Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 4, 2020

N. Kuhn
CNES
G. Fairhurst
University of Aberdeen
J. Border
Hughes Network Systems, LLC
E. Stephan
Orange

July 3, 2019

QUIC for SATCOM
draft-kuhn-quic-4-sat-00

Abstract

QUIC's congestion control is not designed for operating over an Internet path with a high BDP. This limits the user experience. Moreover, a path can combine satellites network segment together with a wide variety of other network technologies (Ethernet, cable modems, WiFi, cellular, radio links, etc): this complicates the characteristics of the end-to-end path. If this is not addressed, the end-to-end quality of experience will be degraded.

This memo identifies the characteristics of a SATCOM link that impact the operation of the QUIC transport protocol and proposes best current practice to ensure acceptable protocol performance.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction 2
2. Operating over a path with a large BDP 3
3. TCP Split Solution 4
4. Mechanisms that improve the performance of QUIC for SATCOM . 4
  4.1. Getting up to speed 4
  4.2. Reliability 5
  4.3. Maximum window 5
  4.4. ACK ratio 5
5. Discussion 6
6. Acknowledgements 6
7. Contributors 6
8. IANA Considerations 6
9. Security Considerations 6
10. References 6
  10.1. Normative References 6
  10.2. Informative References 6
Authors' Addresses 8

1. Introduction

The end-to-end performance of an application using an Internet path can be impacted by the Bandwidth-Delay Product (BDP) of the links and network devices forming the path. For instance, the page load time for a complex page can be much larger when the path includes a satellite link. A significant contribution to this reduced performance arises from the initialisation and design of transport mechanisms. QUIC's congestion control is based on TCP NewReno [I-D.ietf-quic-recovery] and the recommended initial window is defined by [RFC6928].

Moreover, satellite communications (SATCOM) systems have long been used to support point-to-point links and specialised networks. The predominate current use is as a link-layer for Internet Protocols. Typical example applications include: use as an access technology for remote locations, backup and rapid deployment of new services, transit networks and backhaul of various types of IP networks, and provision to mobile (maritime, aircraft, etc.). The satellite IP network segment usually only forms one part of the end-to-end path. This means user traffic can experince a path that includes satellite link together with a wide variety of other network technologies (Ethernet, cable modems, WiFi, cellular, radio links, etc). Although a user can sometimes know the presence of the satellite service, a typical user does not deploy special software or applications because they expect a satellite network is being used. Often a user is unaware of the technologies underpinning the links forming the network path.

This memo identifies the characteristics of a SATCOM link that impact the operation of the QUIC transport protocol and proposes best current practice to ensure acceptable protocol performance.

2. Operating over a path with a large BDP

GEO-satellite based systems characteristics differ from paths only using terrestrial links in their path characteristics:

- A large propagation delay of at least 250ms one-way delay;

- Some systems can exhibit a high loss-rate (e.g. mobile users or users behind a Wi-Fi link);

- Employ radio resource management (often using techniques similar to cellular mobile or DOCSIS cable networks, but differing to accommodate the satellite propagation delay);

- Links can be highly asymmetric (in terms of capacity and one-way delay).

More information on satellite links characteristics can be found in [RFC2488].

These characteristics have an impact on the performance of end-to-end congestion controls:

- Transport initialization: the 3-way handshake takes a long time to complete, reducing the time at which actual data can be transmitted;

- Size of windows required: to fully exploit the bottleneck capacity, the high BDP may induce an important number of in-flights packets;

- Reliability: packet loss detection and correction is slow (the performance of end-to-end retransmission is also impacted when using a high RTT path);

- Getting up to speed: the exponential increase of the data rate during slow start for a channel capacity probing is slowed down when the RTT is high.

3. TCP Split Solution

High BDP networks commonly break the TCP end-to-end paradigm to adapt the transport protocol. Splitting TCP allows adaptations to this specific use-case and assessing the issues discussed in section Section 2. Satellite communications commonly deploy Performance Enhancement Proxy (PEP) for compression, caching and TCP acceleration services [RFC3135]. Their deployment can result in 50% page load time reduction in a SATCOM use-case [ICCRG100].

[NCT13] and [RFC3135] describe the main functions of SATCOM TCP split solutions. Shortly, for traffic originated at a gateway to an endpoint connected via a satellite terminal, the TCP split intercepts TCP SYN packets to act on behalf of the endpoint and adapt the data rate transmission to the SATCOM scenario. The split solution specifically tunes the TCP parameters to the context (latency, available capacity). The tuning can be achieved using a priori information about the satellite system and/or by measuring the properties of the network segment that includes the satellite system.

One important advantage of a TCP split solution is that it does not require any end-to-end modifications and is independent for both client and server sides. That being said, this comes with a drawback: TCP splitters often are unable to track the most recent end-to-end improvements (e.g. ECN or TCP Fast Open support). The methods configured in the split proxy usually continue to be used until a split solution is finally updated. This can delay/negate the benefit of any end-to-end improvements.

4. Mechanisms that improve the performance of QUIC for SATCOM

4.1. Getting up to speed

3-way handshake takes a long time reducing the time at which actual data can be transmitted.

The tuning of the initial window described in [I-D.irtf-iccrg-sallantin-initial-spreading] which has been shown to improve performance both for high BDP and more common BDP [CONEXT15][ICC16].

4.2. Reliability

Packet losses detection and correction is slow and the time needed for the end server to react to a congestion event may not be relevant. This happens when a user uses a Wi-Fi link to access a SATCOM terminal. Although the benefits needed to weighed against the additional capacity in introducing end-to-end FEC and the potential to use link-local ARQ and/or link-adaptive FEC.

Introducing network coding in QUIC could help in recovering from the residual Wi-Fi losses.

4.3. Maximum window

To fully exploit the bottleneck capacity, the high BDP may induce an important number of in-flights packets.

4.4. ACK ratio

Asymmetry in capacity (or in the way capacity is granted to a flow) can lead to cases where the throughput in one direction of communication is restricted by the acknowledgement traffic flowing in the opposite direction. The limitations of specific underlying networks could be in terms of the volume of acknowledgement traffic (limited return path capacity) or in the number of acknowledgement packets (e.g., when a radio-resource management system has to track channel usage) or both.

TCP Performance Implications of Network Path Asymmetry [RFC3449] describes a range of mechanisms that can mitigate the impact of path asymmetry. One simple method is to tell the remote endpoint to send compound acknowledgments less frequently. A rate of one ACK every RTT/4 can significantly reduce this traffic.

Many of these mitigations have been deployed in satellite systems, often as a mechanism within a PEP. Despite their benefits over paths with a high asymmetry of capacity, most mechanisms rely on being able to inspect and/or modify the transport layer header information of TCP ACK packets. This is not possible when the transport layer information is encrypted. The QUIC transport specification may evolve to allow the ACK Ratio to be adjusted.

5. Discussion

Many of the issues identified already exist for any encrypted transport service that uses a path that employs encryption at the IP layer. This includes endpoints that utilise IPsec at the network layer, or use VPN technology over the satellite network segment. These uses are unable to benefit from enhancement within the satellite network segment, and often the user is unaware of the presence of the satellite link on their path, except through observing the impact it has on the performance they experience.

6. Acknowledgements

None.

7. Contributors

None.

8. IANA Considerations

TBD: text is required to register the extension BDP_data field.

9. Security Considerations

This document does not propose changes to the security functions provided by the QUIC protocol. QUIC uses TLS encryption to protect the transport header and its payload. Security is considered in the "Security Considerations" of cited IETF documents.

10. References

10.1. 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, <https://www.rfc-editor.org/info/rfc2119>.

10.2. Informative References

[CONEXT15]
Li, Q., Dong, M., and P. Godfrey, "Halfback: Running Short Flows Quickly and Safely", ACM CoNEXT , 2015.

[I-D.ietf-quic-recovery]
Iyengar, J. and I. Swett, "QUIC Loss Detection and Congestion Control", draft-ietf-quic-recovery-20 (work in progress), April 2019.

[I-D.ietf-quic-tls]
Thomson, M. and S. Turner, "Using TLS to Secure QUIC", draft-ietf-quic-tls-20 (work in progress), April 2019.

[I-D.ietf-quic-transport]
Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", draft-ietf-quic-transport-20 (work in progress), April 2019.

[I-D.ietf-tls-ticketrequests]
Pauly, T., Schinazi, D., and C. Wood, "TLS Ticket Requests", draft-ietf-tls-ticketrequests-01 (work in progress), June 2019.

[I-D.irtf-iccrg-sallantin-initial-spreading]
Sallantin, R., Baudoin, C., Arnal, F., Dubois, E., Chaput, E., and A. Beylot, "Safe increase of the TCP's Initial Window Using Initial Spreading", draft-irtf-iccrg-sallantin-initial-spreading-00 (work in progress), January 2014.

[ICC16] Sallantin, R., Baudoin, C., Chaput, E., Arnal, F., Dubois, E., and A-L. Beylot, "Reducing web latency through TCP IW: Be smart", IEEE ICC , 2016.

[ICCRG100]
Kuhn, N., "MPTCP and BBR performance over Internet satellite paths", IETF ICCRG 100, 2017.

[IJSCN19] Thomas, L., Dubois, E., Kuhn, N., and E. Lochin, "Google QUIC performance over a public SATCOM access", International Journal of Satellite Communications and Networking , 2019.

[NCT13] Pirovano, A. and F. Garcia, "A new survey on improving TCP performances over geostationary satellite link", Network and Communication Technologies , 2013.

[RFC2488] Allman, M., Glover, D., and L. Sanchez, "Enhancing TCP Over Satellite Channels using Standard Mechanisms", BCP 28, RFC 2488, DOI 10.17487/RFC2488, January 1999, <https://www.rfc-editor.org/info/rfc2488>.

[RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z.

Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, DOI 10.17487/RFC3135, June 2001, <https://www.rfc-editor.org/info/rfc3135>.

[RFC3449] Balakrishnan, H., Padmanabhan, V., Fairhurst, G., and M. Sooriyabandara, "TCP Performance Implications of Network Path Asymmetry", BCP 69, RFC 3449, DOI 10.17487/RFC3449, December 2002, <https://www.rfc-editor.org/info/rfc3449>.

[RFC6349] Constantine, B., Forget, G., Geib, R., and R. Schrage, "Framework for TCP Throughput Testing", RFC 6349, DOI 10.17487/RFC6349, August 2011, <https://www.rfc-editor.org/info/rfc6349>.

[RFC6928] Chu, J., Dukkipati, N., Cheng, Y., and M. Mathis, "Increasing TCP's Initial Window", RFC 6928, DOI 10.17487/RFC6928, April 2013, <https://www.rfc-editor.org/info/rfc6928>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

Authors' Addresses

Nicolas Kuhn
CNES

Email: nicolas.kuhn@cnes.fr

Godred Fairhurst
University of Aberdeen

Email: gorry@erg.abdn.ac.uk

John Border
Hughes Network Systems, LLC

Email: border@hns.com

Emile Stephan
Orange

Email: emile.stephan@orange.com