QUIC on Streams

draft-kazuho-quic-quic-on-streams

Kazuho Oku, Lucas Pardue

QUIC is (becoming) a huge success

- better than TCP / TLS:
  - faster handshake
  - path migration
  - no head-of-line blocking
  - better loss recovery
  - better preservation of privacy
  - future proof (prevents ossification by middleboxes)
- broad adoption:
  - support from major web browsers
  - supported by 29.2% of all websites as of Mar 2024 (source: w3techs)

QUIC is (becoming) a huge success

- but TCP continues to be used (at least) as a fallback

Sad state of application protocols

We now have to develop and maintain two different set of stacks.

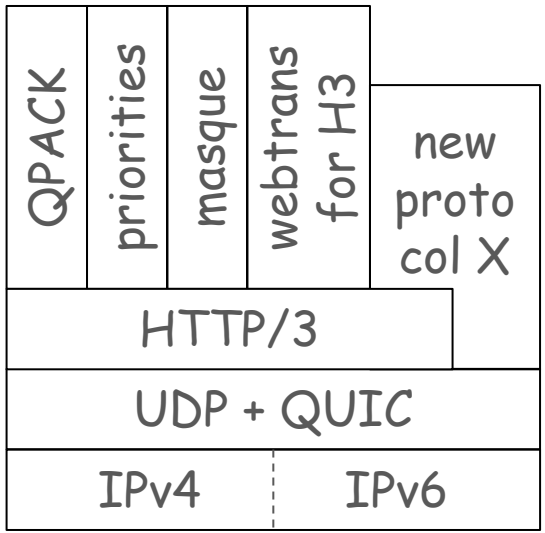

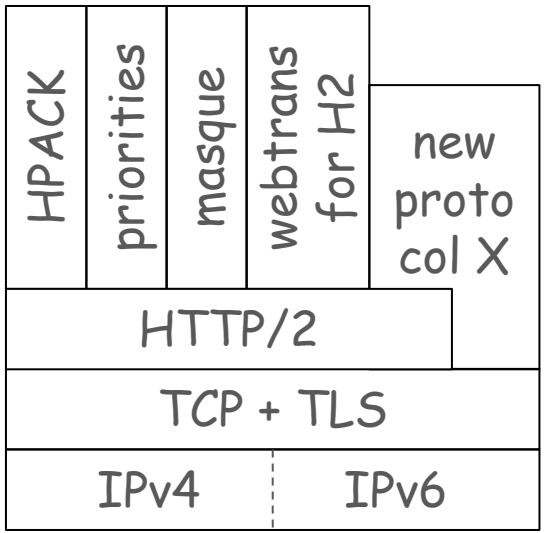

Sad state of application protocols

TCP (UDP) supports both IPv4 and IPv6.
Can QUIC support both UDP and TCP?

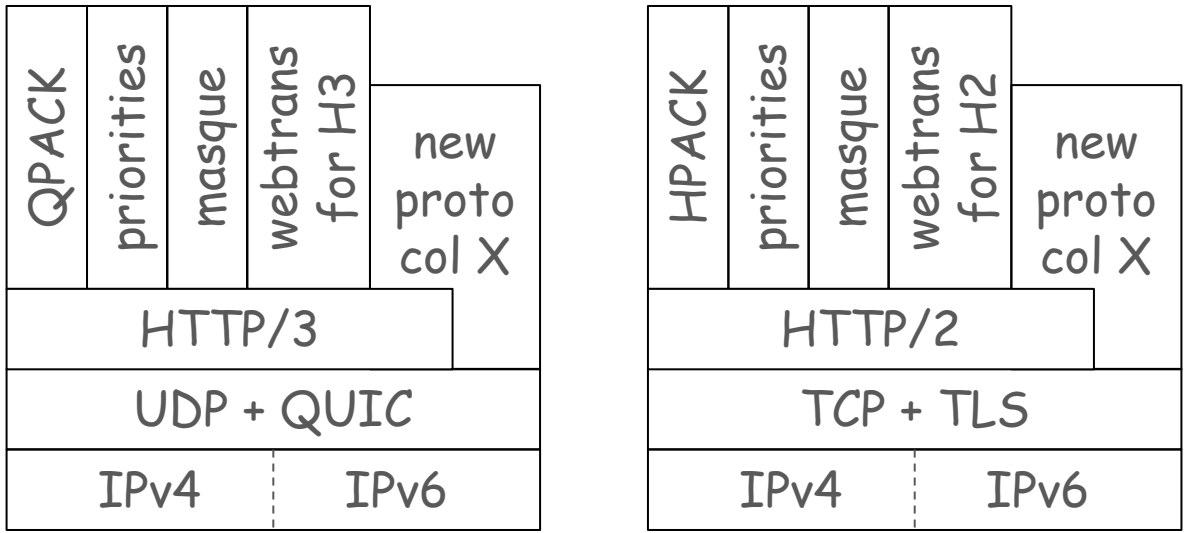

QUIC on Streams

Backport the QUIC API contract (i.e., QUIC streams) to TCP. Then, it is possible to run applications written for QUIC everywhere.

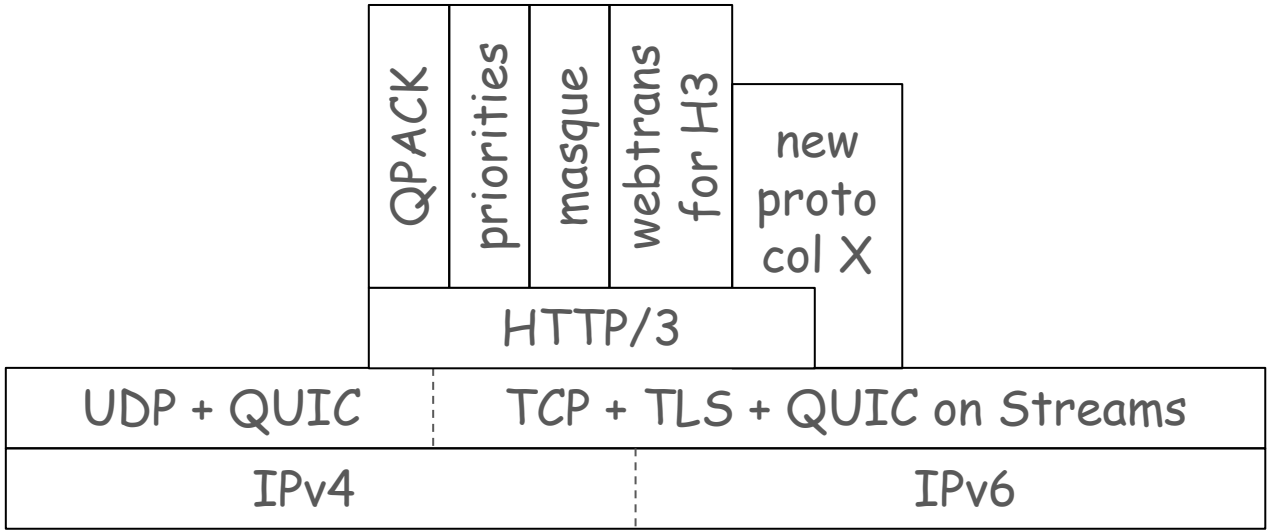

Our goals and non-goals

- Goals
  - eliminate the need to develop new things on top of two protocols
  - eliminate the need to deploy two different protocol stacks, when you control both sides
- Non-goals
  - do not spend time optimizing TCP (e.g., solve HoL blocking) or QUIC frames. QUIC works in most cases and performs better. QUIC on Streams is a fallback.

Design of draft -00

- send QUICv1 frames on top of TCP / TLS
- no ACK frames - all frames are implicitly ACKed
- use of frames unrelated to stream operation are prohibited
  - with datagram extension for message-oriented applications
- Transport Parameters are exchanged using the 1st frame called QS_TRANSPORT_PARAMETERS
- minimum of maximum frame size is 16KB (matches max. TLS record size)

※working PoC for quicly created in 1/2 day. Took another 1/2 day to integrate that into H2O to run H3 client / server over QUIC on Streams.

HTTP datagrams and MASQUE

- HTTP datagrams and capsule protocol - RFC 9297
  - Started as HTTP/3 only - draft-ietf-masque-h3-datagram-00
  - Changed due to non-trivial demand for datagrams over all HTTP versions
    - Despite downsides, unreliable messages over reliable transport
  - Datagrams over reliable streams using capsules
- UDP over HTTP - RFC 9298
- IP over HTTP - RFC 9484
- MASQUE applications can improve reachability by having a TCP fallback
  - QUIC on Streams supports optional datagrams
  - Cloudflare Internal customers have HTTP/3 MASQUE client and server already
    - Asking for QUIC on Streams to:
      - Improve time-to-market
      - Lower development and maintenance costs (by reducing cross-HTTP-version duplication)
      - Avoid runtime capsule encapsulation overhead