

Token-Ring Networks

SOME OF THE MAIN TOPICS IN THIS CHAPTER ARE

Overview of Token-Ring Networking	2
Functions Performed by Active and Standby Monitors	7
Detecting Errors on the Ring	10
What Is Early Token Release?	11
Adding a Station to the Ring	11
The Physical Star Topology	12
Hierarchical Topologies	14
Using a Token-Bus Topology	18
Cabling Rules for Token-Ring Networks	19
Token-Ring Connectors and Media Filters	22
Cabling Scenarios	23
Monitoring Token-Ring Utilization and Errors	24
Using Network and Protocol Analyzers	27
Token-Ring Extensions to the Remote Network Monitoring MIB	27
Troubleshooting Tips	28
Choosing Between Token-Ring and Ethernet	28

Today, most local area networks are based on Ethernet technology. Older Ethernet networks are made up of multiple computers or devices attached to a shared network media, each contending for a share of the bandwidth. Modern Ethernet uses switches to reduce the collision domain, and scales to much larger LANs. What all the different Ethernet topologies and cabling schemes have in common, though, are the frame type and the CSMA/CD (Carrier Sense Multiple Access/Collision Detect) contention method used to gain access to the network medium and to transmit data.

Token-Ring networks use a different media access control mechanism and a different frame format than Ethernet. Token-Ring networking was first developed by IBM and later standardized by the IEEE 802.5 working group. In a Token-Ring network, there are no collisions and there are no back-off algorithms. Network access is not something that each workstation has to contend for. Instead, access to a Token-Ring network is done in a controlled, orderly fashion, using a frame that is passed from workstation to workstation, possession of which is required to begin transmitting data on the network. This “permission” frame is called a *token*. It passes through the network in a logical ring topology; hence, the name Token-Ring is used for this technology.

◀◀ If you want to learn more about Ethernet in its many formats and topologies, see Chapter 14, “Ethernet: The Universal Standard.”

Because the token regularly passes around in a ring formation, it is possible to calculate the longest amount of time it will take until a station can receive the token and begin transmitting data on the network. Thus, Token-Ring networks are called *deterministic* because the worst-case-scenario access time can be calculated. In some environments, such as factory automation, where timing is important, Token-Ring networks can be a good solution.

Overview of Token-Ring Networking

In Figure 1, you can see that the logical layout of a Token-Ring network is a ring. Each computer in the ring connects to two other computers, affectionately referred to as *neighbors*. The upstream neighbor is the computer from which a station will receive the token or data frame. The downstream neighbor is the one to whom the station will pass the frame.

If the workstations were actually wired in this manner, however, it would be quite a problem to add new stations or move them around. Doing so would break the chain of cables, and the network would grind to a halt. To prevent this from happening, each device on the ring is actually connected by two sets of wires to a hublike device called a media access unit (MAU) or a multistation access unit (MSAU). The MSAU looks similar to an Ethernet hub. It is usually tucked away in a wiring closet like a hub and, in some versions, uses RJ-45 connectors and twisted-pair wiring just like Ethernet networks do.

MAUs and MSAUs come in various styles, with capacities generally ranging from 8 to 16 ports, though you can find more expensive MAUs that support more workstations than this. More sophisticated MAUs have an additional feature: remote management. These *controlled access units (CAUs)* allow the network administrator to remotely manage stations connected to the ring. Workstations do not directly connect to the CAU, however; instead, a *lobe attachment module (LAM)* is used to connect workstations to the CAU. Generally, LAMs support 20 individual workstations, and you can connect up to four LAMs to a CAU.

◀◀ At first glance it might seem that Token Ring has a lot in common with ARCnet. However, Token Ring provides a network technology that scales well beyond the 255-node limit of an ARCnet network. In addition, the frame format and other specifications of ARCnet differ from Token-Ring. For more information about ARCnet, see Chapter 13, “The Oldest LAN Protocol Is Still Kicking: ARCnet.”




Figure 1 Token-Ring networks are wired in a star topology to form a logical ring.

Unlike an Ethernet hub, however, the MAU does not repeat incoming traffic on all other ports. Instead, using electrical relays, the MAU connects the receiving wires from one station to the transmitting wires of the previous station in the ring.

A frame called a *token* is continuously passed from one station to another around the ring. When a workstation receives the token frame, it retransmits it to the next station in the ring. When a station needs to transmit data to another station on the ring, it “seizes” the token. That is, instead of immediately retransmitting the token frame, the workstation modifies a certain bit in the frame, appends the data it wants to transmit and the destination address, and transmits the new, larger frame to its neighbor instead.

This data frame is passed from station to station until it reaches its destination or returns to the sender. Each station in the ring checks the destination address to determine whether it is the destination of the frame. If the station recognizes its address as the destination, it sends the frame up the protocol stack for further processing. It also retransmits the data frame to the next workstation on the ring, after changing another bit that indicates that the data was received successfully by the destination station on the ring.

When the data frame circles the ring and gets back to its original sender, it is stripped from the ring. The station that originated the data frame checks to see whether the destination workstation has acknowledged receiving the data and, if it has, starts the token process again by sending a new token frame to its neighbor.

This simple process describes how data is exchanged between computers that are attached to the network. For a complete understanding of how the technology works, a few other questions need to be answered:

- What controls this orderly process? Where does the first token come from when the network is started?
- How does a station join a ring? Won't this disrupt the physical connection made in the MAU or MSAU?
- What does a token frame look like? A data frame? How are they different from one another?
- Token-Ring is described as "self-healing." How are errors detected and isolated to keep the LAN up and running?

Before answering these questions, it will be helpful to first look at the format of the frames used on Token-Ring networks to understand how these problems can be addressed using the types of information that can be sent between stations.

Token-Ring Frames

Three kinds of frames are used in Token-Ring networks. One is the token frame that is passed around the idle network. This frame is only three bytes long and is seized by a computer when it wants to transmit data on the network. The other two frames are longer and contain either data or commands used to manage the network. Token-Ring adapters use the MAC frame when communicating with each other. When sending data that is destined for a higher-level protocol, such as IP or IPX packets, the LLC frame format is used.

The Token Frame

The token frame is a three-byte frame that circles the ring until a station decides that it needs to transmit data. Figure 2 shows the token frame format. Although the start and end delimiter bytes are used to indicate the beginning and end of the frame, only one byte is really worthy of discussion: the access control byte.




Figure 2 The token frame is only three bytes long.

The access control field contains several important values that are all stored in only one byte. These values represent the priority bits, the token bit, the monitor bit, and the reservation bits.

The Token Bit

The *token bit* is obviously used to indicate that the frame is a token frame or another kind of frame. If it is a token frame, the network media is available and the workstation receiving it can transmit data if it wants to when it receives this frame. In this mode, the token bit is set to a value of zero. When a station wants to transmit data, however, it copies this byte to a buffer, changes the token bit to a value of one, adds addressing information to create a larger frame, and then transmits the larger frame to its downstream neighbor. The workstation can continue transmitting until it is finished, or until the *token-holding timer* expires. When the computer has transmitted the last frame that it needs to for the time being and the frame has circled the ring and arrived back at its starting point, a new token is generated—with the token bit set to zero—and released to the ring. The next station that receives the token frame can transmit.

The Monitor Bit

The *monitor bit* is used by a workstation that acts as the *active monitor* for the ring to determine when a frame has been around the ring more than once. The active monitor performs other important functions for the ring, which are discussed later in this chapter. It is important to understand that, although any workstation that generates a token frame sets the value of this bit to zero, only the active monitor can change it to a one, which it does when it sees a data frame for the first time. If the frame continues to circle the ring without ever reaching its destination, the active monitor can detect this, remove the frame from the ring, and start the process over again. The monitor bit, therefore, is used to prevent a data frame from circling the ring endlessly. This might happen if one computer tries to send a frame to another computer that has been taken out of the ring, for example.

So, although the token bit is used by individual stations to indicate that they are transmitting a data frame, the monitor bit is used by the active monitor to set the network straight if the destination computer doesn't exist on the LAN or has been taken offline for some reason.

Priority and Reservation Bits

Token-Ring networks have the capability to enable individual workstations to take priority over others when it comes to claiming a token and beginning a transmission. In some applications, such as real-time factory-floor automation environments, this can be a major advantage over Ethernet's contention method for network access. Using priority bits and reservation bits does this.

The priority bits are the first three bits of the access control byte. Counting in binary, this gives eight possible levels (000–111) of priority values.

The reservation bits are also used to specify a priority value between 000 and 111, in binary. A station that wants to reserve the token for its use will set these bits to the priority it expects to use, provided it is a higher value than the current priority bits represent. When the token is regenerated by a station, it will set the priority bits to this value so that the reserving station will be capable of seizing the token when it passes by again on its way around the ring. The reservation doesn't guarantee that the station will get the token on the next pass, however, because another workstation with a higher priority might intervene.

LLC and MAC Frames

In Figure 3, you can see that the frame type used to send a command or transmit data is much longer than the three-byte token frame. Indeed, because Token-Ring networks do not limit the overall length of the frame to a small value (such as Ethernet frames do), this frame has the potential to be quite large.

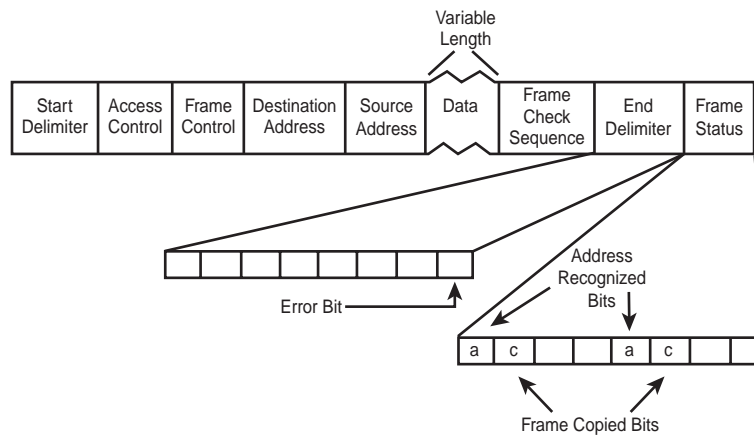


Figure 3 Data and command frames can become quite large.

The access control field for this frame type is copied from the token that the workstation has just seized, but the token bit's value is changed to one to distinguish it from a token frame. The source and destination addresses are each six bytes long. The frame check sequence field is four bytes long and is used to store a cyclic redundancy check (CRC) value that can be used to validate the integrity of the frame as it travels from one workstation to another.

All other fields, except for the data field, are one byte long.

The end delimiter field is one bit (the last bit) that serves an important function. When a station determines that the frame has been damaged or corrupted in transit, it sets this last bit—the error bit—to a one.

The last field is called the frame status field. This is where the station that originates a data frame checks to see whether it was received properly by the destination computer. Remember that the frame will circle the ring until it arrives back at its point of origin.

When a station receives a frame that contains its address in the destination address field, it sets the address-recognized bits to one. If it is capable of copying the frame to a buffer successfully, it also sets the frame-copied bits to one.

The originating station can interpret these bits when the frame circles around to it as described here:

- If the address-recognized bits and the frame-copied bits are all zero, the destination computer never received the frame. This means that the computer does not exist or is not active on the ring at this time.
- If the address-recognized fields are set to ones but the frame-copied fields are still set to zeros, the destination computer was unable, for some reason, to copy the frame into its internal buffer for further processing.
- If both the address-recognized bits and the frame-copied bits are set to ones, the frame was successfully received and copied by the destination computer.

Note

If the frame arrives back at the originating station with the error bit set to one, this does not necessarily mean that the destination did not successfully receive the information. If the frame-copied bits are set to one, the target workstation *did* receive the frame and successfully copied the data. Thus, if the frame-copied bits are set to one and the error flag is set to one, the sending workstation does not consider this an error. It indicates that the frame was corrupted *after* it was received by the target workstation, so there is no need to retransmit the frame.

The frame control byte is used to define the type of frame. If the first two bits of this frame are both zeros, it is a media access control (MAC) frame. If they read 01, it is a logical link control (LLC) frame.

- If the frame is carrying data—it's an LLC frame—the destination computer will pass it up to the next layer in the protocol stack for further processing. The LLC header information, such as the DSAP (destination service access point) and SSAP (source service access point) bits, are found at the beginning of the data field.
- If it is a command frame—it's a MAC frame—the data field will contain a 16-bit field called a vector identifier. The types of commands that can be stored in the vector identifier field and their values appear in Table 1.

Table 1 Access Control Frame Functions

Value	Description
00000000	Duplicate address test
00000010	Beacon frame
00000011	Claim token frame
00000100	Purge ring
00000101	Active monitor present
00000110	Standby monitor present

Each of these commands serves a specific purpose that is used to either notify other stations that everything is functioning well—the active and standby monitor commands—or perform maintenance functions when something appears to have gone awry.

Functions Performed by Active and Standby Monitors

Not every station on a Token-Ring network is equal to the others; some computers serve specific roles in the network. For example, when the network is started, something has to send the first token frame through the ring.

A station called the *active monitor* does this. The active monitor is the station that manages the ring. Although any workstation can be the active monitor, there can be only one active monitor in the ring, and all other stations attached to the ring act as *standby monitors*. If the active monitor fails or is taken off the network, one of the standby monitors automatically takes over the role and becomes the new active monitor. These are the major functions of the active monitor:

- Initializing the first token when the network is powered up
- Initiating *neighbor notification*, whereby all stations in the ring learn the names (addresses) of their upstream and downstream neighbors
- Providing a master clock that all other stations use for timing purposes
- Providing ring delay
- Detecting lost tokens and frames

The clocking function provided by the active monitor is important in a deterministic network such as Token-Ring. All other stations synchronize their clocks to the clocking signal provided by the active monitor. Stations do not have to listen for a set of *preamble bits* as they do in an Ethernet network. The network is precisely clocked so that a computer knows when to expect the token or a data frame. This timing mechanism enables a computer to detect that a problem exists on the network if a token does not arrive on time, and causes error-detection procedures to begin.

Ring Polling (Neighbor Notification)

Each station on the ring must know the address of the station immediately upstream from it. This station is called its *Nearest Active Upstream Neighbor*, or *NAUN* for short. Knowing the address of its NAUN enables a station to detect errors and begin a correction process when it notices that its neighbor has failed in some way.

The active monitor begins a process called neighbor notification when the network is first started. This process also is called *ring polling*. Two types of frames are used in this process: active monitor present and standby monitor present.

To begin the process, the active monitor broadcasts an active monitor present control frame around the network. (Remember that the address-recognized and frame-copied bits in the frame status field are initially set to zeros.) Because this frame is a broadcast frame, it will be recognized by any station that receives it, which will be the station immediately downstream from the workstation that is acting as the active monitor.

The station that receives the broadcast frame will reset its timer, using the value supplied by the active monitor. It will copy the source address—its upstream neighbor's address—into memory for future reference. The address-recognized and frame-copied bits in the frame status field are set to ones, and the frame is sent on its way to the next workstation. Other stations on the ring will not copy this frame because the appropriate bits have already been set. The frame circles the ring until it arrives back at the active monitor, which strips it from the network.

When the second station in the ring retransmits the active monitor present frame with its bits set, it also sets a 20-millisecond timer. This is done to give the active monitor present frame enough time to move around the ring. When this timer expires, it will send out a frame similar to the one it just received. However, this frame is a *standby monitor present* frame. It serves the same type of function as the previous frame to the next workstation downstream because its address-recognized and frame-copied bits are initialized to zero.

When the next station downstream receives the standby monitor present frame, it recognizes it as a broadcast frame and, because the appropriate bits are set to zero, it realizes that it has just received the frame from its immediate upstream neighbor. It stores this neighbor's address in memory for future use, sets the address-recognized and frame-copied bits to ones, and retransmits the frame to the next workstation. It also sets a 20-millisecond timer and, when the timer expires, repeats the process by sending its own standby monitor present frame to the next workstation downstream.

Finally, when the active monitor receives a standby monitor present frame with the address-recognized and frame-copied bits set to zeros, it knows that the neighbor notification process has completed successfully.

So that all members of the ring will be kept updated about potential new members in the ring or other changes, the active monitor sends out an *active monitor present* frame every 7 seconds. If any station in the ring does not receive the frame at least every 15 seconds, it initiates a process called *monitor contention* so that a new active monitor can be designated for the ring.

How the Network Repairs Itself When the Active Monitor Fails

Each station other than the active monitor is designated to be a standby monitor. These stations expect that every 2.6 seconds, either a token frame or a data frame will pass their way. They also expect that an active monitor present frame will arrive at least every 15 seconds. If any of these events fails to occur, it is assumed that the active monitor has been taken offline or has failed in some manner. To remedy the situation, the first standby monitor that notices this situation initiates a process called *monitor contention*, in which a new active monitor is designated to manage the network.

All stations except for the current active monitor will participate in the “election” of a new monitor. If the current active monitor is still present on the network, it assumes that it is not functioning correctly because another station has begun the monitor contention process.

The station that begins the process puts itself into Claim Token Transmit mode and sends out frames called *claim token frames*. These frames contain the station’s own DLC address. The next station downstream receives this frame and compares the address to its own. If its own address is of a higher numerical value, it substitutes its own address for the one put into the frame by its upstream neighbor, retransmits the frame, and puts itself into Claim Token Transmit mode. Otherwise, it puts itself into the Claim Token Repeat mode and forwards the frame to the next station so that it can make the comparison.

After a short time, only one station will remain in Claim Token Transmit mode. When a station receives its own Claim Token Transmit mode back from circling the ring three consecutive times, it will decide that it has won the election and begin operating as the active monitor for the ring.

To begin its new duties, the new active monitor will perform the following tasks:

- Take over the clocking functions for the ring
- Send out a *ring purge frame*
- Begin ring polling (neighbor notification)
- Transmit a free token so that the network can resume normal operation

The ring purge sets the ring back to an initial state, causing stations that receive the frame to reset their timer clocks and to abort any other task that they were in the process of doing. When the frame has circled the ring, the active monitor can send out the new token and stations can resume normal processing.

The Active Monitor Watches for Problems

The active monitor guards the overall health of the network by constantly monitoring activity on the ring. It expects its own active monitor present frame to return, along with the standby monitor present frames generated by other stations (received every seven seconds). It watches to determine whether a station has raised the priority of a token but failed to lower it, thus making the token unusable by other stations. All in all, the active monitor expects to see a good token every 10 milliseconds and will take steps to repair the network if it does not arrive on time.

When a token frame or data frame passes by the active monitor, it sets a timer. If the timer expires before it sees another frame, it assumes that the frame has been lost, purges the ring, and initiates a new token.

The monitor bit is set only by the active monitor when it first sees a frame. If it sees the same frame again—which it can determine because no other station can set this bit to a value of 1—it will assume that the frame is circling the ring endlessly and will strip it from the ring.

When the active monitor detects errors in a frame based on the CRC value, it purges the ring and, if the purged frame returns to the active monitor undamaged, initiates a new token so that processing can resume.

Because Token-Ring networks are so small (compared to other types of networks), the token can travel around from station to station quite quickly. Thus, it is possible for a station to begin stripping bits off a frame before it has finished transmitting the entire frame. To prevent this from happening, a latency factor is introduced by the active monitor. The latency buffer is a minimum of 24 bits and is used to introduce this delay. This value is used because the token frame itself is 24 bits in length (three bytes).

Detecting Errors on the Ring

One of the nice things about Token-Ring networks is that they are somewhat capable of detecting errors—such as a malfunctioning upstream neighbor—and in many cases can make corrections that will keep the network functioning. When a station detects that it has a problem with its upstream neighbor, it starts a process called *beaconing* to inform the other downstream stations so that the process of isolating the error can begin.

Transmitter Failures

Each station has already learned the address of its upstream neighbor because of the ring polling process. When a station determines that it is no longer receiving frames from its NAUN, it begins to send out beaconing frames that contain its own address and the address of the neighbor that it assumes has malfunctioned. Other stations receiving these frames enter an error correction mode, and other traffic on the ring stops until the beaconing process is complete.

Eventually, if there are no other problems in the network, the beacon frames arrive back at the upstream neighbor, which recognizes that its address is present in the frame. Because another station is complaining, the station removes itself from the network ring and performs a self-test, just as it does when it joins a ring. If it does not pass its own self-test, it remains off the ring until an administrator can diagnose the problem and repair it. If it passes the test, it reinserts itself into the ring.

If the problem station stays removed from the ring, the relays in the MSAU reconfigure the wiring connections so that the station is bypassed. Thus, the station that is issuing the beaconing frames is connected to a neighbor that is one farther upstream from it than the one that was removed. If all is functioning normally, it begins to receive the beaconing frames it originated. It then stops transmitting the frames and waits for the active monitor to restore the network to normal functioning, which it does by issuing a new token.

Receiver Failures

In the situation just described, you can see how a station with a faulty network adapter that cannot transmit will be automatically removed from the ring. However, when a station no longer can hear

frames from its upstream neighbor, the problem doesn't have to be in that neighbor's adapter card. Instead, the problem might lie in this station's card, which might be malfunctioning, making it unable to receive the data properly. This type of malfunction will also be detected by other stations on the ring, and the offending station will be isolated.

When a station starts transmitting beaconing frames to let others know that it cannot hear its upstream neighbor, it sets a timer. This timer is of a sufficient value to enable the upstream neighbor to self-test and remove itself from the ring. If the timer expires, the beaconing station suspects its own receiving capability, removes itself from the ring, and performs the self-test. If it fails the test, it knows that it is the malfunctioning unit and stays disconnected from the ring.

In this case, the upstream neighbor will have passed its self-test and reconnected to the ring.

Note

Not all hardware problems can be resolved in this manner. For example, when a cable that connects two MSAUs breaks, the intervention of an administrator is required to correct the problem.

What Is Early Token Release?

Standard Token-Ring networks operate at either 4Mbps or 16Mbps. When operating at 4Mbps, a station will release a new token on the network after it has received a data frame that has finished circling the ring. When the network speed is 16Mbps, this fact introduces a significant delay factor that prevents the available bandwidth from being efficiently used—a frame circling at 16Mbps takes less time to circle the ring. On a ring that holds a large number of stations, very short frames introduce a significant amount of idle time on the ring.

Note

Although 4/16Mbps Token-Ring networks are the most widely deployed, a standard for Gigabit Token-Ring (also called High Speed Token-Ring) was standardized by the IEEE as 802.5t. However, if you go to the Web site, you'll see that not much has been happening to further keep improving the bandwidth of Token-Ring LANs. You can visit the IEEE committee home page at www.8025.org/new.html.

To increase the use of available bandwidth, stations on 16Mbps Token-Ring networks can be configured to release a token shortly after they release their data frame. Thus, a station downstream can begin to transmit data more quickly than if it had to wait for the data frame to circle the ring. This is called *early token release*, and is the default operation for Token-Ring adapters operating at 16Mbps.

Adding a Station to the Ring

Now that I have covered the mechanics of how a Token-Ring network operates, you might wonder how a new workstation joins a ring that is already functioning. The process doesn't require that the entire network be brought down and restarted. However, for a short period, the new workstation checks out its own capability to transmit and receive data, and then announces itself to the network.

Remember that the MSAU contains electrical relays that can be used to add or remove a station from the ring. When a station is first connected, the relay connects the station's transmit wires to its receive wires. Thus, at this point in time, it can only talk to itself.

The process of inserting a new member into the ring consists of the following steps:

1. The Token-Ring card runs its own self-diagnostics tests and sends a signal on its transmit wires. Because the transmit and receive wires are being connected by the MSAU, it receives its own signal back. It compares this to what it sent to determine whether it is functioning correctly.
2. If everything appears to be in working order, the adapter then applies a voltage, called a *phantom voltage*, to the transmit wires. This causes the MSAU to activate the relay that is connecting the transmit and receive wires. The relay then physically connects the station to its neighbor's transmit and receive wires, inserting it into the active ring.
3. The new node waits up to 18 seconds while it monitors incoming signals to determine whether an active monitor present frame is circling the ring. If it does not see one, it initiates the contention process so that a new active monitor will be designated for the ring.
4. Assuming that the active monitor present frame was received, the station then sends out a frame called the Duplicate Address Test frame. This frame contains its own address and circles the ring so that other stations can examine it to determine whether the address already exists on the ring. If the frame arrives back at the station with the address-recognized bits in the frame status field set to ones, it realizes that the address is already in use and removes itself from the ring until an administrator can remedy the problem.
5. If it determines that it has a unique address, the station participates in the next ring poll, learning the address of its upstream neighbor and informing its downstream neighbor of its address.

After these steps, the station is a functioning member of the ring. Other housekeeping functions might be performed, such as the learning of certain ring parameters—for example, ring number and other configurable values.

The Physical Star Topology

As discussed earlier in this chapter, Token-Ring networks are organized in a circular fashion, with one station connected to a neighbor on the left and right, and with the last workstation connecting back to the first, thus forming a ring. This is indeed the actual, *logical* formation for a Token-Ring LAN.

Although at first glance it might appear that the ring is simply a bus topology that has been joined back onto itself to create a ring, this is not the case. You can almost imagine a T-connector attaching the cable to each workstation, with each end of the connector used to attach a cable that joins the station to its two neighbors in the ring.

If this were the *physical* method that was used, however, it would defeat one of the main benefits that Token-Ring has over its rival networking technology. That is, a break anywhere on the cable would bring down the entire LAN because the token, not to mention the data frames, would no longer be capable of completing the circuit. Instead, hublike devices are used to wire the network in a star formation. These devices, which are discussed in the next few sections, are responsible for connecting workstations to their neighbors. There is no single cable to which all workstations are attached, as in a bus topology.

Multistation Access Units

Token-Ring LANs are a *logical ring*, in that communications transpire only between a station and its upstream and downstream neighbors. A workstation receives a frame on one set of wires and retransmits it on another set of wires. The transmit cables from one station carry the signal to the receiving pins on its neighboring network adapter, first passing through a device called a multistation access unit.

The *multistation access unit (MAU)* or *multistation service unit (MSAU)* appears to be very much like a hub, and indeed, these devices are used to centralize wiring. In Figure 4, you can see that Token-Ring LANs appear to be wired as a star, much like a 10BASE-T network.

The main difference between a hub and a MAU, though, is that the MAU does not provide a central location to rebroadcast a frame to every port simultaneously. Instead, the inner workings of the MAU provide a physical connection for the transmit and receive wires from one workstation to the next in the logical ring. You can think of a star-wired Token-Ring as a collapsed “ring in a box” because the ring portion of the wiring is accomplished inside the MAU, and not throughout the physical cable plant of the building.




Figure 4 Token-Ring LANs are wired in a star formation in that all stations are connected to a central wiring concentrator.

The simplest MAU devices do nothing more than provide electrical relays that connect the wires in the proper fashion. More advanced devices can actually perform other functions, such as signal regeneration. The basic function of the MAU, however, remains the same: to provide a central wiring location to connect diverse workstations on the LAN into a logical ring, physical star topology.

Connecting MAUs—Using Ring In and Ring Out Ports to Create a Larger LAN

Getting back to the logical ring concept, let's look at what happens when you need to connect more than one MAU to the network. When two or more MAUs are connected, the logical ring structure still needs to be maintained in the LAN. The token still must be circulated from one station to the next, until it arrives back at the originating station, no matter which MAU a workstation is attached to. In Figure 5, you can see that the MAUs are connected in a ring just as the workstations are.

In some cases, there are primary and secondary paths between MAUs to help reduce network downtime when a problem occurs in the connection between MAUs.

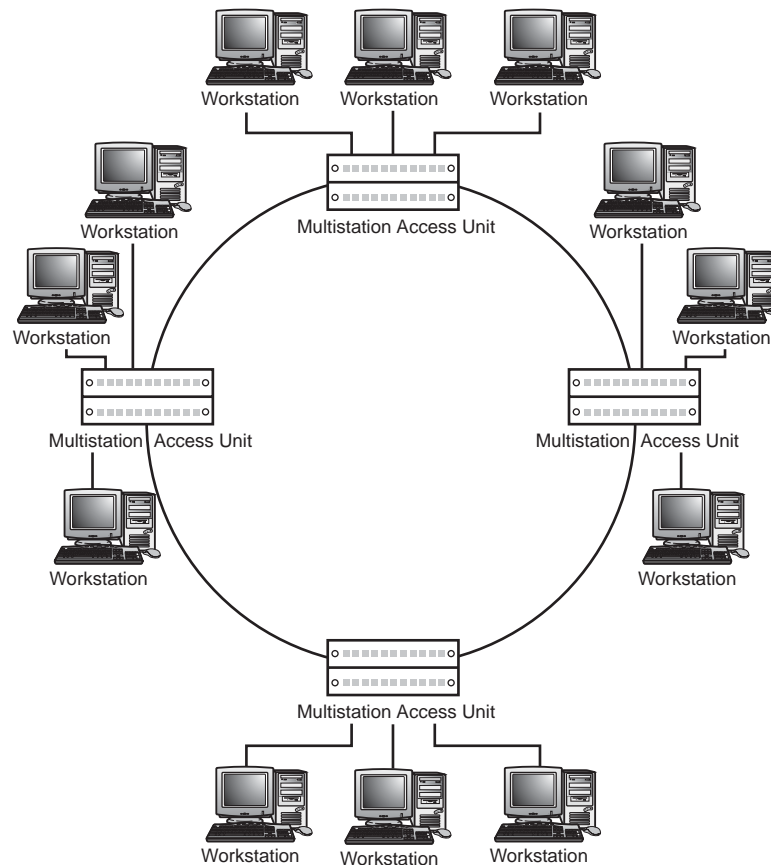


Figure 5 MAUs are joined into a ring.

Hierarchical Topologies

As with all LANs, there comes a time when you have too many workstations to connect, or when the distance between the farthestmost workstations becomes too great. When this happens, you will need to segment the network by using devices such as bridges, switches, or routers. The chapter “Bridges, Repeaters, and Hubs,” also located on the upgradingandrepairingpcs.com Web site, covers these devices in detail, mostly in the context of an Ethernet network. When it comes to bridges and switches, their operation in a Token-Ring environment is a little different.

Source Route Bridging

Bridges are used to combine smaller LANs into larger ones. Bridges are an advanced form of simple repeaters that send a signal from one cable segment out onto another without making any decisions based on information contained in the frame. Source-route bridging is used in Token-Ring networks to connect rings to or from a larger network. As the name implies, the source first determines what route the data exchange should take through the network before beginning to exchange actual data with a remote system.

The basic process can be described in a few simple steps:

1. The source workstation sends out a test frame that is used to determine whether the destination workstation is on the same local ring.
2. If the source workstation doesn't receive a response to its test, it assumes that the destination is on another ring.
3. The source workstation sends out a frame called an *All Routes Explorer (ARE)* frame. This is a broadcast type of frame that bridges typically forward to all other rings in the network.
4. As the ARE frame passes through each bridge, information is added to the frame, indicating the path it is taking. This includes the ring number and an identifier that is used to identify the bridge.
5. The ARE frame eventually makes its way to the destination computer, which sends back a directed reply. It is not necessary to send a broadcast frame for the reply because the destination computer can read the routing information in the ARE frame to see the path it has taken.
6. The originating workstation receives one or more replies from the destination computer. If there is only one route to the destination, only one reply is received. If more than one route exists, a reply is received, showing each route.

The workstation that sent out the original ARE frame makes a judgement on which path to use, and then sends directed frames that use that path for further communications with the remote system. In the simplest case, the workstation uses the path specified by the first reply received.

This type of bridging is different from *transparent bridging* used in Ethernet networks. In transparent bridges, the spanning tree algorithm is used to set up a network of bridges where there is only one possible path a frame can take to its destination, even though there might be multiple physical paths between any two stations. The spanning tree algorithm calculates the path a frame will take and changes it only when a particular path becomes unavailable, at which time the entire bridging tree formation is recalculated. Another difference between Token-Ring ring and Ethernet bridges is that Ethernet bridges make decisions based on MAC addresses, learning which addresses are on which ports. Token-Ring bridges do not need to keep a table of MAC addresses, and instead use a field called the *Routing Information Field (RIF)*. If a frame does not contain this field, it does not need to cross the router.

The ARE frame contains the RIF that is used to store information about each step the route takes. The size of this field generally limits the number of bridges through which a frame can pass. On networks built to IBM specifications, this number is usually 7. For networks based on IEEE 802.4 specifications, this number might be as large as 13. In Figure 6, you can see how the routing information is stored in an 802.5 MAC frame.

The usual fields are still there, such as the source and destination addresses. However, the first bit of the source address is called the *Routing Information Indicator (RII)* bit, and is used to indicate that this is a frame that contains routing information. The other field of importance is the Routing Information Field (RIF), which is divided into several other components. The Routing Control component consists of the following:

- **Type**—The type of routing. This can indicate an ARE frame, a specifically routed frame, or a spanning-tree explorer frame.
- **Length**—Specifies the total length of the RIF field, which can include multiple route descriptor fields.

- **Direction Bit**—The direction of the bit, forward or reverse. When the station that originated the routing request frame receives this frame back, it can reverse this bit so that the routing descriptor fields are read in a reverse manner to reach the destination station.
- **Largest Frame Size**—This field contains the size of the largest frame that can be sent over this route. In other words, it indicates the smallest frame size encountered during its journey, which becomes the common denominator for this route.

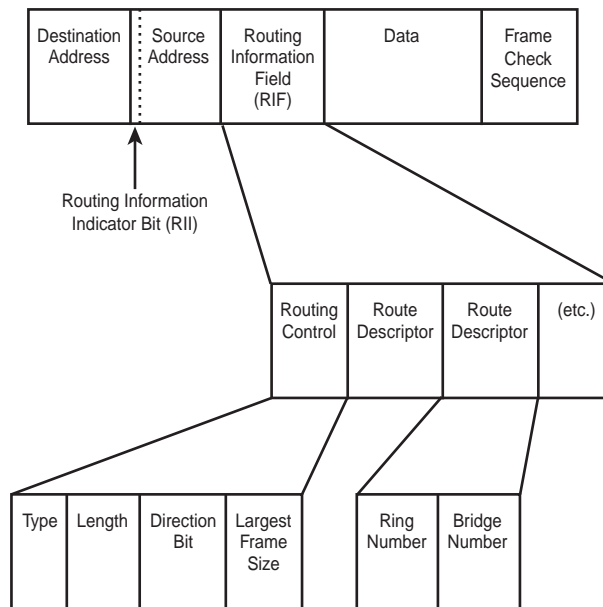


Figure 6 Routing information in an 802.5 MAC frame.

The Routing Control field can be followed by multiple Route Descriptor fields—from 7 to 13, depending on the type of network, IEEE 802.5 or IBM network—each of which contains the ring number and bridge number for each bridge or ring through which the frame passes during route discovery.

As you can see, this format makes it easy for devices between the two workstations to route the packet. The direction bit can be turned on or off by each end of the connection to cause intervening bridges to interpret the route descriptor fields in the opposite direction, depending on the direction of the communication.

When compared with transparent bridges, which are configured with parameters by network administrators that control the various paths through the network, source routing bridges are less complex. They do not have to exchange information among themselves to calculate a spanning tree, because it is the source of the data communication that decides on the path that will be taken. Because of this, however, there is a little more latency on the part of this type of bridge than the other.

Note

IEEE standards define a type of bridge called a *Source-Route Transparent Bridge*. This bridge forwards packets that contain a RIF based on the information in that particular RIF. If the packet does not contain a RIF, the MAC address is used instead.

Gateways—Translating Bridges

When a network is composed of more than one technology, some sort of device is needed that can translate between the different formats and protocols used. A *translational bridge* (also referred to as a *gateway*) can be used to connect Ethernet and Token-Ring LANs. A standard bridge simply forwards the same frame it receives based on the MAC address in the frame. A translational bridge operates up into the Logical Link Control (LLC) layer of the OSI reference model, takes the information it receives in one format, and repackages it into a different format, depending on the kinds of networks it is bridging. Figure 7 shows how bridges can be used to join different network types.

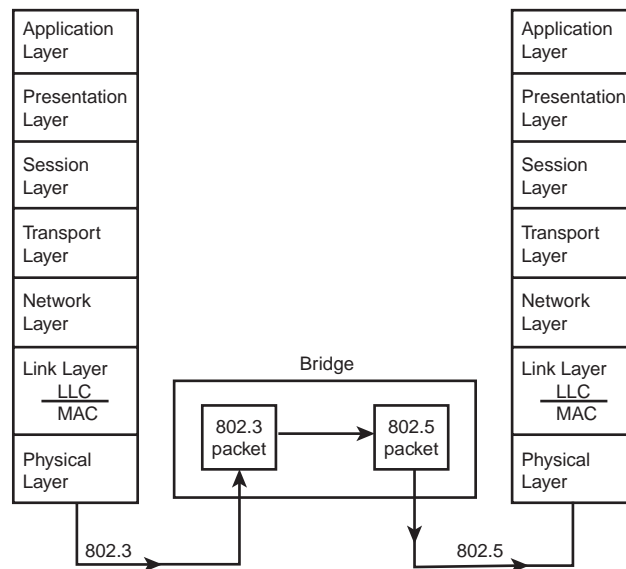


Figure 7 Bridges can be used to join different network types.

In this example, the bridge strips off the 802.3 MAC sublayer information from the incoming Ethernet packet, and at the LLC sublayer decides where the packet should be sent at the LLC sublayer. It then repackages the packet in an 802.5 frame and sends it back down to the physical layer for retransmission on the 802.5 Token-Ring network.

End user workstations do not need to be aware that this translation is even being done. Each continues to receive a frame in the format of the network on which it resides. This is one of the benefits of using a model such as the OSI reference model. Instead of having to write a complex program to allow for the interchange of data between two diverse networks, only a simple repackaging step is needed in the bridge device.

Token-Ring Switching

Switching is a method used to speed up the connection between two network nodes. It provides for a direct, dedicated path that does not involve contention for the network medium, either by CSMA/CD or by possession of a token. The usual means for joining multiple rings in a Token-Ring environment is to use a bridge or router. When you have users who make a lot of use of a particular server, about the only thing you can do to decrease latency and increase access time is to place the server on the ring that the users are on. In a larger environment, you can end up with servers distributed throughout the network, which makes them more difficult to manage.

Unlike bridges, switches enable you to attach a ring segment, an individual workstation, or a server directly. When a server is directly attached to a port on a switch, it does not have to wait for a token to come its way before it can transmit data. Likewise, the switch does not have to wait for a token when it wants to output frames on the port to which the server is connected. Thus, by connecting high-performance servers to a switch, and then using the switch to connect to other workstations or network segments, you can increase the total throughput for the network.

Just as a switch provides a dedicated path for Ethernet networks, allowing for full-duplex Ethernet connections, a switch also can be used to double the bandwidth of a Token-Ring connection. This technique is called *Dedicated Token-Ring (DTR)*. A Token-Ring adapter uses two pairs of wires, one pair to transmit and the other pair to receive. After you eliminate the token-passing function of the protocol, both of these wire pairs can be used at the same time.

Troubleshooting Token-Ring Bridges

Individual Token Rings contain many mechanisms, both physical and software based, that can be used to detect and temporarily fix many problems. When a network adapter on a workstation begins to malfunction, it is possible for other workstations to detect this and force the offending workstation to disconnect itself from the ring and perform a self-test. Administrators usually can quickly diagnose and fix problems that occur within the local ring without a lot of time spent investigating the problem.

When you begin to connect LANs to a larger network, however, connectivity between two workstations that reside on different rings can be troublesome and time-consuming. You must know what types of devices are installed between the workstations, and also understand how they are used.

For example, *SRT bridges*, which can be used for transparent or source route bridging, might be dropping packets that have a Routing Information Field (RIF) when connecting with an Ethernet segment. You can use a network analyzer to determine whether packets are being forwarded through the network. By checking the source address on packets output by the bridge—looking for a high-order bit set to one to indicate RIF—you can determine whether it is passing RIF frames.

You might think a certain path exists that handles all the types of traffic on your network, but you might be unaware that a particular device is misconfigured or is malfunctioning. Check the fields that show the ring numbers on packets that pass through the bridge to be sure they are correct. Also, check the bridge ID itself to be sure you don't have two bridges in the network set to the same value.

When the number of network clients increases suddenly, a bridge might begin to degrade in performance. Remember that bridges used in a small network might not have the capacity to buffer a traffic load when you begin to upgrade your network with higher-capacity servers or additional workstations that generate a lot of traffic. Consider installing a switch to relieve these types of congestion problems. Using a switch, you can isolate high-performance machines from those low-capacity ones while still maintaining connectivity.

Finally, remember that not all vendors' equipment works with everyone else's. When making purchasing decisions, investigate the documentation and sales literature to be sure that the bridge or other device you are planning to add to your network will work as you expect.

Using a Token-Bus Topology

One of the main differences between the star and the bus topologies is that all workstations that are connected to a bus can hear all transmissions that are made. Even when Ethernet networks are connected to a hub in a star-wired fashion, the hub repeats the incoming frames on the other ports so that the basic function of a bus remains.

Note

The IEEE 802.4 Working Group that developed token-bus standards has been inactive since 1997. This does not mean that token bus is a dead technology, just that, for its purpose, the standards have been defined and no further work is anticipated for this technology.

The MAU in a Token-Ring network does not create a bus topology when it concentrates the process of connecting a ring in one location. However, there is a variation on Token-Ring technology that *does* use a bus. This is referred to as a *token-passing bus*. In this type of Token-Ring implementation, the token frame is logically passed from one workstation to another, although all workstations are connected by a bus cable and all can hear every transmission. The MAP (Manufacturing Automation Protocol), which is a dated technology at this point but still in use in many locations, uses this technique for connecting equipment on the factory floor.

The IEEE 802.4 group set forth the standards for the token-bus network. Each station on the bus has a numerical address that determines the order in which it is granted the token. In Figure 8, you can see that the order in which stations are attached to the bus is not important because the token is passed according to station address.




Figure 8 In an 802.4 token-bus network, a logical ring is formed by the stations' addresses.

You can see that the workstations are connected along the bus in no particular order. Workstation 8 is the first on the bus and is followed by workstations 2 and 3, with the end of the bus having workstations 1 and 6 attached side-by-side at the end of the bus. You can move a station from one point in the network to another without having to change its address. The token, however, is passed in order of address from one station to the next—in other words, it starts at workstation 1, moves to 2, then 3, and so on through to workstation 8—and the token is still the method used to grant access to the network medium for transmission of data frames.

Cabling Rules for Token-Ring Networks

Two basic types of cables are used to connect Token-Ring networks. Workstations are connected to the multistation access unit (MAU or MSAU) by means of a *lobe* cable. The MAU is similar to a hub that is used in Ethernet technology but acts to connect workstations into a physical ring and does not repeat signals on all ports simultaneously. To create larger LANs, MAUs are connected to each other by using a Ring In/Ring Out cable.

What Is the Adjusted Ring Length?

The distances you can cover with a Token-Ring network depends on the type of cable used, the number of workstations, and other devices in the ring, as well as the particular vendor who manufactures the equipment.

The term *adjusted ring length (ARL)* is used to specify the worst-case distance between any two stations in the ring. It is calculated in this way:

$$\text{ARL} = \text{lengths of all MAU to MAU cables} - \text{shortest MAU-MAU cable.}$$

In some implementations, you do not have to count cables that are three meters or shorter and used in the same wiring closet to connect MAUs. Other devices, including surge suppressors and punch-down connectors, also might need to be figured into the calculation. Be sure to review the vendor's specifications carefully when making a purchasing decision.

Rules for STP and UTP Cables

Both shielded and unshielded twisted-pair cables can be used with Token-Ring networks. In many cases you will find that cables are classified according to a scheme devised by IBM back in 1982, the *IBM Common Cabling Standard (CCS)*. CCS defined eight levels of cable types, as well as a system of using ducts to conduct cables through buildings, with terminations at faceplates in the user's workspace. This was an early attempt to create a procedure for wiring a building so that various devices used to process data and voice communications could be connected in a standard and consistent manner.

Taking into consideration these cable definitions, several types of cables are now standard in Token-Ring networks. In the discussions that follow, the standard lengths quoted apply when the LAN is composed of more than one MAU. It usually is possible to achieve a greater distance in a small work-group LAN.

Type 1 Cables

Type 1 cables consist of two pairs of solid 22 AWG wires. Each pair is shielded by a layer of foil, and the entire cable is surrounded by a braided metal shield, encased in PVC. This type of cable is usually used for trunk connections—joining together MAUs that are distributed throughout an office building, for example.

These cables are used in places such as the walls of office buildings, and they connect wiring closets and distribution panels that are usually found in the same building. The maximum distance for this type of cable is 101 meters, or 331 feet.

Note

A cable designated as Type 9 is similar to the Type 1 cable, except that it uses a plenum jacket for use in air-handling spaces to satisfy the requirements of fire codes.

This type of cable can be used to connect as many as 260 devices, using 12 or fewer wiring closets and as many as 32 MAUs.

Type 2 Cables

Type 2 cables are similar to Type 1, except that they also include additional unshielded twisted-pair cables that can be used for voice or other transmissions. This means that there are six twisted pairs in the same cable. The cable can be used to provide both voice and data connections, using only one cable strung between end points. The maximum distance for this type of cable is 100 meters, or 328 feet.

Larger Token-Ring networks are usually constructed using STP (either Type 1 or Type 2) and can support as many as 260 workstations and as many as 32 MAUs.

Token-Ring networks of both 4Mbps and 16Mbps can be run on Type 1 and Type 2 cables.

Type 3 Cables

Type 3 cable is much less expensive than Types 1 and 2. Types 1 and 2 are made up of shielded twisted pair, while Type 3 is made up of unshielded twisted pair (UTP) cables with a minimum of two twists per inch, using 22 or 24 AWG solid wires. Remember that the more twists that are used for any given length of cable, the less crosstalk between the wires. The entire cable is encased in PVC.

This UTP cable is more flexible than the other types and, therefore, is easier to route around corners and in tight places. This cable type is generally used only on 4Mbps networks and has a distance of up to 45 meters (148 feet), according to IBM standards. Some vendors publish specifications that use this cable type for distances of up to 150 meters.

This is the most popular type of cable and can support as many as 72 workstations. If repeaters are used, subtract 1 from this value for each repeater. Depending on the vendor, this cable type can be used in two to six wiring closets.

Type 6 Cables

To connect MAUs, a patch cable is used that is usually referred to as Level 6 or Type 6. This is a shielded twisted-pair cable made up of two pairs of stranded wire with both foil and braided shielding. It is encased in a rigid PVC jacket.

These patch cables can be purchased in standard lengths of 8, 30, 75, and 150 feet, and can be special-ordered in custom lengths if required.

This cable is a little more flexible than the Type 1 cable and can be used to connect MAUs and to connect individual stations to the MAU. When used to connect stations, it usually is used only in a small network, with around 12 MAUs, which are usually of the IBM Model 8228 type.

Lobe Cables

The cables that connect individual workstations to the central MAU device are called *lobe cables*. Each cable uses four wires for communications. Two are used for transmitting data to the MAU, and two are used for receiving data. For the standard STP cable (Type 1 and Type 2), an IBM Type A data connector is used at one end to connect the cable to the MAU, and a nine-pin connector is used at the other end to connect to the Token-Ring network card.

If the cable type is UTP (Type 3), the connectors are usually RJ-11 or RJ-45, although the RJ-11 type is an older model that is not often used today. Most Token-Ring adapters that are currently being manufactured have both a nine-pin and an RJ-45 socket to provide for different cable types.

Ring In/Ring Out Cables

Ring In and Ring Out cables are used to connect MAUs to form a larger LAN. When you join MAUs, you connect the Ring In port of one MAU to the Ring Out port of another, with the end result forming a ring so that the last MAU is connected to the first.

Although a malfunction of a lobe cable will only cause the attached workstation to be dropped from the LAN, the failure of a cable that connects two MAUs is a more serious problem. When someone says that a Token-Ring LAN is "self-healing," this refers to its capability to lock out a malfunctioning workstation, not a break between MAUs.

Fiber-Optic Connections

Fiber-optic cable (sometimes called Type 5) has been used in Token-Ring networks for many years. One of the benefits of this cable is that it enables you to greatly extend the distances covered by the LAN. When computing the ARL, for example, fiber-optic cable usually counts as a zero-length cable. This means that if you use fiber-optic cable for all MAUs in the ring, only then do lengths of the lobe cables from the MAU to the workstation need to be figured in.

Token-Ring Connectors and Media Filters

When using IBM Type 1 cables, the standard connector is called an *IBM-type Data Connector (IDC)*, or a *Universal Data Connector (UDC)*. These connectors are not like RJ-45 connectors, in which the cable consists of a male connector that plugs into a female socket. Instead, these connectors are blade types that interlock with each other.

UDC connectors use the color-coding scheme shown in Table 2.

Table 2 Color Codes for UDC Connectors

Wire Color	Use	Polarity
Red	Receive data	+
Green	Receive data	–
Orange	Transmit data	+
Black	Transmit data	–

A DB-9 connector contains nine pins for terminating wires, but only four are used by Token-Ring networks. Table 3 shows the pin-out for this connector and the wire color-coding scheme that is used.

Table 3 Color Codes for DB-9 Connectors

Pin Number	Wire Color	Use	Polarity
1	Red	Receive data	+
5	Black	Transmit data	–
6	Green	Receive data	–
9	Orange	Transmit data	+

You also can use RJ-45 connectors. This is the type of connector most often seen in 10BASE-T Ethernet networks. This type of connector is similar to, but larger than, the modular connector used on most telephones today. The color-coding and pins that are used with this connector for Token-Ring networks appear in Table 4.

Table 4 Color Codes for RJ-45 Connectors

Pin Number	Wire Color	Use	Polarity
3	Blue/White	Transmit data	–
4	White/Orange	Receive data	+
5	Orange/White	Receive data	–
6	White/Blue	Transmit data	+

Another modular connector that is used, although not as often as those discussed here, is the RJ-11 connector that is standard for most telephone connections. This connector type typically is used with older network adapters. Table 5 lists the color codes and pins used for this type.

Table 5 Color Codes and Pin-Out for RJ-11 Connectors

Pin Number	Wire Color	Use	Polarity
2	Blue/White	Transmit data	–
3	White/Orange	Receive data	+
4	Orange/White	Receive data	–
5	White/Blue	Transmit data	+

Media Filters

If you have older Token-Ring adapters that have only a nine-pin connector, you can find media filters that can be used to join them to an RJ-45 type of connector. Various vendors produce these inexpensively. Ethernet cards are, for all practical purposes, a commodity item that can be replaced cheaply. Token-Ring adapters are more expensive, so using a media filter to attach an older card to a new cabling standard for your network might be a good idea, economically, if you have a lot of stations.

Cabling Scenarios

Now let's look at some typical cabling scenarios and the types of cable to use. The general specifications you should analyze are ring speed, maximum distance between any workstation and the MAU, the number of workstations, and whether a voice connection is going to be carried over the LAN.

Scenario #1:

Ring speed: 16Mbps

Maximum distance: 35 meters

Number of workstations: 62

Voice connection: No

Solution:

This scenario requires Type 1 cabling. Even though the number of workstations and maximum distance require only Type 3, the requirement of 16Mbps speed necessitates using Type 1.

Scenario #2:

Ring speed: 16Mbps

Maximum distance: 85 meters

Number of workstations: 185

Voice connection: Yes

Solution:

This scenario requires Type 2 cabling. The ring speed, maximum distance, and number of workstations can be implemented using Type 1 or 2, but the voice requirement can be handled only by using Type 2 cabling.

Scenario #3:

Ring speed: 4Mbps

Maximum distance: 35 meters

Number of workstations: 25

Voice connection: No

Solution:

This scenario requires Type 3 cabling, which is sufficient for the short distance covered by the network and the number of workstations. This type of cabling is the cheapest and easiest to install, making it a good choice for a small LAN.

Monitoring Token-Ring Utilization and Errors

Because of the fundamental difference between Ethernet and Token-Ring technologies, the statistics you use to monitor the network are also quite different. For example, an Ethernet network operating at 10Mbps will begin to degrade rapidly when overall utilization of the network media rises above 40%, mostly due to the access method used, which allows for collisions to occur when more than one station attempts to transmit at about the same time. Token-Ring networks operating at 4Mbps or 16Mbps actually can use a much larger percentage of the network bandwidth because only one station is normally transmitting at a time. The mechanisms that Token-Ring uses to grant access to the network (the token frame) and other functions used for maintaining network functionality (ring polling and beaconing) means that the kinds of information available to the administrator are different than those on an Ethernet network.

In addition, the types of errors that will be encountered will be different. The fact that Token-Ring networks are “self-healing” to a large extent makes tracking down problem workstations much easier than with Ethernet. When a problem is detected on a Token-Ring network, the individual workstation can be isolated and the rest of the LAN will continue to work while you try to determine the cause of the failure of the workstation.

Token-Ring Statistics

As discussed earlier in this chapter, two kinds of frames are used on a Token-Ring network. The LLC frame is used to carry data that a station on the ring wants to send to another station. Most of the frames generated on a Token-Ring network will be of this type. The network adapters, however, use another type of frame, called the MAC (media access control) frame, to communicate with each other and perform ring management functions on the ring.

The most basic MAC frames you will see are the active monitor present (AMP) and standby monitor present (SMP) frames. The AMP frame is normally passed around the ring every seven seconds. It is used to tell other stations which station is the active monitor. Other stations on the ring are considered to be standby monitors and will decide on a new active monitor if the AMP frame fails to arrive on time.

Network analyzers, both hardware and software, usually allow you to collect and display statistics about the functions performed by these ring management frames. They also can show you statistical information about errors that occur on the ring.

In particular, monitoring beacon frames is very important in maintaining a healthy network. These frames are sent to downstream neighbors when an adapter fails to detect any data or MAC frames from its nearest upstream neighbor, usually indicating that the upstream neighbor is malfunctioning

or there is a possible cable fault. Obviously, if you notice that the ring is experiencing a high level of beacon frames, it might be indicative that an adapter or possibly a cable is unstable and should be looked at immediately.

Ring purges are used to set the ring back to a known state when something has gone wrong, and also can happen when the active monitor leaves the ring and a new one is elected, or when the active monitor detects some other kind of error condition and needs to reset the ring. When a station receives a ring purge frame, it stops its current processing and resets its timers. If the ring is experiencing a high level of ring purges, you should study other error or statistical information to determine the cause.

For example, a good monitoring program will show you the current number of active stations on the ring. Usually, it also will show you the maximum number of stations that have been on the ring since the monitor began recording data, and possibly an average number of stations. If the average number of stations on the ring is significantly lower than the maximum, check to determine why stations are leaving and rejoining the ring. If you are moving equipment around or rearranging workstations on the network, this might not present a problem.

The Ring Error Monitor

An optional software component that many administrators will find quite useful for troubleshooting is the Ring Error Monitor software. If you will remember, every ring has an AMP and the remaining workstations are designated as standby monitors. The Ring Error Monitor is a station that does nothing more than receive error reports from other stations as they detect them. It keeps a running list that the administrator can review when trouble is suspected on the ring.

When a station on the ring detects that an error has occurred, it will wait for a few seconds (two seconds by default) and listen to the network to see whether any further errors occur. After the wait interval has expired, the station will send a report of the errors it has seen to the Ring Error Monitor. If you are using a Ring Error Monitor on your network, review the error data it collects on a regular basis. Even simple errors that do not cause downtime for the ring can indicate that a component is unstable and prone to more severe failure in the near future. Taking proactive steps to diagnose and fix a small problem can prevent a much larger problem from disrupting other users later.

Token-Ring Errors

Errors can be classified into two broad categories: hard errors and soft errors. Soft errors are those that can occur during normal ring operation, but they do not bring down the local ring. Hard errors are those caused by such things as broken or malfunctioning cables or ports on the MAU that can possibly stop normal functioning of the ring.

Soft errors can be described as either isolating or nonisolating errors. An *isolating error* is an error that can be traced back to a single station on the ring, and a *nonisolating error* is one that cannot. Simply put, an isolating soft error usually indicates an error condition with a particular station on the ring. That is, it can be isolated to a particular station. A nonisolating error usually indicates a problem with the ring itself, and not necessarily a particular workstation.

The IEEE 802.5 specification, which sets the standards for Token-Ring networks, defines several error types:

- **Burst Error**—This error indicates that there is noise on the network media. It occurs when a station detects three bits with no clock in the middle. The source of the noise can be environmental, such as a strong electrical field near a network cable, or it might result from faulty hardware. This error also can happen when a station is inserting itself into the ring or taking itself out of the ring. The burst error probably is the most frequent error you will see on a Token-Ring.

- **Line Error**—This error is similar to a burst error and usually indicates that one is about to happen. It can be generated when a station receives a frame and calculates that the CRC value does not match the contents of the frame. When a station detects this kind of error, it sets the error bit in the frame to one so that other stations do not report the error. In general, you will see a ratio of 1 line error for every 10 burst errors detected.
- **Lost Frame Error**—When a station transmits a frame, it sets a timer that tells it how long it should expect to wait until the frame travels around the ring back to it. If this timer expires, this kind of error is generated. This error type causes the active monitor to generate a new token.
- **Token Error**—The active monitor sets a timer each time it sees a valid frame on the ring. Because the monitor knows how long it takes for the frame to travel the distance around the ring, it assumes that an error has occurred if this timer expires before it sees another frame. This can result from noise on the line. The active monitor generates a new token when this condition occurs.
- **Internal Error**—A station records this kind of error when it detects an internal parity error when using DMA (direct memory access) to exchange data with the workstation's memory. To determine whether the network adapter card is the problem, you can install it in a different workstation and see whether the error occurs. If it does not, try power-cycling the original workstation to see whether the error can be corrected.
- **Frequency Error**—When the frequency on the network media detected by a station differs significantly from that which is expected, this kind of error is recorded. A standby monitor on the ring can detect this error, which usually results from a problem with the active monitor that generates the clocking signal. To determine whether the active monitor is problematic, remove it from the ring and see whether this problem still occurs.
- **AC Error**—This error type indicates that a station received more than one active monitor present or standby monitor present frame with both the address recognized and the frame copied fields set to zero. This error happens when the nearest upstream neighbor of the station does not properly set these bits. Check the upstream station when troubleshooting.
- **FC or Frame Copied Error**—A station generates this error type when it receives a MAC frame that is addressed to it yet the address recognized bit is already set to 1, indicating that another station also thought this frame was destined for it. This can be the result of a problem on the line or possibly indicate that two stations on the ring have the same address. Note that the duplicate address problem will exist at the frame level and is not indicative of a duplicate address of a higher-level protocol.
- **Abort Delimiter Transmitted**—This error type happens when a station transmits an abort delimiter while it is transmitting. It happens when a station, while transmitting, receives a claim token or beacon frame, which causes it to abort its transmission.
- **Receive Congestion Error**—When a station receives a frame that is destined for it but does not have enough buffer space to copy the frame, this kind of error occurs.

Some of the problems causing these errors can be fixed easily. For example, if you notice a large number of frame copied errors, indicating a possible duplicate address, and you have just installed a new workstation on the ring, recheck the new workstation to be sure you correctly configured its address. Congestion errors can be solved by replacing older network adapters that might not be able to keep up with the traffic on the ring. Congestion errors also can indicate that an adapter is having problems and is about to fail. Replace the adapter to see whether the error persists. Burst errors and token errors might lead you to check the cabling or connectors for all stations on the ring. If cabling has recently been moved or new cabling installed, check for sources of electrical interference.

Internal errors should direct you to examine the adapter or workstation from which they originate. Replace the adapter to see whether the problem subsides and, if not, perform diagnostics on the workstation itself.

Using Network and Protocol Analyzers

Besides the error-recording functions of the Ring Error Monitor, a good protocol analyzer can be a big help in troubleshooting problems on a ring. The protocol decode function provided by the analyzer should be able to decode frames specific to Token-Ring networks (that is, the MAC frames used for ring maintenance), as well as higher-level protocols such as TCP/IP or IPX/SPX.

It is important to understand how your network works during the normal workday when using a protocol analyzer. During certain times of the day, normally there will be peaks during which activity is high on the ring, whereas at other times utilization is quite low. If you are unfamiliar with these times, you might not fully understand the voluminous data you can get using a protocol analyzer. Some soft errors are to be expected, such as those caused when a station is inserted or removed from the ring. If you are not familiar with the pattern of utilization on the ring and the usual error activity, you might judge the ring to be in worse shape than it is!

Token-Ring Extensions to the Remote Network Monitoring MIB

The objects defined in previous RFCs for monitoring network activity, such as those defined in RFC 1271, are separated into nine groups. Some of these groups, such as the statistical groups, are designed with the assumption that the underlying network technology at the Data Link layer is Ethernet. To fully implement SNMP and RMON technology for devices that are used for Token-Ring networks, additions had to be made.

Tip

RFC stands for "Request for Comments" and is a mechanism used on the Internet in the process of defining standards. SNMP stands for the Simple Network Management Protocol, which is a technology that allows for monitoring of network components remotely by the network administrator. RMON (Remote Monitoring protocol) is similar to SNMP but extends its functionality. MIB stands for Management Information Base and is a collection of objects that define the kinds of data collected and functions that can be performed by RMON. For more information, see Chapter 53, "Network Testing and Analysis Tools."

RFC 1513 defines extensions to the MIB that was defined in previous RFCs, and describes how existing object groups will be used for monitoring Token-Ring networks. These are the changes:

- **Host Group**—Only isolating errors will cause error counters in this group to be incremented. These are line errors, burst errors, AC errors, internal errors, and abort errors.
- **Matrix Group**—No error counters will be incremented in this group for Token-Ring networks.
- **Filter Group**—Conditions are defined for how the bitmask is used in this group.

In addition, the RFC provides additional groups of objects that are specific to Token-Ring networks. In particular, RFC 1513 defines four additional object groups:

- **Ring Station Group**—This group provides information about each station on the local ring and also status information about each ring that is being monitored.
- **Ring Station Order Group**—The order of stations in the ring is contained in this group.

- **Ring Station Configuration Group**—This group is used to manage stations by active means. Objects in this group can be used to remove a station from the ring and to download configuration information from a station.
- **Source Routing Statistics Group**—Utilization statistics about source routing are contained in this group.

Although the Ring Error Monitor can provide useful information collected on the local ring, SNMP and RMON can allow you to centralize monitoring for larger networks, making the troubleshooting process much easier. In addition, RMON provides for alarm functions that can be set to trigger when certain events occur, giving you a heads-up before a situation becomes more serious.

When upgrading network components, check to be sure that they provide SNMP and RMON functionality.

Troubleshooting Tips

As you can see, there are many kinds of information you can obtain that can tell you how well your Token-Ring network is functioning. Sometimes, though, too much information can be useless if you don't understand how to interpret it. The best thing you can do to keep yourself prepared is to monitor the network on a frequent basis and to keep track of normal usage patterns. In other words, determine a baseline for the network and use the baseline data when analyzing or troubleshooting.

A few general things that should apply to most rings, however, are listed here:

- Every seven seconds, the active monitor present frame should be seen circling the ring. This ring polling mechanism allows each station to determine who its nearest active upstream neighbor is. It also lets each station know that the ring is functioning normally and that the active monitor is watching for token and other errors. Following shortly after the AMP frame, you should notice a standby monitor present frame.
- Soft errors and ring purge frames are expected events when a station is inserted or removed from the ring, and should not be a cause for alarm. Soft errors that are not associated with insertion or removal should be less than 0.1% of the total number of packets transmitted.
- Insertion and removal of a station can also cause burst errors and line errors, and can cause token errors. If these errors are not associated with the timeframe during which an insertion or a removal was performed, they might indicate a problem unless they are sporadic.
- Utilization on a Token-Ring network above 70% can cause slower response for end users. Normally, response times should be around one-tenth of a second. Peak utilization of 100% does not indicate a problem unless it is happening frequently.

Choosing Between Token-Ring and Ethernet

Token-Ring advocates and those who champion Ethernet have argued for many years over which technology best serves the market. There are advantages and disadvantages to both. It is interesting to note, however, that as Ethernet technologies are increasing in bandwidth, they are incorporating some techniques that are already present in Token-Ring technology.

Because Token-Ring techniques can guarantee an equal access to the network within a specified time limit, they are more suited to environments that require a specific response time. Ethernet networks are fine for many network applications, but can become bogged down when too many workstations try to flood the wire with a lot of information at the same time. A lot of processing power can be spent simply by the contention method that involves collisions and back-off algorithms. However, in a high-bandwidth Ethernet network using switching technology, response time is minimal compared

to older Ethernet implementations. Because of this, it is easy to create an Ethernet network that can satisfy the minimal response time that Token-Ring enables.

Ethernet networks use a relatively small packet size, which can be up to a maximum of only 1,518 bytes.

This size was originally conceived to work in the 10Mbps environment, and has been kept as Ethernet speeds have advanced to 100Mbps. With the exception of Jumbo Frames—proposed by Alteon Networks—most Ethernet networks waste valuable bandwidth due to the overhead of these small frames. In 16Mbps Token-Ring networks, the maximum frame size is controlled by a timer and can reach sizes of up to 18KB, although 4KB is more common in actual practice at this time. This is a considerable improvement over Ethernet!

Other areas in which Token-Ring excels over Ethernet include the priority and reservation fields, which can be used to provide a quality of service functionality. By using source route bridging, Token-Ring clients can specify multiple routes to a particular destination, which is not allowed under the spanning tree routing technologies.

Today, Token-Ring should not be considered for creating a new LAN. It is legacy technology that is offered by IBM and a few other vendors mainly to support customers who still use it. For all practical purposes, Ethernet is the choice today for implementing a new LAN. There are other contenders, such as ATM (which has been revamped for LANs), but the largest installed base of LANs is composed of Ethernet technology.

The main disadvantage to Token-Ring, other than its declining market share, is that Token-Ring usually is a more expensive implementation than Ethernet. Token-Ring adapters are made by very few companies and, because of their complexity, cost more to produce. MSAUs are more complex than their passive hub cousins, and also are more expensive. In addition, because the MSAU is a required component of a Token-Ring network, you cannot simply create a crossover cable to connect two workstations as you can with Ethernet. However, because a two-node network is not a common implementation except in a very small business or a home environment, this is not a large disadvantage.

The two technologies can operate together, each being used for what it does best. Switches, routers, and adapter gateways can be used to connect the two types of networks so that workstations can communicate throughout a larger network.

