Specification of Committed Rate Quality of Service

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet- Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet- Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

This document is a product of the Integrated Services working group of the Internet Engineering Task Force. Comments are solicited and should be addressed to the working group's mailing list at int-serv@isi.edu and/or the author(s).

Abstract

This document describes the network element behavior required to deliver a Committed Rate service in the Internet. The Committed Rate service provides applications with a firm commitment from the network, that at a minimum the transmission rate they requested is available to them at each network element on their path. The commitment of a given transmission rate by a network element is not associated with a specific delay guarantee, but requires that network elements perform admission control to avoid over-allocation of resources.

Introduction

This document defines the requirements for network elements that support a Committed Rate service. This memo is one of a series of documents that specify the network element behavior required to support various qualities of service in IP internetworks. Services described in these documents are useful both in the global Internet and private IP networks.

This document is based on the service specification template given in [1]. Please refer to that document for definitions and additional information about the specification of qualities of service within the IP protocol family.

End-to-End Behavior

The end-to-end behavior provided to an application by a series of network elements that conform to the service described in this document is a committed transmission rate that, when used by a policed flow, ensures that its packets are transmitted with no or minimal queueing losses through the network (assuming no failure of network components or changes in routing during the life of the flow). In addition, while this service does not provide any specific delay guarantees, the provision of a committed transmission rate at each network element should ensure that packets do not experience delays at a network element that are significantly in excess of what they would experienced from a dedicated transmission facility operating at the committed rate.

To ensure that this service is provided, clients requesting Committed Rate service provide the network elements with both the transmission rate they want to have guaranteed and information on their traffic characteristics. Traffic characteristics are specified in the TSpec of the flow, which is defined in the section on Invocation Information below. In return, the network elements perform admission control to allocate enough resources (bandwidth and buffer) to ensure that over some reasonable time interval, a flow with packets waiting to be transmitted sees a transmission rate at least equal to its requested transmission rate and only experiences packet losses very rarely.

Motivation

The Committed Rate service is intended to offer applications a service that provides them with the guarantee that the network will commit a certain amount of bandwidth to them in an attempt to emulate a dedicated circuit of at least that amount of bandwidth. This service is intended for applications that require a given amount of bandwidth in order to operate properly. This bandwidth can be related to an intrinsic rate at which the applications generates data, e.g., the average transmission rate of a video codec, or chosen so as to allow the transmission of a certain amount of data within a reasonable time, e.g., a function of the size of the data object to be transmitted and how fast it needs to be received.

The rate guarantees provided by the Committed Rate service are in a sense similar to those provided by the Guaranteed Service [2]. However, a key difference is that they are not coupled to the same rigorous delay guarantees provided by the Guaranteed Service. This decoupling simplifies the invocation of the service and its support at intermediate network elements as the service is only a function of local resources at each node and, therefore, independent of the end-to-end characteristics of the path itself. In addition, the relaxation of the delay guarantees to be provided, can allow a higher utilization of network resources, e.g., bandwidth. However, note that this greater simplicity and higher efficiency come at a cost, namely

- The lack of hard delay guarantees. This is because the commitment of a transmission rate at a network element can be provided through a range of mechanisms, that correspond to different delay behaviors (see the section on Network Element Data Handling Requirements for additional details). Specifically, depending on the characteristics of the implementation used to support the Committed Rate service at network elements, the worst case delay experienced by packets receiving this service could be much higher than under Guaranteed service, i.e., the delay bounds are relaxed.

- The lack of a priori estimates of the end-to-end delay to be expected. This is because rate guarantees are local to each network element, and hence do not provide any end-to-end delay characterization for the path on which the flow is routed.

- Weaker loss guarantees as the lack of characterization of the behavior of individual network elements also means, that accurate sizing of buffer requirements to ensure lossless operation cannot be provided.

The Committed Rate service is also different from the Controlled-Load service [3] in that it allows policing at the edge of the network and reshaping at intermediated network elements. Note that the emulation of a dedicated circuit at the requested committed rate, can amount to reshaping the flow of packets to this rate. In addition, contrary to the Controlled-Load service that specifies that the transit delay experienced by most packets should be close to the minimum transit delay, the Committed Rate service only guarantees that the transit delay of any packet should not significantly exceed what it would have experienced over a dedicated circuit at the committed rate.

The Committed Rate service can, therefore, be viewed as an intermediate service level in between the Guaranteed Service and the Controlled Load service. It provides weaker service guarantees than the Guaranteed Service, but imposes fewer constraints on network elements. This may facilitate deployment in an heterogeneous environment, where not all network elements may be capable of satisfying the requirements of the Guaranteed Service. Similarly, the provision of a fixed rate service guarantee may be less flexible than the Controlled-Load service for adaptive applications, but may simplify the call admission and scheduling functions at network elements. In particular, the buffer and bandwidth allocation functions may benefit from the stricter traffic (TSpec) specification.

Network Element Data Handling Requirements

The network element must ensure that the service approximates a dedicated circuit of rate at least equal to the requested rate R. This means that the network element must perform call admission to ensure the availability of sufficient bandwidth to accommodate a flow's request for a transmission rate R. This will typically mean allocating an amount of link bandwidth at least equal to R. However, note that this service specification does not require that a network element provide an application with a transmission rate greater than R even when there is excess bandwidth available, i.e., reshaping of the traffic to the rate R is allowed. More generally, approximating a dedicated rate circuit of rate at least R only implies, that if for a period of time T an application has data packets waiting to be transmitted at a network element, i.e., it is backlogged, the amount of data it is able to send during T should approach RT as T grows large. The smaller the value of T for which this is achieved, the better the approximation of a dedicated circuit at rate R.

Specifically, the difference between the service provided at a network element and a dedicated rate circuit is a function of the scheduler used at the service element and also reflects the impact of the packetized nature of transmission units. Many recently proposed schedulers, e.g., Weighted Fair Queueing (WFQ) [4], Virtual Clock [5], Rate Controlled Service Disciplines (RCSD) [6,7], Latency Rate Servers [8], etc., can provide reasonably good approximations for such a service, i.e., typically with a value for T of the order of

L/R, where L is the size of the packet to be transmitted. On the other hand, other simpler schedulers such as FIFO, static priorities, frame based schedulers [9], may only provide relatively coarse approximations. While this service specification does not mandate the use of a particular type of scheduler, the nature of its service definition, i.e., the close approximation of a dedicated rate circuit, means that the use of schedulers that perform well in regards to this measure is recommended.

In addition, the manner in which a network element approximates a dedicated rate circuit will impact the amount of buffering it needs to provide to ensure minimum losses for a compliant flow. For example, a network element for which the scheduler can delay transmission of packets from a given flow for a time period T, may need to buffer up to b+rT amount of data for that flow, where b corresponds to the token bucket depth advertised in the TSpec of the flow and r is the associated token bucket rate (see next section on Invocation Information for details). It is, therefore, expected that each network element allocates, when accepting a new flow, not only enough bandwidth to accommodate the requested rate, but also sufficient buffer space to provide compliant flows with minimal losses. Note that it is possible for a network element to trade bandwidth for buffer space, i.e., allocate to each flow more bandwidth than it requests, so as to ensure a low enough load to keep buffer requirements low. The necessary amount of buffer space is a quantity to be engineered for each network element. For example, edge network elements may need to allocate b or more to each flow (this depends in part on how much larger R is than r), while core network elements may be able to take advantage of statistical multiplexing to allocate much less.

Links are not permitted to fragment packets as part of the Committed Rate service. Packets larger than the MTU of the link must be policed as non-conformant which means that they will be policed according to the rules described in the Policing section below.

Packet losses due to non-congestion-related causes, such as bit errors are not accounted for by this service.

Invocation Information

The Committed Rate service is invoked by specifying the traffic (TSpec) of the flow and the desired Committed Rate (RSpec) to the network element. A service request for an existing flow that has a new TSpec and/or RSpec should be treated as a new invocation, in the sense that admission control must be reapplied to the flow. Flows that reduce their TSpec and/or RSpec (i.e., their new TSpec/RSpec is strictly smaller than the old TSpec/RSpec according to the ordering rules described in the section on Ordering below) should never be denied service.

The TSpec takes the form of a token bucket plus a peak rate (p), a minimum policed unit (m), and a maximum packet size (M).

The token bucket has a bucket depth, b, and a bucket rate, r, which corresponds to the requested rate that the flow is requesting the network to commit. Both b and r must be positive. Note that it is necessary to habe b>=M. The rate, r, is measured in bytes of IP datagrams per second, and can range from 1 byte per second to as large as 40 terabytes per second (or about what is believed to be the maximum theoretical bandwidth of a single strand of fiber). Network elements MUST return an error for requests containing values outside this range. Network elements MUST return an error for any request containing a value within this range which cannot be supported by the element. In practice, only the first few digits of the r parameter are significant, so the use of floating point representations, accurate to at least 0.1%, is encouraged.

The bucket depth, b, is measured in bytes and can range from 1 byte to 250 gigabytes. Network elements MUST return an error for requests containing values outside this range. Network elements MUST return an error for any request containing a value within this range which cannot be supported by the element. In practice, only the first few digits of the r parameter are significant, so the use of floating point representations, accurate to at least 0.1%, is encouraged.

The range of values for these parameters is intentionally large to allow for future network and transmission technologies. This range is not intended to imply that a network element must be capable of supporting the entire range of values.

The peak rate, p, is measured in bytes of IP datagrams per second and has the same range and suggested representation as the bucket rate. The peak rate is the maximum rate at which the source and any reshaping points (reshaping points are defined below) may inject bursts of traffic into the network. More precisely, it is the requirement that for all time periods the amount of data sent cannot exceed M+pT where M is the maximum packet size and T is the length of the time period. Furthermore, p must be greater than or equal to the token bucket rate, r. A peak rate value of 0 means the peak rate is not being used or is unknown.

The minimum policed unit, m, is an integer measured in bytes. All IP datagrams less than size m will be counted, when policed and tested for conformance to the TSpec, as being of size m. The maximum packet size, M, is the biggest packet that will conform to the traffic specification; it is also measured in bytes. A network element must reject a service request, if the requested maximum packet size is larger than the MTU of the link. Both m and M must be positive, and m must be less than or equal to M.

The RSpec consists of the desired service rate R. The motivations for separating the specification of the rate R from the the token bucket rate in the TSpec, are to provide greater flexibility in the level of service a receiver can request and in simplifying support for shared reservations. With shared reservations, a receiver can request a certain Committed Rate R from the network, that may not be directly related to the token bucket rates specified in the TSpec of the different flows that are to share the reservation.

The preferred representation for the TSpec consists of three floating point numbers in single-precision IEEE floating point format followed by two 32-bit integers in network byte order. The first value is the rate (r), the second value is the bucket size (b), the third is the peak rate (p), the fourth is the minimum policed unit (m), and the fifth is the maximum packet size (M).

The preferred representation for the RSpec rate, R, is also in single-precision IEEE floating point format.

For all IEEE floating point values, the sign bit must be zero. (All values must be positive). Exponents less than 127 (i.e., 0) are prohibited. Exponents greater than 162 (i.e., positive 35) are discouraged.

The Committed Rate service is assigned service_name 6.

The Committed Rate traffic specification parameter (TSpec) is assigned parameter_name 1, as indicated in the listing of well-known parameter name assignments given in [1].

Exported Information

The Committed Rate service has no required characterization parameters. Individual implementations may export appropriate implementation-specific measurement and monitoring information.

Policing and Reshaping

Policing and reshaping are two related forms of traffic control, that are meant to limit the amount of traffic that an application can inject into the network. In either cases, the result is that only conformant packets are forwarded. Conformance is determined as a function of the TSpec for the flow. A flow is deemed conformant if the amount of data it sent during any given time period of duration T does not exceed M+min[pT, rT+b-M], where p is the peak rate, r and b are the token bucket parameters, and M is the maximum packet size for that flow. For the purposes of this accounting, links must count packets which are smaller than the minimal policing unit to be of size m. Packets which arrive at an element and cause a violation of the M+min[pT, rT+b-M] bound are considered non-conformant. Additionally, packets bigger than the outgoing link MTU are considered non-conformant. It is expected that such a situation will typically not arise, because flow setup mechanisms are expected to notify the sending application of the appropriate path MTU.

Policing and reshaping differ in their treatment of non-conformant packets. Policing performs a strict control on the traffic of a given flow by either discarding non-conformant packets, or possibly sending them as best effort packets. Note that if and when a marking ability becomes available, non-conformant packets sent as best-effort packets SHOULD be ''marked'' as being non-compliant so that they can be treated as best effort packets at all subsequent network elements. In the context of the Committed Rate service, policing should ONLY be performed at the edge of the network, where it is used to ensure conformance of the user traffic with the TSpec it advertised.

On the other hand, the strict traffic control implied by policing is NOT appropriate inside the network, since the perturbations caused by the queueing and scheduling delays at network elements will often turn an initially conformant flow into a non-conformant one. Instead, it is recommended that reshaping be used at intermediate network elements inside the network. Reshaping amounts to delaying (buffering) non-conformant packets until they are compliant, rather than discard or send them as best-effort. Reshaping, therefore, restores the traffic characteristics of a flow to conform to the specified token bucket and peak rate parameters used by the reshaper. (To avoid delaying unnecessarily the initial packets of a flow, the token bucket at a reshaper should be initialized full).

The benefit of restoring a flow to its original envelope is that it limits the magnitude of the "distortions" that schedulers at network elements can introduce in the initial stream of packets from a flow. As discussed in the section on Network Element Data Handling Requirements, depending on how well a scheduler approximates a dedicated rate circuit, significant bunching up of packets can be introduced. This translates in turn into bigger buffer requirements at downstream network elements. Reshaping the flow ensures that downstream network elements are isolated from the bunching effects introduced by upstream schedulers. However, note that in order to achieve these benefits, the reshapers must provide sufficient buffer space to hold packets until they can be released as compliant with the traffic envelope to which the flow is being reshaped.

Contrary to the Guaranteed Service where the information exported by network elements allows the computation of an upper bound on the amount of buffer needed when reshaping traffic, in the context of the Committed Rate service this quantity can only be estimated. The required amount to ensure minimum packet losses to Committed Rate flows is, therefore, a quantity to be engineered for each network element.

If a packet arrives at a reshaper and finds the reshaping buffer full, the packet can either be discarded or accommodated by forwarding as best effort a packet of the flow.

> NOTE: As with policers, it should be possible to configure how reshapers handle packets that arrive to a full reshaping buffer. If such cases are to be handled by forwarding a packet as best effort, reshaping points may wish to forward a packet from the front of the reshaping queue, in order to minimize packet reordering problems at the receiver(s).

Ordering and Merging

TSpec's are ordered according to the following rule: TSpec A is a substitute ("as good or better than") for TSpec B if

(1) both the token bucket depth and rate for TSpec A are greater than or equal to those of TSpec B,

(2) the minimum policed unit m is at least as small for TSpec A as it is for TSpec B,

(3) the maximum packet size M is at least as large for TSpec A as it is for TSpec B,

(4) the peak rate p is at least as large in TSpec A as it is in TSpec B.

A merged TSpec may be calculated over a set of TSpec's by taking the largest token bucket rate, largest bucket size, largest peak rate, smallest minimal policed unit, and largest maximum packet size across all members of the set. This use of the word "merging" is similar to that in the RSVP protocol; a merged TSpec is one which is adequate to describe the traffic from any one of a number of flows.

RSpec's are merged in a similar manner as the TSpec's, i.e., a set of RSpec's is merged onto a single RSpec by taking the largest rate R of all RSpec's in the set.

NOTE: In case of shared reservations, i.e., a single RSpec whose rate, R, is to be shared between a number of flows, it is important to chose the requested rate R so as to ensure stable operation. Selection of the appropriate value is an application level decision, but two general cases may be considered. In the first case, the token bucket rates advertized by the senders sharing the reservation correspond to their average rate while active, e.g., the average rate for a voice signal when the speaker is talking. In such situations, the requested service rate R may be chosen to be significantly less than the sum of the token bucket rates of all the flows sharing the reservation. For example, this would apply to an audio conference call where only one speaker will typically be active at any time. In the second case, the token bucket rates advertized by the senders sharing the reservation correspond to their true long term average rate. In that case it is important that the requested service rate R be chosen larger than the sum of the token bucket rates of all the flows sharing the reservation.

Guidelines for Implementors

This section reviews two important implementation aspects of a Committed Rate service, that are closely related.

(1) The approximation of a dedicated rate circuit,

(2) The allocation of sufficient buffering to ensure minimal losses.

As mentioned in the section on Network Element Data Handling Requirement, support for the Committed Rate service requires that the network element approximates as well as possible the behavior of a dedicated rate circuit. This means, assuming a requested rate R, that whenever a flow is backlogged (has packets waiting to be transmitted) at a network element, it should ideally be able to transmit the packet at the head of its queue within at most L/R time units, where L is the size in bits of the packet. The ability of a network element to achieve this behavior depends on the type of scheduler it uses. While simple schedulers such as FIFO or priority queues may be used, it is highly recommended that an implementation of Committed Rate rely on a scheduler capable of providing service guarantees to individual connections. As mentioned in the section below on Examples of Use, there are a number of available schedulers that provide such capabilities.

In addition to choosing an appropriate scheduling algorithm, an implementation of the Committed Rate service at a network element also requires the use of an admission control algorithm. Admission control is required to ensure that the network element resources, i.e., bandwidth and buffers, are not over-committed. Admission control is to be performed based on the TSpec and RSpec specified by each flow requesting the Committed Rate service. The RSpec identifies the amount of bandwidth that the flow is requesting, and which should, therefore, at a minimum be available on the corresponding outgoing link from the network element. The exact amount of bandwidth to be allocated to the flow by the call admission algorithm depends on both the scheduler and the buffering scheme used. It is, therefore, a quantity to be engineered for each network element.

The admission control algorithm is also responsible for ensuring that sufficient buffer space is available to accommodate a new request. This is a function of the TSpec of the flow, in particular the peak rate, p, and the token bucket depth, b, but in general depends on a number of factors:

- The token bucket and peak rate parameters and the requested service rate R, i.e., how fast and for how long can the flow be sending at a rate exceeding the transmission rate R that has been allocated to it.

- The amount of statistical multiplexing that is expected at the network element, i.e., how likely is it that all flows simultaneously need the maximum possible amount of buffers.

- The perturbations introduced by schedulers at upstream network elements since the last reshaping point, i.e., how much bunching of packets is likely to have been introduced.

For example, assuming a scheduler that approximates reasonably closely the behavior of a dedicated rate circuit, e.g., a WFQ scheduler, a possible buffer allocation rule for a flow with given TSpec and committed rate R, is to ensure that the network element is able to buffer an amount of data of the order of b(p-R)/(p-r), where b is the token bucket depth, r the token bucket rate, and p the peak rate. Depending on the amount of statistical multiplexing expected at the network element, this does NOT necessarily imply that this amount of buffers has to be dedicated to the flow, i.e., buffer allocation is not necessarily additive in the number of flows even if this represents a simple and somewhat conservative rule.

However, note that the amount of buffering needed for flow at a network element also depends, to some extent, on the behavior of upstream (previous) network elements. For example, assuming that the scheduler at the previous network element did not approximate well the behavior of a dedicated rate circuit, e.g., could delay transmission of packets of a flow for a time period of duration T, the next (downstream) network element may then have to buffer the entire amount b+rT, if at the end of the period T this gets transmitted at at a speed much higher than the rate R. Hence, as mentioned before, even though the Committed Rate service specification does not mandate a particular type of scheduler, it encourages the use of schedulers that approximates as closely as possible a dedicated rate circuit, so as to minimize buffering requirements at downstream network elements.

Evaluation Criteria

The scheduling algorithm and admission control algorithm of the element must ensure that the requested committed rate is provided over some reasonably long time period, and that packets from a compliant flow are rarely lost.

The closer a network element approximates the behavior of a dedicated circuit at the requested committed rate, the better it performs in supporting the Committed Rate service.

This behavior can be evaluated by continuously sending packet into the network at the maximum possible rate allowed while remaining conformant, and by monitoring the delay experienced when traversing a series of network elements. The lower the average delay and its variations, i.e., difference between the maximum and minimum values, as experienced by the packets, the higher the evaluation ratings for the service. In addition, the smaller the value of the time period needed to transmit the amount of data corresponding to the Committed Rate, the higher the evaluation ratings for the service.

This behavior should also be consistently provided to flows accepted by the admission control algorithm, independently of the load levels at network elements. This should be tested by increasing the background best effort traffic on the network elements as well as by increasing, up to the maximum number allowed by the call admission algorithm of each network element, the number of Committed Rate flows being carried. The smaller the worst case values for the delay experienced by Committed Rate service flows across the range of load conditions, the higher the evaluation ratings for the service.

Additionally, users may want to evaluate, when applicable, the behavior of the Committed Rate service at a network element, when provided jointly with some other services whose more rigorous service requirements may affect the level of service given to Committed Rate flows. For example, this may apply to network elements that support both the Guaranteed Service and the Committed Rate service. Evaluation of this behavior can be achieved by loading the network element with a "test" Committed Rate flow and the maximum possible amount of Guaranteed Service traffic that the network element(s) can accept. The delays experienced by the Committed Rate flow should then be compared to those experienced in the other configurations described above. As before, the smaller the delay values, the higher the evaluation ratings for the service.

Examples of Implementation

Several scheduling algorithms and implementations exist that allow a close approximation of a dedicated rate circuit. They include Weighted Fair Queueing (WFQ) [4], Virtual Clock [5], Rate Controlled Service Disciplines [6,7], etc. Additional theoretical results positioning these results in the context of broader classes of algorithms can be found in [8,9].

Examples of Use

Consider an application that requires a specific guaranteed throughput in order to operate properly, but is reasonably tolerant in terms of the delay and delay variations it will experience, so that the delay guarantees of the Guaranteed service may not be warranted. For example, this may consist of an application retrieving a large document including graphics and pictures from a web server. This application wants a large enough rate to ensure that the document is retrieved reasonably fast, but does not require tight delay guarantees as the user will typically start browsing the initial material received and can, therefore, tolerate delay variations in receiving the remainder of the data. Another example, is that of a transaction based application that requires the transfer of reasonably large amounts of data in sufficiently timely fashion to ensure an adequate response time. Such an application will be satisfied by the provision of a large enough transmission rate, but again does not need very tight delay guarantees.

Security Considerations

Security considerations are not discussed in this memo.

Acknowledgments

The authors would like to gratefully acknowledge the help of the INT-SERV working group and the many contributions to its mailing list. In addition, they would like to acknowledge the Guaranteed Service specifications which served as a base for many of the aspects discussed in this draft.

References

[1] S. Shenker and J. Wroclawski. "Network Element Service Specification Template," Internet Draft, June 1995, <draft-ietf-intserv-svc-template-01.txt>

[2] S. Shenker and C. Partridge. "Specification of Guaranteed Quality of Service," Internet Draft, November 1995, <draft-ietf-intserv-guaranteed-svc-03.txt>

[3] J. Wroclawski. "Specification of The Controlled-Load Network Element Service," Internet Draft, November 1995, <draft-ietf-intserv-ctrl-load-svc-01.txt>

[4] A. Demers, S. Keshav and S. Shenker, "Analysis and Simulation of a Fair Queueing Algorithm," in Internetworking: Research and Experience, Vol 1, No. 1., pp. 3-26.

[5] L. Zhang, "Virtual Clock: A New Traffic Control Algorithm for Packet Switching Networks," in Proc. ACM SIGCOMM'90, pp. 19-29.

[6] H. Zhang, and D. Ferrari, "Rate-Controlled Service Disciplines," Journal of High Speed Networks, 3(4):389--412, 1994.

[7] L. Georgiadis, R. Guerin, V. Peris, and K. N. Sivaraja, "Efficient Network QoS Provisioning Based on per Node Traffic Shaping," IEEE/ACM Transactions on Networking, 4(4), August 1996.

[8] D. Stiliadis and A. Varma, "Latency-Rate Servers: A General Model for Analysis of Traffic Scheduling Algorithms," on Proc. INFOCOM'96, pp. 111-119.

[9] P. Goyal, S.S. Lam and H.M. Vin, "Determining End-to-End Delay Bounds in Heterogeneous Networks," in Proc. 5th Intl. Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.

[10] B. Braden, L. Zhang, S. Berson, S. Herzog, and J. Wroclawski. "Resource Reservation Protocol (RSVP) - Version 1 Functional Specification," Internet Draft, November 1995, <draft-ietf-rsvp-spec-08.txt>

[11] J. Wroclawski. "Standard Data Encoding for Integrated Services Objects," Internet Draft, November 1995, <draft-ietf-intserv-data-encoding-01.txt>

[12] S. Shenker. "Specification of General Characterization Parameters," Internet Draft, November 1995, <draft-ietf-intserv-charac-00.txt>

Authors' Address:

Fred Baker
Cisco Systems
519 Lado Drive
Santa Barbara, California 93111
fred@cisco.com
VOICE +1 408 526-4257
FAX +1 805 681-0115

Roch Guerin
IBM T.J. Watson research Center
P.O. Box 704
Yorktown Heights, NY 10598
guerin@watson.ibm.com
VOICE +1 914 784-7038
FAX +1 914 784-6318

Dilip Kandlur
IBM T.J. Watson research Center
P.O. Box 704
Yorktown Heights, NY 10598
kandlur@watson.ibm.com
VOICE +1 914 784-7722
FAX +1 914 784-6625