AVTCORE WG
IETF 119

Hybrid Meeting
March 19, 2024
09:30 - 11:30 Brisbane Time
Session I, P1

Mailing list: avtcore@ietf.org
Notes: https://notes.ietf.org/notes-ietf-119-avtcore
MeetEcho link: Meetecho (ietf.org)

This session is being recorded

IETF 119 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*

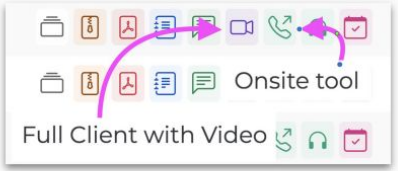

Remote participants

- 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 119 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 119

- Information about IETF 119
  https://www.ietf.org/how/meetings/119
- 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-119-avtcore/
- Notes: https://notes.ietf.org/notes-ietf-119-avtcore
- Secretariat: mtd@jabber.ietf.org
- WG Chairs (Remote): Bernard Aboba
- Onsite: Jonathan Lennox
- Zulip Scribe: Jonathan Lennox
- Note takers:

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.

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-avtext-lrr (EDIT! after 6 years, 8 months in MISSREF)
  - draft-ietf-payload-vp9 (EDIT! after 2 years, 9 months in MISSREF)
  - draft-ietf-avtcore-rtp-evc (EDIT)
  - draft-ietf-avtcore-rtp-scip (EDIT)
  - draft-ietf-avtext-framemarking (EDIT)
- WGLC
  - draft-ietf-avtcore-rtp-v3c (ended November 7, 2023)
    - WGLC summary posted

Draft Status (cont'd)

- Adopted
  - draft-ietf-avtcore-rtp-over-quic
  - draft-ietf-avtcore-rtcp-green-metadata
  - draft-ietf-avtcore-hevc-webrtc
  - draft-ietf-avtcore-rtp-j2k-scl
  - draft-ietf-avtcore-rtp-sframe

Chair Action Items

- **Chairs**
  - Call for adoption of Viewport and Region-of-Interest-Dependent Delivery of Visual Volumetric Media.
  - Figure out WG GitHub repo hierarchy, transfer individual repos there.
  - Find a place for store test vectors and the like

AVTCORE GitHub Setup

- Organization created: https://github.com/ietf-wg-avtcore
- Need to start creating / transferring repositories
- Also need to repoint the “activity this week” script

Author Action Items

- **HEVC-WebRTC**
  - Discuss JSEP one-way codec issues with RTCWEB WG which is not yet closed.
- **RTP over QUIC**
  - Identify the “Experimental” plan (questions and actions) to be followed up after publication of an Experimental RFC.
- **Green metadata**
  - WGLC once reviews are completed and issues are resolved.
- **V3C**
  - WGLC summary posted.
    - Christer found an issue in his review
    - Jonathan posted his review
    - GitHub repo: https://github.com/laurilo/draft-ilola-avtcore-rtp-v3c
    - Authors to respond to review comments (more today)

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
- Magnus Westerlund has written a short doc to update both RFCs on IANA Considerations and close the RTP payload registry.
- Issue tracked by MEDIAMAN WG
  - https://github.com/ietf-wg-mediaman/admin/issues/1

Agenda

1. Preliminaries (Chairs, 15 min)
   Note Well, Note Takers, Agenda Bashing, Draft status, Errata, IANA registries
2. Galois Counter Mode with Secure Short Tags (GCM-SST) (J.P. Mattsson, 15 min)
   https://datatracker.ietf.org/doc/html/draft-mattsson-cfrg-aes-gcm-sst
3. 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
4. RTP over QUIC (M. Engelbart, J. Ott, S. Dawkins, 20 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-over-quic
5. HEVC Profile for WebRTC (B. Aboba, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-hevc-webrtc
6. RTP Payload Format for SFrame (P. Thatcher, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-sframe
7. RTP Payload Format for Haptics (H. Yang, 10 min)
   https://datatracker.ietf.org/doc/html/draft-hsyang-avtcore-rtp-haptics
8. RTP Payload Format for Advance Professional Video (Y. Lim, 10 min)
   https://datatracker.ietf.org/doc/html/draft-lim-apv
9. RTP Payload Format for sub-codestream J2K streaming (P. Lemieux, 10 min)
   draft-ietf-avtcore-rtp-j2k-scl
10. Wrapup and Next Steps (Chairs, 10 min)

Galois Counter Mode with Secure Short Tags (GCM-SST)

draft-mattsson-cfrg-aes-gcm-sst

J. Preuß Mattsson

Start time: 09:45
End time: 10:00

AES with Galois Counter Mode (AES-GCM)

- AES-GCM is widely used due to its attractive performance and its provable security.
- During standardization, Ferguson pointed out two weaknesses in the GCM authentication function. The weaknesses are especially concerning when GCM is used with short tags:
  1. The first weakness significantly increases the probability of successful forgery.
  2. The second weakness reveals the subkey H if the attacker manages to create successful forgeries. With knowledge of the subkey H, the attacker always succeeds with subsequent forgeries. The probability of multiple successful forgeries is therefore significantly increased.
- As a comment to NIST, Nyberg, Gilbert, and Robshaw explained how small changes based on proven theoretical constructions mitigate the weaknesses.
- NIST did not follow the advice of Nyberg et al. and instead specified additional requirements for use with short tags in SP 800-38D Appendix C. Several cryptographers have criticized Appendix C and NIST has recently announced that they will remove Appendix C.
- While AES-CCM with short tags has forgery probabilities close to ideal, CCM has lower performance than GCM.

Every byte matters: the need for short tags

- 32-bit tags are standard in most radio link layers including 5G, 64-bit tags are very common in IoT transport and application layers, and 32-, 64-, and 80-bit tags are common in media-encryption applications.
- Audio packets are small, numerous, and ephemeral, so on the one hand, they are very sensitive in percentage terms to crypto overhead, and on the other hand, forgery of individual packets is not a big concern.
- Due to its weaknesses, GCM is typically not used with short tags. The result is either decreased performance from larger than needed tags, or decreased performance from using much slower constructions such as AES-CTR combined with HMAC.
- Short tags are also useful to protect packets transporting a signed payload such as a firmware and software updates.

Galois Counter Mode with Secure Short Tags (GCM-SST)

- Galois Counter Mode with Secure Short Tags (GCM-SST) is an AEAD algorithm following the recommendations from Nyberg et al.
- GCM-SST is defined with a general interface so that it can be used with any keystream generator, not just a 128-bit block cipher. AES-GCM-SST is a mode of operation of AES.
- The differences compared to GCM are that:
  1. GCM-SST uses an additional subkey Q. This enables short tags with forgery probabilities close to ideal.
  2. Fresh subkeys H and Q are derived for each nonce. This significantly decreases the probability of multiple successful forgeries.
  3. The POLYVAL function from AES-GCM-SIV is used instead of GHASH. POLYVAL is the "little-endian version" of GHASH and is more efficient in software implementations on little-endian architectures. GHASH and POLYVAL can be defined in terms of one another.
- ETSI SAGE and 3GPP have specified GCM-SST as the mode for future mobile networks. 5G Advance and 6G will use AES-256 and SNOW 5G in GCM-SST mode for "256-bit security" (requested by government customers).
  - Provides 10x higher performance on x86 in cloud-native deployments. 32–128-bit integrity tags.
- Strong interest in IETF for solutions like GCM-SST for use in media-encryption applications.

Authenticated encryption function

The performance of GCM-SST is very similar to GCM. The two additional AES invocations are compensated by the use of POLYVAL, the "little-endian version" of GHASH, which is faster on little-endian architectures. GCM-SST maintains the additive encryption characteristic of GCM, which enables efficient implementations on modern processor architectures,

Steps:

1. If the lengths of K, N, A, or P are not supported return error and abort
2. Initiate keystream generator with K and N
3. Let $H = Z[0]$, $Q = Z[1]$, $M = Z[2]$
4. Let $ct = P \oplus \text{truncate}(Z[3: n + 2], \text{len}(P))$
5. Let $S = \text{zeropad}(A) \,||\, \text{zeropad}(ct)$
6. Let $L = \text{LE64}(\text{len}(ct)) \,||\, \text{LE64}(\text{len}(A))$
7. Let $X = \text{POLYVAL}(H, S[0], S[1], \ldots)$
8. Let $full\_tag = \text{POLYVAL}(Q, X \oplus L) \oplus M$
9. Let $tag = \text{truncate}(full\_tag, tag\_length)$
10. Return (ct, tag)

For AES, $Z[i] = \text{AES-ENC}(K, N \,||\, \text{BE32}(i))$

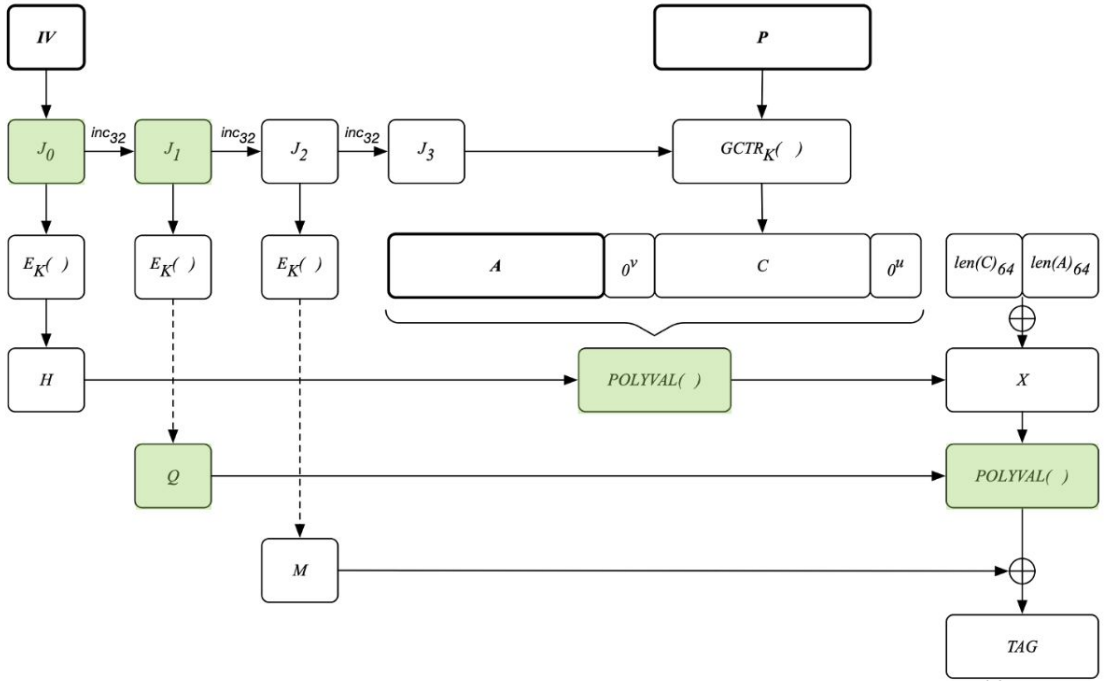

Authenticated decryption function

Steps:

1. If the lengths of K, N, A, or ct are not supported, or if len(tag) ≠ tag_length return error and abort
2. Initiate keystream generator with K and N
3. Let $H = Z[0]$, $Q = Z[1]$, $M = Z[2]$
4. Let S = zeropad(A) || zeropad(ct)
5. Let L = LE64(len(ct)) || LE64(len(A))
6. Let X = POLYVAL($H, S[0], S[1], \ldots$)
7. Let $full\_tag$ = POLYVAL($Q, X \oplus L$) $\oplus M$
8. Let $expected\_tag$ = truncate($full\_tag$, tag_length)
9. If $tag \neq expected\_tag$, return error and abort
10. Let $P = ct \oplus$ truncate($Z[3: n + 2]$, len(ct))
11. Return P

For AES, $Z[i]$ = AES-ENC(K, N || BE32(i))

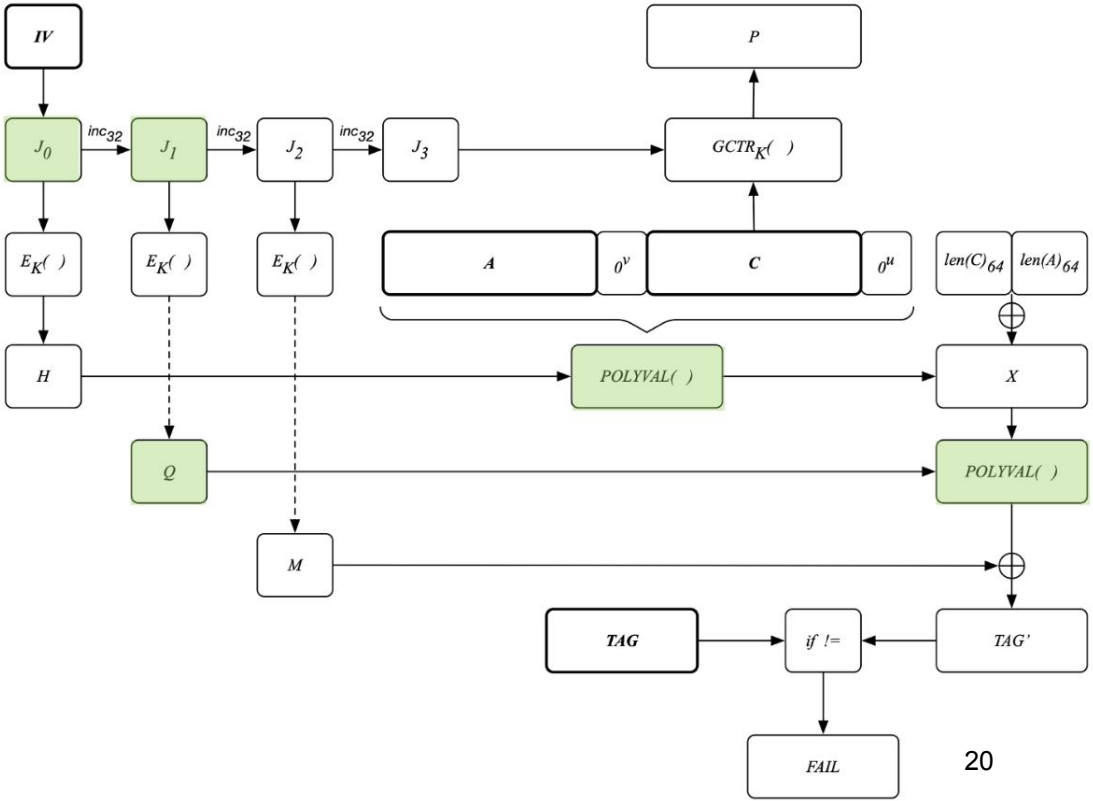

GCM-SST constrains and properties

- Performance is very similar to GCM. Two extra AES invocations are compensated by the faster POLYVAL.
- Tag size t ranges from 32 to 128 bits.
- For short tags of length $t < 128 - \log_2(n + m + 1)$ bits, the worst-case forgery probability is bounded by $\approx 2^{-t}$.
  - This is significantly better than GCM where the security level is only $t - \log_2(n + m + 1)$ bits.
- N_MIN and N_MAX (minimum and maximum size of the nonce) are both 12 octets, i.e., 96-bits.
- P_MAX (maximum size of the plaintext) is 2^{36} - 48 octets $\Rightarrow n \leq 2^{32} - 3$.
  - Adjusted as there are now three subkeys instead of two.
- A_MAX (maximum size of the associated data) is 2^{36} octets $\Rightarrow m \leq 2^{32}$.
  - Lowered to enable a forgery probability close to ideal for larger tags, even with maximum size P and A.
- With these constraints, $n + m + 1 < 2^{33}$ 128-bit blocks, and tags of length up to 95 bits have an almost perfect security level for all allowed plaintext and associated data lengths, i.e., the worst-case forgery probability is bounded by $\approx 2^{-t}$ where t is the tag length in bits.
- For a given key, the nonce must not be reused. Nonce reuse reveals the subkeys.

SRTP tag lengths

- Tags of length t = 32, 80, 96, and 128 bits have been registered for SRTP (MIKEY has variable tag length).
- The fast GCM only provides $t - 0.25 \cdot \log_2(\text{len(plaintext)} + \text{len(aad)})$ "bits of security". Should not be used with short tags.
  - Industry seems to have preferred fast algorithms and larger tags over small tags and slower algorithms (AES-CTR + HMAC-SHA-256 or AES-CCM).
- GCM-SST provides ≈ t "bits of security" for short tag. Allows combining a fast algorithm with short tags.
  - Which tag lengths would people want. Current GCM-SST document registers 32-, 64-, and 80-bit tags to align with SRTP and SFrame.

DTLS-SRTP:

- SRTP_AES128_CM_HMAC_SHA1_32
- SRTP_AES128_CM_HMAC_SHA1_80
- SRTP_AEAD_AES_128_GCM
- SRTP_AEAD_AES_256_GCM
- DOUBLE_AEAD_AES_128_GCM_AEAD_AES_128_GCM
- DOUBLE_AEAD_AES_256_GCM_AEAD_AES_256_GCM

SDP Security Descriptions:

- SEED_128_CCM_80
- SEED_128_GCM_96

Numeric ID	Name	K_LEN (bytes)	tag_length (bits)
TBD1	AEAD_AES_128_GCM_SST_4	16	32
TBD2	AEAD_AES_128_GCM_SST_8	16	64
TBD3	AEAD_AES_128_GCM_SST_10	16	80
TBD4	AEAD_AES_256_GCM_SST_4	32	32
TBD5	AEAD_AES_256_GCM_SST_8	32	64
TBD6	AEAD_AES_256_GCM_SST_10	32	80

Table 1: AEAD Algorithms

Summary

- GCM-SST is a small modification of GCM enabling short tags with forgery probabilities close to ideal.
- The changes are based on proven theoretical constructions and also work for stream ciphers.
- Performance-wise, it closely resembles GCM.
- Strong industry interest in a fast AES encryption mode with secure short tags.
- **Interesting for SRTP?**

Summary by Nick Youngson CC BY-SA 3.0 Alpha Stock Images

RTP Payload Format for Volumetric Video Coding (V3C)

draft-ietf-avtcore-rtp-v3c

L. Ilola

L. Kondrad

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

Update since #118

- Received feedback as part of WGLC - issues resolved
  - [#18] Issues with SDP grouping
    - Erroneous extension of the group attribute with optional parameters - not supported in RFC 5888.
    - The ability to add further parameters in the group attribute was removed from the draft making it compatible with the group attribute definition from RFC 5888.
    - Further clarification was added in the draft to suggest storing these parameters in the media format parameters of the V3C atlas component media.
    - This was already supported in the draft in previous version
  - [#17] Clarification of packetization-modes needed
    - As the packetization modes are not used in the spec, these should indeed not be used in the examples either.
    - Examples fixed
  - [#16] Questions about SDP examples
    - Some errors found in the SDP examples, e.g. missing mid attributes for the media lines and incorrect placing of other examples.
    - Errors in the examples have been resolved.

Update since #118

- Fixed some lingering issues while resolving WGLC comments
  - [#15] Removed unused abbreviations and added missing one
  - [#14] Instead of removing the V3C unit header syntax, its informative role was further clarified.
- Also some additional editorial improvements
  - Indentation of examples improved
  - Bit-alignment in the examples fixed
- As always new issues and suggestions for improvement are welcomed here

Open issues

Open issues can be found here

- [#19] Clarification of 2-D video stream encapsulation
- [#20] Reconsider the need for the three transmission modes (SRST, MRST and MRMT)
- [#21] Evaluate the need for out-of-order decoding (i.e. decoding order number, DON)
- [#22] V3C specific parameters can't be used with the 2D video media lines
- [#23] Clarify the nature V3C specific parameters (send properties vs. receive capabilities)

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: 10:10

End time: 10:30

Important updates since last interim

- Change to Experimental status and add Directions for Future work
- Guidance on Choosing QUIC Streams, QUIC DATAGRAMs, or a Mixture
- New section about Coalescing RTP packets in single QUIC packet
- Updated discussion subsection for 0-RTT
- Remove non-RTP/RTCP multiplexing (see next slides)

Multiplexing with other protocols #159 / #174

- We removed multiplexing with protocols that are not RTP or RTCP but kept the flow ID to simplify multiplexing
- Multiplexing RTP/RTCP with *other protocols* in the same connection was asked for in the past, e.g., to do signaling in the same connection.
- We don't know what *other protocols* should be multiplexed
- Allowing *any* protocol as long as it works with flow ID would bypass QUIC's ALPN requirement and require some custom ALPN on top of RoQ
- Also: Martin Thomson said we shouldn't assume our framework will be important enough to justify something super generic
- Future documents can build on RoQ to allow multiplexing RTP/RTCP with a particular protocol (example on next slides)

Multiplexing with other protocols #159 / #174

- draft-engelbart-quic-data-channels is a simple example protocol to be multiplexed with RoQ in the same QUIC connection
- One message per stream
- Optional sequence number for ordered messages
- Partial reliability by closing QUIC streams
- Open/Close messages to open new data channels with certain properties (ordered/unordered, reliable/unreliable, priority)
- ***Every stream starts with a channel ID that has the same format as RoQ's flow ID***

Multiplexing with other protocols #159 / #174

- draft-engelbart-multiplex-roq-qdc defines multiplexing of data channels and RTP in QUIC
- New ALPN (*roq-qdc-mux*)
- Two modes of multiplexing:
  - Use external signaling to assign flow/channel IDs (e.g., SDP)
  - Use one bit of the flow ID to distinguish RoQ or Data Channels
- Which mode to use is tbd, but could for example use different ALPNs

Merged Pull Requests

- Add considerations for coalescing RTP packets #152
- Explain multiple paths more clearly #154
- Clearly list the three alternatives #155
- Review appendix B #156
- Cleanup of many SHOULDs #158
- Move details about ECN and L4S into Section 7.1 #162
- Add references for actual IANA registries in Appendix B #164
- Convert ambiguous "datagrams" and lower-case "QUIC datagrams" to QUIC DATAGRAMs. #165
- Add guidance on choosing streams and datagrams #166
- Remove double period and space #167
- Clean up some terminology that is not actually used #169

Merged Pull Requests

- Remove out of context cross-reference #170
- Fix grammar #172
- Remove cross reference to removed subsection #173
- Remove non-RTP/RTCP protocol multiplexing #174
- Move motivations section later in the draft #176
- Clarify L4S dependencies for RoQ #177
- Add Futures section and change document category to "exp" (Experimental) #178
- Remove speculations about ICE with QUIC #179
- Clarify use of the word transport parameters #180
- Resolve editor's notes #183
- Describe considerations for early data with 0-RTT #184

Remaining Open Issues

- We have 1 open issues remaining as of 08 March 2024
  - General clean-up of the current draft #163

Who's Interested in These Future Documents?

- Connection Migration and Multipath
  - This relies on experience with RoQ and Multipath Extension for QUIC
- Multiplexing RTP with non-RTP protocols in a single QUIC connection
  - This relies on experience
- ICE and NAT traversal
  - Relying on QUIC working group to address this need
- Use of QUIC multicast
  - Relies on approval of something like Multicast Extension for QUIC

Time to revisit SDP for RTP over QUIC

- Spencer created draft-dawkins-avtcore-sdp-rtp-quic a LONG time ago
  - It seemed better to spend time on the rtp-over-quic specification first
  - Spencer let this expire, planning to circle back to it
- It's time to return to the SDP draft, and align it with RoQ
  - Spencer is interested in collaborators for this, of course

Next Steps

- Do people agree with the directions we've shared here?
- Do people agree RoQ as currently specified will be worth publishing?
  - We've talked about the need for a corresponding SDP specification
  - Are there any other show-stoppers that are in scope for AVTCORE?
- Does WGLC by IETF 120 seem realistic to the working group?
- ***Are We Almost Finished With This Specification???***

HEVC Profile for WebRTC

draft-ietf-avtcore-hevc-webrtc

Bernard Aboba

Philipp Hancke

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

For Discussion Today

- **Open Issues (https://github.com/aboba/hevc-webrtc/issues)**
  - **Issue 22: Issues with Receive-only codecs**

Issue 22: Issues with Receive-only codecs

- Chromium will not include H.265 in an Offer on a send/recv m-line, since it is currently unable to send H.265, only receive it. This results in Chromium sending and receiving another codec (e.g. H.264) when it sends an Offer and Safari Tech Preview answers (case 1). When Safari TP Offers and Chromium Answers, Safari with send H.265 to Chromium and will receive H.264 (case 2).

juberti commented 2 weeks ago

I think both Case 1 and Case 2 should work. While it could be argued that since H.265 is recvonly for Chrome, it shouldn't be offered on its sendrecv transceiver, that seems overly strict. Chrome should offer both H.265 and H.264 and then just choose to send H.264 while receiving H.265, the same way it would if, for some reason, it determined it had insufficient compute to send H.265 and had to fall back to H.264.

Generally, I think that any situation where you get different behavior depending on who is playing the offerer role is going to lead to a suboptimal experience; we should try to avoid such non-deterministic situations.

Issue 22: Issues with Receive-only codecs (cont'd)

- Discussion is occurring here:
  - AVTCORE:
    - WebRTC-HEVC Issue 22
  - W3C WEBRTC WG:
    - w3c/webrtc-pc#2936
    - w3c/webrtc-pc#2933
    - w3c/webrtc-pc#2935
    - w3c/webrtc-pc/#2888

RTP Payload Format for SFrame

draft-ietf-avtcore-rtp-sframe

Peter Thatcher

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

Update since last time

- Created -01 with where the default is to duplicate RTP header extensions, but it's optional:

  "The header extensions of the SFrame RTP packets SHOULD be the same as those of the output of the media-format-specific packetization, but some may be omitted if it is known that the omitted header extensions do not need to be duplicated on each SFrame RTP packet."

- But I haven't confirmed the submission yet. If you want to tweak the text, I can change it before confirming.

Next Steps?

RTP Payload for Haptics

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

H. Yang

X. de Foy

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

History

- Published Draft Version 00 on Oct 16
  - **Initial revision (Presented @ IETF 118)**
- Published Draft Version 01 on Feb 05
  - **Addressed feedback on the mailing list (Presented @ AVTCORE Interim)**
- Published Draft Version 02 on March 04
  - **Filled the security Section**

Status

- Published Draft Version 02 on March 04 (very minor update)
  - **Updated security consideration**
    - Added basic security Issue(RTP/End to End)
    - Added Haptic device(actuator / sensor) based security Issue
  - **Updated reference section**
    - AS-IS :
      - ISO/IEC, "ISO/IEC DIS 23090-31,Information technology - Coded representation of immersive media - Part 31: Haptics coding", ISO/IEC 23090-31, 2023)
    - To-Be :
      - ISO/IEC, "Text of ISO/IEC FDIS 23090-31 MPEG Haptics Coding", ISO/IEC 23090-31, 2024
  - **FDIS(Final Draft International Standard) document is ready**
    - **"Text of ISO/IEC FDIS 23090-31 MPEG Haptics Coding", ISO/IEC 23090-31**

Next steps & Plan

- The current version only updates the HMPG media subtype in section 6, because the top-level media type registration will be done by the draft “'haptics' Top-level Media Type” (https://datatracker.ietf.org/doc/draft-ietf-mediaman-haptics/)

- As pointed out by Jonathan on the mailing list, additionally we need to register the “haptics” Session Description Protocols (SDP) Parameter. **This will be added shortly in section 6.**

Next steps & Plan

- Processing Liaison document
- Suggestions and feedback are welcome
- We are looking for people interested in reviewing, implementing or participating in the draft.
- Waiting for WG adoption call

RTP Payload Format for Advanced Professional Video

draft-lim-apv

Youngkwon Lim

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

Motivation

- Smartphone cameras provide high quality image/video capture (e.g. multiple lenses, 100x zoom, 200M pixel, RAW formats)

- Increasing demands on beyond consumer quality supports by smartphone
  - Smartphones are much easier to carry and convenient to set-up
  - Content creators targeting SNS platform uses smartphone
  - Cloud-based post-production tools are getting popular

- Conventional video codecs such as HEVC, VVC or AV1 are not suitable for professional use as they have been designed for lossy compression of video for mass distribution

Motivation (cont'd)

- Comparison of professional video codec and conventional video codec

	professional video codec	conventional video codec
main purpose	Capturing video for post production (multiple rounds of various type of editing)	Encoding video for mass distribution
coding tools	Intra frame coding only	Intra & Inter frame coding
target quality	Visually lossless	Lossy
target bitrate	200 Mbps ~ 2 Gbps	1Mbps ~ 50Mbps

- Professional video codecs on the market
  - ProRes, REDCODE RAW, BRAW, ARRIRAW, Cinema DNG, SMPTE VC-5
  - No professional video codecs available as open standard, open source SW, royalty free

Advanced Professional Video (APV)

- Professional video codec designed for resource constrained devices
  - very low complexity (less than JPEG)
  - SW implementation on encoders on smartphones and desktop/laptop/cloud computers
- High throughput, high fidelity oriented
  - intra-only coding for supporting easy editing
  - tile structure coding for parallel processing
  - entropy coding for high-bitrate up to several Gbps
  - No loop/post filter to increase pixel precision

Advanced Professional Video (APV) (cont'd)

- Better compression efficiency than other professional video codec
  - around 20% better than well-known professional video codec for 2K, 4K video (4:2:2).

- Open standard-based IPR risk free and open source software
  - Using technologies more than 20-year-old only
  - Specification to be published as informational RFC
    - Review and comments will be appreciated
  - Open source to be made available by an industry forum
    - Any suggestions will be appreciated

RTP payload format for APV

- Use case for RTP transmission of APV data
  - Smartphone connected to cloud-based post-production
  - Wireless cameras

- I-D will be submitted before the IETF#120 for consideration by AVT Core.

High-performance JPEG 2000 RTP payload format

draft-ietf-avtcore-rtp-j2k-scl

P. A. Lemieux, D. Taubman

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

Update

- Specification update (P.-A. Lemieux)
- FPGA implementation update (D. Taubman)

Wrapup and Next Steps

- Action Items
- Next Steps (authors)

Start time: 11:20
End time: 11:30

Thank you

Special thanks to:

The Secretariat, WG Participants & ADs