

SCTP

STATEMENT OF COMPLIANCE

STATEMENT OF COMPLIANCE

Copyright

© Ericsson AB 2009-2016. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.



Contents

1	General	1
1.1	Introduction	1
1.2	Terms	1
1.3	Concept	2
2	Compliance Lists	3
2.1	IETF - Stream Control Transmission Protocol	3
3	Notes	13
	Reference List	25



SCTP



1 General

1.1 Introduction

This document describes to what extent the Ericsson SCTP signaling component conforms with the IETF RFC 4960, see Reference [1].

In the remainder of the document, we refer to the RFC 4960, see Reference [1].

The difference between RFC 2960 and RFC 4960 can be found in Reference [4]

1.2 Terms

CP	Common Parts
CWR	Congestion Window Reduced
DNS	Domain Name System
ECN	Explicit Congestion Notification
ECNE	Explicit Congestion Notification Echo
ESP	Encapsulating Security Payload
IETF	Internet Engineering Task Force
IP	Internet Protocol
HB	Heartbeat Request (SCTP chunk type)
MIS	Maximum Incoming Streams
MTU	Maximum Transmission Unit
OS	Outbound Streams
PMTU	Path MTU
RFC	Request For Comments
RTO	Retransmission Time-Out
RTT	Round-Trip Time
RWND	Receiver Window
SCTP	Stream Control Transmission Protocol



TCB	Transmission Control Block
TCP	Transmission Control Protocol
TLV	Type-Length-Value (coding format)
TSN	Transmission Sequence Number
ULP	Upper Layer Protocol (SCTP user in this document)
V-Tag	Verification Tag of SCTP packet
VPN	Virtual Private Network

1.3 Concept

The terms that are used are:

C	The Ericsson signaling component complies with the specified section in RFC 4960, Reference [1].
N	The Ericsson signaling component does not comply with the specified section in RFC 4960 Reference [1].
P	The Ericsson signaling component complies partly with the specified section in RFC 4960 Reference [1].
-	There is nothing to implement in the referred section (always placed in column C).



2 Compliance Lists

2.1 IETF - Stream Control Transmission Protocol

2.1.1 RFC 4960

Table 1 RFC 4960

Reference		C	N	P	Comments
1	Introduction	-			
1.1	Motivation	-			
1.2	Architectural View of SCTP	X			
1.3	Key terms	X			Note 2 Note 3 Note 116
1.4	Abbreviations	-			
1.5	Functional View of SCTP	X			
1.5.1	Association Startup and Takedown	X			Note 4
1.5.2	Sequenced Delivery within Streams	X			
1.5.3	User Data Fragmentation	X			
1.5.4	Acknowledgements and Congestion Avoidance	X			
1.5.5	Chunk bundling	X			Note 5 Note 64
1.5.6	Packet Validation	X			Note 6
1.5.7	Path management	X			Note 2
1.6	Serial Number Arithmetic	X			
1.7	Changes From RFC 2960	-			
2	Conventions	-			Note 1
3	SCTP packet Format	X			Note 7
3.1	SCTP Common Header Field Descriptions	X			Note 8
3.2	Chunk Field Descriptions	X			Note 9 Note 10



Table 1 RFC 4960

Reference		C	N	P	Comments
3.2.1	Optional/Variable-length Parameter Format	X			
3.2.2	Reporting of Unrecognized Parameters	X			Note 65
3.3	SCTP Chunk Definitions	-			Note 92
3.3.1	Payload Data (DATA) (0)	X			
3.3.2	Initiation (INIT) (1)	X			Note 11 Note 12 Note 7 Note 66 Note 67 Note 101
3.3.2.1	Optional/Variable-Length Parameters in INIT			X	Note 11 Note 12 Note 14 Note 15 Note 69
3.3.3	Initiation Acknowledgement (INIT ACK) (2)			X	Note 16 Note 70 Note 12 Note 65
3.3.3.1	Optional or Variable Length Parameters	X			
3.3.4	Selective Acknowledgement (SACK) (3)	X			Note 89
3.3.5	Heartbeat Request (HEARTBEAT) (4)	X			
3.3.6	Heartbeat Acknowledgement (HEARTBEAT ACK) (5)	X			
3.3.7	Abort Association (ABORT) (6)	X			Note 87
3.3.8	Shutdown Association (SHUTDOWN) (7)	X			



Table 1 RFC 4960

Reference		C	N	P	Comments
3.3.9	Shutdown Acknowledgement (SHUTDOWN ACK) (8)	X			
3.3.10	Operation Error (ERROR) (9)	X			Note 17
3.3.10.1	Invalid Stream Identifier (1)	X			
3.3.10.2	Missing Mandatory Parameter (2)	X			
3.3.10.3	Stale Cookie Error (3)	X			Note 18
3.3.10.4	Out of Resource (4)	X			
3.3.10.5	Unresolvable Address (5)	X			
3.3.10.6	Unrecognized Chunk Type (6)	X			
3.3.10.7	Invalid Mandatory Parameter (7)	X			
3.3.10.8	Unrecognized Parameters (8)	X			
3.3.10.9	No User Data (9)	X			
3.3.10.10	Cookie Received While Shutting Down	X			
3.3.10.11	Restart of an Association with New Addresses (11)	X			
3.3.10.12	User-Initiated Abort (12)	X			Note 72
3.3.10.13	Protocol Violation (13)	X			Note 73 Note 74
3.3.11	Cookie Echo (COOKIE ECHO) (10)	X			Note 7
3.3.12	Cookie Acknowledgement (COOKIE ACK) (11)	X			Note 7
3.3.13	Shutdown Complete (SHUTDOWN COMPLETE) (14)	X			
4	SCTP Association State Diagram	X			Note 19
5	Association Initialization	X			Note 7 Note 4
5.1	Normal Establishment of an Association	X			Note 20 Note 21 Note 22 Note 96
5.1.1	Handle Stream Parameters	X			Note 75

Table 1 RFC 4960

Reference		C	N	P	Comments
5.1.2	Handle Address Parameters	X			Note 14 Note 23
5.1.3	Generating State Cookie	X			Note 25 Note 76 Note 96
5.1.4	State Cookie Processing	X			Note 7
5.1.5	State Cookie Authentication	X			Note 25 Note 7 Note 76
5.1.6	An Example of Normal Association Establishment	-			
5.2	Handle Duplicate or unexpected INIT, INIT ACK, COOKIE ECHO and COOKIE ACK	X			
5.2.1	Handle Duplicate INIT in COOKIE-WAIT or COOKIE-ECHOED States (Item B)	X			Note 77
5.2.2	Unexpected INIT in States Other than CLOSED, COOKIE-ECHOED, COOKIE-WAIT and SHUTDOWN-ACK-SENT	X			Note 77
5.2.3	Unexpected INIT ACK	X			
5.2.4	Handle a COOKIE ECHO when a TCB exists	X			Note 26
5.2.4.1	An Example of an Association Restart	-			
5.2.5	Handle Duplicate COOKIE ACK	X			
5.2.6	Handle Stale COOKIE Error	X			Note 27
5.3	Other Initialization Issues	-			
5.3.1	Selection of V-Tag Value	X			Note 3



Table 1 RFC 4960

Reference		C	N	P	Comments
5.4	Path Verification	X			Note 2 Note 28 Note 29 Note 7 Note 88 Note 94
6	User Data Transfer	X			Note 7 Note 30 Note 4 Note 78
6.1	Transmission of DATA Chunks	X			Note 2 Note 5 Note 31 Note 32 Note 61 Note 79 Note 80 Note 81 Note 98 Note 100 Note 105 Note 107 Note 108



Table 1 RFC 4960

Reference		C	N	P	Comments
6.2	Acknowledgement on Reception of DATA Chunks	X			Note 33 Note 36 Note 37 Note 82 Note 83 Note 84
6.2.1	Processing a Received SACK	X			
6.3	Management Retransmission Timer	X			Note 34
6.3.1	RTO Calculation	X			Note 34
6.3.2	Retransmission Timer Rules	X			
6.3.3	Handle T3-rtx Expiration	X			
6.4	Multi-homed SCTP Endpoints	X			Note 2 Note 35 Note 38 Note 91 Note 93 Note 102
6.4.1	Failover from Inactive Destination Address	X			Note 2 Note 38
6.5	Stream Identifier and Stream Sequence Number	X			
6.6	Ordered and Unordered Delivery	X			
6.7	Report Gaps in Received DATA TSNs			X	Note 106
6.8	CRC32c Checksum Calculation	X			Note 117
6.9	Fragmentation and Reassembly	X			Note 30
6.10	Bundling	X			
7	Congestion Control	X			Note 2
7.1	SCTP Differences from TCP Congestion Control	X			Note 39
7.2	SCTP Slow-Start and Congestion Avoidance	X			Note 39



Table 1 RFC 4960

Reference		C	N	P	Comments
7.2.1	Slow-Start			X	Note 39 Note 60 Note 85 Note 113
7.2.2	Congestion Avoidance	X			Note 39 Note 43 Note 60 Note 103
7.2.3	Congestion Control	X			Note 39
7.2.4	Fast Retransmit on Gap Reports	X			Note 62 Note 104 Note 111 Note 112
7.3	Path MTU Discovery	X			Note 40
8	Fault Management	-			
8.1	Endpoint Failure Detection	X			Note 2 Note 38 Note 41 Note 115
8.2	Path Failure Detection	X			Note 2 Note 42 Note 44 Note 41 Note 114



Table 1 RFC 4960

Reference		C	N	P	Comments
8.3	Path Heartbeat	X			Note 2 Note 45 Note 46 Note 47 Note 95 Note 97 Note 115 Note 118
8.4	Handle “Out of the blue” Packets	X			Note 48 Note 90
8.5	Verification Tag	X			
8.5.1	Exceptions in Verification Tag Rules	X			
9	Termination of Association	X			
9.1	Abort of an Association	X			
9.2	Shutdown of an Association	X			Note 49
10	Interface with Upper Layer	-			
10.1	ULP-to-SCTP	-			Note 45 Note 50 Note 51 Note 63
10.2	SCTP-to-ULP	-			Note 52 Note 99
11.	Security Considerations	-			
11.1	Security Objectives	-			
11.2	SCTP Responses To Potential Threats	-			
11.2.1	Countering Insider Attacks	-			
11.2.2	Protecting against Data Corruption in the Network	X			Note 53
11.2.3	Protecting Confidentiality	-			Note 54



Table 1 RFC 4960

Reference		C	N	P	Comments
11.2.4	Protecting against Blind Denial of Service Attacks	-			
11.2.4.1	Flooding	X			Note 110
11.2.4.2	Blind Masquerade	-			
11.2.4.3	Improper Monopolization of Services	-			
11.3	SCTP Interactions with Firewalls	X			
11.4	Protection of Non-SCTP-Capable Hosts	X			Note 109
12	Network Management Considerations			X	Note 57
13	Recommended Transmission Control Block (TCB) Parameters	-			Note 56
13.1	Parameters Necessary for the SCTP Instance	X			Note 56
13.2	Parameters Necessary per Association (TCB)	X			Note 56
13.3	Per Transport Address Data	X			Note 56
13.4	General Parameters Needed	X			Note 56
14	IANA Considerations	-			
14.1	IETF-defined Chunk Extension			X	Note 92
14.2	IETF-defined Chunk Parameter Extension	-			
14.3	IETF-defined Additional Error Causes	-			
14.4	Payload Protocol Identifiers	-			
14.5	Port Numbers Registry	-			
15	Suggested SCTP Protocol Parameter Values			X	Note 58
16	Acknowledgements	-			
Appendix A. Explicit Congestion Notification		X			Note 59
Appendix B. CRC32c Checksum Calculation		X			
Appendix C. ICMP Handling		X			Note 86
References		-			
Normative references		-			
Informative references		-			





3 Notes

- Note 1** The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT" present mandatory rules. Words "SHOULD", "SHOULD NOT", "RECOMMENDED" mark recommendations. Finally, "MAY", and "OPTIONAL" identifies optional points or alternatives for implementation. All recommendations (SHOULD, SHOULD NOT, RECOMMENDED) in the RFC 4960 are considered as rules (MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT), that is, if the Ericsson SCTP signaling component does not implement an item in a section of RFC 4960 that is stated as recommendation, then it is considered not compliant with this section. Notes explain the chosen alternatives, debatable and tricky places in RFC, reasoning of partial compliance or non-compliance.
- Note 2** BASE SCTP differentiate several IP paths which can be used to reach one destination IP address (see Reference [2] for details). BASE SCTP has active IP path (which includes both active local and remote IP addresses) instead of active destination transport address. Idleness is controlled for all IP paths instead of just destination addresses. Primary path is IP path with configurable local primary IP address. Some destination address is considered as inactive if all IP paths which are linked with the destination address are inactive and vice versa. Network monitoring strategy depends on a configuration option.
- Note 3** Only 24 bits are randomly generated in V-Tag value. 8 bits are reserved for internal needs. However, BASE SCTP layer guarantee unique V-Tag value for each association in scope of distributed stack.
- Note 4** An ASSOCIATE primitive is always needed to open an association (that is SEND primitive can not be used to open an association with default parameters as suggested in the IMPLEMENTATION NOTE).
- Note 5** BASE SCTP has an option to configure delays for DATA chunks bundling before transmission action. If delay is zero, BASE SCTP activates bundling automatically if DATA chunks bundling is enabled and more than one DATA chunk is available in the outgoing queue in time of sending. It is a frequent case if traffic is highly intensive.



- Note 6** Checksum calculation can be deactivated in module configuration.
- Note 7** BASE SCTP is able to bundle DATA with DATA chunks and control chunks with DATA chunks or with other control chunks, but not all allowed cases are supported. For example COOKIE_ECHO, COOKIE_ACK, ERROR, HEARTBEAT and HEARTBEAT_ACK chunks are never bundled with DATA chunks. See **Reference [2]** for more details.
- Note 8** BASE SCTP accepts zero port during endpoint initialization. In this case port number is assigned automatically. Range of available ports can be configured. See Reference [3] for more details.
- Note 9** BASE SCTP implementation is not able to process packets with several bundled chunks which are padded incorrectly.
- Note 10** ECNE and CWR chunk types are used by BASE SCTP.
- Note 11** If OS/MIS field in received INIT chunk carries value 0, the "association request" is aborted.
- Note 12** "Initial TSN" field in outbound INIT/INIT_ACK chunk is set to the same value as "Initiate Tag" field.
- Note 14** "Supported Address Types" parameter is coded in outbound INIT chunks indicating that IPv4, IPv6 and/or DNS Host Name addresses are supported (Address Type = 0x5, 0x6 and 0xB respectively). If IPv4 addresses were included in configuration, BASE SCTP adds IPv4 to "Supported Address Types". If IPv6 addresses were included in configuration, BASE SCTP adds IPv6 to "Supported Address Types". If the support of DNS Host Names in INIT and INIT_ACK chunks was configured (see "DnsSupport" parameter in Reference [3] for more details), BASE SCTP adds DNS Host Name to "Supported Address Types".
- Note 15** "Cookie Preservative" parameter is always coded in outbound INIT chunks (carrying value 0).
- Note 16** BASE SCTP generates ABORT chunk on incoming INIT_ACK chunk in case of faults during negotiation about number of incoming or outgoing streams.
- Note 17** Only one "Error Cause" is coded in outbound ERROR chunks.



- Note 18** "Measure of Staleness" is always zero in "Stale Cookie Error".
- Note 19** BASE SCTP differentiates ESTABLISHED state with regular heartbeats and ESTABLISHED state with activated path probing. However, it does not have any influence on functional behavior of SCTP layer.
- Note 20** Only one "Error Code" is coded in outbound ABORT chunks.
- Note 21** When an association is established BASE SCTP reports COMMUNICATION UP indication which can be used by ULP as trigger to activate transmission of data.
- Note 22** BASE SCTP includes "Cause of ABORT" in ABORT chunk only if the reason of fault is discovered exactly.
- Note 23** IPv4 and IPv6 addresses received in INIT/INIT_ACK chunks are discarded if the IP layer does not support that IP version. DNS Host Name addresses received in INIT/INIT_ACK chunks are discarded if BASE SCTP does not support DNS Host Name resolution, that is a configurable option (see Reference [3]).
- Note 25** BASE SCTP uses several random and periodically changed keys to generate MAC and encode TCB. "Key change period" is configurable value, see Reference [3]. Encoding algorithm is based on MD5 with additional enhancements.
- Note 26** When a "peer restart" is detected, all received user data still not delivered to the user and all sent DATA chunks still not acknowledged from the peer are discarded (dropped).
- Note 27** When a "Stale COOKIE Error" Error Code is received in an ERROR chunk in the "COOKIE-ECHOED" state, the new INIT chunk is sent to the remote endpoint. The new INIT contains Cookie Preservative parameter which requests an extension cookie life time (that is option "3" has been selected from the three different alternatives specified in RFC 4960).
- Note 28** Each IP path is verified during association start up.
- Note 29** Hb.Max.Burst configuration parameter can be reconfigured but the new value will only be used by new associations.



- Note 30** BASE SCTP is able to process messages larger than incoming buffer space (`a_rwnd`, see Reference [1]) or CP message buffer size. In this case partial delivery between BASE SCTP and SCTP user is activated automatically. SCTP user may use partial delivery for outgoing messages if it is necessary but it shall be able to receive and concatenate incoming message parts. See Reference [2] for more information about partial delivery.
- Note 31** Single retransmission timer per IP path is used.
- Note 32** Bundling of DATA chunks is enabled via "Bundling Status" configuration parameter in BASE SCTP (see **Reference [3]** for more details).
- Note 33** Delay for SACK and SACK frequency are configurable options, see Reference [3] for more details. SACK chunk can be reported for each incoming packet if there are outgoing DATA chunks which can be bundled with SACK.
- Note 34** Separate RTO is calculated for each IP path.
- Note 35** BASE SCTP does not vary IP path for SACK chunk which is reported for duplicated DATA.
- Note 36** The maximum delay for generating an acknowledgement is configurable between 0-500 milliseconds. See "SACK Timer" parameter in Reference [3] for more details.
- Note 37** ULP is not allowed to request a SEND operation without any user data to be sent (so no DATA chunks with "no user data" are ever sent). BASE SCTP drops the SEND request from ULP if it does not include some payload. Also BASE SCTP reports `SEND_FAILURE_IND` to the user.
- Note 38** There are several styles to alternate IP path for the retransmission. The style is a configurable option. See "Path Selection Adjustment" parameter in Reference [3] for more details.
- Note 39** BASE SCTP keeps a separate congestion control parameter set for each of the IP paths instead of the destination addresses.



- Note 40** PMTU discovery is a configurable option in BASE SCTP (see Reference [3]). In the case of disabled “PMTU Discovery” common IPv4 PMTU value is shared between all IPv4 paths and common IPv6 PMTU value is shared between all IPv6 paths in scope of one association. In this case BASE SCTP receives PMTU values from the appropriate Configuration Group of the configuration (see “PMTU” and “IPv6 PMTU” parameters in Reference [3] for more details) and shares it between related endpoints, associations and paths.
- Note 41** BASE SCTP is able to detect failures of IP paths, destination addresses and remote endpoints.
- Note 42** Path.Max.Retrans parameter has more complex meaning. Interpretation depends on the chosen path selection algorithm. See Reference [3] for more details.
- Note 43** DATA chunks acknowledged by Gap Ack Blocks (received in a SACK chunk) are taken into account to increase the partial_bytes_acked value of a destination transport address.
- Note 44** Successfully retransmitted DATA chunk does not reset path error counter if it has been sent through two or more different IP paths while transmitting and retransmitting. If the path selection scheme A (“PEER”) is used and SACK chunk was received for DATA chunk successfully retransmitted through the same SCTP path but another IP path (from another local IP address to the same remote IP address) peerT3counter will be reset.
- Note 45** All HEARTBEAT-related primitives are offered to the Management interface (but not to the ULP interface). However, ULP can configure some of association properties which are related to the heartbeat functionality. All configurable parameters can be found in SCTP FS Reference [2].
- Note 46** BASE SCTP monitors each IP path by heartbeats.
- Note 47** There is an option to smooth HB.Interval in a module configuration. See Reference [3] for more details.
- Note 48** BASE SCTP registers SCTP protocol identifier on IP layer, it allows to drop packets from and/or to multicast addresses on IP layer. If such packets are not filtered on IP layer, BASE SCTP processes it as OOTB packets or drops due to incorrect checksum value.
- Note 49** T5-shutdown-guard timer has not been implemented.



- Note 50** External interfaces present the set of primitives/messages from or to BASE SCTP module. All interfaces are non-blocking calls only. Some mandatory parameters are optional for BASE SCTP API, see Reference [2]. INITIALIZE and DESTROY interfaces are available through SCTP_INITIALIZE_req/conf and SCTP_DESTROY_req messages. Registered endpoint is referenced by endpoint ID. ASSOCIATE is presented by SCTP_ASSOCIATE_req/conf pair. Registered association is referenced by association ID. SHUTDOWN and ABORT are SCTP_SHUTDOWN_req/conf and SCTP_ABORT_req messages correspondingly. RECEIVE interface is partly covered by SCTP_RESEND_req message, but it is not necessary to use such request due to BASE SCTP deliveries incoming DATA automatically with help of SCTP_DATA_ARRIVE_ind. STATUS interface is available through the set of primitives and orders, like SCTP_STATUS_req/conf, SCTP_INFO_req/conf, MM_ORDER_req/conf, MM_XORDER_req/conf. Interfaces for layer (re)configuration are covered by management orders and messages.
- Note 51** 'life time' parameter has not been implemented for SEND operation.
- Note 52** SHUTDOWN COMPLETE notification is presented by COMMUNICATION LOST notification message with suitable reason code.
- Note 53** SCTP-AUTH is not supported.
- Note 54** ESP is not part of SCTP, but ESP still could be used by platforms adaptation.
- Note 56** Structures are organized in a different way, but functional behavior is kept according to this document.
- Note 57** All available network management objects (statistics, MM operations, configuration) are described in Reference [2] and Reference [3] for BASE SCTP. There are management objects described in MIB [RFC3873] and not implemented for BASE SCTP.
- Note 58** CFD states recommended values for the protocol parameters. They can differ from recommended by RFC 4960. The list of parameters have differ recommended value: "Minimum RTO", "Maximum RTO", "Initial Retransmission Time-Out", "Assoc.Max.Rtx", "Path.Max.Rtx", "Maximum Shutdown Retransmissions"(not specified in RFC4960), "SACK Timer". See Reference [3] for more details.



- Note 59** BASE SCTP implementation is designed to receive ECNE chunk bundled in packet before SACK chunk or sent before SACK chunk. See Reference [2] for more details.
- Note 60** The first rule is applied for cwnd reduction in case of path's idleness. If path is not used during RTO period, cwnd is adjusted to $\max(\text{cwnd} / 2, \min(4 * \text{MTU}, \max(2 * \text{MTU}, 4380 \text{ bytes})))$.
- Note 61** When ULP sends new data no retransmit performed. Fast retransmit trigger is SACK only.
- Note 62** Fast retransmit is performed through the current Association->property (Data Transfer Path).
- Note 63** Interfaces RECEIVE_UNSENT and RECEIVE_UNACKED are not supported. See Reference [2] for more details.
- Note 64** BASE SCTP performs bundling in the first data transmission if bundling option is enabled in the configuration file. BASE SCTP always performs bundling for retransmissions.
- Note 65** BASE SCTP bundles ERROR chunk with COOKIE_ECHO chunk if error of INIT_ACK processing occurred.
- Note 66** BASE SCTP sends ABORT chunk on INIT chunk in case of an error.
- Note 67** Endpoint a_rwnd value is not changed by incoming SACK.
- Note 69** The 'cookie life-span increase time' parameter is ignored by BASE SCTP.
- Note 70** BASE SCTP sends ABORT chunk if Initiate Tag of INIT_ACK chunk is found to be zero.
- Note 72** User-Initiated Abort error cause is included in ABORT chunks that are sent because of an upper-layer request.
- Note 73** Protocol Violation error cause is included in ABORT chunks that are sent.
- Note 74** BASE SCTP does not provide additional information to specify the kind of protocol violation which has been detected.
- Note 75** BASE SCTP notifies the ULP about actual number of streams in COMMUNICATION UP notification.
- Note 76** The timestamp in the State Cookie is not used to determine which key should be used to verify the MAC.



- Note 77** BASE SCTP sends the ABORT with the error 'Restart of an association with new addresses' if there are new IP addresses in INIT received.
- Note 78** BASE SCTP processes incoming SACK chunk in COOKIE-ECHOED.
- Note 79** The limit of sending new DATA chunks calculated as described in RFC 4960:
- $$\text{if}((\text{flightsize} + \text{Max.Burst} * \text{MTU}) < \text{cwnd}) \text{ cwnd} = \text{flightsize} + \text{Max.Burst} * \text{MTU}$$
- is used as a temporary limitation for particular outgoing packet. It does not impact the true cwnd.
- Note 80** BASE SCTP does not bundle new DATA chunks with chunks to retransmit.
- Note 81** When the transmission windows is full, BASE SCTP still allows the ULP requesting to send new user data (but it is buffered till transmission is allowed again). However, user data is discarded when the transmission buffer is full [2] (chapter 5).
- Note 82** BASE SCTP sends an acknowledgement in the SHUTDOWN chunk.
- Note 83** If one duplicated chunk received then SACK chunk will be sent immediately on the received packet.
- Note 84** BASE SCTP reports list of duplicated TSNs in the SACK chunk.
- Note 85** The initial value of ssthresh is set to the maximum of the Initial RWND and 1073741824 bytes.
- Note 86** ICMP handling depends on SCTP configuration settings. See SCTP FS Reference [2] and SCTP CFD Reference [3] (ICMP status parameter) for more details.
- Note 87** BASE SCTP does not bundle ABORT chunk with other chunks.
- Note 88** BASE SCTP sends HB and HB_ACK chunks to UNCONFIRMED address. HB chunk always includes random nonce. The last successfully transmitted and processed pair of COOKIE_ECHO/COOKIE_ACK chunks defines the primary path which is automatically considered as CONFIRMED. Thus, COOKIE_ECHO and COOKIE_ACK chunks are not sent to UNCONFIRMED address.



- Note 89** In the case of non-zero "Arwnd Update Threshold" configuration parameter BASE SCTP will send an additional SACK chunk bundled with HEARTBEAT chunk which reports the latest received TSN and non-zero Advertised Receiver Window Credit when the incoming buffer (or part of this buffer) is consumed by a user. See SCTP FS Reference [2] and SCTP CFD Reference [3] ("Arwnd Update Threshold" parameter).
- Note 90** BASE SCTP responds to the sender of the OOTB packet with an ABORT. This note is added explicitly in order to comment the rule 8) which includes word 'should' in not capital.
- Note 91** If the primary path state is changed from INACTIVE to ACTIVE, BASE SCTP will use other path as primary for transmitting data if "Activate Threshold" for affected association is greater than 1. See SCTP FS Reference [2] and SCTP CFD Reference [3] ("Minimum Activate Threshold" and "Maximum Activate Threshold" parameters).
- Note 92** BASE SCTP supports a chunk type (170) that is not defined by IETF: ERICSSON_PROPRIETARY. The first parameter (chunk sub-ID) within this chunk defines the actual purpose of the chunk. At this moment there are two sub-IDs used: CHANGE_VTAG (1) and CHANGE_VTAG_ACK (2). The detailed description of the new chunk can be found in SCTP FS Reference [2]. The highest-order 2 bits of the chunk type is encoded to 10 that means "Skip this chunk and continue processing"
- Note 93** When BASE SCTP is going to send outgoing DATA chunks, SCTP may also bundle SACK chunk with these DATA chunks. The SACK will be bundled only if incoming DATA to be SACKed came from the current primary path, otherwise SACK will be sent separately later.
- Note 94** BASE SCTP is not compliant with a rule in ch.5.4 that the CONFIRMED path is the path where the pair of INIT/INIT ACK was sent during association establishment. In BASE SCTP implementation the CONFIRMED path is the IP path where COOKIE ECHO/COOKIE ACK pair was sent.
- Note 95** BASE SCTP starts heartbeat timer per IP path and sends one heartbeat via particular idle IP path each time when the timer is expired independently from other IP paths, when Quicker Failure Detection is enabled. For more details see Reference [2].



- Note 96** BASE SCTP spends more time for cookie generating when Remote Endpoint Blocks are configured (see Reference [3]). It means that if there are a lot of configured Remote Endpoint Blocks, BASE SCTP is more vulnerable to attackers. For more details see Reference [2] and Reference [3].
- Note 97** BASE SCTP can bundle PAD chunks with Heartbeat chunks when Sending Heartbeats with PMTU size is turned on. For more details see Reference [3].
- Note 98** If T3 timer is expired during zero window probing, error counters for the association and the path (SCTP path and IP path) are increased. Sack with the old cumulative TSN ack and nonzero a_rwnd does not reset error counters of the association and the path.
- Note 99** If association closure is caused by DESTROY message then COMMUNICATION LOST notification is not sent to the user.
- Note 100** In BASE SCTP implementation zero probing time could be limited by the Zero RWND Supervision timer. This timer supervises the time a SCTP keeps the calculated peer's rwnd (not only a_rwnd reported in SACK but including already sent data) equal to zero. Association is closed after the Zero RWND Supervision timer expiration. For more details see Reference [2] and Reference [3].
- Note 101** BASE SCTP can be configured to ignore incoming INIT chunks from specific remote sides. So BASE SCTP is able to discard incoming INIT chunks even if they are correctly formed and arrive on the existing endpoint. For more details see Reference [2] and Reference [3].
- Note 102** SCTP performs RTO retransmission of "old" data on the preferred primary path if the switchback to the preferred primary path has been performed while data was in-flight. For more details see Reference [2].
- Note 103** CWND is not increased by more than one MTU bytes per RTT.
- Note 104** BASE SCTP implementation is able to perform Fast Retransmit on reported gaps several times, if configuration option **Multiple Fast Retransmit** is enabled. SCTP performs Fast Retransmit of the same DATA chunk at most only once per RTT.



- Note 105** BASE SCTP supports the "Fast Transfer to Zero Window Probe State" procedure. The procedure guarantees that all gaps are filled and that SCTP (the sender) will not retransmit chunks outside the new announced window. Once all gaps are filled then BASE SCTP will terminate data retransmission process and enter to zero window probe state. For more details see Reference [2].
- Note 106** SCTP will send SACK chunk with size greater than PMTU if there are too many GAP ACK blocks, caused by packet loss.
- Note 107** BASE SCTP implementation adjusts CWND of the current primary path to $\max(\text{CWND}/2, \min(4 \cdot \text{MTU}, \max(2 \cdot \text{MTU}, 4380)))$ per RTO during zero window probing. For more details see Reference [2].
- Note 108** In multihoming endpoint case BASE SCTP implementation can send new DATA chunks even if there are chunks which are marked as ready to be retransmitted (due to t3 time-out or fast retransmission), but has not been sent yet. For more details see Reference [2].
- Note 109** BASE SCTP can not omit ERROR parameters. It will send a big INIT_ACK or abort association establishment.
- Note 110** If the list of IP addresses received from DNS does not include the source IP address of the INIT, BASE SCTP silently discards the INIT.
- Note 111** Additionally to detecting three consecutive miss-indications about TSN, BASE SCTP performs fast retransmission of a DATA chunk indicated as missed in a SACK if three or more IP packets sent later (than the packet with the concerned DATA chunk) are acknowledged by the SACK.
- Note 112** For an association in Fast Recovery BASE SCTP performs fast retransmission of every DATA chunk indicated as missed in a SACK if there are no TSNs from the SACK which fulfil the criteria for Fast Retransmit and there is no new data ready to be sent.
- Note 113** BASE SCTP allows to configure maximum CWND increment in slow start phase in percent from the standard value (1 PMTU). For details see description of the "Slow Start CWND Increment Factor" configuration parameter in **Reference [3]** (the feature is based on **Reference [5]**). By default SCTP follows the RFC 4960.



- Note 114** When all paths within an association become inactive (due to exceeded Path.Max.Rtx threshold or receipt of an ICMP destination unreachable message) but error counter of the association has not reached Assoc.Max.Rtx value yet, BASE SCTP behavior depends on the path selection scheme used for this association. If the path selection scheme A ("PEER") is used, the association will be terminated. If the path selection scheme B ("PATH") is used, the association will be moved to the dormant state and will not be closed. In the dormant state, as long as the Assoc.Max.Rtx threshold has not been exceeded, BASE SCTP will continue HB probing on idle paths, transmission of new data (if any) and retransmission of unacknowledged data.
- For more information on the path selection schemes and on the dormant state of association see **Reference [2]**.
- Note 115** Failed or acknowledged HB has influence on association and path error counters if and only if no reliable path availability indication has been received from data or HB within the current RTO period and there is no outstanding data which can generate reliable path availability indication. Failed or acknowledged HBs on a path different from the data transfer path do not have influence on association error counter.
- Note 116** In BASE SCTP implementation it is possible to create several endpoints with the same port number and the same IP address, if IP addresses have different VPN identifier. It is also possible to create an endpoint with the same IP address listed several times in the SCTP_INITIALIZE_req if IP addresses have different VPN identifiers. See Reference [2] for more details.
- Note 117** In BASE SCTP implementation it is possible to configure which CRC32C checksum calculation algorithm will be used. For more information on the supported calculation algorithms see **Reference [2]**.
- Note 118** BASE SCTP sends regular HB on an IP path if and only if no reliable path availability indication has been received from data or HB within the current heartbeat period and there is no outstanding data which can generate reliable path availability indication (for more information see **Reference [2]**).



Reference List

- [1] R. Stewart, "Stream Control Transmission Protocol", RFC 4960. Standards Track. September 2007.
- [2] Functional Specification for SCTP 155 17-CAA 901 548 Uen
- [3] Configuration File Description for SCTP 19073-CAA 901 548 Uen
- [4] R. Stewart, "Stream Control Transmission Protocol (SCTP) Specification Errata and Issues", RFC4460
- [5] M. Allman, "TCP Congestion Control with Appropriate Byte Counting (ABC)", RFC 3465