AVTCORE WG
IETF 112

Virtual Meeting
Friday, November 12, 2021
12:00 - 14:00 UTC
04:00 - 06:00 Pacific Time

Mailing list: avtcore@ietf.org
Jabber Room: avtcore@jabber.ietf.org
MeetEcho link: https://wws.conf.meetecho.com/conference/?group=avtcore

IETF 112 Meeting Tips

https://www.ietf.org/how/meetings/112
https://datatracker.ietf.org/meeting/agenda

This session is being recorded

- **IETF 112 registration and a datatracker login required to attend**
- **No need to manually fill in blue sheets, it's automatic.**
- **Join the session Jabber room via IETF Datatracker Meeting agenda**
- **Please use headphones when speaking to avoid echo.**
- **Please state your full name before speaking.**

IETF 112 Meeting Tips

https://www.ietf.org/how/meetings/112
https://datatracker.ietf.org/meeting/agenda

This session is being recorded

- **Enter the queue with , leave with**
- **When you are called on, you need to enable your audio to be heard.**
- **Audio is enabled by unmuting and disabled by muting**
- **Video can also be enabled, but it is separate from audio.**
- **Video is encouraged to help comprehension but not required.**

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (https://www.ietf.org/contact/ombudsteam/) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- BCP 9 (Internet Standards Process)
- BCP 25 (Working Group processes)
- BCP 25 (Anti-Harassment Procedures)
- BCP 54 (Code of Conduct)
- BCP 78 (Copyright)
- BCP 79 (Patents, Participation)
- https://www.ietf.org/privacy-policy/(Privacy Policy)

Note really well

- IETF meetings, virtual meetings, and mailing lists are intended for professional collaboration and networking, as defined in the IETF Guidelines for Conduct (RFC 7154), the IETF Anti-Harassment Policy, and the IETF Anti-Harassment Procedures (RFC 7776). If you have any concerns about observed behavior, please talk to the Ombudsteam, who are available if you need to confidentially raise concerns about harassment or other conduct in the IETF.
- The IETF strives to create and maintain an environment in which people of many different backgrounds are treated with dignity, decency, and respect. Those who participate in the IETF are expected to behave according to professional standards and demonstrate appropriate workplace behavior.
- IETF participants must not engage in harassment while at IETF meetings, virtual meetings, social events, or on mailing lists. Harassment is unwelcome hostile or intimidating behavior -- in particular, speech or behavior that is aggressive or intimidates.
- If you believe you have been harassed, notice that someone else is being harassed, or have any other concerns, you are encouraged to raise your concern in confidence with one of the Ombudspersons.

About this meeting

- Agenda: https://datatracker.ietf.org/doc/agenda-112-avtcore/
- Notes: https://notes.ietf.org/s/notes-ietf-112-avtcore
- Jabber Room: avtcore@jabber.ietf.org
- Secretariat: mtd@jabber.ietf.org
- WG Chairs: Jonathan Lennox & Bernard Aboba
- Jabber Scribe:
- Note takers:

Agenda

1. Note Well, Note Takers, Agenda Bashing, Draft status (Chairs, 10 min)
2. Liaisons (Chairs, 10 min)
   https://www.ietf.org/lib/dt/documents/LIAISON/liaison-2021-10-29-iso-iec-jtc1-sc29-wg11-avtcore-liaison-statement-to-ietf-avt-on-mpeg-green-metadata-attachment-1.pdf
3. RTP over QUIC (J. Ott, M. Engelbart, 15 min)
   https://datatracker.ietf.org/doc/html/draft-engelbart-rtp-over-quic
   https://www.in.tum.de/fileadmin/w00bws/cm/papers/epiq21-rtp-over-quic.pdf
4. Cryptex (Sergio Garcia Murillo, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-cryptex
5. RTP Payloads for VVC/EVC (Stephan Wenger, 15 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-vvc
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-evc
6. SCIP (Daniel Hanson, 15 minutes)
   https://datatracker.ietf.org/doc/html/draft-hanson-avtcore-rtp-scip
7. SFrame/SPacket (Sergio Garcia Murillo & Youenn Fablet, 30 min)
   https://datatracker.ietf.org/doc/html/draft-murillo-avtcore-multi-codec-payload-format
8. Wrapup and Next Steps (Chairs, 10 min)

Draft status

- Published
  - RFC 8817: was draft-ietf-payload-tsvcis
  - RFC 8852: was draft-ietf-avtext-rid
  - RFC 8860: was draft-ietf-avtcore-multi-media-rtp-session
  - RFC 8861: was draft-ietf-avtcore-rtp-multi-stream-optimisation
  - RFC 8872: was draft-ietf-avtcore-multiplex-guidelines
  - RFC 8888: was draft-ietf-avtcore-cc-feedback-message
  - RFC 9071: was draft-ietf-avtcore-multi-party-rtt-mix
  - RFC 9134: was draft-ietf-payload-rtp-jpegxs

Draft Status (2)

- RFC Editor Queue
  - draft-ietf-payload-vp9 (MISSREF)
- Publication Requested
  - draft-ietf-avtcore-cryptex
- Need updated draft for writeup (document status changed)
  - draft-ietf-avtext-framemarking
- Expired and dropped from milestones
  - draft-ietf-payload-tetra (expired January 27, 2020)
- Adopted
  - draft-ietf-avtcore-rtp-evc
  - draft-ietf-avtcore-rtp-vvc (First WGLC completed)
  - draft-ietf-avtcore-rfc7983bis

Recent Liaison Statements

- Request from W3C WebTransport WG
  - Presented at the IETF WEBTRANS WG meeting.
  - Not a formal liaison request (yet).
- Liaison from ISO/IEC JTC 1/SC 29/WG 03
  - Subject: MPEG Green Metadata
  - Received: October 22, 2021
  - Link is here.

W3C WebTransport Request

- The W3C WebTransport WG has identified problems with bidirectional realtime audio/video communication over WebTransport, particularly: a client sending media to the server:
    - **Problem:** The client doesn't have enough information to know when it can reapply a multiplicative increase in the media send rate to recover from prior congestion.
    - **Requests**:
        i. To know if RTP over QUIC can satisfy this use case.
        ii. If so, what measurements can a browser make available to a JS client, to assist with this problem.
        iii. Will selectable congestion control be required? And if so, which algorithm(s)?

Title	**Liaison statement to IETF AVT on MPEG Green Metadata**
Source	**WG 03, MPEG Systems**
Status	**Approved**
Serial Number	**21052**

SC 29/WG 03 would like to update you on the progress of the new edition (3rd) of ISO/IEC 23001-11 Energy-Efficient Media Consumption (Green Metadata). SC 29/WG 03 has issued the CD of ISO/IEC 23001-11 3rd edition in October 2021.

ISO/IEC 23001-11 specifies metadata (Green Metadata) that facilitates reduction of energy usage during media consumption. The format of metadata is specified for the following usages:

- reduced decoder power consumption;
- reduced display power consumption;
- media selection for joint decoder and display power reduction;
- quality recovery after low-power encoding.

This metadata facilitates reduced energy usage on the transmitter and the receiver side during media consumption without any degradation in the Quality of Experience (QoE). However, it is also possible to use this metadata to get larger energy savings, but at the expense of some QoE degradation.

The 3rd edition of 23001-11 adds three main aspects to the standard:

- The green metadata for Interactive signalling for remote decoder-power reduction are enhanced by adding new syntax elements. These allow a finer control for reducing decoder complexity on a high variety of platforms implementing different decoders (hardware / software) as well as processor cores. These metadata are mostly targeting point-to-point video applications and can save roughly 30% of the decoder energy and power, independent from the platform.
- The new edition adds the specification of a VVC SEI message carrying green metadata related to Complexity metrics for decoder-power reduction.
- The VVC SEI message can also carry metrics for quality recovery after low-power encoding. These metadata are mostly targeting segmented video delivery mechanism such as DASH.

SC 29/WG 03 hopes that this information is of interest to you, and you may consider using the Green Metadata to offer services for enabling use cases of video delivery with reduced energy usage, in particular concerning interactive signalling for remote decoder-power reduction. We would be interested to get feedback from your organisation on the potential usage of these metadata in your applications.

Screenshot

Liaison from ISO/IEC JTC 1/SC 29/WG03

Discussion/Next Steps

- Comments
- Next steps

RTP over QUIC

https://datatracker.ietf.org/doc/html/draft-engelbart-rtp-over-quic-01

https://www.in.tum.de/fileadmin/w00bws/cm/papers/epiq21-rtp-over-quic.pdf

J. Ott, M. Engelbart

RTP over QUIC (draft-engelbart-rtp-over-quic-01)

- Very similar to draft-hurst-quic-rtp-tunnelling-01 by Sam Hurst (BBC)
- Encapsulation for carrying RTP/RTCP over QUIC
- Uses unreliable datagram extension
- Flow IDs to demultiplex RTP Sessions
- SDP for signaling
- QUIC Interface Requirements:
  - Acks/Loss signaling for DATAGRAM frames
  - Expose RTT statistics to application
- **Congestion Controller Requirements:**
  - **Assumes one of RMCAT algorithms for real-time media**
  - **Input: ACKed packets, delay, RTT estimations, optionally ECN**
  - **Output: bandwidth estimation for media encoder**

Questions regarding Congestion Control

- QUIC and RTP both provide Congestion Control
- QUIC collects connection statistics and initially suggests an algorithm similar to NewReno; further ones under investigation
- RMCAT proposed SCReAM and NADA (and GCC) to use for congestion control for real-time applications
- RTP uses RTCP to signal congestion (e.g. RFC8888)

⇒ How to do proper real-time media CC in QUIC given RTP realizes its own CC?

⇒ How to avoid duplicate signaling in RTCP?

Options for Real-time Congestion Control on top of QUIC

- Use only QUIC congestion control (NewReno) combined with trivial encoder rate control

- Disable QUIC’s internal congestion control and only use real-time congestion control (SCReAM)

- Use real-time and QUIC congestion control simultaneously (SCReAM+NewReno)

Inferring RTCP Feedback

- Instead of using RTCP, feedback can be created at the sender using QUIC connection stats

- QUIC Datagram acknowledgements signal reception of RTP packets in the ACK’ed DATAGRAM frames

- Use QUIC `latest_rtt` to infer the receive time of an RTP packet:
  - `receive-ts = send-ts + latest_rtt / 2`

Test Implementation

- Integrates quic-go with Gstreamer and SCReAM
- quic-go fork to selectively disable NewReno and expose connection statistics

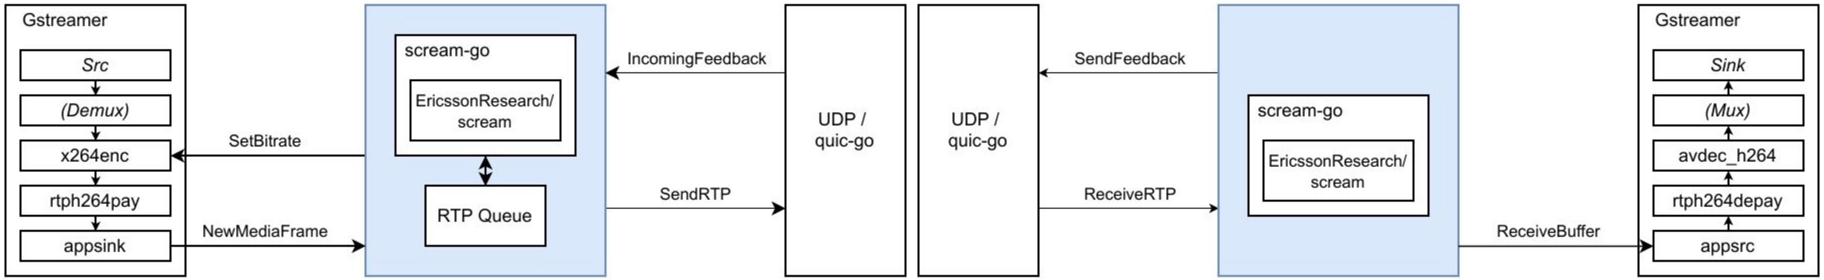

Testbed Setup

- Sender and receiver running in containers on the same host
- Traffic shaping applied to outgoing traffic on virtual interfaces of the containers
- Peers log incoming/outgoing RTP/RTCP and SCReAM target bitrate
- Compute PSNR and SSIM on source and destination raw video files

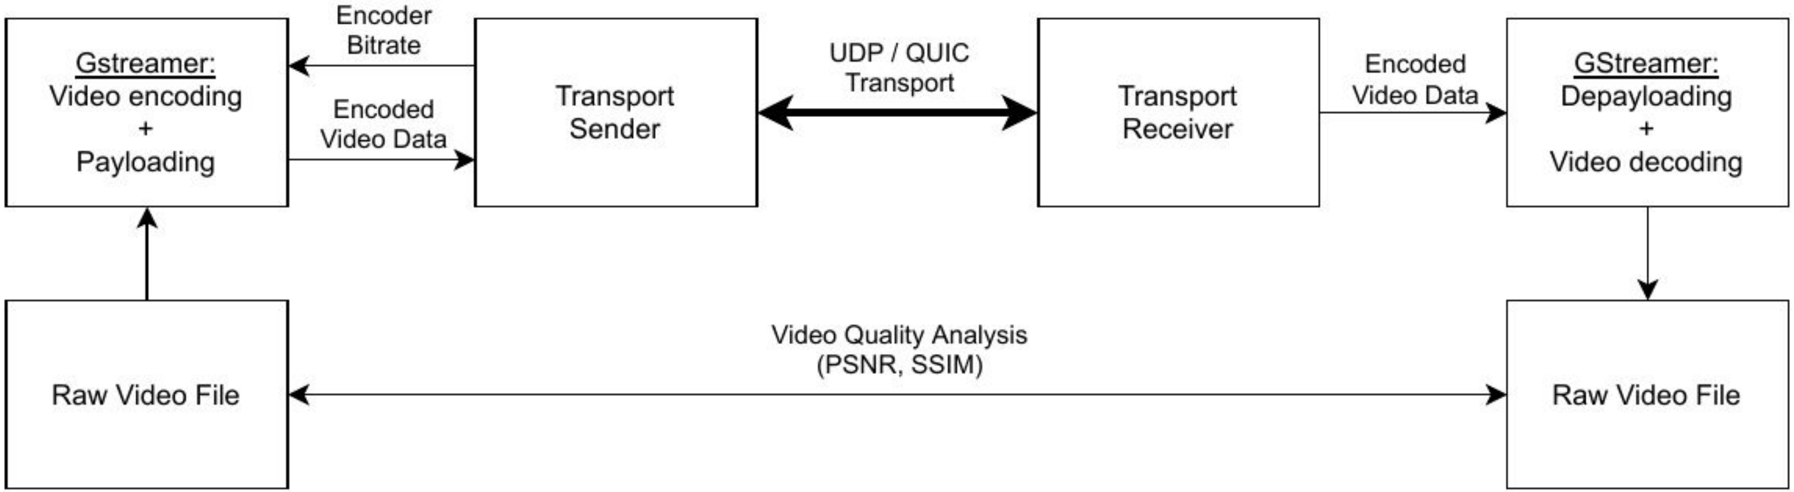

NewReno + Trivial Rate Adaption

- High bitrate fills up RTP queue ⇒ Reduce target bitrate
- Low bitrate drains RTP queue ⇒ Increase target bitrate

⇒ Trivial rate adaption oscillates between high and low

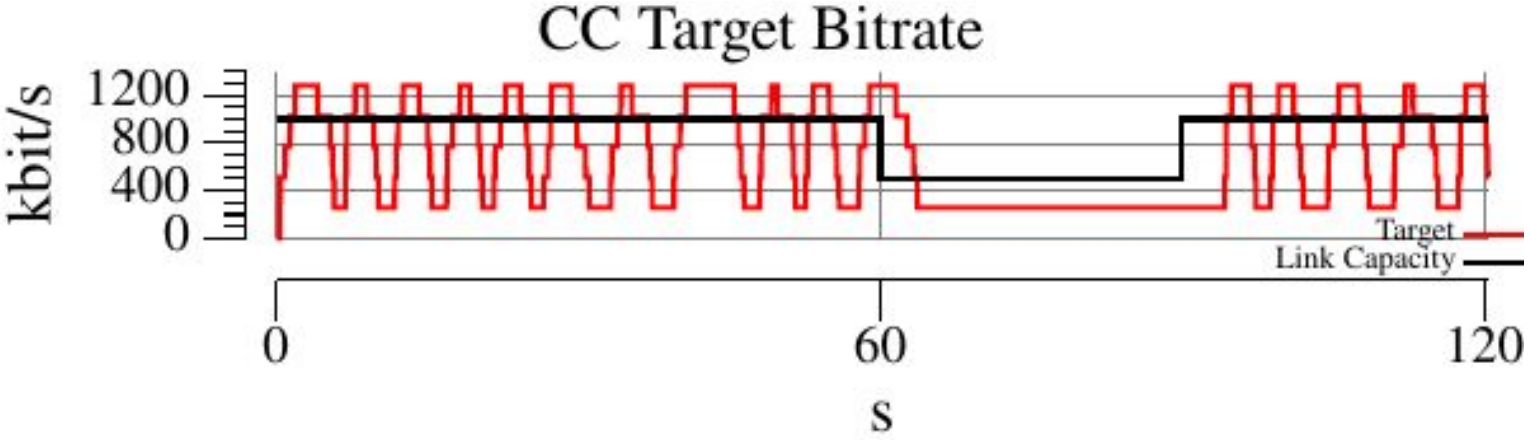

SCReAM with RTP over QUIC and RTP over UDP

- Similar results for QUIC and UDP
- Target bitrate in QUIC stays slightly lower than in UDP

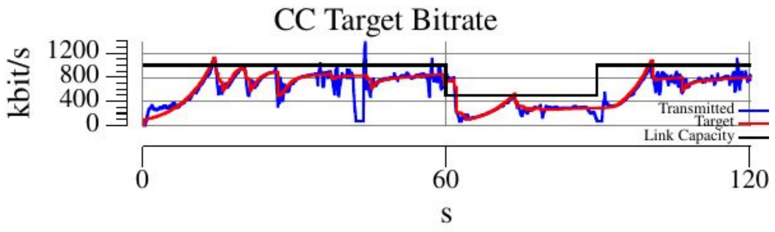

SCReAM + QUIC

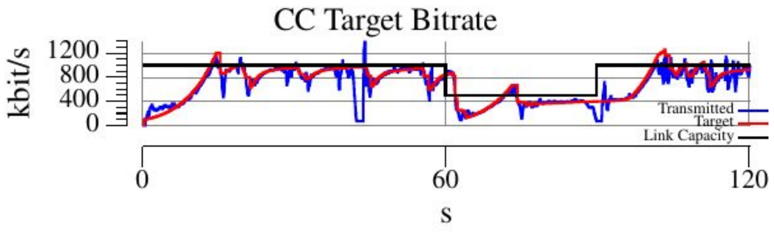

SCReAM + UDP

Combining SCReAM and NewReno

- SCReAM target bitrate similar to RTP over UDP, when application limited
- Problematic, when both congestion controllers try to adapt to bandwidth reduction

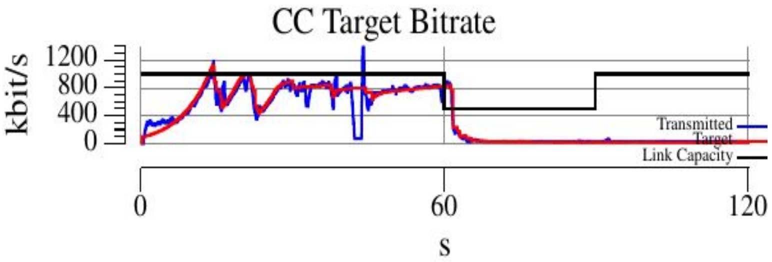

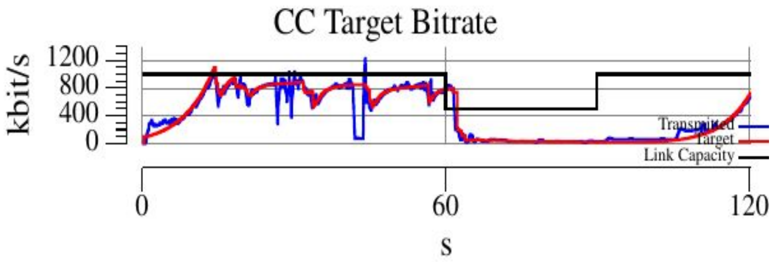

One-way delay: 1ms 50ms

Only SCReAM and Reducing RTCP

OWD	RTCP	QUIC connection stats

1ms

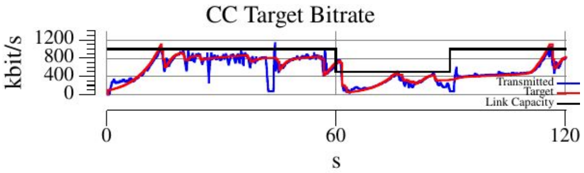

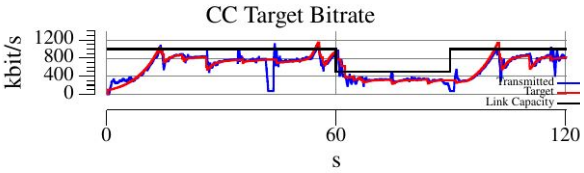

50ms

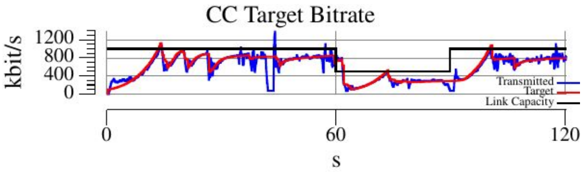

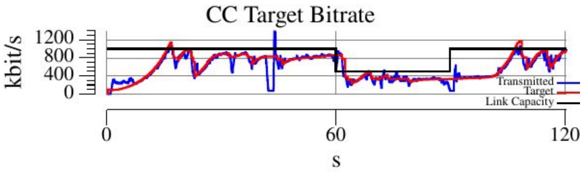

150ms

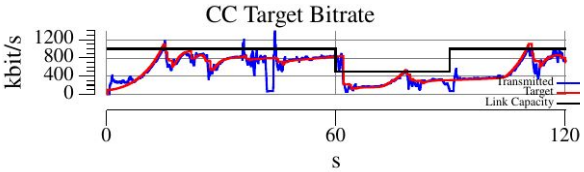

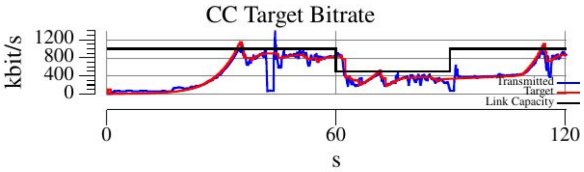

300ms

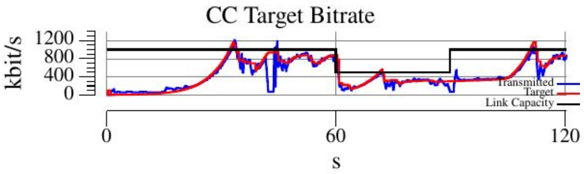

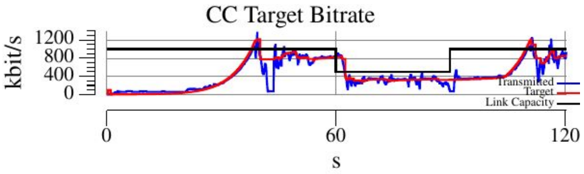

Prioritization

- We ran all experiments again while streaming data on a QUIC stream parallel to RTP in QUIC Datagrams

- Results show, that some form of prioritization is necessary

- Without prioritization, real-time streams may degrade or even starve as a function of the internal operation of the QUIC implementation

- Test cases were rather artificial, more investigation with a more natural form of background traffic needs to be done

Conclusion and Next Steps

Main results:

1. Two separate CC loops at transport and media level are problematic

2. We can reuse QUIC state to reduce RTCP feedback

3. Prioritization is necessary

Next steps:

- Implement more tests with different CC algorithms and other forms of competing traffic
- Come up with some prioritization scheme that gives a reasonable share of bandwidth to each real-time/non-real-time stream

Completely Encrypting RTP Header Extensions and Contributing Sources (Cryptex)

https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-cryptex

Sergio Garcia Murillo

cryptex-03

- Clarify AAD for AEAD modes
- Clarify meaning of cryptex attribute presence in SDP and BUNDLING
- Two implementations (none in production):
  - https://github.com/jitsi/jitsi-srtp/pull/29
  - https://github.com/cisco/libsrtp/pull/511
- Publication request submitted:
  - Archive (ietf.org)
- Proposed WebRTC API changes:
  - Support for cryptex · Issue #88 · w3c/webrtc-extensions (github.com)

Next Steps

- AD Review
- Directorate reviews
- IETF Last Call

RTP Payload Format for VVC

Shuai Zhao, Stephan Wenger (Tencent)

Yago Sanchez (Fraunhofer HHI)

Ye-Kui Wang (Bytedance Inc)

https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-vvc

Summary

- WGLC resulted in many comments from Bernard, Miska, Hendry -- Thanks for those
- All comments were discussed on the avt mailing list, more than 30 emails.
- We believe we have addressed all comments except one (where Miska promised text)
- Very little review by people coming primarily from the IETF side
  - seems normal for payload formats
  - people from security/transport/congestion control will come out of the woods at IETF LC
- Suggested way forward: One more revision including Miska’s text (ETA 11/15), one-week WGLC, publication request

Encouraging examples? When?

Re: [AVTCORE] WGLC on “RTP Payload Format for Versatile Video Coding (VVC)”

Bernard Aboba <bernard.aboba@gmail.com> Mon, 16 August 2021 23:07 UTC

[...]

Section 7

There are no complete examples of SDP Offer/Answer negotiation.

[...]

Re: [AVTCORE] WGLC on “RTP Payload Format for Versatile Video Coding (VVC)”

Stephan Wenger <stewe@stewe.org> Monday, August 30, 2021 at 09:08

[...]

Yes, and in that we follow what RFC7798 has done. At the time, it was remarked that too many examples can be counter-productive towards a full implementation of the spec. All too often, implementers take shortcuts by copy-pasting code sniplets in specs into their production implementation and expect that implementation to work.

I’m not fundamentally against adding an example here, but I think 7798 did the right thing. What does the WG think?

RTP Payload Format for EVC

Shuai Zhao, Stephan Wenger (Tencent)

Youngkwon Lim (Samsung)

https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-evc

Summary

- Draft -01 currently expired
- Will apply “lessons learned” from VVC draft after IETF LC of the VVC draft
- We expect that to be relatively simple, as the payload format mechanisms of the EVC draft are a subset of the VVC draft

- Suggested way forward

    - produce -02 after IETF LC of VVC draft, WGLC on that version

  or

    - produce a keep-alive (perhaps with minor improvements) in late Nov?

RTP Payload Format for the SCIP Codec

Daniel Hanson

and

Michael Faller

https://datatracker.ietf.org/doc/html/draft-hanson-avtcore-rtp-scip

SCIP Background and Purpose

- The Secure Communication Interoperability Protocol (SCIP) began in 1994 in the U.S. as a combined Department of Defense and vendor working group
  - The working group was named the Interoperability Control Working Group (ICWG)
- The goal of this group was to develop the next generation interoperable security protocol supporting U.S. Government and military interests
- The ICWG was later expanded in 2001 to include NATO and NATO partners
  - Name changed to the SCIP Working Group (SCIP WG)
- The SCIP Working Group meets one or two times per year

SCIP Information Access and Awareness Issue

- SCIP standards are currently available to participating government/military communities and select OEMs of network equipment and call management servers that support SCIP
  - Government and business entities must request access to relevant information
  - Access to SCIP standards is based on need-to-know
- Devices that implement the SCIP standards transparently operate over digital carriers
  - Most commercial network and security community personnel are not aware of SCIP
  - Can result in the SCIP media subtype "scip" being removed from the SDP.
  - The lack of awareness among the network and security community has become a larger issue as the use of SCIP grows over more commercial networks, and as network security devices become more restrictive of unknown media

Overview of SCIP RFC (1 of 2)

- The draft RFC submitted to IETF is designed to provide information to network equipment OEMs, network administrators, and security personnel to help SCIP succeed over commercial networks
- The SCIP RFC enables network equipment manufacturers to provide an equipment configuration that supports SCIP as a media subtype
  - Enables network administrators and network security personnel to define and implement a compatible network policy which permits the 'scip' media subtype to traverse the network.
  - End-to-end bit integrity
  - No transcoding on the channel
  - Data streams treated as "clear-channel data"

Overview of SCIP RFC (2 of 2)

- Two media subtypes have been registered with IANA as RTP Payload Format Media Types
  - “audio/scip”
  - “video/scip”
- The RFC is needed to provide additional information for these media subtypes
  - Media Format description
  - Payload Format (RTP Header Fields, Payload Format Parameters)
  - SDP Declaration (Mapping to SDP, Mapping Examples)

SDP Declaration to Support a SCIP Session

- SCIP devices are presently deployed on U.S. and NATO tactical networks, many national networks, and some commercial networks using the following SDP media and submedia types
- Secure Session can use “audio/scip”, “video/scip” or both
- An example mapping for both audio/scip and video/scip is:

```
m=audio 50000 RTP/AVP 96
a=rtpmap:96 scip/8000
m=video 50002 RTP/AVP 97
a=rtpmap:97 scip/90000
```

Summary, Conclusions, and Questions

- Issues have occurred because OEMs of network equipment, network administrators and security personnel are unaware of SCIP and SDP contents necessary to establish a secure session
- The draft RFC increases IETF awareness of the SCIP WG and its efforts to achieve international interoperability
- The purpose of the RFC is to provide global access to information necessary to support SCIP
- A reference to an RFC provides context and a single reference point for the newly defined IANA media subtypes audio/scip and video/scip
    - Provides information about SCIP and the SCIP WG Community to system/network architects, network administrators, security personnel, OEMs, risk analysts, and procurement personnel necessary for SCIP to be included in the system security lifecycle

SFrame RTP payload format

Sergio Garcia Murillo, Youenn Fablet

https://www.ietf.org/id/draft-murillo-avtcore-multi-codec-payload-format-01.html

SFrame Update

- IETF SFrame WG issued a (successful) call for adoption of draft-omara-sframe
  - Interest in using SFrame format at packet or at frame-level
    - Not at arbitrary coding unit level
  - Interest in using SFrame format outside RTP
    - WebTransport/RTCDataChannel + Web Codecs
- W3C WebRTC WG adopted WebRTC encoded transform as FPWD
  - Functionality shipped in Chrome & Safari Tech Preview
  - Used for E2E encryption, app-specific metadata enrichment, RED…
- WebRTC solutions are using all different flavours of SFrame-like formats
  - Google Duo, FaceTime, Webex
  - No interoperability, plus workarounds

We need to make progress!

- Packet vs. frame issue needs additional discussion
- Let's put that aside and focus on other common requirements

SFrame WG mailing list discussions

- Middleware (SFUs, network intermediaries, browsers) cannot inspect encrypted/transformed packet payloads.
- Middleware needs information to route encrypted/transformed packets appropriately, especially in SVC/simulcast cases.

- Questions
  - Should it be possible from RTP packet inspection to determine that content is encrypted/transformed?
  - If so, where should the information be located: payload type, dedicated RTP header extension, within RTP payload?

Solution Requirements

- Ability to identify
  - Whether a RTP packet content is encrypted/transformed or not
    - Prevent packet content inspection from middleware
  - What the RTP packet contains, without inspecting payload content
    - Provide middleware needed information
- Good to have
  - Solution independent of the exact transform/SFrame format
- Support of WebRTC A/V codecs only
  - Not a fully generic solution

Frame Metadata

Potential metadata of interest

- At which packet SFU can switch (SVC and simulcast).
- Resolution and more generally stream 'quality': frame rate, bit rate...
- Codec specific information like profile/levels
- Recovery mechanism required in case of loss (none, RTX, LRR/PLI).
- Opus TOC to know frame length (recording scenarios).

Potential solutions

- Use existing forwarding extensions (e.g. framemarking, AV1 Dependency Descriptor) for supported video codecs.
- Design a new RTP header extension complementing forwarding extensions.

SDP negotiation

```
a=rtpmap:96 vp9/90000
a=rtpmap:97 vp8/90000
a=rtpmap:98 opaque/90000
a=rtpmap:99 rtx/90000
a=fmtp:99 apt=96
a=rtpmap:100 rtx/90000
a=fmtp:100 apt=97
a=rtpmap:101 rtx/90000
a=fmtp:101 apt=98
```

- Encrypted negotiation relies on negotiation of the standard packetizer for the each codec.
- Negotiate an opaque payload type for all codecs in order to avoid duplication or triplication the number of payload types in use.
- Requires sending actual codec payload type, as a RTP header extension or as a prefix to the payload, which will cause minor network overhead.
- Requires negotiating different payload types for each clock rate for audio.

Payload Type Multiplexing

- In order to reduce the number of payload type in the SDP exchange, a single payload type is used
- That requires to identify the original payload type of the negotiated media format, called the associated payload type (APT).
- The APT value is sent in a dedicated header extension:

```
 0                   1
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  ID   | len=0 |S|     APT     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      ID       |     len=1     |S|     APT     |    0 (pad)    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- The APT value is the associated payload type value.
  - The S bit indicates if the media stream can be forwarded safely starting from this RTP packet.
  - Receivers MUST be ready to receive RTP packets with different associated payload types
- The URI for declaring this header extension in an extmap attribute is `urn:ietf:params:rtp-hdrext:associated-payload-type`

Wrapup and Next Steps

- Action items
  - Framemarking: publication request (Chairs)
  - Potential Interim meeting on SFrame (Chairs)
  - Discussion of SCIP publication path (Mailing list)
  - Review of cryptex (AD)
- Next steps
  - VVC: WGLC on -13

Thank you

Special thanks to:

The Secretariat, WG Participants & ADs