RSVP

Petri Jäppilä
Nokia Telecommunications
P.O Box 330
00045 Nokia Group, Finland
Email: petri.jappila@nokia.com

Abstract

Resource ReSerVation Protocol, RSVP, is a protocol to provide resources reservation, with good scaling and robustness. The receiver of data uses RSVP to request a specific Quality of Service (QoS) from the network. RSVP carries the request through the network using same path than actual data. At each node, RSVP attempts to make a resource reservation for the stream [1].

RSVP is designed to scale for large multicast groups. It is on receiver response to reserve resource along the path. This method prevents scalability problems on sender's point of view. Sender can send data using same resource information to entire multicast group.

RSVP is one potential solution to provide QoS in Internet. It has many benefits such as possibility for large multicast groups, but as everything it has also drawbacks. It is not reasonable to make reservations for short lasting flows with RSVP due overhead. RSVP also needs routers do much more processing to every flow which needs processing power. RSVP is still under development and in future it can certainly be used in some specific applications, but it not answer how to provide QoS in all cases.

Introduction

At the moment Internet is used more and more for commercial use and it's essential to be able to satisfy customer needs. The basic need for all customer service in Internet is that sufficient amount of data can be transmitted in reasonable time. Traffic is delivered currently as best effort service, which can't guarantee the quality for service. This is most harmful for real-time services; it's not possible, for example use telephone with used quality over IP, if there is no methods to guarantee at least some transportation level at all time. The real-time services require guaranteed network consistency to avoid problems from jitter, lack of bandwidth, delay variations or information loss.

The Internet integrated services are one method to gain more means to provide quality in the Internet. Integrated services need support from individual network elements and a way to inform network elements what kind of quality is needed [1]. There are two basic concepts to provide quality through integrated services. First one is Controlled-Load method [2]. Second one is to provide guaranteed service.

Controlled load is visible to applications as best-effort service under unloaded conditions, or conditions not heavily loaded or congested. Application can trust that a very high percentage of transmitted packets will be successfully delivered and the transit delay does not greatly exceed the minimum transmit delay experienced by any successfully delivered packet.

The guaranteed service controls the maximum limits for asked service. It can set limits for delays, error limits and so on. The guaranteed service is a promise that all packet are transmitted with asked service.

1. RSVP

Resource ReSerVation Protocol, RSVP, is a resource reservation setup protocol for integrated services Internet. RSVP is a receiver-orientated protocol, which is used to reserve resources for multicast and unicast data flows.

RSVP was developed by the Information Sciences Institute of the University of Southern California and by Xerox. The RSVP specification was submitted to The Internet Engineering Steering Group (IESG) for consideration as a Proposed RFC in November 1994. At moment the development is controlled by IETF's (Internet Engineering Task Forced) Resource Reservation Setup Protocol workgroup. There are already proposed standards, for example RFC 2210 The Use of RSVP with IETF Integrated Services [2] and RFC 2206 RSVP Management Information Base using SMIv2 [3].

RSVP is used to implement quality-of-service (QoS) in the Internet. Reserving resources along the path of data

flow does this. RSVP reserves resources in all nodes along the path. This is challenging, because IP is connectionless protocol and there is not method to know the path of traffic in advance.

There were seven design goals for RSVP [4]:
1. Accommodate heterogeneous receivers.
2. Adapt to changing multicast groups.
3. Aggregate resources for multicast groups
4. The receiver can switch channels
5. Allow routes to change
6. Control protocol overhead
7. Independent from routing protocol.

These goals have been reached by two principles: RSVP is receiver oriented and uses soft state.

1.1. RSVP in action

RSVP is receiver-initiated protocol. It is on receiver's response to initiate and keep reservation active. RSVP reserves resources only for one direction at the time, which means that sender and receiver are logically distinct for each other. The RSVP does know the path for data, although it is not a routing protocol. It interacts with routing protocols and current specifications are meant to especially work with IP versions 4 and 6. RSVP is "signaling" protocol and is in principle independent of routing protocols.

The application in host can request certain QoS and the RSVP is used to inform nodes along the path about needed QoS. The ability to maintain quality for service is based on reservations of resources along the path. Traffic control gives mechanisms to implement that. These mechanisms include policing, classifying, scheduling and controlling admission, figure 1.

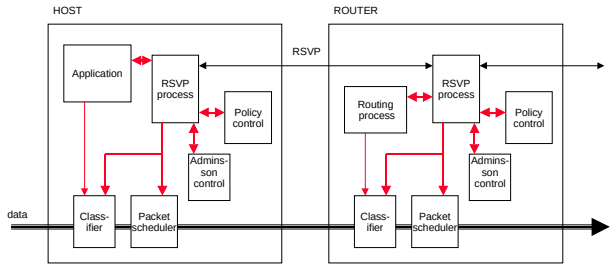

Figure 1. RSVP in Hosts and Routers

The policy control ensures that RSVP users are not able to ask more bandwidth than allowed for that user. The another important part for traffic control is admission control, which checks that there is capacity to provide asked resources. The admission control is essential part for effective performance. It balances between over and under using the capacity of the nodes.

The packet classifier and the packet scheduler implement the actual quality for data flows in nodes. Classifier checks for packet what kind resources are reserved for that packet. The packet scheduler then provides different service for classified packet. The different level QoS packet can have for instance different queues to implement needed quality. The scheduler can in this case decide the queue for classified packet to provide faster service.

RSVP supports guaranteed model, which ensures the delay restrictions are met and controlled-load service model makes no guarantees, but admits new RSVP connections only to the point where service starts to deteriorate. In Internet world even guaranteed model is not absolutely guaranteed because all link layer protocols can not provide in all cases promised quality.

1.2. RSVP messages – Path and Resv

RSVP provides signaling using messages. Typically RSVP messages are sent as IP datagrams, they have own protocol number 46.

There are two main categories in RSVP for messages - Path and Resv, figure 2. The Path message travels using same path as actual data. The sender sends Path message to downstream. Messages' routing is done by provided routing protocols.

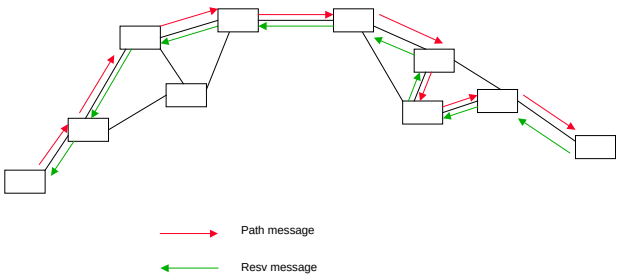

Figure 2. Path and Resv messages.

The Path message includes the information of previous hop. The path message includes also information (Sender Template) to select sender's packet from others in same session, defines traffic characters (Sender TSpec) and optionally One Pass With Advertising (OPWA) in advertisement information (AdSpec). TSpec specifies the nature of sender's traffic; traffic that exceeds the limit can be dropped, handled as best effort traffic or reshaped.

In the node the path state is stored to Path State Block (PSB). Each PSB holds path state for a particular (session, sender) pair. In PSB are stored values from Path message for example: Session, Sender Template, TSpec, Previous hop address, IP TTL and information of non-RSVP hop in route. Each incoming Path messages is

first verified against existing PSB's and if no matching PSB is found new PSB is created.

The receiver can itself decide what kind of QoS is best for it's purposes. The reservation message Resv in created by receiver and sent to upstream towards appropriate senders. These messages travel using exactly same route than the Path messages. The Resv message creates and maintains reservation state to nodes along the path. The Resv message includes flow descriptor that which consists of flowspec (TSpec, RSpec) and a Filterspec. In RSpec receiver defines wanted QoS. RSpec might define for instance the bandwidth, delay, jitter or packet loss probabilities. Filterspec defines which packets may use the reservation in the sender-to-receiver direction.

The reservation state is administered in Reservation State Block (RSB). Each RSB includes for example Session, Next Hop Address, Filter Spec, Flow Spec and Style. When Resv message incomes it's checked that there is matching PSB for session and if RSB is not created then it is created or old RSB is renewed using soft state method.

1.3. Soft State

Soft state is used in RSVP to adapt to changing routes and network topology. The word soft in this context means that state is maintained only for certain amount of time if renewing messages is not received.

In RSVP there are soft states for mentioned paths and reservations. Reservation state is attached to path state. The sender sends periodically Path messages to receiver to sets up reverse path in routers. The receiver send Resv messages to sender and router along the path can renew or establish the resource reservations.

The current routing protocols change the route when a "better" route is found. This is harmful for RSVP because it's possible that new route can not provide same resources. There are some methods to minimize route chancing, such as OPWA.

1.4. RSVP in multicast

RSVP is used for multicast and unicast connections. It was developed to scale for large multicast groups. Each receiver is itself responsible for resource reservations i.e. RSVP is possible to handle large and dynamic groups. Even in same multicast group each receiver can ask specific QoS in RSVP reservation messages upstream towards the senders.

The receivers messages are merged when another reservation in same multicast delivery tree is reached, figure 3.

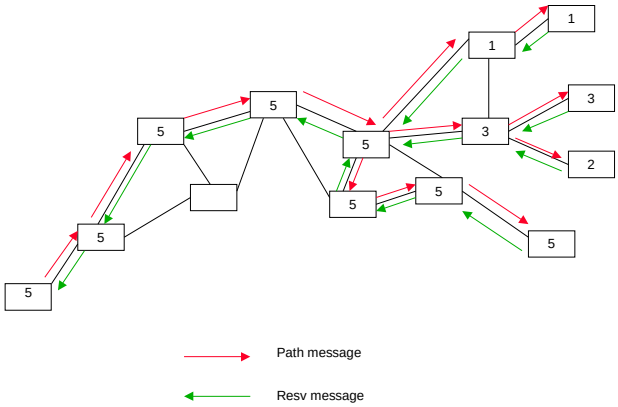

Figure 3. Merging Resv messages in multicast

Merging is needed to be to scale for large multicast groups, but as fine it might sound it is not free of problems. The basic concept in merging is to select the most stringent resource reservation request from downstream.

The request for resources can be heterogeneous and this leads to situation where new group member can cause problems to old group members. It can happen if new request is larger than earlier one; this leads to so called Killer Reservation I, figure 4. [5]

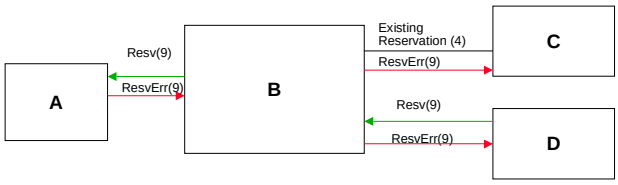

Figure 4. Killer Reservation I. [5]

On the other hand it can happen that there is some receiver asking all the time in same multicast group too much resources figure 5 [5]. The receiver for lower QoS can't get resources reserved because of merging in path. This is called as Killer Reservation II.

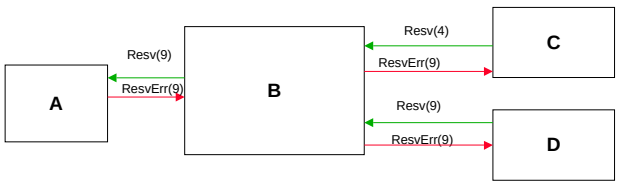

Figure 5. Killer Reservation II. [5]

Killer reservations can be avoided using different methods such as Blockade State Mechanism, which tries

to provide resources for lower request from downstream if merged and larger capacity can't be provided.

Figure 6. represent Blockade State Mechanism in principle, receivers C and D are asking resources for same multicast, C (asks resources 4) and D(9). At merging point B reservations are merged and B asks resources (9). At this time A sends message to B indicate lack of resources. This creates a blockade state at node B for the previous hop A. The merging point B has knowledge that for receiver C smaller amount of resources is enough and it can ask resources for C's purposes B(4). Naturally receiver D is informed that reservation could not be made.

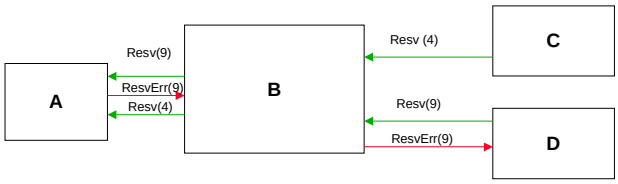

Figure 6. Blockade State Mechanism

2. RSVP in non-RSVP clouds

It is nature of Internet that all routers can't be updated at same time or even in longer period to support all new protocols. The RSVP makes no exception to this. This means that RSVP must work also in networks where along the path might be non-RSVP routers. RSVP is designed to work transparently across non-RSVP clouds, figure 7.

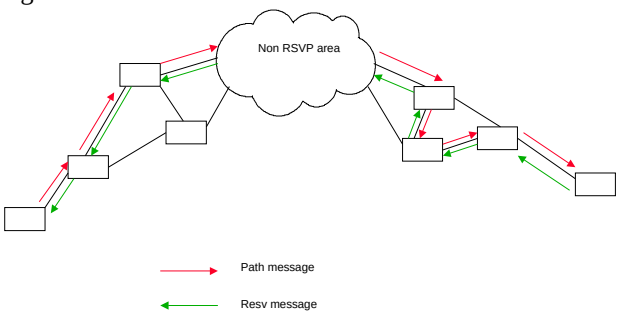

Figure 7. RSVP in non-RSVP clouds

If there is non-RSVP routers along the path the QoS can't be assured. It is needed to inform the receiver that there is non-RSVP router along the path. In normal IP forwarding this quite easy just verifying RSVP-messages Time To Live (TTL) value to normal IP-message TTL-value. In real world this not enough because routing protocol might use tunneling and in this case the routing protocol must inform RSVP when non-RSVP clouds are included.

3. Reservation Styles

RSVP separates resources bandwidth from deciding which packets get to use the resources. RSVP only specifies what amount of packets can use the resources. The packets, which can use reserved resources, can be identified using filtering.

Reservation styles are used to aggregate reservation requests from different receivers of the same multicast group. There are two main options to define the reservation style.

One option defines senders, which can use reserved resources in same session. Another option defines if reservations are distinct or shared. In distinct reservations the sources are explicitly specified.

RSVP uses three reservation styles: Fixed-Filter (FF), shared-explicit (SE) and Wildcard-Filter (WF), figure 8.

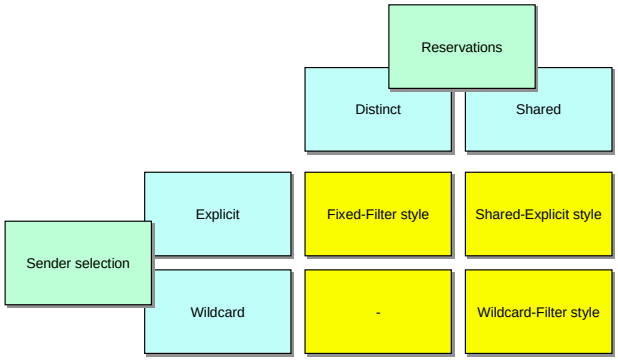

Figure 8. Reservation attributes and styles [1]

Wildcard-Filter Style means that receiver wants to receive all flows from all upstream senders in the session and that its bandwidth reservation is shared among all senders.

Fixed-Filter Style means that receiver specifies a list of senders from which it wants to receive a data flow along with a bandwidth reservation for each of these senders.

Shared-Explicit Style means that receiver specifies a list of senders from which it wants to receive a data flow along with a shared bandwidth.

4. RSVP products available

Many vendors already implement RSVP in their products. RSVP is available for example in Bay Networks, Cisco and 3Com routers [7],[8]. Intel and Sun also implement RSVP in hosts. Microsoft will support RSVP in Windows NT Server 5.0.

5. Problems in RSVP

The big issue in RSVP is how to avoid oversubscribing. The oversubscribing can happen if receivers ask more capacity than they really need, due lack of resources or even by attackers. RSVP needs good policy control mechanisms to be sure that receiver is allowed to make reservations [9].

The routers need in RSVP handle and store information about every flow. This needs more processing power than most of today's router can offer [10]. This leads that RSVP is considered as an Intranet protocol where it can be used to allow important service to have constant bandwidth. The operators could also use their networks so that on the edge of their networks the streams are aggregated to minimize the number of RSVP flows.

The nature of Internet has offered free rides, because of Best Effort. RSVP will need new kinds of accounting principles in Internet. This will rise new and even philosophical problems, but it could also provide more fair mechanisms to provide better quality to paying consumer.

RSVP is not very good for WWW-applications, there is too much negotiation work for short time lasting traffic, which most of current WWW-traffic is. The setting up the reservation causes overhead to the net. It is certainly not wanted to have situation were control traffic is almost as big as data flow.

Plenty of IP security protocols are based on extra headers between IP header and transport layer and RSVP which is tailored to use port numbers (e.g. TCP ports) to identify sessions can't naturally get that information from encapsulated headers. This problem can be avoided using Session Parameter Index to identify session. [11]

Merging of multicast resources allows all receivers to benefit from the accepted reservation. This will be unfair to the highest resource user, which probably is paying quality to rest of multicast users.

6. Conclusions

RSVP tries to provide methods to implement QoS in connectionless Internet. It is partly contradictory to the current infrastructure of Internet. RSVP tries to adapt to chancing routes using soft state. The guaranteed service has limited applicability, because there is not practical methods to provide at least strict guarantees.

It can be seen that the Internet will need some methods to provide QoS and RSVP is one possible way to implement that. RSVP has many benefits such as possibility to support large multicast groups. RSVP needs more processing in routers than routing protocols; it will need also hardware updates. It is presumable that RSVP will be first adapted to Intranet were infrastructure is easier to change and security to policy control be can be provided.

It must be kept on mind that RSVP is the first significant industry-standard protocol in the Internet. RSVP can help us, but it can't do everything.

References

[1] Braden, R., Zhang, L., Berson, S., Herzog, S., Jamin, S., Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification, RFC 2205, September 1997, Proposed Standard. http://www.isi.edu/div7/rsvp/pub.html

[2] Wroclawski, J., The Use of RSVP with IETF Integrated Services, RFC 2210, September 1997, Proposed Standard. http://www.isi.edu/div7/rsvp/pub.html

[3] Baker, F., Krawczyk, and J. Sastry, A., RSVP Management Information Base using SMIv2, RFC 2206, September 1997, Proposed Standard. http://www.isi.edu/div7/rsvp/pub.html

[4] Zhang, L. Deering, S., Estrin, D., Shenker, S., and Zappala, D., RSVP: A New Resource ReSerVation Protocol IEEE Network, September 1993

[5] Talwar,M., RSVP Killer Reservations, http://search.ietf.org/internet-drafts/draft-talwar-rsvp-kr-01.txt

[6] Roberts, E., RSVP: A Priority Problem?, Data Communications On the Web, http://www.data.com/roundups/rsvp.html

[7] Higgins, K., Web Standards: QoS Takes Shape, Trends, September 28, 1998, http://www.internetwk.com/trends/trends092898.htm

[8] Gaines, G., Festa, M., A Survey of RSVP/QoS Implementations, July 1998, RSVP Working Group, http://www.isi.edu/rsvp/DOCUMENTS/ietf_rsvp-qos_survey_02.txt

[9] Higgins, K., Net For Business -- IETF Protocols Boost Network Efficiency, But Are They Just A Short-Term Fix?, InformationWeek, September 28, 1998, Issue: 702, http://www.techweb.com/se/directlink.cgi?IWK19980928S0050

[10] Mankin, A., Baker, F., Braden, B., Bradner, S., O'Dell, M., Romanow, A., Weinrib, A., Zhang, L., Resource ReSerVation Protocol (RSVP) -- Version 1 Applicability Statement RFC 2208, September 1997, Informational.

[11] Berger, L., O'Malley, T., RSVP Extensions for IPSEC Data Flows, RFC 2207, September 1997, Proposed Standard. ftp://ftp.isi.edu/in-notes/rfc2207.txt