INFORMATICA — UNIVERSITEIT VAN AMSTERDAM

FACULTEIT DER NATUURWETENSCHAPPEN, WISKUNDE EN INFORMATICA

User based Policy Control of Attribute Authorities

H.M.F. Lohstroh (marten@stdio.nl)

August 23, 2007

A thesis submitted in partial fulfillment of the requirements for the Degree of Bachelor of Science

Supervisor(s): Dr. L. Gommans, Drs. M. Koot (UvA)

Signed: Dr. A. Pimentel (UvA)

Abstract

Identity fraud has become an increasingly large problem, since people began to use the Internet in order to manage and exchange information with exceedingly greater real-life value. Therefore, the social, legal, or financial consequences that are concerned with the (ab)use of such information became more serious. Since there is yet no mechanism present to verify the genuinity or valid ownership of identity attributes, it is relatively easy for fraudsters to impersonate others. Despite the urge for such mechanism, it will most likely be rejected if it doesn't also protect against privacy violation.

In order to better comprehend the problem, we researched the concepts of identity, interest, trust, authority, reputation, privacy and security — which led to a set of definitions that raised a model for the development of a possible solution. After translating our envisioned solution into a set of requirements, we accordingly specified a secure protocol for attribute verification and privacy preservation. Although the protocol partly relies on infrastructure that still needs to be researched and developed, it has the potential of becoming a significant improvement of the Internet's security.

Contents

CHAPTER 1

Introduction

In the past few decades, computer science and technology have rapidly gained an indispensable role in the everyday life of most people. What is also referred to as the "Information Revolution" is evidently illustrated by the exponential growth of the Internet. Whereas in 1994, merely 3 million people were connected to the Internet, nowadays this number has risen to nearly 1.2 billion people, which equals approximately 18 percent of the world population [1]. Aside the more dense and wider spread use of the Internet, also many more useful applications have been found for the occupation of this network. Websites are no more limited to the disposal of information only, additionally they can provide dynamical content and interactive services. A great variety of Internet based services became available within a very short amount of time, facilitating social and economic activities ranging from personal communication to online shopping; the virtual world of the Internet has largely become a reflection of the society we live in, and as the imagery suggests — virtual actions cannot go without their real-world causes and consequences.

On the Internet, users have the freedom to deliberately adopt fake identities, which in many contexts is purely innocent, but in contexts where actions have real-world social, financial or legal consequences, such can become lucrative for attackers and especially harmfull for their victims. Identity fraud can occur if personal information happens to fall into the hands of someone who intends to use it in order to illegally obtain credit, goods, or other services. A forged identity could even be used to frame somebody for a crime or to facilitate illegal immigration, terrorism, espionage, or permanent change of identity.

Recently, an expert group of public and private sector partners, including the UK Metropolitan Police and other regional police forces reported:

> "A recent survey by credit reference agency Callcredit found that on average one in every 1000 people in the UK had already been a victim of identity theft." [2]

In this fashion, along with the denoted growth of the Internet and its wider range of applications, this already worrisomely high rate will only keep on growing. Nevertheless, people seem to greatly underestimate the magnitude of the problem, since they act carelessly by handing out their personal information to others, probably not realizing that information such as date of birth, address or mother's maiden name can be enough information to fraudulently apply for credit cards, loans and much more. But how could they be blamed? Reasonably enough, registration is after all required to make use of most services. And in the present state of art, the only tool for attribute verification is crossreferencing them throughout different contexts, which is in fact another problem — it violates privacy.

Therefore, the aim of this work is to answer the question if we can develop a secure protocol which reduces the Internet's attack surface[1] with regard to identity fraud, without compromising privacy.

[1] A system's attack surface is the set of ways in which an attacker can enter and potentially cause damage to the system. [3]

CHAPTER 2

Related Work

The problem that we address in this work has already been acknowledged and described by many researchers, among them K. Cameron concisely formulated:

> "The Internet was built without a way to know who and what you are connecting to." [4]

In his paper "The Laws of Identity" he announces the urge for what he refers to as a "unifying identity metasystem" which will provide a reliable way to establish who is connecting with what, anywhere on the Internet. To this end, he formulated seven "laws of identity":

1. User Control and Consent
2. Minimal Disclosure for a Constrained Use
3. Justifiable Parties
4. Directed Identity
5. Pluralism of Operators and Technologies
6. Human Integration
7. Consistent Experience Across Contexts

While it seems relatively easy to argue the correctness of these requirements, it is a lot harder to demonstrate that they are complete. However, this initial concept has found quite a lot of admission. Subsequently Microsoft developed a framework called InfoCard, that would later be renamed to CardSpace. Other than the clear problem statement and requirements, the rationale [5] of this system shows a lot of questionable design decisions. For example,

- usage of *static* role-based identities being represented as a card, whereas in many cases only just a few attributes on a card could be relevant, the obligation remains to hand over the complete set — this contradicts with the second law of identity,
- optional auditing by identity providers[1], which might harm the subject's privacy,
- usage of computational tokens with separate plain text descriptions, which discourages subjects to check the information within the computational token and lets relying parties free to not verify the plain text information,
- introduction of a central trusted third party (VeriSign) for the disposal of identity certificates, hereby centralizing a substantial amount of power, which is not particularly desirable — moreover, this approach raises scalability concerns;

Therefore, we decided to look for an alternative solution instead.

[1] Identity Provider can be used interchangeably with the term Attribute Authority (AA)

CHAPTER 3

Basic Principles & Concepts

Essentially, information- and communication systems are built to serve human needs, inherently meaning that they are subject to human problems. In order to sensibly discuss the adequacy of any suggested measures to decrease or even solve such problems, we first need to know what we are talking about from a computer scientist's point of view. Many of the concepts that are embedded within the issues that we focus on, have multiple notions since they are researched in multiple different fields (e.g. sociology, psychology, economy et cetera). Different notions enlight different aspects that are specifically relevant and applicable within the context of a particular field of research. Nevertheless, it is very useful to have a good look at all these notions since they might be different, but they should be in accordance somehow -or at least not mutually conflicting, since they describe different views on the same concept. This requisite thus sets a framework for our own definitions.

Actually, we what we are looking for are logical or mathematical models that reflect the situations as we experience them, but at the same time can serve as building blocks for a software implementation. The difficulty is that we have to rationalize concepts and phenomena that are actually often perceived as highly irrational. Therefore some might find that the following definitions are compromises since they exclude feeling, but since feeling yet cannot be implemented in computer soft- or hardware (which by the way raises interesting questions about the functioning of our own brain), these critics can be safely disregarded.

In the end, we need to argue the correctness of our design in terms of the model raised by our definitions. Ultimately we of course wish to be able to quantify our results, meaning that rather then proving that one solution is better than the other, we want to know *how* good it is. Where it not that the quantification of the matters we discuss, are *relative*, meaning their quantification is subject to the interpretation of individuals and thus differs from person to person. As B. Schneier[6] intelligibly illustrates the gap between perceptual risk and mathematical risk — it's for example hard to tell exactly how afraid a person is to get killed in a car accident, or how much safer wearing a seat-belt would make him feel, but we can objectively ascertain that seat-belts in fact do improve safety.

We now stress that frankly we don't neglect the aspect of feeling, since it is incorporated in our model by means of the noted relativity. All that is referred to as "feeling" we don't regard as a (missing) component in our model, but as an integral uncertainty factor that descends from the confusion and perceptual moderation that is embedded in human consciousness. Consequently, we have to settle with comparable results, but this is not an issue, since we are not interested in quantitative results on an individual basis; we also don't need to study statistical mechanics (describing micro-level characteristics of individual particles) in order to understand the macro-level characteristics of a material described by thermo dynamics.

After the stage of implementation — similarly to the way car safety is tested using crashtest dummies — statistics could be utilized in order to quantify the effectiveness of the system in terms of reliability (probability of the system working correctly), vulnerability (probability of the system being degraded or damaged by adverse factors or influences), and so forth.

3.1 Identity

The concept of identity basically comes down to any imaginable valid response to the question "Who are you?". Let us say that it's Alice who is asking Bob this rather simple but ambiguous question. Bob's answer will depend on numerous internal and external factors such as the environment and/or context in which the question is asked, Bob's self-perception, his current role and of course his relationship with Alice.

Classically, identity is understood according to the indiscernability of identicals — part of Leibniz's law. It says that if what appear to be two or more objects are in fact identical, there can be no property held by one and not by the others. Some philosophers reject this understanding of identity and state that identity is relative. Peter Geach for example argued that there is no one relation which is the relation expressed by "is identical with", but many different relations expressed in different contexts. He stated:

> When one says "x is identical with y" this, I hold, is an incomplete expression; it is short for "x is the same A as y", where "A" represents some count noun understood from the context of utterance — or else, it is just a vague expression of a half-formed thought.[7]

Note that Leibniz tells us that x cannot have different (absolute) identities since x would then not be identical with x anymore, therefore Geach extended the definition of "identical with" by using predicates to create subsets of properties of x, called relative identities.

Hence,

Definition 1 *In conformity with Leibniz's law, the identity of an object (e.g. a person) consists of the complete collection of all its attributes, whereas the same object's relative identity can be context dependently composed using any possible combination of its attributes.*

Relative identities can also be referred to as "views" of an identity, as some researchers at Hewlett-Packard used this term in correspondence with their graphic illustration[8] of identity.

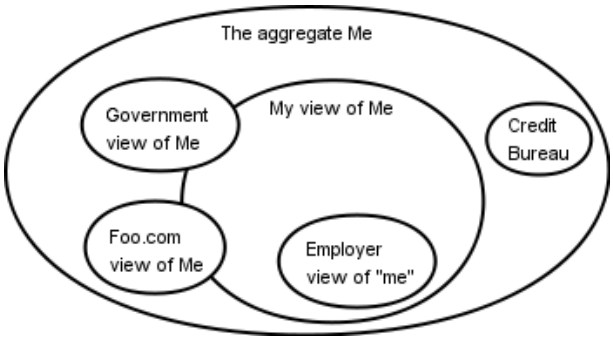

Figure 3.1: HP's views of Identity

The identity of a simple geometric shape as a circle, only holds the numerical attributes radius, surface, and circumference, but the identity of a person is a lot more complex; it does not only include (alpha)numerical attributes such as name or date of birth, but also physical commodities and non-physical characteristics such as desires, opinions, thoughts, abilities, experiences, and intentions. Some attributes can also contain implicit information about other attributes, e.g. the surface of a circle can be derived from its radius and the name Bob is most probably associated with a male individual.

Identifiers are attributes that are specially intended to unambiguously reference to an object within a given context, often used to provide linkage to a set of associated attributes that compose a more extended relative identity. The process in which identifiers are provided and authenticated is appositely referred to as identification.

Unwanted linkage can occur in case different information systems that happen to share a common identifier, are joined together. This way, a larger portion of the absolute identity of a person can be revealed.

3.2 Interest

Interest, also referred to as "well-fare" or "well-being", can be described as a motive for action in order to attain satisfaction of a certain desire. Such desires can be identified as tangible things such as food, money, or shelter, but also as intangible things such as honor, dignity, or recognition. Note that these assets can be described in terms of attributes of an identity.

The desires associated with self-interest are most plausibly restricted to self-regarding desires, but one can of course also care for the well-fare of others. In line with the ideas of the 17th century materialist philosopher T. Hobbes, some argue that even helping others with regard to their desires can be reduced to the satisfaction of self-regarding desires such as display of power, compassion, or kindness. Besides the fact that this postulate is still subject of a lively discussion among philosophers, for our purposes this principle serves as a good "worst-case" assumption.

Common interest is based on desires that are shared with others. However, it's important to note that common interest does not exclude possible conflicts of interest.

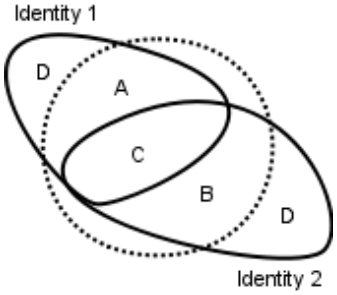

Figure 3.2: Common interest C, disjunct interests A and B, and conflicting interest D

Definition 2 *(Self-)interest is a motive for action to pursue satisfaction of desires that are embedded as attributes within an identity. Common interest is a motive for action to pursue satisfaction of desires within the intersection of two or more associated identities.*

3.3 Trust

Trust is a relationship of reliance involving an expectation, used to bridge uncertainty. Sometimes certainty cannot be obtained (e.g. when dealing with complexities that outgo rationalistic reasoning), but often it's just considered more efficient to establish trust instead of absolute certainty. In other words, trust is expressed by the willingness to take a certain risk.

Grounds for trust can seem very irrational, but in general trust is mainly built on experience. For example, trusting a person for his beautiful blue eyes — besides the fact that there is no proof of any relation between trustworthiness and eye color — most probably has everything to do with pleasant past experiences with blue eyed people. Of course it would be more sensible to evaluate past experiences with this particular person, but such might be scarce or absent. In that case it might be useful to consult others about the person's track record. However, any recommendations coming from untrusted or unknown parties can best be abandoned, since it's not clear what their interests or who their associates are. This brings us to the importance of seeking common interest in trust relationships; people in the same boat face the same fate. But common interest is not sufficient — as you can imagine that if your trustee can swim whilst you cannot, the consequences of being in a sinking boat would all of a sudden be very different

amongst the two of you. Therefore, the stakes, as well as possible disjunct interests that might mutually conflict, must be taken in consideration in order to create a more solid basis for trust.

Similarly to — but more obviously than — identity, trust is a highly dynamical and in most cases also relative concept, noting that absolute trust can only exist unconditionally. In fact, all the criteria that are used to establish a trust relationship with some entity, are attributes of this entity and thus (according to definition 1) can actually be jointly described as a relative identity. The larger the view[8] of a trustee's identity, the more predictable his or her actions will become, making it easier to assess his or her trustworthiness considering an issue at stake. The ultimate trust relationship would then seem to be one of a reflexive[1] kind — having all relevant experience and identity attributes at hand. But ironically enough, conflicting interests and lack of objectivity, still make that people are often not even capable to trust themselves in certain matters. Also note that a trust relationship is not necessarily a symmetrical[2] relation; Alice might trust Bob whilst Bob does not trust Alice.

Along with experience, altered interests, expectations or vulnerabilities, trust can grow, diminish or even completely vanish. Trust growth or decline is two-dimensional; quantatively and qualitatively, respectively related with 1) the number of involved risks and 2) the severity of the consequences associated with those risks. We could also represent trust and trustworthiness in a vectorspace, which according to our model would look like figure 3.3. Then, if we would talk about increase or decrease of either of these notions, we would refer to the length of their representing vector. More ample vector models of trust can be found in the work of Ray and Chakraborty [9].

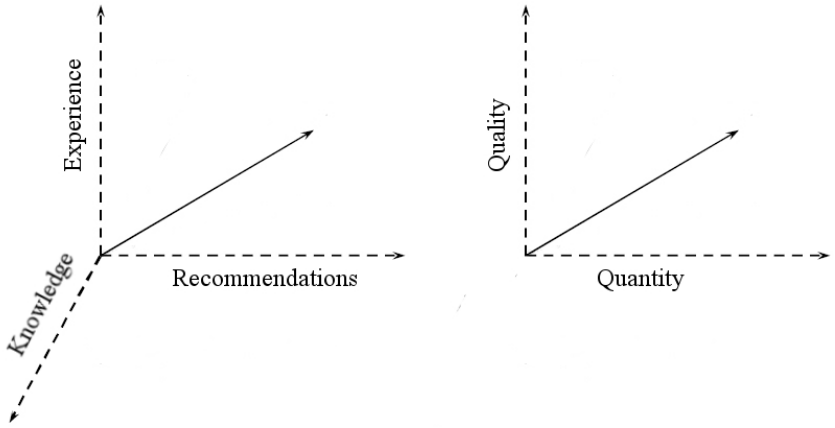

(a) Trustworthiness vector (b) Trust vector

Figure 3.3: Vector representation of trust

To sum this all up,

Definition 3 *Trust is a unidirectional binary relationship of reliance between a trustor and a trustee, expressed by the clemency of the trustor to take a certain preconcidered risk based upon experience, recommendations and knowledge of the trustee's identity.*

Though we already noted the transitive character of trust in the form of recommendations, one might argue that the definition of transitivity[3] conflicts with the characteristic that along with the number of chained links in a relation, trust will drastically deteriorate. Frankly, this phenomenon is dictated by definion 3, since 1) it becomes less probable to have shared experience with all intermediate parties 2) like noted before, recommendations coming from untrusted parties are not valuable — this also holds for intermediately involved parties 3) it becomes less likely to have a proper view of all identities of intermediate parties, they might in fact be completely

[1] Reflexivity: $a \Rightarrow a$
[2] Symmetry: $if\ a \Rightarrow b\ then\ b \Rightarrow a$
[3] Transitivity: $if\ a \Rightarrow b\ and\ b \Rightarrow c\ then\ a \Rightarrow c$

unknown. For example, if the last link in a chain of trust shares the same amount of common interest with its predecessor as the first link does with its successor, it might very well be possible that the first link doesn't share any interest with the last link at all.

Furthermore, we note that trust is vulnerable to abuse, which in fact might go undetected. As long as deception remains secret to the trustor, it won't even affect his trust relationship with a deceitful trustee. Trust is harmed only when the trustor is aware of the fact that his expectations were not met by the trustee.

3.4 Authority

Authority is quite similar to trust since it also is a relationship of reliance that involves certain expectations, but unlike trust it is not necessarily subject to the appraisal of the relying party. Whereas trust is canonically always a one-to-one relation, authority is most often a one-to-many relation. You are for example obliged to obey the laws of the nation you happen to live in, without regard to your individual opinion about these laws, or else you might end up in jail. And unlike most law offenders, the greater part of society will most probably agree with the measures that are taken upon your criminal offense.

Authority resembles power, in the sense that it provides the ability to attain certain ends, but authority also implies a certain legitimacy, justification or right to practice power. The intended function of authority is to preside an organization by representing and protecting common interest within that organization. Externally such authority can bargain with other parties, whilst internally it can maintain a set of rules of conduct that demand members to meet certain criteria (concerning e.g. quality, safety, commitment, or effort) in order to be allowed to take part in the organization. This way something we could describe as a collective relative identity will be established among the organization's members. A great advantage of having such collective relative identity is that it improves the trustworthiness of individual members towards outside parties.

Aside the collective relative identity that arises within an organization, usually also individual relative identities that are associated with certain privileges are administered by the organization's authority. It is of the organization's interest that members live up to its policy and thus it is essential that the right privileges are granted to the right members, which requires authentication.

Individual relative identities kept within an organization may be also relevant outside the organization. You might for example want to supply your bank account number to the telephone company in order to start an automatic payment of your bills. Of course, bank account numbers are administered by your bank, so how can the telephone company ascertain that the number you supplied is truly yours and not someone else's? Even if the telephone company consults the bank, it might conclude that the name and address supplied to the bank correspond with the information it received, but in theory this does not prove that you and the owner of the bank account are the same person. However, if you authenticate at the bank and let the bank vouch for your account ownership, you can indisputably prove to the telephone company that you indeed supplied the correct information. At first sight a lot of practical drawbacks stick to the latter approach, but in a the form of a computer application this construction using a so called trusted third party — which on request provides verifiable attributes that are administered within its own domain — is actually quite common.

Throughout all societies authority plays an important role, often as the cornerstone of complicated hierarchical social systems. Since authority is characteristically a one-to-many relation, it can serve to create typical tree-like hierarchical structures. Examples of such can be found on many different scales and levels, ranging from private households and small companies to governments and multinational enterprises.

Like trust, authority is also vulnerable to abuse. Conflict between an authority's self-interest and the common interest of those he presides can incite abuse of power in the form of corruption or biased decision-making. Distribution of power and authority reduces the probability of abuse since common interest among authorities reduces with the number of authorities, as it

becomes less probable that they are able to unanimously unite their self-interests in the intersection of their identities, disjunct from the common interest of their organization. Note that aside the opportunity, also the willingness to commit abuse is represented as a probability. If multiple authorities have a certain willingness to commit abuse described by the the probabilities $P(A_1), P(A_2), \ldots, P(A_n)$, the probability that they will unanimously agree to actually commit abuse equals $P(A_1 \cap A_2 \ldots \cap A_n) \leq min(P(A_1), P(A_2), \ldots, P(A_n))$.

In terms of minimizing the likelyhood of abuse, it seems clear that an authority must share the maximum amount of common interest with its organization, whilst having minimal disjunct self-interest. However, one might argue that an authority then could commit abuse in order to reduce the likelyhood of future abusiveness, simply by fulfilling personal desire in disadvantage of his organization and thus discarding his disjunct self-interest — which of course is contradictory. Therefore it's important to note that with the removal of the authority's self-interest by means of abuse, this very same interest becomes incorporated in the organization's common interest, but now disjunct from the authority's interest. This notion points out the existence of something like a principle of conservation of interest i.e. you can't take something that you desire, *without* bringing about the same desire to the one(s) you take it from.

In summary,

Definition 4 *Authority is a unidirectional one-to-many relationship that is forced upon individuals though legitimized by general approval of the larger organization. Reliable authorities are characterized by sharing maximum common interest with their organization whilst having minimal disjunct self-interest.*

3.5 Reputation

Reputation is commonly described as a general opinion of the public towards a person, a group or an organization, but if we allow a more technical view on the concept we can define reputation as an aggregation of a distributed question evaluation. It can function as an aid for decision making on a larger basis of experience than one can individually possess. The aggregated characteristic of reputation, which is achieved by mediating information gathered from many different sources, also implies the anonymity of these sources. In retrospect of an earlier statement, in which we argued that recommendations coming from unknown parties are of no value when it comes to the establishment of trust, how can we explain that reputation is such an important factor in so many fields? Why would the characteristic of the denoted transitivity property of trust all of a sudden change when a recommendation would happen to be based on a public instead of individual opinion?

To answer this question let us create our own experiment in which we let an arbitrary but invariable group of people individually and independently evaluate a certain matter, leaving the only variable to be the question that we present to the crowd. The simplest type of question would then be an assertion, that could be answered with "yes", "no" or "I don't know", another would be a multiple choice question with a limited number of valid answers, and the only left alternative would be an open question with an unlimited number of valid answers.

Since the aggregation of reputation is — other then the distribution of authority — not based on unanimous agreement but on the median value obtained from the complete distribution of gathered answers, it is not the intersection but the median of the public interest that is associated with reputation. Therefore we rather need to look at the distribution of gathered answers in order to assess the likelyhood of retrieving abusive or insincere results from our survey. If we look at the answer distribution of the assertion, we can immediate conclude that there is absolutely no possibility to distinguish honest answers from dishonest ones. If we discart all answers from the people who "didn't know" since they after all didn't really answer the question, we are left with a binary distribution that might even result in a worthless 50-50 outcome.

Although the median of the public interest might offer opportunity for deceit, due to the independent and unannounced questioning, conspiracy among the public is actually ruled out. Therefore, individuals might still provide you with wrongful information, but by increasing the number of possible answers the probability that their answers will correspond will of course decrease. All dishonest answers will thus be scattered all over the distribution and the truthful

ones will accumulate in the mid-point. Hence, we propose that open questions are most suitable for reputation-based decision making. However, a word of caution has to be spoken with regard to the resistance of conspiracy; individuals might still gather and arrange a strategy to consistently answer a finite set of questions that are concerned with their common interest. Similar to way we described that opportunity of conspiracy amongst authorities decreases along with the intersection of their interest, this property holds for the individuals involved in an inquiry. For this reason we note the importancy of having a large group to extract such information from.

Now if we hold on to the open question and have a look at the so far invariable group of people, another very important factor is brought to our attention; namely the competence of the public towards answering the specific question. This points out a very important stipulation considering the reliability of the outcome of the survey.

Concludingly we state,

Definition 5 *Reputation is general opinion shared by the public that can be obtained by aggregating the results of an enquiry by means of extracting the median value of its results. In order to obtain reliable results, the question propound to the public must be open and be individually evaluated, the public must be competent to answer the question and large enough to expel conspiracy.*

Many different types of applications could benefit from reputation-based information, for purposes such as authentication or trust establishment.

3.6 Privacy

Privacy is the ability of an entity (e.g. a person, organization or company) to control the flow of sensitive and/or valuable information about itself. Such private information consists of identity attributes and committed or planned actions. Aside unwanted knowledge of private information, especially the linkage between chunks of information is harmful to privacy since it might enable one to reach throughout otherwise separated contexts. Basically three types of linkage are possible, namely:

- between actions
- between attributes
- between actions and attributes

Linkage between actions is almost impossible to prevent, but is generally perceived as quite harmless. Often it's actually used by service providers to improve their service towards their customers. If you ever surfed the web looking for a product, references such as "Customers who bought items like this also bought..." must look familiar to you. Similarly, companies like Google gather personal information by registering your search queries in order to show you so called "relevant" ads. If you ever look at them, you might indeed rather want to see ads that appear to be relevant to you instead of ones that don't. Moreover, personalized ads are considered to be a lot more effective and thus require less views. The reason why linkage between actions generally does not bother people so much, probably has everything to do with the fact that it doesn't involve any identification i.e. it doesn't prevent people from remaining anonymous.

As mentioned earlier, linkage between attributes can occur when different information systems are joined together using a common identifier, which reveals a larger portion of a person's absolute identity than he or she might have intended for others to know. This form of privacy violation is perceived as very harmful since it might cause information to fall into the wrong hands, along with possible consequences like harassment, stigmatization, rejection or other undesirable effects. Also, it can cause loss of anonymity or compromisation of one's efforts to keep a strict separation of information between different contexts.

If we go back to the example of Google, the innocence of linking actions becomes a whole different story if you also happen to have Gmail account. Because in that case, Google is also

able to link your actions with the attributes you provided upon the registration of your e-mail account, let alone the fact that they can obtain additional attributes and actions from reading your correspondence with others. The extend to which Google is able to reconstruct your identity this way, would probably greatly astonish, were it not frighten you. Linkage between actions and attributes is perceived possibly even more harmful then linkage between attributes only, as you can probably think of many things that you would not want to be caught having but certainly not using.

Privacy allows people to have, want, and do things without being associated with their attributes or actions throughout different contexts — in which there might be lack of moral acceptance, approval, or propriety. The latter two forms of linkage we discussed obviously rudely undermine such policies, therefore in order to maintain them, unwanted linkage must be prevented. There is no escape from the need to provide private information to a third party once in a while, but a lot of harm can be prevented by carefully considering what information to provide to whom. The number of attributes obtained by a third party of course determines the magnitude of the view of a person's identity, but it also limits the number of contexts this person is yet able to maintain separately. Therefore, we note the importancy of providing no more attributes than strictly needed. The notion of carefully considering who to provide information to might seem trivial, but this does of course not guarantee that others will restrain themselves from supplying others with information without your permission. Consequently, it will require establishment of trust to reach eligible confidentiality.

Thus,

Definition 6 *Privacy is the ability of an entity to control third party knowledge about its identity and actions throughout different contexts according to a certain policy. Self-maintenace of this policy is trivial, but it requires trust and confidentiality to assure others to respect it.*

3.7 Security

Security is an umbrella term for the probabilities associated with the different risks an entity is exposed to. These risks can be described as potential harm i.e. events with undesirable consequences with regard to one's interest. Thusfar, this description might as well apply to trust or authority, but other than these concepts security does not necessarily involve the uncertainty of reliance. Still many dependencies are of course involved with security, but they are utmost quantifiable restrictions towards security improvement and no cause of uncertainty about the actual probability of risks. Also the effectiveness of security improving measures can be objectively determined in terms of the degree in which they are able to reduce the probability of one or more risks.

Security measures are intended to saveguard certain interest, but doing so they often tend to compromise measures that protect other interests. Therefore, security is all about trade-offs. If we consider intentional efforts to compromise security, this principle actually holds both for the securer and the compromiser; they have to weigh 1) the probability of undesirable consequences, against 2) the magnitude of their investments. Both factors can be independently optimally minimized, but the trade-off between them will merely result in an optimum for both.

Definition 7 *Security is a measure of the probabilities associated with the occurrence of events with undesirable consequences regarding one's interest. Security improvement involves trade-offs between minimizing the probability of undesirable consequences and minimizing the magnitude of the required investments.*

CHAPTER 4

Problem Definition

Identity fraud is not a new phenomenon or some side effect that came along with the rise of the Internet. More traditional methods such as bin raiding, wallet theft or impersonation of the deceased have existed long before the invention of the first computer. Meanwhile, our lives have become exceedingly more complicated and the amount of information we bare with us has grown tremendously — therefore also making it easier to steal some. Additionally, a lot of our administrative affairs are now facilitated remotely, through phone or Internet. Although a lot of these services have implemented some sort of security measures, often they don't require more than a plain text username and password, which can be compared to accepting a signature without a passport. But the problem really lies in the fact that the registration required to obtain such user credentials often doesn't even require proper identification. Currently the only tool at hand for relying parties to verify information is to crossreference it in order to make it more assumable that provided information is indeed correct, but this is not relevant. Correct information doesn't tell anything about who submitted it, therefore what needs to be checked is if a client is rightfully associated with the information he provided. Unfortunately, a mechanism that allows such verification is not yet available.

Not only clients are vulnerable for impersonation, it can also happen to e.g. company or organization representatives. It enables malicious attackers to trick clients into voluntarily releasing valuable information. This strategy can of course be executed in person, by mail or through phone but also by e-mail or even by launching a fake copy of a website — which is also referred to as phishing. To prevent Internet users from such attacks, cryptographic X.509 certificates[1] are used to "prove" the genuinity of a website. However, most Internet users are not capable of judging the genuinity of those X.509 certificates themselves, for the same reason that most Norwegian citizens are not capable of judging the genuinity of a passport issued by the government of Bangladesh; their herefore required frame of reference is absent. So if a user happens to end up on a phishing website, he or she will most probably simply accept any certificate in order to just continue surfing on the website. Then if certificates are not much of an aid, we should perhaps have a better look at the DNS[2], which allowed the user to connect to this website in the first place — noting that it is not a secure system. For more information, see RCF3833 [10] which provides a complete threat analysis of the Domain Name System.

[1] X.509 is a specification for digital certificates that are based on public key encryption and specify information and attributes required for the identification of a person or a computer system.

[2] Domain Name System, a service that translates textual domain names into numerical IP addresses that refer to nodes on the Internet.

CHAPTER 5

Envisioned Solution

The effort of our solution is to improve security on the Internet, as we noted the undesirable consequences of identity fraud. We want to introduce a protocol which reduces the attack surface by minimizing the set of ways in which identity fraud can occur. The protocol shall define control sequences that facilitate unambiguous identification and verifiability of identity attributes, thereby providing the means to make fraud sensitive plain text attributes worthless. What makes it more challenging is that in order to achieve this, we are not willing to trade our privacy.

Let us first define a general use-case, starting by describing its actors.

- subject (S) — wants to provide attributes to RP
- relying party (RP) — wants to verify the attributes provided by S

In order to for RP to be able to verify the attributes provided by S, we introduce a third party that can be trusted by both RP and S.

- attribute authority (AA) — is trusted and capable to verify attributes provided by S

According to definition 4, an AA must share maximum common interest with the subject concerning the privileges attached to the attributes he administers. For example, a bank could have an AA on behalf of their customers concerning their banking related attributes, since both the bank and the customers would be disadvantaged in case a fraud is able to steal from their resources. This notion points out that a person's identity attributes shall inevitably be distributed amongst many different AA's.

As far as goes the trustworthiness of such AA towards an RP as well as an RP towards an S, this can be covered by using either reputation inquiry or authorization by a higher level authority. As you can imagine, it substantially increases the trustworthiness of a bank if it happens to be a member of a national bankers association which is subsequently authorized by the national government. Such hierarchical and/or reputation-based information should be considered as public information, meaning that it should be widely spread and conveniently accessable. In order to provide more transparency in the dynamics of identity, trust, authority, and reputation concerning public entities, all this information should be captured and maintained in a Public Cyberinfrastructure (PCi). Two of the three grounds for trust establishment denoted in definition 3, namely knowledge (authority-based) and recommendations (reputation-based) can be retrieved from such system — making it the perfect aid for assessment of the trustworthiness of public entities including AA's and possibly also RP's.

Then finally, according to definition 6 — in order to safeguard privacy, our protocol must enable subjects to decide when, where, and to whom to release their attributes. Therefore, despite the fact that the administration of attributes is delegated to AA's, subjects must be able to control the policy under which AA's may release their attributes.

5.1 Scope

This work is intended to serve as a specification of the functional requirements regarding our envisioned solution and may serve as basis for a future proof of concept implementation. We shall mainly focus on the control sequences of the protocol, not on the specifics of the involved peripheral systems.

5.2 Perspective

Possible adoption of the protocol for user based control of Attribute Authorities, will depend on the following factors:

- disposal of AA functionality by organizations
- availability of a Public Cyberinfrastructure
- availability of an intuitive and comprehensive user interface

The first dependency will require some investments by the involved organizations, but doing so will reduce their costs of dealing with fraud. In other words, those organizations that suffer most from identity fraud, will be the first to invest in a solution.

The employment of the PCi could be attributed to the Internet Service Providers, similarly like they currently provide DNS service.

When it comes to the implementation of the user interface, enough ideas are in stock already. For example, reputation and authority information from the PCi could be displayed in Internet browsers, right next to the locationbar. Or clumsy HTML registration forms that each time require people their enter their personal information by hand, could be replaced by simple dialogs with a request for certain attributes. In short, there would be ample opportunity to make the websurfing experience a lot more secure and convenient at the same time.

5.3 User characteristics

An implementation of the protocol, along with the development of a Public Cyberinfrastructure would serve virtually all users of the Internet, since really anybody could become a victim of identity theft. It has the potential of offering what all daily Internet users need; a clear view on the hierarchical position and reputation of the companies, organizations and individuals they encounter on the Internet, as well as the means to enforce their own privacy policy.

5.4 Constraints

Besides the presence of a Public Cyberinfrastructure as a substition or extension of contemporary DNS, the protocol requires no structural adaptations of the Internet.

5.5 Assumptions & Dependencies

Both the implementation of the Public Cyberinfrastructure and the Attribute Authorities are not extensively discussed in this work. They fulfill essential roles in the system, but they are rather considered and described as black boxes — as they require further research.

CHAPTER 6

Requirements

This work yet only covers *functional* requirements, therefore implementation specific requirements are not included.

The key words "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED" and "MAY" in this document are to be interpreted as described in RFC2119 [11].

6.1 General

General requirements.

1. There MUST be an extensible standard which unambiguously describes and encodes the set attributes available for use in the system.
2. There MUST be support for optional Subject side security measures such as smartcards or FPGA's with build-in crypto engines or hardware tokens.
3. Subject attributes or credentials MUST always be transfered encryptedly, it is RECOMMENDED also to store then encryptedly.

6.2 Subject

Requirements regarding Subjects.

1. Subjects MUST be protected from impersonation by others.
2. Subjects MUST be able to maintain their own privacy policy, meaning they are able to explicity grant or deny attribute requests from relying parties.
3. It MUST be possible for a Subject to authenticate relying parties.
4. It MAY be possible for Subjects to establish trust towards relying parties on basis of publicly available hierarchical and/or reputational information.
5. Subjects MUST NOT be obliged to inform their Attribute Authorities about the identities of the relying parties.
6. It MUST NOT be possible for Subjects to reuse any signed or validated attributes.
7. It MUST be possible for Subjects to incorporate SSO[1].
8. It is RECOMMENDED for Subjects to use an SSO facility.

[1] Single Sign-On lets a user authenticate only once in order to access multiple applications and systems.

6.3 Relying Party

Requirements regarding Relying Parties.

1. Relying Parties MUST be protected from impersonation by others.
2. Relying Parties MUST be able to verify the authenticity of attributes obtained from Subjects.
3. It MUST be able for Relying Parties to authenticate attribute authorities.
4. It MAY be possible for Relying Parties to establish trust towards Attribute Authorities on basis of publicly available hierarchical and/or reputational information.
5. Relying Parties MUST NOT be required to communicate directly with Attribute Authorities in order to verify the authenticity of a Subject's attributes.

6.4 Attribute Authority

Requirements regarding Attribute Authorities.

1. Attribute Authorities MUST be protected from impersonation by others.
2. Attribute Authorities MUST be able to authenticate Subjects.
3. It MUST NOT be possible for Attribute Authorities to audit or identify Relying Parties.
4. There MUST be a way for Attribute Authorities to enforce their security requirements regarding credential retrieval and storage by Subjects that use an SSO facility.

6.5 Public Cyberinfrastructure

Requirements regarding the Public Cyberinfrastructure.

1. The Public Cyberinfrastructure MUST be a secure and reliable source of information.
2. The Public Cyberinfrastructure MUST be able to resolve the correct host IP address for a given URI[2] and provide its corresponding public key.
3. It MAY also be possible to retrieve additional hierarchical or reputational information about hosts through the Public Cyberinfrastructure.

[2]Uniform Resource Identifier; An alphanumerical address for a resource available on the Internet.

CHAPTER 7

Specification

7.1 Protocol

What we are solving is in fact a distributed third-party authentication and authorization problem that involves privacy constraints and trust issues. Therefore, we first have a look at the AAA (Authentication, Authorization, Accounting) Authorization Framework which is described in RFC2904 [12]. In this document three sequences are denoted — which in our context could be depicted as follows:

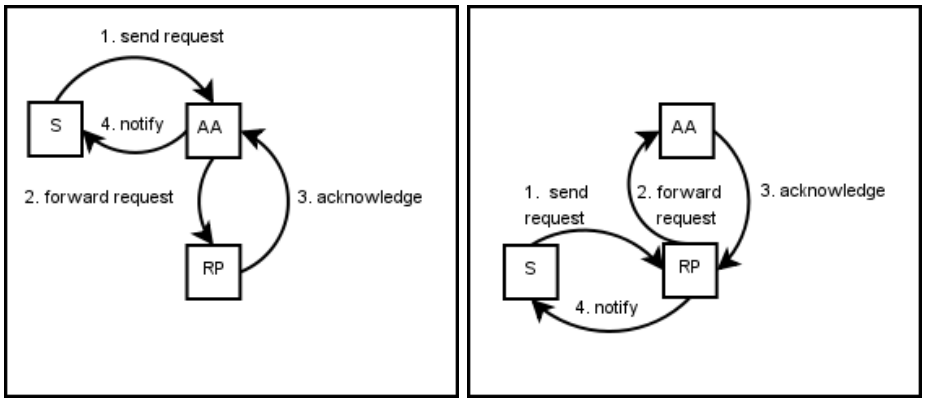

(a) Agent sequence

(b) Pull sequence

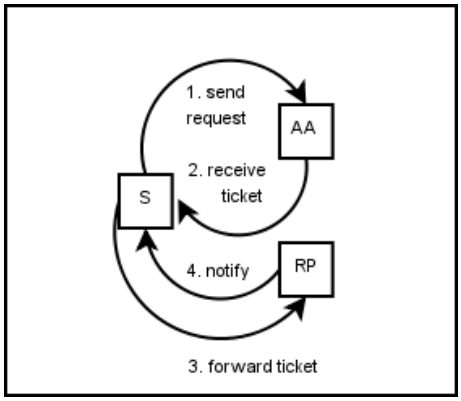

(c) Push sequence

Figure 7.1: AAA Authorization sequences

First of all, requirements **6.3.5** and **6.4.3** rule out the usage of the first two sequences, since they feature direct communication between RP and AA. Then if we take the push sequence as our starting point, we immediately satisfy requirement **6.2.2**, since S regulates all traffic between AA and RP. Now let us extend the sequence as follows:

1. S $\overrightarrow{authenticates\ \&\ trusts}$ AA (former step 1)
2. AA $\overrightarrow{authenticates\ \&\ trusts\ \&\ issues\ cert}$ S (former step 2)
3. S $\overrightarrow{authenticates}$ RP (part of former step 3)
4. $(\dots) \rightarrow$ S $\overrightarrow{trusts\ \&\ forwards\ cert}$ RP (part of former step 3)
5. RP $\overrightarrow{authenticates}$ AA (+)
6. $(\dots) \rightarrow$ RP $\overrightarrow{trusts}$ AA (+)
7. RP $\overrightarrow{authenticates\ cert}$ S (part of former step 4)
8. $(\dots) \rightarrow$ RP $\overrightarrow{trusts\ \&\ authorizes}$ S (part of former step 4)

We first observe the basic assumption that there already exists a trust relationship between S and AA (1); AA is after all the administrator of the attributes of S. After authentication, AA issues a certificate (2) in which a set of requested attributes is embedded, enabling S to verify his claims towards RP. However, it seems more plausible for S to contact RP first (in order to determine which attributes need to be verified), therefore we swap (3,4) and (1, 2). The trust relationship, as well as the authentication between S and RP and RP and AA is less trivial (3, 4, 5, 6); they might not even know eachother, so something is clearly missing here. But if we neglect this for a moment, and move on to (7), we see that RP eventually is able to authenticate the certificate whichafter he may decide to authorize S (8) to make use of his resources, services or whatever.

Besides the fact that in order to protect the system from eavesdropping, **6.1.3** requires the encryption of attributes before they are transmitted, asymmetric encryption[1] will serve as the basis for our authentication method. Attribute Authorities (AA's) are much like Certificate Authorities[2] (CA's) in the sense that they both issue cryptographic certificates under a certain policy. But whereas CA's issue static identity certificates under a certain predefined policy, we define our AA to be able to dynamically compose and issue attribute certificates upon direct request of S.

Furthermore, it's important to note our assumption that AA's are public entities per se, and in many cases RP's are public entities too. For the moment we assume that they are public both. Then, it wouldn't harm if we captured their information in the publicly accessable system, we earlier noted as Public cyberInfratructure (PCi). Such system could then provide the means to bridge the trust en authentication problem involving S and RP, RP and AA in (3, 4, 5, 6), to which requirements (**6.2.4, 6.3.4, 6.2.3, 6.3.3**) are related.

So far, we have introduced the cryptographic certificates and the PCi, but since the certificates are dynamically created and may contain volatile information such as a claims concerning account balances or other commodities, we need to secure the system from replay attacks — as required by **6.2.6**. Finally, the requirements (**6.2.1, 6.3.1, 6.4.1**) regarding impersonation attacks really define the core feature of the system; it must be possible for subjects to refuse to reveal (parts of) their identity, but false identity claims must be ruled out.

In order to satisfy all mentioned requisites, the sequence has to be altered once more, which results in the diagram depicted in fig. 7.2. Using this diagram, we will now demonstrate that this solution indeed solves our problem. We will describe each step in the diagram, along with the arguments that explain the coverage of the designated requirements.

[1] Asymmetric encryption has the property that publicly revealing an encryption key does not thereby reveal the corresponding decryption key. [13]

[2] A Certificate Authority issues digital certificates which contain a public key and identity attributes that the CA attests belong to the designated owner.

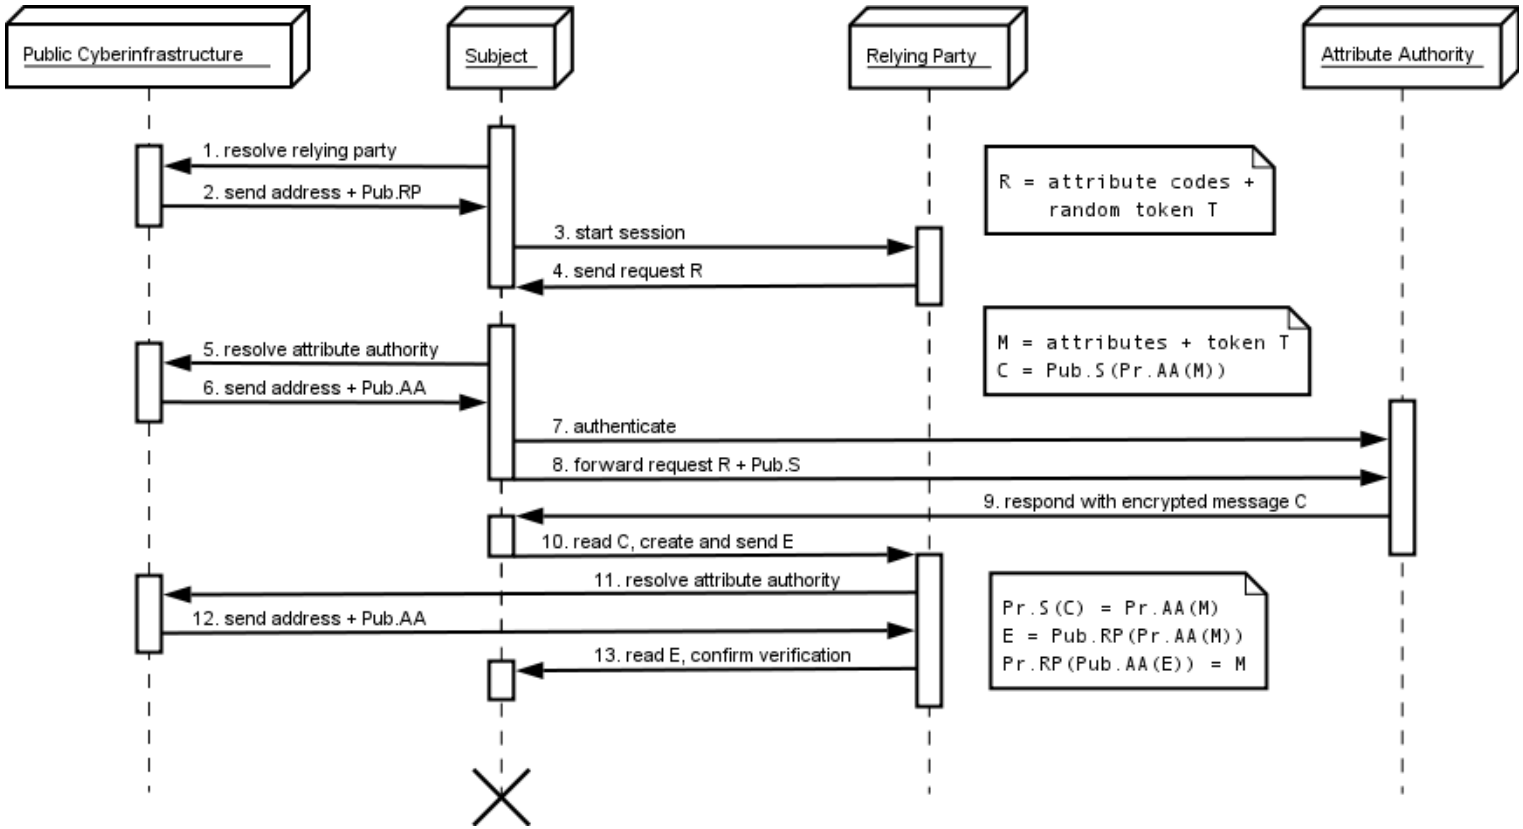

Figure 7.2: UML sequence diagram: communications between S, AA and public RP

1. The PCi is used similar to contemporary DNS in order to resolve hosts on the Internet, but additionally it also provides their corresponding public key. Assuming that the PCi provides correct information, this can be described as S who authenticates RP, which also implies that RP's cannot be impersonated. (**6.2.3, 6.3.1**)

2. In addition to providing the address and public key of RP, the PCi could also provide authority or reputation-based information about RP, which gives S the means to assess the trustworthiness of RP. (**6.2.4**)

3. S starts a session. At this time RP generates a random token (T) which is associated with this particular session between S and RP.

4. RP requests attribute(s) (using an array of standardized attribute codes) from S and sends along token T. S can map the attribute codes to a textual representation in an arbitrary language. At this point S can decide whether to approve or decline the request. (**6.1.1, 6.2.2**)

5. S resolves AA, similar to the first step in this sequence. Although AA is a known host to S, AA might change its name. (**6.4.1**)

6. S obtains the address and public key of AA.

7. S authenticates with AA, the authentication method is arbitrary and solely dependent on an agreement between S and AA. At this time, S generates a keypair (Pri.S, Pub.S). (**6.4.2, 6.1.2**)

8. S forwards the attribute request from RP, along with Pub.S and token T. At this time, AA bundles the requested attributes and token T in message M. Then AA encrypts M using Pub.S and its own private key (Pri.AA), which results in crypted message C, that is signed by AA and can only be read by S using Pri.S and Pub.AA. (**6.1.3**)

9. AA responds to S with encrypted message C, then S decrypts C using Pri.S and Pub.AA. At this point S can still decide not to answer RP's request. S now encrypts Pri.AA(M) with Pub.RP which results in encrypted message E, that is still signed by AA and can now only be read by RP, using Pri.RP and Pub.AA. (**6.2.2, 6.1.3**)

10. S sends E to RP in response to the earlier request.

11. RP resolves AA, similar to steps 1. and 6. (**6.3.3, 6.3.4**)

12. RP obtains the address and public key of AA. In order to establish trust between RP and AA, see step 2. RP decrypts E using Pri.RP and Pub.AA and is able to read the attributes and token T. The fact that RP could encrypt E with Pub.AA, means that M was encrypted with Pri.AA and thus by AA. Token T proves to which particular request M was indeed the answer. Therefore, RP (or another any other party, except AA) cannot impersonate S using M, nor can S itself reissue M without being noticed by RP. (**6.2.1, 6.2.6**)

13. RP decrypts E, reads M and confirms the verification of the attributes provided by S, which was possible without direct communication between RP and AA and thus without enabling AA to audit or identify RP. (**6.3.2, 6.3.5, 6.4.3**)

Then, if we go back to the earlier assumption about RP being a public entity, what would be the consequences if we now assume that RP is a private entity? In that case, RP would not appear in the PCi and would act similar to S, since RP itself would then also need be authenticated trough a third party AA.

The last few remaining requirements (**6.2.7, 6.2.8, 6.2.5, 6.4.4, 6.5.1, 6.5.2, 6.5.3**) only regard specifics concerning the implementation of the Attribute Authority, the Public Cyberinfrastructure and a possible Single Sign-on facility. Therefore, they will be discussed separately.

7.2 Single Sign-On Facility

The protocol does not require an SSO facility, however it's required that the protocol provides support for subjects who wish to use SSO functionality (**6.2.7**). Although an SSO facility would make the sequence more complicated by adding extra communication, it will not alter the ordering of the current steps denoted in fig 7.2. For example, by implementing the AAA agent sequence, extra communication would be added inbetween step 6 and 7 — to let S first retrieve credentials from the SSO facility that otherwise would have already been in his possession. Alternatively, using the AAA push sequence, the SSO facility could authenticate with AA directly and hereafter bridge the communication between S and AA.

The recommendation (**6.2.8**) for subjects to use SSO functionality, results from the rather inconvenient fact that S will have many different AA's that each will administer subsets of all his attributes. This is after all an solution to a real-world problem and it's not reasonable to demand users to simply memorize all the credentials associated with these AA's.

An extra challenge to the implementation of an acceptable SSO service is that since the authentication between S and AA is crucial to the security of the whole system, an AA must be able to prevent it from becoming a weak link due to insecure authentication or communication between S and the SSO service (**6.4.4**). However, finding a solution to this problem is out of the scope of this work.

7.3 Attributes Authorities

An AA basically needs to implement:

- the ability to compose, sign, and encrypt attribute certificates;
- a secure authentication method for subjects (**6.4.2, 6.1.2**);
- the policy not to request information about the relying party (**6.2.5**).

Concerning the format of issued certificates (e.g. header, encryption method and datastructure), an AA would have to comply with a standard, in order for subjects and relying parties to be able to parse and interpret its contents. To this end, for example the X.509 standard for PKI[3] could be extended.

7.4 Public Cyberinfrastructure

The PCi is the part of the entire solution that still needs the most research. The requirements regarding the PCi should therefore rather be considered as prerequisites that serve to fulfill other requirements. Although we cannot — and do not intend to — demonstrate a PCi design that for example is secure and reliable (**6.5.1**), we are confident to say that it's possible to create one. There are already promising projects such as DNSSEC (see RFC4033 [14]) which could serve as a fundament for our ideas. Frankly, if DNSSEC would be widely deployed at the DNS root level, it could already support the distribution of public keys in correspondence with their associated URI's (**6.5.2**).

The dynamics concerning the PCi's hierarchical and reputational information (**6.5.3**), yet remain abstract. At this point we have only explained the use of having such information in order to establish trust, but finding methods to obtain, maintain or distribute it in an integrity preserving way surpasses the scope of this work.

[3]Public Key Infrastructure, a hierarchical arrangement that binds public keys to identities by means of a Certificate Authority

CHAPTER 8

Conclusions

This thesis brings forth a protocol which reduces the Internet's attack surface with regard to identity fraud whilst it enables users to protect their privacy according to their individual policies. A framework is described that allows subjects to delegate arbitrary sets of verifiable identity attributes to relying parties. The verification of these attributes happens by means of third party authentication using an attribute authority. Coming from a situation where the authenticity of attributes cannot be verified at all, this protocol has the potential of becoming a significant improvement of the Internet's security.

Nevertheless, authentication and trust establishment between 1) subjects and attribute authorities, 2) subjects and relying parties, 3) relying parties and attribute authorities, will remain points of failure. In order to reduce these threats, the concept of a Public Cyberinfrastructure is introduced to resolve addresses, public keys and hierarchical and/or reputational information about public entities including attribute authorities and relying parties. This peripheral information system serves as prerequisite for the protocol to work between parties between whom trust establishment is not a trivial case.

Another hazard is that private keys in the system might get compromised if an attacker is able to break a cipher or to simply steal a key. In case the private key of a subject or relying party gets compromised, a third party will be able to read transfered attribute certificates, but he will not be able to (re)issue them in a verifiable manner. However, anybody who gains possession of the private key of an attribute authority is inevitably able to impersonate others. This of course also holds for attribute authorities themselves, but since they control the attributes and resources of subjects directly, subjects basically have to rely on the goodwill of their authorities anyway. Therefore, as in any PKI, proper key management remains of utmost importance.

However, really any security protocol would be based on certain — perhaps unsafe — assumptions, which therefore implies a set of rules of conduct in order preserve its secure properties. To illustrate this, an ultimately secure safe becomes insecure as an open closet in case you would forget to remove the key from its lock after you locked it.

Future work could be dedicated to a more thorough threat analysis of the protocol, the description of a standard for attribute certificates, further research on the Public Cyberinfrastructure, and a proof of concept implementation of the complete system.

Bibliography

[1] Website: "World Internet Usage and Population Statistics"
http://www.internetworldstats.com/stats.htm

[2] Website of UK's National Identity Fraud Prevention Week, "Real Life Stories"
http://www.stop-idfraud.co.uk/public_stories.htm

[3] Carnagie Mellon Cylab, "Attack Surface Measurement"
http://www.cylab.cmu.edu/default.aspx?id=2122

[4] Kim Cameron, "The Laws of Identity",
http://www.identityblog.com/stories/2005/05/13/TheLawsOfIdentity.pdf
Microsoft Corporation, 2005

[5] Kim Cameron, Michael B. Jones, "Design Rationale behind the Identity Metasystem Architecture",
http://www.identityblog.com/wp-content/resources/design_rationale.pdf
Microsoft Corporation

[6] B. Schneier, "The Psychology of Security" (draft)
http://schneier.com/essay-155.pdf
February 28, 2007

[7] P.T. Geach, "Identity", The Review of Metaphysics, Vol. XXI, No. 1 (September 1967), pp. 3-12, See also his "A Reply , The Review of Metaphysics", Vol. XXII, No. 3 (March 1969), pp. 556-9.

[8] Shoaf, Lehigh and Sagamore, "Identity Management: A Developer's View"
(research by Jason Rouault, Joe Pato), pp. 4-1
http://devresource.hp.com/drc/resources/idmgt_devview/idmgt_devview.pdf
Hewlett-Packard Laboratories, March 2004

[9] Indrajit Ray and Sudip Chakraborty, "A Vector Model of Trust for Developing Trustworthy Systems",
http://www.cs.colostate.edu/~indrajit/Security/trust/trust-esorics04.pdf
Colorado State University, 2004

[10] Atkins & Austein [RFC3833] "Threat Analysis of the Domain Name System"
http://tools.ietf.org/html/rfc3833
Network Working Group, August 2004

[11] S. Bradner [RFC2119] "Key words for use in RFCs to Indicate Requirement Levels"
http://tools.ietf.org/html/rfc2119
Network Working Group, March 1997

[12] Vollbrecht, et al. [RFC2904] "AAA Authorization Framework"
http://tools.ietf.org/html/rfc2904
Network Working Group, August 2000

[13] R. L. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems"
MIT Lab. for Computer Science and Department of Mathematics, Cambridge, MA
Association for Computing Machinery, Volume 21, Issue 2, February 1978

[14] Arends, et al. [RFC4033] "DNS Security Introduction and Requirements"
http://tools.ietf.org/html/rfc4033
Network Working Group, March 2005