

QUIC Working Group
Internet-Draft
Intended status: Informational
Expires: July 10, 2021

S. Dawkins, Ed.
Tencent America
January 06, 2021

What To Do With Multiple Active Paths in QUIC
draft-dawkins-quic-what-to-do-with-multipath-03

Abstract

The IETF QUIC working group has been chartered to produce extensions that would "enable ... multipath capabilities" since the working group was formed in 2016, but because multipath was an extension, work on multipath, and the other extensions named in the charter, waited while work proceeded on the core QUIC protocol specifications.

After the QUIC working group chairs requested publication for the core QUIC protocol specifications, they scheduled a virtual interim meeting to understand the use cases that various groups inside and outside the IETF were envisioning for multipath with QUIC.

As part of that discussion, it became obvious that people had a variety of ideas about how multiple paths would be used, because they weren't looking at the same use cases.

This document is intended to capture that variety of ideas, to inform further discussion in the working group.

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 July 10, 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	3
1.1.	Other Voices	3
1.2.	Notes for Readers	4
1.3.	Contribution and Discussion Venues for this draft. . . .	4
2.	Multipath Modes of Operation	4
2.1.	Reference Architecture	4
2.2.	Sending When Multiple Paths Are Available	6
2.2.1.	Traffic Steering	6
2.2.2.	Traffic Switching	6
2.2.3.	Traffic Splitting	7
2.2.4.	Traffic Steering, Traffic Switching, and Traffic Splitting in the Core QUIC Transport Protocol and Datagram Extension	7
2.3.	Sending Strategies with Multiple Paths	8
2.3.1.	Active-Standby	8
2.3.2.	Trading Latency Versus Bandwidth	8
2.3.3.	Bandwidth Aggregation/Load Balancing	8
2.3.4.	Round-Trip-Time Thresholds, Minimum RTT Difference, and RTT Equivalence	9
2.3.5.	Priority-based	9
2.3.6.	Redundant	9
2.3.7.	Combinations of Strategies	10
2.4.	Unified Field Theory of Path Selection	10
3.	IANA Considerations	10
4.	Security Considerations	10
5.	Acknowledgements	11
6.	Document History	12
7.	Informative References	12
	Author's Address	14

Dawkins

Expires July 10, 2021

[Page 2]

1. Introduction

The IETF QUIC working group has been chartered to produce extensions that would "enable ... multipath capabilities" ([[QUIC-charter](#)]) since the working group was formed in 2016, but because multipath was an extension, work on multipath, and the other extensions named in the charter, waited while work proceeded on the core QUIC protocol specifications ([[I-D.ietf-quic-transport](#)] and related specifications).

After the QUIC working group chairs requested publication for the core QUIC protocol specifications, they scheduled a virtual interim meeting ([[QUIC-interim-20-10](#)]) to understand the use cases that various groups inside and outside the IETF were envisioning for multipath with QUIC.

As part of that discussion, it became obvious that people had a variety of ideas about how multiple paths would be used, because they weren't looking at the same use cases. From an architectural point of view, two main multipath scenarios emerged. The first scenario was envisioned by people who were focused on end-to-end service, and the second scenario was envisioned by people who expected some sort of (explicit) intermediary, such as a transport proxy or tunnel endpoint, between the ultimate endpoints.

This document is intended to capture that variety of ideas, to inform further discussion in the working group.

For more details, please see the QUIC working group meeting minutes, at [[QUIC-interim-20-10-minutes](#)], which include pointers to the video recording from the meeting, pointers to each presentation given, questions and answers after each presentation, and overall discussion following the presentations.

1.1. Other Voices

This document grew out of discussions in the QUIC working group, but there are other efforts working on multipath in both the IETF and the IRTF, and their work includes mentions of strategies for using multiple paths. I'll add this work as I find it, if that work is proposing a strategy that hasn't been mentioned in the QUIC discussion. In particular, strategies named in [[I-D.bonaventure-iccr-g-schedulers](#)] have been included in [Section 2.3](#), in the current version of the draft.

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.

It reflects the author's understanding, and contributions that include improvements and clarifications are welcomed, as described in [Section 1.3](#).

1.3. 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/draft-dawkins-quic-what-to-do-with-multipath>. Readers are invited to open issues and send pull requests with contributed text for this document.

Substantial discussion of this document should take place on the QUIC working group mailing list (quic@ietf.org). Subscription and archive details are at <https://www.ietf.org/mailman/listinfo/quic>.

2. Multipath Modes of Operation

For the purposes of this document, it's enough to consider two QUIC endpoints (QUIC Host X and QUIC Host Y) with two disjoint paths between them (Access Net A and Access Net B), as shown in Figure 1.

2.1. Reference Architecture




Figure 1: Simplified Reference Architecture for QUIC Multipath Operation

Some use cases limit themselves to two paths - for example, the use case described in [[I-D.bonaventure-quick-atsss-overview](#)] is limited to one 5G 3GPP access network and one non-3GPP access network. More than two paths are certainly possible - for example, two hosts that are connected by a satellite path, and two different landline paths from different network providers for redundancy. But I believe this simplified architecture is sufficient for discussion purposes.

Please note: as mentioned in [Section 1](#), some readers will be thinking of an end-to-end QUIC connection between QUIC Host X and QUIC Host Y in Figure 1, while other readers will be thinking of QUIC Host Y as some sort of (explicit) transport proxy or tunnel endpoint, e.g. [[I-D.bonaventure-quick-atsss-overview](#)]. This document is only focused on the multiple paths between QUIC Host X and QUIC Host Y, in order to talk about classes of sending and strategies for sending over multiple paths. However, it might be that further implications arose from the motivation of end-to-end multipath or intermediary multipath. In particular the latter will additionally require, according to [[I-D.bonaventure-quick-atsss-overview](#)], QUIC datagram (unreliable transport) and tunneling support. While multipath traffic distribution strategies play a central role in multipath transport, multipath re-ordering strategies become additionally important for multipath over unreliable network protocols. Possible strategies to solve this are described in [[I-D.amend-icrg-multipath-reordering](#)]. Another complexity, not in scope of this document, is the interaction between multipath scheduling and QUIC stream traffic scheduling.

2.2. Sending When Multiple Paths Are Available

This section uses terminology from 3GPP ATSSS {Access Traffic Steering, Switch and Splitting, [[TS23501](#)]} to describe three different classes of packet sending that could be used when multiple disjoint paths are available. Note that these terms describe concepts that are not ATSSS-specific, and could apply to any multipath environment. If terminology more accessible to IETF QUIC working group participants presents itself, I'll change it.

Please note that the relationship between "flows", "packets", and "datagrams" isn't entirely clear yet. In [[TS23501](#)], a "flow" is defined as "a set of packets that share an ordering constraint", which maps reasonably well onto TCP bytestreams and correspondingly onto MP-TCP bytestreams, but the core QUIC transport mechanism [[I-D.ietf-quic-transport](#)] provides multiple bytestreams in a single connection, and leveraging QUIC streams seems to be a use case enabler, whether for redundantly sending overlapping ranges, for efficiently striping non-overlapping ranges for bandwidth, or for strict ordering.

Further, with the QUIC datagram extension [[I-D.ietf-quic-datagram](#)], DATAGRAM frames can be carried in a QUIC connection, but are not part of a stream and do not interact with in-order delivery within a stream.

2.2.1. Traffic Steering

When a sending host has a new "flow", and has more than one path available, a sending host will select a path to begin sending packets on. This can be described as Traffic Steering.

Note that the selection of a path is "per-flow". Different flows may be steered onto different paths, but each flow will be sent on a single path.

The choice of a path can be based on a number of factors - for instance, whether the path will be used for bulk transfer, or whether an unmetered path is available. Those factors are important, but for this document, the point is that a path is selected.

2.2.2. Traffic Switching

When a sending host has more than one path available and is sending a flow on a single path, the sending host can stop using this path for this flow, and start using another path. This can be described as Traffic Switching.

Note that the selection of another path is also per-flow. A sender can switch a flow onto a different path, without switching any other flows to that path.

2.2.3. Traffic Splitting

When a sending host has more than one path available, the sending host may wish to use multiple paths simultaneously for a single flow. This can be described as Traffic Splitting.

Note that splitting a flow onto multiple paths means that the sending decision would be made on a packet-by-packet basis. No guarantee of packet ordering across paths is assumed.

Note that splitting a flow onto multiple paths, packet by packet, means that any ordering constraint must be provided by an upper layer mechanism. In [TS23501], Multipath TCP [RFC8684] is used to provide that ordering constraint.

2.2.4. Traffic Steering, Traffic Switching, and Traffic Splitting in the Core QUIC Transport Protocol and Datagram Extension

[I-D.ietf-quic-transport] and related specifications support the selection and use of multiple paths for a connection between QUIC endpoints, as long as only one path is in active use at any given time. [I-D.ietf-quic-transport] and related specifications do not support sending traffic for a single connection on multiple paths between QUIC endpoints simultaneously.

This level of support for multipath allows Traffic Steering - the selection of a specific path - and Traffic Switching - migration of traffic from one path to another, using QUIC connection migration. It does not allow "Traffic Splitting" - the use of simultaneous paths for a single connection.

One important output from the QUIC interim meeting was the recognition that some participants view "Traffic Switching" as something like "protection switching" from an active path to a standby path, that doesn't happen often, while other participants view QUIC connection migration as a way of responding frequently to changing path conditions. This will be an important part of the conversation about multipath QUIC.

If a QUIC endpoint opens and validates two or more connections, no additional setup or signaling is required for a sender to switch paths - that can begin with the next packet queued for transmission. This would allow "Traffic Splitting" for unordered QUIC DATAGRAMS

[[I-D.ietf-quic-datagram](#)], with an upper-layer mechanism providing flow ordering if that is needed.

2.3. Sending Strategies with Multiple Paths

During the QUIC virtual interim meeting, various presenters described various sending strategies that they have used in the past (especially with TCP [[RFC0793](#)] and Multipath TCP [[RFC8684](#)]). and that they expected to use in the future with some multipath QUIC extension.

This section is an attempt to identify commonalities and gaps that may inform work on future multipath extensions for QUIC.

All of these sending strategies build on Traffic Steering - opening one or more connections for a new flow, and then proceeding to use one connection, switch to another connection, or split traffic across paths. For this reason, Traffic Steering isn't mentioned in the rest of this section.

2.3.1. 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. Traffic Switching is sufficient to support this strategy. [[Alibaba-MPQUIC](#)] and [[Multipath-3GPP-ATSSS](#)] mentioned this strategy at the QUIC Interim meeting.

2.3.2. Trading Latency Versus Bandwidth

Some applications require low latency, and others don't. So using a network path with lower latency for some connections, and using a different network path with higher bandwidth available for others. allows this strategy. Traffic Switching is sufficient for this, if measured path characteristics change enough to justify changing paths. [[Chromium-Multipath](#)] mentioned this strategy at the QUIC Interim Meeting. [[Apple-Multipath](#)] and [[Sat-Terr-Multipath](#)] are actually switching paths at sub-RTT rates.

2.3.3. Bandwidth Aggregation/Load Balancing

Some environments allow the use of multiple paths simultaneously for significant periods of time, because none of the paths are metered, so using all of the bandwidth on all of the paths simultaneously makes sense for bulk transfers. Traffic Splitting is required to support this strategy. [[Alibaba-MPQUIC](#)], [[Hybrid-Access-Multipath](#)], [[Sat-Terr-Multipath](#)], and [[Multipath-3GPP-ATSSS](#)] all mentioned this

strategy at the QUIC Interim Meeting.

[[I-D.bonaventure-iccrq-schedulers](#)] mentions this strategy as "Round-Robin" (not recommended) and "Weighted Round-Robin".

2.3.4. Round-Trip-Time Thresholds, Minimum RTT Difference, and RTT Equivalence

Not all multipath environments include paths with significant differences between path characteristics, but several do.

A sender might select the first available path with a smoothed round-trip-time below a certain threshold. The goal is to keep the RTT of the multipath connection to a small value and avoid having the whole connection impacted by "bad" paths.

[[I-D.bonaventure-iccrq-schedulers](#)] mentions this strategy as "Round-Trip-Time Threshold".

A sender might select the path with the lowest RTT among all available paths, in order to minimize latency for latency-sensitive flows. [[I-D.bonaventure-iccrq-schedulers](#)] mentions this strategy as "Lowest Round-Trip-Time First".

If two or more paths have characteristics that are "sufficiently close" - for example, they may have path RTTs that are equivalent, from the application's point of view - a sender may wish to use both paths as long as that's the case, and then adopt another sending strategy when measured path characteristics diverge. Traffic Splitting is required to support this strategy. [[Apple-Multipath](#)], [[Sat-Terr-Multipath](#)], [[Hybrid-Access-Multipath](#)], and [[Multipath-3GPP-ATSSS](#)] all mentioned this strategy at the QUIC Interim Meeting.

2.3.5. Priority-based

Paths are assigned priority levels that indicate which path will be used first. Concretely, the traffic associated with the matching flow will be steered on the path with a high priority till congestion is detected, then the overflow will be sent over a low priority access. Traffic Splitting is required to support this strategy. [[Multipath-3GPP-ATSSS](#)] mentioned this strategy at the QUIC Interim meeting.

2.3.6. Redundant

Traffic is replicated over two or more paths to increase the probability that the traffic will arrive at the other host undamaged and usable. Redundant traffic may be used for all packets in a connection, or may only be used when measured network conditions

indicate it should be used. Traffic Splitting is required to support this strategy. [[Apple-Multipath](#)] and [[Multipath-3GPP-ATSSS](#)] mentioned this strategy at the QUIC Interim meeting, and Yunfei Ma confirmed in e-mail that [[Alibaba-MPQUIC](#)] also uses this strategy.

2.3.7. Combinations of Strategies

As [[I-D.bonaventure-iccrq-schedulers](#)] points out, a scheduler can use a combination of strategies in sending decisions. For example, a scheduler might use load-balancing over three paths, but when one of the paths becomes unavailable, might switch to the two paths that are still available.

As Lucas Pardue pointed out in e-mail, some strategies may differentiate between "control plane" and "data plane" path selection. For example, an application might send stream data over one or more paths, but might send ACK traffic or retransmissions over a specific path, especially if the path characteristics vary significantly, as would be the case for satellite uplink/downlink stations connected by both higher-bandwidth Low Earth Orbit satellite paths and lower-bandwidth cellular or landline paths.

2.4. Unified Field Theory of Path Selection

The strategies in [Section 2.3](#) don't include all possible strategies, and it's worth noting that few of the presentations given at [[QUIC-interim-20-10](#)] shared a common set of even two or three strategies. Even [[Multipath-3GPP-ATSSS](#)] omitted some strategies under consideration in 3GPP because they were closely tied to 5G network architecture components.

It will be very helpful if we can look for a small number of simple concepts that would allow a small number of schedulers to meet most application requirements. The alternative - adding schedulers every time someone thinks of a new strategy - is too painful to contemplate.

3. IANA Considerations

This document does not make any request to IANA.

4. Security Considerations

QUIC-specific security considerations are discussed in Section 21 of [[I-D.ietf-quic-transport](#)].

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 [Section 8](#) of the individual draft [[I-D.deconinck-quic-multipath](#)].

5. Acknowledgements

I'd like to thank Lars Eggert and Lucas Pardue, the QUIC working group chairs, who called the QUIC virtual interim meeting on multipath.

I'd also like to thank the presenters at the QUIC virtual interim, who put together valuable presentations on short notice.

Thanks to David Allan, for comments resulting in the clarification of Traffic Steering, Traffic Switching, and Traffic Splitting when these concepts are used in contexts outside 3GPP.

Thanks again to Lucas, for providing comments resulting in several clarifications on the mapping of Traffic Splitting onto QUIC protocol operation, and for providing comments about control plane/data plane considerations.

Thanks to Markus Amend, for providing contributions and spotting a whopper of a typo.

Thanks to Yunfei Ma for reviewing my characterization of [[Alibaba-MPQUIC](#)].

Thanks to Joerg Deutschmann for reviewing my characterization of [[Sat-Terr-Multipath](#)].

Thanks to Christian Huitema for suggesting that we try to simplify the large and growing number of strategies described in [Section 2.3](#) into a few simple concepts. This question should come up in QUIC working group discussion of multipath, and will be added to [[I-D.dawkins-quic-multipath-questions](#)] so the author doesn't forget that.

Many thanks to (your name could easily appear here) for reviews and comments.

6. Document History

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

Version -00: initial submission

Version -01: updated for IETF 109 QUIC working group session

Version -02: Updated to add details about end-to-end and tunneling usages

7. Informative References

[Alibaba-MPQUIC]

Yanmei Lei / Yunfei Ma, ., "MPQIIC Use Cases", October 2020, <<https://github.com/quicwg/wg-materials/blob/master/interim-20-10/MPQUIC%20use%20cases.pdf>>.

[Apple-Multipath]

Christoph Paasch, ., "Multipath transports at Apple", October 2020, <<https://github.com/quicwg/wg-materials/blob/master/interim-20-10/Multipath%20transports%20at%20Apple.pdf>>.

[Chromium-Multipath]

Fan Yang/Jana Iyengar, ., "Multipath in Chromium", October 2020, <<https://github.com/quicwg/wg-materials/blob/master/interim-20-10/Multipath%20in%20Chromium.pdf>>.

[Hybrid-Access-Multipath]

Olivier Bonaventure, ., "Hybrid access networks and requirements on MPQUIC", October 2020, <<https://github.com/quicwg/wg-materials/blob/master/interim-20-10/Hybrid%20access%20networks%20and%20requirements%20on%20MPQUIC.pdf>>.

[I-D.amend-iccr-g-multipath-reordering]

Amend, M. and D. Hugo, "Multipath sequence maintenance", [draft-amend-iccr-g-multipath-reordering-01](#) (work in progress), November 2020.

[I-D.bonaventure-iccr-g-schedulers]

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

[I-D.bonaventure-quic-atsss-overview]

Boucadair, M., Bonaventure, O., Piraux, M., Coninck, Q., Dawkins, S., Kuehlewind, M., Amend, M., Kassler, A., An, Q., Keukeleire, N., and S. Seo, "3GPP Access Traffic Steering Switching and Splitting (ATSSS) - Overview for IETF Participants", [draft-bonaventure-quic-atsss-overview-00](#) (work in progress), May 2020.

[I-D.dawkins-quic-multipath-questions]

Dawkins, S., "Questions for Multiple Paths In QUIC", [draft-dawkins-quic-multipath-questions-00](#) (work in progress), December 2020.

[I-D.deconinck-quic-multipath]

Coninck, Q. and O. Bonaventure, "Multipath Extensions for QUIC (MP-QUIC)", [draft-deconinck-quic-multipath-06](#) (work in progress), November 2020.

[I-D.ietf-quic-datagram]

Pauly, T., Kinnear, E., and D. Schinazi, "An Unreliable Datagram Extension to QUIC", [draft-ietf-quic-datagram-01](#) (work in progress), August 2020.

[I-D.ietf-quic-transport]

Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", [draft-ietf-quic-transport-33](#) (work in progress), December 2020.

[Multipath-3GPP-ATSSS]

Spencer Dawkins / Mirja Kuehlewind / Florin Baboescu / Hannu Flinck, ., "Multipath in 3GPP ATSSS", October 2020, <<https://github.com/quicwg/wg-materials/blob/master/interim-20-10/Multipath%20in%203GPP%20ATSSS.pdf>>.

[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>>.

[QUIC-interim-20-10-minutes]

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

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.
- [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>>.
- [Sat-Terr-Multipath]
Joerg Deutschmann, ., "Satellite/Terrestrial Multipath Communication", October 2020, <<https://github.com/quicwg/wg-materials/blob/master/interim-20-10/Satellite-terrestrial%20multipath%20communication.pdf>>.
- [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 (editor)
Tencent America

Email: spencerdawkins.ietf@gmail.com

