

Overview of Token Ring Switching

This chapter provides a brief overview of Token Ring switching, and describes the industry standard functions supported by the Catalyst Token Ring switches as well as several functions that are unique to the Catalyst line of Token Ring switches.

This chapter provides the following information:

- Why Use Token Ring Switches?
- History of Token Ring Switching
- Bridging Modes
- Forwarding Modes
- Port Operation Modes
- Speed Adaptation
- Transmission Priority Queues
- Frame Filtering
- Broadcast Control

Why Use Token Ring Switches?

The traditional method of connecting multiple Token Ring segments is to use a source-routing bridge (SRB). For example, bridges are often used to link workgroup rings to the backbone ring. However, the introduction of the bridge can significantly reduce performance at the user's workstation. Further problems may be introduced by aggregate traffic loading on the backbone ring.

To maintain performance and avoid overloading the backbone ring, you can locate servers on the same ring as the workgroup that needs to access the server. However, dispersing the servers throughout the network makes them more difficult to back up, administer, and secure than if they are located on the backbone ring. Dispersing the servers also limits the number of servers that particular stations can access.

Collapsed backbone routers may offer greater throughput than bridges, and can interconnect a larger number of rings without becoming overloaded. Routers provide both bridging and routing functions between rings and have sophisticated broadcast control mechanisms. These mechanisms become increasingly important as the number of devices on the network increases.

The main drawback of using routers as the campus backbone is the relatively high price per port and the fact that the throughput typically does not increase as ports are added. A Token Ring switch is designed to provide wire speed throughput regardless of the number of ports in the switch. In addition, the Catalyst 3900 Token Ring switch can be configured to provide very low latency between Token Ring ports by using cut-through switching.

As a local collapsed backbone device, a Token Ring switch offers a lower per-port cost and can incur lower interstation latency than a router. In addition, the switch can be used to directly attach large numbers of clients or servers, thereby replacing concentrators. Typically, a Token Ring switch is used in conjunction with a router, providing a high-capacity interconnection between Token Ring segments while retaining the broadcast control and wide-area connectivity provided by the router.

History of Token Ring Switching

The term switching was originally used to describe packet-switch technologies such as Link Access Procedure, Balanced (LAPB), Frame Relay, Switched Multimegabit Data Service (SMDS), and X.25. Today, LAN switching refers to a technology that is similar to a bridge in many ways.

Like bridges, switches connect LAN segments and use information contained in the frame to determine the segment to which a datagram needs to be transmitted. Switches, however, operate at much higher speeds than bridges, and can support new functionality, such as virtual LANs (VLANs).

Token Ring switches first appeared in 1994. The first-generation Token Ring switches can be divided into two basic categories:

- Processor-based switches

These switches use reduced instruction set computer (RISC) processors to switch Token Ring frames. Although they typically have a lot of function, they are slow and relatively expensive. These switches have been deployed mainly as backbone switches because of their high cost.

- Application-specific integrated circuit (ASIC)-based switches with limited functionality

These switches are fast and relatively inexpensive, but have very limited function. Typically, they offer little to no filtering, limited management information, limited support for bridging modes, limited VLANs. Today, although these switches are less expensive than processor-based switches, they are still too expensive and limited for widespread use of dedicated Token Ring to the desktop.

In 1997, a second generation of Token Ring switches was introduced. Cisco's second-generation Token Ring switches use ASIC-based switching, but they provide increased functionality resulting in a higher speed and lower cost. They also provide a wider variety of function than their predecessors, including support for multiple bridging modes, Dedicated Token Ring (DTR) on all ports, high port density, high-speed links, filtering, Remote Monitoring (RMON) management, broadcast control, and flexible VLANs.



The family of second-generation Token Ring switches can be used for backbone switching, workgroup microsegmentation, and dedicated Token Ring to the desktop. Token Ring switches currently being offered include:

- The Catalyst 3900, which is a stackable workgroup switch that provides support for all switching modes, filtering, RMON, DTR, and SNMP management, as well as support for Asynchronous Transmission Mode (ATM) and Inter-Switch Link (ISL).
- The Catalyst 3920, which is also a stackable workgroup switch that provides support for all switching modes, filtering, RMON, DTR, and SNMP management.
- The Catalyst 5000, which is a modular switch that supports Ethernet, Fast Ethernet, Fiber Distributed Data Interface (FDDI), ATM, and now Token Ring.

Bridging Modes

The Catalyst Token Ring switches support the following bridging modes:

- Source-Route Bridging
- Source-Route Transparent Bridging
- Source-Route Switching

Source-Route Bridging

Source-route bridging (SRB) is the original method of bridging used to connect Token Ring segments. A source-route bridge makes all forwarding decisions based upon data in the routing information field (RIF). It does not learn or look up Media Access Control (MAC) addresses. Therefore, SRB frames without a RIF are not forwarded.

With SRB, each port on the switch is assigned a ring number and the switch itself is assigned one or more bridge numbers. This information is used to build RIFs and to search them to determine when to forward a frame.

Clients or servers that support source routing typically send an explorer frame to determine the path to a given destination. There are two types of explorer frames: all-routes explorer (ARE) and spanning-tree explorer (STE). SRB bridges copy ARE frames and add their own routing information. For frames that are received from or sent to ports that are in the spanning-tree forwarding state, bridges copy STE frames and add their own routing information. Because ARE frames traverse all paths between two devices, they are used in path determination. STE frames are used to send datagrams because the spanning tree ensures that only one copy of an STE frame is sent to each ring.

Source-Route Transparent Bridging

Source-route transparent bridging (SRT) bridging is an IEEE standard that combines SRB and transparent bridging. An SRT bridge forwards frames that do not contain a RIF based on the destination MAC address. Frames that contain a RIF are forwarded based on source routing.

There are two possible problems with using SRT:

- Some protocols, such as SNA, attempt to establish a connection using a frame without a RIF. In the SNA case, this test frame is sent to see if the destination is on the same ring as the source. If no response is received from this test frame, then an ARE test frame with a RIF is sent. If SRT bridging is used, the first test frame without a RIF is forwarded through the bridge to the destination. The destination responds, and the spanning-tree path through the bridges is used. Although this path will work, it may be undesirable. The network may be configured with parallel backbones with the intent that traffic is to be distributed across the backbones. This works well if source-routing is used. However, if the spanning-tree path is used, then only one of the backbones will carry traffic. The other backbone will not be used unless there is a failure.
- The use of duplicate SNA gateway MAC addresses can cause a problem. SNA requires the user to enter the destination MAC address of the gateway (for example, IBM 3745 Token Ring interface coupler [TIC]). To prevent the user from having to enter a backup address in the case of a gateway failure, many SNA network designers put another gateway on a different ring with the same MAC address. This works with source routing and allows automatic recovery of a failed gateway. However, SRT does not allow the same MAC address to be on two different rings.

Source-Route Switching

Because standard transparent bridging does not support source-routing information, a new bridging mode, called source-route switching, was created. Source-route switching forwards frames that do not contain routing information based on MAC address, the same way that transparent bridging does. All rings that are source-route switched have the same ring number and the switch learns the MAC addresses of adapters on these rings.

In addition to learning MAC addresses, in source-route switching the switch also learns route descriptors. A route descriptor is a portion of a RIF that indicates a single hop. It is defined as a ring number and a bridge number. When a source-routed frame enters the switch, the switch learns the route descriptor for the hop closest to the switch. Frames received from other ports with the same next-hop route descriptor as their destination are forwarded to that port.

The key difference between SRB and source-route switching is that while a source-route switch looks at the RIF, it never updates the RIF. Therefore, all ports in a source-route switch group have the same ring number.

Source-route switching provides the following benefits:

- The switch does not need to learn the MAC addresses of the devices on the other side of a source-route bridge. Therefore, the number of MAC addresses that the switch must learn and maintain is significantly reduced.
- The switch can support parallel source-routing paths.
- An existing ring can be partitioned into several segments without requiring a change in the existing ring numbers or the source-route bridges.
- The switch can support duplicate MAC addresses if the stations reside on LAN segments with different LAN IDs (ring numbers).

Forwarding Modes

The Catalyst Token Ring switches support one or more of the following forwarding modes:

- Store-and-Forward
- Cut-Through
- Adaptive Cut-Through



Store-and-Forward

Store-and-forward is the traditional mode of operation for a bridge and is one of the modes supported by the Catalyst 3900 and the Catalyst 5000 Token Ring switching card. In store-and-forward mode, the port adapter reads the entire frame into memory and then determines if the frame should be forwarded. At this point, the frame is examined for any errors (frames with errors are not forwarded). If the frame contains no errors, it is sent to the destination port for forwarding.

While store-and-forward mode reduces the amount of error traffic on the LAN, it also causes a delay in frame forwarding that is dependent upon the length of the frame.

Cut-Through

Cut-through mode is a faster mode of forwarding frames and is supported by the Catalyst 3900. In cut-through mode, the switch transfers nonbroadcast packets between ports without buffering the entire frame into memory. When a port on the switch operating in cut-through mode receives the first few bytes of a frame, it analyzes the packet header to determine the destination of the frame, establishes a connection between the input and output ports, and, when the token becomes available, it transmits the frame onto the destination ring.

In accordance with specification ISO/IEC 10038, the Catalyst 3900 uses Access Priority 4 to gain priority access to the token on the output ring if the outgoing port is operating in half-duplex (HDX) mode. This increases the proportion of packets that can be forwarded and makes it possible for the switch to reduce the average interstation latency.

In certain circumstances, however, the cut-through mode cannot be applied and the switch must buffer frames into memory. For example, buffering must be performed in the following circumstances:

- The switch has more than one packet to transmit to the same ring.
- A packet is switched between 4- and 16-Mbps rings.
- The destination ring is beaconing.

Adaptive Cut-Through

Adaptive cut-through mode uses a combination of store-and-forward and cut-through modes and is supported by the Catalyst 3900. With adaptive cut-through mode, the user can configure the switch to automatically use the best forwarding mode based on user-defined thresholds. In adaptive cut-through mode, the ports operate in cut-through mode unless the number of forwarded frames that contain errors exceeds a specified percentage. When this percentage is exceeded, the switch automatically changes the mode of the port to store-and-forward. Then, once the number of frames containing errors falls below a specified percentage, the operation mode of the ports is once again set to cut through.

Port Operation Modes

A port can operate in one of the following modes:

- Half-duplex concentrator port—Port is connected to a single station in HDX. In this case, the port behaves like an active multistation access unit (MAU) port for classic Token Ring.
- Half-duplex station emulation—Port is connected to a port on an MAU. In this case, the port behaves like a station connected to a classic Token Ring segment that contains multiple stations.
- Full-duplex concentrator port—Port is connected to a single station in full-duplex (FDX) mode.
- Full-duplex station emulation—Port is connected to another Token Ring switch in FDX mode.

Ring In/Ring Out

In addition to the port operation modes listed above, certain ports can operate in Ring In/Ring Out (RI/RO) mode. In RI/RO mode, the port is connected to a traditional main ring path coming from either a MAU or a controlled access unit (CAU).

For the Catalyst 3900, ports 19 and 20 and any of the ports on a fiber expansion module can operate in RI/RO mode. For the Catalyst 5000, any of the ports on the fiber Token Ring module can operate in RI/RO mode. The Catalyst 3920 does not support RI/RO mode.

You can use the RI/RO ports to provide redundancy in a segment. For example, let's assume that you have three MAUs that are daisy-chained together (RO of one MAU connected to RI of the next) with the RO of the third MAU being connected back to RI of the first one. To add a Catalyst 3900 to this configuration, remove the cable from the RI port on the first MAU and insert it into port 19 of the Catalyst 3900. Then, insert one end of a new cable into the RI port on the first MAU and insert the other end of the same cable into port 20 of the Catalyst 3900.

Note: The same redundancy can also be accomplished by connecting each of any two normal Token Ring ports to two different MAUs. RI/RO mode enables the use of the MAU RI/RO ports, saving normal MAU ports for attaching other stations.

The result is that port 19 is driving one path through the MAUs that eventually terminates at the receiver of port 20. Port 20 is driving the other path through the MAUs in the opposite direction and terminates at the receiver of port 19. Because both ports must be in the same VLAN (Token Ring Concentrator Relay Function [TrCRF]), the duplicate paths will be detected by the TrCRF's spanning tree and one port will be placed in blocking mode.

If you then removed a different cable from one of the MAUs, the TrCRF spanning tree would detect that the paths are no longer duplicates, the blocked port would be unblocked, and two rings would form. Because the two rings are still in the same TrCRF, the network continues to operate normally.

Dedicated Token Ring

Classic 4- and 16-Mbps Token Ring adapters must be connected to a port on a concentrator. These adapters are also limited to operating in HDX mode. In HDX mode, the adapter can only send or receive a frame; it cannot do both simultaneously.

Dedicated Token Ring (DTR), developed by the IEEE, defines a method in which the switch port can emulate a concentrator port, thereby eliminating the need for an intermediate concentrator. In addition, DTR defines a new FDX data-passing mode called transmit immediate (TXI), which eliminates the need for a token and allows the adapter to transmit and receive simultaneously.

DTR is particularly useful for providing improved access to servers. A server can be attached directly to a switch. This allows the server to take advantage of the full 16 Mbpf available for sending and receiving and results in an aggregate bandwidth of 32 Mbpf.



Speed Adaptation

In addition to supporting 4 Mbps and 16 Mbps, the Catalyst Token Ring switches can automatically configure the speed of a port by sensing the speed of the ring to which a port is connected.

With Token Ring, however, the speed cannot be changed without closing and reopening the port. Therefore, the following rules apply:

- If a port is closed, the speed can be changed without impact to the port or the network.
- If the port is open and running at a speed equal to the new speed specified, no action is taken.
- If the port is open and running at a speed different from the new speed specified, the port closes and reopens at the new speed. Closing and opening a port on an existing ring at a different speed from which the ring is operating will cause the port to issue a beacon on that ring.

Transmission Priority Queues

To address the needs of delay-sensitive data, such as multimedia, the Token Ring ports of the Catalyst switches have two transmission queues: high-priority and low-priority.

The queue for a frame is determined by the value of the priority field in the frame control (FC) byte of the frame. If FC priority is above a configurable level (the default is 3), the frame is put in the high-priority queue. If an output port becomes congested, you can configure the port to transmit all frames at high priority regardless of the FC byte contents.

The switch's CPU software monitors the size of the output queue at each Token Ring port to minimize the effects of congestion at output ports. When port congestion is detected, the switch does the following:

- Raises the transmit priority to a higher level for low-priority frames
- Discards the oldest frames when the output queue is almost full

Frame Filtering

Many bridged networks use filtering to reduce broadcast traffic, block protocols, and provide simple security. Often in Token Ring environments, dedicated gateways and servers are put on their own rings and filters are used to protect them from unnecessary broadcast traffic from other protocols. The Catalyst Token Ring switches allow users to configure filters based on both MAC address (destination and source address) and protocol (destination service access point [DSAP]/Subnetwork Access Protocol [SNAP] type). Because the filters are implemented in the hardware ASICs, filtering can be done at media speed on a per-port basis to control traffic to certain rings.

MAC address filters and broadcast filters can be applied only at input ports. DSAP and SNAP filters can be applied at input ports and output ports.

Broadcast Control

A common design in source-routing networks is parallel backbones. With source routing, the traffic tends to be distributed across both backbones, thereby providing both backup and load distribution. In some cases, these configurations suffer from excessive all-routes explorer (ARE) traffic as the explorer frames are duplicated on the many possible paths through the network. As a result, network managers have had to use hop counts and filters to manage this problem. Second-generation Token Ring switches support the automatic reduction of explorer traffic via the mechanism called ARE reduction.

ARE reduction ensures that the number of ARE frames generated by the switch does not overwhelm the network. The IEEE 802.1d SRT standard specifies the following optional ways of reducing the ARE explosion, which both involve examining the entire RIF to determine where the frame has been:

- The first method is based on whether the frame has been through the bridge before. This is determined by examining the routing information field of the received frame for a ring-bridge-ring combination that contains this bridge's number.
- The second method is based on whether the frame has been on any ring attached to the bridge before. This method is more restrictive than the first. Whether the frame has been on an attached ring is determined by examining the routing information field of the received frame for a LAN ID that matches any of the LAN IDs associated with the rings attached to the bridge.

The Catalyst Token Ring switches use the simpler of the two, which is to discard any ARE frame that has already been on a ring that is attached to the switch.

For example, an ARE frame from ring 1 is sent to switches A and B. The ARE frames are then forwarded to ring 2. When switch B receives the frame from switch A on ring 2, it examines the RIF and determines that this ARE has already been on ring 1. Because switch B is also attached to ring 1, the ARE is discarded.

ARE reduction requires no configuration and ensures that only 2 ARE frames (in this example) are received on each ring. The number of ARE frames will be equal to the number of parallel switches between the rings.

If a port on the switch fails or is disabled, the switch will no longer check for this ring number in the RIF. This allows alternate paths to the ring. Therefore, if there are two failures (for example, switch A to ring 1 and switch B to ring 4), there will still be a path between ring 1 and 4 (ring 1 to switch B to ring 2 to switch A to ring 4).