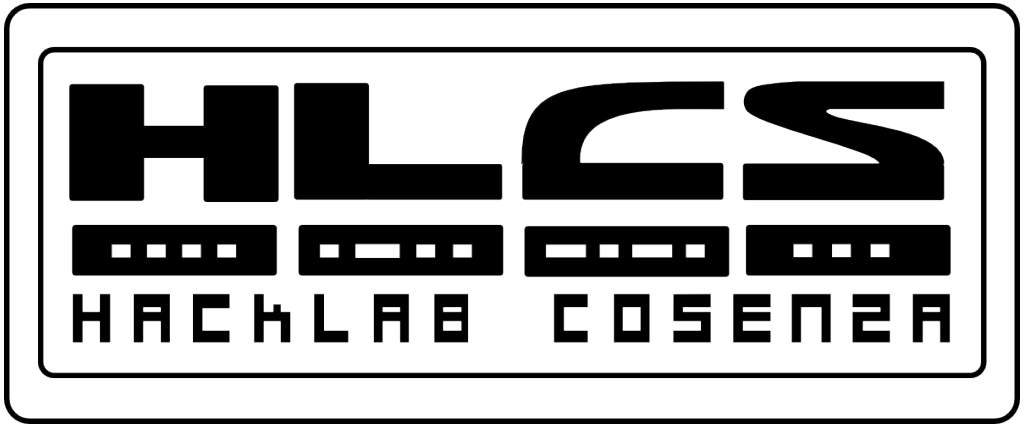

LAN Redundancy

Routing and Switching: Scaling Networks – Chapter 2

MAC Address Table

- A switch **moves traffic** based on MAC addresses.
- Each switch maintains a MAC address table in **high-speed memory**, called ***content addressable memory*** (CAM).
- The switch recreates this table every time it is activated, using both the **source MAC addresses of incoming frames** and the **port number** through which the frame entered the switch.

MAC Address Table			
fa0/1	fa0/2	fa0/3	fa0/4
260d.8c01.0000	260d.8c01.1111	260d.8c01.2222	260d.8c01.3333

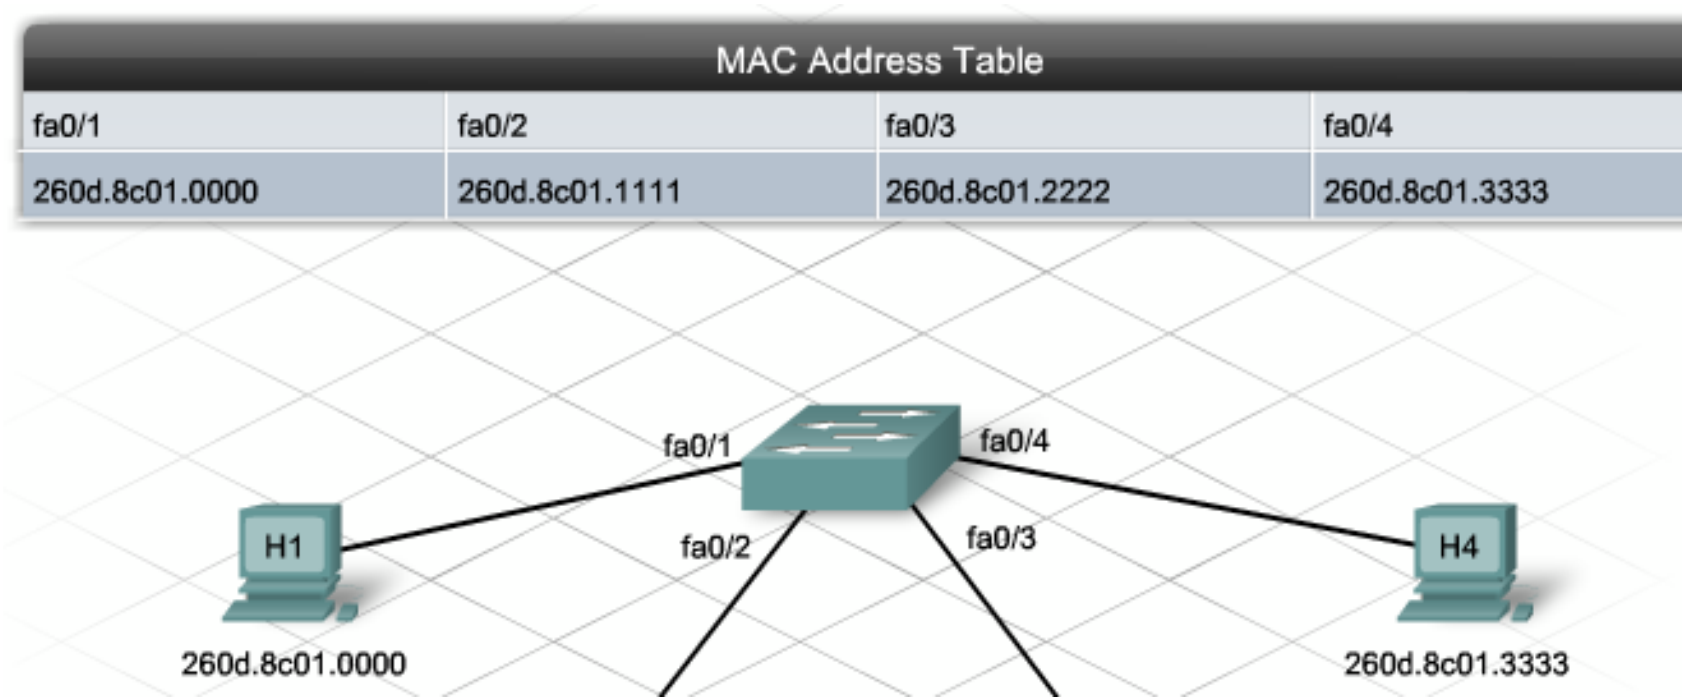

Mechanics of Switching

- As a **unicast frame** enters a port, the switch finds the source MAC address in the frame.
- It then **searches the MAC table**, looking for an entry that matches the address
- If the source MAC address is **not in the table**, the switch adds a MAC address
- Next, the switch **checks the table** for the destination MAC address.
- **If an entry exists**, the switch forwards the frame out the appropriate port number.
- **If the entry does not exist**, the switch floods the frame out every active port except the port upon which it was received

Aging Time

- The switch **deletes entries** from the MAC address table if they are not used within a certain period of time.
- The name given to this period of time is the **aging timer**
- Removal of an entry is called **aging out**.
- When the switch **adds** a MAC address and port number entry and sets the aging timer.
- If the source MAC address **already exists**, the switch resets the aging timer

Network Segmentation

- The **size of broadcast and collision domains** affect the flow of traffic.
- If a switch receives a **broadcast frame**, the switch floods it out every active interface.
- As **more switches** are connected together, the size of the broadcast domain increases.
- **Collision domains** create a similar problem. The more devices participating in a collision domain, the more collisions occur.
  - **Hubs** create large collision domains.
  - **Switches**, however, use a feature called **microsegmentation** to reduce the size of collision domains to a single switch port.

Microsegmentation

- The switch consults the switching table and establishes a virtual connection, or microsegment or virtual circuit, between the ports, maintained until the session terminates

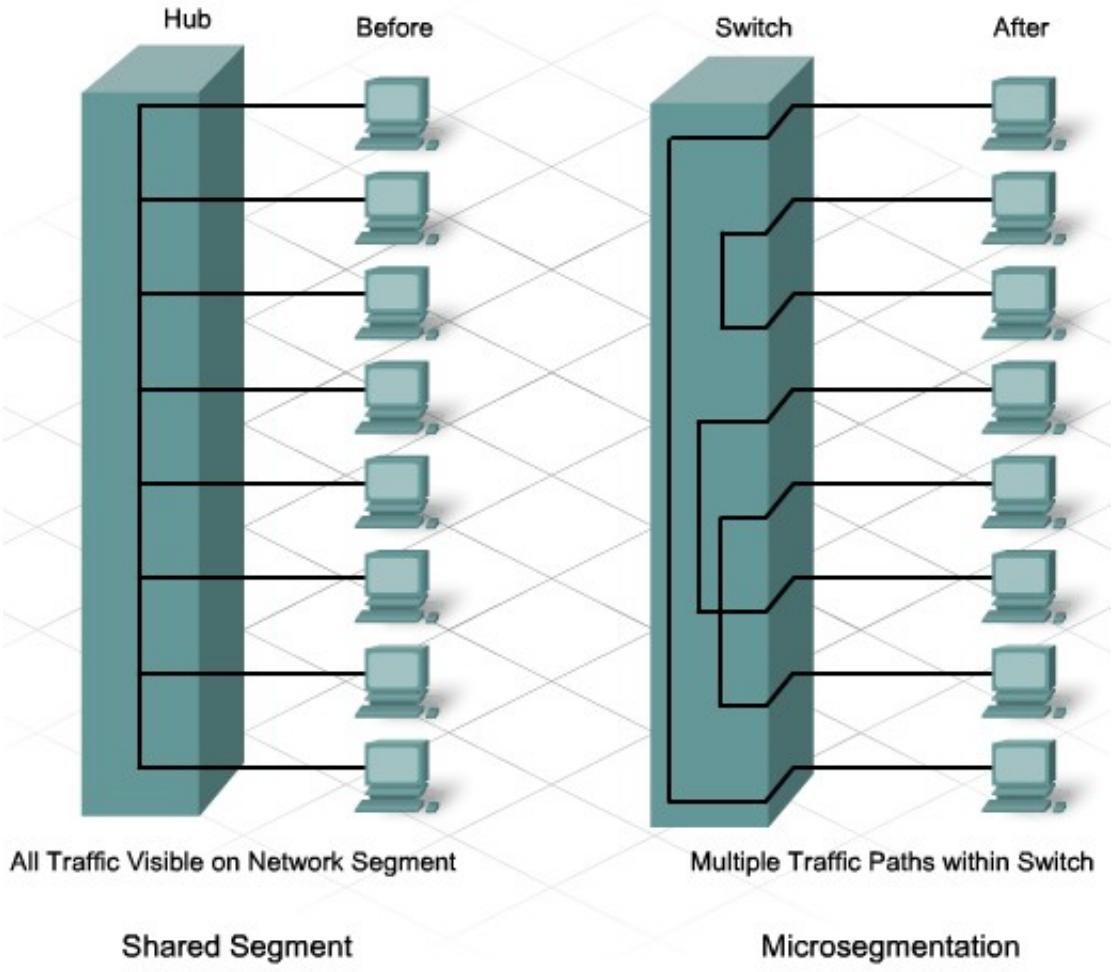

Redundant switching

- Redundant links in a switched network **reduce congestion** and support **high availability** and l**oad balancing**.
- Connecting switches together, however, can cause problems.
  - For example, the broadcast nature of Ethernet traffic creates **switching loops**. The broadcast frames go around and around in all directions, causing a **broadcast storm**

- Multiple Frame Transmissions
- MAC Database Instability

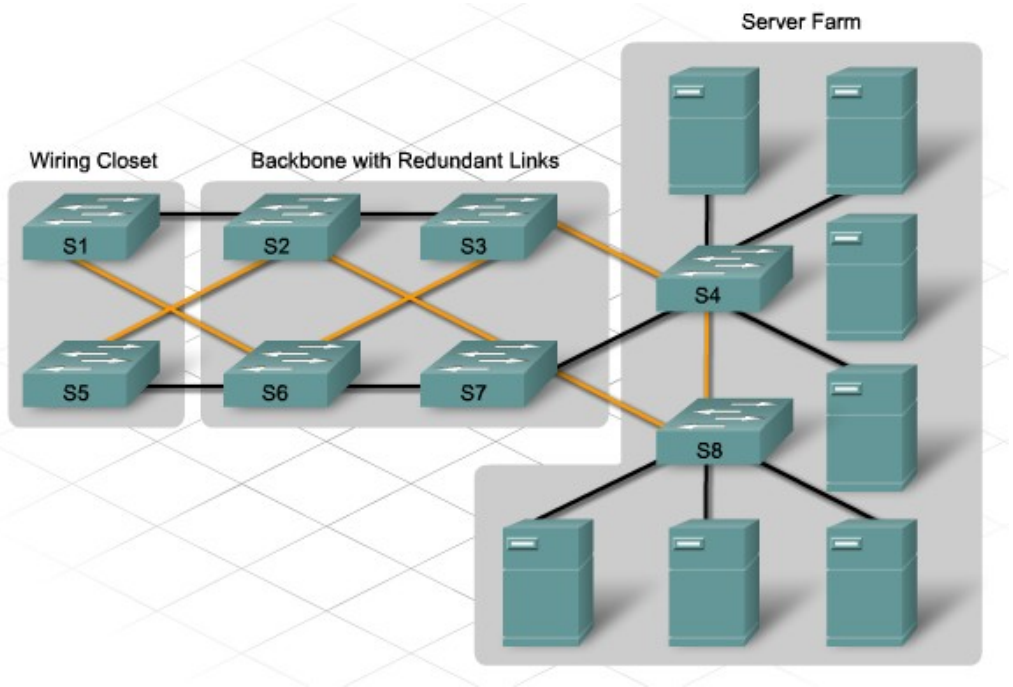

MAC Database Instability

- **Ethernet frames do not have a time to live** (TTL) attribute, like IP packets.
- As a result **they continue to propagate** between switches endlessly, or until a link is disrupted and breaks the loop.
- This **continued propagation between switches** can result in MAC database instability.
- This can occur due to **broadcast frames forwarding**.

Broadcast Storm

- A broadcast storm occurs when there are so **many broadcast frames caught in a Layer 2 loop** that all available bandwidth is consumed.
- Consequently, **no bandwidth is available for legitimate traffic** and the network becomes unavailable for data communication.
- This is an effective **denial of service**.

Multiple Frame Transmissions

- Unicast frames sent onto a looped network can result in **duplicate frames** arriving at the destination device.
- Most **upper layer protocols are not designed to recognize**, or cope with, duplicate transmissions.
- In general, protocols that make use of a sequence-numbering mechanism **assume that the transmission has failed** and that the sequence number has recycled for another communication session.

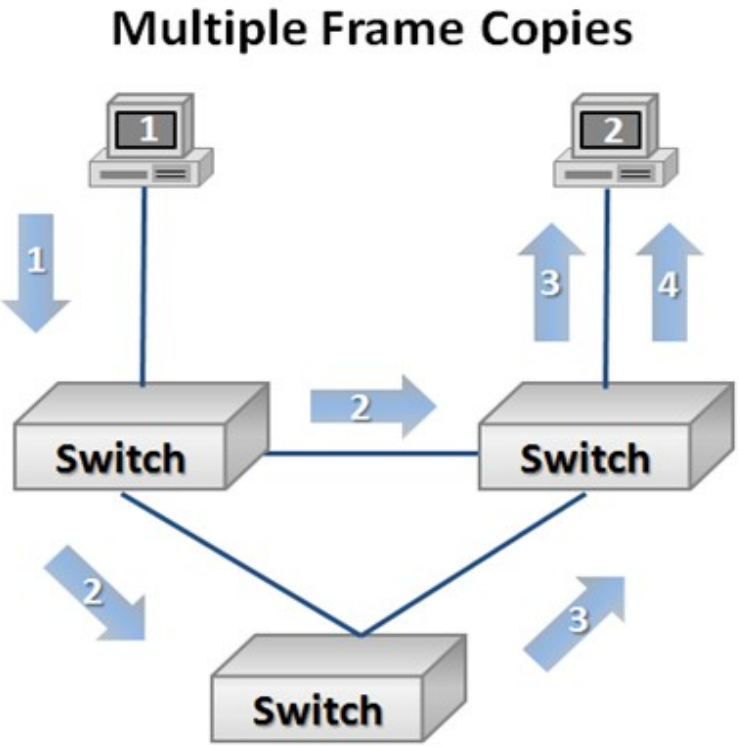

Spanning Tree Protocol (STP)

- Spanning Tree Protocol (STP, *IEEE 802.1D*) provides a mechanism for **disabling redundant links** in a switched network.
- STP is relatively self-sufficient and requires **little configuration**.
- When switches are first powered up with STP enabled, they **check the switched network** for the existence of loops.
- Switches detecting a potential loop **block some of the connecting ports**, while leaving other ports active to forward frames:
  - Forces certain interfaces into a **alternate or blocked state**
  - Leaves other interfaces in a **forwarding state**
  - Reconfigures the network by **activating the appropriate standby path**, if the forwarding path becomes unavailable

The Root of the Tree

- **The root bridge** communicates with the other switches using Bridge Protocol Data Units (BPDUs).
- **BPDUs** are frames that multicast every 2 seconds to all other switches. BPDUs contain information such as:
  - Identity of the source switch
  - Identity of the source port
  - Cumulative cost of path to root bridge
  - Value of aging timers
  - Value of the hello timer

Root bridge

- The root bridge sends out **BPDUs containing network topology information** to all other switches
- There is only one root bridge on each network, and it is elected based on the **lowest bridge ID (BID)**.
- The **bridge priority** value plus the MAC address creates the BID.
- Bridge priority has a **default** value of 32768
- Example BID would be: *32768: AA-11-BB-22-CC-33*

STP root election

- As each switch powers on, it **assumes that it is the root bridge**, and sends out BPDUs containing its BID.
  - For example, if S2 advertises a root ID that is a lower number than S1, S1 stops the advertisement of its root ID and accepts the root ID of S2. S2 is now the root bridge
- **To set priority**:
  - *S3(config)#spanning-tree vlan 1 priority 4096 (*)*
- To restore priority to default:
  - *S3(config)#no spanning-tree vlan 1 priority*

```
(*) Allowed values are:
  0     4096  8192  12288 16384 20480 24576 28672
  32768 36864 40960 45056 49152 53248 57344 61440
Because only the first 4 bits are used in bridge priority
```

STP port types

- STP designates **three types** of ports
  - **Root Port:** The port that provides the least cost path back to the root bridge becomes the root port. Switches calculate the least cost path using the bandwidth cost of each link required to reach the root bridge.
  - **Designated Port**: A designated port is a port that forwards traffic toward the root bridge but does not connect to the least cost path.
  - **Alternate and Blocked Port:** are port that does not forward traffic.

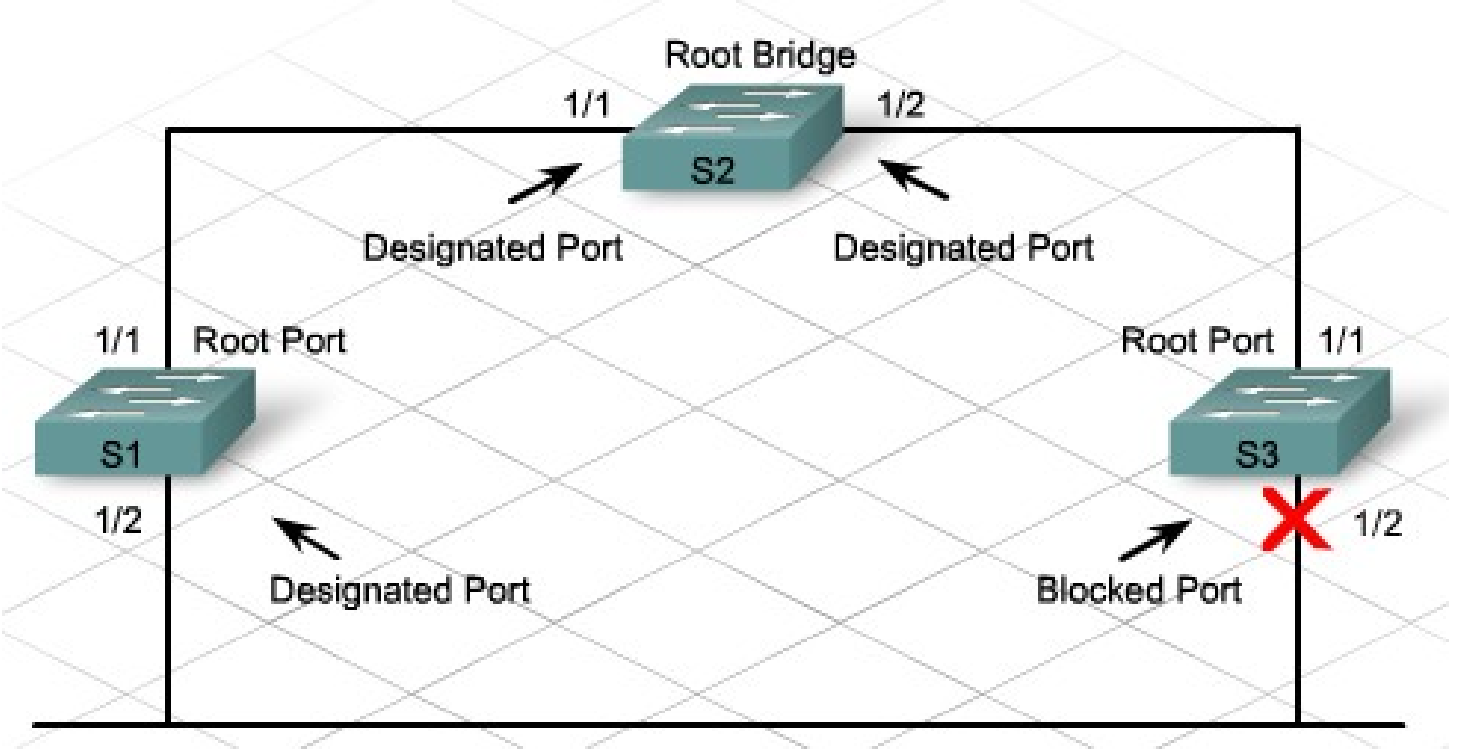

Path cost

- When the root bridge has been elected for the spanning tree instance, the Spanning Tree Algorithm (STA) starts the process of determining the best paths to the root bridge from all destinations in the broadcast domain.
- The path information is determined by **summing up the individual port costs** along the path from the destination to the root bridge.
- The ***default port costs*** are defined by the speed at which the port operates

Data rate	STP Cost (802.1D-1998)	RSTP Cost (802.1D-2004)[3]
4 Mbit/s	250	5,000,000
10 Mbit/s	100	2,000,000
16 Mbit/s	62	1,250,000
100 Mbit/s	19	200,000
1 Gbit/s	4	20,000
2 Gbit/s	3	10,000
10 Gbit/s	2	2,000

Configure Path cost on Cisco

- `Switch#configure terminal`
- `Switch(config)#interface fastEthernet 0/13`
- `Switch(config-if)#spanning-tree cost 55`

Extended System ID

- **Early implementations** of IEEE 802.1D were designed for networks that did not use VLANs.
- 802.1D was **enhanced** to include support for VLANs, requiring the VLAN ID to be included in the BPDU frame.
- **VLAN information is included** in the BPDU frame through the use of the **extended system ID**.
- All newer switches include the use of the **extended system ID by default**.
- the bridge priority field is 2 bytes or 16-bits in length;
  - 4-bits used for the bridge priority and
  - 12-bits for the extended system ID, which identifies the VLAN participating in this particular STP process.

Varieties of Spanning Tree Protocols

- **STP** is the original IEEE 802.1D version (802.1D-1998 and earlier) that provides a loop-free topology in a network with redundant links.
- **Rapid Spanning Tree Protocol (RSTP)** or IEEE 802.1w is an evolution of STP that provides faster convergence than STP
- **802.1D-2004** is an updated version of the STP standard, incorporating IEEE 802.1w.
- **PVST+** is a Cisco enhancement of STP that provides a separate 802.1D spanning tree instance for each VLAN configured in the network.
- **Rapid PVST+** is a Cisco enhancement of RSTP that uses PVST+, it provides a separate instance of 802.1w per VLAN.
- **Multiple Spanning Tree Protocol (MSTP)** is the IEEE 802.1s standard inspired by the earlier Cisco proprietary Multiple Instance STP (MISTP) implementation. MSTP maps multiple VLANs into the same spanning tree instance.

PVST+

- The original IEEE 802.1D standard defines a **Common Spanning Tree (CST)** that assumes only one spanning tree instance for the entire switched network, regardless of the number of VLAN.
- Cisco developed PVST+ so that a network can run an **independent instance of the Cisco implementation of IEEE 802.1D for each VLAN in the network**.
- With PVST+, it is possible for one **trunk port** on a switch to be *blocking* for a VLAN while *not blocking* for other VLANs.

STP switch port states

- **Blocking** - A port that would cause a switching loop, no user data is sent or received. BPDU data is still received in blocking state. Default at power on. *Led is Steady amber*, 20 sec.
- **Listening** - The switch processes BPDUs and awaits possible new information that would cause it to return to the blocking state. It does not populate the MAC address table and it does not forward frames. *Led is Blinking amber*, 15 sec.
- **Learning** - While the port does not yet forward frames, it does learn source addresses from frames received and adds them to the filtering database (switching database). It populates the MAC Address table, but does not forward frames. *Led is Blinking amber*, 15 sec.
- **Forwarding** - A port receiving and sending data, normal operation. STP still monitors incoming BPDUs that would indicate it should return to the blocking state to prevent a loop. *Led is Steady green*.
- **Disabled** - Not strictly part of STP, a network administrator can manually disable a port

PVST+ steps

- For each VLAN in a switched network, PVST+ performs four steps to provide a loop-free logical network topology:
  - **Elects one root bridge -** Only one switch can act as the root bridge for a given VLAN
  - **Selects the root port on each non-root bridge -** The root port is the lowest-cost path from the non-root bridge to the root bridge, indicating the direction of the best path to the root bridge. *Root ports are normally in the forwarding state.*
  - **Selects the designated port on each segment -** On each link, STP establishes one designated port. The designated port is selected on the switch that has the lowest-cost path to the root bridge. *Designated ports are normally in the forwarding state.*
  - **The remaining ports in the switched network are alternate ports -** Alternate ports normally remain in the blocking state, to logically break the loop topology. When a port is in the blocking state, it does not forward traffic, but can still process received BPDU messages.

Rapid Spanning Tree Protocol

- RSTP, defined in **IEEE 802.1w**, significantly s**peeds the recalculation** of the spanning tree.
- Unlike PortFast, UplinkFast, and BackboneFast, **RSTP is not proprietary**.
- Reconfiguration of the spanning tree by RSTP occurs in **less than 1 second**, as compared to 50 seconds in STP.
- To speed up the recalculation process, RSTP reduces the number of port states to three: **discarding, learning and forwarding**.
  - The discarding state is similar to three of the original STP states: blocking, listening, and disabled.
- RSTP also introduces the concept of **active topology**. All ports that are not discarding are part of the active topology and will immediately transit to the forwarding state.

Rapid PVST+

- Rapid PVST+ is the Cisco implementation of RSTP on a **per-VLAN basis**.
- With Rapid PVST+, an independent instance of RSTP runs for each VLAN.

RSTP BPDU

- Because BPDUs are used as a keepalive mechanism, **three consecutively missed BPDUs indicate lost connectivity** between a bridge and its neighboring root or designated bridge.
- Like STP, an RSTP switch sends a BPDU with its current information **every Hello time period** (*two seconds, by default*), even if the RSTP bridge does not receive any BPDUs from the root bridge.

Flag byte of version 2 BPDU

- **Bits 0 and 7** are used for topology change and acknowledgment as they are in the original 802.1D.
- **Bits 1 and 6** are used for the Proposal Agreement process (used for rapid convergence).
- **Bits from 2 to 5** encode the role and state of the port.
- **Bits 4 and 5** are used to encode the port role using a 2-bit code.

Field Bit	Bit
Topology Change	0
Proposal	1
Port Role Unknown Port Alternate or Backup Port Root Port Designated Port	2-3 00 01 10 11
Learning	4
Forwarding	5
Agreement	6
Topology Change Acknowledgment	7

Edge Ports

- An RSTP edge port is a switch port that is never intended to be connected to another switch device.
- It immediately transitions to the forwarding state when enabled skipping the time-consuming original 802.1D listening and learning port states
- As soon as the bridge detects a BPDU coming to an edge port, the port becomes a non-edge port.
- On Cisco switches use the **`spanning-tree portfast`** command for edge port configuration.

Configuring PVST+ Bridge ID

- `Switch(config)#`**`spanning-tree vlan 1 priority ?`**

  `<0-61440>` `bridge priority in increments of 4096`

- `Switch(config)#`**`spanning-tree vlan 1 root ?`**

  `primary` `Configure this switch as primary root for this spanning tree`

  `secondary` `Configure switch as secondary root`

- Primary and secondary root bridges are configured for load balancing (high availability)
- `Switch#`**`show spanning-tree`**

PortFast

- When a switch port is configured with PortFast that port transitions **from blocking to forwarding state immediately,** bypassing the usual 802.1D STP transition states (the listening and learning states)
- Because the purpose of PortFast is to minimize the time that access ports must wait for spanning tree to converge, it should **only be used on access ports**.
- If you enable PortFast on a port connecting to another switch, you **risk creating a spanning tree loop**.
- Configure PortFast with the interface mode command:
- `Switch(config-if)#`**`spanning-tree portfast ?`**

```
disable  Disable portfast for this interface
trunk    Enable portfast on the interface even in trunk mode
<cr>
```

BPDU guard

- In a valid PortFast configuration, **BPDUs should never be received,** because that would indicate that another bridge or switch is connected to the port, potentially causing a spanning tree loop.
- **BPDU guard puts the port in an error-disabled state on receipt of a BPDU**.
- This will effectively **shut down** the port.
- The BPDU guard feature provides a secure response to invalid configurations because ***you must manually put the interface back into service***.

Configuring Rapid PVST+

- The default spanning tree configuration on a Catalyst 2960 Series switch is PVST+.
- The **`spanning-tree mode rapid-pvst`** global configuration mode command is the one required command for the Rapid PVST+ configuration.

Analyze the STP topology

- To analyze the STP topology, follow these steps:
  - **Step 1. Discover the Layer 2 topology**. Use network documentation if it exists or use the **`show cdp neighbors`** command to discover the Layer 2 topology.
  - **Step 2.** After discovering the Layer 2 topology, use STP knowledge to **determine the expected Layer 2 path**. It is necessary to know which switch is the root bridge.
  - **Step 3**. Use the **`show spanning-tree vlan`** command to determine **which switch is the root bridge**.
  - **Step 4**. Use the **`show spanning-tree vlan`** command on all switches to **find out which ports are in blocking or forwarding state** and confirm your expected Layer 2 path.

STP Cisco commands

- *show spanning-tree* - Displays root ID, bridge ID, and port states
- *show spanning-tree summary* - Displays a summary of port states
- *show spanning-tree root (*)* - Displays the status and configuration of the root bridge
- *show spanning-tree detail* - Displays detailed port information
- *show spanning-tree interface* - Displays STP interface status and configuration
- *show spanning-tree blockedports (*)* - Displays blocked ports

(*) Not in Packet Tracer

First Hop Redundancy Protocols

- One way **prevent a single point of failure** at the default gateway, is to implement a virtual router.
- To implement this type of **router redundancy**, multiple routers are configured to work together to present the illusion of a single router to the hosts on the LAN
- By sharing an IP address and a MAC address, two or more routers can act as a single **virtual router**

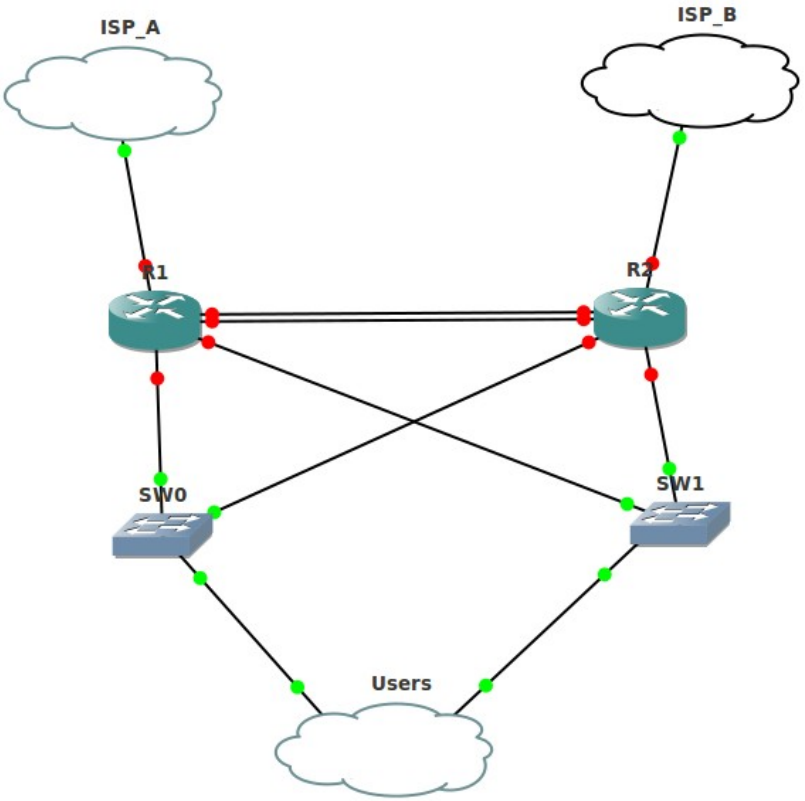

Hot Standby Router Protocol (HSRP)

- A Cisco-proprietary FHRP designed to allow for **transparent failover of a first-hop IPv4 device**.
- HSRP provides **high network availability** by providing first-hop routing redundancy for IPv4 hosts on networks configured with an IPv4 default gateway address.
- HSRP is used in a group of routers for selecting an **active device** and a **standby device**.
- The **standby** device is the device that **takes over when the active device fails**, or when pre-set conditions are met.
- The function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router fails.
- **HSRP for IPv6** supports **IPv6**

Virtual Router Redundancy Protocol version 2 (VRRPv2)

- A **non-proprietary election protocol** that dynamically assigns responsibility for one or more virtual routers to the VRRP routers on an IPv4 LAN.
- This allows several routers on a multiaccess link to use the same virtual IPv4 address.
- A VRRP router is configured to run the VRRP protocol in conjunction with one or more other routers attached to a LAN.
- In a VRRP configuration, one router is elected as the **virtual router master**, with the other routers acting as **backups**, in case the virtual router master fails.
- **VRRPv3** supports **IPv6**

Gateway Load Balancing Protocol (GLBP)

- Cisco-**proprietary** FHRP that protects data traffic from a failed router or circuit, like HSRP and VRRP, while also **allowing load balancing** (also called load sharing) between a group of redundant routers.
- **GLBP for IPv6** supports **IPv6**

HSRP Verification

```
Router# show standby
Ethernet0/1 - Group 1
  State is Active
   2 state changes, last state change 00:30:59
  Virtual IP address is 10.1.0.20
   Secondary virtual IP address 10.1.0.21
  Active virtual MAC address is 0004.4d82.7981
   Local virtual MAC address is 0004.4d82.7981 (bia)
  Hello time 4 sec, hold time 12 sec
   Next hello sent in 1.412 secs
  Gratuitous ARP 14 sent, next in 7.412 secs
  Preemption enabled, min delay 50 sec, sync delay 40 sec
  Active router is local
  Standby router is 10.1.0.6, priority 75 (expires in 9.184 sec)
  Priority 95 (configured 120)
   Tracking 2 objects, 0 up
      Down Interface Ethernet0/2, pri 15
      Down Interface Ethernet0/3
Group name is "HSRP1" (cfgd)
Follow by groups:
Et1/0.3 Grp 2 Active 10.0.0.254 0000.0c07.ac02 refresh 30 secs
(next 19.666)
Et1/0.4 Grp 2 Active 10.0.0.254 0000.0c07.ac02 refresh 30 secs
(next 19.491)
  Group name is "HSRP1", advertisement interval is 34 sec
```

GLBP Verification

```
Router# show glbp
FastEthernet0/1 - Group 1
  State is Active
    1 state change, last state change 00:02:34
  Virtual IP address is 192.168.2.100
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.288 secs
  Redirect time 600 sec, forwarder timeout 14400 sec
  Preemption disabled
  Active is local
  Standby is 192.168.2.2, priority 100 (expires in 8.640 sec)
  Priority 100 (default)
  Weighting 100 (default 100), thresholds: lower 1, upper 100
  Load balancing: round-robin
  Group members:
    001e.7aa3.5e71 (192.168.2.1) local
    001e.7aa3.5f31 (192.168.2.2)
  There are 2 forwarders (1 active)
  Forwarder 1
    State is Active
      1 state change, last state change 00:02:23
    MAC address is 0007.b400.0101 (default)
    Owner ID is 001e.7aa3.5e71
    Redirection enabled
Preemption enabled, min delay 30 sec
    Active is local, weighting 100
```

End of lesson