LAYER 3 SWITCHING

By Mark E. Donaldson

INTRODUCTION

At the heart of the next-generation network is a new breed of switches that allow networks to be designed and run more efficiently than routed or shared-medium networks. Understanding the differences among Layer 2, Layer 3 and Layer 4 switches starts with a definition of switching as it applies to networking. Originally, the term described a device that could make frame-forwarding decisions based on Layer 2 information while building its own forwarding table. Because Layer 2 switches operate at the data link layer and have access to MAC (media access control) addresses, they are able to identify the destination address of each frame and forward it only to that station. Thus, the first Layer 2 switches (back in 1993) resulted in a bandwidth increase that would later open up new possibilities for applications such as graphics-rich intranets.

However, even with additional bandwidth, today's Layer 2 switches, such as Cisco Systems Inc.'s Catalyst 5000, create large, flat networks that can be quickly overwhelmed with ARPs (Address Resolution Protocols), service access point requests, RIP (Routing Information Protocol) updates and so on. Before Layer 3 switching came along, network administrators' only solution to this was to break up the network with VLANs (virtual LANs) or routers. Neither solution is ideal, VLANs are confusing and difficult to implement; routers are slow and expensive. Many administrators worked around the problem by placing servers near clients, thus reducing the traffic volume going through the router. This worked for a while, back when the 80/20 rule was still in force: 80 percent of the traffic stays local, 20 percent travels through the router. When intranets arrived, traffic patterns changed dramatically, with about half of all the traffic going through routers to another network segment.

LAYER THREE TECHNOLOGY

Because traditional routers weren't fast enough to handle the load, in 1995 vendors such as Ipsilon Networks Inc. began offering switches with the first "route once, switch many" technology, known as cut-through switching. Although these switches were an improvement over Layer 2 technology, they can't be considered true Layer 3 switches because the flow technology they use causes some packets, and often entire classes of traffic, to be routed rather than switched.

The next group of products, pioneered last year by Bay Networks Inc. with SwitchNode and Foundry Networks Inc. with NetIron, were router-assisted switches that could forward IP traffic at wire speeds based on a routing table created by an external router. These devices, also known as cache switches or routing switches, are very useful for accelerating IP traffic on medium-size networks with a huge traffic increase due to intranets. They can forward IP traffic at a million packets per second, while allowing legacy protocols to take a slower path through traditional routers. However, because they must be paired with an external router, these switches cannot be considered truly Layer 3.

It was not until Extreme Networks Inc. began shipping its Summit II switch in the fall of 1997 that network administrators could purchase a Layer 3 switch. The Summit II runs at wire speed, makes forwarding decisions based on Layer 3 information, swaps the MAC addresses, runs traditional routing protocols such as RIP and decreases the time-to-live field of each packet.

LAYER FOUR TECHNOLOGY

As if things were not complicated enough by the many varieties of Layer 3 switches, vendors are touting switches that use Layer 4 information (such as port number) to make their forwarding decisions. However, going back to our original definition of a switch, until these switches can build a forwarding database based on port numbers, they cannot be considered Layer 4 switches.

While some network managers question the need for the multitude of switching layers, each mode lends itself to a specific purpose. The simplest types of switching, Layer 2, is best applied at the desktop and workgroup level. The fact is, with the price per 100M-bit port dropping rapidly, almost all network administrators should replace their hubs with Layer 2 switches. Cut-through switches, now known as MPLS (Multi-Protocol Label Switching) switches, are too complex for the LAN and are quickly losing ground in the face of the dropping price in true Layer 3 switches. Cut-through switches, including those using Cisco's Tag Switching and 3Com Corp.'s Fast IP technology, are getting a second chance in the WAN. In situations where there are extensive router hops and large routing tables, MPLS can provide extensive benefits, the most touted of which is route aggregation.

ROUTE SWITCHES

Route switches fit well in the LAN, costing about half as much as full Layer 3 switches and offering many of the same functions. Route switches work best on networks with lots of IP traffic and few routing table updates. Route switches can be used as a boost for existing routers--increasing the overall routing speed without changing the network topology. A true Layer 3 switch is nothing less than a wire speed router, and network administrators should be wary of any vendor data sheet that says otherwise.

The main advantages of wire speed routing are in network design. Whereas slow routers must be carefully placed to reduce router hops and bottlenecks, wire speed routers (Layer 3 switches) can be deployed close to the client, offering a granular level of control over traffic patterns. Since Layer 4 switches do not exist in reality, they don't need to figure into current network design. Furthermore, since there is no advantage to building forwarding tables based on layer port number, it's highly unlikely Layer 4 switches will ever exist.

Layer 4 devices, on the other hand, will prove invaluable in the network. By establishing class-of-service parameters based on port numbers, administrators will be able to quickly prioritize their mission-critical traffic.
These devices will also come in handy in networks requiring heavy resource balancing, where they can be programmed with forwarding tables that correspond to individual servers. By dynamically allocating these forwarding tables based on network utilization, servers will be

able to scale more efficiently. However, network managers considering a move to IPSec (Internet Protocol Security) should be aware that because IPSec encrypts an entire packet, it renders all Layer 4 policies useless.

LAYER THREE SWITCHES

The new Layer 3 switches that emerged in 1997 are more like superfast routers than switches. But regardless of what they're called, these nimble devices give companies the chance to construct their networks in a new way. Layer 3 switches make it possible to flatten the structure of the hierarchical networks to which we've become accustomed. Rather than creating new subsegments and adding routers to isolate them, administrators can use fewer Layer 3 switches attached directly to larger segments. The result is simpler network designs and a correspondingly higher throughput. Also, because Layer 3 switches use standard routing protocols, they can be easily incorporated into networks without the need to replace equipment.

Layer 3 switches have drawbacks. For example, the comparative immaturity of the first generation of these switches means companies will still need traditional routers. In fact, all of the Layer 3 switches introduced by the Spring 1998 were lacking in at least one key area. Most notably, all the new high-speed switches lack WAN interfaces for connections such as T-1, ATM or ISDN. In addition, most of them are limited to routing only TCP/IP, and some Layer 3 switches are further restricted to the Routing Information Protocol.

Network managers will also be unhappy with the lack of redundancy in many Layer 3 switches. Any product sitting at the heart of a network requires a higher than normal set of redundancies and failover capabilities. Some products have redundant power supplies, but most of their competitors don't. Amazingly, we've seen several Layer 3 switches that didn't come with RMON support, which is essential for analyzing backbone traffic when troubleshooting.

The good news is that Layer 3 switches are maturing quickly. Most vendors plan to release downloadable software upgrades in early 1998 to expand their protocol support and management capabilities. However, unit replacements will be necessary to obtain features like WAN connectivity and system redundancy. By the end of 1998, we can expect to see the first Layer 3 boxes that match traditional routers feature for feature.

The main reason for the remarkable performance of Layer 3 switches is a new crop of computer chips called ASICs (application-specific integrated circuits), which have been tailored to handle packet forwarding. Traditionally, routers have used only one or two multifunction processors to execute instructions stored in the downloaded router software or operating system. In contrast, Layer 3 switches have both dedicated ASICs and multifunction CPUs. The ASICs help off-load some of the processing work from the CPU, and the multifunction CPUs ensure that Layer 3 switches are still software upgradable. Along with their routing capabilities, many Layer 3 switches include standard Layer 2 switching capabilities. Just as the distinction between bridges and routers became muddied long ago, so routers and switches are becoming increasingly less defined.

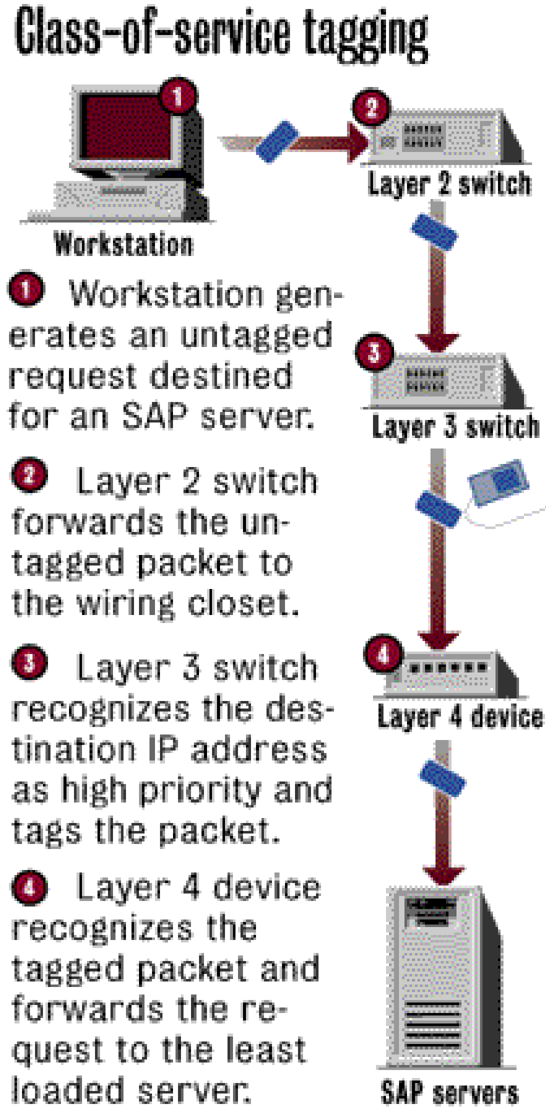

LAYER FOUR SWITCHES

Some new products are reportedly capable of switching at the fourth layer (the application level). These Layer 4 "switches" have extremely limited capabilities, and they do not yet handle the more mundane Layer 2 or 3 packet handling needed by all networks. As such, they will find only a limited use. We should also note that not all Layer 3 switches use ASICs. Some products incorporate an optimized routing technique called cut-through, which avoids the need for ASICs by using some fancy footwork. Cut-through allows Layer 3 switches to improve throughput by processing the routes for only the first packet in a given session. Once the initial route has been found, cut-through switches establish a virtual session between the host and client machines for the rest of the packets. The advantage of this technique is that it conserves processing cycles on router CPUs, but cut-through switching requires creating completely new protocols to handle the virtual sessions and packet forwarding.

Wire speed--that's what the latest Layer 3 switches are striving for. By performing packet-by-packet routing techniques in ASICs, these devices achieve speeds of 4 million, 7 million, even 30 million packets per second. Cabletron's SmartSwitch router uses cut-through Layer 3 switching technology in its current line of SecureFast switches. With packet-by-packet Layer 3 switching, each packet is examined and forwarded to its destination. Cut-through Layer 3 switching examines the first packets to determine the destination and sends the rest through. But Cabletron sees a place for both Layer 3 switching techniques.

One advantage of cut-through Layer 3 switching is lower price. Drawbacks to cut-through Layer 3 switching include networkwide software changes, the chance of errors and the use of proprietary techniques. Packet-by-packet switches read everything into a buffer and make sure there are no defects before forwarding a transmission, while cut-through switches send packets out before the whole sequence is in.

Features to look for in Layer 3 switches

- Both IP and IPX support
- WAN interfaces such as T-1 and T-3
- Redundant power supplies
- Redundant CPU modules
- RMON support
- Web-based management
- Fast Ethernet, Gigabit Ethernet support
- ATM, FDDI support (or plans)