Resource ReSerVation Protocol (RSVP) Extensions for Emergency Services
draft-ietf-tsvwg-emergency-rsvp-11.txt

Status of this Memo

   This Internet-Draft is submitted to IETF in full conformance with the
   provisions of BCP 78 and BCP 79.

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

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

   The list of current Internet-Drafts can be accessed at
   http://www.ietf.org/ietf/1id-abstracts.txt.

   The list of Internet-Draft Shadow Directories can be accessed at
   http://www.ietf.org/shadow.html.

   This Internet-Draft will expire on August 22, 2009.

Copyright Notice

Abstract

   An Emergency Telecommunications Service (ETS) requires the ability to
   provide an elevated probability of session establishment to an
   authorized user in times of network congestion (typically, during a
   crisis). When supported over the Internet Protocol suite, this may
   be facilitated through a network layer admission control solution,
   which supports prioritized access to resources (e.g., bandwidth).
   These resources may be explicitly set aside for emergency services,
   or they may be shared with other sessions.

   This document specifies extensions to the Resource reSerVation
   Protocol (RSVP) that can be used to support such an admission
   priority capability at the network layer. Note that these extensions
   represent one possible solution component in satisfying ETS
   requirements. Other solution components, or other solutions, are
   outside the scope of this document.

   The mechanisms defined in this document are applicable to controlled
   environments.

Table of Contents

1. Introduction

[RFC3689] and [RFC3690] detail requirements for an Emergency
Telecommunications Service (ETS), which is an umbrella term
identifying those networks and specific services used to support
emergency communications. Deployed examples of these types of
networks are the Government Emergency Telecommunications Systems
(GETS) and the Government Telephone Preference System (GTPS) [NCS]
[RFC4190]. Both of these examples represent enhancements to publicly
accessible systems instead of walled garden or private networks.

An underlying goal of [RFC3689] and [RFC3690] is to present
requirements that elevate the probability of session establishment
from an authorized user in times of network congestion (presumably
because of a crisis condition). In some extreme cases, the
requirement for this probability may reach 100%, but that is a topic
subject to policy and most likely local regulation (the latter being
outside the scope of this document).

Solutions to meet this requirement for elevated session establishment
probability may involve session layer capabilities prioritizing
access to resources controlled by the session control function. As
an example, entities involved in session control (such as SIP user
agents, when the Session Initiation Protocol (SIP) [RFC3261], is the
session control protocol in use) can influence their treatment of
session establishment requests (such as SIP requests). This may
include the ability to "queue" session establishment requests when
those can not be immediately honored (in some cases with the notion
of "bumping", or "displacement", of less important session
establishment requests from that queue). It may include additional
mechanisms such as exemption from certain network management
controls, and alternate routing.

Solutions to meet the requirement for elevated session establishment
probability may also take advantage of network layer admission
control mechanisms supporting admission priority. Networks usually
have engineered capacity limits that characterize the maximum load
that can be handled (say, on any given link) for a class of traffic
while satisfying the quality of service requirements of that traffic
class. Admission priority may involve setting aside some network
resources (e.g. bandwidth) out of the engineered capacity limits for
the emergency services only. Or alternatively, it may involve
allowing the emergency related sessions to seize additional resources
beyond the engineered capacity limits applied to normal sessions.
This document specifies the necessary extensions to support such
admission priority when network layer admission control is performed
using the Resource reSerVation Protocol (RSVP) ([RFC2205]).

IP telephony "calls" are one form of "sessions" that can benefit from the elevated session establishment probability discussed in this document. Video over IP and Instant Messaging are other examples. For the sake of generality, we use the term "session" throughout this document to refer to any type of session.

1.1. Applicability

The mechanisms defined in this document are applicable to controlled environments formed by either a single administrative domain or a set of administrative domains that closely coordinate their network policy and network design. The mechanisms defined in this document can be used for a session whose path spans over such a controlled environment where network layer admission control mechanisms are used, in order to elevate the session establishment probability through the controlled environment (thereby elevating the end to end session establishment probability). Let us consider the end to end environment illustrated in Figure 1 that comprises three separate administrative domains, each with 2 endpoints and each with Session Border Controller (SBC) elements ([I-D.ietf-sipping-sbc-funcs]) handling session handover at the domain boundaries.

```
  +----------+              +----------+              +----------+
  |Endpoint 1|              |Endpoint 3|              |Endpoint 5|
  +----------+              +----------+              +----------+
       |                         |                         |
       |                         |                         |
    +----+                    +----+                    +----+
    |SBC |                    |SBC |                    |SBC |
   ,|    |--.                ,-|    |-.                 ,|    |-.
  ,' +----+   `.           ,'  +----+  ` .           ,' +----+  \
 /   ISP       \          /    ISP      \            /    ISP    `.
/ Domain    +----+  +----+    Domain    +----+  +----+   Domain    \
(     A     |+----+  |+----+     B      |+----+  |+----+    C        )
 \(Controlled)||SBC |--||SBC |(Controlled)||SBC |--||SBC |(Controlled)/
  \         +|    | +|    |           +|    | +|    |          /
   `.         +----+   +----+            +----+   +----+      .'
    '+----+--'        `.  +----+   .'       '--+----+--'
      |    |           '--|    |--'            |    |
      |SBC |              |SBC |               |SBC |
      +----+              +----+               +----+
        |                   |                    |
        |                   |                    |
  +----------+        +----------+        +----------+
  |Endpoint 2|        |Endpoint 4|        |Endpoint 6|
  +----------+        +----------+        +----------+
```
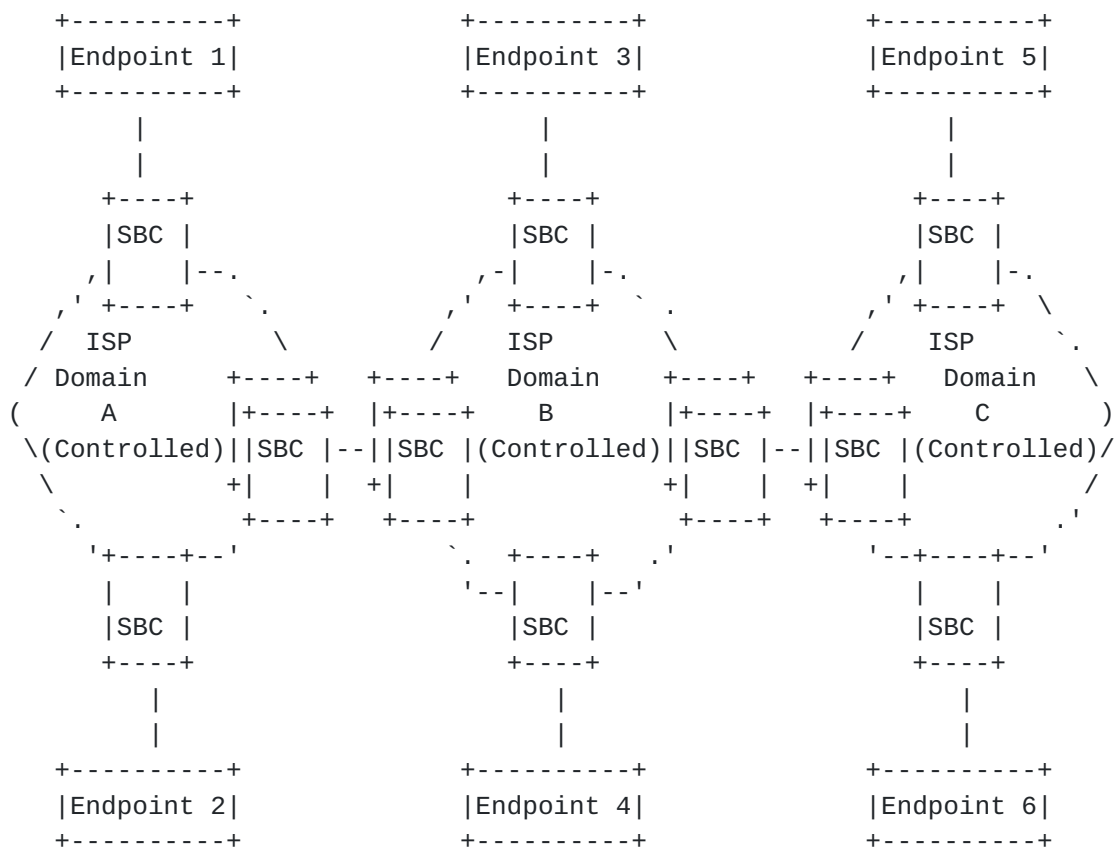
Figure 1: Example End to End Environment

Each domain is operating as a separate controlled environment and may
deploy a given combination of network mechanisms and network policies
within the given domain. For example, ISP Domain A , ISP Domain B
and ISP Domain C may each deploy a different Differentiated Services
([RFC2475]) policy in-between their own SBCs. As another example,
ISP Domain B may elect to deploy MPLS Traffic Engineering ([RFC2702])
within its domain while ISP Domain A and C may not. Similarly, each
domain administrator can make its own decision about whether to
deploy network layer admission control within his domain. If one
domain elects to do so, this can be achieved using RSVP signaling
between the ingress and egress SBC elements of that domain (i.e.,
RSVP signaling operates edge-to-edge and not end-to-end). With this
approach, network layer admission control may be deployed in one
domain regardless of whether it is deployed in the other domains on
the end to end path of sessions. Also, deploying network layer
admission control within one domain does not require any
collaboration or even pre-agreement with other domains since it
operates transparently from other domains (the only externally
visible impact might be on quality of service offered to the sessions
that transit through that domain). The mechanisms defined in this
document are applicable within a controlled environment that elects

to deploy network layer admission control using RSVP and handles
emergency communications. For example, ISP domain A and ISP domain C
may elect to use RSVP and the extensions defined in this document
within their respective domain while ISP domain B may not deploy
network layer admission control within his domain. In that case, a
session between Endpoint 1 and Endpoint 6 would benefit from network
layer admission control and resource reservation through domain A
network and domain C network. If that session is an emergency
session, the extensions defined in this document increase the
probability of admission of that particular session through domain A
and domain C, thereby increasing the end-to-end session establishment
probability.

As another example, all three domains shown in Figure 1 may elect to
deploy RSVP admission control and the extensions defined in this
document within their own domain. This would ensure that emergency
sessions are protected by resource reservation and elevated session
establishment probability through every domain on the end to end
path. But even in that case, RSVP signaling and the extensions
defined in this document need not operate end-to-end; rather they are
expected to operate edge-to-edge within each domain only (with RSVP
being terminated by the egress SBC on the egress edge of one domain
and regenerated by the ingress SBC on the ingress edge of the next
domain).

1.2. Related Technical Documents

[RFC4542] is patterned after [ITU.I.225] and describes an example of
one type of prioritized network layer admission control procedure
that may be used for emergency services operating over an IP network
infrastructure. It discusses initial session set up, as well as
operations after session establishment through maintenance of a
continuing call model of the status of all sessions. [RFC4542] also
describes how these network layer admission control procedures can be
realized using the Resource reSerVation Protocol [RFC2205] along with
its associated protocol suite and extensions, including those for
policy based admission control ([RFC2753], [RFC2750]), for user
authentication and authorization ([RFC3182]) and for integrity and
authentication of RSVP messages ([RFC2747], [RFC3097]). The Diameter
QoS Application ([I-D.ietf-dime-diameter-qos]) allows network
elements to interact with Diameter servers when allocating QoS
resources in the network and thus, is also a possible method for
authentication and authorization of RSVP reservations in the context
of emergency services.

[RFC4542] describes how the RSVP Signaled Preemption Priority Policy
Element specified in [RFC3181] can be used to enforce the session
preemption that may be needed by some emergency services. In

contrast to [RFC4542], this document specifies new RSVP extensions to increase the probability of session establishment without preemption. Engineered capacity techniques in the form of bandwidth allocation models are used to satisfy the "admission priority" required by an RSVP capable ETS network. In particular this document specifies two new RSVP Policy Elements allowing the admission priority to be conveyed inside RSVP signaling messages so that RSVP nodes can enforce selective bandwidth admission control decision based on the session admission priority. Appendix A of this document also provides examples of bandwidth allocation models which can be used by RSVP-routers to enforce such admission priority on every link.

1.3. Terminology

This document assumes the terminology defined in [RFC2753]. For convenience, the definition of a few key terms is repeated here:

o Policy Decision Point (PDP): The point where policy decisions are made.

o Local Policy Decision Point (LPDP): PDP local to the network element.

o Policy Enforcement Point (PEP): The point where the policy decisions are actually enforced.

o Policy Ignorant Node (PIN): A network element that does not explicitly support policy control using the mechanisms defined in [RFC2753].

2. 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 [RFC2119].

3. Overview of RSVP extensions and Operations

Let us consider the case where a session requiring ETS type service is to be established, and more specifically that the preference to be granted to this session is in terms of network layer "admission priority" (as opposed to preference granted through preemption of existing sessions). By "admission priority" we mean allowing that priority session to seize network layer resources from the engineered capacity that have been set-aside and not made available to normal sessions, or alternatively by allowing that session to seize

additional resources beyond the engineered capacity limits applied to normal sessions.

As described in [RFC4542], the session establishment can be conditioned to resource-based and policy-based network layer admission control achieved via RSVP signaling. In the case where the session control protocol is SIP, the use of RSVP-based admission control by SIP is specified in [RFC3312].

Devices involved in the session establishment are expected to be aware of the application-level priority requirements of emergency sessions. Again considering the case where the session control protocol is SIP, the SIP user agents can be made aware of the resource priority requirements in the case of an emergency session using the Resource-Priority Header mechanism specified in [RFC4412]. The end-devices involved in the upper-layer session establishment simply need to copy the application-level resource priority requirements (e.g. as communicated in SIP Resource-Priority Header) inside the new RSVP Application-Level Resource-Priority Policy Element defined in this document.

Conveying the application-level resource priority requirements inside the RSVP message allows this application level requirement to be mapped/remapped into a different RSVP "admission priority" at every administrative domain boundary based on the policy applicable in that domain. In a typical model (see [RFC2753]) where PDPs control PEPs at the periphery of the policy domain (e.g., in border routers), PDPs would interpret the RSVP Application-Level Resource-Priority Policy Element and map the requirement of the emergency session into an RSVP "admission priority" level. Then, PDPs would convey this information inside the new Admission Priority Policy Element defined in this document. This way, the RSVP admission priority can be communicated to downstream PEPs (i.e. RSVP Routers) of the same policy domain, which have LPDPs but no controlling PDP. In turn, this means the necessary RSVP Admission priority can be enforced at every RSVP hop, including all the (many) hops which do not have any understanding of Application-Level Resource-Priority semantics.

As an example of operation across multiple administrative domains, a first domain might decide to provide network layer admission priority to sessions of a given Application Level Resource Priority and map it into a high RSVP admission control priority inside the Admission Priority Policy Element; while a second domain may decide to not provide admission priority to sessions of this same Application Level Resource Priority and hence map it into a low RSVP admission control priority.

As another example of operation across multiple administrative

domains, we can consider the case where the resource priority header enumerates several namespaces, as explicitly allowed by [RFC4412], for support of scenarios where sessions traverse multiple administrative domains using different namespace. In that case, the relevant namespace can be used at each domain boundary to map into an RSVP Admission priority for that domain. It is not expected that the RSVP Application-Level Resource-Priority Header Policy Element would be taken into account at RSVP-hops within a given administrative domain. It is expected to be used at administrative domain boundaries only in order to set/reset the RSVP Admission Priority Policy Element.

The existence of pre-established inter-domain policy agreements or Service Level Agreements may avoid the need to take real-time action at administrative domain boundaries for mapping/remapping of admission priorities.

Mapping/remapping by PDPs may also be applied to boundaries between various signaling protocols, such as those advanced by the NSIS working group.

As can be observed, the framework described above for mapping/ remapping application level resource priority requirements into an RSVP admission priority can also be used together with [RFC3181] for mapping/remapping application level resource priority requirements into an RSVP preemption priority (when preemption is indeed needed). In that case, when processing the RSVP Application-Level Resource-Priority Policy Element, the PDPs at boundaries between administrative domains (or between various QoS signaling protocols) can map it into an RSVP "preemption priority" information. This Preemption priority information comprises a setup preemption level and a defending preemption priority level. This preemption priority information can then be encoded inside the Preemption Priority Policy Element of [RFC3181] and thus, can be taken into account at every RSVP-enabled network hop as discussed [RFC4542]. Appendix B provides examples of various hypothetical policies for emergency session handling, some of them involving admission priority, some of them involving both admission priority and preemption priority. Appendix B also identifies how the Application-Level Resource Priority need to be mapped into RSVP policy elements by the PDPs to realize these policies.

3.1. Operations of Admission Priority

The RSVP Admission Priority policy element defined in this document allows admission bandwidth to be allocated preferentially to an authorized priority service. Multiple models of bandwidth allocation MAY be used to that end.

A number of bandwidth allocation models have been defined in the IETF
for allocation of bandwidth across different classes of traffic
trunks in the context of Diffserv-aware MPLS Traffic Engineering.
Those include the Maximum Allocation Model (MAM) defined in
[RFC4125], the Russian Dolls Model (RDM) specified in [RFC4127] and
the Maximum Allocation model with Reservation (MAR) defined in
[RFC4126]. These same models MAY however be applied for allocation
of bandwidth across different levels of admission priority as defined
in this document. Appendix A provides an illustration of how these
bandwidth allocation models can be applied for such purposes and
introduces an additional bandwidth allocation model that we term the
Priority Bypass Model (PrBM). It is important to note that the
models described and illustrated in Appendix A are only informative
and do not represent a recommended course of action.

We can see in these examples, that the RSVP Admission Priority may
effectively influence the fundamental admission control decision of
RSVP (for example by determining which bandwidth pool is to be used
by RSVP for performing its fundamental bandwidth allocation). In
that sense, we observe that the policy control and admission control
are not separate logics but instead somewhat blended.

4. New Policy Elements

The Framework document for policy-based admission control [RFC2753]
describes the various components that participate in policy decision
making (i.e., PDP, PEP and LPDP).

As described in section 2 of the present document, the Application-
Level Resource Priority Policy Element and the Admission Priority
Policy Element serve different roles in this framework:

o the Application-Level Resource Priority Policy Element conveys
   application level information and is processed by PDPs

o the emphasis of Admission Priority Policy Element is to be simple,
   stateless, and light-weight such that it can be processed
   internally within a node's LPDP. It can then be enforced
   internally within a node's PEP. It is set by PDPs based on
   processing of the Application-Level Resource Priority Policy
   Element.

[RFC2750] defines extensions for supporting generic policy based
admission control in RSVP. These extensions include the standard
format of POLICY_DATA objects and a description of RSVP handling of
policy events.

The POLICY_DATA object contains one or more of Policy Elements, each
representing a different (and perhaps orthogonal) policy. As an
example, [RFC3181] specifies the Preemption Priority Policy Element.
This document defines two new Policy Elements called:

o   the Admission Priority Policy Element

o   the Application-Level Resource Priority Policy Element

4.1. Admission Priority Policy Element

The format of the Admission Priority policy element is as shown in
Figure 2:

```
     0               0 0             1 1             2 2             3
     0   . . .   7 8   . . .   5 6     . . . 3 4   . . .         1
    +-------------+-------------+-------------+-------------+
    |     Length               | P-Type = ADMISSION_PRI    |
    +-------------+-------------+-------------+-------------+
    | Flags       | M. Strategy | Error Code  | Reserved    |
    +-------------+-------------+-------------+-------------+
    |               Reserved                  |Adm. Priority|
    +-------------------------+-------------------------+
```
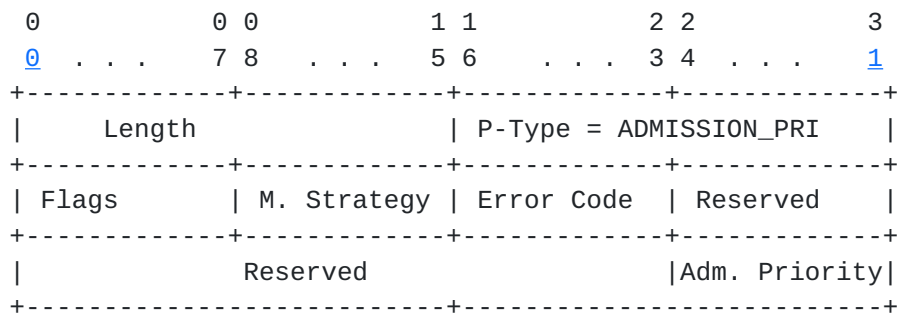
                Figure 2: Admission Priority Policy Element

where:

o   Length: 16 bits

    *   Always 12. The overall length of the policy element, in bytes.

o   P-Type: 16 bits

    *   ADMISSION_PRI = To be allocated by IANA (see "IANA
        Considerations" section)

o   Flags: Reserved

    *   SHALL be set to zero on transmit and SHALL be ignored on
        reception

o   Merge Strategy: 8 bits (only applicable to multicast flows)

    *   values are defined by corresponding registry maintained by IANA
        (see "IANA Considerations" section)

o Error code: 8 bits (only applicable to multicast flows)

   * values are defined by corresponding registry maintained by IANA
      (see "IANA Considerations" section)

o Reserved: 8 bits

   * SHALL be set to zero on transmit and SHALL be ignored on
      reception

o Reserved: 24 bits

   * SHALL be set to zero on transmit and SHALL be ignored on
      reception

o Adm. Priority (Admission Priority): 8 bits (unsigned)

   * The admission control priority of the flow, in terms of access
      to network bandwidth in order to provide higher probability of
      session completion to selected flows. Higher values represent
      higher Priority. A given Admission Priority is encoded in this
      information element using the same value as when encoded in the
      "Admission Priority" field of the "Admission Priority"
      parameter defined in [I-D.ietf-nsis-qspec], or in the
      "Admission Priority" parameter defined in
      [I-D.ietf-dime-qos-parameters]. In other words, a given value
      inside the Admission Priority information element defined in
      the present document, inside the [I-D.ietf-nsis-qspec]
      Admission Priority field or inside the
      [I-D.ietf-dime-qos-parameters] Admission Priority parameter,
      refers to the same admission priority. Bandwidth allocation
      models such as those described in Appendix A are to be used by
      the RSVP router to achieve such increased probability of
      session establishment. The admission priority value
      effectively indicates which bandwidth constraint(s) of the
      bandwidth constraint model in use is(are) applicable to
      admission of this RSVP reservation.

Note that the Admission Priority Policy Element does NOT indicate
that this RSVP reservation is to preempt any other RSVP reservation.
If a priority session justifies both admission priority and
preemption priority, the corresponding RSVP reservation needs to
carry both an Admission Priority Policy Element and a Preemption
Priority Policy Element. The Admission Priority and Preemption
Priority are handled by LPDPs and PEPs as separate mechanisms. They
can be used one without the other, or they can be used both in
combination.

4.1.1. Admission Priority Merging Rules

   This section discusses alternatives for dealing with RSVP admission
   priority in case of merging of reservations. As merging is only
   applicable to multicast, this section also only applies to multicast
   sessions.

   The rules for merging Admission Priority Policy Elements are defined
   by the value encoded inside the Merge Strategy field in accordance
   with the corresponding IANA registry. The merge strategies (and
   associated values) defined by the present document are the same as
   those defined in [RFC3181] for merging Preemption Priority Policy
   Elements (see "IANA Considerations" section).

   The only difference with [RFC3181] is that this document does not
   recommend any merge strategies for Admission Priority, while
   [RFC3181] recommends the first of these merge strategies for
   Preemption Priority. Note that with the Admission Priority (as is
   the case with the Preemption Priority), "Take highest priority"
   translates into "take the highest numerical value".

4.2. Application-Level Resource Priority Policy Element

   The format of the Application-Level Resource Priority policy element
   is as shown in Figure 3:

```
      0             0 0           1 1           2 2           3
      0   . . .     7 8   . . .   5 6   . . . 3 4   . . .     1
     +-------------+-------------+-------------+-------------+
     | Length                    | P-Type = APP_RESOURCE_PRI |
     +-------------+-------------+-------------+-------------+
     //     ALRP List                                       //
     +-------------------------+-------------------------+
```
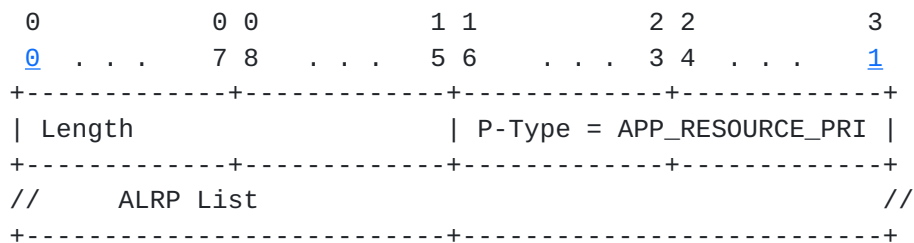
      Figure 3: Application-Level Resource Priority Policy Element

   where:

   o Length:

      * The length of the policy element (including the Length and
         P-Type) is in number of octets (MUST be a multiple of 4) and
         indicates the end of the ALRP list.

   o P-Type: 16 bits

      * APP_RESOURCE_PRI = To be allocated by IANA (see "IANA
         Considerations" section)

o ALRP List:

   * List of ALRP where each ALRP is encoded as shown in Figure 4.

ALRP:
```
    0              0 0           1 1           2 2           3
    0   . . .     7 8   . . .   5 6    . . .  3 4  . . .     1
   +--------------------------+------------+------------+
   |      ALRP Namespace      | Reserved   |ALRP Priority|
   +--------------------------+--------------------------+
```
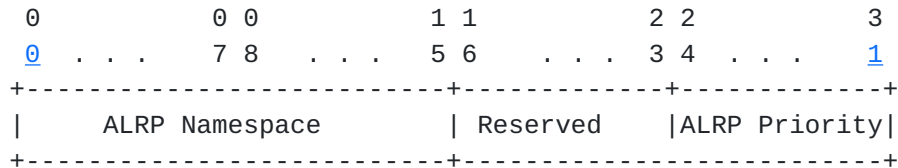
            Figure 4: Application-Level Resource Priority

where:

o ALRP Namespace (Application-Level Resource Priority Namespace): 16
   bits (unsigned)

   * Contains a numerical value identifying the namespace of the
      application-level resource priority. This value is encoded as
      per the "Resource-Priority Namespaces" IANA registry. (See
      IANA Considerations section for the request to IANA to extend
      the registry to include this numerical value).

o Reserved: 8 bits

   * SHALL be set to zero on transmit and SHALL be ignored on
      reception.

o ALRP Priority: (Application-Level Resource Priority Priority): 8
   bits (unsigned)

   * Contains the priority value within the namespace of the
      application-level resource priority. This value is encoded as
      per the "Resource-Priority Priority-Value" IANA registry. (See
      IANA Considerations section for the request to IANA to extend
      the registry to include this numerical value).

4.2.1. Application-Level Resource Priority Modifying and Merging Rules

   When POLICY_DATA objects are protected by integrity, LPDPs should not
   attempt to modify them. They MUST be forwarded as-is to ensure their
   security envelope is not invalidated.

   In case of multicast, when POLICY_DATA objects are not protected by
   integrity, LPDPs MAY merge incoming Application-Level Resource
   Priority elements to reduce their size and number. When they do
   merge those, LPDPs MUST do so according to the following rule:

o The ALRP List in the outgoing APP_RESOURCE_PRI element MUST list
   all the ALRPs appearing in the ALRP List of an incoming
   APP_RESOURCE_PRI element. A given ALRP MUST NOT appear more than
   once. In other words, the outgoing ALRP List is the union of the
   incoming ALRP Lists that are merged.

As merging is only applicable to Multicast, this rule only applies to
Multicast sessions.

4.3. Default Handling

As specified in section 4.2 of [RFC2750], Policy Ignorant Nodes
(PINs) implement a default handling of POLICY_DATA objects ensuring
that those objects can traverse PIN nodes in transit from one PEP to
another. This applies to the situations where POLICY_DATA objects
contain the Admission Priority Policy Element and the ALRP Policy
Element specified in this document, so that those can traverse PIN
nodes.

Section 4.2 of [RFC2750] also defines a similar default behavior for
policy-capable nodes that do not recognized a particular Policy
Element. This applies to the Admission Priority Policy Element and
the ALRP Policy Element specified in this document, so that those can
traverse policy-capable nodes that do not support this document.

5. Security Considerations

As this document defines extensions to RSVP, the security
considerations of RSVP apply. Those are discussed in [RFC2205],
[RFC4230] and [I-D.ietf-tsvwg-rsvp-security-groupkeying].

A subset of RSVP messages are signaled with the Router Alert Option
(RAO)([RFC2113],[RFC2711]). However, some network administrators
activate mechanisms at the edge of their administrative domain to
protect against potential Denial Of Service (DOS) attacks associated
with RAO. This may include hiding of the RAO to downstream interior
routers in the domain (as recommended by default over an MPLS network
in [I-D.dasmith-mpls-ip-options]) or complete blocking of packets
received with RAO at the administrative boundary. As the mechanisms
defined in this document rely on RSVP, their usage assume that such
protection against RAO packets are not activated in a way that
prevents RSVP processing on relevant interfaces or routers of the
controlled environments electing to deploy these mechanisms.
Nonetheless, it is recommended that protection mechanisms be
activated against potential DOS attacks through RAO even when RAO
message are processed. This may include rate limiting of incoming
RAO packets (e.g. at interface and/or router level). This may also

include deploying an RSVP architecture whereby interior routers are
not exposed to any RSVP messages associated with end to end
reservations (such as the architecture defined in
[I-D.ietf-tsvwg-rsvp-l3vpn]). We observe that the risks and security
measures associated with processing of RAO messages at an
administrative domain edge are fundamentally similar to those
involved with other forms of control plane interactions allowed at
administrative domain edges, such as routing or multicast routing
interactions allowed between a customer and his Internet Service
Provider, MPLS VPN ([RFC4364] Service Provider , [RFC4659]) or MPLS
MVPN ([I-D.ietf-l3vpn-2547bis-mcast]) Service Provider.

The ADMISSION_PRI and APP_RESOURCE_PRI Policy Elements defined in
this document are signaled by RSVP through encapsulation in a Policy
Data object as defined in [RFC2750]. Therefore, like any other
Policy Elements, their integrity can be protected as discussed in
section 6 of [RFC2750] by two optional security mechanisms. The
first mechanism relies on RSVP Authentication as specified in
[RFC2747] and [RFC3097] to provide a chain of trust when all RSVP
nodes are policy capable. With this mechanism, the INTEGRITY object
is carried inside RSVP messages. The second mechanism relies on the
INTEGRITY object within the POLICY_DATA object to guarantee integrity
between RSVP Policy Enforcement Points (PEPs) that are not RSVP
neighbors.

5.1. Use of RSVP Authentication between RSVP neighbors

This mechanism can be used between RSVP neighbors that are policy
capable. The RSVP neighbors use shared keys to compute the
cryptographic signature of the RSVP message.
[I-D.ietf-tsvwg-rsvp-security-groupkeying] discusses key types, key
provisioning methods as well as their respective applicability.

5.2. Use of INTEGRITY object within the POLICY_DATA object

The INTEGRITY object within the POLICY_DATA object can be used to
guarantee integrity between non-neighboring RSVP PEPs.

Details for computation of the content of the INTEGRITY object can be
found in Appendix B of [RFC2750]. This states that the Policy
Decision Point (PDP), at its discretion, and based on destination
PEP/PDP or other criteria, selects an Authentication Key and the hash
algorithm to be used. Keys to be used between PDPs can be
distributed manually or via standard key management protocol for
secure key distribution.

Note that where non-RSVP hops may exist in between RSVP hops, as well
as where RSVP capable Policy Ignorant Nodes (PINs) may exist in

between PEPs, it may be difficult for the PDP to determine what is
the destination PDP for a POLICY_DATA object contained in some RSVP
messages (such as a Path message). This is because in those cases
the next PEP is not known at the time of forwarding the message. In
this situation, key shared across multiple PDPs may be used. This is
conceptually similar to the use of key shared across multiple RSVP
neighbors discussed in [I-D.ietf-tsvwg-rsvp-security-groupkeying].
We observe also that this issue may not exist in some deployment
scenarios where a single (or low number of) PDP is used to control
all the PEPs of a region (such as an administrative domain). In such
scenarios, it may be easy for a PDP to determine what is the next hop
PDP, even when the next hop PEP is not known, simply by determining
what is the next region that will be traversed (say based on the
destination address).

6. IANA Considerations

As specified in [RFC2750], Standard RSVP Policy Elements (P-type
values) are to be assigned by IANA as per "IETF Consensus" policy
following the policies outlined in [RFC2434] (this policy is now
called "IETF Review" as per [RFC5226]) .

IANA needs to allocate two P-Types from the Standard RSVP Policy
Element range:

o one P-Type to the Admission Priority Policy Element

o one P-Type to the Application-Level Resource Priority Policy
   Element.

In section 3.1, the present document defines a Merge Strategy field
inside the Admission Priority policy element. IANA needs to create a
registry for this field and allocate the following values:

o 1: Take priority of highest QoS

o 2: Take highest priority

o 3: Force Error on heterogeneous merge

Following the policies outlined in [RFC5226], numbers in the range
4-127 are allocated according to the "IETF Review" policy, numbers in
the range 128-240 as "First Come First Served" and numbers between
241-255 are reserved for "Private Use". Value 0 is Reserved (for
consistency with [RFC3181] Merge Strategy values).

In section 3.1, the present document defines an Error Code field

inside the Admission Priority policy element. IANA needs to create a
registry for this field and allocate the following values:

o 0: NO_ERROR Value used for regular ADMISSION_PRI elements

o 2: HETEROGENEOUS This element encountered heterogeneous merge

Following the policies outlined in [RFC5226], numbers in the range
3-127 are allocated according to the "IETF Review" policy, numbers in
the range 128-240 as "First Come First Served" and numbers between
241-255 are reserved for "Private Use". Value 1 is Reserved (for
consistency with [RFC3181] Error Code values).

The present document defines an ALRP Namespace field in section 3.2
that contains a numerical value identifying the namespace of the
application-level resource priority. The IANA already maintains the
Resource-Priority Namespaces registry (under the SIP Parameters)
listing all such namespace. However, that registry does not
currently allocate a numerical value to each namespace. Hence, this
document requests the IANA to extend the Resource-Priority Namespace
registry in the following ways:

o a new column should be added to the registry

o the title of the new column should be "Namespace Numerical Value
   *"

o in the Legend, add a line saying "Namespace Numerical Value = the
   unique numerical value identifying the namespace"

o add a line at the bottom of the registry stating the following "*
   : [RFCXXX] " where XXX is the RFC number of the present document

o allocate an actual numerical value to each namespace in the
   registry and state that value in the new "Namespace numerical
   Value *" column.

A numerical value should be allocated immediately by IANA to all
existing namespace. Then, in the future, IANA should automatically
allocate a numerical value to any new namespace added to the
registry.

The present document defines an ALRP Priority field in section 3.2
that contains a numerical value identifying the actual application-
level resource priority within the application-level resource
priority namespace. The IANA already maintains the Resource-Priority
Priority-values registry (under the SIP Parameters) listing all such
priorities. However, that registry does not currently allocate a

   numerical value to each priority-value. Hence, this document
   requests the IANA to extend the Resource-Priority Priority-Values
   registry in the following ways:

   o for each namespace, the registry should be structured with two
      columns

   o the title of the first column should read "Priority Values (least
      to greatest)"

   o the first column should list all the values currently defined in
      the registry (e.g. for the drsn namespace: "routine", "priority",
      "immediate", "flash", "flash-override", "flash-override-override"
      for the drsn namespace)

   o the title of the second column should read "Priority Numerical
      Value *"

   o At the bottom of the registry, add a "Legend" with a line saying
      "Priority Numerical Value = the unique numerical value identifying
      the priority within a namespace"

   o add a line at the bottom of the registry stating the following "*
      : [RFCXXX] " where XXX is the RFC number of the present document

   o allocate an actual numerical value to each and state that value in
      the new "Priority Numerical Value *" column.

   A numerical value should be allocated immediately by IANA to all
   existing priority. Then, in the future, IANA should automatically
   allocate a numerical value to any new namespace added to the
   registry. The numerical value must be unique within each namespace.
   For the initial allocation, within each namespace, values should be
   allocated in decreasing order ending with 0 (so that the greatest
   priority is always allocated value 0). For example, in the drsn
   namespace, "routine" would be allocated numerical value 5 and "flash-
   override-override" would be allocated numerical value 0.

7. Acknowledgments

   We would like to thank An Nguyen for his encouragement to address
   this topic and ongoing comments. Also, this document borrows heavily
   from some of the work of S. Herzog on Preemption Priority Policy
   Element [RFC3181]. Dave Oran and Janet Gunn provided useful input
   into this document. Thanks to Magnus Westerlund, Cullen Jennings and
   Ross Callon for helping clarify applicability of the mechanisms
   defined in this document.

8. References

8.1. Normative References

[RFC2113]   Katz, D., "IP Router Alert Option", RFC 2113,
            February 1997.

[RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
            Requirement Levels", BCP 14, RFC 2119, March 1997.

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

[RFC2434]   Narten, T. and H. Alvestrand, "Guidelines for Writing an
            IANA Considerations Section in RFCs", BCP 26, RFC 2434,
            October 1998.

[RFC2711]   Partridge, C. and A. Jackson, "IPv6 Router Alert Option",
            RFC 2711, October 1999.

[RFC2747]   Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic
            Authentication", RFC 2747, January 2000.

[RFC2750]   Herzog, S., "RSVP Extensions for Policy Control",
            RFC 2750, January 2000.

[RFC3097]   Braden, R. and L. Zhang, "RSVP Cryptographic
            Authentication -- Updated Message Type Value", RFC 3097,
            April 2001.

[RFC3181]   Herzog, S., "Signaled Preemption Priority Policy Element",
            RFC 3181, October 2001.

[RFC3261]   Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston,
            A., Peterson, J., Sparks, R., Handley, M., and E.
            Schooler, "SIP: Session Initiation Protocol", RFC 3261,
            June 2002.

[RFC4412]   Schulzrinne, H. and J. Polk, "Communications Resource
            Priority for the Session Initiation Protocol (SIP)",
            RFC 4412, February 2006.

[RFC5226]   Narten, T. and H. Alvestrand, "Guidelines for Writing an
            IANA Considerations Section in RFCs", BCP 26, RFC 5226,
            May 2008.

8.2. Informative References

   [I-D.dasmith-mpls-ip-options]
              Jaeger, W., Mullooly, J., Scholl, T., and D. Smith,
              "Requirements for Label Edge Router Forwarding of IPv4
              Option Packets", draft-dasmith-mpls-ip-options-01 (work in
              progress), October 2008.

   [I-D.ietf-dime-diameter-qos]
              Sun, D., McCann, P., Tschofenig, H., Tsou, T., Doria, A.,
              and G. Zorn, "Diameter Quality of Service Application",
              draft-ietf-dime-diameter-qos-07 (work in progress),
              December 2008.

   [I-D.ietf-dime-qos-parameters]
              Korhonen, J., Tschofenig, H., and E. Davies, "Quality of
              Service Parameters for Usage with Diameter",
              draft-ietf-dime-qos-parameters-09 (work in progress),
              January 2009.

   [I-D.ietf-l3vpn-2547bis-mcast]
              Aggarwal, R., Bandi, S., Cai, Y., Morin, T., Rekhter, Y.,
              Rosen, E., Wijnands, I., and S. Yasukawa, "Multicast in
              MPLS/BGP IP VPNs", draft-ietf-l3vpn-2547bis-mcast-07 (work
              in progress), July 2008.

   [I-D.ietf-nsis-qspec]
              Bader, A., Kappler, C., and D. Oran, "QoS NSLP QSPEC
              Template", draft-ietf-nsis-qspec-21 (work in progress),
              November 2008.

   [I-D.ietf-sipping-sbc-funcs]
              Hautakorpi, J., Camarillo, G., Penfield, B., Hawrylyshen,
              A., and M. Bhatia, "Requirements from SIP (Session
              Initiation Protocol) Session Border Control Deployments",
              draft-ietf-sipping-sbc-funcs-08 (work in progress),
              January 2009.

   [I-D.ietf-tsvwg-rsvp-l3vpn]
              Davie, B., Faucheur, F., and A. Narayanan, "Support for
              RSVP in Layer 3 VPNs", draft-ietf-tsvwg-rsvp-l3vpn-01
              (work in progress), November 2008.

   [I-D.ietf-tsvwg-rsvp-security-groupkeying]
              Behringer, M. and F. Faucheur, "Applicability of Keying
              Methods for RSVP Security",
              draft-ietf-tsvwg-rsvp-security-groupkeying-02 (work in
              progress), November 2008.

   [NCS]   "GETS Home Page", <http://gets.ncs.gov>.

   [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z.,
              and W. Weiss, "An Architecture for Differentiated
              Services", RFC 2475, December 1998.

   [RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J.
              McManus, "Requirements for Traffic Engineering Over MPLS",
              RFC 2702, September 1999.

   [RFC2753] Yavatkar, R., Pendarakis, D., and R. Guerin, "A Framework
              for Policy-based Admission Control", RFC 2753,
              January 2000.

   [RFC3182] Yadav, S., Yavatkar, R., Pabbati, R., Ford, P., Moore, T.,
              Herzog, S., and R. Hess, "Identity Representation for
              RSVP", RFC 3182, October 2001.

   [RFC3312] Camarillo, G., Marshall, W., and J. Rosenberg,
              "Integration of Resource Management and Session Initiation
              Protocol (SIP)", RFC 3312, October 2002.

   [RFC3689] Carlberg, K. and R. Atkinson, "General Requirements for
              Emergency Telecommunication Service (ETS)", RFC 3689,
              February 2004.

   [RFC3690] Carlberg, K. and R. Atkinson, "IP Telephony Requirements
              for Emergency Telecommunication Service (ETS)", RFC 3690,
              February 2004.

   [RFC4125] Le Faucheur, F. and W. Lai, "Maximum Allocation Bandwidth
              Constraints Model for Diffserv-aware MPLS Traffic
              Engineering", RFC 4125, June 2005.

   [RFC4126] Ash, J., "Max Allocation with Reservation Bandwidth
              Constraints Model for Diffserv-aware MPLS Traffic
              Engineering & Performance Comparisons", RFC 4126,
              June 2005.

   [RFC4127] Le Faucheur, F., "Russian Dolls Bandwidth Constraints
              Model for Diffserv-aware MPLS Traffic Engineering",
              RFC 4127, June 2005.

   [RFC4190] Carlberg, K., Brown, I., and C. Beard, "Framework for
              Supporting Emergency Telecommunications Service (ETS) in
              IP Telephony", RFC 4190, November 2005.

   [RFC4230] Tschofenig, H. and R. Graveman, "RSVP Security

                  Properties", RFC 4230, December 2005.

   [RFC4364]   Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private
               Networks (VPNs)", RFC 4364, February 2006.

   [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.

   [RFC4659]   De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur,
               "BGP-MPLS IP Virtual Private Network (VPN) Extension for
               IPv6 VPN", RFC 4659, September 2006.

Appendix A. Examples of Bandwidth Allocation Model for Admission Priority

   Sections A.1 and A.2 respectively illustrate how the Maximum
   Allocation Model (MAM) ([RFC4125]) and the Russian Dolls Model (RDM)
   ([RFC4127]) can be used for support of admission priority. The
   Maximum Allocation model with Reservation (MAR) ([RFC4126]) could
   also be used in a similar manner for support of admission priority.
   Section A.3 illustrates how a simple "Priority Bypass Model" can also
   be used for support of admission priority.

   For simplicity, operations with only a single "priority" level
   (beyond non-priority) are illustrated here; However, the reader will
   appreciate that operations with multiple priority levels can easily
   be supported with these models.

   In all the figures below:

       x represents a non-priority session

       o represents a priority session

A.1. Admission Priority with Maximum Allocation Model (MAM)

   This section illustrates operations of admission priority when a
   Maximum Allocation Model (MAM) is used for bandwidth allocation
   across non-priority traffic and priority traffic. A property of the
   Maximum Allocation Model is that priority traffic can not use more
   than the bandwidth made available to priority traffic (even if the
   non-priority traffic is not using all of the bandwidth available for
   it).

```
                     ------------------------
           ^  ^  ^  |                  |  ^
           .  .  .  |                  |  .
  Total    .  .  .  |                  |  .   Bandwidth
         (1)(2)(3)  |                  |  .   Available
  Engi-    .  .  .  |                  |  .   for non-priority use
 neered   .or.or.   |                  |  .
           .  .  .  |                  |  .
 Capacity. .  .  |                  |  .
           v  .  .  |                  |  v
              .  .  |--------------|  ---
             v  .  |                  |  ^
                .  |                  |  .   Bandwidth available for
                v  |                  |  v   priority use
                   ------------------------
```
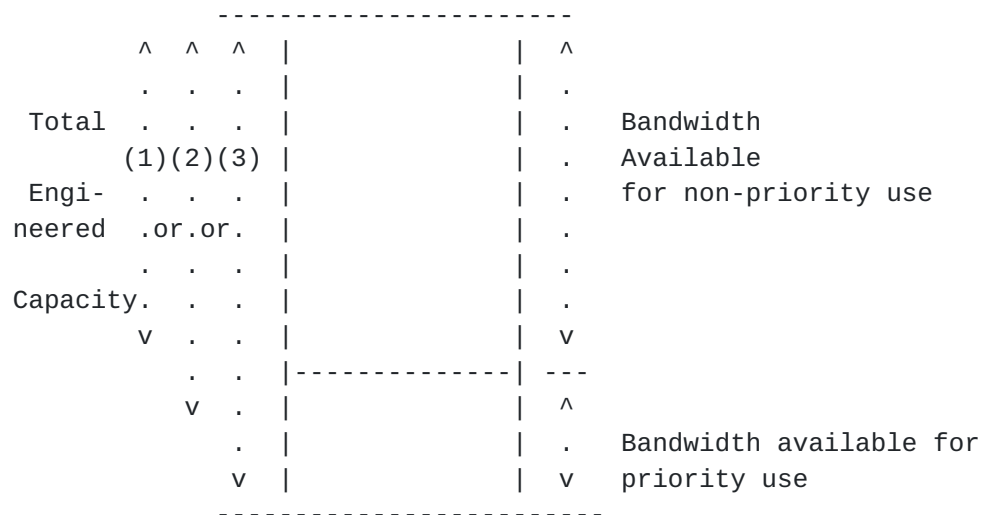
                   Figure 5: MAM Bandwidth Allocation

   Figure 5 shows a link within a routed network conforming to this
   document. On this link are two amounts of bandwidth available to two
   types of traffic: non-priority and priority.

   If the non-priority traffic load reaches the maximum bandwidth
   available for non-priority, no additional non-priority sessions can
   be accepted even if the bandwidth reserved for priority traffic is
   not currently fully utilized.

   With the Maximum Allocation Model, in the case where the priority
   load reaches the maximum bandwidth reserved for priority sessions, no
   additional priority sessions can be accepted.

   As illustrated in Figure 5, an operator may map the MAM model onto
   the Engineered Capacity limits according to different policies. At
   one extreme, where the proportion of priority traffic is reliably
   known to be fairly small at all times and where there may be some
   safety margin factored in the engineered capacity limits, the
   operator may decide to configure the bandwidth available for non-
   priority use to the full engineered capacity limits; effectively
   allowing the priority traffic to ride within the safety margin of
   this engineered capacity. This policy can be seen as an economically
   attractive approach as all of the engineered capacity is made
   available to non-priority sessions. This policy is illustrated as
   (1) in Figure 5. As an example, if the engineered capacity limit on
   a given link is X, the operator may configure the bandwidth available
   to non-priority traffic to X, and the bandwidth available to priority
   traffic to 5% of X. At the other extreme, where the proportion of
   priority traffic may be significant at times and the engineered
   capacity limits are very tight, the operator may decide to configure

the bandwidth available to non-priority traffic and the bandwidth
available to priority traffic such that their sum is equal to the
engineered capacity limits. This guarantees that the total load
across non-priority and priority traffic is always below the
engineered capacity and, in turn, guarantees there will never be any
QoS degradation. However, this policy is less attractive
economically as it prevents non-priority sessions from using the full
engineered capacity, even when there is no or little priority load,
which is the majority of time. This policy is illustrated as (3) in
Figure 5. As an example, if the engineered capacity limit on a given
link is X, the operator may configure the bandwidth available to non-
priority traffic to 95% of X, and the bandwidth available to priority
traffic to 5% of X. Of course, an operator may also strike a balance
anywhere in between these two approaches. This policy is illustrated
as (2) in Figure 5.

Figure 6 shows some of the non-priority capacity of this link being
used.

```
              -----------------------
         ^  ^  ^  |                 | ^
         .  .  .  |                 | .
 Total   .  .  .  |                 | .   Bandwidth
         .  .  .  |                 | .   Available
 Engi-   .  .  .  |                 | .   for non-priority use
neered  .or.or.  |xxxxxxxxxxxxxx|   .
         .  .  .  |xxxxxxxxxxxxxx|   .
Capacity.  .  .  |xxxxxxxxxxxxxx|   .
         v  .  .  |xxxxxxxxxxxxxx|   v
            .  .  |--------------| ---
          v  .  |                 | ^
             .  |                 | .   Bandwidth available for
           v  |                 | v   priority use
              ------------------------
```
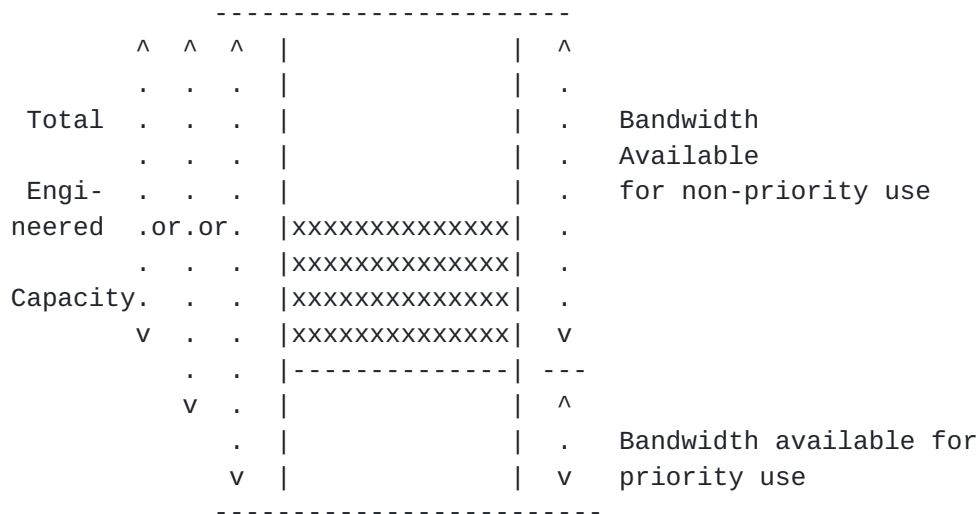
Figure 6: Partial load of non-priority calls

Figure 7 shows the same amount of non-priority load being used at
this link, and a small amount of priority bandwidth being used.

```
                   ----------------------
           ^  ^  ^  |                    | ^
           .  .  .  |                    | .
   Total   .  .  .  |                    | .   Bandwidth
           .  .  .  |                    | .   Available
   Engi-   .  .  .  |                    | .   for non-priority use
  neered  .or.or.   |xxxxxxxxxxxxxx|   .
           .  .  .  |xxxxxxxxxxxxxx|   .
 Capacity. .  .     |xxxxxxxxxxxxxx|   .
           v  .  .  |xxxxxxxxxxxxxx|   v
              .  .  |--------------| ---
           v  .  |                    | ^
              .  |                    | .   Bandwidth available for
              v  |oooooooooooooo|   v   priority use
                   ------------------------
```

Figure 7: Partial load of non-priority calls & partial load of
                    priority calls Calls

Figure 8 shows the case where non-priority load equates or exceeds
the maximum bandwidth available to non-priority traffic. Note that
additional non-priority sessions would be rejected even if the
bandwidth reserved for priority sessions is not fully utilized.

```
                   ----------------------
           ^  ^  ^  |xxxxxxxxxxxxxx|   ^
           .  .  .  |xxxxxxxxxxxxxx|   .
   Total   .  .  .  |xxxxxxxxxxxxxx|   .   Bandwidth
           .  .  .  |xxxxxxxxxxxxxx|   .   Available
   Engi-   .  .  .  |xxxxxxxxxxxxxx|   .   for non-priority use
  neered  .or.or.   |xxxxxxxxxxxxxx|   .
           .  .  .  |xxxxxxxxxxxxxx|   .
 Capacity. .  .     |xxxxxxxxxxxxxx|   .
           v  .  .  |xxxxxxxxxxxxxx|   v
              .  .  |--------------| ---
           v  .  |                    | ^
              .  |                    | .   Bandwidth available for
              v  |oooooooooooooo|   v   priority use
                   ------------------------
```
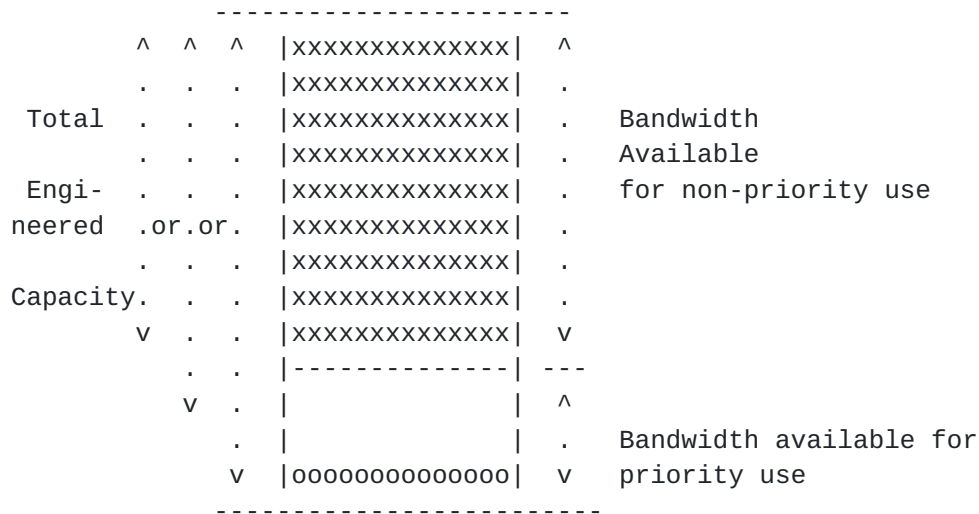
Figure 8: Full non-priority load & partial load of priority calls

Figure 9 shows the case where the priority traffic equates or exceeds
the bandwidth reserved for such priority traffic.

In that case additional priority sessions could not be accepted.
Note that this does not mean that such sessions are dropped
altogether: they may be handled by mechanisms, which are beyond the

scope of this particular document (such as establishment through
preemption of existing non-priority sessions, or such as queuing of
new priority session requests until capacity becomes available again
for priority traffic).

```
                    ----------------------
          ^  ^  ^   |xxxxxxxxxxxxxx|   ^
          .  .  .   |xxxxxxxxxxxxxx|   .
  Total   .  .  .   |xxxxxxxxxxxxxx|   .   Bandwidth
          .  .  .   |xxxxxxxxxxxxxx|   .   Available
  Engi-   .  .  .   |xxxxxxxxxxxxxx|   .   for non-priority use
 neered  .or.or.    |xxxxxxxxxxxxxx|   .
          .  .  .   |xxxxxxxxxxxxxx|   .
Capacity. .  .     |              |   .
          v  .  .   |              |   v
             .  .   |--------------|  ---
          v  .      |oooooooooooooo|   ^
             .      |oooooooooooooo|   .   Bandwidth available for
             v      |oooooooooooooo|   v   priority use
                    ------------------------
```
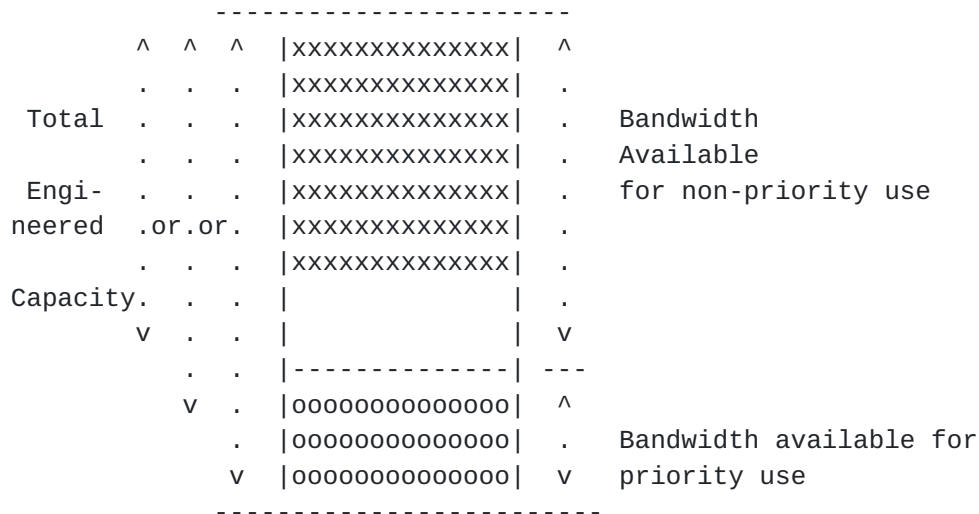
        Figure 9: Partial non-priority load & Full priority load

A.2. Admission Priority with Russian Dolls Model (RDM)

   This section illustrates operations of admission priority when a
   Russian Dolls Model (RDM) is used for bandwidth allocation across
   non-priority traffic and priority traffic. A property of the Russian
   Dolls Model is that priority traffic can use the bandwidth which is
   not currently used by non-priority traffic.

   As with the MAM model, an operator may map the RDM model onto the
   Engineered Capacity limits according to different policies. The
   operator may decide to configure the bandwidth available for non-
   priority use to the full engineered capacity limits; As an example,
   if the engineered capacity limit on a given link is X, the operator
   may configure the bandwidth available to non-priority traffic to X,
   and the bandwidth available to non-priority and priority traffic to
   105% of X.

   Alternatively, the operator may decide to configure the bandwidth
   available to non-priority and priority traffic to the engineered
   capacity limits; As an example, if the engineered capacity limit on a
   given link is X, the operator may configure the bandwidth available
   to non-priority traffic to 95% of X, and the bandwidth available to
   non-priority and priority traffic to X.

   Finally, the operator may decide to strike a balance in between. The

considerations presented for these policies in the previous section in the MAM context are equally applicable to RDM.

Figure 10 shows the case where only some of the bandwidth available to non-priority traffic is being used and a small amount of priority traffic is in place. In that situation both new non-priority sessions and new priority sessions would be accepted.

```
                 ---------------------------------------
                |xxxxxxxxxxxxxx|   .                    ^
                |xxxxxxxxxxxxxx|   . Bandwidth          .
                |xxxxxxxxxxxxxx|   . Available for      .
                |xxxxxxxxxxxxxx|   . non-priority       .
                |xxxxxxxxxxxxxx|   . use                .
                |xxxxxxxxxxxxxx|   .                    . Bandwidth
                |              |   .                    . available for
                |              |   v                    . non-priority
                |--------------|  ---                   . and priority
                |              |                         . use
                |              |                         .
                |oooooooooooooo|                         v
                 ---------------------------------------
```

         Figure 10: Partial non-priority load & Partial Aggregate load

Figure 11 shows the case where all of the bandwidth available to non-priority traffic is being used and a small amount of priority traffic is in place. In that situation new priority sessions would be accepted but new non-priority sessions would be rejected.
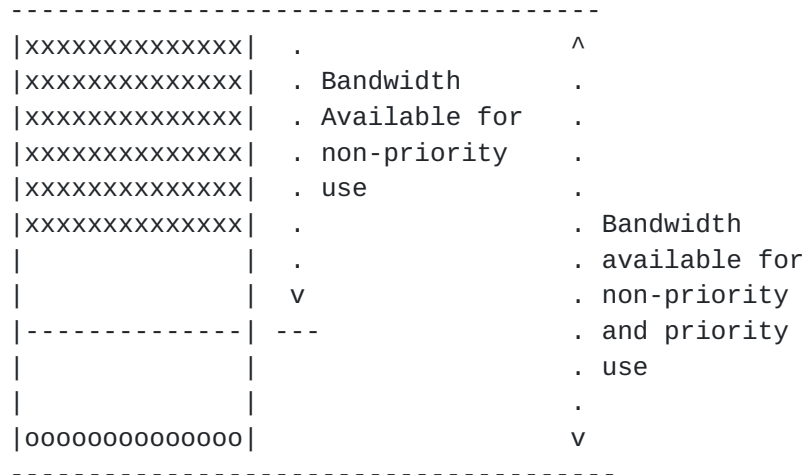
```
             ---------------------------------------
             |xxxxxxxxxxxxxx|   .                   ^
             |xxxxxxxxxxxxxx|   . Bandwidth         .
             |xxxxxxxxxxxxxx|   . Available for     .
             |xxxxxxxxxxxxxx|   . non-priority      .
             |xxxxxxxxxxxxxx|   . use               .
             |xxxxxxxxxxxxxx|   .                   . Bandwidth
             |xxxxxxxxxxxxxx|   .                   . available for
             |xxxxxxxxxxxxxx|   v                   . non-priority
             |--------------| ---                   . and priority
             |              |                       . use
             |              |                       .
             |oooooooooooooo|                       v
             ---------------------------------------
```
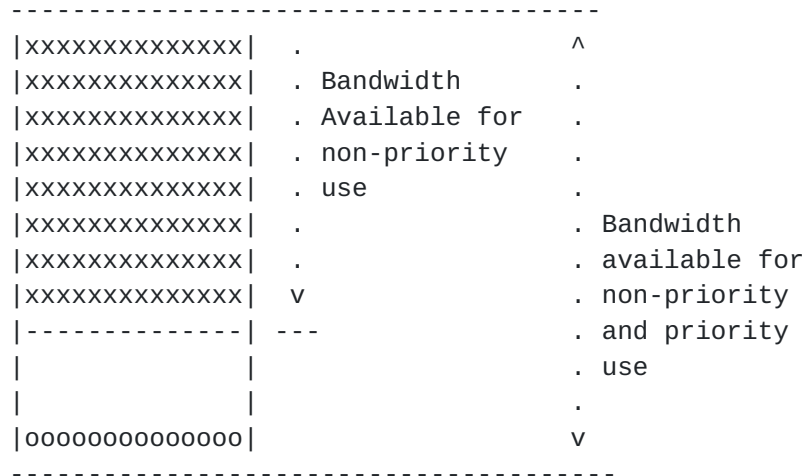
        Figure 11: Full non-priority load & Partial Aggregate load

   Figure 12 shows the case where only some of the bandwidth available
   to non-priority traffic is being used and a heavy load of priority
   traffic is in place. In that situation both new non-priority
   sessions and new priority sessions would be accepted. Note that, as
   illustrated in Figure 11, priority sessions use some of the bandwidth
   currently not used by non-priority traffic.

```
             ---------------------------------------
             |xxxxxxxxxxxxxx|   .                   ^
             |xxxxxxxxxxxxxx|   . Bandwidth         .
             |xxxxxxxxxxxxxx|   . Available for     .
             |xxxxxxxxxxxxxx|   . non-priority      .
             |xxxxxxxxxxxxxx|   . use               .
             |              |   .                   . Bandwidth
             |              |   .                   . available for
             |oooooooooooooo|   v                   . non-priority
             |--------------| ---                   . and priority
             |oooooooooooooo|                       . use
             |oooooooooooooo|                       .
             |oooooooooooooo|                       v
             ---------------------------------------
```
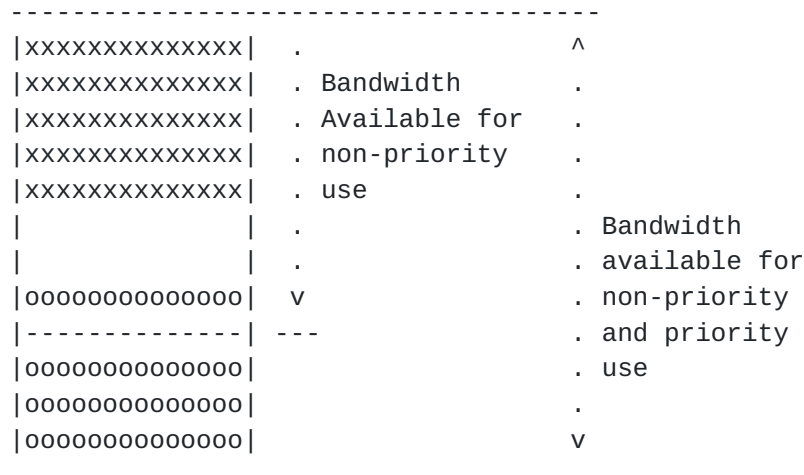
        Figure 12: Partial non-priority load & Heavy Aggregate load

   Figure 13 shows the case where all of the bandwidth available to non-
   priority traffic is being used and all of the remaining available
   bandwidth is used by priority traffic. In that situation new non-
   priority sessions would be rejected. In that situation new priority
   sessions could not be accepted right away. Those priority sessions
   may be handled by mechanisms, which are beyond the scope of this
   particular document (such as established through preemption of

existing non-priority sessions, or such as queuing of new priority
session requests until capacity becomes available again for priority
traffic).

```
               ---------------------------------------
              |xxxxxxxxxxxxxx|   .                  ^
              |xxxxxxxxxxxxxx|   . Bandwidth         .
              |xxxxxxxxxxxxxx|   . Available for     .
              |xxxxxxxxxxxxxx|   . non-priority      .
              |xxxxxxxxxxxxxx|   . use               .
              |xxxxxxxxxxxxxx|   .                   . Bandwidth
              |xxxxxxxxxxxxxx|   .                   . available for
              |xxxxxxxxxxxxxx|   v                   . non-priority
              |--------------| ---                   . and priority
              |oooooooooooooo|                       . use
              |oooooooooooooo|                       .
              |oooooooooooooo|                       v
               ---------------------------------------
```
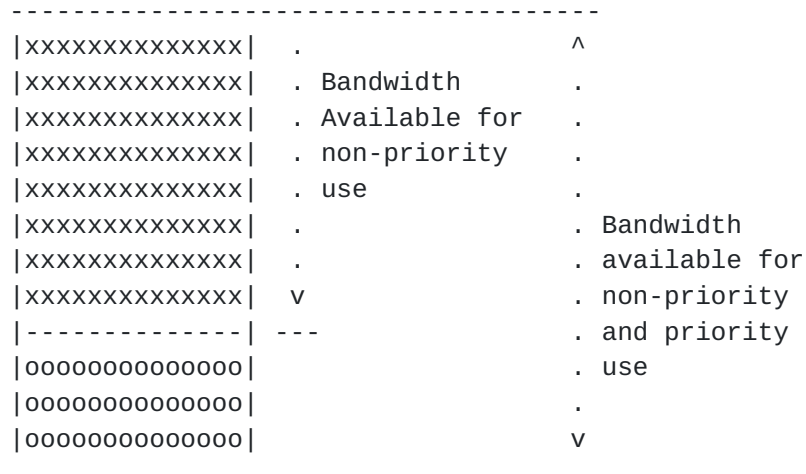
         Figure 13: Full non-priority load & Full Aggregate load

A.3. Admission Priority with Priority Bypass Model (PrBM)

   This section illustrates operations of admission priority when a
   simple Priority Bypass Model (PrBM) is used for bandwidth allocation
   across non-priority traffic and priority traffic. With the Priority
   Bypass Model, non-priority traffic is subject to resource based
   admission control while priority traffic simply bypasses the resource
   based admission control. In other words:

   o when a non-priority session arrives, this session is subject to
      bandwidth admission control and is accepted if the current total
      load (aggregate over non-priority and priority traffic) is below
      the engineered/allocated bandwidth.

   o when a priority session arrives, this session is admitted
      regardless of the current load.

   A property of this model is that a priority session is never
   rejected.

   The rationale for this simple scheme is that, in practice in some
   networks:

   o the volume of priority sessions is very low for the vast majority
      of time, so it may not be economical to completely set aside
      bandwidth for priority sessions and preclude the utilization of
      this bandwidth by normal sessions in normal situations

o even in emergency periods where priority sessions are more heavily
   used, those always still represent a fairly small proportion of
   the overall load which can be absorbed within the safety margin of
   the engineered capacity limits. Thus, even if they are admitted
   beyond the engineered bandwidth threshold, they are unlikely to
   result in noticeable QoS degradation.

As with the MAM and RDM model, an operator may map the Priority
Bypass model onto the Engineered Capacity limits according to
different policies. The operator may decide to configure the
bandwidth limit for admission of non-priority traffic to the full
engineered capacity limits; As an example, if the engineered capacity
limit on a given link is X, the operator may configure the bandwidth
limit for non-priority traffic to X. Alternatively, the operator may
decide to configure the bandwidth limit for non-priority traffic to
below the engineered capacity limits (so that the sum of the non-
priority and priority traffic stays below the engineered capacity);
As an example, if the engineered capacity limit on a given link is X,
the operator may configure the bandwidth limit for non-priority
traffic to 95% of X. Finally, the operator may decide to strike a
balance in between. The considerations presented for these policies
in the previous sections in the MAM and RDM contexts are equally
applicable to the Priority Bypass Model.

Figure 14 illustrates the bandwidth allocation with the Priority
Bypass Model.

```
                    -----------------------
           ^       ^ |                 | ^
           .       . |                 | .
   Total   .       . |                 | .   Bandwidth Limit
          (1)    (2) |                 | .   (on non-priority + priority)
    Engi-  .       . |                 | .   for admission
   neered  . or  . |                 | .   of non-priority traffic
           .       . |                 | .
   Capacity.      . |                 | .
           v       . |                 | v
                   . |--------------| ---
                   . |                 |
                   . |                 |
                   v |                 |
                     |                 |
```
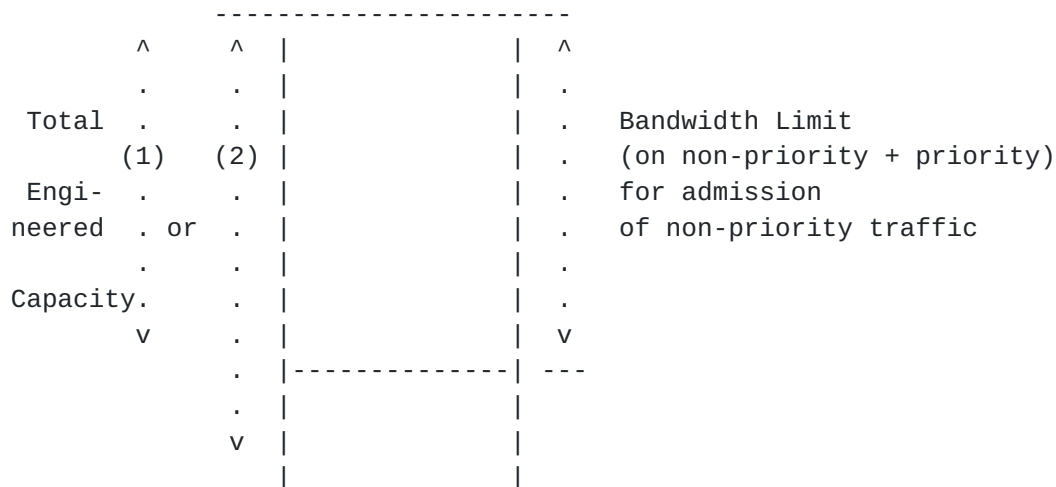
Figure 14: Priority Bypass Model Bandwidth Allocation

Figure 15 shows some of the non-priority capacity of this link being
used. In this situation, both new non-priority and new priority
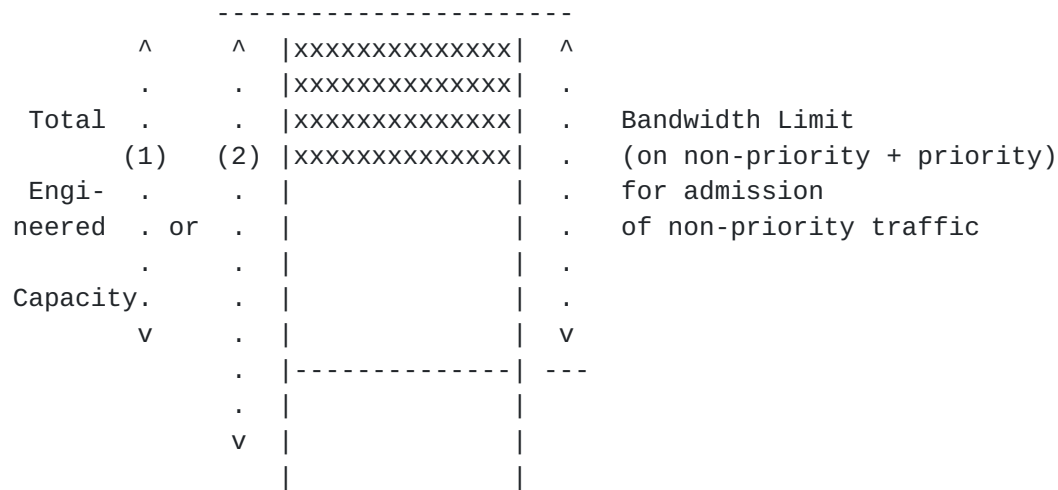sessions would be accepted.

```
                   ----------------------
        ^        ^  |xxxxxxxxxxxxxx|   ^
        .        .  |xxxxxxxxxxxxxx|   .
  Total .        .  |xxxxxxxxxxxxxx|   .   Bandwidth Limit
       (1)  (2)  |xxxxxxxxxxxxxx|   .   (on non-priority + priority)
  Engi-  .     .  |              |   .   for admission
  neered . or .  |              |   .   of non-priority traffic
        .     .  |              |   .
  Capacity.    .  |              |   .
        v     .  |              |   v
              .  |--------------|  ---
              .  |              |
              v  |              |
                 |              |


              Figure 15: Partial load of non-priority calls
```

Figure 16 shows the same amount of non-priority load being used at
this link, and a small amount of priority bandwidth being used. In
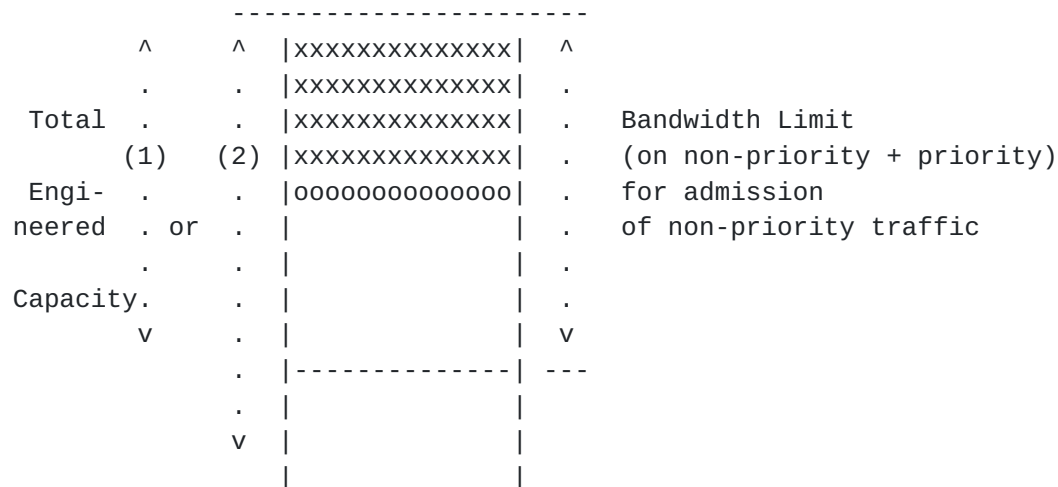this situation, both new non-priority and new priority sessions would
be accepted.

```
                   ----------------------
        ^        ^  |xxxxxxxxxxxxxx|   ^
        .        .  |xxxxxxxxxxxxxx|   .
  Total .        .  |xxxxxxxxxxxxxx|   .   Bandwidth Limit
       (1)  (2)  |xxxxxxxxxxxxxx|   .   (on non-priority + priority)
  Engi-  .     .  |oooooooooooooo|   .   for admission
  neered . or .  |              |   .   of non-priority traffic
        .     .  |              |   .
  Capacity.    .  |              |   .
        v     .  |              |   v
              .  |--------------|  ---
              .  |              |
              v  |              |
                 |              |


       Figure 16: Partial load of non-priority calls & partial load of
                             priority calls
```

Figure 17 shows the case where aggregate non-priority and priority
load exceeds the bandwidth limit for admission of non-priority
traffic. In this situation, any new non-priority session is rejected
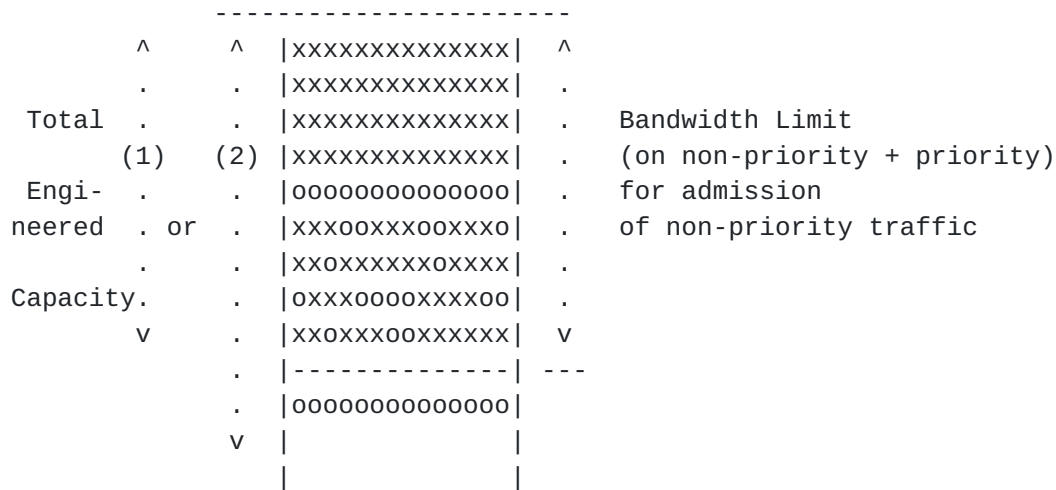while any new priority session is admitted.

```
                  -----------------------
          ^       ^  |xxxxxxxxxxxxxx|   ^
          .       .  |xxxxxxxxxxxxxx|   .
  Total   .       .  |xxxxxxxxxxxxxx|   .   Bandwidth Limit
         (1)     (2) |xxxxxxxxxxxxxx|   .   (on non-priority + priority)
   Engi-  .       .  |oooooooooooooo|   .   for admission
  neered  . or   .   |xxxooxxxooxxxo|   .   of non-priority traffic
          .       .  |xxoxxxxxxoxxxx|   .
  Capacity.       .  |oxxxooooxxxxoo|   .
          v       .  |xxoxxxooxxxxxx|   v
                  .  |--------------|  ---
                  .  |oooooooooooooo|
                  v  |              |
                     |              |
```

                    Figure 17: Full non-priority load

Appendix B. Example Usages of RSVP Extensions

   This section provides examples of how RSVP extensions defined in this
   document can be used (in conjunctions with other RSVP functionality
   and SIP functionality) to enforce different hypothetical policies for
   handling Emergency sessions in a given administrative domain. This
   Appendix does not provide additional specification. It is only
   included in this document for illustration purposes.

   We assume an environment where SIP is used for session control and
   RSVP is used for resource reservation.

   In a mild abuse of language, we refer here to "Call Queueing" as the
   set of "session" layer capabilities that may be implemented by SIP
   user agents to influence their treatment of SIP requests. This may
   include the ability to "queue" session requests when those can not be
   immediately honored (in some cases with the notion of "bumping", or
   "displacement", of less important session requests from that queue).
   It may include additional mechanisms such as exemption from certain
   network management controls, and alternate routing.

   We only mention below the RSVP policy elements that are to be
   enforced by PEPs. It is assumed that these policy elements are set
   at administrative domain boundaries by PDPs. The Admission Priority
   and Preemption Priority RSVP policy elements are set by PDPs as a
   result of processing the Application Level Resource Priority Policy
   Element (which is carried in RSVP messages).

   If one wants to implement an emergency service purely based on Call
   Queueing, one can achieve this by signaling emergency sessions:

o   using "Resource-Priority" Header in SIP

o   not using Admission-Priority Policy Element in RSVP

o   not using Preemption Policy Element in RSVP

If one wants to implement an emergency service based on Call Queueing and on "prioritized access to network layer resources", one can achieve this by signaling emergency sessions:

o   using "Resource-Priority" Header in SIP

o   using Admission-Priority Policy Element in RSVP

o   not using Preemption Policy Element in RSVP

Emergency sessions will not result in preemption of any session. Different bandwidth allocation models can be used to offer different "prioritized access to network resources". Just as examples, this includes strict setting aside of capacity for emergency sessions as well as simple bypass of admission limits for emergency sessions.

If one wants to implement an emergency service based on Call Queueing, on "prioritized access to network layer resources", and ensures that (say) "Emergency-1" sessions can preempt "Emergency-2" sessions, but non-emergency sessions are not affected by preemption, one can do that by signaling emergency sessions:

o   using "Resource-Priority" Header in SIP

o   using Admission-Priority Policy Element in RSVP

o   using Preemption Policy Element in RSVP with:

    *   setup (Emergency-1) > defending (Emergency-2)

    *   setup (Emergency-2) <= defending (Emergency-1)

    *   setup (Emergency-1) <= defending (Non-Emergency)

    *   setup (Emergency-2) <= defending (Non-Emergency)

If one wants to implement an emergency service based on Call Queueing, on "prioritized access to network layer resources", and ensure that "emergency" sessions can preempt regular sessions, one could do that by signaling emergency sessions:

o using "Resource-Priority" Header in SIP

o using Admission-Priority Policy Element in RSVP

o using Preemption Policy Element in RSVP with:

   * setup (Emergency) > defending (Non-Emergency)

   * setup (Non-Emergency) <= defending (Emergency)

If one wants to implement an emergency service based on Call
Queueing, on "prioritized access to network layer resources", and
ensure that "emergency" sessions can partially preempt regular
sessions (i.e. reduce their reservation size), one could do that by
signaling emergency sessions:

o using "Resource-Priority" Header in SIP

o using Admission-Priority Policy Element in RSVP

o using Preemption in Policy Element RSVP with:

   * setup (Emergency) > defending (Non-Emergency)

   * setup (Non-Emergency) <= defending (Emergency)

o activate RFC4495 RSVP Bandwidth Reduction mechanisms

Authors' Addresses

   Francois Le Faucheur
   Cisco Systems
   Greenside, 400 Avenue de Roumanille
   Sophia Antipolis 06410
   France

   Phone: +33 4 97 23 26 19
   Email: flefauch@cisco.com

James Polk
Cisco Systems
2200 East President George Bush Highway
Richardson, TX 75082-3550
United States

Phone: +1 972 813 5208
Email: jmpolk@cisco.com

Ken Carlberg
G11
123a Versailles Circle
Towson, MD 21204
United States

Email: carlberg@g11.org.uk