

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: November 24, 2017

L. Wood
Surrey alumni
W. Eddy
MTI Systems
W. Ivancic
Syzygy
May 23, 2017

Congestion control for the Saratoga protocol
draft-wood-tsvwg-saratoga-congestion-control-11

Abstract

Saratoga is a data transfer protocol designed to carry potentially large volumes of data over difficult network paths, often including only a single high-rate link and only one application flow. As the requirements for use vary across deployment environments, the base Saratoga specification only assumes that an implementation will be able to clock packets out at a configurable rate, and beyond this specifies no inherent or particular congestion-control behaviour. The design of Saratoga deliberately supports the integration of congestion-control algorithms without modification to the base protocol. This document describes how congestion control can be supported in the Saratoga transfer protocol. Saratoga is intended for use in private networks, where its use is engineered as a single flow to fill a link. However, as Saratoga is implemented over UDP, it can be multiplexed, and can be run across the public Internet, in which case congestion control in accordance with the UDP Guidelines becomes necessary.

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 24, 2017.

Copyright Notice

Copyright (c) 2017 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.

This document may not be modified, and derivative works of it may not be created, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Background and Introduction	2
2.	Approaches to congestion control	3
2.1.	TCP-friendly rate control	3
2.2.	Explicit Congestion Notification	3
3.	Security Considerations	4
4.	IANA Considerations	4
5.	Acknowledgements	4
6.	References	4
6.1.	Normative References	4
6.2.	Informative References	4
	Authors' Addresses	5

[1.](#) Background and Introduction

The Saratoga data transfer protocol is described in [[draft-wood-tsvwg-saratoga](#)]. Given that Saratoga was originally developed for scheduled peer-to-peer communications over dedicated links in private networks, where each application has the entire link for the duration of its transfer, many Saratoga implementations deliberately lack any form of congestion control and send at line rate to maximise throughput and link utilisation in their environments. Congestion control is necessary for use in the public Internet, in accordance with the UDP Guidelines [[RFC5405](#)]. Newer Saratoga implementations may use timestamps to perform TCP-Friendly Rate Control (TFRC) [[RFC5348](#)] or other congestion control mechanisms such as LEDBAT [[RFC6817](#)], if appropriate for the environment, and where simultaneous sharing of capacity with other traffic and applications is required. Sender-side TFRC for Saratoga has been shown to be possible without modifications to the Saratoga protocol

specification, using existing timestamps on selective negative acknowledgement messages [[draft-eddy-tsvwg-saratoga-tfrc](#)].

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 [RFC 2119](#). [[RFC2119](#)]

[2.](#) Approaches to congestion control

Saratoga can be implemented to perform congestion control at the sender, based on feedback from acknowledgement STATUS packets, or have the sender configured to use simple open-loop rate control to only use a fixed amount of link capacity. Congestion control is expected to be undesirable for many of Saratoga's use cases and expected environmental conditions, while simple rate control is considered useful for some use cases. Use over the public Internet requires congestion control.

Congestion control **MUST** be supported and used if Saratoga is being used across paths that go over the public Internet, and **SHOULD** be supported in environments where links in the path are shared by traffic flows. Congestion control **MAY NOT** be supported across private, single-flow links engineered for performance: Saratoga's primary use case.

[2.1.](#) TCP-friendly rate control

Sender-side TCP-friendly rate control can be implemented by mirroring timestamps in STATUS messages and using the approach outlined in [[draft-eddy-tsvwg-saratoga-tfrc](#)].

Other approaches to TCP-friendly congestion control are possible, and Saratoga and its selective negative acknowledgements may prove useful as an implementation testbed for developing and refining new congestion-control algorithms.

[2.2.](#) Explicit Congestion Notification

Supporting Explicit Congestion Notification in a UDP-based protocol such as Saratoga requires that ECN events be exposed to userspace applications using UDP via a programming interface. Once such a programming interface becomes available, providing counts of ECN events in STATUS and DATA packets will be straightforward. Until that time, specifying ECN support in more detail is not required.

3. Security Considerations

Use of effective congestion control mechanisms always raises concerns about fairness and spoofing or misleading senders - issues not unique to Saratoga.

4. IANA Considerations

There should be no additional IANA considerations.

5. Acknowledgements

We thank the IETF for reminding us about the importance of congestion for standards-track protocols.

Work on this document at NASA's Glenn Research Center was funded by NASA's Earth Science Technology Office (ESTO).

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

6.2. Informative References

[[draft-eddy-tsvwg-saratoga-tfrc](#)]
Eddy, W., Wood, L., and W. Ivancic, "TFRC-based congestion control for Saratoga", [draft-eddy-tsvwg-saratoga-tfrc-11](#) (work in progress) , May 2017.

[[draft-wood-tsvwg-saratoga](#)]
Wood, L., Eddy, W., Smith, C., Ivancic, W., and C. Jackson, "Saratoga: A Scalable Data Transfer Protocol", [draft-wood-tsvwg-saratoga-21](#) (work in progress) , May 2017.

[RFC5348] Floyd, S., Handley, M., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", [RFC 5348](#), DOI 10.17487/RFC5348, September 2008, <<http://www.rfc-editor.org/info/rfc5348>>.

- [RFC5405] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", [RFC 5405](#), DOI 10.17487/RFC5405, November 2008, <<http://www.rfc-editor.org/info/rfc5405>>.
- [RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", [RFC 6817](#), DOI 10.17487/RFC6817, December 2012, <<http://www.rfc-editor.org/info/rfc6817>>.

Authors' Addresses

Lloyd Wood
University of Surrey alumni
Sydney, New South Wales
Australia

Email: L.Wood@society.surrey.ac.uk

Wesley M. Eddy
MTI Systems
MS 500-ASRC
NASA Glenn Research Center
21000 Brookpark Road
Cleveland, OH 44135
USA

Phone: +1-216-433-6682
Email: wes@mti-systems.com

Will Ivancic
Syzygy Engineering LLC
Westlake, OH 44145
USA

Phone: +1-440-835-8448
Email: ivancic@syzygyengineering.com

