

RTP Splicing Status Update

draft-ietf-avtext-splicing-for-rtp-11

Jinwei Xia

Status

- Documents second IESG telechat on the 11th of October
- IESG raised some remaining issues
- Two draft updates have been submitted since then (version 10 and 11) to resolve issues
- This solved all but two issues.

Stewart Bryant

- Discuss on RTP Payload security mechanism text. The text indicated that this was not compatible with splicing point indication based on RTP payload content
 - Misunderstood at first to refer to security contexts
- NEW Text proposal:
 - Some RTP deployments uses RTP payload security mechanisms (e.g., ISMACryp [ISMACryp]). If any payload internal security mechanisms are used, only the RTP sender and the RTP receiver establishes that security context, in which case, any middlebox (e.g., splicer) between the RTP sender and the RTP receiver will not get such keying material. This may impact the splicers's possibility to perform splicing if it is dependent on RTP payload level hints for finding the splice in and out points. However, other potential solutions exist to specify or mark where the splicing points exist in the media streams. When using RTP payload security mechanisms SRTP or other security mechanism at RTP or lower layers can be used to provide integrity and source authentication between the splicer and the RTP receiver.

Pete Resnick's Issue One

- REQ-7: There at least needs to be a forward reference to section 4.5 in this section. That said, I am still bothered by this section. The first sentence is written in the passive voice, so I can't tell what it's saying: Is it saying that splicing should not be easily detected *by the RTP receiving implementation* (in which case section 4.5 is really all that this is talking about), or is it saying that splicing should not be easily detected *by the RTP stream end-user* (in which case it's talking about section 4.3, which I think is bogus)? The second half of REQ-7 talks about both, but then goes on to say that this memo is only concerned with the RTP layer. Please rewrite this section to make it clear that you are only talking about the protocol, not the user experience.
 - Clarified to be RTP level only
 - Wording was a bit unclear, new text prepared

Issue One Proposed Text

- REQ-7:
 - The RTP receiver should not be able to detect any splicing points in the RTP stream produced by the splicer on RTP protocol level. For the advertisement insertion use case, it is important to make it difficult for the RTP receiver to detect where an advertisement insertion is starting or ending from the RTP packets, and thus avoiding the RTP receiver from filtering out the advertisement content. This memo only focuses on making the splicing undetectable at the RTP layer. The corresponding processing is depicted in section 4.5.

Pete Resnick's Issue Two

- Section 4.3: Similar to the above, this is really all user experience, not protocol clipping. I think this section should simply be dropped. Section 4.5: "User Invisibility" is not the requirement in 4.5; it is "Receiving implementation invisibility."
 - First of all it has been clarified that it is not about implementation invisibility
 - Secondly I have proposed rewording of the text to use other terminology and focus on possible implementation issues.

Issue Two Proposed Text

- 4.3. Considerations for Handling Media Substitution at the RTP Layer

This section provides informative guidelines on how to handle media substitution at both the RTP layer to minimize media impact. Dealing with the media substitution well at the RTP layer is necessary for quality implementations. To perfectly erase any media impact needs more considerations at the higher layers, how the media substitution is erased at the higher layers are outside of the scope of this memo.

Issue Two Proposed Text

If the time duration for any substitutive content mismatches, i.e. shorter or longer, than the duration of the main content to be replaced, then media degradations may occur at the splicing point and thus impact the user's experience.

If the substitutive content has shorter duration from the main content, then there could be a gap in the output RTP stream. The RTP sequence number will be contiguous across this gap, but there will be an unexpected jump in the RTP timestamp. Such a gap would cause the receiver to have nothing to play. This may be unavoidable, unless the mixer can adjust the splice in or splice out point to compensate. This assumes the splicing mixer can send more of the main RTP stream in place of the shorter substitutive stream, or vary the length of the substitutive content. It is the responsibility of the higher layer protocols and the media providers to ensure that the substitutive content is of very similar duration as the main content to be replaced.

Issue Two Proposed Text

If the substitute content have a duration longer than the reserved gap duration, there will be an overlap between the substitutive RTP stream and the main RTP stream at the splicing out point. A straightforward approach is that the mixer performs an ungraceful action, terminating the splicing and switching back to main RTP stream even if this may cause media stuttering on receiver. Alternatively, the mixer may transcode the substitutive content to play at a faster rate than normal, to adjust it to the length of the gap in the main content, and generate a new RTP stream for the transcoded content. This is a complex operation, and very specific to the content and media codec used. Additional approaches exists, these type of issues should be taken into account in both mixer implementors and media generators to enable smooth substitutions.

Pete Resnick's Issue Three

- Section 4.5 does not address user experience invisibility and should be re-titled.
 - Resolved by v11 text

Next Steps

- Awaiting confirmation from ADs on new text resolving issues
- New draft will be submitted when above completed
- 1 Week WG last call on changes
- If all Ok, then AD approves publication