

Encoding of SRP packet types in the DS byte

Werner Almesberger <Werner.Almesberger@epfl.ch>¹,

Tiziana Ferrari <Tiziana.Ferrari@cnaif.infn.it>²,

Jean-Yves Le Boudec <Leboudec@epfl.ch>¹

¹EPFL ICA, CH-1015 Lausanne, Switzerland

²DEIS, University of Bologna, and INFN/CNAF, Italy

March 12, 1998

Abstract

We propose an encoding of the packet types used by SRP (Scalable Reservation Protocol) in the DS byte under study by the Differentiated Services working group.

1 Introduction

SRP [1] is a light-weight reservation protocol designed to provide dependable reservations for the Internet. SRP aggregates flows in routers to achieve good scalability to very large numbers of individual flows.

The semantics of the actual reservation protocol can be expressed on a per-packet and per-hop basis. We describe the desired behaviour of routers and define a possible encoding of the SRP packet types in the “TOS octet” or “DS byte” as defined in the context of the IETF Differentiated Services (diffserv) working group [2].

Although diffserv targets slightly different semantics when using the TOS octet, we believe this octet to be a highly appropriate location for indicating the use of other, packet-based mechanisms to control QoS. With this draft, we intend to advocate that the specification produced by the diffserv working group should allow for such uses, and that it should provide room for the corresponding code points.

We base our description on the initial draft proposal for using the TOS octet [3] and we acknowledge that the exact layout of the TOS octet is still subject to ongoing discussion and may change in the future. We plan to update this document accordingly.

Note that our description of possible uses of the DS byte only reflects our own design ideas and shall not be interpreted as implying any endorsement by the diffserv working group or by related groups.

2 Reservations

SRP sets up reservations for aggregated flows. The exact definition of what constitutes an aggregated flow is a local issue of each router. Simple definitions may be obtained by considering all flows as aggregates, which:

- leave the router at the same port, or
- traverse the router through the same pair of ports.

A reservation means that a router has set aside sufficient resources to process and forward conforming packets using the reservation, and to do so in timely way. A packet conforms to a reservation if

- it traverses the router through the same pair of ports as the packet(s) that have set up the reservation,
- the source has shaped it according to the shaping requirements for sources (TBD) and all routers have forwarded it according to rules specified in this document, and
- it does not request any other service(s) from the router (by whatever means the router supports) than the one(s) used for the packet(s) that have set up the reservation.

3 Packet types

We distinguish three types of packets:

- *reserved* packets use a previously agreed-upon reservation. Sources may emit *reserved* packets only if a reservation was granted and only according to the shaping requirements. After optionally performing policing, routers must include such packets in their estimate of current resource use, and forward them with priority over best-effort.
- flows of *request* packets have the effect of setting up a new reservation or increasing an existing one. This specification does not impose any restrictions on when or how often a source may send *request* packets. When receiving a *request* packet, a router must decide if it has sufficient resources to increase the present reservation. If yes, it updates its estimate accordingly and forwards *request* the packet like it forwards *reserved* packets. Otherwise, it degrades the packet to *denied* and processes it accordingly.
- *denied* packets are treated similar to best-effort, but a router may discard *denied* packets even with higher probability than normal best-effort, because *denied* packets will typically not be congestion-controlled by the source. SRP does not allow a source to emit packets of type *denied* (but, if the encoding of the *denied* packet type coincides with a packet type defined elsewhere, a source may of course emit such packets according to that other specification)

As a local option, routers may entrust other routers to make acceptance decisions for them. Doing so does not remove their obligation to forward *reserved* and *request* packets according to this specification.

Estimator algorithms are not specified in this document. We study examples for such algorithms in [4].

Destinations identify the types of incoming packets and generate feedback for the source. Further details of this are TBD.

4 Encoding of packet types

We propose two possible ways for encoding SRP packet types in the DS byte. Criteria for deciding which of the proposals is appropriate are discussed below.

4.1 Proposal 1

Assuming a PHB (per-hop behaviour) code point is assigned to SRP, we define the encoding of SRP packet types in values of the DS byte as follows:

SRP packet type	IN (in/out of profile)	PHB (per-hop behaviour)	CU (currently unused)
<i>reserved</i>	in(1)	SRP	—
<i>request</i>	out(0)	SRP	—
<i>denied</i>	out(0)	DE	—

Where DE is the default best-effort forwarding described in [3].

4.2 Proposal 2

Recent discussion in the diffserv group has indicated that perhaps only selected routers may be allowed to change the DS byte. Not allowing every router to change the DS byte would conflict with SRP's potential need to perform acceptance control at every router, and consequently to degrade some *request* packets to *denied*. If such restrictions should be put on the use of the DS byte, we propose using the ECN bit (Explicit Congestion Notification, [5]), which – much like SRP packet types – may be set at any router, and which carries similar information. Then, the encoding would be as follows:

SRP packet type	IN	PHB	ECN Capable	ECN
<i>reserved</i>	in(1)	SRP	TBD	TBD
<i>request</i>	out(0)	SRP	TBD	0
<i>denied</i>	out(0)	SRP	TBD	1

Note that the use of the ECN bits as described in this second proposal does not interfere with employing them for implementing congestion control on top of reservations (like it is done in Frame Relay [6]), if a need for this should arise.

Because the exact use of ECN and the relation between DS and ECN are still the subject of vivid discussion, this latter encoding proposal should be considered as even more volatile than the previous one.

References

- [1] Almesberger, Werner; Ferrari, Tiziana; Le Boudec, Jean-Yves. *SRP essentials* (work in progress), Internet Draft `draft-watfjyl-srp-00.txt`, March 1998.
- [2] IETF, Differentiated Services (diffserv) working group. <http://www.ietf.org/html.charters/diffserv-charter.html>
- [3] Nichols, Kathleen (Ed.); Blake, Steven (Ed.). *Differentiated Services Operational Model and Definitions* (work in progress), Internet Draft `draft-nichols-dsopdef-00.txt`, February 1998.
- [4] Almesberger, Werner; Ferrari, Tiziana; Le Boudec, Jean-Yves. *SRP: a Scalable Resource Reservation Protocol for the Internet*, <ftp://lrcftp.epfl.ch/pub/people/almesber/srp/srpMar98.ps.gz>, Technical Report SSC/1998/009, EPFL, March 1998.
- [5] Ramakrishnan, Kadangode K.; Floyd, Sally. *A Proposal to add Explicit Congestion Notification (ECN) to IPv6 and to TCP* (work in progress), Internet Draft `draft-kksjf-ecn-00.txt`, November 1997.
- [6] ITU-T Recommendation X.36. *Interface between data terminal equipment (DTE) and data circuit-terminating equipment (DCE) for public data networks providing frame relay data transmission service by dedicated circuit*, ITU, April 1995.

5 Author's address

Werner Almesberger
Jean-Yves Le Boudec
Institute for computer Communications and Applications
Swiss Federal Institute of Technology (EPFL)
CH-1015 Lausanne
Switzerland
email: {Werner.Amesberger,Leboudec}@epfl.ch

Tiziana Ferrari
DEIS, University of Bologna
viale Risorgimento, 2
I-40136 Bologna
Italy
and
Italian National Inst. for Nuclear Physics/CNAF
viale Berti Pichat, 6/2
I-40127 Bologna
Italy
email: Tiziana.Ferrari@cnaif.infn.it