CHAPTER 3

Bridging Basics

Background

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 emerged as important. *Transparent bridging* is found primarily in Ethernet environments. *Source-route bridging* is found primarily in Token Ring environments. *Translational bridging* provides translation between the formats and transit principles of different media types (usually Ethernet and Token Ring). *Source-route transparent bridging* combines the algorithms of transparent bridging and source-route bridging to allow communication in mixed Ethernet/Token Ring environments.

The diminishing price and the recent inclusion of bridging capability in many routers has taken substantial market share away from pure bridges. Those bridges that have survived include features such as sophisticated filtering, pseudo-intelligent path selection, and high throughput rates. Although an intense debate about the benefits of bridging versus routing raged in the late 1980s, most people now agree that each has its place and that both are often necessary in any comprehensive internetworking scheme.

Internetworking Device Comparison

Internetworking devices offer communication between local area network (LAN) segments. There are four primary types of internetworking devices: *repeaters*, *bridges*, *routers*, and *gateways*. These devices can be differentiated very generally by the *Open System Interconnection* (OSI) layer at which they establish the LAN-to-LAN connection. Repeaters connect LANs at OSI Layer 1; bridges connect LANs at Layer 2; routers connect LANs at Layer 3; and gateways connect LANs at Layers 4 through 7. Each device offers the functionality found at its layer(s) of connection and uses the functionality of all lower layers. This idea is portrayed graphically in Figure 3-1.

Figure 3-1 Internetworking Product Functionality

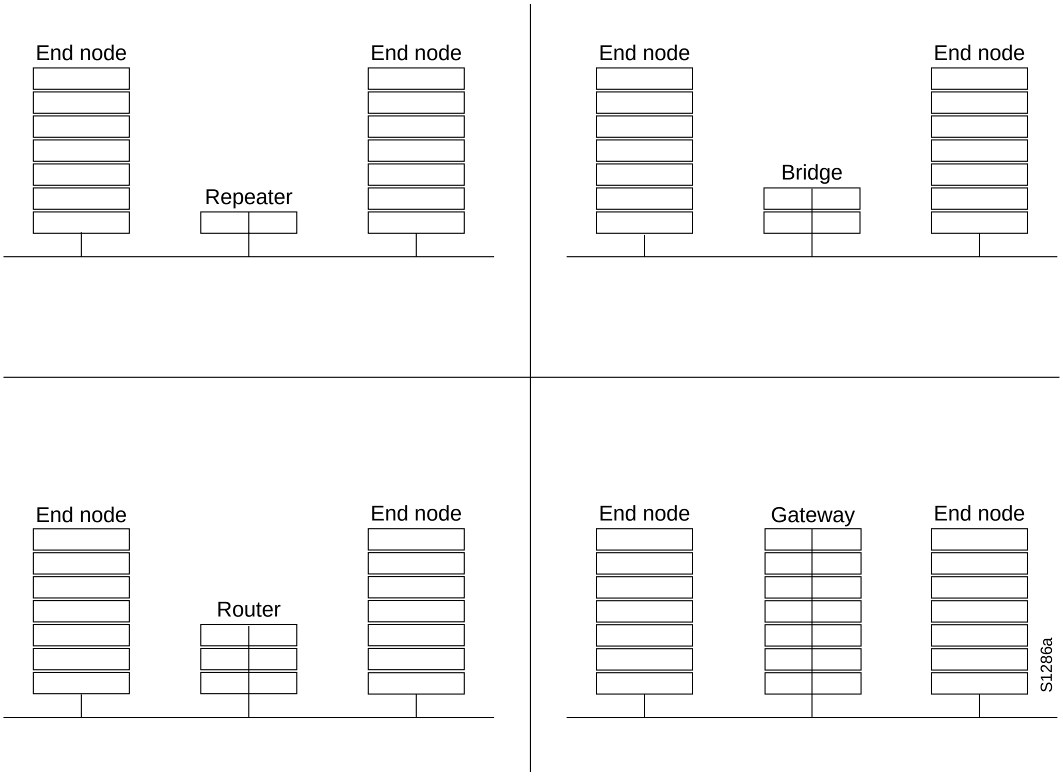

Technology Basics

Bridging occurs 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 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 (for example, source-route bridging), the entire path to the destination is contained in each frame. In other cases (for example, transparent bridging), frames are forwarded one hop at a time toward the destination. For more information on source-route bridging and transparent bridging, see Chapter 30, “Source-Route Bridging,” and Chapter 29, “Transparent Bridging.”

Upper-layer protocol transparency is a primary advantage of bridging. Because bridges 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. For example, a bridge can be programmed to reject (not forward) all frames sourced from a particular network. Since link-layer information often includes a reference to an upper-layer protocol, bridges can usually filter on this parameter. Further, filters can be helpful in dealing with unnecessary broadcast and multicast packets.

By dividing large networks into self-contained units, bridges provide several advantages. First, because only some percentage of traffic is forwarded, the bridge diminishes the traffic experienced by devices on all connected segments. Second, the bridge acts as a firewall for some potentially damaging network errors. Third, bridges allow for communication between a larger number of devices than would be supported on any single LAN connected to the bridge. Fourth, bridges extend the effective length of a LAN, permitting attachment of distant stations that were not previously connected.

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. These two configurations are shown in Figure 3-2.

Figure 3-2 Local and Remote Bridging

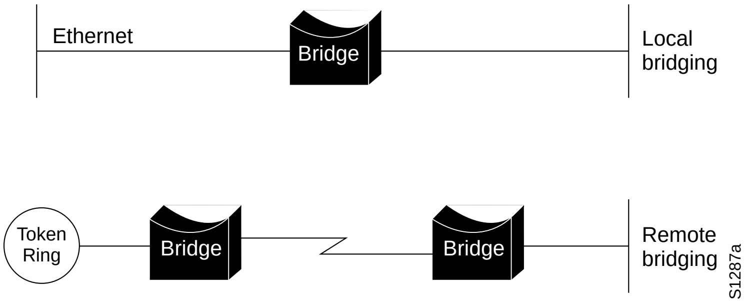

Remote bridging presents several unique internetworking challenges. One of these is the difference between LAN and wide area network (WAN) speeds. Although several fast WAN technologies are now establishing a presence in geographically dispersed internetworks, LAN speeds are often an order of magnitude faster than WAN speeds. Vastly different LAN and WAN speeds sometimes prevent users from running delay-sensitive LAN applications over the WAN.

Remote bridges cannot improve WAN speeds, but can compensate for speed discrepancies through sufficient buffering capability. If a LAN device capable of a 3-Mbps transmission rate wishes 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 the serial link can accommodate. This 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) has divided 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 (for example, contention, token passing, or others), while the LLC sublayer is concerned with framing, flow control, error control, and MAC-sublayer addressing.

Some bridges are *MAC-layer bridges*. These devices bridge between homogeneous networks (for example, IEEE 802.3 and IEEE 802.3). 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 3-3.

Figure 3-3 IEEE 802.3/IEEE 802.5 Bridging

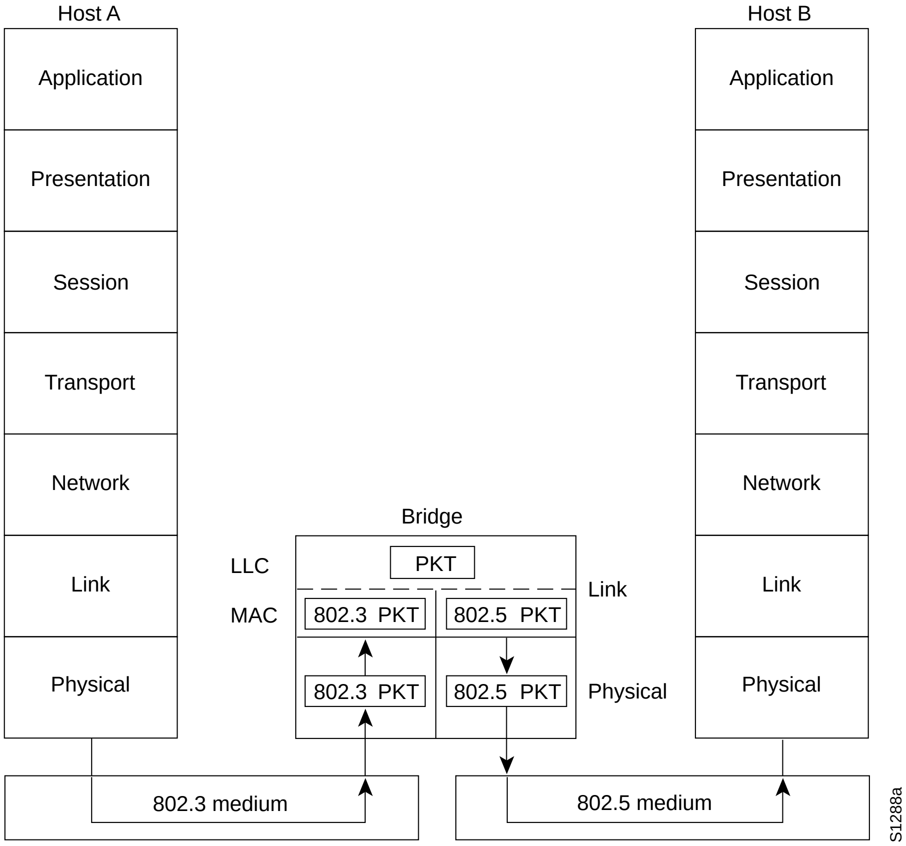

In the figure, the IEEE 802.3 host (Host A) formulates a packet containing application information and encapsulates the packet in an IEEE 802.3-compatible frame for transit over the IEEE 802.3 medium to the bridge. At the bridge, the frame is stripped of its IEEE 802.3 header at the MAC sublayer of the 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 it is likely that one network will support certain frame fields and protocol functions not supported by the other network. Many bridging translation issues are discussed in more detail in Chapter 31, "Mixed-Media Bridging."