

Internet Engineering Task Force (IETF)
Request for Comments: 6084
Category: Experimental
ISSN: 2070-1721

X. Fu
C. Dickmann
University of Goettingen
J. Crowcroft
University of Cambridge
January 2011

General Internet Signaling Transport (GIST)
over Stream Control Transmission Protocol (SCTP)
and Datagram Transport Layer Security (DTLS)

Abstract

The General Internet Signaling Transport (GIST) protocol currently uses TCP or Transport Layer Security (TLS) over TCP for Connection mode operation. This document describes the usage of GIST over the Stream Control Transmission Protocol (SCTP) and Datagram Transport Layer Security (DTLS).

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6084>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

- 1. Introduction 3
- 2. Terminology and Abbreviations 4
- 3. GIST over SCTP 5
 - 3.1. Message Association Setup 5
 - 3.1.1. Overview 5
 - 3.1.2. Protocol-Definition: Forwards-SCTP 5
 - 3.2. Effect on GIST State Maintenance 6
 - 3.3. PR-SCTP Support 6
 - 3.4. API between GIST and NSLP 7
- 4. Bit-Level Formats 7
 - 4.1. MA-Protocol-Options 7
- 5. Application of GIST over SCTP 8
 - 5.1. Multihoming Support of SCTP 8
 - 5.2. Streaming Support in SCTP 8
- 6. NAT Traversal Issue 8
- 7. Use of DTLS with GIST 9
- 8. Security Considerations 9
- 9. IANA Considerations 10
- 10. Acknowledgments 10
- 11. References 10
 - 11.1. Normative References 10
 - 11.2. Informative References 11

1. Introduction

This document describes the usage of the General Internet Signaling Transport (GIST) protocol [1] and Datagram Transport Layer Security (DTLS) [2].

GIST, in its initial specification for Connection mode (C-mode) operation, runs on top of a byte-stream-oriented transport protocol providing a reliable, in-sequence delivery, i.e., using the Transmission Control Protocol (TCP) [9] for signaling message transport. However, some Next Steps in Signaling (NSIS) Signaling Layer Protocol (NSLP) [10] context information has a definite lifetime; therefore, the GIST transport protocol could benefit from flexible retransmission, so stale NSLP messages that are held up by congestion can be dropped. Together with the head-of-line blocking and multihoming issues with TCP, these considerations argue that implementations of GIST should support SCTP as an optional transport protocol for GIST. Like TCP, SCTP supports reliability, congestion control, and fragmentation. Unlike TCP, SCTP provides a number of functions that are desirable for signaling transport, such as multiple streams and multiple IP addresses for path failure recovery. Furthermore, SCTP offers an advantage of message-oriented transport instead of using the byte-stream-oriented TCP where the framing mechanisms must be provided separately. In addition, its Partial Reliability extension (PR-SCTP) [3] supports partial retransmission based on a programmable retransmission timer. Furthermore, DTLS provides a viable solution for securing SCTP [4], which allows SCTP to use almost all of its transport features and its extensions.

This document defines the use of SCTP as the underlying transport protocol for GIST and the use of DTLS as a security mechanism for protecting GIST Messaging Associations and discusses the implications on GIST state maintenance and API between GIST and NSLPs. Furthermore, this document describes how GIST is transported over SCTP and used by NSLPs in order to exploit the additional capabilities offered by SCTP to deliver GIST C-mode messages more effectively. More specifically:

- o How to use the multiple streams feature of SCTP.
- o How to use the PR-SCTP extension of SCTP.
- o How to take advantage of the multihoming support of SCTP.

GIST over SCTP as described in this document does not require any changes to the high-level operation and structure of GIST. However, adding new transport options requires additional interface code and configuration support to allow applications to exploit the additional

transport when appropriate. In addition, SCTP implementations to transport GIST MUST support the optional feature of fragmentation of SCTP user messages.

Additionally, this document also specifies how to establish GIST security using DTLS for use in combination with, e.g., SCTP and UDP.

2. Terminology and Abbreviations

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 [5]. Other terminologies and abbreviations used in this document are taken from related specifications ([1], [2], [3], [6]):

- o SCTP - Stream Control Transmission Protocol
- o PR-SCTP - SCTP Partial Reliability Extension
- o MRM - Message Routing Method
- o MRI - Message Routing Information
- o SCD - Stack-Configuration-Data
- o Messaging Association (MA) - A single connection between two explicitly identified GIST adjacent peers, i.e., between a given signaling source and destination address. A messaging association may use a transport protocol; if security protection is required, it may use a specific network layer security association, or use a transport layer security association internally. A messaging association is bidirectional: signaling messages can be sent over it in either direction, referring to flows of either direction.
- o SCTP Association - A protocol relationship between SCTP endpoints, composed of the two SCTP endpoints and protocol state information. An association can be uniquely identified by the transport addresses used by the endpoints in the association. Two SCTP endpoints MUST NOT have more than one SCTP association between them at any given time.
- o Stream - A unidirectional logical channel established from one to another associated SCTP endpoint, within which all user messages are delivered in sequence except for those submitted to the unordered delivery service.

3. GIST over SCTP

This section defines a new MA-Protocol-ID type, "Forwards-SCTP", for using SCTP as the GIST transport protocol. The use of DTLS in GIST is defined in Section 7.

3.1. Message Association Setup

3.1.1. Overview

The basic GIST protocol specification defines two possible protocols to be used in Messaging Associations, namely Forwards-TCP and TLS. This information is a main part of the Stack Configuration Data (SCD) [1]. This section adds Forwards-SCTP (value 3) as another possible protocol option. In Forwards-SCTP, analog to Forwards-TCP, connections between peers are opened in the forwards direction, from the querying node, towards the responder.

3.1.2. Protocol-Definition: Forwards-SCTP

The MA-Protocol-ID "Forwards-SCTP" denotes a basic use of SCTP between peers. Support for this protocol is OPTIONAL. If this protocol is offered, MA-protocol-options data MUST also be carried in the SCD object. The MA-protocol-options field formats are:

- o in a Query: no information apart from the field header.
- o in a Response: 2-byte port number at which the connection will be accepted, followed by 2 pad bytes.

The connection is opened in the forwards direction, from the querying node towards the responder. The querying node MAY use any source address and source port. The destination for establishing the message association MUST be derived from information in the Response: the address from the interface-address in the Network-Layer-Information object and the port from the SCD object as described above.

Associations using Forwards-SCTP can carry messages with the transfer attribute Reliable=True. If an error occurs on the SCTP connection such as a reset, as can be reported by an SCTP socket API notification [11], GIST MUST report this to NSLPs as discussed in Section 4.1.2 of [1]. For the multihoming scenario, when a destination address of a GIST-over-SCTP peer encounters a change, the SCTP API will notify GIST about the availability of different SCTP endpoint addresses and the possible change of the primary path.

3.2. Effect on GIST State Maintenance

As SCTP provides additional functionality over TCP, this section discusses the implications of using GIST over SCTP on GIST state maintenance.

While SCTP defines unidirectional streams, for the purpose of this document, the concept of a bidirectional stream is used.

Implementations MUST establish both downstream and upstream (unidirectional) SCTP streams and use the same stream identifier in both directions. Thus, the two unidirectional streams (in opposite directions) form a bidirectional stream.

Due to the multi-streaming support of SCTP, it is possible to use different SCTP streams for different resources (e.g., different NSLP sessions), rather than maintaining all messages along the same transport connection/association in a correlated fashion as TCP (which imposes strict (re)ordering and reliability per transport level). However, there are limitations to the use of multi-streaming. When an SCTP implementation is used for GIST transport, all GIST messages for a particular session MUST be sent over the same SCTP stream to assure the NSLP assumption of in-order delivery. Multiple sessions MAY share the same SCTP stream based on local policy.

The GIST concept of Messaging Association re-use is not affected by this document or the use of SCTP. All rules defined in the GIST specification remain valid in the context of GIST over SCTP.

3.3. PR-SCTP Support

A variant of SCTP, PR-SCTP [3] provides a "timed reliability" service, which would be particularly useful for delivering GIST Connection mode messages. It allows the user to specify, on a per-message basis, the rules governing how persistent the transport service should be in attempting to send the message to the receiver. Because of the chunk bundling function of SCTP, reliable and partially reliable messages can be multiplexed over a single PR-SCTP association. Therefore, an SCTP implementation for GIST transport SHOULD attempt to establish a PR-SCTP association using "timed reliability" service instead of a standard SCTP association, if available, to support more flexible transport features for potential needs of different NSLPs.

When using a normally reliable session (as opposed to a partially reliable session), if a node has sent the first transmission before the lifetime expires, then the message MUST be sent as a normal reliable message. During episodes of congestion, this is

particularly unfortunate, as retransmission wastes bandwidth that could have been used for other (non-lifetime expired) messages. The "timed reliability" service in PR-SCTP eliminates this issue and is hence RECOMMENDED to be used for GIST over PR-SCTP.

3.4. API between GIST and NSLP


The GIST specification defines an abstract API between GIST and NSLPs. While this document does not change the API itself, the semantics of some parameters have slightly different interpretations in the context of SCTP. This section only lists those primitives and parameters that need special consideration when used in the context of SCTP. The relevant primitives from [1] are as follows:

- o The Timeout parameter in API "SendMessage": According to [1], this parameter represents the "length of time GIST should attempt to send this message before indicating an error". When used with PR-SCTP, this parameter is used as the timeout for the "timed reliability" service of PR-SCTP.
- o "NetworkNotification": According to [1], this primitive "is passed from GIST to a signalling application. It indicates that a network event of possible interest to the signalling application occurred". Here, if SCTP detects a failure of the primary path, GIST SHOULD also indicate this event to the NSLP by calling this primitive with Network-Notification-Type "Routing Status Change". This notification should be done even if SCTP was able to retain an open connection to the peer due to its multihoming capabilities.

4. Bit-Level Formats

4.1. MA-Protocol-Options

This section provides the bit-level format for the MA-protocol-options field that is used for SCTP protocol in the Stack-Configuration-Data object of GIST.



SCTP port number = Port number at which the responder will accept SCTP connections

The SCTP port number is only supplied if sent by the responder.

5. Application of GIST over SCTP

5.1. Multihoming Support of SCTP

In general, the multihoming support of SCTP can be used to improve fault-tolerance in case of a path or link failure. Thus, GIST over SCTP would be able to deliver NSLP messages between peers even if the primary path is not working anymore. However, for the Message Routing Methods (MRMs) defined in the basic GIST specification, such a feature is only of limited use. The default MRM is path-coupled, which means that if the primary path is failing for the SCTP association, it most likely is also failing for the IP traffic that is signaled for. Thus, GIST would need to perform a refresh to the NSIS nodes to the alternative path anyway to cope with the route change. When the two endpoints of a multihomed SCTP association (but none of the intermediate nodes between them) support NSIS, GIST over SCTP provides a robust means for GIST to deliver NSLP messages even when the primary path fails but at least one alternative path between these (NSIS-enabled) endpoints of the multihomed path is available. Additionally, the use of the multihoming support of SCTP provides GIST and the NSLP with another source to detect route changes. Furthermore, for the time between detection of the route change and recovering from it, the alternative path offered by SCTP can be used by the NSLP to make the transition more smoothly. Finally, future MRMs might have different properties and therefore benefit from multihoming more broadly.

5.2. Streaming Support in SCTP

Streaming support in SCTP is advantageous for GIST. It allows better parallel processing, in particular by avoiding the head-of-line blocking issue in TCP. Since a single GIST MA may be reused by multiple sessions, using TCP as the transport for GIST signaling messages belonging to different sessions may be blocked if another message is dropped. In the case of SCTP, this can be avoided, as different sessions having different requirements can belong to different streams; thus, a message loss or reordering in a stream will only affect the delivery of messages within that particular stream, and not any other streams.

6. NAT Traversal Issue

NAT traversal for GIST over SCTP will follow Section 7.2 of [1] and the GIST extensibility capabilities defined in [12]. This specification does not define NAT traversal procedures for GIST over SCTP, although an approach for SCTP NAT traversal is described in [13].

7. Use of DTLS with GIST

This section specifies a new MA-Protocol-ID "DTLS" (value 4) for the use of DTLS in GIST, which denotes a basic use of datagram transport layer channel security, initially in conjunction with GIST over SCTP. It provides server (i.e., GIST transport receiver) authentication and integrity (as long as the NULL ciphersuite is not selected during ciphersuite negotiation), as well as optionally replay protection for control packets. The use of DTLS for securing GIST over SCTP allows GIST to take the advantage of features provided by SCTP and its extensions. The usage of DTLS for GIST over SCTP is similar to TLS for GIST as specified in [1], where a stack-proposal containing both MA-Protocol-IDs for SCTP and DTLS during the GIST handshake phase.

The usage of DTLS [2] for securing GIST over datagram transport protocols MUST be implemented and SHOULD be used.

GIST message associations using DTLS may carry messages with transfer attributes requesting confidentiality or integrity protection. The specific DTLS version will be negotiated within the DTLS layer itself, but implementations MUST NOT negotiate to protocol versions prior to DTLS v1.0 and MUST use the highest protocol version supported by both peers. NULL authentication and integrity ciphers MUST NOT be negotiated for GIST nodes supporting DTLS. For confidentiality ciphers, nodes can negotiate the NULL ciphersuites. The same rules for negotiating TLS ciphersuites as specified in Section 5.7.3 of [1] apply.

DTLS renegotiation [7] may cause problems for applications such that connection security parameters can change without the application knowing it. Hence, it is RECOMMENDED that renegotiation be disabled for GIST over DTLS.

No MA-protocol-options field is required for DTLS. The configuration information for the transport protocol over which DTLS is running (e.g., SCTP port number) is provided by the MA-protocol-options for that protocol.

8. Security Considerations

The security considerations of [1], [6], and [2] apply. Additionally, although [4] does not support replay detection in DTLS over SCTP, the SCTP replay protection mechanisms [6] [8] should be able to protect NSIS messages transported using GIST over (DTLS over) SCTP from replay attacks.

9. IANA Considerations

According to this specification, IANA has registered the following codepoints (MA-Protocol-IDs) in a registry created by [1]:

MA-Protocol-ID	Protocol
3	SCTP opened in the forwards direction
4	DTLS initiated in the forwards direction

Note that MA-Protocol-ID "DTLS" is never used alone but always coupled with a transport protocol specified in the stack proposal.

10. Acknowledgments

The authors would like to thank John Loughney, Jukka Manner, Magnus Westerlund, Sean Turner, Lars Eggert, Jeffrey Hutzelman, Robert Hancock, Andrew McDonald, Martin Stiernerling, Fang-Chun Kuo, Jan Demter, Lauri Liuhto, Michael Tuexen, and Roland Bless for their helpful suggestions.

11. References

11.1. Normative References

- [1] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", RFC 5971, October 2010.
- [2] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security", RFC 4347, April 2006.
- [3] Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP) Partial Reliability Extension", RFC 3758, May 2004.
- [4] Tuexen, M., Seggelmann, R., and E. Rescorla, "Datagram Transport Layer Security (DTLS) for Stream Control Transmission Protocol (SCTP)", RFC 6083, January 2011.
- [5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [6] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

- [7] Rescorla, E., Ray, M., Dispensa, S., and N. Oskov, "Transport Layer Security (TLS) Renegotiation Indication Extension", RFC 5746, February 2010.
- [8] Tuexen, M., Stewart, R., Lei, P., and E. Rescorla, "Authenticated Chunks for the Stream Control Transmission Protocol (SCTP)", RFC 4895, August 2007.

11.2. Informative References

- [9] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [10] Hancock, R., Karagiannis, G., Loughney, J., and S. Van den Bosch, "Next Steps in Signaling (NSIS): Framework", RFC 4080, June 2005.
- [11] Stewart, R., Poon, K., Tuexen, M., Yasevich, V., and P. Lei, "Sockets API Extensions for Stream Control Transmission Protocol (SCTP)", Work in Progress, January 2011.
- [12] Manner, J., Bless, R., Loughney, J., and E. Davies, "Using and Extending the NSIS Protocol Family", RFC 5978, October 2010.
- [13] Stewart, R., Tuexen, M., and I. Ruengeler, "Stream Control Transmission Protocol (SCTP) Network Address Translation", Work in Progress, December 2010.

Authors' Addresses

Xiaoming Fu
University of Goettingen
Institute of Computer Science
Goldschmidtstr. 7
Goettingen 37077
Germany

EEmail: fu@cs.uni-goettingen.de

Christian Dickmann
University of Goettingen
Institute of Computer Science
Goldschmidtstr. 7
Goettingen 37077
Germany

EEmail: mail@christian-dickmann.de

Jon Crowcroft
University of Cambridge
Computer Laboratory
William Gates Building
15 JJ Thomson Avenue
Cambridge CB3 0FD
UK

EEmail: jon.crowcroft@cl.cam.ac.uk