CHAPTER 2

Catalyst 3000 Theory of Operation

This chapter explains how the Catalyst 3000 improves network performance.

Topics discussed in this chapter:

- Ethernet Network Performance
- How the Catalyst 3000 Works

Ethernet Network Performance

The inherent limitations of Ethernet impede network performance for three main reasons:

- Ethernet supports only one conversation at a time.
- As network utilization increases, throughput decreases.
- High-bandwidth devices and applications quickly consume available bandwidth.

The Catalyst 3000 increases Ethernet throughput by supporting simultaneous, parallel conversations. Full-duplex communication mode is an option for each segment connected to a Catalyst 3000 port.

Rather than storing a whole packet before it is forwarded, the Catalyst 3000 uses error-free cut-through switching so that new connections can be made for the next packet almost immediately.

The Catalyst 3000 can solve congestion problems by dedicating a 10 or 100 Mbps segment to each high-bandwidth device and each powerful application, as well as each user.

The next section explains how the Catalyst 3000 operates and improves network performance.

How the Catalyst 3000 Works

The Catalyst 3000 is an IEEE 802.3-compliant device designed to boost throughput on Ethernet networks. It operates as Media Access Control (MAC)-layer device that is protocol independent; therefore, they are fully compatible with TCP/IP, DECnet, LAT, XNS, AppleTalk, and NetWare.

There are two ways of configuring Catalyst 3000s, either as a single stand-alone unit or as a logical combination of up to eight units. The logical combination of units is called an Catalyst Stack.

This chapter describes how the Catalyst 3000 operates as a single stand-alone unit. See Chapter 3 “Catalyst 3000 Stack Theory of Operation” for an explanation of how to logically combine Catalyst 3000 units into a Stack.

The Catalyst 3000 contains four main elements as listed below:

- Cross-point switch matrix—establishes switched connections between two segments. Each connection lasts only for the duration of the packet transmission.
- 10Base-T Ports—with the 10Base-T ports, multiple conversations are allowed. Now users running basic applications are able to share bandwidth, and users running bandwidth-intensive applications can receive their own dedicated 10 Mbps port. With the optional Enhanced version of the Catalyst 3000, each 10Base-T port can be set up in full-duplex communication mode so that each 10 Mbps port doubles to 20 Mbps.
- AUI connector—the AUI connector can be attached to an external transceiver for connection to another media type such as 10Base-FL.
- 100Base-TX Expansion Modules—each Catalyst 3000 supports two high-speed connections, Fast Ethernet for servers or backbone connectivity. This new Fast Ethernet technology is provided by the Catalyst Stack Port module, which conforms to the IEEE 802.3U 100Base-TX standard. Each Catalyst Stack Port module supports a single 100 Mbps port for Fast Ethernet over Category 5 cabling. They are designed to provide connectivity to servers, routers, and workstations with 100Base-TX standard-compliant interface cards. The Catalyst Stack Port module is user configured to support full-duplex Ethernet, for an effective throughput of 200 Mbps. The Catalyst Stack Port module has expanded buffers that support up to 2000 packets.

Multiple Simultaneous Conversations

A limitation of Ethernet is that it supports only one conversation at a time. The Catalyst 3000 improves data throughput by supporting multiple, simultaneous, full-duplex conversations. Combining fast-packet, circuit-switching technology, and Fast Ethernet technology (also called Fat Pipes), the Catalyst 3000 creates multiple data paths. These switched connections between Ethernet segments last only for the duration of the packet transmission. New connections are made "on-the-fly" between different segments for the next packet.

Figure 2-1 Multiple Conversations Through a Catalyst 3000

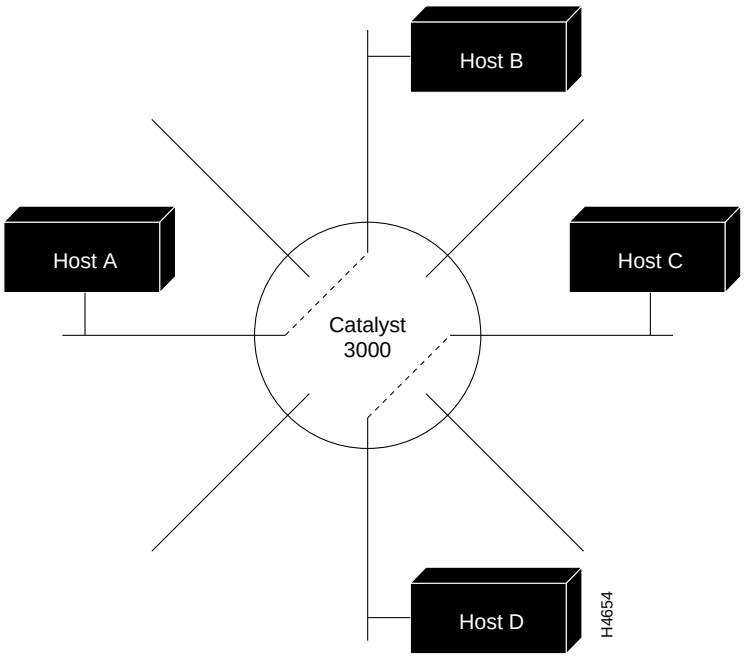

For example, as shown in Figure 2-1, while host A is transmitting a packet to host B, the Catalyst 3000 connects only the lines from A to B since there is no need to send packets to all other ports. At the same time, a second switching circuit can connect host C to host D. The result: two conversations occur simultaneously.

Note The Catalyst 3000 transmits broadcast and multicast packets on all Catalyst 3000 segments simultaneously, except the port of entry.

The increase in throughput is directly proportional to the number of LAN segments that are interconnected through the switch. A Catalyst 3000 with 16 ports interconnected provides eight concurrent paths. With eight simultaneous conversations, the Catalyst 3000 creates 80 Mbps throughput in half-duplex mode, or 160 Mbps throughput in full-duplex mode.

A single segment can be dedicated to a single host or shared by several. To optimize throughput, high-speed servers can be given dedicated Catalyst 3000 ports.

If you think of a standard Ethernet network as a “party-line” telephone system, a network using Catalyst 3000 units is like a PBX that supports multiple telephone conversations. By transporting multiple Ethernet packets simultaneously, it boosts overall network throughput.

Low Latency

The Catalyst 3000 minimizes latency—the time it takes to forward a packet from one Ethernet segment to another—by beginning switching immediately after looking at the first six bytes of the destination address in the packet. If the packet needs to be switched to another LAN segment, its data begins flowing through the destination port before the entire packet has been received. The result: Packets appear at the output port 40 microseconds after entering the input port. Network devices that use store-and-forward technology introduce much longer delays because they wait to receive the entire packet before forwarding it. For example, routers and bridges that use conventional, software-based store-and-forward technology impose a latency of 1,200 and 800 microseconds per packet, respectively.

By minimizing delay, the Catalyst 3000 can move more packets freely throughout the LAN without degrading performance.

Address Management

At power up, the system address tables do not contain any information. Whenever a Catalyst 3000 receives a packet with an unknown source or destination address, before it sends the packet, it learns the new source address by putting its location into memory. Once the Catalyst 3000 has learned the new source address, it then sends the packet out to its destination address. If the destination address is unknown it sends the packet to all of its output ports. When the response packet comes back, the Catalyst 3000 learns its location and adds it to the address table. Once the address table entry is created, the Catalyst 3000 uses all of these learned address to switch all subsequently addressed packets with less processing, and therefore, faster transmissions.

The system address table maintains up to 10,000 entries with 4 MB of DRAM memory, and up to 10,000 entries with 8 MB of DRAM. The port address tables maintain 1,700 active Ethernet addresses. On-demand aging allows users to set a threshold based on a specific time interval or a percentage of address table capacity. This ensures that the port's address table is populated only by the most frequently used address. You can adjust the aging interval to keep the address tables under the limit. Nodes that have not transmitted after the aging interval you specify on the network management console are removed. When new nodes become active, configuration is re-established with only one packet.

This capability allows users to transparently connect to high-volume backbone networks.

Address Filtering

The Catalyst 3000 supports the configuration of MAC layer filters on a per-port basis. This flexibility allows network managers to specify client access only to designated resources for security purposes. Filters can be for source or destination addresses. This allows the network manager to restrict access to certain servers or MAC addresses or to specify that an end user can communicate with only one server.

On-board Buffering

If the destination port is receiving a packet from another Catalyst 3000 port or if the output segment is busy, the Catalyst 3000 stores the packet in one of its internal buffers. Each Catalyst 3000 buffer can hold up to 384 packets in each direction (incoming and outgoing). This helps balance throughput when networks are operating near peak load and more than one packet may be directed to the same port at the same time.

Enhanced Features Provide Scalability

The Catalyst 3000 with the optional Enhanced feature set provides two options for scalability: full-duplex communication and EtherChannel technology. Using these features enables you to increase your network's bandwidth in scale with growing demand.

- Full-Duplex communication.
  - — You can select half-duplex or full-duplex communication for all 10Base-T ports and 100Base-TX modules, doubling throughput capacity. The advantage of using full-duplex is that communication packets can flow in both directions simultaneously, which results in increasing the throughput capacity on the segment.
  - — Full-duplex communication eliminates the performance degradation resulting from packet collisions. Packets cannot collide because they each travel on their own path—like cars going in opposite directions on a two-lane highway. So while the effective bandwidth to a 10Base-T port configured for half-duplex Ethernet is a maximum of 10 Mbps, with full-duplex Ethernet it is doubled to 20 Mbps.

Note You should select full-duplex communication only if the port is connected to another full-duplex device. Full-duplex Ethernet adapters are currently available from IBM, Compaq, National Semiconductor, and Cogent Technologies.

- EtherChannel.
  - — A high-bandwidth connection between two Catalyst units. Using the EtherChannel feature, you can connect the Catalyst 3000 to Catalyst Stack Port modules, or to other Catalyst 3000 units. Each EtherChannel comprises from two to seven ports, for up to 140 Mbps bandwidth in Full-Duplex mode. By connecting three cables between two Catalyst 3000 devices, for example, you would increase throughput to 30 Mbps in Half-Duplex mode, or 60 Mbps in Full-Duplex mode. Network managers can connect multiple 10 Mbps ports to create a single "fat pipe." Existing Catalyst units can gain access to high-speed servers, routers, and backbones connected to the Fast Ethernet ports through the EtherChannel connections to an Catalyst 3000.

- Virtual LAN (VLAN)
  - — Using the Catalyst VLAN feature, you can partition a single Catalyst 3000 into VLANs (also referred to as Catalyst VLANs), each containing its own set of ports. Packets are forwarded only between ports belonging to the same VLAN. The benefit of Catalyst VLAN is to restrict access from one segment to another, either for security purposes or to reduce intersegment traffic.

Error Handling

A major source of errors in Ethernet networks is packet collisions that are due to low network throughput. The Catalyst 3000 uses cut-through switching to reduce latency which increases throughput. The network interface cards of the destination servers or workstations check for any packet errors that the Catalyst 3000 forwards to them, so to maintain low latency, the Catalyst 3000 can be set to not duplicate this function. In contrast, routers and bridges that perform redundant store-and-forward error checking on all packets impose delays of between 51.2 and 1,214 microseconds. The reason for the delay is because they must receive the entire packet before forwarding it to its destination port, since the Frame Check Sequence (FCS) comprises the final 32 bits of an IEEE 802.3 packet.

Individual ports on the Catalyst 3000 can be set to run in a dedicated low latency, error-free cut-through switching mode, or they can be set to run only in the store-and-forward mode, or they can be set to a mode of automatic switching from cut-through to store-and-forward. The automatic mode measures the amount of CRC errors in cut-through mode and if the level rises above a configurable setting, error handling switches to store-and-forward. Error handling will automatically switch back to cut-through when the error level falls below the set amount.

For connections between ports of different speeds (10Base to 100Base for example) the error handling will be in the store-and-forward mode only.

The Catalyst 3000 can also be set to check for incomplete packets and discard them. See Chapter 7, “Console Configuration” for a detailed explanation and configuration of the error handling modes for the Catalyst 3000.