Troubleshooting Campus Switched Solutions

TSHOOT Module 4

Chapter 4 Objectives

- VLAN, VTP, and trunking problems
- Spanning tree and EtherChannel problems
- Problems with SVIs and inter-VLAN routing
- Problems related to first hop redundancy protocols such as HSRP, VRRP, and GLBP

Troubleshooting VLANs

LAN Switch Operation ①

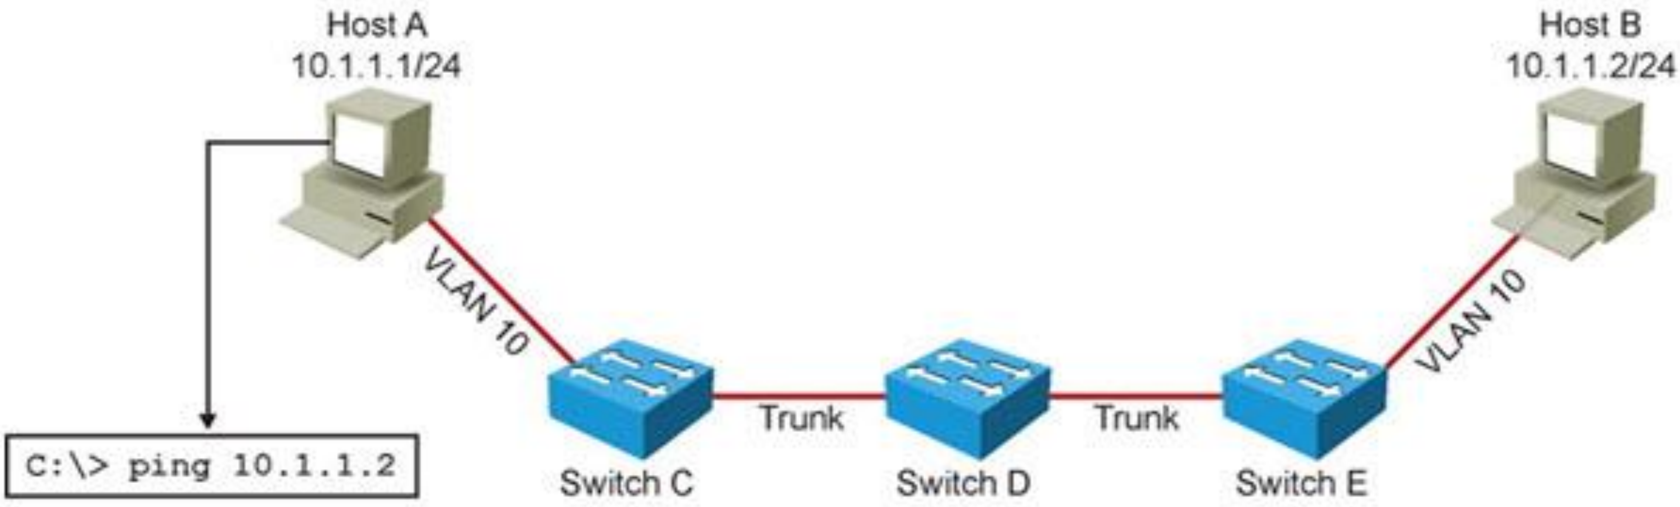

- Host A pings Host B on the same VLAN (subnet).
- Host A determines that destination (Host B) IP is on the same subnet.

LAN Switch Operation ②

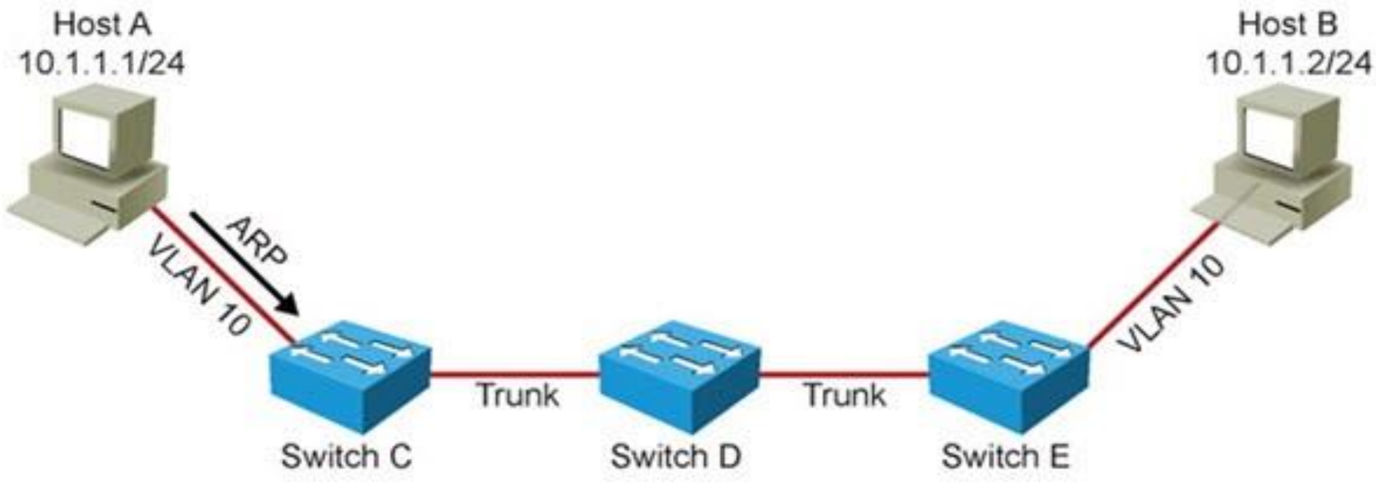

DMAC	SMAC	Type	Data	FCS
BCAST	MAC A	0x0806	ARP Request	CRC

- 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 ③

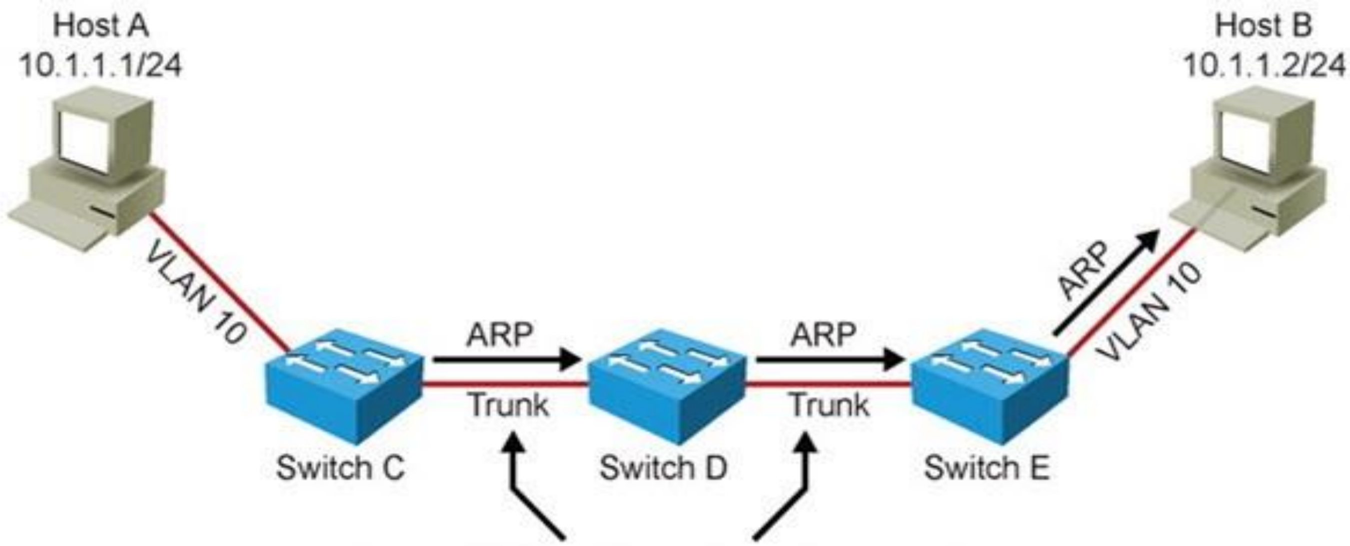

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
BCAST	MAC A	0x8100	VLAN 10	0x0806	ARP Request	CRC

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

LAN Switch Operation ④

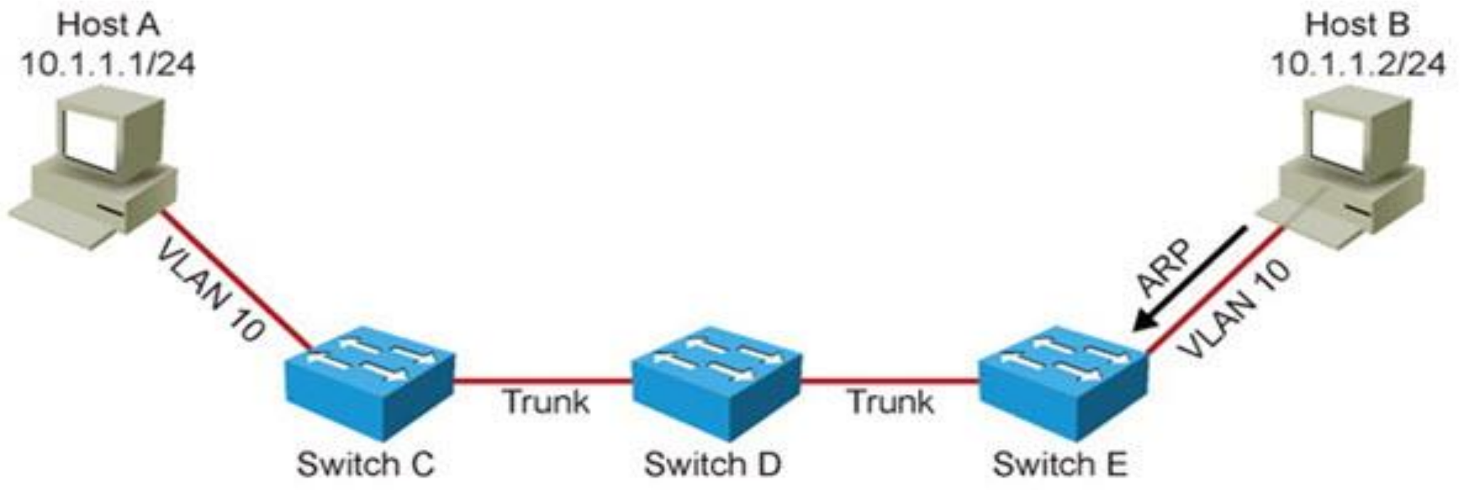

DMAC	SMAC	Type	Data	FCS
MAC A	MAC B	0x0806	ARP Reply	CRC

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

LAN Switch Operation ⑤

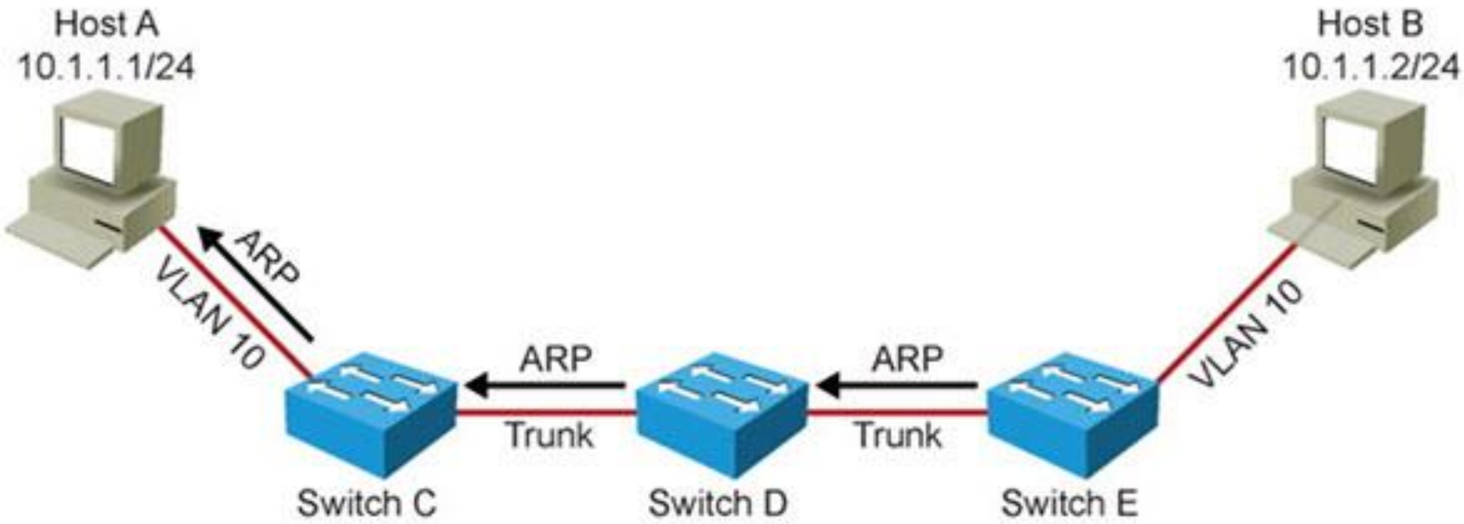

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
MAC A	MAC B	0x8100	VLAN 10	0x0806	ARP Reply	CRC

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

LAN Switch Operation ⑥

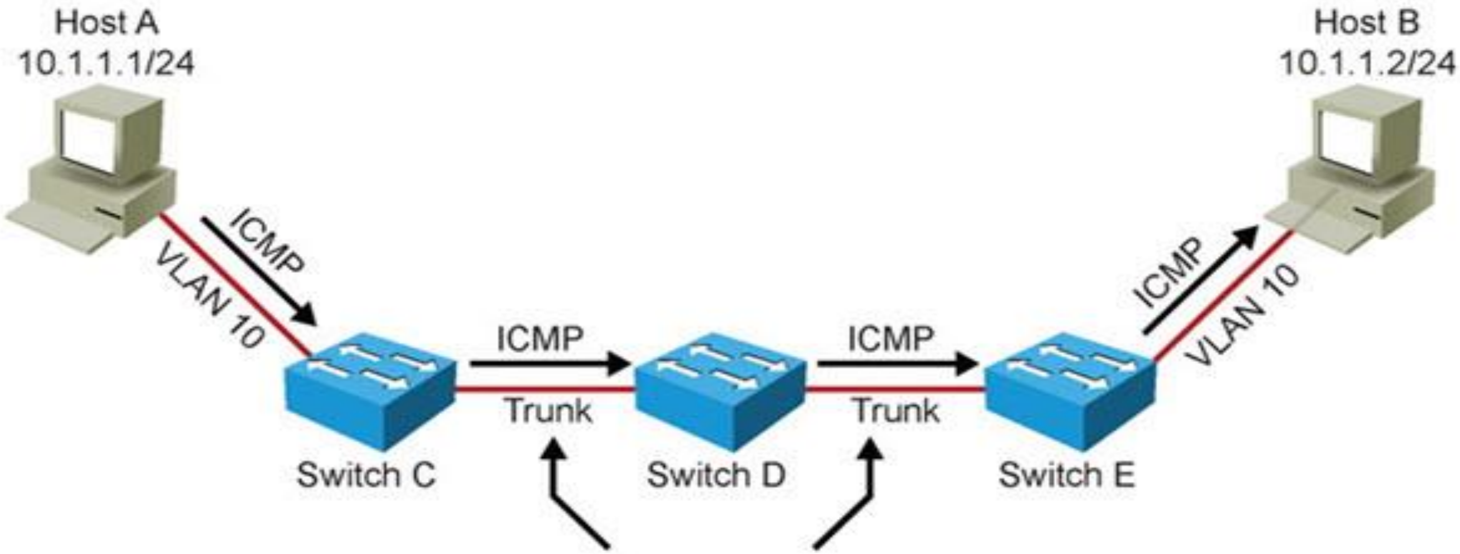

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
MAC B	MAC A	0x8100	VLAN 10	0x0800	ICMP Echo Request	CRC

- 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 ⑦

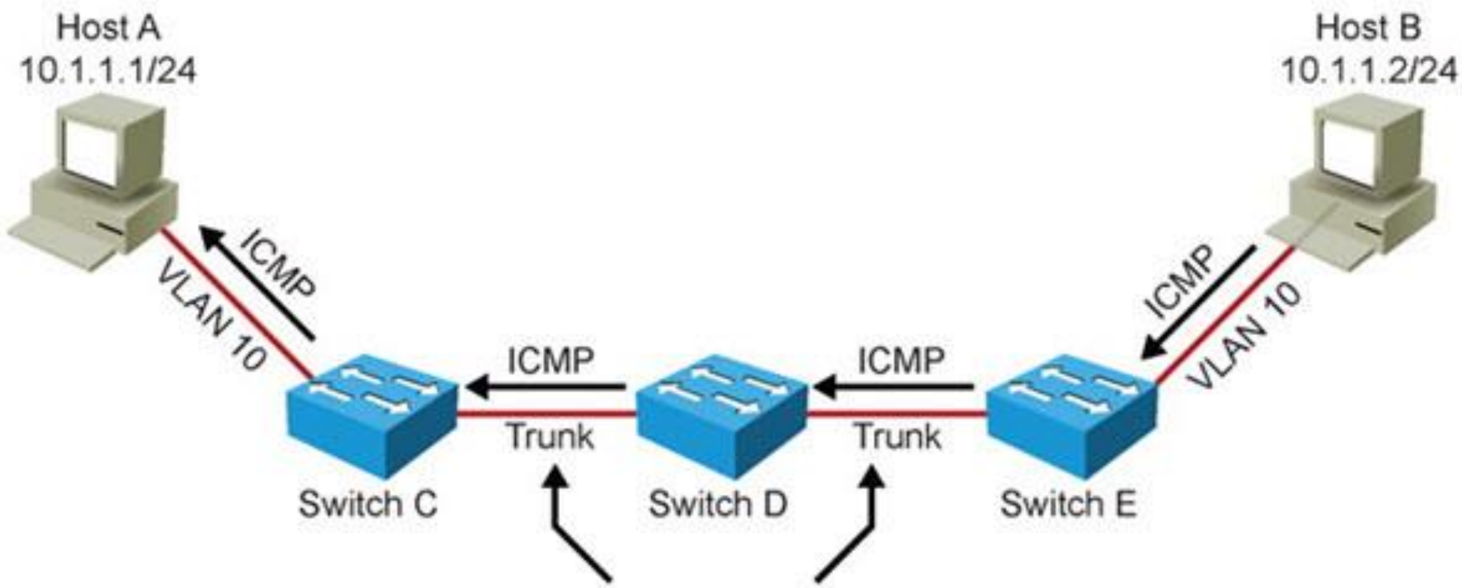

DMAC	SMAC	Type	802.1Q	Type	Data	FCS
MAC A	MAC B	0x8100	VLAN 10	0x0800	ICMP Echo Reply	CRC

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

LAN Switch Operation ⑧

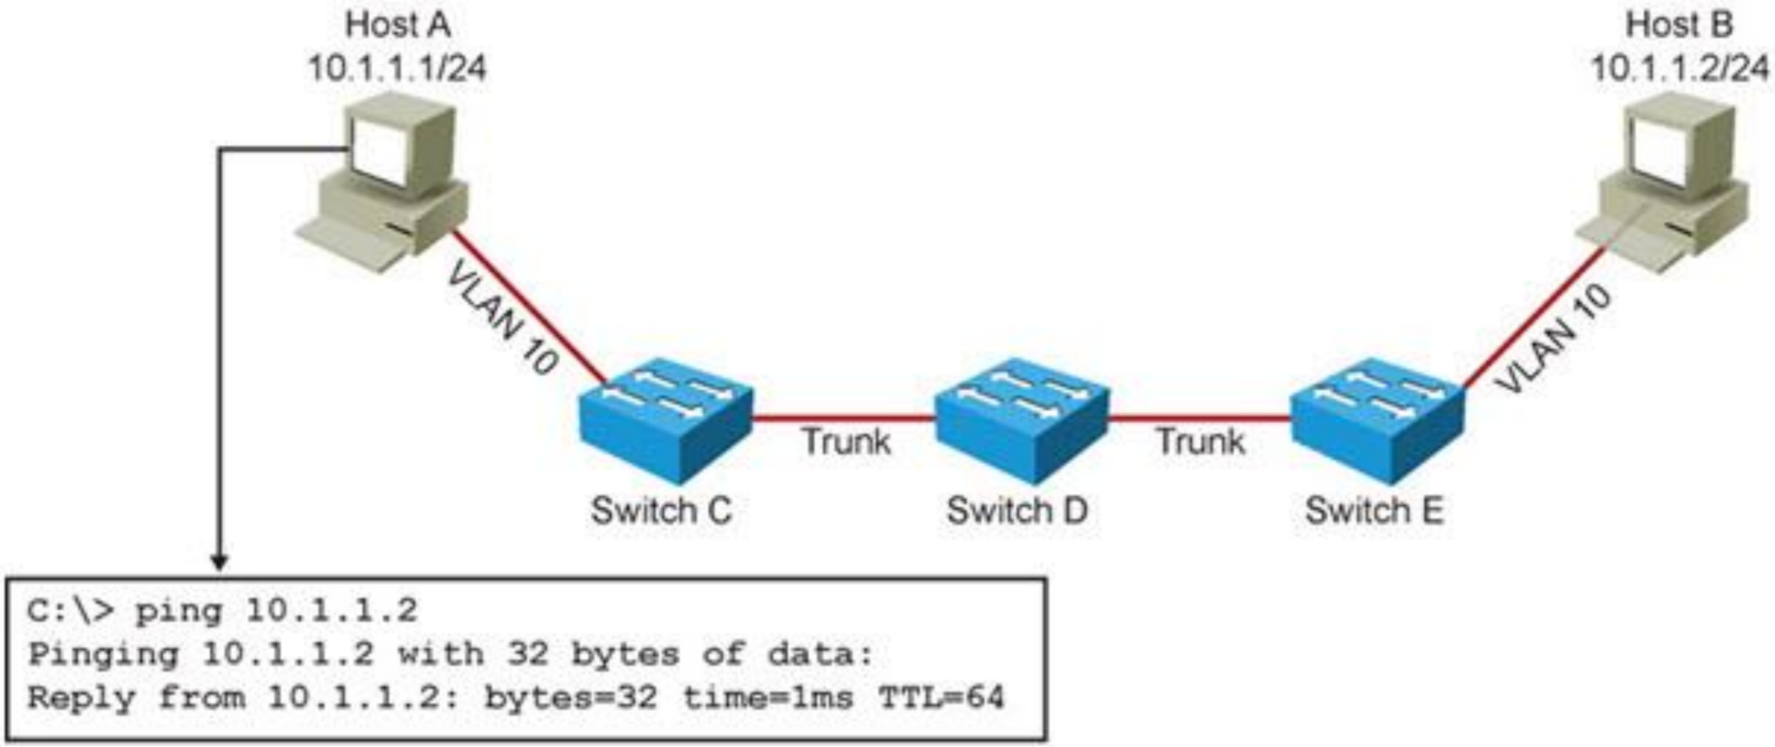

- Host A Receives ICMP Echo Reply Back from Host B.

Layer1 and Layer2 Issues

Issues that could cause the communication to fail:

- Bad, missing, or miswired cables
- Bad ports
- 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 ①

- *Are frames received on the correct VLAN?*
  - This could point to VLAN or trunk misconfiguration as the cause of the problem.
- *Are frames received on a correct port?*
  - This could point to a physical problem, spanning tree issues, a native VLAN mismatch or duplicate MAC addresses.
- *Is the MAC address 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.

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.

The `traceroute mac` Command

```
Router# traceroute mac 0000.0201.0601 0000.0201.0201

Source 0000.0201.0601 found on con6[WS-C2950G-24-EI] (2.2.6.6)
con6 (2.2.6.6)   :    Fa0/1 => Fa0/3
con5 (2.2.5.5)   :    Fa0/3 => Gi0/1
con1 (2.2.1.1)   :    Gi0/1 => Gi0/2
con2 (2.2.2.2)   :    Gi0/2 => Fa0/1
Destination 0000.0201.0201 found on con2[WS-C3550-24] (2.2.2.2)
Layer 2 trace completed

Router# traceroute mac 0001.0000.0204 0001.0000.0304 detail

Source 0001.0000.0204 found on VAYU[WS-C6509] (2.1.1.10)
1 VAYU / WS-C6509 / 2.1.1.10 :
                Gi6/1 [full, 1000M] => Po100 [auto, auto]
2 PANI / WS-C6509 / 2.1.1.12 :
                Po100 [auto, auto] => Po110 [auto, auto]
3 BUMI / WS-C6509 / 2.1.1.13 :
                Po110 [auto, auto] => Po120 [auto, auto]
4 AGNI / WS-C6509 / 2.1.1.11 :
                Po120 [auto, auto] => Gi8/12 [full, 1000M] Destination
0001.0000.0304
found on AGNI[WS-C6509] (2.1.1.11) Layer 2 trace completed.
```

Troubleshooting Spanning Tree

Spanning Tree Operation ①

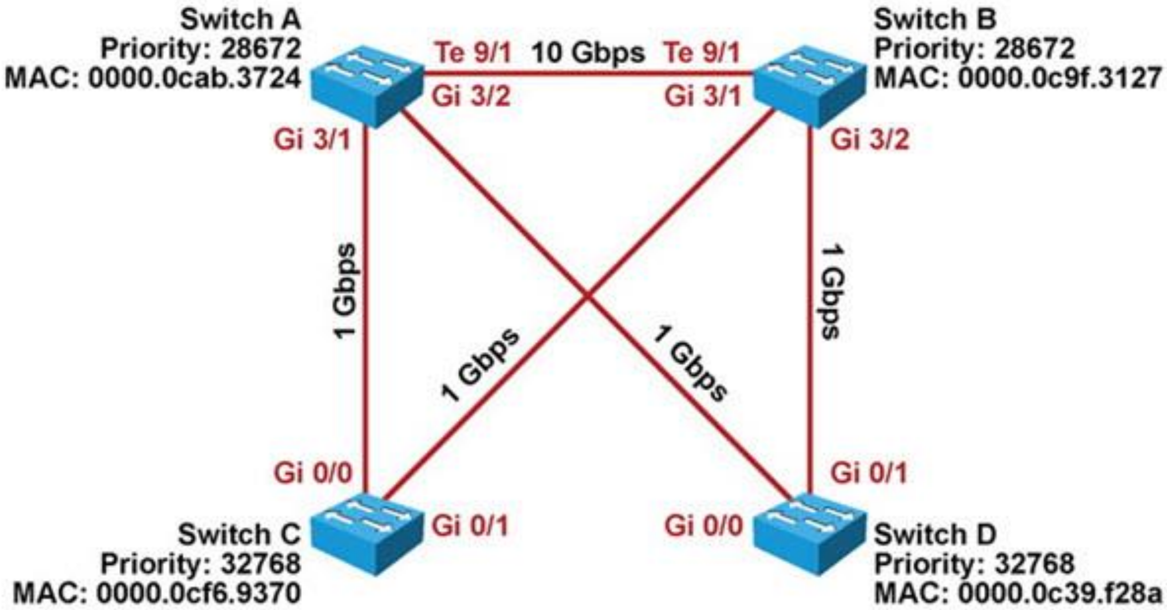

1) Elect a Root Bridge/Switch.
2) Select a Root Port on each Bridge/Switch (except on the Root bridge/switch).
3) Elect a Designated device/port on each network segment.
4) Ports that are neither Root Port nor a Designated Port go into Blocking state.

Spanning Tree Operation ②

1) Elect a Root Bridge/Switch.

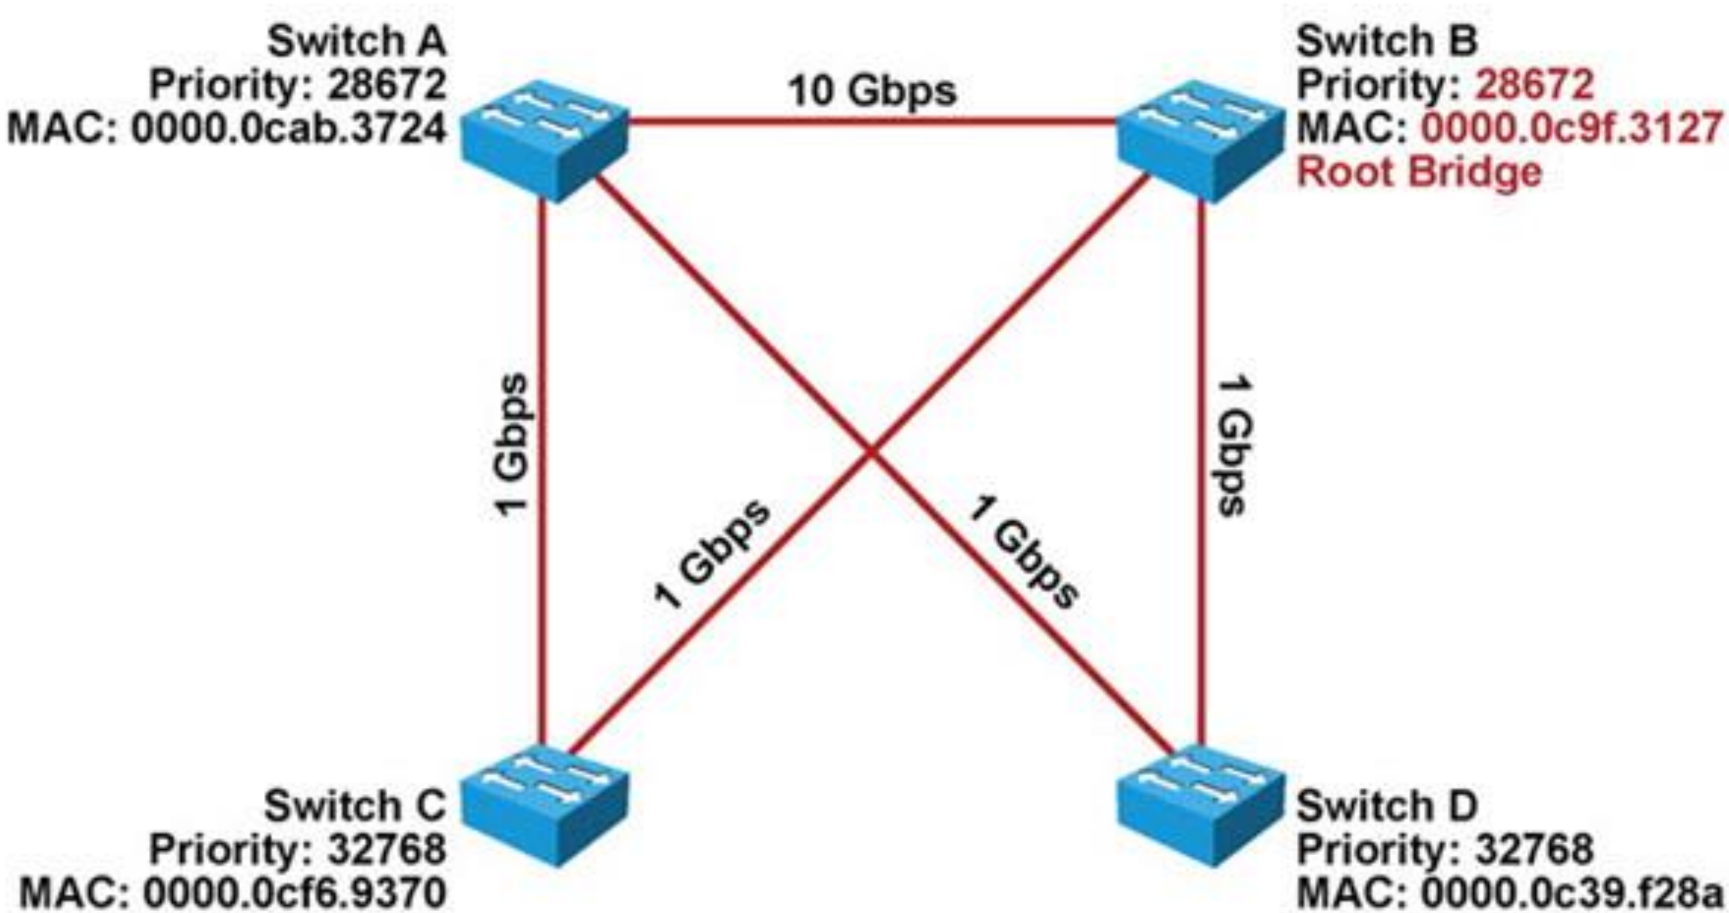

Spanning Tree Operation ③

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

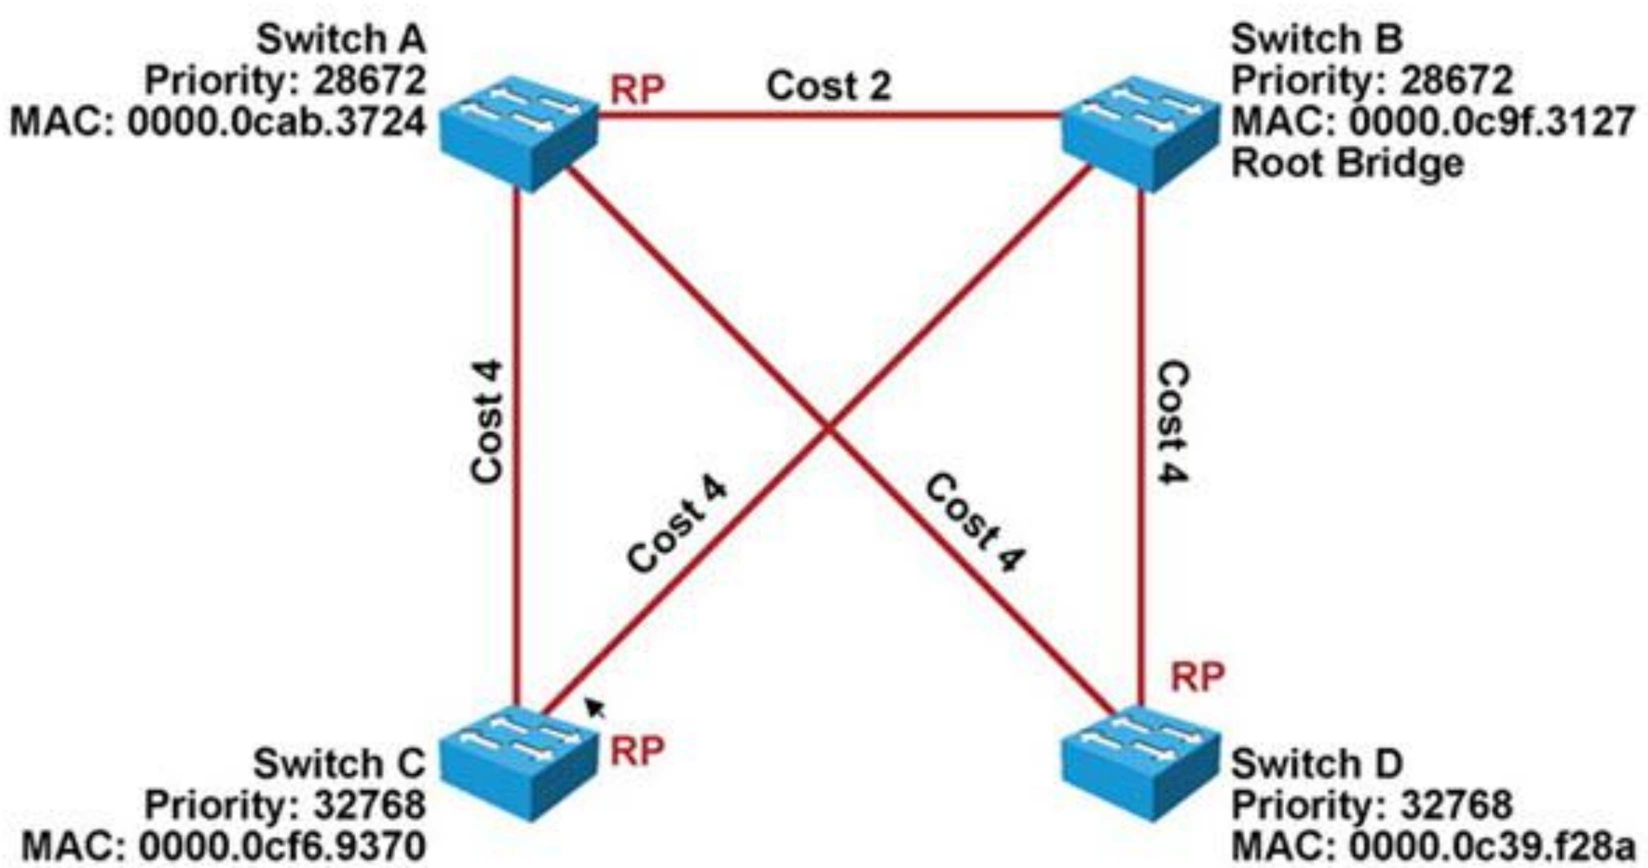

Spanning Tree Operation ④

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

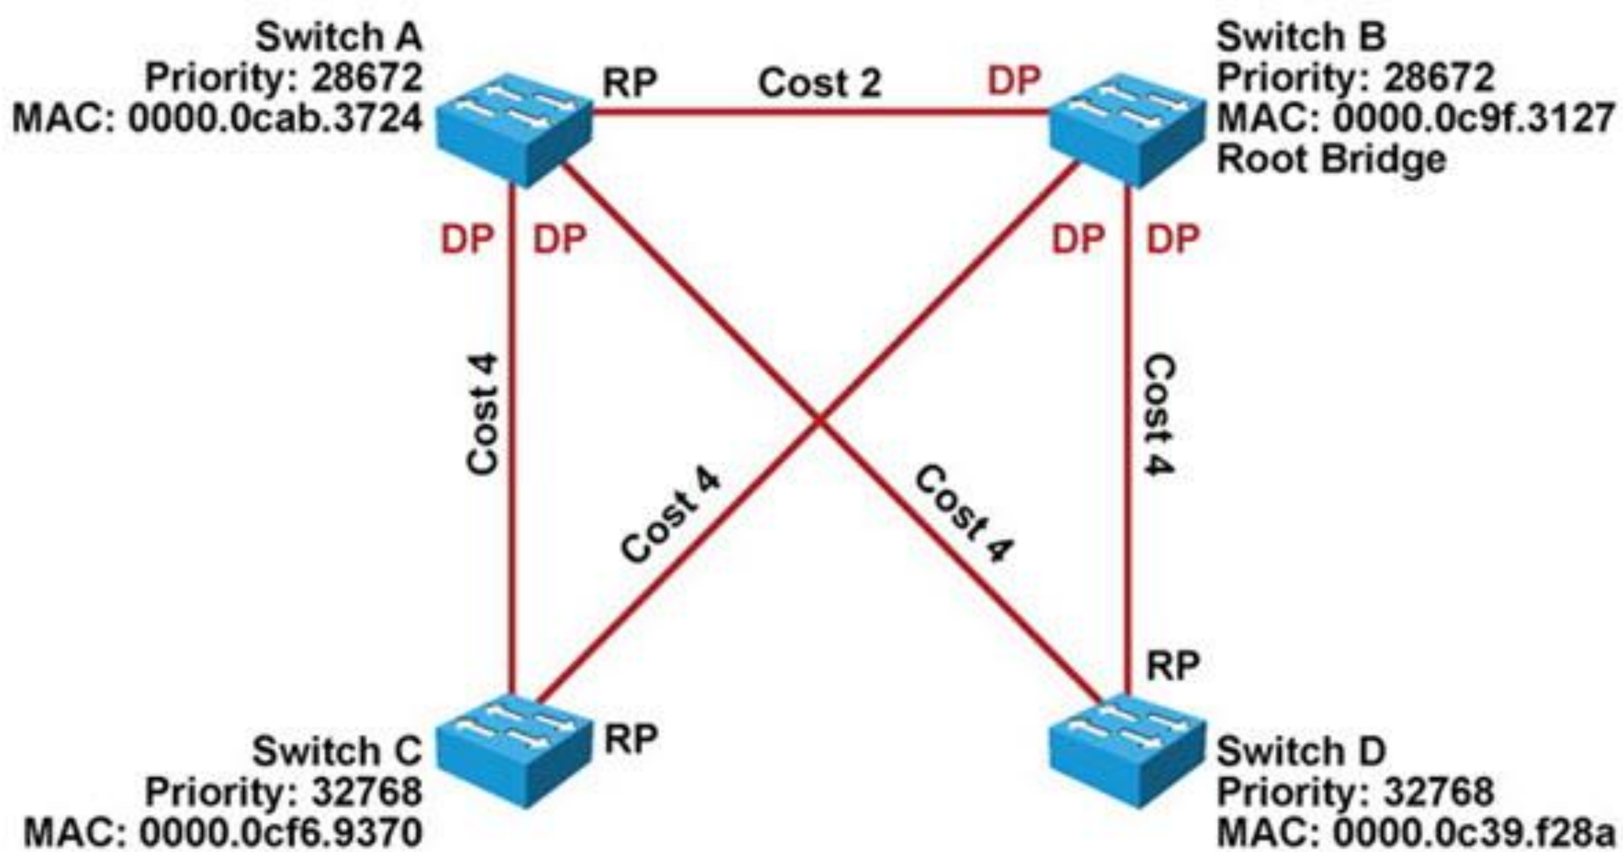

Spanning Tree Operation ⑤

4) Place ports in Blocking state.

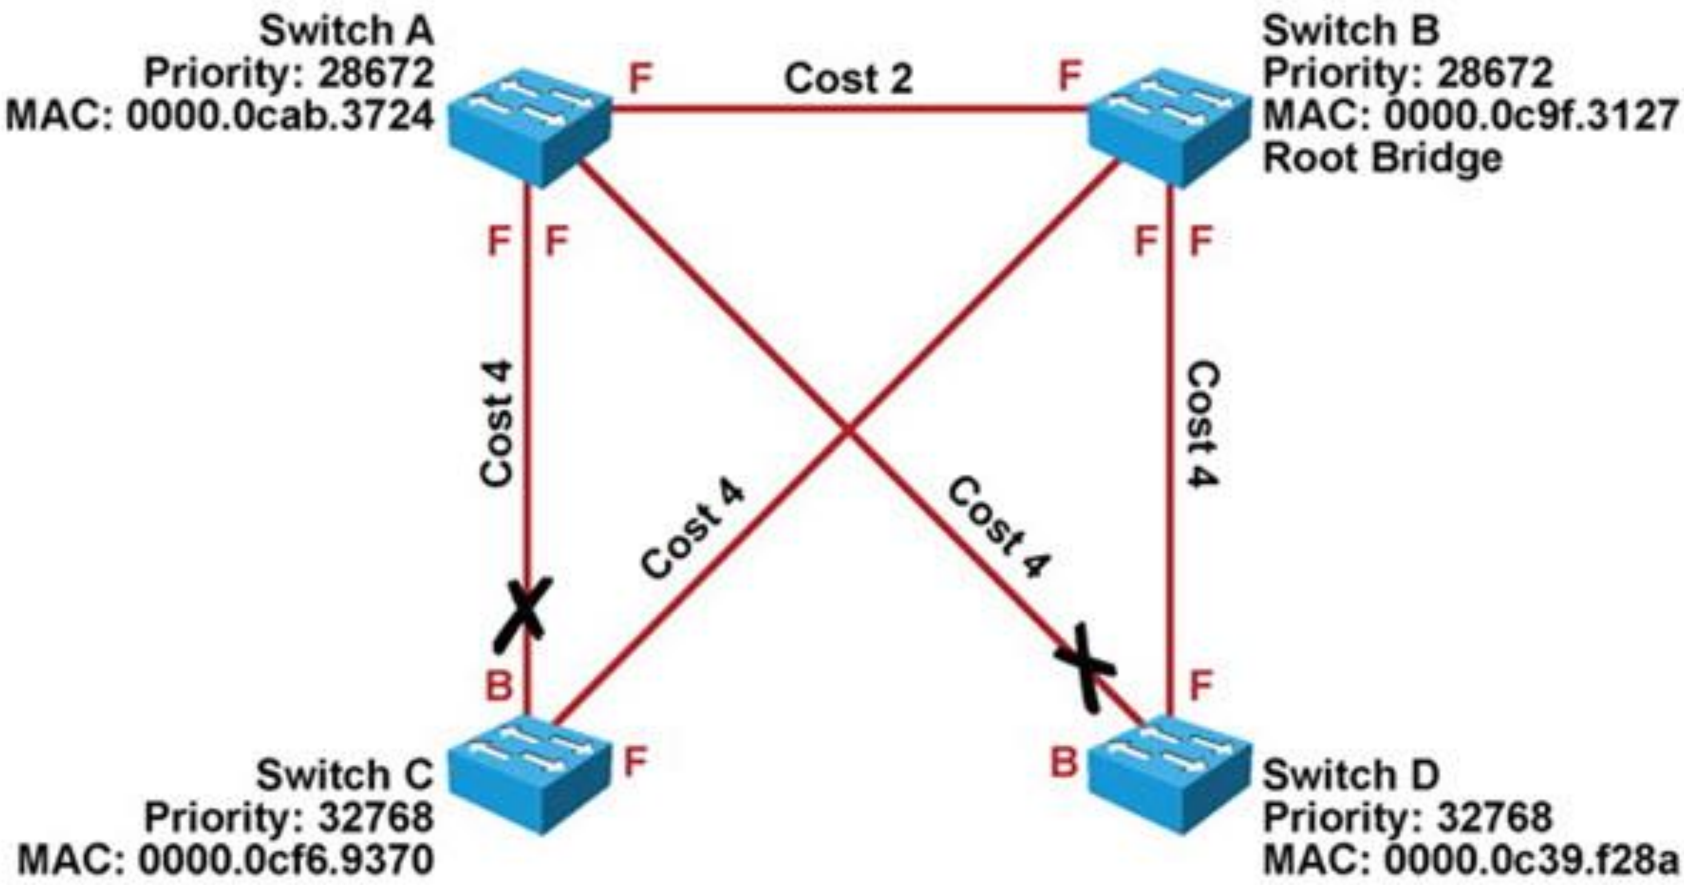

The `show spanning-tree vlan` Command

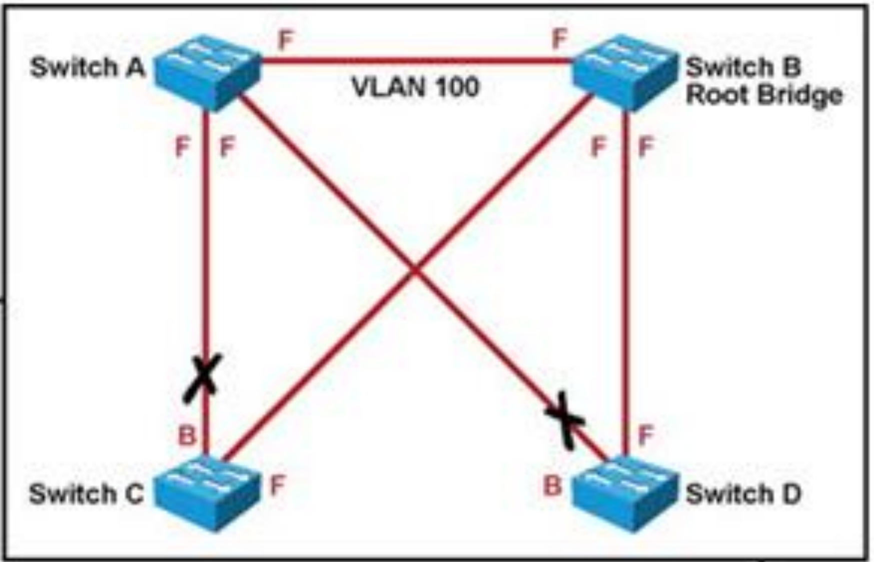


```
SwitchA#show spanning-tree vlan 100

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
```

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
```

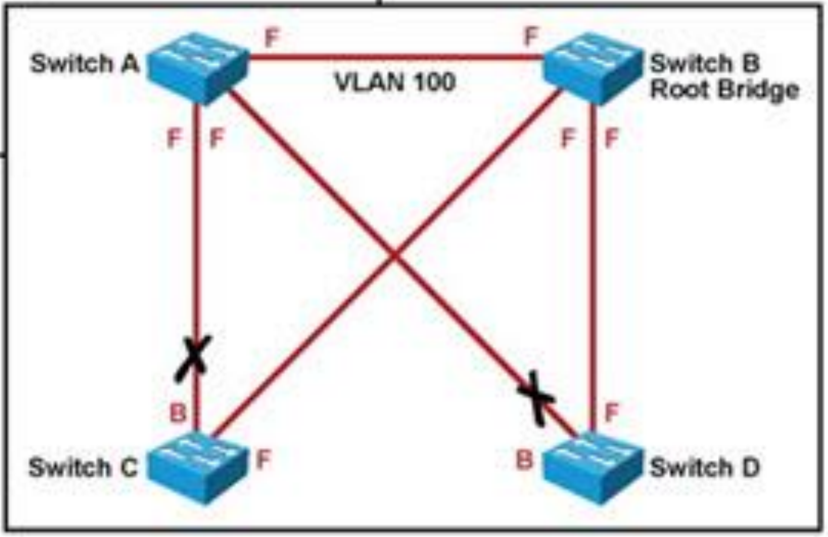

Spanning Tree Failures

- STP is a reliable but not an absolutely failproof protocol.
- `IF` STP fails
  `THEN` there are usually major negative consequences.
- There are two different types of failures.
  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.
  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.

- *Bridging loops 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.

The show spanning-tree Command

```
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
```

Common STP Problems

- Unidirectional or split links
  - Usually on optical medium
- Wrong ACL that blocks BPDU messages
- Duplex mismatch
  - Leads to collisions and BPDUs discarding
- Too much VLANs
  - Technical limitation of the number of running STP instances (e.g., 128). Exceeding VLANs do not run STP and thus are not protected against loops
- Too much CPU utilization, CPU cannot process BPDUs
- Wrong EtherChannel link configuration
  - One side is “on”, other one is not configured at all
- Special problems when mixing MST with PVST+ or RPVST+

Troubleshooting Etherchannel

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
- Packets and frames are routed or switched to the port-channel interface
- STP and routing protocols interact with this single port-channel interface
- A hashing mechanism determines which physical link will be used to transmit them

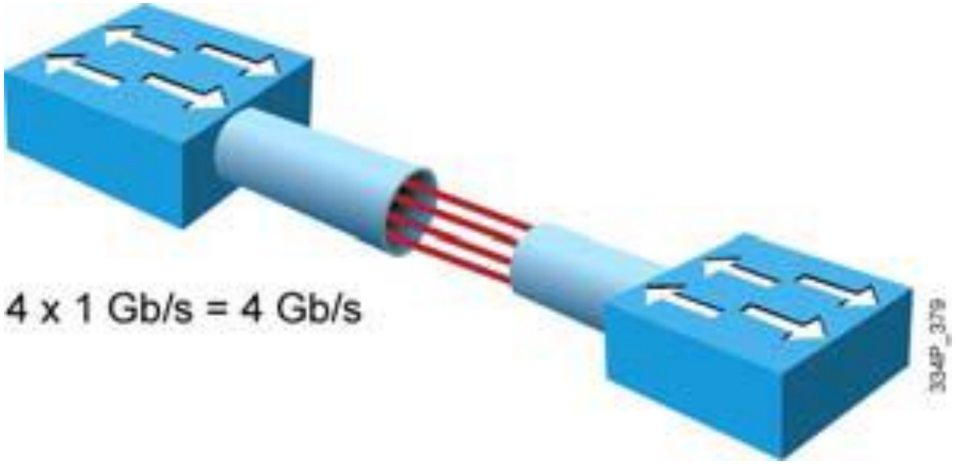

Common EtherChannel Problems

1) Inconsistencies between the physical ports that are members of the channel (a `%EC-5-CANNOT_BUNDLE2` log message is generated)
2) Inconsistencies between the ports on the opposite sides of the EtherChannel link (The switch will generate a `%SPANTREE-2-CHNL_MISCFG` message)
3) Uneven distribution of traffic between EtherChannel bundle members

- The most of the problems could be solved by using EtherChannel management protocols
  - **Link Aggregation Control Protocol (LACP)**, IEEE
  - **Port Aggregation Protocol (PAgP)**, Cisco
  - Avoid „on“ mode

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)
```

The show etherchannel 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>
```

The show spanning-tree Command

```
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
```

L2 Troubleshooting Flow

L2 Troubleshooting Flow ①

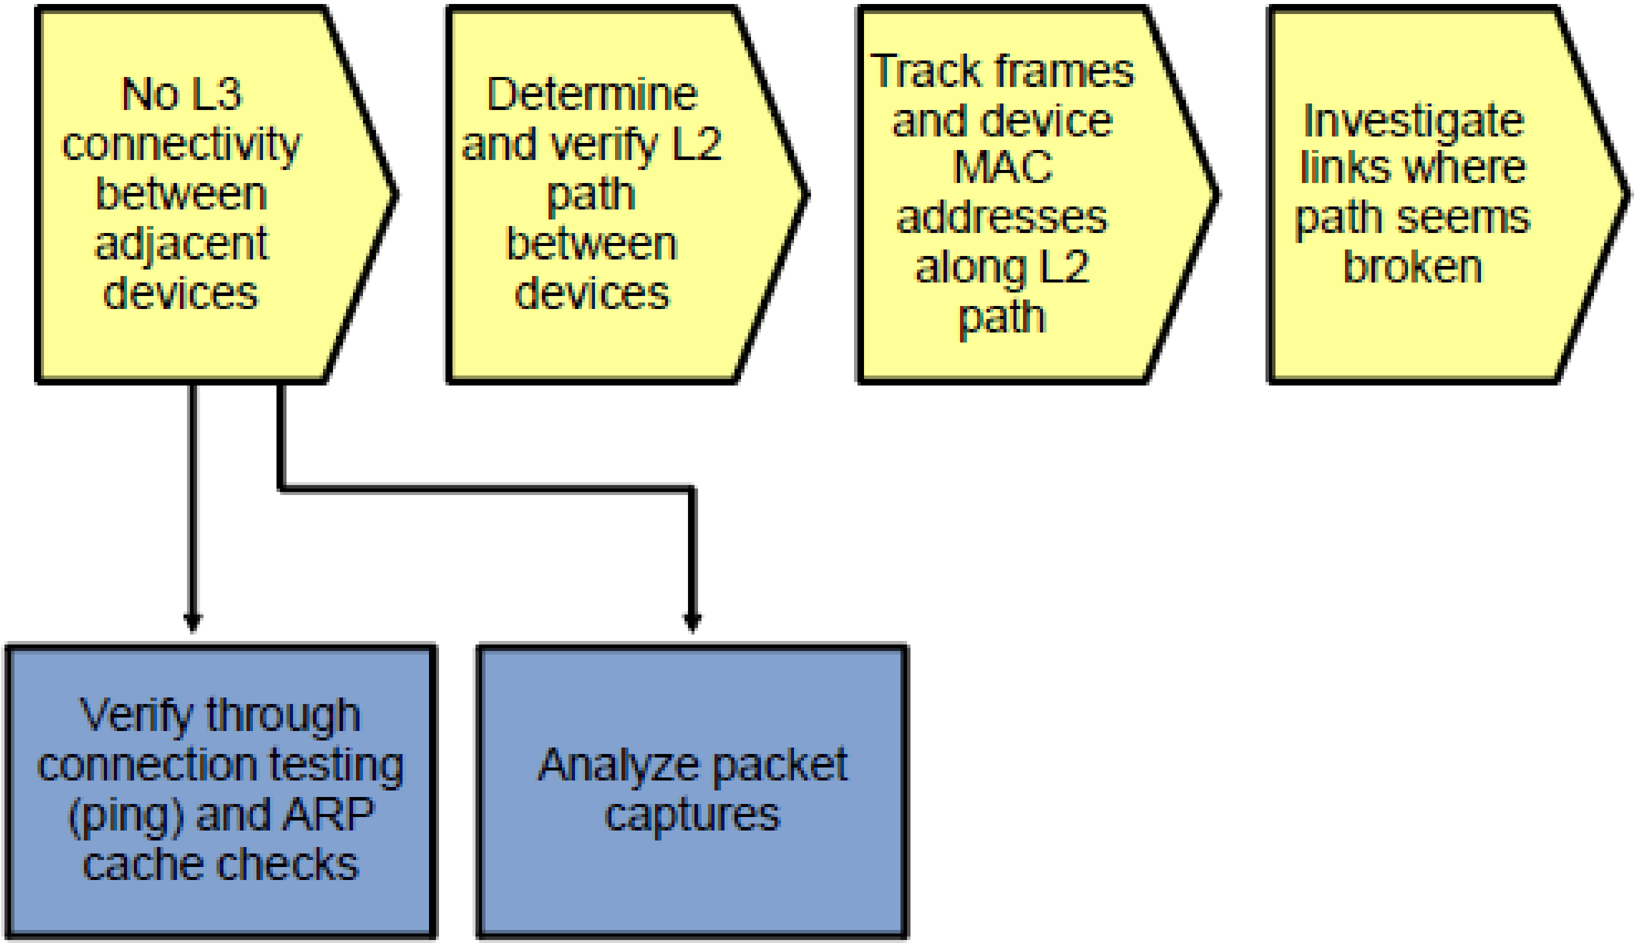

L2 Troubleshooting Flow ②

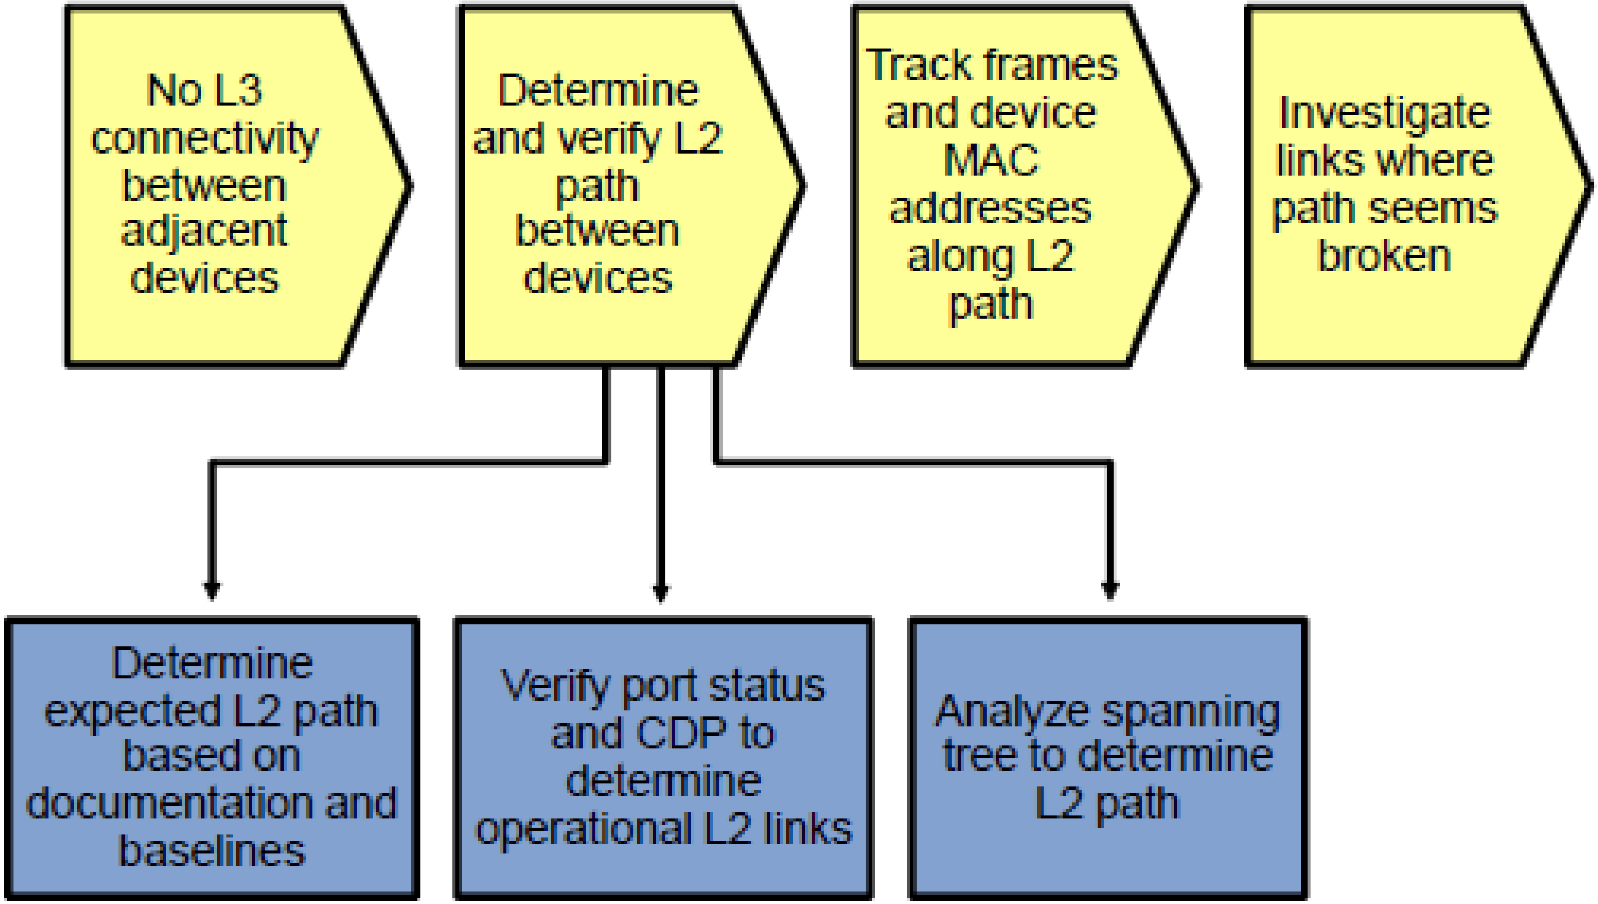

L2 Troubleshooting Flow ③

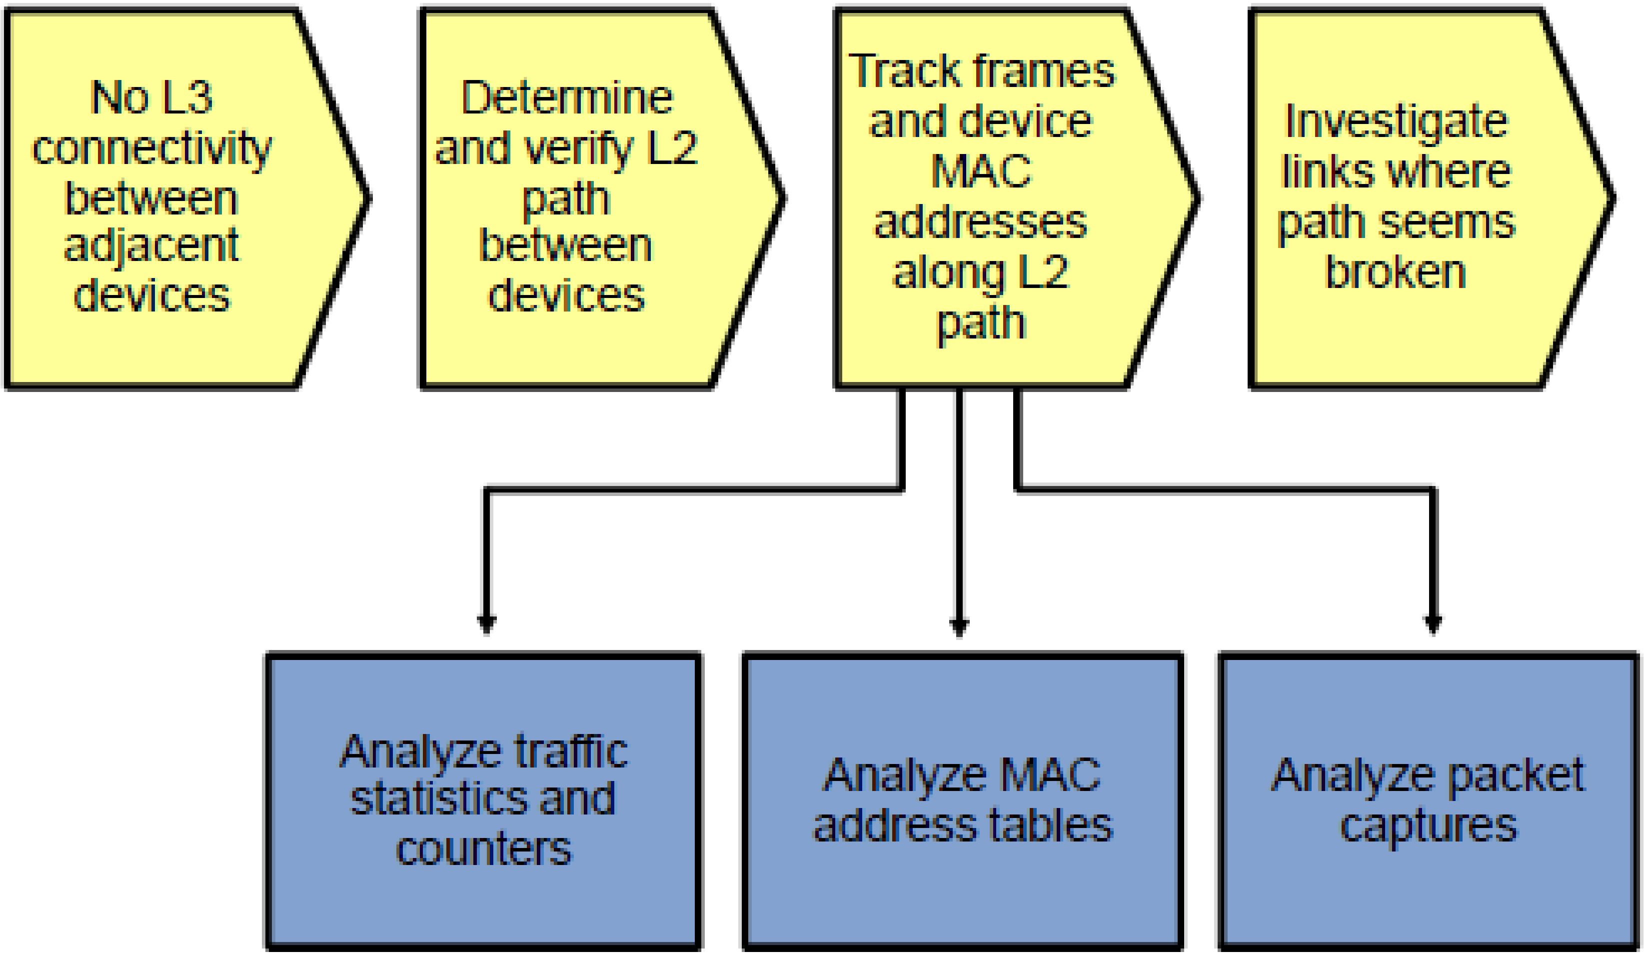

L2 Troubleshooting Flow ④

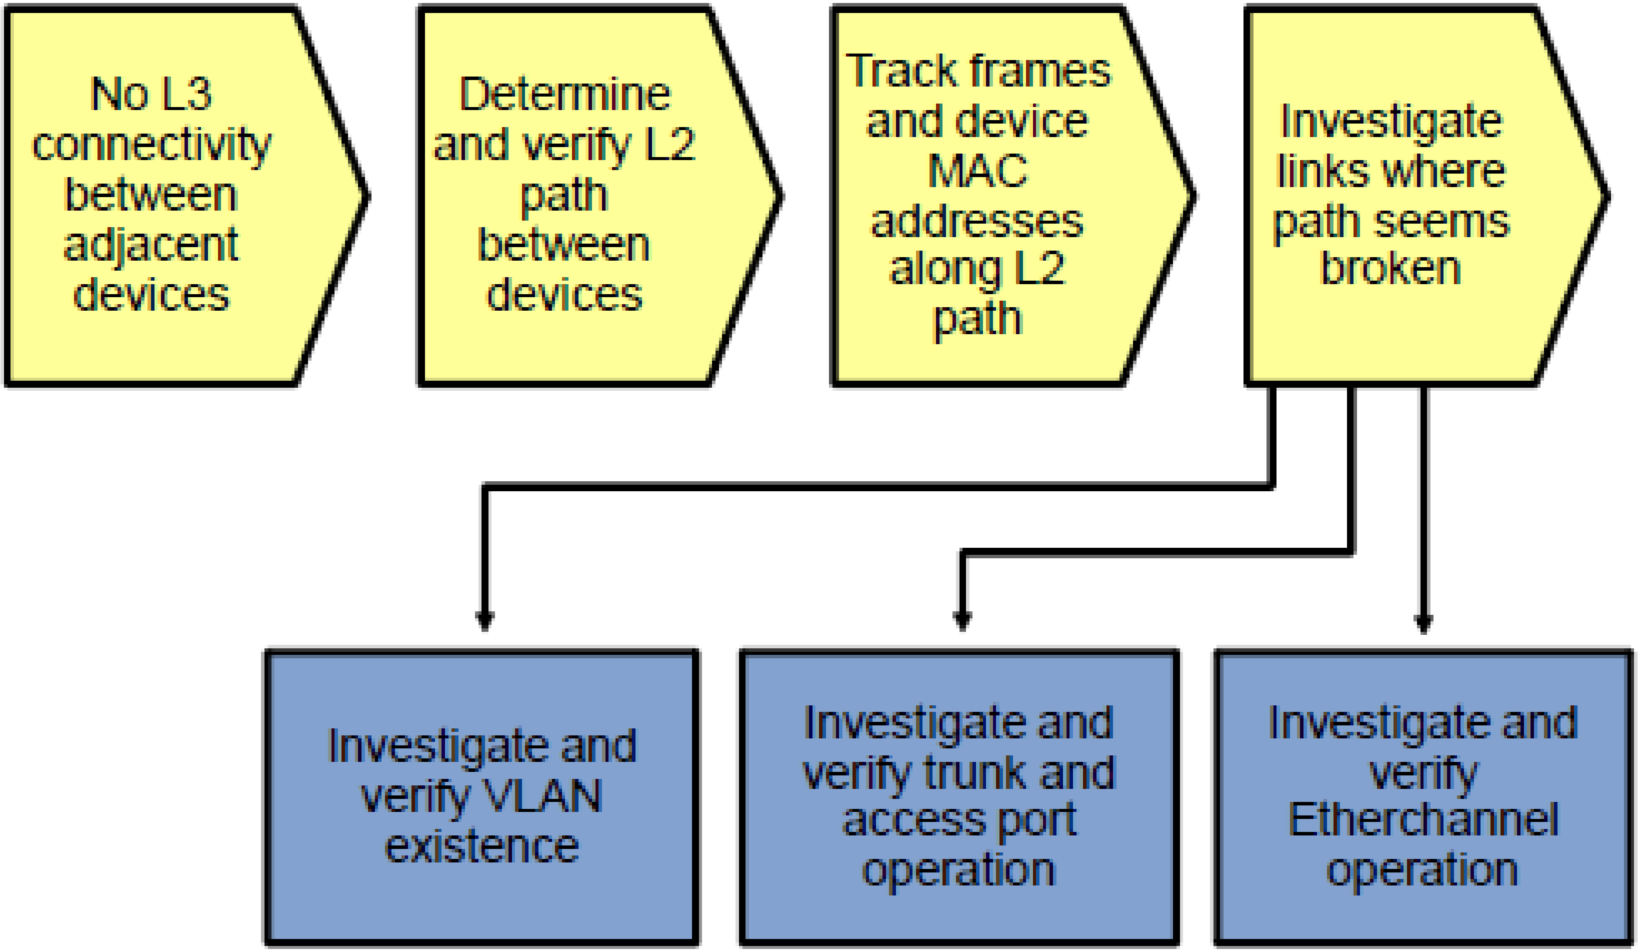

Troubleshooting Switched Virtual Interfaces and Inter-VLAN Routing

Router and MLS Similarities

- Both routers and multilayer switches use routing protocols or static routes to maintain routing information 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

Router and MLS Differences

- Routers connect heterogeneous networks and support a wide variety of media and interfaces
- Multilayer switches typically connect homogenous networks. Most LAN switches are Ethernet only.
- Multilayer switches utilize specialized hardware to achieve wire-speed Ethernet-to-Ethernet packet switching
- Low- to mid-range routers use multi-purpose hardware to perform the packet switching process
- On average, the packet switching throughput of routers is lower than the packet switching throughput of multilayer switches
- Routers usually support a wider range of features, mainly because switches need specialized hardware to be able to support certain data plane features or protocols
- On routers, you can often add features through a software update

Switch Performance

- **Data plane**
  - Ingress interface
  - Forwarding hardware
  - Egress interface
- **Control plane**
  - CPU
  - Memory

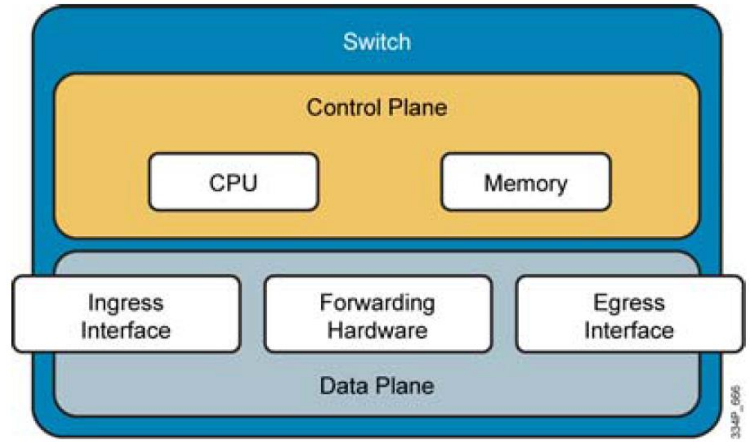

Cisco 7206

Catalyst 6504

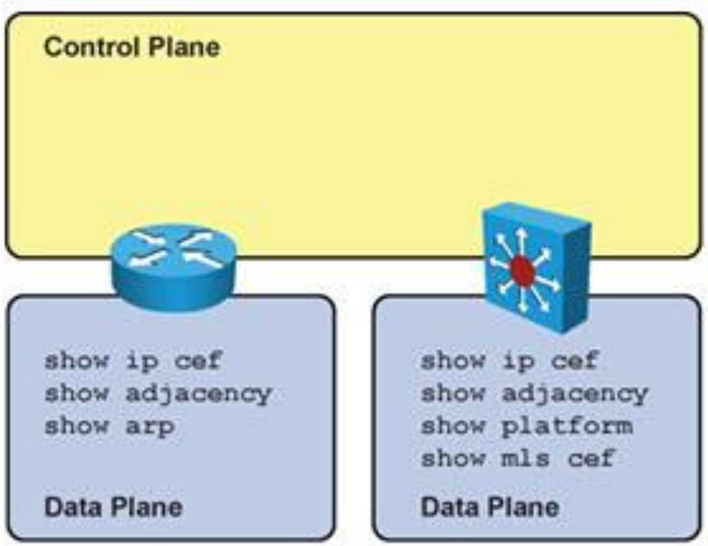

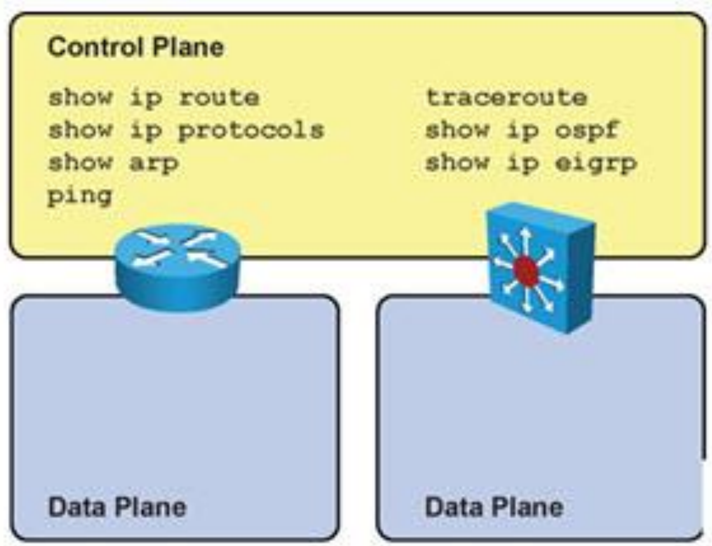

The `show interface counters` Command

```
Router# show interfaces fas 6/1 counters

  Port              InOctets   InUcastPkts   InMcastPkts   InBcastPkts
  Fa6/1             47856076            23        673028           149

  Port             OutOctets  OutUcastPkts  OutMcastPkts  OutBcastPkts
  Fa6/1             22103793            17        255877          3280
```

- Duplex mismatch example

```
DLS1# show interfaces fastEthernet 0/1 | include duplex
  Full-duplex, 100Mb/s, media type is 10/100BaseTX

DLS1# show interfaces fastEthernet 0/1 counters errors

 Port        Align-Err     FCS-Err    Xmit-Err     Rcv-Err UnderSize OutDiscards
 Fa0/1               0       12618           0           0         0           0

 Port      Single-Col Multi-Col  Late-Col Excess-Col Carri-Sen     Runts     Giants
 Fa0/1              0         0         0          0         0         0          0
```

```
DLS2# show interfaces fastEthernet 0/1 | include duplex
  Half-duplex, 100Mb/s, media type is 10/100BaseTX

DLS1# show interfaces fastEthernet 0/1 counters errors

 Port        Align-Err     FCS-Err    Xmit-Err     Rcv-Err UnderSize OutDiscards
 Fa0/1               0           0           0           0         0           0

 Port      Single-Col Multi-Col  Late-Col Excess-Col Carri-Sen     Runts     Giants
 Fa0/1              0         0     12679          0         0         0          0
```

Interface Errors ①

- Following errors indicate cabling, NIC or duplex issues
- **Align-Err**
  - The number of frames with alignnment errors, which means that they not end with an even number and have bad CRC
- **FCS-Err**
  - The number of valid size frames with FCS error but no framing errors
- **Xmit-Err** and **Rcv-Err**
  - Internal Tx or Rx buffers are full
  - Common cause is high utilization of link
- **Undersize**
  - The frames that are smaller than the IEEE 802.3 frame size minimum of 64 bytes long
- **Runts**
  - The frames that are smaller than the IEEE 802.3 frame size minimum of 64 bytes long `AND` with bad CRC
- **Giants**
  - Frames that exceed the IEEE 802.3 frame size minimum of 1518 bytes long `AND` with bad CRC

Interface Errors ②

- Following errors indicate duplex issues
- **Single-Col**
  - The number of times one collision occurs before port transmits the frame to the medium
  - Usual on half-duplex port, should not be seen on full-duplex
- **Multi-Col**
  - The number of times multiple collisions occurs before port transmits the frame to the medium and same conditions as previous
- **Late-Col**
  - The number of frames that a collision is detected on a particular port late in the trasmission process (for 10Mb/s port later than 51.2 usec)
  - For duplex mismatch seen on half-duplex side
- **Excess-Col**
  - This is a count of frames trasmitted on a particular port, which fail due to the excessive collisions (16 times in a row)
  - Typically indicates that a load needs to be split across multple segments
- **Carri-Sen**
  - This occurs every time port wants to send data on half-duplex connection

Duplex and Audo-MDIX Mismatches

- A common cause for performance problems in Ethernet-based networks is duplex mismatch
- *Duplex guidelines*
  - Point-to-Point links should be always full-duplex
  - Half-duplex is not common anymore and is mostly encountered in topologies with hub devices
  - Autonegotiation of speed and duplex is recommended, otherwise setup both ends of the link manually
  - Half-duplex on both ends performs better than duplex mismatch
- The **Automatic Media-Dependent Interface Crossover (Auto-MDIX)** feature detects required connection type
  - Enabled by default on switches
  - Auto-MDIX is dependent on auto-negotiation for speed and duplex
    - `IF` speed and duplex negotiation are turned off
      `THEN` Auto-MDIX is turned off as well

Configuration and Verifying Auto-MDIX

- Setting up Auto-MDIX

```
CSW1(config)#interface FastEthernet 0/10
CSW1(config-if)#mdix auto
CSW1(config-if)#speed auto
CSW1(config-if)#duplex auto
```

- Verifying

```
sw1# show interfaces transceiver properties
Name : Fa0/1
Administrative Speed: auto
Administrative Duplex: auto
Administrative Auto-MDIX: on
Administrative Power Inline: N/A
Operational Speed: auto
Operational Duplex: auto
Operational Auto-MDIX: on
```

Multi-layer Switch Interfaces

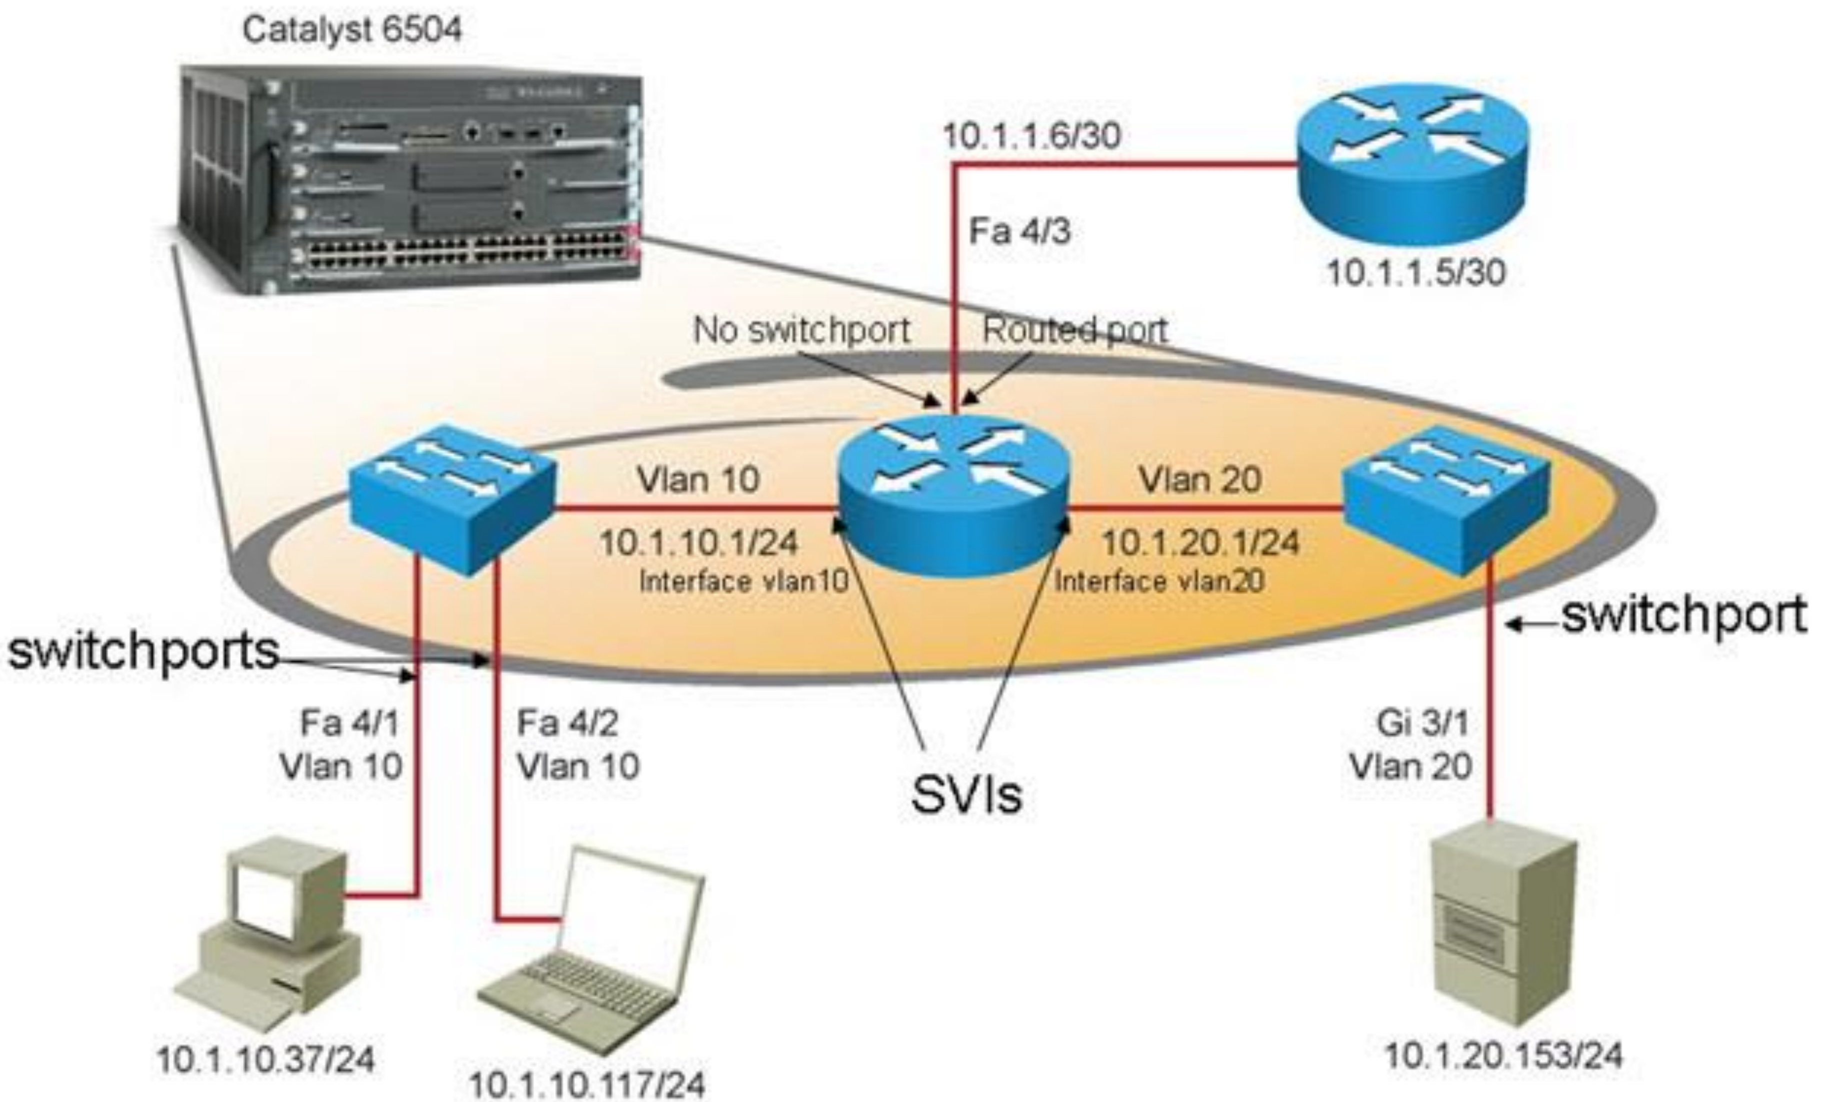

Three MLS Core Functions ①

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.
2) 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.

Three MLS Core Functions ②

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.

SVI vs. Routed Interfaces

- A routed interface is not a L2 port – L2 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 `THEN` 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).

Verifying the status of a VLAN and SVI

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


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
```

Forwarding Hardware

- Forwarding hardware consists of:
  - Decision-making logic
    - L2/L3 switching actions
    - ACL processing
    - QoS processing
  - A backplane to carry data between interfaces

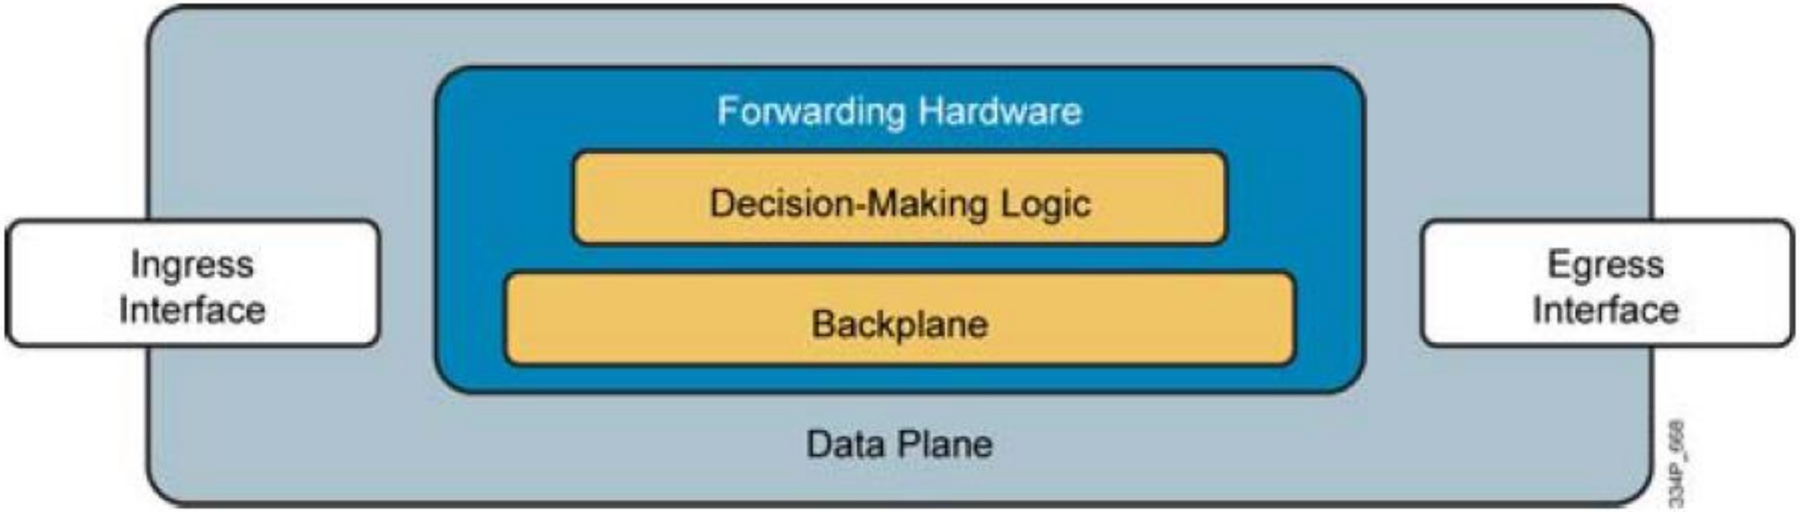

Ternary Content Addressable Memory

- Control plane information that affects packet forwarding is programmed into Ternary Content Addressable Memory (TCAM)
- Packets that cannot be handled by TCAM will be punted to the CPU

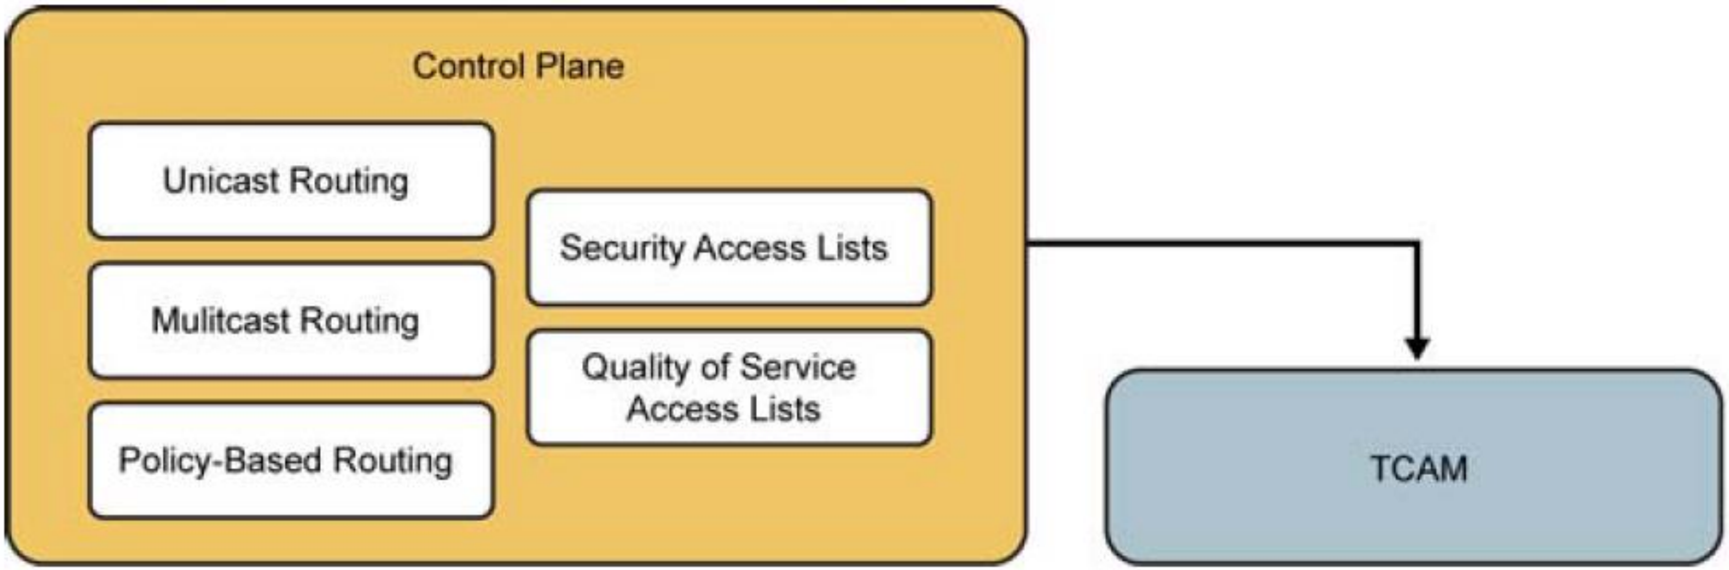

Useful CEF Commands

`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..

Useful Multi-layer Switches Commands

- 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.

The `show platform tcam util` Command

```
Switch# show platform tcam utilization

CAM Utilization for ASIC# 0                         Max              Used
                                               Masks/Values    Masks/values

 Unicast mac addresses:                          784/6272          12/26
 IPv4 IGMP groups + multicast routes:            144/1152           6/26
 IPv4 unicast directly-connected routes:         784/6272          12/26
 IPv4 unicast indirectly-connected routes:       272/2176           8/44
 IPv4 policy based routing aces:                   0/0              0/0
 IPv4 qos aces:                                  528/528           18/18
 IPv4 security aces:                            1024/1024          27/27

Note: Allocation of TCAM entries per feature uses
a complex algorithm. The above information is meant
to provide an abstract view of the current TCAM utilization
```

Traffic Forwarding to the CPU

- Traffic being punted to the CPU is indirect proof of TCAM allocation failures or use of unsupported features
- The **show controllers cpu-interface** displays the statistics for packets that are forwarded by CPU

```
Switch# show controllers cpu-interface
ASIC     Rxbiterr    Rxunder     Fwdctfix    Txbuflos    Rxbufloc    Rxbufdrain
-------------------------------------------------------------------------
ASIC0       0           0           0           0           0           0

cpu-queue-frames  retrieved  dropped    invalid    hol-block  stray
----------------- ---------- ---------- ---------- ---------- ----------
rpc               0          0          0          0          0
stp               1          0          0          0          0
ipc               0          0          0          0          0
routing protocol  28312      0          0          0          0
L2 protocol       0          0          0          0          0
remote console    0          0          0          0          0
sw forwarding     13800556   0          0          0          0
host              7648       0          0          0          0
broadcast         462103     0          0          0          0
cbt-to-spt        0          0          0          0          0
igmp snooping     35916      0          0          0          0
icmp              0          0          0          0          0
logging           0          0          0          0          0
rpf-fail          0          0          0          0          0
dstats            0          0          0          0          0
cpu heartbeat     22302361   0          0          0          0
```

CPU Problems

- First, determine whether interrupts or processes are the major cause of the increased CPU load
  - `IF` case of interrupts `THEN` troubleshoot packet forwarding and TCAM
  - `IF` case of processes `THEN` isolate responsible process and troubleshoot based on outcome
- In general, an average CPU load if 50% is not problematic just same as temporary 100% bursts
- Spikes in load could be caused by
  - Processor-intensive commands such as show tech-support, debug, show running, copy run start
  - Routing protocol updates
  - SNMP Polling

Isolating Process

- Important processes are
  - IP Input
  - IP ARP
  - SNMP Engine
  - IGMPSN

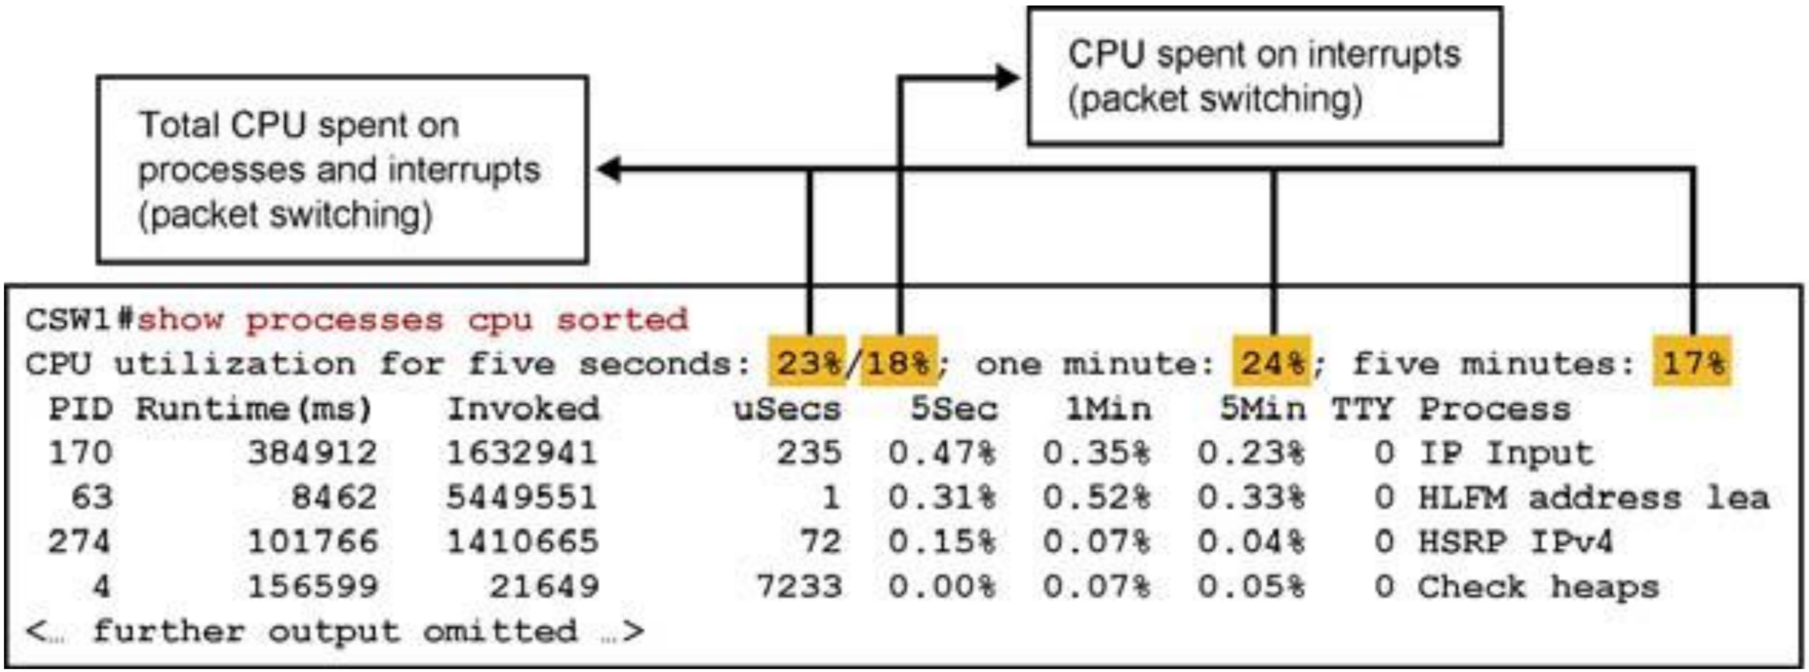

Multilayer Switching Troubleshooting Flow

Multilayer Switching Troubleshooting Flow ①

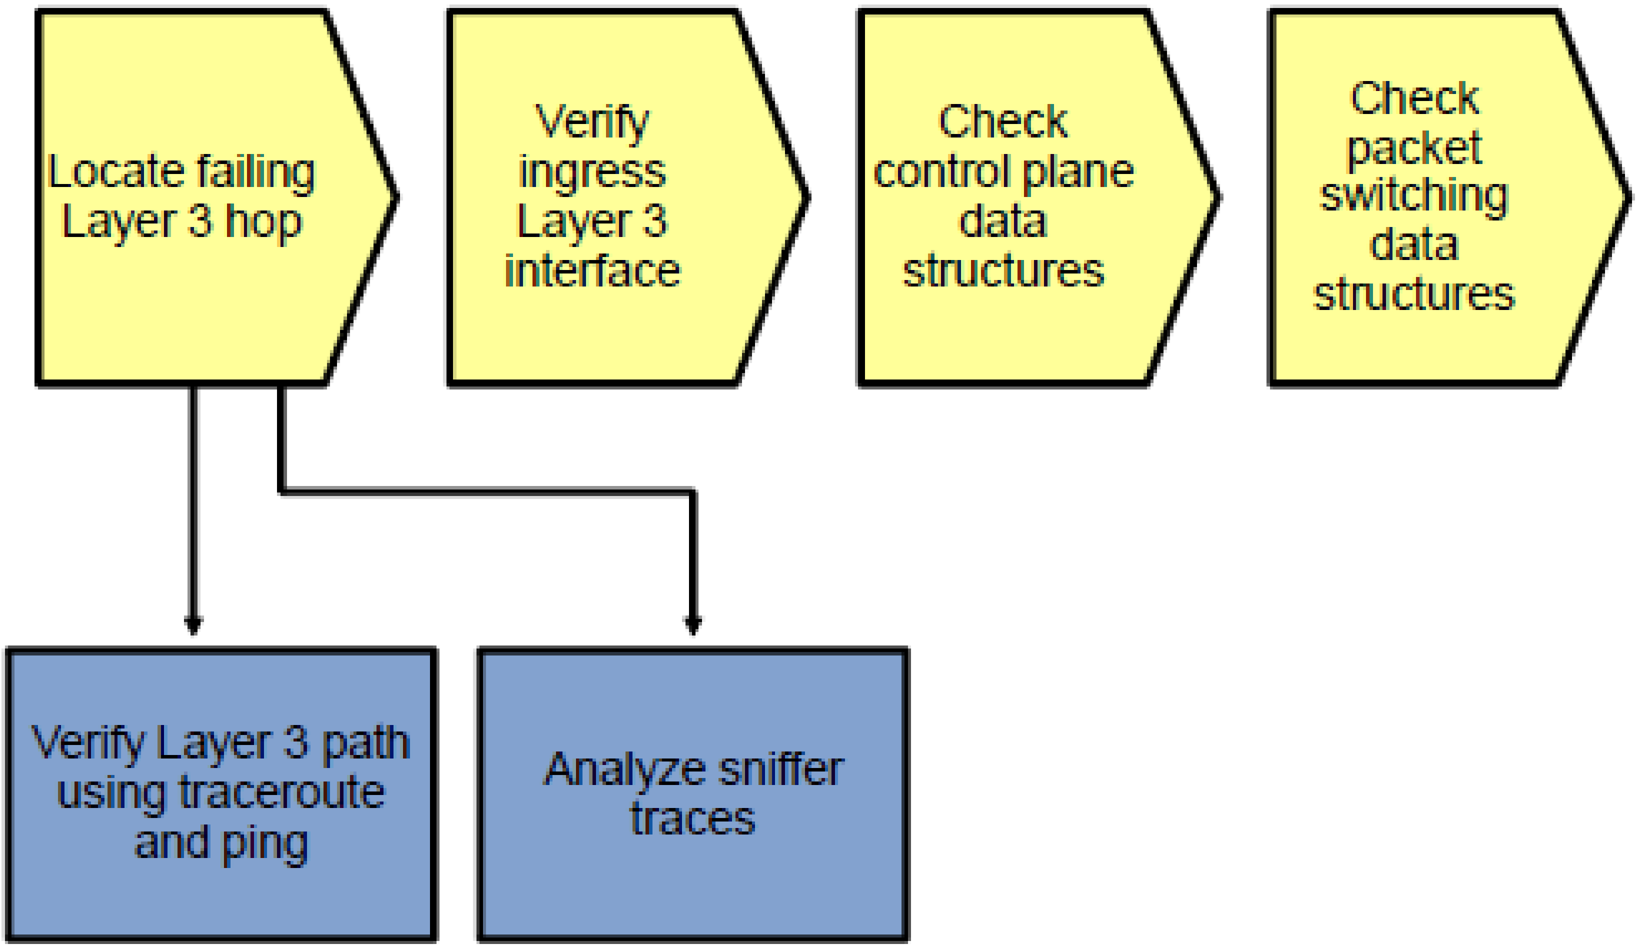

Multilayer Switching Troubleshooting Flow ②

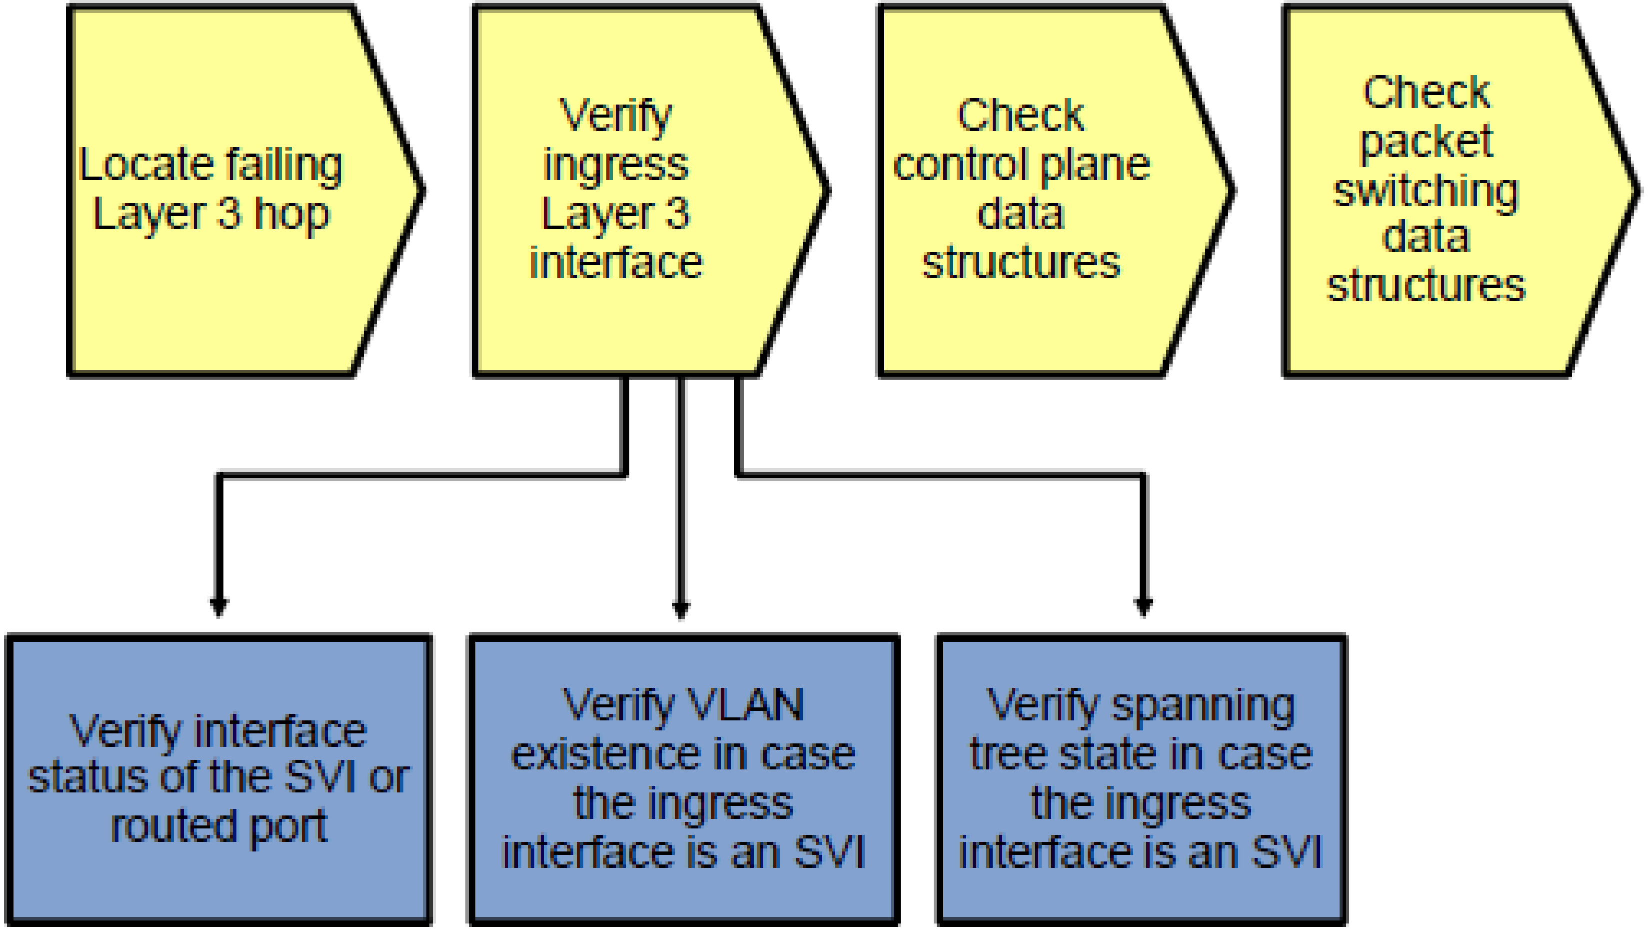

Multilayer Switching Troubleshooting Flow ③

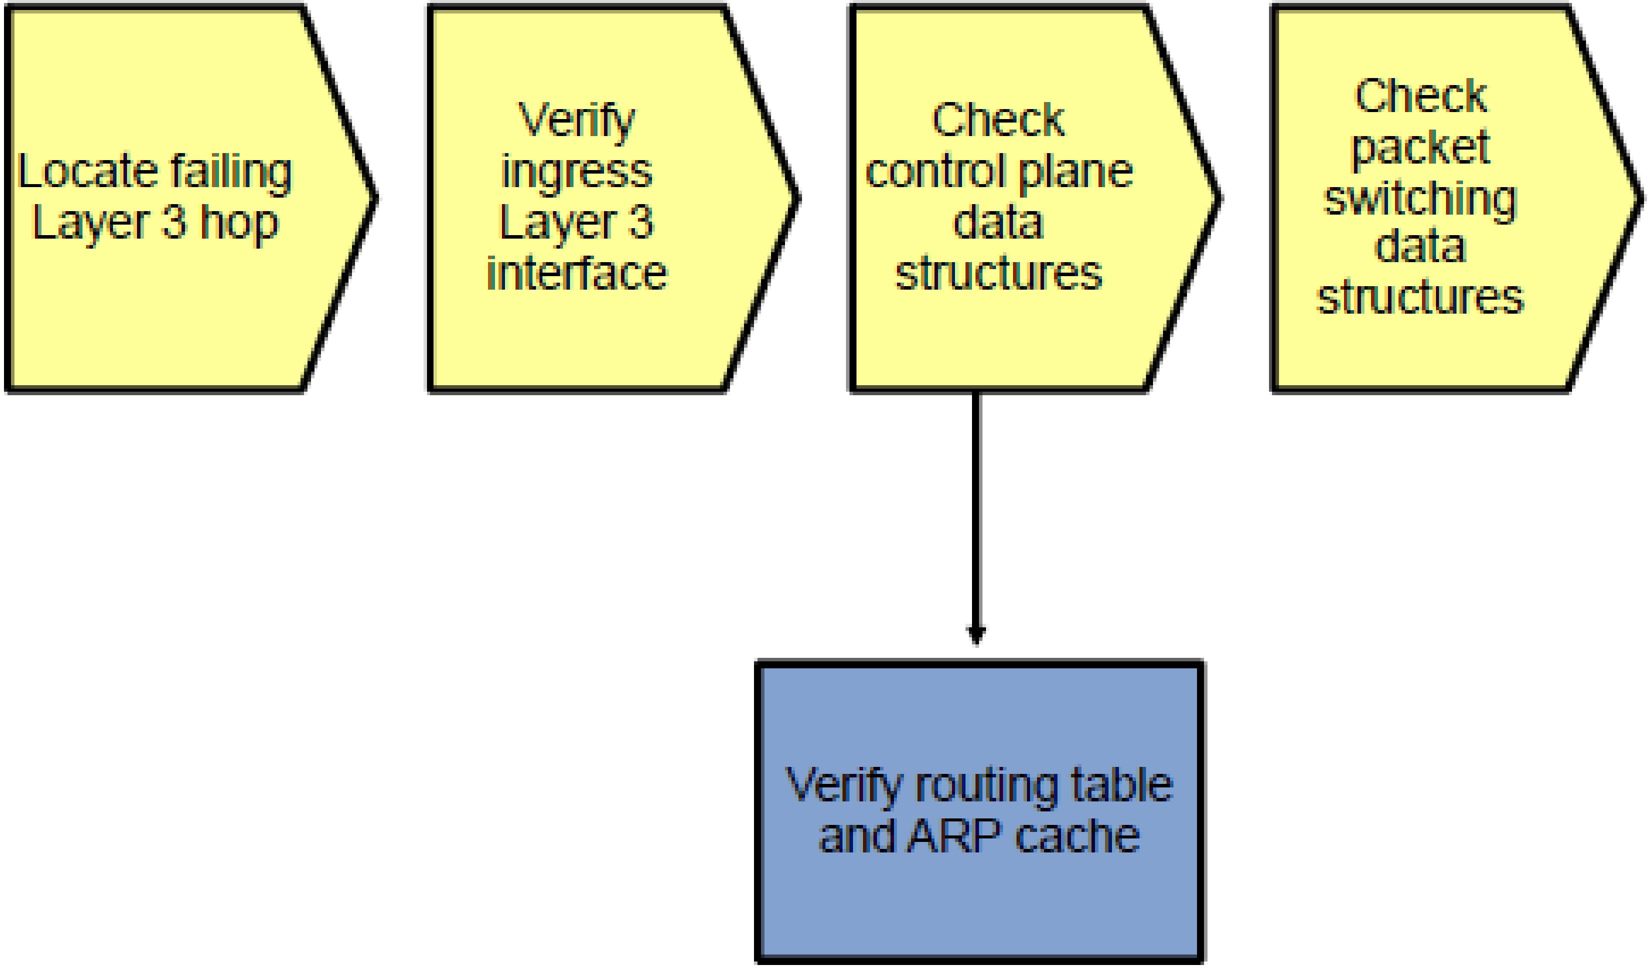

Multilayer Switching Troubleshooting Flow ④

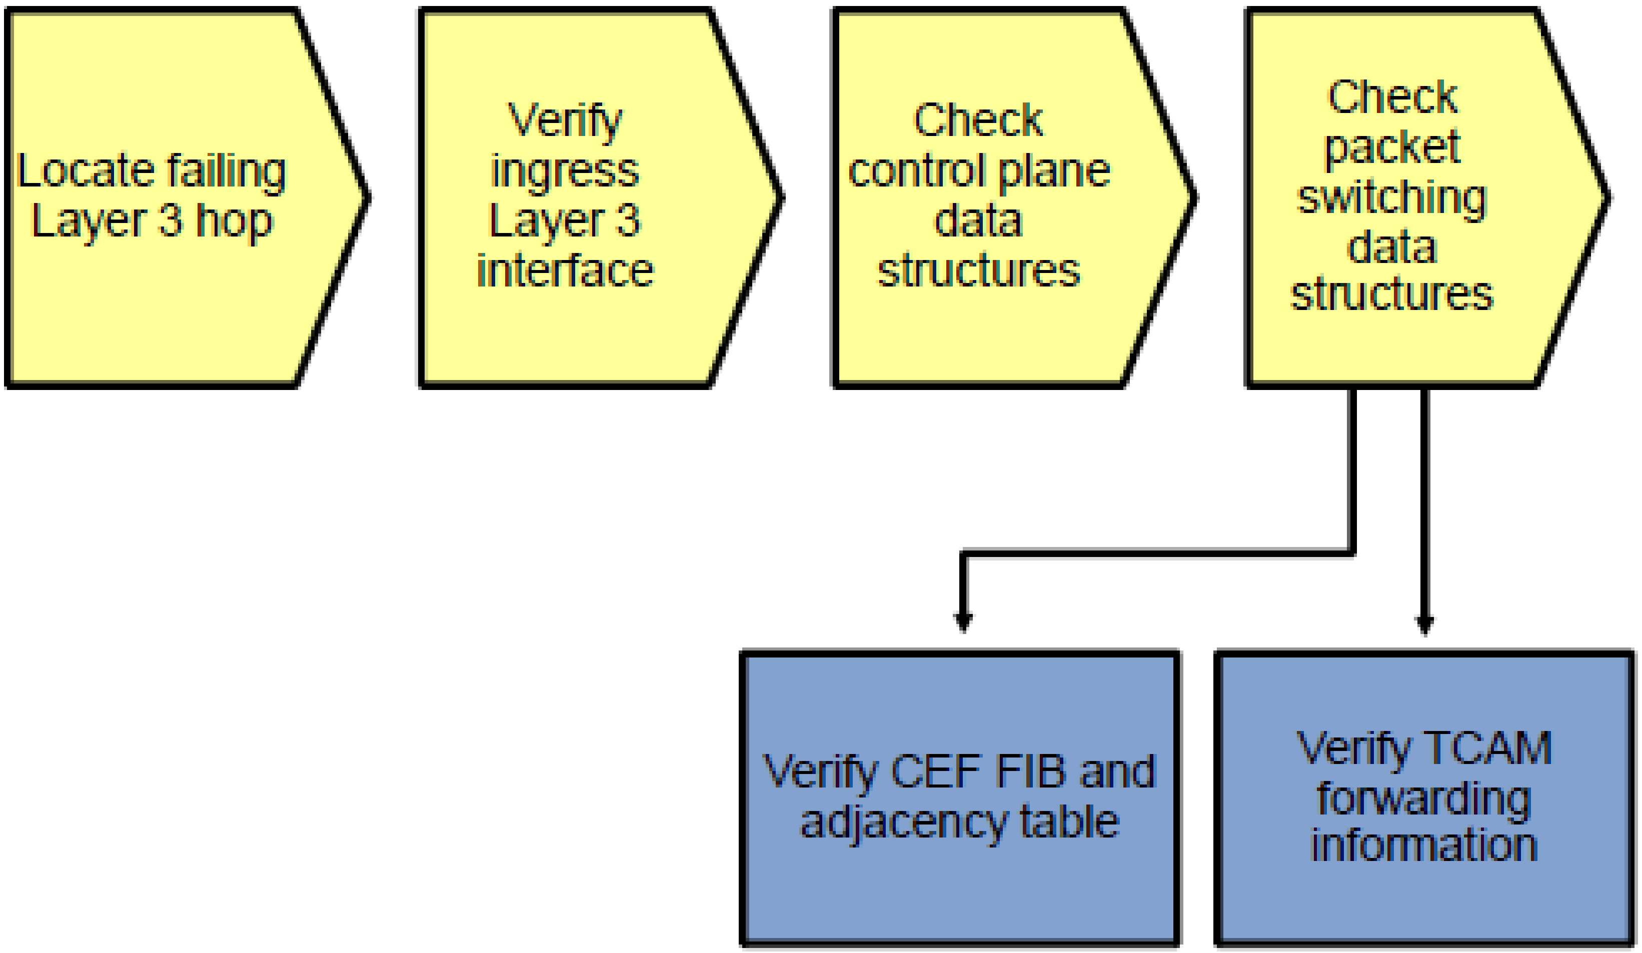

Troubleshooting First Hop Redundancy Protocols

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

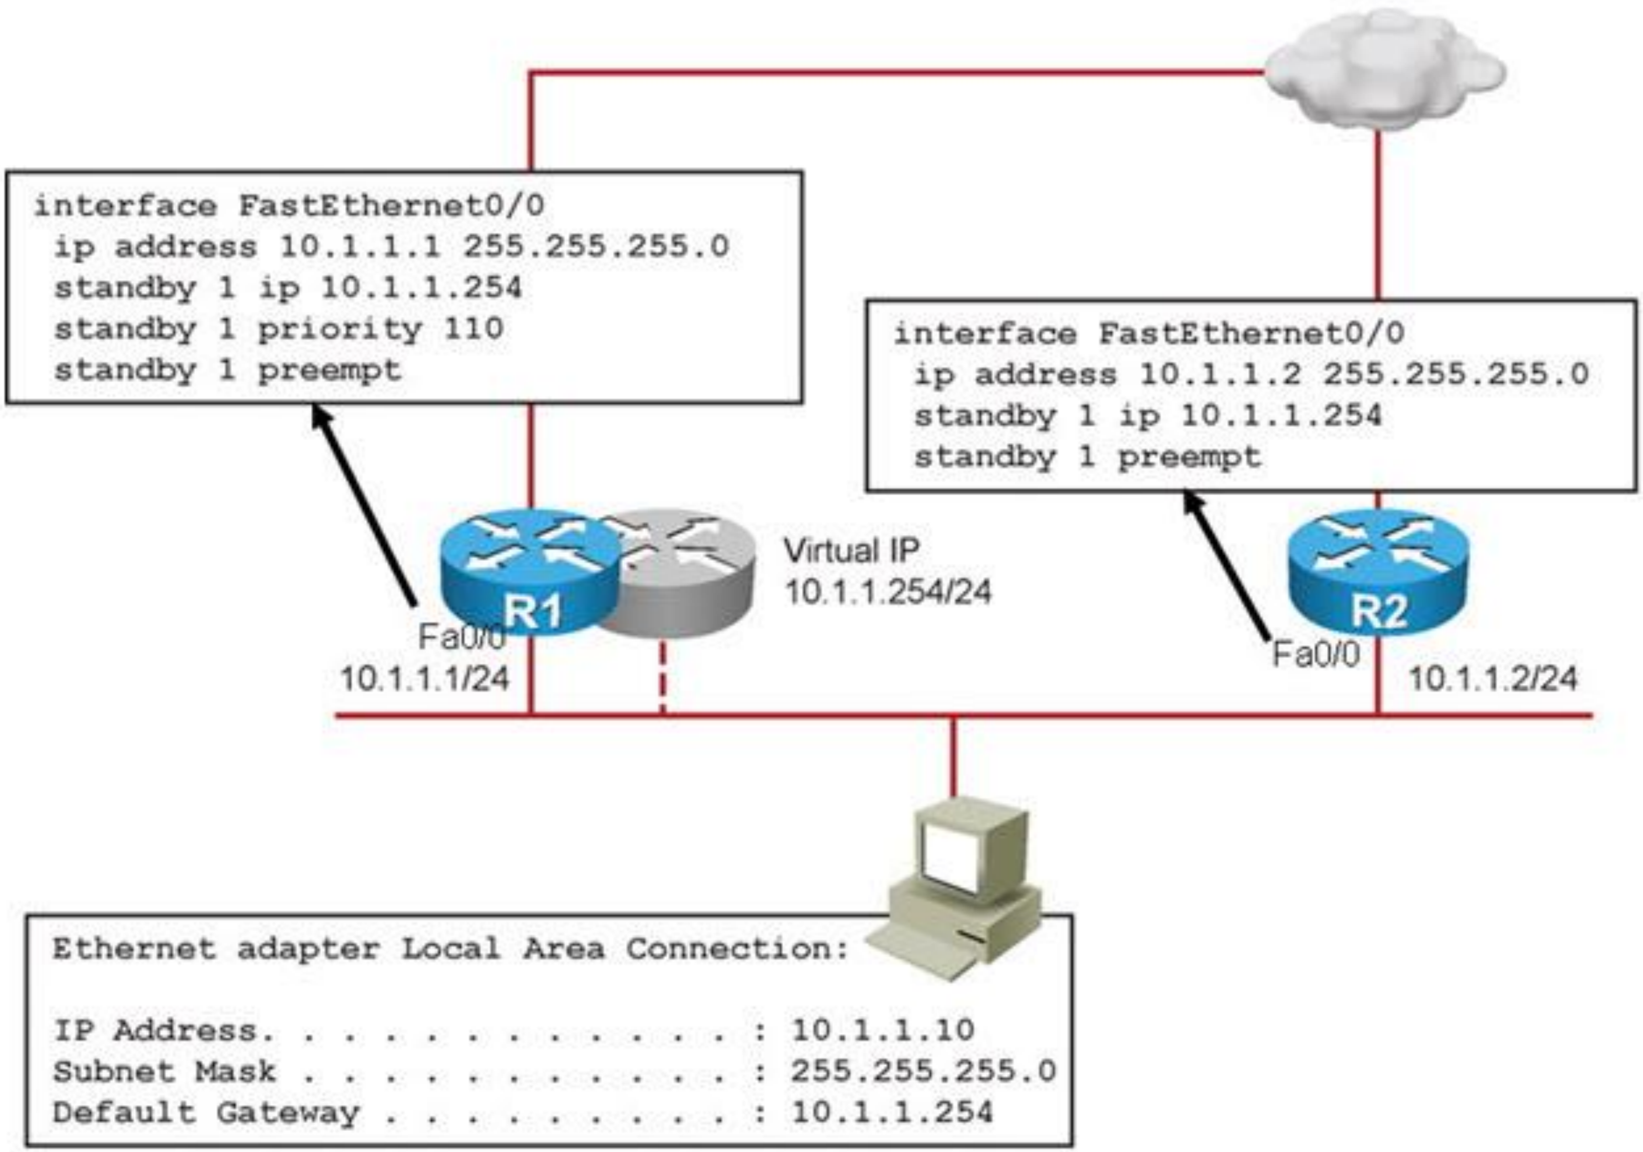

The `show standby brief` Command

```
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
```

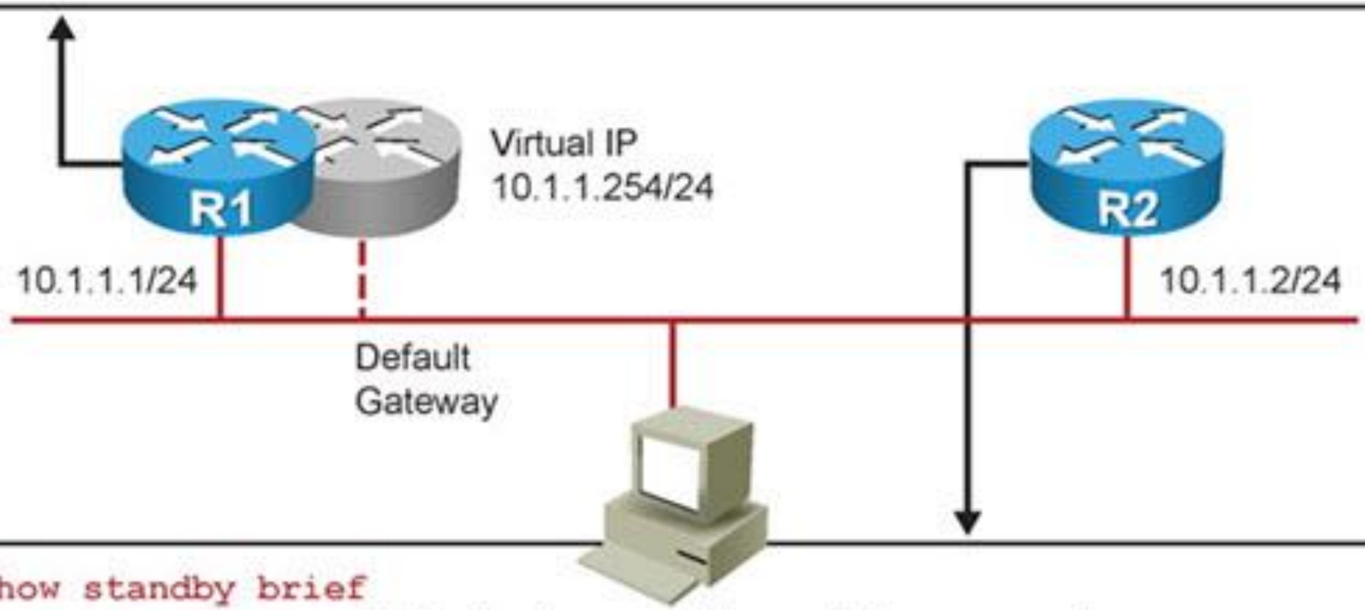


```
R2#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0       1   100  P Standby  10.1.1.1        local           10.1.1.254
```

The `show standby interface-id` Command

```
R1#show standby fa 0/0
FastEthernet0/0 - Group 1
  State is Active
    8 state changes, last state change 01:00:36
  Virtual IP address is 10.1.1.254
  Active virtual MAC address is 0000.0c07.ac01
<…output truncated…>
```

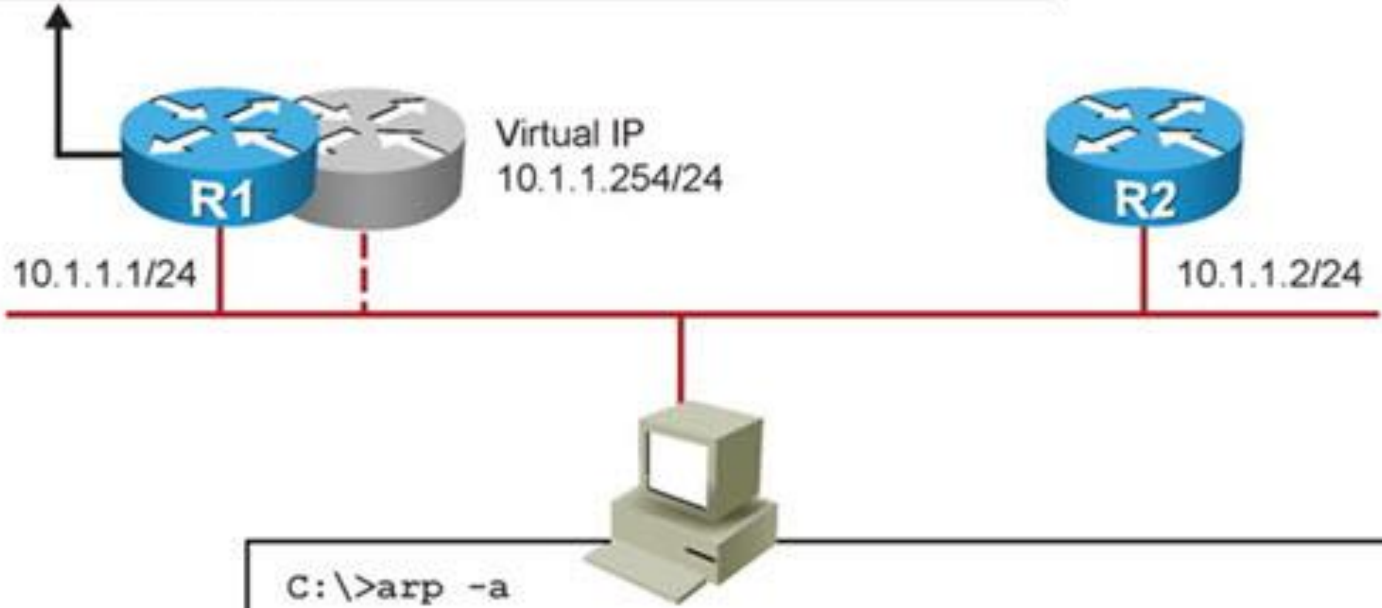


```
C:\>arp -a

Interface: 10.1.1.3 --- 0x4
  Internet Address      Physical Address      Type
  10.1.1.254            00-00-0c-07-ac-01     dynamic
```

Shutting Down FHRP Interface

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
```

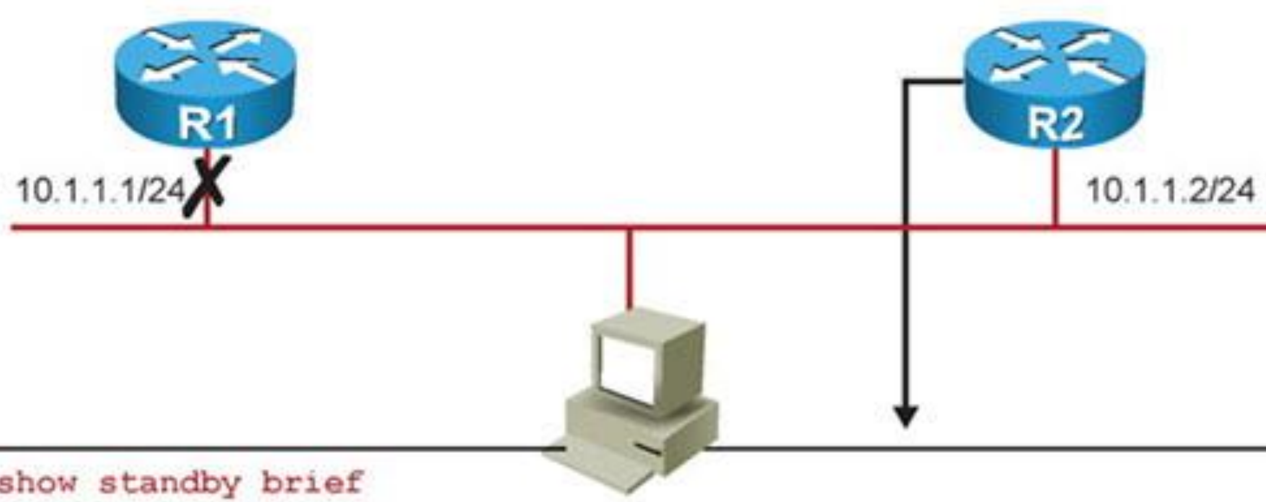


```
R2#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0       1   100  P Active   local           unknown         10.1.1.254
```

The debug standby terse Command ①

```
R2#debug standby terse
HSRP:
  HSRP Errors debugging is on
  HSRP Events debugging is on
    (protocol, redundancy, track)
  HSRP Packets debugging is on
    (Coup, Resign)
R2#
```

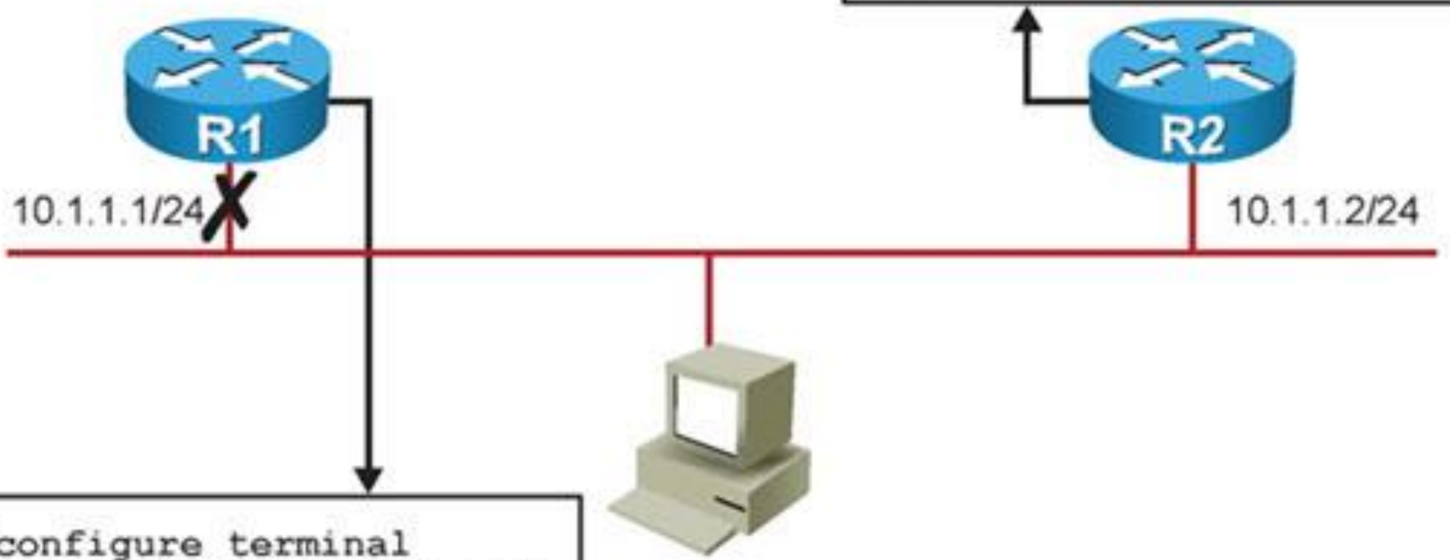


```
R1#configure terminal
R1(config)#interface fa 0/0
R1(config-if)#no shutdown
R1(config-if)#
```

The debug standby terse Command ②

```
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#
```

Operational Differences Between FHRPs

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        -
```

Diagnostic Commands Comparison

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`**

FHRP Troubleshooting Flow

FHRP Troubleshooting Flow ①

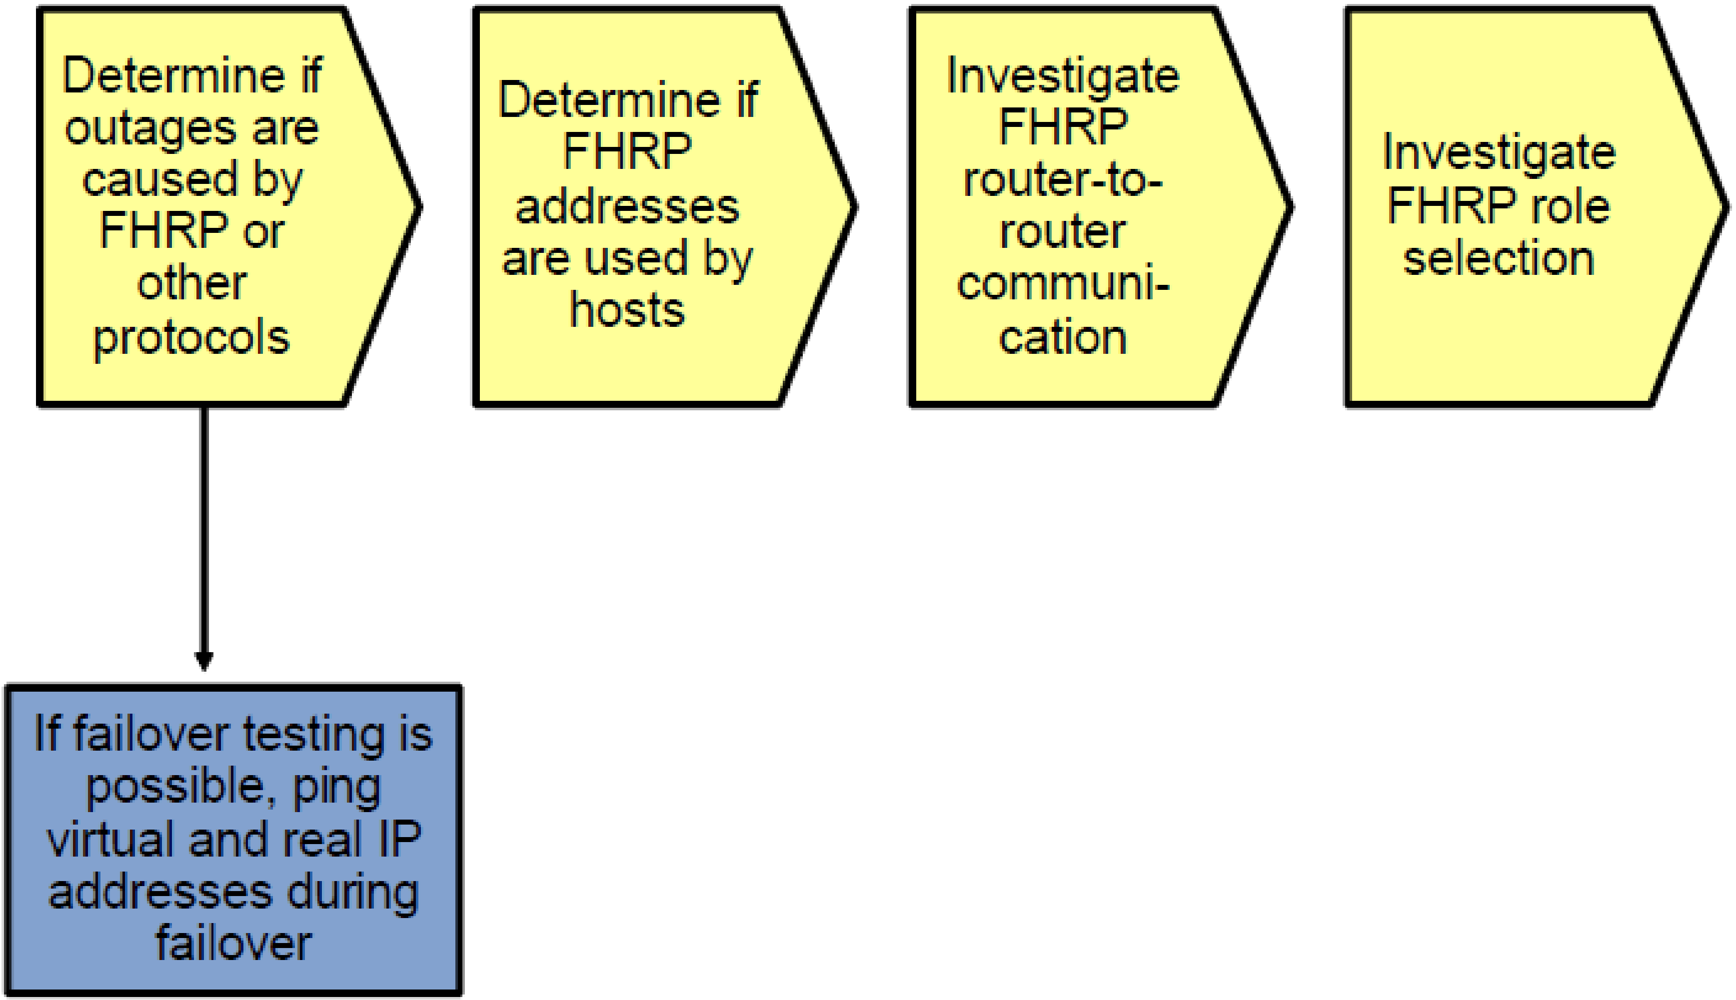

FHRP Troubleshooting Flow ①

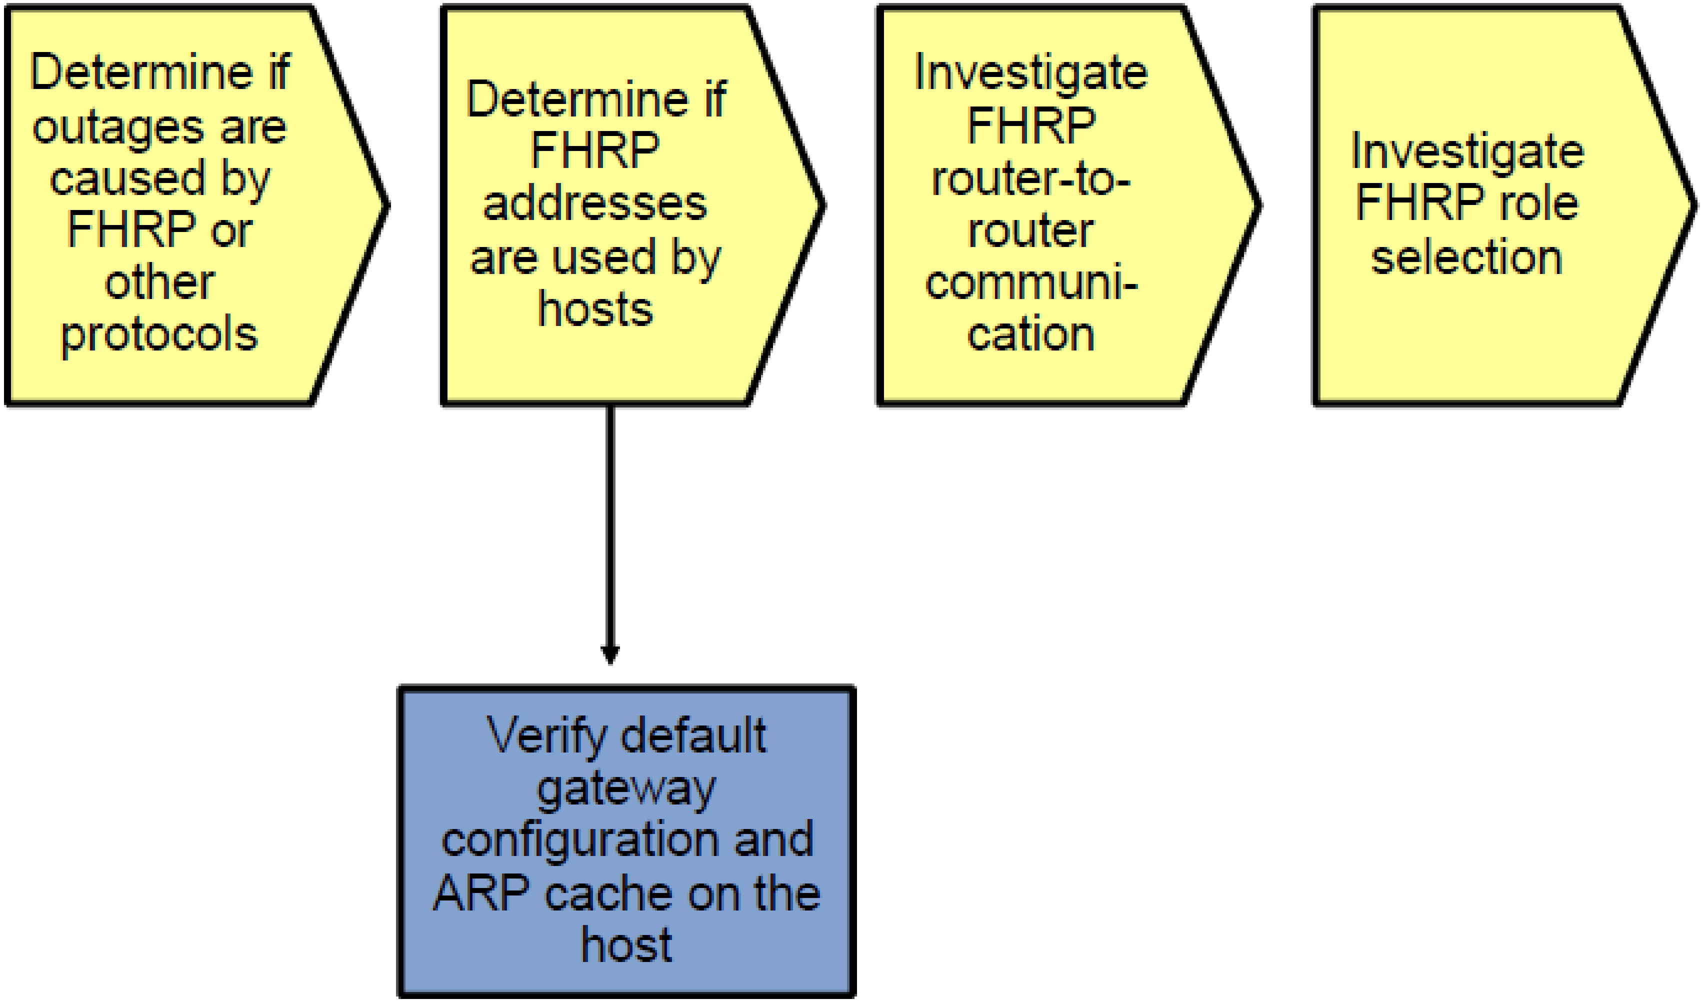

FHRP Troubleshooting Flow ①

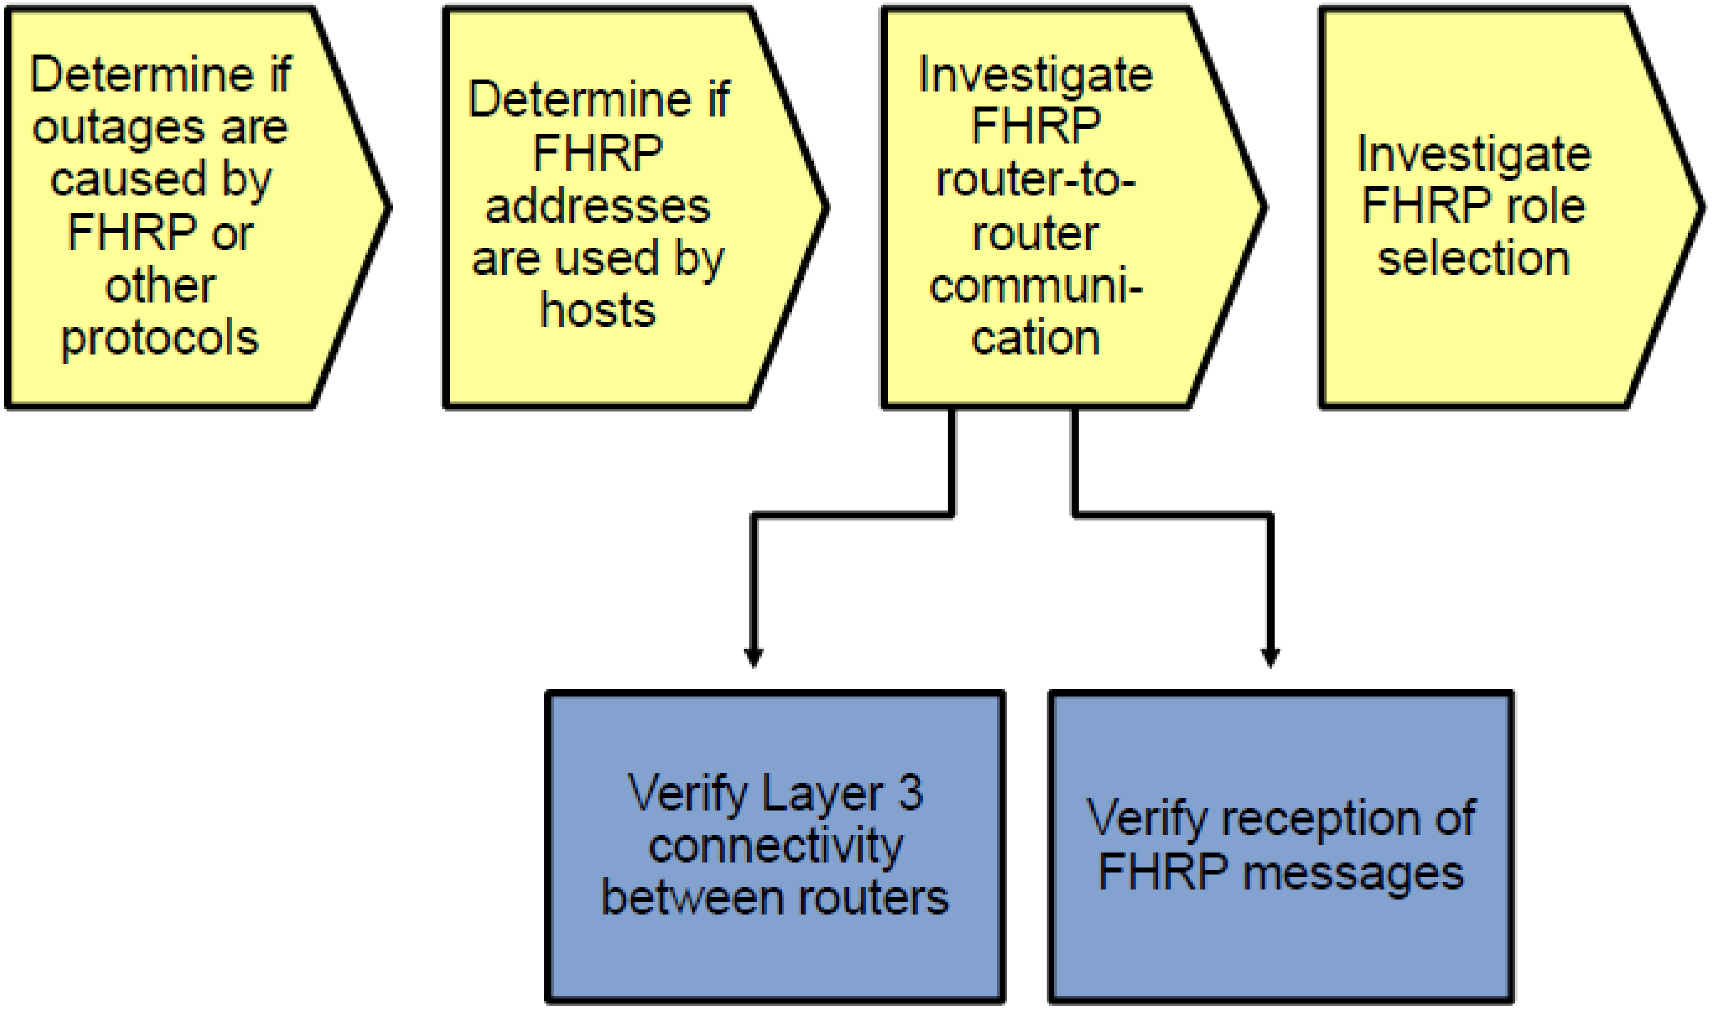

FHRP Troubleshooting Flow ①

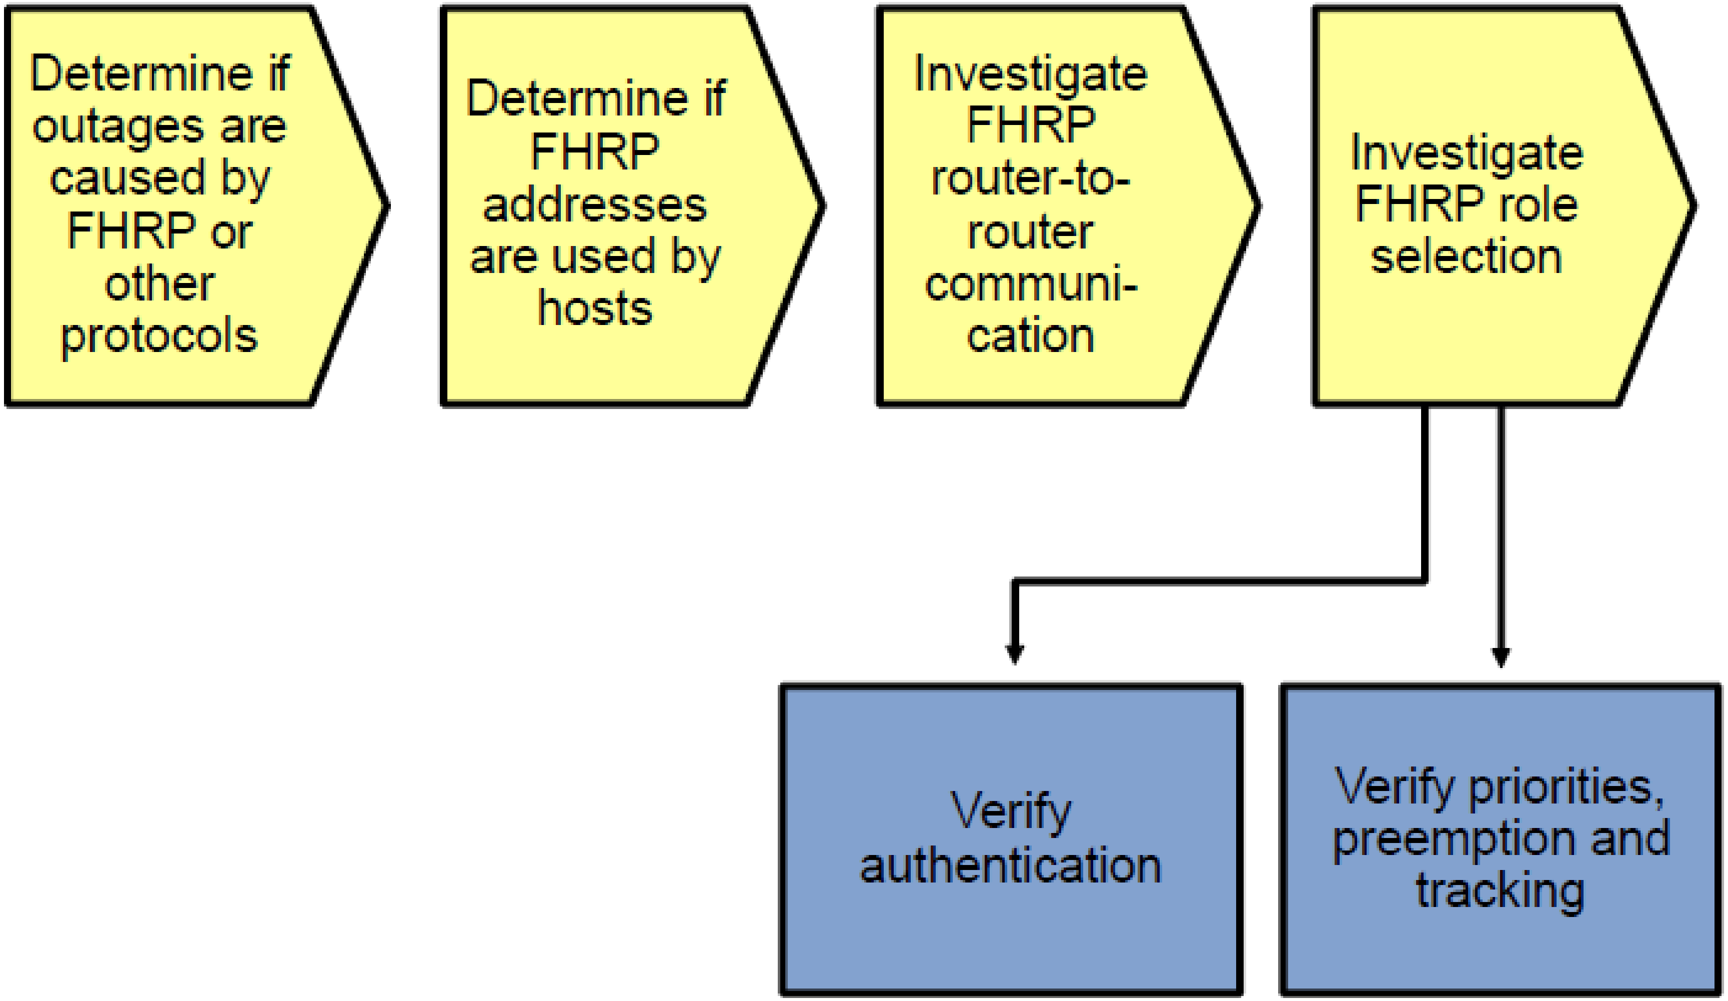

Slides adapted by [Vladimír Veselý] and [Matěj Grégr]
partially from official course materials
but the most of the credit goes to CCIE#23527 Ing. Peter Palúch, Ph.D.

The last update: 2015-10-14