

Javascript Session Establishment Protocol (JSEP)

Justin Uberti (juberti@google.com)
Version 0.2
December 16, 2011

Table of Contents

[Background](#)

[Proposal](#)

[Semantics and Syntax](#)

[Signaling Model](#)

[Session Descriptions](#)

[Session Description Format](#)

[Separation of Signaling and ICE State Machines](#)

[ICE Candidate Format](#)

[Typical Media Setup](#)

[Initiating the Session](#)

[Generating An Offer](#)

[Applying the Offer](#)

[Initiating ICE](#)

[Serializing the Offer and Candidates](#)

[Receiving the Session](#)

[Receiving the Offer](#)

[Initiating ICE](#)

[Handling ICE Messages](#)

[Generating the Answer](#)

[Applying the Answer](#)

[Serializing the Answer](#)

[Completing the Session](#)

[Receiving the Answer](#)

[Updates to the Session](#)

[Proposed WebRTC API changes](#)

[PeerConnection API](#)

[MediaHints](#)

[Example API Flows](#)

[Call using ROAP](#)

[Call using XMPP](#)

[Adding video to a call, using XMPP](#)

[Call using SIP](#)

[Handling early media \(e.g. 1-800-FEDEX\), using SIP](#)

[Security Considerations](#)

[IANA Considerations](#)

[Acknowledgements](#)

[Open Issues](#)

[References](#)

[Normative References](#)

[Informative References](#)

Background

The general thinking behind WebRTC call setup has been to fully specify and control the media plane, but to leave the signaling plane up to the application as much as possible. The rationale is that different applications may prefer to use different protocols, such as the existing SIP or Jingle call signaling protocols, or something custom to the particular application, perhaps for a novel use case. In this approach, the key information that needs to be exchanged is the multimedia session description, which specifies the necessary transport and media configuration information necessary to establish the media plane.

The original spec for WebRTC attempted to implement this protocol-agnostic signaling by providing a mechanism to exchange session descriptions in the form of SDP blobs. Upon starting a session, the browser would generate a SDP blob, which would be passed to the application for transport over its preferred signaling protocol. On the remote side, this blob would be passed into the browser from the application, and the browser would then generate a blob of its own in response. Upon transmission back to the initiator, this blob would be plugged into their browser, and the handshake would be complete.

Experimentation with this mechanism turned up several shortcomings, which generally stemmed from there being insufficient context at the browser to fully determine the meaning of a SDP blob. For example, determining whether a blob is an offer or an answer, or differentiating a new offer from a retransmit.

The ROAP proposal, specified in <http://tools.ietf.org/html/draft-jennings-rtcweb-signaling-01>, attempted to resolve these issues by providing additional structure in the messaging - in essence, to create a generic signaling protocol that specifies how the browser signaling state machine should operate. However, even though the protocol is abstracted, the state machine forces a least-common-denominator approach on the signaling interactions. For example, in Jingle, the call initiator can provide additional ICE candidates even after the initial offer has been sent, which allows the offer to be sent immediately for quicker call startup. However, in the browser state machine, there is no notion of sending an updated offer before the initial offer has been responded to, rendering this functionality impossible.

While specific concerns like this could be addressed by modifying the generic protocol, others would doubtless be discovered later. The main reason this mechanism is inflexible is because it bakes too much behavior into the browser. Since the browser generates the session descriptions on its own, and controls the signaling state machine, modification of the session descriptions or use of alternate state machines becomes difficult or impossible.

Proposal

To resolve these issues, this document proposes the Javascript Session Establishment Protocol (JSEP) that pulls the signaling state machine out of the browser and into Javascript. This mechanism effectively removes the browser almost completely from the core signaling flow; the only interface needed is a way for the application to pass in the local and remote session descriptions negotiated by whatever signaling mechanism is used, and a way to interact with the

ICE state machine.

JSEP's handling of session descriptions is simple and straightforward. Whenever an offer/answer exchange is needed, the initiating side creates an offer by calling the *createOffer(hints)* API on *PeerConnection*. The application can do massaging of that offer, if it wants to, and then sends it off to the remote side over its preferred signaling mechanism (e.g. *WebSockets*). Upon receipt of that offer, the remote party calls the *createAnswer(offer, hints)* to generate an appropriate answer, and sends that back to the initiator. When the offer and answer each have been deemed by the signaling protocol to be "live", the application will apply them to the *PeerConnection* via the *setLocalDescription(type, desc)* and *setRemoteDescription(type, desc)* APIs. This process can be repeated for additional offer/answer exchanges.

Regarding ICE, in this approach we decouple the ICE state machine from the overall signaling state machine; the ICE state machine must remain in the browser, given that only the browser has the necessary knowledge of candidates and other transport info. While transport has typically been lumped in with session descriptions, performing this separation it provides additional flexibility. In protocols that decouple session descriptions from transport, such as *Jingle*, the transport information can be sent separately; in protocols that don't, such as *SIP*, the information can be easily aggregated and recombined. Sending transport information separately can allow for faster ICE and DTLS startup, since the necessary roundtrips can occur while waiting for the remote side to accept the session.

The JSEP approach does come with a minor downside. As the application now is responsible for driving the signaling state machine, slightly more application code is necessary to perform call setup; the application must call the right APIs at the right times, and convert the session descriptions and ICE information into the defined messages of its chosen signaling protocol, instead of simply forwarding the messages emitted from the browser.

One way to mitigate this is to provide a Javascript library that hides this complexity from the developer, which would implement the state machine and serialization of the desired signaling protocol. For example, this library could convert easily adapt the JSEP API into the exact *ROAP* API, thereby implementing the *ROAP* signaling protocol. Such a library could of course also implement other popular signaling protocols, including *SIP* or *Jingle*. In this fashion we can enable greater control for the experienced developer without forcing any additional complexity on the novice developer.

Other approaches

Another approach that was considered for JSEP was to move the mechanism for generating offers and answers out of the browser as well. This approach would add a *getCapabilities* API which would provide the application with the information it needed in order to generate session descriptions. This increases the amount of work that the application needs to do; it needs to know how to generate session descriptions from capabilities, and especially how to

generate the correct answer from an arbitrary offer and available capabilities. While this could certainly be addressed by using a library like the one mentioned above, some experimentation also indicates that coming up with a sufficiently complete `getCapabilities` API is a nontrivial undertaking. Nevertheless, if we wanted to go down this road, JSEP makes it significantly easier; if a `getCapabilities` API is added in the future, the application can generate session descriptions accordingly and pass those to the `setLocalDescription`/`setRemoteDescription` APIs added by JSEP. (Even with JSEP, an application could still perform its own browser fingerprinting and generate approximate session descriptions as a result.)

Note also that while JSEP transfers more control to Javascript, it is not intended to be an example of a "low-level" API. The generally argument against a low-level API is that there are too many necessary API points, and they can be called in any order, leading to something that is hard to specify and test. In the approach proposed here, control is performed via session descriptions; this requires only a few APIs to handle these descriptions, and they are evaluated in a specific fashion, which reduces the number of possible states and interactions.

Semantics and Syntax

Signaling Model

JSEP does not specify a particular signaling model or state machine, other than the generic need to exchange RFC 3264-esque offers and answers in order for both sides of the session to know how to conduct the session. JSEP provides mechanisms to create offers and answers, as well as to apply them to a `PeerConnection`. However, the actual mechanism by which these offers and answers are communicated to the remote side, including addressing, retransmission, forking, and glare handling, is left entirely up to the application.

Session Descriptions

In order to establish the media plane, `PeerConnection` needs specific parameters to indicate what to transmit to the remote side, as well as how to handle the media that is received. These parameters are determined by the exchange of session descriptions in offers and answers, and there are certain details to this process that must be handled in the JSEP APIs.

Whether a session description was sent or received affects the meaning of that description. For example, the list of codecs sent to a remote party indicates what the local side is willing to decode, and what the remote party should send. Not all parameters follow this rule; the SRTP parameters sent to a remote party indicate what the local side will use to encrypt, and thereby how the remote party should expect to receive.

In addition, various RFCs put different conditions on the format of offers versus answers. For example, a offer may propose multiple SRTP configurations, but an answer may only contain a single SRTP configuration.

Lastly, while the exact media parameters are only known only after a offer and an answer have been exchanged, it is possible for the offerer to receive media after they have sent an offer and before they have received an answer. To properly process incoming media in this case, the offerer's media handler must be aware of the details of the offerer before the answer arrives.

Therefore, in order to handle session descriptions properly, PeerConnection needs

1. To know if a session description pertains to the local or remote side.
2. To know If a session description is an offer or an answer.
3. To allow the offer to be specified independently of the answer.

JSEP addresses this by adding both a `setLocalDescription` and a `setRemoteDescription` method, and both these methods take as a first parameter either the value `SDP_OFFER`, or `SDP_ANSWER`. This satisfies the requirements listed above for both the offerer, who first calls `setLocalDescription(SDP_OFFER, sdp)` and then later `setRemoteDescription(SDP_ANSWER, sdp)`, as well as for the answerer, who first calls `setRemoteDescription(SDP_OFFER, sdp)` and then later `setLocalDescription(SDP_ANSWER, sdp)`.

While it could be possible to implicitly determine the value of the offer/answer argument inside of PeerConnection, requiring it to be specified explicitly seems substantially more robust, allowing invalid combinations (i.e. an answer before an offer) to generate an appropriate error.

Session Description Format

In the current WebRTC specification, session descriptions are formatted as SDP messages. While this format is not optimal for manipulation from Javascript, it is widely accepted, and frequently updated with new features. Any alternate encoding of session descriptions would have to keep pace with the changes to SDP, at least until the time that this new encoding eclipsed SDP in popularity. As a result, JSEP continues to use SDP as the format for its session descriptions.

Note that this decision does not prevent us from using an alternate encoding in the future; if we were able to agree on a JSON format for session descriptions, surely it would be easy to add a switch to PeerConnection to tell it to generate/expect JSON.

Separation of Signaling and ICE State Machines

Previously, PeerConnection operated two state machines, referred to in the spec as an "ICE Agent", which handles the establishment of peer-to-peer connectivity, and an "SDP Agent", which handles the state of the offer-answer signaling. The states of these state machines were exposed through the `iceState` and `sdpState` attributes on PeerConnection, with an additional `readyState` attribute that reflected the high-level state of the PeerConnection.

JSEP does away with the SDP Agent within the browser; this functionality is now controlled directly by the application, which uses the `setLocalDescription` and `setRemoteDescription` APIs to tell PeerConnection what SDP has been negotiated. The ICE Agent remains in the browser,

as it still needs to perform gathering of candidates, connectivity checking, and related ICE functionality.

The net effect of this is that *sdpState* goes away, and *processSignalingMessage* becomes *processIceMessage*, which now specifically handles incoming ICE candidates. To allow the application to control exactly when it wants to start ICE negotiation (e.g. either on receipt of the call, or only after accepting the call), a *connect* method has been added.

ICE Candidate Format

As with session descriptions, we choose to use SDP's representation of ICE candidates in this API, specifically *processIceMessage*. For example:

```
a=candidate:1 1 UDP 1694498815 66.77.88.99 10000 typ host
```

While a JSON encoding could have been used, it is probably simplest to stay consistent and use the SDP representation, given the ease with which this string can be parsed.

Currently, *a=candidate* lines are the only thing that are exchanged.

Typical Media Setup

The example here shows a typical call setup using the JSEP model. We assume the following architecture in this example, where UA is synonymous with "browser", and JS is synonymous with "web application"

```
OffererUA <-> OffererJS <-> WebServer <-> AnswererJS <-> AnswererUA
```

Initiating the Session

The initiator creates a *PeerConnection*, installs its *IceCallback*, and adds the desired *MediaStreams* (presumably obtained via *getUserMedia*). The *PeerConnection* is in the *NEW* state.

```
OffererJS->OffererUA: var pc = new PeerConnection(config, iceCallback);  
OffererJS->OffererUA: pc.addStream(stream);
```

Generating An Offer

The initiator then creates a session description to offer to the callee. This description includes the codecs and other necessary session parameters, as well as information about each of the streams that has been added (e.g. *SSRC*, *CNAME*, etc.) The created description includes all parameters that the offerer's UA supports; if the initiator wants to influence the created offer, they can pass in a *MediaHints* object to *createOffer* that allows for customization (e.g. if the initiator wants to receive but not send video). The initiator can also directly manipulate the created session description as well, perhaps if it wants to change the priority of the offered codecs.

```
OffererJS->OffererUA: var offer = pc.createOffer(null);
```

Applying the Offer

The initiator then instructs the PeerConnection to use this offer as the local description for this session, i.e. what codecs it will use for received media, what SRTP keys it will use for sending media (if using SDES), etc. In order that the UA handle the description properly, the initiator marks it as an offer when calling `setLocalDescription`; this indicates to the UA that multiple capabilities have been offered, but this set may be pared back later, when the answer arrives.

```
OffererJS->OffererUA: pc.setLocalDescription(SDP_OFFER, offer);
```

Initiating ICE

The initiator can now start the ICE process of candidate generation and connectivity checking. This results in callbacks to the application's `IceCallback`.

```
OffererJS->OffererUA: pc.connect();
OffererUA->OffererJS: iceCallback(candidates);
```

Serializing the Offer and Candidates

At this point, the offerer is ready to send its offer to the callee using its preferred signaling protocol. Depending on the protocol, it can either send the initial session description first, and then "trickle" the ICE candidates as they are given to the application, or it can wait for all the ICE candidates to be collected, and then send the offer and list of candidates all at once.

Receiving the Session

Through the chosen signaling protocol, the recipient is notified of an incoming session request. It creates a PeerConnection, and installs its own `IceCallback`.

```
AnswererJS->AnswererUA: var pc = new PeerConnection(config, iceCallback);
```

Receiving the Offer

The recipient converts the received offer from its signaling protocol into SDP format, and supplies it to its PeerConnection, again marking it as an offer. As a remote description, the offer indicates what codecs the remote side wants to use for receiving, as well as what SRTP keys it will use for sending. The setting of the remote description causes callbacks to be issued, informing the application of what kinds of streams are present in the offer.

```
AnswererJS->AnswererUA: pc.setRemoteDescription(SDP_OFFER, offer);
AnswererUA->AnswererJS: onAddStream(stream);
```

Initiating ICE

The recipient then starts its own ICE state machine, to allow connectivity to be established as quickly as possible.

```
AnswererJS->AnswererUA: pc.connect();
AnswererUA->AnswererJS: iceCallback(candidates);
```

Handling ICE Messages

If ICE candidates from the remote site were included in the offer, the ICE Agent will automatically start trying to use them. Otherwise, if ICE candidates are sent separately, they are passed into the PeerConnection when they arrive.

```
AnswererJS->AnswererUA: pc.processIceMessage(candidates);
```

Generating the Answer

Once the recipient has decided to accept the session, it generates an answer session description. This process performs the appropriate intersection of codecs and other parameters to generate the correct answer. As with the offer, MediaHints can be provided to influence the answer that is generated, and/or the application can post-process the answer manually.

```
AnswererJS->AnswererUA: pc.createAnswer(offer, null);
```

Applying the Answer

The recipient then instructs the PeerConnection to use the answer as its local description for this session, i.e. what codecs it will use to receive media, etc. It also marks the description as an answer, which tells the UA that these parameters are final. This causes the PeerConnection to move to the ACTIVE state, and transmission of media by the answerer to start.

```
AnswererJS->AnswererUA: pc.setLocalDescription(SDP_ANSWER, answer);
AnswererUA->OffererUA: <media>
```

Serializing the Answer

As with the offer, the answer (with or without candidates) is now converted to the desired signaling format and sent to the initiator.

Completing the Session

Receiving the Answer

The initiator converts the answer from the signaling protocol and applies it as the remote description, marking it as an answer. This causes the PeerConnection to move to the ACTIVE state, and transmission of media by the offerer to start.

```
OffererJS->OffererUA: pc.setRemoteDescription(SDP_ANSWER, answer);
OffererUA->AnswererUA: <media>
```

Updates to the Session

Updates to the session are handled with a new offer/answer exchange. However, since media will already be flowing at this point, the new session descriptions should not be passed into PeerConnection until the changes have been accepted by the remote side, to prevent sending media that the remote side is not prepared to handle.

Note also that in an update scenario, the roles may be reversed, i.e. the update offerer can be different than the original offerer.

Proposed WebRTC API changes

PeerConnection API

The text below indicates the recommended changes to the PeerConnection API to implement the JSEP functionality. Methods and variables in bold are new/proposed; methods in gray are untouched functionality; methods in red/strikethrough have been removed in this proposal.

```
[Constructor (in DOMString configuration, in IceCallback iceCallback)]
interface PeerConnection {
```



```

// creates a blob containing the SDP to be provided as an offer.
DOMString createOffer (MediaHints hints);
// creates a blob containing the SDP to be provided as an answer.
DOMString createAnswer (DOMString offer, MediaHints hints);
// sets the local session description (except ICE info)
void setLocalDescription (Action action, DOMString sdp);
// sets the remote session description (except ICE info)
void setRemoteDescription (Action action, DOMString sdp);
// returns the current local session description
readonly DOMString localDescription;
// returns the current remote session description
readonly DOMString remoteDescription;
void processSignalingMessage (DOMString message);
const unsigned short NEW = 0; // initial state
const unsigned short ACTIVE = 1; // local+remote desc set, live
const unsigned short CLOSED = 2; // ended state
readonly attribute unsigned short readyState;
// starts ICE connection/handshaking
void connect();
// processes received ICE information
void processIceMessage (DOMString message);
const unsigned short ICE_GATHERING = 0x100;
const unsigned short ICE_WAITING = 0x200;
const unsigned short ICE_CHECKING = 0x300;
const unsigned short ICE_CONNECTED = 0x400;
const unsigned short ICE_COMPLETED = 0x500;
const unsigned short ICE_FAILED = 0x600;
const unsigned short ICE_CLOSED = 0x700;
readonly attribute unsigned short iceState;
const unsigned short SDP_IDLE = 0x1000;
const unsigned short SDP_WAITING = 0x2000;
const unsigned short SDP_GLARE = 0x3000;
readonly attribute unsigned short sdpState;
void addStream (MediaStream stream, MediaStreamHints hints);
void removeStream (MediaStream stream);
readonly attribute MediaStream[] localStreams;
readonly attribute MediaStream[] remoteStreams;
void close ();
[...] };

```

MediaHints

MediaHints is an object that can be passed into *createOffer* or *createAnswer* to affect the type of offer/answer that is generated. The following properties can be set on MediaHints:

- `has_audio`: boolean // whether we want to receive audio
// defaults to true if we have audio streams, else false
- `has_video`: boolean // whether we want to receive video
// defaults to true if we have video streams, else false

As an example, `MediaHints` could be used to create a session that transmits only audio, but is able to receive video from the remote side, by forcing the inclusion of a `m=video` line even when no video sources are provided.

Example API Flows

Below are several sample flows for the new `PeerConnection` and library APIs, demonstrating when the various APIs are called in different situations and with various transport protocols.

Call using ROAP

```
// Call is initiated toward Answerer
OffererJS->OffererUA:   pc = new PeerConnection();
OffererJS->OffererUA:   pc.addStream(localStream, null);
OffererJS->OffererUA:   offer = pc.createOffer(null);
OffererJS->OffererUA:   peer.setLocalDescription(SDP_OFFER, offer);

OffererJS->OffererUA:   pc.connect();
OffererUA->OffererJS:   iceCallback(cand); these are added to |offer|
OffererJS->AnswererJS:  {"type":"OFFER", "sdp":"<offer>"}

// OFFER arrives at Answerer
AnswererJS->AnswererUA: pc = new PeerConnection();
AnswererJS->AnswererUA: pc.setRemoteDescription(SDP_OFFER, msg.sdp);
AnswererUA->AnswererJS: onaddstream(remoteStream);
AnswererJS->AnswererUA: pc.processIceMessage(GetCandidates(msg.sdp));
AnswererJS->AnswererUA: pc.connect();
AnswererUA->OffererUA:  iceCallback(cand);

// Answerer accepts call
AnswererJS->AnswererUA: peer.addStream(localStream, null);
AnswererJS->AnswererUA: answer = peer.createAnswer(msg.offer, null);
AnswererJS: candidates are added to |answer|
AnswererJS->AnswererUA: peer.setLocalDescription(SDP_ANSWER, answer);
AnswererJS->OffererJS:  {"type":"ANSWER","sdp":"<answer>"}

// ANSWER arrives at Offerer
OffererJS->OffererUA:   peer.setRemoteDescription(ANSWER, answer);
OffererUA->OffererJS:   onaddstream(remoteStream);
OffererJS->OffererUA:   pc.processIceMessage(GetCandidates(msg.sdp));

// ICE Completes (at Answerer)
AnswererUA->AnswererJS: onopen();
AnswererUA->OffererUA:  Media

// ICE Completes (at Offerer)
OffererUA->OffererJS:   onopen();
OffererJS->AnswererJS:  {"type":"OK" }
OffererUA->AnswererUA:  Media
```

Call using XMPP

```
// Call is initiated toward Answerer
OffererJS->OffererUA: pc = new PeerConnection();
OffererJS->OffererUA: pc.addStream(localStream, null);
OffererJS->OffererUA: offer = pc.createOffer(null);
OffererJS->OffererUA: peer.setLocalDescription(SDP_OFFER, offer);
OffererJS: xmpp = createSessionInitiate(offer);
OffererJS->AnswererJS: <jingle action="session-initiate"/>

OffererJS->OffererUA: pc.connect();
OffererUA->OffererJS: iceCallback(cand);
OffererJS: createTransportInfo(cand);
OffererJS->AnswererJS: <jingle action="transport-info"/>

// session-initiate arrives at Answerer
AnswererJS->AnswererUA: pc = new PeerConnection();
AnswererJS: offer = parseSessionInitiate(xmpp);
AnswererJS->AnswererUA: pc.setRemoteDescription(SDP_OFFER, offer);
AnswererUA->AnswererJS: onaddstream(remoteStream);

// transport-infos arrive at Answerer
AnswererJS->AnswererUA: candidates = parseTransportInfo(xmpp);
AnswererJS->AnswererUA: pc.processIceMessage(candidates);
AnswererJS->AnswererUA: pc.connect();
AnswererUA->AnswererJS: iceCallback(cand)
AnswererJS: createTransportInfo(cand);
AnswererJS->OffererJS: <jingle action="transport-info"/>

// transport-infos arrive at Offerer
OffererJS->OffererUA: candidates = parseTransportInfo(xmpp);
OffererJS->OffererUA: pc.processIceMessage(candidates);

// Answerer accepts call
AnswererJS->AnswererUA: peer.addStream(localStream, null);
AnswererJS->AnswererUA: answer = peer.createAnswer(offer, null);
AnswererJS: xmpp = createSessionAccept(answer);
AnswererJS->AnswererUA: peer.setLocalDescription(SDP_ANSWER, answer);
AnswererJS->OffererJS: <jingle action="session-accept"/>

// session-accept arrives at Offerer
OffererJS: answer = parseSessionAccept(xmpp);
OffererJS->OffererUA: peer.setRemoteDescription(ANSWER, answer);
OffererUA->OffererJS: onaddstream(remoteStream);
OffererJS->OffererUA: pc.processIceMessage(GetCandidates(msg.sdp));

// ICE Completes (at Answerer)
AnswererUA->AnswererJS: onopen();
AnswererUA->OffererUA: Media

// ICE Completes (at Offerer)
OffererUA->OffererJS: onopen();
OffererUA->AnswererUA: Media
```

Adding video to a call, using XMPP

Note that the offerer may be different than the original offerer. In addition, unlike the other local description is not set until the

```
// Offerer adds video stream
OffererJS->OffererUA: pc.addStream(videoStream)
OffererJS->OffererUA: offer = pc.createOffer(null);
OffererJS: xmpp = createContentAdd(offer);
OffererJS->AnswererJS: <jingle action="content-add"/>

// content-add arrives at Answerer
AnswererJS: offer = parseContentAdd(xmpp);
AnswererJS->AnswererUA: pc.setRemoteDescription(SDP_OFFER, offer);
AnswererJS->AnswererUA: answer = pc.createAnswer(offer, null);
AnswererJS->AnswererUA: pc.setLocalDescription(SDP_ANSWER, answer);
AnswererJS: xmpp = createContentAccept(answer);
AnswererJS->OffererJS: <jingle action="content-accept"/>

// content-accept arrives at Offerer
OffererJS: answer = parseContentAccept(xmpp);
OffererJS->OffererUA: pc.setLocalDescription(SDP_OFFER, offer);
OffererJS->OffererUA: pc.setRemoteDescription(SDP_ANSWER, answer);
```

Call using SIP

TODO

Handling early media (e.g. 1-800-FEDEX), using SIP

TODO

```
// normal setup

// 180 is received
var sdp = parseResponse(sip);
pc.setRemoteDescription(ANSWER, sdp);

// 200 is received
var sdp = parseResponse(sip);
pc.setRemoteDescription(ANSWER, sdp);
```

Security Considerations

TODO

IANA Considerations

This document requires no actions from IANA.

Acknowledgements

Harald Alvestrand, Cullen Jennings, Matthew Kaufman, Dan Burnett, Neil Stratford, and Eric Rescorla all provided valuable feedback on this proposal.

Open Issues

- More examples needed
 - Forking
 - Glare

References

Normative References

- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.

Informative References

- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", [RFC 5245](#), April 2010.
- [RFC 4585] TODO
- [webrtc-api] Bergkvist, Burnett, Jennings, Narayanan, "WebRTC 1.0: Real-time Communication Between Browsers", October 2011.

Available at

<http://dev.w3.org/2011/webrtc/editor/webrtc.html>