CAROLINA TECHNOLOGY CONSULTANTS

Tech Brief

March 2009

Architectural Issues in IT and Data Communications

Topic

RTP - The Real-time Transport Protocol

Prepared By

Tyler R. Johnson

RTP / RTCP

---

* Transport of audio / video / data

* Re-ordering of packets due to out of sequence transmissions

* Reporting of packet loss & jitter

* Runs on top of network transport such as UDP

Summary

The Real-time Transport Protocol (RTP) is an IETF standards track protocol for transporting real-time data, such as audio and video, over networks.

Discussion

RTP provides delivery of real-time data over networks, especially IP networks. It provides payload type identification, sequence numbering, time stamping and delivery monitoring. RTP works over both unicast and multicast networks.

A companion protocol defined within RFC-3550, the Real-time Transport Control Protocol (RTCP), provides additional functionality for stream management and monitoring. RTP does not include QoS mechanisms.

RTP itself is a general specification. Application for specific media types requires the specification of payload type codes and format specifications, hence the existence of RFC-3551 for audio and video in RTP.

Sequence numbers are used in RTP to allow receivers to re-order packets whose sequences were altered during transmission. Lost packets can be detected. Each RTP stream carries a single media flow. Multiple media flows can be associated by the application. RTCP timing information can be used to synchronize media streams that have become unlocked in transmission.

Synchronization Sources (SSRC) are represented by 32-bit numeric identifiers unique to each source, so that sources are not identified by IP address. Contributing Sources (CSRC) have identifiers to allow multiple streams to be combined and yet uniquely identified.

Mixers and translators can be used to extract or transcode streams in an organized way.

RTP generally uses the Network Time Protocol (NTP) as a timing source for synchronization.

RTP Headers

RTP packets are preceded by a header which describes the payload type and protocol version. It includes a sequence number for use in stream reconstruction and a time stamp for synchronization. The header also includes SSRC and CSRC data to identify streams.

RTCP

RTCP's main function is to monitor the quality of the media stream connection. The protocol uses canonical names (CNAME) for receivers to identify source participants. RTCP runs on separate transport addresses, outside of the RTP media streams.

RTCP packet types include sender reports (SR), receiver reports (RR), source description items (SDES), BYE messages and application specific messages. Reports indicate identifying information about the stream, percentage loss, absolute packet count statistics and inter-packet arrival jitter measurements. The algorithm for jitter calculation is included in the specification.

RTCP provides a view into the RTP stream performance, but does not manage it. Rather, it makes information available to the overarching application, which can itself modify RTP session parameters in order to react intelligently to network conditions.

RTCP does not provide congestion control mechanisms, since these vary by media type and application. Instead, congestion control is managed through the creation of supplementary profiles, such as RFC 3551.

Security

RTP provides no authentication mechanisms, relying instead on authentication at the application layer. This should not be viewed as a weakness, but rather as a modular feature of the protocol, since its intent is to be used in complex multimedia conferencing environments. RTP includes a weak encryption scheme based on DES in RFC 3550. This is not considered a viable approach for confidentiality.

SRTP

The Secure Real-time Transport Protocol (SRTP) is a version of RTP that uses AES encryption to encode payload data for privacy purposes, while keeping the headers in the clear for access by network-snooping applications. This approach suffers from the key distribution problem common to key-based systems.

Strategy Considerations

RTP can support layered media encodings by including higher frequency data in separate RTP streams. Receivers can then tune into the number of streams appropriate to their bandwidth availability. The use of multiple RTP streams to support scalable video coding is an emergent area and is not yet widely deployed.

RTCP is most appropriate for short term packet loss and jitter measurements. Longer term network analysis should rely on processing of RTCP receiver reports over time to form an aggregate picture of network performance. Additionally, since reports do not provide the precise timing of packet loss, time resolution may be limited to inter-report times. Obtaining resolutions higher than this may require direct access to the RTP sequence numbers and timing data (outside the scope of RTCP).

RTP CNAME resolution over NAT can create a problem for unique stream name resolution.

While RFC 3550 provides some mechanisms for simple conference management, such as user name, host, email, session termination, etc., it should not be viewed as a viable conferencing solution. Rather, it is most appropriately used as a convenient media transport protocol.

For Further Information

1. IETF RFC-3550 *RTP: A Transport Protocol for Real-time Applications*. H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson. July 2003.
2. IETF RFC-3551 *RTP Profile for Audio and Video Conferences with Minimal Control*. H. Schulzrinne, S. Casner. July 2003.
3. IETF RFC-3555 *MIME Type Registration of RTP Payload Formats*. S. Casner, P. Hoschka. July 2003.
4. IETF 3711 *The Secure Real-time Transport Protocol (SRTP)*. M. Baugher, D. McGrew, M. Naslund, E. Carrara, K. Norrman. March 2004