Workgroup: Network Working Group Internet-Draft: draft-tuexen-tsvwg-sctp-init-fwd-03 Published: 10 April 2024 Intended Status: Standards Track Expires: 12 October 2024 Authors: M. Tüxen Münster Univ. of Appl. Sciences T. Völker Münster Univ. of Appl. Sciences INIT Forwarding for the Stream Control Transmission Protocol

Abstract

The Stream Control Transmission Protocol (SCTP) extension described in this document allows the support of a simple mechanism to distribute association requests between a cluster of SCTP end points providing the same service. In particular, this allows the use of anycast addresses in combination with SCTP.

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 <u>https://datatracker.ietf.org/drafts/current/</u>.

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 12 October 2024.

Copyright Notice

Copyright (c) 2024 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

(<u>https://trustee.ietf.org/license-info</u>) 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- <u>1</u>. <u>Introduction</u>
- 2. <u>Conventions</u>
- 3. <u>A New Chunk Parameter</u>
- <u>4</u>. <u>Procedures</u>
- 5. <u>Socket API Considerations</u>
- 5.1. Get or Set Accepting a Zero Checksum (SCTP_INIT_FORWARDING)
- 6. IANA Considerations
- <u>7</u>. <u>Security Considerations</u>
- <u>8.</u> <u>References</u>
 - 8.1. Normative References
 - 8.2. Informative References

<u>Authors' Addresses</u>

1. Introduction

The protocol extension described in this document allows an initiation of an SCTP association to deal with an address change of the peer during the handshake. The extension enables the peer to respond from another address than the one used as destination address in the received packet containing the INIT chunk. The SCTP Dynamic Address Reconfiguration extension described in [RFC5061] can not be used, since it does not apply to the handshake.

2. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. A New Chunk Parameter

The INIT Forwarding Chunk Parameter is defined by the following figure.

Figure 1: INIT Forwarding Chunk Parameter

Type: 16 bits (unsigned integer)

This field holds the IANA defined parameter type for the "INIT Forwarding" chunk parameter. IANA is requested to assign the value 32774 (0x8006) for this parameter type.

Length: 16 bits (unsigned integer)

This field holds the length in bytes of the chunk parameter; the value **MUST** be the length of the parameter included plus 4.

Parameter: variable length

The parameter **MUST** be one of:

*IPv4 Address parameter as specified in [<u>RFC9260</u>].

*IPv6 Address parameter as specified in [<u>RFC9260</u>].

*Padding parameter as specified in [<u>RFC4820</u>]. The length of the Padding parameter **MUST** be either the length of an IPv4 Address parameter or the length of the IPv6 Address parameter.

All transported integer numbers are in "network byte order" a.k.a., Big Endian.

The INIT Forwarding Chunk Parameter **MAY** appear in INIT and INIT ACK chunks and **MUST NOT** appear in any other chunk. If an INIT or INIT ACK chunk contains an INIT Forwarding Chunk Parameter, the INIT Forwarding Chunk Parameter **MUST** be the first optional/variablelength parameter.

If an end point not supporting the extension described in this document receives this parameter in an INIT or INIT ACK chunk, it skips this parameter and continues to process further parameters in the chunk. This behaviour is **REQUIRED** by [<u>RFC9260</u>] because the highest-order 2 bits of the Type are 10.

4. Procedures

If an end point that sends an SCTP packet containing an INIT chunk wants to allow the peer to respond from an address different from the destination address of the packet, **MUST** use the INIT Forwarding Chunk parameter as the first optional/variable-length parameter. The parameter in the INIT Forwarding Chunk parameter **MUST** be a Padding parameter. If the SCTP packet containing the INIT chunk is sent over IPV4, the length of the padding parameter **MUST** be the length of an IPv4 Address parameter, which is 8 bytes. Otherwise, if the SCTP packet containing the INIT chunk is sent over IPV6, the length of the padding parameter **MUST** be the length of an IPv6 Address parameter, which is 20 bytes. If a middlebox receives an SCTP packet containing an INIT chunk with INIT Forwarding Chunk parameter as its first optional/variablelength parameter and wants to change the destination address of the packet, it **MUST** replace the Padding parameter in the INIT Forwarding Chunk parameter with an IPv4 or IPv6 Address parameter containing the original destination address of the SCTP packet containing the INIT chunk. If the INIT Forwarding Chunk parameter does not contain a Padding parameter, but an IPv4 or IPv6 Address parameter, the INIT Forwarding Chunk parameter **MUST NOT** be modified at all.

If an end point receives an SCTP packet containing an INIT chunk and the INIT chunk contains an INIT Forwarding Chunk parameter including an Address parameter as its first optional/variable-length parameter, the end point **MUST** include this INIT Forwarding Chunk parameter as the first optional/variable-length parameter in the INIT ACK chunk, which is sent in response. If the INIT Forwarding Chunk parameter contains a Padding parameter and the end-point does not want to use the destination address, it **MUST** put an INIT Forwarding Chunk parameter containing this address in the INIT ACK chunk sent in response. If the end point wants to use the destination address in the association and the INIT Forwarding Chunk parameter contains a Padding parameter, the INIT Forwarding Chunk parameter **MUST NOT** be included in the INIT ACK chunk.

If an end point receives an SCTP packet containing an INIT ACK chunk and it cannot find the association for this packet using the IP addresses and port numbers, and the INIT ACK chunk contains an INIT Forwarding Chunk parameter as its first optional/variable-length parameter, it **SHOULD** use the IP address contained in the Address parameter of the INIT Forwarding Chunk parameter instead of the source address of the received packet for the association lookup. If an association is then found, the address in the INIT Forwarding Chunk parameter **MUST** be removed as a remote address and the source address of the packet containing the INIT ACK chunk **MUST** be added as an unconfirmed remote address.

5. Socket API Considerations

This section describes how the socket API defined in [RFC6458] needs to be extended to provide a way for the application to control the UDP encapsulation.

Please note that this section is informational only.

A socket API implementation based on [<u>RFC6458</u>] is extended by supporting one new read/write IPPROTO_SCTP level socket option.

5.1. Get or Set Accepting a Zero Checksum (SCTP_INIT_FORWARDING)

This socket option can be used to control the support of INIT forwarding. It applies only to future SCTP associations on the socket.

This option expects an integer boolean flag, where a non-zero value turns on the option, and a zero value turns off the option.

This option is off by default.

6. IANA Considerations

[NOTE to RFC-Editor: "RFCXXXX" is to be replaced by the RFC number you assign this document.]

[NOTE to RFC-Editor: The requested value for the parameter type is tentative and to be confirmed by IANA.]

This document (RFCXXXX) is the reference for the registration described in this section.

A new chunk parameter type has to be assigned by IANA. This requires an additional line in the "Chunk Parameter Types" registry for SCTP:

ID Value	Chunk Parameter Type	Reference
32774	INIT Forwarding (0x8006)	[RFCXXXX]
Table 1: New entry in "Chunk Parameter Types"		
registry		

7. Security Considerations

This document does not change the considerations given in [RFC9260].

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/ RFC2119, March 1997, <<u>https://www.rfc-editor.org/info/</u> rfc2119>.
- [RFC4820] Tuexen, M., Stewart, R., and P. Lei, "Padding Chunk and Parameter for the Stream Control Transmission Protocol

(SCTP)", RFC 4820, DOI 10.17487/RFC4820, March 2007, <https://www.rfc-editor.org/info/rfc4820>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<u>https://www.rfc-editor.org/info/rfc8174</u>>.
- [RFC9260] Stewart, R., Tüxen, M., and K. Nielsen, "Stream Control Transmission Protocol", RFC 9260, DOI 10.17487/RFC9260, June 2022, <<u>https://www.rfc-editor.org/info/rfc9260</u>>.

8.2. Informative References

- [RFC5061] Stewart, R., Xie, Q., Tuexen, M., Maruyama, S., and M. Kozuka, "Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration", RFC 5061, DOI 10.17487/ RFC5061, September 2007, <<u>https://www.rfc-editor.org/</u> <u>info/rfc5061</u>>.
- [RFC6458] Stewart, R., Tuexen, M., Poon, K., Lei, P., and V. Yasevich, "Sockets API Extensions for the Stream Control Transmission Protocol (SCTP)", RFC 6458, DOI 10.17487/ RFC6458, December 2011, <<u>https://www.rfc-editor.org/info/ rfc6458</u>>.

Authors' Addresses

Michael Tüxen Münster University of Applied Sciences Stegerwaldstrasse 39 48565 Steinfurt Germany

Email: tuexen@fh-muenster.de

Timo Völker Münster University of Applied Sciences Stegerwaldstrasse 39 48565 Steinfurt Germany

Email: timo.voelker@fh-muenster.de