

2018

Evaluating a Reference Architecture for Privacy Level Agreement's Management

Vassiliki Diamantopoulou

University of Aegean, Samos, Greece, vdiamant@aegean.gr

Haralambos Mouratidis

University of Brighton, United Kingdom, h.mouratidis@brighton.ac.uk

Follow this and additional works at: <https://aisel.aisnet.org/mcis2018>

Recommended Citation

Diamantopoulou, Vassiliki and Mouratidis, Haralambos, "Evaluating a Reference Architecture for Privacy Level Agreement's Management" (2018). *MCIS 2018 Proceedings*. 28.

<https://aisel.aisnet.org/mcis2018/28>

This material is brought to you by the Mediterranean Conference on Information Systems (MCIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in MCIS 2018 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

EVALUATING A REFERENCE ARCHITECTURE FOR PRIVACY LEVEL AGREEMENTS MANAGEMENT

*Research full-length paper
Track Security and Privacy*

Diamantopoulou, Vasiliki, University of the Aegean, Samos, Greece, vdiamant@aegean.gr
Mouratidis, Haralambos, University of Brighton, Brighton, UK, h.mouratidis@brighton.ac.uk

Abstract

With the enforcement of the General Data Protection Regulation and the compliance to specific privacy- and security-related principles, the adoption of Privacy by Design and Security by Design principles can be considered as a legal obligation for all organisations keeping EU citizens' personal data. A formal way to support Data Controllers towards their compliance to the new regulation could be a Privacy Level Agreement (PLA), a mutual agreement of the privacy settings between a Data Controller and a Data Subject, that supports privacy management, by analysing privacy threats, vulnerabilities and Information Systems' trust relationships. However, the concept of PLA has only been proposed on a theoretical level. In this paper, we propose a novel reference architecture to enable PLA management in practice, and we report on the application and evaluation of PLA management within the context of real-life case studies from two different domains, the public administration and the healthcare, where sensitive data is kept. The results are rather positive, indicating that the adoption of such an agreement promotes the transparency of an organisation while enhances data subjects' trust.

Keywords: Privacy Level Agreement, Security Requirements Engineering, Privacy Requirements Engineering, Practical Evaluation.

1 Introduction

Globally, Public Authorities and private organisations offer an increasing number of e-services to individuals (i.e. service consumers, citizens, patients). As a result, Information Systems (IS) are developed for different areas of e-services (e.g., healthcare, registration services), they operate in a way that supports governmental and industrial initiatives and aim to improve transparency of service consumers' data sharing. However, in developing such IS, Authorities face important challenges related to the privacy of individuals' data. This is becoming even more important with the enforcement of the General Data Protection Regulation (GDPR) that requires compliance to specific privacy- and security-related principles. Within the context of this regulation, organisations that provide e-services or keep data, are known as Data Controllers (and Data Processors, but for simplicity reasons throughout the paper we are using the term of Data Controllers), while those receiving e-services or they have the ownership of the data are the Data Subjects. In this context, it is important that Data Controllers, through their IS, are able to clearly specify Data Subjects' privacy needs, provide them with feedback on how their data is shared and inform them of potential data privacy conflicts. Moreover, Data Controllers should enable Data Subjects to understand potential threats and vulnerabilities to their privacy requirements, as well as trust relationships that might endanger their privacy. It has been argued in the literature (Diamantopoulou, V. et al., 2017b; Diamantopoulou, V., Pavlidis, M., and Mouratidis, H., 2017; Cloud Security Alliance, Privacy Level Agreement Working Group, 2013) that such challenges can be addressed through the use of a Privacy Level Agreement (PLA), which supports a mutual agreement between a Data Controller and a Data Subject, regarding the Data Subject's privacy needs and the transparency of their data sharing.

In the context of cloud services, in order to facilitate the satisfaction of a client's privacy requirements, the notion of PLA has been used (Cloud Security Alliance, Privacy Level Agreement Working Group, 2013; Ahmadian, A. S. et al., 2015) as part of a Service Level Agreement (SLA) (Bouman, J., Trienekens, J. and Van der Zwan, M., 1999; Keller, A. and Ludwig, H., 2003) as a bilateral agreement between cloud service providers and their clients on how and to what degree the data of the latter should be protected. In the area of IS engineering, we have presented in previous work (Diamantopoulou, V. et al., 2017b) a language that supports the modelling of PLAs for Data Controllers' IS services and a theoretical analysis of how the PLA can be used to support implementation of the GDPR (Diamantopoulou, V. et al., 2017b).

In this paper, we build on our previous work, and we make the following contributions: First, we propose a reference architecture for PLA management, both at design time and during runtime. This is important because it allows the creation of PLAs as digital contracts, between Data Subjects and Data Controllers, which can be analysed at design time and used at runtime to guarantee that the Data Subjects' privacy is respected, based on their privacy preferences. Next, we perform an empirical evaluation across two domains (public administration and health care), which is focused on evaluating the applicability of the PLA and on the identification of potential benefits and difficulties raised from its usage. This is the first effort in the literature to practically evaluate PLAs in the context of Information Systems services.

The remainder of the paper is set out as follows: Section 2 discusses related work, while Section 3 presents a brief overview of the PLA structure and introduces the reference architecture. Section 4 illustrates the application of our work on case studies from the Public Administration and Healthcare domains, and Section 5 evaluates the results of this. Finally, Section 6 summarises the conclusions and raises issues for further research.

2 Related Work

Work on PLAs has been limited so far. The Privacy Level Agreement Working Group of the Cloud Security Alliance has defined a PLA in the context of cloud services (Cloud Security Alliance, Privacy Level Agreement Working Group, 2013). Similarly, the concept of PLA has been presented by DErrico and Pearson (2015) as a standardised way for cloud providers to describe their data protection practices. In the same way to the Cloud Security Alliance proposal, this work focuses on the cloud environment

and the PLA is considered as a means for the cloud providers to ensure that their privacy policy is communicated to the service consumers. However, these works are limited only to the examination of privacy aspects of cloud provision and do not provide support for specification of user (e.g., citizen) preferences and needs or ways to define privacy threats and vulnerabilities related to these needs, important aspects of the PLA of this work.

A study of Ahmadian, A. S. et al. (2015) presents a tool-based approach that facilitates the ISO27001 certification process for cloud service providers, appropriate for SMEs. The authors present the ClouDAT framework, a cloud-specific risk assessment process that allows the automatic generation of ISO27001 compliant documentation, based on the outcomes of the risk assessment. The PLA is used as an input to their tool in order to perform security checks. That framework focuses on the security analysis on cloud environment, with little emphasis on privacy requirements been given. Also, this approach does not support data subjects (e.g., citizens, patients) specifying their privacy preferences and needs.

We have also contributed to the State of the Art by formally specifying the concept of the PLA, based on an XML schema, which enables its automated use (Diamantopoulou, V., Pavlidis, M., and Mouratidis, H., 2017). In addition, we have extended our work (Diamantopoulou, V. et al., 2017b) so the PLA can be used to support the GDPR, providing the metamodel of the PLA. Finally, an application of the PLA in healthcare domain is presented in ((Diamantopoulou, V. et al., 2017a), focusing mainly on the functionalities of the PLA management platform.

The idea of a standardised way for web sites to communicate with users about their privacy policies in a standard machine-readable format has been introduced by the Platform of Privacy Preferences (P3P) Project (Platform for Privacy Preferences (P3P) Project, 2016). This standard enables web browsers and other user agents to interpret privacy policies on behalf of their users, assisting them to decide when they exchange data with web sites. However, P3P was designed for static environments where users' privacy preferences are not expected to change, and it also provides limited support for specification of privacy threats and vulnerabilities that might endanger the privacy needs.

A study of Drogkaris, Gritzalis, and Lambrinouidakis (2013) proposes an architecture that promotes the employment of privacy policies and preferences. The authors introduce the Privacy Controller Agent for storing and comparing service providers' privacy policies and user privacy preferences. However, this work does not provide an agreement between two entities (e.g., PA and citizens) but rather an architecture to define privacy policies.

On the other hand, the literature provides many examples of works that focus on the specification of Service Level Agreements (SLAs) which refer to the mutual agreement that defines the obligations and the requirements both of a service provider and a customer (e.g., Bouman, J., Trienekens, J. and Van der Zwan, M., 1999; Keller, A. and Ludwig, H., 2003; García, J.M. et al., 2017; Mohamed, M., et al., 2017). In contrast to the PLA concept, an SLA does not take into account privacy aspects of the agreement between a service provider and a service consumer.

Various approaches have been proposed in the literature for systematically capturing security and privacy requirements. The Privacy Safeguard (PriS) (Kalloniatis, C., Kavakli, E., and Gritzalis, S., 2008) methodology enables the elicitation of privacy requirements in the software design phase, where privacy requirements are modelled as organisational goals. Next, in (Spiekermann, S. and Cranor, L. F., 2009) the authors adopt the concepts of privacy-by-policy and privacy-by-architecture, and propose a three-sphere model of user privacy concerns, relating it to system operations (i.e. data transfer, storage and processing). Additionally, the Modelling and Analysis of Privacy-aware Systems (MAPaS) framework (Colombo, P. and Ferrari, E. (2012) is a framework for modelling requirements for privacy-aware systems. Regarding security requirements methodologies, literature provides numerous works that have been developed. Indicatively, we present SQUARE (Security Quality Requirements Engineering) methodology (Mead, N. R., and Stehney, T., 2005) which is a risk-driven method that supports the elicitation, categorisation, prioritisation and inspection of the security requirements through a number of specific steps. It also supports the performance of risk assessment to verify the tolerance of a system against possible threats. Next in (Faßbender, S., Heisel, M., and Meis, R., 2014a,b) the authors propose the Problem-based Security Requirements Elicitation (PresuRE) Methodology that facilitates the

identification of security needs during requirements analysis of software systems, providing a computer security threat recognition and then the development of security requirements. In (Salini, P., and Kanmani, S., 2013) the authors propose Model Oriented Security Requirements Engineering (MOSRE) framework for Web Applications which considers security requirements at the early stages of the development process, covering all phases of requirements engineering and suggesting the specification of the security requirements in addition to the specification of systems requirements. Differently than these works, our study provides a start-to-end implementation of a security and privacy management approach that takes into account the PbD principles, starting with the elicitation of the user privacy needs and ending with the provision of PA online services, and the security by design principles, by conducting security analysis of the IS of the service provider, allowing the detection of threats and facilitating the selection of suitable security mechanisms to mitigate potential attacks.

3 Reference Architecture for Privacy Level Agreements Management

In the context of our work, we define a PLA as the mutual agreement of the privacy settings between a Data Controller (i.e. Public Administration (PA)) and a Data Subject (i.e. citizen), where the former will commit to provide and maintain these settings throughout the provision of the service. The PLA is delivered in a form of a structured agreement that consists of fields, each of them capturing important and obligatory information with regards to privacy of Data Subjects' data. Moreover, an XML schema has been proposed (Diamantopoulou, V., Pavlidis, M., and Mouratidis, H., 2017) to enable the creation and management of machine-readable PLAs, allowing its utilisation by distributed IS, thus addressing interoperability issues. In this section, we present an overview of the PLA structure, as proposed in (Diamantopoulou, V. et al., 2017b; Diamantopoulou, V., Pavlidis, M., and Mouratidis, H., 2017) to support understanding of the rest of the paper. We then describe the proposed reference architecture for PLA management, across two levels: Design Time and Runtime.

3.1 Brief Overview of PLA structure

The structure of the PLA contains two sections, the first with information related to the Data Controller and the second to the Data Subject. In turn, each section contains a number of fields that include information related to the privacy of the Data Subjects' data.

Data Controller Section The field *Identity* presents the contact details of the Data Controller and the responsible administrator, i.e. name, place of establishment. The field *Data* specifies which personal data the Data Subject needs to provide to the Data Controller. Next, the field *Data Processing Ways* provides information about processing and storing Data Subjects' data. The field *Data Sharing Preferences* contains information about third parties that can have access to Data Subjects' data. The field *Data Privacy Measures* specifies the technical, physical, and organisational measures in place to protect Data Subjects' personal data against any destruction or loss, alteration, unauthorised use, modification, disclosure of access, and any other unlawful form of processing. The field *Privacy Threat Analysis* provides the threat analysis of the Data Controller's privacy requirements. The field *Trust Analysis* accordingly, provides the trust analysis of the Data Controller's privacy requirements. Last, the field *Law Compliance* gives information on whether privacy requirements are compliant with corresponding privacy national or EU laws and regulations.

Data Subject Section The field *National Public Authority* contains the necessary details of the Authority responsible for protecting Data Subjects' personal data rights. The field *Data Subject Privacy Preferences* has the privacy preferences of the Data Subject that have been collected by the Data Controller. The field *History Based Assessment* consists of an analysis of the Data Subjects' privacy preferences and the generation of a prediction of the possible outcomes of subsequent requests. Last, the field *Data Value* contains the average of i) the Data Subject's perspective concerning their data, ii) the valuation that the Data Controller provides, and iii) the average valuation of all the Data Subjects.

3.2 Privacy Level Agreement at Design Time

The IS of a Data Controller which offers e-services has to be set up in a way to facilitate the management of the PLA. To this end, the Data Controller uses appropriate methods and tools that will enable them to capture the relevant information during the design time of the system-to-be. A reference architecture that implements the relevant process during the design time is depicted in Figure 1.

The Data Subject who wishes to use the *Data Controller's e-service* has to provide the necessary data. This data is then stored to the *Data Controller database*. Once the data is captured, the Data Controller needs to specify their processing rights, representing the purpose for which the data is collected. So, the Data Controller captures their *Privacy Requirements*. This can be achieved by any method and tool available in the literature.

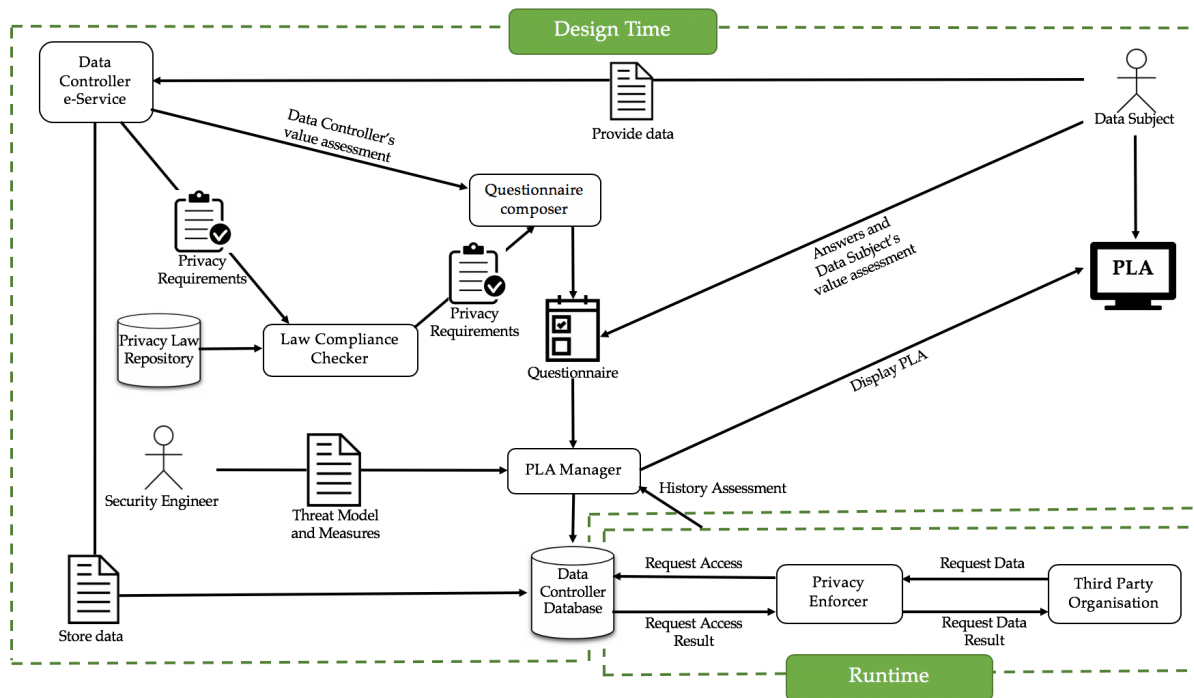


Figure 1. Reference Architecture for PLA Management

Furthermore, the relevant laws and regulations at a national and international level are encoded and stored to the *Privacy Law Repository*. Compliance of the privacy requirements with such laws and regulations has to be checked and verified by the *Law Compliance Checker*. The latter collects i) the specified privacy requirements and ii) the relevant laws, in order to shape the final privacy requirements. Also, at this stage, the Data Controller proceeds to a valuation of the collected Data Subjects' data in order to provide this information later in the PLA.

Next, a Data Controller's *Security Engineer* provides a privacy threat model that depicts the potential threats that can exploit system vulnerabilities and cause privacy breaches. These threats can be mitigated by identifying additional *Security and Privacy Measures*. Furthermore, an appropriate system architecture has to be developed by the Security Engineer along with an examination that this architecture does not violate the privacy requirements.

The last activity of the Data Controller during the design time is to define the questions of the questionnaire which will be used by the *Questionnaire Composer*. After the latter has received this information, they release the *Questionnaire*. Then, the Data Subject is requested to answer this questionnaire, stating their privacy preferences and providing their own assessment, revealing their perception about the value of their data. The completed questionnaire is then collected by the *PLA Manager* which generates the PLA, stores it to the *Data Controller Database* and displays it to the Data Subject.

3.3 Privacy Level Agreement at Runtime

The relevant information during runtime is also depicted in Figure 1. Once the IS of the Data Controller is implemented and put in operation, the Data Subject can register to the system in order to use its e-services. At this point, the PLA between the Data Subject and the Data Controller is generated, since all the relevant information has been captured. To support the PLA generation, we have defined an XML schema¹ which supports machine-readable analysis (Diamantopoulou, V., Pavlidis, M., and Mouratidis, H., 2017).

Access to Data Subject's data will depend on the individual PLA of the Data Subject. Every time a *Third-Party Organisation* is requesting access to the Data Subject data, the request is received by the *Privacy Enforcer*. The privacy enforcer parses the policies defined by the Data Subject's answers, and according to them, the requests are either allowed or denied. These requests are then sent to the PLA manager, creating *History Assessment* information which captures the ratio of allowed/denied requests of Data Subjects' data. This ratio is then depicted in the PLA, keeping Data Subjects updated regarding the level of openness of their data. This information is also useful to the Data Controllers in order to revise their security architecture of their IS.

4 Case Study

The aim of this section is to illustrate the applicability of the PLA and to identify potential benefits and difficulties by its adoption. In order to have objective results, we evaluate the application of the PLA management at real case studies of two different pilots, run in the following two domains: Public Administration and Healthcare domain.

4.1 Public Administration domain

The first pilot examines the applicability of the PLA in the context of a governmental IS. A local PA, the Municipality of Athens (MoA), provides e-government services, using an IS that stores and manages personal data of Athenian citizens, namely MACS (Municipality of Athens Computer Services). The main purpose of MACS is to interconnect with collaborative to MoA organisations, such as hospitals, banks, sport facilities, and many others, and also to store and transmit information necessary for the completion of a citizen's request (e.g., the issue of a birth certificate) without requiring a citizen's physical presence.

Although MACS supports multiple e-services, due to space limitations in this paper, we focus on the e-service related to a citizen's subscription to a local fitness centre. MoA offers 15% discount to Athenian citizens, i.e. the ones who can prove that they permanently live in the city of Athens. As proof of their locality, citizens are required to provide their birth certificate, issued by MoA. An Athenian citizen requests the issue of their birth certificate using the MACS system, which directly transmits the required information to the IS of the fitness centre, as part of the e-government services in MoA. It is also required that the citizen provides a medical certificate regarding their health condition, especially if they suffer from any heart disorders. The citizen has to visit their physician in order to be provided with this certificate. If the citizen has recently received a medical certificate, MACS is able to retrieve it from the hospital's database and, with the consent of the citizen, to forward it to the IS of the fitness centre. In this scenario, the MACS will handle the birth- and the medical certificate of a citizen. Based on the one-stop concept (Tambouris, E. and Wimmer, M., 2008; Sedek, K.A., Sulaiman, S. and Omar, M.A., 2011), i.e. sharing information across multiple IS belonging to different governmental authorities and organisations, the MACS is interconnected with both other PAs and service providers, so the citizen has to provide their privacy preferences regarding the information that they wish to share and how. Moreover, the administrators of MoA must ensure transparency of the data sharing process with other organisations, thus minimising any reluctance of citizens in sharing their personal data, and eventually in using e-services of MoA through MACS. In order for citizens to take the right decision regarding the sharing

¹ http://www.sense-brighton.eu/xml_pla/

of their data, they should be aware of i) the relevant mechanisms MoA administrators apply to their IS in order to protect citizens privacy, and ii) the value of their data.

4.2 Healthcare domain

The second pilot, from the healthcare domain, provides different challenges since it involves cross-border exchange of medical data, which is of more sensitive nature, and also there are emergency cases where the Data Subject (i.e. patient) is not able to get involved in the process of the provision of consent.

This pilot involves two paediatric clinics, namely Ospedale Pediatrico Bambino Gesù (OPBG) and Hospital Infantil Universitario Niño Jesús (HIUNJ), from two different countries, Italy and Spain, respectively. These clinics use a telemedicine platform to exchange medical information of patients. The processes in telemedicine services fall within the sensitive data being processed by electronic instruments, which are currently regulated by the provisions of Directive 2002/58/EC (European Commission: Directive 2002/58/EC of the European Parliament and of the Council, 2002). The methods and the solutions necessary to ensure confidentiality, integrity and availability of data should therefore be adopted in accordance with the security measures explicitly provided in the Directive 95/46/EC (European Commission: Directive 95/46/EC of the European Parliament and of the Council, 1995), covered under the GDPR (European Parliament: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation, 2016)) and the new regulation replacing the Directive 2002/58/EC, which can be found in (European Commission: Proposal for a Regulation of the European Parliament and of the Council, 2012). In our scenario, an Italian patient suffers from a disease which is characterised by doctors as quite rare and complex enough. The physicians who provide health care to the patient, decide to discuss this case with a specialist group in another hospital, HIUNJ, in Spain. Due to several circumstances, it is impossible for the physicians of the two countries to arrange a meeting with physical presence. Thus, tele-consultation is arranged in order for the physicians to decide the most appropriate diagnostic procedures and therapy. The physicians at OPBG conduct their diagnosis, and they produce a medical report which is accompanied with medical images. This data is retrieved by the specialised group at HIUNJ via an OPBG web application. After the HIUNJ have received the necessary data, they are able to provide their assessment for the case, confirming whether the first diagnosis is correct or not.

4.3 Privacy Level Agreement Management in the case of the MoA

The process of the PLA management that is followed is the same in both case studies. Due to space limitation, in this paper we demonstrate, in the rest of this section, only the application of our work on the PA pilot, although we discuss in Section 5 the evaluation of the work on both cases.

To support the management of PLAs both at design and runtime levels, as discussed in Section 3, Data Controllers need to make use of appropriate tools and techniques to capture the organisational structure of their system and perform privacy, security and trust analysis. Although the selection of such tools and techniques depends on each organisation, the MoA pilot made use of tools developed as part of the VisiOn EU project (VisiOn Privacy Platform, 2016) platform.

Design time At the design level, three main activities are performed that form input to the PLA as described in Section 3. In particular, the Questionnaire Creation, the law compliance checks and the privacy and threat analysis. In the MoA pilot, to provide input to the PLA from the privacy and threat analysis, MoA Security Engineers modelled privacy requirements of their MACS system, using the STS modelling language (Salnitri, M., Paja, E., Poggianella, M. and Giorgini, P., 2015) of the VisiOn platform. Based on that analysis, business processes were designed, using the SecBPMN2 tool (Salnitri, M., Paja, E. and Giorgini, P., 2016), that are implemented by the MACS IS, capturing also any identified interaction with other systems. For example, the activity *Allow access* was modelled, in order to provide permission for reading the medical certificate and storing it in the local database, allowing citizens to access the fitness centre's IS.

Then, the MoA Security Engineers made use of the the SecTro tool (Mouratidis, H., Argyropoulos, N. and Shei, S., 2016) to enrich the STS-ml model with privacy and security requirements and relevant privacy objectives. As an example, through that analysis, it was identified that the MACS privacy constraint regarding access to medical certificate can be satisfied through the *undetectability* privacy objective.

Privacy mechanisms stored in a library of the tool were used to identify appropriate privacy solutions, such as smart cards and permission management. The results of this analysis were then used as input in the corresponding field of the PLA, indicating the specific privacy and security mechanisms that MoA has to use, assuring the citizen of the level of their privacy protection, thus increasing their trust in MoA e-services. Next, the MoA Security Engineers constructed a trust model using the JTrust tool (Pavlidis, M., Islam, S., Mouratidis, H. and Kearney, P., 2014) in order to identify dependencies among the system's stakeholders. In our case, the citizen depends on the IS of the fitness centre to keep their medical certificate confidential. This dependency implies a trust relationship between the citizen and the IS of the fitness centre and is justified with *Reported Trust*, i.e. MoA reports that the IS of the fitness centre can be trusted to keep the medical certificate confidential. It is worth noting that in cases where this analysis reports lack of trust, control mechanisms are added.

After that, Security Engineers added security annotations on the model of the system's architecture. The CARiSMA tool (Jürjens, J., 2002) is used for the design of the models by using UMLsec, and performs checks to validate that the architecture satisfies the security requirements. In our scenario, the transmission of a document by the MACS to the IS of the fitness centre was annotated with the security requirement of confidentiality. There, a transmission annotation was added to both actors, to verify that the checks performed by the tools are successful. This ensures the implementation of all the necessary security measures that guarantee the protection of citizens' data.

The last part of the analysis that MoA Security Engineers conducted was compliance with the relevant laws and regulations, at national and European level, to prove to citizens that they are compliant to the law. For the MoA pilot, the LIONoso tool (Battiti, R. and Brunato, M., 2014) was used as the law compliance checker to i) specify how MoA uses/manages citizen's data; ii) specify the constraints imposed by laws and regulations and iii) verify compliance of the data management specified by MoA with the constraints specified in the aforementioned specification of laws and regulations. MoA Security Engineers inserted, in a machine-readable format, the relevant national and EU privacy laws related to the privacy preservation of personal and sensitive data. Then, they specified the operations that are applied over this data, e.g., MoA produces birth certificate. Finally, MoA Security Engineers used the Data Value Tool (DVT) of the VisiOn platform to assess citizen's personal data. This value was then compared with both MoA expectations and citizen's perspective, and the results are visualised to the users through the PLA. When MoA Security Engineers completed the aforementioned analysis, they proceeded with the creation of an appropriate, privacy related, questionnaires so that the Data Subjects (in our case, citizens) using the MoA services could define their privacy preferences. For the MoA pilot, the Dynamic Audit Engine (DAE) tool of the VisiOn platform played the role of the questionnaire composer, assisting MoA in easily creating questionnaires. Once the citizen registered to the MoA e-service, they had to answer this questionnaire. The answers to those questionnaires formed the last input to the PLA.

Runtime The preferences of the citizen, as they have been recorded in the PLA, are parsed by the *Privacy Enforcer*. In this pilot, two tools were used, namely the Privacy Agreement Enforcer (PAE) and the Media Aware Network Element (MANE). PAE creates privacy policies which are formulated by the privacy preferences defined by citizens, after they filled in the questionnaire. Moreover, PAE evaluates requests for accessing private data against these policies, checking the policies that apply to that specific data and enforcing the results. MANE is responsible for monitoring and filtering network traffic, acting as a second layer of data protection by applying access rules according to the data received from PAE. In case that we have a request by a third party organisation to citizen's personal data, PAE receives this request and checks if this activity is allowed, according to the generated PLA (which contains citizen's privacy preferences). If the citizen has denied MoA the right to share their personal data e.g., for commercial purposes, MoA applies this policy. The result of PAE is then forwarded to MANE which

automatically stops future requests to this specific information. Also, a notification (e.g., via SMS) about the attempts to access their data and the corresponding results is sent to the citizen. In this way, the citizen is continuously updated, and this positively affects their trust in MoA e-services. Finally, the results of the actions of both PAE and MANE are forwarded to LIONoso, which updates the number of requests to citizen's data, thus formulating a history-based assessment percentage.

5 Evaluation

This section aims to evaluate the usefulness of the PLA and to provide insights about its applicability. This evaluation has been conducted under two perspectives. On the one hand, at a technical level, we assess PLA's usability and potential challenges that the end users came across during its management. On the other hand, from a social point of view, we evaluate the PLA management reference architecture regarding the impact, for both Data Controllers and Data Subjects, on the protection of their personal data and the level of trust of Data Subjects to Data Controllers during the provision of their data.

5.1 Evaluation Method

PLA management is achieved through the use of a series of tools, each one responsible for a specific part of it. More specifically, the tools that are necessary and contribute to the management of the PLA, during the design time i) conduct privacy analysis on Data Controllers' IS, by analysing threats, vulnerabilities, and trust relationships of their IS with other ones, ii) create the questionnaire that allows Data Subjects to provide/declare their privacy preferences, iii) allow the assessment of the value of Data Subjects' data, and, during runtime, they v) enforce and monitor Data Subjects' privacy preferences. The management of the PLA is accompanied with a novel, quite technologically challenging solution. Such an approach can be either accepted or rejected by the end-users, according to their behaviour towards the solution, after testing its use. For this reason, we decided to support the evaluation process based on the Technology Acceptance Model (TAM) (Davis, F.D., 1989; Venkatesh, V. and Davis, F.D., 2000). TAM approach is an information systems theory that models how users come to accept and use a technology. The model suggests that when users are presented with a new technology, a number of factors influence their decision on how and when they will use it. Notably, the approach introduces two determinants: i) Perceived Usefulness (PU), which is defined as "the degree to which a person believes that using a particular system would enhance their job performance", and ii) Perceived Ease-Of-Use (PEOU), which is defined as "the degree to which a person believes that using a particular system would be free from effort". Therefore, PU provides insight on how useful the technological approach behind PLA management is, while PEOU on how easy it is to use. This evaluation method aims to include parameters that will quantify end-users' perspective, on what extent the release of such technology is necessary and could cover their need. Furthermore, it will indicate the required effort that the end-users have to employ to use the technological solution that supports the management of the PLA, since users tend to adopt technologies that not only are effective but can also be apprehended in an effortless way.

5.2 Results

For the two pilots, the means of the evaluation was an online survey, which was designed under the TAM approach, and took place between February and June 2017. The questions aim to mainly confirm the increased perception of trust in online public services, both for the Data Controllers and for Data Subjects. Thus, the end-users report the level of privacy on their personal data, so they validate their willingness to share data while feeling secure and protected from the adoption of the PLA in the corresponding e-service. The first part of the questionnaire is common for both Data Controllers and Data Subjects, including additionally three demographic questions on sex, age, and education level. Next, there are questions solely for Data Controllers and solely for Data Subjects. The Public Administration questionnaire included 18 questions for PAs and 20 for citizens, while the Healthcare questionnaire included 13 questions for the administrators and 18 for patients. However, we present below, in Table 1 and Table 2, respectively, the questions that triggered significant statistical responses and results, and

may offer valuable findings for this study. All questions are given in multiple choice format, with answers possible in the following 5-point Likert scale; totally disagree, disagree, neither agree/nor disagree, agree, totally agree. Due to space limitations, we provide only the percentages of the three first responds.

Criterion	Question	The Respondents:
PEU, Supportability	Do you find that the process of the PLA management is designed for all levels of users, irrespective of their technical background?	Agree 47,22%, Neither agree/nor disagree 30,56%, Disagree 18,06%
PU, Reliability, Security	Is PLA management platform useful for providing innovative services to Data Subjects, as far as privacy is concerned?	Agree 62,50%, Totally agree 22,22%, Neither agree/nor disagree 15,28%
PU, Reliability, Threat, Information Control	Would PLA management platform improve privacy in public services provided to Data Subjects?	Agree 58,33%, Totally agree 23,61%, Neither agree/nor disagree 15,28%
PU, Reliability	Would PLA increase trust to public services?	Agree 51,39%, Totally agree 25,00% Neither agree/nor disagree 22,22%
PEOU, Perceived Ease of Learning	Are you confident that you can complete a process for the management of a PLA?	Agree 62,50%, Neither agree/nor disagree 22,22%, Totally agree 09,72%
Solely Administrators Questions		
PEOU, Supportability, Performance	Is PLA management platform easy to integrate with existing public services?	Agree 43,24%, Neither agree/nor disagree 43,24%, Disagree 10,81%
Platform requirements	Does PLA management platform communicate securely with other public bodies to execute the request?	Agree 54,05%, Neither agree/nor disagree 29,73%, Totally agree 13,51%
Solely Citizens Questions		
PU, Information Control	Would PLA make it easy to understand the actual values of our personal data?	Agree 51,43%, Neither agree/nor disagree 20,00%, Disagree 17,14%
Platform Requirements	Does PLA assist Data Subjects regarding the preferred level of privacy?	Agree 71,43%, Neither agree/nor disagree 20,00%, Totally agree 08,57%
PU	Would you use PLA for public services?	Agree 71,43%, Neither agree/nor disagree 17,14%, Disagree 08,57%

Table 1. Questionnaire in Public Administration domain

Criterion	Question	The Respondents:
PEOU, Perceived Ease of Learning	Is the PLA management platform easy to use?	Agree 49,50%, Neither agree/nor disagree 25,80%, Totally agree 23,70%
PU, Reliability, Threat, Information Control	Before the simulation, did you know the privacy and security aspects of health data?	Neutral 35,30%, Not really 34,20%, Somewhat 15,80%
PU, Information Control, Reliability	Do you think this process has raised your awareness of the privacy issues by providing you with greater awareness and understanding of the importance of protecting your data?	Agree 46,30%, Neither agree/nor disagree 18,90%, Totally agree 17,40%

PU, Supportability	Do you think that PLA management platform can guarantee greater privacy and security while exchanging health data?	Agree 44,20%, Totally agree 30,50%, Neither agree/nor disagree 15,80%
Solely Patients Questions		
PU, Information Control	Do you think PLA management platform allows patients greater control over privacy using constraints on health data transmission?	Agree 49,50 %, Totally agree 23,70%, Neither agree/nor disagree 14,70%
PU, Information Control, Risks	Do you think this process has made you aware of the potential risks or benefits of consciously compiling the consent for the transmission of health data?	Agree 56,30%, Totally agree 24,20%, Neither agree/nor disagree 12,60%
PEOU, Perceived Ease of Learning, Supportability	Do you think this process can be useful for patients with particular clinical needs or problems, such as reduced mobility?	Agree 42,10%, Totally agree 31,10%, Neither agree/nor disagree 18,40%
PU, Information Control	Do you think the PLA offers a complete insight on privacy and security issues?	Agree 46,80%, Totally agree 36,30%, Disagree 8,40%

Table 2. Questionnaire in Healthcare domain

Public Administration domain In this pilot, the sources for engaging participants – that had already been contacted and recruited during pilot preparation – derived from both private companies as well as in the public sector among the Municipality of Athens, other municipalities and bodies of local administration. The goal was to engage mainly participants that are occupied as Public Administrators but have as well a strong background in Information Technology e.g., Security Engineers, professionals in the IT Department of municipalities, System Administrators, etc. The distribution of participants was in majority males, however the difference with female participants is not so significant (31 out of 72 were female and 41 male). The most populated age-group was between 30-40 years (almost 57%) and 20% were from 40-50 years. Most participants have the age profile of an active working professional, which is also verified by the level of education, where the majority of participants, 19,44% and 63%, have undergraduate or postgraduate studies, respectively.

From a technical perspective, the questions that capture the ease of use of the PLA management platform and the process of the management of a PLA outcome interesting results. The majority of the participants agrees that they can easily complete the process of a PLA management. In addition, this process is designed for all the various users, irrespective of their technical background. These two findings indicate that the PLA management platform is well-designed and does not require possession of specialised technical skills. Of course, there is always room for improvement. PLA management platform is a novel technical approach and requires continuous improvement to be approved by its end-users. The question related to its integration and secure communication with existing systems highlights the interoperability issues that every new system faces during their integration with legacy systems. These difficulties have to be taken into account and individually examined for each system.

Focusing on the usefulness and the impact that the adoption of the PLA has, the findings are rather positive. The PLA management platform is considered quite useful for providing innovative services to individuals, with respect to privacy. Moreover, the adoption of the PLA contributes to the improvement of privacy issues related to public services. Next, the percentage of respondents (51,39%) who indicate the increase of trust to public services provides us the confidence that the adoption of the PLA is actually beneficial for Public Administrators, who increase their trustworthiness, attracting thus more citizens to use their e-services. The last years, public sector has put a lot of effort to promote transparency (Obama, B., 2011), making their processes understandable, accurate, and reusable. PLA can contribute towards the improvement of public sector trustworthiness, and consequently, towards the increase of citizens' trust. Finally, PLA offers the possibility of scalable provision of consent, allowing citizens to define

their privacy preferences, which can be updated, if necessary. This attribute has been assessed positively by individuals, who, hereafter have realised the importance of their data.

Healthcare domain This pilot includes 190 participants from two hospitals, 89 from OPBJ and 101 from HIUNJ. Nearly half of the participants have a job as office worker while nearly 25% of users are freelancers. More specifically, 22,6% are freelancers, 45,8% are office workers, 4,7% have retired, 10,5% are students, and 4,2% are unemployed. Cumulatively, there was a balance in the distribution, since 48,40% were male and 51,60% female. 43,20% of the participants are between 30-39 years old, percentage that again indicates their active working profile. However, there are significant percentages in higher ages (14,70% among 40-49 and 50-59), since the patients in these clinics are minors and are represented by their legal guardians. The participants' educational level is mainly on bachelor degree (57,90%) and upper secondary (36,30%).

Evaluating the management of the PLA from a technical point of view, we can see that the majority of the respondents find this process easy. Of course, this percentage should not be examined solely, but in combination with the high educational level of the majority of participants. The privacy analysis of an organisation's IS is a demanding process that requires a substantial background of knowledge. However, our participants were not familiar with security and privacy engineering methodologies. This parameter indicates that the management platform is user friendly and users who are not familiar with such tools and methodologies are able to use it successfully, delivering insight on privacy and security issues for their organisation.

Regarding the usefulness of the PLA and the impact that it has on the individuals, the results are more than positive. Despite that participants were not aware of the privacy and security aspects of health data (the majority of them (almost 70%)), they mentioned that through their involvement of this process, they were able to deeply understand the importance of protecting the data they provide in such an organisation. Moreover, they feel confident that the hospital can guarantee greater privacy and security while exchanging health data. We highlight here that the healthcare domain pilot contains many peculiarities, such as the provision of not only personal but also sensitive data, the emergency cases that often demand quick responds and decisions, the necessity of cross-border exchange of patients' data, and the representation of the patients by their legal guardians, which also reflects the principle of GDPR related to the parental consent. Thus, it is quite meaningful that with the use of PLA, the patients have realised the criticality behind the provision of their personal and sensitive data or their biometrics. Another important benefit from the adoption of the PLA is that it is useful for patients with particular clinical needs or difficulties. One of the major benefits of an online healthcare service is to simplify some operations for patients living in remote areas or who face particular clinical difficulties that could benefit the most from the use of a web platform, and the PLA can contribute towards this.

6 Conclusions

This paper has presented a novel reference architecture for the management of Privacy Level Agreements and a practical evaluation of this work. The evaluation has been conducted on the public administration domain and on the healthcare domain, where personal and sensitive data is being handled. Moreover, another peculiarity that these domains have is that citizens (and patients, respectively) do not have other option than providing their data in order to deliver the corresponding e-services. The results of the evaluation indicate that appropriate management of PLAs allows Data Controllers (public bodies, private organisations) to demonstrate that they have taken all the necessary actions to mitigate the identified threats, and also to demonstrate transparency with regards to processing and sharing of Data Subjects' data. On the other hand, PLA is a means of enhancing trust of Data Subjects to Data Controllers.

The presented work is part of a larger effort to develop a framework to support compliance with the GDPR. Towards this direction, we envisage future work to focus on the development of tools and methods to deal with aspects of the GDPR that PLAs are not supporting, such as privacy risks assessments, privacy complaints, and breach notifications.

References

- Ahmadian, A. S., Coerschulte, F., and Jürjens, J. (2015, July). Supporting the Security Certification and Privacy Level Agreements in the Context of Clouds. In *International Symposium on Business Modeling and Software Design*, Springer, Cham, 80-95.
- Battiti, R., and Brunato, M. (2014). The LION way. Machine Learning plus Intelligent Optimization. LIONlab, University of Trento, Italy, 94.
- Bouman, J., Trienekens, J., and Van der Zwan, M. (1999). Specification of service level agreements, clarifying concepts on the basis of practical research. In *Software Technology and Engineering Practice*, 1999. STEP'99. Proceedings (pp. 169-178). IEEE.
- Colombo, P. and Ferrari, E. (2012). Towards a modeling and analysis framework for privacy-aware systems. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing (SocialCom)*. IEEE, 81–90.
- CSA: Privacy level agreement outline for the sale of cloud services in the European Union. Tech. rep., Cloud Security Alliance, Privacy Level Agreement Working Group (February 2013)
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS quarterly*, 319-340.
- DErrico, M., and Pearson, S. (2015, March). Towards a formalised representation for the technical enforcement of privacy level agreements. In *Cloud Engineering (IC2E), 2015 IEEE International Conference*, IEEE, 422-427.
- Diamantopoulou, V., Angelopoulos, K., Flake, J., Praitano, A., Ruiz, J. F., Jürjens, J., Pavlidis, M., Bonutto, D., Sanz, A.C., Mouratidis, H., Robles, J. G., and Tozzi, A.E. (2017, June). Privacy data management and awareness for public administrations: a case study from the healthcare domain. In *Annual Privacy Forum*. Springer, Cham, 192-209.
- Diamantopoulou, V., Pavlidis, M., and Mouratidis, H. (2017). Privacy level agreements for public administration information systems. In *Proceedings of the CAiSE Forum 2017 29th International Conference on Advanced Information Systems Engineering*, X. Franh, J. Ralyté, R. Matulevičius, C. Sallinesi, and R. Wieringa, (eds), Essen, Germany, CEUR LNCS, 97-104.
- Drogkaris, P., Gritzalis, S., and Lambrinoudakis, C. (2013). Employing privacy policies and preferences in modern e-government environments. *International Journal of Electronic Governance*, 6(2), 101-116.
- European Commission: Directive 95/46/EC of the European Parliament and of the Council, URL: <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX> (visited on 23/05/2018).
- European Commission: Directive 2002/58/EC of the European Parliament and of the Council (July 2002), URL: http://ec.europa.eu/justice/data-protection/law/files/recast_20091219_en.pdf (visited on 23/05/2018).
- European Commission: Proposal for a regulation of the European Parliament and of the Council (January 2012), URL: <http://eur-lex.europa.eu/legal-content/en/ALL/?uri=CELEX:52012PC0011> (visited on 23/05/2018).
- European Parliament: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EU (General Data Protection Regulation) (2016), URL: <http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679&from=en31995L0046> (visited on 23/05/2018).
- Faßbender, S., Heisel, M., and Meis, R. (2014a, August). Functional requirements under security pressure. In *9th International Conference on Software Paradigm Trends (ICSOFT-PT)* IEEE, 5-16.
- Faßbender, S., Heisel, M., and Meis, R. (2014b, August). Problem-based security requirements elicitation and refinement with pressure. In *International Conference on Software Technologies*. Springer, Cham, 311-330
- García, J. M., Fernandez, P., Pedrinaci, C., Resinas, M., Cardoso, J., and Ruiz-Cortés, A. (2017). Modeling service level agreements with linked USDL agreement. *IEEE Transactions on Services Computing*, 10(1), 52-65.

- Jürjens, J. (2002, September). UMLsec: Extending UML for secure systems development. In *International Conference on The Unified Modeling Language*. Springer, Berlin, Heidelberg, 412-425.
- Kalloniatis, C., Kavakli, E. and Gritzalis, S. (2008). Addressing privacy requirements in system design: the PriS method. *Requirements Engineering* 13, 3, 241–255.
- Keller, A., and Ludwig, H. (2003). The WSLA framework: Specifying and monitoring service level agreements for web services. *Journal of Network and Systems Management*, 11(1), 57-81.
- Mead, N. R., and Stehney, T. (2005). Security quality requirements engineering (SQUARE) methodology (Vol. 30, No. 4, ACM, 1-7.
- Mohamed, M., Anya, O., Tata, S., Mandagere, N., Baracaldo, N., and Ludwig, H. (2017). rSLA: An Approach for Managing Service Level Agreements in Cloud Environments. *International Journal of Cooperative Information Systems*, 26(02), 1742003.
- Mouratidis, H., Argyropoulos, N., and Shei, S. (2016). Security requirements engineering for cloud computing: The secure tropos approach. In *Domain-specific conceptual modelling*. Springer, Cham. 357-380.
- Obama, B. (2011). Transparency and open government. *Presidential Memorandum* (January 21, 2012), <http://www.whitehouse.gov/the-press-office/transparency-and-open-government>.
- Pavlidis, M., Islam, S., Mouratidis, H., and Kearney, P. (2014). Modeling trust relationships for developing trustworthy information systems. *International Journal of Information System Modeling and Design (IJISMD)*, 5(1), 25-48.
- Salini, P., and Kanmani, S. (2013). Model oriented security requirements engineering (MOSRE) framework for web applications. In *Advances in Computing and Information Technology*, Springer, Berlin, Heidelberg, 341-353.
- Salnitri, M., Paja, E., and Giorgini, P. (2016, September). Maintaining secure business processes in light of socio-technical systems' evolution. In *Requirements Engineering Conference Workshops (REW), IEEE International*, IEEE, 155-164
- Salnitri, M., Paja, E., Poggianella, M., and Giorgini, P. (2015). STS-Tool 3.0: Maintaining Security in Socio-Technical Systems. In *CAiSE Forum*, 205-212.
- Sedek, K. A., Sulaiman, S., and Omar, M. A. (2011, December). A systematic literature review of interoperable architecture for e-government portals. In *Software Engineering (MySEC), 2011 5th Malaysian Conference*, IEEE, 82-87.
- Spiekermann, S. and Cranor, L. F. (2009). Engineering privacy. *IEEE Transactions on software engineering* 35, 1, 67–82.
- Tambouris, E., and Wimmer, M. (2008). Online one-stop government: a single point of access to public services. In *Electronic Government: Concepts, Methodologies, Tools, and Applications*, IGI Global, 2805-2829
- Venkatesh, V., and Davis, F. D. (2000). A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management science*, 46(2), 186-204.
- VisiOn: Vision privacy platform (2016), URL: <http://www.visioneuproject.eu/> (visited on 23/05/2018).
- W.W.W.C.: Platform for Privacy Preferences (P3P) Project (2016), URL: <https://www.w3.org/P3P/> (visited on 23/05/2018).