

Media Over QUIC
Internet-Draft
Intended status: Informational
Expires: 23 January 2026

W. Law
Akamai
L. Curley
Twitch
V. Vasiliev
Google
S. Nandakumar
Cisco
K. Pugin
Meta
22 July 2025

WARP Streaming Format
draft-ietf-moq-warp-01

Abstract

This document specifies the WARP Streaming Format, designed to operate on Media Over QUIC Transport.

About This Document

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

The latest revision of this draft can be found at <https://moq-wg.github.io/warp-streaming-format/draft-ietf-moq-warp.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-moq-warp/>.

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

Source for this draft and an issue tracker can be found at <https://github.com/moq-wg/warp-streaming-format>.

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 23 January 2026.

Copyright Notice

Copyright (c) 2025 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	4
2. Conventions and Definitions	4
3. Scope	4
4. Media packaging	6
4.1. LOC packaging	6
4.2. Time-alignment	6
4.3. Content protection and encryption	6
5. Catalog	6
5.1. Catalog Fields	7
5.2. WARP version	9
5.2.1. Delta update	9
5.2.2. Add tracks	9
5.2.3. Remove tracks	9
5.2.4. Clone tracks	10
5.2.5. Generated at	10
5.2.6. Tracks	10
5.2.7. Tracks object	10
5.2.8. Track namespace	10
5.2.9. Track name	10
5.2.10. Packaging	11
5.2.11. Track role	11
5.2.12. Is Live	12
5.2.13. Track label	12
5.2.14. Render group	13
5.2.15. Alternate group	13

5.2.16. Initialization data	13
5.2.17. Dependencies	13
5.2.18. Temporal ID	13
5.2.19. Spatial ID	14
5.2.20. Codec	14
5.2.21. Mimetype	14
5.2.22. Framerate	14
5.2.23. Timescale	14
5.2.24. Bitrate	14
5.2.25. Width	14
5.2.26. Height	15
5.2.27. Audio sample rate	15
5.2.28. Channel configuration	15
5.2.29. Display width	15
5.2.30. Display height	15
5.2.31. Language	15
5.2.32. Parent name	15
5.2.33. Track duration	16
5.3. Delta updates	16
5.4. Catalog Examples	17
5.4.1. Time-aligned Audio/Video Tracks with single quality	17
5.4.2. Simulcast video tracks - 3 alternate qualities along with audio	18
5.4.3. SVC video tracks with 2 spatial and 2 temporal qualities	19
5.4.4. Delta update - adding two tracks	21
5.4.5. Delta update removing tracks	22
5.4.6. Time-aligned Audio/Video Tracks with custom field values	22
5.4.7. Time-aligned VOD Audio/Video Tracks	23
6. Media transmission	24
6.1. Group numbering	24
7. Timeline track	25
7.1. Timeline track payload	25
7.2. Timeline Catalog requirements	26
7.3. Timeline track updating.	26
8. Workflow	26
9. Security Considerations	26
10. IANA Considerations	26
11. Normative References	26
Acknowledgments	28
Authors' Addresses	28

1. Introduction

WARP Streaming Format (WARP) is a media format designed to deliver LOC [LOC] compliant media content over Media Over QUIC Transport (MOQT) [MoQTransport]. WARP works by fragmenting the bitstream into objects that can be independently transmitted. WARP leverages a catalog format to describe the output of the original publisher. WARP specifies how content should be packaged and signaled, defines how the catalog communicates the content, specifies prioritization strategies for real-time and workflows for beginning and terminating broadcasts. WARP also details how end-subscribers may perform adaptive bitrate switching. WARP is targeted at real-time and interactive levels of live latency, as well as VOD content.

This document describes version 1 of the streaming format.

2. Conventions and Definitions

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

This document uses the conventions detailed in Section 1.3 of [RFC9000] when describing the binary encoding.

3. Scope

The purpose of WARP is to provide an interoperable media streaming format operating over [MoQTransport]. Interoperability implies that:

- * An original publisher can package incoming media content into tracks, prepare a catalog and announce the availability of the content to a MOQT relay. Media content refers to audio and video data, as well as ancillary data such as captions, subtitles, accessibility and other timed-text data.
- * A MOQT relay can process the announcement as well as cache and propagate the tracks, both to other relays or to the final subscriber.
- * A final subscriber can parse the catalog, request tracks, decode and render the received media data.

WARP is intended to provide a format for delivering commercial media content. To that end, the following features are within scope:

- * Video codecs - all codecs supported by [LOC]
- * Audio codecs - all audio codecs supported by [LOC]
- * Catalog track - describes the availability and characteristics of content produced by the original publisher.
- * Timeline track - describes the relationship between MOQT Group and Object IDs to media time.
- * Token-based authorization and access control
- * Captions + Subtitles - support for [WEBVTT] and [IMSC1] transmission
- * Latency support across multiple regimes (thresholds are informative only and describe the delay between the original publisher placing the content on the wire and the final subscriber rendering it)
- * Real-time - less than 500ms
- * Interactive - between 500ms and 2500ms
- * Standard - above 2500ms
- * VOD latency - content that was previously produced, is no longer live and is available indefinitely.
- * Content encryption
- * ABR between time-synced tracks - subscribers may switch between tracks at different quality levels in order to maximize visual or audio quality under fconditions of throughput variability.
- * Capable of delivering interstitial advertising.
- * Logs and analytics management - support for the reporting of client-side QoE and relay delivery actions.

Initial verisons of WARP will prioritize basic features necessary to exercise interoperability across delivery systems. Later versions will add commercially necessary features.

4. Media packaging

WARP delivers LOC [LOC] packaged media bitstreams.

4.1. LOC packaging

This specification references Low Overhead Container (LOC) [LOC] to define how audio and video content is packaged. With this packaging mode, each EncodedAudioChunk or EncodedVideoChunk sample is placed in a separate MOQT Object. Samples that belong to the same Group of Pictures (GOP) MUST be placed within the same MOQT Group.

When LOC packaging is used for a track, the catalog packaging attribute (Section 5.2.10) MUST be present and it MUST be populated with a value of "loc".

4.2. Time-alignment

WARP Tracks MAY be time-aligned. Those that are, are subject to the following requirements:

- * Time-aligned tracks MUST be advertised in the catalog as belonging to a common render group.
- * The render duration of the first media object of each equally numbered MOQT Group, after decoding, MUST have overlapping presentation time.

A consequence of this restriction is that a WARP receiver SHOULD be able to cleanly switch between time-aligned media tracks at group boundaries.

4.3. Content protection and encryption

ToDo - content protection for LOC-packaged content.

5. Catalog

A Catalog is a MOQT Track that provides information about the other tracks being produced by a WARP publisher. A Catalog is used by WARP publishers for advertising their output and for subscribers in consuming that output. The payload of the Catalog object is opaque to Relays and can be end-to-end encrypted. The Catalog provides the names and namespaces of the tracks being produced, along with the relationship between tracks, properties of the tracks that consumers may use for selection and any relevant initialization data.

The catalog track MUST have a case-sensitive Track Name of "catalog".

A catalog object MAY be independent of other catalog objects or it MAY represent a delta update of a prior catalog object. The first catalog object published within a new group MUST be independent. A catalog object SHOULD only be published only when the availability of tracks changes.

Each catalog update MUST be mapped to a discreet MOQT Object.

5.1. Catalog Fields

A catalog is a JSON [JSON] document, comprised of a series of mandatory and optional fields. At a minimum, a catalog MUST provide all mandatory fields. A producer MAY add additional fields to the ones described in this draft. Custom field names MUST NOT collide with field names described in this draft. The order of field names within the JSON document is not important.

A parser MUST ignore fields it does not understand.

Table 1 provides an overview of all fields defined by this document.

Field	Name	Definition
WARP version	version	Section 5.2
Delta update	deltaUpdate	Section 5.2.1
Add tracks	addTracks	Section 5.2.2
Remove tracks	removeTracks	Section 5.2.3
Clone tracks	cloneTracks	Section 5.2.4
Generated at	generatedAt	Section 5.2.5
Tracks	tracks	Section 5.2.6
Track namespace	namespace	Section 5.2.8
Track name	name	Section 5.2.9
Packaging	packaging	Section 5.2.10
Is Live	isLive	Section 5.2.12
Track role	role	Section 5.2.11

Track label	label	Section 5.2.13	
+-----+	+-----+	+-----+	+-----+
Render group	renderGroup	Section 5.2.14	
+-----+	+-----+	+-----+	+-----+
Alternate group	altGroup	Section 5.2.15	
+-----+	+-----+	+-----+	+-----+
Initialization data	initData	Section 5.2.16	
+-----+	+-----+	+-----+	+-----+
Dependencies	depends	Section 5.2.17	
+-----+	+-----+	+-----+	+-----+
Temporal ID	temporalId	Section 5.2.18	
+-----+	+-----+	+-----+	+-----+
Spatial ID	spatialId	Section 5.2.19	
+-----+	+-----+	+-----+	+-----+
Codec	codec	Section 5.2.20	
+-----+	+-----+	+-----+	+-----+
Mime type	mimeType	Section 5.2.21	
+-----+	+-----+	+-----+	+-----+
Framerate	framerate	Section 5.2.22	
+-----+	+-----+	+-----+	+-----+
Timescale	timescale	Section 5.2.23	
+-----+	+-----+	+-----+	+-----+
Bitrate	bitrate	Section 5.2.24	
+-----+	+-----+	+-----+	+-----+
Width	width	Section 5.2.25	
+-----+	+-----+	+-----+	+-----+
Height	height	Section 5.2.26	
+-----+	+-----+	+-----+	+-----+
Audio sample rate	samplerate	Section 5.2.27	
+-----+	+-----+	+-----+	+-----+
Channel configuration	channelConfig	Section 5.2.28	
+-----+	+-----+	+-----+	+-----+
Display width	displayWidth	Section 5.2.29	
+-----+	+-----+	+-----+	+-----+
Display height	displayHeight	Section 5.2.30	
+-----+	+-----+	+-----+	+-----+
Language	lang	Section 5.2.31	
+-----+	+-----+	+-----+	+-----+
Parent name	parentName	Section 5.2.32	
+-----+	+-----+	+-----+	+-----+
Track duration	trackDuration	Section 5.2.33	
+-----+	+-----+	+-----+	+-----+

Table 1

Table 2 defines the allowed locations for these fields within the document

Location	Allowed locations for the field
R	The Root of the JSON object
T	Track object

Table 2

5.2. WARP version

Location: R Required: Yes JSON Type: Number

Specifies the version of WARP referenced by this catalog. There is no guarantee that future catalog versions are backwards compatible and field definitions and interpretation may change between versions. A subscriber MUST NOT attempt to parse a catalog version which it does not understand.

5.2.1. Delta update

Location: R Required: Optional JSON Type: Boolean

A Boolean that if true indicates that this catalog object represents a delta (or partial) update. A delta update has a restricted set of fields and special processing rules - see Section 5.3. This value SHOULD NOT be added to a catalog if it is false.

5.2.2. Add tracks

Location: R Required: Optional JSON Type: Array

Indicates a delta processing instruction to add new tracks. The value of this field is an Array of track objects Section 5.2.7.

5.2.3. Remove tracks

Location: R Required: Optional JSON Type: Array

Indicates a delta processing instruction to remove new tracks. The value of this field is an Array of track objects Section 5.2.7. Each track object MUST include a Track Name Section 5.2.9 field, MAY include a Track Namespace Section 5.2.8 field and MUST NOT hold any other fields.

5.2.4. Clone tracks

Location: R Required: Optional JSON Type: Array

Indicates a delta processing instruction to clone new tracks from previously declared tracks. The value of this field is an Array of track objects Section 5.2.7. Each track object MUST include a Parent Name Section 5.2.32 field.

5.2.5. Generated at

Location: R Required: Optional JSON Type: Number

The wallclock time at which this catalog instance was generated, expressed as the number of milliseconds that have elapsed since January 1, 1970 (midnight UTC/GMT). This field SHOULD NOT be included if the isLive field is false.

5.2.6. Tracks

Location: R Required: Yes JSON Type: Array

An array of track objects Section 5.2.7.

5.2.7. Tracks object

A track object is JSON Object containing a collection of fields whose location is specified 'T' in Table 2.

5.2.8. Track namespace

Location: T Required: Optional JSON Type: String

The name space under which the track name is defined. See section 2.3 of [MoQTransport]. The track namespace is optional. If it is not declared within a track, then each track MUST inherit the namespace of the catalog track. A namespace declared in a track object overwrites any inherited name space.

5.2.9. Track name

Location: T Required: Yes JSON Type: String

A string defining the name of the track. See section 2.3 of [MoQTransport]. Within the catalog, track names MUST be unique per namespace.

5.2.10. Packaging

Location: T Required: Yes JSON Type: String

A string defining the type of payload encapsulation. Allowed values are strings as defined in Table 3.

Table 3: Allowed packaging values

Name	Value	Reference
LOC	loc	See RFC XXXX
Timeline	timeline	See Section 7

Table 3

5.2.11. Track role

Location: T Required: Optional JSON Type: String

A string defining the role of content carried by the track. Reserved roles are described in Table 4. These role values are case-sensitive.

This role field MAY be used in conjunction with the Mimetype Section 5.2.21 to fully describe the content of the track.

Table 4: Reserved track roles

Role	Description
audiodescription	An audio description for visually impaired users
video	Visual content
audio	Audio content
timeline	A WARP timeline Section 7
caption	A textual representation of the audio track
subtitle	A transcription of the spoken dialogue
signlanguage	A visual track for hearing impaired users.

Table 4

Custom roles MAY be used as long as they do not collide with the reserved roles.

5.2.12. Is Live

Location: T Required: Required JSON Type: Boolean

True if new Objects will be added to the track. False if no new Objects will be added to the track. This is sent under two possible conditions: * the publisher of a previously live track has ended the track. * the track is Video-On-Demand (VOD) and was never live.

5.2.13. Track label

Location: TF Required: Optional JSON Type: String

A string defining a human-readable label for the track. Examples might be "Overhead camera view" or "Deutscher Kommentar". Note that the [JSON] spec requires UTF-8 support by decoders.

5.2.14. Render group

Location: T Required: Optional JSON Type: Number

An integer specifying a group of tracks which are designed to be rendered together. Tracks with the same group number SHOULD be rendered simultaneously, are usually time-aligned and are designed to accompany one another. A common example would be tying together audio and video tracks.

5.2.15. Alternate group

Location: T Required: Optional JSON Type: Number

An integer specifying a group of tracks which are alternate versions of one-another. Alternate tracks represent the same media content, but differ in their selection properties. Alternate tracks SHOULD have matching framerate Section 5.2.22 and media time sequences. A subscriber typically subscribes to one track from a set of tracks specifying the same alternate group number. A common example would be a set video tracks of the same content offered in alternate bitrates.

5.2.16. Initialization data

Location: T Required: Optional JSON Type: String

A string holding Base64 [BASE64] encoded initialization data for the track.

5.2.17. Dependencies

Location: T Required: Optional JSON Type: Array

Certain tracks may depend on other tracks for decoding. Dependencies holds an array of track names Section 5.2.9 on which the current track is dependent. Since only the track name is signaled, the namespace of the dependencies is assumed to match that of the track declaring the dependencies.

5.2.18. Temporal ID

Location: T Required: Optional JSON Type: Number

A number identifying the temporal layer/sub-layer encoding of the track, starting with 0 for the base layer, and increasing with higher temporal fidelity.

5.2.19. Spatial ID

Location: T Required: Optional JSON Type: Number

A number identifying the spatial layer encoding of the track, starting with 0 for the base layer, and increasing with higher fidelity.

5.2.20. Codec

Location: T Required: Optional JSON Type: String

A string defining the codec used to encode the track. For LOC packaged content, the string codec registrations are defined in Sect 3 and Section 4 of [WEBCODECS-CODEC-REGISTRY].

5.2.21. Mimetype

Location: T Required: Optional JSON Type: String

A string defining the mime type [MIME] of the track.

5.2.22. Framerate

Location: T Required: Optional JSON Type: Number

A number defining the video framerate of the track, expressed as frames per second.

5.2.23. Timescale

Location: T Required: Optional JSON Type: Number

The number of time units that pass per second.

5.2.24. Bitrate

Location: T Required: Optional JSON Type: Number

A number defining the bitrate of track, expressed in bits per second.

5.2.25. Width

Location: T Required: Optional JSON Type: Number

A number expressing the encoded width of the video frames in pixels.

5.2.26. Height

Location: T Required: Optional JSON Type: Number

A number expressing the encoded height of the video frames in pixels.

5.2.27. Audio sample rate

Location: T Required: Optional JSON Type: Number

The number of audio frame samples per second. This property SHOULD only accompany audio codecs.

5.2.28. Channel configuration

Location: T Required: Optional JSON Type: String

A string specifying the audio channel configuration. This property SHOULD only accompany audio codecs. A string is used in order to provide the flexibility to describe complex channel configurations for multi-channel and Next Generation Audio schemas.

5.2.29. Display width

Location: T Required: Optional JSON Type: Number

A number expressing the intended display width of the track content in pixels.

5.2.30. Display height

Location: T Required: Optional JSON Type: Number

A number expressing the intended display height of the track content in pixels.

5.2.31. Language

Location: T Required: Optional JSON Type: String

A string defining the dominant language of the track. The string MUST be one of the standard Tags for Identifying Languages as defined by [LANG].

5.2.32. Parent name

Location: T Required: Optional JSON Type: String

A string defining the parent track name Section 5.2.9 to be cloned. This field MUST only be included inside a Clone tracks Section 5.2.4 object.

5.2.33. Track duration

Location: T Required: Optional JSON Type: Number

The duration of the track expressed in integer milliseconds. This field MUST NOT be included if the isLive Section 5.2.12 field value is false.

5.3. Delta updates

A catalog update might contain incremental changes. This is a useful property if many tracks may be initially declared but then there are small changes to a subset of tracks. The producer can issue a delta update to describe these changes. Changes are described incrementally, meaning that a delta update can itself modify a prior delta update.

A restricted set of operations are allowed with each delta update: * Add a new track that has not previously been declared. * Add a new track by cloning a previously declared track. * Remove a track that has been previously declared.

The following rules are to be followed in constructing and processing delta updates:

- * A delta update MUST include the Delta Update Section 5.2.1 field set to true.
- * A delta update catalog MUST contain at least one instance of Add tracks Section 5.2.2, Remove tracks Section 5.2.3 or Clone Tracks Section 5.2.4 fields and MAY contain more. It MUST NOT contain an instance of a Tracks Section 5.2.6 field or a WARP version Section 5.2 field.
- * The Add, Delete and Clone operations are applied sequentially in the order they are declared in the document. Each operation in the sequence is applied to the target document; the resulting document becomes the target of the next operation. Evaluation continues until all operations are successfully applied.
- * A Cloned track inherits all the attributes of the track defined by the Parent Name Section 5.2.32, except the Track Name which MUST be new. Attributes redefined in the cloning Object overwrite inherited values.

- * The tuple of Track Namespace and Track Name defines a fixed set of Track attributes which MUST NOT be modified after being declared. To modify any attribute, a new track with a different Namespace|Name tuple is created by Adding or Cloning and then the old track is removed.

5.4. Catalog Examples

The following section provides non-normative JSON examples of various catalogs compliant with this draft.

5.4.1. Time-aligned Audio/Video Tracks with single quality

This example shows a catalog for a media producer capable of sending LOC packaged, time-aligned audio and video tracks.

```
{
  "version": 1,
  "generatedAt": 1746104606044,
  "tracks": [
    {
      "name": "1080p-video",
      "namespace": "conference.example.com/conferencel23/alice",
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "renderGroup": 1,
      "codec": "av01.0.08M.10.0.110.09",
      "width": 1920,
      "height": 1080,
      "framerate": 30,
      "bitrate": 1500000
    },
    {
      "name": "audio",
      "namespace": "conference.example.com/conferencel23/alice",
      "packaging": "loc",
      "isLive": true,
      "role": "audio",
      "renderGroup": 1,
      "codec": "opus",
      "samplerate": 48000,
      "channelConfig": "2",
      "bitrate": 32000
    }
  ]
}
```

5.4.2. Simulcast video tracks - 3 alternate qualities along with audio

This example shows catalog for a media producer capable of sending 3 time-aligned video tracks for high definition, low definition and medium definition video qualities, along with an audio track. In this example the namespace is absent, which infers that each track must inherit the namespace of the catalog.

```
{
  "version": 1,
  "generatedAt": 1746104606044,
  "tracks": [
    {
      "name": "hd",
      "renderGroup": 1,
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "codec": "av01",
      "width": 1920,
      "height": 1080,
      "bitrate": 5000000,
      "framerate": 30,
      "altGroup": 1
    },
    {
      "name": "md",
      "renderGroup": 1,
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "codec": "av01",
      "width": 720,
      "height": 640,
      "bitrate": 3000000,
      "framerate": 30,
      "altGroup": 1
    },
    {
      "name": "sd",
      "renderGroup": 1,
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "codec": "av01",
      "width": 192,
      "height": 144,
      "bitrate": 500000,

```


```

    "framerate":30,
    "altGroup":1
  },
  {
    "name": "audio",
    "renderGroup": 1,
    "packaging": "loc",
    "isLive": true,
    "role": "audio",
    "codec":"opus",
    "samplerate":48000,
    "channelConfig":"2",
    "bitrate":32000
  }
]
}

```

5.4.3. SVC video tracks with 2 spatial and 2 temporal qualities

This example shows catalog for a media producer capable of sending scalable video codec with 2 spatial and 2 temporal layers with a dependency relation as shown below:



The corresponding catalog uses "depends" attribute to express the track relationships.

```
{
  "version": 1,
  "generatedAt": 1746104606044,
  "tracks": [
    {
      "name": "480p15",
      "namespace": "conference.example.com/conference123/alice",
      "renderGroup": 1,
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "codec": "av01.0.01M.10.0.110.09",
      "width": 640,
      "height": 480,
      "bitrate": 3000000,
      "framerate": 15
    },
    {
      "name": "480p30",
      "namespace": "conference.example.com/conference123/alice",
      "renderGroup": 1,
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "codec": "av01.0.04M.10.0.110.09",
      "width": 640,
      "height": 480,
      "bitrate": 3000000,
      "framerate": 30,
      "depends": ["480p15"]
    },
    {
      "name": "1080p15",
      "namespace": "conference.example.com/conference123/alice",
      "renderGroup": 1,
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "codec": "av01.0.05M.10.0.110.09",
      "width": 1920,
      "height": 1080,
      "bitrate": 3000000,
      "framerate": 15,
      "depends": ["480p15"]
    },
    {
      "name": "1080p30",
```

```
    "namespace": "conference.example.com/conference123/alice",
    "renderGroup": 1,
    "packaging": "loc",
    "isLive": true,
    "role": "video",
    "codec": "av01.0.08M.10.0.110.09",
    "width": 1920,
    "height": 1080,
    "bitrate": 5000000,
    "framerate": 30,
    "depends": ["480p30", "1080p15"]
  },
  {
    "name": "audio",
    "namespace": "conference.example.com/conference123/alice",
    "renderGroup": 1,
    "packaging": "loc",
    "isLive": true,
    "role": "audio",
    "codec": "opus",
    "samplerate": 48000,
    "channelConfig": "2",
    "bitrate": 32000
  }
]
```

5.4.4. Delta update - adding two tracks

This example shows the catalog delta update for a media producer adding two tracks to an established video conference. One track is newly declared, the other is cloned from a previous track.

```
{
  "deltaUpdate": true,
  "generatedAt": 1746104606044,
  "addTracks": [
    {
      "name": "slides",
      "isLive": true,
      "role": "video",
      "codec": "av01.0.08M.10.0.110.09",
      "width": 1920,
      "height": 1080,
      "framerate": 15,
      "bitrate": 750000,
      "renderGroup": 1
    }
  ],
  "cloneTracks": [
    {
      "parentName": "video-1080",
      "name": "video-720",
      "width": 1280,
      "height": 720,
      "bitrate": 600000
    }
  ]
}
```

5.4.5. Delta update removing tracks

This example shows a delta update for a media producer removing two tracks from an established video conference.

```
{
  "deltaUpdate": true,
  "generatedAt": 1746104606044,
  "removeTracks": [{"name": "video"}, {"name": "slides"}]
}
```

5.4.6. Time-aligned Audio/Video Tracks with custom field values

This example shows catalog for a media producer capable of sending LOC packaged, time-aligned audio and video tracks along with custom fields in each track description.

```
{
  "version": 1,
  "generatedAt": 1746104606044,
  "tracks": [
    {
      "name": "1080p-video",
      "namespace": "conference.example.com/conference123/alice",
      "packaging": "loc",
      "isLive": true,
      "role": "video",
      "renderGroup": 1,
      "codec": "av01.0.08M.10.0.110.09",
      "width": 1920,
      "height": 1080,
      "framerate": 30,
      "bitrate": 1500000,
      "com.example-billing-code": 3201,
      "com.example-tier": "premium",
      "com.example-debug": "h349835bfkjfg82394d945034jsdfn349fns"
    },
    {
      "name": "audio",
      "namespace": "conference.example.com/conference123/alice",
      "packaging": "loc",
      "isLive": true,
      "role": "audio",
      "renderGroup": 1,
      "codec": "opus",
      "samplerate": 48000,
      "channelConfig": "2",
      "bitrate": 32000
    }
  ]
}
```

5.4.7. Time-aligned VOD Audio/Video Tracks

This example shows catalog for a media producer offering VOD (video on-demand) non-live content. The content is LOC packaged, and includes time-aligned audio and video tracks.

```
{
  "version": 1,
  "tracks": [
    {
      "name": "video",
      "namespace": "movies.example.com/assets/boy-meets-girl-season3/episode5",
      "packaging": "loc",
      "isLive": false,
      "trackDuration": 8072340,
      "renderGroup": 1,
      "codec": "av01.0.08M.10.0.110.09",
      "width": 1920,
      "height": 1080,
      "framerate": 30,
      "bitrate": 1500000
    },
    {
      "name": "audio",
      "namespace": "movies.example.com/assets/boy-meets-girl-season3/episode5",
      "packaging": "loc",
      "isLive": false,
      "trackDuration": 8072340,
      "renderGroup": 1,
      "codec": "opus",
      "samplerate": 48000,
      "channelConfig": "2",
      "bitrate": 32000
    }
  ]
}
```

6. Media transmission

The MOQT Groups and MOQT Objects need to be mapped to MOQT Streams. Irrespective of the Section 4 in place, each MOQT Object **MUST** be mapped to a new MOQT Stream.

6.1. Group numbering

The Group ID of the first Group published in a track at application startup **MUST** be a unique integer that will not repeat in the future. One approach to achieve this is to set the initial Group ID to the creation time of the first Object in the group, represented as the number of milliseconds since the Unix epoch, rounded to the nearest millisecond. This ensures that republishing the same track in the future, such as after a loss of connectivity or an encoder restart, will not result in smaller or duplicate Group IDs for the same track name. Note that this method does not prevent duplication if more

than 1000 groups are published per second.

Each subsequent Group ID MUST increase by 1.

If a publisher is able to maintain state across a republish, it MUST signal the gap in Group IDs using the MOQT Prior Group ID Gap Extension header.

7. Timeline track

The timeline track provides data about the previously published groups and their relationship to wallclock time, media time and associated timed-metadata. Timeline tracks allow players to seek to precise points behind the live head in a live broadcast, or for random access in a VOD asset. A timeline track might also be used to insert events at media times which do not correlate with Object boundaries. Timeline tracks are optional. Multiple timeline tracks can exist inside a catalog.

7.1. Timeline track payload

The payload of a timeline track is a UTF-8 encoded CSV text file. This payload is formatted according to RFC4180 "Common Format and MIME Type for Comma-Separated Values (CSV)" Files [RFC4180]. The separator is a comma and each line is separated by a carriage return. The mime-type of a timeline track MUST be specified as "text/csv" in the catalog.

Each timeline track begins with a header row of MEDIA_PTS, GROUP_ID, OBJECT_ID, WALLCLOCK, METADATA. This row defines the 5 columns of data within each record.

- * MEDIA_PTS: a media timestamp rounded to the nearest millisecond. This entry MUST NOT be empty. If the Object ID entry is present, then this value MUST match the media presentation timestamp of the first media sample in the referenced Object.
- * GROUP_ID: the MOQT Group ID. This entry MAY be empty.
- * OBJECT_ID: the MOQT Object ID. This entry MAY be empty.
- * WALLCLOCK: the wallclock time at which the media was encoded, expressed as the number of milliseconds that have elapsed since January 1, 1970 (midnight UTC/GMT). For VOD assets, or if the wallclock time is not known, the value SHOULD be 0.

- * METADATA: a flexible field holding arbitrary string metadata. This field may be empty. If not empty, it MUST be enclosed in double quotes. A double-quote appearing inside this field MUST be escaped by preceding it with another double quote.

7.2. Timeline Catalog requirements

A timeline track MUST carry a 'type' identifier in the Catalog with a value of "timeline". A timeline track MUST carry a 'dependes' attribute which contains an array of all track names to which the timeline track applies.

7.3. Timeline track updating.

The publisher MUST publish a complete timeline in the first MOQT Object of each MOQT Group of a timeline track. The publisher MAY publish incremental updates in the second and subsequent Objects within each GROUP. Incremental updates only contain timeline events since the last timeline Object. Group duration SHOULD not exceed 30 seconds inside a timeline track.

8. Workflow

A WARP publisher MUST publish a catalog track object before publishing any media track objects.

At the completion of a session, a publisher MUST publish a catalog update that removes all currently active tracks. This action SHOULD be interpreted by receivers to mean that the publish session is complete.

9. Security Considerations

ToDo

10. IANA Considerations

This document creates a new entry in the "MoQ Streaming Format" Registry (see [MoQTransport] Sect 8). The type value is 0x001, the name is "WARP Streaming Format" and the RFC is XXX.

11. Normative References

- [BASE64] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/rfc/rfc4648>>.

- [IMSC1] "W3C, TTML Profiles for Internet Media Subtitles and Captions 1.0 (IMSC1)", April 2016, <<https://www.w3.org/TR/ttml-imscl/>>.
- [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>>.
- [LANG] Phillips, A., Ed. and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, DOI 10.17487/RFC5646, September 2009, <<https://www.rfc-editor.org/rfc/rfc5646>>.
- [LOC] Zanaty, M., Nandakumar, S., and P. Thatcher, "Low Overhead Media Container", Work in Progress, Internet-Draft, draft-mzanaty-moq-loc-05, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-mzanaty-moq-loc-05>>.
- [MIME] Freed, N., Klensin, J., and T. Hansen, "Media Type Specifications and Registration Procedures", BCP 13, RFC 6838, DOI 10.17487/RFC6838, January 2013, <<https://www.rfc-editor.org/rfc/rfc6838>>.
- [MoQTransport]
Nandakumar, S., Vasiliev, V., Swett, I., and A. Frindell, "Media over QUIC Transport", Work in Progress, Internet-Draft, draft-ietf-moq-transport-11, 28 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-moq-transport-11>>.
- [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>>.
- [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>>.

[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>>.

[WEBCODECS-CODEC-REGISTRY] "WebCodecs Codec Registry", September 2024, <<https://www.w3.org/TR/webcodecs-codec-registry/>>.

[WEBVTT] "World Wide Web Consortium (W3C), WebVTT: The Web Video Text Tracks Format", April 2019, <<https://www.w3.org/TR/webvtt1/>>.

Acknowledgments

* the MoQ Workgroup and mailing lists.

Authors' Addresses

Will Law
Akamai
Email: wilaw@akamai.com

Luke Curley
Twitch
Email: kixelated@gmail.com

Victor Vasiliev
Google
Email: vasilvv@google.com

Suhas Nandakumar
Cisco
Email: snandaku@cisco.com

Kirill Pugin
Meta
Email: ikir@meta.com