

# Understanding Layer 3 and 4 Classification to Enhance Gigabit Ethernet Deployment

*By Richard Martin*

**Dell® PowerConnect™ switching technologies enable Gigabit Ethernet as well as multilayer prioritization capabilities in network environments of almost any size. This article provides an overview of switching and classification, and suggests deployment designs to maximize networking technologies.**

**M**ore and more organizations are turning to Gigabit Ethernet to meet the increasing demand for bandwidth in their networks. More affordable Gigabit hardware—combined with the fact that Gigabit Ethernet can be deployed without change to the standard Category-5 copper cabling that runs 100 Mbps Ethernet networks—is accelerating Gigabit Ethernet adoption across all areas of the network. In networks with larger distance requirements, Gigabit Ethernet also supports fiber optic cables.

One common application of Gigabit Ethernet is in networks with many 100 Mbps clients where Gigabit Ethernet is required to avoid potential bottlenecks at the network backbone and server level. Many vendors, including Dell, incorporate Gigabit Ethernet into the LOM (LAN-on-Motherboard) on their servers.

In addition to local area network (LAN) deployments, Gigabit Ethernet can provide an alternative to technologies such as SONET (Synchronous Optical Network) in wide and metropolitan area networks (WANs and MANs). The IEEE® 802.3ae standard designating 10 Gigabit Ethernet (GbE), scheduled for ratification before mid-2002, should further expand Ethernet application to MAN and WAN links. The benefit of using the same underlying technology and equipment end-to-end allows for cost savings and increased performance.

Some future storage area network (SAN) technologies, such as Internet SCSI (iSCSI), will require Gigabit Ethernet to transmit data over LANs, WANs, and the Internet. iSCSI is an Internet Engineering Task Force- (IETF-) developed Internet Protocol- (IP-) based storage networking standard that allows servers and clients to be connected to storage devices over low-cost, standard Ethernet equipment.

Gigabit Ethernet deployment has dramatically changed the amount and types of traffic commonly found on LANs and WANs. For example, Gigabit networking lets IT managers converge voice, video, and data across a single network. Although this converged network delivers lower costs and simplified network administration, advanced traffic management and classification features become a necessity to ensure optimal application-level performance.

Dell® PowerConnect™ switches that have Gigabit Ethernet ports and advanced classification features let IT managers maximize the usefulness of their networks. The switches allow for use of the Layer 2, 3, and 4 headers (see ISO®/OSI (International Organization for Standardization/Open Systems Interconnection) reference model) to ensure that high-priority traffic reaches its destination with minimal delay. These devices enable easy-to-use, sophisticated switching capabilities, making commonplace what was previously found only in high-end, proprietary systems.

## Layers 2, 3, and 4 of the ISO/OSI model

The ISO/OSI reference model divides a networking system into seven layers (see Figure 1). Each entity interacts directly with the layer immediately beneath it and provides facilities for the layer above it.

Layer 2, the Data Link Layer, describes the logical organization of data bits transmitted on a particular medium. For example, this layer defines the framing, addressing, and checksumming of Ethernet packets. Media Access Control (MAC) addresses are also resolved at this layer. Ethernet, used on most LANs, is the most common Layer 2 protocol.

Layer 3, the Network Layer, describes how a series of exchanges over various data links can deliver data between any two nodes in a network. For example, Layer 3 defines the addressing and routing structure of the Internet. Logical addressing, message routing, and route determination are executed at this layer. IP is the most common Layer 3 protocol.

Layer 4, the Transport Layer, provides end-to-end, error-free reliable data transfer. It employs sequencing, sliding window, flow control, and error handling as the underlying mechanisms necessary to achieve this service. TCP is the most common Layer 4 protocol, and different applications are bound to TCP by using different port numbers.

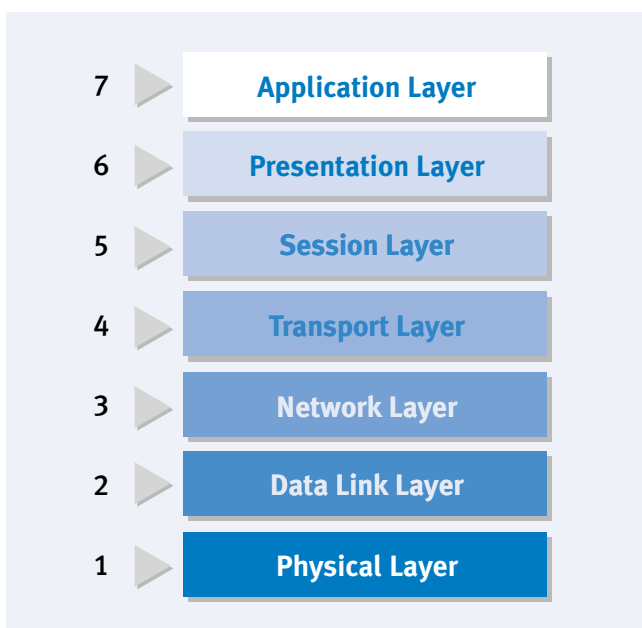


Figure 1. The ISO/OSI reference model

Dell PowerConnect switches enable sophisticated switching capabilities, making commonplace what was previously found only in high-end, proprietary systems.

## Characteristics and gating factors for Layer 2, 3, and 4 switching

The switching decisions made at each layer correspond to the types of activities performed there. Each type of switching has distinct benefits and applications.

### Layer 2 switching reduces congestion

Layer 2 switching is extremely common and is most often implemented in a LAN environment on an Ethernet switch. Ethernet is based on a bus topology, in which every station on a network shares the bandwidth. Layer 2 switching reduces congestion on the network because the switch only forwards packets to the ports for which they are intended, effectively dedicating a full 10 Mbps, 100 Mbps, or 1000 Mbps to each connection. The switch sends out data based on the destination MAC address.

### Layer 3 switching improves performance

Layer 3 switches direct traffic on their ports much like a router, based on the packet's IP address. Although Layer 3 switches usually have many more ports than a typical software-based router, these switches generally accommodate only Ethernet ports and support fewer protocols—often only Routing Information Protocol (RIP) and Open Shortest Path First (OSPF)—to maintain routing tables. Layer 3 switches help network managers segment their networks into IP subnets, providing higher performance, simplified management, and lower per-port costs over traditional software-based routers.

By integrating router functionality into the silicon within a switch, Layer 3 switches provide LAN-based routing at near-switching speeds. Packets that require routing travel across the backplane of the Layer 3 switch faster than a CPU could process the routing information in a traditional, software-based router.

With the LAN router bottleneck removed, switched networks take better advantage of available bandwidth. Desktop users experience high-speed network response, and the network gains stability and reliability. Existing routers remain useful by supporting WAN interfaces and Border Gateway Protocol 4 (BGP4) at the edge of the network, where high port count and ultra-fast packet forwarding are less important.

### Layer 4 switching indicates traffic type

Layer 4 provides TCP port information, which typically indicates the type of traffic a packet carries (such as Hypertext Transfer Protocol (HTTP) or Network File System (NFS)). Layer 4 switching is commonly used in load-balancing devices,

DELL POWERCONNECT SWITCHES				
	Managed switches		Unmanaged switches	
	PowerConnect 3248	PowerConnect 5224	PowerConnect 2124	PowerConnect 2508
<b>Built-in support for Gigabit Ethernet</b>	<ul style="list-style-type: none"> <li>• 48 Fast Ethernet ports</li> <li>• 2 Gigabit Ethernet ports (built-in 1000BaseT or up to two optional fiber modules)</li> </ul>	24 Gigabit Ethernet ports (built-in 1000BaseT, up to four optional fiber modules)	<ul style="list-style-type: none"> <li>• 24 Fast Ethernet ports</li> <li>• 1 Gigabit Ethernet port</li> </ul>	8 Gigabit Ethernet ports
<b>Class of Service</b>	<b>Multilayer</b> <ul style="list-style-type: none"> <li>• Four priority queues per port, adjustable weighted-round-robin scheduling</li> <li>• Layer 2 tag- and port-based priority</li> <li>• Layer 3 IP precedence and IP diffserv code point (DSCP) priority</li> <li>• Layer 4 TCP/UDP-based priority (PowerConnect 3248 only)</li> </ul>		<b>Layer 2</b> <ul style="list-style-type: none"> <li>• Four priority queues per port, weighted-round-robin scheduling</li> <li>• Support of Layer 2 priority tags</li> </ul>	
<b>Advanced switch management</b>	<ul style="list-style-type: none"> <li>• Web-based interface</li> <li>• Industry-standard command-line interface (CLI)</li> </ul>		None	

where all requests for similar traffic types are distributed among servers designated for such requests.

Layer 4 switching examines client requests directed at the Layer 4 switch, multiplexes client requests across any server available to handle those requests, passively measures application health and responsiveness to determine server availability, and provides stateful processing, which monitors the activity.

### Class of Service and Quality of Service

Class of Service (CoS) refers to feature sets, or groups of services, that are assigned to network users based on company policy. In an enterprise network, CoS is used to differentiate high-priority traffic from low-priority traffic.

Quality of Service (QoS) functions, which are implemented in the network devices, are used to guarantee delivery of the packets and determine the packets' priority, based on the network protocols used and the CoS policy. Switches and routers with QoS support enable deployment and activation of CoS policies in the enterprise network.

#### Classification varies by layer

A data packet is defined, or classified, using a traffic descriptor or tag to categorize a packet within a specific group. The packet is then accessible for QoS handling on the network.

Layer 2, 3, and 4 switching differs from Layer 2, 3, and 4 classification: classification affects only prioritization; the switching decision is made independently. Classifications can be made on any combination of Layers 2, 3, and 4, giving IT managers several mechanisms to optimize the network traffic.

**Layer 2.** Source and destination MAC addresses can determine Layer 2 classification, or classification can be based on tag information inserted as a small shim tag header between the Layer 2 and Layer 3 headers. Some Layer 2 protocols, such as Asynchronous Transfer Mode (ATM), carry tag information as part of the Layer 2 header.

**Layer 3.** Layer 3 classification uses the IP Type of Service (TOS) field or the IPv6 flow label field. The 8-bit IP TOS field assigns the three most significant bits to IP precedence, which indicates priority. IPv6 provides a 24-bit flow label field and a 4-bit priority field for this purpose.

**Layer 4.** Layer 4 systems combine the benefits of Layer 4 software on a high-speed Layer 2 switching platform to enable intelligent switching decisions and the prioritization of frames based on TCP/UDP (User Datagram Protocol) and IP source/destination addresses. The information establishes policy controls for traffic management.

Gigabit Ethernet networks and delay-sensitive, high-bandwidth applications such as voice over IP (VoIP), video, and iSCSI will benefit from Layer 4 classification, which allows session-by-session QoS and end-to-end session management.

Most IT managers today implement Layer 2 switching in their campus LAN with Layer 4 classification. Layer 2 Gigabit Ethernet switching allows very high bandwidth links, and Layer 4 classification provides high QoS for delay-sensitive applications across those links.

#### Diffserv aggregates behaviors

Another approach to providing QoS in networks is differentiated services (diffserv), under development by the Internet Engineering Task Force (IETF). Diffserv depends upon a small,

Quality of Service  
functions guarantee  
delivery of the  
packets and actually  
determine which packets  
have priority.

well-designed set of building blocks from which a variety of aggregate behaviors may be built.

Operating at Layer 3, diffserv uses six bits of the IP TOS byte, called the diffserv field (DS field), to specify a small number of per-hop behaviors at each network node. Diffserv aggregates flows in the network so that core routers only need to distinguish a comparably small number of aggregated flows, even if those flows contain thousands or millions of individual flows. Diffserv provides cross-platform, multivendor, and cross-network interoperability for CoS features. Diffserv is supported by Dell PowerConnect switches.

### Switching, classification, and Gigabit Ethernet in deployments

Deployment of Gigabit Ethernet networks over Category-5 cabling using Dell PowerConnect switches can be based on relatively simple designs. Figure 2 illustrates a common small-office architecture using Dell PowerConnect unmanaged switches.

This simple design accommodates up to 72 clients running at 100 Mbps full duplex and a pair of servers with dual GbE connections for load balancing and failover. In this model, IT managers can deploy tag-based Layer 2 priority by configuring the server network interface cards (NICs) to add a priority tag to the data packets. For example, if one server were a VoIP server, an IT manager could ensure that all traffic from that server is classified as high priority, resulting in a low-latency path through the network.

Figure 3 illustrates a deployment model for larger installations. This design allows full redundancy between the network equipment and the servers, provides significant bandwidth




Figure 2. Typical design for a small or a remote office

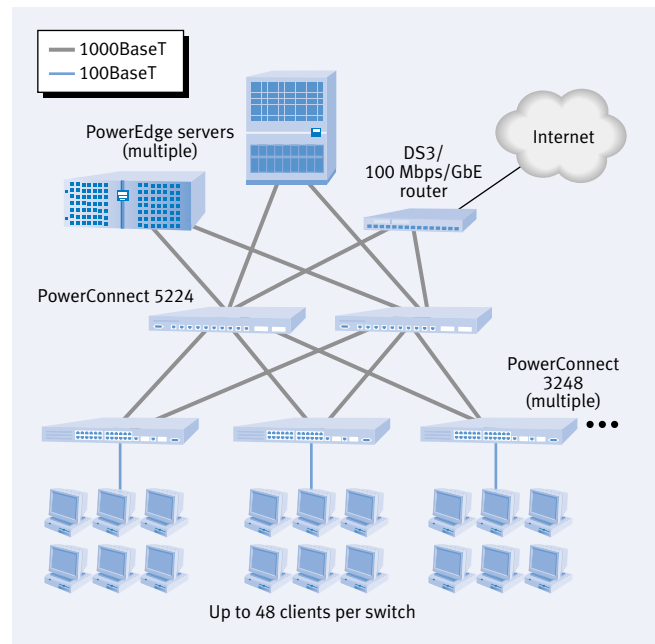


Figure 3. Typical redundant design for medium-sized office

between the servers and the network, and delivers 100 Mbps to 400 client PCs.

Administrators can add client PCs by purchasing another PowerConnect 3248 and running additional, redundant GbE connections back to the PowerConnect 5224 switches. Any single switch and NIC can fail without bringing down the entire network. The switches can accommodate installation of fiber Gigabit Interface Converters (GBICs) to provide additional range if the network spans to another building. New buildings and server farms can be added without changing the design.

### Sophisticated switching for simplified Gigabit Ethernet

The sophisticated switching capabilities in Dell PowerConnect switches simplify migration to Gigabit Ethernet. The devices make Gigabit Ethernet deployments cost-effective, further driving Gigabit Ethernet technology and advanced CoS features into the mainstream. 

Richard Martin ([rmartin@broadcom.com](mailto:rmartin@broadcom.com)) is director of marketing at Broadcom Corporation, an Irvine, California, based provider of integrated circuits enabling broadband communications.

#### FOR MORE INFORMATION

For more information, service, and support, visit <http://www.dell.com>