QUIC Internet-Draft Intended status: Standards Track Expires: March 15, 2019

Version-Independent Properties of QUIC draft-ietf-quic-invariants-02

Abstract

This document defines the properties of the QUIC transport protocol that are expected to remain unchanged over time as new versions of the protocol are developed.

Note to Readers

Discussion of this draft takes place on the QUIC working group mailing list (quic@ietf.org), which is archived at https://mailarchive.ietf.org/arch/search/?email_list=quic [1].

Working Group information can be found at https://github.com/quicwg [2]; source code and issues list for this draft can be found at https://github.com/quicwg/base-drafts/labels/-invariants [3].

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 <u>https://datatracker.ietf.org/drafts/current/</u>.

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 March 15, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

Thomson

Expires March 15, 2019

[Page 1]

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
2. Conventions and Definitions
<u>3</u> . An Extremely Abstract Description of QUIC
<u>4</u> . QUIC Packet Headers
<u>4.1</u> . Long Header
<u>4.2</u> . Short Header
<u>4.3</u> . Connection ID
<u>4.4</u> . Version
5. Version Negotiation
6. Security and Privacy Considerations
<u>7</u> . IANA Considerations
<u>8</u> . References
<u>8.1</u> . Normative References
<u>8.2</u> . Informative References
<u>8.3</u> . URIS
Appendix A. Incorrect Assumptions
Author's Address

1. Introduction

In addition to providing secure, multiplexed transport, QUIC [QUIC-TRANSPORT] includes the ability to negotiate a version. This allows the protocol to change over time in response to new requirements. Many characteristics of the protocol will change between versions.

This document describes the subset of QUIC that is intended to remain stable as new versions are developed and deployed. All of these invariants are IP-version-independent.

The primary goal of this document is to ensure that it is possible to deploy new versions of QUIC. By documenting the properties that can't change, this document aims to preserve the ability to change any other aspect of the protocol. Thus, unless specifically described in this document, any aspect of the protocol can change between different versions.

[Page 2]

<u>Appendix A</u> is a non-exhaustive list of some incorrect assumptions that might be made based on knowledge of QUIC version 1; these do not apply to every version of QUIC.

2. Conventions and Definitions

The key words "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 <u>BCP</u> <u>14</u> [<u>RFC2119</u>] [<u>RFC8174</u>] when, and only when, they appear in all capitals, as shown here.

This document uses terms and notational conventions from [<u>QUIC-TRANSPORT</u>].

3. An Extremely Abstract Description of QUIC

QUIC is a connection-oriented protocol between two endpoints. Those endpoints exchange UDP datagrams. These UDP datagrams contain QUIC packets. QUIC endpoints use QUIC packets to establish a QUIC connection, which is shared protocol state between those endpoints.

<u>4</u>. **QUIC** Packet Headers

A QUIC packet is the content of the UDP datagrams exchanged by QUIC endpoints. This document describes the contents of those datagrams.

QUIC defines two types of packet header: long and short. Packets with long headers are identified by the most significant bit of the first octet being set; packets with a short header have that bit cleared.

Aside from the values described here, the payload of QUIC packets is version-specific and of arbitrary length.

4.1. Long Header

Long headers take the form described in Figure 1. Bits that have version-specific semantics are marked with an X.

Expires March 15, 2019 [Page 3]

0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 1 | X X X X X X X | Version (32) |DCIL(4)|SCIL(4)| Destination Connection ID (0/32..144) . . . Source Connection ID (0/32..144) . . .

Figure 1: QUIC Long Header

A QUIC packet with a long header has the high bit of the first octet set to 1. All other bits in that octet are version specific.

The next four octets include a 32-bit Version field (see <u>Section 4.4</u>).

The next octet contains the length in octets of the two Connection IDs (see <u>Section 4.3</u>) that follow. Each length is encoded as a 4-bit unsigned integer. The length of the Destination Connection ID (DCIL) occupies the high bits of the octet and the length of the Source Connection ID (SCIL) occupies the low bits of the octet. An encoded length of 0 indicates that the connection ID is also 0 octets in length. Non-zero encoded lengths are increased by 3 to get the full length of the connection ID; the final value is therefore either 0 or between 4 and 18 octets in length (inclusive). For example, an octet with the value 0xe0 describes a 17 octet Destination Connection ID and a zero octet Source Connection ID.

The connection ID lengths are followed by two connection IDs. The connection ID associated with the recipient of the packet (the Destination Connection ID) is followed by the connection ID associated with the sender of the packet (the Source Connection ID).

The remainder of the packet contains version-specific content.

4.2. Short Header

Short headers take the form described in Figure 2. Bits that have version-specific semantics are marked with an X.

Expires March 15, 2019

[Page 4]

Figure 2: QUIC Short Header

A QUIC packet with a short header has the high bit of the first octet set to 0.

A QUIC packet with a short header includes an optional Destination Connection ID. The short header does not include the Connection ID Lengths, Source Connection ID, or Version fields.

The remainder of the packet has version-specific semantics.

4.3. Connection ID

A connection ID is an opaque field. A connection ID can be 0 octets in length, or between 4 and 18 octets (inclusive).

The primary function of a connection ID is to ensure that changes in addressing at lower protocol layers (UDP, IP, and below) don't cause packets for a QUIC connection to be delivered to the wrong endpoint. The connection ID is used by endpoints and the intermediaries that support them to ensure that each QUIC packet can be delivered to the correct instance of an endpoint. At the endpoint, the connection ID is used to identify which QUIC connection the packet is intended for.

The connection ID is chosen by each endpoint using version-specific methods. Packets for the same QUIC connection might use different connection ID values.

<u>4.4</u>. Version

QUIC versions are identified with a 32-bit integer, encoded in network byte order. Version 0 is reserved for version negotiation (see <u>Section 5</u>). All other version numbers are potentially valid.

The properties described in this document apply to all versions of QUIC. A protocol that does not conform to the properties described in this document is not QUIC. Future documents might describe

Expires March 15, 2019

[Page 5]

QUIC Invariants

additional properties which apply to a specific QUIC version, or to a range of QUIC versions.

5. Version Negotiation

A QUIC endpoint that receives a packet with a long header and a version it either does not understand or does not support might send a Version Negotiation packet in response. Packets with a short header do not trigger version negotiation.

A Version Negotiation packet sets the high bit of the first octet, and thus it conforms with the format of a packet with a long header as defined in <u>Section 4.1</u>. A Version Negotiation packet is identifiable as such by the Version field, which is set to 0x00000000.

0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 1 | X X X X X X X | Version (32) = 0|DCIL(4)|SCIL(4)| Destination Connection ID (0/32..144) . . . Source Connection ID (0/32..144) Supported Version 1 (32) [Supported Version 2 (32)] [Supported Version N (32)]

Figure 3: Version Negotiation Packet

The Version Negotiation packet contains a list of Supported Version fields, each identifying a version that the endpoint sending the packet supports. The Supported Version fields follow the Version field. A Version Negotiation packet contains no other fields. An endpoint MUST ignore a packet that contains no Supported Version fields, or a truncated Supported Version.

Expires March 15, 2019

[Page 6]

Version Negotiation packets do not use integrity or confidentiality protection. A specific QUIC version might authenticate the packet as part of its connection establishment process.

An endpoint MUST include the value from the Source Connection ID field of the packet it receives in the Destination Connection ID field. The value for Source Connection ID MUST be copied from the Destination Connection ID of the received packet, which is initially randomly selected by a client. Echoing both connection IDs gives clients some assurance that the server received the packet and that the Version Negotiation packet was not generated by an off-path attacker.

An endpoint that receives a Version Negotiation packet might change the version that it decides to use for subsequent packets. The conditions under which an endpoint changes QUIC version will depend on the version of QUIC that it chooses.

See [QUIC-TRANSPORT] for a more thorough description of how an endpoint that supports QUIC version 1 generates and consumes a Version Negotiation packet.

<u>6</u>. Security and Privacy Considerations

It is possible that middleboxes could use traits of a specific version of QUIC and assume that when other versions of QUIC exhibit similar traits the same underlying semantic is being expressed. There are potentially many such traits (see <u>Appendix A</u>). Some effort has been made to either eliminate or obscure some observable traits in QUIC version 1, but many of these remain. Other QUIC versions might make different design decisions and so exhibit different traits.

The QUIC version number does not appear in all QUIC packets, which means that reliably extracting information from a flow based on version-specific traits requires that middleboxes retain state for every connection ID they see.

The Version Negotiation packet described in this document is not integrity-protected; it only has modest protection against insertion by off-path attackers. QUIC versions MUST define a mechanism that authenticates the values it contains.

7. IANA Considerations

This document makes no request of IANA.

Expires March 15, 2019

[Page 7]

Internet-Draft

8. References

8.1. Normative References

[QUIC-TRANSPORT]

Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", <u>draft-ietf-quic-</u> <u>transport-14</u> (work in progress), September 2018.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, DOI 10.17487/RFC2119, March 1997, <<u>https://www.rfc-editor.org/info/rfc2119</u>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in <u>RFC</u> 2119 Key Words", <u>BCP 14</u>, <u>RFC 8174</u>, DOI 10.17487/RFC8174, May 2017, <<u>https://www.rfc-editor.org/info/rfc8174</u>>.

8.2. Informative References

[QUIC-TLS]

Thomson, M., Ed. and S. Turner, Ed., "Using Transport Layer Security (TLS) to Secure QUIC", <u>draft-ietf-quic-</u> <u>tls-14</u> (work in progress), September 2018.

[RFC5116] McGrew, D., "An Interface and Algorithms for Authenticated Encryption", <u>RFC 5116</u>, DOI 10.17487/RFC5116, January 2008, <<u>https://www.rfc-editor.org/info/rfc5116</u>>.

8.3. URIS

- [1] https://mailarchive.ietf.org/arch/search/?email_list=quic
- [2] <u>https://github.com/quicwg</u>
- [3] https://github.com/quicwg/base-drafts/labels/-invariants

<u>Appendix A</u>. Incorrect Assumptions

There are several traits of QUIC version 1 [<u>QUIC-TRANSPORT</u>] that are not protected from observation, but are nonetheless considered to be changeable when a new version is deployed.

This section lists a sampling of incorrect assumptions that might be made based on knowledge of QUIC version 1. Some of these statements are not even true for QUIC version 1. This is not an exhaustive list, it is intended to be illustrative only.

Expires March 15, 2019

[Page 8]

The following statements are NOT guaranteed to be true for every QUIC version:

- QUIC uses TLS [<u>QUIC-TLS</u>] and some TLS messages are visible on the wire
- o QUIC long headers are only exchanged during connection establishment
- o Every flow on a given 5-tuple will include a connection establishment phase
- o The first packets exchanged on a flow use the long header
- o QUIC forbids acknowledgments of packets that only contain ACK frames, therefore the last packet before a long period of quiescence might be assumed to contain an acknowledgment
- o QUIC uses an AEAD (AEAD_AES_128_GCM [<u>RFC5116</u>]) to protect the packets it exchanges during connection establishment
- o QUIC packet numbers appear after the Version field
- o QUIC packet numbers increase by one for every packet sent
- QUIC has a minimum size for the first handshake packet sent by a client
- o QUIC stipulates that a client speaks first
- o A QUIC Version Negotiation packet is only sent by a server
- o A QUIC connection ID changes infrequently
- QUIC endpoints change the version they speak if they are sent a Version Negotiation packet
- o The version field in a QUIC long header is the same in both directions
- Only one connection at a time is established between any pair of QUIC endpoints

Author's Address

Expires March 15, 2019 [Page 9]

Martin Thomson Mozilla

Email: martin.thomson@gmail.com