Chapter 2. Packet Switching Architecture

This chapter covers the following key topics:

- Process Switching
- Fast Switching
- Optimum Switching
- Cisco Express Forwarding

A multiprotocol router's primary purpose is, of course, to switch packets from one network segment to another. If the scheduler and memory manager are the router's software infrastructure, then the IOS switching architecture is the router's very foundation. The switching methods and structures used by IOS essentially dictate how well the router performs its primary task—so, as you can imagine, a great deal of effort has been put into designing and refining this important part of IOS.

In spite of all the fanfare surrounding packet switching, the operation itself is fairly straightforward:

Step 1. A packet comes into an interface.

Step 2. The packet's destination address is inspected and compared against a list of known destinations.

Step 3. If a match is found, the packet is forwarded out the appropriate interface.

Step 4. If a match is not found, the packet is dropped.

Clearly, this isn't rocket science. Indeed, the problem to be solved isn't so much *how* to switch packets, but how to switch them *quickly*. Switching packets is a data-intensive operation as opposed to a computation-intensive one; speeding up the operation isn't as simple as just using a faster CPU. Other factors, such as I/O bus performance and data memory speed, can have a big impact on switching performance. The challenge for IOS developers is to create the highest possible switching performance given the limits of available CPU, I/O bus, and memory technology.

As the size and number of routed networks grow, IOS developers continuously look for better ways to solve this performance challenge, resulting in a continuous redesign and refinement of IOS switching methods. When IOS was first developed, there was only one switching method, known today as process switching. Later releases introduced newer improved methods of switching; some of these methods rely on hardware-specific optimizations, others use software techniques that work well on many platforms. Today, IOS can switch several hundred thousand packets per second using routing tables containing hundreds of thousands of routes.

The following list summarizes the switching methods developed as of Cisco IOS Release 12.0:

- Process switching
- Fast switching
- Autonomous switching
- Silicon switching engine (SSE) switching

- Optimum switching
- Distributed fast switching
- Cisco Express Forwarding (CEF)
- Distributed Cisco Express Forwarding (dCEF)

This chapter covers in detail four of these methods—process switching, fast switching, optimum switching, and Cisco Express Forwarding (CEF). Autonomous and SSE switching are platform-specific and aren't commonly used in today's networks, so they aren't covered here. Distributed fast switching is actually an implementation of the optimum switching method on intelligent line cards, and isn't covered any further than the coverage provided to optimum switching itself.

It's worth noting that although IP routing examples are used here to illustrate the switching methods, many of these methods also apply to other network protocols, such as IPX and bridging. Although the structures used are often independent for each protocol (for example, there is a separate Fast Cache for IP and IPX), the contents of the switching structures are similar, and switching occurs in essentially the same way for each protocol.

Last updated on 12/5/2001

< BACK | Make Note | Bookmark | CONTINUE >

Index terms contained in this section