



Overview of Bridging

The Bridging section of this guide discusses the following software components for bridging and routing protocols in Cisco routers:

- [Transparent and Source-Route Transparent \(SRT\) Bridging, page 1](#)
- [Source-Route Bridging \(SRB\), page 6](#)
- [Token Ring Inter-Switch Link \(TRISL\), page 7](#)
- [Token Ring Route Switch Module \(TRRSM\), page 8](#)

This overview chapter gives a high-level description of each technology. For configuration information, refer to the corresponding chapter in this publication.



Note

All commands supported on the Cisco 7500 series routers are also supported on the Cisco 7000 series routers.

Transparent and Source-Route Transparent (SRT) Bridging

Cisco IOS software supports transparent bridging for Ethernet, Fiber Distributed Data Interface (FDDI), and serial media, and supports source-route transparent (SRT) bridging for Token Ring media. In addition, Cisco supports all the mandatory Management Information Base (MIB) variables specified for transparent bridging in RFC 1286.

This section contains the following topics:

- [Transparent Bridging Features, page 1](#)
- [Integrated Routing and Bridging, page 2](#)
- [SRT Bridging Features, page 5](#)

Transparent Bridging Features

The Cisco transparent bridging software implementation has the following features:

- Complies with the IEEE 802.1D standard.



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

- Provides the ability to logically segment a transparently bridged network into virtual LANs.
- Provides two Spanning Tree Protocols—an older bridge protocol data unit (BPDU) format that is compatible with Digital and other LAN bridges for backward compatibility and the IEEE standard BPDU format. In addition to features standard with these Spanning Tree Protocols, the Cisco proprietary software provides for multiple domains for spanning trees. The spanning-tree parameters are configurable.
- Allows frame filtering based on Media Access Control (MAC) address, protocol type, or the vendor code. Additionally, the bridging software can be configured to selectively filter local-area transport (LAT) multicast service announcements.
- Provides deterministic load distribution while maintaining a loop-free spanning tree.
- Provides the ability to bridge over Asynchronous Transfer Mode (ATM), dial-on-demand routing (DDR), FDDI, Frame Relay, multiprotocol Link Access Procedure, Balanced (LAPB), Switched Multimegabit Data Service (SMDS), and X.25 networks.
- Provides concurrent routing and bridging, which is the ability to bridge a given protocol on some interfaces in a router and concurrently route that protocol on other interfaces in the same router.
- Provides integrated routing and bridging, which is the ability to route a given protocol between routed interfaces and bridge groups, or to route a given protocol between bridge groups.
- Provides fast-switched transparent bridging for Frame Relay encapsulated serial and High-Speed Serial Interface (HSSI) interfaces, according to the format specified in RFC 1490.
- Provides fast-switched transparent bridging for the ATM interface on the Cisco 7000, according to the format specified in RFC 1483.
- Provides for compression of LAT frames to reduce LAT traffic through the network.
- Provides both bridging and routing of VLANs.

Cisco access servers and routers can be configured to serve as both multiprotocol routers and MAC-level bridges, bridging any traffic that cannot otherwise be routed. For example, a router routing the Internet Protocol (IP) can also bridge Digital's LAT protocol or NetBIOS traffic.

Cisco routers also support remote bridging over synchronous serial lines. As with frames received on all other media types, dynamic learning and configurable filtering applies to frames received on serial lines.


Transit bridging of Ethernet frames across FDDI media is also supported. The term *transit* refers to the fact that the source or destination of the frame cannot be on the FDDI media itself. This allows FDDI to act as a highly efficient backbone for the interconnection of many bridged networks. The configuration of FDDI transit bridging is identical to the configuration of transparent bridging on all other media types.

Integrated Routing and Bridging

Although concurrent routing and bridging makes it possible to both route and bridge a specific protocol on separate interfaces within a router, the protocol is not switched between bridged and routed interfaces. Routed traffic is confined to the routed interfaces; bridged traffic is confined to bridged interfaces. A specified protocol may be either routed or bridged on a given interface, but not both.

Integrated routing and bridging makes it possible to route a specific protocol between routed interfaces and bridge groups, or route a specific protocol between bridge groups. Local or unroutable traffic can be bridged among the bridged interfaces in the same bridge group, while routable traffic can be routed to other routed interfaces or bridge groups. [Figure 1](#) illustrates how integrated routing and bridging in a router interconnects a bridged network with a routed network.

Figure 1 *Integrated Routing and Bridging Interconnecting a Bridged Network with a Routed Network*



You can configure the Cisco IOS software to route a specific protocol between routed interfaces and bridge groups or to route a specific protocol between bridge groups. Specifically, local or unroutable traffic is bridged among the bridged interfaces in the same bridge group, while routable traffic is routed to other routed interfaces or bridge groups. Using integrated routing and bridging, you can do the following:


- Switch packets from a bridged interface to a routed interface
- Switch packets from a routed interface to a bridged interface
- Switch packets within the same bridge group

Bridge-Group Virtual Interface

Because bridging operates in the data link layer and routing operates in the network layer, they follow different protocol configuration models. Taking the basic IP model as an example, all bridged interfaces would belong to the same network, while each routed interface represents a distinct network.

In integrated routing and bridging, the bridge-group virtual interface is introduced to avoid confusing the protocol configuration model when a specific protocol is both bridged and routed in a bridge group. [Figure 2](#) illustrates the bridge-group virtual interface as a user-configured virtual interface residing within a router.

Figure 2 *Bridge-Group Virtual Interface in the Router*



The bridge-group virtual interface is a normal routed interface that does not support bridging, but does represent its corresponding bridge group to the routed interface. It has all the network layer attributes (such as a network layer address and filters) that apply to the corresponding bridge group. The interface number assigned to this virtual interface corresponds to the bridge group that this virtual interface represents. This number is the link between the virtual interface and the bridge group.

When you enable routing for a given protocol on the bridge-group virtual interface, packets coming from a routed interface, but destined for a host in a bridged domain, are routed to the bridge-group virtual interface and are forwarded to the corresponding bridged interface. All traffic routed to the bridge-group virtual interface is forwarded to the corresponding bridge group as bridged traffic. All routable traffic received on a bridged interface is routed to other routed interfaces as if it is coming directly from the bridge-group virtual interface.

To receive routable packets arriving on a bridged interface but destined for a routed interface or to receive routed packets, the bridge-group virtual interface must also have the appropriate addresses. MAC addresses and network addresses are assigned to the bridge-group virtual interface as follows:

- The bridge-group virtual interface “borrows” the MAC address of one of the bridged interfaces in the bridge group associated with the bridge-group virtual interface.
- To route and bridge a given protocol in the same bridge group, you must configure the network layer attributes of the protocol on the bridge-group virtual interface. No protocol attributes should be configured on the bridged interfaces, and no bridging attributes can be configured on the bridge-group virtual interface.

Because there can be only one bridge-group virtual interface representing a bridge group, and the bridge group can be made up of different media types configured for several different encapsulation methods, you may need to configure the bridge-group virtual interface with the particular encapsulation methods required to switch packets correctly.

For example, the bridge-group virtual interface has default data link and network layer encapsulations that are the same as those available on Ethernet interfaces, but you can configure the bridge-group virtual interface with encapsulations that are not supported on an Ethernet interface. In some cases, the default encapsulations provide appropriate results; in other cases they do not. For example, with default encapsulation, Advanced Research Projects Agency (ARPA) packets from the bridge-group virtual interface are translated to Subnetwork Access Protocol (SNAP) when bridging IP to a Token Ring- or FDDI-bridged interface. But for Internet Packet Exchange (IPX), Novell-ether encapsulation from the bridge-group virtual interface is translated to raw-token or raw-FDDI when bridging IPX to a Token Ring- or FDDI-bridged interface. Because this behavior is usually not what you want, you must configure IPX SNAP or Service Advertisement Protocol (SAP) encapsulation on the bridge-group virtual interface.

Other Considerations

The following are additional facts regarding the support of integrated routing and bridging:

- Integrated routing and bridging is not supported on cBus platforms (AGS+ and Cisco 7000 series).
- Integrated routing and bridging is supported for transparent bridging, but not for source-route bridging (SRB).
- Integrated routing and bridging is supported on all media interfaces except X.25 and Integrated Services Digital Network (ISDN) bridged interfaces.
- Integrated routing and bridging supports three protocols: IP, IPX, and AppleTalk in both fast-switching and process-switching modes.
- Integrated routing and bridging and concurrent routing and bridging cannot operate at the same time.

SRT Bridging Features

Cisco routers support transparent bridging on Token Ring interfaces that support SRT bridging. Both transparent and SRT bridging are supported on all Token Ring interface cards that can be configured for either 4- or 16-MB transmission speeds.

As with other media, all the features that use **bridge-group** commands can be used on Token Ring interfaces. As with other interface types, the bridge group can be configured to run either the IEEE or Digital Spanning Tree Protocols. When configured for the IEEE Spanning Tree Protocol, the bridge cooperates with other SRT bridges and constructs a loop-free topology across the entire extended LAN.

You can also run the Digital Spanning Tree Protocol over Token Ring. Use it when you have other non-IEEE bridges on other media and you do not have any SRT bridges on Token Ring. In this configuration, all the Token Ring transparent bridges must be Cisco routers. This is because the Digital Spanning Tree Protocol has not been standardized on Token Ring.

As specified by the SRT bridging specification, only packets without a routing information field (RIF) (RII = 0 in the SA field) are transparently bridged. Packets with a RIF (RII = 1) are passed to the SRB module for handling. An SRT-capable Token Ring interface can have both SRB and transparent bridging enabled at the same time. However, with SRT bridging, frames that did not have a RIF when they were produced by their generating host never gain a RIF, and frames that did have a RIF when they were produced never lose that RIF.



Note

Because bridges running only SRT bridging never add or remove RIFs from frames, they do not integrate SRB with transparent bridging. A host connected to a source-route bridge that expects RIFs can *never* communicate with a device across a bridge that does not understand RIFs. SRT bridging cannot tie in existing source-route bridges to a transparent bridged network. To tie in existing bridges, you must use source-route translational bridging (SR/TLB) instead. SR/TLB is described in the “[Configuring Source-Route Bridging](#)” chapter.

Bridging between Token Ring and other media requires certain packet transformations. In all cases, the MAC addresses are bit-swapped because the bit ordering on Token Ring is different from that on other media. In addition, Token Ring supports one packet format, logical link control (LLC), while Ethernet supports two formats (LLC and Ethernet).

The transformation of LLC frames between media is simple. A length field is either created (when the frame is sent to non-Token Ring) or removed (when the frame is sent to Token Ring). When an Ethernet format frame is sent to Token Ring, the frame is translated into an LLC-1 SNAP packet. The destination service access point (DSAP) value is AA, the source service access point (SSAP) value is AA, and the organizational unique identifier (OUI) value is 0000F8. Likewise, when a packet in LLC-1 format is bridged onto Ethernet media, the packet is translated into Ethernet format.



Caution

Bridging between dissimilar media presents several problems that can prevent communication from occurring. These problems include bit order translation (or using MAC addresses as data), maximum transmission unit (MTU) differences, frame status differences, and multicast address usage. Some or all these problems might be present in a multimedia bridged LAN. Because of differences in the way end nodes implement Token Ring, these problems are most prevalent when bridging between Token Ring and Ethernet or between Ethernet and FDDI LANs.

Problems currently occur with the following protocols when bridged between Token Ring and other media: Novell IPX, DECnet Phase IV, AppleTalk, Banyan VINES, Xerox Network Systems (XNS), and IP. Further, problems can occur with the Novell IPX and XNS protocols when bridged between FDDI and other media. We recommend that these protocols be routed whenever possible.

Source-Route Bridging (SRB)


The Cisco IOS bridging software includes SRB capability. A source-route bridge connects multiple physical Token Rings into one logical network segment. If the network segment bridges only Token Ring media to provide connectivity, the technology is termed SRB. If the network bridges Token Ring and non-Token Ring media is introduced into the bridged network segment, the technology is termed remote source-route bridging (RSRB).

SRB enables routers to simultaneously act as a Level 3 router and a Level 2 source-route bridge. Thus, protocols such as Novell's IPX or XNS can be routed on Token Rings, while other protocols such as Systems Network Architecture (SNA) or NetBIOS are source-route bridged.

SRB technology is a combination of bridging and routing functions. A source-route bridge can make routing decisions based on the contents of the MAC frame header. Keeping the routing function at the MAC, or Level 2, layer allows the higher-layer protocols to execute their tasks more efficiently and allows the LAN to be expanded without the knowledge of the higher-layer protocols.

As designed by IBM and the IEEE 802.5 committee, source-route bridges connect extended Token Ring LANs. A source-route bridge uses the RIF in the IEEE 802.5 MAC header of a datagram (Figure 3) to determine which rings or Token Ring network segments the packet must transit.

Figure 3 IEEE 802.5 Token Ring Frame Format



The source station inserts the RIF into the MAC header immediately following the source address field in every frame, giving this style of bridging its name. The destination station reverses the routing field to reach the originating station.

The information in a RIF is derived from explorer packets generated by the source node. These explorer packets traverse the entire source-route bridge network, gathering information on the possible paths the source node might use to send packets to the destination.

Transparent spanning-tree bridging requires time to recompute a topology in the event of a failure; SRB, which maintains multiple paths, allows fast selection of alternate routes in the event of failure. Most importantly, SRB allows the end stations to determine the routes the frames take.

SRB Features

The Cisco SRB implementation has the following features:

- Provides configurable fast-switching software for SRB.
- Provides for a local source-route bridge that connects two or more Token Ring networks.
- Provides *ring groups* to configure a source-route bridge with more than two network interfaces. A ring group is a collection of Token Ring interfaces in one or more routers that are collectively treated as a *virtual ring*.
- Provides two types of explorer packets to collect RIF information—an *all-routes* explorer packet, which follows all possible paths to a destination ring, and a *spanning-tree* explorer packet, which follows a statically configured limited route (spanning tree) when looking for paths.
- Provides a dynamically determined RIF cache based on the protocol. The software also allows you to add entries manually to the RIF cache.

- Provides for filtering by MAC address, link service access point (LSAP) header, and protocol type.
- Provides for filtering of NetBIOS frames either by station name or by a packet byte offset.
- Provides for translation into transparently bridged frames to allow source-route stations to communicate with nonsource-route stations (typically on Ethernet).
- Provides support for the SRB MIB variables as described in the IETF draft “Bridge MIB” document, “Definition of Managed Objects for Bridges,” by E. Decker, P. Langille, A. Rijsinghani, and K. McCloghrie, June 1991. Only the SRB component of the Bridge MIB is supported.
- Provides support for the Token Ring MIB variables as described in RFC 1231, *IEEE 802.5 Token Ring MIB*, by K. McCloghrie, R. Fox, and E. Decker, May 1991. Cisco implements the mandatory tables (Interface Table and Statistics Table), but not the optional table (Timer Table) of the Token Ring MIB. The Token Ring MIB has been implemented for the 4/16-Mb Token Ring cards that can be user adjusted for either 4- or 16-Mb transmission speeds (CSC-1R, CSC-2R, CSC-R16M, or CSC-C2CTR).
- SRB is supported over FDDI on Cisco 7200 series routers.
- Particle-based switching is supported (over FDDI and Token Ring) by default on Cisco 7200 series routers.
- Complies with RFC 1483 in Cisco IOS Release 12.0(3)T and later by offering the ability to encapsulate SRB traffic using RFC 1483 bridged LLC encapsulation. This support enables SRB over ATM functionality that is interoperable with other vendors’ implementations of SRB over ATM.

Token Ring Inter-Switch Link (TRISL)

ISL is a Layer 2 protocol that enables switches and routers to transport Ethernet frames from multiple VLANs across Fast Ethernet or Gigabit Ethernet links. The Cisco TRISL protocol extends the ISL model to include the transport of Token Ring frames from multiple VLANs across these same links.

TRISL support on Cisco routers provides inter-VLAN routing and bridging across a 100 Mb Fast Ethernet link. ISL and TRISL together provide routing and bridging between Token Ring and Ethernet LANs, ELANS, and VLANs.

TRISL is supported on the following platforms with any one of the following port adapters:

- Cisco 7500 or Cisco 7200 series routers
 - Two-port Fast Ethernet/ISL 100BaseTX
 - Two-port Fast Ethernet/ISL 100BaseFX
 - One-port Fast Ethernet 100BaseTX
 - One-port Fast Ethernet 100BaseFX
- Cisco 4500 or 4700 series routers
 - NM-1FE
- Cisco 3600 or 2600 series routers
 - NM-1FE1CT1
 - NM-1FE2CT1
 - NM-1FE1CE1

**Note**

The two-port Fast Ethernet/ISL port adapters support frame sizes up to 17800 bytes and the one-port Fast Ethernet port adapters support a frame size of up to 1500 bytes.

TRISL provides the following new capabilities and features:

- IP routing for source-routed and nonsource-routed frames between TRISL VLANs and any LAN, ELAN, or VLAN.
- IPX routing for source-routed and nonsource-routed frames between TRISL VLANs and any LANs, ELANs, or VLANs.
- SRB between TRISL VLANs and SRB-capable LANs, ELANs, or VLANs.
- SRT between TRISL VLANs and SRT-capable LANs, ELANs, or VLANs.
- SR/TLB between TRISL VLANs and Ethernet LANs, ELANs, or VLANs.
- Duplicate Ring Protocol (DRiP), which prevents external loops that could result if the router's virtual ring number were duplicated elsewhere in the network.

**Note**

VLAN Trunk Protocol (VTP) is not supported for TRISL on the routers in this release.

Token Ring Route Switch Module (TRRSM)

The Token Ring VLAN support on the Route Switch Module (RSM) adds the capability to do multi-protocol routing and bridging for Token Ring VLANs on the RSM. The RSM is a router module running Cisco IOS software that plugs into a Token Ring switch backplane.

This section contains a brief overview of Token Ring switching, which is described in the following topics:

- [Switching Overview, page 8](#)
- [Usability of Switching, page 9](#)
- [VLAN, page 10](#)
- [Token Ring VLANs, page 10](#)
- [Token Ring VLAN Support on the RSM, page 11](#)

Switching Overview

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 sent. Switches, however, operate at much higher speeds than bridges, and can support new functionality, such as virtual LANs (VLANs). See the [“VLAN” section on page 10](#) and the [“Token Ring VLANs” section on page 10](#).

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, and 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. The Cisco 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, and support for ATM and Inter-Switch Link (ISL).
- The Catalyst 3920, which is also a stackable workgroup switch that provides support for all switching modes, filtering, RMON, DTR, and SNMP management.
- The Catalyst 5000, which is a modular switch that supports Ethernet, Fast Ethernet, FDDI, ATM, and now Token Ring.

The Catalyst Token Ring switches support the following bridging modes: SRB, SRT, and source-route switching.

Usability of Switching

The traditional method of connecting multiple Token Ring segments is to use a 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.

VLAN

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.

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 and one level of VLAN. In source routing, however, there are two types of broadcast frames:

- Those that are confined to a single ring
- Those that traverse the bridged domain


Therefore, there are two levels of VLANs in a Token Ring switched network. These two categories of broadcast frames result in a broadcast domain that is hierarchical in nature, just as a local ring domain can exist only within a domain of all the interconnected rings.

The first level is the Token Ring Concentrator Relay Function (TrCRF). 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). At this level, the VLAN is a logical ring and, as such, is assigned a ring number. On a Token Ring switch, the logical ring (TrCRF) contains one or more physical ports. source-route switching is used to forward frames within a TrCRF based on MAC address or Route Descriptor. On an RSM, a logical ring (TrCRF) can be defined that does not contain any physical ports, but rather is used only in processing source-routed traffic to terminate the RIF.

The second level of VLAN is the Token Ring Bridge Relay Function (TrBRF). This is the parent VLAN to which TrCRF VLANs are assigned. The domain of interconnected 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). At this level, the VLAN is a logical bridge and, as such, is assigned a bridge number. The TrBRF is responsible for forwarding frames between groups of ports with the same ring number (TrCRFs) via either SRB or SRT.

Figure 4 depicts the relationship between TrCRF and TrBRF VLANs.

Figure 4 Token Ring VLAN Support on the RSM



Token Ring VLAN Support on the RSM

The Token Ring VLAN support on the RSM adds the capability to do multiprotocol routing and bridging for Token Ring VLANs on the RSM. The RSM can be used alone to do inter-VLAN routing, or it can be paired with a Catalyst VIP2 to provide external network connections with the same port adapters used on Cisco 7500 series routers. The RSM/VIP2 combination provides routing between VLANs and Catalyst VIP2 port adapters. A complete description of the RSM can be found in the *Catalyst 5000 Series Route Switch Module Installation and Configuration Note* and the *Route Switch Module Catalyst VIP2-15 and VIP2-40 Installation and Configuration Note*.

The Token Ring VLAN support on the RSM adds the following functionality to the Catalyst 5000 switch:

- IP routing for source-routed and non-source-routed frames between Token Ring (TrBRF) or Ethernet VLANs and VIP2 interfaces
- IPX routing for source-routed and non-source-routed frames between Token Ring (TrBRF) or Ethernet VLANs and VIP2 interfaces
- SRB between Token Ring (TrBRF) VLANs and VIP2 interfaces
- SR/TLB between Token Ring (TrBRF) VLANs and Ethernet VLANs and VIP2 interfaces
- SRT between Token Ring (TrBRF) VLANs and SRT-capable VLANs and VIP2 interfaces


Both APPN and DLSw+ are supported for Token Ring VLANs on the RSM. However, RSRB is not supported on the RSM.

For information on how Token Ring VLANs are implemented on switches, refer to the *Catalyst Token Ring Switching Implementation Guide*, the *Catalyst 5000 Series Token Ring Configuration Notes*, the *Catalyst 3900 Token Ring Switching User Guide*, and the *Catalyst 3920 Token Ring Switching User Guide*.

The RSM is a router module running Cisco IOS router software that directly interfaces (plugs into) the Catalyst switch backplane. From the Token Ring VLAN perspective, the interface to the RSM is at the Token Ring bridged network (TrBRF) level. With the RSM, it is possible to route or bridge between separate Token Ring and Ethernet domains.

When routing or bridging between TrBRF VLANs that are defined as SRB domains, it is necessary to create a logical ring on the RSM for proper RIF processing. This logical ring is defined as a TrCRF VLAN that does not contain any external Token Ring switch ports. [Figure 5](#) illustrates the logical view of IP routing between two source-route bridged VLANs on the RSM. In this view, the RSM appears to have an interface to both ring 100 and ring 110.

Figure 5 Logical View of VLAN Support on the RSM



The Token Ring RSM feature is supported on the RSM in the Catalyst 5000 platform. Support for the Token Ring RSM feature was first introduced in the Cisco IOS Release 11.3(5)T. The Token Ring RSM feature is supported on all Cisco IOS Release 12.0 T software release images. A list of the supported Cisco IOS releases and software images are located in the *Release Notes for Catalyst 5000 Series RSM/VIP2 Cisco IOS 12.0 T Software Releases* publication. For a complete functional description of the RSM, refer to the *Catalyst 5000 Series RSM Installation and Configuration Note*.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2008 Cisco Systems, Inc. All rights reserved.