

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: August 29, 2013

R. Stewart  
Adara Networks  
M. Tuexen  
I. Ruengeler  
Muenster Univ. of Appl. Sciences  
February 25, 2013

Stream Control Transmission Protocol (SCTP) Network Address Translation  
Support

[draft-ietf-tsvwg-natsupp-05.txt](#)

Abstract

Stream Control Transmission Protocol [[RFC4960](#)] provides a reliable communications channel between two end-hosts in many ways similar to TCP [[RFC0793](#)]. With the widespread deployment of Network Address Translators (NAT), specialized code has been added to NAT for TCP that allows multiple hosts to reside behind a NAT and yet use only a single globally unique IPv4 address, even when two hosts (behind a NAT) choose the same port numbers for their connection. This additional code is sometimes classified as Network Address and Port Translation (NAPT). To date, specialized code for SCTP has not yet been added to most NATs so that only pure NAT is available. The end result of this is that only one SCTP capable host can be behind a NAT.

This document describes the protocol extensions required for the SCTP endpoints to help NATs provide similar features of NAPT in the single-point and multi-point traversal scenario.

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 August 29, 2013.

## 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](http://trustee.ietf.org/license-info) 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

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Conventions . . . . .</a>	<a href="#">5</a>
<a href="#">3.</a>	<a href="#">Terminology . . . . .</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">Data Formats . . . . .</a>	<a href="#">6</a>
<a href="#">4.1.</a>	<a href="#">Modified Chunks . . . . .</a>	<a href="#">6</a>
<a href="#">4.1.1.</a>	<a href="#">Extended ABORT Chunk . . . . .</a>	<a href="#">6</a>
<a href="#">4.1.2.</a>	<a href="#">Extended ERROR Chunk . . . . .</a>	<a href="#">7</a>
<a href="#">4.2.</a>	<a href="#">New Error Causes . . . . .</a>	<a href="#">7</a>
<a href="#">4.2.1.</a>	<a href="#">VTag and Port Number Collision Error Cause . . . . .</a>	<a href="#">7</a>
<a href="#">4.2.2.</a>	<a href="#">Missing State Error Cause . . . . .</a>	<a href="#">8</a>
<a href="#">4.2.3.</a>	<a href="#">Port Number Collision Error Cause . . . . .</a>	<a href="#">8</a>
<a href="#">4.3.</a>	<a href="#">New Parameters . . . . .</a>	<a href="#">9</a>
<a href="#">4.3.1.</a>	<a href="#">Disable Restart Parameter . . . . .</a>	<a href="#">9</a>
<a href="#">4.3.2.</a>	<a href="#">VTags Parameter . . . . .</a>	<a href="#">9</a>
<a href="#">5.</a>	<a href="#">Problem Space and Procedures . . . . .</a>	<a href="#">10</a>
<a href="#">5.1.</a>	<a href="#">Problem Space Overview . . . . .</a>	<a href="#">11</a>
<a href="#">5.2.</a>	<a href="#">Association Setup Considerations . . . . .</a>	<a href="#">11</a>
5.3.	<a href="#">Handling of Internal Port Number and Verification Tag Collisions . . . . .</a>	<a href="#">11</a>
<a href="#">5.4.</a>	<a href="#">Handling of Internal Port Number Collisions . . . . .</a>	<a href="#">12</a>
<a href="#">5.5.</a>	<a href="#">Handling of Missing State . . . . .</a>	<a href="#">13</a>
<a href="#">5.6.</a>	<a href="#">Multi-Point Traversal Considerations . . . . .</a>	<a href="#">15</a>
<a href="#">6.</a>	<a href="#">Socket API Considerations . . . . .</a>	<a href="#">15</a>
<a href="#">6.1.</a>	<a href="#">Get or Set the NAT Friendliness (SCTP_NAT_FRIENDLY) . . .</a>	<a href="#">16</a>
<a href="#">7.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">16</a>
<a href="#">7.1.</a>	<a href="#">New Chunk Flags for Two Existing Chunk Types . . . . .</a>	<a href="#">16</a>
<a href="#">7.2.</a>	<a href="#">Three New Error Causes . . . . .</a>	<a href="#">17</a>
<a href="#">7.3.</a>	<a href="#">Two New Chunk Parameter Types . . . . .</a>	<a href="#">18</a>
<a href="#">8.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">18</a>
<a href="#">9.</a>	<a href="#">Acknowledgments . . . . .</a>	<a href="#">18</a>
<a href="#">10.</a>	<a href="#">References . . . . .</a>	<a href="#">19</a>
<a href="#">10.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">19</a>
<a href="#">10.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">19</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">19</a>



## 1. Introduction

Stream Control Transmission Protocol [RFC4960] provides a reliable communications channel between two end-hosts in many ways similar to TCP [RFC0793]. With the widespread deployment of Network Address Translators (NAT), specialized code has been added to NAT for TCP that allows multiple hosts to reside behind a NAT using private addresses (see [RFC5735]) and yet use only a single globally unique IPv4 address, even when two hosts (behind a NAT) choose the same port numbers for their connection. This additional code is sometimes classified as Network Address and Port Translation (NAPT). To date, specialized code for SCTP has not yet been added to most NATs so that only true NAT is available. The end result of this is that only one SCTP capable host can be behind a NAT.

This document describes SCTP specific packets and procedures to help NATs provide similar features of NAPT in the single-point and multi-point traversal scenario. An SCTP implementation supporting this extension will follow these procedures to assure that in both single-homed and multi-homed cases a NAT will maintain the proper state without needing to change port numbers.

A NAT will need to follow these procedures for generating appropriate SCTP packet formats. NATs should refer to [I-D.ietf-behave-sctpnat] for the BCP in using these formats.

When considering this feature it is possible to have multiple levels of support. At each level, the Internal Host, External Host and NAT may or may not support the features described in this document. The following table illustrates the results of the various combinations of support and if communications can occur between two endpoints.

Internal Host	NAT	External Host	Communication
Support	Support	Support	Yes
Support	Support	No Support	Limited
Support	No Support	Support	None
Support	No Support	No Support	None
No Support	Support	Support	Limited
No Support	Support	No Support	Limited
No Support	No Support	Support	None
No Support	No Support	No Support	None

Table 1: Communication possibilities

From the table we can see that when a NAT does not support the



extension no communication can occur. This is because for the most part of the current situation i. e. SCTP packets sent externally from behind a NAT are discarded by the NAT. In some cases, where the NAT supports the feature but one of the two external hosts does not support the feature, communication may occur but in a limited way. For example only one host may be able to have a connection when a collision case occurs.

## 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. Terminology

This document uses the following terms, which are depicted in Figure 1.

Private-Address (Priv-Addr): The private address that is known to the internal host.

Internal-Port (Int-Port): The port number that is in use by the host holding the Private-Address.

Internal-VTag (Int-VTag): The Verification Tag that the internal host has chosen for its communication. The VTag is a unique 32-bit tag that must accompany any incoming SCTP packet for this association to the Private-Address.

External-Address (Ext-Addr): The address that an internal host is attempting to contact.

External-Port (Ext-Port): The port number of the peer process at the External-Address.

External-VTag (Ext-VTag): The Verification Tag that the host holding the External-Address has chosen for its communication. The VTag is a unique 32-bit tag that must accompany any incoming SCTP packet for this association to the External-Address.

Public-Address (Pub-Addr): The public address assigned to the NAT box which it uses as a source address when sending packets towards the External-Address.





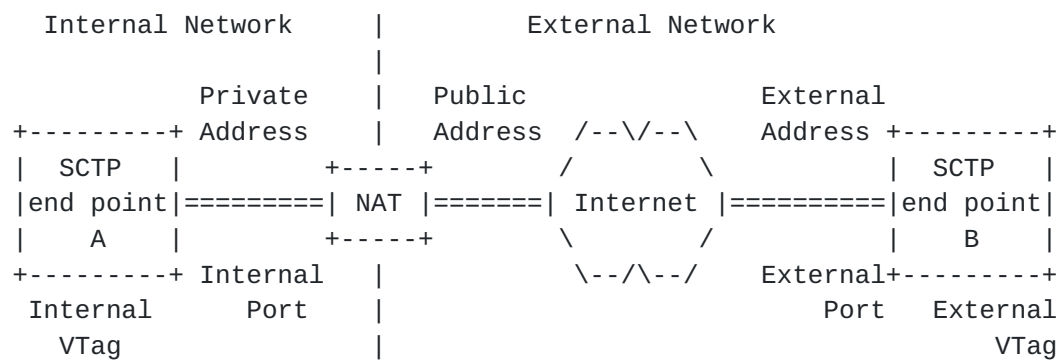


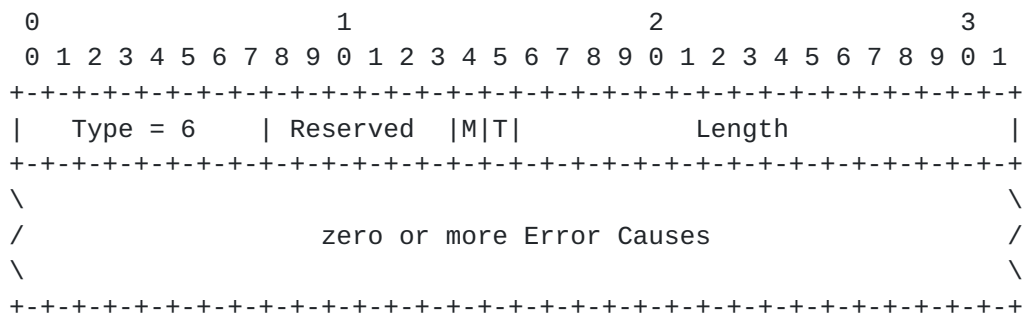
Figure 1: Basic network setup

## 4. Data Formats

### 4.1. Modified Chunks

This section presents existing chunks defined in [RFC4960] that are modified by this document.

#### 4.1.1. Extended ABORT Chunk



The ABORT chunk is extended to add the new 'M-bit'. The M-bit indicates to the receiver of the ABORT chunk that the chunk was not generated by the peer SCTP endpoint, but instead by a middle box.



#### 4.1.2. Extended ERROR Chunk

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Type = 9      | Reserved |M|T|           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
\
/               zero or more Error Causes                    /
\
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The ERROR chunk defined in [RFC4960] is extended to add the new 'M-bit'. The M-bit indicates to the receiver of the ERROR chunk that the chunk was not generated by the peer SCTP endpoint, but instead by a middle box.

#### 4.2. New Error Causes

This section defines the new error causes added by this document.

##### 4.2.1. VTag and Port Number Collision Error Cause

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Cause Code = 0x00B0      | Cause Length = Variable      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
\
/               Chunk                                           /
\
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Cause Code: 2 bytes (unsigned integer)

This field holds the IANA defined cause code for the VTag and Port Number Collision Error Cause. The suggested value of this field for IANA is 0x00B0.

Cause Length: 2 bytes (unsigned integer)

This field holds the length in bytes of the error cause. The value MUST be the length of the Cause-Specific Information plus 4.

Chunk: variable length

The Cause-Specific Information is filled with the chunk that caused this error. This can be an INIT, INIT-ACK, or ASCONF chunk. Note that if the entire chunk will not fit in the ERROR chunk or ABORT chunk being sent then the bytes that do not fit are truncated.



**4.2.2. Missing State Error Cause**

```

      0                               1                               2                               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Cause Code = 0x00B1          |      Cause Length = Variable  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
\                                     Incoming Packet                     /
/                                                                              \
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Cause Code: 2 bytes (unsigned integer)

This field holds the IANA defined cause code for the Missing State Error Cause. The suggested value of this field for IANA is 0x00B1.

Cause Length: 2 bytes (unsigned integer)

This field holds the length in bytes of the error cause. The value MUST be the length of the Cause-Specific Information plus 4.

Incoming Packet: variable length

The Cause-Specific Information is filled with the IPv4 or IPv6 packet that caused this error. The IPv4 or IPv6 header MUST be included. Note that if the packet will not fit in the ERROR chunk or ABORT chunk being sent then the bytes that do not fit are truncated.

**4.2.3. Port Number Collision Error Cause**

```

      0                               1                               2                               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Cause Code = 0x00B2          |      Cause Length = Variable  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
\                                     chunk                               /
/                                                                              \
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Cause Code: 2 bytes (unsigned integer)

This field holds the IANA defined cause code for the Port Number Collision Error Cause. The suggested value of this field for IANA is 0x00B2.

Cause Length: 2 bytes (unsigned integer)

This field holds the length in bytes of the error cause. The value MUST be the length of the Cause-Specific Information plus 4.









```

      0             1             2             3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Parameter Type = 0xC008   |   Parameter Length = 16   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               ASCONF-Request Correlation ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Internal Verification Tag                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               External Verification Tag                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Parameter Type: 2 bytes (unsigned integer)

This field holds the IANA defined parameter type for the VTags Parameter. The suggested value of this field for IANA is 0xC008.

Parameter Length: 2 bytes (unsigned integer)

This field holds the length in bytes of the parameter. The value MUST be 16.

ASCONF-Request Correlation ID: 4 bytes (unsigned integer)

This is an opaque integer assigned by the sender to identify each request parameter. The receiver of the ASCONF Chunk will copy this 32-bit value into the ASCONF Response Correlation ID field of the ASCONF-ACK response parameter. The sender of the ASCONF can use this same value in the ASCONF-ACK to find which request the response is for. Note that the receiver MUST NOT change this 32-bit value.

Internal Verification Tag: 4 bytes (unsigned integer)

The Verification Tag that the internal host has chosen for its communication. The Verification Tag is a unique 32-bit tag that must accompany any incoming SCTP packet for this association to the Private-Address.

External Verification Tag: 4 bytes (unsigned integer) The

Verification Tag that the host holding the External-Address has chosen for its communication. The VTag is a unique 32-bit tag that must accompany any incoming SCTP packet for this association to the External-Address.

This parameter MAY appear in ASCONF chunks and MUST NOT appear in any other chunk.

## 5. Problem Space and Procedures



### **5.1. Problem Space Overview**

When an SCTP endpoint is behind a NAT which supports [\[I-D.ietf-behave-sctpnat\]](#) a number of problems may arise as it tries to communicate with its peer:

- o More than one server behind a NAT may pick the same VTag and source port when talking to the same peer server. This creates a situation where the NAT will not be able to tell the two associations apart. This situation is discussed in [Section 5.3](#).
- o When an SCTP endpoint is a server communicating with multiple peers and the peers are behind the same NAT, then the two endpoints cannot be distinguished by the server. This case is discussed in [Section 5.4](#).
- o A restart of a NAT during a conversation could cause a loss of its state. This problem and its solution is discussed in [Section 5.5](#).
- o An SCTP endpoint may be behind two NATs providing redundancy. The method to set up this scenario is discussed in [Section 5.6](#).

Each of these solutions requires additional chunks and parameters, defined in this document, and possibly modified handling procedures from those specified in [\[RFC4960\]](#).

### **5.2. Association Setup Considerations**

Every association MUST initially be set up single-homed. There MUST NOT be any IPv4 Address parameter, IPv6 Address parameter, or Supported Address Types parameter in the INIT-chunk. The INIT-ACK chunk MUST NOT contain any IPv4 Address parameter or IPv6 Address parameter.

If the association should finally be multi-homed, the procedure in [Section 5.6](#) MUST be used.

The INIT and INIT-ACK chunk SHOULD contain the Disable Restart parameter defined in [Section 4.3.1](#).

### **5.3. Handling of Internal Port Number and Verification Tag Collisions**

Consider the case where two hosts in the Private-Address space want to set up an SCTP association with the same server running on the same host in the Internet. This means that the External-Port and the External-Address are the same. If they both choose the same Internal-Port and Internal-VTag, the NAT box cannot distinguish between incoming packets anymore. But this is very unlikely. The



Internal-VTags are chosen at random and if the Internal-Ports are also chosen from the ephemeral port range at random this gives a 46-bit random number which has to match. In the TCP like NAPT case the NAT box can control the 16-bit Natted Port and therefor avoid collisions deterministically.

The same can happen when an INIT-ACK chunk or an ASCONF chunk is processed by the NAT.

However, in this unlikely event the NAT box MUST send an ABORT chunk with the M-bit set if the collision is triggered by an INIT or INIT-ACK chunk or send an ERROR chunk with the M-bit set if the collision is triggered by an ASCONF chunk. The M-bit is a new bit defined by this document to express to SCTP that the source of this packet is a "middle" box, not the peer SCTP endpoint (see [Section 4.1.1](#)). In a packet containing an INIT-ACK chunk triggers the collision, the corresponding packet containing the ABORT chunk MUST contain the same source and destination address and port numbers as the packet containing the INIT-ACK chunk. In the other two cases, the source and destination address and port numbers MUST be swapped.

The sender of the packet containing the INIT chunk or the receiver of the INIT-ACK chunk, upon reception of an ABORT chunk with M-bit set, SHOULD reinitiate the association setup procedure after choosing a new initiate tag. These procedures SHOULD be followed only if the appropriate error cause code for colliding NAT table state is included AND the association is in the COOKIE-WAIT state (i. e. it is awaiting an INIT-ACK). If the endpoint is in any other state an SCTP endpoint SHOULD NOT respond.

The sender of the ASCONF chunk, upon reception of an ERROR chunk with M-bit set, MUST stop adding the path to the association.

The sender of the ERROR or ABORT chunk MUST include the error cause with cause code 'VTag and Port Number Collision' (see [Section 4.2.1](#)).

#### **5.4. Handling of Internal Port Number Collisions**

When two SCTP hosts are behind a NAT and using the recommendations in [[I-D.ietf-behave-sctpnat](#)] it is possible that two SCTP hosts in the Private-Address space will want to set up an SCTP association with the same server running on the same host in the Internet. For the NAT appropriate tracking may be performed by assuring that the VTags are unique between the two hosts as defined in [[I-D.ietf-behave-sctpnat](#)]. But for the external SCTP server on the internet this means that the External-Port and the External-Address are the same. If they both have chosen the same Internal-Port the server cannot distinguish between both associations based on the



address and port numbers. For the server it looks like the association is being restarted. To overcome this limitation the client sends a Disable Restart parameter in the INIT-chunk.

When the server receives this parameter it MUST do the following:

- o Include a Disable Restart parameter in the INIT-ACK to inform the client that it will support the feature.
- o Disable the restart procedures defined in [[RFC4960](#)] for this association.

Servers that support this feature will need to be capable of maintaining multiple connections to what appears to be the same peer (behind the NAT) differentiated only by the VTags.

The NAT, when processing the INIT-ACK, should note in its internal table that the association supports the Disable Restart extension. This note is used when establishing future associations (i. e. when processing an INIT from an internal host) to decide if the connection should be allowed. The NAT MUST do the following when processing an INIT:

- o If the INIT is destined to an external address and port for which the NAT has no outbound connection, allow the INIT creating an internal mapping table.
- o If the INIT matches the external address and port of an already existing connection, validate that the external server supports the Disable Restart feature, if it does allow the INIT to be forwarded.
- o If the external server does not support the Disable Restart extension the NAT MUST send an ABORT with the M-bit set.

The 'Port Number Collision' error cause (see [Section 4.2.3](#)) MUST be included in the ABORT chunk.

If the collision is triggered by an ASCONF chunk, a packet containing an ERROR chunk with the 'Port Number Collision' error cause MUST be sent back.

## **[5.5.](#) Handling of Missing State**

If the NAT box receives a packet from the internal network for which the lookup procedure does not find an entry in the NAT table, a packet containing an ERROR chunk is sent back with the M-bit set. The source address of the packet containing the ERROR chunk MUST be





the destination address of the incoming SCTP packet. The verification tag is reflected and the T-bit is set. Please note that such a packet containing an ERROR chunk SHOULD NOT be sent if the received packet contains an ABORT, SHUTDOWN-COMPLETE or INIT-ACK chunk. An ERROR chunk MUST NOT be sent if the received packet contains an ERROR chunk with the M-bit set.

When sending the ERROR chunk, the new error cause Missing state (see [Section 4.2.2](#)) MUST be included and the new M-bit of the ERROR chunk MUST be set (see [Section 4.1.2](#)).

Upon reception of this ERROR chunk by an SCTP endpoint the receiver SHOULD take the following actions:

- o Validate that the verification tag is reflected by looking at the VTag that would have been included in the outgoing packet.
- o Validate that the peer of the SCTP association supports the dynamic address extension, if it does not discard the incoming ERROR chunk.
- o Generate a new ASCONF chunk containing the VTags parameter (see [Section 4.3.2](#)) and the Disable Restart parameter if the association is using the disabled restart feature. By processing this packet the NAT can recover the appropriate state. The procedures for generating an ASCONF chunk can be found in [\[RFC5061\]](#).

If the NAT box receives a packet for which it has no NAT table entry and the packet contains an ASCONF chunk with the VTags parameter, the NAT box MUST update its NAT table according to the verification tags in the VTags parameter and the optional Disable Restart parameter.

The peer SCTP endpoint receiving such an ASCONF chunk SHOULD either add the address and respond with an acknowledgment, if the address is new to the association (following all procedures defined in [\[RFC5061\]](#)). Or, if the address is already part of the association, the SCTP endpoint MUST NOT respond with an error, but instead should respond with an ASCONF-ACK chunk acknowledging the address but take no action (since the address is already in the association).

Note that it is possible that upon receiving an ASCONF chunk containing the VTags parameter the NAT will realize that it has an 'Internal Port Number and Verification Tag collision'. In such a case the NAT MUST send an ERROR chunk with the error cause code set to 'VTag and Port Number Collision' (see [Section 4.2.1](#)).

If an SCTP endpoint receives an ERROR with 'Internal Port Number and



Verification Tag collision' as the error cause and the packet in the Error Chunk contains an ASCONF with the VTags parameter, careful examination of the association is required. The endpoint MUST do the following:

- o Validate that the verification tag is reflected by looking at the VTag that would have been included in the outgoing packet.
- o Validate that the peer of the SCTP association supports the dynamic address extension, if it does not discard the incoming ERROR chunk.
- o If the association is attempting to add an address (i. e. following the procedures in [Section 5.6](#)) then the endpoint MUST-NOT consider the address part of the association and SHOULD make no further attempt to add the address (i. e. cancel any ASCONF timers and remove any record of the path), since the NAT has a VTag collision and the association cannot easily create a new VTag (as it would if the error occurred when sending an INIT).
- o If the endpoint has no other path, i. e. the procedure was executed due to missing a state in the NAT, then the endpoint MUST abort the association. This would occur only if the local NAT restarted and accepted a new association before attempting to repair the missing state (Note that this is no different than what happens to all TCP connections when a NAT loses its state).

### **[5.6.](#) Multi-Point Traversal Considerations**

If a multi-homed SCTP endpoint behind a NAT connects to a peer, it SHOULD first set up the association single-homed with only one address causing the first NAT to populate its state. Then it SHOULD add each IP address using ASCONF chunks sent via their respective NATs. The address to add is the wildcard address and the lookup address SHOULD also contain the VTags parameter and optionally the Disable Restart parameter as illustrated above.

## **[6.](#) Socket API Considerations**

This section describes how the socket API defined in [[RFC6458](#)] is extended to provide a way for the application to control NAT friendliness.

Please note that this section is informational only.

A socket API implementation based on [[RFC6458](#)] is extended by supporting one new read/write socket option.



### **6.1. Get or Set the NAT Friendliness (SCTP\_NAT\_FRIENDLY)**

This socket option uses the option\_level IPPROTO\_SCTP and the option\_name SCTP\_NAT\_FRIENDLY. It can be used to enable/disable the NAT friendliness for future associations and retrieve the value for future and specific ones.

```
struct sctp_assoc_value {  
    sctp_assoc_t assoc_id;  
    uint32_t assoc_value;  
};
```

assoc\_id: This parameter is ignored for one-to-one style sockets. For one-to-many style sockets the application may fill in an association identifier or SCTP\_FUTURE\_ASSOC for this query. It is an error to use SCTP\_{CURRENT|ALL}\_ASSOC in assoc\_id.

assoc\_value: A non-zero value indicates a NAT-friendly mode.

## **7. IANA Considerations**

[NOTE to RFC-Editor:

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

]

[NOTE to RFC-Editor:

The suggested values for the chunk type and the chunk parameter types are tentative and to be confirmed by IANA.

]

This document (RFCXXXX) is the reference for all registrations described in this section. The suggested changes are described below.

### **7.1. New Chunk Flags for Two Existing Chunk Types**

As defined in [[RFC6096](#)] two chunk flags have to be assigned by IANA for the ERROR chunk. The suggested value for the T bit is 0x01 and for the M bit is 0x02.

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



## ERROR Chunk Flags

Chunk Flag Value	Chunk Flag Name	Reference
0x01	T bit	[RFCXXXX]
0x02	M bit	[RFCXXXX]
0x04	Unassigned	
0x08	Unassigned	
0x10	Unassigned	
0x20	Unassigned	
0x40	Unassigned	
0x80	Unassigned	

As defined in [[RFC6096](#)] one chunk flag has to be assigned by IANA for the ABORT chunk. The suggested value of the M bit is 0x02.

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

## ABORT Chunk Flags

Chunk Flag Value	Chunk Flag Name	Reference
0x01	T bit	[ <a href="#">RFC4960</a> ]
0x02	M bit	[RFCXXXX]
0x04	Unassigned	
0x08	Unassigned	
0x10	Unassigned	
0x20	Unassigned	
0x40	Unassigned	
0x80	Unassigned	

## 7.2. Three New Error Causes

Three error causes have to be assigned by IANA. It is suggested to use the values given below.

This requires three additional lines in the "Error Cause Codes" registry for SCTP:





## Error Cause Codes

Value	Cause Code	Reference
176	VTag and Port Number Collision	[RFCXXXX]
177	Missing State	[RFCXXXX]
178	Port Number Collision	[RFCXXXX]

**7.3. Two New Chunk Parameter Types**

Two chunk parameter types have to be assigned by IANA. It is suggested to use the values given below. IANA should assign these values from the pool of parameters with the upper two bits set to '11'.

This requires two additional lines in the "Chunk Parameter Types" registry for SCTP:

## Chunk Parameter Types

ID Value	Chunk Parameter Type	Reference
49159	Disable Restart (0xC007)	[RFCXXXX]
49160	VTags (0xC008)	[RFCXXXX]

**8. Security Considerations**

The document does not add any additional security considerations to the ones given in [\[RFC4960\]](#), [\[RFC4895\]](#), and [\[RFC5061\]](#).

**9. Acknowledgments**

The authors wish to thank Jason But, Bryan Ford, David Hayes, Alfred Hines, Henning Peters, Timo Voelker, Dan Wing, and Qiaobing Xie for their invaluable comments.

**10. References**



### **10.1. Normative References**

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4895] Tuexen, M., Stewart, R., Lei, P., and E. Rescorla, "Authenticated Chunks for the Stream Control Transmission Protocol (SCTP)", [RFC 4895](#), August 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.
- [RFC6096] Tuexen, M. and R. Stewart, "Stream Control Transmission Protocol (SCTP) Chunk Flags Registration", [RFC 6096](#), January 2011.

### **10.2. Informative References**

- [RFC5735] Cotton, M. and L. Vegoda, "Special Use IPv4 Addresses", [BCP 153](#), [RFC 5735](#), January 2010.
- [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.
- [I-D.ietf-behave-sctpnat] Stewart, R., Tuexen, M., and I. Ruengeler, "Stream Control Transmission Protocol (SCTP) Network Address Translation", [draft-ietf-behave-sctpnat-07](#) (work in progress), October 2012.



Authors' Addresses

Randall R. Stewart  
Adara Networks  
Chapin, SC 29036  
US

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

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

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

Irene Ruengeler  
Muenster University of Applied Sciences  
Stegerwaldstrasse 39  
48565 Steinfurt  
DE

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

