

SRP essentials

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 11, 1998

Abstract

This paper gives a succinct description of the design of SRP, a highly scalable resource reservation protocol for Internet traffic.

1 About this paper

This paper is a short introduction to the “Scalable Reservation Protocol” (SRP). It aims to provide background information for discussions related to SRP and is, due to focussing only on the most essential issues, necessarily incomplete. Readers interested in all the gory details are kindly referred to [1].

This document is a summary of the current design of SRP. It is not just an update of `draft-almesberger-srp-00.txt`. We also tried to align it better with concepts and terminology currently used in the diffserv [2] working group of IETF.

The Web page of SRP is at <http://lrcwww.epfl.ch/srp/>

2 Overview

SRP provides a light-weight reservation mechanism for adaptive multimedia applications [3]. Our main focus is on good scalability to very large numbers of individual flows. End systems (i.e. senders and destinations) actively participate in maintaining reservations, but routers can still control their conformance. Routers aggregate flows and monitor the aggregate to estimate the resources needed to support present and new reservations. There is no explicit signaling of flow parameters.

3 End-to-end service

Many adaptive multimedia applications require a well-defined fraction of their traffic to reach the destination and to do so in a timely way. We call this fraction the *minimum rate* these applications need in order to operate properly.

SRP aims to allow such applications to make a dependable reservation of their minimum rate.

The sender can expect that, as long as it adheres to the agreed-upon profile, no *reserved* packets will be lost due to congestion. Furthermore, forwarding of *reserved* packets will have priority over best-effort traffic.

4 Reservation mechanism

A source that wishes to make a reservation starts by sending data packets marked as *request* packets to the destination. Packets marked as *request* are subject to packet admission control by routers, based on the following principle. Routers monitor the aggregate flows of *reserved* packets and maintain a running estimate of what level of resources is required to serve them with a good quality of service. The resources are bandwidth and buffer on outgoing links, plus any internal resources as required by the router architecture. Quality of service is loss ratio and delay, and is defined statically. When receiving a *request* packet, a router determines whether hypothetically adding this packet to the flow of *reserved* packets would yield an acceptable value of the estimator. If so, the *request* packet is accepted and forwarded towards the destination, while still keeping the status of a *request* packet; the router must also update the estimator as if the packet had been received as *reserved*. In the opposite case, the *request* packet is degraded and forwarded towards the destination, and the estimator is not updated. Degrading a *request* packet means assigning it a lower traffic class, such as best effort. A packet sent as *request* will reach the destination as *request* only if all routers along the path have accepted the packet as *request*. Note that the choice of an estimation method is local to a router and actual estimators may differ in their principle of operation.

The destination periodically sends feedback to the source indicating the rate at which *request* and *reserved* packets have been received. This feedback does not receive any special treatment in the network (except possibly for policing, see section 7). Upon reception of the feedback, the source can send packets marked as *reserved* according to a profile derived from the rate indicated in the feedback. If necessary, the source may continue to send more *request* packets in an attempt to increase the rate that will be indicated in subsequent feedbacks.

Thus, in essence, a router accepting to forward a *request* packet as *request* allows the source to send a *reserved* packet in the future; it is thus a form of implicit reservation.

5 Aggregation

Routers aggregate flows on output ports, and possibly on any contention point as required by their internal architecture. They use estimator algorithms for each aggregated flow to determine their current reservation levels and to predict the impact of accepting *request* packets. The exact definition of what constitutes an aggregated flow is local to a router.

Likewise, senders and sources treat all flows between each pair of them as a single aggregate and use estimator algorithms for characterizing them. The estimator algorithms in routers and hosts do not need to be the same. In fact, we

expect hosts to implement a fairly simple algorithm, while estimator algorithms in routers may evolve independently over time.

We evaluated example host and router algorithms in [1].

6 Further issues

Further issues, such as support for multicast, are discussed in [1].

7 Security

Denial-of-service conditions may arise if flows can reserve disproportional amounts of resources or if flows can exceed their reservations. We presently consider fairness in accepting reservations a local policy issue (much like billing) which may be addressed at a future time.

Sources violating the agreed upon reservations are a real threat and need to be policed. Policing can be efficiently implemented if policing points filter feedback messages and monitor the accepted reservations for the aggregate flows passing them. We assume that networks will typically perform policing only at their boundaries.

8 Open issues

- Definition of shaping behaviour of the source
- Encoding of *reserved* and *request* packet types
- Format of feedback (use RTCP ?)
- Construction and evaluation of efficient estimator algorithms

References

- [1] 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.
- [2] IETF, Differentiated Services (diffserv) working group. <http://www.ietf.org/html.charters/diffserv-charter.html>
- [3] Diot, Christophe; Huitema, Christian; Turletti, Thierry. *Multimedia Applications should be Adaptive*, <ftp://www.inria.fr/rodeo/diot/nca-hpcs.ps.gz>, HPCS'95 Workshop, August 1995.

9 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.Almesberger,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@cnafl.infn.it