

Workgroup: PANRG
Internet-Draft:
draft-irtf-panrg-path-properties-05
Published: 7 March 2022
Intended Status: Informational
Expires: 8 September 2022
Authors: T. Enhardt C. Krähenbühl
 Netflix ETH Zürich

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 specifies several path properties which might be useful to endpoints or other entities, e.g., for selecting between paths or for invoking some of the provided services.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the "Path-Aware Networking Research Group" mailing list (PANRG), which is archived at <https://mailarchive.ietf.org/arch/browse/panrg/>. Subscription information is at <https://www.ietf.org/mailman/listinfo/panrg/>.

Source for this draft and an issue tracker can be found at <https://github.com/panrg/path-properties/>.

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 8 September 2022.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
 - [2.1. Terminology usage for specific technologies](#)
- [3. Use Cases for Path Properties](#)
 - [3.1. Path Selection](#)
 - [3.2. Protocol Selection](#)
 - [3.3. Service Invocation](#)
- [4. Examples of Path Properties](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

1. Introduction

The current Internet architecture does not explicitly support endpoint discovery of forwarding paths through the network as well as the discovery of properties and services associated with these paths. Path-aware networking, as defined in Section 1.1 of [[I-D.irtf-panrg-questions](#)], 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 [[I-D.irtf-panrg-questions](#)].

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.

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.

2. Terminology

Entity: A physical or virtual device or function, or a collection of devices or functions, which 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 which 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 which 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 which connects an Access Point to stations, or virtual, e.g., a virtual switch which 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 [[I-D.ietf-alto-path-vector](#)].

Path: A sequence of adjacent path elements over which a packet can be transmitted, starting and ending with a node. A path is unidirectional. Paths are time-dependent, i.e., the sequence of path elements over which packets are sent from one node to another may change. A path is defined between two nodes. For multicast or broadcast, a packet may be sent by one node and received by multiple nodes. In this case, the packet is sent over

multiple paths at once, one path for each combination of sending and receiving node; these paths do not have to be disjoint. Note that an entity may have only partial visibility of the path elements that comprise a path and visibility may change over time. Different entities may have different visibility of a path and/or treat 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, or it may be given as a sequence of logical nodes such as a sequence of ASes or an Explicit Route Object (ERO). Similarly, the representation of a path and its properties, as it is known to a specific entity, may be more complex and include details about the physical layer technology, or it may be more abstract and only consist of a specific source and destination which is known to be reachable from that source.

Endpoint: The endpoints of a path are the first and the last node on 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., 1Gbps).

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), e.g.,

the MTU of a path as the minimum MTU of all links on the path, over multiple packets (i.e., a flow), e.g., the median one-way latency of all packets between two nodes, or over both, e.g., 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, it can be logical, e.g., "true if all are true", "true if at least 50\% of values are true", or an arbitrary function which maps multiple input values to an output value.

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

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.

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 an 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 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). Care needs to be taken when selecting paths based on 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.

Entities may select their paths to fulfill a specific goal, e.g., related to security or performance. As an example of security-related path selection, an entity may allow or disallow sending flows over paths involving specific networks or nodes to enforce traffic policies. In an enterprise network where all traffic has to go through a specific firewall, a path-aware entity can implement this policy using path selection. As an example of performance-related path selection, an entity may prefer paths with performance properties that best match application requirements. For example, for sending a small delay sensitive query, the entity may select a path with a short One-Way Delay, while for retrieving a large file, it may select a path with high Link Capacities on all links. Note, 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. As a baseline, a path selection algorithm should aim to not perform worse than the default case 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 [[I-D.ietf-idr-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 on whether a path allows the use of QUIC [[I-D.ietf-quic-transport](#)] and if so, 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-base 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 what nodes are on the path. For example, a 0-RTT Transport Converter [[I-D.ietf-tcpm-converters](#)] will be involved in a path only when invoked by an endpoint; such invocation will lead to the use of MPTCP or TCPinc capabilities while such use is not supported via the default forwarding path. Another example is a connection which 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 which 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 on a low level of abstraction and close, 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 which only changes infrequently. Usefulness over time differs depending on how dynamic a property is: The merit of a momentary measurement of a dynamic path property diminishes greatly as time goes on, e.g. the merit of an RTT measurement from a few seconds ago is quite small, while a non-dynamic path property might stay relevant for a longer period of time, e.g. a NAT typically stays on a specific 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., 2G, 3G, 4G, or 5G cellular, or 802.11a, b, g, n, or ac Wi-Fi. 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 backbone 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 optimizers. 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. A can for example 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 [[RFC2616](#)], and a node listening and reacting to implicit or explicit signals, see [[RFC8558](#)], are not considered transparent.

Administrative Domain: The identity of an individual or an organization that owns 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.

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 which 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.

Path MTU: The maximum size, in octets, of an IP packet 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 cached knowledge.

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 measured or approximated, e.g., by endpoints or by path elements on the path, or they may be available as cost metrics, see [[I-D.ietf-alto-performance-metrics](#)]. Transmission performance properties may be made available in an aggregated form, such as averages or minimums. Properties related to a path element which 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](#)] but 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](#)] but 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 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 not restricted to IP-layer traffic and 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 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 verify the authenticity, integrity, and correctness of path properties received from another entity.

Security related properties such as confidentiality and integrity protection of payloads are difficult to characterize since they are only meaningful with respect to a threat model which 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 draft. Properties related to confidentiality, integrity, and trust are orthogonal to the path terminology and path properties defined in this document. Such properties 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 VPN) while the path is typically oblivious to them. 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

[I-D.ietf-alto-path-vector] Gao, K., Lee, Y., Randriamasy, S., Yang, Y. R., and J. J. Zhang, "An ALTO Extension: Path Vector", Work in Progress, Internet-Draft, draft-ietf-alto-path-vector-24, 7 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-alto-path-vector-24>>.

[I-D.ietf-alto-performance-metrics]

Wu, Q., Yang, Y. R., Lee, Y., Dhody, D., Randriamasy, S., and L. M. C. Murillo, "ALTO Performance Cost Metrics", Work in Progress, Internet-Draft, draft-ietf-alto-performance-metrics-26, 2 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-alto-performance-metrics-26>>.

[I-D.ietf-idr-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, 22 December 2020, <<https://>

datatracker.ietf.org/doc/html/draft-ietf-idr-performance-routing-03>.

[I-D.ietf-quic-transport] Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", Work in Progress, Internet-Draft, draft-ietf-quic-transport-34, 14 January 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-transport-34>>.

[I-D.ietf-tcpm-converters] Bonaventure, O., Boucadair, M., Gundavelli, S., Seo, S., and B. Hesmans, "0-RTT TCP Convert Protocol", Work in Progress, Internet-Draft, draft-ietf-tcpm-converters-19, 22 March 2020, <<https://datatracker.ietf.org/doc/html/draft-ietf-tcpm-converters-19>>.

[I-D.irtf-panrg-questions] Trammell, B., "Current Open Questions in Path Aware Networking", Work in Progress, Internet-Draft, draft-irtf-panrg-questions-12, 25 January 2022, <<https://datatracker.ietf.org/doc/html/draft-irtf-panrg-questions-12>>.

[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/rfc/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/rfc/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/rfc/rfc1940>>.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, DOI 10.17487/

RFC2616, June 1999, <<https://www.rfc-editor.org/rfc/rfc2616>>.

[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/rfc/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/rfc/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/rfc/rfc4271>>.

[RFC5136] Chimento, P. and J. Ishac, "Defining Network Capacity", RFC 5136, DOI 10.17487/RFC5136, February 2008, <<https://www.rfc-editor.org/rfc/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/rfc/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/rfc/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/rfc/rfc7665>>.

[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/rfc/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/rfc/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/rfc/rfc8175>>.

[RFC8558]

Hardie, T., Ed., "Transport Protocol Path Signals", RFC 8558, DOI 10.17487/RFC8558, April 2019, <<https://www.rfc-editor.org/rfc/rfc8558>>.

Acknowledgments

Thanks to the Path-Aware Networking Research Group for the discussion and feedback. Specifically, thanks to Mohamed Boudacair for the detailed review and various text suggestions, thanks to Brian Trammell for suggesting the flow definition, thanks to Adrian Perrig and Matthias Rost for the detailed feedback, thanks to Paul Hoffman for the editorial changes, thanks to Luis M. Contreras and Jake Holland for the reviews, and thanks to Spencer Dawkins for the comments and suggestions.

Authors' Addresses

Theresa Enghardt
Netflix

Email: ietf@tenghardt.net

Cyrill Krähenbühl
ETH Zürich

Email: cyrill.kraehenbuehl@inf.ethz.ch