

Token Ring/IEEE 802.5

Background

The Token Ring network was originally developed by IBM in the 1970s. It is still IBM's primary local-area network (LAN) technology, and is second only to Ethernet/IEEE 802.3 in general LAN popularity. The IEEE 802.5 specification is almost identical to, and completely compatible with, IBM's Token Ring network. In fact, the IEEE 802.5 specification was modeled after IBM Token Ring, and continues to shadow IBM's Token Ring development. The term *Token Ring* is generally used to refer to both IBM's Token Ring network and IEEE 802.5 networks.

Token Ring/IEEE 802.5 Comparison

Token Ring and IEEE 802.5 networks are basically quite compatible, although the specifications differ in relatively minor ways. IBM's Token Ring network specifies a star, with all end stations attached to a device called a *multistation access unit* (MSAU), whereas IEEE 802.5 does not specify a topology (although virtually all IEEE 802.5 implementations also are based on a star). Other differences exist, including media type (IEEE 802.5 does not specify a media type, while IBM Token Ring networks use twisted-pair wire) and routing information field size. Figure 6-1 summarizes IBM Token Ring network and IEEE 802.5 specifications.

Figure 6-1 IBM Token Ring Network/IEEE 802.5 Comparison

	IBM Token Ring network	IEEE 802.5
Data rates	4.16 Mbps	4.16 Mbps
Stations/segment	260 (shielded twisted pair) 72 (unshielded twisted pair)	250
Topology	Star	Not specified
Media	Twisted pair	Not specified
Signaling	Baseband	Baseband
Access method	Token passing	Token passing
Encoding	Differential Manchester	Differential Manchester

S1292a

Token Passing

Token Ring and IEEE 802.5 are the primary examples of token-passing networks. Token-passing networks move a small frame, called a *token*, around the network. Possession of the token grants the right to transmit. If a node receiving the token has no information to send, it simply passes the token to the next end station. Each station can hold the token for a maximum period of time.

If a station possessing the token does have information to transmit, it seizes the token, alters one bit of the token (which turns the token into a start-of-frame sequence), appends the information it wishes to transmit, and finally sends this information to the next station on the ring. While the information frame is circling the ring, there is no token on the network (unless the ring supports *early token release*), so other stations wishing to transmit must wait. Therefore, collisions cannot occur in Token Ring networks. If early token release is supported, a new token can be released when frame transmission is completed.

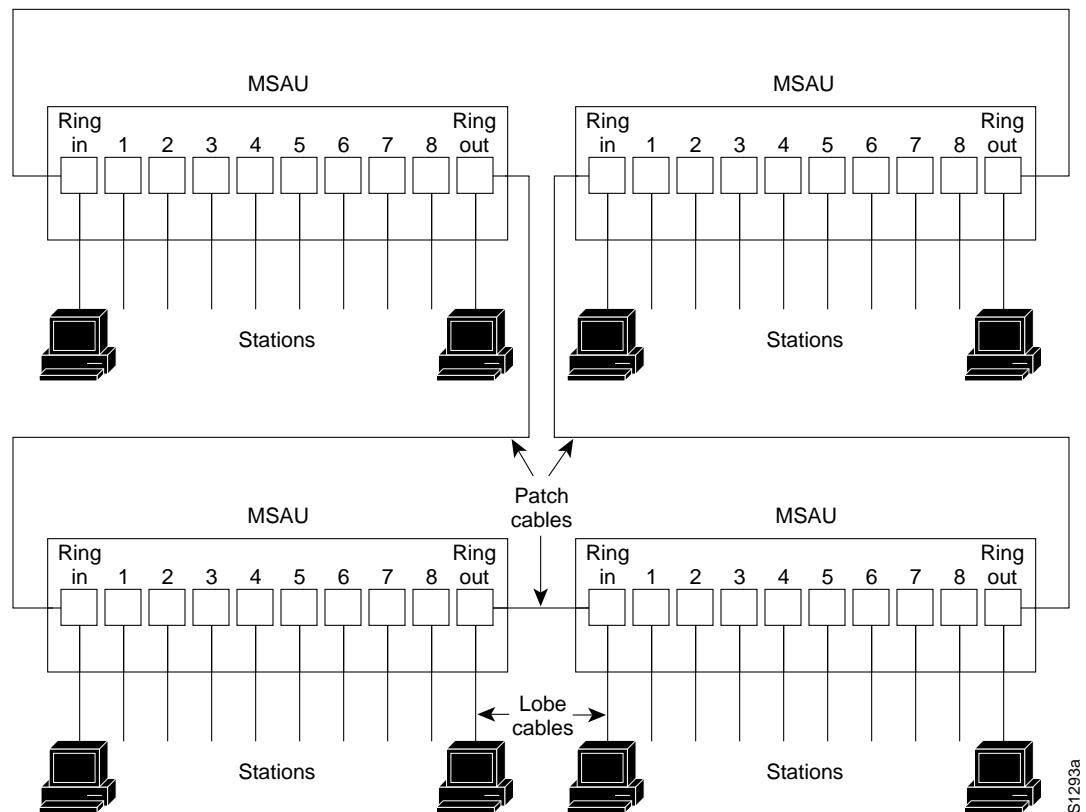
The information frame circulates the ring until it reaches the intended destination station, which copies the information for further processing. The information frame continues to circle the ring and is finally removed when it reaches the sending station. The sending station can check the returning frame to see whether the frame was seen and subsequently copied by the destination.

Unlike CSMA/CD networks (such as Ethernet), token-passing networks are deterministic. In other words, it is possible to calculate the maximum time that will pass before any end station will be able to transmit. This feature and several reliability features, which are discussed in the section “Fault Management Mechanisms” later in this chapter, make Token Ring networks ideal for applications where delay must be predictable and robust network operation is important. Factory automation environments are examples of such applications.

Physical Connections

IBM Token Ring network stations are directly connected to MSAUs, which can be wired together to form one large ring (as shown in Figure 6-2). Patch cables connect MSAUs to adjacent MSAUs. Lobe cables connect MSAUs to stations. MSAUs include bypass relays for removing stations from the ring.

Figure 6-2 IBM Token Ring Network Physical Connections



S1293a

Priority System

Token Ring networks use a sophisticated priority system that permits certain user-designated, high-priority stations to use the network more frequently. Token Ring frames have two fields that control priority: the *priority* field and the *reservation* field.

Only stations with a priority equal to or higher than the priority value contained in a token can seize that token. Once the token is seized and changed to an information frame, only stations with a priority value higher than that of the transmitting station can reserve the token for the next pass around the network. When the next token is generated, it includes the higher priority of the reserving station. Stations that raise a token's priority level must reinstate the previous priority after their transmission is complete.

Fault Management Mechanisms

Token Ring networks employ several mechanisms for detecting and compensating for network faults. For example, one station in the Token Ring network is selected to be the *active monitor*. This station, which can potentially be any station on the network, acts as a centralized source of timing information for other ring stations and performs a variety of ring maintenance functions. One of these functions is the removal of continuously circulating frames from the ring. When a sending device fails, its frame may continue to circle the ring. This can prevent other stations from transmitting their own frames and essentially lock up the network. The active monitor can detect such frames, remove them from the ring, and generate a new token.

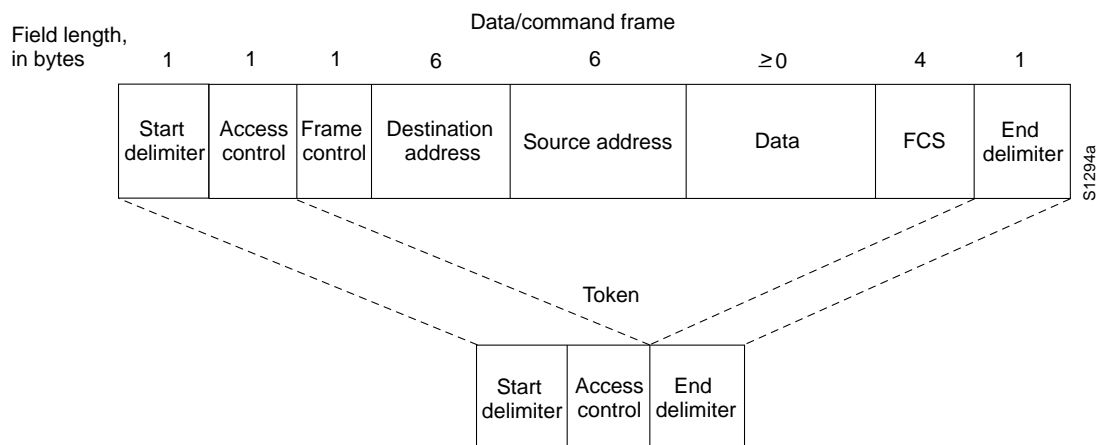
The IBM Token Ring network's star topology also contributes to overall network reliability. Since all information in a Token Ring network is seen by active MSAUs, these devices can be programmed to check for problems and selectively remove stations from the ring if necessary.

A Token Ring algorithm called *beaconing* detects and tries to repair certain network faults. Whenever a station detects a serious problem with the network (such as a cable break), it sends a beacon frame. The beacon frame defines a failure domain, which includes the station reporting the failure, its *nearest active upstream neighbor* (NAUN), and everything in between. Beaconing initiates a process called *autoreconfiguration*, where nodes within the failure domain automatically perform diagnostics in an attempt to reconfigure the network around the failed areas. Physically, the MSAU can accomplish this through electrical reconfiguration.

Frame Format

Token Ring networks define two frame types: tokens and data/command frames. Both formats are shown in Figure 6-3.

Figure 6-3 IEEE 802.5/Token Ring Frame Formats



Tokens

Tokens are 3 bytes in length and consist of a start delimiter, an access control byte, and an end delimiter.

The *start delimiter* serves to alert each station to the arrival of a token (or data/command frame). This field includes signals that distinguish the byte from the rest of the frame by violating the encoding scheme used elsewhere in the frame.

The *access control byte* contains the priority and reservation fields, as well as a token bit (used to differentiate a token from a data/command frame) and a monitor bit (used by the active monitor to determine whether a frame is circling the ring endlessly).

Finally, the *end delimiter* signals the end of the token or data/command frame. It also contains bits to indicate a damaged frame and a frame that is the last in a logical sequence.

Data/Command Frames

Data/command frames vary in size, depending on the size of the information field. Data frames carry information for upper-layer protocols; command frames contain control information and have no data for upper-layer protocols.

In data/command frames, a *frame control* byte follows the access control byte. The frame control byte indicates whether the frame contains data or control information. In control frames, this byte specifies the type of control information.

Following the frame control byte are the two *address* fields, which identify the destination and source stations. As with IEEE 802.3, addresses are 6 bytes in length.

The *data* field follows the address fields. The length of this field is limited by the ring token holding time, which defines the maximum time a station may hold the token.

Following the data field is the *frame check sequence* (FCS) field. This field is filled by the source station with a calculated value dependent on the frame contents. The destination station recalculates the value to determine whether the frame may have been damaged in transit. If so, the frame is discarded.

As with the token, the *end delimiter* completes the data/command frame.

