

RTP Frame Acknowledgement

https://github.com/ietf-wg-avtcore/frame-acknowledgement

Erik Språng

Gurtej Singh Chandok

Shridhar Majali

Draft Status



Progress on GitHub

- Open issues: **7**
- Closed issues: **21**
- Commits: **89**

Notable Issues

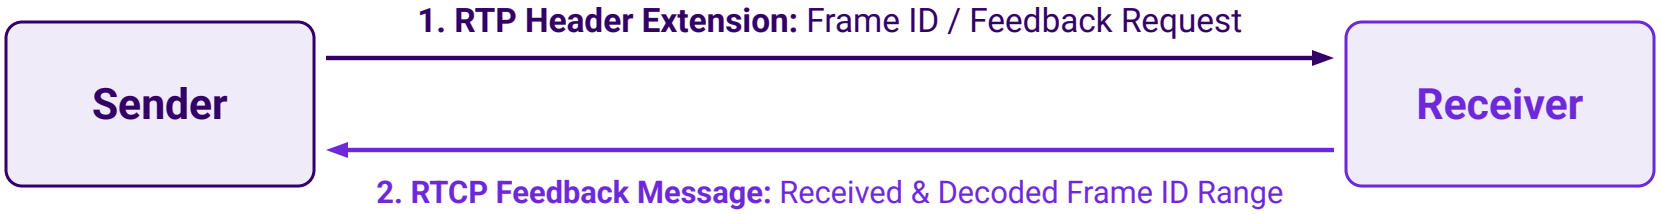
- Add examples of how to identify a frame **#43**
- Decide on American vs British spelling **#53**
- Need an indication on receiver that frame is a resync (LTR-P) frame **#54**

Recap

Frame Acknowledgment consists of two parts:

1. An **RTP header extension** used to mark frames for feedback (with a Frame ID) and/or request feedback about the status of one or more Frame IDs.

2. An **RTCP feedback message** used to indicate whether a receiver has received and decoded frames within a Frame ID range.



Add examples of how to identify a frame #43

Discussed at last IETF (126)



A **frame** is defined as “any independently **decodable unit of bitstream data** that results in an update to the codec state that can be used as a reference for any subsequent decodable unit of bitstream data”.

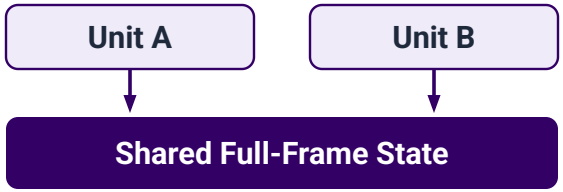
If the decodable unit is smaller than a frame, **should it still be called a frame?**

Add examples of how to identify a frame #43

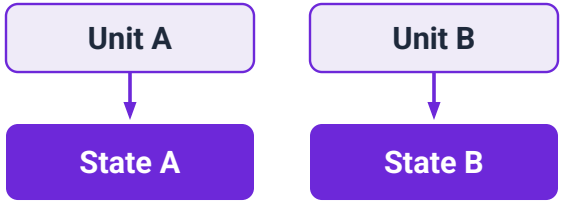
Discussed at last IETF (126)

The current draft still calls it a frame because in practice we seem to have only two cases:

1. The independently decodable data can't actually update the codec state independently, only as **part of decoding the full frame** (the case for e.g., H.265 MCTS) when used in the context of RTP.



2. If the unit *can* be decoded independently, it has a **fully independent state** (e.g., set of reference buffers). H.266 calls this "subpictures"; AV2 calls it "extended layers". No one clear word exists, so calling it a frame in the context of that subset still makes sense.



Decide on American vs British spelling #53



The draft ID is currently `draft-ietf-avtcore-frame-acknowledgement` (British spelling).

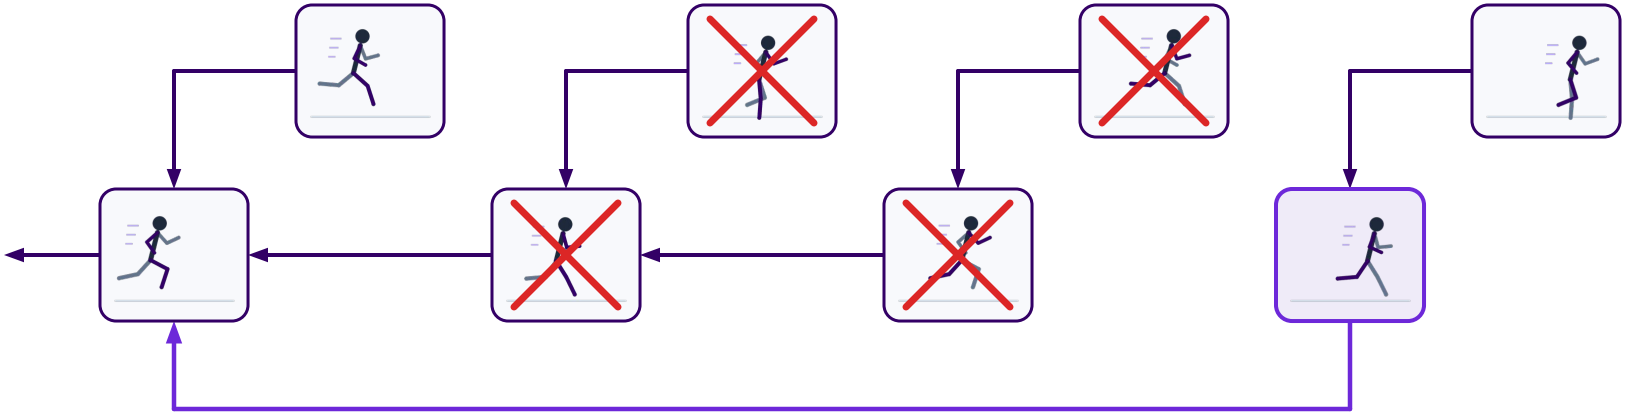
The text of the draft contains a mix of American and British spellings. draft-rpc-rfc7322bis-03 states:

> The RFC publication language is English. Spelling may be either American or British, as long as an individual document is internally consistent. Where both American and British English spelling are used within a document or cluster of documents, the text will be modified to be consistent with American English spelling.

Should we change both the text and draft identifier to American English?

Need an indication on receiver that frame is a resync (LTR-P) frame #54

If an LTR frame is encoded, a receiver must be able to tell that it can fast-forward from a given reference to that frame. The question is, **should this draft specify** how to do that?



Need an indication on receiver that frame is a resync (LTR-P) frame #54



Currently two viable strategies:

Dependency Descriptors

Defined as part of the AV1 RTP profile, but codec-agnostic. Works with E2E encryption.

Codec Specifics, e.g.:

- VP9: RTP payload header
- H.265: Reference Picture Set (Slice Header)

Is defining a new message in scope for this draft?

Hackathon results

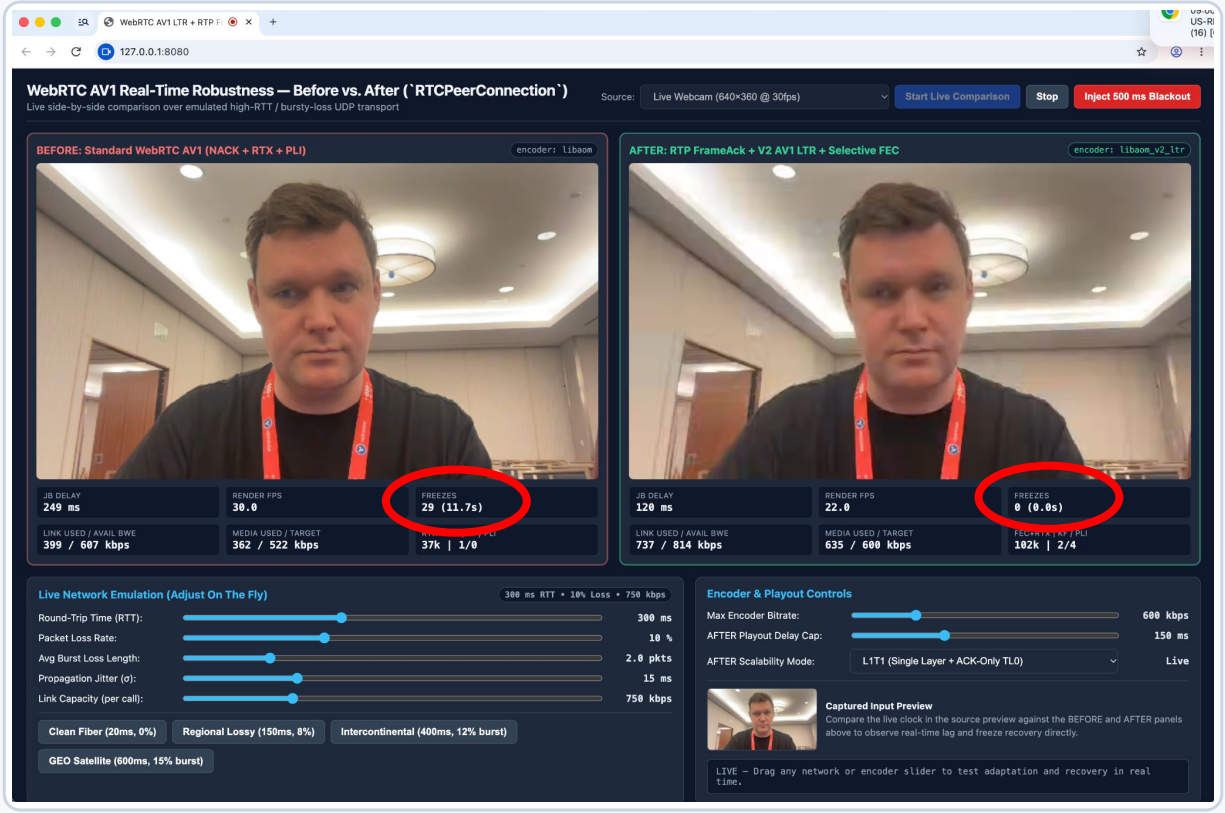


Full implementation in libwebrtc/Chromium:

- Header extension & RTCP message parsing/generation
- SDP negotiation
- Send/receive-side logic
- Simplistic LTR structure implementation (“ping-pong style”, always referencing a known-good state)
- Demo page using two PeerConnection loopbacks and knobs to control network throttling/loss
- Simple Python user-land network throttler



Hackathon results



Hackathon results

Things learned



If the feedback says a frame is NOT decoded, the spec doesn't clearly state what the receiver should do. Filed issue **#55**.

Receive-side logic is relatively straightforward, only the "guaranteed decodability" requirement required some major effort.

As expected, the encoder logic and robustness/quality tradeoff is the really difficult part - but the spec allows this to be implementation specific.

Questions



Next steps:

- Address remaining GitHub issues.
- Broader review by working group.
- Continue implementation work.

Thank you!