

The Use of RSVP with IETF Integrated Services

Abstract

This note describes the use of the RSVP resource reservation protocol with the Controlled-Load and Guaranteed QoS control services. The RSVP protocol defines several data objects which carry resource reservation information but are opaque to RSVP itself. The usage and data format of those objects is given here.

Status of this Memo

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

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

To learn the current status of any Internet-Draft, please check the "l1d-abstracts.txt" listing contained in the Internet- Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

This draft is a product of the Integrated Services Working Group of the Internet Engineering Task Force. Comments are solicited and should be addressed to the working group's mailing list at int-serv@isi.edu and/or the author(s).

1. Introduction

The Internet integrated services framework provides the ability for applications to choose among multiple, controlled levels of delivery service for their data packets. To support this capability, two

things are required:

- Individual network elements (subnets and IP routers) along the path followed by an application's data packets must support mechanisms to control the quality of service delivered to those packets.
- A way to communicate the application's requirements to network elements along the path, and to convey QoS management information between network elements and the application must be provided.

In the integrated services framework the first function is provided by QoS control services such as Controlled-Load [[RFCCL](#)] and Guaranteed [[RFCG](#)]. The second function may be provided in a number of ways, but is frequently implemented by a resource reservation setup protocol such as RSVP [[RFCRSVP](#)].

Because RSVP is designed to be used with a variety of QoS control services, and because the QoS control services are designed to be used with a variety of setup mechanisms, a logical separation exists between the two specifications. The RSVP specification does not define the internal format of those RSVP protocol fields, or objects, which are related to invoking QoS control services. Rather, RSVP treats these objects as opaque. The objects can carry different information to meet different application and QoS control service requirements.

Similarly, interfaces to the QoS control services are defined in a general format, so that the services can be used with a variety of setup mechanisms.

This RFC provides the information required to use RSVP and the integrated service framework's QoS control services together. It defines the usage and contents of three RSVP protocol objects, the FLOWSPEC, ADSPEC, and SENDER_TSPEC, in an environment supporting the Controlled-Load and/or Guaranteed QoS control services. If new services or capabilities are added to the integrated services framework, this note will be revised as required.

2. Use of RSVP

Several types of data must be transported between applications and network elements to correctly invoke QoS control services (1). This

1. In addition to the data used to directly invoke QoS control services, RSVP carries authentication, accounting, and policy information needed to manage the use of these services. This note is concerned only with the RSVP objects

data includes:

- Information generated by each receiver describing the QoS control service desired, a description of the traffic flow to which the resource reservation should apply (the Receiver TSpec), and whatever parameters are required to invoke the service (the Receiver RSpec). This information is carried from the receivers to intermediate network elements and the sender(s) by RSVP FLOWSPEC objects. The information being carried in a FLOWSPEC object may change at intermediate points in the network due to reservation merging and other factors.
- Information generated at each sender describing the data traffic generated by that sender (the Sender TSpec). This information is carried from the sender to intermediate network elements and the receiver(s) by RSVP, but is never modified by intermediate elements within the network. This information is carried in RSVP SENDER_TSPEC objects.
- Information generated or modified within the network and used at the receivers to make reservation decisions. This information might include available services, delay and bandwidth estimates, and operating parameters used by specific QoS control services. This information is collected from network elements and carried towards receivers in RSVP ADSPEC objects. Rather than carrying information from each intermediate node separately to the receivers, the information in the ADSPEC represents a summary, computed as the ADSPEC passes each hop. The size of this summary remains (roughly) constant as the ADSPEC flows through the network, giving good scaling properties.

From the point of view of RSVP objects, the breakdown is as follows:

- The RSVP SENDER_TSPEC object carries the traffic specification (sender TSpec) generated by each data source within an RSVP session. It is transported unchanged through the network, and delivered to both intermediate nodes and receiving applications.
- The RSVP ADSPEC object carries information which is generated at either data sources or intermediate network elements, is flowing downstream towards receivers, and may be used and updated inside the network before being delivered to receiving applications. This information includes both parameters describing the properties of the data path, including the availability of specific QoS control services, and parameters required by specific QoS control services

needed to actually invoke QoS control services, and does not discuss accounting or policy objects.

to operate correctly.

- The RSVP FLOWSPEC object carries reservation request (Receiver_TSpec and RSpec) information generated by data receivers. The information in the FLOWSPEC flows upstream towards data sources. It may be used or updated at intermediate network elements before arriving at the sending application.

NOTE: The existence of both SENDER_TSPEC and ADSPEC RSVP objects is somewhat historical. Using the message format described in this note it would be possible to place all of the service control information carried "downstream" by RSVP in the same object. However, the distinction between data which is not updated within the network (in the SENDER_TSPEC object) and data which is updated within the network (in the ADSPEC object) may be useful to an implementation in practice, and is therefore retained.

2.1 Summary of operation

Operation proceeds as follows:

An application instance participating in an RSVP session as a data sender registers with RSVP. One piece of information provided by the application instance is the Sender TSpec describing the traffic the application expects to generate. This information is used to construct an RSVP SENDER_TSPEC object, which is included in RSVP PATH messages generated for the application.

The sending application also constructs an initial RSVP ADSPEC object. This adspec carries information about the QoS control capabilities and requirements of the sending application itself, and forms the starting point for the accumulation of path properties described below. The ADSPEC is added to the RSVP PATH message created at the sender.

NOTE: For application programmer convenience, a host RSVP implementation may allow the sending application not to provide an initial adspec, instead supplying its own default. This usage is most likely when the application sender does not itself participate in the end-to-end QoS control process (by actively scheduling CPU usage and similar means) and does not itself care which QoS control service is selected by the receivers.

The ADSPEC is modified by subsequent network elements as the RSVP PATH message moves from sender to receiver(s). At each network element, the ADSPEC is passed from RSVP to the traffic control module. The traffic control module updates the ADSPEC, which may

contain data for several QoS control services, by identifying the services mentioned in the ADSPEC and calling each such service to update its portion of the ADSPEC. If the traffic control module discovers a QoS control service mentioned in the ADSPEC but not implemented by the network element, a flag is set to report this to the receiver. The updated ADSPEC is then returned to RSVP for delivery to the next hop along the path.

Upon arrival of the PATH message at an application receiver, the data in the SENDER_TSPEC and ADSPEC objects is passed across the RSVP API to the application. The application (perhaps with the help of a library of common resource-reservation functions) interprets the arriving data, and uses it to guide the selection of resource reservation parameters. Examples of this include use of the arriving "PATH_MTU" composed characterization parameter [[RFCGP](#)] to determine the maximum packet size parameter in the reservation request and use of the arriving Guaranteed service "C" and "D" parameters [[RFCG](#)] to calculate a mathematical bound on delivered packet delay when using the Guaranteed service.

An application receiver wishing to make a resource reservation supplies its local RSVP with the necessary reservation parameters. Among these are the QoS control service desired (Guaranteed or Controlled-Load), the traffic specifier (TSPEC) describing the level of traffic for which resources should be reserved, and, if needed by the selected QoS control service, an RSPEC describing the level of service desired. These parameters are composed into an RSVP FLOWSPEC object and transmitted upstream by RSVP.

At each RSVP-aware point in the network, the SENDER_TSPECS arriving in PATH messages and the FLOWSPECS arriving in RESV messages are used to request an appropriate resource reservation from the desired QoS control service. State merging, message forwarding, and error handling proceed according to the rules of the RSVP protocol.

Finally, the merged FLOWSPEC object arriving at each of an RSVP session's data senders is delivered to the application to inform each sender of the merged reservation request and properties of the data path.

[2.2. RSVP support for multiple QoS control services](#)

The design described in this note supports RSVP sessions in which the receivers choose a QoS control service at runtime.

To make this possible, a receiver must have all the information needed to choose a particular service before it makes the choice. This implies that the RSVP SENDER_TSPEC and ADSPEC objects must

provide the receivers with information for all services which might be chosen.

The Sender TSpec used by the two currently defined QoS control services is identical. This simplifies the RSVP SENDER_TSPEC, which need carry only a single TSpec data structure in this shared format. This SENDER_TSPEC can be used with either Guaranteed or Controlled-Load service.

The RSVP ADSPEC carries information needed by receivers to choose a service and determine the reservation parameters. This includes:

- Whether or not there is a non-RSVP hop along the path. If there is a non-RSVP hop, the application's traffic will receive reservationless best-effort service at at least one point on the path.
- Whether or not a specific QoS control service is implemented at every hop along the path. For example, a receiver might learn that at least one integrated-services aware hop along the path supports the Controlled-Load service but not the Guaranteed service.
- Default or global values for the general characterization parameters described in [[RFCGP](#)]. These values describe properties of the path itself, irrespective of the selected QoS control service. A value reported in this section of the ADSPEC applies to all services unless a different, service-specific value is also present in the ADSPEC.
- A service-specific value for one or more general characterization parameters, if the service-specific value differs from the default value.
- Values of the per-service characterization parameters defined by each supported service.

Data in the ADSPEC is divided into blocks or fragments, each of which is associated with a specific service. This allows the adspec to carry information about multiple services, allows new services to be deployed in the future without immediately updating existing code, and allows an application which will never use a particular service to omit the ADSPEC data for that service. The structure of the ADSPEC is described in detail in [Section 3.3](#).

A sender may indicate that a specific QoS control service should *not* be used by the receivers within an RSVP session. This is done by omitting all mention of that service from the ADSPEC, as described in [Section 3.3](#). Upon arrival at a receiver, the complete absence of

an ADSPEC fragment for a specific service indicates to receivers that the service should not be used.

NOTE: In RSVP Version 1, all receivers within a session are required to choose the same QoS control service. This restriction is imposed by the difficulty of merging reservations requesting different QoS control services, and the current lack of a general service replacement mechanism. The restriction may be eliminated in the future.

Considering this restriction, it may be useful to coordinate the receivers' selection of a QoS control service by having the sender(s) offer only one choice, using the ADSPEC mechanism mentioned above. All receivers must then select the same service. Alternatively, the coordination might be accomplished by using a higher-level session announcement and setup mechanism to inform the receivers of the QoS control service in use, by manual configuration of the receivers, or by an agreement protocol running among the session receivers themselves.

As with the ADSPEC, the FLOWSPEC and SENDER_TSPEC object formats described in [Section 3](#) are capable of carrying TSspecs and RSpecs for more than one QoS control service in separate data fragments. Currently, use of a FLOWSPEC or SENDER_TSPEC containing fragments for more than one QoS control service is not supported. In the future, this capability may be used to implement a more flexible service request and replacement scheme, allowing applications to obtain useful end-to-end QoS control when not all intermediate nodes support the same set of QoS services. RSVP-application APIs should be designed to support passing SENDER_TSPEC, FLOWSPEC, and ADSPEC objects of variable size and containing information about multiple QoS control services between RSVP and its clients.

[2.3. Use of ADSPEC Information](#)

This section gives some details about setting reservation parameters and the use of information conveyed by the RSVP ADSPEC object.

[2.3.1. Determining the availability of a QoS control service](#)

The RSVP ADSPEC carries flag bits telling the application receivers whether or not a completely reservation-capable path exists between each sender and the receiver. These bits are called "break bits", because they indicate breaks in the QoS control along a network path. Break bits are carried within the header which begins each per-service data fragment of an RSVP ADSPEC.

Service number 1 is used within the ADSPEC to carry information about

default or global parameter values [[RFCGP](#)], rather than values for a particular QoS control service. The break bit in Service 1's per-service header tells the receiver(s) whether all of the network elements along the path from sender to receiver support RSVP and integrated services. If a receiver finds this bit set, at least one network element along the data transmission path between the ADSPEC's sender and the receiver can not provide QoS control services at all. This bit corresponds to the global NON_IS_HOP characterization parameter defined in [[RFCGP](#)].

NOTE: If this bit is set, the values of all other parameters in the ADSPEC are unreliable. The bit being set indicates that at least one node along the sender-receiver path did not fully process the ADSPEC.

Service-specific break bits tell the receiver(s) whether all of the network elements along the path from sender to receiver support a particular QoS control service. The break bit for each service is carried within the ADSPEC's per-service header for that service. If a bit is set at the receiver, at least one network element along the data transmission path supports RSVP but does not support the QoS control service corresponding to the per-service header. These bits correspond to the service-specific NON_IS_HOP characterization parameters defined in [[RFCGP](#)].

[Section 3](#) gives more information about break bits.

[2.3.2. Determining Path MTU](#)

Both Guaranteed and Controlled-Load QoS control services place an upper bound on packet size, and require that the application limit the maximum size of packets subject to resource reservation. For both services, the desired maximum packet size is a parameter of the reservation request, and the service will reject (with an admission control error) reservation requests specifying a packet size larger than that supported by the service.

Since RSVP reservation requests are made by receivers, this implies that the *receivers* in an RSVP session, as well as the senders, need to know the MTU supported by the QoS control services along a data path. Further, in some unusual cases the MTU supported by a QoS control service may differ from that supported by the same router when providing best effort service.

A form of MTU negotiation is used to address these problems. MTU negotiation in an RSVP system works as follows:

- Each sending application joining an RSVP session fills in the M

(MTU) parameter in the TSpec (carried from senders to receivers in a SENDER_TSPEC object) with the maximum packet size it wishes to send.

- Each RSVP PATH message from a sending application also carries an ADSPEC object containing at least one PATH_MTU characterization parameter. When it arrives at the receiver, this parameter gives the minimum MTU at any point along the path from sender to receiver. Generally, only the "global" PATH_MTU parameter (service 1, parameter 9) will be present, in which case its value is a legal MTU for all reservation requests. If a service specific PATH_MTU parameter is present, its value will be smaller than that of the global parameter, and should be used for reservation requests for that service.

- Each receiver takes the minimum of all the MTU's (for the desired QoS control service) arriving in ADSPEC messages from different senders and uses that value as the MTU in its reservation requests. This value is used to fill in the M parameter of the TSpec created at the receiver. In the case of a FF style reservation, a receiver may also choose to use the MTU derived from each sender's ADSPEC in the FLOWSPEC generated for that sender, if the receiver is concerned about obtaining the maximum MTU on each data path. To accomodate changes in the data path, the receiver may continue to watch the arriving ADSPECS, and modify the reservation if a newly arriving ADSPEC indicates a smaller MTU than is currently in use.

- As reservation requests (RESV messages) move from receivers to senders, reservation parameters are merged at intermediate nodes. As part of this merging, the smaller of two M parameters arriving at a merge point will be forwarded in the upstream RESV message.

- As reservation requests arrive at intermediate RSVP's, the minimum of the receivers' requested TSpec and the sum of the sender TSpecs is taken, and a reservation for the resulting TSpec is made. The reservation will use the smaller of the actual path MTU value computed by the receivers and the largest maximum packet size declared by any of the sender(s). (The TSpec sum() function result's M parameter is the max of the summed TSpec M parameters).

- When the completely merged RESV message arrives at each sender, the MTU value (M parameter) in the merged FLOWSPEC object will have been set to the smallest acceptable MTU of the data paths from that sender to any session receiver. This MTU should be used by the sending application to size its packets. Any packets larger than this MTU may be delivered as best-effort rather than covered by the session's resource reservation.

Note that the scheme above will allow each sender in a session to use the largest MTU appropriate for that sender, in cases where different data paths or receivers have different acceptable MTU's.

3. RSVP Object Formats

This section specifies the detailed contents and wire format of RSVP SENDER_TSPEC, ADSPEC, and FLOWSPEC objects for use with the Guaranteed and Controlled-Load QoS control services. The object formats specified here are based on the data naming and encoding rules described in [\[RFCDE\]](#).

3.1. RSVP SENDER_TSPEC Object

The RSVP SENDER_TSPEC object carries information about a data source's generated traffic. The required RSVP SENDER_TSPEC object contains a global Token_Bucket_TSpec parameter (service_number 1, parameter 127, as defined in [\[RFCGP\]](#)). This TSpec carries traffic information usable by either the Guaranteed or Controlled-Load QoS control services.

	31	24 23	16 15	8 7	0
1	0 (a)	Unused		6 (b)	
2	1 (c)	5 (d)	127 (e)	0 (f)	
3	Token Bucket Rate [r] (32-bit IEEE floating point number)				
4	Token Bucket Size [b] (32-bit IEEE floating point number)				
5	Peak Data Rate [p] (32-bit IEEE floating point number)				
6	Minimum Policed Unit [m] (32-bit integer)				
7	Maximum Packet Size [M] (32-bit integer)				

- (a) - Message format version number (0)
- (b) - Overall length (6 words not including header)
- (c) - Service header, service number 1 (default/global information)
- (d) - Length of per-service data, 5 words not including header
- (e) - Parameter ID, parameter 127 (Token_Bucket_TSpec)
- (f) - Parameter 127 flags (none set)

In this TSpec, the parameters [r] and [b] are set to reflect the sender's view of its generated traffic. The peak rate parameter [p]

may be set to the sender's peak traffic generation rate (if known and controlled), the physical interface line rate (if known), or positive infinity (if no better value is available). Positive infinity is represented as an IEEE single-precision floating-point number with an exponent of all ones (255) and a sign and mantissa of all zeros. The use of infinity is described further in [\[RFCGP\]](#).

The minimum policed unit parameter [m] should generally be set equal to the smallest packet size generated by the application. Note that smaller values of this parameter may lead to increased likelihood of admission control failure when a receiver tries to make a reservation. In some cases, an application transmitting a small percentage of very small packets may therefore choose to set the value of [m] larger than the actual minimum transmitted packet size. This will increase the likelihood of the reservation succeeding, at the expense of policing packets of size less than [m] as if they were of size [m].

The maximum packet size parameter [M] should be set to the size of the largest packet the application expects to generate.

3.2. RSVP FLOWSPEC Object

The RSVP FLOWSPEC object carries information necessary to make reservation requests from the receiver(s) into the network. This includes an indication of which QoS control service is being requested, and the parameters needed for that service.

The QoS control service requested is indicated by the `service_number` in the FLOWSPEC's per-service header.

3.2.1 FLOWSPEC object when requesting Controlled-Load service

The format of an RSVP FLOWSPEC object originating at a receiver requesting Controlled-Load service is shown below. Each of the TSPEC fields is represented using the preferred concrete representation specified in the 'Invocation Information' section of [[RFCCL](#)]. The value of 5 in the per-service header (field (c), below) indicates that Controlled-Load service is being requested.

	31	24 23	16 15	8 7	0
1	0 (a)	Unused		6 (b)	
2	5 (c)	0	5 (d)	127 (e)	0 (f)
3	Token Bucket Rate [r] (32-bit IEEE floating point number)				


```

4 | Token Bucket Size [b] (32-bit IEEE floating point number) |
  +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
5 | Peak Data Rate [p] (32-bit IEEE floating point number) |
  +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
6 | Minimum Policed Unit [m] (32-bit integer) |
  +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
7 | Maximum Packet Size [M] (32-bit integer) |
  +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- (a) - Message format version number (0)
- (b) - Overall length (6 words not including header)
- (c) - Service header, service number 5 (Controlled-Load)
- (d) - Length of per-service data, 5 words not including per-service header
- (e) - Parameter ID, parameter 127 (Token Bucket TSpec)
- (f) - Parameter 127 flags (none set)

In this object, the TSpec parameters [r] and [b] are set to reflect the traffic parameters of the receiver's desired reservation. The peak rate parameter [p] should be set to the largest peak traffic generation rate specified in an arriving SENDER_TSPEC object. As described above, this value may be infinity.

The maximum packet size parameter [M] should be set to the value of the path MTU, which the receiver learns from information in arriving RSVP ADSPEC objects. However, if the receiving application has built-in knowledge of the maximum packet size in use within the RSVP session, that value may be used to set the [M] parameter of the FLOWSPEC object. See [section 2.3.2](#) for further discussion of the MTU value.

3.2.2. FLOWSPEC Object when Requesting Guaranteed Service

The format of an RSVP FLOWSPEC object originating at a receiver requesting Guaranteed service is shown below. The flowspec object used to request guaranteed service carries a TSpec and RSpec specifying the traffic parameters of the flow desired by the receiver.

Each of the TSpec and RSpec fields is represented using the preferred concrete representation specified in the 'Invocation Information' section of [\[RFCG\]](#). The value of 2 for the service header identifier (field (c) in the picture below) indicates that Guaranteed service is being requested.

```

          31          24 23          16 15          8 7          0
  +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
1 | 0 (a) |      Unused          |          9 (b)          |
  +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```



```

2  |      2 (c)      |0|      8 (d)      | 127 (e)      |      0 (f)      |
   +-+-+-+-+-+-+-+-+
3  |  Token Bucket Rate [r] (32-bit IEEE floating point number)  |
   +-+-+-+-+-+-+-+-+
4  |  Token Bucket Size [b] (32-bit IEEE floating point number)  |
   +-+-+-+-+-+-+-+-+
5  |  Peak Data Rate [p] (32-bit IEEE floating point number)     |
   +-+-+-+-+-+-+-+-+
6  |  Minimum Policed Unit [m] (32-bit integer)                   |
   +-+-+-+-+-+-+-+-+
7  |  Maximum Packet Size [M]  (32-bit integer)                   |
   +-+-+-+-+-+-+-+-+
8  |      130 (g)      |      0 (h)      |      zero pad      |
   +-+-+-+-+-+-+-+-+
9  |  Rate [R]  (32-bit IEEE floating point number)               |
   +-+-+-+-+-+-+-+-+
10 |  Slack Term [S]  (32-bit integer)                             |
   +-+-+-+-+-+-+-+-+

```

- (a) - Message format version number (0)
- (b) - Overall length (9 words not including header)
- (c) - Service header, service number 2 (Guaranteed)
- (d) - Length of per-service data, 9 words not including per-service header
- (e) - Parameter ID, parameter 127 (Token Bucket TSpec)
- (f) - Parameter 127 flags (none set)
- (g) - Parameter ID, parameter 130 (Guaranteed Service RSpec).
- (h) - Parameter 130 flags (none set).

In this object, the TSpec parameters [r] and [b] are set to reflect the traffic parameters of the receiver's desired reservation. The peak rate parameter [p] should be set to the largest peak traffic generation rate specified in an arriving SENDER_TSPEC object. As described above, this value may be infinity.

The maximum packet size parameter [M] should be set to the value of the path MTU, which the receiver learns from information in arriving RSVP ADSPEC objects. However, if the receiving application has built-in knowledge of the maximum packet size in use within the RSVP session, that value may be used to set the [M] parameter of the FLOWSPEC object.

The RSpec terms [R] and [S] are selected to obtain the desired bandwidth and delay guarantees. This selection is described in [\[RFCG\]](#).

3.3. RSVP ADSPEC Object

An RSVP ADSPEC object is constructed from data fragments contributed by each service which might be used by the application. The ADSPEC begins with an overall message header, followed by a fragment for the default general parameters, followed by fragments for every QoS control service which may be selected by application receivers. The size of the ADSPEC varies depending on the number and size of per-service data fragments present and the presence of non-default general parameters (described in [Section 3.3.5](#)).

The complete absence of a data fragment for a particular service means that the application sender does not know or care about that service, and is a signal to intermediate nodes not to add or update information about that service to the ADSPEC. It is also a signal to application receivers that they should not select that service when making reservations.

Each fragment present is identified by a per-service data header. Each header contains a field identifying the service, a break bit, and a length field.

The length field allows the ADSPEC information for a service to be skipped over by a network elements which does not recognize or implement the service. When an element does this, it sets the break bit, indicating that the service's ADSPEC data was not updated at at least one hop. Note that a service's break bit can be set without otherwise supporting the service in any way. In all cases, a network element encountering a per-service data header it does not understand simply sets bit 23 to report that the service is not supported, then skips over the rest of the fragment.

Data fragments must always appear in an ADSPEC in service_number order. In particular, the default general parameters fragment (service_number 1) always comes first.

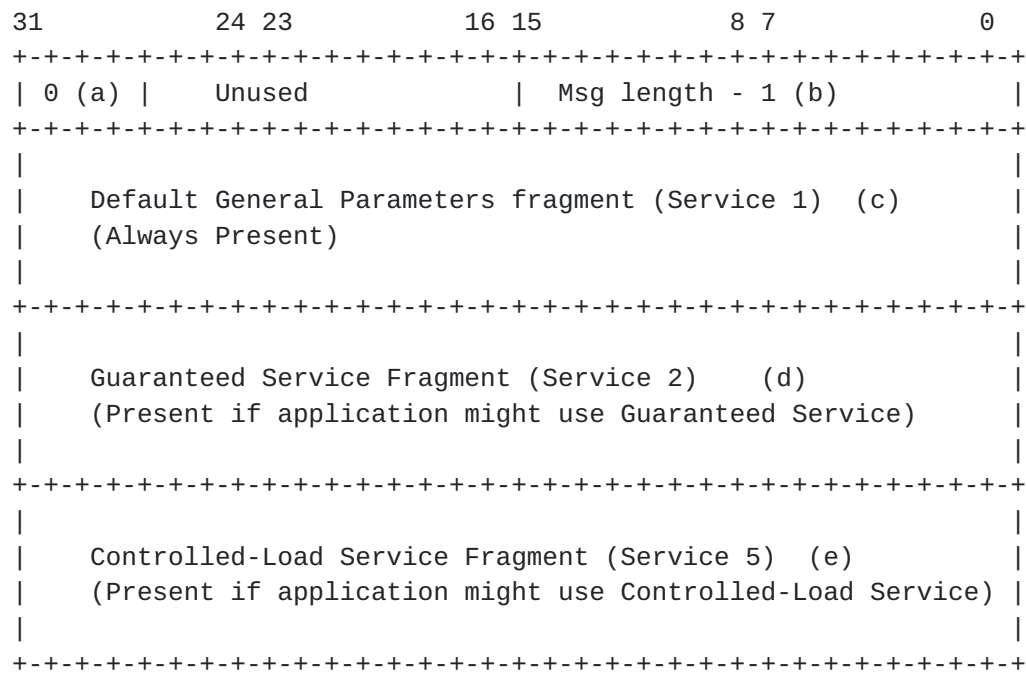
Within a data fragment, the service-specific data must always come first, followed by any non-default general parameters which may be present, ordered by parameter_number. The size and structure of the service-specific data is fixed by the service definition, and does not require run-time parsing. The remainder of the fragment, which carries non-default general parameters, varies in size and structure depending on which, if any, of these parameters are present. This part of the fragment must be parsed by examining the per-parameter headers.

Since the overall size of each data fragment is variable, it is always necessary to examine the length field to find the end of the

fragment, rather than assuming a fixed-size structure.

3.3.1. RSVP ADSPEC format

The basic ADSPEC format is shown below. The message header and the default general parameters fragment are always present. The fragments for Guaranteed or Controlled-Load service may be omitted if the service is not to be used by the RSVP session. Additional data fragments will be added if new services are defined.




(a) - Message format version number (0)

(b) - Overall message length not including header word

(c, d, e) - Data fragments

3.3.2. Default General Characterization Parameters ADSPEC data fragment

All RSVP ADSPECs carry the general characterization parameters defined in [\[RFCGP\]](#). Values for global or default general parameters (values which apply to the all services or the path itself) are carried in the per-service data fragment for service number 1, as shown in the picture above. This fragment is always present, and always first in the message.




```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
3 | Path b/w estimate (32-bit IEEE floating point number) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
4 |      8 (i)      |i      (j)      |      zero pad      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
5 |      Minimum path latency (32-bit integer) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
6 |      10 (k)      |i      (l)      | composed MTU (16 bits) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

(c) - Per-Service header, service number 1 (Default General Parameters)

(d) - Global Break bit ([\[RFCGP\]](#), Parameter 2) (marked x) and length of

General Parameters data block.

(e) - Parameter ID, parameter 4 (Number-of-IS-hops param from [\[RFCGP\]](#))

(f) - Parameter 4 flag byte (bit i = INVALID, others unassigned)

(g) - Parameter ID, parameter 6 (Path-BW param from [\[RFCGP\]](#))

(h) - Parameter 6 flag byte (bit i = INVALID, others unassigned)

(i) - Parameter ID, parameter 8 (minimum path latency from [\[RFCGP\]](#))

(j) - Parameter 8 flag byte (bit i = INVALID, others unassigned)

(k) - Parameter ID, parameter 10 (composed path MTU from [\[RFCGP\]](#))

(l) - Parameter 10 flag byte (bit i = INVALID, others unassigned)

Rules for composing general parameters appear in [\[RFCGP\]](#). The composition rules for some of these parameters allow an implementation to set the parameter's INVALID flag in the ADSPEC, if a local value for that parameter is not available. These bits are marked 'i' in the diagram.

In the above fragment, the global break bit (bit 23 of word 1, marked with (x) in the picture) is used to indicate the existence of a network element not supporting QoS control services somewhere in the data path. This bit is cleared when the ADSPEC is created, and set to one if a network element which does not support RSVP or integrated services is encountered. An ADSPEC arriving at a receiver with this bit set indicates that all other parameters in the ADSPEC may be invalid, since not all network elements along the path support updating of the ADSPEC.

The general parameters are updated at every network node which supports RSVP:

- When a PATH message ADSPEC encounters a network element implementing integrated services, the portion of the ADSPEC associated with service number 1 is passed to the module implementing general parameters. This module updates the global general parameters.

- When a PATH message ADSPEC encounters a network element that does *not* support RSVP or implement integrated services, the break bit in the general parameters service header must be set. In practice, this bit will usually be set by another network element which supports RSVP, but has been made aware of the gap in integrated services coverage.
- In either case, the ADSPEC is passed back to RSVP for delivery to the next hop along the path.

3.3.3. Guaranteed Service ADSPEC data fragment

The Guaranteed service uses the RSVP ADSPEC to carry data needed to compute the C and D terms passed from the network to the application. The minimum size of a non-empty guaranteed service data fragment is 8 32-bit words. The ADSPEC fragment for Guaranteed service has the following format:

	31	24 23	16 15	8 7	0
1	<div> <div>2 (a)</div> <div> x N-1 (b)</div> <div>133 (c)</div> <div>0 (d)</div> </div>				
2	End-to-end composed value for C [Ctot] (32-bit integer)				
3	<div> <div>134 (e)</div> <div>0 (f)</div> <div>zero pad</div> </div>				
4	End-to-end composed value for D [Dtot] (32-bit integer)				
5	<div> <div>135 (g)</div> <div>0 (h)</div> <div>zero pad</div> </div>				
6	Since-last-reshaping point composed C [Csum] (32-bit integer)				
7	<div> <div>136 (i)</div> <div>0 (j)</div> <div>zero pad</div> </div>				
8	Since-last-reshaping point composed D [Dsum] (32-bit integer)				
9	Service-specific general parameter headers/values, if present				
N					

- (a) - Per-Service header, service number 2 (Guaranteed)
- (b) - Break bit and Length of per-service data in 32-bit words not including header word.
- (c) - Parameter ID, parameter 133 (Composed Ctot)
- (d) - Parameter 133 flags (none set).
- (e) - Parameter ID, parameter 134 (Composed Dtot)
- (f) - Parameter 134 flags (none set).

- (g) - Parameter ID, parameter 135 (Composed Csum).
- (h) - Parameter 135 flags (none set).
- (i) - Parameter ID, parameter 136 (Composed Dsum).
- (j) - Parameter 136 flags (none set).

When a node which actually implements guaranteed service creates the guaranteed service adspec fragment, the parameter values are set to the local values for each parameter. When an application or network element which does not itself implement guaranteed service creates a guaranteed service adspec fragment, it should set the values of each parameter to zero, and set the break bit to indicate that the service is not actually implemented at the node.

An application or host RSVP which is creating a guaranteed service adspec fragment but does not itself implement the guaranteed service may create a truncated "empty" guaranteed adspec fragment consisting of only a header word:

```

      31             24 23             16 15             8 7             0
      +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
1  |      2 (a)      |1|    0 (b)    |    zero padding      |
      +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

- (a) - Per-Service header, service number 2 (Guaranteed)
- (b) - Break bit and Length of per-service data in 32-bit words not including header word.

This might occur if the sending application or host does not do resource reservation itself, but still wants the network to do so. Note that in this case the break bit will always be set, since the creator of the adspec fragment does not itself implement guaranteed service.

When a PATH message ADSPEC containing a per-service header for Guaranteed service encounters a network element implementing Guaranteed service, the guaranteed service data fragment is updated:

- If the data block in the ADSPEC is an empty (header-only) block the header-only fragment must first be expanded into the complete data fragment described above, with initial values of Ctot, Dtot, Csum, and Dsum set to zero. An empty fragment can be recognized quickly by checking for a size field of zero. The value of the break bit in the header is preserved when the additional Guaranteed service data is added. The overall message length and the guaranteed-service data fragment size (field (b) in the pictures above) are changed to reflect the increased message length.

The values of Ctot, Csum, Dtot, and Dsum in the ADSPEC data

fragment are then composed with the local values exported by the network element according to the composition functions defined in [RFCG].

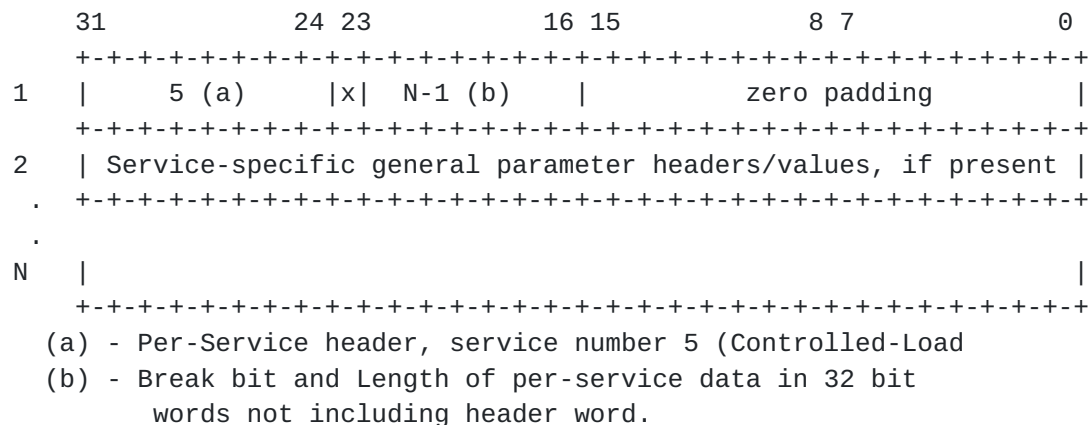
- When a PATH message ADSPEC with a Guaranteed service header encounters a network element that supports RSVP but does **not** implement Guaranteed service, the network element sets the break bit in the Guaranteed service header.
- The new values are placed in the correct fields of the ADSPEC, and the ADSPEC is passed back to RSVP for delivery to the next hop along the path.

When a PATH message ADSPEC containing a Guaranteed service data fragment encounters a network element that supports RSVP but does *not* implement Guaranteed service, the network element sets the break bit in the Guaranteed service header.

When a PATH message ADSPEC **without** a Guaranteed service header encounters a network element implementing Guaranteed service, the Guaranteed service module of the network element leaves the ADSPEC unchanged. The absence of a Guaranteed service per-service header in the ADSPEC indicates that the application does not care about Guaranteed service.

3.3.4. Controlled-Load Service ADSPEC data fragment

Unlike the Guaranteed service, the Controlled-Load service does not require extra ADSPEC data to function correctly. The only ADSPEC data specific to the Controlled-Load service is the Controlled-Load break bit. Therefore the usual Controlled-Load service data block contains no extra information. The minimum size of the controlled-load service data fragment is 1 32-bit word.



The Controlled-Load portion of the ADSPEC is processed according to

the following rules:

- When a PATH message ADSPEC with a Controlled-Load service header encounters a network element implementing Controlled-Load service, the network element makes no changes to the service header.
- When a PATH message ADSPEC with a Controlled-Load service header encounters a network element that supports RSVP but does **not** implement Controlled-Load service, the network element sets the break bit in the Controlled-Load service header.
- In either case, the ADSPEC is passed back to RSVP for delivery to the next hop along the path.

3.3.5. Overriding Global ADSPEC Data with Service-Specific Information

In some cases, the default values for the general parameters are not correct for a particular service. For example, an implementation of Guaranteed service may accept only packets with a smaller maximum size than the link MTU, or the percentage of outgoing link bandwidth made available to the Controlled-Load service at a network element may be administratively limited to less than the overall bandwidth.

In these cases, a service-specific value, as well as the default value, is reported to the receiver receiving the ADSPEC. Service-specific information which overrides general information is carried by a parameter with the same name as the general parameter, placed within the data fragment of the QoS control service to which it applies. These service-specific values are referred to as override or service-specific general parameters.

For example, the following Controlled-Load ADSPEC fragment carries information overriding the global path bandwidth estimate with a different value:

```

      31          24 23          16 15          8 7          0
      +-+-+-+-+-+-+-+-+
1  |      5 (a)      |x|      1 (b)      |      6 (c)      |      0 (d)      |
      +-+-+-+-+-+-+-+-+
2  | Path b/w estimate for C-L service (32b IEEE FP number)      |
      +-+-+-+-+-+-+-+-+
```

- (a) - Per-Service header, service number 5 (Controlled-Load)
- (b) - Break bit and Length of per-service data,
one word not including header word.
- (c) - Parameter ID, parameter 6 (Path-BW param from [\[RFCGP\]](#))
- (d) - Parameter 6 flags (none set)

The presence of override parameters in a data fragment can be quickly detected by examining the fragment's length field, which will be larger than the "standard" length for the fragment. Specific override parameters can be easily identified by examining the parameter headers, because they have parameter_number's from the general parameter portion of the number space (1-127), but are found in service-specific data blocks (those with service_numbers between 2 and 254 in the per_service header field).

The presence of override parameters in a data fragment is optional. A parameter header/value pair is added only when a particular application or QoS control service wishes to override the global value of a general parameter with a service-specific value.

As with IP options, it is only the use of these override parameters that is optional. All implementations must be prepared to receive and process override parameters.

The basic principle for handling override parameters is to use the override value (local or adspec) if it exists, and to use the default value otherwise. If a local node exports an override value for a general parameter, but there is no override value in the arriving adspec, the local node adds it. The following pseudo-code fragment gives more detail:

```
/* Adspec parameter processing rules */

<get arriving ADSPEC from RSVP>

for ( <each service number N with a fragment in the ADSPEC> ) {
  if ( <the local node does not support the service> ) {
    <set the break bit in the service header>
  } else {
    for ( <each parameter in the data fragment for service N> ) {
      if ( < the local service N supplies a value for the parameter> ) {
        <compose the arriving and values and update the adspec>
      } else {
        /* Must be a general parameter, or service N would have
         * supplied a value..
         */
        <compose the arriving value with the local default value
        and update the adspec>
      }
    }
  }
  for ( <any parameters supplied by the local service N implementation
    but not found in the adspec> ) {
    /*
     * Must be an override value for a general parameter,
```



```

    * or the adspec would have contained a value..
    */
    <compose the local override value with the arriving default
      value (from the service 1 data fragment) and add the parameter
      to the adspec's service N fragment in parameter_number order>
  }
}
}

<pass updated ADSPEC back to RSVP>

```

In practice, the two for loops can be combined. Since override parameters within a service's fragment are transmitted in numerical order, it is possible to determine whether a parameter is present without scanning the entire fragment. Also, because the data fragments are ordered by service_number, the default values for general parameters will always be read before they might be needed to update local override values in the second for loop.

3.3.6. Example

The picture below shows the complete adspec for an application which can use either controlled-load or guaranteed service. In the example, data fragments are present for general parameters, guaranteed, and controlled-load services. All fragments are of standard size, and there are no override parameters present.

	31	24	23	16	15	8	7	0
1	0 (a)		Unused				15 (b)	
2	1 (c)		x 5 (d)		4 (e)		i (f)	
3	IS hop cnt (16-bit integer)				6 (g)		i (h)	
4	Path b/w estimate (32-bit IEEE floating point number)							
5	8 (i)		i (j)		zero pad			
6	Minimum path latency (32-bit integer)							
7	10 (k)		i (l)		composed MTU (16 bits)			
8	2 (m)		x 7 (n)		133 (o)		0 (p)	
9	End-to-end composed value for C [Ctot] (32-bit integer)							
10	134 (q)		0 (r)		zero pad			

11	End-to-end composed value for D [Dtot] (32-bit integer)									
12	135 (s)	0 (t)	zero pad							
13	Since-last-reshaping point composed C [Csum] (32-bit integer)									
14	136 (u)	0 (v)	zero pad							
15	Since-last-reshaping point composed D [Dsum] (32-bit integer)									
16	5 (w)	x 0 (x)	zero padding							

Word 1: Message Header:

- (a) - Message header and version number
- (b) - Message length - 15 words not including header

Words 2-7: Default general characterization parameters

- (c) - Per-Service header, service number 1 (Default General Parameters)
- (d) - Global Break bit ([\[RFCGP\]](#), Parameter 2) (marked x) and length of

General Parameters data block (5 words)

- (e) - Parameter ID, parameter 4 (Number-of-IS-hops param from [RFCGP])
- (f) - Parameter 4 flag byte (bit i = INVALID, others unassigned)
- (g) - Parameter ID, parameter 6 (Path-BW param from [RFCGP])
- (h) - Parameter 6 flag byte (bit i = INVALID, others unassigned)
- (i) - Parameter ID, parameter 8 (minimum path latency from [RFCGP])
- (j) - Parameter 8 flag byte (bit i = INVALID, others unassigned)
- (k) - Parameter ID, parameter 10 (composed path MTU from [RFCGP])
- (l) - Parameter 10 flag byte (bit i = INVALID, others unassigned)

Words 8-15: Guaranteed service parameters

- (m) - Per-Service header, service number 2 (Guaranteed)
- (n) - Break bit and Length of per-service data in 32-bit words not including header word (7 words)
- (o) - Parameter ID, parameter 133 (Composed Ctot)
- (p) - Composed Ctot flags (none set).
- (q) - Parameter ID, parameter 134 (Composed Dtot)
- (r) - Composed Dtot flags (none set).
- (s) - Parameter ID, parameter 135 (Composed Csum).
- (t) - Composed Csum flags (none set).
- (u) - Parameter ID, parameter 136 (Composed Dsum).
- (v) - Composed Dsum flags (none set).

Word 16: Controlled-Load parameters

- (w) - Per-Service header, service number 5 (Controlled-Load)
- (x) - Break bit and Length of per-service data in 32-bit words not including header word (0 words)

4. Security Considerations

Security considerations are not discussed in this memo.

5. References

[RFCRSVP] B. Braden, et. al. "Resource Reservation Protocol (RSVP) - Version 1 Functional Specification", Internet Draft, July 1996, <[draft-ietf-rsvp-spec-13.txt](#)>

[RFCTEMPLATE] S. Shenker and J. Wroclawski. "Network Element QoS Control Service Specification Template". Internet Draft, July 1996, <[draft-ietf-intserv-svc-template-03.txt](#)>

[RFCDE] J. Wroclawski. "Data Element Naming and Encoding for Integrated Services Messages", Internet Draft, July 1996, <[draft-ietf-intserv-data-encoding-02.txt](#)>

[RFCG] S. Shenker, C. Partridge, and R Guerin. "Specification of Guaranteed Quality of Service", Internet Draft, July 1996, <[draft-ietf-intserv-guaranteed-svc-05.txt](#)>

[RFCCL] J. Wroclawski. "Specification of the Controlled Load Quality of Service", Internet Draft, July 1996, <[draft-ietf-intserv-ctrl-load-svc-03.txt](#)>

[RFCGP] S. Shenker and J. Wroclawski. "General Characterization Parameters for Integrated Service Network Elements", Internet Draft, July 1996, <[draft-ietf-intserv-charac-02.txt](#)>

Author's Address:

John Wroclawski
MIT Laboratory for Computer Science
545 Technology Sq.
Cambridge, MA 02139
jtw@lcs.mit.edu
617-253-7885
617-253-2673 (FAX)

