Codec-Independent Selective Forwarding

draft-aboba-avtcore-sfu-rtp-00

Bernard Aboba

Microsoft

Imagine if initial routers had only been able to forward telnet packets.

In order to support each new application protocol (FTP, Web, email, etc.), new router code would need to be developed, tested and deployed

This is the situation we face today with Selective Forwarding Units (SFUs) and (new) codecs.

SFUs that originally were developed to support H.264/SVC (or VP8) are now facing demands to add support for HEVC, VP9,

Proposition: It should be possible to build SFUs supporting SVC, with codec-independent forwarding and control planes.

SFUs could interoperate with endpoint implementations supporting RTP-usage. Support for a new codec would only require endpoint

This is an RTP architecture issue, suitable for discussion in AVTCORE.

The draft describes issues that could require new work. We will find more if we continue to look.

Starting a journey with a roadmap is often a good idea.

Forwarding Plane Issues

- Frame marking (discussion later in AVTEXT)
  - Draft describes a core set of payload information used in SFU forwarding plane implementations.
    - Examples: Frame Type (e.g. IDR), Discardable, Layer Identifier, TL0PICIDX, S/E, etc.
  - The draft also describes information that SFUs can do without.
    - Example: PictureID (can use RTP header fields such as timestamp/sequence number instead).
- MTU mismatches
  - Conference participants can have differing MTUs (e.g. due to VPN tunnels).
  - A codec-independent SFU does not perform NAL unit re-packetization, so it cannot address mismatches.
  - Endpoint choices: IP fragmentation (and loss), or MTU discovery.

Control Plane Issues

- Bandwidth estimation
  - This is core to SFU operation, but has not been standardized.
- Adding/dropping layers
  - Today, handled via codec-specific messages (e.g. SEI Layer Drop in H.264/SVC).
  - PAUSE/RESUME only a partial substitute (only applies to MRST SVC implementations).
- Codec-specific RTCP feedback messages.
  - RTP-Usage document describes core feedback messages (NACK, PLI, FIR)
  - Codec-specific feedback messages are problematic:
    - Examples: RPSI (requires PictureId) and SLI (requires macroblock #).
- RTCP SRs and RRs
  - Per-layer information only possible with MRST, not SRST.

Call to action

Consider addition of a work item on this topic.