A4Cloud Workshop: Accountability in the Cloud

Carmen Fernandez-Gago, Siani Pearson, Michela D'errico, Rehab Alnemr,
Tobias Pulls, Anderson Oliveira

► To cite this version:

Carmen Fernandez-Gago, Siani Pearson, Michela D'errico, Rehab Alnemr, Tobias Pulls, et al.. A4Cloud Workshop: Accountability in the Cloud. David Aspinall; Jan Camenisch; Marit Hansen; Simone Fischer-Hübner; Charles Raab. Privacy and Identity Management. Time for a Revolution?: 10th IFIP WG 9.2, 9.5, 9.6/11.7, 11.4, 11.6/SIG 9.2.2 International Summer School, Edinburgh, UK, August 16-21, 2015, Revised Selected Papers, AICT-476, Springer International Publishing, pp.61-78, 2016, IFIP Advances in Information and Communication Technology, 978-3-319-41762-2. 10.1007/978-3-319-41763-9_5. hal-01619741

HAL Id: hal-01619741

<https://hal.inria.fr/hal-01619741>

Submitted on 19 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

A4Cloud Workshop: Accountability in the Cloud^{*}

Carmen Fernandez-Gago¹, Siani Pearson², Michela D’Errico², Rehab Alnemr²,
Tobias Pulls³, and Anderson Santana de Oliveira⁴

¹ Network, Information and Computer Security Lab
University of Malaga, 29071 Malaga, Spain
mcgago@lcc.uma.es

² Hewlett Packard Labs, Bristol, UK
{siani.pearson, michela.derrico, rehab.alnemr}@hpe.com

³ Karlstad University, Sweden
tobias.pulls@kau.se

⁴ SAP Labs France, France
anderson.santana.de.oliveira@sap.com

Abstract. As cloud computing becomes a widely used technology, it is essential to provide mechanisms and tools that enable trust about how personal data is dealt with by cloud providers. The Cloud Accountability (A4Cloud) project tries to solve the problem of ensuring trust in the cloud by providing tools that support the process of achieving accountability. In this paper we will concentrate on some specific tools that were demonstrated and discussed during the A4Cloud workshop held in association with the IFIP Privacy Summer School in Edinburgh in 2015. In particular, we will describe tools that facilitate the appropriate choice of a cloud provider such as the Cloud Offerings Advisory Tool (COAT) and the Data Protection Impact Assessment Tool (DPIAT), tools that are in charge of controlling the data of the users such as the Data Track (DT) tool, and tools that help specify and enforce accountability related policies by using the Accountability-Primelife Policy Language (A-PPL) and an associated enforcement engine.

Keywords: Accountability, Tools, Control Tools, Facilitating Tools

1 Introduction

Cloud computing technology is becoming more and more popular and it is being widely used by companies and users nowadays. However, there are still some concerns about security and how personal data is dealt with by cloud providers. There should be mechanisms and tools in place that help users to have trust in the cloud. The goal of the A4Cloud project [1] is to provide tools and mechanisms that help achieve *accountability* for cloud providers, including demonstration that they are accountable, and helping users to know whether the cloud provider of their choice is accountable. Accountability

^{*} This work has been partially funded by the European Commission through the FP7/2007-2013 project A4Cloud under grant agreement number 317550. The first author is supported by the Ministry of Economy of Spain through the Young Researchers Programme: project PRECISE (TIN2014-54427-JIN). The authors would like to thank Saul Formoso for taking notes during the workshop and all the members of A4Cloud involved in the development of the tools.

consists of defining governance to comply in a responsible manner with internal and external criteria, ensuring implementation of appropriate actions, explaining and justifying those actions and remedying any failure to act properly [2]. The approach followed by A4Cloud is interdisciplinary and includes legal and regulatory, socio-economic and technical aspects.

This paper describes the outcomes of the workshop held during the IFIP summer school 2015 on the topic of accountability for data protection in the cloud. At the time of this workshop some of the A4Cloud prototypes were in a mature enough state to be demonstrated. We chose different kinds of such prototypes to be shown at the workshop. The chosen tools offered complementary functionalities that will be also described in this paper. On the one hand, we concentrated on the Cloud Offering Advisory Tool (COAT) and Data Protection Impact Assessment Tool (DPIAT) that facilitate users in finding and assessing a suitable cloud provider that may fulfill their needs. On the other hand, we also considered tools that aid users having control over their data, such as the Data Track (DT) tool or the Accountability-PrimeLife Policy Language (A-PPL), that is the reference language for representation of accountability policies. These policies can be defined by using the Data Protection Policies Tool (DPPT) and enforced by the associated engine.

The structure of the paper is as follows. Section 2 provides an overview of the A4Cloud project, whereas the following sections provide insights on specific tools demonstrated within the workshop. Thus, Section 3 describes tools for facilitating users' choice of providers: in particular, COAT and DPIAT. Section 4 describes some control and transparency tools, namely DPPT and A-PPL Engine. The feedback given by the participants in the workshop is analysed in Section 5. Finally, Section 6 concludes the paper.

2 Overview of the Accountability Project

The goal of the A4Cloud project is to provide an increased basis for trustworthiness in cloud computing by devising methods and tools, through which cloud stakeholders can be made accountable for the privacy and confidentiality of personal data held in the cloud. These methods and tools combine risk analysis, policy enforcement, monitoring and compliance auditing. They will contribute to the governance of cloud activities, providing transparency and assisting legal, regulatory and socio-economic policy enforcement. The A4Cloud tools are grouped into different categories depending on their functionality or their intended usage. They are classified as follows [3]:

- *Preventive tools* are those that aid mitigation of the consequences of any unauthorised action. These tools include assessing risk, selection of providers or identification of appropriate policies to mitigate the risks. These tools are the Data Protection Impact Assessment Tool (DPIAT), Cloud Offerings Advisory Tool (COAT), Accountability Lab (AccLab), the Data Protection Policies Tool (DPPT), Accountable Primelife Policy Engine (A-PPL Engine) and Assertion Tool (AT).
- *Detective tools* are those that monitor for and detect policy violations. The corresponding A4Cloud tools are the Audit Agent System (AAS), Data Transfer Monitoring Tool (DTMT), Data Track (DT) and Transparency Log (TL).

- *Corrective tools* are those designed to mitigate the consequences of any incident occurring. They are the Incident Management (IMT) and the Remediation Tool (RT).

Among the preventive tools we are going to concentrate in this paper on those designed for facilitating the choice of a cloud provider. They are DPIAT and COAT. We will be also concentrating on the definition of policies and enforcement and will pay special attention to the DPPT tool to define policies and A-PPL engine to enforce and handle policies specified in A-PPL and created by DPPT. Among the detective tools we consider in this paper the DT tool. This tool can be classified as a Data Subject Control tool that provides controls for the management and protection of the personal data of the users.

For the scope of this work no corrective tools were included in the workshop as they were not mature enough to be demonstrated.

3 Tools for Facilitating Choice

3.1 COAT

The Cloud Offerings Advisory Tool (COAT) is a cloud brokerage tool that facilitates evaluation of cloud offerings and contract terms with the goal of enabling more educated decision making about which service and service provider to select. It allows potential cloud customers – with a focus on end users and Small and Medium Sized Enterprises (SMEs) – to make informed choices about data protection, privacy, compliance and governance, based upon making the cloud contracts more transparent to these cloud customers. This type of tool is useful for a number of reasons: reading and interpreting terms and conditions in cloud service offers can be challenging, and indeed non legal experts cannot easily make out the differences between contracts relating to cloud offerings. SMEs and individuals in particular do not typically have enough technical and financial resources to assess cloud offers in this way. However, there is no current brokering service that focuses on the user's data protection requirements.

A number of related factors vary across cloud providers, and are reflected in the contracts. For example, security measures (including the type of encryption and key management solution), processing locations (that determine which law applies), data protection roles (which reflect the capacity of the cloud service provider acting in relation to the personal data), data deletion procedures, notification about changes to terms and conditions, the ways in which subcontracting to a third party is allowed, notification in the event of a security breach involving the customer's data, notification when law enforcement requests a customer's data, the period during which the CSP keeps a customer's data after service termination, and so on. The tool highlights these differences, and explains to the user what implications this might have. The focus of the tool is on providing feedback and advice related to properties that reflect compliance with regulatory obligations rather than providing feedback on qualitative performance aspects (such as availability), although potentially the tool could be integrated with other tools that offer the latter.

From the cloud customer point of view, the tool aims to ease the comparison of alternative cloud offerings, and provide greater transparency and guidance when considering which provider to use. From the cloud service provider point of view, the tool

can provide benefits in terms of decreasing the complexity for customers when choosing a cloud provider, highlighting the unique criteria in the cloud service provider's offer, increasing market exposure, and ultimately matching cloud demands with their offerings.

Fig. 1. COAT High Level Architecture

The overall system is depicted in Figure 1, in which a Web User Interface enables interaction with the target users. During this interaction, potential cloud customers can provide as input to the graphical interface a collection of answers to a questionnaire (that can change according to the inputs already provided), but most of this information is optional and need not be provided although the interactions help guide the users as to their needs and provide a more targeted output. Such information includes the data location, the roles involved in the scenario to be built on the cloud, contact details of those responsible for defining the purpose of use for the involved data, contextual information about the environment setting and the user needs and requirements. Other knowledge used by the system includes the cloud service offerings in structured form, models of cloud contracts and points of attention and reputation information with respect to the agents involved in the offering process. During this process of interaction, guidance is provided on privacy and security aspects to pay attention to when comparing the terms of cloud service offerings. The outcome of COAT is an immediate and dynamically changeable feedback panel that includes an overview of compatible service offerings matching the user requirements and links to further information and analysis. A familiar store-type interface is used to reduce complexity and improve usability. See Figure 2 for an example, which shows the questionnaire on the left hand side, the results on the right hand side and an example of advice that is provided in a dynamic way as the user considers the answer to individual questions. The matching is based on data protection requirements, and transparency is improved about what exactly the service providers are offering from a data protection perspective and how those choices might impact the potential cloud customer.

Ongoing research involves usage of ontologies for more sophisticated reasoning and linkage to PLA terms, and usage of maturity and reputational models to optimise ordering of the outputs. For further information about the system, see [4].

Advisor

Business Questionnaire

Please indicate your requirements

Price Range

From: € 0 To: € 5000

Acceptable Storage Locations including Backup

☐ Europe (EU)
☐ United States
☐ Europe (Non-EU)
☐ China
☐ Local
☐ Any

Acceptable Data processor location

☐ Europe (EU)
☐ United States
☐ Europe (Non-EU)
☐ China
☐ Local
☐ Any

Data transfer in case of emergency? ⓘ

☐ Yes
☐ No
☐ Doesn't Matter

Do you want Encryption?

☐ Yes
☐ No
☐ Doesn't Matter

Is it important that any disputes are resolved in your own country??

☐ Yes
☐ No

Should unlimited backup be included? ⓘ

☐ Yes
☐ No

Notified in case of security breach? ⓘ

☐ Yes
☐ No

8 Matched Offers

Cirrus Thinking

€10.00/Month

Storage Location: United Kingdom

Processor Location: United Kingdom

Client-side encryption: Yes

More info Go to offer

Cloud Corner

€50.00/Month

Storage Location: Europe (EU)

Processor Location: Europe (EU)

Client-side encryption: Yes

More info Go to offer

Dropbox

€7.50/Month

Storage Location: United States

Processor Location: United States

Client-side encryption: No

More info Go to offer

Jottacloud

€6.00/Month

Storage Location: Europe (EU)

Processor Location: Europe (EU)

Client-side encryption: Yes

More info Go to offer

Acceptable Data processor location

This question concerns where personal data is processed and what laws apply to protect it. Personal data is data that relates to identifiable people. In countries within the EU, the data protection laws are similar so transferring and processing data within the EU is treated on the same basis as if you process data locally. Processing data is very wide and it means carrying out any operation or set of operations on the information or data (for example organisation, retrieval, consultation, deletion or use of the information or data).

In countries outside the EU, data protection laws are different. You should not transfer personal data outside the EU without checking whether this data will be adequately protected. This may involve getting contractual guarantees from your Service Provider that this data will be protected. If this data is not adequately protected, you may be in breach of local data protection law.

Fig. 2. Example COAT Screenshot

3.2 DPIAT

The Data Protection Impact Assessment Tool (DPIAT) is a decision support tool focusing on assessment of the risks associated with the proposed usage of cloud computing, involving personal and/or confidential data. It assesses the proposed use of cloud services, helping users to understand, assess, and select CSPs that offer acceptable standards in terms of data protection. The tool is tailored to satisfy the needs of SMEs that intend to process personal data in the cloud; it guides them through the impact assessment and educates them about personal data protection risks, taking into account specific cloud risk scenarios. The approach is based on legal and socio-economic analysis of privacy issues for cloud deployments and takes into consideration the new requirements put forward in the proposed European Union (EU) General Data Protection Regulation (GDPR) [5], which introduces a new obligation on data controllers and/or processors to carry out a Data Protection Impact Assessment prior to risky processing operations (although the requirements differ slightly across the various drafts of the Regulation).

Figure 3 shows the high level approach of DPIAT. The assessment is based on input about the business context gathered within successive questionnaires for an initial screening and for a full screening for a given project, combined with risk assessment of cloud offerings [6] based upon information generated voluntarily by CSPs, and collected from the CSA Security, Trust and Assurance Registry (STAR) [7]. The output

Fig. 3. The High Level Approach of the Data Protection Impact Assessment Tool

of the first phase of the DPIAT reflects advice about whether to proceed to the second phase of assessment. The second phase questionnaire contains a set of 50 questions. An example of part of this input questionnaire, which is responsive to the questions asked in the sense that the next questions to be asked depend partly upon previous answers, is given in Figure 4, which also illustrates how guidance related to the particular answers selected is provided dynamically during the process of answering the questions. The output of this phase is a report that includes: the data protection risk profile, assistance in deciding whether to proceed or not, and the context of usage of this tool within a wider DPIA process. Amongst other things, the tool is able to demonstrate the effectiveness and appropriateness of the implemented practices of a cloud provider helping him to target resources in the most efficient manner to reduce risks. The report from this phase contains three sections. The first, project-based risk assessment, is based on the answers to the questionnaire and contains the risk level associated with: sensitivity, compliance, transborder data flow, transparency, data control, security, and data sharing. An example is shown in Figure 5. The focus here is on assessment of potential harm to individuals and society, which is not typically part of a classic organisational risk assessment process. The second part of the report displays risks associated with the security controls used by the selected CSP. It contains the 35 ENISA risk categories [8] with their associated quantitative and qualitative assessments. The last section highlights additional information that the user needs to know related to requirements associated with GDPR article 33 [5]. The system also logs the offered advice and the users decision for accountability purposes. For further information about the system, including a full list of questions, see [9].

4 Control and Transparency Tools

4.1 Data Track

The Data Track (DT) is a tool that enables data subjects to get an overview of what information they have disclosed to online services [10]. DT provides several different

12: Is the nature of your operations such that you need to comply with rules regarding data processing in more than one set of regulations?

Think for instance specific (data protection) regulation pertaining to you, such as for financial or health services.

☒ Yes
☐ No

13: Are decisions being made on the basis of the information you process?

For instance, information can be collected for historical purposes without being used as part of a decision process.

☒ Yes
☐ No

14: Do the outcomes of these decisions have a direct effect on the individuals whose information is processed?

For instance, are offers based on the characteristics of individuals being collected by your system?

☒ Yes
☐ No

Info

The more rules you have to observe, the higher the likelihood that you breach one these.

Info

The mere collection of information is of different significance than the use of information in decision-making processes.

Info

When the information you handle leads directly to decisions that can affect individuals, the impact of processing is likely to be greater than the one it would have if the processing activities did not have any direct consequence on the individual the information relates to

Fig. 4. Example Section of DPIAT Full Screening Questionnaire

A4 Cloud Data Protection Impact Assessment Tool

Questionnaire Results (selected Cloud Service Provider: [DataSpacer](#))

HIGH Risk Related to Your Proposed Application		
Sensitivity	MEDIUM	Risks related to a sensitive market (i.e. elderly, children, etc.) and/or sensitive data (i.e. health or medical conditions, finance, sexual behaviour)
Compliance	HIGH	Risks related to compliance with external standards, policies, laws, etc.
Trans-Border Data Flow	LOW	Risks related to transfer of information across national borders
Transparency	HIGH	Risks related to transparency in the areas of notice/user messaging and choice/consent
Data Control	HIGH	Risks related to control of the data lifecycle (i.e., collection, usage, quality, and/or retention)
Security	HIGH	Risks related to security of data and data flows
Data Sharing	MEDIUM	Risks related to sharing data with third parties

Risk Related to the selected Cloud Service Provider

Usage of this Report within a Broader Data Protection Impact Assessment (DPIA) Process

Fig. 5. Example DPIAT Report, with First Section Expanded

views on the disclosed data – each view tailored to help data subjects answer different questions – and enables the data subject to exercise some control over the data stored on the service’s side. Data disclosures are collected by DT with the help of compatible service providers. Next, we present one of the views followed by what type of control DT provides for compatible service providers.

Trace View The Trace View is a view for the Data Track tailored for answering the question “What information about me have I sent to which online services?”. Figure 6 shows the Trace View. In the middle of the view there is a picture representing the user him- or her-self. At the top of the view are *attributes*: each attribute has a *type* (like e-mail or username) and one or more *values* (like “alice@example.com” or “bob”). At the bottom of the view are *services* like Spotify or Facebook. The view is called the Trace View because when the user clicks on a service there is a trace (line) drawn from

the service to the user, and from the user to all the attributes the user has sent to the service in question. Furthermore, there are lines drawn from the service to any other *downstream* services used by the service provider to provide the service. In Figure 6, Spotify is as an example using Visa (presumably to accept creditcard payments for subscriptions).

Fig. 6. The Trace View of the Data Track.

Exercising Control For compatible services, the Data Track enables the user to request access to, correction of, or deletion of their personal data stored remotely at services. In A4Cloud, a DT-compatible service is provided by the A-PPL Engine (presented in Section 4.2). Outside A4Cloud, the vast majority of conventional services do not provide data subjects with online APIs to exercise control over their personal data.

4.2 Privacy Policy Definition and Enforcement

Implementation of privacy policies can be thought of as a two phase process, as shown in Figure 7. Tasks carried out in the first phase have the objective to formally specify the policies that describe the different aspects of the practices related to the cloud service provision. During this phase different types of artifacts are designed and produced. For privacy policy statements that can be technically enforced by means of software tools, appropriate technical representation of the statements need to be generated. The specific language for the technical representation of the policies is tied to specific components tasked with the policy enforcement. The A4Cloud Project has designed i) a language, called A-PPL, whose expression capabilities address the need for accountability related policies; ii) an engine, the A-PPL Engine, which is able to translate into actions policies

Fig. 7. Privacy Policy Definition and Enforcement

specified in A-PPL. During the second phase of the policy implementation process, policies are technically enforced by the engine, which also provide features, such as logging and notifications, that contribute to building an accountable enforcement environment. In the following subsections we will describe the tools used in the policy definition and the policy enforcement phases, respectively DPPT and the A-PPL Engine, along with the artifacts produced and used.

Policy Definition: DPPT The Data Protection Policies Tool (DPPT) is a GUI-based tool that cloud service providers (CSPs) can use to specify privacy policy statements to produce two different artifacts: an ontological representation of the policy statements (OWL file); an A-PPL (basically, an xml-based) policy file, that contains a representation in A-PPL of the policy statements that need to be enforced through the A-PPL Engine. The OWL Privacy Policy is linked to the A-PPL policy, in that a subset of the statements will have a corresponding A-PPL policy element.

The DPPT GUI, shown in Figure 8, presents a set of data protection aspects, organised in graphical panes. For each data protection aspect the CSP’s privacy officer can specify a set of statements. When all the choices have been made, an A-PPL policy file can be created. The creation of the different policy elements involves using the data provided through the GUI to customise an appropriate policy element template. In order to bind a privacy policy statement with an A-PPL element template we have carried out an analysis of A-PPL Engine capabilities and mapped those to the data protection statements that can be fulfilled by enforcing that A-PPL element. The different A-PPL elements produced as a result of the templates customisation are then composed according to A-PPL schema to create a valid A-PPL policy file. Once created, the policy can be sent to the A-PPL Engine by clicking the button ”Send to Engine”. The result of the action is that the policy is stored by the A-PPL Engine, which is then ready to enforce them. We clarify that the actual enforcement will be done once data subjects’ personal data will be stored and bound to the policy. This is done at service registration phase, when data subjects are asked to provide data in order to start using the cloud service.

In the following, we will provide details about the options that can be selected to create a set of enforceable privacy policy elements.

Data Processing A CSP can specify the personal data elements that the service will process. In our A4Cloud reference scenario, the privacy officer will select data items be-

Fig. 8. DPPT GUI: configuration of notification sending

longing to two different personal data categories: “User Account” and “Health Data”. “User account” includes typical elements such as email address, username. “Health Data” refers to data elements that are processed by CSPs providing health statistics services. Examples of this category are blood pressure, blood sugar level, heart rate. Personal data categories are bound to ontological classes that describe them as belonging to specific data sensitivity categories. For example, the “Health Data” class is classified as “Sensitive Data”, according to the definition of sensitive in EU privacy laws [11]. The set of data items selected will be referenced in other policy statements to specify additional policies that apply to them.

CSPs need also to specify for how long data will be kept. This can be done by providing the time period value through the GUI. The A-PPL element generated for this statement will enable the engine to automatically manage the deletion of data.

Notification *Notification* is an essential element of accountability and should be used by CSPs to notify (within a predefined timeframe) about relevant events related to data processing operations. Depending on the types of events that can be captured by the components deployed for audit and monitoring purposes, sending of notifications can be automated. DPPT GUI allows the provider to configure the A-PPL Engine for notifications sending. The GUI presents a set of options for the type of events, which reflect the capabilities of the audit and monitoring tools developed within A4Cloud. Examples of event types are: data personal data deletion, data transfer policy violation, access denied. Specific type of events that affect personal data, such as data leakage or data loss, assume particular relevance and are specified in a separate section of the GUI. The information required to be provided for the generation of the corresponding A-PPL policy

element includes: the type of event to notify about, the relevant stakeholder identifier to be notified, the communication channel (basically e-mail or post), the address and the timeframe.

Control Rights Access rights (read, write, delete, update) need to be granted to different actors over different sets of personal data items. The GUI allows the provider to create different access rules for each actor that needs to process personal data collected from data subjects. Data Subjects need to be given access to their own personal data, therefore the CSP should create a specific rule that specifies that a given data subject can access her own data to perform all the available operations. The creation of the A-PPL policy elements for enforcement of access control requires the CSP to provide information such as the set of personal data, the identifier of the actor that should be given access, its role (relevant if the underlying access control system is role-based) and the access rights to grant.

Policy Enforcement: A-PPL Engine Organizational cloud customers usually assume the role of data controller, thus they are held accountable for the way cloud services respond to many regulations, including the EU Data Protection Directive [11]. Appropriate policies mitigate risks to data protection as they clarify in which way obligations regarding personal data are carried out. In particular, machine-readable policy languages, such as A-PPL [12, 13] make organizations accountable, ensure that obligations to protect personal data and data subjects' rights are fulfilled by all who store and process the data, irrespective of where that processing occurs.

In the cloud, mechanisms to automatically enforce organizational and data governance policies are fundamental for compliance management. Using the A4Cloud policy enforcement tools, cloud providers can offer more transparency about the data handling. As far as assurance is given about the deployment and configuration of the tools, which in general can be achieved with independent audits, the policy enforcement will happen in a predictable manner, satisfying the data controller needs and obligations, but also as determined by the data controller, giving back control to the (cloud) data subject. The enforcement engine works in cooperation with further A4Cloud tools to reinforce the assurance about the correct policy execution.

A-PPL Engine High Level Architecture The A-PPL engine has been created as an extension of the PPL engine in order to enforce the new A-PPL language policies. The first step was to extend the PPL XML schema with the new obligations defined in A-PPL. Many PPL engine functionalities have been simplified since the notion of sticky policy is no longer used.

The new A-PPL engine is a simplified version of the PPL engine. Many PPL engine functionalities have been removed since they are no longer needed in A-PPL: the matching engine is no longer used since the notions of sticky policy and data handling preferences do not exist in the A-PPL language.

A new module called **Policy Administration Point (PAP)** is responsible for storing, updating and deleting PII's and their policies in the database.

The access control is performed in the Policy Decision point (PDP) module which uses the Heras XACML engine [14] exactly the way it was implemented in PPL. The usage control is enforced in the obligation handler module which includes the event and action handlers sub-modules.

The Policy Enforcement Point (PEP) is the main module of the engine, which enforces the A-PPL policy of a PII when access is requested to it. For this reason, it communicates with the frontend of the engine (ppl-rest module) as well as the PDP and obligation handler

Upon storing the PII, the engine receives also an A-PPL policy that contains the rules and obligations related to how that piece of personal data has to be handled. For this reason, PII and their associated policy are represented in A-PPLE with the *<PIIType>* element, which has the following attributes:

- *owner*: Denotes the owner of the PIIs
- *attributeName*: It is the name of the PII element.
- *attributeValue*: The value of the PII element
- *policySetOrPolicy*: One or more “sticky” policies describing the access and usage control rights for this PII.
- *creationDate*: The creation date of PII.
- *modificationDate*: The modification date of PII.

When PII is retrieved from the engine, the policy associated with this personal data is analysed by the **Policy Decision Point (PDP)** that takes into account access and usage control rules.

In addition, obligations related to the personal data handling are executed when certain events occur inside the engine. The **Obligation Handler** component is responsible for analysing these A-PPL obligations (pairs of triggers and actions).

All actions happening inside A-PPLE are logged by a central component called **Logging Handler**. It stores the log entries related to decisions and actions (e.g. Action-Log) concerning the PIIs stored inside the engine.

We can see in Figure 9, that A-PPLE adopts a two-layer high-level architecture so that isolation between the engine components and data is performed: The core elements responsible for the policy enforcement functionalities reside in the Business layer. The Persistence layer consists of the PII Store and the Policy Repository where PIIs and the associated sticky policies are stored as well as the Logs produced by the Logging Handler component. Access to the persistence layer is achieved through a Persistence Handler component, which abstracts the underlying location and storage data model to the above business layer functions [15].

5 Discussion

In this section we will describe how we structure the workshop and the the feedback and questions that we received from the audience. Firstly, we introduced a general overview of the A4Cloud project in order to give the audience a general view on it. Then, the group was split into four different subgroups of four or five people each. Each of these

Fig. 9. A-PPL Engine architecture

subgroups approached specific booths for each of the tools (there was a shared booth for DPIAT and COAT), where a member of A4Cloud explained their corresponding tool. They introduced the group to their main features and showed how the tools work in practice. The presentation was interactive and anyone could ask questions at any time.

5.1 DT

Since the nature of the DT tool is mainly how data is dealt with, some concerns were related to mainly how it is stored or visualised, or how to input the information into the tool. The audience was also concerned about the resolution of conflicts or the evaluation of this tool as well as about monitoring aspects. These aspects led to the question of how this tool improves the accountability of the service provider. In general, attendees viewed the tool as an interesting one for auditing companies.

They also pointed out that security aspects should be considered in order to move forward. For the sake of usability, it was mentioned that it would be interesting if the DT tool worked in a mobile environment.

5.2 DPIAT

The comments and concerns raised by the audience with respect to the DPIAT are manifold. They relate to different aspects such as the selection of providers, how risk is considered and how the rating and the evaluation are performed.

Concerning the selection of the providers, some issues arose with respect to the way that the overall flags compute. The main question is whether the algorithm needs tweaking. It was suggested to be convenient for the user if there would be an option to select more than one CSP at the beginning (perhaps, up to three?). Then, on the final output screen it would be a good idea to have the different options selected shown up in the second band of output. This way the user does not have to go through the whole

questionnaire each time if they were not quite sure about which CSPs to use but had a good idea. This would also help highlight the differences between them.

This functionality could indeed be added fairly easily to allow comparison of several providers at once, but there was a difference of opinion amongst the attendees as to whether this would be a good idea, or appropriate for typical contexts in which the tool might be used. Also, it would be interesting to clarify in which category governments fall as they could also be users of the tool.

Attendees' concerns about risk were related to how it is presented for users to understand it. The distinction between the risks to data subjects (the core of a DPIA), and those related to a more classic organisational risk assessment (on the proposed service provider) should be clear. Also, the attendees suggested that the colour code and used signs for the DPIAT overall output, e.g. a high risk should show up as red, not green (this seemed due to a glitch with the system on the day). It was noted that risk mitigation is missing and also other factors that might be relevant for assessing risk, e.g. who owns the key?

There was a difference in opinion about how helpful an overall rating would be for the first section. Perhaps, the consensus was that it would be helpful indeed for decision making, but one would have to be very careful how that was done. It is important that there would be some explanation about what this overall rating would mean, and what the implications would be. It was highly stressed by the more mature researchers and experts that it should be clear how the user should interpret the findings, and what it meant for them in terms of what they should do next, etc.

It was advised to check whether or not the very high and extremely high, etc. of the second part were indeed the same bands (and associated meaning) as those used in the first part - it seemed that there might be a mismatch with different terms being used.

Perhaps it needed to be clearer how DPIAT should be used in the context of an overall DPIA process, and how things like consultation of stakeholders fit in. We do have some information on this, but it still seems not so clear about what else must be done, and when and how often the DPIAT tool could be used. Questions raised related to how and who assess cloud providers. An important issue is also time consuming for running the tools. It would be very interesting to run all the tools at once.

The major concern that attendees had was that there was no justification of why the results were given, or explanation of how the results were obtained, or way of the user or a third party to check the verification of the process used to obtain these results. One aspect of dealing with this is that an explanation of how the results were obtained should be added within the tool in a way that it could be accessed by a user exactly, even if the particular chain of processing leading to that particular result might not be exposable in other words, to be able to click to see how it was done, and provide more evidence about that. Which results from the system can be relied on, and why? There should be more explanation about how the risk was obtained, to what extent probability and impact are taken into account, etc., and whether that differs across the different sections of output. In what way is the output to be regarded as a proof for interpretation and used by the user, and in what way could it be provided as a proof for a data protection authority? The legal implications are different. What is the outcome of the result (single answer, guidelines, etc.)? What is it useful for?

5.3 COAT

Some options in the answers given seem very odd, e.g. in the encryption options there should be a distinction made between data at rest and during transit, and also the types of weaker and stronger encryption should be brought out more accurately. We had already realised this and updated the questionnaire offline, although that was not yet reflected within the demo version. Instead of just having the option of Europe in some of the questions, there should be the possibility (by another question popping up to bring up certain countries if Europe were selected) to allow more specific selection of countries (i.e. local is not enough, it may need to be local and Belgium, or whatever). This is necessary for various reasons (legal and user preferences).

There were also issues relating to what we have been showing for some time in this same demo versus our proposed improvements that were not yet reflected within it, such as more sophisticated ordering of results, display of certification/seals, etc. Questions from the audience were related to the use of standards, in the sense of how everything related to the tool would be documented.

Concerning legal regulations it was pointed out that different regulations depending on the countries should be taken into account.

5.4 A-PPL

One of the first questions that arose from the audience was why the A-PPL language was chosen. If the reason is a question of scalability this seems to be a good reason. The way requirements are dealt with is also another of the questions: are all the requirements translated into the policy? The A-PPL tool is based on the existence of a privacy officer. We have to check what is going to be done in the case where there is not such an officer, as it could be the case for example of SMEs.

6 Conclusion and Future Work

In this paper we have described the tools and the feedback provided by attendees during the A4Cloud workshop held at the IFIP summer school 2015. The A4Cloud project tackles the problem of providing accountability from an interdisciplinary perspective involving technical, legal and socio-economic aspects. A4Cloud provides a set of tools and mechanisms for the different stakeholders involved in the cloud for achieving accountability. Thus, in this paper, we have concentrated on tools for facilitating choice (COAT and DPIAT) and tools for offering control (DT, DPPT and A-PPL Engine). At the time of the workshop all these tool prototypes were in a mature enough state to be demonstrated. Therefore, the attendees could have a very good idea of how they could be used by users or providers.

The feedback from this workshop is being taken into account as the tool prototypes are further modified. Many of the issues raised have already been addressed within improved versions of the tools. We are carrying out further validation of the tools, for example, within a wearables scenario, as well as their socio-economic impact assessment.

References

1. A4Cloud: The Cloud Accountability Project. <http://www.a4cloud.eu/>
2. Catteddu, D., Felici, M., Hogben, G., Holcroft, A., Kosta, E., Leenes, R., Millard, C., Niezen, M., Nuñez, D., Papanikolaou, N., et al.: Towards a model of accountability for cloud computing services. In: Pre-Proceedings of International Workshop on Trustworthiness, Accountability and Forensics in the Cloud (TAFC). (2013)
3. Pearson, S., Wainwright, N.: An interdisciplinary approach to accountability for future internet service provision. *Proc. IJTMCC* **1** (2013) 52–72
4. Alnemr, R., Pearson, S., Leenes, R., Mhungu, R.: Coat: Cloud offerings advisory tool. In: IEEE 6th International Conference on Cloud Computing Technology and Science, Cloud-Com 2014, Singapore, 15-18 December. (2014)
5. EU Parliament and EU Council: Proposal for a regulation of the european parliament and of the council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (general data protection regulation) (2012)
6. Cayirci, E., Garaga, de Oliveira A., S., Roudier, Y.: A cloud adoption risk assessment model. In: IEEE Utility and Cloud Computing (UCC). (2014) 908–913
7. Cloud Security Alliance: Security & Assurance Registry (STAR). <https://cloudsecurityalliance.org/star/>
8. ENISA: Cloud computing - benefits, risks and recommendations for information security (2009)
9. Alnemr, R., Cayirci, E., Corte, L.D., Garaga, A., Leenes, R., Mhungu, R., Pearson, S., Reed, C., de Oliveira, A.S., Stefanatou, D., Tetrimida, K., Vranaki, A.: A data protection impact assessment methodology for cloud. In: Annual Privacy Forum (APF). LNCS, Springer (2015) 908–913
10. Angulo, J., Fischer-Hübner, S., Pulls, T., Wästlund, E.: Usable transparency with the data track: A tool for visualizing data disclosures. In Begole, B., Kim, J., Inkpen, K., Woo, W., eds.: Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, Seoul, CHI 2015 Extended Abstracts, Republic of Korea, April 18 - 23, 2015, ACM (2015) 1803–1808
11. European Parliament and the Council of the European Union: Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data . http://ec.europa.eu/justice/policies/privacy/docs/95-46-ce/dir1995-46_part1_en.pdf (1995)
12. Azraoui, M., Elkhiaoui, K., Önen, M., Bernsmed, K., Santana De Oliveira, A., Sendor, J.: A-PPL: An accountability policy language. In: DPM 2014, 9th International Workshop on Data Privacy Management, Wroclaw, POLAND (September 2014)
13. Benghabrit, W., Grall, H., Royer, J.C., Sellami, M., Azraoui, M., Elkhiaoui, K., Önen, M., de Oliveira, A.S., Bernsmed, K.: A cloud accountability policy representation framework. In Helfert, M., Desprez, F., Ferguson, D., Leymann, F., Muñoz, V.M., eds.: CLOSER 2014 - Proceedings of the 4th International Conference on Cloud Computing and Services Science, Barcelona, Spain, April 3-5, 2014., SciTePress (2014) 489–498
14. HERAS AF team: HERAS AF (Holistic Enterprise-Ready Application Security Architecture Framework). <http://herasaf.org/>
15. Garaga, A., de Oliveira, A.S., Sendor, J., Azraoui, M., Elkhiaoui, K., R. Molva, M.O., Cherrueau, R.A., Douence, R., Grall, H., Royer, J.C., Sellami, M., Südholt, M., Bernsmed, K.: D:C-4.1: Policy Representation Framework. Technical Report D:C-4.1, Accountability for Cloud and Future Internet Services - A4Cloud Project (2013)