

Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 13, 2021

S. Murillo
A. Gouaillard
CoSMo Software
September 09, 2020

WebRTC-HTTP ingestion protocol (WHIP)
draft-murillo-whip-00

Abstract

While WebRTC has been very successful in a wide range of scenarios, its adoption in the broadcasting/streaming industry is lagging behind. Currently there is no standard protocol (like SIP or RTSP) designed for ingesting media in a streaming service, and content providers still rely heavily on protocols like RTMP for it.

These protocols are much older than webrtc and lack by default some important security and resilience features provided by webrtc with minimal delay.

The media codecs used in older protocols do not always match those being used in WebRTC, mandating transcoding on the ingest node, introducing delay and degrading media quality. This transcoding step is always present in traditional streaming to support e.g. ABR, and comes at no cost. However webrtc implements client-side ABR, also called Network-Aware Encoding by e.g. Huavision, by means of simulcast and SVC codecs, which otherwise alleviate the need for server-side transcoding. Content protection and Privacy Enhancement can be achieved with End-to-End Encryption, which preclude any server-side media processing.

This document proposes a simple HTTP based protocol that will allow WebRTC endpoints to ingest content into streaming services and/or CDNs to fill this gap and facilitate deployment.

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 <https://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 March 13, 2021.

Copyright Notice

Copyright (c) 2020 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 (<https://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 | 3 |
| 3. | Overview | 3 |
| 4. | Protocol Operation | 4 |
| 4.1. | ICE and NAT support | 4 |
| 4.2. | Webrtc constrains | 5 |
| 4.3. | Load balancing and redirections | 5 |
| 4.4. | Authentication and authorization | 5 |
| 4.5. | Simulcast and scalable video coding | 5 |
| 5. | Security Considerations | 6 |
| 6. | IANA Considerations | 6 |
| 7. | Acknowledgements | 6 |
| 8. | Normative References | 6 |
| | Authors' Addresses | 6 |

[1.](#) Introduction

WebRTC intentionally does not specify a signaling transport protocol at application level, while RTCWEB standardized the signalling protocol itself (JSEP, SDP O/A) and everything that was going over the wire (media, codec, encryption, ...). This flexibility has allowed for implementing a wide range of services. However, those services are typically standalone silos which don't require interoperability with other services or leverage the existence of tools that can communicate with them.

In the broadcasting/streaming world, the usage of hardware encoders that would make it very simple to plug in (SDI) cables carrying raw media, encoding it in place, and pushing it to any streaming service or CDN ingest is ubiquitous. Having to implement a custom signalling transport protocol for each different webrtc services has hindered adoption.

While some standard signalling protocols are available that can be integrated with WebRTC, like SIP or XMPP, they are not designed to be used in broadcasting/streaming services, and there also is no sign of adoption in that industry. RTSP, which is based on RTP and maybe the closest in terms of features to webrtc, is not compatible with WebRTC SDP offer/answer model.

In the specific case of ingest into a platform, some assumption can be made about the server-side which simplifies the webrtc compliance burden, as detailed in webrtc-gateway document.

<https://tools.ietf.org/html/draft-ietf-rtcweb-gateways-02>

This document propose a simple protocol for supporting WebRTC as ingest method which is:

- Easy to implement,
- As easy to use as current RTMP URI.
- Fully compliant with Webrtc and RTCWEB specs.
- Allow for both ingest in traditionnal media platforms for extention and ingest in webrtc end-to-end platform for lowest possible latency.
- Lowers the requirements on both hardware encoders and broadcasting services to support webrtc.
- Usable both in web browsers and in native encoders.

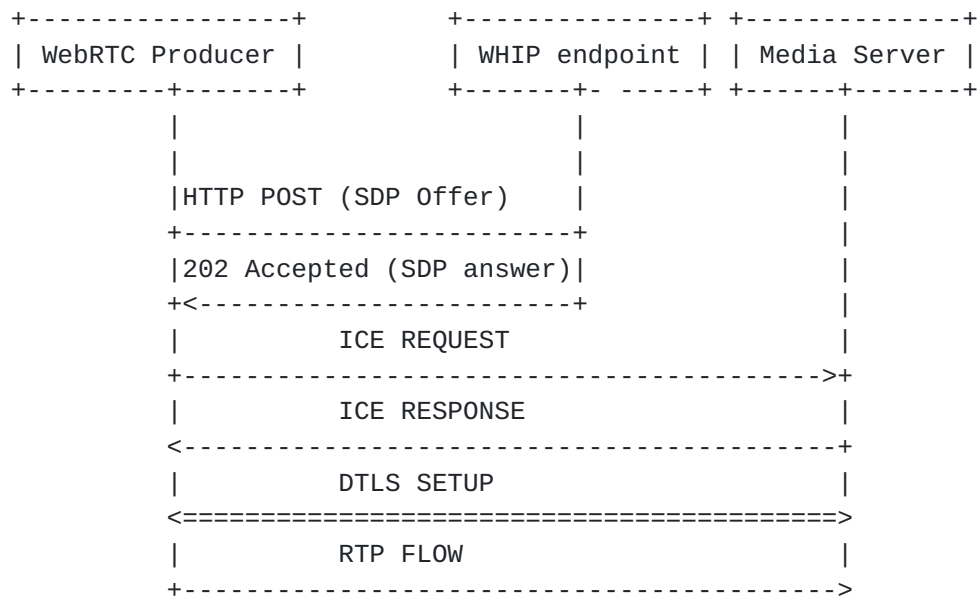
2. Terminology

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. Overview

The WebRTC-HTTP ingest protocol (WHIP) uses an HTTP POST request to perform a single shot SDP offer/answer so an ICE/DTLS session can be established between the encoder/media producer and the broadcasting ingestion endpoint.

Once the ICE/DTLS session is set up, the media will flow unidirectionally from the encoder/media producer broadcasting ingestion endpoint. In order to reduce complexity, no SDP renegotiation is supported, so no tracks or streams can be added or removed once the initial SDP O/A over HTTP is completed.



WHIP session setup

4. Protocol Operation

In order to setup an ingestion session, the WebRTC encoder/media producer will generate an SDP offer according the the JSEP rules and do an HTTP POST request to the WHIP endpoint configured URL.

The HTTP POST request will have a content type of application/sdp and contain the SDP offer as body. The WHIP ingestion endpoint will generate an SDP answer and return it on a 202 Accepted response with content type of application/sdp and the SDP answer as body.

SDP offer SHOULD use the sendonly attribute and the SDP answer MUST use the recvonly attribute.

Once session is setup ICE consent freshness [[RFC7675](#)] will be used to detect abrupt disconnection and DTLS teardown for session termination by either side.

4.1. ICE and NAT support

In order to simplify the protocol, there is no support of exchanging gathered tickle ICE candidates one the SDP offer or answer is sent. So in order to support encoders/media producers behind NAT, the WHIP media server MUST be publicly accessible.

The initial offer by the encoder/media producer MAY be sent after the full ICE gathering is complete containing the full list of ICE candidates, or only contain local candidates or even an empty list of

candidates. The WHIP endpoint SDP answer SHALL contain the full list of ICE candidates publicly accessible of the media server. The media server MAY use ICE lite, while the encoder/media producer MUST implement full ICE.

If the Encoder/Media producer gathers additional candidates (via STUN/TURN) after the SDP offer is sent, it will send directly a STUN request to the ICE candidates received from the media server as per [I-D.[draft-ietf-ice-trickle-21](#)].

4.2. Webrtc constrains

In order to reduce the complexity of implementing WHIP in both encoders and media servers, some restrictions regarding WebRTC usage are made.

SDP bundle SHALL be used by both the encoder/media producer and the media server. The SDP offer created by the encoder/media producer MUST include the bundle-only attribute in all m-lines as per [I-D.[draft-ietf-mmusic-sdp-bundle-negotiation-54](#)]. Also, RTCP muxing SHALL be supported by the both the encoder/media producer and the media server.

4.3. Load balancing and redirections

Encoders/media MAY not be colocated on the same server so it is possible to load balance incoming request to different media server. Encoders/media producers SHALL support HTTP redirection via 307 Temporary Redirect response code.

In case of high load, the WHIP endpoints may return a 503 (Service Unavailable) status code indicating that the server is currently unable to handle the request due to a temporary overload or scheduled maintenance, which will likely be alleviated after some delay. The server MAY send a Retry-After header field indicating the minimum time that the user agent is asked to wait before issuing the redirected request.

4.4. Authentication and authorization

Authntentication and authorization is supported by the Authorization HTTP header with a bearear token as per [[RFC6750](#)].

4.5. Simulcast and scalable video coding

Both simulcast and scalable video coding (including K-SVC modes) MAY be supported by both media servers and encoders/media producers.

5. Security Considerations

HTTPS SHALL be used in order to preserve WebRTC security model.

6. IANA Considerations

7. Acknowledgements

8. Normative References

[I-D.[draft-ietf-ice-trickle-21](#)]

Ivov, E., Rescorla, E., Uberti, J., and P. Saint-Andre, "Trickle ICE: Incremental Provisioning of Candidates for the Interactive Connectivity Establishment (ICE) Protocol", [draft-ietf-ice-trickle-21](#) (work in progress), April 2018.

[I-D.[draft-ietf-mmusic-sdp-bundle-negotiation-54](#)]

Holmberg, C., Alvestrand, H., and C. Jennings, "Negotiating Media Multiplexing Using the Session Description Protocol (SDP)", [draft-ietf-mmusic-sdp-bundle-negotiation-54](#) (work in progress), December 2018.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", [RFC 6750](#), DOI 10.17487/RFC6750, October 2012, <<https://www.rfc-editor.org/info/rfc6750>>.

[RFC7675] Perumal, M., Wing, D., Ravindranath, R., Reddy, T., and M. Thomson, "Session Traversal Utilities for NAT (STUN) Usage for Consent Freshness", [RFC 7675](#), DOI 10.17487/RFC7675, October 2015, <<https://www.rfc-editor.org/info/rfc7675>>.

Authors' Addresses

Sergio Garcia Murillo
CoSMo Software

Email: sergio.garcia.murillo@cosmosoftware.io

Alexandre Gouaillard
CoSMo Software

Email: alex.gouaillard@cosmosoftware.io