UNIT – IV

THE TRANSPORT LAYER

RTP

The Real-time Transport Protocol (RTP) is a network protocol for delivering audio and video over IP networks. RTP is used in communication and entertainment systems that involve streaming media, such as telephony, video teleconference applications including WebRTC, television services and web-based push-to-talk features.

RTP typically runs over User Datagram Protocol (UDP). RTP is used in conjunction with the RTP Control Protocol (RTCP). While RTP carries the media streams (e.g., audio and video), RTCP is used to monitor transmission statistics and quality of service (QoS) and aids synchronization of multiple streams. RTP is one of the technical foundations of voice over IP and in this context is often used in conjunction with a signaling protocol such as the Session Initiation Protocol (SIP) which establishes connections across the network.

RTP was developed by the Audio-Video Transport Working Group of the Internet Engineering Task Force (IETF) and first published in 1996 as RFC 1889, which was then superseded by RFC 3550 in 2003.

Overview

Research on audio and video over packet-switched networks dates back to the early 1970s. The Internet Engineering Task Force (IETF) published RFC 741 in 1977 and began developing RTP in 1992, and would go on to develop Session Announcement Protocol (SAP), the Session Description Protocol (SDP), and the Session Initiation Protocol (SIP).

RTP is designed for end-to-end, real-time transfer of streaming media. The protocol provides facilities for jitter compensation and detection of packet loss and out-of-order delivery, which are common, especially during UDP transmissions on an IP network. RTP allows data transfer to multiple destinations through IP multicast. RTP is regarded as the primary standard for audio/video transport in IP networks and is used with an associated profile and payload format. The design of RTP is based on the architectural principle known as application-layer framing where protocol functions are implemented in the application as opposed to the operating system's protocol stack.

Real-time multimedia streaming applications require timely delivery of information and often can tolerate some packet loss to achieve this goal. For example, loss of a packet in an audio application may result in loss of a fraction of a second of audio data, which can be made unnoticeable with suitable error concealment algorithms. The Transmission Control Protocol (TCP), although standardized for RTP use, is not normally used in RTP applications because TCP favors reliability over timeliness. Instead, the majority of the RTP implementations are built on the User Datagram Protocol (UDP). Other transport protocols specifically designed for multimedia sessions are SCTP and DCCP, although, as of 2012, they were not in widespread use.

RTP was developed by the Audio/Video Transport working group of the IETF standards organization. RTP is used in conjunction with other protocols such as H.323 and RTSP. The RTP specification describes two protocols: RTP and RTCP. RTP is used for the transfer of multimedia data, and the RTCP is used to periodically send control information and QoS parameters.

The data transfer protocol, RTP, carries real-time data. Information provided by this protocol includes timestamps (for synchronization), sequence numbers (for packet loss and reordering detection) and the payload format, which indicates the encoded format of the data. The control protocol, RTCP, is used for quality of service (QoS) feedback and synchronization between the media streams. The bandwidth of RTCP traffic compared to RTP is small, typically around 5%.

RTP sessions are typically initiated between communicating peers using a signaling protocol, such as H.323, the Session Initiation Protocol (SIP), RTSP, or Jingle (XMPP). These protocols may use the Session Description Protocol to specify the parameters for the sessions.

An RTP session is established for each multimedia stream. Audio and video streams may use separate RTP sessions, enabling a receiver to selectively receive components of a particular stream. The RTP and RTCP design is independent of the transport protocol. Applications most typically use UDP with port numbers in the unprivileged range (1024 to 65535). The Stream Control Transmission Protocol (SCTP) and the Datagram Congestion Control Protocol (DCCP) may be used when a reliable transport protocol is desired. The RTP specification recommends even port numbers for RTP and the use of the next odd port number for the associated RTCP session. A single port can be used for RTP and RTCP in applications that multiplex the protocols.

<u>RTCP</u>

RTCP (Real-time Transport Control Protocol) is the companion protocol to RTP, designed to provide feedback, synchronization, and control for real-time media streams such as audio and video. It does not carry media itself but ensures quality monitoring, participant identification, and synchronization across sessions

Purpose

- RTCP works alongside RTP (Real-time Transport Protocol) to manage and monitor media delivery.
- Provides feedback on transmission quality (packet loss, jitter, delay).
- Helps with synchronization between multiple streams (e.g., audio and video).
- Facilitates participant identification in conferencing systems.

Key Features

- Control, not media: RTCP packets are small and periodic, consuming ~5% of session bandwidth.
- Defined in RFC 3550 along with RTP.
- Feedback mechanism: Receivers send reports back to the sender about stream quality.
- Scalability: Works in unicast and multicast sessions.
- Types of RTCP Messages

Message Type	Purpose
SR (Sender Report)	Sent by active senders; includes transmission stats and synchronization info.
RR (Receiver Report)	Sent by receivers; provides feedback on packet loss, jitter, and delay.
SDES (Source Description)	Conveys metadata like CNAME (canonical name), user info.
BYE	Indicates a participant is leaving the session.
APP (Application-specific)	Custom messages defined by applications.

RTCP vs RTP

Aspect	RTP	RTCP
Role	Transports media (audio/video)	Provides control and feedback
Data	Carries actual media packets	Carries reports, metadata, synchronization info
Bandwidth	High (depends on media)	Low (~5% of RTP stream)
Reliability	Media delivery only	Quality monitoring & synchronization

Applications

- VoIP (Voice over IP) – monitors call quality.
- Video conferencing – synchronizes audio/video streams.
- Streaming media – ensures adaptive quality control.
- Multimedia collaboration tools – identifies participants and manages session lifecycle.

SCTP

SCTP (Stream Control Transmission Protocol) is a transport layer protocol designed to combine the reliability of TCP with the message-oriented nature of UDP. It supports multi-streaming and multi-homing, making it ideal for telephony signaling, WebRTC, and applications needing robust, low-latency communication.

Introduction

- Defined in RFC 9260 (updated in 2022).
- Operates at the transport layer (Layer 4) of the OSI model.
- Originally developed for telecommunication signaling (SS7 over IP) but now widely used in VoIP, WebRTC, and mobile networks

Key Features

- Connection-oriented: Like TCP, SCTP establishes a reliable connection between endpoints.
- Message-oriented: Unlike TCP's byte stream, SCTP preserves message boundaries (similar to UDP).

- Reliable delivery: Ensures in-order delivery with retransmission and error detection (CRC32 checksum).
- Multi-streaming: Allows multiple independent streams within a single connection, reducing head-of-line blocking.
- Multi-homing: Supports multiple IP addresses per endpoint for redundancy and failover.
- Congestion control: Similar to TCP, preventing network overload.
- Security: Includes protection against SYN flooding attacks with a four-way handshake.

Types of Chunks in SCTP

SCTP uses chunks instead of segments (TCP) or datagrams (UDP). Common chunk types:

- DATA: Carries user messages.
- INIT / INIT-ACK: Used in the four-way handshake to establish a connection.
- SACK (Selective ACK): Provides acknowledgment of received chunks.
- HEARTBEAT: Tests reachability of an endpoint.
- SHUTDOWN / SHUTDOWN-ACK: Graceful connection termination.
- ERROR: Reports problems in communication
- SCTP vs TCP vs UDP

Feature	TCP	UDP	SCTP
Reliability	✅ Yes	❌ No	✅ Yes
Message-oriented	❌ No	✅ Yes	✅ Yes
Multi-streaming	❌ No	❌ No	✅ Yes
Multi-homing	❌ No	❌ No	✅ Yes
Congestion control	✅ Yes	❌ No	✅ Yes
Security handshake	3-way	None	4-way

Applications

- Telephony signaling (SS7 over IP) – M3UA, M2UA, SUA protocols.
- Mobile networks (3G/4G/5G) – secure signaling and roaming.

- WebRTC – data channels use SCTP for reliable message delivery.
- Military and industrial systems – where redundancy and reliability are critical.