Balancing Trust and Risk in Access Control

Alessandro Armando[1,2], Michele Bezzi[3],
Francesco Di Cerbo[3], and Nadia Metoui[1,4]

[1] Security & Trust Unit, FBK-Irst, Trento, Italy
[2] DIBRIS, University of Genova, Italy
[3] SAP Product Security Research, Sophia-Antipolis, France
[4] DISI, University of Trento, Italy

Abstract. The increasing availability of large and diverse datasets (*big data*) calls for increased flexibility in access control so to improve the exploitation of the data. Risk-aware access control systems offer a natural approach to the problem. We propose a novel access control framework that combines trust with risk and supports access control in dynamic contexts through trust enhancement mechanisms and risk mitigation strategies. This allows to strike a balance between the risk associated with a data request and the trustworthiness of the requester. If the risk is too large compared to the trust level, then the framework can identify adaptive strategies leading to a decrease of the risk (e.g., by removing/obfuscation part of the data through anonymization) or to increase the trust level (e.g., by asking for additional obligation to the requester). We outline a modular architecture to realize our model, and we describe how these strategies can be actually realized in a realistic use case.

Keywords: trust, privacy, risk

1 Introduction

The increasing availability of large and diverse datasets (*big data*) calls for increased flexibility in access control so to improve the exploitation of the data. Indeed, organizations are now in the position to exploit these diverse datasets to create new data-based businesses or optimizing existing process (real-time customization, predictive analytics, etc.). Yet, they are often unable to fully leverage this potential due to the lack of appropriate data release mechanisms ensuring that sensitive information is not disclosed. As a consequence, most organizations still strongly limit (even internally) the sharing and dissemination of data making most of the information unavailable to decision-makers, and thus they do not fully exploit the power of these new data sources.

To overcome the problem, access control systems must weigh the risks against the trustworthiness of the incoming requests. In other words, access control decisions must be based on an estimation of expected cost and benefits, and not (as in traditional access control systems) on a predefined policy that statically defines what accesses are allowed and denied. In other words, in Risk-based Access control for each access request, the corresponding risk is estimated and if the risk is less than a threshold then access is guaranteed, otherwise it is denied. The aim is to be more permissive than in traditional access control system by allowing for a better exploitation of data. Although existing risk-based access control models provide an important step towards a better management and exploitation of data, they have a number of drawbacks which limit their effectiveness. In particular, most of the existing risk-based systems only support binary access decisions: the outcome is *allowed* or *denied*, whereas in real-life we often have exceptions based on additional conditions (e.g., *"I cannot provide this information, unless you sign the following non-disclosure agreement."* or *"I cannot disclose these data, because they contain personal identifiable information, but I can disclose an anonymized version of the data."*). In other words, if the system can propose appropriate risk mitigation measures, and they are accepted by the requester, a relevant part of additional information can be shared.

In this paper we propose a novel access control framework that combines trust with risk and supports access control in dynamic contexts through trust enhancement mechanisms and risk mitigation strategies. This allows us to strike a balance between the risk and the trustworthiness of the data request. If the risk is too large compared to the trust level, then the framework can identify adaptive strategies that can decrease the risk (e.g., by removing/obfuscating part of the data through anonymization) and/or increase the trust level (e.g., by asking for additional obligation to the requester).

Our framework enjoys a number of features:

1. it explicitly models trust and risk, which are the key factors of any business decision;
2. it increases the flexibility of existing risk-aware access control, by introducing trust;
3. it supports complex authorization scenarios by simply changing the configuration (trust and risk configuration modules, and corresponding mitigation/enhancement strategies).

The use case illustrates how the framework can work in practice, addressing access control requirements in a natural way, that would otherwise need complex authorization structure and calibration.

In the next section we provide a motivating use case. In Section 3 we introduce our risk- and trust-based access control model. In Section 4 we provide an architectural view of our access control framework. With the reference to our use case, in Section 5 and in Section 6 we discuss approaches to risk evaluation and mitigation that can be supported by our framework. In Section 8 we discuss the related work and in Section 9 we conclude with some final remarks.

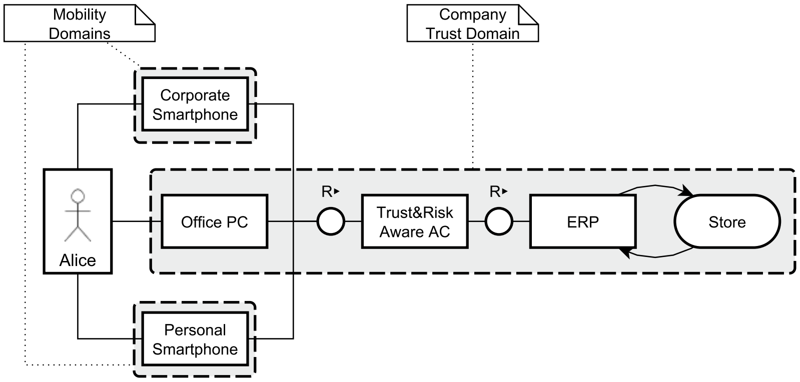

Fig. 1. Use Case: Alice accessing an HR report with personal data covered by EU Directive on Data Protection 95/46/EC.

2 Use Case

Consider a company with an ERP system with a Human Resource (HR) Management module, enabled with the proposed Trust and Risk-aware access control system (see Fig. 1). By using the ERP functionalities corporate user Alice can generate an HR report containing a list of employees with their location and salaries. The report contains sensitive information and personal data, and the company has strict rules for accessing the data such as security measures to minimize the disclosure risk when data are moved outside the company. The risk scenario considered is the leakage of the the salary information associated to a specific employee (re-identification risk). To ensure compliance with EU data protection laws, additional restrictions must be applied if data are accessed outside EU.

In her daily business, Alice may access the report using multiple devices: her office PC at corporate premises, a corporate smartphone and her own smartphone. Access in mobility suffers from a high level of risk, since it is more exposed to external attacks and, depending on the geographical location, different rules may apply. A conservative approach, easily implementable with traditional access control systems, would imply a security policy like that:

- if Alice is on premises, then access is granted
- if Alice is in mobility, access is denied as the security and compliance risks could be too high

Basically, access is limited to corporate premises, where full data can viewed whereas outside no information is available and no reports can be produced. Even though this approach could seem simplistic, many real-life access control systems offer a similar level of functionality [1].

Ideally, Alice would like to get a wider access to the data, and perform her business tasks (e.g., reporting) also in mobility, using different devices in multiple locations, but still keeping security risk under control, as summarized in Table 1.

Table 1. Possible usage scenarios, comprising different devices and locations, and expected utility (i.e., type of reports needed) and security levels

Scenario				Expected	
#	Device	Location	Administration	Utility	Security
1	PC	on premises	corporate	full access	no restriction
2	Smartphone	EU	corporate	grouped by country	medium risk
3	Smartphone	EU	personal	grouped by region	minimal risk
4	Smartphone	no EU	-	no access	no access

In the next section, we will show how these scenarios can be realized in our framework.

3 Model

At an abstract level, a risk- and trust-based access control framework can be represented by a function $Auth(Obj, u, p)$ defined as follows. User u is granted permission p on object obj iff the trustworthiness of the incoming request is larger or equal to the risk, i.e.,

$$Auth(Obj, u, p) = \begin{cases} \textbf{Allow}, & \text{if } (T(u, C) - R(obj, p, C)) \geq 0 \\ \textbf{Deny}, & \text{otherwise} \end{cases} \qquad (1)$$

where $T(u, C)$ is the trustworthiness of the request, which depends on user u and context information C (e.g., location of the requester) and $R(obj, p, C)$ is the risk, which depends on the requested object obj (e.g. a file) to the permission p (e.g., read or write)[1] and context C.

If access is denied, then there are two possible methods for improving the accessibility to the resource: *(i)* applying risk mitigation strategies to decrease R, or *(ii)* increase the trustworthiness T, until the condition condition $T > R$ is granted. We will discuss how risk and trust can be modeled, and possible mechanisms to reduce/increase their values in the next subsections.

[1] In most cases the dependency of risk from permission is mediated by roles. For the sake of simplicity, we do not consider here roles, for an extension of this model including roles, we can follow the lines of the models described in [2,3].

3.1 Modeling Trust

Trust is a wide concept, and different definitions have been proposed in literature [4]. To our scope we can use the definition by McKnight and Chervany [5], which better related to the concepts of utility and risk attitude. [2]

> *Trust is the extent to which one party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible.*

Several definitions have been proposed in for the concept of Trust In our case, we consider trust expressing the level of confidence the resource controller has on the user u not misusing the resource he wants to access. We expect this level to depend on the user u (identity, role, and previous behavior) and on the given context C (e.g., the device or system environment he is using).

Trust values are assigned in various ways depending on the specific use cases. For example in reputation models, trust assessments from other entities are combined to compose a trust evaluation, or in behavioral trust a value is assigned based on the historical records of transactions [4]. Trust can be also derived from assessing a set of *trust indicators* such as security metrics (e.g., level of authentication) and from trust assertions (e.g. stamp of approval) issued by trusted entities (i.e., certification authorities).

From the risk-based system point of view, the identity of the requester heavily depends on the effectiveness of the authentication mechanism employed. To take into account this, the trustworthiness of user u in context C, say $T_{\mathit{eff}}(u, C)$, should take into account the possibility that the authentication is not carried out correctly (e.g., an identity theft scenario). This situation can be modeled in our framework by replacing $T(u, C)$ with $T_{eff}(u, C)$ in (1), where

$$T_{eff}(u) = T(u)(1 - P_{it}) + T(u' \neq u)P_{it}$$

where $T(u' \neq u)$ is the Trust associated to any, not specified, other user that is not u, in practice it should be zero or negligible and P_{it} is the probability of an identity theft. P_{it} represents the strength of the authentication mechanisms.

3.2 Modeling Risk

Risk is defined by the likelihood and the impact of the occurrence of one or more a series of failure scenarios $s \in S$ (also called risk scenarios). Although different quantitative risk methodologies exist, see [7] and references therein, for independent scenarios as risk can be computed by:

$$R(obj, p) = \sum_{s \in S(C)} P(s)I(s)$$

[2] A popular used definition is from Gambetta [6], which stresses the *reliability* aspects of trust. For a discussion see [4].

where S is the set of possible failure scenarios related to the access of p in the context C, $P(s)$ is the probability of occurrence of the failure scenario s, and $I(s)$ the associated impact (often measured as monetary cost).

The risk exposure can be decreased implementing a set of controls and mechanisms, and in this case we refer it as residual risk. In addition, temporary risk mitigation strategies can be applied to further reduce the risk. In case of access control, they include for example, decreasing the probability of failure, by obfuscating (part of) the data (e.g., anonymization) or imposing usage control restrictions (e.g., data retention period); or decreasing the impact, by insurance.

Eq. 1 implies that trust and risk are measured in the same units. Ideally, risk should be measured in monetary units (since the impact is the cost of occurrence of a certain scenario), and, accordingly, trust should have the same units, as in the previous example for financial transactions. Unfortunately, estimating risk in information systems is much less consolidated practice, due to: i) the limited availability of historical data on failure scenarios, which makes difficult to estimate the corresponding probabilities. ii) the difficulty to estimate the impact of a failure to protect an intangible digital assets.[3]

To overcome these problems, existing risk based access control systems use various approaches: they estimate these values from the parameters of traditional (non-risk based) access control models (e.g., see [9] for multi-level security models), they use relative measures for both trust and risk (in practice they normalize these quantities in the interval $[0, 1]$, see [3]), or they use heuristics for estimating these numbers from qualitative risk assessments [7].

In the sequel, to demonstrate our approach, we will consider a single risk factor related to data privacy (re-identification risk). This allows us to compare trust, normalized in the interval $[0, 1]$, directly with the probability of the risk scenario. The model can be clearly include any other security risk factors, as far as a quantitative risk estimation is possible, for example, deriving risk values from the rating of the Common Vulnerability Scoring System (CVSS) [10].

4 Architecture

In this section we present an abstract architecture for our Trust and Risk-aware Access Control Framework. The architecture, depicted in Figure 2, is composed of four main modules that are described in the remaining part of the section. To better illustrate their functionalities, we will use the use case as running example, i.e., we focus on re-identification risk and anonymization as risk mitigation strategies, and obligation as a means for trust enhancement. This architecture, however, is conceived to deal with arbitrary strategies and trust/risk functions.

Risk-Aware Access Control Module. is the entry point to our system, through which users can submit requests to retrieve data from the underlying database.

[3] For these reasons, so far, most of the risk assessments for information system are qualitative, where probability and impacts are classified in broad categories and no explicit numerical values are assigned (e.g., in many application of ISO 27005:2011 [8]).

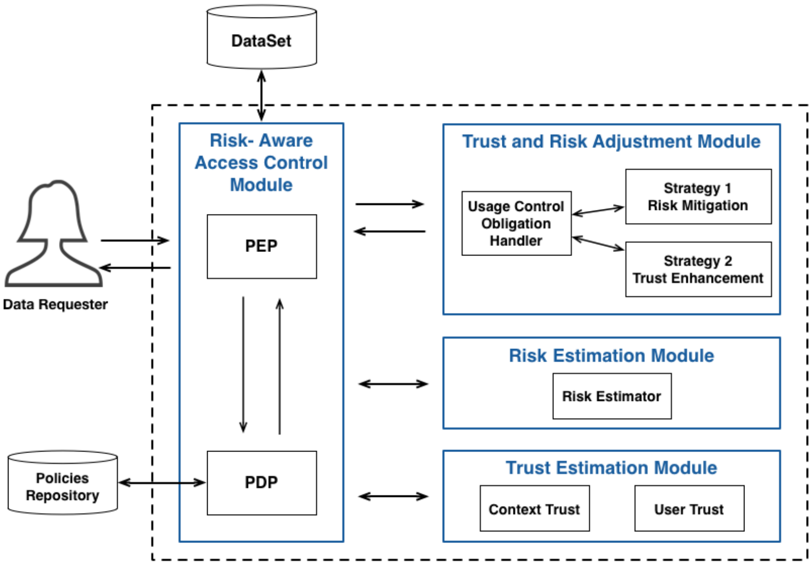
DataSet
Risk- Aware Access Control Module
PEP
PDP
Data Requester
Policies Repository
Trust and Risk Adjustment Module
Usage Control Obligation Handler
Strategy 1 Risk Mitigation
Strategy 2 Trust Enhancement
Risk Estimation Module
Risk Estimator
Trust Estimation Module
Context Trust
User Trust

Fig. 2. Architecture of the Trust and Risk-aware Access Control framework

The module evaluates the access authorizations of the data requester and grants or denies access. To do so, the Risk-Aware Access Control module will call the Risk Estimation Module to determine the risk level of the request and the Trust Estimation Module to determine the requester's and context trust. Trust and Risk Mitigation Module enters into play to increase trust and reduce risk if necessary.

This module is realized internally with a PEP-PDP pair (a Policy Enforcement Point and Policy Decision Point respectively). A PIP (Policy Information Point) is used to provide additional attributes for the requester and the context.

This module is realized internally with a PEP-PDP [4] pair (a Policy Enforcement Point and Policy Decision Point respectively). A PIP [5] (Policy Information Point) is used to provide additional attributes for the requester and the context.

[4] In the XACML (eXtensible Access Control Markup Language) standard [11] the PDP is the point that evaluates an access request against an authorisation policy and issues an access decision and the PEP Policy Enforcement Point is the point that intercept user's request call the PDP for an access decision then enforce this decision by allowing or denying the access.

[5] The PIP (in XACML) is the point that can be called to provide additional informations about resource, requester or envirment.

Risk Estimation module. It is used to determine the level of risk, based on the data requested, context and on the criteria defined in the risk estimator configuration. For example, for re-identification risk, this configuration includes the metrics respect to domain-specific knowledge about what information is to be considered critical or not. Besides the evaluation of the risk, this module produces an estimation of the minimal anonymization level to be applied in order to meet the required risk level. (i.e., in case of k-anonymity metric, the risk estimation module computes the minimal value of k that respects the risk threshold constraint, see Section 5).

Trust Estimation module. It is used to compute the trust aspects of a request. In particular, it can take into account user attributes like role, organizational units, age for instance; as well as context attributes, like the geographic location where a request was created, the software used to issue it as well as the characteristics of a network connection.

Trust and Risk Adjustment module. This module is activated by the Risk-Based Access Control module in order to mitigate risk and/or to increase the trust level for the request, when the access risk to the requested resource exceeds the trust level, in such a case, two possible options are available:

- *Decrease Risk* by applying optimal risk mitigation strategies (e.g., anonymization operations, which decrease risk but minimise the information loss)
- *Enhance Trust* by complementing an authorization decision with access and usage control obligations (explained in the following Sect. 6). These obligations can condition the acceptance of a request to the execution of operations at the moment of the request or when specific events occur; for example, an usage control obligation may prescribe the deletion of a resource after that a retention period expires. If such obligation is guaranteed to be enforced, then the trust estimation can be increased.

The usage control obligations Handler will enforce the selected risk adjastment strategy.

5 Risk Evaluation and Mitigation

The risk of privacy violations is often associated with the concept of *individual identifiability,* used in most privacy laws e.g. EU data protection directive [12], Health Insurance Portability and Accountability Act (HIPAA) [13], etc.). Immunity against individual identifiability can be interpreted as taking measures to prevent the identification or learning of private information about any individual in a dataset with probability or confidence greater than a certain threshold [14], or, in other terms, ensuring a certain level of anonymity to each individual in the dataset. In this section we will discuss how re-identification risk can be measured and how anonymization can be used to mitigate the risk associated with a tabular-data extraction request while querying a privacy-sensitive dataset, as in our use case.

From a privacy perspective, attributes (or columns) in a dataset can be classified as follows:

- *Identifiers.* These are data attributes that can uniquely identify individuals. Examples of identifiers are Name/Last Name, the Social Security Number, and the passport number.
- *Quasi-identifiers (QIs) or key attributes* [15]. These are the attributes that, when combined, can be used to identify an individual. Examples of QIs are the postal code, age, job function, location, and gender.
- *Sensitive attributes.* These attributes contain intrinsically sensitive information about an individual (e.g., diseases, political or religious views, income) or business (e.g., salary figures, restricted financial data or sensitive survey answers).

In presence of identifiers the re-identification risk is clearly maximum (i.e., probability of re-identification $P = 1$), but even if identifiers are removed, combining QIs individuals can be singled out and this implies a high risk. To measure this risk various privacy metrics have been proposed in the literature. These metrics differ in a number of ways, but they all express the risk of disclosing personal-identifiable information when releasing a given dataset (see [16,14] for a review). In the context of re-identification, the most popular one is k-anonymity [17][6]. The k-anonymity condition requires that *every* combination of QIs is shared by at least k records in the dataset. A large k value indicates that the dataset has a low re-identification risk, because, at best, an attacker has a probability $P = 1/k$ to re-identify a record (i.e., associate the sensitive attribute of a record to the identity of a respondent). For example, the (unaltered) table in Table 4, has clearly $k = 1$ (name/lastname are unique identifiers), and $P = 1$ for ri-identification risk (for the sake of simplicity we do not consider impact here).

A possible way to decrease the disclosure risk is anonymization. Anonymization is a commonly used practice to reduce privacy risk, consisting in obfuscating, in part or completely, the personal identifiable information in a dataset. Anonymization methods include [19]:

- *Suppression:* Removal of certain records or part of these records (columns, tuples, etc., such name/last name column);
- *Generalization:* Recoding data into broader classes (e.g., releasing only the first two digits of the zip code or replacing towns with country or regions) or by rounding/clustering numerical data;

Traditionally, anonymization is run offline, but more recently risk-based access control models, which use in-the-fly anonymization as mitigation strategy have been proposed [20]. In practice, to minimize the risk to a certain level (compared to the trust threshold), anonymization methods are applied to decrease the probability of re-identification (or, in other words, by increasing k), common

[6] Other privacy metrics exist (for example, ℓ-diversity, and t-closeness, see [18] for a review), but k-anonymity is still a *de-facto* standard in real applications

Table 2. Re-identification risk for different anonymization methods

Risk	Anonymization Method
Full risk	$k = 1$ ($P = 1.0$)
Medium risk	$k = 2$ ($P = 0.5$)
Minimal risk	$k = 10$ ($P = 0.1$)

values for k in anonymized data are in range between 2 and 30 [21], depending on the use cases.

In our example, we can take $k = 2$ ($P = 0.5$) as *medium* risk (see the 'expected security' column in Table 1) and $k = 10$ ($P = 0.1$) for *minimal* risk. Table 2 summarizes the values for k in relation with the different risk expectations.

6 Risk Mitigation by Obligation

Obligations are actions or operations that must be carried out as result of an authorization decision. In the standard XACML architecture [11], obligations are defined as parts of policies and included in authorization responses created by the PDP; they are enforced by the PEP on behalf of the subject issuing the authorization request. Besides their application as outcome of authorization decisions, obligations may also be applied during or after the consumption of a requested resource or the execution of a requested operation [22,23]: for example, a policy may state a specific retention period for any copy of a resource whose access was granted to the requester. In these cases a trusted component must exist that is able to operate in real time as a PEP. This situation is generally referred as *Usage Control* (UC) [24]. UC models and mechanisms have been proposed to address privacy requirements [25], and applied to both the cloud and the mobile environments [26,27].

AC/UC policy definitions may comprise a broader set of directives, regulating runtime aspects originated from an authorized access; for example, a policy may prescribe to monitor the location where a mobile user consumes a resource and to react with a deletion obligation in case the user leaves the country. Such capabilities are particularly useful to achieve compliance with directives (law requirements or corporate policies): for example, data privacy regulations introduced in Section 5 impose the application of certain principles and UC can enforce automatically some aspects [28].

Therefore, the usage of obligations, when their enforcement is guaranteed, can be considered as a means to enhance a request's trust estimation in our proposed system. In fact, it can be assumed that prescriptions specified by a security policy are applied and that they can regulate how resources or operations are used, thus ensuring their compliance. For instance, in our use case the trust level could change with the context as shown in Table 3. For the sake of simplicity we assume that trust is independent from the specific user, i.e., $T(u, C) = T(u', C)$ for all contexts C and users u and u'. In the for the most trusted environment

Table 3. Trust values in different contexts C

Context	$T(u, C)$
On premise	1.0
Mobility (secure)	0.5
Mobility (standard)	0.1
Mobility (outside EU)	0.0

Table 4. HR report: original view

Name	Job	Location	**Salary**
Timothy Lulic	Senior Developer	London	46200
Alice Salamon	Support	London	45000
Perry Coda	Junior Developer	London	32000
Tom Torreira	Admin	Milan	28000
Ron Savic	Senior Developer	Rome	56000
Omer Regini	Senior Developer	Shanghai	47000
Bob Eramo	Support	Macau	18000
Amber Mesb	Admin	Bangalore	30000
Elise Moisander	Admin	Bangalore	31000

(On premise) we can thus have $T(u, C) = 1$, whereas for requests coming from outside the EU that cannot be trusted and thus $T(u, C) = 0$.

7 Application to the Use Case

We now show how our framework can support the scenarios introduced in Section 2 and achieve the expected utility and security levels. In all scenarios considered, we assume user Alice requests access to the data listed in Table 4.

Scenario #1: Access from business environment. The Risk Estimation module is called to estimate the re-identification risk associated to the dataset: $R(obj, p, C) = 1$, since the report contains personal data with an elevate re-identification risk. The Trust Estimation module in turn computes the trust associated to the context where the request is originated: $T(u, C) = 1$, since Alice is in her office. Therefore, $Auth(Obj, u, p) =$ **Allow** and therefore access is granted.

Scenario #2: Access in mobility from EU using corporate smartphone. Since the request is performed in mobility $T(u, C) = 0.1$ and while $R(obj, p, C) = 1$. The Trust and Risk Adjustment module then triggers the trust enhancement and risk mitigation strategies. Specific AC/UC obligations are thus assigned to the report (e.g., do not share, delete after 2 hours, only usable in EU) to be enforced by an obligation enforcement engine deployed on the corporate smartphone. The application of these measures increases the trust in the context to $T(u, C) = 0.5$. To

Table 5. HR report: anonymized view with $k = 2$.

Name	Job	Location	**Salary**
***	***	UK	46200
***	***	UK	45000
***	***	Italy	32000
***	***	Italy	28000
***	***	Italy	56000
***	***	China	47000
***	***	China	18000
***	***	India	30000
***	***	India	31000

Table 6. HR report: anonymized with $k = 4$.

Name	Job	Location	**Salary**
***	***	EMEA	46200
***	***	EMEA	45000
***	***	EMEA	32000
***	***	EMEA	28000
***	***	EMEA	56000
***	***	APAC	47000
***	***	APAC	18000
***	***	APAC	30000
***	***	APAC	31000

decrease risk, k-anonymity with $k = 2$ allows to reduce the re-identification risk to 0.5. Therefore, $Auth(Obj, u, p) =$ **Allow** and Alice receives the anonymized view of Table 5.

Scenario #3: Access in mobility from EU using personal smartphone. This scenario is similar to the previous one, with the notable exception that now no trust enhancing measures can be enforced on the mobile phone. Therefore, the Trust and Risk Adjustment module can only apply the risk mitigation strategy, by using k-anonymity with an greater value of k, i.e. $k = 10$, that will result in a re-identification risk of 0.1. Thus, $T(u, C) = 0.1$ (access though personal smartphone), $R(obj, p, C) = 0.1$ (after applying k-anonymity with $k = 10$), and thus $Auth(Obj, u, p) =$ **Allow**. The report received by Alice in this scenario is given in Table 6[7].

Scenario #4: Access in mobility from outside EU with personal smartphone. In this case, the risk of violating the regulations is maximum. This means that the trust in the environment is 0, no mitigation strategies may be adopted and

[7] Table 6 is just included as exemplification and depicts the result of k-anonymity for $k = 4$.

therefore $T(u, C) = 0$ (request from outside EU), $R(obj, p, C) = 1$, and thus $Auth(Obj, u, p) =$ **Deny**.

8 Related Work

Several approaches has been recently proposed to address the limitations of traditional access control models in terms of lack of flexibility, inability to handle contextual information, evaluation of the trustworthiness of users and in managing access risk.

Context-aware access control models propose the use of contextual information to determine access to resources, e.g. determining the decision based on temporal[29], or more general, environmental conditions [30], also in combination with risk models [31]. However these models, mostly, define in a static manner the context parameters with which the access to resources will be granted or denied.

A more dynamic approach is taken in risk and trust based access control models (e.g. [32,33,34,35,36]), where for each access request or permission activation, the corresponding risk is estimated and if the risk is less than a threshold (often associated with trust) then access is guaranteed, otherwise it is denied. Cheng et al. [34], following the multi-level-security paradigm, compute risk and trust thresholds from the sensitivity labels of the resource and clearance level of the users. They also consider what we define a trust enhancement mechanism (the authors call it risk mitigation strategy in their paper) that provides users with a limited amount of *tokens*, which allow them to access resources with risk higher than their trust level. The details on how this mechanism can be applied in real cases are not provided.

In another work, Chen et al. [32] introduced an abstract model which allows role activation based on a risk evaluation compared to predefined risk thresholds. Trust values are considered, and they impact (decreasing) risk calculation. If risk is too high, the model includes mitigation strategies, indicated as (system) obligations. The paper does not specify how to compute the risk thresholds, trust, and the structure and impact of obligations. In a derived model [33], mitigation strategies have been explicitly defined in terms of user obligations in addition to system obligation. An user obligation describes some actions that have to be fulfilled by the user to get access. Although the model does not consider explicitly trust, it introduces the concept of *diligence score*, which measured the diligence of the user to fulfill the obligations (as in behavioral trust model), and impact the risk estimation. An extension of the model proposed by Chen et al. [32] has been recently proposed [20], such work focuses on re-identification risk and anonymization is used as mitigation strategy (as in our paper).

Following the original Chen et al. [32] model, these papers consider trust as part of the risk value. As a consequence: *i)* trust enhancement and risk mitigation strategies are mixed, and it becomes difficult to find an optimal set of strategies to increase access, keeping risk under control, *ii)* trust thresholds become dependent on the risk scenario, decreasing the flexibility in presence of multiple risk factors. Our model solves these issues, clearly separating trust aspects from risk.

The impact of obligations on trust is also considered in other studies. We can distinguish between two categories of obligations: *provisions* or *pre-obligations* [37] are actions that must be executed as a pre-condition for authorization decision; *post-obligations* are actions that must be fulfilled after the authorization decision is made. In [38], the trust value of an user is impacted by his previous history of fulfilling or not post-obligations, also considering their level of criticality.

Other approaches have also incorporated user trust in privilege assignment Dimmock et al. [39] propose a framework where users are assigned to roles according to their trust level. Baracaldo et al. extended this idea in [3] and propose to mitigate access risk by lowering the trust level of misbehaving users in order to (temporarily) revoke critical privileges. This model includes separation of duties constraints in risk computation (which we do not consider in our), particularity relevant for addressing insider attack risk scenarios.

These models can be incorporated in the computation of the trust values in our model. Indeed in the scenario we proposed in Sect. 6, obligations increases the trust value, however we do not consider the history of previous obligation fulfillment, since we rely on the secure environment for assuring their enforcement.

9 Future work and Conclusions

Motivated by the need to balance the advantages of big data availability, and stringent security and privacy requirements, novel access control paradigms are emerging. Risk plays a central role, and access control decisions can mimic the business decision process, where risk is assessed relatively to trust. We have proposed an access control framework based on these two factors (trust and risk) and showed that it can address complex authorization requirements by dynamically applying strategies for risk mitigation and trust enhancement. The possibility to play with both risk and trust at the same time and its application to a real use case are the main novelties of our work. Our framework can also be combined with more classical (policy-based) approach, as described in [20] for risk-based access control.

Although promising, our approach presents a number of open issues to be solved for a practical usage. In particular, the overall approach (as for any quantitative risk model) relies on the numerical estimation of risk and trust. These quantities are difficult to compute. Indeed, the diversity of risk scenarios, the intangible nature of trust, and the limited amount of historical data for incidents make an accurate quantitative assessment extremely difficult. As also shown in our paper, using some heuristics it is possible to derive sound relative estimation (i.e., using dimensionless units) for trust and risk, in some specific usage scenarios, but a general approach applicable to multiple use cases is missing. Ideally, we should estimate trust and risk in terms of monetary value, which has several advantages: 1) it provides a common *unit of measure* to combine risk and trust factors of very different nature (e.g., security risk, compliance risk, privacy risk or trust from reputation systems, trust-factors, behavioral analysis), 2) it is easy to understand for non-technical experts 3) it can be easily combined with risk mitigation and trust enhancement strategies that have a clear monetary value (e.g., insurance, certifications, legal contracts, trusted devices). In this respect, it is particularly interesting the emergence of new cyber-insurance models (building on techniques derived by the financial sector, e.g. Value-at-risk, Monte-Carlo simulations) to compute the values of cyber-risk and hence the cost of insurance premiums. [40].

In the short term, we want to validate our model on other use cases, where some quantitative methods are, even partially, available (either using dimensionless units or monetary values). We will also investigate the impact of authentication mechanisms on trust (as hinted in Section 3), and based on estimated probability of authentication success [41], to devise optimal strategies which combine multiple authentication methods according the risk associated to the request.

Acknowledgments The research leading to these results has received funding from the FP7 EU-funded project Coco Cloud (grant No. 610853) and SECENTIS (FP7-PEOPLE-2012-ITN, grant no. 317387).

References

1. Trabelsi, S., Ecuyer, A., y Alvarez, P.C., Di Cerbo, F.: Optimizing access control performance for the cloud. In Helfert, M., Desprez, F., Ferguson, D., Leymann, F., Muñoz, V.M., eds.: CLOSER 2014 - Proceedings of the 4th International Conference on Cloud Computing and Services Science, Barcelona, Spain, April 3-5, 2014., SciTePress (2014) 551–558
2. Chen, L., Crampton, J.: Risk-aware role-based access control. In Meadows, C., Fernandez-Gago, C., eds.: Security and Trust Management. Volume 7170 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2012) 140–156
3. Baracaldo, N., Joshi, J.: An adaptive risk management and access control framework to mitigate insider threats. Computers and Security **39, Part B**(0) (2013) 237 – 254
4. Josang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision Support Systems **43**(2) (2007) 618 – 644 Emerging Issues in Collaborative Commerce.
5. Mcknight, D.H., Chervany, N.L.: The meanings of trust. Technical report (1996)
6. Gambetta, D.: Can we trust trust? In: Trust: Making and Breaking Cooperative Relations, Basil Blackwell (1988) 213–237
7. Celikel, E., Kantarcioglu, M., Thuraisingham, B., Bertino, E.: A risk management approach to RBAC. Risk Decis. Anal. **1**(1) (2009) 21–33
8. ISO: Iec 27005: 2011 (en) information technology–security techniques–information security risk management switzerland. ISO/IEC (2011)
9. Cheng, P.C., Rohatgi, P., Keser, C., Karger, P.A., Wagner, G.M., Reninger, A.S.: Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In: SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy. (2007) 222–230

10. Houmb, S.H., Franqueira, V.N.L., Engum, E.A.: Quantifying security risk level from cvss estimates of frequency and impact. J. Syst. Softw. **83**(9) (September 2010) 1622–1634
11. Moses, T., et al.: extensible access control markup language (xacml) version 2.0. Oasis Standard **200502** (2005)
12. of Europe, C.: Handbook on european data protection law. Technical report (2014)
13. Scholl, M.A., Stine, K.M., Hash, J., Bowen, P., Johnson, L.A., Smith, C.D., Steinberg, D.I.: Sp 800-66 rev. 1. an introductory resource guide for implementing the health insurance portability and accountability act (hipaa) security rule. Technical report (2008)
14. Clifton, C., Tassa, T.: On syntactic anonymity and differential privacy. Trans. Data Privacy **6**(2) (August 2013) 161–183
15. Dalenius, T.: Finding a needle in a haystack-or identifying anonymous census record. Journal of official statistics **2**(3) (1986)
16. Bezzi, M.: An information theoretic approach for privacy metrics. Transactions on Data Privacy **3**(3) (2010) 199–215
17. Samarati, P.: Protecting respondents' identities in microdata release. IEEE Trans. Knowl. Data Eng. **13**(6) (2001) 1010–1027
18. Fung, B.C.M., Wang, K., Chen, R., Yu, P.S.: Privacy-preserving data publishing: A survey of recent developments. ACM Comput. Surv. **42**(4) (June 2010) 14:1–14:53
19. Ciriani, V., De Capitani di Vimercati, S., Foresti, S., Samarati, P.: Theory of privacy and anonymity. In Atallah, M., Blanton, M., eds.: Algorithms and Theory of Computation Handbook (2nd edition). CRC Press (2009)
20. Alessandro Armando, Michele Bezzi, N.M., Sabetta, A.: Risk-aware information disclosure. In: Lecture Notes in Computer Science (LNCS). Volume 8872. (2014)
21. on Strategies for Responsible Sharing of Clinical Trial Data;, C.: Sharing Clinical Trial Data: Maximizing Benefits, Minimizing Risk. National Academies Press (US), Washington (DC) (2015)
22. Mont, M.C., Beato, F.: On parametric obligation policies: Enabling privacy-aware information lifecycle management in enterprises. In: Policies for Distributed Systems and Networks, 2007. POLICY'07. Eighth IEEE International Workshop on, IEEE (2007) 51–55
23. Ali, M., Bussard, L., Pinsdorf, U.: Obligation language for access control and privacy policies. (2010)
24. Sandhu, R., Park, J.: Usage control: A vision for next generation access control. In Gorodetsky, V., Popyack, L., Skormin, V., eds.: Computer Network Security. Number 2776 in Lecture Notes in Computer Science. Springer Berlin Heidelberg (January 2003) 17–31 00208.
25. Ardagna, C.A., Cremonini, M., De Capitani di Vimercati, S., Samarati, P.: A privacy-aware access control system. Journal of Computer Security **16**(4) (2008) 369–397
26. Pretschner, A., Hilty, M., Basin, D.: Distributed usage control. Communications of the ACM **49**(9) (2006) 39–44
27. Di Cerbo, F., Doliere, F., Gomez, L., Trabelsi, S.: Ppl v2.0: Uniform data access and usage control on cloud and mobile. In: Proceedings of the 1st International Workshop on TEchnical and LEgal aspects of data pRIvacy and SEcurity, IEEE (2015)
28. Trabelsi, S., Sendor, J., Reinicke, S.: Ppl: Primelife privacy policy engine. In: Policies for Distributed Systems and Networks (POLICY), 2011 IEEE International Symposium on. (June 2011) 184–185

29. Bertino, E., Bonatti, P.A., Ferrari, E.: Trbac: A temporal role-based access control model. ACM Trans. Inf. Syst. Secur. **4**(3) (2001) 191–233
30. Bonatti, P., Galdi, C., Torres, D.: Erbac: Event-driven rbac. In: Proceedings of the 18th ACM Symposium on Access Control Models and Technologies. SACMAT '13, NY, USA, ACM (2013)
31. Ahmed, A., Zhang, N.: A context-risk-aware access control model for ubiquitous environments. In: IMCSIT, IEEE (2008)
32. Chen, L., Crampton, J.: Risk-aware role-based access control. In Meadows, C., Fernandez-Gago, C., eds.: Security and Trust Management. Volume 7170 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2012) 140–156
33. Chen, L., Crampton, J., Kollingbaum, M.J., Norman, T.J.: Obligations in risk-aware access control. In Cuppens-Boulahia, N., Fong, P., García-Alfaro, J., Marsh, S., Steghöfer, J., eds.: PST, IEEE (2012) 145–152
34. Cheng, P.C., Rohatgi, P., Keser, C., Karger, P.A., Wagner, G.M., Reninger, A.S.: Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In: IEEE Symposium on Security and Privacy, IEEE Computer Society (2007) 222–230
35. Dickens, L., Russo, A., Cheng, P.C., Lobo, J.: Towards learning risk estimation functions for access control. In: In Snowbird Learning Workshop. (2010)
36. Shaikh, R.A., Adi, K., Logrippo, L.: Dynamic risk-based decision methods for access control systems. Volume 31. (2012) 447–464
37. Bettini, C., Jajodia, S., Wang, X.S., Wijesekera, D.: Provisions and obligations in policy management and security applications. In: Proceedings of the 28th International Conference on Very Large Data Bases. VLDB '02, VLDB Endowment (2002) 502–513
38. Baracaldo, N., Joshi, J.: Beyond accountability: Using obligations to reduce risk exposure and deter insider attacks. In: Proceedings of the 18th ACM Symposium on Access Control Models and Technologies. SACMAT '13, New York, NY, USA, ACM (2013) 213–224
39. Dimmock, N., Belokosztolszki, A., Eyers, D., Bacon, J., Moody, K.: Using trust and risk in role-based access control policies. In: Proceedings of the Ninth ACM Symposium on Access Control Models and Technologies. SACMAT '04, New York, NY, USA, ACM (2004) 156–162
40. Shah, A., Dahake, S., J., S.H.H.: Valuing data security and privacy using cyber insurance. SIGCAS Comput. Soc. **45**(1) (February 2015) 38–41
41. Kelley, P., Komanduri, S., Mazurek, M., Shay, R., Vidas, T., Bauer, L., Christin, N., Cranor, L., Lopez, J.: Guess again (and again and again): Measuring password strength by simulating password-cracking algorithms. In: Security and Privacy (SP), 2012 IEEE Symposium on. (May 2012) 523–537