

Workgroup: QUIC  
Internet-Draft:  
draft-ietf-quic-qlog-main-schema-05  
Published: 13 February 2023  
Intended Status: Standards Track  
Expires: 17 August 2023  
Authors: R. Marx, Ed.    L. Niccolini, Ed.    M. Seemann, Ed.  
         Akamai                    Meta                    Protocol Labs  
         L. Pardue, Ed.  
         Cloudflare

## **Main logging schema for qlog**

### **Abstract**

This document describes a high-level schema for a standardized logging format called qlog. This format allows easy sharing of data and the creation of reusable visualization and debugging tools. The high-level schema in this document is intended to be protocol-agnostic. Separate documents specify how the format should be used for specific protocol data. The schema is also format-agnostic, and can be represented in for example JSON, csv or protobuf.

### **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 17 August 2023.

### **Copyright Notice**

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. [Introduction](#)
  - 1.1. [Notational Conventions](#)
    - 1.1.1. [Schema definition](#)
    - 1.1.2. [Serialization examples](#)
2. [Design goals](#)
3. [The high level qlog schema](#)
  - 3.1. [Summary](#)
  - 3.2. [traces](#)
  - 3.3. [Individual Trace containers](#)
    - 3.3.1. [Configuration](#)
    - 3.3.2. [vantage point](#)
  - 3.4. [Field name semantics](#)
    - 3.4.1. [Timestamps](#)
    - 3.4.2. [Category and Event Type](#)
    - 3.4.3. [Data](#)
    - 3.4.4. [protocol type](#)
    - 3.4.5. [Triggers](#)
    - 3.4.6. [group id](#)
    - 3.4.7. [common fields](#)
4. [Guidelines for event definition documents](#)
  - 4.1. [Event design guidelines](#)
  - 4.2. [Event importance indicators](#)
  - 4.3. [Custom fields](#)
5. [Generic events and data classes](#)
  - 5.1. [Raw packet and frame information](#)
  - 5.2. [Generic events](#)
    - 5.2.1. [error](#)
    - 5.2.2. [warning](#)
    - 5.2.3. [info](#)
    - 5.2.4. [debug](#)
    - 5.2.5. [verbose](#)
  - 5.3. [Simulation events](#)
    - 5.3.1. [scenario](#)
    - 5.3.2. [marker](#)
6. [Serializing qlog](#)
  - 6.1. [qlog to JSON mapping](#)
    - 6.1.1. [I-JSON](#)
    - 6.1.2. [Truncated values](#)
  - 6.2. [qlog to JSON Text Sequences mapping](#)
    - 6.2.1. [Supporting JSON Text Sequences in tooling](#)
  - 6.3. [Other optimized formatting options](#)
    - 6.3.1. [Data structure optimizations](#)

- [6.3.2. Compression](#)
  - [6.3.3. Binary formats](#)
  - [6.3.4. Overview and summary](#)
- [6.4. Conversion between formats](#)
- [7. Methods of access and generation](#)
  - [7.1. Set file output destination via an environment variable](#)
  - [7.2. Access logs via a well-known endpoint](#)
- [8. Tooling requirements](#)
- [9. Security and privacy considerations](#)
- [10. IANA Considerations](#)
- [11. References](#)
  - [11.1. Normative References](#)
  - [11.2. Informative References](#)
- [Appendix A. Change Log](#)
  - [A.1. Since draft-ietf-quic-qlog-main-schema-04:](#)
  - [A.2. Since draft-ietf-quic-qlog-main-schema-03:](#)
  - [A.3. Since draft-ietf-quic-qlog-main-schema-02:](#)
  - [A.4. Since draft-ietf-quic-qlog-main-schema-01:](#)
  - [A.5. Since draft-ietf-quic-qlog-main-schema-00:](#)
  - [A.6. Since draft-marx-qlog-main-schema-draft-02:](#)
  - [A.7. Since draft-marx-qlog-main-schema-01:](#)
  - [A.8. Since draft-marx-qlog-main-schema-00:](#)
- [Acknowledgements](#)
- [Authors' Addresses](#)

## 1. Introduction

There is currently a lack of an easily usable, standardized endpoint logging format. Especially for the use case of debugging and evaluating modern Web protocols and their performance, it is often difficult to obtain structured logs that provide adequate information for tasks like problem root cause analysis.

This document aims to provide a high-level schema and harness that describes the general layout of an easily usable, shareable, aggregatable and structured logging format. This high-level schema is protocol agnostic, with logging entries for specific protocols and use cases being defined in other documents (see for example [\[QLOG-QUIC\]](#) for QUIC and [\[QLOG-H3\]](#) for HTTP/3 and QPACK-related event definitions).

The goal of this high-level schema is to provide amenities and default characteristics that each logging file should contain (or should be able to contain), such that generic and reusable toolsets can be created that can deal with logs from a variety of different protocols and use cases.

As such, this document contains concepts such as versioning, metadata inclusion, log aggregation, event grouping and log file size reduction techniques.

The qlog schema can be serialized in many ways (e.g., JSON, CBOR, protobuf, etc). This document describes only how to employ [JSON], its subset [I-JSON], and its streamable derivative [JSON-Text-Sequences].

Note to RFC editor: Please remove the follow paragraphs in this section before publication.

Feedback and discussion are welcome at <https://github.com/quicwg/qlog>. 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/>.

### 1.1. Notational Conventions

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

#### 1.1.1. Schema definition

To define events and data structures, all qlog documents use the Concise Data Definition Language [CDDL]. This document uses the basic syntax, the specific text, uint, float32, float64, bool, and any types, as well as the .default, .size, and .regex control operators, the ~ unwrapping operator, and the \$ extension point syntax from [CDDL].

Additionally, this document defines the following custom types for clarity:

```

; CDDL's uint is defined as being 64-bit in size
; but for many protocol fields we want to be more restrictive
; and explicit
uint8 = uint .size 1
uint16 = uint .size 2
uint32 = uint .size 4
uint64 = uint .size 8

; an even-length lowercase string of hexadecimally encoded bytes
; examples: 82dc, 027339, 4cdbfd9bfc
; this is needed because the default CDDL binary string (bytes/bstr)
; is only CBOR and not JSON compatible
hexstring = text .regexp "([0-9a-f]{2})*"

```

Figure 1: Additional CDDL type definitions

All timestamps and time-related values (e.g., offsets) in qlog are logged as float64 in the millisecond resolution.

Other qlog documents can define their own CDDL-compatible (struct) types (e.g., separately for each Packet type that a protocol supports).

Note to RFC editor: Please remove the following text in this section before publication.

The main general CDDL syntax conventions in this document a reader should be aware of for easy reading comprehension are:

\*? obj : this object is optional

\*TypeName1 / TypeName2 : a union of these two types (object can be either type 1 OR type 2)

\*obj: TypeName : this object has this concrete type

\*obj: [\* TypeName] : this object is an array of this type with minimum size of 0 elements

\*obj: [+ TypeName] : this object is an array of this type with minimum size of 1 element

\*TypeName = ... : defines a new type

\*EnumName = "entry1" / "entry2" / entry3 / ...: defines an enum

\*StructName = { ... } : defines a new struct type

\*; : single-line comment

`** text => any` : special syntax to indicate 0 or more fields that have a string key that maps to any value. Used to indicate a generic JSON object.

All timestamps and time-related values (e.g., offsets) in qlog are logged as float64 in the millisecond resolution.

Other qlog documents can define their own CDDL-compatible (struct) types (e.g., separately for each Packet type that a protocol supports).

### 1.1.2. Serialization examples

Serialization examples in this document use JSON ([\[JSON\]](#)) unless otherwise indicated.

## 2. Design goals

The main tenets for the qlog schema design are:

- \*Streamable, event-based logging
- \*A flexible format that can reduce log producer overhead, at the cost of increased complexity for consumers (e.g. tools)
- \*Extensible and pragmatic
- \*Aggregation and transformation friendly (e.g., the top-level element for the non-streaming format is a container for individual traces, `group_ids` can be used to tag events to a particular context)
- \*Metadata is stored together with event data

## 3. The high level qlog schema

A qlog file should be able to contain several individual traces and logs from multiple vantage points that are in some way related. To that end, the top-level element in the qlog schema defines only a small set of "header" fields and an array of component traces. For this document, the required "qlog\_version" field **MUST** have a value of "0.3".

**Note:** there have been several previously broadly deployed qlog versions based on older drafts of this document (see draft-marx-qlog-main-schema). The old values for the "qlog\_version" field were "draft-00", "draft-01" and "draft-02". When qlog was moved to the QUIC working group, we decided to switch to a new versioning scheme which is independent of individual draft

document numbers. However, we did start from 0.3, as conceptually 0.0, 0.1 and 0.2 can map to draft-00, draft-01 and draft-02.

As qlog can be serialized in a variety of ways, the "qlog\_format" field is used to indicate which serialization option was chosen. Its value **MUST** either be one of the options defined in this document (e.g., [Section 6](#)) or the field must be omitted entirely, in which case it assumes the default value of "JSON".

In order to make it easier to parse and identify qlog files and their serialization format, the "qlog\_version" and "qlog\_format" fields and their values **SHOULD** be in the first 256 characters/bytes of the resulting log file.

An example of the qlog file's top-level structure is shown in [Figure 2](#).

Definition:

```
QlogFile = {  
  qlog_version: text  
  ? qlog_format: text .default "JSON"  
  ? title: text  
  ? description: text  
  ? summary: Summary  
  ? traces: [+ Trace / TraceError]  
}
```

Figure 2: QlogFile definition

JSON serialization example:

```
{  
  "qlog_version": "0.3",  
  "qlog_format": "JSON",  
  "title": "Name of this particular qlog file (short)",  
  "description": "Description for this group of traces (long)",  
  "summary": {  
    ...  
  },  
  "traces": [...]  
}
```

Figure 3: QlogFile example

### 3.1. Summary

In a real-life deployment with a large amount of generated logs, it can be useful to sort and filter logs based on some basic summarized or aggregated data (e.g., log length, packet loss rate, log location, presence of error events, ...). The summary field (if present) **SHOULD** be on top of the qlog file, as this allows for the file to be processed in a streaming fashion (i.e., the implementation could just read up to and including the summary field and then only load the full logs that are deemed interesting by the user).

As the summary field is highly deployment-specific, this document does not specify any default fields or their semantics. Some examples of potential entries are shown in [Section 3.1](#).

Definition:

```
Summary = {  
    ; summary can contain any type of custom information  
    ; text here doesn't mean the type text,  
    ; but the fact that keys/names in the objects are strings  
    * text => any  
}
```

Figure 4: Summary definition

JSON serialization example:

```
{  
  "trace_count": 1,  
  "max_duration": 5006,  
  "max_outgoing_loss_rate": 0.013,  
  "total_event_count": 568,  
  "error_count": 2  
}
```

Figure 5: Summary example

### 3.2. traces

It is often advantageous to group several related qlog traces together in a single file. For example, we can simultaneously perform logging on the client, on the server and on a single point on their common network path. For analysis, it is useful to aggregate these three individual traces together into a single file, so it can be uniquely stored, transferred and annotated.



As such, the "traces" array contains a list of individual qlog traces. Typical qlogs will only contain a single trace in this array. These can later be combined into a single qlog file by taking the "traces" entry/entries for each qlog file individually and copying them to the "traces" array of a new, aggregated qlog file. This is typically done in a post-processing step.

The "traces" array can thus contain both normal traces (for the definition of the Trace type, see [Section 3.3](#)), but also "error" entries. These indicate that we tried to find/convert a file for inclusion in the aggregated qlog, but there was an error during the process. Rather than silently dropping the erroneous file, we can opt to explicitly include it in the qlog file as an entry in the "traces" array, as shown in [Figure 6](#).

Definition:

```
TraceError = {
  error_description: text
  ; the original URI at which we attempted to find the file
  ? uri: text
  ? vantage_point: VantagePoint
}
```

Figure 6: TraceError definition

JSON serialization example:

```
{
  "error_description": "File could not be found",
  "uri": "/srv/traces/today/latest.qlog",
  "vantage_point": { type: "server" }
}
```

Figure 7: TraceError example

Note that another way to combine events of different traces in a single qlog file is through the use of the "group\_id" field, discussed in [Section 3.4.6](#).

### 3.3. Individual Trace containers

The exact conceptual definition of a Trace can be fluid. For example, a trace could contain all events for a single connection, for a single endpoint, for a single measurement interval, for a single protocol, etc. As such, a Trace container contains some metadata in addition to the logged events, see [Figure 8](#).

In the normal use case however, a trace is a log of a single data flow collected at a single location or vantage point. For example, for QUIC, a single trace only contains events for a single logical QUIC connection for either the client or the server.

The semantics and context of the trace can mainly be deduced from the entries in the "common\_fields" list and "vantage\_point" field.

Definition:

```
Trace = {  
  ? title: text  
  ? description: text  
  ? configuration: Configuration  
  ? common_fields: CommonFields  
  ? vantage_point: VantagePoint  
  events: [* Event]  
}
```

Figure 8: Trace definition

JSON serialization example:

```
{  
  "title": "Name of this particular trace (short)",  
  "description": "Description for this trace (long)",  
  "configuration": {  
    "time_offset": 150  
  },  
  "common_fields": {  
    "ODCID": "abcde1234",  
    "time_format": "absolute"  
  },  
  "vantage_point": {  
    "name": "backend-67",  
    "type": "server"  
  },  
  "events": [...]  
}
```

Figure 9: Trace example

### 3.3.1. Configuration

We take into account that a qlog file is usually not used in isolation, but by means of various tools. Especially when aggregating various traces together or preparing traces for a demonstration, one might wish to persist certain tool-based settings

inside the qlog file itself. For this, the configuration field is used.

The configuration field can be viewed as a generic metadata field that tools can fill with their own fields, based on per-tool logic. It is best practice for tools to prefix each added field with their tool name to prevent collisions across tools. This document only defines two optional, standard, tool-independent configuration settings: "time\_offset" and "original\_uris".

Definition:

```
Configuration = {  
    ; time_offset is in milliseconds  
    time_offset: float64  
    original_uris:[* text]  
    * text => any  
}
```

Figure 10: Configuration definition

JSON serialization example:

```
{  
  "time_offset": 150,  
  "original_uris": [  
    "https://example.org/trace1.qlog",  
    "https://example.org/trace2.qlog"  
  ]  
}
```

Figure 11: Configuration example

#### 3.3.1.1. time\_offset

The time\_offset field indicates by how many milliseconds the starting time of the current trace should be offset. This is useful when comparing logs taken from various systems, where clocks might not be perfectly synchronous. Users could use manual tools or automated logic to align traces in time and the found optimal offsets can be stored in this field for future usage. The default value is 0.

#### 3.3.1.2. original\_uris

The original\_uris field is used when merging multiple individual qlog files or other source files (e.g., when converting .pcaps to qlog). It allows to keep better track where certain data came from.

It is a simple array of strings. It is an array instead of a single string, since a single qlog trace can be made up out of an aggregation of multiple component qlog traces as well. The default value is an empty array.

#### 3.3.1.3. custom fields

Tools can add optional custom metadata to the "configuration" field to store state and make it easier to share specific data viewpoints and view configurations.

Two examples from the [qvis toolset](#) are shown in [Figure 12](#).

```
{
  "configuration" : {
    "qvis" : {
      "congestion_graph": {
        "startX": 1000,
        "endX": 2000,
        "focusOnEventIndex": 124
      }
      "sequence_diagram" : {
        "focusOnEventIndex": 555
      }
    }
  }
}
```

Figure 12: Custom configuration fields example

#### 3.3.2. vantage\_point

The vantage\_point field describes the vantage point from which the trace originates, see [Figure 13](#). Each trace can have only a single vantage\_point and thus all events in a trace **MUST** BE from the perspective of this vantage\_point. To include events from multiple vantage\_points, implementers can for example include multiple traces, split by vantage\_point, in a single qlog file.

Definitions:

```

VantagePoint = {
    ? name: text
    type: VantagePointType
    ? flow: VantagePointType
}

; client = endpoint which initiates the connection
; server = endpoint which accepts the connection
; network = observer in between client and server
VantagePointType = "client" / "server" / "network" / "unknown"

```

Figure 13: VantagePoint definition

JSON serialization examples:

```

{
  "name": "aioquic client",
  "type": "client"
}

{
  "name": "wireshark trace",
  "type": "network",
  "flow": "client"
}

```

Figure 14: VantagePoint example

The flow field is only required if the type is "network" (for example, the trace is generated from a packet capture). It is used to disambiguate events like "packet sent" and "packet received". This is indicated explicitly because for multiple reasons (e.g., privacy) data from which the flow direction can be otherwise inferred (e.g., IP addresses) might not be present in the logs.

Meaning of the different values for the flow field: \* "client" indicates that this vantage point follows client data flow semantics (a "packet sent" event goes in the direction of the server). \* "server" indicates that this vantage point follow server data flow semantics (a "packet sent" event goes in the direction of the client). \* "unknown" indicates that the flow's direction is unknown.

Depending on the context, tools confronted with "unknown" values in the vantage\_point can either try to heuristically infer the semantics from protocol-level domain knowledge (e.g., in QUIC, the client always sends the first packet) or give the user the option to switch between client and server perspectives manually.

### 3.4. Field name semantics

Inside of the "events" field of a qlog trace is a list of events logged by the endpoint. Each event is specified as a generic object with a number of member fields and their associated data. Depending on the protocol and use case, the exact member field names and their formats can differ across implementations. This section lists the main, pre-defined and reserved field names with specific semantics and expected corresponding value formats.

Each qlog event at minimum requires the "time" ([Section 3.4.1](#)), "name" ([Section 3.4.2](#)) and "data" ([Section 3.4.3](#)) fields. Other typical fields are "time\_format" ([Section 3.4.1](#)), "protocol\_type" ([Section 3.4.4](#)), "trigger" ([Section 3.4.5](#)), and "group\_id" ([Section 3.4.6](#)). As especially these later fields typically have identical values across individual event instances, they are normally logged separately in the "common\_fields" ([Section 3.4.7](#)).

The specific values for each of these fields and their semantics are defined in separate documents, specific per protocol or use case. For example: event definitions for QUIC, HTTP/3 and QPACK can be found in [[QLOG-QUIC](#)] and [[QLOG-H3](#)].

Other fields are explicitly allowed by the qlog approach, and tools **SHOULD** allow for the presence of unknown event fields, but their semantics depend on the context of the log usage (e.g., for QUIC, the ODCID field is used), see [[QLOG-QUIC](#)].

An example of a qlog event with its component fields is shown in [Figure 15](#).

Definition:

```
Event = {
  time: float64
  name: text
  data: $ProtocolEventBody

  ? time_format: TimeFormat

  ? protocol_type: ProtocolType
  ? group_id: GroupID

  ; events can contain any amount of custom fields
  * text => any
}
```

Figure 15: Event definition

JSON serialization:

```
{
  "time": 1553986553572,

  "name": "transport:packet_sent",
  "data": { ... },

  "protocol_type": ["QUIC", "HTTP3"],
  "group_id": "127ecc830d98f9d54a42c4f0842aa87e181a",

  "time_format": "absolute",

  "ODCID": "127ecc830d98f9d54a42c4f0842aa87e181a"
}
```

Figure 16: Event example

#### 3.4.1. Timestamps

The "time" field indicates the timestamp at which the event occurred. Its value is typically the Unix timestamp since the 1970 epoch (number of milliseconds since midnight UTC, January 1, 1970, ignoring leap seconds). However, qlog supports two more succinct timestamps formats to allow reducing file size. The employed format is indicated in the "time\_format" field, which allows one of three values: "absolute", "delta" or "relative".

Definition:

TimeFormat = "absolute" / "delta" / "relative"

Figure 17: TimeFormat definition

\*Absolute: Include the full absolute timestamp with each event. This approach uses the largest amount of characters. This is also the default value of the "time\_format" field.

\*Delta: Delta-encode each time value on the previously logged value. The first event in a trace typically logs the full absolute timestamp. This approach uses the least amount of characters.

\*Relative: Specify a full "reference\_time" timestamp (typically this is done up-front in "common\_fields", see [Section 3.4.7](#)) and include only relatively-encoded values based on this reference\_time with each event. The "reference\_time" value is

typically the first absolute timestamp. This approach uses a medium amount of characters.

The first option is good for stateless loggers, the second and third for stateful loggers. The third option is generally preferred, since it produces smaller files while being easier to reason about. An example for each option can be seen in [Figure 18](#).

The absolute approach will use:

1500, 1505, 1522, 1588

The delta approach will use:

1500, 5, 17, 66

The relative approach will:

- set the `reference_time` to 1500 in "common\_fields"
- use: 0, 5, 22, 88

Figure 18: Three different approaches for logging timestamps

One of these options is typically chosen for the entire trace (put differently: each event has the same value for the "time\_format" field). Each event **MUST** include a timestamp in the "time" field.

Events in each individual trace **SHOULD** be logged in strictly ascending timestamp order (though not necessarily absolute value, for the "delta" format). Tools CAN sort all events on the timestamp before processing them, though are not required to (as this could impose a significant processing overhead). This can be a problem especially for multi-threaded and/or streaming loggers, who could consider using a separate post-processor to order qlog events in time if a tool do not provide this feature.

Timestamps do not have to use the UNIX epoch timestamp as their reference. For example for privacy considerations, any initial reference timestamps (for example "endpoint uptime in ms" or "time since connection start in ms") can be chosen. Tools **SHOULD NOT** assume the ability to derive the absolute Unix timestamp from qlog traces, nor allow on them to relatively order events across two or more separate traces (in this case, clock drift should also be taken into account).

### 3.4.2. Category and Event Type

Events differ mainly in the type of metadata associated with them. To help identify a given event and how to interpret its metadata in the "data" field (see [Section 3.4.3](#)), each event has an associated "name" field. This can be considered as a concatenation of two other fields, namely event "category" and event "type".



Category allows a higher-level grouping of events per specific event type. For example for QUIC and HTTP/3, the different categories could be "transport", "http", "qpack", and "recovery". Within these categories, the event Type provides additional granularity. For example for QUIC and HTTP/3, within the "transport" Category, there would be "packet\_sent" and "packet\_received" events.

Logging category and type separately conceptually allows for fast and high-level filtering based on category and the re-use of event types across categories. However, it also considerably inflates the log size and this flexibility is not used extensively in practice at the time of writing.

As such, the default approach in qlog is to concatenate both field values using the ":" character in the "name" field, as can be seen in [Figure 19](#). As such, qlog category and type names **MUST NOT** include this character.

JSON serialization using separate fields:

```
{
  "category": "transport",
  "type": "packet_sent"
}
```

JSON serialization using ":" concatenated field:

```
{
  "name": "transport:packet_sent"
}
```

Figure 19: Ways of logging category, type and name of an event.

Certain serializations CAN emit category and type as separate fields, and qlog tools **SHOULD** be able to deal with both the concatenated "name" field, and the separate "category" and "type" fields. Text-based serializations however are encouraged to employ the concatenated "name" field for efficiency.

### 3.4.3. Data

The data field is a generic object. It contains the per-event metadata and its form and semantics are defined per specific sort of event. For example, data field value definitions for QUIC and HTTP/3 can be found in [\[QLOG-QUIC\]](#) and [\[QLOG-H3\]](#).

This field is defined here as a CDDL extension point (a "socket" or "plug") named \$ProtocolEventBody. Other documents **MUST** properly extend this extension point when defining new data field content options to enable automated validation of aggregated qlog schemas.

The only common field defined for the data field is the trigger field, which is discussed in [Section 3.4.5](#).

Definition:

```
; The ProtocolEventBody is any key-value map (e.g., JSON object)
; only the optional trigger field is defined in this document
$ProtocolEventBody /= {
    ? trigger: text
    * text => any
}
; event documents are intended to extend this socket by using:
; NewProtocolEvents = EventType1 / EventType2 / ... / EventTypeN
; $ProtocolEventBody /= NewProtocolEvents
```

Figure 20: ProtocolEventBody definition

One purely illustrative example for a QUIC "packet\_sent" event is shown in [Figure 21](#):

```
TransportPacketSent = {
    ? packet_size: uint16
    header: PacketHeader
    ? frames:[* QuicFrame]
    ? trigger: "pto_probe" / "retransmit_timeout" / "bandwidth_probe"
}
```

could be serialized as

```
{
  "packet_size": 1280,
  "header": {
    "packet_type": "1RTT",
    "packet_number": 123
  },
  "frames": [
    {
      "frame_type": "stream",
      "length": 1000,
      "offset": 456
    },
    {
      "frame_type": "padding"
    }
  ]
}
```

Figure 21: Example of the 'data' field for a QUIC packet\_sent event

#### 3.4.4. protocol\_type

The "protocol\_type" array field indicates to which protocols (or protocol "stacks") this event belongs. This allows a single qlog file to aggregate traces of different protocols (e.g., a web server offering both TCP+HTTP/2 and QUIC+HTTP/3 connections).

Definition:

```
ProtocolType = [+ text]
```

Figure 22: ProtocolType definition

For example, QUIC and HTTP/3 events have the "QUIC" and "HTTP3" protocol\_type entry values, see [\[QLOG-QUIC\]](#) and [\[QLOG-H3\]](#).

Typically however, all events in a single trace are of the same few protocols, and this array field is logged once in "common\_fields", see [Section 3.4.7](#).

#### 3.4.5. Triggers

Sometimes, additional information is needed in the case where a single event can be caused by a variety of other events. In the normal case, the context of the surrounding log messages gives a hint as to which of these other events was the cause. However, in highly-parallel and optimized implementations, corresponding log messages might be separated in time. Another option is to explicitly indicate these "triggers" in a high-level way per-event to get more fine-grained information without much additional overhead.

In qlog, the optional "trigger" field contains a string value describing the reason (if any) for this event instance occurring, see [Section 3.4.3](#). While this "trigger" field could be a property of the qlog Event itself, it is instead a property of the "data" field instead. This choice was made because many event types do not include a trigger value, and having the field at the Event-level would cause overhead in some serializations. Additional information on the trigger can be added in the form of additional member fields of the "data" field value, yet this is highly implementation-specific, as are the trigger field's string values.

One purely illustrative example of some potential triggers for QUIC's "packet\_dropped" event is shown in [Figure 23](#):

```

TransportPacketDropped = {
    ? packet_type: PacketType
    ? raw_length: uint16

    ? trigger: "key_unavailable" / "unknown_connection_id" /
              "decrypt_error" / "unsupported_version"
}

```

Figure 23: Trigger example

#### 3.4.6. group\_id

As discussed in [Section 3.3](#), a single qlog file can contain several traces taken from different vantage points. However, a single trace from one endpoint can also contain events from a variety of sources. For example, a server implementation might choose to log events for all incoming connections in a single large (streamed) qlog file. As such, we need a method for splitting up events belonging to separate logical entities.

The simplest way to perform this splitting is by associating a "group identifier" to each event that indicates to which conceptual "group" each event belongs. A post-processing step can then extract events per group. However, this group identifier can be highly protocol and context-specific. In the example above, we might use QUIC's "Original Destination Connection ID" to uniquely identify a connection. As such, they might add a "ODCID" field to each event. However, a middlebox logging IP or TCP traffic might rather use four-tuples to identify connections, and add a "four\_tuple" field.

As such, to provide consistency and ease of tooling in cross-protocol and cross-context setups, qlog instead defines the common "group\_id" field, which contains a string value. Implementations are free to use their preferred string serialization for this field, so long as it contains a unique value per logical group. Some examples can be seen in [Figure 25](#).

Definition:

```

GroupID = text

```

Figure 24: GroupID definition

JSON serialization example for events grouped by four tuples and QUIC connection IDs:

```

"events": [
  {
    "time": 1553986553579,
    "protocol_type": ["TCP", "TLS", "HTTP2"],
    "group_id": "ip1=2001:67c:1232:144:9498:6df6:f450:110b,
                ip2=2001:67c:2b0:1c1::198,port1=59105,port2=80",
    "name": "transport:packet_received",
    "data": { ... }
  },
  {
    "time": 1553986553581,
    "protocol_type": ["QUIC", "HTTP3"],
    "group_id": "127ecc830d98f9d54a42c4f0842aa87e181a",
    "name": "transport:packet_sent",
    "data": { ... }
  }
]

```

Figure 25: GroupID example

Note that in some contexts (for example a Multipath transport protocol) it might make sense to add additional contextual per-event fields (for example "path\_id"), rather than use the group\_id field for that purpose.

Note also that, typically, a single trace only contains events belonging to a single logical group (for example, an individual QUIC connection). As such, instead of logging the "group\_id" field with an identical value for each event instance, this field is typically logged once in "common\_fields", see [Section 3.4.7](#).

### 3.4.7. common\_fields

As discussed in the previous sections, information for a typical qlog event varies in three main fields: "time", "name" and associated data. Additionally, there are also several more advanced fields that allow mixing events from different protocols and contexts inside of the same trace (for example "protocol\_type" and "group\_id"). In most "normal" use cases however, the values of these advanced fields are consistent for each event instance (for example, a single trace contains events for a single QUIC connection).

To reduce file size and making logging easier, qlog uses the "common\_fields" list to indicate those fields and their values that are shared by all events in this component trace. This prevents these fields from being logged for each individual event. An example of this is shown in [Figure 26](#).

JSON serialization with repeated field values  
per-event instance:

```
{
  "events": [{
    "group_id": "127ecc830d98f9d54a42c4f0842aa87e181a",
    "protocol_type": ["QUIC", "HTTP3"],
    "time_format": "relative",
    "reference_time": 1553986553572,

    "time": 2,
    "name": "transport:packet_received",
    "data": { ... }
  }, {
    "group_id": "127ecc830d98f9d54a42c4f0842aa87e181a",
    "protocol_type": ["QUIC", "HTTP3"],
    "time_format": "relative",
    "reference_time": 1553986553572,

    "time": 7,
    "name": "http:frame_parsed",
    "data": { ... }
  }
]
}
```

JSON serialization with repeated field values instead  
extracted to common\_fields:

```
{
  "common_fields": {
    "group_id": "127ecc830d98f9d54a42c4f0842aa87e181a",
    "protocol_type": ["QUIC", "HTTP3"],
    "time_format": "relative",
    "reference_time": 1553986553572
  },
  "events": [
    {
      "time": 2,
      "name": "transport:packet_received",
      "data": { ... }
    }, {
      "time": 7,
      "name": "http:frame_parsed",
      "data": { ... }
    }
  ]
}
```

Figure 26: CommonFields example

The "common\_fields" field is a generic dictionary of key-value pairs, where the key is always a string and the value can be of any type, but is typically also a string or number. As such, unknown entries in this dictionary **MUST** be disregarded by the user and tools (i.e., the presence of an unknown field is explicitly NOT an error).

The list of default qlog fields that are typically logged in common\_fields (as opposed to as individual fields per event instance) are shown in the listing below:

Definition:

```
CommonFields = {  
    ? time_format: TimeFormat  
    ? reference_time: float64  
  
    ? protocol_type: ProtocolType  
    ? group_id: GroupID  
  
    * text => any  
}
```

Figure 27: CommonFields definition

Tools **MUST** be able to deal with these fields being defined either on each event individually or combined in common\_fields. Note that if at least one event in a trace has a different value for a given field, this field **MUST NOT** be added to common\_fields but instead defined on each event individually. Good example of such fields are "time" and "data", who are divergent by nature.

#### 4. Guidelines for event definition documents

This document only defines the main schema for the qlog format. This is intended to be used together with specific, per-protocol event definitions that specify the name (category + type) and data needed for each individual event. This is with the intent to allow the qlog main schema to be easily re-used for several protocols. Examples include the QUIC event definitions [[QLOG-QUIC](#)] and HTTP/3 and QPACK event definitions [[QLOG-H3](#)].

This section defines some basic annotations and concepts the creators of event definition documents **SHOULD** follow to ensure a measure of consistency, making it easier for qlog implementers to extrapolate from one protocol to another.

#### 4.1. Event design guidelines

TODO: pending QUIC working group discussion. This text reflects the initial (qlog draft 01 and 02) setup.

There are several ways of defining qlog events. In practice, we have seen two main types used so far: a) those that map directly to concepts seen in the protocols (e.g., `packet_sent`) and b) those that act as aggregating events that combine data from several possible protocol behaviors or code paths into one (e.g., `parameters_set`). The latter are typically used as a means to reduce the amount of unique event definitions, as reflecting each possible protocol event as a separate qlog entity would cause an explosion of event types.

Additionally, logging duplicate data is typically prevented as much as possible. For example, packet header values that remain consistent across many packets are split into separate events (for example `spin_bit_updated` or `connection_id_updated` for QUIC).

Finally, we have typically refrained from adding additional state change events if those state changes can be directly inferred from data on the wire (for example flow control limit changes) if the implementation is bug-free and spec-compliant. Exceptions have been made for common events that benefit from being easily identifiable or individually logged (for example `packets_acked`).

#### 4.2. Event importance indicators

Depending on how events are designed, it may be that several events allow the logging of similar or overlapping data. For example the separate QUIC `connection_started` event overlaps with the more generic `connection_state_updated`. In these cases, it is not always clear which event should be logged or used, and which event should take precedence if e.g., both are present and provide conflicting information.

To aid in this decision making, we recommend that each event **SHOULD** have an "importance indicator" with one of three values, in decreasing order of importance and expected usage:

- \*Core

- \*Base

- \*Extra

The "Core" events are the events that **SHOULD** be present in all qlog files for a given protocol. These are typically tied to basic packet and frame parsing and creation, as well as listing basic internal metrics. Tool implementers **SHOULD** expect and add support for these



events, though **SHOULD NOT** expect all Core events to be present in each qlog trace.

The "Base" events add additional debugging options and CAN be present in qlog files. Most of these can be implicitly inferred from data in Core events (if those contain all their properties), but for many it is better to log the events explicitly as well, making it clearer how the implementation behaves. These events are for example tied to passing data around in buffers, to how internal state machines change and help show when decisions are actually made based on received data. Tool implementers **SHOULD** at least add support for showing the contents of these events, if they do not handle them explicitly.

The "Extra" events are considered mostly useful for low-level debugging of the implementation, rather than the protocol. They allow more fine-grained tracking of internal behavior. As such, they CAN be present in qlog files and tool implementers CAN add support for these, but they are not required to.

Note that in some cases, implementers might not want to log for example data content details in the "Core" events due to performance or privacy considerations. In this case, they **SHOULD** use (a subset of) relevant "Base" events instead to ensure usability of the qlog output. As an example, implementations that do not log QUIC packet\_received events and thus also not which (if any) ACK frames the packet contains, **SHOULD** log packets\_acked events instead.

Finally, for event types whose data (partially) overlap with other event types' definitions, where necessary the event definition document should include explicit guidance on which to use in specific situations.

#### 4.3. Custom fields

Event definition documents are free to define new category and event types, top-level fields (e.g., a per-event field indicating its privacy properties or path\_id in multipath protocols), as well as values for the "trigger" property within the "data" field, or other member fields of the "data" field, as they see fit.

They however **SHOULD NOT** expect non-specialized tools to recognize or visualize this custom data. However, tools **SHOULD** make an effort to visualize even unknown data if possible in the specific tool's context. If they do not, they **MUST** ignore these unknown fields.

#### 5. Generic events and data classes

There are some event types and data classes that are common across protocols, applications and use cases that benefit from being

defined in a single location. This section specifies such common definitions.

### 5.1. Raw packet and frame information

While qlog is a high-level logging format, it also allows the inclusion of most raw wire image information, such as byte lengths and byte values. This is useful when for example investigating or tuning packetization behavior or determining encoding/framing overheads. However, these fields are not always necessary, can take up considerable space, and can have a considerable privacy and security impact (see [Section 9](#)). Where applicable, these fields are grouped in a separate, optional, field named "raw" of type RawInfo. The exact definition of entities, headers, trailers and payloads depend on the protocol used.

Definition:

```
RawInfo = {  
    ; the full byte length of the entity (e.g., packet or frame),  
    ; including possible headers and trailers  
    ? length: uint64  
  
    ; the byte length of the entity's payload,  
    ; excluding possible headers or trailers  
    ? payload_length: uint64  
  
    ; the (potentially truncated) contents of the full entity,  
    ; including headers and possibly trailers  
    ? data: hexstring  
}
```

Figure 28: RawInfo definition

The RawInfo:data field can be truncated for privacy or security purposes, see [Section 6.1.2](#). In this case, the length and payload\_length fields should still indicate the non-truncated lengths when used for debugging purposes.

This document does not specify explicit header\_length or trailer\_length fields. In protocols without trailers, header\_length can be calculated by subtracting the payload\_length from the length. In protocols with trailers (e.g., QUIC's AEAD tag), event definition documents **SHOULD** define how to support header\_length calculation.

### 5.2. Generic events

In typical logging setups, users utilize a discrete number of well-defined logging categories, levels or severities to log freeform

(string) data. This generic events category replicates this approach to allow implementations to fully replace their existing text-based logging by qlog. This is done by providing events to log generic strings for the typical well-known logging levels (error, warning, info, debug, verbose).

For the events defined below, the "category" is "generic" and their "type" is the name of the heading in lowercase (e.g., the "name" of the error event is "generic:error").

#### **5.2.1. error**

Importance: Core

Used to log details of an internal error that might not get reflected on the wire.

Definition:

```
GenericError = {  
    ? code: uint64  
    ? message: text  
}
```

Figure 29: GenericError definition

#### **5.2.2. warning**

Importance: Base

Used to log details of an internal warning that might not get reflected on the wire.

Definition:

```
GenericWarning = {  
    ? code: uint64  
    ? message: text  
}
```

Figure 30: GenericWarning definition

#### **5.2.3. info**

Importance: Extra

Used mainly for implementations that want to use qlog as their one and only logging format but still want to support unstructured string messages.

Definition:

```
GenericInfo = {  
    message: text  
}
```

Figure 31: GenericInfo definition

#### 5.2.4. debug

Importance: Extra

Used mainly for implementations that want to use qlog as their one and only logging format but still want to support unstructured string messages.

Definition:

```
GenericDebug = {  
    message: text  
}
```

Figure 32: GenericDebug definition

#### 5.2.5. verbose

Importance: Extra

Used mainly for implementations that want to use qlog as their one and only logging format but still want to support unstructured string messages.

Definition:

```
GenericVerbose = {  
    message: text  
}
```

Figure 33: GenericVerbose definition

### 5.3. Simulation events

When evaluating a protocol implementation, one typically sets up a series of interoperability or benchmarking tests, in which the test

situations can change over time. For example, the network bandwidth or latency can vary during the test, or the network can be fully disable for a short time. In these setups, it is useful to know when exactly these conditions are triggered, to allow for proper correlation with other events.

For the events defined below, the "category" is "simulation" and their "type" is the name of the heading in lowercase (e.g., the "name" of the scenario event is "simulation:scenario").

#### 5.3.1. scenario

Importance: Extra

Used to specify which specific scenario is being tested at this particular instance. This could also be reflected in the top-level qlog's summary or configuration fields, but having a separate event allows easier aggregation of several simulations into one trace (e.g., split by group\_id).

Definition:

```
SimulationScenario = {  
  ? name: text  
  ? details: { * text => any }  
}
```

Figure 34: SimulationScenario definition

#### 5.3.2. marker

Importance: Extra

Used to indicate when specific emulation conditions are triggered at set times (e.g., at 3 seconds in 2% packet loss is introduced, at 10s a NAT rebind is triggered).

Definition:

```
SimulationMarker = {  
  ? type: text  
  ? message: text  
}
```

Figure 35: SimulationMarker definition

## 6. Serializing qlog

This document and other related qlog schema definitions are intentionally serialization-format agnostic. This means that implementers themselves can choose how to represent and serialize qlog data practically on disk or on the wire. Some examples of possible formats are JSON, CBOR, CSV, protocol buffers, flatbuffers, etc.

All these formats make certain tradeoffs between flexibility and efficiency, with textual formats like JSON typically being more flexible but also less efficient than binary formats like protocol buffers. The format choice will depend on the practical use case of the qlog user. For example, for use in day to day debugging, a plaintext readable (yet relatively large) format like JSON is probably preferred. However, for use in production, a more optimized yet restricted format can be better. In this latter case, it will be more difficult to achieve interoperability between qlog implementations of various protocol stacks, as some custom or tweaked events from one might not be compatible with the format of the other. This will also reflect in tooling: not all tools will support all formats.

This being said, the authors prefer JSON as the basis for storing qlog, as it retains full flexibility and maximum interoperability. Storage overhead can be managed well in practice by employing compression. For this reason, this document details how to practically transform qlog schema definitions to [\[JSON\]](#), its subset [\[I-JSON\]](#), and its streamable derivative [\[JSON-Text-Sequences\]](#)s. We discuss concrete options to bring down JSON size and processing overheads in [Section 6.3](#).

As depending on the employed format different deserializers/parsers should be used, the "qlog\_format" field is used to indicate the chosen serialization approach. This field is always a string, but can be made hierarchical by the use of the "." separator between entries. For example, a value of "JSON.optimizationA" can indicate that a default JSON format is being used, but that a certain optimization of type A was applied to the file as well (see also [Section 6.3](#)).

### 6.1. qlog to JSON mapping

When mapping qlog to normal JSON, the "qlog\_format" field **MUST** have the value "JSON". This is also the default qlog serialization and default value of this field.

When using normal JSON serialization, the file extension/suffix **SHOULD** be ".qlog" and the Media Type (if any) **SHOULD** be "application/qlog+json" per [\[RFC6839\]](#).

JSON files by definition ([\[RFC8259\]](#)) **MUST** utilize the UTF-8 encoding, both for the file itself and the string values.

While not specifically required by the JSON specification, all qlog field names in a JSON serialization **MUST** be lowercase.

In order to serialize CDDL-based qlog event and data structure definitions to JSON, the official CDDL-to-JSON mapping defined in Appendix E of [\[CDDL\]](#) **SHOULD** be employed.

#### 6.1.1. I-JSON

For some use cases, it should be taken into account that not all popular JSON parsers support the full JSON format. Especially for parsers integrated with the JavaScript programming language (e.g., Web browsers, NodeJS), users are recommended to stick to a JSON subset dubbed [\[I-JSON\]](#) (or Internet-JSON).

One of the key limitations of JavaScript and thus I-JSON is that it cannot represent full 64-bit integers in standard operating mode (i.e., without using BigInt extensions), instead being limited to the range of  $[-(2^{53})+1, (2^{53})-1]$ . In these circumstances, Appendix E of [\[CDDL\]](#) recommends defining new CDDL types for int64 and uint64 that limit their values to this range.

While this can be sensible and workable for most use cases, some protocols targeting qlog serialization (e.g., QUIC, HTTP/3), might require full uint64 variables in some (rare) circumstances. In these situations, it should be allowed to also use the string-based representation of uint64 values alongside the numerical representation. Concretely, the following definition of uint64 should override the original and (web-based) tools should take into account that a uint64 field can be either a number or string.

```
uint64 = text / uint .size 8
```

Figure 36: Custom uint64 definition for I-JSON

#### 6.1.2. Truncated values

For some use cases (e.g., limiting file size, privacy), it can be necessary not to log a full raw blob (using the hexstring type) but instead a truncated value (for example, only the first 100 bytes of an HTTP response body to be able to discern which file it actually

contained). In these cases, the original byte-size length cannot be obtained from the serialized value directly.

As such, all qlog schema definitions **SHOULD** include a separate, length-indicating field for all fields of type hexstring they specify, see for example [Section 5.1](#). This not only ensures the original length can always be retrieved, but also allows the omission of any raw value bytes of the field completely (e.g., out of privacy or security considerations).

To reduce overhead however and in the case the full raw value is logged, the extra length-indicating field can be left out. As such, tools **MUST** be able to deal with this situation and derive the length of the field from the raw value if no separate length-indicating field is present. The main possible permutations are shown by example in [Figure 37](#).

```
// both the full raw value and its length are present
// (length is redundant)
{
  "raw_length": 5,
  "raw": "051428abff"
}

// only the raw value is present, indicating it
// represents the fields full value the byte
// length is obtained by calculating raw.length / 2
{
  "raw": "051428abff"
}

// only the length field is present, meaning the
// value was omitted
{
  "raw_length": 5,
}

// both fields are present and the lengths do not match:
// the value was truncated to the first three bytes.
{
  "raw_length": 5,
  "raw": "051428"
}
```

Figure 37: Example for serializing truncated hexstrings



## 6.2. qlog to JSON Text Sequences mapping

One of the downsides of using pure JSON is that it is inherently a non-streamable format. Put differently, it is not possible to simply append new qlog events to a log file without "closing" this file at the end by appending "]]]". Without these closing tags, most JSON parsers will be unable to parse the file entirely. As most platforms do not provide a standard streaming JSON parser (which would be able to deal with this problem), this document also provides a qlog mapping to a streamable JSON format called JSON Text Sequences (JSON-SEQ) ([RFC7464](#)).

When mapping qlog to JSON-SEQ, the "qlog\_format" field **MUST** have the value "JSON-SEQ".

When using JSON-SEQ serialization, the file extension/suffix **SHOULD** be ".sqlog" (for "streaming" qlog) and the Media Type (if any) **SHOULD** be "application/qlog+json-seq" per [RFC8091](#).

JSON Text Sequences are very similar to JSON, except that JSON objects are serialized as individual records, each prefixed by an ASCII Record Separator (<RS>, 0x1E), and each ending with an ASCII Line Feed character (\n, 0x0A). Note that each record can also contain any amount of newlines in its body, as long as it ends with a newline character before the next <RS> character.

Each qlog event is serialized and interpreted as an individual JSON Text Sequence record, and can simply be appended as a new object at the back of an event stream or log file. Put differently, unlike default JSON, it does not require a file to be wrapped as a full object with "{ ... }" or "[... ]".

For this to work, some qlog definitions have to be adjusted however. Mainly, events are no longer part of the "events" array in the Trace object, but are instead logged separately from the qlog "header", as indicated by the TraceSeq object in [Figure 38](#). Additionally, qlog's JSON-SEQ mapping does not allow logging multiple individual traces in a single qlog file. As such, the QlogFile:traces field is replaced by the singular QlogFileSeq:trace field, see [Figure 39](#). An example can be seen in [Figure 40](#). Note that the "group\_id" field can still be used on a per-event basis to include events from conceptually different sources in a single JSON-SEQ qlog file.

Definition:

```
TraceSeq = {  
  ? title: text  
  ? description: text  
  ? configuration: Configuration  
  ? common_fields: CommonFields  
  ? vantage_point: VantagePoint  
}
```

Figure 38: TraceSeq definition

Definition:

```
QlogFileSeq = {  
  qlog_format: "JSON-SEQ"  
  
  qlog_version: text  
  ? title: text  
  ? description: text  
  ? summary: Summary  
  trace: TraceSeq  
}
```

Figure 39: QlogFileSeq definition

JSON-SEQ serialization examples:

```

// list of qlog events, serialized in accordance with RFC 7464,
// starting with a Record Separator character and ending with a
// newline.
// For display purposes, Record Separators are rendered as <RS>

<RS>{
  "qlog_version": "0.3",
  "qlog_format": "JSON-SEQ",
  "title": "Name of JSON Text Sequence qlog file (short)",
  "description": "Description for this trace file (long)",
  "summary": {
    ...
  },
  "trace": {
    "common_fields": {
      "protocol_type": ["QUIC", "HTTP3"],
      "group_id": "127ecc830d98f9d54a42c4f0842aa87e181a",
      "time_format": "relative",
      "reference_time": 1553986553572
    },
    "vantage_point": {
      "name": "backend-67",
      "type": "server"
    }
  }
}
<RS>{"time": 2, "name": "transport:parameters_set", "data": { ... } }
<RS>{"time": 7, "name": "transport:packet_sent", "data": { ... } }
...

```

Figure 40: Top-level element

Note: while not specifically required by the JSON-SEQ specification, all qlog field names in a JSON-SEQ serialization **MUST** be lowercase.

In order to serialize all other CDDL-based qlog event and data structure definitions to JSON-SEQ, the official CDDL-to-JSON mapping defined in Appendix E of [[CDDL](#)] **SHOULD** still be employed.

### 6.2.1. Supporting JSON Text Sequences in tooling

Note that JSON Text Sequences are not supported in most default programming environments (unlike normal JSON). However, several custom JSON-SEQ parsing libraries exist in most programming languages that can be used and the format is easy enough to parse with existing implementations (i.e., by splitting the file into its component records and feeding them to a normal JSON parser individually, as each record by itself is a valid JSON object).

### 6.3. Other optimized formatting options

Both the JSON and JSON-SEQ formatting options described above are serviceable in general small to medium scale (debugging) setups. However, these approaches tend to be relatively verbose, leading to larger file sizes. Additionally, generalized JSON(-SEQ) (de)serialization performance is typically (slightly) lower than that of more optimized and predictable formats. Both aspects make these formats more challenging ([though still practical](#)) to use in large scale setups.

During the development of qlog, we compared a multitude of alternative formatting and optimization options. The results of this study are [summarized on the qlog github repository](#). The rest of this section discusses some of these approaches implementations could choose and the expected gains and tradeoffs inherent therein. Tools **SHOULD** support mainly the compression options listed in [Section 6.3.2](#), as they provide the largest wins for the least cost overall.

Over time, specific qlog formats and encodings can be created that more formally define and combine some of the discussed optimizations or add new ones. We choose to define these schemes in separate documents to keep the main qlog definition clean and generalizable, as not all contexts require the same performance or flexibility as others and qlog is intended to be a broadly usable and extensible format (for example more flexibility is needed in earlier stages of protocol development, while more performance is typically needed in later stages). This is also the main reason why the general qlog format is the less optimized JSON instead of a more performant option.

To be able to easily distinguish between these options in qlog compatible tooling (without the need to have the user provide out-of-band information or to (heuristically) parse and process files in a multitude of ways, see also [Section 8](#)), we recommend using explicit file extensions to indicate specific formats. As there are no standards in place for this type of extension to format mapping, we employ a commonly used scheme here. Our approach is to list the applied optimizations in the extension in ascending order of application (e.g., if a qlog file is first optimized with technique A and then compressed with technique B, the resulting file would have the extension ".(s)qlog.A.B"). This allows tooling to start at the back of the extension to "undo" applied optimizations to finally arrive at the expected qlog representation.

### 6.3.1. Data structure optimizations

The first general category of optimizations is to alter the representation of data within an JSON(-SEQ) qlog file to reduce file size.

The first option is to employ a scheme similar to the CSV (comma separated value [\[RFC4180\]](#)) format, which utilizes the concept of column "headers" to prevent repeating field names for each datapoint instance. Concretely for JSON qlog, several field names are repeated with each event (i.e., time, name, data). These names could be extracted into a separate list, after which qlog events could be serialized as an array of values, as opposed to a full object. This approach was a key part of the original qlog format (prior to draft-02) using the "event\_fields" field. However, tests showed that this optimization only provided a mean file size reduction of 5% (100MB to 95MB) while significantly increasing the implementation complexity, and this approach was abandoned in favor of the default JSON setup. Implementations using this format should not employ a separate file extension (as it still uses JSON), but rather employ a new value of "JSON.namedheaders" (or "JSON-SEQ.namedheaders") for the "qlog\_format" field (see [Section 3](#)).

The second option is to replace field values and/or names with indices into a (dynamic) lookup table. This is a common compression technique and can provide significant file size reductions (up to 50% in our tests, 100MB to 50MB). However, this approach is even more difficult to implement efficiently and requires either including the (dynamic) table in the resulting file (an approach taken by for example [Chromium's NetLog format](#)) or defining a (static) table up-front and sharing this between implementations. Implementations using this approach should not employ a separate file extension (as it still uses JSON), but rather employ a new value of "JSON.dictionary" (or "JSON-SEQ.dictionary") for the "qlog\_format" field (see [Section 3](#)).

As both options either proved difficult to implement, reduced qlog file readability, and provided too little improvement compared to other more straightforward options (for example [Section 6.3.2](#)), these schemes are not inherently part of qlog.

### 6.3.2. Compression

The second general category of optimizations is to utilize a (generic) compression scheme for textual data. As qlog in the JSON(-SEQ) format typically contains a large amount of repetition, off-the-shelf (text) compression techniques typically succeed very well in bringing down file sizes (regularly with up to two orders of magnitude in our tests, even for "fast" compression levels). As

such, utilizing compression is recommended before attempting other optimization options, even though this might (somewhat) increase processing costs due to the additional compression step.

The first option is to use GZIP compression ([\[RFC1952\]](#)). This generic compression scheme provides multiple compression levels (providing a trade-off between compression speed and size reduction). Utilized at level 6 (a medium setting thought to be applicable for streaming compression of a qlog stream in commodity devices), gzip compresses qlog JSON files to 7% of their initial size on average (100MB to 7MB). For this option, the file extension `.(s)qlog.gz` **SHOULD** BE used. The "qlog\_format" field should still reflect the original JSON formatting of the qlog data (e.g., "JSON" or "JSON-SEQ").

The second option is to use Brotli compression ([\[RFC7932\]](#)). While similar to gzip, this more recent compression scheme provides a better efficiency. It also allows multiple compression levels. Utilized at level 4 (a medium setting thought to be applicable for streaming compression of a qlog stream in commodity devices), brotli compresses qlog JSON files to 7% of their initial size on average (100MB to 7MB). For this option, the file extension `.(s)qlog.br` **SHOULD** BE used. The "qlog\_format" field should still reflect the original JSON formatting of the qlog data (e.g., "JSON" or "JSON-SEQ").

Other compression algorithms of course exist (for example xz, zstd, and lz4). We mainly recommend gzip and brotli because of their tweakable behaviour and wide support in web-based environments, which we envision as the main tooling ecosystem (see also [Section 8](#)).

### 6.3.3. Binary formats

The third general category of optimizations is to use a more optimized (often binary) format instead of the textual JSON format. This approach inherently produces smaller files and often has better (de)serialization performance. However, the resultant files are no longer human readable and some formats require hard tradeoffs between flexibility for performance.

The first option is to use the CBOR (Concise Binary Object Representation [\[RFC7049\]](#)) format. For our purposes, CBOR can be viewed as a straightforward binary variant of JSON. As such, existing JSON qlog files can be trivially converted to and from CBOR (though slightly more work is needed for JSON-SEQ qlogs to convert them to CBOR-SEQ, see [\[RFC8742\]](#)). While CBOR thus does retain the full qlog flexibility, it only provides a 25% file size reduction (100MB to 75MB) compared to textual JSON(-SEQ). As CBOR support in

programming environments is not as widespread as that of textual JSON and the format lacks human readability, CBOR was not chosen as the default qlog format. For this option, the file extension `.(s)qlog.cbor` **SHOULD** BE used. The "qlog\_format" field should still reflect the original JSON formatting of the qlog data (e.g., "JSON" or "JSON-SEQ"). The media type should indicate both whether JSON or JSON Text Sequences are used, as well as whether CBOR or CBOR Sequences are used (see the table below).

A second option is to use a more specialized binary format, such as [Protocol Buffers](#) (protobuf). This format is battle-tested, has support for optional fields and has libraries in most programming languages. Still, it is significantly less flexible than textual JSON or CBOR, as it relies on a separate, pre-defined schema (a `.proto` file). As such, it is not possible to (easily) log new event types in protobuf files without adjusting this schema as well, which has its own practical challenges. As qlog is intended to be a flexible, general purpose format, this type of format was not chosen as its basic serialization. The lower flexibility does lead to significantly reduced file sizes. Our straightforward mapping of the qlog main schema and QUIC/HTTP3 event types to protobuf created qlog files 24% as large as the raw JSON equivalents (100MB to 24MB). For this option, the file extension `.(s)qlog.protobuf` **SHOULD** BE used. The "qlog\_format" field should reflect the different internal format, for example: "qlog\_format": "protobuf".

Note that binary formats can (and should) also be used in conjunction with compression (see [Section 6.3.2](#)). For example, CBOR compresses well (to about 6% of the original textual JSON size (100MB to 6MB) for both gzip and brotli) and so does protobuf (5% (gzip) to 3% (brotli)). However, these gains are similar to the ones achieved by simply compressing the textual JSON equivalents directly (7%, see [Section 6.3.2](#)). As such, since compression is still needed to achieve optimal file size reductions even with binary formats, we feel the more flexible compressed textual JSON options are a better default for the qlog format in general.

#### 6.3.4. Overview and summary

In summary, textual JSON was chosen as the main qlog format due to its high flexibility and because its inefficiencies can be largely solved by the utilization of compression techniques (which are needed to achieve optimal results with other formats as well).

Still, qlog implementers are free to define other qlog formats depending on their needs and context of use. These formats should be described in their own documents, the discussion in this document mainly acting as inspiration and high-level guidance. Implementers

are encouraged to add concrete qlog formats and definitions to [the designated public repository](#).

The following table provides an overview of all the discussed qlog formatting options with examples:

format	qlog_format	extension	media type
JSON <a href="#">Section 6.1</a>	JSON	.qlog	application/qlog+json
JSON Text Sequences <a href="#">Section 6.2</a>	JSON-SEQ	.sqlog	application/qlog+json-seq
named headers <a href="#">Section 6.3.1</a>	JSON(-SEQ).namedheaders	.(s)qlog	application/qlog+json(-seq)
dictionary <a href="#">Section 6.3.1</a>	JSON(-SEQ).dictionary	.(s)qlog	application/qlog+json(-seq)
CBOR <a href="#">Section 6.3.3</a>	JSON(-SEQ)	.(s)qlog.cbor	application/qlog+json(-seq)+cbor(-seq)
protobuf <a href="#">Section 6.3.3</a>	protobuf	.qlog.protobuf	NOT SPECIFIED BY IANA
gzip <a href="#">Section 6.3.2</a>	no change	.gz suffix	application/gzip
brrotli <a href="#">Section 6.3.2</a>	no change	.br suffix	NOT SPECIFIED BY IANA

Table 1

#### 6.4. Conversion between formats

As discussed in the previous sections, a qlog file can be serialized in a multitude of formats, each of which can conceivably be transformed into or from one another without loss of information. For example, a number of JSON-SEQ streamed qlogs could be combined into a JSON formatted qlog for later processing. Similarly, a captured binary qlog could be transformed to JSON for easier interpretation and sharing.

Secondly, we can also consider other structured logging approaches that contain similar (though typically not identical) data to qlog, like raw packet capture files (for example .pcap files from tcpdump) or endpoint-specific logging formats (for example the NetLog format in Google Chrome). These are sometimes the only options, if an implementation cannot or will not support direct qlog output for any



reason, but does provide other internal or external (e.g., SSLKEYLOGFILE export to allow decryption of packet captures) logging options. For this second category, a (partial) transformation from/to qlog can also be defined.

As such, when defining a new qlog serialization format or wanting to utilize qlog-compatible tools with existing codebases lacking qlog support, it is recommended to define and provide a concrete mapping from one format to default JSON-serialized qlog. Several of such mappings exist. Firstly, [pcap2qlog](https://github.com/quiclog/pcap2qlog) transforms QUIC and HTTP/3 packet capture files to qlog. Secondly, [netlog2qlog](#) converts chromium's internal dictionary-encoded JSON format to qlog. Finally, [quictrace2qlog](#) converts the older quictrace format to JSON qlog. Tools can then easily integrate with these converters (either by incorporating them directly or for example using them as a (web-based) API) so users can provide different file types with ease. For example, the [qvis](#) toolsuite supports a multitude of formats and qlog serializations.

## 7. Methods of access and generation

Different implementations will have different ways of generating and storing qlogs. However, there is still value in defining a few default ways in which to steer this generation and access of the results.

### 7.1. Set file output destination via an environment variable

To provide users control over where and how qlog files are created, we define two environment variables. The first, QLOGFILE, indicates a full path to where an individual qlog file should be stored. This path **MUST** include the full file extension. The second, QLOGDIR, sets a general directory path in which qlog files should be placed. This path **MUST** include the directory separator character at the end.

In general, QLOGDIR should be preferred over QLOGFILE if an endpoint is prone to generate multiple qlog files. This can for example be the case for a QUIC server implementation that logs each QUIC connection in a separate qlog file. An alternative that uses QLOGFILE would be a QUIC server that logs all connections in a single file and uses the "group\_id" field ([Section 3.4.6](#)) to allow post-hoc separation of events.

Implementations **SHOULD** provide support for QLOGDIR and **MAY** provide support for QLOGFILE.

When using QLOGDIR, it is up to the implementation to choose an appropriate naming scheme for the qlog files themselves. The chosen scheme will typically depend on the context or protocols used. For example, for QUIC, it is recommended to use the Original Destination

Connection ID (ODCID), followed by the vantage point type of the logging endpoint. Examples of all options for QUIC are shown in [Figure 41](#).

Command: `QLOGFILE=/srv/qlogs/client.qlog quicclientbinary`

Should result in the the `quicclientbinary` executable logging a single qlog file named `client.qlog` in the `/srv/qlogs` directory. This is for example useful in tests when the client sets up just a single connection and then exits.

Command: `QLOGDIR=/srv/qlogs/ quicserverbinary`

Should result in the the `quicserverbinary` executable generating several logs files, one for each QUIC connection.

Given two QUIC connections, with ODCID values "abcde" and "12345" respectively, this would result in two files:

`/srv/qlogs/abcde_server.qlog`

`/srv/qlogs/12345_server.qlog`

Command: `QLOGFILE=/srv/qlogs/server.qlog quicserverbinary`

Should result in the the `quicserverbinary` executable logging a single qlog file named `server.qlog` in the `/srv/qlogs` directory. Given that the server handled two QUIC connections before it was shut down, with ODCID values "abcde" and "12345" respectively, this would result in event instances in the qlog file being tagged with the "group\_id" field with values "abcde" and "12345".

Figure 41: Environment variable examples for a QUIC implementation

## 7.2. Access logs via a well-known endpoint

After generation, qlog implementers **MAY** make available generated logs and traces on an endpoint (typically the server) via the following .well-known URI:

`.well-known/qlog/IDENTIFIER.extension`

The IDENTIFIER variable depends on the context and the protocol. For example for QUIC, the lowercase Original Destination Connection ID (ODCID) is recommended, as it can uniquely identify a connection. Additionally, the extension depends on the chosen format (see [Section 6.3.4](#)). For example, for a QUIC connection with ODCID "abcde", the endpoint for fetching its default JSON-formatted .qlog file would be:

`.well-known/qlog/abcde.qlog`

Implementers **SHOULD** allow users to fetch logs for a given connection on a 2nd, separate connection. This helps prevent pollution of the logs by fetching them over the same connection that one wishes to observe through the log. Ideally, for the QUIC use case, the logs should also be approachable via an HTTP/2 or HTTP/1.1 endpoint (i.e., on TCP port 443), to for example aid debugging in the case where QUIC/UDP is blocked on the network.

qlog implementers **SHOULD NOT** enable this .well-known endpoint in typical production settings to prevent (malicious) users from downloading logs from other connections. Implementers are advised to disable this endpoint by default and require specific actions from the end users to enable it (and potentially qlog itself). Implementers **MUST** also take into account the general privacy and security guidelines discussed in [Section 9](#) before exposing qlogs to outside actors.

## 8. Tooling requirements

Tools ingestion qlog **MUST** indicate which qlog version(s), qlog format(s), compression methods and potentially other input file formats (for example .pcap) they support. Tools **SHOULD** at least support .qlog files in the default JSON format ([Section 6.1](#)). Additionally, they **SHOULD** indicate exactly which values for and properties of the name (category and type) and data fields they look for to execute their logic. Tools **SHOULD** perform a (high-level) check if an input qlog file adheres to the expected qlog schema. If a tool determines a qlog file does not contain enough supported information to correctly execute the tool's logic, it **SHOULD** generate a clear error message to this effect.

Tools **MUST NOT** produce breaking errors for any field names and/or values in the qlog format that they do not recognize. Tools **SHOULD** indicate even unknown event occurrences within their context (e.g., marking unknown events on a timeline for manual interpretation by the user).

Tool authors should be aware that, depending on the logging implementation, some events will not always be present in all traces. For example, using a circular logging buffer of a fixed size, it could be that the earliest events (e.g., connection setup events) are later overwritten by "newer" events. Alternatively, some events can be intentionally omitted out of privacy or file size considerations. Tool authors are encouraged to make their tools robust enough to still provide adequate output for incomplete logs.

## 9. Security and privacy considerations

Protocols such as TLS [[RFC8446](#)] and QUIC [[RFC9000](#)] provide varying degrees of secure protection for the wire image [[RFC8546](#)]. There is inevitably tension between security and observability, when logging can reveal aspects of the wire image, that would ordinarily be protected. This tension equally applies to any privacy considerations that build on security properties, especially if data can be correlated across data sources.

qlog operators and implementers should be mindful of the security and privacy risks inherent in handling qlog data. This includes but is not limited to logging, storing, or using the data. Data might be considered as non-sensitive, potentially-sensitive, or sensitive; applying the considerations in this section may produce different risks depending on the nature of the data itself, or its handling. However, in many cases the largest risk factors arise from data that can be considered as potentially-sensitive or sensitive.

The following is a non-exhaustive list of such fields and types of data that can be carried in qlog data:

- \*IP addresses and transport protocol port numbers, which can be used to uniquely identify individual connections, endpoints, and potentially users.
- \*Session, Connection, or User identifiers which can be used to correlate nominally separate contexts. For example, QUIC Connection IDs can be used to identify and track users across geographical networks ([Section 9.5](#) of [[RFC9000](#)]).
- \*Stored State which can be used to correlate individual connections or sessions over time. Examples include QUIC address validation and retry tokens, TLS session tickets, and HTTP cookies.
- \*Decryption keys, passwords, and tokens which can be used with other data sources (e.g., captures of encrypted packets) to correlate qlog data to a specific connection or user or leak additional information. Examples include TLS decryption keys and HTTP-level API access or authorization tokens.
- \*Data that can be used to correlate qlogs to other data sources (e.g., captures of encrypted packets). Examples include high-resolution event timestamps or inter-event timings, event counts, packet and frame sizes.
- \*Full or partial encrypted raw packet and frame payloads, which can be used with other data sources (e.g., captures of encrypted

packets) to correlate qlog data to a specific connection or session.

\*Full or partial plaintext raw packet and frame payloads (e.g., HTTP Field values, HTTP response data, TLS SNI field values), which can contain directly sensitive information.

The simplest and most extreme form of protection against abuse of this information is the complete deletion of a given field, which is equivalent to not logging the field(s) in question. While deletion completely protects the data in the deleted fields from the risk of compromise, it also reduces the utility of the dataset as a whole. As such, a balance should be found between logging these fields and the potential risks inherent in their (involuntary) disclosure. This balance depends on the use case at hand (e.g., research datasets might have different requirements to live operational troubleshooting). Capturing the minimal amount of data required for a specific purpose can help to minimize the risks associated with data usage. qlog implementations that provide fine-grained control over the inclusion of data fields, ideally on a per-use-case or per-connection basis, improve the ability to minimize data.

Any data that is determined to be necessary for a use case at hand could be logged or captured. As per [\[RFC6973\]](#), operators must be aware that such data will be at risk of compromise. As such, measures should be taken to firstly reduce the risk of compromise and secondly reduce the risk of abuse of compromised data. While a full discussion of both aspects is out of scope for this document, the following paragraphs discuss high-level considerations that can be applied to qlog data.

To reduce the risk of compromise, operators can take measures such as: limiting the length of time that data is stored, encrypting data in transit and at rest, limiting access rights to the data, and auditing data usage practices. qlog deployments that provide integrated options for automated or manual data deletion and (aggressive) aggregation, improve the ability to minimize the risk of compromise.

To reduce the risk of data abuse after compromise, data can be anonymized, pseudonymized, otherwise permuted/replaced, truncated, (re-)encrypted, or aggregated. A partial discussion of applicable techniques (especially for IP address information) can be found in [Appendix B](#) of [\[DNS-PRIVACY\]](#). Operators should, however, be aware that many of these techniques have been shown to be insufficient to safeguard user privacy and/or to protect user identity, especially if a qlog data set is large or easily correlated against other data sources.

Finally, qlog operators should consider the interplay between their use case needs and end user rights or preferences. While active user participation (as indicated by [RFC6973]) on a per-qlog basis is difficult, as logs are often captured out-of-band to the main user interaction and intent, general user expectations should be taken into account. qlog deployments that provide mechanisms to integrate the capture, storage and removal of qlogs with more general, often pre-existing, user preference and privacy control systems, improve the ability to protect data sensitive or confidential to the end user. In qlog, these data are typically (but not exclusively) contained in fields of the RawInfo type (see [Section 5.1](#)). qlog users should thus be particularly hesitant to include these fields for all but the most stringent use cases.

## 10. IANA Considerations

TODO: primarily the .well-known URI

## 11. References

### 11.1. Normative References

- [CDDL] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/rfc/rfc8610>>.
- [DNS-PRIVACY] Dickinson, S., Overeinder, B., van Rijswijk-Deij, R., and A. Mankin, "Recommendations for DNS Privacy Service Operators", BCP 232, RFC 8932, DOI 10.17487/RFC8932, October 2020, <<https://www.rfc-editor.org/rfc/rfc8932>>.
- [I-JSON] Bray, T., Ed., "The I-JSON Message Format", RFC 7493, DOI 10.17487/RFC7493, March 2015, <<https://www.rfc-editor.org/rfc/rfc7493>>.
- [JSON] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/rfc/rfc8259>>.
- [JSON-Text-Sequences]  
Williams, N., "JavaScript Object Notation (JSON) Text Sequences", RFC 7464, DOI 10.17487/RFC7464, February 2015, <<https://www.rfc-editor.org/rfc/rfc7464>>.
- [RFC1952] Deutsch, P., "GZIP file format specification version 4.3", RFC 1952, DOI 10.17487/RFC1952, May 1996, <<https://www.rfc-editor.org/rfc/rfc1952>>.

**[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/rfc/rfc2119>>.

**[RFC4180]**

Shafranovich, Y., "Common Format and MIME Type for Comma-Separated Values (CSV) Files", RFC 4180, DOI 10.17487/RFC4180, October 2005, <<https://www.rfc-editor.org/rfc/rfc4180>>.

**[RFC6839]**

Hansen, T. and A. Melnikov, "Additional Media Type Structured Syntax Suffixes", RFC 6839, DOI 10.17487/RFC6839, January 2013, <<https://www.rfc-editor.org/rfc/rfc6839>>.

**[RFC6973]**

Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/rfc/rfc6973>>.

**[RFC7049]**

Bormann, C. and P. Hoffman, "Concise Binary Object Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049, October 2013, <<https://www.rfc-editor.org/rfc/rfc7049>>.

**[RFC7464]**

Williams, N., "JavaScript Object Notation (JSON) Text Sequences", RFC 7464, DOI 10.17487/RFC7464, February 2015, <<https://www.rfc-editor.org/rfc/rfc7464>>.

**[RFC7932]**

Alakuijala, J. and Z. Szabadka, "Brotli Compressed Data Format", RFC 7932, DOI 10.17487/RFC7932, July 2016, <<https://www.rfc-editor.org/rfc/rfc7932>>.

**[RFC8091]**

Wilde, E., "A Media Type Structured Syntax Suffix for JSON Text Sequences", RFC 8091, DOI 10.17487/RFC8091, February 2017, <<https://www.rfc-editor.org/rfc/rfc8091>>.

**[RFC8174]**

Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

**[RFC8259]**

Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/rfc/rfc8259>>.

**[RFC8446]**

Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.

**[RFC9000]**

Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/rfc/rfc9000>>.

**11.2. Informative References**

**[QLOG-H3]** Marx, R., Niccolini, L., Seemann, M., and L. Pardue, "HTTP/3 and QPACK qlog event definitions", Work in Progress, Internet-Draft, draft-ietf-quic-qlog-h3-events-03, 24 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-qlog-h3-events-03>>.

**[QLOG-QUIC]** Marx, R., Niccolini, L., Seemann, M., and L. Pardue, "QUIC event definitions for qlog", Work in Progress, Internet-Draft, draft-ietf-quic-qlog-quic-events-03, 24 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-qlog-quic-events-03>>.

**[RFC8546]** Trammell, B. and M. Kuehlewind, "The Wire Image of a Network Protocol", RFC 8546, DOI 10.17487/RFC8546, April 2019, <<https://www.rfc-editor.org/rfc/rfc8546>>.

**[RFC8742]** Bormann, C., "Concise Binary Object Representation (CBOR) Sequences", RFC 8742, DOI 10.17487/RFC8742, February 2020, <<https://www.rfc-editor.org/rfc/rfc8742>>.

**Appendix A. Change Log**

**A.1. Since draft-ietf-quic-qlog-main-schema-04:**

\*Updated RawInfo definition and guidance (#243)

**A.2. Since draft-ietf-quic-qlog-main-schema-03:**

\*Added security and privacy considerations discussion (#252)

**A.3. Since draft-ietf-quic-qlog-main-schema-02:**

\*No changes - new draft to prevent expiration

**A.4. Since draft-ietf-quic-qlog-main-schema-01:**

\*Change the data definition language from TypeScript to CDDL (#143)



#### **A.5. Since draft-ietf-quic-qlog-main-schema-00:**

- \*Changed the streaming serialization format from NDJSON to JSON Text Sequences (#172)
- \*Added Media Type definitions for various qlog formats (#158)
- \*Changed to semantic versioning

#### **A.6. Since draft-marx-qlog-main-schema-draft-02:**

- \*These changes were done in preparation of the adoption of the drafts by the QUIC working group (#137)
- \*Moved RawInfo, Importance, Generic events and Simulation events to this document.
- \*Added basic event definition guidelines
- \*Made protocol\_type an array instead of a string (#146)

#### **A.7. Since draft-marx-qlog-main-schema-01:**

- \*Decoupled qlog from the JSON format and described a mapping instead (#89)
  - Data types are now specified in this document and proper definitions for fields were added in this format
  - 64-bit numbers can now be either strings or numbers, with a preference for numbers (#10)
  - binary blobs are now logged as lowercase hex strings (#39, #36)
  - added guidance to add length-specifiers for binary blobs (#102)
- \*Removed "time\_units" from Configuration. All times are now in ms instead (#95)
- \*Removed the "event\_fields" setup for a more straightforward JSON format (#101, #89)
- \*Added a streaming option using the NDJSON format (#109, #2, #106)
- \*Described optional optimization options for implementers (#30)
- \*Added QLOGDIR and QLOGFILE environment variables, clarified the .well-known URL usage (#26, #33, #51)

\*Overall tightened up the text and added more examples

#### **A.8. Since draft-marx-qlog-main-schema-00:**

\*All field names are now lowercase (e.g., category instead of CATEGORY)

\*Triggers are now properties on the "data" field value, instead of separate field types (#23)

\*group\_ids in common\_fields is now just also group\_id

#### **Acknowledgements**

Much of the initial work by Robin Marx was done at the Hasselt and KU Leuven Universities.

Thanks to Jana Iyengar, Brian Trammell, Dmitri Tikhonov, Stephen Petrides, Jari Arkko, Marcus Ihlar, Victor Vasiliev, Mirja Kuehlewind, and Jeremy Laine for their feedback and suggestions.

#### **Authors' Addresses**

Robin Marx (editor)  
Akamai

Email: [rmarx@akamai.com](mailto:rmarx@akamai.com)

Luca Niccolini (editor)  
Meta

Email: [lniccolini@meta.com](mailto:lniccolini@meta.com)

Marten Seemann (editor)  
Protocol Labs

Email: [marten@protocol.ai](mailto:marten@protocol.ai)

Lucas Pardue (editor)  
Cloudflare

Email: [lucaspardue.24.7@gmail.com](mailto:lucaspardue.24.7@gmail.com)