

Interconnecting Catalyst 3900 Switches

Catalyst 3900s can be interconnected to form large networks. This chapter describes the methods of connecting Catalyst 3900s: using TokenChannels or ISL Channels, as a stack in a back-to-back configuration or using a Catalyst Matrix, or using high-speed uplinks. You can interconnect Catalyst 3920s with Catalyst 3900s. In this chapter, Catalyst 39xx is used to represent both.

This chapter contains the following information about interconnecting Catalyst 3900 switches:

- Using Channels on page 3-1
- Stacking Catalyst 3900s on page 3-2
- Using High-Speed Uplinks on page 3-7

Note If you plan to interconnect your Catalyst 3900 with a Catalyst 5000, the Catalyst 5000 series switch must have a Supervisor Engine II or a Supervisor Engine III installed that is running supervisor engine software Release 3.0(201) or later.

Using Channels

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

Using channels provides the following benefits:

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

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 the 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 an ISL Channel. Failure to disconnect or disable the ports might result in network loops.

TokenChannels

A TokenChannel consists of two to eight parallel channels 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.



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.

For more information on TokenChannels, refer to the “Using Channels” section on page 6-49 of the “Configuring the Catalyst 3900” chapter.

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.

For more information on ISL Channels, refer to the “Using Channels” section on page 6-49 of the “Configuring the Catalyst 3900” chapter.

Stacking Catalyst 3900s

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

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

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

Two Catalyst 39xx 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 39xx. This creates a direct connection between the two Catalyst 39xx 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 39xx 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.

When you power-on a Catalyst 39xx, it runs through a set of self-diagnostics. Immediately after the diagnostics are completed, the Catalyst 39xx runs through a *stack discovery mode*. This discovery mode detects whether the switch is cabled to another Catalyst 39xx. If it is determined during the discovery mode that the Catalyst 39xx is connected to other switches, the switches are automatically combined to form a stack. At the end of the discovery mode, if it is determined that the Catalyst 39xx is not connected to another switch, it will operate as a stand alone 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

Stack Port Revision Levels and Frame Sizes

Catalyst 3900 Release 2.1(1) and later includes support for larger frame sizes. The use of larger frame sizes in a stacked configuration requires Revision B of the Stack Port module. The initial release of the Catalyst 3900 Stack Port module (which is identified as Revision A) did not support MTUs greater than 4472. Revision B of the Catalyst 3900 Stack Port module supports MTUs of up to 18190. When configuring MTUs, keep the following in mind:

- For standalone switches, you can configure MTUs larger than 4472 if there is no Stack Port module installed or if there is a Stack Port module Revision B installed.
- If you create a stack, all the switches in the stack must include the same level of stack port. To use MTUs larger than 4472, all switches in the stack *must* contain stack ports that are Revision B or later.
- If your switch contains an ATM module, the MTU of a TrBRF that contains an ATM port is restricted to 4472 or less. The Catalyst 3900 ATM module does not support MTUs greater than 4472, therefore, configuring an MTU for a TrBRF that contains an ATM port will result in the ATM port being removed from the TrBRF.

If you attempt to create a stack that contains switches that have both Revision A and Revision B stack ports, the stack will not completely form. If the controlling switch (the first switch in the stack to come up) contains a Revision A module, only the switches with Revision A modules will be able to join the stack. Likewise, if the controlling switch contains a Revision B module, only the switches with Revision B modules will be able to join the stack.

To verify the revision level of a stack port module on a Catalyst 3900, see the Module Information panel. The panel has been updated to include the revision level of the stack port module in the HW Rev field. A Stack Port module Revision B will be displayed as 0.B.

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 39xx detects that it is connected to another Catalyst 39xx 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 39xx switch is assigned a box number.
 - The two Catalyst 39xx 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 39xx switch. The Catalyst 39xx with the lower MAC address becomes box 1, and the Catalyst 39xx switch with the higher MAC address becomes box 2.
- The Catalyst 39xx 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 on page 3-5 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 39xx 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 39xx that becomes box 1 also becomes the provider of the interbox parameters.
 - If the configuration information differs between the two Catalyst 39xx 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 eight Catalyst 39xx switches. The following sections describe how this multi-unit stack is formed.

Catalyst Matrix Description

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

Any combination of up to eight Catalyst 39xx 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 39xx 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 39xx.

Forming a Multi-Unit Catalyst Stack

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

As soon as the stack discovery mode is completed, each Catalyst 39xx is assigned a box number. With a Catalyst Matrix configuration, the box number for a Catalyst 39xx is determined by the port number the Catalyst 39xx is connected to on the Catalyst Matrix. For example, the Catalyst 39xx 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 39xx is moved to another port, the box number for that Catalyst 39xx 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 on page 3-5 later in this chapter lists the shared parameters. There are two possible ways of providing configuration information to the Catalyst 39xx switches in a stack. These methods are as follows:

- Preconfigure all the Catalyst 39xx 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 will cause 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 39xx 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 39xx switches have formed a stack and any additional Catalyst 39xx switches 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 39xx switches must combine to form a common configuration (interbox parameters). The stack operates as a single entity when all of the Catalyst 39xx switches in that stack use the same interbox parameters.

The following is a list of these shared interbox parameters. The parameters in this list are accessed through the console configuration menus. The console menus are described in “Configuring the Catalyst 3900” chapter.

- IP Configuration
 - IP address
 - Default gateway
 - Subnet mask
 - IP state
- Spanning-Tree Protocol
 - 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 High-Speed Uplinks

You can interconnect Catalyst 3900s using high-speed uplinks, such as the ATM module or the ISL module.

Using the ATM Module

Figure 3-1 shows how the ATM expansion module can be used to connect Catalyst 3900 switches through an ATM switch.

Figure 3-1 Connection Over an ATM Backbone



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

- ATM backbone for legacy LANs

A Catalyst 3900 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 3900 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 3900, 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 emulated LAN (ELAN), the Catalyst 3900'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 source-route bridging (SRB) or source-route transparent (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 (such as a maximum of 2048 VCCs per ATM expansion module and a maximum data rate of 155 Mbps) 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.

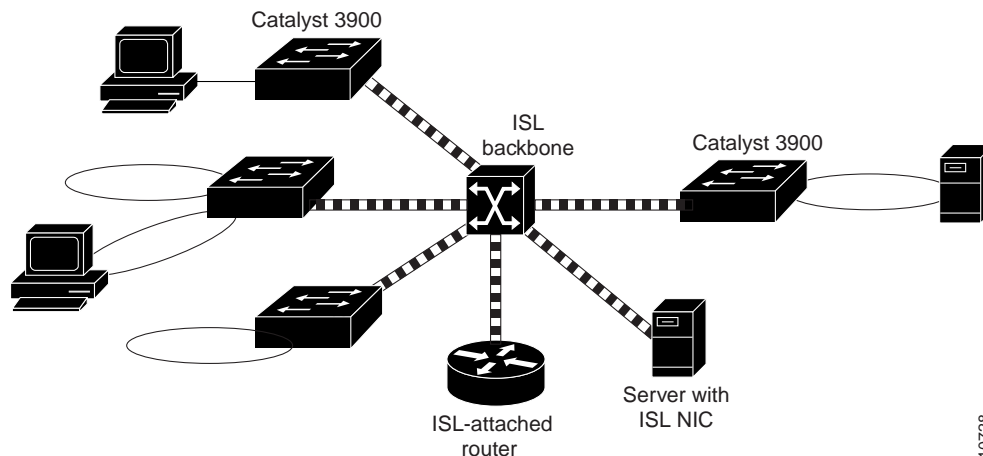
Note If the Catalyst 3900 has multiple ATM modules and each has clients active for the same ELAN, the Catalyst 3900 will not bridge between the ELANs. The Catalyst 3900 acts as an edge device on an ATM cloud.

Using ISL

The ISL backbone design looks very much like an ATM design, however, ISL is less expensive than ATM and avoids the need for 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 route switch module in the Catalyst 5500. Figure 3-2 shows an example of the ISL backbone design.

Figure 3-2 ISL Backbone Design



10728

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 supports both Token Ring and Ethernet VLANs. These NICs can be used for high-speed attachment to servers.

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 mode (using the **set port duplex** command).

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 configurations 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 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 of the ports to be active at a time. When using the default Catalyst 3900 STP values, the path cost is calculated based on a 200-Mbps connection which results in a path cost of 5 and causes the STP to place the ISL port in forwarding mode and the ATM port or Token Ring port in blocked mode.

However, if you modify the Catalyst 3900 port STP values or use devices from other vendors that use different STP default values can, you might cause the ISL port to become blocked. If an ISL port becomes blocked in an ISL-ATM parallel connection, traffic will be passed via the ATM link but VLAN trunking data will not. Also, if your STP configuration causes an ATM or Token Ring port to become 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 all-routes explorers (AREs) to the incoming TrCRF.

Therefore, when modifying STP values, always ensure that the STP port path costs are configured such 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.

