
Stream: Internet Research Task Force (IRTF)
RFC: [9473](#)
Category: Informational
Published: September 2023
ISSN: 2070-1721
Authors: R. Enghardt C. Krähenbühl
Netflix *ETH Zürich*

RFC 9473

A Vocabulary of Path Properties

Abstract

Path properties express information about paths across a network and the services provided via such paths. In a path-aware network, path properties may be fully or partially available to entities such as endpoints. This document defines and categorizes path properties. Furthermore, the document identifies several path properties that might be useful to endpoints or other entities, e.g., for selecting between paths or for invoking some of the provided services. This document is a product of the Path Aware Networking Research Group (PANRG).

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the consensus of the Path Aware Networking Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9473>.

Copyright Notice

Copyright (c) 2023 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.

Table of Contents

| | |
|--|----|
| 1. Introduction | 2 |
| 2. Terminology | 3 |
| 2.1. Terminology Usage for Specific Technologies | 5 |
| 3. Use Cases for Path Properties | 5 |
| 3.1. Path Selection | 5 |
| 3.2. Protocol Selection | 6 |
| 3.3. Service Invocation | 7 |
| 4. Examples of Path Properties | 7 |
| 5. Security Considerations | 10 |
| 6. IANA Considerations | 11 |
| 7. Informative References | 11 |
| Acknowledgments | 13 |
| Authors' Addresses | 13 |

1. Introduction

The current Internet architecture does not explicitly support endpoint discovery of forwarding paths through the network nor the discovery of properties and services associated with these paths. Path-aware networking, as defined in [Section 1.1](#) of [\[RFC9217\]](#), describes "endpoint discovery of the properties of paths they use for communication across an internetwork, and endpoint reaction to these properties that affects routing and/or data transfer". This document provides a generic definition of path properties, addressing the first of the questions in path-aware networking [\[RFC9217\]](#).

As terms related to paths have been used with different meanings in different areas of networking, first, this document provides a common terminology to define paths, path elements, and flows. Based on these terms, the document defines path properties. Then, this document provides some examples of use cases for path properties. Finally, the document lists several path properties that may be useful for the mentioned use cases. This list is intended to be neither exhaustive nor definitive.

Note that this document does not assume that any of the listed path properties are actually available to any entity. The question of how entities can discover and distribute path properties in a trustworthy way is out of scope for this document.

This document represents the consensus of the Path Aware Networking Research Group (PANRG).

2. Terminology

Entity: A physical or virtual device or function, or a collection of devices or functions, that plays a role related to path-aware networking for particular paths and flows. An entity can be on-path or off-path. On the path, an entity may participate in forwarding the flow, i.e., what may be called "data plane functionality". On or off the path, an entity may influence aspects of how the flow is forwarded, i.e., what may be called "control plane functionality", such as path selection or service invocation. An entity influencing forwarding aspects is usually aware of path properties, e.g., by observing or measuring them or by learning them from another entity.

Node: An on-path entity that processes packets, e.g., sends, receives, forwards, or modifies them. A node may be physical or virtual, e.g., a physical device, a service function provided as a virtual element, or even a single queue within a switch. A node may also be an entity that consists of a collection of devices or functions, e.g., an entire Autonomous System (AS).

Link: A medium or communication facility that connects two or more nodes with each other. A link enables a node to send packets to other nodes. Links can be physical, e.g., a Wi-Fi network that connects an Access Point to stations, or virtual, e.g., a virtual switch that connects two virtual machines hosted on the same physical machine. A link is unidirectional. As such, bidirectional communication can be modeled as two links between the same nodes in opposite directions.

Path element: Either a node or a link. For example, a path element can be an Abstract Network Element (ANE) as defined in [\[RFC9275\]](#).

Path: A sequence of adjacent path elements over which a packet can be transmitted, starting and ending with a node.

Paths are unidirectional and time dependent, i.e., there can be a variety of paths from one node to another, and the path over which packets are transmitted may change. A path definition can be strict (i.e., the exact sequence of path elements remains the same) or loose (i.e., the start and end node remain the same, but the path elements between them may vary over time).

The representation of a path and its properties may depend on the entity considering the path. On the one hand, the representation may differ due to entities having partial visibility of path elements comprising a path or their visibility changing over time. On the other hand, the representation may differ due to treating path elements at different levels of abstraction. For example, a path may be given as a sequence of physical nodes and the links connecting these nodes, be given as a sequence of logical nodes such as a sequence of ASes or an Explicit Route Object (ERO), or only consist of a specific source and destination that is known to be reachable from that source.

A multicast or broadcast setting where a packet is sent by one node and received by multiple nodes is described by multiple paths over which the packet is sent, one path for each combination of sending and receiving node; these paths do not have to be disjoint as defined by the disjointness path property, see [Section 4](#).

Endpoint: The endpoints of a path are the start and end node of the path. For example, an endpoint can be a host as defined in [\[RFC1122\]](#), which can be a client (e.g., a node running a web browser) or a server (e.g., a node running a web server).

Reverse Path: The path that is used by a remote node in the context of bidirectional communication.

Subpath: Given a path, a subpath is a sequence of adjacent path elements of this path.

Flow: One or multiple packets to which the traits of a path or set of subpaths may be applied in a functional sense. For example, a flow can consist of all packets sent within a TCP session with the same five-tuple between two hosts, or it can consist of all packets sent on the same physical link.

Property: A trait of one or a sequence of path elements, or a trait of a flow with respect to one or a sequence of path elements. An example of a link property is the maximum data rate that can be sent over the link. An example of a node property is the administrative domain that the node belongs to. An example of a property of a flow with respect to a subpath is the aggregated one-way delay of the flow being sent from one node to another node over this subpath. A property is thus described by a tuple containing the path element(s), the flow or an empty set if no packets are relevant for the property, the name of the property (e.g., maximum data rate), and the value of the property (e.g., 1 Gbps).

Aggregated property: A collection of multiple values of a property into a single value, according to a function. A property can be aggregated over:

- multiple path elements (i.e., a subpath), for example, the MTU of a path as the minimum MTU of all links on the path,
- multiple packets (i.e., a flow), for example, the median one-way latency of all packets between two nodes,
- or both path elements and packets, for example, the mean of the queueing delays of a flow on all nodes along a path.

The aggregation function can be numerical (e.g., median, sum, minimum) or logical (e.g., "true if all are true", "true if at least 50% of values are true"), or it can be an arbitrary function that maps multiple input values to an output value.

Observed property: A property that is observed for a specific path element, subpath, or path. A property may be observed using measurements, for example, the one-way delay of a specific packet transmitted from node to node.

Assessed property: An approximate calculation or assessment of the value of a property. An assessed property includes the reliability of the calculation or assessment. The notion of reliability depends on the property. For example, a path property based on an approximate calculation may describe the expected median one-way latency of packets sent on a path within the next second, including the confidence level and interval. A non-numerical assessment may instead include the likelihood that the property holds.

Target property: An objective that is set for a property over a path element, subpath, or path. Note that a target property can be set for observed properties, such as one-way delay, and also for properties that cannot be observed by the entity setting the target, such as inclusion of certain nodes on a path.

2.1. Terminology Usage for Specific Technologies

The terminology defined in this document is intended to be general and applicable to existing and future path-aware technologies. Using this terminology, a path-aware technology can define and consider specific path elements and path properties on a specific level of abstraction. For instance, a technology may define path elements as IP routers, e.g., in source routing [RFC1940]. Alternatively, it may consider path elements on a different layer of the Internet architecture [RFC1122] or as a collection of entities not tied to a specific layer, such as an AS or ERO. Even within a single path-aware technology, specific definitions might differ depending on the context in which they are used. For example, the endpoints might be the communicating hosts in the context of the transport layer, ASes that contain the hosts in the context of routing, or specific applications in the context of the application layer.

3. Use Cases for Path Properties

When a path-aware network exposes path properties to endpoints or other entities, these entities may use this information to achieve different goals. This section lists several use cases for path properties.

Note that this is not an exhaustive list; as with every new technology and protocol, novel use cases may emerge, and new path properties may become relevant. Moreover, for any particular technology, entities may have visibility of and control over different path elements and path properties and consider them on different levels of abstraction. Therefore, a new technology may implement an existing use case related to different path elements or on a different level of abstraction.

3.1. Path Selection

Nodes may be able to send flows via one (or a subset) out of multiple possible paths, and an entity may be able to influence the decision about which path(s) to use. Path selection may be feasible if there are several paths to the same destination (e.g., in case of a mobile device with two wireless interfaces, both providing a path) or if there are several destinations, and thus several paths, providing the same service (e.g., Application-Layer Traffic Optimization (ALTO) [RFC5693], an application layer peer-to-peer protocol allowing endpoints a better-than-random

peer selection). Entities can express their intent to achieve a specific goal by specifying target properties for their paths, e.g., related to performance or security. Then, paths can be selected that best meet the target properties, e.g., the entity can select these paths from all available paths or express the target properties to the network and rely on the network to select appropriate paths.

Target properties relating to network performance typically refer to observed properties, such as one-way delay, one-way packet loss, and link capacity. Entities then select paths based on their target property and the assessed property of the available paths that best match the application requirements. For such performance-related target properties, the observed property is similar to a Service Level Indicator (SLI), and the assessed property is similar to a Service Level Objective (SLO) for IETF Network Slices [[NETWORK-SLICES](#)]. As an example path-selection strategy, an entity may select a path with a short one-way delay for sending a small delay-sensitive query, while it may select a path with high link capacities on all links for retrieving a large file.

It is also possible for an entity to set target properties that it cannot (directly) observe, similar to Service Level Expectations (SLEs) for IETF Network Slices [[NETWORK-SLICES](#)]. This may apply to security-related target properties (e.g., to mandate that all enterprise traffic goes through a specific firewall) and path selection (e.g., to enforce traffic policies by allowing or disallowing sending flows over paths that involve specific networks or nodes).

Care needs to be taken when selecting paths based on observed path properties, as path properties that were previously measured may not be helpful in predicting current or future path properties, and such path selection may lead to unintended feedback loops. Also, there may be trade-offs between path properties (e.g., one-way delay and link capacity), and entities may influence these trade-offs with their choices. Finally, path selection may impact fairness. For example, if multiple entities concurrently attempt to meet their target properties using the same network resources, one entity's choices may influence the conditions on the path as experienced by flows of another entity.

As a baseline, a path-selection algorithm should aim to do a better job of meeting the target properties, and consequently accommodating the user's requirements, than the default case of not selecting a path most of the time.

Path selection can be done either by the communicating node(s) or by other entities within the network. A network (e.g., an AS) can adjust its path selection for internal or external routing based on path properties. In BGP, the Multi-Exit Discriminator (MED) attribute is used in the decision-making process to select which path to choose among those having the same AS path length and origin [[RFC4271](#)]; in a path-aware network, instead of using this single MED value, other properties such as link capacity or link usage could additionally be used to improve load balancing or performance [[PERFORMANCE-ROUTING](#)].

3.2. Protocol Selection

Before sending data over a specific path, an entity may select an appropriate protocol or configure protocol parameters depending on path properties. For example, an endpoint may cache state if a path allows the use of QUIC [[RFC9000](#)]; if so, it may first attempt to connect using

QUIC before falling back to another protocol when connecting over this path again. A video-streaming application may choose an (initial) video quality based on the achievable data rate or the monetary cost of sending data (e.g., volume-based or flat-rate cost model).

3.3. Service Invocation

In addition to path or protocol selection, an entity may choose to invoke additional functions in the context of Service Function Chaining [RFC7665], which may influence which nodes are on the path. For example, a 0-RTT Transport Converter [RFC8803] will be involved in a path only when invoked by an endpoint; such invocation will lead to the use of Multipath TCP (MPTCP) [RFC8684] or tcpcrypt [RFC8548] capabilities, while such use is not supported via the default forwarding path. Another example is a connection that is composed of multiple streams where each stream has specific service requirements. An endpoint may decide to invoke a given service function (e.g., transcoding) only for some streams while others are not processed by that service function.

4. Examples of Path Properties

This section gives some examples of path properties that may be useful, e.g., for the use cases described in Section 3.

Within the context of any particular technology, available path properties may differ as entities have insight into and are able to influence different path elements and path properties. For example, an endpoint may have some visibility into path elements that are close and on a low level of abstraction (e.g., individual nodes within the first few hops), or it may have visibility into path elements that are far away and/or on a higher level of abstraction (e.g., the list of ASes traversed). This visibility may depend on factors such as the physical or network distance or the existence of trust or contractual relationships between the endpoint and the path element(s). A path property can be defined relative to individual path elements, a sequence of path elements, or "end-to-end", i.e., relative to a path that comprises of two endpoints and a single virtual link connecting them.

Path properties may be relatively dynamic, e.g., the one-way delay of a packet sent over a specific path, or non-dynamic, e.g., the MTU of an Ethernet link that only changes infrequently. Usefulness over time differs depending on how dynamic a property is: the merit of a momentarily observed dynamic path property may diminish greatly as time goes on, e.g., it is possible for the observed values of one-way delay to change on timescales that are shorter than the one-way delay between the measurement point and an entity making a decision such as path selection, which may cause the measurement to be outdated when it reaches the decision-making entity. Therefore, consumers of dynamic path properties need to apply caution when using them, e.g., by aggregating them appropriately or applying a dampening function to their changes to avoid oscillation. In contrast, the observed value of a less dynamic path property might stay relevant for a longer period of time, e.g., a NAT typically stays on a particular path during the lifetime of a connection involving packets sent over this path.

Access Technology:

The physical- or link-layer technology used for transmitting or receiving a flow on one or multiple path elements. If known, the access technology may be given as an abstract link type, e.g., as Wi-Fi, wired Ethernet, or cellular. It may also be given as a specific technology used on a link, e.g., 3GPP cellular or 802.11 Wireless Local Area Network (WLAN). Other path elements relevant to the access technology may include nodes related to processing packets on the physical or link layer, such as elements of a cellular core network. Note that there is no common registry of possible values for this property.

Monetary Cost: The price to be paid to transmit or receive a specific flow across a network to which one or multiple path elements belong.

Service Function: A service function that a path element applies to a flow, see [\[RFC7665\]](#). Examples of abstract service functions include firewalls, Network Address Translation (NAT), and TCP Performance Enhancing Proxies. Some stateful service functions, such as NAT, need to observe the same flow in both directions, e.g., by being an element of both the path and the reverse path.

Transparency: When a node performs an action A on a flow F, the node is transparent to F with respect to some (meta-)information M if the node performs A independently of M. M can, for example, be the existence of a protocol (header) in a packet or the content of a protocol header, payload, or both. For example, A can be blocking packets or reading and modifying (other protocol) headers or payloads. Transparency can be modeled using a function f, which takes as input F and M and outputs the action taken by the node. If a taint analysis shows that the output of f is not tainted (impacted) by M, or if the output of f is constant for arbitrary values of M, then the node is considered to be transparent. An IP router could be transparent to transport protocol headers such as TCP/UDP but not transparent to IP headers since its forwarding behavior depends on the IP headers. A firewall that only allows outgoing TCP connections by blocking all incoming TCP SYN packets regardless of their IP address is transparent to IP but not to TCP headers. Finally, a NAT that actively modifies IP and TCP/UDP headers based on their content is not transparent to either IP or TCP/UDP headers. Note that according to this definition, a node that modifies packets in accordance with the endpoints, such as a transparent HTTP proxy, as defined in [\[RFC9110\]](#), and a node listening and reacting to implicit or explicit signals, see [\[RFC8558\]](#), are not considered transparent.

Transparency only applies to nodes and not to links, as a link cannot modify or perform any other actions on the packets by itself. For example, if the content of a packet is altered when forwarded over a Generic Routing Encapsulation (GRE) tunnel [\[RFC2784\]](#) [\[RFC7676\]](#), per this document the software instances that terminate the tunnel are considered nodes over which the actions are performed; thus, the transparency definition applies to these nodes.

Administrative Domain: The identity of an individual or an organization that controls access to a path element (or several path elements). Examples of administrative domains are an IGP area, an AS, or a service provider network.

Routing Domain Identifier: An identifier indicating the routing domain of a path element. Path elements in the same routing domain are in the same administrative domain and use a common routing protocol to communicate with each other. An example of a routing domain identifier is the globally unique Autonomous System Number (ASN) as defined in [\[RFC1930\]](#).

Disjointness: For a set of two paths or subpaths, the number of shared path elements can be a measure of intersection (e.g., Jaccard coefficient, which is the number of shared elements divided by the total number of elements). Conversely, the number of non-shared path elements can be a measure of disjointness (e.g., $1 - \text{Jaccard coefficient}$). A multipath protocol might use disjointness as a metric to reduce the number of single points of failure. As paths can be defined at different levels of abstraction, two paths may be disjoint at one level of abstraction but not on another.

Symmetric Path: Two paths are symmetric if the path and its reverse path consist of the same path elements on the same level of abstraction, but in reverse order. For example, a path that consists of layer 3 switches and links between them and a reverse path with the same path elements but in reverse order are considered "routing" symmetric, as the same path elements on the same level of abstraction (IP forwarding) are traversed in the opposite direction. Symmetry can depend on the level of abstraction on which the path is defined or modeled. If there are two parallel physical links between two nodes, modeling them as separate links may result in a flow using asymmetric paths, and modeling them as a single virtual link may result in symmetric paths, e.g., if the difference between the two physical links is irrelevant in a particular context.

Path MTU: The maximum size, in octets, of a packet or frame that can be transmitted without fragmentation.

Transport Protocols available: Whether a specific transport protocol can be used to establish a connection over a path or subpath, e.g., whether the path is QUIC-capable or MPTCP-capable, based on input such as policy, cached knowledge, or probing results.

Protocol Features available: Whether a specific protocol feature is available over a path or subpath, e.g., Explicit Congestion Notification (ECN) or TCP Fast Open.

Some path properties express the performance of the transmission of a packet or flow over a link or subpath. Such transmission performance properties can be observed or assessed, e.g., by endpoints or by path elements on the path, or they may be available as cost metrics, see [\[RFC9439\]](#). Transmission performance properties may be made available in an aggregated form, such as averages or minimums. Properties related to a path element that constitutes a single layer 2 domain are abstracted from the used physical- and link-layer technology, similar to [\[RFC8175\]](#).

Link Capacity: The link capacity is the maximum data rate at which data that was sent over a link can correctly be received at the node adjacent to the link. This property is analogous to the link capacity defined in [\[RFC5136\]](#) and [\[RFC9097\]](#) but is not restricted to IP-layer traffic.

Link Usage: The link usage is the actual data rate at which data that was sent over a link is correctly received at the node adjacent to the link. This property is analogous to the link usage defined in [RFC5136] and [RFC9097] but is not restricted to IP-layer traffic.

One-Way Delay: The one-way delay is the delay between a node sending a packet and another node on the same path receiving the packet. This property is analogous to the one-way delay defined in [RFC7679] but is not restricted to IP-layer traffic.

One-Way Delay Variation: The variation of the one-way delays within a flow. This property is similar to the one-way delay variation defined in [RFC3393], but it is not restricted to IP-layer traffic and it is defined for packets on the same flow instead of packets sent between a source and destination IP address.

One-Way Packet Loss: Packets sent by a node but not received by another node on the same path after a certain time interval are considered lost. This property is analogous to the one-way loss defined in [RFC7680] but is not restricted to IP-layer traffic. Metrics such as loss patterns [RFC3357] and loss episodes [RFC6534] can be expressed as aggregated properties.

5. Security Considerations

If entities are basing policy or path-selection decisions on path properties, they need to rely on the accuracy of path properties that other devices communicate to them. In order to be able to trust such path properties, entities may need to establish a trust relationship or be able to independently verify the authenticity, integrity, and correctness of path properties received from another entity.

Entities that reveal their target path properties to the network can negatively impact their own privacy, e.g., if the target property leaks personal information about a user, such as their identity or which (type of) application is used. Such information could then allow network operators to block or reprioritize traffic for certain users and/or applications. Conversely, if privacy-enhancing technologies, e.g., MASQUE proxies [RFC9298], are used on a path, the path may only be partially visible to any single entity. This may diminish the usefulness of path-aware technologies over this path.

The need for, and potential definition of, security- and privacy-related path properties, such as confidentiality and integrity protection of payloads, are the subject of ongoing discussion and research, for example, see [RFC9049] and [RFC9217]. As the discussion of such properties is not mature enough, they are out of scope for this document. One aspect discussed in this context is that security-related properties are difficult to characterize since they are only meaningful with respect to a threat model that depends on the use case, application, environment, and other factors. Likewise, properties for trust relations between entities cannot be meaningfully defined without a concrete threat model, and defining a threat model is out of scope for this document. Properties related to confidentiality, integrity, and trust seem to be orthogonal to the path terminology and path properties defined in this document, since they are tied to the communicating nodes and the protocols they use (e.g., client and server using HTTPS, or client and remote network node using a VPN service) as well as additional context, such as keying material and who has access to such a context. In contrast, the path as defined in this document

is typically oblivious to these aspects. Intuitively, the path describes what function the network applies to packets, while confidentiality, integrity, and trust describe what function the communicating parties apply to packets.

6. IANA Considerations

This document has no IANA actions.

7. Informative References

- [NETWORK-SLICES]** Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "A Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-22, August 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-22>>.
- [PERFORMANCE-ROUTING]** Xu, X., Hegde, S., Talaulikar, K., Boucadair, M., and C. Jacquenet, "Performance-based BGP Routing Mechanism", Work in Progress, Internet-Draft, draft-ietf-idr-performance-routing-03, 21 December 2020, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-performance-routing-03>>.
- [RFC1122]** Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.
- [RFC1930]** Hawkinson, J. and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)", BCP 6, RFC 1930, DOI 10.17487/RFC1930, March 1996, <<https://www.rfc-editor.org/info/rfc1930>>.
- [RFC1940]** Estrin, D., Li, T., Rekhter, Y., Varadhan, K., and D. Zappala, "Source Demand Routing: Packet Format and Forwarding Specification (Version 1)", RFC 1940, DOI 10.17487/RFC1940, May 1996, <<https://www.rfc-editor.org/info/rfc1940>>.
- [RFC2784]** Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [RFC3357]** Koodli, R. and R. Ravikanth, "One-way Loss Pattern Sample Metrics", RFC 3357, DOI 10.17487/RFC3357, August 2002, <<https://www.rfc-editor.org/info/rfc3357>>.
- [RFC3393]** Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, DOI 10.17487/RFC3393, November 2002, <<https://www.rfc-editor.org/info/rfc3393>>.
- [RFC4271]** Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

-
- [RFC5136] Chimento, P. and J. Ishac, "Defining Network Capacity", RFC 5136, DOI 10.17487/RFC5136, February 2008, <<https://www.rfc-editor.org/info/rfc5136>>.
- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, DOI 10.17487/RFC5693, October 2009, <<https://www.rfc-editor.org/info/rfc5693>>.
- [RFC6534] Duffield, N., Morton, A., and J. Sommers, "Loss Episode Metrics for IP Performance Metrics (IPPM)", RFC 6534, DOI 10.17487/RFC6534, May 2012, <<https://www.rfc-editor.org/info/rfc6534>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.
- [RFC7676] Pignataro, C., Bonica, R., and S. Krishnan, "IPv6 Support for Generic Routing Encapsulation (GRE)", RFC 7676, DOI 10.17487/RFC7676, October 2015, <<https://www.rfc-editor.org/info/rfc7676>>.
- [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.
- [RFC7680] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", STD 82, RFC 7680, DOI 10.17487/RFC7680, January 2016, <<https://www.rfc-editor.org/info/rfc7680>>.
- [RFC8175] Ratliff, S., Jury, S., Satterwhite, D., Taylor, R., and B. Berry, "Dynamic Link Exchange Protocol (DLEP)", RFC 8175, DOI 10.17487/RFC8175, June 2017, <<https://www.rfc-editor.org/info/rfc8175>>.
- [RFC8548] Bittau, A., Giffin, D., Handley, M., Mazieres, D., Slack, Q., and E. Smith, "Cryptographic Protection of TCP Streams (tcpcrypt)", RFC 8548, DOI 10.17487/RFC8548, May 2019, <<https://www.rfc-editor.org/info/rfc8548>>.
- [RFC8558] Hardie, T., Ed., "Transport Protocol Path Signals", RFC 8558, DOI 10.17487/RFC8558, April 2019, <<https://www.rfc-editor.org/info/rfc8558>>.
- [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>>.
- [RFC8803] Bonaventure, O., Ed., Boucadair, M., Ed., Gundavelli, S., Seo, S., and B. Hesmans, "0-RTT TCP Convert Protocol", RFC 8803, DOI 10.17487/RFC8803, July 2020, <<https://www.rfc-editor.org/info/rfc8803>>.
- [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>>.

- [RFC9049] Dawkins, S., Ed., "Path Aware Networking: Obstacles to Deployment (A Bestiary of Roads Not Taken)", RFC 9049, DOI 10.17487/RFC9049, June 2021, <<https://www.rfc-editor.org/info/rfc9049>>.
- [RFC9097] Morton, A., Geib, R., and L. Ciavattone, "Metrics and Methods for One-Way IP Capacity", RFC 9097, DOI 10.17487/RFC9097, November 2021, <<https://www.rfc-editor.org/info/rfc9097>>.
- [RFC9110] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/RFC9110, June 2022, <<https://www.rfc-editor.org/info/rfc9110>>.
- [RFC9217] Trammell, B., "Current Open Questions in Path-Aware Networking", RFC 9217, DOI 10.17487/RFC9217, March 2022, <<https://www.rfc-editor.org/info/rfc9217>>.
- [RFC9275] Gao, K., Lee, Y., Randriamasy, S., Yang, Y., and J. Zhang, "An Extension for Application-Layer Traffic Optimization (ALTO): Path Vector", RFC 9275, DOI 10.17487/RFC9275, September 2022, <<https://www.rfc-editor.org/info/rfc9275>>.
- [RFC9298] Schinazi, D., "Proxying UDP in HTTP", RFC 9298, DOI 10.17487/RFC9298, August 2022, <<https://www.rfc-editor.org/info/rfc9298>>.
- [RFC9439] Wu, Q., Yang, Y., Lee, Y., Dhody, D., Randriamasy, S., and L. Contreras, "Application-Layer Traffic Optimization (ALTO) Performance Cost Metrics", RFC 9439, DOI 10.17487/RFC9439, August 2023, <<https://www.rfc-editor.org/info/rfc9439>>.

Acknowledgments

Thanks to the Path Aware Networking Research Group for the discussion and feedback. Specifically, thanks to Mohamed Boucadair for the detailed review, various text suggestions, and shepherding; thanks to Brian Trammell for suggesting the flow definition; and thanks to Luis M. Contreras, Spencer Dawkins, Paul Hoffman, Jake Holland, Colin Perkins, Adrian Perrig, and Matthias Rost for the reviews, comments, and suggestions. Many thanks to Dave Oran for his careful IRSG review.

Authors' Addresses

Reese Enhardt

Netflix

Email: ietf@tenghardt.net

Cyrill Krähenbühl

ETH Zürich

Email: cyrill.kraehenbuehl@inf.ethz.ch