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 a 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 Mbps available for sending and receiving and results in an aggregate bandwidth of 32 Mbps.

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).

Features and Specifications of the Catalyst Token Ring Switches

Cisco offers two options in second-generation Token Ring switching: the Catalyst 3900 series Token Ring switches and the Catalyst 5000 Token Ring module. This chapter provides a brief overview of each switch and a list of the features and specifications of each switch.

This chapter provides the following information:

- Catalyst 3900 Overview
- Catalyst 3920 Overview
- Catalyst 5000 Series Token Ring Module Overview

Catalyst 3900 Overview

The Catalyst 3900 is a Token Ring switch ideally suited for desktop connectivity. The Catalyst 3900 comes standard with 20 fixed ports and an expansion slot that can accommodate two expansion modules. It also offers an optional Stack Port module that enables up to eight units to be stacked together using the Catalyst Matrix switch.

The Catalyst 3900 offers the following options for expanding beyond the base 20 ports:

- Port expansion modules

  The expansion slot can accommodate two port expansion modules (four-port copper or fiber), allowing you to add up to eight additional Token Ring ports to each switch.
- High-speed uplinks

  The expansion slot can accommodate two high-speed uplinks (Token Ring ISL or ATM) for high-speed connectivity between switches and to servers.
- Stack port module

  The stack port module allows you to connect 2 switches in a back-to-back configuration (for a maximum of 56 Token Ring ports) or up to 8 switches via 140-Mbps FDX links to the Catalyst Matrix switch (for a maximum of 224 ports and an aggregate switching capacity of more than 3 Gbps).
- TokenChannels and ISL Channels

  The Catalyst 3900 channel features allows you to configure TokenChannels and ISL Channels. Both types of channel configurations allow you to group up to 8 ports as one logical port for a high-speed connection between switches. These high-speed connections can be up to 256 Mbps for a TokenChannel and up to 800 Mbps for a ISL Channels

An ASIC design results in low-latency, wire-speed switching of unicast, multicast, and broadcast frames at either HDX or FDX speeds, regardless of whether they are source-route bridged, source-route transparently bridged, or source-route switched. Adaptive cut-through mode switching optimizes performance while providing protection from network errors by automatically switching to store-and-forward mode when errors reach a user-defined threshold.

The Catalyst 3900 switch provides a wide range of connectivity options for maximum flexibility. Connecting an MAU, server, or end station is easy because the Catalyst 3900 enables direct station attachment as well as RI/RO connections to scale ring segment size. Each port supports DTR, the IEEE standard that defines direct station attachment at 4, 16, or 32 Mbps. Furthermore, ports 19 and 20 and any of the ports of the fiber expansion module support RI/RO.

The shielded RJ-45 ports support both 150-ohm shielded twisted-pair (STP) and 100-ohm unshielded twisted-pair (UTP). There is no need for external media filters; nor is there a requirement for baluns to do impedance matching for different cable types.

Table 2-1 lists the Catalyst 3900 features and specifications.

Table 2-1 Catalyst 3900 Features and Specification

Performance	Latency: Less than 45 microseconds for all frame sizes
	Throughput: Media speed on all interfaces
Buffers and addressing	Buffers: 1 MB of DRAM per group of 4 ports
	Addressing: 10,000 addresses per system, local cache of up to 6500 addresses per group of 4 ports
System interfaces	20 shielded Token Ring ports for 150-ohm STP or 100-ohm UTP connectivity
	Expansion slot accommodating up to two expansion modules
	Expansion modules include a four-port fiber module, a four-port copper module, a two-port ISL uplink, and an ATM OC-3 uplink
	One rear stack port for an optional stack port module providing a 140-Mbps FDX link between back-to-back switches or between the Catalyst 3900 and the Catalyst Matrix switch for configurations requiring up to 8 units in a stack
	9-pin EIA/TIA-232 interface for local console or modem connectivity
Switching features	SRB, SRT, and source-route switching
	Adaptive cut-through mode switching
	17,848-byte Token Ring frame length support
	Automatic 4/16/32-Mbps speed adaptation
	Automatic shared and dedicated adaptation
	Two priority queues for multimedia traffic
	TokenChannel switch interconnect
	ISL Channel switch interconnect
	MAC address, DSAP, and SNAP type filters
	ARE reduction
	Explorer rate protection
	IEEE and IBM Spanning-Tree Protocols (STPs)

Table 2-1 Catalyst 3900 Features and Specification (Continued)

Standard MIBs supported	Management Information Base (MIB) for network management of TCP/IP-based internets: • MIB-II (RFC 1213) • Definitions of Managed Objects for Bridges (RFC 1493) • Evolution of Interfaces Group of MIB-II (RFC 1573) • Token Ring Extensions to the Managed Objects for Source Routing Bridges (RFC 1525) • IEEE 802.5 Token Ring MIB (RFC 1748) • RMON (RFC 1757) • Statistics, History, Alarm, and Event groups • RMON Token Ring Extensions (RFC 1513) – Token Ring extensions for Statistics, History, Alarm, and Event groups – Ring Station Order Group – Ring Station Control Table – Ring Station Table – Ring Station Config Control Table • IEEE 802.5 DTR Concentrator MIB • IEEE 802.5 DTR MAC MIB
Private MIBs supported	Catalyst 3900 Enterprise MIB
	Cisco VLAN Trunking Protocol MIB v2
	Cisco Discovery Protocol MIB
Monitoring support	CWSI graphical user interface (GUI) management • CiscoView with Threshold Manager • VlanDirector • TrafficDirector
	SPAN
	TFTP and BOOTP
	Menu-driven interface (via console port or telnet)
	Password-level security
Physical Specifications	Dimensions (H x W x D): 3.4 in. x 17.4 in. x 15.3 in. (8.6 cm x 44.2 cm x 38.7 cm)
	Weight: 16-18 lb (6-6.7 kg), depending on configuration
	Mounting: 19-in. (48.26 cm) 2U rack compatible
Power Requirements	Power: 90-264 VAC autosensing (single supply)
	Frequency: 47-63 Hz
	AC current rating: 1.5A at 115V; 0.75A at 230V
	Thermal dissipation: 150W maximum; 512 BTUs/hr
Environmental Conditions	Operating temperature: 50 to 104˚F (10 to 40˚C)
	Nonoperating temperature: -13 to 158˚F (-25 to 70˚C)
	Operating humidity: 8 to 80% (noncondensing)
	Nonoperating humidity: 8 to 90% at 45˚C
	Storage altitude: 40,000 ft
Electromagnetic emissions certifications	FCC Class A/B-UTP
	CE Declaration of Conformity to the EMC Directive-Class B with Unshielded or Shielded Cables
	VCCI Class II (B) Certification (for Japan)
	AS/NRZ 3548 (1992 Class A/B Certification for Australia)
	ICES-003 Class A/B (for Canada)
Safety certifications	UL 1950, third Edition without D3 deviations
	CUL to CAN/CSA 22.2 Number 950
	CE mark to the Low Voltage Directive (EN60 950, 1992 Amendments 1 and 2)
	Certified Body (CB) report to IEC 950, third Edition

Catalyst 3920 Overview

The Catalyst 3920 is a Token Ring switch is also ideally suited for desktop connectivity. The Catalyst 3920 comes standard with 24 fixed ports and an integrated stack port module that enables up to eight units to be stacked together using the Catalyst Matrix switch.

Table 2-2 lists the Catalyst 3920 features and specifications.

Table 2-2 Catalyst 3920 Features and Specifications

Performance	Latency: Less than 45 microseconds for all frame sizes
	Throughput: Media speed on all interfaces
Buffers and addressing	Buffers: 1 MB of DRAM per group of 4 ports
	Addressing: 10,000 addresses per system, local cache of up to 6500 addresses per group of 4 ports
System interfaces	24 shielded Token Ring ports for 150-ohm STP or 100-ohm UTP connectivity
	One rear integrated stack port module providing a 140-Mbps FDX link between back-to-back switches or between the Catalyst 3920 and the Catalyst Matrix switch for configurations requiring up to 8 units in a stack
	9-pin EIA/TIA-232 interface for local console or modem connectivity
Switching features	SRB, SRT, and source-route switching
	Adaptive cut-through mode switching
	17,848-byte Token Ring frame length support
	Automatic 4/16/32-Mbps speed adaptation
	Automatic shared and dedicated adaptation
	Two priority queues for multimedia traffic
	TokenChannel switch interconnect
	MAC address, DSAP, and SNAP type filters
	ARE reduction
	Explorer rate protection
	IEEE and IBM Spanning-Tree Protocols (STPs)
Standard MIBs supported	Management Information Base (MIB) for network management of TCP/IP-based internets: • MIB-II (RFC 1213) • Definitions of Managed Objects for Bridges (RFC 1493) • Evolution of Interfaces Group of MIB-II (RFC 1573) • Token Ring Extensions to the Managed Objects for Source Routing Bridges (RFC 1525) • IEEE 802.5 Token Ring MIB (RFC 1748) • RMON (RFC 1757) • Statistics, History, Alarm, and Event groups • RMON Token Ring Extensions (RFC 1513) – Token Ring extensions for Statistics, History, Alarm, and Event groups – Ring Station Order Group – Ring Station Control Table – Ring Station Table – Ring Station Config Control Table • IEEE 802.5 DTR Concentrator MIB • IEEE 802.5 DTR MAC MIB
Private MIBs supported	Catalyst 3900 Enterprise MIB
	Cisco VLAN Trunking Protocol MIB v2
	Cisco Discovery Protocol MIB

Table 2-2 Catalyst 3920 Features and Specifications (Continued)

Monitoring support	CWSI graphical user interface (GUI) management • CiscoView with Threshold Manager • VlanDirector • TrafficDirector
	SPAN
	TFTP and BOOTP
	Menu-driven interface (via console port or telnet)
	Password-level security
Physical Specifications	Dimensions (H x W x D): 1.7 in. x 17.4 in. x 11.1 in. (4.4 cm x 44.2 cm x 28.2 cm)
	Weight: 10 lb (4.5 kg)
	Mounting: 19-in. (48.26 cm) 1U rack compatible
Power Requirements	Power: 90-264 VAC autosensing (single supply)
	Frequency: 47-63 Hz
	AC current rating: 1.5A at 115V; 0.38A at 230V
	Thermal dissipation: 75W maximum; 256 BTUs/hr
Environmental Conditions	Operating temperature: 50 to 104˚F (10 to 40˚C)
	Nonoperating temperature: -13 to 158˚F (-25 to 70˚C)
	Operating humidity: 8 to 80% (noncondensing)
	Nonoperating humidity: 8 to 90% at 45˚C
	Storage altitude: 40,000 ft
Electromagnetic emissions certifications	FCC Class A/B-UTP
	CE Declaration of Conformity to the EMC Directive-Class B with Unshielded or Shielded Cables
	VCCI Class II (B) Certification (for Japan)
	AS/NRZ 3548 (1992 Class A/B Certification for Australia)
	ICES-003 Class A/B (for Canada)
Safety certifications	UL 1950, third Edition without D3 deviations
	CUL to CAN/CSA 22.2 Number 950
	CE mark to the Low Voltage Directive (EN60 950, 1992 Amendments 1 and 2)
	Certified Body (CB) report to IEC 950, third Edition

Catalyst 5000 Series Token Ring Module Overview

The Catalyst 5000 Series Token Ring module is a switching module you can use with any of the Catalyst 5000 series switches. The Token Ring module is available in fiber or copper. The copper Token Ring module provides 16 RJ-45 ports. The fiber Token Ring module provides 16 ST-type ports. On all Catalyst 5000 series switches interface slot 1 is reserved for the supervisor engine module.

The maximum number of Token Ring ports varies depending on the model of Catalyst 5000 switch as follows:

- Catalyst 5002 contains 2 slots, allowing a maximum configuration of 16 Token Ring ports.
- Catalyst 5000 contains 5 slots, allowing a maximum configuration of 64 Token Ring ports.
- Catalyst 5500 contains 13 slots, however, slot 13 is reserved for the ATM Switch Processor (ASP) module. Therefore, the maximum configuration of Token Ring ports is 176.

As in the Catalyst 3900, an ASIC design results in low-latency, wire-speed switching of unicast, multicast, and broadcast frames at either half- or full-duplex speeds, regardless of whether they are source-route bridged, source-route transparently bridged, or source-route switched.

Like the Catalyst 3900, the Catalyst 5000 Series Token Ring module supports IEEE 802.5r, which defines standards for the direct attachment of end stations to the switch as well as for the transmission of data at half-duplex (4/16 Mbps) and full-duplex (32 Mbps) speeds. The fiber Token Ring module also allows the ports to operate in RI/RO mode.

The shielded RJ-45 ports support both 150-ohm STP and 100-ohm UTP. There is no need for external media filters and there is no requirement for baluns to do impedance matching for different cable types.

Table 2-3 lists the Catalyst 5000 Series Token Ring module features and specifications.

Table 2-3 Catalyst 5000 Token Ring Module Features and Specifications

Performance	Throughput: Media speed on all interfaces
System interfaces	16 Token Ring ports for UTP/STP connectivity
	16 Token Ring ports for multimode 62.5-micron fiber connectivity
	Autosense 4/16/32 Mbps on all ports
	Switch ports can function as concentrator or station ports
Switching features	SRB, SRT, and source-route switching
	Automatic 4/16/32-Mbps speed adaptation
	Automatic shared and dedicated adaptation
	17,848-byte Token Ring frame length support
	Two priority queues for multimedia traffic
	MAC address, DSAP, and SNAP type filters
	ARE reduction
	IEEE and IBM STPs
Standard MIBs supported	MIB for network management of TCP/IP-based internets: • MIB-II (RFC 1213) • Definitions of Managed Objects for Bridges (RFC 1493) • Evolution of Interfaces Group of MIB-II (RFC 1573) • Token Ring Extensions to the Managed Objects for Source Routing Bridges (RFC 1525) • IEEE 802.5 Token Ring MIB (RFC 1748) • RMON (RFC 1757) Statistics, History, Alarm, and Event groups • RMON Token Ring Extensions (RFC 1513) – Token Ring extensions for Statistics, History, Alarm, and Event groups – Ring Station Order Group – Ring Station Control Table – Ring Station Table – Ring Station Config Control Table • AToM MIB (RFC 1695) • LEC MIB (ATM Forum LANE v. 1.0) • LECS, LES, BUS MIB
Private MIBs supported	Cisco VLAN Trunking Protocol MIB v2
	Cisco Discovery Protocol MIB
Monitoring support	CWSI GUI management • CiscoView with Threshold Manager • VlanDirector • TrafficDirector
	SPAN
	TFTP and BOOTP
	Command line interface (via console port or telnet)
	Password-level security and Terminal Access Controller Access Control System (TACACS)
Physical specifications	Single slot dimensions (H x W x D): 1.17 in. x 14.4 in. x 16.0 in. (2.97 cm x 36.58 cm x 40.64 cm)
	Weight: 3.9 lb (1.45 kg)

Table 2-3 Catalyst 5000 Token Ring Module Features and Specifications (Continued)

Electromagnetic emissions certifications	FCC 15J Class A
	VCCI CE II
	CE Mark
	EN 55022 Class B
	CISPR 22 Class B
Safety Certifications	UL 1950
	EN 60950
	CSA to C22.2 No. 950
	IEC 950

Interconnecting Switches

The Catalyst 3900 comes standard with 20 Token Ring ports. Using the two expansion slots and available expansion modules, you can increase the number of Token Ring ports to 28. The Catalyst 3920 comes with 24 Token Ring ports. The Catalyst 5000 series Token Ring switching module comes with 16 Token Ring ports. Using a Catalyst 5500 Switch, you can have up to 176 Token Ring ports.

You can create larger Token Ring port configurations by interconnecting the switches. This chapter discusses the options for interconnecting Catalyst Token Ring switches:

- Using Channel Configurations with Catalyst 3900s
- Stacking Catalyst 3900s
- Using ATM
- Using ISL

You can interconnect Catalyst 3920s with Catalyst 3900s. In this chapter, Catalyst 39*xx* is use to represent both.

Using Channel Configurations with Catalyst 3900s

The Catalyst 3900 allows you to configure two types of channels: ISL Channels and TokenChannels.

Channel configurations consist of two to eight parallel connections. These parallel connections provide the following:

- Logical aggregation of bandwidth of up to 256 Mbps (128 full duplex) for TokenChannel configurations and up to 800 Mbps (400 Mbps full duplex) for ISL Channels.
- Load balancing
- Fault tolerance

The Catalyst 3900 channel configurations are *fault-tolerant*. This feature enables channels to continue to function as long as there is at least one link active within the channel. This capability ensures that large portions of a network are not disrupted in the event a port or cable fails within the channel by transferring the traffic to one or more of the remaining ports in a channel.

You cannot use the ports of an ATM module in a channel.

Note: When the Catalyst 3900 is configured with channels, all broadcast frames use the lowest numbered active port of the channel.

Caution Physically disconnect or disable the Catalyst 3900 ports before configuring a TokenChannel or ISL Channel. Failure to disconnect or disable the ports might result in network loops.

TokenChannels

A TokenChannel consists of two to eight parallel connections between two Catalyst 3900s. These parallel channels provide improved performance between Catalyst 3900s.

A single TokenChannel can consist of a combination of HDX and FDX connections. For example, a TokenChannel consisting of three connections can have one HDX and two FDX connections. However, both ports in each interconnected pair must be either HDX or FDX. In addition, all ports in a single TokenChannel must belong to the same TrCRF on the Catalyst 3900.

For more information about TrCRFs, see the "Token Ring VLANs and Related Protocols" chapter.

Caution While you can use TokenChannels to interconnect Catalyst 3900s and Catalyst 3920s, you cannot use TokenChannels to interconnect other different models of switches. For example, you cannot use a TokenChannel to interconnect a Catalyst 2600 and a Catalyst 3900. Likewise, you cannot use a TokenChannel to interconnect a Catalyst 3900 and a non-Cisco switch.

ISL Channels

A single ISL Channel can consist of two to four Token Ring ISL ports. Configuring an ISL Channel provides Fast EtherChannel connectivity on the Catalyst 3900. You can configure an ISL Channel between two Catalyst 3900 switches or between a Catalyst 3900 switch and a Catalyst 5000, a Token Ring ISL-capable Cisco router, or a Token Ring ISL network adapter. All connections in an ISL Channel must be FDX.

Stacking Catalyst 3900s

A stack of Catalyst 39*xx* switches is not just a connection of several switches. A Catalyst 39*xx* stack of switches combine to form a "virtual" single switch.

A Catalyst stack is configured in one of the following two ways:

- Two Catalyst 39*xx* switches cabled together in a back-to-back configuration.
- A stack of up to 8 Catalyst 39*xx* switches connected together via a Catalyst Matrix.

Two Catalyst 39*xx* switches can be connected to form a stack by using only a stack port cable and an interface card (Catalyst stack port module) plugged into the back of each Catalyst 39*xx*. This creates a direct connection between the two Catalyst 39*xx* switches, which is referred to as a *back-to-back* stack. As an alternative, you can use an 8-port Catalyst Matrix switch to create a stack of up to eight Catalyst 39*xx* switches. The ProStack port operates in FDX mode at speeds of 140 Mbps. It switches packets at wire speeds with low forwarding latency. A proprietary 4-byte header is used to allow the members of the stack to function as one operational system.

Note: The stack port module is an optional feature on the Catalyst 3900 switch. On the Catalyst 3920 switch, the stack port module is an integrated feature.

When you power-on a Catalyst 39*xx*, it runs through a set of self-diagnostics. Immediately after the diagnostics are completed, the Catalyst 39*xx* runs through a *stack discovery mode*. This discovery mode senses whether the switch is cabled to another Catalyst 39*xx*. If it is determined during the discovery mode that the Catalyst 39*xx*

is connected to other switches, the switches automatically combine to form a stack. At the end of the discovery mode, if it is determined that the Catalyst 39*xx* is not connected to another switch, the Catalyst 39*xx* operates as a standalone switch.

Advantages of the stack include the following:

- Manage the entire stack as a single device
- Single SNMP image for entire stack
  - Easier to customize SNMP applications
- Distributed intelligence between the switches of the stack
  - Shared learning
  - Shared management information
- Hot swap of stack switches
  - When a switch is powered off or removed from the stack, the other switches reform as a stack

Forming a Back-to-Back Catalyst Stack

A proprietary shielded cable, 1 meter in length, with 50-pin connectors, is used to connect the Catalyst switches together. After power-on diagnostics, the stack discovery mode runs. If, during this stack discovery mode, a Catalyst 39*xx* detects that it is connected to another Catalyst 39*xx* in a back-to-back configuration, the two switches will begin to form a stack.

As soon as the stack discovery mode is completed, two things happen:

- Each Catalyst 39*xx* is assigned a box number.
  - The two Catalyst 39*xx* switches in a back-to-back stack become box 1 and box 2. The box number is determined by the MAC address of each Catalyst 39*xx*. The Catalyst 39*xx* with the lower MAC address becomes box 1, and the Catalyst 39*xx* with the higher MAC address becomes box 2.
- The Catalyst 39*xx* switches must combine configuration information so that both of the boxes, as a stack, will use certain common parameters. This common information is called the *interbox parameters*. The "Interbox Parameters" section later in this chapter lists the shared parameters. In a stack of switches, one of the switches must become the provider of the interbox parameters.
  - If the Catalyst 39*xx* switches have the *same* configuration information (whether they are new or have been preconfigured to be the same) when they begin to form a stack, the Catalyst 39*xx* that becomes box 1 also becomes the provider of the interbox parameters.
  - If the configuration information differs between the two Catalyst 39*xx* switches, the first switch up provides its configuration information to the other switch. If both switches come up simultaneously, an error message is displayed that instructs the user to briefly press the SYSREQ button on the switch that contains the desired configuration.

After a stack has formed and sets up the interbox parameters, the stack operates the same way whether it is in a back-to-back configuration or is in a multi-unit configuration using the Catalyst Matrix interface.

Creating a Multi-Unit Catalyst Stack with a Catalyst Matrix Interface

Using a Catalyst Matrix, you can create a multi-unit stack of up to 8 Catalyst 39*xx* switches. The following sections describe how this multi-unit stack is formed:

- Catalyst Matrix Description
- Forming a Multi-Unit Catalyst Stack

Catalyst Matrix Description

The Catalyst Matrix is an eight-port switch matrix interface that connects up to eight Catalyst 39*xx* switches. The Catalyst 39*xx* senses if it is connected to a Catalyst Matrix and also senses if there are other Catalyst 39*xx* switches connected to that Catalyst Matrix. The connected Catalyst 39*xx* switches and the Catalyst Matrix combine logically to form a stack.

Any combination of up to eight Catalyst 39*xx* switches can be connected to or disconnected from the Catalyst Matrix while it, or any of the switches, are powered on or powered off. A proprietary shielded cable, 1 meter in length, with 50-pin connectors, is used to connect the Catalyst stack equipment together. The cable has *cross-over* wiring so either end can connect to the Catalyst Matrix, or to the Catalyst 39*xx* switches. The cable is plugged directly into a stack port I/O connector on the back of the Catalyst Matrix. The other end is plugged into a Catalyst stack port module interface card that is installed in the rear expansion slot in the Catalyst 39*xx*.

Forming a Multi-Unit Catalyst Stack

When Catalyst 39*xx* switches first power up, they run through a set of self-diagnostics. Immediately after the diagnostics are completed, the Catalyst 39*xx* switches run through a stack discovery mode. During this stack discovery mode, if two or more Catalyst 39*xx* switches are connected to a Catalyst Matrix, the switches detect the connection and combine logically to create a stack configuration.

As soon as the stack discovery mode is completed, each Catalyst 39*xx* is assigned a box number. With a Catalyst Matrix configuration, the box number for a Catalyst 39*xx* is determined by the port number the Catalyst 39*xx* is connected to on the Catalyst Matrix. For example, the Catalyst 39*xx* plugged into port 3 on the Catalyst Matrix becomes box 3. The box number remains constant as long as that switch is plugged in to that port. If a Catalyst 39*xx* is moved to another port, the box number for that Catalyst 39*xx* will change to the number of the port it is moved to.

Note: The switch with the lowest box number becomes the controlling switch.

For a stack to operate as a single entity, the interbox parameters must be the same in all of the switches in a stack. The "Interbox Parameters" section later in this chapter lists the shared parameters. There are two possible ways of providing configuration information to the Catalyst 39*xx* switches in a stack. These methods are as follows:

- Preconfigure all the Catalyst 39*xx* switches with the same parameters.
- Allow one of the switches to provide the configuration information to the other switches in the stack.

The first switch that comes up provides the initial configuration to the rest of the switches. If the switches come up simultaneously and their configurations differ, a warning message is displayed that instructs the user to briefly press the SYSREQ button on the switch that contains the desired configuration. Pressing the SYSREQ button causes the selected switch to send out its configuration information to the other switches in the stack.

Note: If you press the SYSREQ button for more than a few seconds, the System Request menu is displayed. If this happens, exit the System Request menu and then briefly press the SYSREQ button.

If the Catalyst 39*xx* switches are already powered on and *then* connected together, the same procedure as described above occurs, except that because the switches are already powered up and functioning, they will continue to perform their previous internal switching functions. While the normal internal switching functions are still operating, a split stack is formed. Once the split stack is formed, the console displays the same warning message, instructing the user to press the SYSREQ button of the switch that contains the desired configuration.

If Catalyst 3900s have formed a stack and any additional Catalyst 39*xx*s are added to the stack, the new switches will join the existing stack by altering their interbox parameters to match those of the existing stack.

After a stack has formed and sets up the interbox parameters, the stack operates the same way whether it is in a back-to-back configuration or is in a multi-unit configuration using the Catalyst Matrix interface.

Interbox Parameters

When a stack is formed, certain configuration information within all of the different Catalyst 39*xx* switches must combine to form a common configuration (interbox parameters). The stack operates as a single entity when all of the Catalyst 39*xx* switches in that stack use the same interbox parameters.

The following is a list of these shared interbox parameters:

- IP Configuration
  - IP address
  - Default gateway
  - Subnet mask
  - IP state
- Spanning-Tree Protocol (STP)
  - Participation
  - Switch priority
  - Port priority
  - Port cost
  - Maximum message age
  - Hello time
  - Forward delay
- VLAN information
- Limited Multicast Filters
- System password
- Console time-out
- Telnet configuration
  - Number of Telnet sessions allowed
  - Whether new Telnet sessions are allowed
- TFTP download
  - TFTP VLAN
  - TFTP server address
  - TFTP download filename
- Switch and stack information
  - Stack time-out
  - System name
  - System contact
  - System location
- SNMP configuration
  - Authentication traps
  - Trap table
  - Community name table

Using ATM

ATM is a cell-switching and multiplexing technology that combines the benefits of circuit switching (constant transmission delay and guaranteed capacity) with the benefits of packet switching (flexibility and efficiency for intermittent traffic).

ATM switches transmit small units of data called *cells*. The latency in a cell switch is very small because of the short cell size. Short cells have a tiny store-and-forward delay. In the absence of port contention and buffering, cells are switched quickly in hardware.

Both the Catalyst 3900 and the Catalyst 5000 offer ATM modules that you can use to interconnect your switches.

The Catalyst 3900 ATM expansion module is a single-port expansion module that provides high-speed connectivity between the Catalyst 3900 Token Ring switch and an ATM backbone network. The ATM expansion module has a transmission speed of 155 Mbps over a multimode fiber interface using a duplex subscriber connector (SC). The Catalyst 3900 ATM expansion module supports up to 63 ELANs and up to 2048 virtual channel connections (VCCs).

The Catalyst 5000 offers a variety of ATM expansion modules. These modules offer a transmission speed of 155 Mbps over a single-mode fiber connection, a multimode fiber connection, or a UTP connection. Each module supports up to 63 ELANs and up to 4096 VCCs.

For more information about ATM, see the "ATM and Token Ring LANE" chapter.

Using ISL

ISL was originally developed for Ethernet switches. It uses a Fast Ethernet interface to provide connectivity between switches and extends the VLAN capabilities of the switch by tagging the standard Fast Ethernet frame with the necessary VLAN information. Like ATM, ISL can provide a high-speed link between switches. Unlike ATM, ISL forwards the data across the high-speed link without breaking the frames into cells. The entire frame is sent intact across the ISL connection.

An ISL port is considered a trunk port. A *trunk* is a physical link that carries the traffic of multiple VLANs between two switches or between a switch and a router, thereby allowing the VLANs to be extended across switches. Trunks use high-speed interfaces such as Fast Ethernet, FDDI, or ATM.

The ISL protocol is a packet tagging protocol that contains a standard Ethernet or Token Ring frame and the VLAN information associated with that frame. Some additional information is also present in the frame. For more information about the ISL frame format, see the "Frame Formats" appendix.

The ISL backbone design looks much like an ATM design; however, ISL is less expensive than ATM and avoids the need for LAN emulation (LANE) services. It is primarily intended for network managers who do not want an ATM backbone for the campus. Routing between Token Ring VLANs is provided via an ISL-attached router or the Catalyst 5000 series Route Switch Module (RSM). Figure 3-1 shows an example of the ISL backbone design.

Figure 3-1 ISL Backbone Design

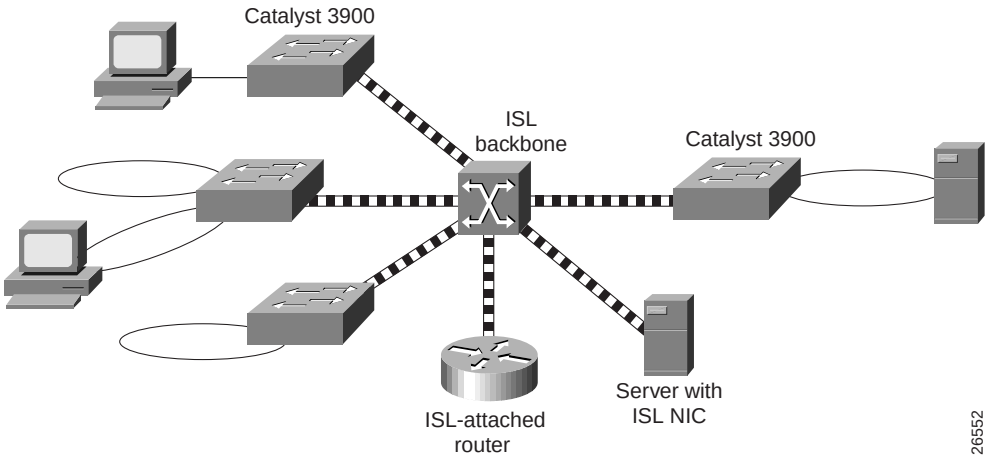

The Catalyst 5000 family of switches should be used to provide the ISL backbone. The Catalyst 3900 switch can then be connected to this backbone via the dual 100-Mbps ISL expansion module. In addition, vendors provide ISL network interface cards (NICs) that support both Token Ring and Ethernet VLANs. These NICs can be used for high-speed attachment to servers.

For the Catalyst 5000, any of the ports on many of its Fast Ethernet modules can be configured as trunk ports that use ISL.

The Catalyst 3900 2-port 100-Mbps Token Ring ISL module supports the encapsulation of Token Ring frames on a standard Fast Ethernet link to allow VLANs to be distributed across multiple platforms and devices. The module is available with a fiber or UTP copper media interface. The ports of the ISL module can be connected to the ports of another ISL module in another router or switch.

If you want to attach the Catalyst 3900 ISL port to the ISL port of a Catalyst 5000, you must manually configure the ISL port on the Catalyst 5000 for 100 Mbps (using the **set port speed** command) and full-duplex (using the **set port duplex** command).

Note: The ISL module does not support MAC or protocol filtering.

Using ISL in Parallel Configurations

While your Catalyst 3900 can contain both an ATM expansion module and an ISL expansion module, use caution when using ISL in a parallel configuration with ATM or Token Ring. Because the Catalyst 3900 supports the propagation of VLAN trunking information via ISL connections only, it is important that the ISL connection be the active path in an ISL-ATM parallel connection.

If the ISL module is configured in parallel connections with ATM or Token Ring, the STP allows only one active port at a time. When using the default Catalyst 3900 STP values, the path cost is calculated based on a 200-Mbps connection that results in a path cost of 5 and causes the STP to place the ISL port in forwarding mode and the ATM port or the Token Ring port in blocked mode.

However, modifying the Catalyst 3900 port STP values or using devices from other vendors that use different STP values can block the ISL port. If an ISL port becomes blocked in an ISL-ATM parallel connection, traffic passes via the ATM link, but VLAN trunking data is not passed. Also, if your STP configuration makes an ATM or Token Ring port the forwarding path to the root switch instead of the ISL link, the switch on the other end of a blocked ISL port might incorrectly limit AREs to the incoming TrCRF.

Therefore, when modifying STP values, always ensure that the STP port path costs are configured so that the ISL port is the preferred path. In an ISL parallel configuration, a Token Ring or ATM link should never have a lower cost to the root bridge than the ISL link.

Token Ring VLANs and Related Protocols

A VLAN is a logical group of LAN segments, independent of physical location, with a common set of requirements. For example, several end stations might be grouped as a department, such as engineering or accounting. If the end stations are located close to one another, they can be grouped into a LAN segment. If any of the end stations are on a different LAN segment, such as different buildings or locations, they can be grouped into a VLAN that has the same attributes as a LAN even though the end stations are not all on the same physical segment. The information identifying a packet as part of a specific VLAN is preserved across a Catalyst switch connection to a router or another switch if they are connected via trunk ports, such as ISL or ATM.

Any VLAN can participate in the STP. The protocol used depends on the type of VLAN and the type of bridging function used.

This chapter provides an overview of the following:
- Token Ring VLANs
- VLAN Trunking Protocol
- Duplicate Ring Protocol
- Spanning-Tree Protocol

Token Ring VLANs

Because a VLAN is essentially a broadcast domain, a Token Ring VLAN is slightly more complex than an Ethernet VLAN. In transparent bridging there is only one type of broadcast frame and therefore only one level of broadcast domain, but in source routing there are multiple types of broadcast frames that fall into two categories:
- Those that are confined to a single ring
- Those that traverse the bridged domain

These two categories of broadcast frames result in a broadcast domain that is hierarchical in nature, as a local ring domain can exist only within a domain of all the inter-connected rings.

In a Token Ring VLAN, logical ring domains are formed by defining groups of ports that have the same ring number. The IEEE calls such a port group a Concentrator Relay Function (CRF). On Catalyst switches, such a grouping of Token Ring ports is called a Token Ring CRF (TrCRF).

The domain of inter-connected rings is formed using an internal multiport bridge function that the IEEE calls a Bridge Relay Function (BRF). On Catalyst switches, such a grouping of logical rings is called a Token Ring BRF (TrBRF).

Figure 4-1 illustrates TrCRFs and a TrBRF within a Catalyst Token Ring switch or module.

Figure 4-1 Token Ring VLANs

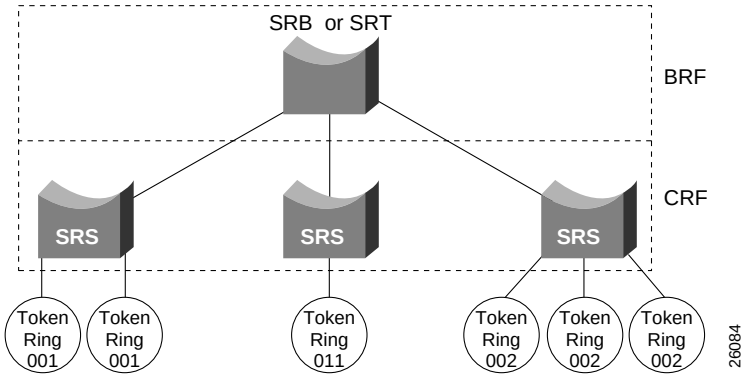

TrCRFs

A TrCRF is a logical grouping of ports. Within the TrCRF, source-route switching is used for forwarding based on either MAC addresses or route descriptors. Frames can be switched between ports within a single TrCRF.

A TrCRF has two global parameters: a ring number and a parent TrBRF identifier. On the Catalyst 3900, the ring number of the TrCRF can be defined or learned from external bridges. On the Catalyst 5000, the ring number must be defined.

As a rule, a TrCRF is limited to the Token Ring ports of a single Catalyst 5000 series switch, the ports of a single Catalyst 3900, or the ports within a stack of Catalyst 3900 switches. This type of TrCRF is called an *undistributed* TrCRF. However, if your switches are connected via ISL, the Cisco Duplicate Ring Protocol (DRiP) allows additional types of TrCRFs to be configured and these types of TrCRFs can have ports of a single TrCRF located on different switches. On the Catalyst 5000 series switch, these types of TrCRFs are the *default*, *distributed*, and the *backup* TrCRF. On the Catalyst 3900, these types of TrCRFs are the default and backup TrCRF.

For more information about DRiP, see the “Duplicate Ring Protocol” section.

Undistributed TrCRF

The undistributed TrCRF is located on one switch and has a logical ring number associated with it. Multiple undistributed TrCRFs located on the same or separate switches can be associated with a single parent TrBRF. The parent TrBRF acts as a multiport bridge, forwarding traffic between the undistributed TrCRFs.

Figure 4-2 illustrates the undistributed TrCRF.

Figure 4-2 Undistributed TrCRF

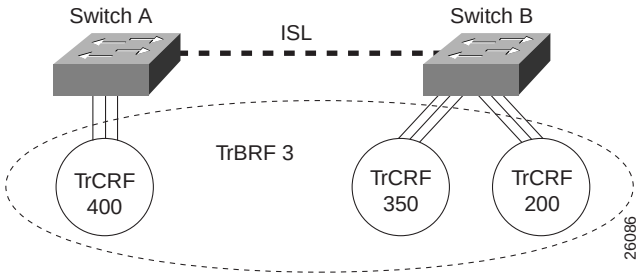

Default and Distributed TrCRFs

As a rule on the Catalyst 3900, TrCRFs cannot span separate switches or stacks of switches. One exception to this rule is the *default* TrCRF. The default TrCRF can contain ports located on separate switches. By default, the Token Ring VLAN configuration on the Catalyst 3900 and the Catalyst 5000 series Token Ring modules has all ports assigned to the default TrCRF (1003). In turn, this default TrCRF is associated with the default TrBRF (1005), which can span switches via ISL. If a user does not configure the ports of a Token Ring module to be associated with a new TrCRF, traffic is passed between the default TrCRFs located on separate switches that are connected via ISL.

Because the default TrCRF is the only TrCRF that can be associated with the default TrBRF, the default TrBRF does not perform any bridging functions, but uses source-route switching to forward traffic between the ports of the TrCRF.

Figure 4-3 illustrates the default TrCRF.

Figure 4-3 Default TrCRF

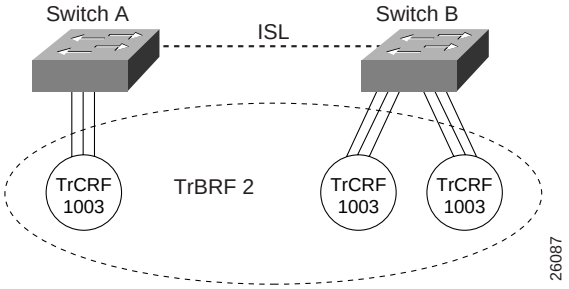

In addition to the default TrCRF, the Catalyst 5000 series Token Ring module supports the configuration of a distributed TrCRF. A distributed TrCRF contains ports on different switches as illustrated in Figure 4-4. While you can distribute the ports of a TrCRF across Token Ring modules on separate Catalyst 5000 series switches, we recommend that you use caution when configuring a distributed TrCRF other than the default TrCRF (1003). Always ensure that there are no loops configured in your network before configuring a distributed TrCRF.

Note: Before you can configure a distributed TrCRF on the Catalyst 5000 series Token Ring module, you must enable the configuration using the **set tokenring distrib-crf** command.

Figure 4-4 Disbributed TrCRF

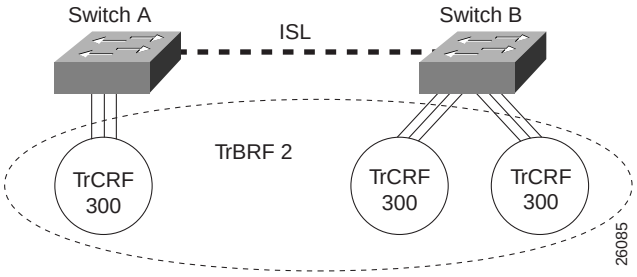

Backup TrCRF

The *backup* TrCRF enables you to configure an alternate route for traffic between undistributed TrCRFs located on separate switches that are connected by a TrBRF. The backup TrCRF is only used if the ISL connection between the switches becomes inactive.

While a TrBRF can contain multiple TrCRFs, it can contain only *one* TrCRF that is configured as a backup TrCRF. The backup TrCRF can contain only *one* port from each related switch. If you have more than one TrBRF defined on a switch, you can have more than one backup TrCRF defined on a switch (one defined for each TrBRF).

To create a backup TrCRF, create the TrCRF, assign it to the TrBRF that traverses the switches, mark it as a backup TrCRF, and then assign one port on each switch to the backup TrCRF.

Caution If the backup TrCRF port is attached to a Token Ring MAU, it will not provide a backup path unless the ring speed and port mode are set by another device. Therefore, we recommend that you manually configure the ring speed and port mode for the port assigned to the backup TrCRF.

Under normal circumstances only one port in the backup TrCRF is active. The active port is the port with the lowest MAC address. If the ISL connection between the switches become inactive, the port that is a part of the backup TrCRF on each affected switch will automatically become active, and will reroute traffic between the undistributed TrCRFs through the backup TrCRF. When the ISL connection is reestablished, all but one port in the backup TrCRF will be disabled.

Figure 4-5 illustrates the backup TrCRF.

Figure 4-5 Backup TrCRF

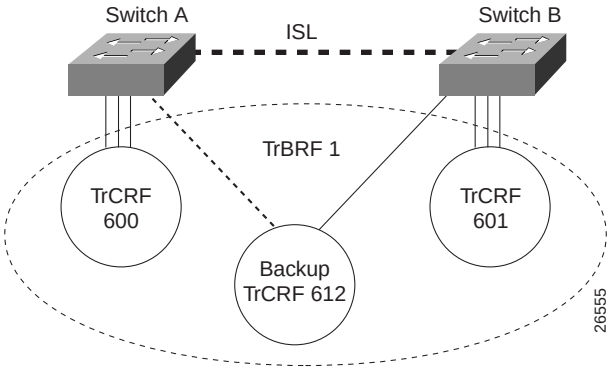

TrBRFs

A TrBRF is a logical grouping of TrCRFs. The TrBRF is used to join different TrCRFs contained within a single Catalyst 3900, a stack of Catalyst 3900s, or the Token Ring modules of a single Catalyst 5000 switch. In addition, the TrBRF can be extended across a network of switches via high-speed ISL uplinks to join TrCRFs configured on different switches.

A TrBRF has two global parameters: a bridge number and a bridge type. The bridge number is used to identify the logical distributed SRB, which interconnects all logical rings that have the same parent TrBRF.

A TrBRF can function as an SRB or SRT bridge running either the IBM or IEEE STP. If SRB is used, duplicate MAC addresses can be defined on different logical rings.

To accommodate SNA traffic, you can use a combination of SRT and SRB modes. In a mixed mode, the TrBRF considers some ports (logical ports connected to TrCRFs) to be operating in SRB mode while others are operating in SRT mode.

VLAN Trunking Protocol

The Cisco VLAN Trunking Protocol (VTP) enables you to set up and manage VLANs across an entire VTP *management domain* (also known as an *administrative domain*). An administrative or management domain is a logical grouping of VLANs used by the VTP for the purpose of administration and management. VTP parameters are propagated throughout the VLANs within a single management domain. While you can have duplicate VLAN names in a network, each VLAN name within a management domain must be unique. A management domain is not device specific. Different devices may belong to the same management domain if the VLANs defined for the devices belong to the same management domain. Likewise, a device may belong to multiple management domains if the VLANs defined for the device belong to different management domains.

When new VLANs are added to a device (a Cisco router or switch) in a management domain, you can use VTP to automatically distribute the information via trunk ports to all of the devices in the management domain. This distribution ensures VLAN naming consistency and connectivity between all devices in the domain by allowing each device in the domain to learn of any new VLANs added to other devices in the domain or to learn of any changes made to existing VLANs in the domain.

Often, VTP is not used in Ethernet environments, but it is important in Token Ring environments because it ensures the distribution of TrCRF information.

VTP advertisements are transmitted on all trunk connections, including the following:
- ISL—Catalyst 5000 and Catalyst 3900
- ATM LANE—Catalyst 5000 only

With the Catalyst 3900 Release 4.1(1), *VTP pruning* is supported on the Catalyst 3900 switch. VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic only to those trunk links that the traffic must use to access the appropriate network devices. When a VLAN is pruned on an ISL trunk link, that trunk does not transmit frames destined for that VLAN.

Note: VTP pruning does not prune traffic for VLANs that are not eligible. VLAN 1, the default TrBRF (1005), and the default TrCRF (1003) cannot be configured to be pruning eligible. Therefore, the traffic from these VLANs cannot be pruned. Pruning eligibility is configured on a TrBRF basis. Therefore, if you configure a TrBRF other than the default TrBRF to be pruning eligible, all TrCRFs associated with the TrBRF are pruning eligible as well.

How VTP Works

VTP sends advertisements to a multicast address so the advertisements are received by all neighboring devices (but they are not forwarded by normal bridging procedures). The advertisement lists the sending device's VTP management domain, its configuration revision number, the VLANs known to the sending device, and parameters for each of the known VLANs. By receiving these advertisements, all devices in the same management domain learn about any new VLANs now configured in the transmitting device. Therefore, a VLAN can be created and configured on one device in the management domain and the information is automatically learned by all the other devices in the same management domain.

The switch can operate in three different VTP modes: server, client, or transparent.

In server mode, the switch permits changes to the management domain's global VLAN configuration from the local device. Redundancy in a network domain is created by using multiple VTP servers.

In client mode, the switch accepts configuration changes from other devices in the management domain, but does not permit local changes to the database.

In transparent mode, the switch forwards any VTP packets received on the default VLANs of any trunk onto the default VLANs of all other trunks. Use VTP transparent mode to enable a Catalyst switch to propagate VTP information even if it is not participating in VTP. In transparent mode, VTP packets received on one trunk are automatically propagated unchanged to all other trunks on the device but are ignored on the device itself.

VTP Clients and Servers

To retain the VLAN information contained in VTP advertisements across reboots and network outages, a subset of the devices must be able to recover all information currently contained in advertisements after they reboot. In large, heterogeneous networks, the amount of information in the advertisements may be beyond the nonvolatile storage capabilities of some devices; however, storage of this same information in every device is normally beyond the required amount of redundancy. Therefore, each device in a VTP management domain is categorized as a VTP client or a VTP server.

VTP servers must be able to recover all the VLAN information in current VTP advertisements from nonvolatile-storage after they reboot. If they cannot, the device ceases being a VTP server and becomes a VTP client.

Under normal circumstances, VTP clients accept changes to the current VLAN information only through VTP advertisements. They do not accept changes via a console interface or SNMP. Upon boot up, the VTP client sends out periodic requests for VTP information on all of its trunks until it receives a summary advertisement from a neighbor. It uses that summary advertisement to determine whether its currently stored configuration is obsolete. If the stored configuration is obsolete, the client requests all VTP information from the neighbor.

If no VTP advertisement is received within a specified time, the VTP client can use the locally configured VLAN information, but will not issue VTP advertisements containing this information. This locally configured information is overridden (but may or may not be overwritten) as soon as the client receives a VTP advertisement. Thus, when a network is partitioned so that there are no VTP servers in a partition and all the VTP clients in that partition are rebooted, then no VTP advertisements are transmitted in that partition.

Upon boot up, the VTP server attempts to recover the information contained in VTP advertisements from nonvolatile-storage. Prior to successful recovery, the device can act only as a VTP client. The nonvolatile-storage used to hold the information can be either:

- The device's own nonvolatile random-access memory (NVRAM), which it must write immediately upon learning of any change in the information.
- A configuration file, which the device downloads via TFTP after a reboot.

In a large heterogeneous network, only a few devices need to be VTP servers. The choice of which devices are servers should be made based on the capabilities of each device and the amount of redundancy required. In a small network, all devices are normally VTP servers.

Caution When a device that is configured to operate in server mode is added to a VTP domain and the configuration of the new device is more current than that of the other devices in the network, all the VLAN information in the other devices will be overwritten. Therefore, exercise care when adding a device that is configured to operate in server mode to a VTP management domain.

VTP Advertisement Messages

In VTP, the following message types are defined:

- Advertisement Request (Advert-Request)—Request for VTP information.
- Summary Advertisement (Summary-Advert)—Message that advertises the existence of new VTP information.
- Subset Advertisement (Subset-Advert)—Message that contains the details of new VTP information.

All messages are sent as link-layer multicast frames.

Advertisement Requests

An Advert-Request is sent after a reboot, and when any of the following occur:

- A Subset-Advert message containing a configuration revision number that is higher than the device's current value is received.
- A Summary-Advert message containing a configuration revision number that is greater than the device's current value and a zero Subset-Adverts is received.
- The expected number of Subset-Advert messages is not received within a time after a Summary-Advert containing a configuration revision number that is greater than the device's current value is received. In this case, the Advert-Request is set to request only those VLANs that were missed. The Start-Value of the frame is set to a value one greater than the ISL VLAN ID of the VLAN contained in the last Subset-Advert received.
- A Summary-Advert containing a configuration revision number that is more than one greater than the device's current value.

After an Advert-Request is sent, a timer is started. The timer has a timeout period of a random value from 0 to 1 second. If the timer expires before a Summary-Advert is received, then another Advert-Request is sent.

An Advert-Request normally requests information on all VLANs; it can, however, request information on only a subset of the VLANs.

Summary and Subset Advertisements

An advertisement is sent in the following situations:

- Immediately upon a change in configuration (via console or SNMP) of VLAN information.
- When no other Summary-Advert with the current configuration revision number has been received for a timeout period. The timeout period is truncated to a small random value by the receipt of an Advert-Request message.

Each advertisement consists of one Summary-Advert immediately followed by zero or more Subset-Adverts:

- A Summary-Advert contains the management domain name, the configuration revision number, the update timestamp and identity, the authentication checksum, and the number of Subset-Advert messages that follow it.
- A Subset-Advert contains all information for one or more VLANs, and indicates its own sequence number with respect to any additional Summary-Advert messages.

The number of Subset-Advert messages that follow a Summary-Advert is determined according to the reason for sending the advertisement as in the following situations:

- When neither this device nor any other device has recently sent an advertisement, the Summary-Advert is followed by zero Subset-Advert messages.
- When a configuration change has been made, the Summary-Advert is followed by the minimum number of Subset-Advert messages required to contain all information on all VLANs (ordered in ascending order of ISL VLAN ID).
- When an Advert-Request for information on all VLANs was received, the Summary-Advert is followed by the minimum number of Subset-Advert messages required to contain all information on all VLANs (ordered in ascending order of ISL VLAN ID).

Configuration Change Indicators

Because the information required for each VLAN could potentially exceed 40 octets, it is impossible to contain all the information for all VLANs in a single MAC frame. It is not necessary for all the information to be sent in each advertisement. Regular, periodic advertisements only need to indicate that nothing has changed.

Occasionally, all information on all VLANs must be transmitted as a sequence of frames. This always occurs after a configuration change. It can also occur when it is requested by one of the devices, either because that device failed to receive one or more of the sequence of frames or because it has just restarted.

The following indications are used in advertisements to indicate whether the configuration has changed:

- Configuration revision number
- Authentication checksum

Configuration Revision Numbers

The device's VTP configuration revision number is incremented each time the device is reconfigured by management (console or SNMP) to do the following:

- Define a new VLAN
- Delete an existing VLAN
- Suspend or resume an existing VLAN
- Modify the parameters of an existing VLAN

The current time and the local device's identity (one of its IP-addresses) are recorded and included in the next VTP advertisement.

When a device receives an advertisement, the following actions occur:

If the configuration revision number in the advertisement is	Then
Less than that of the receiving device	The advertisement is ignored.
The same as that of the receiving device	• If the checksum of the advertisement is exactly the same as the checksum of the current configuration known to the device, then no action is taken. • Otherwise, the device's configuration remains unaffected, but the device indicates to management that a configuration error condition has occurred.
Greater than that of the receiving device and the advertisement's checksum and configuration information match	• If the set of VLANs and their parameters known to the device would be inconsistent if updated based on the information in the advertisement, then the device's configuration is unaffected, and the device indicates to management that a configuration error condition has occurred. • Otherwise: – Any VLAN in the advertisement unknown to the device is learned. – Any VLAN in the advertisement known to the device, but with different parameters, is updated to have the parameters from the advertisement. – Any VLAN known to the device, but not in the advertisement is forgotten by the device. Any static ports currently assigned to that VLAN are disabled. For any dynamic ports currently assigned to that VLAN, the server is queried for a new assignment. – The device's configuration revision number is updated to that of the advertisement. – New values for the "update timestamp" and "updater identity" are obtained from the advertisement. – The VTP advertisement is regenerated on each of the device's trunk ports other than the one on which it was received.

The time required to propagate new information across all devices is typically on the order of milliseconds, or, at most, a few seconds. However, it can be longer if some devices are temporarily partitioned (because of a break in the network). When a set of devices is partitioned for a prolonged period, a device in each partition should be updated. When the partition is repaired, the configuration in the set of devices with the greater configuration revision number takes precedence. If, however, devices in both partitions have the same configuration revision level, a configuration error is indicated.

Checksum

A checksum is defined to ensure that two different configurations with the same configuration revision number (which can occur, for example, after a network partition) are recognized as being different.

The checksum is calculated using an arbitrary security value that is appended to the front end and the back end of the data in a VTP configuration. When a VTP device has received all of the parts of the VTP configuration, it recalculates the checksum using its own security value derived from the password that has been configured locally. The device will not accept the new configuration if the checksums do not match.

On all Cisco VTP devices, the default initial configuration of the security value is all zeroes. Therefore, VTP devices will always accept one another's VLAN configurations as long as none of the security values on any of the devices have been modified. To make use of the security feature, a password needs to be set. The password must be the same for the management domain on all devices in the domain. Neither the password nor the security value itself is ever advertised over the network.

Caution If you use passwords, the same management domain password must be assigned to each Catalyst switch in the domain. Otherwise, the management domain will not function properly.

Transmission of Advertisements

VTP advertisements are transmitted using a multicast destination MAC address (0100.0CCC.CCCC) and are not forwarded using normal bridging techniques. A switch regenerates a VTP advertisement to all of its other trunk ports if the advertisement contains new configuration information.

Advertisements are transmitted on the default VLAN, which corresponds to the type of trunk link. Thus, only one copy is transmitted on a trunk port, no matter how many VLANs are defined.

VTP Advertisement Frame Format

VTP is assigned the Cisco High-level Data Link Control (HDLC) protocol type value of 0x2003. A Cisco-proprietary SNAP value enumerates HDLC protocol type values so VTP can run on all media that support SNAP, such as LAN media, Frame Relay, and ATM.

The SNAP format is as follows:

- LLC—0xAAAA03
- Org ID—0x00000C
- HDLC protocol type—0x2003

VTP sends packets on LANs using the multicast address 0100.0CCC.CCCC.

Because VTP does not run on top of any network layer, but runs only over the data link layer, a switch can learn from an advertisement even if it does not have a Layer 3 address on that VLAN.

For more information about the format of VTP frames, see the "Frame Formats" appendix.

VLAN Status Request and Response Messages

A VLAN Status Request message requests all devices in the management domain to respond if their response will determine some aspect of global status information about a particular VLAN. A device with such information generates a VLAN Status Response message and transmits it directly to the originator of the VLAN Status Request message. Both VLAN Status Requests and VLAN Status Responses are sent on the default VLAN (VLAN1, which is the default Ethernet VLAN). VLAN Status Requests are sent to a multicast address that is a different address than the one to which advertisements are sent, so that they are forwarded via normal bridging procedures.

A VLAN Status Request with the Ports-Assigned code requests a response from any device in the management domain that has at least one port assigned to a particular VLAN. Any device having at least one port assigned to the indicated VLAN generates a VLAN Status Response with the Ports-Assigned code.

VTP Pruning

VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases the available bandwidth by restricting flooded traffic to only those ISL trunk links that the traffic must use to access the appropriate network devices. By default, VTP pruning is disabled on the Catalyst 3900 series switch.

Although VTP pruning can be enabled on a Catalyst 3900 switch that is in VTP Server or Transparent mode, only switches that are in VTP Server or Client mode can participate in VTP pruning. VTP Clients, while they can participate in VTP pruning, cannot alter the pruning mode for the management domain.

Note: Make sure that all devices in the management domain support VTP pruning before you enable it.

Figure 4-6 shows a switched network without VTP pruning enabled. Port 1 on Switch 1 and port 2 on Switch 4 are assigned to the VLAN 200. A broadcast is sent from the host connected to Switch 1. Switch 1 floods the broadcast and every switch in the network receives it, even though Switches 3, 5, and 6 have no ports in the VLAN 200.

Figure 4-6 Flooding Traffic without VTP Pruning

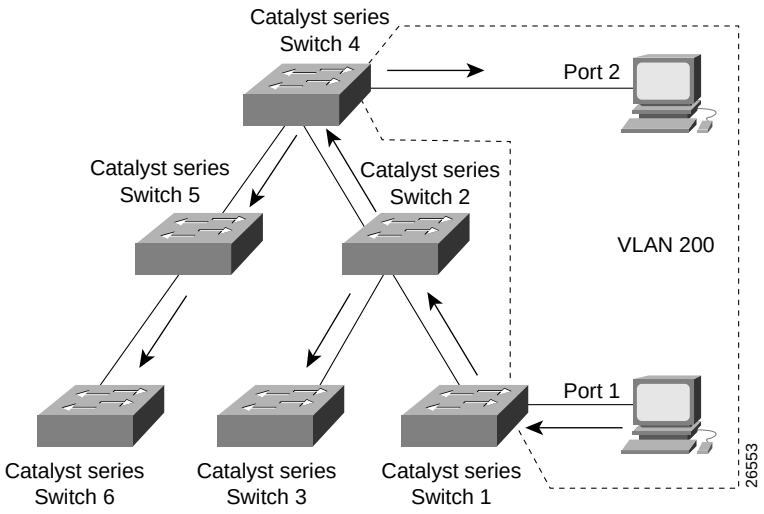

Figure 4-7 shows the same switched network with VTP pruning enabled. The broadcast traffic from Switch 1 is not forwarded to Switches 3, 5, and 6 because traffic for the VLAN 200 has been pruned on the links indicated (port 5 on Switch 2 and port 4 on Switch 4).

Figure 4-7 Flooding Traffic with VTP Pruning

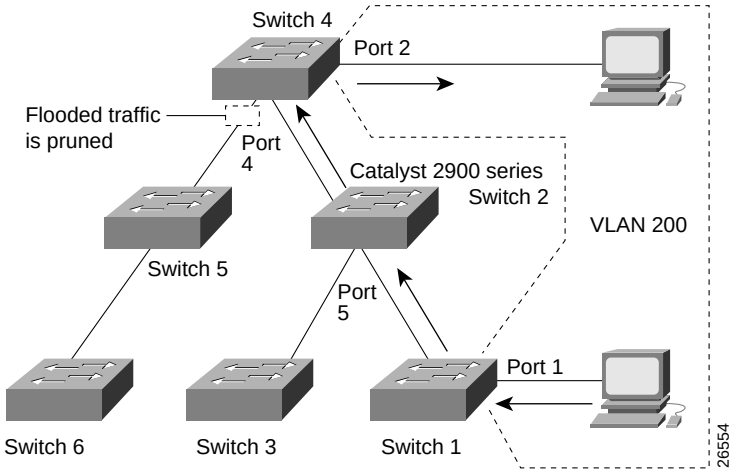

Enabling VTP pruning on a VTP server enables pruning for the entire management domain. VTP pruning takes effect several seconds after you enable it. By default, VLANs 2 through 1000 are pruning eligible. VTP pruning does not prune traffic from VLANs that are pruning ineligible. VLAN 1, the default TrBRF (1005), and the default TrCRF (1003) cannot be configured to be pruning eligible, therefore, the traffic from these VLANs cannot be pruned. Pruning eligibility is configured on a TrBRF basis. Therefore, if you configure a TrBRF other than the default TrBRF to be pruning eligible, all TrCRFs associated with the TrBRF are pruning eligible as well.

Duplicate Ring Protocol

The Cisco Duplicate Ring Protocol (DRiP) runs on Cisco devices that support switched VLAN networking and is used to identify active VLANs and to help prevent the configuration of duplicate rings (TrCRFs) across switches. Through packet advertisements, DRiP maintains the status of TrCRFs. It then uses this information to determine whether there are multiple TrCRFs active in a TrBRF.

DRiP information is used in the following situations:

- ARE filtering.

  To enable the switch to filter out excessive ARE frames, the switch must be aware of the TrCRFs that are attached to the distributed bridge (TrBRF). The DRiP information is used in conjunction with the local configuration to determine which of the TrCRFs configured within a TrBRF have active ports. This information is used on the base switch to correctly filter AREs and on the ISL module to discard AREs that have already been on an attached ring.
- Detecting the configuration of duplicate TrCRFs across switches, which would cause a TrCRF to be distributed across ISL trunks.

  The DRiP information is used in conjunction with the local configuration information to determine which TrCRFs are already active on the switch. If DRiP determines that a TrCRF is configured on more than one switch, it will disable the ports associated with the TrCRF.
- Detecting the failure of an ISL path and enabling a backup path.

  DRiP allows users to configure a unique type of TrCRF that can span multiple switches (the backup TrCRF discussed in the “TrCRFs” section). DRiP monitors the ISL link and when it detects any failure, it activates the necessary ports in the backup TrCRF.

Note: As VLAN IDs are unique throughout the network, DRiP does not need to understand parent–child relationship of a TrCRF to a TrBRF.

How DRiP Works

DRiP sends advertisements to a multicast address so the advertisements are received by all neighboring devices (but they are not forwarded by normal bridging procedures). The advertisement includes VLAN information for the source device only.

When a switch receives a DRiP advertisement, it compares the information in the advertisement with its local configuration to determine which TrCRFs have active ports and then denies any configuration that would allow a TrCRF that is already active on another box to be configured on the local switch.

In the event a TrCRF is believed to not be in use within the TrBRF and ports on two or more separate devices are simultaneously configured for the same TrCRF, each switch will send a DRiP advertisement and the advertisement with the lowest sender's MAC address will be accepted. The switch with the higher MAC address will disable the port that was just configured and will send a DuplicateRingNumberError trap.

If a trunk connection is lost, an aging process ages out all entries associated with that trunk port.

DRiP Advertisements

A DRiP advertisement is sent at periodic intervals (30 seconds). If no change in status or configuration has taken place, then the configuration revision number is not updated. Instead the periodic message will indicate that nothing has changed because the revision number has not changed.

A switch also generates a DRiP advertisement when one of the following situations occur:

- Trunking comes up (for ISL trunks). The DRiP advertisement is sent on all ISL trunks.
- An ISL trunk port is configured for a new TrBRF. The DRiP advertisement is sent on the ISL trunk port on which the TrBRF was configured.
- A TrBRF is created or deleted. The switch updates its configuration revision number and the DRiP advertisement is sent on all ISL trunks.
- A port local to the switch is configured for a TrCRF that is not associated with any other port on the switch. The switch updates its configuration revision number and the DRiP advertisement is sent on all ISL trunk ports.
- A port local to the switch that is the last or only port active on a TrCRF is removed. The device updates its configuration revision number and the DRiP advertisement is sent on all ISL trunk ports.
- The switch receives a DRiP advertisement with a revision number less than its own and the advertisement contains conflicting information about a TrCRF that is in use on that trunk port. The device does not update its revision number. It generates its own advertisement and sends it on the ISL trunk from which the original DRiP advertisement was received.
- The switch receives a DRiP advertisement with a revision number greater than its own and the advertisement contains conflicting information about a TrCRF that is in use on that trunk port. The device updates its configuration revision number and forwards the advertisement on all ISL trunk ports except the one from which the original DRiP advertisement was received.

Transmission of Advertisements

DRiP advertisements are transmitted using the same multicast destination MAC address (0100.0CCC.CCCC) used for VTP and are not forwarded using normal bridging techniques. Like VTP, a switch regenerates DRiP advertisement to all of its other trunk ports if the advertisement contains new configuration information.

Advertisements are transmitted on the default VLAN (VLAN1), which corresponds to the type of trunk link. Thus, only one copy is transmitted on a trunk port, no matter how many VLANs are defined.

DRiP Frame Format

DRiP is assigned the Cisco HDLC protocol type value 0x0102. A Cisco-proprietary SNAP value enumerates HDLC protocol type values so DRiP can run on all media that support SNAP, such as LAN media, Frame Relay, and ATM.

The SNAP format is as follows:

- LLC—0xAAAA03
- Org ID—0x00000C
- HDLC protocol type—0x0102

DRiP sends packets on LANs using the multicast address 0100.0CCC.CCCC.

For more information about the format of DRiP frames, see the "Frame Formats" appendix.

Spanning-Tree Protocol

The STP is a broadcast algorithm used by network bridge connections to dynamically discover a loop-free subset of the network topology while maintaining a path between every pair of LANs or VLANs in the network.

To accomplish this, the STP blocks ports that, if active, would create bridging loops. If the primary link fails, it activates one of the blocked bridge ports to provide a new path through the network.

In a traditional bridged network, there is one STP for each bridge connection. Each bridge maintains its own database of configuration information and transmits and receives only on those ports belonging to the bridge. The type of STP that runs on a bridge depends on the transmission mode of the bridge connection (whether the connection is SRB, source-route switched, or SRT).

As discussed in the "Token Ring VLANs" section, in a switched network, you can configure virtual networks. A switch can have ports that belong to different VLANs, some of which may span several switches. To prevent loops in the bridged connections between the VLANs, you should configure the STP.

In a Token Ring switch, there are two levels of VLANs. Therefore, in a Token Ring switched network, to ensure loops are removed from the topology you must configure a separate STP for the logical bridge (TrBRF) and for the port groups (TrCRF) configured for a VLAN.

The STP that is run at the TrCRF removes the loops in the TrCRF logical ring. The STP that is run at the TrBRF removes the loops in the bridging topology.

How the STP Algorithm Works

In general, the STP eliminates loops in the network as follows:

1. Each bridge is assigned an eight-byte unique bridge identifier.
   The first two bytes are a priority field, and the last six bytes contain one of the bridge's MAC addresses. The bridge with the lowest bridge identifier among all bridges on all LAN segments is the root bridge. The network administrator can assign a lower bridge priority to a selected bridge to control which bridge becomes the root, or the administrator can use default bridge priorities and allow the STP to determine the root.
2. Each bridge port is associated with a path cost.
   The path cost represents the cost of transmitting a frame to a bridged segment through that port. A network administrator typically configures a cost for each port based on the speed of the link (for example, the cost of a port connected to a 16-Mbps LAN could be assigned a lower path cost than a port connected to a 4-Mbps LAN).
3. Each bridge determines its root port and root path cost.
   The root port is the port that represents the shortest path from itself to the root bridge. The root path cost is the total cost to the root. All ports on the root bridge have a zero cost.
4. All participating bridges elect a designated bridge from among the bridges on that LAN segment.
   A designated bridge is the bridge on each LAN segment that provides the minimum root path cost. Only the designated bridge is allowed to forward frames to and from that LAN segment toward the root.
5. All participating bridges select ports for inclusion in the spanning tree.
   The selected ports will be the root port plus the designated ports for the designated bridge. Designated ports are those where the designated bridge has the best path to reach the root. In cases where two or more bridges have the same root path cost, the bridge with the lowest bridge identifier becomes the designated bridge.
6. Using the preceding steps, all but one of the bridges directly connected to each LAN segment are eliminated, thereby removing all multiple LAN loops.

How Spanning-Tree Information is Shared

The STP calculation requires that bridges communicate with other bridges in the network that are running the STP. Each bridge is responsible for sending and receiving configuration messages called bridge protocol data units (BPDUs).

BPDUs are exchanged between neighboring bridges at regular intervals (typically 1 to 4 seconds) and contain configuration information that identifies the:

- Bridge that is presumed to be the main bridge or root (root identifier)
- Distance from the sending bridge to the root bridge (called the root path cost)
- Bridge and port identifier of the sending bridge
- Age of the information contained in the configuration message

If a bridge fails and stops sending BPDUs, the bridges detect the lack of configuration messages and initiate a spanning-tree recalculation.

Note: By default, the functional address used for IEEE STP frames sent by Cisco routers is the same address as the functional address used by the IBM bridges. The Catalyst 3900 and Catalyst 5000 Token Ring module allow you to specify which functional address you want to use.

STPs for Token Ring Switches

The Catalyst Token Ring Switches support the following STPs:

- IEEE 802.1d
- IBM
- Cisco

IEEE 802.1d STP

The IEEE STP can be used at the TrCRF or the TrBRF level. This type of spanning tree supports bridge domains and allows the bridge to construct a loop-free topology across an extended LAN. The IEEE 802.1d STP supports the following bridge modes:

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

The IEEE STP frames use a multicast destination address of x'800143000000.

IBM STP

The IBM STP can be used at the TrBRF level. This type of spanning tree was developed to manage the single-route broadcast path through source-route bridges.

The IBM STP frames use a multicast destination address of x'C00000000100.

Cisco STP

The Cisco STP is designed to be used at the TrCRF level. This type of spanning tree was developed to address a looping problem that can be introduced when you have a ring that spans multiple ports in a Token Ring environment.

One of the rules in processing source-route traffic is that a source-route frame should never be forwarded to a ring that it has previously traversed. If the RIF of a source-route frame already contains the ring number for the next hop, the bridge assumes that the frame has already been on that ring and drops the frame.

With TrCRFs, however, this rule can cause a problem. With the existing STP, a frame that originated on one physical ring of a TrCRF and is processed by an external SRT bridge would not be forwarded to another physical ring of the same TrCRF. Therefore, the IEEE 802.1d STP was used as a basis to create the Cisco STP. The Cisco STP ensures that traffic from one physical ring of a TrCRF is not blocked from the other physical rings that comprise the TrCRF.

The Cisco BPDUs are source-routed frames with two bytes of RIF information. This ensures that BPDUs will not be source routed or transparently routed to other LANs. The Cisco BPDUs use a different multicast destination address (x'800778020200) to ensure that external bridges do not interpret the frames as IEEE or IBM STP frames.

When to Use a Specific STP

Although the Catalyst 3900 switch and Catalyst 5000 Token Ring module support the three STPs, the use of the protocols is implemented slightly different on each switch.

Spanning-Tree Protocol and the Catalyst 3900 Switch

For the Catalyst 3900, you can set the STP for the both TrBRF and the TrCRF.

Possible values for the STP at the TrBRF are no, IBM, IEEE, and Base on Bridging. If you select Base on Bridging (the default), the STP used is determined by the bridge mode. If the bridge mode is SRB, the IBM STP is used. If the bridge mode is SRT, the IEEE STP is used.

Possible values for the STP at the TrCRF are no, IEEE, Cisco, and Base on Bridging Mode. If you select Base on Bridging Mode (the default), the STP used at the TrCRF is determined by the bridging mode.

Recommendations for running the STP are as follows:
- For SRB, run the IBM STP at the TrBRF and the IEEE STP at the TrCRF.
- For SRT, run the IEEE STP at the TrBRF and the Cisco STP at the TrCRF.

Spanning-Tree Protocol and the Catalyst 5000 Token Ring Module

For the Catalyst 5000 series Token Ring module, you can set the STP for the TrBRF only. Possible values are IBM and IEEE.

The STP used at the TrCRF is determined by the bridge mode.
- If the bridging mode for the TrCRF is SRB, the IEEE STP is used at the TrCRF.
- If the bridging mode for the TrCRF is SRT, the Cisco STP is used at the TrCRF.

The ISL module supports the STP at both the TrCRF and the TrBRF level. The STP that is run on the ISL link depends on the type of TrCRF:
- With an undistributed TrCRF, the STP specified for the TrBRF is used.
- With a default TrCRF, the STP specified for the TrCRF is used.

Also, there are some combinations of STP and bridge mode that the Catalyst 5000 series Token Ring module considers incompatible. These combinations are as follows:
- TrBRF STP of IBM and TrCRF bridge mode of SRT
- TrBRF STP of IEEE and TrCRF bridge mode of SRB

If you configure one of these combinations, no STP will be run at the TrBRF level and the logical ports will be placed in a blocked state. You can override the logical port state using the **set spantree portstate** command.

Table 4-1 shows a summary of the STPs used on the Catalyst 5000 Token Ring module.

Table 4-1 Summary of STPs Used on the Catalyst 5000 Token Ring Module

STP Specified for TrBRF	Port Type or Bridge Mode Specified for TrCRF	STP Used at TrBRF	STP Used at TrCRF
IBM	SRB	IBM	IEEE
	SRT	None	Cisco
IEEE	SRB	None	IEEE
	SRT	IEEE	Cisco

ATM and Token Ring LANE

The Catalyst 3900 and the Catalyst 5000 have ATM expansion modules that provide high-speed connectivity between the switch and an ATM backbone network.

This chapter provides the following information:

- Understanding ATM
- Understanding LAN Emulation
- Recommended Environments

Understanding ATM

ATM is a cell-switching and multiplexing technology that combines the benefits of circuit switching (constant transmission delay and guaranteed capacity) with those of packet switching (flexibility and efficiency for intermittent traffic). Like X.25 and Frame Relay, ATM defines the interface between the user equipment (such was workstations and routers) and the network (referred to as the User-Network Interface [UNI]). This definition supports the use of ATM switches (and ATM switching techniques) within both public and private networks.

Because it is an asynchronous mechanism, ATM differs from synchronous transfer mode methods, where time-division multiplexing (TDM) techniques are employed to preassign users to time slots. ATM time slots are made available on demand, with information identifying the source of the transmission contained in the header of each ATM cell. TDM is inefficient relative to ATM because if a station has nothing to transmit when its time slot comes up, that time slot is wasted. The reverse situation, where one station has large amounts of information to transmit, is also less efficient. In this case, that station can only transmit when its turn comes up, even though all the other time slots are empty. With ATM, a station can send cells whenever necessary.

Another critical ATM design characteristic is its star topology. The ATM switch acts as a hub in the ATM network, with all devices attached directly. This provides all the traditional benefits of star-topology networks, including easier troubleshooting and support for network configuration changes and additions.

Furthermore, ATM's switching fabric provides additive bandwidth. As long as the switch can handle the aggregate cell transfer rate, additional connections to the switch can be made. The total bandwidth of the system increases accordingly. If a switch can pass cells among all its interfaces at the full rate of all interfaces, it is described as nonblocking. For example, an ATM switch with 16 ports set at 155 Mbps would require about 2.5 Gbps aggregate throughput to be nonblocking.

ATM switches transmit data in small units called *cells*. The latency in a cell switch is very small because of the short cell size. Short cells have a tiny store-and-forward delay. In the absence of port contention and buffering, cells are switched quickly in hardware. For information about the format of an ATM cell, see the "Frame Formats"appendix.

In addition to the low latency, ATM is beneficial to large networks because it:

- Is a multiplexing and switching technology designed for flexibility and performance.
- Supports Quality of Service (QOS) options for flexibility and high bandwidth options (up to gigabits per second) for performance.
- Offers both permanent virtual circuits (PVCs) that are set up for static connections and switched virtual circuits (SVCs) that are automatically set up and torn down when data needs to be transferred.
- Supports environments where applications with different performance requirements need to be executed on the same computer, multiplexer, router, switch and network. The flexibility of ATM means that voice, video, data, and future payloads can be transported.
- Has worldwide support. The ATM Forum, an industry forum made up of many companies (including Cisco), works with formal standards bodies to specify ATM.

PVC

A PVC is a non-switched connection that is established beforehand (manually pre-provisioned) to satisfy a standing need for network services. It is a logical (not a physical) connection between two communicating ATM peers. This type of connection is typically established by a network administrator.

User applications that require an on-going, specific level of transmission bandwidth typically use PVCs for interconnectivity. The network bandwidth required in this type of application tends to be more predictable and constant, enabling the physical transmission medium to be tailored to an expected volume of traffic, and vice versa.

With a PVC, everything is statically configured and no signaling is involved. The PVC is mapped to a network in a subinterface point-to-point configuration. The logical data link layer can use SNAP encapsulation (as defined in RFC 1483). This allows multiple protocols to be multiplexed over one PVC. Alternatively, the logical data link layer can use LANE Version 1 over PVC.

The PVC is statically mapped at each ATM node. The path of the PVC is identified at each switch by an incoming virtual channel identifier/virtual path identifier (VCI/VPI) and an outgoing VCI/VPI.

Note: The Catalyst 3900 ATM expansion module does not support PVC configuration.

SVC

An SVC is a switched connection that is established by a defined and standardized ATM signaling protocol. This type of connection is set up dynamically (on demand) across the network, as required by the user's communications applications. An SVC is established and torn down using a flexible connection setup protocol that supports various connection types.

The transfer of information between network users by means of SVCs typically occurs through shared network facilities, rather than through dedicated transmission lines or owned physical facilities.

Establishing an ATM SVC involves an agreement between the end nodes and all the switches in between. Each end node has a special signaling channel to the connected switch called the UNI. Switches have a signaling channel between them called the Network-to-Network Interface (NNI). Cells that arrive on the signaling channel are reassembled into frames in the reliable Service Specific Connection Oriented Protocol (SSCOP). The signaling information follows the Q.2931 standard.

Establishing an SVC potentially involves signaling between the following:

- Router and a private ATM switch (private UNI)
- Router and a public ATM switch (public UNI)
- Private ATM switch and a public ATM switch (public UNI)

The UNI is defined by the ATM Forum UNI specification.

Interfaces to public ATM networks are identified by an E.164 address. Interfaces to private ATM networks are identified by a network service access point (NSAP) address. These addresses are contained in different fields of the same 20-octet address.

Once an SVC is established, it functions like a PVC. SVCs can be used in point-to-point subinterface configuration or point-to-multipoint nonbroadcast multiaccess (NBMA) configuration.

ATM Adaptation Layers

The purpose of the ATM adaptation layer (AAL) is to receive the data from the various sources or applications and convert or adapt it to 48-byte segments that will fit into the payload of an ATM cell. The AAL segments upper-layer user information into ATM cells at the transmitting end of a virtual connection and reassembles the cells into a user-compatible format at the receiving end of the connection. These complimentary functions occur between communicating peers in the network at the same level of the ATM architectural model.

The AAL is not a network process. Rather, AAL functions are performed by the user's network terminating equipment on the user side of the UNI. Consequently, the AAL frees the network from concerns about different traffic types.

How AAL processes are carried out depends on the type of traffic to be transmitted. Different types of AALs handle different types of traffic, but all traffic is ultimately packaged by the AAL into 48-byte segments for placement into ATM cell payloads. Consequently, several different AALs have been defined for different types of services.

Table 5-1 lists these AALs.

Table 5-1 ATM Adaption Layers

Traffic Class	Timing Relationship	Connection Mode	Bit Rate	Traffic Description
Class A (AAL1)	Synchronous	Connection- oriented	Constant	This type of traffic typically consists of constant bit rate (CBR) analog signals. Hence, synchronous timing relationships exist between the senders and receivers of this traffic. This type of traffic over an ATM network is often referred to as circuit emulation service, an example of which is fixed bit rate, uncompressed voice, or video data.
Class B (AAL2)	Synchronous	Connection- oriented	Variable	As with Class A traffic, synchronous timing relationships exist between the senders and receivers of Class B traffic. However, Class B relates to variable bit rate (VBR) traffic, typical examples of which are compressed voice and video traffic. Such traffic is typically "bursty" in nature.

Table 5-1 ATM Adaption Layers (Continued)

Traffic Class	Timing Relationship	Connection Mode	Bit Rate	Traffic Description
Class C (AAL3/4)	Asynchronous	Connection- oriented	Variable	No timing relationships exist between the senders and receivers of data. Hence, such traffic is asynchronous. Class C handles VBR connection-oriented traffic. This class provides point-to-point or point-to-multipoint ATM cell relay services over connections established "on the fly" through signaling and routing messages exchanged between data senders and receivers. This service handles multiple traffic types (data, voice, and video) in which user data is arranged into ATM cells for efficient transport through the network. Class C of traffic contains sequencing bits that allows the cells to take different paths and still be reassembled in the correct order at the receiving station. This type of traffic is sensitive to cell loss, but not to cell transport delay (or latency). Latency is the delay between the time a device receives a cell on its input port and the time the cell is forwarded through its output port.
Class D (AAL5)	Asynchronous	Connectionless	Variable	Class D handles unspecified bit rate (UBR) traffic in a connectionless, asynchronous manner.

Because ATM is inherently a connection-oriented transport mechanism and because the current applications of ATM are heavily oriented toward LAN traffic, many of the current ATM products, including the Catalyst 3900 and the Catalyst 5000, support the Class D adaptation layer with AAL5.

AAL Sublayers

The AAL performs two main functions in service-specific sublayers of the AAL:

- A convergence function in the convergence sublayer (CS)
  The CS provides appropriate traffic services to higher-layer protocols. Once a connection is established between communicating ATM entities with an appropriate QOS, the CS accepts higher-layer traffic for transmission through the network. Depending on the traffic type, certain header or trailer fields are added to the user data payload and formed into information packets called CS protocol data units (CS-PDUs).
- A cell segmentation and reassembly function in the segmentation and reassembly (SAR) sublayer.
  The SAR segments each CS-PDU received from the CS into smaller units and adds a header or trailer field, depending on the traffic type, to form 48-byte payloads called SAR sublayer protocol data units (SAR-PDUs). Once the user data is arranged into SAR-PDUs by the AAL layer, they are passed to the ATM layer, which packages the data into 53-byte ATM cells, making them suitable for transport as outgoing ATM cells by the physical layer.
  Upon receipt of incoming ATM cells from the physical layer (that is, cells delivered from a peer physical layer elsewhere in the network), the AAL removes any AAL-specific information from each cell payload and reassembles the packet for presentation to higher layer protocols in a form expected by the user application.

AAL5 Traffic Processing

AAL5 has been designed to process traffic typical of today's LANs. Originally, AAL3/4 was designed to process this kind of traffic. However, the inefficiency of AAL3/4 for handling LAN traffic led to the use of AAL5 for such traffic.

AAL5 provides a streamlined data transport service that functions with less overhead than AAL3/4. AAL5 is typically associated with UBR traffic.

Another AAL5 attribute contributing to its efficiency is that is uses only 5 bytes of header. None of the payload is used for header information. Also, AAL5 calculates a 32-bit cyclic redundancy check (CRC) over the entire AAL5 protocol data unit to detect cell loss and the incorrect ordering or incorrect insertion of cells.

For purposes of AAL5 traffic processing, the CS is divided into the following parts:

- Common part convergence sublayer (CPCS)—Provides the capability to transfer the CPCS protocol data units (CPCS-PDU payloads) from one AAL5 user to another AAL5 peer in the network. The AAL5 traffic type supports both a message-mode service and a streaming mode service.
- Service specific convergence sublayer (SSCS)—Allows different SSCS protocols to be defined to support specific AAL user services or groups of services. The SSCS may also be null because it provides only for the mapping of equivalent AAL primitives to the CPCS, and vice versa.

Components of an ATM Network

The building blocks of an ATM internetwork may consist of the following:

- Routers with ATM interfaces
- Computers with a native ATM NIC
- LightStream 1010 or other ATM switches
- ATM physical layer, supporting Synchronous Optical Network (SONET) OC-3 with single or multimode fiber, Transparent Asynchronous Transmitter/Receiver Interface (TAXI) with multimode fiber, or DS3/E3 with coaxial cable
- LAN switches with ATM interfaces

ATM and VLANs

The Catalyst 3900 ATM expansion module supports up to 63 VLANs (or ELANs). Each ELAN corresponds to a TrCRF. Each association between the ATM expansion module and a TrCRF creates a virtual ATM port. A virtual ATM port is the equivalent of an LAN Emulation Client (LEC).

Understanding LAN Emulation

LANs can use connectionless service. However, ATM is always a connection-oriented service. Devices first use a signaling process to establish a path with an ATM destination. Devices can send cell-based traffic only after the devices have identifiers pointing to the connection path.

LANE uses point-to-multipoint connections to service the connectionless broadcast service that is required by LAN protocols.

Cisco's Token Ring implementation of LANE makes an ATM interface look like one or more Token Ring interfaces. Setting up LECs allows the Catalyst 3900 or Catalyst 5000 Token Ring module to operate in an ATM LAN environment containing Cisco 7000 or Cisco 4500 series routers with ATM Interface Processor (AIP) connected to a LightStream 1010 ATM switch.

Figure 5-1 illustrates the physical layout of an ATM network that uses LANE.

Figure 5-1 Physical View of LANE

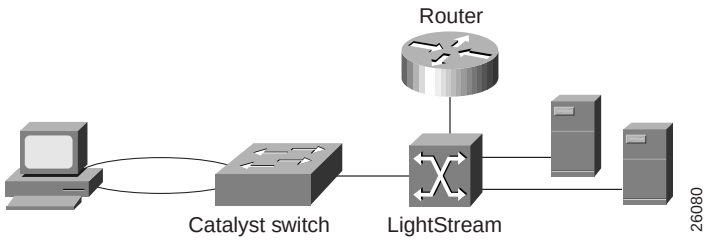

Figure 5-2 illustrates the logical view of the LANE network.

Figure 5-2 Logical View of LANE

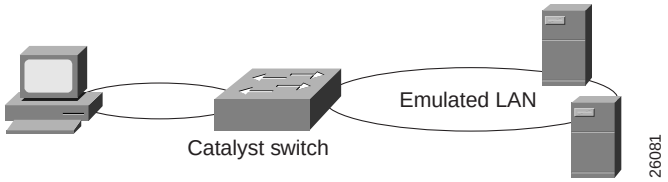

LANE is an ATM service defined by the ATM Forum specification *LAN Emulation over ATM (*ATM_FORUM 94-0035). This service emulates the following LAN-specific characteristics:

- Connectionless services
- Multicast services
- LAN MAC driver services

LANE service provides connectivity between ATM-attached devices and LAN-attached devices. This includes connectivity between ATM-attached stations and LAN-attached stations as well as connectivity between LAN-attached stations across an ATM network.

Because LANE connectivity is defined at the MAC layer, upper protocol-layer functions of LAN applications can continue unchanged when the devices join ELANs. This feature protects corporate investments in legacy LAN applications.

An ATM network can support multiple independent ELANs. Membership of an end system in any of the ELANs is independent of the physical location of the end system. The end systems can move easily from one ELAN to another, regardless of whether or not the hardware is moved.

Components of LANE

A Catalyst 3900 or Catalyst 5000 ATM module can participate in up to 63 of these ELANs.

LANE is defined on a client-server LAN model, as follows:

- LEC

  An LEC emulates a LAN interface to higher layer protocols and applications. It forwards data to other LANE components and performs LANE address resolution functions.

  Each LEC is a member of only one ELAN. However, a router or a Catalyst ATM module can include LECs for multiple ELANs: one LEC for each ELAN of which it is a member.

  If a router has clients for multiple ELANs, the router can route traffic between the ELANs.

Note: If the Catalyst 3900 has multiple ATM modules and each has a client that is active for the same ELAN, the Catalyst 3900 will not bridge between the ELANs on the different modules. The Catalyst 3900 acts as an edge device on an ATM cloud (that is, there are no LANE services in the Catalyst 3900).

- LANE Server (LES)

  The LANE server for an ELAN is the control center. It provides joining, address resolution, and address registration services to the LANE clients in that ELAN. Clients can register destination unicast and multicast MAC addresses with the LANE server. The LANE server also handles LANE ARP (LE_ARP) requests and responses.

  The current configuration is one LES per ELAN.

- LANE Broadcast and Unknown Server (BUS)
  The LANE BUS sequences and distributes multicast and broadcast packets and handles unicast flooding. One combined LES and BUS is required per ELAN.
- LANE Configuration Server (LECS)
  The LECS contains the database that determines which ELAN a device belongs to (each configuration server can have a different named database). Each LEC contacts the LECS once, when it joins an ELAN, to determine which ELAN it should join. The LECS returns the ATM address of the LES for that ELAN.
  One LECS is required per ATM LANE switch cloud.

  The LECS database can have the following four types of entries:
  – ELAN name, ATM address of LANE server pairs
  – LANE client MAC address, ELAN name pairs
  – LANE client ATM template, ELAN name pairs
  – Default ELAN name

Note: ELAN names must be unique on an interface. If two interfaces participate in LANE, the second interface may be in a different switch cloud.

- Simple Server Redundancy Protocol (SSRP)
  The LANE simple server redundancy feature creates fault tolerance using standard LANE protocols and mechanisms. If a failure occurs on the LANE configuration server or on the LES/BUS, the ELAN can continue to operate using the services of a backup LANE server.

The Catalyst 3900 ATM module currently supports only the LEC function. A Catalyst 5000 or a Cisco 7000, Cisco 7200, Cisco 7500, RSP 7000, Cisco 4500, or Cisco 4700 with an AIP can supply all LANE functions.

LANE Operation and Communication

Communication among LANE components is typically handled by several types of VCCs. Some VCCs are unidirectional; others are bidirectional. Some are point-to-point and others are point-to-multipoint. Figure 5-3 illustrates the various types of VCCs.

Figure 5-3 LANE VCC Types

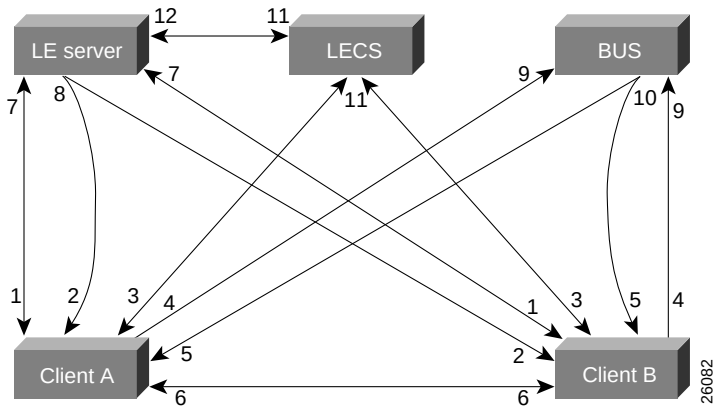

1–7	Control direct	4–9	Multicast send
2–8	Control distribute	5–10	Multicast forward
3–11	Configure direct (client)	6–6	Data direct
		11–12	Configure direct (server)

The following section describes the processes involved with a client requesting to join an ELAN.

Join Process

The following process (illustrated in Figure 5-3) normally occurs after an LEC has been enabled on the ATM module:

Step 1. The client requests to join an ELAN. The client sets up a connection to the LECS to find the ATM address of the LANE server for its ELAN. See the bidirectional, point-to-point link (link 1–7 in Figure 5-3).

An LEC finds the LECS using the following methods in the listed order:

- Locally configured ATM address
- Interim Local Management Interface (ILMI)
- Fixed address defined by the ATM Forum

Step 2. The LECS identifies the LES. Using the same VCC, the LECS returns the ATM address and the name of the LES for the client's ELAN.

Step 3. The client tears down the configure direct VCC.

Step 4. The client contacts the server for its LAN. The client sets up a connection to the LES for its ELAN (bidirectional, point-to-point control direct VCC [link 1–7 in Figure 5-3]) to exchange control traffic. Once a control direct VCC is established between an LEC and LES, it remains up.

Step 5. The LES verifies that the client is allowed to join the ELAN. The server for the ELAN sets up a connection to the LECS to verify that the client is allowed to join the ELAN (bidirectional, point-to-point server configure VCC [link 11–12 in Figure 5-3]).

The server's configuration request contains the client's MAC address, its ATM address, and the name of the ELAN. The LECS checks its database to determine whether the client can join that LAN; then it uses the same VCC to inform the server whether or not the client is allowed to join.

Step 6. The LES allows or disallows the client to join the ELAN. If allowed, the LES adds the LEC to the unidirectional, point-to-multipoint control distribute VCC (link 2–8 in Figure 5-3) and confirms the join over the bidirectional, point-to-point control direct VCC (link 1–7 in Figure 5-3). If disallowed, the LES rejects the join over the bidirectional, point-to-point control direct VCC (link 1–7 in Figure 5-3).

Step 7. The LEC sends LE_ARP packets for the broadcast address, which is all 1s. Sending LE_ARP packets for the broadcast address returns the ATM address of the BUS. Then the client sets up the multicast send VCC (link 4–9 in Figure 5-3) and the BUS adds the client to the multicast forward VCC (link 5–10 in Figure 5-3) to and from the BUS.

Step 8. The LEC registers the ring numbers of all other TrCRFs within its TrBRF that contain active ports on the local switch.

Addressing

On a LAN, packets are addressed by the MAC-layer address of the destination and the source stations. To provide similar functionality for LANE, MAC-layer addressing must be supported. Every LEC must have a MAC address. In addition, every LANE component (server, client, BUS, and configuration server) must have a unique ATM address.

All LECs on the same interface have a different, automatically assigned MAC address. That MAC address is also used as the end-system identifier part of the ATM address as explained in the following sections.

LANE ATM Addresses

A LANE ATM address has the same syntax as an NSAP, but it is not a network-level address. It consists of the following:

- A 13-byte prefix that includes the following fields defined by the ATM Forum: authority and format identifier (AFI) field (1 byte); Data Country Code (DCC) or International Code Designator (ICD) field (2 bytes); Domain Specific Part Format Identifier (DFI) field (1 byte); Administrative Authority (3 bytes); Reserved (2 bytes); Routing Domain (2 bytes); and Area (2 bytes).
- A 6-byte end-system identifier.
- A 1-byte selector field.

ILMI Address Registration

The Catalyst 3900 and Catalyst 5000 ATM modules use ILMI registration to build their ATM addresses and to register the addresses with the ATM switch. To build its ATM address, each module obtains its ATM address prefix from the ATM switch. Then it combines the ATM address prefix with its own MAC address and the selector value of 0 (zero). Once the ATM module has determined its ATM address, it uses ILMI to register this address with the ATM switch.

Address Resolution

As communication occurs on the ELAN, each client dynamically builds a local LE_ARP table. The LE_ARP table maps ELAN MAC addresses (Layer 2) to ATM addresses (also Layer 2). A client's LE_ARP table can also have static, preconfigured entries. The LE_ARP table maps MAC addresses to ATM addresses.

When a client first joins an ELAN, its LE_ARP table has no dynamic entries and the client has no information about destinations on or beyond its ELAN.

To learn about a destination when a packet is to be sent, the client begins the following process to find the ATM address corresponding to the known MAC address:

Step 1. The client sends an LE_ARP request to the LANE server for this ELAN (point-to-point control direct VCC [link 1–7 in Figure 5-3]).

Step 2. If the MAC address is registered with the server, it returns the corresponding ATM address. If not, the LES forwards the LE_ARP request to all clients on the ELAN (point-to-multipoint control distribute VCC [link 2–8 in Figure 5-3]).

Step 3. Any client that recognizes the MAC address responds with its ATM address (point-to-point control direct VCC [link 1–7 in Figure 5-3]).

Step 4. The LES forwards the response (point-to-multipoint control distribute VCC [link 2–8 in Figure 5-3]).

Step 5. The client adds the MAC address-ATM address pair to its LE_ARP cache.

Step 6. Now the client can establish a VCC to the desired destination and proceed to transmit packets to that ATM address (bidirectional, point-to-point data direct VCC [link 6–6 in Figure 5-3]).

For unknown destinations, the client sends a packet to the BUS, which forwards the packet to all clients. The BUS floods the packet because the destination might be behind a bridge that has not yet learned this particular address.

Traffic Handling

The Catalyst 3900 allows you to define up to 64 traffic profiles. These profiles can be used to define the maximum rates for each traffic type.

For each VLAN (or ELAN), a traffic profile must be mapped to each DD-VCC. The process of mapping depends on whether the traffic is incoming or outgoing:

- For incoming calls, the LEC tries to find a traffic profile that best matches the traffic parameters in the call. You can define the maximum discrepancy between the specified parameters and actual values on a per ELAN basis.
- For outgoing calls, you can define up to 10 profiles to use. The destination ATM address is ANDed with the address mask. The resulting ATM address is compared with the ATM address in the mapping table. If there is a match, each defined profile (0 through 9) is used in sequence until a call SETUP is accepted by the destination node.

Multicast Traffic

When an LEC has broadcast or multicast traffic, or unicast traffic with an unknown address to send, the following process occurs:

Step 1. The client sends the packet to the BUS (unidirectional, point-to-point multicast send VCC [link 4–9 in Figure 5-3]).

Step 2. The BUS forwards (floods) the packet to all clients (unidirectional, point-to-multipoint multicast forward VCC [link 5–10 in Figure 5-3]).

This VCC branches at each ATM switch. The switch forwards these packets to multiple outputs. (The switch does not examine the MAC addresses; it simply forwards all packets it receives.)

Recommended Environments

This section describes some scenarios for using the Catalyst Token Ring switches with the ATM expansion module. Figure 5-4 shows how the ATM expansion module can be used to connect Catalyst Token Ring switches through an ATM switch.

Figure 5-4 Token Ring Connection Over an ATM Backbone

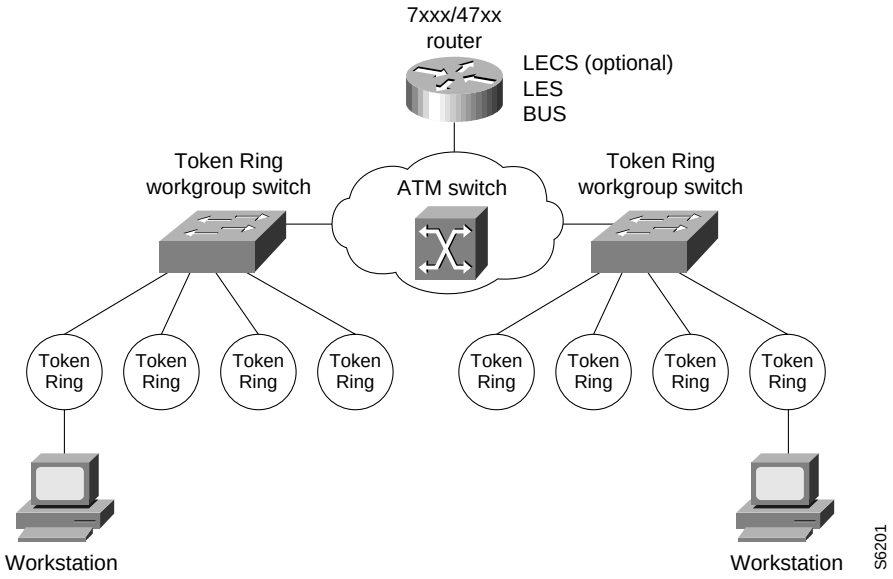

The ATM expansion module is well suited for the following environments:

- ATM backbone for legacy LANs

  A Catalyst Token Ring switch with an ATM expansion module provides a seamless, switched network between legacy LANs communicating over ATM. Therefore, as a first step in migrating from legacy LANs, many users deploy ATM in the backbone or as the WAN technology to connect geographically dispersed legacy LANs.

- LAN-to-ATM interoperability

  A Catalyst Token Ring switch with one or more ATM expansion modules can help protect your legacy LAN investment by providing transparent LAN-to-ATM switching. Therefore, as the next step in a legacy LAN-to-ATM migration, users place high-speed or frequently accessed servers, or both, on the ATM network to take advantage of ATM's scalability. By using LANE over ATM, ATM-based workstations are able to use existing legacy LAN applications with minimal or no upgrade costs for moving to ATM.

Using Multiple ATM Expansion Modules

Because one ATM expansion module can support as many LECs as there are VLANs in a Catalyst Token Ring switch, the question arises: Why use two ATM expansion modules?

The first reason is to provide a backup LEC. By enabling two LECs on two different ATM expansion modules to be members of both the same VLAN and the same ELAN, the Catalyst Token Ring switch's spanning-tree operation will automatically use one LEC for forwarding frames and the other LEC for blocking frames (active standby). For the backup LEC configuration to work, spanning tree *must* be enabled for the related LAN switch domain.

Note: The STP (802.1d or IBM) used depends on whether the bridging mode is configured as SRB or SRT. If the first LEC fails, the second LEC will automatically take over.

The second reason to use two ATM expansion modules is to increase system resources. In some environments the resources associated with a single ATM expansion module might affect individual LEC performance. By moving one or more LECs to a second ATM expansion module, you can significantly increase the resources available to each LEC.

Note: If the Catalyst 3900 has multiple ATM modules and each has clients active for the same ELAN, the Catalyst 3900 will not bridge between the ELANs. The Catalyst 3900 acts as an edge device on an ATM cloud.

Usage Recommendations and Restrictions

Before installing and configuring the ATM module, be aware of the following:

- Because the ATM module cannot internally bridge between TrCRFs, you can configure only one LEC per TrBRF per ATM module. Therefore, while a TrBRF can include only one LEC from a single ATM expansion module, it can contain multiple LECs when each LEC exists on a separate ATM module.
- Configuring a TrBRF to contain an ATM port restricts the MTU of that TrBRF to 4472 or less. The ATM module does not support MTUs greater than 4472. Configuring an MTU size larger than 4472 for a TrBRF that contains an ATM port will result in the ATM port being removed from the TrBRF. Additionally, if a TrBRF is configured for an MTU size larger than 4472, none of the TrCRFs assigned to the TrBRF can be assigned to an ATM port.

Network Management

You can gather status information and statistical data using the console panels of the Catalyst 3900 or the command line interface of the Catalyst 5000. In addition, the Catalyst Token Ring switches support other forms of network management, including:

- Device Management
- Topology Management
- VLAN Management
- Traffic Management
- Soft Error Management

This chapter provides an overview of the network management support provided with the Catalyst Token Ring switches.

Device Management

Use the following two methods to manage the Catalyst Token Ring switches:

- SNMP-based managers
- RMON

SNMP Management

Most Token Ring switches, including the Catalyst Token Ring switches, can be managed using SNMP. There are several SNMP MIB definitions for Token Ring information and switches. The Catalyst Token Ring switches support the following standard MIBs:

- Network Management of TCP/IP-based Internets, MIB-II (RFC 1213)
- Evolution of Interfaces Group of MIB-II (RFC 1573)
- Definitions of Managed Objects for Bridges (RFC 1493)
- Token Ring Extensions to the Managed Objects for Bridges (RFC 1525)
- IEEE 802.5 Token Ring MIB (RFC 1748)
- RMON MIB/Token Ring Extensions (RFC 1757/1513) partial support
- IEEE 802.5 DTR Concentrator MIB (Catalyst 3900 only)
- IEEE 802.5 DTR MAC MIB (Catalyst 3900 only)

In addition, the Catalyst Token Ring switches support the following Cisco-defined private MIBs:

- Cisco VLAN Trunking Protocol MIB
- Cisco Discovery Protocol MIB

For SNMP-managed switches, it is possible to monitor and configure the switch from a network management application, which typically has a GUI that provides a simulated view of the front and rear panels of the switch.

Cisco provides SNMP-based network management applications that can be used to manage switches. For more information about these applications, see the“Configuring and Managing Token Ring Switches Using Cisco’s Network Management Products” chapter.

RMON Management

RMON is an industry-standard method for providing network statistics monitoring using SNMP. It also collects fault, performance, and configuration statistics. RMON can also be used to supplant traffic analyzers by providing packet capture or tracing data through the switch or on a ring.

In typical SNMP management, the SNMP client has to continuously poll the switch for fault, performance, and configuration information while waiting for the value to change. This causes increased traffic through the network. With RMON, you can have the switch monitor a particular statistic internally and when the statistic reaches a threshold the switch sends a trap to the client. This monitoring method reduces traffic between the SNMP client and the switch.

It is expensive to provide full-packet capture in a Token Ring switch because of the amount of memory required to store the information. Therefore, a solution is to use an internal RMON capability to gather traffic statistics and an external RMON probe for packet capture and higher-layer protocol analysis. The external RMON probe can be connected to the switch via a port mirroring port such as Cisco’s SPAN ports.

As an option, the Catalyst Token Ring switches provide RMON support for statistics, history, alarms, and events. They also provide support for the following groups of the Token Ring extensions to the Remote Network Monitoring MIB (RFC 1513):

- MAC-layer Statistics Group—Collection of MAC-layer statistics kept for each Token Ring interface, such as the total number of MAC packets received and the number of times the port entered a beaconing state.
- Promiscuous Statistics Group—Collection of promiscuous statistics kept for non-MAC packets on each Token Ring interface, such as the total number of good non-MAC frames received that were directed to an LLC broadcast address.
- Token Ring Ring Station Group—The Catalyst Token Ring switches support the ringStationControlTable portion of the Token Ring Ring Station Group. This support allows a Catalyst Token Ring switch to gather segment information from each ring segment to which it is attached. This segment information includes Ring State, Beacon Sender, Beacon NAUN, and Active Monitor MAC Address, as well as Station Order Changes.
- Token Ring Ring Station Order Group—List of the order of stations on the monitored rings.
- Token Ring Ring Station Group—List of ring station entries. An entry exists for each station that is currently or has previously been detected as being physically present on the ring.
- Token Ring Ring Station Config Control Group—List of ring station configuration control entries. Each entry controls the management of stations by a probe. One entry exists in this table for each active station in the ring station table.

You can use an external RMON probe for full RMON support.

Access to RMON data is available only via an SNMP management application that supports RFC 1757 and RFC 1513. You cannot access RMON via the switch's console interface; however, the console statistics provide similar information. For full utilization of RMON data, you should use the traffic management services of CWSI. For more information about CWSI, see the "Configuring and Managing Token Ring Switches Using Cisco's Network Management Products" chapter.

Topology Management

To aid in network management, Cisco developed the Cisco Discovery Protocol (CDP). CDP allows the Catalyst Token Ring switches to establish communication with other models of Cisco equipment. CDP support is provided as part of the Cisco IOS software that runs on many types of Cisco equipment.

CDP is a media- and protocol-independent protocol that is intended to be run on Cisco-manufactured equipment including routers, bridges, access servers, and switches. With CDP, Cisco's network management applications and Cisco devices can learn the device type and the SNMP agent address of neighboring devices. This enables applications to send SNMP queries to neighboring devices.
CDP runs on various media that support the SNAP, including LAN, Frame Relay, and ATM media. CDP runs over the data link layer only. Therefore, two systems that support different network-layer protocols can learn about each other.

How CDP Works

All Cisco devices transmit CDP packets periodically. These packets advertise a time-to-live value in seconds, which indicates the length of time that the packet must be retained before it can be discarded. CDP packets are sent with a time-to-live value that is nonzero after an interface is enabled and with a time-to-live value of zero immediately before an interface is idled down. This provides for quick state discovery.

All Cisco devices receive CDP packets and cache the information in the packet. The cached information is available to network management. Cisco devices never forward a CDP packet. If any information changes from the last received packet, the new information is cached and the older information is discarded even if its time-to-live value has not yet expired.

CDP Frame Format

CDP is assigned the Cisco HDLC protocol type value 0x2000. A Cisco-proprietary SNAP value enumerates HDLC protocol type values so CDP can run on all media that support SNAP, such as LAN media, Frame Relay, and ATM.

The SNAP format is as follows:
- LLC—0xAAAA03
- Org ID—0x00000C
- HDLC protocol type—0x2000

CDP sends packets on LANs using the multicast address 0100.0CCC.CCCC.

Because CDP does not run on top of any network layer, but rather runs only over the data link layer, two systems that support different network layer protocols can use CDP to learn about each other.

For more information about the format of CDP frames, see the "Frame Formats" appendix.

VLAN Management

To help VLAN management, Cisco developed VTP. As explained in the "VLAN Trunking Protocol" section of the "Token Ring VLANs and Related Protocols" chapter, VTP is used to set up and manage VLANs across an entire management domain. Using VTP, you can configure and manage the VLANs within a management domain from a single switch that is configured to act as a VTP server.

In addition, Cisco provides the following SNMP-based network management applications that can be used to manage VLANs:

- CiscoView, which allows you to configure and manage VLANs remotely.
- VLAN Director, which provides a graphical mapping utility for viewing and configuring VLANs.
- CWSI, which provides traffic management on a per VLAN basis.

For more information about these applications, see the "Configuring and Managing Token Ring Switches Using Cisco's Network Management Products" chapter.

Traffic Management

To help with traffic management, the Catalyst Token Ring switches support the SPAN.

With SPAN, traffic from any port on the switch can be mirrored or copied to another port, which is designated as the SPAN port. You can then connect the SPAN port to an external RMON probe.

This capability allows you to use the internal RMON to determine where problems might exist, and the external RMON to perform detailed problem analysis. For example, if the internal RMON statistics show high traffic on port 5, you can set up an external RMON probe remotely to capture data from port 5 to obtain more information.

Because forwarding to the SPAN port takes place independently of the normal forwarding, switch performance is not impacted.

As an alternative, you can establish a 16-Mbps monitor ring from the centralized data center that connects to all the SPAN ports on Token Ring switches. Then, you can connect the RMON probe or traffic analyzer at the data center, and via software control, the RMON probe can monitor any port on any switch in the network. Central control of remote monitoring is a powerful tool for the network manager.

The Catalyst 3900 and Catalyst 5000 Token Ring module allow you to configure active monitors. An active port monitor allows you to use a customer-supplied trace tool, such as a Network General Sniffer, to monitor only the LLC traffic that is switched by the monitored port. The MAC frames are not monitored.

On the Catalyst 3900, configuring a port to be a SPAN port removes the port from the TrCRF to which it is currently assigned. On the Catalyst 3900, you can also monitor traffic that is processed by an ISL or ATM port on a per-TrCRF basis (monitoring only one TrCRF at a time). You cannot, however, use an ISL or ATM port to monitor other ports.

Caution Using SPAN on more than one switch at a time may overload the monitoring ring. Also monitoring a TrCRF on a high-speed uplink, such as ATM or ISL, may overload the monitoring ring.

Soft Error Management

The Catalyst Token Ring switches perform error detection and isolation by monitoring the Report Soft Error MAC frames generated by stations on each port. Soft errors occur during normal ring operation and do not typically disrupt traffic on the ring. However, soft errors can occur at a rate that could potentially degrade the performance of the ring.

With both the Catalyst 3900 and the Catalyst 5000 series Token Ring module, you can configure soft error thresholds and sampling intervals for a port. During the interval you define, the Catalyst 3900 monitors the stations on the port and if the threshold is exceeded, can be configured to generate a trap indicating the port number and the station on which the threshold was exceeded. If necessary, you can issue a Remove Ring Station MAC frame to remove the station from the ring.

In summary, the Catalyst Token Ring switches:

- Monitor the Report Soft Error MAC frames generated by stations on each port, collect the data from each soft error frame, and generate a trap containing the port number and station on which the user-defined soft error threshold was exceeded.
- Report the soft error monitoring statistics via the console (for the Catalyst 3900), command line interface (for the Catalyst 5000 series Token Ring module), and SNMP.
- Provide the ability to issue a Remove Ring Station MAC frame to remove a station that is reporting a high level of errors or is not authorized to be on a ring.

Using a Switch for Ring Microsegmentation

The Catalyst 3900 and the Catalyst 5000 Token Ring switching module are shipped with a default configuration that allows you to use the switch without modification in many small networks. One aspect of this default configuration is that the switch is configured as a single VLAN. However, for more complex networks, you can subdivide the Catalyst 3900 or Catalyst 5000 Token Ring switching module into multiple virtual rings (TrCRFs) that can be connected by one or more internal bridges (TrBRFs). Initially, all ports are assigned to the default ring (trcrf-default) and the default ring is associated with the default bridge (trbrf-default).

Note: The Catalyst 5000 series Token Ring module default VLAN configuration requires that VTP V2 be enabled on the switch. VTP V2 is always enabled on the Catalyst 3900.

To assist you in understanding how to subdivide your switch, this chapter provides an example of configuring two additional VLANs for a Catalyst 3900.

This chapter provides the following information:

- Initial Network Configuration
- Before Beginning
- Configuration Steps
- Resulting Network
- Tips
- Microsegmenting the Rings on a Catalyst 5000

Note: Instructions for creating a similar configuration using two Catalyst 5000 Series Token Ring switching modules are included in the "Microsegmenting the Rings on a Catalyst 5000" section.

Initial Network Configuration

In this scenario, you have a small company that is growing. Last year, there were only 10 employees in the human resources and payroll departments. Now there are 34 employees. When there were only 10 employees, they could share a single server that contains a database of records. Now, however, each department needs a dedicated server.

Figure 5-1 illustrates the initial VLAN configuration of the Catalyst 3900.

You want to add a new ring that includes ports 1 and 2 for the employees of the Human Resources department and another ring that includes ports 3 and 4 for the employees of the Payroll department.

Figure 5-1 Initial VLAN Configuration

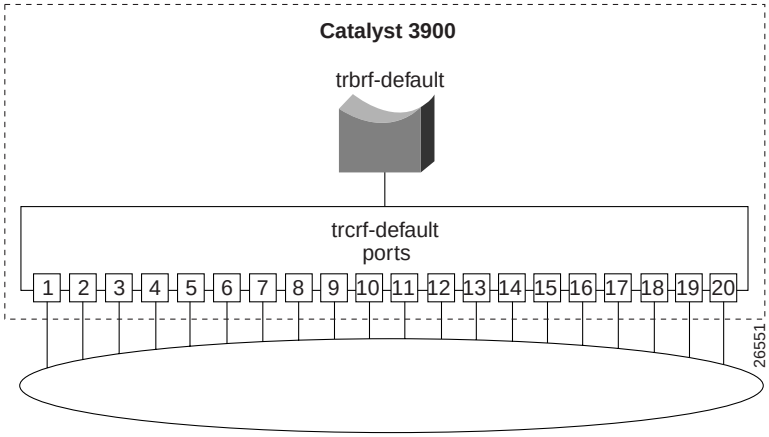

Before Beginning

Only the default ring (TrCRF) can be assigned to the default bridge (TrBRF). You cannot assign new rings to the default bridge. Therefore, you must first define a new bridge (TrBRF) and then you can define the new rings and assign ports to them.

You have met with the IS department and have decided to create two new rings, with ring numbers 11 and 12, and connect them with a bridge, which will have the bridge number of 1. Because the network contains a large number of Cisco devices, you are using VTP to distribute information about the VLANs in the network. You have decided to assign the VLAN IDs as follows:

Ring number	VLAN ID	VLAN Name
11	11	Human Resources Ring 11
12	12	Payroll Ring 12

The bridge will be assigned a VLAN ID of 100 and a VLAN name of BRF100.

Configuration Steps

Microsegmenting the ring involves creating multiple rings, which means you are creating multiple VLANs. You are going to put the users and their servers in separate TrCRFs and join them using a TrBRF.

Separating the Servers from the Users

You have physically separated the servers from the users. Next, you must attach the rings and the servers to separate ports on the Catalyst 3900 switches.

On both switches, do the following:

- Attach port 1 to the Human Resources ring.
- Remove the Human Resources server from the Human Resources ring and attach it to port 2.
- Attach port 3 to the Payroll ring.
- Remove the Payroll server from the Payroll ring and attach it to port 4.

The ports will automatically sense the speed and mode of the connection.

Configuring VLANs

Next, you must define the VLANs. As determined before beginning, you will need a new TrBRF and two TrCRFs; one for the Human Resources users and their server and one for the Payroll users and their server.

Note: For more information about Token Ring VLANs, see the "Token Ring VLANs and Related Protocols" chapter.

Defining the Bridges

To define a bridge (TrBRF), complete the following steps:

Step 1. On the Catalyst 3900 Main Menu, select **Configuration**. The Configuration panel is displayed.

Step 2. On the Configuration panel, select **VLAN and VTP Configuration**. The VLAN and VTP Configuration panel is displayed.

Step 3. On the VLAN and VTP Configuration panel, select **VTP VLAN Configuration**. The VTP VLAN Configuration panel is displayed.

Step 4. On the VTP VLAN Configuration panel, select **Add**.

Step 5. At the prompt, enter a VLAN ID of **100**.

Step 6. At the prompt, select **TrBRF**. The VLAN Parameter Configuration for TrBRF panel (Figure 5-2) is displayed.

Step 7. On the VLAN Parameter Configuration for TrBRF panel, specify:

- VLAN Name of **BRF100**.
- Bridge Number of **1**.

Figure 5-2 VLAN Parameter Configuration for TrBRF Panel

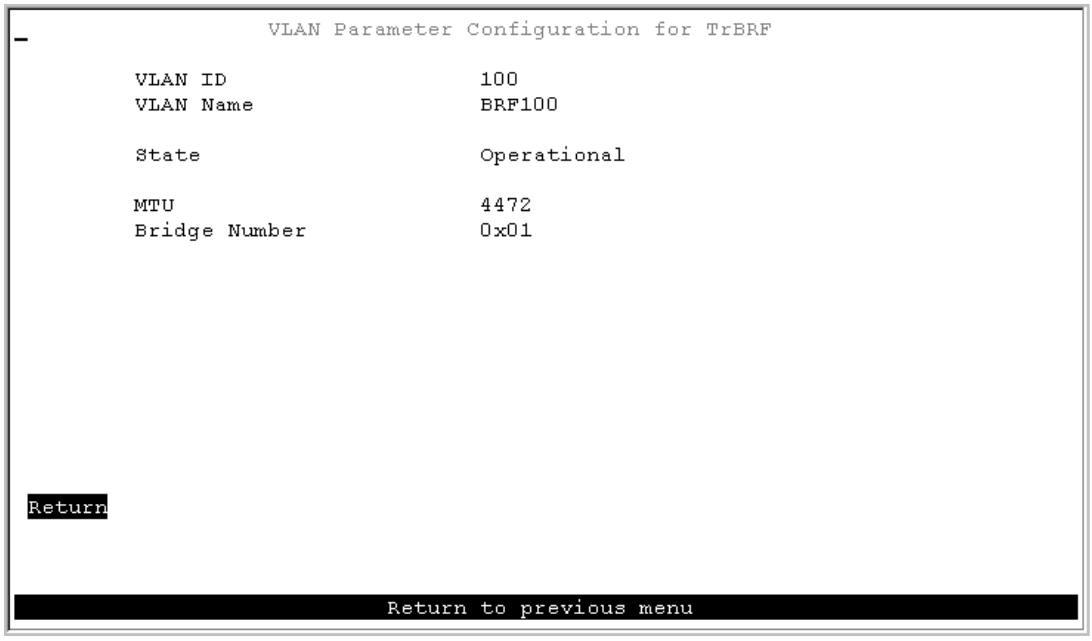

Step 8. Select **Return** to save your changes.

Figure 5-3 illustrates the VLAN configuration of the Catalyst 3900 after the additional bridge has been configured. Notice that no rings are assigned to it yet.

Figure 5-3 Catalyst 3900 with Two Bridges Configured

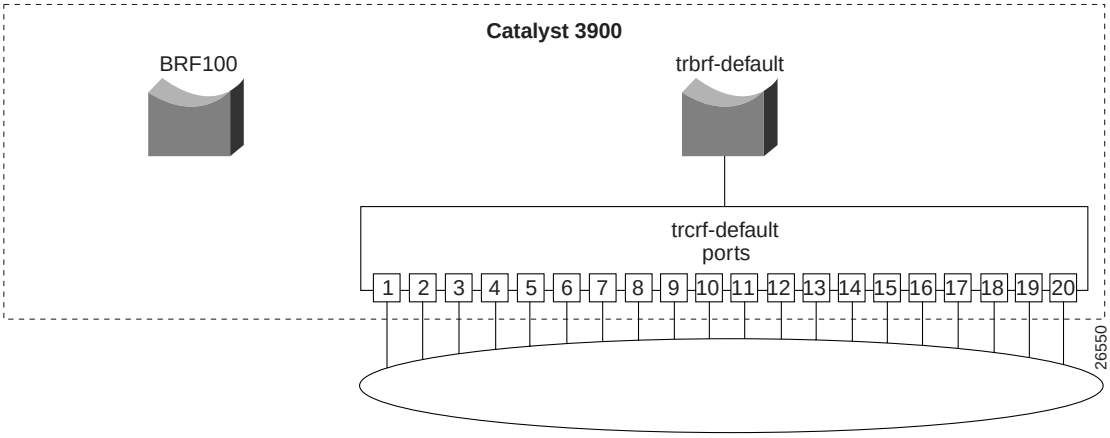

Defining the Rings

To define the ring (TrCRF) for the Human Resources users, complete the following steps:

Step 1. On the VTP VLAN Configuration panel, select **Add**.

Step 2. At the prompt, enter a VLAN ID of **11**.

Step 3. At the prompt, select **TrCRF**. The VLAN Parameter Configuration for TrCRF panel (Figure 5-4) is displayed.

Step 4. On the VLAN Parameter Configuration for TrCRF panel, specify:

- VLAN Name of **Human Resources Ring 11**.
- Parent VLAN of **BRF100**.
- Ring Number of **11**.

Figure 5-4 VLAN Parameter Configuration for TrCRF Panel

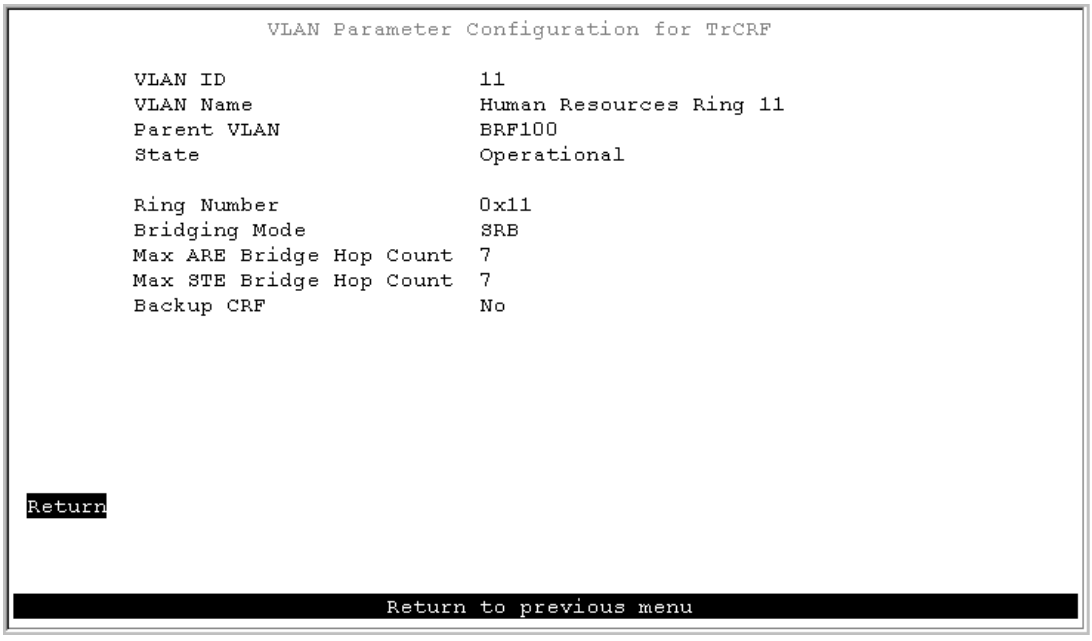

Step 5. Select **Return** to save your changes.

To define the ring (TrCRF) for the Payroll users, repeat Step 1 through Step 4 and use the following values:

- VLAN ID of **12**.
- VLAN Name of **Payroll Ring 12**.
- Parent VLAN of **BRF100**.
- Ring Number of **12**.

Figure 5-5 illustrates the VLAN configuration of the Catalyst 3900 after the additional rings have been configured. Notice that the rings are configured and associated with the bridge, but no ports are assigned to the rings.

Figure 5-5 Catalyst 3900 with Three Rings Configured

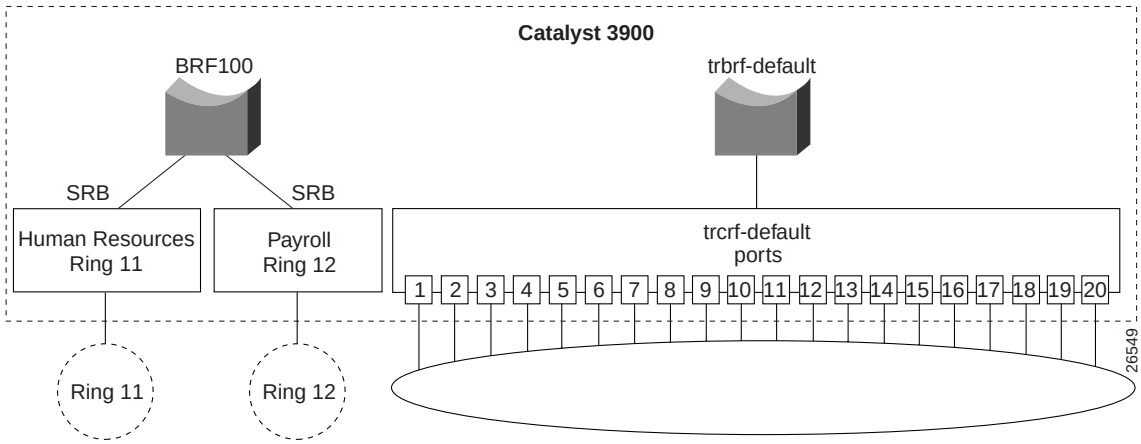

Assigning Ports to the Rings

Next, you must assign the ports to the appropriate rings (TrCRFs). On the Catalyst 3900, do the following:

Step 1. On the VLAN and VTP Configuration panel, select **Local VLAN Port Configuration**. The Local VLAN Port Configuration panel is displayed.

Step 2. On the Local VLAN Port Configuration panel, select **Change**.

Step 3. At the prompt enter port number **1**.

Step 4. Select **Human Resources Ring 11** from the list of possible TrCRFs. To select the TrCRF, use your cursor movement keys to highlight the desired TrCRF, press the space bar to select it, and press **Enter** to implement the change.

Step 5. Repeat Step 2 through Step 4 for port 2.

Step 6. Again, on the Local VLAN Port Configuration panel, select **Change**.

Step 7. At the prompt enter port number **3**.

Step 8. Select **Payroll Ring 12** from the list of possible TrCRFs.

Step 9. Repeat Step 6 through Step 8 for port 4.

Step 10. Select **Return** to save the changes.

Figure 5-6 displays the Local VLAN Port Configuration Panel after you have made your changes.

Figure 5-6 Local VLAN Port Configuration Panel

```
                    Local VLAN Port Configuration

Port  Mode    TrCRF                                TrBRF
  1   Static  Human Resources Ring 11             BRF100
  2   Static  Human Resources Ring 11             BRF100
  3   Static  Payroll Ring 12                     BRF100
  4   Static  Payroll Ring 12                     BRF100
  5   Static  trcrf-default                       trbrf-default
  6   Static  trcrf-default                       trbrf-default
  7   Static  trcrf-default                       trbrf-default
  8   Static  trcrf-default                       trbrf-default
  9   Static  trcrf-default                       trbrf-default
 10   Static  trcrf-default                       trbrf-default
 11   Static  trcrf-default                       trbrf-default
 12   Static  trcrf-default                       trbrf-default
 13   Static  trcrf-default                       trbrf-default
 14   Static  trcrf-default                       trbrf-default
 15   Static  trcrf-default                       trbrf-default

Return        More       Change


                    Return to previous menu
```

Resulting Network

You now have a network with improved performance because the number of users per ring has been reduced and the servers have dedicated bandwidth (Figure 5-7).

Figure 5-7 Final Network Configuration

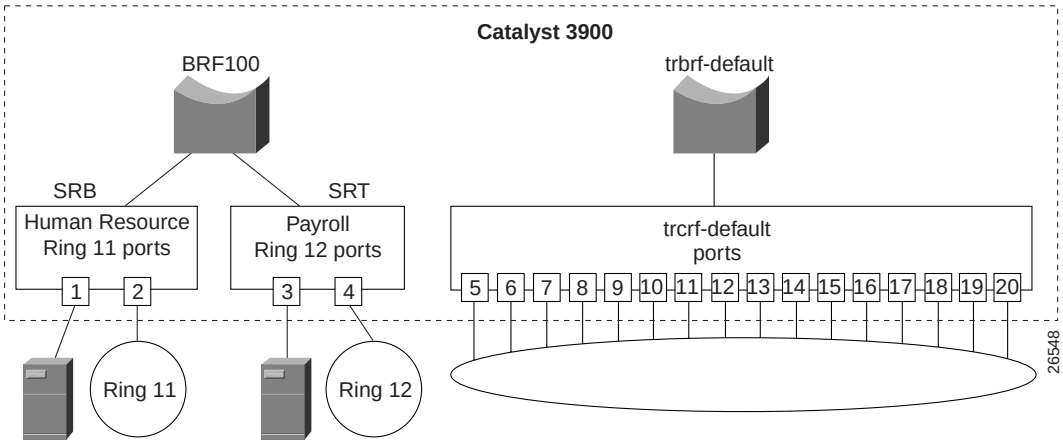

Tips

This section contains tips that may be useful in creating a configuration similar to the one in this scenario.

Configuring the STP

If you install an external bridge to create a backup path between rings 11 and 12, you introduce possible loops into your network. However, STPs prevent these loops. By default, the TrBRF runs the IBM STP. The STP run on the TrCRF can be manually configured, though. By default the TrCRF STP is determined by the bridging mode. TrCRFs with a bridging mode of SRB will run the IEEE STP and TrCRFs with a bridging mode of SRT will run the Cisco STP.

Selecting VLAN Names and IDs

To aid in network management and network identification, we recommend that:

- The VLAN name of a TrCRF should include the ring number.
- The VLAN name for a TrBRF should include the bridge number.
- The VLAN ID of a TrCRF should be the same as the ring number.

Improving Performance

To further improve performance, if you have 16 Mbps connections and the server's NIC supports FDX, you can configure the ports connected to the servers to operate in FDX mode. To configure FDX:

Step 1. Select **Port Configuration** on the Configuration panel.

Step 2. Specify the port to which the server is attached. In this scenario, that would be either port 2 or 4.

Step 3. On the Port Configuration panel, move to the Operation Mode and select a mode of **FDX port**.

Step 4. Select **Return**.

Microsegmenting the Rings on a Catalyst 5000

You can create a similar configuration using two Catalyst 5000 series Token Ring switching modules. The Catalyst 5000 provides a command line interface rather than a menu-driven interface, so the steps are slightly different. This section provides an overview of the configuration steps to achieve a similar configuration using two Catalyst 5000 Token Ring modules.

Defining the Bridge

To define the bridge (TrBRF), complete the following steps:

Step 1. At the Catalyst 5000 command prompt, enter **enable**.

Step 2. At the enable prompt, enter **set vlan 100 name brf100 type trbrf bridge 1**.

Step 3. To verify the configuration of the new VLAN, enter **show vlan**.

The output (Figure 5-8), indicates that brf100 has been added, but it does not have any TrCRFs assigned to it yet.

Figure 5-8 Output for show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
                                                2/1-48
100  brf100                           active
1002 fddi-default                     active
1003 trcrf-default                    active    3/1-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003
```

Defining the Rings

To define the ring (TrCRF) for the Human Resource users, complete the following steps:

Step 1. At the enable prompt, enter **set vlan 11 name hr-ring11 type trcrf ring 11 parent 100 mode srb**.

Step 2. To verify the configuration of the new VLAN, enter **show vlan**.

The output (Figure 5-9) indicates that hr-ring11 has been added, but it does not have any ports assigned to it yet. It also shows that brf100 is the parent of the VLAN with the ID of 11.

Figure 5-9 Output of show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
                                                2/1-48
11   hr-ring11                        active
100  brf100                           active    11
1002 fddi-default                     active
1003 trcrf-default                    active    3/1-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003


VLAN Type  SAID       MTU   Parent RingNo BrdgNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ ------ ---- -------- ------ ------
1    enet  100001     1500  -      -      -      -    -        0      0
11   trcrf 100110     4472  100    0x11   -      -    srb      0      0
100  trbrf 100100     4472  -      -      0x1    ibm  -        0      0
1002 fddi  101002     1500  -      0x0    -      -    -        0      0
1003 trcrf 101003     4472  1005   0xccc  -      -    srb      0      0
1004 fdnet 101004     1500  -      -      0x0    ieee -        0      0
1005 trbrf 101005     4472  -      -      0xf    ibm  -        0      0
```

To define the TrCRF for the Payroll users, do the following:

Step 1. At the enable prompt, enter **set vlan 12 name payroll-ring12 type trcrf ring 12 parent 100 mode srb**.

Step 2. To verify the configuration of the new VLAN, enter **show vlan**.

The output (Figure 5-10) indicates that payroll-ring12 has been added, but it does not have any ports assigned to it yet. It also shows that brf100 is the parent of the VLAN with the ID of 12.

Figure 5-10 Output of show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
                                                2/1-48
11   hr-ring11                        active
12   payroll-ring12                   active
100  brf100                           active    11, 12
1002 fddi-default                     active
1003 trcrf-default                    active    3/1-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003

VLAN Type  SAID       MTU   Parent RingNo BrdgNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ ------ ---- -------- ------ ------
1    enet  100001     1500  -      -      -      -    -        0      0
11   trcrf 100110     4472  100    0x11   -      -    srb      0      0
12   trcrf 100120     4472  100    0x12   -      -    srb      0      0
100  trbrf 100100     4472  -      -      0x1    ibm  -        0      0
1002 fddi  101002     1500  -      0x0    -      -    -        0      0
1003 trcrf 101003     4472  1005   0xccc  -      -    srb      0      0
1004 fdnet 101004     1500  -      -      0x0    ieee -        0      0
1005 trbrf 101005     4472  -      -      0xf    ibm  -        0      0
```

Assigning Ports to the Rings

To assign the ports to the rings (TrCRFs), complete the following steps:

Step 1. At the enable prompt, enter **set vlan 11 3/1-2**.

Step 2. At the enable prompt, enter **set vlan 12 3/3-4**.

The output (Figure 5-11) shows that ports 1 and 2 on module 3 are assigned to crf11 and that ports 3 and 4 on module 3 are assigned to crf12.

Figure 5-11 Output of show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
                                                2/1-48
11   hr-ring11                        active    3/1-2
12   payroll-ring12                   active    3/3-4
100  brf100                           active    11, 12
1002 fddi-default                     active
1003 trcrf-default                    active    3/5-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003
```

Configuring the STP

By default, the TrBRF runs the IBM STP. The STP run on the TrCRFs is determined by the specified bridging mode. TrCRFs with a bridge mode of SRB will run the IEEE STP and TrCRFs with a bridge mode of SRT will run the Cisco STP.

The Catalyst 5000 Token Ring switching module considers the combination of the IBM STP at the TrBRF and the bridge mode of SRT to be incompatible. As a result, if you had configured one of the TrCRFs (for example, payroll-ring12) with a bridge mode of SRT, the Catalyst 500 Token Ring switching module would automatically block the logical port of the TrCRF that is configured for SRT. Use the **show spantree** command to view the state of the logical ports (Figure 5-12).

Figure 5-12 Output of the show spantree Command

```
VLAN 100
Spanning tree enabled
Spanning tree type          ibm

Designated Root             00-e0-1e-2f-6c-63
Designated Root Priority    32768
Designated Root Cost        0
Designated Root Port        1/0
Root Max Age   6  sec    Hello Time 2  sec   Forward Delay 4  sec

Bridge ID MAC ADDR          00-e0-1e-2f-6c-63
Bridge ID Priority          32768
Bridge Max Age 6  sec    Hello Time 2  sec   Forward Delay 4  sec

Port,Vlan Vlan  Port-State     Cost   Priority  Fast-Start  Group-method
--------- ----  -------------  -----  --------  ----------  ------------
 1/2      100   forwarding        19        32   disabled
 11       100   forwarding        80        32   disabled
 12       100   blocking          80        32   disabled
* = portstate set by user configuration
```

You can then use the **set spantree portstate** command to change the forwarding mode of the logical port.

Classic Token Ring Bridged Network Migration

The Catalyst 3900 Token Ring switch and Catalyst 5000 Token Ring switch module can be used to collapse network backbones and floor rings in classic Token Ring environments with redundantly placed source-route bridges. The use of a Catalyst Token Ring switch in this scenario improves network performance by eliminating the need for multiple bridges and by allowing the direct attachment of high-utilization devices, such as servers, front-end processors, and routers.

This section provides an example of using Catalyst Token Ring switches to replace multiple bridges in a classic Token Ring network.

This chapter provides the following information:

- Initial Network Configuration
- Configuration Steps
- Cabling the Network
- Resulting Network Configuration
- Tips

Initial Network Configuration

In your company, you have two backbone rings that service five floor rings. You have a server, router, and front-end processor attached to each of the backbone rings. Because the number of users is growing and there is an increased need to access the devices that are attached to the backbone rings, you need to improve the performance of your network. You have decided to replace one ring with a Catalyst 3900 and one with a Catalyst 5000 with a Token Ring switching module.

Figure 6-1 illustrates the initial network configuration.

Figure 6-1 Initial Network Configuration

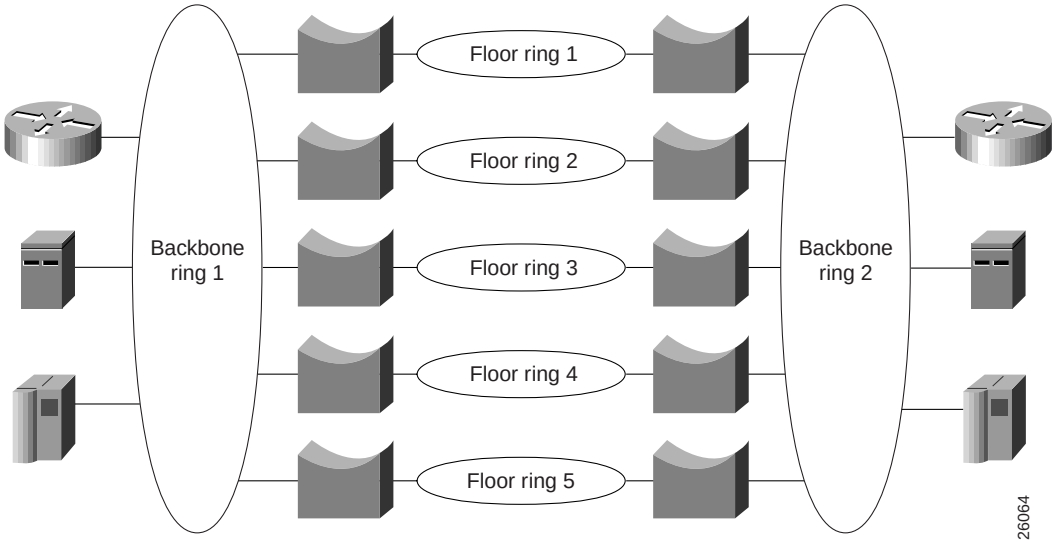

Configuration Steps

To configure your collapsed backbone network, on each switch you will need to configure a new bridge (TrBRF) as you cannot add TrCRFs to the default TrBRF, a ring (TrCRF) for each floor, as well as a TrCRF that will contain the high-utilization devices.

Note: For more information about Token Ring VLANs, see the "Token Ring VLANs and Related Protocols" chapter.

Configuring the Catalyst 3900

On the Catalyst 3900, you must configure a new bridge (TrBRF) and six new rings (TrCRFs).

Defining the Bridges

To define a bridge (TrBRF), complete the following steps:

Step 1. On the Catalyst 3900 Main Menu, select **Configuration**. The Configuration panel is displayed.

Step 2. On the Configuration panel, select **VLAN and VTP Configuration**. The VLAN and VTP Configuration panel is displayed.

Step 3. On the VLAN and VTP Configuration panel, select **VTP VLAN Configuration**. The VTP VLAN Configuration panel is displayed.

Step 4. On the VTP VLAN Configuration panel, select **Add**.

Step 5. At the prompt, enter a VLAN ID of **100**.

Step 6. At the prompt, select **TrBRF**. The VLAN Parameter Configuration for TrBRF panel (Figure 6-2) is displayed.

Step 7. On the VLAN Parameter Configuration for TrBRF panel, specify:

- VLAN Name of **BRF100**.
- Bridge Number of **1**.

Figure 6-2 VLAN Parameter Configuration for TrBRF Panel

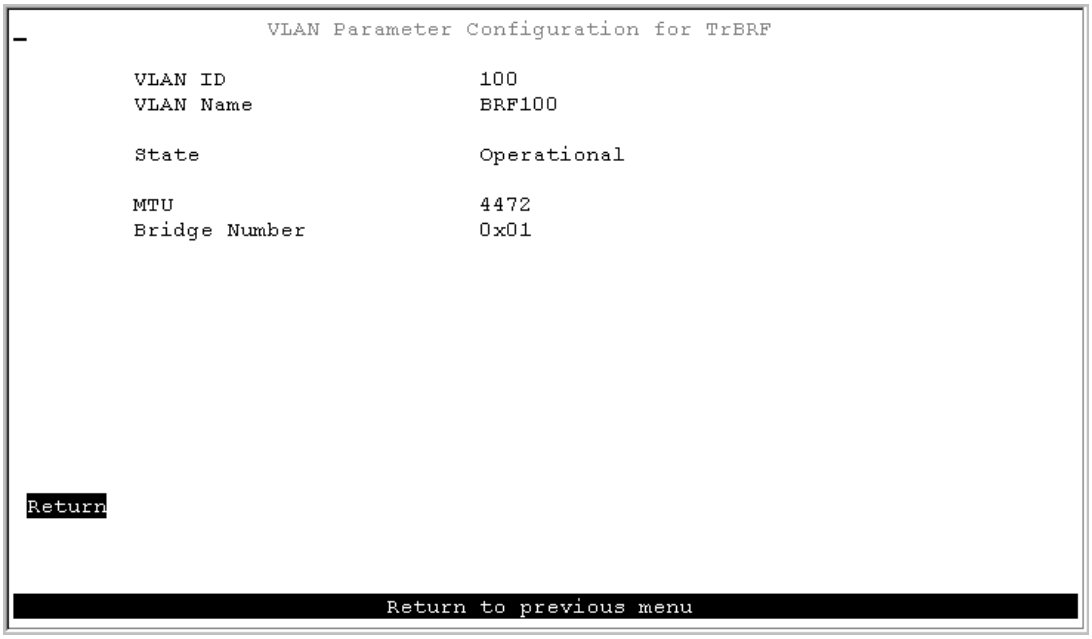

Step 8. Select **Return** to save your changes.

Defining the Rings

To define the rings (TrCRFs) for the first floor, complete the following steps:

Step 1. On the VTP VLAN Configuration panel, select **Add**.

Step 2. At the prompt, enter a VLAN ID of **101**.

Step 3. At the prompt, select **TrCRF**. The VLAN Parameter Configuration for TrCRF (Figure 6-3) panel is displayed.

Step 4. On the VLAN Parameter Configuration for TrCRF panel, specify:

- VLAN Name of **Floor 1**.
- Parent VLAN of **BRF100**.
- Ring Number of **1**.

Figure 6-3 VLAN Parameter Configuration for TrCRF Panel

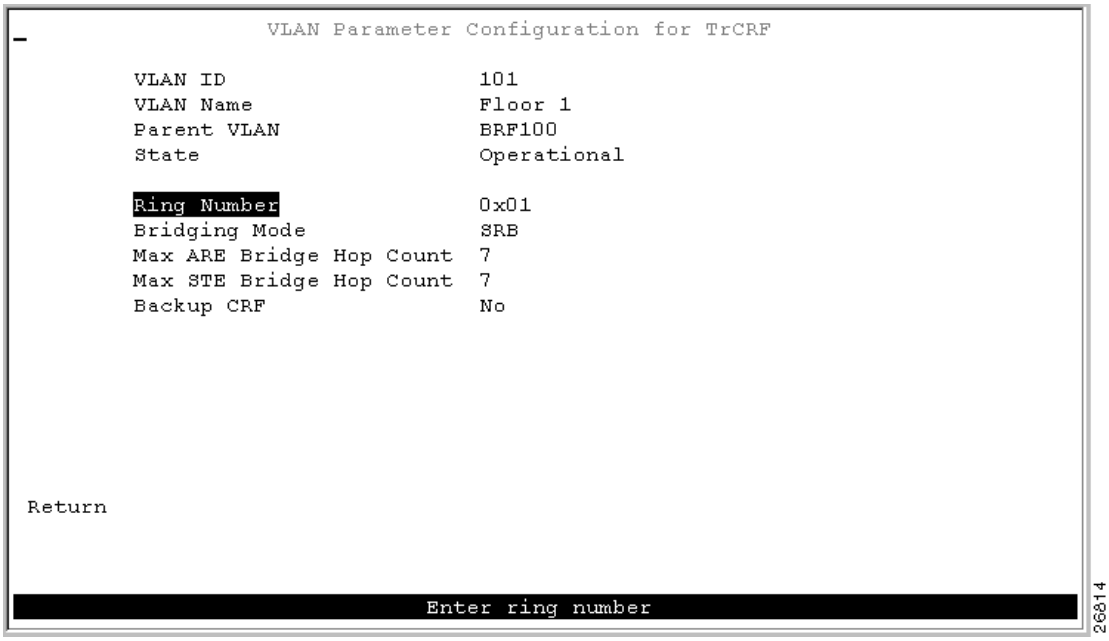

Step 5. Select **Return** to save your changes.

To define the TrCRFs for the remaining floors, repeat Step 1 through Step 5 and use the following values:

- VLAN IDs of **102**, **103**, **104**, and **105**.
- VLAN Names of **Floor 2**, **Floor 3**, **Floor 4**, and **Floor 5**.
- Parent VLAN of **BRF100**.
- Ring Numbers of **2**, **3**, **4**, and **5**.

To define the TrCRF for the high-utilization devices, repeat Step 1 through Step 5 and use the following values:

- VLAN ID of **106**.
- VLAN Name of **Server Ring 1**.
- Parent VLAN of **BRF100**.
- Ring Number of **6**.

Figure 6-4 shows the resulting VTP VLAN Configuration panel.

Figure 6-4 VTP VLAN Configuration Panel

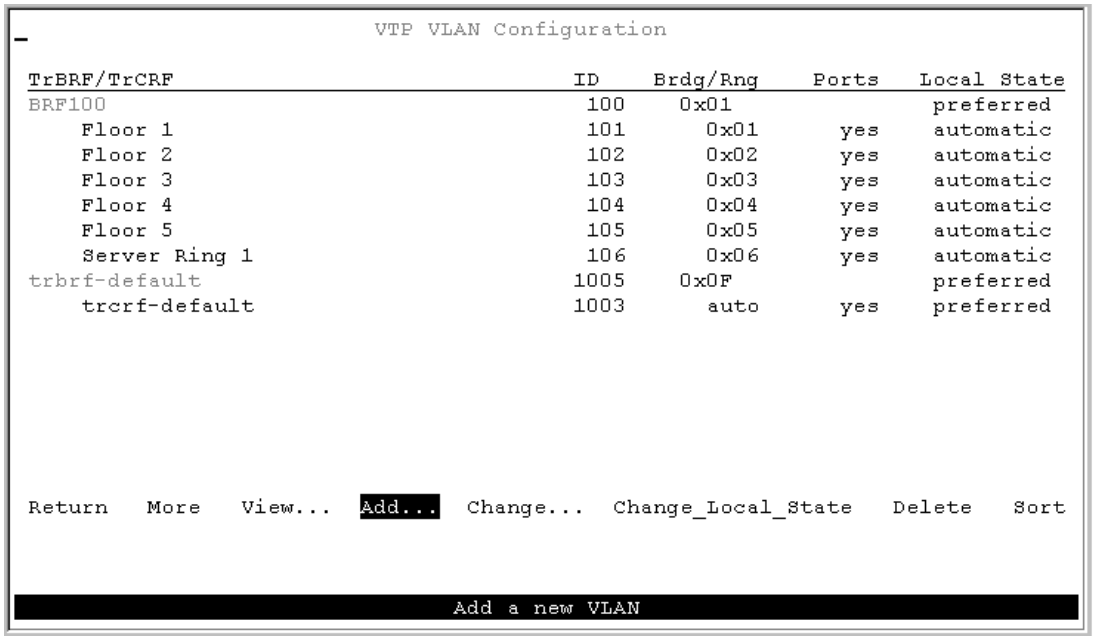

Assigning Ports to the Rings

Next, you must assign the ports to the appropriate TrCRFs. On the Catalyst 3900, complete the following steps:

Step 1. On the VLAN and VTP Configuration panel, select **Local VLAN Port Configuration**. The Local VLAN Port Configuration panel is displayed.

Step 2. On the Local VLAN Port Configuration panel, select **Change**.

Step 3. At the prompt enter port number **1**.

Step 4. Select **Floor 1** from the list of possible TrCRFs. To select the TrCRF, use your arrow keys to highlight the desired TrCRF, press the space bar to select it, and press **Enter** to implement your change (Figure 6-5).

Figure 6-5 Local VLAN Port Configuration Panel

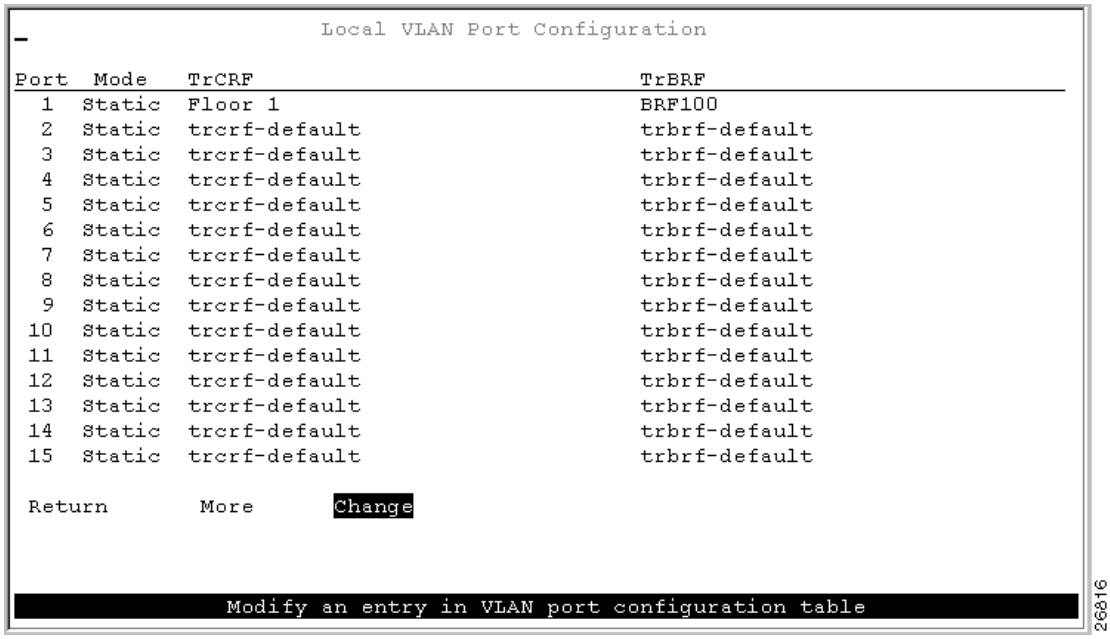

Step 5. Repeat Step 2 through Step 4 to assign the ports to the appropriate TrCRFs as follows:

Ports	TrCRF
2	Floor 1
3, 4	Floor 2
5, 6	Floor 3
7, 8	Floor 4
9, 10	Floor 5
11, 12, 13	Server Ring 1

Step 6. Select **Return** to save your changes.

Configuring the STP

If you install an external bridge to create a backup path, you introduce possible loops into your network. However, STPs prevent these loops. By default, the TrBRF runs the IBM STP. The STP run on the TrCRF can be manually configured, however, by default the TrCRF STP is determined by the bridging mode. TrCRFs with a bridging mode of SRB will run the IEEE STP and TrCRFs with a bridging mode of SRT will run the Cisco STP.

Note: You must assign the ports to the TrCRFs before you can configure spanning-tree parameters for the TrCRFs.

Configuring the Catalyst 5000

On the Catalyst 5000, you must configure a new bridge (TrBRF), 6 new rings (TrCRFs), and the STP. You have inserted the Token Ring module into slot 2 of the Catalyst 5000.

Defining the Bridge

To define the bridge (TrBRF), complete the following steps:

Step 1. At the Catalyst 5000 command prompt, enter **enable**.

Step 2. At the enable prompt, enter **set vlan 200 name brf200 type trbrf bridge 2**.

Step 3. To verify the configuration of the new VLAN, enter **show vlan**.
The output (Figure 6-6) indicates that brf200 has been added, but it does not have any TrCRFs assigned to it yet.

Figure 6-6 Output for show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
200  brf200                           active
1002 fddi-default                     active
1003 trcrf-default                    active    2/1-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003
```

Defining the Rings

To define the ring (TrCRF) for the first floor, complete the following steps:

Step 1. At the enable prompt, enter:

```
set vlan 201 name Floor_1 type trcrf ring 1 parent 200 mode srb
```

Step 2. To verify the configuration of the new VLAN, enter **show vlan**.
The output (Figure 6-7) indicates that FLoor_1 has been added, but it does not have any ports assigned to it yet. It also shows that brf200 is the parent of the VLAN with the ID of 201.

Figure 6-7 Output of Show VLAN Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
200  brf200                           active    201
201  Floor_1                          active
1002 fddi-default                     active
1003 trcrf-default                    active    3/1-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003

VLAN Type  SAID       MTU   Parent RingNo BrdgNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ ------ ---- -------- ------ ------
1    enet  100001     1500  -      -      -      -    -        0      0
200  trbrf 100200     4472  -      -      0x2    ibm  -        0      0
201  trcrf 100201     4472  100    0x01   -      -    srb      0      0
1002 fddi  101002     1500  -      0x0    -      -    -        0      0
1003 trcrf 101003     4472  1005   0xccc  -      -    srb      0      0
1004 fdnet 101004     1500  -      -      0x0    ieee -        0      0
1005 trbrf 101005     4472  -      -      0xf    ibm  -        0      0
```

To define the TrCRFs for the remaining floors, enter the **set vlan** commands as follows:

```
set vlan 202 name Floor_2 type trcrf ring 2 parent 200 mode srb
set vlan 203 name Floor_3 type trcrf ring 3 parent 200 mode srb
set vlan 204 name Floor_4 type trcrf ring 4 parent 200 mode srb
set vlan 205 name Floor_5 type trcrf ring 5 parent 200 mode srb
```

To define the TrCRF for the server ring, enter the **set vlan** commands as follows:

```
set vlan 207 name Server_Ring_2 type trcrf ring 7 parent 200 mode srb
```

The output (Figure 6-8) indicates that the TrCRFs have been added, but there are no ports assigned to them yet. It also shows that brf200 is the parent of the new TrCRFs.

Figure 6-8 Output of show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
200  brf200                           active    201, 202, 203, 204, 205, 207
201  Floor_1                          active
202  Floor_2                          active
203  Floor_3                          active
204  Floor_4                          active
205  Floor_5                          active
207  Server_Ring_2                    active
1002 fddi-default                     active
1003 trcrf-default                    active    2/1-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003


VLAN Type  SAID       MTU   Parent RingNo BrdgNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ ------ ---- -------- ------ ------
1    enet  100001     1500  -      -      -      -    -        0      0
200  trbrf 100200     4472  -      -      0x2    ibm  -        0      0
201  trcrf 100201     4472  200    0x01   -      -    srb      0      0
202  trcrf 100202     4472  200    0x02   -      -    srb      0      0
203  trcrf 100203     4472  200    0x03   -      -    srb      0      0
204  trcrf 100204     4472  200    0x04   -      -    srb      0      0
205  trcrf 100205     4472  200    0x05   -      -    srb      0      0
207  trcrf 100207     4472  200    0x07   -      -    srb      0      0
1002 fddi  101002     1500  -      0x0    -      -    -        0      0
1003 trcrf 101003     4472  1005   0xccc  -      -    srb      0      0
1004 fdnet 101004     1500  -      -      0x0    ieee -        0      0
1005 trbrf 101005     4472  -      -      0xf    ibm  -        0      0
```

Assigning Ports to the Rings

To assign the ports to the rings (TrCRFs), enter the **set vlan** commands at the enable prompt as follows:

```
set vlan 201 23/1-2
set vlan 202 2/3-4
set vlan 203 2/5-6
set vlan 204 2/7-8
set vlan 205 2/9-10
set vlan 207 2/11-13
```

The output (Figure 6-9) shows that two ports on the module are assigned to each of the five TrCRFs that represent each floor and that three ports are assigned to Server_Ring_2.

Figure 6-9 Output of show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
200  brf200                           active    201, 202, 203, 204, 205, 207
201  Floor_1                          active    2/1-2
202  Floor_2                          active    2/3-4
203  Floor_3                          active    2/5-6
204  Floor_4                          active    2/7-8
205  Floor_5                          active    2/9-10
207  Server_Ring_2                    active    2/11-13
1002 fddi-default                     active
1003 trcrf-default                    active    2/14-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003
```

Configuring the STP

By default, the TrBRF runs the IBM STP. The STP run on the TrCRFs is determined by the specified bridging mode. TrCRFs with a bridge mode of SRB will run the IEEE STP and TrCRFs with a bridge mode of SRT will run the Cisco STP.

Cabling the Network

Using the appropriate cabling, attach ports 1 and 2 of the Catalyst 3900 to the existing ring on floor 1. Repeat this for each of the floors. Then attach port 11 to the router, port 12 to the server, and port 13 to the front-end processor. Do the same for the ports on the Catalyst 5000.

Resulting Network Configuration

You now have a faster, more efficient network that includes less hardware to maintain (Figure 6-10).

Figure 6-10 Resulting Network

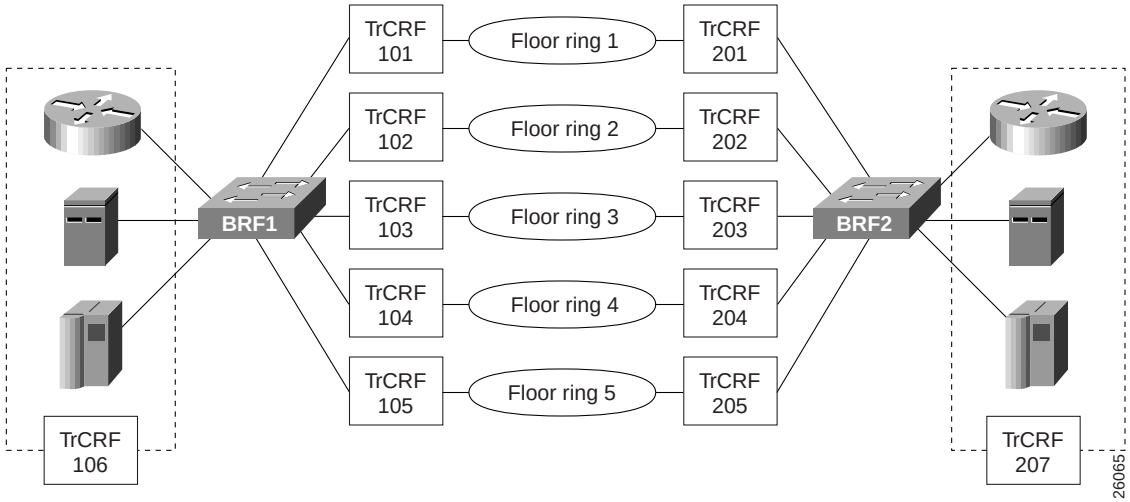

Tips

To further improve performance, if you have 16 Mbps connections and the network interface card (NIC) supports full-duplex, you can configure the ports connected to the servers to operate in FDX mode. To configure FDX:

Step 1. Select **Port Configuration** on the Configuration panel.

Step 2. Specify the port to which the high-utilization device is attached. In this scenario, that would be ports 11, 12, and 13.

Step 3. On the Port Configuration panel, move to the Operation Mode and select the **FDX port** mode.

Step 4. Select **Return**.

Connecting Two Catalyst 3900s via TokenChannel

In this chapter, we are going to configure a TokenChannel connection between two Catalyst 3900s. TokenChannels consist of two to eight parallel connections between two Catalyst 3900s. Because of the increased aggregate speed and the fact that channels provide load balancing by destination address, these parallel channels provide improved performance and are fault-tolerant.

A single TokenChannel can consist of a combination of HDX and FDX connections. For example, a TokenChannel consisting of three connections can have one HDX and two FDX connections. However, both ports in each interconnected pair must be either HDX or FDX. In addition, all ports in a single TokenChannel must belong to the same TrCRF on the Catalyst 3900.

Caution While you can use TokenChannels to interconnect Catalyst 3900s and Catalyst 3920s, you cannot use TokenChannels to interconnect other different models of switches. For example, you cannot use a TokenChannel to interconnect a Catalyst 2600 and a Catalyst 3900. Likewise, you cannot use a TokenChannel to interconnect a Catalyst 3900 and a non-Cisco switch.

Note: Similarly, you could also create an ISL Channel. ISL Channels consist of two to four ISL parallel connections between two Catalyst 3900s, a Catalyst 3900 and a Catalyst 5000, a Token Ring ISL-capable Cisco router, or a Token Ring ISL network adapter. The Catalyst 3900 Token Ring Release 4.1(1) or later supports the configuration of ISL Channels. All connections in an ISL Channel must be FDX.

Note: When the Catalyst 3900 is configured with channels, all broadcast frames use the primary (lowest numbered) port of the channel.

This chapter provides the following information:

- Initial Network Configuration
- Before Beginning
- Configuration Steps
- Resulting Network Configuration
- Tips
- Troubleshooting

Initial Network Configuration

This scenario expands on the microsegmented network described in the "Using a Switch for Ring Microsegmentation" chapter.

Your company has continued to grow. You have leased the floor above your current offices and have moved your Accounts Payable department and your Financial Analysts to that floor. You installed another Catalyst 3900 on that floor (Catalyst 3900 "B"). Similar to your Payroll and Human Resources departments, these two departments required dedicated server support. Therefore, you microsegmented Catalyst 3900 "B" and created a VLAN configuration (Figure 7-1)similar to the one you created in the "Using a Switch for Ring Microsegmentation" chapter. However, these four departments often have a need to share information and so you have decided to connect the two Catalyst 3900s.

Figure 7-1 Initial Network Configuration

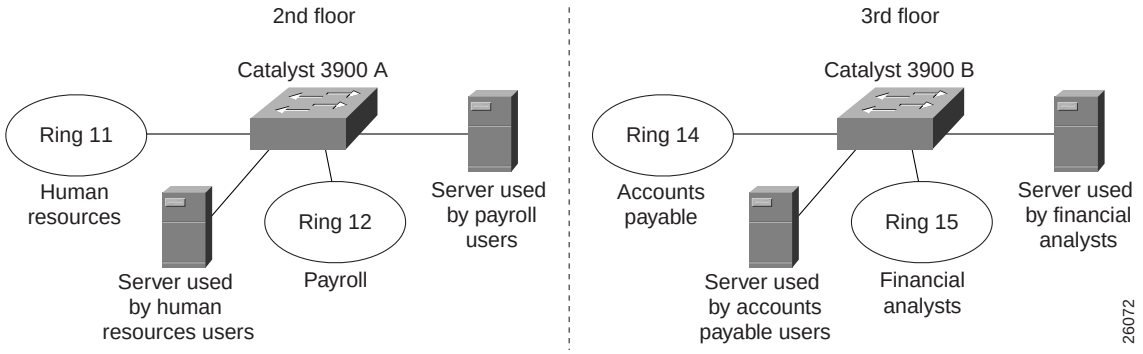

Before Beginning

To create a TokenChannel, you do not need any additional components or equipment. Support for TokenChannels was included in the initial release of the Catalyst 3900. You simply need standard cables with RJ-45 connectors to connect the ports of the two Catalyst switches.

The TrCRFs that were created on Catalyst 3900 "B" were assigned ring numbers of 14 and 15. You assigned the TrBRF that you created on Catalyst 3900 "B" a VLAN ID of 200 and a bridge number of 2.

You have decided to assign the ports of the TokenChannel to a unique TrCRF on each switch. You will assign the TrCRF a ring number of 16. You have decided to create a 3-port TokenChannel using ports 18, 19, and 20 on each switch.

Configuration Steps

To create the TokenChannel, you will first create the new TrCRF on each switch, then add the ports to those TrCRFs, and then define those ports as members of a TokenChannel.

Note: You must define a channel for both connected Catalyst 3900s before physically connecting the linked ports. Therefore, make sure that you have either disabled the ports or disconnected the cables before you configure a channel to avoid creating loops.

Defining the TrCRF

First, you must define a TrCRF for the TokenChannel on each switch. On switch A, complete the following steps:

Step 1. On the Catalyst 3900 Main Menu, select **Configuration**. The Configuration panel is displayed.

Step 2. On the Configuration panel, select **VLAN and VTP Configuration**. The VLAN and VTP Configuration panel is displayed.

Step 3. On the VLAN and VTP Configuration panel, select **VTP VLAN Configuration**. The VTP VLAN Configuration panel is displayed.

Step 4. On the VTP VLAN Configuration panel, select **Add**.

Step 5. At the prompt, enter a VLAN ID of **16**.

Step 6. At the prompt, select **TrCRF**. The VLAN Parameter Configuration for TrCRF panel (Figure 7-2) is displayed.

Step 7. On the VLAN Parameter Configuration for TrCRF panel, specify:

- VLAN Name of **CRF16**.
- Parent VLAN of **BRF100**. (This TrBRF was configured in the "Using a Switch for Ring Microsegmentation" chapter.)
- Ring Number of **16**.

Figure 7-2 VLAN Parameter Configuration for TrCRF Panel

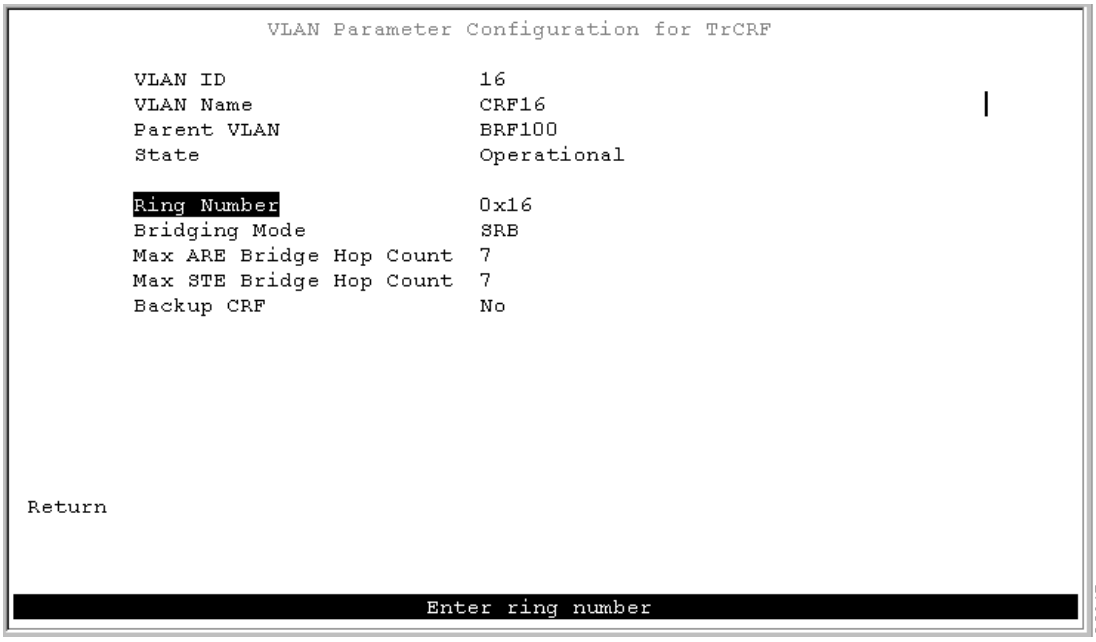

Step 8. Select **Return** to save your changes.

On switch B, repeat Step 1 through Step 7 and use the following values:

- VLAN ID of **16**.
- VLAN Name of **CRF16**.
- Parent VLAN of **BRF200**.
- Ring Number of **16**.

Assigning Ports to the TrCRF

Next, you must assign the ports of the TokenChannel to the TrCRF. On switch A, complete the following steps:

Step 1. On the VLAN and VTP Configuration panel, select **Local VLAN Port Configuration**. The Local VLAN Port Configuration panel is displayed.

Step 2. On the Local VLAN Port Configuration panel, select **Change**.

Step 3. At the prompt enter port number **18**.

Step 4. Select **CRF16** from the list of possible TrCRFs. To select the TrCRF, use your arrow keys to highlight the desired TrCRF, press the space bar to select it, and press **Enter** to implement your change.

Step 5. Repeat Step 2 through Step 4 for ports 19 and 20. The ports are now assigned to CRF16 (Figure 7-3).

Figure 7-3 Local VLAN Port Configuration Panel

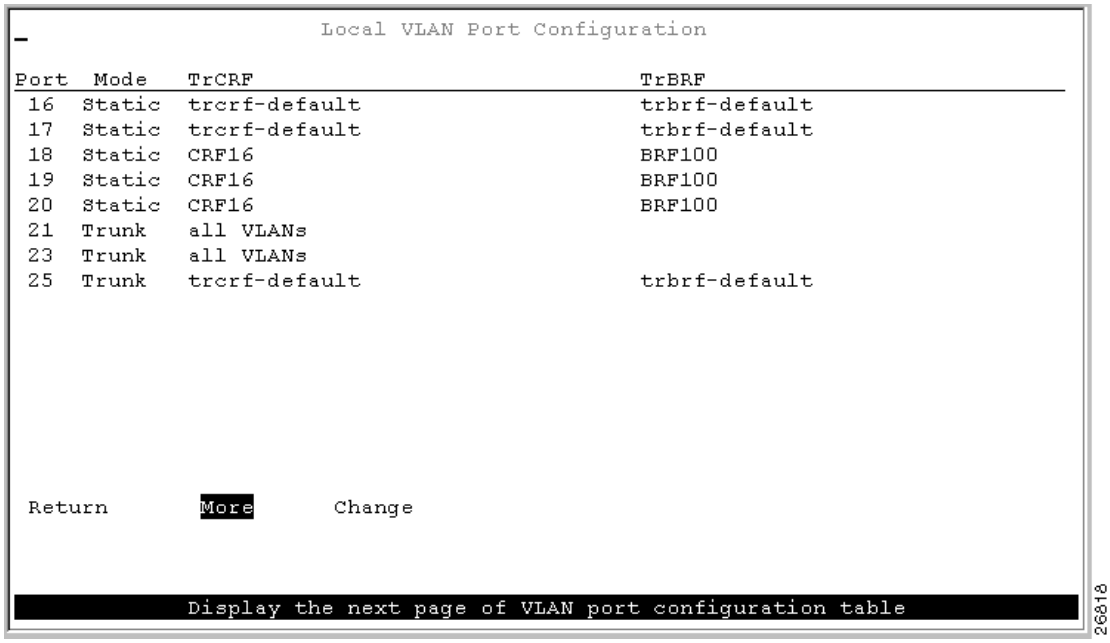

Step 6. Select **Return** to save your changes.

On switch B, repeat Step 1 through Step 6 and associate each of the three ports with CRF16.

Configuring the TokenChannel

Next you must configure the TokenChannel on both switches. On switch A and B, complete the following steps:

Step 1. On the Configuration panel, select **Channel Configuration**. The Channel Configuration panel is displayed.

Step 2. On the Channel Configuration panel, select **Channel Configuration**. The Channel Configuration panel is displayed.

Step 3. On the Channel Configuration panel, select **Add Entry**.

Step 4. At the prompt, enter ports **18 19 20** (separated by spaces). The ports will be assigned to the first available TokenChannel (Figure 7-4).

Figure 7-4 TokenChannel Configuration Panel

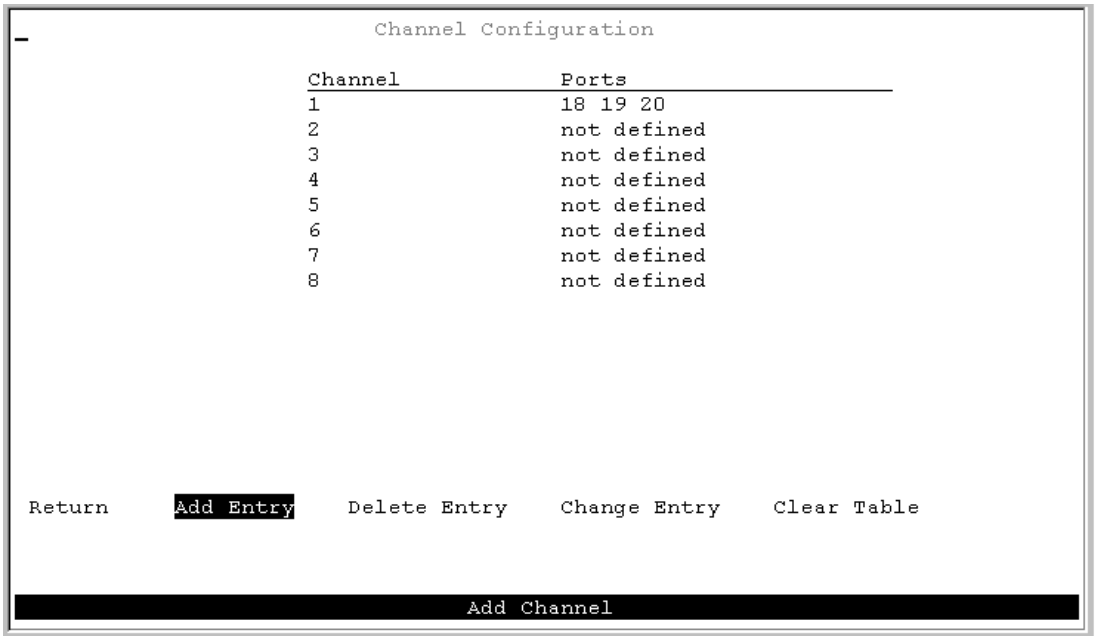

Step 5. Select **Return** to save your changes.

Attaching the Cables

Finally, you must attach the cables to the ports. Using straight-through cables with RJ-45 connectors, attach one end of one cable to port 18 on switch A. Attach the other end of the cable to port 18 on Switch B. Repeat this process to connect port 19 on switch A and B and port 20 on switch A and B.

Note: When you physically connect the linked ports, make sure that the ports with the lowest port numbers are connected. For example, if a TokenChannel links ports 3, 6, and 7 of one Catalyst 3900 and ports 2, 4, and 5 of another Catalyst 3900, the ports must be connected to each other in the following manner: port 3 to port 2, port 6 to port 4, and port 7 to port 5.

Resulting Network Configuration

You now have a network similar to the initial network except the two switches on the third floor can now exchange data with greater throughput over the three ports of their TokenChannel (Figure 7-5).

Figure 7-5 Final Network Configuration

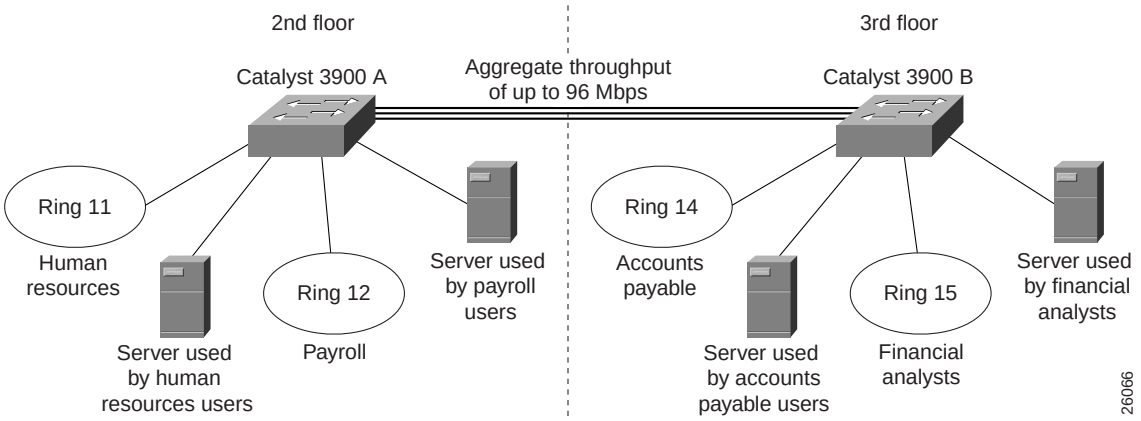

Figure 7-6 shows the final network configuration from a conceptual standpoint.

Figure 7-6 Conceptual Final Network

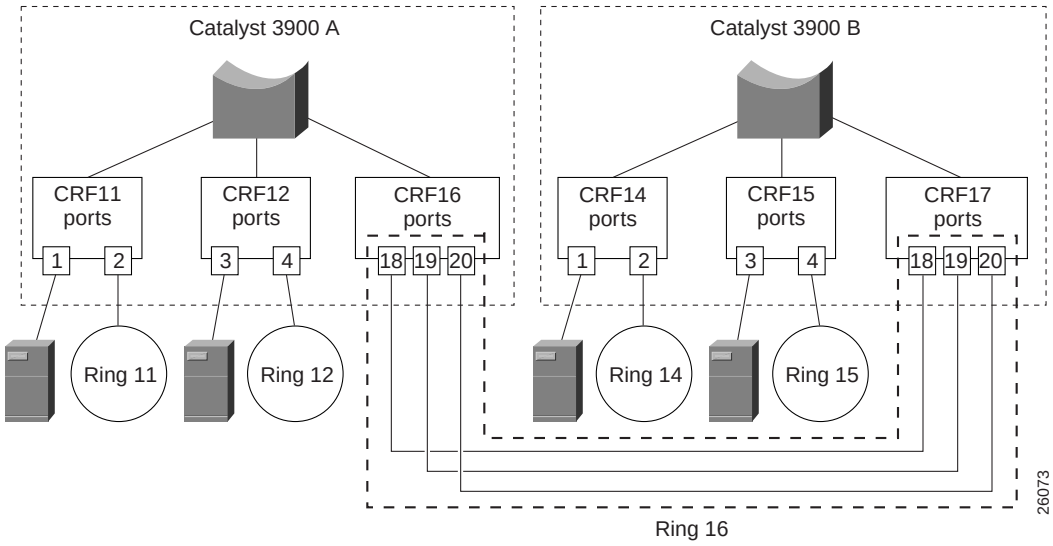

Tips

When configuring your TokenChannels, consider the following:

- If the TokenChannel will forward a combination of SRB and SRT traffic, the bridge mode of the TrCRF containing the TokenChannel ports should be set to SRT. If only SRB is needed, the bridge mode of the TrCRF containing the TokenChannel ports can be set to SRB.
- If SRB is not needed, the TokenChannel ports, the ports that are connected to the end user rings, and the ports that are connected to the servers can be placed in the same TrCRF on each switch.

The ports contained in a TokenChannel should be configured with an Operation Mode of auto. By default, the ports of a TokenChannel will be set to FDX operation. However, if you must configure the ports to anything other than auto, be sure to use FDX station for one port and FDX port for the other port in each port pair.

Troubleshooting

If you have difficulty configuring a TokenChannel, use the Local VLAN Port Configuration panel to verify that all the ports that you intend to include in your TokenChannel are in the same TrCRF. Because all the ports of a TokenChannel must be in the same TrCRF, the Catalyst 3900 will not allow you to define a TokenChannel that contains ports from different TrCRFs. Likewise, after you have defined a TokenChannel, if you assign a port in the TokenChannel to a different TrCRF, all ports in the TokenChannel are assigned to the new TrCRF.

To verify the current state of a channel, select **Channel Configuration** on the Configuration panel, and then select **Current Channel Information** on the Channel Configuration panel. The Current Channel Information panel is displayed. This panel lists the ISL Channels and TokenChannels that are currently defined and their state.

If a channel state is down:

- Make sure the cables are connected properly.
- Verify that all ports of the channel are enabled by looking at the Port Configuration panel.
- If the channel is a TokenChannel, verify that both of the ports in each port pair (port 18 on switch A and port 18 on switch B) are operating in auto mode. If the channel is an ISL Channel, verify that both ports in each ISL Channel pair are operating in FDX mode by looking at the Port Configuration panel.

Connecting Catalyst Token Ring Switches via an ATM Backbone

The Catalyst 3900 and the Catalyst 5000 series provide ATM modules that allow you to connect the switches to an ATM backbone. This chapter provides an example of connecting a Catalyst 3900 and a Catalyst 5000 Token Ring module via an ATM backbone.

This chapter provides the following information:

- Initial Network Configuration
- Before Beginning
- Configuration Steps
- Resulting Network Configuration

Initial Network Configuration

In your company, you have installed a Catalyst 3900 on the fourth floor and a Catalyst 5000 with a Token Ring module in slot 3 on the sixth floor.

There are servers attached to each switch. The users on these floors need to be able to access the servers on both floors. Because your company is in the medical industry (providing telephone support to emergency medical teams), response time is critical. You have decided to utilize the improved performance that ATM offers and to connect your Catalyst Token Ring switches though an ATM backbone.

Figure 8-1 illustrates the initial network configuration. The Catalyst 3900 on the fourth floor contains two TrCRFs (assigned ring numbers 108 and 109) that are joined by a TrBRF that you have named BRF300. The Catalyst 5000 on the sixth floor contains two TrCRFs (assigned ring numbers 110 and 111) that are joined by a TrBRF that you have named BRF400.

Figure 8-1 Initial Network Configuration

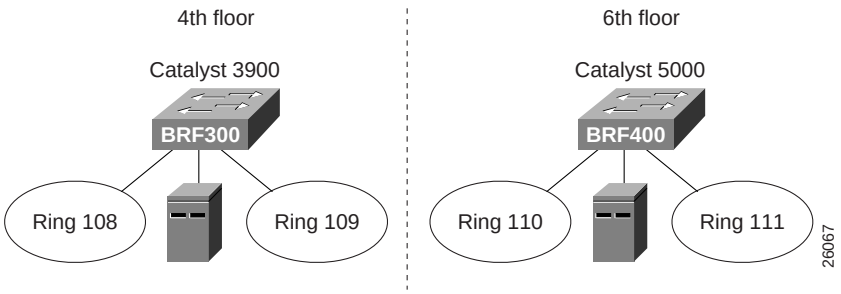

Before Beginning

For your ATM backbone, you have installed a Cisco ATM LS1010 switch. You will be using a Catalyst 5000 to provide the LECS, LES, and BUS. For connectivity between the switches, you have installed an Catalyst 3900 ATM module on the Catalyst 3900 and a Catalyst 5000 series ATM module in slot 4 of the Catalyst 5000. In addition, you downloaded the Token Ring LANE support for the Catalyst 5000 from CCO.

When you join the two switches via the ATM backbone, you need to create an ELAN on each switch. The ELAN is essentially a new TrCRF. You have decided to use 112 for the ELAN name as well as the VLAN ID.

Configuration Steps

To create an ELAN between the two switches, you must configure the ATM module and define an ELAN on both switches.

Configuring the Catalyst 3900

On the Catalyst 3900, you must configure a TrCRF and an LEC. Also, you must assign the port to the appropriate TrCRF.

Configuring the TrCRF

To create an ELAN that will span both switches, you must create a new TrCRF (which will also be defined on the Catalyst 5000) and associate it with the TrBRF.

To define the TrCRF, complete the following tasks:

Step 1. On the Catalyst 3900 Main Menu, select **Configuration**. The Configuration panel is displayed.

Step 2. On the Configuration panel, select **VLAN and VTP Configuration**. The VLAN and VTP Configuration panel is displayed.

Step 3. On the VLAN and VTP Configuration panel, select **VTP VLAN Configuration**. The VTP VLAN Configuration is displayed.

Step 4. On the VTP VLAN Configuration panel, select **Add**.

Step 5. At the prompt, enter a VLAN ID of **112**.

Step 6. At the prompt, select **TrCRF**. The VLAN Parameter Configuration for TrCRF panel (Figure 8-2) is displayed.

Step 7. On the VLAN Parameter Configuration for TrCRF panel, specify:

- VLAN Name of **CRF112**.
- Parent VLAN of **BRF300**.
- Ring Number of **112**.

Figure 8-2 VLAN Parameter Configuration for TrCRF Panel

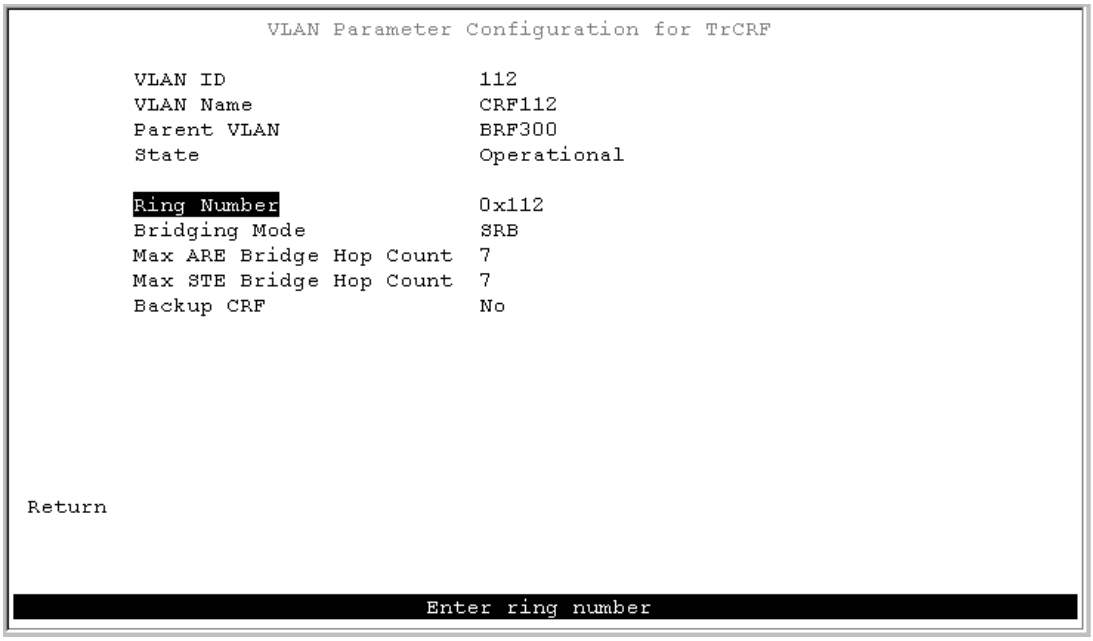

Step 8. Select **Return** to save your changes.

Assigning the ATM Port to the TrCRF

Next, you must assign the ATM port to the new TrCRF. Before you can do that, the ATM module must be inserted into one of the expansion slots on the Catalyst 3900. If you insert the module into the left slot, the ATM port is assigned a port number of 21. If you insert the module into the right slot, the ATM port is assigned a port number of 25.

To assign the ATM port to TrCRF 112, complete the following tasks:

Step 1. On the Configuration panel, select **VLAN and VTP Configuration**.The VLAN and VTP Configuration panel is displayed.

Step 2. On the VLAN and VTP Configuration panel, select **Local VLAN Port Configuration**. The Local VLAN Port Configuration panel is displayed.

Step 3. On the Local VLAN Port Configuration panel, select **Change**.

Step 4. At the prompt, enter port number **25**.

Step 5. Select **CRF112** from the list of possible TrCRFs. To select the TrCRF, use your arrow keys to highlight the desired TrCRF, press the space bar to select it, and press **Enter** to implement your change (Figure 8-3).

Figure 8-3 Local VLAN Port Configuration Panel

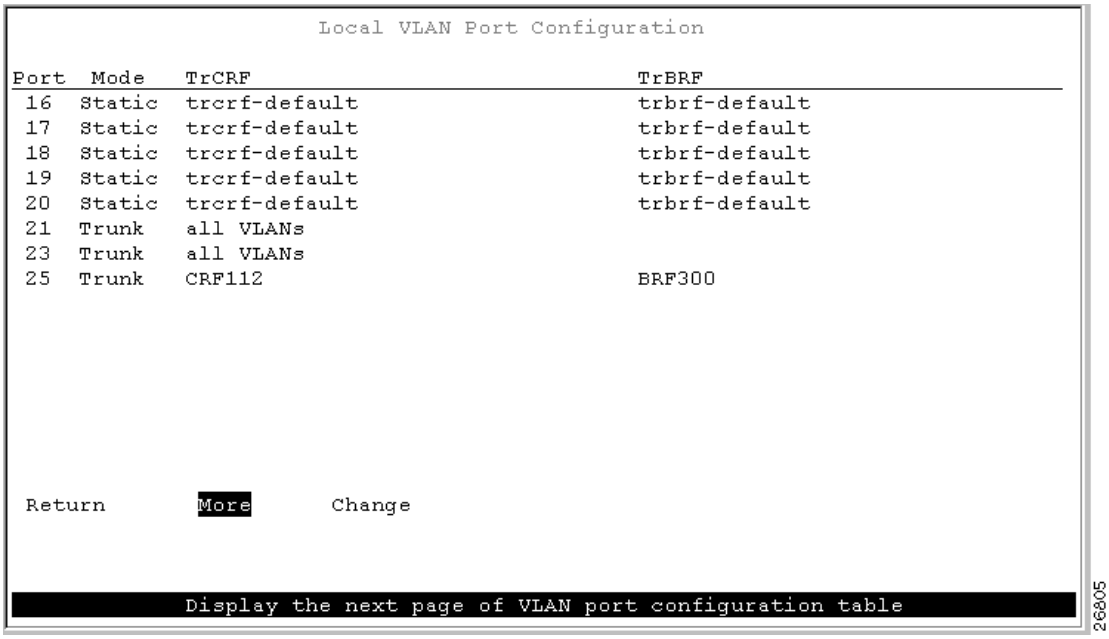

Step 6. Select **Return** to save your changes.

When you assign the ATM port to a TrCRF, an LEC is automatically created with an ELAN name that is the same as the VLAN name of the TrCRF.

Configuring the LEC

Because the LEC is automatically created and given an ELAN name, you do not need to configure anything for the LEC.

To view the configuration of the LEC, complete the following steps:

Step 1. On the Configuration panel, select **Port Configuration** and specify the appropriate port number. Because you have installed the ATM module in the right slot, the port number is 25. The ATM Configuration panel is displayed.

Step 2. On the ATM Configuration panel, select **ATM LEC Setup**. The ATM LEC Setup panel is displayed.

Step 3. On the ATM LEC Setup panel, select **LANE Parameters**. The LANE Parameters panel (Figure 8-4) is displayed.

Figure 8-4 LANE Parameters Panel

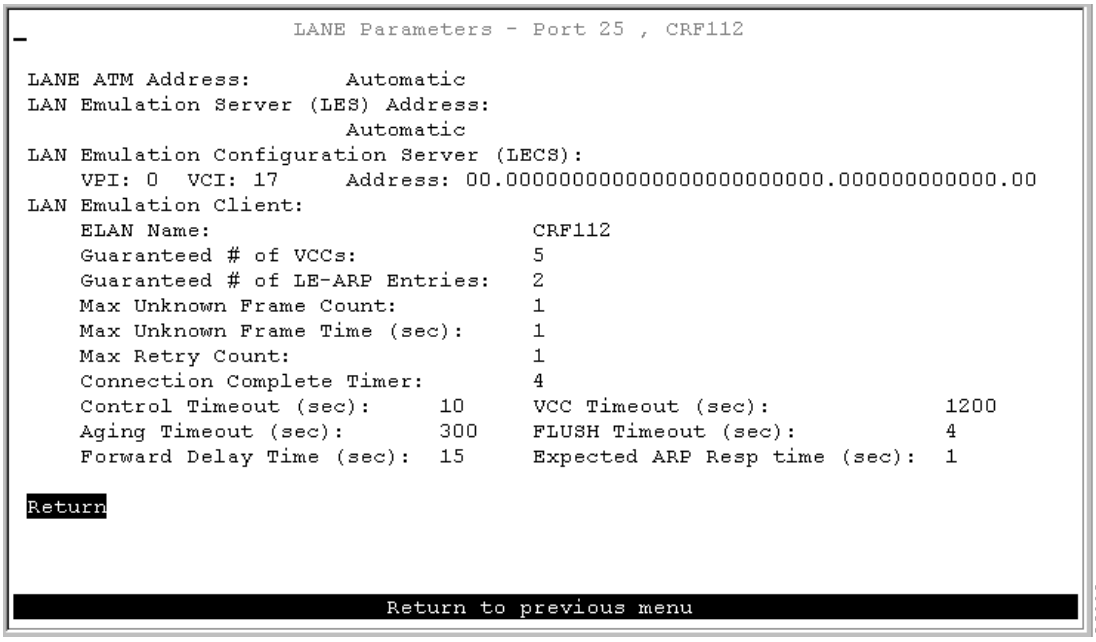

Step 4. Select **Return** to save your changes.

Configuring the Catalyst 5000

On the Catalyst 5000, you must configure a TrCRF and an LEC. Because the Catalyst 5000 ATM module provides a trunk port, you do not need to assign the port to the appropriate TrCRF. Trunk ports are automatically associated with LECs as they are configured.

You have decided to use the Catalyst 5000 for the LES and BUS. Therefore, you will also need to configure the ATM module for the LES/BUS support.

When configuring LECs on a Catalyst 5000, remember the following:

- The VLAN name must match the ELAN name.
- The ring number must match the local segment ID.

The **set vlan** command assumes that any ring number entered is in hexadecimal. Therefore 0x12 or 12 will be stored as the hexadecimal value 0x12. The **name** *elan_name* **local-seg-id** *segment_number* command assumes that any value entered for the local-seg-id is in decimal unless it is entered explicitly in hexadecimal. For example, to define a TrCRF with a ring number of 12, you could enter:

```
set vlan 12 name crf12 type trcrf ring 12 parent 100
```

or

```
set vlan 12 name crf12 type trcrf ring 0x12 parent 100
```

When defining a corresponding LEC, you could enter:

```
name crf12 local-seg-id 0x12
```

or

```
name crf12 local-seg-id 18
```

as 18 is the decimal equivalent of 0x12.

Configuring the TrCRF

To define the TrCRF, complete the following steps:

Step 1. At the enable prompt, enter **set vlan 112 name crf112 type trcrf ring 112 parent 400 mode srb**.

Step 2. To verify the configuration of the new VLAN, enter **show vlan**.

The output (Figure 8-5) indicates that crf112 has been added. It also shows that brf400 is the parent of the VLAN with the ID of 112.

Figure 8-5 Output of show vlan Command

```
VLAN Name                             Status    Mod/Ports, Vlans
---- -------------------------------- --------- ---------------------------
1    default                          active    1/1-2
                                                2/1-48
110  crf110                           active    3/1-4
111  crf111                           active    3/5-8
112  crf112                           active
400  brf400                           active    110,111,112
1002 fddi-default                     active
1003 trcrf-default                    active    3/9-16
1004 fddinet-default                  active
1005 trbrf-default                    active    1003

VLAN Type  SAID       MTU   Parent RingNo BrdgNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ ------ ---- -------- ------ ------
1    enet  100001     1500  -      -      -      -    -        0      0
110  trcrf 100110     4472  300    0x110  -      -    srb      0      0
111  trcrf 100110     4472  300    0x111  -      -    srb      0      0
112  trcrf 100110     4472  300    0x112  -      -    srb      0      0
300  trbrf 100100     4472  -      -      0x3    ibm  -        0      0
400  trbrf 100100     4472  -      -      0x4    ibm  -        0      0
1002 fddi  101002     1500  -      0x0    -      -    -        0      0
1003 trcrf 101003     4472  1005   0xccc  -      -    srb      0      0
1004 fdnet 101004     1500  -      -      0x0    ieee -        0      0
1005 trbrf 101005     4472  -      -      0xf    ibm  -        0      0
```

Configuring the LES, BUS, and LEC

To configure the LEC, BUS, and LEC, complete the following steps:

Step 1. Set up the prefix of the ATM Network Service Access Point (NSAP) address for the switch.

Note: The LightStream 1010 ATM switch provides a default prefix.

Step 2. Start a session to the ATM module that is in slot 4 by entering the **session 4** command. You see the following display:

```
Console> session 4
Trying ATM-4...
Connected to ATM-4.
Escape character is '^]'.
ATM>
```

Step 3. Obtain the addresses of the LES and LES/BUS for later use by entering the **enable** command (to enable configuration mode) and the **show lane default-atm-addresses** command at the ATM prompt. You see the following display:

```
ATM> enable
ATM#
ATM# show lane default-atm-addresses interface atm0

interface ATM0:
LANE Client:        47.009181000000000061705b7701.00400BFF0010.**
LANE Server:        47.009181000000000061705b7701.00400BFF0011.**
LANE Bus:           47.009181000000000061705b7701.00400BFF0012.**
LANE Config Server: 47.009181000000000061705b7701.00400BFF0013.00
ATM#
```

Note: The two asterisks (**) represent the subinterface number byte in hexadecimal.

Step 4. Using the LECS address obtained in Step 3, set the address of the default LECS in the LightStream 1010 switch by entering the **configure terminal** and **atm lecs-address-default** commands on the console of the LightStream 1010 switch. You see the following display:

```
Switch> enable
Switch#
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# atm lecs-address-default
47.009181000000000061705b7701.00400BFF0013.00 1
Switch(config)# end
Switch#
```

The commands shown in this step configure the address of the LECS in the switch. The LECS ATM NSAP address is **47.009181000000000061705b7701.00400BFF0013.00**. The sequence number of this LECS address, which is **1**, means it is the first LECS in this switch.

Step 5. Save the configuration to NVRAM by entering the **write memory** command at the prompt.

Step 6. Start up an LES/BUS pair on the Catalyst 5000 series switch by entering the **interface atm0** and the **lane server-bus tokenring** commands in global configuration mode.

Enter the following commands:

config terminal

interface atm0

lane server-bus tokenring crf112

end

The commands shown in this step start an LES/BUS pair and assign the ATM 0 interface to crf112. The ELAN name is **crf112**, and the interface on which this LES/BUS pair is configured is **atm0**. The ELAN name must be the same as the VLAN name assigned to the TrCRF.

Step 7. Save the configuration in NVRAM by entering the **write memory** command at the prompt.

Step 8. Set up the LECS database on the Catalyst 5000 series switch.

Enter the LES address obtained in Step 3 and replace the ** with the subinterface number of the interface in which the LES/BUS is to be configured. In this example, that number is 00. Enter the **config terminal** command, the **lane database** *database_name* command, the **name** *elan_name* **server-atm-address** *atm_address* command, the **name** *elan_name* **local-seg-id** *segment_number* command, and the **default-name** *elan_name* command at the ATM prompt. You see the following display:

```
ATM# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ATM(config)# lane database test
ATM(lane-config-database)# name crf112 server-atm-address
                              47.009181000000061705b7701.00400BFF0011.00
ATM (lane-config-database) name crf112 local-seg-id 0x112
ATM(lane-config-database)# default-name crf112
ATM(lane-config-database)# exit
ATM#
```

The commands shown in this step create the LECS database. The database name is **test**. The ELAN name is **crf112**. The ELAN segment number is **112**. The LES ATM NSAP address is **47.009181000000061705b7701.00400BFF0011.00**.

Step 9. Save the configuration in NVRAM by entering the **write memory** command at the prompt.

Step 10. Start and bind the LECS on the Catalyst 5000 series switch by entering the **config terminal** command, the **interface atm0** command, the **lane config database** *database_name* command, and the **lane config auto-config-atm-address** command at the ATM prompt. You see the following display:

```
ATM# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ATM(config)# interface atm0
ATM(config-if)# lane config database test
ATM(config-if)# lane config auto-config-atm-address
ATM(config-if)# end
ATM#
```

The commands shown in this step start the LECS. The database to use is **test**. The interface on which the LECS is configured is **atm0**.

Step 11. Save the configuration in NVRAM by entering the **write memory** command at the prompt.

Step 12. Start the LEC on the Catalyst 5000 series Switch by entering the **config terminal** command, the **interface atm0.1** command and the **lane client tokenring 112 crf112** command in configuration mode. The interface on which the LEC is configured is **atm0.1**. The ELAN name is **crf112**, and it is configured to emulate Token Ring. You see this display:

```
ATM# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
ATM(config)# interface atm0.1
ATM(config-subif)# lane client tokenring 112 crf112
ATM(config-subif)# end
ATM#
```

Step 13. Save the configuration in NVRAM by entering the **write memory** command at the prompt.

Attaching the Cables

Finally, you must attach the cables to the ATM ports. Using 62.5/125-micron fiber-optic cables with subscriber connectors, attach one end of one cable to the ATM ports on the Catalyst 3900. Attach the other end of the cable a port on the LightStream 1010. Using the appropriate cabling for the Catalyst 5000 ATM module, repeat this process to connect the ATM ports on the Catalyst 5000 to the LightStream 1010.

Resulting Network Configuration

Because you have bridged the rings across the high-speed ATM network, your users on the fourth and sixth floor are now joined and will be able to access to resources on the different floors with improved response time. The resulting configuration is shown in Figure 8-6.

Figure 8-6 Resulting Network Configuration

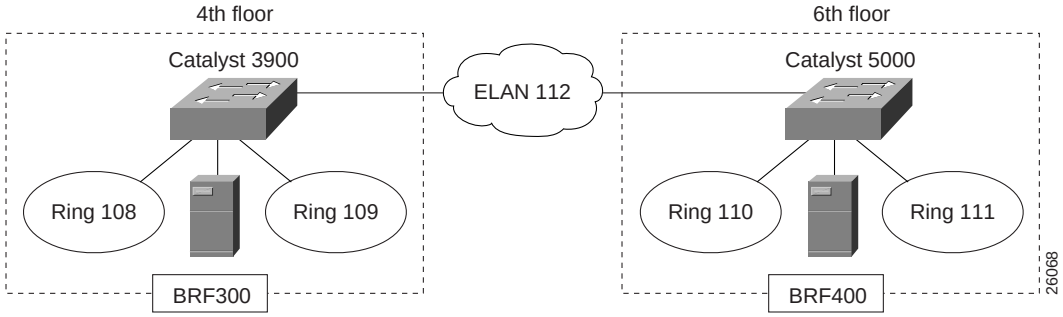

Interconnecting Catalyst Token Ring Switches Using ISL

The Catalyst 3900 and Catalyst 5000 support Cisco's ISL technology, which allows you to interconnect your switches over a 100 Mbps connection.

This section contains three different scenarios; one for ISL between two Catalyst 3900s, one for ISL between two Catalyst 5000s with Token Ring modules, and one for ISL between a Catalyst 3900 and a Catalyst 5000 with a Token Ring module.

This chapter provides the following information:

- Initial Network Configuration
- Before Beginning
- Configuration Steps
- Resulting Network Configuration
- Tips
- Troubleshooting

Initial Network Configuration

In your company offices, you have two Catalyst 3900s on the first floor and two Catalyst 5000 with Token Ring modules on the second floor. One of the Catalyst 5000s is connected to a router though a Fast Ethernet module. The switches are currently interconnected by normal Token Ring connections. To increase the throughput of the connections between the switches, reduce the number of hops between the endstations and the router, and thereby improve network performance, you have decided to interconnect the switches with ISL.

Note: For more information about ISL and how it works, see the "Using ISL" section of the "Interconnecting Switches" chapter.

Figure 9-1 illustrates the initial network configuration. In this configuration, each connection between a switch constitutes another hop in the network. So an endstation attached to Ring 1 must traverse four bridges (or hops) before it reaches the router.

Figure 9-1 Initial Network Configuration

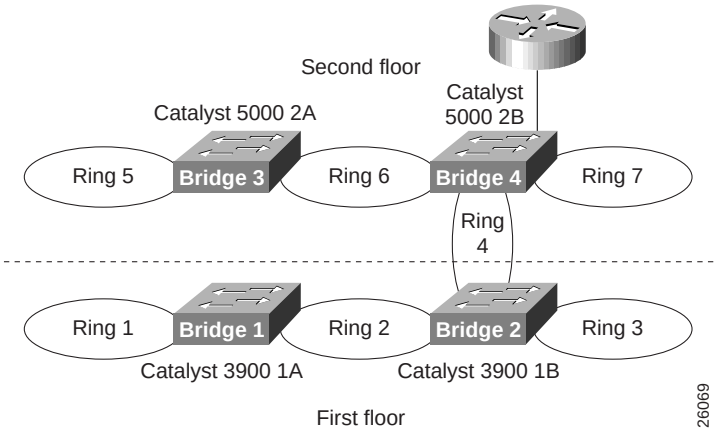

Before Beginning

You have added a Token Ring ISL module to each of the Catalyst 3900s. On the Catalyst 5000, ISL is implemented in software and runs on a Fast Ethernet module. One of the Catalyst 5000s has a Fast Ethernet module and you have added a Fast Ethernet module to the other Catalyst 5000.

With switches that are interconnected via ISL, you can use VTP to propagate VLAN information within the management domain. As explained earlier in the "VLAN Trunking Protocol" section of the "Token Ring VLANs and Related Protocols" chapter, a switch can operate in transparent, server, or client mode. For the purposes of this scenario, you have decided to configure one of the Catalyst 3900s to act as a VTP server and to configure the other switches to operate in client mode. Once your switches are joined in a common VTP management domain (which you have decided to call Domain A) the VLAN definitions that are configured on the server switch will be propagated to the client switches. Because you have existing rings in the network that you want to preserve, you must duplicate the definitions of those rings on the switch that is acting as your VTP server. You will be replacing rings 2, 4, and 6 with ISL connections, but will want to preserve the user rings, which are rings 1, 3, 5 and 7.

Currently, the rings are defined as follows:

Ring Number	VLAN ID	VLAN Name	Parent VLAN
1	101	Ring 1	Bridge 1
3	103	Ring 3	Bridge 2
5	205	Ring 5	Bridge 3
7	207	Ring 7	Bridge 4

Also, because ISL allows us to expand the bridge (TrBRF) across switches over the ISL connection, you will be using a single bridge definition on each switch. You have decided to use bridge 1 for your network and eliminate the definitions of bridges 2 through 4.

You have installed an ISL module in slot 1 of the Catalyst 3900 1A and in slot 1 of Catalyst 3900 1B. You have installed a Fast Ethernet module in slot 2 of Catalyst 5000 2A. You already have a Fast Ethernet module installed in slot 2 of Catalyst 5000 2B, which is already connected to the router.

For the physical connections, you have decided to interconnect the devices as follows:

Switch	Port	Connected to Switch	Port
Catalyst 3900 1A	21	Catalyst 3900 1B	21
Catalyst 3900 1B	23	Catalyst 5000 2B	2/12
Catalyst 5000 2A	2/1	Catalyst 5000 2B	2/11

Configuration Steps

To connect the switches using ISL, you will need to configure an ISL connection between the two Catalyst 3900s, between the two Catalyst 5000s, and between one of the Catalyst 5000s and one of the Catalyst 3900s. You will need to enable the use of VTP in your network, which requires us to assign all the switches to a VTP management domain and configure the switches to act as a server or client.

Configuring the ISL Connections on the Catalyst 3900s

There is little or no configuration that you need to perform on the actual ISL module. By default, the module is enabled and operates in trunking mode. The ports operate at 100 Mbps and in FDX mode. The only possible transmission mode is store-and-forward. However, you will need to configure a VTP management domain, configure one of the Catalyst 3900s to act as the VTP server, duplicate the definitions of all of the existing TrCRFs on the Catalyst 3900 that is acting as the VTP server, and configure the other to operate in client mode.

Configuring one Catalyst 3900 to be a VTP Server

To configure Catalyst 3900 1A to be the VTP server, complete the following steps on each switch:

Step 1. On the Catalyst 3900 Main Menu, select **Configuration**.

Step 2. On the Configuration menu, select **VLAN and VTP Configuration**.

Step 3. On the VLAN and VTP Configuration menu, select **VTP Administrative Configuration**. The VTP Administrative Configuration panel (Figure 9-2) is displayed.

Step 4. On the VTP Administrative Configuration panel, specify the following:

- A Domain Name of **Domain A**.
- A Local Mode of **Server**.

Figure 9-2 VTP Administrative Configuration Panel

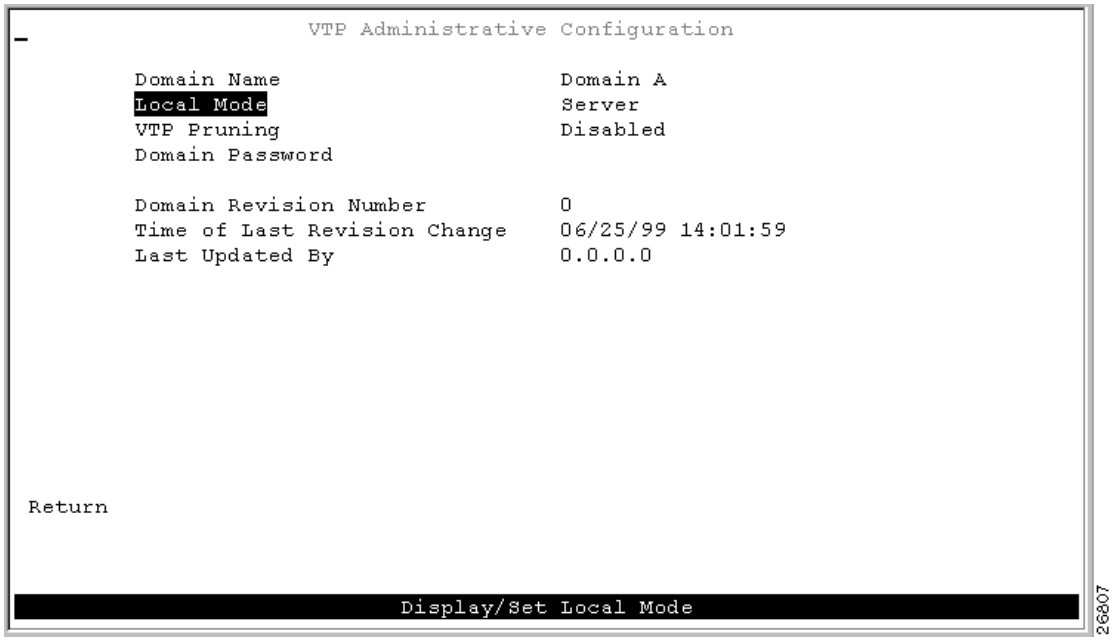

Step 5. Select **Return** to save your changes.

Note: If any of the TrCRFs on the Catalyst 3900 that contain ports are configured to learn their ring numbers (auto), you cannot change the Local Mode to Server or Client. In this case, you would first need to alter the definitions of the existing TrCRFs and specify a ring number for each.

Configuring the TrCRFs on the Server Switch

Because the configuration of a switch that is operating in VTP server mode is propagated to each of the switches in the VTP management domain that are operating in client mode, you must duplicate the configuration of the existing rings on the switch that is acting as the VTP server before propagating the VLAN information throughout the network. As mentioned before, you have decided to use Bridge 1 (TrBRF 100) as the single bridge in your ISL network. Because the Catalyst 3900 1A contains the definition of Bridge 1, you do not need to configure a new bridge. But as you duplicate the configurations of the other TrCRFs, you will need to change their parent to be Bridge 1.

To configure the TrCRFs, complete the following tasks:

Step 1. On the VTP and VLAN Configuration panel, select **VTP VLAN Configuration**.

Step 2. On the VTP VLAN Configuration panel, select **Add**, specify a VLAN ID of **103**, and then select **TrCRF**. The VLAN Parameter Configuration for TrCRF panel (Figure 9-3) is displayed.

Step 3. Specify the following:

- VLAN Name of **Ring 3**.
- Parent VLAN of **Bridge 1**.
- Ring Number of **3**.

Figure 9-3 VLAN Parameter Configuration for TrCRF Panel

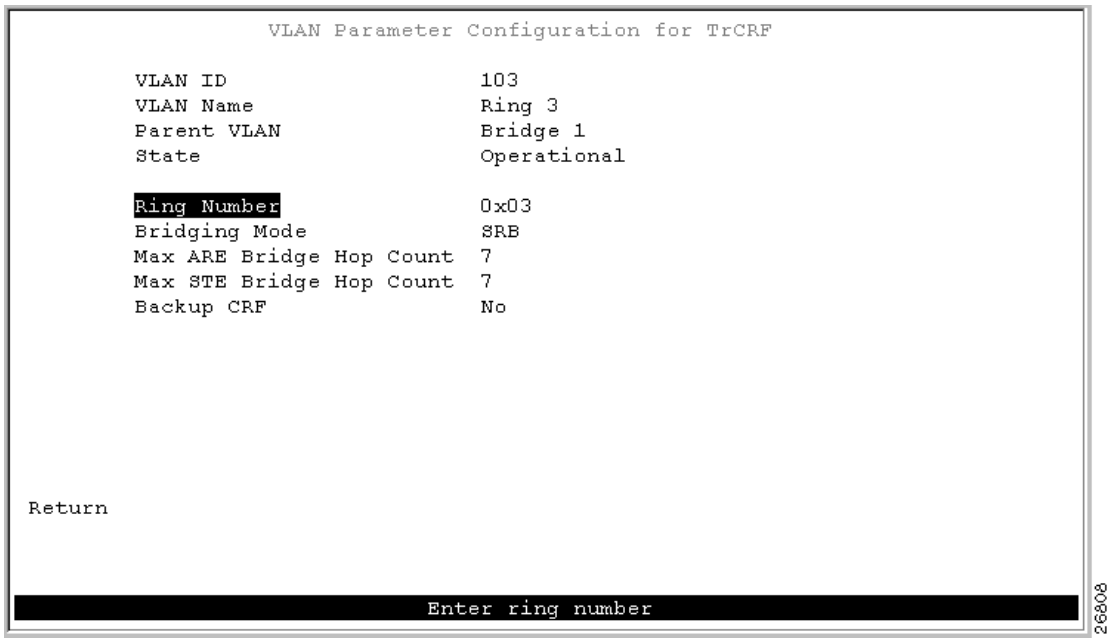

Step 4. Select **Return** to save your changes.

Step 5. On the VTP VLAN Configuration panel, select **Add**, specify a VLAN ID of **105**, and then select **TrCRF**. The VLAN Parameter Configuration for TrCRF panel is displayed.

Step 6. Specify the following:

- VLAN Name of **Ring 5**.
- Parent VLAN of **Bridge 1**.
- Ring Number of **5**.

Step 7. Select **Return** to save your changes.

Step 8. On the VTP VLAN Configuration panel, select **Add**, specify a VLAN ID of **107**, and then select **TrCRF**. The VLAN Parameter Configuration for TrCRF panel is displayed.

Step 9. Specify the following:

- VLAN Name of **Ring 7**.
- Parent VLAN of **Bridge 1**.
- Ring Number of **7**.

Step 10. Select **Return** to save your changes.

As shown in Figure 9-4, the TrCRFs have been added with Bridge 1 as the parent of each. The Local State indicates whether a VLAN (TrCRF or TrBRF) is used on the local switch. Possible values are not-local, preferred, and automatic.

- Not-local indicates that the VLAN is not designated for use on the local switch.
- Preferred indicates that the VLAN is designated for use on the local switch. A preferred VLAN is guaranteed access on the switch. Any TrCRF to which ports on the local switch are assigned is designated as preferred. In addition, you can designate other TrCRFs and TrBRFs as preferred using the Change_Local_State option. Up to 63 TrCRFs and 63 TrBRFs can be designated as preferred.
- Automatic indicates that the VLAN can be used on the local switch if access is available. An automatic VLAN is not guaranteed access on the switch. Automatic VLANs are given access as space is available (if less than 63 VLANs have been designated as preferred).

The new TrCRFs do not have any ports assigned to them on the local switch. Therefore, each has a local mode of Automatic.

Figure 9-4 VTP VLAN Configuration Panel

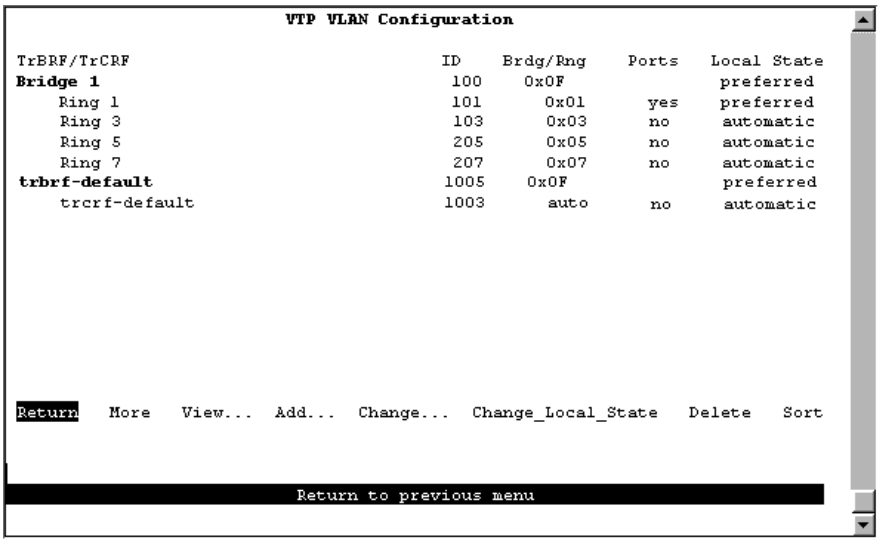

Note: Keep in mind that you are duplicating the definition of the TrCRFs for the purposes of distributing VLAN information through VTP. DRiP does not allow a TrCRF to contain ports on different switches. Therefore, because Rings 3, 5, and 7 contain ports on other switches, you cannot assign ports on the local switch to any of these TrCRFs. If you were to assign ports to any of these TrCRFs, DRiP would disable the related ports on both switches.

Configuring the Other Catalyst 3900 as a VTP Client

To configure Catalyst 3900 1B to participate in the VTP management domain and operate in client mode, complete the following steps on each switch:

Step 1. On the Catalyst 3900 Main Menu, select **Configuration**.

Step 2. On the Configuration menu, select **VLAN and VTP Configuration**.

Step 3. On the VLAN and VTP Configuration menu, select **VTP Administrative Configuration**. The VTP Administrative Configuration panel (Figure 9-5) is displayed.

Step 4. On the VTP Administrative Configuration panel, specify the following:

- A Domain Name of **Domain A**.
- A Local Mode of **Client**.

Figure 9-5 VTP Administrative Configuration Panel

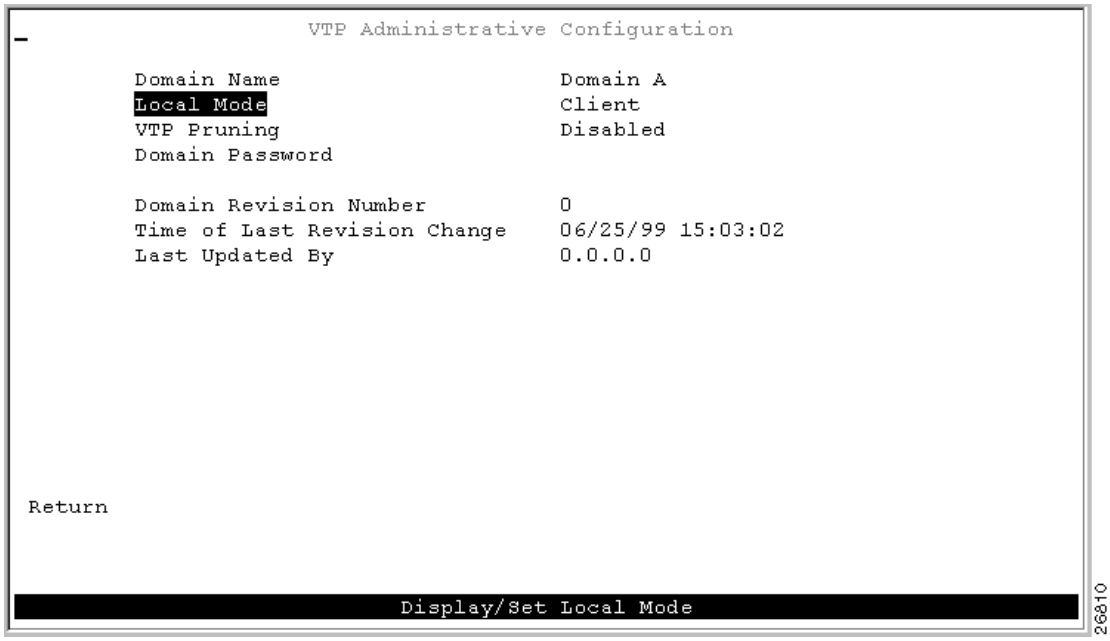

Step 5. Select **Return** to save your changes.

Once the cables are connected, within seconds Catalyst 3900 1B should receive VTP advertisements from Catalyst 3900 1A. Once these advertisements are received, the definition of Ring 3 will be overwritten with the new definition in which Bridge 1 is the parent of Ring 3.

Configuring the Catalyst 5000s

On the Catalyst 5000, ISL is implemented in software and can be run on a Fast Ethernet module. For this scenario, you will need to configure both of the Catalyst 5000s to operate in client mode (as one of the Catalyst 3900s is operating a the VTP server). You will also need to configure an ISL connection between one of the Catalyst 3900s and one of the Catalyst 5000s and between the two Catalyst 5000s.

Configuring the Catalyst 5000s to be VTP Clients

The Catalyst 3900 uses VTP V2. VTP V2 is not downwardly compatible with VTP V1. Therefore, you must configure the Catalyst 5000s to use VTP V2. You must also configure them to participate as clients in your VTP management domain (Domain A). To configure the Catalyst 5000s to be VTP clients, on Catalyst 5000 2A and 2B complete the following steps:

Step 1. At the Catalyst 5000 command prompt, enter **enable**.

Step 2. At the enable prompt, enter **set vtp domain Domain_A mode client v2 enable.**

Once the cables are connected, within seconds Catalyst 5000s should receive VTP advertisements from Catalyst 3900 1A. Once these advertisements are received, the definitions of Ring 5 and Ring 7 will be overwritten with the new definitions in which Bridge 1 is the parent of both.

Configuring the ISL Connections on the Catalyst 5000

To enable ISL connections on the Catalyst 5000s, you must configure the selected ports on the Fast Ethernet module to act as trunk ports. You must configure the port of Catalyst 5000 2B that connects to the Catalyst 3900 (2/12) to be a trunk port. Because Catalyst 5000 switches can detect that a connection to another Catalyst 5000 has been configured as a trunk, you can configure the other port (2/11) to act as a trunk and the port on Catalyst 5000 2A will detect this and automatically become a trunk port.

To configure ports to be trunk ports, complete the following steps:

Step 1. At the Catalyst 5000 command prompt, enter **enable**.

Step 2. At the enable prompt, enter **set trunk 2/11 on**.

Step 3. At the enable prompt, enter **set trunk 2/12 on**

Configuring an ISL Connection Between a Catalyst 5000 and a Catalyst 3900

The Catalyst 5000 Fast Ethernet module is designed to auto-sense the speed and mode of the attached device. The Catalyst 3900 Token Ring ISL module does not support auto-negotiation of speed or mode; it always runs at 100 Mbps in FDX mode. Therefore, when connecting a Catalyst 3900 to a Catalyst 5000 via ISL, you must configure the port of the Catalyst 5000 Fast Ethernet module to operate at the correct speed and mode.

To configure the speed and mode of the Fast Ethernet module on Catalyst 2B, complete the following steps:

Step 1. At the Catalyst 5000 command prompt, enter **enable**.

Step 2. At the enable prompt, enter **set port speed 2/12 100**.

Step 3. At the enable prompt, enter **set port duplex 2/12 full**.

Attaching the Cables

Attach the cables to the switches as described in the “Before Beginning” section. As stated before, once the cables are connected, within seconds client switches should receive VTP advertisements from Catalyst 3900 1A. Once these advertisements are received, the definitions of Ring 3, Ring 5, and Ring 7 will be overwritten with the new definitions in which Bridge 1 is the parent of all.

Configuring the STP

When you connect two Catalyst switches using ISL, it is important that the STP be used. If you do not use the STP, the ports will be placed in a blocked state and create a loop. The spanning-tree settings differ slightly between the Catalyst 3900 and the Catalyst 5000.

Spanning-Tree Protocol on the Catalyst 3900

The Catalyst 3900 ISL module supports the STP at both the TrCRF and the TrBRF level. The STP that is run on the ISL link depends on the type of TrCRF:

- With an undistributed TrCRF, such as in this scenario, the STP specified for the TrBRF is used.
- With a default TrCRF, the STP specified for the TrCRF is used.

By default, the Catalyst 3900 runs the IBM STP on the TrBRF. You can configure the TrCRF STP or configure the STP to be automatically determined by the specified bridging mode. TrCRFs with a bridge mode of SRB will run the IEEE STP and TrCRFs with a bridge mode of SRT will run the Cisco STP.

Spanning-Tree Protocol on the Catalyst 5000

By default, the Catalyst 5000 Token Ring switching module runs the IBM STP on the TrBRF. The STP run on the TrCRFs is determined by the specified bridging mode. TrCRFs with a bridge mode of SRB will run the IEEE STP and TrCRFs with a bridge mode of SRT will run the Cisco STP.

The Catalyst 5000 Token Ring switching module considers the combination of the IBM STP at the TrBRF and the bridge mode of SRT to be incompatible. As a result, if you had configured one of the TrCRFs with a bridge mode of SRT, the Catalyst 500 Token Ring switching module would automatically block the logical port of the TrCRF that is configured for SRT. You can use the **show spantree** command to view the state of the logical ports. The conflict can be addressed in either of the following ways:

- If you configure all the TrCRFs in the TrBRF to use a bridge mode of SRT, you can use the **set vlan** command to change the STP to IEEE. This will eliminate the conflict.
- If you need to have a mixture of TrCRFs that use both bridging modes, leave the STP as IBM, and use the **set spantree portstate** command to change the forwarding mode of the logical ports that are blocked.

Resulting Network Configuration

Figure 9-6 and Figure 9-7 illustrate the resulting network configuration.

Figure 9-6 Resulting Network Configuration (Physical)

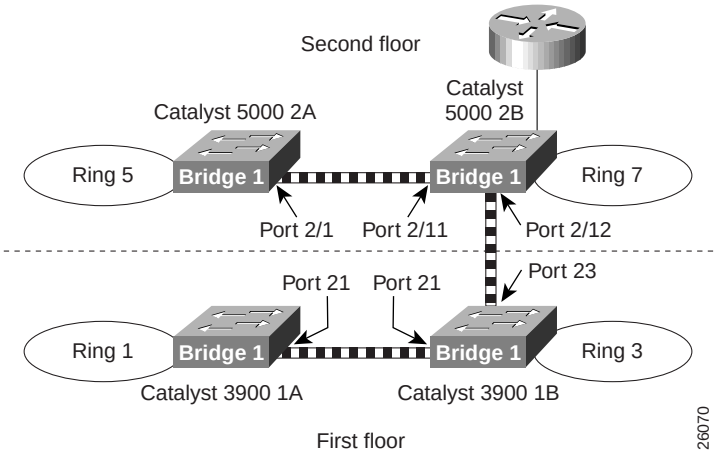

As you can see, rings 2, 4, and 6 have been replaced with the ISL connections. And each switch now contains a ring (TrCRF) that belongs to the same bridge (TrBRF). As a result, logically you now simply have four rings that are joined with the router by a single bridge.

Figure 9-7 Resulting Network Configuration (Logical)

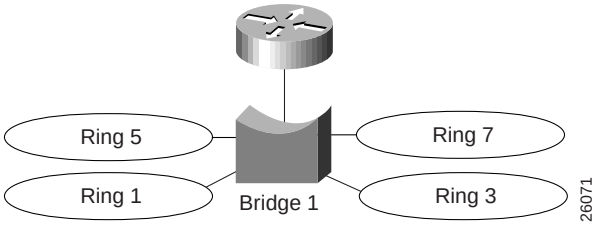

Tips

VTP advertisements propagate only certain configuration information. Other configuration settings must be made on the individual switches. The configuration settings that must be made at the switch include the following:

- Port assignments—Ports cannot be assigned to TrCRFs that are not local. Therefore, if you decide to add other ports to the TrCRFs, you must do so at the individual switch.
- IP information—You cannot configure IP information for TrBRFs that are not-local. Therefore, if you decide to change the IP configuration of any of the switches, you must make the changes at the individual switch.
- Spanning-tree parameters—You cannot configure spanning-tree parameters for TrBRfs that are not-local and, therefore, cannot configure spanning-tree parameters for TrCRFs that are not local.

Troubleshooting

If the VTP information is not being propagated to each of the switches, make sure all switches are configured to be part of the VTP management domain called Domain A. Also, make sure that the Catalyst 5000 switches are running VTP Version 2.

If the Catalyst 3900 does not detect the ISL module when it is inserted (i.e. it does not show up on the listing of modules on the Module Information panel), make sure that you are running revision 3.0(1) or later of the Catalyst 3900 software. You can verify the revision level of the software by selecting Switch Configuration on the Configuration Menu.

If the Catalyst 3900 rejects your attempts to change the VTP mode from Transparent to Client or Server, make sure that you have configured the ring numbers for the associated TrCRFs. The Catalyst 3900 cannot be configured as a VTP client or server if the ring number is set to auto.

Configuring IP Routing between Token Ring VLANs on the Catalyst 5000 RSM

In this chapter, you are going to configure IP routing between Token Ring VLANs on the Catalyst 5000 series Route Switch module (RSM).

As you learned in the "Using a Switch for Ring Microsegmentation" chapter, microsegmenting a network into VLANs enables you to maximize bandwidth and performance in your network. VLANs control the size of broadcast domains and localize traffic. However, in the past, end stations belonging to one bridged domain (TrBRF) could not communicate with network devices in another bridged domain (TrBRF) without an additional piece of equipment—the router.

With the Catalyst 5000 series RSM, you can configure interVLAN communication between bridged domain (TrBRFs) that are configured on multiple switches that are connected via ISL.

The Catalyst 5000 series RSM runs Cisco IOS software and provides multilayer switching and interVLAN routing services between switched VLANs and ELANs. Utilizing the Catalyst 5000 series RSM in your network enables you to bypass the purchase and maintenance of additional equipment in your network and to alleviate the burden being placed on centralized routers that are currently being used for interVLAN communication in your network.

This chapter provides the following information:

- Initial Network Configuration
- Before Beginning
- Configuration Steps
- Resulting Network Configuration

Initial Network Configuration

Currently, you have two bridged domains (TrBRFs) between which IP routing has been configured via a centralized router. The router is a Cisco 7200 in which a 2FEISL port adapter is installed.

One bridged domain (TrBRF) is dedicated to your company's Engineering departments and is configured on a Catalyst 5000 series switch in which a Token Ring module (WS-X5030) is installed in slot 3 and an ISL module is installed in slot 2. There are two rings (TrCRFs) configure to support two engineering departments. Ring 30 supports the engineering department located on the first floor of your R&D building. Ports 2 and 3 of the Token Ring module are assigned to this ring. Ring 50 supports the engineering department located on the second floor of your R&D building. Ports 7 and 8 of the Token Ring module are assigned to Ring 50. The rings have been joined via a logical bridge (TrBRF). The VLAN ID for the Engineering TrBRF is 100.

The second bridged domain is dedicated to your company's Marketing departments and is configured on a Catalyst 3900 series switch in which an ISL module is installed. There are two rings (TrCRFs) configure to support two marketing departments. Ring 20 supports the Public Relations department. Token Ring ports 2 and 3 of the Catalyst 3900 are assigned to this ring. Ring 40 supports the Product Marketing department. Token Ring ports 7 and 8 of the Catalyst 3900 are assigned to Ring 40. The rings have also been joined via a logical bridge (TrBRF). The VLAN ID for the Marketing TrBRF is 200.

Both bridged domains (TrBRFs) are members of the Eng-Marketing VTP management domain. The Catalyst 3900, Catalyst 5000, and Cisco 7200 are connected via ISL.

Figure 10-1 illustrates the initial network configuration.

Figure 10-1 Initial Network Configuration

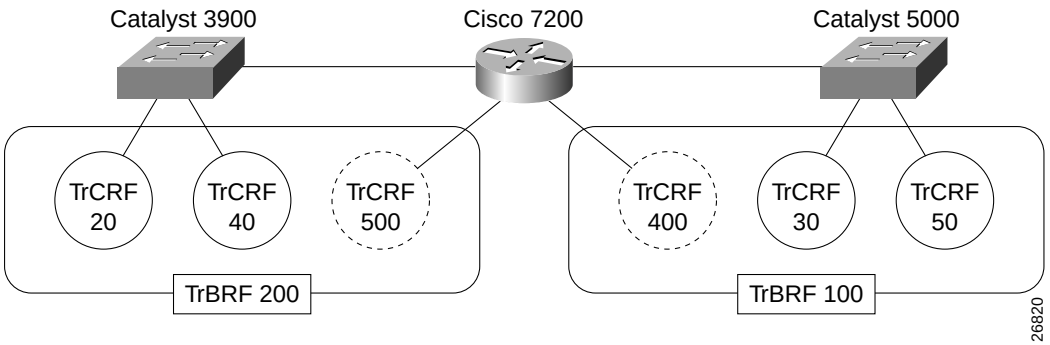

Because the burden of additional traffic that is being placed on the router has affected your network, you have purchased a Catalyst 5000 series RSM and installed it in slot 5 of your Catalyst 5000 switch. Via the RSM, you intend to configure IP routing between the Marketing and Engineering TrBRFs.

Before Beginning

Each TrBRF in a switched network requires an RSM interface and an IP address for IP routing. You have determined these values will be as follows:

VLAN Name	VLAN ID	TrBRF IP Address
TrBRF 100	100	172.122.30.1 255.255.255.0
TrBRF 200	200	172.122.40.1 255.255.255.0

TrBRF 100 and TrBRF 200 are source-route bridges. Therefore, you need to create a unique TrCRF on the RSM for each TrBRF. These TrCRFs are for the pseudo-ring on the RSM. The pseudo-ring is used to terminate RIFs. The RSM TrCRF for TrBRF 100 will be TrCRF 400, ring 100. The RSM TrCRF for TrBRF 200 will be TrCRF 500, ring 100.

Caution A unique TrCRF must be configure on the RSM for each TrBRF in your network. This unique TrCRF is required for the multiring functionality on the RSM. TrCRFs *cannot* be shared on the RSM because the network could be severely affected.

Because IP routing is currently configured between the Marketing and Engineering departments via a centralized router, the Token Ring VLAN configuration as it exists in the switch network should be ready. However, the following is a checklist of the switched network configuration that you use to verify that the configuration is correct:

- The VTP management domain on the Catalyst 3900 and the Catalyst 5000 has been set to **Eng-Marketing**.
- The VTP mode for both switches is **Server**.
- VTP V2 is **enabled**.
- On the Catalyst 5000 the following Token Ring VLAN (TrBRF and TrCRF) configuration exists:
  - TrCRF 30 is configured as ring 30, it has been associated with TrBRF 100, and it has been assigned ports 2 and 3 of the Token Ring module located in slot 3.
  - TrCRF 50 is configured as ring 50, it has also been associated with TrBRF 100, and it has been assigned ports 7 and 8 of the Token Ring module located in slot 3.
  - TrBRF 100 is configured and is associated as the parent VLAN of TrCRF 30 and TrCRF 50.
- On the Catalyst 3900 the following TrBRF and TrCRF configuration exists:
  - TrCRF 20 is configured as ring 20, it has been associated with TrBRF 200, and it has been assigned Token Ring ports 2 and 3.
  - TrCRF 40 is configured as ring 40, it has also been associated with TrBRF 200, and it has been assigned Token Ring ports 7 and 8.
  - TrBRF 200 is configured and is associated as the parent VLAN of TrCRF 30 and TrCRF 50.

Configuration Steps

Once you have verified the switch network configuration, configuring IP routing on the RSM simply involves configuring an RSM interface for each of the bridged domains (TrBRFs) between which you want to configure IP routing and assigning an IP address to each of the RSM TrBRF interfaces.

Because your switch network is SRB, you are also going to configure a unique TrCRF on the RSM for each TrBRF RSM interface. This TrCRF is for the pseudo-ring that is used to cache the RIF for routed protocols.
To configure the TrBRF interfaces on the RSM module (located in slot 5 of the Catalyst 5000 series switch), complete the following tasks.

Accessing Global Configuration Mode on the RSM

To configure the TrBRF RSM interfaces, you must first access the RSM and enter global configuration mode. To access the RSM and enter global interface mode, complete the following tasks:

Step 1. At the Catalyst 5000 prompt, enter **enable** and press **Enter**.
Step 2. At the enable prompt, enter **session 5**.
Step 3. Enter **enable** to access the RSM configuration mode.
Step 4. Enter **configure terminal** to access the RSM global configuration mode.

Enabling IP Routing on the RSM

To enable IP routing on the RSM, complete the following task:

Step 1. Enter **ip routing**.

Configuring the TrBRF 100 RSM Interface

To configure the RSM interface for TrBRF 100, complete the following tasks:

Step 1. At the Router(config)# prompt, enter **interface vlan 100 type trbrf**.

Step 2. To assign an IP address to the TrBRF 100 RSM interface, enter **ip address 172.122.30.1 255.255.255.0**.

Step 3. Create the unique TrCRF for the pseudo-ring to enable the termination of RIF by entering **multiring trcrf-vlan 400 ring 100**.

Step 4. Configure multiring on this interface for all routed protocols by entering **multiring all**.

Step 5. For larger frame size support than the native Ethernet 1500 byte packets, change the MTU of the TrBRF 100 interface by entering **mtu 4472**.

Note: If you don't adjust the MTU of the TrBRF RSM interface, when RIFs pass through the RSM, the RSM adjusts the maximum packet size in the RIF to 1500 bytes.

Step 6. Enter **no shutdown** to administratively bring up the TrBRf 100 RSM interface.

Configuring the TrBRF 200 RSM Interface

To configure the RSM interface for TrBRF 200, complete the following tasks:

Step 1. At the Router(config)# prompt, enter **interface vlan 200 type trbrf**.

Step 2. To assign an IP address to the TrBRF 200 RSM interface, enter **ip address 172.122.40.1 255.255.255.0**.

Step 3. Create the unique TrCRF for the pseudo-ring to enable the termination of RIF by entering **multiring trcrf-vlan 500 ring 100**.

Step 4. Configure multiring on this interface for all routed protocols by entering **multiring all**.

Step 5. For larger frame size support than the native Ethernet 1500 byte packets, change the MTU of the TrBRF 100 interface by entering **mtu 4472**.

Step 6. Enter **no shutdown** to administratively bring up the TrBRF 100 RSM interface.

Exiting Global Configuration Mode and Verifying the RSM TrBRF Interface Configurations

To verify the configuration of the interfaces you created for TrBRF 100 and TrBRF 200, complete the following tasks:

Step 1. Enter **exit** to exit configuration mode.

Step 2. Verify the configuration by entering **show running-config**.

Step 3. Save changes to NVRAM by entering **copy running-config startup-config**.

The output (Figure 10-2) indicates that an RSM interface has been configured for both TrBRF 100 and TrBRF 200. Also, an IP address has been assigned to each RSM TrBRF interface and a unique TrCRF for the pseudo-ring has been created on the RSM for each of the TrBRFs.

Figure 10-2 Output of the show running-config Command

```
Router# show running-config
Building configuration...

Current configuration:
!
version 11.3
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
```

```
!
(additional displayed text omitted)
!
interface vlan100 type trbrf
mtu 4472
ip address 172.122.30.1 255.255.255.0
multiring trcrf-vlan 400 ring 100
multiring all
!
interface vlan200 type trbrf
mtu 4472
ip address 172.122.40.1 255.255.255.0
multiring trcrf-vlan 500 ring 100
multiring all
!
(additional displayed text omitted)
```

RSM Configuration Summary

The following is a summary of the configuration you completed on the RSM:

```
#accessing the rsm in slot 5
session 5
#accessing configuration mode on the RSM
enable
#entering global configuration mode on the RSM
configure terminal
#enabling ip routing on the RSM
ip routing
#configuring the interface for the trbrf 100 on the rsm
interface vlan 100 type trbrf
#assigning an IP parameters to the rsm brf100 interface
ip address 172.122.30.1 255.255.255.0
mtu 4472
#creating the unique TrCRF for the pseudo-ring for the brf100 interface
multiring trcrf-vlan 400 ring 100
#configuring multiring for all routed protocols on the interface
multiring all
#bringing up the brf100 RSM interface
no shutdown
#configuring the interface for the trbrf 200 on the rsm
interface vlan 200 type trbrf
#assigning an IP parameters to the rsm brf200 interface
ip address 172.122.40.1 255.255.255.0
mtu 4472
#creating the unique TrCRF for the pseudo-ring for the brf200 interface
multiring trcrf-vlan 500 ring 100
#configuring multiring for all routed protocols on the interface
multiring all
#bringing up the brf200 RSM interface
no shutdown
#exiting configuration mode
exit
#verifying the interface configuration
show running-config
#saving the configuration changes to NVRAM
copy running-config startup-config
```

Resulting Network Configuration

Figure 10-3 illustrates the resulting network configuration. You have now successfully configured IP routing between the Engineering and Marketing departments of your company via the RSM.

Figure 10-3 Resulting Network Configuration

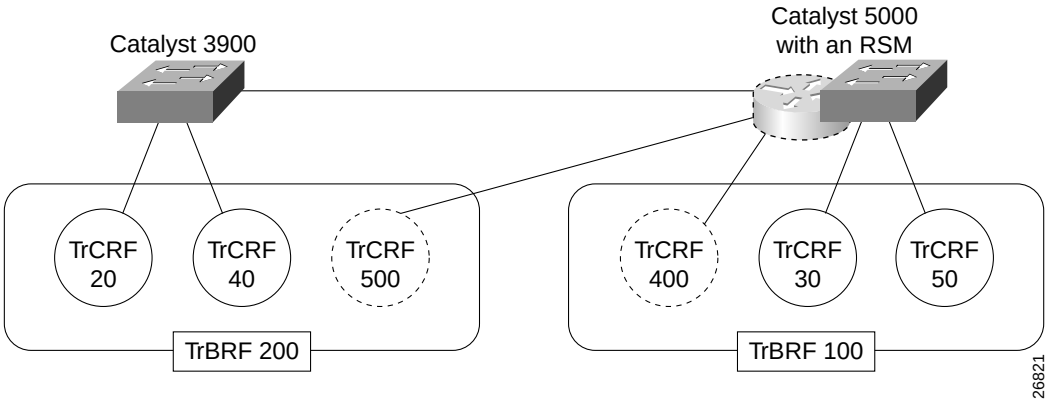

Configuring and Managing Token Ring Switches Using Cisco's Network Management Products

Cisco offers several network management applications that you can use to manage your Catalyst Token Ring switches. The following network management products are described in this chapter:

- CiscoView
- CWSI
- Resource Manager

CiscoView

CiscoView is a GUI-based device management software application that provides dynamic status, statistics, and comprehensive configuration information for Cisco's internetworking products (switches, routers, concentrators, and adapters). CiscoView graphically displays a physical view of Cisco devices. Additionally, this network management tool provides configuring and monitoring functions and offers basic troubleshooting. Using CiscoView, you can more easily understand the tremendous volume of management data available for internetworking devices because CiscoView organizes it into graphical device representations presented in a clear, consistent format.

CiscoView provides the following advantages over using direct SNMP gets and sets:

- Graphically displays Cisco products from a centralized network management location, giving network managers a complete view of Cisco products without physically checking each device at remote sites.
- Oriented for exception reporting, enabling users to access essential inquiry information quickly.
- Shows a continuously updated physical picture of routers, hubs, switches, or adapters.
- Can be invoked several times in the same session to simultaneously support multiple switches, routers, hubs, or adapters.
- Scrollable viewing for use with large devices such as the Catalyst 5500 and Cisco 7500.

CiscoView can run on UNIX workstations as a fully functional, independent management application. It is also available in Windows 95 and NT format on a PC. In addition, CiscoView can be integrated with the following SNMP-based network management systems to provide a seamless, powerful network view:

- Sun Microsystems SunNet Manager
- Hewlett-Packard OpenView
- IBM NetView for AIX

CWSI

CWSI is a GUI-based device management software application that provides a suite of integrated tools to provide topology management, traffic management, VLAN management, management of ATM networks, and user tracking.

Topology Management Services

The topology displays provided by CWSI include both physical and logical layer topology views of Cisco switches and routers. CWSI uses the information contained in CDP advertisements to draw a map of the discovered devices, including Catalyst Token Ring switches. From these maps, you can launch other CWSI functions.

As an integral component of the topology management, CWSI offers logical viewing and configuration services for many of the Catalyst Layer 2 functions, such as spanning tree design and analysis. Network operators can easily select any one of, or a multiple of, these functions with the new topology display service module. This service module customizes the view to a specific feature and then allows the operator to undertake guided configuration functions.

These display and configuration functions are accomplished by simply selecting the Catalyst features to be managed. CWSI highlights the appropriate configuration details on top of the physical map and offers configuration selections in which modifications can be made. Critical information, such as the root bridge within a calculated spanning tree, are all provided as part of these services.

Traffic Management Services

The traffic management services of CWSI allow you to monitor your enterprise networks from a central site, to ensure high network reliability and availability.

The traffic management services work as a distributed system by using a central management console running the software in conjunction with data-gathering agents located at various points on a network. It can simultaneously collect wide-ranging statistical data, display selectively captured and fully decoded network traffic, set user-defined alarm conditions, and get real-time updates from all segments of a widely dispersed internetwork.

The traffic management services are based on two standards that let them operate in a multi-topology, multivendor environment:

- SNMP defines the protocol for all intercommunications between the traffic management services and SwitchProbe devices.
- RMON MIB, which defines the type of information that the agent gathers that is available for you to display for each network segment.

The traffic management services provide functions for the following:

- Monitoring network traffic and measuring the flow of data
- Capturing network traffic and recording it for later examination
- Interpreting raw network data and translating it into a graphic form that you can view and analyze
- Setting limit conditions on network traffic and generating alarms if those limits are exceeded

Network traffic information is collected from RMON agents in devices running Cisco IOS software, SwitchProbe standalone network monitoring probes, or any RMON standards-compliant agent. Performance and fault management are also simplified by multilayer traffic analysis, proactive alarms, remote packet capture, and protocol decode features.

VLAN Management Services

The VLAN management services of CWSI allows you to configure, manage, and monitor interconnected Cisco switches and routers. Integral components of the VLAN management services include graphical mapping utilities for viewing and configuring logically defined workgroups, "drag-and-drop" port-level configuration options for assigning users to VLANs, automated link assignment settings for managing VLANs campus-wide, integration with common SNMP management platforms for consolidating system resources and detailed reporting functions for maintaining audit trails.

The benefits of VLAN management services include:

- Accurate campus view of Cisco switches and routers simplifies VLAN configuration and monitoring functions.
- On-screen reports of VLAN device, link, and port status are quickly obtained using pop-up windows and icon highlighting. Audit trails can then be generated from these reports.
- Drag-and-drop graphical configuration options provided by CiscoView minimize the skill level required to set up and assign ports to VLANs.
- A simplified VLAN naming window and associated directories offer flexibility and easy-to-use search functions for creating and changing VLAN names.
- User-selectable options for configuring VLANs across interswitch backbones include selecting multiple paths, choosing specific paths based on preferences, and manually adding and deleting paths.
- Logical views of configured VLANs include switch, link, and port membership windows, enabling users to audit and verify membership status.

User-selectable color options concurrently display multiple configured VLANs, making it easier to visualize and identify VLAN configurations.

ATM Management Services

The ATM management services of CWSI provides:

- Comprehensive discovery of ATM devices that support ILMI
- Virtual circuit setup and trace analysis of both PVCs and SVCs
- The launching of Cisco's graphical device management application from any of the Cisco icons within the topology map
- A real-time path trace analysis tool for checking end-to-end connectivity

Additional enhancements to the ATM management functionality in CWSI include performance analysis of the LANE components, configurations functions for PNNI routing, and traffic monitoring capabilities with ATM RMON. As part of the LANE management features, network managers can quickly check the status of the LECS database with a graphical representation on the location of this database within the ATM fabric, and they can query this database for ELAN-to-MAC and ELAN-to-network service access point (NSAP) address mappings. This information provides an address list of the end stations within each ELAN. Further, network managers can select the LES database and discover the MAC-to-NSAP address mappings to resolve network addresses. For redundancy, network managers can configure a backup LECS in conjunction with Cisco's SSRP for increased reliability. The ATM management functions provide utilities for synchronizing the database for naming consistency and integrity.

As part of the PNN configuration options, network managers can set administrative weightings and topology metrics to better optimize the communication across the ATM fabric. Further, network managers can view the PNNI-routed topology by selecting this display option within the topology map manager.

User Tracking Services

CWSI also provides utilities for tracking mobile users. The user tracking services of CWSI uses an automated VLAN authentication approach that verifies the user's station address prior to permitting access, and a newly developed tracking tool that maintains user location and identity as part of the embedded database functions. This management functions works seamlessly with Dynamic Host Configuration Protocol (DHCP), which automates TCP/IP addressing, and with the underlying topology services within CWSI that can pinpoint the user location within the topology map.

User tracking services provides:

- Discovery and reporting of end stations by MAC address and IP address
- Spreadsheet-style reporting of addressing information
- Configuration tool for assigning users to VLAN groups
- Pull-down preference settings for scheduling database updates
- Listing of all switch port-attached workstations
- Customized sorting tables for detailed reporting
- Scheduling managers for address tracking updating

Resource Manager

Cisco Resource Manager is a Web-based management solution for enterprise networks, offered on both Solaris and NT. It leverages Internet technologies to provide a flexible framework for simplifying several important tasks critical to network management. Cisco Resource Manager can run alongside CiscoWorks and CWSI, complementing their configurations and diagnostics capabilities with enhanced inventory and software distribution utilities for both routers and switches.

Resource Manager consists of four key management applications: Inventory Manager, Availability Manager, Syslog Analyzer, and Software Image Manager. Together these applications automate the task of finding software updates, speed device software deployment, provide multidevice views of network change, report on Year 2000 compliance, track device availability, and report, categorize, and analyze syslog messages, providing you probable cause and suggested actions.

Currently, the Inventory Manager and Software Image Manager are supported on the Catalyst 5000 series Token Ring switching module.

Inventory Manager

Inventory Manager centrally collects information on all types of network devices (routers, switches, hubs, and SNMP MIB II devices), allowing you to quickly find version and configuration information you need. Inventory Manager allows does the following:

- Quickly collects and displays up-to-date router and switch inventory details
- Notifies users of hardware and software configuration changes to network devices
- Allows users to view high-level device information or drill down to view a device's configuration details
- Allows users to group (create views of) devices by static or dynamic characteristics
- Provides device information to other Resource Manager modules

- Uses a CCO Internet connection to report on your network's compliance with Year 2000 certification information

Inventory Manager takes basic device seed information, entered directly or imported from CiscoWorks, CWW, CWSI, or HPOV, and adds detailed device characteristics to it. Imported data can be filtered to include only Cisco equipment. Once populated, Inventory Manager's database can be exported for uses in other applications.

Inventory Manager uses the concept of views to organize network devices into user-definable groups. A view can be static or dynamic. A static view contains specific devices known by name, whereas a dynamic view contains devices with a particular attribute such as model, location, or configuration characteristic. Dynamic views are powerful tools because they automatically adjust to reflect qualifying devices as they are added or removed from the network. Views created for Inventory Manager are shared with other Resource Manager applications.

Inventory Manager scans the network on a user-defined interval and gathers current hardware information (interface cards, Flash memory, firmware version) for devices it has been setup to manage. Network events such as a device reloading or restarting will automatically cause inventory info to be updated. One of the detailed reports that Inventory Manager provides is a proactive change-management report highlighting changes made to devices over time. Users can save a set number of these reports for historical tracking. Inventory information is displayed using a flexible reporting capability based on the views you have defined.

Software Image Manager

Software Image Manager automates many of the steps associated with scheduling, downloading, and monitoring software upgrades. The process of keeping switch and router images current starts with Software Image Manager scanning its database to check the software version, Flash memory size, and available RAM of each device you want to upgrade. It then notifies you if upgrades to Flash memory or RAM are required to accept the proposed image. Next, it allows you to bring the selected image down from CCO into Software Image Manager's library. Alternatively, software images can be loaded into the library from an existing device or from a local file system.

After you obtain the proper software image, you create a job defining which devices are to receive the new image(s), in what order, and what the script should do if an error is encountered. Software Image Manager can deploy software across the network at the scheduled time, synchronizing the download to multiple devices. To ensure that you know if all updates were successful, Software Image Manager generates detailed job reports which can be E-mailed upon completion, showing the status of each download performed.

Frame Formats

This appendix provides information about the formats of the following types of packets, frames, and cells:

- Token Ring Frame Format
- CDP Packet Format
- DRiP Frame Formats
- VTP Frame Format
- STP BPDU Frame Formats
- ISL Token Ring Frame Format
- ATM Cell Format

Token Ring Frame Format

Figure shows the format of a Token Ring frame.

Figure A-1 Token Ring Frame Format

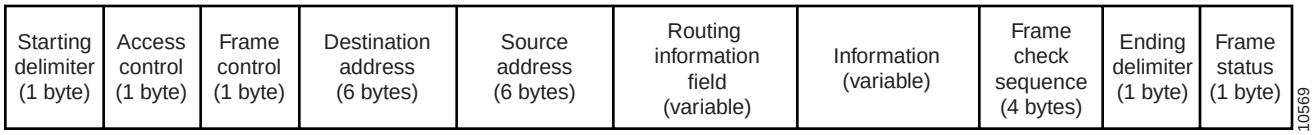

Starting Delimiter

The Starting Delimiter field indicates the arrival of a frame or token. This field includes bits that are set to intentionally violate the Differential Manchester Code to distinguish this field as a delimiter.

Access Control

The Access Control field contains the following bits:

- Priority bit—Used to indicate the priority of the frame or token.
- Reservation bit—Used to indicate the priority required for the next token to gain access to the ring.
- Token bit—Used to differentiate a token from a data or command frame.
- Monitor bit—Used by the active monitor to determine whether a frame is circling the ring endlessly.

Frame Control

The Frame Control field indicates the frame type and contains the following:

- Frame type bit—Used to indicate whether this is a MAC or LLC frame.
- Reserved bit—Reserved for future use.
- Control bits—Used to indicate whether the frame is to be processed by the normal buffer or the high-priority buffer.

Destination Address

The Destination Address field indicates the address of the device or devices for which the frame is intended. The destination address can be one of the following:

- Individual address—Identifies a particular ring station on the Token Ring network. This can be either a universally or locally administered address.
- Group address—Identifies a group of destination ring stations on the Token Ring network. This can be either a locally administered group address or a functional address, such as the functional address of the configuration report server.

Source Address

The Source Address field identifies the station that sent the frame. In the source address the first bit (bit 0) is called the routing information indicator (RII) bit. When this bit is set to one it indicates that the frame contains routing information. If the bit is set to zero then no routing information is included.

Routing Information

Used only in SRB, the Routing Information field indicates the route the frame is to take through the network. The routing information field consists of the following:

- Routing Control field
  - Broadcast indicators—Indicate whether the frame is to be sent along a specified path (nonbroadcast), through all bridges to all segments in a network (all-routes broadcast), or through only certain designated bridges so that the frame will appear only once on every network segment (single-route broadcast).
  - Direction bit—Indicates how the bridge should read the route descriptor when it forwards a frame.
- Route Descriptor field—Indicates the path using a ring number/bridge number/ring number sequence.

Information

The Information field contains the data that is being sent to upper layers.

Frame Check Sequence

The Frame Check Sequence field contains the cyclic redundancy check (CRC) value for all bits from the Frame Control field through the Frame Check Sequence field. The Frame Check Sequence value is checked by a receiving station to determine if errors occurred in transmission.

Ending Delimiter

The Ending Delimiter field indicates the end of the frame or token. It also contains bits to indicate if a frame is damaged or if the frame is the last in a logical sequence.

Frame Status

The Frame Status field indicates to the transmitting station whether this frame has been copied by the destination station.

CDP Packet Format

The CDP packet contains information about the Cisco devices in the network. It consists of a header, followed by a set of variable-length fields consisting of type/length/value triplets.

Figure A-2 shows the format of a CDP packet.

Figure A-2 CDP Packet Format

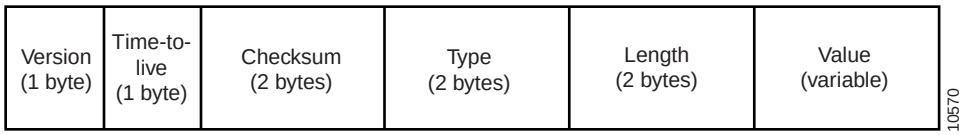

Version

The Version field indicates the version of CDP being used. The value is always 0x01.

Time-to-Live

The Time-to-Live field indicates the amount of time, in seconds, that a receiver should retain the information contained in this packet.

Checksum

The Checksum field indicates the standard IP checksum.

Type

The Type field indicates the type/length/value type. The possible CDP type/length/value types are as follows:

- Device ID
- Address
- Port ID
- Capabilities
- Version
- Platform
- IP Prefix

Length

The Length field indicates the total length, in bytes, of the type, length, and value fields.

Value

The Value field contains the type/length/value value, which depends on the type/length/value type as described below:

- Device ID

  The device ID type/length/value (type 0x0001) identifies the device. This type of type/length/value allows different address references to be associated with the same device.

  By default, the device ID is either the device's fully-qualified host name (including the domain name) or the device's hardware serial number in ASCII.

- Address

  The address type/length/value (type 0x0002) contains a number that indicates how many addresses are contained in the packet, followed by one entry for each address being advertised. The addresses advertised are the ones assigned to the interface on which the CDP message is sent. A device can advertise all addresses for a given protocol suite and, optionally, can advertise one or more loopback IP addresses. If the device can be managed by SNMP, the first entry in the address type/length/value is an address at which the device receives SNMP messages.

  Figure A-3 shows the format of each address contained in the packet.

Figure A-3 Address Type/Length/Value Fields

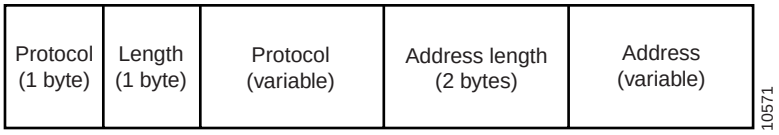

Table A-1 explains the fields in the address type/length/value packet.

Table A-1 Address Type/Length/Value Packet Fields

Field	Description
Protocol	Protocol type. It can be one of the following values: • 1—NLPID format • 2—802.2 format
Length	Length of the protocol field. For protocol type 1, the length is 1. For protocol type 2, the length is either 3 or 8, depending on whether SNAP is used.
Protocol	One of the following values: • 0x81—ISO CLNS (protocol type 3D 1) • 0xCC—IP (protocol type 3D 1) • 0xAAAA03 000000 0800—Pv6 (protocol type 3D 2) • 0xAAAA03 000000 6003—DECNET Phase IV (protocol type 3D 2) • 0xAAAA03 000000 809B—AppleTalk (protocol type 3D 2) • 0xAAAA03 000000 8137—Novell IPX (protocol type 3D 2) • 0xAAAA03 000000 80c4—Banyan VINES (protocol type 3D 2) • 0xAAAA03 000000 0600— XNS (protocol type 3D 2) • 0xAAAA03 000000 8019—Apollo Domain (protocol type 3D 2)
Address length	Length of the address field in bytes.
Address	Address of the interface, or the address of the system if addresses are not assigned to the interface.

- Port ID

  The port ID type/length/value (type 0x0003) contains an ASCII character string that identifies the port on which the CDP message is sent. The type/length/value length determines the length of the string.

- Capabilities

  The capabilities type/length/value (type 0x0004) describes the device's functional capability. It can be set to one of the bits listed in Table A-2.

Table A-2 Capabilities Type/Length/Value Bit Definitions

Bit	Description
0x01	Performs level 3 routing for at least one network layer protocol.
0x02	Performs level 2 transparent bridging.
0x04	Performs level 2 source-route bridging. A source-route bridge would set both this bit and bit 0x02.
0x08	Performs level 2 switching. The difference between this bit and bit 0x02 is that a switch does not run the STP. This device is assumed to be deployed in a physical loop-free topology.
0x10	Sends and receives packets for at least one network layer protocol. If the device is routing the protocol, this bit should not be set.
0x20	The bridge or switch does not forward IGMP Report packets on nonrouter ports.
0x40	Provides level 1 functionality.

- Version

  The version type/length/value (type 0x0005) contains a character string that provides information about the software release version that the device is running. The type/length/value length field determines the length of the string.
- Platform

  The platform type/length/value (type 0x0006) contains an ASCII character string that describes the hardware platform of the device. The type/length/value length field determines the length of the string. The following are the possible string values:
  - Cisco 7000
  - Cisco 7010
  - Cisco 4500
  - Cisco 3100
  - Cisco 3000
  - Cisco 2500
  - Cisco 2000
  - Cisco 1000
  - AGS+
  - AGS
  - MGS
  - CGS
  - IGS
  - cs500
  - Catalyst
  - A100
  - Synergy

- IP Prefix

  The IP Prefix type/length/value (type 0x0007) contains a set of 0 or more IP prefixes in its value field. No prefixes are included when the type/length/value's length field is 0. Otherwise, the length field includes the length of the type and value fields, plus 5 bytes for every IP prefix included. Each IP prefix consists of 4 bytes of IP network number and 1 byte representing the network mask. The network mask can be in the range 0 through 32, and represents the number of bits set in the mask (left contiguous).

  Each IP prefix represents one of the directly connected IP network segments of the local router. This type/length/value enables an IP stub router to communicate IP topology information to a central site router, without requiring the configuration of a full-blown IP routing protocol.

DRiP Frame Formats

The DRiP frame contains information about the VLANs configured in the management domain. It consists of some header information followed by one or more VLAN information fields.

Figure A-4 shows the format of a DRiP frame.

Figure A-4 DRiP Frame Fields

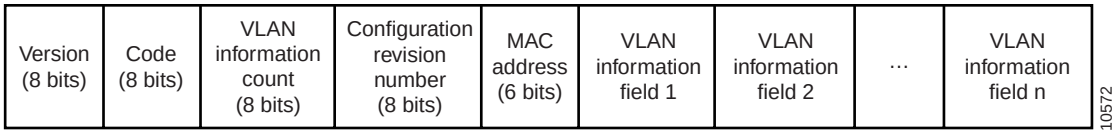

Version

The Version field identifies the version of DRiP being used.

Code

The Code field indicates whether this message is an advertisement that indicates a change (0x01) or no change (0xFF).

VLAN Information Count

The VLAN Information Count field indicates the number of VLAN information fields contained in this advertisement.

Header Length

The Header Length field indicates the size in bytes of the header for this type of advertisement. The header includes all fields from the version up to the first VLAN information field.

Configuration Revision Number

The Configuration Revision Number field indicates the revision number of the configuration information. A configuration revision number starts at zero and increments by one with each modification until it reaches the value 4294947295, at which point it wraps back to zero and starts incrementing again.

Last Changed Revision

The Last Changed Revision field indicates the revision number of the last change associated with the originating MAC address. Switches in the domain compare the value in this field to their current configuration number to determine whether the advertisement contains new information.

MAC Address

The MAC Address field contains the MAC address, in canonical format, of the device that is sending the DRiP advertisement.

VLAN Information

The VLAN Information fields contain information for each active or configured TrCRF on the switch. A TrCRF is considered active if a port associated with the TrCRF is open on the ring. If a TrCRF ceases to be included in the VLAN information field, it indicates there are no longer ports active or configured on the TrCRF. The TrCRF should then be removed from the database.

If a periodic timer triggers an advertisement, regardless of whether there has been configuration revision change, all VLAN information for the device is included.

Figure A-5 shows the format of the VLAN information fields.

Figure A-5 VLAN Information Fields

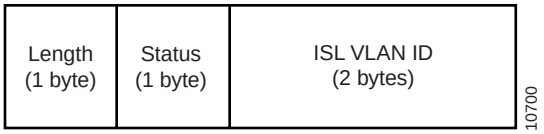

Length

The Length field indicates the length, in bytes, of the VLAN information field (including this length field). This length will be a multiple of 4.

Status

The Status field indicates that the status of the TrCRF has changed. The TrCRF now either has an active port or the last active port that was on the TrCRF has become inactive, leaving the TrCRF with no active ports. Possible values are the following:

- Bit 7 (0x00)—No ports are active on the TrCRF.
- Bit 7 (0x01)—Active ports exist on the TrCRF.
- Bit 6 (0x00)—No ports are configured on the TrCRF.
- Bit 6 (0x01)—Ports are configured on the TrCRF.
- Bit 0 through 5 (0x00)—Reserved.

If the message is triggered by a periodic timer, then the status will indicate the current status of the VLAN.

ISL VLAN ID

The ISL VLAN ID indicates the VLAN ID of this VLAN on ISL trunks. Possible values are 0 through 1023.

VTP Frame Format

There are three types of VTP frames: Advert-Request, Summary-Advert, and Subset-Advert.

Advert-Request Frame Format

An Advertisement Request (Advert-Request) is a request for configuration information. Figure A-6 shows the format of an Advert-Request frame.

Figure A-6 Advert-Request Frame Format

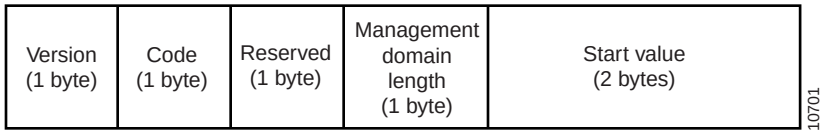

Version

The Version field indicates the VTP version number. This value is always 0x01.

Code

The Code field indicates the message type. Possible values are:

- 0x01—Summary-Advert
- 0x02—Subset-Advert
- 0x03—Advert-Request

Management Domain Length

The Management Domain Length field indicates the length of the name of the management domain.

Start Value

The Start Value field indicates the VLAN ID of the first VLAN for which information is requested. Any response to the request should contain information for all VLANs having an ISL VLAN ID greater than or equal to this value. For example, in a request for information on all VLANs, this value is 0.

Summary-Advert Frame Format

The Summary Advertisement (Summary-Advert) contains information about the sending device and summary information about the advertisement, including the number of subset advertisements to follow. The maximum size of a Summary-Advert is 1492 bytes. Figure A-6 shows the format of a Summary-Advert frame.

Figure A-7 Summary-Advert Frame Format

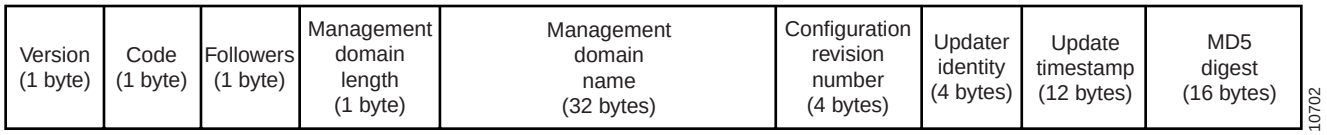

Version

The Version field indicates the VTP version number. This value is always 0x01.

Code

The Code field indicates the message type. Possible values are:

- 0x01—Summary-Advert
- 0x02—Subset-Advert
- 0x03—Advert-Request

Followers

The Followers field indicates the number of Subset-Advert messages that follow this Summary-Advert.

Management Domain Length

The Management Domain Length field indicates the length of the name of the management domain.

Management Domain Name

The Management Domain Name field indicates the name of the management domain.

Configuration Revision Number

The Configuration Revision Number field indicates the revision number of the configuration information. As with CDP configuration revision numbers, a configuration revision number starts at zero and increments by one with each modification until it reaches the value 4294947295, at which point it wraps back to zero and starts incrementing again.

Updater Identity

The Updater Identity field indicates the IP address of the device that received the command that caused the configuration revision number to have its current value.

Update Timestamp

The Update Timestamp field indicates the time at which the configuration revision number was most increased to its current value. The timestamp is in the format “yymmddhhmmss”, where *yymmdd* represents the year, month, and day and *hhmmss* represents the hours, minutes, and seconds.

MD5 Digest

MD5 digest value over the secret value and all VLAN information

Subset-Advert Frame Format

The Subset Advertisement (Subset-Advert) contains information about the VLANs being advertised. Figure A-6 shows the format of a Subset-Advert frame.

Figure A-8 Subset-Advert Frame Format

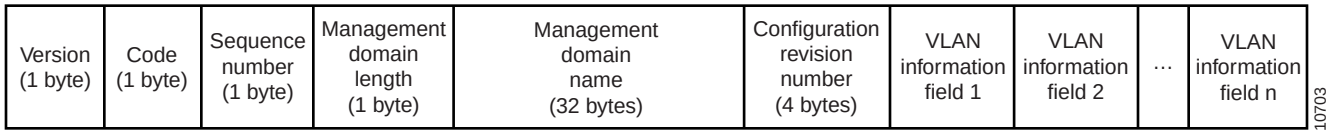

Version

The Version field indicates the VTP version number. This value is always 0x01.

Code

The Code field indicates the message type. Possible values are:

- 0x01—Summary-Advert
- 0x02—Subset-Advert
- 0x03—Advert-Request

Sequence Number

The Sequence Number field indicates the order of this Subset-Advert frame within the series of Subset-Advert frames that follow a Summary-Advert. For the first Subset-Advert frame following a Summary-Advert frame the sequence number is 1.

Management Domain Length

The Management Domain Length field indicates the length of the name of the management domain.

Management Domain Name

The Management Domain Name field indicates the name of the management domain.

Configuration Revision Number

The Configuration Revision field indicates the revision number of the configuration information. As with CDP configuration revision numbers, a configuration revision number starts at zero and increments by one with each modification until it reaches the value 4294947295, at which point it wraps back to zero and starts incrementing again.

VLAN Information Field

Each VLAN Information field contains information for a different VLAN, starting with the VLAN with the lowest ISL VLAN IDs.

Figure A-6 shows the format of the VLAN information field.

Figure A-9 VLAN Information Field

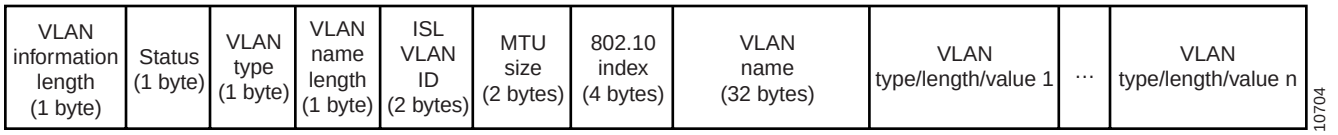

VLAN Information Length

The VLAN Information Length field indicates the length, in bytes, of the VLAN information field for this VLAN in this advertisement. The length is a multiple of 4.

Status

The Status field indicates the status of this VLAN. Possible values are:

- Bit 0 (0x01)—VLAN suspended
- Bits 1 through 7 (0x02 through 0x80)—Reserved

VLAN Type

The VLAN Type field indicates the type of VLAN. Possible values are:

- 0x01—Ethernet
- 0x02—FDDI
- 0x03—TrCRF
- 0x04—FDDI-net
- 0x05—TrBRF

VLAN Name Length

The VLAN Name Length field indicates the length, in bytes, of the VLAN name for this VLAN.

ISL VLAN ID

The ISL VLAN ID field indicates the ID of this VLAN on ISL trunks. Possible values are 0 through 1023.

MTU Size

The MTU Size field indicates the maximum transmission unit (MTU) for this VLAN. Possible values are 1500 through 18190.

802.10 Index

The 802.10 Index field indicates the 802.10 security association identifier (SAID) value for this VLAN.

VLAN Name

The VLAN Name field indicates the VLAN name for this VLAN. The name can be between 1 and 32 bytes in length and is padded with zeros.

VLAN Type/Length/Value Field

The VLAN type/length/value fields are variable in length and contain the a type, length, and value. Figure A-10 shows the format of the VTP variable-length fields.

Figure A-10 VLAN Type/Length/Value Field

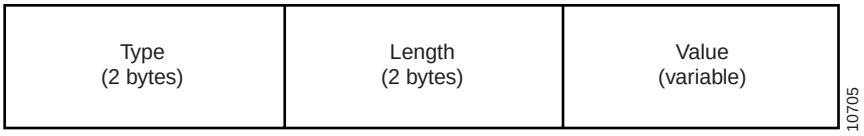

Field	Description
Type	The Type field indicates the type/length/value type. Possible values are: • 0x01—Source-Routing Ring Number • 0x02—Source-Routing Bridge Number • 0x03—STP Type • 0x04—Parent VLAN • 0x05—Translationally bridged VLANs • 0x06—Pruning • 0x07—Bridge Type • 0x08—Max ARP Hop Count • 0x09—Max STE Hop Count • 0x0A—Backup CRF Mode
Length	The Length field indicates the length of this VLAN type/length/value.
Value	The Value field contains the type/length/value value, which depends on the type/length/value type as described below: • Source-Routing Ring Number—Number that uniquely identifies this ring in a source-routed network. • Source-Routing Bridge Number—Number that uniquely identifies this bridge in a source-routed network. • STP Type—Type of STP being used. Possible values are 1 (SRT), 2 (SRB), and 3 (Auto). • Parent VLAN—ISL VLAN ID of the TrBRF to which this TrCRF is assigned. • Translationally Bridged VLANs—ISL VLAN ID of the VLANs to which this VLAN is translational-bridged, formatted as 2 bytes per VLAN appended by 2 bytes of zeros. • Pruning—Whether VTP pruning is enabled. Possible values are 1 (Enabled) and 2 (Disabled). • Bridge Type—Bridging mode of the VLAN. Possible values are 1 (SRT) and 2 (SRB). • Max ARP Hop Count—Maximum number of hops for ARE frames processed by this TrCRF. Possible values are 1 through 13. The default is 7. • Max STE Hop Count—Maximum number of hops for STE frames processed by this TrCRF. Possible values are 1 through 13. The default is 7. • Backup CRF Mode—Whether the TrCRF is configured as a backup. Possible values are 1 (TrCRF is configured as a backup) and 2 (TrCRF is not configured as a backup).

STP BPDU Frame Formats

The format of a STP BPDU frame varies depending on the type of protocol used.

Figure A-11 shows the format of an IEEE 802.1d STP BPDU frame.

Figure A-11 IEEE 802.1d STP BPDU Frame Format

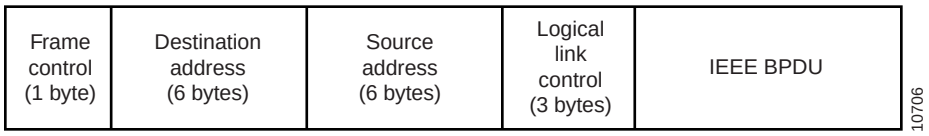

Figure A-12 shows the format of an IBM STP BPDU frame.

Figure A-12 IBM STP BPDU Frame Format

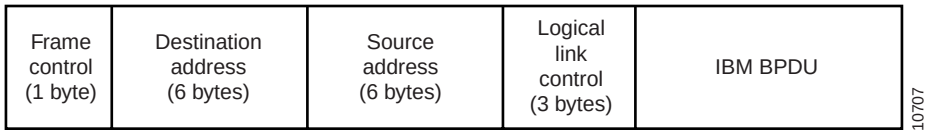

Figure A-13 shows the format of a Cisco STP BPDU frame.

Figure A-13 Cisco STP BPDU Frame Format

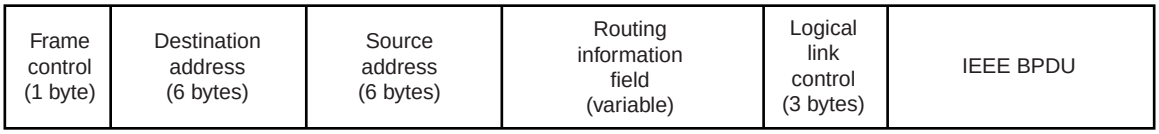

Frame Control

The Frame Control field is always 01.

Destination Address

The Destination Address field indicates the destination address as specified in the Bridge Group Address table. For IEEE STP BPDU frames, the address is 0x800143000000. For IBM STP BPDU frames, the address is 0xC00000000100. For Cisco STP BPDU frames, the address is 0x800778020200.

Source Address

The Source Address field indicates the base MAC address used by the switch. For Cisco STP BPDU frames, the multicast bit is set to indicate the presence of a RIF in the header.

Routing Information Field

Applicable only to Cisco STP BPDU frames, the Routing Information field must be set to 0x0200.

Logical Link Control

For all three types of STP BPDU frames, this field is set to 0x424203.

BPDU

Figure A-13 shows the format of the fields inside a BPDU.

Figure A-14 BPDU Field Formats

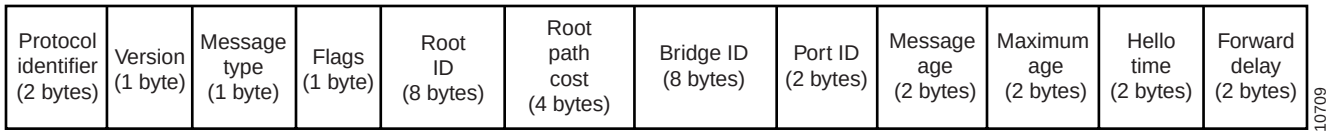

Note: All fields in the BPDU are common to all STPs except for the Port ID field. For IEEE and Cisco STP BPDU frames, the Port ID field specifies the transmitting port number of the originating bridge. For IBM STP BPDU frames, the Port ID field specifies the ring and bridge number through which the message was sent.

Protocol Identifier

The Protocol Identifier Field indicates the type of protocol. This field contains the value zero.

Version

The Version field indicates the version of the protocol. This field contains the value zero.

Message Type

The Message Type field indicates the type of message. This field contains the value zero.

Flags

The Flags field includes one of the following:

- Topology change (TC) bit, which signals a topology change
- Topology change acknowledgment (TCA) bit, which is set to acknowledge receipt of a configuration message with the TC bit set.

Root ID

The Root ID field indicates the root bridge by listing its 2-byte priority followed by its 6-byte ID.

Root Path Cost

The Root Path Cost field indicates the cost of the path from the bridge sending the configuration message to the root bridge.

Bridge ID

The Bridge ID field indicates the priority and ID of the bridge sending the message.

Port ID

The Port ID field indicates the port number (IEEE or Cisco STP BPDU) or the ring and bridge number (IBM STP BPDU) from which the configuration message was sent. This field allows loops created by multiple attached bridges to be detected and corrected.

Message Age

The Message Age field indicates the amount of time that has elapsed since the root sent the configuration message on which the current configuration message is based.

Maximum Age

The Maximum Age field indicates when the current configuration message should be deleted.

Hello Time

The Hello Time field indicates the time between root bridge configuration messages.

Forward Delay

The Forward Delay field indicates the length of time that bridges should wait before transitioning to a new state after a topology change. If a bridge transitions too soon, it is possible that not all network links will be ready to change their state and loops can result.

ISL Token Ring Frame Format

To support Token Ring, another ISL frame format was developed. In addition to the fields found in the original ISL frame format, the ISL Token Ring Frame format includes:

- An extra 6 byte header.
- The routing information field scanning results.
- The source VLAN ID.
- A size indicator.
- A flag for the type of explorer.

Figure A-13 shows the format of an ISL Token Ring frame.

Figure A-15 ISL Token Ring Frame Format

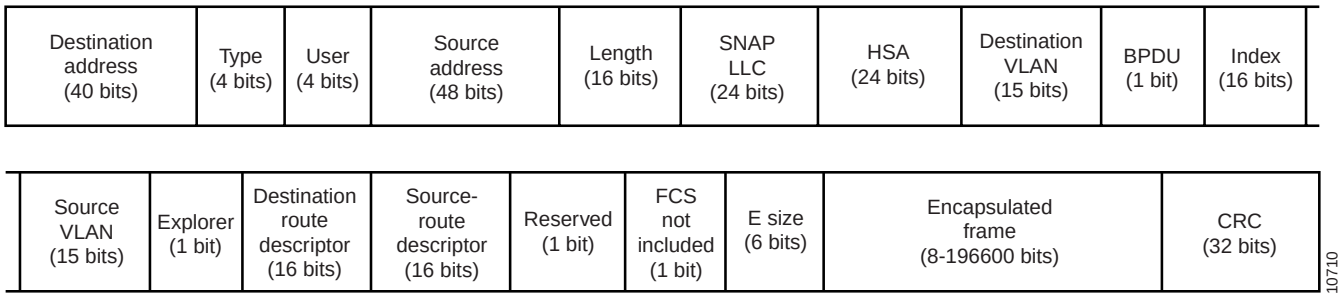

Destination Address

The Destination Address field is a 40-bit multicast address and is set to 0x01000C0000

Type

The Type field indicates the type of frame that is encapsulated. For Token Ring frames, this field is set to 0001.

User

The User field extends the meaning of the Type field. For example, Token Ring frames may have more than one type. The default User field value is 0000.

Source Address

The Source Address field indicates the 802.3 MAC address of the MAC transmitting the frame.

Length

The Length field indicates the length, in bytes, of the frame excluding the Destination Address, Type, User, Source Address, Length, and CRC fields.

SNAP LLC

The SNAP LLC of the frame. For ISL frames this field is set to AAAA03.

HSA

The HSA (high bits of source address) field indicates the upper 3 bytes, which identifies the manufacturer, of the Source Address field.

Destination VLAN

The Destination VLAN field indicates the ID of VLAN for which the packet is destined. This value is used to distinguish frames on different VLANs. This field is often referred to as the *color* of the packet.

BPDU

The BPDU field indicates whether the encapsulated frame is a BPDU. This field is also used to indicate whether the encapsulated frame is a CDP or VTP frame. All frames received with this field set are forwarded to the CPU for processing.

Index

The Index field indicates the port index of the source of the frame as it comes out from the Catalyst switch. It is used for diagnostic purposes only and may be set to any value by other devices.

Source VLAN

The Source VLAN field indicates the ID of VLAN from which the packet was sent.

Explorer

The Explorer field indicates whether the encapsulated frame is a data frame or and explorer (ARE or STE) frame.

Destination Route Descriptor

The Destination Route Descriptor field indicates the route descriptor to be used for forwarding. If there is no route descriptor following the routing information field match in the routing information field or if there is no routing information field present in the frame, this field is set to 0 and the destination address is used for forwarding.

Source-Route Descriptor

The Source Route Descriptor field indicates the route descriptor to be used for source learning. If there is no route descriptor prior to the ring-in in the routing information field or if there is no routing information field present in the frame, this field is set to 0 and the source address is used for source learning.

FCS Not Included

The FCS Not Included field indicates whether the Frame Check Sequence field is included in the Encapsulated Frame field.

E Size

The E Size field indicates the frame size for frames less than 64 bytes. This field is use to account for the case where a frame crosses a router and is padded to 64 bytes (minimum Ethernet frame).

Encapsulated Frame

The actual Token Ring frame. For more information on the format of the Token Ring frame, see the "Token Ring Frame Format" section.

CRC

The CRC field is the frame checksum. This CRC is in addition to the one at the end of the Encapsulate Frame field. It contains a standard 32-bit CRC value calculated on the entire encapsulated frame from the Destination Address field to the Encapsulated Frame field. The receiving device checks this CRC and may discard packets that do not have a valid CRC on them.

ATM Cell Format

The ATM cell is a fixed-length, standard unit of data transmission for all cell relay services in an ATM network. The first five bytes of the ATM cell serve as the cell header. The cell header contains information essential to routing the cell through the network and ensuring that the cell reaches its destination. Figure A-13 shows the format of an ATM cell.

Figure A-16 ATM Cell Format

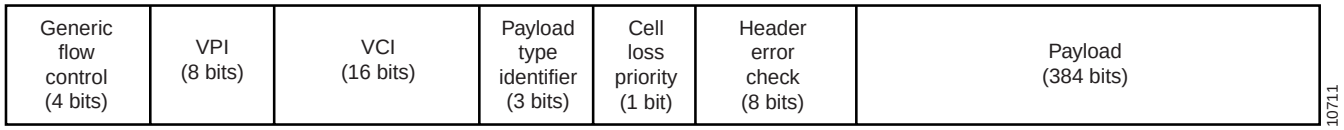

Generic Flow Control

The Generic Flow Control field is used when passing ATM traffic through a user-to-network (UNI) interface to alleviate short-term overload conditions. A network-to-network (NNI) interface does not use this field for generic flow control purposes; rather, an NNI uses this field to define a larger VPI value for trunking purposes.

VPI

The VPI field identifies the virtual paths. In an idle or null cell, the VPI field is set to all zeros. (A cell containing no information in the payload field is either idle or null). VPIs provide a way to bundle ATM traffic being sent to the same destination.

In an ATM UNI header, part of the VPI field (bits 5 through 8 of byte 1) is reserved as a generic flow control field. However, the ATM NNI header provides a larger range of VPI values (using bits 5 through 8 of byte 2, in addition to bits 1 through 4 of byte 1). This larger range of VPI values that can be defined in an ATM NNI cell header reflects the greater use of virtual paths in the network for trunking purposes between ATM inter-switch and ATM inter-network interfaces.

VCI

The VCI field identifies a particular VCC. In an idle or null cell (one containing no payload information), the VCI field is set to all zeros. Other non-zero values in this field are reserved for special purposes. For example, the combined values of VPI = 0 and VCI = 5 are used exclusively for ATM signaling purposes when requesting an ATM connection.

Payload Type Identifier

The Payload Type Identifier indicates the type of payload the cell contains: either user data or special network management data used to perform certain network operation, administration, and maintenance functions in the network.

Cell Loss Priority

The Cell Loss Priority field is set by the AAL to indicate the relative importance of a cell. This field is set to 1 to indicate that a cell can be discarded, if necessary, such as when an ATM switch is experiencing traffic congestion. This field is set to 1 to indicate that the cell should not be discarded, such as when supporting a specified or guaranteed quality of service. This field may also be set by the ATM layer if an ATM connection exceeds the QOS parameters established during connection setup.

Header Error Check

The Header Error Check field is an 8-bit CRC computed on all fields in an ATM UNI/NNI cell header. The header error check is capable of detecting all single-bit errors and certain multiple-bit errors. This field provides protection against incorrect message delivery caused by addressing errors. However, it provides no error protection for the ATM cell payload proper. The physical layer uses this field for cell delineation functions during data transport.

Glossary of Terms

This appendix contains a list of the terms and acronyms used in this document.

Numerics

100BaseT
100-Mbps baseband Fast Ethernet specification using UTP wiring. Like the 10BaseT technology on which it is based, 100BaseT sends link pulses over the network segment when no traffic is present. However, these link pulses contain more information than those used in 10BaseT. Based on the IEEE 802.3 standard. See also *Fast Ethernet* and *IEEE 802.3*.

A

AAL
ATM adaptation layer. Service-dependent sublayer of the data link layer. The AAL accepts data from different applications and presents it to the ATM layer in the form of 48-byte ATM payload segments. AALs consist of two sublayers, convergence sublayer (CS) and segmentation and reassembly (SAR). AALs differ on the basis of the source-destination timing used, whether they use CBR or VBR, and whether they are used for connection-oriented or connectionless mode data transfer. At present, the four types of AAL recommended by the ITU-T are AAL1, AAL2, AAL3/4, and AAL5. See *AAL1*, *AAL2*, *AAL3/4*, *AAL5*, *CS*, and *SAR*. See also *ATM* and *ATM layer*.

AAL1
ATM adaptation layer 1. One of four AALs recommended by the ITU-T. AAL1 is used for connection-oriented, delay-sensitive services requiring constant bit rates, such as uncompressed video and other isochronous traffic. See also *AAL*.

AAL2
ATM adaptation layer 2. One of four AALs recommended by the ITU-T. AAL2 is used for connection-oriented services that support a variable bit rate, such as some isochronous video and voice traffic. See also *AAL*.

AAL3/4

ATM adaptation layer 3/4. One of four AALs (merged from two initially distinct adaptation layers) recommended by the ITU-T. AAL3/4 supports both connectionless and connection-oriented links, but is primarily used for the transmission of SMDS packets over ATM networks. See also *AAL*.

AAL5

ATM adaptation layer 5. One of four AALs recommended by the ITU-T. AAL5 supports connection-oriented, VBR services, and is used predominantly for the transfer of classical IP over ATM and LANE traffic. AAL5 uses SEAL and is the least complex of the current AAL recommendations. It offers low bandwidth overhead and simpler processing requirements in exchange for reduced bandwidth capacity and error-recovery capability. See also *AAL*.

ABR

Available bit rate. QOS class defined by the ATM Forum for ATM networks. ABR is used for connections that do not require timing relationships between source and destination. ABR provides no guarantees in terms of cell loss or delay, providing only best-effort service. Traffic sources adjust their transmission rate in response to information they receive describing the status of the network and its capability to successfully deliver data. Compare with *CBR*, *UBR*, and *VBR*.

active monitor

Device responsible for managing a Token Ring. A network node is selected to be the active monitor if it has the highest MAC address on the ring. The active monitor is responsible for such management tasks as ensuring that tokens are not lost, or that frames do not circulate indefinitely. See also *ring monitor* and *standby monitor*.

active port monitor

A type of monitoring supported by the Switched Port Analyzer (SPAN) that allows you to monitor traffic using a customer-supplied monitoring device, such as an RMON probe, or a trace tool, such as a Network General Sniffer. The trace tool monitors only the LLC traffic that is switched by the monitored port. The MAC frames are not monitored. See also *SPAN*.

adaptive cut-through switching

A switching feature that alternates between cut-through and store-and-forward switching modes based on preset, user-defined error thresholds to optimize performance while providing protection from network errors.

address mask

Bit combination used to describe which portion of an address refers to the network or subnet and which part refers to the host. Sometimes referred to simply as mask. See also *subnet mask*.

address resolution

Generally, a method for resolving differences between computer addressing schemes. Address resolution usually specifies a method for mapping network layer (Layer 3) addresses to data link layer (Layer 2) addresses.

Address Resolution Protocol

See *ARP*.

algorithm

Well-defined rule or process for arriving at a solution to a problem. In networking, algorithms are commonly used to determine the best route for traffic from a particular source to a particular destination.

all-routes explorer
See *ARE*.

ANSI
American National Standards Institute. Voluntary organization comprised of corporate, government, and other members that coordinates standards-related activities, approves U.S. national standards, and develops positions for the United States in international standards organizations. ANSI helps develop international and U.S. standards relating to, among other things, communications and networking. ANSI is a member of the IEC and the ISO.

application-specific integrated circuit
See *ASIC*.

ARE
All-routes explorer. Explorer packet that traverses an entire SRB network, following all possible paths to a specific destination. Sometimes called all-rings explorer packet.

ARP
Address Resolution Protocol. Internet protocol used to map an IP address to a MAC address. Defined in RFC 826.

ASIC
Application-specific integrated circuit. A development process for implementing integrated circuit designs. Integrated circuit designs which are specific to the intended application, as opposed to designs for general purpose use. Both the Quad Token Ring Port chip and the Quad Media Access Control chip are implemented in ASIC.

ASP
ATM switch processor.

ATM
Asynchronous Transfer Mode. A packet-switching technology developed to support both voice and data on a common network infrastructure. ATM uses fixed-length 53-byte cells and can be transported on both LANs and WANs at a variety of operating rates. Because ATM is also application transparent, it is possible for it to be used to transport voice, data, images, and video on the same network.

ATM Forum
International organization jointly founded in 1991 by Cisco Systems, NET/ADAPTIVE, Northern Telecom, and Sprint that develops and promotes standards-based implementation agreements for ATM technology. The ATM Forum expands on official standards developed by ANSI and ITU-T, and develops implementation agreements in advance of official standards.

ATM layer
Service-independent sublayer of the data link layer in an ATM network. The ATM layer receives the 48-byte payload segments from the AAL and attaches a 5-byte header to each, producing standard 53-byte ATM cells. These cells are passed to the physical layer for transmission across the physical medium. See also *AAL*.

ATM UNI
See *UNI*.

ATM user-user connection

Connection created by the ATM layer to provide communication between two or more ATM service users, such as ATMM processes. Such communication can be unidirectional, using one VCC, or bidirectional, using two VCCs. See also *ATM layer* and *VCC*.

B

backbone

The part of a network that acts as the primary path for traffic that is most often sourced from, and destined for, other networks.

backup TrCRF

A type of TrCRF that enables you to configure an alternate route for traffic between undistributed TrCRFs located on separate switches that are connected by a TrBRF, in case the ISL connection between the switches becomes inactive.

balun

Balanced, unbalanced. Device used for matching impedance between a balanced and an unbalanced line, usually twisted-pair and coaxial cable.

bandwidth

The difference between the highest and lowest frequencies available for network signals. The term is also used to describe the rated throughput capacity of a given network medium or protocol.

baud

Unit of signaling speed equal to the number of discrete signal elements transmitted per second. Baud is synonymous with bits per second (bps), if each signal element represents exactly 1 bit.

beacon

Frame from a Token Ring or FDDI device indicating a serious problem with the ring, such as a broken cable. A beacon frame contains the address of the station assumed to be down. See also *failure domain*.

BPDU

Bridge protocol data unit. STP hello packet that is sent out at configurable intervals to exchange information among bridges in the network. See also *PDU*.

bps

Bits per second.

BRF

Bridge relay function. As defined by the IEEE, an internal bridge function on a Token Ring switch that is responsible for forwarding frames between port groupings with the same logical ring number (CRFs). Within a BRF, source-route bridging or source-route transparent bridging can be used to forward frames. See also *CRF*.

bridge

Device that connects and passes packets between two network segments that use the same communications protocol. Bridges operate at the data link layer (Layer 2) of the OSI reference model. In general, a bridge will filter, forward, or flood an incoming frame based on the MAC address of that frame.

bridge forwarding
Process that uses entries in a filtering database to determine whether frames with a given MAC destination address can be forwarded to a given port or ports. Described in the IEEE 802.1 standard. See also *IEEE 802.1*.

bridge group
Bridging feature that assigns network interfaces to a particular spanning-tree group. Bridge groups can be compatible with the IEEE 802.1 or the DEC specification.

bridge number
Number that identifies each bridge in an SRB LAN. Parallel bridges must have different bridge numbers.

bridge protocol data unit
See *BPDU*.

bridge relay function
See *BRF*.

bridge static filtering
Process in which a bridge maintains a filtering database consisting of static entries. Each static entry equates a MAC destination address with a port that can receive frames with this MAC destination address and a set of ports on which the frames can be transmitted. Defined in the IEEE 802.1 standard. See also *IEEE 802.1*.

broadcast
Data packet that will be sent to all nodes on a network. Broadcasts are identified by a broadcast address. Compare with multicast and unicast. See also *broadcast address*.

broadcast address
Special address reserved for sending a message to all stations. Generally, a broadcast address is a MAC destination address of all ones. Compare with multicast address and unicast address. See also *broadcast*.

broadcast and unknown server
See *BUS*.

broadcast domain
The set of all devices that will receive broadcast frames originating from any device within the set. Broadcast domains are typically bounded by routers because routers do not forward broadcast frames.

broadcast search
Propagation of a search request to all network nodes if the location of a resource is unknown to the requester. See also *directed search*.

broadcast storm
Undesirable network event in which many broadcasts are sent simultaneously across all network segments. A broadcast storm uses substantial network bandwidth and, typically, causes network time-outs.

buffer
Storage area used for handling data in transit. Buffers are used in internetworking to compensate for differences in processing speed between network devices. Bursts of data can be stored in buffers until they can be handled by slower processing devices. Sometimes referred to as a packet buffer.

burned-in address
See *universally administered address*.

BUS
Broadcast and unknown server. Multicast server used in ELANs that is used to flood traffic addressed to an unknown destination, and to forward multicast and broadcast traffic to the appropriate clients. See also *ELAN*.

C

cable
Transmission medium of copper wire or optical fiber wrapped in a protective cover.

call admission control
Traffic management mechanism used in ATM networks that determines whether the network can offer a path with sufficient bandwidth for a requested VCC.

CAC
Connection admission control. In ATM, the set of actions taken by the network during the call setup phase (or call renegotiation phase) in order to determine whether a connection request can be accepted or should be rejected.

Category 1 cabling
One of five grades of UTP cabling described in the EIA/TIA-586 standard. Category 1 cabling is used for telephone communications and is not suitable for transmitting data. Compare with Category 2 cabling, Category 3 cabling, Category 4 cabling, and Category 5 cabling. See also *EIA/TIA-586* and *UTP*.

Category 2 cabling
One of five grades of UTP cabling described in the EIA/TIA-586 standard. Category 2 cabling is capable of transmitting data at speeds up to 4 Mbps. Compare with Category 1 cabling, Category 3 cabling, Category 4 cabling, and Category 5 cabling. See also *EIA/TIA-586* and *UTP*.

Category 3 cabling
One of five grades of UTP cabling described in the EIA/TIA-586 standard. Category 3 cabling is used in 10BaseT networks and can transmit data at speeds up to 10 Mbps. Compare with *Category 1 cabling, Category 2 cabling, Category 4 cabling, and Category 5 cabling. See also EIA/TIA-586* and *UTP*.

Category 4 cabling
One of five grades of UTP cabling described in the EIA/TIA-586 standard. Category 4 cabling is used in Token Ring networks and can transmit data at speeds up to 16 Mbps. Compare with Category 1 cabling, Category 2 cabling, Category 3 cabling, and Category 5 cabling. See also *EIA/TIA-586* and *UTP*.

Category 5 cabling
One of five grades of UTP cabling described in the EIA/TIA-586 standard. Category 5 cabling can transmit data at speeds up to 100 Mbps. Compare with Category 1 cabling, Category 2 cabling, Category 3 cabling, and Category 4 cabling. See also *EIA/TIA-586* and *UTP*.

CAU
Controlled access unit. A microprocessor-controlled wiring concentrator that is used to form classical Token Rings and that provides management capabilities not available with unpowered, passive MAUs.

CBR
Constant bit rate. QOS class defined by the ATM Forum for ATM networks. CBR is used for connections that depend on precise clocking to ensure undistorted delivery. Compare with *ABR*, *UBR*, and *VBR*.

CDP
Cisco Discovery Protocol. A protocol that runs on Cisco devices (including routers, bridges, access servers, and switches) that allows Cisco network management applications to learn the device type and SNMP agent address of neighboring devices. CDP runs at Layer 2 and is media- and network-layer independent, allowing network management to be performed from a system that supports a different network-layer protocol from that being managed.

CDV
Cell delay variation. A component of cell transfer delay, which is induced by buffering and cell scheduling. CDV is a QOS delay parameter associated with CBR and VBR service. See also *CBR* and *VBR*.

CDVT
Cell delay variation tolerance. In ATM, a QOS parameter for managing traffic that is specified when a connection is set up. In CBR transmissions, CDVT determines the level of jitter that is tolerable for the data samples taken by the PCR. See also *CBR*.

cell
The basic data unit for ATM switching and multiplexing. Cells contain identifiers that specify the data stream to which they belong. Each cell consists of a 5-byte header and 48 bytes of payload. See also *cell relay.*

cell delay variation
See *CDV.*

cell delay variation tolerance
See *CDVT.*

cell loss priority
See *CLP.*

cell loss ratio
See *CLR*.

cell payload scrambling
Technique used an ATM switch to maintain framing on some medium-speed edge and trunk interfaces.

cell relay
Network technology based on the use of small, fixed-size packets, or cells. Because cells are fixed-length, they can be processed and switched in hardware at high speeds. Cell relay is the basis for many high-speed network protocols including ATM, IEEE 802.6, and SMDS. See also cell.

cell transfer delay
See *CTD*.

CER
Cell error ratio. In ATM, the ratio of transmitted cells that have errors to the total cells sent in a transmission for a specific period of time.

checksum
Method for checking the integrity of transmitted data. A checksum is an integer value computed from a sequence of octets taken through a series of arithmetic operations. The value is recomputed at the receiving end and compared for verification.

Cisco Discovery Protocol
See *CDP*.

CiscoWorks for Switched Internetworks
See *CWSI*.

circuit
Communications path between two or more points.

circuit switching
Switching system in which a dedicated physical circuit path must exist between sender and receiver for the duration of the "call." Used heavily in the telephone company network. Circuit switching can be contrasted with contention and token passing as a channel-access method, and with message switching and packet switching as a switching technique.

CLI
Command line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. The UNIX operating system and DOS provide CLIs. Compare with *GUI*.

client
Node or software program (front-end device) that requests services from a server. See also back end, front end, and server.

client/server computing
Term used to describe distributed computing (processing) network systems in which transaction responsibilities are divided into two parts: client (front end) and server (back end). Both terms (client and server) can be applied to software programs or actual computing devices.

CLP
Cell loss priority. Field in the ATM cell header that determines the probability of a cell being dropped if the network becomes congested. Cells with CLP = 0 are insured traffic, which is unlikely to be dropped. Cells with CLP = 1 are best-effort traffic, which might be dropped in congested conditions in order to free up resources to handle insured traffic.

CLR
Cell loss ratio. In ATM, the ratio of discarded cells to cells that are successfully transmitted. CLR can be set as a QOS parameter when a connection is set up.

collapsed backbone
Nondistributed backbone in which all network segments are interconnected by way of an internetworking device. A collapsed backbone might be a virtual network segment existing in a device such as a hub, a router, or a switch.

community
In SNMP, a logical group of managed devices and NMSs in the same administrative domain.

community string
Text string that acts as a password and is used to authenticate messages sent between a management station and a router containing an SNMP agent. The community string is sent in every packet between the manager and the agent. Also called a community name.

compression
The running of a data set through an algorithm that reduces the space required to store or the bandwidth required to transmit the data set.

concentrator
See *hub*.

concentrator relay function
See *CRF*.

congestion
Traffic in excess of network capacity.

congestion avoidance
The mechanism by which an ATM network controls traffic entering the network to minimize delays. In order to use resources most efficiently, lower-priority traffic is discarded at the edge of the network if conditions indicate that it cannot be delivered.

congestion collapse
A condition in which the re-transmission of frames in an ATM network results in little or no traffic successfully arriving at the destination. Congestion collapse frequently occurs in ATM networks composed of switches that do not have adequate and effective buffering mechanisms complemented by intelligent packet discard or ABR congestion feedback mechanisms.

connectionless
Term used to describe data transfer without the existence of a virtual circuit. Compare with connection-oriented. See also *virtual circuit*.

connection-oriented
Term used to describe data transfer that requires the establishment of a virtual circuit. See also *connectionless* and *virtual circuit*.

console
DTE through which commands are entered into a host.

constant bit rate
See *CBR*.

control direct VCC
In ATM, a bidirectional VCC set up by a LEC to a LES. One of three control connections defined by Phase 1 LANE. Compare with configuration direct VCC and control distribute VCC.

control distribute VCC

In ATM, a unidirectional VCC set up from a LES to a LEC. One of three control connections defined by Phase 1 LANE. Typically, the VCC is a point-to-multipoint connection. Compare with configuration direct VCC and control direct VCC.

convergence

The speed and ability of a group of internetworking devices running a specific routing protocol to agree on the topology of an internetwork after a change in that topology.

cost

Arbitrary value, typically based on hop count, media bandwidth, or other measures, that is assigned by a network administrator and used to compare various paths through an internetwork environment. Cost values are used by routing protocols to determine the most favorable path to a particular destination: the lower the cost, the better the path. Sometimes called path cost.

CRC

Cyclic redundancy check. Error-checking technique in which the frame recipient calculates a remainder by dividing frame contents by a prime binary divisor and compares the calculated remainder to a value stored in the frame by the sending node.

CRF

Concentrator relay function. As defined by the IEEE, a logical grouping of ports on a Token Ring switch with the same ring number. Within a CRF, source-route switching is used to forward frames within a port group. Multiple CRFs may exist within a switch. The BRF forwards frames between CRFs. See also *BRF*.

CRM

Cell rate margin. One of three link attributes exchanged using PTSPs to determine the available resources of an ATM network. CRM is a measure of the difference between the effective bandwidth allocation per traffic class as the allocation for sustainable cell rate.

CS

Convergence sublayer. One of the two sublayers of the AAL CPCS, responsible for padding and error checking. PDUs passed from the SSCS are appended with an 8-byte trailer (for error checking and other control information) and padded, if necessary, so that the length of the resulting PDU is divisible by 48. These PDUs are then passed to the SAR sublayer of the CPCS for further processing.

CTD

Cell transfer delay. In ATM, the elapsed time between a cell exit event at the source UNI and the corresponding cell entry event at the destination UNI for a particular connection. The CTD between the two points is the sum of the total inter-ATM node transmission delay and the total ATM node processing delay.

cut-through switching

Switching approach that streams data through a switch so that the leading edge of a packet exits the switch at the output port before the packet finishes entering the input port. A device using cut-through packet switching reads, processes, and forwards packets as soon as the destination address is looked up, and the outgoing port determined. Also known as on-the-fly packet switching. Compare with *store-and-forward*. See also *adaptive cut-through*.

CWSI

CiscoWorks for Switched Internetworks. A grouping of advanced network management capabilities for switched networks that includes TrafficDirector, CiscoView, and VlanDirector.

D

data direct VCC

In ATM, a bidirectional point-to-point VCC set up between two LECs. One of three data connections defined by Phase 1 LANE. Data direct VCCs do not offer any type of QOS guarantee, so they are typically used for UBR and ABR connections.

data link layer

Layer 2 of the OSI reference model. This layer provides reliable transit of data across a physical link. The data link layer is concerned with physical addressing, network topology, line discipline, error notification, ordered delivery of frames, and flow control. The IEEE has divided this layer into two sublayers: the MAC sublayer and the LLC sublayer. Sometimes simply called *link layer.* Roughly corresponds to the *data link control layer* of the SNA model.

data terminal equipment

See *DTE*.

data terminal ready

See *DTR*.

DB connector

Data bus connector. Type of connector used to connect serial and parallel cables to a data bus. DB connector names are of the format DB-*x*, where *x* represents the number of wires within the connector. Each line is connected to a pin on the connector, but in many cases, not all pins are assigned a function. DB connectors are defined by various EIA/TIA standards.

dedicated Token Ring

See *DTR*.

delay

The time between the initiation of a transaction by a sender and the first response received by the sender. Also, the time required to move a packet from source to destination over a given path.

designated bridge

The bridge that incurs the lowest path cost when forwarding a frame from a segment to the root bridge.

destination address

Address of a network device that is receiving data. See also *source address*.

destination MAC

See *DMAC*.

destination service access point

See *DSAP*.

differential Manchester encoding

Digital coding scheme where a mid-bit-time transition is used for clocking, and a transition at the beginning of each bit time denotes a zero. The coding scheme used by IEEE 802.5 and Token Ring networks.

directed search

Search request sent to a specific node known to contain a resource. A directed search is used to determine the continued existence of the resource and to obtain routing information specific to the node. See also *broadcast search*.

DRAM

Dynamic random-access memory. RAM that stores information in capacitors that must be periodically refreshed. Delays can occur because DRAMs are inaccessible to the processor when refreshing their contents. However, DRAMs are less complex and have greater capacity than SRAMs. See also *SRAM*.

DSAP

Destination service access point. One-byte fields in the LLC protocol data unit of 802.2 frames that specifies the sending (SSAP) and receiving (DSAP) network-layer processes between which the frame is being transferred. Both the DSAP and SSAP numbers are assigned by the IEEE. Compare to *SSAP*. See also *SAP*.

DTE

Data terminal equipment. Device at the user end of a user-network interface that serves as a data source, destination, or both. DTE connects to a data network through a DCE device (for example, a modem) and typically uses clocking signals generated by the DCE. DTE includes such devices as computers, protocol translators, and multiplexers.

DTR

1. Data terminal ready. EIA/TIA-232 circuit that is activated to let the DCE know when the DTE is ready to send and receive data.

2. Dedicated Token Ring. A specification defined by the Token Ring standard (IEEE 802.5r) standard. The DTR standard has two components: concentrator port, or C-port, capability and full-duplex operation. The C-port capability enables any Token Ring NIC to directly attach to a switch port. The full-duplex support enables 802.5r-compliant NICs to simultaneously transmit and receive, for an aggregate of 32 Mbps. When operating in full-duplex mode, the station and switch use a protocol called Transmit Immediate (TXI) and do not use a token. This allows Token Ring NICs to be connected without a concentrator (e.g., a MAU).

DRiP

Duplicate Ring Protocol. A Cisco-developed protocol that allows the management of ring numbers across multiple, interconnected switches.

Duplicate Ring Protocol

See *DRiP*.

E

early token release
Technique used in Token Ring networks that allows a station to release a new token onto the ring immediately after transmitting, instead of waiting for the first frame to return. This feature can increase the total bandwidth on the ring. See also *Token Ring*.

EEPROM
Electrically erasable programmable read-only memory. EPROM that can be erased using electrical signals applied to specific pins. See also *EPROM*.

EIA
Electronic Industries Association. Group that specifies electrical transmission standards. The EIA and TIA have developed numerous well-known communications standards, including EIA/TIA-232 and EIA/TIA-449. See also *TIA*.

EIA/TIA-232
Common physical layer interface standard, developed by EIA and TIA, that supports unbalanced circuits at signal speeds of up to 64 kbps. Closely resembles the V.24 specification. Formerly known as *RS-232*.

EIA/TIA-586
Standard that describes the characteristics and applications for various grades of UTP cabling. See also *Category 1 cabling*, *Category 2 cabling*, *Category 3 cabling*, *Category 4 cabling*, and *Category 5 cabling*.

ELAN
Emulated LAN. ATM network in which an Ethernet or Token Ring LAN is emulated using a client-server model. ELANs are composed of an LEC, an LES, a BUS, and an LECS. Multiple ELANs can exist simultaneously on a single ATM network. ELANs are defined by the LANE specification. See also *BUS*, *LANE*, *LEC*, *LECS*, and *LES*.

electrically erasable programmable read-only memory
See *EEPROM*.

electromagnetic interference
See *EMI*.

Electronic Industries Association
See *EIA*.

electrostatic discharge
See *ESD*.

EMI
Electromagnetic interference. Interference by electromagnetic signals that can cause reduced data integrity and increased error rates on transmission channels.

emulated LAN
See *ELAN*.

encapsulation

The wrapping of data in a particular protocol header. For example, Ethernet data is wrapped in a specific Ethernet header before network transit. Also, when bridging dissimilar networks, the entire frame from one network is simply placed in the header used by the data link layer protocol of the other network.

EPROM

Erasable programmable read-only memory. Nonvolatile memory chips that are programmed after they are manufactured, and, if necessary, can be erased by some means and reprogrammed. Compare with *EEPROM*.

erasable programmable read-only memory

See *EPROM*.

ESD

Electrostatic discharge. Discharge of stored static electricity that can damage electronic equipment and impair electrical circuitry, resulting in complete or intermittent failures.

Ethernet

Baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use CSMA/CD and run over a variety of cable types at 10 Mbps. Ethernet is similar to the IEEE 802.3 series of standards. See also *10Base2*, *10Base5*, *10BaseF*, *10BaseT*, *10Broad36*, *Fast Ethernet*, and *IEEE 802.3*.

excess rate

In ATM, traffic in excess of the insured rate for a given connection. Specifically, the excess rate equals the maximum rate minus the insured rate. Excess traffic is delivered only if network resources are available and can be discarded during periods of congestion. Compare with *maximum rate*.

explorer frame

Frame sent out by a networked device in an SRB environment to determine the optimal route to another networked device. It gathers a hop-by-hop description of a path through the network by being marked (updated) by each bridge that it traverses, thereby creating a complete topological map. See also *all-routes explorer* and *spanning-tree explorer.*

F

failure domain

Area in which a failure has occurred in a Token Ring, defined by the information contained in a beacon. When a station detects a serious problem with the network (such as a cable break), it sends a beacon frame that includes the station reporting the failure, its NAUN, and everything in between. Beaconing in turn initiates a process called autoreconfiguration. See also *beacon* and *NAUN*.

Fast Ethernet

Any of a number of 100-Mbps Ethernet specifications. Fast Ethernet offers a speed increase ten times that of the 10BaseT Ethernet specification, while preserving such qualities as frame format, MAC mechanisms, and MTU. Such similarities allow the use of existing 10BaseT applications and network management tools on Fast Ethernet networks. Based on an extension to the IEEE 802.3 specification. Compare with *Ethernet*.

FC
Frame control. The portion of a frame that indicates the frame type.

FCS
Frame check sequence. Refers to the extra characters added to a frame for error control purposes. Used in HDLC, Frame Relay, and other data link layer protocols.

FDDI
Fiber Distributed Data Interface. LAN standard, defined by ANSI X3T9.5, specifying a 100-Mbps token-passing network using fiber-optic cable, with transmission distances of up to 2 km. FDDI uses a dual-ring architecture to provide redundancy.

FDX
Full duplex. Capability for simultaneous data transmission between a sending station and a receiving station. Compare with *half duplex* and *simplex*.

Fiber Distributed Data Interface
See *FDDI*.

fiber-optic cable
Physical medium capable of conducting modulated light transmission. Compared with other transmission media, fiber-optic cable is more expensive, but is not susceptible to electromagnetic interference, and is capable of higher data rates. Sometimes called *optical fiber.*

filter
Generally, a process or device that screens network traffic for certain characteristics, such as source address, destination address, or protocol, and determines whether to forward or discard that traffic based on the established criteria.

Flash memory
Nonvolatile storage that can be electrically erased and reprogrammed so that software images can be stored, booted, and rewritten as necessary. Flash memory was developed by Intel and is licensed to other semiconductor companies.

flooding
Traffic passing technique used by switches and bridges in which traffic received on an interface is sent out all of the interfaces of that device except the interface on which the information was originally received.

forwarding
Process of sending a frame toward its ultimate destination by way of an internetworking device.

fragmentation
Process of breaking a packet into smaller units when transmitting over a network medium that cannot support the original size of the packet. See also *reassembly*.

frame
Logical grouping of information sent as a data link layer unit over a transmission medium. Often refers to the header and trailer, used for synchronization and error control, that surround the user data contained in the unit. The terms *cell, datagram, message, packet,* and *segment* are also used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.

frame check sequence
See *FCS*.

frame control
See *FC*.

full duplex
See *FDX*.

G

gigabit
Abbreviated *Gb*.

gigabits per second
Abbreviated *Gbps*.

gigabyte
Abbreviated *GB*.

gigabytes per second
Abbreviated *GBps*.

graphical user interface
See *GUI*.

group address
See *multicast address*.

GUI
Graphical user interface. User environment that uses pictorial as well as textual representations of the input and output of applications and the hierarchical or other data structure in which information is stored. Conventions such as buttons, icons, and windows are typical, and many actions are performed using a pointing device (such as a mouse). Microsoft Windows and the Apple Macintosh are prominent examples of platforms utilizing a GUI.

H

half duplex
See *HDX*.

hardware address
See *MAC address*.

HDX
Half duplex. Capability for data transmission in only one direction at a time between a sending station and a receiving station. BSC is an example of a half-duplex protocol. Compare with *full duplex* and *simplex*.

hop count
Routing metric used to measure the distance between a source and a destination. RIP uses hop count as its sole metric.

hot swapping
See *OIR* and *power-on servicing*.

hub
1. Generally, a term used to describe a device that serves as the center of a star-topology network.
2. Hardware or software device that contains multiple independent but connected modules of network and internetwork equipment. Hubs can be active (where they repeat signals sent through them) or passive (where they do not repeat, but merely split, signals sent through them).
3. In Ethernet and IEEE 802.3, an Ethernet multiport repeater, sometimes referred to as a concentrator.

I

IEEE
Institute of Electrical and Electronics Engineers. Professional organization whose activities include the development of communications and network standards. IEEE LAN standards are the predominant LAN standards today.

IEEE 802.1
IEEE specification that describes an algorithm that prevents bridging loops by creating a spanning tree. The algorithm was invented by Digital Equipment Corporation. The Digital algorithm and the IEEE 802.1 algorithm are not exactly the same, nor are they compatible. See also *spanning tree*, *spanning-tree algorithm*, and *Spanning-Tree Protocol*.

IEEE 802.2
IEEE LAN protocol that specifies an implementation of the LLC sublayer of the data link layer. IEEE 802.2 handles errors, framing, flow control, and the network layer (Layer 3) service interface. Used in IEEE 802.3 and IEEE 802.5 LANs. See also *IEEE 802.3* and *IEEE 802.5*.

IEEE 802.3
IEEE LAN protocol that specifies an implementation of the physical layer and the MAC sublayer of the data link layer. IEEE 802.3 uses CSMA/CD access at a variety of speeds over a variety of physical media. Extensions to the IEEE 802.3 standard specify implementations for Fast Ethernet.

IEEE 802.5
IEEE LAN protocol that specifies an implementation of the physical layer and MAC sublayer of the data link layer. IEEE 802.5 uses token passing access at 4 or 16 Mbps over STP cabling and is similar to IBM Token Ring. See also *Token Ring*.

internetwork
Collection of networks interconnected by routers and other devices that functions (generally) as a single network. Sometimes called an *internet*, which is not to be confused with the *Internet*.

internetworking
General term used to refer to the industry that has arisen around the problem of connecting networks together. The term can refer to products, procedures, and technologies.

interoperability
Ability of computing equipment manufactured by different vendors to communicate with one another successfully over a network.

inter-switch link
See *ISL*.

IP address
A 32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, while the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. CIDR provides a new way of representing IP addresses and subnet masks. Also called an *Internet address*. See also *subnet mask*.

ISL
Inter-Switch Link. A Cisco-defined protocol that enables full-length frames from multiple Ethernet or Token Ring VLANs to be transmitted simultaneously across the same 100 Mbps Fast Ethernet link. The ISL protocol is supported between Cisco switches and routers and servers using NICs that support ISL.proprietary link for interconnecting switches. ISL uses 100-Mbps Ethernet and allows the multiplexing of multiple VLANs over a single link.

ISL Channel
A parallel configuration of 2 to 4 ports between two Catalyst 3900s or between a Catalyst 3900 and a Catalyst 5000, a Token Ring ISL-capable Cisco router, or a Token Ring ISL network adapter.

K

KB
Kilobyte. Approximately 1,000 bytes.

Kb
Kilobit. Approximately 1,000 bits.

kBps
Kilobytes per second.

kbps
Kilobits per second.

keepalive interval
Period of time between each keepalive message sent by a network device.

keepalive message
Message sent by one network device to inform another network device that the virtual circuit between the two is still active.

kilobit
Abbreviated *Kb*.

kilobits per second
Abbreviated *kbps*.

kilobyte
Abbreviated *KB*.

kilobytes per second
Abbreviated *kBps*

L

LAA
Locally administered address. A MAC address assigned to an interface that overrides the factory-assigned universally administered address. Assigning an LAA eases network management because the NIC can be replaced without changing the address used by the network to access the station. See also *MAC address*. Compare to *universally administered address*.

LAN
Local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.

LANE
LAN emulation. Technology that allows an ATM network to function as a LAN backbone. The ATM network must provide multicast and broadcast support, address mapping (MAC-to-ATM), SVC management, and a usable packet format. LANE also defines Ethernet and Token Ring ELANs. See also *ELAN*.

LAN emulation
See *LANE*.

LAN Emulation Client
See *LEC*.

LAN Emulation Address Resolution Protocol
See *LE_ARP*.

LAN Emulation Configuration Server
See *LECS*.

LAN Emulation Server
See *LES*.

LAN Network Manager
See *LNM*.

LAN switch

High-speed switch that forwards packets between data-link segments. Most LAN switches forward traffic based on MAC addresses. This variety of LAN switch is sometimes called a *frame switch*. LAN switches are often categorized according to the method they use to forward traffic: cut-through packet switching or store-and-forward packet switching. Multilayer switches are an intelligent subset of LAN switches. Compare with *multilayer switch*.

laser

Light amplification by stimulated emission of radiation. Analog transmission device in which a suitable active material is excited by an external stimulus to produce a narrow beam of coherent light that can be modulated into pulses to carry data. Networks based on laser technology are sometimes run over SONET.

latency

The delay associated with the physical transfer of a frame from one port through a switch to another port, which is based on the switch architecture and adds additional delay above and beyond the delay associated with the physical length of the frame being transported through the switch. The latency of a switch would be the time between the first bit of a frame into the switch and the first bit of that frame out of the switch.

LE_ARP

LAN Emulation Address Resolution Protocol. A protocol that provides the ATM address that corresponds to a MAC address.

LEC

LAN Emulation Client. Entity in an end system that performs data forwarding, address resolution, and other control functions for a single ES within a single ELAN. A LEC also provides a standard LAN service interface to any higher-layer entity that interfaces to the LEC. Each LEC is identified by a unique ATM address, and is associated with one or more MAC addresses reachable through that ATM address. See also *ELAN* and *LES*.

LECS

LAN Emulation Configuration Server. Entity that assigns individual LANE clients to particular ELANs by directing them to the LES that corresponds to the ELAN. There is logically one LECS per administrative domain, and this serves all ELANs within that domain. See also *ELAN*.

LED

Light emitting diode. Semiconductor device that emits light produced by converting electrical energy. Status lights on hardware devices are typically LEDs.

LES

LAN Emulation Server. Entity that implements the control function for a particular ELAN. There is only one logical LES per ELAN, and it is identified by a unique ATM address. See also *ELAN*.

light amplification by stimulated emission of radiation

See *laser*.

light emitting diode

See *LED*.

LLC
Logical Link Control. Higher of the two data link layer sublayers defined by the IEEE. The LLC sublayer handles error control, flow control, framing, and MAC-sublayer addressing. The most prevalent LLC protocol is IEEE 802.2, which includes both connectionless and connection-oriented variants. See also *data link layer* and *MAC*.

LLC2
Logical Link Control, type 2. Connection-oriented OSI LLC-sublayer protocol. See also *LLC*.

LNM
LAN Network Manager. IBM's PC-based Token Ring network management product. LNM establishes LLC2 connections to IBM bridges within the Token Ring network and uses agents in these bridges to monitor and collect MAC-layer information about the ring, as well as to issue MAC-layer commands to the ring-attached NICs for control and to request NIC-level status information.

load balancing
In routing, the ability of a router to distribute traffic over all its network ports that are the same distance from the destination address. Good load-balancing algorithms use both line speed and reliability information. Load balancing increases the utilization of network segments, thus increasing effective network bandwidth.

local-area network
See *LAN*.

locally administered address
See *LAA*.

Logical Link Control
See *LLC*.

Logical Link Control, type 2
See *LLC2*.

LUNI
LAN Emulation User-to-Network Interface. The ATM Forum standard for LAN emulation on ATM networks. LUNI defines the interface between the LAN Emulation Client (LEC) and the LAN Emulation Server components. See also *BUS*, *LES*, and *LECS*.

M

MAC
Media access control. Lower of the two sublayers of the data link layer defined by the IEEE. The MAC sublayer handles access to shared media, such as whether token passing or contention will be used. See also *data link layer* and *LLC*.

MAC address

Standardized data link layer address that is required for every port or device that connects to a LAN. This is, in essence, the address of the NIC. Both the transmitting station's MAC address as well as the destination station's MAC address are contained in all LAN frames. MAC addresses are 6 bytes long and are controlled by the IEEE. Also known as a *hardware address*, a *MAC-layer address*, or a *physical address*. See also *universally administered address* and *locally administered address*.

MAC address learning

Service that characterizes a learning bridge, in which the source MAC address of each received packet is stored so that future packets destined for that address can be forwarded only to the bridge interface on which that address is located. Packets destined for unrecognized addresses are forwarded out every bridge interface. This scheme helps minimize traffic on the attached LANs. MAC address learning is defined in the IEEE 802.1 standard. See also *MAC address*.

management domain

A logical grouping of VLANs used by the VLAN Trunking Protocol (VTP) for the purpose of administration and management. VTP parameters are propagated throughout the VLANs within a single management domain. While you can have duplicate VLAN names in a network, each VLAN name within a management domain must be unique. A management domain is not device specific. Different devices may belong to the same management domain if the VLANs defined for the devices belong to the same management domain. Likewise, a device may belong to multiple management domains if the VLANs defined for the device belong to different management domains.

Management Information Base

See *MIB*.

Manchester encoding

Digital coding scheme, used by IEEE 802.3 and Ethernet, in which a mid-bit-time transition is used for clocking, and a 1 is denoted by a high level during the first half of the bit time.

MAU

Media attachment unit. Device used in Ethernet and IEEE 802.3 networks that provides the interface between the AUI port of a station and the common medium of the Ethernet. The MAU, which can be built into a station or can be a separate device, performs physical layer functions including the conversion of digital data from the Ethernet interface, collision detection, and injection of bits onto the network. Sometimes referred to as a *media access unit*, also abbreviated MAU, or as a transceiver.

In Token Ring, a MAU is known as a *multistation access unit* and is usually abbreviated MSAU to avoid confusion. In Token Ring networks the MAU is a nonpowered device used for forming a classical Token Ring. NICs are attached to the MAU ports via lobe cables. When activated, the NIC provides power (called "phantom drive") to the MAU via its lobe cable to transfer relays in the MAU that cause both the NIC and its lobe cable to be electrically and logically inserted in the ring. When deactivated, the NIC removes the phantom drive voltage, which causes the MAU to electrically bypass the NIC and its lobe cable, allowing the ring to continue via the bypass.

maximum burst
Specifies the largest burst of data above the insured rate that will be allowed temporarily on an ATM PVC, but will not be dropped at the edge by the traffic policing function, even if it exceeds the maximum rate. This amount of traffic will be allowed only temporarily; on average, the traffic source needs to be within the maximum rate. Specified in bytes or cells. See also *maximum rate*.

maximum rate
Maximum total data throughput allowed on a given virtual circuit, equal to the sum of the insured and uninsured traffic from the traffic source. The uninsured data might be dropped if the network becomes congested. The maximum rate, which cannot exceed the media rate, represents the highest data throughput the virtual circuit will ever deliver, measured in bits or cells per second. See also *maximum burst*.

maximum transmission unit
See *MTU*.

MB
Megabyte. Approximately 1,000,000 bytes.

Mb
Megabit. Approximately 1,000,000 bits.

Mbps
Megabits per second.

Media Access Control
See *MAC*.

media access unit
See *MAU*.

media attachment unit
See *MAU*.

megabit
Abbreviated *Mb*. Approximately 1,000,000 bits.

megabits per second
Abbreviated *Mbps*.

megabyte
Abbreviated *MB*. Approximately 1,000,000 bytes.

MIB
Management Information Base. Database of network management information that is used and maintained by a network management protocol such as SNMP or CMIP. The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a GUI network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.

microsegmentation

The process of reconfiguring stations from a shared-media LAN to smaller groups, usually one per segment, using LAN switches. A classic Token Ring network, for example, can have as many as 260 stations attached to, and sharing the bandwidth of, a single 16-Mbps ring. Using Token Ring switching, these stations can be segmented into groups of as few as one per switch port. Each group then has a bandwidth of 16 Mbps and, if desired, the affinity grouping of the prior ring can be maintained using VLAN groupings in the switch.

modem

Modulator-demodulator. Device that converts digital and analog signals. At the source, a modem converts digital signals to a form suitable for transmission over analog communication facilities. At the destination, the analog signals are returned to their digital form. Modems allow data to be transmitted over voice-grade telephone lines.

MPOA

Multiprotocol over ATM. A relatively new standardization effort in the ATM Forum that will specify how existing and future network-layer protocols will exploit the unique benefits of ATM. These benefits include quality of service (QOS) and direct connections between different VLANs.

MSAU

Multistation access unit. See *MAU*.

MTU

Maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle.

multicast

Single packets copied by the network and sent to a specific subset of network addresses. These addresses are specified in the destination address field. Compare with *broadcast* and *unicast*.

multicast address

Single address that refers to multiple network devices. Synonymous with *group address*. Compare with *broadcast address* and *unicast address*. See also *multicast*.

multilayer switch

Switch that filters and forwards packets based on MAC addresses and network addresses. A subset of LAN switch. Compare with *LAN switch*.

multiprotocol over ATM

See *MPOA*.

multistation access unit

See *MAU* or *MSAU*.

N

NADN
Nearest active downstream neighbor. In Token Ring or IEEE 802.5 networks, the closest downstream network device from any given device that is still active.

NAUN
Nearest active upstream neighbor. In Token Ring or IEEE 802.5 networks, the closest upstream network device from any given device that is still active.

NDIS
Network driver interface specification. Microsoft's specification for a generic, hardware- and protocol-independent device driver for NICs.

nearest active downstream neighbor
See *NADN*.

nearest active upstream neighbor
See *NAUN*.

NetBEUI
NetBIOS Extended User Interface. An enhanced version of the NetBIOS protocol used by network operating systems such as LAN Manager, LAN Server, Windows for Workgroups and Windows NT. NetBEUI formalizes the transport frame and adds additional functions. NetBEUI implements the OSI LLC2 protocol. See also *LLC2*.

NetBIOS
Network Basic Input/Output System. API used by applications on an IBM LAN to request services from lower-level network processes. These services might include session establishment and termination, and information transfer.

NetBIOS Extended User Interface
See *NetBEUI*.

network
Collection of computers, printers, routers, switches, and other devices that are able to communicate with each other over some transmission medium.

network analyzer
Hardware or software device offering various network troubleshooting features, including protocol-specific packet decodes, specific preprogrammed troubleshooting tests, packet filtering, and packet transmission.

Network Basic Input/Output System
See *NetBIOS*.

network driver interface specification
See *NDIS*.

network interface card
See *NIC*.

Network-to-Network Interface

See *NNI*.

NIC

Network interface card. Board that provides network communication capabilities to and from a computer system. Also called an *adapter*.

NMS

Network management system. System responsible for managing at least part of a network. An NMS is generally a reasonably powerful and well-equipped computer such as an engineering workstation. NMSs communicate with agents to help keep track of network statistics and resources.

NNI

Network-to-Network Interface. ATM Forum standard that defines the interface between two ATM switches that are both located in a private network or are both located in a public network. The interface between a public switch and private one is defined by the UNI standard. Also, the standard interface between two Frame Relay switches meeting the same criteria. Compare with *UNI*.

nonvolatile random-access memory

See *NVRAM*.

null modem

Small box or cable used to join computing devices directly, rather than over a network.

NVRAM

Nonvolatile RAM. RAM that retains its contents when a unit is powered off.

O

OC

Optical Carrier. Series of physical protocols (OC-1, OC-2, OC-3, and so on), defined for SONET optical signal transmissions. OC signal levels put STS frames onto multimode fiber-optic line at a variety of speeds. The base rate is 51.84 Mbps (OC-1); each signal level thereafter operates at a speed divisible by that number (thus, OC-3 runs at 155.52 Mbps). See also *SONET*.

OIR

Online insertion and removal. Feature that permits the addition, replacement, or removal of cards without interrupting the system power, entering console commands, or causing other software or interfaces to shut down. Sometimes called *hot swapping or power-on servicing*.

online insertion and removal

See *OIR*.

Optical Carrier

See *OC*.

optical fiber

See *fiber-optic cable*.

P

packet
Logical grouping of information that includes a header containing control information and (usually) user data. Packets are most often used to refer to network layer units of data. The terms *datagram*, *frame*, *message*, and *segment* are also used to describe logical information groupings at various layers of the OSI reference model and in various technology circles. See also *PDU*.

packet switching
Networking method in which nodes share bandwidth with each other by sending packets.

PAD
Packet assembler/disassembler. Device used to connect simple devices (like character-mode terminals) that do not support the full functionality of a particular protocol to a network. PADs buffer data and assemble and disassemble packets sent to such end devices.

passive port monitor
A type of monitoring supported by SPAN that monitor allows you to monitor all the frames on a particular ring, including the MAC frames.

payload
Portion of a cell, frame, or packet that contains upper-layer information (data).

payload type identifier
See *PTI*.

PCR
Peak cell rate. Parameter defined by the ATM Forum for ATM traffic management. In CBR transmissions, PCR determines how often data samples are sent. In ABR transmissions, PCR determines the maximum value of the ACR. See also *ABR (available bit rate)* and *CBR*.

PDU
Protocol data unit. OSI term for packet. See also *BPDU* and *packet*.

peak cell rate
See *PCR*.

peak rate
Maximum rate, in kilobits per second, at which a virtual circuit can transmit.

permanent virtual circuit
See *PVC*.

PFP
Proprietary fat pipe. An interface from a switch to a Cisco ProStack port. Switches can be connected together using the 140-Mbps full-duplex ProStack and function as one operational system.

physical address
See *MAC address*.

physical layer

Layer 1 of the OSI reference model. The physical layer defines the electrical, mechanical, procedural and functional specifications for activating, maintaining, and deactivating the physical link between end systems. Corresponds with the *physical control layer* in the SNA model.

PNNI

1. Private Network-Network Interface. ATM Forum specification for distributing topology information between switches and clusters of switches that is used to compute paths through the network. The specification is based on well-known link-state routing techniques and includes a mechanism for automatic configuration in networks in which the address structure reflects the topology.

2. *Private Network Node Interface*. ATM Forum specification for signaling to establish point-to-point and point-to-multipoint connections across an ATM network. The protocol is based on the ATM Forum's UNI specification with additional mechanisms for source routing, crankback, and alternate routing of call setup requests.

point-to-multipoint connection

One of two fundamental connection types. In ATM, a point-to-multipoint connection is a unidirectional connection in which a single source end-system (known as a root node) connects to multiple destination end-systems (known as leaves). Compare with *point-to-point connection*.

point-to-point connection

One of two fundamental connection types. In ATM, a point-to-point connection can be a unidirectional or bidirectional connection between two ATM end-systems. Compare with *point-to-multipoint connection*.

port

Interface on an internetworking device (such as a switch).

POST

Power-on self test. Set of hardware diagnostics that runs on a hardware device when that device is powered up.

power-on self test

See *POST*.

power-on servicing

Feature that allows faulty components to be diagnosed, removed, and replaced while the rest of the device continues to operate normally. Sometimes abbreviated *POS*. Sometimes called *hot swapping*. See also *OIR*.

proprietary fat pipe

See *PFP*.

protocol data unit

See *PDU*.

PTI
Payload type identifier. A 3-bit descriptor in the ATM cell header indicating the type of payload that the cell contains. Payload types include user and management cells; one combination indicates that the cell is the last cell of an AAL5 frame.

PVC
Permanent virtual circuit. Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. Called a *permanent virtual connection* in ATM terminology. Compare with *SVC*.

Q

QMAC
Quad media access controller. An ASIC chip containing that contains four Token Ring protocol handlers. Together with the QTP chip it provides four distinct Token Ring attachment ports.

QOS
Quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.

QOS parameters
quality of service parameters. Parameters that control the amount of traffic the source in an ATM network sends over an SVC. If any switch along the path cannot accommodate the requested QOS parameters, the request is rejected, and a rejection message is forwarded back to the originator of the request.

quad media access controller
See *QMAC*.

quad Token Ring port
See *QTP*.

quality of service
See *QOS*.

QTP
Quad Token Ring port. An ASIC chip that provides the necessary functions for switching directly between the four Token Ring ports of a QMAC, or between these and any other port within the switch.

R

RAM
Random-access memory. Volatile memory that can be read and written by a microprocessor.

random-access memory
See *RAM*.

reassembly
The putting back together of an IP datagram at the destination after it has been fragmented either at the source or at an intermediate node. See also *fragmentation*.

redundant system
Computer, router, switch, or other system that contains two or more of each of the most important subsystems, such as two disk drives, two CPUs, or two power supplies.

remote monitoring
See *RMON*.

repeater
Device that regenerates and propagates electrical signals between two network segments. See also *segment*.

RIF
Routing information field. Field in the IEEE 802.5 header that is used by a source-route bridge to determine through which Token Ring network segments a packet must transit. A RIF is made up of ring and bridge numbers as well as other information.

RII
Routing information identifier. Bit used by SRT bridges to distinguish between frames that should be transparently bridged and frames that should be passed to the SRB module for handling.

ring
Connection of two or more stations in a logically circular topology. Information is passed sequentially between active stations. Token Ring, FDDI, and CDDI are based on this topology.

ring in/ring out
See *RI/RO*.

ring group
Collection of Token Ring interfaces on one or more routers that is part of a one-bridge Token Ring network.

ring latency
Time required for a signal to propagate once around a ring in a Token Ring or IEEE 802.5 network.

ring monitor
Centralized management tool for Token Ring networks based on the IEEE 802.5 specification. See also *active monitor* and *standby monitor.*

ring parameter server
See *RPS*.

ring topology
Network topology that consists of a series of repeaters connected to one another by unidirectional transmission links to form a single closed loop. Each station on the network connects to the network at a repeater. While logically a ring, ring topologies are most often organized in a closed-loop star.

RI/RO
Ring in, ring out. Connectors on a MAU (or CAU) used to cable multiple wiring concentrators in series to form a classical Token Ring of up to 260 NICs.

RMON
Remote monitoring. MIB agent specification described in RFC 1271 that defines functions for the remote monitoring of networked devices. The RMON specification provides numerous monitoring, problem detection, and reporting capabilities.

root bridge
Exchanges topology information with designated bridges in a spanning-tree implementation in order to notify all other bridges in the network when topology changes are required. This prevents loops and provides a measure of defense against link failure.

routing information field
See *RIF*.

routing information identifier
See *RII*.

RPS
Ring parameter server. A network management function that may reside on a Token Ring to provide a ring number, soft error report timer values, and physical location information in response to a Request Parameters MAC frame sent from a NIC during insertion into the ring.

RS-232
Popular physical layer interface. Now known as *EIA/TIA-232*. See *EIA/TIA-232*.

S

SAP
Service access point. Field defined by the IEEE 802.2 specification that is part of an address specification. Thus, the destination plus the DSAP define the recipient of a packet. The same applies to the SSAP. See also *DSAP* and *SSAP*.

SAR
Segmentation and reassembly. One of the two sublayers of the AAL CPCS, responsible for dividing (at the source) and reassembling (at the destination) the PDUs passed from the CS. The SAR sublayer takes the PDUs processed by the CS and, after dividing them into 48-byte pieces of payload data, passes them to the ATM layer for further processing.

SCR
Sustainable cell rate. Parameter defined by the ATM Forum for ATM traffic management. For VBR connections, SCR determines the long-term average cell rate that can be transmitted.

segment
Section of a network that is bounded by bridges, routers, or switches.

segmentation and reassembly
See *SAR*.

service access point
See *SAP*.

service specific convergence sublayer.
See *SSCS*.

shielded twisted-pair
See *STP*.

Simple Network Management Protocol
See *SNMP*.

simple server redundancy protocol
See *SSRP*.

simplex
Capability for transmission in only one direction between a sending station and a receiving station. Broadcast television is an example of a simplex technology. Compare with *full duplex* and *half duplex*.

single-route explorer packet
See *spanning-tree explorer packet*.

SMAC
Source media access control. MAC address specified in the Source Address field of a packet. Compare with *DMAC*. See also *MAC address*.

SMDS
Switched Multimegabit Data Service. High-speed, packet-switched, datagram-based WAN networking technology offered by the telephone companies.

SNAP
Subnetwork Access Protocol. Internet protocol that operates between a network entity in the subnetwork and a network entity in the end system. SNAP specifies a standard method of encapsulating IP datagrams and ARP messages on IEEE networks. The SNAP entity in the end system makes use of the services of the subnetwork and performs three key functions: data transfer, connection management, and QOS selection.

SNMP
Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security.

SNMP communities
Authentication scheme that enables an intelligent network device to validate SNMP requests.

SNMP2
SNMP Version 2. Version 2 of the popular network management protocol. SNMP2 supports centralized as well as distributed network management strategies, and includes improvements in the SMI, protocol operations, management architecture, and security. See also *SNMP*.

SONET
Synchronous Optical Network. High-speed (up to 2.5 Gbps) synchronous network specification developed by Bellcore and designed to run on optical fiber. STS-1 is the basic building block of SONET. Approved as an international standard in 1988.

source address
Address of a network device that is sending data. See also *destination address*.

source-route bridging
See *SRB*.

source-route translational bridging
See *SR/TLB*.

source-route transparent bridging
See *SRT*.

SPAN
Switched Port Analyzer. The SPAN port capability on Cisco switches provide the ability to mirror the traffic from any switch port to the SPAN. Network analyzers and RMON probes can be connected to the port for in-depth troubleshooting. one Token Ring port on a switch on another port, providing a powerful network troubleshooting tool.

spanning tree
Loop-free subset of a network topology. See also *spanning-tree algorithm* and *Spanning-Tree Protocol*.

spanning-tree algorithm
Algorithm used by the STP to create a spanning tree. Sometimes abbreviated *STA*. See also *spanning tree* and *Spanning-Tree Protocol*.

spanning-tree explorer packet
Follows a statically configured spanning tree when looking for paths in an SRB network. Also known as a *limited-route explorer packet* or a *single-route explorer packet*. See also *all-routes explorer.*

Spanning-Tree Protocol
See *STP*.

SRAM
Static random access memory. Type of RAM that retains its contents for as long as power is supplied. SRAM does not require constant refreshing, like DRAM. Compare with *DRAM*.

SRB
Source-route bridging. Method of bridging originated by IBM and popular in Token Ring networks. In a SRB network, the entire route to a destination is predetermined, in real time, prior to the sending of data to the destination. Contrast with *transparent bridging*.

SRT
Source-route transparent bridging. IBM bridging scheme that merges the two most prevalent bridging strategies, SRB and transparent bridging. SRT employs both technologies in one device to satisfy the needs of all ENs. No translation between bridging protocols is necessary. Compare with *SR/TLB*.

SR/TLB
Source-route translational bridging. Method of bridging where source-route stations can communicate with transparent bridge stations with the help of an intermediate bridge that translates between the two bridge protocols. Compare with *SRT*.

SSAP

Source service access point. The SAP of the network node designated in the Source field of a packet. Compare to *DSAP*. See also *SAP*.

SSCS

Service specific convergence sublayer. One of the two sublayers of any AAL. SSCS, which is service dependent, offers assured data transmission. The SSCS can be null as well, in classical IP over ATM or LAN emulation implementations.

SSRP

Simple server redundancy protocol. A Cisco value-add that provides backup capability for LANE 1.0 servers, including the LECS, LES, and BUS. LANE 1.0 did not specify a method for doing this.

standby monitor

Device placed in standby mode on a Token Ring network in case an active monitor fails. See also *active monitor* and *ring monitor*.

star topology

LAN topology in which end points on a network are connected to a common central switch by point-to-point links. A ring topology that is organized as a star implements a unidirectional closed-loop star, instead of point-to-point links. Compare with *ring topology*.

store-and-forward

Frame forwarding technique in which frames are completely processed before being forwarded out the appropriate port. This processing includes calculating the CRC and checking the destination address. In addition, frames must be temporarily stored until network resources (such as an unused link) are available to forward the message. Contrast with *cut-through*.

STP

1. Shielded twisted-pair. Two-pair wiring medium used in a variety of network implementations. STP cabling has a layer of shielded insulation to reduce EMI. Compare with *UTP*. See also *twisted pair*.

2. Spanning-Tree Protocol. Bridge protocol that utilizes the spanning-tree algorithm, enabling a learning bridge to dynamically work around loops in a network topology by creating a spanning tree. Bridges exchange BPDU messages with other bridges to detect loops, and then remove the loops by shutting down selected bridge interfaces. Refers to both the IEEE 802.1 Spanning-Tree Protocol standard and the earlier Digital Equipment Corporation Spanning-Tree Protocol upon which it is based. The IEEE version supports bridge domains and allows the bridge to construct a loop-free topology across an extended LAN. The IEEE version is generally preferred over the Digital version. Sometimes abbreviated *STP*. See also *BPDU*, *MAC address learning*, *spanning tree*, and *spanning-tree algorithm*.

subnet mask

32-bit address mask used in IP to indicate the bits of an IP address that are being used for the subnet address. Sometimes referred to as *mask*.

Subnetwork Access Protocol

See *SNAP*.

SVC
Switched virtual circuit. Virtual circuit that is dynamically established on demand and is torn down when transmission is complete. SVCs are used in situations where data transmission is sporadic. Called a *switched virtual connection* in ATM terminology. Compare with *PVC*.

switch
Network device that filters, forwards, and floods frames based on the destination address of each frame. The switch operates at the data link layer of the OSI model.

switched LAN
LAN implemented with LAN switches. See *LAN switch*.

Switched Multimegabit Data Service
See *SMDS*.

Switched Port Analyzer
See *span*.

switched virtual circuit
See *SVC*.

T

TC
Transmission convergence. A sublayer of the ATM physical layer that transforms the flow of cells into a steady flow of bits for transmission over the physical medium. When transmitting, the TC sublayer maps the cells into the frame format, generates the HEC, and sends idle cells when there is nothing to send. When receiving, the TC sublayer delineates individual cells in the received bit stream and uses HEC to detect and correct errors.

TCP/IP
Transmission Control Protocol/Internet Protocol. Common name for the suite of protocols developed by the U.S. DoD in the 1970s to support the construction of worldwide internetworks. TCP and IP are the two best-known protocols in the suite.

TDM
Time-division multiplexing. Technique in which information from multiple channels can be allocated bandwidth on a single wire based on preassigned time slots. Bandwidth is allocated to each channel regardless of whether the station has data to transmit.

Telnet
Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection, enabling users to log in to remote systems and use resources as if they were connected to a local system. Telnet is defined in RFC 854.

TIA
Telecommunications Industry Association. Organization that develops standards relating to telecommunications technologies. Together, the TIA and the EIA have formalized standards, such as EIA/TIA-232, for the electrical characteristics of data transmission. See also *EIA*.

TIC
Token Ring interface coupler. Controller through which an FEP connects to a Token Ring.

timeout
Event that occurs when one network device expects to hear from another network device within a specified period of time, but does not. The resulting timeout usually results in a retransmission of information or the dissolving of the session between the two devices.

token
Frame that contains control information. Possession of the token allows a network device to transmit data onto the network. See also *token passing*.

TokenChannel
A parallel configuration of 2 to 8 ports between two Catalyst 3900s. Can also be used with the Catalyst 3920.

token passing
Access method by which network devices access the physical medium in an orderly fashion based on possession of a small frame called a token. Contrast with *circuit switching*. See also *token*.

Token Ring
Token-passing LAN developed and supported by IBM. Token Ring runs at 4 or 16 Mbps over a ring topology. Similar to IEEE 802.5. See also *IEEE 802.5*, *ring topology*, and *token passing*.

Token Ring interface coupler
See *TIC*.

topology
Physical arrangement of network nodes and media within an enterprise networking structure.

traffic policing
Process used to measure the actual traffic flow across a given connection and compare it to the total admissible traffic flow for that connection. Traffic outside of the agreed upon flow can be tagged (where the CLP bit is set to 1) and can be discarded en route if congestion develops. Traffic policing is used in ATM, Frame Relay, and other types of networks. Also know as *admission control*, *permit processing*, *rate enforcement*, and *UPC (usage parameter control)*.

traffic profile
Set of COS attribute values assigned to a given port on an ATM switch. The profile affects numerous parameters for data transmitted from the port including rate, cell drop eligibility, transmit priority, and inactivity timer.

translational bridging
Bridging between networks with dissimilar MAC sublayer protocols. MAC information is translated into the format of the destination network at the bridge.

transmission convergence
See *TC*.

transparent bridging

Bridging scheme often used in Ethernet and IEEE 802.3 networks in which bridges pass frames along one hop at a time based on tables associating end nodes with bridge ports. Transparent bridging is so named because the presence of bridges is transparent to network end nodes. Contrast with *SRB*.

TrBRF

See *BRF*.

TrCRF

See *CRF*.

trunk

Physical and logical connection between two switches across which network traffic travels. A backbone is composed of a number of trunks.

twisted pair

Relatively low-speed transmission medium consisting of two insulated wires arranged in a regular spiral pattern. The wires can be shielded or unshielded. Twisted pair is common in telephony applications and is increasingly common in data networks. See also *STP* and *UTP*.

U

UART

Universal Asynchronous Receiver/Transmitter. Integrated circuit, attached to the parallel bus of a computer, used for serial communications. The UART translates between serial and parallel signals, provides transmission clocking, and buffers data sent to or from the computer.

UBR

Unspecified bit rate. QOS class defined by the ATM Forum for ATM networks. UBR allows any amount of data up to a specified maximum to be sent across the network, but there are no guarantees in terms of cell loss rate and delay. Compare with *ABR (available bit rate)*, *CBR*, and *VBR*.

undistributed TrCRF

The standard type of TrCRF. The undistributed TrCRF is located on one switch and has a logical ring number associated with it. Multiple undistributed TrCRFs located on the same or separate switches can be associated with a single parent TrBRF. The parent TrBRF acts as a multiport bridge, forwarding traffic between the undistributed TrCRFs.

UNI

User-Network Interface. ATM Forum specification that defines an interoperability standard for the interface between ATM-based products (a router or an ATM switch) located in a private network and the ATM switches located within the public carrier networks. Also used to describe similar connections in Frame Relay networks. See also *NNI*.

unicast

Message sent to a single network destination. Compare with *broadcast* and *multicast*.

unicast address

Address specifying a single network device. Compare with *broadcast address* and *multicast address*. See also *unicast*.

universally administered address

Factory-assigned MAC address that is unique to each device.

unshielded twisted-pair

See *UTP*.

unspecified bit rate

See *UBR*.

uplinks

A term used to refer to any high-speed connection between two or more switches, between a switch and a router, a channel, and so forth. Examples include ISL, ATM, FDDI, and PFP.

User-Network Interface

See *UNI*.

UTP

Unshielded twisted-pair. Four-pair wire medium used in a variety of networks. UTP does not require the fixed spacing between connections that is necessary with coaxial-type connections. There are five types of UTP cabling commonly used: *Category 1 cabling*, *Category 2 cabling*, *Category 3 cabling*, *Category 4 cabling*, and *Category 5 cabling*. Compare with *STP*. See also *EIA/TIA-586* and *twisted pair*.

V

variable bit rate

See *VBR*.

VBR

Variable bit rate. QOS class defined by the ATM Forum for ATM networks. VBR is subdivided into a real time (RT) class and non-real time (NRT) class. VBR (RT) is used for connections in which there is a fixed timing relationship between samples. VBR (NRT) is used for connections in which there is no fixed timing relationship between samples, but that still need a guaranteed QOS. Compare with *ABR (available bit rate)*, *CBR*, and *UBR*.

VC

See *virtual circuit*.

VCC

Virtual channel connection. Logical circuit, made up of VCLs, that carries data between two end points in an ATM network. Sometimes called a *virtual circuit connection*. See also *VCI* and *VPI*.

VCI

Virtual channel identifier. 16-bit field in the header of an ATM cell. The VCI, together with the VPI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next network VCL that a cell needs to transit on its way to its final destination. The function of the VCI is similar to that of the DLCI in Frame Relay.

VCL
Virtual channel link. Connection between two ATM devices. A VCC is made up of one or more VCLs. See also VCC.

virtual channel connection
See *VCC*.

virtual channel identifier
See *VCI*.

virtual channel link
See *VCL*.

virtual circuit
Logical circuit created to ensure reliable communication between two network devices. A virtual circuit is defined by a VPI/VCI pair, and can be either permanent (a PVC) or switched (an SVC). Virtual circuits are used in Frame Relay and X.25. In ATM, a virtual circuit is called a virtual channel. Sometimes abbreviated VC. See also *PVC*, *SVC*, *VCI*, and *VPI*.

virtual LAN
See *VLAN*.

virtual path identifier
See *VPI*.

virtual path identifier/virtual channel identifier
See *VPI/VCI*.

virtual ring
Entity in an SRB network that logically connects two or more physical rings together either locally or remotely. The concept of virtual rings can be expanded across router boundaries.

VLAN
Virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.

VLAN Trunking Protocol
See *VTP*.

VPC
Virtual path connection. Grouping of VCCs that share one or more contiguous VPLs. See also *VCC*.

VPI
Virtual path identifier. 8-bit field in the header of an ATM cell. The VPI, together with the VCI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next VCL that a cell needs to transit on its way to its final destination. The function of the VPI is similar to that of the DLCI in Frame Relay. See also *VCI*.

VPI/VCI

See *VCI* and *VPI*.

VTP

VLAN Trunking Protocol. Cisco-defined protocol used to configure and manage virtual LANs across a switch network.

W

wiring closet

Specially designed room used for wiring a data or voice network. Wiring closets serve as a central junction point for the wiring and wiring equipment that is used for interconnecting devices.