

QUIC Abstractions

Proposed Abstractions

- QUIC (base)
- Connections
- Streams
- H3 (compression)
- HTTP

QUIC (base)

- Packets
- MTU discovery (packet sizing)
- Version negotiation
- Packet loss detection
- A cryptographic context enabling data encryption within a packet
- Zero-RTT connection establishment with limited data payload
- One-RTT connection establishment

QUIC Connections

- Identification of the connection including a Connection ID in addition to the 5-tuple
- Alternate connection IDs/connID 'renaming' without requiring connection re-establishment
- Multiplexed, non HoL blocking, streams
- Congestion control on a per-path basis
- Data (not packet) retransmission
- Flow control on a per-connection basis
- Mechanisms to prove liveness, measure-RTT

QUIC Streams

- Flow control on a per-stream basis
- Ordered but not necessarily in-order bytestreams
- Grouping: a statement that streams should be delivered to the same endpoint through proxies
- Data frames
- Support for non-data frames

QUIC H3

- Flow-controlled headers frames on streams
- Compression for headers data in a robust way which trades off HoL blocking and compression efficiency

QUIC HTTP

- Maps requests to streams using H3
- Defines restrictions on header/data frame sequencing in line with HTTP semantics

Informing Observations

- HTTP is not reliable; it is request-response
- L4 routing is a must-solve

What is Grouping?

Put simply grouping is:

- A request from sender that all streams in the group terminate at the same L7 location

Why Grouping?

- Notifications either incur latency for long-poll re-request or use $O(n)$ state
- DASH live video serving works... not wonderfully
- Allows stateful overlays on top of HTTP using multiple requests, and potentially taking advantage of caching

Why Grouping?

What happens without grouping for stateful data?

If you don't have grouping (i.e. a guarantee that particular streams go to one endpoint), then you can't safely use multiple streams: The streams may go to separate endpoints behind a proxy and be uninterpretable.

Thus, for stateful data, you can't safely set different priorities, dependencies, etc.

This would be sad.