

expires in six months

February 23, 2001

SCTP Dynamic Addition of IP addresses
<[draft-ietf-sigtran-addip-sctp-01.txt](#)>

Status of This Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#) [RFC2026]. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Abstract

This document describes an extension to the Stream Control Transmission Protocol (SCTP) [RFC2960] to provide a method to add or delete an IP address from an existing association or to request the peer to change its primary destination address. The benefit of this extension is for machines with hot pluggable interface cards the ability to add (and/or delete) IP addresses dynamically without forcing an SCTP association restart.

TABLE OF CONTENTS

1. Introduction.....	2
2. Conventions.....	2
3. New Reliable Request/Acknowledgement Parameters.....	2
3.1 New REL-REQ Parameters.....	2
3.1.1 Add IP Address.....	3
3.1.2 Delete IP Address.....	3
3.1.3 Set Primary IP Address.....	4
3.2 New REL-ACK Parameters.....	4
3.2.1 Error Cause: Request to Delete Last IP Address.....	5
3.2.2 Error Cause: Operation Refused due to Resource Shortage.	5
3.2.3 Error Cause: Request to Delete Source IP Address.....	6

3.3 New Error Causes.....	6
3.3.1 Request to delete last IP address.....	6
4 . Procedures.....	6

4.1	IP address addition and deletion.....	6
4.2	Setting of the Primary Address.....	9
5.	Security Considerations.....	9
6.	IANA considerations.....	10
7.	Authors' Addresses.....	10
8.	References.....	11

1. Introduction

Taking advantage of the extensibility of SCTP, and the generic method for transmitting reliable SCTP control chunks [[RELREQ](#)], this document introduces a use of the reliable control chunk, to allow an existing SCTP association to add or delete IP addresses without the currently required restart of the association.

This enables SCTP to have dynamic IP addresses added and subtracted for those machines that allow addition/removal of an interface card. It also provides a method for a endpoint request that its peer set its primary destination address to an alternate, this can be useful when an address is about to be deleted. As an example application, this will provide the same type of services that exist in telephony signalling, which allow a link set to add an additional link without interfering with the operation of the link set.

2. Conventions

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, NOT RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

3. New Reliable Request/Acknowledgement Parameters

This section describes the addition of three new REL-REQ Parameters to allow for the dynamic addition, deletion and setting of the primary, of IP addresses to an existing SCTP association. The format of these REL-REQ parameters within the Reliable Request Chunk is descibed in [[REL-REQ](#)]. We also describe a REL-ACK parameter that is carried to communicate errors or rejections of REL-REQ parameters.

3.1 New REL-REQ Parameters

The three new REL-REQ Parameters added follow the format defined in [[REL-REQ](#)] [section 3.1.1](#). Table 2 describes the three new REL-REQ Parameters.

REL-REQ Parameter	Type Value
-----	-----
Add IP Address	49153 (0xC001)

Delete IP Address	49154 (0xC002)
Set Primary Address	49159 (0xC007)

Table 1: REL-REQ Parameters

3.1.1 Add IP Address

```

      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 = 49153           | Length = Variable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Address Parameter                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Address Parameter: TLV

This field contains an IPv4 or IPv6 address parameter as described in 3.3.2.1 of [RFC2960](#). The complete TLV is wrapped within this parameter. It informs the receiver that the Address specified is to be added to the existing association. The sender of this request MUST NOT attempt to add an address type that is not supported by its peer.

An example TLV adding the IPv4 address 10.1.1.1 to an existing association would look as follows:

```

+-----+
| Type=49153 | Length = 12 |
+-----+
| Type=5     | Length = 8  |
+-----+
| Value=0x0a010101 |
+-----+

```

3.1.2 Delete IP Address

```

      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 = 49154           | Length = Variable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Address Parameter                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Address Parameter: TLV

This field contains an IPv4 or IPv6 address parameter as described in 3.3.2.1 of [\[RFC2960\]](#). The complete TLV is wrapped within this parameter. It informs the receiver that the Address specified is to be removed from the existing association.

An example TLV deleting the IPv4 address 10.1.1.1 from an existing association would look as follows:


```

+-----+
| Type=49154   | Length = 12   |
+-----+
| Type=5       | Length = 8    |
+-----+-----+
|           Value=0x0a010101           |
+-----+-----+

```

3.1.3 Set Primary IP Address

```

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 = 49159           | Length = Variable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Address Parameter                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Address Parameter: TLV

This field contains an IPv4 or IPv6 address parameter as described in 3.3.2.1 of [[RFC2960](#)]. The complete TLV is wrapped within this parameter. It requests the receiver to mark the specified address as the primary address to send data to (see [section 5.1.2 of \[RFC2960\]](#)). The receiver MAY mark this as its primary upon receiving this request.

An example TLV requesting that the IPv4 address 10.1.1.1 be made the primary destination address would look as follows:

```

+-----+
| Type=49159   | Length = 12   |
+-----+
| Type=5       | Length = 8    |
+-----+-----+
|           Value=0x0a010101           |
+-----+-----+

```

3.2 New REL-ACK Parameters

Three new Error Causes are added to the SCTP Operational Errors, primarily for use as part of the REL-REQ Parameter Error in the REL-ACK (as outlined in [[REL-REQ](#)] [section 3.1.2](#)).

Cause Code

Value	Cause Code
-----	-----
11	Request to delete last IP address.
12	Operation Refused due to resource shortage.
13	Request to delete source IP address.

3.2.1 Error Cause: Request to Delete Last IP Address

Cause of error

Request to delete last IP address: The receiver of this error sent a request to delete the last IP address from its association with its peer. This error indicates that the request is rejected.

```
+-----+
| Cause Code=11 | Cause Length=var |
+-----+
\ TLV-Copied-From-REL-REQ /
/ \
+-----+
```

An example of a failed delete in an Error Cause TLV would look as follows in the response REL-ACK message:

```
+-----+
| Type = 0xC005 | Length = 20 |
+-----+
| Cause=11 | Length = 16 |
+-----+
| Type=49154 | Length = 12 |
+-----+
| Type=5 | Length = 8 |
+-----+
| Value=0x0a010101 |
+-----+
```

3.2.2 Error Cause: Operation Refused due to Resource Shortage

Cause of error

This error cause is used to report a failure by the receiver to perform the requested operation due to a lack of resources. The entire TLV that is refused is copied from the REL-REQ into the error cause.

```
+-----+
| Cause Code=12 | Cause Length=Var |
+-----+
\ TLV-Copied-From-REL-REQ /
/ \
+-----+
```

An example of a failed addition in an Error Cause TLV would look as follows in the response REL-ACK message:

```
+-----+
```

Type = 0xC005	Length = 20	
+-----+		
Cause=12	Length = 16	
+-----+		
Type=49153	Length = 12	
+-----+		

```

|  Type=5          | Length = 8      |
+-----+-----+
|      Value=0x0a010101      |
+-----+-----+

```

3.2.3 Error Cause: Request to Delete Source IP Address

Cause of error

Request to delete source IP address: The receiver of this error sent a request to delete the source IP address of the REL-REQ message. This error indicates that the request is rejected.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Cause Code=13          |      Cause Length=VAR      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
\                               TLV-Copied-From-REL-REQ                /
/                               \                                     \
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

An example of a failed delete in an Error Cause TLV would look as follows in the response REL-ACK message:

```

+-----+-----+
| Type = 0xC005 | Length = 20 |
+-----+-----+
| Cause=13      | Length = 16  |
+-----+-----+
| Type=49154    | Length = 12  |
+-----+-----+
| Type=5        | Length = 8    |
+-----+-----+
|      Value=0x0a010101      |
+-----+-----+

```

4. Procedures

This section will lay out the specific procedures for the add/delete IP address REL-REQ Parameter. The generic procedures for transmitting reliable SCTP control chunks are covered in [\[RELREQ\]](#).

4.1 IP address addition and deletion

When building TLV parameters for the REL-REQ Chunk messages that will add or delete IP addresses the following rules should be applied:

- D1) When adding an IP address to an association, the IP address is NOT considered fully added to the association until the REL-ACK arrives. This means that until such time as the REL-REQ

containing the add is acknowledged the sender MUST NOT use the new IP address as a source for ANY SCTP packet.

D2) After the REL-ACK of an IP address add arrives, the endpoint MAY

begin using the added IP address as a source address.

- D3) If an endpoint receives an Error Cause TLV indicating that the IP address Add, IP address Deletion, or Set Primary IP Address REL-REQ parameters was not understood, the endpoint MUST consider the operation failed and MUST NOT attempt to send any subsequent Add, Delete or Set Primary requests to the peer.
- D4) When deleting an IP address from an association, the IP address MUST be considered a valid source address for DATA chunks until the REL-ACK arrives. This means that any datagrams that arrive before the REL-ACK destined to the IP address being deleted MUST be considered part of the current association. One special consideration is that ABORT chunks arriving destined to the IP address being deleted MUST be ignored (see [Section 4.1](#) for further details).
- D5) An endpoint MUST NOT delete its last IP address from an association. In other words if an endpoint is NOT multi-homed it MUST NOT use the delete IP address. Or if an endpoint sends multiple requests to delete IP addresses it MUST NOT delete all of the IP addresses that the peer has listed for the requester.
- D6) An endpoint MUST NOT use the source address (of the IP packet containing the SCTP packet) for a REL-REQ to delete an IP address from the address being deleted.
- D7) If a request is received to delete the last IP address of a peer endpoint, the receiver MUST send an Error Cause TLV with the error cause set to the new error code 'Request to delete last IP address'. The requested delete MUST NOT be performed or acted upon, other than to send the REL-ACK.
- D8) If a request is received to delete an IP address which is also the source address of the IP packet which contained the REL-REQ chunk, the receiver MUST reject this request. To reject the request the receiver MUST send a Error Cause TLV set to the new error code "Request to Delete Source IP Address" (unless Rule D5 has also been violated, in which case the error code 'Request to delete last IP address' is sent).
- D9) After the REL-ACK of an IP address deletion arrives, the endpoint MUST NOT use the deleted IP address as a source of any SCTP packet.
- D10) If an endpoint receives an ADD IP address request and does not have the local resources to add this new address to the association, it MUST return an Error Cause TLV set to the new


error code "Operation Refused due to Resource Shourtage".

Stewart et.al.

[Page 7]


- D11) If an endpoint receives an 'Out of Resource' error when adding an IP address to an association, it must either ABORT the association or not source any packets from this address. In other words if the endpoint does not ABORT the association, it must consider the add attempt failed and NOT use this address.
- D12) When an endpoint receiving a REL-REQ to add an IP address sends an 'Out of Resource' in its response, it MUST also fail any subsequent add or delete requests bundled in the REL-REQ. The receiver MUST NOT reject an ADD and then accept a subsequent DELETE of an IP address in the same REL-REQ chunk. In other words, once a receiver begins failing any ADD or DELETE request, it must fail all subsequent ADD or DELETE requests contained in that single REL-REQ.
- D13) When an endpoint receives a request to delete an IP address that is the current primary address, it is an implementation decision as to how that endpoint chooses the new primary address.

During the time interval between sending out the REL-REQ and receiving the REL-ACK it MAY be possible to receive DATA chunks out of order. The following examples illustrate these problems:



In the above example we see a new IP address (X) being added to the Endpoint-A. However due to packet re-ordering in the network a new DATA chunk is sent and arrives at Endpoint-A before the REL-ACK confirming the add of the address to the association.

A similar problem exists with the deletion of an IP address as follows:




/
/
<-----/

In this example we see a DATA chunk destined to the IP:X (which is about to be deleted) arriving after the deletion is complete.

For the ADD case an endpoint SHOULD consider the newly adding IP address valid for the association to receive data from during the interval when awaiting the REL-ACK. The endpoint MUST NOT source data from this new address until the REL-ACK arrives but it may receive out of order data as illustrated and MUST NOT treat this data as an OOTB datagram (please see [\[RFC2960\] section 8.4](#)). It MAY drop the data silently or it MAY consider it part of the association but it MUST NOT respond with an ABORT.

For the DELETE case, an endpoint MAY respond to the late arriving DATA packet as an OOTB datagram or it MAY hold the deleting IP address for a small period of time as still valid. If it treats the DATA packet as an OOTB the peer will silently discard the ABORT (since by the time the ABORT is sent the peer will have removed the IP address from this association). If the endpoint elects to hold the IP address valid for a period of time, it MUST NOT hold it valid longer than 2 RTO intervals for the destination being removed.

Another case worth mentioning is illustrated below:



For this case, during the deletion of an IP address, an Abort MUST be ignored if the destination address of the Abort message is that of the destination being deleted.

4.2 Setting of the primary address

A sender of this option may elect to send this combined with a deletion request for an address. A sender MUST only send a set primary request to an address that is considered part of the association. In other words a sender MUST NOT bundle a set primary with an add of a new IP address unless the add request is to be processed BEFORE the set primary.

The request to set an address as the primary path is an option the

receiver MAY perform. It is considered a hint to the receiver of the best destination address to use in sending SCTP packets (in the requestors view). It is possible that the receiver may have other knowledge that it may act upon and NOT set the specified address as

primary. If a request arrives that asks the receiver to set an address as primary that does not exist, the receiver should NOT honor the request, leaving its existing primary address unchanged.

5. Security Considerations

The ADD/DELETE of an IP address to an existing association does provide an additional mechanism by which existing associations can be hijacked. Where the attacker is able to intercept and or alter the packets sent and received in an association the use of this feature MAY increase the ease at which an association may be overtaken. This threat SHOULD be considered when deploying a version of SCTP that use this feature. The IP Authentication Header [[RFC2402](#)] SHOULD be used when the threat environment requires stronger integrity protections, but does not require confidentiality. It should be noted that in the base SCTP specification [[RFC2960](#)], if an attacker is able to intercept and or alter packets, even without this feature it is possible to hijack an existing association, please refer to [Section 11 of RFC2960](#).

6. IANA considerations

Three REL-REQ Parmeter Types and three new SCTP Error Causes are being defined within this document.

7. Authors' Addresses

Randall R. Stewart
Cisco Systems, Inc.
8745 W. Higgins Road, Suite 200
Chicago, Ill 60631
USA

Tel: +1-815-477-2127
EMail: rrs@cisco.com

Qiaobing Xie
Motorola, Inc.
1501 W. Shure Drive, #2309
Arlington Heights, IL 60004
USA

Tel: +1-847-632-3028
EMail: qxie1@email.mot.com

Michael Tuexen
SIEMENS AG
Hofmannstr. 51
81359 Munich
Germany

Tel: +49-89-722-47210
EMail: Michael.Tuexen@icn.siemens.de

Ian Rytina
Ericsson Australia
37/360 Elizabeth Street
Melbourne, Victoria 3000
Australia

Tel: +61-3-9301-6164
EMail: ian.rytina@ericsson.com

8. References

- [RFC2960] R. R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. J. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and, V. Paxson, "Stream Control Transmission Protocol," [RFC 2960](#), October 2000.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [RFC 2026](#), October 1996.
- [RFC2119] Bradner, S. "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2402] S. Kent, R. Atkinson., "IP Authentication Header.", [RFC 2402](#), November 1998.
- [RELREQ] "Generic Method for Transmitting Reliable SCTP Control Chunks", [draft-stewart-relreq-sctp-sigtran-00.txt](#), work in progress.

This Internet Draft expires in 6 months from February, 2001

