

RSVP Extensions for Admission Priority

Abstract

Some applications require the ability to provide an elevated probability of session establishment to specific sessions in times of network congestion. When supported over the Internet Protocol suite, this may be facilitated through a network-layer admission control solution that supports prioritized access to resources (e.g., bandwidth). These resources may be explicitly set aside for prioritized sessions, or 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.

Based on current security concerns, these extensions are intended for use in a single administrative domain.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in [Section 2 of RFC 5741](#).

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6401>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
1.1.	Terminology	4
2.	Applicability Statement	4
3.	Requirements Language	4
4.	Overview of RSVP Extensions and Operations	4
4.1.	Operations of Admission Priority	6
5.	New Policy Elements	7
5.1.	Admission Priority Policy Element	8
5.1.1.	Admission Priority Merging Rules	9
5.2.	Application-Level Resource Priority Policy Element	10
5.2.1.	Application-Level Resource Priority Modifying and Merging Rules	11
5.3.	Default Handling	12
6.	Security Considerations	12
6.1.	Use of RSVP Authentication between RSVP Neighbors	13
6.2.	Use of INTEGRITY object within the POLICY_DATA Object	13
7.	IANA Considerations	14
8.	Acknowledgments	16
9.	References	17
9.1.	Normative References	17
9.2.	Informative References	18
Appendix A.	Examples of Bandwidth Allocation Model for Admission Priority	19
A.1.	Admission Priority with Maximum Allocation Model (MAM)	19
A.2.	Admission Priority with Russian Dolls Model (RDM)	23
A.3.	Admission Priority with Priority Bypass Model (PrBM)	26
Appendix B.	Example Usages of RSVP Extensions	29

1. Introduction

Some applications require the ability to provide an elevated probability of session establishment to specific sessions in times of network congestion.

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 alternate routing and exemption from certain network management controls.

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 (QoS) 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 prioritized sessions only. Or alternatively, it may involve allowing the prioritized 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](#)].

[RFC3181] specifies the Signaled Preemption Priority Policy Element that can be signaled in RSVP so that network node may take into account this policy element in order to preempt some previously admitted low-priority sessions in order to make room for a newer, higher-priority session. In contrast, this document specifies new RSVP extensions to increase the probability of session establishment without preemption of existing sessions. This is achieved by engineered capacity techniques in the form of bandwidth allocation models. 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 a selective bandwidth admission control decision based on the session admission

priority. [Appendix A](#) of this document also provides examples of bandwidth allocation models that can be used by RSVP-routers to enforce such admission priority on every link. A given reservation may be signaled with the admission priority extensions specified in the present document, with the preemption priority specified in [\[RFC3181\]](#), or with both.

1.1. Terminology

This document assumes the terminology defined in [\[RFC2753\]](#). For convenience, the definitions of a few key terms are repeated here:

- o Policy Decision Point (PDP): The point where policy decisions are made.
- o Local Policy Decision Point (LPDP): The 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. Applicability Statement

A subset of RSVP messages are signaled with the Router Alert Option (RAO) ([\[RFC2113\]](#), [\[RFC2711\]](#)). The security aspects and common practices around the use of the current IP Router Alert Option and consequences on the use of IP Router Alert by applications such as RSVP are discussed in [\[RFC6398\]](#). Based on those, the extensions defined in this document are intended for use within a single administrative domain. Thus, in particular, the extensions defined in this document are not intended for end-to-end use on the Internet.

3. 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](#) [\[RFC2119\]](#).

4. Overview of RSVP Extensions and Operations

Let us consider the case where a session requires elevated probability of establishment, 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 the priority session to seize network-layer resources from the engineered capacity that has been set aside for priority sessions (and not made available to normal sessions) or, alternatively, allowing the priority session to seize additional resources beyond the engineered capacity limits applied to normal sessions.

Session establishment can be made conditional on 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 in conjunction with SIP is specified in [\[RFC3312\]](#).

Devices involved in the session establishment are expected to be aware of the application-level priority requirements of prioritized sessions. For example, considering the case where the session control protocol is SIP, the SIP user agents may be made aware of the resource priority requirements of a given 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 the 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 a policy boundary based on the policy applicable in that policy area. In a typical model (see [\[RFC2753\]](#)) where PDPs control PEPs at the periphery of the policy area (e.g., on the first hop router), PDPs would interpret the RSVP Application-Level Resource Priority Policy Element and map the requirement of the prioritized 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 that 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 (possibly many) hops that do not have any understanding of application-level resource priority semantics. 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 policy area. It is expected to be used at policy area boundaries only in order to set/reset the RSVP Admission Priority Policy Element.

Remapping by PDPs of the Admission Priority Policy Element from the Application-Level Resource Priority Policy Element may also be used at boundaries with other signaling protocols, such as the NSIS Signaling Layer Protocol (NSLP) for QoS Signaling [RFC5974].

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 deemed necessary by the prioritized session handling policy). In that case, when processing the RSVP Application-Level Resource Priority Policy Element, the PDPs at policy boundaries (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 that can then be encoded inside the Preemption Priority Policy Element of [RFC3181].

[Appendix B](#) provides examples of various hypothetical policies for prioritized 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 needs to be mapped into RSVP policy elements by the PDPs to realize these policies.

[4.1.](#) Operations of Admission Priority

The RSVP Admission Priority Policy Element defined in this document allows admission bandwidth to be allocated preferentially to prioritized sessions. 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]. However, these same models MAY 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 also 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 how the RSVP Admission Priority can be used by RSVP routers to influence their admission control decision (for example, by determining which bandwidth pool is to be used by RSVP for performing its bandwidth allocation) and therefore to increase the probability of reservation establishment. In turn, this increases the probability of application-level session establishment for the corresponding session.

5. 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 4](#) 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 lightweight 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 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

5.1. Admission Priority Policy Element

The format of the Admission Priority Policy Element is as shown in Figure 1:

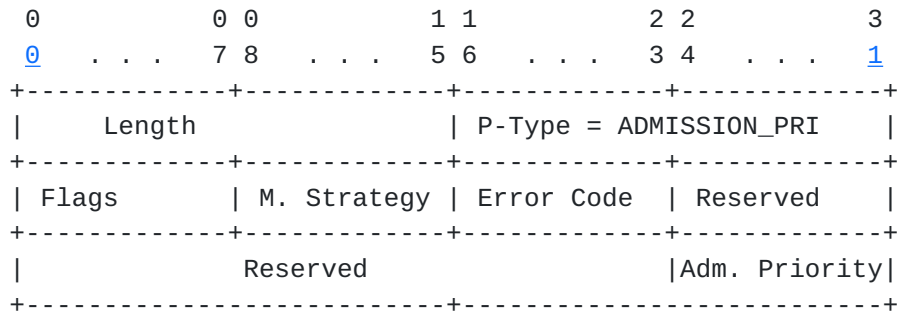


Figure 1: Admission Priority Policy Element

where:

- o Length: 16 bits
 - * Always 12. The overall length of the policy element, in bytes.
- o P-Type: 16 bits
 - * ADMISSI_PRI = 0x05 (see the "IANA Considerations" section).
- o Flags: Reserved
 - * SHALL be set to zero on transmit and SHALL be ignored on reception.
- o Merge Strategy: 8 bits (applicable to multicast flows)
 - * values are defined in the corresponding registry maintained by IANA (see the "IANA Considerations" section).
- o Error code: 8 bits (applicable to multicast flows)
 - * values are defined in the corresponding registry maintained by IANA (see the "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. Bandwidth allocation models such as those described in [Appendix A](#) are to be used by the RSVP router to achieve 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.

5.1.1. Admission Priority Merging Rules

This section discusses alternatives for dealing with RSVP admission priority in case of merging of reservations. As merging applies to multicast, this section also 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. This registry applies both to the Merge Strategy field of the Admission Priority Policy Element defined in the present document and to the Merge Strategy field of the Preemption Priority Policy Element defined in [\[RFC3181\]](#). The registry initially contains the values already defined in [\[RFC3181\]](#) (see the "IANA Considerations" section).

The only difference from [\[RFC3181\]](#) is that this document does not recommend a given merge strategy over the others 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".

5.2. Application-Level Resource Priority Policy Element

The format of the Application-Level Resource Priority Policy Element is as shown in Figure 2:




Figure 2: 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 = 0x06 (see the "IANA Considerations" section).

o ALRP List:

- * List of ALRPs where each ALRP is encoded as shown in Figure 3.

ALRP:




Figure 3: 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 the "IANA Considerations" section; IANA has extended 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 Value (Application-Level Resource Priority Value): 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 Value" IANA registry. (See the "IANA Considerations" section; IANA has extended the registry to include this numerical value).

5.2.1. 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 without modification 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 elements, LPDPs MUST do so according to the following rule:

- o The ALRP List in the outgoing APP_RESOURCE_PRI element MUST contain 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 applies to multicast, this rule also applies to multicast sessions.

5.3. 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 PINs 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 objects can traverse PINs.

[Section 4.2 of \[RFC2750\]](#) also defines a similar default behavior for policy-capable nodes that do not recognize a particular policy element. This applies to the Admission Priority Policy Element and the ALRP Policy Element specified in this document, so that those elements can traverse policy-capable nodes that do not support these extensions defined in the present document.

6. Security Considerations

As this document defines extensions to RSVP, the security considerations of RSVP apply. Those are discussed in [\[RFC2205\]](#), [\[RFC4230\]](#), and [\[RFC6411\]](#). Approaches for addressing those concerns are discussed further below.

A subset of RSVP messages are signaled with the Router Alert Option (RAO) ([\[RFC2113\]](#), [\[RFC2711\]](#)). The security aspects and common practices around the use of the current IP Router Alert Option and consequences on the use of IP Router Alert by applications such as RSVP are discussed in [\[RFC6398\]](#). As discussed in [Section 2](#), the extensions defined in this document are intended for use within a single administrative domain.

[\[RFC6398\]](#) discusses router alert protection approaches for service providers. These approaches can be used to protect a given network against the potential risks associated with the leaking of router alert packets resulting from the use of the present extensions in another domain. Also, where RSVP is not used, by simply not enabling RSVP on the routers of a given network, generally that network can isolate itself from any RSVP signaling that may leak from another network that uses the present extensions (since the routers will then typically ignore RSVP messages). Where RSVP is to be used internally within a given network, the network operator can activate, on the edge of his network, mechanisms that either tunnel or, as a last resort, drop incoming RSVP messages in order to protect the given network from RSVP signaling that may leak from another network that uses the present extensions.

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 PEPs that are not RSVP neighbors.

6.1. Use of RSVP Authentication between RSVP Neighbors

RSVP authentication can be used between RSVP neighbors that are policy capable. RSVP authentication (defined in [RFC2747] and [RFC3097]) SHOULD be supported by an implementation of the present document.

With RSVP authentication, the RSVP neighbors use shared keys to compute the cryptographic signature of the RSVP message. [RFC6411] discusses key types and key provisioning methods as well as their respective applicabilities.

6.2. 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. This is useful only when some RSVP nodes are Policy Ignorant Nodes (PINs). The `INTEGRITY` object within the `POLICY_DATA` object MAY be supported by an implementation of the present document.

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 the 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 a 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 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 a key shared across multiple RSVP neighbors as discussed

in [RFC6411]. 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).

7. IANA Considerations

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

IANA has allocated two P-Types from the standard RSVP policy element range:

- o 0x05 ADMISSION_PRI for the Admission Priority Policy Element
- o 0x06 APP_RESOURCE_PRI for the Application-Level Resource Priority Policy Element

In [Section 5.1](#), the present document defines a Merge Strategy field inside the Admission Priority Policy Element. This registry is to be specified as also applicable to the Merge Strategy field of the Preemption Priority Policy Elements defined in [RFC3181]. Since it is conceivable that, in the future, values will be added to the registry that only apply to the Admission Priority Policy Element or to the Preemption Priority Policy Element (but not to both), IANA has listed the applicable documents for each value. IANA has allocated the following values:

- o 0: Reserved
- o 1: Take priority of highest QoS [RFC3181] [RFC6401]
- o 2: Take highest priority [RFC3181] [RFC6401]
- o 3: Force Error on heterogeneous merge [RFC3181] [RFC6401]

Following the policies outlined in [RFC5226], numbers in the range 0-127 are allocated according to the "IETF Review" policy, numbers in the range 128-240 as "First Come First Served", and numbers in the range 241-255 are "Reserved for Private Use".

In [Section 5.1](#), the present document defines an Error Code field inside the Admission Priority Policy Element. IANA has created 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 0-127 are allocated according to the "IETF Review" policy, numbers in the range 128-240 as "First Come First Served", and numbers in the range 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 5.2](#) 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 namespaces. That registry has been updated to allocate a numerical value to each namespace. To be exact, the IANA has extended the Resource-Priority Namespaces registry in the following ways:

- o A new column has been added to the registry.
- o The title of the new column is "Namespace Numerical Value *".
- o In the Legend, a line has been added stating "Namespace Numerical Value = the unique numerical value identifying the namespace".
- o In the Legend, a line has been added stating "* : [\[RFC6401\]](#)".
- o An actual numerical value has been allocated to each namespace in the registry and is listed in the new "Namespace Numerical Value *" column.

A numerical value has been allocated by IANA for all existing namespaces. 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 5.2](#) 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. That registry has been updated to allocate a numerical value to each priority-value. To be exact, the IANA has extended the Resource-Priority Priority-Values registry in the following ways:

- o For each namespace, the registry is structured with two columns.
- o The title of the first column is "Priority Values (least to greatest)".
- o The first column lists all the values currently defined in the registry (e.g., for the drsn namespace: "routine", "priority", "immediate", "flash", "flash-override", and "flash-override-override")
- o The title of the second column is "Priority Numerical Value *".
- o At the bottom of the registry, a "Legend" has been added with a line stating "Priority Numerical Value = the unique numerical value identifying the priority within a namespace".
- o In the Legend, a line has been added stating "* : [[RFC6401](#)]".
- o An actual numerical value has been allocated to each priority value and is listed in the new "Priority Numerical Value *" column.

A numerical value has been allocated by IANA to all existing priorities. 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. 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" is allocated numerical value 5, and "flash-override-override" is allocated numerical value 0.

8. 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 the Preemption Priority Policy Element [[RFC3181](#)]. Dave Oran and Janet Gunn provided useful input for this document. Ron Bonica, Magnus Westerlund, Cullen Jennings, Ross Callon and Tim Polk provided specific guidance for the applicability statement of the mechanisms defined in this document.

9. References

9.1. Normative References

- [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.
- [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.
- [RFC3312] Camarillo, G., Marshall, W., and J. Rosenberg, "Integration of Resource Management and Session Initiation Protocol (SIP)", [RFC 3312](#), October 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.
- [RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", [BCP 168](#), [RFC 6398](#), October 2011.

[9.2.](#) Informative References

- [RFC2113] Katz, D., "IP Router Alert Option", [RFC 2113](#), February 1997.
- [RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", [RFC 2711](#), October 1999.
- [RFC2753] Yavatkar, R., Pendarakis, D., and R. Guerin, "A Framework for Policy-based Admission Control", [RFC 2753](#), January 2000.
- [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.
- [RFC4230] Tschofenig, H. and R. Graveman, "RSVP Security Properties", [RFC 4230](#), December 2005.
- [RFC4495] Polk, J. and S. Dhesikan, "A Resource Reservation Protocol (RSVP) Extension for the Reduction of Bandwidth of a Reservation Flow", [RFC 4495](#), May 2006.
- [RFC5974] Manner, J., Karagiannis, G., and A. McDonald, "NSIS Signaling Layer Protocol (NSLP) for Quality-of-Service Signaling", [RFC 5974](#), October 2010.
- [RFC6411] Behringer, M., Le Faucheur, F., and B. Weis, "Applicability of Keying Methods for RSVP Security", [RFC 6411](#), October 2011.

Appendix A. Examples of Bandwidth Allocation Model for Admission Priority

Appendices 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] can also be used in a similar manner for support of admission priority. [Appendix 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 cannot 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).




Figure 4: MAM Bandwidth Allocation

Figure 4 shows a link that is within a routed network and conforms 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 fully utilized currently.

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 4, an operator may map the MAM to 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 4. 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 4. 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 4.

Figure 5 shows some of the non-priority capacity of this link being used.

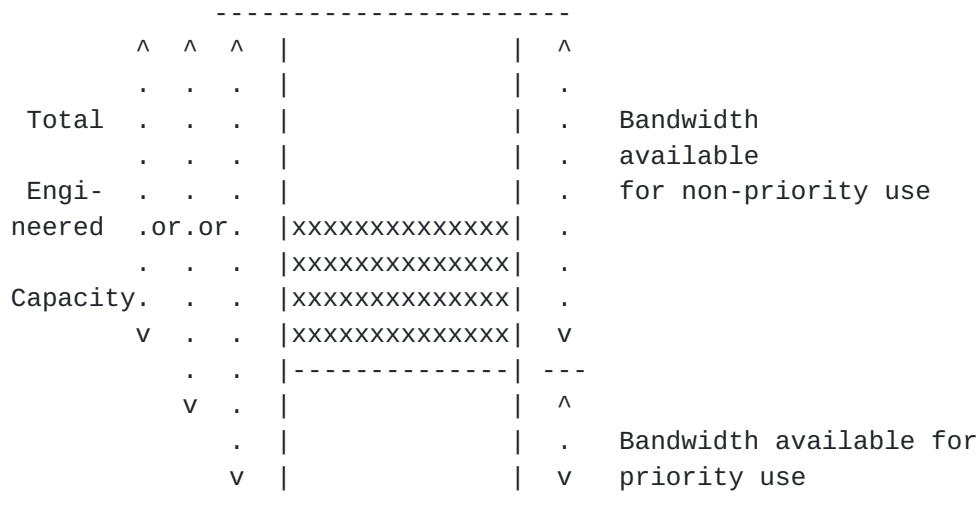


Figure 5: Partial Load of Non-Priority Calls

Figure 6 shows the same amount of non-priority load being used at this link and a small amount of priority bandwidth being used.

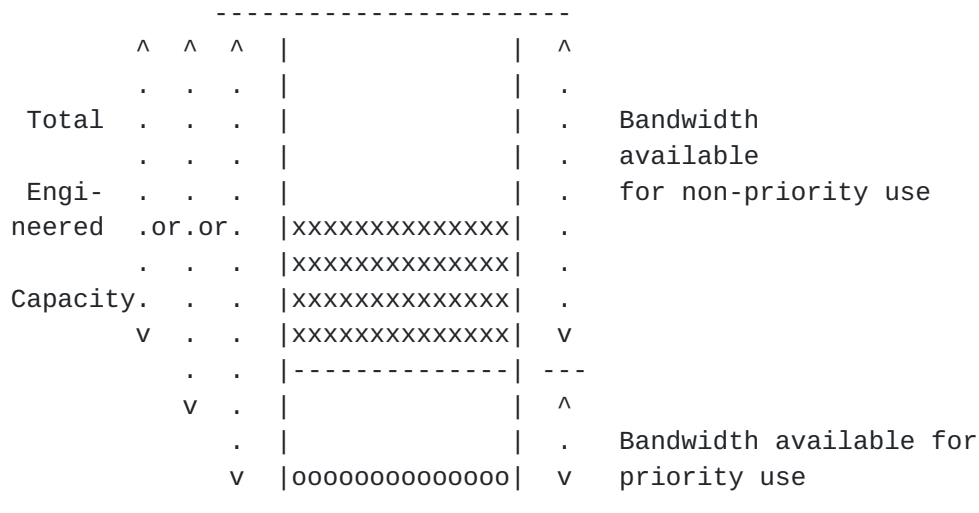


Figure 6: Partial Load of Non-Priority Calls and Partial Load of Priority Calls

Figure 7 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.

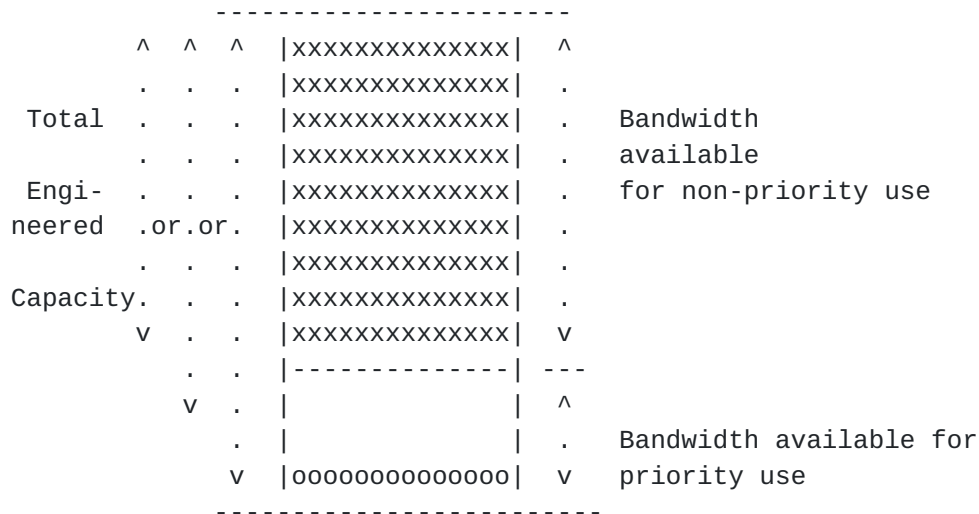


Figure 7: Full Non-Priority Load and Partial Load of Priority Calls

Figure 8 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 queueing of new priority session requests until capacity becomes available again for priority traffic).

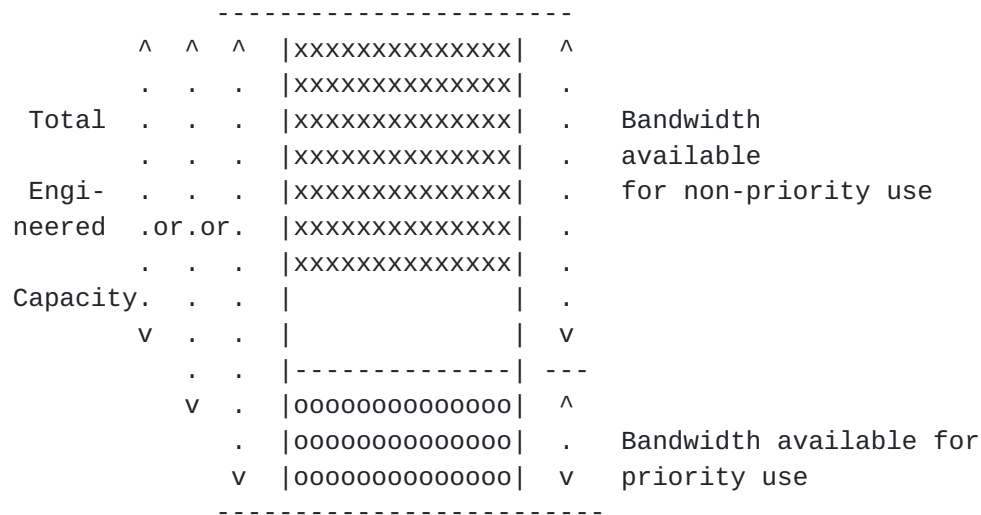


Figure 8: Partial Non-Priority Load and Full Priority Load

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 RDM is that priority traffic can use the bandwidth that is not currently used by non-priority traffic.

As with the MAM, an operator may map the RDM 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 9 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.

```

-----
|xxxxxxxxxxxxxxxx| . ^
|xxxxxxxxxxxxxxxx| . Bandwidth .
|xxxxxxxxxxxxxxxx| . available for .
|xxxxxxxxxxxxxxxx| . non-priority .
|xxxxxxxxxxxxxxxx| . use .
|xxxxxxxxxxxxxxxx| . . Bandwidth
| | . . available for
| | v . non-priority
|-----| --- . and priority
| | . use
| | .
| | v
|oooooooooooooooo|
-----

```

Figure 9: Partial Non-Priority Load and Partial Aggregate Load

Figure 10 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.

```

-----
|xxxxxxxxxxxxxxxx| . ^
|xxxxxxxxxxxxxxxx| . Bandwidth .
|xxxxxxxxxxxxxxxx| . available for .
|xxxxxxxxxxxxxxxx| . non-priority .
|xxxxxxxxxxxxxxxx| . use .
|xxxxxxxxxxxxxxxx| . . Bandwidth
|xxxxxxxxxxxxxxxx| . . available for
|xxxxxxxxxxxxxxxx| v . non-priority
|-----| --- . and priority
| | . use
| | .
| | v
|oooooooooooooooo|
-----

```

Figure 10: Full Non-Priority Load and Partial Aggregate Load

Figure 11 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 10, 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
|ooooooooooooo| v . non-priority
|-----| --- . and priority
|ooooooooooooo| . use
|ooooooooooooo| .
|ooooooooooooo| v
-----

```

Figure 11: Partial Non-Priority Load and Heavy Aggregate Load

Figure 12 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, and 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 queueing 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
|oooooooooooooooo| . use
|oooooooooooooooo| .
|oooooooooooooooo| v
-----

```

Figure 12: Full Non-Priority Load and 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 PrBM, 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 congestion periods where priority sessions may be more heavily used, those sessions always still represent a fairly small proportion of the overall load that 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, an operator may map the PrBM 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 limit. 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 PrBM.

Figure 13 illustrates the bandwidth allocation with the PrBM.




Figure 13: Priority Bypass Model Bandwidth Allocation

Figure 14 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.




Figure 14: Partial Load of Non-Priority Calls

Figure 15 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.

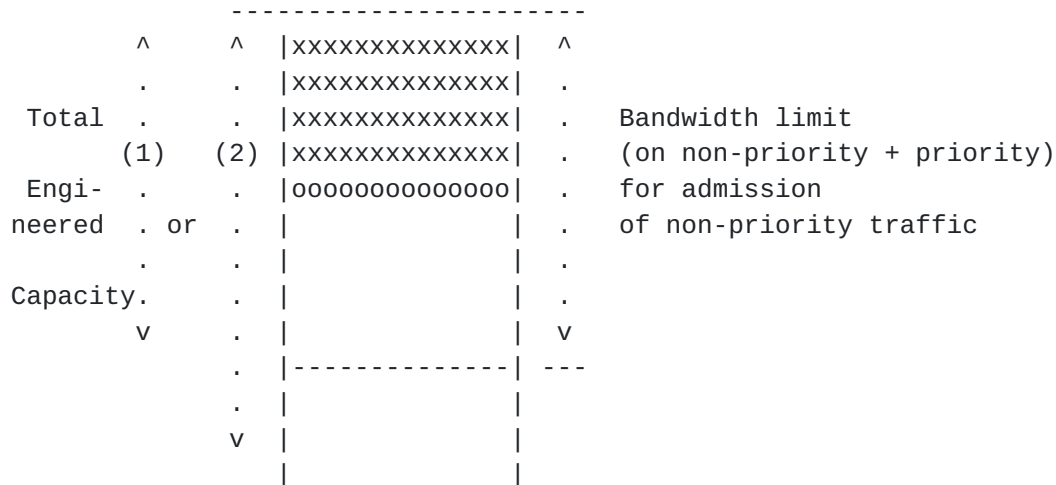


Figure 15: Partial Load of Non-Priority Calls and Partial Load of Priority Calls

Figure 16 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.

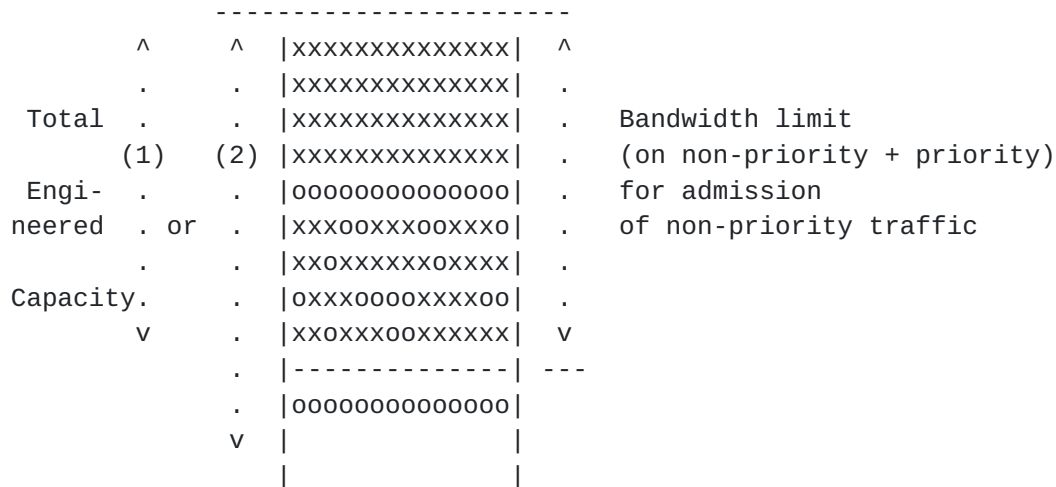


Figure 16: Full Non-Priority Load

Appendix B. Example Usages of RSVP Extensions

This section provides examples of how RSVP extensions defined in this document can be used (in conjunction with other RSVP functionality and SIP functionality) to enforce different hypothetical policies for handling prioritized 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.

We refer here to "Session 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 cannot 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 alternate routing and exemption from certain network management controls.

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 a policy area boundary 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 a prioritized service purely based on Session Queueing, one can achieve this by signaling prioritized sessions:

- o using the "Resource-Priority" header in SIP
- o not using the Admission-Priority Policy Element in RSVP
- o not using the Preemption Policy Element in RSVP

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

- o using the "Resource-Priority" header in SIP
- o using the Admission-Priority Policy Element in RSVP
- o not using the Preemption Policy Element in RSVP

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

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

- o using the "Resource-Priority" header in SIP
- o using the Admission-Priority Policy Element in RSVP
- o using the Preemption Policy Element in RSVP with:
 - * setup (Prioritized-1) > defending (Prioritized-2)
 - * setup (Prioritized-2) <= defending (Prioritized-1)

- * setup (Prioritized-1) <= defending (Non-Prioritized)
- * setup (Prioritized-2) <= defending (Non-Prioritized)

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

- o using the "Resource-Priority" header in SIP
- o using the Admission-Priority Policy Element in RSVP
- o using the Preemption Policy Element in RSVP with:
 - * setup (Prioritized) > defending (Non-Prioritized)
 - * setup (Non-Prioritized) <= defending (Prioritized)

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

- o using the "Resource-Priority" header in SIP
- o using the Admission-Priority Policy Element in RSVP
- o using the Preemption Policy Element in RSVP with:
 - * setup (Prioritized) > defending (Non-Prioritized)
 - * setup (Non-Prioritized) <= defending (Prioritized)
- 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

