# Application of Quality of Service to Voice over IP Deployments

RJ Atkinson
Extreme Networks

January 30, 2004

**Abstract**

There is growing interest in deploying *Voice over IP services*, particularly within single enterprise environments. While applications such as file transfer or common web access do not need quality of service mechanisms in most environments, popular voice encoding algorithms might need support from Quality of Service (QoS) mechanisms in some network environments.

When a network deployment has been carefully engineered and is over-provisioned throughout, then no network congestion is possible and quality of service mechanisms will not be needed. However, over-provisioning of bandwidth is less common in enterprise networks. In enterprise networks where congestion might occur, deploying quality of service mechanisms can provide significant improvements to the perceived quality of the *Voice over IP* service.

This paper discusses the reasons that QoS mechanisms might be important, discusses approaches to deploying the Ethernet precedence and IP Type-of-Service to support QoS, and also discusses potential pitfalls with such deployments.

## 1 Introduction

Voice applications have long been used with datagram networks, such as The Internet. For example, audio/video applications have been used for over a decade on the Multicast Backbone (MBONE). Recently, there has been growing commercial interest in the use of *Voice over IP* as an adjunct to or replacement of traditional telephone service.

With any multimedia application used on the Internet, one needs to find a way to take an analogue natural signal source, such as a human voice, and use a codec algorithm [1] to convert the analogue source into digital format for packetisation and transmission through the network. After the voice has been encoded and had any compression or error-correction coding added, it is placed into a data packet and sent through the network. Commonly, multimedia data is framed using the *Real-Time Protocol (RTP)* [SCFJ03] and then sent via the *User Datagram Protocol (UDP)*[Pos81c].

Depending on the desired multimedia quality and the encoding(s) being used, and the nature of the underlying network between the source and the destination, problems could arise with delay, jitter, and/or packet loss. [2] Network Quality of Service (QoS) mechanisms are one way to help ensure that the desired multimedia quality is actually provided when delay, jitter, and/or packet loss are potential concerns.

[1] Selection of an appropriate codec is an important part of system engineering for a Voice-over-IP system; for example, the G.729 codec has virtually the same voice quality as G.711 but requires only 12% of the bandwidth.[RGW97] The details of codec evaluation and selection are, however, outside the scope of this paper.

[2] Different codecs will have different abilities to tolerate delay, jitter, and data loss.

This paper provides a discussion of circumstances when network QoS mechanisms might be helpful or needed, some candidate network QoS mechanisms to consider deploying, a candidate deployment strategy, and finally the residual issues that one should consider before deploying network QoS mechanisms in one's network. The focus of this paper is a single organisation or enterprise that has its own IP network and is deploying Voice-over-IP services within that network. Deployment scenarios involving more than one organisation are outside the scope of this paper.

## 2 Considerations in Network Design

Many commercial IP backbones have been carefully engineered so that congestion cannot occur within the backbone. Most typically, this involves over-provisioning backbone bandwidth – so that the backbone capacity exceeds the maximum load that could be placed upon it. However, many customers of such backbones have access links connecting to the backbone with less capacity than the maximum offered load on that link. So in modern IP networking, most access links occasionally or frequently experience traffic congestion. While congestion avoidance algorithms built into commonly used transport-layer protocols [3] will automatically detect congestion and reduce the offered load, it typically takes at least one round-trip time for the congestion avoidance algorithms to help reduce the offered load on the congested link.

If a link is experiencing congestion, queuing and packet loss are possible results. Depending upon how the network is configured, periodic congestion can also cause significant variation in the network delay that packets from some source to some destination experience. This variation is commonly known as *jitter*. [4] If one experiences more jitter than the codec in use can tolerate, then one should consider either changing the codec in use or altering the way the network is engineered. The two primary network engineering choices are to increase provisioned bandwidth such that congestion no longer occurs or to deploy one or more Quality of Service mechanisms within the enterprise network.

Traditional network engineering concerns must not be neglected, but instead gain additional importance, when voice or other real-time services are deployed. Network core switches should have high-availability capabilities, such as redundant power, redundant switch fabrics, redundant management modules, and fast failover. Edge switches ought to have at least redundant power options, ideally wired to separate power sources, for example one to a primary power circuit and the other to a separate backup power circuit. Ring-oriented topologies, for example *Ethernet Automatic Protection Switching*[Sha03], offer higher resiliency in the face of fibre cuts or equipment failures than strict tree-and-branch topologies can offer. Also, deploying networking equipment that has lower jitter and lower latency inside the switching/routing fabric will often help provide higher quality VoIP services.

## 3 Network QoS Mechanisms

Quality of Service is a long-standing research topic for the Internet. So far, no QoS mechanism has ever been widely deployed across most of the Internet. In part this is because Internet applications are typically designed to adapt to changing network conditions. In part, this is because inter-domain QoS tends to create significant operational security issues. And in part, this is because most IP backbone operators find it less expensive to over-provision capacity than to deploy and operate more complex network configurations that include network QoS mechanisms.

[3] For example the *Transmission Control Protocol (TCP)*[Pos81b] or the *Stream Control Transport Protocol (SCTP)*[SXM+00].

[4] Many voice codec algorithms will only tolerate modest amounts of jitter.

However, within a single administrative domain it can be practical to deploy some network QoS mechanisms. Two mechanisms that are openly specified and widely available in commercial networking equipment are *Ethernet Precedence*, which was originally specified in IEEE 802.1P, and *IP Type-of-Service* [Pos81a, NBBB98].[5] IP Type-of-Service is defined identically for both IPv4 and IPv6, so this paper will consistently just refer to "IP" and mean either IPv4, IPv6, or both. The general approach outlined here is to mark and police traffic at the edge of the enterprise network, while applying QoS throughout the enterprise network.

Also, it is worth noting that the IETF's work in the middle 1990s to standardise the *Resource Reservation Protocol (RSVP)*[BZB+97] never led to widespread implementation or deployment of that technology. Early commercial implementations of RSVP experienced serious scaling problems, perhaps because those implementations were CPU-based, rather than ASIC-based. In any event, the operational networking community[6] concluded that RSVP lacked sufficient scalability to be practical for per-flow QoS. RSVP remains in use for a very different purpose, as a signalling protocol option [7] for MPLS deployments. However, RSVP for per-flow resource reservation is not a widely-available or a good design option for VoIP deployments today.

Finally, it is important to note that not all implementations of a given standard are equally high quality. There are often significant differences in quality between one implementation and another. So it is important to actually laboratory test equipment before selecting it for deployment in one's network. Subsequent sections of this document will try to highlight particular implementation details that are likely to significantly impact the quality of one's VoIP deployment. It would be prudent to consider each of those points as part of one's overall systems engineering for VoIP services.

## 3.1 Over-Provisioning Bandwidth

This is the oldest QoS mechanism in the Internet community. In a Local Area Network (LAN) environment, fibre-optic backbones have been common since the advent of FDDI.[8] Since Gigabit Ethernet, and now 10 Gigabit Ethernet, have appeared, over-provisioning of LAN and even Metropolitan Area Networks (MANs) has become even more common, in large part because of the significant reduction in cost required to over-provision.

As with anything, there are tradeoffs to adopting over-provisioning as one's QoS mechanism. Capital costs will probably be higher. On the other hand, operational costs are usually lower because it the network design is simpler. One does not need to configure, operate, or trouble-shoot any other QoS mechanism if one has an over-provisioned network. Further, an over-provisioned network core can't be disrupted by a denial-of-service (DoS) attack originating at the edge of the network.

## 3.2 Ethernet Precedence

Ethernet is by far the most widely used technology in enterprise networking and is likely to remain so in the future. In the late 1990s, the IEEE standardised various extensions to Ethernet, for example support for Virtual LANs (VLANs). It also extended Ethernet by adding a QoS mechanism, *Ethernet Precedence*, which was originally specified in IEEE 802.1p.[9] This extension specifies a 3-bit field within the VLAN tag header that is used to carry precedence information. There are 8 precedence values, numbered 0 through 7, with priority 7 being the highest. As we shall see shortly, this scheme maps nicely to the IP Precedence bits.

[5] IP Differentiated Services is backwards compatible with IP Precedence.
[6] For example, within the North American Network Operators Group (NANOG) or the European Operators' Forum (EOF).
[7] *Label Distribution Protocol (LDP)* defined in RFC-3036 is another option for use with MPLS.
[8] Fibre Distributed Data Interface, a 100 Mbps ring technology standardised by ANSI.
[9] Later the contents of IEEE 802.1p were incorporated into a revision of the IEEE 802.1d specification.

Many Ethernet switches now implement support for Ethernet precedence. A variety of queuing algorithms can be found, most commonly strict priority queuing and weighted round-robin. In better implementations, it is possible to guarantee some minimum bandwidth amount for each QoS value. Typically, such implementations will make any guaranteed capacity that is unused available for traffic having other QoS values. Equipment varies widely on which forms of queuing are supported within a 802.1p implementation. Although the IEEE standard specifies 8 different precedence values, not all Ethernet equipment supports a full 8 queues per port. Some equipment that implements the 802.1p specification supports 8 queues/port, others support only 4 queues/port, or even 2 queues/port. As we will discuss later on, 3 queues/port is a practical minimum needed to support 2 different QoS levels for user traffic within a real operational network, though 8 queues/port is ideal.

Further, some Ethernet equipment will evaluate the precedence tag on ingress and use that information in mediating access to the switch backplane, which helps ensure that higher precedence traffic gets higher priority access to the switch backplane. By contrast, some other equipment does not do this, in which case lower precedence traffic might delay backplane access for higher precedence traffic.

However, not all Ethernet switches offer all of these capabilities. So the network engineer should carefully select and test equipment to ensure that it has full QoS capabilities to support all kinds of network traffic.

### 3.3 IP Type-of-Service

IP has long supported a per-packet quality-of-service marking in its *Type of Service* field. This 8-bit field originally used 3 bits to support 8 precedence values, along with some handling flags in the remaining 5 bits. While the precedence values were widely supported in early IP routers, not all products supported them. There were even early deployments of IP precedence, for example in some US military IP networks.

The designers of the RFC-791 precedence model simply adopted a long-standing message handling precedence scheme of the US Department of Defence [DoD83]. In this scheme, there are 6 precedence levels for user-traffic, ranging from *ROUTINE* used for most traffic to *FLASH OVERRIDE* used only in a dire emergency. In addition, the IP precedence model has 2 precedence values higher than those used for any user traffic. The highest precedence value is called *INTERNET CONTROL* and is normally used for control traffic that can affect network availability and stability across multiple administrative domains (e.g. Border Gateway Protocol which carries inter-domain routing information). The second-highest precedence value is called *NETWORK CONTROL* and is normally used for control traffic that can affect network availability and stability within a single administrative domain (e.g. Open Shortest Path First, which carries intra-domain routing information).

The IP precedence model ensures that critical network control traffic is given higher precedence than any user traffic. Should user traffic ever crowd out that network control traffic, the network would probably develop faults that would ultimately prevent user traffic from reaching its intended destination. Further, if the network were to develop non-protocol faults (e.g. a fibre cut), the network control traffic would be crucial to letting the network discover that fault and automatically route around the damaged section(s) of the network.

More recently, the IETF has produced the *Differentiated Services* specification [NBBB98] which provides an alternate set of intepretations for this 8-bit field. Further, the IETF defines some packet handling specifications for use with Differentiated Services. The two IETF standards-track specifications for DiffServ packet processing are known as *Assured Forwarding (AF)* [HBWW99] and *Expedited Forwarding* [D$^+$02].

While EF was originally designed for use in carrying voice traffic in the US Department of Energy's En-

| *JANAP Traffic Type* | *Ethernet Precedence* | *IP Precedence* |
|---|---|---|
| Internet Control | 7 | 7 |
| Network Control | 6 | 6 |
| Critical ECP | 5 | 5 |
| Flash Override | 4 | 4 |
| Flash | 3 | 3 |
| Immediate | 2 | 2 |
| Priority | 1 | 1 |
| Routine | 0 | 0 |

Table 1: IP Precedence Mapping

ergy Sciences Network (ESnet), there is a common misconception that only EF is well suited for handling voice traffic. In fact, experience has shown that AF is also well suited for use with voice traffic. The EF specification [D$^+$02] contains specific suggestions on how to calculate delay and jitter bounds for a given EF implementation. While the AF specification [HBWW99] does not contain specific suggestions, one can also calculate both delay and jitter bounds for AF processing.

Further, better quality DiffServ implementations offer a finer-grained set of queuing configurations to the network operator, rather than merely offering the operator the two coarse-grained options of AF or EF. For example, better implementations offer the operator a choice of queuing algorithms, commonly including Priority Queuing, Weighted Fair Queuing, and Weighted Random Early Drop (WRED). Also, better implementations permit each queue to be allocated a minimum bandwidth that will always be available for traffic in that queue and also a maximum bandwidth. So for the remainder of this paper, we will talk about the QoS configuration in more detail than merely referring to AF or EF would permit. It is recommended that one use equipment that permits such fine-grained QoS configuration, not mere Differentiated Services support, as this helps the resulting deployed network configuration be fully successful.

It is important to select and deploy networking equipment that has all of these capabilities, ideally including 8 queues/port, a variety of queuing algorithms, and the ability to provision minimum and maximum bandwidths for each queue.

### 3.4 Ethernet and IP QoS Comparison

Since both Ethernet Precedence and the IP Precedence defined in the IP Type-of-Service field each specify 8 precedence values or QoS queues, it is straight-forward to use these two mechanisms in tandem to provide end-to-end QoS within the enterprise network. To do so, it is important that the deployed network equipment support both mechanisms. Table 1 is a table showing one mapping between the Ethernet Precedence and IP Precedence, along with original US DoD uses for each QoS value.

### 3.5 QoS Filtering

While the standards define how a QoS marking is represented in an IP packet header or Ethernet frame header, the standards do not define how to ensure that a given packet or frame contains the correct QoS marking. In practice, better implementations of Ethernet Precedence or IP ToS support filtering incoming traffic, often using Access Control Lists, and then marking (or re-marking) the incoming packet or frame with the correct QoS marking. This marking is then used within the switch or router to apply appropriate

packet or frame processing to implement the desired Quality of Service.

In equipment implementations that do not support a full 8 queues/port, the equipment will generally need to be configured so that traffic with the correct set of QoS markings is sent to the correct queue. Equipment having the full 8 queues/port will be more successful in applying the desired QoS handling to packets or frames passing through it.

Also, it is important to select networking equipment that has fully flexible ACL capabilities so that the deployment can ensure that only authorised traffic is able to obtain preferred service quality.

# 4 Deployment

As noted earlier, the simplest QoS deployment consists of carefully engineering the deployed network so that congestion cannot occur because bandwidth has been over-provisioned. If this option is available and economically sensible, it is probably the best approach. On links where congestion might occur, use of other QoS mechanisms might make sense. Here we recommend using both Ethernet Precedence and IP ToS in combination on such links.

There are a variety of different deployment models that one might consider. Due to space considerations, we will only present 3 deployment models here. The first model, which we call the *Simple QoS Model*, can be implemented if equipment on potentially congested links supports IP ToS, has 4 queues/port, and optionally also supports Ethernet Precedence. The second model, which we call the *Fine-Grained QoS Model*, can be implemented if equipment on potentially congested links supports IP ToS, has 8 queues/port, and optionally also supports Ethernet Precedence. In both of these models, we provision a minimum bandwidth for most QoS queues to prevent higher-precedence traffic from totally starving lower-precedence traffic of bandwidth. The third model, which we call the *Strict Priority Model*, differs from the first two in that traffic at lower precedence levels might be starved completely by high demand from traffic at higher precedence levels. This third model might be applicable in situations where there is a well-defined QoS policy that requires starvation. For example, in an emergency situation or in a military context it might be strongly desirable for the most important traffic to be delivered – even if less important traffic were unable to be sent at all.

It is important to keep in mind that these are three deployment examples, not hard and fast design rules. Each organisation ought to consider what kind of Quality of Service policy is appropriate for that organisation – and then deploy a configuration consistent with that locally-designed policy. Each organisation has different needs, different network designs, and so each will probably have a different Quality of Service policy. If an organisation does not mind having lower-priority services starved of network access, then a Strict Precedence queuing approach without any bandwidth guarantee might make sense. If one wants to avoid starving any QoS category, then one probably wants to configure some guaranteed access to network capacity for each QoS value.

There are two fixed design principles for QoS policy. The first is that the network control traffic should be guaranteed a relatively higher amount of capacity than other classes, because if that control traffic cannot get through then the network will likely cease working properly for all traffic. The second is that network control traffic should always be higher precedence than any user traffic, for the same reason.

| *Traffic Type* | *Example Protocols* | *Ethernet Precedence* | *IP Precedence* |
|---|---|---|---|
| Internet Control | BGP, PIM, SNMP | 7 | 7 |
| Network Control | STP, OSPF, RIP | 6 | 6 |
| Voice | SIP, MGCP, RTP | 5 | 5 |
| Other | NFS, SMB, RPC, SQL, IM, HTTP, FTP, SMTP | 0 | 0 |

Table 2: Simple QoS Model

### 4.1 Simple QoS Model

In this model, shown in Table 2, we break network traffic into 4 categories, *Internet Control*, *Network Control*, *Voice*, and *Other*. This break-down requires 4 queues/port in the applicable network equipment. If one's equipment can only support 3 queues/port, one could consolidate *Internet Control* and *Network Control* into a single category without much adverse impact. The *Control* categories are highest precedence and *Other* is the lowest precedence in this scheme.

Inter-domain control traffic, for example BGP, belongs in *Internet Control*. Intra-domain control traffic, for example SNMP, OSPF, or RIP, belongs in *Network Control*. Voice traffic is sorted out next. This includes not only actual voice packets sent using the Real-Time Protocol (RTP), but also any telephony signalling protocols that are deployed, for example SIP. The last category contains all other traffic, probably consisting mainly of HTTP for web access and SMTP, POP, or IMAP for email access.

In this model, we do not want any class of traffic to be starved of bandwidth by other classes of traffic during normal operation. So one should configure a minimum reserved bandwidth for each class. For example, one might want to guarantee 20 percent of bandwidth for Voice, 5 percent for Other, and 10 percent each for each type of Control traffic, with the remaining bandwidth dynamically allocated among the 4 QoS values based on the relative precedence of the QoS values in the traffic being received.

A short-coming of this model is that it lumps all non-voice user traffic into a single QoS class. In most enterprises, not all data traffic is equally important. For example, file server access is typically very important, with database access only slightly less so. Finally, even for mundane data traffic the interactive traffic (e.g. instant messaging) normally should get higher priority than background traffic (e.g. file transfer).

### 4.2 Fine-Grained QoS Model

In this model, shown in Table 3, we break network traffic into 8 categories. This scheme contains the same four categories as in the Simple QoS Model, but adds 4 more categories having precedence greater than *Other* but less than *Voice*. Also, the *Voice* category is broken into two separate categories. The first of these is *Voice Control* which contains only voice control or telephony signalling protocols, such as the Session Initiation Protocol (SIP)[RSC+02]. The second of these, which is lower precedence than the first, is *Voice Traffic* which contains only the actual voice content, typically carried in the Real-Time Protocol (RTP). This leaves three additional categories for higher-precedence non-voice user traffic. A common configuration would use one of these categories for file-server and remote-procedure-call traffic, and the other two for business-critical applications (e.g. remote database access). Web content, electronic mail, instant messaging, and any gaming applications would typically be split between the *Interactive* and *Other* categories as shown.

Again, some categories (e.g. Control, Voice) would probably be given guaranteed capacity, but the percentage that is guaranteed would probably decrease for most categories. For example, the network control traffic

| *Traffic Type* | *Example Protocols* | *Ethernet Precedence* | *IP Precedence* |
|---|---|---|---|
| Internet Control | BGP, PIM, SNMP | 7 | 7 |
| Network Control | STP, OSPF, RIP | 6 | 6 |
| Voice Signalling | SIP, MGCP | 5 | 5 |
| Voice Traffic | RTP | 4 | 4 |
| File Access | NFS, SMB, RPC | 3 | 3 |
| Database | SQL | 2 | 2 |
| Interactive | HTTP, IM, X11 | 1 | 1 |
| Other | FTP, SMTP | 0 | 0 |

Table 3: Fine-Grained QoS Model

categories might each get guaranteed access to 10 percent of capacity, with the next highest five categories each getting guaranteed access to 5 percent of capacity.

### 4.3 Operational Considerations

The general deployment concept outlined here is to mark and police QoS at the edge of the enterprise network, while applying QoS throughout the enterprise network. This approach has been shown to scale well in past deployments and is also straight-forward to deploy.

Experience has shown that it is best if VoIP traffic is segregated from other data traffic. While this is not always practical to deploy on shared WAN links, it can easily be deployed on Ethernet networks by simply using Virtual LANs (VLANs) to provide appropriate separation. For example, VoIP traffic might be on VLAN number 3, while data traffic is on a different VLAN, perhaps VLAN number 0, the default VLAN. Because it is important that such VLANs do not accidentially "leak" traffic onto other VLANs, one should try to select networking equipment that implements VLAN capabilities in dedicated ASIC hardware, rather than on the main switch CPU.

Further, one should lock-down each Ethernet port used for VoIP to the specific MAC Address of the device that is supposed to be connected to that port. This will help reduce the risk of misconfiguration [10] While this is a common feature in Ethernet switches, not all products support this MAC lock-down feature, so one should be careful when selecting one's Ethernet infrastructure.

Also, some Ethernet equipment still relies on the now ancient practice of implementing Access Control Lists (ACLs) on the main switch CPU, rather than having specific hardware support for ACLs. CPU-based ACLs cause switch performance to drop as the number of ACLs increases or the data traffic increases, while hardware-based ACLs can operate at wire-speed regardless of the packet load or the number of configured ACLs. So one should be careful to select networking equipment that implements its ACLs in hardware, not on the main CPU.

Security is one of the larger barriers to deployment of QoS mechanisms in networks. If a network offers differing service quality to different packets, this creates an incentive for users to improperly cause their traffic to get the best service quality. In a best effort only network, there is no incentive to improperly mark traffic to obtain best service quality since all packets are always treated equally.

There are no cryptographic mechanisms available for validating the IP ToS bits or for validating the Ethernet

[10] For example, where a naive user plugs his or her laptop into his VoIP phone's Ethernet port and the phone into the laptop's Ethernet port.

precedence bits. Even if a cryptographic mechanism were available, it likely would be impractical to employ. For example, consider the hardware cost and deployment complexity required to authenticate every frame that is transiting some 10 Gigabit Ethernet link, for example.

Instead, the best security approach is to have Access Control Lists deployed at the edges of the network. These ACLs cause packets that are not marked to become properly marked and also cause packets that are erroneously marked to become properly marked. A side benefit of this approach is that one can use any sort of host on the network and obtain the benefits of differentiated service quality. One does not need to upgrade hosts to implement the QoS mechanisms, nor configure the hosts with one's local QoS policy, nor modify applications to use some new networking API to request a different service quality from the network. The primary issue with this is that one has to ensure that the local QoS policy is implemented consistently at each edge of the network. This can represent a significant operational cost for the network operator, not only for initial configuration but also for configuration maintenance over time. Automated systems for configuration management, such as Extreme's *EpiCentre* product, are a practical requirement for any network of medium or large size.

Further, if one has deployed multiple service qualities, one should monitor the usage of the varying service qualities to ensure that the services actually provided are those that one intends to provide. So there is increased operational cost in network monitoring. Fortunately, this lends itself readily to automation, most commonly using SNMP-based tools[11], perhaps in combination with UNIX scripts or an enterprise network management system.[12]

# 5 Conclusions

Organisations considering deploying Voice over IP service on their enterprise network should consider whether to over-provision capacity or deploy QoS. The organisation should carefully consider which codecs to use in their deployment.

If the decision is to deploy QoS, the network should be outfitted with equipment supporting a full 8 queues/port, flexible QoS queuing and scheduling algorithms, both minimum and maximum bandwidths for each queue, along with flexible ACLs to ensure that the correct QoS marking is applied to each packet and to filter out unauthorised traffic.

It is also important to select and deploy networking equpiment that has ASIC-based implementations of ACLs, VLANs, Ethernet bridging, and IP forwarding. That way enabling important features does not reduce the network performance and the latency through the switch or router is minimised. While CPU-based implementations are becoming less common due to their technical problems, some equipment sold today still places important capabilities inside the main CPU because that equipment does not have ASIC support for these features.

Further, to reduce operational and deployment costs, the enterprise should consider using an enterprise network management system that has the ability to provision policy and other configuration across the entire network.

[11] For example, the *Multi-Router Traffic Grapher (MRTG)*.
[12] For example, Extreme Networks' *EpiCenter* solution.

# References


[BZB+97] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin. Resource Reservation Protocol (RSVP), Version 1, Functional Specification. Technical report, Internet Society, September 1997. RFC-2205.

[D+02] B. Davie et al. An Expedited Forwarding Per-Hop Behaviour (PHB). Technical report, Internet Society, March 2002. RFC-3246.

[DoD83] Automatic Digital Network (AUTODIN) Operating Procedures. Technical report, US Dept of Defense, March 1983. JANAP-128.

[HBWW99] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski. Assured Forwarding PHB Group. Technical report, Internet Society, June 1999. RFC-2597.

[NBBB98] K. Nichols, S. Blake, F. Baker, and D. Black. Definition of the Differentiated Services Field in the IP Headers. Technical report, Internet Society, December 1998. RFC-2474.

[Pos81a] J. Postel. Internet Protocol. Technical report, Internet Society, September 1981. RFC-791.

[Pos81b] J. Postel. Transmission control protocol. Technical report, Internet Society, September 1981. RFC-793.

[Pos81c] J. Postel. User Datagram Protocol. Technical report, Internet Society, September 1981. RFC-792.

[RGW97] S. Rudkin, A. Grace, and M.W. Whybray. Real-Time Applications on the Internet. *BT Journal*, 15(2), April 1997.

[RSC+02] H. Rosenberg, H. Schulzrinne, G. Camarillo, et al. Session Initiation Protocol. Technical report, Internet Society, June 2002. RFC-3261.

[SCFJ03] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP A Transport Protocol for Real-Time Applications. Technical report, Internet Society, July 2003. RFC-3550.

[Sha03] S. Shah. Ethernet Automatic Protection Switching. Technical report, Internet Society, October 2003. RFC-3619.

[SXM+00] R. Stewart, Q. Xie, K. Morneault, et al. Stream Control Transmission Protocol. Technical report, Internet Society, October 2000. RFC-2960.