

Source-Specific SDP Attributes: RTP Issues

Jonathan Lennox

draft-lennox-mmusic-sdp-source-attributes-00.txt

Source-Specific Attributes: Review

- RTP allows multiple sources in an RTP session, but SDP has no way to signal this.
- Solution: define an SDP attribute for characteristics of a source.

```
m=video 49170 RTP/AVP 96
a=rtpmap:96 H264/90000
a=ssrc:12345 cname:user@example.com
a=ssrc:12345 information:Main camera
a=ssrc:67890 cname:user@example.com
a=ssrc:67890 information:Alternate camera
```

- Map SDP “source-specific” attributes into the `ssrc` attribute.
- This generalizes material that was previously in the RTP Single-Source Multicast draft.

Motivation

- Avoid clashes with the SSRC id of a single media sender.
 - This is needed for Single-Source Multicast.
- Make SSRC multiplexing explicit.
 - Describe, and differentiate between, multiple SSRCs from the same participant in the same RTP session.
 - Examples:
 - Multiple cameras
 - FEC
 - Retransmission
 - Layered codecs

Terminology Confusion: SDP vs. RTP

- SDP and RTP terminology is inconsistent.
- An SDP *Media Stream* describes an RTP *Session*.

SDP	RTP
Multimedia Session	–
Media Stream, Media Description, m= line	RTP Session
<i>Media Source</i>	(Audio/Video) Stream, RTP Stream, RTP Source, ssrc

- Introduce the term *Media Source* for SDP.

Implications for RTP: Collisions

- SSRCs can change due to collision.
- This draft: signaled SSRCs **MUST NOT** be re-used by other entities.
- However, collisions can still happen:
 - Race conditions.
 - Endpoints that don't understand this draft.
- Thus, endpoints **MUST** be prepared to handle collisions:
 - Receive packets with non-signaled SSRCs.
 - Receive RTCP mapping a SSRC to a CNAME other than the one signaled.

Implications for RTP: Collisions (2)

- If a signaled source collides with a non-signaled source, the signaled one should “win”:
 - It waits $5 \cdot 1.5 \cdot T_d$ before executing collision resolution.
 - It only changes its SSRC if the other colliding source doesn't go away.
- Colliding signaled sources resolve normally.
 - Shouldn't ever happen, assuming SDP transport protects from offer/answer glare.
- If you do have to change your SSRC, you can send updated SDP with a previous-ssrc attribute.
 - If you're distributing SDP over a channel which supports updated descriptions.

Backward Compatibility

- Many (point-to-point) endpoints won't correctly handle RTP sessions with more than one source.
- Even endpoints which correctly interpret multiple sources might not have the resources to decode them.
- Does there need to be a way to indicate "I can receive multiple sources"?
- If so, recommend that if your receiver doesn't indicate this, you **SHOULD** send only one source at a time.

Source grouping

- Signaled sources are useful for groups of RTP sources.
 - E.g. FEC, RTX, layered codecs.
 - Associate supplemental source with main source.
- Mechanism is analogous to SDP media grouping (a=group).
- Senders still need to be prepared for receivers that don't understand signaled sources.
 - E.g., for SSRC-multiplexed RTP retransmission (RFC 4588) of multiple sources, the mapping between the original source and the retransmission source is initially ambiguous without source signaling.
 - Senders can choose the original sources' RTP sequence numbers so that NACKers won't have NACKs outstanding for the same sequence number in two streams.

RTP Architectural Issues

- Still MUST NOT multiplex (e.g.) audio and video streams on a single RTP session.
- Still a single set of payload type numbers per RTP session.
- If you only know about a source from SDP, it doesn't count for group size estimation, and you don't send SR/RR report blocks for it.

AVT Working Group

- Have any RTP architectural issues been overlooked?
- Are there any AVT objections to this proceeding in MMUSIC?
- AVT will remain responsible for reviewing the draft.
- RTCP-SSM will normatively depend on this.