

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: November 30, 2020

N. Kuhn
CNES
G. Fairhurst
University of Aberdeen
J. Border
Hughes Network Systems, LLC
E. Stephan
Orange
May 29, 2020

QUIC for SATCOM
draft-kuhn-quic-4-sat-05

Abstract

QUIC has been designed for use across Internet paths. Initial designs of QUIC focused on common deployment scenarios for web traffic. This document focuses on the performance when using a path with a large Bandwidth-Delay Product (BDP).

A large BDP path can result from using a satellite communication (SATCOM) system. The BDP is also high for paths where a satellite network segment is combined with other network technologies (Ethernet, cable modems, WiFi, cellular, radio links, etc), resulting in more complex characteristics. These path characteristics have implications on the efficiency of the network traffic and unless considered in a design can impact the end-to-end quality of experience by the transport service.

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 [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 <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."

This Internet-Draft will expire on November 30, 2020.

Copyright Notice

Copyright (c) 2020 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 (<https://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.

Table of Contents

1.	Introduction	2
2.	Operating over a path with a large BDP	3
3.	TCP Split Solution	5
4.	Regression tests	6
4.1.	Medium public satellite broadband access	6
4.2.	Congested medium public satellite broadband access	7
4.3.	Variable medium public satellite broadband access	7
4.4.	Loss-free large public satellite broadband access	8
4.5.	Lossy large public satellite broadband access	9
5.	Mechanisms that improve the performance of QUIC for SATCOM	9
5.1.	Getting up to speed	9
5.2.	Maximum window	10
5.3.	Reliability	10
5.4.	ACK ratio	10
6.	Discussion	11
7.	Acknowledgments	12
8.	Security Considerations	12
9.	Informative References	12
	Authors' Addresses	14

[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 [[RFC6928](#)]. Although QUIC's Congestion Control (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;
- o 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 Telecommunications 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. These systems typically monitor the link quality in real-time, with the help of known symbol sequences, included along with regular packets, which enable an estimation of the current signal-to-noise ratio. This estimation is then feedback allowing the transmitting link to adapt its coding rate and modulation to 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 often uses all the available bandwidth, possibly with several carriers, to communicate with a set of remote terminals. A carrier is a single Time-Division-Multiplexing channel that multiplexes packets addressed to specific terminals. On the return link, satellite resource is shared among the terminals. Two access methods can be distinguished: on-demand access or contention access. In the former, a terminal receives dedicated transmission resources (usually to send to the gateway). In the latter, some resources are reserved for contention access, where a set of terminals are allowed to compete to obtain transmission 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 path RTT. For example, paths using a satellite system can also exhibit a high loss-rate (e.g., a mobile user or a user behind a Wi-Fi link), where additional delay can impact transport mechanisms.

- o Transport initialization: the 3-way handshake takes a longer time to complete, delaying the time to send data (several transport protocol exchanges may be needed, such as TLS);
- o Size of windows required: to fully exploit the bottleneck capacity, a high BDP requires a larger number of in-flight packets;

- o Reliability: packet loss detection and correction is slow (the performance of end-to-end retransmission is also impacted when using a high RTT path);
- o Getting up to speed: many congestion control methods employ an exponential increase in the sending rate during slow start (for path capacity probing), a high RTT will increase the time to reach a specific rate;
- o Asymmetry : when the links are asymmetric, for various reasons, the the return path may modify the rate and/timing of transport acknowledgment traffic, potentially changing behaviour (e.g., limiting the forward sending rate).

3. TCP Split Solution

High BDP networks commonly break the TCP end-to-end paradigm to adapt the transport protocol. Splitting a TCP connection allows adaptation for a specific use-case and to address 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 significant performance improvement (e.g., a 50% page load time reduction in a SATCOM use-case [[ICCRG100](#)]).

[[NCT13](#)] and [[RFC3135](#)] describe the main functions of a SATCOM TCP split solution. For traffic originated at a gateway to an endpoint connected via a satellite terminal, the TCP split proxy intercepts TCP SYN packets, acting on behalf of the endpoint and adapts the sending rate to the SATCOM scenario. The split solution can specifically tune TCP parameters to the satellite link (latency, available capacity).

When a proxy is used on each side of the satellite link, the transport protocol can be replaced by a protocol other than TCP, optimized for the satellite link. This can be tuned 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 can also recover from packet loss that is local to the part of the connection on which the packet losses occur. 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 modification and is independent of both the client and server sides. This comes with a drawback: TCP splitters are often unable to track end-to-end improvements in protocol mechanisms (e.g., RACK, ECN, TCP Fast Open support). 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, contributing to ossification of the transport system.

4. Regression tests

This section proposes a set of regression tests for QUIC that consider high BDP scenarios. We define by:

- o Download path: from Internet to the client endpoint;
- o Upload path: from the client endpoint to a server (e.g., in the Internet).

4.1. Medium public satellite broadband access

The tested scenario has the following path characteristics:

- o Satellite downlink path: 50 Mbps
- o Satellite uplink path: 10 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 upload and the download time for 2 MB, 10 MB and 100 MB.

Initial thoughts of the performance objectives 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
- o TBD s for uploading 2 MB
- o TBD s for uploading 10 MB
- o TBD s for uploading 100 MB

4.2. Congested medium public satellite broadband access

There are cases where the uplink path is congested or where the capacity of the uplink path is not guaranteed.

The tested scenario has the following path characteristics:

- o Satellite downlink path: 50 Mbps
- o Satellite uplink path: 0.5 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 download time for 2 MB, 10 MB and 100 MB.

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

- o TBD s for downloading 2 MB
- o TBD s for downloading 10 MB
- o TBD s for downloading 100 MB

4.3. Variable medium public satellite broadband access

There are cases where the downlink path is congested or where, due to link layer adaptations to rain fading, the capacity of the downlink path is variable.

The tested scenario has the following path characteristics:

- o Satellite downlink path: 50 Mbps - wait 5s - 10 Mbps
- o Satellite uplink path: 10 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 download time for 2 MB, 10 MB and 100 MB.

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

- o TBD s for downloading 2 MB
- o TBD s for downloading 10 MB
- o TBD s for downloading 100 MB

4.4. Loss-free large public satellite broadband access

The tested scenario has the following path characteristics:

- 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 download time for 2 MB, 10 MB and 100 MB. Then, after 10 seconds, repeat the transmission of 100 MB on the download path where the download time for 2 MB, 10 MB and 100 MB is recorded.

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

- o 3 s for the first downloading 2 MB
- o 5 s for the first downloading 10 MB
- o 8 s for the first downloading 100 MB
- o TBD s for the second downloading 2 MB
- o TBD s for the second downloading 10 MB
- o TBD s for the second downloading 100 MB

4.5. Lossy large public satellite broadband access

The tested scenario has the following path characteristics:

- 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 download time for 2 MB, 10 MB and 100 MB.

Initial thoughts of the performance objectives 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)

5. Mechanisms that improve the performance of QUIC for SATCOM

5.1. Getting up to speed

QUIC has an advantage that the TLS and TCP negotiations can be completed during the transport connection handshake. This can reduce the time to transmit the first data. Results of [[IJSCN19](#)] illustrate that it can still take many RTTs for a CC to increase the sending rate to fill the bottleneck capacity. The delay in getting up to speed can dominate performance for a path with a large RTT, and requires the congestion and flow controls to accommodate the impact of path delay.

One relevant solution is tuning of the initial window described in [[I-D.irtf-icrg-sallantin-initial-spreading](#)], which has been shown to improve performance both for high BDP and more common BDP [[CONEXT15](#)][ICC16]. Such a solution requires using sender pacing to avoid generating bursts of packets in a network.

5.2. Maximum window

The number of in-flight packets required to fill a bottleneck capacity, is dependent on the 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 will consume additional time on paths with a larger RTT. The RTT also determines the time needed by a server to react to a congestion event. Both can impact the user experience. For example, when a user uses a Wi-Fi link to access the Internet via SATCOM terminal.

End-to-end packet Forward Error Correction offers an alternative to retransmission with different tradeoffs in terms of utilised capacity and repair capability.

Network coding as proposed in [\[I-D.swett-nwcrq-coding-for-quic\]](#) and [\[I-D.roca-nwcrq-rlc-fec-scheme-for-quic\]](#) could help QUIC recover from link or congestion loss. Another approach could utilise QUIC tunnels [\[I-D.schinazi-masque\]](#) to apply FEC to all or a part of the end-to-end path.

The benefits of introducing FEC need to be weighed against the additional capacity introduced by end-to-end FEC and the opportunity to use link-local ARQ and/or link-adaptive FEC. A transport connections can suffer link-related losses from a particular link (e.g., Wi-Fi), but also congestion loss (e.g. router buffer overflow in a satellite operator ground segment or along an Internet path). Mechanisms have been proposed in [\[I-D.ferrieux-hamchaoui-quic-lossbits\]](#), to identify congestion losses in the ground segment.

5.4. ACK ratio

Asymmetry in capacity (or in the way capacity is granted to a flow) can lead to cases where the transmission in one direction of communication is restricted by the transmission of the acknowledgment traffic flowing in the opposite direction. A network segment could present limitations in the volume of acknowledgment traffic (e.g.,

limited available return path capacity) or in the number of acknowledgment 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 have been used to mitigate the impact of path asymmetry, primarily targeting operation of TCP.

Many mitigations have been deployed in satellite systems, often as a mechanism within a PEP. Despite their benefits over paths with high asymmetry, 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 (e.g., using an IP VPN).

One simple mitigation is for the remote endpoint to send compound acknowledgments less frequently. A rate of one ACK every $RTT/4$ can significantly reduce this traffic. The QUIC transport specification may evolve to allow the ACK Ratio to be adjusted.

6. Discussion

Many of the issues identified for high BDP paths already exist when using an encrypted transport service over a path that employs encryption at the IP layer. This includes endpoints that utilise IPsec at the network layer, or use VPN technology over a 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 different solution could be to deploy and maintain a bespoke protocol tailored to high BDP environments. In the future, we anticipate that fall-back to TCP 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. Acknowledgments

The authors would like to thank Christian Huitema, Igor Lubashev, Alexandre Ferrieux, Francois Michel, Emmanuel Lochin and the participants 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", [draft-ferrieux-hamchaoui-quic-lossbits-00](#) (work in progress), April 2019.

[I-D.ietf-quic-recovery]

Iyengar, J. and I. Swett, "QUIC Loss Detection and Congestion Control", [draft-ietf-quic-recovery-27](#) (work in progress), March 2020.

[I-D.irtf-iccrgr-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-iccrgr-sallantin-initial-spreading-00](#) (work in progress), January 2014.

[I-D.roca-nwcrgr-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", [draft-roca-nwcrgr-rlc-fec-scheme-for-quic-03](#) (work in progress), March 2020.

[I-D.schinazi-masque]

Schinazi, D., "The MASQUE Protocol", [draft-schinazi-masque-02](#) (work in progress), January 2020.

- [I-D.swett-nwcrq-coding-for-quic]
Swett, I., Montpetit, M., Roca, V., and F. Michel, "Coding for QUIC", [draft-swett-nwcrq-coding-for-quic-04](#) (work in progress), March 2020.
- [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", [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>>.

[RFC6928] Chu, J., Dukkupati, 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>>.

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

