

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 23, 2014

M. Tuexen
Muenster Univ. of Appl. Sciences
R. Seggelmann
T-Systems International GmbH
R. Stewart
Adara Networks
S. Loreto
Ericsson
October 20, 2013

Additional Policies for the Partial Delivery Extension of the Stream
Control Transmission Protocol
draft-tuexen-tsvwg-sctp-prpolicies-03.txt

Abstract

This document defines policies for the Partial Reliability Extension of the Stream Control Transmission Protocol (PR-SCTP) allowing to limit the number of retransmissions or to prioritize user messages for more efficient send buffer usage.

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 23, 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
1.1. Overview	2
1.2. Data Types	2
2. Additional PR-SCTP Policies	3
2.1. Limited Retransmissions Policy	3
2.2. Priority Policy	3
3. Socket API Considerations	3
3.1. Support for Added PR-SCTP Policies	3
3.2. Socket Option for Getting the PR-SCTP Status (SCTP_PR_STATUS)	4
4. IANA Considerations	5
5. Security Considerations	5
6. Acknowledgments	5
7. References	5
7.1. Normative References	5
7.2. Informative References	5
Authors' Addresses	6

1. Introduction

1.1. Overview

The SCTP Partial Reliability Extension (PR-SCTP) defined in [RFC3758] provides a generic method for senders to abandon user messages. The decision to abandon a user message is sender side only and the exact condition is called a PR-SCTP policy. [RFC3758] also defines one particular PR-SCTP policy, called Timed Reliability. This allows the sender to specify a timeout for a user message after which the SCTP stack abandons the user message.

This document specifies two additional PR-SCTP policies:

Limited Retransmission Policy: Allows to limit the number of retransmissions.

Priority Policy: Allows to discard lower priority messages if space for higher priority messages is needed in the send buffer.

1.2. Data Types

This documents uses data types from Draft 6.6 (March 1997) of POSIX 1003.1g: `uintN_t` means an unsigned integer of exactly N bits (e.g. `uint16_t`). This is the same as in [RFC6458]

2. Additional PR-SCTP Policies

2.1. Limited Retransmissions Policy

Using the Limited Retransmission Policy allows the sender of a user message to specify an upper limit for the number of retransmissions for each DATA chunk of the given user messages. The sender must abandon a user message if the number of retransmissions of any of the DATA chunks of the user message would exceed the provided limit. Please note that the number of retransmissions includes the fast and the timer based retransmissions.

Limiting the number of retransmissions to 0 is allowed. This provides a service similar to UDP, which also does not send any retransmissions either.

The Limited Retransmissions Policy is used for data channels in the RTCWeb protocol stack.

2.2. Priority Policy

Using the Priority Policy allows the sender of a user message to specify a priority. When storing a user message in the send buffer while there is not enough available space, the SCTP stack may abandon other user messages with a priority lower than the provided one.

After lower priority messages have been abandoned high priority messages can be transferred without blocking the `send()` call.

The Priority Policy can be used in the IPFIX protocol stack. See [RFC7011] for more information.

3. Socket API Considerations

This section describes how the socket API defined in [RFC6458] is extended to support the newly defined PR-SCTP policies and to provide some statistical information.

Please note that this section is informational only.

3.1. Support for Added PR-SCTP Policies

As defined in [RFC6458], the PR-SCTP policy is specified and configured by using the following `sctp_prinfo` structure:

```

struct sctp_prinfo {
    uint16_t pr_policy;
    uint32_t pr_value;
};

```

When the Limited Retransmission Policy described in Section 2.1 is used, `pr_policy` has the value `SCTP_PR_SCTP_RTX` and the number of retransmissions is given in `pr_value`.

For using the Priority Policy described in Section 2.2, `pr_policy` has the value `SCTP_PR_SCTP_PRIO`. The priority is given in `pr_value`. The value of zero is the highest priority and larger numbers in `pr_value` denote lower priorities.

The following table summarizes the possible parameter settings defined in [RFC6458] and this document:

<code>pr_policy</code>	<code>pr_value</code>	Specification
<code>SCTP_PR_SCTP_NONE</code>	Ignored	[RFC6458]
<code>SCTP_PR_SCTP_TTL</code>	Lifetime in ms	[RFC6458]
<code>SCTP_PR_SCTP_RTX</code>	Number of retransmissions	Section 2.1
<code>SCTP_PR_SCTP_PRIO</code>	Priority	Section 2.2

3.2. Socket Option for Getting the PR-SCTP Status (`SCTP_PR_STATUS`)

This socket option uses `IPPROTO_SCTP` as its level and `SCTP_PR_STATUS` as its name. It can only be used with `getsockopt()`, but not with `setsockopt()`. The socket option value uses the following structure:

```

struct sctp_prstatus {
    sctp_assoc_t sprstat_assoc_id;
    uint32_t sprstat_abandoned_unsent;
    uint32_t sprstat_abandoned_sent;
};

```

`sprstat_assoc_id`: This parameter is ignored for one-to-one style sockets. For one-to-many style sockets this parameter indicates for which association the user wants the information. It is an error to use `SCTP_{CURRENT|ALL|FUTURE}_ASSOC` in `sprstat_assoc_id`

`sprstat_abandoned_unsent`: The number of user messages which have been abandoned, before any part of the user message could be sent.

`sprstat_abandoned_sent`: The number of user messages which have been abandoned, after a part of the user message has been sent.

There are separate counters for unsent and sent user messages because the `SCTP_SEND_FAILED_EVENT` supports a similar differentiation. Please note that an abandoned large user messages requiring an SCTP level fragmentation is reported in the `sprstat_abandoned_sent` counter as soon as at least one fragment of it has been sent. Therefore each abandoned user messages is either counted in `sprstat_abandoned_unsent` or `sprstat_abandoned_sent`.

If more detailed information about abandoned user messages is required, the subscription to the `SCTP_SEND_FAILED_EVENT` is recommended.

`sctp_opt_info()` needs to be extended to support `SCTP_PR_STATUS`.

4. IANA Considerations

This document requires no actions from IANA.

5. Security Considerations

This document does not add any additional security considerations in addition to the ones given in [RFC4960], [RFC3758], and [RFC6458].

6. Acknowledgments

The authors wish to thank Irene Ruengeler, Jamal Hadi Salim, and Vlad Yasevich for there invaluable comments.

7. References

7.1. Normative References

- [RFC3758] Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP) Partial Reliability Extension", RFC 3758, May 2004.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

7.2. Informative References

- [RFC6458] Stewart, R., Tuexen, M., Poon, K., Lei, P., and V. Yasevich, "Sockets API Extensions for the Stream Control Transmission Protocol (SCTP)", RFC 6458, December 2011.

[RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013.

Authors' Addresses

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

Email: tuexen@fh-muenster.de

Robin Seggelmann
T-Systems International GmbH
Fasanenweg 5
70771 Leinfelden-Echterdingen
DE

Email: robin.seggelmann@t-systems.com

Randall R. Stewart
Adara Networks
Chapin, SC 29036
US

Email: randall@lakerest.net

Salvatore Loreto
Ericsson
Hirsalantie 11
Jorvas 02420
FI

Email: Salvatore.Loreto@ericsson.com