

Internet Engineering Task Force (IETF)
Request for Comments: 7195
Category: Standards Track
ISSN: 2070-1721

M. Garcia-Martin
Ericsson
S. Veikkolainen
Nokia
May 2014

Session Description Protocol (SDP) Extension for
Setting Audio and Video Media Streams over Circuit-Switched Bearers in
the Public Switched Telephone Network (PSTN)

Abstract

This memo describes use cases, requirements, and protocol extensions for using the Session Description Protocol (SDP) offer/answer model for establishing audio and video media streams over circuit-switched bearers in the Public Switched Telephone Network (PSTN).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7195>.

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	3
2. Conventions Used in This Document	4
3. Requirements	5
4. Overview of Operation	5
4.1. Example Call Flow	6
5. Protocol Description	7
5.1. Level of Compliance	7
5.2. Extensions to SDP	7
5.2.1. Connection Data	7
5.2.2. Media Descriptions	9
5.2.3. Correlating the PSTN Circuit-Switched Bearer with SDP	10
5.2.3.1. The "cs-correlation" Attribute	11
5.2.3.2. Caller ID Correlation Mechanism	12
5.2.3.3. User-User Information Element Correlation Mechanism	13
5.2.3.4. DTMF Correlation Mechanism	14
5.2.3.5. External Correlation Mechanism	15
5.2.3.6. Extensions to Correlation Mechanisms	16
5.3. Negotiating the Correlation Mechanisms	17
5.3.1. Determining the Direction of the Circuit-Switched Bearer Setup	17
5.3.2. Populating the "cs-correlation" Attribute	18
5.3.3. Considerations for Correlations	18
5.4. Considerations for Usage of Existing SDP	19
5.4.1. Originator of the Session	19
5.4.2. Contact Information	20
5.5. Considerations for Usage of Third Party Call Control (3PCC)	20
5.6. Offer/Answer Mode Extensions	20
5.6.1. Generating the Initial Offer	21
5.6.2. Generating the Answer	23
5.6.3. Offerer Processing the Answer	26
5.6.4. Modifying the Session	27
5.7. Formal Syntax	28
6. Examples	30
6.1. Single PSTN Audio Stream	30
6.2. Advanced SDP Example: Circuit-Switched Audio and Video Streams	32
7. Security Considerations	33
8. IANA Considerations	35
8.1. Registration of the New "cs-correlation" SDP Attribute	35
8.2. Registration of a New "nettype" Value	36
8.3. Registration of a New "addrtype" Value	36
8.4. Registration of a New "proto" Value	36
9. Acknowledgments	37

10. References	37
10.1. Normative References	37
10.2. Informative References	38

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 is most commonly used for describing media streams that are transported over the Real-Time Transport Protocol (RTP) [RFC3550], using the profiles for audio and video media defined in "RTP Profile for Audio and Video Conferences with Minimal Control" [RFC3551].

However, SDP can be used to describe media transport protocols other than RTP. Previous work includes SDP conventions for describing ATM bearer connections [RFC3108] and the Message Session Relay Protocol [RFC4975].

SDP is commonly carried in Session Initiation Protocol (SIP) [RFC3261] messages in order to agree on a common media description among the endpoints. "An Offer/Answer Model with the Session Description Protocol (SDP)" [RFC3264] defines a framework by which two endpoints can exchange SDP media descriptions and come to an agreement as to which media streams should be used, along with the media-related parameters.

In some scenarios, it might be desirable to establish the media stream over a circuit-switched bearer connection even if the signaling for the session is carried over an IP bearer. An example of such a scenario is illustrated with two mobile devices capable of both circuit-switched and packet-switched communication over a low-bandwidth radio bearer. The radio bearer may not be suitable for carrying real-time audio or video media, and using a circuit-switched bearer would offer a better perceived quality of service. So, according to this scenario, SDP and its higher-layer session control protocol (e.g., the Session Initiation Protocol (SIP) [RFC3261]) are used over regular IP connectivity, while the audio or video is received through the classical circuit-switched bearer.

This document addresses only the use of circuit-switched bearers in the PSTN, not a generic circuit-switched network. The mechanisms presented below require a call signaling protocol of the PSTN to be used (such as ITU-T Q.931 [ITU.Q931.1998] or 3GPP TS 24.008 [TS.24.008]).

Setting up a signaling relationship in the IP domain instead of just setting up a circuit-switched call also offers the possibility of negotiating, in the same session, other IP-based media that is not sensitive to jitter and delay, for example, text messaging or presence information.

At a later point in time, the mobile device might move to an area where a high-bandwidth packet-switched bearer, for example, a Wireless Local Area Network (WLAN) connection, is available. At this point, the mobile device may perform a handover and move the audio or video media streams over to the high-speed bearer. This implies a new exchange of SDP offer/answer that leads to a renegotiation of the media streams.

Other use cases exist. For example, an endpoint might have at its disposal circuit-switched and packet-switched connectivity, but the same audio or video codecs are not feasible for both access networks. For example, the circuit-switched audio or video stream supports narrow-bandwidth codecs, while the packet-switched access allows any other audio or video codec implemented in the endpoint. In this case, it might be beneficial for the endpoint to describe different codecs for each access type and get an agreement on the bearer together with the remote endpoint.

There are additional use cases related to third party call control where the session setup time is improved when the circuit-switched bearer in the PSTN is described together with one or more codecs.

The rest of the document is structured as follows: Section 2 provides the document conventions, Section 3 introduces the requirements, Section 4 presents an overview of the proposed solutions, and Section 5 contains the protocol description. Section 6 provides examples of circuit-switched audio or video streams in SDP. Sections 7 and 8 contain the Security and IANA considerations, respectively.

2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "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. Requirements

This section presents the general requirements that are specific for the audio or video media streams over circuit-switched bearers.

- REQ-1: A mechanism for endpoints to negotiate and agree on an audio or video media stream established over a circuit-switched bearer MUST be available.
- REQ-2: The mechanism MUST allow the endpoints to combine circuit-switched audio or video media streams with other complementary media streams, for example, text messaging.
- REQ-3: The mechanism MUST allow the endpoint to negotiate the direction of the circuit-switched bearer, i.e., which endpoint is active when initiating the circuit-switched bearer.
- REQ-4: The mechanism MUST be independent of the type of the circuit-switched access (e.g., Integrated Services Digital Network (ISDN), Global System for Mobile Communication (GSM), etc.)
- REQ-5: There MUST be a mechanism that helps an endpoint to correlate an incoming circuit-switched bearer with the one negotiated in SDP, as opposed to another incoming call that is not related to that. In case correlation by programmatic means is not possible, correlation may also be performed by the human user.
- REQ-6: It MUST be possible for endpoints to advertise different lists of audio or video codecs in the circuit-switched audio or video stream from those used in a packet-switched audio or video stream.
- REQ-7: It MUST be possible for endpoints to not advertise the list of available codecs for circuit-switched audio or video streams.

4. Overview of Operation

The mechanism defined in this memo extends SDP [RFC4566] and allows describing an audio or video media stream established over a circuit-switched bearer. A new network type ("PSTN") and a new protocol type ("PSTN") are defined for the "c=" and "m=" lines to be able to describe a media stream over a circuit-switched bearer. These SDP extensions are described in Section 5.2. Since circuit-switched bearers are connection-oriented media streams, the mechanism reuses

the connection-oriented extensions defined in RFC 4145 [RFC4145] to negotiate the active and passive sides of a connection setup. This is further described in Section 5.3.1.

4.1. Example Call Flow

Consider the example presented in Figure 1. In this example, Endpoint A is located in an environment where it has access to both IP and circuit-switched bearers for communicating with other endpoints. Endpoint A decides that the circuit-switched bearer offers a better perceived quality of service for voice and issues an SDP offer containing the description of an audio media stream over a circuit-switched bearer.

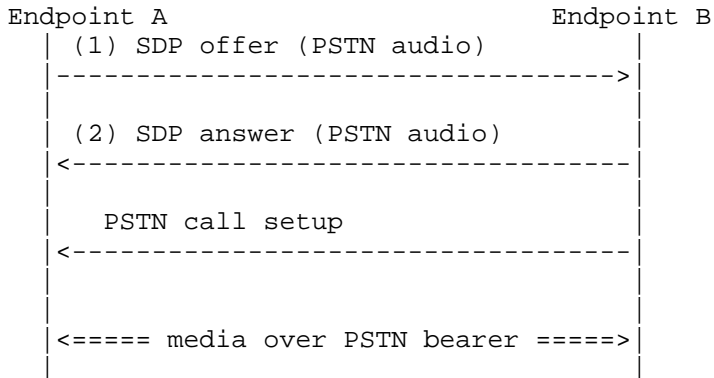


Figure 1: Example Flow

Endpoint B receives the SDP offer and determines that it is located in an environment where the IP-based bearer is not suitable for real-time audio media. However, Endpoint B also has a PSTN circuit-switched bearer available for audio. Endpoint B generates an SDP answer containing a description of the audio media stream over a circuit-switched bearer.

During the offer/answer exchange, Endpoints A and B also agree upon the direction in which the circuit-switched bearer should be established. In this example, Endpoint B becomes the active party; in other words, it establishes the circuit-switched call to the other endpoint. The offer/answer exchange contains identifiers or references that can be used on the circuit-switched network for addressing the other endpoint, as well as information that is used to determine that the incoming circuit-switched bearer establishment is related to the ongoing session between the two endpoints.

Endpoint B establishes a circuit-switched bearer towards Endpoint A using whatever mechanisms are defined for the network type in question. When receiving the incoming circuit-switched connection attempt, Endpoint A is able to determine that the attempt is related to the session it is just establishing with B.

Endpoint A accepts the circuit-switched connection; the circuit-switched bearer setup is completed. The two endpoints can now use the circuit-switched connection for two-way audio media.

If, for some reason, Endpoint B would like to reject the offered stream, it would set the port number of the specific stream to zero, as specified in RFC 3264 [RFC3264]. Also, if B does not understand some of the SDP attributes specified in this document, it would ignore them, as specified in RFC 4566 [RFC4566].

5. Protocol Description

5.1. Level of Compliance

Implementations that are compliant with this specification MUST implement the SDP extensions described in Section 5.2 and MUST implement the considerations discussed in Sections 5.3, 5.4, and 5.6.

5.2. Extensions to SDP

This section provides the syntax and semantics of the extensions required for providing a description of audio or video media streams over circuit-switched bearers in SDP.

5.2.1. Connection Data

According to SDP [RFC4566], the connection data line in SDP has the following syntax:

```
c=<nettype> <addrtype> <connection-address>
```

where <nettype> indicates the network type, <addrtype> indicates the address type, and <connection-address> is the connection address, which is dependent on the address type.

At the moment, the only network type defined is "IN", which indicates Internet network type. The address types "IP4" and "IP6" indicate the type of IP addresses.

This memo defines a new network type for describing a circuit-switched bearer network type in the PSTN. The mnemonic "PSTN" is used for this network type.

For the address type, we initially considered the possibility of describing E.164 telephone numbers. We define a new "E164" address type to be used within the context of a "PSTN" network type. The "E164" address type indicates that the connection address contains an E.164 number represented according to the ITU-T E.164 [ITU.E164.2010] recommendation.

It is a common convention that an international E.164 number contains a leading '+' sign. For consistency's sake, we also require the E.164 telephone is prepended with a '+', even if that is not necessary for routing of the call in the PSTN network.

There are cases, though, when the endpoint is merely aware of a circuit-switched bearer, without having further information about the E.164 number allocated to it. In these cases, a dash ("-") is used to indicate an unknown connection address. This makes the connection data line consistent with SDP syntax.

Please note that the "E164" address type defined in this memo is exclusively defined to be used in conjunction with the "PSTN" network type in accordance with regular offer/answer procedures [RFC4566].

Note: RFC 3108 [RFC3108] also defines address type "E.164". This definition is distinct from the one defined by this memo and shall not be used with <nettype> "PSTN".

This memo exclusively uses the international representation of E.164 numbers, i.e., those including a country code and, as described above, prepended with a '+' sign. Implementations conforming to this specification and using the "E164" address type together with the "PSTN" network type MUST use the 'global-number-digits' construction specified in RFC 3966 [RFC3966] for representing international E.164 numbers. This representation requires the presence of the '+' sign and additionally allows for the presence of one or more 'visual-separator' constructions for easier human readability (see Section 5.7).

Note that <connection-address> MUST NOT be omitted when unknown since this would violate basic syntax of SDP [RFC4566]. In such cases, it MUST be set to a "-".

The following are examples of the extension to the connection data line:

```
c=PSTN E164 +441134960123
```

```
c=PSTN E164 -
```


When the <addrtype> is E164, the connection address is defined as follows:

- o an international E.164 number (prepended with a '+' sign)
- o the value "-", signifying that the address is unknown
- o any other value resulting from the production rule of connection-address in RFC 4566 [RFC4566], but in all cases any value encountered will be ignored.

5.2.2. Media Descriptions

According to SDP [RFC4566], the media description line in SDP has the following syntax:

```
m=<media> <port> <proto> <fmt> ...
```

The <media> subfield carries the media type. For establishing an audio bearer, the existing "audio" media type is used. For establishing a video bearer, the existing "video" media type is used.

The <port> subfield is the transport port to which the media stream is sent. Circuit-switched access lacks the concept of a port number; therefore, the <port> subfield does not carry any meaningful value. In order to be compliant with SDP syntax, implementations SHOULD set the <port> subfield to the discard port value "9" and MUST ignore it on reception.

According to RFC 3264 [RFC3264], a port number of zero in the offer of a unicast stream indicates that the stream is offered but must not be used. If a port number of zero is present in the answer of a unicast stream, it indicates that the stream is rejected. These rules are still valid when the media line in SDP represents a circuit-switched bearer.

The <proto> subfield is the transport protocol. The circuit-switched bearer uses whatever transport protocol it has available. This subfield SHOULD be set to the mnemonic "PSTN" to be syntactically correct with SDP [RFC4566] and to indicate the usage of circuit-switched protocols in the PSTN.

The <fmt> subfield is the media format description. In the classical usage of SDP to describe RTP-based media streams, when the <proto> subfield is set to "RTP/AVP" or "RTP/SAVP", the <fmt> subfield contains the payload types as defined in the RTP audio profile [RFC3551].

When "RTP/AVP" is used in the <proto> field, the <fmt> subfield contains the RTP payload type numbers. We use the <fmt> subfield to indicate the list of available codecs over the circuit-switched bearer, by reusing the conventions and payload type numbers defined for RTP / AVP. The RTP audio and video media types, when applied to PSTN circuit-switched bearers, represent merely an audio or video codec. If the endpoint is able to determine the list of available codecs for circuit-switched media streams, it MUST use the corresponding payload type numbers in the <fmt> subfield.

In some cases, the endpoint is not able to determine the list of available codecs for circuit-switched media streams. In this case, in order to be syntactically compliant with SDP [RFC4566], the endpoint MUST include a single dash ("-") in the <fmt> subfield.

As per RFC 4566 [RFC4566], the media format descriptions are listed in priority order.

Examples of media descriptions for circuit-switched audio streams are:

```
m=audio 9 PSTN 3 0 8
```

```
m=audio 9 PSTN -
```

Similarly, an example of a media description for circuit-switched video stream is:

```
m=video 9 PSTN 34
```

```
m=video 9 PSTN -
```

5.2.3. Correlating the PSTN Circuit-Switched Bearer with SDP

The endpoints should be able to correlate the circuit-switched bearer with the session negotiated with SDP in order to avoid ringing for an incoming circuit-switched bearer that is related to the session controlled with SDP (and SIP).

Several alternatives exist for performing this correlation. This memo provides three mutually non-exclusive correlation mechanisms. Additionally, we define a fourth mechanism where correlation may be performed by external means, typically by the human user, in case using other correlation mechanisms is not possible or does not succeed. Other correlation mechanisms may exist, and their usage will be specified when need arises.

All mechanisms share the same principle: some unique information is sent in the SDP and in the circuit-switched signaling protocol. If these pieces of information match, then the circuit-switched bearer is part of the session described in the SDP exchange. Otherwise, there is no guarantee that the circuit-switched bearer is related to such session.

The first mechanism is based on the exchange of PSTN Caller ID between the endpoints. The Caller ID is also available as the Calling Party Number in the circuit-switched signaling.

The second mechanism is based on the inclusion in SDP of a value that is also sent in the User-User Information Element that is part of the bearer setup signaling in the PSTN.

The third mechanism is based on sending in SDP a string that represents Dual-Tone Multi-Frequency (DTMF) digits that will be later sent right after the circuit-switched bearer is established.

The fourth correlation mechanism declares support for cases where correlation is done by external means. Typically, this means that the decision is left to the human user. This is how some current conferencing systems operate: after logging on to the conference, the system calls back to the user's phone number to establish audio communications, and it is up to the human user to accept or reject the incoming call. By declaring explicit support for this mechanism, endpoints can use it only when such a possibility exists.

Endpoints may opt to implement any combination of the correlation mechanisms specified in Sections 5.2.3.2, 5.2.3.3, 5.2.3.4, and 5.2.3.5, including the option to implement none at all.

5.2.3.1. The "cs-correlation" Attribute

In order to provide support for the correlation mechanisms, we define a new media-level SDP attribute called "cs-correlation". There **MUST** be at most one "cs-correlation" attribute per media description.

This "cs-correlation" attribute **MAY** contain zero or more subfields -- "callerid", "uuie", "dtmf", or "external" to specify additional information required by the Caller ID, User-User Information Element, DTMF, or external correlation mechanisms, respectively. The list of correlation mechanisms may be extended by other specifications; see Section 5.2.3.6 for more details.

The following sections provide more detailed information about these subfields.

The values "callerid", "uuie", "dtmf", and "external" refer to the correlation mechanisms defined in Sections 5.2.3.2, 5.2.3.3, 5.2.3.4, and 5.2.3.5, respectively. The formal Augmented Backus-Naur Format (ABNF) syntax of the "cs-correlation" attribute is presented in Section 5.7.

5.2.3.2. Caller ID Correlation Mechanism

The Caller ID correlation mechanism consists of an exchange of the Calling Party Number as an international E.164 number in SDP, followed by the availability of the Calling Party Number Information Element in the call setup signaling of the circuit-switched connection. If both pieces of information match, the circuit-switched bearer is correlated to the session described in SDP.

An example of inclusion of an international E.164 number in the "cs-correlation" attribute is:

```
a=cs-correlation:callerid:+441134960123
```

The presence of the "callerid" subfield indicates that the endpoint supports use of the Calling Party Number as a means of correlating a PSTN call with the session being negotiated. The "callerid" subfield MAY be accompanied by the international E.164 number of the party inserting the parameter.

Note that there are no guarantees that this correlation mechanism works or is even available, due a number of problems:

- * The endpoint might not be aware of its own E.164 number, in which case it cannot populate the SDP appropriately.
- * The Calling Party Number Information Element in the circuit-switched signaling might not be available, e.g., due to policy restrictions of the network operator or caller restriction due to privacy.
- * The Calling Party Number Information Element in the circuit-switched signaling might be available, but the digit representation of the E.164 number might differ from the one expressed in the SDP, due to, e.g., lack of country code. To mitigate this problem, implementations should consider only some of the rightmost digits from the E.164 number for correlation. For example, the numbers +44-113-496-0123 and 0113-496-0123 could be considered as the same number. This is also the behavior of some cellular phones, which correlate the incoming calling party with a number stored in the phone book,

for the purpose of displaying the caller's name. Please refer to ITU-T E.164 recommendation [ITU.E164.2010] for consideration of the relevant number of digits to consider.

5.2.3.3. User-User Information Element Correlation Mechanism

A second correlation mechanism is based on including in SDP a string that represents the User-User Information Element that is part of the call setup signaling of the circuit-switched bearer. The User-User Information Element is specified in ITU-T Q.931 [ITU.Q931.1998] and 3GPP TS 24.008 [TS.24.008], among others. The User-User Information Element has a maximum size of 35 or 131 octets, depending on the actual message of the PSTN protocol where it is included and the network settings.

The mechanism works as follows. An endpoint creates a User-User Information Element, according to the requirements of the call setup signaling protocol. The same value is included in the SDP offer or SDP answer, in the "uuie" subfield of the "cs-correlation" attribute. When the SDP offer/answer exchange is completed, each endpoint has become aware of the value that will be used in the User-User Information Element of the call setup message of the PSTN protocol. The endpoint that initiates the call setup attempt includes this value in the User-User Information Element. The recipient of the call setup attempt can extract the User-User Information Element and correlate it with the value previously received in the SDP. If both values match, then the call setup attempt corresponds to that indicated in the SDP.

According to ITU-T Q.931 [ITU.Q931.1998], the User-User Information Element (UUIE) identifier is composed of a first octet identifying this as a User-User Information Element, a second octet containing the length of the user-user contents, a third octet containing a Protocol Discriminator, and a value of up to 32 or 128 octets (depending on network settings) containing the actual User Information (see Figure 4-36 in [ITU.Q931.1998]). The first two octets of the UUIE MUST NOT be used for correlation; only the octets carrying the Protocol Discriminator and the User Information value are input to the creation of the value of the "uuie" subfield in the "cs-correlation" attribute. Therefore, the value of the "uuie" subfield in the "cs-correlation" attribute MUST start with the Protocol Discriminator octet, followed by the User Information octets. The value of the Protocol Discriminator octet is not specified in this document; it is expected that organizations using this technology will allocate a suitable value for the Protocol Discriminator.

Once the binary value of the "uuie" subfield in the "cs-correlation" attribute is created, it MUST be base 16 (also known as "hex") encoded before it is inserted in SDP. Please refer to RFC 4648 [RFC4648] for a detailed description of base 16 encoding. The resulting encoded value needs to have an even number of hexadecimal digits and MUST be considered invalid if it has an odd number.

Note: The encoding of the "uuie" subfield of the "cs-correlation" attribute is largely inspired by the encoding of the same value in the User-to-User header field in SIP, according to "A Mechanism for Transporting User to User Call Control Information in SIP" [SIP-UUI].

As an example, an endpoint willing to send a UUIE containing a Protocol Discriminator with the hexadecimal value of %x56 and an hexadecimal User Information value of %xA390F3D2B7310023 would include an "a=cs-correlation" attribute line as follows:

```
a=cs-correlation:uuie:56A390F3D2B7310023
```

Note that the value of the User-User Information Element is considered as an opaque string and only used for correlation purposes. Typically, call signaling protocols impose requirements on the creation of a User-User Information Element for end-user protocol exchange. The details regarding the generation of the User-User Information Element are outside the scope of this specification.

Please note that there are no guarantees that this correlation mechanism works. On one side, policy restrictions might not make the User-User information available end to end in the PSTN. On the other hand, the generation of the User-User Information Element is controlled by the PSTN circuit-switched call protocol, which might not offer enough freedom for generating different values from one endpoint to another one or from one call to another in the same endpoint. This might result in the same value of the User-User Information Element for all calls.

5.2.3.4. DTMF Correlation Mechanism

We introduce a third mechanism for correlating the circuit-switched bearer with the session described with SDP. This is based on agreeing on a sequence of digits that are negotiated in the SDP offer/answer exchange and sent as DTMF tones as described in ITU-T Recommendation Q.23 [ITU.Q23.1988] over the circuit-switched bearer once this bearer is established. If the DTMF digit sequence received through the circuit-switched bearer matches the digit string negotiated in the SDP, the circuit-switched bearer is correlated with

the session described in the SDP. The mechanism is similar to many voice conferencing systems that require the user to enter a PIN code using DTMF tones in order to be accepted in a voice conference.

The mechanism works as follows. An endpoint selects a DTMF digit sequence. The same sequence is included in the SDP offer or SDP answer, in a "dtmf" subfield of the "cs-correlation" attribute. When the SDP offer/answer exchange is completed, each endpoint has become aware of the DTMF sequence that will be sent right after the circuit-switched bearer is set up. The endpoint that initiates the call setup attempt sends the DTMF digits according to the procedures defined for the circuit-switched bearer technology used. The recipient (passive side of the bearer setup) of the call setup attempt collects the digits and compares them with the value previously received in the SDP. If the digits match, then the call setup attempt corresponds to that indicated in the SDP.

Note: Implementations are advised to select a number of DTMF digits that provide enough assurance that the call is related but do not prolong the bearer setup time unnecessarily. A number of 5 to 10 digits is a good compromise.

As an example, an endpoint willing to send DTMF tone sequence "14D*3" would include an "a=cs-correlation" attribute line as follows:

```
a=cs-correlation:dtmf:14D*3
```

If the endpoints successfully agree on the usage of the DTMF digit correlation mechanism but the passive side does not receive any DTMF digits after successful circuit-switched bearer setup or receives a set of DTMF digits that do not match the value of the "dtmf" attribute (including receiving too many digits), the passive side SHOULD consider that this DTMF mechanism has failed to correlate the incoming call.

5.2.3.5. External Correlation Mechanism

The fourth correlation mechanism relies on external means for correlating the incoming call to the session. Since endpoints can select which correlation mechanisms they support, it may happen that no other common correlation mechanism is found or that the selected correlation mechanism does not succeed due to the required feature not being supported by the underlying PSTN network. In these cases, the human user can make the decision to accept or reject the incoming call, thus "correlating" the call with the session. Since not all endpoints are operated by a human user and since there may be no

other external means implemented by the endpoint for the correlation function, we explicitly define support for such an external correlation mechanism.

Endpoints wishing to use this external correlation mechanism would use the "external" subfield in the "cs-correlation" attribute. Unlike the other three correlation mechanisms, the "external" subfield does not accept a value. The following is an example of an "a=cs-correlation" attribute line:

```
a=cs-correlation:external
```

Endpoints that are willing to only use the three explicit correlation mechanisms defined in this document ("callerid", "uuie", and/or "dtmf") would not include the "external" mechanism in the offer/answer exchange.

The external correlation mechanism typically relies on the human user to make the decision on whether or not the call is related to the ongoing session. After the user accepts the call, that bearer is considered as related to the session. There is a small chance that the user receives at the same time another circuit-switched call that is not related to the ongoing session. The user may reject this call if he is able to determine (e.g., based on the calling line identification) that the call is not related to the session and continue waiting for another call attempt. If the user accepts the incoming circuit-switched call, but it turns out to be not related to the session, the endpoints need to rely on the human user to take appropriate action (typically, the user would hang up).

5.2.3.6. Extensions to Correlation Mechanisms

New values for the "cs-correlation" attribute may be specified. The registration policy for new values is "Specification Required"; see Section 8. Any such specification MUST include a description of how the SDP offer/answer mechanism is used to negotiate the use of the new values, taking into account how endpoints determine which side will become active or passive (see Section 5.3 for more details).

If, during the offer/answer negotiation, either endpoint encounters an unknown value in the "cs-correlation" attribute, it MUST consider that mechanism as unsupported and MUST NOT include that value in subsequent offer/answer negotiation.

5.3. Negotiating the Correlation Mechanisms

The four correlation mechanisms presented above (based on Called Party Number, User-User Information Element, DTMF digit sending, and external) are non-exclusive and can be used independently of each other. In order to know how to populate the "cs-correlation" attribute, the endpoints need to agree which endpoint will become the active party, i.e., the one that will set up the circuit-switched bearer.

5.3.1. Determining the Direction of the Circuit-Switched Bearer Setup

In order to avoid a situation where both endpoints attempt to initiate a connection simultaneously, the direction in which the circuit-switched bearer is set up MUST be negotiated during the offer/answer exchange.

The framework defined in RFC 4145 [RFC4145] allows the endpoints to agree which endpoint acts as the active endpoint when initiating a TCP connection. While RFC 4145 [RFC4145] was originally designed for establishing TCP connections, it can be easily extrapolated to the connection establishment of circuit-switched bearers. This specification uses the concepts specified in RFC 4145 [RFC4145] for agreeing on the direction of establishment of a circuit-switched bearer.

RFC 4145 [RFC4145] defines two new attributes in SDP: "setup" and "connection". The "setup" attribute indicates which of the endpoints should initiate the connection establishment of the PSTN circuit-switched bearer. Four values are defined in Section 4 of RFC 4145 [RFC4145]: "active", "passive", "actpass", and "holdconn". Please refer to Section 4 of RFC 4145 [RFC4145] for a detailed description of this attribute.

The "connection" attribute indicates whether a new connection is needed or an existing connection is reused. The attribute can take the values "new" or "existing". Please refer to Section 5 of RFC 4145 [RFC4145] for a detailed description of this attribute.

Implementations that are compliant with this specification MUST support the "setup" and "connection" attributes specified in RFC 4145 [RFC4145], but applied to circuit-switched bearers in the PSTN.

We define the active party as the one that initiates the circuit-switched bearer after the offer/answer exchange. The passive party is the one receiving the circuit-switched bearer. Either party may indicate its desire to become the active or passive party during the offer/answer exchange using the procedures described in Section 5.6.

5.3.2. Populating the "cs-correlation" Attribute

By defining values for the subfields in the "cs-correlation" attribute, the endpoint indicates that it is willing to become the active party and that it can use those values in the Calling Party Number, in the User-User Information Element, or as DTMF tones during the circuit-switched bearer setup.

Thus, the following rules apply:

- o An endpoint that can only become the active party in the circuit-switched bearer setup MUST include all correlation mechanisms it supports in the "cs-correlation" attribute and MUST also specify values for the "callerid", "uuie", and "dtmf" subfields. Notice that the "external" subfield does not accept a value.
- o An endpoint that can only become the passive party in the circuit-switched bearer setup MUST include all correlation mechanisms it supports in the "cs-correlation" attribute but MUST NOT specify values for the subfields.
- o An endpoint that is willing to become either the active or passive party (by including the "a=setup:actpass" attribute in the offer) MUST include all correlation mechanisms it supports in the "cs-correlation" attribute and MUST also specify values for the "callerid", "uuie", and "dtmf" subfields. Notice that the "external" subfield does not accept a value.

5.3.3. Considerations for Correlations

Passive endpoints should expect an incoming circuit-switched (CS) call for setting up the audio bearer. Passive endpoints MAY suppress the incoming CS alert during certain time periods. Additional restrictions can be applied, such as the passive endpoint not alerting incoming calls originated from the number that was observed during the offer/answer negotiation.

There may be cases when an endpoint is not willing to include one or more correlation mechanisms in the "a=cs-correlation" attribute line even if it supports it. For example, some correlation mechanisms can be omitted if the endpoint is certain that the PSTN network does not support carrying the correlation identifier. Also, since using the DTMF-based correlation mechanism requires the call to be accepted before DTMF tones can be sent, some endpoints may enforce a policy restricting this due to, for example, cost associated with received calls, making the DTMF-based mechanism unusable.

Note that it cannot be guaranteed that the correlation mechanisms relying on caller identification, User-User Information Element, and DTMF sending will succeed even if the usage of those was agreed beforehand. This is due to the fact that correlation mechanisms require support from the circuit-switched bearer technology used.

Therefore, even a single positive indication using any of these mechanisms SHOULD be interpreted by the passive endpoint so that the circuit-switched bearer establishment is related to the ongoing session, even if the other correlation mechanisms fail.

If, after successfully negotiating any of the "callerid", "uuie", or "dtmf" correlation mechanisms in the SDP offer/answer exchange, an endpoint receives an incoming establishment of a circuit-switched bearer with no correlation information present, the endpoint first checks whether or not the offer/answer exchange was also used to successfully negotiate the "external" correlation mechanism. If it was, the endpoint should let the decision be made by external means, typically the human user. If the "external" correlation mechanism was not successfully negotiated, the endpoint should treat the call as unrelated to the ongoing session in the IP domain.

5.4. Considerations for Usage of Existing SDP

5.4.1. Originator of the Session

According to SDP [RFC4566], the origin line in SDP has the following syntax:

```
o=<username> <sess-id> <sess-version> <nettype> <addrtype>
<unicast-address>
```

Of interest here are the <nettype> and <addrtype> fields, which indicate the type of network and type of address, respectively. Typically, this field carries the IP address of the originator of the session. Even if the SDP was used to negotiate an audio or video media stream transported over a circuit-switched bearer, the originator is using SDP over an IP bearer. Therefore, <nettype> and <addrtype> fields in the "o=" line should be populated with the IP address identifying the source of the signaling.

5.4.2. Contact Information

SDP [RFC4566] defines the "p=" line, which may include the phone number of the person responsible for the conference. Even though this line can carry a phone number, it is not suited for the purpose of defining a connection address for the media. Therefore, we have selected to define the PSTN-specific connection addresses in the "c=" line.

5.5. Considerations for Usage of Third Party Call Control (3PCC)

"Best Current Practices for Third Party Call Control (3PCC) in the Session Initiation Protocol (SIP)" [RFC3725] outlines several flows that are possible in third party call control scenarios and recommends some flows for specific situations.

One of the assumptions in [RFC3725] is that an SDP offer may include a "black hole" connection address, which has the property that packets sent to it will never leave the host that sent them. For IPv4, this "black hole" connection address is 0.0.0.0 or a domain name within the .invalid DNS top level domain.

When using an E.164 address scheme in the context of third party call control, when the User Agent needs to indicate an unknown phone number, it MUST populate the <addrtype> of the SDP "c=" line with a "-" string.

Note: This may result in the recipient of the initial offer rejecting such offer if the recipient of the offer was not aware of its own E.164 number. Consequently, it will not be possible to establish a circuit-switched bearer, since neither party is aware of its E.164 number.

5.6. Offer/Answer Mode Extensions

In this section, we define extensions to the offer/answer model defined in "An Offer/Answer Model with the Session Description Protocol (SDP)" [RFC3264] to allow for PSTN addresses to be used with the offer/answer model.

5.6.1. Generating the Initial Offer

The offerer, wishing to use PSTN audio or video stream, MUST populate the "c=" and "m=" lines as follows.

The endpoint MUST set the <nettype> in the "c=" line to "PSTN" and the <addrtype> to "E164". Furthermore, the endpoint SHOULD set the <connection-address> field to its own international E.164 number (with a leading "+"). If the endpoint is not aware of its own E.164 number, it MUST set the <connection-address> to "-".

In the "m=" line, the endpoint MUST set the <media> subfield to "audio" or "video", depending on the media type, and the <proto> subfield to "PSTN". The <port> subfield SHOULD be set to "9" (the discard port). The values "audio" or "video" in the <media> subfield MUST NOT be set by the endpoint unless it has knowledge that these bearer types are available on the circuit-switched network.

The <fmt> subfield carries the payload type number(s) the endpoint is wishing to use. Payload type numbers in this case refer to the codecs that the endpoint wishes to use on the PSTN media stream. For example, if the endpoint wishes to use the GSM codec, it would add payload type number 3 in the list of codecs. The list of payload types MUST only contain those codecs the endpoint is able to use on the PSTN bearer. In case the endpoint is not aware of the codecs available for the circuit-switched media streams, it MUST include a dash ("-") in the <fmt> subfield.

The mapping table of static payload types numbers to payload types is initially specified in [RFC3551] and maintained by IANA. For dynamic payload types, the endpoint MUST define the set of valid encoding names and related parameters using the "a=rtpmap" attribute line. See Section 6 of RFC 4566 [RFC4566] for details.

When generating the offer, the offerer MUST include an "a=cs-correlation" attribute line in the SDP offer. The offerer MUST NOT include more than one "cs-correlation" attribute per media description. The "a=cs-correlation" line SHOULD contain an enumeration of all the correlation mechanisms supported by the offerer, in the format of subfields. See Section 5.3.3 for more information on usage of the correlation mechanisms.

The current list of subfields include "callerid", "uuie", "dtmf", and "external", and they refer to the correlation mechanisms defined in Sections 5.2.3.2, 5.2.3.3, 5.2.3.4, and 5.2.3.5, respectively.

If the offerer supports any of the correlation mechanisms defined in this memo and is willing to become the active party, the offerer MUST add the "callerid", "uuie", "dtmf", and/or "external" subfields and MUST specify values for them as follows:

- o The international E.164 number as the value in the "callerid" subfield.
- o The contents of the User-User Information Element as the value of the "uuie" subfield.
- o The DTMF tone string as the value of the "dtmf" subfield.
- o The endpoint MUST NOT specify any value for the "external" subfield.

If the offerer is only able to become the passive party in the circuit-switched bearer setup, it MUST add at least one of the possible correlation mechanisms but MUST NOT specify values for those subfields.

For example, if the offerer is willing to use the User-User Information Element and DTMF digit-sending mechanisms but can only become the passive party, and is also able to let the human user decide whether the correlation should be done or not, it includes the following lines in the SDP:

```
a=cs-correlation:uuie dtmf external  
  
a=setup:passive
```

If, on the other hand, the offerer is willing to use the User-User Information Element and the DTMF correlation mechanisms and is able to become the active or passive side, and is also able to let the human user decide whether the correlation should be done or not, it includes the following lines in the SDP:

```
a=cs-correlation:uuie:56A390F3D2B7310023 dtmf:14D*3 external  
  
a=setup:actpass
```

The negotiation of the value of the "setup" attribute takes place as defined in Section 4.1 of RFC 4145 [RFC4145].

The offerer states which role or roles it is willing to perform; the answerer, taking the offerer's willingness into consideration, chooses which roles both endpoints will actually perform during the circuit-switched bearer setup.

By "active" endpoint, we refer to an endpoint that will establish the circuit-switched bearer; by "passive" endpoint, we refer to an endpoint that will receive a circuit-switched bearer.

If an offerer does not know its international E.164 number, it **MUST** set the "setup" attribute to the value "active". If the offerer knows its international E.164 number, it **SHOULD** set the value to either "actpass" or "passive".

Also "holdconn" is a permissible value in the "setup" attribute. It indicates that the connection should not be established for the time being.

The offerer uses the "connection" attribute to decide whether a new circuit-switched bearer is to be established or not. For the initial offer, the offerer **MUST** use value "new".

5.6.2. Generating the Answer

If the offer contained a circuit-switched audio or video stream, the answerer first determines whether it is able to accept and use such streams on the circuit-switched network. If the answerer does not support or is not willing to use circuit-switched media for the session, it **MUST** construct an answer where the port number for such media stream(s) is set to zero, according to Section 6 of [RFC3264]. If the answerer is willing to use circuit-switched media for the session, it **MUST** ignore the received port number (unless the port number is set to zero).

If the offer included a "-" as the payload type number, it indicates that the offerer is not willing or able to define any specific payload type. Most often, a "-" is expected to be used instead of the payload type when the endpoint is not aware of or not willing to define the codecs that will eventually be used on the circuit-switched bearer. The circuit-switched signaling protocols have their own means of negotiating or indicating the codecs; therefore, an answerer **SHOULD** accept such offers and **SHOULD** set the payload type to "-" in the answer.

If the answerer explicitly wants to specify a codec for the circuit-switched media, it **MAY** set the respective payload numbers in the <fmt> subfield in the answer. This behavior, however, is **NOT RECOMMENDED**.

When receiving the offer, the answerer **MUST** determine whether it becomes the active or passive party.

If the SDP in the offer indicates that the offerer is only able to become the active party, the answerer needs to determine whether it is able to become the passive party. If this is not possible, e.g., due to the answerer not knowing its international E.164 number, the answerer MUST reject the circuit-switched media by setting the port number to zero on the answer. If the answerer is aware of its international E.164 number, it MUST include the "setup" attribute in the answer and set it to value "passive" or "holdconn". The answerer MUST also include its E.164 number in the "c=" line.

If the SDP in the offer indicates that the offerer is only able to become the passive party, the answerer MUST verify that the offerer's E.164 number is included in the "c=" line of the offer. If the number is included, the answerer MUST include the "setup" attribute in the answer and set it to value "active" or "holdconn". If the number is not included, the recipient of the offer is not willing to establish a connection the E.164 based on a priori knowledge of cost, or other reasons, call establishment is not possible, and the answerer MUST reject the circuit-switched media by setting the port number to zero in the answer.

If the SDP in the offer indicates that the offerer is able to become either the active or passive party, the answerer determines which role it will take. If the offer includes an international E.164 number in the "c=" line, the answerer SHOULD become the active party. If the answerer does not become the active party and if the answerer is aware of its E.164 number, it MUST become the passive party. If the answerer does not become the active or the passive party, it MUST reject the circuit-switched media by setting the port number to zero in the answer.

For each media description where the offer includes a "cs-correlation" attribute, the answerer MUST select from the offer those correlation mechanisms it supports and include in the answer one "a=cs-correlation" attribute line containing those mechanisms it is willing to use. The answerer MUST only add one "cs-correlation" attribute in those media descriptions where also the offer included a "cs-correlation" attribute. The answerer MUST NOT add any mechanisms that were not included in the offer. If there is more than one "cs-correlation" attribute per media description in the offer, the answerer MUST discard all but the first for any media description. Also, the answerer MUST discard all unknown "cs-correlation" attribute values.

If the answerer becomes the active party, it MUST add a value to any of the possible subfields.

If the answerer becomes the passive party, it MUST NOT add any values to the subfields in the "cs-correlation" attribute.

After generating and sending the answer, if the answerer became the active party, it

- o MUST extract the E.164 number from the "c=" line of the offer and MUST establish a circuit-switched bearer to that address.
- o if the SDP answer contained a value for the "callerid" subfield, MUST set the Calling Party Number Information Element to that number.
- o if the SDP answer contained a value for the "uuie" subfield, MUST send the User-User Information Element according to the rules defined for the circuit-switched technology used and set the value of the Information Element to that received in the SDP offer.
- o if the SDP answer contained a value for the "dtmf" subfield, MUST send those DTMF digits according to the circuit-switched technology used.

If, on the other hand, the answerer became the passive party, it

- o MUST be prepared to receive a circuit-switched bearer,
- o if the offer contained a value for the "callerid" subfield, MUST compare that value to the Calling Party Number Information Element of the circuit-switched bearer. If the received Calling Party Number Information Element matches the value of the "callerid" subfield, the call SHOULD be treated as correlated to the ongoing session.
- o if the offer contained a value for the "dtmf" subfield, MUST be prepared to receive and collect DTMF digits once the circuit-switched bearer is set up. The answerer MUST compare the received DTMF digits to the value of the "dtmf" subfield. If the received DTMF digits match the value of the "dtmf" subfield in the "cs-correlation" attribute, the call SHOULD be treated as correlated to the ongoing session.
- o if the offer contained a value for the "uuie" subfield, MUST be prepared to receive a User-User Information Element once the circuit-switched bearer is set up. The answerer MUST compare the received UUIE to the value of the "uuie" subfield. If the value of the received UUIE matches the value of the "uuie" subfield, the call SHOULD be treated as correlated to the ongoing session.

- o if the offer contained an "external" subfield, MUST be prepared to receive a circuit-switched call and use the external means (typically, the human user) for accepting or rejecting the call.

If the answerer becomes the active party, generates an SDP answer, and then it finds out that the circuit-switched call cannot be established, then the answerer MUST create a new SDP offer where the circuit-switched stream is removed from the session (actually, by setting the corresponding port in the "m=" line to zero) and send it to its counterpart. This is to synchronize both parties (and potential intermediaries) on the state of the session.

5.6.3. Offerer Processing the Answer

When receiving the answer, if the SDP does not contain an "a=cs-correlation" attribute line, the offerer should take that as an indication that the other party does not support or is not willing to use the procedures defined in the document for this session and MUST revert to normal processing of SDP.

When receiving the answer, the offerer MUST first determine whether it becomes the active or passive party, as described in Section 5.3.1.

If the offerer becomes the active party, it

- o MUST extract the E.164 number from the "c=" line and MUST establish a circuit-switched bearer to that address.
- o if the SDP answer contained a value for the "uuie" subfield, MUST send the User-User Information Element according to the rules defined for the circuit-switched technology used and set the value of the Information Element to that received in the SDP answer.
- o if the SDP answer contained a value for the "dtmf" subfield, MUST send those DTMF digits according to the circuit-switched technology used.

If the offerer becomes the passive party:

- o It MUST be prepared to receive a circuit-switched bearer.
- o Note that if delivery of the answer is delayed for some reason, the circuit-switched call attempt may arrive at the offerer before the answer has been processed. In this case, since the correlation mechanisms are negotiated as part of the offer/answer exchange, the answerer cannot know whether or not the incoming

circuit-switched call attempt is correlated with the session being negotiated; thus, the offerer SHOULD answer the circuit-switched call attempt only after it has received and processed the answer.

- o If the answer contained a value for the "dtmf" subfield, the offerer MUST be prepared to receive and collect DTMF digits once the circuit-switched bearer is set up. The offerer SHOULD compare the received DTMF digits to the value of the "dtmf" subfield. If the received DTMF digits match the value of the "dtmf" subfield in the "cs-correlation" attribute, the call SHOULD be treated as correlated to the ongoing session.
- o If the answer contained a value for the "uuie" subfield, the offerer MUST be prepared to receive a User-User Information Element once the circuit-switched bearer is set up. The offerer SHOULD compare the received UUIE to the value of the "uuie" subfield. If the value of the received UUIE matches the value of the "uuie" subfield, the call SHOULD be treated as correlated to the ongoing session.
- o If the answer contained an "external" subfield, the offerer MUST be prepared to receive a circuit-switched call and use the external means (typically, the human user) for accepting or rejecting the call.

According to the "An Offer/Answer Model with the Session Description Protocol (SDP)" [RFC3264], the offerer needs to be ready to receive media as soon as the offer has been sent. It may happen that the answerer, if it became the active party, will initiate a circuit-switched bearer setup that will arrive at the offerer before the answer has arrived. However, the offerer needs to receive the answer and examine the information about the correlation mechanisms in order to successfully perform correlation of the circuit-switched call to the session. Therefore, if the offerer receives an incoming circuit-switched call, it MUST NOT accept the call before the answer has been received. If no answer is received during an implementation-specific time, the offerer MUST either modify the session according to [RFC3264] or terminate it according to the session signaling procedures in question (for terminating a SIP session, see Section 15 of [RFC3261]).

5.6.4. Modifying the Session

If, at a later time, one of the parties wishes to modify the session, e.g., by adding a new media stream or by changing properties used on an existing stream, it may do so via the mechanisms defined in "An Offer/Answer Model with the Session Description Protocol (SDP)" [RFC3264].

If there is an existing circuit-switched bearer between the endpoints and the offerer wants to reuse that, the offerer MUST set the value of the "connection" attribute to "existing".

If either party removes the circuit-switched media from the session (by setting the port number to zero), it MUST terminate the circuit-switched bearer using whatever mechanism is appropriate for the technology in question.

If either party wishes to drop and reestablish an existing call, that party MUST first remove the circuit-switched media from the session by setting the port number to zero and then use another offer/answer exchange where it MUST set the "connection" attribute to "new". If the media types are different (for example, a different codec will be used for the circuit-switched bearer), the media descriptions for terminating the existing bearer and the new bearer can be in the same offer.

If either party would like to remove existing RTP-based media from the session and replace that with a circuit-switched bearer, it would create a new offer to add the circuit-switched media as described in Section 5.6.1 above, replacing the RTP-based media description with the circuit-switched media description, as specified in RFC 3264 [RFC3264].

Once the offer/answer exchange is done, but the circuit-switched bearer is not yet established, there may be a period of time when no media is available. Also, it may happen that correlating the circuit-switched call fails for reasons discussed in Section 5.3.3. In this case, even if the offer/answer exchange was successful, endpoints are not able to receive or send media. It is up to the implementation to decide the behavior in this case; if nothing else is done, the user most likely hangs up after a while if there is no other media in the session. Note that this may also happen when switching from one RTP media to another RTP media (for example, when firewall blocks the new media stream).

If either party would like to remove existing circuit-switched media from the session and replace that with RTP-based media, it would modify the media description as per the procedures defined in RFC 3264 [RFC3264]. The endpoint MUST then terminate the circuit-switched bearer using whatever mechanism is appropriate for the technology in question.

5.7. Formal Syntax

The following is the formal Augmented Backus-Naur Form (ABNF) [RFC5234] syntax that supports the extensions defined in this

specification. The syntax is built above the SDP [RFC4566] and the tel URI [RFC3966] grammars. Implementations that are compliant with this specification MUST be compliant with this syntax.

Figure 2 shows the formal syntax of the extensions defined in this memo.

```

; extension to the connection field originally specified
; in RFC 4566

connection-field = [%x63 "=" nettype SP addrtype SP
connection-address CRLF]
; CRLF defined in RFC 5234

;nettype and addrtype are defined in RFC 4566

connection-address =/ global-number-digits / "-"
; global-number-digits specified in RFC 3966

;subrules for correlation attribute
attribute =/ cs-correlation-attr
; attribute defined in RFC 4566
cs-correlation-attr = "cs-correlation:" corr-mechanisms
corr-mechanisms = corr-mech *(SP corr-mech)
corr-mech = caller-id-mech / uuie-mech /
dtmf-mech / external-mech /
ext-mech
caller-id-mech = "callerid" [":" caller-id-value]
caller-id-value = "+" 1*15DIGIT
; DIGIT defined in RFC 5234
uuie-mech = "uuie" [":" uuie-value]
uuie-value = 1*65(HEXDIG HEXDIG)
;This represents up to 130 HEXDIG
; (65 octets)
;HEXDIG defined in RFC 5234
;HEXDIG defined as 0-9, A-F
dtmf-mech = "dtmf" [":" dtmf-value]
dtmf-value = 1*32(DIGIT / %x41-44 / %x23 / %x2A )
;0-9, A-D, '#' and '*'
external-mech = "external"
ext-mech = ext-mech-name [":" ext-mech-value]
ext-mech-name = token
ext-mech-value = token
; token is specified in RFC 4566

```

Figure 2: Syntax of the SDP Extensions

6. Examples

In the examples below, where an SDP line is too long to be displayed as a single line, a breaking character "\" indicates continuation in the following line. Note that this character is included for display purposes only. Implementations MUST write a single line without breaks.

6.1. Single PSTN Audio Stream

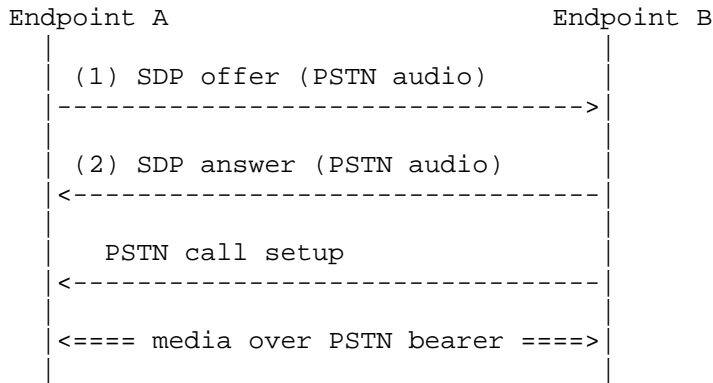


Figure 3: Basic Flow

Figure 3 shows a basic example that describes a single audio media stream over a circuit-switched bearer. Endpoint A generates an SDP offer, which is shown in Figure 4. The offer describes a PSTN circuit-switched bearer in the "m=" and "c=" line where it also indicates its international E.164 number format. Additionally, Endpoint A expresses that it can initiate the circuit-switched bearer or be the recipient of it in the "a=setup" attribute line. The SDP offer also includes correlation identifiers that this endpoint will insert in the Calling Party Number and/or User-User Information Element of the PSTN call setup if eventually this endpoint initiates the PSTN call. Endpoint A also includes "external" as one correlation mechanism, indicating that it can use the human user to perform correlation in case other mechanisms fail.

```

v=0
o=alice 2890844526 2890842807 IN IP4 192.0.2.5
s=
t=0 0
m=audio 9 PSTN -
c=PSTN E164 +441134960123
a=setup:actpass
a=connection:new
a=cs-correlation:callerid:+441134960123 \
  uuie:56A390F3D2B7310023 external

```

Figure 4: SDP Offer (1)

Endpoint B generates an SDP answer (Figure 5), describing a PSTN audio media on port 9 without information on the media subtype on the "m=" line. The "c=" line contains B's international E.164 number. In the "a=setup" line, Endpoint B indicates that it is willing to become the active endpoint when establishing the PSTN call, and it also includes the "a=cs-correlation" attribute line containing the values it is going to include in the Calling Party Number and User-User Information Element of the PSTN call establishment. Endpoint B is also able to perform correlation by external means, in case other correlation mechanisms fail.

```

v=0
o=- 2890973824 2890987289 IN IP4 192.0.2.7
s=
t=0 0
m=audio 9 PSTN -
c=PSTN E164 +441134960124
a=setup:active
a=connection:new
a=cs-correlation:callerid:+441134960124 \
  uuie:74B9027A869D7966A2 external

```

Figure 5: SDP Answer with Circuit-Switched Media

When Endpoint A receives the answer, it examines that B is willing to become the active endpoint when setting up the PSTN call. Endpoint A temporarily stores B's E.164 number and the User-User IE value of the "cs-correlation" attribute and waits for a circuit-switched bearer establishment.

Endpoint B initiates a circuit-switched bearer using whatever circuit-switched technology is available for it. The Called Party Number is set to A's number, and the Calling Party Number is set to B's own number. Endpoint B also sets the User-User Information Element value to the one contained in the SDP answer.

When Endpoint A receives the circuit-switched bearer establishment, it examines the UUIE and the Calling Party Number and, by comparing those received during the offer/answer exchange, determines that the call is related to the SDP session.

It may also be that neither the UUIE nor the Calling Party Number is received by the called party, or the format of the Calling Party Number is changed by the PSTN. Implementations may still accept such call establishment attempts as being related to the session that was established in the IP network. As it cannot be guaranteed that the values used for correlation are always passed intact through the network, they should be treated as additional hints that the circuit-switched bearer is actually related to the session.

6.2. Advanced SDP Example: Circuit-Switched Audio and Video Streams

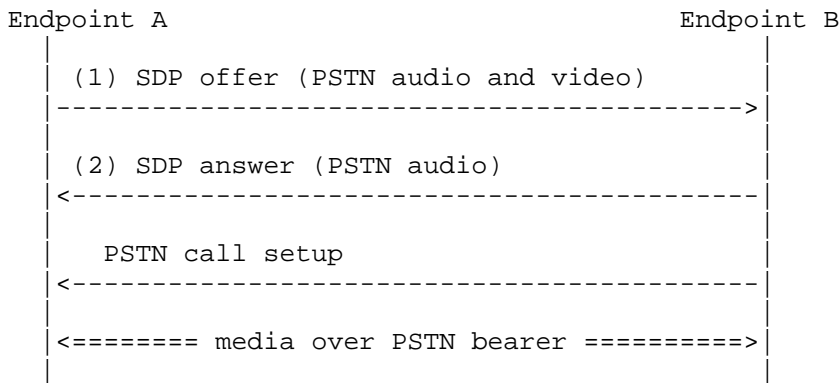


Figure 6: Circuit-Switched Audio and Video Streams

Figure 6 shows an example of negotiating audio and video media streams over circuit-switched bearers.


```

v=0
o=alice 2890844526 2890842807 IN IP4 192.0.2.5
s=
t=0 0
a=setup:actpass
a=connection:new
c=PSTN E164 +441134960123
m=audio 9 PSTN -
a=cs-correlation:dtmf:1234536
m=video 9 PSTN 34
a=rtpmap:34 H263/90000
a=cs-correlation:callerid:+441134960123

```

Figure 7: SDP Offer with Circuit-Switched Audio and Video (1)

Upon receiving the SDP offer described in Figure 7, Endpoint B rejects the video stream as the device does not currently support video, but it accepts the circuit-switched audio stream. As Endpoint A indicated that it is able to become either the active or passive party, Endpoint B gets to select which role it would like to take. Since the offer contained the international E.164 number of Endpoint A, Endpoint B decides that it becomes the active party in setting up the circuit-switched bearer. B includes a new value in the "dtmf" subfield of the "cs-correlation" attribute, which it is going to send as DTMF tones once the bearer setup is complete. The answer is described in Figure 8.

```

v=0
o=- 2890973824 2890987289 IN IP4 192.0.2.7
s=
t=0 0
a=setup:active
a=connection:new
c=PSTN E164 +441134960124
m=audio 9 PSTN -
a=cs-correlation:dtmf:654321
m=video 0 PSTN 34
a=cs-correlation:callerid:+441134960124

```

Figure 8: SDP Answer with Circuit-Switched Audio and Video (2)

7. Security Considerations

This document provides an extension to RFC 4566 [RFC4566] and RFC 3264 [RFC3264]. As such, the security considerations of those documents apply.

This memo provides mechanisms to agree on a correlation identifier or identifiers that are used to evaluate whether an incoming circuit-switched bearer is related to an ongoing session in the IP domain. If an attacker replicates the correlation identifier and establishes a call within the time window the receiving endpoint is expecting a call, the attacker may be able to hijack the circuit-switched bearer. These types of attacks are not specific to the mechanisms presented in this memo. For example, Caller ID spoofing is a well-known attack in the PSTN. Users are advised to use the same caution before revealing sensitive information as they would on any other phone call. Furthermore, users are advised that mechanisms that may be in use in the IP domain for securing the media, like Secure RTP (SRTP) [RFC3711], are not available in the CS domain.

For the purposes of establishing a circuit-switched bearer, the active endpoint needs to know the passive endpoint's phone number. Phone numbers are sensitive information, and some people may choose not to reveal their phone numbers when calling using supplementary services like Calling Line Identification Restriction (CLIR) in GSM. Implementations should take the caller's preferences regarding calling line identification into account if possible, by restricting the inclusion of the phone number in the SDP "c=" line if the caller has chosen to use CLIR. If this is not possible, implementations may present a prompt informing the user that their phone number may be transmitted to the other party.

As with IP addresses, if there is a desire to protect the SDP containing phone numbers carried in SIP, implementers are advised to follow the security mechanisms defined in [RFC3261].

It is possible that an attacker creates a circuit-switched session whereby the attacked endpoint should dial a circuit-switched number, perhaps even a premium-rate telephone number. To mitigate the consequences of this attack, endpoints MUST authenticate and trust remote endpoints users who try to remain passive in the circuit-switched connection establishment. It is RECOMMENDED that endpoints have local policies precluding the active establishment of circuit-switched connections to certain numbers (e.g., international, premium, and long distance). Additionally, it is strongly RECOMMENDED that the end user is asked for consent prior to the endpoint initiating a circuit-switched connection.

8. IANA Considerations

IANA has registered a number of SDP tokens according to the following data.

8.1. Registration of the New "cs-correlation" SDP Attribute

Contact: Miguel Garcia <miguel.a.garcia@ericsson.com>

Attribute name: cs-correlation

Long-form attribute name: PSTN Correlation Identifier

Type of attribute: media level only

Subject to charset: No

Description: This attribute provides the Correlation Identifier used in PSTN signaling

Appropriate values: see Section 5.2.3.1

Specification: RFC 7195

The IANA has created a subregistry for the "cs-correlation" attribute under the "Session Description Protocol (SDP) Parameters" registry. The initial values for the subregistry are presented in the following; IANA has registered these values accordingly:

Value of "cs-correlation" attribute	Reference	Description
callerid	RFC 7195	Caller ID
uuie	RFC 7195	User-User Information Element
dtmf	RFC 7195	Dual-Tone Multi-Frequency
external	RFC 7195	External

As per the terminology in [RFC5226], the registration policy for new values of the "cs-correlation" attribute is "Specification Required".

8.2. Registration of a New "nettype" Value

IANA has registered a new "nettype" in the "Session Description Protocol (SDP) Parameters" registry [IANA]. The registration data, according to RFC 4566 [RFC4566], is as follows.

Type	SDP Name	Reference
nettype	PSTN	RFC 7195

8.3. Registration of a New "addrtype" Value

IANA has registered a new "addrtype" in the "Session Description Protocol (SDP) Parameters" registry [IANA]. The registration data, according to RFC 4566 [RFC4566], is as follows.

Type	SDP Name	Reference
addrtype	E164	RFC 7195

Note: This document defines the "E164" addrtype in the context of the "PSTN" nettype only. RFC 3108 [RFC3108] also defines address type "E.164". This definition is distinct from the one defined by this memo and shall not be used with <nettype> "PSTN".

8.4. Registration of a New "proto" Value

IANA has registered a new "proto" in the "Session Description Protocol (SDP) Parameters" registry [IANA]. The registration data, according to RFC 4566 [RFC4566], is as follows.

Type	SDP Name	Reference
proto	PSTN	RFC 7195

The related "fmt" namespace reuses the conventions and payload type number defined for RTP/AVP. In this document, the RTP audio and video media types, when applied to PSTN circuit-switched bearers, represent merely an audio or video codec in its native format directly on top of a single PSTN bearer.

In some cases, the endpoint is not able to determine the list of available codecs for circuit-switched media streams. In this case, in order to be syntactically compliant with SDP [RFC4566], the endpoint MUST include a single dash ("-") in the <fmt> subfield.

9. Acknowledgments

The authors want to thank Paul Kyzivat, Flemming Andreasen, Thomas Belling, John Elwell, Jari Mutikainen, Miikka Poikselka, Jonathan Rosenberg, Ingemar Johansson, Christer Holmberg, Alf Heidermark, Tom Taylor, Thomas Belling, Keith Drage, and Andrew Allen for providing their insight and comments on this document.

10. References

10.1. Normative References

- [ITU.Q931.1998] International Telecommunications Union, "Digital Subscriber Signalling System No. 1 - ISDN User-Network Interface Layer 3 Specification for Basic Call Control", ITU-T Recommendation Q931, May 1998.
- [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.
- [RFC3966] Schulzrinne, H., "The tel URI for Telephone Numbers", RFC 3966, December 2004.
- [RFC4145] Yon, D. and G. Camarillo, "TCP-Based Media Transport in the Session Description Protocol (SDP)", RFC 4145, September 2005.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

10.2. Informative References

- [IANA] IANA, "Session Description Protocol (SDP) Parameters Registry", <<http://www.iana.org/assignments/sdp-parameters>>.
- [ITU.E164.2010] International Telecommunications Union, "The International Public Telecommunication Numbering Plan", ITU-T Recommendation E.164, 2010.
- [ITU.Q23.1988] International Telecommunications Union, "Technical features of push-button telephone sets", ITU-T Technical Recommendation Q.23, 1988.
- [RFC3108] Kumar, R. and M. Mostafa, "Conventions for the use of the Session Description Protocol (SDP) for ATM Bearer Connections", RFC 3108, May 2001.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, 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.
- [RFC3725] Rosenberg, J., Peterson, J., Schulzrinne, H., and G. Camarillo, "Best Current Practices for Third Party Call Control (3pcc) in the Session Initiation Protocol (SIP)", BCP 85, RFC 3725, April 2004.
- [RFC4975] Campbell, B., Mahy, R., and C. Jennings, "The Message Session Relay Protocol (MSRP)", RFC 4975, September 2007.
- [SIP-UUI] Johnston, A. and J. Rafferty, "A Mechanism for Transporting User to User Call Control Information in SIP", Work in Progress, April 2014.

[TS.24.008] 3GPP, "Mobile radio interface Layer 3 specification; Core network protocols; Stage 3", 3GPP TS 24.008 3.20.0, December 2005.

Authors' Addresses

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

EEmail: miguel.a.garcia@ericsson.com

Simo Veikkolainen
Nokia
P.O. Box 226
NOKIA GROUP, FI 00045
Finland

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

