Conceptual Framework and Architecture for Privacy Audit

Ksenya Kveler[2], Kirsten Bock[1], Pietro Colombo[3], Tamar Domany[2], Elena Ferrari[3], and Alan Hartman[2]

[1] Unabhaengiges Landeszentrum fuer Datenschutz (ULD)
[2] IBM Israel - Science and Technology LTD
[3] University of Insubria, Department of Theoretical and Applied Science

Abstract. Many ICT applications involve the collection of personal information or information on the behaviour of customers, users, employees, citizens, or patients. The organisations that collect this data need to manage the privacy of these individuals. In many organisations there are insufficient data protection measures and a low level of trust among those whose data are concerned. It is often difficult and burdensome for organisations to prove privacy compliance and accountability especially in situations that cross national boundaries and involve a number of different legal systems governing privacy. In response to these obstacles, we describe instruments facilitating accountability, audit, and meaningful certification. These instruments are based on a set of fundamental **data protection goals** (DPG): availability, integrity, confidentiality, transparency, intervenability, and unlinkability. By using the data protection goals instead of focusing on fragmented national privacy regulations, a well defined set of privacy metrics can be identified recognising **privacy by design** requirements and widely accepted certification criteria. We also describe a novel conceptual framework and architecture for defining **comprehensive privacy compliance metrics** and providing **assessment tools for ICT applications and services** using as much automation as possible. The proposed metrics and tools will identify gaps, provide clear suggestions and will assist audit and certification to support informed decisions on the trustworthiness of ICT for citizens and businesses.

1 Introduction

Rapid developments in IT technology are constantly offering new IT products and services that involve personal data processing. An enormous amount of digital information is collected, stored, and shared all over the world. Alongside the benefits, new risks arise when privacy concerns are not properly addressed during the development process. The worldwide exchange of personal data, electronic surveillance possibilities, and the discriminatory use of personal information for actions such as profiling and identity theft impose advanced privacy concerns for individuals and significant economic and reputational risks for businesses.

B. Preneel and D. Ikonomou (Eds.): APF 2012, LNCS 8319, pp. 17–40, 2014.

The fundamental right to the protection of personal data is recognized in Article 8 of the Charter of Fundamental Rights of the European Union and is set forth in the national data protection acts implementing the European Data Protection Directive 95/46/EC and the ePrivacy Directive 2002/58/EC. European national data protection acts recognize the rights of data subjects and impose obligations on data controllers, providing sanctions and remedies in cases of law infringements.

Privacy regulations apply to all the processes that relate to personal information. Personal information is any information relating to an identified or identifiable natural person (data subject) including secondary personal information such as log data. Public and private organisations are not always aware of the amount and existence of such personal identifiable data and therefore do not fully apply data protection regulations. Thus data protection non-compliance is a common problem in the EU. Data protection authorities who are responsible to safeguard the protection of the individual under the Data Protection Directive have neither the financial means, staff, nor powers to ensure full data protection compliance.

The processing of personal data is subject to rapid technical innovation. Law makers are therefore more and more refraining from regulating specific technical requirements for the processing of personal data. Instead principles of fair data processing and technical and organisational measures have been identified and put into legislation. Recently, regulators have introduced data protection goals (DPG) as a more comprehensive and adequate way to ensure data protection and privacy [DSK].

Compliance audits are one of the key mechanisms of the accountability principle and data protection regulations enforcement. Moreover, voluntary product audit privacy certifications are becoming more prevalent, providing competitive advantages and fostering user trust. Instead of sanctioning violations of privacy, promising market advantages offer positive incentives for implementing and observing privacy. For example, the EuroPriSe privacy seal [EuroPriSe] certifies that a product or service is compliant to regulations, based on an evaluation provided by privacy experts using a variety of time consuming legal and technical validation steps. Automatic tools are needed to assist auditors in assessing data protection compliance in an efficient and reliable way, improving the likelihood that the data protection goals defined by regulations are indeed met.

The difficulty in data protection is to produce comparable results in the absence of reliable privacy compliance indicators. This is partly due to the fact that data protection measures and compliance requirements are subject to legal decisions produced by the weighing of contrary principles such as data availability and data confidentiality. Thus, the validation and final result whether a specific data processing application is compliant with data protection law cannot be achieved automatically using information technologies alone; however, the evaluation process can be supported by privacy compliance indicators and an automatic toolset that allows for better and faster assessment of data protection compliance prerequisites and metrics. In doing so privacy compliance indicators support any data protection monitoring or assessment, such as Privacy Impact Assessments (PIA) [PIA] or an accountability program resulting in third party certification.

Our objective is to define a comprehensive set of privacy compliance metrics and create a set of assessment tools. Those tools will enable an audit of how an organisation performs based on those metrics, using as much automation as possible, and will provide clear suggestions for improvements. These metrics will be mapped to the protection goals, thus providing a means to assess which data protection principles are violated and why. The metrics will be defined and implemented with a set of privacy preserving techniques for their computation using a privacy by design approach.

Our privacy compliance assessment initially focuses on four main contexts: privacy policy compliance in general, and then specific compliance in the data storage, data sharing, and web sites operated by the organisation under audit. Other areas will be considered in future work.

The transformation of legal data protection requirements into technical metrics and the assessment of compliance poses a serious research challenge. New technologies are challenging lawmakers by introducing more complex systems and services which may be in conflict with the law. Current approaches in PIA and other assessment methodologies lack systematics and focus on the risks, based on the perspective of the organization either from a technology, an economic, or a legal point of view with each one of them demanding priority. They lack an explicit and systematic coverage of the protection of the interests of the data subjects. Data protection and data security operate from different perspectives and consider different risk sources: the attacker-model of data security aims to protect the operation of the organization primarily from persons who can pose special risks because they act as (former) employees, unfair or fraudulent citizens, as customers, or hackers. Data protection, by contrast, models organisations and their processing operations as potential attackers on the integrity and privacy of persons who are the data subjects in their roles of citizen, customer, client, patient, etc. Audit tools incorporating data protection goals go beyond data security assessment by operationalisation of data protection requirements which allow focusing on a common approach towards legal requirements, technical implementation and economic calculation without one of these domains dominating the other [Rost2012].

The paper is organised into three main sections where we discuss Data Protection Goals, Privacy Assessment, and the proposed Assessment Tool Architecture.

2 Data Protection Goals

The right to protection of personal data is established by Article 8 of the Charter and Article 16 TFEU and in Article 8 of the European Convention on Human Rights (ECHR) as well as by national laws. It imposes an obligation on private and public organisations to observe the right to privacy when processing personal data. The European Data Protection Directive only provides very general guidance on how data protection shall be implemented by technical and organisational measures. Due to the rapid developments in ICT, regulations on specific requirements are quickly outdated. The latest approach of privacy legislators [DSK] is to refrain from regulating specific technical security requirements; instead the regulators introduce Data Protection

Goals (DPG) as a more comprehensive and adequate way to ensure the protection of the individual [LDSG-SH]. The classification and applicability of protection goals has recently been elaborated in several articles and studies [AAL][GB2012][ZH2012]. The specific function and merit of "goals" is their ability to express a compulsory directive (normative ought), and their ability to address aspects of rules of operation, particularly of system applications. Having the same goals the different domains may be addressed coherently; experts may pursue the same direction and thus the same goal. The DPG provide a standardised approach to data protection investigations and audits [Rost 2012]. Data protection can be refined into specific fundamental protection goals: availability, integrity, confidentiality, transparency, intervenability, and unlinkability [RP2009]. The "classic" security goals of data security, availability, integrity, and confidentiality, focus primarily on guaranteeing the safe and secure maintenance of operation and infrastructure of an organisation. Data protection, by contrast, specifies these demands from the perspective of the data subjects (more precisely: citizens, customers, users, and patients) and augments this perspective with further demands derived from the basic rights of individuals. The specific demands can also be shaped into protection goals.

The following definitions are from Secttion 5 paragraph 1 of LDSG-SH. Availability is ensured if processes are timely available and can be used according to the rules. Integrity is ensured if data remain undamaged, complete, attributable, and up to date. Confidentiality requires that only authorized access is possible. Transparency means that the processing of personal data can be reproduced, verified and reviewed with reasonable. Unlinkability is ensured if personal data cannot or can only with unreasonably high efforts be collected, processed or used for another than its defined purpose. Intervenability requires a process to be designed in such a way that the data subject can exercise her rights effectively.

Availability, integrity, and confidentiality are classic, best practice IT-security protection goals since the 1990s.Data protection- goals also address transparency – as a prerequisite for the governance and regulation of technical-organisational processes– unlinkability – as an operationalisation of purpose bindingness/purpose separation – and intervenability – to operationalise data subject rights and the requirement on operators of systems to demonstrate that the data subjects have control over their information and are not dominated by the system. These goals comprehensively and explicitly address all relevant data protection aspects in a processing operation with respect to the data itself, the system and the procedures implemented [BM2012]. Using the same best practice methodology as in IT-security reduces translation errors between legal requirements and technical implementation [RB2012] and provides the methodology for privacy by design [RB2011].The data protection goals allow for the implementation of objective-specific protection measures which are technically and organisationally viable and controllable [Probst 2012].

With respect to governance the COBIT-framework [COBIT] offers a best practice to address regulation and controlling of processes. Key performance indicators and key risk indicators offered by this framework can be utilized as a regulative variable to implement and enforce data protection compliant processes in organisations. Some of the risk indicators have been specified for the RFID PIA Framework [PIA] and

were further put into more concrete form by the German Federal Office for Information Security (BSI) "Privacy Impact Assessment Guideline" [BSI]. The PIA Framework strives to address potential security and privacy risks and proposes measures to mitigate risks in the context of RFID. Nevertheless, this framework does not provide a systematic approach to specific privacy indicators (e.g., use of encryption to ensure confidentiality) or metrics (e.g., to determine the encryption level of a concept to secure confidentiality: in an organisation, how much information is sent encrypted?) to actually determine the scope of legal compliance. The PIA risk approach fails to identify those legal requirements which have not been implemented and whose non-implementation causes a potential threat. Privacy compliance indicators developed from DPG will cover not only risk based indicators but also performance and requirement based indicators. The privacy indicators focus on the performance of an IT product or service as compared to user and provider requirements and values, as defined in the legal requirements.

Data protection goals guide the systematic assessment of all privacy aspects of data processing for emerging technologies as well as for audits of running systems. Any audit will start with a description of its target of evaluation (ToE). This requires a comprehensive analysis of the processing operations, the players involved, and the data that is processed. Only an accurate analysis of who collects and processes personal data at what moment allows determining the applicable regulations and requirements.

The protection goals have been tested first on an implementation of ambient assisted living [AAL] technologies which aims to assist elderly or challenged individuals to live a more independent life. The use of AAL-technology provides a good example to illustrate a data protection audit scenario. To determine all relevant aspects a cube-model is used. The DPG CUBE (see [DSK], below) allows integrating and considering all parties involved and determines the work space which is the target of privacy evaluation (ToE).

In trying to determine the ToE one has to consider the data, the IT-systems, and the processes used. Relevant data groups in an AAL-scenario are intervention data (e.g., remote-medication by setting an injection or securing doors), vital signs (e.g., blood sugar, weight, temperature), behavioural data (sleeping-, eating-, working-, and resting times), technical infrastructure data, measurement and environmental data (e.g., temperature, lighting, humidity, sound level), triggered data (e.g., alarm contact, on-off switch). With these data in mind the IT-technology producing the data comes into focus as well as the processing, transmitting, archiving, and deletion.

In any use-case we see three process domains that are important in the processing of data: 1. Processes involving the data subject e.g. patient, 2. Processes involving the organisation or service provider (e.g., doctors office, hospital, public administration, insurance company), and 3. Infrastructural processes involving service providers of the organisation under 2. (e.g., data centres, access and content providers, but also controlling authorities and research institutes). Each process belongs to a process-owner who needs to be responsible and accountable for its design. This is where most often responsibility gaps are detected.

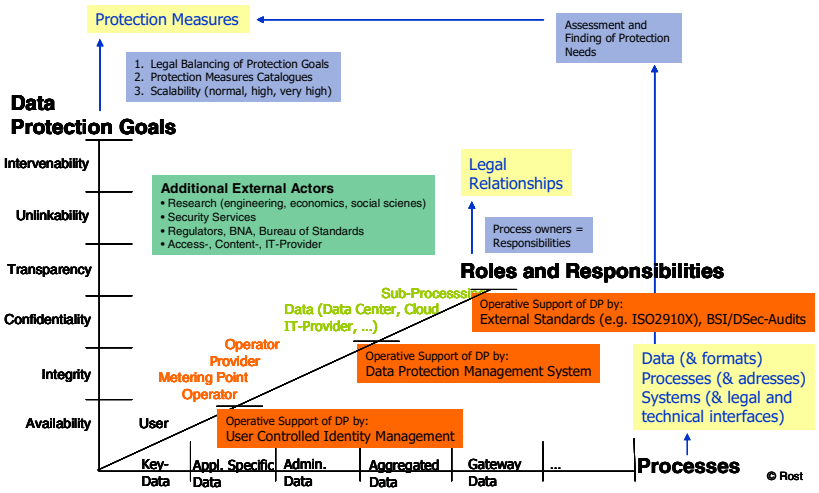

Fig. 1. DPG CUBE - differentiated risk assessment [Rost2011]

For example, if a person suffers from an early stage of Alzheimer and keeps forgetting to turn off the stove, close and lock the door at night or receives remote medication, AAL-technology can assist this person to control electrical items or the door by the use of sensor-systems which monitor status. Monitoring facilitates remote caretaking and allows the person to live at home and postpone admission to an institution. However, monitoring may entail a harmful degree of surveillance and deprivation of freedom if data protection is not observed. Lack of acceptance and trust may deter people from using such assistance.

The process-domains addressed in our AAL-example distinguish between the person concerned who needs to understand the system and must be able to control it (e.g., set-off an alarm); The process-owner, e.g. a home care organisation which is responsible for the functioning of the care system and its control mechanisms and interventions (e.g., what is to be done if an alarm is set off?); and on the infrastructure domain level a data centre which is processing the data collected at the home. By focusing on processes, the organisational structures, legal relationships and their justification, and also the responsibilities come into focus. By applying the DPG the extent of processing and access to data becomes visible and facilitates finding the appropriate legal basis or identifying an illegal process.

The three descriptive dimensions of the cube in our example are: 1) the processed data: e.g. door open – door closed, time stamp; 2) the basic processes and their owners that play a role in data processing: the person concerned opening and (not) closing the door, the care service monitoring the activities and offering the intervention service, the insurance paying for the service and the data centre hosting the IT; and 3) the protection goals to be applied to the complete use-case operationalising data protection norms [BM2012].

In our example the home would be equipped with technology to monitor the door status (closed/open) and send out an alarm in case the door stays open longer than

usual or for more than a specified time. With regard to the DPG we consider what kind of data is accessible/sent (transparency) to whom (integrity, confidentiality) for what reason (unlinkability) and how the person concerned and the organisation can or cannot intervene (intervenability). The care taking organisation needs to decide whether to send to the control room:

1. the exact time when the door was opened and closed (highly intrusive because habits and behaviour can be monitored),
2. the amount of time the door is open without a time stamp (less intrusive because no tracking of behaviour is logged), or
3. only the alarm-signal from a box taking the time the door is open at the home of the patient (privacy-friendly because no monitoring of behaviour takes place).

The care taking organisation also needs to decide whether the same access rights should be granted to the insurance company or whether it is sufficient to report the number of alarms or no data at all; and also how the IT-system integrity and confidentiality is ensured and whether redundant system architecture is provided. Finally, how to deactivate the system in case, e.g. during repair work, the craftsmen need an open door so that an alarm should not be set off.

The cube provides a model which enables us to describe the three dimensions relevant for data protection compliance analysis. The dimensions allow different view- and entry points. The advantage of the model consists of its ability to identify process responsibilities crucial for determining a ToE. All too often, only data security is addressed when IT-technology is assessed. The three dimensions of the cube ensure and enable all relevant stakeholders in an IT application to identify data types, processes and systems relevant for an assessment of the data protection performance of an IT application. The positioning of DPGs in a differentiated risk assessment for technical and organisational measures is illustrated above in [DSK].

The cube-model approach clarifies that data is always produced by a specific process which is carried out by specific application of IT-technology. When it comes to compliance every process must be covered by a legal basis which governs the activities of the person or technology involved. This is often overlooked when designing and evaluating ICT. The protection goals address compliance requirements in a comprehensive yet abstract manner. A catalogue of technical and organisational protection measures is deployed to implement the protection goals, as in the German Federal Office for Information Security [BSI] baseline protection [Probst2012].

On the dimensions of data and processes, we analyse which functional objectives are to be achieved by the process, what kind of data is required, which technical measures should have been chosen, and who is responsible in which role. Based on the processes, we can distinguish functions and determine the purposes that will generate the necessary data. In an assessment of a technical system, demonstrating what kind of data can be generated is of great importance. The data must be concretely defined and categorised in a manner that is relevant to its content.

The main focus of the third dimension of the DPG cube is the protection goals—the regulation and controllability of organisational processes. Here, process-organisational paradigms such as [ITIL, CoBIT], or processes based on ISO-oriented quality management are fairly well known. DPGs are expected to augment these paradigms. The approach, framework and toolset described in this paper will help determine the status of a system and match it to the target state, determined in the DPG-CUBE model. This procedure will support the data protection management of an organisation in continually monitoring data protection compliance.

In summary, contemplating data and processes is necessary to determine the purpose and necessity of a data processing. Contemplating data and protection goals leads to the analysis and determination of the protection demand of the data at hand, and governs the choice of technical and organisational protection measures. Contemplating processes and protection goals visualises processes and their regulation in the data protection management of the organisations involved. The generic DPG CUBE allows determining the protection demands of data and IT systems, measured processes, legal relationships and responsibilities, as well as the legally-weighted protection goals and protection measures. By addressing these relationships in a systematic manner, we address privacy compliance in a holistic way.

3 Privacy Assessment

The first step towards the development of privacy metrics and the associated assessment tools is the availability of a conceptual model for privacy onto which the assessed system can be mapped. Roughly speaking, it should define how the privacy-aware system is ideally supposed to be implemented. Privacy is a complex property characterised by numerous structural features, such as processing actions, data, purposes, obligations, users, authorisations to perform processing actions, and so on. All these concepts have to be properly formalised and composed to form the conceptual model for the privacy domain.

The conceptual model should be built around the principles of the DPG CUBE, and thus it has to support data specification, organisation and aggregation, as well as role-based data manipulation processes, protection goals and privacy policies, data protection mechanisms and adversary models. Although some proposals exist for languages to specify privacy policies, such as XACML [XACML], our analysis of the literature revealed the absence of a conceptual model that comprehensively considers all the concepts needed. Therefore, we first introduce a proper conceptual model. We approached this task by considering a subset of core privacy elements, originally formalised in [BL08], centred on the concept of purpose and related purpose-based access control policies. The preliminary version of the conceptual model, which is represented in the UML Class diagram in Figure 2, is discussed in [CF2012].

A key element of the model in Figure 2 is the concept of Purpose, which specifies the reasons for data collection and use. The other main components of the model are explained through our running example. Suppose that the AAL system manages data for *assisting*, *marketing* and *analysis* purposes. This can be modelled by *PS1,* an

instance of element PurposeSet, which groups these purposes. At any point in time the system administrator may decide to add a new Purpose to a PurposeSet, to remove or modify an existing Purpose. Therefore, an instance of PurposeSet can change dynamically over time. Data owners accessing the AAL system grant consent to use their data for specific purposes and prohibit their processing for other purposes. This is reflected in our conceptual model through the IntendedPurpose element which models a collection of allowed (intended) purposes (*aip*) and prohibited (intended) purposes (*pip*), which are bound to data (for simplicity, in Figure 2 we assume that data are organised according to the relational model). The purposes collected by an IntendedPurpose element must belong to the same PurposeSet. In our example, suppose that Bob, who suffers from Alzheimer, granted consent to process the state of the door of his apartment, which is modelled by means of an instance of element Data called *doorState,* for *assisting* purposes only. Bob also specified that his data cannot be processed for *marketing* purposes. These privacy requirements can be modelled through an IntendedPurpose *IP1*, including *assisting* in the *aip* component and *marketing* in the *pip* component, respectively, which is then assigned to *doorState.* In contrast, the element AccessPurpose collects the access purposes that are assigned to ProcessingActions that access and manipulate data. For instance, in our example, suppose we have *AP1*, an instance of AccessPurpose including purpose *assisting*, and that *AP1* is assigned to the ProcessingActions *monitoring, openDoor*, *closeDoor*, and *sendingAlarm*. A required (but not sufficient) condition to allow a ProcessingAction to process a set of data is that the purposes grouped by the IntendedPurposes assigned to the data and those collected by the AccessPurpose associated with the ProcessingAction belong to the same PurposeSet and are compliant.

Let us suppose that at a given point in time the system administrator introduces the new ProcessingAction *homeMonitoring*, which checks the state of windows and doors of the apartment where a patient lives for *security* reasons. Accordingly, the administrator 1) adds Purpose *security* to PurposeSet *PS1*, 2) introduces AccessPurpose *AP2* which includes *security* in the declared purposes, and 3) assigns *AP2* to *homeMonitoring*. As a consequence, Bob, who wants to benefit from the homeMonitoring service, 1) includes *security* in the *aip* set of *IP1* assigned to *doorState*, 2) introduces *IP2*, which is an instance of IntendedPurpose whose *aip* set consists of the purpose *security*, and 3) assigns *IP2* to *windowState*. The required condition for the execution of *homeMonitoring* on *doorState* and *windowState* is satisfied. On the other hand, *monitoring, openDoor*, *closeDoor*, and *sendingAlarm* cannot process *windowState* since the access purposes assigned to the ProcessingActions and the IntendedPurposes assigned to the data are not compliant (the *aip* set of *IP2* assigned to *windowState* does not include the Purpose *assisting*).Besides the straightforward concepts of User, the model in Figure 2 introduces the concepts of Role and ConditionalRole as a way to fine tune the administration of access rights. Role is composed of a set of attributes that characterise the role properties. For instance, a Role *domiciliaryAssistant* may be specified to model employees of the AAL service provider characterised by the *city* where they work, and the *zone* of the city. The attributes of a Role are initialized when the Role is assigned to a User. For instance, if the Role *domiciliaryAssistant* is assigned to user Mary, *zone* is set to 'Rosemont', whereas if it is assigned to Alice,

zone is set to 'Mont Royal'. To allow a more fine-grained management of access rights, element ConditionalRole extends Role with a condition that constrains when the Role can be assigned to a User. Such a constraint is a Boolean predicate defined in terms of the Role's attributes. For instance, the ConditionalRole *domiciliaryAssistantCR* extends *domiciliaryAssistant* with the constraint *"zone= 'Mont Royal'"*. This allows the assignment of the role to Alice, and prevents the assignment to Mary, whose zone of competence is Rosemont.

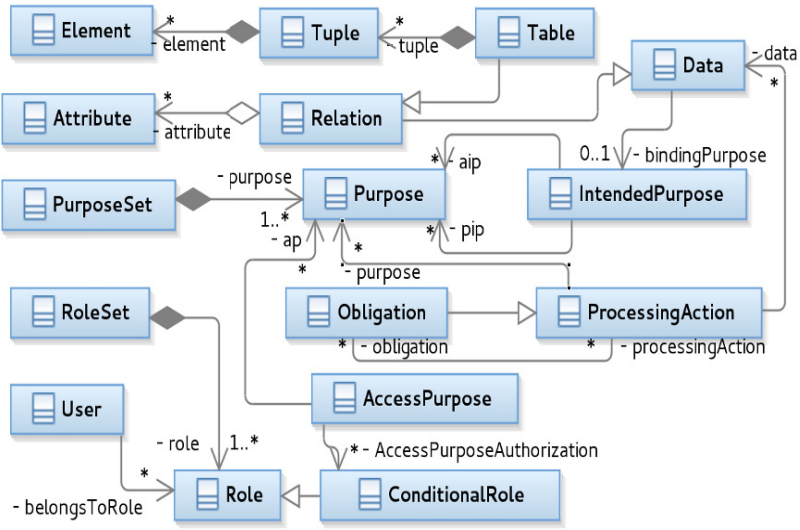

Fig. 2. A preliminary version of the privacy conceptual model

The conceptual model supports the specification of various types of privacy requirements that will then be used to drive the assessment phase. Privacy requirements are specified by means of a set of predicates that constrain the properties of the model elements. Such constraints capture all the properties that must be assessed for the system under analysis. For instance, an example of a privacy requirement is the privacy policy authorisation, which authorises access to data for actions associated with access purposes that comply with the intended purposes for which data owners granted their consent. This implies that, when an AccessPurposes AP is authorised for a ConditionalRole CR that extends a Role R, users that are authorised to play R and satisfy the constraint of the conditional role are allowed to execute processing actions with AccessPurpose AP on data D, only if AP complies with the IntendedPurpose IP specified for D. For instance, we can authorise domiciliaryAssistantCR to execute processing actions with AccessPurpose AP1. The authorisation is implicitly granted to all users who are authorised to play the Role domiciliaryAssistant and satisfy the domiciliaryAssistantCR constraint. Based on previous assumptions, in our scenario Alice is the only User with those characteristics. Let us consider the effect of the authorisation for the processing of Bob's data doorState. AP1 includes assisting as

access purpose and IP1, assigned to doorState, includes assisting in the set of allowed purposes. Therefore, AP1 complies with IP1. As such, Alice is allowed to execute the ProcessingActions monitoring, openDoor, closeDoor and sendingAlarm (associated with AP1) on Bob's doorState.

The conceptual model also supports the specification of data minimisation requirements, i.e., requirements that specify the largest set of data that can be collected and processed by the system under analysis. The elements *Relation* and *Attribute* (see Figure 2) are used to define the scheme of the data that can be collected. The minimisation requirement specifies that the tables schemes of the system under analysis must not have more attributes than those specified in the corresponding Relation element. The specification is performed by setting the Boolean attribute *minimisation* of element *Relation* to true. A similar mechanism is used to require the minimisation of the processed data. More specifically, the element *ProcessingAction* allows the specification of the *Relations* and the *Attributes* that can be involved in the execution of the *ProcessingAction*.

Requirements belonging to different categories (e.g., privacy policies and data minimisation) are integrated into a unified instance of the conceptual model (that is, as constraints on the corresponding entities). According to the DPG cube principles, multiple views can then be defined on the expressed requirements, allowing analysts to look at the requirements from the perspective of different data protection goals.

After specifying the requirements, the second step for the definition of the assessment mechanisms is the mapping of the components of the system under analysis to elements of the privacy conceptual model. This practice requires the analysis of the access control component, the database structure, and the configuration options of the assessed system.

Metrics are defined to evaluate the compliance of the mapped system to the reference requirements and, in case of non-compliance, to determine the non-compliant components (e.g., obligation support, role and purpose management). The output of such metrics evaluation will then be used by the assessment toolset to provide a set of recommendations on how to improve the system under assessment. A privacy metric expresses the similarity measure of the expected system state with the actual state of the system. Every policy involves a distinct set of conceptual model elements; therefore, a metric is required for every significant combination of conceptual elements.

For instance, suppose the AAL system supports Role-based access control. In this case, the mapping of AAL roles to conceptual roles is straightforward. In contrast, if roles are not supported, a possible countermeasure is to introduce a guest role with basic authorisations, and to assign it to all the users.

As another example, suppose that *doorState* is a column of the table *Patient* that collects patient's data. Suppose that the access to *doorState* is performed by the SQL query *extractDoorState*. Based on the conceptual model, an access purpose must be assigned to each processing action, whereas allowed and prohibited intended purposes must be assigned to data. If the AAL system does not record this information for *doorState* and/or *extractDoorState*, a warning message is returned by the assessment toolset along with proper recommendations to improve the system (e.g., information on the missing purposes).

The requirements constrain how the mapped elements must be related and configured. For instance, constraints can require

1. every action to be associated with an access purpose,
2. every data with an intended purpose, and
3. every user with a role.

Each constraint will have a weight that specifies its relevance. For instance, suppose that constraints 1 and 2 have weight 5, since the purpose is the key concept of the system, while the weight of constraint 3 is 1, since if no role is assigned to a user, he/she cannot be authorised to perform any action. These constraints can be used to achieve a compliance measure of the mapped model with the conceptual model. The number of satisfied constraints along with their weights will provide a compliance measure.

Furthermore, several dimensions of the mapped model can be measured and used to analyse the effectiveness of privacy protection. For instance, an analyst may be interested in counting and deriving:

1. all access purposes that comply with an intended purpose (e.g., in our example, considering *IP1* the resulting measure is *<1,{AP1}>*),
2. all roles that are authorised for an access purpose (e.g., considering *AP1*, we derive *<1,{ domiciliaryAssistant }>*),
3. all users that belong to a conditional role (e.g., for *domiciliaryAssistantCR*, we get *<1,{Alice}>*),
4. all access purposes that are granted to a user based on the conditional role the user belongs to (e.g., in case of Mary we derive *<0, Ø >*).

Dedicated metrics will also be defined to analyse the compliance of the system with data minimisation requirements. For instance, suppose the mapped elements include the Table Patient, which collects personal and sensitive data of all the patients, and the Relation PatientRM, which represents the scheme of Patient. Let us suppose that the AAL requirements specify the Relation PatientRR, whose minimisation attribute is set to true, and the ProcessingAction extractDoorStateR, which requires the processing of Attribute doorState of PatientRR under a minimisation constraint. Since the attribute minimisation of PatientRR is set to true, the relation PatientRM, which is derived from table Patient, can include only Attributes corresponding to those of PatientRR. A metric can check this constraint and count the number of attributes of PatientRM that are not included in PatientRR. Similarly, the SQL query extractDoorState is traced back to the processing action extractDoorStateR for which a data minimisation requirement is defined. Therefore, it is required to check that the set of data fields of extractDoorState is a subset of extractDoorStateR, and to count the number of possibly exceeding fields.

Even in the case of 100% compliance, further measures may be needed, because having all the necessary components for enforcing privacy-preserving access control, does not necessarily mean that the current access control configuration correctly enforces the desired privacy requirements.

Our conceptual model includes key concepts that are required for specifying general privacy requirements and supporting the assessment of existing systems with respect to these requirements. Moreover, the proposed metrics provide a quantifiable measure of different privacy aspects of the system under analysis. However, specific application domains may require additional conceptual elements that are not included in the current version of the conceptual model. Therefore, we cannot argue that the proposed conceptual model is complete. A direct parall can be traced with software testing. Testing cannot prove that the developed system satisfies the specification, but helps developers to increase the quality of the developed systems.

We support the assessment of implemented software systems both at run time and post execution. The metrics that can be evaluated at runtime are those associated with privacy policies expressed in terms of the current system state and/or previous states, whereas post-execution metrics involve events and states that refer to current, past and future points in times.

As far as the run-time assessment is concerned, metrics computation requires the analysis of the system behaviour with respect to the privacy requirements that express invariant properties of the system, such as the policy Authorisation introduced above. The derived measures are used to constrain the system execution by allowing, forbidding, or obligating the execution of operations associated with the involved events.

For instance, in case of invocation of an SQL query (e.g., openDoor), the run-time assessment metrics verify its authorisations. In case of non-compliance, the metric assessor will determine the non-compliant components. As an example, suppose that Mary executed action openDoor accessing Bob's data. The system will inform the proper controller that Mary does not satisfy the ConditionalRole constraint and therefore she is not authorised to execute the processing action openDoor. The output of the query evaluation will also provide a set of recommendations on how to improve/correct the configuration of the access control mechanism in place.

In contrast, privacy policies that can be checked by post-execution metrics are those that refer to current, past and future points in times. This allows for specifying complex trace execution constraints that involve retention conditions and obligations that refer to a future point in time.

For instance, suppose that the AAL system is used to monitor diabetic patients. Patients are required to periodically measure their glycaemia with a device that automatically informs the AAL system of the measured value. The entire measuring and notification process is modelled by means of the ProcessingAction monitoringNotification. The policy GlycaemiaAlarm states 'if the glycaemia exceeds a certain value, a physician must be contacted by phone within 5 minutes and informed of the measured value of the involved patient'. The called user should have activated the ConditionalRole doctor, and this role should be among those the patients gave the consent to access their personal data for assisting purposes. In this case, referring to the components of our conceptual model, glycaemia data are sensitive Patient data collected by the system for the Purpose of assisting the patient, doctor is a Role, and the phone-call is an Obligation associated with monitoringNotification. GlycaemiaAlarm is a policy that can be evaluated only by the post execution assessment, since checking the obligation requires delaying the analysis till after the query execution. Therefore, in our

example a phone call trace must be included in the AAL system log that should be checked five minutes after the invocation of monitoringNotification.

4 Assessment Tool Architecture

This section discusses the proposed architecture for the assessment toolset that will support the privacy audit process. The toolset will facilitate privacy metrics computation, identify gaps and provide compliance improvement recommendations. The intended users of the toolset are experienced external privacy auditors and internal privacy reviewers, so we assume a certain level of expertise and maturity. We also do not intend the tools to be totally automated, they are intended to make the work of these specialists more efficient, but not to replace them.

The assessment tool architecture which we propose is driven by two guiding principles, transparency and extensibility. It is designed to provide easy plugging of privacy compliance metric assessor components, each providing user interface (UI), analysis, and reporting capabilities for the particular technical metrics. This approach enables a gradual delivery process, starting with a limited set of assessor plugins, while targeting additional assessment techniques, areas, and privacy goals at a later stage. The modular approach also allows customized toolset packaging depending on particular target customer needs and audit type. The transparency principle aims to show the users how assessment decisions have been made, and what evidence has been collected during the analysis. The transparency is enforced by a required interface for all assessor plugins. This interface assures that each assessor provides evidence of the compliance or non-compliance and also advises for the improvement of the metric performance of the assessed artefact.

A conceptual architectural diagram of assessment tooling is shown in Figure 3. It comprises three main modules, the Audit Engine, the set of Metric Assessor Plugins, and the Administration module.

The Audit Engine is the core of the system, responsible for audit planning, privacy compliance analysis execution and report generation. The assessment process begins with the Planning component, which collects all the necessary information needed to understand the scope, plan and execute the analysis. Users might be asked to provide target privacy requirements for the system under assessment, in terms of the conceptual model described above, specify the desired assessment categories and supply any other metric-specific inputs. The Execution component performs the actual compliance analysis, according to the selections made and to information collected during the planning stage. Depending on the particular assessment logic and needs, if complete automation cannot be achieved, execution might be interrupted to collect additional user inputs. The Reports and Analysis component generates and presents detailed reports for the completed audit together with recommendations for improvements. In particular, Evidence Reports provide a record of all non-compliance evidence found. The Advisor generates recommendations for compliance improvements based on analysis results. Recommendations might be derived from a particular set of metrics that has been executed together with the higher level conclusions drawn from the

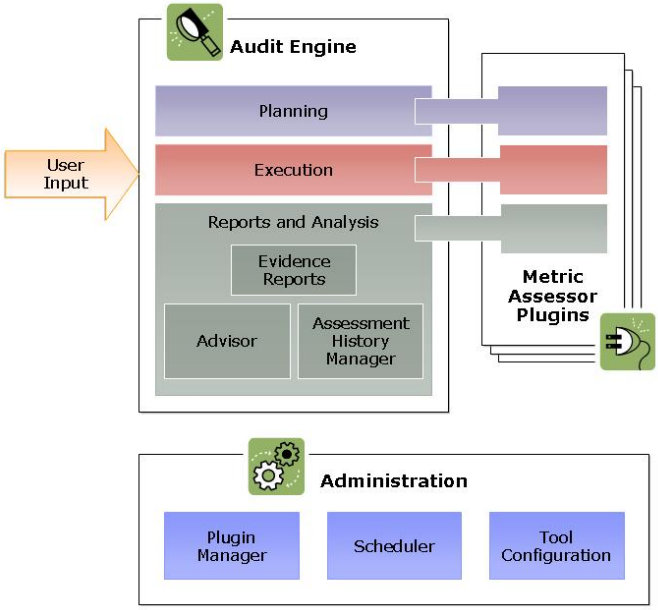

Fig. 3. Conceptual Privacy Audit tooling architecture

accumulative results of multiple metrics. The Assessment History Manager tool provides capabilities for viewing and for the analysis of previously executed audits. This enables tracking of compliance progress and improvement history.

A simplified sample audit report is shown in Figure 4. In reality, metrics will be more complicated and the resulting report will provide more details and capabilities. We plan to present audit results in a display compatible with the DPG CUBE model, providing different views on privacy compliance, from the perspectives of data protection goals, roles, processes and data. This visualisation capability is an item for future research.

The main audit report will show a general summary of the compliance analysis with grades and brief details for each metric that has been assessed. It will also allow rerunning of certain assessments to re-evaluate the metric after fixes have been applied; enable comparison to previous assessment results and drilling down into more detailed reports. The detailed reports will include more information about the data protection goals the particular metric is linked to and the concrete analysis steps that have been performed. They will also present the collected evidence of any non-compliance found and recommendations, such as guidelines on concrete measures to implement for compliance improvement.

The Metric Assessor Plugins are a set of pluggable components, each encapsulating everything that is needed in order to plan for, execute and report about a particular technical metric. Each metrics assessor contains:

1. its contribution to the planning UI,
2. its assessment execution code together with related user input UI if needed, and
3. its metrics-specific report generation capabilities.

The latter will include support for assembling and presenting metrics-specific assessment details, evidence and recommendations. Therefore, as shown on the conceptual architecture diagram in Figure 3, each metrics assessor possibly contributes to each of the three Audit Engine main modules. Each metrics assessor will also include specific information to facilitate its later use within an audit report, such as the assessment category it relates to (e.g., data store assessment) and the particular protection goals it is linked to.

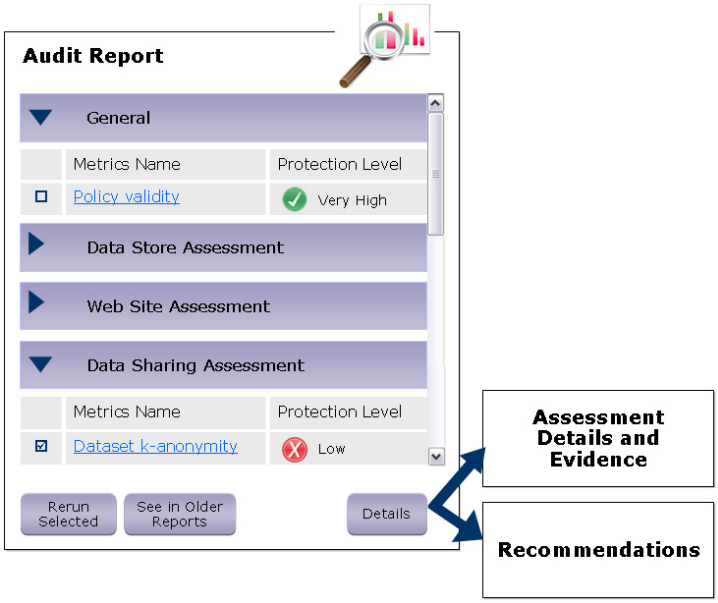

Fig. 4. Sample audit report

For example, one of the possible metrics can deal with overall validity, consistency and usability of the target privacy policy, specified in terms of the conceptual model described in the previous section. In other words, this metric should validate that the target privacy requirements make sense and would lead to a usable, conflict-free system. Following the AAL scenario above, the related assessor plugin can, for example, perform quantitative analysis and warn if there are too many users with a ConditionalRole doctor and too few with the ConditionalRole domiciliaryAssistantCR. It can also detect that according to the policy no AccessPurpose will be granted to some user based on his ConditionalRole (e.g., the case of Mary and the privacy policy Authorisation) or that there are no roles authorised for a certain AccessPurpose.

As described in the previous section, the system under analysis is mapped to the elements of the privacy conceptual model, by analysing the access control component, the database structure, and the configuration options of the assessed system. Thus another possible metric assessor plugin could examine the access policies that are actually used by the system (e.g., written using XACML) and verify their consistency with the desired target conceptual privacy policies and model. Moreover, yet another plugin could verify that the elements of the mapped system are related among them and configured according to the conceptual model constraints, for example, that every user is associated with a role.

Another metric assessor plugin could provide anomaly detection capabilities for identifying unauthorised or non-compliant database access by potential hackers, privileged insiders or other end-users. Patterns that do not conform to an established normal behaviour, and are thus considered suspicious, can be extracted by examining logs (post-execution assessment) and SQL queries intercepted at runtime. For example, in the AAL scenario, we could detect that while most of the time Alice triggers *extractDoorState* SQL query for accessing the database table *Patient*, which holds personal and sensitive data of patients, she also occasionally sends *extractPhoneNumber* query without being restricted to do so by the system. It might indicate, for example, that Alice is using patient's data (including the data of Bob) for marketing purposes without having an appropriate consent for that. This plugin would need to contribute to the planning UI, by providing plugin-specific screens where log location and database access parameters can be specified by the assessment tool users.

In our example, the AAL system manages data not only for *assisting*, but also for the purpose of *analysis*. In this case, it is necessary to assess anonymity of the datasets being shared with external parties (e.g., for statistical analysis). The Privacy Audit tool could provide a plugin that performs such an assessment based on anonymity metrics for privacy-preserving microdata release, for example *k*-anonymity [Sweeney2002]. It will contribute a dataset upload screen to the planning UI and its assessment result will appear under the Data Sharing Assessment category, as shown in the sample report in Figure 4. The plugin will respond with the result protection level ("Low" in Figure 4) according to the value of *k* for which k-anonymity is guaranteed and the particular value of *k* will be shown within the plugin-specific assessment details view.

A plugin for data minimisation assessment could support cases like the AAL door status and alert sending scenario described in the previous sections. The plugin will check that the system collects only the minimal amount of information needed, like alarm-signal events, but not the exact times the door was opened or closed. The assessment can be made by inspecting database schema and SQL queries for the presence of legitimate data elements only, according to the privacy requirements defined in the conceptual model. For example, let's say that privacy requirements specify that the attribute *alarmTimestamp* is part of the Relation *PatientRR,* whose *minimisation* attribute is set to true. The plugin will then analyse the schema of the database table Patient in the system under assessment, detect and warn about any fields beyond the above-mentioned permitted minimal set of attributes, such as *doorOpenTimestamp* or *doorClosedTimestamp*.

Yet another metric assessor plugin can deal with detection of obligation events and give an assessment of their compliance with the privacy requirements. For example, the plugin will analyse AAL log files, look for and correlate between records related to glycemia alarm events and events of establishing phone calls. The plugin will then verify that the user being contacted by phone has a ConditionalRole *doctor* and that this role is among the roles to which this particular patient gave the consent for accessing his *bloodSugarMeasurement* data.

The last module of the assessment tooling, the **Administration** module, provides management and configuration capabilities. For example, the Plugin Manager component will be responsible for the management of the Metrics Assessor Plugins repository, including viewing, editing and capabilities of adding new assessors or removing others. The Scheduler component allows the scheduling of automatic runs of pre-configured audits. The Tool Configuration component will enable adjustments of any other tool configuration, e.g., user interface options, general reporting options and any other settings.

We do not address the implementation details of any of the plugins at this stage of our work, Any discussion of algorithms for the privacy preserving computation of the metrics is beyond the scope of this conceptual design.

5 Related Work

In this section we describe relevant related work concerning privacy metrics and assessment tools.

5.1 Privacy Assessment

The systematic development of security and privacy assessment techniques is recognized today as a paramount requirement to assess the quality of any system with respect to its security and privacy guarantees [Jaquith2007]. However, most of the efforts developed so far focus on security, rather than on privacy [SA2009]. Savola [Savola2006] provides some high-level guidelines for the development of a framework for security evaluation based on security behaviour modelling and security evidence collection. The use of an ontology-based approach in support of run-time security monitoring is presented in [EOS09], and [HSHJ08] presents a security metrics framework, in which security metrics are associated with security patterns as a way to facilitate the interpretation of measurements. Information assurance metrics are described in [SPMNLH04], in which a review of existing metrics is performed, along with the proposal for a new taxonomy for information assurance metrics. A logic-based approach for reasoning about system security properties has been presented in [DFGK09] and applied to trusted computing.

Research on privacy assessment, however, is still in early stages, due mainly to the fact that despite recent advances in the field, privacy is still not a clearly defined concept. A discussion of measurements of compliance with security and privacy regulations and standards is presented in [Herrmann07]. Additionally, the preliminary study

reported in [Savola2010] presented a high level risk-driven methodology for privacy metrics development.

Much of the research efforts in the field have been instead devoted to anonymity metrics for privacy-preserving microdata releasing. Examples of such metrics are k-anonymity [Sweeney2002], l-diversity [MGKV06], t-closeness [LTV07], and differential privacy [Dwork2008]. These metrics capture different aspects of the disclosure risk, imposing some requirements on the association of an individual with the released sensitive private attributes, by making different assumptions on the attacker's background knowledge. Other works [Bezzi2010, RFD09] attempt to define an aggregated anonymity metric, based on information theory.

The need for a formal approach to privacy preservation was recognized by [Datta2011]. In this work, a logic-based model was defined with the aim of facilitating privacy polices specifications, and enforcement and compliance analysis. That model has been complemented with algorithms to check audit logs for compliance with privacy policies. It was also applied to several US privacy laws and resulted in the first complete logical specification and audit of all disclosure-related clauses of the HIPAA Privacy Rule [DGLKD10].

5.2 Existing Assessment Tools

Today privacy compliance analysis is still mainly done by specially trained experts, using reference documents, templates, forms and guidelines about how the audit should be conducted, rather than by applying automatic or semi-automatic analysis tools. Several commercial resources exist, for example the Privacy Management Toolkit [InfoShield] providing templates, forms, regulations library and expert commentary, and the Compliance Meter [CompMeter] which assigns privacy compliance scores based on the expert review of the templates filled by the customers. Following the European Committee for Standardisation (CEN) workshop on data protection & privacy (WS/DPP) in 2005, several workshop agreement (CWA) reference documents have been developed, including those on “Personal Data Protection Audit Framework" defining standard practices, templates, questionnaires and processes for audits.

The research on privacy metrics is still in its incipient stage and therefore not many tools exist. One of the more researched areas is the anonymity metrics of datasets, where some metrics are already available and thus naturally more tools exist to measure them. For example, the Privacy Analytics Risk Assessment Tool [PARAT] measures the risk of re-identification in different scenarios. There are tools providing capabilities both for anonymised dataset creation and for evaluation of the anonymisation status, like [CAT] and UTD Anonymisation ToolBox [UTDToolBox].

Another area where some metrics exist is on access control policy. Although access control is usually associated with security, it has also privacy aspects, mainly related to the policy adequateness and conflicts. For example, if the policy allows all the users to access all the data, there is probably a privacy violation. There are works that analyse the access control policy for finding conflicts and dominance [Vaniea08, Martin07], but little exists in terms of assessment of the quality and quantity of the privileges given to the various roles for accessing various data objects.

Protecting individual privacy on the web draws a lot of research attention. Web site privacy seal solutions, such as TRUSTe [TRUSTe], provide certain web site assessment capabilities. For example, TRUSTe is able to verify the site against its privacy policy but not compliance, and to scan the site for potential threats but not towards compliance with legal regulations. However, there is not enough transparency in terms of how exactly these capabilities are achieved, what particular assessment steps are performed and what techniques are used. Moreover, the TRUSTe seal does not address EU regulations, especially in terms of data collection.

Protecting individual privacy on a web site requires first of all that the site itself is secured from any type of hacking. Several commercial tools for testing web site vulnerability exist, for example, Acunetix Web Vulnerability Scanner [Acunetix] and IBM Rational AppScan [AppScan]. While these tools assess web site ability to resist various types of known attacks, it is not enough from the privacy preservation perspective. Web sites should be also examined in terms of the privacy policy existence and relevance, the limitation of the private data that the users are required to supply for clearly specified purposes, the processing that this data undergoes, the level of data protection within the data store, open sessions separation and more. There are no automatic tools with such wide assessment capabilities.

6 Conclusions

We have presented a conceptual framework for privacy auditing based on the legal concept of data protection goals, supported by a formal definition of technical privacy metrics. We described the privacy compliance assessment tools architecture, based on transparency and extensibility principles.

Our goal is to define a set of technical privacy metrics, by using a sound and formal approach to privacy quantification. This will represent a significant advance to the state-of-the-art for many reasons. First, the majority of previous proposals focused mainly on security. Those addressing privacy only considered data sharing by proposing a set of metrics to quantify the degree of anonymity of the released data. In contrast, we plan to develop a more general framework, in which data sharing is only one of the considered dimensions. Moreover, our ambitious goal is to combine both a sound and theoretical foundation of the developed metrics with an easy way of computing them and presenting the results to users.

We have designed a framework to allow easy plugging of privacy compliance metrics assessment components, each providing its specific user interface, analysis and reporting capabilities for the particular technical metrics. This approach enables gradual development, addition of more assessment techniques and areas in the future by other users and certification bodies, and the creation of a customised toolset packaging depending on particular target customer needs and audit type. Use of the framework will provide assessment transparency, by clearly showing how the decisions have been made, and what evidence has been collected during the analysis.

In this paper we have discussed the initial results we have achieved with our privacy preserving framework. The work is still in its early phases and a lot of work

remains to be done. In the near future we plan to work both on theoretical and implementation aspects. We plan to investigate the completeness and effectiveness of the conceptual model. We plan to identify new privacy metrics that can provide a measure of robustness to inference and statistical privacy attacks. We also plan to assess the capabilities and the scalability of our framework with case studies of realistic complexity. We also plan to work with users with the aim of getting feedback and suggestions on how it can be enhanced.

References

[AAL] Unabhaengiges Landeszentrum fuer Datenschutz (ULD). Juristische Fragen im Bereich Altersgerechter Assistenzsysteme, pre-study on behalf of VDI/VDE-IT, funded by the German Bundesministerium fuer Bildung und Forschung, https://www.datenschutzzentrum.de/projekte/aal/

[Acunetix] Acunetix Web Vulnerability Scanner, http://www.acunetix.com/vulnerability-scanner/

[AppScan] IBM Rational AppScan, http://www-01.ibm.com/software/awdtools/appscan/

[Article29] The Article 29 Data Protection Working Party was set up under Article 29 of Directive 95/46/EC, http://ec.europa.eu/justice/policies/privacy/index_en.htm

[Bezzi2010] Bezzi, M.: Expressing privacy metrics as one-symbol information. In: Proc. of the 2010 EDBT/ICDT Workshops (2010)

[BL08] Byun, J.-W., Li, N.: Purpose based access control for privacy protection in relational database systems. VLDB J. 17(4), 603–619 (2008)

[BM2012] Bock, K., Meissner, S.: Datenschutz-Schutzziele im Recht. DuD – Datenschutz und Datensicherheit 36(6), 425–431 (2012)

[BSI] German Federal Office for Information Security, http://www.bsi.bund.de

[CAT] Xiao, X., Wang, G., Gehrke, J.: Interactive Anonymization of Sensitive Data. In: SIGMOD 2009 (2009)

[COBIT] ISACA: COBIT Framework for IT Governance and Control, http://www.isaca.org/Knowledge-Center/COBIT/Pages/Overview.aspx

[CompMeter] The Compliance Meter, http://www.compliancehelper.com/compliance-meter/

[CF2012] Colombo, P., Ferrari, E.: Towards a modeling and analysis framework for privacy aware systems. Technical report, University of Insubria (2012) (submitted for publication)

[Datta2011] Datta, A., et al.: Understanding and Protecting Privacy: Formal Semantics and Principled Audit Mechanisms. In: Proc. of the International Conference on Information Systems Security (2011)

[DFGK09] Datta, A., Franklin, J., Garg, D., Kaynar, D.K.: A Logic of Secure Systems and its Application to Trusted Computing. In: Proc. of the IEEE Symposium on Security and Privacy (2009)

[DGLKD10] DeYoung, H., Garg, D., Jia, L., Kaynar, D., Datta, A.: Experiences in the Logical Specification of the HIPAA and GLBA Privacy Laws. In: Proc. of 9th ACM Workshop on Privacy in the Electronic Society (October 2010)

[DSK] Ein modernes Datenschutzrecht fuer das 21. Jahrhundert, Eckpunkte; Konferenz der Datenschutzbeauftragten des Bundes und der Laender, http://www.lfd.m-v.de/dschutz/beschlue/Eckpunkte.pdf (presented on March 18, 2010)

[Dwork2008] Dwork, C.: Differential Privacy: A Survey of Results. In: Agrawal, M., Du, D.-Z., Duan, Z., Li, A. (eds.) TAMC 2008. LNCS, vol. 4978, pp. 1–19. Springer, Heidelberg (2008)

[EOS09] Evesti, A., Ovaska, E., Savola, R.: From Security Modelling to Run-time Security Monitoring. In: Proc. of the Fifth European Conference on Model-driven Architecture Foundations and Applications, Enchede, The Netherlands (June 2009)

[EuroPriSe] EuroPriSe, the European Privacy Seal for IT Products and IT-Based Services, http://www.european-privacy-seal.eu

[GB2012] Geisberger, E., Broy, M. (eds.): AgendaCPS, Integrierte Forschungsagenda Cyber-Physical Systems, acatech Studie, Deutsche Akademie der Technikwissenschaften (2012)

[HDB] IBM Hippocratic Database (HDB) Technology Projects, http://www.almaden.ibm.com/cs/projects/iis/hdb/hdb_projects.shtml

[Herrmann07] Herrmann, D.S.: Complete guide to security and privacy metrics – measuring regulatory compliance, operational resilience and ROI. Auerbach Publications (2007)

[HSHJ08] Heyman, T., Scandariato, R., Huygens, C., Joosen, W.: Using security patterns to combine security metrics. In: Proc. of the 3rd Int. Conf. on Availability, Reliability and Security (ARES) (2008)

[InfoShield] The Privacy Management Toolkit, http://www.informationshield.com/privacy_main.html

[ITIL] Arraj, V.: ITIL - IT Infrastructure Library, The Basics, White Paper, http://www.itil-officialsite.com/AboutITIL/WhatisITIL.aspx (downloaded January 1, 2012)

[Jaquith2007] Jaquith, A.: Security metrics: replacing fear, uncertainty and doubt. Addison-Wesley (2007)

[JABK2008] Jouault, F., Allilaire, F., Bézivin, J., Kurtev, I.: Atl: A model transformation tool. Science of Computer Programming 72(1-2) (2008)

[LDSG-SH] Schleswig-Holstein Act on the Protection of Personal Information of February 9, 2000 last amended by Article 1 of the Act to amend the State Data Protection Act (January 11, 2012) (GVOBl. Schl.-H. p. 78)

[LTV07] Li, N., Li, T., Venkatasubramanian, S.: t-closeness: Privacy beyond k-anonymity and l-diversity. In: Proc. of the 23rd IEEE International Conference on Data Engineering (ICDE 2007). IEEE Computer Society (April 2007)

[Martin07] Martin, E.: Testing and Analysis of Access Control Policies. In: ICSE 2007 (2007)

[MASTER] Managing Assurance, Security and Trust for Services, European research project, http://cordis.europa.eu/fetch?CALLER=PROJ_ICT&ACTION=D&CAT=PROJ&RCN=85559

[MGKV06] Machanavajjhala, A., Gehrke, J., Kifer, D., Venkitasubramaniam, M.: l-diversity: Privacy beyond k-anonymity. In: Proc. of the 22nd IEEE International Conference on Data Engineering (ICDE 2006). IEEE Computer Society, Washington, DC (2006)

[OCL] OMG, Object Constraint Language (OCL) (2012), http://www.omg.org/spec/OCL/2.3.1

[PARAT] PARAT, http://www.privacyanalytics.ca/products.asp

[PIA] European Commission (EC): The Privacy Impact Assessment Framework for RFID Applications: PIA Framework (January 2011), http://ec.europa.eu/information_society/policy/rfid/pia/index_en.htm

[PICOS] Privacy and Identity Management for Community Services, European research project, http://cordis.europa.eu/fetch?CALLER=PROJ_ICT&ACTION=D&CAT=PROJ&RCN=85533

[PRBAC] Ni, Q., Bertino, E., Lobo, J., Brodie, C., Karat, C.-M., Karat, J., Trombeta, A.: Privacy-aware role-based access control. ACM Trans. Inf. Syst. Secur. 13(3), Article 24 (July 2010)

[Probst 2012] Generische Schutzmassnahmen für Datenschutz-Schutzziele. DuD – Datenschutz und Datensicherheit 36(6), 439–444 (2012), https://www.european-privac-seal.eu/results/articles/201206-DuD-Probst.pdf

[QVT] OMG, Meta Object Facility (MOF) 2.0 Query/View/Transformation (QVT) (2011), http://www.omg.org/spec/QVT/1.1

[RFD09] Rebollo-Monedero, D., Forne, J., Domingo-Ferrer, J.: From t-closeness-like privacy to postrandomization via information theory. IEEE Transactions on Knowledge and Data Engineering 99(1) (2009)

[RP2009] Rost, M., Pfitzmann, A.: Datenschutz-Schutzziele – revisited. DuD – Datenschutz und Datensicherheit 33(6), 353–358 (2009)

[Rost2011] Rost, M.: Datenschutz in 3D. DuD – Datenschutz und Datensicherheit 35(5), 351–353 (2011)

[RB2011] Rost, M., Bock, K.: Privacy by Design und die neuen Schutzziele. DuD – Datenschutz und Datensicherheit 35(1), 30–35 (2011)

[SA2009] Savola, R., Abie, H.: Development of Measurable Security for a Distributed Messaging System. International Journal on Advances in Security 2(4), 358–380 (2010) ISSN 1942-2636

[Savola2006] Savola, R.: A Requirement Centric Framework for Information Security Evaluation. In: Yoshiura, H., Sakurai, K., Rannenberg, K., Murayama, Y., Kawamura, S.-i. (eds.) IWSEC 2006. LNCS, vol. 4266, pp. 48–59. Springer, Heidelberg (2006)

[Savola2010] Savola, R.: Towards a Risk-Driven Methodology for Privacy Metrics Development. In: Proc. of the Symposium on Privacy and Security Applications (PSA 2010) (August 2010)

[Schmidt2006] Schmidt, D.C.: Model-Driven Engineering. IEEE Computer 39(2) (2006)

[SPMNLH04] Seddigh, N., Pieda, P., Matrawy, A., Nandy, B., Lambadaris, J., Hatfield, A.: Current trends and advances in information assurance metrics. In: Proc. of the 2nd Annual Conference on Privacy Security and Trust (2004)

[Sweeney2002] Sweeney, L.: k-anonymity: a model for protecting privacy. International Journal on Uncertainty, Fuzziness and Knowledge-based Systems 10(5), 557–570 (2002)

[TRUSTe] TRUSTe, http://www.truste.com/privacy_seals_and_services/enterprise_privacy/web_privacy_seal

[UML] OMG, Unified Modeling Language, v2.4.1 (2011), http://www.omg.org/spec/UML/2.4.1/

[UTDToolBox] UTD Anonymization ToolBox, http://cs.utdallas.edu/dspl/cgi-bin/toolbox/index.php

[Vaniea08] Vaniea, K., Ni, Q., Cranor, L., Bertino, E.: Access control policy analysis and visualization tools for security professionals. In: USM 2008: Workshop on Usable IT Security Management (2008)

[XACML] OASIS eXtensible Access Control Markup Language (XACML), http://www.oasis-open.org/committees/xacml/

[ZH2012] Zwingelberg, H., Hansen, M.: Privacy Protection Goals and Their Implications for eID Systems. In: Camenisch, J., Crispo, B., Fischer-Hübner, S., Leenes, R., Russello, G. (eds.) Privacy and Identity Management for Life – 7th IFIP WG 9.2, 9.6/11.7, 11.4, 11.6 International Summer School Trento, Italy (September 2011); Revised Selected Papers. Springer, Boston (2012) (to appear)