

Network Working Group  
Request For Comments: 2209  
Category: Informational

R. Braden  
ISI  
L. Zhang  
UCLA  
September 1997

Resource ReSerVation Protocol (RSVP) --  
Version 1 Message Processing Rules

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This memo contains an algorithmic description of the rules used by an RSVP implementation for processing messages. It is intended to clarify the version 1 RSVP protocol specification.

This memo provides a generic description of the rules for the operation of Version 1 of RSVP [RFC 2205]. It is intended to outline a set of algorithms that will accomplish the needed function, omitting many details.

1. GENERIC DATA STRUCTURES

This memo assumes the generic interface calls defined in [RFC 2005] and the following data structures. An actual implementation may use additional or different data structures and interfaces. The data structure fields that are shown are required unless they are explicitly labelled as optional.

o PSB -- Path State Block

Each PSB holds path state for a particular (session, sender) pair, defined by SESSION and SENDER\_TEMPLATE objects, respectively, received in a PATH message.

PSB contents include the following values from a PATH message:

- Session
- Sender\_Template
- Sender\_Tspec
- The previous hop IP address and the Logical Interface Handle (LIH) from a PHOP object
- The remaining IP TTL
- POLICY\_DATA and/or ADSPEC objects (optional)
- Non\_RSVP flag
- E\_Police flag
- Local\_Only flag

In addition, the PSB contains the following information provided by routing: OutInterface\_list, which is the list of outgoing interfaces for this (sender, destination), and IncInterface, which is the expected incoming interface. For a unicast destination, OutInterface\_list contains one entry and IncInterface is undefined.

Note that there may be more than one PSB for the same (session, sender) pair but different incoming interfaces. At most one of these, which will have the Local\_Only flag off, will be the PSB used for forwarding PATH messages downstream; we will refer to it as the "forwarding PSB" in the following. The other PSB's will have the Local\_Only flag on and an empty OutInterface\_list. The Local\_Only flag is needed to correctly match PSB's against RSB's, by the rules of [RFC 2205].

o RSB -- Reservation State Block

Each RSB holds a reservation request that arrived in a particular RESV message, corresponding to the triple: (session, next hop, Filter\_spec\_list). Here "Filter\_spec\_list" may be a list of FILTER\_SPECS (for SE style), a single FILTER\_SPEC (FF style), or empty (WF style). We define the virtual object type "FILTER\_SPEC\*" for such a data structure.

RSB contents include:

- Session specification
- Next hop IP address
- Filter\_spec\_list
- The outgoing (logical) interface OI on which the reservation is to be made or has been made.
- Style
- Flowspec
- A SCOPE object (optional, depending upon style)
- RESV\_CONFIRM object that was received (optional)

o TCSB -- Traffic Control State Block

Each TCSB holds the reservation specification that has been handed to traffic control for a specific outgoing interface. In general, TCSB information is derived from RSB's for the same outgoing interface. Each TCSB defines a single reservation for a particular triple: (session, OI, Filter\_spec\_list). TCSB contents include:

- Session
- OI (Outgoing Interface)
- Filter\_spec\_list
- TC\_Flowspec, the effective flowspec, i.e., the LUB over the corresponding FLOWSPEC values from matching RSB's. TC\_Flowspec is passed to traffic control to make the actual reservation.
- Fwd\_Flowspec, the updated object to be forwarded after merging.
- TC\_Tspec, equal to Path\_Te, the effective sender Tspec.
- Police Flags

The flags are E\_Police\_Flag, M\_Police\_Flag, and B\_Police\_Flag.

- Rhandle, F\_Handle\_list

Handles returned by the traffic control interface, corresponding to a flowspec and perhaps a list of filter specs.

- A RESV\_CONFIRM object to be forwarded.

- o BSB -- Blockade State Block

Each BSB contains an element of blockade state. Depending upon the reservation style in use, the BSB's may be per (session, sender\_template) pair or per (session, PHOP) pair. In practice, an implementation might embed a BSB within a PSB; however, for clarity we describe BSB's independently.

The contents of a BSB include:

- Session
- Sender\_Template (which is also a filter spec)
- PHOP
- FLOWSPEC Qb
- Blockade timer Tb

The following Boolean Flag variables are used in this section: Path\_Refresh\_Needed, Resv\_Refresh\_Needed, Tear\_Needed, Need\_Scope, B\_Merge, and NeworMod. Refresh\_PHOP\_list is a variable-length list of PHOPs to be refreshed.

## 2. PROCESSING RULES

### MESSAGE ARRIVES

Verify version number and RSVP checksum, and discard message if any mismatch is found.

If the message type is not PATH or PTEAR or RACK and if the IP destination address does not match any of the addresses of the local interfaces, then forward the message to IP destination address and return.

Parse the sequence of objects in the message. If any required objects are missing or the length field of the common header does not match an object boundary, discard the message and return.

Verify the INTEGRITY object, if any. If the check fails, discard the message and return.

Verify the consistent use of port fields. If the DstPort in the SESSION object is zero but the SrcPort in a SENDER\_TEMPLATE or FILTER\_SPEC object is non-zero, then the message has a "conflicting source port" error; silently discard the message and return.

Processing of POLICY\_DATA objects will be specified in the future.

Further processing depends upon message type.

#### PATH MESSAGE ARRIVES

Assume the PATH message arrives on interface InIf.

Process the sender descriptor object sequence in the message as follows. The Path\_Refresh\_Needed and Resv\_Refresh\_Needed flags are initially off.

- o Search for a path state block (PSB) whose (session, sender\_template) pair matches the corresponding objects in the message, and whose IncInterface matches InIf.

During this search:

1. If a PSB is found whose session matches the DestAddress and Protocol Id fields of the received SESSION object, but the DstPorts differ and one is zero, then build and send a "Conflicting Dst Port" PERR message, drop the PATH message, and return.
  2. If a PSB is found with a matching sender host but the Src Ports differ and one of the SrcPorts is zero, then build and send an "Ambiguous Path" PERR message, drop the PATH message, and return.
  3. If a forwarding PSB is found, i.e., a PSB that matches the (session, sender\_template) pair and whose Local\_Only flag is off, save a pointer to it in the variable fPSB. If none is found, set fPSB to NULL.
- o If there was no matching PSB, then:
    1. Create a new PSB.
    2. Copy contents of the SESSION, SENDER\_TEMPLATE, SENDER\_TSPEC, and PHOP (IP address and LIH) objects

into the PSB.

3. If the sender is from the local API, set OutInterface\_List to the single interface whose address matches the sender address, and make IncInterface undefined. Otherwise, turn on the Local\_Only flag.
  4. Turn on the Path\_Refresh\_Needed flag.
- o Otherwise (there is a matching PSB):
    - If the PHOP IP address, the LIH, or Sender\_Tspec differs between the message and the PSB, copy the new value into the PSB and turn on the Path\_Refresh\_Needed flag. If the PHOP IP address or the LIH differ, also turn on the Resv\_Refresh\_Needed flag.
  - o Call the resulting PSB the "current PSB" (cPSB). Update the cPSB, as follows:
    - Start or Restart the cleanup timer for the PSB.
    - If the message contains an ADSPEC object, copy it into the PSB.
    - Copy E\_Police flag from SESSION object into PSB.
    - Store the received TTL into the PSB. If the received TTL differs from Send\_TTL in the RSVP common header, set the Non\_RSVP flag on in the PSB.
  - o If the PSB is new or if there is no route change notification in place, then perform the following routing manipulations, but not if the cPSB is from the local API.
    1. Invoke the appropriate Route\_Query routine using DestAddress from SESSION and (for multicast routing) SrcAddress from Sender\_Template.  
  
Call the results (Rt\_OutL, Rt\_InIf).
    2. If the destination is multicast and Rt\_InIf differs from IncInterface in the cPSB, but fPSB points to the cPSB, then do the following.
      - Turn on the Local\_Only flag and clear the OutInterface\_list of the fPSB. Set the fPSB

pointer to NULL.

- Search for a PSB for the same (session, sender\_template) pair whose IncInterface matches Rt\_InIf. If one is found, set FPSB to point to it.
3. If the destination is multicast and Rt\_InIf is the same as IncInterface in the cPSB, but FPSB does not point to the cPSB, then do the following.
    - Copy into the cPSB the OutInterface\_list from the PSB, if any, pointed to by FPSB. Clear OutInterface\_list and turn on the Local\_Only flag in the PSB pointed to by FPSB, if any.
    - Turn off the Local\_Only flag in the cPSB and set FPSB to point to cPSB.
  4. If Rt\_OutL differs from OutInterface\_list of the PSB pointed to by FPSB, then:
    - Update the OutInterface\_list of the PSB from Rt\_OutL, and then execute the PATH LOCAL REPAIR event sequence below.
- o If the Path\_Refresh\_Needed flag is now off, drop the PATH message and return.

Otherwise (the path state is new or modified), do refreshes, upcalls, and state updates as follows.

1. If this PATH message came from a network interface and not from a local application, make a Path Event upcall for each local application for this session:
 

```
Call: <Upcall_Proc>( session-id, PATH_EVENT,
                      flags, sender_tspec, sender_template
                      [ , ADSPEC] [ , POLICY_DATA] )
```
2. If OutInterface\_list is not empty, execute the PATH REFRESH event sequence (below) for the sender defined by the PSB.
3. Search for any matching reservation state, i.e., an RSB whose Filter\_spec\_list includes a FILTER\_SPEC matching the SENDER\_TEMPLATE and whose OI appears in the OutInterface\_list, and make this the 'active RSB'.

If none is found, drop the PATH message and return.

4. Execute the RESV REFRESH sequence (below) for the PHOP in the PSB.
  5. Execute the event sequence UPDATE TRAFFIC CONTROL to update the local traffic control state if necessary. This sequence will turn on the Resv\_Refresh\_Needed flag if the traffic control state has been modified in a manner that should trigger a reservation refresh. If so, execute the RESV REFRESH sequence for the PHOP in the PSB.
- o Drop the PATH message and return.

#### PTEAR MESSAGE ARRIVES

- o Search for a PSB whose (Session, Sender\_Template) pair matches the corresponding objects in the message. If no matching PSB is found, drop the PTEAR message and return.
- o Forward a copy of the PTEAR message to each outgoing interface listed in OutInterface\_list of the PSB.
- o Find each RSB that matches this PSB, i.e., whose Filter\_spec\_list matches Sender\_Template in the PSB and whose OI is included in OutInterface\_list.
  1. If the RSB style is explicit, then:
    - Delete from Filter\_spec\_list the FILTER\_SPEC that matches the PSB.
    - if Filter\_spec\_list is now empty, delete the RSB.
  2. Otherwise (RSB style is wildcard) then:
    - If this RSB matches no other PSB, then delete the RSB.
  3. If an RSB was found, execute the event sequence UPDATE TRAFFIC CONTROL (below) to update the traffic control state to be consistent with the current reservation and path state.
- o Delete the PSB.
- o Drop the PTEAR message and return.



## PERR MESSAGE ARRIVES

- o Search for a PSB whose (SESSION, SENDER\_TEMPLATE) pair matches the corresponding objects in the message. If no matching PSB is found, drop the PERR message and return.
- o If the previous hop address in the PSB is the local API, make an error upcall to the application:

```
Call: <Upcall_Proc>( session-id, PATH_ERROR,
                    Error_code, Error_value, Node_Addr,
                    Sender_Template [ , Policy_Data] )
```

Any SENDER\_TSPEC or ADSPEC object in the message is ignored.

Otherwise, send a copy of the PERR message to the PHOP IP address.

- o Drop the PERR message and return.

## RESV MESSAGE ARRIVES

Initially, Refresh\_PHOP\_list is empty and the Resv\_Refresh\_Needed and NeworMod flags are off. These variables are used to control immediate reservation refreshes.

- o Determine the Outgoing Interface OI

The logical outgoing interface OI is taken from the LIH in the NHOP object. (If the physical interface is not implied by the LIH, it can be learned from the interface matching the IP destination address).

- o Check the path state
  1. If there are no existing PSB's for SESSION then build and send a RERR message (as described later) specifying "No path information", drop the RESV message, and return.
  2. If a PSB is found with a matching sender host but the SrcPorts differ and one of the SrcPorts is zero, then build and send an "Ambiguous Path" PERR message, drop the RESV message, and return.

- o Check for incompatible styles.

If any existing RSB for the session has a style that is incompatible with the style of the message, build and send a RERR message specifying "Conflicting Style", drop the RESV message, and return.

Process the flow descriptor list to make reservations, as follows, depending upon the style. The following uses a filter spec list struct Filtss of type FILTER\_SPEC\* (defined earlier).

For FF style: execute the following steps independently for each flow descriptor in the message, i.e., for each (FLOWSPEC, Filtss) pair. Here the structure Filtss consists of the FILTER\_SPEC from the flow descriptor.

For SE style, execute the following steps once for (FLOWSPEC, Filtss), with Filtss consisting of the list of FILTER\_SPEC objects from the flow descriptor.

For WF style, execute the following steps once for (FLOWSPEC, Filtss), with Filtss an empty list.

- o Check the path state, as follows.

1. Locate the set of PSBs (senders) that route to OI and whose SENDER\_TEMPLATES match a FILTER\_SPEC in Filtss.

If this set is empty, build and send an error message specifying "No sender information", and continue with the next flow descriptor in the RESV message.

2. If the style has explicit sender selection (e.g., FF or SE) and if any FILTER\_SPEC included in Filtss matches more than one PSB, build and send a RERR message specifying "Ambiguous filter spec" and continue with the next flow descriptor in the RESV message.
3. If the style is SE and if some FILTER\_SPEC included in Filtss matches no PSB, delete that FILTER\_SPEC from Filtss.
4. Add the PHOP from the PSB to Refresh\_PHOP\_list, if the PHOP is not already on the list.

- o Find or create a reservation state block (RSB) for (SESSION, NHOP). If the style is distinct, Filtss is also used in the selection. Call this the "active RSB".
- o If the active RSB is new:
  - 1. Set the session, NHOP, OI and style of the RSB from the message.
  - 2. Copy Filtss into the Filter\_spec\_list of the RSB.
  - 3. Copy the FLOWSPEC and any SCOPE object from the message into the RSB.
  - 4. Set NeworMod flag on.
- o If the active RSB is not new, check whether Filtss from the message contains FILTER\_SPECS that are not in the RSB; if so, add the new FILTER\_SPECS and turn on the NeworMod flag.
- o Start or restart the cleanup timer on the active RSB, or, in the case of SE style, on each FILTER\_SPEC of the RSB that also appears in Filtss.
- o If the active RSB is not new, check whether STYLE, FLOWSPEC or SCOPE objects have changed; if so, copy changed object into RSB and turn on the NeworMod flag.
- o If the message contained a RESV\_CONFIRM object, copy it into the RSB and turn on NeworMod flag.
- o If the NeworMod flag is off, continue with the next flow descriptor in the RESV message, if any.
- o Otherwise (the NeworMod flag is on, i.e., the active RSB is new or modified), execute the UPDATE TRAFFIC CONTROL event sequence (below). If the result is to modify the traffic control state, this sequence will turn on the Resv\_Refresh\_Needed flag and make a RESV\_EVENT upcall to any local application.

If the UPDATE TRAFFIC CONTROL sequence fails with an error, then delete a new RSB but restore the original reservation in an old RSB.
- o Continue with the next flow descriptor.

- o When all flow descriptors have been processed, check the Resv\_Refresh\_Needed flag. If it is now on, execute the RESV REFRESH sequence (below) for each PHOP in Refresh\_PHOP\_list.
- o Drop the RESV message and return.

If processing a RESV message finds an error, a RERR message is created containing flow descriptor and an ERRORS object. The Error Node field of the ERRORS object is set to the IP address of OI, and the message is sent unicast to NHOP.

#### RTEAR MESSAGE ARRIVES

Processing of a RTEAR message roughly parallels the processing of the corresponding RESV message

A RTEAR message arrives with an IP destination address matching outgoing interface OI. Flag Resv\_Refresh\_Needed is initially off and Refresh\_PHOP\_list is empty.

- o Determine the Outgoing Interface OI

The logical outgoing interface OI is taken from the LIH in the NHOP object. (If the physical interface is not implied by the LIH, it can be learned from the interface matching the IP destination address).

- o Process the flow descriptor list in the RTEAR message to tear down local reservation state, as follows, depending upon the style. The following uses a filter spec list struct Filtss of type FILTER\_SPEC\* (defined earlier).

For FF style: execute the following steps independently for each flow descriptor in the message, i.e., for each (FLOWSPEC, Filtss) pair. Here the structure Filtss consists of the FILTER\_SPEC from the flow descriptor.

For SE style, execute the following steps once for (FLOWSPEC, Filtss), with Filtss consisting of the list of FILTER\_SPEC objects from the flow descriptor.

For WF style, execute the following steps once for (FLOWSPEC, Filtss), with Filtss an empty list.

1. Find an RSB matching (SESSION, NHOP). If the style is distinct, Filtss is also used in the selection. Call this the "active RSB". If no active RSB is found, continue with next flow descriptor.
  2. Check the style  
  
If the active RSB has a style that is incompatible with the style of the message, drop the RTEAR message and return.
  3. Delete from the active RSB each FILTER\_SPEC that matches a FILTER\_SPEC in Filtss.
  4. If all FILTER\_SPECS have now been deleted from the active RSB, delete the active RSB.
  5. Execute the UPDATE TRAFFIC CONTROL event sequence (below) to update the traffic control state to be consistent with the reservation state. If the result is to modify the traffic control state, the Resv\_Refresh\_Needed flag will be turned on and a RESV\_EVENT upcall will be made to any local application.
  6. Continue with the next flow descriptor.
- o All flow descriptors have been processed.

Build and send any RTEAR messages to be forwarded, in the following manner.

1. Select each PSB that routes to the outgoing interface OI, and, for distinct style, that has a SENDER\_TEMPLATES matching Filtss.
2. Select a flow descriptor (Qj,Fj) (where Fj may be a list) in the RTEAR message whose FILTER\_SPEC matches the SENDER\_TEMPLATES in the PSB. If not match is found, return for next PSB.
  - Search for an RSB (for any outgoing interface) to which the PSB routes and whose Filter\_spec\_list includes the SENDER\_TEMPLATES from the PSB.
  - If an RSB is found, add the PHOP of the PSB to the Refresh\_PHOP\_list.

- Otherwise (no RSB is found), add the flow descriptor (Qj,Fj) to the new RTEAR message being built, in a manner appropriate to the style.
  - Continue with the next PSB.
3. If the next PSB is for a different PHOP or the last PSB has been processed, forward any RTEAR message that has been built.
- o If any PSB's were found in the preceding step, and if the Resv\_Refresh\_Needed flag is now on, execute the RESV REFRESH sequence (below) for each PHOP in Refresh\_PHOP\_list.
  - o Drop the RTEAR message and return.

#### RERR MESSAGE ARRIVES

A RERR message arrives through the (real) incoming interface In\_If.

- o If there is no path state for SESSION, drop the RERR message and return.
- o If the Error Code = 01 (Admission Control failure), do special processing as follows:
  - 1. Find or create a Blockade State Block (BSB), in the following style-dependent manner.

For WF (wildcard) style, there will be one BSB per (session, PHOP) pair.

For FF style, there will be one BSB per (session, filter\_spec) pair. Note that an FF style RERR message carries only one flow descriptor.

For SE style, there will be one BSB per (session, filter\_spec), for each filter\_spec contained in the filter spec list of the flow descriptor.
  - 2. For each BSB in the preceding step, set (or replace) its FLOWSPEC Qb with FLOWSPEC from the message, and set (or reset) its timer Tb to Kb\*R seconds. If the BSB is new, set its PHOP value, and set its Sender\_Template equal to the appropriate filter\_spec from the message.

3. Execute the RESV REFRESH event sequence (shown below) for the previous hop PHOP, but only with the B\_Merge flag off. That is, if processing in the RESV REFRESH sequence reaches the point of turning the B\_Merge flag on (because all matching reservations are blockaded), do not turn it on but instead exit the REFRESH sequence and return here.
- o Execute the following for each RSB for this session whose OI differs from In\_If and whose Filter\_spec\_list has at least one filter spec in common with the FILTER\_SPEC\* in the RERR message. For WF style, empty FILTER\_SPEC\* structures are assumed to match.
1. If Error\_Code = 01 and the InPlace flag in the ERROR\_SPEC is 1 and one or more of the BSB's found/created above has a Qb that is strictly greater than Flowspec in the RSB, then continue with the next matching RSB, if any.
  2. If NHOP in the RSB is the local API, then:
    - If the FLOWSPEC in the RERR message is strictly greater than the RSB Flowspec, then turn on the NotGuilty flag in the ERROR\_SPEC.
    - Deliver an error upcall to application:
 

```
Call: <Upcall_Proc>( session-id, RESV_ERROR,
                      Error_code, Error_value,
                      Node_Addr, Error_flags,
                      Flowspec, Filter_Spec_List
                      [ , Policy_data] )
```

and continue with the next RSB.
  3. If the style has wildcard sender selection, use the SCOPE object SC.In from the RERR message to construct a SCOPE object SC.Out to be forwarded. SC.Out should contain those sender addresses that appeared in SC.In and that route to OI, as determined by scanning the PSB's. If SC.Out is empty, continue with the next RSB.
  4. Create a new RERR message containing the error flow descriptor and send to the NHOP address specified by the RSB. Include SC.Out if the style has wildcard sender selection.

5. Continue with the next RSB.

- o Drop the RERR message and return.

#### RESV CONFIRM ARRIVES

- o If the (unicast) IP address found in the RESV\_CONFIRM object in the RACK message matches an interface of the node, a confirmation upcall is made to the matching application:

```
Call: <Upcall_Proc>( session-id, RESV_CONFIRM,
                    Error_code, Error_value, Node_Addr,
                    LUB-Used, nlist, Flowspec,
                    Filter_Spec_List, NULL, NULL )
```

- o Otherwise, forward the RACK message to the IP address in its RESV\_CONFIRM object.

Drop the RACK message and return.

#### UPDATE TRAFFIC CONTROL

The sequence is invoked by many of the message arrival sequences to set or adjust the local traffic control state in accordance with the current reservation and path state. An implicit parameter of this sequence is the 'active' RSB.

If the result is to modify the traffic control state, this sequence notifies any matching local applications with a RESV\_EVENT upcall. If the state change is such that it should trigger immediate RESV refresh messages, it also turns on the Resv\_Refresh\_Needed flag.

- o Compute the traffic control parameters using the following steps.
  1. Initially the local flag Is\_Biggest is off.
  2. Consider the set of RSB's matching SESSION and OI from the active RSB. If the style of the active RSB is distinct, then the Filter\_spec\_list must also be matched.
    - Compute the effective kernel flowspec, TC\_Flowspec, as the LUB of the FLOWSPEC values in these RSB's.



- Compute the effective traffic control filter spec (list) `TC_Filter_Spec*` as the union of the `Filter_spec_lists` from these RSB's.
  - If the active RSB has a `FLOWSPEC` larger than all the others, turn on the `Is_Biggest` flag.
3. Scan all RSB's matching session and `Filtss`, for all OI. Set `TC_B_Police_flag` on if `TC_Flowspec` is smaller than, or incomparable to, any `FLOWSPEC` in those RSB's.
  4. Locate the set of PSBs (senders) whose `SENDER_TEMPLATES` match `Filter_spec_list` in the active RSB and whose `OutInterface_list` includes OI.
  5. Set `TC_E_Police_flag` on if any of these PSBs have their `E_Police` flag on. Set `TC_M_Police_flag` on if it is a shared style and there is more than one PSB in the set.
  6. Compute `Path_Te` as the sum of the `SENDER_TSPEC` objects in this set of PSBs.
- o Search for a TCSB matching `SESSION` and OI; for distinct style (FF), it must also match `Filter_spec_list`.  
  
If none is found, create a new TCSB.
  - o If TCSB is new:
    1. Store `TC_Flowspec`, `TC_Filter_Spec*`, `Path_Te`, and the police flags into TCSB.
    2. Turn the `Resv_Refresh_Needed` flag on and make the traffic control call:
 

```
TC_AddFlowspec( OI, TC_Flowspec,
                  Path_Te, police_flags)
                  -> Rhandle, Fwd_Flowspec
```
    3. If this call fails, build and send a RERR message specifying "Admission control failed" and with the `InPlace` flag off. Delete the TCSB, delete any `RESV_CONFIRM` object from the active RSB, and return.

4. Otherwise (call succeeds), record Rhandle and Fwd\_Flowspec in the TCSB. For each filter\_spec F in TC\_Filter\_Spec\*, call:

```
TC_AddFilter( OI, Rhandle, Session, F)
              -> Fhandle
```

and record the returned Fhandle in the TCSB.

- o Otherwise, if TCSB is not new but no effective kernel flowspec TC\_Flowspec was computed earlier, then:
  1. Turn on the Resv\_Refresh\_Needed flag.
  2. Call traffic control to delete the reservation:

```
TC_DelFlowspec( OI, Rhandle )
```
  3. Delete the TCSB and return.
- o Otherwise, if TCSB is not new but the TC\_Flowspec, Path\_Te, and/or police flags just computed differ from corresponding values in the TCSB, then:
  1. If the TC\_Flowspec and/or Path\_Te values differ, turn the Resv\_Refresh\_Needed flag on.
  2. Call traffic control to modify the reservation:

```
TC_ModFlowspec( OI, Rhandle, TC_Flowspec,
                  Path_Te, police_flags )
                  -> Fwd_Flowspec
```
  3. If this call fails, build and send a RERR message specifying "Admission control failed" and with the InPlace bit on. Delete any RESV\_CONFIRM object from the active RSB and return.
  4. Otherwise (the call succeeds), update the TCSB with the new values and save Fwd\_Flowspec in the TCSB.
- o If the TCSB is not new but the TC\_Filter\_Spec\* just computed differs from the FILTER\_SPEC\* in the TCSB, then:
  1. Make an appropriate set of TC\_DelFilter and TC\_AddFilter calls to transform the Filter\_spec\_list in the TCSB into the new TC\_Filter\_Spec\*.

2. Turn on the Resv\_Refresh\_Needed flag.
- o If the active RSB contains a RESV\_CONFIRM object, then:
    1. If the Is\_Biggest flag is on, move the RESV\_CONFIRM object into the TCSB and turn on the Resv\_Refresh\_Needed flag. (This will later cause the RESV\_REFRESH sequence to be invoked, which will either forward or return the RESV\_CONFIRM object, deleting it from the TCSB in either case).
    2. Otherwise, create and send a RACK message to the address in the RESV\_CONFIRM object. Include the RESV\_CONFIRM object in the RACK message. The RACK message should also include an ERROR\_SPEC object whose Error\_Node parameter is IP address of OI from the TCSB and that specifies "No Error".
  - o If the Resv\_Refresh\_Needed flag is on and the RSB is not from the API, make a RESV\_EVENT upcall to any matching application:

```
Call: <Upcall_Proc>( session-id, RESV_EVENT,
                    style, Flowspec, Filter_spec_list [ ,
                    POLICY_DATA] )
```

where Flowspec and Filter\_spec\_list come from the TCSB and the style comes from the active RSB.

- o Return to the event sequence that invoked this one.

#### PATH REFRESH

This sequence sends a path refresh for a particular sender, i.e., a PSB. This sequence may be entered by either the expiration of a refresh timer or directly as the result of the Path\_Refresh\_Needed flag being turned on during the processing of a received PATH message.

- o Insert TIME\_VALUES object into the PATH message being built. Compute the IP TTL for the PATH message as one less than the TTL value received in the message. However, if the result is zero, return without sending the PATH message.
- o Create a sender descriptor containing the SENDER\_TEMPLATE, SENDER\_TSPEC, and POLICY\_DATA objects, if present in the PSB, and pack it into the PATH message being built.

- o Send a copy of the PATH message to each interface OI in OutInterface\_list. Before sending each copy:
  1. If the PSB has the E\_Police flag on and if interface OI is not capable of policing, turn the E\_Police flag on in the PATH message being built.
  2. Pass the ADSPEC object and Non\_RSVP flag present in the PSB to the traffic control call TC\_Advertise. Insert the modified ADSPEC object that is returned into the PATH message being built.
  3. Insert into its PHOP object the interface address and the LIH for the interface.

#### RESV REFRESH

This sequence sends a reservation refresh towards a particular previous hop with IP address PH. This sequence may be entered by the expiration of a refresh timer, or invoked from the PATH MESSAGE ARRIVES, RESV MESSAGE ARRIVES, RTEAR MESSAGE ARRIVES, or RERR MESSAGE ARRIVES sequence.

In general, this sequence considers each of the PSB's with PHOP address PH. For a given PSB, it scans the TCSBs for matching reservations and merges the styles, FLOWSPECs and Filter\_spec\_list's appropriately. It then builds a RESV message and sends it to PH. The details depend upon the attributes of the style(s) included in the reservations.

Initially the Need\_Scope flag is off and the new\_SCOPE object is empty.

- o Create an output message containing INTEGRITY (if configured), SESSION, RSVP\_HOP, and TIME\_VALUES objects.
- o Determine the style for these reservations from the first RSB for the session, and move the STYLE object into the proto-message. (Note that the present set of styles are never themselves merged; if future styles can be merged, these rules will become more complex).
- o If style is wildcard and if there are PSB's from more than one PHOP and if the multicast routing protocol does not use shared trees, set the Need\_Scope flag on.

- o Select each sender PSB whose PHOP has address PH. Set the local flag B\_Merge off and execute the following steps.
  1. Select all TCSB's whose Filter\_spec\_list's match the SENDER\_TEMPLATE object in the PSB and whose OI appears in the OutInterface\_list of the PSB.
  2. If the PSB is from the API, then:
    - If TCSB contains a CONFIRM object, then create and send a RACK message containing the object and delete the CONFIRM object from the TCSB.
    - Continue with next PSB.
  3. If B\_Merge flag is off then ignore a blockaded TCSB, as follows.
    - Select BSB's that match this TCSB. If a selected BSB is expired, delete it. If any of the unexpired BSB's has a Qb that is not strictly larger than TC\_Flowspec, then continue processing with the next TCSB.

However, if steps 1 and 2 result in finding that all TCSB's matching this PSB are blockaded, then:

- If this RESV REFRESH sequence was invoked from RESV ERROR RECEIVED, then return to the latter.
  - Otherwise, turn on the B\_Merge flag and restart at step 1, immediately above.
4. Merge the flowspecs from this set of TCSB's, as follows:
    - If B\_Merge flag is off, compute the LUB over the flowspec objects. From each TCSB, use the Fwd\_Flowspec object if present, else use the normal Flowspec object.

While computing the LUB, check for a RESV\_CONFIRM object in each TCSB. If a RESV\_CONFIRM object is found:

- If the flowspec (Fwd\_Flowspec or Flowspec) in that TCSB is larger than all other (non-blockaded) flowspecs being compared, then save this RESV\_CONFIRM object for forwarding and delete from the TCSB.
  - Otherwise (the corresponding flowspec is not the largest), create and send a RACK message to the address in the RESV\_CONFIRM object. Include the RESV\_CONFIRM object in the RACK message. The RACK message should also include an ERROR\_SPEC object whose Error\_Node parameter is IP address of OI from the TCSB and specifying "No Error".
  - Delete the RESV\_CONFIRM object from the TCSB.
- Otherwise (B\_Merge flag is on), compute the GLB over the Flowspec objects of this set of TCSB's.

While computing the GLB, delete any RESV\_CONFIRM object object in any of these TCSB's.

5. (All matching TCSB's have been processed). The next step depends upon the style attributes.

Distinct reservation (FF) style

Use the Sender\_Template as the merged FILTER\_SPEC. Pack the merged (FLOWSPEC, FILTER\_SPEC, F\_POLICY\_DATA) triplet into the message as a flow descriptor.

Shared wildcard reservation (WF) style

There is no merged FILTER\_SPEC. Merge (compute the LUB of) the merged FLOWSPECS from the TCSB's, across all PSB's for PH.

## Shared distinct reservation (SE) style

Using the Sender\_Template as the merged FILTER\_SPEC, form the union of the FILTER\_SPECS obtained from the TCSB's. Merge (compute the LUB of) the merged FLOWSPECS from the TCSB's, across all PSB's for PH.

6. If the Need\_Scope flag is on and the sender specified by the PSB is not the local API:
  - Find each RSB that matches this PSB, i.e., whose Filter\_spec\_list matches Sender\_Template in the PSB and whose OI is included in OutInterface\_list.
  - If the RSB either has no SCOPE list or its SCOPE list includes the sender IP address from the PSB, insert the sender IP address into new\_SCOPE.
- o (All PSB's for PH have been processed). Finish the RESV message.
  1. If Need\_Scope flag is on but new\_SCOPE is empty, no RESV message should be sent; return. Otherwise, if Need\_Scope is on, move new\_SCOPE into the message.
  2. If a shared reservation style is being built, move the final merged FLOWSPEC object and filter spec list into the message.
  3. If a RESV\_CONFIRM object was saved earlier, move it into the new RESV message.
  4. Set the RSVP\_HOP object in the message to contain the IncInterface address through which it will be sent and the LIH from (one of) the PSB's.
- o Send the message to the address PH.

## ROUTE CHANGE NOTIFICATION

This sequence is triggered when routing sends a route change notification to RSVP.

- o Each PSB is located whose SESSION matches the destination address and whose SENDER\_TEMPLATE matches the source address (for multicast).

1. If the OutInterface\_list from the notification differs from that in the PSB, execute the PATH LOCAL REPAIR sequence.
2. If the IncInterface from the notification differs from that in the PSB, update the PSB.

#### PATH LOCAL REPAIR

The sequence is entered to effect local repair after a route change for a given PSB.

- o Wait for a delay time of W seconds.
- o Execute the PATH REFRESH event sequence (above) for the PSB.

#### References

- [Baker96] Baker, F., "RSVP Cryptographic Authentication", Work in Progress.
- [RFC 2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 FunctionalSpecification", RFC 2205, September 1997.
- [RFC 2207] Berger, L. and T. O'Malley, "RSVP Extensions for IPSEC IPv4 Data Flows", RFC 2207, September 1997.
- [RSVP93] Zhang, L., Deering, S., Estrin, D., Shenker, S., and D. Zappala, "RSVP: A New Resource ReSerVation Protocol", IEEE Network, September 1993.

#### Security Considerations

Processing the RSVP INTEGRITY object [Baker96] is only mentioned in this memo, because the processing rules are described here only in general terms. The RSVP support for IPSEC [RFC 2207] will imply modifications that have not yet been incorporated into these processing rules.



## Authors' Addresses

Bob Braden  
USC Information Sciences Institute  
4676 Admiralty Way  
Marina del Rey, CA 90292

Phone: (310) 822-1511  
EMail: Braden@ISI.EDU

Lixia Zhang  
UCLA Computer Science Department  
4531G Boelter Hall  
Los Angeles, CA 90095-1596 USA

Phone: 310-825-2695  
EMail: lixia@cs.ucla.edu

