

Network Working Group
Internet-Draft
Expires: November 14, 2002

R. Stewart
Cisco Systems, Inc.
M. Tuexen
Siemens AG
May 16, 2002

Stream Control Transmission Protocol (SCTP) IPv4 Address Scoping
draft-stewart-tsvwg-sctp-ipv4-00.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of 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.

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."

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>.

This Internet-Draft will expire on November 14, 2002.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

Stream Control Transmission Protocol [RFC2960](#) [5] provides transparent multi-homing to its upper layer users. This multi-homing is accomplished through the passing of address parameters in the initial setup message used by SCTP. In an IPv4 network addresses SHOULD NOT be passed without consideration of their routeability. This document defines considerations and enumerates general rules that an SCTP endpoint MUST use in formulating both the INIT and INIT-ACK chunks when including IPv4 addresses.

Table of Contents

1.	Introduction	3
2.	Conventions	4
3.	IPv4 address scoping	5
3.1	Classification of IPV4 addresses	5
3.2	Black-hole scenario	5
3.3	Address handling for INIT chunks	5
3.4	Address handling for INIT-ACK chunks	6
4.	Security considerations	7
	References	8
	Authors' Addresses	8
	Full Copyright Statement	9

1. Introduction

Stream Control Transmission Protocol [RFC2960](#) [5] provides transparent multi-homing to its upper layer users. This multi-homing is accomplished through the passing of address parameters in the initial setup message used by SCTP. In an IPv4 network addresses SHOULD NOT be passed without consideration of their routeability. This document defines considerations and enumerates general rules that an SCTP endpoint MUST use in formulating both the INIT and INIT-ACK chunks when including IPv4 addresses.

The emphasis in the rules laid out in this document are to prevent an SCTP endpoint from listing an IPv4 address that is not routeable to a peer endpoint. This will minimize black-hole conditions that may cause the unexpected failure of SCTP associations.

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 [RFC2119](#) [4].

3. IPv4 address scoping

3.1 Classification of IPV4 addresses

Several blocks of IP-addresses have been assigned by IANA for special use. See IANA-SPECIAL-IPV4 [[1](#)] for further details.

In this document the IPv4 addresses are divided into several different levels:

Level 0: Addresses unusable with SCTP: 0.0.0.0/8, 224.0.0.0/4, 198.18.0.0/24, 192.88.99.0/24.

Level 1: Loopback addresses: 127.0.0.0/8.

Level 2: Link-local addresses: 169.254.0.0/16.

Level 3: Private addresses: 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16.

Level 4: Global addresses.

Addresses of Level 0 MUST not be used

- o as a source address of a SCTP packet.
- o as a destination address of a SCTP packet.
- o within an address parameter of an INIT, INIT-ACK chunk.

3.2 Black-hole scenario

A black-hole condition is where some other host is using the same address. In a IPv4 network this COULD happen if an INIT was sent to a global address that listed private addresses. If the peer also has a separate private based addressing it MAY send a heartbeat to an internal peer using the address listed. This causes the internal peer to send an ABORT thus destroying the association.

The rules given in the next two sections for address handling will minimize the risk of having a black-hole condition.

3.3 Address handling for INIT chunks

When the ULP requests establishment of an SCTP association to a IPv4 destination address, the following considerations apply:

- o Let L be the level of the requested destination address. Therefore $L > 0$ holds.
- o The sender of the INIT chunk SHOULD include all of its addresses with level greater than or equal to L in the outgoing INIT chunk.
- o The sender of the INIT chunk SHOULD NOT include all of its addresses with level smaller than L in the outgoing INIT chunk.

Note that by listing both private and global addresses to a peer that does NOT have any global address the peer may find the senders global address unreachable. This is not a problem however since it would NOT cause a black-hole condition.

3.4 Address handling for INIT-ACK chunks

The receiver of an INIT will identify the relevant address level by examining the source address of the SCTP packet. In choosing addresses to place in the INIT-ACK the following considerations apply:

- o Let L be the level of the received source address of the INIT chunk. Therefore $L > 0$ holds.
- o The sender of the INIT-ACK chunk SHOULD include all of its addresses with level greater than or equal to L in the outgoing INIT-ACK chunk.
- o The sender of the INIT-ACK chunk SHOULD NOT include all of its addresses with level smaller than L in the outgoing INIT-ACK chunk.

Note that it is possible that a sender of an INIT incorrectly places addresses within its INIT. To protect against this the receiver of the INIT SHOULD examine carefully each address. If the level of an address listed is less than the level of the received source address, the address SHOULD be discarded and not put into the cookie parameter.

4. Security considerations

This document does not add any security risks other than those already found in [RFC2960](#) [5]

References

- [1] IANA, I., "Special-Use IPv4 Addresses", [draft-iana-special-ipv4-03](#) (work in progress), April 2002.
- [2] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.
- [3] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [5] Stewart, R., Xie, Q., Morneault, K., Sharp, C., Schwarzbauer, H., Taylor, T., Rytina, I., Kalla, M., Zhang, L. and V. Paxson, "Stream Control Transmission Protocol", [RFC 2960](#), October 2000.

Authors' Addresses

Randall R. Stewart
Cisco Systems, Inc.
8725 West Higgins Road
Suite 300
Chicago, IL 60631
USA

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

Michael Tuexen
Siemens AG
ICN WN CC SE 7
D-81359 Munich
Germany

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

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

