

Multihoming IPV6 Working group
Internet-Draft
Expires: January 16, 2005

L. Coene
Siemens
J. Loughney
Nokia
July 18, 2004

Multihoming: the SCTP solution
<[draft-coene-multi6-sctp-01.txt](#)>

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with [RFC 3668](#).

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 January 16, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This document describes the multihoming solution used in SCTP. It compares the SCTP solution with the goals set out in "Goals for IPv6 Site-Multihoming Architectures" [[11](#)]. The document also tries to answer the questions posed in "Things MULTI6 developers should think about" [[1](#)].

Internet-Draft

Multi6 SCTP solution

July 2004

Table of Contents

1.	INTRODUCTION	4
1.1	Terminology	4
2.	SCTP and Multi6 Goals(RFC3582)	5
2.1	Capabilities of IPv4 Multihoming	5
2.1.1	Redundancy	5
2.1.2	Load Sharing	5
2.1.3	Performance	5
2.1.4	Policy	5
2.1.5	Simplicity	5
2.1.6	Transport Layer Survivability	5
2.1.7	Impact on DNS	6
2.1.8	Packet filtering	6
2.2	Additional Requirements	6
2.2.1	Scalability	6
2.2.2	Impact on routers	6
2.2.3	Impact on Hosts	6
2.2.4	Interaction between Hosts and the routing system	6
2.2.5	Operation and Management	6
2.2.6	Cooperation between Transit Providers	7
2.2.7	Multiple solutions	7
3.	Answer to Multi6 solution Questions	8
3.1	On the wire behaviour	8
3.1.1	How does SCTP solve the multihoming problem	8
3.1.2	At what layer is SCTP applied to?	8
3.1.3	Why is this layer the correct one?	8
3.1.4	Does SCTP address mobility?	9
3.1.5	Does SCTP expand the size of a IP packet?	9
3.1.6	Does SCTP add additional latency?	9
3.1.7	Can SCTP negotiate the multihoming capabilities end-to-end during a connection?	9
3.1.8	Does SCTP change the way fragmenting is handled?	9
3.1.9	Implications of SCTP with layer2?	9
3.2	Identifiers and locators	9
3.2.1	Uniqueness	9
3.2.2	Does SCTP provide a split between identifier and locator?	10
3.2.3	What is the lifetime of a binding from locator to identifier?	10
3.2.4	How is the binding updated?	10
3.2.5	How does the host know its identity?	10

3.2.6	Can a host have multiple identities?	10
3.2.7	Mapping between locators and identifiers.	10
3.2.8	Does SCTP create an alternative DNS-like service? . .	10
3.2.9	Authentication & authorisation	11
3.2.10	Is the mechanism hierarchical?	11
3.2.11	Middlebox interactions.	11

3.2.12	Implications of SCTP for scoped addressing	11
3.3	Routing System interactions	11
3.3.1	Does SCTP change existing aggregation methods?	11
3.3.2	SCTP and new name space aggregation?	11
3.3.3	Are there any changes to ICMP error semantics?	11
3.4	Names service interactions	11
3.4.1	Relation of SCTP to DNS	11
3.4.2	Interaction of SCTP with 2-faced DNS.	12
3.4.3	Does SCTP require a centralized registration?	12
3.4.4	Has SCTP checked for DNS circular dependencies? . . .	12
3.4.5	What happens if the DNS server itself is multihomed? .	12
3.4.6	What additional load will be placed on DNS servers? .	12
3.4.7	Any upstream provider support required?	12
3.4.8	How do you debug connectivity?	12
3.5	Application concerns and backwards compatibility	12
3.5.1	What application/API changes are needed?	12
3.5.2	Is this backward compatible with IPv6?	13
3.5.3	Is this backward compatible with IPV4?	13
3.5.4	Can IPv4 devices take advantage of this solution? . .	13
3.5.5	What is the impact of SCTP on different types of sites?	13
3.5.6	What are the interactions with other middleboxes? . .	13
3.5.7	SCTP and referrals?	13
4.	Legal concerns	15
5.	Security considerations	16
6.	Acknowledgments	17
7.	References	17
	Authors' Addresses	18
	Intellectual Property and Copyright Statements	19

1. INTRODUCTION

SCTP is a transport protocol which among its features offers support for multihoming. The mechanism is described in detail in "[RFC2960](#)" [[2](#)]. A more general description of its uses can be found in "[RFC3257](#)" [[4](#)] and "SCTP multihoming Issues" [[3](#)].

1.1 Terminology

The terms are commonly identified in related work "[RFC2960](#)" [[2](#)], "[RFC3257](#)" [[4](#)] and "SCTP multihoming Issues" [[3](#)] .

[2.](#) SCTP and Multi6 Goals([RFC3582](#))

This chapter compare the features of SCTP with the goals set forth in "Goals for IPv6 Site-Multihoming Architectures" [[11](#)].

[2.1](#) Capabilities of IPv4 Multihoming

[2.1.1](#) Redundancy

If paths belonging to a single SCTP association are distinct through the network, SCTP will retain connectivity in case of physical, logical link, routing protocol, transit provider or Exchange failure.

[2.1.2](#) Load Sharing

Loadsharing is at present not implemented in SCTP. However applications may try to loadshare their traffic over the different paths by changing the primary path(= primary address) for each user data send to SCTP. This area is for further study.

[2.1.3](#) Performance

The endpoints using SCTP multihoming may be using the paths within the SCTP association on the performance of the paths through the network. This can be based on RT0, the long term congestion of a path, throughput, etc..

[2.1.4](#) Policy

No support for policy in SCTP multihoming.

[2.1.5](#) Simplicity

The SCTP solution is simple in that it only impact endpoints which want to actually use it, that it does not impact software anywhere else in the network. Given at least 2 addresses and cable and SCTP will do the job.

[2.1.6](#) Transport Layer Survivability

The SCTP multihoming solution is a transport layer solution. It checks every path within the association and changes to a different working path(=rehoming) if at any point during the lifetime of the association, a certain path fails. The association remains in service if at least a single path remains in service. The path change is transparent to the layers above SCTP.

[2.1.7](#) Impact on DNS

The SCTP multihoming solution does not depend on DNS for its operation. It requires only a single IP address of the remote peer.

[2.1.8](#) Packet filtering

Packet filtering will work without any additions. Only packets with incorrect source address(= source IP address used in packet is NOT the address of the interface on which the packet is sent) may be discarded by the packet filtering.

[2.2](#) Additional Requirements

[2.2.1](#) Scalability

The scalability of the SCTP solution depends on getting multiple addresses for the 2 endpoints. The 2 endpoints are the only ones to keep state about the different paths between the endpoint, so the solution will scale up very easily when adding new endpoints. The number of paths available to a endpoint is equal to the number of IP addresses assigned to the endpoint.

[2.2.2](#) Impact on routers

None

[2.2.3](#) Impact on Hosts

Host needs to be fitted with multiple IP addresses and should be using SCTP to take advantage of the multihoming.

[2.2.4](#) Interaction between Hosts and the routing system

The interaction needed between the host and the gateway router is described in "SCTP multihoming Issues" [\[3\]](#). The rest of the routing system is not involved.

[2.2.5](#) Operation and Management

Monitoring of the SCTP multihoming is possible as SCTP uses multiple addresses, so monitoring based on addresses should be possible. The configuration needed for SCTP is restricted to provide the SCTP software and provisioning multiple addresses and a correct routing table on the host.

[2.2.6](#) Cooperation between Transit Providers

None required.

[2.2.7](#) Multiple solutions

SCTP is a solution in itself. It does not prevent other solutions to

work on the same problem. If a solution works on the underlying layer of SCTP, SCTP will view itself as being singled homed, albeit the underlying solution is actually multihomed. Solutions defined in other transport protocols cannot be used by SCTP.

[3.](#) Answer to Multi6 solution Questions

[3.1](#) On the wire behaviour

[3.1.1](#) How does SCTP solve the multihoming problem

A general overview of the solution can be found in "SCTP multihoming Issues" [[3](#)] in paragraph 2.1. The detail message elements with their syntax and semantics can be found in "[RFC2960](#)" [[2](#)].

The present solution allows for the exchange of the multiple addresses of each endpoint at the start of the association. Once the association has been set up, then heartbeat messages are used to check the reachability of each address. If the reachability test fails(because the heartbeat went unanswered for X times(with $X = 1..n$)), then that particular address is deemed not reachable and will NOT be used to send data on. If the reachability test is successful, then the address may be used to send data to. If changeover is requested(by the application or by SCTP itself), then this address will be used to send data on. No IP address can be added or deleted from to association once it has been setup.

A extension to SCTP is in the works which allows a already active SCTP association to add or delete a IP address [[5](#)] to a association.(Thus new "paths" are added or removed). The association will use this addresses based on the reachability information obtained by the use of the SCTP heartbeat just as mentioned above.

An additional extension allows to secure this sort of ADDDELIP msg exchange via the use of Purpose Built keys(PBK) [[7](#)]. If a more secure association is required, then TLS or IPSEC are recommended.

[3.1.2](#) At what layer is SCTP applied to?

It is a layer 4 solution(=transport layer).

[3.1.3](#) Why is this layer the correct one?

Every IP address corresponds to a single path through the network. Each path can have different delay, loss and so forth, characteristics. The congestion control algorithm depends on some of this info to perform its congestion control. Thus the transport layer has to measure this himself so that it internal variables are updated. Otherwise the info may be distributed and/or duplicated accross multiple layers. Therefore decisions about using or changing of path are taken by the transport layer.

Internet-Draft

Multi6 SCTP solution

July 2004

[3.1.4](#) Does SCTP address mobility?

SCTP does NOT solve the "where is the endpoint?" problem. It assumes that the location of the mobile user is known, because it has a IP address(which is the locator). It will try to setup a association with that IP address and exchange IP addresses between the two endpoints of the association at the start of the association(as in [RFC 2960](#)) or during the association lifetime(ADDELIP) [[5](#)].

SCTP does solve the "handover" problem, namely the problem of moving the traffic through the association from one IP address to another IP address. The new address can be the result of a DHCP request by the lower layers, renumbering in IPv6...

Mobility in SCTP is only a byproduct of putting in multihoming in SCTP. SCTP can be used for mobility if add or delete a IP address [[5](#)] is implemented.

[3.1.5](#) Does SCTP expand the size of a IP packet?

SCTP contains its own header just as other transport protocols. It comes in the place of the header of other transport protocols.

[3.1.6](#) Does SCTP add additional latency?

No.

[3.1.7](#) Can SCTP negotiate the multihoming capabilities end-to-end during a connection?

Yes, see add or delete a IP address [[5](#)].

[3.1.8](#) Does SCTP change the way fragmenting is handled?

No. It leaves IP fragmentation alone and uses its own fragmenting and reassembly code.

[3.1.9](#) Implications of SCTP with layer2?

None.

[3.2](#) Identifiers and locators

[3.2.1](#) Uniqueness

Not needed.

[3.2.2](#) Does Sctp provide a split between identifier and locator?

Not really. Sctp uses the IP address as the locator but the identifier is assumed to be implicit. Sctp do NOT exchange any identifier between the peer endpoints, only IP addresses are exchanged. The association ID used between the application and Sctp may be regarded as the identifier, but this identifier is completely local.

Sctp allows endpoints to be addressed by multiple IP addresses, the concept of an Sctp endpoint is much broader than in TCP. In this way, a Sctp association can use multiple interfaces and multiple addresses for upper layer protocols.

[3.2.3](#) What is the lifetime of a binding from locator to identifier?

The lifetime of a binding from locator to identifier is equal to the lifetime of a Sctp association([RFC 2960](#)) or less(in case of ADDELIP).

[3.2.4](#) How is the binding updated?

A control message(called a chunk in Sctp) is used to exchanged the IP addresses between the endpoints. It can be done at setup of the association(see [RFC 2960](#)) or during the lifetime of the association(see ADDELIP).

[3.2.5](#) How does the host know its identity?

The hosts determines which IP addresses it is going to use with the association, thus forming its identity implicit. The easiest way is to bind to all present interfaces, but the application above Sctp can decide to use all or part of the addresses present in the host.

[3.2.6](#) Can a host have multiple identities?

The host can have multiple identifiers, by having distinct sets of

addresses for each of the identifiers.

[3.2.7](#) Mapping between locators and identifiers.

The mapping is done within SCTP and it is really implementation dependant. The identity itself goes never over the wire.

[3.2.8](#) Does SCTP create an alternative DNS-like service?

No

Coene & Loughney

Expires January 16, 2005

[Page 10]

Internet-Draft

Multi6 SCTP solution

July 2004

[3.2.9](#) Authentication & authorisation

SCTP uses Purpose Built keys to authenticate the bindings. See add or delete a IP address [\[5\]](#) , Authenticated Chunks for Stream Control Transmission Protocol (SCTP) [\[6\]](#) and Purpose built keys [\[7\]](#)..

[3.2.10](#) Is the mechanism hierarchical?

No.

[3.2.11](#) Middlebox interactions.

Middleboxes are NOT part of the SCTP solution. If middleboxes have to rewrite information in the packets(especially in SCTP), they have to be updated for SCTP. Middleboxes will in general limit the use of multihoming via SCTP, because all traffic(=all paths) have to pass through the middlebox, thus creating a single point of failure. For further information see "SCTP multihoming Issues" [\[3\]](#).

[3.2.12](#) Implications of SCTP for scoped addressing

If the address is reachable, the communication will get through. It is however suggested to use globally scoped addresses first and descend from there. It is suggested not to mix global, link or site scope addresses within a single association.

[3.3](#) Routing System interactions

[3.3.1](#) Does SCTP change existing aggregation methods?

No.

[3.3.2](#) SCTP and new name space aggregation?

Not needed.

SCTP does NOT introduce a new naming space, thus no aggregation of a new name space is needed.

[3.3.3](#) Are there any changes to ICMP error semantics?

No.

[3.4](#) Names service interactions

[3.4.1](#) Relation of SCTP to DNS

SCTP has no direct interface to DNS. It however uses the result that

come back from a DNS query by the application software on the host, to setup a association to the peer with the returned IP address. If DNS returns a non-reachable address, then SCTP will not be able to reach the peer. If the DNS returns a reachable address, then SCTP can start its association and figure out if the peer is multihomed via a appropriate message exchange. It already knows for his own endpoint if it is multihomed, yes or no.

[3.4.2](#) Interaction of SCTP with 2-faced DNS.

SCTP has no direct interaction with DNS, so it does not need direct interaction with 2 faced DNS either.

[3.4.3](#) Does SCTP require a centralized registration?

NO.

[3.4.4](#) Has SCTP checked for DNS circular dependencies?

As SCTP does not rely on the DNS for any functionality of its multihoming solution, no dependency exists on DNS and as a result, no circular dependencies are possible.

[3.4.5](#) What happens if the DNS server itself is multihomed?

No dependency exists on the DNS, so DNS multihoming is invisible to SCTP in the host. If naturally the communication between the DNS resolver and the DNS server itself uses SCTP then there is still no problem as only SCTP internal mechanism are used for doing the multihoming.

[3.4.6](#) What additional load will be placed on DNS servers?

None.

[3.4.7](#) Any upstream provider support required?

None.

[3.4.8](#) How do you debug connectivity?

No present day tools need to be enhanced.

[3.5](#) Application concerns and backwards compatibility

[3.5.1](#) What application/API changes are needed?

The application software has to be ported on a socket api very similar to the already present socketapi of TCP. The application

will use multihoming unknowingly as No specific API change is needed to activate multihoming on the own endpoint.

If the application wishes to actively control the multihoming of the association, new socketapi [\[8\]](#) options exists to do that but then this must be considered as adding new features to applications, not porting old applications.

It should be noted that SCTP is a connection-oriented, congestion control protocol. Therefore, traffic running over UDP is not considered at this time. A UDP style socket is present in SCTP but requires more changes to the application. UDP traffic can also use the partial reliability feature of SCT [\[9\]](#) if required.

[3.5.2](#) Is this backward compatible with IPv6?

Yes, it is even backward compatible with IPv4. The SCTP association can be multihomed across a ipv4 and ipv6 network(meaning the single association will use Ipv4 and Ipv6 address within the same association). No change is require to present IPv6 code.

[3.5.3](#) Is this backward compatible with IPV4?

Yes. see also paragraph above.

[3.5.4](#) Can IPv4 devices take advantage of this solution?

Yes, see also paragraphs above.

[3.5.5](#) What is the impact of SCTP on different types of sites?

None. SCTP does not need to know how big the sites should be. It only depends on having IP addresses and being informed by the IP layer in its own host of new or retracted IP addresses(example: Ad-hoc sites). Other hosts or routers are not involved.

[3.5.6](#) What are the interactions with other middleboxes?

Middleboxes which do not change or drop SCTP chunks, do not impact the multihoming. Only NAT boxes have to do their work in the INIT and INIT-ACK chunks as addresses are transported in those chunks. If ADDELIP is used, the the add and delete IP chunks must also be screwed around by the NAT box. The NAT box will very likely be the single point-of-failure in the association.

[3.5.7](#) SCTP and referrals?

If a referral is a new IP address, then the application can setup a

new association via SCTP with the new endpoint and be multihomed again(if the new endpoint is also multihomed).

None.

[5](#). Security considerations

SCTP has mechanisms for reducing the risk of blind denial-of-service attacks and/or masquerade attacks. If such measures are required by the applications, then it is advised to check the SCTP applicability statement "[RFC3257](#)" [\[4\]](#) for guidance on this issue.

Additional work on securing the ADDELIP [\[5\]](#) via the use of Purpose Built keys(PBK) [\[6\]](#) in SCTP is going on.

6. Acknowledgments

The authors wish to thank x, Y, and many others for their invaluable comments.

7 References

- [1] Lear, E., "Things MULTI6 Developers should think about", Draft in progress , May 2004.
- [2] 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.
- [3] Coene, L., "SCTP multihoming issues", Draft in progress , June 2003.
- [4] Coene, L., "Stream Control Transmission Protocol Applicability statement", [RFC 3257](#), April 2002.
- [5] Stewart, R., Ramalho, M., Xie, Q., Tuxen, M., Rytina, I., Belinchon, M. and P. Conrad, "Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration", Draft in progress , September 2003.
- [6] Tuxen, M. and R. Stewart, "Authenticated Chunks for Stream Control Transmission Protocol (SCTP)", Draft in progress , October 2003.
- [7] Bradner, S., Mankin, Allison. and J. Schiller, "A Framework for Purpose-Built Keys (PBK)", Draft in progress , June 2003.
- [8] Stewart, R., Xie, Q., Yarroll, L., Wood, J., Poon, K., Fujita, K. and M. Tuxen, "Sockets API Extensions for Stream Control Transmission Protocol (SCTP)", Draft in progress , August 2003.
- [9] Stewart, R., Ramalho, M., Xie, Q., Tuxen, M. and P. Conrad,

"SCTP Partial Reliability Extension", Draft in progress ,
January 2004.

[10] Stewart, R., Tuxen, M. and G. Camarillo, "Stream Control
Transmission Protocol (SCTP) Security Threats", Draft in
progress , April 2004.

[11] Abley, J., Black, B. and V. Gill, "Goals for IPv6
Site-Multihoming Architectures ", [RFC 3582](#), August 2003.

Coene & Loughney

Expires January 16, 2005

[Page 17]

Internet-Draft

Multi6 SCTP solution

July 2004

Authors' Addresses

Lode Coene
Siemens
Atealaan 32
Herentals 2200
Belgium

Phone: +32-14-252081
EMail: lode.coene@siemens.com

John Loughney
Nokia
Itämerenkatu 11-13
Espoo 00180
Finland

Phone: +???????
EMail: john.loughney@nokia.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

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