

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 13, 2012

A. Cooper
CDT
H. Tschofenig
Nokia Siemens Networks
B. Aboba
Microsoft Corporation
J. Peterson
NeuStar, Inc.
J. Morris
March 12, 2012

Privacy Considerations for Internet Protocols
draft-iab-privacy-considerations-02.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 September 13, 2012.

Copyright Notice

Copyright (c) 2012 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

Internet-Draft

Privacy Considerations

March 2012

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	3
2.	Scope	4
3.	Internet Privacy Threat Model	5
3.1.	Communications Model	5
3.2.	Privacy Threats	6
3.2.1.	Combined Security-Privacy Threats	6
3.2.2.	Privacy-Specific Threats	8
4.	Internet Privacy Goals	11
4.1.	Data Minimization	11
4.2.	User Participation	12
4.3.	Accountability	12
4.4.	Security	12
5.	Guidelines	13
5.1.	General	13
5.2.	Data Minimization	13
5.3.	User Participation	14
5.4.	Accountability	15
5.5.	Security	15
6.	Example	16
7.	Glossary	17
8.	Security Considerations	20
9.	IANA Considerations	21
10.	Acknowledgements	22
11.	References	23
11.1.	Normative References	23
11.2.	Informative References	23
	Authors' Addresses	25

1. Introduction

[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. Documents whose entire focus is privacy may not merit a separate section (for example, [RFC3325]). For certain specifications, privacy considerations are a subset of security considerations and can be discussed explicitly in the security considerations section. The guidance provided here can and should be used to assess the privacy considerations of protocol, architectural, and operational specifications and to decide whether those considerations are to be documented in a stand-alone section, within the security considerations section, or throughout the document.

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), a baseline set of privacy protections pertaining to the collection and use of data about individuals (see [OECD] for one example), and the Privacy by Design concept, which provides high-level privacy guidance for systems design (see [PbD] for one example). 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.

This document is organized as follows. [Section 2](#) describes the extent to which the guidance offered is applicable within the IETF. [Section 3](#) discusses threats to privacy as they apply to Internet protocols. [Section 4](#) outlines privacy goals. [Section 5](#) provides the guidelines for analyzing and documenting privacy considerations within IETF specifications. [Section 6](#) examines the privacy characteristics of an IETF protocol to demonstrate the use of the guidance framework. [Section 7](#) provides a concise glossary of terms used in this document, with a more complete discussion of some of the terms available in [[I-D.iab-privacy-terminology](#)].

[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 5](#) 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 may be useful in making choices about those details, but its primary aim is to assist with the design, implementation, and operation of protocols.

Privacy issues, even those related to protocol development, go beyond the technical guidance discussed herein.

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.](#) Internet Privacy Threat Model

Privacy harms come in a number of forms, including harms to financial standing, reputation, solitude, autonomy, and safety. A victim of identity theft or blackmail, for example, may suffer a financial loss as a result. Reputational harm can occur when disclosure of information about an individual, whether true or false, subjects that individual to stigma, embarrassment, or loss of personal dignity. Intrusion or interruption of an individual's life or activities can harm the individual's ability to be left alone. When individuals or their activities are monitored, exposed, or at risk of exposure, those individuals may be stifled from expressing themselves, associating with others, and generally conducting their lives freely. In cases where such monitoring is for the purpose of stalking or violence, it can put individuals in physical danger.

This section lists common privacy threats (drawing liberally from [[Solove](#)]), showing how each of them may cause individuals to incur privacy harms and providing examples of how these threats can exist on the Internet.

[3.1.](#) Communications Model

To understand attacks in the privacy-harm sense, it is helpful to consider the overall communication architecture and different actors' roles within it. Consider a protocol element that initiates communication with some recipient (an "initiator"). Privacy analysis is most relevant for protocols with use cases in which the initiator acts on behalf of a natural person (or different people at different times). It is this natural person -- the data subject -- whose privacy is potentially threatened.

Communications may be direct between the initiator and the recipient, or they may involve an intermediary (such as a proxy or cache) that is necessary for the two parties to communicate. 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.

Some communications tasks require multiple protocol interactions with different entities. For example, a request to an HTTP server may be preceded by an interaction between the initiator and an Authentication, Authorization, and Accounting (AAA) server or DNS resolver. In this case, the HTTP server is the recipient and the other entities are enablers of the initiator-to-recipient communication. Similarly, a single communication with the recipient may generate further protocol interactions between either the

initiator or the recipient and other entities. For example, an HTTP request might trigger interactions with an authentication server or with other resource servers.

As a general matter, recipients, intermediaries, and enablers are usually assumed to be authorized to receive and handle data from initiators. As [[RFC3552](#)] explains, "we assume that the end-systems engaging in a protocol exchange have not themselves been compromised."

Although they may not generally be considered as attackers, recipients, intermediaries, and enablers may all pose privacy threats (depending on the context) because they are able to observe and collect privacy-relevant data. These entities are collectively described below as "observers" to distinguish them from traditional

attackers. From a privacy perspective, one important type of attacker is an eavesdropper: an entity that passively observes the initiator's communications without the initiator's knowledge or authorization.

The threat descriptions in the next section explain how observers and attackers might act to harm data subjects' privacy. Different kinds of attacks may be feasible at different points in the communications path. For example, an observer could mount surveillance or identification attacks between the initiator and intermediary, or instead could surveil an enabler (e.g., by observing DNS queries from the initiator).

[3.2.](#) Privacy Threats

Some privacy threats are already considered in IETF protocols as a matter of routine security analysis. Others are more pure privacy threats that existing security considerations do not usually address. The threats described here are divided into those that may also be considered security threats and those that are primarily privacy threats.

Note that an individual's knowledge and authorization of the practices described below can greatly affect the extent to which they threaten privacy. If a data subject authorizes surveillance of his own activities, for example, the harms associated with it may be significantly mitigated.

[3.2.1.](#) Combined Security-Privacy Threats

[3.2.1.1.](#) Surveillance

Surveillance is the observation or monitoring of an individual's communications or activities. The effects of surveillance on the individual can range from anxiety and discomfort to behavioral changes such as inhibition and self-censorship to the perpetration of violence against the individual. The individual need not be aware of the surveillance for it to impact privacy -- the possibility of

surveillance may be enough to harm individual autonomy.

Surveillance can be conducted by observers or eavesdroppers at any point along the communications path. Confidentiality protections (as discussed in [\[RFC3552\] Section 3](#)) are necessary to prevent surveillance of the content of communications. To prevent traffic analysis or other surveillance of communications patterns, other measures may be necessary, such as [\[Tor\]](#).

[3.2.1.2.](#) Stored Data Compromise

End systems that do not take adequate measures to secure stored data from unauthorized or inappropriate access expose individuals to potential financial, reputational, or physical harm.

By and large, protecting against stored data compromise is outside the scope of IETF protocols. However, a number of common protocol functions -- key management, access control, or operational logging, for example -- require the storage of data about initiators of communications. When requiring or recommending that information about initiators or their communications be stored or logged by end systems (see, e.g., [RFC 6302](#)), it is important to recognize the potential for that information to be compromised and for that potential to be weighed against the benefits of data storage. Any recipient, intermediary, or enabler that stores data may be vulnerable to compromise.

[3.2.1.3.](#) Intrusion

Intrusion consists of invasive acts that disturb or interrupt one's life or activities. Intrusion can thwart individuals' desires to be let alone, sap their time or attention, or interrupt their activities.

Unsolicited mail and denial-of-service attacks are the most common types of intrusion on the Internet. Intrusion can be perpetrated by any attacker that is capable of sending unwanted traffic to the initiator.

[3.2.2.](#) Privacy-Specific Threats

[3.2.2.1.](#) Correlation

Correlation is the combination of various pieces of information about an individual. Correlation can defy people's expectations of the limits of what others know about them. It can increase the power that those doing the correlating have over individuals as well as correlators' ability to pass judgment, threatening individual autonomy and reputation.

Correlation is closely related to identification. Internet protocols can facilitate correlation by allowing data subjects' activities to be tracked and combined over time. The use of persistent or infrequently refreshed identifiers at any layer of the stack can facilitate correlation. For example, an initiator's persistent use of the same device ID, certificate, or email address across multiple interactions could allow recipients to correlate all of the initiator's communications over time.

In theory any observer or attacker that receives an initiator's communications can engage in correlation. The extent of the potential for correlation will depend on what data the entity receives from the initiator and has access to otherwise. Often, 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 initiators' data than is strictly necessary.

[3.2.2.2.](#) Identification

Identification is the linking of information to a particular individual. In some contexts it is perfectly legitimate to identify individuals, whereas in others identification may potentially stifle individuals' activities or expression by inhibiting their ability to be anonymous or pseudonymous. Identification also makes it easier for individuals to be explicitly controlled by others (e.g., governments).

Many protocol identifiers, such as those used in SIP or XMPP, may allow for the direct identification of data subjects. Protocol identifiers may also contribute indirectly to identification via correlation. For example, a web site that does not directly authenticate users may be able to match its HTTP header logs with

logs from another site that does authenticate users, rendering users on the first site identifiable.

As with correlation, any observer or attacker may be able to engage in identification depending on the information about the initiator that is available via the protocol mechanism or other channels.

[3.2.2.3.](#) Secondary Use

Secondary use is the use of collected information without the data subject's consent for a purpose different from that for which the information was collected. Secondary use may violate people's expectations or desires. The potential for secondary use can generate uncertainty over how one's information will be used in the future, potentially discouraging information exchange in the first place.

One example of secondary use would be a network access server that uses an initiator's access requests to track the initiator's location. Any observer or attacker could potentially make unwanted secondary uses of initiators' data.

[3.2.2.4.](#) Disclosure

Disclosure is the revelation of truthful information about a person that affects the way others judge the person. Disclosure can violate people's expectations of the confidentiality of the data they share. The threat of disclosure may deter people from engaging in certain activities for fear of reputational harm.

Any observer or attacker that receives data about an initiator may choose to engage in disclosure. In most cases, there is nothing done at the protocol level to influence or limit disclosure, although there are some exceptions. For example, the GEOPRIV architecture [[RFC6280](#)] provides a way for users to express a preference that their location information not be disclosed beyond the intended recipient.

[3.2.2.5.](#) Exclusion

Exclusion is the failure to allow the data subject to know about the data that others have about him or her and to participate in its handling and use. Exclusion reduces accountability on the part of entities that maintain information about people and creates a sense of vulnerability about individuals' ability to control how information about them is collected and used.

The most common way for Internet protocols to be involved in limiting exclusion is through access control mechanisms. The presence

Cooper, et al.

Expires September 13, 2012

[Page 9]

Internet-Draft

Privacy Considerations

March 2012

architecture developed in the IETF is a good example where data subjects are included in the control of information about them. Using a rules expression language (e.g., Presence Authorization Rules [[RFC5025](#)]), presence clients can authorize the specific conditions under which their presence information may be shared.

Exclusion is primarily considered problematic when the recipient fails to involve the initiator in decisions about data collection, handling, and use. Eavesdroppers engage in exclusion by their very nature since their data collection and handling practices are covert.

[4.](#) Internet Privacy Goals

Privacy is notoriously difficult to measure and quantify. The extent to which a particular protocol, system, or architecture "protects" or "enhances" privacy is dependent on a large number of factors relating to its design, use, and potential misuse. However, there are certain widely recognized privacy properties against which designs may be assessed for their potential to impact privacy. This section adapts these properties into four privacy goals for Internet protocols: (1) data minimization, (2) user participation, (3) accountability, and (4) security.

[4.1.](#) Data Minimization

Data minimization refers to collecting, using, and storing the minimal data necessary to perform a task. The less data about data subjects that gets exchanged in the first place, the lower the chances of that data being used for privacy invasion.

Data minimization is comprised of a number of mutually exclusive sub-goals:

- o Use limitation: Limiting the uses to which data is put helps contain the spread of data to third parties and protects against uses that may violate data subjects' expectations.
- o Retention limitation: Limiting the duration of data storage reduces the risk of stored data compromise.
- o Identifiability limitation: Minimization pertains not only to the amount of data exchanged, but also the extent to which it can be used to identify data subjects. Reducing the identifiability of

data by using pseudonymous or anonymous identifiers helps to weaken the link between a data subject and his or her communications. Refreshing or recycling identifiers reduces the possibility that multiple protocol interactions or communications can be correlated back to the same data subject.

- o Sensitivity limitation: The sensitivity of data is another property that can be minimized. For example, the street address of a building that an individual visits may be considered to be a more sensitive piece of information than the city and postal code in which the building is located. Collecting, using, and storing less sensitive data may mitigate the damage caused by secondary use, disclosure, stored data compromise, and correlation.

Cooper, et al.

Expires September 13, 2012

[Page 11]

Internet-Draft

Privacy Considerations

March 2012

[4.2.](#) User Participation

As explained in [Section 3.2.2.5](#), data collection and use that happens "in secret," without the data subject's knowledge, is apt to violate the subject's expectation of privacy and may create incentives for misuse of data. As a result, privacy regimes tend to include provisions to support informing data subjects about data collection and use and involving them in decisions about the treatment of their data. In an engineering context, supporting the goal of user participation usually means providing ways for users to control the data that is shared about them.

[4.3.](#) Accountability

An entity that collects, uses, or stores data can undergird its commitments to the other privacy goals by providing mechanisms by which data subjects and third parties can hold the entity accountable for those commitments. These mechanisms usually allow for verification of what data is collected or stored and with whom it is shared, again helping to mitigate the threat of exclusion.

[4.4.](#) Security

Keeping data secure at rest and in transit is another important component of privacy protection. As they are described in [\[RFC3552\]](#)

Section 2, a number of security goals also serve to enhance privacy:

- o Confidentiality: Keeping data secret from unintended listeners.
- o Peer entity authentication: Ensuring that the endpoint of a communication is the one that is intended (in support of maintaining confidentiality).
- o Unauthorized usage: Limiting data access to only those users who are authorized, helping to prevent stored data compromise.
- o Inappropriate usage: Limiting how authorized users can use data. (Note that this goal also falls within data minimization.)

5. 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 each of the

goals from [Section 4](#), plus a general section. Security is not fully elaborated since substantial guidance already exists in [[RFC3552](#)].

[5.1.](#) 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.

[5.2.](#) Data Minimization

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

b. Personal data. What information does the protocol expose about data subjects and/or their devices (other than the identifiers discussed in (a))? To what extent is this information linked to the identities of data subjects? How does the protocol combine personal data with the identifiers discussed in (a)?

c. Observers. Which information discussed in (a) and (b) is exposed to each other protocol entity (i.e., recipients, intermediaries, and enablers)? Are there ways for protocol implementers to choose to limit the information shared with each entity? Are there operational controls available to limit the information shared with each entity?

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

e. 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?

f. Correlation. 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 fingerprinting of a user, device, or application? Are there expected combinations or correlations with outside data that will make users of the protocol more identifiable?

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

5.3. User Participation

a. User control. What controls or consent mechanisms does the protocol define or require before personal data 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 individual recipients. Does the protocol provide ways for initiators to share different information with different recipients? If not, are there mechanisms that exist outside of the protocol to provide initiators with such control?

c. Control over sharing with intermediaries. Does the protocol provide ways for initiators 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 initiators to express data subjects' preferences to recipients or intermediaries with regard to the use or disclosure of their personal data?

5.4. Accountability

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

5.5. Security

a. Surveillance. How do the protocol's security considerations prevent surveillance, including eavesdropping and traffic analysis?

b. Stored data compromise. How do the protocol's security considerations prevent or mitigate stored data compromise?

c. Intrusion. How do the protocol's security considerations prevent or mitigate intrusion, including denial-of-service attacks and unsolicited communications more generally?

[6.](#) Example

[To be provided in a future version once the guidance is settled.]

[7.](#) Glossary

\$ Anonymity

The state of being anonymous. See [[I-D.iab-privacy-terminology](#)].

\$ Anonymous

A property of a data subject in which an observer or attacker cannot identify the data subject within a set of other subjects (the anonymity set).

\$ Attacker

An entity that intentionally works against some protection goal.

\$ Attribute

A property of a data subject or initiator.

\$ Correlation

The combination of various pieces of information about a data subject.

\$ Data Subject

An identified natural person or a natural person who can be identified, directly or indirectly.

\$ Eavesdropper

An entity that passively observes an initiator's communications without the initiator's knowledge or authorization. See [[RFC4949](#)].

\$ Enabler

A protocol entity that facilitates communication between an initiator and a recipient without being directly in the

communications path.

\$ Fingerprint

A set of information elements that identifies a device, application, or initiator.

Cooper, et al.

Expires September 13, 2012

[Page 17]

Internet-Draft

Privacy Considerations

March 2012

\$ Fingerprinting

The process of an observer or attacker partially or fully identifying a device, application, or initiator based on multiple information elements communicated to the observer or attacker. See [[EFF](#)].

\$ Identifiable

A state in which a data subject's identity is known.

\$ Identifiability

The extent to which a data subject is identifiable. See [[I-D.iab-privacy-terminology](#)].

\$ Identifier

A data object that represents a specific identity of a protocol entity or data subject. See [[RFC4949](#)].

\$ Identification

The linking of information to a particular data subject to infer the subject's identity.

\$ Identity

Any subset of a data subject's attributes that identifies the subject within a given context. Data subjects usually have multiple identities for use in different contexts.

\$ Initiator

A protocol entity that initiates communications with a recipient.

\$ Intermediary

A protocol entity that sits between the initiator and the recipient and is necessary for the initiator and recipient to communicate.

\$ Item of Interest (IOI)

Any data item that an observer or attacker might be interested in. This includes attributes, identifiers, identities, communications, or actions (such as the sending or receiving of a communication). See [[I-D.iab-privacy-terminology](#)].

Cooper, et al.

Expires September 13, 2012

[Page 18]

Internet-Draft

Privacy Considerations

March 2012

\$ Observer

An entity that is authorized to receive and handle data from an initiator and thereby is able to observe and collect information, potentially posing privacy threats depending on the context. Recipients, intermediaries, and enablers can all be observers.

\$ Personal Data

Any information relating to a data subject.

\$ (Protocol) Interaction

A unit of communication within a particular protocol. A single interaction may be comprised of a single message between an initiator and recipient or multiple messages, depending on the protocol.

\$ Pseudonym

An identifier of a data subject other than the subject's real name.

\$ Pseudonymity

The state of being pseudonymous. See

[\[I-D.iab-privacy-terminology\]](#).

\$ Pseudonymous

A property of a data subject in which the subject is identified by a pseudonym.

\$ Recipient

A protocol entity that receives communications from an initiator.

\$ Traffic Analysis

The inference of information from observation of traffic flows (presence, absence, amount, direction, and frequency). See [\[RFC4949\]](#).

Cooper, et al.

Expires September 13, 2012

[Page 19]

Internet-Draft

Privacy Considerations

March 2012

[8.](#) Security Considerations

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

Cooper, et al.

Expires September 13, 2012

[Page 20]

Internet-Draft

Privacy Considerations

March 2012

[9.](#) IANA Considerations

This document does not require actions by IANA.

[10](#). 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.

11. References

11.1. Normative References

[I-D.iab-privacy-terminology]

Hansen, M., Tschofenig, H., and R. Smith, "Privacy Terminology", [draft-iab-privacy-terminology-00](#) (work in progress), January 2012.

11.2. Informative References

[EFF] Electronic Frontier Foundation, "Panopticlick", 2011.

[OECD] Organization for Economic Co-operation and Development, "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.

[PbD] Office of the Information and Privacy Commissioner, Ontario, Canada, "Privacy by Design", 2011.

[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.

[RFC3325] Jennings, C., Peterson, J., and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks", [RFC 3325](#), November 2002.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), July 2003.

[RFC4101] Rescorla, E. and 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.

[RFC5025] Rosenberg, J., "Presence Authorization Rules", [RFC 5025](#), December 2007.

Internet-Draft

Privacy Considerations

March 2012

[RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J.,
Tschofenig, H., and H. Schulzrinne, "An Architecture for
Location and Location Privacy in Internet Applications",
[BCP 160](#), [RFC 6280](#), July 2011.

[Solove] Solove, D., "Understanding Privacy", 2010.

[Tor] The Tor Project, Inc., "Tor", 2011.

Internet-Draft

Privacy Considerations

March 2012

Authors' Addresses

Alissa 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
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
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
US

Email: bernarda@microsoft.com

Jon Peterson
NeuStar, Inc.

1800 Sutter St Suite 570
Concord, CA 94520
US

Email: jon.peterson@neustar.biz

John B. Morris, Jr.

Email: ietf@jmorris.org