

Bridging and Switching Basics

This chapter introduces the technologies employed in devices loosely referred to as *bridges* and *switches*. Topics summarized here include general link-layer device operations, local and remote bridging, ATM switching, and LAN switching. Chapters in Part 4, “Bridging and Switching,” of this book address specific technologies in more detail.

What are Bridges and Switches?

Bridges and switches are data communications devices that operate principally at Layer 2 of the OSI reference model. As such, they are widely referred to as data link layer devices.

Bridges became commercially available in the early 1980s. At the time of their introduction, bridges connected and enabled packet forwarding between homogeneous networks. More recently, bridging between different networks has also been defined and standardized.

Several kinds of bridging have proven important as internetworking devices. *Transparent bridging* is found primarily in Ethernet environments, while *source-route bridging* occurs primarily in Token Ring environments. *Translational bridging* provides translation between the formats and transit principles of different media types (usually Ethernet and Token Ring). Finally, *source-route transparent bridging* combines the algorithms of transparent bridging and source-route bridging to enable communication in mixed Ethernet/Token Ring environments.

Today, switching technology has emerged as the evolutionary heir to bridging based internetworking solutions. Switching implementations now dominate applications in which bridging technologies were implemented in prior network designs. Superior throughput performance, higher port density, lower per-port cost, and greater flexibility have contributed to the emergence of switches as replacement technology for bridges and as complements to routing technology.

Link-Layer Device Overview

Bridging and switching occur at the link layer, which controls data flow, handles transmission errors, provides physical (as opposed to logical) addressing, and manages access to the physical medium. Bridges provide these functions by using various link-layer protocols that dictate specific flow control, error handling, addressing, and media-access algorithms. Examples of popular link-layer protocols include Ethernet, Token Ring, and FDDI.

Bridges and switches are not complicated devices. They analyze incoming frames, make forwarding decisions based on information contained in the frames, and forward the frames toward the destination. In some cases, such as source-route bridging, the entire path to the destination is contained in each frame. In other cases, such as transparent bridging, frames are forwarded one hop at a time toward the destination.

Upper-layer protocol transparency is a primary advantage of both bridging and switching. Because both device types operate at the link layer, they are not required to examine upper-layer information. This means that they can rapidly forward traffic representing any network-layer protocol. It is not uncommon for a bridge to move AppleTalk, DECnet, TCP/IP, XNS, and other traffic between two or more networks.

Bridges are capable of filtering frames based on any Layer 2 fields. A bridge, for example, can be programmed to reject (not forward) all frames sourced from a particular network. Because link-layer information often includes a reference to an upper-layer protocol, bridges usually can filter on this parameter. Furthermore, filters can be helpful in dealing with unnecessary broadcast and multicast packets.


By dividing large networks into self-contained units, bridges and switches provide several advantages. Because only a certain percentage of traffic is forwarded, a bridge or switch diminishes the traffic experienced by devices on all connected segments. The bridge or switch will act as a firewall for some potentially damaging network errors, and both accommodate communication between a larger number of devices than would be supported on any single LAN connected to the bridge. Bridges and switches extend the effective length of a LAN, permitting the attachment of distant stations that were not previously permitted.

Although bridges and switches share most relevant attributes, several distinctions differentiate these technologies. Switches are significantly faster because they switch in hardware, while bridges switch in software and can interconnect LANs of unlike bandwidth. A 10-Mbps Ethernet LAN and a 100-Mbps Ethernet LAN, for example, can be connected using a switch. Switches also can support higher port densities than bridges. Some switches support cut-through switching, which reduces latency and delays in the network, while bridges support only store-and-forward traffic switching. Finally, switches reduce collisions on network segments because they provide dedicated bandwidth to each network segment.

Types of Bridges

Bridges can be grouped into categories based on various product characteristics. Using one popular classification scheme, bridges are either *local* or *remote*. Local bridges provide a direct connection between multiple LAN segments in the same area. Remote bridges connect multiple LAN segments in different areas, usually over telecommunications lines. Figure 4-1 illustrates these two configurations.

Figure 4-1 Local and remote bridges connect LAN segments in specific areas.



Remote bridging presents several unique internetworking challenges, one of which is the difference between LAN and WAN speeds. Although several fast WAN technologies now are establishing a presence in geographically dispersed internetworks, LAN speeds are often an order of magnitude faster than WAN speeds. Vast differences in LAN and WAN speeds can prevent users from running delay-sensitive LAN applications over the WAN.

Remote bridges cannot improve WAN speeds, but they can compensate for speed discrepancies through a sufficient buffering capability. If a LAN device capable of a 3-Mbps transmission rate wants to communicate with a device on a remote LAN, the local bridge must regulate the 3-Mbps data stream so that it does not overwhelm the 64-kbps serial link. This is done by storing the incoming data in on-board buffers and sending it over the serial link at a rate that the serial link can accommodate. This buffering can be achieved only for short bursts of data that do not overwhelm the bridge's buffering capability.

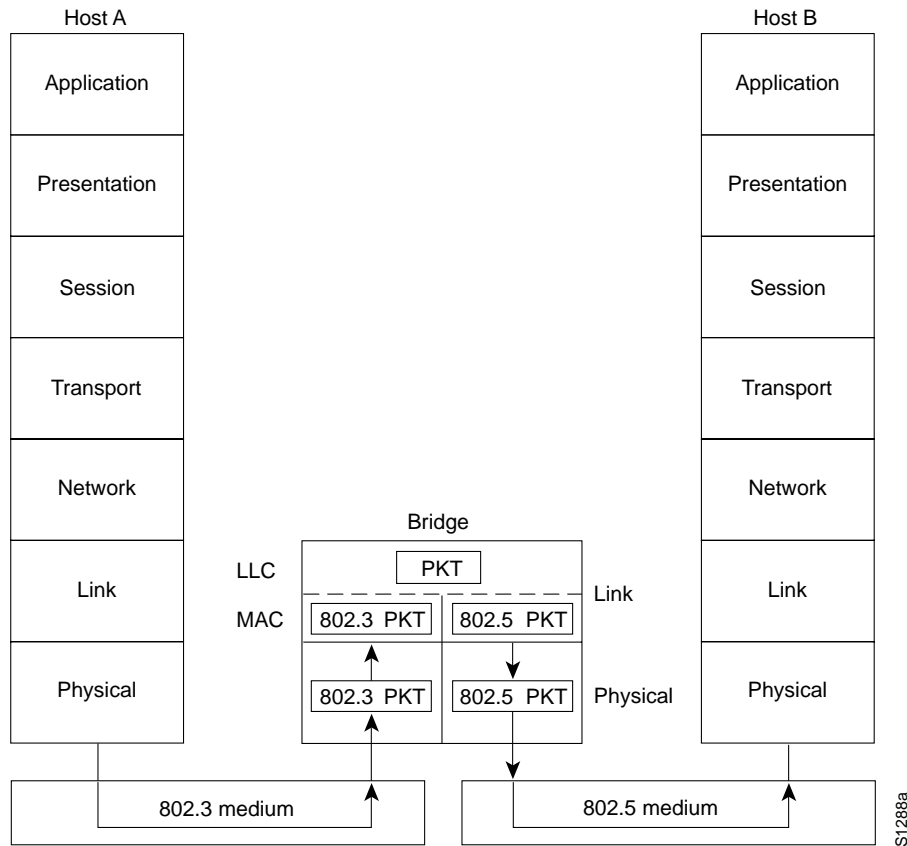
The Institute of Electrical and Electronic Engineers (IEEE) differentiates the OSI link layer into two separate sublayers: the *Media Access Control* (MAC) sublayer and the *Logical Link Control* (LLC) sublayer. The MAC sublayer permits and orchestrates media access, such as contention and token passing, while the LLC sublayer deals with framing, flow control, error control, and MAC-sublayer addressing.

Some bridges are *MAC-layer bridges*, which bridge between homogeneous networks (for example, IEEE 802.3 and IEEE 802.3), while other bridges can translate between different link-layer protocols (for example, IEEE 802.3 and IEEE 802.5). The basic mechanics of such a translation are shown in Figure 4-2.

Figure 4-2 illustrates an IEEE 802.3 host (Host A) formulating a packet that contains application information and encapsulating the packet in an IEEE 802.3-compatible frame for transit over the IEEE 802.3 medium to the bridge. At the bridge, the frame is stripped of its IEEE 802.3 header at the MAC sublayer of the link layer and is subsequently passed up to the LLC sublayer for further processing. After this processing, the packet is passed back down to an IEEE 802.5 implementation, which encapsulates the packet in an IEEE 802.5 header for transmission on the IEEE 802.5 network to the IEEE 802.5 host (Host B).

A bridge's translation between networks of different types is never perfect because one network likely will support certain frame fields and protocol functions not supported by the other network.

Figure 4-2 A MAC-layer bridge connects the IEEE 802.3 and IEEE 802.5 networks.



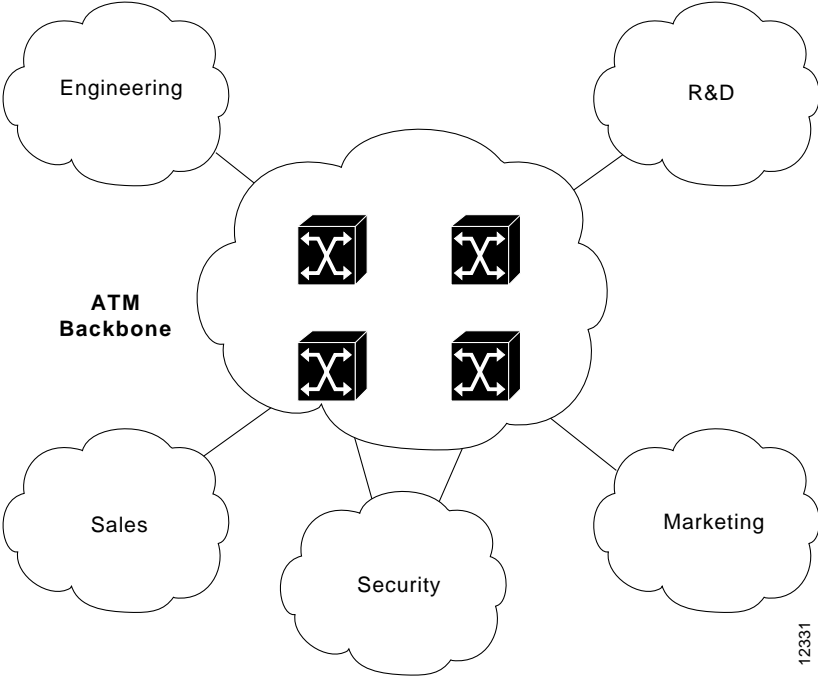
Types of Switches

Switches are data link layer devices that, like bridges, enable multiple physical LAN segments to be interconnected into a single larger network. Similar to bridges, switches forward and flood traffic based on MAC addresses. Because switching is performed in hardware instead of in software, however, it is significantly faster. Switches use either store-and-forward switching or cut-through switching when forwarding traffic. Many types of switches exist, including ATM switches, LAN switches, and various types of WAN switches.

ATM Switch

Asynchronous Transfer Mode (ATM) switches provide high-speed switching and scalable bandwidths in the workgroup, the enterprise network backbone, and the wide area. ATM switches support voice, video, and data applications and are designed to switch fixed-size information units called *cells*, which are used in ATM communications. Figure 4-3 illustrates an enterprise network comprised of multiple LANs interconnected across an ATM backbone.

Figure 4-3 Multi-LAN networks can use an ATM-based backbone when switching cells.



LAN Switch

LAN switches are used to interconnect multiple LAN segments. LAN switching provides dedicated, collision-free communication between network devices, with support for multiple simultaneous conversations. LAN switches are designed to switch data frames at high speeds. Figure 4-4 illustrates a simple network in which a LAN switch interconnects a 10-Mbps and a 100-Mbps Ethernet LAN.

Figure 4-4 A LAN switch can link 10-Mbps and 100-Mbps Ethernet segments.

