RSVP and Integrated Services in the Internet: a tutorial

Author: Paul P.White

Department of Computer Science
University College London
Gower Street
London
WC1E 6BT

Tel: 44 171 419 3701
Email: p.white@cs.ucl.ac.uk

Abstract

The growing use of multimedia communication applications with specific bandwidth and real-time delivery requirements has created the need for an Integrated Services Internet in which traditional best-effort datagram delivery can coexist with additional enhanced Quality Of Service(QOS) delivery classes. Such classes provide data flows with QOS commitments with regard to bandwidth, packet loss and delay through the reservation of network resources along the data path which can be done using the Resource ReSerVation Protocol (RSVP). This paper is a tutorial on how RSVP can be used by end-applications to ensure that they receive the end-to-end QOS that they require.

Keywords

Controlled-Load Service
Data Flow
Guaranteed Service
Integrated Services (IS)
One Pass with Advertising (OPWA)
Quality of Service (QOS)
ReSerVation setup Protocol (RSVP)

This paper appeared in the May 1997 Issue of IEEE Communications magazine

1.0 Introduction

The current Internet consists of a multitude of networks built from various link layer technologies and relies on the Internet Protocol (IP) to interwork between them. IP makes no assumptions about the underlying protocol stacks and offers an unreliable, connectionless network layer service that is subject to packet loss, reordering and packet duplication, all of which, together with queuing delay in router buffers will increase with network load. Because of the lack of any firm guarantees the traditional IP delivery model is often referred to as 'best-effort' with an additional higher layer end-to-end protocol such as the Transmission Control Protocol (TCP) being required to provide end-to-end reliability. TCP does this through the use of such mechanisms as packet re-transmission which further adds to the overall information transfer delay.

For traditional non-real-time Internet traffic such as File Transfer Protocol (FTP) data the best-effort delivery model of IP has not been a problem. However as we move further into the age of multimedia communications many real-time applications are being developed that are delay-sensitive to the point where the best-effort delivery model of IP can be inadequate even under modest network loads. Although the problem has been alleviated somewhat through making certain applications adaptive to network load where possible, there is still a firm need to provide many applications with additional service classes offering enhanced QOS with regard to bandwidth, packet queuing delay and loss. These additional enhanced QOS delivery classes would supplement the best-effort delivery service in what could be described as an Integrated Services Internet[2].

2.0 IETF Integrated Services

In response to the growing demand for an Integrated Services Internet the Internet Engineering Task Force (IETF)[14] set up an Integrated Services (intserv) Working Group[15] which has since defined several service classes that if supported by the routers traversed by a data flow[1] can provide the data flow with certain QOS commitments. By contrast best-effort traffic entering a router will receive no such service commitment and will have to make do with whatever resources are available. The level of QOS provided by these enhanced QOS classes is programmable on a per-flow basis according to requests from the end applications. These requests can be passed to the routers by network management procedures or, more commonly, using a reservation protocol such as RSVP which is described in section 3. The requests dictate the level of resources(e.g. bandwidth, buffer space) that must be reserved along with the transmission scheduling behaviour that must be installed in the routers to provide the desired end-to-end QOS commitment for the data flow.

In determining the resource allocations necessary to satisfy a request the router needs to take account of the QOS support provided by the link layer in the data forwarding path. Furthermore, in the case of a QOS-active link layer such as ATM or certain types of LAN the router is responsible for negotiations with the link layer to ensure that the link layer installs appropriate QOS support should the request be accepted.

---

[1]A data flow identifies the set of packets to receive special QOS. It is defined by a 'Session' comprising the IP address, transport layer protocol type and port number of the destination along with a list of specific senders to that Session that are entitled to receive the special QOS. Each sender is identified by source address, and port number while it's protocol type must be the same as for the Session.

This mapping to link-layer QOS is medium-dependent and the mechanisms for doing so are currently being defined by the Integrated Services over Specific Lower Layers (issll) Working Group of the IETF[16]. In the case of a QOS-passive link-layer such as a leased-line the mapping to the link-layer QOS is trivial since transmission capacity is handled entirely by the router's packet scheduler.

Each router must apply admission control to requests to ensure that they are only accepted if sufficient local resources are available. In making this check, admission control must consider information supplied by end applications regarding the traffic envelope that their data flow will fall within. One of the parameters in the traffic envelope that must be supplied is the maximum datagram size of the data flow, and should this be greater than the MTU of the link then admission control will reject the request since the Integrated Services models rely on the assumption that datagrams receiving an enhanced QOS class are never fragmented.

Once an appropriate reservation has been installed in each router along the path, the data flow can expect to receive an end-to-end QOS commitment provided no path changes or router failures occur during the lifetime of the flow, and provided the data flow conforms to the traffic envelope supplied in the request. Service-specific policing and traffic reshaping actions as described in sections 2.1 and 2.2 will be employed within the network to ensure that non-conforming data flows do not affect the QOS commitments for behaving data flows. The IETF has considered various QOS classes such as [1][6][10][12] although to date only two of these, Guaranteed Service[10] and Controlled-Load Service[12], have been formally specified for use with RSVP[13].

2.1 Guaranteed Service

Guaranteed Service[10] provides an assured level of bandwidth, a firm end-to-end delay bound and no queuing loss for conforming packets of a data flow. It is intended for applications with stringent real-time delivery requirements such as certain audio and video applications that use "play-back" buffers and are intolerant of any datagram arriving after their playback time.

Each router characterises the Guaranteed Service for a specific flow by allocating a bandwidth, R and buffer space, B that the flow may consume. This is done by approximating the "fluid model" of service[8][9] so that the flow effectively sees a dedicated wire of bandwidth, R between source and receiver. In a perfect fluid model a flow conforming to a token bucket of rate, r and depth, b will have it's delay bounded by b/R provided R >= r. To allow for deviations from this perfect fluid model in the router's approximation [2], two error terms, C and D are introduced. Consequently the delay bound now becomes b/R + C/R + D. However with Guaranteed Service a limit is imposed on the peak rate, p of the flow which results in a reduction of the delay bound. In addition, the packetisation effect of the flow needs to be taken into account by considering the maximum packet size, M. These additional factors result in a more precise bound on the end to end queuing delay as follows

$$Qdelay_{end2end} = \frac{(b-M)(p-R)}{R(p-r)} + \frac{(M+Ctot)}{R} + Dtot \quad \text{(case p > R >= r)} \qquad (1)$$

[2] Among other things the router's approximation must take account of the medium-dependent behaviour of the link layer of the data forwarding path.

$$Qdelay_{end2end} = \frac{(M + Ctot)}{R} + Dtot \qquad \text{(case R >= p >= r)} \quad (2)$$

Where Ctot and Dtot represent the summation of the C and D error terms respectively for each router along the end-to-end data path.

In order for a router to invoke Guaranteed Service for a specific data flow it needs to be informed of the traffic characteristics, Tspec of the flow along with the reservation characteristics, Rspec. Furthermore to enable the router to calculate sufficient local resources to guarantee a lossless service it requires the terms Csum and Dsum which represent the summation of the C and D error terms respectively for each router along the path since the last re-shaping point(see below).

Tspec parameters

p = peak rate of flow (bytes/second)
b = bucket depth (bytes)
r = token bucket rate (byes/second)
m = minimum policed unit (bytes)[3]
M = maximum datagram size (bytes)

Rspec parameters

R = bandwidth, i.e. service rate (bytes/second)
S = Slack Term (ms) (see section 3.6)

Guaranteed Service traffic must be policed at the network access points to ensure conformance to the Tspec. The usual enforcement policy is to forward non-conforming packets as best-effort datagrams[4] and if and when a marking facility becomes available these non-conforming datagrams should be marked to ensure that they are treated as best-effort datagrams at all subsequent routers.

In addition to policing of data flows at the edge of the network Guaranteed Service also requires reshaping of traffic to the token bucket of the reserved Tspec at certain points on the distribution tree. Any packets failing the reshaping are treated as best-effort and marked accordingly if such a facility is available. Reshaping must be applied at any points where it is possible for a data flow to exceed the reserved Tspec even when all senders associated with the data flow conform to their individual Tspecs. Such an occurrence is possible in the following 2 cases.

Firstly, at branch points in the distribution tree where the reserved Tspecs of the outgoing branches are not the same. In this case the reserved Tspec of the incoming branch is given by the 'maximum'[5] of the reserved Tspecs on each of the outgoing branches. Consequently some of the outgoing branches will have a reserved Tspec that is less than the reserved Tspec of the incoming branch and so it is possible that in the absence of reshaping, traffic that conforms to the Tspec of the incoming branch

---

[3] Policing will treat any IP datagram less than size m as being of size m.

[4] Action with regard to non-conforming datagrams should be configurable to allow for situations such as traffic-sharing where the preferred action might be to discard non-conforming datagrams. This configuration requirement also applies to reshaping.

[5] Maximum according to rules defined in [3].

might not conform when routed through to an outgoing branch with a smaller reserved Tspec. As a result, reshaping must be performed at each such outgoing branch to ensure that the traffic is within this smaller reserved Tspec.

Secondly, at merge points in the distribution tree for sources sharing the same reservation since in these cases the sum of the Tspecs relating to the incoming branches will be greater than the Tspec reserved on the outgoing branch. Consequently when multiple incoming branches are each simultaneously active with traffic conforming to their respective Tspecs it is possible that when this traffic is merged onto the outgoing branch it will violate the reserved Tspec of the outgoing branch. Hence reshaping to the reserved Tspec of the outgoing branch is necessary.

2.2 Controlled-Load Service

Unlike Guaranteed Service the Controlled-Load Service[12] provides no firm quantitative guarantees. A Tspec for the flow desiring Controlled-Load Service must be submitted to the router as for the case of Guaranteed Service although it is not necessary to include the peak rate parameter. If the flow is accepted for Controlled-Load Service then the router makes a commitment to offer the flow a service equivalent to that seen by a best-effort flow on a lightly loaded network. The important difference is that the Controlled-Load flow does not noticeably deteriorate as the network load increases. This will be true regardless of the level of load increase. By contrast, a best-effort flow would experience progressively worse service (higher delay and loss) as the network load increased. The Controlled-Load Service is intended for those classes of applications that can tolerate a certain amount of loss and delay provided it is kept to a reasonable level. Examples of applications in this category include adaptive real-time applications.

Routers implementing the Controlled-Load Service must check for conformance of Controlled-Load data flows to their appropriate reserved Tspecs. Any non-conforming Controlled-Load data flows must not be allowed to affect the QOS offered to conforming Controlled-Load data flows or to unfairly affect the handling of best-effort traffic. Within these constraints the router should attempt to forward as many of the packets of the non-conforming Controlled-Load data flow as possible. This might be done by dividing the packets into conforming and non-conforming groups and forwarding the non-conforming group on a best effort basis. Alternatively, the router may choose to degrade the QOS of all packets of a non-conforming Controlled-Load data flow equally.

Resource ReSerVation Protocol (RSVP)

The Resource ReSerVation Protocol RSVP[3] was designed to enable the senders, receivers and routers of communication Sessions(either multicast or unicast) to communicate with each other in order to set up the necessary router state to support the services described in sections 2.1 and 2.2. It is worth noting that RSVP is not the only IP reservation protocol that has been designed for this purpose. Others include ST-II[11] and ST-II+[5] which incidently contain some interesting architectural differences to RSVP such as the use of hard-state and sender-initiated[6] reservations

[6] ST-II+ permits both sender and receiver-initiated reservations, ST-II permits sender-initiated reservations only.

rather than soft-state[7] and receiver-initiated reservations as in RSVP. However for the rest of this tutorial the only reservation protocol we consider is RSVP since currently this has the most industry support. For further discussion on the mentioned alternatives the interested reader can refer to [7].

RSVP identifies a communication Session by the combination of destination address, transport layer protocol type and destination port number. It is important to note that each RSVP operation only applies to packets of a particular Session and as such every RSVP message must include details of the Session to which it applies. For the remainder of this tutorial it will be assumed that any discussion is for a single Session only. In addition, although RSVP is applicable to both unicast and multicast Sessions we concentrate on the more complicated multicast case. Also we do not discuss the security issues of RSVP or any billing that may be necessary to exert backpressure on the use of reservations.

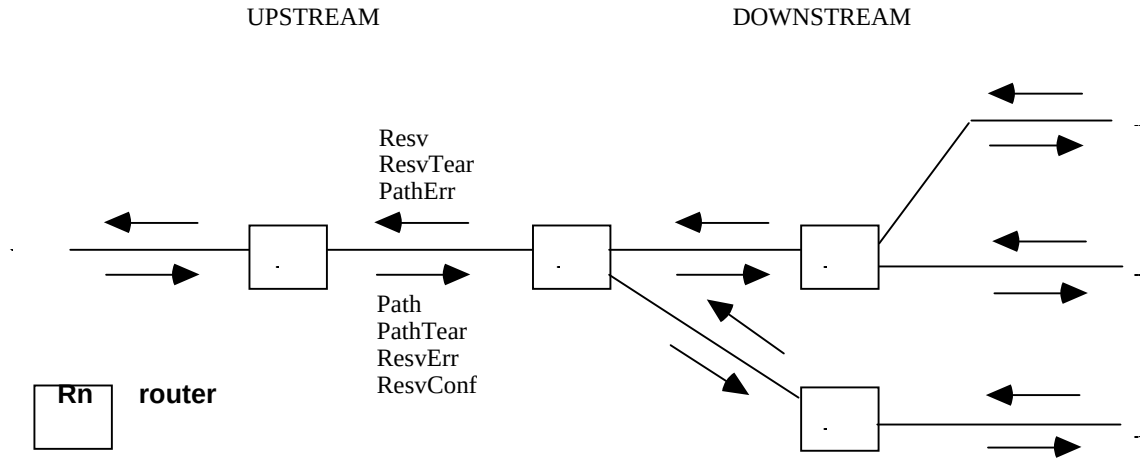

Figure 1: Direction of RSVP messages.

RSVP is not a routing protocol; it is merely used to reserve resources along the existing route set up by whichever underlying routing protocol is in place. Figure 1 shows an example of RSVP for a multicast Session involving one sender, S1 and 3 receivers, RCV1 - RCV3. The primary messages used by RSVP are the Path message which originates from the traffic sender and the Resv message which originates from the traffic receivers. The primary roles of the Path message are firstly, to install reverse routing state in each router along the path and secondly, to provide receivers with information about the characteristics of the sender traffic and end-to-end path so that they can make appropriate reservation requests. The primary role of the Resv message is to carry reservation requests to the routers along the distribution tree between receivers and senders. Returning now to Figure 1, as soon as S1 has data to send it begins periodically forwarding RSVP Path messages to the next hop, R1 down the distribution tree. RSVP messages can be transported 'raw' within IP datagrams using protocol number 46 although hosts without this raw I/O capability may first encapsulate the RSVP messages within a UDP header.

[7] With hard-state the network is responsible for reliably maintaining router state whereas with soft-state the responsibility is passed to the end-systems which must generate periodic refreshes to prevent state timeout.

3.1 Path Messages

Each Path message includes the following information

- Phop, the address of the last RSVP-capable node to forward this Path message. This address is updated at every RSVP-capable router along the path.
- The Sender Template, a filter specification identifying the sender. It contains the IP address of the sender and optionally the sender port(in the case of IPv6 a flow label may be used in place of the sender port).
- The Sender Tspec defining the sender traffic characteristics.
- An optional Adspec containing OPWA information(see sections 3.3 and 3.4) which is updated at every RSVP-capable router along the path to attain end-to-end significance before being presented to receivers to enable them to calculate the level of resources that must be reserved to obtain a given end-to-end QOS.

3.2 Processing and Propagation of Path Messages by Network Routers

Each intermediate RSVP-capable router along the distribution tree intercepts Path messages and checks them for validity. If an error is detected then the router will drop the Path message and send a PathErr message upstream to inform the sender who can then take appopriate action. Assuming the Path message is valid the router does the following:

- Update the path state entry for the sender identified by the Sender Template. If no path state exists then create it. Path state includes the Sender Tspec, the address, Phop of the previous hop upstream router and optionally an Adspec. The Phop address needs to be stored in order to route Resv messages in the reverse direction up the tree. The Sender Tspec provides a ceiling to clip any inadvertently over-specified Tspecs subsequently received in Resv messages
- Set cleanup timer equal to cleanup timeout interval and restart timer.

Associated with each path state entry is a cleanup timer, the expiration of which triggers deletion of the path state. Expiration of the timer will be prevented if a Path message for the entry is received at least once every cleanup timeout interval. This is the so-called RSVP "soft state" mechanism and ensures that state automatically times out if routing changes while subsequent Path messages install state along the new routing path. In this way the use of soft-state rather than hard-state helps to maintain much of the robustness of the initial Internet design concepts whereby all flow-related state was restricted to the end systems[4].

The router is also responsible for generating Path messages based on the stored path state and forwarding them down the routing tree making sure that for each outgoing interface the Adspec(see section 3.3) and Phop objects are updated accordingly. Path messages will be generated and forwarded whenever RSVP detects any changes to stored path state or is informed by the underlying routing protocol of a change in the set of outgoing interfaces in the data forwarding path. Otherwise, a Path message for each specific path state entry is created and forwarded every refresh period timeout interval in order to refresh downstream path state.

The refresh period timeout interval is several times smaller than the cleanup timeout interval so that occasional lost Path messages can be tolerated without triggering unnecessary deletion of path state. However it is still recommended that a minimum network bandwidth be configured for RSVP messages to protect them from congestion losses.

Although all path state would eventually timeout in the absence of any refreshes via Path messages, RSVP includes an additional message, PathTear to expedite the process. PathTear messages travel across the same path as Path messages and are used to explicitly tear down path state. PathTear messages are generated whenever a path state entry is deleted and so a PathTear message generated by a sender will result in deletion of all downstream path state for that sender. It is recommended that senders do this as soon as they leave the communications Session. Also, deletion of any path state entry triggers deletion of any dependent reservation state(see section 3.5).

3.3 Adspec

The Adspec is an optional object that the sender may include in it's generated Path messages in order to advertise to receivers the characteristics of the end to end communications path. This information can be used by receivers to determine the level of reservation required in order to achieve their desired end to end QOS. The Adspec consists of a message header, a Default General Parameters fragment and at least one of the following; Guaranteed Service fragment, Controlled-Load Service fragment. Omission of either the Guaranteed or Controlled-Load Service fragment is an indication to receivers that the omitted service is not available. This feature can be used in a multicast Session to force all receivers to select the same service. (At present RSVP does not accommodate heterogeneity of services between receivers within a given multicast Session).

The Default General Parameters fragment includes the following fields which are updated at each RSVP-capable router along the path in order to present end-to-end values to the receivers.

- Minimum Path Latency (summation of individual link latencies). This parameter represents the end-to-end latency in the absence of any queuing delay. In the case of Guaranteed Service, receivers can add this value to the bounded end-to-end queuing delay to obtain the overall bounded end-to-end delay.

- Path Bandwidth (minimum of individual link bandwidths along the path)
- Global Break bit - This bit is cleared when the Adspec is created by the sender. Encountering any routers that do not support RSVP will result in this bit being set to one in order to inform the receiver that the Adspec may be invalid.
- Integrated Services(IS) Hop count - incremented by one at every RSVP/IS-capable router along the path.
- PathMTU - Path Maximum Transmission Unit (minimum of MTUs of individual links along the path).

Correct functioning of IETF Integrated Services requires that packets of a data flow to receive the special QOS are never fragmented. This also means that the value of M in the Tspec of a reservation request must never exceed the MTU of any link that the reservation request applies to. A receiver can ensure that this requirement is met by

setting the value of M in the Tspec of it's reservation request to the minimum of the PathMTU values received in 'relevant' Path messages. A Path message is relevant if it originated from a sender that is captured in the intended reservation request in accordance with the reservation Styles described in section 3.5. The value of M in each generated reservation request may be further reduced on the way to each sender if merging of Resv messages occurs(see section 3.5). The minimum value of M from the Tspec of each Resv message[8] received by the sender should then be used by the sending application as the upper limit on the size of packets to receive special QOS. In this way fragmentation of these packets will never occur. It is worth noting that [13] recommends that the value of M in the Sender Tspec, which has played no part in the above MTU negotiation process, should be set equal to the maximum packet size that the sender is capable of generating rather than what it is currently sending.

The Guaranteed Service fragment of the Adspec includes the following fields which are updated at each RSVP-capable router along the path in order to present end-to-end values to the receivers.

- Ctot - end to end composed value for C.
- Dtot - end to end composed value for D.
- CSum - composed value for C since last reshaping point.
- DSum - composed value for D since last reshaping point (CSum and Dsum values are used by reshaping processes at certain points along the distribution tree).
- Guaranteed Service Break Bit - This bit is cleared when the Adspec is created by the sender. Encountering any routers that do support RSVP/IS but do NOT support Guaranteed Service will result in this bit being set to one in order to inform the receiver that the Adspec may be invalid and the service cannot be guaranteed.
- Guaranteed Service General Parameters Headers/Values - These are optional but if any are included then each one overrides the corresponding value given in the Default General Parameters fragment as far as a receiver wishing to make a Guaranteed Service reservation is concerned. These override parameters could for example be added by routers along the path that have certain service-specific requirements. For example a router may have been configured by network management so that Guaranteed Service reservations can only take up a certain amount, Bgs of the outgoing link bandwidth. Consequently if the Default Path Bandwidth value in the Adspec to be sent out of this interface is greater than Bgs then a Guaranteed Service Specific Path Bandwidth header and value equal to Bgs may be included in the Adspec. As for Default General Parameters, any Service-Specific General Parameters must be updated at each RSVP hop.

The Controlled-Load Service fragment of the Adspec includes the following fields which are updated at each RSVP-capable router along the path in order to present end-to-end values to the receivers.

- Controlled-Load Service Break Bit - This bit is cleared when the Adspec is created by the sender. Encountering any routers that do support RSVP/IS but do NOT support Controlled-Load will result in this bit being set to one in order to inform the receiver that the Adspec may be invalid and the service cannot be guaranteed.

---

[8] In cases where the last hop to a sender is a shared medium LAN the sender may receive Resv messages across the same interface from multiple next hop routers.

- Controlled-Load Service General Parameters Headers/Values - As for the Guaranteed Service fragment, override Service-Specific General Parameters may be added to the Controlled-Load Service fragment.

3.4 Making a Reservation using One Pass with Advertising (OPWA)

One Path With Advertising(OPWA) refers to the reservation model for the case where the sender includes an Adspec in it's Path messages to enable the receiver to determine the end-to-end service that will result from a given reservation request. If the sender omits the Adspec from it's Path messages then the reservation model is referred to simply as One Pass in which case there is no easy way for the receiver to determine the resulting end-to-end service. Here we consider the OPWA case.

Let us assume that the sender omits the Controlled-Load Service data fragment from the Adspec thereby restricting each receiver to reservation of Guaranteed Service only. Upon receiving Path messages the receiver extracts the following parameters from the Sender Tspec contained therein: r, b, p, m. In addition the following are extracted from the Adspec: Minimum Path Latency, Ctot, Dtot, PathMTU, Path Bandwidth.

The required bound on end-to-end queuing delay, Qdelreq is now calculated by subtracting the Minimum Path Latency from the value of end-to-end delay required by the receiver's application. Typically, the receiver would then perform an initial check by evaluating equation (2) for R equal to the peak rate, p. If the resultant delay was greater than or equal to Qdelreq then equation (2) would be used for calculation of the minimum value of R necessary to satisfy Qdelreq. Otherwise equation (1) would be used for this purpose. This minimum value of R is then obtained by inserting Qdelreq into either equation (1) or (2) along with M(given by PathMTU), Ctot, Dtot, r, b, p, as appropriate. If the obtained value of R exceeds the Path Bandwidth value as obtained from the Adspec of the received Path message then it must be reduced accordingly. The receiver can now create a reservation specification, Rspec comprising firstly the calculated value, R of bandwidth to be reserved in each router, and secondly a Slack Term that is initialised to zero[9]. The Rspec can now be used in the creation of a Resv message which also includes the following:

- An indication of the reservation style which can be FF, SE or WF (see section 3.5)
- A filter specification, Filterspec (omitted for the case of WF reservation style). This is used to identify the sender(s) and the format is identical to that of the Sender Template in a Path message.
- A flow specification, Flowspec comprising the Rspec and a traffic specification, Tspec. Tspec is usually set equal to the Sender Tspec except M will be given by PathMTU obtained from the received Adspec.
- Optionally a reservation confirm object, ResvConf containing the IP address of the receiver. If present this object indicates that the node accepting this reservation request at which propagation of the message up the distribution tree finishes should return a ResvConf message to the receiver to indicate that there is a high probability[10] that the end-to-end reservation has been successfully installed.

---

[9]In some cases even with R set to the minimum permissible value of r the resultant end-to-end queuing delay as given by eqs (1) and (2) will still be less than Qdelreq in which case the difference can be represented in a non-zero slack term. In addition there are other scenarios explained in section 3.6 in which the slack term may not be initialised to zero.

[10]In practice there are certain scenarios in which a ResvConf message might be received by a receiver only for the request to be rejected shortly afterwards.

The Resv message is now sent to the previous hop upstream as obtained from the stored path state. Upon reaching the next upstream router the Resv messages can be merged with other Resv messages arriving on the same interface according to certain rules as described in section 3.5 to obtain an effective Flowspec and Filterspec. The following action is then taken.

- The effective Flowspec is passed to the traffic control module within the router which applies both admission control and policy control to determine whether the reservation can be accepted. Admission control is concerned solely about whether enough capacity exists to satisfy the request while policy control also takes into account any additional factors that need to be considered (e.g. certain policies may limit a users reserved bandwidth even if spare bandwidth exists).
- If the reservation attempt is denied then any existing reservations are left unaltered and the router must send a ResvErr message downstream.
- If the reservation request is accepted then reservation state is set up in accordance with the effective Flowspec and Filterspec as described in section 3.5. In accepting the request it may be permissible to alter the Rspec associated with the reservation from (Rin, Sin) to (Rout, Sout) in accordance with the rules described in section 3.6. The resultant reservation may then be merged with other reservations in accordance with the rules in section 3.5 to obtain a new Resv message that is sent to the next router upstream, the address of which is obtained from the stored path state.

3.5 Reservation Styles and Merging

Associated with each reservation made at a router's interface is a Filterspec describing the packets to which the reservation applies along with an effective Flowspec. Both the Filterspec and effective Flowspec are obtained from a merging process applied to selected Resv messages arriving on the router's interface. The rules for merging are dependent upon the reservation Style of each Resv message as described below. In addition the router calculates the Filterspec and Flowspec of Resv messages to be sent to the previous hop(s) upstream by applying Style-dependent merging of stored reservation state. Any changes to stored reservation state that result in changes to the Resv messages to be sent upstream will cause an updated Resv message to be sent upstream immediately. Otherwise Resv messages are created based on stored reservation state and sent upstream periodically. As for path state all reservation state is stored in routers using soft-state and consequently relies on periodic refreshes via Resv messages to prevent state timeout. In addition just as a PathTear message exists to explicitly tear down path state, a ResvTear message exists to explicitly tear down reservation state. Currently 3 reservation Styles are permissible as described below and illustrated in Figures 2-6 where the convention Style(Filterspec{Flowspec}) is used to summarise the requests made by the Resv messages. It should be noted that the merging processes described below apply only to packets of the same Session(This is true of any RSVP process). Also merging can only occur between messages with the same reservation style. Details of the reservation styles are as follows where it is assumed that each interface, I in Figures 2-4 is routable to each of the router's other interfaces.

Fixed Filter (FF) **(distinct reservation and explicit sender selection)**

The Filterspec of each FF reservation installed at an interface consists of a single sender only. The effective Flowspec of the reservation installed is the maximum of all FF reservation requests received[11] on that interface for that particular sender. The Flowspec of the FF Resv message unicast to the previous hop of a particular sender is given by the maximum Flowspec of all reservations installed in the router for that particular sender.

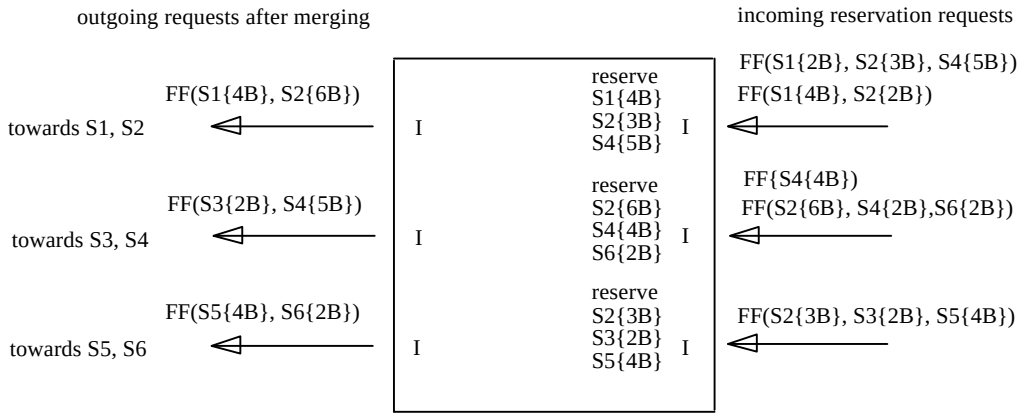

Figure 2: Fixed Filter Reservation Example.

Wildcard Filter (WF) **(shared reservation and wildcard sender selection)**

The Filterspec of each WF reservation installed at an interface is wildcard and matches on any sender from upstream. The effective Flowspec installed is the maximum from all WF reservation requests received on that particular interface. The Flowspec of each WF Resv message unicast to a previous hop upstream is given by the maximum Flowspec of all WF reservations installed in the router[12].

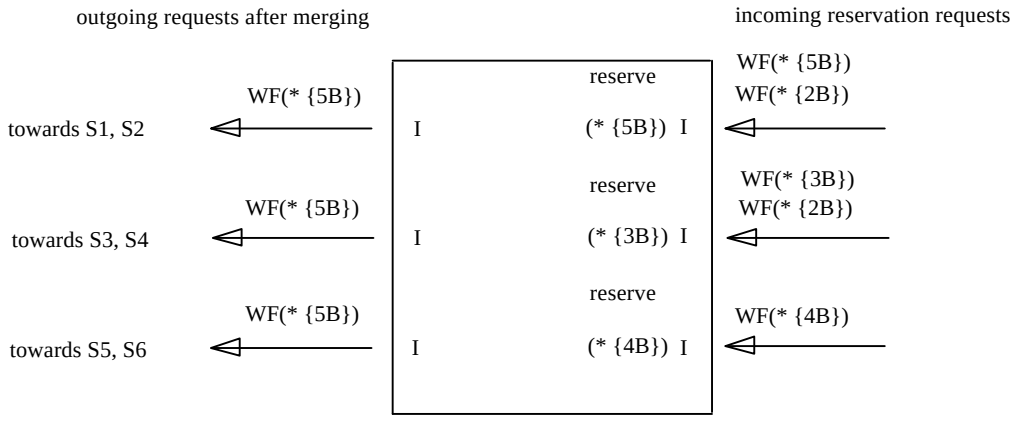

[11] In cases where the interface connects to a shared medium LAN Resv messages from multiple next hops may be received.

[12]Strictly speaking, only WF reservations whose 'Scope' applies to the interface out of which the Resv message is sent are considered for this second merging process. Scope details are required for WF reservations on non-shared trees to prevent looping. Further details can be found in [3].

Figure 3: Wildcard Filter Reservation Example.

Shared Explicit (SE) **(shared reservation and explicit sender selection)**

The Filterspec of each SE reservation installed at an interface contains a specific set of senders from upstream and is obtained by taking the union of the individual Filterspecs from each SE reservation request received on that interface. The effective Flowspec installed is the maximum from all SE reservation requests received on that particular interface. The Filterspec of a SE Resv message unicast out of an interface to a previous hop upstream is the union of all senders whose previous hop is via that interface and who are contained in the Filterspec of at least one SE reservation in the router. Likewise the Flowspec of this SE Resv message is given by the maximum Flowspec of all SE reservations whose Filterspecs contain at least one sender whose previous hop is via that interface.

Figure 4: Shared Explicit Reservation Example.

SE and WF styles are useful for conferencing applications where only one sender is likely to be active at once in which case reservation requests for say twice the sender bandwidth could be reserved in order to allow an amount of over-speaking.

Although RSVP is unaware of which service(Controlled-Load or Guaranteed) reservations refer to, RSVP is able to identify those points in the distribution tree that require reshaping in the event that the reservations are for Guaranteed Service as described in section 2.1. Consequently at all such points RSVP informs the traffic control mechanisms within the appropriate router accordingly although such action will only result in reshaping if the reservation is actually for Guaranteed Service.

3.6 Slack term

When a receiver generates an Rspec for a Resv message to be sent for a Guaranteed Service reservation request it must include a Slack Term, S(ms) as well as the amount of bandwidth, R to be installed in each router along the path. S represents the amount by which the end-to-end delay bound will be below the end-to-end delay required by the application assuming each router along the path reserves R bandwidth according to the Guaranteed Service fluid approximation. Inclusion of a non-zero Slack Term offers the individual routers greater flexibility in making their local reservations. In certain circumstances this greater flexibility could increase the chance of an end-to-

end reservation being successful. Some routers have deadline based schedulers that decouple rate and delay guarantees. Such a scheduler may sometimes be unable to meet it's deadline requirement for Guaranteed Service in which case it might still be able to accept the reservation providing the Slack Term is at least as large as the excess delay. The excess delay would then be subtracted from the Slack Term before unicasting the Resv message to the previous hop upstream. Similarly a rate based scheduler might be able to admit a reservation request by reserving less than the requested bandwidth and unicasing the reduced reservation request to a previous hop upstream provided it could extract enough slack. Any router using available slack to reduce it's reservation must conform to the rules in equation (3) to ensure that the end-to-end delay bound remains satisfied.

$$Sout + \frac{b}{Rout} + \frac{Ctoti}{Rout} \leq Sin + \frac{b}{Rin} + \frac{Ctoti}{Rin} \qquad r \leq Rout \leq Rin \qquad (3)$$

where:

Ctoti is the cumulative sum of the error terms, C for all the routers that are upstream of, and including the current element, i.

(Rin, Sin) is the reservation request received by router, i.

(Rout, Sout) is the modified reservation request unicast to the previous hop router upstream.

An example of how intelligent use of the Slack Term can increase the probability of an end-to-end reservation request being accepted is illustrated in Figures 5 and 6. Suppose the token bucket rate of the data to be sent is 1.5Mb/s and the receiver has calculated from the Tspec and Adspec parameters in received Path messages that the desired end-to-end delay can be achieved by a reservation of (R=2.5Mb/s, S=0) which is then requested in Figure 5. However because R3 only has 2Mb/s of unused bandwidth and there is no slack available the reservation is denied. In Figure 6 the reservation is increased to R=3Mb/s and the amount by which such a reservation would be within the required delay bound is put in the Slack Term(S>0). R5 and R6 reserve the requested 3Mb/s. R3 can only reserve a value of 2Mb/s which if used as the new reservation value in the propagated Resv message will cause an increase in the end to end delay bound. R3 can calculate this increase, di and if it is less than the value of the Slack Term, S1 in the received Resv message then the request can be accepted and a reservation of 2Mb/s installed in R3. R3 will then set the Rspec in the Resv message to (R=2Mb/s, S2=S1-di) before unicasting it to the next hop upstream which results in R2 and R1 also reserving 2Mb/s.The end-to-end delay bound of the reserved path is now no greater than for a reservation of 2.5Mb/s in every router if that were possible.

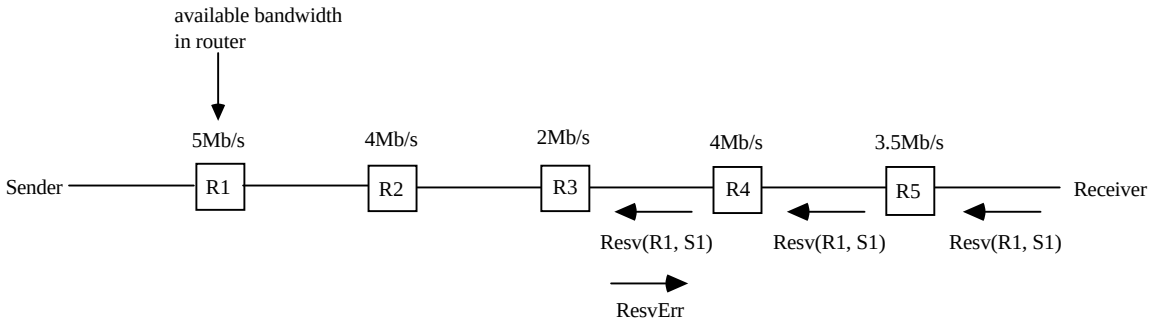

Figure 5 R=2.5Mb/s, S1=0. Reservation request denied.

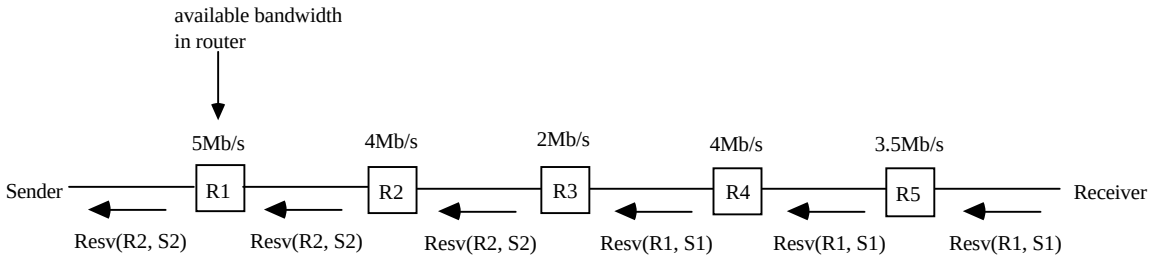

Figure 6 R1=3Mb/s, S1>0, R2=2Mb/s, S2<S1. Reservation accepted.

6.0 Summary

In this tutorial we have looked at the Controlled-Load and Guaranteed Service classes that if supported by the routers along an end-to-end data path can provide end applications with enhanced QOS commitments over conventional best-effort delivery. RSVP can be used by end applications to select and invoke the appropriate class and QOS level. In addition if the OPWA reservation model is used with RSVP then the requesting application is able to determine the resultant end-to-end QOS in advance of making the reservation.

5.0 References

[1] F. Baker, R. Guerin and D. Kandlur. Specification of Committed Rate Quality of Service, Internet Draft, June 1996, ftp://ds.internic.net/internet-drafts/draft-ietf-intserv-commit-rate-svc-00.txt.

[2] R. Braden, D. Clark and S. Shenker. Integrated Services in the Internet Architecture: an Overview, Request for Comments, July 1994, ftp://ds.internic.net/rfc/rfc1633.txt.

[3] R. Braden, L. Zhang, S. Berson, S. Herzog and S. Jamin. Resource Reservation Protocol (RSVP) - Version 1 Functional Specification, August 12, 1996. Available via http://www.ietf.org/html.charters/intserv-charter.html.

[4] D. Clark. The Design Philosophy of the DARPA Internet Protocols, ACM SIGCOMM '88, August 1988.

[5] L. Delgrossi and L. Berger. Internet Stream Protocol Version 2 (ST2) Protocol Specification - Version ST2+, Request for Comments. August 1995, RFC1819, ftp://ds.internic.net/rfc/rfc1190.txt.

[6] J. Heinanen. Protected Best Effort Service, Internet Draft, February 1996, ftp://ds.internic.net/internet-drafts/draft-heinanen-pbe-svc-01.txt.

[7] D. Mitzel, D. Estrin, S. Schenker, L. Zhang. An Architectural Comparison of ST-II and RSVP, Proceedings of Infocom '94, http://www.isi.edu/div7/rsvp/pub.html.

[8] A. Parekh and R. Gallagher. A Generalized Processor Sharing Approach to Flow Control-The Single Node Case, In IEEE/ACM Transactions on Networking, 1(3), pp. 366-357, 1993.

[9] A. Parekh and R. Gallagher. A Generalized Processor Sharing Approach to Flow Control-The Multiple Node Case, In IEEE/ACM Transactions on Networking, 2(2), pp. 137-150, 1996.

[10] S. Schenker, C.Partridge, R.Guerin. Specification of Guaranteed Quality of Service, Internet Draft, August 1996, ftp://ds.internic.net/internet-drafts/draft-ietf-intserv-guaranteed-svc-06.txt.

[11] C. Topolcic. Experimental Internet Stream Protocol, Version 2 (ST-II), Request for Comments, October 1990, RFC1190, ftp://ds.internic.net/rfc/rfc1190.txt.

[12] J. Wroclawski. Specification of the Controlled-Load Network Element Service, Internet Draft, August 1996, ftp://ds.internic.net/internet-drafts/draft-ietf-intserv-ctrl-load-svc-03.txt.

[13] J. Wroclaski. The Use of RSVP with IETF Integrated Services, Internet Draft, August 1996, ftp://ds.internic.net/internet-drafts/draft-ietf-intserv-rsvp-use-00.txt.

[14] IETF Home Page, http://www.ietf.cnri.reston.va.us.

[15] Integrated Services Charter, http://www.ietf.org/html.charters/intserv-charter.html.

[16] Integrated Services over Specific Link Layers (issl) Charter, http://www.ietf.org/html.charters/issl-charter.html.