WHITE PAPER

Turning Packets into Profits with Intelligent QoS

Dave Olisar
Staff Marketing Manager

Tellabs Intelligent Service Edge™ Solution enables Service Providers to take advantage of new high-margin IP-based services.

Abstract
The Internet Protocol (IP) has always been a 'best effort' technology, and most people dealing with IP applications such as the Web or streaming media support this fact. How fast Web pages download and how smoothly video streams play are at the mercy of packet loss and routing decisions. In contrast, traditional voice communications systems have always been extremely reliable in terms of quality and response — no one expects to lose one out of 10 words during a telephone conversation. The ability to converge these services and offer new premium business data services from a single platform requires the capability to deliver hard Quality of Service (QoS) guarantees. Tellabs delivers these capabilities with an Intelligent Service Edge solution. This solution brings service intelligence to the edge of the network — closer to the end-customer for service speed and innovation, with a simpler, less complex network to maximize profitability.

Business Drivers for Intelligent Quality of Service
The majority of business applications today rely on IP protocols. However, the majority of services that are mission critical do not touch a carrier's IP network due to the lack of service quality and reliability of IP backbones. The IP backbone is consumed largely by best-effort Internet traffic. According to industry observers RHK and Probe Research, the growth of Internet traffic has most recently been tracking at more than 100% and will continue at a rate of 60 to 80% annually for the next couple of years.

This surging growth, combined with the current flat-rate pricing structure for "all you can eat" Internet service, is squeezing all profitability from the IP network. In fact, IP revenues are decreasing even as IP traffic booms (see Figure 1) resulting in an increasingly lower price-per-bit in revenue. The current IP business model is clearly broken.

The challenge is to find a way to extend the applicability of IP-based services beyond best-effort Internet service and bring service guarantees to IP. Without them, the future of a profitable IP-rich service environment envisioned by so many is neither practical nor realistic.

Next generation data services require intelligent Service Level Agreement (SLA) support from service providers. Intelligent SLAs provide more than the availability and mean time-to-repair guarantees typical of today's IP network SLAs. Intelligent SLAs deliver strict QoS guarantees, including bandwidth and packet delay guarantees, and advanced usage-based pricing structures to support new business data services with seamless support for existing services. As businesses increasingly rely on electronic connections to customers, employees, suppliers and investors, they must depend on the next generation of Internet Protocol Multi-Protocol Label Switching (IP/MPLS) networks to deliver mission-critical business data, voice and video, as well as Internet communications.

Moving to a more profitable IP network model requires a new means of service delivery — one that uses QoS capabilities to enable service providers to offer differentiated services reliably over a common network. By consolidating overlay networks onto a single QoS-enabled (IP/MPLS) network, service providers can reduce capital and operational expenditures and increase revenues and profitability with new differentiated services. QoS enables a service provider to offer enforceable SLAs, which in turn leads to additional revenue — turning performance promises into profits.

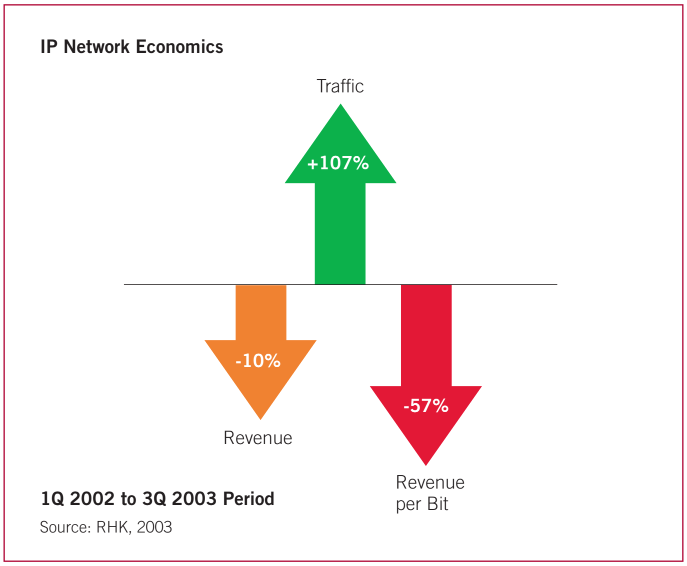

Figure 1. Business Model for IP Traffic is Broken

For example, Virtual Private Network (VPN) services based on IP/MPLS technology can match the QoS, availability and predictability characteristics of traditional Asynchronous Transfer Mode (ATM), Frame Relay (FR) and leased line services.

End-users Demand Quality IP Services

Since IP networks have become business tools and not just the backbone of a best-effort Internet, QoS has gone from being a nice bonus to a necessity. Businesses are demanding SLAs as a standard part of a data services contract. If a mission-critical business application is delayed due to poor network quality, users may incur significant loss in business revenue and will expect the service provider to live up to the financial terms of the SLA. Given the competitive nature of today's networks, a service provider who fails to meet SLA commitments or who cannot offer SLAs with tight guarantees, is going to be hard pressed to compete.

By contrast, service providers offering QoS can now differentiate themselves by providing varying levels of QoS for different types of business applications. For example, a VoIP service receives bandwidth and delay guarantees, while mission-critical business data may receive a bandwidth guarantee and burst allowance, but no delay guarantee. Businesses will appreciate being able to pay for what they need, and not oversubscribing bandwidth to ensure the success of all applications. In addition, service providers will realize cost savings by moving voice from legacy delivery systems into a single IP/MPLS network that delivers multiple network services.

From a technical perspective, enterprise networks have changed significantly since the early '90s. Today the dominant enterprise networking technology is IP over Ethernet, and the best solution for the connectivity needs of such networks, from both a technical and an economic point of view, is a technology with Ethernet service interfaces and an IP-oriented control plane. In other words, an MPLS-based Layer 2 and Layer 3 VPN infrastructure with hard QoS support.

This creates an opportunity for network convergence with a single IP/MPLS infrastructure serving as the traffic growth platform for both LAN and Internet services. In addition, voice can also be expected to migrate to IP networks, driven by the fact that the cost of IP-based voice connectivity is approaching zero. Finally, new value-added services are also IP-based, and thus naturally delivered over an IP-based connectivity layer.

QoS Requirements for Applications

The architecture of an IP/MPLS network supports two types of QoS: hard and soft; however, today's networks only support soft QoS. Hard QoS is based on per-flow resource reservations that enable mutually independent QoS guarantees among flows. Soft QoS (a.k.a. Class of Service) is based on aggregation of flows into classes and service that is defined in terms of the relationship between classes, e.g. share of available bandwidth. A flow is a string of packets associated with a particular application. Both types of QoS have value, and in particular service providers require hard QoS in order to deliver the next generation of services such as VoIP and video applications.

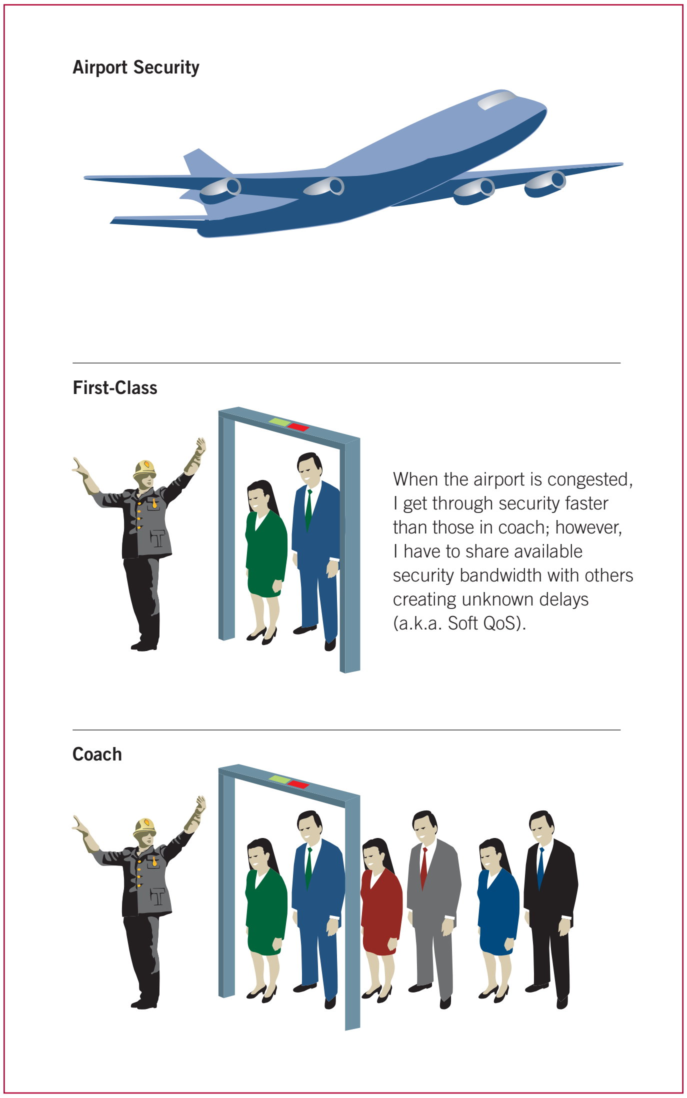

Figure 2. Soft QoS Example: Airport Security

The service needs of different applications can be represented as a set of QoS parameters that include delay, jitter, guaranteed bandwidth and packet loss rate. Delay refers to the time difference between the transmission and reception of a packet end-to-end, and can significantly impact the QoS for real-time applications. Jitter refers to the variation in packet delay, and has an adverse effect on real-time applications such as VoIP. Each QoS parameter is measured against an objective, such as a minimum bandwidth guarantee of 10 Mbps. These objectives, appropriately tuned to a customer's needs, form the basis of an SLA, which is the key to turning promises into profits. Table 1 depicts the application performance characteristics required by a set of typical applications.

Let's take a real-life example to illustrate the differences between hard QoS and soft QoS. Imagine your favorite airport with its separate security lines for coach and first-class passengers. When the airport is busy there are typically multiple first-class passengers waiting in the first-class line. The first-class passengers are served faster than coach passengers; however, they still have to wait. This is an example of soft QoS — high-priority traffic is handled faster than low-priority traffic; however, during congestion, all traffic slows down. This is not acceptable for applications such as voice and video conferencing, which require guaranteed bandwidth plus very little delay and jitter.

Now, imagine an airport that offers a unique service for first-class passengers — immediate security check service no matter how busy it is at airport security. During the busiest flying day of the year, every first-class passenger is checked immediately while coach passengers have a long wait. This special handling of first-class passengers is an example of hard QoS — high-priority traffic serviced independently from all other traffic, even from other high-priority traffic. This is the type of treatment required by real-time applications such as voice and video conferencing — each real-time flow receives a guaranteed amount of bandwidth with minimal delay and jitter.

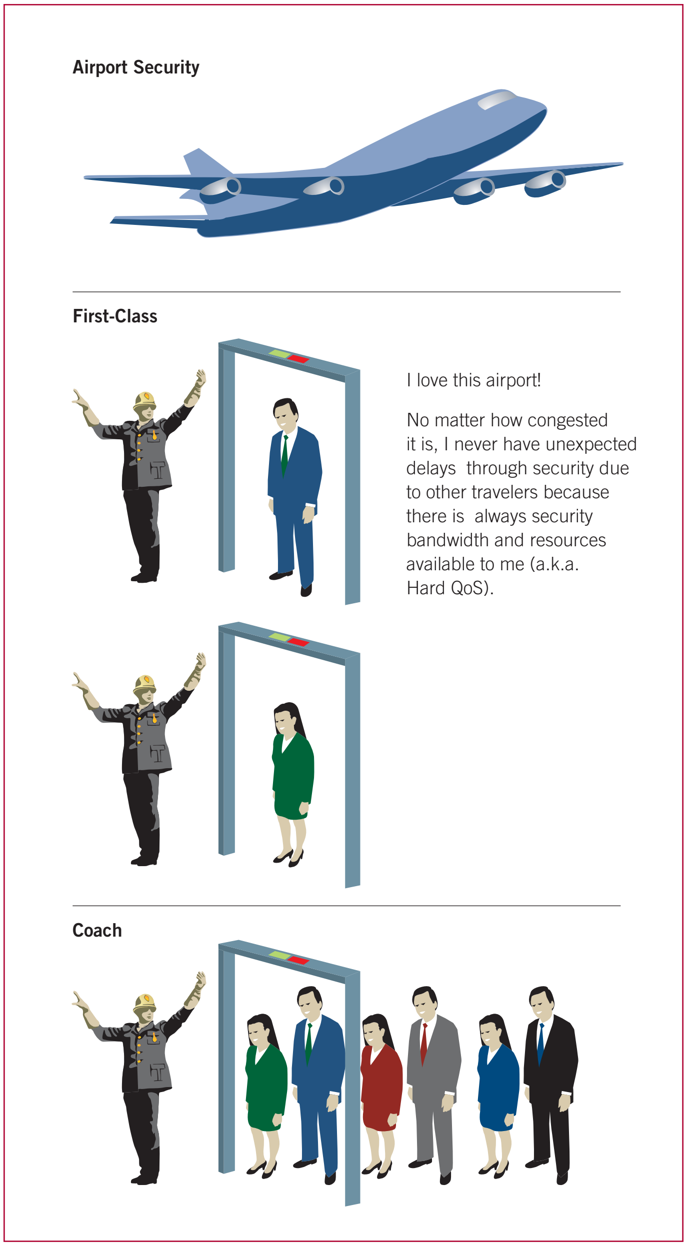

Figure 3. Hard QoS Applied to Airport Security

	Bandwidth		Sensitivity		
Services	**Guarantee**	**Amount**	**Jitter**	**Delay**	**Loss**
VoIP	Yes	Low	Strict	Strict	Moderate
Video Conference	Yes	High	Strict	Strict	Moderate
Streaming Video	No	High	Moderate	Moderate	Moderate
Mission-critical Business Data	Yes	Varies	Tolerant	Tolerant	Strict
Best Effort	No	Varies	Tolerant	Tolerant	Tolerant

Table 1. Application QoS Requirements

Now imagine an airport that offers a unique service for first-class passengers — immediate security check service no matter how busy it is at airport security. During the busiest flying day of the year, every first-class passenger is checked immediately while coach passengers have a very long time to wait. This special handling of first-class passengers is an example of hard QoS — high-priority traffic serviced independently from all other traffic, even from other high-priority traffic. This is the type of treatment required by real-time applications such as voice and video conferencing — each real-time flow receives a guaranteed amount of bandwidth with minimal delay and jitter.

Example: An Enterprise Network Benefiting from Hard QoS

Today, enterprise customers may turn to a number of technologies to accomplish their networking needs. The existing ATM/FR/Leased Line infrastructure is able to support SLAs that guarantee specific levels of QoS. To match the capabilities of the legacy connectivity infrastructure, the IP/MPLS infrastructure also requires tools to support these SLAs.

Figure 4 illustrates a business customer with multiple applications supported by a single IP/MPLS network. The Tellabs® 8800 Multi-Service Routers Series provides the ideal solution to achieve this multi-service support over a single IP/MPLS network because of its hard QoS capabilities. In support of the applications identified in Figure 4, a service provider can create three service classes for its customers:

- **Real Time,** similar to the existing leased line services or ATM CBR/VBRrt services with their own hard QoS requirements;
- **Business Data,** with a committed rate and the possibility of bursting above the committed rate, similar to FR with graceful discard; and
- **Best Effort,** for Internet service.

The Real-Time service class supports the business's voice and video applications using hard QoS to match the leased line or ATM SLAs that were previously established. The Business Data service class supports the mission-critical business applications using hard QoS to match the Frame Relay SLAs. Lastly, the Best Effort service class uses soft QoS to provide Internet services.

Expanding upon this example, the Tellabs 8800 Series of Intelligent MSR support of an IP/Ethernet service interface with hard QoS enables service providers to market new IP/Ethernet services to their existing Frame Relay and ATM customers. If the business customer depicted in Figure 4 opens a new branch office, the customer can install a low cost Fast Ethernet (FE) connection instead of a more expense frame relay connection. This FE connection, using the Business Data service class, will communicate seamlessly with a regional office that has an existing FR connection. The Tellabs 8800 Series enables a business customer to easily migrate to IP/Ethernet services which include new capabilities such as interface rates up to 10 Gbps and multipoint-to-multipoint connectivity.

This example illustrates the convergence of various networks into a single IP/MPLS network while supporting the SLA requirements of each application. Businesses will appreciate being able to pay for what they need, and not oversubscribing bandwidth to ensure the success of all applications. In addition, service providers will realize cost savings by moving voice from legacy delivery systems into a single IP/MPLS network that delivers multiple network services.

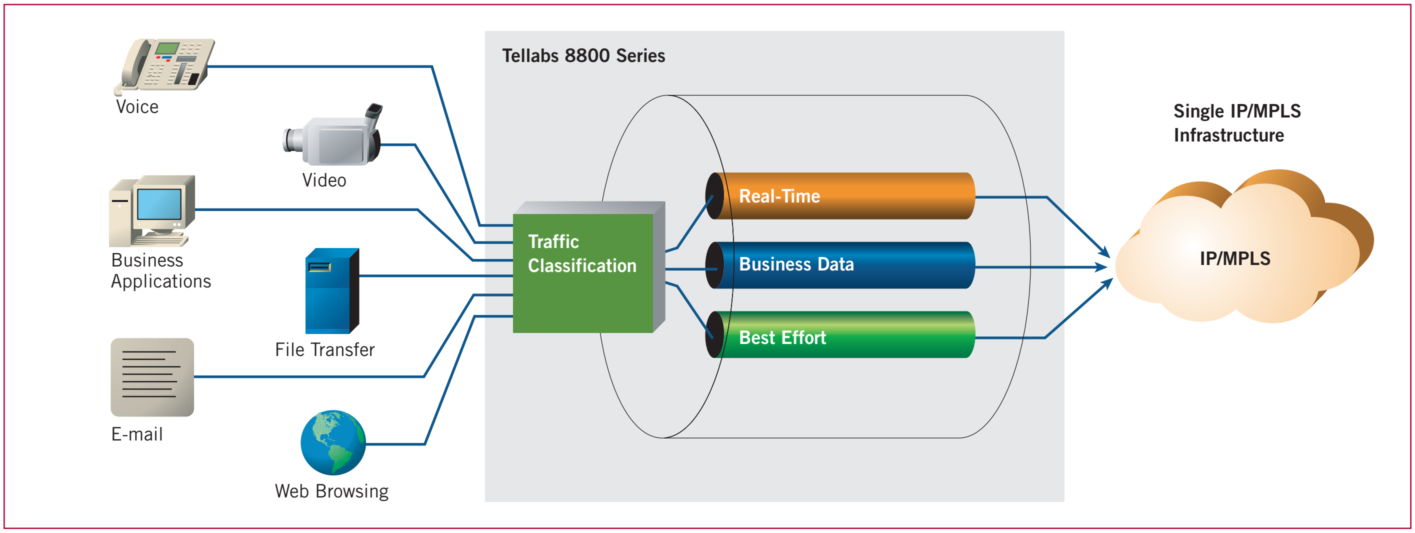

Figure 4. Multiple Services Supported Over a Single IP/MPLS Network

Requirements for Delivering Intelligent QoS

Hard QoS support is the key to deliver multiple services in a common IP/MPLS network. To ensure QoS for various service classes, the edge router must be capable of handling packets based on their associated priority indicator, such as ATM Class of Service for PVCs, FR Priority PVCs, IP ToS bits, Diff-Serv DS field, MPLS EXP bits and so on.

The Tellabs 8800 Series of Intelligent MSRs has the four key components needed to provide hard QoS:

- **Per-flow / Per-group Traffic Management**
  The QoS requirement of the flow and the location of the router in the network (e.g., edge, core, etc.) are the keys to what granularity of traffic management should be applied. At the edge of the network, a large number of microflows, each with a highly fluctuated traffic rate, are typically carried on the same trunk. Short-term congestion tends to happen more frequently on this type of link. A fine granularity of traffic management (per flow) is required to ensure QoS for well-behaved flows. At the core network, most of the microflows have been policed, shaped and aggregated into macroflows. Therefore, short-term congestion is less likely to happen and a coarse granularity of traffic management (per group) is sufficient to process the mass traffic.
- **Priority Queuing**
  Priority Queuing is a straightforward way to deliver QoS for each service class. It is composed of two components: a multiple queuing mechanism specific to each particular traffic class and a scheduler that places high-priority traffic out first over low-priority traffic.
- **Deficit Round Robin (DRR) Queuing Algorithm**
  DRR is a Per Flow Queuing (PFQ) mechanism that supports Weighted Fair Queuing (WFQ). DRR ensures that the bandwidth which is not being allocated is fairly shared by other flows according to their preconfigured weights.
- **Per-flow / Per-group Call Admission Control (CAC)**
  CAC plays an important role in resource management. Both priority queuing and DRR cannot function as expected without careful resource allocation. The CAC is used to ensure that a router has sufficient resource (bandwidth, buffers) to deliver all subscribed services with a satisfactory QoS. When an end user requests a new connection, all routers along the route need to be checked to see if they have sufficient remaining resources to support the new connection. If they do, the new connection is accepted; otherwise, it is rejected.

Conclusion

The growth of Internet traffic forces investments in IP infrastructure. Moving to a more profitable IP network model requires a new means of intelligent service delivery — one that uses QoS capabilities and enables service providers to offer differentiated services reliably over a common network. By consolidating overlay networks onto a single QoS-enabled IP/MPLS network, service providers can reduce capital and operational expenditures and increase revenues and profitability. QoS enables a service provider to offer enforceable SLAs, which in turn leads to additional revenue — turning performance promises into profits. The Tellabs 8800 Series of Intelligent Multi-Service Routers delivers this Hard QoS-enabled IP/MPLS solution.

One Tellabs Center
1415 West Diehl Road
Naperville, IL 60563 U.S.A.
Tel: +1 630 798 8800
Fax: +1 630 798 2000