Yes	Yes
Virtual IP address can also be a real address	No	Yes	No
IETF standard	No	Yes	No
Preempt is enabled by default	No	Yes	No
Multiple forwarding routers per group	No	No	Yes
Default Hello timer (seconds)	3	1	3

HSRP, VRRP, and GLBP Diagnostic Commands

Output of basic **show** commands for HSRP, VRRP, and GLBP

```
R1# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0       1   110  P Active   local           10.1.1.2        10.1.1.254
…
R1# show vrrp brief
Interface          Grp Pri Time  Own Pre State   Master addr     Group addr
Fa0/0              1   110 3570        Y  Master  10.1.1.1        10.1.1.254
…
R1# show glbp brief
Interface   Grp  Fwd Pri State    Address         Active router   Standby
router
Fa0/0       1    -   110 Active   10.1.1.254      local           10.1.1.2
Fa0/0       1    1   -   Active   0007.b400.0101  local           -
Fa0/0       1    2   -   Listen   0007.b400.0102  10.1.1.2        -
```

HSRP, VRRP, and GLBP Diagnostic Commands

Main Troubleshooting Commands for HSRP, VRRP, and GLBP

HSRP	VRRP	GLBP
`show standby brief`	`show vrrp brief`	`show glbp brief`
`show standby` *interface-id*	`show vrrp interface` *interface-id*	`show glbp` *interface-id*
`debug standby terse`	**No real equivalent option exits. Multiple `debug` options must be used simultaneously.**	`debug glbp terse`