

QUIC  
Internet-Draft  
Intended status: Standards Track  
Expires: November 18, 2021

R. Marx  
KU Leuven  
L. Niccolini, Ed.  
Facebook  
M. Seemann, Ed.  
Protocol Labs  
May 17, 2021

**QUIC event definitions for qlog  
draft-marx-quick-qlog-quick-events-00**

Abstract

This document describes concrete qlog event definitions and their metadata for QUIC events. These events can then be embedded in the higher level schema defined in [[QLOG-MAIN](#)].

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 18, 2021.

Copyright Notice

Copyright (c) 2021 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

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">4</a>
<a href="#">1.1.</a>	<a href="#">Notational Conventions</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Overview</a>	<a href="#">4</a>
<a href="#">2.1.</a>	<a href="#">Links to the main schema</a>	<a href="#">5</a>
<a href="#">2.1.1.</a>	<a href="#">Raw packet and frame information</a>	<a href="#">5</a>
<a href="#">2.1.2.</a>	<a href="#">Events not belonging to a single connection</a>	<a href="#">6</a>
<a href="#">3.</a>	<a href="#">QUIC event definitions</a>	<a href="#">6</a>
<a href="#">3.1.</a>	<a href="#">connectivity</a>	<a href="#">6</a>
<a href="#">3.1.1.</a>	<a href="#">server_listening</a>	<a href="#">6</a>
<a href="#">3.1.2.</a>	<a href="#">connection_started</a>	<a href="#">7</a>
<a href="#">3.1.3.</a>	<a href="#">connection_closed</a>	<a href="#">7</a>
<a href="#">3.1.4.</a>	<a href="#">connection_id_updated</a>	<a href="#">8</a>
<a href="#">3.1.5.</a>	<a href="#">spin_bit_updated</a>	<a href="#">9</a>
<a href="#">3.1.6.</a>	<a href="#">connection_retried</a>	<a href="#">9</a>
<a href="#">3.1.7.</a>	<a href="#">connection_state_updated</a>	<a href="#">9</a>
<a href="#">3.1.8.</a>	<a href="#">MIGRATION-related events</a>	<a href="#">11</a>
<a href="#">3.2.</a>	<a href="#">security</a>	<a href="#">11</a>
<a href="#">3.2.1.</a>	<a href="#">key_updated</a>	<a href="#">11</a>
<a href="#">3.2.2.</a>	<a href="#">key_retired</a>	<a href="#">12</a>
<a href="#">3.3.</a>	<a href="#">transport</a>	<a href="#">12</a>
<a href="#">3.3.1.</a>	<a href="#">version_information</a>	<a href="#">12</a>
<a href="#">3.3.2.</a>	<a href="#">alpn_information</a>	<a href="#">13</a>
<a href="#">3.3.3.</a>	<a href="#">parameters_set</a>	<a href="#">14</a>
<a href="#">3.3.4.</a>	<a href="#">parameters_restored</a>	<a href="#">16</a>
<a href="#">3.3.5.</a>	<a href="#">packet_sent</a>	<a href="#">16</a>
<a href="#">3.3.6.</a>	<a href="#">packet_received</a>	<a href="#">17</a>
<a href="#">3.3.7.</a>	<a href="#">packet_dropped</a>	<a href="#">18</a>
<a href="#">3.3.8.</a>	<a href="#">packet_buffered</a>	<a href="#">19</a>
<a href="#">3.3.9.</a>	<a href="#">packets_acked</a>	<a href="#">20</a>
<a href="#">3.3.10.</a>	<a href="#">datagrams_sent</a>	<a href="#">20</a>
<a href="#">3.3.11.</a>	<a href="#">datagrams_received</a>	<a href="#">21</a>
<a href="#">3.3.12.</a>	<a href="#">datagram_dropped</a>	<a href="#">21</a>
<a href="#">3.3.13.</a>	<a href="#">stream_state_updated</a>	<a href="#">22</a>
<a href="#">3.3.14.</a>	<a href="#">frames_processed</a>	<a href="#">24</a>
<a href="#">3.3.15.</a>	<a href="#">data_moved</a>	<a href="#">25</a>
<a href="#">3.4.</a>	<a href="#">recovery</a>	<a href="#">26</a>
<a href="#">3.4.1.</a>	<a href="#">parameters_set</a>	<a href="#">26</a>
<a href="#">3.4.2.</a>	<a href="#">metrics_updated</a>	<a href="#">26</a>
<a href="#">3.4.3.</a>	<a href="#">congestion_state_updated</a>	<a href="#">27</a>
<a href="#">3.4.4.</a>	<a href="#">loss_timer_updated</a>	<a href="#">28</a>
<a href="#">3.4.5.</a>	<a href="#">packet_lost</a>	<a href="#">29</a>
<a href="#">3.4.6.</a>	<a href="#">marked_for_retransmit</a>	<a href="#">29</a>
<a href="#">4.</a>	<a href="#">Security Considerations</a>	<a href="#">30</a>



<a href="#">5.</a>	<a href="#">IANA Considerations</a>	<a href="#">30</a>
<a href="#">6.</a>	<a href="#">References</a>	<a href="#">30</a>
<a href="#">6.1.</a>	<a href="#">Normative References</a>	<a href="#">30</a>
<a href="#">6.2.</a>	<a href="#">Informative References</a>	<a href="#">31</a>
<a href="#">6.3.</a>	<a href="#">URIs</a>	<a href="#">31</a>
<a href="#">Appendix A.</a>	<a href="#">QUIC data field definitions</a>	<a href="#">31</a>
<a href="#">A.1.</a>	<a href="#">IPAddress</a>	<a href="#">31</a>
<a href="#">A.2.</a>	<a href="#">PacketType</a>	<a href="#">31</a>
<a href="#">A.3.</a>	<a href="#">PacketNumberSpace</a>	<a href="#">32</a>
<a href="#">A.4.</a>	<a href="#">PacketHeader</a>	<a href="#">32</a>
<a href="#">A.5.</a>	<a href="#">Token</a>	<a href="#">32</a>
<a href="#">A.6.</a>	<a href="#">KeyType</a>	<a href="#">33</a>
<a href="#">A.7.</a>	<a href="#">QUIC Frames</a>	<a href="#">33</a>
<a href="#">A.7.1.</a>	<a href="#">PaddingFrame</a>	<a href="#">33</a>
<a href="#">A.7.2.</a>	<a href="#">PingFrame</a>	<a href="#">34</a>
<a href="#">A.7.3.</a>	<a href="#">AckFrame</a>	<a href="#">34</a>
<a href="#">A.7.4.</a>	<a href="#">ResetStreamFrame</a>	<a href="#">35</a>
<a href="#">A.7.5.</a>	<a href="#">StopSendingFrame</a>	<a href="#">35</a>
<a href="#">A.7.6.</a>	<a href="#">CryptoFrame</a>	<a href="#">35</a>
<a href="#">A.7.7.</a>	<a href="#">NewTokenFrame</a>	<a href="#">35</a>
<a href="#">A.7.8.</a>	<a href="#">StreamFrame</a>	<a href="#">36</a>
<a href="#">A.7.9.</a>	<a href="#">MaxDataFrame</a>	<a href="#">36</a>
<a href="#">A.7.10.</a>	<a href="#">MaxStreamDataFrame</a>	<a href="#">36</a>
<a href="#">A.7.11.</a>	<a href="#">MaxStreamsFrame</a>	<a href="#">36</a>
<a href="#">A.7.12.</a>	<a href="#">DataBlockedFrame</a>	<a href="#">37</a>
<a href="#">A.7.13.</a>	<a href="#">StreamDataBlockedFrame</a>	<a href="#">37</a>
<a href="#">A.7.14.</a>	<a href="#">StreamsBlockedFrame</a>	<a href="#">37</a>
<a href="#">A.7.15.</a>	<a href="#">NewConnectionIDFrame</a>	<a href="#">37</a>
<a href="#">A.7.16.</a>	<a href="#">RetireConnectionIDFrame</a>	<a href="#">37</a>
<a href="#">A.7.17.</a>	<a href="#">PathChallengeFrame</a>	<a href="#">38</a>
<a href="#">A.7.18.</a>	<a href="#">PathResponseFrame</a>	<a href="#">38</a>
<a href="#">A.7.19.</a>	<a href="#">ConnectionCloseFrame</a>	<a href="#">38</a>
<a href="#">A.7.20.</a>	<a href="#">HandshakeDoneFrame</a>	<a href="#">38</a>
<a href="#">A.7.21.</a>	<a href="#">UnknownFrame</a>	<a href="#">38</a>
<a href="#">A.7.22.</a>	<a href="#">TransportError</a>	<a href="#">39</a>
<a href="#">A.7.23.</a>	<a href="#">CryptoError</a>	<a href="#">39</a>
<a href="#">Appendix B.</a>	<a href="#">Change Log</a>	<a href="#">39</a>
<a href="#">B.1.</a>	<a href="#">Since <a href="#">draft-marx-qlog-event-definitions-quic-h3-02</a>:</a>	<a href="#">40</a>
<a href="#">B.2.</a>	<a href="#">Since <a href="#">draft-marx-qlog-event-definitions-quic-h3-01</a>:</a>	<a href="#">40</a>
<a href="#">B.3.</a>	<a href="#">Since <a href="#">draft-marx-qlog-event-definitions-quic-h3-00</a>:</a>	<a href="#">41</a>
<a href="#">Appendix C.</a>	<a href="#">Design Variations</a>	<a href="#">42</a>
<a href="#">Appendix D.</a>	<a href="#">Acknowledgements</a>	<a href="#">42</a>
	<a href="#">Authors' Addresses</a>	<a href="#">42</a>



## 1. Introduction

This document describes the values of the qlog name ("category" + "event") and "data" fields and their semantics for the QUIC protocol. This document is based on [draft-34](#) of the QUIC I-Ds QUIC-TRANSPORT [QUIC-RECOVERY] [QUIC-TLS]. HTTP/3 and QPACK events are defined in a separate document [QLOG-H3].

Feedback and discussion are welcome at <https://github.com/quiclog/internet-drafts> [1]. Readers are advised to refer to the "editor's draft" at that URL for an up-to-date version of this document.

Concrete examples of integrations of this schema in various programming languages can be found at <https://github.com/quiclog/qlog/> [2].

### 1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The examples and data definitions in this document are expressed in a custom data definition language, inspired by JSON and TypeScript, and described in [QLOG-MAIN].

## 2. Overview

This document describes the values of the qlog "name" ("category" + "event") and "data" fields and their semantics for the QUIC protocol.

This document assumes the usage of the encompassing main qlog schema defined in [QLOG-MAIN]. Each subsection below defines a separate category (for example connectivity, transport, recovery) and each subsubsection is an event type (for example "packet\_received").

For each event type, its importance and data definition is laid out, often accompanied by possible values for the optional "trigger" field. For the definition and semantics of "importance" and "trigger", see the main schema document.

Most of the complex datastructures, enums and re-usable definitions are grouped together on the bottom of this document for clarity.



## **2.1. Links to the main schema**

This document re-uses all the fields defined in the main qlog schema (e.g., name, category, type, data, group\_id, protocol\_type, the time-related fields, importance, RawInfo, etc.).

One entry in the "protocol\_type" qlog array field MUST be "QUIC" if events from this document are included in a qlog trace.

When the qlog "group\_id" field is used, it is recommended to use QUIC's Original Destination Connection ID (ODCID, the CID chosen by the client when first contacting the server), as this is the only value that does not change over the course of the connection and can be used to link more advanced QUIC packets (e.g., Retry, Version Negotiation) to a given connection. Similarly, the ODCID should be used as the qlog filename or file identifier, potentially suffixed by the vantagepoint type (For example, abcd1234\_server.qlog would contain the server-side trace of the connection with ODCID abcd1234).

### **2.1.1. Raw packet and frame information**

This document re-uses the definition of the RawInfo data class from [\[QLOG-MAIN\]](#).

Note: QUIC packets always include an AEAD authentication tag ("trailer") at the end. As this tag is always the same size for a given connection (it depends on the used TLS cipher), this document does not define a separate "RawInfo:aead\_tag\_length" field here. Instead, this field is reflected in "transport:parameters\_set" and can be logged only once.

Note: As QUIC uses trailers in packets, packet header\_lengths can be calculated as:

$$\text{header\_length} = \text{length} - \text{payload\_length} - \text{aead\_tag\_length}$$

For UDP datagrams, the calculation is simpler:

$$\text{header\_length} = \text{length} - \text{payload\_length}$$

Note: In some cases, the length fields are also explicitly reflected inside of packet headers. For example, the QUIC STREAM frame has a "length" field indicating its payload size. Similarly, the QUIC Long Header has a "length" field which is equal to the payload length plus the packet number length. In these cases, those fields are intentionally preserved in the event definitions. Even though this can lead to duplicate data when the full RawInfo is





logged, it allows a more direct mapping of the QUIC specifications to qlog, making it easier for users to interpret.

### **2.1.2. Events not belonging to a single connection**

For several types of events, it is sometimes impossible to tie them to a specific conceptual QUIC connection (e.g., a `packet_dropped` event triggered because the packet has an unknown `connection_id` in the header). Since qlog events in a trace are typically associated with a single connection, it is unclear how to log these events.

Ideally, implementers SHOULD create a separate, individual "endpoint-level" trace file (or `group_id` value), not associated with a specific connection (for example a "server.qlog" or `group_id = "client"`), and log all events that do not belong to a single connection to this grouping trace. However, this is not always practical, depending on the implementation. Because the semantics of most of these events are well-defined in the protocols and because they are difficult to mis-interpret as belonging to a connection, implementers MAY choose to log events not belonging to a particular connection in any other trace, even those strongly associated with a single connection.

Note that this can make it difficult to match logs from different vantage points with each other. For example, from the client side, it is easy to log connections with version negotiation or retry in the same trace, while on the server they would most likely be logged in separate traces. Servers can take extra efforts (and keep additional state) to keep these events combined in a single trace however (for example by also matching connections on their four-tuple instead of just the connection ID).

## **3. QUIC event definitions**

Each subheading in this section is a qlog event category, while each sub-subheading is a qlog event type. Concretely, for the following two items, we have the category "connectivity" and event type "server\_listening", resulting in a concatenated qlog "name" field value of "connectivity:server\_listening".

### **3.1. connectivity**

#### **3.1.1. server\_listening**

Importance: Extra

Emitted when the server starts accepting connections.

Data:



```
{
  ip_v4?: IPAddress,
  ip_v6?: IPAddress,
  port_v4?: uint32,
  port_v6?: uint32,

  retry_required?:boolean // the server will always answer client initials
with a retry (no 1-RTT connection setups by choice)
}
```

Note: some QUIC stacks do not handle sockets directly and are thus unable to log IP and/or port information.

### **3.1.2. connection\_started**

Importance: Base

Used for both attempting (client-perspective) and accepting (server-perspective) new connections. Note that this event has overlap with `connection_state_updated` and this is a separate event mainly because of all the additional data that should be logged.

Data:

```
{
  ip_version?: "v4" | "v6",
  src_ip?: IPAddress,
  dst_ip?: IPAddress,

  protocol?: string, // transport layer protocol (default "QUIC")
  src_port?: uint32,
  dst_port?: uint32,

  src_cid?: bytes,
  dst_cid?: bytes,
}
```

Note: some QUIC stacks do not handle sockets directly and are thus unable to log IP and/or port information.

### **3.1.3. connection\_closed**

Importance: Base

Used for logging when a connection was closed, typically when an error or timeout occurred. Note that this event has overlap with `connectivity:connection_state_updated`, as well as the `CONNECTION_CLOSE` frame. However, in practice, when analyzing large



deployments, it can be useful to have a single event representing a `connection_closed` event, which also includes an additional `reason` field to provide additional information. Additionally, it is useful to log closures due to timeouts, which are difficult to reflect using the other options.

In QUIC there are two main connection-closing error categories: connection and application errors. They have well-defined error codes and semantics. Next to these however, there can be internal errors that occur that may or may not get mapped to the official error codes in implementation-specific ways. As such, multiple error codes can be set on the same event to reflect this.

```
{
  owner?: "local" | "remote", // which side closed the connection

  connection_code?: TransportError | CryptoError | uint32,
  application_code?: ApplicationError | uint32,
  internal_code?: uint32,

  reason?: string
}
```

Triggers: \* `clean` \* `handshake_timeout` \* `idle_timeout` \* `error` // this is called the "immediate close" in the QUIC specification \* `stateless_reset` \* `version_mismatch` \* `application` // for example HTTP/3's GOAWAY frame

#### [3.1.4.](#) `connection_id_updated`

Importance: Base

This event is emitted when either party updates their current Connection ID. As this typically happens only sparingly over the course of a connection, this event allows loggers to be more efficient than logging the observed CID with each packet in the `.header` field of the `"packet_sent"` or `"packet_received"` events.

This is viewed from the perspective of the one applying the new id. As such, if we receive a new connection id from our peer, we will see the `dst_` fields are set. If we update our own connection id (e.g., `NEW_CONNECTION_ID` frame), we log the `src_` fields.

Data:



```
{
  owner: "local" | "remote",

  old?:bytes,
  new?:bytes,
}
```

#### [3.1.5.](#) spin\_bit\_updated

Importance: Base

To be emitted when the spin bit changes value. It SHOULD NOT be emitted if the spin bit is set without changing its value.

Data:

```
{
  state: boolean
}
```

#### [3.1.6.](#) connection\_retried

TODO

#### [3.1.7.](#) connection\_state\_updated

Importance: Base

This event is used to track progress through QUIC's complex handshake and connection close procedures. It is intended to provide exhaustive options to log each state individually, but also provides a more basic, simpler set for implementations less interested in tracking each smaller state transition. As such, users should not expect to see -all- these states reflected in all qlogs and implementers should focus on support for the SimpleConnectionState set.

Data: ~~~ { old?: ConnectionState | SimpleConnectionState, new: ConnectionState | SimpleConnectionState }

```
enum ConnectionState { attempted, // initial sent/received
peer_validated, // peer address validated by: client sent Handshake
packet OR client used CONNID chosen by the server. transport-draft-
32, section-8.1 handshake_started, early_write, // 1 RTT can be sent,
but handshake isn't done yet handshake_complete, // TLS handshake
complete: Finished received and sent. tls-draft-32, section-4.1.1
handshake_confirmed, // HANDSHAKE_DONE sent/received (connection is
now "active", 1RTT can be sent). tls-draft-32, section-4.1.2 closing,
```





```
draining, // connection_close sent/received closed // draining period
done, connection state discarded }
```

```
enum SimpleConnectionState { attempted, handshake_started,
handshake_confirmed, closed } ~~~
```

These states correspond to the following transitions for both client and server:

**\*Client:\***

- o send initial
  - \* state = attempted
- o get initial
  - \* state = validated \_(not really "needed" at the client, but somewhat useful to indicate progress nonetheless)\_
- o get first Handshake packet
  - \* state = handshake\_started
- o get Handshake packet containing ServerFinished
  - \* state = handshake\_complete
- o send ClientFinished
  - \* state = early\_write (1RTT can now be sent)
- o get HANDSHAKE\_DONE
  - \* state = handshake\_confirmed

**\*Server:\***

- o get initial
  - \* state = attempted
- o send initial \_(don't think this needs a separate state, since some handshake will always be sent in the same flight as this?)\_
- o send handshake EE, CERT, CV, ...
  - \* state = handshake\_started



- o send ServerFinished
  - \* state = early\_write (1RTT can now be sent)
- o get first handshake packet / something using a server-issued CID of min length
  - \* state = validated
- o get handshake packet containing ClientFinished
  - \* state = handshake\_complete
- o send HANDSHAKE\_DONE
  - \* state = handshake\_confirmed

Note: connection\_state\_changed with a new state of "attempted" is the same conceptual event as the connection\_started event above from the client's perspective. Similarly, a state of "closing" or "draining" corresponds to the connection\_closed event.

### **3.1.8. MIGRATION-related events**

e.g., path\_updated

TODO: read up on the draft how migration works and whether to best fit this here or in TRANSPORT TODO: integrate  
<https://tools.ietf.org/html/draft-deconinck-quic-multipath-02>

For now, infer from other connectivity events and path\_challenge/  
path\_response frames

## **3.2. security**

### **3.2.1. key\_updated**

Importance: Base

Note: secret\_updated would be more correct, but in the draft it's called KEY\_UPDATE, so stick with that for consistency

Data:



```
{
  key_type:KeyType,
  old?:bytes,
  new:bytes,
  generation?:uint32 // needed for 1RTT key updates
}
```

Triggers:

- o "tls" // (e.g., initial, handshake and 0-RTT keys are generated by TLS)
- o "remote\_update"
- o "local\_update"

### **3.2.2. key\_retired**

Importance: Base

Data:

```
{
  key_type:KeyType,
  key?:bytes,
  generation?:uint32 // needed for 1RTT key updates
}
```

Triggers:

- o "tls" // (e.g., initial, handshake and 0-RTT keys are dropped implicitly)
- o "remote\_update"
- o "local\_update"

## **3.3. transport**

### **3.3.1. version\_information**

Importance: Core

QUIC endpoints each have their own list of of QUIC versions they support. The client uses the most likely version in their first initial. If the server does support that version, it replies with a version\_negotiation packet, containing supported versions. From this, the client selects a version. This event aggregates all this



information in a single event type. It also allows logging of supported versions at an endpoint without actual version negotiation needing to happen.

Data:

```
{
  server_versions?:Array<bytes>,
  client_versions?:Array<bytes>,
  chosen_version?:bytes
}
```

Intended use:

- o When sending an initial, the client logs this event with `client_versions` and `chosen_version` set
- o Upon receiving a client initial with a supported version, the server logs this event with `server_versions` and `chosen_version` set
- o Upon receiving a client initial with an unsupported version, the server logs this event with `server_versions` set and `client_versions` to the single-element array containing the client's attempted version. The absence of `chosen_version` implies no overlap was found.
- o Upon receiving a version negotiation packet from the server, the client logs this event with `client_versions` set and `server_versions` to the versions in the version negotiation packet and `chosen_version` to the version it will use for the next initial packet

### **3.3.2. alpn\_information**

Importance: Core

QUIC implementations each have their own list of application level protocols and versions thereof they support. The client includes a list of their supported options in its first initial as part of the TLS Application Layer Protocol Negotiation (alpn) extension. If there are common option(s), the server chooses the most optimal one and communicates this back to the client. If not, the connection is closed.

Data:





```
{
  server_alpns?:Array<string>,
  client_alpns?:Array<string>,
  chosen_alpn?:string
}
```

Intended use:

- o When sending an initial, the client logs this event with `client_alpns` set
- o When receiving an initial with a supported alpn, the server logs this event with `server_alpns` set, `client_alpns` equalling the client-provided list, and `chosen_alpn` to the value it will send back to the client.
- o When receiving an initial with an alpn, the client logs this event with `chosen_alpn` to the received value.
- o Alternatively, a client can choose to not log the first event, but wait for the receipt of the server initial to log this event with both `client_alpns` and `chosen_alpn` set.

### **3.3.3. parameters\_set**

Importance: Core

This event groups settings from several different sources (transport parameters, TLS ciphers, etc.) into a single event. This is done to minimize the amount of events and to decouple conceptual setting impacts from their underlying mechanism for easier high-level reasoning.

All these settings are typically set once and never change. However, they are typically set at different times during the connection, so there will typically be several instances of this event with different fields set.

Note that some settings have two variations (one set locally, one requested by the remote peer). This is reflected in the "owner" field. As such, this field **MUST** be correct for all settings included a single event instance. If you need to log settings from two sides, you **MUST** emit two separate event instances.

In the case of connection resumption and 0-RTT, some of the server's parameters are stored up-front at the client and used for the initial connection startup. They are later updated with the server's reply. In these cases, utilize the separate "parameters\_restored" event to



indicate the initial values, and this event to indicate the updated values, as normal.

Data:

```
{
  owner?: "local" | "remote",

  resumption_allowed?: boolean, // valid session ticket was received
  early_data_enabled?: boolean, // early data extension was enabled on the TLS
layer
  tls_cipher?: string, // (e.g., "AES_128_GCM_SHA256")
  aead_tag_length?: uint8, // depends on the TLS cipher, but it's easier to be
explicit. Default value is 16

  // transport parameters from the TLS layer:
  original_destination_connection_id?: bytes,
  initial_source_connection_id?: bytes,
  retry_source_connection_id?: bytes,
  stateless_reset_token?: Token,
  disable_active_migration?: boolean,

  max_idle_timeout?: uint64,
  max_udp_payload_size?: uint32,
  ack_delay_exponent?: uint16,
  max_ack_delay?: uint16,
  active_connection_id_limit?: uint32,

  initial_max_data?: uint64,
  initial_max_stream_data_bidi_local?: uint64,
  initial_max_stream_data_bidi_remote?: uint64,
  initial_max_stream_data_uni?: uint64,
  initial_max_streams_bidi?: uint64,
  initial_max_streams_uni?: uint64,

  preferred_address?: PreferredAddress
}

interface PreferredAddress {
  ip_v4: IPAddress,
  ip_v6: IPAddress,

  port_v4: uint16,
  port_v6: uint16,

  connection_id: bytes,
  stateless_reset_token: Token
}
```



Additionally, this event can contain any number of unspecified fields. This is to reflect setting of for example unknown (greased) transport parameters or employed (proprietary) extensions.

#### **3.3.4. parameters\_restored**

Importance: Base

When using QUIC 0-RTT, clients are expected to remember and restore the server's transport parameters from the previous connection. This event is used to indicate which parameters were restored and to which values when utilizing 0-RTT. Note that not all transport parameters should be restored (many are even prohibited from being re-utilized). The ones listed here are the ones expected to be useful for correct 0-RTT usage.

Data:

```
{
  disable_active_migration?:boolean,

  max_idle_timeout?:uint64,
  max_udp_payload_size?:uint32,
  active_connection_id_limit?:uint32,

  initial_max_data?:uint64,
  initial_max_stream_data_bidi_local?:uint64,
  initial_max_stream_data_bidi_remote?:uint64,
  initial_max_stream_data_uni?:uint64,
  initial_max_streams_bidi?:uint64,
  initial_max_streams_uni?:uint64,
}
```

Note that, like `parameters_set` above, this event can contain any number of unspecified fields to allow for additional/custom parameters.

#### **3.3.5. packet\_sent**

Importance: Core

Data:



```
{
  header:PacketHeader,

  frames?:Array<QuicFrame>, // see appendix for the definitions

  is_coalesced?:boolean, // default value is false

  retry_token?:Token, // only if header.packet_type === retry

  stateless_reset_token?:bytes, // only if header.packet_type ===
stateless_reset. Is always 128 bits in length.

  supported_versions:Array<bytes>, // only if header.packet_type ===
version_negotiation

  raw?:RawInfo,
  datagram_id?:uint32
}
```

Note: We do not explicitly log the `encryption_level` or `packet_number_space`: the `header.packet_type` specifies this by inference (assuming correct implementation)

Triggers:

- o "retransmit\_reordered" // [draft-23](#) 5.1.1
- o "retransmit\_timeout" // [draft-23](#) 5.1.2
- o "pto\_probe" // [draft-23](#) 5.3.1
- o "retransmit\_crypto" // [draft-19](#) 6.2
- o "cc\_bandwidth\_probe" // needed for some CCs to figure out bandwidth allocations when there are no normal sends

Note: for more details on "datagram\_id", see [Section 3.3.10](#). It is only needed when keeping track of packet coalescing.

### **[3.3.6](#). packet\_received**

Importance: Core

Data:





```
{
  header:PacketHeader,

  frames?:Array<QuicFrame>, // see appendix for the definitions

  is_coalesced?:boolean,

  retry_token?:Token, // only if header.packet_type === retry

  stateless_reset_token?:bytes, // only if header.packet_type ===
stateless_reset. Is always 128 bits in length.

  supported_versions:Array<bytes>, // only if header.packet_type ===
version_negotiation

  raw?:RawInfo,
  datagram_id?:uint32
}
```

Note: We do not explicitly log the `encryption_level` or `packet_number_space`: the `header.packet_type` specifies this by inference (assuming correct implementation)

Triggers:

- o "keys\_available" // if packet was buffered because it couldn't be decrypted before

Note: for more details on "datagram\_id", see [Section 3.3.10](#). It is only needed when keeping track of packet coalescing.

### **3.3.7. packet\_dropped**

Importance: Base

This event indicates a QUIC-level packet was dropped after partial or no parsing.

Data:

```
{
  header?:PacketHeader, // primarily packet_type should be filled here, as
other fields might not be parseable

  raw?:RawInfo,
  datagram_id?:uint32
}
```

For this event, the "trigger" field SHOULD be set (for example to one

of the values below), as this helps tremendously in debugging.

Triggers:

- o "key\_unavailable"
- o "unknown\_connection\_id"
- o "header\_parse\_error"
- o "payload\_decrypt\_error"
- o "protocol\_violation"
- o "dos\_prevention"
- o "unsupported\_version"
- o "unexpected\_packet"
- o "unexpected\_source\_connection\_id"
- o "unexpected\_version"
- o "duplicate"
- o "invalid\_initial"

Note: sometimes packets are dropped before they can be associated with a particular connection (e.g., in case of "unsupported\_version"). This situation is discussed more in [Section 2.1.2](#).

Note: for more details on "datagram\_id", see [Section 3.3.10](#). It is only needed when keeping track of packet coalescing.

### **[3.3.8](#). packet\_buffered**

Importance: Base

This event is emitted when a packet is buffered because it cannot be processed yet. Typically, this is because the packet cannot be parsed yet, and thus we only log the full packet contents when it was parsed in a packet\_received event.

Data:



```
{
  header?:PacketHeader, // primarily packet_type and possible packet_number
  should be filled here, as other elements might not be available yet

  raw?:RawInfo,
  datagram_id?:uint32
}
```

Note: for more details on "datagram\_id", see [Section 3.3.10](#). It is only needed when keeping track of packet coalescing.

Triggers:

- o "backpressure" // indicates the parser cannot keep up, temporarily buffers packet for later processing
- o "keys\_unavailable" // if packet cannot be decrypted because the proper keys were not yet available

### [3.3.9](#). **packets\_acked**

Importance: Extra

This event is emitted when a (group of) sent packet(s) is acknowledged by the remote peer for the first time. This information could also be deduced from the contents of received ACK frames. However, ACK frames require additional processing logic to determine when a given packet is acknowledged for the first time, as QUIC uses ACK ranges which can include repeated ACKs. Additionally, this event can be used by implementations that do not log frame contents.

Data: ~~~ { packet\_number\_space?:PacketNumberSpace,  
packet\_numbers?:Array<uint64> } ~~~

Note: if packet\_number\_space is omitted, it assumes the default value of PacketNumberSpace.application\_data, as this is by far the most prevalent packet number space a typical QUIC connection will use.

### [3.3.10](#). **datagrams\_sent**

Importance: Extra

When we pass one or more UDP-level datagrams to the socket. This is useful for determining how QUIC packet buffers are drained to the OS.

Data:



```
{
  count?:uint16, // to support passing multiple at once
  raw?:Array<RawInfo>, // RawInfo:length field indicates total length of the
  datagrams, including UDP header length

  datagram_ids?:Array<uint32>
}
```

Note: QUIC itself does not have a concept of a "datagram\_id". This field is a purely qlog-specific construct to allow tracking how multiple QUIC packets are coalesced inside of a single UDP datagram, which is an important optimization during the QUIC handshake. For this, implementations assign a (per-endpoint) unique ID to each datagram and keep track of which packets were coalesced into the same datagram. As packet coalescing typically only happens during the handshake (as it requires at least one long header packet), this can be done without much overhead.

#### **3.3.11. datagrams\_received**

Importance: Extra

When we receive one or more UDP-level datagrams from the socket. This is useful for determining how datagrams are passed to the user space stack from the OS.

Data:

```
{
  count?:uint16, // to support passing multiple at once
  raw?:Array<RawInfo>, // RawInfo:length field indicates total length of the
  datagrams, including UDP header length

  datagram_ids?:Array<uint32>
}
```

Note: for more details on "datagram\_ids", see [Section 3.3.10](#).

#### **3.3.12. datagram\_dropped**

Importance: Extra

When we drop a UDP-level datagram. This is typically if it does not contain a valid QUIC packet (in that case, use packet\_dropped instead).

Data:





```
{  
  raw?:RawInfo  
}
```

### **3.3.13. stream\_state\_updated**

Importance: Base

This event is emitted whenever the internal state of a QUIC stream is updated, as described in QUIC transport [draft-23 section 3](#). Most of this can be inferred from several types of frames going over the wire, but it's much easier to have explicit signals for these state changes.

Data:



```
{
  stream_id:uint64,
  stream_type?:"unidirectional"|"bidirectional", // mainly useful when
opening the stream

  old?:StreamState,
  new:StreamState,

  stream_side?:"sending"|"receiving"
}

enum StreamState {
  // bidirectional stream states, draft-23 3.4.
  idle,
  open,
  half_closed_local,
  half_closed_remote,
  closed,

  // sending-side stream states, draft-23 3.1.
  ready,
  send,
  data_sent,
  reset_sent,
  reset_received,

  // receive-side stream states, draft-23 3.2.
  receive,
  size_known,
  data_read,
  reset_read,

  // both-side states
  data_received,

  // qlog-defined
  destroyed // memory actually freed
}
```

Note: QUIC implementations SHOULD mainly log the simplified bidirectional (HTTP/2-alike) stream states (e.g., idle, open, closed) instead of the more finegrained stream states (e.g., data\_sent, reset\_received). These latter ones are mainly for more in-depth debugging. Tools SHOULD be able to deal with both types equally.



### **3.3.14. frames\_processed**

Importance: Extra

This event's main goal is to prevent a large proliferation of specific purpose events (e.g., `packets_acknowledged`, `flow_control_updated`, `stream_data_received`). We want to give implementations the opportunity to (selectively) log this type of signal without having to log packet-level details (e.g., in `packet_received`). Since for almost all cases, the effects of applying a frame to the internal state of an implementation can be inferred from that frame's contents, we aggregate these events in this single "frames\_processed" event.

Note: This event can be used to signal internal state change not resulting directly from the actual "parsing" of a frame (e.g., the frame could have been parsed, data put into a buffer, then later processed, then logged with this event).

Note: Implementations logging "packet\_received" and which include all of the packet's constituent frames therein, are not expected to emit this "frames\_processed" event. Rather, implementations not wishing to log full packets or that wish to explicitly convey extra information about when frames are processed (if not directly tied to their reception) can use this event.

Note: for some events, this approach will lose some information (e.g., for which encryption level are packets being acknowledged?). If this information is important, please use the `packet_received` event instead.

Note: in some implementations, it can be difficult to log frames directly, even when using `packet_sent` and `packet_received` events. For these cases, this event also contains the direct `packet_number` field, which can be used to more explicitly link this event to the `packet_sent/received` events.

Data:

```
{
  frames:Array<QuicFrame>, // see appendix for the definitions

  packet_number?:uint64
}
```



### [3.3.15.](#) data\_moved

Importance: Base

Used to indicate when data moves between the different layers (for example passing from the application protocol (e.g., HTTP) to QUIC stream buffers and vice versa) or between the application protocol (e.g., HTTP) and the actual user application on top (for example a browser engine). This helps make clear the flow of data, how long data remains in various buffers and the overheads introduced by individual layers.

For example, this helps make clear whether received data on a QUIC stream is moved to the application protocol immediately (for example per received packet) or in larger batches (for example, all QUIC packets are processed first and afterwards the application layer reads from the streams with newly available data). This in turn can help identify bottlenecks or scheduling problems.

Data:

```
{
  stream_id?:uint64,
  offset?:uint64,
  length?:uint64, // byte length of the moved data

  from?:string, // typically: use either of
"user","application","transport","network"
  to?:string, // typically: use either of
"user","application","transport","network"

  data?:bytes // raw bytes that were transferred
}
```

Note: we do not for example use a "direction" field (with values "up" and "down") to specify the data flow. This is because in some optimized implementations, data might skip some individual layers. Additionally, using explicit "from" and "to" fields is more flexible and allows the definition of other conceptual "layers" (for example to indicate data from QUIC CRYPTO frames being passed to a TLS library ("security") or from HTTP/3 to QPACK ("qpack")).

Note: this event type is part of the "transport" category, but really spans all the different layers. This means we have a few leaky abstractions here (for example, the stream\_id or stream offset might not be available at some logging points, or the raw data might not be in a byte-array form). In these situations, implementers can decide to define new, in-context fields to aid in manual debugging.





### [3.4.](#) recovery

Note: most of the events in this category are kept generic to support different recovery approaches and various congestion control algorithms. Tool creators SHOULD make an effort to support and visualize even unknown data in these events (e.g., plot unknown congestion states by name on a timeline visualization).

#### [3.4.1.](#) parameters\_set

Importance: Base

This event groups initial parameters from both loss detection and congestion control into a single event. All these settings are typically set once and never change. Implementation that do, for some reason, change these parameters during execution, MAY emit the parameters\_set event twice.

Data:

```
{
  // Loss detection, see recovery draft-23, Appendix A.2
  reordering_threshold?:uint16, // in amount of packets
  time_threshold?:float, // as RTT multiplier
  timer_granularity?:uint16, // in ms
  initial_rtt?:float, // in ms

  // congestion control, Appendix B.1.
  max_datagram_size?:uint32, // in bytes // Note: this could be updated after
  pmtud
  initial_congestion_window?:uint64, // in bytes
  minimum_congestion_window?:uint32, // in bytes // Note: this could change
  when max_datagram_size changes
  loss_reduction_factor?:float,
  persistent_congestion_threshold?:uint16 // as PTO multiplier
}
```

Additionally, this event can contain any number of unspecified fields to support different recovery approaches.

#### [3.4.2.](#) metrics\_updated

Importance: Core

This event is emitted when one or more of the observable recovery metrics changes value. This event SHOULD group all possible metric updates that happen at or around the same time in a single event (e.g., if min\_rtt and smoothed\_rtt change at the same time, they should be bundled in a single metrics\_updated entry, rather than



split out into two). Consequently, a `metrics_updated` event is only guaranteed to contain at least one of the listed metrics.

Data:

```
{
  // Loss detection, see recovery draft-23, Appendix A.3
  min_rtt?:float, // in ms or us, depending on the overarching qlog's
configuration
  smoothed_rtt?:float, // in ms or us, depending on the overarching qlog's
configuration
  latest_rtt?:float, // in ms or us, depending on the overarching qlog's
configuration
  rtt_variance?:float, // in ms or us, depending on the overarching qlog's
configuration

  pto_count?:uint16,

  // Congestion control, Appendix B.2.
  congestion_window?:uint64, // in bytes
  bytes_in_flight?:uint64,

  ssthresh?:uint64, // in bytes

  // qlog defined
  packets_in_flight?:uint64, // sum of all packet number spaces

  pacing_rate?:uint64 // in bps
}
```

Note: to make logging easier, implementations MAY log values even if they are the same as previously reported values (e.g., two subsequent `METRIC_UPDATE` entries can both report the exact same value for `min_rtt`). However, applications SHOULD try to log only actual updates to values.

Additionally, this event can contain any number of unspecified fields to support different recovery approaches.

### **[3.4.3](#). congestion\_state\_updated**

Importance: Base

This event signifies when the congestion controller enters a significant new state and changes its behaviour. This event's definition is kept generic to support different Congestion Control algorithms. For example, for the algorithm defined in the Recovery draft ("enhanced" New Reno), the following states are defined:

- o slow\_start
- o congestion\_avoidance

- o application\_limited
- o recovery

Data:

```
{
  old?:string,
  new:string
}
```

The "trigger" field SHOULD be logged if there are multiple ways in which a state change can occur but MAY be omitted if a given state can only be due to a single event occurring (e.g., slow start is exited only when ssthresh is exceeded).

Some triggers for ("enhanced" New Reno):

- o persistent\_congestion
- o ECN

#### [3.4.4. loss\\_timer\\_updated](#)

Importance: Extra

This event is emitted when a recovery loss timer changes state. The three main event types are:

- o set: the timer is set with a delta timeout for when it will trigger next
- o expired: when the timer effectively expires after the delta timeout
- o cancelled: when a timer is cancelled (e.g., all outstanding packets are acknowledged, start idle period)

Note: to indicate an active timer's timeout update, a new "set" event is used.

Data:



```
{
  timer_type?: "ack" | "pto", // called "mode" in draft-23 A.9.
  packet_number_space?: PacketNumberSpace,

  event_type: "set" | "expired" | "cancelled",

  delta?: float // if event_type === "set": delta time in ms or us (see
configuration) from this event's timestamp until when the timer will trigger
}
```

TODO: how about CC algo's that use multiple timers? How generic do these events need to be? Just support QUIC-style recovery from the spec or broader?

TODO: read up on the loss detection logic in [draft-27](#) onward and see if this suffices

#### [3.4.5.](#) packet\_lost

Importance: Core

This event is emitted when a packet is deemed lost by loss detection.

Data:

```
{
  header?: PacketHeader, // should include at least the packet_type and
packet_number

  // not all implementations will keep track of full packets, so these are
optional
  frames?: Array<QuicFrame> // see appendix for the definitions
}
```

For this event, the "trigger" field SHOULD be set (for example to one of the values below), as this helps tremendously in debugging.

Triggers:

- o "reordering\_threshold",
- o "time\_threshold"
- o "pto\_expired" // [draft-23 section 5.3.1](#), MAY

#### [3.4.6.](#) marked\_for\_retransmit

Importance: Extra



This event indicates which data was marked for retransmit upon detecting a packet loss (see packet\_lost). Similar to our reasoning

for the "frames\_processed" event, in order to keep the amount of different events low, we group this signal for all types of retransmittable data in a single event based on existing QUIC frame definitions.

Implementations retransmitting full packets or frames directly can just log the constituent frames of the lost packet here (or do away with this event and use the contents of the packet\_lost event instead). Conversely, implementations that have more complex logic (e.g., marking ranges in a stream's data buffer as in-flight), or that do not track sent frames in full (e.g., only stream offset + length), can translate their internal behaviour into the appropriate frame instance here even if that frame was never or will never be put on the wire.

Note: much of this data can be inferred if implementations log packet\_sent events (e.g., looking at overlapping stream data offsets and length, one can determine when data was retransmitted).

Data:

```
{
  frames:Array<QuicFrame>, // see appendix for the definitions
}
```

#### **4. Security Considerations**

TBD

#### **5. IANA Considerations**

TBD

#### **6. References**

##### **6.1. Normative References**

[QLOG-H3] Marx, R., Ed., Niccolini, L., Ed., and M. Seemann, Ed., "HTTP/3 and QPACK event definitions for qlog", [draft-marx-quic-qlog-h3-events-00](#) (work in progress).

[QLOG-MAIN] Marx, R., Ed., Niccolini, L., Ed., and M. Seemann, Ed., "Main logging schema for qlog", [draft-marx-qlog-main-schema-03](#) (work in progress).



**[QUIC-RECOVERY]**

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

**[QUIC-TLS]**

Thomson, M., Ed. and S. Turner, Ed., "Using Transport Layer Security (TLS) to Secure QUIC", [draft-ietf-quic-tls-latest](#) (work in progress).

**[QUIC-TRANSPORT]**

Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", [draft-ietf-quic-transport-latest](#) (work in progress).

**6.2. Informative 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>>.

**6.3. URIs**

- [1] <https://github.com/quiclog/internet-drafts>
- [2] <https://github.com/quiclog/qlog/>

**Appendix A. QUIC data field definitions****A.1. IPAddress**

```
class IPAddress : string | bytes;
```

```
// an IPAddress can either be a "human readable" form (e.g., "127.0.0.1" for v4  
or "2001:0db8:85a3:0000:0000:8a2e:0370:7334" for v6) or use a raw byte-form (as  
the string forms can be ambiguous)
```

**A.2. PacketType**



```
enum PacketType {
    initial,
    handshake,
    zerortt = "0RTT",
    onertt = "1RTT",
    retry,
    version_negotiation,
    stateless_reset,
    unknown
}
```

### [A.3.](#) PacketNumberSpace

```
enum PacketNumberSpace {
    initial,
    handshake,
    application_data
}
```

### [A.4.](#) PacketHeader

```
class PacketHeader {
    // Note: short vs long header is implicit through PacketType

    packet_type: PacketType;
    packet_number: uint64;

    flags?: uint8; // the bit flags of the packet headers (spin bit, key update
    bit, etc. up to and including the packet number length bits if present)
    interpreted as a single 8-bit integer

    token?: Token; // only if packet_type == initial

    length?: uint16, // only if packet_type == initial || handshake || 0RTT.
    Signifies length of the packet_number plus the payload.

    // only if present in the header
    // if correctly using transport:connection_id_updated events,
    // dcid can be skipped for 1RTT packets
    version?: bytes; // e.g., "ff00001d" for draft-29
    scil?: uint8;
    dcil?: uint8;
    scid?: bytes;
    dcid?: bytes;
}
```

### [A.5.](#) Token



```
class Token {  
    type?: "retry" | "resumption" | "stateless_reset";  
  
    length?: uint32; // byte length of the token  
    data?: bytes; // raw byte value of the token  
  
    details?: any; // decoded fields included in the token (typically: peer's IP  
    address, creation time)  
}
```

The token carried in an Initial packet can either be a retry token from a Retry packet, a stateless reset token from a Stateless Reset packet or one originally provided by the server in a NEW\_TOKEN frame used when resuming a connection (e.g., for address validation purposes). Retry and resumption tokens typically contain encoded metadata to check the token's validity when it is used, but this metadata and its format is implementation specific. For that, this field includes a general-purpose "details" field.

#### [A.6.](#) KeyType

```
enum KeyType {  
    server_initial_secret,  
    client_initial_secret,  
  
    server_handshake_secret,  
    client_handshake_secret,  
  
    server_0rtt_secret,  
    client_0rtt_secret,  
  
    server_1rtt_secret,  
    client_1rtt_secret  
}
```

#### [A.7.](#) QUIC Frames

```
type QuicFrame = PaddingFrame | PingFrame | AckFrame | ResetStreamFrame |  
StopSendingFrame | CryptoFrame | NewTokenFrame | StreamFrame | MaxDataFrame |  
MaxStreamDataFrame | MaxStreamsFrame | DataBlockedFrame |  
StreamDataBlockedFrame | StreamsBlockedFrame | NewConnectionIDFrame |  
RetireConnectionIDFrame | PathChallengeFrame | PathResponseFrame |  
ConnectionCloseFrame | HandshakeDoneFrame | UnknownFrame;
```

##### [A.7.1.](#) PaddingFrame

In QUIC, PADDING frames are simply identified as a single byte of value 0. As such, each padding byte could be theoretically interpreted and logged as an individual PaddingFrame.



However, as this leads to heavy logging overhead, implementations SHOULD instead emit just a single PaddingFrame and set the payload\_length property to the amount of PADDING bytes/frames included in the packet.

```
class PaddingFrame{
    frame_type:string = "padding";

    length?:uint32; // total frame length, including frame header
    payload_length?:uint32;
}
```

#### [A.7.2.](#) PingFrame

```
class PingFrame{
    frame_type:string = "ping";

    length?:uint32; // total frame length, including frame header
    payload_length?:uint32;
}
```

#### [A.7.3.](#) AckFrame

```
class AckFrame{
    frame_type:string = "ack";

    ack_delay?:float; // in ms

    // first number is "from": lowest packet number in interval
    // second number is "to": up to and including // highest packet number in
interval
    // e.g., looks like [[1,2],[4,5]]
    acked_ranges?:Array<[uint64, uint64]|[uint64]>;

    // ECN (explicit congestion notification) related fields (not always
present)
    ect1?:uint64;
    ect0?:uint64;
    ce?:uint64;

    length?:uint32; // total frame length, including frame header
    payload_length?:uint32;
}
```

Note: the packet ranges in AckFrame.acked\_ranges do not necessarily have to be ordered (e.g., [[5,9],[1,4]] is a valid value).

Note: the two numbers in the packet range can be the same (e.g., [120,120] means that packet with number 120 was ACKed). However, in that case, implementers SHOULD log [120] instead and tools MUST be able to deal with both notations.



**[A.7.4.](#) ResetStreamFrame**

```
class ResetStreamFrame{
    frame_type:string = "reset_stream";

    stream_id:uint64;
    error_code:ApplicationError | uint32;
    final_size:uint64; // in bytes

    length?:uint32; // total frame length, including frame header
    payload_length?:uint32;
}
```

**[A.7.5.](#) StopSendingFrame**

```
class StopSendingFrame{
    frame_type:string = "stop_sending";

    stream_id:uint64;
    error_code:ApplicationError | uint32;

    length?:uint32; // total frame length, including frame header
    payload_length?:uint32;
}
```

**[A.7.6.](#) CryptoFrame**

```
class CryptoFrame{
    frame_type:string = "crypto";

    offset:uint64;
    length:uint64;

    payload_length?:uint32;
}
```

**[A.7.7.](#) NewTokenFrame**

```
class NewTokenFrame{
    frame_type:string = "new_token";

    token:Token
}
```



**[A.7.8.](#) StreamFrame**

```
class StreamFrame{
  frame_type:string = "stream";

  stream_id:uint64;

  // These two MUST always be set
  // If not present in the Frame type, log their default values
  offset:uint64;
  length:uint64;

  // this MAY be set any time, but MUST only be set if the value is "true"
  // if absent, the value MUST be assumed to be "false"
  fin?:boolean;

  raw?:bytes;
}
```

**[A.7.9.](#) MaxDataFrame**

```
class MaxDataFrame{
  frame_type:string = "max_data";

  maximum:uint64;
}
```

**[A.7.10.](#) MaxStreamDataFrame**

```
class MaxStreamDataFrame{
  frame_type:string = "max_stream_data";

  stream_id:uint64;
  maximum:uint64;
}
```

**[A.7.11.](#) MaxStreamsFrame**

```
class MaxStreamsFrame{
  frame_type:string = "max_streams";

  stream_type:string = "bidirectional" | "unidirectional";
  maximum:uint64;
}
```



**[A.7.12.](#) DataBlockedFrame**

```
class DataBlockedFrame{
  frame_type:string = "data_blocked";

  limit:uint64;
}
```

**[A.7.13.](#) StreamDataBlockedFrame**

```
class StreamDataBlockedFrame{
  frame_type:string = "stream_data_blocked";

  stream_id:uint64;
  limit:uint64;
}
```

**[A.7.14.](#) StreamsBlockedFrame**

```
class StreamsBlockedFrame{
  frame_type:string = "streams_blocked";

  stream_type:string = "bidirectional" | "unidirectional";
  limit:uint64;
}
```

**[A.7.15.](#) NewConnectionIDFrame**

```
class NewConnectionIDFrame{
  frame_type:string = "new_connection_id";

  sequence_number:uint32;
  retire_prior_to:uint32;

  connection_id_length?:uint8;
  connection_id:bytes;

  stateless_reset_token?:Token;
}
```

**[A.7.16.](#) RetireConnectionIDFrame**

```
class RetireConnectionIDFrame{
  frame_type:string = "retire_connection_id";

  sequence_number:uint32;
}
```





**[A.7.17.](#) PathChallengeFrame**

```
class PathChallengeFrame{
  frame_type:string = "path_challenge";

  data?:bytes; // always 64-bit
}
```

**[A.7.18.](#) PathResponseFrame**

```
class PathResponseFrame{
  frame_type:string = "path_response";

  data?:bytes; // always 64-bit
}
```

**[A.7.19.](#) ConnectionCloseFrame**

raw\_error\_code is the actual, numerical code. This is useful because some error types are spread out over a range of codes (e.g., QUIC's crypto\_error).

```
type ErrorSpace = "transport" | "application";
```

```
class ConnectionCloseFrame{
  frame_type:string = "connection_close";

  error_space?:ErrorSpace;
  error_code?:TransportError | ApplicationError | uint32;
  raw_error_code?:uint32;
  reason?:string;

  trigger_frame_type?:uint64 | string; // For known frame types, the
  appropriate "frame_type" string. For unknown frame types, the hex encoded
  identifier value
}
```

**[A.7.20.](#) HandshakeDoneFrame**

```
class HandshakeDoneFrame{
  frame_type:string = "handshake_done";
}
```

**[A.7.21.](#) UnknownFrame**



```
class UnknownFrame{
    frame_type:string = "unknown";
    raw_frame_type:uint64;

    raw_length?:uint32;
    raw?:bytes;
}
```

#### [A.7.22.](#) **TransportError**

```
enum TransportError {
    no_error,
    internal_error,
    connection_refused,
    flow_control_error,
    stream_limit_error,
    stream_state_error,
    final_size_error,
    frame_encoding_error,
    transport_parameter_error,
    connection_id_limit_error,
    protocol_violation,
    invalid_token,
    application_error,
    crypto_buffer_exceeded
}
```

#### [A.7.23.](#) **CryptoError**

These errors are defined in the TLS document as "A TLS alert is turned into a QUIC connection error by converting the one-byte alert description into a QUIC error code. The alert description is added to 0x100 to produce a QUIC error code from the range reserved for CRYPTO\_ERROR."

This approach maps badly to a pre-defined enum. As such, we define the crypto\_error string as having a dynamic component here, which should include the hex-encoded value of the TLS alert description.

```
enum CryptoError {
    crypto_error_{TLS_ALERT}
}
```

#### [Appendix B.](#) **Change Log**



**B.1. Since [draft-marx-qlog-event-definitions-quic-h3-02](#):**

- o These changes were done in preparation of the adoption of the drafts by the QUIC working group (#137)
- o Split QUIC and HTTP/3 events into two separate documents
- o Moved RawInfo, Importance, Generic events and Simulation events to the main schema document.
- o Changed to/from value options of the "data\_moved" event

**B.2. Since [draft-marx-qlog-event-definitions-quic-h3-01](#):**

Major changes:

- o Moved data\_moved from http to transport. Also made the "from" and "to" fields flexible strings instead of an enum (#111,#65)
- o Moved packet\_type fields to PacketHeader. Moved packet\_size field out of PacketHeader to RawInfo:length (#40)
- o Made events that need to log packet\_type and packet\_number use a header field instead of logging these fields individually
- o Added support for logging retry, stateless reset and initial tokens (#94,#86,#117)
- o Moved separate general event categories into a single category "generic" (#47)
- o Added "transport:connection\_closed" event (#43,#85,#78,#49)
- o Added version\_information and alpn\_information events (#85,#75,#28)
- o Added parameters\_restored events to help clarify 0-RTT behaviour (#88)

Smaller changes:

- o Merged loss\_timer events into one loss\_timer\_updated event
- o Field data types are now strongly defined (#10,#39,#36,#115)
- o Renamed qpack instruction\_received and instruction\_sent to instruction\_created and instruction\_parsed (#114)



- o Updated `qpack:dynamic_table_updated.update_type`. It now has the value "inserted" instead of "added" (#113)
- o Updated `qpack:dynamic_table_updated`. It now has an "owner" field to differentiate encoder vs decoder state (#112)
- o Removed `push_allowed` from `http:parameters_set` (#110)
- o Removed explicit trigger field indications from events, since this was moved to be a generic property of the "data" field (#80)
- o Updated `transport:connection_id_updated` to be more in line with other similar events. Also dropped importance from Core to Base (#45)
- o Added length property to `PaddingFrame` (#34)
- o Added `packet_number` field to `transport:frames_processed` (#74)
- o Added a way to generically log packet header flags (first 8 bits) to `PacketHeader`
- o Added additional guidance on which events to log in which situations (#53)
- o Added "simulation:scenario" event to help indicate simulation details
- o Added "packets\_acked" event (#107)
- o Added "datagram\_ids" to the `datagram_X` and `packet_X` events to allow tracking of coalesced QUIC packets (#91)
- o Extended `connection_state_updated` with more fine-grained states (#49)

**B.3. Since [draft-marx-qlog-event-definitions-quic-h3-00](#):**

- o Event and category names are now all lowercase
- o Added many new events and their definitions
- o "type" fields have been made more specific (especially important for `PacketType` fields, which are now called `packet_type` instead of `type`)
- o Events are given an importance indicator (issue #22)





- o Event names are more consistent and use past tense (issue #21)
- o Triggers have been redefined as properties of the "data" field and updated for most events (issue #23)

#### [Appendix C.](#) Design Variations

TBD

#### [Appendix D.](#) Acknowledgements

Much of the initial work by Robin Marx was done at Hasselt University.

Thanks to Marten Seemann, Jana Iyengar, Brian Trammell, Dmitri Tikhonov, Stephen Petrides, Jari Arkko, Marcus Ihlar, Victor Vasiliev, Mirja Kuehlewind, Jeremy Laine, Kazu Yamamoto, Christian Huitema, and Lucas Pardue for their feedback and suggestions.

#### Authors' Addresses

Robin Marx  
KU Leuven

Email: robin.marx@kuleuven.be

Luca Niccolini (editor)  
Facebook

Email: lniccolini@fb.com

Marten Seemann (editor)  
Protocol Labs

Email: marten@protocol.ai

