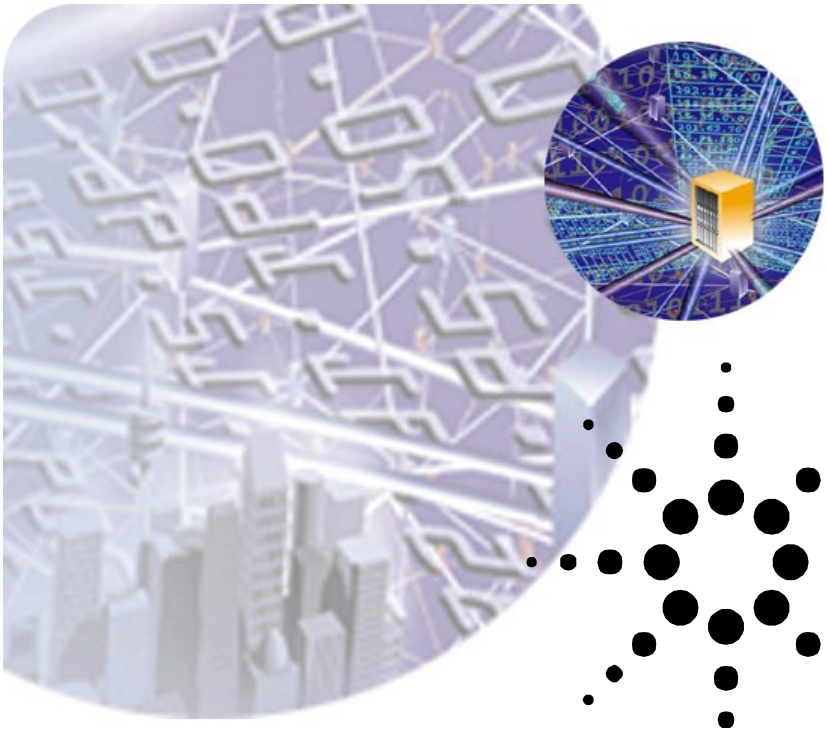

# IP Quality of Service

Agilent Technologies RouterTester
Application Note

## Introduction

IP-based networks will become the next century's public network infrastructure for time-sensitive services such as voice and multimedia, and value-added applications such as financial transactions and just-in-time inventory tracking. However, the successful delivery of these services is dependent upon the ability to provide reliable, predictable, and class-aware IP transport. With mission-critical traffic currently being carried across a network infrastructure that is best effort by design, IP Quality of Service (QoS) has become a very hot topic.

In order to attract and retain customers in today's highly competitive environment, service providers must offer new IP-based services with an associated quality of service. The ability to differentiate and guarantee service offerings enables service providers to charge customers according to the quality of service that is delivered. By offering unique, value-added and customized services, service providers are better able to differentiate themselves from competitors and leverage their networks to increase revenues.

Agilent Technologies
Innovating the HP Way

To deliver IP QoS, the Internet Engineering Task Force (IETF) has a number of initiatives that augment the IP protocol to provide reliable, classful delivery of IP traffic. Network equipment manufacturers are also implementing architectural improvements and support for new technologies within core network routers that enable these devices to deliver quality of service guarantees.

This paper reviews the emerging technologies for implementing IP QoS as well as the critical test methodologies for verifying the QoS capabilities of core network routers.

## QoS Defined

Within the context of this paper, QoS refers to the successful delivery of an agreed upon level, or class of service. A class of service is characterized by a set of performance parameters including:

- delay (otherwise known as “latency”), which refers to the time interval it takes a packet to be forwarded between two reference points;
- delay variation (otherwise known as “jitter”), which refers to the variation in transit time for all packets in a stream taking the same route;
- throughput; which refers to the rate at which packets go through or transit a network or network device, expressed as an average or peak rate;
- and packet loss, which is the maximum rate at which packets are discarded during transfer through a network.

These performance metrics are used to differentiate the QoS level provided by a service. Implicit in the concept of QoS is the ability to differentiate traffic into distinct and distinguishable service classes, which can be treated individually and predictably by network devices.

## The Problem with IP

The problem with the legacy IP infrastructure is that is was not designed to deliver traffic with different service requirements, nor predictable service. The IP infrastructure is based on a "best-effort" model where all network traffic is created equal, and service is based on availability rather than guarantees.

In the absence of QoS, service providers have commonly opted to simply over-provision bandwidth. In the absence of congestion, traffic can be forwarded through a network with minimal latency, jitter, and loss. Service providers have traditionally attempted to avoid congestion by provisioning more bandwidth than they expect will be needed. This solution is acceptable for transporting voice over the public switched telephony network. However, with Internet traffic doubling every 4 months, it is practically impossible to precisely match data traffic volumes to bandwidth provisioning on an IP network. Because an uncongested network wastes a certain fraction of its throughput capacity, throwing bandwidth at the QoS problem is obviously not a viable long-term solution.

## Bringing QoS to IP

The Differentiated Services (DiffServ) Working Group has been established to define new IP QoS mechanisms that can be formalized in a set of industry standards. The following section of this paper reviews the IETF initiatives for delivering IP-QoS, including the mechanisms for differentiating traffic into distinct service classes, as well as the architectures and technologies that enable core network routers to recognize and manage different traffic classes efficiently.

## Differentiating the Traffic

### Class of Service

Before QoS can be delivered, mechanisms that can differentiate network traffic into different classes of service must be provided. For example, the traffic handing requirements of mission-critical and real-time applications such as voice over IP (VoIP) differ from fax and e-mail applications, which are less sensitive to bandwidth and delay issues. As such, traffic with different service handling requirements must be sorted into different categories or service classes that can be treated individually. This concept of traffic classification is referred to as *Class of Service* (CoS).

In order to define and deliver CoS, user and application requirements must be known to the network, and in turn, the network must be capable of providing the mechanisms that can deliver the service levels approximated by these requirements. It is important to note that CoS is just a small part of the larger QoS picture. QoS encompasses CoS, as well as the all the mechanisms required to recognize and manage CoS.

Because IP is connectionless, and without traffic contract concepts, marking IP packets with CoS and traffic handling information is a challenging endeavor. Evolving techniques that address this challenge include manipulating the TOS field of the IP header, or encapsulating the IP packet. These techniques are utilized in the Differentiated Services (DiffServ) and MultiProtocol Label Switching (MPLS) initiatives currently under revision with the IETF. However, before reviewing these emerging QoS initiatives, it is useful to explain the limitations of the existing technologies that have been employed to deliver QoS.

## Existing IP QoS Delivery

### Asynchronous Transfer Mode (ATM)

Asynchronous Transfer Mode (ATM) is a connection-oriented technology that relies on the pre-provisioning of bandwidth to deliver QoS. ATM uses a signaling mechanism to set up and establish virtual circuits with specific QoS parameters. Network nodes respond to the signaled request by reserving the resources necessary for the ATM connection. Unfortunately, ATM is exclusively a circuit-switched technology, which doesn't map well to the modern packet-switched Internet.

### Integrated Services (IntServ)

The Integrated Services (IntServ) was introduced to apply the QoS concepts employed in ATM to the connectionless IP world. Like ATM, IntServ relies on the reservation and control of network resources in order to deliver QoS. IntServ defines different service levels that are characterized by quantifiable QoS parameters, such as the amount of required bandwidth, and allowable latency, jitter and loss. IntServ relies on the Resource Reservation Protocol (RSVP) to signal the QoS parameters for a specific traffic flow through the network. Network nodes respond to the signaled request by reserving the requested resources, and keeping state-information for each traffic flow.

This pre-provisioning of network resources burdens the IntServ QoS model with many scaling and implementation challenges. First, the signaling required to reserve service parameters at each network node takes a considerable amount of time. Thus, there is a certain amount of delay overhead associated with each QoS negotiation.

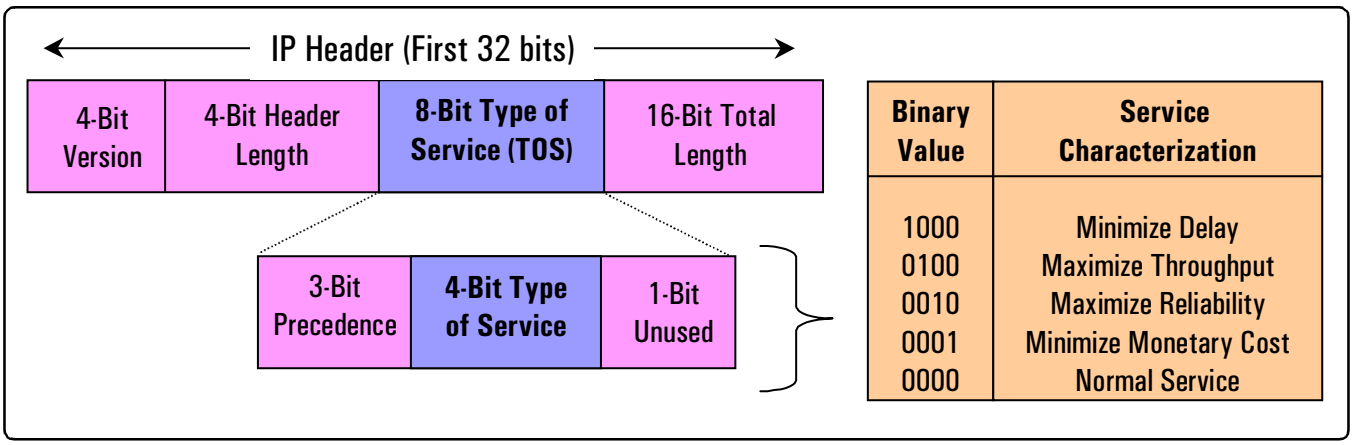


Figure 1: TOS service characterizations provide no quantifiable or relative parameters.

Next, the IntServ model requires that each switch/router along the signaled path maintain state information for potentially thousands of traffic flows, which consumes a substantial amount of processing power, and limits network scalability. Using RSVP in the network core inhibits a router's ability to maximize speed and minimize computational resources because it must store and look up pathway information for each packet for the duration of a traffic flow. Lastly, IntServ requires the traffic source to know exactly how much bandwidth to reserve. This can be an almost impossible feat when the dealing with bursty traffic over an IP network, so bandwidth is either wasted through over-provisioning, or QoS is jeopardized by bandwidth under-provisioning.

## Type of Service (TOS)

The Type of Service (TOS) field in the IP header was designed to deliver QoS by tagging IP packets with different service characterizations. These service characterizations describe the how network nodes reading the IP header should treat the packet. The most recent TOS specification, RFC 1349, defines the TOS field as a set of bits to be considered collectively. The TOS values, shown in Figure 1, denote how the network should treat the packet with respect to tradeoffs between throughput, delay, reliability, and cost. The first 3 precedence bits of the TOS octet are intended to denote the importance or priority of the packet. A network router can use the TOS field when choosing a path over which to forward the packet, and when making queuing decisions.

Although the TOS field has been a part of the IP specification since its implementation, it has been little used in the past. This lack of use can be attributed to poorly defined service characterizations. The ambiguous nature of these service characterizations provided no quantifiable or even relative service parameters, thus making it extremely difficult to differentiate traffic with a consistent level of service quality. The lack of TOS implementation, coupled with its inability to allow an application to quantify the level of service that it desires, makes TOS an inappropriate mechanism for delivering service guarantees.

## New IP QoS Delivery

### Differentiated Services (DiffServ)

The Differentiated Services (DiffServ) QoS model addresses the scalability and implementation issues that burden previous QoS models. DiffServ defines a more scalable and flexible way to apply IP QoS in the network core. While IntServ and ATM are characterized by end-to-end signaling and stateful forwarding decisions, DiffServ eliminates signaling, handles flow aggregates, and employs standard markings in each packet that routers can quickly examine without reference to processing-intensive session lookups.

### Traffic Classification

DiffServ marks each packet with specific service-level requirements, thus enabling routing decisions to be made on a per-packet rather than per-session basis. This process makes more efficient use of bandwidth than previous QoS mechanisms because it eliminates the need to reserve bandwidth without knowing exactly how much is needed.

DiffServ service-level markings take place in the type of service (TOS) field of the IPv4 header. Under DiffServ, the eight-bit TOS field has been renamed the DS (differentiated services) field, which embodies a six-bit DS code point (DSCP) and two currently unused (CU) bits. The DSCP carries information about the service requirements, or relative priority of the IP packet. Using the six bits, DiffServ is capable of defining 64 service levels, enabling a higher degree of service granularity than ever achievable before. The DSCP corresponds to a Per-Hop-Behavior (PHB) that defines the relative priority and QoS parameters that a packet should be given by each node in a DS domain. A Differentiated Services Domain (DS domain) is defined as "a contiguous portion of the Internet over which a consistent set of differentiated services policies are administered in a coordinated fashion" (RFC 2474). Specifications of DiffServ policies and their administration are determined by service level agreements and network administrators, and are outside the scope of this paper. Because all network nodes within a DS domain apply PHBs in a consistent manner, DiffServ PHB classifications are much easier to implement than the ambiguous TOS values.

### Per Hop Behaviors (PHBs)

Standardized PHBs enable service providers to design services from a well-known set of packet forwarding treatments that can be implemented in the equipment of many vendors. The DiffServ Working Group has defined several standard PHBs ranging from best effort to guaranteed delivery, as discussed below.

### Best Effort (RFC 2474)

Best effort is defined as the default class of service, and is mapped with the DSCP 000000. Traffic with this DiffServ mapping has no specific traffic contract, and thus receives whatever bandwidth remains after traffic with other PHBs has been processed. The DSCP 000000 is exactly the same as the old TOS 'normal service' characterization, so it achieves backwards compatibility with the previous usage of this field.

### Expedited Forwarding (RFC 2598)

In contrast to the default best effort delivery, the expedited forwarding (EF) PHB is used to establish a guaranteed bandwidth service for an IP packet traversing DS domains. Often referred to as Premium service, EF delivers a guaranteed amount of

bandwidth, while minimizing delay and packet loss to traffic marked with the DSCP 101110. Operating as a 'virtual leased line' service over the shared IP network, EF is able to leverage the benefits of a traditional leased line service while minimizing costs. Network administrators specify the maximum allowable rate and burst size for EF PHB traffic aggregates traversing their DS domain.

## Assured Forwarding (RFC 2597)

The Assured Forwarding (AF) PHB defines a class of service with a delivery guarantee that is better than best effort, but inferior to EF. Assured forwarding delivers traffic at a guaranteed sustained rate, with bursts up to a maximum as specified by network administrators. AF currently defines four distinct traffic classes, each coupled with three possible drop probabilities (low, medium, high). Within DS network node, each AF class is allocated a certain amount of bandwidth and buffer space. The drop precedence of a packet determines the relative importance of the packet within the AF class. Under congestion, the forwarding guarantee of an IP packet within a DS-capable router is determined by:

- the amount of bandwidth and buffering allocated to the AF class
- the existing load of the AF Class
- the drop precedence of the packet

The recommended DSCPs for the AF PHBs are shown in Table 1. The AF class is indicated by the three most significant bits of the DSCP, and the drop precedence by the three least significant bits.

**Table 1: Recommended AF PHBs**

| | 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 |
| Olympic Service Model | Bronze | Silver | Gold | Network Control |

*Source: RFC 2597*

RFC 2597 describes how the AF PHB can be used to implement the Olympic service model, consisting of a bronze, silver and gold service class. Packets are assigned to these classes so that packets in the gold class are forwarded ahead of packets in the silver class, and packets in the silver class are forwarded ahead of those in the bronze service class. It is suggested that the bronze, gold and silver service classes in the network be mapped to AF classes 1, 2 and 3. The low, medium, and high drop precedence values may be mapped to AF drop precedence levels 1, 2, or 3. Using this model, delay and loss-sensitive voice traffic could be assigned to the gold traffic class with low drop precedence to ensure timely forwarding with minimal packet loss.

Because DiffServ works at Layer 3, its IP level marking has the advantage that the requested QoS can be used end-to-end. The DSCP maps onto the existing TOS value, and DiffServ QoS specifications can be recognized by any network device that reads the IP header and DS byte. While DiffServ's flexibility enables users to classify traffic at the source, or network managers to apply classification at precise points in the LAN, traffic classification is most likely to occur at the LAN/WAN edge router to address the critical need for QoS in the network core.

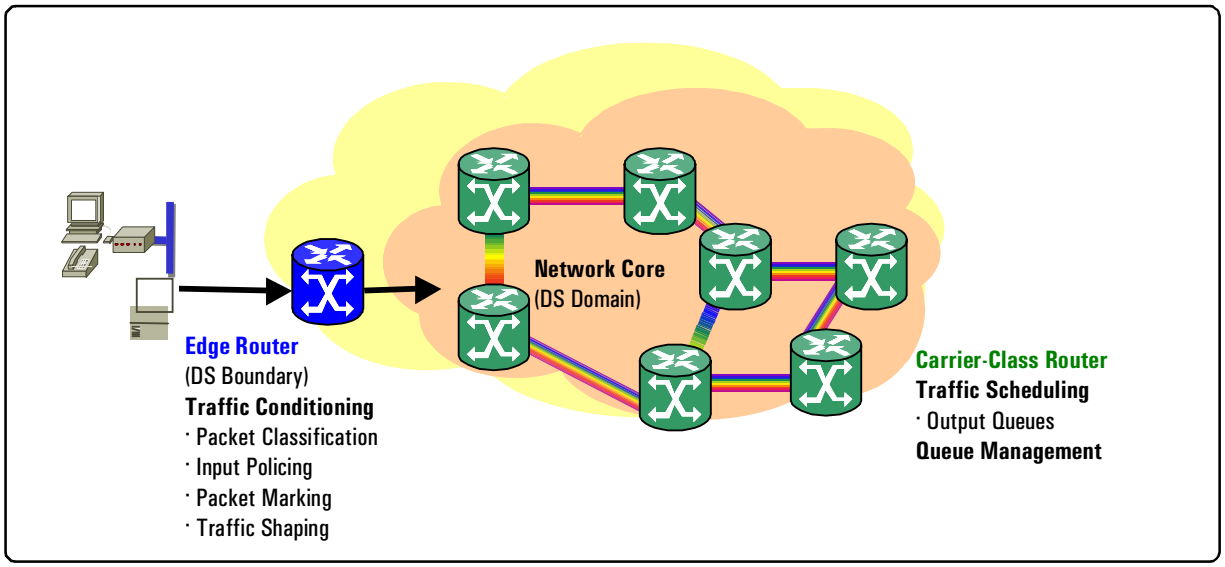


Figure 2: Traffic Management within Network Routers. The edge router performs processing-intensive tasks such as packet classification. The core router performs high-bandwidth traffic management tasks such as queue management.

## Traffic Management

DiffServ enables all packets with the same PHB to be grouped into the same flow for efficient transport across the Internet. DiffServ's per-hop QoS model allows routers to easily manage different traffic classes by assigning traffic flows to standard service levels. Because packets with similar priorities can be aggregated into a limited and manageable set of class flows, DiffServ easily scales to support larger environments. DiffServ also makes efficient use of bandwidth because QoS is implemented on a per-hop basis, eliminating the need to reserve bandwidth without knowing exactly how much is needed.

While DiffServ's QoS per-hop model simplifies the amount of the work that core routers must do, there is a number of additional mechanisms that routers must support in order to manage different traffic classes and deliver QoS. As illustrated in Figure 2, the extent of a router's traffic managing responsibilities depends on where within the network the router is located. We will now review the traffic conditioning functions that take place within DS boundary routers, as well as the scheduling and congestion management mechanisms that enable core network routers to effectively manage and deliver different classes of traffic.

### Traffic Conditioning

Traffic conditioning plays an integral role in managing different traffic classes. In the DiffServ architecture, traffic conditioning is used to:

- enforce service agreements between DS domains
- classify traffic to receive a differentiated service within a domain by marking packets with the appropriate codepoint in the DS field
- police and modify the traffic distribution characteristics where necessary

Traffic conditioning is typically deployed in DS boundary routers called *traffic conditioners*. Because these edge routers perform all of the process-intensive multi-field classification, policing and marking of DiffServ packets, core routers are able to simplify their processes significantly.

As mentioned earlier, packets are classified with a PHB in accordance with some service specification determined by the network administrator. Traffic conditioners use input policers to measure input traffic rates and determine

**Table 2: Implementation of DiffServ PHBs**

| PHB | Input Policing Schedule | Output Management | Congestion |
|---|---|---|---|
| Best Effort | None | Lowest priority queuing | Most likely to be dropped |
| Assured Forwarding | Police on sustained and burst rates<br>Burst: packets are dropped<br>Out-of-contract: Packets are dropped | In-Contract: Better-priority queuing<br>Burst: Same as in-contract<br>Out-of-contract: Treated as best effort | In-contract: Won't be dropped<br>Burst: May be dropped<br>Out-of-contract: Same as best-effort |
| Expedited Forwarding | Police on sustained rate<br>Out-of-contract: Packets are dropped | Highest priority queuing (traffic is also shaped in edge routers) | Won't be dropped |

*Source: Kaufman, 1999*

whether the traffic entering the DS domain complies with the service specification, or pre-negotiated traffic-shaping policies used to control the volume and transmission rate of traffic entering the network. A packet's physical port, IP source address or destination address, or TCP/UDP port ID can be used to verify a particular PHB for the packet.

Burst and out-of contract packets are either immediately dropped, or marked with a different PHB so they can be dropped later if congestion occurs. The packet's treatment within the router is determined by its PHB, as indicated in Table 2.

Traffic shapers are also implemented in hardware to control the volume and transmission rate of traffic exiting the edge router and entering the network core. By scheduling traffic flows for forwarding, traffic shapers attempt to smooth out bursty traffic streams so that they to fall within the parameters of a service contract.

## Scheduling

Interior nodes within the DS domain are not required to perform traffic conditioning. Core network routers simply modify their behavior according to a packet's PHB. There is a significant difference between conditioning traffic, which can be very resource intensive, and simply administering traffic according to its DSCP. Conditioning the traffic consumes a significant amount of overhead because it requires both determining whether the traffic complies with the service contract, as well as remarking or dropping the traffic when necessary.

Core network routers administer different service classes by applying scheduling techniques. IP output scheduling in core network routers is provided through multiple priority queues that are managed by a queue management mechanism.

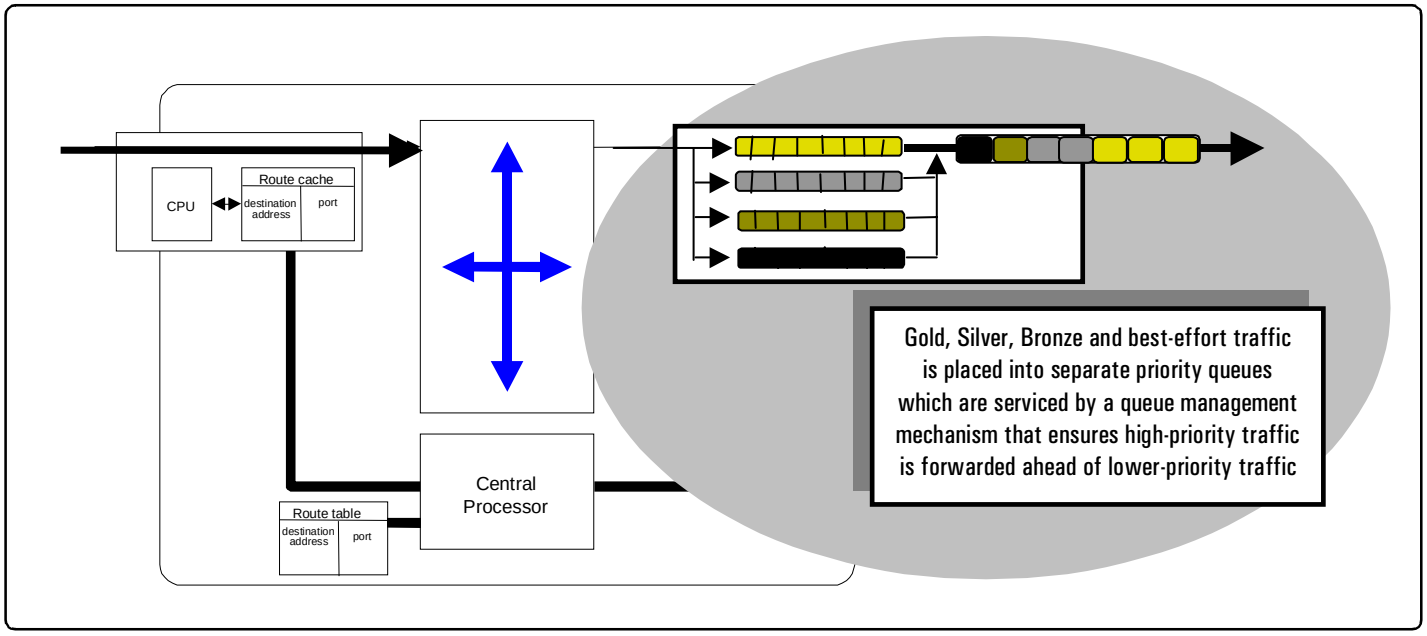


Figure 3: Multiple output queues residing on each output interface of the router provide different QoS levels to different classes of traffic.

## Output Queues

Output queues residing on each interface of a core network router, shown in Figure 3, provide a fine degree of control over the quality of service levels provided to different traffic classes. When congestion in the router creates contention for an output interface, output queues are used to hold the excess traffic as it awaits delivery. Traffic in these queues is forwarded from the output interfaces in an orderly fashion as congestion diminishes. In order to provide different Quality of Service levels to different types of traffic, output queues can be prioritized, so that traffic with different service handling requirements can be placed into separate queues. For example, traffic belonging to the 'gold service class' (AF PHB class 3), can be placed in a high priority queue, while best-effort traffic (default PHB) can be placed in the lowest priority queue. Using airline service as analogy, a congested output interface can be though of as an airplane, and the different seating sections (i.e. first class, business class and economy) on the plane represent the separate output queues that reside on the interface. Passengers, representative of the network traffic, are assigned to a seating section based on the type of ticket that they possess, just like traffic is assigned to an output queue based on the PHB that they posses. On the airplane, passengers sitting in first class receive better service than passengers sitting in business class, and similarly, the service provided to passengers seated in business class is better than that provided to passengers sitting in economy. This same type of relationship exists between the output queues, as traffic placed in a high priority queue is forwarded ahead of traffic waiting in lower priority queues.

## Queue Management Mechanisms

Output queues must be managed carefully to prevent the occurrence of unpredictable packet loss and excessive latency. Queue management mechanisms are used to ensure that high priority traffic, such as delay and loss sensitive voice, is forwarded ahead of low-priority traffic, while preventing occurrences of buffer starvation. Buffer starvation causes traffic waiting in a queue to be excessively delayed, or

dropped completely. This situation can occur for a number of reasons, but typically results from there being too many packets waiting to be queued, and not enough room on the queue (buffer space) to accommodate them. Weighted Random Early Detection (WRED) and Weighted Fair Queuing (WFQ) are queuing disciplines that algorithmically help routers cope with traffic congestion on the Internet.

## RED/WRED

Random early detection (RED) works with the transport control protocol (TCP) to detect and avoid congestion in the network core. When RED identifies that traffic is entering buffers faster than it can be forwarded, an algorithm is used to randomly discard packets. These intentionally dropped packets cause connection-oriented TCP to throttle back and slow the sender's transmission. By randomly discarding packets before congestion occurs, and forcing sources to reduce their transmit rate, the network gets much better overall throughput.

Weighted RED (WRED) is used to protect high-priority traffic from being randomly discarded when congestion occurs. WRED adjusts the discard parameters in the packet-dropping algorithm so packets belonging to high-priority flows (identified by the PHB aggregate under DiffServ) are far less likely to be dropped when congestion occurs. This enables specific traffic classes to bypass the arbitrary discard process so that QoS levels can be maintained for high-priority traffic.

## Weighted Fair Queuing (WFQ)

Weighted Fair Queuing (WFQ) uses a queue-servicing algorithm that provides preferential treatment to low-volume traffic flows, and allows higher-volume traffic flows to obtain equity in the remaining amount of queuing capacity. This process is used to prevent large traffic flows from consuming excessive bandwidth and starving smaller traffic flows. WFQ thus provides fair treatment to network traffic by ensuring that larger traffic flows do not arbitrarily starve smaller flows.

The weighted aspect of WFQ is dependent on the way in which the servicing algorithm is affected by other extraneous criteria. Under DiffServ, the servicing algorithm uses the DiffServ markings in the DS field to weight the method of handling individual traffic flows. The amount of queue resources given to a flow depends on the PHB class to which the flow belongs. Thus, WFQ can provide high priority traffic with more queue resources than a lower-priority traffic.

In summary, DiffServ provides a standard, highly scalable traffic classification model that is easily managed by core network routers. While it may take some time to determine the success of DiffServ, its simplicity, flexibility and initial wide acceptance in the user, vendor and ISP communities could finally make end-to-end IP QoS a reality. However, this discussion would not be complete without briefly reviewing another recent IP QoS initiative - Multiprotocol Label Switching (MPLS).

## Multiprotocol Label Switching (MPLS)

Multiprotocol Label Switching (MPLS) is an emerging technology that aggregates traffic into class flows that can be handled differently according to specific traffic conditions. While MPLS can be employed to deliver different classes of traffic, it is primarily implemented and recognized for its traffic engineering capabilities. Traffic engineering can be used to help deliver QoS - but its underlying premise is much different from the IP QoS technologies discussed above. While QoS revolves around the concept of traffic classification, prioritization, and management (as illustrated in the DiffServ model), traffic engineering revolves around the concept of path determination and flow manipulation.

MPLS engineers traffic by adding a label to each packet that enables the packet to be routed along a specific path through the network. Under conventional routing, IP traffic follows the shortest path through a network. In contrast, under MPLS, the route taken by IP traffic can be pre-determined by configuring explicit paths through the network. By moving traffic flows away from the shortest path determined by conventional routing, and onto less congested paths through the network, MPLS can better balance a network's traffic load and improve IP routing efficiency. By preventing the over- or under-utilization of network components, overall network response time and traffic throughput can be maximized.

Like ATM and IntServ, MPLS employs a signaling mechanism, such as the Resource Reservation Protocol (RSVP) or the Constraint-based Routing Label Distribution Protocol (CR-LDP), to reserve resources and establish traffic paths across the network. However, MPLS is not burdened by the scalability limitations of RSVP because it doesn't employ signaling to set up each individual flow as in IntServ, but rather uses RSVP to establish the traffic path over which an aggregate of traffic flows will traverse. Furthermore, MPLS

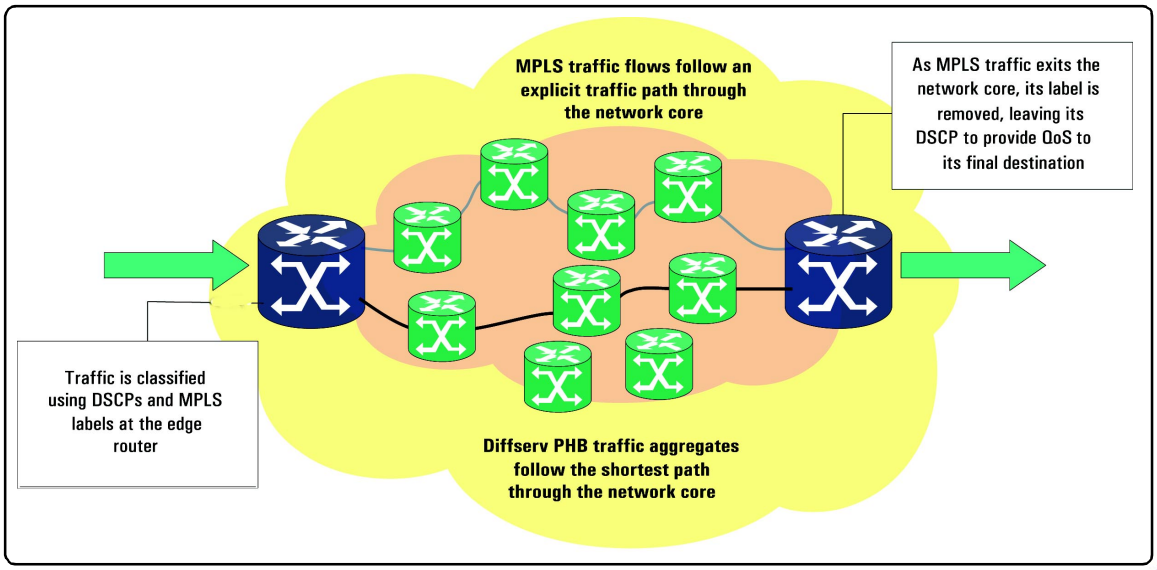


Figure 4: DiffServ and MPLS QoS models work in tandem.

introduces a set of IETF-approved extensions to RSVP (RFC 2205 and RFC 2209) which reduce the number of refresh messages and processing requirements of the protocol.

The traffic engineering capabilities of MPLS equips service providers with an unprecedented level of control over the flow of traffic through IP networks. However, MPLS is just one piece of the IP QoS puzzle. As illustrated in Figure 4, not all network traffic will follow the same configured path, and traffic that is engineered still requires QoS provisioning once it exits the MPLS network. Therefore, MPLS must work in tandem with DiffServ to provide end-to-end IP QoS. Some traffic will be routed along a specific path using MPLS, and other traffic will rely on DiffServ's per-hop QoS model to traverse the same network. Once the MPLS traffic exits the MPLS network, it can revert to its DiffServ markings to provide QoS to its final destination.

## Putting QoS to the Test

IP QoS is still a relatively new technology. The IP QoS technologies discussed above are still evolving, and new mechanisms continue to be introduced. This presents service providers with a considerable challenge as they begin to introduce new IP-based services. In order to bill customers according to a specified class of service, service providers need to be certain that their networks can deliver the negotiated QoS. It is therefore imperative to ascertain whether network equipment can actually deliver QoS guarantees before these services are deployed. The area most prone to congestion and performance degradation is the network core. The weakest link in the network core is routing. To address bottlenecks in the network core, router manufacturers have implemented significant improvements in router architectures. However, with Internet traffic doubling every 4 months, the ability of these new 'carrier-class' routers to effectively manage different classes of traffic under increasing network congestion must be determined. The remainder of the paper discusses the critical test methodologies for testing the QoS capabilities of a core Internet router.

### Testing QoS in Carrier-Class Router

Router performance can be measured in several ways. Within the scope of this paper, we are only concerned with router performance as it pertains to QoS delivery. As discussed earlier, QoS refers to the successful delivery of an agreed upon level, or class of service. A class of service is characterized by a set of performance parameters including delay, delay variation, throughput, and packet loss.

Testing the QoS performance of a carrier-class router requires making comparative measurements of the above performance metrics, for different classes of traffic, at increasing traffic loads. The test requirements can therefore be divided into two overall areas: traffic generation and performance measurements.

## Test Requirements

### Traffic Generation

The CoS aspect of QoS requires that certain traffic classes must be provided a predictable level of service. Testing a router's capability to provide a consistent level of service to a traffic class is twofold:

1. different traffic classes must be generated into the router under test
2. the conditions that force the router to prioritize these different traffic classes must be created

### Class Differentiation

The generation of multiple IP packet streams with complex traffic parameters is necessary to ascertain the router's ability to effectively manage many different traffic classes. Specifically, using different values of the TOS field, or different settings of the DS byte, IP packet streams representing different traffic classes must be simultaneously generated into the router under test.

Because the size and burst profile packets within a traffic flow can have a significant effect on a router's performance, the traffic classes generated into the router should also consist of packets with varying lengths and profiles, so the effect of these parameters on the successful delivery of the traffic class can be determined.

### Forced Prioritization

As previously discussed, a router manages different traffic classes through a series of prioritized output queues, coupled with a queue management mechanism that provides service to these queues. In order to determine the successful operation of these algorithmically managed queues, the conditions that force the router to prioritize and manage different traffic classes must be created - namely, the over-subscription of output ports.

The over-subscription of output ports requires both the ability to define explicit traffic paths through the router under test, as well as the ability to generate different traffic classes into the router under test at wire-speed. Different traffic classes must be simultaneously directed to the same destination port on the router under test, with the aggregate load of these streams exceeding the load capacity of the output interface. This means that the test traffic generated into the router under test must be of sufficient speed to fully congest the output ports on the router.

It is crucial to determine how successfully the router prioritizes and manages different traffic classes at different loads. Therefore, in addition to generating traffic at wire-speed, testing QoS performance also requires the ability to manipulate the traffic load of each traffic stream on the fly, and measure in real-time the effects of this manipulation on the delivered QoS.

### Performance Measurements

As mentioned earlier, delay, delay variation, throughput, and loss measurements must be provided for the different traffic streams traversing the router under test. In order to understand the interaction effects of different traffic classes and packet parameters, side-by-side stream measurements must be provided in real-time. For example, to determine whether the throughput of a high priority traffic class (e.g. voice) is adversely effected when a lower-priority traffic class (e.g. bursty web data) shows an increase in load requires that throughput statistics for both streams be provided in real-time, as the load of the low-priority stream is incremented.

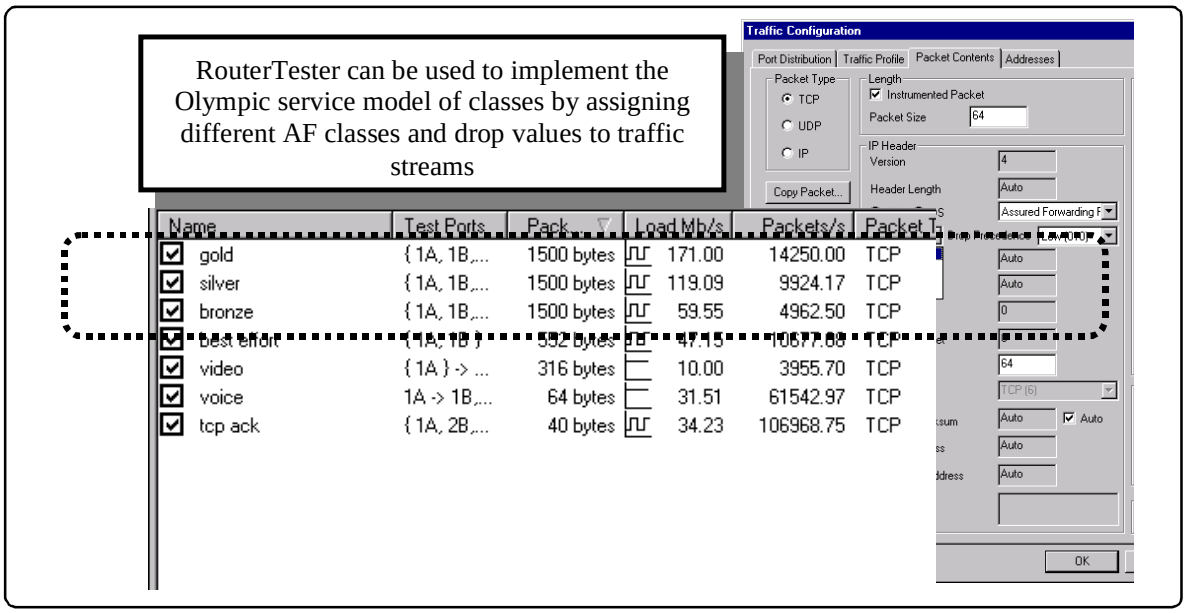


Figure 5: RouterTester fills the router under test with a rich mix of traffic classes and types.

## Using RouterTester to Test QoS

Designed specifically to test the performance of carrier-class routers, the Agilent Technologies RouterTester is able to fully satisfy each of the above test requirements.

### Traffic Generation

The CoS aspect of QoS requires that certain traffic classes must be provided a predictable level of service. Testing a router's capability to provide a consistent level of service to a traffic class is twofold: first, different traffic classes must be generated into the router under test; second, the conditions that force the router to prioritize these different traffic classes must be created.

### Class Differentiation

RouterTester generates up to 255 IP traffic streams, each with thousands of IP addresses and complex traffic parameters, from each port. This wire-speed traffic generation fills the router under test with a rich mix of traffic classes and types. As shown in Figure 5, RouterTester can be used to implement the 'Olympic Service' model described earlier in this paper. Using different classes and drop precedence values of the AF PHB, RouterTester can create bronze, silver, and gold service classes. RouterTester can assign a specific load to each traffic class, and then compare the measured throughput of each class against the intended load. Using different TOS values, RouterTester can also generate streams of different traffic types with varying service priorities. Different packet lengths and profiles can be applied to each of the generated streams so that the effects of these parameters on a router's QOS performance can be determined.

### Forced Prioritization

RouterTester can then define the explicit traffic path that each traffic stream will take through the router under test to create contention on the output ports. As shown in Figure 6, RouterTester has configured a high priority 'gold' stream of traffic, as well as a low priority 'best-effort' stream of traffic destined to the same output port. The aggregate load of these traffic streams (700 Mb/s) exceeds the load capacity of the output interface (622 Mb/s). When RouterTester generates these traffic streams into the router under test at wire-speed, over-subscription of the output interface should cause the

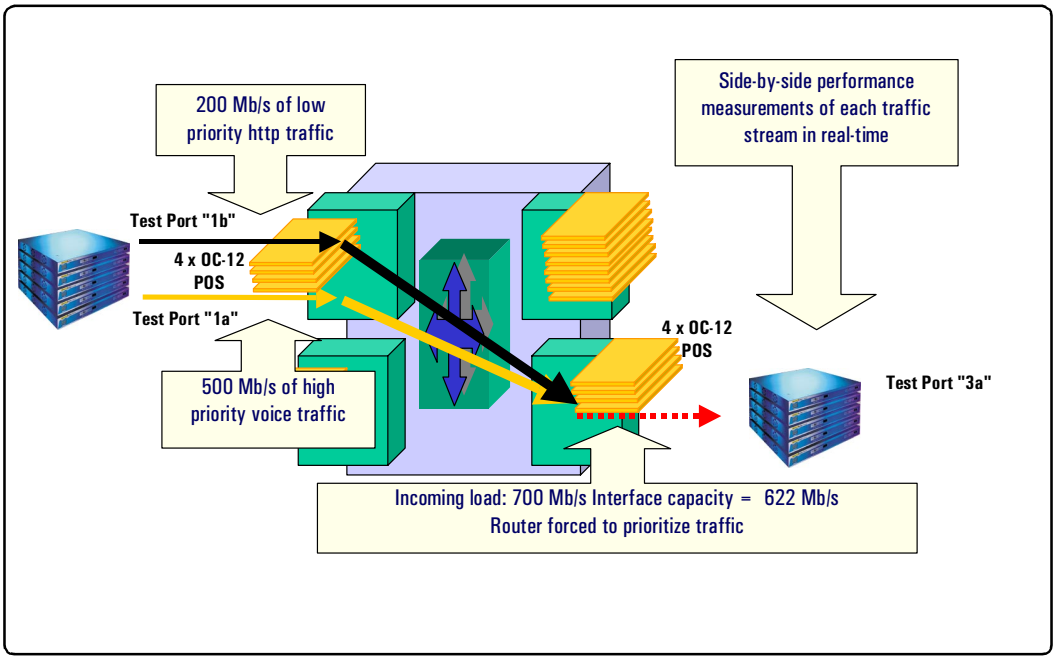


Figure 6: RouterTester creates the conditions under which a router is forced to prioritize and manage traffic.

router to prioritize the traffic within the respective output queues, and the queue management mechanism to provide preferential service to the high-priority queues.

To determine how successfully the router prioritizes and manages the different traffic classes at different loads, RouterTester enables users to manipulate the traffic load of each traffic stream on the fly, and measure in real-time the effects of this manipulation on the delivered QoS.

## Performance Measurements

RouterTester delivers correlated performance measurements in real time to identify the router's QoS performance. Comparison of packet throughput, latency and loss metrics between the streams in real-time reveals the router's ability to effectively manage different traffic classes under increasing load. RouterTester provides graphical and tabular output on both a per-stream and per-port basis.

Figure 7 shows the throughput performance of the two traffic streams identified in Figure 6. RouterTester's side-by-side graphical output reveals that the router has forwarded all 200 Mb/s of the high priority voice traffic, while only 390.04 Mb/s of the low priority http traffic (31.94 Mb/s is consumed by SONET overhead). RouterTester shows that the router under test is able to effectively mange these two traffic classes under an aggregate load of 700 Mb/s.

To isolate specific performance criteria and events, RouterTester's defines performance thresholds and triggers that can capture packets to memory for detailed packet analysis. RouterTester's powerful data reduction tools can be used to identify performance patterns and event sequences that reveal how the router responds to

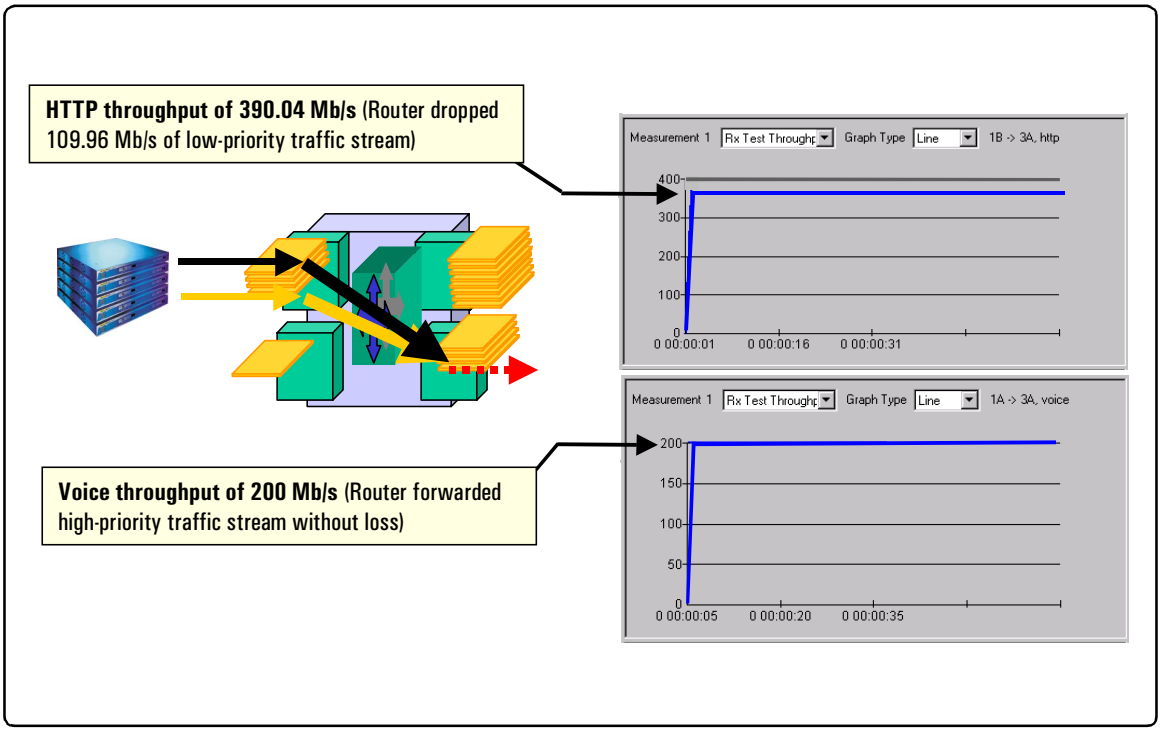


Figure 7: RouterTester provides correlated performance measurements in real time.

different traffic configurations and load structures. For example, Figure 8 shows the average latency across a stream of high-priority voice traffic. All packets exceeding an explicitly defined threshold of 400 microseconds are captured to memory for detailed IP analysis. Delay variation within the voice traffic can be examined to identify the underlying reasons for the performance degradation. RouterTester's powerful data reduction tools, protocol decodes and visual interpretations can be used to drill-down and thoroughly understand why the router was unable to deliver a consistent service level to the traffic.

By generating up to 255 traffic streams per port at wire-speed, RouterTester is able to fully stress the QoS capabilities of the router under test and provide valuable insight into the router's QoS performance capabilities and limitations.

## Conclusion

The introduction of new IP-based standards coupled with improvements in router architectures is paving the way for end-to-end IP QoS. Testing these emerging technologies before network deployment is essential to ensure their ability to deliver QoS guarantees for new IP-based services. RouterTester has been specifically designed to not only measure the quality of service capabilities of core network routers, but to also identify the underlying reasons for QoS performance limitations, so that QoS delivery can be continually improved.

## Acronymns

| | |
|---|---|
| AF | Assured Forwarding (QoS) |
| ATM | Asynchronous Transfer Mode |
| CoS | Class of Service |
| CR-LDP | Constraint-based Routing Label Distribution Protocol (MPLS) |
| CU | Currently Unused (IP header field) |
| DiffServ | Differentiated Services (IETF QoS model) |
| DSCP | Differentiated Service Code Point (IP header field) |
| EF | Expedited Forwarding (QoS) |
| HTTP | HyperText Transfer Protocol |
| IETF | Internet Engineering Task Force |
| IntServ | Integrated Services (IETF QoS model) |
| IP | Internet Protocol |
| ISP | Internet Service Provider |
| MPLS | MultiProtocol Label Switching |
| PHB | Per-Hop Behavior (QoS) |
| QoS | Quality of Service |
| RSVP | Resource Reservation Protocol |
| TCP | Transmission Control Protocol |
| TOS | Type of Service (IP header field) |
| VoIP | Voice over IP |
| WFQ | Weighted Fair Queuing (QoS) |
| WRED | Weighted Random Early Detection (QoS) |

# References

## Internet Drafts

Bernet,Y., S. Blake, and A. Smith, A Conceptual Model for DiffServ Routers, draft-ietf-diffserv-model-00.txt, June 1999.

## Request for Comments

RFC 2474, Definition of the Differentiates Services Field (DS Field) in the IPv4 and IPv6 Headers, F. Baker, D. Black, S. Blake, and K. Nichols, December 1998.

RFC 2475, An Architecture for Differentiated Services, F. Baker, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, December 1998.

RFC 2598, An Expedited Forwarding PHB, V. Jacobson, K. Nichols, and K. Poduri, June 1999.

RFC 2597, Assured Forwarding PHB Group, F. Baker, J. Heinanen, W. Weiss, and J. Wroclawski, June 1999.

## Textbooks

Ferguson, Paul, and Geoff Huston, Quality of Service: Delivering QoS on the Internet and in Corporate Networks, John Wiley & Sons, New York, 1998 (ISBN 0-471-24358-2).

## Whitepapers

Semeria, Chuck (1999), Multiprotocol Label Switching: Enhancing Routing in the New Public Network, Whitepaper. Juniper Networks.

Semeria, Chuck (1999), Traffic Engineering for the New Public Network, Whitepaper. Juniper Networks.

Trotter, Guy (1999), Testing Carrier-Class Routers with Internet-Scale Simulation, Whitepaper. Agilent Technologies.

## Other References

Kaufman, David H. (1999), "Delivering Quality of Service on the Internet," Telecommunications (Feb), 35-42.

Steinke, Steve (1999), "ATM and alternatives in the wide area backbone," Network Magazine (Jul).

Passmore, David (1999), "Classifying the Traffic," Business Communications Review (Aug), 18-19.

Stephenson, Ashley (1998), "DiffServ and MPLS a Quality Choice," Data Communications (Nov), 73-77.

Klessig, Bob and Mick Seaman (1999), "Going the Distance With QoS," Data Communications (Feb) http://www/data/com/issue/990207/distance.html.

DeVeaux, Paul and Annie Lindstrom (1999), "Follow the QoS road," America's Network (Jun).

This page intentionally left blank.

## Agilent RouterTester

RouterTester provides true Internet-scale testing through realistic routing protocol support, multi-stream wire-speed traffic generation and real-time analysis, and multi-port scalability. RouterTester is set to grow as the testing needs of the carrier class router industry evolve to meet the challenges of scale and Quality of Service within the Internet.

# www.Agilent.com/comms/RouterTester


**United States:**
Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

**Canada:**
Agilent Technologies Canada Inc.
5150 Spectrum Way
Mississauga, Ontario
L4W 5G1
1-877-894-4414

**Europe:**
Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-9999

**Japan:**
Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

**Latin America:**
Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 267-4245
Fax: (305) 267-4286

**Asia Pacific:**
Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 2599-7889
Fax: (852) 2506-9233

**Australia/New Zealand:**
Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881




**Specifications subject to change.**


5968-9722E 03/00 Rev A


**Agilent Technologies**
Innovating the HP Way