Network Working Group
Internet-Draft
Expires: December 20, 2006

X. Fu
C. Dickmann
University of Goettingen
J. Crowcroft
University of Cambridge
June 18, 2006

General Internet Signaling Transport (GIST) over SCTP
draft-ietf-nsis-ntlp-sctp-00.txt

Status of this Memo

Copyright Notice

Abstract

The General Internet Signaling Transport (GIST) protocol currently uses TCP or TLS over TCP for connection mode operation. This document describes the usage of GIST over the Stream Control Transmission Protocol (SCTP). The use of SCTP can take the advantage of features provided by SCTP, namely streaming-based transport, support of multiple streams to avoid head of line blocking, and the

support of multi-homing to provide network level fault tolerance. Additionally, the support for some extensions of SCTP is also discussed, namely its Partial Reliability Extension and the usage of TLS over SCTP.

Table of Contents

1. Introduction

This document describes the usage of the General Internet Signaling Transport (GIST) protocol [1] over the Stream Control Transmission Protocol (SCTP) [2].

GIST, in its initial specification for connection 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) [4] for signaling message transport. However, some NSLP context information has a definite lifetime, therefore, the GIST transport protocol must accommodate flexible retransmission, so stale NSLP messages that are held up by congestion can be dropped. Together with the head-of-line blocking issue and other issues with TCP, these considerations argue that implementations of GIST should support the Stream Control Transport Protocol (SCTP)[2] as an optional transport protocol for GIST, especially if deployment over the public Internet is contemplated. Like TCP, SCTP supports reliability, congestion control, 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. In addition, its Partial Reliability extension (PR-SCTP) [5] supports partial retransmission based on a programmable retransmission timer.

This document shows how GIST should be used with SCTP to provide these additional features to deliver the GIST C-mode messages (which can in turn carry NSIS Signaling Layer Protocol (NSLP) [6] messages as payload). More specifically:

- how to use the multiple streams feature of SCTP.
- how to handle the message oriented nature of SCTP.
- how to take the advantage of multi-homing support of SCTP.

Additionally, this document also discusses how to support two extensions of SCTP, namely PR-SCTP [5] and TLS over SCTP [7].

The method described in this document does not require any changes of GIST or SCTP. It is only required that SCTP implementations support the optional feature of fragmentation of SCTP user messages.

2. Terminology and Abbreviations

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL", in this document are to be interpreted as described in BCP 14, RFC 2119 [3]. Other terminologies and abbreviations used in this document are taken from related specifications (e.g., [1] and

[2]) as follows:

- TLS - Transport Layer Security
- SCTP - Stream Control Transmission Protocol
- PR-SCTP - SCTP Partial Reliability Extension
- MRM - Message Routing Method
- MRI - Message Routing Information
- MRS - Message Routing State
- MA - A GIST Messaging Association is a single connection between two explicitly identified GIST adjacent peers on the data path. A messaging association may use a specific transport protocol and known ports. 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, and can refer to flows of either direction.
- 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.
- Stream - A sequence of user messages that are to be delivered to the upper-layer protocol in order with respect to other messages within the same stream.

3. GIST Over SCTP

3.1. Message Association Setup

The basic GIST protocol specification defines two possible protocols to be used in message associations, namely Forwards-TCP and TLS. This document adds Forwards-SCTP as another possible protocol. In Forwards-SCTP, analog to Forwards-TCP, connections between peers are opened in the forwards direction, from the querying node, towards the responder. SCTP connections may carry NSLP messages with the transfer attribute 'reliable'.

A new MA-Protocol-ID type, "Forwards-SCTP", is defined in this document for using SCTP as GIST transport protocol.

3.2. Stack-Configuration-Data information for SCTP

In order to run GIST over SCTP, the Stack-Proposal and Stack-Configuration-Data objects need to recognize the Forwards-SCTP MA-Protocol-ID type, and interpret it for the transport protocol negotiation during the GIST MA setup handshake (e.g., whether SCTP runs alone or together with TLS).

In turn, the "MA-protocol-options" field for Forwards-SCTP needs to be defined for the Stack-Configuration-Data object defined by GIST. This "MA-protocol-options" contains proposed values for the initial and maximum retransmission timeout (RTO) as well as a port number in the case of Response messages. The proposed values for RTO are only suggestions to the peer and may be overridden by local policy. In fact, in order to avoid denial of service attacks, the minimum RTO value is not included in the proposal and in addition implementations should only accept reasonable RTO proposals.

The MA-protocol-options formats are:

- in a Query: 4 byte RTO initial value and 4 byte RTO maximum value
- in a Response: 4 byte RTO initial value, 4 byte RTO maximum value and 2 byte port number at which the connection will be accepted.

3.3. Effect on GIST State Maintenance

A GIST MA is established over an SCTP association, which comprises one or more SCTP streams. Each of such streams can be used for one or multiple sessions (i.e., one or more MRSs). After completing a GIST MA setup, which implicitly establishes a bi-directional SCTP stream, C-mode messages can be sent over the SCTP association in either direction. Due to multi-streaming support of SCTP, it is easy to maintain sequencing of messages that affect the same resource (e.g., the same NSLP session), 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).

3.4. PR-SCTP Support

A variant of SCTP, PR-SCTP [5] provides a "timed reliability" service. 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 partial reliable messages can be multiplexed over a single PR-SCTP association. Therefore, a GIST over SCTP implementation SHOULD attempt to establish a PR-SCTP association instead of a standard SCTP association, if available, to support more flexible transport features for potential needs of different NSLPs.

3.5. API between GIST and NSLP

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 interpretation 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 are repeatet from [1] to improve readability, but [1] remains authoritative.

3.5.1. SendMessage

The SendMessage primitive is used by the NSLP to initiate sending of messages.

```
SendMessage ( NSLP-Data, NSLP-Data-Size, NSLP-Message-Handle,
              NSLP-Id, Session-ID, MRI,
              SSI-Handle, Transfer-Attributes, Timeout, IP-TTL, GHC )
```

The following parameter has changed semantics:

Timeout: According to [1] this parameter represents the "length of time GIST should attempt to send this message before indicating an error". When used with SCTP, this parameter is also used as the timeout for the "timed reliability" service of PR-SCTP.

3.5.2. NetworkNotification

The NetworkNotification primitive is passed from GIST to an NSLP. It indicates that a network event of possible interest to the NSLP occurred.

```
NetworkNotification ( MRI, Network-Notification-Type )
```

If SCTP detects a failure of the primary path, GIST should indicate this event to the NSLP by calling the NetworkNotification primitive with Network-Notification-Type "Routing Status Change". This notification should be done even if SCTP was able to remain an open connection to the next peer due to its multi-homing capabilities.

3.6. TLS over SCTP Support

GIST using TLS over SCTP is analog to GIST using TLS over TCP. Thus, TLS over SCTP is triggered by a protocol stack consisting of the Forwards-SCTP MA-protocol-ID and the TLS MA-protocol-ID ([1], Section 5.7.3). The GIST specification defines the versions of TLS that can be used, as well as the authentication model. All these aspects are not changed by this document and remain valid for TLS over SCTP. Regarding GIST implementations, no special treatment is required in the case of TLS over SCTP in contrast to the existing TLS over TCP case. However, the SCTP and TLS implementations need to provide a TLS over SCTP service as descriped in [7]. One should note that an

SCTP association with TLS support takes advantages of SCTP, such as multi-streaming and multi-homing.

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 (see Section 3.2).

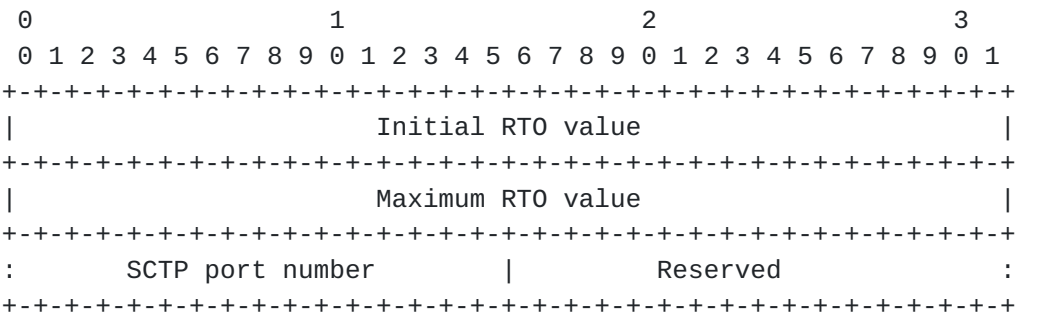

Initial RTO value = Initial RTO value (SCTP configuration) in msec
Maximum RTO value = Maximum RTO value (SCTP configuration) in msec
SCTP port number = Port number at which the responder will accept SCTP connections

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

5. Security Considerations

The security considerations of both [1] and [2] apply. Further security analysis is needed to consider any additional security vulnerabilities, and will be included in an updated draft.

6. IANA Considerations

A new MA-Protocol-ID (Forwards-SCTP) needs to be assigned, with a recommended value of 3.

7. Acknowledgments

The authors would like to thank John Loughney, Robert Hancock and Jan

Demter for their helpful suggestions.

8. References

8.1. Normative References

[1] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signaling Transport", draft-ietf-nsis-ntlp-09 (work in progress), February 2006.

[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] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

[4] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[5] Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP) Partial Reliability Extension", RFC 3758, May 2004.

[6] Hancock, R., Karagiannis, G., Loughney, J., and S. Van den Bosch, "Next Steps in Signaling (NSIS): Framework", RFC 4080, June 2005.

[7] Jungmaier, A., Rescorla, E., and M. Tuexen, "Transport Layer Security over Stream Control Transmission Protocol", RFC 3436, December 2002.

Authors' Addresses

Xiaoming Fu
University of Goettingen
Institute for Informatics
Lotzestr. 16-18
Goettingen 37083
Germany

Email: fu@cs.uni-goettingen.de

Christian Dickmann
University of Goettingen
Institute for Informatics
Lotzestr. 16-18
Goettingen 37083
Germany

Email: mail@christian-dickmann.de

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

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

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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