

moq
Internet-Draft
Intended status: Informational
Expires: 11 October 2026

L. Curley
9 April 2026

Media over QUIC - Lite
draft-lcurley-moq-lite-04

Abstract

moq-lite is designed to fanout live content 1->N across the internet. It leverages QUIC to prioritize important content, avoiding head-of-line blocking while respecting encoding dependencies. While primarily designed for media, the transport is payload agnostic and can be proxied by relays/CDNs without knowledge of codecs, containers, or encryption keys.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Media Over QUIC Working Group mailing list (moq@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/moq/>.

Source for this draft and an issue tracker can be found at <https://github.com/kixelated/moq-drafts>.

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 11 October 2026.

Copyright Notice

Copyright (c) 2026 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. Conventions and Definitions	3
2. Rationale	3
3. Concepts	4
3.1. Session	4
3.2. Broadcast	5
3.3. Track	5
3.4. Group	6
3.5. Frame	6
4. Flow	6
4.1. Connection	6
4.2. Termination	7
4.3. Handshake	7
5. Streams	7
5.1. Bidirectional Streams	7
5.1.1. Announce	8
5.1.2. Subscribe	8
5.1.3. Fetch	9
5.1.4. Probe	9
5.1.5. Goaway	10
6. Delivery	10
6.1. Prioritization	10
6.1.1. Priority	10
6.1.2. Ordered	11
6.2. Expiration	11
6.3. Unidirectional Streams	12
6.3.1. Group	12
7. Encoding	13
7.1. Message Length	13
7.2. STREAM_TYPE	13
7.3. ANNOUNCE_INTEREST	13
7.4. ANNOUNCE	14
7.5. SUBSCRIBE	15

7.6.	SUBSCRIBE_UPDATE	16
7.7.	SUBSCRIBE_OK	16
7.8.	SUBSCRIBE_DROP	17
7.9.	FETCH	17
7.10.	PROBE	18
7.11.	GOAWAY	18
7.12.	GROUP	18
7.13.	FRAME	19
8.	Appendix A: Changelog	19
8.1.	moq-lite-04	19
8.2.	moq-lite-03	19
8.3.	moq-lite-02	20
8.4.	moq-lite-01	20
9.	Appendix B: Upstream Differences	20
9.1.	Deleted Messages	21
9.2.	Renamed Messages	22
9.3.	Deleted Fields	22
10.	Security Considerations	23
11.	IANA Considerations	23
12.	Normative References	23
	Acknowledgments	23
	Author's Address	23

1. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Rationale

This draft is based on MoqTransport [moqt]. The concepts, motivations, and terminology are very similar and when in doubt, refer to existing MoqTransport literature. A few things have been renamed (ex. object -> frame) to better align with media terminology.

I absolutely believe in the motivation and potential of Media over QUIC. The layering is phenomenal and addresses many of the problems with current live media protocols. I fully support the goals of the working group and the IETF process.

But it's been difficult to design such an experimental protocol via committee. MoqTransport has become too complicated.

There are too many messages, optional modes, and half-baked features. Too many hypotheses, too many potential use-cases, too many diametrically opposed opinions. This is expected (and even desired) as compromise gives birth to a standard.

But I believe the standardization process is hindering practical experimentation. The ideas behind MoQ can be proven now before being cemented as an RFC. We should spend more time building an actual application and less time arguing about a hypothetical one.

moq-lite is the bare minimum needed for a real-time application aiming to replace WebRTC. Every feature from MoqTransport that is not necessary (or has not been implemented yet) has been removed for simplicity. This includes many great ideas (ex. group order) that may be added as they are needed. This draft is the current state, not the end state.

3. Concepts

moq-lite consists of:

- * ***Session***: An established QUIC connection between a client and server.
- * ***Broadcast***: A collection of Tracks from a single publisher.
- * ***Track***: A series of Groups, each of which can be delivered and decoded out-of-order.
- * ***Group***: A series of Frames, each of which must be delivered and decoded in-order.
- * ***Frame***: A sized payload of bytes within a Group.

The application determines how to split data into broadcast, tracks, groups, and frames. The moq-lite layer provides fanout, prioritization, and caching even for latency sensitive applications.

3.1. Session

A Session consists of a connection between a client and a server. There is currently no P2P support within QUIC so it's out of scope for moq-lite.

The moq-lite version identifier is moq-lite-xx where xx is the two-digit draft version. For bare QUIC, this is negotiated as an ALPN token during the QUIC handshake. For WebTransport over HTTP/3, the QUIC ALPN remains h3 and the moq-lite version is advertised via the WT-Available-Protocols and WT-Protocol CONNECT headers.

The session is active immediately after the QUIC/WebTransport connection is established. Extensions are negotiated via stream probing: an endpoint opens a stream with an unknown type and the peer resets it if unsupported.

While moq-lite is a point-to-point protocol, it's intended to work end-to-end via relays. Each client establishes a session with a CDN edge server, ideally the closest one. Any broadcasts and subscriptions are transparently proxied by the CDN behind the scenes.

3.2. Broadcast

A Broadcast is a collection of Tracks from a single publisher. This corresponds to a MoqTransport's "track namespace".

A publisher may produce multiple broadcasts, each of which is advertised via an ANNOUNCE message. The subscriber uses the ANNOUNCE_INTEREST message to discover available broadcasts. These announcements are live and can change over time, allowing for dynamic origin discovery.

A broadcast consists of any number of Tracks. The contents, relationships, and encoding of tracks are determined by the application.

3.3. Track

A Track is a series of Groups identified by a unique name within a Broadcast.

A track consists of a single active Group at any moment, called the "latest group". When a new Group is started, the previous Group is closed and may be dropped for any reason. The duration before an incomplete group is dropped is determined by the application and the publisher/subscriber's latency target.

Every subscription is scoped to a single Track. A subscription starts at a configurable Group (defaulting to the latest) and continues until a configurable end Group or until either the publisher or subscriber cancels the subscription.

The subscriber and publisher both indicate their delivery preference:

- Priority indicates if Track A should be transmitted instead of Track B.
- Ordered indicates if the Groups within a Track should be transmitted in order.
- Max Latency indicates the maximum duration before a Group is abandoned.

The combination of these preferences enables the most important content to arrive during network degradation while still respecting encoding dependencies.

3.4. Group

A Group is an ordered stream of Frames within a Track.

Each group consists of an append-only list of Frames. A Group is served by a dedicated QUIC stream which is closed on completion, reset by the publisher, or cancelled by the subscriber. This ensures that all Frames within a Group arrive reliably and in order.

In contrast, Groups may arrive out of order due to network congestion and prioritization. The application SHOULD process or buffer groups out of order to avoid blocking on flow control.

3.5. Frame

A Frame is a payload of bytes within a Group.

A frame is used to represent a chunk of data with an upfront size. The contents are opaque to the moq-lite layer.

4. Flow

This section outlines the flow of messages within a moq-lite session. See the Messages section for the specific encoding.

4.1. Connection

moq-lite runs on top of WebTransport. WebTransport is a layer on top of QUIC and HTTP/3, required for web support. The API is nearly identical to QUIC with the exception of stream IDs.

How the WebTransport connection is authenticated is out-of-scope for this draft.

4.2. Termination

QUIC bidirectional streams have an independent send and receive direction. Rather than deal with half-open states, moq-lite combines both sides. If an endpoint closes the send direction of a stream, the peer **MUST** also close their send direction.

moq-lite contains many long-lived transactions, such as subscriptions and announcements. These are terminated when the underlying QUIC stream is terminated.

To terminate a stream, an endpoint may:

- close the send direction (STREAM with FIN) to gracefully terminate (all messages are flushed).
- reset the send direction (RESET_STREAM) to immediately terminate.

After resetting the send direction, an endpoint **MAY** close the recv direction (STOP_SENDING). However, it is ultimately the other peer's responsibility to close their send direction.

4.3. Handshake

See the Section 3.1 section for ALPN negotiation and session activation details.

5. Streams

moq-lite uses a bidirectional stream for each transaction. If the stream is closed, potentially with an error, the transaction is terminated.

5.1. Bidirectional Streams

Bidirectional streams are used for control streams. There's a 1-byte STREAM_TYPE at the beginning of each stream.

ID	Stream	Creator
0x1	Announce	Subscriber
0x2	Subscribe	Subscriber
0x3	Fetch	Subscriber
0x4	Probe	Subscriber
0x5	Goaway	Either

Table 1

5.1.1.1. Announce

A subscriber can open an Announce Stream to discover broadcasts matching a prefix.

The subscriber creates the stream with a `ANNOUNCE_INTEREST` message. The publisher replies with `ANNOUNCE` messages for any matching broadcasts and any future changes. Each `ANNOUNCE` message contains one of the following statuses:

- * `active`: a matching broadcast is available.
- * `ended`: a previously active broadcast is no longer available.

Each broadcast starts as ended and MUST alternate between active and ended. The subscriber MUST reset the stream if it receives a duplicate status, such as two active statuses in a row or an ended without active. When the stream is closed, the subscriber MUST assume that all broadcasts are now ended.

Path prefix matching and equality is done on a byte-by-byte basis. There MAY be multiple Announce Streams, potentially containing overlapping prefixes, that get their own `ANNOUNCE` messages.

5.1.1.2. Subscribe

A subscriber opens Subscribe Streams to request a Track.

The subscriber MUST start a Subscribe Stream with a SUBSCRIBE message followed by any number of SUBSCRIBE_UPDATE messages. The publisher replies with a SUBSCRIBE_OK message followed by any number of SUBSCRIBE_DROP and additional SUBSCRIBE_OK messages. The first message on the response stream MUST be a SUBSCRIBE_OK; it is not valid to send a SUBSCRIBE_DROP before SUBSCRIBE_OK.

The publisher closes the stream (FIN) when every group from start to end has been accounted for, either via a GROUP stream (completed or reset) or a SUBSCRIBE_DROP message. Unbounded subscriptions (no end group) stay open until the publisher closes the stream to indicate the track has ended, or either endpoint resets. Either endpoint MAY reset/cancel the stream at any time.

5.1.3. Fetch

A subscriber opens a Fetch Stream (0x3) to request a single Group from a Track.

The subscriber sends a FETCH message containing the broadcast path, track name, priority, and group sequence. Unlike Group Streams (which MUST start with a GROUP message), the publisher responds with FRAME messages directly on the same bidirectional stream — there is no preceding GROUP header. The Subscribe ID and Group Sequence for the returned FRAME messages are implicit, taken from the original FETCH request. The publisher FINs the stream after the last frame, or resets the stream on error.

Fetch behaves like HTTP: a single request/response per stream.

5.1.4. Probe

A subscriber opens a Probe Stream (0x4) to measure the available bitrate of the connection.

The subscriber sends a PROBE message with a target bitrate on the bidirectional stream. The subscriber MAY send additional PROBE messages on the same stream to update the target bitrate; the publisher MUST treat each PROBE as a new target to attempt. The publisher SHOULD pad the connection to achieve the most recent target bitrate. The publisher periodically replies with PROBE messages on the same bidirectional stream containing the current estimated bitrate and smoothed RTT.

If the publisher does not support PROBE (e.g., congestion controller is not exposed), it MUST reset the stream.

5.1.5. Goaway

Either endpoint can open a Goaway Stream (0x5) to initiate a graceful session shutdown.

The sender sends a GOAWAY message containing an optional new session URI. If the URI is non-empty, the peer SHOULD establish a new session at the provided URI and migrate any active subscriptions. The peer MUST NOT open new streams on the current session after receiving a GOAWAY.

The sender closes the stream (FIN) when it is ready to terminate the session. The peer SHOULD close all streams and the session after migrating or when it no longer needs the session.

6. Delivery

The most important concept in moq-lite is how to deliver a subscription. QUIC can only improve the user experience if data is delivered out-of-order during congestion. This is the sole reason why data is divided into Broadcasts, Tracks, Groups, and Frames.

moq-lite consists of multiple groups being transmitted in parallel across separate streams. How these streams get transmitted over the network is very important, and yet has been distilled down into a few simple properties:

6.1. Prioritization

The Publisher and Subscriber both exchange Priority and Ordered values: - Priority determines which Track should be transmitted next. - Ordered determines which Group within the Track should be transmitted next.

A publisher SHOULD attempt to transmit streams based on these fields. This depends on the QUIC implementation and it may not be possible to get fine-grained control.

6.1.1. Priority

The Subscriber Priority is scoped to the connection. The Publisher Priority SHOULD be used to resolve conflicts or ties.

A conflict can occur when a relay tries to serve multiple downstream subscriptions from a single upstream subscription. Any upstream subscription SHOULD use the publisher priority, not some combination of different subscriber priorities.

Rather than try to explain everything, here's an example:

Example: There are two people in a conference call, Ali and Bob.

We subscribe to both of their audio tracks with priority 2 and video tracks with priority 1. This will cause equal priority for Ali and Bob while prioritizing audio. ali/audio + bob/audio:
subscriber_priority=2 publisher_priority=2 ali/video + bob/video:
subscriber_priority=1 publisher_priority=1

If Bob starts actively speaking, they can bump their publisher priority via a SUBSCRIBE_OK message. This would cause tracks be delivered in this order: bob/audio: subscriber_priority=2
publisher_priority=3 ali/audio: subscriber_priority=2
publisher_priority=2 bob/video: subscriber_priority=1
publisher_priority=2 ali/video: subscriber_priority=1
publisher_priority=1

The subscriber priority takes precedence, so we could override it if we decided to full-screen Ali's window: ali/audio
subscriber_priority=4 publisher_priority=2 ali/video
subscriber_priority=3 publisher_priority=1 bob/audio
subscriber_priority=2 publisher_priority=3 bob/video
subscriber_priority=1 publisher_priority=2

6.1.2. Ordered

The Subscriber Ordered field signals if older (0x1) or newer (0x0) groups should be transmitted first within a Track. The Publisher Ordered field MAY likewise be used to resolve conflicts.

An application SHOULD use ordered when it wants to provide a VOD-like experience, preferring to buffer old groups rather than skip them. An application SHOULD NOT use ordered when it wants to provide a live experience, preferring to skip old groups rather than buffer them.

Note that Section 6.2 is not affected by ordered. An old group may still be cancelled/skipped if it exceeds max_latency set by either peer. An application MUST support gaps and out-of-order delivery even when ordered is true.

6.2. Expiration

The Publisher and Subscriber both transmit a Max Latency value, indicating the maximum duration before a group is expired.

It is not crucial to aggressively expire groups thanks to Section 6.1. However, a lower priority group will still consume RAM, bandwidth, and potentially flow control. It is RECOMMENDED that an application set conservative limits and only resort to expiration when data is absolutely no longer needed.

A subscriber SHOULD expire groups based on the Subscriber Max Latency in SUBSCRIBE/SUBSCRIBE_UPDATE. A publisher SHOULD expire groups based on the Publisher Max Latency in SUBSCRIBE_OK. An implementation MAY use the minimum of both when determining when to expire a group.

Group age is computed relative to the latest group by sequence number. A group is never expired until at least the next group (by sequence number) has been received or queued. Once a newer group exists, a group is considered expired if the time between its arrival and the latest group's arrival exceeds Max Latency. The arrival time is when the first byte of a group is received (subscriber) or queued (publisher). An expired group SHOULD be reset at the QUIC level to avoid consuming flow control.

6.3. Unidirectional Streams

Unidirectional streams are used for data transmission.

ID	Stream	Creator
0x0	Group	Publisher

Table 2

6.3.1. Group

A publisher creates Group Streams in response to a Subscribe Stream.

A Group Stream MUST start with a GROUP message and MAY be followed by any number of FRAME messages. A Group MAY contain zero FRAME messages, potentially indicating a gap in the track. A frame MAY contain an empty payload, potentially indicating a gap in the group.

Both the publisher and subscriber MAY reset the stream at any time. This is not a fatal error and the session remains active. The subscriber MAY cache the error and potentially retry later.

7. Encoding

This section covers the encoding of each message.

7.1. Message Length

Most messages are prefixed with a variable-length integer indicating the number of bytes in the message payload that follows. This length field does not include the length of the varint length itself.

An implementation SHOULD close the connection with a `PROTOCOL_VIOLATION` if it receives a message with an unexpected length. The version and extensions should be used to support new fields, not the message length.

7.2. `STREAM_TYPE`

All streams start with a short header indicating the stream type.

```
STREAM_TYPE {  
    Stream Type (i)  
}
```

The stream ID depends on if it's a bidirectional or unidirectional stream, as indicated in the Streams section. A receiver MUST reset the stream if it receives an unknown stream type. Unknown stream types MUST NOT be treated as fatal; this enables extension negotiation via stream probing.

7.3. `ANNOUNCE_INTEREST`

A subscriber sends an `ANNOUNCE_INTEREST` message to indicate it wants to receive an `ANNOUNCE` message for any broadcasts with a path that starts with the requested prefix.

```
ANNOUNCE_INTEREST Message {  
    Message Length (i)  
    Broadcast Path Prefix (s),  
    Exclude Hop (i),  
}
```

***Broadcast Path Prefix*:** Indicate interest for any broadcasts with a path that starts with this prefix.

***Exclude Hop*:** If non-zero, the publisher SHOULD skip `ANNOUNCE` messages for broadcasts whose Hop ID entries contain this value. This is used by relays to avoid routing loops in a cluster.

The publisher MUST respond with ANNOUNCE messages for any matching and active broadcasts, followed by ANNOUNCE messages for any future updates. Implementations SHOULD consider reasonable limits on the number of matching broadcasts to prevent resource exhaustion.

7.4. ANNOUNCE

A publisher sends an ANNOUNCE message to advertise a change in broadcast availability. Only the suffix is encoded on the wire, as the full path can be constructed by prepending the requested prefix.

The status is relative to all prior ANNOUNCE messages on the same stream. A publisher MUST ONLY alternate between status values (from active to ended or vice versa).

```
ANNOUNCE Message {  
  Message Length (i)  
  Announce Status (i),  
  Broadcast Path Suffix (s),  
  Hop Count (i),  
  Hop ID (i) ...,  
}
```

Announce Status: A flag indicating the announce status.

* ended (0): A path is no longer available.

* active (1): A path is now available.

Broadcast Path Suffix: This is combined with the broadcast path prefix to form the full broadcast path.

Hop Count: The number of Hop ID entries that follow. A value of 0 means no Hop ID entries are present, indicating a fully unknown path (e.g. the announcement originated locally or was received from a peer that does not support hop tracking). A receiver MUST close the stream with a `PROTOCOL_VIOLATION` if the Hop Count does not match the number of subsequent Hop ID entries.

Hop ID: A unique identifier for each relay in the path from the origin publisher. Each relay MUST append its own Hop ID and increment the Hop Count when forwarding an announcement, even if no prior Hop ID entries are present (e.g. when receiving from an older protocol version). The Hop Count is used as a tiebreaker when there are multiple paths to the same broadcast. A Hop ID value of 0 indicates an unknown or bridged relay hop (e.g. when bridging from an older protocol version that does not assign Hop IDs); the Hop Count still reflects the total number of entries including unknown hops.

7.5. SUBSCRIBE

SUBSCRIBE is sent by a subscriber to start a subscription.

```
SUBSCRIBE Message {  
  Message Length (i)  
  Subscribe ID (i)  
  Broadcast Path (s)  
  Track Name (s)  
  Subscriber Priority (8)  
  Subscriber Ordered (8)  
  Subscriber Max Latency (i)  
  Start Group (i)  
  End Group (i)  
}
```

***Subscribe ID*:** A unique identifier chosen by the subscriber. A Subscribe ID MUST NOT be reused within the same session, even if the prior subscription has been closed.

***Subscriber Priority*:** The priority of the subscription within the session, represented as a u8. The publisher SHOULD transmit *_higher_* values first during congestion. See the Section 6.1 section for more information.

***Subscriber Ordered*:** A single byte representing whether groups are transmitted in ascending (0x1) or descending (0x0) order. The publisher SHOULD transmit *_older_* groups first during congestion if true. See the Section 6.1 section for more information.

***Subscriber Max Latency*:** This value is encoded in milliseconds and represents the maximum age of a group relative to the latest group. The publisher SHOULD reset old group streams when the difference in arrival time between the group and the latest group exceeds this duration. See the Section 6.2 section for more information.

***Start Group*:** The first group to deliver. A value of 0 means the latest group (default). A non-zero value is the absolute group sequence + 1.

***End Group*:** The last group to deliver (inclusive). A value of 0 means unbounded (default). A non-zero value is the absolute group sequence + 1.

7.6. SUBSCRIBE_UPDATE

A subscriber can modify a subscription with a SUBSCRIBE_UPDATE message. A subscriber MAY send multiple SUBSCRIBE_UPDATE messages to update the subscription. The start and end group can be changed in either direction (growing or shrinking).

```
SUBSCRIBE_UPDATE Message {  
  Message Length (i)  
  Subscriber Priority (8)  
  Subscriber Ordered (8)  
  Subscriber Max Latency (i)  
  Start Group (i)  
  End Group (i)  
}
```

See Section 5.1.2 for information about each field.

7.7. SUBSCRIBE_OK

A SUBSCRIBE_OK message is sent in response to a SUBSCRIBE. The publisher MAY send multiple SUBSCRIBE_OK messages to update the subscription. The first message on the response stream MUST be a SUBSCRIBE_OK; a SUBSCRIBE_DROP MUST NOT precede it.

```
SUBSCRIBE_OK Message {  
  Type (i) = 0x0  
  Message Length (i)  
  Publisher Priority (8)  
  Publisher Ordered (8)  
  Publisher Max Latency (i)  
  Start Group (i)  
  End Group (i)  
}
```

Type: Set to 0x0 to indicate a SUBSCRIBE_OK message.

Start Group: The resolved absolute start group sequence. A value of 0 means the start group is not yet known; the publisher MUST send a subsequent SUBSCRIBE_OK with a resolved value. A non-zero value is the absolute group sequence + 1.

End Group: The resolved absolute end group sequence (inclusive). A value of 0 means unbounded. A non-zero value is the absolute group sequence + 1.

See Section 5.1.2 for information about the other fields.

7.8. SUBSCRIBE_DROP

A SUBSCRIBE_DROP message is sent by the publisher on the Subscribe Stream when groups cannot be served.

```
SUBSCRIBE_DROP Message {  
  Type (i) = 0x1  
  Message Length (i)  
  Start Group (i)  
  End Group (i)  
  Error Code (i)  
}
```

***Type*:** Set to 0x1 to indicate a SUBSCRIBE_DROP message.

***Start Group*:** The first absolute group sequence in the dropped range.

***End Group*:** The last absolute group sequence in the dropped range (inclusive).

***Error Code*:** An application-specific error code. A value of 0 indicates no error; the groups are simply unavailable.

7.9. FETCH

FETCH is sent by a subscriber to request a single group from a track.

```
FETCH Message {  
  Message Length (i)  
  Broadcast Path (s)  
  Track Name (s)  
  Subscriber Priority (8)  
  Group Sequence (i)  
}
```

***Broadcast Path*:** The broadcast path of the track to fetch from.

***Track Name*:** The name of the track to fetch from.

***Subscriber Priority*:** The priority of the fetch within the session, represented as a u8. See the Section 6.1 section for more information.

***Group Sequence*:** The sequence number of the group to fetch.

The publisher responds with FRAME messages on the same stream. The publisher FINs the stream after the last frame, or resets on error.

7.10. PROBE

PROBE is used to measure the available bitrate of the connection.

```
PROBE Message {  
  Message Length (i)  
  Bitrate (i)  
  RTT (i)  
}
```

Bitrate: When sent by the subscriber (stream opener): the target bitrate in bits per second that the publisher should pad up to. When sent by the publisher (responder): the current estimated bitrate in bits per second. A value of 0 means unknown.

RTT: The smoothed round-trip time in milliseconds, as defined in [RFC9002]. A value of 0 means unknown.

NOTE: RTT is included in the PROBE message because not all QUIC implementations and browser WebTransport APIs expose RTT statistics directly. This field may be deprecated once RTT is universally available via the underlying transport API.

7.11. GOAWAY

A GOAWAY message is sent to initiate a graceful session shutdown with an optional redirect.

```
GOAWAY Message {  
  Message Length (i)  
  New Session URI (s)  
}
```

New Session URI: A URI for the peer to reconnect to. An empty string indicates no redirect; the peer should simply close the session. A recipient **MUST** validate the URI against local policy before reconnecting, including verifying the scheme, authority, and port are permitted. If validation fails, the recipient **MUST** close the session without reconnecting.

7.12. GROUP

The GROUP message contains information about a Group, as well as a reference to the subscription being served.

```
GROUP Message {  
  Message Length (i)  
  Subscribe ID (i)  
  Group Sequence (i)  
}
```

Subscribe ID: The corresponding Subscribe ID. This ID is used to distinguish between multiple subscriptions for the same track.

Group Sequence: The sequence number of the group. This SHOULD increase by 1 for each new group. A subscriber MUST handle gaps, potentially caused by congestion.

7.13. FRAME

The FRAME message is a payload within a group.

```
FRAME Message {  
  Message Length (i)  
  Payload (b)  
}
```

Payload: An application specific payload. A generic library or relay MUST NOT inspect or modify the contents unless otherwise negotiated.

8. Appendix A: Changelog

8.1. moq-lite-04

- * Renamed ANNOUNCE_PLEASE to ANNOUNCE_INTEREST.
- * ANNOUNCE Hops count replaced with explicit Hop ID list for loop detection.
- * Added Exclude Hop to ANNOUNCE_INTEREST for relay loop avoidance.
- * Added GOAWAY stream for graceful session shutdown and migration.
- * Added RTT to PROBE message. Bitrate and RTT use 0 for unknown.

8.2. moq-lite-03

- * Version negotiated via ALPN (moq-lite-xx) instead of SETUP messages.
- * Removed Session, SessionCompat streams and SESSION_CLIENT/SESSION_SERVER/SESSION_UPDATE messages.

- * Unknown stream types reset instead of fatal; enables extension negotiation via stream probing.
- * Added FETCH stream for single group download.
- * Added Start Group and End Group to SUBSCRIBE, SUBSCRIBE_UPDATE, and SUBSCRIBE_OK.
- * Added SUBSCRIBE_DROP on Subscribe stream.
- * Subscribe stream closed (FIN) when all groups accounted for.
- * Added PROBE stream replacing SESSION_UPDATE bitrate.
- * Removed ANNOUNCE_INIT message.
- * Added Hops to ANNOUNCE.
- * Added Subscriber Max Latency and Subscriber Ordered to SUBSCRIBE and SUBSCRIBE_UPDATE.
- * Added Publisher Priority, Publisher Max Latency, and Publisher Ordered to SUBSCRIBE_OK.
- * SUBSCRIBE_OK may be sent multiple times.

8.3. moq-lite-02

- * Added SessionCompat stream.
- * Editorial stuff.

8.4. moq-lite-01

- * Added Message Length (i) to all messages.

9. Appendix B: Upstream Differences

A quick comparison of moq-lite and moq-transport-14:

- * Streams instead of request IDs.
- * Pull only: No unsolicited publishing.
- * FETCH is HTTP-like (single request/response) vs MoqTransport FETCH (multiple groups).
- * Extensions negotiated via stream probing instead of parameters.

- * Both moq-lite and MoqTransport use ALPN for version identification.
- * Names use utf-8 strings instead of byte arrays.
- * Track Namespace is a string, not an array of any array of bytes.
- * Subscriptions default to the latest group, not the latest object.
- * No subgroups
- * No group/object ID gaps
- * No object properties
- * No datagrams
- * No paused subscriptions (forward=0)

9.1. Deleted Messages

- * MAX_SUBSCRIBE_ID
- * REQUESTS_BLOCKED
- * SUBSCRIBE_ERROR
- * UNSUBSCRIBE
- * PUBLISH_DONE
- * PUBLISH
- * PUBLISH_OK
- * PUBLISH_ERROR
- * FETCH_OK
- * FETCH_ERROR
- * FETCH_CANCEL
- * FETCH_HEADER
- * TRACK_STATUS
- * TRACK_STATUS_OK

- * TRACK_STATUS_ERROR
- * PUBLISH_NAMESPACE
- * PUBLISH_NAMESPACE_OK
- * PUBLISH_NAMESPACE_ERROR
- * PUBLISH_NAMESPACE_CANCEL
- * SUBSCRIBE_NAMESPACE_OK
- * SUBSCRIBE_NAMESPACE_ERROR
- * UNSUBSCRIBE_NAMESPACE
- * OBJECT_DATAGRAM

9.2. Renamed Messages

- * SUBSCRIBE_NAMESPACE -> ANNOUNCE_INTEREST
- * SUBGROUP_HEADER -> GROUP

9.3. Deleted Fields

Some of these fields occur in multiple messages.

- * Request ID
- * Track Alias
- * Group Order
- * Filter Type
- * StartObject
- * Expires
- * ContentExists
- * Largest Group ID
- * Largest Object ID
- * Parameters

- * Subgroup ID
- * Object ID
- * Object Status
- * Extension Headers

10. Security Considerations

TODO Security

11. IANA Considerations

This document has no IANA actions.

12. Normative References

- [moqt] Nandakumar, S., Vasiliev, V., Swett, I., and A. Frindell, "Media over QUIC Transport", Work in Progress, Internet-Draft, draft-ietf-moq-transport-17, 2 March 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-moq-transport-17>>.
- [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>>.
- [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>>.
- [RFC9002] Iyengar, J., Ed. and I. Swett, Ed., "QUIC Loss Detection and Congestion Control", RFC 9002, DOI 10.17487/RFC9002, May 2021, <<https://www.rfc-editor.org/rfc/rfc9002>>.

Acknowledgments

TODO acknowledge.

Author's Address

Luke Curley
Email: kixelated@gmail.com