Risk-Based Privacy-Aware Information Disclosure

Alessandro Armando
DIBRIS, University of Genova, Italy
Security & Trust Unit, FBK-Irst, Trento, Italy

Michele Bezzi
Product Security Research, SAP Labs,
Sophia-Antipolis, France

Nadia Metoui
Security & Trust Unit, FBK-Irst, Trento, Italy
DISI, University of Trento, Italy

Antonino Sabetta
Product Security Research, SAP Labs,
Sophia-Antipolis, France

Risk-aware access control systems grant or deny access to resources based on the notion of risk. It has many advantages compared to *classical* approaches, allowing for more flexibility, and ultimately supporting for a better exploitation of data. We propose and demonstrate a risk-aware access control framework for information disclosure, which supports run-time risk assessment. In our framework access-control decisions are based on the disclosure-risk associated with a data access request and, differently from existing models, adaptive anonymization operations are used as risk-mitigation method. The inclusion of on-the-fly anonymization allows for extending access to data, still preserving privacy below the maximum tolerable risk. Risk thresholds can be adapted to the trustworthiness of the requester role, so a single access control framework can support multiple data access use cases, ranging from sharing data among a restricted (highly trusted) group to public release (low trust value). We have developed a prototype implementation of our framework and we have assessed it by running a number of queries against the Adult Data Set from the UCI Machine Learning Repository, a publicly available dataset that is widely used by the research community. The experimental results are encouraging and confirm the feasibility of the proposed approach.

1 Introduction

The increase in the amount of data generated by today's digital society is astonishing. According to IDC estimate [11], the global volume of digital data will double every two years, reaching 44 trillion gigabytes by 2020. Potentially organizations are now in the position to fully exploit these huge amount of diverse datasets to create new data-based businesses as well as optimizing existing process (e.g., real-time customization). On the other hand, often, organizations are not fully leveraging this potential due to the lack of appropriate dissemination and control mechanisms, which supports risk-based decision making, balancing the advantages of information access with the security. Personal information is particularly critical, since they are subject to strict regulations, and enterprises will have to comply with it to avoid significant fines and impact on reputation. As a result, most organizations strongly limit (even internally) the sharing and dissemination of data making most of the information unavailable to decision-makers and therefore do not exploit the power of existing data sources.

Already a few years ago, the JASON report [19] indicated that the inflexibility of existing access control mechanisms is a major obstacle with dealing with diverse data sources in dynamic environments. To address this issue, access control mechanisms based on risk estimation (i.e., risk-aware access control) have been put forward [6]. In a nutshell, in risk-aware access control access control decisions are based on an estimation of expected cost and benefits and only not (as in traditional access control systems) on a policy statically defining stating which requests should be allowed and which should be denied. In a risk-aware access control system, for each access request, the corresponding risk is estimated and compared with a risk-threshold. If the risk is less than a given risk threshold, then access is granted, otherwise it is denied. This allows for a better exploitation of the data than in traditional access control system while controlling risk. Although existing risk-aware access control models enjoy many advantages and allow for a better management and exploitation of the data, they suffer from a number of drawbacks which limit its effectiveness. For instance, most existing risk-based access control models only support binary access decision (i.e., the outcome is either *allowed* or *denied*), whereas in real-life we often have exceptions based on additional conditions (e.g., *I cannot disclose these data, because they contain personal identifiable information, but I can disclose an anonymized version of the data*). In other words, the system should enforce appropriate risk mitigation measures, and relevant part of additional information could be shared. From a risk perspective, such mitigation measures have the effect of decreasing the risk associated with the release of the data.

Anonymization is a commonly used practice to reduce privacy risk, obfuscating, in part or completely, the personal identifiable information in a dataset. Anonymization methods include [7]: suppressing part of or entire records; generalizing the data, i.e., recoding variables into broader classes (e.g., releasing only the first two digits of the zip code) or rounding/clustering numerical data; replacing identifiers with random values (e.g., replacing a real name with a randomly chosen one). To quantify the level of anonymity, several metrics have been proposed in the literature (see [3, 8] for a review). These metrics differ in a number of ways, but they all express the risk of disclosing personal-identifiable information when releasing a given dataset. Anonymization increases protection, by lowering the privacy risk, and enables a wider exploitation of the data, but it clearly impacts the utility of the data. Accordingly, different level of anonymization should be considered depending on a number of factors, often known at run-time only, such as the trustworthiness of the requester or security context of the query.

In this paper, we propose and demonstrate a risk-aware access control framework for information disclosure, which addresses the concerns described above. In our framework access-control decisions are based on the disclosure-risk associated with a data access request and, differently from existing models, we include adaptive anonymization operations as risk-mitigation methods. The inclusion of on-the-fly anonymization allows for extending access to the data, still preserving privacy below the maximum tolerable risk. Risk thresholds can be adapted to the trustworthiness of the requester role, so a single access control framework can support multiple data access use cases, ranging from sharing data among a restricted (highly trusted) group to public release (low trust value). To evaluate the effectiveness of the proposed approach we have developed a prototype implementation of our risk-aware access control framework and we have assessed it by running a number of queries against the Adult Data Set from the UCI Machine Learning Repository, a publicly available dataset that is widely used by the research community. The experimental results are encouraging and confirm the feasibility of our proposed approach.

Structure of the paper. In the next Section we provide a simple but realistic scenario that illustrates the main features of risk-aware information disclosure. We then recall some background notions on risk-aware access control (Section 3) and privacy preserving information disclosure (Section 4). In Section 5 we present our access control model for risk-aware information disclosure and in Section 6 we illustrate its application on the scenario introduced in Section 2. In Section 7 we present an architecture for Risk-Aware Access Control Framework. In Section 8 we discuss an experimental evaluation of the proposed approach. We discuss the related work in Section 9 and we conclude in Section 10 with some final remarks.

2 Scenario

Employee surveys are a widely used instrument for organizations to assess job satisfaction, quality of management, people motivation, etc. Considering the possible sensitivity of data, surveys should be anonymous, meaning that the organization and management should not be able to identify how a specific employee responded. Usually, the organization—say, a large company—conducting the survey outsources the data collection to a third-party. When processing the data, the third-party has access to individual-level information, whereas the same data is not accessible to the company. To protect the anonymity of the survey, the company can access the data under the condition that *(i)* identifiers are removed and *(ii)* the number of respondents is larger than a certain threshold (usually between 10 and 25). Different splits of data can be requested (e.g., per organization, per job profile, etc.), but data are accessible only if the query results contains a number of respondents that is larger than the fixed threshold. On top of that, additional access control rule can be enforced, e.g., a manager would only see data referring to his/her team or department (provided that conditions *(i)* and *(ii)* are also fulfilled); an employee would be allowed to see overall (company results) only. As an example, consider a question like "Do you respect your manager as a competent professional?" with a five points scale (1 to 5). A manager could see the response of his/her team if at least, say, 10 people answered to it. If the manager decides to refine the analysis asking for data related to the people in his/her team AND with a "developer" role, again the response should be made available only if at least 10 respondents with that role answered to the question.[1] Current systems typically do not provide any data if the number of respondents is below the defined thresholds (for the specific role). In other words, in order to avoid the risk of disclosing too much information, an overly conservative approach is taken and risky queries are not permitted altogether. Ideally, the access control system should be able to provide the largest possible amount of information (still preserving anonymity) for any query. In practice, in presence of queries that might cause anonymity issues (i.e., not enough respondents, or more generally, too small a result set), the system should be able to quantify the disclosure risk associated with the query and compare it with whatever risk level has been set as the acceptable threshold. If the threshold is exceeded, the system could apply, for example, a "generalization" operation (making the query less specific), thus increasing the cardinality of the result set and reducing the risk of disclosing the identity of respondents. Of course, applying such an operation would not yield the *exact*

[1] In real surveys single records are actually never shown, but just percentages, in this example it would be something like 10% answered 1, 25% answered 2, etc. Since the number of respondents is known, in practice, for one question, this equivalent of getting the data with no identifiers.

data set the user asked for, but this method would: 1) provide some relevant (i.e., as close as possible to the original query) information to the user, and 2) preserve anonymity according to some pre-defined disclosure-risk levels (possibly linked to the requestor trust or role).

In the next section,we discuss how to implement such a system using risk-based access control, and anonymization mitigation strategies.

3 Risk-Aware Access Control

We provide a brief presentation of the formal model for Risk-Aware Access Control (RAAC) that has been introduced in [5]. We use this model as the basis of our access control model for risk-aware information disclosure that will be presented later.

The RAAC model consists of the following components:

- a set of users U;
- a set of permissions P, usually representing action-object pairs;
- a set of access requests Q, modeled as pairs of the form (u, p) for $u \in U$ and $p \in P$;
- a set of *risk mitigation methods* $\mathcal{M}$, i.e., actions that are required to be executed to mitigate risk;
- a function π mapping permissions into *risk mitigation strategies*, i.e., lists of the form $[(l_0, M_0), (l_1, M_1), \ldots, (l_{n-1}, M_{n-1}), (l_n, M_n)]$, where $0 = l_0 < l_1 < \cdots < l_{n-1} < l_n \leq 1$ and $M_i \subseteq \mathcal{M}$ for $i = 0, \ldots, n$;
- a set of *states* Σ, i.e., tuples of the form (U, P, π, τ) where τ abstracts further specific features of the state; for instance, in the Risk-Aware Role-Based Access Control (R^{2}BAC) model [4], τ comprises the set of roles R, the user-role assignment relation $UA \subseteq U \times R$, the role-permission assignment relation $PA \subseteq P \times R$, the role hierarchy $\geq \subseteq R \times R$, and the user trustworthiness $\alpha : U \to (0..1]$, the user-role competence function $\beta : U \times R \to (0..1]$, and the role-permission appropriateness function $\gamma : R \times P \to (0..1]$;
- a *risk function risk* $: Q \times \Sigma \to [0..1]$ such that $risk(q, \sigma)$ denotes the risk associated to granting q in state σ;
- an *authorization decision function Auth* $: Q \times \Sigma \to D \times 2^{\mathcal{M}}$ with $D = \{\mathsf{allow}, \mathsf{deny}\}$ such that if $q = (u, p)$ and $\pi(p) = [(l_0, M_0), \ldots, (l_n, M_n)]$, and σ the current state, then

$$Auth(q, \sigma) = \begin{cases} (d_i, M_i) & \text{if } risk(q, \sigma) \in [l_i, l_{i+1}), i < n, \\ (d_n, M_n) & \text{otherwise} \end{cases}$$

where $d_i \in D$. Intuitively, if the risk associated with access request (u, p) is l, then $Auth$ returns an authorization decision and a set of risk mitigation methods corresponding to the interval containing l.

4 Privacy Preserving Information Disclosure

We assume that the data is represented as a relational table, called *private table*. Each record in the table is relative to a specific *respondent*. The attributes (columns) in the table can be classified as follows:

- *Identifiers*. These are data attributes that can uniquely identify individuals. Examples of identifiers are the Social Security Number, the passport number, the complete name.
- *Quasi-identifiers (QIs) or key attributes* [9]. These are the attributes that, when combined, can be used to identify an individual. Examples of QIs are the postal code, age, job function, gender, etc.
- *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).

Various anonymity metrics have been proposed in the literature (see [3, 8] for a review), the most popular being k-anonymity [20], ℓ-diversity [15], and t-closeness [13]. 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 identity privacy risk, because, at best, an attacker has a probability $1/k$ to re-identify a record (i.e., associate the sensitive attribute of a record to the identity of a respondent). Consider now a table with a group of k records sharing the same combination of quasi-identifiers have the same sensitive attribute. Even if the attacker is unable to re-identify the record, he can discover the sensitive information (attribute disclosure). The ℓ-diversity metrics was introduced to capture this type of risk. It requires that for *every* combination of key attributes there should be at least ℓ values for each confidential attribute. Although, the ℓ-diversity condition prevents the attacker from inferring exactly the sensitive attributes, he may still learn a considerable amount of probabilistic information: if the distribution of confidential attributes within a group sharing the same key attributes is very dissimilar from the distribution over the whole set, an attacker may increase his knowledge on sensitive attributes (*skewness attack*, see [13] for details). To overcome the problem, t-closeness estimates this risk by computing the distance between the distribution of confidential attributes within the group and in the entire dataset. These measures provide a quantitative assessment of the different risks associated to data release, and each of them (or a combination thereof) can be applied to estimate privacy risk depending on the use case at hand.

In this paper we will use k-anonymity as anonymity metrics to present our ideas, but it must be emphasized that the approach can readily adapted to use alternative metrics (including ℓ-diversity and t-closeness).

5 Risk-Aware Information Disclosure

We now refine the RAAC model of Section 3 into our model for Risk-Aware Information Disclosure. Let P be a set of database views (or virtual tables). If p is a view, then $|p|$ denotes the anonymity of p according to some given metrics (e.g. k-anonymity). The higher is the value of $|p|$, the smaller is the risk to disclose sensitive information by releasing p. Thus, for instance, we can define *the (privacy) risk of disclosing p* to be $1/|p|$ and *the (privacy) risk of disclosing p to u in* $\sigma = (U, P, \pi, \tau)$ to be

$$risk((u, p), \sigma) = \begin{cases} 1 & \text{if not } granted_\tau(u, p) \\ 1/|p| & \text{otherwise} \end{cases}$$

where $granted_\tau(u, p)$ holds if and only if u is granted access to p according to τ. For instance, if τ is an RBAC policy $(U, R, P, UA, RA, \geq)$, then $granted_\tau(u, p)$ holds if and only if there exist $r, r' \in R$ such that $(u, r) \in UA$, $r \geq r'$, and $(p, r') \in PA$.

When the risk associated to the disclosure of a certain view p is greater than the maximal accepted risk t, we can use obligations for obfuscating or redacting the view and thus bring the risk below t. In this paper we consider k-anonymization functions $\phi_k : P \to P$ for $k \in \mathbb{N}$ as risk mitigation methods, but functions based on other metrics can be used as well. Clearly $|\phi_k(p)| \geq k$ for all $p \in P$. We then consider risk mitigation strategies of the form $\pi(p) = [(0, \iota), (t, \phi_{\lceil 1/t \rceil}(.))]$, where $\iota : P \to P$ is the identity function (i.e. such that $\iota(p) = p$ for all $p \in P$) and the following authorization decision function:

$$Auth((u, p), \pi) = \begin{cases} (\textsf{deny}) & \text{if not } granted_\tau(u, p) \\ (\textsf{allow}, \phi_{\lceil 1/t \rceil}(\cdot)) & \text{if } risk(u, p) \geq t \\ (\textsf{allow}, \iota) & \text{if } risk(u, p) < t \end{cases} \tag{1}$$

that always grants access but yields an anonymized version of the requested view if the risk is greater that the maximal accepted risk t. In other words, if user u asks to access p, then access to p is granted unconditionally if $risk(u, p) < t$, otherwise an anonymized version of p, say $\phi_{\lceil 1/t \rceil}(p)$, is computed and returned to u.

Example 5.1. To illustrate assume Alice asks for a view p_1 such that $|p_1| = 4$ and that $\pi(p_1) = [(0, \iota), (t, \phi_{\lceil 1/t \rceil}(.))]$ with $t = 0.1$, i.e. $\pi(p_1) = [(0, \iota), (0.1, \phi_{10}(.))]$. It is easy to see that $risk(Alice, p_1) = 0.25$ and that $Auth((Alice, p_1), \pi) = \phi_{10}(p_1)$. Alice then asks for a view p_2 such that $|p_2| = 20$ and that $\pi(p_2) = \pi(p_1) = [(0, \iota), (t, \phi_{\lceil 1/t \rceil}(.))]$ with $t = 0.1$, i.e. $\pi(p_2) = [(0, \iota), (0.1, \phi_{10}(.))]$. It is easy to see that now $risk(Alice, p_2) = 0.05$ and therefore that $Auth((Alice, p_2), \pi) = \iota(p_1) = p_1$.

The following results state that the risk of disclosing the view returned by our authorization decision function is never greater than the maximum accepted risk.

Proposition 5.1. Let $(D, M) = Auth((u, p), \pi)$. Then $risk(u, M(p)) \leq t$.

In many situations of practical interest, we want the risk of a query $q = (u, p)$ to depend also on the trustworthiness of the user u. This can be done by (re)defining the risk function as follows:

$$risk((u, p), \sigma) = \begin{cases} 1 & \text{if not } granted_\tau(u, p) \\ \max\{0, \frac{1}{|p|} - \alpha(u)\} & \text{otherwise} \end{cases} \tag{2}$$

where $\alpha : U \to (0..1]$ is a function that assigns a trust value to users.

When roles correspond to job functions, it is natural to assign trust to roles and to derive the trust of a user from the trust assigned to the roles assigned to that user in the following way:

$$\alpha(u) = \max\{\alpha(r') : (p, r') \in PA \text{ and } \exists r \geq r' \text{ s.t. } (u, r) \in UA\}.$$

6 Application of Risk-Aware Role-Based Access Control

We now show how our risk-aware information disclosure model can be used to support the scenario of Section 2. This will be done by setting appropriate values to the parameters occurring in the definition of the risk function (2).

For sake of simplicity we consider a small company, with 8 employees and one manager. The company runs an employee survey, with one single question with answer ranging in a five points scale (from 1 to 5) (*sensitive attribute*, cf. Section 4), and collecting user names[2] (the *identifiers*), as well as the job title and the location of the office (the *quasi-identifiers*). The actual dataset is in Table 1(a). To preserve privacy we set the maximal acceptable risk to $t = 0.125$.

The outsourcing company collecting the data is considered fully trusted and will therefore have access to all the information. We model this by setting the trust of the `admin` role to 1, i.e. $\alpha(\texttt{admin}) = 1$. Thus, an administrator can access the original dataset, say p_{all} with anonymity $|p_{all}| = 1$ (i.e., all distinct values, see Table 1(a)), since $\alpha(\texttt{admin}) = 1$ and the risk value is smaller than the threshold, i.e., $1 - 1 = 0 < 0.125$. If we set the trust value of the `manager` role to

[2] In real cases they are typically user IDs

Table 1
The Employee Survey Example

(a) Original dataset

<table>
<tr><th colspan="4">Survey Administrator view
$|p_{all}| = 1$</th></tr>
<tr><th>Name</th><th>Job</th><th>Location</th><th><b>Answer</b></th></tr>
<tr><td>Timothy</td><td>SeniorDeveloper</td><td>Houston</td><td>4</td></tr>
<tr><td>Alice</td><td>Support</td><td>Houston</td><td>5</td></tr>
<tr><td>Perry</td><td>JuniorDeveloper</td><td>Rome</td><td>5</td></tr>
<tr><td>Tom</td><td>Admin</td><td>Rome</td><td>3</td></tr>
<tr><td>Ron</td><td>SeniorDeveloper</td><td>London</td><td>4</td></tr>
<tr><td>Omer</td><td>JuniorDeveloper</td><td>London</td><td>4</td></tr>
<tr><td>Bob</td><td>Support</td><td>Houston</td><td>5</td></tr>
<tr><td>Amber</td><td>Admin</td><td>Houston</td><td>3</td></tr>
</table>

(b) Anonymized version: *identifiers* and *quasi-identifiers* are suppressed

<table>
<tr><th colspan="4">Employee View
$|p_{supp}| = 8$</th></tr>
<tr><th>Name</th><th>Job</th><th>Location</th><th><b>Answer</b></th></tr>
<tr><td>***</td><td>***</td><td>***</td><td>4</td></tr>
<tr><td>***</td><td>***</td><td>***</td><td>5</td></tr>
<tr><td>***</td><td>***</td><td>***</td><td>5</td></tr>
<tr><td>***</td><td>***</td><td>***</td><td>3</td></tr>
<tr><td>***</td><td>***</td><td>***</td><td>4</td></tr>
<tr><td>***</td><td>***</td><td>***</td><td>4</td></tr>
<tr><td>***</td><td>***</td><td>***</td><td>5</td></tr>
<tr><td>***</td><td>***</td><td>***</td><td>3</td></tr>
</table>

Table 2
Views of the employee survey for the Rome location

(a) Before generalization.

<table>
<tr><th colspan="4">View: Location=Rome, $|p_{Rome}| = 2$</th></tr>
<tr><th>Name</th><th>Job</th><th>Location</th><th><b>Answer</b></th></tr>
<tr><td>***</td><td>***</td><td>Rome</td><td>5</td></tr>
<tr><td>***</td><td>***</td><td>Rome</td><td>3</td></tr>
</table>

(b) After generalization

<table>
<tr><th colspan="4">View: Location=Rome
Anoymized $|p_{EMEA}| = 4$</th></tr>
<tr><th>Name</th><th>Job</th><th>Location</th><th><b>Answer</b></th></tr>
<tr><td>***</td><td>***</td><td>EMEA</td><td>5</td></tr>
<tr><td>***</td><td>***</td><td>EMEA</td><td>3</td></tr>
<tr><td>***</td><td>***</td><td>EMEA</td><td>4</td></tr>
<tr><td>***</td><td>***</td><td>EMEA</td><td>4</td></tr>
</table>

Table 3
Views of the employee survey for Rome and JuniorDeveloper

(a) Before generalization of location and job

<table>
<tr><th colspan="4">Loc=Rome AND Job=JuniorDeveloper
$|p_{Rome+JuniorDeveloper}| = 1$</th></tr>
<tr><th>Name</th><th>Job</th><th>Location</th><th><b>Answer</b></th></tr>
<tr><td>***</td><td>JuniorDeveloper</td><td>Rome</td><td>5</td></tr>
</table>

(b) After generalization of location and job

<table>
<tr><th colspan="4">View Loc=Rome AND Job=JuniorDeveloper
Anonymized $|p_{EMEA+Dev}| = 3$</th></tr>
<tr><th>Name</th><th>Job</th><th>Location</th><th><b>Answer</b></th></tr>
<tr><td>***</td><td>Dev</td><td>EMEA</td><td>5</td></tr>
<tr><td>***</td><td>Dev</td><td>EMEA</td><td>4</td></tr>
<tr><td>***</td><td>Dev</td><td>EMEA</td><td>4</td></tr>
</table>

0.21, i.e. $\alpha(\texttt{manager}) = 0.21$ (corresponding to access views with anonymity $k \geq 3$), than a manager cannot access p_{all} as is, since $1 - 0.21 > 0.125$ and some anonymization, as risk mitigation strategy, must be carried out on the data to decrease the risk. For example, if we suppress the identifier attribute (*Name*) and the quasi-identifiers (*Job* and *Location*), we obtain the view p_{supp} shown in Table 1(b). The view p_{supp} corresponds to an anonymity level $|p_{supp}| = 8$ and since $0.125 - 0.21 < 0.125$, access is granted to the manager.[3] The manager can also ask for more granular views of the results. For example, if she wants to know the distribution of the answers in one location, say Houston, $|p_{Houst}| = 4$, the risk $0.25 - 0.21 = 0.04$ is still smaller than $t = 0.125$. On the other hand, if she asks for the result in Rome, $|p_{Rome}| = 2$, then the risk associated with the view for the manager is $0.5 - 0.21 > 0.125$ and the access is granted only if appropriate anonymization is performed. In this case, location could be generalized from Rome to EMEA (so including London workforce), as shown in Table 2(b). The resulting view has anonimity $|p_{EMEA}| = 4$ and since the risk is smaller than $t = 0.125$, then the manager is allowed to see the view.

Similarly, if the manager wants to see the results per location and per job function (say in Rome for JuniorDeveloper only, see Table 3(a)), the anonymity level is low, $|p_{Rome+JuniorDeveloper}| = 1$, and the associated risk is greater than $t = 0.125$. Again, instead of simply denying access, the system can perform generalization on both the quasi-identifiers, *Job* (generalized to the job family developer) and *Location*, thereby increasing the anonymity ($|p_{EMEA+Dev}| = 3$) and decreasing the risk ($risk(manager, p_{EMEA+Dev}) = 0.123$) to an acceptable level for a manager (see Table 3(b)).

Finally, employees should have access to the global results only. The trust value is therefore set to $\alpha(\texttt{employee}) = 0.125$ and the only view permitted is with suppression of all identifiers and quasi-identifiers, which has $|p_{supp}| = 8$, see Table 3(b).

[3] In real surveys the result will appear as a report like: 37.5% answered 5, 37.5% answered 4 and 25% answered 3. For a single question this is equivalent to the view in Table 1(b).

7 Risk-Aware Access Control Framework

This section presents an abstract architecture for our Risk-Aware Access Control Framework. The architecture, depicted in Figure1, is composed of three main modules whose role is described in the following.

Risk-Aware Access Control module. This module is the entry point to our system, through which users can submit requests to retrieve data from the underlying database. The module evaluates the access authorisations of the data requestor and grants or denies access. if the access is denied by the access policy the request is rejected if it is granted the Risk-Aware Access Control module will call the Risk Estimation Module to determine the risk level of the query and the Risk Mitigation Module to reduce risk as 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 (such as, the user's role and trustworthiness, and acceptable risk threshold) that are needed to determine the risk level. These additional attributes are passed to the Risk Estimation module to compute the risk associated with a particular query for a given user.

Risk Estimation module. The Risk Estimation module receives the user attributes and determines the level of disclosure risk, based on the data that is requested and on the criteria defined in the risk estimator configuration. This configuration includes the metrics used to estimate disclosure risk with respect to domain-specific knowledge about what information is to be considered critical or not.

Besides the evaluation of the risk, the Risk Estimation module produces an estimation of the minimal anonymization level to be applied in order to meet the risk threshold (i.e., in case of k-anonymity, the risk estimation module computes the minimal value of k that respects the risk threshold constraint).

Risk Mitigation module. The Risk Mitigation module is activated by the Risk-based Access Control module when the disclosure risk exceeds the acceptable risk threshold for the requested resource. In such a case, the Risk Mitigation Module applies the optimal anonymization operation (e.g., generalisation, suppression) that is needed to reduce the disclosure risk down to an acceptable level (that is, a level that is equal or less than the threshold) while minimising information loss.

8 Evaluation

This section documents the results of an initial evaluation of our approach. The two questions we investigate are (A) whether the approach described in this paper can be realized in practice and (B) whether the performance that can be expected under typical workloads matches the needs of real-time (more precisely: online) operation.

In order to address question A, we realized a prototype system that we have used to run sample scenarios. We use the same prototype also to study the response time under several representative conditions (queries of varying complexity, different levels of user trust and therefore, different loads for the anonymizer module).

In the following, we first describe our prototype implementation, then we present the dataset we used for the evaluation and outline the results of the experiments we run on that dataset.

8.1 Prototype Implementation

In order to evaluate the practical feasibility of our approach, we developed a proof-of-concept implementation of our framework (see Section 7) that we used to run the experiments described in the following.

Our prototype is implemented in Java 7 and uses MySQL Server version 5.6.20 to store the dataset. The Risk Aware Access Control module mimics a typical XACML data flow, providing a basic implementation of the PDP, the PEP, and the PIP functionality as well as a set of authorization policies. The Risk Mitigation module is implemented using the ARX[4] anonymization framework [12]. The ARX toolkit offers a Java API supporting data de-identification. ARX is capable of altering input data in a way that guarantees minimal information loss while ensuring that the transformed data adheres to well-defined privacy criteria, expressed in such metrics as k-anonymity, ℓ-diversity, t-closeness, etc. ARX also offers several reporting features allowing to collect metrics such as execution time, information loss, etc. We evaluated other available anonymization libraries (e.g., Cornell Anonymization Toolkit[5], University of Texas Anonymisation Toolbox[6]). We eventually adopted ARX because we found it easy to integrate and considering that it is a well-documented, actively developed, and well maintained project.

8.2 Dataset

To test the performance of our framework, we used a dataset that is widely used in the research community, namely the Adult Data Set [7] from the UCI Machine Learning

[4] http://arx.deidentifier.org/overview/

[5] http://anony-toolkit.sourceforge.net/

[6] http://cs.utdallas.edu/dspl/cgi-bin/toolbox/index.php

[7] Available at http://archive.ics.uci.edu/ml/datasets/Adult

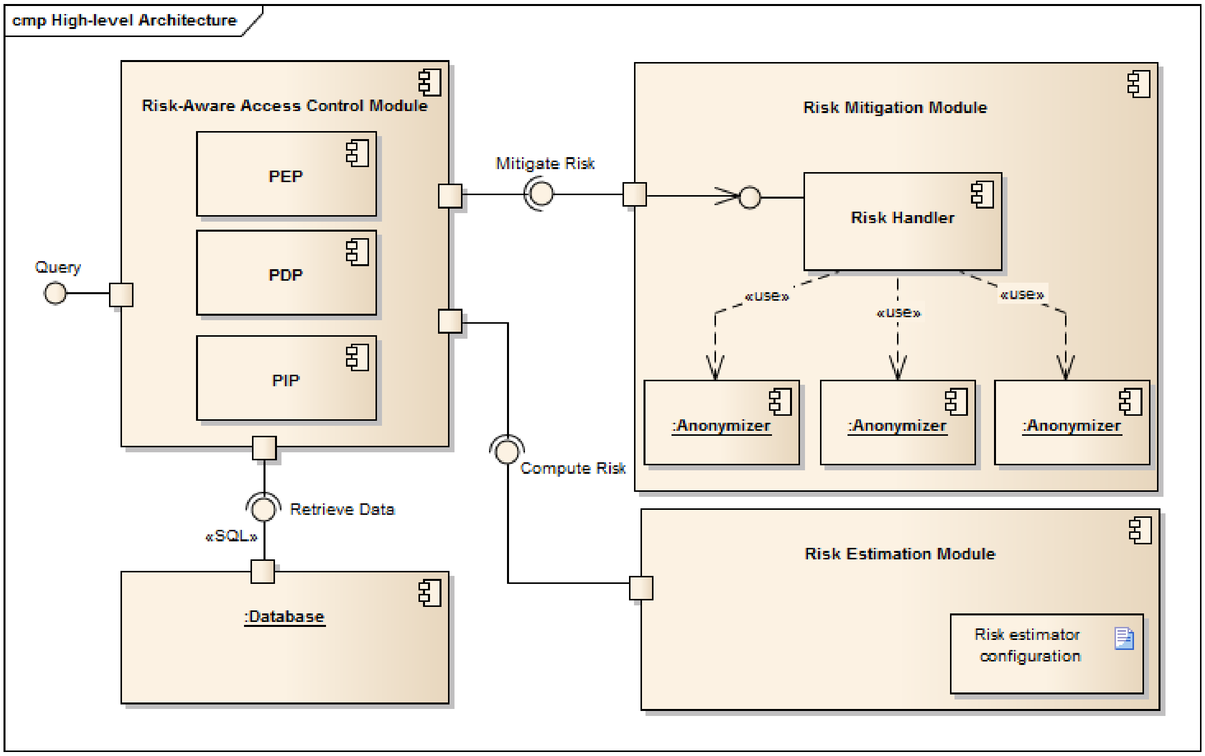

Figure 1. Architecture of the Risk-Aware Access Control framework

Repository. This dataset contains 32561 records from the US Census dataset with 15 demographic and employment-related variables. We removed records with missing values, ending with 30, 162 usable records, and we reduced the number of fields to nine, as shown in Table 4.

The choice of the identifiers, QIs and sensitive attribute set, typically, depends on the specific domain. QIs should include the attributes a possible attacker is likely to have access to (e.g., using a phonebook or a census database), whereas sensitive attributes depend on the application the anonymized data are used for.

Generally speaking increasing the number of QIs increases the risk, or results in strong anonymization impacting the usefulness of the resulting view. In our experiments we set $QI \equiv \{\text{AGE}, \text{NATIVE-COUNTRY}\}$. In the census data, the SALARY-CLASS attribute is typically chosen as a sensitive attribute. We also classified RACE as a sensitive attribute because of its discriminatory nature.

QIs will be generalized according to the generalisation scheme of Figure 2 (for the attribute AGE) and Figure 3 (for the attribute NATIVE-COUNTRY).

8.3 Experiment and Results

In order to evaluate the performance of our tool, including the computational overhead caused by the anonymization engine, we used a number of queries of increasing complexity in terms of the size of the returned views and the disclosure risk. The queries are given in Table 5 and the corresponding size and anonymity level of the views returned by our tool are reported in Table 6. In the following we will indicate both the queries and the corresponding views as **Q1**, **Q2**, **Q3**, **Q4**.

Table 4
Summary of the dataset columns, number of distinct values, and nature of each column

UCI Adult Dataset		
Attribute	Values	Nature
AGE	72	QI
NATIVE-COUNTRY	41	QI
EDUCATION	16	not Sensitive
OCCUPATION	14	not Sensitive
WORKCLASS	7	not Sensitive
MARITAL-STATUS	7	not Sensitive
GENDER	2	not Sensitive
RACE	5	Sensitive
SALARY-CLASS	2	Sensitive

For our experiments, we want to investigate the impact of risk mitigation, anonymization, on *(i)* the performance of the access control system and *(ii)* the quality of the resulting data. For case *(i)* we focus on the views with the largest sizes (namely, **Q1** and **Q2**, with more than 20,000 tuples each as shown in Table 6). For case *(ii)* we focus on the views with the highest risk profiles (namely, **Q1**, **Q3**, and

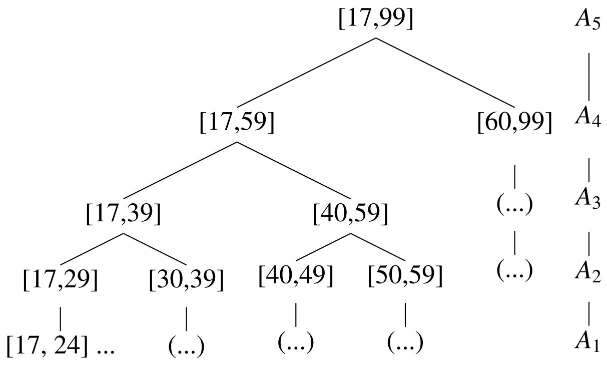

Figure 2. Generalization hierarchy for the attribute AGE [17, 99]. Level A_1: Age is generalized in 5 year range. Level A_2 in 10 year range. Level A_3 in 20 years. Level A_4 in 40 year range. In level A_5 the age is fully generalized. Age is not generalized in level A_0 (not shown).

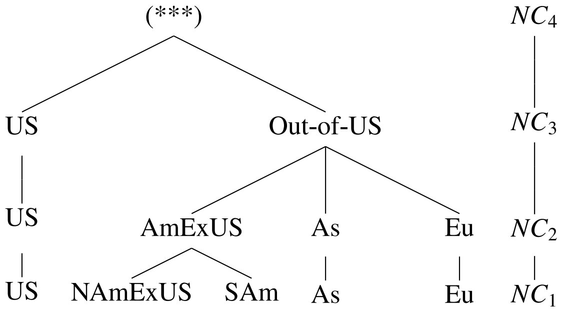

Figure 3. Generalization hierarchy for the attribute NATIVE-COUNTRY: Level NC_1: NATIVE-COUNTRY is generalized to US (United States), AmExUS (America Excluding United States), Asia (As), or Europe (Eu). Level NC_2: NAmExUS (North America Excluding United States) and SAm are generalized to AmExUS (America Excluding United States). Level NC_3: All countries excluding United States are generalized to Out-of-US. Level NC_4: native countries are suppressed. Level NC_0: native countries are not generalized (not shown).

Table 5
Queries

Q1: Data about male respondents

```
SELECT * FROM ADULT
WHERE SEX = ‘‘Male’’;
```

Q2: Data about adults between 30 and 75 years old born in the United States

```
SELECT * FROM ADULT
WHERE AGE BETWEEN 30 AND 75
AND NATIVE-COUNTRY = ‘‘United-States’’;
```

Q3: Data about adults between 30 and 35 years old working in the private sector and originally from the american continent excluding United States

```
SELECT * FROM ADULT
WHERE WORKCLASS = ‘‘Private’’
AND AGE BETWEEN 30 AND 35
AND NATIVE-COUNTRY IN
(<America Excluding the United-States>);
```

Q4: Data about adults without-pay

```
SELECT * FROM ADULT
WHERE WORKCLASS = ‘‘Without-pay’’
```

Table 6
Size and disclosure risk level of the views returned in response to the queries

Query	Size	Anonimity	Risk level
Q1	20,380	1	High
Q2	19,392	32	Low
Q3	215	1	High
Q4	14	1	High

Q4, with the lowest possible anonymity), whose computation is significantly affected by anonymization. We consider five risk thresholds α i.e. users/role with different trustworthiness level, as shown in Table 7, and each experiment is run 100 times to average out the variance of the response time. In Figure 4 we report the results of the experiments for the four queries, panels **Q1**, **Q2**, **Q3**, and **Q4**, respectively, for the five different trustworthiness levels. Figure 5 shows the (possible) impact of generalization on the data accuracy, as measured by the Precision metric (Prec) [22], which counts the average number of generalization steps performed on the generalization trees (cf. Figure 2 and Figure 3).

For **Q1**, we observe that the anonymization process increases significantly the response time. Indeed the query is carried our by the most trusted user ($\alpha = 1$), with no anonymization needed, takes on average 8ms (see Figure 4.Q1, horizontally striped bar corresponding to $\alpha = 1$). By decreasing the trustworthiness of the requester the view must be anonymized and the average response time increases to 27ms (cf. Figure 4.Q1, horizontally striped bar corresponding to $\alpha = 0.52$). This time difference is entirely due to the anonymization time (19 ms, as shown in Figure 4.Q1, diagonally striped bar corresponding to $\alpha = 0.52$). Decreasing further the trust level results in additional anonymization. Also the attribute NATIVE-COUNTRY (NC) gets anonymized (cf. Figure 5.Q1), but this does not significantly affect the response time (see Figure 4.Q1).

We can observe a similar behavior in the other queries (see Figure 4.Q2, Q3, and Q4), with an increase of response time when anoymization takes place and no significant variations

Table 7
User roles and trustworthiness

User Name	Role	Trustworthiness
Alice	SuperUser	1
Megha	Admin	0.52
Dana	SeniorDataAnalyst	0.1
Frida	JuniorDataAnalyst	0.028
Eliyes	IT	0.015

in performance for different levels of anonymization. For instance, for **Q2** we have a view with an already high level of anonymity ($k = 32$), and a *small* anonymization (a single level of generalization for the Age attribute, see Figure 5.Q2 for $\alpha = 0.015$) still significantly impacts the performance. In case of **Q3** we see that, despite different combinations of anonymization strategies for different values of α (Figure 5.Q3), the response time is not affected (Figure 4.Q3), except for $\alpha = 1$ where we have no anonymization. We should note that for **Q3** (as well as **Q4**) the difference in the average response time with and without anonymization is relevant ($\alpha = 1$ has response time of 0.16ms, and $\alpha = 0.52$ of 1.6 ms) but these views have few tuples and these times are small in absolute value, with large fluctuations, as shown by the high standard deviations.

Q4 is characterized by a low cardinality and (consequently) by high anonymity. Except for the maximum trust value the data are strongly anonymized and for low trust levels $\alpha = 0.28$ and $\alpha = 0.015$ access is denied in spite of the anonymization, see Figure 5.Q4. Note that in these cases, the anonymization engine tries to minimize the risk (anonymization time is not zero, see Figure 4.Q4), but due to the low cardinality no solution is found.

From these experiments, we observe that when anonymization is applied the response time increases significantly, but, even in the worst cases, the increase is far less than one order of magnitude with no impact on the real-time response of the system. Moreover, the application of different anonymization strategies have no impact on the response time.

The experiments were carried out using a MacBook Air with the operating system OS X 10.8.5, processor 1.3GHz Intel Core i5, memory 8GB 1600Mhz DDR3 and flash storage 120GB.

9 Related Work

Risk-aware access control (see, e.g., [4, 5, 6, 10, 21]) has received a growing attention in the last few years. However, little attention is given to privacy aspects. The approaches that address privacy (see, e.g., [18, 16]) do so by adding privacy policy enforcement on top of the access control evaluation process. In our approach privacy risk as well as access risk are evaluated for every access request.

Risk Aware Access Control Models generally determine the risk as a function of the likelihood of a permission misuse and the cost of the permission authorized and misused. The likelihood of misuse can depend on the user trustworthiness and competence [4], the user behavior [1], and the uncertainty of the access decision [17]. The quantification of the cost of permission misuse has been addressed by several researches. Cheng et al. [6], in their assign a sensitivity label to every resource. The value of a resource is then determined according to its sensitivity. The cost of a misused permission depends on the resource's value. Molloy et al. [17] and Baracaldo et al. [1] propose to evaluate the cost in term of financial gain and damage. Chen and Crampton [4] do not explicitly calculate the permission misuse cost in their model, but mention that the cost of misuse is valued and used to define risk thresholds and risk mitigation strategies for every permission. In our model the risk results from the likelihood of identity disclosure which depends on the sensitivity of the requested information and the requestor trustworthiness.

Chen et al. [5, 14] propose to use, both user and system obligations as risk mitigation methods. An obligation describes some actions that have to be fulfilled by the subject, the system or a third part (e.g.an administrator), in a specific time window. In the literature we can distinguish between two categories of obligations: *provisions* or *pre-obligations* [2] are actions that must be executed prior to making an authorization decision; *post-obligations* are actions that must be fulfilled after the authorization decision is made. Unlike Chen et al. models that use post-obligations, monitor the fulfillment of these obligations after granting access and reward or punish users according to whether they have succeed or not to fulfill the required action, in our model we use provisions to enforce the risk mitigation strategy at run-time.

10 Conclusions

We have presented a model for information disclosure where access-control decisions are based on the risk associated with a data access request. Anonymization operations are used as risk-mitigation methods to compute views satisfy the accepted level of risk. This allows for granting access to requests that would otherwise be rejected. Our model leverages existing modes for Risk-Aware Access Control (most notably [5, 4]) but it also shows how they can be adapted so to support the controlled disclosure of privacy-sensitive information. We confirmed the feasibility of our approach by developing a prototype implementation of the proposed model and assessing it against a dataset widely used by the research community.

Acknowledgments

This work has been partly supported by the EU under grant 317387 SECENTIS (FP7-PEOPLE-2012-ITN).

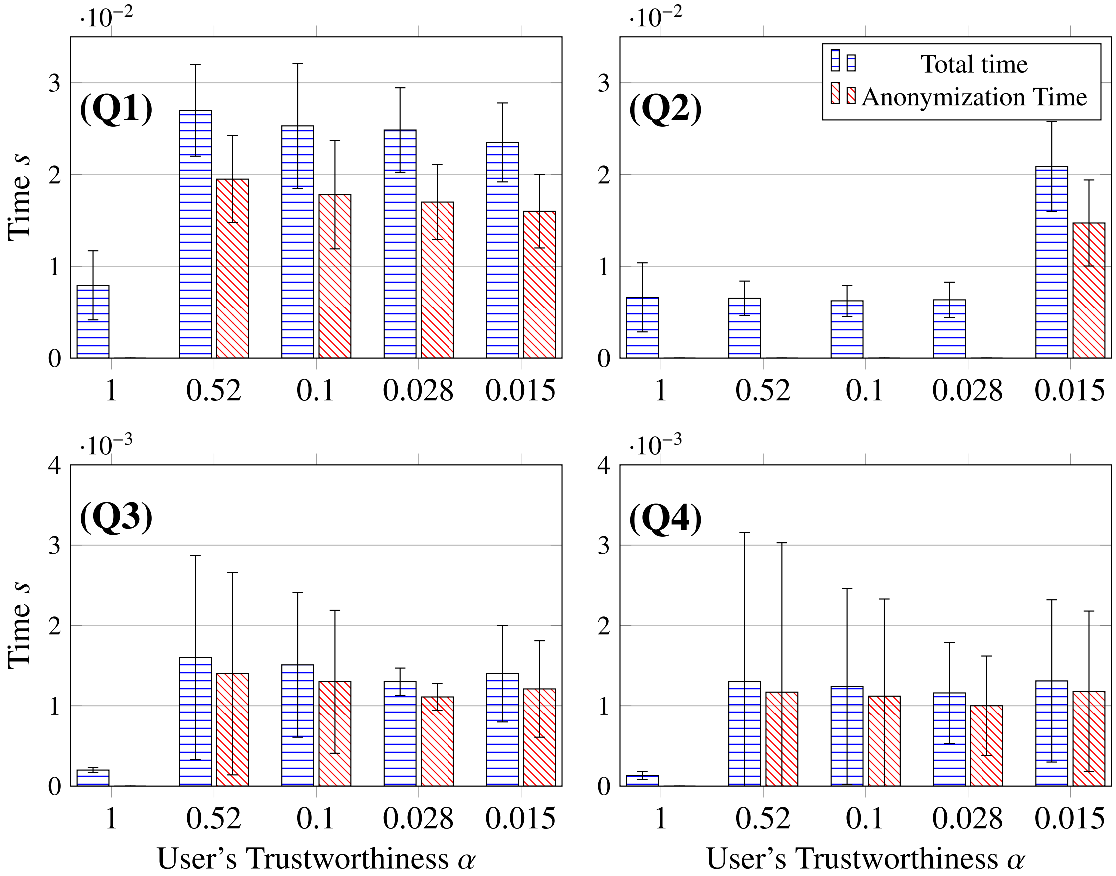

Figure 4. Average total response time (horizontal striped bars) and average anonymisation time (diagonally striped bars) for the four views and different trust levels.

References

[1] Nathalie Baracaldo and James Joshi. A trust-and-risk aware rbac framework: Tackling insider threat. In *Proceedings of the 17th ACM Symposium on Access Control Models and Technologies*, SACMAT'12, pages 167–176, New York, NY, USA, 2012. ACM.

[2] Claudio Bettini, Sushil Jajodia, X. Sean Wang, and Duminda Wijesekera. Provisions and obligations in policy management and security applications. In *Proceedings of the 28th International Conference on Very Large Data Bases*, VLDB '02, pages 502–513. VLDB Endowment, 2002.

[3] Michele Bezzi. An information theoretic approach for privacy metrics. *Transactions on Data Privacy*, 3(3):199–215, 2010.

[4] Liang Chen and Jason Crampton. Risk-aware role-based access control. In Catherine Meadows and Carmen Fernandez-Gago, editors, *Security and Trust Management*, volume 7170 of *Lecture Notes in Computer Science*, pages 140–156. Springer Berlin Heidelberg, 2012.

[5] Liang Chen, Jason Crampton, Martin J. Kollingbaum, and Timothy J. Norman. Obligations in risk-aware access control. In Nora Cuppens-Boulahia, Philip Fong, Joaquín García-Alfaro, Stephen Marsh, and Jan-Philipp Steghöfer, editors, *PST*, pages 145–152. IEEE, 2012.

[6] Pau-Chen Cheng, Pankaj Rohatgi, Claudia Keser, Paul A. Karger, Grant M. Wagner, and Angela Schuett Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In *IEEE Symposium on Security and Privacy*, pages 222–230. IEEE Computer Society, 2007.

[7] V. Ciriani, S. De Capitani di Vimercati, S. Foresti, and P. Samarati. Theory of privacy and anonymity. In M. Atallah and M. Blanton, editors, *Algorithms and Theory of Computation Handbook (2nd edition)*. CRC Press, 2009.

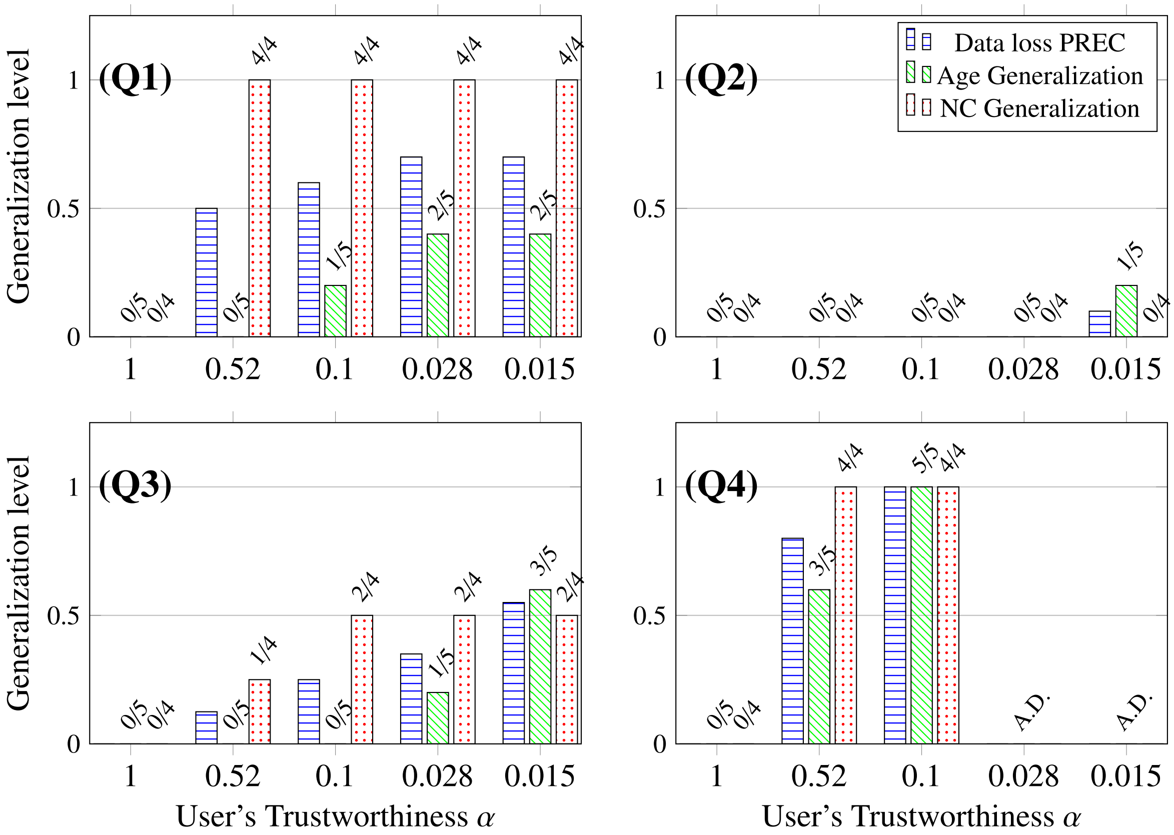

Figure 5. Generalization levels for the four views. Horizontal striped bar shows PREC metric (see text), diagonally striped bar the level of generalization for Age attribute and dotted bar the level of generalization for Native Country attribute. A.D. stands for Access Denied.

[8] Chris Clifton and Tamir Tassa. On syntactic anonymity and differential privacy. *Trans. Data Privacy*, 6(2):161–183, August 2013.

[9] T. Dalenius. Finding a needle in a haystack-or identifying anonymous census record. *Journal of official statistics*, 2(3):329–336, 1986.

[10] Luke Dickens, Alessandra Russo, Pau-Chen Cheng, and Jorge Lobo. Towards learning risk estimation functions for access control. In *In Snowbird Learning Workshop*, 2010.

[11] International Data Corporation (IDC). The digital universe of opportunities: Rich data and the increasing value of the internet of thingsr, April 2014.

[12] Florian Kohlmayer, Fabian Prasser, Claudia Eckert, Alfons Kemper, and Klaus A. Kuhn. Flash: Efficient, stable and optimal k-anonymity. In *ASE/IEEE International Conference on Privacy, Security, Risk and Trust*, SOCIALCOM-PASSAT '12, DC, USA, 2012. IEEE Computer Society.

[13] Ninghui Li, Tiancheng Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*, pages 106–115, April 2007.

[14] Timothy J. Norman Liang Chen, Luca Gasparini. XACML and risk-aware access control. Technical report, 2013.

[15] Ashwin Machanavajjhala, Johannes Gehrke, Daniel Kifer, and Muthuramakrishnan Venkitasubramaniam. l-diversity: Privacy beyond k-anonymity. In *ICDE '06: Proceedings of the 22nd International Conference on Data Engineering (ICDE'06)*, page 24, Washington, DC, USA, 2006. IEEE Computer Society.

[16] L.D. Martino, Q. Ni, D. Lin, and E. Bertino. Multidomain and privacy-aware role based access control in ehealth. In *Pervasive Computing Technologies for Healthcare, 2008. PervasiveHealth 2008. Second International Conference on*, pages 131–134, Jan 2008.

[17] Ian Molloy, Luke Dickens, Charles Morisset, Pau-Chen Cheng, Jorge Lobo, and Alessandra Russo. Risk-based security decisions under uncertainty. In *Proceedings of the Second ACM Conference on Data and Application Security and Privacy*, CODASPY '12, pages 157–168, New York, NY, USA, 2012. ACM.

[18] Qun Ni, Alberto Trombetta, Elisa Bertino, and Jorge Lobo. Privacy-aware role based access control. In *Proceedings of the 12th ACM Symposium on Access Control Models and Technologies*, SACMAT'07, pages 41–50, New York, NY, USA, 2007. ACM.

[19] M. C. Jason Programm Office. Horizontal integration: Broader access models for realizing information dominance, jsr-04-132. Technical report, 2004.

[20] Pierangela Samarati. Protecting respondents' identities in microdata release. *IEEE Trans. Knowl. Data Eng.*, 13(6):1010–1027, 2001.

[21] Riaz Ahmed Shaikh, Kamel Adi, and Luigi Logrippo. Dynamic risk-based decision methods for access control systems. volume 31, pages 447–464, 2012.

[22] Latanya Sweeney. Achieving k-anonymity privacy protection using generalization and suppression. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):571–588, October 2002.