

MMUSIC WG
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2011

M. Garcia-Martin
Ericsson
S. Veikkolainen
Nokia
R. Gilman
March 8, 2011

Title and Bandwidth Capabilities Negotiation in the Session Description
Protocol (SDP)
[draft-mmusic-sdp-icap-bcap-01](#)

Abstract

SDP has been extended with a capability negotiation mechanism framework that allows the endpoints to negotiate transport protocols and attributes. This framework has been extended with a media capabilities negotiation mechanism that allows endpoints to negotiate additional media-related capabilities. This negotiation is embedded into the widely-used SDP offer/answer procedures.

This memo extends the SDP capability negotiation framework to allow endpoints to negotiate two additional SDP capabilities. In particular, this memo provides a mechanism to negotiate titles ("i=" line for each session or media) and bandwidth ("b=" line).

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 September 9, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
2.	Conventions Used in This Document	4
3.	Protocol Description	4
3.1.	Extensions to SDP	4
3.1.1.	Bandwidth Capability	6
3.1.2.	Title Capability	8
3.2.	Session Level versus Media Level	11
3.3.	Offer/Answer model extensions	11
3.3.1.	Generating the Initial Offer	12
3.3.2.	Generating the Answer	12
3.3.3.	Offerer Processing of the Answer	12
3.3.4.	Modifying the Session	12
4.	Field Replacement Rules	12
5.	IANA Considerations	13
5.1.	New SDP Attributes	13
5.2.	New Option Tags	14
5.3.	New SDP Capability Negotiation Configuration Parameters	14
6.	Security Considerations	14
7.	Acknowledgments	14
8.	References	15
8.1.	Normative References	15
8.2.	Informative References	15
	Authors' Addresses	15

1. Introduction

The Session Description Protocol (SDP) [[RFC4566](#)] is intended for describing multimedia sessions for the purposes of session announcement, session invitation, and other forms of multimedia session initiation. SDP has been extended with a capability negotiation mechanism framework [[RFC5939](#)] which allows the endpoints to negotiate capabilities, such as support for Real-time Transport Protocol (RTP) [[RFC3550](#)] and Secure Real-time Transport Protocol (SRTP) [[RFC3711](#)]. The SDP media capabilities [[I-D.ietf-mmusic-sdp-media-capabilities](#)] provides negotiation capabilities to media lines as well.

The capability negotiation is embedded into the widely used SDP offer/answer procedure [[RFC3264](#)]. This memo provides the means to negotiate further capabilities than those specified in the SDP capability negotiation mechanism framework [[RFC5939](#)] and the SDP media capabilities negotiation [[I-D.ietf-mmusic-sdp-media-capabilities](#)]. In particular, this memo provides a mechanism to negotiate session or media titles ("i=") and bandwidth ("b=").

Since the two added capabilities are highly unconnected, it is not expected that implementations will support both at the same time. Instead, it is expected that applications will choose their needed capability for their specific purpose. Due to this, we are writing the normative part pertaining to both capabilities in a self-contained section: [Section 3.1.1](#) describes the bandwidth capability extension, and [Section 3.1.2](#) describes the title capability extension. Separate option tags are defined for both capabilities.

2. Conventions Used in This Document

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 [BCP 14](#), [RFC 2119](#) [[RFC2119](#)] and indicate requirement levels for compliant implementations.

3. Protocol Description

3.1. Extensions to SDP

The SDP Capability Negotiation Framework [[RFC5939](#)] and the SDP media capabilities negotiation [[I-D.ietf-mmusic-sdp-media-capabilities](#)] specify attributes for negotiating SDP capabilities. These documents

specify new attributes (e.g., 'acap', 'tcap', 'mcap') for achieving their purpose. In this document we define two new additional capability attributes for SDP lines of the the general form:

type=value

for types "i" and "b". The corresponding capability attributes are defined as "icap" for title capability, and "bcap" for bandwidth capability, respectively.

From the sub-rules of "a=" line in SDP [[RFC4566](#)], SDP attributes are of the form:

```
attribute      = (att-field ":" att-value) / att-field
att-field      = token
att-value      = byte-string
```

Capability attributes use only the 'att-field:att-value' form.

The new attributes may be referenced in potential configurations ("a=pcfg") or in latent configurations ("a=lcfg"), as productions conforming to the extension-config-list as defined in [[RFC5939](#)].

```
extension-config-list = ["+"] ext-cap-name "=" ext-cap-list
ext-cap-name          = 1*(ALPHA / DIGIT)
                      ; ALPHA and DIGIT defined in RFC5234
ext-cap-list          = 1*VCHAR ; VCHAR defined in RFC5234
```

The optional "+" is used to indicate that the extension is mandatory and MUST be supported in order to use that potential configuration.

The attributes may be referenced in actual configurations ("a=acfg") as productions conforming to the sel-extension-config defined in [[RFC5939](#)].

```
sel-extension-config = ext-cap-name "=" 1*VCHAR
```

The specific parameters are defined in the individual description of each capability, below.

The "icap" and "bcap" capability attributes can be provided either at the session or media level. According to the SDP Capability Negotiation [[RFC5939](#)], each extension capability must specify the implication of making it part of a configuration at the media level.

According to SDP [[RFC4566](#)], "b=" and "i=" lines may appear either at session or media level. In line with this, the "bcap" and "icap" capability attributes, when declared at session level, are to be

interpreted as-if that attribute was provided with that value at the session level. The "bcap" and "icap" capability attributes declared at media level, are to be interpreted as-if that capability attribute was declared at the session level.

For example, extending the example in [\[I-D.ietf-mmusic-sdp-media-capabilities\]](#) with "icap" and "bcap" capability attributes, we get the following SDP:

```
v=0
a=bcap:1 CT:200
a=icap:1 Video conference
m=audio 54320 RTP/AVP 0
a=mcap:1 L16/8000/1
a=mcap:2 L16/16000/2
a=pcfg:1 m=1|2, pt=1:99,2:98
m=video 66544 RTP/AVP 100
a=mcap:3,4 H263-1998/90000
a=rtpmap:100 H264/90000
a=pcfg:10 m=3 pt=3:101 b=1 i=1
```

Example SDP offer with bcap and icap defined at session level

The above SDP defines one PCMU audio stream and one H.264 video stream. It also defines two Media Format capabilities (numbered 1 and 2), using L16 audio at 8 kbps and 16 kbps, respectively, as well as Media Format capabilities for H.263 video (numbered 3 and 4). The Media Format capabilities all appear at the media level. The example also contains a single bandwidth capability and a single title capability at session level. According to the definition above, when the capabilities defined in "bcap", and "icap" attributes are referenced from the potential configuration, in the resulting SDP they are to be interpreted as session level attributes (but the Media Format capabilities are to be interpreted as media level attributes).

[3.1.1.](#) Bandwidth Capability

According to [RFC 4566](#) [[RFC4566](#)] the bandwidth field denotes the proposed bandwidth to be used by the session or media. In this memo, we specify the bandwidth capability attribute which can also appear either at session or media level. The bandwidth field is specified in [RFC 4566](#) [[RFC4566](#)] with the following syntax:

```
b=<bwtype>:<bandwidth>
```

where <bwtype> is an alphanumeric modifier giving the meaning of the <bandwidth> figure.

In this document, we define a new capability attribute: the bandwidth capability attribute "bcap". This attribute lists bandwidth as capabilities according to the following definition:

"a=bcap:" bw-cap-num 1*WSP bwtype ":" bandwidth CRLF

where <bw-cap-num> is a unique integer between 1 and $2^{31}-1$ (both included) user to number the bandwidth capability, and the other elements are as defined for the "b=" field in SDP [[RFC4566](#)].

This format satisfies the general attribute production rules in SDP [[RFC4566](#)] according to the following Augmented Backus-Naur Form (ABNF) [[RFC5234](#)] syntax:

```
att-field      = "bcap"
att-value      = bw-cap-num 1*WSP bwtype ":" bandwidth
bw-cap-num     = 1*10(DIGIT) ; DIGIT defined in RFC5234
```

Negotiation of bandwidth per media stream can be useful when negotiating media encoding capabilities with different bandwidths.

[3.1.1.1](#). Configuration Parameters

The SDP capability negotiation framework [[RFC5939](#)] provides for the existence of the "pcfg" and "acfg" attributes. The concept is extended by the SDP media capabilities negotiation [[I-D.ietf-mmusic-sdp-media-capabilities](#)] with an "lcfg" attribute that conveys latent configurations.

Extensions to the "pcfg" and "lcfg" attributes are defined through <extension-config-list>, and extensions to the "acfg" attribute are defined through the <sel-extension-config> as defined in the SDP Capability Negotiation [[RFC5939](#)].

In this document we extend the <extension-config-list> field to be able to convey lists of bandwidth capabilities in latent or potential configurations, according to the following Augmented Backus-Naur Form (ABNF) [[RFC5234](#)] syntax:

```
extension-config-list = bandwidth-config-list
bandwidth-config-list = ["+"] "b=" bw-cap-list *(BAR bw-cap-list)
                        ; BAR defined in RFC5939
bw-cap-list           = bw-cap-num *(", " bw-cap-num)
bw-cap-num            = 1*10(DIGIT) ; DIGIT defined in RFC5234
```

Figure 1: Syntax of the bandwidth parameter in lcfg and pcfg attributes

Each bandwidth capability configuration is a comma-separated list of bandwidth capability attribute numbers where 'bw-cap-num' refers to the bw-cap-num bandwidth capability numbers defined explicitly earlier in this document, and hence must be between 1 and $2^{31}-1$ (both included). Alternative bandwidth configurations are separated by a vertical bar ("|").

The above syntax is very flexible, allowing referencing to multiple "b=" lines per configuration, even for the same bwtype. While the need for such definitions is not seen, we have not restricted this, as it is not restricted in SDP [RFC4566] either.

The bandwidth parameter to the actual configuration attribute ("a=acfg") is formulated as a sel-extension-config with

```
ext-cap-name = "b"
```

hence

```
sel-extension-config = sel-bandwidth-config  
sel-bandwidth-config = "b=" bw-cap-list ; bw-cap-list as above.
```

Figure 2: Syntax of the bandwidth parameter in acfg attributes

3.1.1.2. Option tag

The SDP Capability Negotiation Framework [RFC5939] allows for capability negotiation extensions to be defined. Associated with each such extension is an option tag that identifies the extension in question. Hereby, we define a new option tag "bcap-v0" that identifies support for the bandwidth capability. The endpoints using the "bcap" capability attribute SHOULD add the option tag to other existing option tags present in the "csup" and "creq" attributes in SDP, according to the procedures defined in the SDP Capability Negotiation Framework [RFC5939].

3.1.2. Title Capability

SDP [RFC4566] provides for the existence of an information field expressed in the format of the "i=" line, which can appear either at the session level or at the media level. An "i=" line that is present at the session level is known as the "session name", and its purpose is to convey a human-readable textual information about the session.

The "i=" line in SDP can also appear at the media level, in which case it is used to provide human-readable information about the media stream to which it is related, e.g., it may indicate the purpose of

the media stream. The "i=" line is not to be confused with the label attribute ("a=label:", [RFC4574](#)) which provides a machine-readable tag. It is foreseen that applications declaring capabilities related to different configurations of a media stream may need to provide different identifying information for each of those configurations. That is, a party might offer alternative media configurations for a stream, each of which represents a different presentation of the same or similar information. For example, an audio stream might offer English or Spanish configurations, or a video stream might offer a choice of video source such as speaker camera, group camera, or document viewer. The title capability is needed to inform the answering user in order to select the proper choice, and the label is used to inform the offering machine which choice the answerer has selected. Hence, there is value in defining a mechanism to provide titles of media streams as capabilities.

According to SDP [RFC4566](#), the session information ("i=") line has the following syntax:

```
"i="text
```

where "text" represents a human-readable text indicating the purpose of the session or media stream.

In this document we define a new capability attribute: the Title capability, "icap". This attribute lists session or media titles as capabilities, according to the following definition:

```
"a=icap:" title-cap-num 1*WSP text
```

where <title-cap-num> is a unique integer between 1 and $2^{31}-1$ (both included) user to number the unique ordinal identifier of the particular title capability and <text> is a human-readable text that indicates the purpose of the session or media stream it is supposed to characterize.

As an example, one might use:

```
a=icap:1 Document Camera
```

to define a title capability number 1 to identify a particular source of a media stream.

The title capability attribute satisfies the general attribute production rules in SDP [RFC4566](#) according to the following Augmented Backus-Naur Form (ABNF) [RFC5234](#) syntax:


```
att-field      = "icap"
att-value      = title-cap-num 1*WSP text
                  ; text defined in RFC4566
title-cap-num  = 1*10(DIGIT)  ; DIGIT defined in RFC5234
```

3.1.2.1. Configuration Parameters

The SDP Capability Negotiation Framework [[RFC5939](#)] provides for the existence of the "pcfg" and "acfg" attributes. The concept is extended by the SDP media capabilities negotiation [[I-D.ietf-mmusic-sdp-media-capabilities](#)] with an "lcfg" attribute that conveys latent configurations.

In this document, we define an <title-config-list> parameter to be used to convey title capabilities in a potential or latent configuration. This parameter is defined as an <extension-config-list> with the following associations:

```
ext-cap-name = "i"

ext-cap-list = title-cap-list
```

This leads to the following definition for the title capability parameter:

```
extension-config-list = title-config-list
title-config-list     = ["+"] "i=" title-cap-list
title-cap-list        = title-cap-num *(BAR title-cap-num)
                        ; BAR defined in RFC5939
title-cap-num         = 1*10(DIGIT) ; DIGIT defined in RFC5234
```

Figure 3: Syntax of the title capability parameter in lcfg and pcfg attributes

Each potential capability configuration contains a single title capability attribute number where 'title-cap-num' is the title capability number defined explicitly earlier in this document, and hence must be between 1 and $2^{31}-1$ (both included). The title capability allows the expression of only a single capability in each alternative, since no more than a single title field is permitted per block. Nevertheless, it is still allowed to express alternative potential title configurations separated by a vertical bar ("|").

An endpoint includes a plus sign ("+") in this configuration attribute to mandate support for this extension. An endpoint that receives this attribute prefixed with a plus sign and does not support this extension MUST treat that potential configuration as not valid.

3.1.2.2. Option Tag

At present, it is difficult to envision a scenario in which the "icap" attribute must be supported or the offer must be rejected. In most cases, if the icap attribute or its contents were to be ignored, an offered configuration could still be chosen based on other criteria such as configuration numbering. However, one might imagine an SDP offer that contained English and Spanish potential configurations for an audio stream. The session might be unintelligible if the choice is based on configuration numbering, rather than informed user selection. Based on such considerations, it may well prove useful to announce the ability to use the icap attribute and its contents to select media configurations, or to inform the user about the selected configuration(s). Therefore, we define a new option tag of "icap-v0" that identifies support for the title capability. This option tag SHOULD be added to other existing option tags present in the "csup" and/or "creq" attributes in SDP, according to the procedures defined in the SDP Capability Negotiation Framework [[RFC5939](#)]. The discussion above suggests that "icap-v0" will typically appear in a "csup" attribute, but rarely in a "creq" attribute.

3.2. Session Level versus Media Level

The "icap" and "bcap" attributes can appear at the session level and/or at the media level. Endpoints MUST interpret capabilities declared at session level as part of the session level in the resulting SDP for that particular configuration. Endpoints MUST interpret capabilities declared at media level as part of the media level in the resulting SDP for that particular configuration.

If an "icap" or "bcap" capability for the same bwtype is declared at both session and media level, the media level attribute overrides the value of the session level attribute.

To avoid confusion, the <type-attr-num> for each "a=bcap" and "a=icap" line must be unique across all capability attributes of the same type within the entire session description.

3.3. Offer/Answer model extensions

In this section, we define extensions to the offer/answer model defined in SDP Offer/Answer Model [[RFC3264](#)] and extended in the SDP Capability Negotiation [[RFC5939](#)] to allow for bandwidth and title capabilities to be used with the SDP Capability Negotiation framework.

3.3.1. Generating the Initial Offer

When an endpoint generates an initial offer and wants to use the functionality described in the current document, it first defines appropriate values for the bandwidth and title capability attributes according to rules defined in [[RFC4566](#)] for "b=" and "i=" lines. The endpoint then MUST include the respective capability attributes and associated values in the SDP offer. The preferred configurations for each media stream are identified following the media line in a "pcfg" attribute. Bandwidth and title capabilities may also be referenced in latent configurations, defined in [[I-D.ietf-mmusic-sdp-media-capabilities](#)].

The offer SHOULD include the level of capability negotiation extensions needed to support this functionality in a "creq" attribute.

3.3.2. Generating the Answer

When the answering party receives the offer, and if it supports the required capability negotiation extensions, it SHOULD select the most preferred configuration it can support for each media stream, and build the answer accordingly, as defined in [Section 3.6.2](#) of the SDP Capability Negotiation [[RFC5939](#)].

3.3.3. Offerer Processing of the Answer

When the offerer receives the answer, it MUST process the media lines according to normal SDP processing rules to identify the media stream(s) accepted by the answer, if any. The "acfg" attribute, if present, may be used to verify the proposed configuration used to form the answer, and to infer the lack of acceptability of higher-preference configurations that were not chosen.

3.3.4. Modifying the Session

If, at a later time, one of the parties wishes to modify the operating parameters of a session, e.g. by adding a new media stream, or by changing the properties used on an existing stream, it may do so via the mechanisms defined for SDP offer/answer [[RFC3264](#)].

4. Field Replacement Rules

To simplify the construction of SDP records, given the need to include fields within the media description in question for endpoints that do not support capabilities negotiation, we define some simple field-replacement rules for those fields invoked by potential or

latent configurations. In particular, any "i=" line invoked by a configuration MUST replace the corresponding line, if present within the media description in question. Any "b=" line invoked by a configuration MUST replace any "b=" of the same bandwidth type at the media level.

5. IANA Considerations

5.1. New SDP Attributes

IANA is hereby requested to register new attributes in the "att-field (both session and media level)" of the "Session Description Protocol (SDP) Parameteres" registry, according to the following registration form:

Attribute name: icap

Long form name: Title Capability

Type of attribute: Both media and session level

Subject to charset: Yes

Purpose: Negotiate human-readable information describing the session or media

Appropriate values: See [Section 3.1.2](#) of RFC XXXX

[Note to the RFC Editor: Please replace the above RFC XXXX with the RFC number of this specification.]

Contact name: Miguel A. Garcia, Miguel.A.Garcia@ericsson.com

Attribute name: bcap

Long form name: Bandwidth Capability

Type of attribute: Both media and session level

Subject to charset: No

Purpose: Negotiate session or media-level bandwidths

Appropriate values: See [Section 3.1.1](#) of RFC XXXX

[Note to the RFC Editor: Please replace the above RFC XXXX with the RFC number of this specification.]

Contact name: Miguel A. Garcia, Miguel.A.Garcia@ericsson.com

5.2. New Option Tags

IANA is hereby requested to add the new option tags "bcap-v0" and "icap-v0", defined herein, to the "SDP Capability Negotiation Option Tag subregistry" of the "Session Description Protocol (SDP) Parameters" registry.

5.3. New SDP Capability Negotiation Configuration Parameters

IANA is hereby requested to add the new parameter identifiers "i" for "title" and "b" for "bandwidth" to the "SDP Capability Negotiation Potential Configuration Parameters" subregistry of the "Session Description Protocol (SDP) Parameters" registry. These parameters are permitted in 'lcfg', 'acfg', and 'pcfg' attributes.

6. Security Considerations

This document provides an extension on top of [RFC 4566](#) [[RFC4566](#)], [RFC 3264](#) [[RFC3264](#)], SDP Capability Negotiation Framework [[RFC5939](#)], and SDP media capabilities negotiation [[I-D.ietf-mmusic-sdp-media-capabilities](#)]. As such, the security considerations of those documents apply.

The bandwidth capability attribute may be used for reserving resources at endpoints and intermediaries which inspect the SDP. Modification of the bandwidth value by an attacker can lead to the network being underutilized (too high bandwidth value) or congested (too low bandwidth value). In case it is essential to protect the bandwidth value, one of the security mechanisms proposed in [[RFC5939](#)] should be used.

The "i=" line and thus the value carried in the title capability attribute is intended for human-readable description only. It should not be parsed programmatically.

7. Acknowledgments

Thanks to Christer Holmberg, Alf Heidermark, and Ingemar Johansson for arguing for the existence of this document and early reviewing it. Thanks to Flemming Andreassen for a detailed review and many improvement suggestions.

8. References

8.1. Normative References

- [I-D.ietf-mmusic-sdp-media-capabilities]
Gilman, R., Even, R., and F. Andreassen, "SDP Media Capabilities Negotiation", [draft-ietf-mmusic-sdp-media-capabilities-11](#) (work in progress), February 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [RFC5939] Andreassen, F., "Session Description Protocol (SDP) Capability Negotiation", [RFC 5939](#), September 2010.

8.2. Informative References

- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", [RFC 3711](#), March 2004.
- [RFC4574] Levin, O. and G. Camarillo, "The Session Description Protocol (SDP) Label Attribute", [RFC 4574](#), August 2006.

Authors' Addresses

Miguel A. Garcia-Martin
Ericsson
Calle Via de los Poblados 13
Madrid, 28033
Spain

Phone: +34 91 339 1000
Email: miguel.a.garcia@ericsson.com

Simo Veikkolainen
Nokia
P.O. Box 407
NOKIA GROUP, FI 00045
Finland

Phone: +358 50 486 4463
Email: simo.veikkolainen@nokia.com

Robert R. Gilman
3243 W. 11th Ave. Dr.
Broomfield, Colorado 80020
U.S.A.

Phone: +1 303 898 9780
Email: bob_gilman@comcast.net

