

moqt://

Alan Frindell, Will Law

Transport Selection

At Boulder Interim, we designed (on the fly): PR #1486

moqt:// → send ALPNs: h3, moqt-16, moqt-15

If server selects:

moqt-NN → native QUIC + MOQT version

H3 → Send WebTransport CONNECT with WT-Available-Protocols

But we also added...

moqt+q:// → don't send h3

moqt+wt:// → only send h3

On second thought...

moqt://foo.com/bar and moqt+q://foo.com/bar represent *the same resource*

URI compares exactly on the scheme – we're fighting the system

And...what's the point?

Server Supports → ClientSupports ↓	QUIC	WT	BOTH
QUIC	QUIC	✗	QUIC
WT	✗	WT	WT
BOTH	QUIC	WT	Server Pref

MSF/CMSF URLs

MSF currently inherits the MOQT URI syntax (as defined by PR#1486) and then extends it with a fragment definition to convey full track name along with KV parameters intended for the client.

```
moqt://example.com/relay-app?a=1#namespace1-namespace2--name&token=123
```

```
moqt://example.com/relay-app#customerID-broadcastID--catalog
```

Architecturally, the fragment data is intended only for the client and is not conveyed directly to the server. We could add definitions for reserved client parameters which define the transport to be used:

```
moqt://example.com/relay-app#customerID-broadcastID--catalog&transport=wt
```

```
moqt://example.com/relay-app#customerID-broadcastID--catalog&transport=q
```

Or we can always allow the client to negotiate the transport it prefers.

Fragments

Since MSF / CMSF are planning to use the moqt:// fragment to convey instructions to the client

What, if anything, do we need to say in MOQT about that?

If each streaming format is allowed to define the fragment, how does a MOQT client know which one?