Network Working Group	A. Cooper
Internet-Draft	CDT
Intended status: Informational	H. Tschofenig
Expires: April 25, 2012	Nokia Siemens Networks
	B. Aboba
	Microsoft Corporation
	J. Peterson
	NeuStar, Inc.
	J. Morris
	October 23, 2011

Privacy Considerations for Internet Protocols
draft-iab-privacy-considerations-00.txt

Abstract

This document offers guidance for developing privacy considerations for
IETF documents and aims to make protocol designers aware of privacy-
related design choices.
Discussion of this document is taking place on the IETF Privacy
Discussion mailing list (see https://www.ietf.org/mailman/listinfo/
ietf-privacy).

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 25, 2012.

Copyright Notice

Table of Contents

1. Introduction

All IETF specifications are required by [RFC2223] to contain a security considerations section. [RFC3552] provides detailed guidance to protocol designers about both how to consider security as part of protocol design and how to inform readers of IETF documents about security issues. This document intends to provide a similar set of guidance for considering privacy in protocol design. Whether any individual document will require a specific privacy considerations section will depend on the document's content. The guidance provided here can and should be used to assess the privacy considerations of protocol and architectural specifications regardless of whether those considerations are documented in a stand-alone section.

Privacy is a complicated concept with a rich history that spans many disciplines. Many sets of privacy principles and privacy design frameworks have been developed in different forums over the years. These include the Fair Information Practices (FIPs) [OECD], a baseline set of privacy protections pertaining to the collection and use of data about individuals, and the Privacy by Design framework [PbD], which provides high-level privacy guidance for systems design. The guidance provided in this document is inspired by this prior work, but it aims to be more concrete, pointing protocol designers to specific engineering choices that can impact the privacy of the individuals that make use of Internet protocols.

Privacy as a legal concept is understood differently in different jurisdictions. The guidance provided in this document is generic and can be used to inform the design of any protocol to be used anywhere in the world, without reference to specific legal frameworks.

The document is organized as follows: Section 2 describes the extent to which the guidance offered in this document is applicable within the IETF, Section 3 discusses a generic threat model to motivate the need for privacy considerations, Section 4 provides the guidelines for analyzing and documenting privacy considerations within IETF specifications, and Section 5 examines the privacy characteristics of an IETF protocol to demonstrate the use of the guidance framework.

2. Scope

The core function of IETF activity is building protocols. Internet protocols are often built flexibly, making them useful in a variety of architectures, contexts, and deployment scenarios without requiring significant interdependency between disparately designed components. Although some protocols assume particular architectures at design time, it is not uncommon for architectural frameworks to develop later, after implementations exist and have been deployed in combination with other protocols or components to form complete systems.
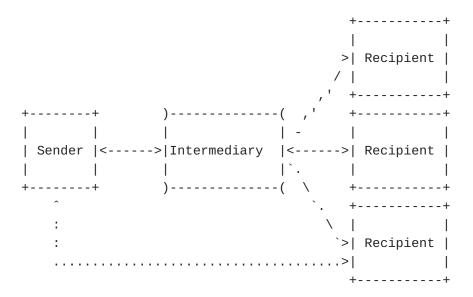
As a consequence, the extent to which protocol designers can foresee all of the privacy implications of a particular protocol at design time is significantly limited. An individual protocol may be relatively benign on its own, but when deployed within a larger system or used in a way not envisioned at design time, its use may create new privacy risks. The guidelines in Section 4 ask protocol designers to consider how their protocols are expected to interact with systems and information that exist outside the protocol bounds, but not to imagine every possible deployment scenario.

Furthermore, in many cases the privacy properties of a system are dependent upon API specifics, internal application functionality, database structure, local policy, and other details that are specific to particular instantiations and generally outside the scope of the work conducted in the IETF. The guidance provided here only reaches as far as protocol design can go.

As an example, consider HTTP [RFC2616], which was designed to allow the exchange of arbitrary data. A complete analysis of the privacy considerations for uses of HTTP might include what type of data is exchanged, how this data is stored, and how it is processed. Hence the analysis for an individual's static personal web page would be different than the use of HTTP for exchanging health records. A protocol designer working on HTTP extensions (such as WebDAV [RFC4918]) is not expected to describe the privacy risks derived from all possible usage scenarios, but rather the privacy properties specific to the extensions and any particular uses of the extensions that are expected and foreseen at design time.

3. Threat Model

To consider privacy in protocol design it is helpful to consider the overall communication architecture and different actors' roles within it. This analysis is similar to a threat analysis found in the security considerations sections of IETF documents. Figure 1 presents a communication model found in many of today's protocols where a sender wants to establish communication with some recipient and thereby uses some form of intermediary. In some cases this intermediary stays in the communication path for the entire duration of the communication and sometimes it is only used for communication establishment, for either inbound or outbound communication. In rare cases there may be a series of intermediaries that are traversed.

```
                                       +-----------+
                                       |           |
                                      >| Recipient |
                                     / |           |
                                   ,'  +-----------+
+--------+        )--------------(  ,'    +-----------+
|        |        |              | -      |           |
| Sender |<------>|Intermediary  |<------>| Recipient |
|        |        |              |`.      |           |
+--------+        )--------------(  \     +-----------+
    ^                              `.   +-----------+
    :                               \ |           |
    :                               `>| Recipient |
    ......................................>|           |
                                       +-----------+
```

Legend:

<....> End-to-End Communication
<----> Hop-by-Hop Communication

This model is vulnerable to three types of adversaries:

Eavesdropper:

> [RFC4949] describes the act of 'eavesdropping' as

> > *"Passive wiretapping done secretly, i.e., without the
> > knowledge of the originator or the intended recipients of the
> > communication."

> Eavesdropping is often considered by IETF protocols in the context
> of a security analysis. Confidentiality protection is often employed
> to defend against attacks based on eavesdropping, and
> [RFC3552] demands that confidentiality be incorporated as a security
> consideration. While IETF protocols offer guidance on how to secure
> communication against eavesdroppers, deployments sometimes choose
> not to enable such security.

Intermediary: Many protocols developed today show a more complex
communication pattern than simple client-server or peer-to-peer
communication, as motivated in Figure 1. Store-and-forward protocols
are examples where entities participate in the message delivery even
though they are not the final recipients. Often, these
intermediaries only require a small amount of information for
message routing and/or security. In theory, protocol mechanisms
could ensure that end-to-end information is not made accessible to
these entities, but in practice the difficulty of deploying end-to-
end security procedures, additional messaging or computational
overhead, and other business or legal requirements often slow or
prevent the deployment of end-to-end security mechanisms, giving
intermediaries greater exposure to communication patterns and
payloads than is strictly necessary.

Recipient: It may not seem intuitive to treat the recipient as an
adversary since the entire purpose of the communication interaction
is to provide information to the recipient. However, the recipient
can act as the sender's privacy foe in two respects. First, the
sender may be unintentionally communicating with the recipient,
whether because of a lack of access control or because the sender
was not properly informed about what data it would be communicating
to the recipient. Second, the recipient may choose to use the
sender's data in ways that contravene the sender's wishes, whether
by putting it to some purpose that the sender opposes, sharing it
with other entities, or storing it after the communication session
has ended. Whether the recipient becomes an adversary depends on
whether it makes use of mechanisms that reduce these risks,
including informing the sender about how his or her data will be
used, offering choices, and obtaining authorization to receive and
use the sender's data.

4. Guidelines

This section provides guidance for document authors in the form of a questionnaire about a protocol being designed. The questionnaire may be useful at any point in the design process, particularly after document authors have developed a high-level protocol model as described in [RFC4101].
Note that the guidance does not recommend specific practices. The range of protocols developed in the IETF is too broad to make recommendations about particular uses of data or how privacy might be balanced against other design goals. However, by carefully considering the answers to each question, document authors should be able to produce a comprehensive analysis that can serve as the basis for discussion of whether the protocol adequately protects against privacy threats.
The framework is divided into four sections that address different aspects of privacy -- data minimization, user participation, security, and accountability -- plus a general section. Security is not elaborated since substantial guidance already exists in [RFC3552]. Privacy-specific terminology used in the framework is [will be] defined in [I-D.hansen-privacy-terminology].

4.1. General

   a. Trade-offs. Does the protocol make trade-offs between privacy and usability, privacy and efficiency, privacy and implementability, or privacy and other design goals? Describe the trade-offs and the rationale for the design chosen.

4.2. Data Minimization

   a. Identifiers. What identifiers does the protocol use for distinguishing endpoints? Does the protocol use identifiers that allow different protocol interactions to be correlated?

   b. User information. What information does the protocol expose about end users and/or their devices (other than the identifiers discussed in (a))? How identifiable is this information? How does the protocol combine user information with the identifiers discussed in (a)?

   c. Fingerprinting. In many cases the specific ordering and/or occurrences of information elements in a protocol allow devices using the protocol to be uniquely fingerprinted. Is this protocol vulnerable to fingerprinting? If so, how?

   d. Persistence of identifiers. What assumptions are made in the protocol design about the lifetime of the identifiers discussed in (a)? Does the protocol allow implementers or users to delete or

recycle identifiers? How often does the specification recommend to delete or recycle identifiers by default?

e. Leakage. Are there expected ways that information exposed by the protocol will be combined or correlated with information obtained outside the protocol? How will such combination or correlation facilitate user or device fingerprinting? Are there expected combinations or correlations with outside data that will make the information exposed by the protocol more identifiable?

f. Recipients. In the protocol design, what information discussed in (a) and (b) is exposed to other endpoints (i.e., recipients)? Are there ways for protocol implementers to choose to limit the information shared with other endpoints?

g. Intermediaries. In the protocol design, what information discussed in (a) and (b) is exposed to intermediaries? Are there ways for protocol implementers to choose to limit the information shared with intermediaries?

h. Retention. Do the protocol or its anticipated uses require that the information discussed in (a) or (b) be retained by recipients or intermediaries? Is the retention expected to be persistent or temporary?

4.3. User Participation

a. Control over initial sharing. What user controls or consent mechanisms does the protocol define or require before user information or identifiers are shared or exposed via the protocol? If no such mechanisms are specified, is it expected that control and consent will be handled outside of the protocol?

b. Control over sharing with recipients. Does the protocol provide ways for users to limit which information is shared with recipients? If not, are there mechanisms that exist outside of the protocol to provide users with such control?

c. Control over sharing with intermediaries. Does the protocol provide ways for users to limit which information is shared with intermediaries? If not, are there mechanisms that exist outside of the protocol to provide users with such control? Is it expected that users will have relationships (contractual or otherwise) with intermediaries that govern the use of the information?

d. Preference expression. Does the protocol provide ways for users to express their preferences to recipients or intermediaries with regard to the use or disclosure of their information?

4.4. Security

a. Communication security. Do the protocol's security considerations account for communication security, per RFC 3552?

4.5. Accountability

a. User verification. If the protocol provides for user preference expression, does it also define or require mechanisms that allow users to verify that their preferences are being honored? If not, are there mechanisms that exist outside of the protocol that allow for user verification?

5. Example

[To be provided in a future version.]

6. Security Considerations

This document describes privacy aspects that protocol designers should consider in addition to regular security analysis.

7. IANA Considerations

This document does not require actions by IANA.

8. Acknowledgements

We would like to thank the participants for the feedback they provided during the December 2010 Internet Privacy workshop co-organized by MIT, ISOC, W3C and the IAB.

9. References

9.1. Normative References

[I-D.hansen-privacy-terminology]	Hansen, M, Tschofenig, H and R Smith, "Privacy Terminology", Internet-Draft draft-hansen-privacy-terminology-03, October 2011.

9.2. Informative References

[RFC2223]	Postel, J. and J.K. Reynolds, "Instructions to RFC Authors", RFC 2223, October 1997.
[RFC2616]	Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P. and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
[RFC3552]	Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.
[RFC4101]	

	Rescorla, E., IAB, "Writing Protocol Models", RFC 4101, June 2005.
[RFC4918]	Dusseault, L., "HTTP Extensions for Web Distributed Authoring and Versioning (WebDAV)", RFC 4918, June 2007.
[RFC4949]	Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.
[OECD]	, , "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data", available at (September 2010) , http://www.oecd.org/EN/document/0,,EN-document-0-nodirectorate-no-24-10255-0,00.html, 1980.
[browser-fingerprinting]	Eckersley, P., "How Unique Is Your Browser?", Springer Lecture Notes in Computer Science , Privacy Enhancing Technologies Symposium (PETS 2010), 2010.
[PbD]	Office of the Information and Privacy Commissioner, Ontario, Canada, "Privacy by Design", 2011.

Authors' Addresses

Alissa Cooper Cooper CDT 1634 Eye St. NW, Suite 1100 Washington, DC 20006 US Phone: +1-202-637-9800 EMail: acooper@cdt.org URI: http://www.cdt.org/

Hannes Tschofenig Tschofenig Nokia Siemens Networks Linnoitustie 6 Espoo, 02600 Finland Phone: +358 (50) 4871445 EMail: Hannes.Tschofenig@gmx.net URI: http://www.tschofenig.priv.at

Bernard Aboba Aboba Microsoft Corporation One Microsoft Way Redmond, WA 98052 US EMail: bernarda@microsoft.com

Jon Peterson Peterson NeuStar, Inc. 1800 Sutter St Suite 570 Concord, CA 94520 US EMail: jon.peterson@neustar.biz

John B. Morris, Jr. Morris EMail: ietf@jmorris.org