

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2011

J. Peterson
NeuStar, Inc.
C. Jennings
Cisco Systems
March 14, 2011

The Advertisement/Proposal Model of Session Description
draft-peterson-sipcore-advprop-01

Abstract

In common SIP practice, a two-phase "offer/answer" exchange of session description documents negotiates preferences, capabilities and requested sessions. However, the structure of the session description greatly confuses the disambiguation of these elements and thus the clear characterization of sessions. The current work proposes an alternative to the offer/answer model which leverages pre-association between user agents to recast those two phases into a less ambiguous form: an Advertisement of capabilities and preferences which occurs in non-real-time before a session is ever requested, which is followed during session establishment by a unidirectional and complete Proposal of a session.

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 15, 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](#) 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.	Terminology	4
2.	Introduction	4
3.	Limitations of the offer/answer model	5
4.	Overview of the Advertisement/Proposal Model	6
5.	Benefits of the Advertisement/Proposal Model	8
6.	The Advertisement	9
7.	The Proposal	10
8.	Backward Compatibility	10
9.	Security Considerations	10
10.	IANA Considerations	11
11.	Informative References	11
	Authors' Addresses	12

1. Terminology

In this document, the key words "MAY", "MUST", "MUST NOT", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [\[RFC2119\]](#). This document additionally uses [\[RFC5226\]](#) language to describe IANA registrations. The terms "watcher" and "presentity" and related presence terms are defined in [RFC2778](#).

2. Introduction

The Session Description Protocol (SDP) [\[RFC4566\]](#) enjoys widespread use in the Internet for the description of sessions established by protocols such as the Session Initiation Protocol (SIP) [\[RFC3261\]](#). Most uses of SDP assume an offer/answer model as described in [\[RFC3264\]](#) for interactive session establishment, where one side sends an offer and the other side replies with an answer. However, many important matters relating to the management of media streams in SDP are inherently ambiguous in the offer/answer model. In particular, distinguishing the expression of capabilities from preferences from actual desired media streams has proven extremely difficult for more advanced applications of SDP. Moreover, as the offer and answer reflect only the media sinks of their respective creators, neither offers a complete picture of the session, and the degree to which the answer qualifies the offer admits of similar crippling ambiguities. These core difficulties have motivated numerous proposals to reform SDP, most of which radically increase its syntactical complexity, and thus have found mixed acceptance in the deployment community.

Perhaps the root of these problems lies not in SDP itself, however, but in application of SDP to the offer/answer model. SDP significantly predates the codification of the offer/answer model, and while [RFC3264](#) added a great deal of value by formalizing the previously tacit principles of using SDP in SIP, nonetheless those principles themselves have intrinsic limitations.

This document therefore proposes a reexamination of the fundamental requirements underlying the use of SDP in session-establishment protocols such as SIP. It recognizes, first of all, the desirability of conveying in a SIP INVITE message a document that describes the proposed session completely, hereafter a Proposal. In order to formulate such a Proposal, however, the originator of a session must have access to the capabilities and preferences of the target. Thus, in an alternative model entities willing to receive a Proposal might promulgate Advertisements to would-be proposers. Advertisements contain all of the information about the advertiser required by proposers to formulate a Proposal, including details such as the IP address and port where they hope to receive media. One way that

advertisements might deliver Advertisement to proposers is through presence.

Several previous efforts have recognized the need to provide capability and preference information prior to commencing the session establishment process. Notably, the work on expressing user agent capability within presence information [[RFC5196](#)] (which in turn builds upon the user agent capabilities work in [[RFC3840](#)]) provides one possible framework in which, with some enhancements, an Advertisement might be articulated. Work has also been done in [[RFC3407](#)] to express capability information with less ambiguity in SDP; by going down that path, an Advertisement might even be rendered in SDP. Rendering a Proposal in SDP would probably return that format to something like its original purpose. This document, however, proposes no specific syntactical format for Advertisements or Proposals - it merely describes the semantics of these objects as a set of requirements that future format proposals could meet.

While the work in this document is most immediately relevant to SIP, and discusses its applicability to SIP throughout, the underlying model may have further applicability to other protocol using session descriptions in a similar manner.

3. Limitations of the offer/answer model

There are several limitations of the offer/answer model that cause issues with the operation of SIP. Many have been the target of prior work intended to repair the problem by adding additional syntactical structure to SDP. However, these problems arise from the semantics of SDP in the offer/answer context - not only because an offer represents an incomplete picture of the session, and at that a picture filled with conditional elements, but because an answer does not describe a session completely either.

In the traditional offer/answer model, in order to maximize the chances of the negotiation process arriving at an acceptable session, an offerer must propose the largest set of session features possible. This includes the various combinations of RTP profiles, transport addresses, security properties and so forth that would be acceptable to the offerer. Disambiguating the semantics of these features - distinguishing, for example, the session features the offerer would like to invoke concurrently as opposed to those offered as alternatives - has required the addition of significant complexity to SDP. Because SDP's underlying syntax is not designed to carry structured or hierarchical data, the resulting documents can be large and confusing. Indeed, the term "offer" no longer seems to fit what that sort of document describes - it is a menu rather than a

selection. The resulting negotiation process, in which an answerer evaluates the alternatives described in the offer and arrives at a response, has become complex enough to call into question the wisdom of executing it in real-time, such as during the alerting phase of a telephone call.

Architecturally, an offerer moreover cannot anticipate what entity, or entities, their offer will reach, due to the potential in SIP for retargeting and/or forking. The fact that an answer typically arrives in a message confirming session establishment (the 200 OK), yet that it can originate from a surprising source, and make surprising stipulations, is a weakness in offer/answer. The offer/answer model places all of the onus, and the authority, on the answerer to complete a session within the very broad constraints that can be specified by the offerer. It is unclear why this power is better invested in the answerer than the offerer, but ideally, whichever side develops a finalizes proposal for the session should pass it to the other for approval, rather than assuming it will be acceptable and turning on the session. That somewhat philosophical point is not purely theoretical - it is one of the major drivers for classic problems in SIP such as early media, where due to forking several answerers may unknowingly operate on the same offer at once.

The fact that neither an offer nor an answer alone contains a complete picture of the session also compels intermediaries interested in SDP, for example application-layer gateways assisting with NAT traversal, to hear and modify, potentially, both requests and responses (e.g. a 200 carrying SDP) during session establishment. Unlike requests, responses cannot be rejected at the SIP layer, for security reasons or any other reasons, which effectively prevents endpoints from cleanly negotiating away from any unwanted "suggestions" made by intermediaries.

Some of this problem space in SDP is addressed by ongoing work of the MMUSIC working group on capability negotiation in SDP. The Advertisement/Proposal model differs from the capneg work in so far as capneg retains the offer/answer model and provides sufficient syntactical equipment in SDP to differentiate the various functions of expressing capability, proposing the properties of a session and confirming those proposals. However, despite their differences the approaches here and in are not mutually exclusive; in order to express Advertisements or Proposals in SDP, something very much like the attributes defined in the capneg work is required.

4. Overview of the Advertisement/Proposal Model

An Advertisement is a document detailing the parameters of sessions

that a user is willing to accept. A user agent who aspires to initiate a session uses an Advertisement as the basis for a Proposal, a complete description of a session, by selecting a common set of desired capabilities from the Advertisement. This Proposal is then carried in a session initiation message such as an INVITE whose recipient will either accept or reject it in its entirety. A recipient who rejects a proposal may search for an Advertisement from the proposer, and in turn make their own Proposal based on it; this counter-Proposal has no necessary relationship with the rejected Proposal. Unlike the offer/answer model, in the Advertisement/Proposal model a Proposal is a complete description of a session - it does not specify preferences or capabilities, instead it proposes the media types, sources and sinks that shall be used by both endpoints in a session.

Advertisements incorporate the session-related capabilities of the user agent(s) they describe, as modified by user policy. A user agent may elect to share a customized Advertisement with a would-be proposer, and thus depending on the proposer, user policy may dictate an Advertisement offer different media types or what have. Advertisements may contain various proposer-specific information that is of use to endpoints prior to setting up a session, including information used in connectivity checks or keying media-layer security.

Promulgating Advertisements requires some administrative pre-association between user agents as well as a means of requesting and transmitting the Advertisement itself outside context of a session. The existing availability of presence in many systems that establish multimedia sessions permits the sharing of capability and preference expression, as per the work in [RFC5196](#). Presence establishes a channel for communicating availability or capability information between presence user agents which can also share additional information relevant to a later stage of session establishment. Many presence systems moreover have the capability to supply customized presence data to particular watchers, a desirable quality for the Advertisement/Proposal model. An Advertisement may be shared as a component of an existing presence document, including a Presence Information Data Format (PIDF) document, or any other presence format. Note that the use of presence does not necessarily entail the maintenance of long-term subscriptions, and might rely in some cases on fetch operations conducted immediately prior to session initiation.

The Session Description Protocol provides numerous other expressions of capability (for example, the language of speakers in audio sessions) which are not considered here as requirements for adaptation to the Advertisement/Proposal model, though they may be

incorporated into applications if desired.

5. Benefits of the Advertisement/Proposal Model

In addition to addressing some existing problems in common SIP deployments, the use of the advertisement/proposal model also creates opportunities for new features. Most of these new features rely on the availability of the Advertisement prior to the time of session establishment, which allows user agents or applications to consider or act upon this information, in non-real-time.

The simplest application of this model is to allow the proposer the opportunity to browse the capabilities of an advertiser prior to making a session initiation attempt. In the traditional offer/answer model, the user must express any preferences to his user agent prior to attempting to initiate a session, and only when the INVITE is received, and one is on the proverbial clock, does the answerer begin the process of determining how the offerer's capabilities and preferences match to its own. While some potential methods in SIP of providing this "browsing" user experience prior to initiating a session have already been specified, as detailed above, those methods might turn out to be either partially redundant with the negotiation of the offer/answer phase or even contradictory to it - the Advertisement/Proposal model makes this user experience part and parcel of the session initiation process.

Another example is the use of an Advertisement to predict whether or not connectivity will be possible behind endpoints. Due to various network-layer obstacles (include NATs or firewalls), one SIP user agent may not be capable of forming a media session with another endpoint, despite their ability to rendez-vous successfully at the SIP layer. In an Advertisement, a user agent might express enough information to allow a would-be proposer to perform ICE-style checks prior to initiating a session, and thus to learn whether or not a potential target would be reachable. This information could be rendered to the proposer's user as a sort of presence information, for example ("red" signifying a media session cannot be formed with the target, "green" suggesting that it could). This is just one example of a whole class of information useful during session establishment that could be shared during the Advertisement phase - encryption keys for the bodies of SIP requests are another bit of information that might be difficult to acquire otherwise.

Advertisements are also useful in bunches. For a case where a third-party call control system is attempting to figure out how two or more parties could share some sort of common session, acquiring Advertisements from all parties ahead of time greatly facilitates the

process of assembling Proposals that would be acceptable to the group. Conducting this process during session establishment itself (by inviting each party serially to a session, for example) could yield far less optimal results.

Any application that consumes Advertisement data prior to session establishment in order to make predictions about the session faces a trade-off resulting from the potential for that data to become stale. For an application performing a pre-session connectivity check, for example, changes in both the state of the present entity and the state of the network could render the results of a check stale at any time; however, the application simply can't perform connectivity checks constantly to mitigate this difficulty. Moreover, since that results of that connectivity check are rendered to a human, nor do the checks even need to be performed if applications are aware that no human currently requires them. Setting sensible thresholds for these sorts of pre-session application activities is a subject for future work, as are the scalability concerns arising from a watcher who follows Advertisements from an enormous number of present entities.

6. The Advertisement

This document articulates the high-level qualities of the Advertisement/Proposal model. Document formats for Advertisements are therefore out of scope; it is possible that even the Session Description Protocol could be adapted to express an Advertisement or Proposal, provided that it can articulate the necessary elements unambiguously.

Purely for exemplary purposes, an Advertisement might contain the following elements:

- o Codecs and payload types supported by the originator of the advertisement
- o Protocol support (RTP, DCCP, TCP, UDP, HTTP, etc)
- o IP Addresses and ports where the originator may send and receive media
- o Layer coding schemes, redundant encoding and forward error correction
- o Bit-rate limits, frame rates, display capabilities
- o Security parameters of the session, including any information needed to key media
- o Content meta-data and information about repositories (for file-transfer sessions)

7. The Proposal

The Proposal is a document which describes the entire state of the desired session. The construction of the Proposal may be informed by the Advertisement, but the interpretation of the Proposal is unambiguous without the Advertisement. Again, this section merely sketches the high-level properties of a Proposal as an indication of the direction of future work.

A Proposal may contain:

- o A specific five tuple for each media session proposed by its creator
- o A specific codec, payload type and SSRC for said media where applicable
- o Codec parameters (bit rates, picture sizes, etc)

8. Backward Compatibility

The use of the Advertisement/Proposal model integrates nicely with an existing deployed base of offer/answer model compliance, for the simple reason that one cannot send a Proposal without an Advertisement, and any entity that promulgates Advertisements supports Proposals. For deployments where both models are in simultaneous use, baseline mandatory compliance with offer/answer would be required, though when support for Advertisement/Proposal is detected (via the presence of an Advertisement), it could be then invoked when preferred.

At a syntactical level, any path to implementing an Advertisement/Proposal model requires disambiguating Proposals from offers or answers. For Proposals written in SDP, perhaps a separate MIME media type (e.g., "application/sdp-prop") might be defined which could in turn be negotiated in the usual SIP matter, via OPTIONS requests, 415 response codes, and so forth. Any new document format for Proposals could of course define a new MIME type and operate in the same fashion.

9. Security Considerations

The Advertisement/Proposal model assumes the secure distribution of Advertisements from a presentity to one or more watchers with whom the presentity might want to share a session. An Advertisement could impart secrets for keying media security, for example, which are intended only for a particular watcher. The secure distribution of presence documents is therefore essential to this model. Fortunately, this problem is shared with most other sensitive

presence data, and thus this document need only point to the existing guidance for securing presence documents in [RFC3863](#).

10. IANA Considerations

This document contains no considerations for the IANA.

11. Informative References

- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [RFC3265] Roach, A., "Session Initiation Protocol (SIP)-Specific Event Notification", [RFC 3265](#), June 2002.
- [RFC3325] Jennings, C., Peterson, J., and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks", [RFC 3325](#), November 2002.
- [RFC3407] Andreassen, F., "Session Description Protocol (SDP) Simple Capability Declaration", [RFC 3407](#), October 2002.
- [RFC3840] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", [RFC 3840](#), August 2004.
- [RFC3863] Sugano, H., Fujimoto, S., Klyne, G., Bateman, A., Carr, W., and J. Peterson, "Presence Information Data Format (PIDF)", [RFC 3863](#), August 2004.
- [RFC3968] Camarillo, G., "The Internet Assigned Number Authority (IANA) Header Field Parameter Registry for the Session Initiation Protocol (SIP)", [BCP 98](#), [RFC 3968](#),

December 2004.

- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.
- [RFC5111] Aboba, B. and L. Dondeti, "Experiment in Exploratory Group Formation within the Internet Engineering Task Force (IETF)", [RFC 5111](#), January 2008.
- [RFC5196] Lonnfors, M. and K. Kiss, "Session Initiation Protocol (SIP) User Agent Capability Extension to Presence Information Data Format (PIDF)", [RFC 5196](#), September 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

Authors' Addresses

Jon Peterson
NeuStar, Inc.

Email: jon.peterson@neustar.biz

Cullen Jennings
Cisco Systems

Email: fluffy@cisco.com

