

Network Working Group
Request for Comments: 2815
Category: Standards Track

M. Seaman
Telseon
A. Smith
Extreme Networks
E. Crawley
Unisphere Solutions
J. Wroclawski
MIT LCS
May 2000

Integrated Service Mappings on IEEE 802 Networks

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This document describes mappings of IETF Integrated Services over LANs built from IEEE 802 network segments which may be interconnected by IEEE 802.1D MAC Bridges (switches). It describes parameter mappings for supporting Controlled Load and Guaranteed Service using the inherent capabilities of relevant IEEE 802 technologies and, in particular, 802.1D-1998 queuing features in switches.

These mappings are one component of the Integrated Services over IEEE 802 LANs framework.

Table of Contents

1	Introduction	2
2	Flow Identification and Traffic Class Selection	3
3	Choosing a flow's IEEE 802 user_priority class	5
3.1	Context of admission control and delay bounds	6
3.2	Default service mappings	7
3.3	Discussion	9
4	Computation of integrated services characterization parameters by IEEE 802 devices	10
4.1	General characterization parameters	10
4.2	Parameters to implement Guaranteed Service	11
4.3	Parameters to implement Controlled Load	11
4.4	Parameters to implement Best Effort	12
5	Merging of RSVP/SBM objects	12
6	Applicability of these service mappings	13
7	References	14
8	Security Considerations	15
9	Acknowledgments	15
10	Authors' Addresses	16
11	Full Copyright Statement	17

1. Introduction

The IEEE 802.1 Interworking Task Group has developed a set of enhancements to the basic MAC Service provided in Bridged Local Area Networks (a.k.a. "switched LANs"). As a supplement to the original IEEE MAC Bridges standard, IEEE 802.1D-1990 [802.1D-ORIG], the updated IEEE 802.1D-1998 [802.1D] proposes differential traffic class queuing in switches. The IEEE 802.1Q specification [802.1Q] extends the capabilities of Ethernet/802.3 media to carry a traffic class indicator, or "user_priority" field, within data frames.

The availability of this differential traffic queuing, together with additional mechanisms to provide admission control and signaling, allows IEEE 802 networks to support a close approximation of the IETF Integrated Services capabilities [CL][GS]. This document describes methods for mapping the service classes and parameters of the IETF model into IEEE 802.1D network parameters. A companion document [SBM] describes a signaling protocol for use with these mappings. It is recommended that readers be familiar with the overall framework in which these mappings and signaling protocol are expected to be used; this framework is described fully in [IS802FRAME].

Within this document, Section 2 describes the method by which end systems and routers bordering the IEEE Layer-2 cloud learn what traffic class should be used for each data flow's packets. Section 3 describes the approach recommended to map IP-level traffic flows to

IEEE traffic classes within the Layer 2 network. Section 4 describes the computation of Characterization Parameters by the layer 2 network. The remaining sections discuss some particular issues with the use of the RSVP/SBM signaling protocols, and describe the applicability of all of the above to different layer 2 network topologies.

2. Flow Identification and Traffic Class Selection

One model for supporting integrated services over specific link layers treats layer-2 devices very much as a special case of routers. In this model, switches and other devices along the data path make packet handling decisions based on the RSVP flow and filter specifications, and use these specifications to classify the corresponding data packets. The specifications could either be used directly, or could be used indirectly by mapping each RSVP session onto a layer-2 construct such as an ATM virtual circuit.

This approach is inappropriate for use in the IEEE 802 environment. Filtering to the per-flow level becomes expensive with increasing switch speed; devices with such filtering capabilities are likely to have a very similar implementation complexity to IP routers, and may not make use of simpler mechanisms such as 802.1D user priority.

The Integrated Services over IEEE 802 LANs framework [IS802FRAME] and this document use an "aggregated flow" approach based on use of layer-2 traffic classes. In this model, each arriving flow is assigned to one of the available classes for the duration of the flow and traverses the 802 cloud in this class. Traffic flows requiring similar service are grouped together into a single class, while the system's admission control and class selection rules ensure that the service requirements for flows in each of the classes are met. In many situations this is a viable intermediate point between no QoS control and full router-type integrated services. The approach can work effectively even with switches implementing only the simplest differential traffic classification capability specified in the 802.1D model. In the aggregated flow model, traffic arriving at the boundary of a layer-2 cloud is tagged by the boundary device (end host or border router) with an appropriate traffic class, represented as an 802.1D "user_priority" value. Two fundamental questions are "who determines the correspondence between IP-level traffic flows and link-level classes?" and "how is this correspondence conveyed to the boundary devices that must mark the data frames?"

One approach to answering these questions would be for the meanings of the classes to be universally defined. This document would then standardize the meanings of a set of classes; e.g., 1 = best effort, 2 = 100 ms peak delay target, 3 = 10 ms peak delay target, 4 = 1 ms

peak delay target, etc. The meanings of these universally defined classes could then be encoded directly in end stations, and the flow-to-class mappings computed directly in these devices.

This universal definition approach would be simple to implement, but is too rigid to map the wide range of possible user requirements onto the limited number of available 802.1D classes. The model described in [IS802FRAME] uses a more flexible mapping: clients ask "the network" which user_priority traffic class to use for a given traffic flow, as categorized by its flow-spec and layer-2 endpoints. The network provides a value back to the requester that is appropriate considering the current network topology, load conditions, other admitted flows, etc. The task of configuring switches with this mapping (e.g., through network management, a switch-switch protocol or via some network-wide QoS-mapping directory service) is an order of magnitude less complex than performing the same function in end stations. Also, when new services (or other network reconfigurations) are added to such a network, the network elements will typically be the ones to be upgraded with new queuing algorithms etc. and can be provided with new mappings at this time.

In the current model it is assumed that all data packets of a flow are assigned to the same traffic class for the duration of the flow: the characteristics of the MAC service, as defined by Clause 6 of [802.1D], then ensure the ordering of the data packets of the flow between adjacent Layer 3 routers. This is usually desirable to avoid potential re-ordering problems as discussed in [IS802FRAME] and [CL]. Note that there are some scenarios where it might be desirable to send conforming data traffic in one traffic class and non-conforming traffic for the same flow in a different, lower traffic class: such a division into separate traffic classes is for future study. When a new session or "flow" requiring QoS support is created, a client must ask "the network" which traffic class (IEEE 802 user_priority) to use for a given traffic flow, so that it can label the packets of the flow as it places them into the network. A request/response protocol is needed between client and network to return this information. The request can be piggy-backed onto an admission control request and the response can be piggy-backed onto an admission control acknowledgment. This "one pass" assignment has the benefit of completing the admission control transaction in a timely way and reducing the exposure to changing conditions that could occur if clients cached the knowledge for extensive periods. A set of extensions to the RSVP protocol for communicating this information have been defined [SBM].

The network (i.e., the first network element encountered downstream from the client) must then answer the following questions:

1. Which of the available traffic classes would be appropriate for this flow?

In general, a newly arriving flow might be assigned to a number of classes. For example, if 10ms of delay is acceptable, the flow could potentially be assigned to either a 10ms delay class or a 1ms delay class. This packing problem is quite difficult to solve if the target parameters of the classes are allowed to change dynamically as flows arrive and depart. It is quite simple if the target parameters of each class is held fixed, and the class table is simply searched to find a class appropriate for the arriving flow. This document adopts the latter approach.

2. Of the appropriate traffic classes, which if any have enough capacity available to accept the new flow?

This is the admission control problem. It is necessary to compare the level of traffic currently assigned to each class with the available level of network resources (bandwidth, buffers, etc), to ensure that adding the new flow to the class will not cause the class's performance to go below its target values. This problem is compounded because in a priority queuing system adding traffic to a higher-priority class can affect the performance of lower-priority classes. The admission control algorithm for a system using the default 802 priority behavior must be reasonably sophisticated to provide acceptable results.

If an acceptable class is found, the network returns the chosen `user_priority` value to the client.

Note that the client may be an end station, a router at the edge of the layer 2 network, or a first switch acting as a proxy for a device that does not participate in these protocols for whatever reason. Note also that a device e.g., a server or router may choose to implement both the "client" as well as the "network" portion of this model so that it can select its own `user_priority` values. Such an implementation would generally be discouraged unless the device has a close tie-in with the network topology and resource allocation policies. It may, however, work acceptably in cases where there is known over-provisioning of resources.

3. Choosing a flow's IEEE 802 `user_priority` class

This section describes the method by which IP-level flows are mapped into appropriate IEEE `user_priority` classes. The IP-level services considered are Best Effort, Controlled Load, and Guaranteed Service.

The major issue is that admission control requests and application requirements are specified in terms of a multidimensional vector of parameters e.g., bandwidth, delay, jitter, service class. This multidimensional space must be mapped onto a set of traffic classes whose default behavior in L2 switches is unidimensional (i.e., strict priority default queuing). This priority queuing alone can provide only relative ordering between traffic classes. It can neither enforce an absolute (quantifiable) delay bound for a traffic class, nor can it discriminate amongst Int-Serv flows within the aggregate in a traffic class. Therefore, it cannot provide the absolute control of packet loss and delay required for individual Int-Serv flows.

To provide absolute control of loss and delay three things must occur:

- (1) The amount of bandwidth available to the QoS-controlled flows must be known, and the number of flows admitted to the network (allowed to use the bandwidth) must be limited.
- (2) A traffic scheduling mechanism is needed to give preferential service to flows with lower delay targets.
- (3) Some mechanism must ensure that best-effort flows and QoS controlled flows that are exceeding their Tspecs do not damage the quality of service delivered to in-Tspec QoS controlled flows. This mechanism could be part of the traffic scheduler, or it could be a separate policing mechanism.

For IEEE 802 networks, the first function (admission control) is provided by a Subnet Bandwidth Manager, as discussed below. We use the link-level `user_priority` mechanism at each switch and bridge to implement the second function (preferential service to flows with lower delay targets). Because a simple priority scheduler cannot provide policing (function three), policing for IEEE networks is generally implemented at the edge of the network by a layer-3 device. When this policing is performed only at the edges of the network it is of necessity approximate. This issue is discussed further in [IS802FRAME].

3.1. Context of admission control and delay bounds

As described above, it is the combination of priority-based scheduling and admission control that creates quantified delay bounds. Thus, any attempt to quantify the delay bounds expected by a given traffic class has to be made in the context of the admission control elements. Section 6 of the framework [IS802FRAME] provides for two different models of admission control - centralized or distributed Bandwidth Allocators.

It is important to note that in this approach it is the admission control algorithm that determines which of the Int-Serv services is being offered. Given a set of priority classes with delay targets, a relatively simple admission control algorithm can place flows into classes so that the bandwidth and delay behavior experienced by each flow corresponds to the requirements of the Controlled-Load service, but cannot offer the higher assurance of the Guaranteed service. To offer the Guaranteed service, the admission control algorithm must be much more stringent in its allocation of resources, and must also compute the C and D error terms required of this service.

A delay bound can only be realized at the admission control element itself so any delay numbers attached to a traffic class represent the delay that a single element can allow for. That element may represent a whole L2 domain or just a single L2 segment.

With either admission control model, the delay bound has no scope outside of a L2 domain. The only requirement is that it be understood by all Bandwidth Allocators in the L2 domain and, for example, be exported as C and D terms to L3 devices implementing the Guaranteed Service. Thus, the end-to-end delay experienced by a flow can only be characterized by summing along the path using the usual RSVP mechanisms.

3.2. Default service mappings

Table 1 presents the default mapping from delay targets to IEEE 802.1 user_priority classes. However, these mappings must be viewed as defaults, and must be changeable.

In order to simplify the task of changing mappings, this mapping table is held by *switches* (and routers if desired) but generally not by end-station hosts. It is a read-write table. The values proposed below are defaults and can be overridden by management control so long as all switches agree to some extent (the required level of agreement requires further analysis).

In future networks this mapping table might be adjusted dynamically and without human intervention. It is possible that some form of network-wide lookup service could be implemented that serviced requests from clients e.g., `traffic_class = getQoSbyname("H.323 video")` and notified switches of what traffic categories they were likely to encounter and how to allocate those requests into traffic classes. Alternatively, the network's admission control mechanisms might directly adjust the mapping table to maximize the utilization of network resources. Such mechanisms are for further study.

The delay bounds numbers proposed in Table 1 are for per-Bandwidth Allocator element delay targets and are derived from a subjective analysis of the needs of typical delay-sensitive applications e.g., voice, video. See Annex H of [802.1D] for further discussion of the selection of these values. Although these values appear to address the needs of current video and voice technology, it should be noted that there is no requirement to adhere to these values and no dependence of IEEE 802.1 on these values.

user_priority	Service
0	Default, assumed to be Best Effort
1	reserved, "less than" Best Effort
2	reserved
3	reserved
4	Delay Sensitive, no bound
5	Delay Sensitive, 100ms bound
6	Delay Sensitive, 10ms bound
7	Network Control

Table 1 - Example user_priority to service mappings

Note: These mappings are believed to be useful defaults but further implementation and usage experience is required. The mappings may be refined in future editions of this document.

With this example set of mappings, delay-sensitive, admission controlled traffic flows are mapped to user_priority values in ascending order of their delay bound requirement. Note that the bounds are targets only - see [IS802FRAME] for a discussion of the effects of other non-conformant flows on delay bounds of other flows. Only by applying admission control to higher-priority classes can any promises be made to lower-priority classes.

This set of mappings also leaves several classes as reserved for future definition.

Note: this mapping does not dictate what mechanisms or algorithms a network element (e.g., an Ethernet switch) must perform to implement these mappings: this is an implementation choice and does not matter so long as the requirements for the particular service model are met.

Note: these mappings apply primarily to networks constructed from devices that implement the priority-scheduling behavior defined as the default in 802.1D. Some devices may implement more complex scheduling behaviors not based only on priority. In that circumstance these mappings might still be used, but other, more

specialized mappings may be more appropriate.

3.3. Discussion

The recommendation of classes 4, 5 and 6 for Delay Sensitive, Admission Controlled flows is somewhat arbitrary; any classes with priorities greater than that assigned to Best Effort can be used. Those proposed here have the advantage that, for transit through 802.1D switches with only two-level strict priority queuing, all delay-sensitive traffic gets "high priority" treatment (the 802.1D default split is 0-3 and 4-7 for a device with 2 queues).

The choice of the delay bound targets is tuned to an average expected application mix, and might be retuned by a network manager facing a widely different mix of user needs. The choice is potentially very significant: wise choice can lead to a much more efficient allocation of resources as well as greater (though still not very good) isolation between flows.

Placing Network Control traffic at class 7 is necessary to protect important traffic such as route updates and network management. Unfortunately, placing this traffic higher in the `user_priority` ordering causes it to have a direct effect on the ability of devices to provide assurances to QoS controlled application traffic. Therefore, an estimate of the amount of Network Control traffic must be made by any device that is performing admission control (e.g., SBMs). This would be in terms of the parameters that are normally taken into account by the admission control algorithm. This estimate should be used in the admission control decisions for the lower classes (the estimate is likely to be a configuration parameter of SBMs).

A traffic class such as class 1 for "less than best effort" might be useful for devices that wish to dynamically "penalty tag" all of the data of flows that are presently exceeding their allocation or `Tspec`. This provides a way to isolate flows that are exceeding their service limits from flows that are not, to avoid reducing the QoS delivered to flows that are within their contract. Data from such tagged flows might also be preferentially discarded by an overloaded downstream device.

A somewhat simpler approach would be to tag only the portion of a flow's packets that actually exceed the `Tspec` at any given instant as low priority. However, it is often considered to be a bad idea to treat flows in this way as it will likely cause significant re-ordering of the flow's packets, which is not desirable. Note that the default 802.1D treatment of `user_priorities` 1 and 2 is "less than" the default class 0.

4. Computation of integrated services characterization parameters by IEEE 802 devices

The integrated service model requires that each network element that supports integrated services compute and make available certain "characterization parameters" describing the element's behavior. These parameters may be either generally applicable or specific to a particular QoS control service. These parameters may be computed by calculation, measurement, or estimation. When a network element cannot compute its own parameters (for example, a simple link), we assume that the device sending onto or receiving data from the link will compute the link's parameters as well as its own. The accuracy of calculation of these parameters may not be very critical; in some cases loose estimates are all that is required to provide a useful service. This is important in the IEEE 802 case, where it will be virtually impossible to compute parameters accurately for certain topologies and switch technologies. Indeed, it is an assumption of the use of this model by relatively simple switches (see [IS802FRAME] for a discussion of the different types of switch functionality that might be expected) that they merely provide values to describe the device and admit flows conservatively. The discussion below presents a general outline for the computation of these parameters, and points out some cases where the parameters must be computed accurately. Further specification of how to export these parameters is for further study.

4.1. General characterization parameters

There are some general parameters [GENCHAR] that a device will need to use and/or supply for all service types:

- * Ingress link
- * Egress links and their MTUs, framing overheads and minimum packet sizes (see media-specific information presented above).
- * Available path bandwidth: updated hop-by-hop by any device along the path of the flow.
- * Minimum latency

Of these parameters, the MTU and minimum packet size information must be reported accurately. Also, the "break bits" must be set correctly, both the overall bit that indicates the existence of QoS control support and the individual bits that specify support for a particular scheduling service. The available bandwidth should be reported as accurately as possible, but very loose estimates are acceptable. The minimum latency parameter should be determined and reported as

accurately as possible if the element offers Guaranteed service, but may be loosely estimated or reported as zero if the element offers only Controlled-Load service.

4.2. Parameters to implement Guaranteed Service

A network element supporting the Guaranteed Service [GS] must be able to determine the following parameters:

- * Constant delay bound through this device (in addition to any value provided by "minimum latency" above) and up to the receiver at the next network element for the packets of this flow if it were to be admitted. This includes any access latency bound to the outgoing link as well as propagation delay across that link. This value is advertised as the 'C' parameter of the Guaranteed Service.
- * Rate-proportional delay bound through this device and up to the receiver at the next network element for the packets of this flow if it were to be admitted. This value is advertised as the 'D' parameter of the Guaranteed Service.
- * Receive resources that would need to be associated with this flow (e.g., buffering, bandwidth) if it were to be admitted and not suffer packet loss if it kept within its supplied Tspec/Rspec. These values are used by the admission control algorithm to decide whether a new flow can be accepted by the device.
- * Transmit resources that would need to be associated with this flow (e.g., buffering, bandwidth, constant- and rate-proportional delay bounds) if it were to be admitted. These values are used by the admission control algorithm to decide whether a new flow can be accepted by the device.

The exported characterization parameters for this service should be reported as accurately as possible. If estimations or approximations are used, they should err in whatever direction causes the user to receive better performance than requested. For example, the C and D error terms should overestimate delay, rather than underestimate it.

4.3. Parameters to implement Controlled Load

A network element implementing the Controlled Load service [CL] must be able to determine the following:

- * Receive resources that would need to be associated with this flow (e.g., buffering) if it were to be admitted. These values are used by the admission control algorithm to decide whether a new flow can be accepted by the device.

- * Transmit resources that would need to be associated with this flow (e.g., buffering) if it were to be admitted. These values are used by the admission control algorithm to decide whether a new flow can be accepted by the device.

The Controlled Load service does not export any service-specific characterization parameters. Internal resource allocation estimates should ensure that the service quality remains high when considering the statistical aggregation of Controlled Load flows into 802 traffic classes.

4.4. Parameters to implement Best Effort

For a network element that implements only best effort service there are no explicit parameters that need to be characterized. Note that an integrated services aware network element that implements only best effort service will set the "break bit" described in [RSVPINTSERV].

5. Merging of RSVP/SBM objects

Where reservations that use the SBM protocol's TCLASS object [SBM] need to be merged, an algorithm needs to be defined that is consistent with the mappings to individual user_priority values in use in the Layer-2 cloud. A merged reservation must receive at least as good a service as the best of the component reservations.

There is no single merging rule that can prevent all of the following side-effects:

- * If a merger were to demote the existing branch of the flow into a higher-delay traffic class then this is a denial of service to the existing flow which would likely receive worse service than before.
- * If a merger were to promote the existing branch of the flow into a new, lower-delay, traffic class, this might then suffer either admission control failures or may cost more in some sense than the already-admitted flow. This can also be considered as a denial-of-service attack.
- * Promotion of the new branch may lead to rejection of the request because it has been re-assigned to a traffic class that has not enough resources to accommodate it.

Therefore, such a merger is declared to be illegal and the usual SBM admission control failure rules are applied. Traffic class selection is performed based on the TSpec information. When the first RESV for

a flow arrives, a traffic class is chosen based on the request, an SBM TCLASS object is inserted into the message and admission control for that traffic class is done by the SBM. Reservation succeeds or fails as usual.

When a second RESV for the same flow arrives at a different egress point of the Layer-2 cloud the process starts to repeat. Eventually the SBM-augmented RESV may hit a switch with an existing reservation in place for the flow i.e., an L2 branch point for the flow. If so, the traffic class chosen for the second reservation is checked against the first. If they are the same, the RESV requests are merged and passed on towards the sender(s).

If the second TCLASS would have been different, an RSVP/SBM ResvErr error is returned to the Layer-3 device that launched the second RESV request into the Layer-2 cloud. This device will then pass on the ResvErr to the original requester according to RSVP rules. Detailed processing rules are specified in [SBM].

6. Applicability of these service mappings

Switches using layer-2-only standards (e.g., 802.1D-1990, 802.1D-1998) need to inter-operate with routers and layer-3 switches. Wide deployment of such 802.1D-1998 switches will occur in a number of roles in the network: "desktop switches" provide dedicated 10/100 Mbps links to end stations and high speed core switches often act as central campus switching points for layer-3 devices. Layer-2 devices will have to operate in all of the following scenarios:

- * every device along a network path is layer-3 capable and intrusive into the full data stream
- * only the edge devices are pure layer-2
- * every alternate device lacks layer-3 functionality
- * most devices lack layer-3 functionality except for some key control points such as router firewalls, for example.

Where int-serv flows pass through equipment which does not support Integrated Services or 802.1D traffic management and which places all packets through the same queuing and overload-dropping paths, it is obvious that some of a flow's desired service parameters become more difficult to support. In particular, the two integrated service classes studied here, Controlled Load and Guaranteed Service, both assume that flows will be policed and kept "insulated" from misbehaving other flows or from best effort traffic during their passage through the network. This cannot be

done within an IEEE 802 network using devices with the default user_priority function; in this case policing must be approximated at the network edges.

In addition, in order to provide a Guaranteed Service, *all* switching elements along the path must participate in special treatment for packets in such flows: where there is a "break" in guaranteed service, all bets are off. Thus, a network path that includes even a single switch transmitting onto a shared or half-duplex LAN segment is unlikely to be able to provide a very good approximation to Guaranteed Service. For Controlled Load service, the requirements on the switches and link types are less stringent although it is still necessary to provide differential queuing and buffering in switches for CL flows over best effort in order to approximate CL service. Note that users receive indication of such breaks in the path through the "break bits" described in y [RSVPINTSERV]. These bits must be correctly set when IEEE 802 devices that cannot provide a specific service exist in a network.

Other approaches might be to pass more information between switches about the capabilities of their neighbours and to route around non-QoS-capable switches: such methods are for further study. And of course the easiest solution of all is to upgrade links and switches to higher capacities.

7. References

- [802.1D-ORIG] "MAC Bridges", ISO/IEC 10038, ANSI/IEEE Std 802.1D-1993
- [802.1D] "Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Common specifications - Part 3: Media Access Control (MAC) Bridges: Revision. This is a revision of ISO/IEC 10038: 1993, 802.1j-1992 and 802.6k-1992. It incorporates P802.11c, P802.1p and P802.12e." ISO/IEC 15802-3:1998"
- [INTSERV] Braden, R., Clark, D. and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.
- [RSVP] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource Reservation Protocol (RSVP) - Version 1 Functional Specification", RFC 2205, September 1997.
- [CL] Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, September 1997.

- [GS] Schenker, S., Partridge, C. and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212 September 1997.
- [802.1Q] ANSI/IEEE Standard 802.1Q-1998, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks", 1998.
- [GENCHAR] Shenker, S., and J. Wroclawski, "General Characterization Parameters for Integrated Service Network Elements", RFC 2215, September 1997.
- [IS802FRAME] Ghanwani, A., Pace, W., Srinivasan, V., Smith, A. and M. Seaman, "A Framework for Providing Integrated Services Over Shared and Switched LAN Technologies", RFC 2816, May 2000.
- [SBM] Yavatkar, R., Hoffman, D., Bernet, Y., Baker, F. and M. Speer, "SBM (Subnet Bandwidth Manager): A Protocol for Admission Control over IEEE 802-style Networks", RFC 2814, May 2000.
- [RSVPINTSERV] Wroclawski, J., "The use of RSVP with IETF Integrated Services", RFC 2210, September 1997.
- [PROCESS] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

8. Security Considerations

Any use of QoS requires examination of security considerations because it leaves the possibility open for denial of service or theft of service attacks. This document introduces no new security issues on top of those discussed in the companion ISSLL documents [IS802FRAME] and [SBM]. Any use of these service mappings assumes that all requests for service are authenticated appropriately.

9. Acknowledgments

This document draws heavily on the work of the ISSLL WG of the IETF and the IEEE P802.1 Interworking Task Group.

10. Authors' Addresses

Mick Seaman
Telseon
480 S. California Ave
Palo Alto, CA 94306
USA

Email: mick@telseon.com

Andrew Smith
Extreme Networks
3585 Monroe St.
Santa Clara, CA 95051
USA

Phone: +1 408 579 2821
EMail: andrew@extremenetworks.com

Eric Crawley
Unisphere Solutions
5 Carlisle Rd.
Westford, MA 01886

Phone: +1 978 692 1999
Email: esc@unispheresolutions.com

John Wroclawski
MIT Laboratory for Computer Science
545 Technology Sq.
Cambridge, MA 02139
USA

Phone: +1 617 253 7885
EMail: jtw@lcs.mit.edu

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.