

MSF & CMSF Updates & Issues

<https://github.com/moq-wg/msf>

Will Law & Suhas Nandakumar

IETF #125, Shenzhen, March 2026

Brain-storming side meeting - Feb 5, 2026

We held a brain-storming meeting to gather input and comments on MSF & CMSF.

32 attendees, 17 issues filed.

Good discussions with representatives of companies across the OTT landscape.

Agenda [available](#).

Thanks to the room sponsor



Rename complete (PR#71)

At IETF #124 we decided to rename the the two adopted streaming formats:

WARP -> MOQT Streaming Format (**MSF**)

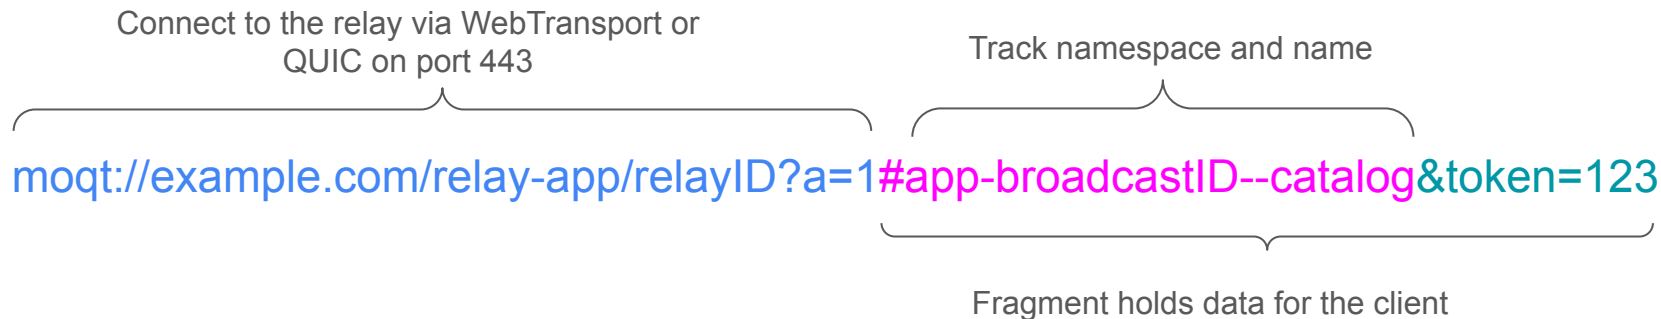
CARP -> CMAF-Compliant MOQT Streaming Format (**CMSF**)

I-D drafts are updated:

<https://datatracker.ietf.org/doc/draft-ietf-moq-msf/>

<https://datatracker.ietf.org/doc/draft-ietf-moq-cmsf/>

URL format #87 - fragment -based scheme



Alan's PR#1571 - fragment identifier

moqt://example.com/relay-app/relayID?a=1#msf:app-broadcastID--catalog&token=123

Note:

1. This proposal uses the namespace+name syntax defined in MOQT.
2. The fragment is not sent to the relay and is only intended for the client.
3. The fragment can carry parameters and some of these are reserved.

Name	Description
wallclock-range	A subclip defined by a wallclock time range
mediatime-range	A subclip defined by a media time range
location-range	A subclip defined by a MOQT Location range
c4m	A base64 encoded C4M token

Add support for variable substitution #123

For scalable delivery, we want to cache catalogs and provide the same content to all viewers, while still allowing each to be customized for **advertising**, A/B **watermarking**, **QoE reporting** or **logging** purposes. The % character is prohibited for non-variable usage.

```
moqt://example.com/relay#app-live--catalog& id=12345&event=baseball&token=h487cn  
32
```

```
{  
  "name": "cmcdv2-%id%",  
  "namespace": "advertising-decisions/live-sports/%event%",  
  "packaging": "eventtimeline",  
  "eventTimelineType": "com.example.iab.vast",  
  "c4m": "%token%"  
}, ...
```

Support Encryption Scheme #116

MSF supports end-to-end encryption of media content using **MoQ Secure Objects** (<https://datatracker.ietf.org/doc/draft-ietf-moq-secure-objects-00>)

This example shows a catalog for encrypted LOC-packaged audio and video tracks using MoQ Secure Objects with AES-128-GCM.

This is not DRM. It is E2E encryption of MOQT Objects across a distribution network.

```
{
  "version": 1,
  "generatedAt": 1746104606044,
  "tracks": [
    {
      "name": "1080p-video",
      "namespace":
"conference.example.com/conference123/alice",
      "packaging": "loc",
      "isLive": true,
      "targetLatency": 2000,
      "role": "video",
      "renderGroup": 1,
      "codec": "av01.0.08M.10.0.110.09",
      "width": 1920,
      "height": 1080,
      "framerate": 30,
      "bitrate": 1500000,
      "encryptionScheme": "moq-secure-objects",
      "cipherSuite": "aes-128-gcm-sha256",
      "keyId": "key-2024-q1-premium",
      "trackBaseKey": "dGhpc2lzYXNhbXBsZWJhc2VrZXk="
    },
    {
      "name": "audio",
      "namespace":
"conference.example.com/conference123/alice",
      "packaging": "loc",
      "isLive": true,
      "targetLatency": 2000,
      "role": "audio",
      "renderGroup": 1,
      "codec": "opus",
      "samplerate": 48000,
      "channelConfig": "2",
      "bitrate": 32000,
      "encryptionScheme": "moq-secure-objects",
      "cipherSuite": "aes-128-gcm-sha256",
      "keyId": "key-2024-q1-premium",
      "trackBaseKey": "dGhpc2lzYXNhbXBsZWJhc2VrZXk="
    }
  ]
}
```

Define template for Media Timeline track #127

When the relationship between media time, MOQT Location and wallclock time follows a **regular, predictable pattern**, a media timeline **template** MAY be used instead of an explicit media timeline track.


Best suited for content with **fixed-duration Groups**, such as VOD assets or live broadcasts with constant segment durations.

```
{
  "name": "1080p-video",
  "namespace":
"conference.example.com/conference123/alice",
  "packaging": "loc",
  "isLive": true,
  "targetLatency": 2000,
  "role": "video",
  "renderGroup": 1,
  "template": [0, 2002, [0, 0], [1, 0], 1759924158381, 2002],
  "codec": "av01.0.08M.10.0.110.09",
  "width": 1920,
  "height": 1080,
  "framerate": 30,
  "bitrate": 1500000
},
```

Equivalent media timeline JSON

```
0, [0, 0], 1759924158381
2002, [1, 0], 1759924160383
4004, [2, 0], 1759924162385
6006, [3, 0], 1759924164387
8008, [4, 0], 1759924166389
10010, [5, 0], 1759924168391
12012, [6, 0], 1759924170393
14014, [7, 0], 1759924172395
16016, [8, 0], 1759924174397
18018, [9, 0], 1759924176399
20020, [10, 0], 1759924178401
22022, [11, 0], 1759924180403
24024, [12, 0], 1759924182405
26026, [13, 0], 1759924184407
28028, [14, 0], 1759924186409
30030, [15, 0], 1759924188411
32032, [16, 0], 1759924190413
34034, [17, 0], 1759924192415
36036, [18, 0], 1759924194417
38038, [19, 0], 1759924196419
40040, [20, 0], 1759924198421
42042, [21, 0], 1759924200423
44044, [22, 0], 1759924202425
46046, [23, 0], 1759924204427
48048, [24, 0], 1759924206429
50050, [25, 0], 1759924208431
52052, [26, 0], 1759924210433
54054, [27, 0], 1759924212435
56056, [28, 0], 1759924214437
58058, [29, 0], 1759924216439
60060, [30, 0], 1759924218441
62062, [31, 0], 1759924220443
64064, [32, 0], 1759924222445
66066, [33, 0], 1759924224447
...
```

PR for discussion: initial text on zapping #122



This provides **fast start** and **fast switching** between tracks with dissimilar keyframe intervals. Both tracks are **time aligned at Object boundaries, not only Group boundaries**. Track A then has a keyframe every N (3) objects, Track B every M (6) objects.

The player can switch from A to B when $(ID_A - 1) * N / M + 1$ is an integer.

Try to switch at Group 2 (ID_A) gives 1.5 , so cannot switch

Try to switch at Group 3 gives 2, so can switch and the target group ID in B is 2.

We can enable switching across arbitrary intervals by adding a **Keyframe Spacing Parameter** to the catalog entry for each track. Default is 1 when every group is aligned.

PR for discussion: Pub tracks, logs and metrics #121

We want the catalog to be able to **instruct players where to publish log and QOE data.**

New **publishTracks** array defines a set publish.

For each, it defines

- Namespace
- Name
- Packaging type
- Role
- Optional token

Useful when combined with **variable injection** to allow a cached catalog to still define custom destinations.

```
"publishTracks": [  
  {  
    "namespace": "moq://metrics.moq.arpa/v1%resourceId%",  
    "name": "4",  
    "packaging": "moqmetrics",  
    "role": "metrics",  
    "token": "eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.."  
  },  
  {  
    "namespace": "moq://moq-syslog.arpa/logs-v1%resourceId%",  
    "name": "6",  
    "packaging": "moqlog",  
    "role": "log",  
    "connectionUri": "moqt://logs.example.com:4443",  
    "token": "eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.."  
  }  
]
```

PR for discussion: Add support for InitTracks #141

Currently init data is specified in the catalog as a base64 track attribute.

What if we want to be able to change the init data midway through a long broadcast? For example - a 1080p movie followed by 2160p sports segment?

Introducing InitTracks - a track holding init updates.

Define like any other track, but with a defined role of 'init'.

Each group in the track holds a new init update.

Current challenge - how to sync these updates with the media tracks they are initializing?

```
"tracks":[
  {
    "name": "video-init",
    "namespace": "streaming.example.com/channel1",
    "packaging": "loc",
    "role": "init"
  },
  {
    "name": "hd",
    "namespace": "streaming.example.com/channel1",
    "renderGroup": 1,
    "packaging": "loc",
    "isLive": true,
    "targetLatency": 1500,
    "role": "video",
    "codec": "avc1.640028",
    "width": 1920,
    "height": 1080,
    "bitrate": 5000000,
    "framerate": 30,
    "altGroup": 1,
    "initTrack": "video-init"
  },

```

PR for discussion: Add SCTE-35 support and CEA-608/708 accessibility fields #133 (1 of 3)

- Allows defining that an event timeline track carries **SCTE-35** data.
- Defines a packaging format for carriage of SCTE-35 in an event timeline.

```
"tracks": [  
  {  
    "name": "scte35",  
    "packaging": "eventtimeline",  
    "eventType": "urn:scte:scte35:2013:bin",  
    "mimeType": "application/json",  
    "isLive": true,  
    "role": "eventtimeline",  
    "depends": ["video"]  
  }  
]
```

Example of a event timeline payload carrying **SCTE-35** data

```
[  
  {  
    "m": 120000,  
    "data": {  
      "type": "splice_insert",  
      "eventId": 12345,  
      "outOfNetwork": true,  
      "duration": 30000,  
      "bin": "/DA0AAAA..."  
    }  
  },  
  {  
    "m": 150000,  
    "data": {  
      "type": "splice_insert",  
      "eventId": 12345,  
      "outOfNetwork": false  
    }  
  }  
]
```

PR for discussion: Add SCTE-35 support and CEA-608/708 accessibility fields #133 (2 of 3)

For **out-of-band captions**, publishers can generate event timelines with one of these two types:

```
`urn:msf:captions:webvtt`  
`urn:msf:captions:imsc1`
```

The IMSC1 carriage is still TBD.

Example of a event timeline payload carrying **webvtt** data

```
[  
  {  
    "m": 0,  
    "data": {  
      "start": 0,  
      "end": 2500,  
      "text": "Welcome to the show."  
    }  
  },  
  {  
    "m": 2500,  
    "data": {  
      "start": 2500,  
      "end": 5000,  
      "text": "Today we will be discussing..."  
    }  
  }  
]
```

PR for discussion: Add SCTE-35 support and CEA-608/708 accessibility fields #133 (3 of 3)

- Allows defining that a track contains **embedded CEA-608/708** captions

```
"tracks": [  
  {  
    "name": "video",  
    "packaging": "loc",  
    "isLive": true,  
    "codec": "avc1.4d401f",  
    "width": 1920,  
    "height": 1080,  
    "framerate": 30,  
    "bitrate": 5000000,  
    "accessibility": [  
      {  
        "scheme": "urn:scte:dash:cc:cea-608:2015",  
        "value": "CC1=eng;CC3=spa"  
      }  
    ]  
  },  
]
```

PR for discussion: add details of authorization flows #118

We need to be able to tell a client when it needs to supply a token in order to access a track.

MSF/CMSF supports two token schemes:

Scheme	Value	Reference
Privacy Pass	privacy-pass	[PrivacyPassAuth]
CAT	cat	[C4M]

Question:

1. Should we declare tokens in one place and then be able to refer to them by ID?

```
"tracks": [  
  {  
    "name": "premium-4k-video",  
    "namespace": "streaming.example.com/live/sports",  
    "packaging": "loc",  
    "isLive": true,  
    ...  
    "width": 3840,  
    "height": 2160,  
    "framerate": 60,  
    "bitrate": 15000000,  
    "authInfo": {  
      "cat": {%cat-token%}  
    }  
  },  
  {  
    "name": "standard-720p-video",  
    "namespace": "streaming.example.com/live/sports",  
    "packaging": "loc",  
    "isLive": true,  
    ...  
    "width": 1280,  
    "height": 720,  
    "framerate": 30,  
    "bitrate": 2500000,  
    "authInfo": {  
      "privacy-pass": {%pp-token%}  
    }  
  },  
]
```

Thank you for your time.