Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 4, 2012

J. Uberti
Google
C. Jennings
Cisco Systems, Inc.
March 3, 2012

Javascript Session Establishment Protocol
draft-ietf-rtcweb-jsep-00

Abstract

This document proposes a mechanism for allowing a Javascript application to fully control the signaling plane of a multimedia session, and discusses how this would work with existing signaling protocols.

This document is an input document for discussion. It should be discussed in the RTCWEB WG list, rtcweb@ietf.org.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 26, 2012.

Copyright Notice

Table of Contents

1. Introduction

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 likely be discovered later. The main reason this mechanism is inflexible is because it embeds a signaling state machine within the browser. Since the browser generates the session descriptions on its own, and fully controls the possible

states and advancement of the signaling state machine, modification of the session descriptions or use of alternate state machines becomes difficult or impossible.

The browser environment also has its own challenges that cause problems for an embedded signaling state machine. One of these is that the user may reload the web page at any time. If this happens, and the state machine is being run at a server, the server can simply push the current state back down to the page and resume the call where it left off. If instead the state machine is run at the browser end, and is instantiated within, for example, the PeerConnection object, that state machine will be reinitialized when the page is reloaded and the JavaScript re-executed. This actually complicates the design of any interoperability service, as all cases where an offer or answer has already been generated but is now "forgotten" must now be handled by trying to move the client state machine forward to the same state it had been in previously in order to match what has already been delivered to and/or answered by the far side, or handled by ensuring that aborts are cleanly handled from every state and the negotiation rapidly restarted.

Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. JSEP Approach

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 a createOffer() API on PeerConnection. The application can do massaging of that offer, if it wants to, and then uses it to set up its local config via a setLocalDescription() API. The offer is then sent off to the remote side over its preferred signaling mechanism (e.g. WebSockets); upon receipt of that offer, the remote party installs it using a setRemoteDescription() API.

When the call is accepted, the callee uses a createAnswer() API to generate an appropriate answer, applies it using setLocalDescription(), and sends the answer back to the initiator over the signaling channel. When the offerer gets that answer, it installs it using setRemoteDescription(), and initial setup is complete. 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 desciptions 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.

3. Other Approaches Considered

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 general 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.

4. Semantics and Syntax

4.1. Signaling Model

JSEP does not specify a particular signaling model or state machine, other than the generic need to exchange RFC 3264 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.

4.2. 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 [RFC4568] 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, SDP_PRANSWER (for a non-final answer) or SDP_ANSWER (for a final 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.

4.3. 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 internal representation for its session descriptions.

However, to simplify Javascript processing, and provide for future flexibility, the SDP syntax is encapsulated within a SessionDescription object, which can be constructed from SDP, and be serialized out to SDP. If we were able to agree on a JSON format for session descriptions, we could easily enable this object to generate/expect JSON.

Other methods may be added to SessionDescription in the future to simplify handling of SessionDescriptions from Javascript.

4.4. 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 startIce method has been added.

4.5. ICE Candidate Trickling

Candidate trickling is a technique through which a caller may incrementally provide candidates to the callee after the initial offer has been dispatched. This allows the callee to begin acting upon the call and setting up the ICE (and perhaps DTLS) connections

immediately, without having to wait for the caller to allocate all possible candidates, resulting in faster call startup in many cases.

JSEP supports optional candidate trickling by providing APIs that provide control and feedback on the ICE candidate gathering process. Applications that support candidate trickling can send the initial offer immediately and send individual candidates when they get a callback with a new candidate; applications that do not support this feature can simply wait for the callback that indicates gathering is complete, and simply create and send their offer, with all the candidates, at this time.

To be clear, aplications that do not make use of candidate tricking can ignore processIceMessage entirely, and use IceCallback solely to indicate when candidate gathering is complete.

4.6. ICE Candidate Format

As with session descriptions, we choose to provide an IceCandidate object that provides some abstraction, but can be easily converted to/from SDP a=candidate lines.

The IceCandidate object has a field to indicate which m= line it should be associated with, and a method to convert to a SDP representation, ex:

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

Currently, a=candidate lines are the only thing that are contained within IceCandidate, as this is the only information that is needed that is not present in the initial offer (i.e. for trickle candidates).

5. Media Setup Overview

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

5.1. 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, iceCb);
OffererJS->OffererUA: pc.addStream(stream);
```


5.1.1. 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);
```

5.1.2. 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.

Since the local user agent must be prepared to receive media upon applying the offer, this operation will cause local decoder resources to be allocated, based on the codecs indicated in the offer.

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

5.1.3. 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. Candidates are provided to the IceCallback as they are allocated, with the |moreToFollow| argument set to true if there are still allocations pending; when the last allocation completes or times out, this callback will be invoked with |moreToFollow| set to false.

```
OffererJS->OffererUA: pc.startIce();
OffererUA->OffererJS: iceCallback(candidate, ...);
```

5.1.4. 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.

5.2. 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, iceCb);
```

5.2.1. 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.

This step will also cause encoder resources to be allocated, based on

the codecs specified in |offer|.

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

5.2.2. Initiating ICE

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

AnswererJS->AnswererUA: pc.startIce();
AnswererUA->AnswererJS: iceCallback(candidate, ...);

5.2.3. 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(candidate);

5.2.4. 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);

5.2.5. 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>
```

5.2.6. 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.

5.3. Completing the Session

5.3.1. 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>
```

5.4. 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 offerer needs to support both its old session description as well as the new one it has offered, until the change is accepted by the remote side.

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

6. Proposed WebRTC API changes

6.1. PeerConnection API

The text below indicates the recommended changes to the PeerConnection API to implement the JSEP functionality. Methods marked with a [+] are new/proposed; methods marked with a [-] have been removed in this proposal.

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

```
    // creates a blob of SDP to be provided as an offer.
[+] SessionDescription createOffer (MediaHints hints);
    // creates a blob of SDP to be provided as an answer.
[+] SessionDescription createAnswer (DOMString offer,
                                     MediaHints hints);
    // actions, for setLocalDescription/setRemoteDescription
[+] const unsigned short SDP_OFFER = 0x100;
[+] const unsigned short SDP_PRANSWER = 0x200;
[+] const unsigned short SDP_ANSWER = 0x300;
    // sets the local session description
[+] void setLocalDescription (unsigned short action,
                              SessionDescription desc);
    // sets the remote session description
[+] void setRemoteDescription (unsigned short action,
                               SessionDescription desc);
    // returns the current local session description
[+] readonly SessionDescription localDescription;
    // returns the current remote session description
[+] readonly SessionDescription remoteDescription;
[-] void processSignalingMessage (DOMString message);
    const unsigned short NEW = 0;     // initial state
[+] const unsigned short OPENING = 1; // local or remote desc set
    const unsigned short ACTIVE = 2;  // local and remote desc set
    const unsigned short CLOSED = 3;  // ended state
    readonly attribute unsigned short readyState;
    // starts ICE connection/handshaking
[+] void startIce (optional IceOptions options);
    // processes received ICE information
[+] void processIceMessage (IceCandidate candidate);
    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 ();
        [ rest of interface omitted ]
    };

    [Constructor (in DOMString sdp)]
    interface SessionDescription {
      // adds the specified candidate to the description
      void addCandidate(IceCandidate candidate);
      // serializes the description to SDP
      DOMString toSdp();
    };

    [Constructor (in DOMString label, in DOMString candidateLine)]
    interface IceCandidate {
      // the m= line this candidate is associated with
      readonly DOMString label;
      // creates a SDP-ized form of this candidate
      DOMString toSdp();
    };
```

6.1.1 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

Indicates whether we want to receive audio; defaults to true if we have audio streams, else false

has_video: boolean

Indicates 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.

6.1.2 createOffer

The createOffer method generates a blob of SDP that contains a RFC 3264 offer with the supported configurations for the session, including descriptions of the local MediaStreams attached to this PeerConnection, the codec/RTP/RTCP options supported by this implementation, and any candidates that have been gathered by the ICE Agent. The |hints| parameter may be supplied to provide additional control over the generated offer.

As an offer, the generated SDP will contain the full set of capabilities supported by the session (as opposed to an answer, which will include only a specific negotiated subset to use); for each SDP line, the generation of the SDP must follow the appropriate process for generating an offer. In the event createOffer is called after the session is established, createOffer will generate an offer that is compatible with the current session, incorporating any changes that have been made to the session since the last complete offer-answer exchange, such as addition or removal of streams. If no changes have been made, the offer will be identical to the current local description.

Session descriptions generated by createOffer must be immediately usable by setLocalDescription; if a system has limited resources (e.g. a finite number of decoders), createOffer should return an offer that reflects the current state of the system, so that setLocalDescription will succeed when it attempts to acquire those resources.

Calling this method does not change the state of the PeerConnection; its use is not required.

A TBD exception is thrown if the |hints| parameter is malformed.

6.1.3 createAnswer

The createAnswer method generates a blob of SDP that contains a RFC 3264 SDP answer with the supported configuration for the session that is compatible with the parameters supplied in |offer|. Like createOffer, the returned blob contains descriptions of the local MediaStreams attached to this PeerConnection, the codec/RTP/RTCP options negotiated for this session, and any candidates that have been gathered by the ICE Agent. The |hints| parameter may be supplied to provide additional control over the generated answer.

As an answer, the generated SDP will contain a specific configuration

that specifies how the media plane should be established. For each SDP line, the generation of the SDP must follow the appropriate process for generating an answer.

Session descriptions generated by createAnswer must be immediately usable by setLocalDescription; like createOffer, the returned description should reflect the current state of the system.

Calling this method does not change the state of the PeerConnection; its use is not required.

A TBD exception is thrown if the |hints| parameter is malformed, or the |offer| parameter is missing or malformed.

6.1.4 SDP_OFFER, SDP_PRANSWER, and SDP_ANSWER

The SDP_XXXX enums serve as arguments to setLocalDescription and setRemoteDescription. They provide information as to how the |description| parameter should be parsed, and how the media state should be changed.

SDP_OFFER indicates that a description should be parsed as an offer; said description may include many possible media configurations. A description used as a SDP_OFFER may be applied anytime the PeerConnection is in a stable state, or as an update to a previously sent but unanswered SDP_OFFER.

SDP_PRANSWER indicates that a description should be parsed as an answer, but not a final answer, and so should not result in the starting of media transmission. A description used as a SDP_PRANSWER may be applied as a response to a SDP_OFFER, or an update to a previously sent SDP_PRANSWER.

SDP_ANSWER indicates that a description should be parsed as an answer, and the offer-answer exchange should be considered complete. A description used as a SDP_ANSWER may be applied as a response to a SDP_OFFER, or an update to a previously send SDP_PRANSWER.

6.1.5 setLocalDescription

The setLocalDescription method instructs the PeerConnection to apply the supplied SDP blob as its local configuration. The |type| parameter indicates whether the blob should be processed as an offer (SDP_OFFER), provisional answer (SDP_PRANSWER), or final answer (SDP_ANSWER); offers and answers are checked differently, using the various rules that exist for each SDP line.

This API changes the local media state; among other things, it sets

up local resources for receiving and decoding media. In order to successfully handle scenarios where the application wants to offer to change from one media format to a different, incompatible format, the PeerConnection must be able to simultaneously support use of both the old and new local descriptions (e.g. support codecs that exist in both descriptions) until a final answer is received, at which point the PeerConnection can fully adopt the new local description, or roll back to the old description if the remote side denied the change.

Changes to the state of media transmission will only occur when a final answer is successfully applied.

A TBD exception is thrown if |description| is invalid. A TBD exception is thrown if there are insufficient local resources to apply |description|.

6.1.6 setRemoteDescription

The setRemoteDescription method instructs the PeerConnection to apply the supplied SDP blob as the desired remote configuration. As in setLocalDescription, the |type| parameter indicates how the blob should be processed.

This API changes the local media state; among other things, it sets up local resources for sending and encoding media.

Changes to the state of media transmission will only occur when a final answer is successfully applied.

A TBD exception is thrown if |description| is invalid. A TBD exception is thrown if there are insufficient local resources to apply |description|.

6.1.7 localDescription

The localDescription method returns a copy of the current local configuration, i.e. what was most recently passed to setLocalDescription, plus any local candidates that have been generated by the ICE Agent.

A null object will be returned if the local description has not yet been established.

6.1.8 remoteDescription

The remoteDescription method returns a copy of the current remote configuration, i.e. what was most recently passed to setRemoteDescription, plus any remote candidates that have been

supplied via processIceMessage.

A null object will be returned if the remote description has not yet been established.

6.1.9 IceOptions

IceOptions is an object that can be passed into startIce to restrict the candidates that are provided to the application and used for connectivity checks. This can be useful if the application wants to only use TURN candidates for privacy reasons, or only local + STUN candidates for cost reasons.

The following properties can be set on IceOptions:

use_candidates: "all", "no_relay", "only_relay"

Indicates what types of local candidates should be used; defaults to "all"

6.1.10 startIce

The startIce method starts or updates the ICE Agent process of gathering local candidates and pinging remote candidates. The |options| argument can be used to restrict which types of local candidates are provided to the application and used for pinging; this

can be used to limit the use of TURN candidates by a callee to avoid leaking location information prior to the call being accepted.

This call may result in a change to the state of the ICE Agent, and may result in a change to media state if it results in connectivity being established.

A TBD exception will be thrown if |options| is malformed.

6.1.11 processIceMessage

The processIceMessage method provides a remote candidate to the ICE Agent, which will be added to the remote description. If startIce has been called, connectivity checks will be sent to the new candidates.

This call will result in a change to the state of the ICE Agent, and may result in a change to media state if it results in connectivity being established.

A TBD exception will be thrown if |candidate| is missing or malformed.

7. 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.

7.1. Call using ROAP

This example demonstrates a ROAP call, without the use of trickle candidates.

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

   // Answerer accepts call
   AnswererJS->AnswererUA: peer.addStream(localStream, null);
   AnswererJS->AnswererUA: answer = peer.createAnswer(msg.offer, null);
   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);

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

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

7.2. Call using XMPP

This example demonstrates an XMPP call, making use of trickle candidates.

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

   OffererJS->OffererUA:   pc.startIce();
```

```
   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.startIce();
   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: pc.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);

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



```
   AnswererUA->OffererUA:  Media

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

7.3. Adding video to a call, using XMPP

This example demonstrates an XMPP call, where the XMPP content-add mechanism is used to add video media to an existing session. For simplicity, candidate exchange is not shown.

Note that the offerer for the change to the session may be different than the original call offerer.

```
// Offerer adds video stream
OffererJS->OffererUA:   pc.addStream(videoStream)
OffererJS->OffererUA:   offer = pc.createOffer(null);
OffererJS:              xmpp = createContentAdd(offer);
OffererJS->OffererUA:   pc.setLocalDescription(SDP_OFFER, 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.setRemoteDescription(SDP_ANSWER, answer);
```

7.4. Simultaneous add of video streams, using XMPP

This example demonstrates an XMPP call, where new video sources are added at the same time to a call that already has video; since adding these sources only affects one side of the call, there is no conflict. The XMPP description-info mechanism is used to indicate the new sources to the remote side.

```
// Offerer and "Answerer" add video streams at the same time
OffererJS->OffererUA:   pc.addStream(offererVideoStream2)
OffererJS->OffererUA:   offer = pc.createOffer(null);
OffererJS:              xmpp = createDescriptionInfo(offer);
OffererJS->OffererUA:   pc.setLocalDescription(SDP_OFFER, offer);
```

```
OffererJS->AnswererJS:  <jingle action="description-info"/>

AnswererJS->AnswererUA: pc.addStream(answererVideoStream2)
AnswererJS->AnswererUA: offer = pc.createOffer(null);
AnswererJS:             xmpp = createDescriptionInfo(offer);
AnswererJS->AnswererUA: pc.setLocalDescription(SDP_OFFER, offer);
AnswererJS->OffererJS:  <jingle action="description-info"/>

// description-info arrives at "Answerer", and is acked
AnswererJS:             offer = parseDescriptionInfo(xmpp);
AnswererJS->OffererJS:  <iq type="result/>  // ack

// description-info arrives at Offerer, and is acked
OffererJS:              offer = parseDescriptionInfo(xmpp);
OffererJS->AnswererJS:  <iq type="result/>  // ack

// ack arrives at Offerer; remote offer is used as an answer
OffererJS->OffererUA:   pc.setRemoteDescription(SDP_ANSWER, offer);

// ack arrives at "Answerer"; remote offer is used as an answer
AnswererJS->AnswererUA: pc.setRemoteDescription(SDP_ANSWER, offer);
```

7.5. Call using SIP

This example demonstrates a simple SIP call (e.g. where the client talks to a SIP proxy over WebSockets).

```
// Call is initiated toward Answerer
OffererJS->OffererUA:   pc = new PeerConnection();
OffererJS->OffererUA:   pc.addStream(localStream, null);
OffererJS->OffererUA:   pc.startIce();
OffererUA->OffererJS:   iceCallback(candidate, false);
OffererJS->OffererUA:   offer = pc.createOffer(null);
OffererJS->OffererUA:   pc.setLocalDescription(SDP_OFFER, offer);
OffererJS:              sip = createInvite(offer);-
OffererJS->AnswererJS:  SIP INVITE w/ SDP

// INVITE arrives at Answerer
AnswererJS->AnswererUA: pc = new PeerConnection();
AnswererJS:             offer = parseInvite(sip);
AnswererJS->AnswererUA: pc.setRemoteDescription(SDP_OFFER, offer);
AnswererUA->AnswererJS: onaddstream(remoteStream);
AnswererJS->AnswererUA: pc.startIce();
AnswererUA->OffererUA:  iceCallback(candidate, false);

// Answerer accepts call
AnswererJS->AnswererUA: peer.addStream(localStream, null);
AnswererJS->AnswererUA: answer = peer.createAnswer(offer, null);
```

```
AnswererJS:              sip = createResponse(200, answer);
AnswererJS->AnswererUA: peer.setLocalDescription(SDP_ANSWER, answer);
AnswererJS->OffererJS:  200 OK w/ SDP

// 200 OK arrives at Offerer
OffererJS:               answer = parseResponse(sip);
OffererJS->OffererUA:    peer.setRemoteDescription(ANSWER, answer);
OffererUA->OffererJS:    onaddstream(remoteStream);
OffererJS->AnswererJS:   ACK

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

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

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

This example demonstrates how early media could be handled; for simplicity, only the offerer side of the call is shown.

```
// Call is initiated toward Answerer
OffererJS->OffererUA:    pc = new PeerConnection();
OffererJS->OffererUA:    pc.addStream(localStream, null);
OffererJS->OffererUA:    pc.startIce();
OffererUA->OffererJS:    iceCallback(candidate, false);
OffererJS->OffererUA:    offer = pc.createOffer(null);
OffererJS->OffererUA:    pc.setLocalDescription(SDP_OFFER, offer);
OffererJS:               sip = createInvite(offer);
OffererJS->AnswererJS:   SIP INVITE w/ SDP

// 180 Ringing is received by offerer, w/ SDP
OffererJS:               answer = parseResponse(sip);
OffererJS->OffererUA:    pc.setRemoteDescription(SDP_PRANSWER, answer);
OffererUA->OffererJS:    onaddstream(remoteStream);

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

```
// 200 OK arrives at Offerer
OffererJS:              answer = parseResponse(sip);
OffererJS->OffererUA:   pc.setRemoteDescription(SDP_ANSWER, answer);
OffererJS->AnswererJS:  ACK
```

8. Example Application

The following example demonstrates a simple video calling application, roughly corresponding to the flow in Example 7.1.

```
var signalingChannel = createSignalingChannel();
var pc = null;
var hasCandidates = false;

function start(isCaller) {
  // create a PeerConnection and hook up the IceCallback
  pc = new webkitPeerConnection(
         "", function (candidate, moreToFollow) {
    if (!moreToFollow) {
      hasCandidates = true;
      maybeSignal(isCaller);
    }
  });

  // get the local stream and show it in the local video element
  navigator.webkitGetUserMedia(
        {"audio": true, "video": true}, function (localStream) {
    selfView.src = webkitURL.createObjectURL(localStream);
    pc.addStream(localStream);
    maybeSignal(isCaller);
  }

  // once remote stream arrives, show it in the remote video element
  pc.onaddstream = function(evt) {
    remoteView.src = webkitURL.createObjectURL(evt.stream);
  };

  // if we're the caller, create and install our offer,
  // and start candidate generation
  if (isCaller) {
    offer = pc.createOffer(null);
    pc.setLocalDescription(SDP_OFFER, offer);
```

```
      pc.startIce();
    }
  }

  function maybeSignal(isCaller) {
    // only signal once we have a local stream and local candidates
    if (localStreams.size() == 0 || !hasCandidates) return;
    if (isCaller) {
      offer = pc.localDescription;
      signalingChannel.send(
          JSON.stringify({ "type": "offer", "sdp": offer }));
    } else {
      // if we're the callee, generate, apply, and send the answer
```



```
      answer = pc.createAnswer(pc.remoteDescription, null);
      pc.setLocalDescription(SDP_ANSWER, answer);
      signalingChannel.send(
          JSON.stringify({ "type": "answer", "sdp": answer }));
    }
  }

  signalingChannel.onmessage = function(evt) {
    var msg = JSON.parse(evt.data);
    if (msg.type == "offer") {
      // create the PeerConnection
      start(false);
      // feed the received offer into the PeerConnection and
      // start candidate generation
      pc.setRemoteDescription(PeerConnection.SDP_OFFER, msg.sdp);
      pc.startIce();
    } else if (msg.type == "answer") {
      // feed the answer into the PeerConnection to complete setup
      pc.setRemoteDescription(PeerConnection.SDP_ANSWER, msg.sdp);
    }
```

9. Security Considerations

TODO

10. IANA Considerations

This document requires no actions from IANA.

11. Acknowledgements

Harald Alvestrand, Dan Burnett, Neil Stratford, Eric Rescorla, and Anant Narayanan all provided valuable feedback on this proposal. Matthew Kaufman provided the observation that keeping state out of the browser allows a call to continue even if the page is reloaded. Adam Bergvist provided a code example that served as the basis for the example in Section 8.

12. References

12.1. Normative References

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

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.

12.2. Informative References

[RFC4568] Andreasen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", RFC 4568, July 2006.

[RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.

[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

Appendix A. Open Issues

- Determine list of exceptions that can be thrown by each method. Leaning toward something like a PCException, a la https://developer.mozilla.org/en/IndexedDB/IDBDatabaseException

- Need callback to indicate that the transport is down, e.g. ICE_DISCONNECTED or ondisconnected().

Appendix B. Change log

00: Migrated from draft-uberti-rtcweb-jsep-02.

Authors' Addresses

Justin Uberti
Google
5 Cambridge Center
Cambridge, MA 02142

Email: justin@uberti.name

Cullen Jennings
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: fluffy@cisco.com