Transport Working Group	F. Baker	
Internet-Draft	J. Polk	
Updates: 4542, 4594	Cisco Systems	
(if approved)	M. Dolly	
Intended status: Standards Track	AT&T Labs	
Expires: May 7, 2009	November 03, 2008	

DSCP for Capacity-Admitted Traffic
draft-ietf-tsvwg-admitted-realtime-dscp-05

Status of this Memo

Abstract

This document requests one Differentiated Services Code Point (DSCP) from the Internet Assigned Numbers Authority (IANA) for real-time traffic classes similar to voice conforming to the Expedited Forwarding Per Hop Behavior, and admitted using a call admission procedure involving authentication, authorization, and capacity admission.
This document also recommends that certain classes of video traffic described in RFC 4594 and which have similar requirements be changed to require admission using a Call Admission Control (CAC) procedure involving authentication, authorization, and capacity admission.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.) [RFC2119].

Table of Contents

---

1.   Introduction   TOC

This document requests one Differentiated Services Code Point (DSCP) from the Internet Assigned Numbers Authority (IANA) for a class of real-time traffic. This class conforms to the Expedited Forwarding (Davie, B., Charny, A., Bennet, J., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)," March 2002.) [RFC3246] [RFC3247] (Charny, A., Bennet, J., Benson, K., Boudec, J., Chiu, A., Courtney, W., Davari, S., Firoiu, V., Kalmanek, C., and K. Ramakrishnan, "Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior)," March 2002.) Per Hop Behavior. It is also admitted using a CAC procedure involving authentication, authorization, and capacity admission. This differs from a real-time traffic class conforming to the Expedited Forwarding

Per Hop Behavior but not subject to capacity admission or subject to very coarse capacity admission.

It also recommends that certain classes of video described in [RFC4594] (Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes," August 2006.) be treated as requiring capacity admission as well.

These applications have one or more potential congestion points between the endpoints. Reserving capacity for them is important to application performance. All of these applications have low tolerance to jitter (aka delay variation) and loss, as summarized in Section 2 (Candidate Implementations of the Admitted Telephony Service Class), and most (except for multimedia conferencing) have inelastic flow behavior from Figure 1 of [RFC4594] (Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes," August 2006.). Inelastic flow behavior and low jitter/loss tolerance are the service characteristics that define the need for admission control behavior.

One of the reasons behind this is the need for classes of traffic that are handled under special policies, such as the non-preemptive Emergency Telecommunication Service, the US Department of Defense's Assured Service (which is similar to Multi-Level Precedence and Preemption (International Telecommunications Union, "Multilevel Precedence and Preemption Service," 1990.) [ITU.MLPP.1990] procedure), or e-911, in addition to normal routine calls that use call admission. It is possible to use control plane protocols to generally restrict session admission such that admitted traffic should receive the desired service, and the policy (e.g. Routine, National Security or Emergency Preparedness [NS/EP] communications, e-911, etc) need not be signaled in a DSCP. However, service providers need to distinguish between special-policy traffic and other classes, particularly the existing VoIP services that perform no capacity admission or only very coarse capacity admission and can exceed their allocated resources.

The requested DSCP applies to the Telephony Service Class described in [RFC4594] (Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes," August 2006.).

Since video classes have not had the history of mixing admitted and non-admitted traffic in the same Per-Hop Behavior (PHB) as has occurred for EF, an additional DSCP code point is not recommended. Instead, the recommended "best practice" is to perform admission control for all traffic in three of [RFC4594] (Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes," August 2006.)'s video classes: the

   *Interactive Real-Time Traffic (CS4, used for Video conferencing
    and Interactive gaming),

   *Broadcast TV (CS3) for use in a video on demand context, and

   *AF4 Multimedia Conferencing (video conferencing).

Other video classes are believed to not have the current problem of
confusion with unadmitted traffic and therefore would not benefit from
the notion of a separate DSCP for admitted traffic. Within an ISP and
on inter-ISP links (i.e. within networks whose internal paths are
uniform at hundreds of megabits or faster), one would expect all of
this traffic to be carried in the Real Time Traffic Class described in
[RFC5127] (Chan, K., Babiarz, J., and F. Baker, "Aggregation of
DiffServ Service Classes," February 2008.).

1.1. Definitions

The following terms and acronyms are used in this document.

PHB: A Per-Hop-Behavior (PHB) is the externally observable
forwarding behavior applied at a Differentiated Services
compliant node to a DS behavior aggregate [RFC2475] (Blake, S.,
Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An
Architecture for Differentiated Services," December 1998.). It
may be thought of as a program configured on the interface of an
Internet host or router, specified drop probabilities, queuing
priorities or rates, and other handling characteristics for the
traffic class.

DSCP: The Differentiated Services Code Point (DSCP), as defined in
[RFC2474] (Nichols, K., Blake, S., Baker, F., and D. Black,
"Definition of the Differentiated Services Field (DS Field) in
the IPv4 and IPv6 Headers," December 1998.), is a value which is
encoded in the DS field, and which each DS Node MUST use to
select the PHB which is to be experienced by each packet it
forwards [RFC3260] (Grossman, D., "New Terminology and
Clarifications for Diffserv," April 2002.). It is a 6-bit number
embedded into the 8-bit TOS field of an IPv4 datagram or the
Traffic Class field of an IPv6 datagram.

CAC: Call Admission Control includes concepts of authorization and
capacity admission. "Authorization" refers to any procedure that
identifies a user, verifies the authenticity of the
identification, and determines whether the user is authorized to
use the service under the relevant policy. "Capacity Admission"
refers to any procedure that determines whether capacity exists
supporting a session's requirements under some policy.
In the Internet, these are separate functions, while in the PSTN
they and call routing are carried out together.

UNI: A User/Network Interface (UNI) is the interface (often a
physical link or its virtual equivalent) that connects two
entities that do not trust each other, and in which one (the

user) purchases connectivity services from the other (the network). Figure 1 (UNI and NNI interfaces) shows two user networks connected by what appears to each of them to be a single network ("The Internet", access to which is provided by their service provider) that provides connectivity services to other users.
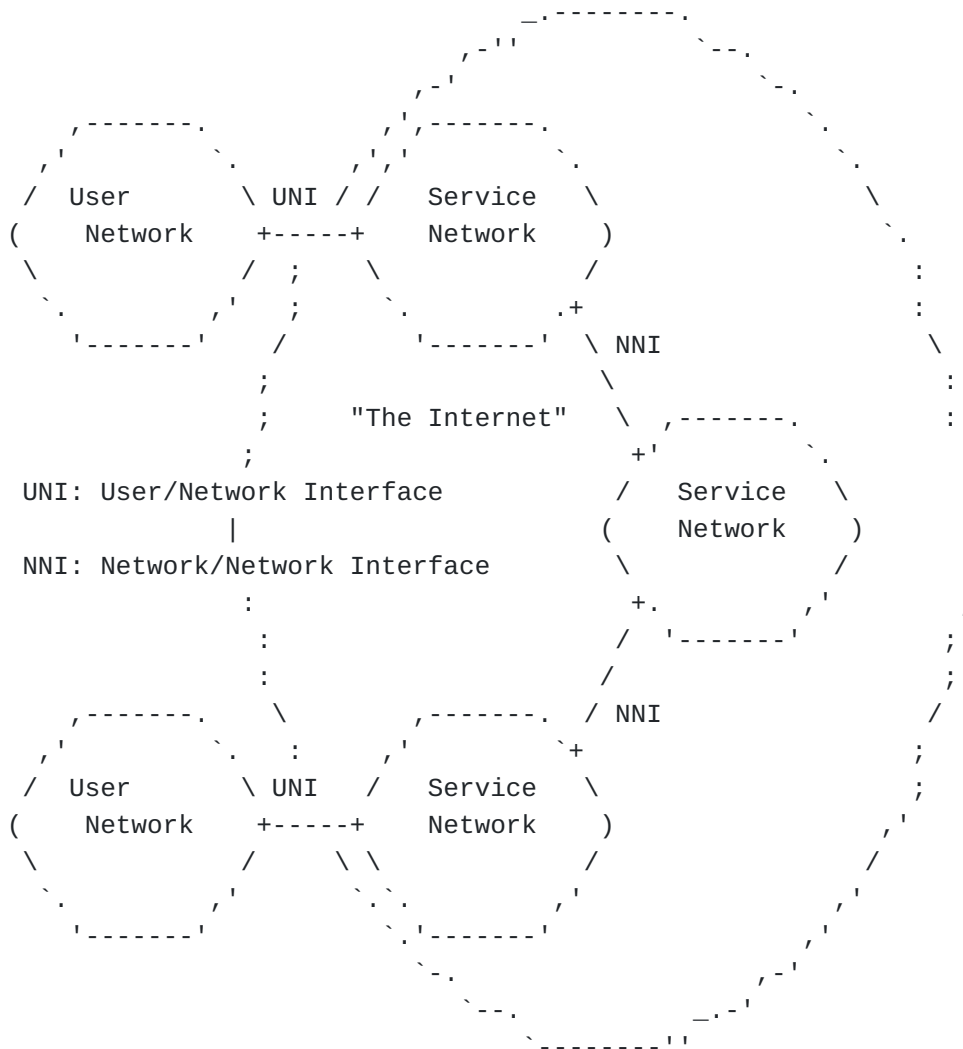
UNIs tend to be the bottlenecks in the Internet, where users purchase relatively low amounts of bandwidth for cost or service reasons, and as a result are most subject to congestion issues and therefore issues requiring traffic conditioning and service prioritization.

NNI: A Network/Network Interface (NNI) is the interface (often a physical link or its virtual equivalent) that connects two entities that trust each other within limits, and in which the two are seen as trading services for value. Figure 1 (UNI and NNI interfaces) shows three service networks that together provide the connectivity services that we call "the Internet". They are different administrations and are very probably in competition, but exchange contracts for connectivity and capacity that enable them to offer specific services to their customers.

NNIs may not be bottlenecks in the Internet if service providers contractually agree to provision excess capacity at them, as they commonly do. However, NNI performance may differ by ISP, and the performance guarantee interval may range from a month to a much shorter period. Furthermore, a peering point NNI may not have contractual performance guarantees or may become overloaded under certain conditions. They are also policy-controlled interfaces, especially in BGP. As a result, they may require traffic prioritization policy.

Queue: There are multiple ways to build a multi-queue scheduler. Weighted Round Robin (WRR) literally builds multiple lists and visits them in a specified order, while a calendar queue (often used to implement Weighted Fair Queuing, or WFQ) builds a list for each time interval and enqueues at most a stated amount of data in each such list for transmission during that time interval. While these differ dramatically in implementation, the external difference in behavior is generally negligible when they are properly configured. Consistent with the definitions used in the Differentiated Services Architecture (Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services," December 1998.) [RFC2475], these are treated as equivalent in this document, and the lists of WRR and the classes of a calendar queue will be referred to uniformly as "queues".

```
                              _.---------.
                          ,-''            `--.
                        ,-'                    `-.
                      ,-'                          `.
          ,-------.          ,',-------.              `.
        ,'         `.      ,','         `.              `.
      ,'   User       \ UNI / /   Service   \             \
     (    Network     +-----+     Network    )            `.
      \              /  ;    \               /             :
       `.          ,'   ;     `.           .+              :
        '-------'     /        '-------'  \ NNI             \
                     ;                     \                :
                     ;       "The Internet"  \  ,-------.    :
                     ;                        +'         `.   :
       UNI: User/Network Interface          /   Service   \   |
                     |                      (    Network    )   |
       NNI: Network/Network Interface        \             /    |
                     :                        +.         ,'     ;
                     :                        /  '-------'      ;
                     :                       /                 ;
          ,-------.      \        ,-------. / NNI             /
        ,'         `.     :     ,'         `+                ;
       /   User       \ UNI  /   Service    \                ;
      (    Network     +-----+    Network     )            ,'
       \              /    \ \               /             /
        `.          ,'      `.`.           ,'            ,'
          '-------'            `.'-------'             ,'
                                  `-.              ,-'
                                     `--.      _.-'
                                         `--------''
```

Figure 1: UNI and NNI interfaces

1.2. Problem

In short, the Telephony Service Class described in [RFC4594] (Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes," August 2006.) permits the use of capacity admission in implementing the service, but present implementations either provide no capacity admission services or do so in a manner that depends on specific traffic engineering. In the context of the Internet backbone, the two are essentially equivalent; the edge network depends on specific engineering by the service provider that may not be present, especially in a mobile environment.

However, services are being requested of the network that would specifically make use of capacity admission, and would distinguish among users or the uses of available Voice-over-IP or Video-over-IP capacity in various ways. Various agencies would like to provide services as described in section 2.6 of [RFC4504] (Sinnreich, H., Lass, S., and C. Stredicke, "SIP Telephony Device Requirements and Configuration," May 2006.) or in [RFC4190] (Carlberg, K., Brown, I., and C. Beard, "Framework for Supporting Emergency Telecommunications Service (ETS) in IP Telephony," November 2005.). This requires the use of capacity admission to differentiate among users (which might be 911 call centers, other offices with preferential service contracts, or individual users gaining access with special credentials) to provide services to them that are not afforded to non-capacity admitted customer-to-customer IP telephony sessions.

---

2. Candidate Implementations of the Admitted Telephony Service Class [TOC]

2.1. Potential implementations of EF in this model [TOC]

There are at least two possible ways to implement isolation between the Capacity Admitted PHB and the Expedited Forwarding PHB in this model. They are to implement separate classes as a set of
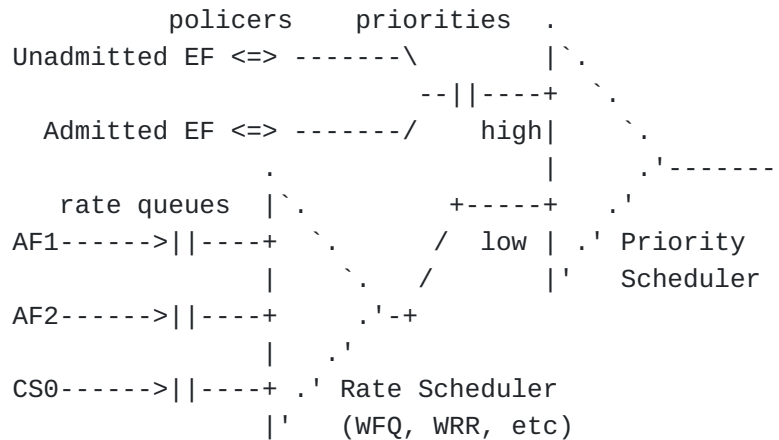
> *Multiple data plane traffic classes, each consisting of a policer and a queue, and the queues enjoying different priorities, or

> *Multiple data plane traffic classes, each consisting of a policer but feeding into a common queue or multiple queues at the same priority.

We will explain the difference, and describe in what way they differ in operation. The reason this is necessary is that there is current confusion in the industry.
The multi-priority model is shown in Figure 2 (Implementation as a data plane priority). In this model, traffic from each service class is placed into a separate priority queue. If data is present in both queues, traffic from one of them will always be selected for transmission. This has the effect of transferring jitter from the higher priority queue to the lower priority queue, and reordering traffic in a way that gives the higher priority traffic a smaller average queuing delay. Each queue must have its own policer, however, to protect the network from errors and attacks; if a traffic class

thinks it is carrying a certain data rate but an abuse sends
significantly more, the effect of simple prioritization would not
preserve the lower priorities of traffic, which could cause routing to
fail or otherwise impact an SLA.

```
                                          .
                    policers    priorities  |`.
         Unadmitted EF <=> ----------||----+  `.
                                       high|    `.
           Admitted EF <=> ----------||----+     .'-----------
                      .                 medium  .'
            rate queues  |`.          +-----+ .' Priority
         AF1------>||----+  `.      /  low |'   Scheduler
                    |      `.   /
         AF2------>||----+     .'-+
                    |    .'
         CS0------>||----+ .' Rate Scheduler
                    |'   (WFQ, WRR, etc)
```

Figure 2: Implementation as a data plane priority

The multi-policer model is shown in <u>Figure 3 (Implementation as a data
plane policer)</u>. In this model, traffic from each service class is
policed according to its SLA requirements, and then placed into a
common priority queue. Unlike the multi-priority model, the jitter
experienced by the traffic classes in this case is the same, as there
is only one queue, but the sum of the traffic in this higher priority
queue experiences less average jitter than the elastic traffic in the
lower priority.

```
                  policers    priorities  .
         Unadmitted EF <=> -------\        |`.
                              --||----+  `.
           Admitted EF <=> -------/    high|     `.
                         .              |     .'--------
            rate queues  |`.          +-----+   .'
         AF1------>||----+  `.     /  low | .' Priority
                         |    `.  /       |'  Scheduler
         AF2------>||----+     .'-+
                         |   .'
         CS0------>||----+ .' Rate Scheduler
                         |'   (WFQ, WRR, etc)
```

Figure 3: Implementation as a data plane policer

The difference between the two operationally is, as stated, the issues
of loss due to policing and distribution of jitter.
If the two traffic classes are, for example, voice and video, datagrams
conaining video data can be relatively large (often of variable sizes
up to the path MTU) while datagrams containing voice are relatively
small, on the order of only 40 to 200 bytes, depending on the codec. On
lower speed links (less than 10 MBPS), the jitter introduced by video
to voice can be disruptive, while at higher speeds the jitter is
nominal compared to the jitter requirements of voice. At access network
speeds, therefore, [RFC4594] (Babiarz, J., Chan, K., and F. Baker,
"Configuration Guidelines for DiffServ Service Classes," August 2006.)
recommends separation of video and voice into separate queues, while at
optical speeds [RFC5127] (Chan, K., Babiarz, J., and F. Baker,
"Aggregation of DiffServ Service Classes," February 2008.) recommends
that they use a common queue.
If, on the other hand, the two traffic classes are carrying the same
type of application with the same jitter requirements, then giving one
preference in this sense does not benefit the higher priority traffic
and may harm the lower priority traffic. In such a case, using separate
policers and a common queue is a superior approach.

2.2. Capacity admission control

There are at least six major ways that capacity admission is done or
has been proposed to be done for real-time applications. Each will be
described below, then Section 3 will judge which ones are likely to

meet the requirements of the Admitted Telephony service class. These
include:

    *Drop Precedence used to force sessions to voluntarily exit,

    *Capacity admission control by assumption or engineering,

    *Capacity admission control by call counting,

    *End-point capacity admission performed by probing the network,

    *Centralized capacity admission control via bandwidth broker, and

    *Distributed capacity admission control using protocols such as
     RSVP or NSIS.

The problem with capacity admission by assumption, which is widely
reployed in today's VoIP environment, is that it depends on the
assumptions made. One can do careful traffic engineering to ensure
needed bandwidth, but this can also be painful, and has to be revisited
when the network is changed or network usage changes.
The problem with dropping traffic to force users to hang up is that it
affects a broad class of users - if there is capacity for N calls and
the N+1 calls are active, data is dropped randomly from all sessions to
ensure that offered load doesn't exceed capacity. On very fast links,
that is acceptable, but on lower speed links it can seriously affect
call quality.
There are two fundamental problems with depending on the endpoint to
perform capacity admission; it may not be able to accurately measure
the impact of the traffic it generates on the network, and it tends to
be greedy (eg., it doesn't care). If the network operator is providing
a service, he must be able to guarantee the service, which means that
he can't trust systems that are not controlled by his network.
The problem with mechanisms that don't enable the association of a
policy with the request is that they don't allow for multi-policy
services, which are becoming important.
The operator's choice of admission procedure must, for this DSCP,
ensure the following:

    *The actual links that a session uses have enough bandwidth to
     support it.

    *New sessions are refused admission if there is inadequate
     bandwidth under the relevant policy.

    *If multiple policies are in use in a network, that the user is
     identified and the correct policy applied.

*Under periods of network stress, the process of admission of new
sessions does not disrupt existing sessions, unless the service
explicitly allows for disruption of calls.

2.3. Recommendations on implementation of an Admitted Telephony Service Class

It is the belief of the authors that either PHB implementation
described in Section 2.1 (Potential implementations of EF in this
model), if coupled with adequate AAA and capacity admission procedures
as described in Section 2.2 (Capacity admission control), are
sufficient to provide the services required for an Admitted Telephony
service class. If preemption is required, as described in section
2.3.5.2 of [RFC4542] (Baker, F. and J. Polk, "Implementing an Emergency
Telecommunications Service (ETS) for Real-Time Services in the Internet
Protocol Suite," May 2006.), this provides the tools for carrying out
the preemption. If preemption is not in view, or if used in addition to
preemptive services, the application of different thresholds depending
on call precedence has the effect of improving the probability of call
completion by admitting preferred calls at a time that other calls are
being refused. Routine and priority traffic can be admitted using the
same DSCP value, as the choice of which calls are admitted is handled
in the admission procedure executed in the control plane, not the
policing of the data plane.
On the point of what protocols and procedures are required for
authentication, authorization, and capacity admission, we note that
clear standards do not at this time exist for bandwidth brokers, NSIS
has not at this time been finalized and in any event is limited to
unicast sessions, and that RSVP has been standardized and has the
relevant services. We therefore recommend the use of RSVP at the UNI.
Procedures at the NNI are business matters to be discussed between the
relevant networks, and are recommended but not required.

3. Summary: changes from RFC 4594

To summarize, there are two changes to [RFC4594] (Babiarz, J., Chan,
K., and F. Baker, "Configuration Guidelines for DiffServ Service
Classes," August 2006.) discussed in this document:

Telephony class: The Telephony Service Class in RFC 4594 does not
involve capacity admission, but depends on application layer
admission that only estimates capacity, and that through static
engineering. In addition to that class, a separate Admitted

> Telephony Class is added which performs capacity admission
> dynamically.

Video classes Capacity admission is added to three video classes.
These are the Interactive Real-Time Traffic class, Broadcast TV
class when used for video on demand, and the Multimedia
Conferencing class.

---

4. IANA Considerations

This note requests that IANA assign a DSCP value to a second EF traffic
class consistent with [RFC3246] (Davie, B., Charny, A., Bennet, J.,
Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D.
Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior),"
March 2002.) and [RFC3247] (Charny, A., Bennet, J., Benson, K., Boudec,
J., Chiu, A., Courtney, W., Davari, S., Firoiu, V., Kalmanek, C., and
K. Ramakrishnan, "Supplemental Information for the New Definition of
the EF PHB (Expedited Forwarding Per-Hop Behavior)," March 2002.) in
the "Differentiated Services Field Codepoints" registry. It implements
the Telephony Service Class described in [RFC4594] (Babiarz, J., Chan,
K., and F. Baker, "Configuration Guidelines for DiffServ Service
Classes," August 2006.) at lower speeds and is included in the Real
Time Treatment Aggregate (Chan, K., Babiarz, J., and F. Baker,
"Aggregation of DiffServ Service Classes," February 2008.) [RFC5127] at
higher speeds. The recommended value for the code point is 101100,
paralleling the EF code point, which is 101110. The code point should
be referred to as VOICE-ADMIT.
This traffic class requires the use of capacity admission such as RSVP
services together with AAA services at the User/Network Interface
(UNI); the use of such services at the NNI is at the option of the
interconnected networks.

---

5. Security Considerations

A major requirement of this service is effective use of a signaling
protocol such as RSVP, with the capabilities to identify its user
either as an individual or as a member of some corporate entity, and
assert a policy such as "routine" or "priority".
This capability, one has to believe, will be abused by script kiddies
and others if the proof of identity is not adequately strong or if
policies are written or implemented improperly by the carriers. This
goes without saying, but this section is here for it to be said...

6. Acknowledgements

7. References   TOC

7.1. Normative References

TOC

[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC2474]	Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," RFC 2474, December 1998 (TXT, HTML, XML).
[RFC3246]	Davie, B., Charny, A., Bennet, J., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)," RFC 3246, March 2002 (TXT).
[RFC4594]	Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes," RFC 4594, August 2006 (TXT).

7.2. Informative References

TOC

[ITU.MLPP. 1990]	International Telecommunications Union, "Multilevel Precedence and Preemption Service," ITU-T Recommendation I.255.3, 1990.
[RFC2475]	Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services," RFC 2475, December 1998 (TXT, HTML, XML).
[RFC3247]	Charny, A., Bennet, J., Benson, K., Boudec, J., Chiu, A., Courtney, W., Davari, S., Firoiu, V., Kalmanek, C., and K. Ramakrishnan, "Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior)," RFC 3247, March 2002 (TXT).
[RFC3260]	Grossman, D., "New Terminology and Clarifications for Diffserv," RFC 3260, April 2002 (TXT).

[RFC4190]	Carlberg, K., Brown, I., and C. Beard, "Framework for Supporting Emergency Telecommunications Service (ETS) in IP Telephony," RFC 4190, November 2005 (TXT).
[RFC4504]	Sinnreich, H., Lass, S., and C. Stredicke, "SIP Telephony Device Requirements and Configuration," RFC 4504, May 2006 (TXT).
[RFC4542]	Baker, F. and J. Polk, "Implementing an Emergency Telecommunications Service (ETS) for Real-Time Services in the Internet Protocol Suite," RFC 4542, May 2006 (TXT).
[RFC5127]	Chan, K., Babiarz, J., and F. Baker, "Aggregation of DiffServ Service Classes," RFC 5127, February 2008 (TXT).

Authors' Addresses

	Fred Baker
	Cisco Systems
	Santa Barbara, California 93117
	USA
Phone:	+1-408-526-4257
Email:	fred@cisco.com
	James Polk
	Cisco Systems
	Richardson, Texas 75082
	USA
Phone:	+1-817-271-3552
Email:	jmpolk@cisco.com
	Martin Dolly
	AT&T Labs
	Middletown Township, New Jersey 07748
	USA
Phone:	+1-732-420-4574
Email:	mdolly@att.com

Full Copyright Statement

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.
Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.
The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.