

QUIC
Internet-Draft
Intended status: Standards Track
Expires: February 4, 2018

C. Krasic
Google
August 3, 2017

Header Compression for HTTP over QUIC
draft-krasic-quic-qcram-02

Abstract

The design of the core QUIC transport and the mapping of HTTP semantics over it subsume many HTTP/2 features, prominent among them stream multiplexing and HTTP header compression. A key advantage of the QUIC transport is it provides stream multiplexing free of HoL blocking between streams, while in HTTP/2 multiplexed streams can suffer HoL blocking primarily due to HTTP/2's layering above TCP. However if HPACK is used for header compression, HTTP over QUIC is still vulnerable to HoL blocking, because of how HPACK exploits header redundancies between multiplexed HTTP transactions. This draft defines QCRAM, a variation of HPACK and mechanisms in the QUIC HTTP mapping that allow QUIC implementations the flexibility to avoid header-compression induced HoL blocking.

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 <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 February 4, 2018.

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

Internet-Draft

QCRAM

August 2017

(<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. 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.	QCRAM overview	3
2.1.	Example of HoL blocking	3
2.2.	How QCRAM minimizes HoL blocking	3
3.	HPACK extensions	4
3.1.	Header Block Prefix	4
3.2.	Hybrid absolute-relative indexing	4
3.3.	Preventing Eviction Races	5
3.3.1.	Blocked Evictions	5
3.4.	Handling Stream Resets	5
3.5.	Refreshing Entries with Duplication	6
3.5.1.	Mandatory Entry De-duplication	6
4.	Performance considerations	6
4.1.	Speculative table updates	6
4.2.	Fixed overhead.	7
4.3.	Co-ordinated Packetization	7
5.	Security Considerations	7
6.	IANA Considerations	7
7.	Acknowledgments	7
8.	References	8
8.1.	Normative References	8
8.2.	URIs	8
	Author's Address	8

[1.](#) Introduction

The QUIC transport protocol was designed from the outset to support HTTP semantics, and its design subsumes most of the features of HTTP/2. Two of those features, stream multiplexing and header compression come into some conflict in QUIC. A key goal of the design of QUIC is to improve stream multiplexing relative to HTTP/2, by eliminating HoL (head of line) blocking that can occur in HTTP/2. HoL blocking can happen because HTTP/2 streams are multiplexed onto a

single TCP connection with its in-order semantics. QUIC can maintain independence between streams because it implements core transport functionality in a fully stream-aware manner. However, the HTTP over QUIC mapping is still subject to HoL blocking if HPACK is used directly as in HTTP/2. HPACK exploits multiplexing for greater

compression, shrinking the representation of headers that have appeared earlier on the same connection. In the context of QUIC, this imposes a vulnerability to HoL blocking as will be described more below ([Section 2.1](#)).

QUIC is described in [[QUIC-TRANSPORT](#)]. The HTTP over QUIC mapping is described in [[QUIC-HTTP](#)]. For a full description of HTTP/2, see [[RFC7540](#)]. The description of HPACK is [[RFC7541](#)].

[2.](#) QCRAM overview

Readers may wish to refer to [[RFC7541](#)] [Section 1.3](#) to review HPACK terminology, and [[QUIC-HTTP](#)], [Sections 4](#) on "HTTP over QUIC stream mapping" and [4.2.1](#) on "Header Compression". QCRAM extensions to HPACK allow correctness in the presence of out-of-order delivery, with flexibility to balance between resilience against HoL blocking and compression ratio.

QCRAM is intended to be a relatively non-intrusive extension to HPACK, an implementation should be easily shared within stacks supporting both HTTP/2 over (TLS+)TCP and HTTP over QUIC.

[2.1.](#) Example of HoL blocking

The following is an example of how HPACK can induce HoL blocking in QUIC. Assume two HTTP message exchange streams "A" and "B", and corresponding header blocks "HA" and "HB". Stream "B" experiences HoL blocking due to "A" as follows:

1. HPACK encodes header field "HB[i]" using an index that refers to a table entry that resulted from header field "HA[j]".
2. "HA" and "HB" are delivered via distinct packets that are inflight in the same round trip.
3. "HB"'s packet is delivered but "HA"'s is dropped. HPACK can not

decode "HB" until "HA"'s packet is successfully retransmitted.

2.2. How QCRAM minimizes HoL blocking

Continuing the example, QCRAM's approach is as follows.

1. "HB[i]" will not introduce HoL blocking if "HA[j]" was delivered in a prior round trip. To identify this case, QCRAM assumes that QUIC transport surfaces acknowledgment notifications to the HTTP layer, and that the QCRAM encoder can rely that acknowledged headers have been received by the decoder.

2. "HB[i]" may be represented with one of the Literal variants (see [\[RFC7541\] Section 6.2](#)), trading lower compression ratio for HoL resilience.
3. "HB[i]" may be represented with an Indexed Representation. This favors compression ratio, but the decoder MUST ensure that HB is not decoded until after HA (see blocking in [Section 3.2](#)).

3. HPACK extensions

3.1. Header Block Prefix

In HEADERS and PUSH_PROMISE frames, HPACK Header data should be prefixed by a pair of integers: "Fill" and the "Evictions". "Fill" is the number of entries in the table, and "Evictions" is the cumulative number entries that have been evicted from the table. Their sum is the cumulative number of entries inserted. Each is encoded as a single HPACK integer (8-bit prefix):




Figure 1: Absolute indexing

[Section 3.2](#) describes the role of "Fill" and [Section 3.3](#) covers the

role of "Evictions".

[3.2.](#) Hybrid absolute-relative indexing

HPACK indexed entries refer to an entry by its current position in the dynamic table. As Figure 1 of [RFC7541](#) [1] illustrates, newest entries have smallest indices, and oldest entries are evicted first if the table is full. Under this scheme, each insertion to the table causes the index of all existing entries to change (implicitly). Implicit index updates are acceptable for HTTP/2 because TCP is totally ordered, but it is problematic in the out-of-order context of QUIC.

QCRAM uses a hybrid absolute-relative indexing approach. The prefix defined in [Section 3.1](#) is used by the decoder to interpret all subsequent HPACK instructions at absolute positions for indexed lookups and insertions. It is also used for evictions ([Section 3.3](#)).

As was defined in [Section 2.2](#) case 3, the encoder has the option to select indexed representations that are vulnerable to HoL blocking. Decoder processing of indexed header fields MUST block the encompassing header block if the referenced entry has not been added to the table yet.

To protect against buggy or malicious peers, a timer should be used to set an upper bound on such blocking and in treat expiration of the timer as a decoding error. However, if the implementation chooses not to abort the connection, the remainder of the header block MUST be decoded and output discarded.

[3.3.](#) Preventing Eviction Races

Due to out of order arrival, QCRAM's eviction algorithm requires changes (relative to HPACK) to avoid the possibility that an indexed representation is decoded after the referenced entry is already evicted. QCRAM employs a two-phase eviction algorithm, in which the encoder will not evict entries that have outstanding (unacknowledged) references. The QCRAM encoder maintains a counter as entries are evicted, which is the cumulative number of evictions so far, "Evictions" ([Section 3.1](#)). On arrival at the decoder, if "Evictions"

is higher than previously seen, the decoder MUST evict all entries at or below. Unlike HPACK where the decoder follows the same logic as the encoder to perform evictions, in QCRAM the decoder evicts exclusively based on the encoder's explicit guidance.

[3.3.1.](#) Blocked Evictions

In some cases, the encoder must forgo eviction by selecting a literal representation (blocked eviction), namely in the event that the entry subject to eviction `_is_` referenced by one or more unacknowledged header frames. To assure that the blocked eviction case is rare, a form of thresholding MAY be applied that constrains selection of Indexed representations, such that the oldest entries in the dynamic table will largely be evictable. The constraint is applied when encoding header fields: comparing the cumulative position (in bytes) of the matching entry to a threshold, categorizing oldest entries (past threshold) as at-risk. Avoiding references to at-risk entries, the encoder SHOULD use an Indexed-Duplicate representation instead (see [Section 3.5](#)).

[3.4.](#) Handling Stream Resets

The QCRAM encoder has the option to select representations that might require blocking ([Section 2.2](#) case 3), but the decoder must be prevented from becoming hung if the stream associated with the referenced entry is reset. On stream reset, the QCRAM encoder MUST

check if the stream has unacknowledged headers, and if so resend them on the Control Stream ([\[QUIC-HTTP\] Section 4.1](#)). If header blocks are resent on the control stream, duplicate arrivals are possible due to reset-acknowledgment races. The decoder MUST ignore duplicate header block arrivals, which is straightforward because of unambiguous indexing (see [Section 3.2](#)).

[3.5.](#) Refreshing Entries with Duplication

```
 0 1 2 3 4 5 6 7
+---+---+---+---+
|0|0|1|Index(5+)|
+---+---+-----+
```

Figure 2: Indexed Header Field with Duplication

`_Indexed-Duplicates_` are treated as an Indexed Header Field Representation (see [\[RFC7541\] Section 6.1](#)), additionally inserting a new duplicate entry. [\[RFC7541\]](#) allows duplicate HPACK table entries, that is entries that have the same name and value.

`_Figure 2` annexes the representation for HPACK Dynamic Table Size Update (see [Section 6.3 of RFC7541](#)), which is not supported by HTTP over QUIC.

[3.5.1](#). Mandatory Entry De-duplication

To help mitigate memory consumption due to duplicate entries, HPACK for QCRAM is required to de-duplicate strings in the dynamic table. The table insertion logic should check if the new entry matches any existing entries (name and value), and if so, table accounting **MUST** charge only the overhead portion ([\[RFC7541\] Section 4.1](#)) to the new entry.

Specific de-duplication mechanisms are left to implementations, but using a map in conjunction with reference counted pointers to strings would be typical.

[4](#). Performance considerations

[4.1](#). Speculative table updates

Implementations can `_speculatively_` send header frames on the HTTP Connection Control Stream. Such headers would not be associated with any HTTP transaction, but could be used strategically to improve performance. For instance, the encoder might decide to `_refresh_` by sending Indexed-Duplicate representations for popular header fields

([Section 3.1](#)), ensuring they have small indices and hence minimal size on the wire.

[4.2](#). Fixed overhead.

HPACK defines overhead as 32 bytes ([\[RFC7541\] Section 4.1](#)). QCRAM adds some per-entry state, to track acknowledgment status and eviction rank, and requires mechanisms to de-duplicate strings. A

larger value than 32 might be more accurate for QCRAM.

[4.3.](#) Co-ordinated Packetization

In [Section 2.2](#) case 3, an exception exists when the representation of "HA[i]" and "HB[j]" are delivered within the same transport packet. If so, there is no risk of HoL blocking and using an indexed representation is strictly better than using a literal. An implementation could exploit this exception by employing co-ordination between QCRAM compression and QUIC transport packetization.

[5.](#) Security Considerations

TBD.

[6.](#) IANA Considerations

This document currently makes no request of IANA, and might not need to.

[7.](#) Acknowledgments

This draft draws heavily on the text of [\[RFC7541\]](#). The indirect input of those authors is gratefully acknowledged, as well as ideas from:

- o Mike Bishop
- o Patrick McManus
- o Biren Roy
- o Alan Frindell
- o Ian Swett
- o Ryan Hamilton

[8.](#) References

8.1. Normative References

[QUIC-HTTP]

Bishop, M., Ed., "Hypertext Transfer Protocol (HTTP) over QUIC", August 2017.

[QUIC-TRANSPORT]

Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", August 2017.

[RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", [RFC 7540](#), DOI 10.17487/RFC7540, May 2015, <<http://www.rfc-editor.org/info/rfc7540>>.

[RFC7541] Peon, R. and H. Ruellan, "HPACK: Header Compression for HTTP/2", [RFC 7541](#), DOI 10.17487/RFC7541, May 2015, <<http://www.rfc-editor.org/info/rfc7541>>.

8.2. URIs

[1] <https://tools.ietf.org/html/rfc7541#section-2.3.3>

Author's Address

Charles 'Buck' Krasnic
Google

Email: ckrasnic@google.com