

Path Selection for Multiple Paths In QUIC
draft-dawkins-quic-multipath-selection-01

Abstract

In QUIC working group discussions about proposals to use multiple paths, an obvious question came up - given multiple paths, how does QUIC select paths to send packets over?

The answer to that question may inform decisions in the QUIC working group about the scope of any multipath extensions considered for experimentation and adoption.

This document is intended to summarize aspects of path selection from those contributions and conversations.

It is recognized that path selection is not the only important open question about QUIC Multipath, but other open questions are out of scope for this document.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 4, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Why We Should Look at Path Selection Strategies Now . . .	3
1.2.	Notes for Readers	4
1.3.	Minimal Terminology	4
1.4.	Contribution and Discussion Venues for this draft. . . .	5
2.	Background for this document	5
3.	Overview of Proposed Path Selection Strategies	6
3.1.	Active-Standby	7
3.2.	Latency Versus Bandwidth	7
3.3.	Bandwidth Aggregation/Load Balancing	7
3.4.	Minimum RTT Difference	7
3.5.	Round-Trip-Time Thresholds	7
3.6.	RTT Equivalence	7
3.7.	Priority-based	7
3.8.	Redundant	8
3.9.	Control Plane Versus Data Plane	8
3.10.	Combinations of Strategies	8
4.	Implications for QUIC Multipath	8
4.1.	Selecting a Single Path Among Multiple Validated Paths ("Traffic Switching")	8
4.2.	Selecting Multiple Active Paths ("Traffic Splitting") . .	9
4.3.	Arbitrary Combinations	9
5.	Next Steps	10
6.	IANA Considerations	10
7.	Security Considerations	10
8.	Acknowledgments	10
9.	Informative References	10
	Author's Address	13

[1.](#) Introduction

In QUIC working group [[QUIC-charter](#)] discussions about proposals to use multiple paths, an obvious question came up - given multiple paths, how does QUIC select paths to send packets over?

Dawkins

Expires December 4, 2021

[Page 2]

The answer to that question may inform decisions in the QUIC working group about the scope of any multipath extensions considered for experimentation and adoption.

This document is intended to summarize aspects of path selection from those contributions and conversations.

It is recognized that path selection is not the only important open question about QUIC Multipath, but other open questions are out of scope for this document.

1.1. Why We Should Look at Path Selection Strategies Now

One of the first questions that's come up in discussions about multiple paths for QUIC has been about path selection. As soon as an implementation has multiple paths available, it must decide how to use these paths. The [\[RFC9000\]](#) answer, relying on connection migration, is "if you have multiple paths available, you can validate more than one connection at a time, but you only send on one connection at a time, and you migrate to another connection when you decide sending on the current connection is no longer appropriate. How you decide to migrate to another connection is up to you".

That has been a fine answer for many of the implementers who have worked on the first version of QUIC, and have deployed it in their networks. For other implementers, targeting other use cases and other networking environments, it may not be sufficient.

To take only one example, one of several presentations at [\[QUIC-interim-20-10\]](#) described aspects of 3GPP Access Traffic Steering, Switch and Splitting support (ATSSS), which contained four "Steering Modes" as part of Rel-16 in 2019 [\[TS23501\]](#), each of which corresponding roughly to a path selection strategy described in [Section 3](#) of this document. A study on "ATSSS Phase 2" [\[TR23700-93\]](#) included four more Steering Modes for Rel-17, expected to be finalized in mid-2021, and none of these eight (so far) Steering Modes are based on QUIC - they are based on Multipath TCP ([\[RFC8684\]](#)) or simple IP packet forwarding. And if that were not enough, a proposal for a study on "ATSSS Phase 3" [\[S2-2104599\]](#) was provided to the SA2 145-e meeting in May 2021. Some of the ATSSS strategies rely in 5G network internals and don't translate to the broader Internet, but most do translate, and 3GPP participants certainly aren't the only people thinking about path selection strategies.

Since the various proposals presented at [\[QUIC-interim-20-10\]](#) were developed in isolation from each other, the path selection strategies that they reflect may be similar between proposals, but not quite the

same, and none of the proposals presented had more than two strategies in common with any other proposal.

Given the number of path selection strategies being considered, implemented, and even deployed in any number of venues, and the potential for combinatorial explosion as described in [Section 3.10](#), it seems that identifying common aspects of path selection strategies, sooner rather than later, is important.

[1.2.](#) Notes for Readers

This document is an informational Internet-Draft, not adopted by any IETF working group, and does not carry any special status within the IETF.

Please note well that this document reflects the author's current understanding of past working group discussions and proposals. Contributions that add or improve considerations are welcomed, as described in [Section 1.4](#).

[1.3.](#) Minimal Terminology

In this document, "QUIC multipath" is only used as shorthand for "QUIC using multiple paths". It does not refer to a specific proposal.

In this document, "path selection strategy" means the policy that a QUIC sender uses to guide its choice between multiple interfaces of a QUIC connection for "the next packet".

This document adopts three terms, stolen from [\[TS23501\]](#), that seemed helpful in previous discussions about multipath in the QUIC working group.

- o Traffic Steering - selecting an initial path (in [\[RFC9000\]](#), this would be "validating a connection, and then using it". Although an [\[RFC9000\]](#) client can validate multiple connections, the client will only use one validated connection at a time.
- o Traffic Switching - selecting a different validated path (in [\[RFC9000\]](#), this is something like "migrating to a new validated connection", although whether connection migration as defined in [\[RFC9000\]](#)) would be sufficient is discussed in [Section 4](#)).
- o Traffic Splitting - using multiple validated paths simultaneously (this would almost certainly require an extension beyond connection migration as defined in [\[RFC9000\]](#)).

"Traffic Steering" does not require any extension to [RFC9000], and is not discussed further in this document. The focus will be on "Traffic Switching" and "Traffic Splitting".

1.4. Contribution and Discussion Venues for this draft.

(Note to RFC Editor - if this document ever reaches you, please remove this section)

This document is under development in the Github repository at <https://github.com/SpencerDawkins/quic-multipath-selection>.

Readers are invited to open issues and send pull requests with contributed text for this document, but since the document is intended to guide discussion for the QUIC working group, substantial discussion of this document should take place on the QUIC working group mailing list (quic@ietf.org). Mailing list subscription and archive details are at <https://www.ietf.org/mailman/listinfo/quic>.

2. Background for this document

A number of individual draft proposals for "QUIC over multiple paths" have been submitted to the IETF QUIC and INTAREA working groups, dating back as far as 2017. The author thinks that the complete list is as follows (and reminders for proposals he missed are welcomed):

- o [[I-D.an-multipath-quic](#)]
- o [[I-D.an-multipath-quic-application-policy](#)]
- o [[I-D.an-multipath-quic-traffic-distribution](#)]
- o [[I-D.chan-quic-owl](#)]
- o [[I-D.deconinck-quic-multipath](#)]
- o [[I-D.deconinck-quic-multipath](#)],
- o [[I-D.huitema-quic-mpath-req](#)]
- o [[I-D.huitema-quic-mpath-option](#)])
- o [[I-D.liu-multipath-quic](#)]
- o [[I-D.piraux-intarea-quic-tunnel-session](#)]

[I-D.bonaventure-iccrs-schedulers] has also been submitted to the Internet Congestion Control Research Group [[ICCRG-charter](#)] in the

Internet Research Task Force. It contains specific proposals for implementing some multipath schedulers, and includes some discussion of path selection relevant to this document.

One point of confusion in QUIC working group discussions was that the various proposals (also using Multipath TCP [[RFC8684](#)], so not all proposals were QUIC-specific) discussed in working group meetings and on the QUIC mailing list were from various proponents who weren't solving the same problem. This meant that no two of the use cases presented at the QUIC working group virtual interim on QUIC Multipath [[QUIC-interim-20-10](#)] were relying on the same strategies.

It seemed useful to collect the path selection strategies described in those proposals, to look for common elements, and to write them down in one place, to allow more focused discussion. [[I-D.dawkins-quic-what-to-do-with-multipath](#)] was intended to summarize, at a high level, various proposals for the use of multipath capabilities in QUIC, both inside the IETF and outside the IETF, in order to identify elements that were common across proposals. This draft tries to describe the impact of these various strategies on potential QUIC Multipath extensions.

One element that is certainly worth considering is whether the path selection strategies that have been proposed can be satisfied using a small number of "building block" strategies.

3. Overview of Proposed Path Selection Strategies

The following strategies were discussed at [[QUIC-interim-20-10](#)], and afterwards on the QUIC mailing list. These are summarized in this section, are described in more detail in [[I-D.dawkins-quic-what-to-do-with-multipath](#)], and are attributed to various proposals in that document.

- o Active-Standby - described in [Section 3.1](#)
- o Latency Versus Bandwidth - described in [Section 3.2](#)
- o Bandwidth Aggregation/Load Balancing - described in [Section 3.3](#)
- o Minimum RTT Difference - described in [Section 3.4](#)
- o Round-Trip-Time Thresholds - described in [Section 3.5](#)
- o RTT Equivalence - described in [Section 3.6](#)
- o Priority-based - described in [Section 3.7](#)

- o Redundant - described in [Section 3.8](#)
- o Control Plane Versus Data Plane - described in [Section 3.9](#)
- o Combinations of Strategies - described in [Section 3.10](#)

[3.1.](#) Active-Standby

The traffic associated with a specific flow will be sent via a specific path (the 'active path') and switched to another path (called 'standby path') when the active access is unavailable.

[3.2.](#) Latency Versus Bandwidth

Some traffic might be sent over a network path with lower latency and other traffic might be sent over a different network path with higher bandwidth.

[3.3.](#) Bandwidth Aggregation/Load Balancing

Traffic is sent using all available paths simultaneously, so that all available bandwidth is utilized, likely based on something like weighted round-robin path selection. This strategy is often used for bulk transfers.

[3.4.](#) Minimum RTT Difference

Traffic is sent over the path with the lowest smoothed RTT among all available paths, in order to minimize latency for latency-sensitive flows.

[3.5.](#) Round-Trip-Time Thresholds

Traffic is sent over the first path with a smoothed RTT that meets a certain threshold.

[3.6.](#) RTT Equivalence

When multiple paths each have sufficiently similar smoothed RTTs, traffic is sent over all of these paths. This is similar to "Bandwidth Aggregation/Load Balancing", with the additional qualification that not all available paths are used for this traffic.

[3.7.](#) Priority-based

Priorities are assigned to each path (often by association with network interfaces). Traffic is sent on a highest-priority path

until it becomes congested, and then "overflows" onto a lower-priority path.

[3.8.](#) Redundant

Traffic is replicated over two or more paths. This strategy could be used continuously, but is more commonly used when measured network conditions indicate that redundant sending may be necessary to increase the likelihood that at least one copy of each packet will arrive at the receiver.

[3.9.](#) Control Plane Versus Data Plane

An application might stream media over one or more available paths (based on one of the other strategies named in this section), but might send ACK traffic or retransmission over a path specifically chosen for that purpose. This is more likely to be beneficial if the path characteristics differ significantly between available paths - for example, satellite uplink/downlink stations connected by both higher-bandwidth Low Earth Orbit satellite paths and lower-bandwidth cellular or landline paths.

[3.10.](#) Combinations of Strategies

In addition to the strategies described above, it is also possible to combine these strategies. For example, a scheduler might use load-balancing over three paths, but when one of the paths becomes unavailable, the scheduler might switch to the two paths that are still available, in a way similar to Active-Standby. This is very much an example chosen at random - potentially, there are many combinations that could be useful.

[4.](#) Implications for QUIC Multipath

This section summarizes potential implications for "Multipath QUIC" of path selection strategies described in [Section 3](#), dividing them between "Traffic Switching" ([Section 4.1](#)) and "Traffic Splitting" ([Section 4.2](#)).

[4.1.](#) Selecting a Single Path Among Multiple Validated Paths ("Traffic Switching")

If a sender using Active-Standby (described in [Section 3.1](#)) does not perform frequent path switching, it can likely be supported using connection migration as defined in [\[RFC9000\]](#) without change.

- o The caveat here is that connection migration can include the also-implicit assumption that an endpoint can free up resources

associated with the previously-active path. If connection migration happens often enough, the endpoint may spend considerable time "thrashing" between allocating resources and quickly freeing them. Of course, if a sender is frequently selecting a new path for connection migration, this probably degenerates into one of the other path selection strategies.

Some path selection strategies could be supported by a mechanism as simple as the one proposed in [[I-D.huitema-quic-mpath-option](#)], which replaces "the implicit signaling of path migration through data transmission, by means of a new PATH_OPTION frame" (this isn't intended to imply the proposal is simple, only the explicit signaling), if the receiver uses this option to notify the sender of the preferred path. For example, Minimum RTT Difference (described in [Section 3.4](#)) and Round-Trip-Time Thresholds (described in [Section 3.5](#)) likely fall into this category.

Some path selection strategies are exploiting a relatively long-lived difference between paths - for example, Latency Versus Bandwidth (described in [Section 3.2](#)), Priority-based (described in [Section 3.7](#)), and Control Plane Versus Data Plane (described in [Section 3.9](#)) may fall into this category. One might wonder why these senders would need to use a single "multipath connection", rather than multiple [[RFC9000](#)] connections, for these cases, but if there is a reason to use a single multipath connection, a mechanism similar to the one proposed in [[I-D.huitema-quic-mpath-option](#)] could also be used in these cases.

[4.2.](#) Selecting Multiple Active Paths ("Traffic Splitting")

Some path selection strategies are treating more than one path as a set of active paths, whether the sender is performing "Traffic Splitting" (as defined in [Section 1.3](#)), as is the case for Bandwidth Aggregation/Load Balancing (described in [Section 3.3](#)) and RTT Equivalence (described in [Section 3.6](#)), or simply transmitting the same packet across multiple paths, as is the case for Redundant (described in [Section 3.8](#)).

For these cases, a more complex mechanism is likely required.

[4.3.](#) Arbitrary Combinations

Because it is simple enough to imagine various combinations of strategies (as described in [Section 3.10](#)), it seems important to understand what basic building blocks are required in order to support the strategies that seem common across a variety of use cases, because interactions between strategies may have significant

implications for QUIC Multipath that might not arise when considering strategies in isolation.

This seems especially important because existing proposals for QUIC Multipath don't use the same vocabulary to describe path selection strategies, so implementations may not behave in the same way, even if they are each using a strategy that seems to be common.

5. Next Steps

If this discussion is useful, it may also be useful to take the next step, and identify potential building blocks that can be used to construct the path selection strategies described in [Section 4.1](#) and [Section 4.2](#).

6. IANA Considerations

This document does not make any request to IANA.

7. Security Considerations

QUIC-specific security considerations are discussed in [Section 21 of \[RFC9000\]](#).

Section 6 of [\[I-D.ietf-quic-datagram\]](#) discusses security considerations specific to the use of the Unreliable Datagram Extension to QUIC.

Some "Multipath QUIC"-specific security considerations can be found in the corresponding section of [\[I-D.deconinck-quic-multipath\]](#).

Having said that, it may be best to repeat the security considerations section from [\[I-D.huitema-quic-mpath-option\]](#): "TBD. There are probably ways to abuse this."

8. Acknowledgments

Your name could appear here. Please comment and contribute, as per [Section 1.4](#).

9. Informative References

[I-D.an-multipath-quic]

An, Q., Liu, Y., Ma, Y., and Z. Li, "Multipath Extension for QUIC", [draft-an-multipath-quic-00](#) (work in progress), October 2020.

[I-D.an-multipath-quic-application-policy]

An, Q., Li, Z., and Y. Liu, "Enabling application policy-awareness in Multipath QUIC", [draft-an-multipath-quic-application-policy-00](#) (work in progress), July 2020.

[I-D.an-multipath-quic-traffic-distribution]

An, Q., Liu, D., and Y. Liu, "Key Components for Multipath QUIC Traffic Distribution", [draft-an-multipath-quic-traffic-distribution-00](#) (work in progress), March 2020.

[I-D.bonaventure-iccrq-schedulers]

Bonaventure, O., Piraux, M., Coninck, Q. D., Baerts, M., Paasch, C., and M. Amend, "Multipath schedulers", [draft-bonaventure-iccrq-schedulers-01](#) (work in progress), September 2020.

[I-D.chan-quic-owl]

Chan, H. A., Wei, A., Song, F., and H. Zhang, "One Way Latency Considerations for Multipath in QUIC", [draft-chan-quic-owl-01](#) (work in progress), March 2017.

[I-D.dawkins-quic-what-to-do-with-multipath]

Dawkins, S., "What To Do With Multiple Active Paths in QUIC", [draft-dawkins-quic-what-to-do-with-multipath-03](#) (work in progress), January 2021.

[I-D.deconinck-quic-multipath]

Coninck, Q. D. and O. Bonaventure, "Multipath Extensions for QUIC (MP-QUIC)", [draft-deconinck-quic-multipath-07](#) (work in progress), May 2021.

[I-D.huitema-quic-mpath-option]

Huitema, C., "QUIC Multipath Negotiation Option", [draft-huitema-quic-mpath-option-00](#) (work in progress), October 2020.

[I-D.huitema-quic-mpath-req]

Huitema, C., "QUIC Multipath Requirements", [draft-huitema-quic-mpath-req-01](#) (work in progress), January 2018.

[I-D.ietf-quic-datagram]

Pauly, T., Kinnear, E., and D. Schinazi, "An Unreliable Datagram Extension to QUIC", [draft-ietf-quic-datagram-02](#) (work in progress), February 2021.

[I-D.liu-multipath-quic]

Liu, Y., Ma, Y., Huitema, C., An, Q., and Z. Li,
"Multipath Extension for QUIC", [draft-liu-multipath-quic-03](#) (work in progress), March 2021.

[I-D.piraux-intarea-quic-tunnel-session]

Piraux, M., Bonaventure, O., and A. Masputra, "Session mode for multiple QUIC Tunnels", [draft-piraux-intarea-quic-tunnel-session-00](#) (work in progress), November 2020.

[ICCRG-charter]

"IRTF Internet Congestion Control Research Group Charter", n.d., <<https://datatracker.ietf.org/rg/iccrq/about/>>.

[QUIC-charter]

"IETF QUIC Working Group Charter", n.d., <<https://datatracker.ietf.org/wg/quic/about/>>.

[QUIC-interim-20-10]

"IETF QUIC Working Group Virtual Interim Meeting - October 2020", October 2020, <<https://github.com/quicwg/wg-materials/tree/master/interim-20-10>>.

[RFC8684] Ford, A., Raiciu, C., Handley, M., Bonaventure, O., and C. Paasch, "TCP Extensions for Multipath Operation with Multiple Addresses", [RFC 8684](#), DOI 10.17487/RFC8684, March 2020, <<https://www.rfc-editor.org/info/rfc8684>>.

[RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", [RFC 9000](#), DOI 10.17487/RFC9000, May 2021, <<https://www.rfc-editor.org/info/rfc9000>>.

[S2-2104599]

Lenovo, Motorola Mobility, ., "Study on Access Traffic Steering, Switching and Splitting support in the 5G system architecture; Phase 3", 2021, <https://www.3gpp.org/ftp/tsg_sa/WG2_Arch/TSGS2_145E_Electronic_2021-05/Docs/S2-2104599.zip>.

[TR23700-93]

3GPP (3rd Generation Partnership Project), ., "Technical Specification Group Services and System Aspects; Study on access traffic steering, switch and splitting support in the 5G System (5GS) architecture; Phase 2 (Release 17)", 2021, <https://www.3gpp.org/ftp/Specs/archive/23_series/23.700-93/>.

[TS23501] 3GPP (3rd Generation Partnership Project), ., "Technical Specification Group Services and System Aspects; System Architecture for the 5G System; Stage 2 (Release 16)", 2019, <https://www.3gpp.org/ftp/Specs/archive/23_series/23.501/>.

Author's Address

Spencer Dawkins
Tencent America LLC

Email: spencerdawkins.ietf@gmail.com