

Network Working Group
Internet-Draft
Intended status: Informational
Expires: February 25, 2016

M. Nottingham

J. Hall
CDT
N. ten Oever
Article19
W. Seltzer
W3C
August 24, 2015

User Impact of Transport Metadata
draft-nottingham-transport-metadata-impact-00

Abstract

This draft attempts to identify potential impacts associated with new, extensible metadata facilities in Internet protocols, and suggests possible mitigations. Its goal is to have the discussion of these tradeoffs up-front, rather than after the development of such mechanisms.

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 February 25, 2016.

Copyright Notice

Copyright (c) 2015 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 (<http://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	2
2.	Potential Impact	3
2.1.	Security and Privacy	3
2.2.	Network Neutrality	4
3.	Possible Mitigations	4
3.1.	Constrained Vocabulary	4
3.2.	Transparency	5
4.	Security Considerations	5
5.	Informative References	5
	Authors' Addresses	7

[1.](#) Introduction

Recently, there has been an increasing amount of discussion in the IETF about adorning protocol flows with metadata about the network's state for consumption by applications, as well as that of the application in order to inform decisions in the network. For examples, see [[I-D.nottingham-gin](#)], [[I-D.sprecher-mobile-tg-exposure-req-arch](#)] and [[I-D.hildebrand-spud-prototype](#)].

These discussions are being at least partially motivated by the increasing use of encryption, both in deployment (thanks to the Snowden revelations) and standards (thanks in some part to [[RFC7258](#)], [[IAB-confidentiality](#)], and [[TAG-securing-web](#)]); while it's becoming widely accepted that networks don't have legitimate need to access the content of flows in most cases, they still wish to meet certain use cases that require more information.

For example, networks may wish to communicate their state to applications, so that link limitations and transient problems can be accounted for in applications, by doing things like degrading (or improving) video streaming quality.

Applications also need to give enough information to networks to

enable proper function; e.g., packets in UDP flows need to be associated to be able to cleanly transit NAT and firewalls. See [[I-D.trammell-stackevo-newtea](#)] and [[I-D.hardie-spud-use-cases](#)] for more discussion.

At the same time, it has been widely noted that "metadata" in various forms can be profoundly sensitive information, particularly when aggregated into large sets over extensive periods of time.

Indeed, much of the effort in combatting pervasive monitoring (as per [[RFC7258](#)]) has focused on minimizing metadata in existing, known protocols (such as TLS and HTTP).

Any new metadata facility, then - whether it be introduced to an existing protocol, or as part of a new one - needs to be carefully scrutinized and narrowly tailored to conservatively emit metadata. Of particular concern is an observed trend towards arbitrarily extensible metadata.

This draft attempts to identify potential impacts associated with new metadata facilities in Internet protocols, and suggest possible mitigations. Its goal is to initiate a discussion of these tradeoffs up-front, rather than waiting until after the development of such mechanisms.

Adding metadata to protocols is not an inherent harm - i.e., there are some legitimate uses of metadata, particularly if it eases the adoption of encrypted protocols or aligns well with both the interests of users and service or network operations, e.g., traffic management on mobile networks. However, the balance between the interests of constituents like end users, content providers and network operators needs to be carefully considered.

[2.](#) Potential Impact

[2.1.](#) Security and Privacy

It's been established [[Injection](#)] that many network operators inject HTTP headers into requests, in order to identify their customers using a unique identifier, thereby allowing "third-party advertisers and websites to assemble a deep, permanent profile of visitors' web

browsing habits without their consent." [\[X-UIDH\]](#)

In doing so, these networks are taking advantage of a relatively unconstrained extension point in the HTTP protocol - header fields. While HTTP header fields do require registration [\[RFC3864\]](#), the requirements are lax, and fields are often used without registration, because there is no technical enforcement of the requirements, due to HTTP's policy of ignoring unrecognized header fields [\[RFC7230\]](#).

HTTP header fields can be made a protected end-to-end facility by using HTTPS, avoiding the risk of such injection. A new transport

metadata facility that explicitly allows any node on the path to add arbitrary metadata cannot.

Well-intentioned metadata can also put the user at substantial risk without careful consideration. For example, if a Web browser "labels" flows based upon what they contain (e.g., "video", "image", "interactive"), an observer on the network path - including pervasive ones - can more effectively perform traffic analysis to determine what the user is doing. Similarly, metadata adornment might reveal sensitive information; for example the Server Name Indicator (SNI) in the TLS handshake would reveal if a web visitor intends to go to "bannedcontent.github.com" versus "kitties.github.com".

Standardizing an extensible transport metadata mechanism could also trigger various jurisdictions to define and require insertion of in-band metadata, an extension of current practices [\[AU-data-retention\]](#). While the IETF would not be directly responsible for such an outcome, it is notable that in the past we've explicitly said we won't serve conceptually similar use cases [\[RFC1984\]](#).

[2.2.](#) Network Neutrality

There is obvious potential for network neutrality impact from a mechanism that allows networks to communicate with endpoints about flows.

For example, if a network can instruct content servers to throttle back bandwidth available to users for video based upon a commercial arrangement (or lack thereof), the network can achieve their goals

without directly throttling traffic, thereby offering the potential to circumvent a regulatory regime that's designed to effect network neutrality.

While the IETF has not taken as firm a stance on network neutrality as it has for Pervasive Monitoring (for good reasons, since network neutrality problems are at their heart a sign of market failure, not a technical issue), new metadata facilities that enable existing regulatory regimes – thereby upsetting "the tussle" – must be carefully considered.

[3.](#) Possible Mitigations

[3.1.](#) Constrained Vocabulary

Much of the potential for harm above comes about because a transport-level metadata mechanism effectively becomes a side channel for arbitrary data, for use by any node on the path. The risks of

Nottingham, et al.

Expires February 25, 2016

[Page 4]

Internet-Draft

Transport Metadata Impact

August 2015

questionable use could be mitigated by constraining the data that's allowed in this side channel.

In other words, if the network doesn't have a means of inserting a unique identifier for customers, they won't be able to do so. If notification of constrained network conditions takes place using well-defined terms, regulatory regimes can be adjusted to achieve desired outcomes. And, information about application semantics can be carefully vetted for security considerations before being included in transport metadata.

One way to technically enforce such constraints would be to require nodes to silently drop non-standard metadata. Another would be to not make metadata extensible at all.

Naturally, this would constrain the ability of networks and applications to add new terms to metadata, thereby requiring more careful thought to go into the metadata that is standardised "up front."

[3.2.](#) Transparency

Many proposals for transport metadata assert that it will be encrypted, to improve security. While well-intentioned, it also creates an opaque side channel with a third party (the first and second being the endpoints).

The effect of of such designs should be carefully considered before standardisation; it may be that the community is better served by keeping this metadata "in the clear", albeit possibly with some form of authentication and integrity available (or required).

4. Security Considerations

This document describes security and privacy aspects of metadata adornment to internet protocols that protocol designers should consider.

5. Informative References

[AU-data-retention]

Keane, B., "Your guide to the data retention debate: what it is and why it's bad", March 2015, <http://www.crikey.com.au/2015/03/18/your-guide-to-the-data-retention-debate-what-it-is-and-why-it's-bad/>.

Nottingham, et al.

Expires February 25, 2016

[Page 5]

Internet-Draft

Transport Metadata Impact

August 2015

[I-D.hardie-spud-use-cases]

Hardie, T., "Use Cases for SPUD", [draft-hardie-spud-use-cases-01](#) (work in progress), February 2015.

[I-D.hildebrand-spud-prototype]

Hildebrand, J. and B. Trammell, "Substrate Protocol for User Datagrams (SPUD) Prototype", [draft-hildebrand-spud-prototype-03](#) (work in progress), March 2015.

[I-D.nottingham-gin]

Nottingham, M., "Granular Information about Networks", [draft-nottingham-gin-00](#) (work in progress), July 2014.

[I-D.sprecher-mobile-tg-exposure-req-arch]

Jain, A., Terzis, A., Sprecher, N., Swaminathan, S.,

Smith, K., and G. Klas, "Requirements and reference architecture for Mobile Throughput Guidance Exposure", [draft-sprecher-mobile-tg-exposure-req-arch-01](#) (work in progress), February 2015.

[I-D.trammell-stackevo-newtea]

Trammell, B., "Thoughts a New Transport Encapsulation Architecture", [draft-trammell-stackevo-newtea-01](#) (work in progress), May 2015.

[IAB-confidentiality]

Internet Architecture Board, "IAB Statement on Internet Confidentiality", November 2014, <<http://www.iab.org/2014/11/14/iab-statement-on-internet-confidentiality/>>.

[Injection]

Ammari, N., Björksten, G., Micek, P., and D. Olukotun, "The Rise of Mobile Tracking Headers: How Telcos Around the World Are Threatening Your Privacy", August 2015, <<https://www.accessnow.org/blog/2015/08/17/read-our-new-report-on-the-troubling-rise-of-tracking-headers-worldwide2>>.

[RFC1984] IAB and , "IAB and IESG Statement on Cryptographic Technology and the Internet", [RFC 1984](#), DOI 10.17487/[RFC1984](#), August 1996, <<http://www.rfc-editor.org/info/rfc1984>>.

[RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), DOI 10.17487/RFC3864, September 2004, <<http://www.rfc-editor.org/info/rfc3864>>.

[RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", [RFC 7230](#), DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), DOI 10.17487/RFC7258, May 2014, <<http://www.rfc-editor.org/info/rfc7258>>.

[TAG-securing-web]

W3C Technical Architecture Group, "Securing the Web",
January 2015, <<http://www.w3.org/2001/tag/doc/web-https>>.

[X-UIDH]

Hoffman-Andrews, J., "Verizon Injecting Perma-Cookies to
Track Mobile Customers, Bypassing Privacy Controls",
November 2014, <<https://www.eff.org/deeplinks/2014/11/verizon-x-uidh>>.

Authors' Addresses

Mark Nottingham

Email: mnot@mnot.net

URI: <https://www.mnot.net/>

Joseph Lorenzo Hall

CDT

Email: joe@cdt.org

URI: <https://cdt.org/>

Niels ten Oever

Article19

Email: niels@article19.org

URI: <https://www.article19.org/>

Wendy Seltzer

W3C

Email: wseltzer@w3.org

URI: <http://wendy.seltzer.org/>