Network Working Group
Internet Draft
Document: draft-ietf-rap-modify-sender-behavior-00.txt
Category: Standards Track
Expiration: December 2001

R. Santitoro
Nortel Networks

R. Pabbati
Y. Bernet
Microsoft

July 2001

RSVP ErrorValues Used to Modify Sender Behavior

draft-ietf-rap-modify-sender-behavior-00.txt

Status of this Memo

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [2].

1. Abstract

This draft defines several mechanisms by which network policies can use RSVP signaling to control the behavior of compliant sending applications. Specifically, two new error codes are defined for use in the RSVP Policy Data object [RFC 2752]. In addition, a new use of the DCLASS object [RFC 2996] is defined.

2. Introduction

The initial focus of RSVP was to offer enhanced service to quantitative, multimedia applications. Over the past few years, we have learned that this focus is inconsistent with the priorities of many network managers. Most network managers are primarily concerned with maintaining control of

their network resources to protect qualitative, mission critical traffic. Ironically, this usually means either disallowing or severely restricting the deployment of just the type of multimedia application that RSVP was

originally intended to serve. If it were possible to better control the behavior of these applications, they would become more deployable. RSVP can be used to do so.

From the perspective of a sender and a receiver of application traffic, the conventional usage of RSVP is as follows:

1. Application sends PATH message.

2. Receiver uses RESV message to request reservation of resources for transmitted traffic.

3. Network either dedicates resources to the transmitted traffic or not (makes an admission control decision).

4. The admission control decision is indicated to the *receiver*.

5. Sender sends traffic regardless of the admission control decision.

RSVP signaling focuses on allocating network resources rather than controlling the behavior of the sending application. (Certain applications may use out of band signaling between receiver and sender. This signaling can be used to convey the network's admission control decision from the receiver to the sender, in order to impact the behavior of the sender. However, there is no explicit mechanism by which network policies can use RSVP to control the behavior of the sender).

RSVP has been adapted to address pragmatic concerns. Its integration with DiffServ [RFC 2998] addresses scalability concerns. The definition of the null service [RFC 2997] applies it to the type of qualitative mission critical applications that network managers deem most important to protect. Finally, the specification of the DCLASS object [RFC 2996] provides a mechanism by which network policies can control the behavior of sending applications (by using RSVP signaling to tell the sending application or host which DSCP to use in marking its traffic).

The ErrorValues and the DCLASS usage proposed in this draft provide the ability for network policies to explicitly control the behavior of sending applications.

The first ErrorValue informs the sending application that it MUST NOT send the traffic described in its PATH message. The second ErrorValue

informs the sending application that prioritized resources are not available but that it may proceed to send with no resource guarantees. (The ErrorValues are intended to be included in PATH_ERR messages, in response to corresponding PATH messages).

Finally, the inclusion of a DCLASS object in a PATH_ERR message is also discussed. In RFC 2996, the DCLASS object is discussed primarily in the context of RESV messages that are returned to the sender when a reservation request is admitted in the network. By including a DCLASS object in PATH_ERR messages as described in this document, it is possible

to control the behavior of the sender when a reservation request is not admitted.

3. New ErrorValues for Policy Error Object

The Policy Error Object from the Policy Data Object (P-Type=1 or 2, A-Type=4, SubType=0) contains the ErrorValue field. The following two new ErrorValues are defined below.

ErrorValue	Description
7 DO_NOT_SEND	The network cannot accommodate the traffic described in the sender's PATH message. The application must not transmit this traffic.
6 NO_QOS_PROVIDED	The application may send the traffic described in the PATH message but the network will not offer any service assurances.

3.1. DO_NOT_SEND (ErrorValue=7)

This ErrorValue is used to instruct sending applications not to send the traffic described in the PATH messages for the corresponding session. Note that this ErrorValue does not preclude sending altogether. Rather, it precludes sending per the profile described in the PATH messages for the session. Compliant applications respond either by refraining from sending altogether, or by modifying their traffic profile (typically, to a less demanding profile). The new traffic profile will be reflected in subsequent PATH messages and is less likely to elicit the DO_NOT_SEND response from the network.

This ErrorValue enables network managers (via a policy management system) to limit the impact of certain applications on the network resources. It is important to distinguish this mechanism from alternate control mechanisms available to the network manager.

One alternative is to simply deny QoS to the sending application (as in rejecting RESV messages on the corresponding session). Although this approach prevents the traffic transmitted on the session from interfering with other prioritized traffic, it does not prevent it from consuming best-effort resources. Thus, the transmitted traffic will compete with all other best-effort applications.

Another alternative is available to the network manager in certain cases. If the network is capable, it may force the traffic transmitted on the session to a less than best-effort (LBE) queue. To do so, the network must identify the traffic on the session. It may do so either by extracting the classification information for the session from the corresponding RSVP messages, or by causing the sending host to mark the traffic with a DSCP corresponding to LBE service. (The latter approach uses the DCLASS object as described in section 4).

While forcing traffic to an LBE queue does protect best-effort traffic, it requires functionality that may or may not be available in different parts of the network. The DO_NOT_SEND ErrorValue makes it possible to deploy compliant applications on networks that do not support this functionality.

Finally, an often-used alternative is to simply refuse to allow the deployment of aggressive applications on the network.

3.2 NO_QOS_PROVIDED (ErrorValue=6)

This ErrorValue indicates to the sending application that it will receive prioritized service for the traffic described in the PATH message for the corresponding session. Unlike the DO_NOT_SEND ErrorValue, it does not preclude the application from sending. Rather, it warns the application that transmitted traffic will not be assured any particular QoS.

From the network perspective, this response is similar to rejecting RESV messages for the corresponding session. However, unlike RESV message rejection (which is not indicated to the sender and may or may not be indicated to the receiver), the NO_QOS_PROVIDED ErrorValue gives immediate and explicit indication to the sender.

The sender may respond to the NO_QOS_PROVIDED ErrorValue by either:

- Not sending traffic on the corresponding session,

- Proceeding to send the traffic described by the corresponding PATH message (with no QoS assurances) or

- Modifying its traffic profile (typically, to a less demanding profile). The new traffic profile will be reflected in subsequent PATH messages and is less likely to elicit the NO_QOS_PROVIDED response from the network.

4. Use of the DCLASS Object in PATH_ERR Messages

As discussed in section 3.1, it is often desirable to force certain traffic to an LBE queue in the network. To do so, PEPs must either store classification information to be used in identifying the traffic (typically in the form of a 5-tuple), or the traffic must be marked explicitly for LBE service. One way to mark traffic for LBE service is by marking the transmitted packets with an LBE DSCP. The DCLASS object [RFC 2996] can be used by policy management systems to tell senders the DSCP with which to mark their traffic flows. RFC 2996 focuses on the use of the DCLASS object in RSVP RESV messages.

However, senders only receive RESV messages if the network has admitted the RSVP request. If the network rejects the RSVP request, no RESV message will arrive at the sender and there is no mechanism by which to force the sender to mark the rejected traffic with a specific DSCP. In the absence of alternate mechanisms, rejected traffic is either sent with

the best-effort DSCP (DSCP=0) or is not sent at all (DO_NOT_SEND response). We therefore propose that the DCLASS object be used in PATH_ERR messages when it is necessary to mark traffic on a session for which the corresponding RSVP request was rejected. A DCLASS object in a PATH_ERR message can specify a DSCP that is interpreted by the network PEPs to correspond to an LBE service.

5. Policies and Policy Server Support

The mechanisms described are expected to be supported by policy servers (PDPs) that are COPS/RSVP [COPS-RSVP] conversant. The following examples illustrate the types of policies that may be authored.

5.1 Prevent Streaming Video App. from Compromising Best-Effort Services

A policy would be created in a PDP that controls PEPs in the affected part of the network where streaming video applications are to be blocked. The policy would apply to all PATH messages including the application ID [RFC 2872] corresponding to the streaming video application. It would respond to each such PATH message with a PATH_ERR message specifying the DO_NOT_SEND ErrorValue.

5.2 Allocate Prioritized Service to a Limited Volume of Streaming Video Application traffic while Preventing Excess Traffic from Compromising

Best-Effort Service

A policy would be created in a PDP that controls PEPs in the affected part of the network. The policy would admit RSVP resource requests including the application ID corresponding to the streaming video application, up to a maximum allowed bandwidth. Once the maximum bandwidth is reached, additional resource requests will be rejected (using the conventional RESV_ERR message). The network will preclude additional traffic by responding to the sender's PATH messages with a PATH_ERR message specifying the DO_NOT_SEND ErrorValue.

5.3 Allocate Prioritized Service to a Limited Volume of Streaming Audio Traffic while Forcing Excess Traffic to LBE Service

A policy would be created in a PDP that controls PEPs in the affected part of the network. The policy would admit RSVP resource requests including the application ID corresponding to the streaming audio application, up to a maximum bandwidth. Once the maximum bandwidth is reached, additional resource requests will be rejected (using the conventional RESV_ERR message). However, additional traffic will not be precluded but rather, relegated to an LBE service. PATH_ERR messages specifying the NO_QOS_PROVIDED ErrorValue and a DCLASS object (specifying a DSCP corresponding to LBE service) will be sent in response to PATH messages corresponding to the additional traffic. These will provide explicit and immediate notification to the sending application indicating that its traffic will not receive prioritized service and that it must be marked for LBE service.

6. Security Considerations

Security mechanisms defined in [RFC 2752] apply to this draft.

7. References

[RFC 2753] Yavatkar R., et. al. "A Framework for Policy-based Admission Control", RFC 2753, January 2000.

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

[RFC 2752] Yadav S., et. al. "Identity Representation for RSVP", RFC 2752, January 2000.

[RFC 2872] Bernet Y., Pabbati R. "Application and Sub Application Identity Policy Element for Use with RSVP", RFC 2872, June 2000.

[RFC 2996] Bernet Y. "Format of the RSVP DCLASS Object", RFC 2996, November 2000.

[COPS-RSVP] Herzog S., et. al. "COPS usage for RSVP", RFC 2749, January 2000.

[RFC 2997] Bernet Y., et. al. "Specification of the Null Service Type", RFC 2997, November 2000.

[RFC 2998] Bernet Y., et. al. "A Framework for Integrated Services Operation over DiffServ Networks", RFC 2998, November 2000.

8. Acknowledgements

The authors would like to thank Kwok-Ho Chan, Ron Pashby, Eric Edwards and Nabil Seddigh for their input into the creation of this document.

9. Author's Addresses

Ralph Santitoro
Nortel Networks
4100 Guardian Street
Simi Valley, CA 93063
Phone: 805-527-3024
Email: rsantito@nortelnetworks.com

Ramesh Pabbati
Microsoft
1 Microsoft Way
Redmond, WA 98054
Phone: 425-936-9438
Email: rameshpa@microsoft.com

Yoram Bernet

Microsoft
1 Microsoft Way
Redmond, WA 98054
Phone: 425-936-9568
Email: yoramb@microsoft.com