

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2014

J. Uberti
Google
C. Jennings
Cisco
October 22, 2013

Javascript Session Establishment Protocol
draft-ietf-rtcweb-jsep-05

Abstract

This document describes the mechanisms for allowing a Javascript application to control the signaling plane of a multimedia session via the interface specified in the W3C RTCPeerConnection API, and discusses how this relates to existing signaling protocols.

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 April 25, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. General Design of JSEP	3
1.2. Other Approaches Considered	5
2. Terminology	6
3. Semantics and Syntax	6
3.1. Signaling Model	6
3.2. Session Descriptions and State Machine	7
3.3. Session Description Format	9
3.4. ICE	10
3.4.1. ICE Candidate Trickling	10
3.4.1.1. ICE Candidate Format	10
3.5. Interactions With Forking	11
3.5.1. Sequential Forking	11
3.5.2. Parallel Forking	12
3.6. Session Rehydration	13
4. Interface	13
4.1. Methods	14
4.1.1. createOffer	14
4.1.2. createAnswer	15
4.1.3. SessionDescriptionType	16
4.1.3.1. Use of Provisional Answers	16
4.1.3.2. Rollback	17
4.1.4. setLocalDescription	18
4.1.5. setRemoteDescription	18
4.1.6. localDescription	19
4.1.7. remoteDescription	19
4.1.8. updateIce	19
4.1.9. addIceCandidate	19
5. SDP Interaction Procedures	20
5.1. Requirements Overview	20
5.1.1. Implementation Requirements	20
5.1.2. Usage Requirements	21
5.2. Constructing an Offer	22
5.2.1. Initial Offers	22
5.2.2. Subsequent Offers	26
5.2.3. Constraints Handling	28
5.2.3.1. OfferToReceiveAudio	28
5.2.3.2. OfferToReceiveVideo	28
5.2.3.3. VoiceActivityDetection	28
5.2.3.4. IceRestart	29
5.3. Generating an Answer	29
5.3.1. Initial Answers	29
5.3.2. Subsequent Answers	33
5.3.3. Constraints Handling	33
5.4. Parsing an Offer	33
5.5. Parsing an Answer	33

5.6. Applying a Local Description	33
5.7. Applying a Remote Description	33
6. Configurable SDP Parameters	33
7. Security Considerations	34
8. IANA Considerations	34
9. Acknowledgements	34
10. References	35
10.1. Normative References	35
10.2. Informative References	37
Appendix A. JSEP Implementation Examples	38
A.1. Example API Flows	38
A.1.1. Call using ROAP	38
A.1.2. Call using XMPP	39
A.1.3. Adding video to a call, using XMPP	40
A.1.4. Simultaneous add of video streams, using XMPP	41
A.1.5. Call using SIP	41
A.1.6. Handling early media (e.g. 1-800-GO FEDEX), using SIP	42
A.2. Example Session Descriptions	43
A.2.1. createOffer	43
A.2.2. createAnswer	45
A.2.3. Call Flows	46
Appendix B. Change log	46
Authors' Addresses	48

1. Introduction

This document describes how the W3C WEBRTC `RTCPeerConnection` interface[W3C.WD-webrtc-20111027] is used to control the setup, management and teardown of a multimedia session.

1.1. General Design of JSEP

The 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 browser environment also has its own challenges that pose problems for an embedded signaling state machine. One of these is that the user may reload the web page at any time. If the browser is fully in charge of the signaling state, this will result in the loss of the call when this state is wiped by the reload. However, if the

state can be stored at the server, and pushed back down to the new page, the call can be resumed with minimal interruption.

With these considerations in mind, this document describes the Javascript Session Establishment Protocol (JSEP) that allows for full control of the signaling state machine from 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.

In this document, the use of JSEP is described as if it always occurs between two browsers. Note though in many cases it will actually be between a browser and some kind of server, such as a gateway or MCU. This distinction is invisible to the browser; it just follows the instructions it is given via the API.

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. The application optionally modifies that offer, and then uses it to set up its local config via the `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 the `setRemoteDescription()` API.

When the call is accepted, the callee uses the `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 [RFC5245], JSEP decouples the ICE state machine from the overall signaling state machine, as the ICE state machine must remain in the browser, because only the browser has the necessary knowledge of candidates and other transport info. Performing this separation also 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 used in the aggregated form. 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.

Through its abstraction of signaling, the JSEP approach does require the application to be aware of the signaling process. While the application does not need to understand the contents of session descriptions to set up a call, the application must call the right APIs at the right times, convert the session descriptions and ICE information into the defined messages of its chosen signaling protocol, and perform the reverse conversion on the messages it receives from the other side.

One way to mitigate this is to provide a Javascript library that hides this complexity from the developer; said library would implement a given signaling protocol along with its state machine and serialization code, presenting a higher level call-oriented interface to the application developer. For example, this library could easily adapt the JSEP API into the API that was proposed for the ROAP signaling protocol [I-D.jennings-rtcweb-signaling], which would perform a ROAP call setup under the covers, interacting with the application only when it needs a signaling message to be sent. In the same fashion, one could also implement other popular signaling protocols, including SIP or Jingle. This allow JSEP to provide greater control for the experienced developer without forcing any additional complexity on the novice developer.

1.2. Other Approaches Considered

One approach that was considered instead of JSEP was to include a lightweight signaling protocol. Instead of providing session descriptions to the API, the API would produce and consume messages from this protocol. While providing a more high-level API, this put more control of signaling within the browser, forcing the browser to have to understand and handle concepts like signaling glare. In addition, it prevented the application from driving the state machine to a desired state, as is needed in the page reload case.

A second approach that was considered but not chosen was to decouple the management of the media control objects from session descriptions, instead offering APIs that would control each component directly. This was rejected based on a feeling that requiring exposure of this level of complexity to the application programmer would not be beneficial; it would result in an API where even a simple example would require a significant amount of code to orchestrate all the needed interactions, as well as creating a large API surface that needed to be agreed upon and documented. In addition, these API points could be called in any order, resulting in a more complex set of interactions with the media subsystem than the JSEP approach, which specifies how session descriptions are to be evaluated and applied.

One variation on JSEP that was considered was to keep the basic session description-oriented API, but to move the mechanism for generating offers and answers out of the browser. Instead of providing createOffer/createAnswer methods within the browser, this approach would instead expose a getCapabilities API which would provide the application with the information it needed in order to generate its own 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 the supported capabilities. While this could certainly be addressed by using a library like the one mentioned above, it basically forces the use of said library even for a simple example. Providing createOffer/createAnswer avoids this problem, but still allows applications to generate their own offers/answers (to a large extent) if they choose, using the description generated by createOffer as an indication of the browser's capabilities.


2. 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 [RFC2119].

3. Semantics and Syntax

3.1. Signaling Model

JSEP does not specify a particular signaling model or state machine, other than the generic need to exchange SDP media descriptions in the fashion described by [RFC3264] (offer/answer) 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 session. However, the browser is totally decoupled from the actual mechanism by which these offers and answers are communicated to the remote side, including addressing, retransmission, forking, and glare handling. These issues are left entirely up to the application; the application has complete control over which offers and answers get handed to the browser, and when.



+-----+

+-----+

Figure 1: JSEP Signaling Model

3.2. Session Descriptions and State Machine

In order to establish the media plane, the user agent 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 applies to the local side or the remote side 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 receive, which, when intersected with the set of codecs the remote side supports, specifies what the remote side should send. However, not all parameters follow this rule; for example, the SRTP parameters [RFC4568] sent to a remote party indicate what the local side will use to encrypt, and thereby what the remote party should expect to receive; the remote party will have to accept these parameters, with no option to choose a different value.

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 offer before the answer arrives.

Therefore, in order to handle session descriptions properly, the user agent needs:

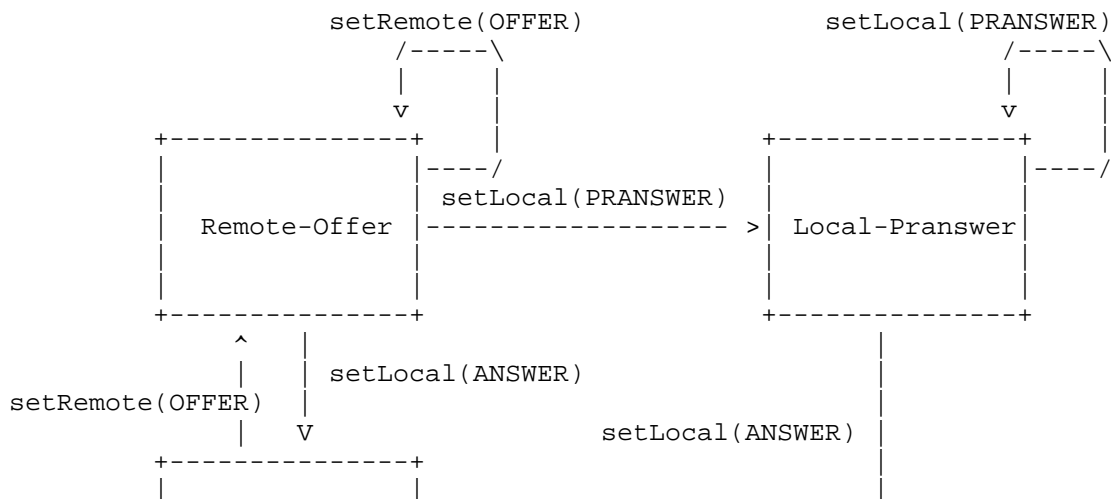
1. To know if a session description pertains to the local or remote side.
2. To know if a session description is an offer or an answer.
3. To allow the offer to be specified independently of the answer.

JSEP addresses this by adding both a `setLocalDescription` and a `setRemoteDescription` method and having session description objects

contain a type field indicating the type of session description being supplied. This satisfies the requirements listed above for both the offerer, who first calls `setLocalDescription(sdp [offer])` and then later `setRemoteDescription(sdp [answer])`, as well as for the answerer, who first calls `setRemoteDescription(sdp [offer])` and then later `setLocalDescription(sdp [answer])`.

JSEP also allows for an answer to be treated as provisional by the application. Provisional answers provide a way for an answerer to communicate initial session parameters back to the offerer, in order to allow the session to begin, while allowing a final answer to be specified later. This concept of a final answer is important to the offer/answer model; when such an answer is received, any extra resources allocated by the caller can be released, now that the exact session configuration is known. These "resources" can include things like extra ICE components, TURN candidates, or video decoders. Provisional answers, on the other hand, do no such deallocation results; as a result, multiple dissimilar provisional answers can be received and applied during call setup.

In [RFC3264], the constraint at the signaling level is that only one offer can be outstanding for a given session, but from the media stack level, a new offer can be generated at any point. For example, when using SIP for signaling, if one offer is sent, then cancelled using a SIP CANCEL, another offer can be generated even though no answer was received for the first offer. To support this, the JSEP media layer can provide an offer whenever the Javascript application needs one for the signaling. The answerer can send back zero or more provisional answers, and finally end the offer-answer exchange by sending a final answer. The state machine for this is as follows:






Figure 2: JSEP State Machine

Aside from these state transitions, there is no other difference between the handling of provisional ("pranswer") and final ("answer") answers.

3.3. Session Description Format

In the 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 currently uses 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 future specifications agree on a JSON format for session descriptions, we could easily enable this object to generate and consume that JSON.

Other methods may be added to SessionDescription in the future to simplify handling of SessionDescriptions from Javascript. In the meantime, Javascript libraries can be used to perform these manipulations.

Note that most applications should be able to treat the SessionDescriptions produced and consumed by these various API calls as opaque blobs; that is, the application will not need to read or change them. The W3C API will provide appropriate APIs to allow the application to control various session parameters, which will provide the necessary information to the browser about what sort of SessionDescription to produce.

3.4. ICE

When a new ICE candidate is available, the ICE Agent will notify the application via a callback; these candidates will automatically be added to the local session description. When all candidates have been gathered, the callback will also be invoked to signal that the gathering process is complete.

3.4.1. 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; the semantics of "Trickle ICE" are defined in [I-D.ietf-mmusic-trickle-ice]. This process 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 gather all possible candidates. This results in faster call startup in cases where gathering is not performed prior to initiating the call.

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 the notified of a new candidate; applications that do not support this feature can simply wait for the indication that gathering is complete, and then create and send their offer, with all the candidates, at this time.

Upon receipt of trickled candidates, the receiving application will supply them to its ICE Agent. This triggers the ICE Agent to start using the new remote candidates for connectivity checks.

3.4.1.1. ICE Candidate Format

As with session descriptions, the syntax of the IceCandidate object provides some abstraction, but can be easily converted to and from the SDP candidate lines.

The candidate lines are the only SDP information that is contained within IceCandidate, as they represent the only information needed that is not present in the initial offer (i.e. for trickle candidates). This information is carried with the same syntax as the "candidate-attribute" field defined for ICE. For example:

```
candidate:1 1 UDP 1694498815 192.0.2.33 10000 typ host
```

The IceCandidate object also contains fields to indicate which m-line it should be associated with. The m line can be identified in one of two ways; either by a m-line index, or a MID. The m-line index is a zero-based index, referring to the Nth m-line in the SDP. The MID uses the "media stream identification", as defined in [RFC5888], to identify the m-line. WebRTC implementations creating an ICE Candidate object MUST populate both of these fields. Implementations receiving an ICE Candidate object SHOULD use the MID if they implement that functionality, or the m-line index, if not.

3.5. Interactions With Forking

Some call signaling systems allow various types of forking where an SDP Offer may be provided to more than one device. For example, SIP [RFC3261] defines both a "Parallel Search" and "Sequential Search". Although these are primarily signaling level issues that are outside the scope of JSEP, they do have some impact on the configuration of the media plane which is relevant. When forking happens at the signaling layer, the Javascript application responsible for the signaling needs to make the decisions about what media should be sent or received at any point of time, as well as which remote endpoint it should communicate with; JSEP is used to make sure the media engine can make the RTP and media perform as required by the application. The basic operations that the applications can have the media engine do are:

Start exchanging media to a given remote peer, but keep all the resources reserved in the offer.

Start exchanging media with a given remote peer, and free any resources in the offer that are not being used.

3.5.1. Sequential Forking

Sequential forking involves a call being dispatched to multiple remote callees, where each callee can accept the call, but only one active session ever exists at a time; no mixing of received media is performed.

JSEP handles sequential forking well, allowing the application to easily control the policy for selecting the desired remote endpoint. When an answer arrives from one of the callees, the application can choose to apply it either as a provisional answer, leaving open the possibility of using a different answer in the future, or apply it as a final answer, ending the setup flow.

In a "first-one-wins" situation, the first answer will be applied as a final answer, and the application will reject any subsequent answers. In SIP parlance, this would be ACK + BYE.

In a "last-one-wins" situation, all answers would be applied as provisional answers, and any previous call leg will be terminated. At some point, the application will end the setup process, perhaps with a timer; at this point, the application could reapply the existing remote description as a final answer.

3.5.2. Parallel Forking

Parallel forking involves a call being dispatched to multiple remote callees, where each callee can accept the call, and multiple simultaneous active signaling sessions can be established as a result. If multiple callees send media at the same time, the possibilities for handling this are described in Section 3.1 of [RFC3960]. Most SIP devices today only support exchanging media with a single device at a time, and do not try to mix multiple early media audio sources, as that could result in a confusing situation. For example, consider having a European ringback tone mixed together with the North American ringback tone - the resulting sound would not be like either tone, and would confuse the user. If the signaling application wishes to only exchange media with one of the remote endpoints at a time, then from a media engine point of view, this is exactly like the sequential forking case.

In the parallel forking case where the Javascript application wishes to simultaneously exchange media with multiple peers, the flow is slightly more complex, but the Javascript application can follow the strategy that [RFC3960] describes using UPDATE. (It is worth noting that use cases where this is the desired behavior are very unusual.) The UPDATE approach allows the signaling to set up a separate media flow for each peer that it wishes to exchange media with. In JSEP, this offer used in the UPDATE would be formed by simply creating a new PeerConnection and making sure that the same local media streams

have been added into this new PeerConnection. Then the new PeerConnection object would produce a SDP offer that could be used by the signaling to perform the UPDATE strategy discussed in [RFC3960].

As a result of sharing the media streams, the application will end up with N parallel PeerConnection sessions, each with a local and remote description and their own local and remote addresses. The media flow from these sessions can be managed by specifying SDP direction attributes in the descriptions, or the application can choose to play out the media from all sessions mixed together. Of course, if the application wants to only keep a single session, it can simply terminate the sessions that it no longer needs.

3.6. Session Rehydration

In the event that the local application state is reinitialized, either due to a user reload of the page, or a decision within the application to reload itself (perhaps to update to a new version), it is possible to keep an existing session alive, via a process called "rehydration". The explicit goal of rehydration is to carry out this session resumption with no interaction with the remote side other than normal call signaling messages.

With rehydration, the current signaling state is persisted somewhere outside of the page, perhaps on the application server, or in browser local storage. The page is then reloaded, the saved signaling state is retrieved, and a new PeerConnection object is created for the session. The previously obtained MediaStreams are re-acquired, and are given the same IDs as the original session; this ensures the IDs in use by the remote side continue to work. Next, a new offer is generated by the new PeerConnection; this offer will have new ICE and possibly new DTLS-SRTP certificate fingerprints (since the old ICE and SRTP state has been lost). Finally, this offer is used to re-initiate the session with the existing remote endpoint, who simply sees the new offer as an in-call renegotiation, and replies with an answer that can be supplied to setRemoteDescription. ICE processing proceeds as usual, and as soon as connectivity is established, the session will be back up and running again.

[OPEN ISSUE: EKR proposed an alternative rehydration approach where the actual internal PeerConnection object in the browser was kept alive for some time after the web page was killed and provided some way for a new page to acquire the old PeerConnection object.]

4. Interface

This section details the basic operations that must be present to implement JSEP functionality. The actual API exposed in the W3C API

may have somewhat different syntax, but should map easily to these concepts.

4.1. Methods

4.1.1. createOffer

The createOffer method generates a blob of SDP that contains a [RFC3264] 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. A constraints parameters may be supplied to provide additional control over the generated offer. This constraints parameter should allow for the following manipulations to be performed:

- o To indicate support for a media type even if no MediaStreamTracks of that type have been added to the session (e.g., an audio call that wants to receive video.)
- o To trigger an ICE restart, for the purpose of reestablishing connectivity.
- o For re-offer cases, to request an offer that contains the full set of supported capabilities, as opposed to just the currently negotiated parameters.

In the initial offer, the generated SDP will contain all desired functionality for the session (certain parts that are supported but not desired by default may be omitted); for each SDP line, the generation of the SDP will follow the process defined for generating an initial offer from the document that specifies the given SDP line. The exact handling of initial offer generation is detailed in Section 5.2.1. below.

In the event `createOffer` is called after the session is established, `createOffer` will generate an offer to modify the current session based on any changes that have been made to the session, e.g. adding or removing `MediaStreams`, or requesting an ICE restart. For each existing stream, the generation of each SDP line must follow the process defined for generating an updated offer from the document that specifies the given SDP line. For each new stream, the generation of the SDP must follow the process of generating an initial offer, as mentioned above. If no changes have been made, or for SDP lines that are unaffected by the requested changes, the offer will only contain the parameters negotiated by the last offer-answer exchange. The exact handling of subsequent offer generation is detailed in Section 5.2.2. below.

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. Because this method may need to inspect the system state to determine the currently available resources, it may be implemented as an async operation.

Calling this method may do things such as generate new ICE credentials, but does not result in candidate gathering, or cause media to start or stop flowing.

4.1.1.2. `createAnswer`

The `createAnswer` method generates a blob of SDP that contains a [RFC3264] SDP answer with the supported configuration for the session that is compatible with the parameters supplied in the 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. A constraints 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 process defined for generating an answer from the document that specifies the given SDP line. The exact handling of answer generation is detailed in Section 5.3. below.

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

this method may need to inspect the system state to determine the currently available resources, it may need to be implemented as an async operation.

Calling this method may do things such as generate new ICE credentials, but does not trigger candidate gathering or change media state.

4.1.3. SessionDescriptionType

Session description objects (RTCSessionDescription) may be of type "offer", "pranswer", and "answer". These types provide information as to how the description parameter should be parsed, and how the media state should be changed.

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

"pranswer" indicates that a description should be parsed as an answer, but not a final answer, and so should not result in the freeing of allocated resources. It may result in the start of media transmission, if the answer does not specify an inactive media direction. A description used as a "pranswer" may be applied as a response to an "offer", or an update to a previously sent "answer".

"answer" indicates that a description should be parsed as an answer, the offer-answer exchange should be considered complete, and any resources (decoders, candidates) that are no longer needed can be released. A description used as an "answer" may be applied as a response to a "offer", or an update to a previously sent "pranswer".

The only difference between a provisional and final answer is that the final answer results in the freeing of any unused resources that were allocated as a result of the offer. As such, the application can use some discretion on whether an answer should be applied as provisional or final, and can change the type of the session description as needed. For example, in a serial forking scenario, an application may receive multiple "final" answers, one from each remote endpoint. The application could choose to accept the initial answers as provisional answers, and only apply an answer as final when it receives one that meets its criteria (e.g. a live user instead of voicemail).

4.1.3.1. Use of Provisional Answers

Most web applications will not need to create answers using the "pranswer" type. The preferred handling for a web application would be to create and send an "inactive" answer more or less immediately after receiving the offer, instead of waiting for a human user to physically answer the call. Later, when the human input is received, the application can create a new "sendrecv" offer to update the previous offer/answer pair and start the media flow. This approach is preferred because it minimizes the amount of time that the offer-answer exchange is left open, in addition to avoiding media clipping by ensuring the transport is ready to go by the time the call is physically answered. However, some applications may not be able to do this, particularly ones that are attempting to gateway to other signaling protocols. In these cases, "pranswer" can still allow the application to warm up the transport.

Consider a typical web application that will set up a data channel, an audio channel, and a video channel. When an endpoint receives an offer with these channels, it could send an answer accepting the data channel for two-way data, and accepting the audio and video tracks as inactive or receive-only. It could then ask the user to accept the call, acquire the local media streams, and send a new offer to the remote side moving the audio and video to be two-way media. By the time the human has accepted the call and sent the new offer, it is likely that the ICE and DTLS handshaking for all the channels will already be set up.

4.1.3.2. Rollback

In certain situations it may be desirable to "undo" a change made to `setLocalDescription` or `setRemoteDescription`. Consider a case where a call is ongoing, and one side wants to change some of the session parameters; that side generates an updated offer and then calls `setLocalDescription`. However, the remote side, either before or after `setRemoteDescription`, decides it does not want to accept the new parameters, and sends a reject message back to the offerer. Now, the offerer, and possibly the answerer as well, need to return to a stable state and the previous local/remote description. To support this, we introduce the concept of "rollback".

A rollback returns the state machine to its previous state, and the local or remote description to its previous value. Any resources or candidates that were allocated by the new local description are discarded; any media that is received will be processed according to the previous session description.

A rollback is performed by supplying a session description of type "rollback" to either `setLocalDescription` or `setRemoteDescription`, depending on which needs to be rolled back (i.e. if the new offer was

supplied to `setLocalDescription`, the rollback should be done on `setLocalDescription` as well.)

4.1.4. `setLocalDescription`

The `setLocalDescription` method instructs the `PeerConnection` to apply the supplied SDP blob as its local configuration. The `type` field indicates whether the blob should be processed as an offer, provisional answer, or final 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.

This API indirectly controls the candidate gathering process. When a local description is supplied, and the number of transports currently in use does not match the number of transports needed by the local description, the `PeerConnection` will create transports as needed and begin gathering candidates for them.

If `setRemoteDescription` was previously called with an offer, and `setLocalDescription` is called with an answer (provisional or final), and the media directions are compatible, and media are available to send, this will result in the starting of media transmission.

4.1.5. `setRemoteDescription`

The `setRemoteDescription` method instructs the `PeerConnection` to apply the supplied SDP blob as the desired remote configuration. As in `setLocalDescription`, the `type` field of the 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.

If `setRemoteDescription` was previously called with an offer, and `setLocalDescription` is called with an answer (provisional or final), and the media directions are compatible, and media are available to send, this will result in the starting of media transmission.

4.1.6. 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.

TODO: Do we need to expose accessors for both the current and proposed local description?

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

4.1.7. 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.

TODO: Do we need to expose accessors for both the current and proposed remote description?

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

4.1.8. updateIce

The updateIce method allows the configuration of the ICE Agent to be changed during the session, primarily for changing which types of local candidates are provided to the application and used for connectivity checks. A callee may initially configure the ICE Agent to use only relay candidates, to avoid leaking location information, but update this configuration to use all candidates once the call is accepted.

Regardless of the configuration, the gathering process collects all available candidates, but excluded candidates will not be surfaced in onIceCandidate callback or used for connectivity checks.

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.

4.1.9. addIceCandidate

The addIceCandidate method provides a remote candidate to the ICE Agent, which, if parsed successfully, will be added to the remote

description according to the rules defined for Trickle ICE. Connectivity checks will be sent to the new candidate.

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.

5. SDP Interaction Procedures

This section describes the specific procedures to be followed when creating and parsing SDP objects.

5.1. Requirements Overview

JSEP implementations must comply with the specifications listed below that govern the creation and processing of offers and answers.

The first set of specifications is the "mandatory-to-implement" set. All implementations must support these behaviors, but may not use all of them if the remote side, which may not be a JSEP endpoint, does not support them.

The second set of specifications is the "mandatory-to-use" set. The local JSEP endpoint and any remote endpoint must indicate support for these specifications in their session descriptions.

5.1.1. Implementation Requirements

This list of mandatory-to-implement specifications is derived from the requirements outlined in [I-D.ietf-rtcweb-rtp-usage].

R-1 [RFC4566] is the base SDP specification and MUST be implemented.

R-2 [RFC5764] MUST be supported for signaling the UDP/TLS/RTP/SAVPF RTP profile.

R-3 [RFC5245] MUST be implemented for signaling the ICE credentials and candidate lines corresponding to each media stream. The ICE implementation MUST be a Full implementation, not a Lite implementation.

R-4 [RFC5763] MUST be implemented to signal DTLS certificate fingerprints.

R-5 [RFC4568] MUST NOT be implemented to signal SDES SRTP keying information.

- R-6 The [RFC5888] grouping framework MUST be implemented for signaling grouping information, and MUST be used to identify m= lines via the a=mid attribute.
- R-7 [I-D.ietf-mmusic-msid] MUST be supported, in order to signal associations between RTP objects and W3C MediaStreams and MediaStreamTracks in a standard way.
- R-8 The bundle mechanism in [I-D.ietf-mmusic-sdp-bundle-negotiation] MUST be supported to signal the ability to multiplex RTP streams on a single UDP port, in order to avoid excessive use of port number resources.
- R-9 The SDP attributes of "sendonly", "recvonly", "inactive", and "sendrecv" from [RFC4566] MUST be implemented to signal information about media direction.
- R-10 [RFC5576] MUST be implemented to signal RTP SSRC values.
- R-11 [RFC4585] MUST be implemented to signal RTCP based feedback.
- R-12 [RFC5761] MUST be implemented to signal multiplexing of RTP and RTCP.
- R-13 [RFC5506] MUST be implemented to signal reduced-size RTCP messages.
- R-14 [RFC3556] with bandwidth modifiers MAY be supported for specifying RTCP bandwidth as a fraction of the media bandwidth, RTCP fraction allocated to the senders and setting maximum media bit-rate boundaries.

As required by [RFC4566], Section 5.13, JSEP implementations MUST ignore unknown attribute (a=) lines.

5.1.2. Usage Requirements

All session descriptions handled by JSEP endpoints, both local and remote, MUST indicate support for the following specifications. If any of these are absent, this omission MUST be treated as an error.

- R-1 Either the UDP/TLS/RTP/SAVP or the UDP/TLS/RTP/SAVPF RTP profile, as specified in [RFC5764], MUST be used.
- R-2 ICE, as specified in [RFC5245], MUST be used. Note that the remote endpoint MAY use a Lite implementation.
- R-3 DTLS-SRTP, as specified in [RFC5763], MUST be used.

5.2. Constructing an Offer

When `createOffer` is called, a new SDP description must be created that includes the functionality specified in [I-D.ietf-rtcweb-rtp-usage]. The exact details of this process are explained below.

5.2.1. Initial Offers

When `createOffer` is called for the first time, the result is known as the initial offer.

The first step in generating an initial offer is to generate session-level attributes, as specified in [RFC4566], Section 5. Specifically:

- o The first SDP line MUST be "v=0", as specified in [RFC4566], Section 5.1
- o The second SDP line MUST be an "o=" line, as specified in [RFC4566], Section 5.2. The value of the <username> field SHOULD be "-". The value of the <sess-id> field SHOULD be a cryptographically random number. To ensure uniqueness, this number SHOULD be at least 64 bits long. The value of the <sess-version> field SHOULD be zero. The value of the <nettype> <addrtype> <unicast-address> tuple SHOULD be set to a non-meaningful address, such as IN IP4 0.0.0.0, to prevent leaking the local address in this field. As mentioned in [RFC4566], the entire o= line needs to be unique, but selecting a random number for <sess-id> is sufficient to accomplish this.
- o The third SDP line MUST be a "s=" line, as specified in [RFC4566], Section 5.3; to match the "o=" line, a single dash SHOULD be used as the session name, e.g. "s=-".
- o Session Information ("i="), URI ("u="), Email Address ("e="), Phone Number ("p="), Bandwidth ("b="), Repeat Times ("r="), and Time Zones ("z=") lines are not useful in this context and SHOULD NOT be included.
- o Encryption Keys ("k=") lines do not provide sufficient security and MUST NOT be included.
- o A "t=" line MUST be added, as specified in [RFC4566], Section 5.9; both <start-time> and <stop-time> SHOULD be set to zero, e.g. "t=0 0".

- o An "a=msid-semantic:WMS" line MUST be added, as specified in [I-D.ietf-mmusic-msid], Section 4.

The next step is to generate m= sections for each MediaStreamTrack that has been added to the PeerConnection via the addStream method. (Note that this method takes a MediaStream, which can contain multiple MediaStreamTracks, and therefore multiple m= sections can be generated even if addStream is only called once.) When generating m= sections, the ordering is based on (1) the order in which the MediaStreams were added to the PeerConnection, and (2) the alphabetical ordering of the media type for the MediaStreamTrack. For example, if a MediaStream containing both an audio and a video MediaStreamTrack is added to a PeerConnection, the resultant m=audio section will precede the m=video section.

Each m= section, provided it is not being bundled into another m= section, MUST generate a unique set of ICE credentials and gather its own unique set of ICE candidates. Otherwise, it MUST use the same ICE credentials and candidates that were used in the m= section that it is being bundled into.

For DTLS, all m= sections MUST use the certificate for the identity that has been specified for the PeerConnection; as a result, they MUST all have the same [RFC4572]fingerprint value.

Each m= section should be generated as specified in [RFC4566], Section 5.14. For the m= line itself, the following rules MUST be followed:

- o The port value is set to the port of the default ICE candidate for this m= section; if this m= section is not being bundled into another m= section, the port value MUST be unique. If no candidates have yet been gathered, and a 'null' port value is being used, as indicated in [I-D.ietf-mmusic-trickle-ice], Section 5.1., this port MUST still be unique.
- o To properly indicate use of DTLS, the <proto> field MUST be set to "UDP/TLS/RTP/SAVPF", as specified in [RFC5764], Section 8.

Each m= section MUST include the following attribute lines:

- o An "a=mid" line, as specified in [RFC5888], Section 4.
- o An "a=msid" line, as specified in [I-D.ietf-mmusic-msid], Section 2.
- o [OPEN ISSUE: Use of App Token versus stream-correlator]

- o An "a=sendrecv" line, as specified in [RFC3264], Section 5.1.
- o For each supported codec, "a=rtpmap" and "a=fmtp" lines, as specified in [RFC4566], Section 6. For audio, the codecs specified in [I-D.ietf-rtcweb-audio], Section 3, MUST be supported.
- o For each primary codec where RTP retransmission should be used, a corresponding "a=rtpmap" line indicating "rtx" with the clock rate of the primary codec and an "a=fmtp" line that references the payload type for the primary codec, as specified in [RFC4588], Section 8.1.
- o For each supported FEC mechanism, a corresponding "a=rtpmap" line indicating the desired FEC codec.
- o "a=ice-ufrag" and "a=ice-passwd" lines, as specified in [RFC5245], Section 15.4.
- o An "a=ice-options" line, with the "trickle" option, as specified in [I-D.ietf-mmusic-trickle-ice], Section 4.
- o For each candidate that has been gathered during the most recent gathering phase, an "a=candidate" line, as specified in [RFC5245], Section 4.3., paragraph 3.
- o For the current default candidate, a "c=" line, as specific in [RFC5245], Section 4.3., paragraph 6. If no candidates have been gathered yet, the default candidate should be set to the 'null' value defined in [I-D.ietf-mmusic-trickle-ice], Section 5.1.
- o An "a=fingerprint" line, as specified in [RFC4572], Section 5; the algorithm used for the fingerprint MUST match that used in the certificate signature.
- o An "a=setup" line, as specified in [RFC4145], Section 4, and clarified for use in DTLS-SRTP scenarios in [RFC5763], Section 5. The role value in the offer MUST be "actpass".
- o An "a=rtcp-mux" line, as specified in [RFC5761], Section 5.1.1.
- o An "a=rtcp-rsize" line, as specified in [RFC5506], Section 5.

- o For each supported RTP header extension, an "a=extmap" line, as specified in [RFC5285], Section 5. The list of header extensions that SHOULD/MUST be supported is specified in [I-D.ietf-rtcweb-rtp-usage], Section 5.2. Any header extensions that require encryption MUST be specified as indicated in [RFC6904], Section 4.
- o For each supported RTCP feedback mechanism, an "a=rtcp-fb" mechanism, as specified in [RFC4585], Section 4.2. The list of RTCP feedback mechanisms that SHOULD/MUST be supported is specified in [I-D.ietf-rtcweb-rtp-usage], Section 5.1.
- o An "a=ssrc" line, as specified in [RFC5576], Section 4.1, indicating the SSRC to be used for sending media, along with the mandatory "cname" source attribute, as specified in Section 6.1, indicating the CNAME for the source. The CNAME must be generated in accordance with draft-rescorla-random-cname-00. [OPEN ISSUE: How are CNAMEs specified for MSTs? Are they randomly generated for each MediaStream? If so, can two MediaStreams be synced?]
- o If RTX is supported for this media type, another "a=ssrc" line with the RTX SSRC, and an "a=ssrc-group" line, as specified in [RFC5576], section 4.2, with semantics set to "FID" and including the primary and RTX SSRCs.
- o If FEC is supported for this media type, another "a=ssrc" line with the FEC SSRC, and an "a=ssrc-group" line, as specified in [RFC5576], section 4.2, with semantics set to "FEC" and including the primary and FEC SSRCs.
- o [OPEN ISSUE: Handling of a=imageattr]
- o [TODO: bundle-only]

Lastly, if a data channel has been created, a m= section MUST be generated for data. The <media> field MUST be set to "application" and the <proto> field MUST be set to "DTLS/SCTP", as specified in [I-D.ietf-mmusic-sctp-sdp], Section 3; the "fmt" value MUST be set to the SCTP port number, as specified in Section 4.1.

Within the data m= section, the "a=mid", "a=ice-ufrag", "a=ice-passwd", "a=ice-options", "a=candidate", "a=fingerprint", and "a=setup" lines MUST be included as mentioned above, along with an "a=sctpmap" line referencing the SCTP port number and specifying the application protocol indicated in [I-D.ietf-rtcweb-data-protocol]. [OPEN ISSUE: the -00 of this document is missing this information.]

Once all m= sections have been generated, a session-level "a=group" attribute MUST be added as specified in [RFC5888]. This attribute MUST have semantics "BUNDLE", and identify the m= sections to be bundled. [OPEN ISSUE: Need to determine exactly how this decision is made.]

Attributes that are common between all m= sections MAY be moved to session-level, if desired.

Attributes other than the ones specified above MAY be included, except for the following attributes which are specifically incompatible with the requirements of [I-D.ietf-rtcweb-rtp-usage], and MUST NOT be included:

- o "a=crypto"
- o "a=key-mgmt"
- o "a=ice-lite"

Note that when BUNDLE is used, any additional attributes that are added MUST follow the advice in [I-D.nandakumar-mmusic-sdp-mux-attributes] on how those attributes interact with BUNDLE.

5.2.2. Subsequent Offers

When createOffer is called a second (or later) time, or is called after a local description has already been installed, the processing is somewhat different than for an initial offer.

If the initial offer was not applied using setLocalDescription, meaning the PeerConnection is still in the "stable" state, the steps for generating an initial offer should be followed, subject to the following restriction:

- o The fields of the "o=" line MUST stay the same except for the <session-version> field, which MUST increment if the session description changes in any way, including the addition of ICE candidates.

If the initial offer was applied using setLocalDescription, but an answer from the remote side has not yet been applied, meaning the PeerConnection is still in the "local-offer" state, an offer is generated by following the steps in the "stable" state above, along with these exceptions:

- o The "s=" and "t=" lines MUST stay the same.

- o Each "a=mid" line MUST stay the same.
- o Each "a=ice-ufrag" and "a=ice-pwd" line MUST stay the same.
- o For MediaStreamTracks that are still present, the "a=msid", "a=ssrc", and "a=ssrc-group" lines MUST stay the same.
- o If any MediaStreamTracks have been removed, either through the removeStream method or by removing them from an added MediaStream, their m= sections MUST be marked as recvonly by changing the value of the [RFC3264] directional attribute to "a=recvonly". The "a=msid", "a=ssrc", and "a=ssrc-group" lines MUST be removed from the associated m= sections.

If the initial offer was applied using setLocalDescription, and an answer from the remote side has been applied using setRemoteDescription, meaning the PeerConnection is in the "remote-pranswer" or "stable" states, an offer is generated based on the negotiated session descriptions by following the steps mentioned for the "local-offer" state above, along with these exceptions: [OPEN ISSUE: should this be permitted in the remote-pranswer state?]

- o If a m= section was rejected, i.e. has had its port set to zero in either the local or remote description, it MUST remain rejected and have a zero port in the new offer, as indicated in RFC3264, Section 5.1.
- o If a m= section exists in the current local description, but has its state set to inactive or recvonly, and a new MediaStreamTrack is added, the previously existing m= section MUST be recycled instead of creating a new m= section. [OPEN ISSUE: Nail down exactly what this means. Should the codecs remain the same? (No.) Should ICE restart? (No.) Can the "a=mid" attribute be changed? (Yes?)]
- o If a m= section exists in the current local description, but does not have an associated MediaStreamTrack (i.e. it is inactive or recvonly), a corresponding m= section MUST be generated in the new offer, but without "a=msid", "a=ssrc", or "a=ssrc-group" attributes, and the appropriate directional attribute must be specified.

In addition, for each previously existing, non-rejected m= section in the new offer, the following adjustments are made based on the contents of the corresponding m= section in the current remote description:

- o The m= line and corresponding "a=rtpmap" and "a=fmtp" lines MUST only include codecs present in the remote description.
- o The RTP header extensions MUST only include those that are present in the remote description.
- o The RTCP feedback extensions MUST only include those that are present in the remote description.
- o The "a=rtcp-mux" line MUST only be added if present in the remote description.
- o The "a=rtcp-rsize" line MUST only be added if present in the remote description.

5.2.3. Constraints Handling

The `createOffer` method takes as a parameter a `MediaConstraints` object. Special processing is performed when generating a SDP description if the following constraints are present.

5.2.3.1. OfferToReceiveAudio

If the "OfferToReceiveAudio" constraint is specified, with a value of "true", the offer MUST include a non-rejected m= section with media type "audio", even if no audio `MediaStreamTrack` has been added to the `PeerConnection`. This allows the offerer to receive audio even when not sending it; accordingly, the directional attribute on the audio m= section MUST be set to `recvonly`. If this constraint is specified when an audio `MediaStreamTrack` has already been added to the `PeerConnection`, or a non-rejected m= section with media type "audio" previously existed, it has no effect.

5.2.3.2. OfferToReceiveVideo

If the "OfferToReceiveAudio" constraint is specified, with a value of "true", the offer MUST include a m= section with media type "video", even if no video `MediaStreamTrack` has been added to the `PeerConnection`. This allows the offerer to receive video even when not sending it; accordingly, the directional attribute on the video m= section MUST be set to `recvonly`. If this constraint is specified when an video `MediaStreamTrack` has already been added to the `PeerConnection`, or a non-rejected m= section with media type "video" previously existed, it has no effect.

5.2.3.3. VoiceActivityDetection

If the "VoiceActivityDetection" constraint is specified, with a value of "true", the offer MUST indicate support for silence suppression by including comfort noise ("CN") codecs for each supported clock rate, as specified in [RFC3389], Section 5.1.

5.2.3.4. IceRestart

If the "IceRestart" constraint is specified, with a value of "true", the offer MUST indicate an ICE restart by generating new ICE ufrag and pwd attributes, as specified in RFC5245, Section 9.1.1.1. If this constraint is specified on an initial offer, it has no effect (since a new ICE ufrag and pwd are already generated).

5.3. Generating an Answer

When createAnswer is called, a new SDP description must be created that is compatible with the supplied remote description as well as the requirements specified in [I-D.ietf-rtcweb-rtp-usage]. The exact details of this process are explained below.

5.3.1. Initial Answers

When createAnswer is called for the first time after a remote description has been provided, the result is known as the initial answer. If no remote description has been installed, an answer cannot be generated, and an error MUST be returned.

Note that the remote description SDP may not have been created by a JSEP endpoint and may not conform to all the requirements listed in Section 5.2. For many cases, this is not a problem. However, if any mandatory SDP attributes are missing, or functionality listed as mandatory-to-use above is not present, this MUST be treated as an error, and MUST cause the affected m= sections to be marked as rejected.

The first step in generating an initial answer is to generate session-level attributes. The process here is identical to that indicated in the Initial Offers section above.

The next step is to generate m= sections for each m= section that is present in the remote offer, as specified in [RFC3264], Section 6. For the purposes of this discussion, any session-level attributes in the offer that are also valid as media-level attributes SHALL be considered to be present in each m= section.

If any of the offered m= sections have been rejected, by stopping the associated remote MediaStreamTrack, the corresponding m= section in the answer MUST be marked as rejected by setting the port in the m=

line to zero, as indicated in [RFC3264], Section 6., and processing continues with the next m= section.

For each non-rejected m= section of a given media type, if there is a local `MediaStreamTrack` of the specified type which has been added to the `PeerConnection` via `addStream` and not yet associated with a m= section, the `MediaStreamTrack` is associated with the m= section at this time. If there are more m= sections of a certain type than `MediaStreamTracks`, some m= sections will not have an associated `MediaStreamTrack`. If there are more `MediaStreamTracks` of a certain type than m= sections, only the first N `MediaStreamTracks` will be able to be associated in the constructed answer. The remainder will need to be associated in a subsequent offer.

Each m= section should then generated as specified in [RFC3264], Section 6.1. Because use of DTLS is mandatory, the <proto> field MUST be set to "UDP/TLS/RTP/SAVPF". If the offer supports BUNDLE, all m= sections to be BUNDLED must use the same ICE credentials and candidates; all m= sections not being BUNDLED must use unique ICE credentials and candidates. Each m= section MUST include the following:

- o If present in the offer, an "a=mid" line, as specified in [RFC5888], Section 9.1. The "mid" value MUST match that specified in the offer.
- o If a local `MediaStreamTrack` has been associated, an "a=msid" line, as specified in [I-D.ietf-mmusic-msid], Section 2.
- o [OPEN ISSUE: Use of App Token versus stream-correlator]
- o If a local `MediaStreamTrack` has been associated, an "a=sendrecv" line, as specified in [RFC3264], Section 6.1. If no local `MediaStreamTrack` has been associated, an "a=recvonly" line. [TODO: handle non-sendrecv offered m= sections]
- o For each supported codec that is present in the offer, "a=rtpmap" and "a=fmtp" lines, as specified in [RFC4566], Section 6, and [RFC3264], Section 6.1. For audio, the codecs specified in [I-D.ietf-rtcweb-audio], Section 3, MUST be supported. Note that for simplicity, the answerer MAY use different payload types for codecs than the offerer, as it is not prohibited by Section 6.1.

- o If "rtx" is present in the offer, for each primary codec where RTP retransmission should be used, a corresponding "a=rtpmap" line indicating "rtx" with the clock rate of the primary codec and an "a=fmtp" line that references the payload type for the primary codec, as specified in [RFC4588], Section 8.1.
- o For each supported FEC mechanism that is present in the offer, a corresponding "a=rtpmap" line indicating the desired FEC codec.
- o "a=ice-ufrag" and "a=ice-passwd" lines, as specified in [RFC5245], Section 15.4.
- o If the "trickle" ICE option is present in the offer, an "a=ice-options" line, with the "trickle" option, as specified in [I-D.ietf-mmusic-trickle-ice], Section 4.
- o For each candidate that has been gathered during the most recent gathering phase, an "a=candidate" line, as specified in [RFC5245], Section 4.3., paragraph 3.
- o For the current default candidate, a "c=" line, as specific in [RFC5245], Section 4.3., paragraph 6. If no candidates have been gathered yet, the default candidate should be set to the 'null' value defined in [I-D.ietf-mmusic-trickle-ice], Section 5.1.
- o An "a=fingerprint" line, as specified in [RFC4572], Section 5; the algorithm used for the fingerprint MUST match that used in the certificate signature.
- o An "a=setup" line, as specified in [RFC4145], Section 4, and clarified for use in DTLS-SRTP scenarios in [RFC5763], Section 5. The role value in the answer MUST be "active" or "passive"; the "active" role is RECOMMENDED.
- o If present in the offer, an "a=rtcp-mux" line, as specified in [RFC5761], Section 5.1.1.
- o If present in the offer, an "a=rtcp-rsize" line, as specified in [RFC5506], Section 5.
- o For each supported RTP header extension that is present in the offer, an "a=extmap" line, as specified in [RFC5285], Section 5. The list of header extensions that SHOULD/MUST be supported is specified in [I-D.ietf-rtcweb-rtp-usage], Section 5.2. Any header extensions that require encryption MUST be specified as indicated in [RFC6904], Section 4.

- o For each supported RTCP feedback mechanism that is present in the offer, an "a=rtcp-fb" mechanism, as specified in [RFC4585], Section 4.2. The list of RTCP feedback mechanisms that SHOULD/MUST be supported is specified in [I-D.ietf-rtcweb-rtp-usage], Section 5.1.
- o If a local MediaStreamTrack has been associated, an "a=ssrc" line, as specified in [RFC5576], Section 4.1, indicating the SSRC to be used for sending media.
- o If a local MediaStreamTrack has been associated, and RTX has been negotiated for this m= section, another "a=ssrc" line with the RTX SSRC, and an "a=ssrc-group" line, as specified in [RFC5576], section 4.2, with semantics set to "FID" and including the primary and RTX SSRCs.
- o If a local MediaStreamTrack has been associated, and FEC has been negotiated for this m= section, another "a=ssrc" line with the FEC SSRC, and an "a=ssrc-group" line, as specified in [RFC5576], section 4.2, with semantics set to "FEC" and including the primary and FEC SSRCs.
- o [OPEN ISSUE: Handling of a=imageattr]
- o [TODO: bundle-only]

If a data channel m= section has been offered, a m= section MUST also be generated for data. The <media> field MUST be set to "application" and the <proto> field MUST be set to "DTLS/SCTP", as specified in [I-D.ietf-mmusic-sctp-sdp], Section 3; the "fmt" value MUST be set to the SCTP port number, as specified in Section 4.1.

Within the data m= section, the "a=mid", "a=ice-ufrag", "a=ice-passwd", "a=ice-options", "a=candidate", "a=fingerprint", and "a=setup" lines MUST be included as mentioned above, along with an "a=sctpmap" line referencing the SCTP port number and specifying the application protocol indicated in [I-D.ietf-rtcweb-data-protocol]. [OPEN ISSUE: the -00 of this document is missing this information.]

[TODO: processing of BUNDLE group]

Attributes that are common between all m= sections MAY be moved to session-level, if desired.

The attributes prohibited in the creation of offers are also prohibited in the creation of answers.

- 5.3.2. Subsequent Answers
- 5.3.3. Constraints Handling
- 5.4. Parsing an Offer
- 5.5. Parsing an Answer
- 5.6. Applying a Local Description
- 5.7. Applying a Remote Description
- 6. Configurable SDP Parameters

Note: This section is still very early and is likely to significantly change as we get a better understanding of a) the use cases for this b) the implications at the protocol level c) feedback from implementors on what they can do.

The following elements of the SDP media description MUST NOT be changed between the createOffer and the setLocalDescription, since they reflect transport attributes that are solely under browser control, and the browser MUST NOT honor an attempt to change them:

- o The number, type and port number of m-lines.
- o The generated ICE credentials (a=ice-ufrag and a=ice-pwd).
- o The set of ICE candidates and their parameters (a=candidate).

The following modifications, if done by the browser to a description between createOffer/createAnswer and the setLocalDescription, MUST be honored by the browser:

- o Remove or reorder codecs (m=)

The following parameters may be controlled by constraints passed into createOffer/createAnswer. As an open issue, these changes may also be performed by manipulating the SDP returned from createOffer/createAnswer, as indicated above, as long as the capabilities of the endpoint are not exceeded (e.g. asking for a resolution greater than what the endpoint can encode):

- o disable BUNDLE (a=group)
- o disable RTCP mux (a=rtcp-mux)
- o change send resolution or frame rate

- o change desired recv resolution or frame rate
- o change maximum total bandwidth (b=) [OPEN ISSUE: need to clarify if this is CT or AS - see section 5.8 of [RFC4566]]
- o remove desired AVPF mechanisms (a=rtcp-fb)
- o remove RTP header extensions (a=extmap)
- o change media send/recv state (a=sendonly/recvonly/inactive)

For example, an application could implement call hold by adding an a=inactive attribute to its local description, and then applying and signaling that description.

The application can also modify the SDP to reduce the capabilities in the offer it sends to the far side in any way the application sees fit, as long as it is a valid SDP offer and specifies a subset of what the browser is expecting to do.

As always, the application is solely responsible for what it sends to the other party, and all incoming SDP will be processed by the browser to the extent of its capabilities. It is an error to assume that all SDP is well-formed; however, one should be able to assume that any implementation of this specification will be able to process, as a remote offer or answer, unmodified SDP coming from any other implementation of this specification.

7. Security Considerations

The intent of the WebRTC protocol suite is to provide an environment that is securable by default: all media is encrypted, keys are exchanged in a secure fashion, and the Javascript API includes functions that can be used to verify the identity of communication partners.

8. IANA Considerations

This document requires no actions from IANA.

9. Acknowledgements

Significant text incorporated in the draft as well and review was provided by Harald Alvestrand and Suhas Nandakumar. Dan Burnett, Neil Stratford, Eric Rescorla, Anant Narayanan, Andrew Hutton, Richard Ejzak, and Adam Bergkvist 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.

10. References

10.1. Normative References

- [I-D.ietf-mmusic-msid]
Alvestrand, H., "Cross Session Stream Identification in the Session Description Protocol", draft-ietf-mmusic-msid-01 (work in progress), August 2013.
- [I-D.ietf-mmusic-sctp-sdp]
Loreto, S. and G. Camarillo, "Stream Control Transmission Protocol (SCTP)-Based Media Transport in the Session Description Protocol (SDP)", draft-ietf-mmusic-sctp-sdp-04 (work in progress), June 2013.
- [I-D.ietf-mmusic-sdp-bundle-negotiation]
Holmberg, C., Alvestrand, H., and C. Jennings, "Multiplexing Negotiation Using Session Description Protocol (SDP) Port Numbers", draft-ietf-mmusic-sdp-bundle-negotiation-04 (work in progress), June 2013.
- [I-D.ietf-rtcweb-audio]
Valin, J. and C. Bran, "WebRTC Audio Codec and Processing Requirements", draft-ietf-rtcweb-audio-02 (work in progress), August 2013.
- [I-D.ietf-rtcweb-rtp-usage]
Perkins, C., Westerlund, M., and J. Ott, "Web Real-Time Communication (WebRTC): Media Transport and Use of RTP", draft-ietf-rtcweb-rtp-usage-09 (work in progress), September 2013.
- [I-D.nandakumar-mmusic-sdp-mux-attributes]
Nandakumar, S., "A Framework for SDP Attributes when Multiplexing", draft-nandakumar-mmusic-sdp-mux-attributes-03 (work in progress), July 2013.
- [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.
- [RFC4145] Yon, D. and G. Camarillo, "TCP-Based Media Transport in the Session Description Protocol (SDP)", RFC 4145, September 2005.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC4572] Lennox, J., "Connection-Oriented Media Transport over the Transport Layer Security (TLS) Protocol in the Session Description Protocol (SDP)", RFC 4572, July 2006.
- [RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, July 2006.
- [RFC5124] Ott, J. and E. Carrara, "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", RFC 5124, February 2008.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC5285] Singer, D. and H. Desineni, "A General Mechanism for RTP Header Extensions", RFC 5285, July 2008.
- [RFC5761] Perkins, C. and M. Westerlund, "Multiplexing RTP Data and Control Packets on a Single Port", RFC 5761, April 2010.
- [RFC5888] Camarillo, G. and H. Schulzrinne, "The Session Description Protocol (SDP) Grouping Framework", RFC 5888, June 2010.
- [RFC6904] Lennox, J., "Encryption of Header Extensions in the Secure Real-time Transport Protocol (SRTP)", RFC 6904, April 2013.

10.2. Informative References

- [I-D.ietf-mmusic-trickle-ice]
Ivov, E., Rescorla, E., and J. Uberti, "Trickle ICE: Incremental Provisioning of Candidates for the Interactive Connectivity Establishment (ICE) Protocol", draft-ietf-mmusic-trickle-ice-00 (work in progress), March 2013.
- [I-D.ietf-rtcweb-data-protocol]
Jesup, R., Loreto, S., and M. Tuexen, "WebRTC Data Channel Protocol", draft-ietf-rtcweb-data-protocol-04 (work in progress), February 2013.
- [I-D.jennings-rtcweb-signaling]
Jennings, C., Rosenberg, J., and R. Jesup, "RTCWeb Offer/Answer Protocol (ROAP)", draft-jennings-rtcweb-signaling-01 (work in progress), October 2011.
- [I-D.nandakumar-rtcweb-sdp]
Nandakumar, S. and C. Jennings, "SDP for the WebRTC", draft-nandakumar-rtcweb-sdp-02 (work in progress), July 2013.
- [RFC3389] Zopf, R., "Real-time Transport Protocol (RTP) Payload for Comfort Noise (CN)", RFC 3389, September 2002.
- [RFC3556] Casner, S., "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", RFC 3556, July 2003.
- [RFC3960] Camarillo, G. and H. Schulzrinne, "Early Media and Ringing Tone Generation in the Session Initiation Protocol (SIP)", RFC 3960, December 2004.
- [RFC4568] Andreasen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", RFC 4568, July 2006.
- [RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R. Hakenberg, "RTP Retransmission Payload Format", RFC 4588, July 2006.
- [RFC5506] Johansson, I. and M. Westerlund, "Support for Reduced-Size Real-Time Transport Control Protocol (RTCP): Opportunities and Consequences", RFC 5506, April 2009.

- [RFC5576] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", RFC 5576, June 2009.
- [RFC5763] Fischl, J., Tschofenig, H., and E. Rescorla, "Framework for Establishing a Secure Real-time Transport Protocol (SRTP) Security Context Using Datagram Transport Layer Security (DTLS)", RFC 5763, May 2010.
- [RFC5764] McGrew, D. and E. Rescorla, "Datagram Transport Layer Security (DTLS) Extension to Establish Keys for the Secure Real-time Transport Protocol (SRTP)", RFC 5764, May 2010.
- [W3C.WD-webrtc-20111027]
Bergkvist, A., Burnett, D., Narayanan, A., and C. Jennings, "WebRTC 1.0: Real-time Communication Between Browsers", World Wide Web Consortium WD WD-webrtc-20111027, October 2011,
<<http://www.w3.org/TR/2011/WD-webrtc-20111027>>.

Appendix A. JSEP Implementation Examples

A.1. 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. For clarity and simplicity, the createOffer/createAnswer calls are assumed to be synchronous in these examples, whereas the actual APIs are async.

A.1.1. 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);
OffererUA->OffererJS: iceCallback(candidate);
OffererJS->OffererUA: offer = pc.createOffer(null);
OffererJS->OffererUA: pc.setLocalDescription("offer", offer);
OffererJS->AnswererJS: { "type": "OFFER", "sdp": offer }

// OFFER arrives at Answerer
AnswererJS->AnswererUA: pc = new PeerConnection();
AnswererJS->AnswererUA: pc.setRemoteDescription("offer", msg.sdp);
AnswererUA->AnswererJS: onaddstream(remoteStream);
AnswererUA->OffererUA: iceCallback(candidate);
```

```
// Answerer accepts call
AnswererJS->AnswererUA: pc.addStream(localStream, null);
AnswererJS->AnswererUA: answer = pc.createAnswer(msg.sdp, null);
AnswererJS->AnswererUA: pc.setLocalDescription("answer", answer);
AnswererJS->OffererJS:  {"type":"ANSWER","sdp":answer }

// ANSWER arrives at Offerer
OffererJS->OffererUA:   pc.setRemoteDescription("answer", answer);
OffererUA->OffererJS:   onaddstream(remoteStream);

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

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

A.1.2. 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("offer", offer);
OffererJS:             xmpp = createSessionInitiate(offer);
OffererJS->AnswererJS:  <jingle action="session-initiate"/>

OffererJS->OffererUA:   pc.startIce();
OffererUA->OffererJS:   onicecandidate(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("offer", offer);
AnswererUA->AnswererJS: onaddstream(remoteStream);

// transport-infos arrive at Answerer
AnswererJS->AnswererUA: candidate = parseTransportInfo(xmpp);
AnswererJS->AnswererUA: pc.addIceCandidate(candidate);
AnswererUA->AnswererJS: onicecandidate(cand)
AnswererJS:             createTransportInfo(cand);
AnswererJS->OffererJS:  <jingle action="transport-info"/>
```

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

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

// session-accept arrives at Offerer
OffererJS:           answer = parseSessionAccept(xmpp);
OffererJS->OffererUA:  pc.setRemoteDescription("answer", answer);
OffererUA->OffererJS:  onaddstream(remoteStream);

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

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

A.1.1.3. 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("offer", offer);
OffererJS->AnswererJS: <jingle action="content-add"/>

// content-add arrives at Answerer
AnswererJS:           offer = parseContentAdd(xmpp);
AnswererJS->AnswererUA: pc.setRemoteDescription("offer", offer);
AnswererJS->AnswererUA: answer = pc.createAnswer(offer, null);
AnswererJS->AnswererUA: pc.setLocalDescription("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("answer", answer);
```

A.1.4. 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("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("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("answer", offer);

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

A.1.5. 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);
```

```
OffererUA->OffererJS:  onicecandidate(candidate);
OffererJS->OffererUA:  offer = pc.createOffer(null);
OffererJS->OffererUA:  pc.setLocalDescription("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("offer", offer);
AnswererUA->AnswererJS: onaddstream(remoteStream);
AnswererUA->OffererUA:  onicecandidate(candidate);

// Answerer accepts call
AnswererJS->AnswererUA: pc.addStream(localStream, null);
AnswererJS->AnswererUA: answer = pc.createAnswer(offer, null);
AnswererJS:             sip = createResponse(200, answer);
AnswererJS->AnswererUA: pc.setLocalDescription("answer", answer);
AnswererJS->OffererJS:  200 OK w/ SDP

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

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

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

A.1.6. Handling early media (e.g. 1-800-GO 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);
OffererUA->OffererJS:  onicecandidate(candidate);
OffererJS->OffererUA:  offer = pc.createOffer(null);
OffererJS->OffererUA:  pc.setLocalDescription("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("pranswer", answer);
OffererUA->OffererJS: onaddstream(remoteStream);

// ICE Completes (at Offerer)
OffererUA->AnswererUA: Media

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

```

A.2. Example Session Descriptions

A.2.1. createOffer

This SDP shows a typical initial offer, created by `createOffer` for a `PeerConnection` with a single audio `MediaStreamTrack`, a single video `MediaStreamTrack`, and a single data channel. Host candidates have also already been gathered. Note some lines have been broken into two lines for formatting reasons.

```

v=0
o=- 4962303333179871722 1 IN IP4 0.0.0.0
s=-
t=0 0
a=msid-semantic:WMS
a=group:BUNDLE audio video data
m=audio 56500 UDP/TLS/RTP/SAVPF 111 0 8 126
c=IN IP4 192.0.2.1
a=rtcp:56501 IN IP4 192.0.2.1
a=candidate:3348148302 1 udp 2113937151 192.0.2.1 56500
      typ host generation 0
a=candidate:3348148302 2 udp 2113937151 192.0.2.1 56501
      typ host generation 0
a=ice-ufrag:ETEnlv9DoTMB9J4r
a=ice-pwd:OtSK0WpNtpUjkY4+86js7ZQl
a=ice-options:trickle
a=mid:audio
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level
a=sendrecv
a=rtcp-mux
a=rtcp-rsize
a=fingerprint:sha-256
      19:E2:1C:3B:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73:04
      BB:05:2F:70:9F:04:A9:0E:05:E9:26:33:E8:70:88:A2
a=setup:actpass

```

```
a=rtpmap:111 opus/48000/2
a=fmtp:111 minptime=10
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:126 telephone-event/8000
a=maxptime:60
a=ssrc:1732846380 cname:EocUGlf0fcg/yvY7
a=msid:47017fee-b6c1-4162-929c-a25110252400
      f83006c5-a0ff-4e0a-9ed9-d3e6747be7d9
m=video 56502 UDP/TLS/RTP/SAVPF 100 115 116 117
c=IN IP4 192.0.2.1
a=rtcp:56503 IN IP4 192.0.2.1
a=candidate:3348148302 1 udp 2113937151 192.0.2.1 56502
      typ host generation 0
a=candidate:3348148302 2 udp 2113937151 192.0.2.1 56503
      typ host generation 0
a=ice-ufrag:BGKkWnG5GmiUpdIV
a=ice-pwd:mqyWSAjvtKwTGnvHPztQ9mIf
a=ice-options:trickle
a=mid:video
a=extmap:2 urn:ietf:params:rtp-hdext:toffset
a=extmap:3 http://www.webrtc.org/experiments/rtp-hdext/abs-send-time
a=sendrecv
a=rtcp-mux
a=rtcp-rsize
a=fingerprint:sha-256
      19:E2:1C:3B:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73:04
      :BB:05:2F:70:9F:04:A9:0E:05:E9:26:33:E8:70:88:A2
a=setup:actpass
a=rtpmap:100 VP8/90000
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 goog-remb
a=rtpmap:115 rtx/90000
a=fmtp:115 apt=100
a=rtpmap:116 red/90000
a=rtpmap:117 ulpfec/90000
a=ssrc:1366781083 cname:EocUGlf0fcg/yvY7
a=ssrc:1366781084 cname:EocUGlf0fcg/yvY7
a=ssrc:1366781085 cname:EocUGlf0fcg/yvY7
a=ssrc-group:FID 1366781083 1366781084
a=ssrc-group:FEC 1366781083 1366781085
a=msid:61317484-2ed4-49d7-9eb7-1414322a7aae
      f30bdb4a-5db8-49b5-bcdc-e0c9a23172e0
m=application 56504 DTLS/SCTP 5000
c=IN IP4 192.0.2.1
a=candidate:3348148302 1 udp 2113937151 192.0.2.1 56504
      typ host generation 0
```

```
a=ice-ufrag:VD5v2BnbZm3mgP3d
a=ice-pwd:+Jlkuox+VVIUDqxcfIDuTZMH
a=ice-options:trickle
a=mid:data
a=fingerprint:sha-256 19:E2:1C:3B:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73:04
                        :BB:05:2F:70:9F:04:A9:0E:05:E9:26:33:E8:70:88:A2
a=setup:actpass
a=sctpmap:5000 webrtc-datachannel 16
```

A.2.2. createAnswer

This SDP shows a typical initial answer to the above offer, created by createAnswer for a PeerConnection with a single audio MediaStreamTrack, a single video MediaStreamTrack, and a single data channel. Host candidates have also already been gathered. Note some lines have been broken into two lines for formatting reasons.

```
v=0
o=- 6729291447651054566 1 IN IP4 0.0.0.0
s=-
t=0 0
a=msid-semantic:WMS
a=group:BUNDLE audio video data
m=audio 20000 UDP/TLS/RTP/SAVPF 111 0 8 126
c=IN IP4 192.0.2.2
a=candidate:2299743422 1 udp 2113937151 192.0.2.2 20000
                typ host generation 0
a=ice-ufrag:6sFvz2gdLkEwjZEr
a=ice-pwd:cOTZKZNVlO9RSGsEGM63JXT2
a=fingerprint:sha-256 6B:8B:F0:65:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B:35
                        :DC:B8:5F:64:1A:24:C2:43:F0:A1:58:D0:A1:2C:19:08
a=setup:active
a=mid:audio
a=extmap:1 urn:ietf:params:rtp-hdrext:ssrc-audio-level
a=sendrecv
a=rtcp-mux
a=rtpmap:111 opus/48000/2
a=fmtp:111 minptime=10
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:126 telephone-event/8000
a=maxptime:60
a=ssrc:3429951804 cname:Q/NWslao1HmN4Xa5
a=msid:PI39StLS8W7ZbQ1lsJsWUXkr3Zf12fJUvzQ1
                PI39StLS8W7ZbQ1lsJsWUXkr3Zf12fJUvzQ1a0
m=video 20000 UDP/TLS/RTP/SAVPF 100 115 116 117
c=IN IP4 192.0.2.2
```

```
a=candidate:2299743422 1 udp 2113937151 192.0.2.2 20000
      typ host generation 0
a=ice-ufrag:6sFvz2gdLkEwjZEr
a=ice-pwd:cOTZKZNVlO9RSGsEGM63JXT2
a=fingerprint:sha-256 6B:8B:F0:65:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B:35
      :DC:B8:5F:64:1A:24:C2:43:F0:A1:58:D0:A1:2C:19:08

a=setup:active
a=mid:video
a=extmap:2 urn:ietf:params:rtp-hdrext:toffset
a=extmap:3 http://www.webrtc.org/experiments/rtp-hdrext/abs-send-time
a=sendrecv
a=rtcp-mux
a=rtpmap:100 VP8/90000
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 goog-remb
a=rtpmap:115 rtx/90000
a=fmtp:115 apt=100
a=rtpmap:116 red/90000
a=rtpmap:117 ulpfec/90000
a=ssrc:3229706345 cname:Q/NWslao1HmN4Xa5
a=ssrc:3229706346 cname:Q/NWslao1HmN4Xa5
a=ssrc:3229706347 cname:Q/NWslao1HmN4Xa5
a=ssrc-group:FID 3229706345 3229706346
a=ssrc-group:FEC 3229706345 3229706347
a=msid:PI39StLS8W7ZbQ1lsJsWUXkr3Zf12fJUvzQ1
      PI39StLS8W7ZbQ1lsJsWUXkr3Zf12fJUvzQ1v0
m=application 20000 DTLS/SCTP 5000
c=IN IP4 192.0.2.2
a=candidate:2299743422 1 udp 2113937151 192.0.2.2 20000
      typ host generation 0
a=ice-ufrag:6sFvz2gdLkEwjZEr
a=ice-pwd:cOTZKZNVlO9RSGsEGM63JXT2
a=fingerprint:sha-256 6B:8B:F0:65:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B:35
      :DC:B8:5F:64:1A:24:C2:43:F0:A1:58:D0:A1:2C:19:08

a=setup:active
a=mid:data
a=sctpmap:5000 webrtc-datachannel 16
```

A.2.3. Call Flows

Example SDP for WebRTC call flows can be found in [I-D.nandakumar-rtcweb-sdp]. [TODO: should these call flows be merged into this section?]

Appendix B. Change log

Changes in draft-05:

- o Fixed several issues identified in the createOffer/Answer sections during document review.
- o Updated references.

Changes in draft-04:

- o Filled in sections on createOffer and createAnswer.
- o Added SDP examples.
- o Fixed references.

Changes in draft-03:

- o Added text describing relationship to W3C specification

Changes in draft-02:

- o Converted from nroff
- o Removed comparisons to old approaches abandoned by the working group
- o Removed stuff that has moved to W3C specification
- o Align SDP handling with W3C draft
- o Clarified section on forking.

Changes in draft-01:

- o Added diagrams for architecture and state machine.
- o Added sections on forking and rehydration.
- o Clarified meaning of "pranswer" and "answer".
- o Reworked how ICE restarts and media directions are controlled.
- o Added list of parameters that can be changed in a description.
- o Updated suggested API and examples to match latest thinking.
- o Suggested API and examples have been moved to an appendix.

Changes in draft -00:

- o Migrated from draft-uberti-rtcweb-jsep-02.

Authors' Addresses

Justin Uberti
Google
747 6th Ave S
Kirkland, WA 98033
USA

Email: justin@uberti.name

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

Email: fluffy@iii.ca