

Multimedia Over IP: RSVP, RTP, RTCP, RTSP

[Chunlei Liu, liu.223@osu.edu](mailto:liu.223@osu.edu)

The future Integrated Services Internet will provide means to transmit real-time multimedia data across networks. RSVP, RTP, RTCP and RTSP are the foundation of real-time services. This paper is a detailed survey of the four related protocols.

[Other Reports on Recent Advances in Networking](#)

[Back to Raj Jain's Home Page](#)

Table of the Contents

- [Multimedia Networking: Goals and Challenges](#)
 - [The real-time challenge](#)
 - [Multimedia over Internet](#)
 - [The solution](#)
- [RSVP](#)
 - [Development](#)
 - [How does RSVP work](#)
 - [RSVP features](#)
 - [RSVP interfaces](#)
- [RTP](#)
 - [Development](#)
 - [How does RTP work](#)
 - [RTP fixed header fields](#)
 - [RTCP](#)
 - [RTP features](#)
 - [RTP Implementation Resources](#)
- [RTSP](#)
 - [Development](#)
 - [RTSP operations and methods](#)

- [RTSP features](#)
- [RTSP implementation resources](#)
- [Summary](#)
- [Annotated Bibliography](#)
 - [RSVP bibliography](#)
 - [RTP and RTCP bibliography](#)
 - [RTSP bibliography](#)

[Other Reports on Recent Advances in Networking](#)

[Back to Raj Jain's Home Page](#)

Multimedia Networking: Goals and Challenges

Computer networks were designed to connect computers on different locations so that they can share data and communicate. In the old days, most of the data carried on networks was textual data. Today, with the rise of multimedia and network technologies, multimedia has become an indispensable feature on the Internet. Animation, voice and video clips become more and more popular on the Internet. Multimedia networking products like Internet telephony, Internet TV, video conferencing have appeared on the market. In the future, people would enjoy other multimedia products in distance learning, distributed simulation, distributed work groups and other areas.

For networkers, multimedia networking is to build the hardware and software infrastructure and application tools to support multimedia transport on networks so that users can communicate in multimedia. Multimedia networking will greatly boost the use of computer as a communication tool. We believe someday multimedia networks will replace telephone, television and other inventions that had dramatically changed our life.

The real-time challenge

However, multimedia networking is not a trivial task. We can expect at least three difficulties. First, compared with traditional textual applications, multimedia applications usually require much higher bandwidth. A typical piece of 25 second 320x240 QuickTime movie could take 2.3MB, which is equivalent to about 1000 screens of textual data. This is unimaginable in the old days when only textual data is transmitted on the net.

Second, most multimedia applications require the real-time traffic. Audio and video data must be played back continuously at the rate they are sampled. If the data does not arrive in time, the playing back process will stop and human ears and eyes can easily pick up the artifact. In Internet telephony, human beings can tolerate a latency of about 250 millisecond. If the latency exceeds this limit, the voice will sound like a call routed over a long satellite circuit and users will complain about the quality of the call. In addition to the delay, network congestion also has more serious effects on real-time traffic. If the

network is congested, the only effect on non-real-time traffic is that the transfer takes longer to complete, but real-time data becomes obsolete and will be dropped if it doesn't arrive in time. If no proper reaction is not taken, the retransmission of lost packets would aggravate the situation and jam the network.

Third, multimedia data stream is usually bursty. Just increasing the bandwidth will not solve the burstiness problem. For most multimedia applications, the receiver has a limited buffer. If no measure is taken to smooth the data stream, it may overflow or underflow the application buffer. When data arrives too fast, the buffer will overflow and the some data packets will be lost, resulting in poor quality. When data arrives too slow, the buffer will underflow and the application will starve.

Contrary to the high bandwidth, real-time and bursty traffic of multimedia data, in real life, networks are shared by thousands and millions of users, and have limited bandwidth, unpredictable delay and availability. How to solve these conflicts is a challenge multimedia networking must face.

The possibility of answering this challenge comes from the existing network software architecture and fast developing hardware. The basis of Internet, TCP/IP and UDP/IP, provides a range of services that multimedia applications can use. Fast networks like Gigabit Ethernet, FDDI, and ATM provide high bandwidth required by digital audio and video.

So the design of real-time protocols for multimedia networking becomes imperative before the multimedia age comes.

Multimedia over Internet

There are other ways to transmit multimedia data, like dedicated links, cables and ATM. However, the idea of running multimedia over Internet is extremely attractive.

Dedicated links and cables are not practical because they require special installation and new software. Without an existing technology like LAN, WAN, the software development will be extremely expensive. ATM was said to be the ultimate solution for multimedia because it supports very high bandwidth, is connection-oriented and can tailor different level of quality of service to different type of applications. But at this moment, very few users have ATM networks reaching their organization, even fewer have ATM connections to their desktops.

On the other hand, the Internet is growing exponentially. The well established LAN and WAN technologies based on IP protocol suite connect bigger and bigger networks all over the world to the Internet. In fact, Internet has become the platform of most networking activities. This is the primary reason to develop multimedia protocols over Internet. Another benefit of running multimedia over IP is that users can have integrated data and multimedia service over one single network, without investing on another network hardware and building the interface between two networks.

At current time, IP and Ethernet seem to be more favored in the desktops and LANs, with ATM in wide area networks.

As a shared datagram network, Internet is not naturally suitable for real-time traffic. To run multimedia over Internet, several issues must be solved. First, multimedia means extremely dense data and heavy traffic. The hardware has to provide enough bandwidth.

Second, multimedia applications are usually related to multicast, i.e., the same data stream, not multiple copies, is sent a group of receivers. For example, in video conference, the video data need to be sent to all participants at the same time. Live video can be sent to thousands of recipients. The protocols designed for multimedia applications must take into account multicast in order to reduce the traffic.

Third, the price tag attached shared network resources is unpredictable availability. But real-time applications require guaranteed bandwidth when the transmission takes place. So there must be some mechanisms for real-time applications to reserve resources along the transmission path.

Fourth, Internet is a packet-switching datagram network where packets are routed independently across shared networks. The current technologies cannot guarantee that real-time data will reach the destination without being jumbled and jerky. Some new transport protocols must be used to take care of the timing issues so that audio and video data can be played back continuously with correct timing and synchronization.

Fifth, there should be some standard operations for applications to manage the delivery and present the multimedia data.

The answers to the above issues are the protocols that will be discussed in this paper.

The solution

The Internet carries all types of traffic, each type has different characteristics and requirements. For example, a file transfer application requires that some quantity of data is transferred in an acceptable amount of time, while Internet telephony requires that most packets get to the receiver in less than 0.3 seconds. If enough bandwidth is available, best-effort service fulfills all of these requirements. When resources are scarce, however, real-time traffic will suffer from the congestion.

The solution for multimedia over IP is to classify all traffic, allocate priority for different applications and make reservations. The Integrated Services working group in the IETF (Internet Engineering Task Force) developed an enhanced Internet service model called Integrated Services that includes best-effort service and real-time service, see RFC 1633. The real-time service will enable IP networks to provide quality of service to multimedia applications. Resource ReReservation Protocol (RSVP), together with Real-time Transport Protocol (RTP), Real-Time Control Protocol (RTCP), Real-Time Streaming Protocol (RTSP), provides a working foundation for real-time services. Integrated Services allows applications to configure and manage a single infrastructure for multimedia applications and traditional applications. It is a comprehensive approach to provide applications with the type of service they need and in the quality they choose.

This paper, which takes many materials from corresponding Internet Drafts and RFCs, is a detailed review of the four protocols.

[Go back to the table of contents.](#)

RSVP --- Resource ReSerVation Protocol

RSVP is the network control protocol that allows data receiver to request a special end-to-end quality of service for its data flows. Real-time applications use RSVP to reserve necessary resources at routers along the transmission paths so that the requested bandwidth can be available when the transmission actually takes place. RSVP is a main component of the future Integrated Services Internet which can provide both best-effort and real-time service[ISInt93].

Development

RSVP design has been a joint effort of [Xerox Corp.'s Palo Alto Research Center \(PARC\)](#), [MIT](#), and [Information Sciences Institute of University of California \(ISI\)](#). The RSVP specification was submitted to the [Internet Engineering Steering Group \(IESG\)](#) for consideration as a Proposed RFC in November 1994. In September 1997, [RSVP Version 1 Functional Specification](#) and several other related internet drafts were approved as Proposed Standards They are:

- [RFC 2205](#), Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification
- [RFC 2206](#), RSVP Management Information Base using SMIV2 (RFC 2206)
- [RFC 2207](#), RSVP Extensions for IPSEC Data Flows
- [RFC 2208](#), RSVP Version 1 Applicability Statement Some Guidelines on Deployment
- [RFC 2209](#), RSVP Version 1 Message Processing Rules

The [RSVP working group](#) of the [IETF](#) is developing other protocols to be used with RSVP.

How does RSVP work?

RSVP is used to set up reservations for network resources. When an application in a host (the data stream receiver) requests a specific quality of service (QoS) for its data stream, it uses RSVP to deliver its request to routers along the data stream paths. RSVP is responsible for the negotiation of connection parameters with these routers. If the reservation is setup, RSVP is also responsible for maintaining router and host states to provide the requested service.

Each node capable of resource reservation has several local procedures for reservation setup and enforcement (see Figure 1). *Policy control* determines whether the user has administrative permission to make the reservation. In the future, authentication, access control and accounting for reservation will also be implemented by policy control. *Admission control* keeps track of the system resources and determines whether the node has sufficient resources to supply the requested QoS.

The RSVP daemon checks with both procedures. If either check fails, the RSVP program returns an error notification to the application that originated the request. If both checks succeed, the RSVP daemon sets parameters in the packet classifier and packet scheduler to obtain the requested QoS. The *packet classifier* determines the QoS class for each packet and the *packet scheduler* orders packet transmission to achieve the promised QoS for each stream.

RSVP daemon also communicate with the routing process to determine the path to send its reservation

requests and to handle changing memberships and routes.

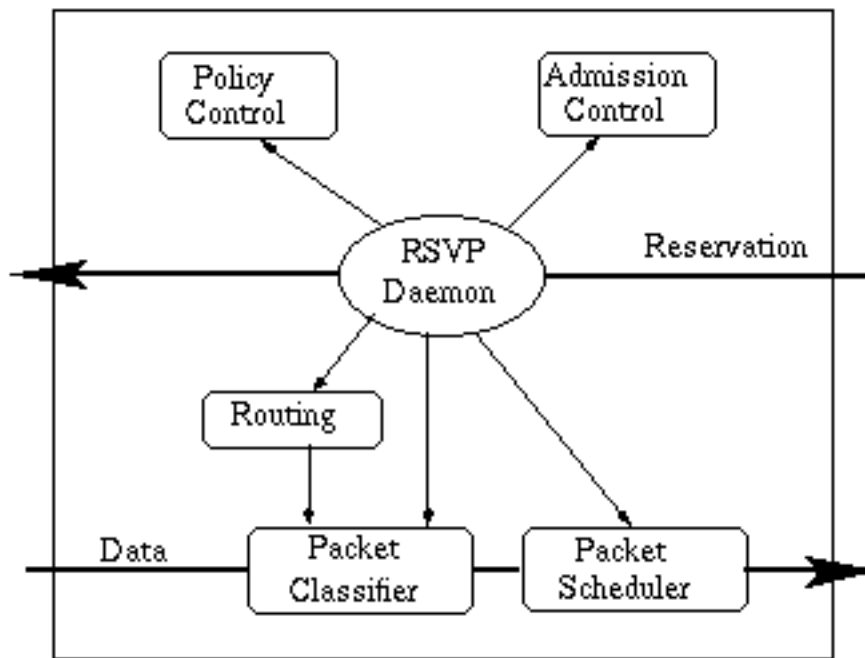


Figure 1: reservation at a node on the data flow path

This reservation procedure is repeated at routers along the reverse data stream path until the reservation merges with another reservation for the same source stream.

Reservations are implemented through two types of RSVP messages: PATH and RESV. The PATH messages are sent periodically from the sender to the multicast address. A PATH message contains *flow spec* to describe sender template (data format, source address, source port) and traffic characteristics. This information is used by receivers to find the reverse path to the sender and to determine what resources should be reserved. Receivers must join the multicast group in order to receive PATH messages.

RESV messages are generated by the receivers and contains reservation parameters including *flow spec* and *filter spec*. The filter spec defines what packets in the flow should be used by the packet classifier. The flow spec is used in packet scheduler and its content depends on the service. RESV messages follow the exact reverse path of PATH messages, setting up reservations for one or more senders at every node.

The reservation states RSVP builds at the routers are *soft states*. The RSVP daemon needs to send refresh messages periodically to maintain the reservation states. The absence of refresh message within a certain time will destroy the reservation state. By using soft states, RSVP can easily handle changing memberships and routes.

The reservation requests are initiated by the receivers. They do not need to travel all the way to the source of the sender. In stead, it travels upstream until it meets another reservation request for the same source stream, then merges with that reservation. Figure 2 shows how the reservation requests merge as they progress up the multicast tree.




Figure 2: reservation merging.

This reservation merging leads to the primary advantage of RSVP: scalability---a large number of users can be added to a multicast group without increasing the data traffic significantly. So RSVP can scale to large multicast groups and the average protocol overhead decreases as the number of participants increases.

The reservation process does not actually transmit the data and provide the requested quality of service. But through reservation, RSVP guarantees the network resources are available when the transmission actually takes place.

Although RSVP sits on top of IP in the protocol stack, it is not a routing protocol, but rather an internet control protocol. Actually, RSVP relies on the underlying routing protocols to find where it should deliver the reservation requests. RSVP is also intended to cooperate with unicast and multicast routing protocols. When the RSVP-managed flow changes its path, the routing module will notify the RSVP module of the route changes. Therefore, RSVP can quickly adjust the resource reservation to new routes.

The delivery of reservation parameters is different from the determination of these parameters. How to set the connection parameters to achieve the requested QoS is the task of QoS control devices, the role of RSVP is just a general facility to distribute these parameters. Since different applications may have different QoS control devices, RSVP is designed to treat these QoS parameters as opaque data to be delivered to and interpreted by the control modules at the routers. This logical separation of QoS control devices and distribution facility simplifies the design of RSVP and makes it more adaptive to new network technologies and applications.

RSVP Features

- RSVP flows are simplex.

RSVP distinguishes senders and receivers. Although in many cases, a host can act both as a sender and as a receiver, one RSVP reservation only reserves resources for data streams in one direction.

- RSVP supports both multicast and unicast, and adapts to changing memberships and routes.

RSVP is designed for both multicast and unicast. Since the reservations are initiated

by the receivers and the reservation states are soft, RSVP can easily handle changing memberships and routes. A host can send IGMP (Internet Group Management Protocol) messages to join a multicast group. Reservation merging enables RSVP to scale to large multicast groups without causing heavy overhead for the sender.

- RSVP is receiver-oriented and handles heterogeneous receivers.

In heterogeneous multicast groups, receivers have different capacities and levels of QoS. The receiver oriented RSVP reservation requests facilitate the handling of heterogeneous multicast groups. Receivers are responsible for choosing its own level of QoS, initiating the reservation and keeping it active as long as it wants. The senders divide traffic in several flows, each is a separate RSVP flow with different level of QoS. Each RSVP flow is homogeneous and receivers can choose to join one or more flows. This approach makes it possible for heterogeneous receivers to request different QoS tailored to their particular capacities and requirements.

- RSVP has good compatibility.

Efforts have been made to run RSVP over both IPv4 and IPv6. It provides opaque transport of traffic control and policy control messages in order to be more adaptive to new technologies. It also provides transparent operation through non-supporting regions.

Further discussion on the objectives and general justification for RSVP design are presented in [[RSVP93](#), ISInt93].

RSVP Interfaces

RSVP communicates with both applications at the end hosts and various components inside network routers. For application programmers, the most important part is the client application programming interface. In an implementation developed by the [ISI](#) and Sun Microsystems, the following fundamental functions make up an RSVP client library called RAPI (RSVP Application Programming Interface):

- `rapi_session()` creates and initiates an API session and return a session handle for further reference.
- `rapi_sender()` is called by a sender application to specifies the parameters of its data flow. The RSVP daemon at the sender will use this flow spec to create PATH messages for this flow.
- `rapi_reserve()` is used to make, modify, or delete a reservation.
- `rapi_release()` asks the RSVP daemon to tear down the reservation.

[Go back to the table of contents.](#)

RTP --- Real-time Transport Protocol

Realtime transport protocol (RTP) is an IP-based protocol providing support for the transport of real-time data such as video and audio streams. The services provided by RTP include time reconstruction, loss detection, security and content identification. RTP is primarily designed for multicast of real-time data, but it can be also used in unicast. It can be used for one-way transport such as video-on-demand as well as interactive services such as Internet telephony.

RTP is designed to work in conjunction with the auxiliary control protocol RTCP to get feedback on quality of data transmission and information about participants in the on-going session.

Development

Attempts to send voice over networks began in early 70's. Several patents on packet transmission of speech, time stamp and sequence numbering were granted in 70's and 80's. In 1991, a series of voice experiments were completed on DARTnet. In August 1991, the [Network Research Group](#) of [Lawrence Berkeley National Laboratory](#) released an audio conference tool [vat](#) for DARTnet use. The protocol used was referred later as *RTP version 0*.

In December 1992 Henning Schulzrinne, GMD Berlin, published *RTP version 1*. It went through several states of Internet Drafts and was finally approved as an Proposed Standard on November 22, 1995 by the [IESG](#). This version was called *RTP version 2* and was published as

- [RFC 1889](#), *RTP: A Transport Protocol for Real-Time Applications*
- [RFC 1890](#), *RTP Profile for Audio and Video Conferences with Minimal Control*

On January 31, 1996, [Netscape](#) announced "Netscape LiveMedia" based on RTP and other standards. Microsoft claims that their NetMeeting Conferencing Software supports RTP. The latest extensions have been made by an [industry alliance around Netscape Inc.](#), who uses RTP as the basis of the Real Time Streaming Protocol RTSP.

How does RTP work

As discussed in the first section, Internet is a shared datagram network. Packets sent on the Internet have unpredictable delay and jitter. But multimedia applications require appropriate timing in data transmission and playing back. RTP provides timestamping, sequence numbering, and other mechanisms to take care of the timing issues. Through these mechanisms, RTP provides end-to-end transport for real-time data over datagram network.

Timestamping is the most important information for real-time applications. The sender sets the timestamp according to the instant the first octet in the packet was sampled. Timestamps increase by the time covered by a packet. After receiving the data packets, the receiver uses the timestamp to reconstruct the original timing in order to play out the data in correct rate. Timestamp is also used to synchronize different streams with timing properties, such as audio and video data in MPEG. However, RTP itself is

not responsible for the synchronization. This has to be done in the application level.

UDP does not deliver packets in timely order, so *sequence numbers* are used to place the incoming data packets in the correct order. They are also used for packet loss detection. Notice that in some video format, when a video frame is split into several RTP packets, all of them can have the same timestamp. So just timestamp is not enough to put the packets in order.

The *payload type identifier* specifies the payload format as well as the encoding/compression schemes. From this payload type identifier, the receiving application knows how to interpret and play out the payload data. Default payload types are defined in RFC 1890. Example specifications include PCM, MPEG1/MPEG2 audio and video, JPEG video, Sun CellB video, H.261 video streams, et al. More payload types can be added by providing a profile and payload format specification. At any given time of transmission, an RTP sender can only send one type of payload, although the payload type may change during transmission, for example, to adjust to network congestion.

Another function is *source identification*. It allows the receiving application to know where the data is coming from. For example, in an audio conference, from the source identifier a user could tell who is talking.

The above mechanisms are implemented through the RTP header. Figure 3 shows a RTP packet encapsulated in a UDP/IP packet.




Figure 3: RTP data in an IP packet

RTP is typically run on top of UDP to make use of its multiplexing and checksum functions. TCP and UDP are two most commonly used transport protocols on the Internet. TCP provides a connection-oriented and reliable flow between two hosts, while UDP provides a connectionless but unreliable datagram service over the network. UDP was chosen as the target transport protocol for RTP because of two reasons. First, RTP is primarily designed for multicast, the connection-oriented TCP does not scale well and therefore is not suitable. Second, for real-time data, reliability is not as important as timely delivery. Even more, reliable transmission provided by retransmission as in TCP is not desirable. For example, in network congestion, some packets might get lost and the application would result in lower but acceptable quality. If the protocol insists a reliable transmission, the retransmitted packets could possibly increase the delay, jam the network, and eventually starve the receiving application.

RTP and RTCP packets are usually transmitted using UDP/IP service. However, efforts have been made to make them transport-independent so they can be also run on CLNP(Connectionless Network Protocol), IPX (Internetwork Packet Exchange), AAL5/ATM or other protocols.

In practice, RTP is usually implemented within the application. Many issues, such as lost packet recovery

marker (M): 1 bit. Defined by a profile, the marker is intended to allow significant events such as frame boundaries to be marked in the packet stream.

payload type (PT): 7 bits. Identifies the format of the RTP payload and determines its interpretation by the application.

sequence number: 16 bits. Increments by one for each RTP data packet sent, may be used by the receiver to detect packet loss and to restore packet sequence. The initial value is randomly set.

timestamp: 32 bits. The sampling instant of the first octet in the RTP data packet. May be used for synchronization and jitter calculations. The initial value is randomly set.

SSRC: 32 bits. Randomly chosen number to distinguish synchronization sources within the same RTP session. It indicates where the data was combined, or the source of the data if there is only one source.

CSRC list: 0 to 15 items, 32 bits each. Contributing sources for the payload contained in this packet. The number of identifiers is given by the CC field.

RTCP---Real-Time Control Protocol

RTCP is the control protocol designed to work in conjunction with RTP. It is standardized in RFC 1889 and 1890. In an RTP session, participants periodically send RTCP packets to convey feedback on quality of data delivery and information of membership. RFC 1889 defines five RTCP packet types to carry control information. These five types are:

- **RR:** receiver report. Receiver reports are generated by participants that are not active senders. They contain reception quality feedback about data delivery, including the highest packets number received, the number of packets lost, inter-arrival jitter, and timestamps to calculate the round-trip delay between the sender and the receiver.
- **SR:** sender report. Sender reports are generated by active senders. In addition to the reception quality feedback as in RR, they contain a sender information section, providing information on inter-media synchronization, cumulative packet counters, and number of bytes sent.
- **SDES:** source description items. They contains information to describe the sources.
- **BYE:** indicates end of participation.
- **APP:** application specific functions. It is now intended for experimental use as new applications and new features are developed.

Through these control information packets, RTCP provides the following services:

- QoS monitoring and congestion control

This is the primary function of RTCP. RTCP provides feedback to an application about the quality of data distribution. The control information is useful to the senders, the receivers and third-party monitors. The sender can adjust its transmission based on the receiver report feedback. The receivers can determine whether a congestion is local, regional or global. Network managers can evaluate the network performance for multicast distribution.

- source identification

In RTP data packets, sources are identified by randomly generated 32-bit identifiers. These identifiers are not convenient for human users. RTCP SDES (source description) packets contain textual information called *canonical names* as globally unique identifiers of the session participants. It may include user's name, telephone number, email address and other information.

- inter-media synchronization

RTCP sender reports contain an indication of real time and the corresponding RTP timestamp. This can be used in inter-media synchronization like lip synchronization in video.

- control information scaling

RTCP packets are sent periodically among participants. When the number of participants increases, it is necessary to balance between getting up-to-date control information and limiting the control traffic. In order to scale up to large multicast groups, RTCP has to prevent the control traffic from overwhelming network resources. RTP limits the control traffic to at most 5% of the overall session traffic. This is enforced by adjusting the RTCP generating rate according to the number of participants.

RTP features

- RTP provides end-to-end delivery services for data with real-time characteristics, such as interactive audio and video. But RTP itself does not provide any mechanism to ensure timely delivery. It needs support from lower layers that actually have control over resources in switches and routers. RTP depends on RSVP to reserve resources and to provide the requested quality of service.
- RTP doesn't assume anything about the underlying network, except that it provides framing. RTP is typically run on the top of UDP to make use of its multiplexing and checksum service, but efforts have been made to make RTP compatible with other transport protocols, such as ATM AAL5 and IPv6.
- Unlike usual data transmission, RTP does not offer any form of reliability or flow/congestion control. It provides timestamps, sequence numbers as hooks for adding reliability and flow/congestion control, but how to implement is totally left to the application.
- RTP is a protocol framework that is deliberately not complete. It is open to new payload formats and new multimedia software. By adding new profile and payload format specifications, one can tailor RTP to new data formats and new applications.
- RTP/RTCP provides functionality and control mechanisms necessary for carrying real-time content. But RTP/RTCP itself is not responsible for the higher-level tasks like assembly and synchronization. These have to be done at application level.
- The flow and congestion control information of RTP is provided by RTCP sender and receiver reports.

RTP Implementation resources

RTP is an open protocol that does not provide pre-implemented system calls. Implementation is tightly coupled to the application itself. Application developers have to add the complete functionality in the application layer by themselves. However, it is always more more efficient to share and reuse code rather than starting from scratch. The RFC 1889 specification itself contains numerous code segments that can be borrowed directly to the applications. There are some implementations with source code available on the web for evaluation and educational purposes. Many modules in the source code can be usable with minor modifications. The following is a list of useful resources:

- **vat** (<http://www-nrg.ee.lbl.gov/vat/>)
- **rtptools**(<ftp://ftp.cs.columbia.edu/pub/schulzrinne/rtptools/>)
- **NeVoT**(<http://www.cs.columbia.edu/~hgs/rtp/nevot.html>)
- **RTP page** (<http://www/cs.columbia.edu/~hgs/rtp>) maintained by [Henning Schulzrinne](#) is a very complete reference.
- **RTP Library** (http://www.cs.columbia.edu/~jdrosen/rtp_api.html) provides a high level interface for developing applications that make use of the Real Time Transport Protocol (RTP).

[Go back to the table of contents.](#)

RTSP---Real-Time Streaming Protocol

Instead of storing large multimedia files and playing back, multimedia data is usually sent across the network in streams. Streaming breaks data into packets with size suitable for transmission between the servers and clients. The real-time data flows through the transmission, decompressing and playing back pipeline just like a water stream. A client can play the first packet, decompress the second, while receiving the third. Thus the user can start enjoying the multimedia without waiting to the end of transmission.

RTSP, the Real Time Streaming Protocol, is a client-server multimedia presentation protocol to enable controlled delivery of streamed multimedia data over IP network. It provides "VCR-style" remote control functionality for audio and video streams, like pause, fast forward, reverse, and absolute positioning. Sources of data include both live data feeds and stored clips.

RTSP is an application-level protocol designed to work with lower-level protocols like RTP, RSVP to provide a complete streaming service over internet. It provides means for choosing delivery channels (such as UDP, multicast UDP and TCP), and delivery mechanisms based upon RTP. It works for large audience multicast as well as single-viewer unicast.

Development

RTSP was jointly developed by [RealNetworks](#), [Netscape Communications](#) and [Columbia University](#). It was developed from the streaming practice and experience of RealNetworks' RealAudio and Netscape's LiveMedia. The first draft of RTSP protocol was submitted to IETF on October 9, 1996 for consideration

as an Internet Standard. Since then, it has gone through significant changes. The latest version is [draft-ietf-mmusic-rtsp-07](#).

Although RTSP is still an internet draft at IETF and is likely to have significant changes, products using RTSP are available today. The major online players, Netscape, Apple, IBM, Silicon Graphics, VXtreme, Sun and other companies, had claimed their support to RTSP, though Microsoft's absence is conspicuous.

RTSP operations and methods

RTSP establishes and controls streams of continuous audio and video media between the media servers and the clients. A *media server* provides playback or recording services for the media streams while a *client* requests continuous media data from the media server. RTSP is the "network remote control" between the server and the client. It provides the following operations:

- **Retrieval of media from media server:** The client can request a presentation description, and ask the server to setup a session to send the requested data.
- **Invitation of a media server to a conference:** The media server can be invited to the conference to play back media or to record a presentation.
- **Adding media to an existing presentation:** The server or the client can notify each other about any additional media becoming available.

RTSP aims to provide the same services on streamed audio and video just as HTTP does for text and graphics. It is designed intentionally to have similar syntax and operations so that most extension mechanisms to HTTP can be added to RTSP.

In RTSP, each presentation and media stream is identified by an RTSP URL. The overall presentation and the properties of the media are defined in a presentation description file, which may include the encoding, language, RTSP URLs, destination address, port, and other parameters. The presentation description file can be obtained by the client using HTTP, email or other means.

But RTSP differs from HTTP in several aspects. First, while HTTP is a stateless protocol, an RTSP server has to maintain "session states" in order to correlate RTSP requests with a stream. Second, HTTP is basically an asymmetric protocol where the client issues requests and the server responds, but in RTSP both the media server and the client can issue requests. For example the server can issue a request to set playing back parameters of a stream.

In the current version, the services and operations are supported through the following methods:

- **OPTIONS:** The client or the server tells the other party the options it can accept.
- **DESCRIBE:** The client retrieves the description of a presentation or media object identified by the request URL from the server.
- **ANNOUNCE:** When sent from client to server, ANNOUNCE posts the description of a presentation or media object identified by the request URL to a server. When sent from server to client, ANNOUNCE updates the session description in real-time.
- **SETUP:** The client asks the server to allocate resources for a stream and start an RTSP session.
- **PLAY:** The client asks the server to start sending data on a stream allocated via SETUP.

- **PAUSE:** The client temporarily halts the stream delivery without freeing server resources.
- **TEARDOWN:** The client asks the server to stop delivery of the specified stream and free the resources associated with it.
- **GET_PARAMETER:** Retrieves the value of a parameter of a presentation or a stream specified in the URI.
- **SET_PARAMETER:** Sets the value of a parameter for a presentation or stream specified by the URI.
- **REDIRECT:** The server informs the clients that it must connect to another server location. The mandatory location header indicates the URL the client should connect to.
- **RECORD:** The client initiates recording a range of media data according to the presentation description.

Note that some of these methods can be sent either from the server to the client or from the client to the server, but others can only be sent in one direction. Not all these methods are necessary in a fully functional server. For example, a media server with live feeds may not support the PAUSE method.

RTSP requests are usually sent on a channel independent of the data channel. They can be transmitted in persistent transport connections, or as a one connection per request/response transaction, or in connectionless mode.

RTSP features

- RTSP is an application level protocol with syntax and operations similar to HTTP, but works for audio and video. It uses URLs like those in HTTP.
- An RTSP server needs to maintain states, using SETUP, TEARDOWN and other methods.
- RTSP messages are be carried out-of-band. The protocol for RTSP may be different from the data delivery protocol.
- Unlike HTTP, in RTSP both servers and clients can issue requests.
- RTSP is implemented on multiple operating system platforms, it allows interoperability between clients and servers from different manufacturers.

RTSP implementation resources

Although RTSP is still an IETF draft, there are a few implementations already available on the web. The following is a collection of useful implementation resources:

- [RTSP Reference Implementation](http://www.real.com/devzone/library/fireprot/rtsp/reference.html)(<http://www.real.com/devzone/library/fireprot/rtsp/reference.html>)

This is a source code testbed for the standards community to experiment with RTSP compatibility.

- [RealMedia SDK](http://www.realityworks.com/devzone/tools/index.html)(<http://www.realityworks.com/devzone/tools/index.html>)

This is an open, cross platform, client-server system that implementors can create RTSP-based streaming applications. It includes a working RTSP client and server, as

well as the components to quickly create RTSP-based applications which stream arbitrary data types and file formats.

- [W3C's Jigsaw](http://www.w3.org/Jigsaw/)(<http://www.w3.org/Jigsaw/>)

A Java-based web server. The RTSP server in the latest beta version was written in Java.

- [IBM's RTSP Toolkit](http://www.research.ibm.com/rtsptoolkit/)(<http://www.research.ibm.com/rtsptoolkit/>)

IBM's toolkits derived from from tools developed for ATM/video research and other applications in 1995-1996. Its shell-based implementation illustrates usefulness of the RTSP protocol for non-multimedia applications.

[Go back to the table of contents.](#)

Summary

This paper discusses the four related protocols for real-time multimedia data in the future Integrated Services Internet.

RSVP is the protocol that deals with lower layers that have direct control over network resources to reserve resources for real-time applications at the routers on the path. It does not deliver the data.

RTP is the transport protocol for real-time data. It provides timestamp, sequence number and other means to handle the timing issues in real-time data transport. It relies on RVSP for resource reservation to provide quality of service.

RTCP is the control part of RTP that helps with quality of service and membership management.

RTSP is a control protocol that initiates and directs delivery of streaming multimedia data from media servers. It is the "Internet VCR remote control protocol". Its role is to provide the remote control, the actual data delivery is done separately, most likely by RTP.

[Go back to the table of contents.](#)

Annotated Bibliography

RSVP bibliography

Describes version 1 of RSVP, a resource reservation setup protocol designed for an integrated services Internet.

1. [RSVP93] Zhang, L., Deering, S., Estrin, D., Shenker, S., and D. Zappala, "RSVP: A New

Resource ReSerVation Protocol", <ftp://parcftp.xerox.com/pub/net-research/rsvp.ps.Z>, IEEE Network, Vol 7, no 5, pp 8-18, September 1993.

Novel design features lead to an Internet Protocol that is flexible and scalable. This article presents the proposal of Resource ReSerVation Protocol: its design goals, principles, and operation overview.

2. [MITZ94] Mitzel, D., Estrin, D., Shenker, S., and L. Zhang, "An Architectural Comparison of ST-II and RSVP", <ftp://catarina.usc.edu/pub/mitzel/Infocom94/infocom94.ps>, Proceedings of Infocomm '94.

Presents a comparative analysis of two resource reservation protocols, ST-II and RSVP, in support of an Integrated Services Packet Network.

3. [ISInt93] Braden, R., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview" <ftp://ftp.isi.edu/in-notes/rfc1633.txt>, RFC1633, ISI, MIT, and PARC, June 1994.

Discusses a proposed extension to the Internet architecture and protocols to provide integrated services, i.e., to support real-time as well as the current non-real-time service of IP.

4. Ursula Schwantag, An Analysis of the Applicability of RSVP, http://network-services.uoregon.edu/~ursula/diplom_lt.ps, Diploma Thesis at the Advanced Network Technology Center, University of Oregon, July 15, 1997.

A comprehensive study of RSVP and Integrated Services.

More Internet-Drafts:

Describes the format and use of RSVP's INTEGRITY object to provide hop-by-hop integrity and authentication of RSVP messages.

Presents extensions to Version 1 of RSVP that permit support of individual data flows using RFC 1826 IP Authentication Header (AH) or RFC 1827 IP Encapsulating Security Payload (ESP).

Presents a set of extensions for supporting generic policy based admission control in RSVP, including the standard format of POLICY_DATA objects, a generic RSVP/Policy-Control interface, and a description of RSVP's handling of policy events.

An algorithmic description of the rules used by an RSVP implementation for processing messages.

Presents extensions to Version 1 of the Resource Reservation Setup Protocol (RSVP), that ameliorate RSVP support of Large Scale Multicast Applications (LSMA) and Virtual Private

Networks (VPN).

Provides information for designers of tunneling protocols that use IP-in-IP encapsulation.

5. Shai Herzog, Dimitrios Pendarakis, Raju Rajan, Roch Guerin, Open Outsourcing Policy Service (OOPS) for RSVP, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-policy-oops-01.txt> (75099 bytes, 7/30/97), or PostScript version: <ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-policy-oops-01.ps> (154944 bytes).

A protocol for exchanging policy information and decisions between an RSVP-capable router (client) and a policy server.

6. Allison Mankin, Fred Baker, Bob Braden, Scott Bradner, Mike O'Dell, Allyn Romanow, Abel Weinrib, Lixia Zhang, Resource ReSerVation Protocol (RSVP) Version 1 Applicability Statement Some Guidelines on Deployment, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-intsrv-analysis-01.txt> (16313 bytes, March 1997)

Describes the applicability of RSVP along with the Integrated Services protocols and other components of resource reservation and offers guidelines for deployment of resource reservation at this time.

7. L. Breslau, S. Shenker, Partial Service Deployment in the Integrated Services Architecture, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-partial-service-00.txt> (16471 bytes, 4/23/1997)

A mechanism for dealing with heterogeneity in the set of services offered by different network elements which enables end-to-end service to be composed of different per-hop services while not requiring end systems to be aware of the details of the service provided at each hop.

8. R. Braden, D. Hoffman, RAPI -- An RSVP Application Programming Interface, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-rapi-00.txt> (50578 bytes, 6/16/1997), or PostScript version: <ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-rapi-00.ps> (135952 bytes).

Version 5 of RAPI, a specific API (application programming interface) for RSVP.

9. Jim Boyle, Ron Cohen, Laura Cunningham, David Durham, Arun Sastry, Raj Yavatkar, Protocol for Exchange of PoliCy Information (PEPCI), <ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-pepci-00.txt> (35845 bytes, 7/25/1997)
10. L. Berger, RSVP over ATM Implementation Requirements, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-issll-atm-imp-req-00.txt> (29819 bytes, 7/11/1997), or PostScript version: <ftp://ftp.ietf.org/internet-drafts/draft-ietf-issll-atm-imp-req-00.ps> (206075 bytes).

Presents specific implementation requirements for running RSVP over ATM switched virtual circuits (SVCs).

11. E. Crawley, L. Berger, S. Berson, F. Baker, M. Borden, J. Krawczyk, A Framework for Integrated

Services and RSVP over ATM,

<ftp://ftp.ietf.org/internet-drafts/draft-ietf-issll-atm-framework-00.txt> (78623 bytes, 7/24/1997)

RTP and RTCP bibliography

1. H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, rfc1889: "RTP: A Transport Protocol for Real-Time Applications", 01/25/1996. <http://www.cis.ohio-state.edu/htbin/rfc/rfc1889.html>

Describes RTP, the real-time transport protocol. RTP provides end-to-end network transport functions suitable for applications transmitting real-time data, such as audio, video or simulation data, over multicast or unicast network services.

2. H. Schulzrinne, rfc1890: "RTP Profile for Audio and Video Conferences with Minimal Control";, 01/25/1996. <http://www.cis.ohio-state.edu/htbin/rfc/rfc1890.html>

Describes a profile for the use of the real-time transport protocol (RTP), version 2, and the associated control protocol, RTCP, within audio and video multiparticipant conferences with minimal control.

3. Bettati, D. Ferrari, A. Gupta, W. Heffner, W. Howe, M. Moran, Q. Nguyen, R. Yavatkar, "Connection establishment for multi-party real-time communication", <http://http.cs.berkeley.edu/~amit/research/Postscript/estab-nosdav.ps>, Proceedings of the Fifth International Workshop on Network and Operating System Support for Digital Audio and Video, Durham, NH, April 1995.

This paper outlines the Suite 2 of the Tenet protocols that provide QoS guarantees for real-time traffic on packet switching networks.

4. D. Ferrari, A. Gupta, G. Ventre, "Distributed advance reservation of real-time connections", <http://http.cs.berkeley.edu/~amit/research/Postscript/adv-nosdav.ps>, Proceedings of the Fifth International Workshop on Network and Operating System Support for Digital Audio and Video, Durham, NH, April 1995.

This paper discusses the requirements of the clients of an advance reservation service, and a distributed design for such a service.

5. A. Gupta, W. Howe, M. Moran, and Q. Nguyen, "Resource sharing for multi-party real-time communication" <http://http.cs.berkeley.edu/~amit/research/Postscript/infocom95.ps>, Proc. IEEE INFOCOM'95, Boston, MA, April 1995.

Presents a fully distributed technique for using resource sharing to provide guaranteed performance communication in a heterogeneous in internetwork.

6. Henning Schulzrinne, RTP, <http://www.cs.columbia.edu/~hgs/rtp/>.

A very complete collection of RTP maintained by Henning Schulzrinne: overview, RFCs, drafts, papers, faq, implementations, and more.

7. Stardust Technologies, Inc, Higher level protocols used with IP multicast: an IP multicast initiative white paper, <http://www.ipmulticast.com/community/whitepapers/highprot.html>.

A fairly good introduction to selected protocols used with IP Multicast to support multimedia and reliable data transmission.

More RFCs and Internet-Drafts:

8. L. Berc, W. Fenner, R. Frederick, S. McCanne, RTP Payload Format for JPEG-compressed Video (RFC 2035), <ftp://ftp.isi.edu/in-notes/rfc2035.txt> (30079 bytes, October 1996)

The RTP payload format for JPEG video streams.

9. T. Turetti, C. Huitema, RTP payload format for H.261 video streams (RFC 2032), <ftp://ftp.isi.edu/in-notes/rfc2032.txt> (27408 bytes, October 1996)

A scheme to packetize an H.261 video stream for transport using the Real-time Transport Protocol, RTP, with any of the underlying protocols that carry RTP.

10. D. Hoffman, G. Fernando, V. Goyal, RTP Payload Format for MPEG1/MPEG2 Video (RFC 2038), <ftp://ftp.isi.edu/in-notes/rfc2038.txt> (23266 bytes, October 1996)

A packetization scheme for MPEG video and audio streams.

11. M. Speer, D. Hoffman, RTP Payload Format of Sun's CellB Video Encoding (RFC 2029), <ftp://ftp.isi.edu/in-notes/rfc2029.txt> (11216 bytes, October 1996)

A packetization scheme for the CellB video encoding.

12. Michael F. Speer, Steven McCanne, RTP usage with Layered Multimedia Streams, <ftp://ftp.ietf.org/internet-drafts/draft-speer-avt-layered-video-02.txt> (8361 bytes, 12/20/1996)

Describes how one should make use of RTP (rfc1889) when employing layered media streams.

13. C. Zhu, RTP Payload Format for H.263 Video Streams, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-rtp-payload-04.txt> (25789 bytes, 6/15/1997)

Specifies the payload format for encapsulating an H.263 bitstream in the Real-Time Transport Protocol (RTP).

14. Colin Perkins, Isidor Kouvelas, Orion Hodson, Vicky Hardman, Mark Handley, RTP Payload for Redundant Audio Data, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-rtp-redundancy-01.txt> (24118 bytes, 7/25/1997), or PostScript version:

<ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-rtp-redundancy-01.ps> (66953 bytes).

A payload format for use with the real-time transport protocol (RTP), version 2, for encoding redundant audio data.

15. M. Reha Civanlar, Glenn L. Cash, Barry G. Haskell, RTP Payload Format for Bundled MPEG, <ftp://ftp.ietf.org/draft-civanlar-bmpeg-01.txt> (12748 bytes, February, 1997)

A payload type for bundled, MPEG-2 encoded video and audio data to be used with RTP, version 2.

16. S. Casner, V. Jacobson, Compressing IP/UDP/RTP Headers for Low-Speed Serial Links, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-crtp-03.txt> (51365 bytes, 7/11/1997)

A method for compressing the headers of IP/UDP/RTP datagrams to reduce overhead on low-speed serial links. In many cases, all three headers can be compressed to 2-4 bytes.

17. Bernard Aboba, Alternatives for Enhancing RTCP Scalability, <ftp://ftp.ietf.org/internet-drafts/draft-aboba-rtpscale-02.txt> (29889 bytes, 1/28/1997)

Discusses current issues with RTCP scalability as well as describing the advantages and disadvantages of possible solutions.

18. Mark Baugher, John Du, Stan Naudus, Real-Time Transport Protocol Management Information Base, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-rtp-mib-00.txt> (47338 bytes, 3/24/1997)

Defines an experimental Management Information Base(MIB) for use with network management protocols in TCP/IP-based internets.

19. Schulzrinne, RTP Profile for Audio and Video Conferences with Minimal Control, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-profile-new-01.txt> (70722 bytes, 7/29/1997), or PostScript version, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-profile-new-01.ps> (123633 bytes).

Describes a profile called "RTP/AVP" for the use of the real-time transport protocol (RTP), version 2, and the associated control protocol, RTCP, within audio and video multiparticipant conferences with minimal control.

20. D. Hoffman, G. Fernando, V. Goyal, M. Reha Civanlar, RTP Payload Format for MPEG1/MPEG2 Video, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-mpeg-new-01.txt> (31349 bytes, June 1997)

Describes a packetization scheme for MPEG video and audio streams.

21. A. Jones, A. Periyannan, D. Singer, RTP Payload Format for QuickTime Media Streams, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-qt-rtp-00.txt> (31261 bytes, 7/22/1997)

Specifies the payload format for encapsulating QuickTime media streams in the Realtime

Transport Protocol (RTP).

22. Schulzrinne, RTP Payload for DTMF Digits, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-dtmf-00.txt> (12477 bytes, 7/8/1997), or PostScript version: <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-dtmf-00.ps> (29633 bytes).

A payload type for carrying dual-tone multifrequency (DTMF) digits in RTP packets.

23. J. Rosenberg, H. Schulzrinne, An A/V Profile Extension for Generic Forward Error Correction in RTP, <ftp://ftp.ietf.org/internet-drafts/draft-ietf-avt-fec-00.txt> (26839 bytes, July 1997)

An extension to RFC 1890 which allows for forward correction (FEC) of continuous media encapsulated in RTP.

RTSP bibliography

1. H. Schulzrinne, A. Rao, R. Lanphier: Real Time Streaming Protocol (RTSP), <http://ds.internic.net/internet-drafts/draft-ietf-mmusic-rtsp-03.txt> (135617 bytes, 7/30/97)

Describes RTSP, an application-level protocol for control over the delivery of data with real-time properties, which provides an extensible framework to enable controlled, on-demand delivery of real-time data, such as audio and video.

2. H. Schulzrinne, A real-time stream control protocol (RTSP'), <http://www.cs.columbia.edu/~hgs/rtsp/draft/draft-ietf-mmusic-stream-00.txt> (44787 bytes, 11/26/96).

A revised version of the RTSP proposal put forward to the MMUSIC group, which is intended to control multiple data delivery sessions, provide a means for choosing delivery channels such as UDP, multicast UDP and TCP, and delivery mechanisms based upon RTP (RFC 1889).

3. RTSP Resource Center : <http://www.realnetworks.com/devzone/library/rtsp/index.html>

A quite complete collection of RTSP papers, implementations, links, and events.

4. Dr. Henning Schulzrinne's RTSP page : <http://www.cs.columbia.edu/~hgs/rtsp/>

Drafts, slides and related work maintained by Dr. Henning Schulzrinne.

[Go back to the table of contents.](#)

[Other Reports on Recent Advances in Networking](#)

[Back to Raj Jain's Home Page](#)

cliu@cis.ohio-state.edu

Last Modified: January 15, 1998