Network Working Group   M. Nottingham
Internet-Draft
Intended status: Informational   July 04, 2014
Expires: January 05, 2015

                  Granular Information about Networks
                        draft-nottingham-gin-00

Abstract

   Protocol endpoints often want to adapt their behavior based upon the
   current properties of the network path, but have limited information
   available to aid these decisions. This document motivates discussion
   of protocol work to make this information available.

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 http://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 January 05, 2015.

Table of Contents

1. Introduction

   Protocol endpoints often want to adapt their behavior based upon the
   current properties of the network path.

   For example, it has become common practice for HTTP [RFC7230] servers
   to adapt the responses they give based upon the IP address of the
   client, client "fingerprinting" (e.g., using the User-Agent request
   header field), and other properties.

   Likewise, client using HTTP sometimes adapt their behavior in a
   similar fashion; for example, a mobile client on a 3G network might
   download a different video file then if it were on a wifi network.
   Often, the goal of these adaptations is to improve user experience by
   making content more suitable for the properties of the network it is
   traversing, whilst utilizing the network resources more optimally.

   There are currently a number of sources of information to inform
   these decisions, but they share a few limitations. For example, it
   is possible to measure delay to a given server using ICMP, but the
   results are ephemeral, and may change if a different server has
   changed.

   There have also been attempts to provide relevant information in
   APIs; for example, [netinfo]. Doing so has proven to be difficult,
   because of the limited information available to the client.

   To address these issues, network operators have been deploying
   infrastructure that uses the information available to them to modify
   content; e.g., [bytemobile], [verizon], [syniverse], [flashnet].

However, at the same time, encryption has become more prevalent on
the Internet, with many prominent (and heavily traffic'd) Web sites
going HTTPS-only. This frustrates attempts to adapt content in the
network.

This document proposes an alternative approach. By making the
information that the network operators have available to the
endpoints, it allows them to make more informed choices about
content, thereby allowing the user experience to be improved and the
network to be used more optimally without requiring that the end-to-
end nature of encryption (e.g., in HTTPS) to be compromised.

1.1. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in [RFC2119].

2. Requirements

To be useful to endpoints, the information a network exposes needs to
be:

o Specific to the client - General information about network
   properties is often an improvement over current practice, but to
   be truly useful, it should be able to be tailored to a specific
   client IP address.

o Reasonably current - Information from fifteen minutes ago is often
   useless; when necessary, an endpoint ought to be able to get
   information that is fresh (on the granularity of a few seconds).

o Scalable - The overhead of conveying information to clients ought
   to be minimal, and it needs to be usable on the scale of a large
   Web site.

o Private - The protocol ought not expose details of private
   networks, or any personally identifying information beyond that
   already available.

A protocol for exposing this information necessarily must choose the
scope of its applicability. Due to the nature of the Internet, it is
not practical to meet the goals above for the end-to-end path(s)
between any given pair of IP addresses; the permutations are
impractical, and discovering meaningful information on this scale is
likewise unlikely.

However, it is comparatively easy for a network operator to expose what it considers to be the "last mile" properties of an IP address. For example, an ISP providing ADSL access to its subscribers could advertise the properties of those end links, whereas a mobile operator could use the information available to advertise the properties of individual subscriber handset IP addresses (whether they be globally routable, or behind NAT).

This partial information is not a complete picture, of course, but it is information that's difficul to aquire now, and often has a disproprionate impact upon the delivery of content.

A protocol for such information ought to expose a minimum of:

o Bandwidth - an approximation of the bandwidth currently unused on the "last mile" connection, in bits per second.

o Delay - an approximation of delay seen on the "last mile" connection, in milliseconds.

o Packet Loss - the current packet loss seen on "last mile" connections from the client, expressed as a percentage.

Additional metrics (including some that are operator-specific) might also be useful, and ought to be accommodated.

This information, in turn, could be used by Web servers, browsers and other tools to optimize both the responses and requests made. For example, MPEG-DASH clients could use the information about their own address to better choose an encoding; servers could re-encode images and HTML to account for slow networks, based upon the requesting client's IP address.

3. Granular Information about Networks: Straw-Men

NOTE: the technical mechanisms discussed are straw-men, and might not be the "real" approach. Readers are encouraged to consider and discuss the overall viability of the idea expressed above before focusing too much upon the details below.

3.1. DNS

One approach would be to using DNS [RFC1035] to convey this information. This has several advantages:

o DNS works at the granularity of an IP address

o Reverse DNS for a public IP address is often administered by the
   access network that provisions it

o DNS is lightweight and has a built-in caching mechanism

Potential disadvantages include:

o Servers receiving requests from clients that are unknown (or where
   there is only stale information available) will need to either
   wait for the lookup, or act without information for such requests

o Additional load on DNS infrastructure may be considerable

This would require a new RRTYPE to be defined to carry the
information outlined above.

3.2. HTTP

It might be possible to provide such information with a lightweight
HTTP [RFC7230] service exposed by the network operator. However,
discovery of that service would still need to be established; this
might be possible through DNS.

This approach's advantages include:

o Built-in caching and scaling mechanisms

o Rich extensibility

o Familiarity for developers and ops

Potential disadvantages include:

o Servers receiving requests from clients that are unknown (or where
   there is only stale information available) will need to either
   wait for the lookup, or act without information for such requests

o Comparatively high overhead

3.3. TLS

Another approach would be to add another channel in TLS [RFC5246]
that does not form part of teh encrypted session, to allow the
network to annotate connections directly.

This has a few advantages:

o Immediate availability of network information in-channel

o Direct binding to a single connection

o Annotations could be added on subsequent hops

However:

o Doing so is likely to be technically invasive, my have interop
   problems with deployed infrastructure

o May be seen as a layering violation / security issue

3.4. TCP

Yet another approach would be to define simliar side-channel
mechanisms in TCP [RFC0793].

The advantages and disadvantages of this approach are similar to
those around TLS; however, there is an additional disadvantage, in
that TCP extensibility is even more constrained than TLS'.

4. Security Considerations

This document is only exploratory now, but there are already clearly
evident security and privacy implications, including:

o Whether the information exposed can be used to identify a user

o Whether denial of service attacks are possible using this
   mechanism

5. References

5.1. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and
           specification", STD 13, RFC 1035, November 1987.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
           Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC7230] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol
           (HTTP/1.1): Message Syntax and Routing", RFC 7230, June
           2014.

5.2. Informative References

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC
           793, September 1981.

   [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security
              (TLS) Protocol Version 1.2", RFC 5246, August 2008.

   [bytemobile]
              Citrix, "ByteMobile Adaptive Traffic Management", 2014,
              <http://www.citrix.com/products/bytemobile-adaptive-
              traffic-management/overview.html>.

   [flashnet]
              Flash Networks, "Optimization Overview", 2014,
              <http://www.flashnetworks.com/Optimization-Overview>.

   [netinfo] W3C, "The Network Information API", 2014,
              <http://www.w3.org/TR/netinfo-api/>.

   [syniverse]
              Syniverse, "Hosted Data Optimization", 2014, <http://
              www.syniverse.com/products-services/product/Hosted-Data-
              Optimization>.

   [verizon] Verizon, "VERIZON WIRELESS OPTIMIZED VIEW FOR MOBILE WEB",
              2014, <http://www.vzwdevelopers.com/aims/downloads/
              wapoptout/
              Optimized_View_for_Mobile_Website_Developers_Guide.pdf>.

Author's Address

   Mark Nottingham

   Email: mnot@mnot.net
   URI: http://www.mnot.net/