

RTP Usage for CLUE

draft-lennox-clue-rtp-usage-04
Clue WG Interim — 7 June 2012

Jonathan Lennox jonathan@vidyo.com

Allyn Romanow allyn@cisco.com

Paul Witty paul.witty@balliol.oxon.org

RTP Usage: What's New since Paris

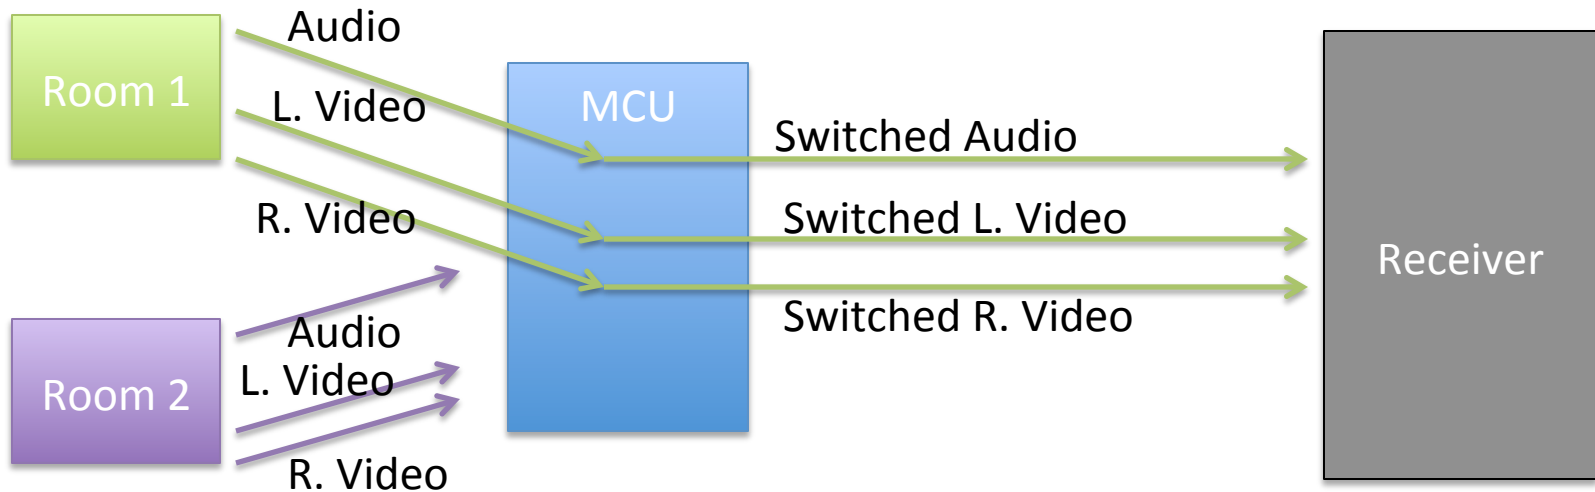
- Added one new media requirement.
- Expanded discussion of multiple transport flows.
- Clarified several other requirements, following the Paris discussion.
- In this presentation, I'll describe the new requirement, then describe the architecture I envision and how it emerges from the requirements.

New Requirement: Media-12

“If multiple sources from a single synchronization context are being sent simultaneously, it must be possible for a receiver to associate and synchronize them properly, even for sources that are mapped to switched captures.”


- RTCP SR and CNAME messages provide synchronization of RTP sources (SSRC values). To do correct synchronization, we need to be able to correlate the sources that are switched, not just the switched captures.

Synchronization: example



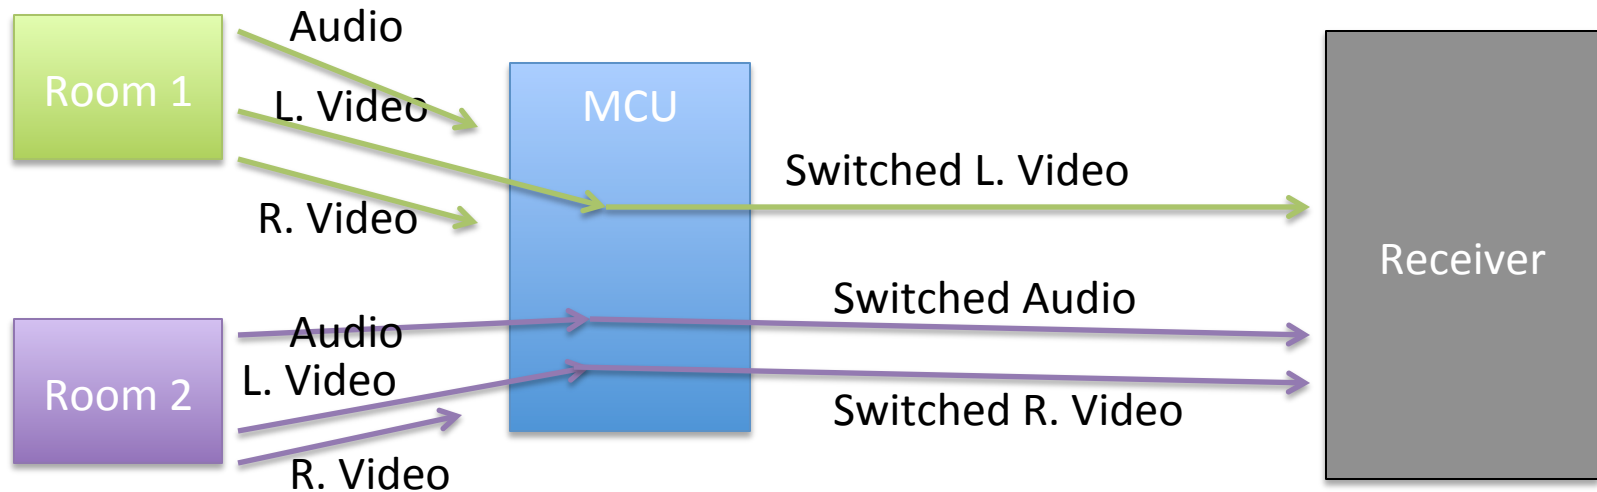
- Synchronize media from the same room.

Synchronization: example



- Don't synchronize media from different rooms.

Synchronization: example



- Don't synchronize media from different rooms.

Multiple transport flows

- There are cases where you want to use multiple transport flows
 - Quality-of-Service
 - Disaggregated media
 - Backward compatibility
- This isn't a problem for media-level, but has implications for CLUE signaling
 - Providers might say "I can only provide this capture encoding on these transports."
 - Consumers might say "I want to receive this capture encoding on this transport."

Proposed Architecture (1)

- Single RTP session, on a single UDP transport flow (per media type), between two participants; use RTP source multiplexing to send different sources as separate SSRCs.
 - Unless you need multiple transport flows: then, use one RTP session per transport flow.
 - Requirements Media-1 (constant number of flows) and Media-2 (can use multiple flows when needed); also, to some extent, Media-10 (don't redefine existing protocols).

Proposed Architecture (2)

- MCU is an RTP Media Translator for static and switched captures; RTP Mixer for (locally-generated) composited captures.
 - This preserves the identity of streams switched into switched captures.
 - Necessary for Req. Media-8 (move sources among switched captures); Media-9 (send a source only once, for multiple capture encodings); Media-12 (identify sources for synchronization).
- Endpoints generating switched captures behave the same.

Proposed Architecture (3)

- Because SSRCs are associated to actual sources, need additional information to associate them with the requested capture encodings.
- Provide the association (encoding ID) in an RTP header extension, so this information is available immediately (Req. Media-6).

Why doesn't all-Mixer work?

- It would be possible for an MCU to assign a single SSRC per capture encoding.
- The MCU then acts as a mixer, in Topo-Mixer (or Topo-Video-Switched-MCU), with a single CSRC per switched-capture SSRC.
- However:
 - You can't do SR-based synchronization between switched captures, without all switched captures always being synchronized (i.e., a single CNAME / SR timestamp space).
 - Decoder state can't persist when sources move among switched captures.
 - You can't send a single source to satisfy multiple capture encodings.

BACKUP SLIDES: REQUIREMENTS

Req. Media-1

“It must not be necessary for a Clue session to use more than a single transport flow for transport of a given media type (video or audio).”

Req. Media-2

“It must, however, be possible for a Clue session to use multiple transport flows for a given media type where it is considered valuable (for example, for distributed media, or differential quality-of-service).”

Req. Media-3

“It must be possible for a Clue endpoint or MCU to simultaneously send sources corresponding to static, to composited, and to switched captures, in the same transport flow. (Any given device might not necessarily be able send all of these source types; but for those that can, it must be possible for them to be sent simultaneously.)”

Req. Media-4

“It must be possible for an original source to move among switched captures (i.e. at one time be sent for one switched capture, and at a later time be sent for another one).”

Req. Media-5

“It must be possible for a source to be placed into a switched capture even if the source is a ‘late joiner’, i.e. was added to the conference after the receiver requested the switched source.”

Req. Media-6

“Whenever a given source is assigned to a switched capture, it must be immediately possible for a receiver to determine the switched capture it corresponds to, and thus that any previous source is no longer being mapped to that switched capture.”

Req. Media-7

“It must be possible for a receiver to identify the actual source that is currently being mapped to a switched capture, and correlate it with out-of-band (non-Clue) information such as rosters.”

Req. Media-8

“It must be possible for a source to move among switched captures without requiring a refresh of decoder state (e.g., for video, a fresh I-frame), when this is unnecessary. However, it must also be possible for a receiver to indicate when a refresh of decoder state is in fact necessary.”

Req. Media-9

“If a given source is being sent on the same transport flow for more than one reason (e.g. if it corresponds to more than one switched capture at once, or to a static capture), it should be possible for a sender to send only one copy of the source.”

Req. Media-10

“On the network, media flows should, as much as possible, look and behave like currently-defined usages of existing protocols; established semantics of existing protocols must not be redefined.”

Req. Media-11

“The solution should seek to minimize the processing burden for boxes that distribute media to decoding hardware.”

- In many architectures, a single box receives media traffic, then farms it out to decoding hardware.
- The protocol should make this as light-weight as possible (subject to other constraints).

Req. Media-12

“If multiple sources from a single synchronization context are being sent simultaneously, it must be possible for a receiver to associate and synchronize them properly, even for sources that are mapped to switched captures.”