OFFICE OF THE MANAGER
NATIONAL COMMUNICATIONS SYSTEM
TECHNICAL NOTES

APRIL 2000 | *TECHNOLOGY AND PROGRAMS DIVISION* | VOLUME 7, NUMBER 1

Differentiated Services—One Solution for Priority Over the Internet

by Ray Young

The Internet is becoming increasingly important for national security and emergency preparedness (NS/EP) communications as the circuit-switched telecommunication infrastructure is integrated with emerging Internet Protocol (IP) based networks. It is now more likely that mission-critical communications could be delayed as a result of congested IP networks.

One solution in development that could help NS/EP communications traverse congested IP networks is Differentiated Services (DS). This technical note explores DS and projects how this suite of protocols could be used to ensure NS/EP communications reach their destination.

TYPE OF SERVICE FIELD

Currently the IP version 4 (IPv4) header includes a field called "Type of Service" (TOS) intended to indicate the quality of service (QOS) Internet packets should receive. QOS, a measurement for Internet traffic, quantifies delay, throughput, and reliability that an Internet packet or packet stream is receiving. Different kinds of Internet traffic require differing QOS levels. Some types of data (e.g., electronic mail) can tolerate a lower QOS because packets can be buffered and, if lost, resent with little impact on the user. For both voice and video, which are real-time services, even a small delay in communications produces a noticeable degradation of service for users.

As shown in Figure 1, the TOS field[1] is defined as an eight-bit field. It provides two levels of service, normal and high, for each of the three QOS measurements (delay, throughput, and reliability). The TOS field designates the zeroth through second bits for internal use within networks. Initially, the sixth and seventh bits were set aside for future use and set to zero. With only three bits available to prioritize packet traffic, TOS has proved itself limited in accommodating the growing number and types of services being added to the Internet. Later, the Internet Engineering Task Force (IETF) defined the sixth bit as cost.[2] When this bit is set to 1, the network tries to minimize the monetary cost of routing the packet.

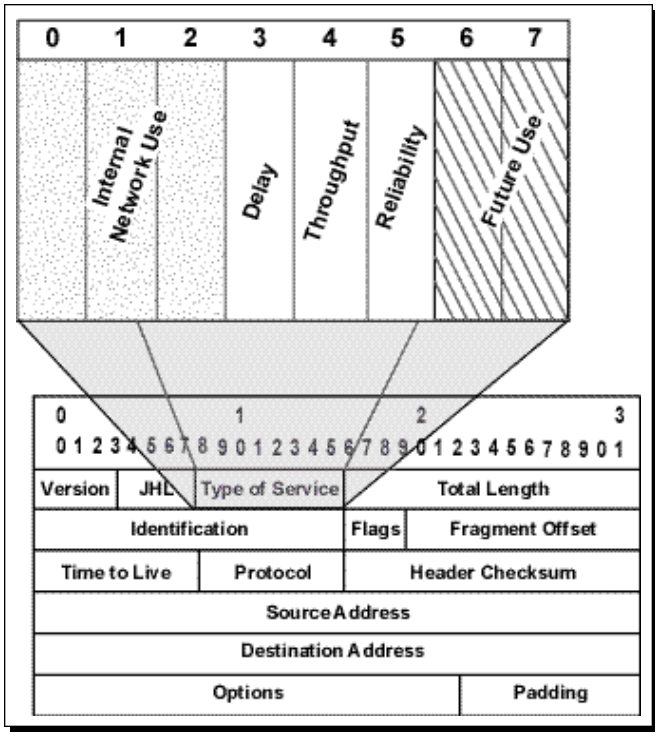

Figure 1. The Type of Service Field

The DS field[3] supersedes the TOS field as defined for IPv4. DS can also be applied to the Traffic Class octet within the emerging IPv6. Like the TOS field described in Request for Comments (RFC) 791,[1] the DS field is eight bits long and reserves the sixth and seventh bits for future use. (See Figure 2.) However, with DS, the zeroth through fifth bits are available to define services.

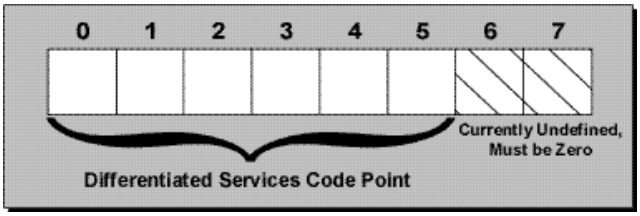

Figure 2. Structure of Differentiated Services Field

Differentiated Service Code Points

A code point is a specific binary value that has special meaning to the network based on rules and policies. For DS, a code point tells a network node (router) how to prioritize, treat, and route a particular packet within the general traffic the node is receiving.

As stated earlier, the DS field is eight bits, six of which are used. Therefore, there are 64 possible code points (2^6=64). The code point space is divided into three pools.[3] (See Table 1.) Thirty-two code points are for regular use, while the remaining 32 code points are divided into 2 experimentation pools. Sixteen code points are for experimentation within networks. The remaining 16 code points are also set aside for experimentation, but their designation changes to regular use when the first pool of 32 regular codes is exhausted. Each code point is mapped to the particular forwarding treatment that nodes should provide for a packet at each hop along the packet's path (per hop behavior).

Pool	Code Point Space
Regular Use	xxxxx0
Experimental	xxxx11
Experimental (but may be reassigned to Regular Use)	xxxx01

Table 1. PHB Pools

Per Hop Behaviors

Packets reach their destination by being forwarded from one node to another. As they are routed, packets that require similar service are grouped by their per hop behavior (PHB). This ability to group packets provides network engineers with the tools (rules and policies) to develop differentiated services. Currently, two types of PHBs are under development, expedited forwarding (EF)[4] and assured forwarding (AF).[5]

Expedited Forwarding

The EF approach (see Figure 3) to forwarding packets divides packets into two classes. Most packets are forwarded using the best effort available under current network conditions. A small subset of network traffic is given a special EF code

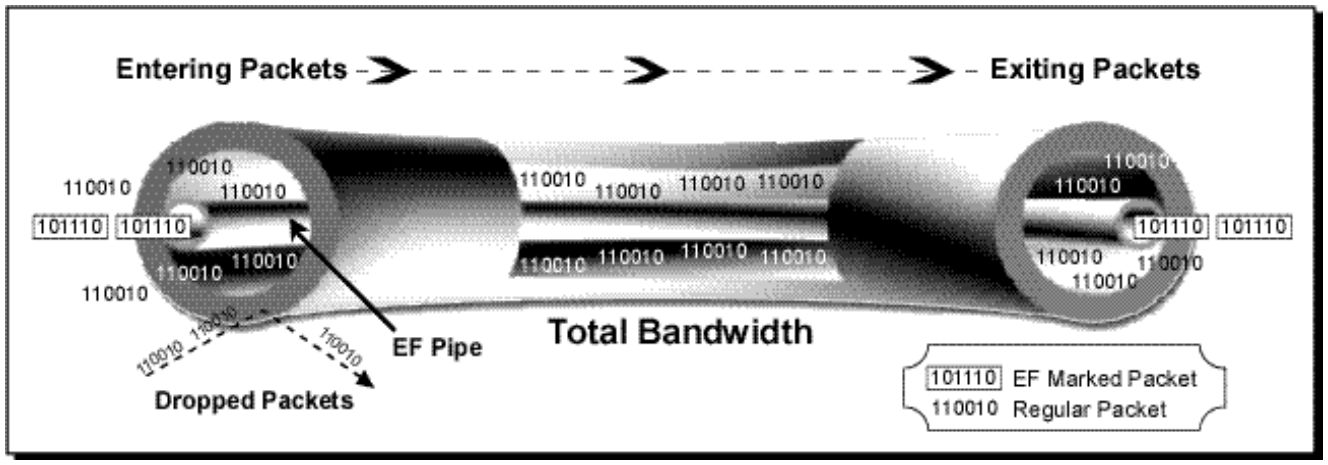

Figure 3. EF Description

point designation. This subset of traffic provides the highest QOS possible. EF could be viewed as a virtual leased-line connection, because EF defines a particular level of assured bandwidth, low loss, low latency, or low jitter end-to-end service. Other packets on the network are preempted to make room for these packets when congestion arises.

While at the node, the EF-designated packets can be queued or delayed to fit the service requirements of the network. Therefore, EF compensates for the delay by receiving packets at a higher rate within a network node than is needed by the next or destination nodes. Code point 101110 is recommended for the EF PHB.[4] EF PHB can coexist on a network that uses other PHB schemes. EF PHB focuses only on the service thresholds that EF PHB-marked packets receive. If non-EF PHB-marked packets receive treatment from another PHB with the remaining bandwidth, and it does not impact the EF marked packets, then the protocol is not violated.

ASSURED FORWARDING

Assured forwarding groups packets into one of four classes. Each class has three drop precedence levels, low, medium, and high. Each class or type of traffic is independent of the other classes and can have its own unique drop precedence. Table 2 describes the code points defined for AF PHB.

The best way to visualize AF is to think of the seating in an airplane. Generally, there are first class, business class, coach, and stand-by tickets.

Better Service	Class 1	Class 2	Class 3	Class 4
Low Drop Precedence	001010	010010	011010	100010
Medium Drop Precedence	001100	010100	011100	100100
High Drop Precedence	001110	010110	011110	100110
Default Drop Precedence: 000000				

Table 2. AF PHB Code Points

This is analogous to dividing the classes of communications service by type (e.g., video, telephony, or data). Within each airline class of service, “flyers” can have “Gold” cards, be preferred passengers, or be general customers, which determines who gets seated when the flight is oversold. The categories within each class of service are analogous to drop precedence, which determines treatment if there is network congestion.

DIFFERENTIATED SERVICES OPERATION

For DS to work, the network administrator must define PHBs. Next, the users of the network resources must negotiate service level agreements (SLA) with the network service provider. An SLA is a contract between the user and service provider that specifies the level of service (i.e., bandwidth, loss rate, and delays), as well as the type of treatment and routing the Internet traffic will receive. It also includes times of service availability and describes how the service will be measured and billed.

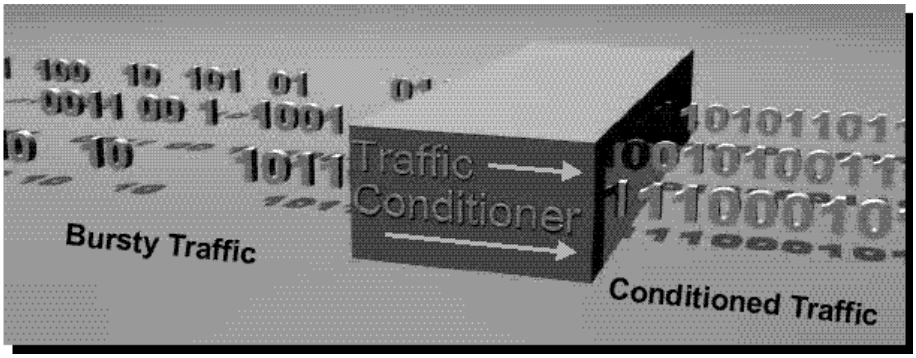

Figure 4. Packet Conditioning

For DS packets, applying an SLA to a PHB is called classifying. DS typically is implemented at the boundary nodes of a network. Once the packets arrive at a boundary node and are classified, the node conditions the traffic as specified by the rules and policies contained in the SLA. The traffic conditioner is a set of functional elements within a node that meters, marks, shapes, and polices the incoming traffic. Metering measures agreed-to network properties over a period of time. Marking adds the code points specified by the SLA and PHB. Shaping delays packets within a packet stream to conform with a specified traffic profile. Traffic conditioning usually occurs when there is a burst of network traffic beyond the norms described in the SLA. Figure 4 illustrates the concept of traffic conditioning. Policing is the process of deciding which packets should be dropped based on the traffic profile. This traffic profile, a predefined set of rules, governs the Internet traffic and includes maximum burst size and data rate. Figure 5 illustrates DS operation, including functional elements.[8]

DS Security Considerations

Because a DS code point can be altered or an improper code point can be inserted into a network, DS is vulnerable to denial of service

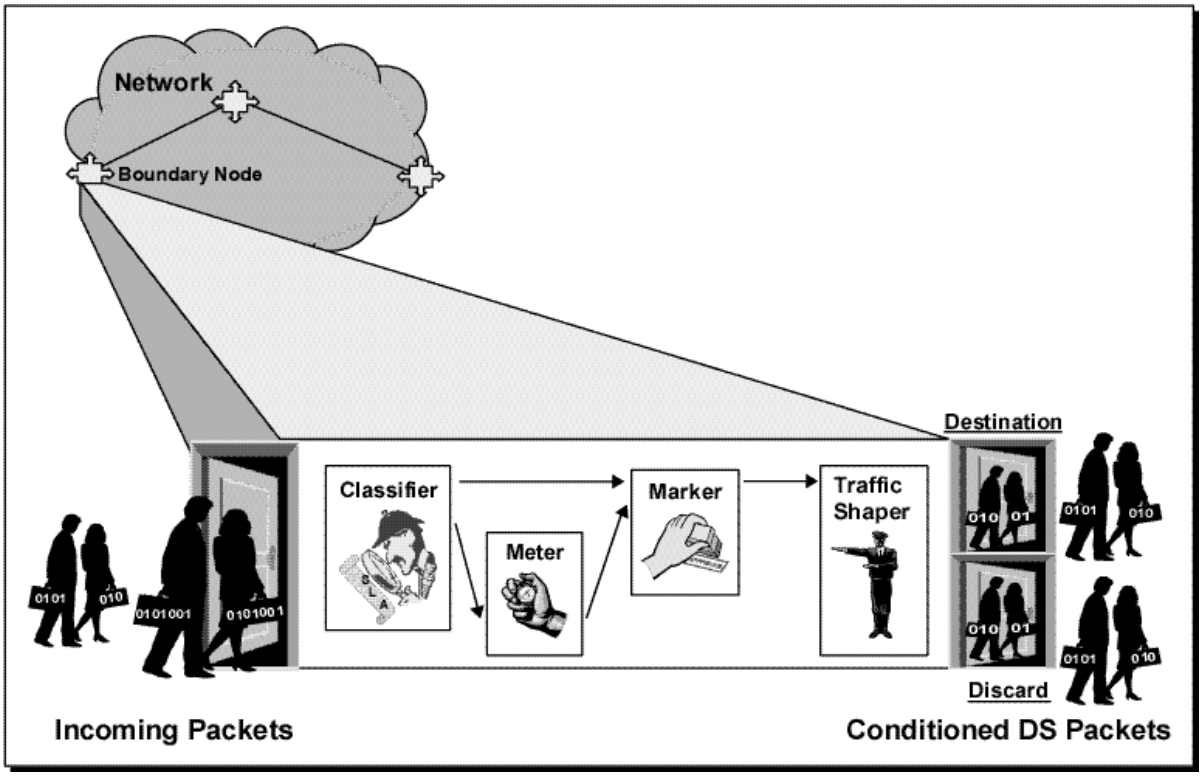

Figure 5. DS Operation

attacks. Because DS provides a better level of service to one class of packets over another, when the network is congested an adversary could deny service to authorized network traffic.[3] Also, because different kinds of traffic are marked with different code points, an adversary might be able to acquire information about a packet or network by observing traffic patterns.

IPSEC, the IP protocol that provides encryption and end-to-end security for packets, does not encrypt the DS field of a packet. However, when IPSEC is used to create a secure tunnel between two endpoints and reasonably strong encryption is used, an integrity check of the packet can be used to assure that the packet was not modified. The most important aspect of securing DS is authenticating who is sending traffic to edge routers within a network before the traffic is introduced into the network.

DIFFERENTIATED SERVICE–A PIECE OF THE PUZZLE

DS is one tool among many being developed to overcome network congestion. No single tool is likely to provide the optimal solution. Therefore, the various tools in development will need to work together to provide the best service for priority communications. Two services that complement DS are Resource Reservation Setup Protocol (RSVP) and Multiprotocol Label Switching (MPLS).

RESOURCE RESERVATION SETUP PROTOCOL

The RSVP is a very complex protocol that mimics circuit-switched communications. Before communication occurs, the network sends out a request for bandwidth to all the nodes in the communications path. The request includes all the parameters required, including QOS level and load. After all the required resources are reserved, the communication can proceed.

RSVP attempts to convert the best-effort signaling of the IP world into circuit switching. This technique is probably best suited for establishing connections from the network edges to the parties at both ends of the communication. However, its high overhead is not well suited for backbone connections. RSVP is the best method in development

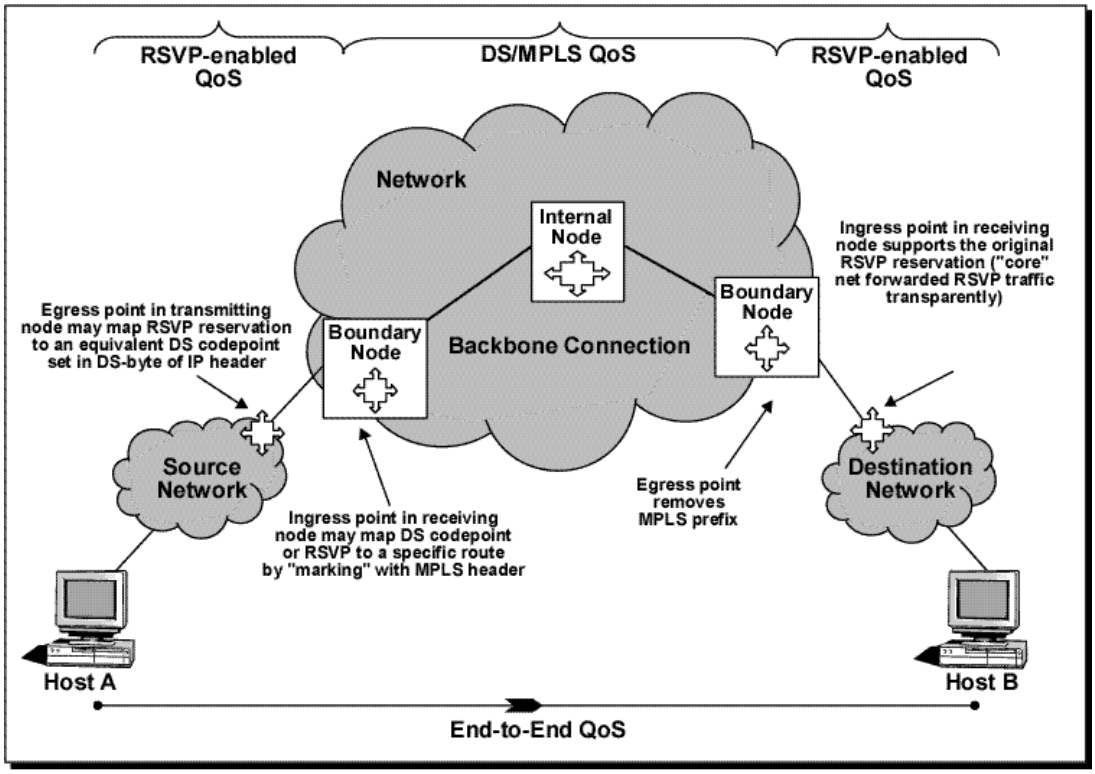

Figure 6. DS Integrated with RSVP and MPLS

for capturing the QOS requirements of a packet stream. These requirements could then be translated into DS PHBs.

MULTIPROTOCOL LABEL SWITCHING

MPLS and RSVP are similar in that they both try to emulate circuit-switched networks. However, unlike RSVP, MPLS works by predefining paths from one endpoint to another. When a packet arrives, its endpoint is noted and a label is added to the packet. The label is called a label switched path (LSP). As the packet is forwarded, each node reads the LSP to determine where to forward the packet and updates the LSP for the next node.[7]

MPLS, once fully standardized, should provide an efficient means to emulate circuit-switched networks over backbone connections within a network. MPLS can capitalize on DS's use of classes to determine the best path to route a particular class of packets or packet stream.

PUTTING THE PIECES TOGETHER

DS, RSVP, and MPLS could be pieced together to build a priority service. Figure 6 describes one possible implementation of the protocols. Data are sent from Host A through its network. Using RSVP, Host A's network requests provision of a certain prenegotiated QOS. The requirements for the QOS are contained within RSVP and are mapped to a DS class. Between the boundary nodes, MPLS uses the DS class mark to select a route for the packet stream to follow. RSVP is used again to deliver the packets to the final destination, Host B.

NS/EP IMPLICATIONS

The National Communications System (NCS) is studying the impact of emerging Internet technologies on NS/EP communications. The NCS has formed an Internet Program Office to ensure that emerging Internet technologies can be used for NS/EP communications as the circuit-switched network merges with packet-switched networks.

DS is one of the emerging priority handling techniques the Office is investigating. As demonstrated in the previous section, several technologies must work in concert for the network to provide end-to-end priority treatment for NS/EP communications. DS will be an important building block in the development of future services.

REFERENCES:

1. Postel, J., editor, "Internet Protocol," RFC 791, September 1981.
2. Almquist, P. "Type of Service in the Internet Protocol Suite," RFC 1349, July 1992.
3. Nichols, K.; Blake, S.; Baker, F.; and Black, D. "Definition of the Differentiated Services in IPv4 and IPv6 Headers," RFC 2474, December 1998.
4. Jacobson, V.; Nichols, K.; and Poduri, K. "An Expedited Forwarding PHB," RFC 2598, June 1999.
5. Heinanen, J.; Baker, F.; Weiss, W.; and Wroclawski, J. "Assured Forwarding PHB Group," RFC 2597, June 1999.
6. Bernet, Y. "A Framework for Differentiated Services," Internet Draft, February 1999. (Drafts are subject to change or withdrawal.)
7. Awduche, D.; Malcolm, J.; Agogbua, J.; O'Dell, M.; and McManus, J. "Requirements for Traffic Engineering Over MPLS," RFC 2702, September 1999.
8. Blake, S.; Black, D.; Carlson, M.; Davies, E.; Wang, Z.; and Weiss, W. "An Architecture for Differentiated Services," RFC 2475, December 1998.

Requests for Comments can be viewed on line from the Internet Engineering Task Force (IETF) website, http://www.ietf.org/

For more information on Differentiated Services, contact:

Ray Young
National Communications System
Technology and Programs Division (N2)
701 South Court House Road
Arlington, VA 22204-2198
(703) 607-6200