Internet Draft

Francois Le Faucheur
James Polk
Cisco Systems, Inc.

Ken Carlberg
G11

draft-ietf-tsvwg-emergency-rsvp-01.txt
Expires: July 2007

January 2007

Resource ReSerVation Protovol (RSVP) Extensions for Emergency Services

Status of this Memo

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 RSVP extensions 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.

Copyright Notice

Specification of Requirements

In this document, 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 [KEYWORDS] and indicate requirement levels for compliant implementations.

Table of Contents

1. Introduction

[EMERG-RQTS] and [EMERG-TEL] detail requirements for an Emergency Telecommunications Service (ETS), which is an umbrella term identifying those networks and specific services used to support emergency communications. An underlying goal of these documents 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 [SIP], 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" call requests when those can not be immediately honored (in some cases with the notion of "bumping", or "displacement", of less important call request 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 calls.

Note: Below, this document references several examples of IP telephony and its use of "calls", which is one form of the term "sessions" (Video over IP and Instant Messaging being other examples that rely on session establishment). For the sake of simplicity, we shall use the widely known term "call" for the remainder of this document.

1.1. Related Technical Documents

[EMERG-IMP] 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 call set up, as well as operations after call establishment through maintenance of a continuing call model of the status of all calls. [EMERG-IMP] also describes how these network layer admission control procedures can be realized using the Resource reSerVation Protocol [RSVP] along with its associated protocol suite and extensions, including those for policy based admission control ([FW-POLICY], [RSVP-POLICY]), for user authentication and authorization ([RSVP-ID]) and for integrity and authentication of RSVP messages ([RSVP-CRYPTO-1], [RSVP-CRYPTO-2]).

Furthermore, [EMERG-IMP] describes how the RSVP Signaled Preemption Priority Policy Element specified in [RSVP-PREEMP] can be used to enforce the call preemption that may be needed by some emergency services.

In contrast to [EMERG-IMP], this document specifies new RSVP extensions to increase the probability of call completion 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 call admission priority. Appendix A of this document also provides three examples of a bandwidth allocation model, which can be used by RSVP-routers to enforce such admission priority on every link.

1.2. Terminology

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

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

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

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

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

1.3. Changes from previous versions

[Note to RFC Editor: This section is to be removed before publication]

Changes from ietf-tsvwg-emergency-rsvp-00 to ietf-tsvwg-emergency-rsvp-01

The most significant changes are:

- editorial change (correction in description of "Take highest priority" in section 3.1.1).
- expanded Security Considerations section

Changes from lefaucheur-rsvp-emergency-01 to ietf-tsvwg-rsvp-emergency-00

The most significant change is:

- Extended the Admission Priority field from 3 to 8 bits and inverted the encoding order, in particular for better alignment with NSIS Qspec.

Changes from lefaucheur-rsvp-emergency-01 to lefaucheur-rsvp-emergency-02

The most significant changes are:

- modified the Introduction to add additional clarity and to place related work in a better context to the extensions proposed in this draft
- Moved bandwidth allocation models to an appendix
- Allowed multiple Application-Level Resource Priority inside ALRP Policy Element
- Added a 2nd appendix providing examples of RSVP extensions usage

Changes from lefaucheur-rsvp-emergency-00 to lefaucheur-rsvp-emergency-01

The most significant changes were:

- o adding a second RSVP Policy Element that contains the application-level resource priority requirements (for example as communicated in the SIP Resource-Priority Header) for scenarios where priority calls transits through multiple administrative domains.

- o adding description of a third bandwidth allocation model example: the Priority Bypass Model

- o adding discussion on policies for mapping the various bandwidth allocation model over the engineered capacity limits.

2. Overview of RSVP extensions and Operations

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

As described in [EMERG-IMP], 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 [SIP-RESOURCE].

Devices involved in the session establishment are expected to be aware of the application-level priority requirements of emergency calls. 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 call using the Resource-Priority Header mechanism specified in [SIP-PRIORITY]. 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 [FW-POLICY]) 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 (ie 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 calls 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 calls 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 [SIP-PRIORITY], for support of scenarios where calls 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 [RSVP-PREEMP] 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 [RSVP-PREEMP] and thus, can be taken into account at every RSVP-enabled network hop as discussed [EMERG-IMP]. Appendix B provides examples of various hypothetical policies for emergency call 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.

2.1. 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 [DSTE-MAM] and the Russian Dolls Model (RDM) specified in [DSTE-RDM]. 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 (PBM). 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.

3. New Policy Elements

The Framework document for policy-based admission control [FW-POLICY] 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:

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

- 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.

[RSVP-POLICY] 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, [RSVP-PREEMP] specifies the Preemption Priority Policy Element.

This document defines two new Policy Elements called:

- the Admission Priority Policy Element
- the Application-Level Resource Priority Policy Element

3.1. Admission Priority Policy Element

The format of the Admission Priority policy element is as follows:

```
      +-------------+-------------+-------------+-------------+
      |     Length                | P-Type = ADMISSION_PRI    |
      +-------------+-------------+-------------+-------------+
      | Flags       | M. Strategy | Error Code  | Reserved    |
      +-------------+-------------+-------------+-------------+
      |Adm. Priority|            Reserved                     |
      +---------------------------+---------------------------+
```

Length: 16 bits
Always 12. The overall length of the policy element, in bytes.

P-Type: 16 bits
ADMISSION_PRI = To be allocated by IANA
(see "IANA Considerations" section)

Flags: Reserved (MUST be set to zero on transmit and ignored on receive)

Merge Strategy: 8 bit (only applicable to multicast flows)
1   Take priority of highest QoS
2   Take highest priority

3	Force Error on heterogeneous merge

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

0	NO_ERROR	Value used for regular ADMISSION_PRI elements
2	HETEROGENEOUS	This element encountered heterogeneous merge

Reserved: 8 bits
Always 0.

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 call completion to selected flows. Higher values represent higher 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 call completion. 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.

Reserved: 16 bits
Always 0.

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 orthogonal and independent mechanisms.

3.1.1. 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 the same as those defined in [RSVP-PREEMP] for merging Preemption Priority Policy Elements. In particular, the following merging strategies are supported:

- Take priority of highest QoS
- Take highest priority

- Force Error on heterogeneous merge.

The only difference with [RSVP-PREEMP] is that this document does not recommend any merge strategies for Admission Priority while [RSVP-PREEMP] 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".

3.2. Application-Level Resource Priority Policy Element

The format of the Application-Level Resource Priority policy element is as follows:

```
      +-------------+-------------+-------------+-------------+
      | Length                    | P-Type = APP_RESOURCE_PRI |
      +-------------+-------------+-------------+-------------+
      //     ALRP List                                       //
      +---------------------------+---------------------------+
```

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.

P-Type: 16 bits
    APP_RESOURCE_PRI = To be allocated by IANA
   (see "IANA Considerations" section)

ARLP:

```
      +---------------------------+---------------------------+
      |     ALRP Namespace        |ALRP Priority| Reserved    |
      +---------------------------+---------------------------+
```

ALRP Namespace (Application-Level Resource Priority Namespace):
16 bits (unsigned)
Contains the namespace of the application-level resource priority. This is encoded as a numerical value which represents the position of the namespace in the "Resource-Priority Namespace" IANA registry, starting with 0. Creation of this registry has been requested to IANA in [SIP-PRIORITY].
For example, as "drsn", "dsn", "q735", "ets" and "wps" are currently the first, second, third, fourth and fifth namespaces defined in the "Resource-Priority Namespace" registry, those are respectively encoded as value 0, 1, 2, 3 and 4.

ALRP Priority: (Application-Level Resource Priority Priority):
8 bits (unsigned)
Contains the priority value within the namespace of the application-level resource priority. This is encoded as a numerical value which represents the priority defined in the "Resource-Priority Namespace" IANA registry for the considered namespace, starting from 0 for the highest priority and increasing as priority decreases.
For example, as "flash-override", "flash", "immediate", "priority" and "routine" are the priorities in decreasing order of priority registered for the "dsn" namespace, those are respectively encoded as value 0, 1, 2, 3 and 4. As another example, as "flash-override-override", "flash-override", "flash", "immediate", "priority" and "routine" are the priorities in decreasing order of priority registered for the "drsn" namespace, those are respectively encoded as value 0, 1, 2, 3, 4 and 5.

Reserved: 16 bits
Always 0.

3.2.1. 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:

> 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_PR element. A given ALRP MUST NOT appear more than once. In other words, the outgoing ALRP List is the reunion of the incoming ARLP Lists that are merged.

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

4. Security Considerations

The ADMISSION_PRI and APP_RESOURCE_PRI are Policy Elements that can be signaled by RSVP through encapsulation in a Policy Data object as defined in [RSVP-POLICY]. Therefore, like any other Policy Elements, their integrity can be protected as discussed in section 6 of [RSVP-

POLICY] by two optional security mechanisms. The first mechanism relies on RSVP Authentication as specified in [RSVP-CRYPTO-1] and [RSVP-CRYPTO-2] to provide a chain of trust when all RSVP nodes are policy capable. 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.

4.1. Use of RSVP Authentication

[RSVP-CRYPTO-1] discusses several approaches for distribution of keys to be used for RSVP Authentication. First, the RSVP Authentication shared keys can be distributed manually. This is the base option and its support is mandated for any implementation. However, in some environments, this approach may become a burden if keys frequently change over time. Alternatively, a standard key management protocol for secure key distribution can be used. However, existing key distribution protocols may not be appropriate in all environments because of the complexity or operational burden they involve. Finally, [RSVP-CRYPTO-1] specifies how Kerberos [KERBEROS] may be used to generate the RSVP Authentication keys. Kerberos allows for the use of trusted third party keying relationships between security principals (RSVP sender and receivers) where the Kerberos key distribution center (KDC) establishes an ephemeral session key to be shared between RSVP sender and receivers.

The use of RSVP Authentication in parts of the network where there may be one or more IP hops in between two RSVP neighbors raises an additional challenge. This is because, with some RSVP messages such as a Path message, an RSVP router does not know the RSVP next hop for that message at the time of forwarding it. In fact, part of the role of a Path message is precisely to discover the RSVP next hop (and to dynamically re-discover it when it changes, say because of a routing change). Hence, the RSVP router may not know which security association to use when forwarding such a message.
In that situation, one approach is to share the same RSVP Authentication shared key across all the RSVP routers of a part of the network where there may be RSVP neighbors with IP hops in between. For example, all the RSVP routers of an administrative domain could share the same RSVP Authentication key, while different per-neighbor keys could be used between any RSVP router pair straddling the boundary between two administrative domains that have agreed to use RSVP signaling.

When the same RSVP Authentication shared key is to be shared among multiple RSVP neighbors, manual key distribution may be used. For situations where RSVP is being used for multicast flows, it might also be possible, in the future, to adapt a multicast key management method (e.g. from IETF Multicast Security Working Group) for key distribution with such multicast RSVP usage. For situations where

RSVP is being used for unicast flows within a single administrative domain, the Kerberos technique described in Section 7 of [RSVP-CRYPTO-1] might be considered. For situations where RSVP is being used for unicast flows across domain boundaries, it is not currently clear how one might provide automated key management. Specification of a specific automated key management technique is outside the scope of this document. Operators should consider these key management issues when contemplating deployment of this specification.

4.2. 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 [RSVP-POLICY]. 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. This issue is similar to the one discussed in section 4.1, except it now applies to PDP neighbors instead of RSVP neighbors. Hence similar approaches could be used, such as the use of a key shared across multiple PDPs. We observe 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).

5. IANA Considerations

As specified in [RSVP-POLICY], Standard RSVP Policy Elements (P-type values) are to be assigned by IANA as per "IETF Consensus" following the policies outlined in [IANA-CONSIDERATIONS].

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

- one P-Type to the Admission Priority Policy Element

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

6. 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 [RSVP-PREEMP]. Dave Oran and Janet Gunn provided useful input into this document.

7. Normative References

[DSTE-MAM] Le Faucheur & Lai, "Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering", RFC 4125, June 2005.

[DSTE-RDM] Le Faucheur et al, Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering, RFC 4127, June 2005

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

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

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

[IANA-CONSIDERATIONS] Alverstrand et al., "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[KEYWORDS] "Key words for use in RFCs to Indicate Requirement Levels", Bradner, RFC2119, BCP14

[KERBEROS] Neuman et al., "The Kerberos Network Authentication Service (V5)", RFC 4120, July 2005.

[RSVP] Braden, R., ed., et al., "Resource ReSerVation Protocol (RSVP)- Functional Specification", RFC 2205, September 1997.

[RSVP-CRYPTO-1] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.

[RSVP-CRYPTO-2] Braden, R. and L. Zhang, "RSVP Cryptographic Authentication -- Updated Message Type Value", RFC 3097, April 2001.

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

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

[SIP] Rosenberg et al., "SIP: Session Initiation Protocol", RFC3261, June 2002

[SIP-PRIORITY] H. Schulzrinne & J. Polk. "Communications Resource Priority for the Session Initiation Protocol (SIP)", RFC4412, February 2006.

8. Informative References

[EMERG-IMP] F. Baker & J. Polk, "Implementing an Emergency Telecommunications Service for Real Time Services in the Internet Protocol Suite", RFC 4542, May 2006.

[ITU.I.225] ITU, "Multi-Level Precedence and Preemption Service, ITU, Recommendation, I.255.3, July, 1990.

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

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

Appendix A: Examples of Bandwidth Allocation Model for Admission Priority

Sections A.1 and A.2 respectively illustrate how the Maximum Allocation Model [DSTE-MAM] and the Russian Dolls Model (RDM) [DSTE-RDM] can be used 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 charts 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
             -------------------------
```

Chart 1. MAM Bandwidth Allocation

Chart 1 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 calls, no additional priority sessions can be accepted.

As illustrated in Chart 1, 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 calls. This policy illustrated as (1) in Chart 1. 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 calls from using the full engineered capacity, even when there is no or little priority load, which is the majority of time. This policy illustrated as (3) in Chart 1. 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 illustrated as (2) in Chart 1.

Chart 2 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
```

```
             -------------------------
```

Chart 2. Partial load of non-priority calls

Chart 3 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
             -------------------------
```

Chart 3. Partial load of non-priority calls
& partial load of priority calls

Chart 4 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
             -------------------------
```

Chart 4. Full non-priority load

& partial load of priority calls

Chart 5 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 calls 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
             -------------------------
```

Chart 5. 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.

Chart 6 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
            ---------------------------------------
```

Chart 6. Partial non-priority load & Partial Aggregate load

Chart 7 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
            ---------------------------------------
```

Chart 7. Full non-priority load & Partial Aggregate load

Chart 8 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 Chart 7, priority calls 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
            ---------------------------------------
```

Chart 8. Partial non-priority load & Heavy Aggregate load

Chart 9 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
---------------------------------------
```

Chart 9. Full non-priority load & Full Aggregate load

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

This section illustrates operations of admission priority when a simple Priority Bypass Model (PBM) 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:

- when a non-priority call arrives, this call 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.
- when a priority call arrives, this call is admitted regardless of the current load.

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

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

- the volume of priority calls is very low for the vast majority of time, so it may not be economical to completely set aside bandwidth for priority calls and preclude the utilization of this bandwidth by normal calls in normal situations
- even in emergency periods where priority calls 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.

Chart 10 shows 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  |              |
                 |              |
```

Chart 10. Priority Bypass Model Bandwidth Allocation

Chart 11 shows some of the non-priority capacity of this link being used. In this situation, both new non-priority and new priority calls 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  |              |
                 |              |
```

Chart 11. Partial load of non-priority calls

Chart 12 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 calls 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  |              |
                 |              |
```

Chart 12. Partial load of non-priority calls
& partial load of priority calls

Chart 13 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 call is rejected while any new priority call 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  |              |
                 |              |
```

Chart 13. 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. The content of this appendix may be moved into a future applicability statement document.

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" call requests when those can not be immediately honored (in some cases with the notion of "bumping", or "displacement", of less important call request 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 calls:

* using "Resource-Priority" Header in SIP
* not using Admission-Priority Policy Element in RSVP
* 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 calls:

* using "Resource-Priority" Header in SIP
* using Admission-Priority Policy Element in RSVP
* not using Preemption Policy Element in RSVP

Emergency calls 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 calls:

- using "Resource-Priority" Header in SIP
- using Admission-Priority Policy Element in RSVP
- 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 calls:

- using "Resource-Priority" Header in SIP
- using Admission-Priority Policy Element in RSVP
- 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 (ie reduce their reservation size), one could do that by signaling emergency calls:

- using "Resource-Priority" Header in SIP
- using Admission-Priority Policy Element in RSVP
- using Preemption in Policy Element RSVP with:
    - setup (Emergency) > defending (Non-Emergency)
    - setup (Non-Emergency) <= defending (Emergency)
- activate RFC4495 RSVP Bandwidth Reduction mechanisms

Authors' Address

Francois Le Faucheur
Cisco Systems, Inc.
Village d'Entreprise Green Side - Batiment T3
400, Avenue de Roumanille

06410 Biot Sophia-Antipolis
France
Email: flefauch@cisco.com

James Polk
Cisco Systems, Inc.
2200 East President George Bush Turnpike
Richardson, Texas 75082
USA
Email: jmpolk@cisco.com

Ken Carlberg
G11
123a Versailles Circle
Towson, MD. 21204
USA
email: carlberg@g11.org.uk

Intellectual Property

Full Copyright Statement

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.