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

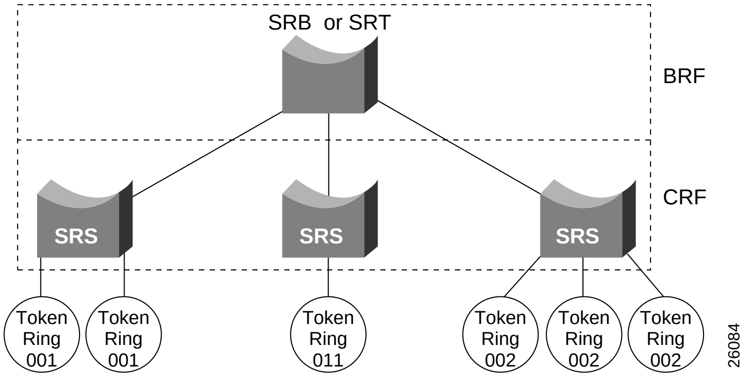

TrCRFs

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

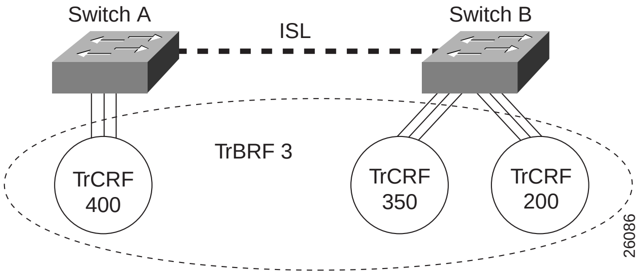

Default and Distributed TrCRFs

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

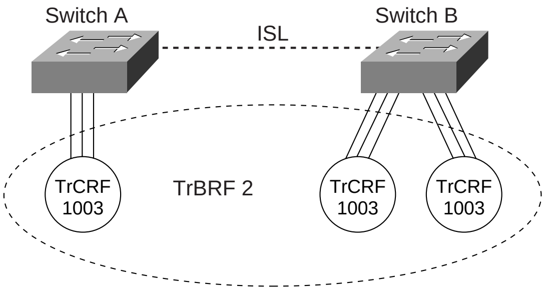

In addition to the default TrCRF, the Catalyst 5000 series Token Ring module supports the configuration of a distributed TrCRF. A distributed TrCRF contains ports on different switches as illustrated in Figure 4-4. While you can distribute the ports of a TrCRF across Token Ring modules on separate Catalyst 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

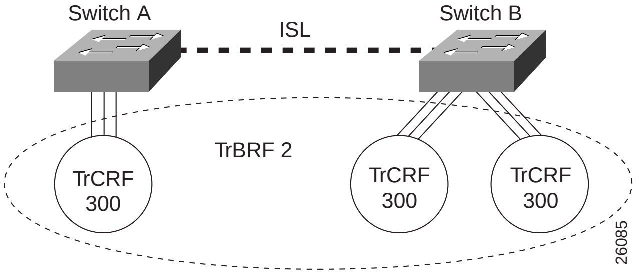

Backup TrCRF

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

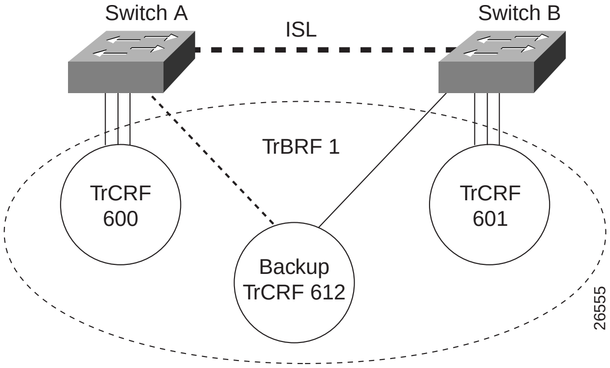

TrBRFs

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

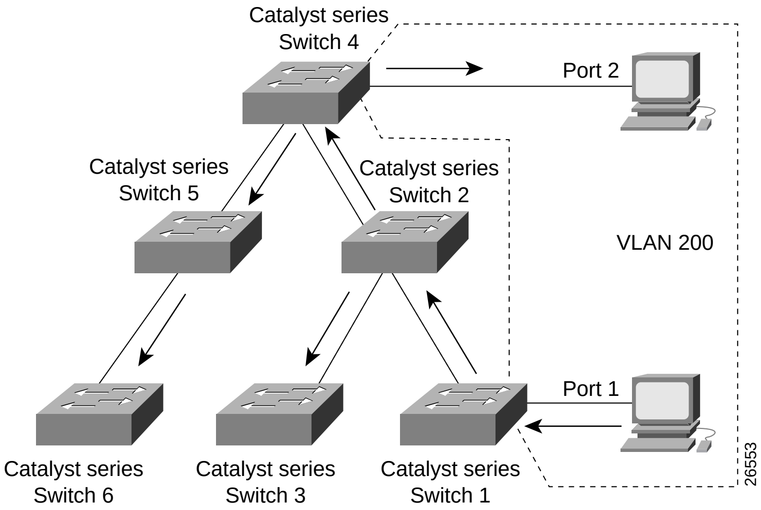

Figure 4-7 shows the same switched network with VTP pruning enabled. The broadcast traffic from Switch 1 is not forwarded to Switches 3, 5, and 6 because traffic for the VLAN 200 has been pruned on the links indicated (port 5 on Switch 2 and port 4 on Switch 4).

Figure 4-7 Flooding Traffic with VTP Pruning

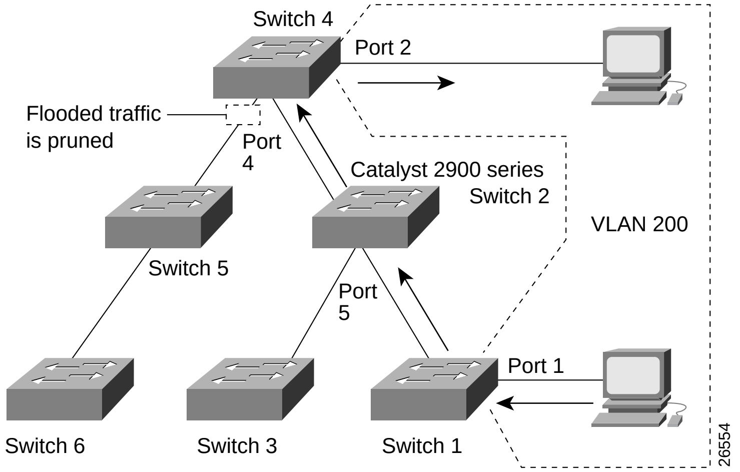

Enabling VTP pruning on a VTP server enables pruning for the entire management domain. VTP pruning takes effect several seconds after you enable it. By default, VLANs 2 through 1000 are pruning eligible. VTP pruning does not prune traffic from VLANs that are pruning ineligible. VLAN 1, the default TrBRF (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