Internet Draft
Expires September 26, 1997
File: draft-berson-classy-approach-00.ps

S. Berson
S. Vincent
ISI
March 1997

A "Classy" Approach to Aggregation for Integrated Services

March 26, 1997

Status of Memo

Abstract

The current Integrated Services (IS) architecture has a fundamental scaling problem in that per flow state is maintained everywhere. Our approach is to use aggregation as a technique to address this problem. There is a wide spectrum of approaches towards aggregation of IS state, with different tradeoffs in the amount of state required. This draft describes one promising approach to aggregating IS state within defined regions of a network. Service classes are configured on all of the routers in a region. At the edges of the region, routers keep detailed IS state and map packets into the configured traffic service classes. In the interior of this region, routers need keep only a bounded amount of IS state.

1 Introduction

In order to deploy Integrated Services (IS) in the Internet, resources are needed in the routers to keep state and to process packets. As currently defined [1], this state is on a per session basis, where a session is a unicast destination or a multicast group. Thus, widespread use of IS would mean a large amount of state at routers in the core of the internet. This large amount of state and related processing could overwhelm routers, so aggregation models that do not require per-session state need to be explored. This memo describes a particular approach to reducing this IS state.

Several types of per-session state are used with integrated services. First, there is scheduling state that consists of the different traffic service queues for packet forwarding. Second, there is packet classifier state. Classifier state is used to assign each arriving packet to a packet scheduling queue for forwarding. Third, there is the state for the setup protocol, e.g. state carried in RSVP PATH and RESV messages. Finally, current multicast routing algorithms also store state, either per source and session (e.g. DVMRP, PIM-DM) or per session shared-tree (e.g. CBT, PIM-SM). While approaches to aggregation of routing state are similar to aggregation of RSVP state, they are beyond the scope of this note.

Any aggregation scheme entails some tradeoffs. Keeping full integrated services state at all routers allows the resources dedicated to the flow of packets from each admitted session to be isolated from the resources for packets from all other sessions. This isolation means that resources reserved for one session, if needed, will be used only on the reserving session. The disadvantage of doing aggregation is the lesser degree of flow isolation, since with aggregation many flows would share the same traffic service class. The benefits of additional flow isolation achievable from keeping full IS state is unknown.

2 Framework

This draft proposes an approach to reducing the amount of IS state based on configured traffic service classes within aggregating regions. Configuring a fixed number of traffic service classes bounds the amount of scheduler state. Establishing aggregating regions (described below) and storing full IS state only at the edges reduces the classifier state and setup protocol state.

We make several assumptions:

1. We assume that an explicitly definable region (e.g. one or more contiguous autonomous systems), called an *aggregating region*, will aggregate and will implement supporting mechanisms at its boundary points.
2. We assume that the choice to aggregate and the mechanism used to do so is an intra-domain issue, not an inter-domain issue and therefore there can be variability across regions so long as at the boundaries routers continue to pass along adequate messages to support RSVP upstream and downstream.
3. We assume current multicast and unicast routing within the aggregating region; but with no other enhancements.

3 Approach

In our approach, traffic service classes are defined at all routers in an aggregating region. While the exact form of these traffic service classes is the subject of continuing research, the scheduling algorithms used by these service classes could include features such as priority queueing and/or weighted fair queueing. Each traffic service class is subject to admission control and traffic policing (additional subjects of continuing research) but only at the ingress to the region. Note that no changes to packet forwarding are needed in the interior of the aggregating region. Similarly, no RSVP processing is needed in the interior of the region. When an RSVP reservation message arrives at an ingress to the region (from an egress), and that reservation passes admission control, then packets from that flow are assigned to a traffic service class. Data packets belonging to a traffic class are *tagged* on entry to the aggregating region with some identifier indicating which class of service the packets should receive. This identifier can consist of some bits in the packet header (e.g. type of service bits) or the packet could be encapsulated. The tag (or encapsulation) encodes the traffic service class that this packet will receive across the aggregating region. Tagging or encapsulating each packet minimizes the classifier state in the interior of the aggregating region.

In summary our approach provides for reduction of packet scheduler state, packet classifier state, and setup protocol (e.g. RSVP) state in the following ways. A bounded amount of packet scheduler state at all routers in an aggregating region is achieved by using preconfigured traffic service classes. A bounded amount of packet classifier state at all interior routers in the aggregating region is accomplished by tagging or encapsulating each packet with its traffic service class. No setup protocol state at interior routers is achieved by doing admission control and policing only at the edges of the aggregating region.

4 Open Issues

Aggregation provides opportunities for increased multiplexing and resource usage efficiency, but it also brings up several challenges in admission control and policing. First there is an issue of how many service classes are to be used and of how to configure the service classes on the routers. Second, there is an issue of how and when to assign flows to the service classes, i.e. aggregate admission control. Finally, there is the issue of how and when to police an aggregate reservation. Good solutions to these problems will provide integrated services across aggregating regions with greatly reduced state requirements.

5 Acknowledgements

This draft is the result of discussions with many people particularly Bob Braden, Deborah Estrin, and Daniel Zappala.

References

[1] Braden, R., Zhang, L., Berson, S., Herzog, S., and Jamin, S., "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," Internet Draft, November 1996.

[2] Floyd, S., and Jacobson, V., "Link-sharing and Resource Management Models for Packet Networks," IEEE/ACM Transactions on Networking, Vol. 3 No. 4, pp. 365-386, August 1995.

[3] Rampal, S., "Flow Grouping for Reducing Reservation Requirements for Guaranteed Delay Service," Internet Draft, December 1996.

Security Considerations

Security considerations have not been addressed in this draft.

Author's Address

Steven Berson
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292

Phone: +1 310 822 1511
EMail: berson@isi.edu

Subramaniam Vincent
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292

Phone: +1 310 822 1511
EMail: svincent@isi.edu