digital™

*Networks*

# *IP Switching from Digital*

## *Optimized Switching for the Internet Protocol with the Industry's Fastest IP Switching Solution*

**Program Backgrounder**


Digital Equipment Corporation
Network Product Business
550 King Street, Littleton, Massachusetts 01460 USA

http://www.networks.digital.com/

# Why IP Switching?

The major challenge for network managers today is improving network throughput. Switches deliver higher performance at lower cost than traditional routers by implementing selected key tasks in hardware. However, often switches are limited in the network topologies they can support. As a result, network layer routing technology has been added to some switching architectures, such as Digital's *en*VISN network architecture, to allow more robust, scalable networks to be constructed. This approach combines the cost/performance advantages of switches with the topological flexibility of routing. (See Figure 1.)

For corporate backbone environments that are deploying IP-based applications (in particular, Internet/intranet), maximizing network throughput cost-effectively has become increasingly crucial. It was also inevitable that a solution that optimizes the network for the Internet Protocol (IP) would emerge. IP Switching provides a high-performance, cost-effective solution to build robust IP-based networks.

# What is an IP Switch?

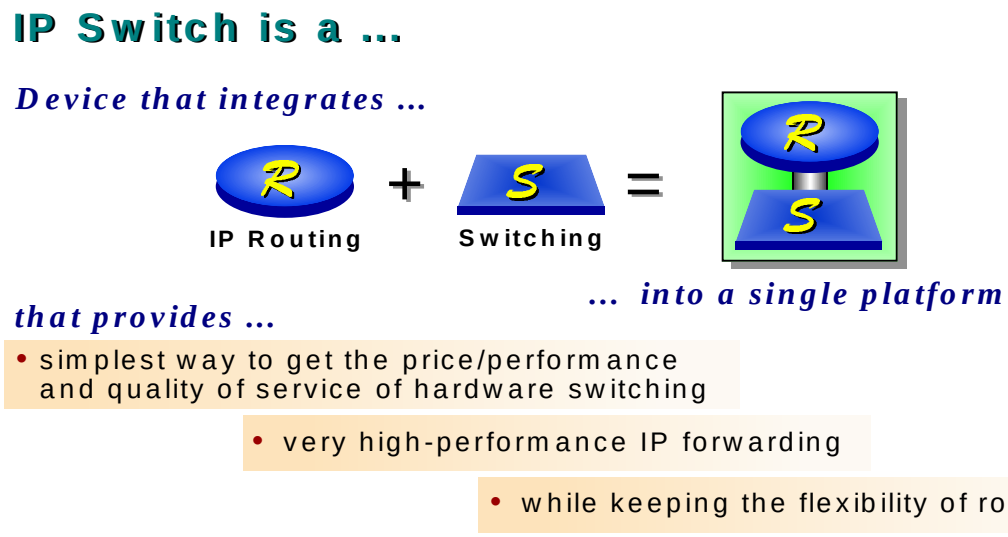


**Figure 1: An IP Switch Defined**

## IP Switching Operation

An IP Switch combines all the benefits of a router with the high performance, low latency and price benefits of a switch. Packets enter an IP Switch from either another IP Switch or an end station, as shown in Figure 2.
IP Switch routing software in the IP Switch examines each packet and makes two decisions. The first decision routes the IP packet just as traditional routers do today. The second decision, called *flow classification*, determines whether this packet represents a potential *flow* that can be switched directly in hardware.

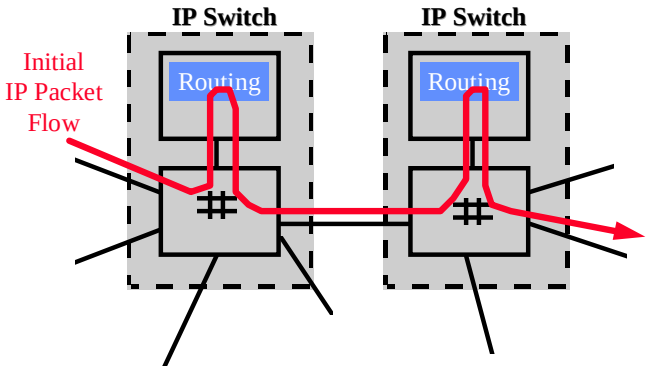


**Figure 2: Initial IP Packet Handled by IP Routing Software in IP Switch**

A flow is an extended IP conversation. It could be an ftp file copy, a download of a web image, a multimedia session, or even an "aggregation" of all traffic between an IP source and destination. The network manager has the flexibility to set the policies that determine how flow classification is made in each IP Switch. The IP Switch optimizes its performance by providing direct hardware switching of packets identified as a flow. Short-lived intermittent traffic that does not qualify as a flow (such as name look-ups or SNMP packets) continues to receive hop-by-hop store-and-forward routing.

Once a flow is identified in an IP Switch, that IP Switch asks the upstream device to re-label the traffic so it can be switched directly in hardware. (See Figure 3). If the upstream device agrees, then subsequent traffic directed to the IP Switch is then redirected. A *control protocol* called *IP Flow Management Protocol (IFMP),* is used for communication between the IP Switches for this purpose.

As shown in Figure 4, in addition to requesting the upstream node to redirect the incoming traffic, the IP Switch resets its internal forwarding hardware to bypass the switch controller and switch the flow directly in the ATM hardware. This design allows IP Switches to forward packets at rates limited only by the aggregate throughput of the underlying switch engine and the datalink speed of the port while maintaining all the benefits of Layer 3 routing.

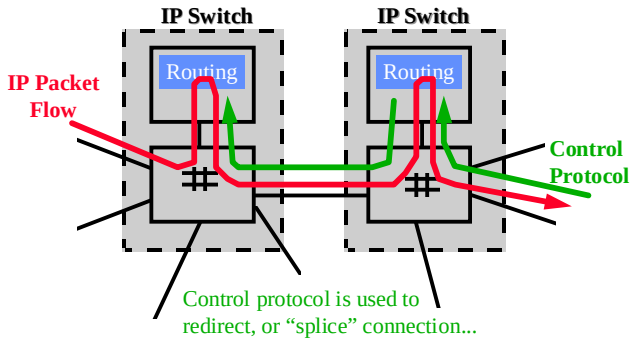


**Figure 3: Flow Classification Determines IP Packets to Switch in Hardware**

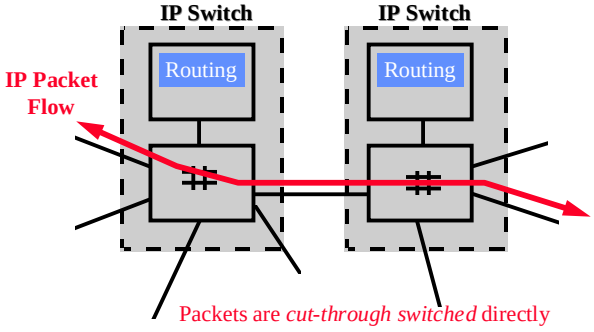


**Figure 4: IP Packets of a Flow Switched in Hardware**

## Flow Classification Locally within a Switch

An important aspect of IP Switching is that flow classification and switching are “soft-state” decisions local to individual IP Switches; they have no impact on the operation of the network. The policies in a specific IP Switch may classify a session as a flow and bypass the routing software, regardless of the policies, decisions or state of other IP Switches in the end to end path of that session. This attribute of IP Switching preserves the connectionless nature of IP and allows each IP Switch to maintain an independent network identity. Network state remains dynamic, allowing IP Switches to route around failures without reestablishing the circuit from the host point. IP Switches use the normal, proven and scalable IP routing protocols (OSPF, RIP, BGP) for this routing operation.

Flow classification policies may reflect differing implementation and/or deployment requirements of each IP Switch. For example, a telecom or Internet service provider may wish to classify flows such as “all traffic between this IP address source/destination pair” or “all traffic to/from this IP subnet” or “only one flow allowed for a multimedia video conference at your station”. Alternatively, an intranet campus network manager may wish to express policies such as “any long lived stream of IP packets may be a flow, but no more than N flows for any station in this subnet,” or “all bulk transfers (ftp, web, etc.) are flows.” IP Switching is inherently a very flexible technology because these flow classification policies can be set by the network manager depending on the requirements for that node at any instant in time. Flow classification can also accommodate new types of applications and traffic as they emerge.

## IP Switching and Multicast Traffic

IP Switches support normal IP multicast without any modification to standard IP multicast protocols. This includes the Distance Vector Multicast Routing Protocol (DVMRP) and the Internet Group Management Protocol (IGMP). From the point of view of an IP Switch, multicast traffic is handled naturally by flow classification. Continuous multimedia streams, such as live video and audio from a group video conference, map cleanly into the IP flow-classification model, allowing the switch to handle multicast replication at hardware speeds. Short-lived multicast traffic is forwarded automatically to multiple destinations by the IP Switch routing software.

# The Business Case for IP Switching

Within both LANs and wide area networks, there is a need for intelligent devices that know the network topology and capacities and can dynamically identify an optimal route through the network. Traditionally, that requirement has been handled by routers, yet routers have typically been relatively slow in both throughput and latency and relatively expensive compared to the switches.

Building large networks such as the Internet or a corporate enterprise network backbone using Internet technology (i.e. an *intranet),* requires network layer routing. However, routers are too slow and costly. Nevertheless, there is a large body of practice (management tools, operational knowledge, etc.) which deals with routing networks – IP in particular. The consolidation towards IP and Internet technologies makes it clear how to optimize the solution – switching IP frames.

IP Switch technology combines the intelligence and control of IP routing with the high speed and capacity of switching hardware to advance the state of the art in internetworking price/performance. Based on scalable, robust, and proven technologies, IP Switches provide millions of IP packets-per-second (PPS) throughput while maintaining full compatibility with existing IP networks, applications, and network management tools.

Digital's strategy is to incorporate IP Switching directly into its full range of department and backbone switches, such as GIGAswitch/ATM, DEChub 900 MultiSwitch, as well as host NIC drivers such as Windows/NT and Digital UNIX.

### Business Case Summary:

Traditional backbone routers can't keep upas a result of:

- Rapidly increasing business use of Internet and related applications
- Intranets becoming the organizing principle for corporate networks

IP Switching solves the problem:

- Delivers significant price/performance advantages of switches with the topological flexibility of routing

## When to Deploy IP Switching

IP Switching is particularly beneficial for:

Applications that demand maximum IP throughput

- Internet or intranet backbones

LAN environment where LAN Emulation cannot support traffic patterns

- High levels of broadcast or IP Multicast traffic

High-performance IP routing between VLANs

Wide Area IP Networks

- High performance and maximum efficiency of link capacity

## Digital's IP Switching Strategy

IP-based networks and applications are pervasive. Digital's product strategy is to optimize the cost-effective transport of IP data, while supporting an increasing broad range of legacy LAN applications.

Digital is also sensitive to the need for customers to protect their investments in their network infrastructure. Therefore, Digital will incorporate IP Switching technology into existing products and platforms, such as the GIGAswitch/ATM, the DEChub 900 MultiSwitch, and ATMworks host adapter drivers such as Windows/NT, Digital UNIX, and SunOS. Digital will also continue to support and extend its industry leading LAN emulation-based ATM solutions. This combination of IP Switching and LANE will leverage the best of both worlds — the protocol independence of LANE in a campus environment and the optimized end-to-end performance of IP Switching.

## Digital's GIGAswitch/IP Solution

The GIGAswitch/IP Solution is the first in a planned series of IP Switching product offerings from Digital . The GIGAswitch/IP Solution easily integrates into existing IP environments and can be purchased as a complete package or as individual components. The components of the GIGAswitch/IP Solution include the GIGAswitch/IP System, consisting of a GIGAswitch/ATM high-performance switch with General Switch Management Protocol (GSMP) routing software and the GIGAswitch/IP Switch Controller; IP Switch Gateway; ATMworks adapters; and GIGAswitch/IP management application.

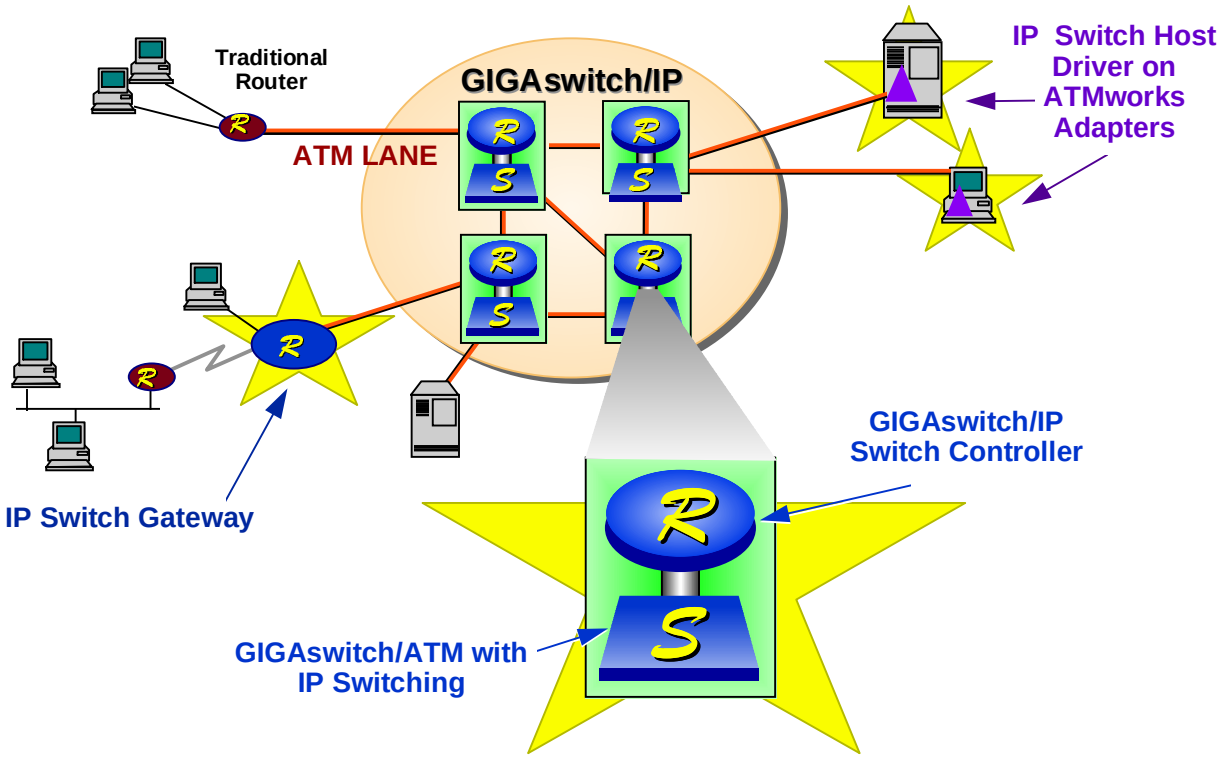


**Figure 5: GIGAswitch/IP Solution**

## *GIGAswitch/IP* Solution Characteristics

- Industry's highest-performance IP throughput (up to 18 million packets per second )
- Over 10 Gb/s of IP Switching bandwidth
- Dynamic store-and-forward routing and cut-through IP Switching
- Support for Ipsilon's *GSMP* and *IFMP* protocols
- Support for standard IP routing protocols, such as OSPF and RIP
- ATM standards compliance with LANE, UNI 3.0/3.1 and PNNI
- Support for IP unicast and multicast traffic
- Transparent quality-of-service guarantees
- Support for SNMP MIBs and network management tools — no new learning curve
- Web-based device management
- Scalable architecture with link speeds from T1 (1.5 Mb/s) to OC12c (622 Mb/s) and port density up to 52, 155 Mb/s, T1/E1 or T3/E3 ATM ports, up to 13, 622 Mb/s ATM
- Proven, reliable switch hardware (redundant power and fans, 48 volt, etc.)

## IP Switch Gateway and ATMworks Host Adapters

The IP Switch Gateway that supports IFMP is an edge device to connect legacy LANs to the IP Switch backbone. The IP Switch Gateway has a 155 Mb/s ATM link to connect to the GIGAswitch/IP system and the ability to support a

range of LAN adapters including Ethernet, Fast Ethernet and FDDI. The gateway is a 19-inch rackmountable router supporting OSPF, SNMP, and other standard IP protocols.

The ATMworks family of adapters will be upgraded with software support for the IFMP protocol allowing end stations to connect to the GIGAswitch/IP. In the first release of IP Switching products, Digital will support:

- Windows NT on Intel
- Windows NT on Alpha
- Digital UNIX on Alpha
- SunOS on Sbus-based systems

# Summary

The two most important trends in networking today are the use of switches and the expansion of Internet technologies for the Internet/intranets. IP Switching combines these two important technologies. It provides the cost and performance advantages of switches (millions of packets/second at low cost) and the richness and flexibility of proven and extensively deployed IP routing, management, and resource control protocols. Furthermore, it is uniquely positioned to take direct advantage of the natural evolution of the Internet Protocol (IP), including IPv6, RSVP, BGP, MOSPF, and other protocols.

Digital's strategy is to incorporate IP Switching into its *enVISN* network architecture. This provides optimization for IP users and applications while also allowing coexistence and migration for legacy environments.

Digital's initial offering, the GIGAswitch/IP Solution, is the price/performance industry leader in connectivity for corporate networks and

- Combines the best of routing and ATM switching at significant cost savings over traditional backbone routers
- Provides a complete solution and simple integration to typical IP environments

Digital delivers IP Switching PLUS ATM Forum standards

- Dynamic support for IP Switching and LAN Emulation
- Allows customers to exploit inherent QOS ability of ATM between all groups within an enterprise

# More Information

For more information or for the name and location of a reseller near you to serve all your networking needs, please visit our web site:

http://www.networks.digital.com/



Digital Equipment Corporation
Network Product Business
550 King Street, Littleton, Massachusetts 01460 USA