

CLUE
Internet-Draft
Intended status: Informational
Expires: July 30, 2014

C. Groves, Ed.
W. Yang
R. Even
Huawei
January 26, 2014

CLUE and latent configurations
draft-groves-clue-latent-config-00

Abstract

This document proposes to use Latent Configurations as described by the SDP media capability negotiation framework [RFC6871] for the description and negotiation of CLUE encodings.

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 July 30, 2014.

Copyright Notice

Copyright (c) 2014 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	2
2.	Latent Configurations and CLUE	3
2.1.	Semantics	3
2.2.	Messaging	3
2.3.	Correlation	4
2.4.	Returning an Answer	5
2.5.	Interworking	5
2.6.	Relation to BUNDLE	6
2.7.	Open Issues	6
3.	Example	7
4.	Summary	11
5.	Acknowledgements	11
6.	IANA Considerations	11
7.	Security Considerations	11
8.	References	11
8.1.	Normative References	11
8.2.	Informative References	11
	Authors' Addresses	12

1. Introduction

One of the issues faced in CLUE is how to describe the encodings associated with the Advertised Captures. It was recently decided that this encoding information would not be described in CLUE itself. This means that other methods such as the use of SDP are required to transmit this encoding information. When considering the use of SDP (and in particular the use of SDP Offer/Answer) it should be recognized that there is a semantic difference between a CLUE encoding and an SDP media stream description. Given the nature of a CLUE exchange the encoding represents a Capture/Stream that may occur in the future (i.e. no resources are reserved) whereas a SDP Offer typically relates to resources that are available at that point in time. This does lead to complications when trying to describe CLUE encodings using standard SDP O/A mechanisms.

An alternate approach to using the standard SDP O/A mechanisms for describing CLUE encodings is to use the "SDP Capability Negotiation framework" [RFC5939] and in particular to use the SDP Media Capabilities Negotiation [RFC6871]. [RFC6871] defines "Latent Configurations" as a means to describe media streams that may be used in the future.

2. Latent Configurations and CLUE

The sections below discuss different aspects and benefits of using Latent Configurations to describe CLUE encodings.

2.1. Semantics

[RFC6501] defines a Latent Configuration as:

A latent configuration indicates which combinations of capabilities could be used in a future negotiation for the session and its associated media stream components. Latent configurations are neither ready for use nor offered for actual or potential use in the current offer/answer exchange. Latent configurations merely inform the other side of possible configurations supported by the entity. Those latent configurations may be used to guide subsequent offer/answer exchanges, but they are not offered for use as part of the current offer/answer exchange.

From the above description it can be seen that the semantic of a Latent Configurations closely matches a CLUE message flow. I.e. A set of possible Captures/Encodings (e.g. configurations) are Advertised, the receiver can choose particular Captures/Encodings and then the actual media stream is subsequently established. Therefore the authors believe that use of latent configurations is a good semantic fit with CLUE to describe the encodings.

2.2. Messaging

It has been recognized that CLUE Advertisements may contain a large number of Captures and as there may be multiple encodings per capture potentially a larger number of encodings.

Section 4.2 of CLUE signaling draft [I-D.kyzivat-clue-signaling] indicates the CLUE Provider uses an "m" line for each separate encoding and utilizes the "a=inactive", "a=sendonly" and "a=recvonly" to manage when the media flows are instantiated.

This means that ports must be allocated for these "m" lines and the SDP Offer/Answer [RFC3264] rules regarding maintaining these "m" lines must be followed.

This results in potentially very large SDP descriptions containing superfluous information that must be maintained for the life of the session.

Latent Configurations allow a Provider to advertise potential media without allocating multiple "m" lines or allocating ports for the configurations. The SDP O/A model doesn't apply to Latent Configurations which means that less data is sent over the life of a session.

This allows a Provider to offer a basic media stream for immediate use (i.e. an audio "m" line) and a list of latent configurations for later use without the need for additional m-lines. This is described by [RFC6871]:

A new attribute ("a=lcfg") specifies latent media stream configurations when no corresponding media line ("m=") is offered. An example is the offer of latent configurations for video even though no video is currently offered. If the peer indicates support for one or more offered latent configurations, the corresponding media stream(s) may be added via a new offer/answer exchange.

AND

The "lcfg" attribute is defined as a media-level attribute since it specifies a possible future media stream. However, the "lcfg" attribute is not necessarily related to the media description within which it is provided. Each media line in an SDP description represents an offered simultaneous media stream, whereas each latent configuration represents an additional stream that may be negotiated in a future offer/answer exchange.

The use of Latent Configurations also means that resources aren't tied up and can be allocated when they are needed. i.e. from [RFC6871]:

A latent configuration represents a future capability; hence, the "pt=" parameter is not directly meaningful in the "lcfg" attribute because no actual media session is being offered or accepted. It is permitted in order to tie any payload type number parameters within attributes to the proper media format.

Therefore the authors believe that Latent Configurations provide a clear benefit in terms of messaging size and complexity over normal SDP Offer/Answer mechanism for Advertising CLUE encodings.

2.3. Correlation

One of the issues recognized in the CLUE work is that there needs to be a correlation between the Captures signaled in CLUE and the encodings defined in SDP. The encoding identity (EncodeID) is used

as an identity in CLUE. This same identity is then assigned to the encoding in SDP. Currently [I-D.kyzivat-clue-signaling] indicates that the label attribute [RFC4574] is used to identify the encoding and that it is set per "m" line.

This same method can be used with Latent configurations as they allow the use of SDP attributes in the configurations' description. Section 3.3.8/[RFC6871] shows an example of the use of "label" with a latent configuration, e.g.

```
a=lcfg:4 mt=video t=1 m=1 a=41,42
a=acap:41 label:13
a=acap:42 content:slides
```

The use of Latent Configurations does not require any new SDP attributes to be defined in order for it to be used for CLUE encodings.

2.4. Returning an Answer

A consumer upon receiving an SDP Offer containing CLUE related latent configurations from the Provider could immediately send an SDP answer with the configurations that it could support, i.e. section 3.3.6.1/[RFC6871]:

```
Potential and/or latent configuration attributes may be returned
within an answer SDP to indicate the ability of the answerer to
support alternative configurations of the corresponding
stream(s).
```

The consumer would then send a CLUE Configure indicating the Capture Encodings it wants. The Provider can then subsequently offer actual media streams for the encodings.

2.5. Interworking

One aspect to be considered is the level of adoption of the SDP media capabilities negotiation framework in network entities. Whilst the framework is not widely deployed it is supported by 3GPP (e.g. [SDO-3GPP.24.292]) and is supported by the ITU-T (e.g. [ITU.H248.80.2014] and [ITU.T38.2010]).

It must also be recognized that CLUE itself is something completely new and clients and network equipment must be upgraded to support CLUE signaling. Thus this equipment could also be upgraded to support Latent Configurations at the same time.

In cases where a CLUEfull client sends SDP requesting a CLUE channel and a number of latent configurations to a client that doesn't support CLUE or the media capability framework, the receiving client will ignore the attributes associated with the latent configuration as per normal SDP behavior. Thus there are no interworking issues in this case.

In cases where a CLUEfull client sends SDP requesting a CLUE channel and a number of latent configurations to a client that doesn't support CLUE but DOES support the media capability framework, the receiving client would ignore the CLUE related attributes but could respond with what latent configurations it could support. This would allow the sender to decide if it wanted to offer subsequent media streams. Again there are no interworking issues in this case.

In either of the above cases the session between the clients wouldn't be forced to maintain "m" lines for media that would never be used.

2.6. Relation to BUNDLE

At its core BUNDLE is about using SDP to describe that several "m" lines use the same IP Address/Port for the transport of RTP media. If SDP O/A is being used to describe CLUE encodings then there is a potential interaction in that the CLUE encoding "m" lines would all be subject to the BUNDLE procedures whether or not they were actually used for media.

The use of Latent Configurations would simplify this interaction because Latent Configurations do not allocate or specify ports. They would not be subject to BUNDLE procedures. Once the use of BUNDLE is established (i.e. for the base media streams), then only the media streams (Capture Encodings) that have been chosen by the Consumer can be added to the BUNDLE.

2.7. Open Issues

There are several issues that need to be clarified in [RFC6871] with respect to latent configurations.

1. Latent configurations are specified as media level attributes and thus may be associated with any m-line in an SDP O/A as they don't really pertain to a particular media. There appears to be no guidance as to with m-line they should be associated with in the case of multiple m-line in a SDP Offer.
2. It's not clear from [RFC6871] what happens to latent configurations when an endpoint decides to instantiate the latent configuration as an m-line through an updated SDP Offer.

Section 3.4.4/[RFC6871] discusses modifying the session but has minimal information. It is assumed by the authors that the latent configuration is removed once instantiated.

3. It needs to be clarified whether [RFC6871] indicates that the SDP Answerer SHOULD reply with the latent configurations it supports or whether this is optional. If it's optional what does it mean?
 4. [RFC6871] allows an SDP Answerer to reply with additional latent configurations. However it doesn't specify what action the SDP Offerer should take. This needs to be clarified.
3. Example

This section describes an example session establishment utilizing CLUE and latent configurations. The Provider offers a base audio stream, a CLUE channel (according to [I-D.ietf-mmusic-sctp-sdp] and several latent configurations related to video captures.

```
v=0
o=alice 2890844526 2890844526 IN IP4 host.anywhere.com
s=
c=IN IP4 host.anywhere.com
t=0 0
m=audio 49170 RTP/AVP 0 ; Base audio stream
a=rtpmap:0 PCMU/8000

; CLUE channel establishment request
m=application 49172 SCTP 49172
a=setup:actpass
a=connection:new
a=sctpmmap:49172 CLUE 1

;Offered configurations
a=rmcap:1 H264/90000 ; this is needed for the m=
a=rmcap:2 VP8/90000
a=tcap:1 RTP/AVPF ; for the t=
a=acap:5 sendonly
a=acap:1 label:encoding1
a=lcfg:1 mt=video t=1 m=1|2 a=1,5 pt=1:100,2:101
a=acap:2 label:encoding2
a=lcfg:2 mt=video t=1 m=1|2 a=2,5 pt=1:102,2:103
a=acap:3 label:encoding3
a=lcfg:3 mt=video t=1 m=1|2 a=3,5 pt=1:104,2:105
a=acap:4 label:encoding4
a=lcfg:4 mt=video t=1 m=1|2 a=4,5 pt=1:106,2:107
```

Figure 1: Initial Offer

The receiver is CLUE capable and responds indicating support of the CLUE channel and indicates the IP Address/Port where the Provider should establish a connection to. It also indicates that it only supports H264 encoding.

```
v=0
o=bob 2890844730 2890844730 IN IP4 host.example.com
s=
c=IN IP4 host.example.com
t=0 0
m=audio 49920 RTP/AVP 0
a=rtpmap:0 PCMU/8000

; Accepted SCTP CLUE connection
m=application 54321 SCTP 54321
c=IN IP4 192.0.2.1
a=setup:passive
a=connection:new
a=sctpmap:54321 CLUE 1

;Accepted configurations
a=rmcap:1 H264/90000 ; this is needed for the m=
a=tcap:1 RTP/AVPF ; for the t=
a=acap:5 sendonly
a=acap:1 label:encoding1
a=lcfg:1 mt=video t=1 m=1|2 a=1,5 pt=1:100
a=acap:2 label:encoding2
a=lcfg:2 mt=video t=1 m=1|2 a=2,5 pt=1:102
a=acap:3 label:encoding3
a=lcfg:3 mt=video t=1 m=1|2 a=3,5 pt=1:104
a=acap:4 label:encoding4
a=lcfg:4 mt=video t=1 m=1|2 a=4,5 pt=1:106
```

Figure 2: Answer

Author's note: In the signaling document the grouping framework [RFC5888] is used to indicate the "m" lines that are under CLUE control. It's not clear whether a=group:CLUE and a=mid is needed at this stage or during the updated-Offer. It could be assumed that the above latent configurations are related to CLUE due to the fact they appear under the m=application SCTP/CLUE line.

On receipt of the SDP Answer the Provider establishes the SCTP connection, performs CLUE version and capability negotiation (not shown) and then sends the initial CLUE Advertisement. In it the Provider advertises a single Capture Scene described by three video captures (i.e. Left,Centre,Right) or a video capture of the entire scene.

Author's note: According to the current CLUE protocol work [I-D.presta-clue-protocol], it's the consumer that sends the first Advertisement. The author believes that the Provider should send the first Advertisement to align with the Offer.

Capture Scene #1	
VC1	CaptureArea=Left
VC2	EncodingGroup=EG1
VC3	CaptureArea=Centre
VC4	EncodingGroup=EG2
CSE(VC1,VC2,VC3)	CaptureArea=Right
CSE(VC4)	EncodingGroup=EG3
EncodingGroups	
EG1	EncodeID=encoding1
EG2	EncodeID=encoding2
EG3	EncodeID=encoding3
EG4	EncodeID=encoding4

Figure 3: Advertisement

On receipt of the Advertisement the Consumer determines that it wants the three video captures representing left/centre/right. It sends a CLUE Configure to the Provider:

VC1	encoding1
VC2	encoding2
VC3	encoding3

Figure 4: Configure

On receipt of the CLUE Configure the Provider determines that the Consumer wants to see VC1, VC2 and VC3 according to encoding1, encoding2 and encoding3 respectively. The Provider then issues an updated SDP Offer for three additional video streams. Note: The latent configurations have been removed but the latent configuration related to VC4/encoding4 could also be maintained if still available.

The grouping framework [RFC5888] is used to indicate that the additional video streams are under CLUE control.

```
v=0
o=alice 2890844526 2890844526 IN IP4 host.anywhere.com
s=
a=group:CLUE 1 2 3
c=IN IP4 host.anywhere.com
t=0 0
m=audio 49170 RTP/AVP 0 ; Base audio stream
a=rtpmap:0 PCMU/8000

; CLUE channel estab. Req.
m=application 49172 SCTP 49172
a=setup:actpass
a=connection:new
a=sctpmap:49172 CLUE 1

;Additional video streams
m=video 49174 RTP/AVPF 100
a=rtpmap:100 H264/90000
a=label:encoding1
a=mid:1
a=sendonly

m=video 49176 RTP/AVPF 102
a=rtpmap:102 H264/90000
a=label:encoding2
a=mid:2
a=sendonly

m=video 49178 RTP/AVPF 104
a=rtpmap:104 H264/90000
a=label:encoding3
a=mid:3
a=sendonly
```

Figure 5: Updated Offer

The Consumer then receives the Updated Offer and confirms with an updated Answer. Media flow for the 3 video streams then starts to flow.

4. Summary

The authors believe that the use of Latent Configurations is an ideal way to indicate CLUE encoding information. It is proposed that the use of Latent Configuration is the preferred way of describing CLUE encoding information.

5. Acknowledgements

This template was derived from an initial version written by Pekka Savola and contributed by him to the xml2rfc project.

6. IANA Considerations

It is not expected that the proposed changes present the need for any IANA registrations.

7. Security Considerations

It is not expected that the proposed changes present any addition security issues to the current framework.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

[I-D.ietf-mmusic-sctp-sdp]
Loreto, S. and G. Camarillo, "Stream Control Transmission Protocol (SCTP)-Based Media Transport in the Session Description Protocol (SDP)", draft-ietf-mmusic-sctp-sdp-05 (work in progress), October 2013.

[I-D.kyzivat-clue-signaling]
Kyzivat, P., Xiao, L., Groves, C., and R. Hansen, "CLUE Signaling", draft-kyzivat-clue-signaling-05 (work in progress), September 2013.

[I-D.presta-clue-protocol]
Presta, R. and S. Romano, "CLUE protocol", draft-presta-clue-protocol-03 (work in progress), November 2013.

- [ITU.H248.80.2014]
International Telecommunications Union, "Usage of the revised SDP offer / answer model with H.248", ITU-T Recommendation H.248.80, 2014.
- [ITU.T38.2010]
International Telecommunications Union, "Procedures for real time Group 3 facsimile communication between terminals using IP Networks", ITU-T Recommendation T.38, 2010.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.
- [RFC5888] Camarillo, G. and H. Schulzrinne, "The Session Description Protocol (SDP) Grouping Framework", RFC 5888, June 2010.
- [RFC5939] Andreasen, F., "Session Description Protocol (SDP) Capability Negotiation", RFC 5939, September 2010.
- [RFC6871] Gilman, R., Even, R., and F. Andreasen, "Session Description Protocol (SDP) Media Capabilities Negotiation", RFC 6871, February 2013.
- [SDO-3GPP.24.292]
3GPP, "IP Multimedia (IM) Core Network (CN) subsystem Centralized Services (ICS); Stage 3", 3GPP TS 24.292 10.11.0, June 2013.

Authors' Addresses

Christian Groves (editor)
Huawei
Melbourne
Australia

Email: Christian.Groves@nteczone.com

Weiwei Yang
Huawei
P.R.China

Email: tommy@huawei.com

Roni Even
Huawei
Tel Aviv
Israel

Email: roni.even@mail01.huawei.com