

Describing multiple RTP media streams in SDP

draft-even-mmusic-multiple-streams-02
IETF 86

Roni Even

Jonathan Lennox

Qin Wu



Introduction

- RFC3550 allows multiplexing of multiple RTP streams in a single RTP session.
- SDP RFC4566 specifies that when a list of payload type numbers is given, this implies that all of these payload formats MAY be used in the session, but the first of these formats SHOULD be used as the default format for the session.
- RFC3264 uses the term “media stream” stating that “In SDP, a media stream is described by an "m=" line and its associated attributes.” but an m-line describes an RTP session which can include multiple RTP streams.

Introduction

- When looking at the following offer
 - m=video 10000 RTP/AVP 31 32
 - a=rtpmap:31 H261/90000
 - a=rtpmap:32 MPV/90000
- Is this an offer for:
 - one RTP session with H.261 or MPV codecs for the same content – this is what many systems today expect to happen.
 - one RTP session is offered with multiple H.261 and MPV codecs each with different content.
- This offer should really mean "arbitrarily many streams, with potentially different content, any of which could use either H.261 or MPV, potentially switching dynamically between them."

Introduction

- SSRCs usage in different RTP multipoint topologies:
 - The SSRCs are static (assigned by the MCU/Mixer), and there is an SSRC for each media encoding. Source information may be conveyed using CSRC, or, in the case of topo-RTCP-Terminating MCU, is not conveyed.
 - The SSRCs are dynamic, representing the original source and are relayed by the Mixer/MCU to the participants.
- The second SSRC usage make it difficult to specify the SSRCs as part of SDP. It brings the following requirement
 - It should be possible to support receipt of multiple RTP sources without explicit per-source signaling or negotiation.

Proposal

- Use multiple m-lines each describing one or multiple RTP media encodings.
 - Each m-line will specify the maximum number of SSRCs that can be sent or received using this m-line.
 - The semantics of the different media encoding may be sent using other mechanism, for example, CLUE WG propose to send information about the content of the stream using a separate data channel.
 - Similar streams can be added or removed implicitly without requiring more signaling.
 - This option allows the SDP to use a single RTP media encoding per m-line if there is a need to specify specific attributes that cannot be described in a single m-line and bundle the m-lines.