

Bridging and Switching

Bridging is used to connect two network segments. This alleviates congestion problems on a single Ethernet segment and extends allowed cabling distances because the segments on each side of the bridge conformed to the same distance limitation as a single segment. This bridge is called "transparent bridging" because the end-point devices do not need to know that the bridge exists.

Transparent bridges forward frames only when necessary and, thus, reduces network overhead. To accomplish this, transparent bridges learn MAC addresses by examining the source MAC address of each frame received by the bridge; decides when to forward a frame or when to filter a frame, based on the destination MAC address; and creates a loop-free environment with other bridges by using the Spanning Tree Protocol.

Generally, broadcasts and multicast frames are forwarded by the bridge in networks that use bridges. In addition, transparent bridges perform switching of frames using Layer 2 headers and Layer 2 logic and are Layer 3 protocol-independent. Store-and-forward operation, which means that the entire frame is received before the first bit of the frame is forwarded, is also typical in transparent bridging devices. However, the transparent bridge must perform processing on the frame, which also can increase latency.

A transparent bridge operates in the following manner:

- . The bridge has no initial knowledge of the location of any end device; therefore, the bridge must listen to frames coming into each of its ports to figure out on which network a device resides.

- . The bridge constantly updates its bridging table upon detecting the presence of a new MAC address or upon detecting a MAC address that has changed location from one bridge port to another. The bridge is then able to forward frames by looking at the destination address, looking up the address in the bridge table, and sending the frame out the port where the destination device is located.

the frame out all available ports. However, the frame is not forwarded out the port that initially received the frame. Hence, broadcasts are able to reach all available networks. A bridge only segments collision domains but does not segment broadcast domains.

- . If a frame arrives with a destination address that is not found in the bridge table, the bridge is unable to determine which port to forward the frame to for transmission. This is known as an unknown unicast. In this case, the bridge treats the frame as if it was a broadcast and forwards it out all remaining ports. After a reply to that frame is received, the bridge will learn the location of the unknown station and add it to the bridge table.

- . Frames that are forwarded across the bridge cannot be modified.

An Ethernet switch uses the same logic as a transparent bridge, but performs more functions, has more features, and has more physical ports. Switches use hardware to learn MAC addresses and to make forwarding and filtering decisions, whereas bridges use software.

A switch listens for frames that enter all its interfaces. After receiving a frame, a switch decides whether to forward a frame and out which port(s). To perform these functions, switches perform three tasks:

- . Learning, which means that the switch learns MAC addresses by examining the source MAC address of each frame the bridge receives. Switches dynamically learn the MAC addresses in the network to build its MAC address table. With a full, accurate MAC address table, the switch can make accurate forwarding and filtering decisions. Switches build the MAC address table by listening to incoming frames and examining the frame's source MAC address. If a frame enters the switch, and the source MAC address is not in the address table, the switch creates an entry in the table. The MAC address is placed in the table, along with the interface in which the frame arrived. This allows the switch to make good forwarding choices in the future. Switches also forward unknown unicast

frames, which are frames whose destination MAC addresses are not yet in the bridging table, out all ports, which is called flooding, with the hope that the unknown device will be on some other Ethernet segment and will reply. When the unknown device does reply, the switch will build an entry for that device in the address table

. Forwarding or filtering, which means that the switch decides when to forward a frame or when to filter it, i.e., not to forward it, based on the destination MAC address. Switches reduce network overhead by forwarding traffic from one segment to another only when necessary. To decide whether to forward a frame, the switch uses a dynamically built table called a bridge table or MAC address table. The switch looks at the previously learned MAC addresses in an address table to decide where to forward the frames.

. Loop prevention, which means that the switch creates a loop-free environment with other bridges by using Spanning Tree Protocol (STP). Having physically redundant links helps LAN availability, and STP prevents the switch logic from letting frames loop around the network indefinitely, congesting the LAN.

Frames sent to unicast addresses are destined for a single device; frames sent to a broadcast address are sent to all devices on the LAN. Frames sent to multicast addresses are meant for all devices that care to receive the frame. Thus, when a switch receives a frame, it checks if the address is a unicast address, a broadcast address or a multicast address. If the address is unicast, and the address is in the address table, and if the interface connecting the switch to the destination device is not the same interface on which the frame arrived, the switch forwards the frame to the destination device. If the address is not in the address table, the switch forwards the frame on all ports. If the address is a broadcast or multicast address, the switch also forwards the frame on all ports.

The internal processing on a switch can decrease latency for frames. Switches can use store-and-forward processing as well as cut-through processing logic. With cut-through processing, the first bits of the frame are sent out the outbound port before the last bit of the incoming frame is received. However, because the frame check sequence (FCS) is in the Ethernet trailer, a cut-through forwarded frame might have bit errors that the switch will not notice before sending most of the frame.