Reference: SCHLÄGER, C. & NOWEY, T. (2007) On the Effects of Authentication and Authorisation Infrastructures on E-Commerce Environments. International Journal of Computer Systems Science & Engineering (CSSE).

On the Effects of Authentication and Authorisation Infrastructures on E-Commerce Environments

Christian Schläger and Thomas Nowey

University of Regensburg, Germany

christian.schlaeger@wiwi.uni-regensburg.de, thomas.nowey@wiwi.uni-regensburg.de

Abstract: Authentication and Authorisation Infrastructures (AAIs) support service providers on the Internet in outsourcing security services. AAIs influence and change the process of e-commerce transactions on multiple points. Changes influence users, service providers, and provider federations likewise. This work analyses the alterations implied by an AAI, comparing various AAI paradigms and traditional service providing. Effects explored comprise new functionalities, influenced trust relationships, risk management, economic factors, and user privacy. Compared to traditional service providing on the Internet, AAI usage is generally advantageous for all involved parties. However, allocations of services and AAI architectures have a direct impact on the benefits and risks. This work enables vendors and clients to assess security infrastructures and decide on the usage.

Index Terms: Authentication and Authorisation Infrastructures, Risk Assessment, e-commerce, IT-Security, IT-Outsourcing

1. Introduction

Service providing on the Internet has been a huge success story. Although ease of use is proclaimed in many advertisements, the usage of a service on the Internet – e.g. to buy a book or to use a geographic routing service – is not trivial at all. The purchase of a book is not simply a link to click on but it rather stands at the end of a sequel of security and data intensive processes – most of them hidden from the user. The complexity of doing business over the Internet has increased both for customers and vendors.

Concentrating on security-connected processes we find a chain of distinctive services linking the user's request with the service providing as shown in Fig. 1. The shown chain of security services is enhanced by an attribute infrastructure like the ones deduced from OASIS' XACML (eXtensible Access Control Markup Language) and SAML (Security Assertion Markup Language) standards in (Schläger, Nowey et al. 2006).

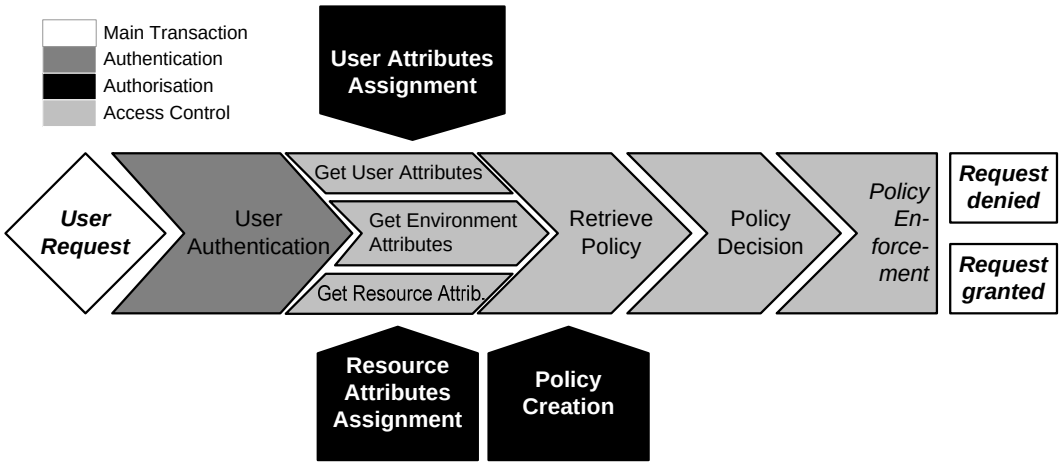

Fig. 1: Security services for accessing a resource in an attribute enhanced infrastructure

Service providers (SPs) are familiar with such infrastructures supporting security services necessary to conduct business. From a commercial point of view the main focus of Authentication and Authorisation Infrastructures (AAIs) lies on the authorisation of customers or, more precisely, on the additional information (attributes) available. This point was first made by Joan Feigenbaum in 1998 (Feigenbaum 1998) and has led to many research activities in the field (e.g. Oppliger 2003). Providers can chose between prominent frameworks like Liberty's Identity Federation Framework ID-FF (Cantor 2004), Microsoft's .NET Passport (Microsoft Inc. 2003) or Internet2's Shibboleth (Cantor 2004). Depending on various factors these architectures are suited differently to address topical demands of service providers and customers in Internet-based transactions.

For SPs these demands include a higher level of security through fine grained access control (AC) and additional information about customers and their reputation as well as the possibility to outsource security services to third party providers, providing cheaper or better services (Schläger, Sojer et al. 2006). Users require better usability with a single sign-on (SSO), central maintenance of account data, and the possibility to pass their reputation and trustworthiness from one provider on to another as well as the protection of their privacy, e.g. through the usage of pseudonyms. These contradictory requirements are depicted in Fig. 2. Of course this list of user and provider demands is not complete. For an in-depth discussion of stakeholders' demands see (Schläger and Pernul 2005) and (Schläger, Sojer et al. 2006).

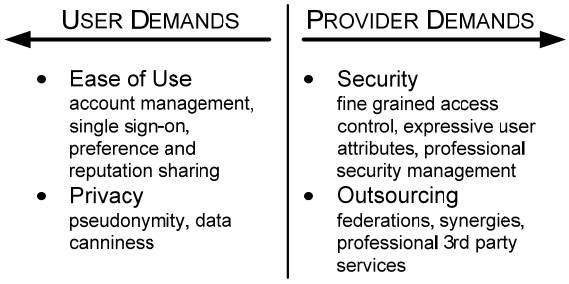

Fig. 2: User and provider demands on AAIs

Risk is an omnipresent factor in Internet transactions. Risks have to be identified and valuated to decide upon appropriate controls and monitoring mechanisms. Basically, one has the option to avoid, reduce, shift, or accept a certain risk. In (Schläger, Nowey et al. 2006) and (Schläger and Pernul 2005) it is argued that AAIs can be used to source out security services in order for the service provider to concentrate on core competencies, raise the overall level of security, provide new, flexible, and more powerful access control services like ABAC (attribute-based access control), and strengthen the usability through user's single sign-on (SSO). However, the usage of a new architecture, especially if not entirely under the control of a service provider, raises questions about risk assessment in comparison to traditional methods of service providing on the Internet. Different approaches to AAIs are available, each with inherent benefits and shortcomings. Differences result from service allocations, the architecture, the level of outsourcing, and specific use cases (Schläger and Ganslmayer 2007).

The second group influenced by an AAI are the users of electronic transactions. If the AAI is able to provide additional functionalities the ease-of-use can be enhanced. Most importantly, the AAI will be evaluated if they can meet the criteria of privacy and reputation sharing. Both issues are strongly interwoven. It is the task of the AAI to provide enough information for a SP to trust the customer and at the same time protect the user's privacy as good as possible. The AAI needs to mediate between the sometimes competing requirements. Usable methods include so called Privacy Enhanced Technologies (PET) (Federrath 2005).

The third group to be considered is the federation or alliance a SP is conducting business in. Following the idea of outsourcing services to the infrastructure (Tanenbaum and Steen 2002) we treat the SP as a part of a federation sharing resources or using third party resources together with others. In this dimension the AAI's role is to guarantee trust. Information about customers has to be certified and secured. Information about the provider's detailed portfolio and his security policies needs to be protected.

In this article effects of Authentication and Authorisation Infrastructures on the involved parties are identified and measured. We explain and evaluate risk factors, functionalities, privacy, and trust relationships in traditional e-commerce surroundings versus e-commerce applications with an AAI. Three AAI paradigms are used for the assessment. Although AAIs are by nature generic architectures, e-commerce has been taken as an example.

2. Related Work

This work combines mechanisms and methods from different research areas to assess the effects of AAIs on the individual user groups. Key concepts for the assessment are Trust and Risk.

2.1. Authentication and Authorisation Infrastructures

AAIs make it possible to combine service outsourcing strategies with strengthened security and more flexible and suitable techniques. A special benefit lies in the accumulation of user data over a federation, such as user profiles, buying patterns, and earned privileges. Identities could be transferred from one service provider to another making it possible to always use up-to-date address data or proof a good reputation acquired at one federation member.

The topic of AAIs as a tool for service providers on the Internet has been discussed on a technical level by Lopez, Oppliger, and Pernul in 2004 (Lopez et al. 2004) and in more detail by Schläger, Pernul in 2005 (Schläger and Pernul 2005). Various architectures of research projects and products have been analysed and motivation for the stakeholders in such infrastructures has been given. In 2006, Schläger, Nowey, and Montenegro (Schläger, Nowey et al. 2006) proposed a reference architecture for an AAI respecting privacy and flexibility through the usage of attribute-based access control. Data canniness, user control, and privacy is guaranteed by the usage of a trusted identity provider using inferring mechanisms and XACML technology. Access rights are derived through policies and user, environment and resource attributes. A part of AAIs – mainly the idea of a single sign-on – has been discussed in the field of identity management. A classification of architectures can be seen at Jøsang and Pope at AusCERT 2005 (Jøsang and Pope 2005).

Due to its origin the effects on user groups, especially the ones concerning security, risks, trust, and usability, have been neglected. A first approach to assess and evaluate new risks through AAI usage has been presented by (Schläger and Nowey 2006) in 2006.

2.2. Trust

The idea of trust - or rather „Trust Management“ – is closely connected to Risk Management. For our further analysis it is important to distinguish the customer’s trust in the service provider and the service provider’s trust in the customer.

2.2.1. The Notion of Customer’s Trust in the Service Provider

Usually, trust is the adversary of risk (Gefen et al. 2003). Risk management primarily analyses the provider’s security threats and management methods. Trust in e-commerce usually is defined as a customer’s willingness to conduct business with a provider or a service. Generally, we can say that the higher the trust of a customer in a provider the higher the likeability of a purchase.

If we talk about customer trust we imply, for this article, that a positive purchase decision depends on the client’s trust in the service provider and all interacting sub providers and services.

2.2.2. The Notion of Service Provider’s Trust in the Customer

From a service provider’s perspective trust relies on dependable predictions on the outcome of a transaction. This dependability is not gained through a PKI or through mere authentication but through authorisation as (Feigenbaum 1998) has stated in 1998. The ability of ABAC models to take into consideration the historical behaviour (e.g. buying patterns) of a customer together with descriptive metadata builds the basis for trusted decisions. The more information is analysed about a client the better the outcome of a policy decision. Consistent, ensured, and dependable user attributes enable stability of trust. Using trust calculating algorithms from, for example, McKnight and Chervany (McKnight and Chervany 2001) and calculating a trust value via Fuzzy Cognitive Maps, the client’s trust value stabilises after a certain number of transactions and a certain amount of metadata used in the process.

2.3. Risk

Risk management techniques have become a vital element of modern security management. A risk is an unwanted event that has negative consequences (Pfleeger 2000) and can be described as the "*combination of the probability of an event and its consequence*" (see ISO/IEC Guide 73:2002).

Systematic risk assessment is especially helpful when it comes to the economic evaluation of information security investments (Nowey et al. 2005). In literature and in practice numerous methodologies and frameworks for conducting risk analyses exist (Vidalis 2004). Virtually all of today's existing approaches use qualitative metrics to assess risks. Quantification is regarded

an important issue but due to many challenges in this field there is to our knowledge no methodology for our purpose up until now. Thus we are going to use a qualitative scale.

The concept of Annual Loss Expectancy has a long tradition in risk management (National Bureau of Standards 1974) and can be seen as the basis for the Return on Security Investment that has gained popularity in recent times (Wei et al. 2001, Berinato 2002).

Cremonini and Martini (Cremonini and Martini 2005) and Cavusoglu et al. (Cavusoglu et al. 2004) pointed out that a holistic cost-benefit evaluation of security investments should also consider the motivation for and possible return of an attacker (*RoA*). Kilger et al. (Kilger et al. 2004) show that there are other motives besides monetary ones which make it almost impossible to quantify this aspect.

3. Methodology

Besides the two obvious user groups of customers and vendors we also analyse effects on the federation or alliance of SPs. The idea is that, apart from individual effects on users and providers, global effects can be assessed as well.

3.1. Service Providers

For the risk assessment we have opted for the approach from (Pfleeger and Pfleeger 2003) with slight adoptions. Risks distinguish themselves from other events due to a loss associated with the event, a probability of the occurrence, and by a chance to change the outcome of the event. Consequently, we are going to divide the holistic process of accessing a resource into separate steps, distinguishing between several forms of implementation for different architectures (traditional e-commerce providing, usage of a fully developed AAI, and hybrid forms). Following to this, we measure the impact for service providers and the according frequency. We evaluate architectural decisions on their impact and suggest the usage of an AAI. The Return on Security Investment (RoSI) is used to economically justify an AAI usage.

Let l_i be the frequency of successful attacks on i in one year. L_i is defined as the expected loss for i in the case of a successful attack. Consequently, the Annual Loss Expectancy for i is defined as

$$ALE_i = l_i \cdot L_i \tag{1}$$

3.2. Users

To assess the effects of AAIs on the end user, two main questions need to be answered:

- Is the AAI able to provide the required functionalities?
- Is the service respecting the user's privacy demands?

The first question will be dealt with by analysing the protocols of selected AAI approaches. A single sign-on functionality should be part of every authentication infrastructure. AAIs will differentiate themselves by how authorisation is handled. Infrastructures should be able to support reputation sharing, central account management, and easy-to-use services.

The second question deals with data security and privacy. A common definition of privacy has been given by Alan Westin (Westin 1967): "*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*." When talking about privacy on the Internet, we usually mean informational privacy, which can be defined as "*Self-determination of what information is known about a person and how it is used*" (Federrath et al. 2002).

A privacy-aware AAI should have the following three basic functionalities. These will be tested in the course of this work.

- Identity and privileges need to be separated.
- Data gained in the process of access needs to be stored in a distributed manner. Involved parties should use and gain as little information as possible.
- Users must not be forced to trust a predefined identity provider but should rather be able to choose among various providers the one of his liking.

3.3. Federation

The integration of an AAI into the process of an electronic transaction changes the role of a single SP. The additional parties make it possible to realise synergies but they also introduce new security problems. The mechanisms of protection and assurance of user and provider information will be assessed comparing the three selected AAI approaches directly. A special focus is put on the changed motivation for an attack in the changed environment.

4. AAI Architectures and Differences

Usually, the usage of AAIs is motivated from a software engineering and point economic of view. The first referring to outsourcing non-functional activities into an infrastructure (Tanenbaum and Steen 2002), the latter to outsourcing security services to concentrate on core competencies gaining competitive advantages and raising the security level through third party know-how (Schläger and Pernul 2005). Referring back to Fig. 1 we see that several security services build on each other to take an access decision. Again, taken the SAML and XACML termini (OASIS Security Services Technical Committee, OASIS eXtensible Access Control Markup Language Technical Committee) as a guideline we can deduct four separate steps of services (Fig. 3). One, all of them, or combinations can be outsourced by a service provider into an AAI.

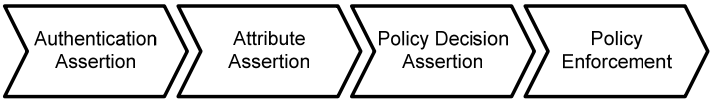

Fig. 3: Outsource-able security services to an AAI, SAML terminology

The characteristics of an AAI can be determined with the help of the given four sub-services in combination with the two prevailing architectural paradigms. AAIs are to this day built either centrally, with a central database or provider in the middle, or as a federation where service providers act as AAI providers themselves. The best examples of these two archetypes are Microsoft's centralised .NET Passport or the Spanish PAPI Software (Castro-Rojo and López 2001) versus the distributed Liberty Identity Federation Framework. The proposed architecture by (Schläger, Sojer et al. 2006) mediates between these paradigms.

AAIs not supporting authorisation and access control services can only stay superficial. They neglect the momentum gained through an integration of all security sub-services. AAIs need to integrate an attribute-based approach. ABAC guarantees fine grained control over privileges and rights as well as it subsumes the traditional access control approaches. A holistic AAI consists of the four sub-services depicted in Fig. 1: Authentication, attribute and policy collection, policy decision, and policy enforcement.

Each sub-service can be provided autonomously. AAI designers need to decide for each sub service whether to provide it in a centralised manner, totally distributed over the federation, or in a manner in between. Service providers have the choice to integrate a sub-service from one central third party provider, from a variety of providers, or not at all, resulting in an in-house realisation of the service. A detailed explanation of the allocation issues in AAIs is given in (Schläger and Ganslmayer 2007).

For the analysis of AAI effects we have restricted ourselves to three approaches:

- A centralised approach with the representatives Microsoft .Net Passport and PAPI,

- a completely distributed approach like Liberty's ID-FF, and
- the partly federated and centralised approach by Schläger et al. from 2006.

4.1. Microsoft .Net Passport and PAPI

Microsoft .NET Passport, although often criticised, has been the first and the largest commercial AAI so far. Passport concentrates on single sign-on for the user who gets its Passport account with every Hotmail account. Passport relies heavily on the usage of cookies imitating to some extend Kerberos's ticketing functionalities. The login to a SP is redirected to Passport requiring his username and password. The SP's ID is transmitted via URL encoding and thus enabling Passport to redirect the client and storing several cookies (see Fig. 4). At the SP a software agent is needed – the so called Passport-Manager. This software reads URL encoded data and stores additional cookies into the SP's domain permitting an access control decision. At another vendor the Passport cookies are used to enable a SSO (Microsoft Inc. 2003). The vendor decides about access of resources using his authorisation and access control mechanism of choice. Passport is a centrally organised SSO system meaning that it only asserts the user's authentication.

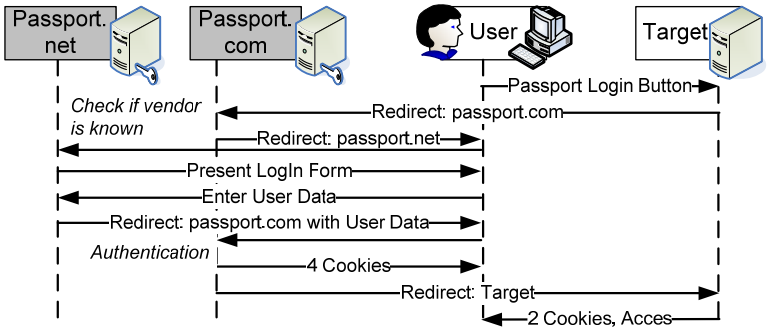

Fig. 4: Microsoft .NET Passport Login Procedure - SSO

Passport uses a central database to keep all client information.

Following the paradigm of centralising all security sub-services a proxy solution like PAPI (Point of Access to Providers of Information) evolves.

Developed in 2001 by RedIRIS, a Spanish research network, it forms a distributed access control system for digital resources accessible over an institution's intranet or the Internet. The user has to authenticate at the authentication server (AS) of his home domain as depicted in Fig. 5.

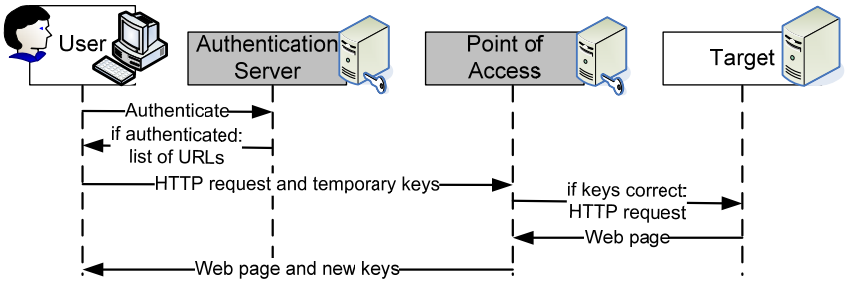

Fig. 5: SSO procedure and access control in PAPI

As PAPI is agnostic to the form of authentication the user's domain is responsible to supply a distinguished name. After a successful authentication a website is given back to the user containing all accessible digital resources. Clicking on a link, he/she is redirected to the Point of Access (PoA) taking with him an encrypted key identifying the AS. The PoA fetches the resource and delivers it to the user. PAPI acts as a Proxy server and handles any interaction for the associated clients and servers. Consequentially, PAPI incorporates all four steps of Fig. 3, forming the maximal AAI (Castro-Rojo and López 2001).

4.2. Liberty Alliance Identity Federation Framework (ID-FF)

Liberty was the open source community's answer to Microsoft Passport in 2001. In Liberty a Circle of Trust (CoT) establishes a Liberty system (Liberty Alliance 2004a, Liberty Alliance 2004b). Each partner provides the authentication for his users with his own methods while they themselves can login to all other partners with a SSO. The user authenticates at his Identity Provider (IdP) and, if he wishes, a cookie is stored under a common domain where every member hosts a server so all of them can access the cookie. If a user moves to a CoT member the cookie is read, the IdP asks for appropriate authentication, and an assertion is awaited. Communication is based on the SAML protocol.

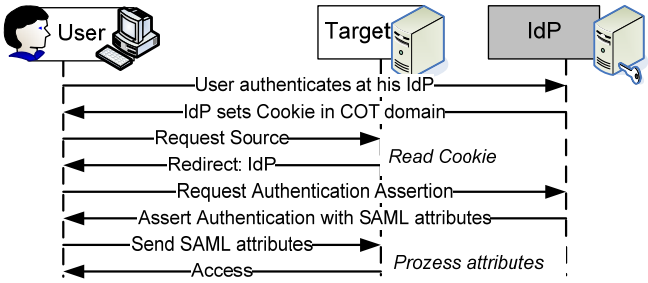

Fig. 6: Liberty SSO procedure

Liberty resembles more a framework than an actual AAI. A CoT has to decide on the implementations. The SAML assertions can carry any attribute the CoT agrees upon. Liberty's architecture is distributed. The IdP is not fixed like in Shibboleth or centralised like in Microsoft .NET Passport. It is possible to login at different points of the CoT thus resulting e.g. in different user names or attributes that are transferred. The identity of the user is not revealed in the process of requests and assertions. For risk assessment purposes we call Liberty identity and attribute federated.

4.3. Attribute-based AAI by Schläger et al.

(Schläger, Sojer et al. 2006) propose a reference architecture for an AAI including attribute-based access control (ABAC). ABAC functionalities are integrated via the open standards SAML and XACML. The idea bases upon the Liberty ID-FF. However, Liberty's distributed paradigm is not followed. In order to mediate between provider and customer demands the architecture is only partly distributed. The user chooses his IdP among a list of third party Identity Providers. He uses the one he trusts most. An external Policy Decision Point (PDP) evaluates the user's access request based on policies as well as user, environment, and resource attributes.

The initial request is directed at the SP. The SP requests an authentication and access control decision from the AAI. The IdP authenticates the user and requests related attributes from all members. The PDP queries the SP for the resource attributes and uses the respective policies loaded at its initialisation. Subsequently, the access decision is computed. The access control decision is forwarded to the SP. Complying with the idea of a generic architecture the SP enforces this decision with its own means.

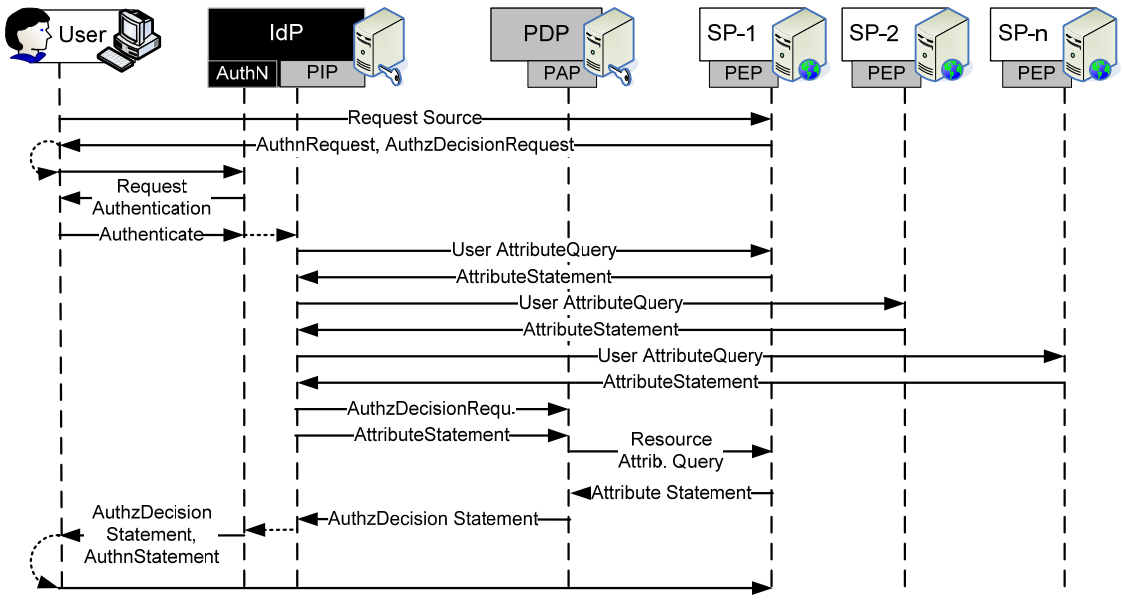

Fig. 7: Schläger et al. attribute-based AAI Processes

5. AAI Effects on Service Providers

Assets under risk are the identities of service providers, attributes about resources and users, the service or the good requested, as well as the system itself. All assets are prone to lose the three major security goals: Integrity, confidentiality, and availability. (Pfleeger and Pfleeger 2003) have shown the types of vulnerabilities one might find for hardware, software, and data. Adopting that notion the vulnerabilities are interruption, for example via a Denial-of-Service-attack, interception of the communication, for example via a Man-in-the-Middle-attack, the modification of the asset, for example attributes granting access to the resource only if the user is over 18 could be changed to access under 16, and finally fabrication of new identities. Fabrication would occur if a bogus merchant is created luring the client to log in with his SSO credentials.

To assess the risk of each asset we make use of the introduced frequency for a successful attack. The frequency is affected twofold – firstly by the technical barrier one raises to prevent an attack *T* and secondly by the motivation of the attacker, the so called return of attack – *RoA*.

$$l_i = f(T_i, RoA_i) \tag{2}$$

$$ALE_i = f(T_i, RoA_i) \cdot L_i \tag{3}$$

The following assumptions are made. Firstly, all equations only consider one period of conducting business. Secondly, the loss at one service provider for a security incident is constant. Thirdly, if one technical barrier is compromised the whole system is corrupted. The effect for the SP is the same regardless of which technical barrier is breached.

The higher *T* the less likely a successful attack occurs; the higher *RoA* the higher the attacker's motivation and the resources employed and consequently the likelier an attack. The RoA_i can be seen as more or less stable as a service provider per se is doing business by offering something of value. He will not stop providing services to minimise risks. However, *T* is completely in the hands of the service provider. Outsourcing security services to an AAI can inflect on *T* and therefore on the frequency of a successful attack.

If the outsourcing of security services inflects the *ALE* the question remains which security services should be outsourced and to what extent. Different AAI approaches and architectures

are able to perform one, all, or combinations of these services. We will take each sub-service and analyse the risks associated as can be seen in Table 1. Each sub-service can be interrupted via a Denial of Service. As the effect is devastating but trivial - no provider or user can conduct business – it is not shown explicitly.

Table 1. Security sub-services with associated risks and consequences for SPs

Service	Risks for provider
Authentication Assertion (AuthN)	**Identity theft:** User identity is forged or intercepted resulting in delivery without access rights. The theft of the provider's identity results in a loss of reputation.
Attribute Assertion (Attrib)	**Attribute forging or modification:** If user attributes are modified, not complete, or added, the following decision cannot be trusted. False denies result in loss of business or the user might be motivated to change the provider. False access can be used for fraud.
	Attribute interception: an attacker could gain knowledge about processes or products.
Policy Decision Assertion (PD)	**Decision forging or modification:** Access could be wrongly denied or granted.
Policy Enforcement (PE)	**Enforcement modification:** Access could be wrongly denied or granted.

5.1. General Implications of AAI Usage

With the usage of an AAI various changes occur in the business surrounding. For once the potential number of customers for one provider enlarges. The number of users (n) of an AAI that merges N service providers is the union of each SP's customer base, at most the sum of all users (4).

$$n_{AAI} \leq \sum_{i=1}^{N} n_i \tag{4}$$

The technical barrier for an attack T is no longer just T_i but has to be seen as a combination of all barriers for the given sub-services, each potentially outsourced: T^{AuthN}, T^{Attrib}, T^{PD}, and T^{PE}. T_i can't be computed by the sum of all barriers but is determined by the minimum: the weakest link in the chain determines its overall strength (5).

$$T_i = \min(T^{AuthN}, T^{Attrib}, T^{PD}, T^{PE}) \tag{5}$$

5.2. AAI Architectures and their Implications

If using an AAI like Microsoft's .NET Passport the authentication of the user is relayed to Passport. The provider uses Passport's technical barrier to prevent misuse of the authentication sub-service for his business. His ALE_i, consequently, depends on the following equation (6):

$$ALE_i = f(\min(T_{AAI}^{AuthN}, T_i^{Attrib}, T_i^{PD}, T_i^{PE}), RoA_i) \cdot L_i \tag{6}$$

For one single provider the loss and supplied return of attack stays the same. However, in the case of a fully developed AAI – like in PAPI – where all security services are outsourced and the AAI provider acts as a proxy for all N service providers, a successful attack on one security service results in a breach of all N vendors. T_i is substituted by T_{PAPI}. The AAI resembles a middleman. Consequently, the RoA is the sum of all returns (7).

$$ALE_i = f(T_{PAPI}, \sum_{i=1}^{N} RoA_i) \cdot L_i \quad \textbf{(7)}$$

(6) is true if the barrier T is set by one AAI provider like Passport or PAPI. However, if the AAI is distributed like the Liberty ID-FF the technical barrier can't be estimated as easily. As *N* service providers act also as identity and attribute providers for other service providers in a CoT and use their own means of authentication the notion of the weakest link, once more takes effect (8):

$$ALE_i = f\begin{pmatrix} \min(\min(T_1^{AuthN},...,T_N^{AuthN}), \\ \min(T_1^{Attrib},...,T_N^{Attrib}), T_i^{PD}, T_i^{PE}), RoA_i \end{pmatrix} \cdot L_i \quad \textbf{(8)}$$

Finally, in the case of the Schläger et al. approach, the weakest of the *j* third party providers for the authentication and the policy decision point sets the barrier for an attack. As the attributes are collected from all members the potential to foist false attributes is leveraged. Therefore, the average technical barrier determines the frequency (9):

$$ALE_i = f\begin{pmatrix} \min(\min(T_1^{AuthN},...,T_j^{AuthN}), \\ \frac{1}{N}\sum_{i=1}^{N} T_i^{Attrib}, T_{Schl.}^{PD}, T_i^{PE}), RoA_i \end{pmatrix} \cdot L_i \quad \textbf{(9)}$$

5.3. Risk Assessment in AAIs

To correctly assess the risks of the usage of an AAI the authors make use of a qualitative method. With the exception of a proxy solution the L_i and RoA_i are not affected by the integration of an AAI. Therefore one can narrow the effect down to the technical barrier of the security sub-service – see (5) and (6). The technical barriers of the four sub-services, when provided by SP_i himself, are taken as the normalised value. The outsourced value T_{AAI} is either higher (↑), less (↓), or equal (↔). Table 2 states the risk assessment. A distinction is being made if the sub-service is federated or centralised.

Table 2.

	$T_{AAI,\ centralised}$	$T_{AAI,\ federated}$	$T_{AAI,\ Schläger\ et\ al.}$
T^{AuthN}	↑: Although n_{AAI} exceeds n_i, the potential of granting SSO with a hard password in a controlled environment argues for a stronger authentication. The usage of complex identification methods like a PKI is more preferable.	↓: As SP_i is in no control of all other SP the weakest link in the chain dictates the barrier for identity theft and alike. A controlled, standardised approach for each SP is not mandatory with the Liberty guidelines.	↑: The idea to use professional IdM provider argues for a stronger authentication, especially through the SSO functionality.
T^{Attrib}	↑: Merging of all attributes balances modified or forged information. With a pattern of the user's behaviour suspicious behaviour can be detected.	↑: Same as centralised approach.	↑: Same as centralised approach.
T^{PD}	↑: Centralised policy decision enables complex, flexible, and specialised access control like	↔: As policy decision making has to be provided by every SP no synergies can be utilised.	↑: Centralised policy decision enables complex, flexible, and specialised access control like

	XACML policies through synergies and a broader information base.		XACML policies through synergies and a broader information base.
T^{PE}	↓: The usage of a central proxy strongly affects the potential *RoA* resulting in a higher l_i (PAPI, (7)).	Not feasible.	Not feasible.

In a distributed architecture there is always the risk of information being altered or intercepted. The usage of SSL, SAML, XML encryption etc. can reduce these risks. In this work we assume that the appropriate techniques have been used and implemented correctly. Users and members of an AAI have to make sure that transferred information stays secure. Table 2 defines consequences for service providers.

5.4. Deciding on AAI Usage

In section 5.3 we have assessed risks depending on different AAI structures and services. Question remains, whether an AAI is economically useful. To determine the cost effectiveness of security investments the RoSI approach has been widely accepted. ALE_{old} depicts the expected loss without additional security investment. C are security costs to reach ALE_{new}, R is additional revenue through strengthened security r^+, e.g. through wider adoption of the service and through cost savings through outsourcing.

$$ALE_{old} - ALE_{new} - C + R = RoSI \qquad \textbf{(10)}$$

For economic reasons the RoSI must be at least positive. Taking into consideration that *ALE* is defined by the weakest point of the security sub-services T^{min} and the cost for sub-service x at SP_i is c_i^x we can deduct two reasons for outsourcing security sub-services.

First, if the outsourced sub-service T_i^x is not $T^{min,}$ but $T_i^x \leq T_{AAI}^x$, then

$$\begin{aligned} ALE_{old} - ALE_{new} &= \Delta ALE = 0 \text{ from (9)} \\ -C + R \geq 0 &\rightarrow C \leq R \rightarrow c_{AAI}^x \leq c_i^x + r^+ \end{aligned} \qquad \textbf{(11)}$$

Meaning that if no strengthened barrier against an attack results out of the decision to use the AAI's service it can be economically reasonable to use the AAI if cost-savings are realised or additional revenue is gained for example through a larger customer base.

Second, if the outsourced sub-service T_i^x is $T^{min,}$ and $T_i^x \leq T_{AAI}^x$, then

$$\begin{aligned} ALE_{old} - ALE_{new} &= \Delta ALE > 0 \text{ from (9)} \\ c_{AAI}^x &< c_i^x + r^+ + \Delta ALE \end{aligned} \qquad \textbf{(12)}$$

The amount to be invested in an AAI is the sum of reduced costs through outsourcing, additional revenue through a larger customer base, and the saved *ALE*. Please note, when changing more than one sub-service in (12) the additional revenue r^+ is not affected proportionally.

Using AAI services does not automatically change the risk assessment of a business process. As seen in Table 2 the decision has to be carefully evaluated if additional risks are worth the enhancements. Furthermore, the decision of outsourcing doesn't have to depend on risks but can be seen as an entirely economic decision (11).

However, as shown in (12) the implication of fewer risks – or lesser *ALE* - motivates higher investments for the AAI usage as well as it sums up to potential savings.

6. AAI Effects on Users

As stated in section 5, an AAI has to meet the three major security goals of integrity, confidentiality, and availability. Identity theft and fraud are the two main concerns of users in e-commerce. For the different security sub-services potential threats are given in Table 3.

Table 3. Security sub-services with associated risks and consequences for users

Service	Risks for users
Authentication Assertion	**Identity theft**: Identity is intercepted and/or misused. Provider's identity is forged and the user lured into signing-in or paying for services never to be received.
Attribute Assertion	**Attribute forging or modification**: If resource attributes are modified, not complete, or added, the following decision can't be trusted. It might be that access or privileges are not granted.
	Attribute interception: A bogus merchant could use attributes to misuse credentials, like a credit card number, conduct illegal profiling, or sell the information.
Policy Decision Assertion	**Decision forging or modification:** Access could be wrongly denied.
Policy Enforcement	**Enforcement modification:** Access could be wrongly denied.

Assuming the service provider himself is acting trustworthily, an attacker could only harm the user if a technical barrier fails. Consequently, the user has a strong interest in high security but is in no position of influencing the barriers directly. An exception has to be made as far as the authentication is concerned. Using weak passwords makes identity theft easy. Service providers usually shy at demanding strong passwords or the usage of a PKI, fearing increased help desk costs or shrinking user acceptance. With a SSO these disadvantages could be reduced. However, as far as privacy and data canniness are concerned the user has to trust the identity provider not to misuse his data. The discussion about .Net Passport and the development of the Liberty ID-FF shows an interest in privacy and a missing user acceptance for doubtable providers. In the end a user has to evaluate privacy aspects versus an increased ease-of-use through SSO and a potentially higher and transparent security system for himself.

6.1. AAI Functionalities

The requested functionalities encompass single sign-on, central account management, reputation sharing over a federation, and the use of pseudonyms. The three AAIs under evaluation and their contribution for the four required functionalities are given in Table 4.

Table 4. General AAI functionalities

Functionality	**Centralised AAI**	**Federated AAI**	**Schläger et al. AAI**
SSO	via one single entry point	via every federation member	via a permanent selection from a list of 3rd party IdPs
Account	at single AAI	only for specific profile	at the selected IdP

Management	provider	at the current entry point into the federation	
Reputation Sharing	possible but not realised	hard – reputation is not combined but only relayed from the entry point to the target	integrated through the accumulation of user data from all providers
Pseudonyms	possible, realised e.g. in PAPI	ID-FF standard procedure (opaque ID)	standard procedure
Sum	**3 of 4**	**2 of 4**	**4 of 4**

Obviously the centralised and the federated approach have been developed in contrast to each other. This comes as no surprise considering the origin of the Liberty Alliance. Centralised approaches make it easy to manage profile data due to the central database. It would also be feasible for centralised approaches to support reputation sharing. However, none of the existing frameworks and products has this feature. In (Schläger and Ganslmayer 2007) this evolution of a centralised AAI has been described. The ID-FF, on the other hand, puts emphasis on privacy protection. It was made deliberately hard to compute one single customer profile or buying pattern in the federation. Since the aim of the last AAI version of Schläger et al. was a compromise between user privacy and profile management this AAI is able to fulfil all four requirements.

6.2. AAI Privacy

In chapter 3.2 we have raised questions about privacy in AAIs. All three approaches have been studied and tested for PETs. The criteria were the separation of identity and privileges (using e.g. pseudonyms), a distributed storage of user attributes, data canniness, and no dictated trust relationships. The separation of identity and privileges is in line with (Feigenbaum 1998). Furthermore it enhances the user's privacy. If the separation of sub-services as depicted in Fig. 1 is used, the step "policy decision" is able to compute an access control decision solely on the attributes of users and resources (subjects and objects). There is no need to identify either of them. The results are given in Table 5.

Table 5. PET AAI functionalities

Functionality	**Centralised AAI**	**Federated AAI**	**Schläger et al. AAI**
Separation of ID and Privileges	no separation	yes, an opaque ID is used together with attributes	yes, the PDP decides only on subject, environment, object attributes and a policy
Distributed Storage	no, all data is centrally stored	yes, for each access only attributes stored at the entry point are used	partly, some data is stored at the IdP, other attributes are scattered over the federation
Data Canniness	no, all user interactions can be logged and used for the AC decision	partly, the user's consent is needed for attribute exchanges	no, all user interactions can be logged and used for the AC decision
No Dictated Trust Relation	No	Yes	Yes
Sum	**0 of 4**	**4 of 4**	**3 of 4**

Again we see a trisection between the analysed AAIs. The centralised approach can be privacy-aware if the provider is a trusted entity. If he secures user information from unauthorised or unwanted access, privacy could be guaranteed. The federated approach on the contrary is privacy-aware even if an IdP can not be trusted. This is realised especially through the distribution of data. The architecture of Schläger et al. ranges in between. User data is accumulated over the federation but due to the separation of identities and attributes not relayed to the SP.

6.3. Deciding on AAI Usage

Referring back to chapter 2.2.1 we see that trust is needed in order to commit a purchase. Trust depends on all AAI elements, the IdP, the technology, and the SP. If we consider system security, we have to refer to section 5.2 and the given formal evaluation of security risks.

A user valuing privacy over functionalities will probably decide on the usage of a Liberty ID-FF based, distributed AAI. If functionalities and ease-of-use are the key argument for AAIs centralised or semi-centralised approaches will prevail. The non-adoption of Microsoft .Net Passport shows, however although trust and privacy may not be the sole decision criterion, doubt on their existence can be enough to boycott the service.

7. Global Effects in a Federation

Besides the individual loss expectancy of one SP the global effects can be assessed and compared as well.

If the same metrics and ratios as in chapter 5.2 are used, the ALE_{global} for all N federation members is (13). ALE_{global} equals the sum of all losses due to successful attacks.

$$ALE_{global} = \sum_{i=1}^{N} \left(f(\min(T_i^{AuthN}, T_i^{Attrib}, T_i^{PD}, T_i^{PE}), RoA_i) \cdot L_i \right) \qquad \textbf{(13)}$$

As can be seen again in (13) the individual technical barriers influence the probability of a successful attack as well as the RoA. It was shown in chapter 5 which technical barriers are influenced by the introduction of an AAI. This chapter will now concentrate on the global effects on the return of attack for an intruder.

For Microsoft's .Net Passport the authentication changes to the same T^{AuthN} for all SPs. An attacker compromising .Net Passport consequently can access all resources. The RoA is therefore the sum of all assets. If he or she attacks any other technical barrier he or she will only get access to the individual resource of the SP in question. For PAPI any of the compromised barriers results in a global access. With Liberty's ID-FF the weakest authentication and attribute provider set the barrier for each a global access. Using the AAI proposed by Schläger et al. the successful attack on the weakest IdP's authentication mechanism and the single PDP results in a global access. For the attribute assertion mechanism the successful attack on one provider results in access to one of the SPs. After the incidence is detected, the computation of the forged attributes together with the additional attributes from the incident forces the attacker to attack anew.

Consequently, we have to distinguish for the AAIs if the attack is performed on a technical barrier of the AAI or the individual one. The explained effects are subsumed in Table 6. The successfully attacked SPs in one period are *k*.

Table 6. Global RoA effects of attacked barrier using an AAI

AAI	RoA of breaking T_{AAI}^{x}	RoA of breaking T_{SPi}^{x}
.Net Passport	$RoA_{AAI}^{AuthN} = \sum_{i=1}^{N} RoA_i$	$RoA_{AAI}^{Attrib} = \sum_{i=1}^{k} RoA_i$

		$RoA_{AAI}^{PD} = \sum_{i=1}^{k} RoA_i$
		$RoA_{AAI}^{PE} = \sum_{i=1}^{k} RoA_i$
PAPI	$RoA_{AAI}^{AuthN,Attrib,PD,PE} = \sum_{i=1}^{N} RoA_i$	N.A.
Liberty ID-FF	$RoA_{weakestSP}^{AuthN} = \sum_{i=1}^{N} RoA_i$	$RoA_{AAI}^{PD} = \sum_{i=1}^{k} RoA_i$
	$RoA_{weakestSP}^{Attrib} = \sum_{i=1}^{N} RoA_i$	$RoA_{AAI}^{PE} = \sum_{i=1}^{k} RoA_i$
Schläger et al.	$RoA_{weakestIdP}^{AuthN} = \sum_{i=1}^{N} RoA_i$	$RoA_{AAI}^{Attrib} = \sum_{i=1}^{k} RoA_i$
	$RoA_{AAI}^{PD} = \sum_{i=1}^{N} RoA_i$	$RoA_{AAI}^{PE} = \sum_{i=1}^{k} RoA_i$

Through the change in the attackers motivation SPs joining an AAI have to carefully consider the depicted global effects for themselves. If the motivation to attack a single SP is perceived as very low, the SP can suddenly become a victim in the AAI due to his or her higher motivation to breach the AAIs security. An SP with a former low probability to get attacked can be attacked along with a highly valuable SP. Special attention needs to be given to Liberty's ID-FF where not only the RoA for a breach of the authentication and the attribute assertion mechanism increases but simultaneously the barrier to do so is set to the weakest mechanism in the whole federation. Based on these findings a participation in an ID-FF federation can not be recommended from a SP's perspective.

8. Conclusion

Our approach permits the analysis of risks in each sub-service in authentication, authorisation and access control deducting formally the factors which are influencing an outsourcing decision. Exclusively motivating AAIs from a technical perspective is not sufficient. It is of high importance to identify the four security sub-services for a system and measure its costs and risks. Accordingly, a service provider can decide on a suitable AAI.

From a user perspective privacy is most respected in the Liberty approach. However, centralised approaches and mediating AAIs offer more functionality. If trust relations are not foisted on the user, adoption is more likely, even if the decision of a suitable IdP is restricted to a limited number.

As derived in chapter 7, the usage of Liberty's ID-FF from a global perspective cannot be recommended. All solutions have the problem of a single point of failure where a single incident can affect all federation members simultaneously. Both, the Schläger et al. approach and Microsoft's .Net Passport, depend on the agreed and contracted security level of the AAI provider. A single service provider needs to respect global effects in order to evaluate if the changed RoA in the AAI increases his individual ALE.

It is a common misconception that anonymity and data canniness are adversaries to secure and dependable business providing on the internet. Even high valuable goods can be sold anonymously if the right security infrastructure is in place. Of course, to make such a scenario work all parties need to trust such an infrastructure and see considerable benefits for them. This research shows that extreme forms of AAIs might not be able to provide this momentum for all user groups.

Acknowledgements

This research benefited massively from the lively and constructive discussions at the Institute for Information Systems at the University of Regensburg. We especially thank our advisors Günther Pernul and Hannes Federrath for their support, input, and ideas. This work also owes thank to the TrustBus 2006 Chairs Simone Fischer-Hübner, Kostas Lambrinoudakis, and Steven Furnell. It was because of their encouragement this article developed and our research was brought forward. Last but not least we thank Nina Benischke for her tireless efforts in proof reading this work.

References

Berinato, S. (2002). Finally, a Real Return on Security Spending. *CIO Magazin*, **Feb. 2002**.

Cantor, S. (2004). *Shibboleth Architecture Protocols and Profiles Working Draft 05*. Accessible at http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-arch-protocols-02.pdf.

Castro-Rojo, R. and López, D. R. (2001). The PAPI system: point of access to providers of information. *Computer Networks*, **37**(6) 703-710.

Cavusoglu, H., Mishra, B. and Raghunathan, S. (2004). A model for evaluating IT security investments. *Communications of the ACM*, **47**(7) 87-92.

Cremonini, M. and Martini, P. (2005). *Evaluating Information Security Investments from Attackers Perspective: the Return-On-Attack (ROA).* Proc. of the Fourth Workshop on the Economics of Information Security, Harvard, USA.

Federrath, H. (2005). *Privacy Enhanced Technologies: Methods, Markets, Misuse.* Proc. of the International Conference on Trust, Privacy & Security in Digital Business (TrustBus '05), Lecture Notes in Computer Science (LNCS), Vol. 3592, Copenhagen, Denmark, Springer.

Federrath, H., Köpsell, S. and Langos, H. (2002). *Anonyme und unbeobachtbare Kommunikaton im Internet.* Proc. of the 32. Jahrestagung der Gesellschaft für Informatik e.v., LNI 19, Dortmund, GI.

Feigenbaum, J. (1998). *Towards an Infrastructure for Authorization - Position Paper.* Proc. of the 3rd USENIX Workshop on Electronic Commerce - Invited Talks Supplement, Boston, Massachusetts, USA, USENIX.

Gefen, D., Rao, V. S. and Tractinsky, N. (2003). *The Conceptualization of Trust, Risk and Their Relationship in Electronic Commerce: The Need for Clarifications.* Proc. of the 36th Annual Hawaii International Conference on System Sciences (HICSS'03), Kohala Coast, HI, USA, IEEE.

Jøsang, A. and Pope, S. (2005). *User Centric Identity Management.* Proc. of AusCERT Asia Pacific Information Technology Security Conference, Gold Coast.

Kilger, M., Arkin, O. and Stutzman, a. J. (2004) In Know Your Enemy: Learning about Security Threats(Ed, Project, T. H.) Addison-Wesley Professional, pp. 505-556.

Liberty Alliance (2004a). *Liberty ID-FF Bindings and Profiles Specification*. Accesible at http://www.projectliberty.org/liberty/content/download/319/2369/file/draft-liberty-idff-bindings-profiles-1.2-errata-v2.0.pdf.

Liberty Alliance (2004b). *Liberty ID-FF Protocols Schema Specification*. Accesible at http://www.projectliberty.org/liberty/content/download/2197/14625/file/draft-liberty-idff-protocols-schema-1.2-errata-v3.0.pdf.

Lopez, J., Oppliger, R. and Pernul, G. (2004). Authentication and authorization infrastructures (AAIs): a comparative survey. *Computers & Security*, **23**(7) 578-590.

McKnight, D. H. and Chervany, N. L. (2001). *Conceptualizing Trust: A Typology and E-Commerce Customer Relationships Model.* Proc. of the 34th Annual Hawaii International Conference on System Sciences (HICSS'01), Kohala Coast, HI, USA, IEEE.

Microsoft Inc. (2003). *Microsoft Passport Review Guide*. Accesible at http://download.microsoft.com/download/a/f/4/af49b391-086e-4aa2-a84b-ef6d916b2f08/passport_reviewguide.doc.

National Bureau of Standards (1974). *Guidelines for Automatic Data Processing, Physical Security and Risk Management, FIPS PUB 31*. Accessible at http://www.governmentsecurity.org/articles/articles2/fips-31.pdf_fl/.

Nowey, T., Federrath, H., Klein, C. and Plößl, K. (2005). *Ansätze zur Evaluierung von Sicherheitsinvestitionen.* Proc. of the Annual "Sicherheit" Conference 2005, Lecture Notes in Informatics, Regensburg, Germany, GI.

OASIS eXtensible Access Control Markup Language Technical Committee *eXtensible Access Control Markup Language (XACML)*. Accesible at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml.

OASIS Security Services Technical Committee *Security Assertion Markup Language (SAML)*. Accesible at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security.

Oppliger, R. (2003). *Security technologies for the World Wide Web*. Boston, Mass., USA, Artech House.

Pfleeger, C. P. and Pfleeger, S. L. (2003). *Security in computing*. Upper Saddle River, NJ, USA, Prentice Hall PTR.

Pfleeger, S. L. (2000). Risky business: what we have yet to learn about risk management. *Journal of Systems and Software*, **53**(3) 265-273.

Schläger, C. and Ganslmayer, M. (2007). *Effects of Architectural Decisions in Authentication and Authorisation Infrastructures*. Proc. of the 2nd International Conference on Availability, Reliability and Security (ARES'07), Vienna, Austria, IEEE.

Schläger, C. and Nowey, T. (2006). *Towards a Risk Management Perspective on AAIs.* Proc. of the International Conference on Trust, Privacy & Security in Digital Business (TrustBus '06), Lecture Notes in Computer Science (LNCS), Vol. 4083, Krakow, Poland, Springer.

Schläger, C., Nowey, T. and Montenegro, J. A. (2006). *A reference model for Authentication and Authorisation Infrastructures respecting privacy and flexibility in b2c E-Commerce.* Proc. of the 1st International Conference on Availability, Reliability and Security (ARES 2006), Vienna, Austria, Vienna, Austria, IEEE.

Schläger, C. and Pernul, G. (2005). *Authentication and Authorisation Infrastructures in b2c E-Commerce.* Proc. of the International Conference on E-Commerce and Web Technologies (EC-Web'05), Copenhagen, Denmark, Lecture Notes in Computer Science (LNCS), Vol. 3590, Copenhagen, Denmark, Springer.

Schläger, C., Sojer, M., Muschall, B. and Pernul, G. (2006). *Attribute-Based Authentication and Authorisation Infrastructures for E-Commerce Providers.* Proc. of the International Conference on E-Commerce and Web Technologies (EC-Web'06), Krakow, Poland, Lecture Notes in Computer Science (LNCS), Vol. 4082, Krakow, Poland, Springer.

Tanenbaum, A. S. and Steen, M. v. (2002). *Distributed systems: principles and paradigms*. Upper Saddle River, N.J., USA, Prentice Hall.

Vidalis, S. (2004). *A Critical Discussion of Risk and Threat Analysis Methods and Methodologie*. Accessible at http://www.glam.ac.uk/socschool/research/publications/technical/CS-04-03.pdf.

Wei, H., Frinke, D., Carter, O. and Ritter, C. (2001). *Cost-Benefit Analysis for Network Intrusion Detection Systems.* Proc. of the CSI 28th Annual Computer Security Conference, Washington, D.C., USA.

Westin, A. F. (1967). *Privacy and freedom*. New York, USA, Atheneum.