

Path Aware Networking RG
Internet-Draft
Intended status: Informational
Expires: June 10, 2018

B. Trammell
ETH Zurich
December 07, 2017

Open Questions in Path Aware Networking
draft-trammell-panrg-questions-02

Abstract

This document poses open questions in path-aware networking, as a background for framing discussions in the Path Aware Networking proposed Research Group (PANRG). These are split into making properties of Internet paths available to endpoints, and allowing endpoints to select paths through the Internet for their traffic.

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 June 10, 2018.

Copyright Notice

Copyright (c) 2017 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 to Path-Aware Networking	2
2. Questions	3
2.1. A Vocabulary of Path Properties	3
2.2. Discovery, Distribution, and Trustworthiness of Path Properties	3
2.3. Supporting Path Selection	4
2.4. Interfaces for Path Awareness	4
2.5. Implications of Path Awareness for the Data Plane	4
2.6. What is an Endpoint?	5
2.7. Operating a Path Aware Network	5
2.8. Deploying a Path Aware Network	6
3. Acknowledgments	6
4. References	6
4.1. Normative References	6
4.2. Informative References	7
Author's Address	7

1. Introduction to Path-Aware Networking

In the current Internet architecture, the network layer provides an unverifiable, best-effort service: an application can assume that a packet with a given destination address will eventually be forwarded toward that destination, but little else. A transport layer protocol such as TCP can provide reliability over this best-effort service, and a protocol above the network layer such as IPsec AH [RFC4302] or TLS [RFC5246] can authenticate the remote endpoint. However, no explicit information about the path is available, and assumptions about that path sometimes do not hold, sometimes with serious impacts on the application, as in the case with BGP hijacking attacks.

By contrast, in a path-aware networking architecture, endpoints have the ability to select or influence the path through the network used by any given packet, and the network layer explicitly exposes information about the path or paths available between two endpoints to those endpoints so that they can make this selection. Path control at the packet level enables new transport protocols that can leverage multipath connectivity across maximally-disjoining paths through the Internet, even over a single interface. It also provides transparency and control for applications and end-users to specify constraints on the paths its traffic should traverse, for instance to confound pervasive passive surveillance in the network core.

2. Questions

Realizing path-aware networking requires answers to a set of open research questions. This document poses these questions, as a starting point for discussions about how to realize path awareness in the Internet, and to direct future research efforts within the Path Aware Networking Research Group.

2.1. A Vocabulary of Path Properties

In order for information about paths to be exposed to the endpoints, and for those endpoints to be able to use that information, it is necessary to define a common vocabulary for path properties. The elements of this vocabulary could include relatively static properties, such as the presence of a given node on the path; as well as relatively dynamic properties, such as the current values of metrics such as loss and latency.

This vocabulary must be defined carefully, as its design will have impacts on the expressiveness of a given path-aware internetworking architecture. This expressiveness also exhibits tradeoffs. For example, a system that exposes node-level information for the topology through each network would maximize information about the individual components of the path at the endpoints at the expense of making internal network topology universally public, which may be in conflict with the business goals of each network's operator.

The first question is therefore: how are path properties defined and represented?

2.2. Discovery, Distribution, and Trustworthiness of Path Properties

Once endpoints and networks have a shared vocabulary for expressing path properties, the network must have some method for distributing those path properties to the endpoint. Regardless of how path property information is distributed to the endpoints, the endpoints require a method to authenticate the properties - to determine that they originated from and pertain to the path that they purport to. The end goal of authentication is not necessarily to establish that a given property is actually bound to a given path, but to ensure that the information is trustworthy, that actions taken based on it will have the predicted result.

Choices in an distribution and authentication methods will have impacts on the scalability of a path-aware architecture. Possible dimensions in the space of distribution methods include in-band versus out-of-band, push versus pull versus publish-subscribe, and so on. There are temporal issues with path property dissemination as

well, especially with dynamic properties, since the measurement or elicitation of dynamic properties may be outdated by the time that information is available at the endpoints, and interactions between the measurement and dissemination delay may exhibit pathological behavior for unlucky points in the parameter space.

The second question: how do endpoints get access to trustworthy path properties?

2.3. Supporting Path Selection

Access to trustworthy path properties is only half of the challenge in establishing a path-aware architecture. Endpoints must be able to use this information in order to select paths for traffic they send. As with path property distribution, choices made in path selection methods will also have an impact on the scalability and expressiveness of a path-aware architecture, and dimensions included in-band versus out-of-band, as well. Paths may also be selected on multiple levels of granularity - per packet, per flow, per aggregate - and this choice also has impacts on the scalability/expressiveness tradeoff.

The third question: how can endpoints select paths to use for traffic in a way that can be trusted by the network?

2.4. Interfaces for Path Awareness

In order for applications to make effective use of a path-aware networking architecture, the interfaces presented by the network and transport layers must also expose path properties to the application in a useful way, and provide a useful selection for path selection. Path selection must be possible based not only on the preferences and policies of the application developer, but of end-users as well.

The fourth question: how can interfaces to the transport and application layers support the use of path awareness?

2.5. Implications of Path Awareness for the Data Plane

In the current Internet, the basic assumption that at a given time t all traffic for a given flow will traverse a single path, for some definition of path, generally holds. In a path aware network, this assumption no longer holds. The failure of this assumption has implications for the design of protocols above a path-aware network layer.

For example, one advantage of multipath communication is that a given end-to-end flow can be "sprayed" along multiple paths in order to

confound attempts to collect data or metadata from those flows for pervasive surveillance purposes [RFC7624]. However, the benefits of this approach are reduced if the upper-layer protocols use linkable identifiers on packets belonging to the same flow across different paths. Clients may mitigate linkability by opting to not re-use cleartext connection identifiers, such as TLS session IDs or tickets, on separate paths. The privacy-conscious strategies required for effective privacy in a path-aware Internet are only possible if higher-layer protocols such as TLS permit clients to obtain unlinkable identifiers.

The fifth question: how should transport-layer and higher layer protocols be redesigned to work most effectively over a path-aware networking layer?

2.6. What is an Endpoint?

The vision of path-aware networking articulated so far makes an assumption that path properties will be disseminated to endpoints on which applications are running (terminals with user agents, servers, and so on). However, incremental deployment may require that a path-aware network "core" be used to interconnect islands of legacy protocol networks. In these cases, it is the gateways, not the application endpoints, that receive path properties and make path selections for that traffic. The interfaces provided this gateway are necessarily different than those a path-aware networking layer provides to its transport and application layers, and the path property information the gateway needs and makes available over those interfaces may also be different.

The sixth question: how is path awareness (in terms of vocabulary and interfaces) different when applied to tunnel and overlay endpoints?

2.7. Operating a Path Aware Network

The network operations model in the current Internet architecture assumes that traffic flows are controlled by the decisions and policies made by network operators, as expressed in interdomain routing protocols. In a path-aware network with effective path selection, however, this assumption no longer holds, as endpoints may react to path properties by selecting alternate paths. Competing control inputs from path-aware endpoints and the interdomain routing control plane may lead to more difficult traffic engineering or nonconvergent routing, especially if the endpoints' and operators' idea of the "best" path for given traffic differs significantly.

The seventh question: how can a path aware network in a path aware internetwork be effectively operated, given control inputs from the network administrator as well as from the endpoints?

2.8. Deploying a Path Aware Network

The vision presented in the introduction discusses path aware networking from the point of view of the benefits accruing at the endpoints, to designers of transport protocols and applications as well as to the end users of those applications. However, this vision requires action not only at the endpoints but within the interconnected networks offering path aware connectivity. While the specific actions required are a matter of the design and implementation of a specific realization of a path aware protocol stack, it is clear that any path aware architecture will require network operators to give up some control of their networks over to endpoint-driven control inputs. The incentives for network operators and equipment vendors to do this must be made clear.

The eighth question: how can the incentives of network operators and end-users be aligned to realize the vision of path aware networking?

3. Acknowledgments

Many thanks to Adrian Perrig, Jean-Pierre Smith, Mirja Kuehlewind, Olivier Bonaventure, Martin Thomson, Shwetha Bhandari, and Chris Wood, for discussions leading to questions in this document.

This work is partially supported by the European Commission under Horizon 2020 grant agreement no. 688421 Measurement and Architecture for a Middleboxed Internet (MAMI), and by the Swiss State Secretariat for Education, Research, and Innovation under contract no. 15.0268. This support does not imply endorsement.

4. References

4.1. Normative References

- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.

4.2. Informative References

[RFC7624] Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C., and D. Borkmann, "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement", RFC 7624, DOI 10.17487/RFC7624, August 2015, <<https://www.rfc-editor.org/info/rfc7624>>.

Author's Address

Brian Trammell
ETH Zurich
Gloriastrasse 35
8092 Zurich
Switzerland

Email: ietf@trammell.ch