

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 14, 2016

T. Hardie, Ed.
October 12, 2015

Design considerations for Metadata Insertion
draft-hardie-privsec-metadata-insertion-00

Abstract

The IAB has published [[RFC7624](#)] in response to several revelations of pervasive attack on Internet communications. In this document we consider the implications of protocol designs which associate metadata with encrypted flows.

In particular, we assert that designs which do so by explicit actions of the end system are preferable to designs in which middleboxes insert them.

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 April 14, 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

Internet-Draft

privsec-metadata-insertion

October 2015

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	2
3.	Design patterns	4
4.	Deployment considerations	5
5.	IANA Considerations	6
6.	Security Considerations	6
7.	Contributors {Contributors}	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	7
	Author's Address	8

[1.](#) Introduction

To ensure that the Internet can be trusted by users, it is necessary for the Internet technical community to address the vulnerabilities exploited in the attacks document in [[RFC7258](#)] and the threats described in [[RFC7624](#)]. The goal of this document is to address a common design pattern which emerges from the increase in encryption: explicit association of metadata which would previously have been inferred from the plaintext protocol.

[2.](#) Terminology

This document makes extensive use of standard security and privacy terminology; see [[RFC4949](#)] and [[RFC6973](#)]. Terms used from [[RFC6973](#)] include Eavesdropper, Observer, Initiator, Intermediary, Recipient, Attack (in a privacy context), Correlation, Fingerprint, Traffic Analysis, and Identifiability (and related terms). In addition, we use a few terms that are specific to the attacks discussed in this document. Note especially that "passive" and "active" below do not refer to the effort used to mount the attack; a "passive attack" is any attack that accesses a flow but does not modify it, while an "active attack" is any attack that modifies a flow. Some passive attacks involve active interception and modifications of devices, rather than simple access to the medium. The introduced terms are:

Pervasive Attack: An attack on Internet communications that makes

use of access at a large number of points in the network, or otherwise provides the attacker with access to a large amount of Internet traffic; see [[RFC7258](#)].

Passive Pervasive Attack: An eavesdropping attack undertaken by a pervasive attacker, in which the packets in a traffic stream between two endpoints are intercepted, but in which the attacker does not modify the packets in the traffic stream between two endpoints, modify the treatment of packets in the traffic stream (e.g. delay, routing), or add or remove packets in the traffic stream. Passive pervasive attacks are undetectable from the endpoints. Equivalent to passive wiretapping as defined in [[RFC4949](#)]; we use an alternate term here since the methods employed are wider than those implied by the word "wiretapping", including the active compromise of intermediate systems.

Active Pervasive Attack: An attack undertaken by a pervasive attacker, which in addition to the elements of a passive pervasive attack, also includes modification, addition, or removal of packets in a traffic stream, or modification of treatment of packets in the traffic stream. Active pervasive attacks provide more capabilities to the attacker at the risk of possible detection at the endpoints. Equivalent to active wiretapping as defined in [[RFC4949](#)].

Observation: Information collected directly from communications by an eavesdropper or observer. For example, the knowledge that <alice@example.com> sent a message to <bob@example.com> via SMTP taken from the headers of an observed SMTP message would be an observation.

Inference: Information derived from analysis of information collected directly from communications by an eavesdropper or observer. For example, the knowledge that a given web page was accessed by a given IP address, by comparing the size in octets of measured network flow records to fingerprints derived from known sizes of linked resources on the web servers involved, would be an inference.

Collaborator: An entity that is a legitimate participant in a

communication, and provides information about that communication to an attacker. Collaborators may either deliberately or unwittingly cooperate with the attacker, in the latter case because the attacker has subverted the collaborator through technical, social, or other means.

Key Exfiltration: The transmission of cryptographic keying material for an encrypted communication from a collaborator, deliberately or unwittingly, to an attacker.

Content Exfiltration: The transmission of the content of a communication from a collaborator, deliberately or unwittingly, to an attacker.

Data Minimization: With respect to protocol design, refers to the practice of only exposing the minimum amount of data or metadata necessary for the task supported by that protocol to the other endpoint(s) and/or devices along the path.

3. Design patterns

One of the core mitigations for the loss of confidentiality in the presence of pervasive surveillance is data minimization, which limits the amount of data disclosed to those elements absolutely required to complete the relevant protocol exchange. When data minimization is in effect, some information which was previously available may be removed from specific protocol exchanges. The information may be removed explicitly (by a browser suppressing cookies during private modes, as an example) or by other means. As noted in [[RFC7624](#)], some topologies which aggregate or alter the network path also acted to reduce the ease with which metadata is available to eavesdroppers.

In some cases, other actors within a protocol context will continue to have access to the information which has been thus withdrawn from specific protocol exchanges. If those actors attach the information as metadata to those protocol exchange, the confidentiality effect of data minimization is lost.

The restoration of information is particularly tempting for systems

whose primary function is not to provide confidentiality. A proxy providing compression, for example, may wish to restore the identity of the requesting party; similarly a VPN system used to provide channel security may believe that origin IP should be restored. Actors considering restoring metadata may believe that they understand the relevant privacy considerations or believe that, because the primary purpose of the service was not privacy-related, none exist. Examples of this design pattern include [[RFC7239](#)] and [[I-D.ietf-dnsop-edns-client-subnet](#)].

While it may be tempting to restore previous information flows, this design pattern should be avoided, as it contributes to the overall loss of confidentiality for the Internet. This document recommends against restoration in these cases unless a positive affirmation of approval for restoration has been received from the actor whose data will be added. In general, we recommend that the actor add such metadata themselves so that it flows end-to-end, rather than requiring the action of other parties. Where this is not possible, opt-in methods for consent are strongly recommended; opt-out systems,

especially for previously deployed systems, may provide sufficient targeting that the most vulnerable users would be reluctant to employ them.

[4.](#) Deployment considerations

There are two common tensions associated with the deployment of systems which restore metadata. The first is the trade-off in speed of deployment for different actors. The "Forward-for" method cited above provides an example of this. When used with a proxy, Forwarded-for restores the original identity of the requesting party, thus allowing a responding server to tailor responses according to the original party's region, network, or other characteristics associated with the identity. It would, of course, be possible for the originating client to add this data itself, using STUN [[RFC5389](#)] or a similar mechanism to first determine the identity to declare. This would require, however, full specification and adoption of this mechanism by the end systems. It would not be available at all during this period, and would thereafter be limited to those systems which have been upgraded to include it. The long tail of browser deployments indicates that many systems might go without upgrades for a significant period of time. The proxy infrastructure, in contrast,

is commonly under more active management and represents a much smaller number of elements; this impacts both the general deployment difficulty and the number of systems which the origin server must trust.

The second common tension is between the metadata minimization and the desire to tailor content responses. For origin servers whose content is common across users, the loss of metadata may have limited impact on the system's functioning. For other systems, which commonly tailor content by region or network, the loss of metadata may imply a loss of functionality. Where the user desires this functionality, restoration can commonly be achieved by the use of other identifiers or login procedures. Where the user does not desire this functionality, but it is a preference of the server or a third party, adjustment is more difficult. At the extreme, content blocking by network origin may be a regulatory requirement. Trusting a network intermediary to provide accurate data is, of course, fragile in this case, but it may be a part of the regulatory framework.

These tensions do not change the basic recommendation, but they suggest that the parties who are introducing encryption and data minimization for existing protocols consider carefully whether the work also implies introducing mechanisms for the end-to-end provisioning of metadata when a user has actively consented to provide it.

[5.](#) IANA Considerations

This memo makes no request of IANA.

[6.](#) Security Considerations

This memorandum describes a design pattern related emerging from responses to the attacks described in [[RFC7258](#)]. Continued use of this design pattern lowers the impact of mitigations to that attack.

[7.](#) Contributors {Contributors}

This document is derived in part from the work initially done on the Perpass mailing list and at the STRINT workshop. It has been discussed with the IAB's Privacy and Security program, whose review

is gratefully acknowledged.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, [RFC 4949](#), DOI 10.17487/RFC4949, August 2007, <<http://www.rfc-editor.org/info/rfc4949>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", [RFC 6973](#), DOI 10.17487/RFC6973, July 2013, <<http://www.rfc-editor.org/info/rfc6973>>.
- [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>>.
- [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, <<http://www.rfc-editor.org/info/rfc7624>>.

Hardie

Expires April 14, 2016

[Page 6]

Internet-Draft

privsec-metadata-insertion

October 2015

8.2. Informative References

- [I-D.ietf-dnsop-edns-client-subnet] Contavalli, C., Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", [draft-ietf-dnsop-edns-client-subnet-04](#) (work in progress), September 2015.
- [RFC2015] Elkins, M., "MIME Security with Pretty Good Privacy

(PGP)", [RFC 2015](#), DOI 10.17487/RFC2015, October 1996, <<http://www.rfc-editor.org/info/rfc2015>>.

- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), DOI 10.17487/RFC4301, December 2005, <<http://www.rfc-editor.org/info/rfc4301>>.
- [RFC4306] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", [RFC 4306](#), DOI 10.17487/RFC4306, December 2005, <<http://www.rfc-editor.org/info/rfc4306>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", [RFC 5389](#), DOI 10.17487/RFC5389, October 2008, <<http://www.rfc-editor.org/info/rfc5389>>.
- [RFC5750] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Certificate Handling", [RFC 5750](#), DOI 10.17487/RFC5750, January 2010, <<http://www.rfc-editor.org/info/rfc5750>>.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", [RFC 6698](#), DOI 10.17487/RFC6698, August 2012, <<http://www.rfc-editor.org/info/rfc6698>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", [RFC 6962](#), DOI 10.17487/RFC6962, June 2013, <<http://www.rfc-editor.org/info/rfc6962>>.
- [RFC7239] Petersson, A. and M. Nilsson, "Forwarded HTTP Extension", [RFC 7239](#), DOI 10.17487/RFC7239, June 2014, <<http://www.rfc-editor.org/info/rfc7239>>.

<https://www.w3.org/2014/strint/draft-iab-strint-report.html>.

Author's Address

Ted Hardie (editor)

Email: ted.ietf@gmail.com