

Network Working Group
Internet Draft
[draft-ietf-rsvp-proxy-03.txt](#)
Expiration Date: September 2002

Silvano Gai
Dinesh G Dutt
Nitsan Elfassy
Cisco Systems Inc.
Yoram Bernet
Microsoft

March 2002

RSVP Proxy

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>. The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Gai, Dutt, Elfassy, Bernet

[Page 1]

RSVP Proxy

March 2002

Abstract

RSVP has been extended in several directions [[POLICY](#)], [[RSVP-APPID](#)],

[[DCLASS](#)], [[AGGRRSVP](#)], [[RSVPDIFF](#)]. These extensions have broadened the applicability of RSVP characterizing it as a signaling protocol usable both inside and outside the Integrated Services [[INTSERV](#)] model.

With the addition of the "Null Service Type" [[NULLSERV](#)], RSVP is also being adopted by mission critical applications that require some form of prioritized service, but cannot quantify their resource requirements. In cases where RSVP cannot travel end-to-end, these applications may still benefit from reservations that are not truly end-to-end, but that are 'proxied' by a network node on the data path between the sender and the receiver(s).

RSVP Receiver Proxy is an extension to the RSVP message processing (not to the protocol itself) in which an intermediate network node originates the Resv message on behalf of the receiver(s) identified by the Path message.

RSVP Sender Proxy involves generating a Path message based on some match criteria at a router. For example, a packet filter can be installed at a router and the action associated with a match in the filter could be to generate a Path message.

1. Introduction

Network administrators and application developers would sometimes like to provide QoS to a flow based on information such as:

- o the type of application to which the flow belongs
- o a specific transaction within an application to which the flow belongs
- o the user running the application to which the flow belongs

Typically, such flows belong to applications that cannot quantify their traffic characteristics.

Since the data packets themselves do not usually carry information such as application or user id, an alternative approach is to signal this information separately.

RSVP [[RFC2205](#)] is a well established, standard IETF protocol that is used by applications to signal their QoS requirements to the network and obtain feedback about the network's ability to provide the requested QoS. An existing RFC [[RSVP-APPID](#)] defines the objects that can be used to carry the application id/sub-id and user-id in an RSVP message. Also, ISSLL has defined a new service type called Null Service Type [[NULLSERV](#)] for use within the IntServ framework. This service is intended for applications whose QoS requirements are better left to the discretion of the network administrators.

However, RSVP as currently defined travels end-to-end i.e. from the sender to the receiver and back. For the applications discussed above, this end-to-end nature of RSVP is not always applicable. For example, it might be that the application has been modified only on the sender side to support RSVP; there is no use in forwarding this message to the receiver since it does not support RSVP. Another example is where RSVP is used only within an administrative domain and a provisioned core is used outside of this domain. In such situations, RSVP is beneficial only within the administrative domain it has been enabled in. An example of this situation is QoS in PacketCable[DQOS], where resource reservations are used only within the access portion of the network and the core of the network is provisioned.

RSVP Receiver Proxy is proposed to address such situations.

[2.](#) RSVP Receiver Proxy

RSVP Receiver Proxy is a functionality provided by a network device, such as a switch or a router, in which the network device originates the Resv message in response to an incoming Path message, on behalf of the receiver(s) identified by the Path message.

The generation of the Resv message is done under policy control. Policy control can be performed using policy that has either been locally specified or specified by a policy server via a protocol such as COPS for RSVP [[COPS-RSVP](#)].

The proxy functionality does not imply merely generating a single Resv message. Proxying the Resv involves installing state in the node doing the proxy i.e. the proxying node should act as if it had received a Resv from the true endpoint. This involves reserving resources (if required), sending periodic refreshes of the Resv message and tearing down the reservation if the Path is torn down.

Optionally, the network device can also be configured to classify the packets and mark them with an appropriate DSCP. The codepoint used to mark these packets can also be communicated to the sender of the Path message via the DCLASS[DCLASS] object carried in the proxy Resv message.

RSVP Receiver Proxy does not change the "on-the-wire" RSVP protocol. It entails only a modification in the processing of the RSVP messages.

RSVP Receiver Proxy can be used with all the service types - Controlled Load [[CLSVC](#)], Guaranteed Service[GUSVC] and Null Service -

V

H1 <---- R1

Hx: Host x
Ry: Router y

Figure 1: Possible Message Forwarding Behaviors in RSVP

In Figure 1, case A illustrates the normal RSVP message processing. The Path message is generated by H1, is destined to H2, and it gets to H2 from H1 via R1, R2, R3 and R4. The Resv message uses the

reverse of the path setup by the Path message and goes hop-by-hop from H2 to H1.

With RSVP Receiver Proxy (case B) the RSVP Path message is terminated by the router R1 acting as a proxy for H2.

A possible sequence of steps is:

- o An application on H1 indicates to the RSVP subsystem that it is a sender wishing to use RSVP. It might specify additional parameters such as traffic characteristics and application specific information.
- o This causes the RSVP subsystem on H1 to start transmitting RSVP Path messages in accordance with normal RSVP/SBM rules.
- o The first hop network device (R1) receives this message and applies policy control to decide how to process this message.
- o The policy specifies a decision to not forward the Path message, but instead to proxy a Resv on behalf of H2. Additionally, the policy could specify the list of objects that need to be sent in the Resv message. One such additional object is the DCLASS object. Further, the policy could specify a DSCP that the network device (R1) must mark the flow identified by the Path message.
- o On receiving the Resv message, if the DCLASS object is specified the message, H1 can mark the packets of the traffic flow signaled, according to the DSCP specified in the DCLASS object.

[3.](#) RSVP Sender Proxy

Just as a network device can proxy a Resv message on behalf of a receiver, it can also be made to proxy a Path message on behalf of a sender. However the trigger that determines when a network device must generate a proxy Path message is potentially outside the RSVP subsystem. One mechanism for example, would be to install filter entries in the network device such that if an incoming flow matched one of the filters, the device would start generating a proxy Path message. At this point, it could potentially contact a policy server or use local policy in determining the behavior and contents of the proxy Path message.

The device generating the Path message must correctly terminate the Resv, ResvTear and PathErr messages.

[4.](#) Where To Proxy

In the example described in [section 3](#), the Receiver Proxy functionality was placed in the network device that was the first hop

Gai, Dutt, Elfassy, Bernet

[Page 5]

RSVP Proxy

March 2002

from the sender of the Path message. This is one possibility, not a requirement. While designing a network, the following trade-offs should be considered:

- o In case of Receiver Proxy, proxying farther from the sender of the Path message enables additional downstream network elements to benefit from the information carried in the signaling messages, and to participate in the response. For example, if some receivers are located off low-bandwidth links and other receivers off high-bandwidth links, the QoS to be applied could be different for the different receivers.
- o The proxying might be done at the boundary of an access network and a core network as in the case of PacketCable.
- o In case of Receiver Proxy, proxying closer to the sender results in a lower the latency experienced by the sender between the generation of the Path message and the receipt of the Resv message. This lower latency might be desirable to some applications.

The network administrator must take into account such factors in deciding where to place the proxy.

[5.](#) Security Considerations

The security considerations related to proxying are similar to those raised with respect to RSVP ([section 2.8 in \[RFC2205\]](#)). Specifically, the main concern in using a proxy is to ensure that unauthorized nodes do not mount a denial of service attack or cause theft of service by generating proxy RSVP messages on behalf of either a receiver or a sender. The problem is addressed via router to router authentication and using the INTEGRITY object [[RFC 2747](#)] in RSVP messages. Security policy enforcement can further prevent such attacks.

[6.](#) Intellectual Property Considerations

The IETF is being notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

[7.](#) References

- [INTSERV] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview," June 1994.
- [RFC2205] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource Reservation Protocol (RSVP) Version 1 Functional Specification", [RFC 2205](#), September 1997.

Gai, Dutt, Elfassy, Bernet

[Page 6]

RSVP Proxy

March 2002

- [DIFFSERV] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," [RFC 2474](#), December 1998.
- [CLSVC] J. Wroclawski, "Specification of the Controlled-Load Network Element Service," [RFC 2211](#), September 1997.
- [GUSVC] S. Shenker, C. Partridge, R. Guerin, "Specification of Guaranteed Quality of Service," [RFC 2212](#), September 1997.
- [COPS-RSVP] J. Boyle, R. Cohen, D. Durham, S. Herzog, R. Rajan, A. Sastry, "COPS usage for RSVP," [RFC 2749](#), January 2000.
- [POLICY] Shai Herzog, "RSVP Extensions for Policy Control," [RFC 2750](#), January 2000.
- [RSVPDIFF] Y. Bernet, R. Yavatkar, et. al., "A Framework For Integrated Services Operation Over Diffserv Networks, " [RFC 2998](#), November 2000.

- [RSVP-APPID] Y. Bernet, R. Pabbati, "Application and Sub Application Identity Policy Element for Use with RSVP," [RFC 2872](#), June 2000.
- [AGGRRSVP] F. Baker, C. Iturralde, F. Le Faucheur, B. Davie, "Aggregation of RSVP for IP4 and IP6 Reservations," [RFC 3175](#), September 2001.
- [DCLASS] Y. Bernet, "Format of the RSVP DCLASS Object", [RFC 2996](#), November 2000.
- [NULLSERV] Y. Bernet, A. Smith, B. Davie, "Specification of the Null Service Type," [RFC 2997](#), November 2000.
- [DQOS] PacketCable Dynamic Quality Of Service Specification, Interim Version,
<http://www.packetcable.com/specs/pkt-sp-dqos-I01-991201.pdf>.
- [RFC2747] F. Baker, B. Lindell, M.Talwar, "RSVP Cryptographic Authentication", [RFC 2747](#), January 2000.

8. Author Information

Silvano Gai
Cisco Systems, Inc.
[170](#) Tasman Dr.
San Jose, CA 95134-1706
Phone: (408) 527-2690
email: sgai@cisco.com

Gai, Dutt, Elfassy, Bernet

[Page 7]

RSVP Proxy

March 2002

Dinesh G Dutt
Cisco Systems, Inc.
[170](#) Tasman Dr.
San Jose, CA 95134-1706
Phone: (408) 527-0955
email: ddutt@cisco.com

Nitsan Elfassy
Cisco Systems, Inc.
Cisco Systems, Inc.
[170](#) Tasman Dr.
San Jose, CA 95134-1706
Phone: +972 9 970 0066

email: nitsan@cisco.com

Yoram Bernet
Microsoft
One Microsoft Way,
Redmond, WA 98052
Phone: (425) 936-9568
Email: yoramb@microsoft.com

9. Full Copyright Statement

Copyright (C) The Internet Society (1997). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined

in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.