Multipath Extension for QUIC

draft-lmbdhk-quic-multipath-00

QUIC session @ IETF-112, Nov 10, 2021

Yanmei Liu, Yunfei Ma, Quentin De Coninck, Olivier Bonaventure,
Christian Huitema, **Mirja Kühlewind**

What happened so far

- QUIC interim meeting Oct 2020 on multipath QUIC use cases
  - https://datatracker.ietf.org/meeting/interim-2020-quic-02/session/quic
- QUIC side meeting Oct 18, 2021 on unifying the proposed QUIC extension
  - https://github.com/mirjak/draft-lmbdhk-quic-multipath/tree/master/presentations
- New draft submitted that unifies components of all three previous proposals:
  - draft-deconinck-quic-multipath-07
  - draft-liu-multipath-quic-04
  - draft-huitema-quic-mpath-option-01

draft-lmbdhk-quic-multipath-00

- Focus on core components
  - Negotiation
  - Path management (setup/closure)
  - Basic scheduling
  - Packet transmission and retransmission
- Other drafts may cover
  - Advanced Scheduling
  - Multipath extensions, such as
    - Unidirectional paths
    - Address discovery and selection

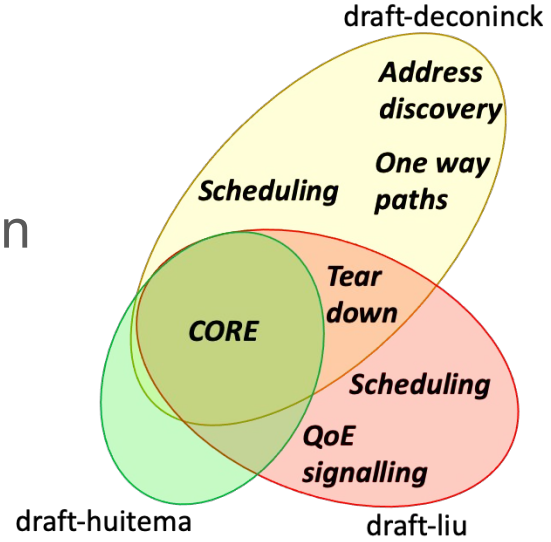

Design Principles

- Re-use as much as possible from RFC9000
  - Path validation is unchanged
  - Per-path congestion control
  - Header format is unchanged
  - Multipath usage only for 1-RTT packets
- Path is defined as 4-tuple (bidirectional)
  - At most one active path/CID per 4-tuple

Changes from RFC9000

- Replace “migration” by “simultaneous use”
  - Sending of non-probing frames on multiple paths
  - Additional signaling for removal of abandoned paths
- Additional considerations on
  - Efficient loss recovery and RTT estimation
  - ACKing and Packet Numbers (see next slides)

Handshake negotiation

New transport parameter: **`enable_multipath`**

Option	Definition
0x0	No multipath support
0x1	Only support for one PN space
0x2	Only support for multiple PN spaces
0x3	Support for both - multiple PN spaces is selected if both endpoint set 0x3

More evaluation and implementation experience needed to select on approach for final publication!

Use of one or more Packet Number (PN) spaces

Single PN Space

Pros

- Support of zero-length CID allows for minimal transmission overhead
- Implementation complexity: Fewer code changes
- Fewer crypto stack requirements: Does not require 96 bit nonce

Cons

- Potential increases ACK size, especially for paths with different latencies
- Higher complexity in packet scheduling and/or ACK logic

Multiple PN Spaces

Pros

- Smaller ACK ranges: Works well, even with large CWND
- No ambiguity about per-path packet loss and RTTs
- Simple logic: Per path version of RFC 9002 algorithms

Cons

- Currently requires use of CIDs in both directions
- More code changes needed

Path Management

Path Initiation

- New paths are only initiated by the client
- Use of RFC9000 path validation before non-probing packets can be sent

Path Removal

- New PATH_ABANDON frame indicates to peer that path should not be used anymore
- RETIRE_CONNECTION_ID frames indicates that resources can be released
- Idle timeout also causes path closure and removal of resources

Two new frame types

PATH_ABANDON

- Carries path identifier, error code, and reason phrase
- Three path identifier types to indicate either use of source or destination CID as identifier, or to refer to the current path used
- If CID(s) are used this frame can be sent over any path

ACK_MP (for use with multiple PN spaces only)

- Like ACK frames but additional packet number space identifier

Ready for working group adoption?

- Draft focus on core components only
- Agreement on design principles of all draft authors
- Negotiation option for PN space selection enables experimentation
- Side meeting has indicated interest and planned implementation work