Global Knowledge®

Expert Reference Series of White Papers

Spanning Tree Protocol Essentials for the CCNA Exam

Spanning Tree Protocol Essentials for the CCNA Exam

John Pherson, Global Knowledge Course Director, Cisco Certified Instructor

The Need for Spanning Tree

The Spanning Tree Protocol (STP) was developed to solve the problems caused by loops in a bridged, or switched, topology. We can use the terms bridge and switch interchangeably in this discussion because they act in a similar fashion; switches are essentially very fast multi-port bridges.

The underlying cause of bridging loops is the fact that switches flood broadcast frames and frames with unknown destination MAC addresses. When there are loops in the topology, switches receive copies of the same frame on multiple ports which confuses the MAC table. It also causes switches to multiply the number of frame copies, since each switch floods the frames it is receiving. Each switch then starts receiving multiple copies of the frames and flooding the copies. This leads to broadcast storms and duplicate unicast frames on the network.

STP stops this process by first identifying a bridge, called a **root bridg**e, to serve as a reference point for the topology. Once the root bridge is selected, each non-root bridge determines the lowest-cost path back to the root, and identifies the port associated with that path as the root port. Once each bridge determines its **root port**, it then evaluates any other paths to the root that will cause loops. For each of these paths a single bridge, called the **designated bridg**e, is selected to provide connectivity for that segment back to the root. Any other switches on the same segment will be **non-designated**, and the ports on these bridges will be put in blocking state.

Blocking state does not mean the port is shut down; it simply means that no frames are allowed to be sent or received through that port. Spanning tree information continues to be received via that port from the designated bridge on that segment, which allows STP to communicate a change in topology should one occur.

The Bridge ID (BID)

Spanning Tree Priority (16 bits)	Switch MAC Address (48 bits)

Table 1. IEEE Spanning Tree Bridge ID

Each switch has a spanning-tree bridge ID composed of two parts, the priority and the MAC address. The MAC address is burned into the switch at the time of manufacture, and can't be changed. But the bridge ID as a whole can be modified by changing the priority assigned to the switch. Since STP always prefers the switch with the lowest bridge ID, lowering the priority value of a switch makes it more likely to become the root, and makes it more likely to be the designated switch for a segment. The **spanning-tree vlan 1 root [primary, secondary, # value]** command modifies the priority setting.

The default priority setting is 32768.

Spanning Tree Priority (4 bits)	VLAN ID (12 bits)	Switch MAC Address (48 bits)

Table 2. Extended System-ID Bridge ID

Cisco's implementation of spanning-tree is done, by default, at the VLAN level. Each VLAN runs its own instance of spanning-tree. This is called **PVST** (per-vlan spanning-tree). To support this, Cisco modified the bridge ID to include a 12-bit VLAN ID field in the bridge ID which is called the **Extended System I**D. The rationale for this was that the priority field did not need the granularity provided by a 16-bit field, but instead could get by with the top four bits of the field, leaving 12 bits to express a VLAN ID. This is consistent with the VLAN ID field in a dot1q frame, which also has a 12-bit VLAN ID field. It means then, that the priority bits have decimal values of 4096, 8192, 16384, and 32768. Using this form of the bridge ID means that priorities are always expressed as some multiple of 4096.

Path and Port Cost

Port Speed	Port Cost
10Mb	100
100Mb	19
1 Gb	4
10 Gb	2

Table 3. Port Costs

The costs used by STP are derived from the above table, which assigns cost based on the port speed. A switch directly attached to the root looks at the port by which it is attached, and derives its **path cost** to the root from the speed of that port. These switches then advertise this path cost to their downstream neighbors. Each downstream switch adds the path cost advertised by the upstream switch to the cost associated with the local port receiving the information. The sum of these costs becomes their own path cost back to the root.

The table shows the current cost scheme used by IEEE (traditional 802.1d) Spanning Tree and Rapid Spanning-Tree protocols. This cost scheme supersedes an earlier cost scheme that ran out of headroom as port speeds became faster. It has also been superseded by a more recent cost scheme used on a variant of spanning-tree called Multiple Spanning-Tree, which is outside the scope of the CCNA content. You should be familiar with the costs listed above.

BPDUs

Field	Size (Bytes)
Protocol	2
Version	1
Type	1
Flags	1
Root BID	**8**
Path Cost	**4**
BID	**8**
Port ID	**2**
Message Age	2
Max Age Time	2
Hello Time	2

Table 4. BPDU Fields

Bridge Protocol Data Units (BPDUs) are the frames sent by STP to convey information. They are sent at an interval determined by the **hello-time** parameter, which is set to two seconds by default. The most important fields of the BPDU are highlighted in bold in the table above. The Root BID is the bridge ID of the root switch. The BID is the bridge ID of the switch sending the BPDU. The path cost is this switch's path cost back to the root. The port ID is the identifier of the port transmitting this BPDU.

A BPDU is preferred, or superior to others, if it advertises the true root (the switch with the lowest BID in the topology), and the lowest path cost. It must come from the switch with the lowest BID, and must be transmitted from the lowest port ID on that switch. Therefore, the path along which these BPDUs travel becomes the preferred path back to the root.

Spanning Tree Port Roles

Port Role	**Behavior**	**Definition**
Root Port	Forwarding	Port by which frames LEAVE switch to reach the root
Designated Port	Forwarding	Port by which frames ENTER switch to reach the root
Non-Designated Port	Blocking	Port blocking frames to prevent a loop in the topology

Table 5. Port Roles

As a result of the spanning-tree calculation, ports will assume various roles in the topology. A root port is a port facing towards the root that is connected to the best path back to the root. Best path means:

1) The path with the lowest cost back to the root
2) The path going through the switch with the lowest BID if there is more than switch advertising the lowest cost
3) The lowest port ID on that switch if there is more than one connection to the switch

Each segment must have a designated port, which is a port facing away from the root that frames will enter to reach the root. Each segment in the topology will have only one designated port. The designated port is on the switch attached to that segment with the lowest BID (designated switch).

All switches attached to a segment that are not attached via their root port, and not designated for that segment, must put their ports in blocking state (non-designated port) to prevent loops.

Spanning Tree States and Timers

State	Timer	Default Duration
Blocking	Max Age	20 sec
Listening	Forward Delay	15 sec
Learning	Forward Delay	15 sec
Forwarding	N/A	N/A

Table 6. STP States and Timers

STP (IEEE 802.1d) proceeds through various states to establish the topology. Ports start out in blocking state and then proceed into listening state where they exchange BPDUs with other switches to establish the root switch, root ports, and designated ports. They then enter a learning phase where they listen to frames hitting their ports to build the MAC tables on the switch. Finally, the root and designated ports enter forwarding state, while the non-designated ports are set to blocking state.

The timers associated with this process are the **Max Age** parameter which controls the initial blocking time, and the **Forward Delay** which controls the duration of both the listening and learning intervals. These timers are what contribute to the relatively slow convergence time of STP, which can be up to 50 seconds if the full blocking interval is observed. Cisco implementations prematurely end the blocking time to get convergence time down to about 30 seconds.

Spanning Tree Rules

1) The first rule of STP is that the lowest always wins.
2) The lowest BID becomes the root.
3) The lowest cost path is preferred.
4) If more than one switch advertises the lowest path cost, the one with the lowest BID is preferred.
5) If the lowest BID switch advertises the path through multiple connected ports, the port with the lowest port ID is preferred.

The Four-Switch Example

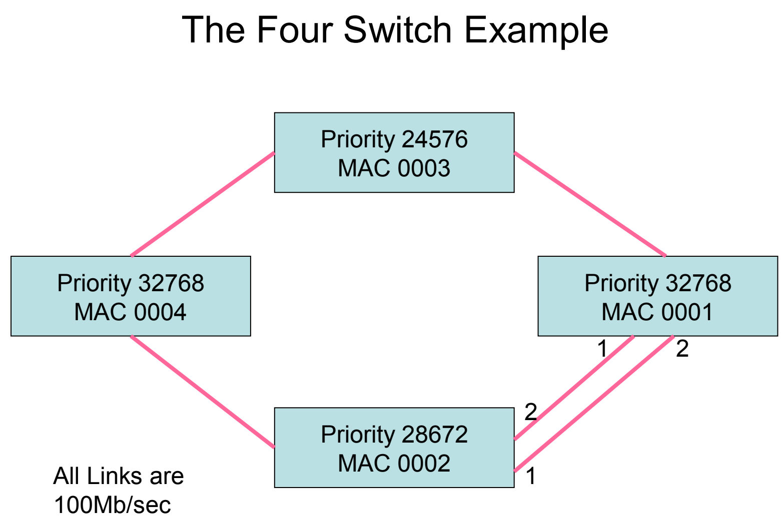

Figure 1. The Four Switch Example

To illustrate how the rules work, consider the topology above. All links are Fast Ethernet, and therefore have a port cost of 19.

The first task is selecting the root switch. Because Switch 0003 has the lowest priority, it will become the root. Its ports will all become designated ports.

The two switches attached to the root, Switches 0001 and 0004, will evaluate their path costs to the root. They each have a direct path with a cost of 19, and an alternate path with a cost of 57 (3 x 19). Therefore, they select their directly connected ports as their root ports.

Switch 0002 has a dilemma, since it has two equal cost paths back to the root, each with a cost of 38. Since it prefers the path advertized by the switch with the lowest bridge ID, it will prefer the path through Switch 0001. Since the priorities of Switches 0001 and 0004 are the same, the lowest BID is determined by the MAC addresses.

In choosing the path through Switch 0001, Switch 0002 still has a dilemma, since there are two connections between the switches. The path advertised by Switch 0001 through its Port 1 is preferred over the one advertised through Port 2, since lowest port always breaks the tie. Therefore Switch 0002 will make its Port 2 the root port, and Port 1 on Switch 0001 becomes designated.

There are still two loops left in the topology to deal with. The first is the redundant connection between Switch 0001 and Switch 0002. On this link, between Port 1 on Switch 0002 and Port 2 on Switch 0001, one switch must be designated. Since Switch 0002 has a lower bridge ID, by virtue of its lower priority, it becomes designated for that segment and Port 1 is placed in forwarding state, while Switch 0001 Port 2 becomes non-designated and goes into blocking state.

On the link between Switch 0002 and 0004, Switch 0002 has the lower BID, and is therefore designated for that segment. It will forward on the port attached to this segment, while Switch 0004 puts its port in blocking state as a non-designated port for that segment.

"Show Spanning-Tree" Output

```
Switch# show spanning-tree vlan 1
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority     24577
              Address      0001.4332.1230
              Cost         38
              Port         1 (FastEthernet0/1)
              Hello Time    2 sec   Max Age 20 sec   Forward Delay 15
sec

Bridge ID  Priority   32769   (priority 32768 sys-id-ext 1)
              Address      0003.9997.7650
  Hello Time    2 sec   Max Age 20 sec   Forward Delay 15 sec
  Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------
------------
Fa0/1            Root FWD    19       128.1    P2p
Fa0/2            Altn BLK    19       128.2    P2P
```

Knowing the output of this command is essential to achieving success on the spanning tree portion of the CCNA exam. Notice that the output is specific to VLAN 1. Since Cisco implements STP on a VLAN basis this output relates only to the instance of STP running on VLAN 1.

The first thing to notice is the version of STP running. The output says "Spanning tree enabled protocol ieee". This means the switch is running traditional 802.1D Spanning Tree, not Rapid Spanning Tree or MST.

Notice that the Root ID identifies the BID of the root bridge. This is broken out separately into the priority portion and the MAC portion of the BID. Under that you will notice a cost listed of 38. This is the path cost from your switch to the root. A cost of 38 implies that there are probably two fast Ethernet links between you and the root, which means there is probably another switch between your switch and the root.

The next section lists the Bridge ID of your switch. You will also notice that it is likewise broken out into the priority portion and the MAC portion, and that the priority is further broken out into the priority (top 4 bits) and the sys-id-ext of 1 (bottom 12 bits), which represents VLAN 1.

The interface section lists the ports on your switch that are participating in this STP instance. Notice that port fa0/1 is the root port, meaning that you are connected via fa0/1 to the root. There is also another port listed, fa0/2, in blocking state. It is has a role defined as **Altn** which is actually a Rapid Spanning Tree term. Most switches these days are Rapid Spanning Tree capable, but they run in IEEE mode (traditional STP) by default, so although this is actually a non-designated, blocking port, it uses the newer role terminology associated with Rapid Spanning Tree Protocol (RSTP).

Since there is a port on the switch in blocking state, it indicates that this switch is configured in a looped topology and has a redundant path back to the root. Your switch must also have a higher bridge ID than the switch connected to your port fa0/2 or else they would be blocking, not you.

You will also notice that the interface list contains a column that says Prio/Nbr. This is indicating the port ID. The port ID is comprised of both the port number and a priority. Like the bridge priority, this can be manipulated to modify the spanning tree preference for the lowest numbered port. By lowering the port priority, we can make spanning-tree favor a higher numbered port. The port ID is an 8-bit field with values of 0-255 and a default setting of 128.

Other Things You Need To Know about Spanning Tree

Beyond this discussion of traditional Spanning Tree, there are a few other things you will have to know about STP that we don't have time for in this paper. These include the enhancements to STP that are Cisco-specific. Probably the most important of these to know is the **portfast** feature. This is used on access ports (ports connected to end devices) to force them directly into forwarding mode while still affording the protection of STP if they are accidentally connected in a loop.

The other topic you should know about is the basic operation of Rapid Spanning Tree Protocol (RSTP-802.1w). The CCNA treatment of this topic is fairly general but includes things like the differences between this and traditional STP, new terminology, the "proposal and acknowledgement" process, and the improved convergence time. These, however, are topics for another day.

Here's wishing you success on your pursuit of CCNA certification, and I hope this white paper will be helpful in those endeavors.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge courses:

ICND1 – Interconnecting Cisco Network Devices 1

ICND2 – Interconnecting Cisco Network Devices 2

CCNA Boot Camp v2.0

BCMSN – Building Cisco Multilayer Switched Networks v3.0

For more information or to register, visit **www.globalknowledge.com** or call **1-800-COURSES** to speak with a sales representative.

Our courses and enhanced, hands-on labs offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 1,200 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and business training needs.

About the Author

John Pherson is a senior instructor and course director for Global Knowledge with 30 years' experience in the computer field and over 13 years of teaching experience. He is certified to teach over a dozen Cisco courses in the Routing and Switching, Data Center, and Voice areas. He has also been a Microsoft and Novell Certified Instructor. In addition to numerous technical certifications, he holds a B.S. in Business Administration and a Masters degree in adult education.