

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 4, 2017

J. Uberti
Google
October 31, 2016

WebRTC Forward Error Correction Requirements
draft-ietf-rtcweb-fec-04

Abstract

This document provides information and requirements for how Forward Error Correction (FEC) should be used by WebRTC applications.

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 May 4, 2017.

Copyright Notice

Copyright (c) 2016 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	2
2. Terminology	2
3. Types of FEC	2
3.1. Separate FEC Stream	3
3.2. Redundant Encoding	3
3.3. Codec-Specific In-band FEC	3
4. FEC for Audio Content	4
4.1. Recommended Mechanism	4
4.2. Negotiating Support	4
5. FEC for Video Content	5
5.1. Recommended Mechanism	5
5.2. Negotiating Support	5
6. FEC for Application Content	6
7. Implementation Requirements	6
8. Adaptive Use of FEC	6
9. Security Considerations	7
10. IANA Considerations	7
11. Acknowledgements	7
12. References	7
12.1. Normative References	7
12.2. Informative References	8
Appendix A. Change log	8
Author's Address	10

1. Introduction

In situations where packet loss is high, or perfect media quality is essential, Forward Error Correction (FEC) can be used to proactively recover from packet losses. This specification provides guidance on which FEC mechanisms to use, and how to use them, for WebRTC client implementations.

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. Types of FEC

By its name, FEC describes the sending of redundant information in an outgoing packet stream so that information can still be recovered even in the face of packet loss. There are multiple ways in which this can be accomplished; this section enumerates the various mechanisms and describes their tradeoffs.

3.1. Separate FEC Stream

This approach, as described in [RFC5956], Section 4.3, sends FEC packets as an independent SSRC-multiplexed stream, with its own SSRC and payload type. While by far the most flexible, each FEC packet will have its own IP+UDP+RTP+FEC header, leading to additional overhead of the FEC stream.

3.2. Redundant Encoding

This approach, as described in [RFC2198], allows for redundant data to be piggybacked on an existing primary encoding, all in a single packet. This redundant data may be an exact copy of a previous packet, or for codecs that support variable-bitrate encodings, possibly a smaller, lower-quality representation. In certain cases, the redundant data could include multiple prior packets.

Since there is only a single set of packet headers, this approach allows for a very efficient representation of primary + redundant data. However, this savings is only realized when the data all fits into a single packet (i.e. the size is less than a MTU). As a result, this approach is generally not useful for video content.

3.3. Codec-Specific In-band FEC

Some audio codecs, notably Opus [RFC6716] and AMR [RFC4867] support their own in-band FEC mechanism, where redundant data is included in the codec payload.

For Opus, packets deemed as important are re-encoded at a lower bitrate and added to the subsequent packet, allowing partial recovery of a lost packet. This scheme is fairly efficient; experiments performed indicate that when Opus FEC is used, the overhead imposed is about 20-30%, depending on the amount of protection needed. Note that this mechanism can only carry redundancy information for the immediately preceding packet; as such the decoder cannot fully recover multiple consecutive lost packets, which can be a problem on wireless networks. See [RFC6716], Section 2.1.7 for complete details.

For AMR/AMR-WB, packets can contain copies or lower-quality encodings of multiple prior audio frames. This mechanism is similar to the [RFC2198] mechanism described above, but as it adds no additional framing, it can be slightly more efficient. See [RFC4867], Section 3.7.1 for details on this mechanism.

4. FEC for Audio Content

The following section provides guidance on how to best use FEC for transmitting audio data. As indicated in Section 8 below, FEC should only be activated if network conditions warrant it, or upon explicit application request.

4.1. Recommended Mechanism

When using the Opus codec, use of the built-in Opus FEC mechanism is RECOMMENDED. This provides reasonable protection of the audio stream against typical losses, with modest overhead. Note that as indicated above the built-in Opus FEC only provides single-frame redundancy; if multi-packet protection is needed, the built-in FEC should be combined with [RFC2198] redundancy to protect the N-2th, N-3rd, etc. packets.

When using the AMR/AMR-WB codecs, use of their built-in FEC mechanism is RECOMMENDED. This provides slightly more efficient protection of the audio stream than [RFC2198].

When using variable-bitrate codecs without an internal FEC, [RFC2198] redundant encoding with lower-fidelity version(s) of previous packet(s) is RECOMMENDED. This provides reasonable protection of the payload with moderate overhead.

When using constant-bitrate codecs, e.g. PCMU, use of [RFC2198] redundant encoding MAY be used, but note that this will result in a potentially significant bitrate increase, and that suddenly increasing bitrate to deal with losses from congestion may actually make things worse.

Because of the lower packet rate of audio encodings, usually a single packet per frame, use of a separate FEC stream comes with a higher overhead than other mechanisms, and therefore is NOT RECOMMENDED.

4.2. Negotiating Support

Support for redundant encoding MUST be indicated by offering "red" as a supported payload type in the offer. Answerers can reject the use of redundant encoding by not including "red" as a supported payload type in the answer.

Support for codec-specific FEC mechanisms are typically indicated via "a=fmtp" parameters.

For Opus, a receiver MUST indicate that it is prepared to use incoming FEC data with the "useinbandfec=1" parameter, as specified

in [RFC7587]. This parameter is declarative and can be negotiated separately for either media direction.

For AMR/AMR-WB, support for redundant encoding, and the maximum supported depth, are controlled by the 'max-red' parameter, as specified in [RFC4867], Section 8.1. Receivers MUST include this parameter, and set it to an appropriate value, as specified in [3GPP.26.114], Table 6.3.

5. FEC for Video Content

The following section provides guidance on how to best use FEC for transmitting video data. As indicated in Section 8 below, FEC should only be activated if network conditions warrant it, or upon explicit application request.

5.1. Recommended Mechanism

For video content, use of a separate FEC stream with the RTP payload format described in [I-D.ietf-payload-flexible-fec-scheme] is RECOMMENDED. The receiver can demultiplex the incoming FEC stream by SSRC and correlate it with the primary stream via the SSRC field present in the FEC header.

Support for protecting multiple primary streams with a single FEC stream is complicated by WebRTC's 1-m-line-per-stream policy, which does not allow for a m-line dedicated specifically to FEC.

5.2. Negotiating Support

To offer support for a SSRC-multiplexed FEC stream that is associated with a given primary stream, the offerer MUST offer the formats supported for the primary stream, as well as one of the formats described in [I-D.ietf-payload-flexible-fec-scheme], Section 5.1.

Use of FEC-only m-lines, and grouping using the SDP group mechanism as described in [RFC5956], Section 4.1 is not currently defined for WebRTC, and SHOULD NOT be offered.

Answerers can reject the use of SSRC-multiplexed FEC, by not including FEC formats in the answer.

Answerers SHOULD reject any FEC-only m-lines, unless they specifically know how to handle such a thing in a WebRTC context (perhaps defined by a future version of the WebRTC specifications). This ensures that implementations will not malfunction when said future version of WebRTC enables offers of FEC-only m-lines.

6. FEC for Application Content

WebRTC also supports the ability to send generic application data, and provides transport-level retransmission mechanisms to support full and partial (e.g. timed) reliability. See [I-D.ietf-rtcweb-data-channel] for details.

Because the application can control exactly what data to send, it has the ability to monitor packet statistics and perform its own application-level FEC, if necessary.

As a result, this document makes no recommendations regarding FEC for the underlying data transport.

7. Implementation Requirements

To support the functionality recommended above, implementations **MUST** be able to receive and make use of the relevant FEC formats for their supported audio codecs, and **MUST** indicate this support, as described in Section 4. Use of these formats when sending, as mentioned above, is **RECOMMENDED**.

The general FEC mechanism described in [I-D.ietf-payload-flexible-fec-scheme] **SHOULD** also be supported, as mentioned in Section 5.

Implementations **MAY** support additional FEC mechanisms if desired, e.g. [RFC5109].

8. Adaptive Use of FEC

Since use of FEC always causes redundant data to be transmitted, this will lead to less bandwidth available for the primary encoding when in a bandwidth-constrained environment. This is in contrast to methods like RTX [RFC4588], which only transmits redundant data when necessary, at the cost of an extra roundtrip.

Given this, WebRTC implementations **SHOULD** consider using RTX instead of FEC when RTT is low, and **SHOULD** only transmit the amount of FEC needed to protect against the observed packet loss (which can be determined, e.g., by monitoring transmit packet loss data from RTCP Receiver Reports [RFC3550]), unless the application indicates it is willing to pay a quality penalty to proactively avoid losses.

When using FEC with layered codecs, e.g., [RFC6386], where only base layer frames are critical to the decoding of future frames, implementations **SHOULD** only apply FEC to these base layer frames.

9. Security Considerations

This document makes recommendations regarding the use of FEC. Generally, it should be noted that although applying redundancy is often useful in protecting a stream against packet loss, if the loss is caused by network congestion, the additional bandwidth used by the redundant data may actually make the situation worse, and can lead to significant degradation of the network.

Additional security considerations for each individual FEC mechanism are enumerated in their respective documents.

10. IANA Considerations

This document requires no actions from IANA.

11. Acknowledgements

Several people provided significant input into this document, including Bernard Aboba, Jonathan Lennox, Giri Mandyam, Varun Singh, Tim Terriberry, Magnus Westerlund, and Mo Zanaty.

12. References

12.1. Normative References

- [I-D.ietf-payload-flexible-fec-scheme] Singh, V., Begen, A., Zanaty, M., and G. Mandyam, "RTP Payload Format for Flexible Forward Error Correction (FEC)", draft-ietf-payload-flexible-fec-scheme-03 (work in progress), October 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2198] Perkins, C., Kouvelas, I., Hodson, O., Hardman, V., Handley, M., Bolot, J., Vega-Garcia, A., and S. Fosse-Parisis, "RTP Payload for Redundant Audio Data", RFC 2198, DOI 10.17487/RFC2198, September 1997, <<http://www.rfc-editor.org/info/rfc2198>>.
- [RFC5956] Begen, A., "Forward Error Correction Grouping Semantics in the Session Description Protocol", RFC 5956, DOI 10.17487/RFC5956, September 2010, <<http://www.rfc-editor.org/info/rfc5956>>.

12.2. Informative References

- [I-D.ietf-rtcweb-data-channel]
Jesup, R., Loreto, S., and M. Tuexen, "WebRTC Data Channels", draft-ietf-rtcweb-data-channel-13 (work in progress), January 2015.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<http://www.rfc-editor.org/info/rfc3550>>.
- [RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R. Hakenberg, "RTP Retransmission Payload Format", RFC 4588, DOI 10.17487/RFC4588, July 2006, <<http://www.rfc-editor.org/info/rfc4588>>.
- [RFC4867] Sjoberg, J., Westerlund, M., Lakaniemi, A., and Q. Xie, "RTP Payload Format and File Storage Format for the Adaptive Multi-Rate (AMR) and Adaptive Multi-Rate Wideband (AMR-WB) Audio Codecs", RFC 4867, DOI 10.17487/RFC4867, April 2007, <<http://www.rfc-editor.org/info/rfc4867>>.
- [RFC5109] Li, A., Ed., "RTP Payload Format for Generic Forward Error Correction", RFC 5109, DOI 10.17487/RFC5109, December 2007, <<http://www.rfc-editor.org/info/rfc5109>>.
- [RFC6386] Bankoski, J., Koleszar, J., Quillio, L., Salonen, J., Wilkins, P., and Y. Xu, "VP8 Data Format and Decoding Guide", RFC 6386, DOI 10.17487/RFC6386, November 2011, <<http://www.rfc-editor.org/info/rfc6386>>.
- [RFC6716] Valin, JM., Vos, K., and T. Terriberry, "Definition of the Opus Audio Codec", RFC 6716, DOI 10.17487/RFC6716, September 2012, <<http://www.rfc-editor.org/info/rfc6716>>.
- [RFC7587] Spittka, J., Vos, K., and JM. Valin, "RTP Payload Format for the Opus Speech and Audio Codec", RFC 7587, DOI 10.17487/RFC7587, June 2015, <<http://www.rfc-editor.org/info/rfc7587>>.

Appendix A. Change log

Changes in draft -04:

- o Discussion of layered codecs.
- o Discussion of RTX.

- o Clarified implementation requirements.
- o FlexFEC MUST -> SHOULD.
- o Clarified AMR max-red handling.
- o Updated references.

Changes in draft -03:

- o Added overhead stats for Opus.
- o Expanded discussion of multi-packet FEC for Opus.
- o Added discussion of AMR/AMR-WB.
- o Removed discussion of ssrc-group.
- o Referenced the data channel doc.
- o Referenced the RTP/RTCP RFC.
- o Several small edits based on feedback from Magnus.

Changes in draft -02:

- o Expanded discussion of FEC-only m-lines, and how they should be handled in offers and answers.

Changes in draft -01:

- o Tweaked abstract/intro text that was ambiguously normative.
- o Removed text on FEC for Opus in CELT mode.
- o Changed RFC 2198 recommendation for PCMU to be MAY instead of NOT RECOMMENDED, based on list feedback.
- o Explicitly called out application data as something not addressed in this document.
- o Updated flexible-fec reference.

Changes in draft -00:

- o Initial version, from sidebar conversation at IETF 90.

Author's Address

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

Email: justin@uberti.name

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 18, 2017

J. Uberti
G. Shieh
Google
January 14, 2017

WebRTC IP Address Handling Requirements
draft-ietf-rtcweb-ip-handling-03

Abstract

This document provides information and requirements for how IP addresses should be handled by WebRTC applications.

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 18, 2017.

Copyright Notice

Copyright (c) 2017 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	2
2. Problem Statement	2
3. Goals	3
4. Detailed Design	4
5. Application Guidance	6
6. Security Considerations	6
7. IANA Considerations	6
8. Acknowledgements	6
9. Informative References	6
Appendix A. Change log	8
Authors' Addresses	8

1. Introduction

As a technology that supports peer-to-peer connections, WebRTC may send data over different network paths than the path used for HTTP traffic. This may allow a web application to learn additional information about the user, which may be problematic in certain cases. This document summarizes the concerns, and makes recommendations on how best to handle the tradeoff between privacy and media performance.

2. Problem Statement

WebRTC enables real-time peer-to-peer communications by enumerating network interfaces and discovering the best route through the ICE [RFC5245] protocol. During the ICE process, the peers involved in a session gather and exchange all the IP addresses they can discover, so that the connectivity of each IP pair can be checked, and the best path chosen. The addresses that are gathered usually consist of an endpoint's private physical/virtual addresses, and its public Internet addresses.

These addresses are exposed upwards to the web application, so that they can be communicated to the remote endpoint. This allows the application to learn more about the local network configuration than it would from a typical HTTP scenario, in which the web server would only see a single public Internet address, i.e. the address from which the HTTP request was sent.

The information revealed falls into three categories:

1. If the client is behind a NAT, the client's private IP addresses, typically [RFC1918] addresses, can be learned.

2. If the client tries to hide its physical location through a VPN, and the VPN and local OS support routing over multiple interfaces (i.e., a "split-tunnel" VPN), WebRTC will discover the public address for the VPN as well as the ISP public address that the VPN runs over.
3. If the client is behind a proxy (a client-configured "classical application proxy", as defined in [RFC1919], Section 3), but direct access to the Internet is also supported, WebRTC's STUN [RFC5389] checks will bypass the proxy and reveal the public address of the client.

Of these three concerns, #2 is the most significant concern, since for some users, the purpose of using a VPN is for anonymity. However, different VPN users will have different needs, and some VPN users (e.g. corporate VPN users) may in fact prefer WebRTC to send media traffic directly, i.e., not through the VPN.

#3 is a less common concern, as proxy administrators can control this behavior through organization firewall policy if desired, coupled with the fact that forcing WebRTC traffic through a proxy will have negative effects on both the proxy and on media quality. For situations where this is an important consideration, use of a RETURN proxy, as described below, can be an effective solution.

#1 is considered to be the least significant concern, given that the local address values often contain minimal information (e.g. 192.168.0.2), or have built-in privacy protection (e.g. [RFC4941]IPv6 addresses).

Note also that these concerns predate WebRTC; Adobe Flash Player has provided similar functionality since the introduction of RTMFP [RFC7016] in 2008.

3. Goals

Being peer-to-peer, WebRTC represents a privacy-enabling technology, and therefore we want to avoid solutions that disable WebRTC or make it harder to use. This means that WebRTC should be configured by default to only reveal the minimum amount of information needed to establish a performant WebRTC session, while providing options to reveal additional information upon user consent, or further limit this information if the user has specifically requested this. Specifically, WebRTC should:

- o Provide a privacy-friendly default behavior which strikes the right balance between privacy and media performance for most users and use cases.

- o For users who care more about one versus the other, provide a means to customize the experience.

4. Detailed Design

The key principles for the design are listed below:

1. By default, WebRTC should follow normal IP routing rules, to the extent that this is easy to determine (i.e., not considering proxies). This can be accomplished by binding local sockets to the wildcard addresses (0.0.0.0 for IPv4, :: for IPv6), which allows the OS to route WebRTC traffic the same way as it would HTTP traffic, and allows only the 'typical' public addresses to be discovered.
2. By default, support for direct connections between hosts (i.e., without traversing a NAT or relay server) should be maintained. To accomplish this, the local IPv4 and IPv6 addresses of the interface used for outgoing STUN traffic should still be surfaced as candidates, even when binding to the wildcard addresses as mentioned above. The appropriate addresses here can be discovered by the common trick of binding sockets to the wildcard addresses, connect()ing those sockets to some well-known public IP address (one particular example being "8.8.8.8"), and then reading the bound local addresses via getsockname(). This approach requires no data exchange; it simply provides a mechanism for applications to retrieve the desired information from the kernel routing table.
3. Determining whether a web proxy is in use is a complex process, as the answer can depend on the exact site or address being contacted. Furthermore, web proxies that support UDP are not widely deployed today. As a result, when WebRTC is made to go through a proxy, it typically needs to use TCP, either ICE-TCP [RFC6544] or TURN-over-TCP [RFC5766]. Naturally, this has attendant costs on media quality as well as proxy performance, and should be avoided where possible.
4. RETURN [I-D.ietf-rtcweb-return] is a proposal for explicit proxying of WebRTC media traffic. When RETURN proxies are deployed, media and STUN checks will go through the proxy, but without the performance issues associated with sending through a typical web proxy.

Based on these ideas, we define four specific modes of WebRTC behavior, reflecting different media quality/privacy tradeoffs:

- Mode 1: Enumerate all addresses: WebRTC MUST bind to all interfaces individually and use them all to attempt communication with STUN servers, TURN servers, or peers. This will converge on the best media path, and is ideal when media performance is the highest priority, but it discloses the most information.
- Mode 2: Default route + associated local addresses: WebRTC MUST follow the kernel routing table rules (e.g., by binding solely to the wildcard address), which will typically cause media packets to take the same route as the application's HTTP traffic. In addition, any private IPv4 and IPv6 addresses associated with the kernel-chosen interface MUST be discovered through `getsockname`, as mentioned above, and provided to the application. This ensures that direct connections can still be established in this mode.
- Mode 3: Default route only: This is the the same as Mode 2, except that the associated private address MUST NOT be provided. This may cause traffic to hairpin through a NAT, fall back to the application TURN server, or fail altogether, with resulting quality implications.
- Mode 4: Force proxy: This forces all WebRTC media traffic through a proxy, if one is configured. If the proxy does not support UDP (as is the case for all HTTP and most SOCKS [RFC1928]proxies), or the WebRTC implementation does not support UDP proxying, the use of UDP will be disabled, and TCP will be used to send and receive media through the proxy. Use of TCP will result in reduced quality, in addition to any performance considerations associated with sending all WebRTC media through the proxy server.

Mode 1 MUST only be used when user consent has been provided; this thwarts the typical drive-by enumeration attacks. The details of this consent are left to the implementation; one potential mechanism is to tie this consent to `getUserMedia` consent.

In cases where user consent has not been obtained, Mode 2 SHOULD be used. This allows applications to still achieve direct connections in many cases, even without consent (e.g., data channel applications). However, user agents MAY choose a stricter default policy in certain circumstances.

Note that when a RETURN proxy is configured for the interface associated with the default route, Mode 2 and 3 will cause any external media traffic to go through the RETURN proxy. While the RETURN approach gives the best performance, a similar result can be achieved for non-RETURN proxies via an organization firewall policy

that only allows external WebRTC traffic to leave through the proxy (typically, over TCP). This provides a way to ensure the proxy is used for any external traffic, but avoids the performance issues of Mode 4, where all media is forced through said proxy, for intra-organization traffic.

5. Application Guidance

The recommendations mentioned in this document may cause certain WebRTC applications to malfunction. In order to be robust in all scenarios, the following guidelines are provided for applications:

- o Applications SHOULD deploy a TURN server with support for both UDP and TCP connections to the server. This ensures that connectivity can still be established, even when Mode 3 or 4 are in use, assuming the TURN server can be reached.
- o Applications SHOULD detect when they don't have access to the full set of ICE candidates by checking for the presence of host candidates. If no host candidates are present, Mode 3 or 4 above is in use; this knowledge can be useful for diagnostic purposes.
- o Future versions of browsers may present an indicator to signify that the page is using WebRTC to set up a peer-to-peer connection. Applications MUST only use WebRTC in a fashion that is consistent with user expectations.

6. Security Considerations

This document is entirely devoted to security considerations.

7. IANA Considerations

This document requires no actions from IANA.

8. Acknowledgements

Several people provided input into this document, including Harald Alvestrand, Ted Hardie, Matthew Kaufmann, Eric Rescorla, Adam Roach, and Martin Thomson.

9. Informative References

[I-D.ietf-rtcweb-return]

Schwartz, B. and J. Uberti, "Recursively Encapsulated TURN (RETURN) for Connectivity and Privacy in WebRTC", draft-ietf-rtcweb-return-01 (work in progress), January 2016.

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<http://www.rfc-editor.org/info/rfc1918>>.
- [RFC1919] Chatel, M., "Classical versus Transparent IP Proxies", RFC 1919, DOI 10.17487/RFC1919, March 1996, <<http://www.rfc-editor.org/info/rfc1919>>.
- [RFC1928] Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., and L. Jones, "SOCKS Protocol Version 5", RFC 1928, DOI 10.17487/RFC1928, March 1996, <<http://www.rfc-editor.org/info/rfc1928>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, DOI 10.17487/RFC5245, April 2010, <<http://www.rfc-editor.org/info/rfc5245>>.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, DOI 10.17487/RFC5389, October 2008, <<http://www.rfc-editor.org/info/rfc5389>>.
- [RFC5766] Mahy, R., Matthews, P., and J. Rosenberg, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)", RFC 5766, DOI 10.17487/RFC5766, April 2010, <<http://www.rfc-editor.org/info/rfc5766>>.
- [RFC6544] Rosenberg, J., Keranen, A., Lowekamp, B., and A. Roach, "TCP Candidates with Interactive Connectivity Establishment (ICE)", RFC 6544, DOI 10.17487/RFC6544, March 2012, <<http://www.rfc-editor.org/info/rfc6544>>.
- [RFC7016] Thornburgh, M., "Adobe's Secure Real-Time Media Flow Protocol", RFC 7016, DOI 10.17487/RFC7016, November 2013, <<http://www.rfc-editor.org/info/rfc7016>>.

Appendix A. Change log

Changes in draft -03:

- o Clarified when to use which modes.
- o Use 2119 qualifiers to make normative statements.
- o Defined 'proxy'.
- o Mentioned split tunnels in problem statement.

Changes in draft -02:

- o Recommendations -> Requirements
- o Updated text regarding consent.

Changes in draft -01:

- o Incorporated feedback from Adam Roach; changes to discussion of cam/mic permission, as well as use of proxies, and various editorial changes.
- o Added several more references.

Changes in draft -00:

- o Published as WG draft.

Authors' Addresses

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

Email: justin@uberti.name

Guo-wei Shieh
Google
747 6th St S
Kirkland, WA 98033
USA

Email: guoweis@google.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2017

J. Uberti
Google
C. Jennings
Cisco
E. Rescorla, Ed.
Mozilla
October 21, 2016

Javascript Session Establishment Protocol
draft-ietf-rtcweb-jsep-17

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 24, 2017.

Copyright Notice

Copyright (c) 2016 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	4
1.1.	General Design of JSEP	4
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	10
3.4.	Session Description Control	10
3.4.1.	RtpTransceivers	10
3.4.2.	RtpSenders	11
3.4.3.	RtpReceivers	11
3.5.	ICE	11
3.5.1.	ICE Gathering Overview	11
3.5.2.	ICE Candidate Trickling	12
3.5.2.1.	ICE Candidate Format	12
3.5.3.	ICE Candidate Policy	13
3.5.4.	ICE Candidate Pool	14
3.6.	Video Size Negotiation	14
3.6.1.	Creating an imageattr Attribute	15
3.6.2.	Interpreting an imageattr Attribute	16
3.7.	Simulcast	17
3.8.	Interactions With Forking	18
3.8.1.	Sequential Forking	18
3.8.2.	Parallel Forking	19
4.	Interface	20
4.1.	PeerConnection	20
4.1.1.	Constructor	20
4.1.2.	addTrack	22
4.1.3.	addTransceiver	22
4.1.4.	createDataChannel	23
4.1.5.	createOffer	23
4.1.6.	createAnswer	24
4.1.7.	SessionDescriptionType	25
4.1.7.1.	Use of Provisional Answers	25
4.1.7.2.	Rollback	26
4.1.8.	setLocalDescription	27
4.1.9.	setRemoteDescription	28
4.1.10.	currentLocalDescription	28
4.1.11.	pendingLocalDescription	28
4.1.12.	currentRemoteDescription	28
4.1.13.	pendingRemoteDescription	29
4.1.14.	canTrickleIceCandidates	29

4.1.15.	setConfiguration	30
4.1.16.	addIceCandidate	30
4.2.	RtpTransceiver	31
4.2.1.	stop	31
4.2.2.	stopped	31
4.2.3.	setDirection	31
4.2.4.	setCodecPreferences	32
5.	SDP Interaction Procedures	32
5.1.	Requirements Overview	32
5.1.1.	Implementation Requirements	33
5.1.2.	Usage Requirements	34
5.1.3.	Profile Names and Interoperability	34
5.2.	Constructing an Offer	35
5.2.1.	Initial Offers	35
5.2.2.	Subsequent Offers	41
5.2.3.	Options Handling	44
5.2.3.1.	IceRestart	44
5.2.3.2.	VoiceActivityDetection	45
5.3.	Generating an Answer	45
5.3.1.	Initial Answers	45
5.3.2.	Subsequent Answers	50
5.3.3.	Options Handling	51
5.3.3.1.	VoiceActivityDetection	51
5.4.	Modifying an Offer or Answer	51
5.5.	Processing a Local Description	52
5.6.	Processing a Remote Description	53
5.7.	Parsing a Session Description	53
5.7.1.	Session-Level Parsing	54
5.7.2.	Media Section Parsing	55
5.7.3.	Semantics Verification	58
5.8.	Applying a Local Description	59
5.9.	Applying a Remote Description	60
5.10.	Applying an Answer	64
6.	Processing RTP/RTCP packets	66
7.	Examples	68
7.1.	Simple Example	68
7.2.	Normal Examples	72
8.	Security Considerations	81
9.	IANA Considerations	81
10.	Acknowledgements	81
11.	References	82
11.1.	Normative References	82
11.2.	Informative References	85
Appendix A.	Appendix A	86
Appendix B.	Change log	87
Authors' Addresses		94

1. Introduction

This document describes how the W3C WEBRTC RTCPeerConnection interface [W3C.WD-webrtc-20140617] 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.

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. JSEP removes the browser almost entirely from the core signaling flow, which is instead handled by the Javascript making use of two interfaces: (1) passing in local and remote session descriptions and (2) interacting 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.

To complete the offer/answer exchange, the remote party uses the `createAnswer()` API to generate an appropriate answer, applies it using the `setLocalDescription()` API, and sends the answer back to the initiator over the signaling channel. When the initiator gets that answer, it installs it using the `setRemoteDescription()` API, 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 session description can be sent immediately and the transport information can be sent when available. 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 ICE checks can start as soon as any transport information is available rather than waiting for all of it.

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, libraries exist to adapt the JSEP API into an API suitable for a SIP or XMPP. Thus, JSEP provides 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 session 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 DTLS-SRTP parameters [RFC5763] sent to a remote party indicate what certificate the local side will use in DTLS setup, and thereby what the remote party should expect to receive; the remote party will have to accept these parameters, with no option to choose different values.

In addition, various RFCs put different conditions on the format of offers versus answers. For example, an offer may propose an arbitrary number of media streams (i.e. m sections), but an answer must contain the exact same number as the offer.

Lastly, while the exact media parameters are only known only after an 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 `setLocalDescription` and `setRemoteDescription` methods 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 at 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 via the `createOffer()` method 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:

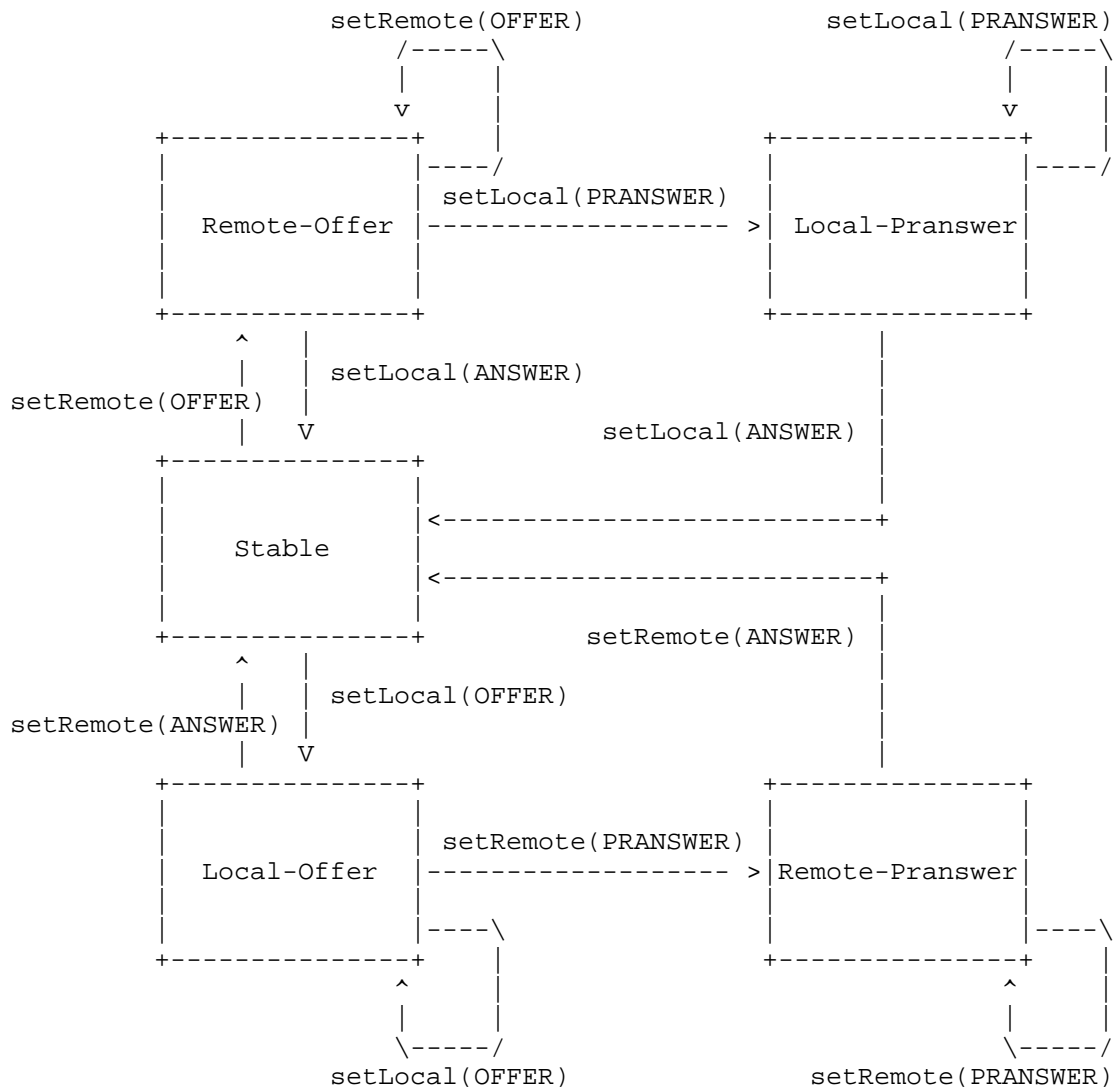


Figure 2: JSEP State Machine

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.

3.4. Session Description Control

In order to give the application control over various common session parameters, JSEP provides control surfaces which tell the browser how to generate session descriptions. This avoids the need for Javascript to modify session descriptions in most cases.

Changes to these objects result in changes to the session descriptions generated by subsequent createOffer/Answer calls.

3.4.1. RtpTransceivers

RtpTransceivers allow the application to control the RTP media associated with one m= section. Each RtpTransceiver has an RtpSender and an RtpReceiver, which an application can use to control the sending and receiving of RTP media. The application may also modify the RtpTransceiver directly, for instance, by stopping it.

RtpTransceivers generally have a 1:1 mapping with m= sections, although there may be more RtpTransceivers than m= sections when

RtpTransceivers are created but not yet associated with a m= section, or if RtpTransceivers have been stopped and disassociated from m= sections. An RtpTransceiver is never associated with more than one m= section, and once a session description is applied, a m= section is always associated with exactly one RtpTransceiver.

RtpTransceivers can be created explicitly by the application or implicitly by calling setRemoteDescription with an offer that adds new m= sections.

3.4.2. RtpSenders

RtpSenders allow the application to control how RTP media is sent.

3.4.3. RtpReceivers

RtpReceivers allows the application to control how RTP media is received.

3.5. ICE

3.5.1. ICE Gathering Overview

JSEP gathers ICE candidates as needed by the application. Collection of ICE candidates is referred to as a gathering phase, and this is triggered either by the addition of a new or recycled m= line to the local session description, or new ICE credentials in the description, indicating an ICE restart. Use of new ICE credentials can be triggered explicitly by the application, or implicitly by the browser in response to changes in the ICE configuration.

When the ICE configuration changes in a way that requires a new gathering phase, a 'needs-ice-restart' bit is set. When this bit is set, calls to the createOffer API will generate new ICE credentials. This bit is cleared by a call to the setLocalDescription API with new ICE credentials from either an offer or an answer, i.e., from either a local- or remote-initiated ICE restart.

When a new gathering phase starts, the ICE Agent will notify the application that gathering is occurring through an event. Then, when each new ICE candidate becomes available, the ICE Agent will supply it to the application via an additional event; these candidates will also automatically be added to the current and/or pending local session description. Finally, when all candidates have been gathered, an event will be dispatched to signal that the gathering process is complete.

Note that gathering phases only gather the candidates needed by new/recycled/restarting m= lines; other m= lines continue to use their existing candidates. Also, when bundling is active, candidates are only gathered (and exchanged) for the m= lines referenced in BUNDLE-tags, as described in [I-D.ietf-mmusic-sdp-bundle-negotiation].

3.5.2. 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-ice-trickle]. 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 media setup in cases where gathering is not performed prior to initiating the call.

JSEP supports optional candidate trickling by providing APIs, as described above, 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.

3.5.2.1. 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, with index N referring to the N+1th m= line in the SDP sent by the entity which sent the IceCandidate. The MID uses the "media stream identification" attribute, as defined in [RFC5888], Section 4, to identify the m= line. JSEP implementations creating an ICE Candidate object MUST populate both of these fields, using the MID of the associated RtpTransceiver object (which may be locally generated by the answerer when interacting with a non-JSEP remote endpoint that does not support the MID attribute, as discussed in Section 5.9 below). Implementations receiving an ICE Candidate object MUST use the MID if present, or the m= line index, if not (the non-JSEP remote endpoint case).

3.5.3. ICE Candidate Policy

Typically, when gathering ICE candidates, the browser will gather all possible forms of initial candidates - host, server reflexive, and relay. However, in certain cases, applications may want to have more specific control over the gathering process, due to privacy or related concerns. For example, one may want to suppress the use of host candidates, to avoid exposing information about the local network, or go as far as only using relay candidates, to leak as little location information as possible (note that these choices come with corresponding operational costs). To accomplish this, the browser MUST allow the application to restrict which ICE candidates are used in a session. Note that this filtering is applied on top of any restrictions the browser chooses to enforce regarding which IP addresses are permitted for the application, as discussed in [I-D.ietf-rtcweb-ip-handling].

There may also be cases where the application wants to change which types of candidates are used while the session is active. A prime example is where a callee may initially want to use only relay candidates, to avoid leaking location information to an arbitrary caller, but then change to use all candidates (for lower operational cost) once the user has indicated they want to take the call. For this scenario, the browser MUST allow the candidate policy to be changed in mid-session, subject to the aforementioned interactions with local policy.

To administer the ICE candidate policy, the browser will determine the current setting at the start of each gathering phase. Then, during the gathering phase, the browser MUST NOT expose candidates disallowed by the current policy to the application, use them as the source of connectivity checks, or indirectly expose them via other fields, such as the raddr/rport attributes for other ICE candidates.

Later, if a different policy is specified by the application, the application can apply it by kicking off a new gathering phase via an ICE restart.

3.5.4. ICE Candidate Pool

JSEP applications typically inform the browser to begin ICE gathering via the information supplied to `setLocalDescription`, as this is where the app specifies the number of media streams, and thereby ICE components, for which to gather candidates. However, to accelerate cases where the application knows the number of ICE components to use ahead of time, it may ask the browser to gather a pool of potential ICE candidates to help ensure rapid media setup.

When `setLocalDescription` is eventually called, and the browser goes to gather the needed ICE candidates, it SHOULD start by checking if any candidates are available in the pool. If there are candidates in the pool, they SHOULD be handed to the application immediately via the ICE candidate event. If the pool becomes depleted, either because a larger-than-expected number of ICE components is used, or because the pool has not had enough time to gather candidates, the remaining candidates are gathered as usual.

One example of where this concept is useful is an application that expects an incoming call at some point in the future, and wants to minimize the time it takes to establish connectivity, to avoid clipping of initial media. By pre-gathering candidates into the pool, it can exchange and start sending connectivity checks from these candidates almost immediately upon receipt of a call. Note though that by holding on to these pre-gathered candidates, which will be kept alive as long as they may be needed, the application will consume resources on the STUN/TURN servers it is using.

3.6. Video Size Negotiation

Video size negotiation is the process through which a receiver can use the "a=imageattr" SDP attribute [RFC6236] to indicate what video frame sizes it is capable of receiving. A receiver may have hard limits on what its video decoder can process, or it may wish to constrain what it receives due to application preferences, e.g. a specific size for the window in which the video will be displayed.

Note that certain codecs support transmission of samples with aspect ratios other than 1.0 (i.e., non-square pixels). JSEP implementations will not transmit non-square pixels, but SHOULD receive and render such video with the correct aspect ratio. However, sample aspect ratio has no impact on the size negotiation described below; all dimensions assume square pixels.

3.6.1.1. Creating an imageattr Attribute

In order to determine the limits on what video resolution a receiver wants to receive, it will intersect its decoder hard limits with any mandatory constraints that have been applied to the associated `MediaStreamTrack`. If the decoder limits are unknown, e.g. when using a software decoder, the mandatory constraints are used directly. For the answerer, these mandatory constraints can be applied to the remote `MediaStreamTracks` that are created by a `setRemoteDescription` call, and will affect the output of the ensuing `createAnswer` call. Any constraints set after `setLocalDescription` is used to set the answer will result in a new offer-answer exchange. For the offerer, because it does not know about any remote `MediaStreamTracks` until it receives the answer, the offer can only reflect decoder hard limits. If the offerer wishes to set mandatory constraints on video resolution, it must do so after receiving the answer, and the result will be a new offer-answer to communicate them.

If there are no known decoder limits or mandatory constraints, the "a=imageattr" attribute SHOULD be omitted.

Otherwise, an "a=imageattr" attribute is created with "recv" direction, and the resulting resolution space formed by intersecting the decoder limits and constraints is used to specify its minimum and maximum x= and y= values. If the intersection is the null set, i.e., there are no resolutions that are permitted by both the decoder and the mandatory constraints, this SHOULD be represented by x=0 and y=0 values.

The rules here express a single set of preferences, and therefore, the "a=imageattr" q= value is not important. It SHOULD be set to 1.0.

The "a=imageattr" field is payload type specific. When all video codecs supported have the same capabilities, use of a single attribute, with the wildcard payload type (*), is RECOMMENDED. However, when the supported video codecs have differing capabilities, specific "a=imageattr" attributes MUST be inserted for each payload type.

As an example, consider a system with a HD-capable, multiformat video decoder, where the application has constrained the received track to at most 360p. In this case, the implementation would generate this attribute:

```
a=imageattr:* recv [x=[16:640],y=[16:360],q=1.0]
```

This declaration indicates that the receiver is capable of decoding any image resolution from 16x16 up to 640x360 pixels.

3.6.2. Interpreting an imageattr Attribute

[RFC6236] defines "a=imageattr" to be an advisory field. This means that it does not absolutely constrain the video formats that the sender can use, but gives an indication of the preferred values.

This specification prescribes more specific behavior. When a sender of a given MediaStreamTrack, which is producing video of a certain resolution, receives an "a=imageattr recv" attribute, it MUST check to see if the original resolution meets the size criteria specified in the attribute, and adapt the resolution accordingly by scaling (if appropriate). Note that when considering a MediaStreamTrack that is producing rotated video, the unrotated resolution MUST be used. This is required regardless of whether the receiver supports performing receive-side rotation (e.g., through CVO), as it significantly simplifies the matching logic.

For the purposes of resolution negotiation, only size limits are considered. Any other values, e.g. picture or sample aspect ratio, MUST be ignored.

When communicating with a non-JSEP endpoint, multiple relevant "a=imageattr recv" attributes may be received. If this occurs, attributes other than the one with the highest "q=" value MUST be ignored.

If an "a=imageattr recv" attribute references a different video codec than what has been selected for the MediaStreamTrack, it MUST be ignored.

If the original resolution matches the size limits in the attribute, the track MUST be transmitted untouched.

If the original resolution exceeds the size limits in the attribute, the sender SHOULD apply downscaling to the output of the MediaStreamTrack in order to satisfy the limits. Downscaling MUST NOT change the track aspect ratio.

If the original resolution is less than the size limits in the attribute, upscaling is needed, but this may not be appropriate in all cases. To address this concern, the application can set an upscaling policy for each sent track. For this case, if upscaling is permitted by policy, the sender SHOULD apply upscaling in order to provide the desired resolution. Otherwise, the sender MUST NOT apply

upsampling. The sender SHOULD NOT upscale in other cases, even if the policy permits it. Upsampling MUST NOT change the track aspect ratio.

If there is no appropriate and permitted scaling mechanism that allows the received size limits to be satisfied, the sender MUST NOT transmit the track.

If the attribute includes a "sar=" (sample aspect ratio) value set to something other than "1.0", indicating the receiver wants to receive non-square pixels, this cannot be satisfied and the sender MUST NOT transmit the track.

In the special case of receiving a maximum resolution of [0, 0], as described above, the sender MUST NOT transmit the track.

3.7. Simulcast

JSEP supports simulcast of a `MediaStreamTrack`, where multiple encodings of the source media can be transmitted within the context of a single `m=` section. The current JSEP API is designed to allow applications to send simulcasted media but only to receive a single encoding. This allows for multi-user scenarios where each sending client sends multiple encodings to a server, which then, for each receiving client, chooses the appropriate encoding to forward.

Applications request support for simulcast by configuring multiple encodings on an `RTPSender`, which, upon generation of an offer or answer, are indicated in SDP markings on the corresponding `m=` section, as described below. Receivers that understand simulcast and are willing to receive it will also include SDP markings to indicate their support, and JSEP endpoints will use these markings to determine whether simulcast is permitted for a given `RTPSender`. If simulcast support is not negotiated, the `RTPSender` will only use the first configured encoding.

Note that the exact simulcast parameters are up to the sending application. While the aforementioned SDP markings are provided to ensure the remote side can receive and demux multiple simulcast encodings, the specific resolutions and bitrates to be used for each encoding are purely a send-side decision in JSEP.

JSEP currently does not provide an API to configure receipt of simulcast. This means that if simulcast is offered by the remote endpoint, the answer generated by a JSEP endpoint will not indicate support for receipt of simulcast, and as such the remote endpoint will only send a single encoding per `m=` section. In addition, when the JSEP endpoint is the answerer, the permitted encodings for the `RTPSender` must be consistent with the offer, but this information is

currently not surfaced through any API. This means that established simulcast streams will continue to work through a received re-offer, but setting up initial simulcast by way of a received offer requires out-of-band signaling or SDP inspection. Future versions of this specification may add additional APIs to provide this control.

When using JSEP to transmit multiple encodings from a RTPSender, the techniques from [I-D.ietf-mmusic-sdp-simulcast] and [I-D.ietf-mmusic-rid] are used. Specifically, when multiple encodings have been configured for a RTPSender, the m= section for the RTPSender will include an "a=simulcast" attribute, as defined in [I-D.ietf-mmusic-sdp-simulcast], Section 6.2, with a "send" simulcast stream description that lists each desired encoding, and no "recv" simulcast stream description. The m= section will also include an "a=rid" attribute for each encoding, as specified in [I-D.ietf-mmusic-rid], Section 4; the use of RID identifiers allows the individual encodings to be disambiguated even though they are all part of the same m= section.

3.8. 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 that 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:

- o Start exchanging media with a given remote peer, but keep all the resources reserved in the offer.
- o Start exchanging media with a given remote peer, and free any resources in the offer that are not being used.

3.8.1. 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 pending remote description as a final answer.

3.8.2. 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. 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.

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. PeerConnection

4.1.1. Constructor

The PeerConnection constructor allows the application to specify global parameters for the media session, such as the STUN/TURN servers and credentials to use when gathering candidates, as well as the initial ICE candidate policy and pool size, and also the bundle policy to use.

If an ICE candidate policy is specified, it functions as described in Section 3.5.3, causing the browser to only surface the permitted candidates (including any internal browser filtering) to the application, and only use those candidates for connectivity checks. The set of available policies is as follows:

all: All candidates permitted by browser policy will be gathered and used.

relay: All candidates except relay candidates will be filtered out. This obfuscates the location information that might be ascertained by the remote peer from the received candidates. Depending on how the application deploys its relay servers, this could obfuscate location to a metro or possibly even global level.

The default ICE candidate policy MUST be set to "all" as this is generally the desired policy, and also typically reduces use of application TURN server resources significantly.

If a size is specified for the ICE candidate pool, this indicates the number of ICE components to pre-gather candidates for. Because pre-gathering results in utilizing STUN/TURN server resources for

potentially long periods of time, this must only occur upon application request, and therefore the default candidate pool size MUST be zero.

The application can specify its preferred policy regarding use of bundle, the multiplexing mechanism defined in [I-D.ietf-mmusic-sdp-bundle-negotiation]. Regardless of policy, the application will always try to negotiate bundle onto a single transport, and will offer a single bundle group across all media section; use of this single transport is contingent upon the answerer accepting bundle. However, by specifying a policy from the list below, the application can control exactly how aggressively it will try to bundle media streams together, which affects how it will interoperate with a non-bundle-aware endpoint. When negotiating with a non-bundle-aware endpoint, only the streams not marked as bundle-only streams will be established.

The set of available policies is as follows:

balanced: The first media section of each type (audio, video, or application) will contain transport parameters, which will allow an answerer to unbundle that section. The second and any subsequent media section of each type will be marked bundle-only. The result is that if there are N distinct media types, then candidates will be gathered for for N media streams. This policy balances desire to multiplex with the need to ensure basic audio and video can still be negotiated in legacy cases. When acting as answerer, if there is no bundle group in the offer, the implementation will reject all but the first m= section of each type.

max-compat: All media sections will contain transport parameters; none will be marked as bundle-only. This policy will allow all streams to be received by non-bundle-aware endpoints, but require separate candidates to be gathered for each media stream.

max-bundle: Only the first media section will contain transport parameters; all streams other than the first will be marked as bundle-only. This policy aims to minimize candidate gathering and maximize multiplexing, at the cost of less compatibility with legacy endpoints. When acting as answerer, the implementation will reject any m= sections other than the first m= section, unless they are in the same bundle group as that m= section.

As it provides the best tradeoff between performance and compatibility with legacy endpoints, the default bundle policy MUST be set to "balanced".

The application can specify its preferred policy regarding use of RTP/RTCP multiplexing [RFC5761] using one of the following policies:

negotiate: The browser will gather both RTP and RTCP candidates but also will offer "a=rtcp-mux", thus allowing for compatibility with either multiplexing or non-multiplexing endpoints.

require: The browser will only gather RTP candidates. This halves the number of candidates that the offerer needs to gather. When acting as answerer, the implementation will reject any m= section that does not contain an "a=rtcp-mux" attribute.

The default multiplexing policy MUST be set to "require". Implementations MAY choose to reject attempts by the application to set the multiplexing policy to "negotiate".

4.1.2. addTrack

The addTrack method adds a MediaStreamTrack to the PeerConnection, using the MediaStream argument to associate the track with other tracks in the same MediaStream, so that they can be added to the same "LS" group when creating an offer or answer. addTrack attempts to minimize the number of transceivers as follows: If the PeerConnection is in the "have-remote-offer" state, the track will be attached to the first compatible transceiver that was created by the most recent call to setRemoteDescription() and does not have a local track. Otherwise, a new transceiver will be created, as described in Section 4.1.3.

4.1.3. addTransceiver

The addTransceiver method adds a new RTPTransceiver to the PeerConnection. If a MediaStreamTrack argument is provided, then the transceiver will be configured with that media type and the track will be attached to the transceiver. Otherwise, the application MUST explicitly specify the type; this mode is useful for creating recvonly transceivers as well as for creating transceivers to which a track can be attached at some later point.

At the time of creation, the application can also specify a transceiver direction attribute, a set of MediaStreams which the transceiver is associated with (allowing LS group assignments), and a set of encodings for the media (used for simulcast as described in Section 3.7).

4.1.4. createDataChannel

The createDataChannel method creates a new data channel and attaches it to the PeerConnection. If no data channel currently exists for this PeerConnection, then a new offer/answer exchange is required. All data channels on a given PeerConnection share the same SCTP/DTLS association and therefore the same m= section, so subsequent creation of data channels does not have any impact on the JSEP state.

The createDataChannel method also includes a number of arguments which are used by the PeerConnection (e.g., maxPacketLifetime) but are not reflected in the SDP and do not affect the JSEP state.

4.1.5. 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 media added to this PeerConnection, the codec/RTP/RTCP options supported by this implementation, and any candidates that have been gathered by the ICE Agent. An options parameter may be supplied to provide additional control over the generated offer. This options parameter allows an application to trigger an ICE restart, for the purpose of reestablishing connectivity.

In the initial offer, the generated SDP will contain all desired functionality for the session (functionality that is 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 stopping RtpTransceivers, 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 RFC 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.

4.1.6. `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 most recent call to `setRemoteDescription`, which **MUST** have been called prior to calling `createAnswer`. Like `createOffer`, the returned blob contains descriptions of the media added to this `PeerConnection`, the codec/RTP/RTCP options negotiated for this session, and any candidates that have been gathered by the ICE Agent. An options 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.

4.1.7. SessionDescriptionType

Session description objects (RTCSessionDescription) may be of type "offer", "pranswer", "answer" or "rollback". 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 "pranswer".

"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 an "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).

"rollback" is a special session description type implying that the state machine should be rolled back to the previous stable state, as described in Section 4.1.7.2. The contents MUST be empty.

4.1.7.1. Use of Provisional Answers

Most web applications will not need to create answers using the "pranswer" type. While it is good practice to send an immediate response to an "offer", in order to warm up the session transport and prevent media clipping, the preferred handling for a web application

would be to create and send an "inactive" final answer immediately after receiving the offer. Later, when the called user actually accepts the call, the application can create a new "sendrecv" offer to update the previous offer/answer pair and start the media flow. While this could also be done with an inactive "pranswer", followed by a sendrecv "answer", the initial "pranswer" leaves the offer-answer exchange open, which means that neither side can send an updated offer during this time.

As an example, 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 triggered the new offer, it is likely that the ICE and DTLS handshaking for all the channels will already have finished.

Of course, some applications may not be able to perform this double offer-answer exchange, particularly ones that are attempting to gateway to legacy signaling protocols. In these cases, "pranswer" can still provide the application with a mechanism to warm up the transport.

4.1.7.2. 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 discards any proposed changes to the session, returning the state machine to the stable state, and setting the pending local and/or remote description back to null. Any resources or candidates that were allocated by the abandoned local description are discarded; any media that is received will be processed according to the previous local and remote descriptions. Rollback can only be used to cancel proposed changes; there is no support for rolling back from a stable state to a previous stable state. Note that this implies that

once the answerer has performed `setLocalDescription` with his answer, this cannot be rolled back.

A rollback will disassociate any `RtpTransceivers` that were associated with `m=` sections by the application of the rolled-back session description (see Section 5.9 and Section 5.8). This means that some `RtpTransceivers` that were previously associated will no longer be associated with any `m=` section; in such cases, the value of the `RtpTransceiver`'s `mid` attribute MUST be set to null. `RtpTransceivers` that were created by applying a remote offer that was subsequently rolled back MUST be removed. However, a `RtpTransceiver` MUST NOT be removed if the `RtpTransceiver`'s `RtpSender` was activated by the `addTrack` method. This is so that an application may call `addTrack`, then call `setRemoteDescription` with an offer, then roll back that offer, then call `createOffer` and have a `m=` section for the added track appear in the generated offer.

A rollback is performed by supplying a session description of type "rollback" with empty contents to either `setLocalDescription` or `setRemoteDescription`, depending on which was most recently used (i.e. if the new offer was supplied to `setLocalDescription`, the rollback should be done using `setLocalDescription` as well).

4.1.8. `setLocalDescription`

The `setLocalDescription` method instructs the `PeerConnection` to apply the supplied session description as its local configuration. The `type` field indicates whether the description 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 current and pending 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 pending local description, or roll back to the current 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.

4.1.9. `setRemoteDescription`

The `setRemoteDescription` method instructs the `PeerConnection` to apply the supplied session description as the desired remote configuration. As in `setLocalDescription`, the `type` field of the description indicates how it should be processed.

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

If `setLocalDescription` was previously called with an offer, and `setRemoteDescription` 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.

4.1.10. `currentLocalDescription`

The `currentLocalDescription` method returns a copy of the current negotiated local description - i.e., the local description from the last successful offer/answer exchange - in addition to any local candidates that have been generated by the ICE Agent since the local description was set.

A null object will be returned if an offer/answer exchange has not yet been completed.

4.1.11. `pendingLocalDescription`

The `pendingLocalDescription` method returns a copy of the local description currently in negotiation - i.e., a local offer set without any corresponding remote answer - in addition to any local candidates that have been generated by the ICE Agent since the local description was set.

A null object will be returned if the state of the `PeerConnection` is "stable" or "have-remote-offer".

4.1.12. `currentRemoteDescription`

The `currentRemoteDescription` method returns a copy of the current negotiated remote description - i.e., the remote description from the last successful offer/answer exchange - in addition to any remote

candidates that have been supplied via processIceMessage since the remote description was set.

A null object will be returned if an offer/answer exchange has not yet been completed.

4.1.13. pendingRemoteDescription

The pendingRemoteDescription method returns a copy of the remote description currently in negotiation - i.e., a remote offer set without any corresponding local answer - in addition to any remote candidates that have been supplied via processIceMessage since the remote description was set.

A null object will be returned if the state of the PeerConnection is "stable" or "have-local-offer".

4.1.14. canTrickleIceCandidates

The canTrickleIceCandidates property indicates whether the remote side supports receiving trickled candidates. There are three potential values:

null: No SDP has been received from the other side, so it is not known if it can handle trickle. This is the initial value before setRemoteDescription() is called.

true: SDP has been received from the other side indicating that it can support trickle.

false: SDP has been received from the other side indicating that it cannot support trickle.

As described in Section 3.5.2, JSEP implementations always provide candidates to the application individually, consistent with what is needed for Trickle ICE. However, applications can use the canTrickleIceCandidates property to determine whether their peer can actually do Trickle ICE, i.e., whether it is safe to send an initial offer or answer followed later by candidates as they are gathered. As "true" is the only value that definitively indicates remote Trickle ICE support, an application which compares canTrickleIceCandidates against "true" will by default attempt Half Trickle on initial offers and Full Trickle on subsequent interactions with a Trickle ICE-compatible agent.

4.1.15. setConfiguration

The setConfiguration method allows the global configuration of the PeerConnection, which was initially set by constructor parameters, to be changed during the session. The effects of this method call depend on when it is invoked, and differ depending on which specific parameters are changed:

- o Any changes to the STUN/TURN servers to use affect the next gathering phase. If an ICE gathering phase has already started or completed, the 'needs-ice-restart' bit mentioned in Section 3.5.1 will be set. This will cause the next call to createOffer to generate new ICE credentials, for the purpose of forcing an ICE restart and kicking off a new gathering phase, in which the new servers will be used. If the ICE candidate pool has a nonzero size, any existing candidates will be discarded, and new candidates will be gathered from the new servers.
- o Any change to the ICE candidate policy affects the next gathering phase. If an ICE gathering phase has already started or completed, the 'needs-ice-restart' bit will be set. Either way, changes to the policy have no effect on the candidate pool, because pooled candidates are not surfaced to the application until a gathering phase occurs, and so any necessary filtering can still be done on any pooled candidates.
- o Any changes to the ICE candidate pool size take effect immediately; if increased, additional candidates are pre-gathered; if decreased, the now-superfluous candidates are discarded.
- o The bundle and RTCP-multiplexing policies MUST NOT be changed after the construction of the PeerConnection.

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.

4.1.16. addIceCandidate

The addIceCandidate method provides a remote candidate to the ICE Agent, which, if parsed successfully, will be added to the current and/or pending remote description according to the rules defined for Trickle ICE. The pair of MID and ufrag is used to determine the m= section and ICE candidate generation to which the candidate belongs. If the MID is not present, the m= line index is used to look up the locally generated MID (see Section 5.9), which is used in place of a supplied MID. If these values or the candidate string are invalid, an error is generated.

The purpose of the ufrag is to resolve ambiguities when trickle ICE is in progress during an ICE restart. If the ufrag is absent, the candidate MUST be assumed to belong to the most recently applied remote description. Connectivity checks will be sent to the new candidate.

This method can also be used to provide an end-of-candidates indication to the ICE Agent, as defined in [I-D.ietf-ice-trickle]). The MID and ufrag are used as described above to determine the m= section and ICE generation for which candidate gathering is complete. If the ufrag is not present, then the end-of-candidates indication MUST be assumed to apply to the relevant m= section in the most recently applied remote description. If neither the MID nor the m= index is present, then the indication MUST be assumed to apply to all m= sections in the most recently applied remote description.

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.

4.2. RtpTransceiver

4.2.1. stop

The stop method stops an RtpTransceiver. This will cause future calls to createOffer to generate a zero port for the associated m= section. See below for more details.

4.2.2. stopped

The stopped method returns "true" if the transceiver has been stopped, either by a call to stopTransceiver or by applying an answer that rejects the associated m= section, and "false" otherwise.

A stopped RtpTransceiver does not send any outgoing RTP or RTCP or process any incoming RTP or RTCP. It cannot be restarted.

4.2.3. setDirection

The setDirection method sets the direction of a transceiver, which affects the direction attribute of the associated m= section on future calls to createOffer and createAnswer.

When creating offers, the transceiver direction is directly reflected in the output, even for reoffers. When creating answers, the transceiver direction is intersected with the offered direction, as explained in the Section 5.3 section below.

4.2.4. setCodecPreferences

The setCodecPreferences method sets the codec preferences of a transceiver, which in turn affect the presence and order of codecs of the associated m= section on future calls to createOffer and createAnswer. Note that setCodecPreferences does not directly affect which codec the implementation decides to send. It only affects which codecs the implementation indicates that it prefers to receive, via the offer or answer. Even when a codec is excluded by setCodecPreferences, it still may be used to send until the next offer/answer exchange discards it.

The codec preferences of an RtpTransceiver can cause codecs to be excluded by subsequent calls to createOffer and createAnswer, in which case the corresponding media formats in the associated m= section will be excluded. The codec preferences cannot add media formats that would otherwise not be present. This includes codecs that were not negotiated in a previous offer/answer exchange that included the transceiver.

The codec preferences of an RtpTransceiver can also determine the order of codecs in subsequent calls to createOffer and createAnswer, in which case the order of the media formats in the associated m= section will match. However, the codec preferences cannot change the order of the media formats after an answer containing the transceiver has been applied. At this point, codecs can only be removed, not reordered.

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.

5.1.1.1. 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 [RFC5764], TCP/DTLS/RTP/SAVPF [I-D.nandakumar-mmusic-proto-iana-registration], "UDP/DTLS/SCTP" [I-D.ietf-mmusic-sctp-sdp], and "TCP/DTLS/SCTP" [I-D.ietf-mmusic-sctp-sdp] RTP profiles.
- 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 and grouping semantics.
- 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 [RFC4588] MUST be implemented to signal RTX payload type associations.
- R-15 [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.
- R-16 TODO: any others?

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 ICE, as specified in [RFC5245], MUST be used. Note that the remote endpoint may use a Lite implementation; implementations MUST properly handle remote endpoints which do ICE-Lite.
- R-2 DTLS [RFC6347] or DTLS-SRTP [RFC5763], MUST be used, as appropriate for the media type, as specified in [I-D.ietf-rtcweb-security-arch]

5.1.3. Profile Names and Interoperability

For media m= sections, JSEP endpoints MUST support both the "UDP/TLS/RTP/SAVPF" and "TCP/DTLS/RTP/SAVPF" profiles and MUST indicate one of these two profiles for each media m= line they produce in an offer. For data m= sections, JSEP endpoints must support both the "UDP/DTLS/SCTP" and "TCP/DTLS/SCTP" profiles and MUST indicate one of these two profiles for each data m= line they produce in an offer. Because ICE can select either TCP or UDP transport depending on network conditions, both advertisements are consistent with ICE eventually selecting either either UDP or TCP.

Unfortunately, in an attempt at compatibility, some endpoints generate other profile strings even when they mean to support one of these profiles. For instance, an endpoint might generate "RTP/AVP"

but supply "a=fingerprint" and "a=rtcp-fb" attributes, indicating its willingness to support "(UDP,TCP)/TLS/RTP/SAVPF". In order to simplify compatibility with such endpoints, JSEP endpoints MUST follow the following rules when processing the media m= sections in an offer:

- o The profile in any "m=" line in any answer MUST exactly match the profile provided in the offer.
- o Any profile matching the following patterns MUST be accepted: "RTP/[S]AVP[F]" and "(UDP/TCP)/TLS/RTP/SAVP[F]"
- o Because DTLS-SRTP is REQUIRED, the choice of SAVP or AVP has no effect; support for DTLS-SRTP is determined by the presence of one or more "a=fingerprint" attribute. Note that lack of an "a=fingerprint" attribute will lead to negotiation failure.
- o The use of AVPF or AVP simply controls the timing rules used for RTCP feedback. If AVPF is provided, or an "a=rtcp-fb" attribute is present, assume AVPF timing, i.e., a default value of "trr-int=0". Otherwise, assume that AVPF is being used in an AVP compatible mode and use AVP timing, i.e., "trr-int=4".
- o For data m= sections, JSEP endpoints MUST support receiving the "UDP/DTLS/SCTP", "TCP/DTLS/SCTP", or "DTLS/SCTP" (for backwards compatibility) profiles.

Note that re-offers by JSEP endpoints MUST use the correct profile strings even if the initial offer/answer exchange used an (incorrect) older profile string.

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 "-". [RFC3264] requires that the <sess-id> be representable as a 64-bit signed integer. It is RECOMMENDED that the <sess-id> be generated as a 64-bit quantity with the high bit being sent to zero and the remaining 63 bits being cryptographically random. 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=-". Note that this differs from the advice in [RFC4566] which proposes a single space, but as both "o=" and "s=" are meaningless, having the same meaningless value seems clearer.
- 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=00".
- o An "a=ice-options" line with the "trickle" option MUST be added, as specified in [I-D.ietf-ice-trickle], Section 4.

The next step is to generate m= sections, as specified in [RFC4566] Section 5.14. An m= section is generated for each RtpTransceiver that has been added to the PeerConnection. This is done in the order that their associated RtpTransceivers were added to the PeerConnection and excludes RtpTransceivers that are stopped and not associated with an m= section (either due to an m= section being recycled or an RtpTransceiver having been stopped before being associated with an m= section) .

Each m= section, provided it is not marked as bundle-only, MUST generate a unique set of ICE credentials and gather its own unique

set of ICE candidates. Bundle-only m= sections MUST NOT contain any ICE credentials and MUST NOT gather any candidates.

For DTLS, all m= sections MUST use all the certificate(s) that have been specified for the PeerConnection; as a result, they MUST all have the same [I-D.ietf-mmusic-4572-update] fingerprint value(s), or these value(s) MUST be session-level attributes.

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, but given that no candidates have yet been gathered, the "dummy" port value of 9 (Discard) MUST be used, as indicated in [I-D.ietf-ice-trickle], Section 5.1.
- o To properly indicate use of DTLS, the <proto> field MUST be set to "UDP/TLS/RTP/SAVPF", as specified in [RFC5764], Section 8, if the default candidate uses UDP transport, or "TCP/DTLS/RTP/SAVPF", as specified in [I-D.nandakumar-mmusic-proto-iana-registration] if the default candidate uses TCP transport.
- o If codec preferences have been set for the associated transceiver, media formats MUST be generated in the corresponding order, and MUST exclude any codecs not present in the codec preferences.
- o Unless excluded by the above restrictions, the media formats MUST include the mandatory audio/video codecs as specified in [I-D.ietf-rtcweb-audio](see Section 3) and [I-D.ietf-rtcweb-video](see Section 5).

The m= line MUST be followed immediately by a "c=" line, as specified in [RFC4566], Section 5.7. Again, as no candidates have yet been gathered, the "c=" line must contain the "dummy" value "IN IP4 0.0.0.0", as defined in [I-D.ietf-ice-trickle], Section 5.1.

[I-D.ietf-mmusic-sdp-mux-attributes] groups SDP attributes into different categories. To avoid unnecessary duplication when bundling, Section 8.1 of [I-D.ietf-mmusic-sdp-bundle-negotiation] specifies that attributes of category IDENTICAL or TRANSPORT should not be repeated in bundled m= sections.

The following attributes, which are of a category other than IDENTICAL or TRANSPORT, MUST be included in each m= section:

- o An "a=mid" line, as specified in [RFC5888], Section 4. When generating mid values, it is RECOMMENDED that the values be 3

bytes or less, to allow them to efficiently fit into the RTP header extension defined in [I-D.ietf-mmusic-sdp-bundle-negotiation], Section 11.

- o A direction attribute which is the same as that of the associated transceiver.
- o For each media format on the m= line, "a=rtpmap" and "a=fmtp" lines, as specified in [RFC4566], Section 6, and [RFC3264], Section 5.1.
- o If this m= section is for media with configurable frame sizes, e.g. audio, an "a=maxptime" line, indicating the smallest of the maximum supported frame sizes out of all codecs included above, as specified in [RFC4566], Section 6.
- o If this m= section is for video media, and there are known limitations on the size of images which can be decoded, an "a=imageattr" line, as specified in Section 3.6.
- 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 of the primary codec, as specified in [RFC4588], Section 8.1.
- o For each supported FEC mechanism, "a=rtpmap" and "a=fmtp" lines, as specified in [RFC4566], Section 6. The FEC mechanisms that MUST be supported are specified in [I-D.ietf-rtcweb-fec], Section 6, and specific usage for each media type is outlined in Sections 4 and 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 If the bundle policy for this PeerConnection is set to "max-bundle", and this is not the first m= section, or the bundle policy is set to "balanced", and this is not the first m= section for this media type, an "a=bundle-only" line.

- o If the RtpTransceiver has a sendrecv or sendonly direction:
 - * An "a=msid" line, as specified in [I-D.ietf-mmusic-msid], Section 2.
- o If the RtpTransceiver has a sendrecv or sendonly direction, and the application has specified RID values or has specified more than one encoding in the RtpSenders's parameters, an "a=rid" line for each encoding specified. The "a=rid" line is specified in [I-D.ietf-mmusic-rid], and its direction MUST be "send". If the application has chosen a RID value, it MUST be used as the rid-identifier; otherwise a RID value MUST be generated by the implementation. When generating RID values, it is RECOMMENDED that the values be 3 bytes or less, to allow them to efficiently fit into the RTP header extension defined in [I-D.ietf-avtext-rid], Section 11. If no encodings have been specified, or only one encoding is specified but without a RID value, then no "a=rid" lines are generated.
- o If the RtpTransceiver has a sendrecv or sendonly direction and more than one "a=rid" line has been generated, an "a=simulcast" line, with direction "send", as defined in [I-D.ietf-mmusic-sdp-simulcast], Section 6.2. The list of RIDs MUST include all of the RID identifiers used in the "a=rid" lines for this m= section.

The following attributes, which are of category IDENTICAL or TRANSPORT, MUST appear only in "m=" sections which either have a unique address or which are associated with the bundle-tag. (In initial offers, this means those "m=" sections which do not contain an "a=bundle-only" attribute.

- o "a=ice-ufraq" and "a=ice-pwd" lines, as specified in [RFC5245], Section 15.4.
- o An "a=fingerprint" line for each of the endpoint's certificates, as specified in [RFC4572], Section 5; the digest 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=dtls-id" line, as specified in [I-D.ietf-mmusic-dtls-sdp] Section 5.2.

- o An "a=rtcp" line, as specified in [RFC3605], Section 2.1, containing the dummy value "9 IN IP4 0.0.0.0", because no candidates have yet been gathered.
- 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.

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 "UDP/DTLS/SCTP" if the default candidate uses UDP transport, or "TCP/DTLS/SCTP" if the default candidate uses TCP transport [I-D.ietf-mmusic-sctp-sdp]. The "fmt" value MUST be set to "webrtc-datachannel" as specified in [I-D.ietf-mmusic-sctp-sdp], Section 4.1.

Within the data m= section, the "a=mid", "a=ice-ufrag", "a=ice-pwd", "a=fingerprint", "a=dtls-id", and "a=setup" lines MUST be included as mentioned above, along with an "a=fmtp:webrtc-datachannel" line and an "a=sctp-port" line referencing the SCTP port number as defined in [I-D.ietf-mmusic-sctp-sdp], Section 4.1.

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 MUST include the mid identifiers of each m= section. The effect of this is that the browser offers all m= sections as one bundle group. However, whether the m= sections are bundle-only or not depends on the bundle policy.

The next step is to generate session-level lip sync groups as defined in [RFC5888], Section 7. For each MediaStream referenced by more than one RtpTransceiver (by passing those MediaStreams as arguments to the addTrack and addTransceiver methods), a group of type "LS" MUST be added that contains the mid values for each RtpTransceiver.

Attributes which SDP permits to either be at the session level or the media level SHOULD generally be at the media level even if they are identical. This promotes readability, especially if one of a set of initially identical attributes is subsequently changed.

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.ietf-mmusic-sdp-mux-attributes] on how those attributes interact with bundle.

Note that these requirements are in some cases stricter than those of SDP. Implementations MUST be prepared to accept compliant SDP even if it would not conform to the requirements for generating SDP in this specification.

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 by one on each call to createOffer if the offer might differ from the output of the previous call to createOffer; implementations MAY opt to increment <session-version> on every call. The value of the generated <session-version> is independent of the <session-version> of the current local description; in particular, in the case where the current version is N, an offer is created with version N+1, and then that offer is rolled back so that the current version is again N, the next generated offer will still have version N+2.

Note that if the application creates an offer by reading currentLocalDescription instead of calling createOffer, the returned SDP may be different than when setLocalDescription was originally called, due to the addition of gathered ICE candidates, but the <session-version> will not have changed. There are no known scenarios in which this causes problems, but if this is a concern, the solution is simply to use createOffer to ensure a unique <session-version>.

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 If any RtpTransceiver has been added, and there exists an m= section with a zero port in the current local description or the current remote description, that m= section MUST be recycled by generating an m= section for the added RtpTransceiver as if the m= section were being added to the session description, placed at the same index as the m= section with a zero port.
- o If an RtpTransceiver is stopped and is not associated with an m= section, an m= section MUST NOT be generated for it. This prevents adding back RtpTransceivers whose m= sections were recycled and used for a new RtpTransceiver in a previous offer/answer exchange, as described above.
- o If an RtpTransceiver has been stopped and is associated with an m= section, and the m= section is not being recycled as described above, an m= section MUST be generated for it with the port set to zero and the "a=msid" line removed.
- o For RtpTransceivers that are not stopped, the "a=msid" line MUST stay the same if they are present in the current description.
- o Each "m=" and "c=" line MUST be filled in with the port, protocol, and address of the default candidate for the m= section, as described in [RFC5245], Section 4.3. If ICE checking has already completed for one or more candidate pairs and a candidate pair is in active use, then that pair MUST be used, even if ICE has not yet completed. Note that this differs from the guidance in [RFC5245], Section 9.1.2.2, which only refers to offers created when ICE has completed. In each case, if no RTP candidates have yet been gathered, dummy values MUST be used, as described above.
- o Each "a=mid" line MUST stay the same.
- o Each "a=ice-ufrag" and "a=ice-pwd" line MUST stay the same, unless the ICE configuration has changed (either changes to the supported STUN/TURN servers, or the ICE candidate policy), or the "IceRestart" option (Section 5.2.3.1 was specified. If the m= section is bundled into another m= section, it still MUST NOT contain any ICE credentials.
- o If the m= section is not bundled into another m= section, an "a=rtcp" attribute line MUST be added with of the default RTCP candidate, as indicated in [RFC5761], section 5.1.3.
- o If the m= section is not bundled into another m= section, for each candidate that has been gathered during the most recent gathering

phase (see Section 3.5.1), an "a=candidate" line MUST be added, as defined in [RFC5245], Section 4.3., paragraph 3. If candidate gathering for the section has completed, an "a=end-of-candidates" attribute MUST be added, as described in [I-D.ietf-ice-trickle], Section 9.3. If the m= section is bundled into another m= section, both "a=candidate" and "a=end-of-candidates" MUST be omitted.

- o For RtpTransceivers that are still present, the "a=msid" line MUST stay the same.
- o For RtpTransceivers that are still present, the "a=rid" lines MUST stay the same.
- o For RtpTransceivers that are still present, any "a=simulcast" line MUST stay the same.
- o If any RtpTransceiver has been stopped, the port MUST be set to zero and the "a=msid" line MUST be removed.
- o If any RtpTransceiver has been added, and there exists a m= section with a zero port in the current local description or the current remote description, that m= section MUST be recycled by generating a m= section for the added RtpTransceiver as if the m= section were being added to session description, except that instead of adding it, the generated m= section replaces the m= section with a zero port.

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.

In addition, for each non-recycled, 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 most recent answer which have not been excluded by the codec preferences of the associated transceiver.
- o The media formats on the m= line MUST be generated in the same order as in the current local description.

- o The RTP header extensions MUST only include those that are present in the most recent answer.
- o The RTCP feedback extensions MUST only include those that are present in the most recent answer.
- o The "a=rtcp" line MUST only be added if the most recent answer did not include an "a=rtcp-mux" line.
- o The "a=rtcp-mux" line MUST only be added if present in the most recent answer.
- o The "a=rtcp-mux-only" line MUST only be added if present in the most recent answer.
- o The "a=rtcp-rsize" line MUST only be added if present in the most recent answer.

The "a=group:BUNDLE" attribute MUST include the mid identifiers specified in the bundle group in the most recent answer, minus any m= sections that have been marked as rejected, plus any newly added or re-enabled m= sections. In other words, the bundle attribute must contain all m= sections that were previously bundled, as long as they are still alive, as well as any new m= sections.

The "LS" groups are generated in the same way as with initial offers.

5.2.3. Options Handling

The createOffer method takes as a parameter an RTCOfferOptions object. Special processing is performed when generating a SDP description if the following options are present.

5.2.3.1. IceRestart

If the "IceRestart" option 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 option is specified on an initial offer, it has no effect (since a new ICE ufrag and pwd are already generated). Similarly, if the ICE configuration has changed, this option has no effect, since new ufrag and pwd attributes will be generated automatically. This option is primarily useful for reestablishing connectivity in cases where failures are detected by the application.

5.2.3.2. VoiceActivityDetection

If the "VoiceActivityDetection" option is specified, with a value of "true", the offer MUST indicate support for silence suppression in the audio it receives by including comfort noise ("CN") codecs for each offered audio codec, as specified in [RFC3389], Section 5.1, except for codecs that have their own internal silence suppression support. For codecs that have their own internal silence suppression support, the appropriate fntp parameters for that codec MUST be specified to indicate that silence suppression for received audio is desired. For example, when using the Opus codec, the "usedtx=1" parameter would be specified in the offer. This option allows the endpoint to significantly reduce the amount of audio bandwidth it receives, at the cost of some fidelity, depending on the quality of the remote VAD algorithm.

If the "VoiceActivityDetection" option is specified, with a value of "false", the browser MUST NOT emit "CN" codecs. For codecs that have their own internal silence suppression support, the appropriate fntp parameters for that codec MUST be specified to indicate that silence suppression for received audio is not desired. For example, when using the Opus codec, the "usedtx=0" parameter would be specified in the offer.

Note that setting the "VoiceActivityDetection" parameter when generating an offer is a request to receive audio with silence suppression. It has no impact on whether the local endpoint does silence suppression for the audio it sends.

The "VoiceActivityDetection" option does not have any impact on the setting of the "vad" value in the signaling of the client to mixer audio level header extension described in [RFC6464], Section 4.

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, except that the "a=ice-options" line, with the "trickle" option as specified in [I-D.ietf-ice-trickle], Section 4, is only included if such an option was present in the offer.

The next step is to generate session-level lip sync groups as defined in [RFC5888], Section 7. For each group of type "LS" present in the offer, determine which of the local RtpTransceivers identified by that group's mid values reference a common local MediaStream (as specified in the addTrack and addTransceiver methods). If at least two such RtpTransceivers exist, a group of type "LS" with the mid values of these RtpTransceivers MUST be added. Otherwise, this indicates a difference of opinion between the offerer and answerer regarding lip sync status, and as such, the offered group MUST be ignored and no corresponding "LS" group generated.

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.

The next step is to go through each offered m= section. Each offered m= section will have an associated RtpTransceiver, as described in Section 5.9. If there are more RtpTransceivers than there are m= sections, the unmatched RtpTransceivers will need to be associated in a subsequent offer.

For each offered m= section, if any of the following conditions are true, 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 further processing for this m= section can be skipped:

- o The associated RtpTransceiver has been stopped.
- o No supported codec is present in the offer.

- o The bundle policy is "max-bundle", and this is not the first m= section or in the same bundle group as the first m= section.
- o The bundle policy is "balanced", and this is not the first m= section for this media type or in the same bundle group as the first m= section for this media type.
- o The RTP/RTCP multiplexing policy is "require" and the m= section doesn't contain an "a=rtcp-mux" attribute.

Otherwise, each m= section in the answer should then be generated as specified in [RFC3264], Section 6.1. For the m= line itself, the following rules must be followed:

- o The port value would normally be set to the port of the default ICE candidate for this m= section, but given that no candidates have yet been gathered, the "dummy" port value of 9 (Discard) MUST be used, as indicated in [I-D.ietf-ice-trickle], Section 5.1.
- o The <proto> field MUST be set to exactly match the <proto> field for the corresponding m= line in the offer.
- o If codec preferences have been set for the associated transceiver, media formats MUST be generated in the corresponding order, and MUST exclude any codecs not present in the codec preferences or not present in the offer.
- o Unless excluded by the above restrictions, the media formats MUST include the mandatory audio/video codecs as specified in [I-D.ietf-rtcweb-audio](see Section 3) and [I-D.ietf-rtcweb-video](see Section 5).

The m= line MUST be followed immediately by a "c=" line, as specified in [RFC4566], Section 5.7. Again, as no candidates have yet been gathered, the "c=" line must contain the "dummy" value "IN IP4 0.0.0.0", as defined in [I-D.ietf-ice-trickle], Section 5.1.

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 contain the following attributes (which are of attribute types other than IDENTICAL and TRANSPORT):

- o If and only 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 A direction attribute, determined by applying the rules regarding the offered direction specified in [RFC3264], Section 6.1, and then intersecting with the direction of the associated RtpTransceiver. For example, in the case where an m= section is offered as "sendonly", and the local transceiver is set to "sendrecv", the result in the answer is a "recvonly" direction.
- o For each media format on the m= line, "a=rtpmap" and "a=fmtp" lines, as specified in [RFC4566], Section 6, and [RFC3264], Section 6.1.
- o If this m= section is for media with configurable frame sizes, e.g. audio, an "a=maxptime" line, indicating the smallest of the maximum supported frame sizes out of all codecs included above, as specified in [RFC4566], Section 6.
- o If this m= section is for video media, and there are known limitations on the size of images which can be decoded, an "a=imageattr" line, as specified in Section 3.6.
- 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 of the primary codec, as specified in [RFC4588], Section 8.1.
- o For each supported FEC mechanism, "a=rtpmap" and "a=fmtp" lines, as specified in [RFC4566], Section 6. The FEC mechanisms that MUST be supported are specified in [I-D.ietf-rtcweb-fec], Section 6, and specific usage for each media type is outlined in Sections 4 and 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 the RtpTransceiver has a sendrecv or sendonly direction:

- * An "a=msid" line, as specified in [I-D.ietf-mmusic-msid], Section 2.

Each m= section which is not bundled into another m= section, MUST contain the following attributes (which are of category IDENTICAL or TRANSPORT):

- o "a=ice-ufrag" and "a=ice-pwd" lines, as specified in [RFC5245], Section 15.4.
- o An "a=fingerprint" line for each of the endpoint's certificates, as specified in [RFC4572], Section 5; the digest 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. The role value MUST be consistent with the existing DTLS connection, if one exists and is being continued.
- o An "a=dtls-id" line, as specified in [I-D.ietf-mmusic-dtls-sdp] Section 5.3.
- o If present in the offer, an "a=rtcp-mux" line, as specified in [RFC5761], Section 5.1.1. Otherwise, an "a=rtcp" line, as specified in [RFC3605], Section 2.1, containing the dummy value "9 IN IP4 0.0.0.0" (because no candidates have yet been gathered).
- o If present in the offer, an "a=rtcp-rsize" line, as specified in [RFC5506], Section 5.

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> and "fmt" fields MUST be set to exactly match the fields in the offer.

Within the data m= section, the "a=mid", "a=ice-ufrag", "a=ice-pwd", "a=candidate", "a=fingerprint", "a=dtls-id", and "a=setup" lines MUST be included under the conditions described above, along with an "a=fmtp:webrtc-datachannel" line and an "a=sctp-port" line referencing the SCTP port number as defined in [I-D.ietf-mmusic-sctp-sdp], Section 4.1.

If "a=group" attributes with semantics of "BUNDLE" are offered, corresponding session-level "a=group" attributes MUST be added as specified in [RFC5888]. These attributes MUST have semantics

"BUNDLE", and MUST include the all mid identifiers from the offered bundle groups that have not been rejected. Note that regardless of the presence of "a=bundle-only" in the offer, no m= sections in the answer should have an "a=bundle-only" line.

Attributes that are common between all m= sections MAY be moved to session-level, if explicitly defined to be valid at session-level.

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

5.3.2. Subsequent Answers

When createAnswer 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 answer.

If the initial answer was not applied using setLocalDescription, meaning the PeerConnection is still in the "have-remote-offer" state, the steps for generating an initial answer 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 from the previously generated answer.

If any session description was previously supplied to setLocalDescription, an answer is generated by following the steps in the "have-remote-offer" state above, along with these exceptions:

- o The "s=" and "t=" lines MUST stay the same.
- o Each "m=" and "c=" line MUST be filled in with the port and address of the default candidate for the m= section, as described in [RFC5245], Section 4.3. Note, however, that the m= line protocol need not match the default candidate, because this protocol value must instead match what was supplied in the offer, as described above.
- o The media formats on the m= line MUST be generated in the same order as in the current local description.
- o Each "a=ice-ufraq" and "a=ice-pwd" line MUST stay the same, unless the m= section is restarting, in which case new ICE credentials must be created as specified in [RFC5245], Section 9.2.1.1. If the m= section is bundled into another m= section, it still MUST NOT contain any ICE credentials.

- o If the m= section is not bundled into another m= section and RTCP multiplexing is not active, an "a=rtcp" attribute line MUST be filled in with the port and address of the default RTCP candidate. If no RTCP candidates have yet been gathered, dummy values MUST be used, as described in the initial answer section above.
- o If the m= section is not bundled into another m= section, for each candidate that has been gathered during the most recent gathering phase (see Section 3.5.1), an "a=candidate" line MUST be added, as defined in [RFC5245], Section 4.3., paragraph 3. If candidate gathering for the section has completed, an "a=end-of-candidates" attribute MUST be added, as described in [I-D.ietf-ice-trickle], Section 9.3. If the m= section is bundled into another m= section, both "a=candidate" and "a=end-of-candidates" MUST be omitted.
- o For RtpTransceivers that are not stopped, the "a=msid" line MUST stay the same.

5.3.3. Options Handling

The createAnswer method takes as a parameter an RTCAnswerOptions object. The set of parameters for RTCAnswerOptions is different than those supported in RTCOfferOptions; the IceRestart option is unnecessary, as ICE credentials will automatically be changed for all m= lines where the offerer chose to perform ICE restart.

The following options are supported in RTCAnswerOptions.

5.3.3.1. VoiceActivityDetection

Silence suppression in the answer is handled as described in Section 5.2.3.2, with one exception: if support for silence suppression was not indicated in the offer, the VoiceActivityDetection parameter has no effect, and the answer should be generated as if VoiceActivityDetection was set to false. This is done on a per-codec basis (e.g., if the offerer somehow offered support for CN but set "usedtx=0" for Opus, setting VoiceActivityDetection to true would result in an answer with CN codecs and "usedtx=0").

5.4. Modifying an Offer or Answer

The SDP returned from createOffer or createAnswer MUST NOT be changed before passing it to setLocalDescription. If precise control over the SDP is needed, the aforementioned createOffer/createAnswer options or RTPSender APIs MUST be used.

Note that the application MAY modify the SDP to reduce the capabilities in the offer it sends to the far side (post-setLocalDescription) or the offer that it installs from the far side (pre-setRemoteDescription), as long as it remains a valid SDP offer and specifies a subset of what was in the original offer. This is safe because the answer is not permitted to expand capabilities, and therefore will just respond to what is present in the offer.

The application SHOULD NOT modify the SDP in the answer it transmits, as the answer contains the negotiated capabilities, and this can cause the two sides to have different ideas about what exactly was negotiated.

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.

5.5. Processing a Local Description

When a SessionDescription is supplied to setLocalDescription, the following steps MUST be performed:

- o First, the type of the SessionDescription is checked against the current state of the PeerConnection:
 - * If the type is "offer", the PeerConnection state MUST be either "stable" or "have-local-offer".
 - * If the type is "pranswer" or "answer", the PeerConnection state MUST be either "have-remote-offer" or "have-local-pranswer".
- o If the type is not correct for the current state, processing MUST stop and an error MUST be returned.
- o Next, the SessionDescription is parsed into a data structure, as described in the Section 5.7 section below. If parsing fails for any reason, processing MUST stop and an error MUST be returned.
- o Finally, the parsed SessionDescription is applied as described in the Section 5.8 section below.

5.6. Processing a Remote Description

When a SessionDescription is supplied to setRemoteDescription, the following steps MUST be performed:

- o First, the type of the SessionDescription is checked against the current state of the PeerConnection:
 - * If the type is "offer", the PeerConnection state MUST be either "stable" or "have-remote-offer".
 - * If the type is "pranswer" or "answer", the PeerConnection state MUST be either "have-local-offer" or "have-remote-pranswer".
- o If the type is not correct for the current state, processing MUST stop and an error MUST be returned.
- o Next, the SessionDescription is parsed into a data structure, as described in the Section 5.7 section below. If parsing fails for any reason, processing MUST stop and an error MUST be returned.
- o Finally, the parsed SessionDescription is applied as described in the Section 5.9 section below.

5.7. Parsing a Session Description

When a SessionDescription of any type is supplied to setLocal/RemoteDescription, the implementation must parse it and reject it if it is invalid. The exact details of this process are explained below.

The SDP contained in the session description object consists of a sequence of text lines, each containing a key-value expression, as described in [RFC4566], Section 5. The SDP is read, line-by-line, and converted to a data structure that contains the deserialized information. However, SDP allows many types of lines, not all of which are relevant to JSEP applications. For each line, the implementation will first ensure it is syntactically correct according to its defining ABNF, check that it conforms to [RFC4566] and [RFC3264] semantics, and then either parse and store or discard the provided value, as described below.

If any line is not well-formed, or cannot be parsed as described, the parser MUST stop with an error and reject the session description, even if the value is to be discarded. This ensures that implementations do not accidentally misinterpret ambiguous SDP.

5.7.1. Session-Level Parsing

First, the session-level lines are checked and parsed. These lines MUST occur in a specific order, and with a specific syntax, as defined in [RFC4566], Section 5. Note that while the specific line types (e.g. "v=", "c=") MUST occur in the defined order, lines of the same type (typically "a=") can occur in any order, and their ordering is not meaningful.

The following non-attribute lines are not meaningful in the JSEP context and MAY be discarded once they have been checked.

The "c=" line MUST be checked for syntax but its value is not used. This supersedes the guidance in [RFC5245], Section 6.1, to use "ice-mismatch" to indicate mismatches between "c=" and the candidate lines; because JSEP always uses ICE, "ice-mismatch" is not useful in this context.

The "i=", "u=", "e=", "p=", "t=", "r=", "z=", and "k=" lines are not used by this specification; they MUST be checked for syntax but their values are not used.

The remaining non-attribute lines are processed as follows:

The "v=" line MUST have a version of 0, as specified in [RFC4566], Section 5.1.

The "o=" line MUST be parsed as specified in [RFC4566], Section 5.2.

The "b=" line, if present, MUST be parsed as specified in [RFC4566], Section 5.8, and the bwtype and bandwidth values stored.

Finally, the attribute lines are processed. Specific processing MUST be applied for the following session-level attribute ("a=") lines:

- o Any "a=group" lines are parsed as specified in [RFC5888], Section 5, and the group's semantics and mids are stored.
- o If present, a single "a=ice-lite" line is parsed as specified in [RFC5245], Section 15.3, and a value indicating the presence of ice-lite is stored.
- o If present, a single "a=ice-ufrag" line is parsed as specified in [RFC5245], Section 15.4, and the ufrag value is stored.

- o If present, a single "a=ice-pwd" line is parsed as specified in [RFC5245], Section 15.4, and the password value is stored.
- o If present, a single "a=ice-options" line is parsed as specified in [RFC5245], Section 15.5, and the set of specified options is stored.
- o Any "a=fingerprint" lines are parsed as specified in [RFC4572], Section 5, and the set of fingerprint and algorithm values is stored.
- o If present, a single "a=setup" line is parsed as specified in [RFC4145], Section 4, and the setup value is stored.
- o If present, a single "a=dtls-id" line is parsed as specified in [I-D.ietf-mmusic-dtls-sdp] Section 5, and the dtls-id value is stored.
- o Any "a=extmap" lines are parsed as specified in [RFC5285], Section 5, and their values are stored.

Once all the session-level lines have been parsed, processing continues with the lines in media sections.

5.7.2. Media Section Parsing

Like the session-level lines, the media session lines MUST occur in the specific order and with the specific syntax defined in [RFC4566], Section 5.

The "m=" line itself MUST be parsed as described in [RFC4566], Section 5.14, and the media, port, proto, and fmt values stored.

Following the "m=" line, specific processing MUST be applied for the following non-attribute lines:

- o As with the "c=" line at the session level, the "c=" line MUST be parsed according to [RFC4566], Section 5.7, but its value is not used.
- o The "b=" line, if present, MUST be parsed as specified in [RFC4566], Section 5.8, and the bwtype and bandwidth values stored.

Specific processing MUST also be applied for the following attribute lines:

- o If present, a single "a=ice-ufrag" line is parsed as specified in [RFC5245], Section 15.4, and the ufrag value is stored.
- o If present, a single "a=ice-pwd" line is parsed as specified in [RFC5245], Section 15.4, and the password value is stored.
- o If present, a single "a=ice-options" line is parsed as specified in [RFC5245], Section 15.5, and the set of specified options is stored.
- o Any "a=candidate" attributes MUST be parsed as specified in [RFC5245], Section 15.1, and their values stored.
- o Any "a=remote-candidates" attributes MUST be parsed as specified in [RFC5245], Section 15.2, but their values are ignored.
- o If present, a single "a=end-of-candidates" attribute MUST be parsed as specified in [I-D.ietf-ice-trickle], Section 8.2, and its presence or absence flagged and stored.
- o Any "a=fingerprint" lines are parsed as specified in [RFC4572], Section 5, and the set of fingerprint and algorithm values is stored.

If the "m=" proto value indicates use of RTP, as described in the Section 5.1.3 section above, the following attribute lines MUST be processed:

- o The "m=" fmt value MUST be parsed as specified in [RFC4566], Section 5.14, and the individual values stored.
- o Any "a=rtpmap" or "a=fmtp" lines MUST be parsed as specified in [RFC4566], Section 6, and their values stored.
- o If present, a single "a=ptime" line MUST be parsed as described in [RFC4566], Section 6, and its value stored.
- o If present, a single "a=maxptime" line MUST be parsed as described in [RFC4566], Section 6, and its value stored.
- o If present, a single direction attribute line (e.g. "a=sendrecv") MUST be parsed as described in [RFC4566], Section 6, and its value stored.
- o Any "a=ssrc" or "a=ssrc-group" attributes MUST be parsed as specified in [RFC5576], Sections 4.1-4.2, and their values stored.

- o Any "a=extmap" attributes MUST be parsed as specified in [RFC5285], Section 5, and their values stored.
- o Any "a=rtcp-fb" attributes MUST be parsed as specified in [RFC4585], Section 4.2., and their values stored.
- o If present, a single "a=rtcp-mux" attribute MUST be parsed as specified in [RFC5761], Section 5.1.1, and its presence or absence flagged and stored.
- o If present, a single "a=rtcp-mux-only" attribute MUST be parsed as specified in [I-D.ietf-mmusic-mux-exclusive], Section 3, and its presence or absence flagged and stored.
- o If present, a single "a=rtcp-rsize" attribute MUST be parsed as specified in [RFC5506], Section 5, and its presence or absence flagged and stored.
- o If present, a single "a=rtcp" attribute MUST be parsed as specified in [RFC3605], Section 2.1, but its value is ignored, as this information is superfluous when using ICE.
- o If present, a single "a=msid" attribute MUST be parsed as specified in [I-D.ietf-mmusic-msid], Section 3.2, and its value stored.
- o Any "a=imageattr" attributes MUST be parsed as specified in [RFC6236], Section 3, and their values stored.
- o Any "a=rid" lines MUST be parsed as specified in [I-D.ietf-mmusic-rid], Section 10, and their values stored.
- o If present, a single "a=simulcast" line MUST be parsed as specified in [I-D.ietf-mmusic-sdp-simulcast], and its values stored.

Otherwise, if the "m=" proto value indicates use of SCTP, the following attribute lines MUST be processed:

- o The "m=" fmt value MUST be parsed as specified in [I-D.ietf-mmusic-sctp-sdp], Section 4.3, and the application protocol value stored.
- o An "a=sctp-port" attribute MUST be present, and it MUST be parsed as specified in [I-D.ietf-mmusic-sctp-sdp], Section 5.2, and the value stored.

- o If present, a single "a=max-message-size" attribute MUST be parsed as specified in [I-D.ietf-mmusic-sctp-sdp], Section 6, and the value stored. Otherwise, use the specified default.

5.7.3. Semantics Verification

Assuming parsing completes successfully, the parsed description is then evaluated to ensure internal consistency as well as proper support for mandatory features. Specifically, the following checks are performed:

- o For each m= section, valid values for each of the mandatory-to-use features enumerated in Section 5.1.2 MUST be present. These values MAY either be present at the media level, or inherited from the session level.
 - * ICE ufrag and password values, which MUST comply with the size limits specified in [RFC5245], Section 15.4.
 - * dtls-id value, which MUST be set according to [I-D.ietf-mmusic-dtls-sdp] Section 5. If this is a re-offer and the dtls-id value is different from that presently in use, the DTLS connection is not being continued and the remote description MUST be part of an ICE restart, together with new ufrag and password values. If this is an answer, the dtls-id value, if present, MUST be the same as in the offer.
 - * DTLS setup value, which MUST be set according to the rules specified in [RFC5763], Section 5 and MUST be consistent with the selected role of the current DTLS connection, if one exists and is being continued.
 - * DTLS fingerprint values, where at least one fingerprint MUST be present.
- o All RID values referenced in an "a=simulcast" line MUST exist as "a=rid" lines.
- o Each m= section is also checked to ensure prohibited features are not used. If this is a local description, the "ice-lite" attribute MUST NOT be specified.

If this session description is of type "pranswer" or "answer", the following additional checks are applied:

- o The session description must follow the rules defined in [RFC3264], Section 6, including the requirement that the number of

m= sections MUST exactly match the number of m= sections in the associated offer.

- o For each m= section, the media type and protocol values MUST exactly match the media type and protocol values in the corresponding m= section in the associated offer.

5.8. Applying a Local Description

The following steps are performed at the media engine level to apply a local description.

First, the parsed parameters are checked to ensure that they have not been altered after their generation in createOffer/createAnswer, as discussed in Section 5.4; otherwise, processing MUST stop and an error MUST be returned.

Next, media sections are processed. For each media section, the following steps MUST be performed; if any parameters are out of bounds, or cannot be applied, processing MUST stop and an error MUST be returned.

- o If this media section is new, begin gathering candidates for it, as defined in [RFC5245], Section 4.1.1, unless it has been marked as bundle-only.
- o Or, if the ICE ufrag and password values have changed, and it has not been marked as bundle-only, trigger the ICE Agent to start an ICE restart, and begin gathering new candidates for the media section as described in [RFC5245], Section 9.1.1.1. If this description is an answer, also start checks on that media section as defined in [RFC5245], Section 9.3.1.1.
- o If the media section proto value indicates use of RTP:
 - * If there is no RtpTransceiver associated with this m= section (which should only happen when applying an offer), find one and associate it with this m= section according to the following steps:
 - + Find the RtpTransceiver that corresponds to the m= section with the same MID in the created offer.
 - + Set the value of the RtpTransceiver's mid attribute to the MID of the m= section.
 - * If RTCP mux is indicated, prepare to demux RTP and RTCP from the RTP ICE component, as specified in [RFC5761],

Section 5.1.1. If RTCP mux is not indicated, but was indicated in a previous description, this MUST result in an error.

- * For each specified RTP header extension, establish a mapping between the extension ID and URI, as described in section 6 of [RFC5285]. If any indicated RTP header extension is not supported, this MUST result in an error.
- * If the MID header extension is supported, prepare to demux RTP data intended for this media section based on the MID header extension, as described in [I-D.ietf-mmusic-msid], Section 3.2.
- * For each specified media format, establish a mapping between the payload type and the actual media format, as described in [RFC3264], Section 6.1. If any indicated media format is not supported, this MUST result in an error.
- * For each specified "rtx" media format, establish a mapping between the RTX payload type and its associated primary payload type, as described in [RFC4588], Sections 8.6 and 8.7. If any referenced primary payload types are not present, this MUST result in an error.
- * If the directional attribute is of type "sendrecv" or "recvonly", enable receipt and decoding of media.

Finally, if this description is of type "pranswer" or "answer", follow the processing defined in the Section 5.10 section below.

5.9. Applying a Remote Description

If the answer contains any "a=ice-options" attributes where "trickle" is listed as an attribute, update the PeerConnection canTrickle property to be true. Otherwise, set this property to false.

The following steps are performed at the media engine level to apply a remote description.

The following steps MUST be performed for attributes at the session level; if any parameters are out of bounds, or cannot be applied, processing MUST stop and an error MUST be returned.

- o For any specified "CT" bandwidth value, set this as the limit for the maximum total bitrate for all m= sections, as specified in Section 5.8 of [RFC4566]. The implementation can decide how to allocate the available bandwidth between m= sections to simultaneously meet any limits on individual m= sections, as well as this overall session limit.

- o For any specified "RR" or "RS" bandwidth values, handle as specified in [RFC3556], Section 2.
- o Any "AS" bandwidth value MUST be ignored, as the meaning of this construct at the session level is not well defined.

For each media section, the following steps MUST be performed; if any parameters are out of bounds, or cannot be applied, processing MUST stop and an error MUST be returned.

- o If the ICE ufrag or password changed from the previous remote description, then an ICE restart is needed, as described in Section 9.1.1.1 of [RFC5245] If the description is of type "offer", mark that an ICE restart is needed. If the description is of type "answer" and the current local description is also an ICE restart, then signal the ICE agent to begin checks as described in Section 9.3.1.1 of [RFC5245]. An answer MUST change the ufrag and password in an answer if and only if ICE is restarting, as described in Section 9.2.1.1 of [RFC5245].
- o Configure the ICE components associated with this media section to use the supplied ICE remote ufrag and password for their connectivity checks.
- o Pair any supplied ICE candidates with any gathered local candidates, as described in Section 5.7 of [RFC5245] and start connectivity checks with the appropriate credentials.
- o If an "a=end-of-candidates" attribute is present, process the end-of-candidates indication as described in [I-D.ietf-ice-trickle] Section 11.
- o If the media section proto value indicates use of RTP:
 - * If the m= section is being recycled (see Section 5.2.2), dissociate the currently associated RtpTransceiver by setting its mid attribute to null.
 - * If the m= section is not associated with any RtpTransceiver (possibly because it was dissociated in the previous step), either find an RtpTransceiver or create one according to the following steps:
 - + If the m= section is sendrecv or recvonly, and there are RtpTransceivers of the same type that were added to the PeerConnection by addTrack and are not associated with any m= section and are not stopped, find the first (according to

the canonical order described in Section 5.2.1) such RtpTransceiver.

- + If no RtpTransceiver was found in the previous step, create one with a recvonly direction.
- + Associate the found or created RtpTransceiver with the m= section by setting the value of the RtpTransceiver's mid attribute to the MID of the m= section. If the m= section does not include a MID (i.e., the remote side does not support the MID extension), generate a value for the RtpTransceiver mid attribute, following the guidance for "a=mid" mentioned in Section 5.2.1.
- * For each specified media format that is also supported by the local implementation, establish a mapping between the specified payload type and the media format, as described in [RFC3264], Section 6.1. Specifically, this means that the implementation records the payload type to be used in outgoing RTP packets when sending each specified media format, as well as the relative preference for each format that is indicated in their ordering. If any indicated media format is not supported by the local implementation, it MUST be ignored.
- * For each specified "rtx" media format, establish a mapping between the RTX payload type and its associated primary payload type, as described in [RFC4588], Section 4. If any referenced primary payload types are not present, this MUST result in an error.
- * For each specified fntp parameter that is supported by the local implementation, enable them on the associated media formats.
- * For each specified RTP header extension that is also supported by the local implementation, establish a mapping between the extension ID and URI, as described in [RFC5285], Section 5. Specifically, this means that the implementation records the extension ID to be used in outgoing RTP packets when sending each specified header extension. If any indicated RTP header extension is not supported by the local implementation, it MUST be ignored.
- * For each specified RTCP feedback mechanism that is supported by the local implementation, enable them on the associated media formats.

- * For any specified "TIAS" bandwidth value, set this value as a constraint on the maximum RTP bitrate to be used when sending media, as specified in [RFC3890]. If a "TIAS" value is not present, but an "AS" value is specified, generate a "TIAS" value using this formula:

$$\text{TIAS} = \text{AS} * 1000 * 0.95 - 50 * 40 * 8$$

The 50 is based on 50 packets per second, the 40 is based on an estimate of total header size, the 1000 changes the unit from kbps to bps (as required by TIAS), and the 0.95 is to allocate 5% to RTCP. If more accurate control of bandwidth is needed, "TIAS" should be used instead of "AS".

- * For any "RR" or "RS" bandwidth values, handle as specified in [RFC3556], Section 2.
- * Any specified "CT" bandwidth value MUST be ignored, as the meaning of this construct at the media level is not well defined.
- * If the media section is of type audio:
 - + For each specified "CN" media format, enable DTX for all supported media formats with the same clockrate, as described in [RFC3389], Section 5, except for formats that have their own internal DTX mechanisms. DTX for such formats (e.g., Opus) is controlled via fntp parameters, as discussed in Section 5.2.3.2.
 - + For each specified "telephone-event" media format, enable DTMF transmission for all supported media formats with the same clockrate, as described in [RFC4733], Section 2.5.1.2. If the application attempts to transmit DTMF when using a media format that does not have a corresponding telephone-event format, this MUST result in an error.
 - + For any specified "ptime" value, configure the available media formats to use the specified packet size. If the specified size is not supported for a media format, use the next closest value instead.

Finally, if this description is of type "pranswer" or "answer", follow the processing defined in the Section 5.10 section below.

5.10. Applying an Answer

In addition to the steps mentioned above for processing a local or remote description, the following steps are performed when processing a description of type "pranswer" or "answer".

For each media section, the following steps MUST be performed:

- o If the media section has been rejected (i.e. port is set to zero in the answer), stop any reception or transmission of media for this section, and discard any associated ICE components, as described in Section 9.2.1.3 of [RFC5245].
- o If the remote DTLS fingerprint has been changed or the dtls-id has changed, tear down the DTLS connection. If a DTLS connection needs to be torn down but the answer does not indicate an ICE restart, an error MUST be generated. If an ICE restart is performed without a change in dtls-id or fingerprint, then the same DTLS connection is continued over the new ICE channel.
- o If no valid DTLS connection exists, prepare to start a DTLS connection, using the specified roles and fingerprints, on any underlying ICE components, once they are active.
- o If the media section proto value indicates use of RTP:
 - * If the media section references any media formats, RTP header extensions, or RTCP feedback mechanisms that were not present in the corresponding media section in the offer, this indicates a negotiation problem and MUST result in an error.
 - * If the media section has RTCP mux enabled, discard any RTCP component, and begin or continue muxing RTCP over the RTP component, as specified in [RFC5761], Section 5.1.3. Otherwise, prepare to transmit RTCP over the RTCP component; if no RTCP component exists, because RTCP mux was previously enabled, this MUST result in an error.
 - * If the media section has reduced-size RTCP enabled, configure the RTCP transmission for this media section to use reduced-size RTCP, as specified in [RFC5506].
 - * If the directional attribute in the answer is of type "sendrecv" or "sendonly", choose the media format to send as the most preferred media format from the remote description that is also present in the answer, as described in [RFC3264], Sections 6.1 and 7, and start transmitting RTP media once the underlying transport layers have been established. If a SSRC

has not already been chosen for this outgoing RTP stream, choose a random one.

- * The payload type mapping from the remote description is used to determine payload types for the outgoing RTP streams, including the payload type for the send media format chosen above. Any RTP header extensions that were negotiated should be included in the outgoing RTP streams, using the extension mapping from the remote description; if the RID header extension has been negotiated, and RID values are specified, include the RID header extension in the outgoing RTP streams, as indicated in [I-D.ietf-mmusic-rid], Section 4.
 - * If simulcast has been negotiated, send the number of Source RTP Streams as specified in [I-D.ietf-mmusic-sdp-simulcast], Section 6.2.2.
 - * If the send media format chosen above has a corresponding "rtx" media format, or a FEC mechanism has been negotiated, establish a Redundancy RTP Stream with a random SSRC for each Source RTP Stream, and start or continue transmitting RTX/FEC packets as needed.
 - * If the send media format chosen above has a corresponding "red" media format of the same clockrate, allow redundant encoding using the specified format for resiliency purposes, as discussed in [I-D.ietf-rtcweb-fec], Section 3.2. Note that unlike RTX or FEC media formats, the "red" format is transmitted on the Source RTP Stream, not the Redundancy RTP Stream.
 - * Enable the RTCP feedback mechanisms referenced in the media section for all Source RTP Streams using the specified media formats. Specifically, begin or continue sending the requested feedback types and reacting to received feedback, as specified in [RFC4585], Section 4.2. When sending RTCP feedback, use the SSRC of an outgoing Source RTP Stream as the RTCP sender SSRC; if no outgoing Source RTP Stream exists, choose a random one.
 - * If the directional attribute is of type "recvonly" or "inactive", stop transmitting all RTP media, but continue sending RTCP, as described in [RFC3264], Section 5.1.
- o If the media section proto value indicates use of SCTP:
 - * If no SCTP association yet exists, prepare to initiate a SCTP association over the associated ICE component and DTLS connection, using the local SCTP port value from the local

description, and the remote SCTP port value from the remote description, as described in [I-D.ietf-mmusic-sctp-sdp], Section 10.2.

If the answer contains valid bundle groups, discard any ICE components for the m= sections that will be bundled onto the primary ICE components in each bundle, and begin muxing these m= sections accordingly, as described in [I-D.ietf-mmusic-sdp-bundle-negotiation], Section 8.2.

6. Processing RTP/RTCP packets

Note: The following algorithm does not yet have WG consensus but is included here as something concrete for the working group to discuss.

When an RTP packet is received by a transport and passes SRTP authentication, that packet needs to be routed to the correct RtpReceiver. For each transport, the following steps MUST be followed to prepare to route packets:

Construct a table mapping MID to RtpReceiver for each RtpReceiver configured to receive from this transport.

Construct a table mapping incoming SSRC to RtpReceiver for each RtpReceiver configured to receive from this transport and for each SSRC that RtpReceiver is configured to receive. Some of the SSRCs may be present in the m= section corresponding to that RtpReceiver in the remote description.

Construct a table mapping outgoing SSRC to RtpSender for each RtpSender configured to transmit from this transport and for each SSRC that RtpSender is configured to use when sending.

Construct a table mapping payload type to RtpReceiver for each RtpReceiver configured to receive from this transport and for each payload type that RtpReceiver is configured to receive. The payload types of a given RtpReceiver are found in the m= section corresponding to that RtpReceiver in the local description. If any payload type could map to more than one RtpReceiver, map to the RtpReceiver whose m= section appears earliest in the local description.

As RtpTransceivers (and, thus, RtpReceivers) are added, removed, stopped, or reconfigured, the tables above must also be updated.

For each RTP packet received, the following steps MUST be followed to route the packet:

If the packet has a MID and that MID is not in the table mapping MID to RtpReceiver, drop the packet and stop.

If the packet has a MID and that MID is in the table mapping MID to RtpReceiver, update the incoming SSRC mapping table to include an entry that maps the packet's SSRC to the RtpReceiver for that MID.

If the packet's SSRC is in the incoming SSRC mapping table, deliver the packet to the associated RtpReceiver and stop.

If the packet's payload type is in the payload type table, update the the incoming SSRC mapping table to include an entry that maps the packet's SSRC to the RtpReceiver for that payload type. In addition, deliver the packet to the associated RtpReceiver and stop.

Otherwise, drop the packet.

For each RTCP packet received (including each RTCP packet that is part of a compound RTCP packet), the following type-specific handling MUST be performed to route the packet:

If the packet is of type SR, and the sender SSRC for the packet is found in the incoming SSRC table, deliver a copy of the packet to the RtpReceiver associated with that SSRC. In addition, for each report block in the report whose SSRC is found in the outgoing SSRC table, deliver a copy of the RTCP packet to the RtpSender associated with that SSRC.

If the packet is of type RR, for each report block in the packet whose SSRC is found in the outgoing SSRC table, deliver a copy of the RTCP packet to the RtpSender associated with that SSRC.

If the packet is of type SDDES, and the sender SSRC for the packet is found in the incoming SSRC table, deliver the packet to the RtpReceiver associated with that SSRC. In addition, for each chunk in the packet that contains a MID that is in the table mapping MID to RtpReceiver, update the incoming SSRC mapping table to include an entry that maps the SSRC for that chunk to the RtpReceiver associated with that MID. (This case can occur when RTCP for a source is received before any RTP packets.)

If the packet is of type BYE, for each SSRC indicated in the packet that is found in the incoming SSRC table, deliver a copy of the packet to the RtpReceiver associated with that SSRC.

If the packet is of type RTPFB or PSFB, as defined in [RFC4585], and the media source SSRC for the packet is found in the outgoing SSRC table, deliver the packet to the RtpSender associated with that SSRC.

After packets are routed to the RtpReceiver, further processing of the RTP packets is done at the RtpReceiver level. This includes using [I-D.ietf-mmusic-rid] to distinguish between multiple Encoded Streams, as well as determine which Source RTP stream should be repaired by a given Redundancy RTP stream. If the RTP packet's PT does not match any codec in use by the RtpReceiver, the packet will be dropped.

7. Examples

Note that this example section shows several SDP fragments. To format in 72 columns, some of the lines in SDP have been split into multiple lines, where leading whitespace indicates that a line is a continuation of the previous line. In addition, some blank lines have been added to improve readability but are not valid in SDP.

More examples of SDP for WebRTC call flows can be found in [I-D.nandakumar-rtcweb-sdp].

7.1. Simple Example

This section shows a very simple example that sets up a minimal audio / video call between two browsers and does not use trickle ICE. The example in the following section provides a more realistic example of what would happen in a normal browser to browser connection.

The flow shows Alice's browser initiating the session to Bob's browser. The messages from Alice's JS to Bob's JS are assumed to flow over some signaling protocol via a web server. The JS on both Alice's side and Bob's side waits for all candidates before sending the offer or answer, so the offers and answers are complete. Trickle ICE is not used. Both Alice and Bob are using the default policy of balanced.

```
//          set up local media state
AliceJS->AliceUA: create new PeerConnection
AliceJS->AliceUA: addTrack with two tracks: audio and video
AliceJS->AliceUA: createOffer to get offer
AliceJS->AliceUA: setLocalDescription with offer
AliceUA->AliceJS: multiple onicecandidate events with candidates

//          wait for ICE gathering to complete
AliceUA->AliceJS: onicecandidate event with null candidate
AliceJS->AliceUA: get |offer-A1| from pendingLocalDescription

//          |offer-A1| is sent over signaling protocol to Bob
AliceJS->WebServer: signaling with |offer-A1|
WebServer->BobJS: signaling with |offer-A1|

//          |offer-A1| arrives at Bob
BobJS->BobUA: create a PeerConnection
BobJS->BobUA: setRemoteDescription with |offer-A1|
BobUA->BobJS: onaddstream event with remoteStream

//          Bob accepts call
BobJS->BobUA: addTrack with local tracks
BobJS->BobUA: createAnswer
BobJS->BobUA: setLocalDescription with answer
BobUA->BobJS: multiple onicecandidate events with candidates

//          wait for ICE gathering to complete
BobUA->BobJS: onicecandidate event with null candidate
BobJS->BobUA: get |answer-A1| from currentLocalDescription

//          |answer-A1| is sent over signaling protocol to Alice
BobJS->WebServer: signaling with |answer-A1|
WebServer->AliceJS: signaling with |answer-A1|

//          |answer-A1| arrives at Alice
AliceJS->AliceUA: setRemoteDescription with |answer-A1|
AliceUA->AliceJS: onaddstream event with remoteStream

//          media flows
BobUA->AliceUA: media sent from Bob to Alice
AliceUA->BobUA: media sent from Alice to Bob
```

The SDP for |offer-A1| looks like:

```
v=0
o=- 4962303333179871722 1 IN IP4 0.0.0.0
```

```
s=-
t=0 0
a=group:BUNDLE a1 v1
a=ice-options:trickle
m=audio 56500 UDP/TLS/RTP/SAVPF 96 0 8 97 98
c=IN IP4 192.0.2.1
a=mid:a1
a=rtcp:56501 IN IP4 192.0.2.1
a=msid:47017fee-b6c1-4162-929c-a25110252400
      f83006c5-a0ff-4e0a-9ed9-d3e6747be7d9
a=sendrecv
a=rtpmap:96 opus/48000/2
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 telephone-event/8000
a=rtpmap:98 telephone-event/48000
a=maxptime:120
a=ice-ufrag:ETEnlv9DoTMB9J4r
a=ice-pwd:OtSK0WpNtpUjkY4+86js7ZQl
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=rtcp-mux
a=rtcp-rsize
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid
a=candidate:3348148302 1 udp 2113937151 192.0.2.1 56500
      typ host
a=candidate:3348148302 2 udp 2113937151 192.0.2.1 56501
      typ host
a=end-of-candidates

m=video 56502 UDP/TLS/RTP/SAVPF 100 101
c=IN IP4 192.0.2.1
a=rtcp:56503 IN IP4 192.0.2.1
a=mid:v1
a=msid:61317484-2ed4-49d7-9eb7-1414322a7aae
      f30bdb4a-5db8-49b5-bcdc-e0c9a23172e0
a=sendrecv
a=rtpmap:100 VP8/90000
a=rtpmap:101 rtx/90000
a=fmtp:101 apt=100
a=ice-ufrag:BGKkWnG5GmiUpdIV
a=ice-pwd:mqyWsAjvtKwTGnvhPztQ9mIf
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=rtcp-mux
a=rtcp-rsize
a=extmap:3 urn:ietf:params:rtp-hdext:sdes:mid
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 nack pli
a=candidate:3348148302 1 udp 2113937151 192.0.2.1 56502
      typ host
a=candidate:3348148302 2 udp 2113937151 192.0.2.1 56503
      typ host
a=end-of-candidates
```

The SDP for |answer-A1| looks like:

```
v=0
o=- 6729291447651054566 1 IN IP4 0.0.0.0
s=-
t=0 0
a=group:BUNDLE a1 v1
m=audio 20000 UDP/TLS/RTP/SAVPF 96 0 8 97 98
c=IN IP4 192.0.2.2
a=mid:a1
a=rtcp:20000 IN IP4 192.0.2.2
a=msid:PI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1
      PI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1a0
a=sendrecv
a=rtpmap:96 opus/48000/2
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 telephone-event/8000
a=rtpmap:98 telephone-event/48000
a=maxptime:120
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=rtcp-mux
a=rtcp-rsize
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level
a=candidate:2299743422 1 udp 2113937151 192.0.2.2 20000
      typ host
a=end-of-candidates

m=video 20000 UDP/TLS/RTP/SAVPF 100 101
```

```

c=IN IP4 192.0.2.2
a=rtcp 20001 IN IP4 192.0.2.2
a=mid:v1
a=msid:PI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1
      PI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1v0
a=sendrecv
a=rtpmap:100 VP8/90000
a=rtpmap:101 rtx/90000
a=fmtp:101 apt=100
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=rtcp-mux
a=rtcp-rsize
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 nack pli

```

7.2. Normal Examples

This section shows a typical example of a session between two browsers setting up an audio channel and a data channel. Trickle ICE is used in full trickle mode with a bundle policy of max-bundle, an RTCP mux policy of require, and a single TURN server. Later, two video flows, one for the presenter and one for screen sharing, are added to the session. This example shows Alice's browser initiating the session to Bob's browser. The messages from Alice's JS to Bob's JS are assumed to flow over some signaling protocol via a web server.

```

//                               set up local media state
AliceJS->AliceUA:                create new PeerConnection
AliceJS->AliceUA:                addTrack with an audio track
AliceJS->AliceUA:                createDataChannel to get data channel
AliceJS->AliceUA:                createOffer to get |offer-B1|
AliceJS->AliceUA:                setLocalDescription with |offer-B1|

//                               |offer-B1| is sent over signaling protocol to Bob
AliceJS->WebServer:              signaling with |offer-B1|
WebServer->BobJS:                signaling with |offer-B1|

//                               |offer-B1| arrives at Bob
BobJS->BobUA:                    create a PeerConnection
BobJS->BobUA:                    setRemoteDescription with |offer-B1|
BobUA->BobJS:                    onaddstream with audio track from Alice

//                               candidates are sent to Bob

```

```
AliceUA->AliceJS: onicecandidate event with |candidate-B1| (host)
AliceJS->WebServer: signaling with |candidate-B1|
AliceUA->AliceJS: onicecandidate event with |candidate-B2| (srflx)
AliceJS->WebServer: signaling with |candidate-B2|

WebServer->BobJS: signaling with |candidate-B1|
BobJS->BobUA: addIceCandidate with |candidate-B1|
WebServer->BobJS: signaling with |candidate-B2|
BobJS->BobUA: addIceCandidate with |candidate-B2|

// Bob accepts call
BobJS->BobUA: addTrack with local audio
BobJS->BobUA: createDataChannel to get data channel
BobJS->BobUA: createAnswer to get |answer-B1|
BobJS->BobUA: setLocalDescription with |answer-B1|

// |answer-B1| is sent to Alice
BobJS->WebServer: signaling with |answer-B1|
WebServer->AliceJS: signaling with |answer-B1|
AliceJS->AliceUA: setRemoteDescription with |answer-B1|
AliceUA->AliceJS: onaddstream event with audio track from Bob

// candidates are sent to Alice
BobUA->BobJS: onicecandidate event with |candidate-B3| (host)
BobJS->WebServer: signaling with |candidate-B3|
BobUA->BobJS: onicecandidate event with |candidate-B4| (srflx)
BobJS->WebServer: signaling with |candidate-B4|

WebServer->AliceJS: signaling with |candidate-B3|
AliceJS->AliceUA: addIceCandidate with |candidate-B3|
WebServer->AliceJS: signaling with |candidate-B4|
AliceJS->AliceUA: addIceCandidate with |candidate-B4|

// data channel opens
BobUA->BobJS: ondatachannel event
AliceUA->AliceJS: ondatachannel event
BobUA->BobJS: onopen
AliceUA->AliceJS: onopen

// media is flowing between browsers
BobUA->AliceUA: audio+data sent from Bob to Alice
AliceUA->BobUA: audio+data sent from Alice to Bob

// some time later Bob adds two video streams
// note, no candidates exchanged, because of bundle
BobJS->BobUA: addTrack with first video stream
BobJS->BobUA: addTrack with second video stream
BobJS->BobUA: createOffer to get |offer-B2|
```

```
BobJS->BobUA:      setLocalDescription with |offer-B2|

//                |offer-B2| is sent to Alice
BobJS->WebServer:  signaling with |offer-B2|
WebServer->AliceJS: signaling with |offer-B2|
AliceJS->AliceUA:  setRemoteDescription with |offer-B2|
AliceUA->AliceJS:  onaddstream event with first video stream
AliceUA->AliceJS:  onaddstream event with second video stream
AliceJS->AliceUA:  createAnswer to get |answer-B2|
AliceJS->AliceUA:  setLocalDescription with |answer-B2|

//                |answer-B2| is sent over signaling protocol to Bob
AliceJS->WebServer: signaling with |answer-B2|
WebServer->BobJS:   signaling with |answer-B2|
BobJS->BobUA:      setRemoteDescription with |answer-B2|

//                media is flowing between browsers
BobUA->AliceUA:    audio+video+data sent from Bob to Alice
AliceUA->BobUA:    audio+video+data sent from Alice to Bob
```

The SDP for |offer-B1| looks like:

```
v=0
o=- 4962303333179871723 1 IN IP4 0.0.0.0
s=-
t=0 0
a=group:BUNDLE a1 d1
a=ice-options:trickle
m=audio 9 UDP/TLS/RTP/SAVPF 96 0 8 97 98
c=IN IP4 0.0.0.0
a=rtcp:9 IN IP4 0.0.0.0
a=mid:a1
a=msid:57017fee-b6c1-4162-929c-a25110252400
      e83006c5-a0ff-4e0a-9ed9-d3e6747be7d9
a=sendrecv
a=rtpmap:96 opus/48000/2
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 telephone-event/8000
a=rtpmap:98 telephone-event/48000
a=maxptime:120
a=ice-ufrag:ATEnlv9DoTMB9J4r
a=ice-pwd:AtSK0WpNtpUjkY4+86js7ZQ1
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=rtcp-mux
a=rtcp-rsize
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid

m=application 0 UDP/DTLS/SCTP webrtc-datachannel
c=IN IP4 0.0.0.0
a=bundle-only
a=mid:d1
a=fmtp:webrtc-datachannel max-message-size=65536
a=sctp-port 5000
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
```

The SDP for |candidate-B1| looks like:

```
candidate:109270923 1 udp 2122194687 192.168.1.2 51556 typ host
```

The SDP for |candidate-B2| looks like:

```
candidate:4036177503 1 udp 1685987071 11.22.33.44 52546 typ srflx
raddr 192.168.1.2 rport 51556
```

The SDP for |answer-B1| looks like:

```
v=0
o=- 7729291447651054566 1 IN IP4 0.0.0.0
s=-
t=0 0
a=group:BUNDLE a1 d1
a=ice-options:trickle
m=audio 9 UDP/TLS/RTP/SAVPF 96 0 8 97 98
c=IN IP4 0.0.0.0
a=rtcp:9 IN IP4 0.0.0.0
a=mid:a1
a=msid:QI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1
      QI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1a0
a=sendrecv
a=rtpmap:96 opus/48000/2
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 telephone-event/8000
a=rtpmap:98 telephone-event/48000
a=maxptime:120
a=ice-frag:7sFvz2gdLkEwjZEr
a=ice-pwd:DOTZKZNV109RSGsEGM63JXT2
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=rtcp-mux
a=rtcp-rsize
a=extmap:1 urn:ietf:params:rtp-hdrext:ssrc-audio-level
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid
m=application 9 UDP/DTLS/SCTP webrtc-datachannel
c=IN IP4 0.0.0.0
a=mid:d1
a=fmtp:webrtc-datachannel max-message-size=65536
a=sctp-port 5000
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
```

The SDP for |candidate-B3| looks like:

```
candidate:109270924 1 udp 2122194687 192.168.2.3 61665 typ host
```

The SDP for |candidate-B4| looks like:

```
candidate:4036177504 1 udp 1685987071 55.66.77.88 64532 typ srflx
raddr 192.168.2.3 rport 61665
```

The SDP for |offer-B2| looks like: (note the increment of the version number in the o= line, and the c= and a=rtcp lines, which indicate the local candidate that was selected)

```
v=0
o=- 7729291447651054566 2 IN IP4 0.0.0.0
s=-
t=0 0
a=group:BUNDLE a1 d1 v1 v2
a=ice-options:trickle
m=audio 64532 UDP/TLS/RTP/SAVPF 96 0 8 97 98
c=IN IP4 55.66.77.88
a=rtcp:64532 IN IP4 55.66.77.88
a=mid:a1
a=msid:QI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1
      QI39StLS8W7ZbQ11sJsWUXkr3Zf12fJUvzQ1a0
a=sendrecv
a=rtpmap:96 opus/48000/2
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 telephone-event/8000
a=rtpmap:98 telephone-event/48000
a=maxptime:120
a=ice-frag:7sFvz2gdLkEwjZEr
a=ice-pwd:DOTZKZNVlO9RSGsEGM63JXT2
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:actpass
a=rtcp-mux
a=rtcp-rsize
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid
a=candidate:109270924 1 udp 2122194687 192.168.2.3 61665 typ host
a=candidate:4036177504 1 udp 1685987071 55.66.77.88 64532 typ srflx
raddr 192.168.2.3 rport 61665
a=candidate:3671762467 1 udp 41819903 66.77.88.99 50416 typ relay
raddr 55.66.77.88 rport 64532
```

```
a=end-of-candidates

m=application 64532 UDP/DTLS/SCTP webrtc-datachannel
c=IN IP4 55.66.77.88
a=mid:d1
a=fmtp:webrtc-datachannel max-message-size=65536
a=sctp-port 5000
a=ice-ufraq:7sFvz2gdLkEwjZEr
a=ice-pwd:DOTZKZNVlO9RSGsEGM63JXT2
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:actpass
a=candidate:109270924 1 udp 2122194687 192.168.2.3 61665 typ host
a=candidate:4036177504 1 udp 1685987071 55.66.77.88 64532 typ srflx
                        raddr 192.168.2.3 rport 61665
a=candidate:3671762467 1 udp 41819903 66.77.88.99 50416 typ relay
                        raddr 55.66.77.88 rport 64532
a=end-of-candidates

m=video 0 UDP/TLS/RTP/SAVPF 100 101
c=IN IP4 55.66.77.88
a=bundle-only
a=rtcp:64532 IN IP4 55.66.77.88
a=mid:v1
a=msid:61317484-2ed4-49d7-9eb7-1414322a7aae
                        f30bdb4a-5db8-49b5-bcdc-e0c9a23172e0
a=sendrecv
a=rtpmap:100 VP8/90000
a=rtpmap:101 rtx/90000
a=fmtp:101 apt=100
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=rtcp-mux
a=rtcp-rsize
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 nack pli

m=video 0 UDP/TLS/RTP/SAVPF 100 101
c=IN IP4 55.66.77.88
a=bundle-only
a=rtcp:64532 IN IP4 55.66.77.88
a=mid:v1
a=msid:71317484-2ed4-49d7-9eb7-1414322a7aae
                        f30bdb4a-5db8-49b5-bcdc-e0c9a23172e0
```



```
a=sendrecv
a=rtpmap:100 VP8/90000
a=rtpmap:101 rtx/90000
a=fmtp:101 apt=100
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=rtcp-mux
a=rtcp-rsize
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 nack pli
```

The SDP for |answer-B2| looks like: (note the use of setup:passive to maintain the existing DTLS roles, and the use of a=recvonly to indicate that the video streams are one-way)

```
v=0
o=- 4962303333179871723 2 IN IP4 0.0.0.0
s=-
t=0 0
a=group:BUNDLE a1 d1 v1 v2
a=ice-options:trickle
m=audio 52546 UDP/TLS/RTP/SAVPF 96 0 8 97 98
c=IN IP4 11.22.33.44
a=rtcp:52546 IN IP4 11.22.33.44
a=mid:a1
a=msid:57017fee-b6c1-4162-929c-a25110252400
    e83006c5-a0ff-4e0a-9ed9-d3e6747be7d9
a=sendrecv
a=rtpmap:96 opus/48000/2
a=rtpmap:0 PCMU/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 telephone-event/8000
a=rtpmap:98 telephone-event/48000
a=maxptime:120
a=ice-frag:ATEnlv9DoTMB9J4r
a=ice-pwd:AtSK0WpNtpUjkY4+86js7ZQ1
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:passive
a=rtcp-mux
a=rtcp-rsize
```

```
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid
a=candidate:109270923 1 udp 2122194687 192.168.1.2 51556 typ host
a=candidate:4036177503 1 udp 1685987071 11.22.33.44 52546 typ srflx
    raddr 192.168.1.2 rport 51556
a=candidate:3671762466 1 udp 41819903 22.33.44.55 61405 typ relay
    raddr 11.22.33.44 rport 52546
a=end-of-candidates

m=application 52546 UDP/DTLS/SCTP webrtc-datachannel
c=IN IP4 11.22.33.44
a=mid:d1
a=fmtp:webrtc-datachannel max-message-size=65536
a=sctp-port 5000
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:passive

m=video 52546 UDP/TLS/RTP/SAVPF 100 101
c=IN IP4 11.22.33.44
a=rtcp:52546 IN IP4 11.22.33.44
a=mid:v1
a=recvonly
a=rtpmap:100 VP8/90000
a=rtpmap:101 rtx/90000
a=fmtp:101 apt=100
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:passive
a=rtcp-mux
a=rtcp-rsize
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 nack pli

m=video 52546 UDP/TLS/RTP/SAVPF 100 101
c=IN IP4 11.22.33.44
a=rtcp:52546 IN IP4 11.22.33.44
a=mid:v2
a=recvonly
a=rtpmap:100 VP8/90000
a=rtpmap:101 rtx/90000
a=fmtp:101 apt=100
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:passive
a=rtcp-mux
a=rtcp-rsize
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid
a=rtcp-fb:100 ccm fir
a=rtcp-fb:100 nack
a=rtcp-fb:100 nack pli
```

8. Security Considerations

The IETF has published separate documents [I-D.ietf-rtcweb-security-arch] [I-D.ietf-rtcweb-security] describing the security architecture for WebRTC as a whole. The remainder of this section describes security considerations for this document.

While formally the JSEP interface is an API, it is better to think of it is an Internet protocol, with the JS being untrustworthy from the perspective of the browser. Thus, the threat model of [RFC3552] applies. In particular, JS can call the API in any order and with any inputs, including malicious ones. This is particularly relevant when we consider the SDP which is passed to `setLocalDescription()`. While correct API usage requires that the application pass in SDP which was derived from `createOffer()` or `createAnswer()`, there is no guarantee that applications do so. The browser MUST be prepared for the JS to pass in bogus data instead.

Conversely, the application programmer MUST recognize that the JS does not have complete control of browser behavior. One case that bears particular mention is that editing ICE candidates out of the SDP or suppressing trickled candidates does not have the expected behavior: implementations will still perform checks from those candidates even if they are not sent to the other side. Thus, for instance, it is not possible to prevent the remote peer from learning your public IP address by removing server reflexive candidates. Applications which wish to conceal their public IP address should instead configure the ICE agent to use only relay candidates.

9. IANA Considerations

This document requires no actions from IANA.

10. Acknowledgements

Significant text incorporated in the draft as well and review was provided by Peter Thatcher, Taylor Brandstetter, Harald Alvestrand and Suhas Nandakumar. Dan Burnett, Neil Stratford, Anant Narayanan,

Andrew Hutton, Richard Ejzak, Adam Bergkvist and Matthew Kaufman all provided valuable feedback on this proposal.

11. References

11.1. Normative References

[I-D.ietf-avtext-rid]

Roach, A., Nandakumar, S., and P. Thatcher, "RTP Stream Identifier (RID) Source Description (SDS)", draft-ietf-avtext-rid-00 (work in progress), February 2016.

[I-D.ietf-ice-trickle]

Ivov, E., Rescorla, E., Uberti, J., and P. Saint-Andre, "Trickle ICE: Incremental Provisioning of Candidates for the Interactive Connectivity Establishment (ICE) Protocol".

[I-D.ietf-mmusic-4572-update]

Holmberg, C., "Updates to RFC 4572", draft-ietf-mmusic-4572-update-05 (work in progress), June 2016.

[I-D.ietf-mmusic-dtls-sdp]

Holmberg, C. and R. Shpount, "Using the SDP Offer/Answer Mechanism for DTLS", draft-ietf-mmusic-dtls-sdp-14 (work in progress), July 2016.

[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-mux-exclusive]

Holmberg, C., "Indicating Exclusive Support of RTP/RTCP Multiplexing using SDP", draft-ietf-mmusic-mux-exclusive-08 (work in progress), June 2016.

[I-D.ietf-mmusic-rid]

Thatcher, P., Zanaty, M., Nandakumar, S., Burman, B., Roach, A., and B. Campen, "RTP Payload Format Constraints", draft-ietf-mmusic-rid-04 (work in progress), February 2016.

[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-mmusic-sdp-mux-attributes]
Nandakumar, S., "A Framework for SDP Attributes when
Multiplexing", draft-ietf-mmusic-sdp-mux-attributes-01
(work in progress), February 2014.
- [I-D.ietf-mmusic-sdp-simulcast]
Burman, B., Westerlund, M., Nandakumar, S., and M. Zanaty,
"Using Simulcast in SDP and RTP Sessions", draft-ietf-
mmusic-sdp-simulcast-04 (work in progress), February 2016.
- [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-fec]
Uberti, J., "WebRTC Forward Error Correction
Requirements", draft-ietf-rtcweb-fec-00 (work in
progress), February 2015.
- [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.ietf-rtcweb-security]
Rescorla, E., "Security Considerations for WebRTC", draft-
ietf-rtcweb-security-06 (work in progress), January 2014.
- [I-D.ietf-rtcweb-security-arch]
Rescorla, E., "WebRTC Security Architecture", draft-ietf-
rtcweb-security-arch-09 (work in progress), February 2014.
- [I-D.ietf-rtcweb-video]
Roach, A., "WebRTC Video Processing and Codec
Requirements", draft-ietf-rtcweb-video-00 (work in
progress), July 2014.

- [I-D.nandakumar-mmusic-proto-iana-registration]
Nandakumar, S., "IANA registration of SDP 'proto' attribute for transporting RTP Media over TCP under various RTP profiles.", September 2014.
- [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.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.
- [RFC3605] Huitema, C., "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", RFC 3605, October 2003.
- [RFC3890] Westerlund, M., "A Transport Independent Bandwidth Modifier for the Session Description Protocol (SDP)", RFC 3890, DOI 10.17487/RFC3890, September 2004, <<http://www.rfc-editor.org/info/rfc3890>>.
- [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.

- [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.
- [RFC6236] Johansson, I. and K. Jung, "Negotiation of Generic Image Attributes in the Session Description Protocol (SDP)", RFC 6236, May 2011.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.
- [RFC6904] Lennox, J., "Encryption of Header Extensions in the Secure Real-time Transport Protocol (SRTP)", RFC 6904, April 2013.

11.2. Informative References

- [I-D.ietf-rtcweb-ip-handling] Uberti, J. and G. Shieh, "WebRTC IP Address Handling Recommendations", draft-ietf-rtcweb-ip-handling-01 (work in progress), March 2016.
- [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] Andreassen, 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.
- [RFC4733] Schulzrinne, H. and T. Taylor, "RTP Payload for DTMF Digits, Telephony Tones, and Telephony Signals", RFC 4733, DOI 10.17487/RFC4733, December 2006, <<http://www.rfc-editor.org/info/rfc4733>>.
- [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.
- [RFC6464] Lennox, J., Ed., Ivov, E., and E. Marocco, "A Real-time Transport Protocol (RTP) Header Extension for Client-to-Mixer Audio Level Indication", RFC 6464, DOI 10.17487/RFC6464, December 2011, <<http://www.rfc-editor.org/info/rfc6464>>.
- [W3C.WD-webrtc-20140617]
Bergkvist, A., Burnett, D., Narayanan, A., and C. Jennings, "WebRTC 1.0: Real-time Communication Between Browsers", World Wide Web Consortium WD WD-webrtc-20140617, June 2014, <<http://www.w3.org/TR/2011/WD-webrtc-20140617>>.

Appendix A. Appendix A

For the syntax validation performed in Section 5.7, the following list of ABNF definitions is used:

Attribute	Reference
ptime	[RFC4566] Section 9
maxptime	[RFC4566] Section 9
rtpmap	[RFC4566] Section 9
recvonly	[RFC4566] Section 9
sendrecv	[RFC4566] Section 9
sendonly	[RFC4566] Section 9
inactive	[RFC4566] Section 9
framerate	[RFC4566] Section 9
fntp	[RFC4566] Section 9
quality	[RFC4566] Section 9
rtcp	[RFC3605] Section 2.1
setup	[RFC4145] Sections 3, 4, and 5
connection	[RFC4145] Sections 3, 4, and 5
fingerprint	[RFC4572] Section 5
rtcp-fb	[RFC4585] Section 4.2
candidate	[RFC5245] Section 15.1
remote-candidates	[RFC5245] Section 15.2
ice-lite	[RFC5245] Section 15.3
ice-ufrag	[RFC5245] Section 15.4
ice-pwd	[RFC5245] Section 15.4
ice-options	[RFC5245] Section 15.5
extmap	[RFC5285] Section 7
mid	[RFC5888] Section 4 and 5
group	[RFC5888] Section 4 and 5
imageattr	[RFC6236] Section 3.1
extmap (encrypt option)	[RFC6904] Section 4
msid	[I-D.ietf-mmusic-msid] Section 2
rid	[I-D.ietf-mmusic-rid] Section 10
simulcast	[I-D.ietf-mmusic-sdp-simulcast]Section 6.1
dtls-id	[I-D.ietf-mmusic-dtls-sdp]Section 4

Table 1: SDP ABNF References

Appendix B. Change log

Note: This section will be removed by RFC Editor before publication.

Changes in draft-17:

- o Split createOffer and createAnswer sections to clearly indicate attributes which always appear and which only appear when not bundled into another m= section.

- o Add descriptions of RtpTransceiver methods.
- o Describe how to process RTCP feedback attributes.
- o Clarify transceiver directions and their interaction with 3264.
- o Describe setCodecPreferences.
- o Update RTP demux algorithm. Include RTCP.
- o Update requirements for when a=rtcp is included, limiting to cases where it is needed for backward compatibility.
- o Clarify SAR handling.
- o Updated addTrack matching algorithm.
- o Remove a=ssrc requirements.
- o Handle a=setup in reoffers.
- o Discuss how RTX/FEC should be handled.
- o Discuss how telephone-event should be handled.
- o Discuss how CN/DTX should be handled.
- o Add missing references to ABNF table.

Changes in draft-16:

- o Update addIceCandidate to indicate ICE generation and allow per-m= section end-of-candidates.
- o Update fingerprint handling to use draft-ietf-mmusic-4572-update.
- o Update text around SDP processing of RTP header extensions and payload formats.
- o Add sections on simulcast, addTransceiver, and createDataChannel.
- o Clarify text to ensure that the session ID is a positive 63 bit integer.
- o Clarify SDP processing for direction indication.
- o Describe SDP processing for rtcp-mux-only.

- o Specify how SDP session version in o= line.
- o Require that when doing an re-offer, the capabilities of the new session are mostly required to be a subset of the previously negotiated session.
- o Clarified ICE restart interaction with bundle-only.
- o Remove support for changing SDP before calling setLocalDescription.
- o Specify algorithm for demuxing RTP based on MID, PT, and SSRC.
- o Clarify rules for rejecting m= lines when bundle policy is balanced or max-bundle.

Changes in draft-15:

- o Clarify text around codecs offered in subsequent transactions to refer to what's been negotiated.
- o Rewrite LS handling text to indicate edge cases and that we're living with them.
- o Require that answerer reject m= lines when there are no codecs in common.
- o Enforce max-bundle on offer processing.
- o Fix TIAS formula to handle bits vs. kilobits.
- o Describe addTrack algorithm.
- o Clean up references.

Changes in draft-14:

- o Added discussion of RtpTransceivers + RtpSenders + RtpReceivers, and how they interact with createOffer/createAnswer.
- o Removed obsolete OfferToReceiveX options.
- o Explained how addIceCandidate can be used for end-of-candidates.

Changes in draft-13:

- o Clarified which SDP lines can be ignored.

- o Clarified how to handle various received attributes.
- o Revised how attributes should be generated for bundled m= lines.
- o Remove unused references.
- o Remove text advocating use of unilateral PTs.
- o Trigger an ICE restart even if the ICE candidate policy is being made more strict.
- o Remove the 'public' ICE candidate policy.
- o Move open issues/TODOs into GitHub issues.
- o Split local/remote description accessors into current/pending.
- o Clarify a=imageattr handling.
- o Add more detail on VoiceActivityDetection handling.
- o Reference draft-shieh-rtcweb-ip-handling.
- o Make it clear when an ICE restart should occur.
- o Resolve reference TODOs.
- o Remove MSID semantics.
- o ice-options are now at session level.
- o Default RTCP mux policy is now 'require'.

Changes in draft-12:

- o Filled in sections on applying local and remote descriptions.
- o Discussed downscaling and upscaling to fulfill imageattr requirements.
- o Updated what SDP can be modified by the application.
- o Updated to latest datachannel SDP.
- o Allowed multiple fingerprint lines.
- o Switched back to IPv4 for dummy candidates.

- o Added additional clarity on ICE default candidates.

Changes in draft-11:

- o Clarified handling of RTP CNAMEs.
- o Updated what SDP lines should be processed or ignored.
- o Specified how a=imageattr should be used.

Changes in draft-10:

- o TODO

Changes in draft-09:

- o Don't return null for {local,remote}Description after close().
- o Changed TCP/TLS to UDP/DTLS in RTP profile names.
- o Separate out bundle and mux policy.
- o Added specific references to FEC mechanisms.
- o Added canTrickle mechanism.
- o Added section on subsequent answers and, answer options.
- o Added text defining set{Local,Remote}Description behavior.

Changes in draft-08:

- o Added new example section and removed old examples in appendix.
- o Fixed <proto> field handling.
- o Added text describing a=rtcp attribute.
- o Reworked handling of OfferToReceiveAudio and OfferToReceiveVideo per discussion at IETF 90.
- o Reworked trickle ICE handling and its impact on m= and c= lines per discussion at interim.
- o Added max-bundle-and-rtcp-mux policy.
- o Added description of maxptime handling.

- o Updated ICE candidate pool default to 0.
- o Resolved open issues around AppID/receiver-ID.
- o Reworked and expanded how changes to the ICE configuration are handled.
- o Some reference updates.
- o Editorial clarification.

Changes in draft-07:

- o Expanded discussion of VAD and Opus DTX.
- o Added a security considerations section.
- o Rewrote the section on modifying SDP to require implementations to clearly indicate whether any given modification is allowed.
- o Clarified impact of IceRestart on CreateOffer in local-offer state.
- o Guidance on whether attributes should be defined at the media level or the session level.
- o Renamed "default" bundle policy to "balanced".
- o Removed default ICE candidate pool size and clarify how it works.
- o Defined a canonical order for assignment of MSTs to m= lines.
- o Removed discussion of rehydration.
- o Added Eric Rescorla as a draft editor.
- o Cleaned up references.
- o Editorial cleanup

Changes in draft-06:

- o Reworked handling of m= line recycling.
- o Added handling of BUNDLE and bundle-only.
- o Clarified handling of rollback.

- o Added text describing the ICE Candidate Pool and its behavior.
- o Allowed OfferToReceiveX to create multiple recvonly m= sections.

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 St S
Kirkland, WA 98033
USA

Email: justin@uberti.name

Cullen Jennings
Cisco
400 3rd Avenue SW
Calgary, AB T2P 4H2
Canada

Email: fluffy@iii.ca

Eric Rescorla (editor)
Mozilla
331 Evelyn Ave
Mountain View, CA 94041
USA

Email: ekr@rtfm.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 7, 2017

S. Nandakumar
C. Jennings
Cisco
January 3, 2017

SDP for the WebRTC
draft-ietf-rtcweb-sdp-03

Abstract

The Web Real-Time Communication [WebRTC] working group is charged to provide protocol support for direct interactive rich communication using audio, video and data between two peers' web browsers. With in the WebRTC framework, Session Description protocol (SDP) [RFC4566] is used for negotiating session capabilities between the peers. Such a negotiation happens based on the SDP Offer/Answer exchange mechanism described in [RFC3264].

This document provides an informational reference in describing the role of SDP and the Offer/Answer exchange mechanism for the most common WebRTC use-cases.

This SDP examples provided in this document is still a work in progress, but it aims to align closest to the evolving standards work.

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 7, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
2. Terminology	3
3. SDP and the WebRTC	3
4. Offer/Answer and the WebRTC	5
5. WebRTC Session Description Examples	6
5.1. Some Conventions	7
5.2. Basic Examples	9
5.2.1. Audio Only Session	9
5.2.2. Audio/Video Session	12
5.2.3. Data Only Session	17
5.2.4. Audio Call On Hold	19
5.2.5. Audio with DTMF Session	22
5.2.6. One Way Audio/Video Session - Document Camera	26
5.2.7. Audio, Video Session with BUNDLE Support Unknown	30
5.2.8. Audio, Video and Data Session	35
5.2.9. Audio, Video Session with BUNDLE Unsupported	40
5.2.10. Audio, Video BUNDLED, but Data (Not BUNDLED)	45
5.2.11. Audio Only, Add Video to BUNDLE	50
5.3. MultiResolution, RTX, FEC Examples	57
5.3.1. Sendonly Simulcast Session with 2 cameras and 2 encodings per camera	57
5.3.2. Successful SVC Video Session	64
5.3.3. Successful Simulcast Video Session with Retransmission	69
5.3.4. Successful 1-way Simulcast Session with 2 resolutions and RTX - One resolution rejected	74
5.3.5. Simulcast Video Session with Forward Error Correction	79
5.4. Others	84
5.4.1. Audio Session - Voice Activity Detection	84
5.4.2. Audio Conference - Voice Activity Detection	87
5.4.3. Successful legacy Interop Fallback with bundle-only	90

5.4.4. Legacy Interop with RTP/AVP profile	95
6. IANA Considerations	100
7. Acknowledgments	100
8. Change Log	100
9. Informative References	102
Authors' Addresses	107

1. Introduction

Javascript Session Exchange Protocol(JSEP) [I-D.ietf-rtcweb-jsep] specifies a generic protocol needed to generate [RFC3264] Offers and Answers negotiated between the WebRTC peers for setting up, updating and tearing down a WebRTC session. For this purpose, SDP is used to construct [RFC3264] Offers/Answers for describing (media and non-media) streams as appropriate for the recipients of the session description to participate in the session.

The remainder of this document is organized as follows: Sections 3 and 4 provides an overview of SDP and the Offer/Answer exchange mechanism. Section 5 provides sample SDP generated for the most common WebRTC use-cases.

2. Terminology

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

3. SDP and the WebRTC


The purpose of this section is to provide a general overview of SDP and its components. For a more in-depth understanding, the readers are advised to refer to [RFC4566].

The Session Description Protocol (SDP) [RFC4566] describes multimedia sessions, which can contain audio, video, whiteboard, fax, modem, and other streams. SDP provides a general purpose, standard representation to describe various aspects of multimedia session such as media capabilities, transport addresses and related metadata in a transport agnostic manner, for the purposes of session announcement, session invitation and parameter negotiation.

As of today SDP is widely used in the context of Session Initiation Protocol [RFC3261], Real-time Transport Protocol [RFC3550] and Real-time Streaming Protocol applications [RFC2326].

Below figure introduces high-level breakup of SDP into components that semantically describe a multimedia session, in our case, a

WebRTC session [WebRTC]. It by no means captures everything about SDP and hence, should be used for informational purposes only.



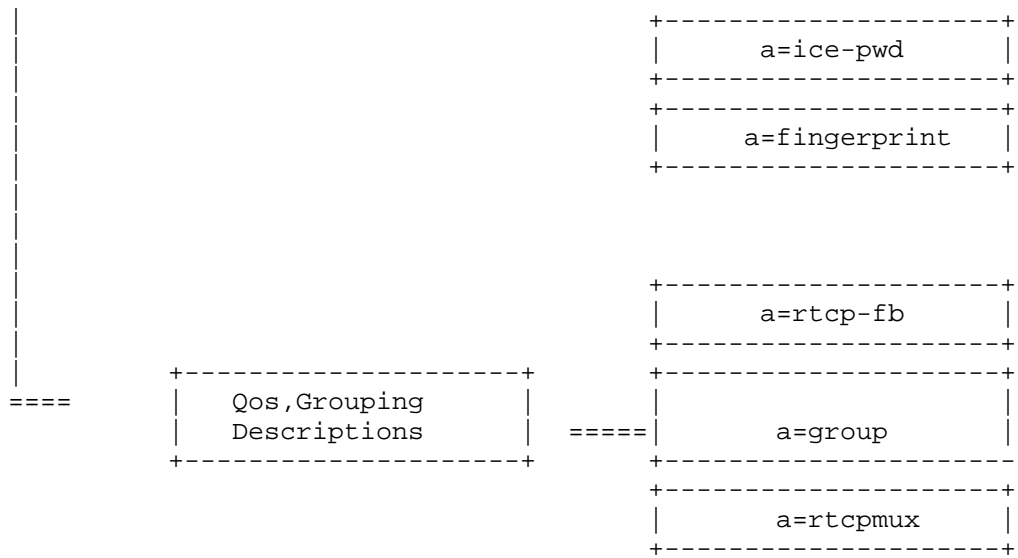


Figure 1: Semantic Components of SDP

[WebRTC] proposes JavaScript application to fully specify and control the signaling plane of a multimedia session as described in the JSEP specification [I-D.ietf-rtcweb-jsep]. JSEP provides mechanisms to create session characterization and media definition information to conduct the session based on SDP exchanges.

In this context, SDP serves two purposes:

1. Provide grammatical structure syntactically.
 2. Semantically convey participant's intention and capabilities required to successfully negotiate a session.
4. Offer/Answer and the WebRTC

This section introduces SDP Offer/Answer Exchange mechanism mandated by WebRTC for negotiating session capabilities while setting up, updating and tearing down a WebRTC session. This section is intentionally brief in nature and interested readers are recommended to refer [RFC3264] for specific details on the protocol operation.

The Offer/Answer [RFC3264] model specifies rule for the bilateral exchange of Session Description Protocol (SDP) messages for creation of multimedia streams. It defines protocol with involved participants exchanging desired session characteristics from each others perspective constructed as SDP to negotiate the session between them.

In the most basic form, the protocol operation begins by one of the participants sending an initial SDP Offer describing its intent to start a multimedia communication session. The participant receiving the offer MAY generate an SDP Answer accepting the offer or it MAY reject the offer. If the session is accepted the Offer/Answer model guarantees a common view of the multimedia session between the participants.

At any time, either participant MAY generate a new SDP offer that updates the session in progress.

With in the context of WebRTC, the Offer/Answer model defines the state-machinery for WebRTC peers to negotiate session descriptions between them during the initial setup stages as well as for eventual session updates. Javascript Session Establishment Protocol specification [I-D.ietf-rtcweb-jsep] for WebRTC provides the mechanism for generating [RFC3264] SDP Offers and Answers in order for both sides of the session to agree upon details such as list of media formats to be sent/received, bandwidth information, crypto parameters, transport parameters, for example.

5. WebRTC Session Description Examples

A typical web based real-time multimedia communication session can be characterized as below:

- o It has zero or more Audio only, Video only or Audio/Video RTP Sessions,
- o MAY contain zero or more non-media data sessions,
- o All the sessions are secured with DTLS-SRTP,
- o Supports NAT traversal using ICE mechanism,
- o Provides RTCP based feedback mechanisms,
- o Sessions can be over IPv4-only, IPv6-only, dual-stack based clients.

5.1. Some Conventions

The examples given in this document follow the conventions listed below:

- o In all the examples, Alice and Bob are assumed to be the WebRTC peers.
- o It is assumed that for most of the examples, the support for [I-D.ietf-mmusic-sdp-bundle-negotiation] is established a priori either out-of-band or as a consequence of successful Offer/Answer negotiation between Alice and Bob, unless explicitly stated otherwise.
- o Call-flow diagrams that accompany the use-cases capture only the prominent aspects of the system behavior and intentionally is not detailed to improve readability.
- o Eventhough the call-flow diagrams shows SDP being exchanged between the parties, it doesn't represent the only way an WebRTC setup is expected to work. Other approaches may involve WebRTC applications to exchange the media setup information via non-SDP mechanisms as long as they confirm to the [I-D.ietf-rtcweb-jsep] API specification.
- o The SDP examples deviate from actual on-the-wire SDP notation in several ways. This is done to facilitate readability and to conform to the restrictions imposed by the RFC formatting rules.
 - * Visual markers/Empty lines in any SDP example are inserted to make functional divisions in the SDP clearer, and are not actually part of the SDP syntax.
 - * Any SDP line that is indented (compared to the initial line in the SDP block) is a continuation of the preceding line. The line break and indent are to be interpreted as a single space character.
 - * Excepting the above two conventions, line endings are to be interpreted as <CR><LF> pairs (that is, an ASCII 13 followed by an ASCII 10).
- o Against each SDP line, pointers to the appropriate RFCs are provided for further informational reference. Also an attempt has been made to provide explanatory notes to enable better understanding of the SDP usage, wherever appropriate.

- o Following SDP details are common across all the use-cases defined in this document unless mentioned otherwise.
 - * DTLS fingerprint for SRTP (a=fingerprint)
 - * RTP/RTCP Multiplexing (a=rtcp-mux)
 - * RTCP Feedback support (a=rtcp-fb)
 - * Host and server-reflexive candidate lines (a=candidate)
 - * SRTP Setup framework parameters (a=setup)
 - * RTCP attribute (a=rtcp)
 - * RTP header extension indicating audio-levels from client to the mixer

For specific details, readers must refer to [I-D.ietf-rtcweb-jsep] specification.

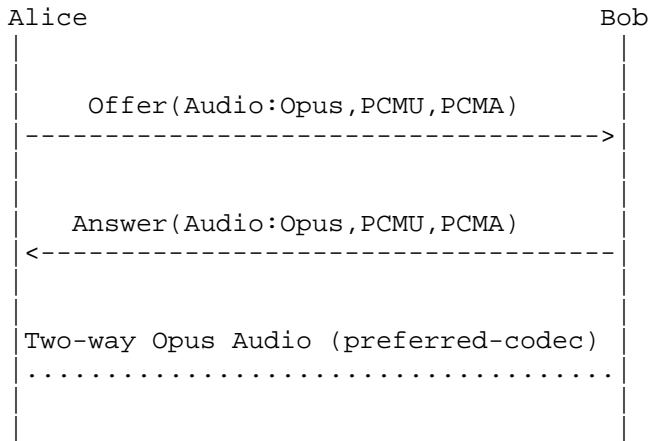
- o The term "Session" is used rather loosely in this document to refer to either a "Communication Session" or a "RTP Session" or a "RTP Stream" depending on the context.
- o Payload type 109 is usually used for OPUS, 0 for PCMU, 8 for PCMA, 99 for H.264 and 120 for VP8 in most of the examples to maintain uniformity.
- o The IP Address:Port combinations '192.168.1.4:61665' (host) and '24.23.204.141:54609' (Server Reflexive) is typically used for Alice.
- o The IP Address:Port combinations '192.168.1.7:51556' (host) and '98.248.92.77:49203' (Server Reflexive) is typically used for Bob.
- o SSRC cname 'cname:EocUG1f0fcg/yvY7' applies to Alice and the cname 'cname:Q/NWslaolHmN4Xa5' corresponds to Bob.
- o The SSRCs values '12345', '56789' and '11111' till '66666' usually represent Alice's RTP Streams and the values '54321', '98765' and '77777' till '99999' represent Bob's RTP Streams.
- o In the actual use the values that represent SSRCs, ICE candidate foundations, WebRTC Mediastream and MediaStreamTrack Ids shall be much larger and/or random than the ones shown in the examples.

5.2. Basic Examples

5.2.1. Audio Only Session

This common scenario shows SDP for secure two-way audio session with Alice offering Opus, PCMU, PCMA and Bob accepting all the offered audio codecs.

2-Way Audio Only Session



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109 0 8	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:60065 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Alice can send

a=rtpmap:109 opus/48000/2	and recv audio
a=rtpmap:0 PCMU/8000	[I-D.ietf-payload-rtp-opus] -
a=rtpmap:8 PCMA/8000	Opus Codec 48khz, 2 channels
a=maxptime:120	[RFC3551] PCMU Audio Codec
a=ice-ufrag:074c6550	[RFC3551] PCMA Audio Codec
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC4566]
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	[RFC5245] - ICE user fragment
a=setup:actpass	[RFC5245] - ICE password
a=rtcp-mux	[RFC5245] - DTLS Fingerprint for SRTP
a=rtcp-rsize	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-fb:109 nack	[RFC5761] - Alice can perform RTP/RTCP Muxing
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC5506] - Alice intends to use reduced size RTCP for this session
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[RFC5104] - Indicates NACK RTCP feedback support
a:ssrc:12345	[RFC6464] Alice supports RTP header extension to indicate audio levels
cname:EocUG1f0fcg/yvY7	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=candidate:0 1 UDP 2122194687 192.168.1.4 61665 typ host	[RFC5576] - Alice's audio stream SSRC.
a=candidate:1 1 UDP 1685987071 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245] - RTP Host Candidate
a=candidate:0 2 UDP 2122194687 192.168.1.4 61667 typ host	[RFC5245] - RTP Server Reflexive ICE Candidate
a=candidate:1 2 UDP 1685987071 24.23.204.141 60065 typ srflx raddr 192.168.1.4 rport 61667	[RFC5245] - RTCP Host Candidate
a=end-of-candidates	[RFC5245] - RTCP Server Reflexive ICE Candidate
	[I-D.ietf-mmusic-trickle-ice]

Table 1: 5.2.1 SDP Offer

Answer SDP Contents	RFC#/Notes
---------------------	------------

v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109 0 8	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Bob can send and recv audio
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] Opus Codec
a=rtpmap:0 PCMU/8000	[RFC3551] PCMU Audio Codec
a=rtpmap:8 PCMA/8000	[RFC3551] PCMA Audio Codec
a=maxptime:120	[RFC4566]
a=ice-ufrag:05067423	[RFC5245] - ICE user fragment
a=ice-pwd:1747dlee3474a28a397a4c3f3af08a068	[RFC5245] - ICE password parameter
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	[RFC5245] - DTLS Fingerprint for SRTP
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761] - Bob can perform RTP/RTCP Muxing on port 49203
a=rtcp-rsize	[RFC5506] - Bob intends to use reduced size RTCP for this session
a=rtcp-fb:109 nack	[RFC5104] - Indicates NACK RTCP feedback support
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464] Bob supports audio level RTP header extension as well
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:54321	[RFC5576] - Bob's audio stream SSRC.
cname:Q/NWslaolHmN4Xa5	
a=candidate:0 1 UDP 2122194687	[RFC5245] - RTP/RTCP Host ICE


192.168.1.7 51556 typ host	Candidate
a=candidate:1 1 UDP 1685987071	[RFC5245] - RTP/RTCP Server
98.248.92.77 49203 typ srflx	Reflexive ICE Candidate
raddr 192.168.1.7 rport 51556	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 2: 5.2.1 SDP Answer

5.2.2. Audio/Video Session

Alice and Bob establish a two-way audio and video session with Opus as the audio codec and H.264 as the video codec.

2-Way Audio, Video Session



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109 0 8	[RFC4566]

c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Alice can send and recv audio
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] - Opus Codec 48khz, 2 channels
a=rtpmap:0 PCMU/8000	[RFC3551] PCMU Audio Codec
a=rtpmap:8 PCMA/8000	[RFC3551] PCMA Audio Codec
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245] - ICE user fragment
a=ice-pwd:a28a397a4c3f31747d1ee3474af08a068	[RFC5245] - ICE password parameter
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	[RFC5245] - DTLS Fingerprint for SRTP
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761] - Alice can perform RTP/RTCP Muxing
a=rtcp-rsize	[RFC5506] - Alice intends to use reduced size RTCP for this session
a=rtcp-fb:109 nack	[RFC5104] - Indicates NACK RTCP feedback support
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:12345	[RFC5576]
cname:EocUGlf0fcg/yvY7	
a=candidate:0 1 UDP 2122194687 192.168.1.4 61665 typ host	[RFC5245] - RTP/RTCP Host Candidate
a=candidate:1 1 UDP 1685987071 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245] - RTP/RTCP Server Reflexive ICE Candidate
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 54609 UDP/TLS/RTP/SAVPF 99 120	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID

a=sendrecv	(ma) and RTCMediaStreamTrack ID (tb) [RFC3264] - Alice can send and recv video
a=rtpmap:99 H264/90000	[RFC3984] - H.264 Video Codec
a=fmtp:99 profile-level-id=4d0028	[RFC3984]
;packetization-mode=1	
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8] - VP8 video codec
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	[RFC5245] - DTLS Fingerprint for SRTP
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761] - Alice can perform RTP/RTCP Muxing
a=rtcp-rsize	[RFC5506] - Alice intends to use reduced size RTCP for this session
a=rtcp-fb:99 nack	[RFC5104] - Indicates NACK RTCP feedback support
a=rtcp-fb:99 nack pli	[RFC5104] - Indicates support for Picture loss Indication and NACK
a=rtcp-fb:99 ccm fir	[RFC5104] - Full Intra Frame Request-Codec Control Message support
a=rtcp-fb:120 nack	[RFC5104] - Indicates NACK RTCP feedback support
a=rtcp-fb:120 nack pli	[RFC5104] - Indicates support for Picture loss Indication and NACK
a=rtcp-fb:120 ccm fir	[RFC5104] - Full Intra Frame Request-Codec Control Message support
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	

Table 3: 5.2.2 SDP Offer

Answer SDP Contents	RFC#/Notes
---------------------	------------

v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ICE]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Bob can send and recv audio
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] - Bob accepts only Opus Codec
a=maxptime:120	[RFC4566]
a=ice-ufraq:c300d85b	[RFC5245] - ICE username frag
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245] - ICE password
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	[RFC5245] - DTLS Fingerprint for SRTP
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761] - Bob can perform RTP/RTCP Muxing
a=rtcp-rsize	[RFC5506] - Bob intends to use reduced size RTCP for this session
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:54321	[RFC5576]
cname:Q/NWslao1HmN4Xa5	
a=candidate:0 1 UDP 3618095783 192.168.1.7 49203 typ host	[RFC5245] - RTP/RTCP Host ICE Candidate
a=candidate:1 1 UDP 565689203 98.248.92.77 49203 typ srflx	[RFC5245] - RTP/RTCP Server Reflexive ICE Candidate
raddr 192.168.1.7 rport 51556	

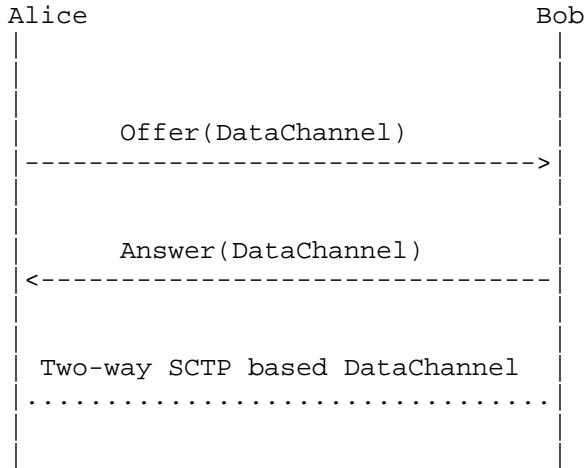
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF	[RFC4566]
99	
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendrecv	[RFC3264] - Bob can send and recv video
a=rtpmap:99 H264/90000	[RFC3984] - Bob accepts H.264 Video Codec.
a=fmtp:99 profile-level-id=4d0028	[RFC3984]
;packetization-mode=1	
a=fingerprint:sha-256 6B:8B:F0:65	[RFC5245] - DTLS Fingerprint for SRTP
:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B	
:35 :DC:B8:5F:64:1A:24:C2:43:F0:A	
1:58:D0:A1:2C:19:08	
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761] - Bob can perform RTP/RTCP Muxing
a=rtcp-rsize	[RFC5506] - Bob intends to use reduced size RTCP for this session
a=rtcp-fb:99 nack	[RFC5104] - Indicates support for NACK based RTCP feedback
a=rtcp-fb:99 nack pli	[RFC5104] - Indicates support for Picture loss Indication and NACK
a=rtcp-fb:99 ccm fir	[RFC5104] - Full Intra Frame Request- Codec Control Message support
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:98765	[RFC5576]
cname:Q/NWslaolHmN4Xa5	

Table 4: 5.2.2 SDP Answer

5.2.3. Data Only Session

This scenario illustrates SDP negotiated to setup a data-only session based on SCTP Data Channel, thus enabling use-cases such as file-transfer for example.

2-Way DataChannel Session



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE data	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Application m=line	*****

m=application 54609 UDP/DTLS/SCTP webrtc-datachannel	[I-D.ietf-rtcweb-data-channel]
c=IN IP4 24.23.204.141	[RFC4566]
a=mid:data	[RFC5888]
a=sendrecv	[RFC3264] - Alice can send and recv non-media data
a=sctp-port:5000	[I-D.ietf-mmusic-sctp-sdp]
a=max-message-size:100000	[I-D.ietf-mmusic-sctp-sdp]
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=connection:new	[RFC4145]
a=ice-ufrag:074c6550	[RFC5245] - Session Level ICE parameter
a=ice-pwd:a28a397a4c3f31747d1ee3474af08a068	[RFC5245] - Session Level ICE parameter
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	[RFC5245] - Session DTLS Fingerprint for SRTP
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 5: 5.2.3 SDP Offer


Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE data	[I-D.ietf-mmusic-sdp-bundle-negotiation]
***** Application m=line *****	*****
m=application 49203 UDP/DTLS/SCTP webrtc-datachannel	[I-D.ietf-mmusic-sctp-sdp]
c=IN IP4 98.248.92.771	[RFC4566]
a=mid:data	[RFC5888]
a=sendrecv	[RFC3264] - Bob can send and recv non-media data
a=sctp-port:5000	[I-D.ietf-mmusic-sctp-sdp]
a=max-message-size:100000	[I-D.ietf-mmusic-sctp-sdp]
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=ice-ufrag:c300d85b	[RFC5245] - Session Level ICE username frag
a=ice-pwd:de4e99bd291c325921d5d47 efbabd9a2	[RFC5245] - Session Level ICE password
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:A 1:58:D0:A1:2C:19:08	[RFC5245] - Session DTLS Fingerprint for SRTP
a=candidate:0 1 UDP 2113667327 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 6: 5.2.3 SDP Answer

5.2.4. Audio Call On Hold

Alice calls Bob, but when Bob answers he places Alice on hold by setting the SDP direction attribute to a=inactive in the Answer.

Audio On Hold



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Alice can send and recv audio
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] - Opus Codec 48khz, 2 channels
a=maxptime:120	[RFC4566]
a=ice-ufrag:074c6550	[RFC5245] - ICE user fragment
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245] - ICE password
a=fingerprint:sha-256 19:E2:1C:3B	[RFC5245] - DTLS Fingerprint

:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73	for SRTP
:04 :BB:05:2F:70:9F:04:A9:0E:05:E	
9:26:33:E8:70:88:A2	
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761] - Alice can perform RTP/RTCP Muxing
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104] - Indicates NACK RTCP feedback support
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 1685987071 24.23.204.141 54609 typ srflx	[RFC5245]
raddr 192.168.1.4 rport 61665	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 7: 5.2.4 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=inactive	[RFC3264] - Bob puts call On Hold


a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] - Bob accepts Opus Codec
a=maxptime:120	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245] - ICE username frag
a=ice-pwd:de4e99bd291c325921d5d47 efbabd9a2	[RFC5245] - ICE password
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:A 1:58:D0:A1:2C:19:08	[RFC5245] - DTLS Fingerprint for SRTP
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761] - Bob can perform RTP/RTCP Muxing
a=rtcp-rsize	[RFC5506]
a=extmap:1 urn:ietf:params:rtp- hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp- hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a:ssrc:54321	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
a=candidate:0 1 UDP 2113667327 192.168.1.7 51556 typ host	[RFC5245] - Host candidate
a=candidate:1 1 UDP 1685987071 24.23.204.141 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245] - Server Reflexive candidate
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 8: 5.2.4 SDP Answer

5.2.5. Audio with DTMF Session

In this example, Alice wishes to establish two separate audio streams, one for normal audio and the other for telephone-events. Alice offers first audio stream with three codecs and the other with [RFC2833] tones (for DTMF). Bob accepts both the audio streams by choosing Opus as the audio codec and telephone-event for the other stream.

Audio Session with DTMF



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio dtmf	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109 0 8	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Alice can send and recv audio
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] - Opus Codec 48khz, 2 channels
a=rtpmap:0 PCMU/8000	[RFC3551] PCMU Audio Codec
a=rtpmap:8 PCMA/8000	[RFC3551] PCMA Audio Codec
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245] - ICE user fragment
a=ice-pwd:a28a397a4c3f31747dlee34	[RFC5245] - ICE password

<pre> 74af08a068 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:E 9:26:33:E8:70:88:A2 a=setup:actpass a=rtcp-mux a=rtcp-rsize a=rtcp-fb:109 nack a=extmap:1 urn:ietf:params:rtp- hdnext:ssrc-audio-level a=extmap:2 urn:ietf:params:rtp- hdnext:sdes:mid a:ssrc:12345 cname:EocUG1f0fcg/yvY7 a=candidate:0 1 UDP 2122194687 192.168.1.4 61665 typ host a=candidate:1 1 UDP 1685987071 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665 a=end-of-candidates ***** DTMF m=line ***** m=audio 54609 UDP/TLS/RTP/SAVPF 126 c=IN IP4 24.23.204.141 a=rtcp:54609 IN IP4 24.23.204.141 a=mid:dtmf a=msid:ma tb a=sendonly a=rtpmap:126 telephone-event/8000 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:E 9:26:33:E8:70:88:A2 a=setup:actpass a=rtcp-mux a=rtcp-rsize a=rtcp-fb:109 nack a=extmap:2 urn:ietf:params:rtp- hdnext:sdes:mid </pre>	<pre> parameter [RFC5245] - DTLS Fingerprint for SRTP [RFC4145] - Alice can perform DTLS before Answer arrives [RFC5761] - Alice can perform RTP/RTCP Muxing [RFC5506] [RFC5104] - Indicates NACK RTCP feedback support [RFC6464] [I-D.ietf-mmusic-sdp-bundle-n egotiation] [RFC5576] [RFC5245] [RFC5245] [I-D.ietf-mmusic-trickle-ice] ***** [RFC4566] [RFC4566] [RFC3605] [RFC5888] Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb) [RFC3264] - Alice can send DTMF Events [RFC2833] [RFC5245] - DTLS Fingerprint for SRTP [RFC4145] - Alice can perform DTLS before Answer arrives [RFC5761] [RFC5506] [RFC5104] - Indicates NACK RTCP feedback support [I-D.ietf-mmusic-sdp-bundle-n egotiation] </pre>
---	--

a=ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	

Table 9: 5.2.5 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio dtmf	[I-D.ietf-mmusic-sdp-bundle-negotiation]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Bob can send and receive Opus audio
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] - Bob accepts Opus Codec
a=maxptime:120	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245] - ICE username frag
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245] - ICE password
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	[RFC5245] - Fingerprint for SRTP
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761] - Bob can perform RTP/RTCP Muxing on port 49203
a=rtcp-rsize	[RFC5506] - Alice intends to use reduced size RTCP for this session
a=extmap:1 urn:ietf:params:rtp-hdnext:ssrc-audio-level	[RFC6464]

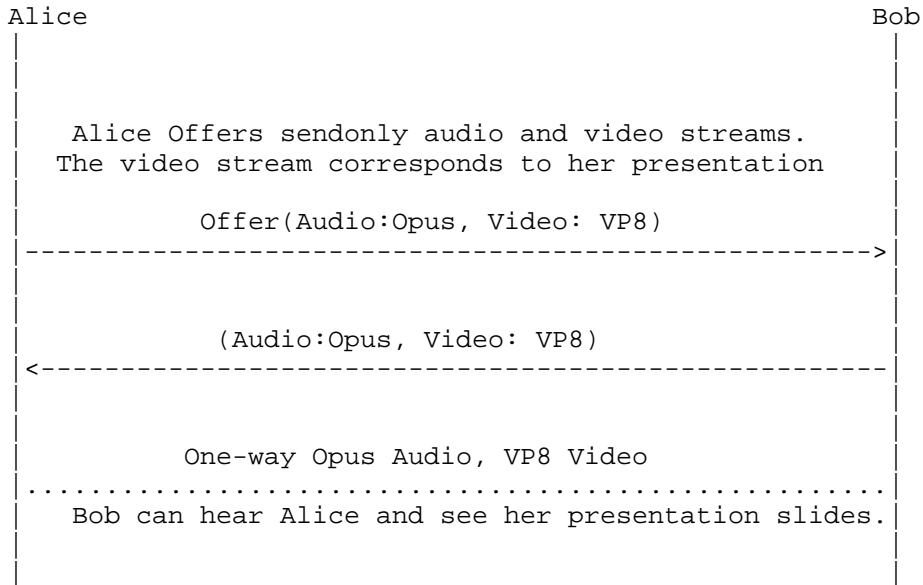
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:54321	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
a=candidate:0 1 UDP 2122194687 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1685987071 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** DTMF m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 126	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 24.23.204.141	[RFC3605]
a=mid:dtmf	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=recvonly	[RFC3264] - Alice can receive DTMF events
a=rtpmap:126 telephone-event/8000	[RFC2833]
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	[RFC5245] - Fingerprint for SRTP
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506] - Bob intends to use reduced size RTCP for this session
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:98765	[RFC5576]
cname:Q/NWslaolHmN4Xa5	

Table 10: 5.2.5 SDP Answer

5.2.6. One Way Audio/Video Session - Document Camera

In this scenario Alice and Bob engage in a 1 way audio and video session with Bob receiving Alice's audio and her presentation slides as video stream.

One Way Audio & Video Session - Document Camera



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendonly	[RFC3264] - Send only audio stream
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]

a=ice-ufrag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2122194687 24.23.204.141 54609 typ host	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 54609 UDP/TLS/RTP/SAVPF 120	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendonly	[RFC3264] - Send only video stream
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=content:slides	[RFC4796] -Alice's presentation video stream
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	[RFC5245]
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]

a=ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	

Table 11: 5.2.6 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=recvonly	[RFC3264] - Receive only audio stream
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdrext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:54321	[RFC5576]
cname:Q/NWslaolHmN4Xa5	

a=candidate:0 1 UDP 2113667327 98.248.92.77 49203 typ host	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF 120	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=recvonly	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=content:slides	[RFC4796]presentation stream
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:A 1:58:D0:A1:2C:19:08	[RFC5245]
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp- hdnext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a=ssrc:98765	[RFC5576]
cname:Q/NWslao1HmN4Xa5	

Table 12: 5.2.6 SDP Answer

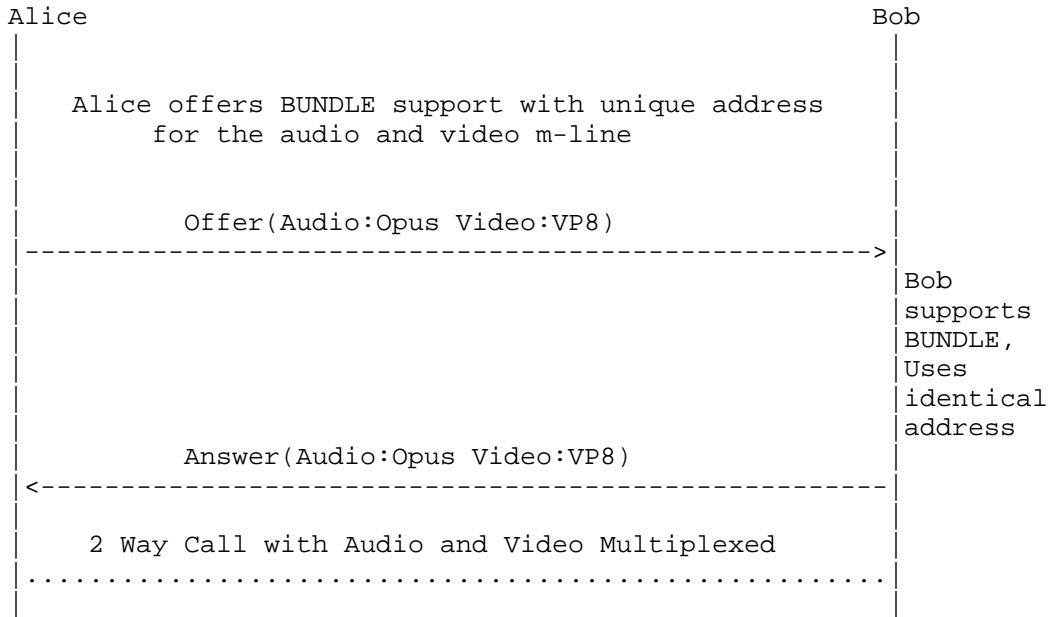
5.2.7. Audio, Video Session with BUNDLE Support Unknown

In this example, since Alice is unsure of the Bob's support of the BUNDLE framework, following steps are performed in order to negotiate and setup a BUNDLE Address for the session

- o An SDP Offer, in which the Alice assigns unique addresses to each "m=" line in the BUNDLE group, and requests the Answerer to select the Offerer's BUNDLE address.
- o An SDP Answer, in which the Bob indicates its support for BUNDLE, selects the offerer's BUNDLE address, selects its own BUNDLE address and associates it with each BUNDLED m=line within the BUNDLE group.

Once the Offer/Answer exchange completes, both Alice and Bob each end up using single RTP Session for both the Media Streams.

Two-Way Secure Audio,Video with BUNDLE support unknown



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION] Alice supports grouping of m-lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ICE]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54610 IN IP4 24.23.204.141	[RFC3605] - RTCP port different from RTP Port

a=mid:audio	[RFC5888] Audio m=line part of BUNDLE group with a unique port number
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-ufrag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2122194687 192.168.1.4 61665 typ host	[RFC5245] - RTP host candidate
a=candidate:1 1 UDP 1685987071 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245] - RTP Server Reflexive candidate
a=candidate:0 2 UDP 2122194687 192.168.1.4 61666 typ host	[RFC5245] - RTCP host candidate
a=candidate:1 2 UDP 1685987071 24.23.204.141 54610 typ srflx raddr 192.168.1.4 rport 61666	[RFC5245] - RTCP Server Reflexive candidate
***** Video m=line *****	*****
m=video 62537 UDP/TLS/RTP/SAVPF 120	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:62538 IN IP4 24.23.204.141	[RFC3605]
a=mid:video	[RFC5888] Video m=line part of the Bundle group with a unique port number
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)

a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=ice-ufrag:6550074c	[RFC5245]
a=ice-pwd:74af08a068a28a397a4c3f31747dlee34	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2122194687 192.168.1.4 61886 typ host	[RFC5245] - RTP Host candidate
a=candidate:1 1 UDP 1685987071 24.23.204.141 62537 typ srflx raddr 192.168.1.4 rport 61886	[RFC5245] - RTP Server Reflexive candidate
a=candidate:0 2 2122194687 192.168.1.4 61888 typ host	[RFC5245] - RTCP host candidate
a=candidate:1 2 UDP 1685987071 24.23.204.141 62538 typ srflx raddr 192.168.1.4 rport 61888	[RFC5245] - RTCP Server Reflexive candidate

Table 13: 5.2.7 SDP Offer w/BUNDLE

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION] Bob supports BUNDLE semantics.
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ICE]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF	[RFC4566]

109	
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:audio	[RFC5888] Audio m=line part of the BUNDLE group
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-frag:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47 efbabb9a2	[RFC5245]
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:A 1:58:D0:A1:2C:19:08	[RFC5245]
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp- hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp- hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a:ssrc:54321	[RFC5576]
cname:Q/NWslao1HmN4Xa5	
a=candidate:0 1 UDP 2122194687 192.168.1.7 49203 typ host	[RFC5245]
a=candidate:1 1 UDP 1685987071 98.248.92.77 51556 typ srflx raddr 192.168.1.7 rport 49203	[RFC5245]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF 120	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:video	[RFC5888] Video m=line part of the BUNDLE group with the port from audio line repeated
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=fingerprint:sha-256 6B:8B:F0:65	[RFC5245]

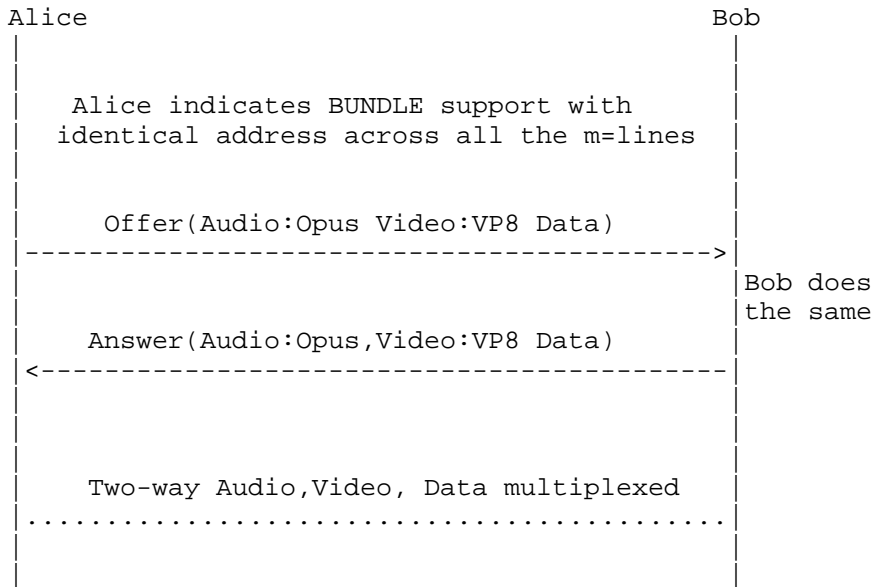
:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B	
:35 :DC:B8:5F:64:1A:24:C2:43:F0:A	
1:58:D0:A1:2C:19:08	
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdext:sdes:mid	egotiation]
a=ssrc:98765	[RFC5576]
cname:Q/NWslao1HmN4Xa5	

Table 14: 5.2.7 SDP Answer w/BUNDLE

5.2.8. Audio, Video and Data Session

This example shows SDP for negotiating a session with Audio, Video and data streams between Alice and Bob with BUNDLE support known.

Audio,Video,Data with BUNDLE support known



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video data	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=mid:audio	[RFC5888]
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747d1ee34	[RFC5245]

74af08a068	
a=fingerprint:sha-256 19:E2:1C:3B	[RFC5245]
:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73	
:04 :BB:05:2F:70:9F:04:A9:0E:05:E	
9:26:33:E8:70:88:A2	
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-	[RFC6464]
hdext:ssrc-audio-level	
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdext:sdes:mid	egotiation]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2122194687	[RFC5245]
192.168.1.4 61665 typ host	
a=candidate:1 1 UDP 1685987071	[RFC5245]
24.23.204.141 54609 typ srflx	
raddr 192.168.1.4 rport 61665	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 54609 UDP/TLS/RTP/SAVPF	[RFC4566]
120	
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID
	(ma) and RTCMediaStreamTrack
	ID (tb)
a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=fingerprint:sha-256 19:E2:1C:3B	[RFC5245]
:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73	
:04 :BB:05:2F:70:9F:04:A9:0E:05:E	
9:26:33:E8:70:88:A2	
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdext:sdes:mid	egotiation]
a:ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	
***** Application m=line	*****

m=application 54609 UDP/DTLS/SCTP webrtc-datachannel c=IN IP4 24.23.204.141 a=mid:data a=sctp-port:5000 a=max-message-size:100000 a=sendrecv a=setup:actpass a=connection:new a=ice-ufrag:074c6550 a=ice-pwd:a28a397a4c3f31747dlee34 74af08a068 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:E 9:26:33:E8:70:88:A2	[I-D.ietf-rtcweb-data-channel] [RFC4566] [RFC5888] [I-D.ietf-mmusic-sctp-sdp] [I-D.ietf-mmusic-sctp-sdp] [RFC3264] [RFC4145] [RFC4145] [RFC5245] [RFC5245] [RFC5245]
--	--

Table 15: 5.2.8 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0 o=- 16833 0 IN IP4 0.0.0.0 s=- t=0 0 a=group:BUNDLE audio video data a=ice-options:trickle ***** Audio m=line ***** m=audio 49203 UDP/TLS/RTP/SAVPF 109 c=IN IP4 98.248.92.77 a=rtcp:49203 IN IP4 98.248.92.77 a=msid:ma ta a=mid:audio a=sendrecv a=rtpmap:109 opus/48000/2 a=maxptime:120 a=ice-ufrag:c300d85b a=ice-pwd:de4e99bd291c325921d5d47 efbabd9a2 a=fingerprint:sha-256 6B:8B:F0:65	[RFC4566] [RFC4566] - Session Origin Information [RFC4566] [RFC4566] [I-D.ietf-mmusic-sdp-bundle-n egotiation] [I-D.ietf-mmusic-trickle-ice] ***** [RFC4566] [RFC4566] [RFC3605] Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta) [RFC5888] [RFC3264] [I-D.ietf-payload-rtp-opus] [RFC4566] [RFC5245] [RFC5245] [RFC5245]

:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B	
:35 :DC:B8:5F:64:1A:24:C2:43:F0:A	
1:58:D0:A1:2C:19:08	
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-	[RFC6464]
hdext:ssrc-audio-level	
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdext:sdes:mid	egotiation]
a:ssrc:54321	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
a=candidate:0 1 UDP 2122194687	[RFC5245]
192.168.1.7 51556 typ host	
a=candidate:1 1 UDP 1685987071	[RFC5245]
98.248.92.77 49203 typ srflx	
raddr 192.168.1.7 rport 51556	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF	[RFC4566]
120	
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID
	(ma) and RTCMediaStreamTrack
	ID (tb)
a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=fingerprint:sha-256 6B:8B:F0:65	[RFC5245]
:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B	
:35 :DC:B8:5F:64:1A:24:C2:43:F0:A	
1:58:D0:A1:2C:19:08	
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdext:sdes:mid	egotiation]
a:ssrc:98765	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
***** Application m=line	*****

m=application 49203 UDP/DTLS/SCTP	[I-D.ietf-mmusic-sctp-sdp]
webrtc-datachannel	

c=IN IP4 98.248.92.771	[RFC4566]
a=mid:data	[RFC5888]
a=sctp-port:5000	[I-D.ietf-mmusic-sctp-sdp]
a=max-message-size:100000	[I-D.ietf-mmusic-sctp-sdp]
a=setup:active	[RFC4145]
a=sendrecv	[RFC3264]
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:A 1:58:D0:A1:2C:19:08	[RFC5245]

Table 16: 5.2.8 SDP Answer

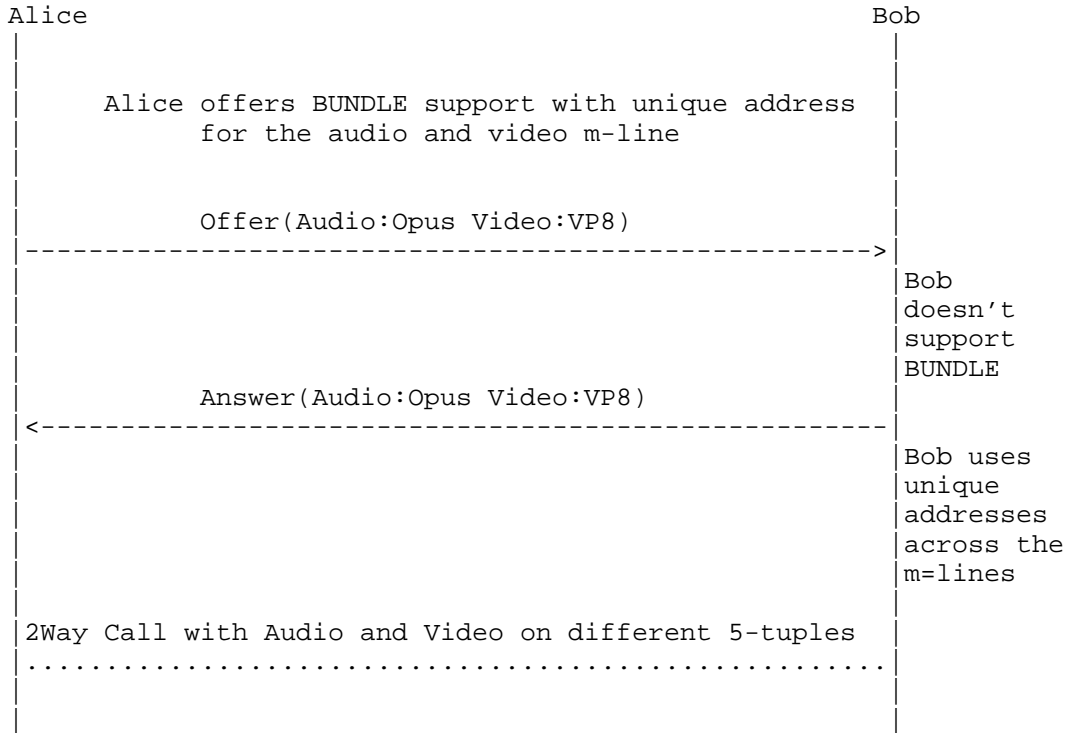
5.2.9. Audio, Video Session with BUNDLE Unsupported

This use-case illustrates SDP Offer/Answer exchange where the far-end (Bob) either doesn't support media bundling or doesn't want to group m=lines over a single 5-tuple.

The same is indicated by dropping the "a=group:BUNDLE" line and BUNDLE RTP header extension in the Answer SDP.

On successful Offer/Answer exchange, Alice and Bob each end up using unique 5-tuple for audio and video media streams respectively.

Two-Way Secure Audio,Video with BUNDLE Unsupported



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:55232 IN IP4 24.23.204.141	[RFC3605] - RTCP port different from RTP port
a=mid:audio	[RFC5888] Audio m=line part

a=msid:ma ta	of BUNDLE group with a unique port number Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta) [RFC3264]
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-ufrag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdrext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2122194687 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 1685987071 24.23.204.141 54609 typ srflx	[RFC5245]
raddr 192.168.1.4 rport 61665	
a=candidate:0 2 UDP 2122194687 192.168.1.4 61666 typ host	[RFC5245]
a=candidate:1 2 UDP 1685987071 24.23.204.141 55232 typ srflx	[RFC5245]
raddr 192.168.1.4 rport 61666	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 54332 UDP/TLS/RTP/SAVPF 120	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:60052 IN IP4 24.23.204.141	[RFC3605]
a=mid:video	[RFC5888] Video m=line part of the BUNDLE group with a unique port number
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)

a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=ice-ufrag:7872093	[RFC5245]
a=ice-pwd:ee3474af08a068a28a397a4c3f31747d1	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145] - Alice can perform DTLS before Answer arrives
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2122194687 192.168.1.4 71775 typ host	[RFC5245]
a=candidate:1 1 UDP 1685987071 24.23.204.141 54332 typ srflx	[RFC5245]
raddr 192.168.1.4 rport 71775	
a=candidate:0 2 2122194687 192.168.1.4 71776 typ host	[RFC5245]
a=candidate:1 2 UDP 1685987071 24.23.204.141 60052 typ srflx	[RFC5245]
raddr 192.168.1.4 rport 71776	

Table 17: 5.2.9 SDP Offer w/BUNDLE

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 53214 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]

<pre> a=rtcp:60065 IN IP4 98.248.92.77 a=msid:ma ta a=sendrecv a=rtpmap:109 opus/48000/2 a=maxptime:120 a=ice-ufrag:c300d85b a=ice- pwd:de4e99bd291c325921d5d47efbabd9a2 a=fingerprint:sha-256 6B:8B:F0:65:5F:7 8:E2:51:3B:AC:6F:F3:3F:46:1B:35 :DC:B8 :5F:64:1A:24:C2:43:F0:A1:58:D0:A1:2C:1 9:08 a=setup:active a=rtcp-rsize a=rtcp-fb:109 nack a=extmap:1 urn:ietf:params:rtp-hdext :ssrc-audio-level a:ssrc:54321 cname:Q/NWslaolHmN4Xa5 a=candidate:0 1 UDP 2122194687 192.168.1.7 51556 typ host a=candidate:1 1 UDP 1685987071 98.248.92.77 53214 typ srflx raddr 192.168.1.7 rport 51556 a=candidate:0 2 UDP 2122194687 192.168.1.7 51558 typ host a=candidate:1 2 UDP 1685987071 98.248.92.77 60065 typ srflx raddr 192.168.1.7 rport 51558 ***** Video m=line ***** m=video 58679 UDP/TLS/RTP/SAVPF 120 c=IN IP4 98.248.92.77 a=rtcp:56507 IN IP4 98.248.92.77 a=msid:ma tb a=sendrecv a=rtpmap:120 VP8/90000 a=ice-ufrag:85bc300 a=ice- pwd:325921d5d47efbabd9a2de4e99bd291c </pre>	<pre> [RFC3605] Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta) [RFC3264] [I-D.ietf-payload-rtp-op us] [RFC4566] [RFC5245] [RFC5245] [RFC5245] [RFC5245] [RFC4145] - Bob carries out DTLS Handshake in parallel [RFC5506] [RFC5104] [RFC6464] [RFC5576] [RFC5245] [RFC5245] [RFC5245] [RFC5245] ***** ***** [RFC4566] [RFC4566] [RFC3605] Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb) [RFC3264] [I-D.ietf-payload-vp8] [RFC5245] [RFC5245] </pre>
---	--


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	[RFC5245]
a=setup:active	[RFC4145] - Bob carries out DTLS Handshake in parallel
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=ssrc:98765 cname:Q/NWslao1HmN4Xa5	[RFC5576]
a=candidate:0 1 UDP 2122194687 192.168.1.7 61556 typ host	[RFC5245]
a=candidate:1 1 UDP 1685987071 98.248.92.77 58679 typ srflx raddr 192.168.1.7 rport 61556	[RFC5245]
a=candidate:0 1 UDP 2122194687 192.168.1.7 61558 typ host	[RFC5245]
a=candidate:1 1 UDP 1685987071 98.248.92.77 56507 typ srflx raddr 192.168.1.7 rport 61558	[RFC5245]

Table 18: 5.2.9 SDP Answer without BUNDLE

5.2.10. Audio, Video BUNDLED, but Data (Not BUNDLED)

This example show-cases SDP for negotiating a session with Audio, Video and data streams between Alice and Bob with data stream not being part of the BUNDLE group. This is shown by assigning unique port for data media section and not adding the "mid" identification tag to the BUNDLE group.

Audio, Video, with Data (Not in BUNDLE)



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice wants to BUNDLE only audio and video media.
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747d1ee34	[RFC5245]

74af08a068	
a=fingerprint:sha-256 19:E2:1C:3B	[RFC5245]
:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73	
:04 :BB:05:2F:70:9F:04:A9:0E:05:E	
9:26:33:E8:70:88:A2	
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-	[RFC6464]
hdext:ssrc-audio-level	
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdext:sdes:mid	egotiation]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2113667327	[RFC5245]
192.168.1.4 54609 typ host	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 54609 UDP/TLS/RTP/SAVPF	[RFC4566]
120	
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID
	(ma) and RTCMediaStreamTrack
	ID (tb)
a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=fingerprint:sha-256 19:E2:1C:3B	[RFC5245]
:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73	
:04 :BB:05:2F:70:9F:04:A9:0E:05:E	
9:26:33:E8:70:88:A2	
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdext:sdes:mid	egotiation]
a:ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	
***** Application m=line	*****

m=application 10000 UDP/DTLS/SCTP	[I-D.ietf-rtcweb-data-channel
webrtc-datachannel]
c=IN IP4 24.23.204.141	[RFC4566]

a=mid:data	[RFC5888]
a=sctp-port:5000	[I-D.ietf-mmusic-sctp-sdp]
a=max-message-size:100000	[I-D.ietf-mmusic-sctp-sdp]
a=sendrecv	[RFC3264]
a=setup:actpass	[RFC4145]
a=connection:new	[RFC4145]
a=ice-ufrag:89819013	[RFC5245]
a=ice-pwd:1747dlee3474af08a068a28 a397a4c3f3	[RFC5245]
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:E 9:26:33:E8:70:88:A2	[RFC5245]
a=candidate:0 1 UDP 2113667327 192.168.1.4 10000 typ host	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 19: 5.2.10 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47 efbabd9a2	[RFC5245]
a=fingerprint:sha-256 6B:8B:F0:65	[RFC5245]

:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B	
:35 :DC:B8:5F:64:1A:24:C2:43:F0:A	
1:58:D0:A1:2C:19:08	
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-	[RFC6464]
hdrext:ssrc-audio-level	
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdrext:sdes:mid	egotiation]
a=ssrc:54321	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
a=candidate:0 1 UDP 2113667327	[RFC5245]
192.168.1.7 49203 typ host	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF	[RFC4566]
120	
c=IN IP4 98.248.92.771	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID
	(ma) and RTCMediaStreamTrack
	ID (tb)
a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
a=fingerprint:sha-256 6B:8B:F0:65	[RFC5245]
:5F:78:E2:51:3B:AC:6F:F3:3F:46:1B	
:35 :DC:B8:5F:64:1A:24:C2:43:F0:A	
1:58:D0:A1:2C:19:08	
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-	[I-D.ietf-mmusic-sdp-bundle-n
hdrext:sdes:mid	egotiation]
a=ssrc:98765	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
***** Application m=line	*****

m=application 20000 UDP/DTLS/SCTP	[I-D.ietf-mmusic-sctp-sdp]
webrtc-datachannel	
c=IN IP4 98.248.92.77	[RFC4566]
a=mid:data	[RFC5888]
a=sctp-port:5000	[I-D.ietf-mmusic-sctp-sdp]


a=max-message-size:100000	[I-D.ietf-mmusic-sctp-sdp]
a=setup:active	[RFC4145]
a=sendrecv	[RFC3264]
a=ice-ufrag:991Ca2a5e	[RFC5245]
a=ice-pwd:921d5d47efbabd9a2de4e99 bd291c325	[RFC5245]
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:A 1:58:D0:A1:2C:19:08	[RFC5245]
a=candidate:0 1 UDP 2113667327 192.168.1.7 20000 typ host	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 20: 5.2.10 SDP Answer

5.2.11. Audio Only, Add Video to BUNDLE

This example involves 2 Offer/Answer exchanges. First one is used to negotiate and setup BUNDLE support for Audio-only session followed by an updated Offer/Answer exchange to add video stream to the ongoing session. Also the newly added video stream is BUNDLED with the audio stream.

Audio Only , Add Video and BUNDLE



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice adds audio m=line to the BUNDLE group
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack

a=sendrecv	ID (ta)
a=rtpmap:109 opus/48000/2	[RFC3264]
a=maxptime:120	[I-D.ietf-payload-rtp-opus]
a=ice-frag:074c6550	[RFC4566]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdrext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 21: 5.2.11 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]

a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-ufraq:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdrext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:54321	[RFC5576]
cname:Q/NWslao1HmN4Xa5	
a=candidate:0 1 UDP 2113667327 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 22: 5.2.10 SDP Answer

Updated Offer SDP Contents	RFC#/Notes
v=1	Version number incremented [RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF	[RFC4566]

109	
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 54609 UDP/TLS/RTP/SAVPF 120	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
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	[RFC5245]

a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:56789	[RFC5576]
cname:EocUG1f0fcg/yvY7	

Table 23: 5.2.11 SDP Updated Offer

Updated Answer SDP Contents	RFC#/Notes
v=1	[RFC4566] Version number incremented
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio video	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]

a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle- negotiation]
a=ssrc:54321	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
a=candidate:0 1 UDP 2113667327 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF 120	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:video	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendrecv	[RFC3264]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:120 nack	[RFC5104]
a=rtcp-fb:120 nack pli	[RFC5104]
a=rtcp-fb:120 ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle- negotiation]
a=ssrc:98765	[RFC5576]
cname:Q/NWslaolHmN4Xa5	

Table 24: 5.2.11 SDP Updated Answer

5.3. MultiResolution, RTX, FEC Examples

This section deals with scenarios related to multi-source, multi-stream negotiation such as layered coding, simulcast, along with techniques that deal with providing robustness against transmission errors such as FEC and RTX. Also to note, mechanisms such as FEC and RTX could be envisioned in the above basic scenarios as well.

5.3.1. Sendonly Simulcast Session with 2 cameras and 2 encodings per camera

The SDP below shows Offer/Answer exchange with one audio and two video sources. Each of the video source can be sent at two different resolutions.

One video source corresponds to VP8 encoding, while the other corresponds to H.264 encoding.


[I-D.ietf-mmusic-rid] framework is used to further constrain the media format encodings and map the payload types (PT) to the 'rid' identifiers.

[I-D.ietf-mmusic-sdp-simulcast] framework identifies the simulcast streams via their 'rid' identifiers.

bundle-only attribute is used for the video sources in the Offer to ensure enabling video sources in the context of BUNDLE alone.

BUNDLE grouping framework enables multiplexing of all the 5 streams (1 audio stream + 4 video streams) over a single RTP Session.

1 Way Successful Simulcast w/BUNDLE



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1 m2	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m0	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)

a=sendonly	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdrext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:11111 cname:EocUG1f0fcg/yvY7	[RFC5576]
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video-1 m=line *****	*****
m=video 0 UDP/TLS/RTP/SAVPF 98 100	**
c=IN IP4 24.23.204.141	bundle-only video line with port number set to zero
a=bundle-only	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=mid:m1	[RFC3605]
a=msid:ma tb	[RFC5888] Video m=line part of BUNDLE group
a=sendonly	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=rtpmap:98 VP8/90000	[RFC3264] - Send only video stream
a=fmtp:98 max-fr=30	[I-D.ietf-payload-vp8]
a=rtpmap:100 VP8/90000	[RFC4566]
a=fmtp:100 max-fr=15	[I-D.ietf-payload-vp8]
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	[RFC4566]
	[RFC5245]

<pre> 3:E8:70:88:A2 a=setup:actpass a=rtcp-mux a=rtcp-mux-only a=rtcp-rsize a=rtcp-fb:* nack a=rtcp-fb:* nack pli a=rtcp-fb:* ccm fir a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid a=ssrc:22222 cname:EocUG1f0fcg/yvY7 a=ssrc:33333 cname:EocUG1f0fcg/yvY7 a=rid:1 send pt=98;max-width=1280 ;max-height=720; a=rid:2 send pt=100;max-width=640 ;max-height=480; a=simulcast: send 1;~2 ***** Video-2 m=line ***** m=video 0 UDP/TLS/RTP/SAVPF 101 102 c=IN IP4 24.23.204.141 a=bundle-only a=rtcp:54609 IN IP4 24.23.204.141 a=mid:m2 a=msid:ma tc a=sendonly a=rtpmap:101 H264/90000 a=rtpmap:102 H264/90000 </pre>	<pre> [RFC4145] [RFC5761] [I-D.ietf-mmusic-mux-exclusive] [RFC5506] [RFC5104] [RFC5104] [RFC5104] [I-D.ietf-mmusic-sdp-bundle-negotiation] [RFC5576] [RFC7022] Camera-1,Encoding-1 SSRC with Session CNAME [RFC5576] [RFC7022] Camera-1,Encoding-2 SSRC with Session CNAME [I-D.ietf-mmusic-rid] 1:1 rid mapping to payload type and specify resolution constraints [I-D.ietf-mmusic-rid] 1:1 rid mapping to payload type and specify resolution constraints [I-D.ietf-mmusic-sdp-simulcast] Alice can send 2 resolutions identified by the 'rid' identifiers Also, the second stream is initially paused. ***** ** bundle-only video line with port number set to zero [RFC4566] [I-D.ietf-mmusic-sdp-bundle-negotiation] [RFC3605] [RFC5888] Video m=line part of BUNDLE group Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tc) [RFC3264] - Send only video stream [RFC3984] [RFC3984] </pre>
--	--

a=fmtp:101 profile-level-id=42401f ;packetization-mode=0;max-fr=30	[RFC3984]Camera-2,Encoding-1
a=fmtp:102 profile-level-id=42401f ;packetization-mode=1;max-fr=15	[RFC3984]Camera-2,Encoding-2
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-mux-only	[I-D.ietf-mmusic-mux-exclusive]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:44444 cname:EocUG1f0fcg/yvY7	[RFC5576] [RFC7022] Camera-2,Encoding-1 SSRC with Session CNAME
a=ssrc:55555 cname:EocUG1f0fcg/yvY7	[RFC5576] [RFC7022] Camera-2,Encoding-2 SSRC with Session CNAME
a=rid:3 send pt=101;max-width=1280;max-height=720;	[I-D.ietf-mmusic-rid] 1:1 rid mapping to payload type and specify resolution constraints
a=rid:4 send pt=102;max-width=640;max-height=360;	[I-D.ietf-mmusic-rid] 1:1 rid mapping to payload type and specify resolution constraints
a=simulcast: send 3;4	[I-D.ietf-mmusic-sdp-simulcast] Alice can send 2 resolutions identified by the 'rid' identifiers

Table 25: 5.3.1 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]

<pre> a=group:BUNDLE m0 m1 m2 a=ice-options:trickle ***** Audio m=line ***** m=audio 49203 UDP/TLS/RTP/SAVPF 109 c=IN IP4 98.248.92.77 a=rtcp:49203 IN IP4 98.248.92.77 a=mid:m0 a=msid:ma ta a=recvonly a=rtpmap:109 opus/48000/2 a=rtcp-fb:109 nack a=maxptime:120 a=ice-ufraq:c300d85b a=ice-pwd:de4e99bd291c325921d5d47ef babb9a2 a=fingerprint:sha-256 6B:8B:F0:65:5 F:78:E2:51:3B:AC:6F:F3:3F:46:1B:35 :DC:B8:5F:64:1A:24:C2:43:F0:A1:58:D 0:A1:2C:19:08 a=setup:active a=rtcp-mux a=rtcp-rsize a=extmap:1 urn:ietf:params:rtp- hdnext:ssrc-audio-level a=extmap:2 urn:ietf:params:rtp- hdnext:sdes:mid a=ssrc:77777 cname:Q/NWslao1HmN4Xa5 a=candidate:0 1 UDP 2113667327 192.168.1.7 61665 typ host a=candidate:1 1 UDP 694302207 98.248.92.77 49203 typ srflx raddr 192.168.1.4 rport 61665 a=end-of-candidates ***** Video-1 m=line ***** m=video 49203 UDP/TLS/RTP/SAVPF 98 100 c=IN IP4 98.248.92.77 </pre>	<pre> [I-D.ietf-mmusic-sdp-bundle -negotiation] Alice supports grouping of m=lines under BUNDLE semantics [I-D.ietf-mmusic-trickle-ic e] ***** ** [RFC4566] [RFC4566] [RFC3605] [RFC5888] Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta) [RFC3264] [I-D.ietf-payload-rtp-opus] [RFC5104] [RFC4566] [RFC5245] [RFC5245] [RFC5245] [RFC4145] [RFC5761] [RFC5506] [RFC6464] [I-D.ietf-mmusic-sdp-bundle -negotiation] [RFC5576] [RFC5245] [RFC5245] [I-D.ietf-mmusic-trickle-ic e] ***** ** BUNDLE accepted with port repeated from the audio port [RFC4566] </pre>
--	---

<pre> a=rtcp:49203 IN IP4 98.248.92.77 a=mid:m1 a=msid:ma tb a=recvonly a=rtpmap:98 VP8/90000 a=rtpmap:100 VP8/90000 a=fmtp:98 max-fr=30 a=fmtp:100 max-fr=15 a=fingerprint:sha-256 6B:8B:F0:65:5 F:78:E2:51:3B:AC:6F:F3:3F:46:1B:35 :DC:B8:5F:64:1A:24:C2:43:F0:A1:58:D 0:A1:2C:19:08 a=setup:active a=rtcp-mux a=rtcp-rsize a=rtcp-fb:* nack a=rtcp-fb:* nack pli a=rtcp-fb:* ccm fir a=extmap:2 urn:ietf:params:rtp- hdext:sdes:mid a=ssrc:88888 cname:Q/NWslao1HmN4Xa5 a=rid:1 recv pt=98;max-width=1280 ;max-height=720; a=rid:2 recv pt=100;max-width=640 ;max-height=480; a=simulcast: recv 1;2 ***** Video-2 m=line ***** m=video 49203 UDP/TLS/RTP/SAVPF 101 102 c=IN IP4 98.248.92.77 a=rtcp:49203 IN IP4 98.248.92.77 a=mid:m2 </pre>	<pre> [RFC3605] [RFC5888] Video m=line part of BUNDLE group Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb) [RFC3264] - receive only video stream [I-D.ietf-payload-vp8] [I-D.ietf-payload-vp8] [RFC4566] [RFC4566] [RFC5245] [RFC4145] [RFC5576] [RFC5506] [RFC5104] [RFC5104] [RFC5104] [RFC5104] [I-D.ietf-mmusic-sdp-bundle -negotiation] [RFC5576] - SSRC associated with the receiver report from Bob [I-D.ietf-mmusic-rid] Bob accepts the offered payload format constraints [I-D.ietf-mmusic-rid] Bob accepts the offered payload format constraints [I-D.ietf-mmusic-sdp-simulc ast] Bob accepts the offered simulcast streams and removes the paused state of stream with 'rid' value 2. ***** ** BUNDLE accepted with port repeated from the audio port [RFC4566] [RFC3605] [RFC5888] Video m=line part of BUNDLE group </pre>
---	---

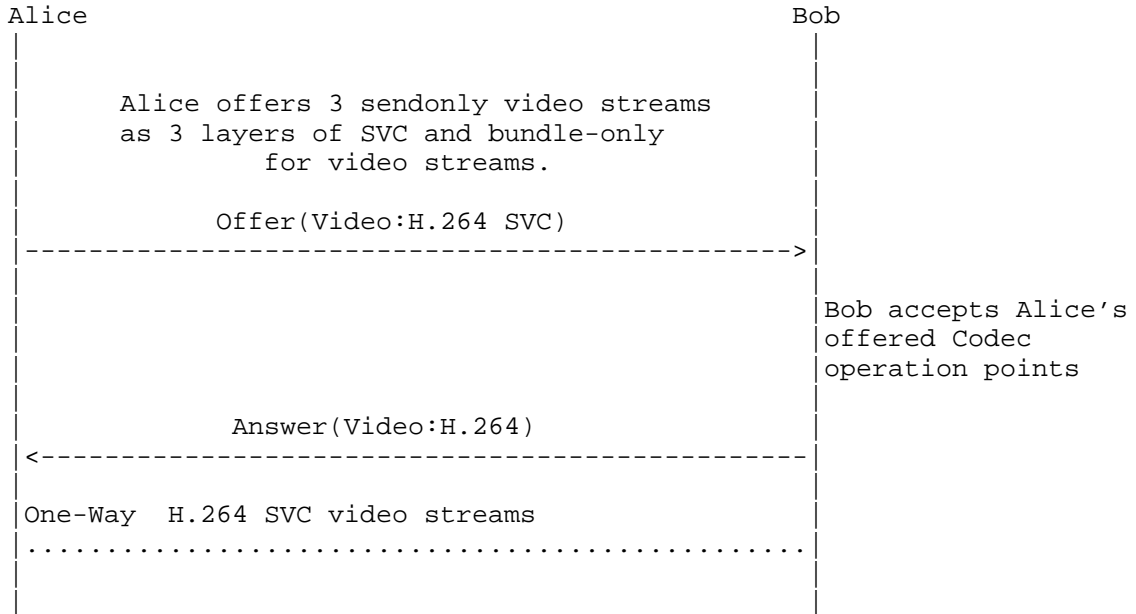
a=msid:ma tc	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tc)
a=recvonly	[RFC3264]
a=rtpmap:101 H264/90000	[RFC3984]
a=rtpmap:102 H264/90000	[RFC3984]
a=fmtp:101 profile-level-id=42401f ;packetization-mode=1;max-fr=30	[RFC3984]
a=fmtp:102 profile-level-id=42401f ;packetization-mode=1;max-fr=15	[RFC3984]
a=fingerprint:sha-256 6B:8B:F0:65:5 F:78:E2:51:3B:AC:6F:F3:3F:46:1B:35 :DC:B8:5F:64:1A:24:C2:43:F0:A1:58:D 0:A1:2C:19:08	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5576]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:99999 cname:Q/NWslaolHmN4Xa5	[RFC5576] - SSRC associated with the receiver report from Bob
a=rid:3 recv pt=101;max-width=1280 ;max-height=720;	[I-D.ietf-mmusic-rid] Bob accepts the offered payload format constraints
a=rid:4 recv pt=102;max-width=640 ;max-height=360;	[I-D.ietf-mmusic-rid] Bob accepts the offered payload format constraints
a=simulcast: recv 3;4	[I-D.ietf-mmusic-sdp-simulcast] Bob accepts the offered simulcast streams.

Table 26: 5.3.1 SDP Answer

5.3.2. Successful SVC Video Session

This section shows an SDP Offer/Answer for a session with an audio and a single video source. The video source is encoded as layered coding at 3 different resolutions based on [RFC5583]. The video m=line shows 3 streams with last stream (payload 100) dependent on streams with payload 96 and 97 for decoding.

SVC Session - 3 Layers w/BUNDLE



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m0	[RFC5888] Audio m=line part of BUNDLE group with a unique port number
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack

a=sendonly	ID (ta)
a=rtpmap:109 opus/48000/2	[RFC3264]
a=maxptime:120	[I-D.ietf-payload-rtp-opus]
a=ice-frag:074c6550	[RFC4566]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:11111	[RFC5576]
cname:EocUGlf0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 0 UDP/TLS/RTP/SAVPF 96 97 100	bundle-only video line with port number set to zero
c=IN IP4 24.23.204.141	[RFC4566]
a=bundle-only	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m1	[RFC5888] Video m=line part of BUNDLE group
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tc)
a=sendonly	[RFC3264] - Send only video stream
a=rtpmap:96 H264/90000	[RFC3984]
a=fmtp:96 profile-level-id=4d0028; packetization-mode=1;max-fr=30;max-fs=8040	[RFC3984]H.264 Layer 1
a=rtpmap:97 H264/90000	[RFC3984]
a=fmtp:97 profile-level-id=4d0028;packetization-mode=1; max-fr=15	[RFC3984] H.264 Layer 2

<pre> ;max-fs=1200 a=rtpmap:100 H264-SVC/90000 a=fmtp:100 profile-level- id=4d0028;packetization-mode=1; max-fr=30;max-fs=8040 a=depend:100 lay m1:96,97; 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:E 9:26:33:E8:70:88:A2 a=setup:actpass a=rtcp-mux a=rtcp-rsize a=rtcp-fb:* nack a=rtcp-fb:* nack pli a=rtcp-fb:* ccm fir a=extmap:2 urn:ietf:params:rtp- hdnext:sdes:mid a=ssrc:22222 cname:EocUG1f0fcg/yvY7 a=ssrc:33333 cname:EocUG1f0fcg/yvY7 a=ssrc:44444 cname:EocUG1f0fcg/yvY7 </pre>	<pre> [RFC3984] [RFC3984] [RFC5583]Layer 3 dependent on layers 1 and 2 [RFC5245] [RFC4145] [RFC5761] [RFC5506] [RFC5104] [RFC5104] [RFC5104] [I-D.ietf-mmusic-sdp-bundle-n egotiation] [RFC5576] [RFC5576] [RFC5576] </pre>
---	---

Table 27: 5.3.2 SDP Offer with SVC

Answer SDP Contents	RFC#/Notes
<pre> v=0 o=- 20519 0 IN IP4 0.0.0.0 s=- t=0 0 a=group:BUNDLE m0 m1 a=ice-options:trickle ***** Audio m=line ***** m=audio 49203 UDP/TLS/RTP/SAVPF 109 c=IN IP4 98.248.92.77 a=rtcp:49203 IN IP4 98.248.92.77 a=mid:m0 a=msid:ma ta </pre>	<pre> [RFC4566] [RFC4566] [RFC4566] [RFC4566] [I-D.ietf-mmusic-sdp-bundle-n egotiation] [I-D.ietf-mmusic-trickle-ice] ***** [RFC4566] [RFC4566] [RFC3605] [RFC5888] Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack </pre>

a=recvonly	ID (ta)
a=rtpmap:109 opus/48000/2	[RFC3264]
a=maxptime:120	[I-D.ietf-payload-rtp-opus]
a=ice-frag:074c6550	[RFC4566]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:88888	[RFC5576]
cname:Q/NWslaolHmN4Xa5	
a=candidate:0 1 UDP 2113667326 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302206 98.248.92.77 49203 typ srflx raddr 192.168.1.5 rport 51556	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF 96 100	BUNDLE accepted Bundle address same as audio m=line.
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:m1	[RFC5888] Video m=line part of BUNDLE group
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=recvonly	[RFC3264] - Receive only video stream
a=rtpmap:96 H264/90000	[RFC3984]
a=fmtp:96 profile-level-id=4d0028;packetization-mode=1; max-fr=30;max-fs=8040	[RFC3984]H.264 Layer 1
a=rtpmap:100 H264-SVC/90000	[RFC3984]
a=fmtp:100 profile-level-id=4d0028;packetization-mode=1;max-fr=30;max-fs=8040	[RFC3984]
a=depend:100 lay m1:96;	[RFC5583] Bob chooses 2 Codec

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:A 1:58:D0:A1:2C:19:08	Operation points [RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp- hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a=ssrc:99999	[RFC5576]
cname:Q/NWslaolHmN4Xa5	

Table 28: 5.3.2 SDP Answer with SVC


5.3.3. Successful Simulcast Video Session with Retransmission

This section shows an SDP Offer/Answer exchange for a simulcast scenario with 3 resolutions and has [RFC4588] style re-transmission flows.

[I-D.ietf-mmusic-rid] framework is used to specify all the (3) resolution constraints mapped to a single Payload Type (98).

[I-D.ietf-mmusic-sdp-simulcast] framework identifies the simulcast streams via their 'rid' identifiers.

Simulcast Streams with Retransmission



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m0	[RFC5888] Audio m=line part of BUNDLE group with a unique port number
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendonly	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]

a=maxptime:120	[RFC4566]
a=ice-ufrag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747d1ee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:11111	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 0 UDP/TLS/RTP/SAVPF 98 103	bundle-only video line with port number set to zero
c=IN IP4 24.23.204.141	[RFC4566]
a=bundle-only	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=mid:m1	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendonly	[RFC3264]
a=rtpmap:98 VP8/90000	[I-D.ietf-payload-vp8]
a=fmtp:98 max-fr=30	[RFC4566]
a=rtpmap:103 rtx/90000	[RFC4588]
a=fmtp:103 apt=98;rtx-time=200	[RFC4588]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:22222	[RFC5576] - Encoding-1 SSRC

cname:EocUG1f0fcg/yvY7	
a:ssrc:33333	[RFC5576] - Encoding-1 RTX
cname:EocUG1f0fcg/yvY7	SSRC
a:ssrc-group:FID 22222 33333	[RFC5888]
a:ssrc:44444	[RFC5576] - Encoding-2 SSRC
cname:EocUG1f0fcg/yvY7	
a:ssrc:55555	[RFC5576] - Encoding-2 RTX
cname:EocUG1f0fcg/yvY7	SSRC
a:ssrc-group:FID 44444 55555	[RFC5888]
a:ssrc:66666	[RFC5576] - Encoding-3 SSRC
cname:EocUG1f0fcg/yvY7	
a:ssrc:77777	[RFC5576] - Encoding-3 RTX
cname:EocUG1f0fcg/yvY7	SSRC
a:ssrc-group:FID 66666 77777	[RFC5888]
a:rid:1 send pt=98;max-fs=921600	[I-D.ietf-mmusic-rid]
;max-fr=30;	
a:rid:2 send pt=98;max-fs=614400	[I-D.ietf-mmusic-rid]
;max-fr=15;	
a:rid:3 send pt=98;max-fs=230400	[I-D.ietf-mmusic-rid]
;max-fr=30;	
a=simulcast: send 1;2;3	[I-D.ietf-mmusic-sdp-simulcast] Alice can send all the simulcast streams

Table 29: 5.3.3 SDP Offer w/Simulcast, RTX

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1	[I-D.ietf-mmusic-sdp-bundle-negotiation] Bob supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:m0	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack

a=recvonly	ID (ta)
a=rtpmap:109 opus/48000/2	[RFC3264]
a=maxptime:120	[I-D.ietf-payload-rtp-opus]
a=ice-frag:074c6550	[RFC4566]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:54321	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2113667326 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302206 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 49203 UDP/TLS/RTP/SAVPF 98 100 101 103	BUNDLE accepted with Bundle address identical to audio m-line
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:m1	[RFC5888] Video m=line part of BUNDLE group
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=recvonly	[RFC3264]
a=rtpmap:98 VP8/90000	[I-D.ietf-payload-vp8]
a=fmtp:98 max-fr=30	[RFC4566]
a=rtpmap:103 rtx/90000	[RFC4588]
a=fmtp:103 apt=98;rtx-time=200	[RFC4588]
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	[RFC5245]
a=setup:active	[RFC4145]

a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle- negotiation]
a=ssrc:98765	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=rid:1 recv pt=98;max-fs=921600 ;max-fr=30;	[I-D.ietf-mmusic-rid]
a=rid:2 recv pt=98;max-fs=614400 ;max-fr=15;	[I-D.ietf-mmusic-rid]
a=rid:3 recv pt=98;max-fs=230400 ;max-fr=30;	[I-D.ietf-mmusic-rid]
a=simulcast: recv 1;2;3	[I-D.ietf-mmusic-sdp-simulcast] Bob accepts the offered simulcast streams


Table 30: 5.3.3 SDP Answer w/Simulcast, RTX

5.3.4. Successful 1-way Simulcast Session with 2 resolutions and RTX - One resolution rejected

This section shows an SDP Offer/Answer exchange for a simulcast scenario with 2 two resolutions.

It also showcases where Bob rejects one of the Simulcast Video Stream which results in the rejection of the associated repair stream implicitly.

Simulcast Streams with Retransmission Rejected



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m0	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack

a=sendonly	ID (ta)
a=rtpmap:109 opus/48000/2	[RFC3264]
a=maxptime:120	[I-D.ietf-payload-rtp-opus]
a=ice-frag:074c6550	[RFC4566]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:11111	[RFC5576]
cname:EocUGlf0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 0 UDP/TLS/RTP/SAVPF 98 100 101 103	bundle-only video line with port number set to zero
c=IN IP4 24.23.204.141	[RFC4566]
a=bundle-only	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m1	[RFC5888]
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendonly	[RFC3264]
a=rtpmap:98 VP8/90000	[I-D.ietf-payload-vp8]
a=rtpmap:100 VP8/90000	[I-D.ietf-payload-vp8]
a=rtpmap:101 rtx/90000	[RFC4588]
a=rtpmap:103 rtx/90000	[RFC4588]
a=fmtp:98 max-fr=30;max-fs=8040	[RFC4566]
a=fmtp:100 max-fr=15;max-fs=1200	[RFC4566]
a=fmtp:101 apt=98;rtx-time=200	[RFC4588]
a=fmtp:103 apt=100;rtx-time=200	[RFC4588]
a=setup:actpass	[RFC4145]

a=rtcp-mux	[RFC5761]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=rtcp-rsize	[RFC5506]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle- negotiation]
a=ssrc:22222	[RFC5576] - Encoding-1 SSRC
cname:EocUG1f0fcg/yvY7	
a=ssrc:33333	[RFC5576] - Encoding-1 RTX SSRC
cname:EocUG1f0fcg/yvY7	
a=ssrc-group:FID 22222 33333	[RFC5888]
a=ssrc:44444	[RFC5576] - Encoding-2 SSRC
cname:EocUG1f0fcg/yvY7	
a=ssrc:55555	[RFC5576] - Encoding-2 RTX SSRC
cname:EocUG1f0fcg/yvY7	
a=ssrc-group:FID 44444 55555	[RFC5888]
a=rid:1 send pt=98;	[I-D.ietf-mmusic-rid] 1:1 mapping between the PT and the 'rid' identifier
a=rid:2 send pt=100;	[I-D.ietf-mmusic-rid] 1:1 mapping between the PT and the 'rid' identifier
a=simulcast: send 1;2	[I-D.ietf-mmusic-sdp-simulcas t]

Table 31: 5.3.4 SDP Offer w/Simulcast, RTX

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1	[I-D.ietf-mmusic-sdp-bundle- negotiation] Bob supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:m0	[RFC5888]

<pre> a=msid:ma ta a=recvonly a=rtpmap:109 opus/48000/2 a=maxptime:120 a=ice-ufrag:074c6550 a=ice-pwd:a28a397a4c3f31747dlee34 74af08a068 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:A 1:58:D0:A1:2C:19:08 a=setup:active a=rtcp-mux a=rtcp-rsize a=extmap:1 urn:ietf:params:rtp- hdnext:ssrc-audio-level a=extmap:2 urn:ietf:params:rtp- hdnext:sdes:mid a:ssrc:54321 cname:Q/NWslao1HmN4Xa5 a=candidate:0 1 UDP 2113667326 192.168.1.7 51556 typ host a=candidate:1 1 UDP 1694302206 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556 a=end-of-candidates ***** Video m=line ***** m=video 49203 UDP/TLS/RTP/SAVPF 98 101 c=IN IP4 98.248.92.77 a=rtcp:49203 IN IP4 98.248.92.77 a=mid:m1 a=msid:ma tb a=recvonly a=rtpmap:98 VP8/90000 a=rtpmap:101 VP8/90000 a=fmtp:98 max-fr=30;max-fs=8040 a=fmtp:101 apt=98;rtx-time=200 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:A 1:58:D0:A1:2C:19:08 a=setup:active </pre>	<pre> Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta) [RFC3264] [I-D.ietf-payload-rtp-opus] [RFC4566] [RFC5245] [RFC5245] [RFC5245] [RFC4145] [RFC5761] [RFC5506] [RFC6464] [I-D.ietf-mmusic-sdp-bundle-n egotiation] [RFC5576] [RFC5245] [RFC5245] [I-D.ietf-mmusic-trickle-ice] ***** BUNDLE accepted with Bundle address identical to audio m-line [RFC4566] [RFC3605] [RFC5888] Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb) [RFC3264] [I-D.ietf-payload-vp8] [I-D.ietf-payload-vp8] [RFC4566] [RFC4588] [RFC5245] [RFC4145] </pre>
---	--

a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ssrc:98765	[RFC5576]
cname:Q/NWslao1HmN4Xa5	
a=rid:1 recv pt=98;	[I-D.ietf-mmusic-rid]
a=simulcast: recv 1	[I-D.ietf-mmusic-sdp-simulcast] Bob rejects the second simulcast stream and the associated rtx stream.


Table 32: 5.3.4 SDP Answer (one Simulcast Rejected)

5.3.5. Simulcast Video Session with Forward Error Correction

This section shows an SDP Offer/Answer exchange for Simulcast video stream at two resolutions and has [RFC5956] style FEC flows.

On completion of the Offer/Answer exchange mechanism we end up one audio stream, 2 simulcast video streams and 2 associated FEC streams are sent over a single 5-tuple.

Simulcast Streams with Forward Error Correction



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1	[I-D.ietf-mmusic-sdp-bundle-negotiation] Alice supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m0	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack

a=sendonly	ID (ta)
a=rtpmap:109 opus/48000/2	[RFC3264]
a=maxptime:120	[I-D.ietf-payload-rtp-opus]
a=ice-frag:074c6550	[RFC4566]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:11111	[RFC5576]
cname:EocUGlf0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]
***** Video m=line *****	*****
m=video 0 UDP/TLS/RTP/SAVPF 98 100 101 103	bundle-only video line with port number set to zero
c=IN IP4 24.23.204.141	[RFC4566]
a=bundle-only	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:m1	[RFC5888] Video m=line part of BUNDLE group
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=sendonly	[RFC3264]
a=rtpmap:98 VP8/90000	[I-D.ietf-payload-vp8]
a=rtpmap:100 VP8/90000	[I-D.ietf-payload-vp8]
a=rtpmap:101 flexfec/90000	[I-D.ietf-payload-flexible-fec-scheme]
a=rtpmap:103 flexfec/90000	[I-D.ietf-payload-flexible-fec-scheme]
a=fmtp:98 max-fr=30;max-fs=8040	[RFC4566]
a=fmtp:100 max-fr=15;max-fs=1200	[RFC4566]

a=fmtp:101 L=5; D=10; ToP=2; repair-window=200000	[I-D.ietf-payload-flexible-fe c-scheme]
a=fmtp:103 L=5; D=10; ToP=2; repair-window=200000	[I-D.ietf-payload-flexible-fe c-scheme]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp- hdnext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a=ssrc:22222	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=ssrc:33333	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=ssrc-group:FEC-FR 22222 33333	[RFC5956]
a=ssrc:44444	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=ssrc:55555	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=ssrc-group:FEC-FR 44444 55555	[RFC5956]
a=rid:1 send pt=98;	[I-D.ietf-mmusic-rid] 1:1 mapping between the PT and the 'rid' identifier
a=rid:2 send pt=100;	[I-D.ietf-mmusic-rid] 1:1 mapping between the PT and the 'rid' identifier
a=simulcast: send 1;2	[I-D.ietf-mmusic-sdp-simulcas t]

Table 33: 5.3.5 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]

c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:m0	[RFC5888] Audio m=line part of BUNDLE group with a unique port number
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=recvonly	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747dlee34 74af08a068	[RFC5245]
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:A 1:58:D0:A1:2C:19:08	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp- hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp- hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a:ssrc:54321	[RFC5576]
cname:Q/NWslao1HmN4Xa5	
a=candidate:0 1 UDP 2113667326 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302206 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245]
a=end-of-candidates	
***** Video m=line *****	[I-D.ietf-mmusic-trickle-ice] *****
m=video 49203 UDP/TLS/RTP/SAVPF 98 100 101 103	BUNDLE accepted with Bundle Address identical to audio m=line.
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:m1	[RFC5888] Video m=line part of BUNDLE group
a=msid:ma tb	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb)
a=recvonly	[RFC3264]
a=rtpmap:98 VP8/90000	[I-D.ietf-payload-vp8]
a=rtpmap:100 VP8/90000	[I-D.ietf-payload-vp8]

a=rtpmap:101 flexfec/90000	[I-D.ietf-payload-flexible-fe c-scheme]
a=rtpmap:103 flexfec/90000	[I-D.ietf-payload-flexible-fe c-scheme]
a=fmtp:98 max-fr=30;max-fs=8040	[RFC4566]
a=fmtp:100 max-fr=15;max-fs=1200	[RFC4566]
a=fmtp:101 L=5; D=10; ToP=2; repair-window=200000	[I-D.ietf-payload-flexible-fe c-scheme]
a=fmtp:103 L=5; D=10; ToP=2; repair-window=200000	[I-D.ietf-payload-flexible-fe c-scheme]
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:A 1:58:D0:A1:2C:19:08	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=rtcp-fb:* nack pli	[RFC5104]
a=rtcp-fb:* ccm fir	[RFC5104]
a=extmap:2 urn:ietf:params:rtp- hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-n egotiation]
a=ssrc:98765	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=rid:1 recv pt=98;	[I-D.ietf-mmusic-rid]
a=rid:2 recv pt=100;	[I-D.ietf-mmusic-rid]
a=simulcast: recv 1;2	[I-D.ietf-mmusic-sdp-simulcas t]

Table 34: 5.3.5 SDP Answer

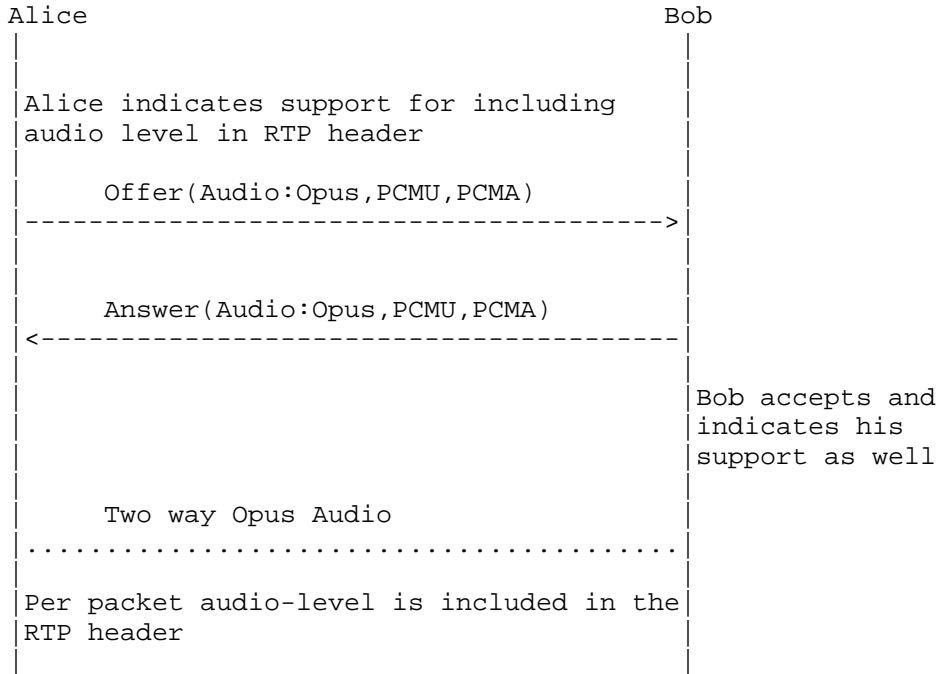
5.4. Others

The examples in the section provide SDP Offer/Answer exchange for a variety of scenarios related to RTP Header extension for conference usages, Legacy Interop scenarios and more.

5.4.1. Audio Session - Voice Activity Detection

This example shows Alice indicating the support of the RTP header extension to include the audio-level of the audio sample carried in the RTP packet.

2-Way Audio with VAD



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109 0 8	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]

a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=rtpmap:0 PCMU/8000	[RFC3551]
a=rtpmap:8 PCMA/8000	[RFC3551]
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747dlee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ssrc:12345	[RFC5576]
cname:EocUGlf0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 35: 5.4.1 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109 0 98	[RFC4566]
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:audio	[RFC5888]

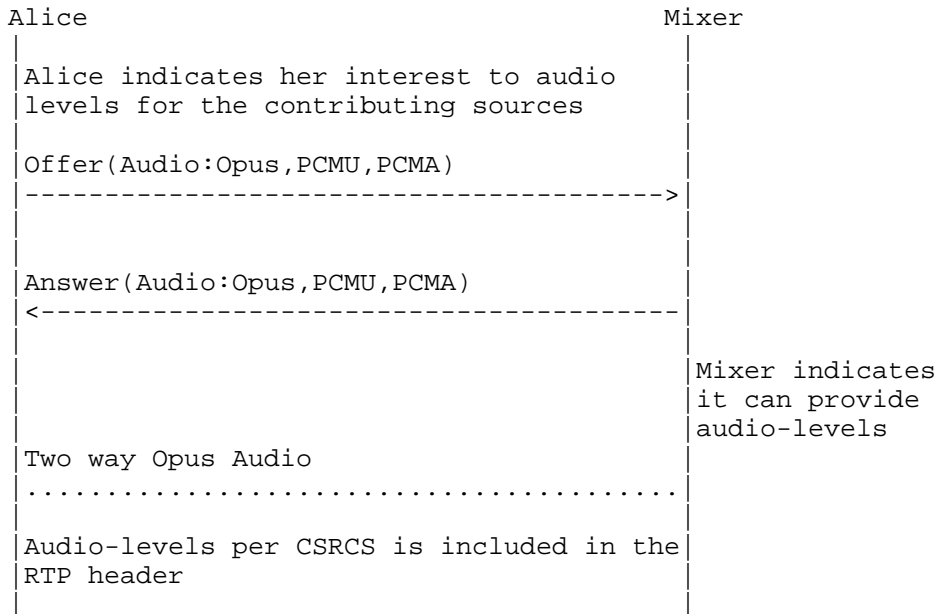
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Bob can send and recv audio
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus] - Bob accepts only Opus Codec
a=rtpmap:0 PCMU/8000	[RFC3551] PCMU Audio Codec
a=rtpmap:0 PCMA/8000	[RFC3551] PCMA Audio Codec
a=maxptime:120	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761] - Bob can perform RTP/RTCP Muxing on port 49203
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:54321	[RFC5576]
cname:Q/NWslao1HmN4Xa5	
a=candidate:0 1 UDP 2113667327 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 98.248.92.77 49203 typ srflx	[RFC5245]
raddr 192.168.1.7 rport 51556	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 36: 5.4.1 SDP Answer

5.4.2. Audio Conference - Voice Activity Detection

This example shows SDP for RTP header extension that allows RTP-level mixers in audio conferences to deliver information about the audio level of individual participants.

Audio Conference with VAD Support



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 54609 UDP/TLS/RTP/SAVPF 109 0 8	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:54609 IN IP4 24.23.204.141	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264] - Alice can send and recv audio

a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=rtpmap:0 PCMU/8000	[RFC3551] PCMU Audio Codec
a=rtpmap:0 PCMA/8000	[RFC3551] PCMA Audio Codec
a=maxptime:120	[RFC4566]
a=ice-frag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747d1ee3474af08a068	[RFC5245]
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	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=extmap:1/recvonly	[RFC6465]
urn:ietf:params:rtp-hdext:csrc-audio-level	
a=extmap:2 urn:ietf:params:rtp-hdext:ssrc-audio-level	[RFC6464]
a=extmap:3 urn:ietf:params:rtp-hdext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a:ssrc:12345	[RFC5576]
cname:EocUG1f0fcg/yvY7	
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 37: 5.4.2 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566] - Session Origin Information
s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE audio	[I-D.ietf-mmusic-sdp-bundle-NEGOTIATION]
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ice]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF	[RFC4566]

109 0 98	
c=IN IP4 98.248.92.77	[RFC4566]
a=rtcp:49203 IN IP4 98.248.92.77	[RFC3605]
a=mid:audio	[RFC5888]
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=rtpmap:0 PCMU/8000	[RFC3551] PCMU Audio Codec
a=rtpmap:0 PCMA/8000	[RFC3551] PCMA Audio Codec
a=maxptime:120	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:* nack	[RFC5104]
a=extmap:1/sendonly	[RFC6465]
urn:ietf:params:rtp-hdrext:csrc-audio-level	
a=extmap:2 urn:ietf:params:rtp-hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle-negotiation]
a:ssrc:54321	[RFC5576]
cname:Q/NWslao1HmN4Xa5	
a=candidate:0 1 UDP 2113667327 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 98.248.92.77 49203 typ srflx	[RFC5245]
raddr 192.168.1.7 rport 51556	
a=end-of-candidates	[I-D.ietf-mmusic-trickle-ice]

Table 38: 5.4.2 SDP Answer

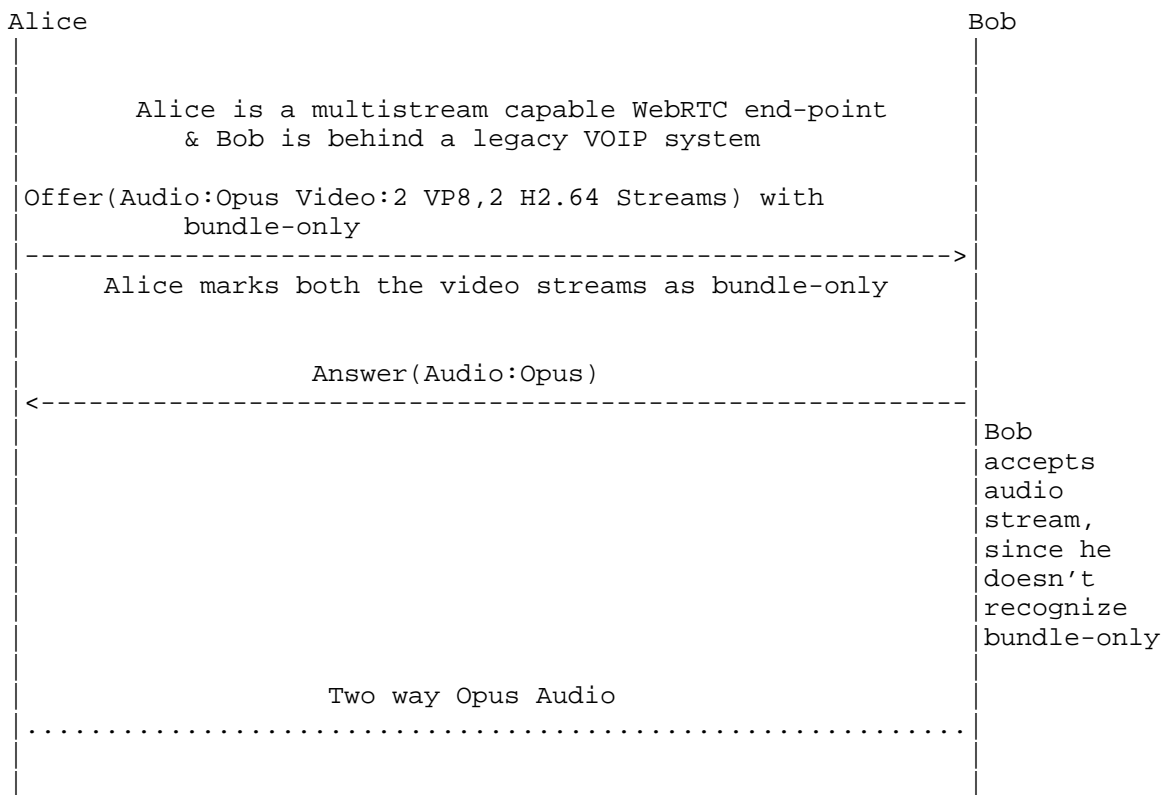
5.4.3. Successful legacy Interop Fallback with bundle-only

In the scenario described below, Alice is a multi-stream capable WebRTC endpoint while Bob is a legacy VOIP end-point. The SDP Offer/Answer exchange demonstrates successful session setup with fallback to audio only stream negotiated via bundle-only framework between the end-points. Specifically,

- o Offer from Alice describes 2 cameras via 2 video m=lines with both marked as bundle-only.
- o Since Bob doesnot recognize either the BUNDLE mechanism or the bundle-only attribute, he accepts only the audio stream from Alice.

NOTE: Since Alice is unaware of Bob's support for BUNDLE framework, Alice ensures to include separate RTP/RTCP ports and candidate information.

Successful 2-Way WebRTC <-> VOIP Interop



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]

s=-	[RFC4566]
t=0 0	[RFC4566]
a=group:BUNDLE m0 m1 m2	[I-D.ietf-mmusic-sdp-bundle -negotiation] Alice supports grouping of m=lines under BUNDLE semantics
a=ice-options:trickle	[I-D.ietf-mmusic-trickle-ic e]
***** Audio m=line *****	***** **
m=audio 54609 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:64678 IN IP4 24.23.204.141	[RFC3605]
a=mid:m0	[RFC5888] Audio m=line part of BUNDLE group with a unique port number
a=msid:ma ta	Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (ta)
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=rtcp-fb:109 nack	[RFC5104]
a=maxptime:120	[RFC4566]
a=ice-ufrag:074c6550	[RFC5245]
a=ice-pwd:a28a397a4c3f31747d1ee3474 af08a068	[RFC5245]
a=fingerprint:sha-256 19:E2:1C:3B:4 B:9F:81:E6:B8:5C:F4:A5:A8:D8:73:04 :BB:05:2F:70:9F:04:A9:0E:05:E9:26:3 3:E8:70:88:A2	[RFC5245]
a=setup:actpass	[RFC4145]
a=rtcp-mux	[RFC5761]
a=rtcp-rsize	[RFC5506]
a=extmap:1 urn:ietf:params:rtp- hdrext:ssrc-audio-level	[RFC6464]
a=extmap:2 urn:ietf:params:rtp- hdrext:sdes:mid	[I-D.ietf-mmusic-sdp-bundle -negotiation]
a=ssrc:12345 cname:EocUG1f0fcg/yvY7	[RFC5576]E
a=candidate:0 1 UDP 2113667327 192.168.1.4 61665 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 61665	[RFC5245]
a=candidate:0 1 UDP 2113667326 192.168.1.4 61667 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302206 24.23.204.141 64678 typ srflx raddr	[RFC5245]

<pre> 192.168.1.4 rport 61667 ***** Video-1 m=line ***** m=video 0 UDP/TLS/RTP/SAVPF 98 100 c=IN IP4 24.23.204.141 a=bundle-only a=mid:m1 a=msid:ma tb a=sendrecv a=rtpmap:98 VP8/90000 a=imageattr:98 [x=1280,y=720] a=fmtp:98 max-fr=30 a=setup:actpass a=rtcp-mux a=rtcp-rsize a=rtcp-fb:* nack a=rtcp-fb:* nack pli a=rtcp-fb:* ccm fir a=extmap:2 urn:ietf:params:rtp-hdext:sdes:mid a:ssrc:56789 cname:EocUG1f0fcg/yvY7 ***** Video-2 m=line ***** m=video 0 UDP/TLS/RTP/SAVPF 101 103 c=IN IP4 24.23.204.141 a=bundle-only a=mid:m2 a=msid:ma tc a=sendrecv a=rtpmap:101 H264/90000 a=rtpmap:103 H264/90000 a=fmtp:101 profile-level-id=4d0028 ;packetization-mode=1;max-fr=30 a=rtcp-mux a=rtcp-rsize a=rtcp-fb:* nack a=rtcp-fb:* nack pli a=rtcp-fb:* ccm fir </pre>	<pre> ***** ** bundle-only video line with port number set to zero [RFC4566] [I-D.ietf-mmusic-sdp-bundle- negotiation] [RFC5888] Video m=line part of BUNDLE group Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tb) [RFC3264] [I-D.ietf-payload-vp8] [RFC6236] [RFC4566] [RFC4145] [RFC5761] [RFC5506] [RFC5104] [RFC5104] [RFC5104] [I-D.ietf-mmusic-sdp-bundle- negotiation] [RFC5576] ***** ** bundle-only video line with port number set to zero [RFC4566] [I-D.ietf-mmusic-sdp-bundle- negotiation] [RFC5888] Video m=line part of BUNDLE group Identifies RTCMediaStream ID (ma) and RTCMediaStreamTrack ID (tc) [RFC3264] [RFC3984] [RFC3984] [RFC3984]Camera-2,Encoding- 1 Resolution [RFC5761] [RFC5506] [RFC5104] [RFC5104] [RFC5104] </pre>
--	--

a=extmap:2 urn:ietf:params:rtp-hdrex	[I-D.ietf-mmusic-sdp-bundle
hdrex:sdes:mid	-negotiation]
a:ssrc:67890 cname:EocUG1f0fcg/yvY7	[RFC5576]

Table 39: 5.4.3 SDP Simulcast bundle-only

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20519 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
***** Audio m=line *****	*****
m=audio 49203 UDP/TLS/RTP/SAVPF 109	[RFC4566]
c=IN IP4 24.23.204.141	[RFC4566]
a=rtcp:60065 IN IP4 24.23.204.141	[RFC3605]
a=sendrecv	[RFC3264]
a=rtpmap:109 opus/48000/2	[I-D.ietf-payload-rtp-opus]
a=maxptime:120	[RFC4566]
a=ice-ufrag:ufrag:c300d85b	[RFC5245]
a=ice-pwd:de4e99bd291c325921d5d47efbabd9a2	[RFC5245]
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	[RFC5245]
a=setup:active	[RFC4145]
a=rtcp-rsize	[RFC5506]
a=rtcp-fb:109 nack	[RFC5104]
a=extmap:1 urn:ietf:params:rtp-hdrex:sdes:audio-level	[RFC6464]
a:ssrc:54321 cname:Q/NWslao1HmN4Xa5	[RFC5576]
a=candidate:0 1 UDP 2113667327 192.168.1.7 51556 typ host	[RFC5245]
a=candidate:1 1 UDP 694302207 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 51556	[RFC5245]
a=candidate:0 2 UDP 2113667326 192.168.1.7 51558 typ host	[RFC5245]
a=candidate:1 2 UDP 1694302206 98.248.92.77 60065 typ srflx raddr 192.168.1.7 rport 51558	[RFC5245]
***** Video m=line *****	*****

m=video 0 UDP/TLS/RTP/SAVPF 98 100	***** Bob doesn't recognize bundle-only and hence the m=line is rejected implicitly due to port 0 *****
***** Video m=line *****	*****
m=video 0 UDP/TLS/RTP/SAVPF 98 100	Bob doesn't recognize bundle-only and hence the m=line is rejected implicitly due to port 0

Table 40: 5.4.3 SDP Answer

5.4.4. Legacy Interop with RTP/AVP profile

In this section, we attempt to provide session descriptions showcasing inter-operability between a WebRTC end-point and a Legacy VOIP end-point. The ideas included in here are not fully baked into the standards and might be controversial in nature. The hope here is to demonstrate a plausible SDP composition to enhance seamless inter-operability between the aforementioned communication systems.

In the scenario described below, Alice is a legacy end-point which sends [RFC3264] Offer with two sets of media descriptions per media type.

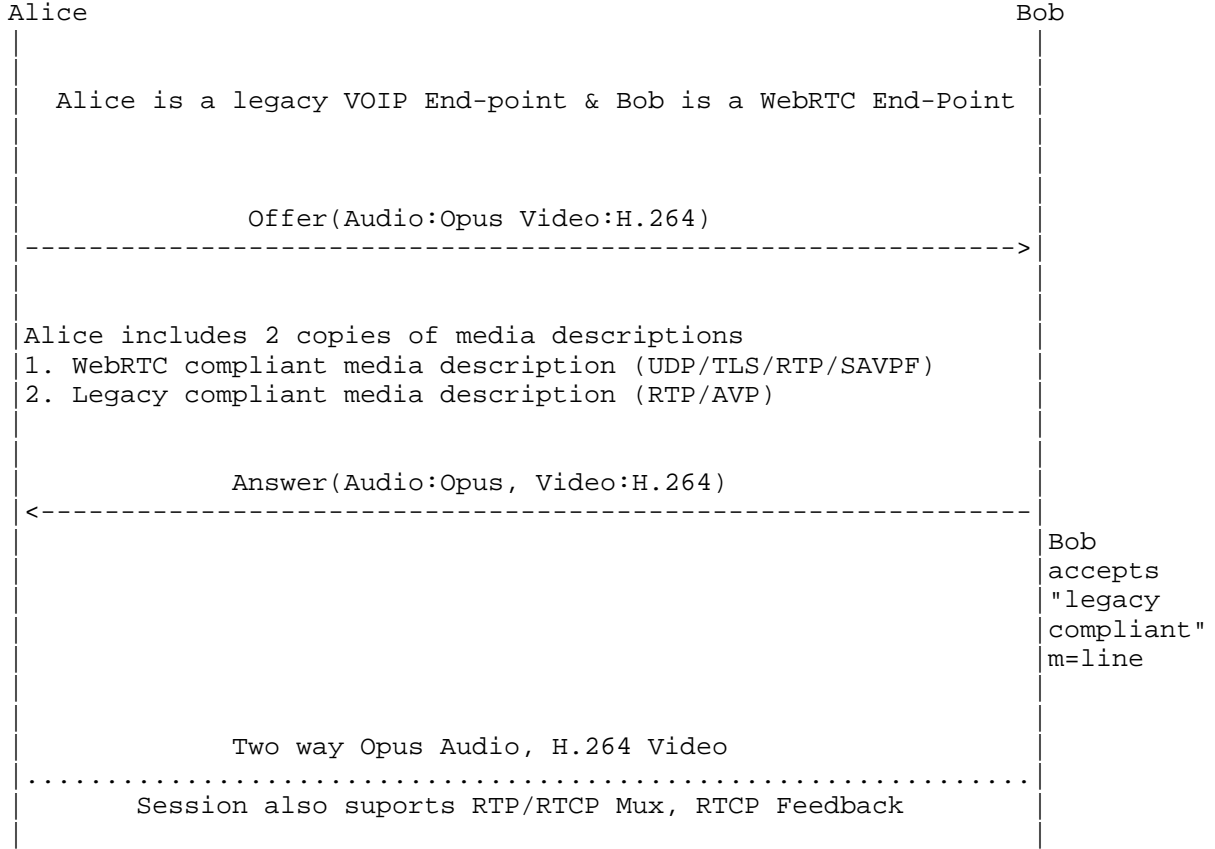
One set that corresponds to [WebRTC] compliant UDP/TLS/RTP/SAVPF based audio and video descriptions.

Another set with RTP/AVP based audio and video descriptions for the legacy Interop purposes.

Also to note, Alice includes session level DTLS information and media level RTCP feedback information as applicable to both the sets of media descriptions

On the other hand, Bob being a WebRTC end-point, recognizes accepts the media descriptions with RTP/AVP profile. The security and feedback requirements for the session are either handled by a intermediate gateway or with some combination of Alice's capabilities and the intermediate gateway.

Successful 2-Way WebRTC <-> VOIP Interop



Offer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 20518 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=ice-ufrag:074c6550	[RFC5245]
a=ice-	[RFC5245]
pwd:a28a397a4c3f31747d1ee3474af08a068	
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	[RFC5245]
a=rtcp-rsize	[RFC5506]

<pre> ***** Audio m=line ***** m=audio 54609 UDP/TLS/RTP/SAVPF 109 c=IN IP4 24.23.204.141 a=rtpmap:109 opus/48000 a=ptime:20 a=sendrecv a=rtcp-mux a=candidate:0 1 UDP 2113667327 192.168.1.4 54609 typ host a=candidate:1 1 UDP 694302207 24.23.204.141 54609 typ srflx raddr 192.168.1.4 rport 54609 a=candidate:0 2 UDP 2113667326 192.168.1.4 64678 typ host a=candidate:1 2 UDP 1694302206 24.23.204.141 64678 typ srflx raddr 192.168.1.4 rport 64678 a=rtcp-fb:109 nack ***** Audio m=line ***** m=video 62537 UDP/TLS/RTP/SAVPF 120 c=IN IP4 24.23.204.141 a=rtpmap:120 VP8/90000 a=sendrecv a=rtcp-mux a=candidate:0 1 UDP 2113667327 192.168.1.4 62537 typ host a=candidate:1 1 UDP 1694302207 24.23.204.141 62537 typ srflx raddr 192.168.1.4 rport 62537 a=candidate:0 2 2113667326 192.168.1.4 54721 typ host a=candidate:1 2 UDP 1694302206 24.23.204.141 54721 typ srflx raddr 192.168.1.4 rport 54721 a=rtcp-fb:120 nack pli a=rtcp-fb:120 ccm fir ----- ***** Audio m=line ***** m=audio 54732 RTP/AVP 109 c=IN IP4 24.23.204.141 a=fingerprint:sha-256 19:E2:1C:3B:4B:9 </pre>	<pre> ***** [RFC4566] [RFC4566] [RFC3264] [RFC5761] [RFC5245] [RFC5245] [RFC5245] [RFC5245] [RFC5104] ***** ***** [RFC4566] [RFC4566] [I-D.ietf-payload-vp8] [RFC3264] [RFC5761] [RFC5245] [RFC5245] [RFC5245] [RFC5104] [RFC5104] These set of media descriptions are for Legacy Inter-op purposes ***** ***** [RFC4566]Alice includes RTP/AVP audio stream description [RFC4566] [RFC5245] </pre>
--	--

<pre> F:81:E6:B8:5C:F4:A5:A8:D8:73:04 :BB:05 :2F:70:9F:04:A9:0E:05:E9:26:33:E8:70:8 8:A2 a=rtpmap:109 opus/48000 a=ptime:20 a=sendrecv a=rtcp-mux a=candidate:0 1 UDP 2113667327 192.168.1.4 54732 typ host a=candidate:1 1 UDP 694302207 24.23.204.141 54732 typ srflx raddr 192.168.1.4 rport 54732 a=candidate:0 2 UDP 2113667326 192.168.1.4 64678 typ host a=candidate:1 2 UDP 1694302206 24.23.204.141 64678 typ srflx raddr 192.168.1.4 rport 64678 a=rtcp-fb:109 nack ***** Video m=line ***** m=video 62445 RTP/AVP 120 c=IN IP4 24.23.204.141 a=fingerprint:sha-256 DC:B8:5F:64:1A:2 4:C2:43:F0:A1:58:D0:A1:2C:19:08 :6B:8B :F0:65:5F:78:E2:51:3B:AC:6F:F3:3F:46:1 B:35 a=rtpmap:120 VP8/90000 a=sendrecv a=rtcp-mux a=candidate:0 1 UDP 2113667327 192.168.1.4 62445 typ host a=candidate:1 1 UDP 1694302207 24.23.204.141 62537 typ srflx raddr 192.168.1.4 rport 62445 a=candidate:0 2 2113667326 192.168.1.4 54721 typ host a=candidate:1 2 UDP 1694302206 24.23.204.141 54721 typ srflx raddr 192.168.1.4 rport 54721 a=rtcp-fb:120 nack pli </pre>	<pre> [RFC3264] [RFC5761]Alice still includes RTP/RTCP Mux support [RFC5245] [RFC5245] [RFC5245] [RFC5245] [RFC5104]She adds her intent for NACK RTCP feedback support ***** ***** [RFC4566]Alice includes RTP/AVP video stream description [RFC4566] [RFC5245] [I-D.ietf-payload-vp8] [RFC3264] [RFC5761]Alice intends to perform RTP/RTCP Mux [RFC5245] [RFC5245] [RFC5245] [RFC5104] Alice indicates support for </pre>
---	--

a=rtcp-fb:120 ccm fir	Picture loss Indication and NACK RTCP feedback [RFC5104]
-----------------------	--

Table 41: 5.4.5 SDP Offer

Answer SDP Contents	RFC#/Notes
v=0	[RFC4566]
o=- 16833 0 IN IP4 0.0.0.0	[RFC4566]
s=-	[RFC4566]
t=0 0	[RFC4566]
a=ice-ufrag:c300d85b	[RFC5245]
a=ice-	[RFC5245]
pwd:de4e99bd291c325921d5d47efbabd9a2	
a=fingerprint:sha-256 BB:05:2F:70:9F:04:A9:0E:05:E9:26:33:E8:70:88:A2 :19:E2:1C:3B:4B:9F:81:E6:B8:5C:F4:A5:A8:D8:73:04	[RFC5245]
***** Audio m=line *****	*****
m=audio 49203 RTP/AVP 109	[RFC4566] Bob accepts RTP/AVP based audio stream
c=IN IP4 98.248.92.77	[RFC4566]
a=rtpmap:109 opus/48000	
a=ptime:20	
a=sendrecv	[RFC3264]
a=candidate:0 1 UDP 2113667327 192.168.1.7 49203 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 98.248.92.77 49203 typ srflx raddr 192.168.1.7 rport 49203	[RFC5245]
a=candidate:0 2 UDP 2113667326 192.168.1.7 60065 typ host	[RFC5245]
a=candidate:1 2 UDP 1694302206 98.248.92.77 60065 typ srflx raddr 192.168.1.7 rport 60065	[RFC5245]
***** Video m=line *****	*****
m=video 63130 RTP/SAVP 120	[RFC4566] Bob accepts RTP/AVP based video stream
c=IN IP4 98.248.92.771	[RFC4566]
a=rtpmap:120 VP8/90000	[I-D.ietf-payload-vp8]

a=sendrecv	[RFC3264]
a=candidate:0 1 UDP 2113667327 192.168.1.7 63130 typ host	[RFC5245]
a=candidate:1 1 UDP 1694302207 98.248.92.77 63130 typ srflx raddr 192.168.1.7 rport 63130	[RFC5245]
a=candidate:0 2 UDP 2113667326 192.168.1.7 56607 typ host	[RFC5245]
a=candidate:1 2 UDP 1694302206 98.248.92.77 56607 typ srflx raddr 192.168.1.7 rport 56607	[RFC5245]

Table 42: 5.4.5 SDP Answer

6. IANA Considerations

This document requires no actions from IANA.

7. Acknowledgments

We would like to thank Justin Uberti, Chris Flo, Paul Kyzivat for their detailed review and inputs.

8. Change Log

[RFC EDITOR NOTE: Please remove this section when publishing]

Changes from draft-ietf-rtcweb-sdp-02

- o Version increment to avoid expiry

Changes from draft-ietf-rtcweb-sdp-01

- o Complete face-lift
- o Added visual markers around m=lines to indicate their type, added spacing between tables for aiding readers
- o Updated table names to indicate offer vs answer
- o Attempted to align to latest versions of SCTP, BUNDLE, MSID drafts
- o Added mid header extensions to all the lines
- o Harmonized BUNDLE semantics and conventions updated.

Changes from draft-ietf-rtcweb-sdp-00

- o Updated Simulcast/FEC/RTX examples to use RID framework
- o Fixed BUNDLE references for a=bundle-only

Changes from draft-nandakumar-rtcweb-sdp-08

- o Fixed typos
- o Moved to a WG version

Changes from draft-nandakumar-rtcweb-sdp-06 and draft-nandakumar-rtcweb-sdp-07

- o Added clarification on Call-Flow diagram usage
- o More cleanups

Changes from draft-nandakumar-rtcweb-sdp-05

- o Added Ascii chart for all the SDP Examples
- o Improved text and updated SDP Examples for Simulcast and FEC
- o Fixed MediaStream ID Semantics SDP Errors

Changes from draft-nandakumar-rtcweb-sdp-04

- o Interim version of the draft to avert expiry
- o Corrected placement of c= line as per RFC4566
- o Updated simulcast SDP to reflect draft-westerlund-avtcore-rtp-simulcast-04

Changes from draft-nandakumar-rtcweb-sdp-03

- o Aligned more closely with JSEP version -05
- o Added Conventions to help readability
- o Add more examples to clarify BUNDLE use-cases

Changes from draft-nandakumar-rtcweb-sdp-02

- o Major refactoring was done to group the examples in to categories

- o SDP was updated through out to reflect JSEP-04 style of defining attributes per m=line than at the session level.
- o Added 8 new examples.
- o Updated references for Trickle, Unified Plan
- o Add section to explain the syntax conventions followed in the examples.

Changes from draft-nandakumar-rtcweb-sdp-01

- o Updated references to OPUS RTP Payload Specification.
- o Updated BUNDLE examples based on the latest draft-ietf-mmusic-sdp-bundle-negotiation.
- o Added examples for multiple audio and video flows based on Unified Plan.
- o Added new examples for RTX and FEC streams
- o Updated Simulcast and SVC examples

Changes from draft-nandakumar-rtcweb-sdp-00

- o Fixed editorial comments on the mailing list.
- o Updated Data-channel SDP information based on draft-ietf-mmusic-sctp-sdp.
- o Updated BUNDLE examples based on draft-ietf-mmusic-sdp-bundle-negotiation.
- o Added examples for few more BUNDLE variants
- o Added new examples for Simulcast and SVC

9. Informative References

- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, DOI 10.17487/RFC3264, June 2002, <<http://www.rfc-editor.org/info/rfc3264>>.

- [RFC4145] Yon, D. and G. Camarillo, "TCP-Based Media Transport in the Session Description Protocol (SDP)", RFC 4145, DOI 10.17487/RFC4145, September 2005, <<http://www.rfc-editor.org/info/rfc4145>>.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, DOI 10.17487/RFC4566, July 2006, <<http://www.rfc-editor.org/info/rfc4566>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, DOI 10.17487/RFC5245, April 2010, <<http://www.rfc-editor.org/info/rfc5245>>.
- [RFC5506] Johansson, I. and M. Westerlund, "Support for Reduced-Size Real-Time Transport Control Protocol (RTCP): Opportunities and Consequences", RFC 5506, DOI 10.17487/RFC5506, April 2009, <<http://www.rfc-editor.org/info/rfc5506>>.
- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, DOI 10.17487/RFC3551, July 2003, <<http://www.rfc-editor.org/info/rfc3551>>.
- [RFC3952] Duric, A. and S. Andersen, "Real-time Transport Protocol (RTP) Payload Format for internet Low Bit Rate Codec (iLBC) Speech", RFC 3952, DOI 10.17487/RFC3952, December 2004, <<http://www.rfc-editor.org/info/rfc3952>>.
- [RFC4796] Hautakorpi, J. and G. Camarillo, "The Session Description Protocol (SDP) Content Attribute", RFC 4796, DOI 10.17487/RFC4796, February 2007, <<http://www.rfc-editor.org/info/rfc4796>>.
- [RFC5761] Perkins, C. and M. Westerlund, "Multiplexing RTP Data and Control Packets on a Single Port", RFC 5761, DOI 10.17487/RFC5761, April 2010, <<http://www.rfc-editor.org/info/rfc5761>>.

- [RFC3556] Casner, S., "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", RFC 3556, DOI 10.17487/RFC3556, July 2003, <<http://www.rfc-editor.org/info/rfc3556>>.
- [RFC5104] Wenger, S., Chandra, U., Westerlund, M., and B. Burman, "Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)", RFC 5104, DOI 10.17487/RFC5104, February 2008, <<http://www.rfc-editor.org/info/rfc5104>>.
- [RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R. Hakenberg, "RTP Retransmission Payload Format", RFC 4588, DOI 10.17487/RFC4588, July 2006, <<http://www.rfc-editor.org/info/rfc4588>>.
- [RFC5956] Begen, A., "Forward Error Correction Grouping Semantics in the Session Description Protocol", RFC 5956, DOI 10.17487/RFC5956, September 2010, <<http://www.rfc-editor.org/info/rfc5956>>.
- [RFC5888] Camarillo, G. and H. Schulzrinne, "The Session Description Protocol (SDP) Grouping Framework", RFC 5888, DOI 10.17487/RFC5888, June 2010, <<http://www.rfc-editor.org/info/rfc5888>>.
- [RFC6236] Johansson, I. and K. Jung, "Negotiation of Generic Image Attributes in the Session Description Protocol (SDP)", RFC 6236, DOI 10.17487/RFC6236, May 2011, <<http://www.rfc-editor.org/info/rfc6236>>.
- [RFC3984] Wenger, S., Hannuksela, M., Stockhammer, T., Westerlund, M., and D. Singer, "RTP Payload Format for H.264 Video", RFC 3984, DOI 10.17487/RFC3984, February 2005, <<http://www.rfc-editor.org/info/rfc3984>>.
- [RFC5583] Schierl, T. and S. Wenger, "Signaling Media Decoding Dependency in the Session Description Protocol (SDP)", RFC 5583, DOI 10.17487/RFC5583, July 2009, <<http://www.rfc-editor.org/info/rfc5583>>.
- [RFC5576] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", RFC 5576, DOI 10.17487/RFC5576, June 2009, <<http://www.rfc-editor.org/info/rfc5576>>.

- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<http://www.rfc-editor.org/info/rfc3550>>.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC2326] Schulzrinne, H., Rao, A., and R. Lanphier, "Real Time Streaming Protocol (RTSP)", RFC 2326, DOI 10.17487/RFC2326, April 1998, <<http://www.rfc-editor.org/info/rfc2326>>.
- [RFC3605] Huitema, C., "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", RFC 3605, DOI 10.17487/RFC3605, October 2003, <<http://www.rfc-editor.org/info/rfc3605>>.
- [RFC2833] Schulzrinne, H. and S. Petrack, "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals", RFC 2833, DOI 10.17487/RFC2833, May 2000, <<http://www.rfc-editor.org/info/rfc2833>>.
- [RFC6464] Lennox, J., Ed., Ivov, E., and E. Marocco, "A Real-time Transport Protocol (RTP) Header Extension for Client-to-Mixer Audio Level Indication", RFC 6464, DOI 10.17487/RFC6464, December 2011, <<http://www.rfc-editor.org/info/rfc6464>>.
- [RFC6465] Ivov, E., Ed., Marocco, E., Ed., and J. Lennox, "A Real-time Transport Protocol (RTP) Header Extension for Mixer-to-Client Audio Level Indication", RFC 6465, DOI 10.17487/RFC6465, December 2011, <<http://www.rfc-editor.org/info/rfc6465>>.
- [RFC7022] Begen, A., Perkins, C., Wing, D., and E. Rescorla, "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMEs)", RFC 7022, DOI 10.17487/RFC7022, September 2013, <<http://www.rfc-editor.org/info/rfc7022>>.
- [I-D.ietf-mmusic-sdp-bundle-negotiation] Holmberg, C., Alvestrand, H., and C. Jennings, "Negotiating Media Multiplexing Using the Session Description Protocol (SDP)", draft-ietf-mmusic-sdp-bundle-negotiation-24 (work in progress), January 2016.

- [I-D.ietf-mmusic-sdp-simulcast]
Burman, B., Westerlund, M., Nandakumar, S., and M. Zanaty, "Using Simulcast in SDP and RTP Sessions", draft-ietf-mmusic-sdp-simulcast-03 (work in progress), October 2015.
- [I-D.ietf-mmusic-rid]
Thatcher, P., Zanaty, M., Nandakumar, S., Burman, B., Roach, A., and B. Campen, "RTP Payload Format Constraints", draft-ietf-mmusic-rid-04 (work in progress), February 2016.
- [I-D.ietf-payload-rtp-opus]
Spittka, J., Vos, K., and J. Valin, "RTP Payload Format for the Opus Speech and Audio Codec", draft-ietf-payload-rtp-opus-11 (work in progress), April 2015.
- [I-D.ietf-payload-vp8]
Westin, P., Lundin, H., Glover, M., Uberti, J., and F. Galligan, "RTP Payload Format for VP8 Video", draft-ietf-payload-vp8-17 (work in progress), September 2015.
- [I-D.ietf-rtcweb-jsep]
Uberti, J., Jennings, C., and E. Rescorla, "Javascript Session Establishment Protocol", draft-ietf-rtcweb-jsep-12 (work in progress), October 2015.
- [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-02 (work in progress), January 2015.
- [I-D.ietf-mmusic-msid]
Alvestrand, H., "WebRTC MediaStream Identification in the Session Description Protocol", draft-ietf-mmusic-msid-11 (work in progress), October 2015.
- [I-D.ietf-mmusic-sctp-sdp]
Holmberg, C., Loreto, S., and G. Camarillo, "Stream Control Transmission Protocol (SCTP)-Based Media Transport in the Session Description Protocol (SDP)", draft-ietf-mmusic-sctp-sdp-15 (work in progress), September 2015.
- [I-D.ietf-rtcweb-data-channel]
Jesup, R., Loreto, S., and M. Tuexen, "WebRTC Data Channels", draft-ietf-rtcweb-data-channel-13 (work in progress), January 2015.

- [I-D.ietf-payload-flexible-fec-scheme]
Singh, V., Begen, A., Zanaty, M., and G. Mandyam, "RTP Payload Format for Flexible Forward Error Correction (FEC)", draft-ietf-payload-flexible-fec-scheme-01 (work in progress), October 2015.
- [I-D.ietf-mmusic-mux-exclusive]
Holmberg, C., "Indicating Exclusive Support of RTP/RTCP Multiplexing using SDP", draft-ietf-mmusic-mux-exclusive-08 (work in progress), June 2016.
- [WebRTC] W3C, "WebRTC 1.0: Real-time Communication Between Browsers",
<<http://dev.w3.org/2011/webrtc/editor/webrtc.html>> , .

Authors' Addresses

Suhas Nandakumar
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: snandaku@cisco.com

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

Phone: +1 408 421-9990
Email: fluffy@cisco.com