

Spanning Tree Protocol

Overview

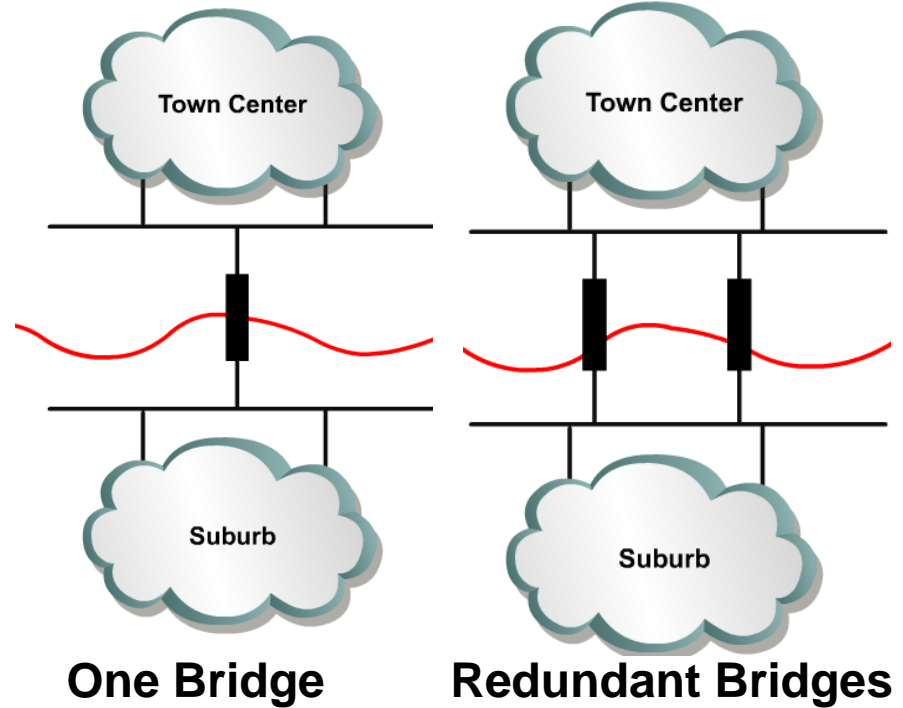
- Define redundancy and its importance in networking
- Describe the key elements of a redundant networking topology
- Define broadcast storms and describe their impact on switched networks
- Define multiple frame transmissions and describe their impact on switched networks
- Identify causes and results of MAC address database instability
- Identify the benefits and risks of a redundant topology
- Describe the role of spanning tree in a redundant-path switched network
- Identify the key elements of spanning tree operation
- Describe the process for root bridge election
- List the spanning-tree states in order
- Compare Spanning-Tree Protocol and Rapid Spanning-Tree Protocol

Redundancy



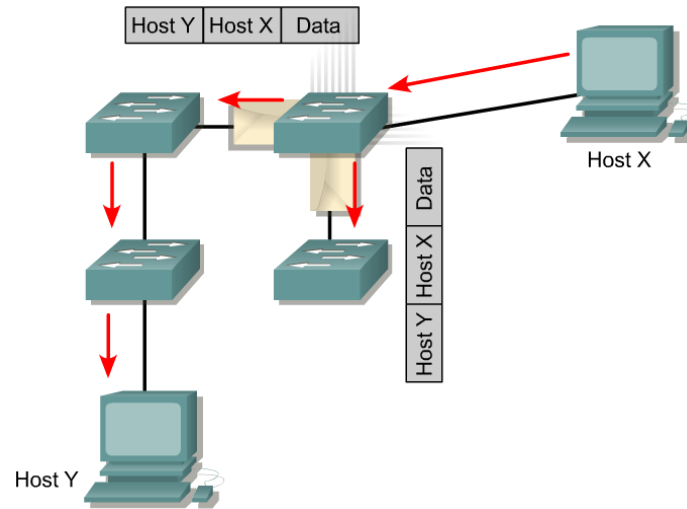
- Achieving such a goal requires extremely reliable networks.
- Reliability in networks is achieved by reliable equipment and by designing networks that are tolerant to failures and faults.
- The network is designed to reconverge rapidly so that the fault is bypassed.
- Fault tolerance is achieved by redundancy.
- Redundancy means to be in excess or exceeding what is usual and natural.

Redundant topologies



- A network of roads is a global example of a redundant topology.
- If one road is closed for repair there is likely an alternate route to the destination

Types of Traffic




Unknown Unicast

Types of traffic (Layer 2 perspective)


- **Known Unicast:** Destination addresses are in Switch Tables
- **Unknown Unicast:** Destination addresses are not in Switch Tables
- **Multicast:** Traffic sent to a group of addresses
- **Broadcast:** Traffic forwarded out all interfaces except incoming interface.

Redundant switched topologies



- **Switches learn the MAC addresses of devices on their ports so that data can be properly forwarded to the destination.**
- **Remember: switches use the Source MAC address to learn where the devices are, and enters this information into their MAC address tables.**
- Switches will flood frames for unknown destinations until they learn the MAC addresses of the devices.
- Broadcasts and multicasts are also flooded. (Unless switch is doing Multicast Snooping or IGMP)
- A redundant switched topology **may** (STP disabled) cause broadcast storms, multiple frame copies, and MAC address table instability problems.

Broadcast Storm




Broadcast storm: “A state in which a message that has been broadcast across a network results in even more responses, and each response results in still more responses in a snowball effect.”
www.webopedia.com

A broadcast storm because Spanning Tree Protocol is not enabled:


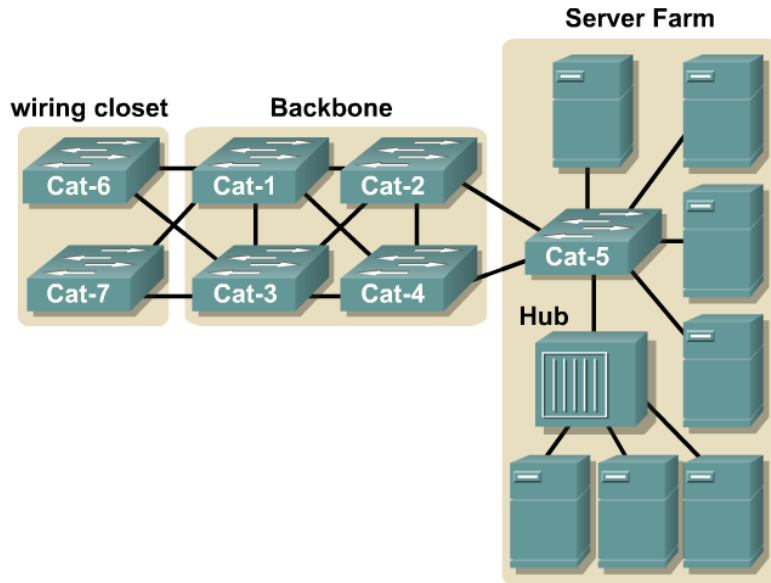
- Broadcasts and multicasts can cause problems in a switched network.
- If Host X sends a broadcast, like an ARP request for the Layer 2 address of the router, then Switch A will forward the broadcast out all ports.
- Switch B, being on the same segment, also forwards all broadcasts.
- Switch B sees all the broadcasts that Switch A forwarded and Switch A sees all the broadcasts that Switch B forwarded.
- Switch A sees the broadcasts and forwards them.
- Switch B sees the broadcasts and forwards them.
- The switches continue to propagate broadcast traffic over and over.
- This is called a broadcast storm.

Multiple frame transmissions

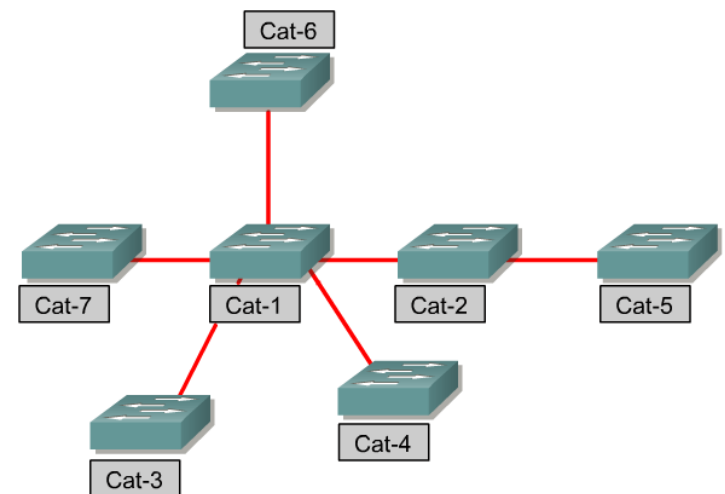


- In a redundant switched network it is possible for an end device to receive multiple frames.
- Assumptions:
 - Spanning Tree Protocol is not enabled
 - MAC address of Router Y has been timed out by both switches.
 - Host X still has the MAC address of Router Y in its ARP cache
- Host X sends a unicast frame to Router Y.



Redundant topology and spanning tree



- Unlike IP, in the Layer 2 header there is **no Time To Live (TTL)**.
- The solution is to **allow physical loops (redundant physical connections) but create a loop free logical topology**.
- The loop free **logical topology created is called a tree**.
- This topology is a star or extended star logical topology, the **spanning tree of the network**.




Redundant topology and spanning tree

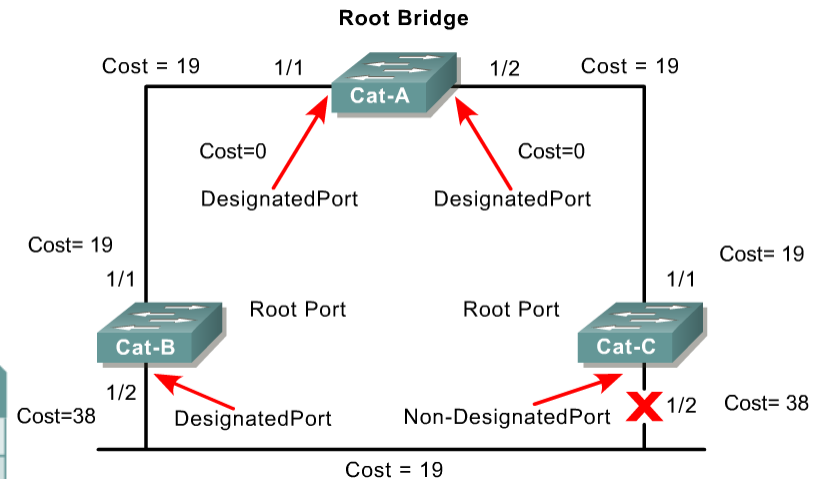


- It is a spanning tree because all devices in the network are reachable or spanned.
- The algorithm used to create this loop free logical topology is the **spanning-tree algorithm**.
- This algorithm can take a “relatively” long time to converge.
- A new algorithm called the **rapid spanning-tree algorithm** is being introduced to reduce the time for a network to compute a loop free logical topology. (later)

Spanning-Tree Protocol (STP)



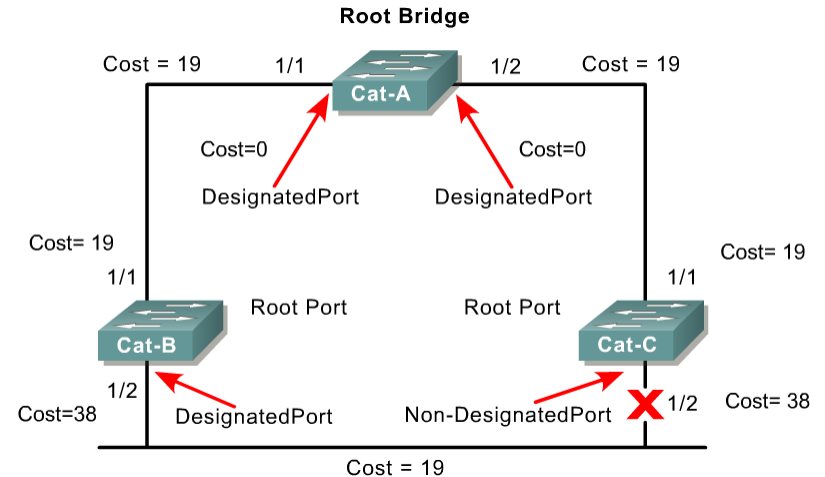
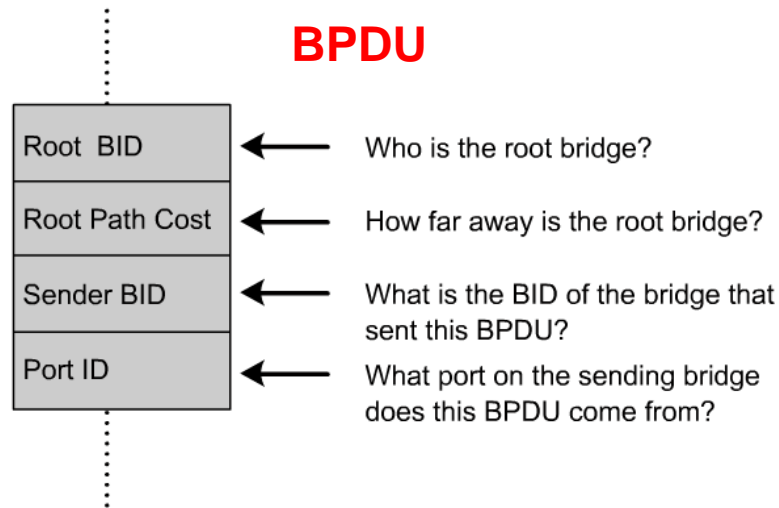
Link Speed	Cost(Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100



We will see how this works in a moment.

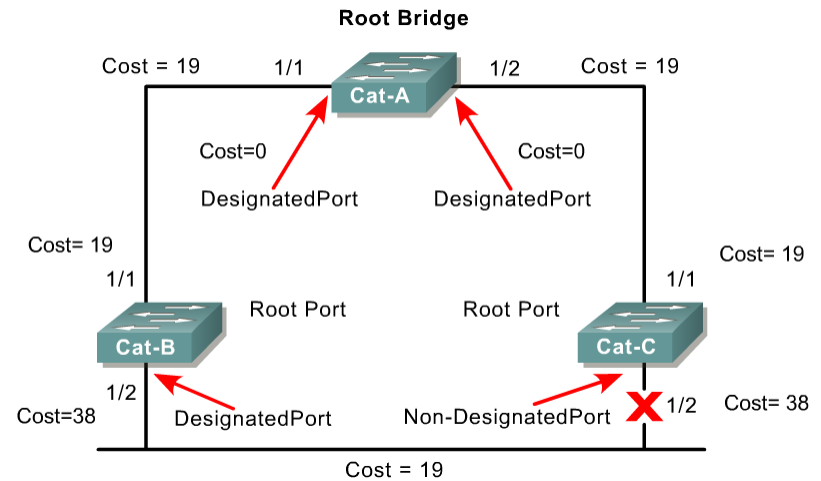
- **Shortest path is based on cumulative link costs.**
- Link costs are **based on the speed of the link.**
- The Spanning-Tree Protocol establishes a root node, called the root bridge.
- The Spanning-Tree Protocol constructs a topology that has one path for reaching every network node.
- The resulting tree originates from the **root bridge**.
- **Redundant links** that are not part of the shortest path tree are **blocked**.

Spanning-Tree Protocol (STP)



- **It is because certain paths are blocked that a loop free topology is possible.**
- **Data frames received on blocked links are dropped.**
- **The Spanning-Tree Protocol requires network devices to exchange messages to prevent bridging loops, called Bridge Protocol Data Unit (BPDU).**
- Links that will cause a loop are put into a blocking state.
- BPDUs continue to be received on blocked ports.
- This ensures that if an active path or device fails, a new spanning tree can be calculated.

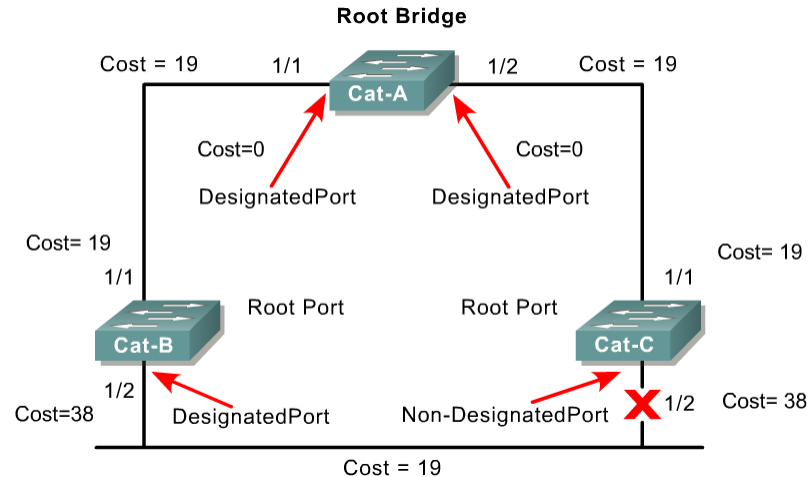
Spanning-Tree Protocol (STP)



BPDUs contain enough information so that all switches can do the following:

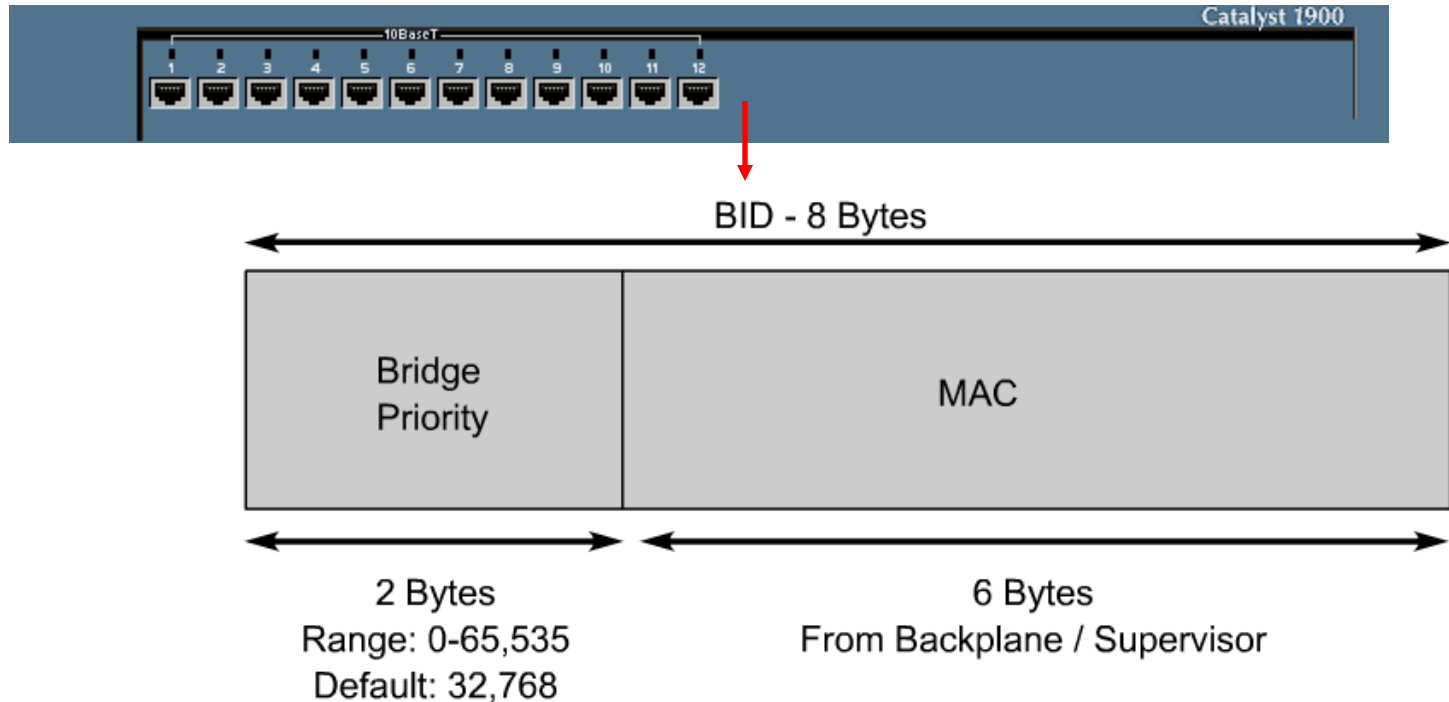
- Select a **single switch that will act as the root** of the spanning tree
- Calculate the **shortest path from itself to the root switch**
- **Designate one of the switches as the closest one to the root**, for each LAN segment. This bridge is called the “**designated switch**”.
 - The designated switch handles all communication from that LAN towards the root bridge.
- Choose one of its ports as its root port, for each non-root switch.
 - This is the interface that gives the best path to the root switch.
- Select ports that are part of the spanning tree, the designated ports. Non-designated ports are blocked.

Two Key Concepts: BID and Path Cost




- STP executes an algorithm called Spanning Tree Algorithm (STA).
- STA chooses a reference point, called a root bridge, and then determines the available paths to that reference point.
 - If more than two paths exists, STA picks the best path and blocks the rest
- STP calculations make extensive use of two key concepts in creating a loop-free topology:
 - **Bridge ID**
 - **Path Cost**

Bridge ID (BID)



- **Bridge ID (BID)** is used to identify each bridge/switch.
- The BID is **used in determining the center of the network**, in respect to **STP**, known as the root bridge.
- Consists of two components:
 - **A 2-byte Bridge Priority:** Cisco switch defaults to **32,768** or 0x8000.
 - **A 6-byte MAC address**

Bridge ID (BID)



- **Bridge Priority** is usually expressed in **decimal format** and the **MAC address** in the BID is usually expressed in **hexadecimal format**.
- BID is used to elect a root bridge (coming)
- **Lowest Bridge ID is the root.**
- If all devices have the same priority, the bridge with the lowest MAC address becomes the root bridge. (Yikes!)

Bridge ID (BID)

```
ALSwitch#show spanning-tree
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID      Priority      32768
             Address      0003.e334.6640
             Cost        19
             Port        23 (FastEthernet0/23)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Bridge ID  Priority      32769  (priority 32768 sys-id-ext 1)
             Address      000b.fc28.d400
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time   300
```

Interface	Port ID	Designated	Port ID
Name	Prio.Nbr	Cost Sts	Prio.Nbr

Fa0/23	128.23	19 FWD	128.25

```
ALSwitch#
```

Path Cost

Link Speed	Cost(Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

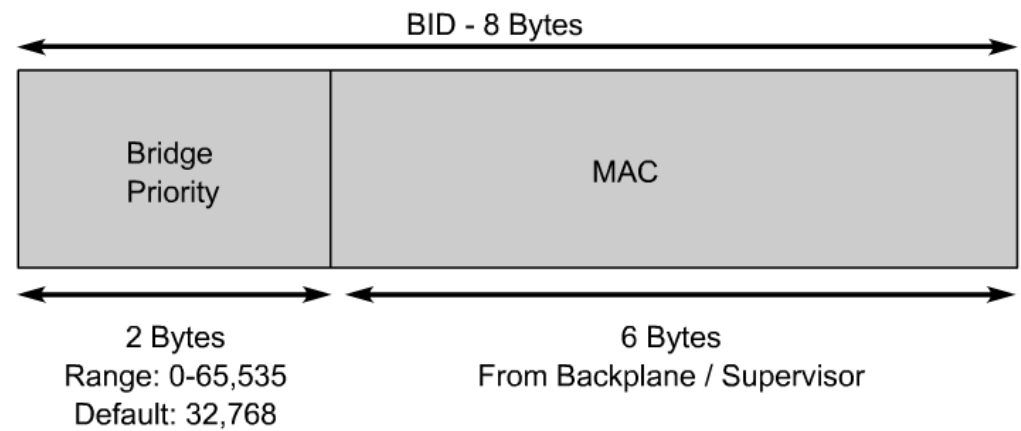
- Bridges use the concept of cost to evaluate how close they are to other bridges.
- This will be used in the STP development of a loop-free topology .
- **Originally, 802.1d** defined cost as **1000/bandwidth** of the link in Mbps.
 - Cost of 10Mbps link = 100 or $1000/10$
 - Cost of 100Mbps link = 10 or $1000/100$
 - Cost of 1Gbps link = 1 or $1000/1000$
- Running out of room for faster switches including 10 Gbps Ethernet.

Path Cost

Link Speed	Cost(Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

- IEEE modified the most to use a **non-linear scale** with the new values of:
 - 4 Mbps 250 (cost)
 - 10 Mbps 100 (cost)
 - 16 Mbps 62 (cost)
 - 45 Mbps 39 (cost)
 - 100 Mbps 19 (cost)
 - 155 Mbps 14 (cost)
 - 622 Mbps 6 (cost)
 - 1 Gbps 4 (cost)
 - 10 Gbps 2 (cost)

Path Cost



Link Speed	Cost(Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

- You can modify the path cost by modifying the cost of a port.
 - **Switch(config-if) # spanning-tree cost *value***
 - Exercise caution when you do this!
- BID and Path Cost are used to develop a loop-free topology .
- Coming very soon!
- But first the **Four-Step STP Decision Sequence**

Four-Step STP Decision Sequence

- When creating a loop-free topology, STP always uses the same four-step decision sequence:

Four-Step decision Sequence

Step 1 - Lowest BID

Step 2 - Lowest Path Cost to Root Bridge

Step 3 - Lowest Sender BID

Step 4 - Lowest Port ID

- Bridges use Configuration BPDUs during this four-step process.
 - There is another type of BPDU known as Topology Change Notification (TCN) BPDU (later)

Three Steps of Initial STP Convergence

- The STP algorithm uses three simple steps to converge on a loop-free topology.
- Switches go through three steps for their initial convergence:

STP Convergence

Step 1 Elect one Root Bridge

Step 2 Elect Root Ports

Step 3 Elect Designated Ports

- All STP decisions are based on a the following predetermined sequence:

Four-Step decision Sequence

Step 1 - Lowest BID

Step 2 - Lowest Path Cost to Root Bridge

Step 3 - Lowest Sender BID

Step 4 - Lowest Port ID

Three Steps of Initial STP Convergence

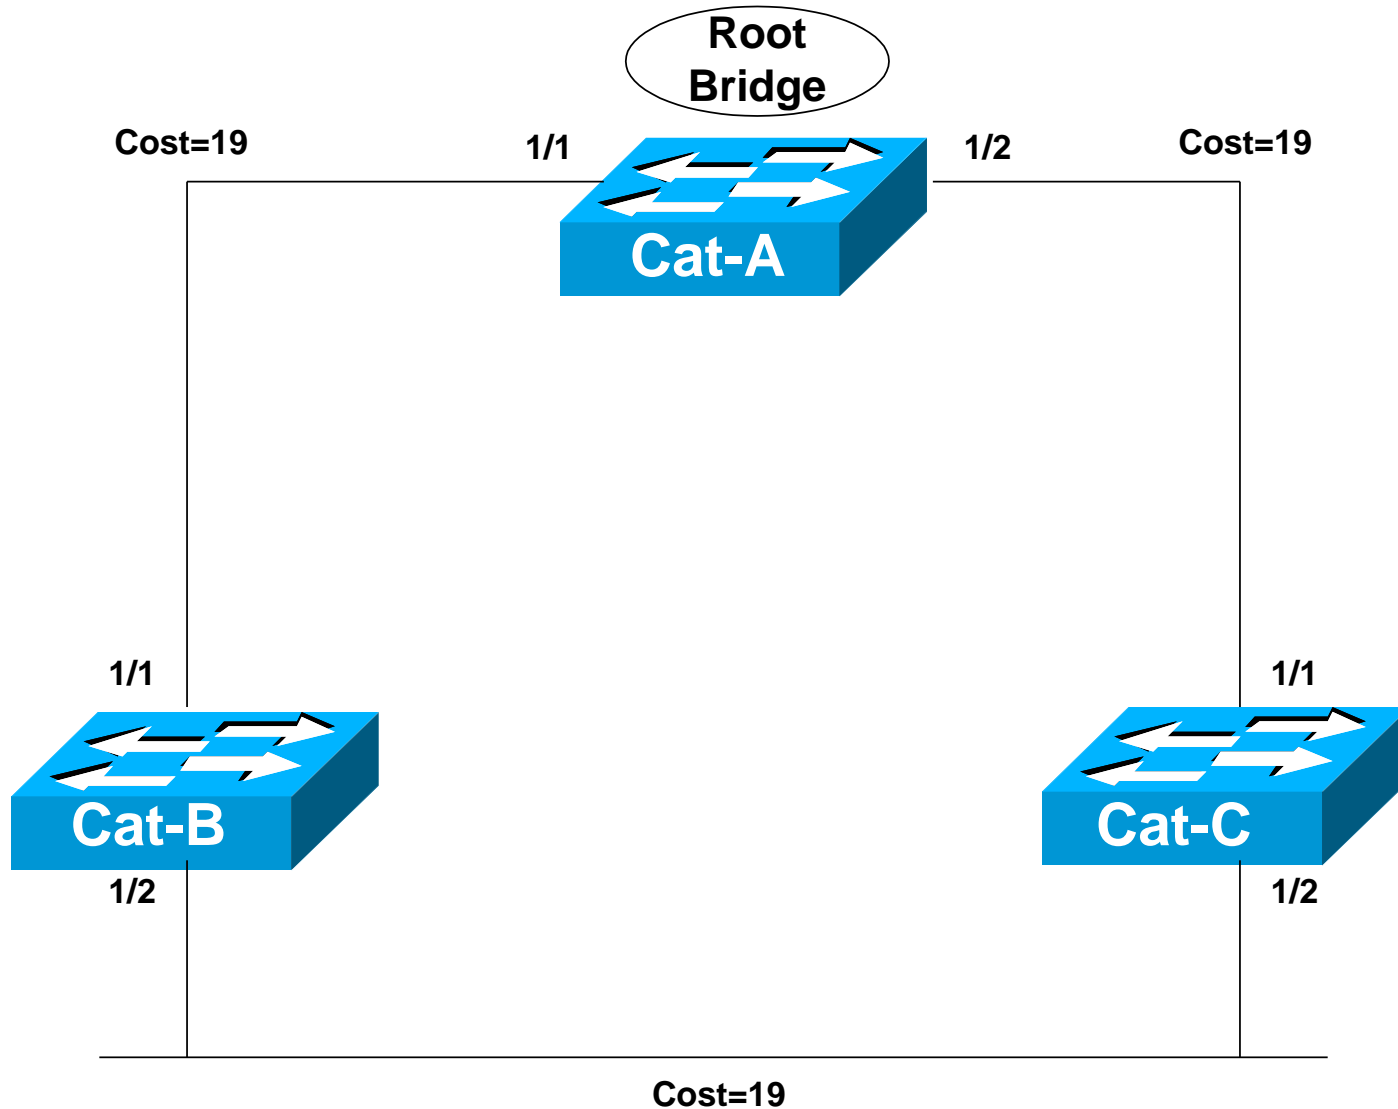
STP Convergence

Step 1 Elect one Root Bridge

Step 2 Elect Root Ports

Step 3 Elect Designated Ports

Step 1 Elect one Root Bridge



Step 1 Elect one Root Bridge

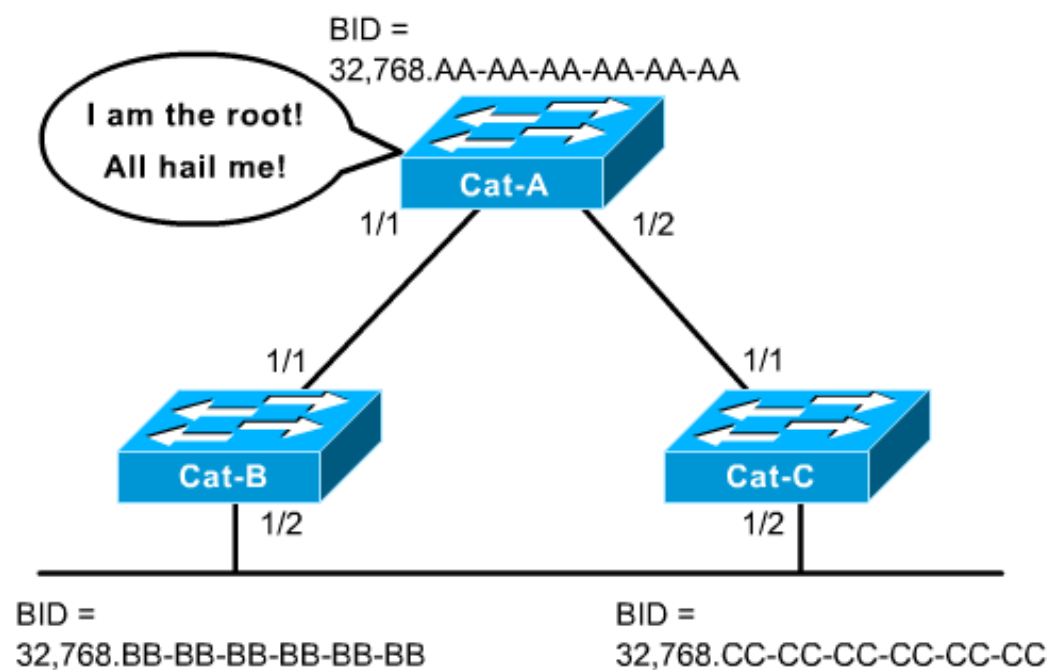
Four-Step decision Sequence

Step 1 - Lowest BID

Step 2 - Lowest Path Cost to Root Bridge

Step 3 - Lowest Sender BID

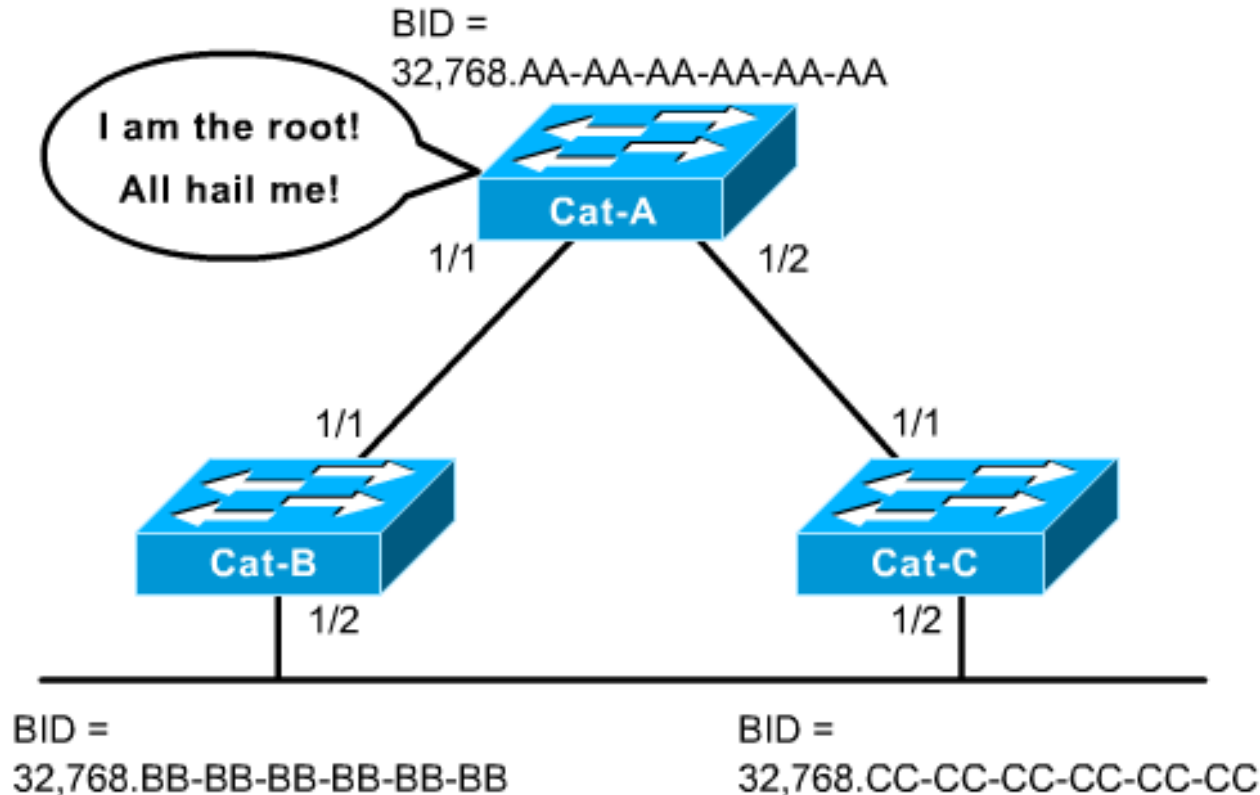
Step 4 - Lowest Port ID



- When the network first starts, all bridges are announcing a chaotic mix of BPDUs.
- All bridges immediately begin applying the four-step sequence decision process.
- Switches need to elect a single Root Bridge.
- Switch with the **lowest BID** wins!
- Note: Many texts refer to the term “highest priority” which is the “lowest” BID value.
- This is known as the “Root War.”

Step 1 Elect one Root Bridge

Cat-A has the lowest Bridge MAC Address, so it wins the Root War!



All 3 switches have the same default Bridge Priority value of 32,768

Step 1 Elect one Root Bridge

BPDUs

802.3 Header

Destination: 01:80:C2:00:00:00 *Mcast 802.1d Bridge group*
Source: 00:D0:C0:F5:18:D1
LLC Length: 38

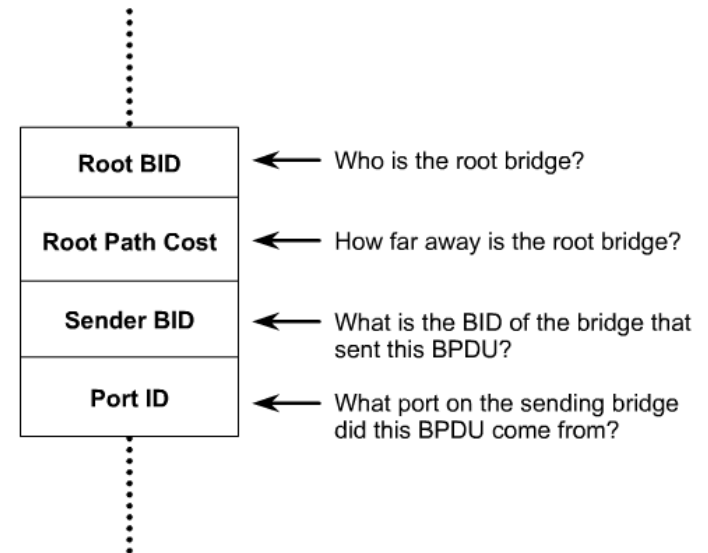
802.2 Logical Link Control (LLC) Header

Dest. SAP: 0x42 *802.1 Bridge Spanning Tree*
Source SAP: 0x42 *802.1 Bridge Spanning Tree*
Command: 0x03 *Unnumbered Information*

802.1 - Bridge Spanning Tree

Protocol Identifier: 0
Protocol Version ID: 0
Message Type: 0 *Configuration Message*
Flags: %00000000
Root Priority/ID: 0x8000/ 00:D0:C0:F5:18:C0
Cost Of Path To Root: 0x00000000 *(0)*
Bridge Priority/ID: 0x8000/ 00:D0:C0:F5:18:C0
Port Priority/ID: 0x80/ 0x1D
Message Age: 0/256 seconds *(exactly 0 seconds)*
Maximum Age: 5120/256 seconds *(exactly 20 seconds)*
Hello Time: 512/256 seconds *(exactly 2 seconds)*
Forward Delay: 3840/256 seconds *(exactly 15 seconds)*

Its all done with BPDUs!



Configuration BPDUs are sent every 2 seconds by default.

Step 1 Elect one Root Bridge

- In a real network, you do not want the placement of the root bridge to rely on the random placement of the switch with the lowest MAC address.
- A misplaced root bridge can have significant effects on your network including less than optimum paths within the network.
- It is better to configure a switch to be the root bridge:

```
Switch(config)# spanning-tree [vlan vlan-list] priority priority *
```

- Priority
 - Default = 32,768
 - Range 0-65,535
 - Lowest wins

Step 1 Elect one Root Bridge

```
2950#show spanning-tree
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID      Priority    32768  
            
Address      0003.e334.6640  
Cost          19  
Port          23 (FastEthernet0/23)  
Hello Time    2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Bridge ID    Priority    32769  (priority 32768 sys-id-ext 1)  
Address      000b.fc28.d400  
Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec  
Aging Time   300
```

Interface	Port ID	Designated	Port ID
Name	Prio.Nbr	Cost Sts	Prio.Nbr

Fa0/23	128.23	19 FWD	128.25

```
ALSwitch#
```

Step 1 Elect one Root Bridge

```
2900#show spanning-tree
```

```
Spanning tree 1 is executing the IEEE compatible Spanning Tree protocol
```

```
Bridge Identifier has priority 32768, address 0003.e334.6640
```

```
Configured hello time 2, max age 20, forward delay 15
```

```
We are the root of the spanning tree
```

```
Topology change flag not set, detected flag not set, changes 1
```

```
Times: hold 1, topology change 35, notification 2
```

```
hello 2, max age 20, forward delay 15
```

```
Timers: hello 0, topology change 0, notification 0
```

```
Interface Fa0/1 (port 13) in Spanning tree 1 is down
```

```
Port path cost 19, Port priority 128
```

```
Designated root has priority 32768, address 0003.e334.6640
```

```
Designated bridge has priority 32768, address 0003.e334.6640
```

```
Designated port is 13, path cost 0
```

```
Timers: message age 0, forward delay 0, hold 0
```

```
BPDUs: sent 1, received 0
```

Three Steps of Initial STP Convergence

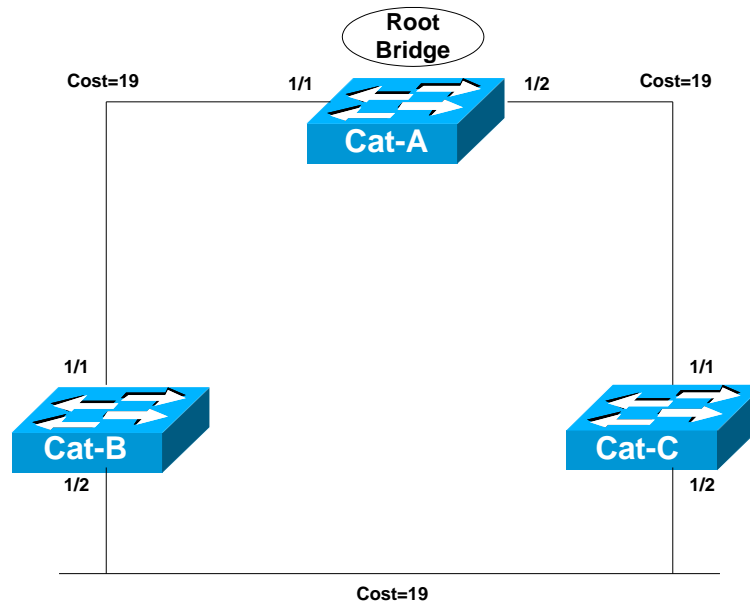
STP Convergence

Step 1 Elect one Root Bridge

Step 2 Elect Root Ports

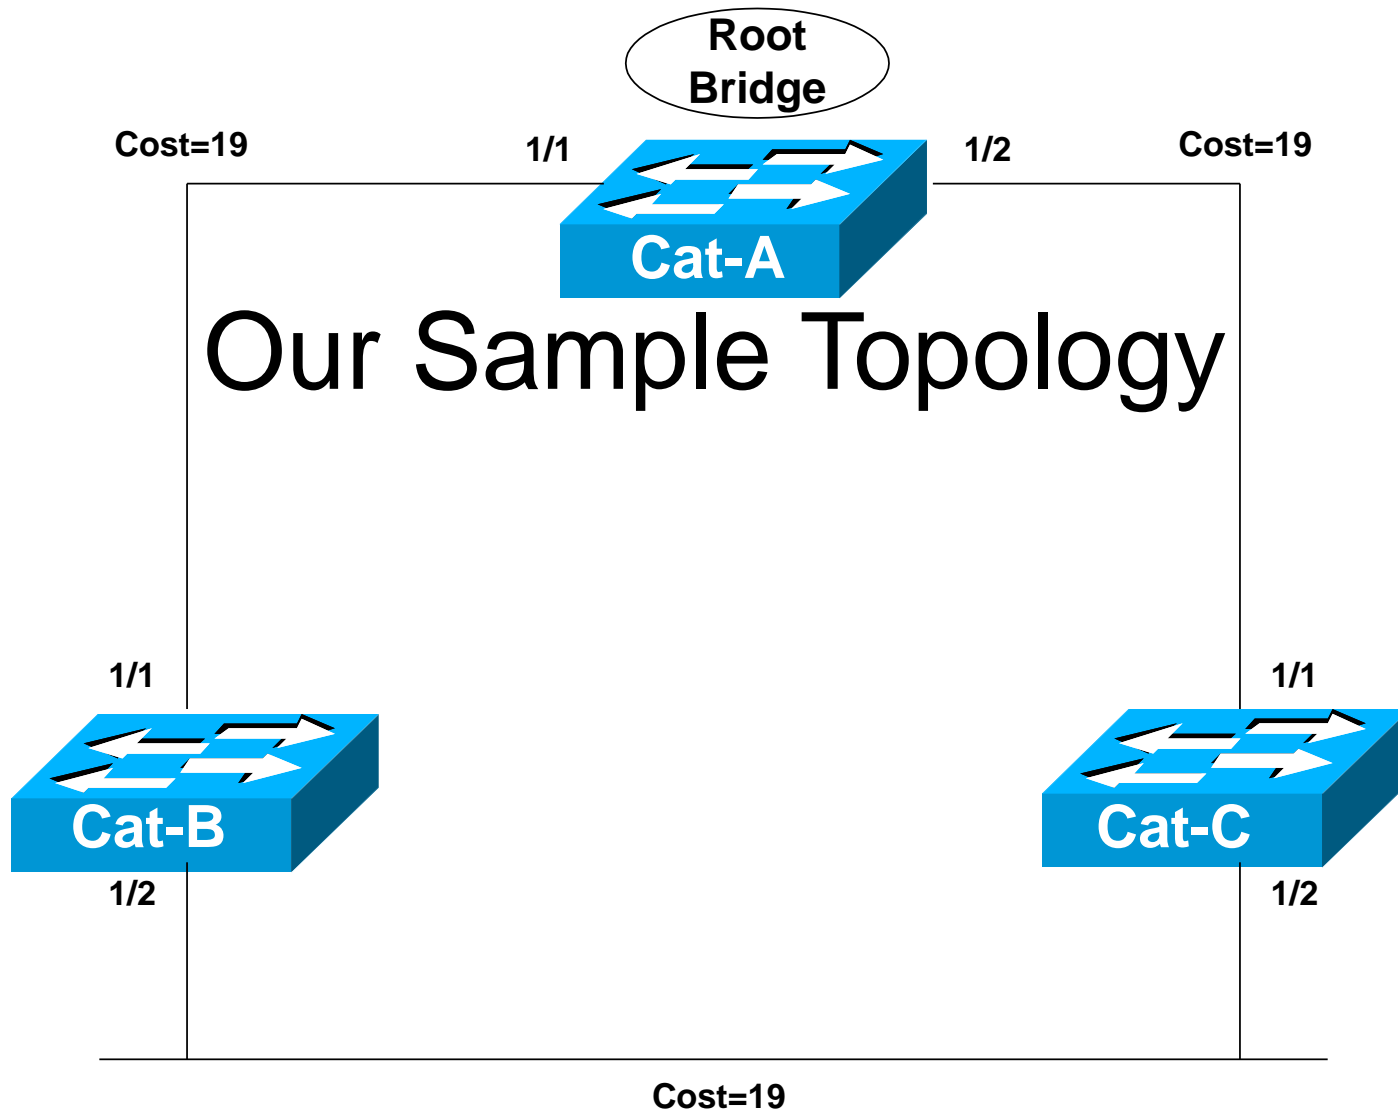
Step 3 Elect Designated Ports

Step 2 Elect Root Ports

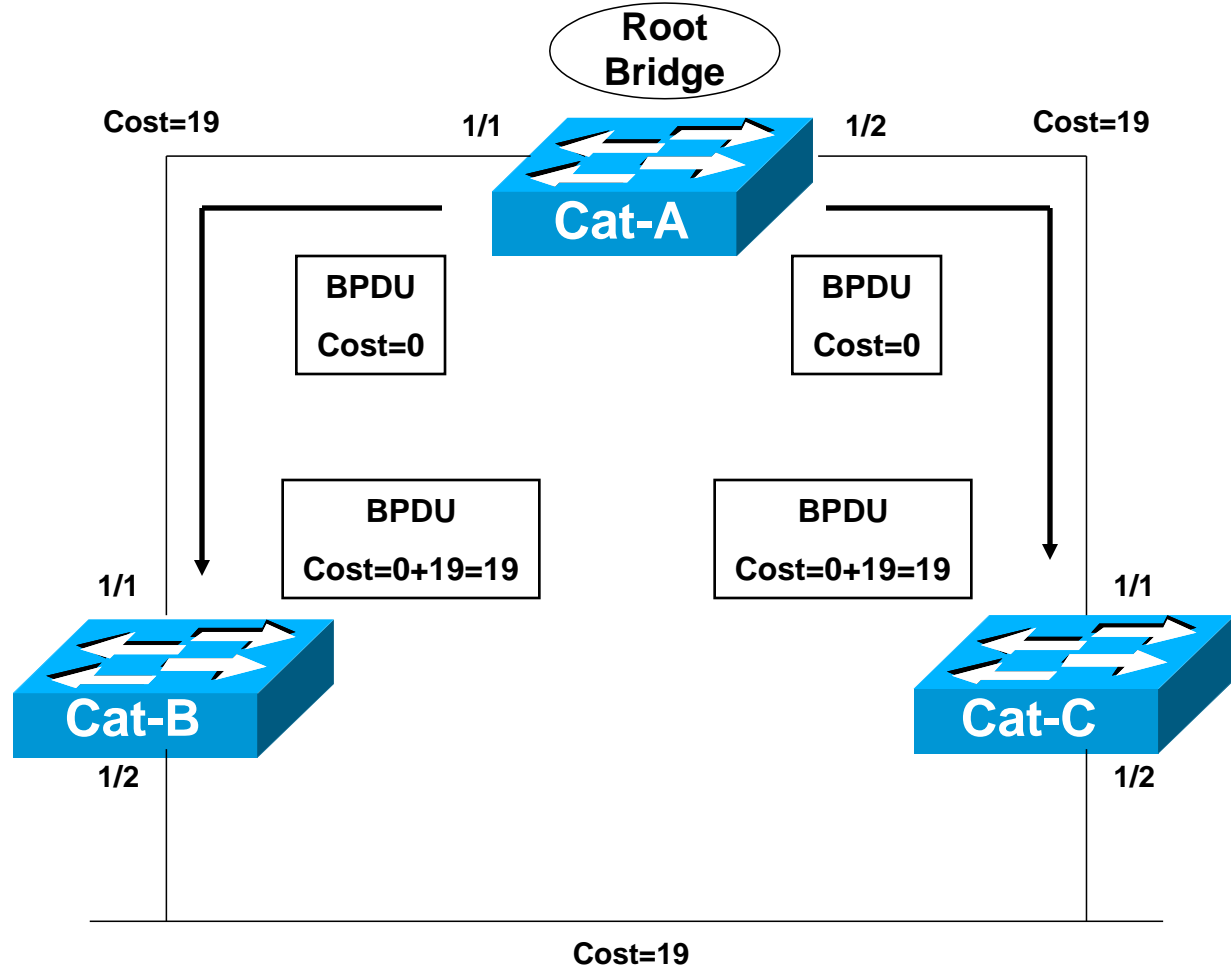


- Now that the Root War has been won, switches move on to selecting **Root Ports**.
- A bridge's **Root Port** is the *port closest to the Root Bridge*.
- Bridges use the **cost** to determine closeness.
- **Every non-Root Bridge will select one Root Port!**
- Specifically, bridges track the **Root Path Cost**, the cumulative cost of all links to the Root Bridge.

Step 2 Elect Root Ports



Step 2 Elect Root Ports



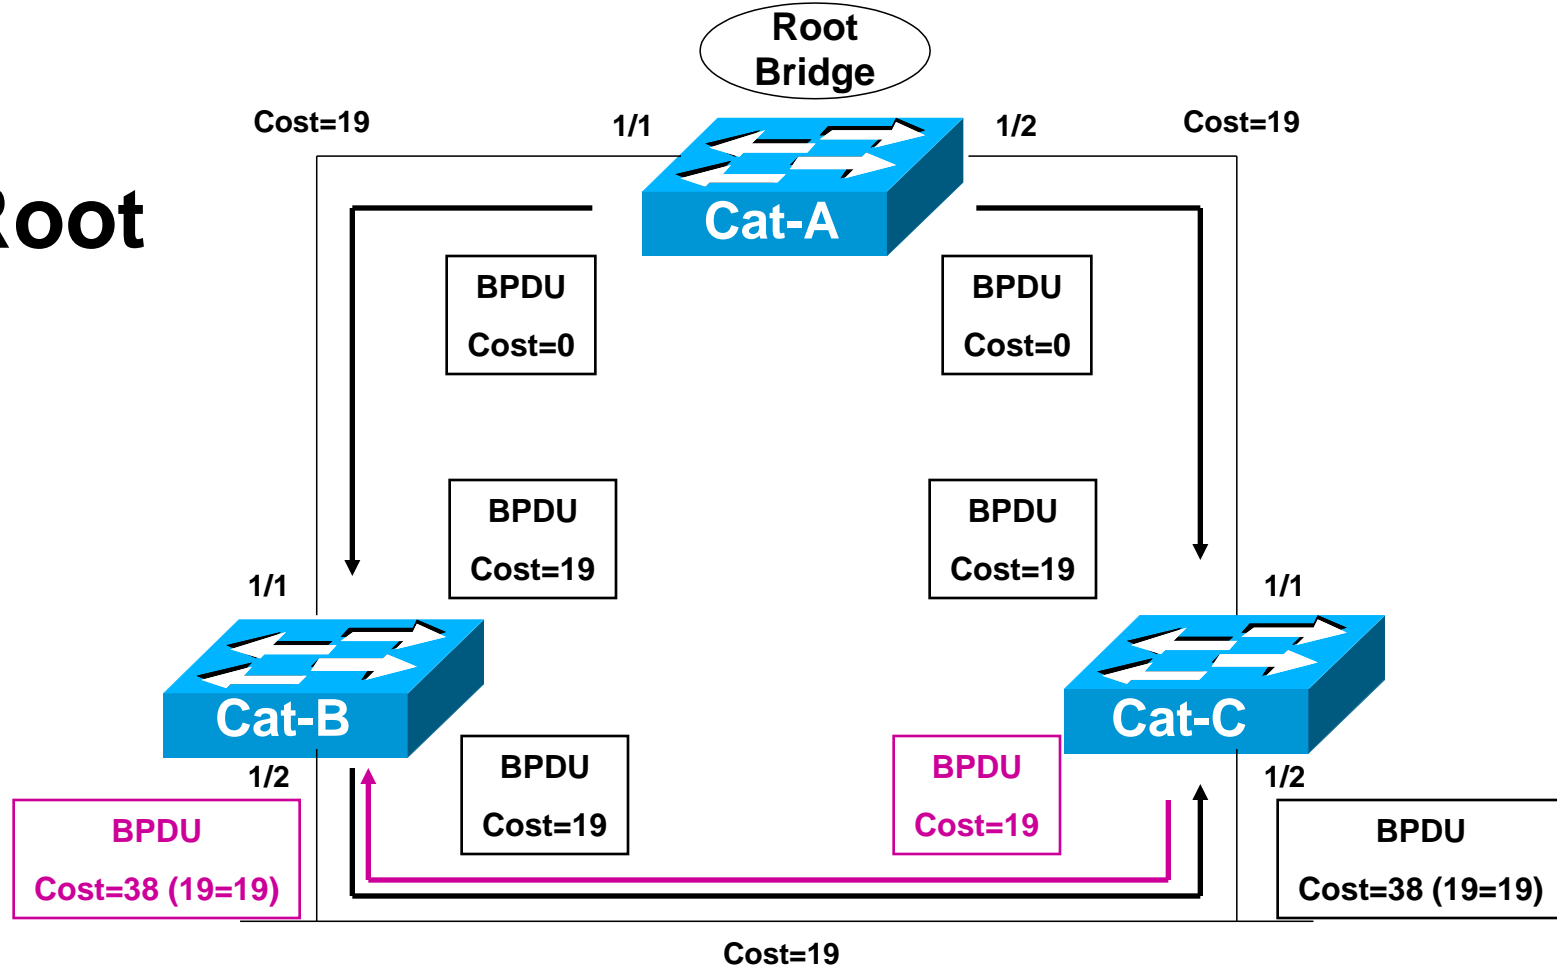
Step 1

- Cat-A sends out BPDUs, containing a Root Path Cost of 0.
- Cat-B receives these BPDUs and adds the Path Cost of Port 1/1 to the Root Path Cost contained in the BPDU.

Step 2

- Cat-B adds Root Path Cost 0 PLUS its Port 1/1 cost of 19 = 19

Step 2 Elect Root Ports



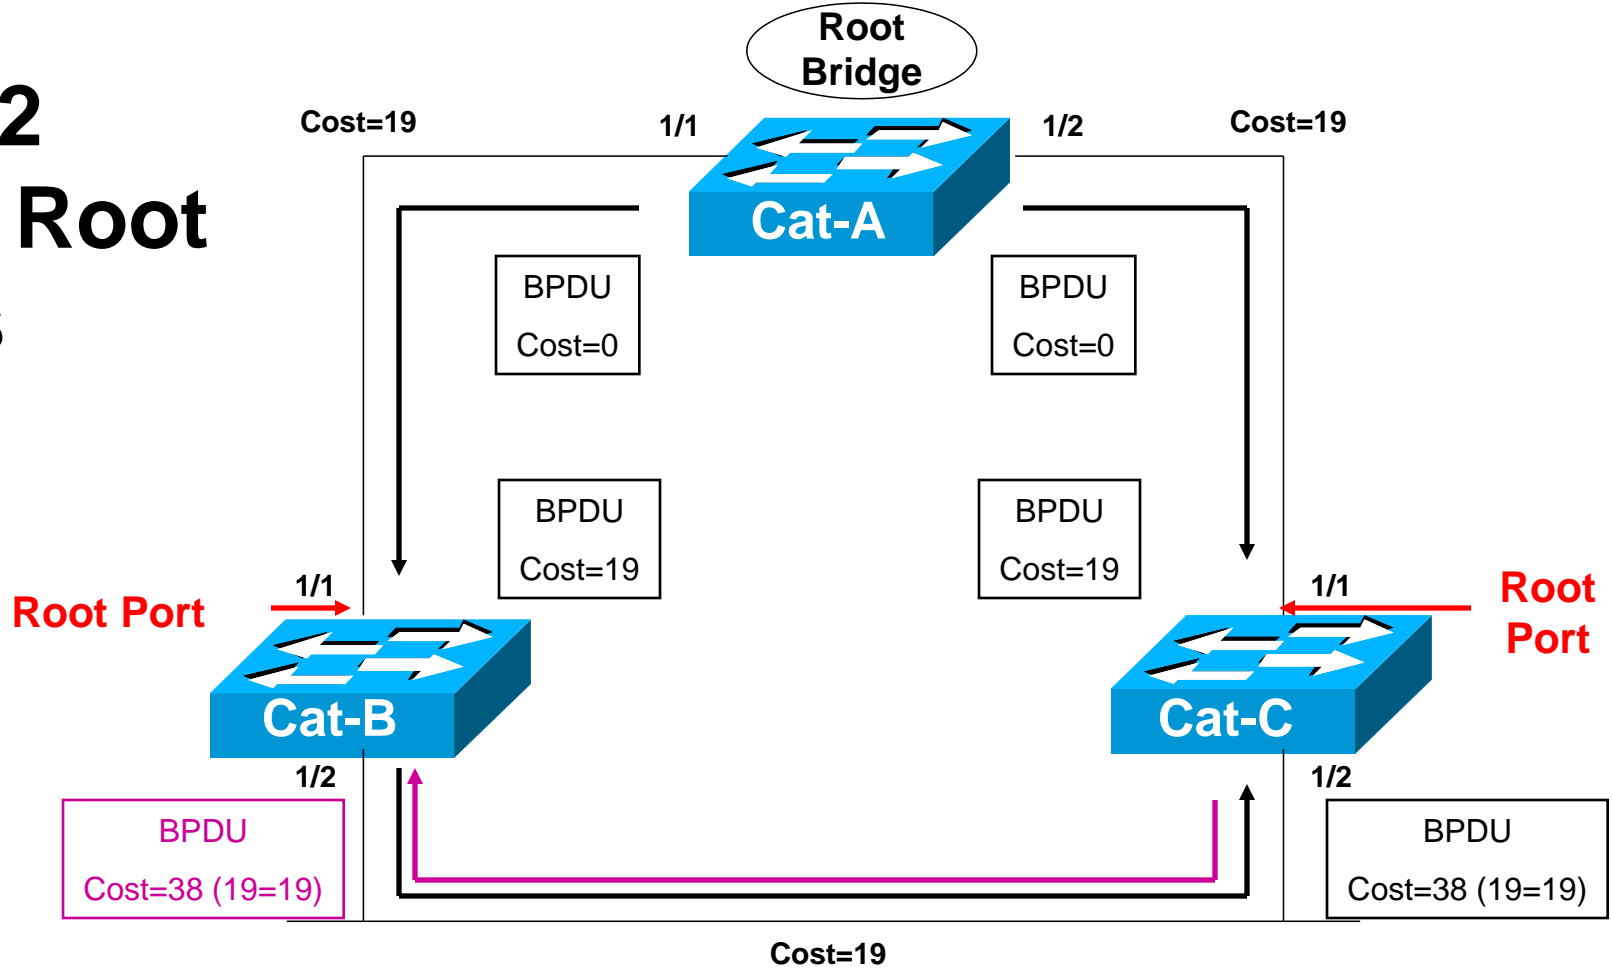
Step 3

- Cat-B uses this value of 19 internally and sends BPDUs with a Root Path Cost of 19 out Port 1/2.

Step 4

- Cat-C receives the BPDU from Cat-B, and increased the Root Path Cost to 38 (19+19). (Same with Cat-C sending to Cat-B.)

Step 2 Elect Root Ports



Step 5

- Cat-B calculates that it can reach the Root Bridge at a cost of 19 via Port 1/1 as opposed to a cost of 38 via Port 1/2.
- Port 1/1 becomes the Root Port for Cat-B, the port closest to the Root Bridge.
- Cat-C goes through a similar calculation. Note: Both Cat-B:1/2 and Cat-C:1/2 save the best BPDU of 19 (its own).

Step 2 Elect Root Ports

```
2950#show spanning-tree
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID      Priority      32768  
Address      0003.e334.6640  
Cost         19  
Port         23 (FastEthernet0/23)  
Hello Time   2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID    Priority      32769 (priority 32768 sys-id-ext 1)  
Address      000b.fc28.d400  
Hello Time   2 sec Max Age 20 sec Forward Delay 15 sec  
Aging Time   300
```

Interface Name	Port ID Prio.Nbr	Cost	Sts	Designated Cost Bridge ID	Port ID Prio.Nbr
-----	-----	-----	-----	-----	-----
Fa0/23	128.23	19	FWD	0 32768 0003.e334.6640	128.25

```
ALSwitch#
```

Three Steps of Initial STP Convergence

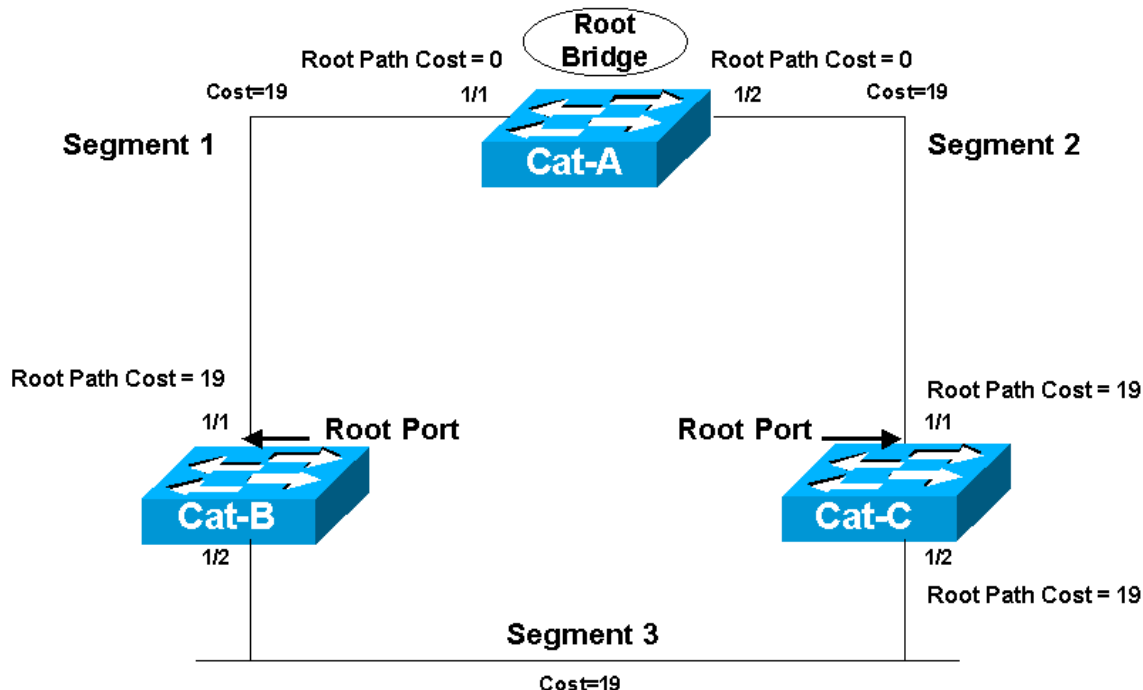
STP Convergence

Step 1 Elect one Root Bridge

Step 2 Elect Root Ports

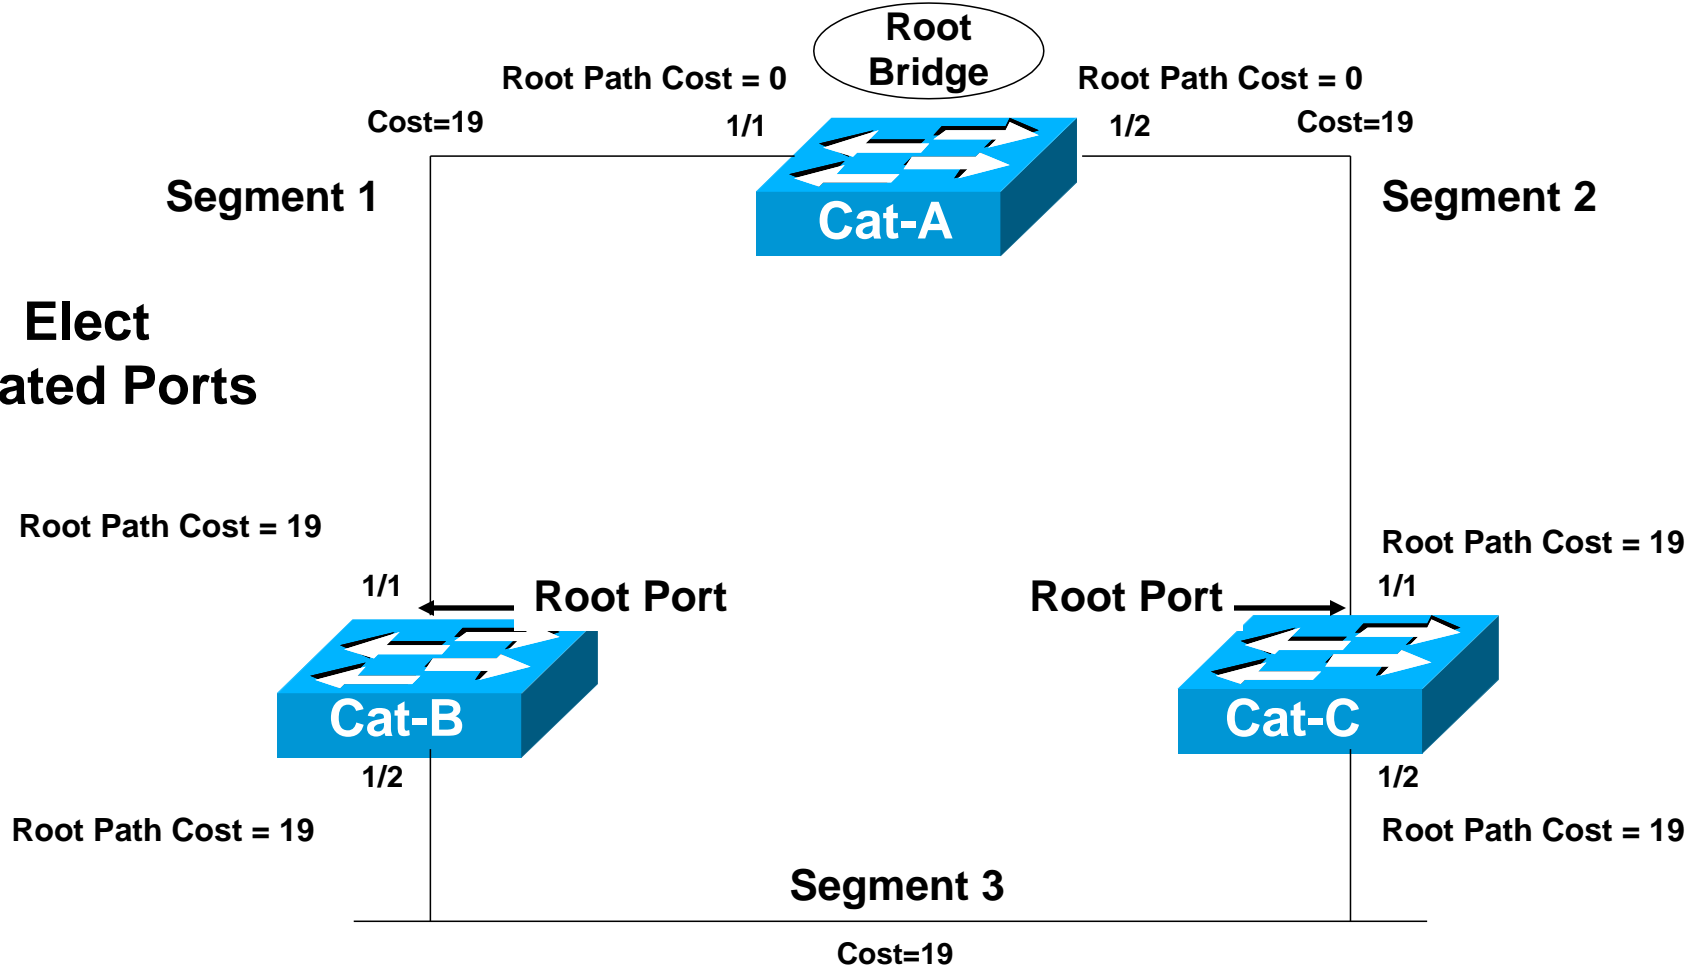
Step 3 Elect Designated Ports

Step 3 Elect Designated Ports

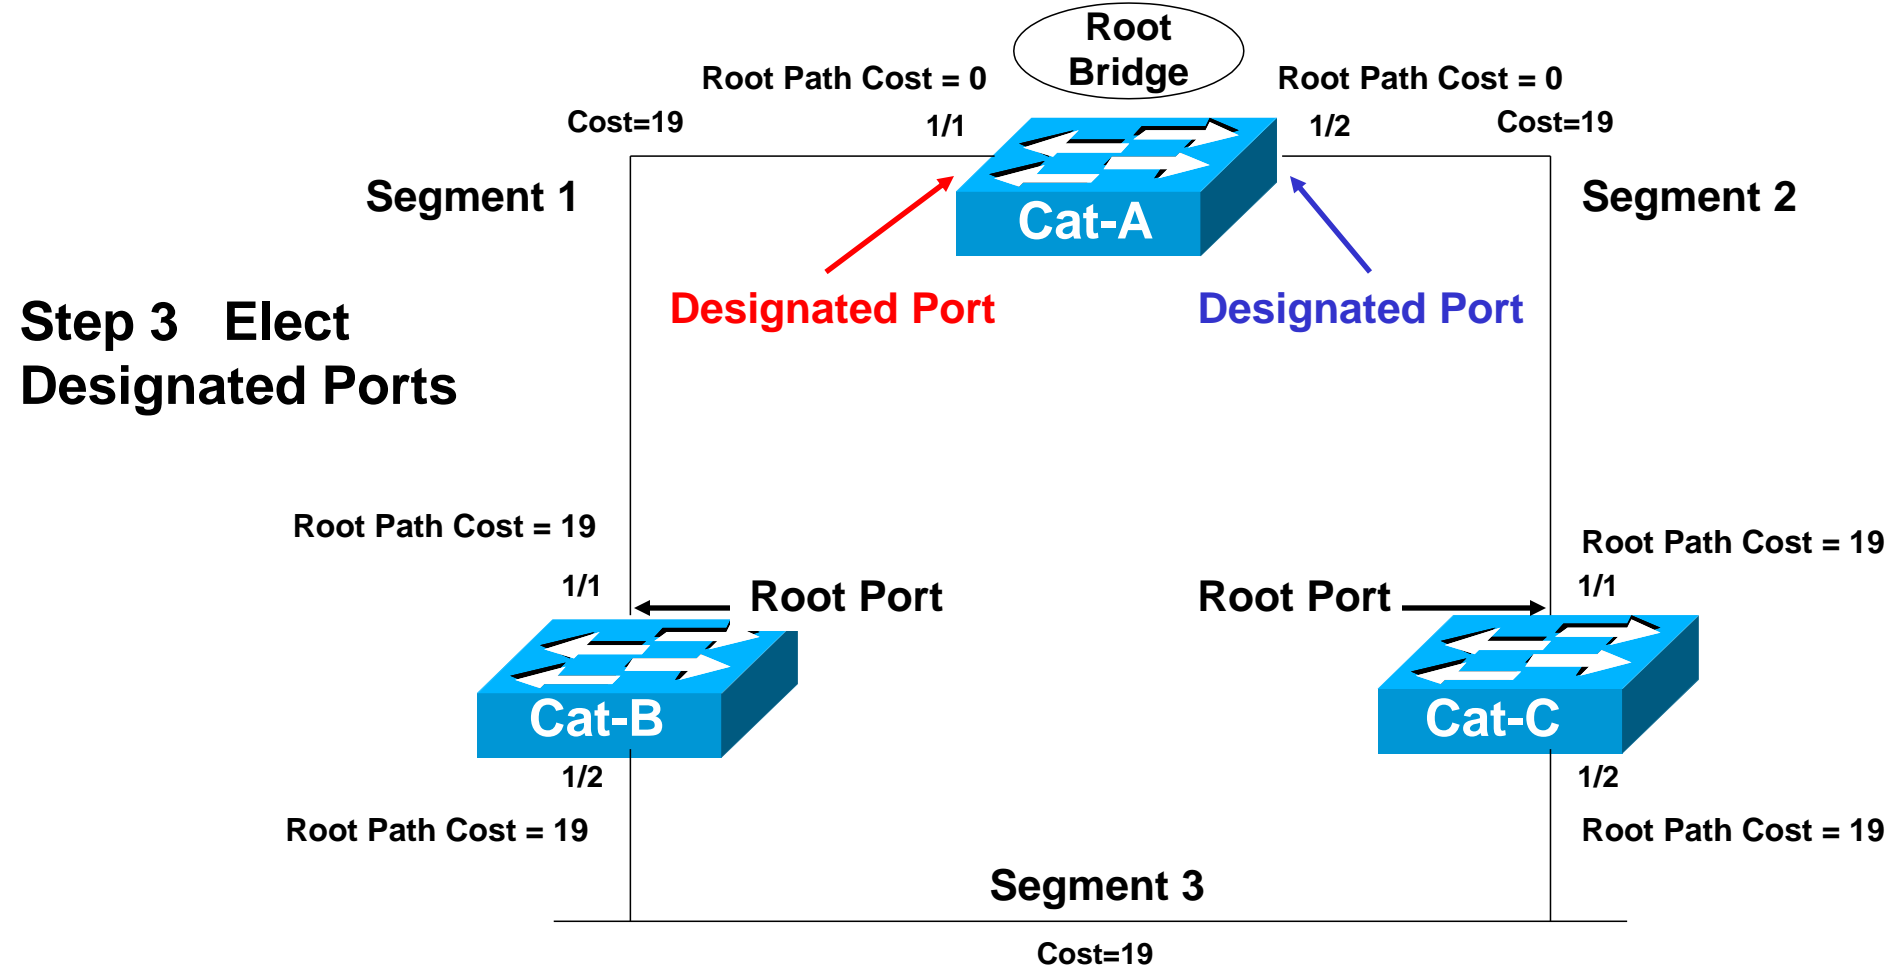


- The loop prevention part of STP becomes evident during this step, electing designated ports.
- A **Designated Port** functions as *the single bridge port that both sends and receives traffic to and from that segment and the Root Bridge*.
- **Each segment in a bridged network has one Designated Port, chosen based on cumulative Root Path Cost to the Root Bridge.**
- The switch containing the Designated Port is referred to as the **Designated Bridge** for that segment.
- To locate Designated Ports, let's take a look at each segment.
- **Root Path Cost**, the cumulative cost of all links to the Root Bridge.

Step 3 Elect Designated Ports



- **Segment 1:** Cat-A:1/1 has a Root Path Cost = 0 (after all it has the Root Bridge) and Cat-B:1/1 has a Root Path Cost = 19.
- **Segment 2:** Cat-A:1/2 has a Root Path Cost = 0 (after all it has the Root Bridge) and Cat-C:1/1 has a Root Path Cost = 19.
- **Segment 3:** **Cat-B:1/2 has a Root Path Cost = 19 and Cat-C:1/2 has a Root Path Cost = 19. It's a tie!**

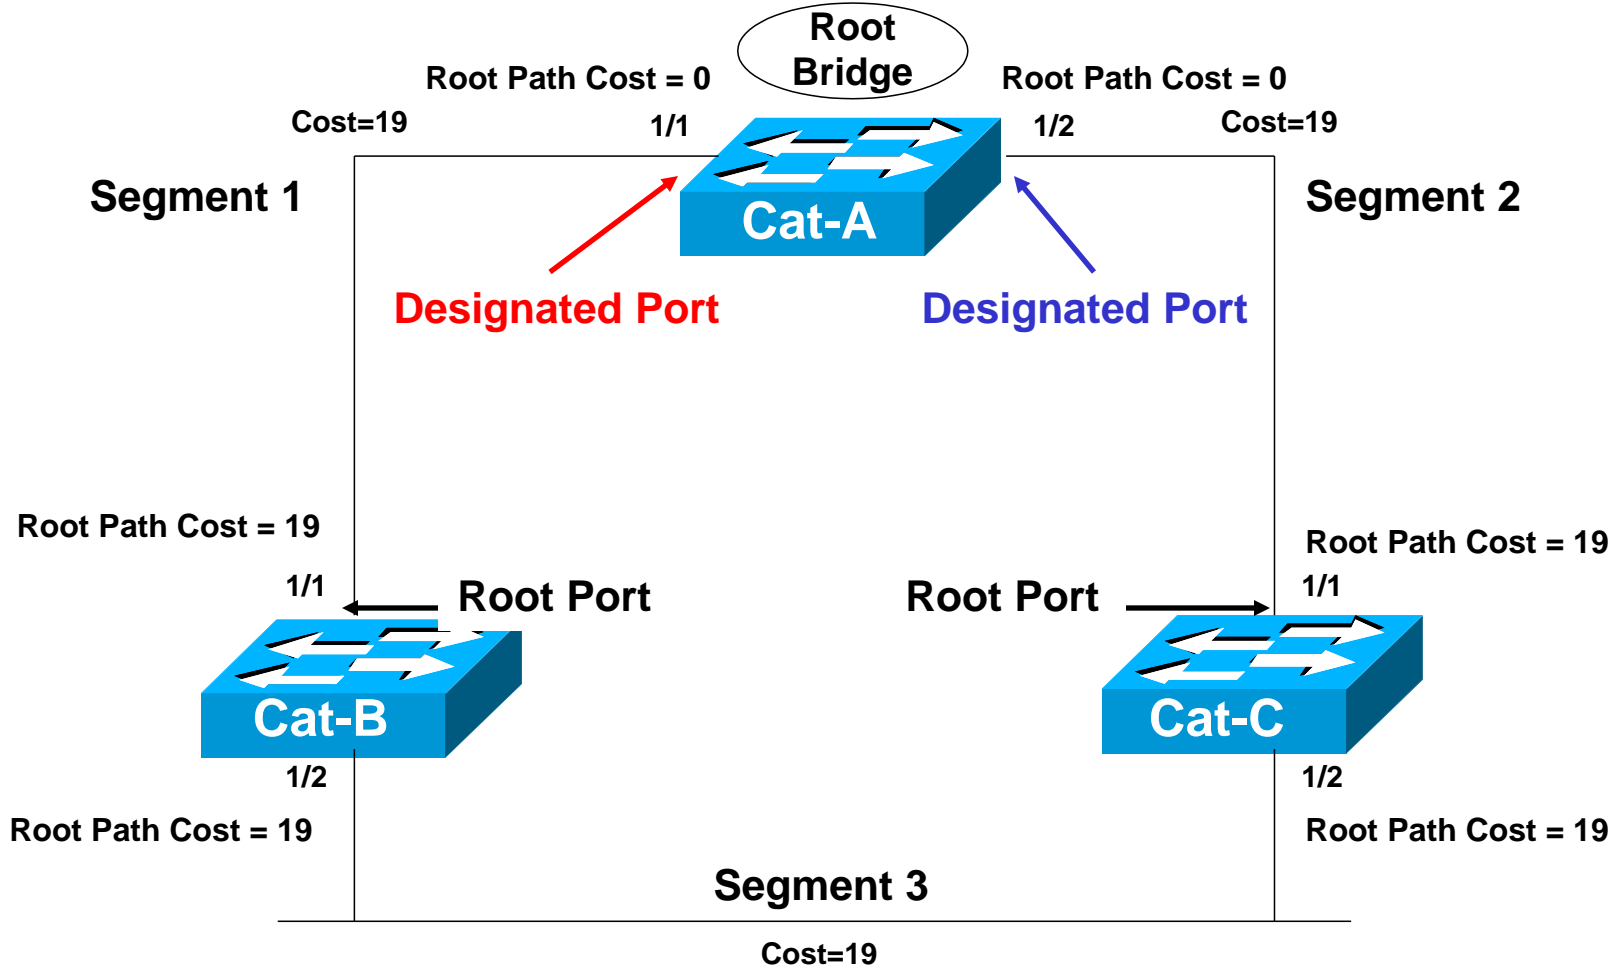


Segment 1

- Because Cat-A:1/1 has the lower Root Path Cost it becomes the **Designate Port for Segment 1**.

Segment 2

- Because Cat-A:1/2 has the lower Root Path Cost it becomes the **Designate Port for Segment 2**.



Segment 3

- **Both Cat-B and Cat-C have a Root Path Cost of 19, a tie!**
- When faced with a tie (or any other determination) STP always uses the four-step decision process:

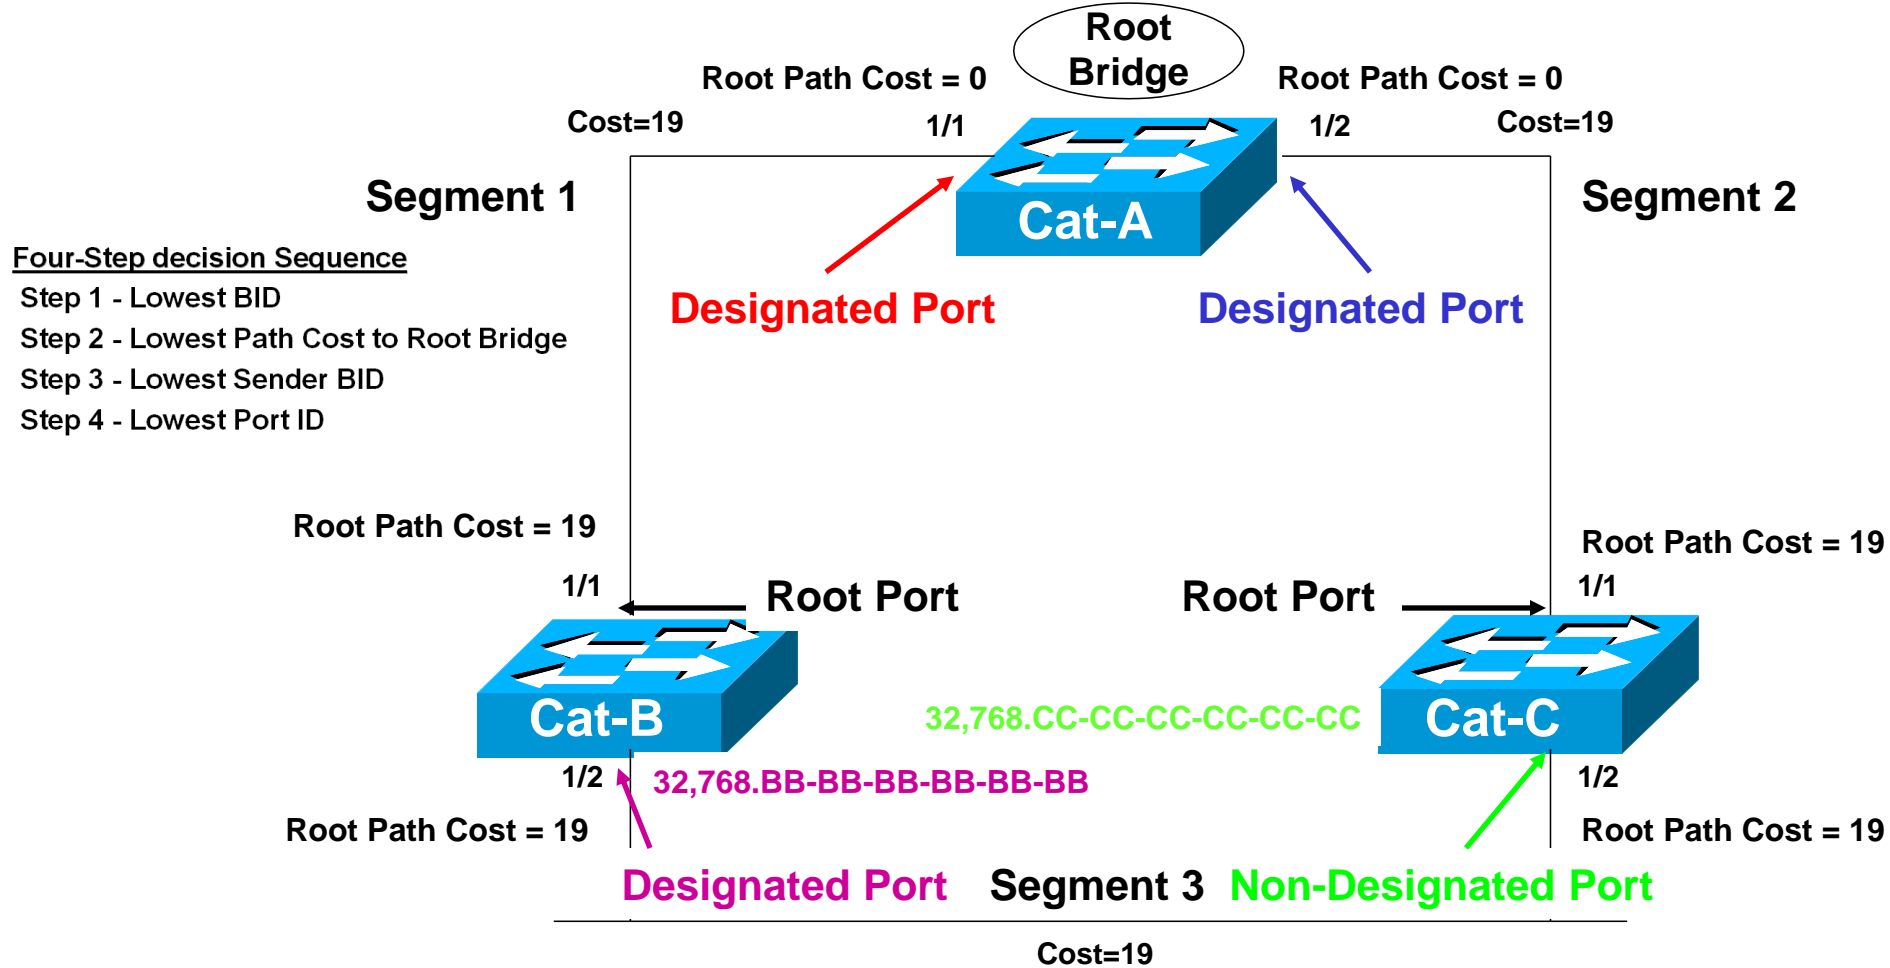
Four-Step decision Sequence

Step 1 - Lowest BID

Step 2 - Lowest Path Cost to Root Bridge

Step 3 - Lowest Sender BID

Step 4 - Lowest Port ID



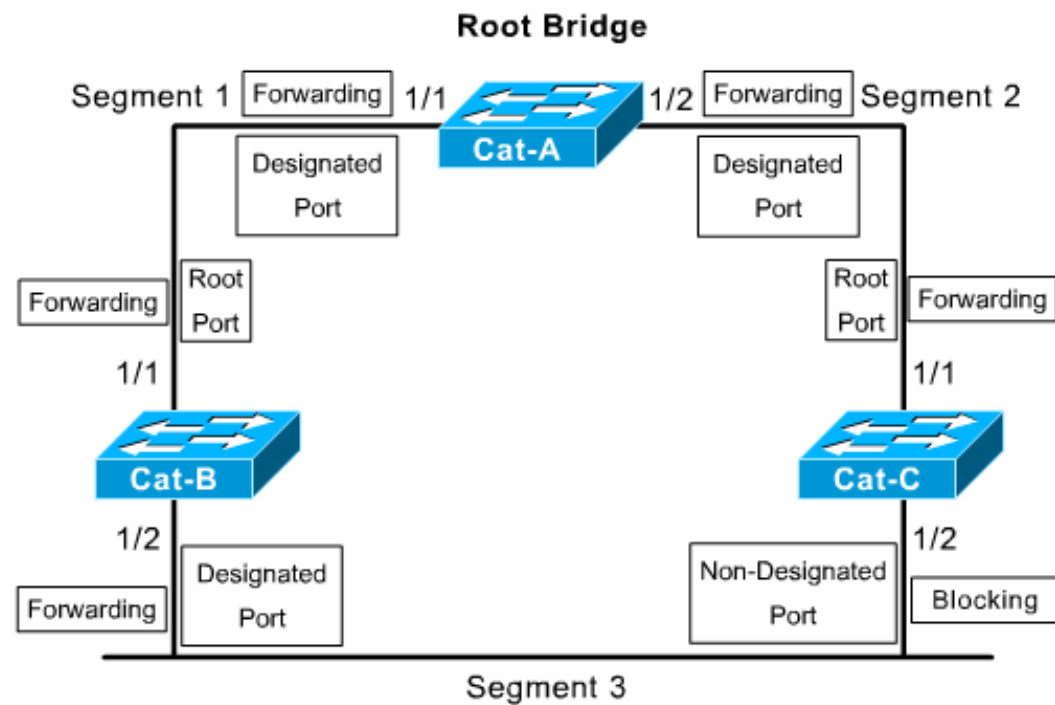
Segment 3 (continued)

- 1) All three switches agree that Cat-A is the Root Bridge, so this is a tie.
- 2) Root Path Cost for both is 19, also a tie.
- 3) The sender's BID is lower on Cat-B, than Cat-C, so Cat-B:1/2 becomes the **Designated Port for Segment 3**.
- Cat-C:1/2 therefore becomes the **non-Designated Port for Segment 3**.

State	Purpose
Forwarding	Sending / receiving user data
Learning	Building bridging table
Listening	Building "active" topology
Blocking	Receives BPDUs only
Disabled	Administratively down

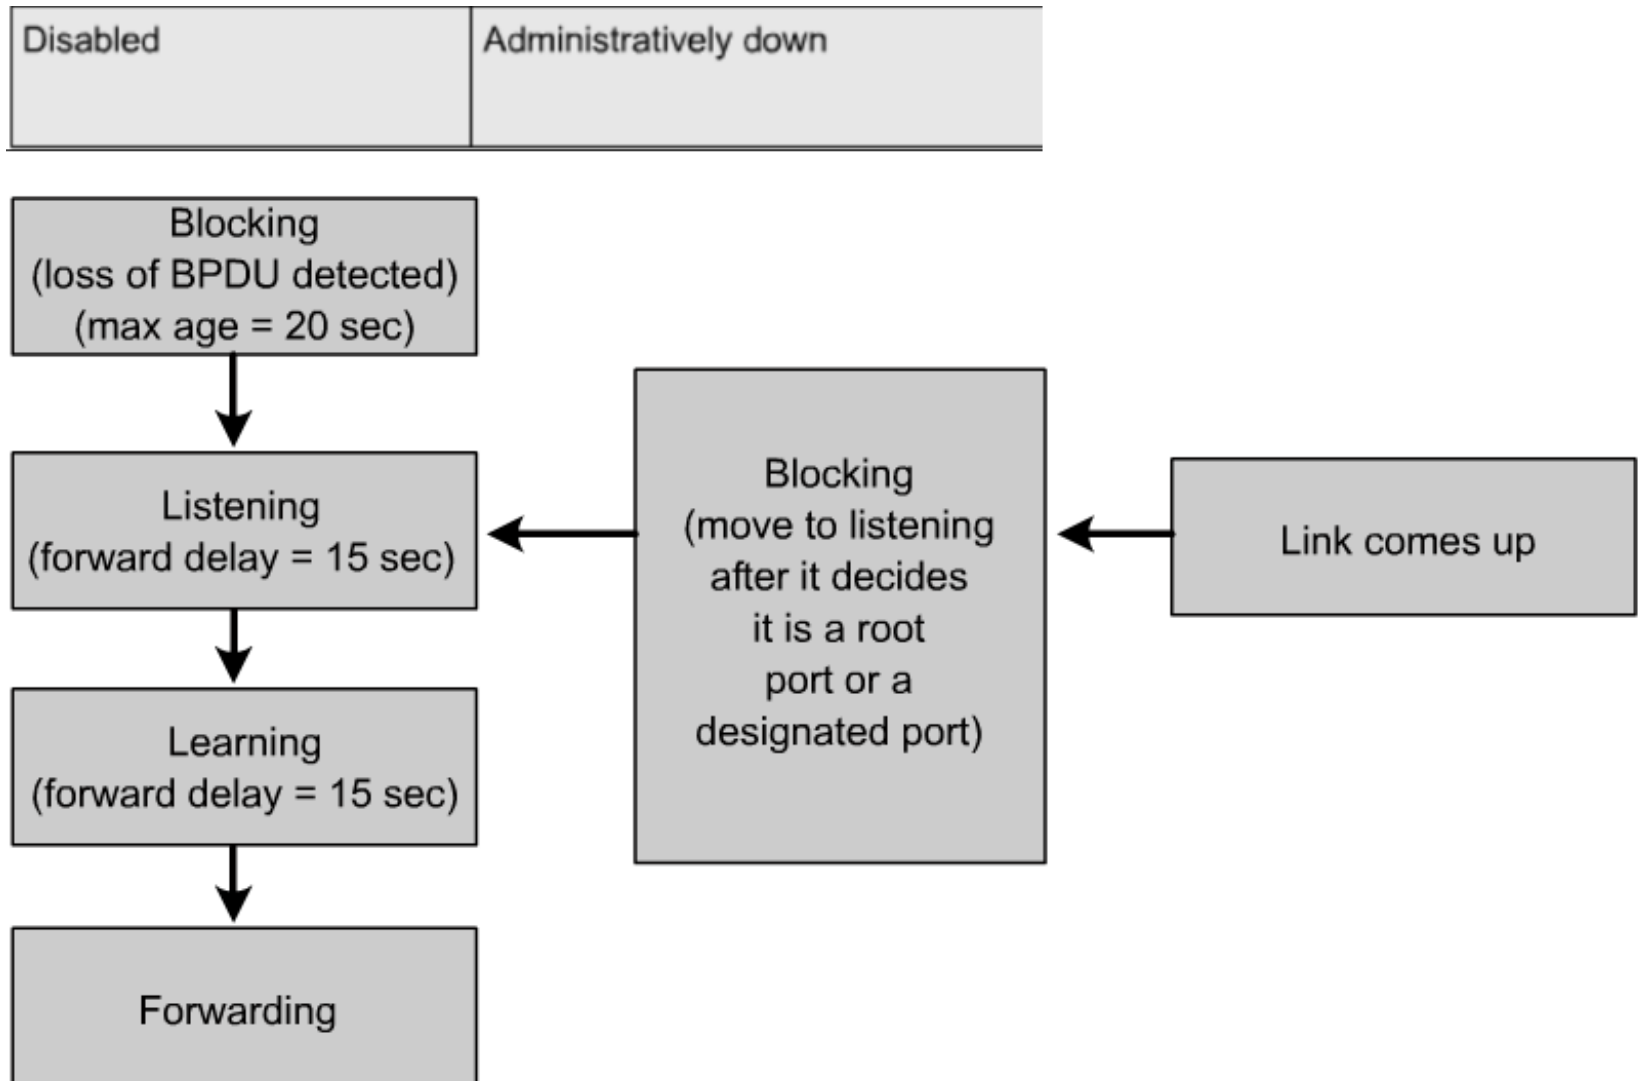
Designated Ports & Root Ports

Non-Designated Ports

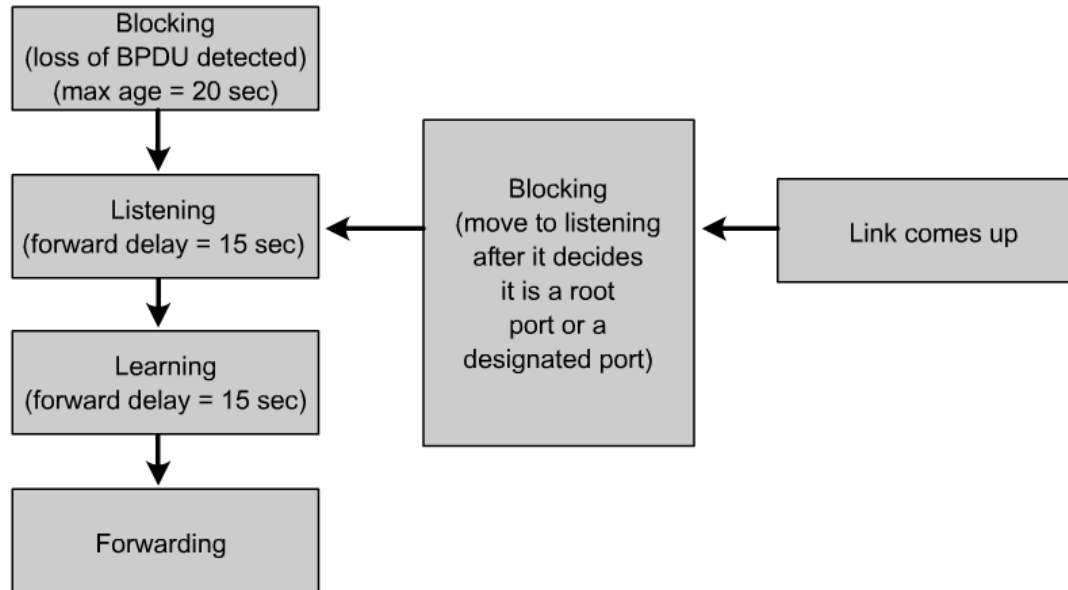


We will first only look at switch ports that have connections to other switches.

Stages of spanning-tree port states

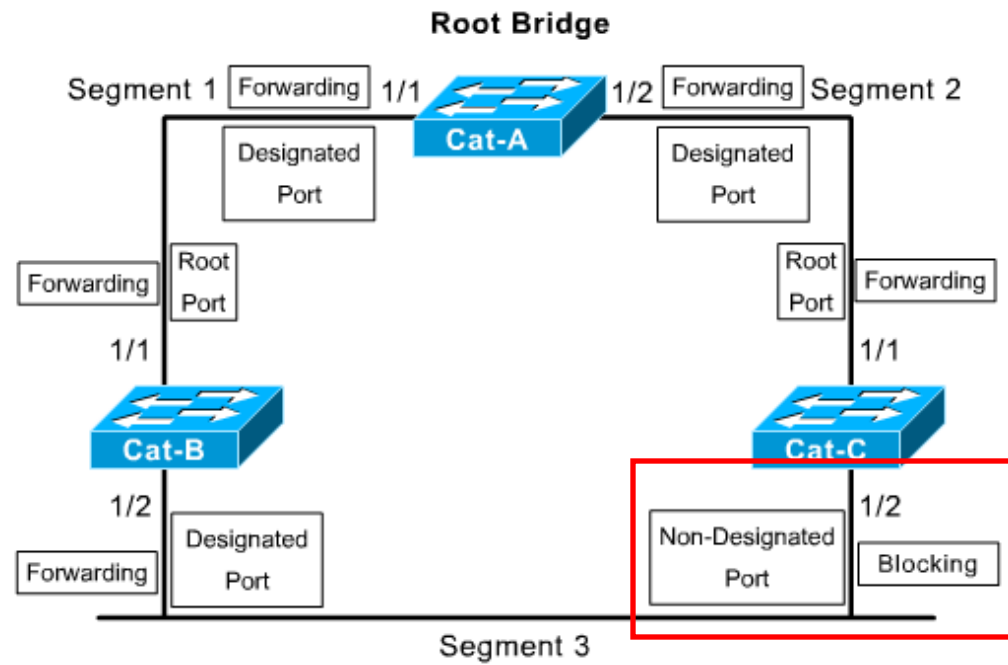
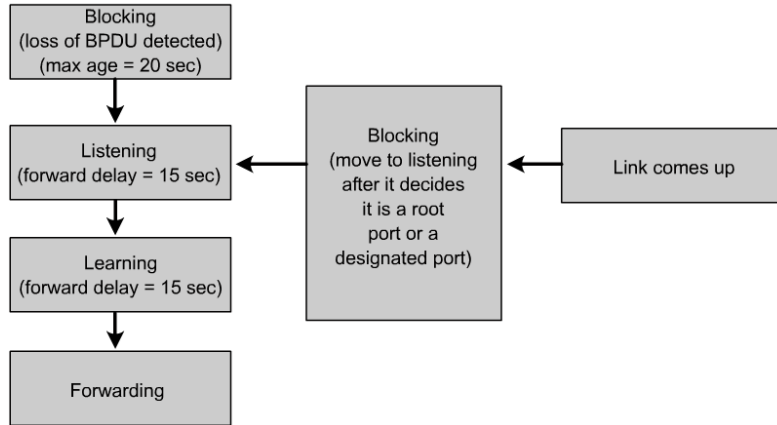


Stages of spanning-tree port states



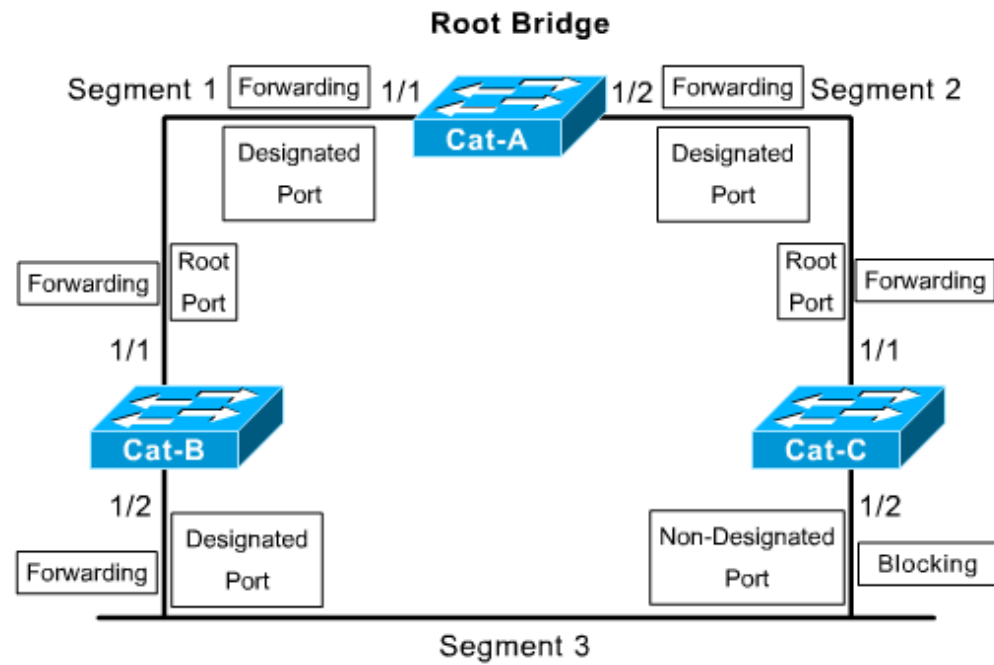

- Time is required for (BPDU) protocol information to propagate throughout a switched network.
- Topology changes in one part of a network are not instantly known in other parts of the network.
- There is propagation delay.
- **A switch should not change a port state from inactive (Blocking) to active (Forwarding) immediately, as this may cause data loops.**
- Each port on a switch that is using the Spanning-Tree Protocol has one of five states,

STP Port States



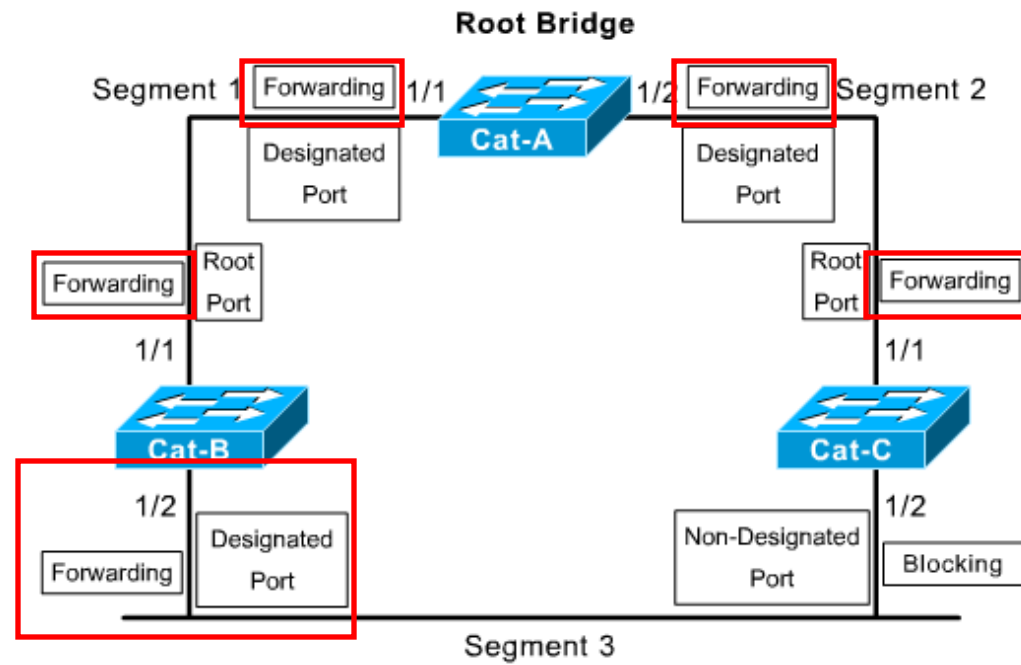
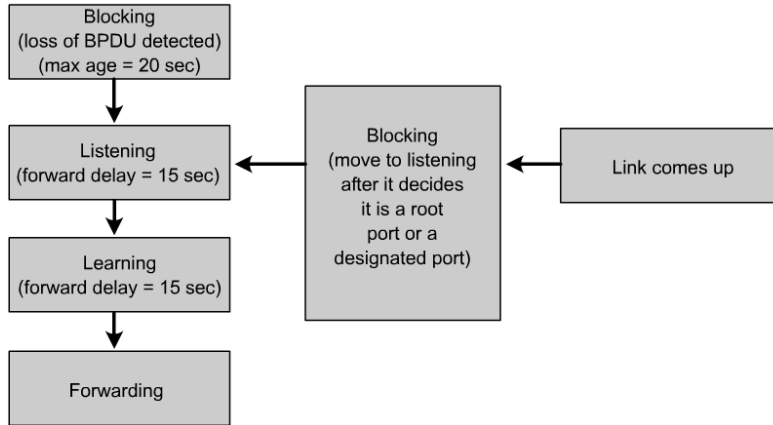
- In the **blocking state**, ports can **only receive BPDUs**.
 - **Data frames are discarded and no addresses can be learned.**
 - It may take up to 20 seconds to change from this state.
- Ports go from the blocked state to the **listening state**.
 - Switch **determines** if there are any other paths to the root bridge.
 - The **path that is not the least cost path to the root bridge goes back to the blocked state.**
 - The **listening period** is called the **forward delay** and lasts for **15 seconds**.
 - In the listening state, **user data is not being forwarded** and **MAC addresses are not being learned.**
 - **BPDUs are still processed.**

STP Port States




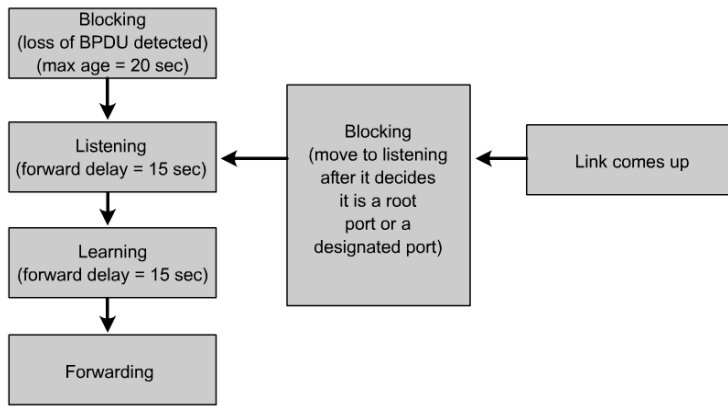
- Ports transition from the listening to the **learning state**.
 - In this state **user data is not forwarded, but MAC addresses are learned** from any traffic that is seen.
 - The **learning state lasts for 15 seconds** and is also called the **forward delay**.
 - **BPDUs are still processed.**

STP Port States





- A port goes from the learning state to the **forwarding state**.
 - In this state user data is forwarded and MAC addresses continue to be learned.
 - BPDUs are still processed.
- **Remember** – A switch port is allowed to transition to the Forwarding state only if no redundant links (loops) are detected and if the port has the best path to the Root Bridge as the **Root Port** or **Designated Port**.

Access ports



- When a device is connected to a port, the port normally moves from Blocking State to Listening state, for 15 seconds.
- When the Forward Delay timer expires, the port enters the Learning state, for 15 seconds.
- When the Forward Delay timer expires a second time, the port is transitioned to the Forwarding or Blocking state.
- This 30 seconds delay can cause a problem with computers asking for an IP address (DHCP) before the switch port has transitioned to Forwarding State.
- This causes the DHCP to fail, and the host to configure a default IP address.

Access ports



- When PortFast is enabled on a switch or trunk port, the port is immediately transitioned to the Forwarding state.
- As soon as the switch detects the link, the port is transitioned to the Forwarding state (less than 2 seconds after the cable is plugged in).
- This should only be enabled on switch ports where there are only hosts and not any switches.

```
Switch(config-if) # spanning-tree portfast
```

ALSwitch#show spanning-tree (Connecting a host without Portfast on)

VLAN0001

Spanning tree enabled protocol ieee

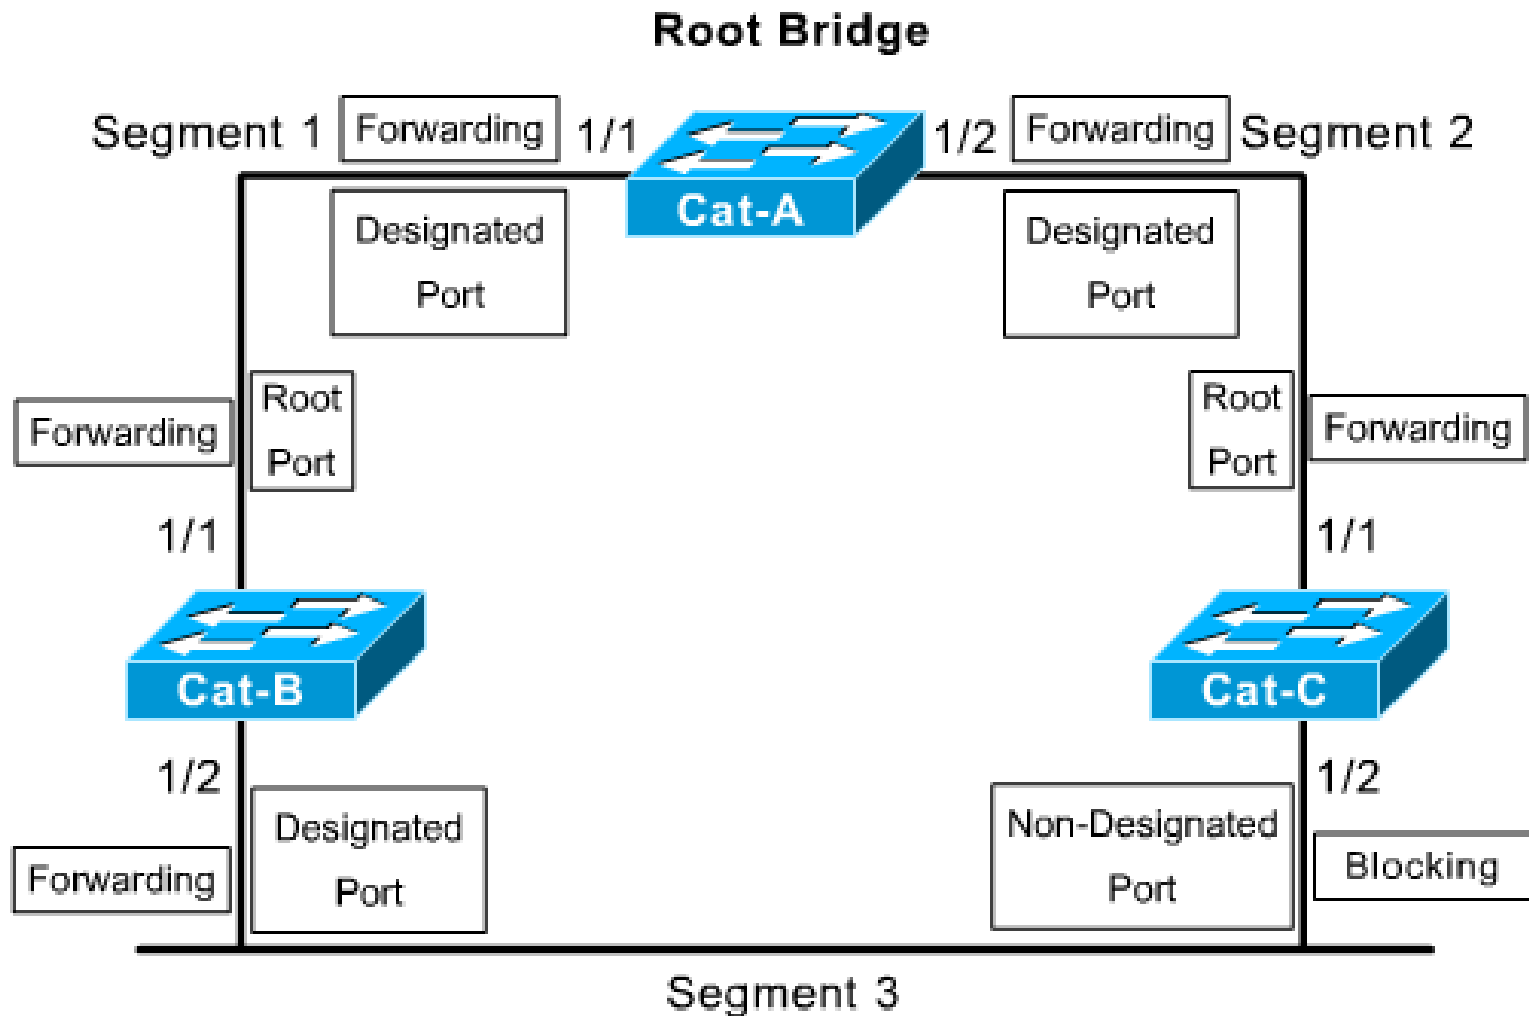
Root ID Priority 32768
 Address 0003.e334.6640
 Cost 19
 Port 23 (FastEthernet0/23)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

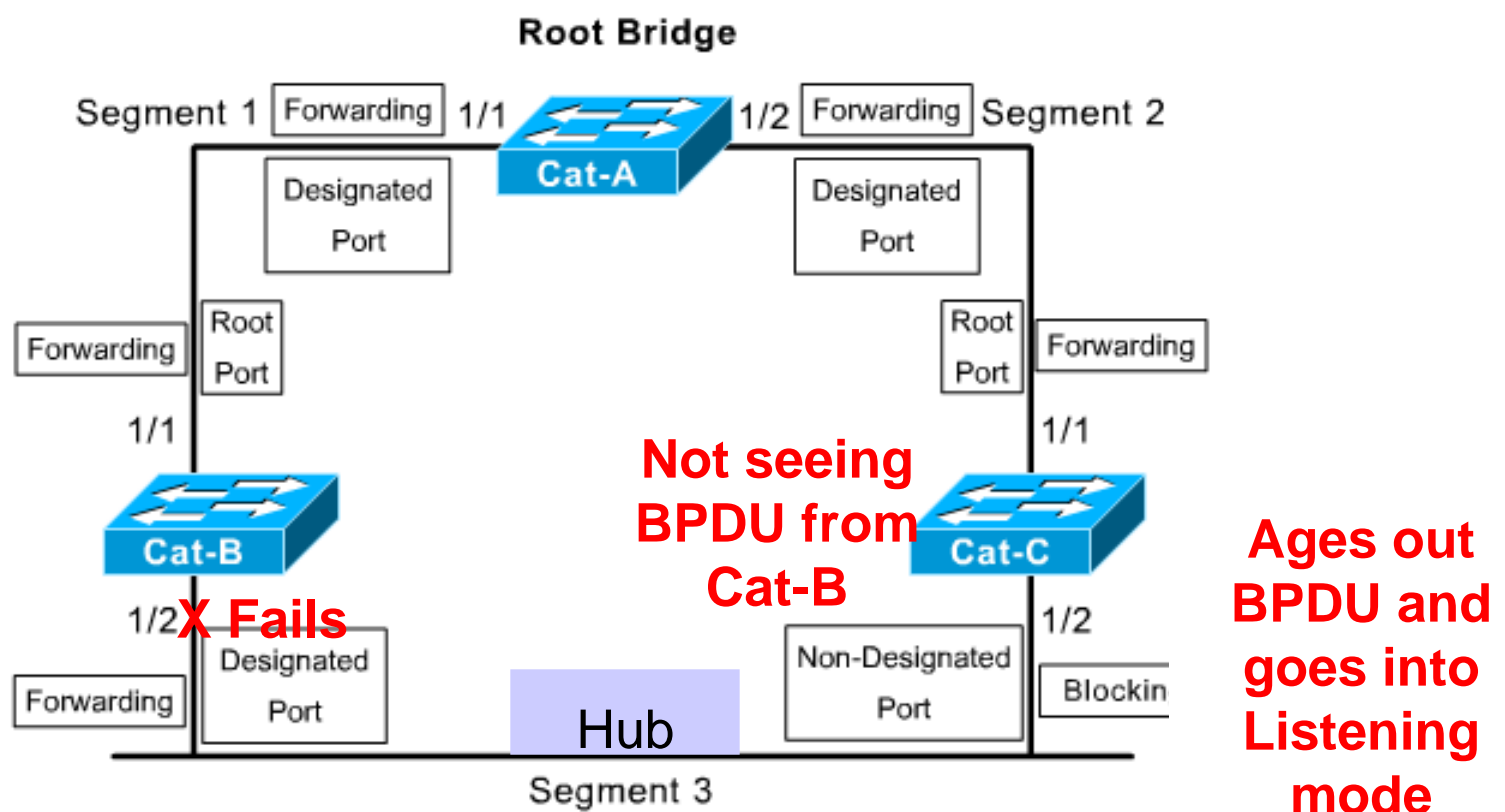
Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
 Address 000b.fc28.d400
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Aging Time 15

Interface Name	Port ID	Priority	Cost	Sts	Designated Cost	Bridge ID	Port ID
Fa0/8	128.8	19	19	LIS	19	32769 000b.fc28.d400	128.8
Fa0/23	128.23	19	19	FWD	0	32768 0003.e334.6640	128.25
Fa0/8	128.8	19	19	LRN	19	32769 000b.fc28.d400	128.8
Fa0/23	128.23	19	19	FWD	0	32768 0003.e334.6640	128.25
Fa0/8	128.8	19	19	FWD	19	32769 000b.fc28.d400	128.8
Fa0/23	128.23	19	19	FWD	0	32768 0003.e334.6640	128.25

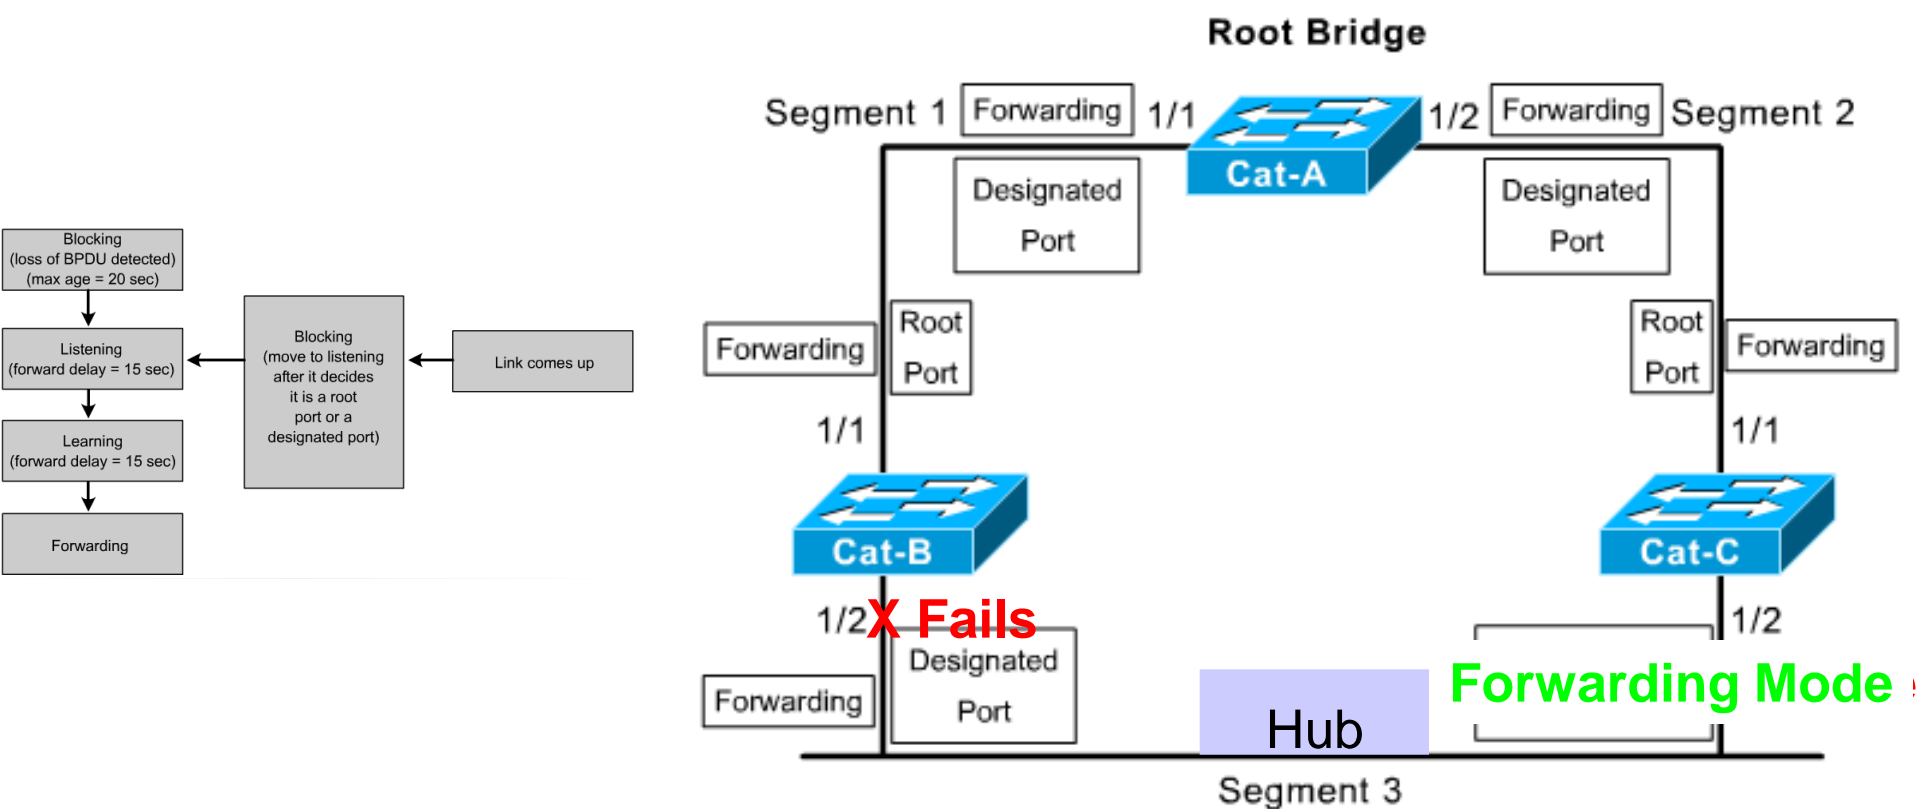
ALSwitch#

Example of redundant links





- Cat-B:1/2 fails.
- Cat-C has no immediate notification because it's still receiving a link from the hub.
- Cat-C notices it is not receiving BPDUs from Cat-B.
- **20 seconds (max age)** after the failure, Cat-C ages out the BPDU that lists Cat-B as having the DP for segment 3.
- This causes **Cat-C:1/2 to transition into the Listening state (15 seconds)** in an effort to become the DP.




- Because Cat-C:1/2 now offers the most attractive access from the Root Bridge to this link, it **eventually transitions to Learning State (15 seconds), then all the way into Forwarding mode.**
- In practice this will take **50 seconds (20 max age + 15 Listening + 15 Learning)** for Cat-C:1/2 to take over after the failure of Cat-B:1/2.

Port Cost/Port ID

Four-Step decision Sequence

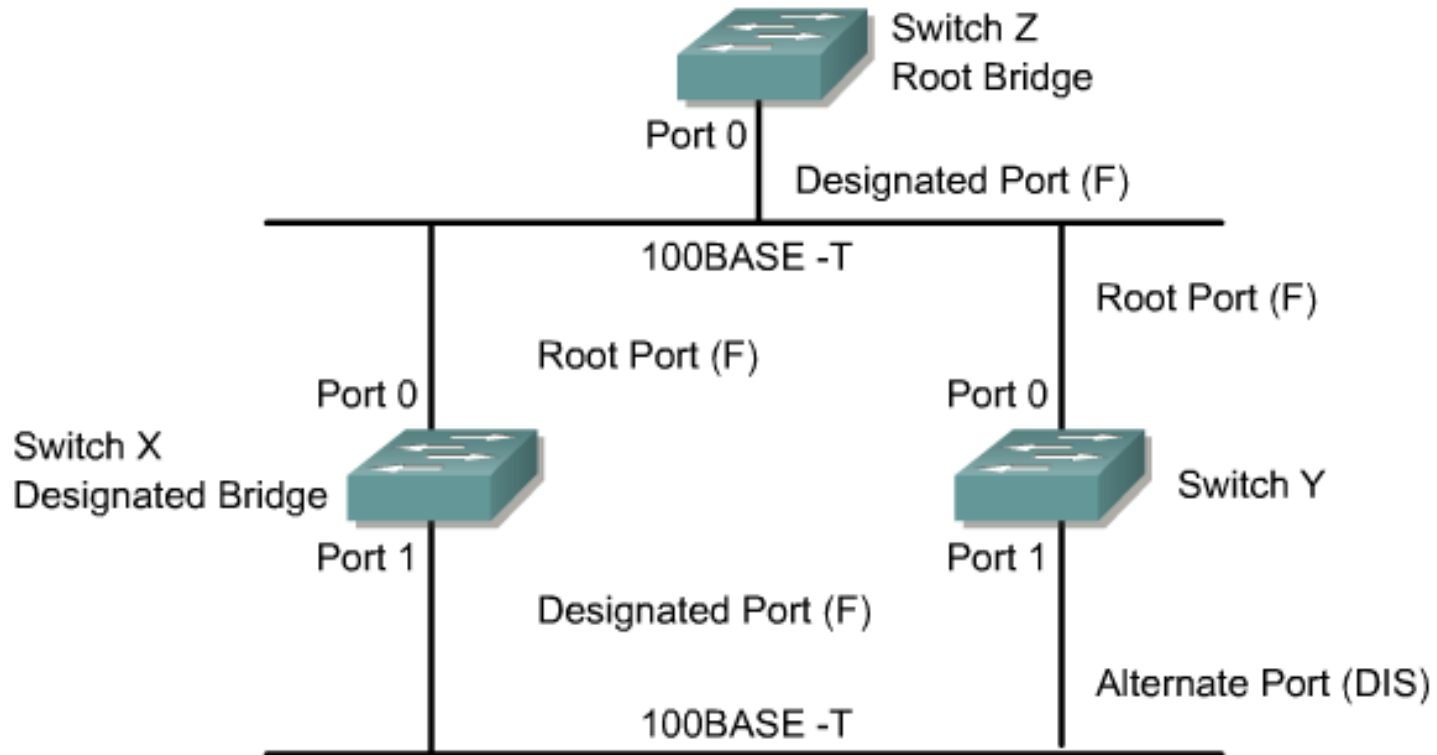
- Step 1 - Lowest BID
- Step 2 - Lowest Path Cost to Root Bridge
- Step 3 - Lowest Sender BID
- Step 4 - Lowest Port ID ←



Assume path cost and port priorities are default (32). Port ID used in this case. Port 0/1 would forward because it's the lower than Port 0/2.


- If the path cost and bridge IDs are equal (as in the case of parallel links), the switch goes to the port priority as a tiebreaker.
- Lowest port priority wins (all ports set to 32).
- You can set the priority from 0 – 63.
- If all ports have the same priority, the port with the lowest port number forwards frames.

Rapid Spanning Tree Protocol (RSTP)



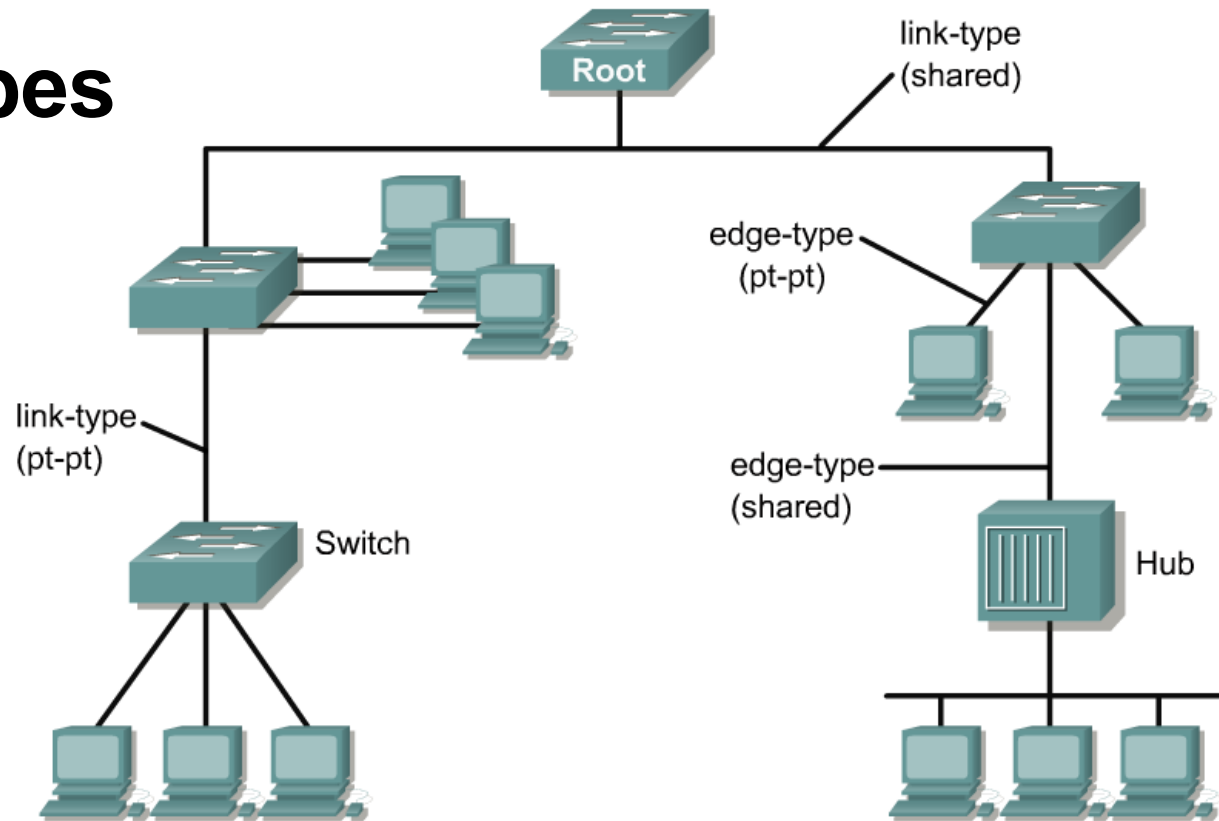
It is difficult to explain RSTP in just a few slides. RSTP is discussed in detail in CCNP 3.

Rapid Spanning Tree Protocol (RSTP)



- The Rapid Spanning-Tree Protocol is defined in the **IEEE 802.1w** LAN standard.
- The standard and protocol introduce the following:
 - Clarification of port states and roles
 - Definition of a set of link types that can go to forwarding state rapidly
 - Concept of allowing switches, in a converged network, to generate their own BPDUs rather than relaying root bridge BPDUs
- The “blocked” state of a port has been renamed as the “discarding” state.

RSTP Link Types



- **Link types** have been defined as **point-to-point**, **edge-type**, and **shared**.
- These changes allow failure of links in switched network to be learned rapidly.
- **Point-to-point links** and edge-type links can go to the forwarding state immediately.
- Network convergence does not need to be any longer than 15 seconds with these changes.
- The Rapid Spanning-Tree Protocol, IEEE 802.1w, will eventually replace the Spanning-Tree Protocol, IEEE 802.1D

RSTP Port States

STP (802.1D) Port State	RSTP (802.1w) Port State	Is Port Included in Active Topology?	Is Port Learning Mac Addresses?
Disabled	Discarding	No	No
Blocking	Discarding	No	No
Listening	Discarding	No	No
Learning	Learning	No	Yes
Forwarding	Forwarding	Yes	Yes