Network Working Group
Request for Comments: 2996
Category: Standards Track

Y. Bernet
Microsoft

November 2000

Format of the RSVP DCLASS Object

Status of this Memo

Copyright Notice

Abstract

Resource Reservation Protocol (RSVP) signaling may be used to request Quality of Service (QoS) services and enhance the manageability of application traffic's QoS in a differentiated service (diff-serv or DS) network. When using RSVP with DS networks it is useful to be able to carry carry Differentiated Services Code Points (DSCPs) in RSVP message objects. One example of this is the use of RSVP to arrange for the marking of packets with a particular DSCP upstream from the DS network's ingress point, at the sender or at a previous network's egress router.

The DCLASS object is used to represent and carry DSCPs within RSVP messages. This document specifies the format of the DCLASS object and briefly discusses its use.

1. Introduction

This section describes the mechanics of using RSVP [RSVP] signaling and the DCLASS object for effecting admission control and applying QoS policy within a Differentiated Service network [DS]. It assumes standard RSVP senders and receivers, and a diff-serv network somewhere in the path between sender and receiver. At least one RSVP aware network element resides in the diff-serv network. This network element may be a policy enforcement point (PEP) [RAP] or may simply act as an admission control agent for the network, admitting or denying resource requests based on the availability of resources. In either case, this network element interacts with RSVP messages arriving from outside the DS network, accepting resource requests

from RSVP-aware senders and receivers, and conveying the DS network's admission control and resource allocation decisions to the higher-level RSVP. The network element is typically a router and will be considered to be so for the purpose of this document. This model is described fully in [INTDIFF].

1.1 Use of the DCLASS Object to Carry Upstream Packet Marking Information

A principal usage of the DCLASS object is to carry DSCP information between a DS network and upstream nodes that may wish to mark packets with DSCP values. Briefly, the sender composes a standard RSVP PATH message and sends it towards the receiver. At some point the PATH message reaches the DS network. The PATH message traverses one or more network elements that are PEPs and/or admission control agents for the diff-serv network. These elements install appropriate state and forward the PATH message towards the receiver. If admission control is successful downstream of the diff-serv network, then a RESV message will arrive from the direction of the receiver. As this message arrives at the PEPs and/or admission control agents that are RSVP enabled, each of these network elements must make a decision regarding the admissibility of the signaled flow to the diff-serv network.

If the network element determines that the request represented by the PATH and RESV messages is admissible to the diff-serv network, the appropriate diff-serv service level (or behavior aggregate) for the traffic represented in the RSVP request is determined. Next, a decision is made to mark arriving data packets for this traffic locally using MF classification, or to request upstream marking of the packets with the appropriate DSCP(s). This upstream marking could occur anywhere before the DS network's ingress point. Two likely candidates are the originating sender and the egress boundary router of some upstream (DS or non-DS) network. The decision about where the RSVP request's packets should be marked can be made by agreement or through a negotiation protocol; the details are outside the scope of this document.

If the packets for this RSVP request are to be marked upstream, information about the DSCP(s) to use must be conveyed from the RSVP-aware network element to the upstream marking point. This information is conveyed with the DCLASS object. To do this, the network element adds a DCLASS object containing one or more DSCPs corresponding to the behavior aggregate, to the RESV message. The RESV message is then sent upstream towards the RSVP sender.

If the network element determines that the RSVP request is not admissible to the diff-serv network, it sends a RESV error message

towards the receiver. No DCLASS is required.

1.1 Additional Uses of the DCLASS Object

The DCLASS object is intended to be a general tool for conveying DSCP information in RSVP messages. This may be useful in a number of situations. We give one further example here as motivation.

In this example, we assume that the decision about the appropriate behavior aggregate for a RSVP-mediated traffic flow is made at the DS network egress router (or a related Policy Decision Point) by observing RSVP PATH and RESV messages and other necessary information. However, the actual packet marking must be done at the ingress of the network. The DCLASS object can be used to carry the needed marking information between egress and ingress routers.

2. Format of the DCLASS Object

The DCLASS object has the following format:

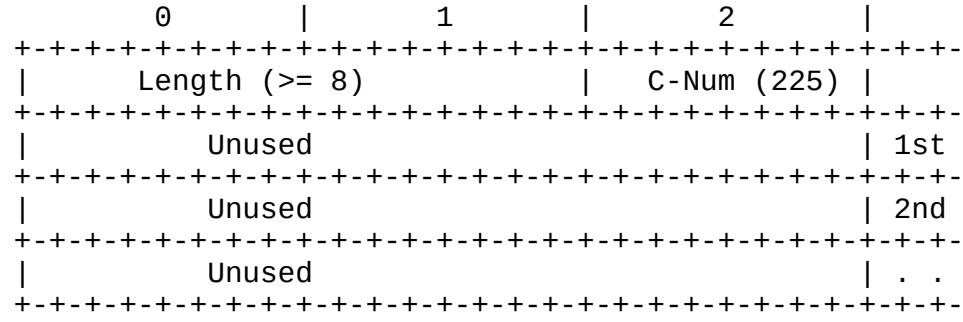

The first word contains the standard RSVP object header (the Class Num for the DCLASS object is 225). The length field indicates the total object length in bytes. The object header is followed by one or more 32-bit words, each containing a DSCP in the six high-order bits of the least significant byte. The length field in the object header indicates the number of DSCPs included in the object. Specifically, the number of DCLASS objects present is equal to (Length - 4) / 4.

The network may return multiple DSCPs in the DCLASS object in order to enable the host to discriminate sub-flows within a behavior aggregate. For example, in the case of the AF PHB group [AF], the network may return the DSCPs 001010, 001100, and 001110 corresponding to increasing levels of drop precedence within Class 1 of the AF PHB group. Note that this document makes no statements regarding the significance of the order of the returned DSCPs. Further interpretation of DSCP sets is dependent on the specific service

requested by the host and is beyond the scope of this document.

Note that the Class-Num for the DCLASS object is chosen from the space of unknown class objects that should be ignored and forwarded by nodes that do not recognize it. This is to assure maximal backward compatibility.

3. Admission Control Functionality

From a black-box perspective, admission control and policy functionality amounts to the decision whether to accept or reject a request and the determination of the DSCPs that should be used for the corresponding traffic. The specific details of admission control are beyond the scope of this document. In general the admission control decision is based both on resource availability and on policies regarding the use of resources in the diff-serv network. The admission control decision made by RSVP aware network elements represents both considerations.

In order to decide whether the RSVP request is admissible in terms of resource availability, one or more network elements within or at the boundary of the diff-serv network must understand the impact that admission would have on specific diff-serv resources, as well as the availability of these resources along the relevant data path in the diff-serv network.

In order to decide whether the RSVP request is admissible in terms of policy, the network element may use identity objects describing users and/or applications that may be included in the request. The router may act as a PEP/PDP and use data from a policy database or directory to aid in this decision.

See Appendix A for a simple mechanism for configurable resource based admission control.

4. Security Considerations

The DCLASS object conveys information that can be used to request enhanced QoS from a DS network, so inappropriate modification of the object could allow traffic flows to obtain a higher or lower level of QoS than appropriate. Particularly, modification of a DCLASS object by a third party inserted between the DS network ingress node and the upstream marker constitutes a possible denial of service attack. This attack is subtle because it is possible to reduce the received QoS to an unacceptably low level without completely cutting off data flow, making the attack harder to detect.

The possibility of raising the received level of QoS by inappropriate

modification of the DCLASS object is less significant because it a subclass of a larger class of attacks that must already be detected by the system. Protection must already be in place to prevent a host raising its received level of QoS by simply guessing "good" DSCP's and marking packets accordingly. If this protection is at the boundary of the DS network, it will detect inappropriate marking of arriving packets caused by modified DCLASS objects as well. If, however, the protection function as well as the marking function has been pushed upstream (perhaps to a trusted third party or intermediate node), correct transmission of the DCLASS object must be ensured to prevent a possible theft of service attack.

Simple observation of the DCLASS object in a RSVP message raises several issues which may be seen as security concerns. Correlation of observed DCLASS object values with RSVP requests or MF classification parameters allows the observer to determine that different flows are receiving different levels of QoS, which may be knowledge that should be protected in some environments. Similarly, observation of the DCLASS object can allow the observer to determine that a single flow's QoS has been promoted or demoted, which may signal significant events in the life of that flow's application or user. Finally, observation of the DCLASS object may reveal information about the internal operations of a DS network that could be useful to observers interested in theft-of-services attacks.

5. References

[INTDIFF] Bernet, Y., Yavatkar, R., Ford, P., Baker, F., Zhang, L., Speer, M., Braden, R., Davie, B. and J. Wroclawski, "A Framework for Integrated Services Operation over Diffserv Networks", RFC 2998, November 2000.

[DS] Blake, S., Carlson, M., Davies, D., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[RSVP] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[RAP] Yavatkar, R., Pendarakis, D. and R. Guerin, "A Framework for Policy Based Admission Control", RFC 2753, January 2000.

[AF] Heinanen, J., Baker, F., Weiss, W. and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.

6. Acknowledgments

Thanks to Fred Baker and Carol Iturralde for reviewing this document. Thanks to Ramesh Pabbati, Tim Moore, Bruce Davie and Kam Lee for input.

7. Author's Address

Yoram Bernet
Microsoft
One Microsoft Way,
Redmond, WA 98052

Phone: (425) 936-9568
EMail: yoramb@microsoft.com

Appendix A - Simple Configurable Resource Based Admission Control

Routers may use quite sophisticated mechanisms in making the admission control decision, including policy considerations, various intra-domain signaling protocols, results of traffic monitoring and so on. It is recommended that the following basic functionality be provided to enable simple resource based admission control in the absence of more sophisticated mechanisms. This functionality can be used with configurable, standalone routers. It applies to standard RSVP/Intserv requests. This minimal functionality assumes only a single DSCP is included in the DCLASS object, but may readily be extended to support multiple DSCPs.

It must be possible to configure two tables in the router. These are described below.

A.1 Service Type to DSCP Mapping

One table provides a mapping from the intserv service-type specified in the RSVP request to a DSCP that can be used to obtain a corresponding service in the diff-serv network. This table contains a row for each intserv service type for which a mapping is available. Each row has the following format:

Intserv service type : DSCP

The table would typically contain at least three rows; one for Guaranteed service, one for Controlled Load service and one for Best-Effort service. (The best-effort service will typically map to DSCP 000000, but may be overridden). It should be possible to add rows for as-yet-undefined service types.

This table allows the network administrator to statically configure a DSCP that the router will return in the DCLASS object for an admitted RSVP request. In general, more sophisticated and likely more dynamic mechanisms may be used to determine the DSCP to be returned in the DCLASS object. Also, it is likely that a real mapping for some services would use more than one DSCP, with the DSCP depending on the invocation parameters of a specific service request. In this case, these mechanisms may override or replace the static table based mapping described here.

A.2 Quantitative Resource Availability

Standard intserv requests are quantitative in nature. They include token bucket parameters describing the resources required by the traffic for which admission is requested. The second table enables the network administrator to statically configure quantitative

parameters to be used by the router when making an admission control decision for quantitative service requests. Each row in this table has the following form:

DSCP : Token bucket profile

The first column specifies those DSCPs for which quantitative admission control is applied. The second column specifies the token bucket parameters which represent the total resources available in the diff-serv network to accommodate traffic in the service class specified by the DSCP.

Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.