Token Ring Solutions

Jonathan Follows

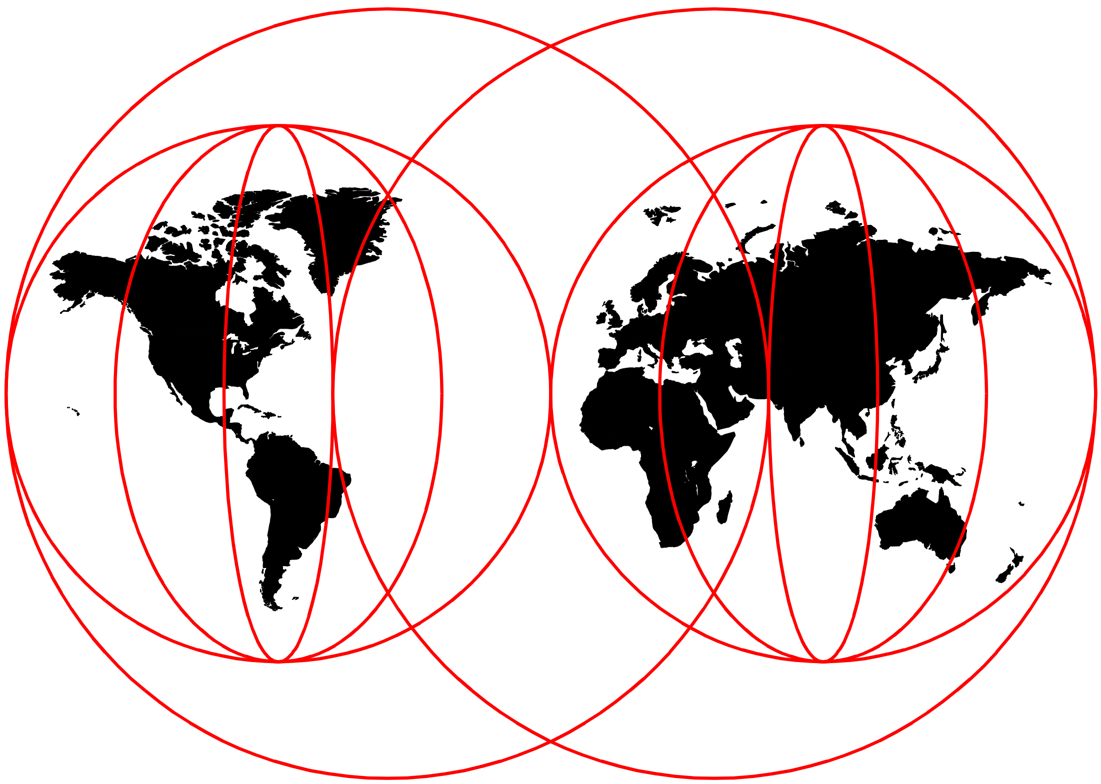

International Technical Support Organization

www.redbooks.ibm.com

Preface

This white paper describes the token ring products and technologies offered by IBM, concentrating on adapters, hubs and switches. It provides some basic training in token ring technology and the various features of its token ring products. It explains how token ring continues to be used in the 21st century and how existing token ring networks can be enhanced by the use of IBM's token ring products.

Who wrote this white paper

Jonathan Follows is an IBM-certified networking specialist at the International Technical Support Organization, Raleigh. He writes redbooks on all areas of IBM networking hardware and software, most recently on Cisco/IBM router interoperability. Before joining the ITSO in 1998, he worked as a technical specialist providing sales and marketing support in the United Kingdom and has 15 years' experience in all types of networking. Jonathan read Mathematics at Oxford University, England, and holds a degree in Computing Science from London University, England.

Thanks to the following people for their invaluable contributions to this project:

Mark De Cain, Jeff Huxtable, Terry Brest, Greg Zeph, Jon Bitner, John Botjer, Frank Tutone, Bob Love, Sun-Ling Meckley, Gene Osten.
IBM Networking Business Unit

Shawn Walsh, Gail Christensen
International Technical Support Organization, Raleigh Center

Contents

Executive summary

The robustness of token ring has won it a place in the heart of many large enterprise network managers. IDC estimates that more than 20 million network nodes exist, with growth foreseen through 2001. Network managers are now seeking ways to leverage their investments as they expand their networks to support ever-increasing numbers of bandwidth-hungry users and applications. This paper describes the current token ring product set offered by IBM and will also describe the many solutions still available to our token ring customers. Deployment of the methodologies and principles presented in this paper, when combined with the fundamental strengths and advantages of token ring technology, will provide the great majority of current token ring users with the tools and sound justification for continued investment in this proven technology.

IBM’s Networking Business Unit (NBU) plans to gradually phase out new development and marketing of Ethernet routing and switching products. IBM’s NBU will, however, continue to support and invest in its token ring adapter products. We are committed to our customers and will continue to sell our routing and switching products to those with existing contracts, and will provide ongoing support for the full length of that contract with those customers. IBM is working closely with Cisco to develop customer migration plans to allow interoperability between IBM and Cisco products and realize a smooth transition. For customers who have investments in our token ring technology, IBM continues to improve and invest in the token ring family of adapters.

It is the practice of IBM’s Networking Business Unit to assure customer satisfaction by delivering excellent solutions, products and services that meet or exceed our customers’ expectations.

Introduction

Network managers usually focus on the near term but recognize the need for longer term strategies. They must always have one eye on today's requirements and all new solutions which satisfy them, and the other on the new problems looming on their horizon. How will today's networks have to change? What plans must be put in place to create an environment where network capacity will increase smoothly and comfortably to meet the requirements of the organization? Strategies evolve from requirements. Plans must be firmly established to continue meeting the ongoing requirements of the network, while preparing for stepping up to its new challenges. Many network managers chose token ring initially for its ability to flexibly accommodate their initial requirements. The choice has been validated over time and proven to be a good one. Token ring has delivered on its promise to be the most robust, reliable LAN in the industry. Frequent innovations that resulted in product announcements and enhancements have provided token ring users with a technology that has continued to surpass network requirements. This paper will address the virtues of why token ring technology was chosen, its present status, and outline a smooth, step-by-step migration path to increased network performance. This migration path starts with protection of customer investments in their token ring network components and capabilities available today and continues by revealing how to exploit newer product functions for extending the life of a proven network design.

The remainder of this section provides a short history of token ring's development and leads to a summary of the current state of the business.

IBM leadership role in Token Ring

How and why IBM pioneered the development of token ring in the early 1980s as a highly robust and efficient network protocol for mission-critical business applications is now considered ancient history. Nevertheless the underlying customer requirements that drove the development effort have never been more pressing than in the Internet era. Token ring technology has served that business well, becoming one of the two major LAN types and enjoying widespread industry acceptance. Token ring is an indispensable foundation in many corporate infrastructures. Some of the inherent capabilities of token ring were developed specifically to provide superior reliability, availability, and serviceability (RAS) characteristics unavailable at the time in other LAN types (such as Ethernet). A few examples of the RAS features built into token ring are:

- End stations perform comprehensive pre-insertion checks of both their electronics and their attachment cabling before becoming active on the network.
- Each station maintains a record of its nearest active upstream neighbor's address, and of the number of frames it received that were corrupted after being transmitted by its upstream neighbor. If the number of corrupted frames exceeds a settable threshold, messages are transmitted that can be used by network management to locate a problematic station. In the event of hard errors caused by stations malfunctioning, special frames are transmitted that will automatically cause the failing station to remove itself from the ring, allowing the ring to heal.

- Access to the ring is via a protocol that is inherently “fair,” thereby allowing all stations the availability they require. No advantage accrues based on topological position in the ring. In addition, because access is controlled, there are no collisions, and ring utilization rates near 100% are achievable.
- For special messages, and devices (like bridges), high-priority access can be assigned, thereby improving the overall performance of the LAN.
- Source-route bridging allows for multiple bridges and parallel bridges, improving network reliability and availability, and significantly improving performance. Source-route bridging also allows less-expensive bridges to be used because the route information does not have to be stored in every bridge.
- Token ring was designed without the topology restrictions inherent in Ethernet. Single rings may span tens of kilometers, or be confined to a small office. There is no trade-off between ring size and minimum frame size either. Padding of frames is not required.
- Because token ring stations transmit idle frames when there is no data to be transmitted, each station's receiver is continually locked onto the incoming signal. Therefore, stations preparing to transmit do not need to start each transmission by sending bandwidth-wasteful synchronization characters.

Users who required highly robust LANs for their mission-critical applications eagerly embraced these and other benefits of token ring.

In the mid-1980s, when 4 Mbps token ring LANs were first being deployed on a trial basis, their reliable operation and superior loading characteristics were quickly proven and rapidly accepted. Within a few years, these first network trials evolved into sprawling campus networks interconnecting hundreds or even thousands of end stations. Users were delighted to have a network that easily met their increased demands, worked harder and performed better the more they took advantage of its full capabilities. Token ring performance remains stable and predictable, and edges ever closer to its theoretical maximum capacity under increasing network load. Laboratory measurements have shown that a 16 Mbps token ring can “…deliver 15.99 Mbps of throughput running actual application traffic.”[1]

Early network applications and deployments were hard-pressed to strain the capabilities of 4 Mbps operation at the local ring level. Only in the backbones, where traffic is concentrated, was network capacity considered an issue. By the time network demands were able to push backbone requirements to the limits of 4 Mbps token ring, IBM had developed its 16 Mbps token ring technology. With the introduction of 16/4 token ring adapters, IBM protected customer investments by providing the capability of running at either speed. This four-fold increase in capacity of a highly efficient LAN protocol continues to serve token- ring users today, satisfying all but the most demanding network deployments.

In each succeeding year since its introduction, customers have seen a steady flow of product and application developments. The following list is a summary of IBM’s investments in its token ring line of products and an example of IBM’s commitment to product enhancement and new technology:

In the past five years alone, IBM has introduced:

[1] Kevin Tolly, “The Time Is Right for Fast token ring,” Network World, 5/26/97

- Its fifth generation token ring chip
- Token ring for PCMCIA and CardBus
- Token ring for PCI
- Token ring with Wake on LAN technology
- Token ring with Alert on LAN
- Switched token ring
- Dedicated token ring (token ring that runs in full-duplex mode)
- Low memory token ring solutions for DOS and Windows
- Enhanced management via Desktop Management Interface and SNMP for IBM token ring adapters
- Numerous improvements for performance, ease-of-use, and installation improvements

Application developers have had over a decade to become comfortable with the high bandwidths available on modern LANs. As a consequence, many new applications exploit this resource. Perhaps the best example is the networking paradigm shift to the Internet and intranets. This shift has turned on its head the old “80/20 rule” that stated 80% of LAN communication should be local (within a local LAN or a department), with only 20% of the traffic going to the wider network. Today large files are commonly transferred across the network and users rely on the backbone capacity to handle it, requiring the capacity of the backbone to be greater than the 16 Mbps that shared media environments can provide.

A careful look at the full costs of networking shows that wholesale replacements of existing LAN networks are time-consuming, problematic and usually extremely costly. If a path exists, it is far more acceptable financially, and far safer from a business standpoint, to migrate to higher performance gradually and non-intrusively. That migration path has been established and is now widely accepted.

It is the innovation and arrival of modern switched token ring networks, led by IBM, which points to the future for token ring. IBM chairs the IEEE 802.5 working group and worked with ASTRAL, the Alliance for Strategic Token Ring Advancement and Leadership (which is now known as the High-Speed Token Ring Alliance (or HSTRA)) to develop token ring standards and networks for the 21st century.

Dedicated Token Ring (DTR) technology is fully compatible with end stations and concentrators used in today's shared-media token ring environments. In addition, it allows dedicated, switched access links and thus, efficient micro-segmentation of rings. Token ring switches such as IBM's 8270 Nways Token Ring switch have throughput capacities of the order of hundreds of megabits per second. The use of these switches as collapsed LAN backbones provides the network with a multi-hundred megabits per second, high speed switching fabric capable of addressing even extreme bandwidth-intensive environments. These switches can be used with existing end-station adapters and existing concentrators with no disruption to the network. Because IBM switches are non-blocking in design, they have the power and capacity to carry all the traffic that can be presented to them. This paper provides examples that show how to build on the capabilities of existing token ring network installations to increase the overall network performance dramatically as a means of satisfying the ever increasing user load for years to come.

State of the business

According to IDC there are more than 20 million token ring end stations installed worldwide; Dell'Oro quotes a larger figure of 30.5 million. Most of these installations are in Fortune 1000 companies (primarily in the financial, banking, insurance, and investment business sector). The majority of the installed base is comprised of shared media hubs: some 1.7 million ports out of the total number are switched ports, the rest being shared hub ports.

While it appears that the business is declining, it is still a large one: 1.4 million token ring ports were sold in 1999, compared with the 2.8 million ports that were sold in 1997, according to Network World. And, even though the number of switched token ring ports continues to increase, only 24% of the ports sold in the first half of 1999 were switched token ring ports, so about 1 million shared token ring ports were still sold in 1999.

The term "declining" should be qualified: the number of ports sold each year continues to fall but the size of the installed base continues to grow, and will continue to grow in the years 2000 and 2001 according to IDC.

The Dell'Oro report goes on to say that IBM sold more than 40% of all ports sold in 1999; Madge (which has acquired Olicom) was next with 34% of this business. But the report states that IBM had 51% of the token ring adapter business in 1999, a 7% increase over 1998.

It is a proven fact that token ring products are sold to customers who already have them installed in their enterprise. It is also unlikely that there will be any brand new token ring installations; instead, these products are sold to replace, upgrade or add to existing token ring installations. Customers using token ring are highly dependent on mission-critical applications and are sensitive to all outages and have little use for system downtime. Of course, this is a true statement for all customers today, but it does mean that an existing token ring infrastructure will not automatically be a candidate for replacement. There is a perception that Ethernet is both "just as good" and "substantially less expensive", but customers contemplating complete replacement of token ring by Ethernet also need to consider well-established facts such as:

- The initial attraction of Ethernet was the fact that it was less expensive than token ring on a simple "price per port" comparison.
- 16 Mbps token ring has nominally 60% more effective capacity than 10 Mbps Ethernet.
- 16 Mbps shared token ring networks actually support much more than 60% aggregate throughput than 10 Mbps shared Ethernet networks.
- The majority of existing 16 Mbps adapters can be modified (by microcode) to operate at 32 Mbps full-duplex (dedicated token ring).
- A price comparison of "price per megabit per workstation" may actually prove to be less expensive for token ring than for Ethernet.

The statistics in this section originate from IDC Corporation and Dell'Oro published reports. The actual figures here are taken from a snapshot in time and therefore may be open to discussion, but they are not really the point. The point is that a lot of existing token ring customers are going to keep token ring, perhaps using Ethernet in branch offices where it is "good enough" at a lower price per port. The business is declining, but it is still a significant one, and is also one in which IBM intends to invest and remain the leader.

All statistics derive from studies by industry analysts such as the Dell'Oro Group and IDC Corporation, and in some cases they were obtained from trade press articles such as *Network World*, which quotes these two organizations.

Technology

This section discusses the key concepts behind token ring, and provides a basic knowledge of the features described in the following chapter, "Products" on page 19.

Why Token Ring?

An understanding of the fundamental architecture of token ring can help in the understanding of why token ring continues to handle the constant increase and growth of local-area networks, and handle this traffic well. In many environments, shared token ring hubs continue to deliver mission-critical networking traffic without the need for expensive and painful upgrades.

What is Token Ring?

The IBM Token Ring network is a general purpose Local Area Network (LAN) with a star-wired ring topology, using baseband signalling and token-passing protocols conforming with IEEE 802.5 standards. Device attachments conforming with IEEE 802.2 and 802.5 standards may communicate over an IBM Token Ring network.

The token-passing protocol for ring access control is based on a predefined 24-bit pattern, called a token, which continuously circulates around the ring.

When a station has data to transmit, it waits until its station adapter receives a free token (token bit=0). Upon capturing the free token, the station creates a frame by setting the token bit to 1. It then inserts source and destination addresses, certain control information and the data to be sent to the destination station, and starts frame transmission.

During the time the frame is being transmitted, no token is available on the ring and no other station can initiate a transmission. Thus, collisions on the ring are avoided. The frame is passed (received, regenerated and retransmitted) from one station to another on the ring until it is received by a station with a matching destination address.

The destination station copies the data to its internal buffers, sets control bits to indicate that it recognized the address and successfully copied the data, and retransmits the frame.

When the frame returns to the source station following successful transmission and receipt, it is removed from the ring. The source station creates a new free token and transmits it on the ring, thereby allowing other stations access. Until the source station releases a free token, the rest of the stations are unable to transmit.

To reduce the amount of time a station has to wait for a free token, a function, known as Early Token Release is available. With Early Token Release, a sending station releases a free token following frame transmission without waiting for the transmitted frame to return. This enhances the utilization of the ring by allowing one token and one or more frames to circulate on the network at the same time.

Token ring has certain important architectural considerations:

- Token ring frames may contain as many as 17,800 bytes of information, which is significantly more than the 1500 bytes of user data in an Ethernet frame. A negotiation process is required by token ring stations to determine the maximum frame size they can use, but most implementations end up using larger frames than Ethernet implementations, which is more efficient when transferring large volumes of data.
- A single token ring may have up to 260 devices connected to it.
- Timing considerations affect the maximum frame size (only 4472 bytes for a 4 Mbps ring) and the cable lengths comprising a single ring.
- Most bridged token ring networks implement source-route bridging. Stations that want to communicate with each other across a bridged network first send discovery frames through the network to discover the best route. The stations then save this routing information and include it in every frame they subsequently send. The token ring frame itself contains an indicator bit to denote that this source-routing information is present in the frame. Source-route bridging is especially useful in SNA networks, because it allows multiple paths between two points across the network and also allows configuration of duplicate MAC addresses, both of which increase reliability and availability of the network. Maybe this has led to the perception that SNA and token ring are somehow inextricably linked, but the facts are that:
    1. Token ring networks can implement transparent bridging; it's just less common for them to do so.
    2. Whatever bridging method is used has no direct impact on the layer 3 traffic being transported in the frame: SNA, TCP/IP or any other protocol.
- Hubs which provide *active* token ring ports can effectively increase the total cable length of a single ring because they re-drive and re-synchronize the signals on these ports.
- There are some minor differences between the IEEE 802.5 and the IBM Token Ring standards, but for all practical purposes they are identical. IEEE 802.5 is the formal standard based on IBM's original implementation.

Token Ring versus Ethernet

There is no simple and objective answer as to which technology is “better”; any comparisons between token ring and Ethernet end up being emotional and subjective. The only points which need to be made here are:

1. Shared media token ring copes much better under heavy load than shared media Ethernet. As traffic increases, Ethernet performance degrades significantly whereas 16 Mbps token ring has been shown to be capable of delivering an aggregate data throughput of 15.9 Mbps. The vast majority of end users, however, do not themselves require more than 10 Mbps of network bandwidth, and a switched Ethernet infrastructure delivering switched 10 Mbps dedicated bandwidth to each user can meet their needs. What this means, though, is that existing 16 Mbps shared token ring users may well be able to continue using their existing networks when 10 Mbps shared Ethernet users are required to upgrade to switched Ethernet or 100 Mbps Ethernet.
2. Many existing token ring users, most especially those with S/390 mainframes, will want to remain with token ring because of its ability to support duplicate identical MAC addresses. Duplicate addresses are used on mainframe token ring gateways to provide a measure of automatic load-balancing and

switch-over in the case of failure, which requires the implementation of a source-route bridging network which is not normally possible to implement with Ethernet LANs.

Shared media networks

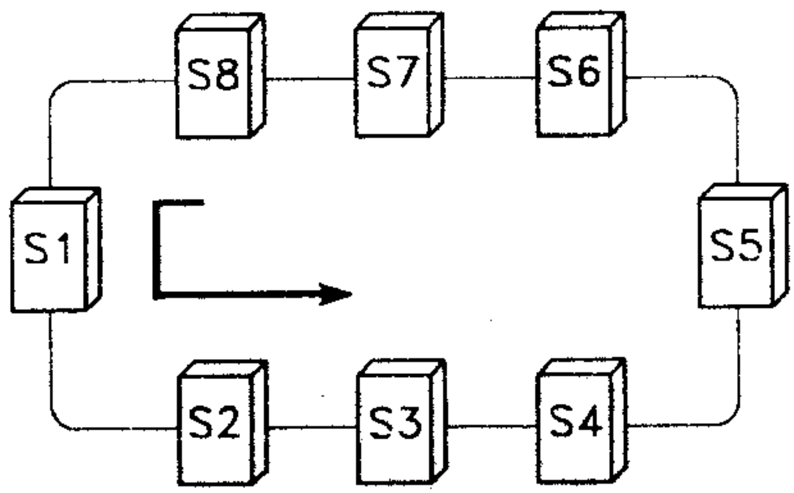

Figure 1. Single token ring with eight ring stations

In its simplest form a single token ring can be viewed precisely as above: a ring to which all network stations attach. The adapters on the network stations themselves provide the intelligence in the network: an "active monitor" is elected and then serves to monitor the status of the ring, for example.

Unmanaged hubs

The simplest token ring hub does not even have to have a power supply. The network adapters themselves provide all the intelligence in the network, even to the extent of determining the ring speed. A complete ring is established by "daisy chaining" hubs together: a cable connects the "ring out" port of one hub to the "ring in" port on the next one, with the last hub connecting back to the first one to complete the ring. A complete ring can be implemented on just one hub since the ring-in and ring-out ports will be wrapped internally to complete the ring.

Most unmanaged hubs will have a power supply, but this will only serve to illuminate indicator LEDs in the hub itself.

Even an unmanaged hub such as the 8228 (see "8228 Token Ring Network Multistation Access Unit (MAU)" on page 32) provides some resilience; if a cable between one hub's ring-out port and the next hub's ring-in port should break or be removed, the 8228 will *wrap* the connection internally and a single ring will be preserved, although this recovery requires that both ends of the failing cable be removed from the hubs manually.

Managed hubs

Managed hubs differ from unmanaged hubs in two ways:

1. Their configuration can be specified and changed from a management station.
2. They report status and error indications to the management station.

For more information on the subject of network management see “Network management” on page 47.

Bridges

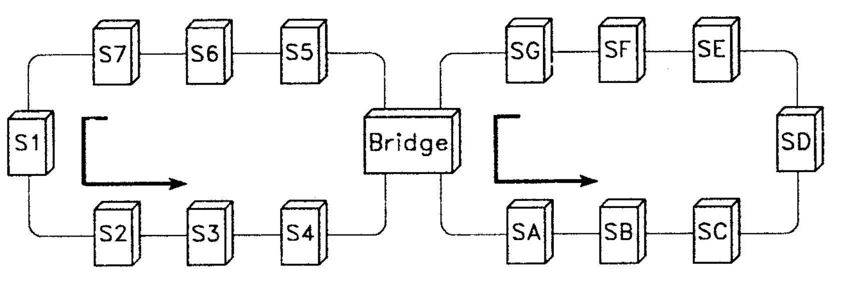

Figure 2. Two rings connected using a single bridge

The original bridge definition for the token ring was for a source routing bridge. This differs from the transparent bridge operation normally seen on Ethernet networks; transparent bridge operation requires that the bridges maintain tables of MAC addresses whereas in source routing each frame carries information about the route it is to follow through the network. In a source routing environment, this routing information is acquired through a search process that originates at the source station, using commands such as TEST or XID (Exchange ID).

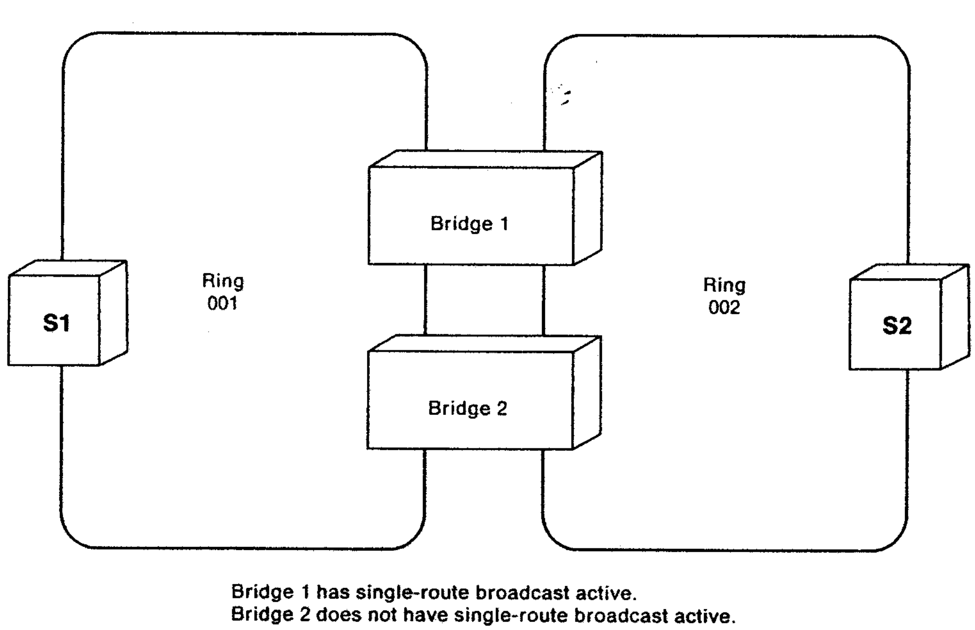

Figure 3. Parallel bridge configuration

Source-route bridging allows the configuration of parallel bridge paths through the network, which has to be blocked in a transparent bridging environment. It allows for multiple (redundant) routes through a network, with automatic discovery of the quickest path through the network without the necessity of maintaining a record of network topology (which may change rapidly requiring extensive updates). There are several ways that this route discovery process can proceed, and the following paragraphs describe two of these methods.

Consider a set of interconnected (bridged) token rings, of which a very simple example is shown in Figure 3 on page 12, where node S1 is ready to transmit a series of messages (or set up a session with) to node S2 (which may be in another city or another country). The location of the ring that includes S2 as well as the interconnections of the various rings are initially unknown to S1.

In the first method, S1 will send an ARB (All Routes Broadcast) frame to S2. This frame or frames will cross all available bridges in the network on the way to S2 (unless the “hop count” expires), thus exploring all possible routes between the two nodes. It is perfectly acceptable to have two bridges actively forwarding frames from any source ring to the same destination ring simultaneously because the bridges will have different bridge numbers assigned (1 and 2 in our example). While doing this, each frame records its route in the Routing Information Fields included in the 802.5 frame header. (Each RIF contains a source ring number, bridge number, and destination ring number). Upon receiving each frame, S2 will send a specific unicast reply back to S1 with a bit set in the header indicating that the RIF route is to be inverted (traveled backwards). Now, this is where it gets interesting! S1 will receive these replies and discard all but the first (which, by definition, is the one that won the race and which has traversed the quickest and most optimal path through the network). It can then use the routing information stored in the RIFs for future messages to that destination (generally for the duration of the session). It has found the most efficient path through the network without having to be aware of its topology.

In the second method, S1 can send an SRB (Single Route Broadcast) frame to S2. This frame crosses only a designated set of bridges (only bridge 1 in our simple example) resulting in only one copy of the frame appearing on any given ring in the network. When the SRB frame gets to S2, S2 will reply using ARB, which will generate multiple frames (with the route inverted) which will explore all of the possible return paths, with the first frame returning to S1 again being the fastest. The result is the same - the source node has discovered the most efficient path between itself and the desired destination node without having to know where it is or how the network was set up (or changed).

Source-route bridges still require the configuration of a spanning tree, which represents the network topology, in order to handle single-route broadcast frames. This will normally be handled by having the bridges communicate with each other and dynamically maintain a single-route broadcast path through the network, just as is normally the case with transparent bridges. Many broadcasts in a token ring environment are all-routes broadcast frames, which will be forwarded by all bridges without consideration of spanning tree topology. If appropriate, as an alternative to implementing any kind of spanning tree configuration, bridges can be configured simply to discard single-route broadcast frames. Bridges need to be configured with a bridge number which has to be a unique number for the rings to which it attaches; ring numbers also need to be defined and are usually configured in one or more of the bridges as well.

Priority and class of service

Token ring implements *access priority* by including three bits of priority information in every token or LAN frame. A ring station can transmit a frame of a given priority using an available token with a priority less than or equal to that of the frame. If no token is available (because another station is transmitting data) then the ring station may reserve a token of the required priority by setting bits in a passing frame. The next free token will then be sent at that priority and will not be used by any other ring station, allowing our station to transmit a frame at the required priority. Two stations may not make a reservation for the same priority value, but a station may make a higher-level priority reservation than another station's existing priority reservation and effectively pre-empt it.

Token ring access priority relates to two emerging and complementary layer 2 standards which can be implemented in a bridged environment:

a. 802.1p, which is a component of the 802.1D standard for LAN bridges. 802.1p implements a priority queueing mechanism inside a LAN bridge which allows certain LAN frames to be given higher forwarding and transmission priority than other LAN frames. The access priority in a received token ring frame may be mapped to an internal priority value inside the LAN bridge, and could be used - for example - to allow voice traffic to "overtake" data traffic inside a bridge.

b. 802.1Q, which provides the ability of all LAN media to include user priority information as part of the MAC frame. So a translational bridge could map token ring access priority in received frames into a priority marking for frames transmitted over Ethernet using 802.1Q.

For more information on these two standards, see *Application-Driven Networking: Class of Service in IP, Ethernet, and ATM Networks*, SG24-5384.

On the face of it, this priority value seems to map nicely to the 802.1Q standard for LAN frames. 802.1Q adds 3 bits of priority information to a LAN frame, and a related standard (802.3ac) allows the maximum frame size for Ethernet frames to be increased to allow the addition of this information without the need to reduce the data payload in the frame. So a growing implementation and use of 802.1Q (such as in Windows 98 and Windows 2000) even on Ethernet networks might seem to allow the priority indication in the LAN frame to map directly to this access priority on token ring.

In reality, things are not quite as simple. The default action of most token ring bridges and switches is to forward user frames with an access priority value of x'100', and the highest value, x'111', is reserved for vital network management frames such as "Active Monitor Present" frames. Although the two next higher values of access priority, x'101' and x'110', are described as "reserved" by IBM, these are in fact the values which need to be used for more important traffic such as mission-critical data traffic and voice traffic.

What this means, though, is that the token ring access priority can be used to differentiate between different types of user traffic, which can be placed in separate transmission queues by token ring devices, including bridges which implement 802.1p along the following lines:

Table 1. Access priority and traffic types

Priority bits	Traffic type
x’000’	Normal data traffic
x’001’	Not used
x’010’	Not used
x’011’	Not used
x’100’	Normal data traffic forwarded by bridges and switches
x’101’	Time-sensitive data traffic (SNA traffic perhaps)
x’110’	Real-time critical traffic (voice traffic, for example)
x’111’	Station management

In order to support this sort of quality of service mechanism, LAN bridges and switches need to support both the following IEEE standards:

- 802.1Q to preserve packet priority values *across the entire network*
- 802.1p to implement multiple transmission queues to enable higher priority packets to be sent ahead of normal priority packets during congestion

Neither of these standards is required to implement “class of service” on a hop-by-hop basis, and this function is implemented for IP traffic by IBM token ring adapters.

The use of token ring access priority in this way is just not possible on Ethernet LANs where all stations contend equally for use of the shared medium. 802.1Q provides a means of indicating a frame’s priority across the entire network but plays no part in prioritizing the transmission of LAN frames over a single LAN segment; this is something only token ring can do.

Switched networks

Switched LAN networks are implemented by LAN switches in which traffic is not sent to all stations which attach to each port in the switch. LAN switches learn MAC addresses in much the same way as transparent bridges do, and use this information to transmit specifically-addressed LAN frames to only those ports over which it is necessary to transmit them. LAN switches do have to transmit broadcast frames (those with MAC address “FFFF FFFF FFFF”) over all ports.

LAN switches originated in the Ethernet world, where there was the greatest need to increase network capacity coupled with the relative simplicity of the Ethernet design. Now LAN switches also exist for token ring networks.

Backbone, workgroup, desktop

On one level, LAN switches *are* translational bridges, and they can be used in the same way as bridges to segment a single token ring LAN into multiple LAN segments. One way in which switches differ from bridges is that they can operate in “cut through” mode, in which a LAN frame is being transmitted on an outbound port before it has been received on an inbound port. Bridges need to receive a frame in its entirety before re-transmitting it, but this does have the advantage that the frame check sequence (FCS) at the end of the received frame can be verified. If the FCS is found not to match the frame itself, the bridge will discard the frame. The switch which doesn't check FCS will needlessly transmit a corrupted frame, wasting network bandwidth; at some point further in the network the FCS *will* be checked and the frame will ultimately be discarded anyway.

Because of the reduction of traffic on each port of a LAN switch, the ultimate design point is to have end stations themselves directly attached to a switch. Under these circumstances there is little contention for use of the connection: the only traffic apart from broadcast traffic flowing between the end station and the switch is real traffic to or from the end station itself.

Because of the extra cost and implementation effort/disruption, a mid-point in a network designed around LAN switches will be on in which the important servers in the network are connected directly to the LAN switches but the less demanding end users will continue to connect to shared LAN hubs.

Full-duplex dedicated Token Ring

If an end station is directly connected to a LAN switch in an Ethernet network, it knows that no other end stations will be attempting to use the connection and it therefore no longer has to attempt to detect their presence before transmitting data as it would have to do in a shared hub environment. It also allows the switch and end station to "break the rules" and implement a full-duplex connection where both devices can be transmitting data to each other simultaneously. This is not permitted in a shared Ethernet environment.

Token ring is different, but a similar enhancement is possible by the use of dedicated token ring (DTR) connections.

In the standard RJ-45 token ring implementation ("Classic Token Ring" or CTR) a connection between a device and a hub/switch is made with two twisted pairs of wires. Although one pair is used for the receipt of data and the other used for the transmission of data, the standard token-passing protocol (TKP) means that these will never happen simultaneously: a token or frame will first be received by the station and then a token or frame will be transmitted around the ring.

DTR provides a dedicated connection between end station and switch, but TKP can still be used over it, in which case tokens and frames flow over a two-port LAN. Little advantage has been gained here, but this mode of operation is required in DTR devices for downward compatibility. Of significant advantage, however, is the TKI (Transmit Immediate Access) protocol, where tokens are no longer used and both devices can transmit at will.

The use of TKI allows two DTR stations to transmit and receive in "full duplex" mode, which doubles the bandwidth available to them to 32 Mbps. Of course, this means that 16 Mbps of bandwidth is simultaneously available in both directions and not that 32 Mbps is available in any one direction.

Unlike Ethernet adapters, most token ring adapters can be configured for DTR operation with, at most, a microcode upgrade. It offers a significant upgrade to network servers.

ATM, MSS, MPOA

ATM can be used as a high-speed backbone connection between token ring switches, typically using a 155 Mbps up-link from the switch and defining an emulated token ring LAN across the ATM network. IBM's Multiprotocol Switched Services (MSS) Server can be used to define the emulated LAN, and Multiprotocol Over ATM (MPOA) provides one way of defining and using short-cut direct connections across complex ATM networks.

Higher speeds

What technologies exist and are under development for the transport of token ring traffic at speeds greater than 16 Mbps?

100 Mbps and gigabit Token Ring

IBM continues to chair the IEEE 802.5 committee in the development and ratification of token ring standards. In particular, the committee developed the 802.5t standard for 100 Mbps token ring, which is now an approved IEEE standard. The 802.5v standard for gigabit token ring is going through the final balloting stages (as of March 2000), and should be approved as a standard by July 2000. The purpose of these standards is to ensure interoperability between different vendors' high speed token ring implementations. IBM supports 100 Mbps token ring on its 100/16/4 PCI adapter.

IBM does not view high-speed token ring as a requirement for the majority of its customers, and therefore the decision has been made not to provide 100 Mbps high-speed token ring uplinks on its products - primarily the 8270 token ring switch. That is not to say, however, that IBM does not provide high-speed uplinks at all; the rest of this paper will serve to describe how FasTR, TokenPipe and ATM can each be used to provide 100 Mbps or faster uplinks from IBM's token ring switch. High-speed server attachment can be made using one or more dedicated token ring connections between a server and LAN switch, each offering 32 Mbps of full-duplex network bandwidth.

More information on high-speed token ring can be found at the Web site of the High-Speed Token Ring Alliance, http://www.hstra.com.

FasTR

FasTR is a special name for a special implementation; it's a form of high-speed token ring which is actually implemented using ATM adapters on the 2216 router, the 8270-800 switch, the 3-slot 8272 blade in the 8260 hub and the Multiaccess Enclosure (MAE) inside the 3746-9x0.

FasTR encapsulates the token ring frames using ATM AAL5; it is intended to be used *without* the use of any actual ATM network but simply as a point-to-point link using singlemode or multimode fiber between two devices. It *can* be established using an ATM permanent virtual circuit (PVC) but in all likelihood it would not make much sense to do so; if there's really an ATM network then the adapters should be configured as real ATM adapters and not as FasTR adapters.

FasTR can be used to relieve congestion in a token ring backbone connection to a mainframe server. It is configured as a token ring to the devices which support it. Tests have shown throughput of greater than 100 Mbps using 4,000-byte packets.

FasTR is actually supported on two basic adapters only: the MSS Client UFC for the 8270-800 and the 8272 blade, and the ATM-155 adapter for the 2216 and MAE.

The devices which support FasTR can simply be used for source-route bridging, or they can configure themselves as IP routers and use it as a fast router-router link. In addition, the 2216 and MAE can route APPN traffic over FasTR links.

TokenPipe

A TokenPipe is another type of high-speed backbone connection between token ring switches; it is a collection or bundle of between 2 and 4 connections between two token ring switches.

Each of the connections in a token-pipe is a 16 Mbps connection and is capable of operating as full-duplex (DTR, TKI). This gives a theoretical maximum capacity of a TokenPipe of 128 Mbps (16 x 2 x 4).

The ports in a TokenPipe should be configured in ascending order of the actual port number on the switch; the lowest-numbered port is known as the *primary* port and is used for transmission of all broadcast frames and spanning-tree frames.

Traffic on any one connection between devices on separate token ring switches which are connected using TokenPipe will flow on just one of the links which comprise the TokenPipe. The switch will use its internal table of MAC addresses to assign each target address to one of the TokenPipe links, but it will assign many addresses in a round-robin or similar fashion to achieve load balancing across the set of links.

If a single link configured as part of a TokenPipe should be disconnected or fail for any reason, the entire TokenPipe is made inoperable.

Load balancing

Other solutions for higher speed token ring connectivity include load balancing between multiple token ring LAN adapters. This approach is one that should not be undertaken lightly: there are many complications possible using multiple LAN adapters connecting to the same switch or backbone LAN segment. One possible approach would be to configure multiple LAN adapters with different IP addresses and use layer 3 load balancing techniques such as OSPF equal-cost multipath to spread the load across multiple LAN adapters; this requires some kind of OSPF “gated” implementation in the machine. Multiple LAN adapters can cause problems in NetBIOS environments because of the requirement for a computer’s name to be unique in a network.

Load balancing approaches are fraught with danger and can often, at best, lead to a configuration in which only one LAN adapter is used for the significant majority of network traffic anyway.

Products

This section lists the current IBM token ring product range: the switches, adapters, hubs and associated products which either provide token ring connectivity or use token ring themselves. A brief technical overview of each product is given with some of the key features.

Each section concludes with a table showing model numbers and feature codes which remain orderable.

The chapter concludes with a brief discussion on products which are no longer marketed by IBM but which may still be seen in customer environments.

Servers

S/390

Token ring connectivity to the S/390 mainframe has been available for many years now: support for token ring was announced for the 3725 communications controller back in 1986, for example. What follows is a brief overview of the IBM products that provide token ring connectivity for the S/390 mainframe server.

3745

The 3745 Communications Controller is the latest in a range of controllers dating back to the early 1970s, and from which large SNA networks were built in the 1970s and 1980s. It is mainly seen today in computer data centers and provides many different types of network connectivity for the mainframe environment. 3745 customers pay a monthly license fee for the software (usually Network Control Program - NCP), which runs in the 3745. A single 3745 can contain up to 8 token ring interfaces.

3746

The 3746 Model 900 is attached to a 3745 and provides additional processing power to “off-load” the 3745’s central processor. Many customers invested in the 3746-900 because of its token ring processing capabilities; a fully-configured 3746-900 can provide an additional 32 or 33 token ring ports, but its real benefit is in the ability of the combined 3745 + 3746 to handle much greater traffic volumes and number of SNA sessions as on the 3745 alone. The 3746 Model 950 is the same unit “stand-alone”, in other words not connected to a partner 3745.

OSA

A single S/390 Open Systems Adapter (OSA-2) is now provided automatically in the base configuration of any S/390 processor bought from IBM, and by default this adapter will be the “ENTR” version which provides two LAN ports, each of which can be configured as Ethernet or token ring ports. Each OSA-2 port will automatically determine the LAN type to which it is initially connected (Ethernet or token ring) and then sense token ring speed (4 Mbps or 16 Mbps) and duplex operation (full-duplex or half-duplex). A single S/390 mainframe can support up to 24 token ring OSA-2 ports.

3174

The 3174 started life as a controller for 3270 screens and printers in SNA networks. The 3174-63R is the only model now shipped with a token ring adapter, and connects 3270 screens and printers to a token ring LAN.

Other models of 3174 can be fitted with a token ring adapter (feature #3044) to serve the same role. Additional memory and upgraded software may be required in the 3174 to support the new configuration. Note, however, that no token ring adapter can be installed in the 3174-91R model.

3174s may also be found configured as token ring gateways. There are two features for providing a 3174 token ring gateway: features #3026 and #3044. Both in fact provide the same hardware - a 16/4 token ring adapter. Feature #3026 included microcode which supports a limited subset of gateway features but was withdrawn in 1998; feature #3044 requires that the customer purchase microcode separately and may also require a memory upgrade on the 3174.

The token ring gateway is rarely used these days; many customers used it in the past as a relatively low-cost method of connecting token ring LANs to S/390 mainframes for SNA traffic.

The 3174 can also act as a Telnet client over TCP/IP networks: this allows existing 3270 terminals to log on to UNIX or other IP host systems. Again, this requires appropriate microcode, which might also require a memory upgrade for the 3174.

The following table shows the feature codes applicable to the different models of the 3174 that are likely to be required for token ring use:

Table 2. Subset of 3174 models and features

Model / feature	Description
3174-21L	Rack-mounted Establishment Controller with parallel channel adapter
3174-21R	Rack-mounted Establishment Controller with V.24 or V.35 adapter
3174-22L	Rack-mounted Establishment Controller with ESCON adapter
3174-61R	Shelf Establishment Controller with V.24 or V.35 adapter
3174-63R	Shelf Establishment Controller with token ring adapter
3174-91R	Small shelf Establishment Controller with V.24 or V.35 interface
#1014	2 MB memory expansion (to a total of 4 MB)
#1016	4 MB memory expansion (to a total of 6 MB)
#1048	Second diskette drive (required for CS-C microcode)
#1056	20 MB disk drive (alternative to #1048 for CS-C)
#3044	16/4 token ring adapter
#3026	Token ring network gateway (for 3174-61R only); withdrawn in 1998
#5010	Configuration Support-B microcode
#6010	Configuration Support-C microcode for -21L, -21R, -22L models
#6015	Configuration Support-C upgrade for -21L, -21R, -22L models
#6060	Configuration Support-C microcode for -61R, -63R models
#6065	Configuration Support-C upgrade for -61R, -63R models

Stop press

IBM has announced that the 3174-21R, 3174-61R, 3174-63R and 3174-91R will be withdrawn from marketing with effect from June 30, 2000. The only models which remain are those with a direct channel attachment to the S/390 mainframe, which is still required today to provide an operating system console.

AS/400

Token Ring adapters

There are different token ring adapters for different AS/400 models based on the different architecture and configuration of the AS/400. Some of the latest AS/400 models provide a PCI bus driven by a Multifunction I/O Processor (MFIOP). Each adapter type has two feature numbers. The #9xxx feature number indicates the identical component but one which is included as part of an overall AS/400 package rather than being a separately-charged feature. The different AS/400 token ring adapters include:

Table 3. AS/400 token ring adapters

Feature	Description
#2724 / #9724	PCI 16/4 Mbps Token ring IOA
#6149 / #9249	16/4 Mbps Token ring IOA
#2619 / #9619	*16/4 Mbps Token ring Adapter/HP*
#2626	*16/4 Mbps Token ring Adapter/A*

The first two features in this table, #2724 and #6149, both support full-duplex dedicated token ring and should be ordered in preference to the last two when possible.

5494 and 9401 Model 150

The 5494 remote control unit has been withdrawn; it was used to provide remote token ring attachment for AS/400 5250 twinaxial terminals. The reason it has been withdrawn is because the AS/400 9401 Model 150 is now available at much the same cost. This is a “portable” AS/400 system unit and a token ring adapter is included in the base order for Models #0593 and #0594; up to 28 twinaxial terminals can be connected to the 9401.

RS/6000

There are two token ring adapters currently sold for the RS/6000 family, one of which connects to the Micro Channel bus in RS/6000 SP machines. Both adapters are capable of full-duplex dedicated token ring operation, possibly requiring a microcode update on feature #2972.

Table 4. RS/6000 token ring features

Feature	Description
#4959	IBM Token ring PCI Adapter
#2972	Auto Token ring LANStreamer 32 MC Adapter (RS/6000 SP only)
#2970	Token ring High-performance Network Adapter (withdrawn; RS/6000 SP)
#2920	PCI Auto LANStreamer Token ring Adapter (withdrawn)

Netfinity

There are no unique adapters or special considerations for Netfinity servers; they use the same token ring adapters discussed in “Desktop” on page 23. The 100/16/4 PCI token ring adapter is particularly suited for use in a Netfinity server, and Redundant NIC support (on certain operating systems) can be used to increase server availability.

Server summary

Figure 4 on page 23 depicts one of the most attractive features of using Token Ring technology: the ability to attach any/all of IBM’s host server systems.

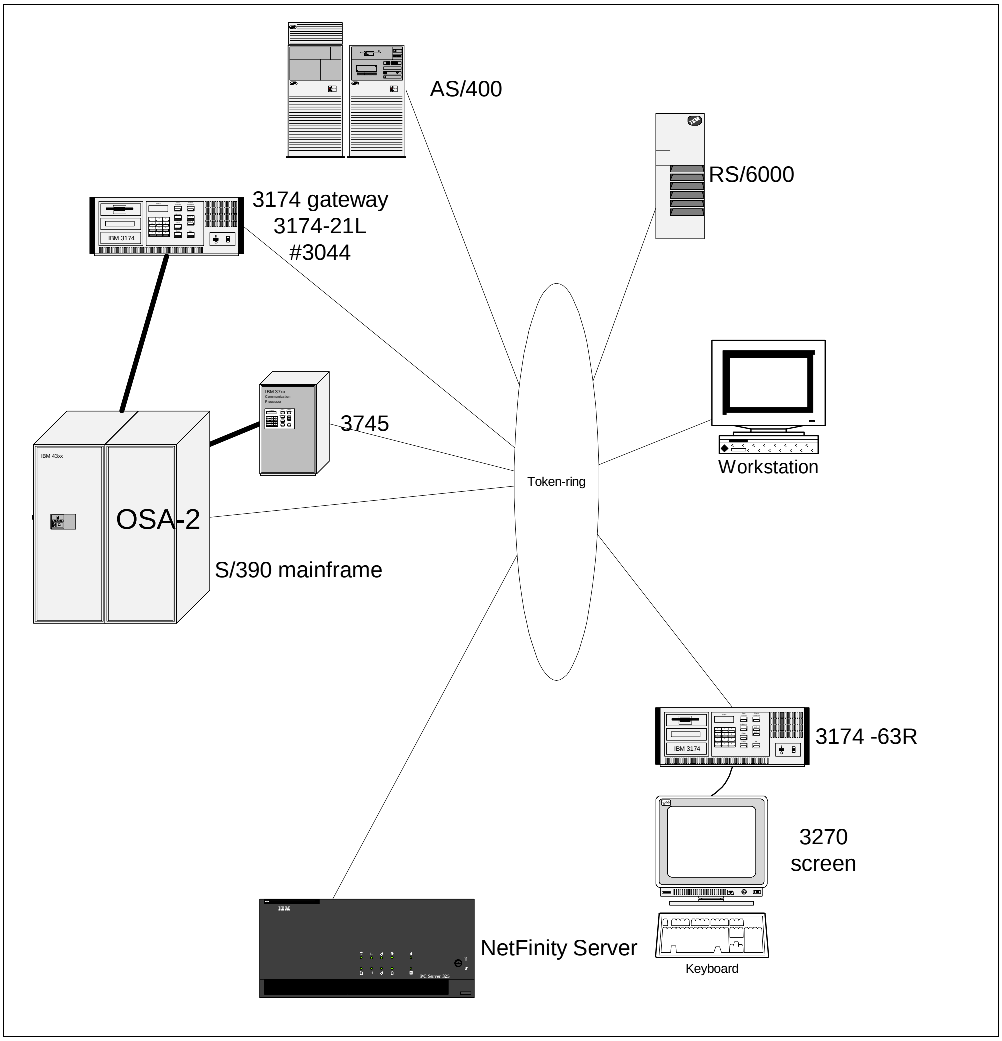

Figure 4. Servers all together using token ring

Desktop

This section describes the IBM token ring adapters available for PC workstations and Netfinity servers. The technology section introduces the various functions provided on each of the products and the second section describes the products and which of the previously-described functions are supported on each.

Technology

IBM Token Ring adapters implement some or all of the following functions:

Wake on LAN

A token ring adapter's Wake on LAN feature requires a PC planar board that supports Wake on LAN as well. This support either requires the use of a 3-pin connector between the adapter and the motherboard or else a motherboard which recognizes when the adapter asserts the PME (Power Management Event) signal on a PCI bus; all IBM token ring adapters which support Wake on LAN support both of these methods of operation.

When a PC which supports Wake on LAN is powered off, provided that the power cord remains attached and connected to a live power supply, an auxiliary power supply continues to power the token ring adapter. Although the adapter will initially de-insert from the token ring, it will then reset itself and re-connect to the token ring. In this state it will repeat token ring frames in just the same manner as any other token ring adapter connected to the ring.

When in this auxiliary-powered state, the Wake on LAN token ring adapter monitors LAN traffic looking for special wake-up frames. These frames can either be broadcast or unicast frames at the MAC level, but they contain a data field which starts with 6 bytes of x'FF' followed by the actual MAC address of the Wake on LAN adapter repeated at least 8 times. When the token ring adapter detects this sequence it turns the PC on, just as if the user had switched it on using the On/Off switch.

Wake on LAN can be used in a variety of different ways: to upgrade software on unattended PCs in the middle of the night, or to turn on printers and servers remotely at the start of the day. Additional software (not related to the LAN adapter) can also be used to shut down PCs completely, so the two functions can be used together to control devices remotely and save power.

Alert on LAN

Alert on LAN was developed jointly by IBM and Intel; it requires a Wake on LAN adapter and it provides the capability for a token ring adapter to send an alert across the network even when the PC itself is powered off. It can be used to notify a network administrator (using a pager, perhaps, or a cellular telephone) if a PC is tampered with: if the case is removed in an attempt to steal some of the PC's contents, for example. It can also be used to send an alert when a PC's configuration is changed or when the power-on self-test (POST) fails, for example, if the PC's hard drive fails.

Alert on LAN can be used in conjunction with Wake on LAN. For example, Alert on LAN can report a fault and Wake on LAN can then be used to reconfigure and reboot the machine remotely.

Alert on LAN is part of IBM's Universal Manageability Initiative (UMI), and when used in conjunction with IBM's UM Services, Alert on LAN can be integrated into most network environments, including Tivoli, Microsoft SMS and Intel LANDesk.

DMI

The Distributed Management Interface (DMI) was one of the first results of the activities of the Distributed Management Task Force (DMTF). DMI defines a standard framework for managing and tracking components in a computer. All IBM's token ring adapters are compliant with DMI standards.

The DMI architecture includes a Service Layer, a Management Information Format (MIF) database, a Management Interface (MI), and a Component Interface (CI). The DMI Service Layer acts as an information broker between manageable products and management applications. The MIF database defines the standard manageable attributes of PC and server products.

The MI allows DMI-enabled management applications to access, manage and control desktop computers, components, and peripherals, while the CI allows components to be seen and managed by applications that call the DMI Service

Layer. The CI gets real-time dynamic instrumentation information from manageable products and passes it to the MI via the Service Layer. It shields component vendors from decisions about management applications, allowing them to focus on providing competitive management features and functions for their products.

Alert on LAN monitors the system variables measured by a DMI-compliant manageability chip: temperature, voltage and chassis intrusion, for example. Alert on LAN will generate an alert when any of these variables falls outside an allowed range.

SNMP

The Simple Network Management Protocol (SNMP) has grown to be the main standard for management and monitoring of network equipment such as hubs and routers. It is used to retrieve status and configuration information from devices (GET commands), send information and configuration changes to a device (SET) and for unsolicited alerts (TRAPs) to be sent to a network management station. It is similar in concept and purpose to DMI, but comes from a different product set. A PC workstation may implement agents for both DMI and SNMP simultaneously, and a management station may manage both at the same time, but there is no inherent interoperability between the two: they act in parallel but not together.

LAN Adapter Management Agent

The LAN Adapter Management Agent provides support for DMI and SNMP agents and therefore makes IBM LAN adapters visible to network management applications using either or both protocols.

Redundant NIC

Redundant NIC provides a high availability solution for your NetWare 4.11, Windows NT 3.51 or 4.0 server using a token ring network connection. The Redundant NIC function ensures that network connectivity is maintained in the event of a NIC-related failure. The user can assign a backup NIC to take control of the network connection upon failure of the active NIC.

The benefits of Redundant NIC include the ability to:

- Maintain your network connection after a NIC failure such as a cable fault or a hard error on the adapter.
- Remove the need to reboot your server to re-establish connectivity to the network.
- Minimize network connectivity downtime for critical server applications.
- Create multiple Redundant NIC pairs for complex server environments.
- Provide LAN Adapter Management Agent support to monitor Redundant NIC status.

Route switching peer/client

Route switching improves IP performance in certain network environments: those in which a layer 2 connection exists in parallel with a layer 3 connection. One environment where this applies is a one-armed router in which two different devices may be configured on different IP networks but are in fact connected to the same physical LAN. Another environment is where an ATM connection exists somewhere in the connection path between two devices.

Route switching peer applies to the first configuration: a one-armed router. When two devices which both have route switching enabled in peer mode but are in different IP subnets attempt to communicate, they will exchange layer 2 information. If they find that a layer 2 connection can be established between them, they will use this for transport of layer 3 IP traffic. This means that although traffic is logically being routed through the one-armed router, it is in fact being transmitted directly between the two devices over a shortcut. The initial exchange of capability information is conducted using IP layer 3 packets. This means that any access controls or filters in the router can still be used to prevent communication between these devices. In other words, this feature cannot be used to bypass normal router security features.

Route switching client is more specialized and requires an appropriately configured MSS Server acting as an interface between token ring and ATM networks. Again, a layer 2 connection can be set up between two route switch clients across an intervening ATM network. The actual data transfer then uses switching across the network and does not require handling by an intervening router. Note that route switching peer does not require the presence or use of MSS, MSS Client or MSS Domain Client anywhere in the network.

Route switching only functions between two “like” LAN adapters, meaning that both must be token ring adapters and both must be using either source-route or transparent bridging. In particular, route switching is not possible between token ring and Ethernet devices, even if both have the function enabled.

Most token ring drivers allow the configuration of “automatic” mode, which means that the adapter will attempt to operate both as a peer and a client and will adopt the first mode of operation at which it is successful.

Implementation of route switching is performed by the token ring adapter’s device driver and is therefore operating-system dependent. See Figure 17 on page 44 for an example of a configuration that implements route switching with a one-armed router.

Class of service

IBM token ring adapters provide an implementation of a class of service mechanism for IP traffic which allows the mapping of layer 3 IP settings to token ring’s access priority. Access priority is discussed in more depth in “Priority and class of service” on page 14. “CoS for IP” allows certain types of IP traffic to be selected based on UDP or TCP port numbers and then assigned a specific token ring access priority which is different from the default value.

One application of CoS for IP would be in a workstation which is transmitting voice traffic encapsulated in IP: the adapter would be configured to identify this traffic by the range of UDP ports used (which may be the range 16384 to 32767) and would assign a token ring access priority of 6 (the highest possible value) to this traffic.

Troubleshooting Utility

Troubleshooting Utility is an application that you can use to collect information about your system to help diagnose problems with your token ring PCI adapter. It is available for Windows 95 OSR2, Windows 98, Windows NT 4.0, and Windows 2000. Troubleshooting Utility helps you identify and solve network problems that involve the adapter. It analyzes your system and all token ring PCI adapters in your system. Using the available information, it offers suggestions for solving your networking problems. Troubleshooting Utility also generates an adapter information report that provides detailed information about your system to enable your network administrator or other support personnel to efficiently assist you. You can also use it to create a Diagnostic diskette.

Managed Driver Upgrade

Managed Driver Upgrade is for Windows 95, Windows 98, Windows NT, and Windows 2000 operating systems that assists local users and system administrators in upgrading their systems to the latest level of the driver. The package consists of the new driver, all necessary support files, and an executable program that can upgrade the driver without any user intervention. Managed driver upgrade can be run locally by the user, or can be distributed remotely by a system administrator using management software such as Tivoli TME 10 Software Distribution.

BootROM

In the past, we referred to this feature simply as Remote Program Load (RPL) or Remote Initial Program Load (RIPL). The RPL feature enables an adapter to boot and install a computer using files that the computer receives from a LAN server. RPL/RIPL was initially only a NetBIOS-based protocol (nonroutable). However, recently the industry has adopted a (routable) TCP/IP-based version for accomplishing essentially the same thing, known as DHCP/PXE. Since these cards are capable of performing either traditional NetBIOS-based RPL or the newer TCP/IP-based DHCP/PXE, it's a bit misleading to continue to refer to the feature simply as “RPL”. So, we are now referring to it as the BootROM feature, and the BootROM feature can support either traditional RPL or the newer DHCP/PXE, depending on how you configure the BootROM to behave. Details are documented in the *IBM Token ring Adapter Features d*ocument.

Products

The complete current range of IBM Token Ring adapters follows. Where support for a particular function is noted on a particular adapter, this should be taken to mean that the function is supported in conjunction with necessary hardware (such as appropriate planar board) and software (such as appropriate operating system). These co-requisite requirements are noted in the description of the function itself in the previous section and in the operating system compatibility matrices that follow the description of the adapters themselves.

Part numbers given below are for the “Single Pack (SP)” adapters which are packaged with documentation and device drivers on CD-ROM, and “Card Packs (CP)” which are single adapters with no documentation; card packs can be ordered as single units but they are designed to ship in an overpack which holds 10 adapters and therefore order quantities should normally be in multiples of 10.

IBM 16/4 Token Ring PCI Adapter 2

Part number 34L0601(SP), 34L0610(CP): requires a computer with a PCI Version 2.0 or later 32-bit or 64-bit busmaster-enabled slot. Supports auto-sensing ring speed (4 Mbps or 16 Mbps) and full-duplex dedicated token ring (DTR).

IBM 16/4 Token Ring PCI Adapter 2 with Wake on LAN

Part number 34L0701(SP), part number 34L0710(CP): requires a computer with a PCI Version 2.0 or later 32-bit or 64-bit busmaster-enabled slot. Supports auto-sensing ring speed (4 Mbps or 16 Mbps) and full-duplex dedicated token ring (DTR). Supports Alert on LAN in addition to Wake on LAN.

IBM 16/4 Token Ring PCI Management Adapter

Part number 34L5001(SP), 34L5010(CP): will be available from 28 April 2000. Supports Wake on LAN and PXE 2.0-compliant remote boot. Supports auto-sensing ring speed (4 Mbps or 16 Mbps) and full-duplex dedicated token ring (DTR). Ships with the Tivoli Management Agent and provides support for SNMP and DMI 2.0. Has received the Microsoft PC99 logo and has been certified for use with Novell NetWare. Is a replacement product for both 34L0601 and 34L0701.

IBM High-Speed 100/16/4 Token Ring PCI Adapter

Part number 34L0501(SP), 34L0510(CP): requires a computer with a PCI Version 2.0 or later 32-bit or 64-bit busmaster-enabled slot. Supports auto-sensing ring speed (4 Mbps or 16 Mbps or 100 Mbps) and full-duplex dedicated token ring (DTR). Supports Wake on LAN but not Alert on LAN.

IBM Turbo 16/4 Token Ring PC Card 2

Part number 34L1401(SP), 34L1410(CP): requires a computer with a PCMCIA Type II or Type III slot. Supports auto-sensing ring speed (4 Mbps or 16 Mbps) and full-duplex dedicated token ring (DTR).

IBM 16/4 Token Ring CardBus Adapter

Part number 34L4801(SP), 34L4810(CP): will be available from 28 April 2000. Requires a computer with a Type II CardBus slot, which is a 32-bit interface and enables higher throughput than cards with a 16-bit PCMCIA interface. Supports auto-sensing ring speed (4 Mbps or 16 Mbps) and full-duplex dedicated token ring (DTR). Ships with the Tivoli Management Agent and provides support for SNMP and DMI 2.0. Has received the Microsoft PC99 logo and has been certified for use with Novell NetWare. Also supports ACPI power management in ACPI-compliant notebook computers with power management features enabled. Is a replacement product for 34L1401.

IBM Turbo 16/4 Token Ring ISA Adapter

Part number 72H3482(SP), 72H3514(CP): supports auto-sensing ring speed (4 Mbps or 16 Mbps).

Compatibility and operating system support

Supported environments for some of the features listed earlier are operating-system dependent; specific information is provided in the following tables. Note that some of the Token Ring adapters mentioned in these tables are no longer marketed by IBM; these adapters appear in *italic type* and have no part number after their name.

Table 5. Support for route switching peer and client

Route Switching	**Windows NT 3.51, 4.0**	**Windows 95, 98, 2000**	**Windows 3.x**	**OS/2 Warp 3.0 and later**	**Novell NetWare Server**
IBM 16/4 Token Ring CardBus Adapter (34L4801)	Supported (NT 4.0 with third party card /socket services)	Supported	Not supported	Supported	Not supported
IBM 16/4 Token Ring PCI Management Adapter (34L5001)	Supported	Supported	Not supported	Supported	Supported
IBM High-Speed 100/16/4 Token Ring PCI Adapter (34L0501) IBM 16/4 Token Ring PCI Adapter 2 (34L0601) IBM 16/4 Token Ring PCI Adapter 2 with Wake on LAN (34L0701) *IBM PCI Token Ring Adapter* *IBM PCI Wake on LAN Token Ring Adapter*	Supported	Supported	Supported (using LAN Client)	Supported	Supported
IBM Turbo 16/4 Token Ring ISA Adapter (72H3482) *IBM Auto Wake 16/4 Token Ring Network Adapter*	Supported	Not supported	Not supported	Not supported	Not supported
IBM Turbo 16/4 Token Ring PC Card 2 (34L1401) *IBM Turbo 16/4 Token Ring PC Card*	Supported	Supported	Not supported	Not supported	Supported

Table 6. Support for LAN Adapter Management Agent

LAN Adapter Management Agent	**Windows NT 3.51, NT 4.0, 95, 98, 2000**	**OS/2 Warp 3.0 or later**
All adapters	DMI Version 2.0 and SNMP Version 2	DMI Version 1.0 and SNMP Version 2

Table 7. Support for CoS for IP

CoS for IP	Windows NT 3.51, 4.0	Windows 95, 98, 2000	Windows 3.x	OS/2 Warp 3.0 and later	Novell NetWare Server 4.11 and higher
IBM 16/4 Token Ring CardBus Adapter (34L4801)	Supported (NT 4.0 only)	Supported	Not supported	Supported	Not supported
IBM 16/4 Token Ring PCI Management Adapter (34L5001)	Supported	Supported	Not supported	Supported	Supported
IBM High-Speed 100/16/4 Token Ring PCI Adapter (34L0501) IBM 16/4 Token Ring PCI Adapter 2 (34L0601) IBM 16/4 Token Ring PCI Adapter 2 with Wake on LAN (34L0701) *IBM PCI Token Ring Adapter* *IBM PCI Wake on LAN Token Ring Adapter*	Supported	Supported	Supported (using LAN Client)	Supported	Supported
IBM Turbo 16/4 Token Ring PC Card 2 (34L1401) *IBM Turbo 16/4 Token Ring PC Card*	Supported	Supported	Not supported	Not supported	Supported

Table 8. Support for Redundant NIC

Redundant NIC	**Windows NT Server 3.51 and 4.0**	**NetWare 4.11, 4.2 and 5.0**
IBM 16/4 Token Ring PCI Management Adapter (34L5001) IBM High-Speed 100/16/4 Token Ring PCI Adapter (34L0501) IBM 16/4 Token Ring PCI Adapter 2 (34L0601) IBM 16/4 Token Ring PCI Adapter 2 with Wake on LAN (34L0701) *IBM PCI Wake on LAN Token Ring Adapter*	Supported Quick Failover extension supported on Windows NT 4.0 SP5, potentially requires adapter microcode upgrade	Supported Quick Failover extension supported on NetWare 4.11 and 4.2 with IWSP6A and NetWare 5.0 with NW5SP2A, potentially requires adapter microcode upgrade
IBM PCI Token Ring Adapter	Supported; Quick Failover not supported	Supported; Quick Failover not supported

Hubs

Overview

All token ring hubs provide a shared token ring connection in which the users of the hub all share the use of a single token ring. The only exception to this is that a single stack of 8239 units can be configured so that each unit provides its own token ring; in this case all devices attached to a single unit are connected to the same shared token ring. The simplest hubs do not even require an external power supply and simply provide internal relays which make and break electrical connections between the user devices.

Products

8228 Token Ring Network Multistation Access Unit (MAU)

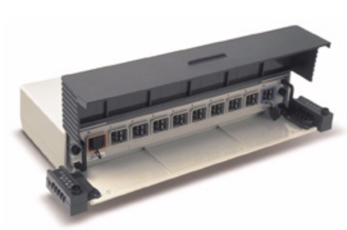

Figure 5. The 8228-001

The 8228 - an unmanaged hub - has 8 token ring user ports which use the IBM Cabling System (ICS) connectors; a media filter can be used to attach an RJ-45 token ring cable to the 8228. The 8228 has no external power supply; all its power requirements are provided by the connecting token ring adapters themselves which cause relays inside the 8228 to operate when a connection is made. There are no indicator lights on the 8228, so there is no way to know whether or not any particular port is in use or not.

The 8228 also has ring-in and ring-out ports, which can be used to link to other 8228s to form a complete ring. Two 8228s *can* be connected using a single cable between one unit's RO port and the other unit's RI port; greater resilience will be provided for networks of two or more 8228s if all RI/RO ports are attached, with the RO port of the last 8228 being connected back to the RI port of the first 8228. Doing this will allow the ring to remain active and complete even if any one of the RI/RO cables should fail or be disconnected. If the 8228 is to be used on its own then it is not necessary to connect the ring-out port to the ring-in port to complete the ring.

In a ring formed solely of 8228s, the speed at which the ring operates (4 Mbps or 16 Mbps) is determined by the first device to connect to the ring. This sometimes causes problems when the first device is set to auto-sense the speed of the ring: because there is not yet any ring activity it is unable to determine which speed to use and is unable to connect. This sometimes requires pre-configuration of ring speed in one or more devices. Watch out, however, because if the ring is operating at one speed (say 16 Mbps) and a new device is pre-configured with the other speed (4 Mbps in this case), attaching the new device to the 8228 will cause the ring to break. No recovery is possible until the offending device is disconnected from the 8228. The only totally fool-proof mode of operation is one in which all connecting devices are pre-configured with the same ring speed.

Table 9. 8228 models and features

Model	Description
8228-001	IBM Token Ring Network Multistation Access Unit

8226 Token Ring RJ45 Connection

Figure 6. The 8226-001

The 8226 is essentially an RJ-45 version of the 8228; it requires an external power supply and is therefore capable of showing port activity using indicator LEDs. It is still an unmanaged hub, meaning that it cannot report errors or be configured across the network.

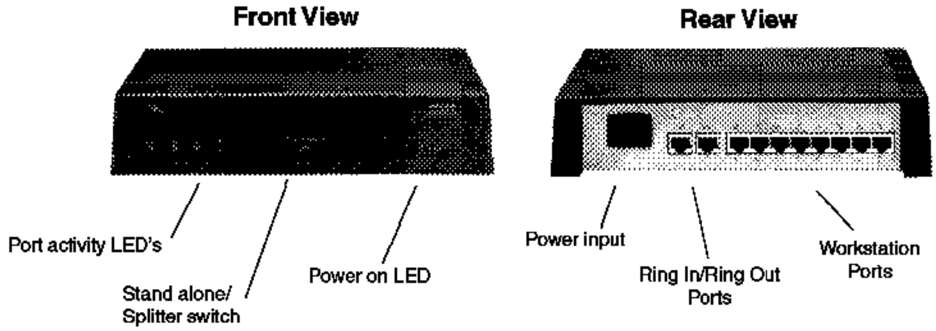

Figure 7. The 8226 indicator LEDs and ports

Although RJ-45 connectors can be used with the 8228 when appropriate converters are installed, the 8226 does offer one function not possible with the 8228: the capability to act as a *splitter*. In this mode, which is selected by the switch on the 8226, the 8226's ring-in port is connected to a normal user port in another hub (such as an 8228 or another 8226). The 8226 itself provides the power required to activate this link, which then makes the devices attached to the 8226 acting as a splitter appear as part of the ring provided by the upstream hub. This can be very useful if a small group of terminals requires to be connected to a ring when the hub for the ring is not immediately adjacent to the terminals; the use of the 8226 means that only a single cable needs to be run to the main hub instead of a cable from each workstation separately.

There is no such thing as a "16/4 switch" on the 8226: the speed of operation of the token ring is determined by the connecting network adapters themselves, just as for the 8228.

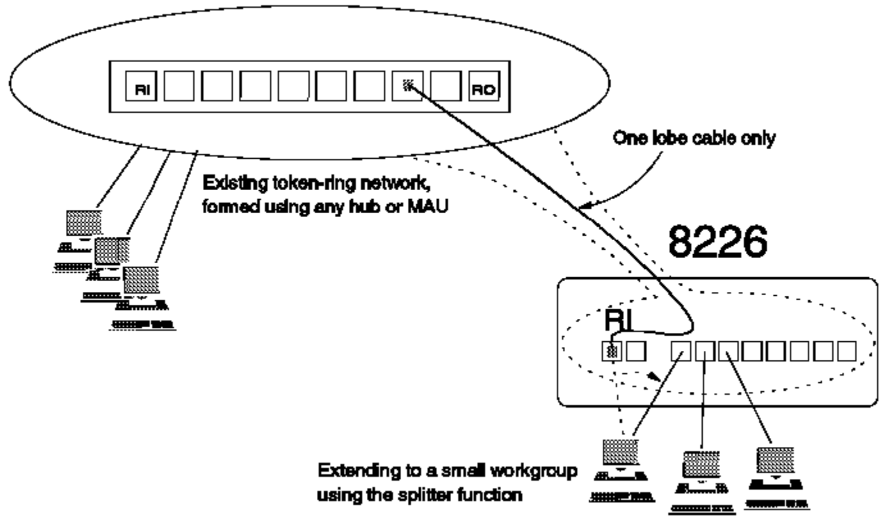

Figure 8. The 8226 using the splitter function

Table 10. 8226 models and features

Model	Description
8226-001	IBM Token Ring RJ45 Connection

8230 Controlled Access Unit (CAU)

Figure 9. The 8230 base unit

The 8230 Controlled Access Unit (CAU) - a managed hub - provides a base unit as shown in Figure 9 in which up to 5 Lobe Insertion Units (LIUs) can be fitted. In addition, up to three Lobe Attachment Modules (LAMs) can also be connected using connections on the rear of the 8230 base unit. A single 8230 can provide token ring hub ports for up to 92 devices, all of which connect to the same token ring.

If the 8230 is required to connect to another hub in the ring using ring-in and ring-out connections, an RI/RO LIU must be installed in slot 1, the right-most LIU slot in the base unit. Apart from this requirement, any type of LIU can be installed in any slot.

The 8230 itself has internal token ring ports which are used for management purposes, and since these ports connect to the same single token ring as the rest of the user devices they need to be configured with an appropriate ring speed (4 Mbps or 16 Mbps). This is done using the switch on the front of the 8230 base unit, and requires that all devices which connect to the 8230 are either configured for this same ring speed or are auto-sensing devices. Attempting to connect a device set with the wrong ring speed, however, will not affect other ports - the 8230 will simply refuse to allow it to connect.

The MAC addresses of the internal ports are shown on the front of the base unit itself. There are in fact three of them (PO, PI & S) and they are consecutively numbered in this order. The presence of three MAC addresses allows either the ring-in or ring-out port to wrap in the case of some kind of failure, allowing the ring to remain active even with a single physical break in a ring cable. This feature is known as “dual ring redundancy”. Previous models of the 8230 required the addition of the “dual ring redundancy feature”, feature #2029, but this feature is now supplied as part of the base 8230-013 or 8230-213.

User ports on the 8230 can either be “active” or “passive”. Active ports provide re-timing and re-generation of the signal and allow longer cable lengths to be used. The only difference between the two base models of 8230 is that the 8230-213 supports active UTP ports in its LAMs; the 8230-013 only supports passive LAM ports.

LAM and LIU ports offer a variety of media formats: primarily RJ-45 and the IBM Cabling System (ICS). The RI/RO ports also offer the option of using fiber, which can be used to interconnect 8230s over long distances.

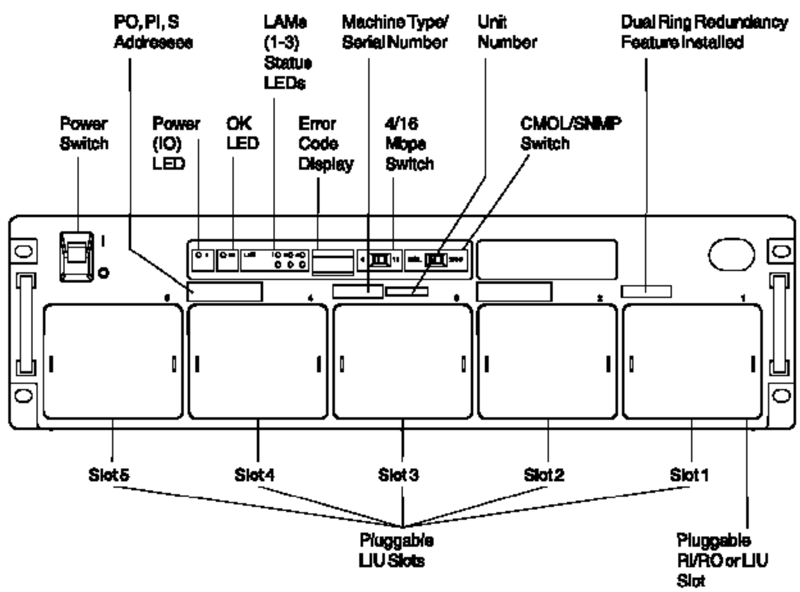

Figure 10. Layout of the front of the 8230 base unit

The 8230 LIU ports also offer a *fan-out* function. This allows up to 8 MAC addresses to be connected to a single LIU port, by using a connection to the ring-in port of another hub such as the 8228. No more than 8 user devices can be connected to a single LIU port in this manner.

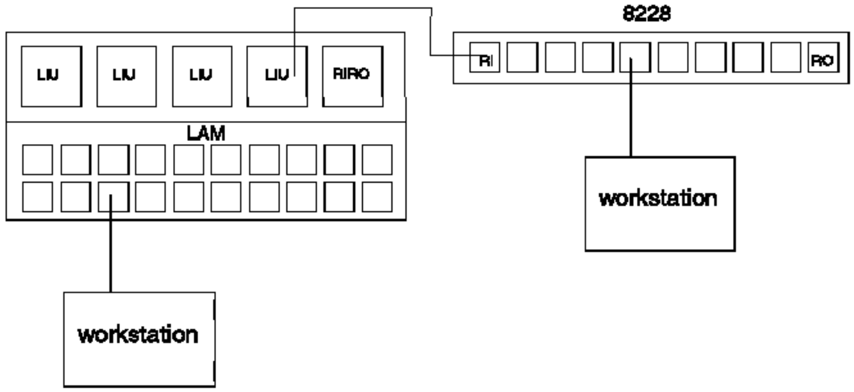

Figure 11. Fan-out using LIU ports

Fan-out is not possible using LAM ports. The other option here is the use of a hub such as the 8226 which provides the splitter function. There is no 8-port restriction imposed by the 8230 on the number of workstations which can connect using the splitter function. It just happens in this example that the 8226 only has 8 ports.

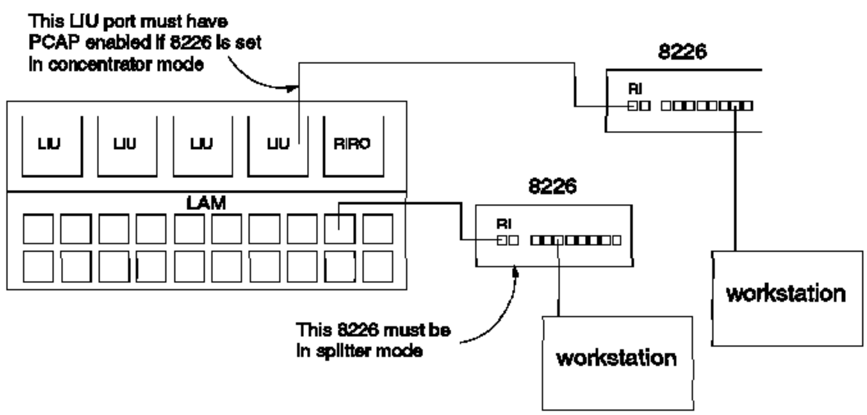

Figure 12. Fan-out using LIU ports, splitter using LAM ports

Both 8230 models support a choice of remote management through either CMOL (Common Management Information Protocol over Logical Link Control) or SNMP (Simple Network Management Protocol).

Management in CMOL mode is supported LAN Network Manager (LNM) for OS/2. Diskettes with the CMOL and SNMP microcode currently ship with the product; the reason for this is that the 8230 hardware only supports one mode of management at a time and the management software needs to be reloaded to change from CMOL to SNMP operating mode or vice-versa. A planned future hardware change will allow both code loads to reside in the 8230 simultaneously, at which point it will no longer be necessary to ship diskettes containing microcode. LNM for OS/2 V2 provides full management of the Controlled Access Units, including time-of-day access control, password security for configuration changes, enabling/disabling of ports, and remote microcode update. LNM offers a graphical display of the hub with color-coded indications of port health. This view is available on the LNM local ring, on token rings bridged to LNM, or across a routed environment if the router supports the flows of LLC type 1.

SNMP management of the 8230 is provided IBM Lan Network Manager, which is part of IBM's Nways Manager for AIX, HP-UX or Windows NT operating systems. LNM provides full management of these concentrators, including time-of-day access control, enabling/disabling of ports, remote microcode update, and token ring segment utilization. LNM offers a graphical display of the concentrators with color-coded indications of port health.

Table 11. 8230 models and features

Model / feature	Part number	Description
8230-013		IBM Controlled Access Unit
8230-213		IBM Controlled Access Unit
#2011	73G2011	ICS 2-port LIU
#2009	73G2009	UTP 4-port LIU
#2008	73G2008	UTP 3-port LIU
#6738	13H6738	Shielded RJ-45 LAM (passive)
#5501	53F5501	ICS LAM
#6748	13H6748	Shielded UTP LAM (active - model 213 only)
#7737	59G7737	ICS RI/RO
#2007	73G2007	Shielded RJ-45 RI/RO
#2010	73G2010	Optical fiber RI/RO
#7751	59G7751	Combination ICS/fiber RI/RO
#7754	59G7754	Combination fiber/ICS RI/RO

8239 Token Ring Stackable Hub

Figure 13. A fully-configured 8239-001

The 8239 Token Ring Stackable Hub - a managed hub - provides a logical concentration of a single management point for several token ring segments. Unlike the 8230, in which all devices connecting to a single managed unit are connected to the same token ring segment, a stack of 8239s can provide up to 8 separate token ring segments from a single management point.

The primary building block of this hub offering is the Model 1. It is a fully functional concentrator or hub, with the additional capability of providing Ring-In/Ring-Out (RI/RO) connections to expand the LAN segment, stack management functions, and network management functions. The Model 2 is a fully functional hub that can be used as a stand-alone hub or as an expansion hub to increase the number of device attachments in the stack. For example, the 8239-002 will isolate beaconing token ring ports and will show the error using its status LEDs but will not be able to communicate this problem to a network management station if there is no 8239-001 on the same ring segment.

An optional 16-port Expansion Feature for both models and, for Model 1 only, RJ-45 copper or ST fiber RI/RO modules complete the offering. The RI/RO modules would only be used to connect to another hub such as the 8228 or 8230. Fan-out to devices such as the 8226 or 8228 is also supported; more than 8 ports may be provided through the fan-out device but only the first 8 are reported for network management purposes using RMON.

A single stack of 8239 units is constructed by interconnecting the 8239 units through the stack-in and stack-out ports. Standard Category 5 UTP cables are required for this.

All hubs in the stack are manageable from a single point using out-of-band access, available on both models, or in-band access via the Model 1. A Model 1 is required for network management functions. Information is accessed either through the terminal interface or an SNMP-compatible application. The 8239 is supported by IBM's Nways Workgroup Manager for Windows NT, Nways Manager for AIX and Nways Manager for HP-UX. The 8239 offers many or more of the same network management features as multifunction switching hubs but at a much lower price per port. Management functions include the ability to identify which port in the stack is being used by a particular MAC address and the ability to configure ports so that only specific MAC addresses can use them.

In the 8239 stack, segmentation is provided at the hub level, by wrapping a given hub in or out of the default single segment. This isolates the data path of one or more hubs from the data paths of the other hubs in the stack. This segmentation does not extend to segmenting ports in a single hub: all ports in the hub, including additional ports if the optional feature card is installed, are treated as a single entity or hub when segmenting a ring. Every hub in a stack must be assigned to a segment; the factory default has all hubs assigned to a single segment. If, for example, a stack consisted of four hubs (any combination of models), it could be segmented into a maximum of four segments. If a stack consisted of the maximum eight hubs, from one to eight segments could be created. Hubs in the same segment must be adjacent to each other in ring order when you connect them with the intra-stack cables. Management of the stack is maintained because the control path of the stack is not affected when the data path is altered because of segmentation. If network management functions, like Ring-In/Ring-Out or RMON, are important for a particular segment, that segment must consist of at least one Model 1.

The 8239 offers an impressive array of network management functions as well as stack management capability. Functions like Ring-In/Ring-Out, RMON, 802.5, Surrogate MIB for media management, etc. are all programmed into the Model 1. Accordingly, a Model 1 must be included in each segment where this activity is required.

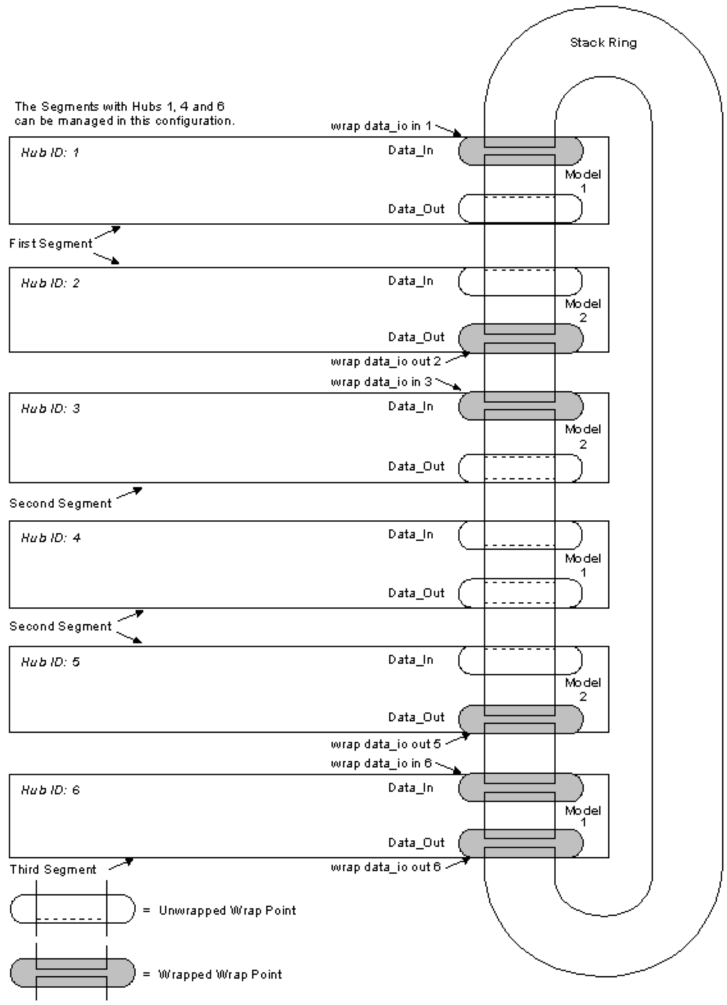

Figure 14. 8239 stack with three segments

As an example of segmentation, Figure 14 shows a single stack containing three segments and six hubs; each segment can be managed because an 8239-001 is included in each.

Table 12. 8239 models and features

Model / feature	Part number	Description
8239-001	08L3033	IBM Token Ring Stackable Hub
8239-002	08L3034	IBM Token Ring Stackable Hub
#3035	08L3035	16-port Expansion Feature
#3036	08L3036	RJ-45 Copper RI/RO (Model -001 only)
#3037	08L3037	Optical fiber RI/RO (Model -001 only)

Switches

Overview

A token ring switch is effectively a combination of a hub and a bridge. It differs from a hub because it does not repeat all traffic on all ports; it maintains an internal table of MAC addresses and only sends directed frames over the ports on which it is required to do so. It does, however, need to send broadcast frames over all ports. It is similar to a bridge in the sense that it can be used to interconnect token ring LAN segments, and in this way it usually operates just like a transparent bridge. It is also like a bridge because it does not have ring-in and ring-out ports; instead it connects to normal user ports on token ring hubs in an identical manner to a bridge connecting to a shared token ring hub.

Products

8270 Nways Token Ring LAN Switch

Figure 15. The IBM 8270 Nways Token Ring Switch Model 600

Figure 16. The IBM 8270 Nways Token Ring Switch Model 800

The IBM 8270 Token Ring Switch provides up to 30 token ring ports, each of which can provide a 32 Mbps dedicated full-duplex connection. It is designed to be “plug and play” because all ports are configured by default to be auto-sensing: they will detect whether they are connected to a shared hub or to a single LAN device and will adjust ring speed and mode of operation automatically. The 8270 operates internally at 480 Mbps, and buffers are provided on ports to permit data bursts to exceed this capacity.

The 8270 maintains an internal table of the MAC addresses that can be reached through each port. The 8270 supports up to 1,790 MAC addresses per 16 Mbps port and up 10,000 MAC addresses for the entire 8270. Even in a full configuration of 30 token ring ports it is unlikely that more than a small number of these ports will connect to some kind of backbone environment through which many MAC addresses are reachable; most ports will connect directly to servers or to workgroup environments with a limited number of MAC addresses.

The 8270 does not need to maintain information on every possible MAC address reachable through all ports; frames that are destined for a target station more than one source-routed bridge hop away are forwarded based on the source-route descriptor alone. This also means that duplicate MAC addresses are still supported in the switched environment. The figures above for the number of MAC addresses per port and per unit are actually the numbers for combined MAC address and source-route descriptor entries.

By default, the ports on the 8270 will operate in “adaptive cut-through” mode, which means that they will operate in cut-through switching mode for highest performance unless large numbers of errors are detected, in which case the port will change to “store and forward” mode which includes verification of frame checksum (FCS) values.

The 8270 supports a maximum frame length of 4,540 bytes. When operating in cut-through mode, frames larger than this will be truncated; when operating in store-and-forward mode they will simply be rejected.

The operating mode of each port can also be configured so that it is required to operate in a particular mode or at a specified speed. This configuration can be performed using the console port, using a remote Telnet terminal or by using SNMP from a remote network management station. The 8270 SNMP agent also includes RMON support.

In order to reduce the amount of broadcast traffic, which by default is transmitted on all ports, the switch can be divided up into 16 broadcast domains. Broadcasts are then restricted to the ports which are members of a particular domain. The 8270 can also be configured as a source-route bridge itself, and this allows for the internal connection of different 8270 domains.

A single port on the 8270 can be designated to mirror traffic on any other token ring port; this allows tuning and troubleshooting using an external network protocol analyzer.

The 8270 would typically be used to increase network performance by segmenting a single shared token ring network into separate segments, in much the same manner as a token ring bridge. Since it is effectively a multi-port bridge, a single 8270 can perform the role of many separate dedicated LAN bridges. It would also be used to act as a network backbone hub, with major servers being connected directly to the 8270 and operating in full-duplex mode.

Although the 8270 does not provide 100 Mbps token ring ports, there are several ways of interconnecting more than one 8270 to provide a high-speed backbone:

- Up to four ports each on a pair of 8270s can be configured to provide a TokenPipe connection between them, with effectively 128 Mbps capacity provided by the pipe.
- The ATM UFC may be used to provide a 155 Mbps connection to a high-speed backbone ATM network. The ATM uplink will be configured as a token ring port connecting to an emulated token ring ATM LAN segment, and traffic will be switched to this port in the same manner as all actual token ring ports on the switch.

- The MSS Client UFC may be used to provide a FasTR connection to another device, including another 8270. In this mode of operation no ATM network or MSS Server is actually required and throughput of the order of 100 Mbps is achievable over the fiber connection.
- The MSS Client UFC may be used to connect to an MSS Server in an ATM network and provide higher-performance configurations than with the ATM UFC alone; the MSS Client can be configured either as a router or as a source-route bridge and can establish shortcuts across the ATM network using either NHRP (Next Hop Resolution Protocol) or MPOA (Multiprotocol Over ATM). The MSS Client is compatible with any NHRP or MPOA server in the ATM network, although IBM has implemented some extensions to the standard to increase performance with the use of IBM's MSS Server.

The MSS Domain Client UFC provides no external connection (in other words, there is no ATM or token ring port on the device) but implements an IP routing function for the 8270 switch. This may be used as a normal IP router, but the main advantage is in conjunction with route switching peer devices in the token ring network. It allows a single 8270 to subdivide the entire token ring network into separate IP networks without the need for a separate IP router, but with the use of route switching peer devices the actual data traffic between devices on separate IP subnets will flow on a *switched* path between the devices and will not flow through the IP router provided by the MSS Domain Client. This provides the high performance of switching coupled with the convenience of routing.

The 8270-800 provides 8 feature slots and the base unit comes with a single power supply unit. A token ring processor card must be ordered in addition to the base unit and, optionally, a second power supply may be ordered for backup/redundancy purposes. The token ring processor card does not take up one of the feature slots, it is installed in a dedicated slot in the base unit. The 8270-600 provides 6 feature slots, does not allow the installation of a second power supply but comes with a built-in token ring processor.

Although at first glance it might appear that 32 token ring ports can be provided by the 8270-800, in fact if a 4-port token ring module is installed in UFC slot 1 only the first two of the token ring ports will become active. An "invalid slot" message will appear in the status display, but the first two ports will still operate correctly. This explains the limit of 30 token ring ports on the 8270-800. The 8270-600 is limited to 24 token ring ports with no special restrictions on module placement.

The fiber optic token ring ports are compatible with the ring-in/ring-out fiber ports on machines such as the 8230 and 8265. This is the only exception to the statement that the 8270 does not provide RI/RO ports. They will only operate in full-duplex mode when connected to similar ports on another 8270, under which circumstances they will not be operating in “RI/RO compatibility mode” but form part of a TokenPipe configuration.

Table 13. 8270 models and features

Model / feature	Part number	Description
8270-600	25L4970	IBM 8270 Nways LAN Switch Model 600
8270-800	85H6584	IBM 8270 Nways LAN Switch Model 800
#4650	72H4650	Token Ring Processor Card (Model 800 only)
#4648	72H4648	Redundant Power Supply (Model 800 only)
#5092	85H5092	4-Port Token Ring Enhanced UTP/STP
#5087	85H5087	2-Port Token Ring Enhanced Fiber
#2762	86H2762	1-Port ATM Token Ring II Multimode Fiber
#5205	85H4596	MSS Client Multimode Fiber
#5206	85H4599	MSS Client Singlemode Fiber
#5207	85H9303	MSS Domain Client

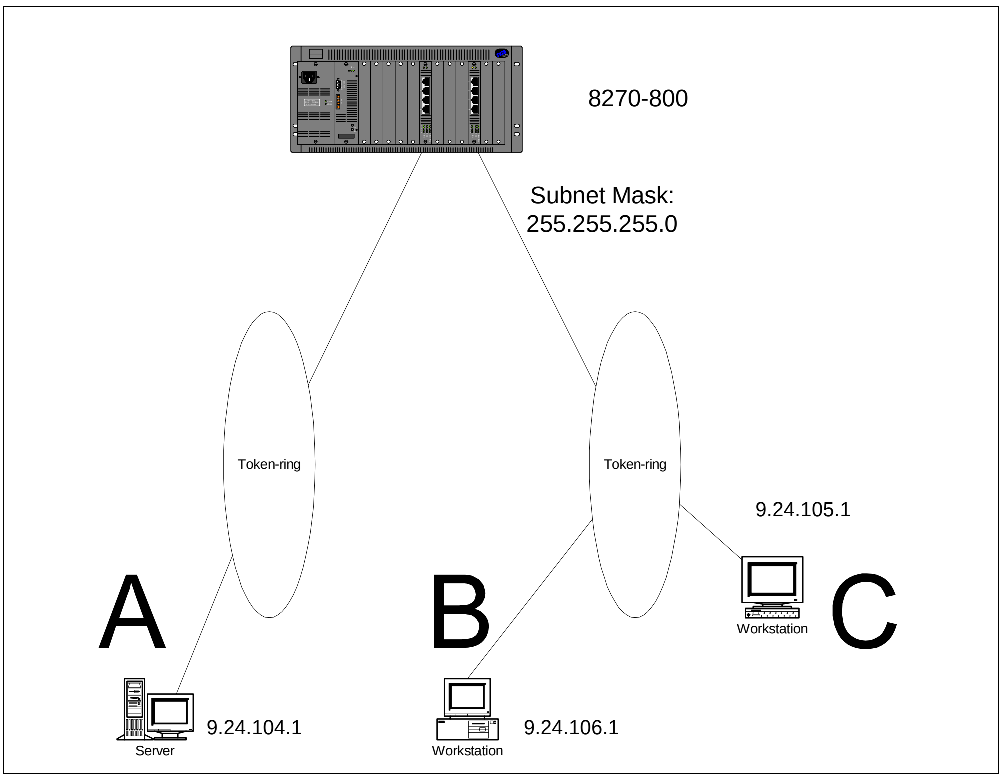

Figure 17. Route switching peer configuration

Figure 17 shows a network of three computers, each configured to be in a different IP subnetwork. If all three have route switching peer enabled on their token ring adapters, IP traffic will be switched directly between them. Traffic from A to either B or C still needs to pass through the token ring switch, but it will be switched and not routed. There is no router shown in this diagram; the router could in fact be installed inside the 8270 using the MSS Domain Client UFC or else it could be connected to one of the token rings above with multiple IP addresses configured on a single physical LAN interface. In either case, traffic between the three LAN stations shown will not pass through the router itself once the layer 2 shortcuts have been established.

Withdrawn products

Table 14 is a table summarizing some token ring products which may still be encountered in customer environments even though they have now been withdrawn from marketing.

Table 14. Withdrawn token ring products

Product	Date withdrawn	Replacement product
8272	12/98	8270
8238	7/98	8239
8229 bridge	7/98	2210 router
3172-003	12/98	S/390 OSA
8230-001, 8230-002	4/96	8230-x13 or 8239
8230-04A, 8230-04P	9/99	8230-x13 or 8239
8230-003	2/00	8230-x13 or 8239

Network management

Network management is essentially concerned with two issues: the remote configuration and monitoring of network-attached devices. The first network management implementations were used to control the building blocks of the network: the hubs, bridges and routers. Specific users could be permitted or denied access at specified times of day or days of the week, but this was accomplished by configuring the network devices themselves. Network devices reported faults and errors back to a management station to improve problem determination in complex network environments. More recently, distributed intelligence has been provided in workstations themselves, not just in complex servers but in more basic PCs, to the extent that a PC can report the removal of its cover or a basic hardware fault. This increases the controllability and manageability of the total network environment.

Overview

Network management relies on standard protocols and architectures for the creation and transmission of network management information. In essence, network devices and network management stations communicate using standard LAN packets conforming to different specifications.

CMOL

CMOL stands for Common Management Information Protocol over LLC (or CMIP over LLC) and is an OSI standard for the transport of network management information. It was much favored in the public telecommunications carrier environment and provides a rich set of management functions; it was originally used by IBM to provide network management of the 8230 Controlled Access Unit. The management station for the 8230 was typically an OS/2 workstation running LAN Network Manager for OS/2, and many customer 8230 installations continue to run this today.

The 8230 uses one of its internal MAC addresses as the source address for network management communication to the network management station, and hence CMOL information flows at the layer 2 level, using MAC addresses across a bridged LAN environment. The related CMOT (CMIP over TCP/IP) standard allows this information to be transported using layer 3 IP packets, but this is not implemented by the 8230.

SNMP

The Simple Network Management Protocol (SNMP) grew up in much the same way as the rest of the TCP/IP protocol stack: rather than being designed by committee and providing a rich set of functions it grew up as a pragmatic implementation based on what was really required in the network. SNMP has since grown to be the effective industry standard for management of network devices, and provides a relatively simple “vocabulary” of commands which effectively comprises just three - the ability to SET commands on the remote device, the ability to GET information from the remote device and the ability of the remote device to send unsolicited TRAPs alerting the network management station of status changes and problems in the network.

Most hubs, bridges, routers and similar devices in networks today offer the capability of being managed using SNMP, which requires them to be configured with an internal IP address (if they don't already have one) and requires IP connectivity between the device and the management station. Even the 8230 now offers this capability, selected by a switch on the front of the base unit.

RMON

Remote Network Monitoring (RMON) is concerned with the implementation of *probes* which can be used to troubleshoot and monitor remote LANs. RMON probes are typically accessed across the network using IP protocols and allow the retrieval of status information - percentage utilization, for example - rather than the application of configuration information and retrieval of error information. So RMON serves a separate purpose to both SNMP and CMOL/CMOT and both would normally be expected to operate in parallel.

DMI

The Desktop Management Interface standard is concerned with the monitoring and retrieval of status and error information from workstations. It serves much the same purpose as either SNMP or CMOL/CMOT but usually applies to a different set of network devices - user workstations themselves rather than the building blocks of the network.

Applications

LAN Network Manager

LAN Network Manager for OS/2 continues to provide CMOL management of the 8230 Controlled Access Unit, although it is most likely that new 8230 installations would today choose SNMP management.

LAN Adapter Management Agent

The IBM LAN Adapter Management Agent supports any IBM LAN adapter operating on Windows NT Server 3.51 or later, Windows 95, Windows 98, Windows 2000 or OS/2 Version 3.0 or later. It implements DMI Version 2.0 and SNMP Version 1 on Windows platforms, DMI Version 1.0 and SMNP Version 2 on OS/2.

The IBM LAN Adapter Management Agent makes IBM LAN adapters visible to management applications - ones which use either DMI or SNMP or both.

The IBM LAN Adapter Management Agent can be coupled with the Nways Management Application to provide remote management of IBM LAN adapters from an Nways management workstation such as Nways Workgroup Manager for Windows NT, Nways Manager for AIX or Nways Manager for HP-UX - the agent can also be managed by any SNMP-compliant management application.

The software package shipped by IBM provides DMI browsers which can be used for remote management using DMI.

Tivoli Management Agent

The Tivoli Management Agent (TMA) is supported on all IBM token ring adapters which operate on Windows NT, Windows 95, Windows 98, NetWare 3/4/5, OS/2 and Windows 3.x operating systems. TMA provides a framework which allows remote management, software distribution, inventory and user administration and is controlled from a Tivoli Management Gateway such as Tivoli Enterprise.

TCP/IP and SNA

There has always been a perception that the SNA networking protocol is in some way linked to the token ring LAN medium. This probably comes about because SNA and token ring were both developed by IBM. It is also true that most large SNA networking customers tend to use token ring, in the data center at least, and there are certain features inherent in token ring which are specifically used in token ring environments: the efficiencies of transporting large frames and the resilience of duplicate MAC addresses coupled with source-route bridging.

The problem with this perception is that the assumption is sometimes made that because token ring is somehow *good* for SNA traffic that it is also somehow *bad* for IP traffic. Nothing could be further from the truth.

Indeed, IBM has made many innovations and enhancements, specifically for TCP/IP, which apply to the token ring environment just as for any other. Route switching allows token ring devices to establish shortcuts to increase IP throughput. Token ring devices are monitored and managed using IP protocols. Indeed, some of the advantages of token ring can be used to the advantage of IP traffic: “CoS for IP” allows the marking of important IP traffic with higher transmission priority than other IP traffic in a way which is just not possible using Ethernet. Almost nothing described in this paper on token ring applies exclusively to SNA traffic or, indeed to any other layer 3 protocol. Customers with existing token ring networks are going to want to continue using it as their layer 2 transport mechanism of choice for all network traffic.

Conclusions

The purpose of this white paper has been to show that IBM remains committed to token ring and that IBM continues to sell and to enhance its range of token ring products. Further, IBM intends to remain the leader for token ring products, and the business continues to be one of considerable size. The installed base of token ring ports is expected to grow in each of the next two years and this means a worldwide business of the order of a million token ring ports per year in the short term.

IBM token ring products cover a wide range: from token ring adapters on workstations and servers to token ring switched hubs with high-speed backbone uplinks. While IBM does not plan to develop standalone high-end network routers with token ring interfaces, IBM is committed to the rest of the token ring business, including the mainframe channel-attached controller which continues to be used to provide high-capacity token ring access using both SNA and TCP/IP.

This paper explains how and why all of IBM's token ring products can be used and gives technical information on the features and functions of these products as well as listing model, part and feature numbers.

Token ring is not about to be overtaken and superseded by other technologies such as high-speed Ethernet and ATM, primarily due to the fact that token ring has delivered a reliable and stable infrastructure for many years and will continue to do so in the future. Many customers are not going to replace their token ring infrastructures simply because it may be fashionable: they realize that their investment in token ring is capable of catering to their networking needs for years to come. Token ring offers significant architectural advantages which are discussed in this paper, and in many cases significant improvements to existing token ring infrastructures can be made by microcode updates and replacement of key network components rather than by wholesale replacement of the entire network. As an example, key bottlenecks in token ring networks can be eliminated by replacing a few hubs with switches and by leaving the rest of the network unchanged.

Abbreviations and acronyms

ACPI	Advanced Configuration and Power Interface
APPN	Advanced Peer-to-Peer Networking
ASTRAL	Alliance for Strategic Token Ring Advancement
ATM	Asynchronous Transfer Mode
CAU	Controlled Access Unit
CI	Component Interface (part of DMI)
CMIP	Common Management Information Protocol
CMOL	CMIP Over LLC
CMOT	CMIP Over TCP/IP
CTR	Classic Token Ring
DHCP	Dynamic Host Configuration Protocol
DMI	Desktop Management Interface
DMTF	Desktop Management Task Force
DTR	Dedicated Token Ring
FCS	Frame Check Sequence
HSTRA	High-Speed Token Ring Alliance
IBM	International Business Machines Corporation
ICS	IBM Cabling System
IDC	International Data Corporation
IEEE	Institute of Electrical and Electronics Engineers
ITSO	International Technical Support Organization
LAM	Lobe Attachment Module
LAN	Local Area Network
LIU	Lobe Insertion Unit
LNM	LAN Network Manager
LLC	Logical Link Control
MAC	Medium Access Control
MAE	Multi-Access Enclosure (inside 3746-9x0)
MI	Management Interface (part of DMI)
MIB	Management Information Base (part of SNMP)
MIF	Management Information Format (part of DMI)
MFIOP	Multi-Function I/O Processor (on AS/400 systems)
MPOA	Multiprotocol Over ATM
MSS	Multiprotocol Switched Services
NBU	IBM Networking Business Unit
NCP	Network Control Program (runs on 3745)
NetBIOS	Network Basic Input/Output System
NIC	Network Interface Card
OSA	Open Systems Adapter (on S/390 mainframes)
OSPF	Open Shortest Path First
PCI	Peripheral Component Interconnect
PCMCIA	Personal Computer Memory Card International Association
PME	Power Management Event
PXE	Preboot eXecution Environment
RAS	Reliability, Availability, Serviceability
RI	Ring In
RIPL	Remote Initial Program Load
RMON	Remote MONitor
RO	Ring Out
RPL	Remote Program Load
SMS	System Management Server (Microsoft)
SNA	Systems Network Architecture
SNMP	Simple Network Management Protocol
STP	Shielded Twisted Pair
TCP/IP	Transmission Control Protocol/ Internet Protocol
TKI	Transmit Immediate Access (“full duplex” token ring)
TKP	Token Passing Protocol
TMA	Tivoli Management Agent

UDP	User Datagram Protocol
UFC	Universal Feature Card
UMI	Universal Manageability Initiative (IBM strategy)
UTP	Unshielded Twisted Pair
XID	eXchange ID

Special notices

This publication is intended to help customers, business partners and IBM sales and technical people understand the scope and range of token ring products currently available from IBM.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS, without warranty. The information about non-IBM ("vendor") products in this paper has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX	Alert on LAN
APPN	AS/400
AT	CT
ESCON	IBM
LANStreamer	Manage. Anything. Anywhere.
Micro Channel	Netfinity
Nways	OS/2
RS/6000	S/390
SP	SP2
Streamer	System/390
Wake on LAN	400
Lotus	Tivoli
Planet Tivoli	TME
NetView	Cross-Site
Tivoli Ready	Tivoli Certified

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere.,The Power To Manage., Anything. Anywhere.,TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this white paper.

IBM Redbooks

For information on ordering these publications see “How to get IBM Redbooks” on page 61.

- *MSS 2.1, Including the MSS Client and Domain Client,* SG24-5231
- *Layer 3 Switching Using MSS and MSS Release 2.2 Enhancements*, SG24-5311
- *3174 Establishment Controller/Networking Server Installation Guide*, GG24-3061
- *OSA-2 Implementation Guide (Update),* SG24-4770
- *AS/400 System Builder*, SG24-2155
- *An Inside Look at IBM Workgroup Hubs and Switches*, GG24-2528

IBM Redbooks collections

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at `http://www.redbooks.ibm.com/` for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	**Collection Kit Number**
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr Format)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

Other resources

These publications are also relevant as further information sources:

- *Local Area Network Technical Reference*, SC30-3383
- *Token-Ring Network Architecture Reference*, SC30-3374

Referenced Web sites

These Web sites are also relevant as further information sources:

- `http://www.networkworld.com/` Network World magazine
- `http://www.idc.com/` IDC Corporation

- `http://www.delloro.com/` Dell'Oro Group
- `http://www.hstra.com/` High-Speed Token Ring Alliance
- `http://www.networking.ibm.com/` IBM Networking Home Page
- `http://www.tivoli.com/` Tivoli Systems

How to get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** http://www.redbooks.ibm.com/

  Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

  Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the IBM Redbooks fax order form to:

	e-mail address
In United States or Canada	pubscan@us.ibm.com
Outside North America	Contact information is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at http://w3.itso.ibm.com/ and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at http://w3.ibm.com/ for redbook, residency, and workshop announcements.

IBM Redbooks fax order form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

☐ Invoice to customer number

☐ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.