CERIAS Tech Report 2000-06

# Security Considerations in Quality of Service Architectures

**Stephanie A. Miller**
Center for Education and Research in
Information Assurance and Security
Purdue University, West Lafayette, IN 47907

# Contents

# 1 INTRODUCTION

## 1.1 Motivation and Goals

The notions of Quality of Service (QoS) and Security have existed for many years as separate areas of research. Security is a topic that has been studied extensively in academia and industry throughout history. Quality of Service is a topic of great importance as user demands for network performance and utilization grows faster than available bandwidth. In modern networks both security and QoS are topics that play a significant role in deployment. Each area of interest contribute many advantages to the overall design of emerging network architectures. The integration of QoS and security, however, brings about new issues to consider and further possibilities for the evolution of network protocols. Additionally, there is a strong business case for integration of the two areas from the onset. Building a safe and efficient network architecture satisfies the needs of everybody in an organization.

Unfortunately, there is very little literature in the research community to address the need for merging security with QoS protocols. In this paper I provide overview material on each topic based on current research. My goal for delving into this area centers around the need for increased awareness in the both QoS and security communities about combining the two topics successfully. The contribution of this survey paper is to highlight areas of weakness in current QoS schemes as well as to report on the current strategies for protection. The intent of this paper is to begin to formulate where the challenges for the protocols exist. We do not have all of the answers, but we do raise important areas for consideration in the studey of evolving QoS standards.

Using the knowledge gained from addressing security concerns in other network protocols, such as TCP/IP, we can begin to understand where security fits into the QoS framework.

## 1.2 Basic Concepts

### 1.2.1 Quality of Service

The current state of the Internet is to offer best effort delivery of information. The goal of QoS is to improve the predictability of packet delivery. It provides a means to ensure better network service to selected traffic. This can be

accomplished in varying degress based on traffic characteristics and cost. Some of the specific benefits of QoS [QoS99a] include:

- Support for dedicated bandwidth
- Improved loss characteristics
- Avoidance and Management of network congestion
- Shaping of network traffic
- Setting traffic priorities across the network

There are many ways to think about the term Quality of Service. The phrases has become overloaded in its usage. In a generic sense, QoS adds predictability to a network [Bak99b]. Predictability is a desirable network feature. This is especially true for real-time applications that often need to deliver time sensitive data over the network.

Quality of Service is at times a controversial subject. Arguments have been made that increasing bandwidth will solve the problem of congestion and improve service. While increasing bandwidth does accommodate those real-time application needs in the short term, it is not the only answer. Bandwidth does have limits and costs. Additionally, delivery delays can still occur on an unloaded network. [QoS99d]

The focus for this paper is that QoS is a viable (and necessary) mechanism for improving network performance using existing bandwidth. As such, QoS architectures should be designed and implemented with security in mind.

### 1.2.2 Security

Information security is an area of growing significance in the technology community and beyond. As we have transitioned from an industrial to an information based society, new areas of crime have emerged using the newly developed technology. Considerable money and time is being spent by organizations of all varieties and sizes to ensure they are protecting their most valued asset, information.

The fundamental goals of information security rely on trusting a computer system to preserve and protect its data and resources. Terms like security protection and privacy often carry different meanings to different people.

However the underlying definition is that a secure system should be dependable and behave as expected.[SG96] Some of the most common concerns in security protection are confidentiality, integrity, and availability. However additional requirements in authentication, access control, authorization, and non-repudiation are also significant when implementing security controls. All project developers must account for the unique blend of requirements they must manage to adequately protect valuable resources and information in a computer system.

Like QoS, security is a very broadly defined term which contains many specialized disciplines. There are a variety of topics to discuss at a high level when dealing with information protection. Some of those fundamental principles in security include [LS87]:

- Confidentiality: The computer security characteristic that ensures that individuals are given access to computer resources based on security clearance and need-to-know. This characteristic protects against compromise and inadvertent disclosure.

- Integrity: The computer security characteristic which ensures that computer resources operate correctly and that the data in the databases is correct. THis characteristic protects against deliberate or inadvertent unauthorized manipulation of the system.

- Availability: The characteristic that ensures the computer resources will be available to authorized users when they need them. This characteristic protects against denial of service.

- Authentication: The act of identifying or verifying the eligibility of a station, originator or individual to access specific categories of information.

- Non-Repudiation: A property that prevents denial by one of the entities involved in a communication of having participated in all of part of the communication.

- Authorization: The granting to a user, a program or a process the right of access.

- Audit: To conduct the independent review and examination of system records and activities in order to test for adequacy of system controls,

to ensure compliance with established policy and operational procedures, and to recommend any indicated changes in controls, policy or procedures.

Criminals will always try to find ways to cheat a system. Implementing controls to address these security requirements offers protection against such attempts. Security is not just a technology problem. Process issues related to audit, management, and policy along with human issues including awareness and accountability are also important to understand when developing a security framework.

## 1.3 Document Organization

The remainder of this document will be organized as follows. Section 2 will focus on Quality of Service and the protocols that exist in that area. That section will provide a general background about these protocols and how they extend existing network technologies. That section will also provide examples of how end-to-end QoS can be accomplished. Section 3 will delve into specific security issues that arise within network communications. That section will address the challenges of establishing trust in communication and how those challenges have been handled in the past. Section 4 will offer ideas on how an attacker may be able to compromise QoS mechanisms. Section 5 will highlight existing efforts to unite security with QoS protocols. Finally, conclusions and thoughts on future work will be given in Section 6.

# 2 Overview: Quality of Service

Quality of Service is not a property that can be precisely measured. Some applications require more stringent QoS guarantees than others. The fundamental QoS goal of providing some degree of assurance for consistent network delivery can be met in varying degrees to match the necessary application requirements. Network topology, policy, and the applications dictate which blend of QoS is most appropriate for an individual flow of data or for an aggregate. [QoS99d]

There are several parameters to analyze when devising a QoS model. [Joh99, Hus98]

- Delay/Latency: Delay in a transmission path or in a device within a transmission path. In a router, latency is the amount of time between when a data packet is received and when it is retransmitted.
- Jitter: The distortion of a signal as it is propogated through the network, where the signal varies from its original reference timing and packets do not arrive at its destination in consecutive order or on a timely basis, i.e. they vary in latency.
- Throughput: In data transmission, throughput is the amount of data moved from one place to another in a given time period.
- Error Rate: The rate in which packets are dropped, get lost, or become corrupted while traversing the network.

Each QoS protocol and standard that is defined must address these topics and allow for them to be tuned for a desirable configuration. Absolute guarantees in the Internet are unrealistic. QoS aims to provide a means to manage expectations that can be defined using these parameters.

## 2.1 The Dominant Protocols

When the research for this survey paper began, I expected to find that there would be just a single protocol that would dominate the network. It seemed as if the different protocols designed to offer QoS were in competition and mutually exclusive. This was a poor assumption. In reality, Quality of Service entails blending different varieties of mechanisms to provide an overall

end-to-end QoS architecture. Different locations in the network, such as the core, the edges, or the endpoints, have different requirements and computing resources. The result is that merging the different schemes for improving QoS at the appropriate locations will offer the best solution. In particular the Differentiated Services Protocol is designed to work best within the core of a network, and Integrated Services can be utilized at the edges. Along the data path mappings can occur to combine the functionality of each along with faster routing schemes into a true end-to-end quality of service.

Quality of Service can be considered in these ways:

- Best Effort (No guaranteed quality)
- Resource Reservation (Integrated Services)
- Prioritization (Differentiated Services)

There are many protocols and algorithms that exist to offer the above functionality. Some examples include ReSerVation Protocol (RSVP), Differentiated Services (DiffServ), Multi Protocol Labeling Switching (MPLS), along with queuing algorithms such as Weighted Fair Queuing, and Random Early Detection.

An unfortunate disadvantage of QoS is that any non-QoS enabled device within a network domain imposes unpredictable behavior on the traffic. This unpredictability can lead to a loss of end-to-end QoS.

This section will offer specific information on DiffServ, MPLS, and RSVP. The area of queue management, while important to QoS in general, will not be covered. Those algorithms are generally transparent to applications and are not explicitly considered to be QoS protocols. The section closes by demonstrating how an architecture can be designed that takes advantage of the best properties of each protocol.

### 2.1.1 Best Effort

The model for quality of service in use on the Internet has been Best Effort Service. Best Effort is synonmous with the absense of QoS. There are no guarantees offered for delivery characteristics. Applications can send data whenever they must, in any quantity, and without requesting permission or first informing the network. [cis99c] This area of discussion is provided to give context for what we have today and where QoS can help. Best Effort

service has many problems. The network, in light of congestion, may drop any packet or cause it to be delayed longer than expected. This level of service has worked well for a wide range of network applications, including email and file tranfer. However, with the increase in demand for real-time data better network guarantees are necessary.

### 2.1.2 Integrated Services

Integrated services is a framework that supports controlled sharing of network links. It does so by allowing an application to request specific guarantees from the network before sending data packets. The request is made using a signaling protocol that allows the application to tell the network it's traffic profile and desired service. The network takes into account network availability, admission control, and other parameters to make the determination to grant the application's request. Once the network confirms the request, the application is free to start sending data within that profile. Integrated Services allows for single, robust, integrated-service communications infrastructure that can support the transport of audio, video, real-time, and classical data traffic. [QoS99a]

Integrated Services can be provided on a per-flow basis according to application requests. A flow is defined as a data stream between an application sender and receiver. A flow can be identified as an individual, uni-directional, data stream between two applications (sender and receiver), uniquely identified by a 5-tuple (transport protocol, source address, source port number, destination address, and destination port number). [QoS99d]

There are two service categories in Integrated Services. These categories are made possible by intelligent queueing mechanisms in network devices along with the other Integrated Services functions of classification and admission control.

- Guaranteed Service: Provide an assured level of bandwidth, a firm end-to-end delay bound, and no queuing loss for conforming packets of a data flow. This service is intended for applications with stringent real-time delivery requirements, which are intolerant of any datagram arriving after their playback time. [Whi97]

- Controlled Load: A commitment to offer the flow a service equivalent to that seen by a best-effort flow on a lightly loaded network with no noticeable deterioration of service as the network load increases. This

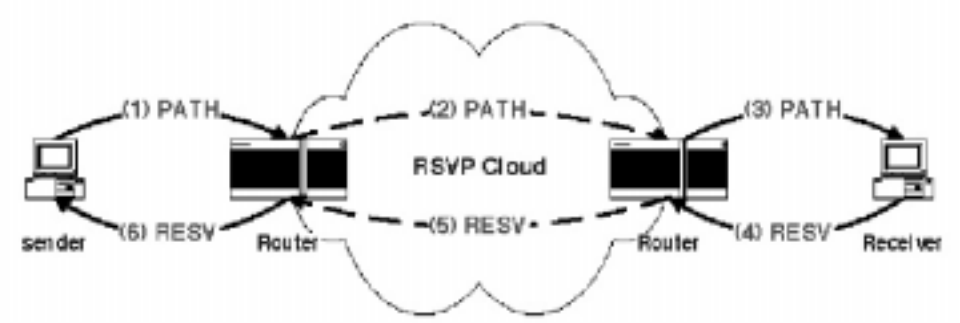


Figure 1: RSVP Signaling [XN99]

entails offering the application low delay and high throughput even during times of congestion. This class of service is intended for applications that can tolerate a certain amount of loss and delay provided it is kept to a reasonable level. [Whi97]

**RSVP: Resource Reservation Protocol** The primary signaling mechanism associated with Integrated Services is the Resource Reservation Protocol. RSVP is not a routing protocol. RSVP allows for receiver-controlled reservation requests. The protocol supports Integrated Services by communicating soft state information within each router along a flow's path, therefore the reservation must be periodically updated. Figure 1 shows the signaling process. The sender application sends a PATH message to the receiver identifying its traffic characteristics (bandwidth bounds, delay jitter, etc.) Each router downstream from the sender to the receiver stores "path-state" for this flow [QoS99a] and sends the PATH message to the next hop determined by the routing protocol. The receiver then generates a RESV message to be sent along the same path back to the sender. This RESV contains the requested QoS level to be allocated by the routers along with a description of the packets to be processed. Each RSVP router along the path then uses admission control and resource availability information to determine whether to accept the request. An error message is sent to the receiver if any router on the path denies the request. Otherwise a confirmation message is sent to the receiver once all routers accept.

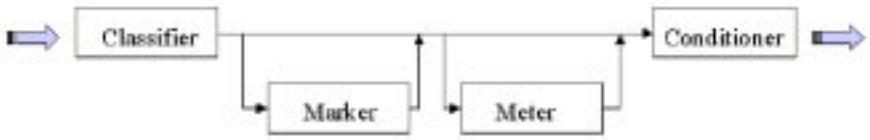


Figure 2: Differentiated Services Traffic Conditioning Functionality

Further details on RSVP and its use with Integrated Services can be found in [Whi97, She94]. Similar signaling protocols have been designed but RSVP has the strongest backing for use with Integrated Services.

### 2.1.3 Differentiated Services

As noted in [Zha98], the reliance of IntServ on per-flow state and per-flow processing is an impediment to its deployment in the Internet at large, and in particular in large carrier networks. Differentiated services (DiffServ) is the scalable answer to this QoS deployment barrier. DiffServ moves the complexity out of each router on a path and into the edge nodes. This leaves the core of the network unencumbered from heavy QoS state maintenance.

Differentiated services works by aggregating transmission flows and defining one or many per-hop behaviors (PHB) to be associated with the aggregates in the network core. At the edge of the network, traffic is conditioned using the actions of classification, marking, shaping, and policing as depicted in figure 2. (For specific descriptions of these functions, reference [Wei98, cis99c]) The conditioning enforced at these boundary nodes is based on Traffic Conditioning Agreements (TCAs) that are established between adjacent DiffServ domains.

The DiffServ marking for each packet is maintained in what used to be the 8 bit IP TOS (Type of Service) header field, now called the Differentiated Services Code Point (DSCP). Figure 3 shows the usage of this field. The DSCP will be referenced by DiffServ capable devices within the core of the network to determine the forwarding treatment to give that class of packets.

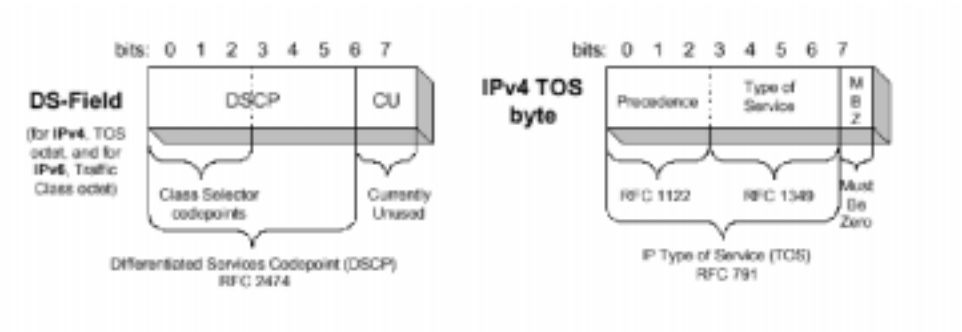


Figure 3: Differentiated Services Code Points redefine the IPv4 TOS byte [QoS99d]

Forwarding behavior within the network core is defined by the PHBs. A PHB is a description of the externally observable forwarding behavior of a DS node applied to a particular DS behavior aggregrate [Wei98]. The PHB is the means by which a node allocates resources to behavior aggregates and is implemented through buffer management and packet scheduling mechanisms.

There are two globally recognized PHBs available. Customized PHBs can also be defined for local administrative domains.

- Expedited Forwarding (EF): Has a single codepint (DiffServ value). EF minimized delay and jitter and provides the highest level of aggregate quality of service. Any traffic that exceeds the traffic profile (which is defined by local policy) is discarded. [QoS99d]

- Assured Forwarding (AF): Has four classes and three drop-precedences within each class (a total of twelve codepoints). Excess AF traffic is not delivered with as high probability as the traffic "within profile", which means it may be demoted but not necessarily dropped. [QoS99d]

A Service Level Agreement (SLA) is a contract between a service provider and customer defining provider responsibilities in terms of network levels and times of availability, method of measurement, consequences if service levels aren't met or the defined traffic levels are exceeded by the customer. The

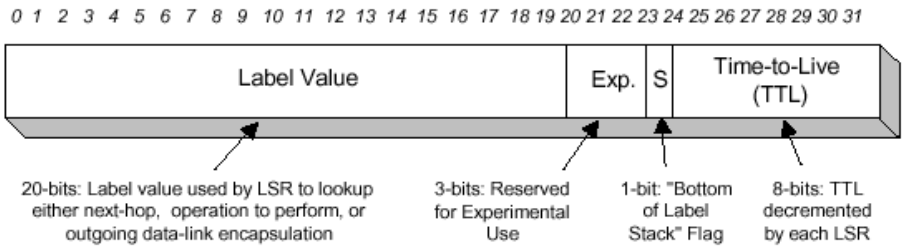


Figure 4: MPLS Header Layout

customer may be a user organization or another domain. [Joh99] SLAs can be static or dynamic.

In addition to traffic conditioning, and PHBs, the Differentiated Services framework includes a component called the Bandwidth Broker (BB), which is described in section 5.1.4 of this paper.

### 2.1.4 Multiprotocol Label Switching (MPLS)

MPLS is a protocol that is defined for routers only. Communicating host endpoints do not need to be MPLS aware. A router that supports the MPLS protocol is called a Label Switching Router (LSR). The MPLS technique makes an IP routed network more connection oriented in nature where traffic is routed along a labeled path in the topology. [Joh99]

The goal of MPLS is to speed up the routing function by using table lookup (indexing) based on a label instead of address matching as traditional routing protocols do. The technique is to assign each packet a label upon entry to a MPLS domain. While MPLS is often considered to be more of a traffic engineering protocol than QoS, the label defines the QoS level that a packet receives. The benefit of MPLS is that once a label has been assigned to a packet, the overhead of complex address matching algorithms is removed. The packet is forwarded to the next hop simply based on the label it contains.

The MPLS header is encapsulated between the link layer header and the network layer header. [XN99] The label itself is 20-bits long and can be stacked. Figure 4 illustrates an MPLS header.

There is some overhead with MPLS. It relies on a means of communicating

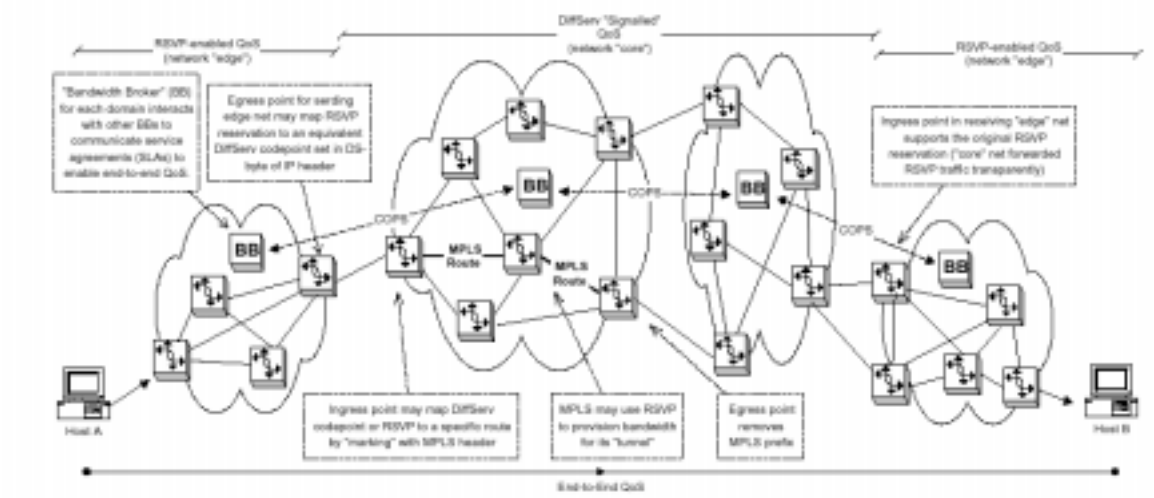


Figure 5: Possible use of different QoS technologies for an End-to-End Architecture

label mapping to forwarding behavior. Each LSR must manage an agreement of how the label is defined between peers. A benefit of MPLS is that a Label Switched Path (LSP) can be used for tunneling since a path can be completely determined by the label assigned by the ingress LSR. [XN99] This ability is expected to reduce the cost of implementing VPN technology.

## 2.2 A Sample Architecture

Figure 4 is taken from [QoS99d]. This figure highlights the complementary nature of the many QoS protocols. The goal is to use Differentiated Services within the core of a network where it can be put to good use as a lightweight prioritization protocol, while using Integrated Services on the stub networks to offer finer granularity to applications in requesting their service needs.

DiffServ is a perfect complement to RSVP as the combination can enable end-to-end QoS. Border routers at backbone ingress points can map RSVP reservations to a class of service indicated by a DS-byte (details on the mapping can be found in [Zha98]). At the backbone egress point, the RSVP

provisioning may be honored again, to the final destination. [QoS99d]

## 2.3 Current State of Affairs

Quality of Service is a topic that is building momentum. Many vendors are developing products with QoS functionality builtin. Cisco Systems, Inc. has developed a significant amount of QoS technology into their IOS system. Similarily many smaller companies with a focus on QoS are forming. [Ray99] surveys the product offerings available in today's market. Deployment of QoS capable devices will be a gradual process. In the near term we should begin to see increasing use of the protocols described, and likewise heightened awareness for the ensuing challenges for securing the differing levels of service.

Standards bodies are actively investigating the protocols and issues related to QoS. Many of the papers referenced in this survey paper were written within the last 18 months. The result of such activity is that QoS building blocks are maturing every day. The task of solving the end-to-end QoS equation is yet another milestone for QoS engineers. End-to-end is defined as a system that enforces consistent Quality of Service policies throughout a network. [IPH99a] Many discussions within the QoS community have taken place concerning this goal. Several documents have also been prepared to propose likely end-to-end QoS architectures using specific protocols. [QoS99d, Zha98] are among the proposals.

# 3 Overview: Security in Network Protocols

## 3.1 Pitfalls in TCP/IP

The Transport Control Protocol (TCP) and Internet Protocol (IP) handle the majority of traffic on the Internet today. These protocols are over 30 years old and have been studied extensively. [Cla88]

The TCP/IP protocols can protect against some threats. Checksums protect against packet header modification. For TCP, sequence numbers protect against lost and duplicated packets, and other measures protect against reuse of packets. [Sum97]

Unfortunately fundamental problems in the TCP/IP architecture are identified in [Bel89]. Often packet authentication relies on the source IP address and other information that can be easily manipulated. Many Internet attacks are successful because they can "spoof" an IP address. This effectively allows an attacker to masquerade as another system. Using IP spoofing along with the source routing feature of IP provides an attacker with the weapons to launch potent and stealthy attacks.

TCP, like IP, suffers from problems inherent to its very design. The technique of using sequence numbers to provide reliable communication can also be abused to attack a system. "Syn flooding" is a denial of service attack that works by opening a flood of TCP connections with the victim host and not completing the 3-way handshake. This attack depletes the network resources for that host and makes it unresponsive to legitimate connections. Sequence number guessing is another well-known attack. This strategy is to predict the sequence number to be used during a communication and inject forged packets into the connection that will be accepted by the target.

The underlying difficulty for establishing secure protocols is the notion of trust. Most authentication schemes rely on some ability to trust the entities in communication and the information being provided for the authentication. In the human world trust is established between people on a daily basis. It is very difficult to translate this notation to the machine world. A stated in [Lam92], our formal understanding of trust in distributed systems is at best inadequate.

## 3.2 A Collection of Common Attacks

There are a variety of classifications that exist for grouping attacks and vulnerabilities. We define a vulnerability as some weakness of a system that could allow security to be violated. An attack is a circumstance or event that could cause harm by violating security. An attack often exploits a vulnerability. [Sum97]

The focus for our study is in the type of attacks that can be considered in a QoS model. This section provides general background on the issues related to network threats. The specific attacks associated with QoS will be identified later in this survey paper.

- Network Denial of Service: Denial of Service is an attack in which legitimate users are prevented from using the network. According to [SG96] there are several methods for causing a network DoS. These include
  - service overloading: Occurs when floods of network requests are made to a server daemon on a single computer.
  - message flooding: Occurs when a user slows down the processing of a system on the network to prevent the system from processing its normal workload, by "flooding" the machine with network messages addressed to it.
  - signal grounding: Physical methods that can be used to disable a network such as grounding the signal on a network cable, introducing some other signal, or removing an Ethernet terminator.
  - clogging: Occurs when an attacker uses up the limit of partially open connections. Particularily harmful to TCP.
- Session Hijacking: The principle of session hijacking is to seige control of a network connection. Once an attacker has successfully hijacked the connection he is able to supply user commands on behalf of the legitimate user.
- Masquerading: Identity theft is the misuse of another user's identity with the objective of taking actions permitted to the owner of the identity. Authentication and access control services are useless since an established session has been hijacked. [Sum97]

- Information Leakage: Failures in the protocol or implementation can lead to an attacker gathering information about a session that she would not otherwise be able to deduce.

- Unauthorized Resource Use: Compromise of any devise on a network constitutes unauthorized resource use.

## 3.3 Interoperability Considerations

Protocols are the building blocks for networks. They are defined and implemented as modules that will be used together to create an operational environment as we have learned to be the case with QoS design.

Individually each protocol poses concerns for security weaknesses. When a full network architecture has been built, the interaction between these protocols can lead to new areas for concern. This concern is addressed in [AF] where examples of multiprotocol vulnerabilities are given. The author shows that individually an authentication protocol can be demonstrated to be correct. But when used in conjunction with another protocol during the exchange of messages in a Public Key Infrastructure, new attacks can be introduced. PKI is just one area of many where multiprotocol vulnerabilities should be studied.

The same ideas can be applied to QoS protocols that are being designed to work together for an end-to-end QoS architecture. Careful attention needs to be given to how movement of data packets and control messages through IntServ and DiffServ devices happens and the possible ambiguities that can arise in the communication structures.

# 4 Possible Attacks to Consider in QoS

## 4.1 Setting the Stage

In order to visualize the attacks about to be discussed, we will consider a practical QoS environment. For the purpose of example, consider a global corporate network with sites in every major geography worldwide. This network carries the standard email and http traffic patterns. Bandwidth is abundant, however congestion still arises during various time intervals. During these times of congestion the transaction software for the sales unit, which is QoS aware, is given preferential treatment for packet delivery. The QoS aware network is configured to give better service to mission-critical, delay-sensitive applications over email communication and http traffic.

There are a number of VPN (Virtual Private Network) [vdM98] connections attaching the corporate network to many supplier networks. The connections between this corporation and the various suppliers include routes through a couple of large ISP (Internet Service Provider) [Joh99] backbones.

While this is a basic network description, it serves as a model for thinking of the high level purpose of QoS. The goal of this section is to shift our thinking about disjoint protocols toward a more complete and connected architecture with all of the complexities being handled by today's corporations.

Given this context, we will now propose a subset of attack scenarios that focus specifically on QoS enabled networks.

## 4.2 QoS Attack Scenarios

QoS, as presented in this paper, builds on the functionality of the Internet Protocol (IP). Attacks that have been proven to work against IP will undoubtedly work within a QoS enabled environment. Additionally, many new areas of concern arise in light of the features introduced by QoS. This section identifies weaknesses in the QoS scheme and methods to attack the protocols.

In October 1999, an Internet draft was released entitled *Security Issues for Differentiated Service Framework*. [Gon99] That draft addresses many similar concerns and is recommended to support the proposed attacks in this section.

### 4.2.1 Theft of Service

The biggest concern when designing technology that offers any degree of preferential treatment is theft. By manipulating the QoS signals and codepoints, an attacker may be able to obtain better service than was purchased. This crime is one that entails no cost to the perpetrator, unlike later attacks to be discussed.

An improvement in service is realized when the QoS parameters are set to match an existing service agreement. According to RFC 2475 for Differentiated Services [Wei98], the mapping of network traffic to the specific behaviors that result in different (better or worse) service is indicated primarily by the DS field. An adversary may be able to obtain better service by modifying the DS field to codepoints indicating behaviors used for enhanced services or by injecting packets with the DS field set to such codepoints.

This same line of thought can be applied to theft of service in the RSVP context. RSVP uses two types of messages, PATH and RESV, when making a reservation. The PATH message contains information about the senders traffic characteristics and it is also used to set up the reverse routing path. A RESV message flows in the reverse direction making the reservation request. An attacker can take advantage of the RSVP protocol by claiming false traffic characteristics and forcing unnecessary reservations to be made.

The bottom line is that since enabling QoS on an IP network effectively means that some users will get better network service than others, it creates some incentive to steal. Some users inevitable want the better service, but they don't want to have to pay the (likely) higher costs involved. Hence, there is a need to authenticate those that request the better service levels. [QoS99b]

Additional defenses to this attack include good traffic conditioning at DiffServ boundary nodes (ensuring codepoints conform to the applicable TCA(s) and the domain's service provisioning policy) along with security and integrity of the network infrastructure within a DiffServ domain. [Wei98]

### 4.2.2 Denial or Degregation of Service

Denial of service in QoS is a method of attack that removes availability of resources from providing the QoS that was offered. Degradation is decreasing the level of QoS to a noticably poor level (i.e. longer delay, bigger jitter, high drop rate), but not completely destroying resource availability.

Excessive theft of service can eventually lead to denial or degradation of service. When malicious users are receiving higher QoS than they pay for, they deplete the resources available to forward other traffic streams.

There are many other opportunities for an attacker to cause a denial or degradation of service in the network. The following list highlights a few techniques that cause concern.

- A compromised router can be reconfigured to drop or add deliberate delay to packets requiring high QoS in the DiffServ domain. Likewise if RSVP state is being maintained in a router, that state information can be corrupted or erased by the attacker. The consequence is even greater if a border router is compromised as those devices handle much higher traffic volume than interior routers. [Gon99]

- A device along a transmission path may be under the control of an attacker to remark packets with undefined codepoints or with codepoints indicating a lower level of QoS than they should receive.

- A flooding tool can be used within a domain to generate so much traffic that congestion is unavoidable. Such artificial congestion can have adverse affects on many QoS schemes.

### 4.2.3 Session Hijacking and Identity Masquerading

Session hijacking is a threat to IP networks. Identity masquerading is a related problem. Given the design principle of QoS protocols to build on the IP layer of the network stack, we identify session hijacking as a threat for QoS as well. The technique of hijacking a session in IP applies in the QoS architecture with the additional step of marking the appropriate QoS level in the forged packets. The concern is heightened in QoS if connections between policy entities become hijacked. In this manner an attacker would be able to issue decisions on behalf of the policy decision points in the network. Also worth noting is that this type of attack is feasible even when using the strongest authentication techniques available.

### 4.2.4 Other Issues

The majority of problems identified in this paper and others are based on analysis of protocol design as described in the RFCs. There are other chal-

lenges awaiting security and network professionals as QoS goes into deployment. Many exploits that have been waged in the past prey on problems in the implementation of the protocols. Buffer overflows result from poor boundary and parameter checking by programmers. Subtleties in the protocol specifications (the RFCs) can lead to differing interpretations by implementers leading to inconsistency in software. Such inconsistencies are fertile ground for exploitation.

An issue raised in Section 3.4 is that of protocol interoperability. QoS, as we have learned, is being designed as a set of modules to be used together to form an end-to-end QoS framework. Security flaws can arise when message formats are tangled in such a way that slight modification of one message can lead to a breach of information another. This is an area requiring further study in QoS architecture designs.

Finally, protection needs to be considered for the technology that supports QoS. Accounting data must be secured so that fair billing is achievable. As we will learn in the next section, mechanisms used to create and distribute policy and resource allocation must also be safeguarded. Less obvious but also of importance, the statistical data used to maintain audit capabilities must be protected. [ea99]

### 4.2.5 Summary

Each of these attacks are realistic if one considers the interaction between compliant and non-compliant QoS domains. If an attacker is positioned within a non-compliant domain or has subverted an edge router, the injection or modification of packets will not be caught by the traffic conditioners. In DiffServ, ingress nodes are the primary lines of defense against such activity.

# 5 Existing Efforts for Integration

The good news is that protocol designers are not ignoring the importance of establishing security principles in emerging QoS standards. Many RFCs and Internet drafts do contain notes on security considerations. A strong message found in many documents is the need for good system and network administration. [Zha98] states that network administrators are expected to protect network resources by configuring secure policers at interfaces with untrusted customers. Great care should be taken in deploying QoS technologies. This section also highlights efforts underway to aid administrators in the protection of their networks as they drive the inevitable deployment of QoS.

## 5.1 IETF Activities and Drafts

There are efforts underway within the standards bodies, such as the IETF and the QoS forum, to address the issues of security in QoS architectures. The IETF (Internet Engineering Task Force) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. [iet]

This section provides a synopsis of the IETF drafts that have been written thus far related to the issues identified in this paper.

### 5.1.1 Common Open Policy Standard: COPS

The COPS protocol provides a client/server model for distributed policy management in a network. COPS can be used within a domain for router policy enforcement points (PEPs) to retrieve policy from policy distribution points. [QoS99c] TCP is used as the transport protocol for reliable exchange of messages between policy clients and a server. [Sas99]

The COPS protocol provides an outsourcing mechanism for policy-based admission control so that a cohesive policy is achieved. COPS distinguishes between three request types [Pau99]:

- Admission Control Request: If a packet is just received by a PEP, it asks the PDP for an admission control decision on it.

- Resource Allocation Request: The PEP queries the PDP for a decision on whether, how, and when to reserve local resources for the request.
- Forwarding Request: The PEP asks the PDP how to modify a request and forward it to the other network devices.

State management is a large component of COPS. PDPs maintain state for all PEP requests until informed to delete that state. (This feature could lead to a clogging attack as described in section 3.2) PEPs periodically report status information to the PDP related to accounting and monitoring of requests.

Security considerations are also relevant to the COPS protocol. The COPS specification discusses an Integrity object that must be supported by all COPS implementations. The specification also highlights the use of IPSEC to secure communication between the PDP and PEPs. [Sas99] Currently COPS is designed to complement RSVP.

### 5.1.2 RSVP Integrity Object

The RSVP protocol was described in section 2.1.2 as a protocol for establishing distributed state in routers and hosts related to resource reservation. To ensure the integrity of the admission control mechanism, RSVP requires the ability to protect its messages against corruption and spoofing. [Tal99] An Internet Draft was written to propose an Integrity Object for RSVP. The decision was made not to use the IPSEC authentication header for this purpose. The rationale for this decision is documented in [Tal99].

The RSVP Integrity Object contains a message digest (HMAC-MD5 is recommended but not required) along with a sequence number. These two elements protect against message forgery and replay attacks. Confidentiality is not offered by this mechanism. The Internet draft provides details on algorithms to be used for generating and using the sequence numbers along with message handling procedures.

Key management for the Integrity Object mechanism is an area requiring further investigation. The requirements for a key management system are presented in the specification along with ideas of possible integration with the Kerberos. [Sch]

### 5.1.3 Working with IPSEC

**RSVP Extensions for IPSEC:** RSVP was originally specified to carry IP packets with TCP or UDP-like ports. The IPSEC protocol does not fit that model. The result was that RSVP offered limited functionality for IPSEC packets. In an effort to overcome those limitations, the IETF Network Working Group published a draft on extending RSVP for IPSEC flows. [O'M97] The extension takes advantage of the Security Parameter Index (SPI) in IPSEC to provide similar functionality as the TCP/UDP-like ports. Consequences of this extension for RESV messages is that processing will need to be slightly modified based on the new FILTER_SPEC ( which is the description of the packets to which the reservation applies) Additionally PATH message processing will be also be modified as described in the specification [O'M97]

**DiffServ Considerations:** RFC 2475 discusses the interaction of DiffServ with the IPSEC protocol. The two protocols have been developed such that modification of the DiffServ field by a network node has no effect on IPSEC's end-to-end security, because it cannot cause any IPSEC integrity check to fail. As a consequence, IPSEC does not provide any defense against an adversary's modification of the DS field. [Wei98] The solution is to take advantage of the IPSEC tunneling capabilities to encapsulate a DiffServ packet so that the DS codepoint is then protected by the encapsulation.. The requirement for this scenario is that the tunnel ingress and egress nodes be capable of performing the traffic conditioning functionality required of Differentiated Services. Further details on this scheme are provided in section 6.2 of RFC 2475.

### 5.1.4 Bandwidth Broker

Bandwidth brokers (BB) for handling QoS policies are being defined in association with Differentiated Services. A Bandwidth Broker is associated with a particular trust domain. Their purpose is to allocate bandwidth for end-to-end connections with state and simpler trust relationships than deploying per flow guarantees in all network elements, as is the case in Integrated Services. [QoS99a] In essence, the BB decides how applications should share services specified by the SLAs (see section 2.1.3).

A Bandwidth Broker has two primary responsibilities:

1. To parcel out their region's Marked traffic allocations and set up the leaf routers within the local domain.

2. To manage the messages that are sent across boundaries to adjacent region's BBs

These responsibilities are accomplished by maintaining a policy database and establishing relationships of limited trust with their peers in adjacent DiffServ domains. [QoS99a] Typically agreements are static once defined, but mechanisms are being developed to accommodate more dynamic service arrangements.

## 5.2 The Importance of Policy

A fundamental component of most of the IETF activities for securing QoS is policy. Many of the areas identified in the previous section are focused on various aspects of policy and communication. Policy provides a mechanism to establish what is allowed and disallowed in a network. Any compromise of security is in effect a breach of policy. It is therefore very important to build good and reliable policy tools to be used in QoS environments.

As an example, policy can be used to define how bandwidth is to be allocated among mission-critical financial software and multimedia entertainment feeds from the Internet. A policy should be defined and enforced to avoid "bandwidth hijacking" by those multimedia applications.

A policy framework has been proposed and is described in [QoS99b]. That document focuses on all aspects of policy. The framework (see figure 5) identifies the functional elements and protocols required to support QoS policy in the network. Many technologies will be needed to store, access, update, and monitor policy.

The IETF working groups focused on QoS and related subjects continuously point to policy frameworks as a means of supporting fair use of the technology.

## 5.3 Contribution of VPN Technology

Virtual Private Networking is a technology that has been around for some time. The motivation driving its continued development is the cost savings it offers over maintaining private leased lines. A VPN provides many other

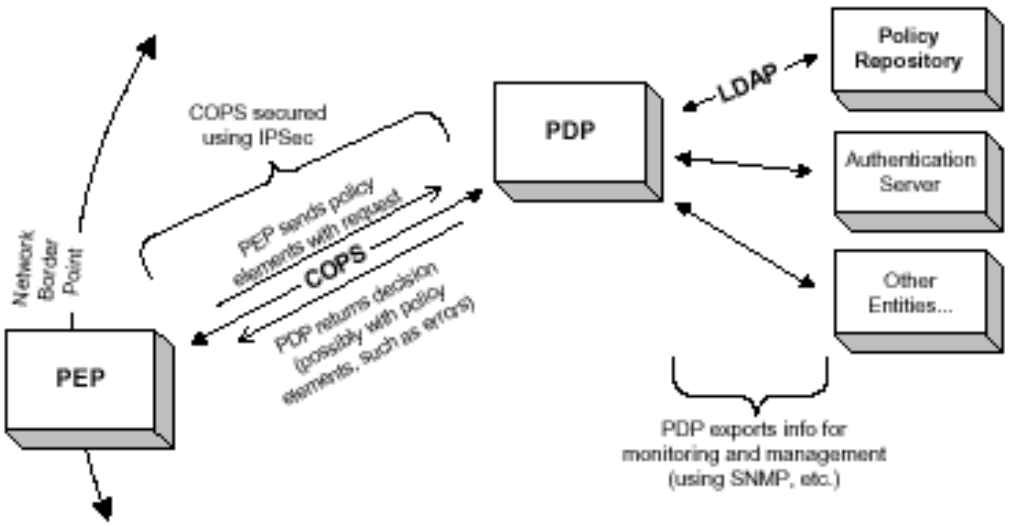


Figure 6: Functional elements and protocols required to support QoS policy in the network

business benefits which make it an attractive deployment option for all types of organizations. [IPH99b]

VPN technology is understood to be a good complement for QoS strategies. Both techniques aim to replace private leased line capabilities. One of the important requirements for IP based VPNs is to obtain differentiated and dependable Quality of Service for flows belonging to a VPN. Such VPNs will be capable of supporting a wide range of QoS guarantees as well as multiple traffic classes. [vdM98] The relationship between VPN customers can be used to tailor the service level agreements that drive traffic characteristics and applications that can take advantage of the VPN.

# 6 Conclusions

## 6.1 Areas for Future Work

A significant amount of research still remains in the area of securing QoS mechanisms. The study of analyzing the overall security of end-to-end QoS initiated by this paper should be continued as the protocol standards continue to be defined. Testing of proposed threats should be formalized to support the vulnerability claims.

In a larger scope, the notion of defining a policy and implementing it within the appropriate technology devices continues to be an active area of development. Every variety of policy must be defined, constructed, stored, accessed, applied, and enforced. [QoS99b] There are many challenges related to enforcing and updating policy controls within networks. The Policy working group at the IETF is chartered to define a scalable, interoperable, and vendor-independent framework for policy definition and administration. Their intermediate goal is to support QoS, however many other technologies will be able to take advantage of the outcomes from this group. [pWG99]

Security imposed techniques always come at a price of usability and performance. Each proposal for improving security in a protocol should be carefully studied in the effects to performance and usability. Every effort should be taken to prevent loss of satisfaction in the protocol after security is introduced.

## 6.2 Timing is Critical

A struggle that has existed as long as technology is in the inclusion of security controls in the initial design phase. So many project efforts wait until the end of their development cycle before they begin to explore the security issues effecting them. This survey paper has focused on ways to introduce security topics early in the design of new protocols and standards. Security is a topic that will not disappear, therefore it is critical that developers start every new project with a sense of security awareness. The time for building security controls into QoS standards is now. The protocols we analyzed in this paper are far enough in their evolution that implementation efforts are now underway.

A great deal of work has gone into encompassing QoS with policy frameworks as foundation along with admission control, and policing activities.

The work for securing QoS enabled networks does not stop at protocol design, however. It must be a continuous effort so that QoS standards are embraced as the networking intelligence tools they are designed to be.

## 6.3 Lessons from the Past

Our approach in this study was to review vulnerabilities that have been discovered in traditional network protocols (such as TCP and IP) and to let history guide the way for uncovering similar vulnerabilities in QoS standards. A lot of research has already taken place in the study of security in networks, so it seems a worthwhile activity to apply past results to new technologies. In this survey of security for QoS architectures we did encounter many of the historical issues related to protecting networks still applies. Those issues include trust, policy, and authentication.

Bringing those topics to bear in the QoS framework has led to a lot of activity currently underway in a variety of working groups at the IETF. Having this past experience to build upon gives QoS protocols the opportunity to be designed effectively so that when they are deployed and exploits become a reality, strong defenses will already be in place.

Harry Truman states our point well, "The only thing new in this world is the history you don't know."

# References


[AF] Jim Alves-Foss. Multi-Protocol attacks and the Public Key Infrastructure. Center for Secure and Dependable Software, Department of Computer Sciences, University of Idaho.

[Atk95] R. Atkinson. *Security Architecture for the Internet Protocol*, August 1995. RFC 1825.

[Bak99a] Fred Baker. Email communications, 1999.

[Bak99b] Fred Baker. Toward a Global Quality of Service Architecture, 1999. Presentation given at Purdue University.

[Bel89] S. M. Bellovin. Security problems in the TCP/IP protocol suite. *Computer Communication Review*, 19(2):32–48, April 1989.

[Bha] Bharat Bhargava. *Network Behavior and Quality of Service in Emerging Applications*. Purdue University.

[Bla99] Y. Bernet, A. Smith, S. Blake. *A Conceptual Model for Diffserv Routers*. Internet Engineering Task Force Draft, June 1999.

[Cal99] Eric Rosen, Arun Viswanathan, Ross Callon. *Multiprotocol Label Switching Architecture*. Internet Engineering Task Force Draft, August 1999.

[cis99a] Benefits of quality of service in intelligent networks, 1999. Cisco Systems Whitepaper.

[cis99b] End-to-end quality of service for WAN and campus networks, 1999. Cisco Systems Whitepaper.

[cis99c] Introduction: Quality of service overview. URL `http://www.cisco.com/univercd/home/home.htm`, 1999. Cisco Systems Documentation.

[Cla88] David D. Clark. The Design Philosophy of the DARPA Internet Protocols. *Computer Communication Review*, 18(4):106–114, August 1988.

[ea99] Yoram Bernet , et. al. *A Framework for Differentiated Services*. Internet Engineering Task Force Draft, Feburary 1999.

[Gon99] Zhi Fu, S. Felix Wu, T.S. Wu, He Huang, Fengmin Gong. *Security Issues for Differentiated Service Framework*. Internet Engineering Task Force Draft, October 1999.

[Hus98] Paul Ferguson, Geoff Huston. Quality of service in the internet: Fact, fiction, or compromise? In *INET 1998 Conference Proceedings*, July 1998.

[iet] Overview of the IETF. URL `http://www.ietf.org/overview.html`.

[IPH99a] IPHighway. Glossary of QoS terms. URL `http://www.iphighway.com`, 1999.

[IPH99b] IPHighway. Virtual private networking(VPN). URL `http://www.iphighway.com`, 1999.

[Jai99] Raj Jain. Quality of Service in Data Networks: Problems, Solutions, and Issues. URL `http://www.cis.ohio-state.edu/~jain/talks.html`, 1999. Presentation Slides.

[Joh99] Vicki Johnson. Quality of service – glossary of terms. URL `http://www.qosforum.com`, May 1999.

[Lam92] Thomas Y.C. Woo, Simon S. Lam. Authentication for distributed systems. *Computer*, 25(1):39–52, January 1992.

[LS87] Dennis Longley and Michael Shain. *Data and Computer Security: Dictionary of standards concepts and terms*. M Stockton Press, 1987.

[nor99] IP QoS – A Bold New Network, 1999. Nortel and Bay Networks Whitepaper.

[O'M97] L. Berger , T. O'Malley. *RSVP Extensions for IPSEC Data Flows*, September 1997. RFC 2207.

[Pau99] Arindam Paul. QoS in Data Networks: Protocols and Standards. URL `http://www.cis.ohio-state.edu/~jain/cis788-99/qos_protocols/index.html`, 1999.

[pWG99] IETF policy Working Group. Policy framework (policy). URL `http://www.ietf.org/html.charters/policy-charter.html`, October 1999.

[QoS99a] QoSforum.com. *Frequently Asked Qustions about IP Quality of Service*, May 1999.

[QoS99b] QoSforum.com. Introduction to qoS Policies. URL `http://www.qosforum.com/tech_resources.htm`, 1999.

[QoS99c] QoSforum.com. The Need for QoS. URL `http://www.qosforum.com/tech_resources.htm`, 1999.

[QoS99d] QoSforum.com. QoS protocols & architectures. URL `http://www.qosforum.com/tech_resources.htm`, 1999.

[Ray99] Gautam Ray. QoS in Data Networks: Products. URL `http://www.cis.ohio-state.edu/~jain/cis788-99/qos_products/index.html`, 1999.

[Rek98] T. Li, Y. Rekhter. *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*, October 1998. RFC 2430.

[San98] E. Crawley, R. Nair, B. Rajagopalan, H. Sandick. *A Framework for QoS-based Routing in the Internet*, August 1998. RFC 2386.

[Sas99] Jim Boyle , Ron Cohen , David Durham , Shai Herzog , Raju Rajan , Arun Sastry. *The COPS (Common Open Policy Service) Protocol*. Internet Engineering Task Force Draft, August 1999.

[Sch] Jennifer Steiner, Clifford Neuman, Jeffrey Schiller. Kerberos: An authentication service for open network systems.

[SG96] Gene Spafford and Simson Garfinkel. *Practical Unix and Internet Security*. O'Reilly & Associates, Inc, second edition, 1996.

[She94] R. Braden , D. Clark , S. Shenker. *Integrated Services in the Internet Architecture: an Overview*, June 1994. RFC 1633.

[Sum97] Rita C. Summers. *Secure Computing: Threats and Safeguards.* McGraw-Hill, 1997.

[Tal99] Fred Baker, Bob Lindell, Mohit Talwar. *RSVP Cryptographic Authentication.* Internet Engineering Task Force Draft, August 1999.

[TS98] Mahesh V. Tripunitara and Eugene H. Spafford. Issues in the incorporation of security services into a protocol reference model. Technical Report 98-03, Center for Education and Research in Information Assurance and Security, CERIAS, 1998.

[vdM98] Nick Duffield, Pawan Goyal, Albert Greenberg, Partho Mishra, K.K. Ramakrishnan, Jacobus van der Merwe. *A Performance Oriented Service Interface for Virtual Private Networks.* Internet Engineering Task Force Draft, November 1998.

[Wei98] S. Blake , D. Black , M. Carlson , E. Davies , Z. Wang , W. Weiss. *An Architecture for Differentiated Services*, December 1998. RFC 2475.

[Whi97] Paul P. White. RSVP and Integrated Services in the Internet: A Tutorial. *IEEE Communications Magazine*, May 1997.

[XN99] Xipeng Xiao and Lionel M. Ni. Internet QoS: A Big Picture. *IEEE Network Magazine*, March 1999.

[Zha98] Y. Bernet, R. Yavatkar, P. Ford, F. Baker, L. Zhang. *A Framework for End-to-End QoS Combining RSVP/Intserv and Differentiated Services.* Internet Engineering Task Force Draft, March 1998.