Open Questions in Path Aware Networking
draft-irtf-panrg-questions-01

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 April 20, 2019.

Table of Contents

1. Introduction to Path-Aware Networking

   In the current Internet architecture, the interdomain 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 internetworking 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-disjoint 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 that traffic should traverse, for instance
   to confound pervasive passive surveillance in the network core.

   We note that this property of "path awareness" already exists in many
   Internet-connected networks in an intradomain context. Indeed, much
   of the practice of network engineering using encapsulation at layer 3

can be said to be "path aware", in that it explicitly assigns traffic
at tunnel endpoints to a given path within the network. Path-aware
internetworking seeks to extend this awareness across domain
boundaries without resorting to overlays, except as a transition
technology.

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: 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.

Choices in 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 scalabilty/expressiveness
tradeoff. Path selection must, like path property information, be
trustworthy, such that the result of a path selection at an endpoint
is predictable.

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

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 set of paths among which the
application can select. Path selection must be possible based not
only on the preferences and policies of the application developer,
but of end-users as well. Also, the path selection interfaces
presented to applications and end users will need to support multiple
levels of granularity. Most applications' requirements can be
satisfied with the expression path selection policies in terms of
properties of the paths, while some applications may need finer-
grained, per-path control.

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 absence of this assumption has
   implications for the design of protocols above any 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 by 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 network providing path selection to the
   endpoints, 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'
   notion of the "best" path for given traffic diverges significantly.

   A concept for path aware network operations will need to have clear
   methods for the resolution of apparent (if not actual) conflicts of
   intent between the network's operator and the path selection at an
   endpoint. It will also need set of safety principles to ensure that
   increasing path control does not lead to decreasing connectivity; one
   such safety principle could be "the existence of at least one path
   between two endpoints guarantees the selection of at least one path
   between those endpoints."

   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 than any path aware architecture will require
   network operators to give up some control of their networks over to
   endpoint-driven control inputs.

   Here the question of apparent versus actual conflicts of intent
   arises again: certain network operations requirements may appear
   essential, but are merely accidents of the interfaces provided by
   current routing and management protocols. Incentives for deployment
   must show how existing network operations requirements are met
   through new path selection and property dissemination mechanisms.

The incentives for network operators and equipment vendors to do
provide be made clear, in terms of a plan to transition [RFC8170] an
internetwork to path-aware operation, one network and facility at a
time.

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, Chris Wood,
and Lee Howard, 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>.

[RFC8170] Thaler, D., Ed., "Planning for Protocol Adoption and
           Subsequent Transitions", RFC 8170, DOI 10.17487/RFC8170,
           May 2017, <https://www.rfc-editor.org/info/rfc8170>.

Author's Address

    Brian Trammell
    ETH Zurich
    Gloriastrasse 35
    8092 Zurich
    Switzerland

    Email: ietf@trammell.ch