Internet Engineering Task Force Internet-Draft Intended status: Informational Expires: September 7, 2020 N. Kuhn CNES G. Fairhurst University of Aberdeen J. Border Hughes Network Systems, LLC E. Stephan Orange March 6, 2020

QUIC for SATCOM draft-kuhn-quic-4-sat-04

Abstract

QUIC has been designed for use across Internet paths. Initial designs of QUIC have focussed on common deployment scenarios for web traffic and have not focussed on the performance when using a path with a large Bandwidth-Delay Product (BDP). 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. One example of such a scenario occurs when a satellite communication (SATCOM) system is used to provide all or a part of the end-to-end path. If this is not addressed, the end-to-end quality of experience can be degraded.

This memo identifies the characteristics of a SATCOM link that impact the operation of the QUIC transport protocol. It proposes regression tests to evaluate QUIC over SATCOM links. It discusses how to ensure acceptable protocol performance.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

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 https://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."

Kuhn, et al.

Expires September 7, 2020

This Internet-Draft will expire on September 7, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>https://trustee.ietf.org/license-info</u>) 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.

Table of Contents

<u>1</u> .	Introduction	<u>2</u>
<u>2</u> .	Operating over a path with a large BDP	<u>3</u>
<u>3</u> .	TCP Split Solution	<u>5</u>
<u>4</u> .	Regression tests	<u>6</u>
4	<u>1</u> . Small public satellite broadband access	<u>6</u>
4	<u>2</u> . Medium public satellite broadband access	<u>6</u>
4	<u>3</u> . Loss-free large public satellite broadband access	<u>7</u>
4	<u>4</u> . Lossy large public satellite broadband access	<u>8</u>
5.	Mechanisms that improve the performance of QUIC for SATCOM .	8
<u>5</u>	<u>1</u> . Getting up to speed	<u>8</u>
<u>5</u>	<u>.2</u> . Maximum window	<u>9</u>
5	<u>.3</u> . Reliability	<u>9</u>
<u>5</u>	<u>4</u> . ACK ratio	<u>9</u>
<u>6</u> .	Discussion	.0
<u>7</u> .	Acknowledgements	.0
<u>8</u> .	Security Considerations	.0
<u>9</u> .	Informative References	.1
Autl	nors' Addresses	.3

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 default congestion control is based on TCP NewReno [I-D.ietf-quic-recovery] and the recommended initial window

is defined by [<u>RFC6928</u>]. Although QUIC's CC and recovery have been designed for use across Internet Paths, the initial design could not optimise for the wide diversity of path characteristics that can occur. This document therefore considers the specific implications of paths with a significant BDP.

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.). In most scenarios, the satellite IP network segment usually only forms one part of the end-to-end path. This means user traffic can experience 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. It proposes regression tests to evaluate QUIC over SATCOM links. It discusses how to ensure acceptable protocol performance.

2. Operating over a path with a large BDP

Satellite communications systems have been deployed using many space orbits, including low earth orbit, medium earth orbits, geosynchronous orbits, elliptical orbits and more. This document focuses on Geosynchronous Earth Orbit (GEO) satellite systems.

The characteristics of systems using Geosynchronous Earth Orbit (GEO) satellites differ from paths only using terrestrial links in their path characteristics:

- o A large propagation delay of at least 250ms one-way delay;
- Employ radio resource management (often using techniques similar to cellular mobile or DOCSIS cable networks, but differing to accommodate the satellite propagation delay);
- o Links can be highly asymmetric (in terms of capacity, one-way delay and in their cost of operation).

Many systems use the DVB-S2 specifications, published by the European Telecomunications Standards Institute (ETSI), where the key concept is to ensure both a good usage of the satellite resource and a Quasi Error Free (QEF) link. It consists of monitoring the link quality in real-time, with the help of known symbol sequences, included along with regular packets, on which an estimation of the current signalto-noise ratio can be done. Then, this estimation is sent back to the transmitter that can adapt its coding rate and modulation in order to best fit the actual transmission conditions.

It is common to consider the satellite network segment composed of a forward link and a return link. The two links can have different capacities and employ different technologies to carry the IP packets. On the forward link, the satellite gateway uses all the available bandwidth, possibly with several carriers, to communicate with the remote terminals. A carrier is a single Time-Division-Multiplexing channel where packets addressed to terminals are multiplexed. On the return link, the satellite resource is shared among the users. Two access methods can be distinguished: on-demand access or contention access. In the former, a terminal receives dedicated resources on its own to communicate with the gateway. In the latter, some resources are reserved for contention access, where several terminals can compete to obtain the resource. Dedicated access, which is more common in currently deployed systems, can be through a Demand Assigned Multiple Access (DAMA) mechanism, while contention access techniques are usually based on Slotted Aloha (SA) and its numerous derivatives. More information on satellite links characteristics can be found in [RFC2488][IJSCN17].

Beyond that, even for characteristics shared with terrestrial links, the impact on a satellite link could be more and can be amplified by the large RTT. For example, systems can exhibit a high loss-rate (e.g. mobile users or users behind a Wi-Fi link) which would impact loss recovery mechanisms and congestion control reaction to such loss events. The characteristics of a GEO SATCOM system impact the performance of congestion control:

- o Transport initialization: the 3-way handshake takes a long time to complete, delaying the time at which actual data can be transmitted (there could be other transport protocol exchanges, such as TLS);
- Size of windows required: to fully exploit the bottleneck capacity, a high BDP will increase the 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;
- o Asymmetry : when the links are asymmetric, for various reasons, the sizing of the return link may induce modifications of the transport level acknowledgement traffic.

<u>3</u>. 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 <u>Section 2</u>. Satellite communications commonly deploy Performance Enhancement Proxy (PEP) for compression, caching and TCP acceleration services [<u>RFC3135</u>]. Their deployment can result in 50% page load time reduction in a SATCOM use-case [<u>ICCRG100</u>].

[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) of each link. When a PEP is used on each side of the satellite link, a protocol other than TCP, optimized for the satellite link, may be used. 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. Split connections also allow for recovery from packet losses that is local to the part of the connection on which the packet losses occurred. This eliminates the need for end-to-end recovery of lost packets.

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 in protocol mechanisms (e.g., RACK, ECN, TCP Fast Open support) contributing to ossification of the transport system. 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.

Internet-Draft

QUIC 4 SAT

4. Regression tests

This section proposes regression tests for QUIC. We define by:

- o Download path: from Internet to host
- o Upload path: from the host to Internet

4.1. Small public satellite broadband access

The characteristics of the architecture under test are the following:

- o Satellite downlink path: 10 Mbps
- o Satellite uplink path: 2 Mbps
- o No emulated packet loss
- o RTT: 650 ms
- o Buffer size : BDP

During the transmission of 100 MB on both download and upload paths, the test should report the downloading time of 2 MB, 10 MB and 100 MB.

Initial thoughts of the performance obectives for QUIC are the following:

- o 3 s for downloading 2 MB
- o 10 s for downloading 10 MB
- o 85 s for downloading 100 MB
- o 10 s for uploading 2 MB
- o 50 s for uploading 10 MB
- o 420 s for uploading 100 MB

4.2. Medium public satellite broadband access

The characteristics of the architecture under test are the following:

- o Satellite downlink path: 50 Mbps
- o Satellite uplink path: 10 Mbps

Kuhn, et al.Expires September 7, 2020[Page 6]

Internet-Draft

- o No emulated packet loss
- o RTT: 650 ms
- o Buffer size : BDP

During the transmission of 100 MB on the download path, the test should report the downloading time of 2 MB, 10 MB and 100 MB.

Initial thoughts of the performance obectives for QUIC are the following:

- o 3 s for downloading 2 MB
- o 5 s for downloading 10 MB
- o 20 s for downloading 100 MB

4.3. Loss-free large public satellite broadband access

The characteristics of the architecture under test are the following:

- o Satellite downlink path: 250 Mbps
- o Satellite uplink path: 3 Mbps
- o No emulated packet loss
- o RTT: 650 ms
- o Buffer size : BDP

During the transmission of 100 MB on the download path, the test should report the downloading time of 2 MB, 10 MB and 100 MB.

Initial thoughts of the performance obectives for QUIC are the following:

o 3 s for downloading 2 MB

- o 5 s for downloading 10 MB
- o 8 s for downloading 100 MB

Kuhn, et al.Expires September 7, 2020[Page 7]

Internet-Draft

4.4. Lossy large public satellite broadband access

The characteristics of the architecture under test are the following:

- o Satellite downlink path: 250 Mbps
- o Satellite uplink path: 3 Mbps
- o Emulated packet loss on both downlink and uplink paths:
 - * Uniform random transmission link losses: 1%
- o RTT: 650 ms
- o Buffer size : BDP

During the transmission of 100 MB on the download path, the test should report the downloading time of 2 MB, 10 MB and 100 MB.

Initial thoughts of the performance obectives for QUIC are the following:

- o 3 s for downloading 2 MB (uniform transmission link losses)
- o 6 s for downloading 10 MB (uniform transmission link losses)
- o 10 s for downloading 100 MB (uniform transmission link losses)

$\underline{5}$. Mechanisms that improve the performance of QUIC for SATCOM

5.1. Getting up to speed

The advantage of using QUIC is that it includes the TLS and TCP negotiations that reduce the time at which the data can be transmitted. That being said, results of [IJSCN19] illustrate that it will take many RTTs for the congestion controller to increase the rate before it fills the bottleneck capacity. This dominates performance when a path has a large RTT (as in GEO SATCOM networks). There can be an issue in QUIC getting up to speed in a SATCOM context if the congestion and flow controls are not adapted.

The tuning of the initial window described in [<u>I-D.irtf-iccrg-sallantin-initial-spreading</u>] which has been shown to improve performance both for high BDP and more common BDP [<u>CONEXT15</u>][ICC16] could be a relevant solution. That being said, such solution requires the usage of pacing to avoid important bursts of packets in a network that does not have a large BDP.

Kuhn, et al.Expires September 7, 2020[Page 8]

5.2. Maximum window

A large number of in-flight packets are prewired to fully exploit the bottleneck capacity, when there is a large BDP. Default values of maximum windows may not be suitable for a SATCOM context.

Such as presented in [PANRG105], only increasing the initial congestion window is not the only way that can improve QUIC performance in a SATCOM context: increasing maximum congestion windows can also result in much better performance. Other protocol mechanisms also need to be considered, such as flow control at the stream level in QUIC.

5.3. Reliability

Packet loss detection and loss repair take additional time on paths with a larger RTT. This increases the time that a server needs to react to a congestion event. Both can impact the user experience. This happens when a user uses a Wi-Fi link to access a SATCOM terminal. Although the benefits need 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. End-to-end connections may suffer from losses not only in the Wi-Fi segment, but also from congestion losses in the satellite operator ground segment (and even on the Internet itself). Using the mechanisms proposed in [<u>I-D.ferrieux-hamchaoui-quic-lossbits</u>], congestion losses have been identified on the ground segment.

Introducing network coding in QUIC such as proposed in
[I-D.swett-nwcrg-coding-for-quic] and
[I-D.roca-nwcrg-rlc-fec-scheme-for-quic] could help in recovering
from the residual Wi-Fi or congestion losses. Another solution would
be the usage of QUIC tunnels [I-D.schinazi-masque].

5.4. 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 [<u>RFC3449</u>] 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.

6. 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.

One solution would be to provide PEP functions at the termination of the security association (e.g., in a VPN client). Another solution could be to fall-back to using TCP (possibly with TLS or similar methods being used on the transport payload). A final solution could be to deploy and maintain a bespoke protocol tailored to high BDP environments. In the future, we anticipate that fall-back will become less desirable, and methods that rely upon bespoke configurations or protocols will be unattractive. In parallel, new methods such as QUIC will become widely deployed. The opportunity therefore exists to ensure that the new generation of protocols offer acceptable performance over high BDP paths without requiring operating tuning or specific updates by users.

7. Acknowledgements

The authors would like to thank Christian Huitema, Igor Lubashev, Alexandre Ferrieux, Francois Michel, Emmanuel Lochin and the partipants of the IETF106 side-meeting on QUIC for high BDP for their useful feedbacks.

8. 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.

9. Informative References

[CONEXT15]

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

[I-D.ferrieux-hamchaoui-quic-lossbits]

Ferrieux, A. and I. Hamchaoui, "The QUIC Loss Bits", <u>draft-ferrieux-hamchaoui-quic-lossbits-00</u> (work in progress), April 2019.

[I-D.ietf-quic-recovery]

Iyengar, J. and I. Swett, "QUIC Loss Detection and Congestion Control", <u>draft-ietf-quic-recovery-26</u> (work in progress), February 2020.

[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", <u>draft-irtf-iccrg-</u> <u>sallantin-initial-spreading-00</u> (work in progress), January 2014.

[I-D.roca-nwcrg-rlc-fec-scheme-for-quic]

Roca, V., Michel, F., Swett, I., and M. Montpetit, "Sliding Window Random Linear Code (RLC) Forward Erasure Correction (FEC) Schemes for QUIC", <u>draft-roca-nwcrg-rlc-</u> <u>fec-scheme-for-quic-02</u> (work in progress), November 2019.

[I-D.schinazi-masque]

Schinazi, D., "The MASQUE Protocol", <u>draft-schinazi-</u> <u>masque-02</u> (work in progress), January 2020.

[I-D.swett-nwcrg-coding-for-quic]

Swett, I., Montpetit, M., Roca, V., and F. Michel, "Coding for QUIC", <u>draft-swett-nwcrg-coding-for-quic-03</u> (work in progress), July 2019.

[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.

- [IJSCN17] Ahmed, T., Dubois, E., Dupe, JB., Ferrus, R., Gelard, P., and N. Kuhn, "Software-defined satellite cloud RAN", International Journal of Satellite Communications and Networking, 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.

[PANRG105]

Kuhn, N., Stephan, E., Border, J., and G. Fairhurst, "QUIC Over In-sequence Paths with Different Characteristics", IRTF PANRG 105, 2019.

- [RFC2488] Allman, M., Glover, D., and L. Sanchez, "Enhancing TCP Over Satellite Channels using Standard Mechanisms", <u>BCP 28</u>, <u>RFC 2488</u>, DOI 10.17487/RFC2488, January 1999, <<u>https://www.rfc-editor.org/info/rfc2488></u>.
- [RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", <u>RFC 3135</u>, DOI 10.17487/RFC3135, June 2001, <<u>https://www.rfc-editor.org/info/rfc3135</u>>.
- [RFC3449] Balakrishnan, H., Padmanabhan, V., Fairhurst, G., and M. Sooriyabandara, "TCP Performance Implications of Network Path Asymmetry", <u>BCP 69</u>, <u>RFC 3449</u>, DOI 10.17487/RFC3449, December 2002, <<u>https://www.rfc-editor.org/info/rfc3449</u>>.
- [RFC6534] Duffield, N., Morton, A., and J. Sommers, "Loss Episode Metrics for IP Performance Metrics (IPPM)", <u>RFC 6534</u>, DOI 10.17487/RFC6534, May 2012, <<u>https://www.rfc-editor.org/info/rfc6534</u>>.
- [RFC6928] Chu, J., Dukkipati, N., Cheng, Y., and M. Mathis, "Increasing TCP's Initial Window", <u>RFC 6928</u>, DOI 10.17487/RFC6928, April 2013, <<u>https://www.rfc-editor.org/info/rfc6928</u>>.

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

Kuhn, et al.Expires September 7, 2020[Page 13]