Madge Networks
TechNote

Number: 08104-02
Date: March 1995
Total pages: 7

How Token Rings work and maintain themselves

1 Introduction

This document describes the principles of Token Ring operation and introduces the system of self-maintenance that is built into the design of Token Ring.

The document begins by describing the ring topology (Section 2) and the method by which stations transmit data onto the ring (Sections 3 and 4). It then describes the main processes by which a Token Ring maintains itself (Section 5). This description introduces the Media Access Control (MAC) protocol on which the system of self-management in Token Ring depends.

The document is based on the description of Token Ring to be found in *Novell's Guide to NetWare LAN Analysis* by L. Chappell and D. Hakes. We recommend that you refer to this book which is cited in full in Section 6 below.

2 How Token Ring stations are connected together

The stations of a Token Ring network are connected in a ring formation, and all signals pass around the entire ring. The stations are connected one-to-another by transmit and receive cables. Figure 1 shows a Token Ring comprising six stations. Station 1's transmit cable is connected to station 2's receive cable. Station 2's transmit cable is connected to station 3's receive cable, and so on, until the ring is complete. In each case, the transmitting station is the Upstream Neighbour of the station that receives the signal. For example, station 1 is the Upstream Neighbour of station 2. The concept of the Upstream Neighbour is an important part of the ring's system for isolating problems and recovering from them (see Section 5.5).

Figure 1 represents the network as a logical ring. Remember that, physically, the stations are connected in a star formation, to a central wiring concentrator.

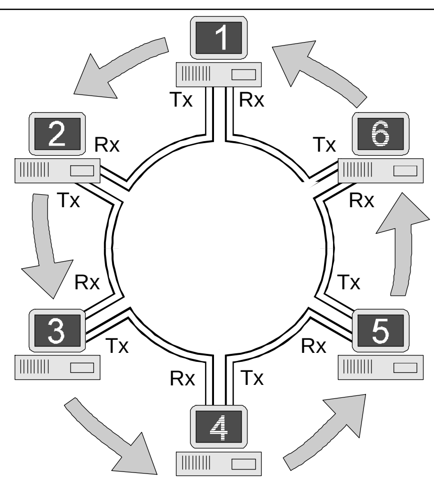

Figure 1 Each station on a ring receives frames from its Upstream Neighbour

3 How a station transmits onto a ring

When a Token Ring station has data to transmit, it must perform the following tasks (of course, these take only milliseconds to execute):

Capture the token To send data, the station must monitor the ring until it detects and captures the token. This is a bit-pattern which circulates on the ring until a station removes it. Because a station can only transmit when it has the token, every other station is denied access to the ring until the transmitting station releases it again.

Transmit the data Once it has captured the token, the station transmits a data frame onto the ring. Each station on the ring repeats the frame until the transmitting station receives it back again. Each station also checks to see whether a frame is addressed to itself. If it is, that station sets the frame's Address Recognized Indicator (ARI) and Frame Copied Indicator (FCI) bits to indicate that it has received the frame. If it detects an error, it sets a bit called the Error Detected Indicator (EDI).

Strip the frame When the frame arrives back at the transmitting station, the station removes it from the ring. This process is called stripping the frame. The station can tell from the ARI, FCI, and EDI, whether the frame arrived at its destination intact.

Release the token When it has stripped the data from the ring, the transmitting station releases the token back onto the ring. This gives other stations the opportunity to transmit frames.

4 The concept of Early Token Release

Token Ring networks operating at 16Mbps benefit from reduced latency because they permit Early Token Release. What this means is that stations on a 16Mbps network can release the token immediately after transmitting a frame instead of waiting until they have stripped the frame. This increases the performance of the ring by enabling:

- **More than one station to have frames on the ring at the same time.** Because a station releases the token as soon as it has finished transmitting a frame (instead of after stripping it) another station can capture the token and transmit before the original station has stripped its frame from the ring. This means that there can be frames from more than one station on the ring at the same time.
- **Stations to transmit more often.** A station that is transmitting large frames (for example, 4kbyte frames) can begin to receive a large frame back for stripping before it has finished transmitting it. If no other stations need to transmit, the station receives the token straight back and can transmit again immediately after stripping the frame. This allows the station to take maximum advantage of intervals when no other stations need to transmit. On a 4Mbps Token Ring a station cannot release the token until it has finished stripping the frame. While it is doing this another station might become ready to transmit. If the other station captures the token when the transmitting station does release it, then the original station will have to wait even longer to finish sending the data in its buffers.

5 Self-maintenance on the ring and the role of the MAC protocol

This section describes the main processes by which a Token Ring maintains itself in operation. Each process involves communication between stations by means of the MAC protocol. The different processes constitute a self-recovery cycle which is the ring's response to breaks in communication. By means of this cycle the ring is able to isolate and recover from problems dynamically and with minimal disruption.

5.1 Monitor Contention: electing an Active Monitor

The station with the most important management role on the ring is called the Active Monitor. It is this station that starts the ring after an interruption (see Section 5.2) and that ensures that all the stations on the network know the address of their Upstream Neighbour (see Section 5.3).

The first station to join a ring detects that the ring has no Active Monitor. It then initiates a process that will produce one. This process is called Monitor Contention. (In practice, some stations can be configured not to participate in Monitor Contention.)

To initiate the process, the station transmits a MAC frame called a Claim Token frame, and it continues to do this every ten-to-twenty milliseconds (depending on the implementation of the MAC software). When a station receives a Claim Token frame whose source address is:

- Higher than its own, the receiving station repeats the Claim Token frame, and takes itself out of contention.
- Lower than its own, the receiving station strips the frame and transmits a Claim Token frame of its own. (If it has already started transmitting Claim Token frames, then it just strips the frame from the ring.)

When a station that is releasing Claim Token frames receives back one-to-three (depending on the implementation of the MAC software) of its own Claim Token frames, it makes itself Active Monitor. In Figure 2, station 2 has the highest MAC address. Therefore, all the other stations repeat station 2's Claim Token frame until station 2 makes itself Active Monitor.

Although the Active Monitor is effectively the manager of the ring, all of the stations that are configured for Monitor Contention participate in the management process. Once an Active Monitor has emerged, the other stations begin to monitor the network for evidence of a failure of the Active Monitor. Each one that is ready to initiate Monitor Contention itself is known as a Standby Monitor (see Section 5.3).

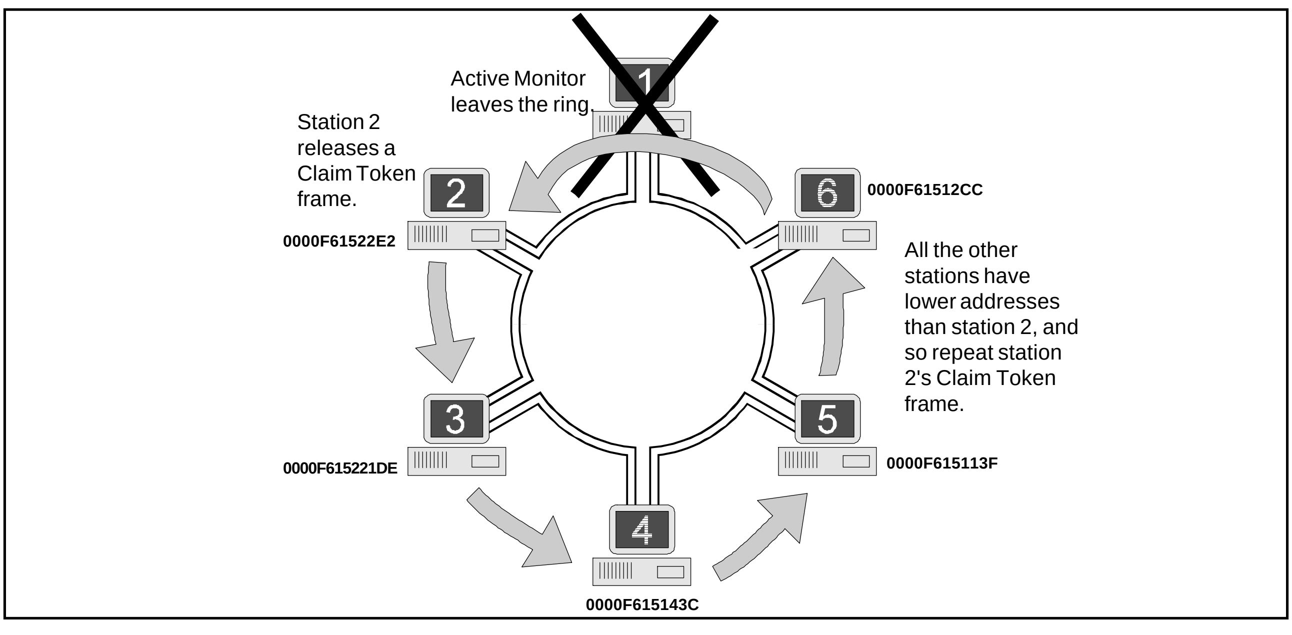

Figure 2 Monitor contention when the departing AM's downstream neighbour (station 2) has the highest MAC address and will, therefore, take over as the new AM

5.2 Ring Purges: starting the ring after an interruption

The first thing that the Active Monitor does when it emerges from Monitor Contention is to perform a Ring Purge. This means that it transmits Ring Purge MAC frames approximately every few milliseconds until it receives one back. It then stops transmitting them, and starts the ring by releasing a token.

Whenever the Active Monitor detects that the token, or a frame, has been lost or corrupted, it performs a Ring Purge again. Once a ring is up and working Ring Purges are the first stage in its self-recovery cycle. They are normally successful because they are normally responses to temporary breaks caused, for example, by stations entering or leaving the ring.

However, if a Ring Purge fails (that is, if the Active Monitor does not receive any of its Ring Purge frames back within a timeout period of approximately one second), then the Active Monitor starts Monitor Contention itself. If Monitor Contention fails to produce an Active Monitor that can start the ring up again, the station that was the Active Monitor starts Beaconing (see Section 5.5).

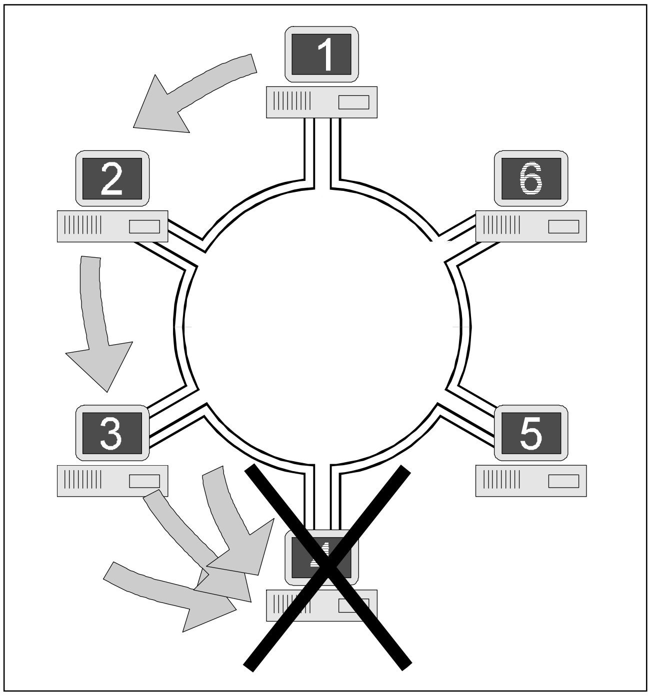

Figure 3 Frames are lost when station 4 leaves the ring

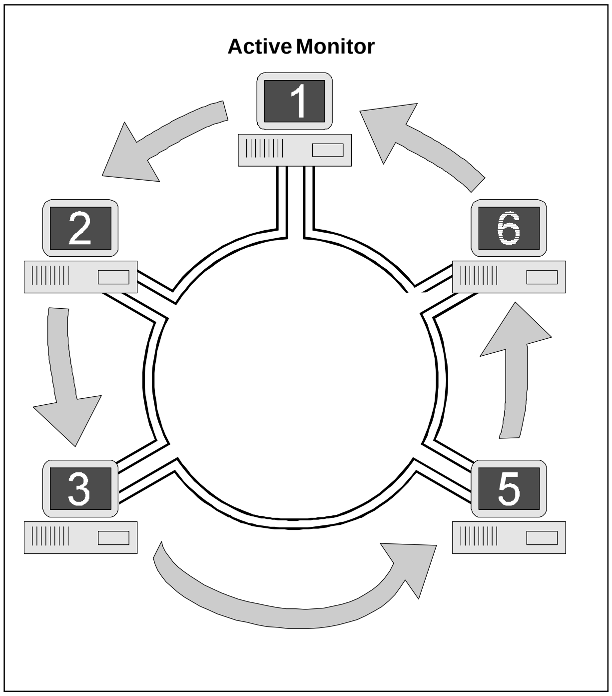

Figure 4 Station 1, the Active Monitor, sends Purge frames until the break disappears

5.3 Ring Polling: the neighbour notification process

After releasing the token, the Active Monitor can perform another of its ongoing management functions: every seven seconds it performs a Ring Poll. This enables all the stations on the ring to learn the address of their Upstream Neighbour (information that changes when stations enter or leave the ring). It also informs them that the ring is working properly and that there is an Active Monitor present.

To initiate Ring Polling the Active Monitor releases an Active Monitor Present frame whose Address Received Indicator (ARI) and Frame Copied Indicator (FCI) bits are set to 0. This indicates to the next station downstream that it is the first station to receive the frame. The receiving station, therefore, stores the source address of the frame as the address of its Upstream Neighbour. Before repeating the Active Monitor Present frame, it sets the ARI/FCI bits to 1 to ensure that all stations downstream are aware that the frame has already been received.

Having repeated the Active Monitor Present frame, the station sends a Standby Monitor Present frame whose ARI/FCI bits are set to 0. This indicates to the next station that no other station has received the frame. The receiving station, therefore, stores the source address of the Standby Monitor Present frame as that of its Upstream Neighbour. Before repeating the frame, it sets the ARI/FCI bits to 1 to ensure that all stations downstream are aware that the frame has already been received. It then sends its own Standby Monitor Present frame with the ARI/FCI bits set to 0.

The Ring Poll is complete when the Active Monitor receives a Standby Monitor Present frame whose ARI/FCI bits are set to 0. At this point all the stations know the address of their Upstream Neighbour. This helps them to isolate and recover from hardware problems on the network (see Section 5.5).

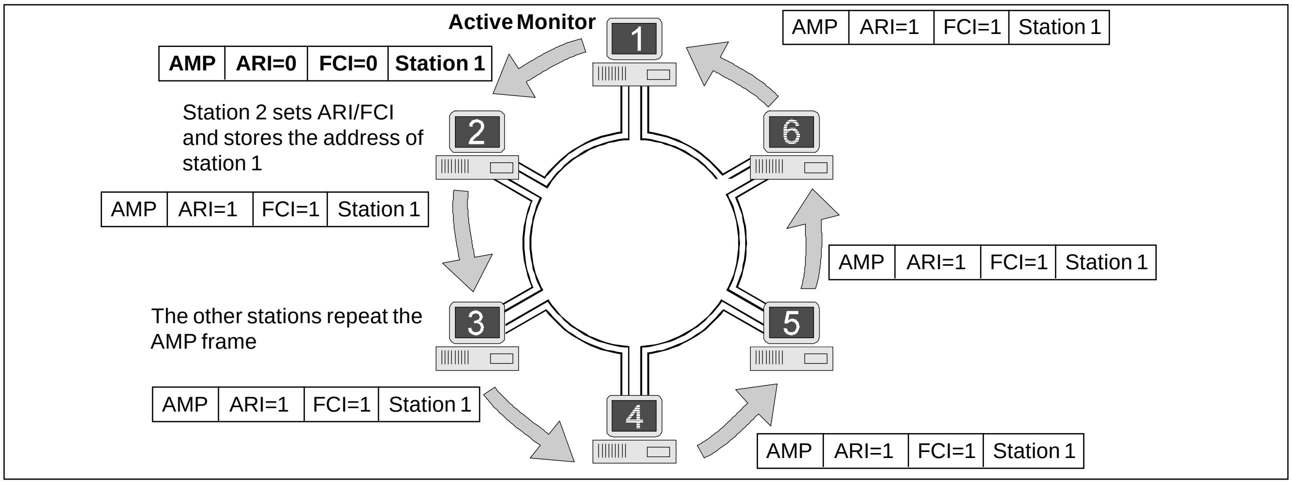

Figure 5 Station 2 stores the address of station 1 as that of its Upstream Neighbour

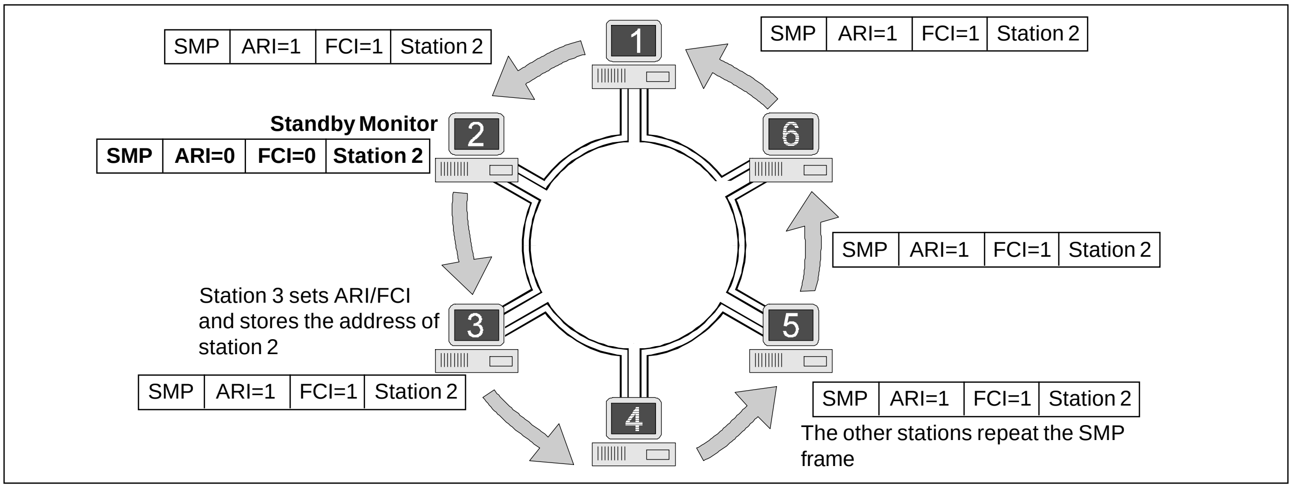

Figure 6 Station 3 stores the address of station 2 as that of its Upstream Neighbour

5.4 Entering the ring: tests performed by the station

Before attempting to join the ring, or in response to a Beacon frame (see Section 5.5), a station performs a lobe test. During this test, the concentrator connects the transmit and receive cables of the adapter together (which isolates it from the ring, if the lobe test is a consequence of Beaconing) and the station transmits a number of Lobe Media Check frames and a Duplicate Address Test frame. If it receives all of these back, it places a voltage on the lobe cable and when the wiring concentrator detects this, it inserts the station into the ring.

The station now waits to see that the ring is working. To check this it looks for Active Monitor Present frames (see Section 5.3), Standby Monitor Present frames (see Section 5.3), or Ring Purge frames (see Section 5.2). If it has to wait more than eighteen seconds, it assumes that there is no Active Monitor present, and initiates Monitor Contention (see Section 5.1).

The station also checks that no other station is using its MAC address. It does this by sending a set of Duplicate Address Test frames. If it receives two of these frames back and both have their ARI and FCI bits set to 1 (indicating that another station has received the frames), then it assumes that another station is using its address and it removes itself from the ring. Otherwise, it participates in the Ring Poll process (see Section 5.3) to learn the address of its Upstream Neighbour on the ring.

5.5 Beaconing: isolating and recovering from hardware errors

Beaconing is a process that enables the ring to isolate hardware errors. It can be initiated by any station and is another example of the distribution of network management functions in the design of Token Ring.

Any station that detects a loss of signal on its receive cable first starts Monitor Contention (see Section 5.1). If this fails, it sends Beacon frames every 20 milliseconds. In them, it identifies its Upstream Neighbour as the likely source of a problem. (When a station has not received a signal for some time, it always assumes, first of all, that the problem lies with its Upstream Neighbour.) The first consequence of Beaconing is that all activity on the network is interrupted and each station goes into Beacon repeat mode.

When the station upstream of the Beaconing station receives Beacon frames identifying it as the possible cause of a problem, it removes itself from the ring to perform a lobe test. This is a test of its transmit and receive cables and of its adapter (see Section 5.4).

If the station upstream of the Beaconing station:

- **Fails its lobe test**, it remains off the ring, and, when the Beaconing station starts receiving its own Beacon frames back, it knows that the problem has been removed (because it can now receive frames again). It then stops Beaconing and goes into Monitor Contention, and the ring recovers.
- **Passes its lobe test**, it re-enters the ring and enters Beacon repeat mode. However, the Beaconing station will still not receive any of its Beacon frames back, because the hardware problem still exists even though it has not been caused by the Beaconing station's Upstream Neighbour. The Beaconing station, therefore, removes itself from the ring to perform a lobe test. When it does this, its downstream neighbour starts Monitor Contention. If the original Beaconing station fails its lobe test and stays out of the ring, Monitor Contention will succeed and the ring will recover. If the Beaconing station passes its lobe test and re-enters the ring, manual intervention may be required to restore the network. This is because the hardware fault lies not with either of the two stations but in an area between them: the Fault Domain (see Figure 7). Intelligent Wiring Concentrators (for example, Madge Smart CAUs), can often isolate the problem in the fault domain so that the ring is not interrupted.

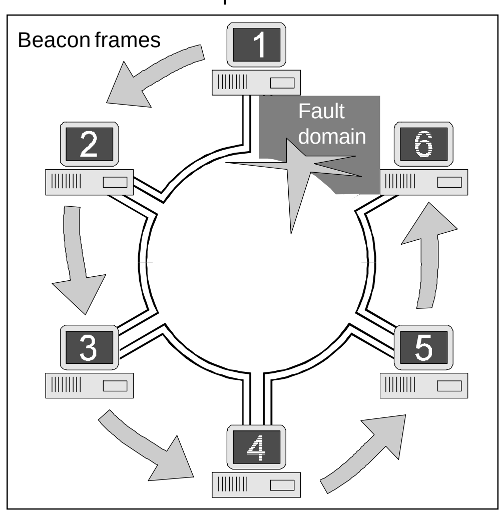

Figure 7 Manual intervention required between stations 1 and 6

6 Further reading

L Chappell & D Hakes, *Novell's Guide to NetWare LAN Analysis*, (2nd Edition, Novell Press, 1994)

D Nasser, *Token Ring Troubleshooting Guide* (New Riders Publishing, 1992)

IBM Corporation, *IBM Token-Ring Network Architecture Reference*

Europe:

Madge Europe
Madge Networks Ltd.
Knaves Beech Business Park
Loudwater
High Wycombe
Bucks HP10 9QZ
England

Tel: +44 1628 858000
Fax: +44 1628 858011
BBS: +44 1628 858008
Technical Support: +44 1628 858700

Disclaimer of warranty and limitation of liability.
Madge makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Madge shall not be liable for incidental, consequential, or special damages in connection with the furnishing, performance, or use of this material which can be changed without notice.

United States:

Madge Americas
Madge Networks Inc.
2310 North First Street
San Jose
CA 95131 - 1011
USA

Tel: +1 408 955 0700
Fax: +1 408 955 0970
BBS: +1 408 955 0262
Tech Supp: +1 800 876 2343

International:

Madge Asia
Madge International Ltd.
64/F Central Plaza
18 Harbour Road
Wanchai
Hong Kong

Tel: + 852 2593 9888
Fax: + 852 2519 8022
BBS: + 852 2593 9829
Tech Supp: + 852 2593 9888

Japan:

Madge Japan
Madge Japan KK.
Believe Mita
43-16, Shiba 3-chome
Minato-ku
Tokyo 105
Japan

Tel: + 81 3 5232 3281
Fax: + 81 3 5232 3208
Tech Supp: + 81 3 5232 3281