AVTCORE WG
IETF 118

Hybrid Meeting
Wednesday, November 8, 2023
Session I, Berlin 1/2
09:30 - 11:30 Prague time
00:30 - 02:30 Pacific time

Mailing list: avtcore@ietf.org
Notes: https://notes.ietf.org/notes-ietf-118-avtcore
MeetEcho link: avtcore (ietf.org)

This session is being recorded

IETF 118 Meeting Tips

In-person participants

- Make sure to sign into the session using the Meetecho (usually the “Meetecho lite” client) from the Datatracker agenda
- Use Meetecho to join the mic queue
- *Keep audio and video off if not using the onsite version*

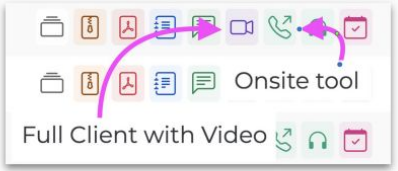

Remote participants

- Make sure your audio and video are off unless you are chairing or presenting during a session
- Use of a headset is strongly recommended

This session is being recorded

IETF 118 Remote Meeting Tips

- 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.

Resources for IETF 118

- Information about IETF 118
  https://www.ietf.org/how/meetings/118
- Agenda
  https://datatracker.ietf.org/meeting/agenda
- If you need technical assistance, see the Reporting Issues page:
  http://www.ietf.org/how/meetings/issues/

About this meeting

- Agenda: https://datatracker.ietf.org/doc/agenda-118-avtcore/
- Notes: https://notes.ietf.org/notes-ietf-118-avtcore
- Secretariat: mtd@jabber.ietf.org
- WG Chairs (Remote): Bernard Aboba
- Onsite: Jonathan Lennox
- Zulip Scribe: Jonathan Lennox
- Note takers: Mo Zanaty, Spencer Dawkins

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.

Agenda

1. Preliminaries (Chairs, 15 min)
   Note Well, Note Takers, Agenda Bashing, Draft status, Errata, IANA registries
2. RTP Payload Format for sub-codestream J2K streaming (P. Lemieux, 10 min)
   https://datatracker.ietf.org/doc/html/draft-lemieux-avtcore-rtp-j2k-scl
3. RTP Payload Format for the SCIP Codec (D. Hanson & M. Faller, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-scip
4. RTP Payload Format for Visual Volumetric Video-based Coding (V3C) (L. Ilola, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-v3c
5. RTP over QUIC (M. Engelbart, J. Ott, S. Dawkins, 15 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-over-quic
6. HEVC Profile for WebRTC (B. Aboba, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-hevc-webrtc
7. RTP Payload Format for SFrame (P. Thatcher, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-sframe
8. Viewport and Region-of-Interest-Dependent Delivery of Visual Volumetric Media (S. Gudumasu, 10 min)
   https://datatracker.ietf.org/doc/html/draft-gudumasu-avtcore-rtp-volumetric-media-roi
9. RTP Payload for Haptics (H. Yang, 10 min)
   https://datatracker.ietf.org/doc/html/draft-hsyang-avtcore-rtp-haptics
10. RTP Payload Format for Geometry-based Point Cloud Compression (M. Engelbart, 10 min)
    https://datatracker.ietf.org/doc/html/draft-engelbart-avtcore-rtp-gpcc
11. Wrapup and Next Steps (Chairs, 10 min)

Draft Status

- Published
  - RFC 9071: was draft-ietf-avtcore-multi-party-rtt-mix
  - RFC 9134: was draft-ietf-payload-rtp-jpegxs
  - RFC 9328: was draft-ietf-avtcore-rtp-vvc
  - RFC 9335: was draft-ietf-avtcore-cryptex
  - RFC 9443: was draft-ietf-avtcore-rfc7983bis
- RFC Editor Queue
  - draft-ietf-payload-vp9 (877 days, MISSREF)
- IESG: AD Followup (303 days, 3 DISCUSS positions)
  - draft-ietf-avtcore-rtp-scip
  - Ballot statements: https://datatracker.ietf.org/doc/draft-ietf-avtcore-rtp-scip/ballot/
  - More discussion today.
- IESG Evaluation::Revised I-D Needed (100 days)
  - draft-ietf-avtext-framemarking
- In IETF Last Call (Ends November 21, 2023)
  - draft-ietf-avtcore-rtp-evc

Draft Status (cont'd)

- WGLC
  - draft-ietf-avtcore-rtp-v3c (ended November 7, 2023)
- Adopted
  - draft-ietf-avtcore-rtp-over-quic
  - draft-ietf-avtcore-rtp-green-metadata
  - draft-ietf-avtcore-rtp-v3c
  - draft-ietf-avtcore-hevc-webrtc
  - draft-ietf-avtcore-rtp-j2k-scl
    - Adoption summary
  - draft-ietf-avtcore-rtp-sframe
    - Adoption summary

Errata (from Francesca)

- Eid 4873 for RFC5764:

  https://mailarchive.ietf.org/arch/msg/avt/Sc1fRHdSh9PDAm4PeMxl2lfwrz0/
  Francesca: This seems like it should be verified.

- Eid 4938 for RFC7714:

  https://mailarchive.ietf.org/arch/msg/avt/MDLct0Lo1j18vbU4gpF00M7z3Kw/
  Francesca: I want to make sure the WG signs off on this change - was this always the intent? If that is not completely clear, then maybe this should rather be "hold for document update".

- Eid 6752 for RFC9134:

  https://mailarchive.ietf.org/arch/msg/avt/0yXojWIOp99HfPMZpoxSU0hNkCI/
  Francesca: This also seems like it should be verified as well.

IANA Registries Background

- A post to the W3C public-webrtc mailing list pointed out an issue with IANA RTP payload format type registrations: https://lists.w3.org/Archives/Public/public-webrtc/2023Aug/0033.html
  - RTP payload types registry is missing VP8, AV1, HEVC, VVC: https://www.iana.org/assignments/rtp-parameters/rtp-parameters.xhtml#rtp-parameters-2
- IANA mime-types registry (see “video”) is more complete: https://www.iana.org/assignments/media-types/media-types.xhtml

  Spreadsheet: https://www.iana.org/assignments/media-types/video.csv
- Issue tracked by MEDIAMAN WG
  - https://github.com/ietf-wg-mediaman/admin/issues/1

RFC 8088, Section 7.4 (IANA Considerations)

Since all RTP payload formats contain a media type specification, they also need an IANA Considerations section. The media type name must be registered, and this is done by requesting that IANA register that media name. When that registration request is written, **it shall also be requested that the media type is included under the "RTP Payload Format media types" subregistry of the RTP registry (http://www.iana.org/assignments/rtp-parameters).**

Parameters for the payload format need to be included in this registration and can be specified as required or optional ones. The format of these parameters should be such that they can be included in the SDP attribute "a=fmtp" string (see Section 6 [RFC4566]), which is the common mapping. Some parameters, such as "Channel" are normally mapped to the rtpmap attribute instead; see Section 3 of [RFC4855].

Response from Stephan

Stephan Wenger <stewe@stewe.org> Wed, 04 October 2023 13:43 UTCShow header

Copying specifically Harald Alvestrand (as Mediaman chair) and Magnus Westerlund (as RFC 8088 author)

I'm an arguably experienced payload format author, but know little about IANA mechanics, and care even less. I'm sure I'm not alone with respect to the latter. People like me will need a template and guidance from which we can copy-paste-adapt in the future. I don't know who can provide that guidance, but frankly, I don't believe AVTcore is the right group to provide it, simply because we didn't get it right so many times in the past. I also don't think we can rely on guidance by the IESG or the media type review people without specifically asking. They didn't spot problems in the past, either.

I have no idea on how to arrive at that suggestion, though W3C identifying and Bernard and Jonathan taking ownership of the problem by at least describing it clearly is an excellent first step. Maybe move the discussion over to the media-types list? Or Mediaman could help? Once we have the sentence or two needed, the EVC payload could be used as a test case as it is in the right stage of the process to see results quickly.

Then, I would suggest revising RFC 8088 and specifically section 7.4. (In a revised RFC 8088, we should also insert a section dealing with normative references to media codec specs, as this has been another recent pain point.)

Stephan

Response from Magnus

From: Magnus Westerlund <magnus.westerlund@ericsson.com> Mon, 09 October 2023 10:04 UTC

The RTP Payload Format Media Types is a mostly redundant registry, in that its only purpose is to help determine which media types that has RTP Payload formats. Already prior to RFC 8088 was written we had discussion about this registry and if we should kill it. We didn't do that at that time, and the best we could do in RFC 8088 was to be very clear that you need to ensure that it is registered there also. However, as noted it has been forgotten several times, as the media type registry is what is actually relevant to register.

I would think a reasonable way forward is to actually close the registry and update RFC4855 and RFC 8088. When closing this, one could consider to clean up the registry up to the date of closing by adding the missing RTP payload format media types that have been missed. But, unless someone knows of someone externally depending on this registry, to avoid future issues, and also be able to make a clear comment that this registry is not relevant in a note, writing a document killing it is likely an easy way forward.

Cheers
Magnus

Response from IANA

---------- Forwarded message ---------
From: Amanda Baber via RT <iana-issues@iana.org>
Date: Tue, Oct 3, 2023 at 18:54
Subject: [IANA #1283165] FW: [Ext] RTP payload format type registry vs. MIME-type registry
To: <amanda.baber@iana.org>
CC: <bernard.aboba@gmail.com>, <hta@google.com>, <jib@mozilla.com>, <jonathan.lennox42@gmail.com>

Hi Bernard, all,

Sabrina Tanamal (sabrina.tanamal@iana.org, bcc'd on all IETF-related ticketing system messages) and I are the people to contact about this, but you can reach anyone in IANA via iana@iana.org, which is heavily monitored throughout the day.

Before we forward this to an ART AD, could you give us the list of registrations/references (along with any information that should be added to the entries' "Clock Rate (Hz)" and Channels (audio)" fields) that need to be added to https://www.iana.org/assignments/rtp-parameters/rtp-parameters.xhtml#rtp-parameters-2?

Response from IANA (cont'd)

In the registry itself, the registration procedures are listed as "Standards Action or Expert Review." (Given that RFC 4855 doesn't mention an expert, it's not clear when the latter policy is meant to be applied.) The ADs haven't had a chance to replace Steve Casner as expert, but an AD can tell us at least how this registry can be updated (e.g. can we make the RFC-related registrations now?), if not actually fill that expert role.

The issue is that we haven't been asked to add these, and we don't appear to have been asked to look for triggers that would tell us to initiate a registration (which we aren't really equipped to do; here in operations, we maintain 3500 registries and are liberal arts majors).

And the authors would have been aware that we weren't making those RTP registrations. When we reviewed RFC 7741 during IETF Last Call, for example, we told the authors/chairs/IESG that we understood from the IANA Considerations section that the only action was to add the media type to the registry at https://www.iana.org/assignments/media-types, and asked them to respond if that was incorrect. After we made that registration, as with the other RFCs mentioned below, the authors reviewed it and confirmed that the registry actions were complete and correct.

You also mentioned https://datatracker.ietf.org/doc/html/draft-ietf-payload-vp9 below. That IANA Considerations section is clear, but we won't make the registrations
until the IESG approves the document.

Response from IANA (cont'd)

For any media type registrations made by SDOs, we also would have needed to receive an explicit request to add an RTP registration or an instruction from a media type expert who reviewed the request (followed by the RTP expert approval). However, I don't know whether the media type expert would have been aware that they were the only person who might flag it, or that the SDO wouldn't have already been submitting a separate request.

I can ask the AD if IANA should be looking for a trigger here, given that media types registered by RFCs aren't subject to an official expert review. Except in rare cases that require escalation (for example, an issue related to .arpa or LGRs, which would involve IANA directly), those of us in operations aren't involved in looking at technical content (unless given specific non-technical instructions).

Also, Harald: is there anything here that should be brought to the attention of the mediaman WG?

thanks,

Amanda Baber
IANA Operations Manager

High-performance JPEG 2000 RTP payload format

draft-lemieux-avtcore-rtp-j2k-scl

Start time: 09:45
End time: 09:55

Status

- Call for adoption (successfully) closed 2023-11-07.
  - Summary
- New I-D expected shortly to address two issues identified during 2023-09-26 WG call
  - Update packet structure diagrams to match common RFC style
  - Discuss PTSTAMP relationship with RFC 5450 (“Transmission Time Offsets in RTP Streams)
- An FPGA implementation using the proposed payload format was demonstrated by Kakadu Software at SMPTE MTS (packet captures available)

RFC 5450 and PTSTAMP address different issues

RFC 5450: transmission time of a packet is significantly different (> 1 video frame) from the presentation time of its payload

- Frame reordering, retransmission, large variations in video frame size…
- 24-bit offset (32-bit field)

PTSTAMP: improves clock recovery by providing additional clock samples between successive video frames

- 12-bit offset
- Fits in a hardware-friendly 48-byte (16×3) header (IPv4/UDP/RTP/SCL)

The scenarios contemplated by RFC 5450 do not generally apply to J2K SCL (intra-frame coding) but individual applications could specify the use of both

Plan: add informative text to J2K SCL

RTP Payload Format for SCIP

draft-ietf-avtcore-rtp-scip

Dan Hanson

Mike Faller

Start time: 09:55
End time: 10:05

Status

- Published Draft Version 06 on September 19
  - Updated Abstract
  - Added Key Points section up front
  - Included link to old public version of SCIP-210 (from 2013)
  - Added SDP example of prioritized codecs
  - Note: no comments from IESG reviewers for this revision
- Version 07 is in progress
  - Additional updates to Abstract
  - Will probably remove Key Points section since most of the information is now conveyed in the revised Abstract
  - Awaiting feedback and internal review before publishing

Next Steps

- At SCIP Working Group meeting in Warsaw, Poland October 16-18, management expressed concern that this effort has stalled
- Both U.S. Government and NATO need this document to move to RFC for the next generation of SCIP products
- What can we do to help?

Suggested Additions to -07

- Section 5 (RTP Payload)
  - More detail on RTP packetization/depacketization needed, to make it clear why parsing SCIP is unnecessary and potentially harmful.
  - Material can be structured as text explaining the figure.
- Section TBD (Transport)
  - Purpose is to explain interaction of SCIP with multiplexing.
  - SCIP negotiation is orthogonal to aspects of RTP/RTCP, such as multiplexing (RFC 5761) and symmetric RTP (RFC 4961).
    - SCIP implementations may or may not negotiate RTP/RTCP mux.
    - Can SCIP BUNDLE audio/video (RFC 8843)?

Suggested Additions (cont'd)

- Section TBD (Profiles and Feedback)
  - Describe which profiles can be used with SCIP.
  - SCIP provides security services including confidentiality, replay and integrity protection for RTP payloads.
    - SCIP does not provide security services for RTCP
    - Additional RTP security provided by SRTP
      - Authentication services and (with cryptex) encryption of RTP header extensions and CSRCs.
      - Does use of SRTP mitigate any important attacks?
  - SCIP handles retransmission, so no reason to negotiate NACK or RTX.
    - Negotiation of FEC is orthogonal to SCIP.
    - Is it useful to negotiate PLI? FIR? RPSI?
      - If RPSI is useful, what are the codec-specific parameters?

Suggested Additions (cont'd)

- Section 7 (Security Considerations)
  - Summary of SCIP security services (e.g. payload confidentiality, replay and integrity protection for RTP).
  - If SRTP and cryptex is not mandatory, explain the additional attack surface and potential mitigations.
    - No SRTP: Attacks on RTP header
    - No cryptex: Snooping on RTP header extensions and CSRCs
      - Can determine the active speakers and their locations.
  - Summary of potential RTCP vulnerabilities
    - For audio (is RTCP feedback provided?)
    - For video (e.g. H.264)

RTP Payload Format for Volumetric Video Coding (V3C)

draft-ietf-avtcore-rtp-v3c

L. Ilola

L. Kondrad

Start time: 10:05
End time: 10:15

Update since #117

- Working Group Last Call issued
  - running until 7th of November
  - results to be discussed
- Open source implementation released
  - Provides Gstreamer plugins for atlas data payloader and deplayloader
  - Can be used to packetize atlas data NAL units into RTP packets
- As always issues can be submitted here. Three open editorial issues:
  - #12 Aggregation packet figure
  - #14 Removal of V3C unit header syntax structure
  - #15 Unused terms & undefined acronyms

RTP over QUIC

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

https://datatracker.ietf.org/doc/draft-dawkins-avtcore-sdp-rtp-quic/

https://datatracker.ietf.org/doc/draft-dawkins-avtcore-sdp-rtp-quic-issues/

Mathis Engelbart, Jörg Ott, Spencer Dawkins

Start time: 10:15

End time: 10:30

Done since September Interim (1)

- #13, #112 / PR#144: Strategies for datagram extension non-support
- #75 / PR#125: Improve consideration for multi-hop topologies
- #76 / PR#140: Define error codes and related IANA considerations
- #117 / PR#124: Suppressing QUIC signaling in favor of RTCP signaling
- #127 / PR#133: Remove normative references to individual drafts
- #128 / PR#134: Clarify congestion and rate control

Done since September Interim (2)

- #129 / PR#132: CONNECTION_CLOSE does have an error phrase
- #131 / PR#138: Remove normative API requirements
- #136 / PR#139: Allow to continue sending media on QUIC datagrams after receiving STOP_SENDING
- #141 / PR#145: Move rate adaptation mention to scope
- #142 / PR#143: Configure streams (flow control)
- #146 / PR#146: Improve congestion control and rate adaptation in motivation section

Closed as *won't fix*

- #29: SFrame/SPacket
  - See draft-thatcher-avtcore-rtp-sframe, but isn't RoQ-specific
- #52: Possible use of QUIC Multicast?
  - Would require draft-jholland-quic-multicast to progress in QUIC WG
- #112: RESET_STREAM and media dependencies
  - Discussed at last interim - Section 5.2.2 describes this now
- #113: Given that we are doing STOP_SENDING, do we need CLOSE_STREAM?
  - Discussed at last interim - not needed for loss-tolerant RTP traffic

Labeled NextDoc / Not Yet, won't fix?

- #50: QUIC interaction with ICE
  - Waiting for QUIC NAT traversal / P2P QUIC in QUIC WG
- #51: Possible use of QUIC multipath
  - Waiting for MP-QUIC to clear QUIC WGLC
  - Should do #115: Motivation: Exploiting Multiple Connections first

Question or Discussion / Review required

- #84: Add more detail to Section 7.3 on congestion control when sharing connections with non-RTP streams
- #87: Are real-time congestion controllers tied to RTP or usable in QUIC stacks?
- #114: Reality check on RoQ DATAGRAM use, especially for RoQ over WebTransport
- #126: Review Appendix B for completeness

#84: Add more detail to Section 7.3 on congestion control when sharing connections with non-RTP streams

Continued on next slide

- QUIC connections can be shared between RTP/RTCP and non-real-time data streams
- Need to make sure not to starve media or data streams
- Artificially rate limiting the non-real-time data sounds like a good idea, but is that sufficient?
- Toy protocols for testing this kind of multiplexing:
  - draft-engelbart-quic-data-channels
  - draft-engelbart-multiplex-roq-qdc

#84: Add more detail to Section 7.3 on congestion control when sharing connections with non-RTP streams

Continued from previous slide

- Leave it to the application
  - rather than defining complex data structures to express policies
  - E.g., ratio vs. minimum
- Example
  - Assume an application can gain insight into how much data is queued within QUIC and what the current congestion control window is.
  - Then it can determine how much data vs. real-time media to send at any given instant.
  - If the application keeps the QUIC buffers low, decisions result in instant “actions”.
  - Allows dynamically adjusting decisions at any point in time
  - Helps deferring decisions to when it’s clearer what is actually needed
- Another API question?

#87: Are real-time congestion controllers tied to RTP or usable in QUIC stacks?

- Section 7 on Congestion Control changed **a lot** recently
- It still has informative references to SCReAM/NADA/GCC as examples
- As far as we know there are no QUIC implementations of these algorithms
  - Implementation would require *some* QUIC timestamp extension
- Does the document need to say anything else about the applicability of these algorithms in QUIC?
  - There's an ongoing conversation about ***part*** of this topic in GitHub
- L4S is a complementary technology that is also referenced in the draft

#114: Reality check on RoQ DATAGRAM use, especially for RoQ over WebTransport

- The draft has always included both QUIC streams and datagrams
  - This makes the draft longer and more complex
  - We asked if people expect to use datagrams (they do, this is fine)
  - Should we provide guidance on choosing between them?
- RoQ endpoints can use either streams or datagrams, **or both**
  - RoQ receiver uses STOP_SENDING to "catch up" on streams (fine)
  - RoQ sender can resume sending using datagrams (is this **also** fine?)
  - Strawman SDP requires a different SDP proto (STREAM vs DGRAM)
  - Would anyone want/need to do this? Should we allow this?

#126: Review Appendix B for completeness

- Contains most of the RTCP analysis we did to identify packet types that can be replaced by QUIC state
- We will review the appendix again
- ***Additional reviews would be appreciated***

Next Steps (target PRs for next interim)

- #65: RTP and non-RTP traffic sharing multiple QUIC connections
- #86: Coalescing RTP packets in single QUIC packet
- #111: Check SHOULD requirements for ROQ receivers
- #115: Motivation: Exploiting Multiple Connections
- #147: Ensure *Datagram* terminology is used consistently

HEVC Profile for WebRTC

draft-ietf-avtcore-hevc-webrtc

Bernard Aboba

Philipp Hancke

Start time: 10:30
End time: 10:40

Implementation Update

- Chromium
  - WebRTC Tracking bug (13485)
    - Last CL submitted on September 20, 2023.
      - Includes support for SDP negotiation, RTP payload format.
    - Issues found with RPSI CL
- Safari
  - WebRTC Tracking bug (242921)
    - PR 19703 (enable HEVC by default, not yet merged)
  - Discovery
    - WebCodecs decoding bug (262950)
    - Media Capabilities Accuracy (236444)

For Discussion Today

https://github.com/aboba/hevc-webrtc/issues

- Issues & PRs
  - Issue 3/PR 18: **tx-mode**
  - Issue 12/PR 16: **PACI packet**
  - Issue 13/PR 17: **RPSI RTCP feedback support**

Issue 3/PR 18: tx-mode

- Issue 3: Need to cover tx-mode SDP parameter.
- PR 18 Proposed text:

“tx-mode: Implementations SHOULD NOT include this parameter within SDP. If no tx-mode parameter is present, a value of "SRST" MUST be inferred. Implementations MUST support "SRST"; support for "MRST" and "MRMT" is OPTIONAL. Implementations that do not support "MRST" or "MRMT" MUST NOT include these tx-mode values in SDP.”

Issue 12: PACI packet

- PACI extensions are defined in RFC7798 Section 4.4.4.
  - A new payload type 50 is introduced for the PACI packet.
  - Section 4.4.4.2 defines a payload header extension for Temporal Scalability Control Information (TSCI).
- In WebRTC implementations, TSCI information can also be communicated using RTP header extensions such as the Dependency Descriptor or Generic Frame Descriptor.
  - If TSCI is supported via negotiated RTP header extensions, what should implementations do with TSCI information received in PACI extensions?

Issue 12/PR 16: Add paragraph on TSCI

- PR 16 Proposed text:

  "[RFC7798] Section 4.5 defines how Temporal Scalability Control Information (TSCI) is communicated using PACI Extensions defined in [RFC7798] Section 4.4.4.2. A WebRTC implementation that has negotiated use of RTP header extensions containing TSCI information (such as the Dependency Descriptor [DD]) SHOULD NOT send TSCI information within the PACI. If TSCI information is being received in an RTP header extension, implementations MUST ignore TSCI information contained in the PACI."

Issue 13: RPSI RTCP feedback support

- RFC 7798 Section 8.3 defines support for RPSI feedback in HEVC.
- “Use of RTP in WebRTC” RFC 8834 Section 5.1.4 says:
  - “Reference Picture Selection Indication (RPSI) messages are defined in Section 6.3.3 of the RTP/AVPF profile [RFC4585]. Some video-encoding standards allow the use of older reference pictures than the most recent one for predictive coding. If such a codec is in use, and if the encoder has learned that encoder-decoder synchronization has been lost, then a known-as-correct reference picture can be used as a base for future coding. The RPSI message allows this to be signaled. **Receivers that detect that encoder-decoder synchronization has been lost SHOULD generate an RPSI feedback message if the codec being used supports reference-picture selection. An RTP packet-stream sender that receives such an RPSI message SHOULD act on that messages to change the reference picture, if it is possible to do so within the available bandwidth constraints and with the codec being used.**”
- Clarifications in RFC 8082 Section 6.3.
  - “no implementations are known in conjunction of layered codecs. The current understanding is that the reception of an RPSI message on any layer indicating a missing reference picture forces the encoder to appropriately handle that missing reference picture in the layer indicated, and in all dependent layers.”

Issue 13/PR 17: RPSI

- PR 17 Proposed text:

  “[RFC7798] Section 8.3 specifies the use of the Reference Picture Selection Indication (RPSI) in H.265. Implementations MUST use the RPSI feedback message only as a reference picture selection request, and MUST NOT use it as positive acknowledgement. Receivers that detect that H.265 encoder-decoder synchronization has been lost SHOULD generate an RPSI feedback message if support for RPSI has been negotiated, unless the receiver has knowledge that the sender does not support RPSI. Such knowledge can be established during capability exchange or through previously sent RPSI requests that were not replied to by the sender through the use of a non-IRAP picture. An RTP packet-stream sender that receives an RPSI message MUST act on that message, and SHOULD change the reference picture.”

Comments on PR 17: RPSI

- Philipel-Webrtc:

“unless the receiver has knowledge that the sender does not support RPSI”

  - Should this be removed? If it was negotiated then the sender supports it, if it wasn't then the assumption is that the sender does not support it. I feel like the "unless" is implied in the first part of the sentence.

RPSI Implementation Issues

- There are two usages of RPSI defined by RFC 4585:
  a. Indication of requested reference picture (kept in RFC 7798).
  b. Positive feedback of successfully decoded picture (deprecated in RFC 7798 Section 8.3).
- RPSI CL review (https://webrtc-review.googlesource.com/c/src/+/320167) found issues in supporting usage a.
  a. SFU forwards RPSI to encoder, who generates a new P-frame with dependency info (e.g. DD)
  b. Can SFU determine whether the new P-frame is decodable by all participants?
     - In RPSI usage b, SFU knows which participants have decoded LTR
     - For usage a, SFU needs last seqno received (transport-cc) ***and decoded***. See:
       - Issues with LTR support
       - Addition of fields to PLI
       - LNTF RTCP message

RTP Payload Format for SFrame

draft-ietf-avtcore-rtp-sframe

Peter Thatcher

Start time: 10:40
End time: 10:50

Since Last Meeting

- draft-ietf-avtcore-rtp-sframe adopted and submitted

Next Steps

- Please review it!
- Maybe try and implement it?
- Maybe discuss the following questions mentioned previously:
  - Do all SFrame packets contain all of the inner-codec's header extensions?
    - If yes: duplicates data (Dependency Descriptor)
    - If no: have to figure out which ones go when and which don't
  - Do all SFrame packets contain the inner-codec's payload type?
    - If yes: duplicates a byte per SFrame packet after the first of each frame
    - If no: save a byte, but don't know what the inner payload type is until you get the first packet of each frame

Reference

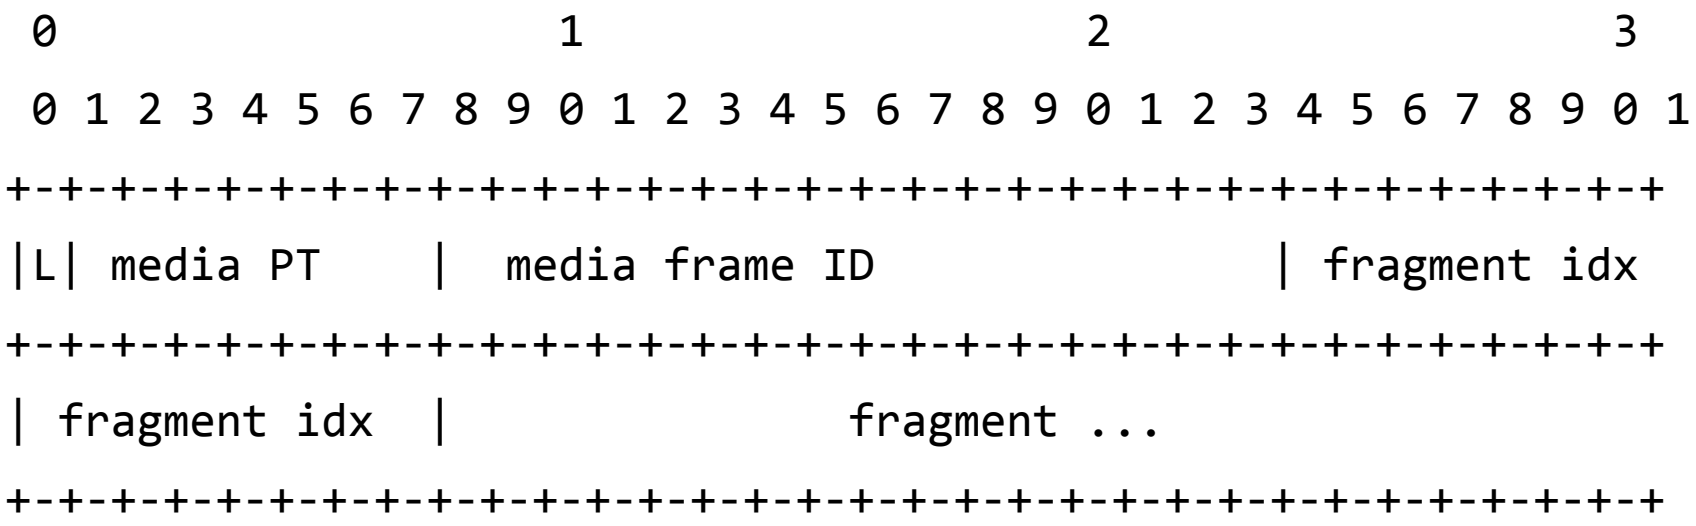
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
|L| media PT | media frame ID | fragment idx
| fragment idx | fragment ...

Viewport and Region-of-Interest-Dependent Delivery of Visual Volumetric Media

draft-guduasu-avtcore-rtp-volumetric-media-roi

S. Gudumasu

Start time: 10:50
End time: 11:00

Status since last meeting

- Published the new draft version on 25th September 2023
- Addressed the comments received in the earlier meeting
  - Syntax and semantics of Spatial region and Viewport are aligned with MPEG specifications (23090-5, 23090-10 and 23090-7).
  - Added definitions for Coordinate systems and others
  - Anchor point and sizes are defined in terms of volumetric pixels
  - 3D spatial region anchor, and size formats are defined
  - Viewport FCI is updated to align with MPEG specifications
  - Security considerations, IANA considerations
  - Normative references and Editorial updates

Next Steps

- Suggestions and feedback
  - IETF mailing list
- Working group adoption

RTP Payload for Haptics

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

H. Yang

X. de Foy

Start time: 11:00
End time: 11:10

What is Haptics?

- Science and technology related to the sense of touch, for example with the modalities:
  - **Vibrotactile perception (most used today)**
  - **force**
  - texture
  - temperature
  - … (e.g., 14 listed in the specs)

- Science on haptics is several decades old
- Technology is catching up fast: more and more devices, used along with audio/video.

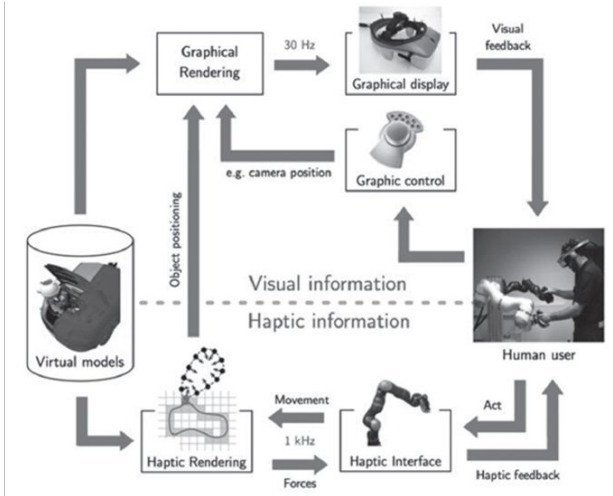

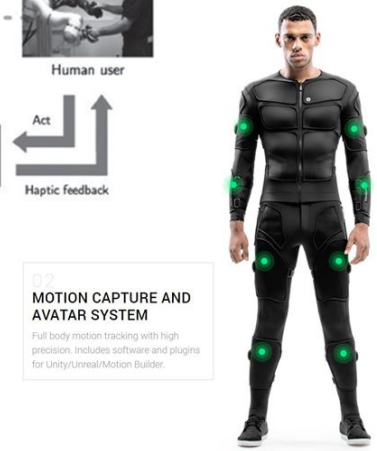

Background on Haptic Standard

We are participating in the development of the MPEG standard for haptics:

- ***Haptics Phase 1 – Haptic Coding Standard (ISO/IEC 23090-31) – towards completion (ISO standard ~2024)***
  - Coding of time-dependent haptic signals
  - Suitable for timed-haptic experiences that may be synchronized with audio and/or video media.
  - Moving to Final Draft International Standard (FDIS) – January, 2024
  - Reference software is available for the haptics encoder/decoder (only MPEG member access now)

- ***Haptics Phase 2 – Haptic Interactions, Avatars, XR, Scene Description – ongoing***
  - Started with use cases, and requirements

MPEG-I Haptic Stream (MIHS)

- MIHS units are the units of streaming defined in MPEG-I haptic code standard.

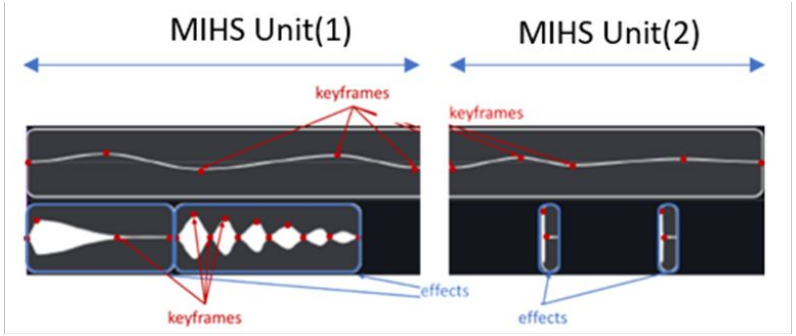

- 4 types of MIHS units
  - Initialization units (periodic timing and metadata)
  - Temporal units (covers a duration of time, carry haptic effects that vary over time)
  - Spatial units (carry haptic effects which do not vary continuously, e.g., texture)
  - Silent units (replace temporal units during “haptic silence”)

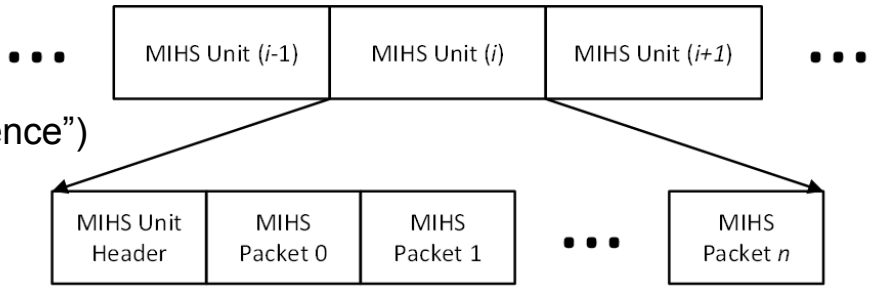

RTP Payload Header

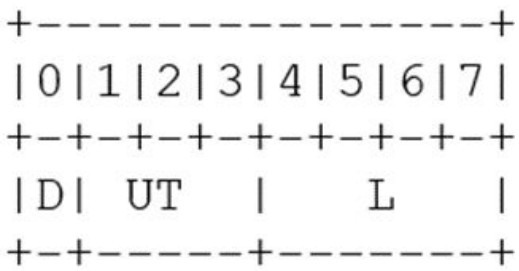

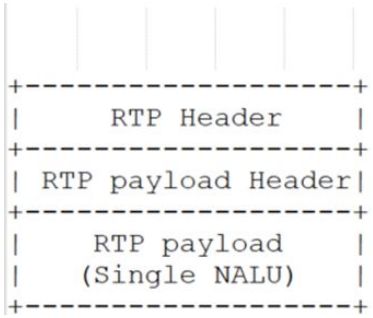

Single MIHS unit

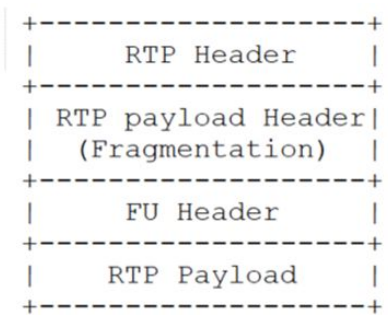

Fragmented MIHS unit

- D (Dependency) indicates whether the packet contained in the RTP Payload is a sync packet or not.
- UT (Unit Type)
- L (MIHS Layer) MIHS unit layer set by the sender based on application-specific needs.

Unit Type	Payload Structure	Name
0	N/A	Reserved
1	Single	Initialization MIHS Unit
2	Single	Temporal MIHS Unit
3	Single	Spatial MIHS Unit
4	Single	Silent MIHS Unit
7	Frag	Fragmented Packet*

** aggregation is not included for now – TBD*

Example of SDP parameter

- The media name in the "m=" line of SDP MUST be haptics.
- The encoding name in the "a=rtpmap" line of SDP MUST be hmpg
- The clock rate in the "a=rtpmap" line may be any sampling rate, typically 8000.
- The OPTIONAL parameters, when present, MUST be included in the "a=fmtp" line of SDP.
- An example of media representation corresponding to the hmpg RTP payload in SDP is as follows:
  - `m=haptics 43291 UDP/TLS/RTP/SAVPF 115`
  - `a=rtpmap:115 hmpg/8000`
  - `a=fmtp:115 hmpg-profile=1;hmpg-lvl=1;hmpg-ver=2023`

Conclusion

- **Conclusion**
  - This draft is a first step to describe an RTP payload format for the MPEG-I haptics coding standard.
  - A reference software implementation is available for the encoder/decoder.
  - We developed an in-house implementation of the RTP payload format (for future interop testing).

- **Next Steps**
  - We are looking for people interested in reviewing, implementing or participating in the draft.
  - The MPEG-I haptics coding standard is nearly completed. We need to follow it and adapt the draft if need be.
  - We can ask MPEG to make some information available to IETF members though a Liaison Statement.

RTP Payload Format for Geometry-Based Point Cloud Compression

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

Mathis Engelbart, Jörg Ott, Lukasz Kondrad

Start time: 11:10
End time: 11:20

Background: Point Clouds

- Data Structure used to represent three dimensional data
- List of points in three dimensional space
- Each point may be associated with zero or more attributes
  e.g., Color, Reflectance, …
- Acquired by, e.g., LiDAR, Radar, Multiple Camera setups
- Example use case: 3D representation of a vehicle's surrounding environment

Background: G-PCC

- ISO/IEC 23090-9: *Information technology — Coded representation of immersive media — Part 9: Geometry-based point cloud compression*
- Geometry data and attributes are coded independently
- 10 Types of Data Units (DUs)
  - Parameter Sets
  - Geometry Data Units
  - Attribute Data Units
  - …
- Annex-B describes Type-Length-Value Bitstream format

RTP Payload Format

- Standard RTP header usage
- Timestamp defined as the earliest sampling time of the point cloud frame
- Payload Header includes two type fields:
  - Single Unit / Aggregation Unit / Part of Fragmentation Unit
  - Type of Data Unit in packet payload
- Aggregation Unit packets additionally prefix each Data Unit with a variable length integer length field describing the length of the following Data Unit in bytes

Signaling

- Parameters defined in the -00 draft:
  - Profile/Level ID
  - Resolution
  - Coverage (e.g. Field of View)
  - Anchor (an anchor point to relate multiple point clouds together)
  - Orientation (of the sensor)
  - Position (of the sensor)
  - Attributes

Wrapup and Next Steps

- Action Items
  - VC3: Bernard to summarize WGLC and pick a shepard
  - SCIP: Bernard to work with authors and IESG to address DISCUSS comments.
- Next Steps (authors)
  - Submission of draft-ietf-avtcore-rtp-j2k-scl

Start time: 11:20

End time: 11:30

Thank you

Special thanks to:

The Secretariat, WG Participants & ADs