The Session Description Protocol (SDP) 'trafficclass' Attribute
draft-ietf-mmusic-traffic-class-for-sdp-02

Abstract

   This document proposes a new Session Description Protocol (SDP)
   attribute to identify the traffic class a session is requesting
   in its offer/answer exchange.

Status of this Memo

Copyright Notice

Table of Contents

   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 [RFC2119].

1. **Introduction**

   The Session Description Protocol (SDP) [RFC4566] provides a means
   for an offerer to describe the specifics of a session to an
   answerer, and for the answerer to respond back with its session
   specifics to the offerer. These session specifics include offering
   the codec or codecs to choose from, the specific IP address and port
   number the offerer wants to receive the RTP stream(s) on/at, the
   particulars about the codecs the offerer wants considered or
   mandated, and so on.

   There are many facets within SDP to determine the Real-time
   Transport Protocol (RTP) [RFC3550] details for the session
   establishment between one or more endpoints, but identifying how the
   underlying network should process each stream still remains
   under-specified.

   The ability to identify a traffic flow by port number gives an
   indication to underlying network elements to treat traffic with
   dissimilar ports in a different way, the same or in groups the same
   - but different from other ports or groups of ports.

   Within the context of realtime communications, the labeling of an
   RTP session based on media descriptor lines as just a voice and/or
   video session is insufficient, and provides no guidelines to the

underlying network on how to treat the traffic. A more granular
labeling helps on several fronts to

- inform application layer elements in the signaling path the
  intent of this session.

- inform the network on how to treat the traffic if the network is
  configured to differentiate session treatments based on the type
  of session the RTP is, including the ability to provide call
  admission control based on the type of traffic in the network.

- allow network monitoring/management of traffic types realtime and
  after-the-fact analysis.

Some network operators want the ability to guarantee certain traffic
gets a minimum amount of network bandwidth per link or through a
series of links that make up a network such as a campus or WAN, or a
backbone. For example, a call center voice application might get at
least 20% of the available link bandwidth.

Some network operators want the ability to allow certain users or
devices access to greater bandwidth during non-busy hours, than
during busy hours of the day. For example, all desktop video might
operate at 1080p during non-peak times, but a similar session might
be curtailed between the same users or devices to 720p or 360p
during peak hours. Another example would be to reduce the frames
per second (fps) rate, say from 30fps to 15fps. This case is not as
clear as accepting or denying similar sessions during different
times of the day, but tuning the access to the bandwidth based on
the type of session. In other words, tune down the bandwidth for
desktop video during peak hours to allow a 3-screen Telepresence
session that would otherwise look like the same type of traffic
(RTP, and more granular, video).

RFC 4594 established a guideline for classifying the various flows
in the network and the Differentiated Services Codepoints (DSCP)
that apply to many traffic types (table 3 of [RFC4594]), including
RTP based voice and video traffic sessions. The RFC also defines the
per hop network behavior that is strongly encouraged for each of
these application traffic types based on the traffic characteristics
and tolerances to delay, loss and jitter within each traffic class.

Video was broken down into 4 categories in that RFC, and voice into
another single category. We do not believe this satisfies the
technical and business requirements to accomplish sufficiently
unique labeling of RTP traffic.

If the application becomes aware of traffic labeling,

- this can be coded into layer 3 mechanisms.

- this can be coded into layer 4 protocols and/or mechanisms.

- this can be coded into a combination of mechanisms and protocols.

The layer 3 mechanism for differentiating traffic is either the port
number or the Differentiated Services Codepoint (DSCP) value
[RFC2474]. Within the public Internet, if the application is not
part of a managed service, the DSCP likely will be best effort (BE),
or reset to BE when ingressing a provider's network. Within the
corporate LAN, this is usually completely configurable and a local
IT department can take full advantage of this labeling to shape and
manage their network as they see fit.

Within a network core, DiffServ typically does not apply. That said,
DiffServ can be used to identify which traffic goes into which MPLS
tunnel [RFC4124].

Labeling realtime traffic types using a layer 4 protocol would
likely involve RSVP [RFC2205] or NSIS [RFC4080]. RSVP has an
Application Identifier (app-ID) defined in [RFC2872] that provides a
means for carrying a traffic class label along the media path. An
advantage of this mechanism is that the label can inform each domain
along the media path what type of traffic this traffic flow is, and
allow each domain to adjust the appropriate DSCP value (set by each
domain for use within that domain). Meaning, if a DSCP value is set
by an endpoint or a router in the first domain and gets reset by a
service provider, the far-end domain will be able to reset the DSCP
value to the intended traffic class. There is a proposed extension
to RSVP which creates individual profiles for what goes into each
app-ID field to describe these traffic classes [ID-RSVP-PROF], which
will take advantage of what is described in this document.

There are several proprietary mechanisms that can take advantage of
this labeling, but none of those will be discussed here.

The idea of traffic - or service - identification is not new; it has
been described in [RFC5897]. If that RFC is used as a guideline,
identification that leads to stream differentiation can be quite
useful. One of the points within RFC 5897 is that users cannot be
allowed to assign any identification (fraud is one reason given). In
addition, RFC 5897 recommends that service identification should be
done in signaling, rather than guessing or deep packet inspection.
Any network currently would have to currently guess or perform deep
packet inspection to classify traffic and offer the service as per
RFC 4594 since such service identification information is currently
not available in SDP and therefore to the network elements. Since
SDP understands how each stream is created (i.e., the particulars of
the RTP stream), this is the right place to have this service
differentiated. Such service differentiation can then be
communicated to and leveraged by the network.

[Editor's Note: the words "traffic" and "service" are similar enough

that the above paragraph talks about RFC 5897's
"service identification", but this document only
discuss and propose traffic indications in SDP.]

This document proposes a simple attribute line to identify the
application a session is requesting in its offer/answer exchange.
This document uses previously defined service class strings for
consistency between IETF documents.

This document modifies the traffic classes originally created in RFC
4594 in Section 2, incrementing each class with application
identifiers and optional adjective strings. Section 3 defines the
new SDP attribute "trafficclass". Section 4 discusses the offerer
and answerer behavior when generating or receiving this attribute.

2. Traffic Class Framework and String Definitions

The framework of the traffic class attribute will have at least two
parts, allowing for several more to be included. The intention is to
have a category class (e.g., Conversational) that merely serves as
the anchor point for an application component that when paired
together, form the highest level traffic class. An adjective
component provides further granularity for the application. There
can be more than one adjective within a traffic class label to
further refine the uniqueness of a traffic class being described.

The traffic class label will have the following structure,

    category.application(.adjective)(.adjective)

[Editor's Note: the above is not exactly the ABNF to be used.
                The order is right. The category and application
                MUST appear first (each only once) and zero or more
                adjectives can appear following the application
                component.]

    Where
    1) the 1st component is the human understandable category;
    2) the 2nd component is the application;
    3) an optional 3rd component or series of components are
       adjective(s) used to further refine the application component;

The construction of the traffic class label for Telepresence video
would follow the minimum form of:

    Conversational.video.immersive

where there might be one or more adjective after '.immersive'.

There is no traffic class or DSCP value associated with just
"Conversational". There is a traffic class associated with
"Conversational.video", creating a differentiation between it and a

"Conversational.video.immersive" traffic class, which would have
DSCP associated with the latter traffic class, depending on local

   policy. Each category component is defined below, as are several of
   application and adjective strings.

   [Editor's Note: We're not yet sure how much of what's below will be
                   proposed for IANA registration, but the 5 category
                   components will be, as well as at least some
                   application components per category component. Some
                   adjective components will also likely be proposed
                   for IANA registration.

   The 5 category components of the traffic class attribute are as
   follows:

   o Conversational
   o Multimedia-Conferencing
   o Realtime-Interactive
   o Multimedia-Streaming
   o Broadcast

   The following application components of the traffic class attribute
   are as follows:

   o Audio
   o Video
   o Text
   o application-sharing
   o Presentation-data
   o Whiteboarding
   o Webchat/IM
   o Gaming
   o Virtual-desktop (interactive)
   o Remote-desktop
   o Telemetry (e.g., NORAD missile control)
   o Multiplex (i.e., combined streams)
   o Webcast
   o IPTV
   o Live-events (one-way, in realtime)
   o surveillance

   The following adjective components of the traffic class attribute
   are as follows:

   o Immersive
   o avconf
   o Realtime-Text
   o web

   Each of the above 3 lists will be defined in the following

subsections.

2.1 Conversational Category Traffic Class

   The Conversational traffic class is best suited for applications
   that require very low delay variation and generally intended to
   enable realtime, bi-directional person-to-person or
   multi-directional via an MCU communication. The following
   application components are appropriate for use with the
   Conversational category:

   o Audio (voice)**

   o Video**

   o Text (i.e., real-time text required by deaf users)

   **The above applications will also be used within Multimedia
     Streaming and Broadcast

   With adjective substrings to the above

   Immersive (TP) - An interactive audio-visual communications
         experience between remote locations, where the users enjoy a
         strong sense of realism and presence between all participants
         by optimizing a variety of attributes such as audio and video
         quality, eye contact, body language, spatial audio,
         coordinated environments and natural image size.

   Avconf - An interactive audio-visual communication experience
         that is not immersive in nature, though can have a high
         resolution video component.

   Realtime-Text (RTT) - a term for real-time transmission of text in
         a character-by-character fashion for use in conversational
         services, often as a text equivalent to voice-based
         conversational services. Conversational text is defined in the
         ITU-T Framework for multimedia services, Recommendation F.700
         [RFC5194].

   Web - for realtime aspects of web conferencing; mutually exclusive
         of both Immersive and Desktop video experiences

```
+--------------------------------------------------------------------+
|Traffic Class  |                                | Tolerance to      |
|    Name       |   Traffic Characteristics      | Loss |Delay |Jitter|
|===============+================================+======+======+======|
|               | High priority, typically       | Very | Very | Very |
|Conversational | small packets (large video     | Low  | Low  | Low  |
|               | frames produce large packets), |      |      |      |
|               | generally sustained high       |      |      |      |
```

```
|               | packet rate, low inter-packet |    |    |    |
|               | transmission interval,        |    |    |    |
```

```
|                 | usually UDP framed in (S)RTP  |      |      |      |
+---------------+------------------------------+------+------+------+
```

Figure 1. Conversational Traffic Characteristics

2.2 Multimedia-Conferencing Category Traffic Class

Multimedia-Conferencing traffic class is best suited for
applications that are generally intended for communication between
human users, but are less demanding in terms of delay, packet loss,
and jitter than what Conversational applications require. These
applications require low to medium delay and may have the ability to
change encoding rate (rate adaptive) or transmit data at varying
rates. The following application components are appropriate for use
with the Multimedia-Conferencing category:

o application-sharing (that webex does or protocols like T.128) -
    An application that shares the output of one or more running
    applications or the desktop on a host. This can utilize
    vector graphics, raster graphics or video.

o Presentation-data - can be a series of still images or motion
    video.

o Whiteboarding - an application enabling the exchange of graphical
    information including images, pointers and filled and
    unfilled parametric drawing elements (points, lines,
    polygons and ellipses).

o (RTP-based) file transfer

o Web (conference) chat/instant messaging

```
+-----------------------------------------------------------------+
|Traffic Class  |                              |   Tolerance to    |
|    Name       |   Traffic Characteristics    | Loss |Delay |Jitter|
|===============+==============================+======+======+======|
|  Multimedia   | Variable size packets,       | Low  | Low  | Low  |
| Conferencing  | Variable transmit interval,  |  -   |  -   |  -   |
|               | rate adaptive, reacts to     |Medium|Medium|Medium|
|               | loss, usually TCP-based      |      |      |      |
+---------------+------------------------------+------+------+------+
```

Figure 2. Multimedia Conferencing Traffic Characteristics

2.3 Realtime-Interactive Category Traffic Class

Realtime-Interactive traffic class is intended for interactive

variable rate inelastic applications that require low jitter and
loss and very low delay. The following application components are

   appropriate for use with the Realtime-Interactive category:

   o Gaming - interactive player video games with other users on other
       hosts (e.g., Doom)

   o Virtualized desktop (interactive) - similar to an X-windows
       station, has no local hard drive, or is operating an
       application with no local storage

   o Remote Desktop - controlling a remote node with local peripherals
       (i.e., monitor, keyboard and mouse)

   o Telemetry - a communication that allows remote measurement and
       reporting of information (e.g., post launch missile status or
       energy monitoring)

```
+------------------------------------------------------------------+
|Traffic Class  |                                  |   Tolerance to   |
|    Name       |   Traffic Characteristics        | Loss |Delay |Jitter|
|===============+==================================+======+======+======|
|   Realtime    | Inelastic, mostly variable       | Low  | Very | Low  |
|  Interactive  | rate, rate increases with        |      | Low  |      |
|               | user activity                    |      |      |      |
+---------------+----------------------------------+------+------+------+
```

           Figure 3. Realtime Interactive Traffic Characteristics

2.4 Multimedia-Streaming Category Traffic Class

   Multimedia-Streaming traffic class is best suited for variable rate
   elastic streaming media applications where a human is waiting for
   output and where the application has the capability to react to
   packet loss by reducing its transmission rate. The following
   application components are appropriate for use with the
   Multimedia-Streaming category:

   o Audio (see Section 2.1)

   o Video (see Section 2.1)

   o Multiplex (i.e., combined a/v streams)

   With adjective substrings to the above (which may or may not get
   IANA registered)

   Webcast

   The primary difference from the Multimedia-streaming category class
   and the Broadcast category class is about the length of time for

buffering. Buffered streaming audio and/or video which are initiated
   by SDP, and not HTTP. Buffering here can be from many seconds to

hours, and is typically at the destination end (as opposed to
Broadcast buffering which is minimal at the destination). The
buffering aspect is what differentiates this category class from the
Broadcast class (which has minimal or no buffering).

```
+------------------------------------------------------------------+
|Traffic Class  |                              |    Tolerance to    |
|    Name       |  Traffic Characteristics     | Loss |Delay |Jitter|
|===============+==============================+======+======+======|
|  Multimedia   | Variable size packets,       |Low - |Medium| High |
|   Streaming   | elastic with variable rate   |Medium|- High|      |
|               |                              |      |      |      |
+---------------+------------------------------+------+------+------+
```

        Figure 4. Multimedia Streaming Traffic Characteristics

2.5 **Broadcast Category Traffic Class**

   Broadcast traffic class is best suited for inelastic streaming media
   Applications, which might have a 'wardrobe malfunction' delay at or
   near the source but not typically at the destination, that may be of
   constant or variable rate, requiring low jitter and very low packet
   loss. The following application components are appropriate for use
   with the Broadcast category:

   o Audio (see Section 2.1)

   o Video (see Section 2.1)

   o Multiplex (i.e., combined a/v streams)

   With adjective substrings to the above:

   o IPTV

   o Live events (non-buffered)

   o surveillance - one way audio from a microphone or video from a
           camera (e.g., observing a parking lot or building exit),
           typically enabled for long periods of time, usually stored
           at the destination.

```
+------------------------------------------------------------------+
|Traffic Class  |                              |    Tolerance to    |
|    Name       |  Traffic Characteristics     | Loss |Delay |Jitter|
|===============+==============================+======+======+======|
|   Broadcast   | Constant and variable rate,  | Very |Low - |Low - |
|               | inelastic, generally         | Low  |Medium|Medium|
|               | non-bursty flows, generally  |      |      |      |
```

```
|                     | sustained high packet rate,   |      |      |      |
|                     | low inter-packet transmission |      |      |      |
```

```
|               | interval, usually UDP framed |      |      |      |
|               | in (S)RTP                    |      |      |      |
+---------------+------------------------------+------+------+------+
```

Figure 5. Broadcast Traffic Characteristics

3. SDP Attribute Definition

This document proposes the 'trafficclass' session and media-level
SDP [RFC4566] attribute. The following is the Augmented
Backus-Naur Form (ABNF) [RFC5234] syntax for this attribute, which
is based on the SDP [RFC4566] grammar:

```
attribute              =/ traffic-class-label

traffic-class-label    = "trafficclass" ":" [SP] category
                          "." application *( "." adjective )

category               = "Broadcast" /
                          "Realtime-Interactive" /
                          "Multimedia-Conferencing" /
                          "Multimedia-Streaming" /
                          "Conversational" / tcl-token

application            = tcl-token

adjective              = classified-adjective /
                          unclassified-adjective

classified-adjective   = tcl-token ":" tcl-token

unclassified-adjective = tcl-token

tcl-token              = %2D / %30-%39 / %41-%5A / %61-7A
```


The attribute is named "trafficclass", for traffic classification,
identifying which one of the five categories applies to the
media stream associated with this m-line. There MUST NOT be more
than one category component per media line.

The category in this document are an augmented version of the
application labels introduced by table 3 of RFC 4595 (which will be
rewritten based on the updated labels and treatments expected for
each traffic class defined in this document).

```
+-------------------------+------------------------------+
| Application Labels       |  Category Classes Defined   |
| Defined in RFC 4594      |    in this document         |
```

```
+========================+============================+
| Broadcast-video        |   Broadcast                |
```

```
+========================+============================+
| Broadcast-video        |   Broadcast                |
```

```
+------------------------+-----------------------------+
| Realtime-Interactive   |  Realtime-Interactive       |
+------------------------+-----------------------------+
| Multimedia-Conferencing|  Multimedia-Conferencing    |
+------------------------+-----------------------------+
| Multimedia-Streaming   |  Multimedia-Streaming       |
+------------------------+-----------------------------+
| Telephony              |  Conversational             |
+------------------------+-----------------------------+
```

Figure 6. Label Change Differences from RFC 4594

As is evident from the changes above, from left to right, two labels
are different and each of the meanings are different in this
document relative to how RFC 4594 defined them. These differences
are articulated in Section 2 of this document.

A category is a human understandable categorization, and MUST NOT be
the only component of the traffic class label present in the
attribute. The category string MUST always be paired with an
application component, with a "." as the component separator.

The application types define the application of a particular traffic
flow, for example, audio or video. The application types are listed
and defined in Section 2 of this document. Not every category is
paired with application each listed, at least as defined in this
document. Section 2.1 through 2.5 list many of the expected
combinations.

For additional application type granularity, adjective components
can be added (also listed in Section 2). One or more adjectives can
be within the same traffic class attribute. It is also permitted to
include one or more non-IANA registered adjective component, but
these MUST be prefaced by the additional delimiter "_", creating a
possibility such as

  category.application-type.adjective._non-standard-adjective
                                      ^^^^
                              See the underscore

For example, this is valid:

    m=video 50000 RTP/AVP 112
    a=trafficclass Conversational.video.immersive._foo._bar

where both "foo" and "bar" are not IANA registered adjectives, but
"immersive" is IANA registered. However, including non-registered
adjectives without the "_" delimiter MUST NOT occur, such as the
following:

```
m=video 50000 RTP/AVP 112
a=trafficclass Conversational.video.immersive.foo.bar
```

There is no limit to the number of adjectives allowed, without regard for whether they are registered or not. These non-registered adjectives can be vendor generated, or merely considered to be proprietary in nature.

It is important to note that the order of component types matter, but not the order of the adjective components. There might be local significance to the ordering of adjectives though. In other words, the category class component MUST be before the application component, which MUST be before any and all adjective component(s).

Some algorithm such as alphabetizing the list and matching the understood strings SHOULD be used.

Adjectives can be either unqualified or qualified. Qualified adjectives have a delimiter ":" after the previous "." separating the string component into two parts.

The tcl-token "aq" is the first part of an adjective if it is qualified, and either the "admitted", "non-admitted" or "none" tcl-token is the second part of the qualified adjective allowable according to this specification. In the form

    aq:admitted|non-admitted|none

The only valid use of the tcl-token "aq" is to pair with either the "admitted", "non-admitted" or "none" tcl-token. At the same time, the tcl-tokens "admitted", "non-admitted" or "none" MUST NOT appear without a preceding "aq:".

Like all adjectives, it is OPTIONAL to include this adjective in any trafficclass attribute, and has the following meanings:

- aq -   for 'admission qualifier' to indicate the purpose of
                the following adjective parts with respect to the
                capacity admission status of this traffic flow
                described by this m-line.

- admitted - capacity admission mechanisms or protocols are to be or
                were used for the full amount of bandwidth in relation
                to this m= line.

- non-admitted - capacity admission mechanisms or protocols were
                attempted but failed in relation to this m= line. This
                does not mean the flow described by this m= line
                failed. It just failed to attain the capacity admission
                mechanism or protocol necessary for a predictable
                quality of service, and is likely to continue with only
                a class of service marking or best effort.

- none -   no capacity admission mechanisms or protocols are or

were attempted in relation to this m= line.

The default for any flow generated from an m-line not having a
trafficclass adjective of 'aq:admitted' or 'aq:non-admitted' MUST be
the equivalent of 'aq:none', whether or not it is present.

Any category class, application, or adjective string component
within this attribute that is not understood MUST be ignored,
leaving all that is understood to be processed. Ignored string
components SHOULD NOT be deleted, as a downstream entity could
understand the component(s) and use it/them during processing.

Not understanding the category class string SHOULD mean that this
attribute is ignored.

The following is an example of media level description with a
'trafficclass' attribute:

    m=video 50000 RTP/AVP 112
    a=trafficclass conversational.video.immersive.aq:admitted

The above indicates a Telepresence session that has had capacity
admission process applied to its media flow.

4. Offer/Answer Behavior

Through the inclusion of the 'trafficclass' attribute, an
offer/answer exchange identifies the application type for use by
endpoints within a session. Policy elements can use this attribute
to determine the acceptability and/or treatment of that session
through lower layers. One specific use-case is for setting of the
DSCP specific for that application type (say a Broadcast instead
of a Conversational video), decided on a per domain basis -
instead of exclusively by the offering domain.

4.1 Offer Behavior

Offerers include the 'trafficclass' attribute with a single string
comprised of two or more components (from the list in Section 2) to
obtain configurable and predictable classification between the
answerer and the offerer. The offerer can also include a private set
of components, or a combination of IANA registered and private
components within a single domain (e.g., enterprise networks).

Offerers of this 'trafficclass' attribute MUST NOT change the label
in transit (e.g., wrt to B2BUAs). Session Border Controllers (SBC)
at domain boundaries can change this attribute through local policy.

Offers containing a 'trafficclass' label not understood are ignored
   by default (i.e., as if there was no 'trafficclass' attribute in the

offer).

4.2 Answer Behavior

Upon receiving an offer containing a 'trafficclass' attribute, if
the offer is accepted, the answerer will use this attribute to
classify the session or media (level) traffic accordingly towards
the offerer. This answer does not need to match the traffic class in
the offer, though this will likely be the case most of the time.

In order to understand the traffic class attribute, the answerer
MUST check several components within the attribute, such as

1 - does the answerer understand the category component?

    If not, the attribute SHOULD be ignored.

    If yes, it checks the application component.

2 - does the answerer understand the application component?

    If not, the answerer needs to check if it has a local policy to
    proceed without an application component. The default for this
    situation is as if the category component was not understand,
    the attribute SHOULD be ignored.

    If yes, it checks to see if there are any adjective components
    present in this attribute to start its classification.

3 - does the answerer understand the adjective component or
    components if any are present?

    If not present, process and match the trafficclass label value
    as is.

    If yes, determine if there is more than one. Search for each
    that is understood. Any adjectives not understood are to be
    ignored, as if they are not present. Match all remaining
    understood components according to local policy and process
    attribute.

The answerer will answer the offer with its own 'trafficclass'
attribute, which will likely be the same value, although this is not
mandatory (at this time). The Offerer will process the received
answer just as the answerer processed the offer. In other words, the
processing steps and rules are identical for each end.

The answerer should expect to receive RTP packets marked as
indicated by its 'trafficclass' attribute in the answer itself.

An Answer MAY have a 'trafficclass' attribute when one was not in

the offer. This will at least aid the local domain, and perhaps
each domain the session transits, to categorize the application type
of this RTP session.

Answerers that are middleboxes can use the 'trafficclass' attribute
to classify the RTP traffic within this session however local policy
determines. In other words, this attribute can help in deciding
which DSCP an RTP stream is assigned within a domain, if the
answerer were an inbound SBC to a domain.

5. Security considerations

RFC 5897 [RFC5897] discusses many of the pitfalls of service
classification, which is similar enough to this idea of traffic
classification to apply here as well. That document highly
recommends the user not being able to set any classification.
Barring a hack within an endpoint (i.e., to intentionally
misclassifying (i.e., lying) about which classification an RTP
stream is), this document's solution makes the classification part
of the signaling between endpoints, which is recommended by RFC
5897.

6. IANA considerations

6.1 Registration of the SDP 'trafficclass' Attribute

This document requests IANA to register the following SDP att-field
under the Session Description Protocol (SDP) Parameters registry:

Contact name: jmpolk@cisco.com

Attribute name: trafficclass

Long-form attribute name: Traffic Classification

Type of attribute: Session and Media levels

Subject to charset: No

Purpose of attribute: To indicate the Traffic Classification
                        application for this session

Allowed attribute values: IANA Registered Tokens

Registration Procedures: Specification Required

Type	SDP Name	Reference

att-field (both session and media level)

                    trafficclass   [this document]

6.2 The Traffic Classification Category Registration

   This document requests IANA to create a new registry for the
   traffic Category classes similar to the following table within
   the Session Description Protocol (SDP) Parameters registry:

   Registry Name: "trafficclass" SDP Category Attribute Values
   Reference: [this document]
   Registration Procedures: Standards-Track document Required

Category Values	Reference
Broadcast	[this document]
Realtime-Interactive	[this document]
Multimedia-Conferencing	[this document]
Multimedia-Streaming	[this document]
Conversational	[this document]

6.3 The Traffic Classification Application Type Registration

   This document requests IANA to create a new registry for the
   traffic application classes similar to the following table
   within the Session Description Protocol (SDP) Parameters registry:

   Registry Name: "trafficclass" Attribute Application Type Values
   Reference: [this document]
   Registration Procedures: Specification Required

Application Values	Reference
Audio	[this document]
Video	[this document]
Text	[this document]
Application-sharing	[this document]
Presentation-data	[this document]
Whiteboarding	[this document]
Webchat/IM	[this document]
Gaming	[this document]
Virtualized-desktop	[this document]
Remote-desktop	[this document]
Telemetry	[this document]
Multiplex	[this document]
Webcast	[this document]
IPTV	[this document]
Live-event	[this document]
surveillance	[this document]

6.4 The Traffic Classification Unqualified Adjective Registration

   This document requests IANA to create a new registry for the
   traffic adjective values similar to the following table
   within the Session Description Protocol (SDP) Parameters registry:

   Registry Name: "trafficclass" Attribute Adjective Values
   Reference: [this document]
   Registration Procedures: Specification Required

   Adjective Values   Reference
   ------------------   ---------
   Immersive   [this document]
   Desktop-video   [this document]
   Realtime-Text   [this document]
   web   [this document]
   aq   [this document]
   admitted   [this document]
   non-admitted   [this document]
   none   [this document]

7. Acknowledgments

   To Dave Oran, Toerless Eckert, Henry Chen, David Benham, David
   Benham, Mo Zanty, Michael Ramalho, Glen Lavers, Charles Ganzhorn,
   Paul Kyzivat and Greg Edwards for their comments and suggestions.

8. References

8.1. Normative References

 [RFC2119] S. Bradner, "Key words for use in RFCs to Indicate
           Requirement Levels", RFC 2119, March 1997

 [RFC2205] R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, S. Jamin,
           "Resource ReSerVation Protocol (RSVP) -- Version 1
           Functional Specification", RFC 2205, September 1997

 [RFC2474] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the
           Differentiated Services Field (DS Field) in the IPv4 and
           IPv6 Headers ", RFC 2474, December 1998

 [RFC2872] Y. Bernet, R. Pabbati, "Application and Sub Application
           Identity Policy Element for Use with RSVP", RFC 2872,
           June 2000

 [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V.

Jacobson, "RTP: A Transport Protocol for Real-Time
Applications", STD 64, RFC 3550, July 2003.

  [RFC4080] R. Hancock, G. Karagiannis, J. Loughney, S. Van den Bosch,
            "Next Steps in Signaling (NSIS): Framework", RFC 4080, June
            2005

  [RFC4124] F. Le Faucheur, Ed., " Protocol Extensions for Support of
            Diffserv-aware MPLS Traffic Engineering ", RFC 4124,
            June 2005

  [RFC4566] M. Handley, V. Jacobson, C. Perkins, "SDP: Session

            Description Protocol", RFC 4566, July 2006

  [RFC5234] Crocker, D., Ed., and P. Overell, "Augmented BNF for Syntax
            Specifications: ABNF", STD 68, RFC 5234, January 2008.

  [RFC5865] F. Baker, J. Polk, M. Dolly, "A Differentiated Services Code
            Point (DSCP) for Capacity-Admitted Traffic", RFC 5865,
            May 2010

  [RFC5897] J. Rosenberg, "Identification of Communications Services in
            the Session Initiation Protocol (SIP)", RFC 5897, June 2010

8.2. Informative References

  [RFC4594] J. Babiarz, K. Chan, F Baker, "Configuration Guidelines for
            Diffserv Service Classes", RFC 4594, August 2006

  [ID-RSVP-PROF] J. Polk, S. Dhesikan, "Resource Reservation Protocol
            (RSVP) Application-ID Profiles for Voice and Video Streams",
            work in progress, Mar 2011

Author's Addresses

   James Polk
   3913 Treemont Circle
   Colleyville, Texas, USA
   +1.818.271.3552

   mailto: jmpolk@cisco.com

   Subha Dhesikan
   170 W Tasman St
   San Jose, CA, USA
   +1.408-902-3351

   mailto: sdhesika@cisco.com

Paul E. Jones
7025 Kit Creek Rd.
Research Triangle Park, NC, USA
+1 919 476 2048

mailto: paulej@packetizer.com

Appendix - Changes from Previous Versions

A.1 **From -01 to -02**

These are the following changes made between the WG -01 version and
the -02 version:

- converged the use of terms 'parent' and 'category' to just
  'category' for consistency.

- changed ABNF to reflect extensibility by not having applications
  and adjectives named in the ABNF, rather have them merely IANA
  registered.

- merged the qualified and unqualified adjective sections into a
  single section on adjectives, but allowing some to have a
  preceding qualifier.

- text clean-up

A.2 **From -00 to -01**

These are the following changes made between the WG -00 version and
the -01 version:

- removed the non-SDP applications Netflix and VOD

- switched the adjective 'desktop' to 'avconf'

- Labeled each of the figures.

- clarified the differences between Multimedia-Streaming and
  Broadcast category categories.

- defined Video surveillance

- added the concept of a 'qualified' adjective, and modified the
  ABNF.

- deleted the idea of a 'cac-class' as a separate component, and
  made the equivalent a qualified adjective.

- modified the answerer behavior because of the removal of the

'cac-class' component.

- created an IANA registry for qualified adjectives

- general clean-up of the doc.

Did *not* do the following in this version:

- add the ability to have more than one trafficclass attribute based
  on the codec chosen, as feedback indicated this was a bad idea.

- no swap of the Multimedia-Conferencing category with the
  offered Collaboration category, as doing this did not solve
  any perceived problems.

- add more to the 'how does this get processed' portion of Section
  3. That will come in the next revision.