Global Knowledge®

Expert Reference Series of White Papers

How a Spanning Tree Works

How a Spanning Tree Works

Carol Kavalla, Global Knowledge Instructor, BS, CCSI, CCDP

Introduction

This paper focuses on Spanning Tree Protocol (STP), IEEE Standard 802.1D. But first, it's important to understand the functions of an Ethernet LAN switch:

- Address learning
- Address filtering
- Forwarding of frames
- Loop avoidance

The way a switch learns Ethernet (MAC) addresses is by inspecting the Ethernet frame and recording the source MAC address in a dynamic table. The switch will also associate a learned MAC address with a port. It can then make intelligent forwarding decisions based on the destination MAC address. This white paper illustrates the process.

Ethernet Header

First, the Ethernet LAN header contains information about the source MAC address and the destination MAC address.

Preamble 101010...11	**Destination MAC ADD**	**Source MAC Add**	**Type/Length**	**Data**	**FCS**

Ethernet Header

When host A sends a frame to host B on the same LAN, the switch learns the MAC address and port number, then stores it in the MAC-Address-Table (sometimes called the Content Addressable Memory (CAM) Table in larger switches).

Mac-Address-Table
Before Host A sends to Host B

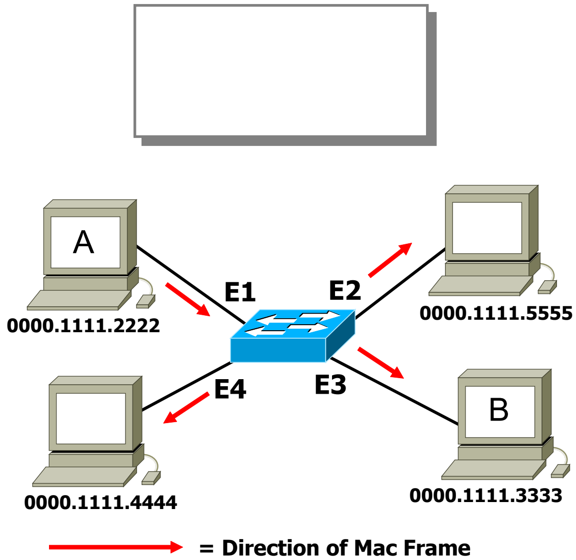

Before any activity, the MAC-Address-Table is empty. Once Host A sends data to Host B, the MAC-Address-Table gets populated with Host A's MAC address. Because the switch does not know which port the destination MAC address is associated with, it floods the frame out all of the ports except the one it arrived on (in this case Port E1).

Mac-Address-Table
After Host A sends to Host B

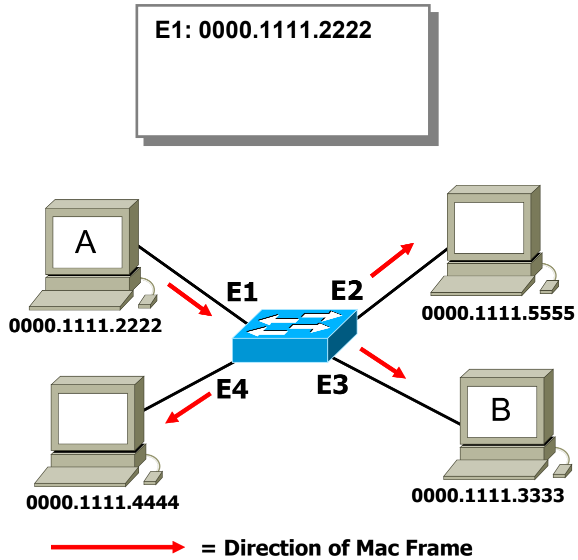

When Host B replies to Host A, the switch first learns the MAC address associated with Host B and the port to which it is attached. At this point, the switch will only forward the frame out of Port E1 because the MAC-address-table has already been populated with Host A's MAC and its associated port.

Forwarding Frames

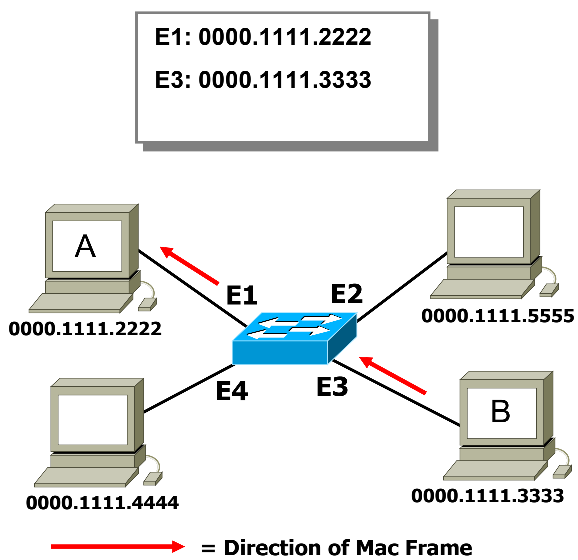

Eventually, all host MAC addresses will be learned in the same way, and the MAC-Address-Table will be populated.

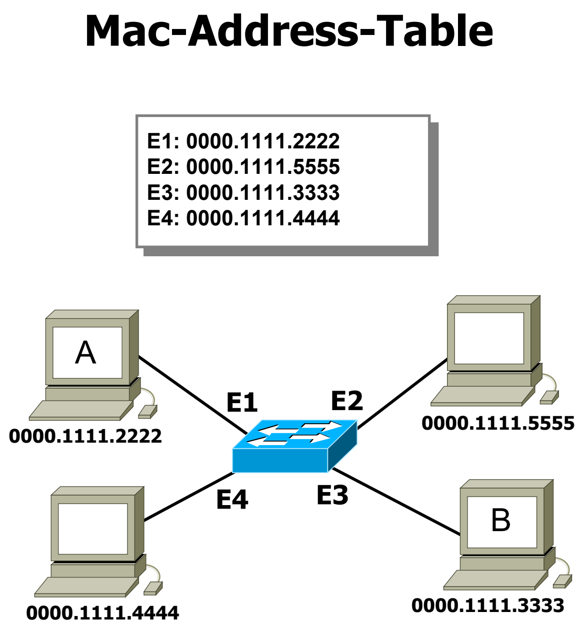

Besides showing how a switch forwards frames, the preceding illustrates two other important switching concepts. One, if a switch does not have the destination MAC address in its table, or if the destination address is a broadcast (all ones in the destination MAC), the switch will flood the frame out all ports except the port it came in on. Second (and a partial follow-up on one), the default behavior for a switch is to flood.

Loops

Loop avoidance is another function of the switch. How do loops occur and how can they be prevented? One way for loops to occur is when a broadcast occurs on the LAN. Any dynamic discovery protocol may generate a broadcast Ethernet frame, for example, DHCP (Dynamic Host Configuration Protocol) DNS (Domain Name Service), or an ARP (Address Resolution Protocol). (ARP was discussed in a previous white paper, Router Vulnerabilities.) The switch will forward the broadcast frame out of all ports except the port that it came in on. If there is redundancy in the network, there will be a loop and, consequently, a broadcast storm, which is the endless forwarding of the same frame.

In the following example, Host A is ARPing for the MAC address of B. Because the switches have redundant links, the broadcast will continue to get propagated until the network can no longer function due to high bandwidth utilization of the links and high processor utilization on the hosts.

Broadcast Storm

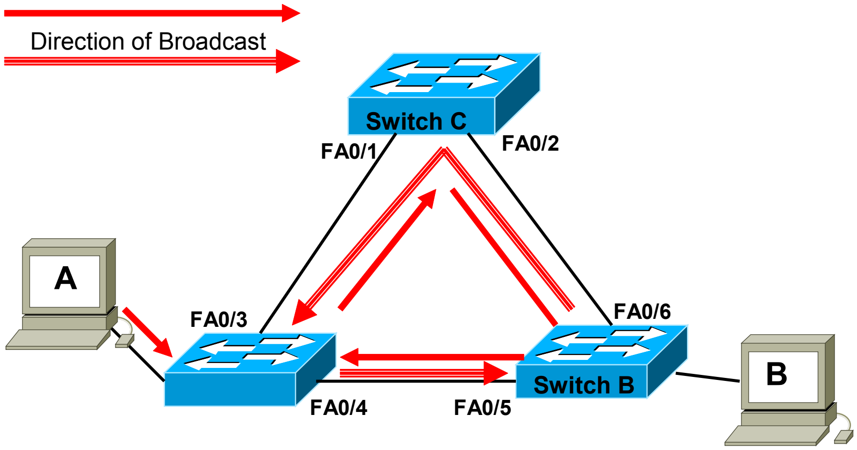

- **Host A sends a broadcast.**
- **Switches continue to propagate broadcast traffic over and over**

The solution for this problem is Spanning Tree Protocol (STP). STP is on by default on all Cisco switches. STP chooses a point of reference called the Root Bridge. By default the bridge with the lowest Bridge ID is chosen as the Root. The Bridge ID is a 64-bit field made up of the priority (the default is 32768) and the MAC address. If the default priority has not been changed, the Root will be chosen based on the lowest MAC address. The Root Bridge then sends Bridge Protocol Data Units (BPDUs) out all of its designated (forwarding) ports. If a bridge is not the Root (called a Non-Root Bridge), it will receive regular BPDU on the port that is closest to the Root, called the Root Port.

In the diagram below, the Root Bridge sends BPDUs out ports FA0/1 and FA0/2 (its designated ports). Switch A receives the BPDU on root port FA0/3 and then sends a BPDU out of port FA0/4. Likewise, Switch B receives a BPDU on root port FA0/6 and sends a BPDU out on port FA0/5. Both Switch A and Switch B have received BPDUs on two different ports. This indicates that there is a loop in the network. Since each non-root switch can have only one root port, the loop is identified by the BPDUs on the non-root ports showing a longer distance to the root, indicating that the link must not be used for user frames.

Root Election & BPDU Propagation

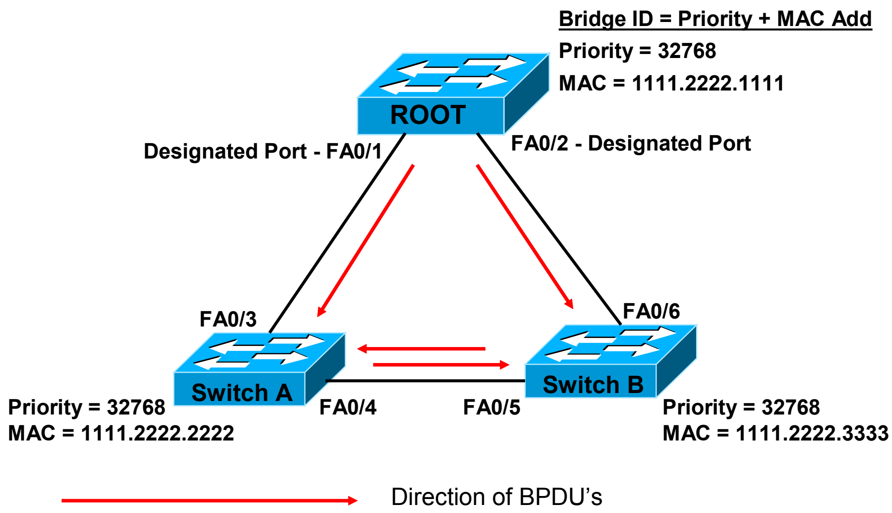

In this example, STP will block one of the ports on one of the non-root bridges. To determine this, the following criteria have to be taken into consideration.

- Each non-Root Bridge needs a Root Port (a forwarding port).
- The Root Port is the port that is closest to the Root Bridge.
- The port that is farthest away from the Root Bridge will be blocked.
- In the case of a tie, the switch with the highest Bridge ID (BID) will be blocked.

In this case, Switch A's Root Port would be port Fa0/3, and Switch B's Root Port would be Fa0/6. Assuming that all links are 100 Mbps, Switch A's path back to the root from Fa0/4 is equal to the path from FA0/5 of Switch B back to the root. Because the path would be equal, the tie breaker will be the MAC address. The MAC address of Switch B's port FA0/5 is higher than Switch A's MAC address for port Fa0/4; therefore, Switch B will block port Fa0/5. That would make FA0/4 on Switch A a Designated (forwarding) Port.

Final STP Topology

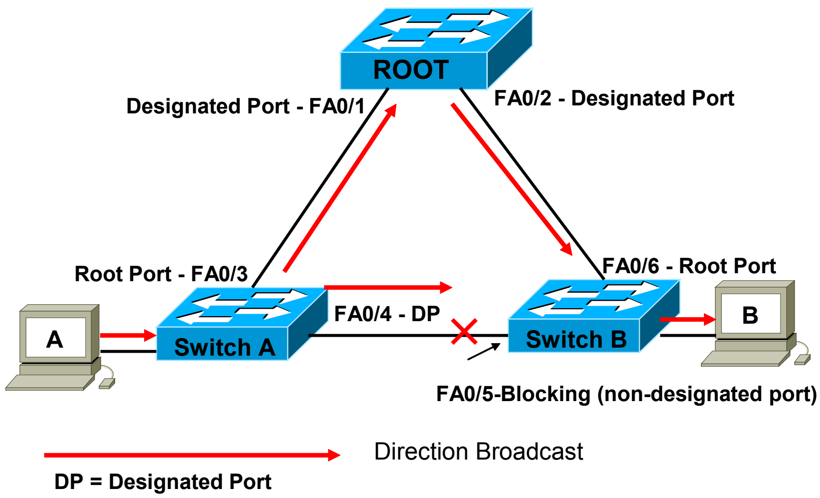

The preceding diagram is the converged topology for STP. With STP, when Host A ARPs for the MAC address of Host B, Switch A forwards the broadcast out port FA0/3 toward the Root and Fa0/4 toward Switch B. The Root receives the broadcasts on port Fa0/1 and forwards it out port Fa0/2 toward Switch B. Switch B drops the broadcast on port Fa0/5 as it is a blocking port, and Switch B receives a copy of the broadcast on port Fa0/6 from the Root. Each switch has received the broadcast one time, and there is no loop.

One of the problems with IEEE 802.1d is the time it takes for convergence; or the time it takes for all ports to be in either the forwarding or blocking state after a network topology change. This is because the spanning tree algorithm was timer-based with lengthy timers. At the very least, after a topology change, a switch would wait twice the forward delay (30 seconds) before it began forwarding traffic.

This is considered to be an unnecessarily long time for reconvergence on a LAN, especially when having to support real-time applications like IP Telephony, for example. The IEEE implemented a newer version of spanning tree protocol in 1999 called Rapid Spanning Tree Protocol (RSTP) or 802.1w. Now the recommended protocol for loop avoidance, it basically works the same as 802.1D. But this protocol no longer relies on lengthy timers for reconvergence; rather, it uses a new bridge-bridge handshake mechanism and shorter timers, which allow ports to move to forwarding much more quickly.

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

CCDA Boot Camp

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 700 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and management training needs.

About the Author

Carol Kavalla's background includes teaching at Rockland Community College in New York, managing networks and being a consultant for the NYS small business development center. For the last eight and a half years Carol has taught for Global Knowledge and is certified to teach nine Cisco Courses: ICND1; ICND2; CCDA; BSCI; BCMSN; TCN; ICMI; BGP; and ARCH. She also has a consulting firm in Charleston, South Carolina, where she works with small companies (100-200 nodes) installing, configuring routers and switches, and troubleshooting network problems.