

RTP

Real-time Transport Protocol

The Real-time Transport Protocol (or RTP) defines a standardized packet format for delivering audio and video over the Internet. It was developed by the Audio-Video Transport Working Group of the IETF and first published in 1996 as RFC 1889 which was made obsolete in 2003 by RFC 3550.

RTP does not have a standard TCP or UDP port on which it communicates. The only standard that it obeys is that UDP communications are done via an even port and the next higher odd port is used for RTP Control Protocol (RTCP) communications. Although there are no standards assigned, RTP is generally configured to use ports 16384-32767. RTP can carry any data with real-time characteristics, such as interactive audio and video. Call setup and tear-down is usually performed by the SIP protocol. The fact that RTP uses a dynamic port range makes it difficult for it to traverse firewalls. In order to get around this problem, it is often necessary to set up a STUN server.

It was originally designed as a multicast protocol, but has since been applied in many unicast applications. It is frequently used in streaming media systems (in conjunction with RTSP) as well as videoconferencing and push to talk systems (in conjunction with H.323 or SIP), making it the technical foundation of the Voice over IP industry. It goes along with the RTCP and it's built on top of the User Datagram Protocol (UDP). Applications using RTP are less sensitive to packet loss, but typically very sensitive to delays, so UDP is a better choice than TCP for such applications.

According to RFC 1889, the services provided by RTP include:

- Payload-type identification - Indication of what kind of content is being carried
- Sequence numbering - PDU sequence number
- Time stamping - allow synchronization and jitter calculations
- Delivery monitoring

The protocols themselves do not provide mechanisms to ensure timely delivery. They also do not give any Quality of Service (QoS) guarantees. These things have to be provided by some other mechanism.

Also, out of order delivery is still possible, and flow and congestion control are not supported directly. However, the protocols do deliver the necessary data to the application to make sure it can put the received packets in the correct order. Also, RTCP provides information about reception quality which the application can use to make local adjustments. For example if a congestion is forming, the application could decide to lower the data rate.

RTP was also published by the ITU-T as H.225.0, but later removed once the IETF had a stable standards-track RFC published. It exists as an Internet Standard (STD 64) defined in RFC 3550 (which obsoletes RFC 1889). RFC 3551 (STD 65) (which obsoletes RFC 1890) defines a specific profile for Audio and Video Conferences with Minimal Control. RFC 3711 defines the Secure Real-time Transport Protocol (SRTP) profile (actually an extension to RTP Profile for Audio and Video Conferences) which can be used (optionally) to provide confidentiality, message authentication, and replay protection for audio and video streams being delivered. Packet structure + Bits 0-1 2 3 4-7 8 9-15 16-31 0
Ver. P X CC M PT Sequence Number 32 Timestamp 64 SSRC identifier 96 ... CSRC identifiers ... 96+(CCx32)
Additional header (optional), indicates length "AHL" 96+(CCx32)
+ (Xx(AHL+16))
Data

Ver. (2 bits) indicates the version of the protocol. Current version is 2. P (one bit) is used to indicate if there are extra padding bytes at the end of the RTP packet. X (one bit) indicates if the extensions to the protocol are being used in the packet. CC (four bits) contains the number of CSRC identifiers that follow the fixed header. M (one bit) is used at the application level and is defined by a profile. If it's set, it means that the current data has some special relevance for the application. PT (7 bits) indicates the format of the payload and determines its interpretation by the application. SSRC indicates the synchronization source.

[edit] Potential further development of RTP & RTCP

The Real-time Transport Protocol (RTP) and the Real-time Transport Control Protocol (RTCP) are commonly used together. RTP is used to transmit data (e.g. audio and video) and RTCP is used to monitor QoS. The monitoring of quality of service is very important for modern applications. In large scale applications (e.g. IPTV), there is an unacceptable delay between RTCP reports, which can cause quality of service related problems.

For more information read about problems and potential further development of RTCP

Mathematical background

The equations for RTCP protocol are explained in section I. and II.A in the Optimization of Large-Scale RTCP Feedback Reporting in Fixed and Mobile Networks paper.

Structure of RTP/RTCP applications

RTP/RTCP protocols are commonly used to transport audio or audio/video data. Separate sessions are used for each media content (e.g. audio and video). The main advantage of this separation is to make it possible to receive only one part of the transmission, commonly audio data, which lowers the total bandwidth.

See also

- Real time control protocol
- Real Time Streaming Protocol (RTSP)
- Secure Real-time Transport Protocol
- Stream Control Transmission Protocol
- ZRTP

References

- Henning Schulzrinne and Stephen Casner. RTP: A Transport Protocol for Real-Time Applications. (1993) Internet Engineering Task Force, Internet Draft, October 20, 1993. The memo originating RTP; only an early draft, does not describe the current standard.
- Perkins, Colin (2003). RTP: Audio and Video for the Internet (1st ed.) Addison-Wesley. ISBN 0-672-32249-8

External links

- An Introduction to RTP for Beginners
- Henning Schulzrinne's RTP page
- GNU ccRTP
- Good explanations about RTP packet structure

RFCs

- RFC 3551, Standard 65, RTP Profile for Audio and Video Conferences with Minimal Control
- RFC 3550, Standard 64, RTP : A Transport Protocol for Real-Time Applications
- RFC 1890, Obsolete, RTP Profile for Audio and Video Conferences with Minimal Control
- RFC 1889, Obsolete, RTP : A Transport Protocol for Real-Time Applications