ACCEPTED FROM OPEN CALL

# Toward IP Virtual Private Network Quality of Service: A Service Provider Perspective

*Jingdi Zeng and Nirwan Ansari, New Jersey Institute of Technology*


## ABSTRACT

To complement classical enterprise wide area network infrastructures, IP (based) virtual private networks have been gaining ground, with the capability of offering cost-effective, secure, and private-network-like services. In order to provision the equivalent quality of service of legacy connection-oriented layer 2 VPNs, IP VPNs have to overcome the intrinsically best effort characteristics of the Internet in this multimedia era. This article discusses the IP VPN QoS issue from a service provider point of view, where QoS guarantees are carried out at the network level as well as at the node level. It presents the whole picture by highlighting and stitching together various QoS enabling technologies from recent research and engineering work.


## INTRODUCTION

Virtual private networks (VPNs) complement classical enterprise wide area network (WAN) infrastructures, aiming to accommodate mushrooming telecommuters, road warriors, and business partners dispersed around the world. They carve public WAN links out of the rest of the network, and thus connect sites through WANs or provide the remote access to enterprise networks, all in a private-network-like manner, that is, the same policies for security, manageability, quality of service (QoS), and so on. The VPN hype will continue in years to come, due to the rising desire for economical, reliable, and secure communications. Cahners In-Stat Group estimated that VPN services would hold a $23.7 billion strong share of the $104.4 billion worldwide IP service revenues in 2005.

A downside shared by legacy layer 2 VPN strategies, such as frame relay and asynchronous transfer mode (ATM) virtual networks, is the connection-oriented characteristic; in the network core, the mesh of the permanent virtual circuits required to provision redundancy becomes costly and does not scale well. For a bigger market share, a scalable and cheaper VPN solution is sought; this is where the Internet, with its global reachability and cost effectiveness, comes into play. Enabling a low-cost secure IP solution to replace expensive dedicated WANs, IP VPNs can be broadly classified into three categories: remote access VPNs connect remote users to the enterprise WAN; intranet VPNs connect branch offices and home offices within the enterprise WAN; and extranet VPNs supply business partners with limited access to the enterprise WAN.

There are two typical VPN deployment strategies. First, taking control of their VPN services, enterprises adopt and manage their own VPN-enabled customer premises equipment (CPE) devices. Second, enterprises outsource part or all of their VPNs to a service provider (SP). The VPN management complexity is then shifted to provider edge (PE) devices. The second strategy, an SP perspective solution that will be addressed in this article, is becoming fairly popular. It gives SPs a foothold in the enterprise networks for new revenues, and minimizes/eliminates enterprises' in-house need for network management expertise as well.

## SCOPE

This article assumes the following premises. First, peer-to-peer VPNs, all of whose routers have the capability to forward VPN traffic to appropriate destinations, are addressed. Overlay VPNs, alternative implementations that only take VPN tunnel endpoints into consideration, have no control over intermediate routers; they cannot deliver end-to-end QoS, and therefore are of no interest here. Second, the term VPN SP is used in the rest of the article to represent an Internet SP that provisions VPN services. Third, technical approaches to the IP VPNs discussed in this article utilize IP over IP (IETF RFC 1853), the IP security (IPSec) suite (IETF RFC 2402, 2406), and generic routing encapsulation (GRE) (IETF RFC 1701) protocols. Fourth, end-to-end QoS in this article means QoS enforcement between SP PE devices. The last mile from the subscriber edge to the SP edge is under the control of the subscriber.

Without drilling into the details of various enabling techniques for VPN tunneling, encryp-

*This work is resulted from the effort on incorporating IP VPN services into a 10G Ethernet product for an industry vendor.*

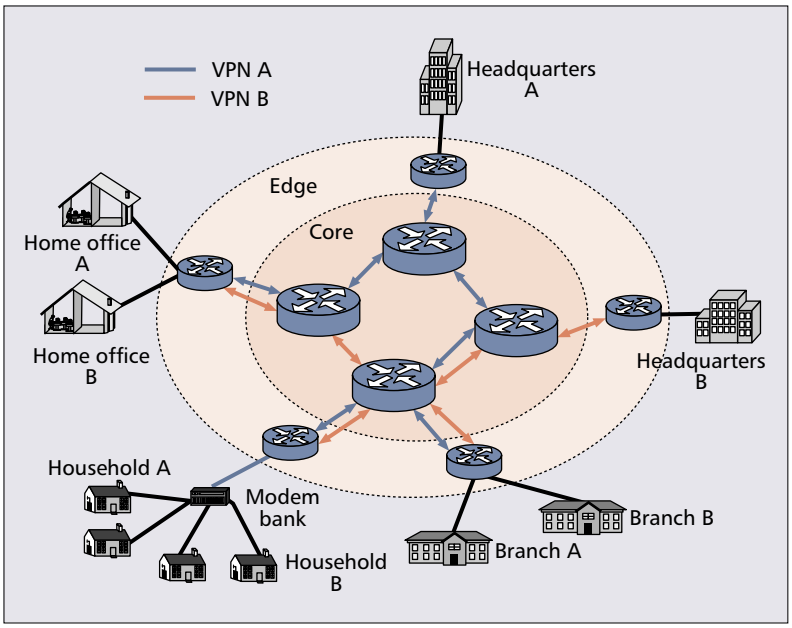


**Figure 1.** *The general IP-based VPN architecture.*

tion, authentication, and network management, the typical IP VPN deployment architecture is depicted in Fig. 1. It illustrates one (or multiple) SP network(s) gluing together enterprise networks, SP edge routers, and core routers to accommodate overlapping VPNs. In the rest of the article, SP edge routers and core routers will be referred to as PE routers and P routers, respectively.

## The QoS Issue in IP VPNs

Compared to layer 2 strategies (e.g., frame relay and ATM), IP VPNs inherit the flexibility and simplicity of connectionless IP networks. Utilizing IP VPNs can save users more than 50 percent of the connectivity cost of the corresponding frame relay deployment [1]. However, the Internet is a two-edged sword: its ubiquity offers VPNs more potential to grow, and yet it is not the "right" network to support QoS, due to its intrinsically best effort characteristics. What are called for, therefore, are standards-based appropriate QoS mechanisms for IP VPNs.

The QoS guarantee is the capability of a network infrastructure to deliver different levels of services. Its metrics include, but are not limited to, packet loss, delay, delay jitter, bandwidth guarantee, and throughput. In addition to information security, VPN services have various QoS requirements. For instance, executive videoconferencing may need stringent QoS as well as security requirements, whereas a secure database transaction may tolerate a certain QoS downgrade when the network resource is in short supply. In general, the VPN QoS can be delivered on a VPN subscriber and/or application type basis, as depicted in Fig. 2; the whole issue can be viewed as handling multiple traffic classes/aggregates with different QoS criteria.

To yield the equivalent end-to-end QoS of connection-oriented layer 2 VPNs, IP VPNs fulfill QoS control in a hierarchical manner. First, following service level agreements (SLAs) with subscribers, VPN SPs identify a route (or routes) capable of offering the required QoS and provision appropriate resources (e.g., bandwidth). Second, VPN QoS parameters are pushed down to router interfaces along the identified routes, by utilizing a certain centralized or signaling-based mechanism. QoS is then enforced by queuing and scheduling mechanisms in the routers. Bearing in mind this hierarchical framework, the rest of the article will provide a glimpse into QoS enabling technologies of IP VPNs.

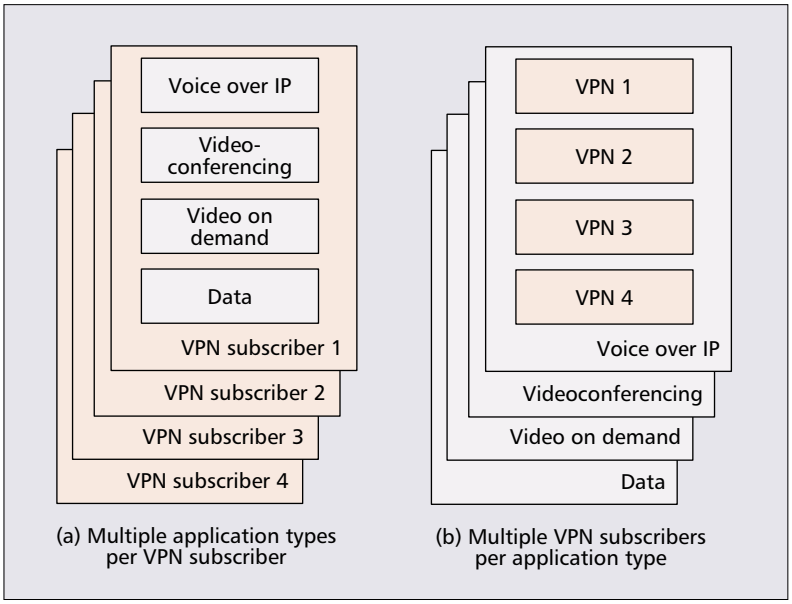


**Figure 2.** *Two examples of how VPN SPs deliver QoS.*

## IP QoS Architectures

IP VPNs may adopt a number of IP QoS architectures whose differences, in terms of SLA policies, are shown in Fig. 3. Different architectures often use different mechanisms to establish network routes and enforce QoS guarantees.

### Integrated Services

Integrated services (IntServ), along with the Resource Reservation Protocol (RSVP) (IETF RFC 2205), provides hard, end-to-end, fine-grained service guarantees; all routers in the network participate in RSVP signaling to reserve, tear down, and manage appropriate resources. RSVP signaling often implies a per-flow resource allocation identified by a five-tuple (transport protocol, source address and port, destination address and port).

IntServ/RSVP leads to a severe scalability difficulty because it is impossible for a core router to maintain the state of all application flows routed through it. However, it may be implemented on a limited scale, for instance, in an enterprise network; or in the core network where RSVP is under the control of a network management system to set up QoS-capable routes for traffic aggregates. The multiprotocol label switching (MPLS) working group likewise proposed use of an extended version of RSVP (IETF RFC 3209) to set up explicit routes in the core network.

### DIFFERENTIATED SERVICES

Differentiated services (DiffServ) define three types of per-hop behaviors (PHBs): expedited forwarding (EF), assured forwarding (AF), and best effort (BE); they specify how packets will be forwarded. With certain specifications in the packet header, customers indicate which type of service they require for an application. The philosophy of "move the complexity toward the edge" has led to a widely accepted concept that the DiffServ architecture should be implemented in the core, pushing IntServ to the edge.

The DiffServ infrastructure has been rather favored in IP VPN implementations due to the following facts: DiffServ handles traffic aggregates, and is thus capable of differentiating QoS per VPN or per application within a VPN; DiffServ QoS operations become fairly straightforward when handling VPN traffic with explicit destinations; the scalability advantage of DiffServ benefits multi-SP VPN deployments.

As will be noticed, the majority of strategies in this article are based on DiffServ, taking the mainstream technologies into consideration.

## THE VPN NETWORK PERSPECTIVE

Requiring the comprehensive information of a network, QoS operations at the VPN network level include resource provisioning, admission control, and routing. They can also be referred to as control plane functionalities.

An SLA between a VPN subscriber and its SP is a fundamental component. In addition to the charging and compensation matters in the event of an agreement violation, it defines conventional specifications such as the service availability and offered service (e.g., bandwidth, latency, packet loss, hop count, and cost); other VPN-specific criteria, such as VPN tunnel start time, duration, and redundancy, are included as well. VPN SPs are therefore challenged to provide services that meet this quantifiable commitment (i.e., SLA).

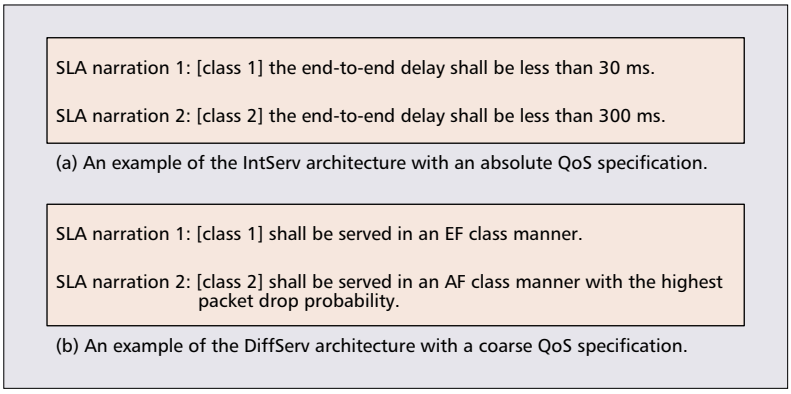


**Figure 3.** *Example differences, in terms of the SLA policy, of IP QoS architectures.*

### MANAGEMENT INFRASTRUCTURE

While the time-consuming, prone-to-error manual/static resource provisioning is still in practice, notable efforts have been made to bring more automation and intelligence into the VPN network operations.

An automated software agent, namely the VPN service broker, has been under intensive discussion for VPN QoS management. It monitors and enforces the service as specified in an SLA by carrying out the functionality of a system administrator, such as dynamic service configuration, VPN tunnel admission, and capacity provisioning. This concept can be implemented as an internal entity that does interdomain resource allocation and pushes the configuration information down to routers within an SP domain. It can also be adapted as an external entity that han-

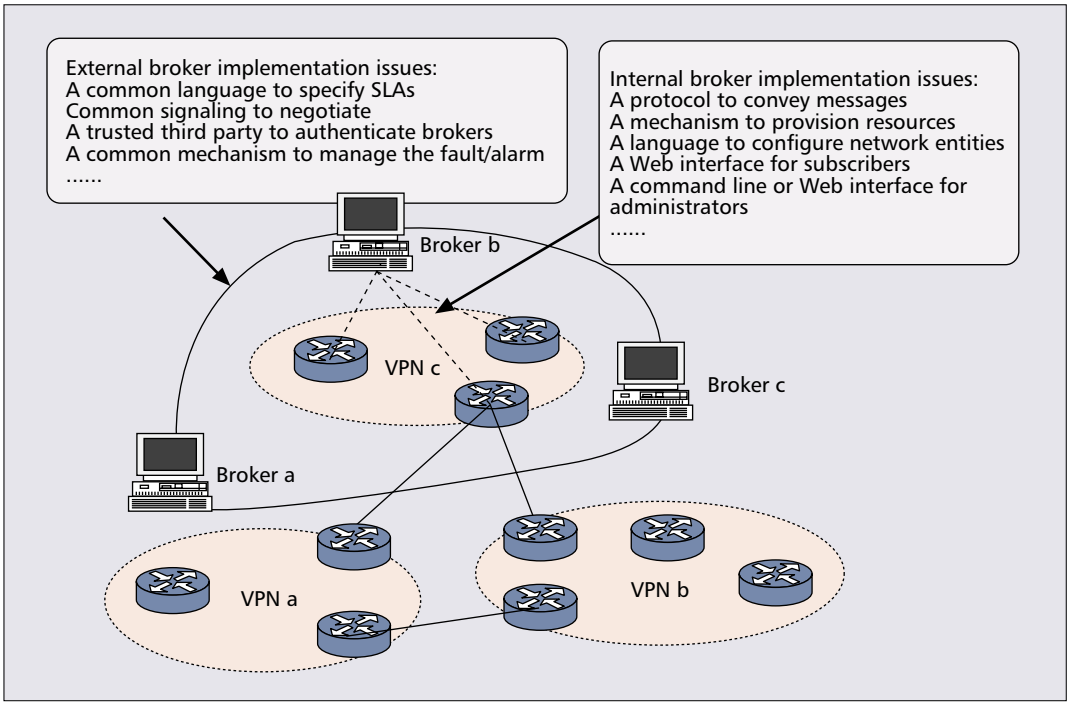


**Figure 4.** *The general VPN service broker infrastructure.*

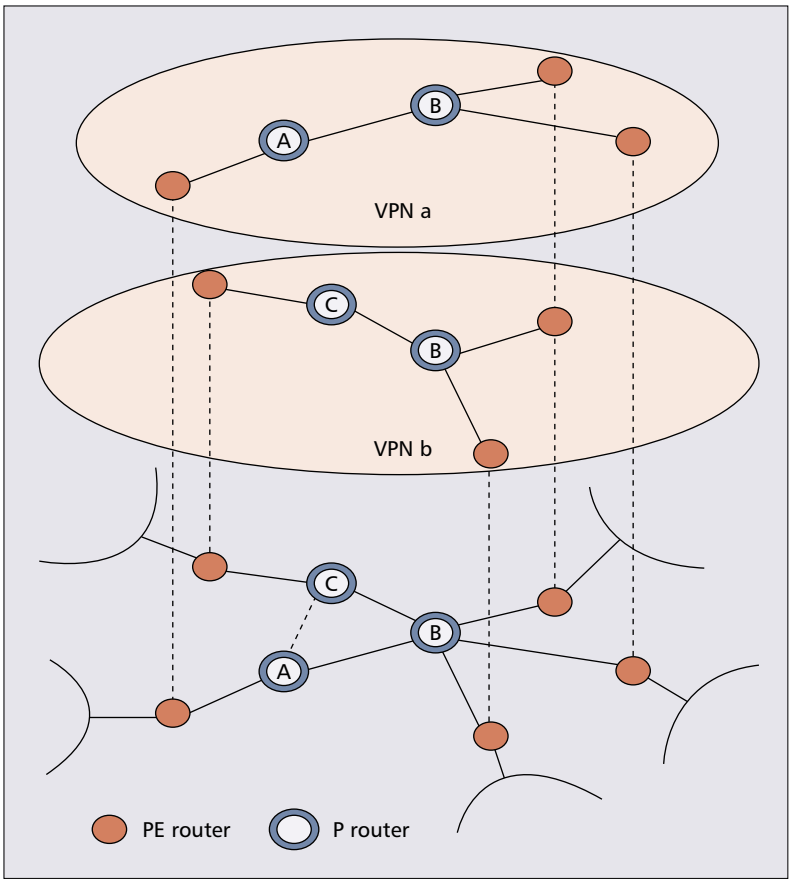


**Figure 5.** *An example of VPN resource provisioning, where QoS-warrant routes (topologies) of VPNs a and b have been reserved from the original network.*

dles VPN SLA requests and agreements with peer external brokers of adjacent VPN SPs.

There are full-fledged standards available for the intradomain service broker implementation, such as the policy-controlled network structure (IETF RFC 3060), Simple Network Management Protocol (SNMP) (IETF RFC 2571), common open policy services (COPS) (IETF RFC 2748), and Lightweight Directory Protocol (LDAP). When several SPs collectively provide VPN services, however, the interdomain broker implementation poses a new challenge. Heterogeneous operation support systems (OSSs) of different domains demand a means of exchanging accounting, billing, and resource provisioning information. For this interdomain federation, therefore, an open and standardized framework as well as interfaces between OSSs are under intensive investigation. The general view of the service broker system is depicted in Fig. 4, taking both the current status and future expectations into account.

As for today, although there is no complete standard suite available, a big amount of work has been done by project groups to tackle the interdomain federation issue. A generic and high-level interdomain prototype system and a general QoS-enabled VPN management system [2] were developed in the CATI [3] project. Adopting the generic network model [4], a telecommunication management network (TMN) (International Telecommunication Union — Telecommunication Standardization Sector, ITU-T, M3010) compatible infrastructure was suggested; it utilizes a cross-domain VPN manager to handle end-to-end VPN service activation and provisioning. Aligning with the telecommunications information networking architecture (TINA) [5] principle, a software platform [6] for VPN connection management, VPN service management, and SLA monitoring was developed; it is based on the common object request broker architecture (CORBA), a de facto middleware standard for interface and service definitions.

Industry vendors, furthermore, have already put their proprietary broker products into practice; for instance, Alcatel has implemented the VPN bandwidth broker solution and the dynamic call admission control (CAC) module.

### Resource Provisioning

VPN resource provisioning can be viewed as searching for the cheapest network route or topology that satisfies a subscriber's QoS constraints. By generalizing the whole network into a weighted directed graph, searching for one or multiple subgraphs (i.e., the topology of a QoS-warrant VPN) with the least cost improves network resource utilization. The cost of a route can be defined as a function of the hop count, residual bandwidth, VPN redundancy, and other QoS associated parameters. Figure 5 illustrates an example of VPN resource provisioning where resources for two VPNs are stipulated. With the knowledge of individual VPNs, the problem is modeled as the optimization of an objective function with particular constraints. It can become an NP problem, and certain heuristic approximations (e.g., relaxing certain constraints) will be required to make the problem tractable.

Searching routes for VPNs can be deployed in either a centralized or a distributed way. A typical example of the first case is the service broker that is in charge of admitting, setting up, and tearing down VPN connections. In the previous broker implementations, often network routes are determined without involving any routing intelligence. This is the very reason that centralized databases have to be consulted for tunnel management. However, VPN SPs shall endeavor to accommodate more automation into their network infrastructures, targeting more diverse and flexible services, such as short-lived or highly dynamic VPNs. QoS (constraint-based) routing [7], with routers themselves searching for eligible network routes with sufficient resources to meet the QoS requirements in a distributed manner, can be a potential complement to the VPN router functionality. Its general goals are twofold: every admitted VPN connection has its QoS requirements satisfied; the total cost of all connections on a path is minimized.

VPN resource provisioning and utilization optimization have been undergoing intensive study, taking SLAs, VPN topologies, VPN policies, and available resources into consideration. VServ [8], a comprehensive architecture, presented a set of automated functionalities to support intra- and inter-VPN resource provisioning. It utilizes a VPN description language to trans-

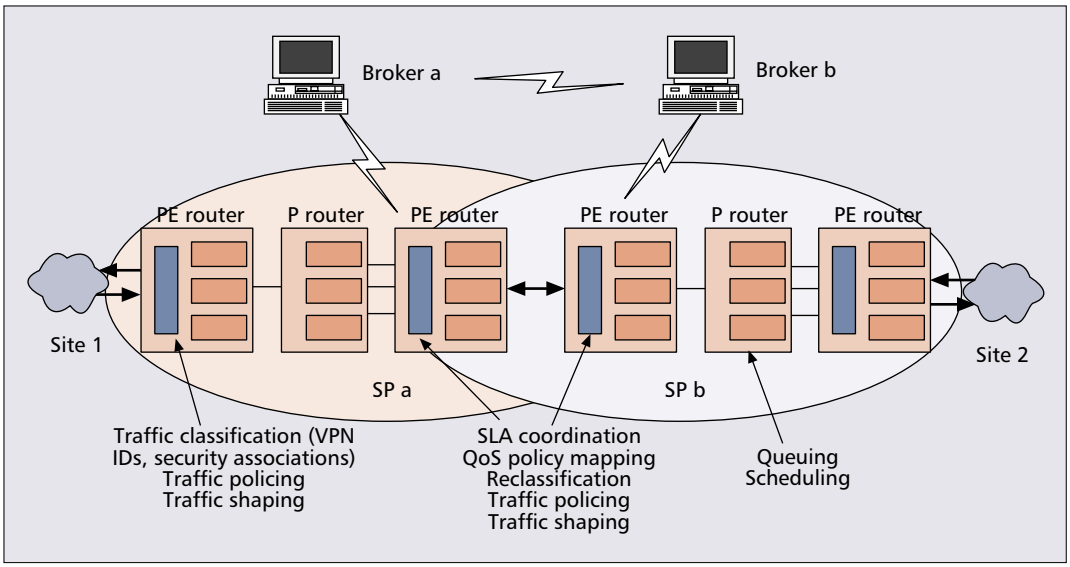


**Figure 6.** *The VPN data flow across multiple SP domains between sites 1 and 2.*

*Looking at a VPN tunnel as just another type of link, many existing QoS mechanisms can be applied to the VPN traffic with VPN-specific parameters; so can the techniques adopted for IntServ and DiffServ.*

late high-level customer criteria into lower-level specifications, constructs a search space according to VPN requirements, and then looks for the optimal topology to complete the resource allocation.

### Ongoing Issues

It will not be rare for a VPN to geographically extend over multiple autonomous domains, or functionally multiple SPs. The VPN QoS issue, as discussed above, has to deal with federation among independent management entities. While a generic management infrastructure is being intensively pursued, how to accommodate different IP QoS architectures (e.g., IntServ and DiffServ) is also on the agenda.

As an example, the service broker can be designed to aggregate per-VPN IntServ messages from enterprise networks at PE routers, so the core network (often a DiffServ domain) can process IntServ messages without any hassle. Aiming to eventually fix the problem, standards bodies have been working on the issue of handling RSVP signaling in a DiffServ domain that is either RSVP-aware or RSVP-unaware. For seamless interoperation, follow-up standardization work, such as mapping IntServ service specifications into DiffServ PHBs, defining a certain functionality for IntServ signaling to deliver aggregate traffic control, and designing a dynamic mechanism for DiffServ resource provisioning, is required (IETF RFC 2998).

Originally as a software module in VPN PE nodes/devices, the concept of a virtual router was proposed to handle control plane operations on a per-VPN basis, thereby restricting the effect of a single misbehaving VPN. Each virtual router is expected to partition individualized service definitions of bandwidth, priority, and security on either a per-subscriber or per-traffic-aggregate basis. Attributes that distinguish VPNs from each other could be topology, duration, and the service they carry. PE routers then maintain separate routing tables and make forwarding decisions for each distinctive VPN. To match packets to the corresponding VPN routing table, PE routers could use a certain tag, such as a VPN ID (IETF RFC 2685), with global significance. Other issues being addressed include the scalability of the number of routing instances, the processing power, and the separation between different VPN routing instances.

A typical industry implementation of the virtual router concept is the IPSX service processing switch family released by CoSine Communications [9]. Although it delivers finer-grained control of the routing topology, the virtual router implementation consumes extra bandwidth and router resources; it may not be cost effective for simple VPN topologies.

## The VPN Node Perspective

VPN SP proprietary routers must act in concert with network-level operations to complete end-to-end QoS enforcement. Therefore, the data plane operations in VPN nodes, which involve shaping, policing, queuing, and scheduling, must be configured according to the QoS parameters determined by the network-level operations. Looking at a VPN tunnel as just another type of link, many existing QoS mechanisms can be applied to VPN traffic with VPN-specific parameters; so can the techniques adopted for IntServ and DiffServ. In association with the VPN data flow illustrated in Fig. 6, VPN router implementations from industry vendors are selectively touched on in the following sections.

### Classification

Classification at the SP edge is the foundation of all other QoS operations in VPN routers. Its purpose is to subject the traffic for future specific treatments; for instance, smaller delay for videoconferencing applications, lower dropping probability for mission-critical services, or faster forwarding for "golden" VPN subscribers. The edge router groups the incoming VPN traffic according to predefined criteria from an SLA and/or a policy server, such as the IP address and application type. It then marks packets, ensuring that the classification will be honored

*In the network core, the SLA conformable VPN traffic classes/aggregates are placed into different queues that are either logically or physically separated. Scheduling strategies then determine the transmission order of enqueued packets, using priorities assigned to the packets by diverse schemes.*

all the way to the other end of the VPN tunnel. An implementation example is Cisco's committed access rate (CAR), one of whose features is to partition the VPN traffic into multiple priority levels or service classes.

### Conditioning

To enforce subscribers following their SLAs, traffic conditioning (shaping/policing) takes place on boundary nodes between VPN subscribers and SPs. As implied by their names, traffic shaping queues the bursty traffic and smooths the stream to a certain degree, and traffic policing simply drops the excess traffic, and lost data have to be retransmitted. Depending on the application type of the VPN traffic, these two mechanisms can be correspondingly deployed. Note that an SP may need to condition the traffic leaving its core too, depending on the SLA negotiation at that boundary. One example of an industry implementation is Cisco's generic traffic shaping (GTS). It regulates the data sending rate and drops the last packet in the queue once the queue is full.

### Queuing and Scheduling

In the network core, the SLA conformable VPN traffic classes/aggregates are placed into different queues that are either logically or physically separated. Scheduling strategies then determine the transmission order of enqueued packets, using priorities assigned to the packets by diverse schemes. A number of scheduling mechanisms adopted by industry vendors, such as Cisco, Alcatel, and Nortel, are selectively listed below. Note that certain industry implementations may be slightly different from their academic counterparts.

**Class-Based Weighted Fair Queuing (CBWFQ):** This mechanism extends weighted fair queuing (WFQ) by supporting user-defined VPN classes (e.g., a mission-critical application class). Traffic belonging to a certain class is then assigned an appropriate bandwidth, buffer length, or drop policy.

**Low-Latency Queuing (LLQ):** Serving packets based on weights, CBWFQ grants no class of packets strict priority. This could introduce delay, especially delay jitter to voice applications. By adding a priority queue to CBWFQ, therefore, LLQ is designed to provide explicit priority to delay-sensitive voice applications.

**Hierarchical Class-Based Queuing (HCBQ):** HCBQ divides the traffic into classes and subclasses as well. One subclass can take the bandwidth from other subclasses of the same class. Different scheduling methods can be accordingly adopted.

**Modified Deficit Round Robin (MDRR):** Regular deficit round robin (DRR) provides every queue equal scheduling opportunities in a round robin way. As an approximation of LLQ, MDRR has one of its queues defined as the priority queue, thereby providing low delay and jitter to delay-sensitive applications such as voice over IP (VoIP).

### Congestion Management

Congestion avoidance recognizes and acts on congestion to relieve or eliminate its negative effects on QoS guarantees. Among a variety of strategies, two classical congestion avoidance mechanisms adopted by leading VPN industry vendors are briefly described here.

**Random Early Detection (RED):** The average queue size is calculated to compare with two thresholds: one minimum queue size and one maximum queue size. Below the minimum limit, no packet is marked; above the maximum threshold, every packet is marked. In between these two, packets are marked with a probability that is a function of the average queue size. The packets are then randomly dropped at the moment of congestion, attempting to avoid global synchronization when multiple TCP streams change their rates [10].

**Weighted RED (WRED):** Combining the RED mechanism and different classification scenarios, it provides the preferential traffic handling and thus differentiated performance for service classes by selectively discarding lower-priority traffic at the moment of congestion. As in RED, network engineers have the flexibility to configure the minimum and maximum queue length thresholds as well as drop probabilities of each service class.

## MPLS-Based VPNs

Envisioning a backbone that supports QoS, MPLS entails significant changes in existing IP network architectures. As a hybrid of the layers 3 and 2 structures, it forwards layer 3 packets like a layer 2 switch, thereby taking advantage of layer 3 routing intelligence and layer 2 fast forwarding capabilities.

As one of the technical approaches for IP-based VPN implementations, MPLS is more than another innovative paradigm due to its unique characteristics. First, MPLS-enabled routers or switches attach labels to packets according to the forwarding equivalence class (FEC), and then forward packets based on the MPLS label instead of conventional IP address lookup. Second, instead of routing the packets through the network, MPLS passes on packets to the destination by swapping or peeling away their labels hop by hop. Third, forwarding packets based on labels and distributing the labels with routing protocols, MPLS-enabled devices separate these two functionalities. It introduces more implementation flexibility than IP routers that couple the forwarding decisions with the generation of routing information [11].

Although MPLS does address the QoS issue, its original motivation was more on improving Internet scalability through better traffic engineering. Nevertheless, this has not hindered MPLS-based VPNs from phenomenally gaining momentum. For instance, in June 2002 AT&T announced MPLS-based IP VPN services in Australia. The QoS issue of MPLS VPNs, however, needs to be investigated from another, if not totally different, angle, and thus is beyond the scope of this article. As a matter of fact, since SPs will probably prefer retaining existing enterprise subscribers and gradually attracting new ones, both types of VPNs will exist alongside one another in years to come. To furnish a general rather than an

| | |
|---|---|
| Packet forwarding speed | MPLS VPNs tend to have a faster forwarding speed than IP VPNs, by avoiding the IP header lookup and instead using the information in MPLS labels. |
| Traffic engineering signaling | Except for the centralized management architecture, IP VPN implementations work on the adaptation of IntServ signaling in the DiffServ domain; a candidate signaling for MPLS VPNs, RSVP-TE is under development by the IETF MPLS working group. Note that the MPLS working group has decided to stop implementing constraint routing LDP (CR-LDP). |
| Scalability | PE routers of IP VPNs maintain a full mesh of tunnels among all sites of a particular VPN, and P routers hold the information for all accommodated VPNs. However, no single router in the MPLS VPN backbone has to maintain the routing information for all supported VPNs [12]. Besides, by using route reflectors in MPLS VPNs, the scalability hazard of maintaining a full mesh of intersite VPN connectivity is eliminated. |
| IP address space | IP VPN traffic needs a globally unique IP address to cross the IP core, whereas MPLS VPN subscribers can use a globally unique address space, private IP address space, or even overlapping address spaces. |
| Security | IP VPNs can support strict information confidentiality by configuring IPSec security associations in PE routers among VPN sites. MPLS VPNs, by themselves, provide equivalent security to layer 2 VPNs. There is no direct support for authentication and confidentiality. Besides SP PE routes, the intermediate routers belonging to differentadministrative domains must be trusted. |
| Multiprovider environment | While a notable amount of work on interdomain federation has been done for IP VPNs, the same issue in MPLS VPNs has not yet found a firm basis due to the lack of interoperable standards. |

**Table 1.** *A comparison of IP VPNs and MPLS VPNs.*

exhaustive comparison, the differences between IP VPNs addressed in this article and MPLS VPNs are listed in Table 1.

## Conclusions

Although several technologies for delivering IP VPNs are still in the "cloud," this booming service is being adapted and gaining ground at a surprising speed as standard bodies, industry vendors, and research communities push each other ahead. This article presents a QoS guarantee framework for IP VPNs. QoS operations from the VPN network perspective determine the QoS configuration parameters; routers at the node level are then configured in concert to enforce end-to-end QoS. Diverse VPN QoS enabling technologies as well as development progress in recent research and engineering work have been addressed to further complete the whole picture of the IP VPN QoS issue.


### Acknowledgment

The authors would like to acknowledge the anonymous referees' constructive suggestions to help improve the quality of this article.


## References


[1] "A Primer for Implementing a Cisco Virtual Private Network," Cisco ref. guide, Aug. 2000, http://www.cisco.com/warp/public/cc/so/neso/vpn/vpne/vpn21_rg.htm.

[2] T. Braun, M. Guenter, and I. Khalil, "Management of Quality of Service Enabled VPNs," *IEEE Commun. Mag.*, vol. 39, no. 5, May 2001, pp. 90–98.

[3] Charging and Accounting Technologies for the Internet (CATI) project, http://www.tik.ee.ethz.ch/~cati/home.html

[4] Q. Kong, I. Rose, and D. Cameron, "Towards Technology Independent and Automated Service Activation and Provisioning," *Proc. Net. Ops. and Mgmt. Symp.*, Florence, Italy, Apr. 2002, pp. 931–33.

[5] TINA Consortium, http://www.tinac.com

[6] F. De Turck *et al.*, "Design and Implementation of a Generic Connection Management and Service Level Agreement Monitoring Platform Supporting the Virtual Private Network Service," *Proc. IFIP/IEEE Int'l. Symp. Integ. Net. Mgmt*, Seattle, WA, May 2001, pp. 153–66.

[7] S. Chen and K. Nahrstedt, "An Overview of Quality-of-Service Routing for the Next Generation High-Speed Networks: Problems and Solutions," *IEEE Network*, vol. 12, no. 6, Nov./Dec. 1998, pp. 64–79.

[8] R. Isaacs and I. Leslie, "Support for Resource-Assured and Dynamic Virtual Private Networks," *JSAC*, vol. 19, no. 3, Mar. 2001, pp. 460–72.

[9] "Managed Network-based Site-to-site IP VPN Service," app. note, CoSine Commun., Aug. 2002, http://www.cosinecom.com/vitrualipservices

[10] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Trans. Net.*, vol. 1, no. 4, Aug. 1993, pp. 397–413. http://www.icir.org/floyd/red.html.

[11] "Nest Generation VPNs," white paper, Lucent Technologies, Nov. 2001, http://www.lucent.com/knowledge/documentdetail

[12] "RFC 2547bis: BGP/MPLS VPN Fundamentals," white paper, Juniper Networks, 2001, http://www.juniper.net/techcenter/techpapers/200012.html


## Biographies


JINGDI ZENG [M] received B.E. and M.E. degrees in communications engineering and computer applications from Hunan University, P.R. China, in 1995 and 1998, respectively. She is currently pursuing her doctoral studies at NJIT. Her current research interests include network resource management, virtual private networks, and other issues related to Internet QoS. She is a member of Alpha Epsilon Lambda.

NIRWAN ANSARI (nirwan.ansari@njit.edu) received a B.S.E.E. (summa cum laude), an M.S.E.E., and a Ph.D. from New Jersey Institute of Technology (NJIT), University of Michigan, and Purdue University in 1982, 1983, and 1988, respectively. He joined the Department of Electrical and Computer Engineering, NJIT, in 1988, and has been a professor since 1997. He is a technical editor of *IEEE Communications Magazine* as well as the *Journal of Computing and Information Technology*; is General Chair of ITRE2003; was instrumental while serving as Chapter Chair in rejuvenating the North Jersey Chapter of IEEE ComSoc, which received the 1996 Chapter of the Year Award, served as Chair of the IEEE North Jersey Section and on the IEEE Region 1 Board of Governors, 2001–2002, and currently serves on various IEEE committees. He was the 1998 recipient of the NJIT Excellence Teaching Award in Graduate Instruction, and a 1999 IEEE Region 1 Award. His current research focuses on various aspects of high-speed networks and multimedia communications. He authored with E.S.H. Hou *Computational Intelligence for Optimization* (1997, translated into Chinese in 2000), and edited with B. Yuhas *Neural Networks in Telecommunications* (1994), both published by Kluwer. He has also contributed over 50 refereed journal articles and over 15 book chapters.