

RTCWeb Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 24, 2014

R. Jesup  
Mozilla  
S. Loreto  
Ericsson  
M. Tuexen  
Muenster Univ. of Appl. Sciences  
October 21, 2013

RTCWeb Data Channels  
draft-ietf-rtcweb-data-channel-06.txt

## Abstract

The Real-Time Communication in WEB-browsers (RTCWeb) working group is charged to provide protocol support for direct interactive rich communication using audio, video, and data between two peers' web-browsers. This document specifies the non-media data transport aspects of the RTCWeb framework. It provides an architectural overview of how the Stream Control Transmission Protocol (SCTP) is used in the RTCWeb context as a generic transport service allowing WEB-browsers to exchange generic data from peer to peer.

## 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, 2014.

## Copyright Notice

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

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction . . . . .	2
2. Conventions . . . . .	3
3. Use Cases . . . . .	3
3.1. Use Cases for Unreliable Data Channels . . . . .	3
3.2. Use Cases for Reliable Data Channels . . . . .	3
4. Requirements . . . . .	4
5. SCTP over DTLS over UDP Considerations . . . . .	5
6. The Usage of SCTP in the RTCWeb Context . . . . .	8
6.1. SCTP Protocol Considerations . . . . .	8
6.2. Association Setup . . . . .	9
6.3. SCTP Streams . . . . .	9
6.4. Channel Definition . . . . .	10
6.5. Opening a Channel . . . . .	10
6.6. Transferring User Data on a Channel . . . . .	10
6.7. Closing a Channel . . . . .	11
7. Security Considerations . . . . .	11
8. IANA Considerations . . . . .	12
9. Acknowledgments . . . . .	12
10. References . . . . .	12
10.1. Normative References . . . . .	12
10.2. Informative References . . . . .	14
Authors' Addresses . . . . .	14

## 1. Introduction

Non-media data types in the context of RTCWeb are handled by using SCTP [RFC4960] encapsulated in DTLS [RFC6347].




Figure 1: Basic stack diagram

The encapsulation of SCTP over DTLS (see [I-D.ietf-tsvwg-sctp-dtls-encaps]) over ICE/UDP (see [RFC5245]) provides a NAT traversal solution together with confidentiality, source authentication, and integrity protected transfers. This data transport service operates in parallel to the media transports, and all of them can eventually share a single transport-layer port number.

SCTP as specified in [RFC4960] with the partial reliability extension defined in [RFC3758] provides multiple streams natively with reliable, and partially-reliable delivery modes for user messages. Using the reconfiguration extension defined in [RFC6525] allows to increase the number of streams during the lifetime of an SCTP association and to reset individual SCTP streams.

The remainder of this document is organized as follows: Section 3 and Section 4 provide use cases and requirements for both unreliable and reliable peer to peer data channels; Section 5 arguments SCTP over DTLS over UDP; Section 6 provides the specification of how SCTP should be used by the RTCWeb protocol framework for transporting non-media data between WEB-browsers.

## 2. Conventions

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. Use Cases

This section defined use cases specific to data channels. For general use cases see [I-D.ietf-rtcweb-use-cases-and-requirements].

### 3.1. Use Cases for Unreliable Data Channels

U-C 1: A real-time game where position and object state information is sent via one or more unreliable data channels. Note that at any time there may be no media channels, or all media channels may be inactive, and that there may also be reliable data channels in use.

U-C 2: Providing non-critical information to a user about the reason for a state update in a video chat or conference, such as mute state.

### 3.2. Use Cases for Reliable Data Channels

U-C 3: A real-time game where critical state information needs to be transferred, such as control information. Such a game may have no media channels, or they may be inactive at any given time, or may only be added due to in-game actions.

U-C 4: Non-realtime file transfers between people chatting. Note that this may involve a large number of files to transfer sequentially or in parallel, such as when sharing a folder of images or a directory of files.

U-C 5: Realtime text chat during an audio and/or video call with an individual or with multiple people in a conference.

U-C 6: Renegotiation of the set of media streams in the PeerConnection.

U-C 7: Proxy browsing, where a browser uses data channels of a PeerConnection to send and receive HTTP/HTTPS requests and data, for example to avoid local internet filtering or monitoring.

#### 4. Requirements

This section lists the requirements for P2P data channels between two browsers.

- Req. 1: Multiple simultaneous data channels MUST be supported. Note that there may 0 or more media streams in parallel with the data channels, and the number and state (active/inactive) of the media streams may change at any time.
- Req. 2: Both reliable and unreliable data channels MUST be supported.
- Req. 3: Data channels MUST be congestion controlled; either individually, as a class, or in conjunction with the media streams, to ensure that data channels don't cause congestion problems for the media streams, and that the RTCWeb PeerConnection as a whole is fair with competing traffic such as TCP.
- Req. 4: The application SHOULD be able to provide guidance as to the relative priority of each data channel relative to each other, and relative to the media streams. [ TBD: how this is encoded and what the impact of this is. ] This will interact with the congestion control algorithms.

- Req. 5: Data channels MUST be secured; allowing for confidentiality, integrity and source authentication. See [I-D.ietf-rtcweb-security] and [I-D.ietf-rtcweb-security-arch] for detailed info.
- Req. 6: Data channels MUST provide message fragmentation support such that IP-layer fragmentation can be avoided no matter how large a message the JavaScript application passes to be sent. It also MUST ensure that large data channel transfers don't unduly delay traffic on other data channels.
- Req. 7: The data channel transport protocol MUST NOT encode local IP addresses inside its protocol fields; doing so reveals potentially private information, and leads to failure if the address is depended upon.
- Req. 8: The data channel transport protocol SHOULD support unbounded-length "messages" (i.e., a virtual socket stream) at the application layer, for such things as image-file-transfer; Implementations might enforce a reasonable message size limit.
- Req. 9: The data channel transport protocol SHOULD avoid IP fragmentation. It MUST support PMTU (Path MTU) discovery and MUST NOT rely on ICMP or ICMPv6 being generated or being passed back, especially for PMTU discovery.
- Req. 10: It MUST be possible to implement the protocol stack in the user application space.

## 5. SCTP over DTLS over UDP Considerations

The important features of SCTP in the RTCWeb context are:

- o Usage of a TCP-friendly congestion control.
- o The congestion control is modifiable for integration with media stream congestion control.
- o Support of multiple unidirectional streams, each providing its own notion of ordered message delivery.
- o Support of ordered and out-of-order message delivery.
- o Supporting arbitrary large user message by providing fragmentation and reassembly.
- o Support of PMTU-discovery.

- o Support of reliable or partially reliable message transport.

SCTP multihoming will not be used in RTCWeb. The SCTP layer will simply act as if it were running on a single-homed host, since that is the abstraction that the lower layer (a connection oriented, unreliable datagram service) exposes.

The encapsulation of SCTP over DTLS defined in [I-D.ietf-tsvwg-sctp-dtls-encaps] provides confidentiality, source authenticated, and integrity protected transfers. Using DTLS over UDP in combination with ICE enables NAT traversal in IPv4 based networks. SCTP as specified in [RFC4960] MUST be used in combination with the extension defined in [RFC3758] and provides the following interesting features for transporting non-media data between browsers:

- o Support of multiple unidirectional streams.
- o Ordered and unordered delivery of user messages.
- o Reliable and partial-reliable transport of user messages.

Each SCTP user message contains a so called Payload Protocol Identifier (PPID) that is passed to SCTP by its upper layer and sent to its peer. This value can be used to multiplex multiple protocols over a single SCTP association. The sender provides for each protocol a specific PPID and the receiver can demultiplex the messages based on the received PPID.

The encapsulation of SCTP over DTLS, together with the SCTP features listed above satisfies all the requirements listed in Section 4.

The layering of protocols for WebRTC is shown in the following Figure 2.




Figure 2: WebRTC protocol layers

This stack (especially in contrast to DTLS over SCTP [RFC6083] in combination with SCTP over UDP [RFC6951]) has been chosen because it

- o supports the transmission of arbitrary large user messages.
- o shares the DTLS connection with the media channels.
- o provides privacy for the SCTP control information.

Considering the protocol stack of Figure 2 the usage of DTLS over UDP is specified in [RFC6347], while the usage of SCTP on top of DTLS is specified in [I-D.ietf-tsvwg-sctp-dtls-encaps].

Since DTLS is typically implemented in user-land, the SCTP stack also needs to be a user-land stack.

When using DTLS as the lower layer, only single homed SCTP associations MUST be used, since DTLS does not expose any address management to its upper layer. The ICE/UDP layer can handle IP address changes during a session without needing to notify the DTLS and SCTP layers, though it would be advantageous to retest Path MTU on an IP address change.

DTLS implementations used for this stack SHOULD support controlling fields of the IP layer like the Don't Fragment (DF)-bit in case of IPv4 and the Differentiated Services Code Point (DSCP) field required for supporting [I-D.ietf-rtcweb-qos]. Being able to set the (DF)-bit in case of IPv4 is required for performing path MTU discovery. The DTLS implementation SHOULD also support sending user messages exceeding the Path MTU.

Incoming ICMP or ICMPv6 messages can't be processed by the SCTP layer, since there is no way to identify the corresponding

association. Therefore SCTP MUST support performing Path MTU discovery without relying on ICMP or ICMPv6 as specified in [RFC4821] using probing messages specified in [RFC4820]. The initial Path MTU at the IP layer MUST NOT exceed 1200 bytes for IPv4 and 1280 for IPv6. Taking an overhead of 20 bytes for IPv4, 40 bytes for IPv6, 8 bytes for UDP, 13 + X for DTLS and 28 bytes for SCTP into account, this results in an SCTP payload of 1131 - X when IPv4 is used and 1192 - X bytes when IPv6 is used.

In general, the lower layer interface of an SCTP implementation SHOULD be adapted to address the differences between IPv4 and IPv6 (being connection-less) or DTLS (being connection-oriented).

When protocol stack of Figure 2 is used, DTLS protects the complete SCTP packet, so it provides confidentiality, integrity and source authentication of the complete SCTP packet.

This protocol stack MUST support the usage of multiple SCTP streams. A user message can be sent ordered or unordered and with partial or full reliability. The partial reliability extension MUST support policies to limit

- o the transmission and retransmission by time.
- o the number of retransmissions.

Limiting the number of retransmissions to zero combined with unordered delivery provides a UDP-like service where each user message is sent exactly once and delivered in the order received.

SCTP provides congestion control on a per-association base. This means that all SCTP streams within a single SCTP association share the same congestion window. Traffic not being sent over SCTP is not covered by the SCTP congestion control. Using a congestion control different from the standard one might improve the impact on the parallel SRTP media streams. Since SCTP does not support the negotiation of a congestion control algorithm, alternate congestion controls SHOULD only require a different sender side behavior using existing information carried in the association.

## 6. The Usage of SCTP in the RTCWeb Context

### 6.1. SCTP Protocol Considerations

The DTLS encapsulation of SCTP packets as described in [I-D.ietf-tsvwg-sctp-dtls-encaps] MUST be used.

The following SCTP protocol extensions are required:



- o The stream reset extension defined in [RFC6525] MUST be supported. It is used for closing channels.
- o The dynamic address reconfiguration extension defined in [RFC5061] MUST be used to signal the support of the stream reset extension defined in [RFC6525], other features of [RFC5061] MUST NOT be used.
- o The partial reliability extension defined in [RFC3758] MUST be supported. In addition to the timed reliability PR-SCTP policy defined in [RFC3758], the limited retransmission policy defined in [I-D.tuexen-tsvwg-sctp-prpolicies] MUST be supported.

Once support for message interleaving as currently being discussed in [I-D.stewart-tsvwg-sctp-ndata] is available, it SHOULD be supported.

## 6.2. Association Setup

The SCTP association will be set up when the two endpoints of the WebRTC PeerConnection agree on opening it, as negotiated by JSEP (typically an exchange of SDP) [I-D.ietf-rtcweb-jsep]. Additionally, the negotiation SHOULD include some type of congestion control selection. It will use the DTLS connection selected via SDP; typically this will be shared via BUNDLE or equivalent with DTLS connections used to key the DTLS-SRTP media streams.

The application SHOULD indicate the initial number of streams required when opening the association, and if no value is supplied, the implementation SHOULD provide an appropriate default. If more simultaneous streams are needed, [RFC6525] allows adding additional (but not removing) streams to an existing association. Note there can be up to 65536 SCTP streams per SCTP association in each direction.

## 6.3. SCTP Streams

SCTP defines a stream as a unidirectional logical channel existing within an SCTP association one to another SCTP endpoint. The streams are used to provide the notion of in-sequence delivery and for multiplexing. Each user message is sent on a particular stream, either order or unordered. Ordering is preserved only for ordered messages sent on the same stream.

#### 6.4. Channel Definition

The W3C has consensus on defining the application API for WebRTC DataChannels to be bidirectional. They also consider the notions of in-sequence, out-of-sequence, reliable and unreliable as properties of Channels. One strong wish is for the application-level API to be close to the API for WebSockets, which implies bidirectional streams of data and waiting for onopen to fire before sending, a textual label used to identify the meaning of the stream, among other things. Each data channel also has a priority. These priorities MUST NOT be strict priorities.

The realization of a bidirectional Data Channel is a pair of one incoming stream and one outgoing SCTP stream.

Note that there's no requirement for the SCTP streams used to create a bidirectional channel have the same number in each direction. How stream values are selected is protocol and implementation dependent.

#### 6.5. Opening a Channel

Data channels can be opened by using internal or external negotiation. The details are out of scope of this document.

A simple protocol for internal negotiation is specified in [I-D.ietf-rtcweb-data-protocol] and MUST be supported.

When one side wants to open a channel using external negotiation, it picks a Stream. This can be based on the DTLS role (the client picks even stream identifiers, the server odd stream identifiers) or done in a different way. However, the application is responsible for avoiding collisions with existing Streams. If it attempts to re-use a Stream which is part of an existing Channel, the addition SHOULD fail. In addition to choosing a Stream, the application SHOULD also inform the protocol of the options to use for sending messages. The application MUST ensure in an application-specific manner that the other side will also inform the protocol that the selected Stream is to be used, and the parameters for sending data from that side.

#### 6.6. Transferring User Data on a Channel

All data sent on a Channel in both directions MUST be sent over the underlying Stream using the reliability defined when the Channel was opened unless the options are changed, or per-message options are specified by a higher level.

No more than one message should be put into an SCTP user message.

The SCTP Payload Protocol Identifiers (PPIDs) are used to signal the interpretation of the "Payload data". For identifying a JavaScript string the PPID "DOMString Last" MUST be used, for JavaScript binary data (ArrayBuffer or Blob) the PPID "Binary Data Last" MUST be used (see Section 8).

The SCTP base protocol specified in [RFC4960] does not support the interleaving of user messages. Therefore sending a large user message can monopolize the SCTP association. To overcome this limitation, [I-D.stewart-tsvwg-sctp-ndata] defines an extension to support message interleaving. Once such an extension is available, it SHOULD be used.

As long as message interleaving is not supported, the sending application SHOULD fragment large user messages for reliable and ordered data channels. For sending large JavaScript strings, it uses the PPID "DOMString Partial" for all but the last fragments and the PPID "DOMString Last" for the last one. For JavaScript binary data the PPIDs "Binary Data Partial" and "Binary Data Last" are used. The reassembly based on the PPID MUST be supported. For data channel which are not reliable and ordered, the sender MAY limit the maximum message size to avoid monopolization.

It is recommended that message size be kept within certain size bounds (TBD) as applications will not be able to support arbitrarily-large single messages.

The sender MAY disable the Nagle algorithm to minimize the latency.

## 6.7. Closing a Channel

Closing of a Data Channel MUST be signaled by resetting the corresponding outgoing streams [RFC6525]. Resetting a stream set the Stream Sequence Numbers (SSNs) of the stream back to 'zero' with a corresponding notification to the application layer that the reset has been performed. Streams are available to reuse after a reset has been performed.

[RFC6525] also guarantees that all the messages are delivered (or abandoned) before resetting the stream.

## 7. Security Considerations

This document does not add any additional considerations to the ones given in [I-D.ietf-rtcweb-security] and [I-D.ietf-rtcweb-security-arch].

## 8. IANA Considerations

[NOTE to RFC-Editor:

"RFCXXXX" is to be replaced by the RFC number you assign this document.

]

This document uses four already registered SCTP Payload Protocol Identifiers (PPIDs). [RFC4960] creates the registry "SCTP Payload Protocol Identifiers" from which these identifiers were assigned. IANA is requested to update the reference of these four assignments to point to this document. Therefore these four assignments should be updated to read:

Value	SCTP PPID	Reference
DOMString Last	51	[RFCXXXX]
Binary Data Partial	52	[RFCXXXX]
Binary Data Last	53	[RFCXXXX]
DOMString Partial	54	[RFCXXXX]

## 9. Acknowledgments

Many thanks for comments, ideas, and text from Harald Alvestrand, Adam Bergkvist, Cullen Jennings, Eric Rescorla, Randall Stewart, Justin Uberti, and Magnus Westerlund.

## 10. References

### 10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3758] Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP) Partial Reliability Extension", RFC 3758, May 2004.
- [RFC4820] Tuexen, M., Stewart, R., and P. Lei, "Padding Chunk and Parameter for the Stream Control Transmission Protocol (SCTP)", RFC 4820, March 2007.

- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, March 2007.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [RFC5061] Stewart, R., Xie, Q., Tuexen, M., Maruyama, S., and M. Kozuka, "Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration", RFC 5061, September 2007.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.
- [RFC6525] Stewart, R., Tuexen, M., and P. Lei, "Stream Control Transmission Protocol (SCTP) Stream Reconfiguration", RFC 6525, February 2012.
- [I-D.stewart-tsvwg-sctp-ndata]  
Stewart, R., Tuexen, M., Loreto, S., and R. Seggelmann, "A New Data Chunk for Stream Control Transmission Protocol", draft-stewart-tsvwg-sctp-ndata-03 (work in progress), October 2013.
- [I-D.ietf-rtcweb-data-protocol]  
Jesup, R., Loreto, S., and M. Tuexen, "WebRTC Data Channel Protocol", draft-ietf-rtcweb-data-protocol-00 (work in progress), July 2013.
- [I-D.ietf-tsvwg-sctp-dtls-encaps]  
Tuexen, M., Stewart, R., Jesup, R., and S. Loreto, "DTLS Encapsulation of SCTP Packets", draft-ietf-tsvwg-sctp-dtls-encaps-02 (work in progress), October 2013.
- [I-D.ietf-rtcweb-security]  
Rescorla, E., "Security Considerations for WebRTC", draft-ietf-rtcweb-security-05 (work in progress), July 2013.
- [I-D.ietf-rtcweb-security-arch]  
Rescorla, E., "WebRTC Security Architecture", draft-ietf-rtcweb-security-arch-07 (work in progress), July 2013.
- [I-D.ietf-rtcweb-jsep]

Uberti, J. and C. Jennings, "Javascript Session Establishment Protocol", draft-ietf-rtcweb-jsep-04 (work in progress), September 2013.

[I-D.ietf-rtcweb-qos]

Dhesikan, S., Druta, D., Jones, P., and J. Polk, "DSCP and other packet markings for RTCWeb QoS", draft-ietf-rtcweb-qos-00 (work in progress), October 2012.

## 10.2. Informative References

[RFC6083] Tuexen, M., Seggelmann, R., and E. Rescorla, "Datagram Transport Layer Security (DTLS) for Stream Control Transmission Protocol (SCTP)", RFC 6083, January 2011.

[RFC6951] Tuexen, M. and R. Stewart, "UDP Encapsulation of Stream Control Transmission Protocol (SCTP) Packets for End-Host to End-Host Communication", RFC 6951, May 2013.

[I-D.ietf-rtcweb-use-cases-and-requirements]

Holmberg, C., Hakansson, S., and G. Eriksson, "Web Real-Time Communication Use-cases and Requirements", draft-ietf-rtcweb-use-cases-and-requirements-12 (work in progress), October 2013.

[I-D.tuexen-tsvwg-sctp-prpolicies]

Tuexen, M., Seggelmann, R., Stewart, R., and S. Loreto, "Additional Policies for the Partial Delivery Extension of the Stream Control Transmission Protocol", draft-tuexen-tsvwg-sctp-prpolicies-03 (work in progress), October 2013.

## Authors' Addresses

Randell Jesup  
Mozilla  
US

Email: randell-ietf@jesup.org

Salvatore Loreto  
Ericsson  
Hirsalantie 11  
Jorvas 02420  
FI

Email: salvatore.loreto@ericsson.com

Michael Tuexen  
Muenster University of Applied Sciences  
Stegerwaldstrasse 39  
Steinfurt 48565  
DE

Email: [tuexen@fh-muenster.de](mailto:tuexen@fh-muenster.de)