

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 6, 2016

P. Kyzivat
L. Xiao
C. Groves
Huawei
R. Hansen
Cisco Systems
August 5, 2015

CLUE Signaling
draft-ietf-clue-signaling-06

Abstract

This document specifies how CLUE-specific signaling such as the CLUE protocol [[I-D.ietf-clue-protocol](#)] and the CLUE data channel [[I-D.ietf-clue-datachannel](#)] are used with each other and with existing signaling mechanisms such as SIP and SDP to produce a telepresence call.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 6, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Media Feature Tag Definition	4
4.	SDP Grouping Framework CLUE Extension Semantics	4
4.1.	General	4
4.2.	The CLUE data channel and the CLUE grouping semantic	4
4.3.	CLUE-controlled media and the CLUE grouping semantic	5
4.4.	SDP semantics for CLUE-controlled media	5
4.4.1.	Signalling CLUE Encodings	5
4.4.1.1.	Referencing Encodings in the CLUE protocol	6
4.4.1.2.	Media line directionality	7
4.4.2.	Negotiating receipt of CLUE Capture Encodings in SDP	7
4.5.	SDP Offer/Answer Procedures	7
4.5.1.	Generating the Initial Offer	7
4.5.2.	Generating the Answer	8
4.5.2.1.	Negotiating use of CLUE and the CLUE data channel	8
4.5.2.2.	Negotiating CLUE-controlled media	8
4.5.2.3.	Negotiating non-CLUE controlled media	9
4.5.3.	Processing the initial Offer/Answer negotiation	9
4.5.3.1.	Successful CLUE negotiation	9
4.5.3.2.	CLUE negotiation failure	10
4.5.4.	Modifying the session	10
4.5.4.1.	Adding and removing CLUE-controlled media	10
4.5.4.2.	Enabling CLUE mid-call	10
4.5.4.3.	Disabling CLUE mid-call	11
5.	Interaction of CLUE protocol and SDP negotiations	11
5.1.	Independence of SDP and CLUE negotiation	11
5.2.	Constraints on sending media	12
5.3.	Recommendations for operating with non-atomic operations	12
6.	Interaction of CLUE protocol and RTP/RTCP CaptureID	13
6.1.	CaptureID reception during MCC redefinition	14
7.	Multiplexing of CLUE-controlled media using BUNDLE	14
7.1.	Overview	14
7.2.	Usage of BUNDLE with CLUE	15
7.2.1.	Generating the Initial Offer	15
7.2.2.	Bundle Address Synchronization	15
7.2.3.	Multiplexing of the data channel and RTP media	15
8.	Example: A call between two CLUE-capable Endpoints	16
9.	Example: A call between a CLUE-capable and non-CLUE Endpoint	24
10.	Acknowledgements	25
11.	IANA Considerations	25
11.1.	New SDP Grouping Framework Attribute	25

11.2.	New SIP Media Feature Tag	26
12.	Security Considerations	26
13.	Change History	27
14.	References	32
14.1.	Normative References	32
14.2.	Informative References	33
Authors'	Addresses	34

[1.](#) Introduction

To enable devices to participate in a telepresence call, selecting the sources they wish to view, receiving those media sources and displaying them in an optimal fashion, CLUE involves two principal and inter-related protocol negotiations. SDP, conveyed via SIP, is used to negotiate the specific media capabilities that can be delivered to specific addresses on a device. Meanwhile, a CLUE protocol [[I-D.ietf-clue-protocol](#)], transported via a CLUE data channel [[I-D.ietf-clue-datachannel](#)], is used to negotiate the Capture Sources available, their attributes and any constraints in their use, along with which Captures the far end provides a device wishes to receive.

Beyond negotiating the CLUE channel, SDP is also used to negotiate the details of supported media streams and the maximum capability of each of those streams. As the CLUE Framework [[I-D.ietf-clue-framework](#)] defines a manner in which the Media Provider expresses their maximum encoding capabilities, SDP is also used to express the encoding limits for each potential Encoding.

Backwards-compatibility is an important consideration of the document: it is vital that a CLUE-capable device contacting a device that does not support CLUE is able to fall back to a fully functional non-CLUE call. The document also defines how a non-CLUE call may be upgraded to CLUE in mid-call, and similarly how CLUE functionality can be removed mid-call to return to a standard non-CLUE call.

[2.](#) Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

This document uses terminology defined in the CLUE Framework [[I-D.ietf-clue-framework](#)].

A few additional terms specific to this document are defined as follows:

non-CLUE device: A device that supports standard SIP and SDP, but either does not support CLUE, or that does but does not currently wish to invoke CLUE capabilities.

CLUE-controlled media: A media "m" line that is under CLUE control; the Capture Source that provides the media on this "m" line is negotiated in CLUE. See [Section 4](#) for details of how this control is signalled in SDP. There is a corresponding "non-CLUE-controlled" media term.

[3.](#) Media Feature Tag Definition

The "sip.clue" media feature tag indicates support for CLUE. A CLUE-capable device SHOULD include this media feature tag in its REGISTER requests and OPTION responses. It SHOULD also include the media feature tag in INVITE and UPDATE [[RFC3311](#)] requests and responses.

Presence of the media feature tag in the contact field of a request or response can be used to determine that the far end supports CLUE.

[4.](#) SDP Grouping Framework CLUE Extension Semantics

[4.1.](#) General

This section defines a new SDP Grouping Framework extension, CLUE.

The CLUE extension can be indicated using an SDP session-level 'group' attribute. Each SDP media "m" line that is included in this group, using SDP media-level mid attributes, is CLUE-controlled, by a CLUE data channel also included in this CLUE group.

Currently only support for a single CLUE group is specified; support for multiple CLUE groups in a single session is beyond the scope of this document. A device MUST NOT include more than one CLUE group in its SDP unless it is following a specification that defines how multiple CLUE channels are signalled, and is either able to determine that the other side of the SDP exchange supports multiple CLUE channels, or is able to fail gracefully in the event it does not.

[4.2.](#) The CLUE data channel and the CLUE grouping semantic

The CLUE data channel [[I-D.ietf-clue-datachannel](#)] is a bidirectional SCTP over DTLS channel used for the transport of CLUE messages. This channel must be established before CLUE protocol messages can be exchanged and CLUE-controlled media can be sent.

The data channel is negotiated over SDP as described in the relevant document. A CLUE-capable device wishing to negotiate CLUE MUST also

include a CLUE group in the SDP and include the "mid" of the "m" line for the data channel in that group. A CLUE group MUST include the "mid" of the "m" line for one (and only one) data channel.

Presence of the data channel in a CLUE group in an SDP offer or answer also serves, along with the "sip.clue" media feature tag, as an indication that the device supports CLUE and wishes to upgrade the call to include CLUE-controlled media. A CLUE-capable device SHOULD include a data channel "m" line in offers and, when allowed by [\[RFC3264\]](#), answers.

[4.3.](#) CLUE-controlled media and the CLUE grouping semantic

CLUE-controlled media lines in an SDP are "m" lines in which the content of the media streams to be sent is negotiated via the CLUE protocol [\[I-D.ietf-clue-protocol\]](#). For an "m" line to be CLUE-controlled, its "mid" value MUST be included in a CLUE group. CLUE-controlled media is controlled by the CLUE protocol as negotiated on the CLUE data channel with an "mid" included in the CLUE group.

"m" lines not specified as under CLUE control follow normal rules for media streams negotiated in SDP as defined in documents such as [\[RFC3264\]](#).

The restrictions on CLUE-controlled media always apply to "m" lines in an SDP offer or answer, even if negotiation of the data channel in SDP failed due to lack of CLUE support by the remote device or for any other reason, or in an offer if the recipient does not include the "mid" of the corresponding "m" line in their CLUE group.

[4.4.](#) SDP semantics for CLUE-controlled media

[4.4.1.](#) Signalling CLUE Encodings

The CLUE Framework [\[I-D.ietf-clue-framework\]](#) defines the concept of "Encodings", which represent the sender's encode ability. Each Encoding the Media Provider wishes to signal is signalled via an "m" line of the appropriate media type, which MUST be marked as sendonly with the "a=sendonly" attribute or as inactive with the "a=inactive" attribute.

The encoder limits of active (eg, "a=sendonly") Encodings can then be expressed using existing SDP syntax. For instance, for H.264 see Table 6 in [\[RFC6184\]](#) for a list of valid parameters for representing encoder sender stream limits.

These Encodings are CLUE-controlled and hence MUST include an "mid" in a CLUE group as defined above.

As well as the normal restrictions defined in [\[RFC3264\]](#) the stream MUST be treated as if the "m" line direction attribute had been set to "a=inactive" until the Media Provider has received a valid CLUE Configure message specifying the Capture to be used for this stream. This means that media packets MUST NOT be sent until configuration is complete, while non-media packets such as STUN and DTLS MUST be sent as normal if negotiated.

Every "m" line representing a CLUE Encoding MUST contain a "label" attribute as defined in [\[RFC4574\]](#). This label is used to identify the Encoding by the sender in CLUE Advertisement messages and by the receiver in CLUE Configure messages. Each label used for a CLUE-controlled "m" line MUST be different from the label on all other "m" lines in the same CLUE group in the SDP message, unless an "m" line represents a dependent stream related to another "m" line (such as a FEC stream), in which case it MUST have the same label value as the "m" line on which it is dependent.

[4.4.1.1](#). Referencing Encodings in the CLUE protocol

CLUE Encodings are defined in SDP, but can be referenced from CLUE protocol messages - this is how the protocol defines which Encodings are part of an Encoding group (in Advertisement messages) and which Encoding with which to encode a specific Capture (in Configure messages). The labels on the CLUE-controlled "m" lines are the references that are used in the CLUE protocol.

Each <encID> (in encodingIDListType) in a CLUE Advertisement message SHOULD represent an Encoding defined in SDP; the specific Encoding referenced is a CLUE-controlled "m" line in the most recent SDP sent by the sender of the Advertisement message with a label value corresponding to the text content of the <encID>.

Similarly, each <encodingID> (in captureEncodingType) in a CLUE Configure message SHOULD represent an Encoding defined in SDP; the specific Encoding referenced is a CLUE-controlled "m" line in the most recent SDP received by the sender of the Configure message with a label value corresponding to the text content of the <encodingID>.

Note that the non-atomic nature of SDP/CLUE protocol interaction may mean that there are temporary periods where an <encID>/<encodingID> in a CLUE message does not reference an SDP "m" line, or where an Encoding represented in SDP is not referenced in a CLUE protocol message. See [Section 5](#) for specifics.

4.4.1.2. Media line directionality

Presently, this specification mandates that CLUE-controlled "m" lines must be unidirectional. This is because setting "m" lines to "a=sendonly" allows the encoder limits to be expressed, whereas in other cases codec attributes express the receive capabilities of a media line.

It is possible that in future versions of this draft or its successor this restriction will be relaxed. If a device does not feel there is a benefit to expressing encode limitations, or if there are no meaningful codec-specific limitations to express (such as with many audio codecs) there are benefits to allowing bidirectional "m" lines. With bidirectional media lines recipients do not always need to create a new offer to add their own "m" lines to express their send capabilities; if they can produce an equal or lesser number of streams to send then they may not need additional "m" lines.

However, at present the need to express encode limitations and the wish to simplify the offer/answer procedure means that for the time being only unidirectional media lines are allowed for CLUE-controlled media. The highly asymmetric nature of CLUE means that the probability of the recipient of the initial offer needing to make their own offer to add additional "m" lines is significantly higher than it is for most other SIP call scenarios, in which there is a tendency for both sides to have similar numbers of potential audio and video streams they can send.

4.4.2. Negotiating receipt of CLUE Capture Encodings in SDP

A receiver who wishes to receive a CLUE stream via a specific Encoding requires an "a=recvonly" "m" line that matches the "a=sendonly" Encoding.

These "m" lines are CLUE-controlled and hence MUST include their "mid" in the CLUE group corresponding to the CLUE group of the Encoding they wish to receive.

4.5. SDP Offer/Answer Procedures

4.5.1. Generating the Initial Offer

A CLUE-capable device sending an initial SDP offer of a SIP session SHOULD include an "m" line for the data channel to convey the CLUE protocol, along with a CLUE group containing the "mid" of the data channel "m" line.

For interoperability with non-CLUE devices a CLUE-capable device sending an initial SDP offer SHOULD NOT include any "m" line for CLUE-controlled media beyond the "m" line for the CLUE data channel, and SHOULD include at least one non-CLUE-controlled media "m" line.

If the device has evidence that the receiver is also CLUE-capable, for instance due to receiving an initial INVITE with no SDP but including a "sip.clue" media feature tag, the above recommendation is waived, and the initial offer MAY contain "m" lines for CLUE-controlled media.

With the same interoperability recommendations as for Encodings, the sender of the initial SDP offer MAY also include "a=recvonly" media lines to preallocate "m" lines to receive media. Alternatively, it MAY wait until CLUE protocol negotiation has completed before including these lines in a new offer/answer exchange - see [Section 5](#) for recommendations.

[4.5.2.](#) Generating the Answer

[4.5.2.1.](#) Negotiating use of CLUE and the CLUE data channel

If the recipient is CLUE-capable and the initial offer contains both an "m" line for a data channel and a CLUE group containing the "mid" for that "m" line, they SHOULD negotiate data channel support for an "m" line, and include the "mid" of that "m" line in a corresponding CLUE group.

A CLUE-capable recipient that receives an "m" line for a data channel but no corresponding CLUE group containing the "mid" of that "m" line MAY still include a corresponding data channel "m" line if there are any other non-CLUE protocols it can convey over that channel, but MUST NOT negotiate use of the CLUE protocol on this channel.

[4.5.2.2.](#) Negotiating CLUE-controlled media

If the initial offer contained "a=recvonly" CLUE-controlled media lines the recipient SHOULD include corresponding "a=sendonly" CLUE-controlled media lines, up to the maximum number of Encodings it wishes to advertise. As CLUE-controlled media, the "mid" of these "m" lines must be included in the corresponding CLUE group.

If the initial offer contained "a=sendonly" CLUE-controlled media lines the recipient MAY include corresponding "a=recvonly" CLUE-controlled media lines, up to the maximum number of Capture Encodings it wishes to receive. Alternatively, it MAY wait until CLUE protocol negotiation has completed before including these lines in a new offer/answer exchange - see [Section 5](#) for recommendations.

4.5.2.3. Negotiating non-CLUE controlled media

A CLUE-controlled device implementation may prefer to render initial, single-stream audio and/or video for the user as rapidly as possible, transitioning to CLUE-controlled media once that has been negotiated. Alternatively, an implementation may wish to suppress initial media, only providing media once the final, CLUE-controlled streams have been negotiated.

The receiver of the initial offer, if making the call CLUE-enabled with their SDP answer, can make their preference clear by their action in accepting or rejecting non-CLUE-controlled media lines. Rejecting these "m" lines will ensure that no non-CLUE-controlled media flows before the CLUE-controlled media is negotiated. In contrast, accepting one or more non-CLUE-controlled "m" lines in this initial answer will enable initial media to flow.

If the answerer chooses to send initial non-CLUE-controlled media in a CLUE-enabled call, [Section 4.5.4.1](#) addresses the need to disable it once CLUE-controlled media is fully negotiated.

4.5.3. Processing the initial Offer/Answer negotiation

In the event that both offer and answer include a data channel "m" line with a mid value included in corresponding CLUE groups CLUE has been successfully negotiated and the call is now CLUE-enabled, otherwise the call is not CLUE-enabled.

4.5.3.1. Successful CLUE negotiation

In the event of successful CLUE-enablement of the call, devices MUST now begin negotiation of the CLUE channel, see [\[I-D.ietf-clue-datachannel\]](#) for negotiation details. If negotiation is successful, sending of CLUE protocol [\[I-D.ietf-clue-protocol\]](#) messages can begin.

A CLUE-capable device MAY choose not to send media on the non-CLUE-controlled channels during the period in which control of the CLUE-controlled media lines is being negotiated. However, a CLUE-capable device MUST still be prepared to receive media on non-CLUE-controlled media lines that have been successfully negotiated as defined in [\[RFC3264\]](#).

If either side of the call wishes to add additional CLUE-controlled "m" line to send or receive CLUE-controlled media they MAY now send a SIP request with a new SDP offer. Note that if BUNDLE has been successfully negotiated and a Bundle Address Synchronization offer is

required, the device to receive that offer SHOULD NOT generate a new SDP offer until it has received that BAS offer.

4.5.3.2. CLUE negotiation failure

In the event that the negotiation of CLUE fails and the call is not CLUE-enabled in the initial offer/answer then CLUE is not in use in the call, and the CLUE-capable devices MUST either revert to non-CLUE behaviour or terminate the call.

4.5.4. Modifying the session

4.5.4.1. Adding and removing CLUE-controlled media

Subsequent offer/answer exchanges MAY add additional "m" lines for CLUE-controlled media; in most cases at least one additional exchange will be required before both sides have added all the Encodings and ability to receive Encodings that they desire. Devices MAY delay adding "a=recvonly" CLUE-controlled m-lines until after CLUE protocol negotiation completes - see [Section 5](#) for recommendations.

Subsequent offer/answer exchanges MAY also deactivate "m" lines for CLUE-controlled media.

Once CLUE media has been successfully negotiated devices SHOULD ensure that non-CLUE-controlled media is deactivated in cases where it corresponds to the media type of CLUE-controlled media that has been successfully negotiated. This deactivate may require an additional SDP exchange, or may be incorporated into one that is part of the CLUE negotiation.

4.5.4.2. Enabling CLUE mid-call

A CLUE-capable device that receives an initial SDP offer from a non-CLUE device SHOULD include a new data channel "m" line and corresponding CLUE group in any subsequent offers it sends, to indicate that it is CLUE-capable.

If, in an ongoing non-CLUE call, an SDP offer/answer exchange completes with both sides having included a data channel "m" line in their SDP and with the "mid" for that channel in corresponding CLUE groups then the call is now CLUE-enabled; negotiation of the data channel and subsequently the CLUE protocol begin.

4.5.4.3. Disabling CLUE mid-call

If, in an ongoing CLUE-enabled call, an SDP offer-answer negotiation completes in a fashion in which either the CLUE data channel was not successfully negotiated or one side did not include the data channel in a matching CLUE group then CLUE for this channel is disabled. In the event that this occurs, CLUE is no longer enabled and sending of all CLUE-controlled media associated with the corresponding CLUE group MUST stop. If the data channel is still present but not included in the CLUE group semantic CLUE protocol messages MUST no longer be sent.

Note that this is distinct to cases where the CLUE data channel fails or an error occurs on the CLUE protocol; see [[I-D.ietf-clue-protocol](#)] for details of media and state preservation in this circumstance.

5. Interaction of CLUE protocol and SDP negotiations

Information about media streams in CLUE is split between two message types: SDP, which defines media addresses and limits, and the CLUE channel, which defines properties of Capture Devices available, scene information and additional constraints. As a result certain operations, such as advertising support for a new transmissible Capture with associated stream, cannot be performed atomically, as they require changes to both SDP and CLUE messaging.

This section defines how the negotiation of the two protocols interact, provides some recommendations on dealing with intermediary stages in non-atomic operations, and mandates additional constraints on when CLUE-configured media can be sent.

5.1. Independence of SDP and CLUE negotiation

To avoid the need to implement interlocking state machines with the potential to reach invalid states if messages were to be lost, or be rewritten en-route by middle boxes, the state machines in SDP and CLUE operate independently. The state of the CLUE channel does not restrict when an implementation may send a new SDP offer or answer, and likewise the implementation's ability to send a new CLUE Advertisement or Configure message is not restricted by the results of or the state of the most recent SDP negotiation (unless the SDP negotiation has removed the CLUE channel).

The primary implication of this is that a device may receive an SDP with a CLUE Encoding it does not yet have capture information for, or receive a CLUE Configure message specifying a Capture Encoding for which the far end has not negotiated a media stream in SDP.

CLUE messages contain an <encID> (in encodingIDListType) or <encodingID> (in captureEncodingType), which is used to identify a specific encoding or captureEncoding in SDP; see [\[I-D.ietf-clue-data-model-schema\]](#) for specifics. The non-atomic nature of CLUE negotiation means that a sender may wish to send a new Advertisement before the corresponding SDP message. As such the sender of the CLUE message MAY include an <encID> which does not currently match a CLUE-controlled "m" line label in SDP; A CLUE-capable implementation MUST NOT reject a CLUE protocol messages solely because it contains <encID> elements that do not match an id in SDP.

The current state of the CLUE participant or Media Provider/Consumer state machines do not affect compliance with any of the normative language of [\[RFC3264\]](#). That is, they MUST NOT delay an ongoing SDP exchange as part of a SIP server or client transaction; an implementation MUST NOT delay an SDP exchange while waiting for CLUE negotiation to complete or for a Configure message to arrive.

Similarly, a device in a CLUE-enabled call MUST NOT delay any mandatory state transitions in the CLUE Participant or Media Provider/Consumer state machines due to the presence or absence of an ongoing SDP exchange.

A device with the CLUE Participant state machine in the ACTIVE state MAY choose not to move from ESTABLISHED to ADV (Media Provider state machine) or from ESTABLISHED to WAIT FOR CONF RESPONSE (Media Consumer state machine) based on the SDP state. See [\[I-D.ietf-clue-protocol\]](#) for CLUE state machine specifics. Similarly, a device MAY choose to delay initiating a new SDP exchange based on the state of their CLUE state machines.

5.2. Constraints on sending media

While SDP and CLUE message states do not impose constraints on each other, both impose constraints on the sending of media - CLUE-controlled media MUST NOT be sent unless it has been negotiated in both CLUE and SDP: an implementation MUST NOT send a specific CLUE Capture Encoding unless its most recent SDP exchange contains an active media channel for that Encoding AND the far end has sent a CLUE Configure message specifying a valid Capture for that Encoding.

5.3. Recommendations for operating with non-atomic operations

CLUE-capable devices MUST be able to handle states in which CLUE messages make reference to EncodingIDs that do not match the most recently received SDP, irrespective of the order in which SDP and CLUE messages are received. While these mis-matches will usually be

transitory a device **MUST** be able to cope with such mismatches remaining indefinitely. However, this document makes some recommendations on message ordering for these non-atomic transitions.

CLUE-capable devices **SHOULD** ensure that any inconsistencies between SDP and CLUE signalling are temporary by sending updated SDP or CLUE messages as soon as the relevant state machines and other constraints permit.

Generally, implementations that receive messages for which they have incomplete information **SHOULD** wait until they have the corresponding information they lack before sending messages to make changes related to that information. For instance, an implementation that receives a new SDP offer with three new "a=sendonly" CLUE "m" lines that has not received the corresponding CLUE Advertisement providing the capture information for those streams **SHOULD NOT** include corresponding "a=recvonly" lines in its answer, but instead should make a new SDP offer when and if a new Advertisement arrives with Captures relevant to those Encodings.

Because of the constraints of offer/answer and because new SDP negotiations are generally more 'costly' than sending a new CLUE message, implementations needing to make changes to both channels **SHOULD** prioritize sending the updated CLUE message over sending the new SDP message. The aim is for the recipient to receive the CLUE changes before the SDP changes, allowing the recipient to send their SDP answers without incomplete information, reducing the number of new SDP offers required.

6. Interaction of CLUE protocol and RTP/RTCP CaptureID

[I-D.ietf-clue-framework] allows for Multiple Content Captures (MCCs): Captures which contain multiple source Captures, whether composited into a single stream or switched based on some metric.

The Captures that constitute these MCCs may or may not be defined in the Advertisement message. If they are defined and the MCC is providing them in a switched format the recipient may wish to determine which originating source Capture is currently being provided, so that they can apply geometric corrections based on that Capture's geometry, or take some other action based on the original Capture information.

To do this, [I-D.ietf-clue-rtp-mapping] allows for the CaptureID of the originating Capture to be conveyed via RTP or RTCP. A Media Provider sending switched media from an MCC with defined originating sources **MUST** send the CaptureID in both RTP and RTCP, as described in the mapping document.

6.1. CaptureID reception during MCC redefinition

Because the RTP/RTCP CaptureID is delivered via a different channel to the Advertisement in which the contents of the MCC are defined there is an intrinsic race condition in cases in which the contents of an MCC are redefined.

When a Media Provider redefines an MCC which involves CaptureIDs, the reception of the relevant CaptureIDs by the recipient will either lead or lag reception and processing of the new Advertisement by the recipient. As such, a media recipient **MUST** not be disrupted by any of the following in any CLUE- controlled media stream it is receiving, whether that stream is for a static Capture or for an MCC (as any static Capture may be redefined to an MCC in a later Advertisement):

- o Receiving RTP or RTCP containing a CaptureID when the most recently processed Advertisement means that none are expected.
- o Receiving RTP or RTCP without CaptureIDs when the most recently processed Advertisement means that media CaptureIDs are expected.
- o Receiving a CaptureID in RTP or RTCP for a Capture defined in the most recently processed Advertisement, but which the same Advertisement does not include in the MCC.
- o Receiving a CaptureID in RTP or RTCP for a Capture not defined in the most recently processed Advertisement.

7. Multiplexing of CLUE-controlled media using BUNDLE

7.1. Overview

A CLUE call may involve sending and/or receiving significant numbers of media streams. Conventionally, media streams are sent and received on unique ports. However, each separate port used for this purpose may impose costs that a device wishes to avoid, such as the need to open that port on firewalls and NATs, the need to collect ICE candidates [[RFC5245](#)], etc.

The BUNDLE [[I-D.ietf-mmusic-sdp-bundle-negotiation](#)] extension can be used to negotiate the multiplexing of multiple media lines onto a single 5-tuple for sending and receiving media, allowing devices in calls to another BUNDLE-supporting device to potentially avoid some of the above costs.

While CLUE-capable devices MAY support the BUNDLE extension for this purpose supporting the extension is not mandatory for a device to be CLUE-compliant.

7.2. Usage of BUNDLE with CLUE

This specification imposes no additional requirements or restrictions on the usage of BUNDLE when used with CLUE. There is no restriction on combining CLUE-controlled media lines and non-CLUE-controlled media lines in the same BUNDLE group or in multiple such groups. However, there are several steps an implementation may wish to ameliorate the cost and time requirements of extra SDP offer/answer exchanges between CLUE and BUNDLE.

7.2.1. Generating the Initial Offer

BUNDLE mandates that the initial SDP offer MUST use a unique address for each m-line with a non-zero port. Because CLUE implementations generallly will not include CLUE-controlled media lines with the exception of the data channel CLUE devices that support large numbers of streams can avoid ever having to open large numbers of ports if they successfully negotiate BUNDLE.

7.2.2. Bundle Address Synchronization

When using BUNDLE the initial offerer may be mandated to send a Bundle Address Synchronisation offer. If the initial offerer also followed the recommendation of not including CLUE-controlled media lines in their offer, they MAY choose to include them in this subsequent offer. In this circumstance the BUNDLE specification recommends that the offerer does not "modify SDP parameters that could get the answerer to reject the BAS offer". Including new CLUE-controlled media lines using codecs and other attributes used in existing media lines should not increase the chance of the answerer rejecting the BAS offer; implementations should consider carefully before including new codecs or other new SDP attributes in these CLUE-controlled media lines.

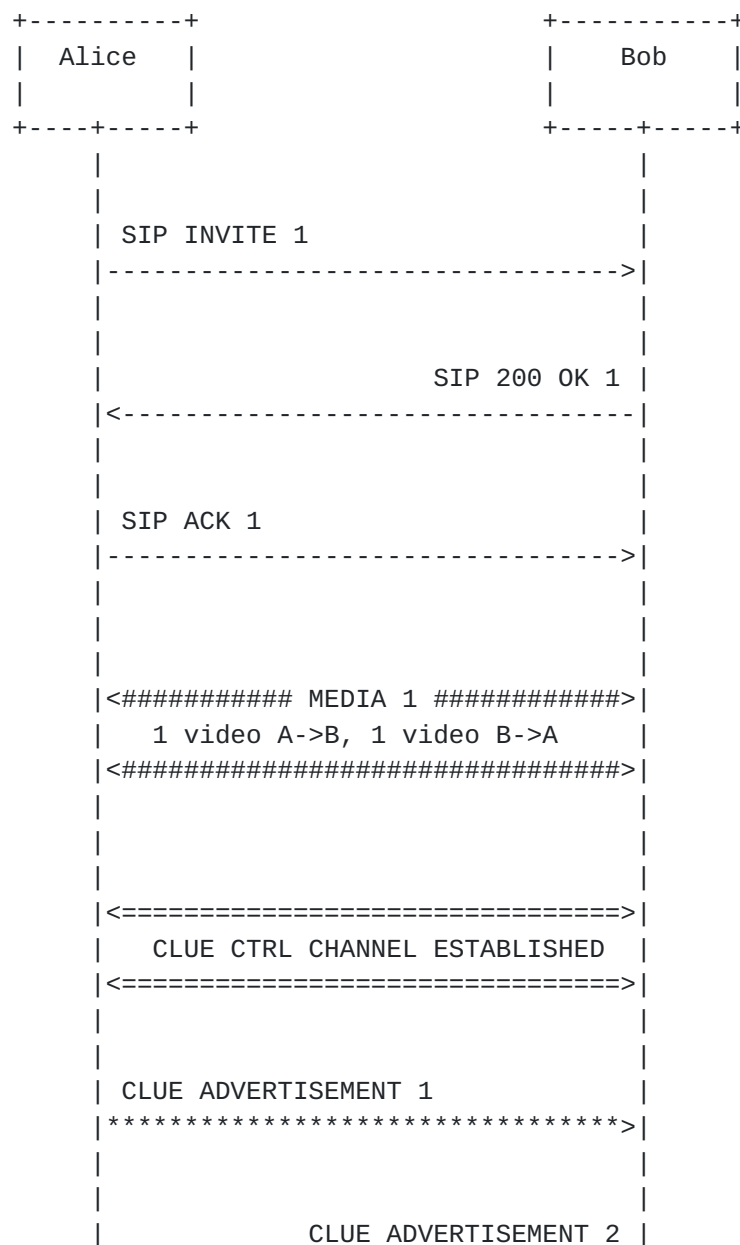
7.2.3. Multiplexing of the data channel and RTP media

BUNDLE-supporting CLUE-capable devices MAY include the data channel in the same BUNDLE group as RTP media. In this case the device MUST be able to demultiplex the various transports - see [section 7.2](#) of the BUNDLE draft [[I-D.ietf-mmusic-sdp-bundle-negotiation](#)]. If the BUNDLE group includes other protocols than the data channel transported via DTLS the device MUST also be able to differentiate the various protocols.

8. Example: A call between two CLUE-capable Endpoints

This example illustrates a call between two CLUE-capable Endpoints. Alice, initiating the call, is a system with three cameras and three screens. Bob, receiving the call, is a system with two cameras and two screens. A call-flow diagram is presented, followed by an summary of each message.

To manage the size of this section SDP snippet only illustrate video 'm' lines. ACKs are not discussed. Note that BUNDLE is not in use.




```

|<*****|
|
|
|SIP INVITE 2 (+3 sendonly)|
|----->|
|
|
|CLUE CONFIGURE 1|
|<*****|
|
|
|CLUE RESPONSE 1|
|*****>|
|
|
|SIP 200 OK 2 (+2 recvonly)|
|<-----|
|
|SIP ACK 2|
|----->|
|
|
|<##### MEDIA 2 #####>|
|2 video A->B, 1 video B->A|
|<#####>|
|
|
|SIP INVITE 3 (+2 sendonly)|
|<-----|
|
|
|CLUE CONFIGURE 2|
|*****>|
|
|
|CLUE RESPONSE 2|
|<*****|
|
|
|SIP 200 OK 3 (+2 recvonly)|
|----->|
|
|
|SIP ACK 3|
|<-----|

```



```

|
|
|
|<##### MEDIA 3 #####>|
|  2 video A->B, 2 video B->A  |
|<#####>|
|
|
|
|
v                               v

```

In INVITE 1, Alice sends Bob a SIP INVITE including in the SDP body the basilar audio and video capabilities and the information needed for opening a control channel to be used for CLUE protocol messages exchange, according to what is envisioned in the COMEDIA approach for DTLS/SCTP channel [[I-D.ietf-mmusic-sctp-sdp](#)]. A snippet of the SDP showing the grouping attribute and the video m-line are shown below. Alice has included a "CLUE" group, and included the mid corresponding to a data channel in the group (3). Note that Alice has chosen not to include any CLUE-controlled media in the initial offer - the mid value of the video line is not included in the "CLUE" group.

```

...
a=group:CLUE 3
...
m=video 6002 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016;max-mps=108000;max-fs=3600
a=sendrecv
a=mid:2
...
m=application 6100 UDP/DTLS/SCTP webrtc-datachannel
a=sctp-port: 5000
a=dcmap:2 subprotocol="CLUE";ordered=true
a=mid:3

```

Bob responds with a similar SDP (200 OK 1), which also has a "CLUE" group including the mid value of a data channel; due to their similarity no SDP snippet is shown here. Bob wishes to receive initial media, and so includes corresponding non-CLUE-controlled audio and video lines. Alice and Bob are each now able to send a single audio and video stream. This is illustrated as MEDIA 1.

With the successful initial O/A Alice and Bob are also free to negotiate the CLUE channel. Once this is successfully established

CLUE negotiation can begin. This is illustrated as CLUE CTRL CHANNEL ESTABLISHED.

Alice now sends her CLUE Advertisement (ADVERTISEMENT 1). She advertises three static Captures representing her three cameras. She also includes switched Captures suitable for two- and one-screen systems. All of these Captures are in a single Capture Scene, with suitable Capture Scene entries to tell Bob that he should either subscribe to the three static Captures, the two switched Captures or the one switched Capture. Alice has no simultaneity constraints, so includes all six Captures in one simultaneous set. Finally, Alice includes an Encoding Group with three Encoding IDs: "enc1", "enc2" and "enc3". These Encoding IDs aren't currently valid, but will match the next SDP offer she sends.

Bob received ADVERTISEMENT 1 but does not yet send a Configure message, because he has not yet received Alice's Encoding information, so as yet he does not know if she will have sufficient resources to send him the two streams he ideally wants at a quality he is happy with.

Bob also sends his CLUE Advertisement (ADVERTISEMENT 2). He advertises two static Captures representing his cameras. He also includes a single composed Capture for single-screen systems, in which he will composite the two camera views into a single video stream. All three Captures are in a single Capture Scene, with suitable Capture Scene entries to tell Alice that she should either subscribe to the two static Captures, or the single composed Capture. Bob also has no simultaneity constraints, so includes all three Captures in one simultaneous set. Bob also includes a single Encoding Group with two Encoding IDs: "foo" and "bar".

Similarly, Alice receives ADVERTISEMENT 2 but does not yet send a Configure message, because she has not yet received Bob's Encoding information.

Alice now sends INVITE 2. She maintains the sendrecv audio, video and CLUE m-lines, and she adds three new sendonly m-lines to represent the three CLUE-controlled Encodings she can send. Each of these m-lines has a label corresponding to one of the Encoding IDs from ADVERTISEMENT 1. Each also has its mid added to the grouping attribute to show they are controlled by the CLUE channel. A snippet of the SDP showing the grouping attribute, data channel and the video "m" lines are shown below:


```
...
a=group:CLUE 3 4 5 6
...
m=video 6002 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016;max-mps=108000;max-fs=3600
a=sendrecv
a=mid:2
...
m=application 6100 UDP/DTLS/SCTP webrtc-datachannel
a=sctp-port: 5000
a=dcmap:2 subprotocol="CLUE";ordered=true
a=mid:3
...
m=video 6004 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016
a=sendonly
a=mid:4
a=label:enc1
m=video 6006 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016
a=sendonly
a=mid:5
a=label:enc2
m=video 6008 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016
a=sendonly
a=mid:6
a=label:enc3
```

Bob now has all the information he needs to decide which streams to configure. As such he now sends CONFIGURE 1. This requests the pair of switched Captures that represent Alice's scene, and he configures them with encoder ids "enc1" and "enc2". This also serves as an ack for Alice's ADVERTISEMENT 1.

Alice receives Bob's message CONFIGURE 1 and sends RESPONSE 1 to ack its reception. She does not yet send the Capture Encodings specified, because at this stage Bob hasn't negotiated the ability to receive these streams in SDP.

Bob now sends his SDP answer as part of 200 OK 2. Alongside his original audio, video and CLUE m-lines he includes two active recvonly m-lines and a zeroed m-line for the third. He adds their

mid values to the grouping attribute to show they are controlled by the CLUE channel. A snippet of the SDP showing the grouping attribute and the video m-lines are shown below (mid 100 represents the CLUE channel, not shown):

```
...
a=group:CLUE 11 12 100
...
m=video 58722 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=sendrecv
a=mid:10
...
m=video 58724 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=recvonly
a=mid:11
m=video 58726 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=recvonly
a=mid:12
m=video 0 RTP/AVP 96
```

On receiving 200 OK 2 from Bob Alice is now able to send the two streams of video Bob requested - this is illustrated as MEDIA 2.

The constraints of offer/answer meant that Bob could not include his encoder information as new m-lines in 200 OK 2. As such Bob now sends INVITE 3 to generate a new offer. Along with all the streams from 200 OK 2 Bob also includes two new sendonly streams. Each stream has a label corresponding to the Encoding IDs in his ADVERTISEMENT 2 message. He also adds their mid values to the grouping attribute to show they are controlled by the CLUE channel. A snippet of the SDP showing the grouping attribute and the video m-lines are shown below (mid 100 represents the CLUE channel, not shown):


```
...
a=group:CLUE 11 12 13 14 100
...
m=video 58722 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=sendrecv
a=mid:10
...
m=video 58724 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=recvonly
a=mid:11
m=video 58726 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=recvonly
a=mid:12
m=video 0 RTP/AVP 96
m=video 58728 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016
a=sendonly
a=label:foo
a=mid:13
m=video 58730 RTP/AVP 96
a=rtpmap:96 H264/900000
a=fmtp:96 profile-level-id=42e016
a=sendonly
a=label:bar
a=mid:14
```

Having received this Alice now has all the information she needs to send CONFIGURE 2. She requests the two static Captures from Bob, to be sent on Encodings "foo" and "bar".

Bob receives Alice's message CONFIGURE 2 and sends RESPONSE 2 to ack its reception. Bob does not yet send the Capture Encodings specified, because Alice hasn't yet negotiated the ability to receive these streams in SDP.

Alice now sends 200 OK 3, matching two recvonly m-lines to Bob's new sendonly lines. She includes their mid values in the grouping attribute to show they are controlled by the CLUE channel. Alice also now deactivates the initial non-CLUE-controlled media, as bidirectional CLUE-controlled media is now available. A snippet of

the SDP showing the grouping attribute and the video m-lines are shown below (mid 3 represents the data channel, not shown):


```
...
a=group:CLUE 3 4 5 7 8
...
m=video 0 RTP/AVP 96
a=mid:2
...
m=video 6004 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016
a=sendonly
a=mid:4
a=label:enc1
m=video 6006 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016
a=sendonly
a=mid:5
a=label:enc2
m=video 0 RTP/AVP 96
m=video 6010 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=recvonly
a=mid:7
m=video 6012 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016;max-mbps=108000;max-fs=3600
a=recvonly
a=mid:8
```

Finally, on receiving 200 OK 3 Bob is now able to send the two streams of video Alice requested - this is illustrated as MEDIA 3.

Both sides of the call are now sending multiple video streams with their sources defined via CLUE negotiation. As the call progresses either side can send new Advertisement or Configure message or new SDP negotiation to add, remove or change what they have available or want to receive.

9. Example: A call between a CLUE-capable and non-CLUE Endpoint

In this brief example Alice is a CLUE-capable Endpoint making a call to Bob, who is not CLUE-capable ((i.e. is not able to use the CLUE protocol)).



In INVITE 1, Alice sends Bob a SIP INVITE including in the SDP body the basilar audio and video capabilities and the information needed for opening a control channel to be used for CLUE protocol messages exchange, according to what is envisioned in the COMEDIA approach for a DTLS/SCTP channel [[I-D.ietf-mmusic-sctp-sdp](#)]. A snippet of the SDP showing the grouping attribute, data channel and the video m-line are shown below:


```
...
a=group:CLUE 3
...
m=video 6002 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42e016;max-mps=108000;max-fs=3600
a=sendrecv
a=mid:2
...
m=application 6100 UDP/DTLS/SCTP webrtc-datachannel
a=sctp-port: 5000
a=dcmap:2 subprotocol="CLUE";ordered=true
a=mid:3
```

Bob is not CLUE-capable, and hence does not recognize the "CLUE" semantic for grouping attribute, nor does he support the data channel. He responds with an answer with audio and video, but with the data channel zeroed.

From the lack of the data channel and grouping framework Alice understands that Bob does not support CLUE, or does not wish to use it. Both sides are now able to send a single audio and video stream to each other. Alice at this point begins to send her fallback video: in this case likely a switched view from whichever camera shows the current loudest participant on her side.

10. Acknowledgements

The team focusing on this draft consists of: Roni Even, Rob Hansen, Christer Holmberg, Paul Kyzivat, Simon Pietro-Romano, Roberta Presta.

Christian Groves and Jonathon Lennox have contributed detailed comments and suggestions.

11. IANA Considerations

11.1. New SDP Grouping Framework Attribute

This document registers the following semantics with IANA in the "Semantics for the "group" SDP Attribute" subregistry (under the "Session Description Protocol (SDP) Parameters" registry: Semantics Token Reference -----
----- CLUE controlled m-line CLUE [this draft]

11.2. New SIP Media Feature Tag

This specification registers a new media feature tag in the SIP [[RFC3264](#)] tree per the procedures defined in [[RFC2506](#)] and [[RFC3840](#)]. Media feature tag name: sip.clue ASN.1 Identifier: 1.3.6.1.8.4.26 Summary of the media feature indicated by this tag: This feature tag indicates that the device supports CLUE controlled media. Values appropriate for use with this feature tag: Boolean. The feature tag is intended primarily for use in the following applications, protocols, services, or negotiation mechanisms: This feature tag is most useful in a communications application for describing the capabilities of a device which uses multiple media streams.

12. Security Considerations

CLUE makes use of a number of protocols and mechanism, either defined by CLUE or long-standing. The security considerations section of the CLUE Framework [[I-D.ietf-clue-framework](#)] addresses the need to secure these mechanisms by following the recommendations of the individual protocols.

Beyond the need to secure the constituent protocols, the use of CLUE does impose additional security concerns. One area of increased risk involves the potential for a malicious party to subvert a CLUE-capable device to attack a third party by driving large volumes of media (particularly video) traffic at them by establishing a connection to the CLUE-capable device and directing the media to the victim. While this is a risk for all media devices, a CLUE-capable device may allow the attacker to configure multiple media streams to be sent, significantly increasing the volume of traffic directed at the victim.

This attack can be prevented by ensuring that the media recipient intends to receive the media packets. As such all CLUE-capable devices MUST support key negotiation and receiver intent assurance via DTLS [[RFC5763](#)] on CLUE-controlled RTP "m" lines. All CLUE-controlled RTP "m" lines must be secured and implemented using mechanisms such as SRTP [[RFC3711](#)]; no specific security mechanisms are made mandatory to use due to the issues addressed in [[RFC7202](#)]. Due to the requirements of backwards compatibility, there is not a mandatory requirement for non-CLUE-controlled "m" lines.

CLUE also defines a new media feature tag that indicates CLUE support. This tag may be present even in non-CLUE calls, which increases the metadata available about the sending device, which can help an attacker differentiate between multiple devices and help them identify otherwise anonymised users via the fingerprint of features

their device supports. To prevent this, SIP signalling SHOULD always be encrypted using TLS [[RFC5630](#)].

13. Change History

Revision by Rob Hansen

- o State machine interactions updated to match versions in -04 of protocol doc.
- o Section on encoding updated to specify both encID and encodingID from data model doc.
- o Removed the limitations on describing H264 encoding limits using SDP syntax as an open issue.
- o Previous draft had SRTP and DTLS mandatory to implement and to use on CLUE- controlled m lines. Current version has DTLS mandatory to implement, and 'security' mandatory to use but does not define what that security is.
- o Terminology reference to framework doc reinforced. All terminology that duplicates framework removed. All text updated with capitalisation that matches framework document's terminology.
- o SDP example syntax updated to match that of ietf-clue-datachannel and hence ietf-mmusic-data-channel-sdpneg.

Revision by Rob Hansen

- o SRTP/DTLS made mandatory for CLUE-controlled media lines.
- o IANA consideration section added (text as proposed by Christian Groves).
- o Includes provision for dependent streams on separate "m" lines having the same encID as their parent "m" line.
- o References to putting CLUE-controlled media and data channels in more than one CLUE group removed, since the document no longer supports using more than one CLUE group.
- o Section on CLUE controlled media restrictions still applying even if the call does not end up being CLUE enabled being rewritten to hopefully be clearer.
- o Other minor syntax improvements.

Revision by Rob Hansen

- o Updated DTLS/SCTP channel syntax in examples to fix errors and match latest format defined in [draft-ietf-mmusic-sctp-sdp-07](#).
- o Clarified the behaviour if an SDP offer includes a CLUE-controlled "m" line and the answer accepts that "m" line but without CLUE control of that line.
- o Added a new section on the sending and receiving of CaptureIDs in RTP and RTCP. Includes a section on the necessity of the receiver coping with unexpected CaptureIDs (or the lack thereof) due to MCCs being redefined in new Advertisement messages.
- o Added reminder on IANA section on registering grouping semantic and media feature tag, removed the less formal sections that did the same job.
- o Fixed and clarified issues raised by Christian's document review.
- o Added a number of security considerations.

Revision by Rob Hansen

- o Clarified text on not rejecting messages because they contain unknown encIDs.
- o Removed normative language in section on accepting/rejecting non-CLUE-controlled media in the initial answer.
- o Example SDP updated to include the data channel "m" lines.
- o Example call flow updated to show disablement of non-CLUE-controlled media once CLUE-controlled media is flowing.

-02: Revision by Rob Hansen

- * Added section on not accepting non-CLUE-controlled "m" lines in the initial answer when CLUE is to be negotiated.
- * Removed previous language attempting to describe media restrictions for CLUE-controlled "m" lines that had not been configured, and replaced it with much more accurate 'treat as "a=inactive" was set'.
- * Made label element mandatory for CLUE-controlled media (was previously "SHOULD include", but there didn't seem a good reason for this - anyone wishing to include the "m" line but

not immediately use it in CLUE can simply leave it out of the <encodingIDList>.)

- * Added a section on the specifics of relating encodings in SDP to <encID> elements in the CLUE protocol, including the fact that both Advertisement and Configure messages reference the *encoding* (eg, in the Configure case the sender of the Configure message includes the labels of the recipient's "m" lines as their <encID> contents).
- * Minor revisions to the section on complying with normative SDP/CLUEstate machine language to clarify that these were not new normative language, merely that existing normative language still applies.
- * Removed appendices which previously contained information to be transferred to the protocol and data channel drafts. Removed other text that discussed alternatives to the current approach.
- * Cleaned up some 'todo' text.

-01: Revision by Rob Hansen

- * Revised terminology - removed the term 'CLUE-enabled' device as insufficiently distinct from 'CLUE-capable' and instead added a term for 'CLUE-enabled' calls.
- * Removed text forbidding RTCP and instead added text that ICE/DTLS negotiation for CLUE controlled media must be done as normal irrespective of CLUE negotiation.
- * Changed 'sip.telepresence' to 'sip.clue' and 'TELEPRESENCE' grouping semantic back to CLUE.
- * Made it mandatory to have exactly one mid corresponding to a data channel in a CLUE group
- * Forbade having multiple CLUE groups unless a specification for doing so is published.
- * Refactored SDP-related text; previously the encoding information had been in the "initial offer" section despite the fact that we recommend that the initial offer doesn't actually include any encodings. I moved the specifications of encodings and how they're received to an earlier, separate section.

- * Added text on how the state machines in CLUE and SDP are allowed to affect one another, and further recommendations on how a device should handle the sending of CLUE and SDP changes.

-00: Revision by Rob Hansen

- * Submitted as -00 working group document

[draft-kyzivat-08](#): Revisions by Rob Hansen

- * Added media feature tag for CLUE support ('sip.telepresence')
- * Changed grouping semantic from 'CLUE' to 'TELEPRESENCE'
- * Restructured document to be more centred on the grouping semantic and its use with O/A
- * Lots of additional text on usage of the grouping semantic
- * Stricter definition of CLUE-controlled m lines and how they work
- * Some additional text on defining what happens when CLUE supports is added or removed
- * Added details on when to not send RTCP for CLUE-controlled "m" lines.
- * Added a section on using BUNDLE with CLUE
- * Updated data channel references to point at new WG document rather than individual draft

[draft-kyzivat-07](#): Revisions by Rob Hansen

- * Removed the text providing arguments for encoding limits being in SDP and encoding groups in the CLUE protocol in favor of the specifics of how to negotiate encodings in SDP
- * Added normative language on the setting up of a CLUE call, and added sections on mid-call changes to the CLUE status.
- * Added references to [[I-D.ietf-clue-datachannel](#)] where appropriate.
- * Added some terminology for various types of CLUE and non-CLUE states of operation.

- * Moved language related to topics that should be in [[I-D.ietf-clue-datachannel](#)] and [[I-D.ietf-clue-protocol](#)], but that has not yet been resolved in those documents, into an appendix.

[draft-kyzivat-06](#): Revisions by Rob Hansen

- * Removed CLUE message XML schema and details that are now in [draft-presta-clue-protocol](#)
- * Encoding limits in SDP section updated to note that this has been investigated and discussed and is the current working assumption of the WG, though consensus has not been fully achieved.
- * A section has also been added on the current mandation of unidirectional "m" lines.
- * Updated CLUE messaging in example call flow to match [draft-presta-clue-protocol-03](#)

[draft-kyzivat-05](#): Revisions by pkyzivat:

- * Specified versioning model and mechanism.
- * Added explicit response to all messages.
- * Rearranged text to work with the above changes. (Which rendered diff almost useless.)

[draft-kyzivat-04](#): Revisions by Rob Hansen: ???

[draft-kyzivat-03](#): Revisions by pkyzivat:

- * Added a syntax section with an XML schema for CLUE messages. This is a strawhorse, and is very incomplete, but it establishes a template for doing this based on elements defined in the data model. (Thanks to Roberta for help with this!)
- * Did some rewording to fit the syntax section in and reference it.
- * Did some relatively minor restructuring of the document to make it flow better in a logical way.

[draft-kyzivat-02](#): A bunch of revisions by pkyzivat:

- * Moved roberta's call flows to a more appropriate place in the document.
- * New section on versioning.
- * New section on NAK.
- * A couple of possible alternatives for message acknowledgment.
- * Some discussion of when/how to signal changes in provider state.
- * Some discussion about the handling of transport errors.
- * Added a change history section.

These were developed by Lennard Xiao, Christian Groves and Paul, so added Lennard and Christian as authors.

[draft-kyzivat-01](#): Updated by roberta to include some sample call flows.

[draft-kyzivat-00](#): Initial version by pkyzivat. Established general outline for the document, and specified a few things thought to represent wg consensus.

14. References

14.1. Normative References

[I-D.ietf-clue-framework]

Duckworth, M., Pepperell, A., and S. Wenger, "Framework for Telepresence Multi-Streams", [draft-ietf-clue-framework-22](#) (work in progress), April 2015.

[I-D.ietf-clue-data-model-schema]

Presta, R. and S. Romano, "An XML Schema for the CLUE data model", [draft-ietf-clue-data-model-schema-10](#) (work in progress), June 2015.

[I-D.ietf-clue-protocol]

Presta, R. and S. Romano, "CLUE protocol", [draft-ietf-clue-protocol-04](#) (work in progress), April 2015.

[I-D.ietf-clue-datachannel]

Holmberg, C., "CLUE Protocol data channel", [draft-ietf-clue-datachannel-09](#) (work in progress), March 2015.

[I-D.ietf-clue-rtp-mapping]

Even, R. and J. Lennox, "Mapping RTP streams to CLUE media captures", [draft-ietf-clue-rtp-mapping-04](#) (work in progress), March 2015.

[I-D.ietf-mmusic-sctp-sdp]

Holmberg, C., Loreto, S., and G. Camarillo, "Stream Control Transmission Protocol (SCTP)-Based Media Transport in the Session Description Protocol (SDP)", [draft-ietf-mmusic-sctp-sdp-14](#) (work in progress), March 2015.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

[RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", [RFC 3711](#), DOI 10.17487/RFC3711, March 2004, <<http://www.rfc-editor.org/info/rfc3711>>.

[RFC4574] Levin, O. and G. Camarillo, "The Session Description Protocol (SDP) Label Attribute", [RFC 4574](#), DOI 10.17487/RFC4574, August 2006, <<http://www.rfc-editor.org/info/rfc4574>>.

[RFC5763] Fischl, J., Tschofenig, H., and E. Rescorla, "Framework for Establishing a Secure Real-time Transport Protocol (SRTP) Security Context Using Datagram Transport Layer Security (DTLS)", [RFC 5763](#), DOI 10.17487/RFC5763, May 2010, <<http://www.rfc-editor.org/info/rfc5763>>.

14.2. Informative References

[RFC2506] Holtman, K., Mutz, A., and T. Hardie, "Media Feature Tag Registration Procedure", [BCP 31](#), [RFC 2506](#), DOI 10.17487/RFC2506, March 1999, <<http://www.rfc-editor.org/info/rfc2506>>.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), DOI 10.17487/RFC3264, June 2002, <<http://www.rfc-editor.org/info/rfc3264>>.

[RFC3311] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method", [RFC 3311](#), DOI 10.17487/RFC3311, October 2002, <<http://www.rfc-editor.org/info/rfc3311>>.

- [RFC3840] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", [RFC 3840](#), DOI 10.17487/RFC3840, August 2004, <<http://www.rfc-editor.org/info/rfc3840>>.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", [RFC 5245](#), DOI 10.17487/RFC5245, April 2010, <<http://www.rfc-editor.org/info/rfc5245>>.
- [RFC5630] Audet, F., "The Use of the SIPS URI Scheme in the Session Initiation Protocol (SIP)", [RFC 5630](#), DOI 10.17487/RFC5630, October 2009, <<http://www.rfc-editor.org/info/rfc5630>>.
- [RFC6184] Wang, Y., Even, R., Kristensen, T., and R. Jesup, "RTP Payload Format for H.264 Video", [RFC 6184](#), DOI 10.17487/RFC6184, May 2011, <<http://www.rfc-editor.org/info/rfc6184>>.
- [RFC7202] Perkins, C. and M. Westerlund, "Securing the RTP Framework: Why RTP Does Not Mandate a Single Media Security Solution", [RFC 7202](#), DOI 10.17487/RFC7202, April 2014, <<http://www.rfc-editor.org/info/rfc7202>>.
- [I-D.ietf-mmusic-sdp-bundle-negotiation] Holmberg, C., Alvestrand, H., and C. Jennings, "Negotiating Media Multiplexing Using the Session Description Protocol (SDP)", [draft-ietf-mmusic-sdp-bundle-negotiation-23](#) (work in progress), July 2015.

Authors' Addresses

Paul Kyzivat
Huawei

Email: pkyzivat@alum.mit.edu

Lennard Xiao
Huawei

Email: lennard.xiao@huawei.com

Christian Groves
Huawei

Email: Christian.Groves@nteczone.com

Robert Hansen
Cisco Systems

Email: rohanse2@cisco.com