AVTCORE WG

Virtual Interim

October 4, 2022

08:00 - 10:00 AM Pacific Time

Mailing list: avtcore@ietf.org

Meeting info:

https://datatracker.ietf.org/meeting/interim-2022-avtcore-03/session/avtcore

This session is being recorded

Virtual Interim 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.
- Keep audio and video off unless you are chairing or presenting.
- Use of a headset is strongly recommended.

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 and other meeting info:
  https://datatracker.ietf.org/meeting/interim-2022-avtcore-03/session/avtcore
- Jabber Room: avtcore@jabber.ietf.org
- Secretariat: mtd@jabber.ietf.org
- WG Chairs: Jonathan Lennox & Bernard Aboba
- Zulip Scribe: Jonathan Lennox
- Note takers: Spencer Dawkins, ?

Agenda

1. Note Well, Note Takers, Agenda Bashing, Draft status (Chairs, 10 min)
2. RTP Payload Format for SCIP (D. Hanson, 10 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-scip
3. RTP over QUIC Sandbox (B. Aboba, 25 min)
4. RTP over QUIC (J. Ott, M. Engelbart, 25 min)
   https://datatracker.ietf.org/doc/html/draft-ietf-avtcore-rtp-over-quic
5. RTP Control Protocol (RTCP) Messages for Green Metadata (Y. He, 15 min)
   https://datatracker.ietf.org/doc/html/draft-he-avtcore-rtcp-green-metadata
6. Wrapup and Next Steps (Chairs, 15 min)

Draft Status

- Published
  - RFC 9071: was draft-ietf-avtcore-multi-party-rtt-mix
  - RFC 9134: was draft-ietf-payload-rtp-jpegxs
- RFC Editor Queue
  - draft-ietf-payload-vp9 (MISSREF)
  - draft-ietf-avtcore-cryptex
    - IANA Review state: In progress
  - draft-ietf-avtcore-rtp-vvc
    - IANA Review state: Waiting on RFC Editor

Draft Status (cont'd)

- Waiting for AD Go-Ahead::Revised I-D Needed
  - draft-ietf-avtext-framemarking
- WGLC Completed, Waiting for WG Chair Go-Ahead : Proposed Standard
  - draft-ietf-avtcore-rtp-scip (WGLC completed May 8, 2022)
    - Gen-Art and Art-Art reviews posted
    - Secdir review pending
  - draft-ietf-avtcore-rfc7983bis (WGLC completed June 6, 2022)
    - WGLC announcement: https://mailarchive.ietf.org/arch/msg/avt/k8r222c7e06_5XjYNueFXQWejQI/
    - Revised I-D submitted: draft-ietf-avtcore-rfc7983bis-06
- Adopted
  - draft-ietf-avtcore-rtp-over-quic
  - draft-ietf-avtcore-rtp-evc

CfA on “RTP Payload Format for V3C”

- CfA announcement: https://mailarchive.ietf.org/arch/msg/avt/4SZNSxg6IjcAl00bNOUdwvCkYqM/
- Draft available here: https://datatracker.ietf.org/doc/html/draft-ilola-avtcore-rtp-v3c
- CfA runs until October 31, 2022

CfA on “Game State over RTP”

- Initial CfA Completed on May 8, 2022
- CfA summary: https://mailarchive.ietf.org/arch/msg/avt/w80E9ihE4rrJyMzU6F9eOYgvYVE/
- Two responses:
  - In favor:
    - Suhas Nandakumar: https://mailarchive.ietf.org/arch/msg/avt/c__oH22Gg-bImOAzVuYj3Ub_liU/
  - In favor of adopting the RTP payload format (but not the format itself):
    - Stephan Wenger: https://mailarchive.ietf.org/arch/msg/avt/R7iI2WZ1k2xz1jIjiQ-MNoPI9Xg/
- IETF 114 Chair proposal
  - Proponents to provide a plan for obtaining responses outside the IETF (e.g. game developer forums)
  - WG to extend the CfA
- Current status: draft expired. Do authors wish to continue?

RTP Payload Format for for SCIP

Dan Hanson

Mike Faller

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

SCIP Draft RFC – Status

- Revision 02 was submitted August 2
  - Incorporated comments from GENART and ARTART
- SECDIR reviewer comments received September 7
  - More comments regarding “opaque” standards
  - Security Considerations section: “may be adequate”
    - Basically boilerplate text
    - What needs to change?

SCIP Draft RFC – Status (2)

- We have an XML version finally available
  - Required intervention from ietf-tools id2xml group
  - Only editorial changes

Actions and Questions

- What actions are required to resolve issues with SECDIR, GENART, ARTART?
- Next steps?
- [Chairs] Authors need to respond to the reviewers:
  - Quoting the review, indicate the changes that have been made in response to specific comments (CC’ing the WG and the relevant directorate).
  - Reviewer response indicating acceptance is desirable.
  - Review/response record will be cited in the Publication Request.

RTP over QUIC Sandbox

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

Bernard Aboba

Why a Sandbox?

- To allow exploration of RTP over QUIC transport performance.
  - Goal is to ***visualize*** performance of RTP over QUIC.
  - Focus is on ***finding*** potential issues, not necessarily ***fixing*** them.
- To enable rapid experimentation:
  - Rapid turnaround enabled by pipeline modularity, small code size (JS only, no WASM)
- To explore issues that require complete systems. Examples:
  - Spec gaps, implementation bugs or unexpected behavior
  - Interactions between codecs and transport
  - E2E support for partial reliability
  - Multiplexing of data and media
  - Forwarder behavior (e,g. store/forward vs. cut-through)

What Can You Do With It?

- Experimenters can vary encoding parameters, codecs, bitrates, resolutions, etc.
  - Can visually compare local and remote (echo’d) video
  - Enables measurement of glass-glass latency
- Post-experiment diagnostics
  - Metrics: RTT stats, loss, reordering, etc.
    - Calculated at application layer
    - Chrome does not surface QUIC stack metrics yet
  - Graphs: RTT versus frame length

How was it built?

- Based on “next generation” Web media APIs: WHATWG Streams, Webcodecs, WebTransport, media-capture transform.
- Media pipeline (both send and receive) implemented in a (single) worker.
- Some features (such as “Bring Your Own Buffer”) not yet supported in Chrome Stable, so separate version for Chrome Canary.
- Not a complete implementation
  - Frame/stream transport only (no datagrams)
  - No application layer congestion control, RTCP support or jitter buffer
    - Experimenter can vary the average bitrate, look at measured RTT and loss.
  - Works surprisingly well, considering…

Two Sandbox Applications

1. Sandbox #1 encodes and decodes video in a WHATWG Streams pipeline without transport.
   a. Live site: https://webrtc.internaut.com/wc/wcWorker/
   b. Github repo: https://github.com/aboba/wc-demo/
2. Sandbox #2 adds network transport to the sending and receiving pipelines, bouncing encoded frames off an echo server in the cloud. Comparison with sandbox #1 can help isolate network effects.
   a. Live site (Chrome Stable): https://webrtc.internaut.com/wc/wtSender2/
   b. Live site (Chrome Canary): https://webrtc.internaut.com/wc/wtSender4/
   c. GitHub repo (Canary version): https://github.com/aboba/wt-demo

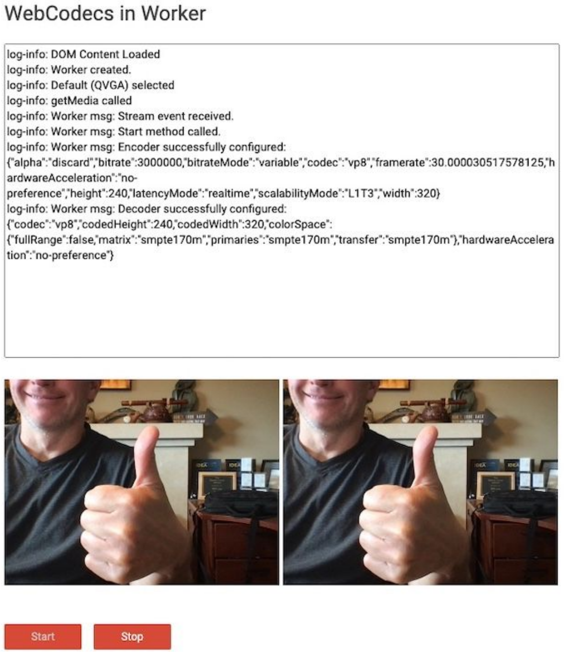

Parameters to Select

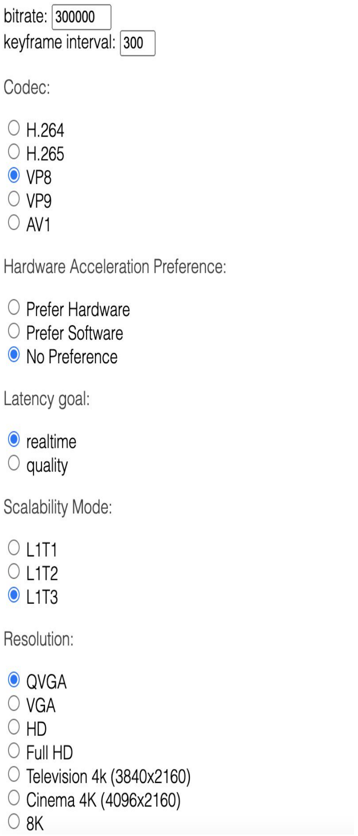

- Bitrate: “Average Target Bitrate” target provided to the encoder.
  - Actual bandwidth consumption is typically lower.
- Keyframe interval: number of frames between each keyframe.
- Codec: VP8, VP9, H.264 or AV1
  - Some oddities noted with VP9 (large frame size with “realtime”)
  - AV1 most solid on MacOS
  - H.265 not supported currently.
- Hardware Acceleration Preference: hw accelerated versus software encode/decode. Hw acceleration often not available.
- Latency goal: “quality” produces smaller frame sizes, but takes (marginally) longer than “realtime”.
- Scalability mode: how many temporal layers to use. Enables differential protection for the base layer.
- Resolution: reflected in getUserMedia constraints. If your camera doesn’t support the requested resolution, window will be blacked out.

High Level Observations

- Video quality
  - Quality highly dependent on device and camera.
  - Good quality possible with desktop/high quality notebook and appropriate settings.
  - Full-HD video (talking head) consumes < 1 Mbps.
- CPU Utilization
  - Higher resolutions (e.g. full-HD) or complex codecs can result in high CPU utilization.
- Resilience
  - QUIC reliable transport + temporal scalability provides good resilience.
    - QUIC stream/frame transport provides retransmission.
    - Temporal scalability enables partial reliability.
      - Non-base layer frames can be considered discardable.

High Level Observations (cont'd)

- Latency
  - Observed glass-glass latency ***considerably*** higher than measured frame RTT.
  - P-frames are typically small (a few packets) and exhibit low frame RTT.
  - I-frames are ***much*** larger (10X or more) and exhibit frame RTT multiple times higher (though not always).
    - Effect most pronounced with high GoP sizes (only a few I-frames per experiment)
    - Effect seen even under conditions of low bandwidth utilization and low loss.

Example

- AV1 @full-hd
  - Target average bitrate = 1 Mbps, GoP = 300
  - P-frame RTT ~ 100 ms with low jitter/no frame loss
  - For large frames, frame RTT multiple times higher
  - Glass-glass latency ~ 630 ms
  - Much lower glass-glass latency with “no network” sandbox

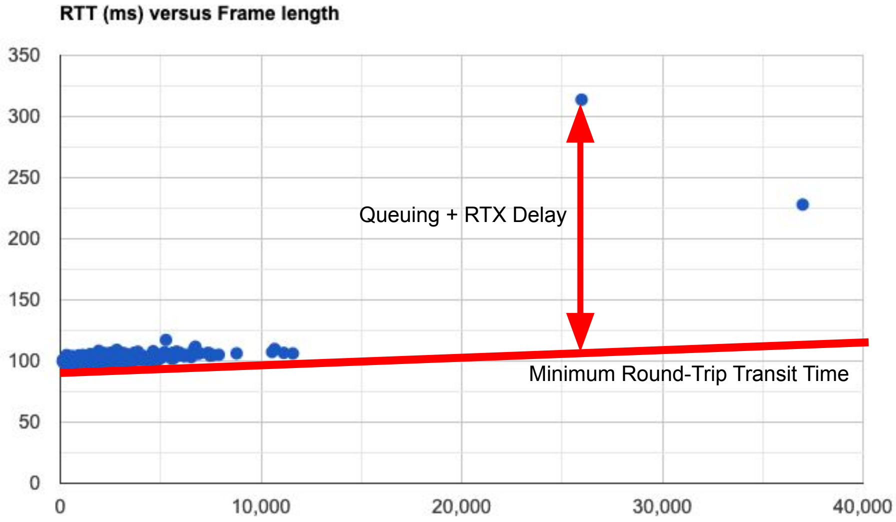

Local Video

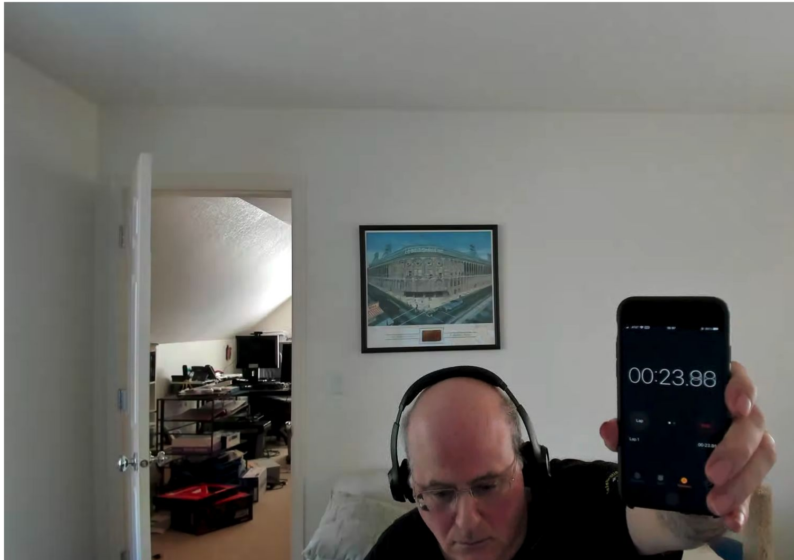

High Level Observations (cont’d)

- Potential cause for higher I-frame latency:
  - Poor Bandwidth Estimates
    - P-frames are “application limited”, preventing application from accurately estimating available bandwidth
    - Lack of a good BWE results in:
      - Congestion window too small to send the I-frame in a single RTT.
      - Application unable to adjust the I-frame size to fit BWE via average target bitrate or per-frame QP.
    - Probing can improve estimates (implemented in WebRTC)
- Anomalies
  - Latency preference of “quality” reduces large frame RTT compared with “realtime” due to reduction in frame sizes.
  - Subsequent I-frames show lower frame RTT due to increases in the congestion window.
  - Lowering GoP size decreases highest frame RTTs while increasing bandwidth utilization.

High Level Observations (cont’d)

- CPU utilization and threading
  - High resolutions and/or complex codecs can result in high CPU utilization
  - High glass-glass latency only observed when network transport is added.
    - Hypothesis: Potential interaction between encode/decode and transport?
      - Potential solution: support for multiple threads (e.g. separate threads for sender and receiver pipelines)?

Specification Issues

- Support for partial Reliability
  - Forwarding of FIN and RESET_STREAM frames.
  - Issue 39: Length field.
  - Issue 41: Allowing mixing of streams and datagrams
- Support for multiplexing of data and media
  - Use cases
    - Negotiation and configuration.
    - Communication of state or device input.
  - Issue 31: ALPN allocation.
- Issue 25: RTP topologies
  - Issues in translation between RTP/QUIC and RTP/UDP
    - Codec-specific knowledge required in translators??
    - How does SFrame to SPacket conversion work??
      - Middleboxes can only do “generic packetization”

Implications of Partial Reliability

- Temporal scalability enables “differential reliability” at the sender.
  - Sender can set timer, send RESET_STREAM if timer expires.
  - Timer set lower for “discardable” frames (extension layers).
  - High timer set for “non-discardable” frames (base layer) due to high loss penalty (e.g. generation of a 10X I-frame).
- Implications
  - Receiver needs to verify receiving a complete frame.
    - Useful to have a Length field at the beginning of the packet.
    - RFC 4571 Length field (16-bit) may not be large enough for high resolution I-frames (could be 200KB+)
  - Forwarder needs to forward RESET_STREAM frames.
    - Sender may send RESET_STREAM after FIN.
    - Forwarder may receive RESET_STREAM before or after FIN.

Support for Data + Media

- Multiple uses for multiplexing data and media over QUIC:
  - Negotiation: In client-server use cases (e.g. conferencing) can use the same QUIC connection for negotiation/configuration and media.
  - Data: Can send data over the same QUIC connection. Similar use cases to WebRTC data channel (state, device input, etc.)
- Supportable without multiple ALPNs. Precedents:
  - WebTransport ALPN can be used by the application, regardless of content (media, data, etc.). No mechanism needed in the API for setting the ALPN.
  - WebRTC uses a single (D)TLS ALPN for media and data.

RTP over QUIC

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

Mathis Engelbart, Jörg Ott

Updates since IETF 114

- Clarify stream usage (use unidirectional streams, immediately close stream after sending packet)
- Clarify flow identifier usage for datagram retransmissions
- Add exposing of bandwidth estimation to API considerations
- Remove some obsolete editor notes
- WIP:
  - Topology section
  - Stream concurrency
  - Expressing congestion control requirements instead of specifying algorithms

ALPN

- Issue #31
- Current draft specifies ALPN token “rtp-mux-quic” indicating possibility to multiplex RTP and other protocols over QUIC
- Other protocols will (have to) define their own ALPN token

  => Incompatible ALPN tokens
- Cannot define multiplexing with just *any* other protocol within “rtp-mux-quic”

Proposal:

- Define “rtp-quic”
- Future documents can specify multiplexing RTP/RTCP with other protocols and define new ALPNs

  (similar approach as in RFC 8833: ALPN for WebRTC)

Alternative

- Define “rtp-mux-quic”
- How to multiplex RTP/RTCP and other protocols?

Multiplexing RTP/RTCP/…

- Issue #24
- Current draft has *Flow ID* for multiplexing many RTP/RTCP/… flows
- Flow ID may be incompatible with other protocols that should run over the same QUIC connection
- Depends on ALPN (Issue #31)

“rtp-quic”:

- Only allow RTP/RTCP
- Multiplexing:
  - RTP/RTCP (RFC 5761)
  - Multiple types of media in a single RTP session (RFC 8860)
  - Still need a flow identifier?

“rtp-mux-quic”

- Allow RTP/RTCP and other protocols
- How to define Multiplexing?
  - Flow ID?
  - RFC 7983bis?
  - …?

Length field in QUIC Streams

- Issue #39
- Length field in RFC 4571 is not only useful for framing
    - Buffer allocation
    - Could in QUIC be used to identify incomplete frames
- 16 bit Length field will likely be too short for large ADUs
- Proposal: Add length field before RTP packet in QUIC streams
    - 32 bit Field?
    - 16 bit Field but denote size in units of 4 octets?
    - QUIC variable length integer?

Mixing Streams and Datagrams

- Issue #41
- Current draft only allows to choose between streams and datagrams but not mixing both in a single RTP session
- Mixing would allow partial reliability for example:
  - I-Frames in streams, P-Frames in datagrams
  - SVC base layer in streams, additional layers in datagrams
- Mixing may introduce synchronization issues
  - e.g., need to wait for the referenced I-Frame after receiving a P-Frame

- Pros? Cons?

Next Steps

- Resolve open issues
- Submit new draft
- Working further with Spencer on SDP signaling

RTP Control Protocol Messages (RTCP) for Green Metadata

https://datatracker.ietf.org/doc/html/draft-he-avtcore-rtcp-green-metadata

Yong He (Qualcomm), Waqar Zia (Qualcomm), Christian Herglotz (FAU), Edouard Francois (InterDigital)

Overview

- ISO/IEC 23001-11, Energy Efficient Media Consumption (green metadata) is specified by MPEG
- Interactive metadata for decoder-power reduction
  - A format is needed for this to carry these messages
    - updated draft version: https://www.ietf.org/archive/id/draft-he-avtcore-rtcp-green-metadata-01.html
    - Specifies a new RTCP payload format for
      - spatial and temporal resolution request and
      - notification feedback message

New messages

- AVPF [RFC4585][RFC5104] defines seven payload-specific feedback messages and one application layer feedback message
- This document specifies 2 new payload-specific FB messages
- Message may be sent
  - in a regular full compound RTCP packet or
  - in an early RTCP packet

```
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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Seq nr.     |           Reserved            |   Frame Rate  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Picture Width           |     Picture Height          |0 0 0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
        Syntax of an FCI Entry in the TSRR Message
```

Temporal-Spatial Resolution Request

```
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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Seq nr.     |           Reserved            |   Frame Rate  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Picture Width           |     Picture Height          |0 0 0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
        Syntax of an FCI Entry in the TSRN Message
```

Temporal-Spatial Resolution Notification (TSRN)

Draft updates

- New FMT value 11 and 12 are assigned to Temporal-Spatial Resolution Request message and Temporal-Spatial Resolution Notification message (9 and 10 were already assigned to ROI and LRR).
- Semantic update to prohibit the value of spatial resolution or frame rate equal to 0
- Text on existing FMT value assignment was removed
- Typos in green metadata reference URL and author's email address were corrected

Value	Name	Long Name
1	PLI	Picture Loss Indication
2	SLI	Slice Loss Indication
3	RPSI	Reference Picture Selection Indication
4	FIR	Full Intra Request Command
5	TSTR	Temporal-Spatial Trade-off Request
6	TSTN	Temporal-Spatial Trade-off Notification
7	VBCM	Video Back Channel Message
8	PSLEI	Payload-Specific Third-Party Loss Early Indication
9	ROI	Video region-of-interest (ROI)
10	LRR	Layer Refresh Request Command
11-14		Unassigned
15	AFB	Application Layer Feedback
16-30		Unassigned
31	Extension	Reserved for future extensions

Next step

- We would like to ask for WG adoption of the draft

Wrapup and Next Steps

- Action Items
  - Item 1
  - Item 2
- Next Steps
  - Step 1
  - Step 2

Thank you

Special thanks to:

The Secretariat, WG Participants & ADs