Evaluation of User Specific Privacy Policy Architecture for Collaborative BPaaS on the Example of Logistics

Björn Schwarzbach[1](✉), Michael Glöckner[1], Bogdan Franczyk[1,2], and André Ludwig[3]

[1] Leipzig University, Leipzig, Germany
{schwarzbach, gloeckner, franczyk}@wifa.uni-leipzig.de
[2] Wrocław University of Economics, Wrocław, Poland
[3] Kühne Logistics University, Hamburg, Germany
andre-ludwig@the-klu.org

Abstract. Nowadays, collaboration between multiple companies along the supply chain is one of the key factors for ensuring sustainable success. Although this fact is known by almost all companies the actual collaboration is quite low because of the fear of losing sensitive and critical data to competitors. To solve this problem an architecture for modeling and execution of privacy preserved business processes and a privacy modeling approach have been developed. This paper evaluates both artifacts. The used method is framework for evaluation in design science (FEDS).

Keywords: Privacy · BPaaS · Evaluation · Cloud computing · XACML

1 Introduction

For many years, cloud computing has been very successful since it is being applied to an increasing number of use cases [1]. The most important key factors for this success are its dynamics, its decentralized nature, and the abstraction and outsourcing of physical IT systems. These factors make cloud computing indispensable for the logistics sector with its complex supply chains and collaborative processes. In the year 2008 Thomas J. Bittman, vice president of Gartner Research, published his thoughts on the future of Cloud Computing. He stated that there are three phases of Cloud Computing. While the first phase was focused on providing mostly infrastructure services with proprietary interfaces, the second phase introduced an ecosystem of smaller cloud providers offering services for the vertical supply chain based on the dominant providers of the first phase. During the third phase, these smaller providers form horizontal federations that lead to new interoperability standards of service communication.

The next logical phase is moving the management of collaborations into the Cloud. The stakeholders of the collaboration can not only access the information needed to operate their tasks but the services provided by the stakeholders can also be easily connected and combined with each other to form complex processes. Such a process consists of several IT services of different partners which interact with each other and

E. Ziemba (Ed.): AITM 2016/ISM 2016, LNBIP 277, pp. 137–154, 2017.
DOI: 10.1007/978-3-319-53076-5_8

exchange data. While the process has a global scope, the individual services are executed by the partners' local control, namely the service providers. This is called collaborative Business Process as a Service. Due to the collaborative nature with multiple involved partners, privacy, security, and confidentiality related requirements gain more and more importance [2]. Cloud service providers have to ensure that data, which is generated by a cloud service, is treated confidentially and is only stored or transmitted to other services or providers accordingly to the service providers privacy policies [3]. Therefore, a centralized trusted platform is needed that copes with these tasks. Further, it has to provide interfaces that are able to regulate and control the data flow. In addition, this platform must provide the functionality to model and execute use-case and user-specific privacy policies.

The chaining of services in a collaborative business process and the therewith resulting data exchange have to be defined in detail and have to be monitored at any time. Only that way, a high level of privacy, security, and confidentiality can be ensured [4]. Not only does confidentiality of the data have to be ensured but also the access of each process partner to the appropriate data has to comply with a specific time- and role-dependent set of policies. In [5] an architecture of a platform that enables companies to consume cloud services of providers has been introduced. The platform offers features of a business process management system in terms of orchestrating individual cloud services in business processes that have been modeled by consumers, i.e. the companies. Hence, the term for this approach and the service provided by the platform is *Business Process as a Service* (BPaaS).

Especially for companies participating in collaborative business processes, privacy is a very important topic in terms of risk and compliance [6]. In an interview, we discovered that most of the companies that do not consume cloud services are reluctant because they are afraid of losing control of their data. This is becoming even more important because of recent hack attacks on global players. Also, companies that consume at least one cloud service are concerned because of privacy issues [7].

In order to encourage companies to participate in collaborative processes, one of the main challenges is the preservation of data privacy and the compliance with privacy laws during business process execution. This is getting even more important when multiple competitors are collaborating in one business process. These competitors need to share data that is usually kept secret with each other. In [2], we have proposed an approach for secure service interaction, which has shown its feasibility in multiple tests. The architecture proposed in [5] also provides a component to annotate business processes and individual activities with privacy policies. These are evaluated and enforced by the platform during business processes execution. This paper focuses on the evaluation of these two artifacts.

The paper is structured as follows: After this introduction, the architecture of the platform for privacy preserving collaborative BPaaS and the modeling approach for privacy policies in collaborative business processes are presented. The third section outlines the applied evaluation framework, while section four discusses the artifacts' evaluation and the resulting findings. The paper closes with a conclusion featuring limitations, implications, and an outlook on future research.

2 Theoretical Background

2.1 Architecture for Privacy Preserved Business Processes

In this section we propose an architecture consisting of a central platform, third party services, and gateways to enable secure communication between these components. This architecture is depicted in Fig. 1.

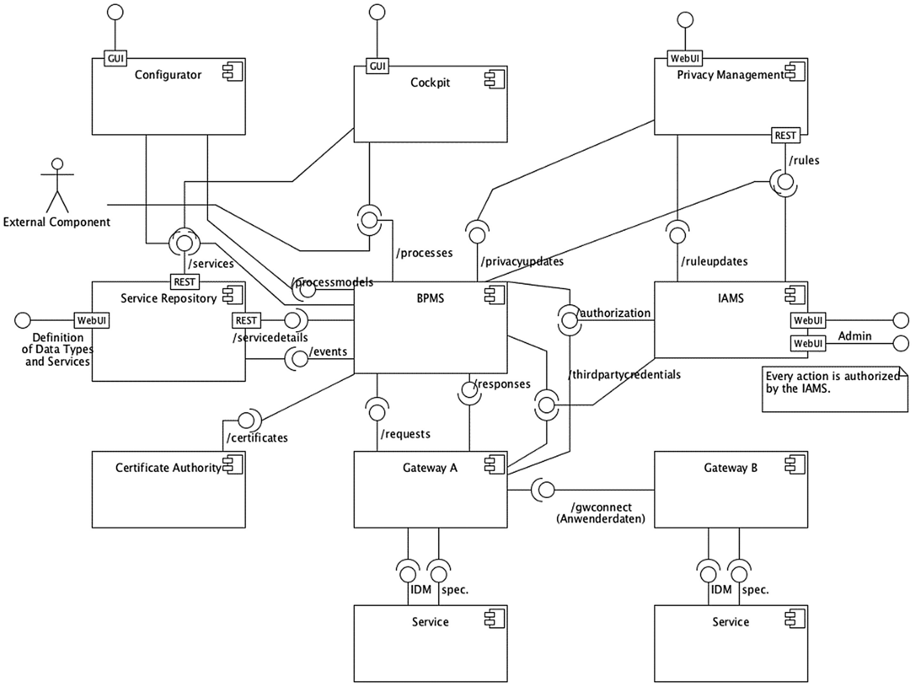

Fig. 1. Architecture of the platform and the gateways for privacy-preserving collaborative business processes

The platform comprises a user interface for the design of collaborative business processes. These collaborative business processes are composed of third party services. The third party service descriptions are stored in and managed by the service repository. The business process management system (BPMS) stores the business processes and provides features to instantiate and execute the processes. The user defines his privacy requirements in form of privacy rules with a set of privacy management tools. These privacy rules are used by the identity and access management system (IAMS) to determine whether an entity (e.g. user) is allowed to access an asset or not.

While executing a business process the BPMS has to invoke third party services. To ensure privacy during the service invocation and the whole business process we propose the use of gateways, which are similar to the proxies proposed by [8]. The core architecture of these gateways will be described in the following.

A gateway consists of four components that are organized as a stack. At the lower end of the stack the service adapter handles the current third party service consumption. To consume a third party service, the service adapter first requests authentication and authorization on behalf of the user of the platform by the service provider's identity and access management (IAM). The second step is to call the current service by using the interfaces provided by the service provider. To enable the gateway to handle multiple service providers there are multiple implementations of the service adapter, one for each IAM and service.

The platform uses a domain specific data scheme, which is not compatible with the third party service's input and output schemes. This domain specific data scheme helps to provide a common basis for communication. The domain specific data scheme is implemented as an ontology. Hence, the data need to be transformed between the global domain specific and the service specific schemes. This task is carried out through the data adapter by transforming the data from domain specific to service specific scheme and forwarding the result to the service adapter. After the service adapter finished the service call the result is forwarded to the data adapter. The data adapter transforms the result from service specific scheme to domain specific scheme. Since there are multiple service specific schemes multiple data adapters are implemented. The actual transformation done by the data adapter is configuration based. Hence, one data adapter can be used for multiple transformations, given that input and output schemes are the same.

Prior to transforming the data between the schemas the data need to be checked for potential privacy issues. This task is realized by the privacy guard. The privacy guard retrieves corresponding privacy rules for the service that need to be invoked and the data that need to be transferred. In addition, the privacy guard loads service meta data, e.g. number of service calls, and privacy data, e.g. number of hits on the data by the current user, that resulted from previous service invocations. After retrieving all required information on the privacy situation the privacy gateway processes the payload that need to be sent to the third party service according to the rules and meta data. This procession may include pseudonymization and anonymization, but also projections, e.g. scale conversion, may be applied. The privacy guard ensures that only data is transferred to the third party services that need to be transferred in order to fulfil the purpose of the current action of the business process. If the privacy guard is not able to apply all privacy rules an alert is raised in the cockpit and the business process is paused. This ensures that privacy is not violated at any time.

The communicator ensures that all payloads that need to be transferred between the actions of the business process, i.e. from one gateway to another, are encrypted with unique asymmetric keys. It also ensures that the payload of the services, both input and output data, is never loaded into the platform itself, neither encrypted nor unencrypted. The communicator's third task is to load privacy data and service meta data of the current service invocation into the platform to provide feedback for the next actions. This data is encrypted, too. To ensure that no data is kept in memory between two service invocations, gateways are for one time use only. Due to the gateway's modular structure it can be composed for different third party services and privacy rules. This ensures the maximal flexibility. Hence, the reusability, flexibility, easy adaptation and implementation are the key advantages of the gateways.

All the gateway components communicate with the neighbor units via encrypted web services. The encryption is based on a public key infrastructure (PKI) that is part of the platform. When the BPMS creates a gateway the PKI issues public and private keys to all four components. When a message reaches a component, the component checks the authenticity of the message and whether the sender is allowed to send this type of message, i.e. the sender is part of the same gateway and the component is the successor of the sender according. To enable the check of the gateway the issues keys will contain the ID of the gateway in their data. Since all components use a unique set of keys all information flows inside of the gateways are secured as long as the PKI is not compromised.

2.2 Privacy Policies for Collaborative Business Processes

This section describes in detail our approach for defining privacy policies in the context of collaborative business processes. One of our main requirements for the approach was to provide the companies with a tool that they could understand.

To define privacy policies that can be evaluated automatically and be used to decide whether a service is allowed to access some data or not we rely on use access control approaches. Basically there are four different types of access control. The mandatory access control and discretionary access control where applied in computer systems in the 70 s of the last century. While mandatory access control describes security from the system itself by policies like “access is only granted from localhost”, discretionary access control assigns each identity the appropriate access rights [9]. Mandatory access control is still used nowadays, e.g. SElinux is applying this approach [10].

In the late 80 s and early 90 s more and more users where using computer systems, hence assigning each individual user, i.e. identity, the correct access rights was not feasible any more. So in the beginning of the 90 s role based access control emerged [11]. Role based access control assigns roles to identities and access rights are assigned to roles. This approach is used in Linux and Windows file systems and almost every modern software. Roles can be organized hierarchically as shown in Fig. 2 [12–14]. Because of the well established application of role based access control our first approach for defining privacy policies was to apply role based access control.

During multiple workshops with local companies we discovered that the companies do not think about privacy identically. One common thing is that all companies separated the actors who want to access data into groups. But while some companies have had a very easy and strict approach for group setup, others could not clearly tell us which companies are member of which group. Instead they used phrases like “The driver of the truck while he is in the destination city is allowed to get the recipient’s phone number to call the recipient to tell him his arrival time”. This simple phrase contains the following information:

- ‘Driver’ has to be the role of the person requesting access to the recipient’s number.
- The requesting person has to be located in the area (city) of destiny.
- In addition, the driver is only permitted access to the number to announce the arrival.

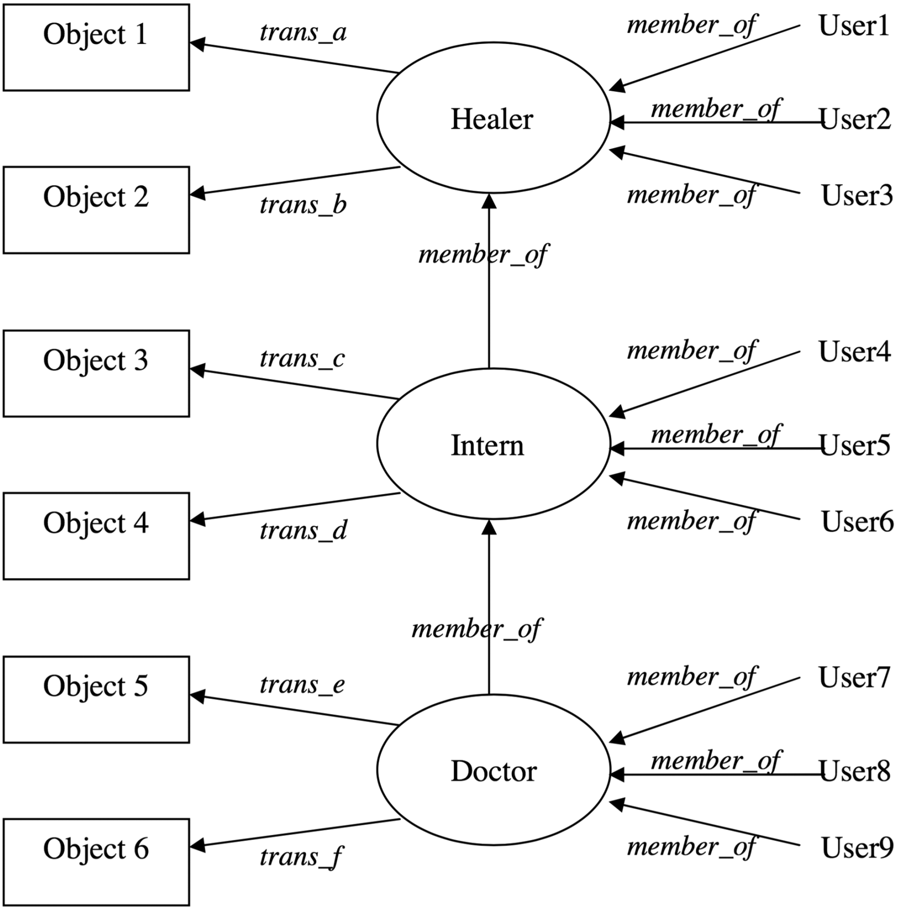

Fig. 2. Role based access control [14]

This simple policy cannot be represented easily with roles because the location of the driver is changing over time. To tackle such requirements, the research community followed two core approaches: extend role based access control with additional features, e.g. context or attributes, and creating a new access control model. [13].

Following the role based access control [15] has developed an access control model, which extends role based access control for virtual organizations. Unfortunately, this model does not cover business processes, workflows, and cloud computing. Other approaches in this direction do cover business processes but leave out the cooperation aspect. References [16–18] proposed and evaluated an extended role based access control model for team collaboration and workflows in the health sector.

All of the proposed models do not provide the flexibility in policy definition language that was needed by the participants of our workshops. To achieve a maximum flexibility, the research community developed a novel approach, the attribute based access control. In attribute access control, policies are based on attributes of subjects and objects. According to [19] attribute based access control is:

"An access control method where subject requests to perform operations on objects are granted or denied based on assigned attributes of the subject, assigned attributes of the object, environment conditions, and a set of policies that are specified in terms of those attributes and conditions." [19].

The entity requesting access is called a subject. Typical attributes of subjects are their id, e.g. username, company name, and name of the department. The data that subjects want to access is called object or resources. A policy in attribute based access control is a triple of a subject, a resource, an operation, where the operation describes what access type the subject wants to have, and a result, e.g. grant or deny access. A policy can also comprise one or more conditions. The policy "The driver of the truck while he is in the destination city is allowed to get the recipient's phone number to call the recipient to tell him his arrival time" consists of:

- *Subject:* The driver of the truck who is located in the destination city;
- *Resource:* Recipient's phone number;
- *Operation:* Read;
- *Condition:* Current workflow activity is "call recipient for dispatch notification";
- *Result:* Permit.

The remainder of this section presents our approach on applying attribute based access control for privacy preservation to collaborative business process as a service. Privacy of data is always specified to the owner of the data, i.e. the creating entity.

First of all, in our platform business processes consist of activities that call external cloud based web services. Hence, there are two very basic roles in our platform. A process designer is an entity that models the business process, that is responsible for the correctness of the process itself, and that offers the resulting business process as a service to its customers. The second role is the service provider. A service provider is an entity that provides the external services that are being orchestrated in the business process by the process designer.

Our approach enables both roles to define their privacy policies independently from each other. It also includes privacy policies defined by law. Hence, the combined privacy policy consists of three columns as shown in Fig. 3 that can be evaluated independently. The combined privacy policy results in permit if all three columns result in permit, else it results in deny.

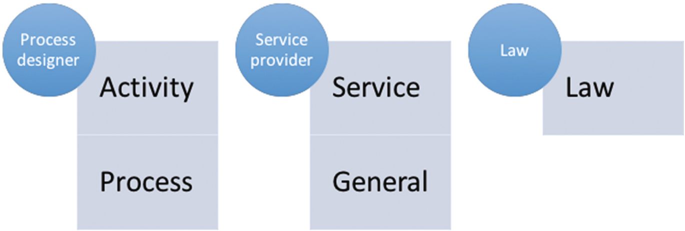

Fig. 3. Three columns of privacy policy process designer, service provider, and law

To simplify the process of policy definition and to reduce redundancy in policies we provide each role with two levels of policies. First a process designer can specify general privacy policies that are valid for the whole business process. E.g. a process designer may restrict access to all data to companies that are located in the Europe Union to ensure no data is transferred to other countries. Such privacy policies are visualized as tables where the objects are in the rows while the subjects are located in the columns. The cells contain either permit or deny depending on whether or not the subject is allowed to access the object. The subjects are defined by filters using attributes. So in this example the subject filter would be:

Companies meeting the condition: all locations have an attribute country with a value that is in a list of the countries of the European Union.

The relevant section of the table for the privacy policy "Data can only be accessed by European companies." is shown in Fig. 4. Apart from the groups created by the process designer, every table does have an additional column *Default*. The algorithm to select the correct column when the evaluation of a policy, i.e. table, takes place is select the rightmost column whose filter does accept the subject, where Default accepts every subject.

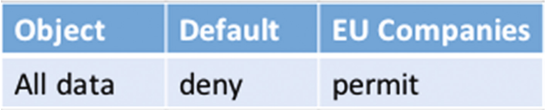

Object	Default	EU Companies
All data	deny	permit

Fig. 4. Privacy policy "Data can only be accessed by European companies" in table form

The rows of the table represent the objects, i.e. the data the policy is about. The data is organized in an object hierarchy. Objects can be expanded to define policies for child elements as shown in Fig. 5. This table also states that EU Companies have unspecified access to the object Child. The evaluation algorithm handles *t/s* as if this column does not exist. The only column that is not allowed to have *t/s* is the *Default* column since else there would be no result for the evaluation of the policy.

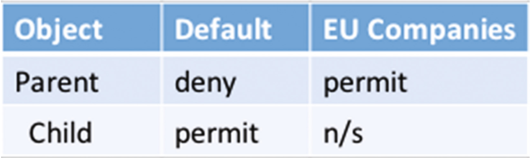

Object	Default	EU Companies
Parent	deny	permit
Child	permit	n/s

Fig. 5. Expandable objects in table to define a privacy policy for a child element different from the privacy policy for the parent element

The process level privacy policy applies to all data created by activities of the business process, i.e. it is assigned to all activities. If the process designer wants to define a different policy for a specific activity, he defines a privacy policy on activity level. Privacy policies on activity level are evaluated before the process level privacy policies, i.e. activity level overrides process level. On activity level even the *Default* column can be set to *n/s*. If the evaluation of activity level policy results in *n/s* the policy on process level is evaluated.

The groups of subjects of a process designer's privacy policies can use both, companies and roles of the business process, as target. In case the process designer wants to use a business process role as the subject's filter, the systems show up a list of the names of all swim lanes of the process. The process designer selects the appropriate entries and specifies the access rights as he does for company based filters.

The second role, i.e. the service providers, can define privacy roles that are applied to all data generated by their services in any business process. This is done on the level *General*. The definition of the policies follows the same concepts as for the process designer's policies. A service provider can override his general privacy policies by setting up a service specific privacy policy.

The third type of privacy policies are laws. Laws are provided by the platform provider as is and are not represented in an easy to read form as the process designer's and service provider's policy are.

3 Research Methodology

The evaluation of the presented artifacts is based on the Framework for Evaluation in Design Science Research (FEDS), developed by Venable [20]. Artifacts that have been created by a methodology based on Design Science in Information Systems Research [21] can be evaluated by this framework in order to ensure rigor. The framework offers several strategies the could be pursued depending on the characteristics of the designed artifacts. Generally, the FEDS regards evaluation as an ongoing process during design science research in order to improve the artifacts iteratively. Several characteristics influence the evaluation's purpose (why?), point of progress of the design process (when?), strategy (how?) and the artifact itself (what?). The characteristics and resulting strategies are briefly introduced. By outlining the characteristics of the current research, a strategy is chosen and the resulting methodological steps are described.

The framework distinguishes between formative and summative evaluation [22]. Formative evaluation has the main purpose to improve the results of an artifact in the ongoing research process. On the contrary, summative evaluations have the purpose to create a shared meaning of the artifact concerning distinct contexts of application. The question about the point of progress of design evaluation can be chosen ex-ante or ex-post [22] during the continuous design process. While ex-ante evaluations are more predictive in order to e.g. select a certain technology alternative, ex-post approaches are used to assess developed artifacts in terms of applicability or degree of achievements of objectives. With this, a greater likelihood of ex-post evaluation can be expected for summative evaluations but is not obligatory [20]. Goals of evaluations can be for different purpose: either achievement of environmental utility, or usefulness of solving a specific problem, or comparative advantage over existing solutions, or a complex composite of criteria (e.g. functionality, completeness, consistency), or other impacts (side effects), or reason artifact's functioning.

The framework is displayed in Fig. 6, it comprises two dimensions. On the x-axis the distinction between already described formative and summative evaluation purpose is located. The y-axis contains a distinction on how to evaluate with either artificial or naturalistic setup. While artificial setup is used to prove general functionality of a concept, naturalistic evaluations prove an artifacts functionality in real environments, i.e. real people, real systems, and real settings [23]. Different *strategies* can be pursued that are displayed in Fig. 6 as well. Depending on the needs, available resources and circumstances, a strategy is chosen for and possibly changed during evaluation. The fastest strategy with the lowest costs is found in the 'quick\&simple' approach with a very limited number of iterations bears the risk of being not reasonable. A *'Purely Technical'* approach is suitable if naturalistic data and behavior is irrelevant and human users are not focus of the artifact. The other two strategies are used for either facing *'Human Risk & Effectiveness'* or *'Technical Risk & Efficacy'*. A more detailed description of selecting a suitable strategy depending on specific circumstances can be found in Table 1.

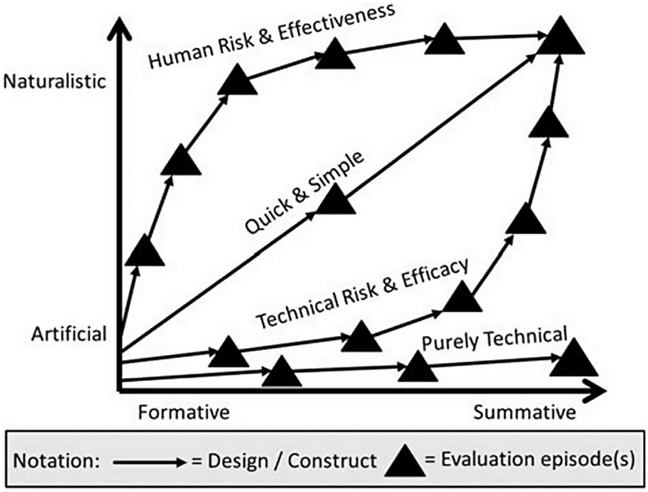
Human Risk & Effectiveness
Naturalistic
Quick & Simple
Technical Risk & Efficacy
Artificial
Purely Technical
Formative
Summative
Notation: → = Design / Construct ▲ = Evaluation episode(s)

Fig. 6. Framework for Evaluation in Design Science (FEDS) with evaluation strategies [20]

The FEDS proposes 4 particular steps during the evaluation process [20]:

1. *explicate the goals:* 4 goals of the evaluation could be distinguished:
   (a) *rigor* focuses on confirming that the artifact directly produced a certain effect (more likely to be shown with artificial evaluation) or that an instantiation of the artifact works thoroughly in a real situation (more likely to be shown with naturalistic evaluation). A summative evaluation provides the greatest rigor and reliability of the produced knowledge [20].
   (b) *uncertainty and risk reduction* focuses on reducing either human and social risks or on reducing technical risks, which influences the choice of strategy (see Table 1).

Table 1. Circumstances for selecting a relevant DSR evaluation strategy [20]

DSR evaluation strategy	Circumstances selection criteria
Quick & Simple	If small and simple construction of design, with low social and technical risk and uncertainty
Human Risk & Effectiveness	If the major design risk is social or user oriented *and/or* If it is relatively cheap to evaluate with real users in their real context *and/or* If a critical goal of the evaluation is to rigorously establish that the utility/benet will continue in real situations and over the long run
Technical Risk & Efficacy	If the major design risk is technically oriented *and/or* If it is prohibitively expensive to evaluate with real users and real systems in the real setting *and/or* If a critical goal of the evaluation is to rigorously establish that the utility/benet is due to the artifact, not something else
Purely Technical Artefact	If artifact is purely technical (no social aspects) or artifact use will be well in future and not today

(c) *ethics* focuses on reduction of potential risks to animals, people, or the public society. With this especially potential stakeholders should not be put into risk.

(d) *efficiency* focuses on balancing the aforementioned goals in case of resource shortage Hence, a more formative evaluation is proposed.

2. *choose a strategy or strategies for the evaluation:* Depending on the aforementioned goals and the described circumstances of Table 1 one or more strategies have to be chosen. This can be done with a 4-step heuristic: (1) evaluate and prioritize design risks (either social/user oriented or technical or both). (2) Estimation of costs for real users, real systems and real settings. If human feedback is available for a reasonable price the *'Human Risk & Effectiveness'* strategy is suitable. If the price is to high or serious health concerns exist for users, the *'Technical Risk & Efficacy'* strategy is favorable. (3) If the artifact is purely technical and potential usage lies in remote future, the 'Purely Technical' strategy appears to be suitable or a naturalistic evaluation is just impossible. (4) If the construction that is to be evaluated is of rather small and simple extent, and none of the above mentioned risks apply, the *'Quick & Simple'* strategy is the best choice.
3. *determine the properties to evaluate:* the general set of features, goals and requirements of the artifacts that are to be evaluated are chosen. Again, a heuristic with 4 steps is proposed: (1) determine a list of potential evaluands (examples are given in [23–26]), (2) evaluands are to be aligned with the chosen goals, (3) depending on the chosen strategy of step 2, the evaluands should be of rather naturalistic or technical character and (4) determine the final list of evaluands.
4. *design the individual evaluation episode:* the 3 heuristic sub-steps comprise: (1) derived from the environmental constraints, availability of resources determines their usage. (2) Priority shall be given to essential and more important aspects and

resource are to be (re-)allocated. (3) Determination of number and structure of evaluation episodes and the according responsibility.

4 Research Findings and Discussion

This section presents the evaluation of the research artifacts. As shown in the previous section, the preparative work for the applied Framework for Evaluation in Design Science (FEDS) is structured in four steps: explicate the goals, choose a strategy or strategies for the evaluation, determine the properties to evaluate, and design the individual evaluation episodes. The evaluands are (1) the architecture and (2) the modeling approach for privacy policies. Since both artifacts can be evaluated separately the remainder of this section is split into two parts, one for each artifact.

4.1 Evaluation of the Architecture

Preparation. The first step of the evaluation is to explicate the core goals. The first goal of the architecture's evaluation is to show the architecture's *ability to model and execute* the collaborative business processes with privacy preservation. The second goal of the evaluation is *efficiency*.

The second step focuses on the selection of strategies for the architecture's evaluation. Since naturalistic evaluation involves real data from real businesses and the architectures purpose is to transfer these data between multiple partners of a business process, there is the risk of exposing data illegally. In order to avoid the exposition of real sensitive data in case of non-successful evaluation, only artificial data is used. To reduce uncertainty and risks, a formative evaluation needs to be applied first, especially to ensure that the technology is able to fulfill privacy preservation and that technical risks are reduced during design time. The architecture consists of several components which are connected with each other, as shown in Fig. 1. This leads to a high complexity. Further, the major risks are technology-driven as no manual input is necessary for the interpretation of the privacy rules. Additionally, the evaluation with real companies and real data does not provide further or deeper insight. According to Table 1, in case of absence of naturalistic evaluation, a *Purely Technical* strategy is selected to evaluate the architecture's ability to model and execute privacy preservation.

In the third step the architecture's properties that are to be evaluated, are determined. Those are partly inspired by the properties proposed by [25]. First, essential properties of the architecture are the *existence of corresponding components* that enable a user to individually model privacy rules and also components that enable the automatic execution and interpretation of privacy rules. The platform's main purpose is to ensure privacy, hence *security* and *reliability* are further crucial properties. Interoperability is crucial but also given due to the cloud computing nature of the architecture.

The last step is the design of the individual evaluation episode. The evaluation of the architecture was done after substantial changes of the components and their links. For this purpose, the architecture was discussed with teams of expert to ensure the appropriate definition, composition, and implementation of each component. The experts were selected from different research groups. Further, attendants of conferences were contributing to the architecture by discussions after the presentations of the actual status. Additionally, system tests were applied to evaluate the architecture.

Evaluation Episodes. The first part of the evaluation was done on a regular basis throughout the first iterations of the architecture's design by expert discussion with experts from the fields cloud computing, service orientation, logistics, and privacy. The discussions were initialized by brainstorming. The main ideas of the participants have been summarized and were used to improve the architecture. Examples are given in the following. In an earlier version of the architecture the same gateways were reused for the next service interactions and business data was stored in the business process management system. The experts raised some concerns that both could lead to a potential security risk in terms of data leakage or data remaining in cache. Hence, these issues were addressed in later iterations of the platform's design process. After the architecture reached a first stable state with only minor changes it was presented at two scientific conferences where additional feedback was collected during the discussion with the conferences' attendants.

After the completion of the implementation of the architecture's prototype, the technical evaluation was done mostly based on automated and semi-automated tests with artificial test cases. First, the evaluation of the components focused on the privacy related ones, which are privacy management system, identity management system, business process management system, service repository, certificate authority, and gateways. A test suite has been implemented, which executes specific test cases towards these components. The test cases comprise the execution of a task of a simulated or emulated, respectively, business process. This included the instantiation of a gateway and the third party service invocation as well as all the interpretation and testing of all related privacy policies. The tests have been executed on demand. The configurator was able to model all possible aspects according to the modeling approach for both, business processes and services. Hence, the configurator is able to model privacy policies for collaborative business processes (The evaluation of modeling approach itself by researchers and practitioners is described in the next subsection). The results confirmed the prevention of privacy violations, timing information, input messages, and output messages of the components. This way the ability of those components of the architecture to model and execute privacy rules could be positively evaluated. Further, important timing information could be collected in order to prove the efficiency of the architecture's service invocation process. Due to the certificates' characteristic of only one-time use, the time needed to create certificates is a significant portion of the tests' total time. In average the tests took 22 s each, which was evaluated as a reasonable time by the experts. Due to the use of one-time-use certificates the architecture executes the business processes slower than a normal work flow engine. The additional time needed for privacy checking is irrelevantly lower compared to the time needed for creating those certificates. On the other hand, the use of such one-time certificates results in a substantial increase of security. So efficiency of the architecture could be improved by using long-term-certificates, but this has to be paid by the price of a significantly lower level of security. The evaluation has shown that the developed architecture enables the modeling and execution of collaborative business processes with privacy preservation. It has also proven that the components of the architecture are well selected and modeled. The continually performed automated and semi-automated system tests have shown that the architecture is able to handle service calls during the business process execution with acceptable additional resource consumption under the strict condition of privacy preservation. In summary, the evaluation confirmed that the architecture fulfills the requirements and goals that were imposed. This way it was proven that the architecture possesses the ability to model and execute privacy policies of the stakeholders of the collaborative business processes (Goal 1). Further, the experts evaluated the efficiency as reasonable (Goal 2).

4.2 Evaluation of the Modeling Approach for Privacy Policies

Preparation. The overarching goal comprises the *user acceptance* as the sum of two sub goals. The first sub goal for the evaluation of the modeling approach is to prove that ability of *modeling individual privacy policies* of companies. The second sub goal is to prove that the architecture's components can understand and execute the *privacy policies as intended by the user* throughout the execution of the collaborative business processes.

The evaluation strategy implemented considers the aspect of user interaction. Again, naturalistic evaluation data appears to be not appropriate for evaluation in order not to expose sensible and confidential data. User acceptance requires to present not only incremental progress but of course the final result in order to confirm user acceptance of the final version. Hence, formative as well as subsequent summative evaluation is conducted. In order to fully analyze user acceptance, near-naturalistic data and cases are used. The modeling approach requires human interaction, hence the 'Purely Technical' strategy is not appropriate. Because of the importance of the privacy policies the 'Quick & Simple' evaluation strategy is rejected. So the 'Human Risk & Effectiveness' strategy was selected. The modification of using only a near-naturalistic setup, as real confidential data would be too sensible to be exposed in case of non-positive evaluation.

The evaluated properties comprise the privacy and not-exposition of modeled policies of companies. If the policies itself would be visible to other companies, competitors could extract confidential data. For instance, if company A would be aware of the companies that are allowed to view certain data or if the geographical zone of certain business partners would be exposed, then competitors could identify or under certain circumstances estimate a list of critical business partners of their competitors. Further, the users' satisfaction with the modeling tool is a crucial property of this evaluation. The most important properties of the modeling approach are usability, maintainability, flexibility, and comprehensibility to ensure that the modeling approach is easy to use and understandable to the companies.

The evaluation episodes again comprise tests with users after significant changes of the user interface and/or the 'look & feel' of the prototype. For this purpose, the concept was presented to and discussed with teams of experts from different fields of research and from different companies in order to ensure the appropriate usability.

Evaluation Episodes. The first iterations of the modeling approach were evaluated by researchers and practitioners in workshops. During those workshops the participants had the task to model some artificial privacy policies. After the results became satisfying and a common understanding of the principles of the modeling approach was established, the participants were asked to model some near-naturalistic cases that are close to real privacy policies from their daily business. The modeling approach was also evaluated in the course of one research project. Privacy policies were modeled according to a realistic collaborative business process. These privacy policies were interpreted successfully by the prototype of the architecture. The first workshops have shown that there is no common understanding about what exactly the concept of privacy is and although a definition was given to the participants, they had difficulties in expressing their ideas. This got even more important as the first iterations of the modeling approach offered a maximum of flexibility at the expense of big complexity. It turned out that a reduction of complexity increased the comprehensibility of the resulting modeled privacy policies but reduced the possible flexibility. As an example, the first version of the approach comprised the modeling objects of 'Subjects', 'Resources', 'Actions', 'Environment', and 'Information'. The participants could not understand these terms and their accurate meaning easily. Hence, the terms were renamed to more common ones of the logistics sector. After renaming subjects to companies or group of companies, resources to attributes (data), actions to access granted, and skipping environment and information the participants could create the privacy policies more easily. This renaming has to be applied accordingly to the specific domain where the modeling approach will be applied to.

Of course this reduction in complexity and flexibility lead to some privacy policies that could not be modeled anymore, because the environmental and object based information were not available any more. This can be addressed by adding new sub-services to the platform which contain that information so the information can be used by other tasks of the process. After this modification the participants were able to model their privacy policies easily for their business processes and their services offered to their partners with near-naturalistic data.

It turned out that the participants' privacy understanding of business processes was different to their privacy understanding of services. While the privacy policies for business process were modeled in a way that the data flow along the intended flow of goods and information was always possible, the privacy policies for services were modeled much more restrictive. This lead to tasks of business processes that could not be aligned to appropriate services. This problem was addressed by a change of the architecture. The configurator was changed to check if there is at least one possible service for every task so that the process as a whole can be executed successfully under the consideration of all privacy policies modeled. The configurator presents an alert if this condition is broken and a continuous alignment is not achievable.

Although there is a big gap between the understanding of privacy policies of researchers and practitioners, all participants were able to model their policies individually. Hence, the modeling approach is able to individually model privacy policies of different companies. Those modeled policies can be interpreted and executed by the architecture's components successfully to ensure a privacy aware collaborative business process. Users were able to conduct those steps on their own, hence, the usability was confirmed (Goal 1). General feedback of the participants was the request for a more graphical user-friendly interface.

5 Conclusion

5.1 Summary

In this paper, an architecture for modeling and execution of privacy preserving collaborative business processes and a modeling approach for privacy policies for such business processes have been introduced briefly. Both artifacts have been evaluated successfully by applying the framework for evaluation in design science (FEDS). The feasibility of the architecture has been evaluated positively by expert discussions as well as automated and semi-automated tests. Multiple workshops during the design process of the modeling approach have shown the architecture's suitability and usability.

5.2 Limitations

A high number of the practitioners that took part in the evaluation process came from the logistics sector. Also, some of the participating researchers came from the fields of cloud and logistics. Hence, the evaluation's findings are valid for logistics use cases. However, it can be assumed that the developed artifacts can be applied to other fields as well. Yet, this is still an open task for future research. During our workshops, we did not only involve experts of cloud computing and privacy but also practitioners from the logistics sector. Another limitation is the rather small number of 25 companies that participated in the interviews and the evaluation workshops. However, it is assumed that the evaluation results are valid because of their homogenous nature.

5.3 Managerial and Scientific Implications

The evaluation has proven that the modeling of privacy policies for collaborative business processes by users is reasonably practicable. Hence, there are no major obstacles for the application of the modeling approach in real use cases. Also, the architecture for modeling and execution of privacy preserved collaborative business processes was evaluated positively. Hence, the architecture is suitable for application by cloud service providers. Both artifacts are important steps towards privacy preserving collaborative business process modeling and execution in cloud environments. This will enable small companies to take part in complex supply chains.

5.4 Outlook and Further Research Steps

As the users requested in their feedback during the evaluation of the modeling approach, a graphical interface with a higher user friendliness and usability appears to be a meaningful further research step. First steps in this direction looked promising, e.g. a graphical editor for XACML was found in [27]. Additionally, further large scale experiments with real use cases and real data will be helpful to identify potential shortcomings of the artefacts' design. Finally, application to other domains than logistics is a promising aspect.

Acknowledgement. The work presented in this paper was funded by the German Federal Ministry of Education and Research under the projects PREsTiGE (BMBF 16KIS0082K) and LSEM (BMBF 03IPT504X).

References

1. Wolf, M.-B., Rahn, J., Hompel, M.T.: Cloud Computing für Logistik 2: Akzeptanz und Nutzungsbereitschaft der Logistics Mall bei Anwendern und Anbietern: [eine qualitative und quantitative empirische Analyse des Fraunhofer-Institutes für Materialfluss und Logistik IML. Fraunhofer Verlag (2013)
2. Schwarzbach, B., Pirogov, A., Schier, A., Franczyk, B.: Inter-cloud architecture for privacy-preserving collaborative BPaaS. QUIS14 (2015)
3. Takabi, H., Joshi, J.B.D., Ahn, G.-J.: Security and privacy challenges in cloud computing environments. IEEE Secur. Priv. **8**(6), 24–31 (2010)
4. Bélanger, F., Crossler, R.E.: Privacy in the digital age: a review of information privacy research in information systems. MIS Q. **35**(4), 1017–1042 (2011)
5. Schwarzbach, B., Glöckner, M., Pirogov, A., Röhling, M.M., Franczyk, B.: Secure service interaction for collaborative business processes in the inter-cloud. In: 2015 Federated Conference on Computer Science and Information Systems, pp. 1377–1386. IEEE (2015). doi:10.15439/2015F282
6. Pearson, S.: Taking account of privacy when designing cloud computing services. In: Proceedings of the 2009 ICSE Workshop on Software Engineering Challenges of Cloud Computing, pp. 44–52 (2009)
7. Bundesamt, S.: 12% der Unternehmen setzen auf Cloud Computing. https://www.destatis.de/DE/PresseService/Presse/Pressemitteilungen/2014/12/PD14\textunderscore467\textunderscore52911.html(2014)
8. Singhal, M., Chandrasekhar, S., Ge, T., Sandhu, R., Krishnan, R., Ahn, G.-J., Bertino, E.: Collaboration in multicloud computing environments: framework and security issues. Computer (2013). doi:10.1109/MC.2013.46
9. Cuppens-Boulahia, N., Cuppens, F., Garcia-Alfaro, J. (eds.): DBSec 2012. LNCS, vol. 7371. Springer, Heidelberg (2012). doi:10.1007/978-3-642-31540-4
10. Lindqvist, H.: Mandatory access control. Master's Thesis in Computing Science, Umea University, Department of Computing Science, SE-901, vol. 87 (2006)
11. Ferraiolo, D., Cugini, J., Kuhn, D.R.: Role-Based Access Control (RBAC): features and motivations. In: Proceedings of 11th Annual Computer Security Application Conference, pp. 241–248 (1995)
12. Zahid, I., Josef, N.: Towards semantic-enhanced attribute-based access control for cloud services. In: 2012 IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications, pp. 1223–1230 (2012). doi:10.1109/TrustCom.2012.280
13. Jin, X., Krishnan, R., Sandhu, R.: A unified attribute-based access control model covering DAC, MAC and RBAC. In: Cuppens-Boulahia, N., Cuppens, F., Garcia-Alfaro, J. (eds.) DBSec 2012. LNCS, vol. 7371, pp. 41–55. Springer, Heidelberg (2012). doi:10.1007/978-3-642-31540-4_4

14. Ferraiolo, D.F., Kuhn, D.R.: Role-based access controls. arXiv preprint arXiv:0903.2171 (2009)
15. Gouglidis, A., Mavridis, I.: domRBAC: an access control model for modern collaborative systems. Comput. Secur. **31**(4), 540–556 (2012)
16. Le, X.H., Wang, D.: Development of a system framework for implementation of an enhanced role-based access control model to support collaborative processes. In: Proceedings of 3rd USENIX Workshops on Health Security and Privacy (2012)
17. Le, X.H., Doll, T., Barbosu, M., Luque, A., Wang, D.: An enhancement of the role-based access control model to facilitate information access management in context of team collaboration and workflow. J. Biomed. Inform. **45**(6), 1084–1107 (2012)
18. Le, X.H., Doll, T., Barbosu, M., Luque, A., Wang, D.: Evaluation of an enhanced role-based access control model to manage information access in collaborative processes for a statewide clinical education program. J. Biomed. Inf. (2014). doi:10.1016/j.jbi.2013.11.007
19. Hu, V.C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., Scarfone, K.: Guide to Attribute Based Access Control (ABAC) definition and considerations. national institute of standards and technology (2014)
20. Venable, J., Pries-Heje, J., Baskerville, R.: FEDS: a framework for evaluation in design science research. Eur. J. Inf. Syst. (2014). doi:10.1057/ejis.2014.36
21. Hevner, A., March, S., Park, J., Ram, S.: Design science in information systems research. MIS Q. **28**(1), 75–105 (2004)
22. Wiliam, D., Black, P.: Meanings and consequences: a basis for distinguishing formative and summative functions of assessment? Brit. Educ. Res. J. **22**(5), 537–548 (1996)
23. Sun, Y., Kantor, P.B.: Cross-evaluation: a new model for information system evaluation. J. Am. Soc. Inf. Sci. Technol. (2006). doi:10.1002/asi.20324
24. Stufflebeam, D.L.: The CIPP model for evaluation. In: Kellaghan, T., Stufflebeam, D.L. (eds.) International Handbook of Educational Evaluation, vol. 9, pp. 31–62. Springer, Dordrecht (2003). Kluwer International Handbooks of Education
25. Mathiassen, L., Munk-Madsen, A., Nielsen, P.A., Stage, J., Jacksen, M.: Object-Oriented Analysis and Design. Marko, Aalborg (2000)
26. Smithson, S., Hirschheim, R.: Analysing information systems evaluation: another look at an old problem. Eur. J. Inf. Syst. (1998). doi:10.1057/palgrave.ejis.3000304
27. Nergaard, H., Ulltveit-Moe, N., Gjøsæter, T.: A scratch-based graphical policy editor for XACML. In: ICISSP 2015 Proceedings of the 1st International Conference on Information Systems Security and Privacy ESEO, Angers, Loire Valley, France, pp. 182–191 (2015)