

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 15, 2011

M. Hansen, Ed.
ULD Kiel
H. Tschofenig
Nokia Siemens Networks
March 14, 2011

Terminology for Talking about Privacy by Data Minimization: Anonymity,
Unlinkability, Undetectability, Unobservability, Pseudonymity, and
Identity Management

[draft-hansen-privacy-terminology-02.txt](#)

Abstract

This document is an attempt to consolidate terminology in the field privacy by data minimization. It motivates and develops definitions for anonymity/identifiability, (un)linkability, (un)detectability, (un)observability, pseudonymity, identity, partial identity, digital identity and identity management. Starting the definitions from the anonymity and unlinkability perspective reveals some deeper structures in this field.

Note: This document is discussed at
<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 15, 2011.

Copyright Notice

Copyright (c) 2011 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	3
2.	Anonymity	3
3.	Unlinkability	6
4.	Anonymity in Terms of Unlinkability	8
5.	Undetectability and Unobservability	10
6.	Pseudonymity	13
7.	Identity Management	19
8.	Contributors	20
9.	Acknowledgments	21
10.	Security Considerations	21
11.	IANA Considerations	21
12.	References	21
12.1.	Normative References	21
12.2.	Informative References	21
Appendix A.	Overview of Main Definitions and their Opposites	22
Appendix B.	Relationships between Terms	23

Internet-Draft

Privacy Terminology

March 2011

1. Introduction

Early papers from the 1980ies about privacy by data minimization already deal with anonymity, unlinkability, unobservability, and pseudonymity. These terms are often used in discussions about privacy properties of systems.

Data minimization means that first of all, the ability for others to collect personal data should be minimized. Often, however, the collection of personal data cannot not be prevented entirely. In such a case, the goal is to minimize the collection of personal data. The time how long collected personal data is stored should be minimized.

Data minimization is the only generic strategy to enable anonymity, since all correct personal data help to identify if we exclude providing misinformation (inaccurate or erroneous information, provided usually without conscious effort at misleading, deceiving, or persuading one way or another) or disinformation (deliberately false or distorted information given out in order to mislead or deceive).

Furthermore, data minimization is the only generic strategy to enable unlinkability, since all correct personal data provide some linkability if we exclude providing misinformation or disinformation.

This document does not aim to collect all terms used in the area of privacy. Even the definition of the term 'privacy' itself difficult due to the contextual nature of it; the understanding of privacy has changed over time. For the purpose of this document we refer to one fairly well established definition by Alan Westin from 1967 [[West67](#)]:

"Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others. Viewed in terms of the relation of the individual to social participation, privacy is the

voluntary and temporary withdrawal of a person from the general society through physical or psychological means, either in a state of solitude or small-group intimacy or, when among larger groups, in a condition of anonymity or reserve.", see page 7 of [[West67](#)].

2. Anonymity

To enable anonymity of a subject, there always has to be an appropriate set of subjects with potentially the same attributes.

Definition: Anonymity of a subject means that the subject is not identifiable within a set of subjects, the anonymity set.

Note:

"not identifiable within the anonymity set" means that only using the information the attacker has at his discretion, the subject is not distinguishable from the other subjects within the anonymity set.

In order to underline that there is a possibility to quantify anonymity for some applications (instead to treating it purely as a binary value it is possible to use the following variation of the previous definition: "Anonymity of a subject from an attacker's perspective means that the attacker cannot sufficiently identify the subject within a set of subjects, the anonymity set."

The anonymity set is the set of all possible subjects. The set of possible subjects depends on the knowledge of the attacker. Thus, anonymity is relative with respect to the attacker. With respect to actors, the anonymity set consists of the subjects who might cause an action. With respect to actees, the anonymity set consists of the subjects who might be acted upon. Therefore, a sender may be anonymous (sender anonymity) only within a set of potential senders, his/her sender anonymity set, which itself may be a subset of all subjects who may send a message. The same for the recipient means that a recipient may be anonymous (recipient anonymity) only within a set of potential recipients, his/her recipient anonymity set. Both anonymity sets may be disjoint, be the same, or they may overlap.

The anonymity sets may vary over time. Since we assume that the attacker does not forget anything he knows, the anonymity set cannot increase w.r.t. a particular IOI. Especially subjects joining the system in a later stage, do not belong to the anonymity set from the point of view of an attacker observing the system in an earlier stage. (Please note that if the attacker cannot decide whether the joining subjects were present earlier, the anonymity set does not increase either: It just stays the same.) Due to linkability, cf. below, the anonymity set normally can only decrease.

Anonymity of a set of subjects within an anonymity set means that all these individual subjects are not identifiable within this anonymity set. In this definition, "set of subjects" is just taken to describe that the anonymity property holds for all elements of the set. Another possible definition would be to consider the anonymity property for the set as a whole. Then a semantically quite different definition could read: Anonymity of a set S of subjects within a larger anonymity set A means that it is not distinguishable whether the subject S whose anonymity is at stake (and which clearly is

within A) is within S or not.

Anonymity in general as well as the anonymity of each particular subject is a concept which is very much context dependent (on, e.g., subjects population, attributes, time frame, etc). In order to quantify anonymity within concrete situations, one would have to describe the system in sufficient detail, which is practically not always possible for large open systems. Besides the quantity of anonymity provided within a particular setting, there is another aspect of anonymity: its robustness. Robustness of anonymity characterizes how stable the quantity of anonymity is against changes in the particular setting, e.g., a stronger attacker or different probability distributions. We might use quality of anonymity as a term comprising both quantity and robustness of anonymity. To keep this text as simple as possible, we will mainly discuss the quantity of anonymity in the following, using the wording "strength of anonymity".

The above definitions of anonymity and the mentioned measures of quantifying anonymity are fine to characterize the status of a subject in a world as it is. If we want to describe changes to the anonymity of a subject if the world is changed somewhat, e.g., the

subject uses the communication network differently or uses a modified communication network, we need another definition of anonymity capturing the delta. The simplest way to express this delta is by the observations of "the" attacker.

Definition: An anonymity delta (regarding a subject's anonymity) from an attacker's perspective specifies the difference between the subject's anonymity taking into account the attacker's observations (i.e., the attacker's a-posteriori knowledge) and the subject's anonymity given the attacker's a-priori knowledge only.

Note:

In some publications, the a-priori knowledge of the attacker is called "background knowledge" and the a-posteriori knowledge of the attacker is called "new knowledge".

As we can quantify anonymity in concrete situations, so we can quantify the anonymity delta. This can be done by just defining:
quantity(anonymity delta) := quantity(anonymity_a-posteriori) - quantity(anonymity_a-priori)

If anonymity_a-posteriori and anonymity_a-priori are the same, their quantification is the same and therefore the difference of these quantifications is 0. If anonymity can only decrease (which usually is quite a reasonable assumption), the maximum of quantity(anonymity

delta) is 0.

Since anonymity cannot increase, the anonymity delta can never be positive. Having an anonymity delta of zero means that anonymity stays the same. This means that if the attacker has no a-priori knowledge about the particular subject, having no anonymity delta implies anonymity. But if the attacker has an a-priori knowledge covering all actions of the particular subject, having no anonymity delta does not imply any anonymity at all. If there is no anonymity from the very beginning, even preserving it completely does not yield any anonymity. To be able to express this conveniently, we use wordings like "perfect preservation of a subject's anonymity". It might be worthwhile to generalize "preservation of anonymity of single subjects" to "preservation of anonymity of sets of subjects", in the limiting case all subjects in an anonymity set. An important

special case is that the "set of subjects" is the set of subjects having one or several attribute values A in common. Then the meaning of "preservation of anonymity of this set of subjects" is that knowing A does not decrease anonymity. Having a negative anonymity delta means that anonymity is decreased.

3. Unlinkability

Definition: Unlinkability of two or more items of interest (IOIs, e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system (comprising these and possibly other items), the attacker cannot sufficiently distinguish whether these IOIs are related or not.

Linkability is the negation of unlinkability:

Definition: Linkability of two or more items of interest (IOIs, e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system (comprising these and possibly other items), the attacker can sufficiently distinguish whether these IOIs are related or not.

For example, in a scenario with at least two senders, two messages sent by subjects within the same anonymity set are unlinkable for an attacker if for him, the probability that these two messages are sent by the same sender is sufficiently close to $1/(\text{number of senders})$.

Definition: An unlinkability delta of two or more items of interest (IOIs, e.g., subjects, messages, actions, ...) from an attacker's perspective specifies the difference between the unlinkability of these IOIs taking into account the attacker's observations and the unlinkability of these IOIs given the attacker's a-priori knowledge only.

Since we assume that the attacker does not forget anything, unlinkability cannot increase. Normally, the attacker's knowledge cannot decrease (analogously to Shannon's definition of "perfect secrecy"). An exception of this rule is the scenario where the use of misinformation (inaccurate or erroneous information, provided usually without conscious effort at misleading, deceiving, or persuading one way or another [Wils93]) or disinformation (deliberately false or distorted information given out in order to

mislead or deceive [Wils93]) leads to a growing uncertainty of the attacker which information is correct. A related, but different aspect is that information may become wrong (i.e., outdated) simply because the state of the world changes over time. Since privacy is not only about to protect the current state, but the past and history of a data subject as well, we will not make use of this different aspect in the rest of this document. Therefore, the unlinkability delta can never be positive. Having an unlinkability delta of zero means that the probability of those items being related from the attacker's perspective stays exactly the same before (a-priori knowledge) and after the attacker's observations (a-posteriori knowledge of the attacker). If the attacker has no a-priori knowledge about the particular IOIs, having an unlinkability delta of zero implies unlinkability. But if the attacker has a-priori knowledge covering the relationships of all IOIs, having an unlinkability delta of zero does not imply any unlinkability at all. If there is no unlinkability from the very beginning, even preserving it completely does not yield any unlinkability. To be able to express this conveniently, we use wordings like "perfect preservation of unlinkability w.r.t. specific items" to express that the unlinkability delta is zero. It might be worthwhile to generalize "preservation of unlinkability of two IOIs" to "preservation of unlinkability of sets of IOIs", in the limiting case all IOIs in the system.

For example, the unlinkability delta of two messages is sufficiently small (zero) for an attacker if the probability describing his a-posteriori knowledge that these two messages are sent by the same sender and/or received by the same recipient is sufficiently (exactly) the same as the probability imposed by his a-priori knowledge. Please note that unlinkability of two (or more) messages of course may depend on whether their content is protected against the attacker considered. In particular, messages may be unlinkable if we assume that the attacker is not able to get information on the sender or recipient from the message content. Yet with access to their content even without deep semantical analysis the attacker can notice certain characteristics which link them together - e.g. similarities in structure, style, use of some words or phrases, consistent appearance of some grammatical errors, etc. In a sense, content of messages may play a role as "side channel" in a similar

way as in cryptanalysis - i.e., content of messages may leak some

information on their linkability.

Roughly speaking, no unlinkability delta of items means that the ability of the attacker to relate these items does not increase by observing the system or by possibly interacting with it.

The definitions of unlinkability, linkability and unlinkability delta do not mention any particular set of IOIs they are restricted to. Therefore, the definitions of unlinkability and unlinkability delta are very strong, since they cover the whole system. We could weaken the definitions by restricting them to part of the system:

"Unlinkability of two or more IOIs from an attacker's perspective means that within an unlinkability set of IOIs (comprising these and possibly other items), the attacker cannot sufficiently distinguish whether these IOIs are related or not."

4. Anonymity in Terms of Unlinkability

To describe anonymity in terms of unlinkability, we have to augment the definitions of anonymity given in [Section 2](#) by making explicit the attributes anonymity relates to. For example, if we choose the attribute "having sent a message" then we can define:

A sender s sends a set of messages M anonymously, iff s is anonymous within the set of potential senders of M , the sender anonymity set of M .

If the attacker's focus is not on the sender, but on the message, we can define:

A set of messages M is sent anonymously, iff M can have been sent by each set of potential senders, i.e., by any set of subjects within the cross product of the sender anonymity sets of each message m within M .

When considering sending and receiving of messages as attributes, the items of interest (IOIs) are "who has sent or received which message", then, anonymity of a subject w.r.t. an attribute may be defined as unlinkability of this subject and this attribute. In the wording of the definition of unlinkability: a subject s is related to the attribute value "has sent message m " if s has sent message m . s is not related to that attribute value if s has not sent message m . Same for receiving. Unlinkability is a sufficient condition of anonymity, but it is not a necessary condition. Thus, failing unlinkability w.r.t. some attribute value(s) does not necessarily eliminate anonymity as defined in [Section 2](#); in specific cases (i.e., depending on the attribute value(s)) even the strength of anonymity

may not be affected.

Definition: Sender anonymity of a subject means that to this potentially sending subject, each message is unlinkable.

Note:

The property unlinkability might be more "fine-grained" than anonymity, since there are many more relations where unlinkability might be an issue than just the relation "anonymity" between subjects and IOIs. Therefore, the attacker might get to know information on linkability while not necessarily reducing anonymity of the particular subject - depending on the defined measures. An example might be that the attacker, in spite of being able to link, e.g., by timing, all encrypted messages of a transactions, does not learn who is doing this transaction.

Correspondingly, recipient anonymity of a subject means that to this potentially receiving subject, each message is unlinkable.

Relationship anonymity of a pair of subjects, the potentially sending subject and the potentially receiving subject, means that to this potentially communicating pair of subjects, each message is unlinkable. In other words, sender and recipient (or each recipient in case of multicast) are unlinkable. As sender anonymity of a message cannot hold against the sender of this message himself nor can recipient anonymity hold against any of the recipients w.r.t. himself, relationship anonymity is considered w.r.t. outsiders only, i.e., attackers being neither the sender nor one of the recipients of the messages under consideration.

Thus, relationship anonymity is a weaker property than each of sender anonymity and recipient anonymity: The attacker might know who sends which messages or he might know who receives which messages (and in some cases even who sends which messages and who receives which messages). But as long as for the attacker each message sent and each message received are unlinkable, he cannot link the respective senders to recipients and vice versa, i.e., relationship anonymity holds. The relationship anonymity set can be defined to be the cross product of two potentially distinct sets, the set of potential senders and the set of potential recipients or - if it is possible to exclude some of these pairs - a subset of this cross product. So the relationship anonymity set is the set of all possible sender-recipient(s)-pairs. In case of multicast, the set of potential recipients is the power set of all potential recipients. If we take the perspective of a subject sending (or receiving) a particular

message, the relationship anonymity set becomes the set of all potential recipients (senders) of that particular message. So fixing

one factor of the cross product gives a recipient anonymity set or a sender anonymity set.

Note:

The following is an explanation of the statement made in the previous paragraph regarding relationship anonymity: For all attackers it holds that sender anonymity implies relationship anonymity, and recipient anonymity implies relationship anonymity. This is true if anonymity is taken as a binary property: Either it holds or it does not hold. If we consider quantities of anonymity, the validity of the implication possibly depends on the particular definitions of how to quantify sender anonymity and recipient anonymity on the one hand, and how to quantify relationship anonymity on the other. There exists at least one attacker model, where relationship anonymity does neither imply sender anonymity nor recipient anonymity. Consider an attacker who neither controls any senders nor any recipients of messages, but all lines and - maybe - some other stations. If w.r.t. this attacker relationship anonymity holds, you can neither argue that against him sender anonymity holds nor that recipient anonymity holds. The classical MIX-net [[Chau81](#)] without dummy traffic is one implementation with just this property: The attacker sees who sends messages when and who receives messages when, but cannot figure out who sends messages to whom.

[5.](#) Undetectability and Unobservability

In contrast to anonymity and unlinkability, where not the IOI, but only its relationship to subjects or other IOIs is protected, for undetectability, the IOIs are protected as such. Undetectability can be regarded as a possible and desirable property of steganographic systems. Therefore it matches the information hiding terminology (see [[Pfit96](#)], [[ZFKP98](#)]). In contrast, anonymity, dealing with the relationship of discernible IOIs to subjects, does not directly fit into that terminology, but independently represents a different dimension of properties.

Definition: Undetectability of an item of interest (IOI) from an

attacker's perspective means that the attacker cannot sufficiently distinguish whether it exists or not.

If we consider messages as IOIs, this means that messages are not sufficiently discernible from, e.g., "random noise". A slightly more precise formulation might be that messages are not discernible from no message. A quantification of this property might measure the number of indistinguishable IOIs and/or the probabilities of distinguishing these IOIs.

Undetectability is maximal iff whether an IOI exists or not is completely indistinguishable. We call this perfect undetectability.

Definition: An undetectability delta of an item of interest (IOI) from an attacker's perspective specifies the difference between the undetectability of the IOI taking into account the attacker's observations and the undetectability of the IOI given the attacker's a-priori knowledge only.

The undetectability delta is zero iff whether an IOI exists or not is indistinguishable to exactly the same degree whether the attacker takes his observations into account or not. We call this "perfect preservation of undetectability".

Undetectability of an IOI clearly is only possible w.r.t. subjects being not involved in the IOI (i.e., neither being the sender nor one of the recipients of a message). Therefore, if we just speak about undetectability without spelling out a set of IOIs, it goes without saying that this is a statement comprising only those IOIs the attacker is not involved in.

As the definition of undetectability stands, it has nothing to do with anonymity - it does not mention any relationship between IOIs and subjects. Even more, for subjects being involved in an IOI, undetectability of this IOI is clearly impossible. Therefore, early papers describing new mechanisms for undetectability designed the mechanisms in a way that if a subject necessarily could detect an IOI, the other subject(s) involved in that IOI enjoyed anonymity at least. The rationale for this is to strive for data minimization: No subject should get to know any (potentially personal) data - except this is absolutely necessary. This means that

1. Subjects being not involved in the IOI get to know absolutely nothing.
2. Subjects being involved in the IOI only get to know the IOI, but not the other subjects involved - the other subjects may stay anonymous.

The attributes "sending a message" or "receiving a message" are the only kinds of attributes considered, 1. and 2. together provide data minimization in this setting in an absolute sense. Undetectability by uninvolved subjects together with anonymity even if IOIs can necessarily be detected by the involved subjects has been called unobservability:

Definition: Unobservability of an item of interest (IOI) means

- * undetectability of the IOI against all subjects uninvolved in it and
- * anonymity of the subject(s) involved in the IOI even against the other subject(s) involved in that IOI.

As we had anonymity sets of subjects with respect to anonymity, we have unobservability sets of subjects with respect to unobservability. Mainly, unobservability deals with IOIs instead of subjects only. Though, like anonymity sets, unobservability sets consist of all subjects who might possibly cause these IOIs, i.e. send and/or receive messages.

Sender unobservability then means that it is sufficiently undetectable whether any sender within the unobservability set sends. Sender unobservability is perfect iff it is completely undetectable whether any sender within the unobservability set sends.

Recipient unobservability then means that it is sufficiently undetectable whether any recipient within the unobservability set receives. Recipient unobservability is perfect iff it is completely undetectable whether any recipient within the unobservability set receives.

Relationship unobservability then means that it is sufficiently undetectable whether anything is sent out of a set of could-be senders to a set of could-be recipients. In other words, it is sufficiently undetectable whether within the relationship unobservability set of all possible sender-recipient(s)-pairs, a message is sent in any relationship. Relationship unobservability is perfect iff it is completely undetectable whether anything is sent out of a set of could-be senders to a set of could-be recipients.

All other things being equal, unobservability is the stronger, the larger the respective unobservability set is.

Definition: An unobservability delta of an item of interest (IOI) means

- * undetectability delta of the IOI against all subjects uninvolved in it and
- * anonymity delta of the subject(s) involved in the IOI even against the other subject(s) involved in that IOI.

Since we assume that the attacker does not forget anything,

unobservability cannot increase. Therefore, the unobservability delta can never be positive. Having an unobservability delta of zero w.r.t. an IOI means an undetectability delta of zero of the IOI against all subjects uninvolved in the IOI and an anonymity delta of zero against those subjects involved in the IOI. To be able to express this conveniently, we use wordings like "perfect preservation of unobservability" to express that the unobservability delta is zero.

6. Pseudonymity

Having anonymity of human beings, unlinkability, and maybe unobservability is superb w.r.t. data minimization, but would prevent any useful two-way communication. For many applications, we need appropriate kinds of identifiers:

Definition: A pseudonym is an identifier of a subject other than one of the subject's real names.

Note:

An identifier is defined in [[id](#)] as "a lexical token that names entities".

In our setting 'subject' means sender or recipient.

The term 'real name' is the antonym to "pseudonym". There may be multiple real names over lifetime, in particular the legal names, i.e., for a human being the names which appear on the birth certificate or on other official identity documents issued by the State; for a legal person the name under which it operates and which is registered in official registers (e.g., commercial register or register of associations). A human being's real name typically comprises their given name and a family name. In the realm of identifiers, it is tempting to define anonymity as "the attacker cannot sufficiently determine a real name of the subject". But despite the simplicity of this definition, it is severely restricted: It can only deal with subjects which have at least one real name. It presumes that it is clear who is authorized to attach real names to subjects. It fails to work if the relation to real names is irrelevant for the application at hand. Therefore, we stick to the definitions given in [Section 2](#). Note that from a mere technological perspective it cannot always be determined whether an identifier of a subject is a pseudonym or a real name.

Additional useful terms are:

Definition: The subject which the pseudonym refers to is the holder of the pseudonym.

Definition: A subject is pseudonymous if a pseudonym is used as identifier instead of one of its real names.

Definition: Pseudonymity is the use of pseudonyms as identifiers.

So sender pseudonymity is defined as the sender being pseudonymous, recipient pseudonymity is defined as the recipient being

pseudonymous.

In order to be useful in the context of Internet communication we use the term digital pseudonym and declare it as a pseudonym that is suitable to be used to authenticate the holder's IOIs.

Defining the process of preparing for the use of pseudonyms, e.g., by establishing certain rules how and under which conditions civil identities of holders of pseudonyms will be disclosed by so-called identity brokers or how to prevent uncovered claims by so-called liability brokers, leads to the more general notion of pseudonymity, as defined below.

Note:

Identity brokers have for the pseudonyms they are the identity broker for the information who is their respective holder. Therefore, identity brokers can be implemented as a special kind of certification authorities for pseudonyms. Since anonymity can be described as a particular kind of unlinkability, cf. [Section 4](#), the concept of identity broker can be generalized to linkability broker. A linkability broker is a (trusted) third party that, adhering to agreed rules, enables linking IOIs for those entities being entitled to get to know the linking.

To authenticate IOIs relative to pseudonyms usually is not enough to achieve accountability for IOIs.

Therefore, in many situations, it might make sense to let identity brokers authenticate digital pseudonyms (i.e., check the civil identity of the holder of the pseudonym and then issue a digitally signed statement that this particular identity broker has proof of the identity of the holder of this digital pseudonym and is willing to divulge that proof under well-defined circumstances) or both.

Note:

If the holder of the pseudonym is a natural person or a legal person, civil identity has the usual meaning, i.e. the identity attributed to that person by a State (e.g., a natural person being represented by the social security number or the combination of

name, date of birth, and location of birth etc.). If the holder is, e.g., a computer, it remains to be defined what "civil identity" should mean. It could mean, for example, exact type and serial number of the computer (or essential components of it) or even include the natural person or legal person responsible for its operation.

If the digitally signed statement of a trusted identity broker is checked before entering into a transaction with the holder of that pseudonym, accountability can be realized in spite of anonymity.

Whereas anonymity and accountability are the extremes with respect to linkability to subjects, pseudonymity is the entire field between and including these extremes. Thus, pseudonymity comprises all degrees of linkability to a subject. Ongoing use of the same pseudonym allows the holder to establish or consolidate a reputation. Establishing and/or consolidating a reputation under a pseudonym is, of course, insecure if the pseudonym does not enable to authenticate messages, i.e., if the pseudonym is not a digital pseudonym. Then, at any moment, another subject might use this pseudonym possibly invalidating the reputation, both for the holder of the pseudonym and all others having to do with this pseudonym. Some kinds of pseudonyms enable dealing with claims in case of abuse of unlinkability to holders: Firstly, third parties (identity brokers) may have the possibility to reveal the civil identity of the holder in order to provide means for investigation or prosecution. To improve the robustness of anonymity, chains of identity brokers may be used [[Chau81](#)]. Secondly, third parties may act as liability brokers of the holder to clear a debt or settle a claim. [[BuPf90](#)] presents the particular case of value brokers.

There are many properties of pseudonyms which may be of importance in specific application contexts. In order to describe the properties of pseudonyms with respect to anonymity, we limit our view to two aspects and give some typical examples:

The knowledge of the linking may not be a constant, but change over time for some or even all people. Normally, for non-transferable pseudonyms the knowledge of the linking cannot decrease (with the exception of misinformation or disinformation, which may blur the attacker's knowledge.). Typical kinds of such pseudonyms are:

Public Pseudonym: The linking between a public pseudonym and its holder may be publicly known even from the very beginning. E.g., the linking could be listed in public directories such as the entry of a phone number in combination with its owner.

Initially non-Public Pseudonym: The linking between an initially non-public pseudonym and its holder may be known by certain parties, but is not public at least initially. E.g., a bank account where the bank can look up the linking may serve as a non-public pseudonym. For some specific non-public pseudonyms, certification authorities acting as identity brokers could reveal the civil identity of the holder in case of abuse.

Initially Unlinked Pseudonym: The linking between an initially unlinked pseudonym and its holder is - at least initially - not known to anybody with the possible exception of the holder himself/herself. Examples for unlinked pseudonyms are (non-public) biometrics like DNA information unless stored in databases including the linking to the holders.

Public pseudonyms and initially unlinked pseudonyms can be seen as extremes of the described pseudonym aspect whereas initially non-public pseudonyms characterize the continuum in between.

Anonymity is the stronger, the less is known about the linking to a subject. The strength of anonymity decreases with increasing knowledge of the pseudonym linking. In particular, under the assumption that no gained knowledge on the linking of a pseudonym will be forgotten and that the pseudonym cannot be transferred to other subjects, a public pseudonym never can become an unlinked pseudonym. In each specific case, the strength of anonymity depends on the knowledge of certain parties about the linking relative to the chosen attacker model.

If the pseudonym is transferable, the linking to its holder can change. Considering an unobserved transfer of a pseudonym to another subject, a formerly public pseudonym can become non-public again.

With respect to the degree of linkability, various kinds of pseudonyms may be distinguished according to the kind of context for their usage:

Person pseudonym: A person pseudonym is a substitute for the holder's name which is regarded as representation for the holder's civil identity. It may be used in many different contexts, e.g., a number of an identity card, the social security number, DNA, a nickname, the pseudonym of an actor, or a mobile phone number.

Internet-Draft

Privacy Terminology

March 2011

Role pseudonym: The use of role pseudonyms is limited to specific roles, e.g., a customer pseudonym or an Internet account used for many instantiations of the same role "Internet user". The same role pseudonym may be used with different communication partners. Roles might be assigned by other parties, e.g., a company, but they might be chosen by the subject himself/herself as well.

Relationship pseudonym: For each communication partner, a different relationship pseudonym is used. The same relationship pseudonym may be used in different roles for communicating with the same partner. Examples are distinct nicknames for each communication partner. In case of group communication, the relationship pseudonyms may be used between more than two partners.

Role-relationship pseudonym: For each role and for each communication partner, a different role-relationship pseudonym is used. This means that the communication partner does not necessarily know, whether two pseudonyms used in different roles belong to the same holder. On the other hand, two different communication partners who interact with a user in the same role, do not know from the pseudonym alone whether it is the same user. As with relationship pseudonyms, in case of group communication, the role-relationship pseudonyms may be used between more than two partners.

Transaction pseudonym: Apart from "transaction pseudonym" some employ the term "one-time-use pseudonym", taking the naming from "one-time pad". For each transaction, a transaction pseudonym unlinkable to any other transaction pseudonyms and at least initially unlinkable to any other IOI is used, e.g., randomly generated transaction numbers for online-banking. Therefore, transaction pseudonyms can be used to realize as strong anonymity as possible. In fact, the strongest anonymity is given when there is no identifying information at all, i.e., information that would allow linking of anonymous entities, thus transforming the anonymous transaction into a pseudonymous one. If the transaction pseudonym is used exactly once, we have the same strength of anonymity as if no pseudonym is used at all. Another possibility to achieve strong anonymity is to prove the holdership of the pseudonym or specific attribute values (e.g., with zero-knowledge

Figure 1: Lattice of pseudonyms according to their use across different contexts

In general, unlinkability of both role pseudonyms and relationship pseudonyms is stronger than unlinkability of person pseudonyms. The strength of unlinkability increases with the application of role-relationship pseudonyms, the use of which is restricted to both the same role and the same relationship. If a role-relationship pseudonym is used for roles comprising many kinds of activities, the danger arises that after a while, it becomes a person pseudonym in the sense of: "A person pseudonym is a substitute for the holder's name which is regarded as representation for the holder's civil identity." This is even more true both for role pseudonyms and

relationship pseudonyms. Ultimate strength of unlinkability is obtained with transaction pseudonyms, provided that no other information, e.g., from the context or from the pseudonym itself, enabling linking is available.

Anonymity is the stronger, ...

- o the less personal data of the pseudonym holder can be linked to the pseudonym;
- o the less often and the less context-spanning pseudonyms are used and therefore the less data about the holder can be linked;
- o the more often independently chosen, i.e., from an observer's perspective unlinkable, pseudonyms are used for new actions.

The amount of information of linked data can be reduced by different subjects using the same pseudonym (e.g., one after the other when pseudonyms are transferred or simultaneously with specifically created group pseudonyms) or by misinformation or disinformation. The group of pseudonym holders acts as an inner anonymity set within a, depending on context information, potentially even larger outer anonymity set.

[7. Identity Management](#)

Identity can be explained as an exclusive perception of life, integration into a social group, and continuity, which is bound to a body and - at least to some degree - shaped by society. This concept of identity distinguishes between "I" and "Me" [[Mead34](#)] : "I" is the instance that is accessible only by the individual self, perceived as an instance of liberty and initiative. "Me" is supposed to stand for the social attributes, defining a human identity that is accessible by communications and that is an inner instance of control and consistency (see [[ICPP03](#)] for more information). In this terminology, we are interested in identity as communicated to others and seen by them. Therefore, we concentrate on the "Me".

Motivated by identity as an exclusive perception of life, i.e., a psychological perspective, but using terms defined from a computer science, i.e., a mathematical perspective (as we did in the sections before), identity can be explained and defined as a property of an entity in terms of the opposite of anonymity and the opposite of unlinkability. In a positive wording, identity enables both to be identifiable as well as to link IOIs because of some continuity of life. Here we have the opposite of anonymity (identifiability) and the opposite of unlinkability (linkability) as positive properties. So the perspective changes: What is the aim of an attacker w.r.t.

anonymity, now is the aim of the subject under consideration, so the attacker's perspective becomes the perspective of the subject. And again, another attacker (attacker2) might be considered working against identifiability and/or linkability. I.e., attacker2 might try to mask different attributes of subjects to provide for some kind of anonymity or attacker2 might spoof some messages to interfere with the continuity of the subject's life.

Definition: An identity is any subset of attribute values of an individual person which sufficiently identifies this individual person within any set of persons. So usually there is no such thing as "the identity", but several of them.

Definition: Identity management means managing various identities (usually denoted by pseudonyms) of an individual person, i.e., administration of identity attributes including the development and choice of the partial identity and pseudonym to be (re-)used in a specific context or role. Establishment of reputation is possible when the individual person re-uses partial identities. A

prerequisite to choose the appropriate partial identity is to recognize the situation the person is acting in.

Of course, attribute values or even attributes themselves may change over time. Therefore, if the attacker has no access to the change history of each particular attribute, the fact whether a particular subset of attribute values of an individual person is an identity or not may change over time as well. If the attacker has access to the change history of each particular attribute, any subset forming an identity will form an identity from his perspective irrespective how attribute values change. Any reasonable attacker will not just try to figure out attribute values per se, but the point in time (or even the time frame) they are valid (in), since this change history helps a lot in linking and thus inferring further attribute values. Therefore, it may clarify one's mind to define each "attribute" in a way that its value cannot get invalid. So instead of the attribute "location" of a particular individual person, take the set of attributes "location at time x". Depending on the inferences you are interested in, refining that set as a list ordered concerning "location" or "time" may be helpful.

Identities may of course comprise particular attribute values like names, identifiers, digital pseudonyms, and addresses - but they don't have to.

8. Contributors

The authors would like to thank Andreas Pfitzmann for all his work on this document.

9. Acknowledgments

Before this document was submitted to the IETF it already had a long history starting at 2000 and a number of people helped to improve the quality of the document with their feedback. A number of persons contributed to the original writeup and they are acknowledged in http://dud.inf.tu-dresden.de/Anon_Terminology.shtml.

10. Security Considerations

This document introduces terminology for talking about privacy by data minimization. Since privacy protection relies on security

mechanisms this document is also related to security in a broader context.

11. IANA Considerations

This document does not require actions by IANA.

12. References

12.1. Normative References

12.2. Informative References

- [BuPf90] Buerk, H. and A. Pfitzmann, "Value Exchange Systems Enabling Security and Unobservability", Computers & Security , 9/8, 715-721, January 1990.
- [Chau81] Chaum, D., "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms", Communications of the ACM , 24/2, 84-88, 1981.
- [ICPP03] Independent Centre for Privacy Protection & Studio Notarile Genghini, "Identity Management Systems (IMS): Identification and Comparison Study", Study commissioned by the Joint Research Centre Seville, Spain , <http://www.datenschutzzentrum.de/projekte/idmanage/study.htm>, September 2003.
- [Mead34] Mead, G., "Mind, Self and Society", Chicago Press , 1934.
- [Pfit96] Pfitzmann, B., "Information Hiding Terminology -- Results of an informal plenary meeting and additional proposals", Information Hiding , NCS 1174, Springer, Berlin 1996, 347-350, 1996.
- [ReRu98] Reiter, M. and A. Rubin, "Crowds: Anonymity for Web

Transactions", ACM Transactions on Information and System Security , 1(1), 66-92, November 1998.

- [West67] Westin, A., "Privacy and Freedom", Atheneum, New York , 1967.

[Wils93] Wilson, K., "The Columbia Guide to Standard American English", Columbia University Press, New York , 1993.

[ZFKP98] Zoellner, J., Federrath, H., Klimant, H., Pfitzmann, A., Piotraschke, R., Westfeld, A., Wicke, G., and G. Wolf, "Modeling the security of steganographic systems", 2nd Workshop on Information Hiding , LNCS 1525, Springer, Berlin 1998, 345-355, 1998.

[id] "Identifier - Wikipedia", Wikipedia , 2011.

[Appendix A](#). Overview of Main Definitions and their Opposites

o

o

Definition	Negation
Anonymity of a subject from an attacker's perspective means that the attacker cannot sufficiently identify the subject within a set of subjects, the anonymity set.	Identifiability of a subject from an attacker's perspective means that the attacker can sufficiently identify the subject within a set of subjects, the identifiability set.
Unlinkability of two or more items of interest (IOIs, e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system (comprising these and possibly other items), the attacker cannot sufficiently distinguish whether these IOIs are related or not.	Linkability of two or more items of interest (IOIs, e.g., subjects, messages, actions, ...) from an attacker's perspective means that within the system (comprising these and possibly other items), the attacker can sufficiently distinguish whether these IOIs are related or not.

Undetectability of an item of interest (IOI) from an attacker's perspective means that the attacker cannot sufficiently distinguish whether it exists or not.	Detectability of an item of interest (IOI) from an attacker's perspective means that the attacker can sufficiently distinguish whether it exists or not.
-----	-----
Unobservability of an item of interest (IOI) means undetectability of the IOI against all subjects uninvolved in it and anonymity of the subject(s) involved in the IOI even against the other subject(s) involved in that IOI.	Observability of an item of interest (IOI) means "many possibilities to define the semantics".
-----	-----

[Appendix B](#). Relationships between Terms

With respect to the same attacker, unobservability reveals always only a subset of the information anonymity reveals. [ReRu98] propose a continuum for describing the strength of anonymity. They give names: "absolute privacy" (the attacker cannot perceive the presence of communication, i.e., unobservability) - "beyond suspicion" - "probable innocence" - "possible innocence" - "exposed" - "provably exposed" (the attacker can prove the sender, recipient, or their relationship to others). Although we think that the terms "privacy" and "innocence" are misleading, the spectrum is quite useful. We might use the shorthand notation

unobservability => anonymity

for that (=> reads "implies"). Using the same argument and notation, we have

sender unobservability => sender anonymity

recipient unobservability => recipient anonymity

relationship unobservability => relationship anonymity

As noted above, we have

sender anonymity => relationship anonymity

recipient anonymity => relationship anonymity

Internet-Draft

Privacy Terminology

March 2011

sender unobservability => relationship unobservability

recipient unobservability => relationship unobservability

With respect to the same attacker, unobservability reveals always only a subset of the information undetectability reveals

unobservability => undetectability

Authors' Addresses

Marit Hansen (editor)

ULD Kiel

E-Mail: marit.hansen@datenschutzzentrum.de

Hannes Tschofenig

Nokia Siemens Networks

Linnoitustie 6

Espoo 02600

Finland

Phone: +358 (50) 4871445

E-Mail: Hannes.Tschofenig@gmx.net

URI: <http://www.tschofenig.priv.at>

