

Network Working Group	M. Tuexen
Internet-Draft	I. Ruengeler
Updates: 4960 (if approved)	Muenster Univ. of Appl. Sciences
Intended status: Standards Track	R. Stewart
Expires: March 01, 2014	Adara Networks
	August 28, 2013

SACK-IMMEDIATELY Extension for the Stream Control Transmission Protocol  
draft-ietf-tsvwg-sctp-sack-immediately-04.txt

Abstract

This document updates RFC 4960 by defining a method for the sender of a DATA chunk to indicate that the corresponding SACK chunk should be sent back immediately and not be delayed. It is done by specifying a bit in the DATA chunk header, called the I-bit, which can get set either by the SCTP implementation or by the application using an SCTP stack. Since unknown flags in chunk headers are ignored by SCTP implementations, this extension does not introduce any interoperability problems.

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 March 01, 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. The I-bit in the DATA Chunk Header . . . . .	3
4. Use Cases . . . . .	4
4.1. Triggering at the Application Level . . . . .	4
4.2. Triggering at the SCTP Level . . . . .	4
5. Procedures . . . . .	4
5.1. Sender Side Considerations . . . . .	5
5.2. Receiver Side Considerations . . . . .	5
6. Interoperability Considerations . . . . .	5
7. Socket API Considerations . . . . .	5
8. IANA Considerations . . . . .	6
9. Security Considerations . . . . .	7
10. Acknowledgments . . . . .	7
11. References . . . . .	7
11.1. Normative References . . . . .	7
11.2. Informative References . . . . .	7
Authors' Addresses . . . . .	7

## 1. Introduction

According to [RFC4960] the receiver of a DATA chunk should use delayed SACKs. This delaying is completely controlled by the receiver of the DATA chunk and remains the default behavior.

In specific situations the delaying of SACKs results in reduced performance of the protocol:

1. If such a situation can be detected by the receiver, the corresponding SACK can be sent immediately. For example, [RFC4960] recommends the immediate sending if the receiver has detected message loss or message duplication.

2. However, if the situation can only be detected by the sender of the DATA chunk, [RFC4960] provides no method of avoiding a delay in sending the SACK. Examples of these situations include ones which require interaction with the application (e.g. applications using the SCTP\_SENDER\_DRY\_EVENT, see Section 4.1) and ones which can be detected by the SCTP stack itself (e.g. closing the association, hitting window limits or resetting streams, see Section 4.2).

To overcome the limitation described in the second case, this document describes a simple extension of the SCTP DATA chunk by defining a new flag, the I-bit. The sender of a DATA chunk indicates by setting this bit that the corresponding SACK chunk should not be delayed.

## 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. The I-bit in the DATA Chunk Header

The following Figure 1 shows the extended DATA chunk.




Figure 1: Extended DATA chunk format

The only difference between the DATA chunk in Figure 1 and the DATA chunk defined in [RFC4960] is the addition of the I-bit in the flags field of the DATA chunk header.

This bit was Reserved in [RFC4960]. [RFC4960] specified that this bit should be set to 0 by the sender and ignored by the receiver.

#### 4. Use Cases

The setting of the I-bit can either be triggered by the application using SCTP or by the SCTP stack itself. The following two subsections provide a non-exhaustive list of examples.

##### 4.1. Triggering at the Application Level

One example of a situation in which it may be desirable for an application to trigger setting of the I-bit involves the SCTP\_SENDER\_DRY\_EVENT in the SCTP socket API [RFC6458]. Upper layers of SCTP using the socket API as defined in [RFC6458] may subscribe to the SCTP\_SENDER\_DRY\_EVENT for getting a notification as soon as no user data is outstanding anymore. To avoid an unnecessary delay while waiting for such an event, the application can request the setting of the I-Bit when sending the last user message before waiting for the event. This results in setting the I-bit of the last DATA chunk corresponding to the user message and is possible using the extension of the socket API described in Section 7.

##### 4.2. Triggering at the SCTP Level

There are also situations in which the SCTP implementation can set the I-bit without interacting with the upper layer.

If the association is in the SHUTDOWN-PENDING state, setting the I-bit reduces the number of simultaneous associations for a busy server handling short living associations.

Another case is where the sending of a DATA chunk fills the congestion or receiver window. Setting the I-bit in these cases improves the throughput of the transfer.

If an SCTP association supports the SCTP Stream Reconfiguration extension defined in [RFC6525], the performance can be improved by setting the I-bit when there are pending reconfiguration requests that require that there be no outstanding DATA chunks.

#### 5. Procedures

### 5.1. Sender Side Considerations

Whenever the sender of a DATA chunk can benefit from the corresponding SACK chunk being sent back without delay, the sender MAY set the I-bit in the DATA chunk header. Please note that it is irrelevant to the receiver why the sender has set the I-bit.

Reasons for setting the I-bit include, but are not limited to, the following (see Section 4 for the benefits):

- o The application requests to set the I-bit of the last DATA chunk of a user message when providing the user message to the SCTP implementation (see Section 7).
- o The sender is in the SHUTDOWN-PENDING state.
- o The sending of a DATA chunk fills the congestion or receiver window.
- o The sending of an Outgoing SSN Reset Request Parameter or an SSN/TSN Reset Request Parameter is pending, if the association supports the Stream Reconfiguration extension defined in [RFC6525].

### 5.2. Receiver Side Considerations

On reception of an SCTP packet containing a DATA chunk with the I-bit set, the receiver SHOULD NOT delay the sending of the corresponding SACK chunk, i.e., the receiver SHOULD immediately respond with the corresponding SACK chunk.

## 6. Interoperability Considerations

According to [RFC4960] the receiver of a DATA chunk with the I-bit set should ignore this bit when it does not support the extension described in this document. Since the sender of the DATA chunk is able to handle this case, there is no requirement for negotiating the support of the feature described in this document.

## 7. Socket API Considerations

This section describes how the socket API defined in [RFC6458] is extended to provide a way for the application to set the I-bit.

Please note that this section is informational only.

A socket API implementation based on [RFC6458] needs to be extended to allow the application to set the I-bit of the last DATA chunk when sending each user message.

This can be done by setting a flag called `SCTP_SACK_IMMEDIATELY` in the `snd_flags` field of the struct `sctp_sndinfo` structure when using `sctp_sendv()` or `sendmsg()`. If the deprecated struct `sctp_sndrcvinfo` structure is used instead when calling `sctp_send()`, `sctp_sendx()`, or `sendmsg()`, the `SCTP_SACK_IMMEDIATELY` flag can be set in the `sinfo_flags` field. When using the deprecated function `sctp_sendmsg()` the `SCTP_SACK_IMMEDIATELY` flag can be in the `flags` parameter.

## 8. IANA Considerations

[NOTE to RFC-Editor:

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

]

Following the chunk flag registration procedure defined in [RFC6096], IANA should register a new bit, the I-bit, for the DATA chunk. The suggested value is 0x08 and the reference should be RFCXXXX.

This requires an update of the "DATA Chunk Flags" registry for SCTP:

DATA Chunk Flags

Chunk Flag Value	Chunk Flag Name	Reference
0x01	E bit	[RFC4960]
0x02	B bit	[RFC4960]
0x04	U bit	[RFC4960]
0x08	I Bit	[RFCXXXX]
0x10	Unassigned	
0x20	Unassigned	
0x40	Unassigned	
0x80	Unassigned	

## 9. Security Considerations

See [RFC4960] for general security considerations for SCTP. In addition, a malicious sender can force its peer to send packets containing a SACK chunk for each received packet containing DATA chunks instead of every other. This could impact the network, resulting in more packets sent on the network, or the peer because the generating and sending of the packets has some processing cost. However, the additional packets can only contain the most simplest SACK chunk (no gap reports, no duplicate TSNs), since in case of packet drop or reordering in the network a SACK chunk would be sent immediately anyway. Therefore this does neither introduce a significant additional processing cost on the receiver side. This does not result in more traffic in the network than a receiver that sends a SACK for every packet, which is already permitted.

## 10. Acknowledgments

The authors wish to thank Mark Allmann, Brian Bidulock, David Black, Anna Brunstrom, Gorry Fairhurst, Janardhan Iyengar, Kacheong Poon, and Michael Welzl for their invaluable comments.

## 11. References

### 11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [RFC6096] Tuexen, M. and R. Stewart, "Stream Control Transmission Protocol (SCTP) Chunk Flags Registration", RFC 6096, January 2011.

### 11.2. 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.
- [RFC6525] Stewart, R., Tuexen, M., and P. Lei, "Stream Control Transmission Protocol (SCTP) Stream Reconfiguration", RFC 6525, February 2012.

## Authors' Addresses

Michael Tuexen  
Muenster University of Applied Sciences  
Stegerwaldstr. 39  
48565 Steinfurt  
DE

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

Irene Ruengeler  
Muenster University of Applied Sciences  
Stegerwaldstr. 39  
48565 Steinfurt  
DE

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

Randall R. Stewart  
Adara Networks  
Chapin, SC 29036  
US

Email: [randall@lakerest.net](mailto:randall@lakerest.net)