

Concepts

This chapter explains some of the concepts you'll need to understand to effectively configure and manage your Catalyst 2800. The Catalyst 2800 features described include:

- Address learning
- Forwarding, filtering, and flooding
- Switching modes
- Securing Ports
- Full-duplex 100Base-TX or 100Base-FX
- Port monitoring
- Virtual LANs
- Spanning-Tree Protocol

Address Learning

With multiple MAC address support on all ports, you can connect individual workstations, repeaters, switches, routers, or other network devices to any Catalyst 2800 port.

The Catalyst 2800 maintains an address table associating the address of each device with the port the device resides on.

The Catalyst 2800 can automatically learn the source address of packets received on each port. It adds these dynamic addresses and associated port numbers to the address table. As stations are added or removed from the network, the Catalyst 2800 automatically updates the address table by adding new entries and aging out ones currently not in use. Dynamic addressing simplifies the management of the switch and requires no configuration.

A network administrator can also manually enter addresses into the Catalyst 2800 address table. These static addresses do not age and must be added or removed by the administrator. Static addressing also allows for a measure of security in that access to a port can be restricted. See the “Securing Ports” section in this chapter for more information.

Forwarding, Filtering, and Flooding

The Catalyst 2800 forwards, filters, and floods packets in accordance with the IEEE 802.1d specification. It transfers, or forwards, packets between any combination of ports based on the destination address of the received packet. By maintaining an address table that determines which port the address resides on, the Catalyst 2800 is able to forward the packet only to that port. If it is determined that the destination address resides on the port the packet was received from, the packet is filtered and not forwarded.

If a destination address is unknown to the Catalyst 2800, it will flood the packet to all ports. Flooding ensures that the packet will arrive at its destination even when the Catalyst 2800 does not know which port the destination address resides on. The Catalyst 2800 also floods multicast and broadcast packets. There are times when it may be useful to disable flooding. See the “Flooding Controls” section in this chapter for more information.

The Catalyst 2800 also supports source port filtering. This enhanced filtering capability only forwards packets to destinations when received on specified ports. Packets to these destinations received on other ports are filtered. These destinations are referred to as restricted static addresses. See the “Port Addressing” section in the “Out-of-Band Management” chapter for more information. If you are using SNMP, see the “Standard MIBs and MIB Extensions” section in the “In-Band Management” chapter.

Switching Modes

The switching mode determines how quickly the Catalyst 2800 can forward a packet and therefore how much latency the packet will experience. Latency is the delay between the time a packet is received on one port and the time it is transmitted out the appropriate destination port. Selecting a switching mode is often a choice between enhanced error checking and lower latency.

The Catalyst 2800 offers three switching modes:

- FastForward (the default)
- FragmentFree
- Store-and-Forward

FastForward and FragmentFree are two forms of *cut-through* switching. The switching mode is defined on the System Menu of the Catalyst 2800 management console, as described in the “System Configuration” section of the “Out-of-Band Management” chapter, or using the appropriate MIB objects listed in the “Standard MIBs and MIB Extensions” section in the “In-Band Management” chapter.

FastForward

FastForward offers the lowest level of latency by immediately forwarding a packet after receiving the destination address. Because FastForward starts forwarding before the entire packet is received, there may be times when packets are relayed with errors. Although this occurs infrequently and the destination network adapter discards the faulty packet upon receipt, the superfluous traffic may be deemed unacceptable in certain environments. Packets forwarded with errors can be reduced by using the FragmentFree option. In FastForward mode, latency is measured first bit received to first bit transmitted or FIFO.

FragmentFree

FragmentFree switching filters out collision fragments, the majority of packet *errors*, before forwarding begins. In a properly functioning network, collision fragments must be less than 64 bytes. Anything greater than 64 bytes is a valid packet and is usually received

Switching Modes

without error. FragmentFree switching waits until the received packet has been determined not to be a collision fragment before forwarding the packet. In FragmentFree mode, latency is measured as FIFO.

Store-and-Forward

The third switching mode supported by the Catalyst 2800 is the traditional Store-and-Forward mode. Complete packets are stored and checked for errors prior to transmission. In Store-and-Forward mode, latency is measured last bit received to first bit transmitted or LIFO. This does not include the time it takes to receive the entire packet, which can vary according to packet size, from 65 microseconds to 1.3 milliseconds. Store-and-Forward is the most error-free form of switching, but the forwarding latency is higher than either of the two cut-through switching modes. This can be seen in Table 3-1.

Table 3-1 Catalyst 2800 Switching Latencies

Switching Mode	10 Mbps to 10 Mbps	10 Mbps to 100 Mbps	100 Mbps to 100 Mbps	100 Mbps to 10 Mbps
FastForward (FIFO) ¹	31 microsec.	na	7 microsec.	7 microsec.
FragmentFree (FIFO)	70 microsec.	na	9 microsec.	10 microsec.
Store-and-Forward (LIFO) ²	7 microsec.	7 microsec.	3 microsec.	3 microsec.

1. First In, First Out

2. Last In, First Out

Note Although Table 3-1 shows Store-and-Forward experiencing the lowest latency, the figures do not include the time it takes to receive the packet, which varies according to the packet size.

Selecting a Switching Mode

FastForward is the default switching mode and provides the lowest latency packet switching. If, however, your attached networks experience a significant number of collisions, FragmentFree can be used to eliminate the chance of forwarding collision fragments. If you are experiencing frame check sequence (FCS) or alignment errors, Store-and-Forward could be used to ensure that packets with errors are filtered and not propagated to the rest of the network.

You select a switching mode for the entire switch; the chosen switching mode is then applied to traffic flowing between all ports with the following exceptions:

- Store-and-Forward is always used for transfers from 10-Mbps to 100-Mbps ports.
- Store-and-Forward is always used for broadcast packets; multicast packets can be Store-and-Forward or the switching mode set for the Catalyst 2800.

To define the switching mode with the Catalyst 2800 management console, turn to the “System Configuration” section in the “Out-of-Band Management” chapter for more information. You can also define the switching mode in-band using any SNMP-compatible management station. Refer to the *Catalyst 2000 MIB Reference Manual* for a description of the Catalyst 2800 MIB objects. The switching mode is set with the MIB objects listed in Table 6-1 in the “In-Band Management” chapter.

Securing Ports

Secured ports restrict the use of a port to a user-defined group of stations. The number of devices on a secured port can range from 1 to 132. The addresses for the devices on a secure port are statically assigned by an administrator or *sticky learned*. Sticky learning takes place when the address table for a port that is set as secured does not contain a full complement of static addresses. The port sticky-learns the source address of incoming packets and automatically assigns them as static addresses.

Secured ports generate address-security violations under the following conditions:

- When the address table of a secured port is full and the address of an incoming packet is not found in the table
- When an incoming packet has a source address statically assigned to another port.

Securing Ports

When a security violation occurs, the port may be suspended or disabled, as described in the “Port Status” section in this chapter, and SNMP traps may be generated. You can also choose to ignore the violation and keep the port enabled. What action is taken by the Catalyst 2800 is defined by the administrator using the MIB objects listed in Table 6-1 in the “In-Band Management” chapter, or with the Management Console menu described in the “System Configuration” section in the “Out-of-Band Management” chapter.

Note To fully secure a port, you must also disable flooding to the port. See the “Flooding Controls” section in this chapter for more information.

Applications of Secured Ports

There are several reasons you may want to secure ports. If you define a secure port’s address table to contain only one address, the workstation or server attached to that port is guaranteed the full bandwidth of the port. This is the private Ethernet configuration.

Address security can also be used to ensure that only members of a workgroup have access to the switch on a given port. By assigning static addresses to a secure port, the Catalyst 2800 will not forward any packets with source addresses outside the group. Secured ports used in conjunction with VLANs allow for a completely secure switch.

See the “Port Addressing” section in the “Out-of-Band Management” chapter for more on securing ports with the management console. If you are using SNMP, Table 6-1 in the “In-Band Management” chapter lists the MIB objects used for securing ports.

Port Status

Port status is a system-wide indicator of the state of a port. A port's status can change in response to security violations, by management intervention, or by actions of the Spanning-Tree Protocol. At any given time, each Catalyst 2800 port will be in one of three states:

Enabled	The port is active and receiving and transmitting packets.
Suspended	The port is not active but will be automatically returned to the enabled state when the condition causing its suspension is removed.
Disabled	The port is inactive and must be manually returned to an enabled state.

No packets are forwarded to or from a disabled or suspended port. Suspended ports do monitor incoming packets, however, to look for an activating condition. If a linkbeat returns, for example, a port suspended due to linkbeat failure returns to the enabled state.

Full-Duplex 100Base-TX or 100Base-FX

The Catalyst 2800 CollisionFree feature provides full-duplex operation on switched 100Base-T ports for up to 200 Mbps of bandwidth. Full-duplex is the simultaneous transmission and reception of 100-Mbps streams. As both ends of the link must be configured for full-duplex, a full-duplex port cannot be connected to a repeater. A likely scenario would be to connect a one-port 100Base-TX or 100Base-FX module to a server with a 100Base-T adapter configured for full-duplex. You could also connect it to another Catalyst 2800 or other 100Base-TX or 100Base-FX switch configured for full-duplex operation. 100Base-FX full-duplex links can span distances of up to 2 km.

See the "Port Configuration" section in the "Out-of-Band Management" chapter for instructions on configuring ports for full-duplex. If you are using SNMP, Table 6-1 in the "In-Band Management" chapter lists the MIB objects required to configure full-duplex operation.

Port Monitoring

The Catalyst 2800's design has the added security feature of precluding eavesdropping or monitoring of traffic destined for other ports. There are times, however, when you will want to use an RMON probe or sniffer to troubleshoot network problems by examining traffic on some or all of the Catalyst 2800 ports. The Catalyst 2800 provides a port-monitoring mode for this purpose.

Port-monitoring mode forwards frames received from ports that are assigned to a capture list and forwards them to the port designated as the monitor port. This list can include any number of Catalyst 2800 ports from none to all twenty-seven. To enable port-monitoring mode with the management console, see the "Monitoring Configuration" section in the "Out-of-Band Management" chapter. If you are using SNMP, the MIB objects for configuring port monitoring are listed in Table 6-1 in the "In-Band Management" chapter.

Flooding Controls

Flooding is the forwarding of broadcast, multicast, and unicast packets with unknown destination addresses to all ports. In certain applications, this flooding may be unnecessary and undesirable. Note that when you create a VLAN, all traffic—including broadcast traffic—is kept within the boundaries of the VLAN, regardless of whether flooding controls are used.

Unicast Filtering

When the Catalyst 2800 receives a unicast packet with a destination address that it has not learned, the default is to flood it to all ports. However, on ports with only statically assigned addresses or single stations attached, there are no unknown destinations on these ports and flooding would serve no purpose. You can disable flooding in this case on a per-port basis. See the "Port Addressing" section in the "Out-of-Band Management" chapter for more information on disabling the flooding of packets with the management console. If you are using SNMP, the MIB objects for configuring this function are listed in Table 6-1 in the "In-Band Management" chapter.

Multicast Registration and Filtering

When the Catalyst 2800 receives a multicast or broadcast packet, the default is to flood it to all ports. You can use the Catalyst 2800 management console or SNMP to register multicast addresses and list the ports these packets are to be forwarded to. You can also disable the normal flooding of unregistered multicast packets on a per-port basis. Besides reducing unnecessary traffic, these features open up the possibility of using multicast packets for dedicated groupcast applications such as broadcast video. See the “Multicast Registration” section in the “Out-of-Band Management” chapter for more information about registering multicast addresses with the management console. If you are using SNMP, the MIB objects for configuring this function are listed in Table 6-1 in the “In-Band Management” chapter.

Virtual LANs


The Catalyst 2800 supports up to four VLANs. This capability allows ports on the switch to be grouped into separate logical networks. Unicast, broadcast, and multicast packets are forwarded (and flooded) only to those stations within a VLAN, creating a virtual firewall between VLANs. A port can be assigned to a VLAN using the Catalyst 2800 management console or in-band using any SNMP-compatible management station.

Because a VLAN is considered a separate logical network, it contains its own bridge MIB information and supports its own implementation of some of the features described in this section, including Spanning-Tree Protocol.

Configuring Simple VLANs

One of the primary benefits of VLANs is simplified adds, moves, and changes. Figure 3-1 illustrates a simple VLAN configuration that groups workstations attached to the Catalyst 2800 ports 1 through 12, and a server attached to port F2, into VLAN1. Workstations attached to ports 13 through 24 are grouped with a server attached to port F1 into VLAN2. The traffic on each VLAN is completely isolated from the other.

Figure 3-1 Workgroups Defined as Virtual LANs



Users can change their physical location in the network without changing their associated VLAN membership. Consider the example in Figure 3-1 where a user attached to port 3 in VLAN1 moves to be co-located with users in VLAN2 but needs to retain network access to resources in VLAN1. The physical network topology has changed because this user is now attached to a different port on the Catalyst 2800. However, the administrator maintains the logical network structure by assigning the new port to VLAN1. See the “Virtual LAN Configuration” section in the “Out-of-Band Management” chapter for details on assigning ports to VLANs. If you are using SNMP, the MIB objects for configuring this function are listed in Table 6-1 in the “In-Band Management” chapter.

Overlapping VLANs


The Catalyst 2800 supports overlapping VLANs so resources can be shared between logical workgroups. Figure 3-2 shows two VLANs, VLAN1 and VLAN2, sharing a server connected to port F1. Stations in both VLAN1 and VLAN2 can access the server, but otherwise they are completely isolated from each other.

Note Two nonoverlapping VLANs can be connected by two router interfaces. IP routers can route between two overlapped VLANs by connecting one interface to an overlapped port. Such IP routers must support multiple subnets on one interface.



Caution Spanning Tree may not prevent network loops in overlapping VLANs if the VLANs are connected in other ways besides the overlapping port.

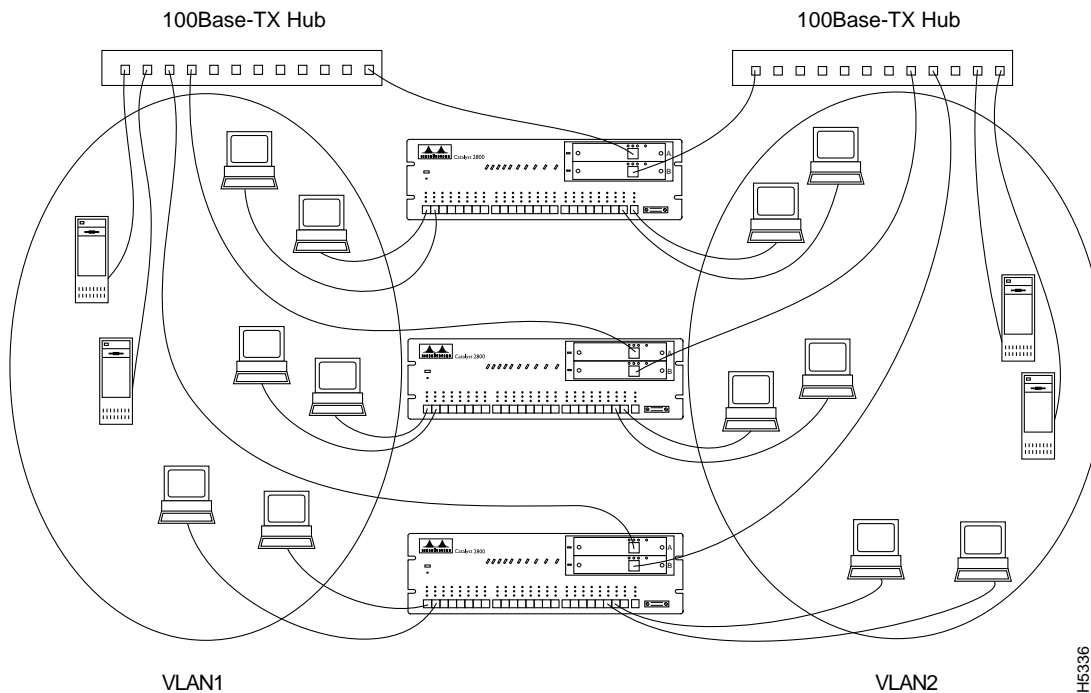
Figure 3-2 Overlapping VLANs Sharing a High-Speed Port



VLANs Spanning Multiple Switches

VLANs spanning multiple Catalyst 2800s require a dedicated port on each switch for each VLAN interconnection. The configuration in Figure 3-3 shows two VLANs spanning multiple Catalyst 2800s. In this example, ports F1 and F2 on each Catalyst 2800 provide the link to interconnect VLAN1 and VLAN2, respectively. An external 100Base-T hub (or FDDI concentrator) then connects each link to form the overall VLAN. If one or more of the switched 100Base-T ports on the Catalyst 2800 are used for other connections, or if more than two VLANs span multiple Catalyst 2800s, then a dedicated 10Base-T port on each switch is required to interconnect the VLAN.

Figure 3-3 VLANs Spanning Multiple Switches



Spanning-Tree Protocol

Spanning-Tree Protocol is a standardized mechanism for maintaining a network of multiple bridges or switches. As part of the IEEE 802.1d standard, Spanning-Tree Protocol will interoperate with compliant bridges and switches from other vendors. It transparently reconfigures bridges when the topology changes to avoid the creation of loops and to establish redundant paths in event of lost connections. All twenty-seven ports are included in Catalyst 2800 Spanning-Tree Protocol support, and management of Spanning-Tree Protocol is through the standard bridge MIB.

Note Each Catalyst 2800 VLAN is treated as a separate bridge, and a separate instance of VLAN and bridge MIB is applied to each VLAN.

Using Spanning-Tree Protocol to Support Redundant Connectivity

You can create a redundant backbone with Spanning-Tree Protocol by connecting two of the switch's ports to another device, or to two different devices. Spanning-Tree Protocol will automatically disable one port but enable it if the other port is lost. If one link is high-speed and the other low-speed, the low-speed link is always disabled. If the speed of the two links is the same, the port priority and port ID are added together and the link with the lowest value is disabled.

Spanning-Tree Protocol and Accelerated Address Aging

Dynamic addresses are aged and dropped from the address table after a configurable period of time. The default for aging dynamic addresses is five minutes. A reconfiguration of the spanning tree, however, can cause many station locations to change. As this would mean that many stations were unreachable for five minutes or more, the address-aging time is accelerated so that station addresses can be dropped from the address table and then relearned. The accelerated aging is the same as the forward-delay parameter value when Spanning-Tree Protocol reconfigures. This parameter is described in the "Spanning-Tree Configuration" section in the "Out-of-Band Management" chapter. If you are using SNMP, Table 6-3 in the "In-Band Management" chapter lists the MIB objects used to configure this function.

Spanning-Tree Protocol
