Scenario-based Evaluation of Privacy-enhancing Service Composition Solution

Editors:	Uli Pinsdorf, (EMIC)
	Stuart Short, (SAP)
Reviewer:	Jan Schallaböck, (ULD)
Identifier:	D6.3.3
Type:	Deliverable
Class:	Public
Date:	May 18, 2011

Abstract

In this deliverable we validate the technology for privacy-aware service composition against the requirements that have been formulated in the beginning of the project. The requirements, defined in the public deliverable H6.3.1 expressed how privacy-aware services should be built from a legal, technical, and user experience point of view. For the evaluation of the WP 6.3 demonstrator, a workshop with participants from EMIC, SAP and ULD was held at ULD in Kiel, Germany from May 5th to May 6th 2011. Within this workshop, EMIC and SAP presented their current work on the eCV scenario and illustrated how this would work in practice. All in all the evaluation showed that the WP6.3 demonstrator achieves the goal of being a working system that is in line with the privacy requirements developed. The current status of the demonstrator also illustrates that the requirements for privacy-enhancing Service-oriented architectures are application oriented and not too difficult to be addressed.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n^o^ 216483 for the project PrimeLife.

Members of the PrimeLife Consortium

1.	IBM Research GmbH	IBM	Switzerland
2.	Unabhängiges Landeszentrum für Datenschutz	ULD	Germany
3.	Technische Universität Dresden	TUD	Germany
4.	Karlstads Universitet	KAU	Sweden
5.	Università degli Studi di Milano	UNIMI	Italy
6.	Johann Wolfgang Goethe – Universität Frankfurt am Main	GUF	Germany
7.	Stichting Katholieke Universiteit Brabant	TILT	Netherlands
8.	GEIE ERCIM	W3C	France
9.	Katholieke Universiteit Leuven	K.U.Leuven	Belgium
10.	Università degli Studi di Bergamo	UNIBG	Italy
11.	Giesecke & Devrient GmbH	GD	Germany
12.	Center for Usability Research & Engineering	CURE	Austria
13.	Europäisches Microsoft Innovations Center GmbH	EMIC	Germany
14.	SAP AG	SAP	Germany
15.	Brown University	UBR	USA

Disclaimer: *The information in this document is provided "as is", and no guarantee or warranty is given that the information is fit for any particular purpose. The above referenced consortium members shall have no liability for damages of any kind including without limitation direct, special, indirect, or consequential damages that may result from the use of these materials subject to any liability which is mandatory due to applicable law. Copyright 2011 by ULD, EMIC, SAP.*

List of Contributors

This deliverable has been jointly authored by multiple PrimeLife partner organisations. The following list presents the contributors for the individual parts of this deliverable.

Chapter	**Author(s)**
Executive Summary	Ulrich Pinsdorf (EMIC), Stuart Short (SAP)
Introduction	Ulrich Pinsdorf (EMIC)
eCV Scenario	Ulrich Pinsdorf (EMIC)
Evaluation	Leif-Erik Holtz, Harald Zwingelberg, Uli König (ULD), Ulrich Pinsdorf (EMIC), Stuart Short (SAP)
Closing Remarks	Stuart Short (SAP)

Executive Summary

In this deliverable we validate the technology for privacy-aware service composition against the requirements that have been formulated in the beginning of the project. The requirements, defined in the public deliverable H6.3.1 expressed how privacy-aware services should be built from a legal, technical, and user experience point of view.

The requirements are derived from an analysis of two central regulatory instruments within the European Union, namely the data protection directives 95/46/EC and 2002/58/EC. This analysis had previously been conducted in the context of SOA in Virtual Organization by project partners ULD who are also the evaluators of the eCV demonstrator.

We selected a scenario-based approach by defining a concrete showcase for cross-domain service composition. Mechanisms such as sticky policies expressed in PPL, policy-based dynamic binding, trusted mobile device, and policy composition help maintaining user's privacy. We put a strong emphasize on legal evaluation. The scenario-based approach supports this emphasize, since the concrete scenario defines the case law. The validation this closes the lifecycle loop for this work-package.

The evolution of the scenario was influenced in many ways by the nature of the partnership structure; with EMIC, SAP and G&D being separate legal entities it was clear from the start that this would shape the overall structure of the cross-domain service composition both from a scenario and technical perspective. The clear division of tasks and application development ensured rich collaboration amongst the partners and enabled a role-playing environment.

For the evaluation of the WP 6.3 demonstrator, a workshop with participants from EMIC, SAP and ULD was held at ULD in Kiel, Germany from May 5th to May 6th 2011. Within this workshop, EMIC and SAP presented their current work on the eCV scenario and illustrated how this would work in practice. The options of the eCV scenario were illustrated and its potential impact on existing Service-oriented infrastructures was discussed. The differences between the eCV scenario in the WP 6.3 demonstrator and existing career platforms (especially regarding privacy aspects) were pointed out and the option of implementing parts of the eCV scenario as part of the WP 6.3 demonstrator functionality also into mobile devices (especially also for a more thorough description cf. PrimeLife Deliverables D6.3.1 and D6.3.2) was illustrated.

All in all the evaluation showed that the WP 6.3 demonstrator achieves the goal of being a working system that is in line with the privacy requirements developed. The current status of the demonstrator also illustrates that the requirements for privacy-enhancing Service-oriented architectures defined in [MeS09] are application oriented and not too difficult to be addressed.

Contents

Chapter 1

Introduction

In this deliverable we validate the technology for privacy-aware service composition against the requirements that have been formulated in the beginning of the project. The requirements, defined in the public deliverable H6.3.1 expressed how privacy-aware services should be built from a legal, technical, and user experience point of view. This collection of 39 requirements, assigned to 5 groups, was very influential in course of PrimeLife. Of course they have served as guideline for the work in this work-package. But more important the requirements were given to Activity 5 and helped defining the requirement for the PrimeLife Policy Language (PPL).

This work-package built a demonstrator to showcase research results. We selected a scenario-based approach by defining a concrete showcase for cross-domain service composition. This so called "eCV scenario" is motivated by a job application service that processes user data and communicates them to several involved parties. Mechanisms such as sticky policies expressed in PPL, policy-based dynamic binding, trusted mobile device, and policy composition help maintaining user's privacy. As part of that the PPL policy engine was integrated.

This deliverable validates the implementation against the requirements from the beginning of the project. We put a strong emphasize on legal evaluation. The scenario-based approach supports this emphasize, since the concrete scenario defines the case law. The validation this closes the lifecycle loop for this work-package.

The document is structured as follows. First we will recap the eCV scenario. Although this scenario has been described several times in various levels of detail and technical depth, it may be convenient for the reader to have a self-contained evaluation document. The third chapter contains the assessment itself. We first describe the setting of the evaluation. Then analyze each requirement one-by-one. The evaluation is always against the implemented prototype. Since WP6.3 worked out a generalization of privacy-aware service composition, the "Abstract Privacy Lifecycle" (cf. D6.3.2), we will give an outlook how the requirement can be addressed if it was not addressed in our specific scenario or scope of implementation. The last chapter summarizes the key findings.

Chapter 2

eCV Scenario

PrimeLife experimented with a demonstrator prototype in order to evaluate the challenges in turning a common SOA application into a privacy-preserving SOA application. Moreover, the demonstrator scenario was shaped in a way that it gives room to showcase many privacy-enhancing extensions, especially w.r.t. the PPL policy engine from Activity 5 [Pri09a]. The electronic CV scenario was already presented it in PrimeLife reports H6.3.1 [MS09], D6.1.1 [Pri09b], and D6.3.2 [Pri11a]. In fact, this chapter is to large extent taken from H6.3.2 in order to make this report a self-contained and thus make it easier for the interested reader to follow the reasoning behind the evaluation.

2.1 Persona and the typical Use-Case

Inga Vainstein is 46 years old and is currently working as journalist. As a part of her job she is traveling to various countries. Inga makes heavy use of online applications for new job opportunities.

She uses a platform to get job offers and apply for new positions in a convenient and easy way. In fact, this is one of her main motivations to collect all certificates and testimonials (also) in electronic form. She uses this platform to collect all her digital certificates and documents. She attends professional trainings on a regular basis. For each completed course she gets a certification. Moreover she collects testimonials from former employers. Last year she won an award for her outstanding press story on identity theft.

The eCV platform allows creating many profiles based on these claims, e.g. one profile with an emphasis on her academic achievements and another profile with journalistic achievements. For each job offer she can decide which profile to send to the employer.

Altogether the eCV scenario features five roles, which we will describe in more detail now. Please notice that only the roles of User and Employer are needed to showcase the scenario. The remaining roles act more or less autonomously.

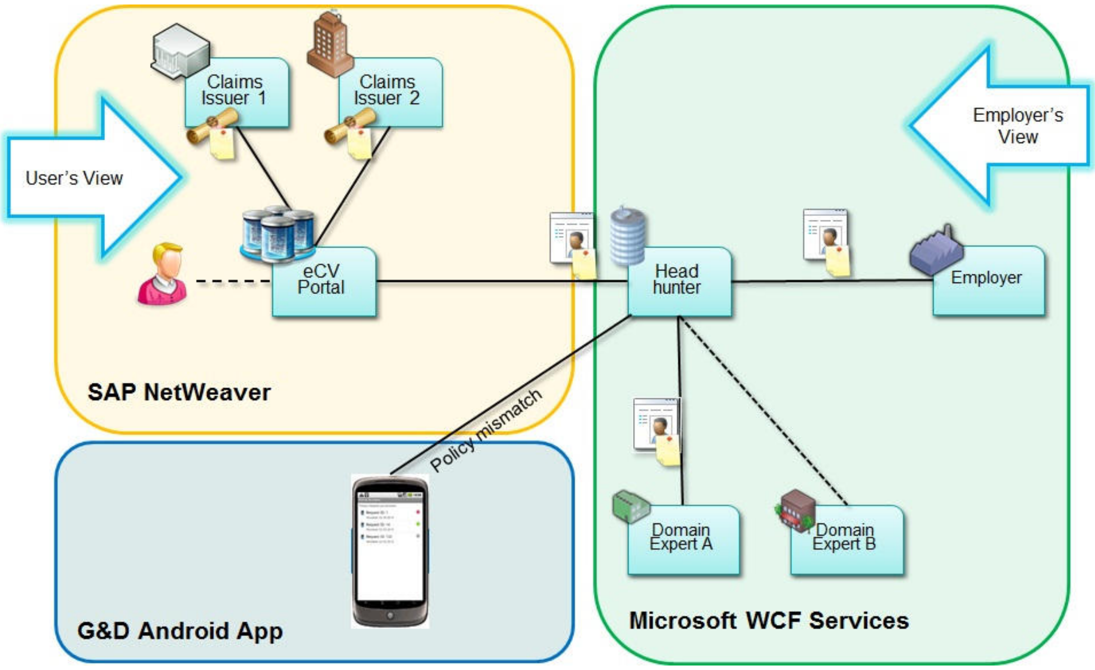

Fig. 1: The eCV use case

2.2 Roles and Workflow

Claims issuer: The claims issuer certifies an attribute to the user. This attribute comes in the form of a digital claim to the user and is accompanied by a privacy policy. This policy is crafted by the claims issuer and defines how long the claim is valid and what rights and obligations are associated with it. The claim policy is expressed in the PrimeLife Policy Language (PPL). For instance, a former professor might write a recommendation letter (claim), but express the obligation that the user may not use it for an application outside this university. A former employer could state for another letter of recommendation that this latter may not be used for an application at the employer's main competitor. The claims come in the form of a "sticky policy" attached to the claims. This attaching is done by means of an API or schema that features two slots, one for the claim and one for the policy.

User: Main goal for the user is to enjoy the benefits of a privacy-friendly online job application portal. For this purpose she collects claims from the claims issuers. She stores these in the portal and combines them to profiles. One profile consists of a collection of claims, additional information, and a privacy policy. The eCV portal composes this policy from all individual polices attached to the claims. Moreover the user can scope Section 7.2: End-to-End Workflow 57 this automatically generated policy further down. She can add more obligations or grant fewer rights.

Employer: The employer generates a job offer. The intention of the employer is to hire somebody for an open position. Thus, the employer generates a description of the open position. The employer communicates through the headhunter with the eCV portal. That means, the employer gives the open position to the headhunter and leaves it to him to find a suitable candidate. The headhunter in turn uses the eCV portal (potentially as one of many means) to find this candidate. This setting allows us to simulate and experiment with a longer chain of downstream data controllers. Besides a description of the open position, the job offer document features a policy that the employer promises to adhere to in course of this hiring transaction. It is thus a service provider policy. Unlike the general PPL scenario this policy is not submitted via a kind of web service meta-data request, but is attached to the job offer document and communicated upstream through the headhunter to the eCV portal.

Headhunter: The headhunter is the central turning point in this scenario. The reason to introduce this extra layer of communication in the scenario is to make the scenario richer in terms of downstream data usage. Moreover, the user is not aware of this instance, which creates some interesting use cases for the reasoning on policies. The headhunter receives a job offer (with an attached sticky policy) from the employer and forwards it to the eCV portal. The portal in turn sends a job application to the headhunter. The headhunter is now responsible to evaluate the capabilities of the applicant on behalf of the employer. We assume that this evaluation needs domain-specific knowledge. Hence, the headhunter dynamically looks up a suitable domain expert service and hands over the job application for evaluation. Of course, the headhunter needs to make sure that the domain expert complies with the policy attached to the job application.

eCV Portal: The eCV portal is the interface to the user. It allows for requesting claims, administrating issued claims, creating of profiles, and definition of user's privacy policy. Furthermore it utilizes policy composition and policy matching.

Domain Expert: The domain expert receives a job application from the headhunter. Its job is to evaluate the skills of an applicant according the skill set demanded by the employer. It stores the data as maximally as long as the obligation allows it to. Primary objective is to showcase the automatic execution of obligations. Moreover it is a dynamically bound service. That means that neither the user nor the employer know about its existence beforehand. The headhunter knows *that* a domain expert is part of his business process, but not *which instance* will be invoked. Certainly, the policy and the mechanism shall be able to deal with this situation.

Chapter 3

Evaluation

This chapter comprises the evaluation of the current state of the demonstrator. The current status of the demonstrator was evaluated by ULD together with EMIC and SAP. Below, we will illustrate the modus operandi of the evaluation workshop as well as its settings. Thereafter, the findings of this evaluation will be displayed as well as some more abstract results will be described as a summary.

3.1 Modus operandi & workshop settings

For the evaluation of the WP 6.3 demonstrator, a workshop with participants from EMIC, SAP and ULD was held at ULD in Kiel, Germany from May 5th to May 6th 2011. Within this workshop, EMIC and SAP presented their current work on the eCV scenario and illustrated how this would work in practice. The options of the eCV scenario were illustrated and its potential impact on existing Service-oriented infrastructures was discussed. The differences between the eCV scenario in the WP 6.3 demonstrator and existing career platforms (especially regarding privacy aspects) were pointed out and the option of implementing parts of the eCV scenario as part of the WP 6.3 demonstrator functionality also into mobile devices (especially also for a more thorough description cf. PrimeLife Deliverables D6.3.1 and D6.3.2) was illustrated.

Hereafter, the 39 requirements for privacy-enhancing Service-oriented architectures defined in and derived from PrimeLife Heartbeat H6.3.1 [MeS09] were addressed one by one and their implementation into the current demonstrator was evaluated. The evaluation was supported by the authors of [MeS09] to grant a maximum of significance. The findings of this evaluation will be illustrated below in section 1.2.

3.2 Evaluation of Requirements

The 39 requirements for privacy-enhancing and privacy-compliant SOAs [MeS09] can be categorized as "core policy requirements" (No.1-9), "privacy logging requirements" (No.10-18), "requirements on access to primary information" (No. 19-24), "cross-domain-specific requirements" (No. 25-30) and "requirements for additional mechanisms" (No. 31-39).

The evaluation workshop aimed at summarizing if and how these requirements have been met by the current eCV scenario within the WP 6.3 demonstrator. The various results, if and how these requirements are met, will be illustrated in the remainder of this Chapter.

3.2.1 Core Policy Requirements

No	Requirement	Achieved
1	Policies should be available in an unambiguous formalization. Thereby, the content of policies should be machine-interpretable.	Yes

This requirement was met by the current demonstrator. Using XML and the PrimeLife policy language, a machine interpretable policy is available. More detailed information about the XML schema of PPL, the design of the PPL language, and the implementation of the PPL policy engine is provided in other PrimeLife reports H5.2.2, D5.3.4, [CaVSa10], and others.

No	Requirement	Achieved
2	It must be ensured that communicated policies cannot be argued by the ensuring entity.	No

This requirement was not implemented in the scenario. Digital signatures are well understood mechanisms and thus were not prioritized for this research prototype. In our case we would utilize signatures for policies and not only for the claims. Thus, the ensuring entity will hardly be able to argue the communicated policy. One thing that is special in this setting is that a signature shall remain valid and verifiable even after the user's data has been deleted. This requires some storage and logging at least of the policies and a "hash value" of the stored personal data. The hash values serves as input for the signature check even if the personal data had been purged.

No	Requirement	Achieved
3	Policies must be easily accessible to users. The way of accessing the policies should be determined by a clear specification.	Yes

This requirement has been met. So far there is no duty to even publish (classical) privacy policies in machine readable way not to mention policies for individual services / actions such as hiring procedures even broken down to details such as retention periods and rights to pass on. This issue also relates to the first reviewers comment after Y2 EC-review of the PrimeLife Project. The goals for solving this issue could look like this: Only make the statements on a "human readable" level of abstraction. Also parse the machine readable version.

No	Requirement	Achieved
4	Policies should be presented to users in an easily comprehensible manner.	Partially

This requirement has been met partially, yet. The policy is written in an XML, which is human-readable but certainly not suitable for the end-user. However, it allows for easy transformation and presentation. A written policy is necessary, but the written policy could additionally be supported by the usage of privacy approaches like privacy icons (cf. PrimeLife deliverable D4.1.5), Privicons (cf. PrimeLife deliverable D4.1.5) or layered policies. These approaches can help to visualize certain aspects of a (written) policy and stress certain purposes for data handling.

No	Requirement	Achieved
5	It must be explicitly determined who is responsible for the policy. This determination must be visible for users.	No

This requirement has not been met, yet. Adding a line to the policy with the name and address of the entity issuing the policy would be sufficient, to meet the requirement. Jurisdiction and exact information on the responsible entity will be part of the user-readable policy (see Requirement 4) anyhow. The name of the entity could be additionally part of the X509 signature.

No	Requirement	Achieved
6	It must be explicitly determined what data are covered by a policy. This determination must be visible for users.	Partially

This requirement has partially been met since the visualization missing. PPL supports to assign different rights to different types of data, e.g. postal address and e-mail address. Strong (cryptographic) binding is not strictly necessary. In fact data may always be copied circumventing policies (see DRM problems) as evil entities will just not implement the necessary enforcement mechanisms. However, the presumption that companies and data controllers want to behave in accordance with the law (compliance) one can assume that also loosely bound policies will be adhered to.

No	Requirement	Achieved
7	Policies should cover all aspects of data processing with regard to privacy legislation.	Partially

This requirement could not be met within the PrimeLife project. The emphasize of the evaluation was on “all aspects”. PPL is able to express many arbitrary purposes. Besides, there is some P3P engine available that creates text from a P3P policy. However, currently, it is hardly possible to fine-grained predefine all possible purposes for data handling in a policy, especially, if data processing could be based on many different purposes that might occur within the same data processing body. Ontology would be needed define all purposes and set them in relation, but so far no product or research project was able to come up with a complete ontology. This topic was evaluated within Work Package 5.2 (cf. PrimeLife Heartbeat H5.2.2).

No	Requirement	Achieved
8	Recipients or categories of recipients to which the data will be passed on to, must be explicitly determined. This must include a reference to the applicable jurisdiction for the recipient.	Yes

This requirement has been met, as the applicant could define a policy for the intended receiver (the employer) and in addition a downstream policy (for the domain expert). In addition to that the implementation allows calling back the user on the mobile device in case of a policy mismatch. A missing part is the explicit reference to the concerning domain experts since they are bound only at runtime. In general it is not possible to disclose all potential controllers by name beforehand without giving away business secrets of the head-hunter (which employers are customers or which experts does a head-hunter work with). However, the head-hunter must declare that the data will be passed on (which also is not a problem here as it is exactly the interest of the customer that her data is passed on to potential employers). For the eCV scenario the requirement is met in relation to the employer whenever a concrete job is applied for, thus the employer is known as a downstream recipient.

No	Requirement	Achieved
9	It should be explicitly determined under what policies data is passed on to other parties.	Yes

This requirement has been met by the demonstrator. In the eCV scenario, a sticky policy is attached to the personal data. PPL allows specifying rules for individual data fields.

3.2.2 Privacy Logging Requirements

No	Requirement	Achieved
10	Log files should be available in an unambiguous formalization and their content should be machine interpretable.	No

This requirement has not yet been met by the demonstrator. Machine interpretability is achieved by determination of a formal language whereas unambiguousness is obtained by an agreement about a joint ontology. Thus, the requirement was addressed for the WP6.3 demonstrator in general but not in the eCV scenario.

No	Requirement	Achieved
11	It must be possible to check the compliance of processing operations with communicated policies on the basis of log files afterwards.	Partially

This requirement has partially been met. The demonstrator uses formalized log files. Partner EMIC experimented in WP6.3 and WP5.3 with formal analysis of log files to verify compliance. Although aiming at cross-domain service compositions, this work has not been applied to the logfiles from the eCV demonstrator though.

No	Requirement	Achieved
12	It must be ensured that log files cannot be argued by their originating entity in charge of the processing.	N/A

Like requirement 2 this this requirement was deliberately not addressed since digital signature of content is a well understood mechanism. PrimeLife addressed privacy-aware logging in Activity 4. This could certainly be integrated within the eCV demonstrator. Yet, both developments were planned as parallel activities.

No	Requirement	Achieved
13	The fact that data are logged must be visible to the user.	N/A

This requirement has not yet been addressed, see Requirement 12 for details. PrimeLife addressed privacy-aware logging in Activity 4. Thus, the requirement can be met with little refinement.

No	Requirement	Achieved
14	The originator of a logging entry must be clearly visible. In particular, it must be visible which service provider a cross-domain service composition is the originator of a certain logging entry.	N/A

This requirement has not yet been addressed, see Requirement 12 for details.

No	Requirement	Achieved
15	A simple methodology must allow access for the users to those logs or parts thereof, to which she has a legal right to access, or to which the service provider wants to grant access to.	N/A

This requirement has not yet been addressed, see Requirement 12 for details. The log files serve as a core repository of relevant data, but need to be supported by access mechanisms. Thus, the requirement can be met with little refinement.

No	Requirement	Achieved
16	It must be clearly visible to what data a log entry refers.	N/A

This requirement has not yet been addressed, see Requirement 12 for details.

No	Requirement	Achieved
17	Log files should describe all contractual and further legally relevant aspects of data processing. Beyond that, technical aspects should only be described in case they are relevant.	N/A

This requirement has not yet been addressed, see Requirement 12 for details.

No	Requirement	Achieved
18	Log files must contain explicit information on recipients or categories of recipients, data have been passed on to. This includes a reference to the applicable jurisdiction.	N/A

This requirement has not yet been addressed, see Requirement for details. Recipients or categories of recipients can be included in as discussed in Requirement 8. Thus, the requirement can be met with little refinement.

3.2.3 Access to primary information

No	Requirement	Achieved
19	Access to personal information should be provided in an unambiguous formalization. The content of the information should be machine interpretable.	Yes

This requirement has been met. Data is communicated as instance of an XML schemes. The user sees the personal data she possesses and even data she communicated to employers. She cannot access data that has already been disclosed though. This is standard technology you find for instance in most web shops today (alter user account, revisit order history).

No	Requirement	Achieved
20	It must be ensured that access to information that has been granted cannot be argued by the granting entity.	Partially

This requirement has partially been met. The data provider sends data to the data consumer along with a customized sticky policy, and cannot dispute that. However, digital signature should be used (compare Requirement 2). Besides, the privacy policy should contain the fact that the data consumer records granting of access.

No	Requirement	Achieved
21	A simple methodology with regard to request and granting of access to information should be provided to users.	Partially

This requirement has been partially met. A standardised interface for subject access requests is supported by every service. This interface is mentioned in the description of the service. For the eCV platform, it is implemented as the eCV platform stores to whom which profile has been sent. For the head-hunter and employer, it is not yet implemented.

No	Requirement	Achieved
22	Users accessing information must be enabled to easily recognize what data are covered by what policy and have been disclosed to what third parties.	Yes

This requirement has been met. The policy is attached to the data set. PPL supports access control for individual data fields. Moreover, the same considerations apply as in requirements 20 and 21.

No	Requirement	Achieved
23	Accessed information should cover only contractual or further legally relevant aspects of data processing.	Yes

This requirement has been met. As repository for accessible data, sticky logs are used. To avoid too much complexity, a filter is applied to the logs that is developed through an expert system and supported by experts with the necessary legal and technical background (see requirement 17).

No	Requirement	Achieved
24	Users must be enabled to access explicit information on recipients or categories of recipients, data have been passed on to. This includes a reference to the applicable jurisdiction.	Partially

This requirement has partially been met. Each data controller can be asked to whom data was passed on in the downstream service chain. For this a log file exists in the eCV demonstrator that can be accessed by the data controller to answer such a request. However, there is no interface for the user to directly access this information. On the other hand some identification / authentication would be required to assure that the requesting person is actually the data subject concerned.

3.2.4 Cross-Domain-specific Requirements

No	Requirement	Achieved
25	It must be possible to maintain communicated policies even if the Service Oriented Architecture is dynamically adapted (refers to the constellation of a SOA being established by several entities).	Yes

This requirement has been met. The policy is attached as sticky policy to the communicated data. Each data consumer stores the policy next to the personal data (e.g. through a link in the database between both entries). Later changes in the service policy will thus not be reflected in already communicated policies. Even when data is passed on, this happens with the sticky policy that refers to the original agreement between data provider and data consumer.

No	Requirement	Achieved
26	It is not possible to maintain (all) communicated policies in case of an adaption of the virtual organization and it must be possible to adapt the communicated policies (builds on requirement 25) through renegotiation. If this fails the service must be stopped.	Yes

This requirement has been met. In case of a mismatch between the user's privacy preferences and the employer's policy, a special notification will be sent to the user. Only if the user overrules her original preference, data will be passed on. This requirement enables users to “renegotiate” original policies. The “negotiation” is a simple overruling of user’s policy by the user herself, so there is no multi-step communication process involved. If the user disagrees with overruling her policy, the processing is stopped. If she agrees to it the new policy is used. The solution is deployed on a specially tailored G&D mobile device developed in WP6.2, but could potentially be adopted to work on other devices, as well.

No	Requirement	Achieved
27	A service provider whose service is a downstream part (those that process data later) of the overall workflow must adhere to policies given by service providers whose service are upstream parts (those that process data first) of the workflow.	Yes

This requirement has been met. In order to achieve that common policies do not have to be negotiated in advance, a mechanism is applied that generates new preferences from existing preferences and policies: At the first service of a workflow user preferences and policies of the service are matched. The result of the matching process is the sticky policy which travels with the data. In the next step of the processing the sticky policy is matched with the policy of the second service. Further information about downstream policies was provided within the work packages 5.2 and 5.3 [CaVSa10], [Tra10]. Moreover PrimeLife deliverable D6.3.2 provides an abstraction of this procedure.

No	Requirement	Achieved
28	Multi-level-matching within a Service Oriented Architecture must be supported.	Yes

This requirement has been met. Multi-level-matching of policies is enabled by the concept of sticky policies. The policy travels with the data. Origin in the user’s preferences the sticky policy will be matched each time before the data is communicated to a new service. The demonstrator supports both PPL’s normal matching and "lazy matching". In case of mismatch the transaction is ended.

No	Requirement	Achieved
29	The ability of the data subject to have access to information must be ensured for the future.	No

This requirement has not been met, yet. That would require history mechanism and a federation/delegation mechanism. We address the history aspect in the Abstract Privacy Lifecycle (D6.3.2).

No	Requirement	Achieved
30	A ex post notice must be enabled by appropriate mechanisms.	Yes

This requirement has been met. Policies stick to data already disclosed, thus a change of policies only impacts future data disclosure. Besides, the user has to be informed about changed policies, no matter, what partial service of the eCV platform (or of the head-hunter) she uses. This can be achieved by writing a policy that obliges a data provider to send a notification to the data subject when her data is being sent. The obligation enforcement engine defines the trigger “TriggerPersonalDataAccessedForPurpose” for this.

3.2.5 Additional Mechanisms

No	Requirement	Achieved
31	It must be ensured that correction and erasure of users' data are feasible.	Partially

This requirement has partially been met. Customer data and logs are stored in a database or a data format that allows manipulations by the data controller. It is not necessary that the user is able to trigger the deletion himself (cf. Requirement 33). It is sufficient if the user can ask the data controller to trigger the deletion. The data controller can accomplish this in the database. The communication is through an out-of-band communication.

No	Requirement	Achieved
32	It must be ensured that blocking of user data is feasible.	Partially

This requirement has partially been met. The obligation enforcement engine allows for time based deletion. But a deletion upon request is only possible through out-of-band communication. However, sine requirement 31 has partially been met there is no need to meet this requirement, too.

No	Requirement	Achieved
33	It should be made easy for users to exercise their rights of correction, erasure and blocking.	No

This requirement has not been implemented, yet. This requirement is based on legal guidelines (cf. § 35 of the German federal data protection act (*Bundesdatenschutzgesetz*)). Within the eCV scenario it would be possible to allow applicants to withdraw their application - triggering a deletion of their data. Also updating their CV until the deadline of the job offer might be possible (made with credentials such as login for a customer account). Although the requirement is not yet met, it is still possible to retreat from a job application by requesting deletion (see requirement 31).

No	Requirement	Achieved
34	It should be possible to guarantee compliance with communicated policies.	Yes

This requirement was met. The use of DRM-like-mechanisms could guarantee control of data usage and compliance with communicated policies. In the demonstrator, the obligation enforcement engine is used for the obligation part. EMIC experimented with formal reasoning on collected policies and logging data (cf. Requirement 11). This allows the data controller to verify if all promised obligation have been executed. However, in any case an intention to comply with policies and the law is necessary on part of the data controller.

No	Requirement	Achieved
35	There should be a possibility to support trust between user and service provider.	No

This requirement has not been met, yet. There are no trust-establishing mechanisms and no reputation management systems. But here third party systems appear to be better trust mediators, which was not within scope of the conducted research. A clear identification of the service provider as is required for German service providers wanting to access the new German eID could be an approach to address this requirement. Getting certificates from a third party could ensure trustworthiness.

No	Requirement	Achieved
36	The user shall have the possibility to express her preferences in an easy manner.	Yes

This requirement has been met. A provision of a well-defined ontology that is limited to concepts is necessary for the explanation of users' preferences. The requirement is addressed by nature of the scenario and PPL. Nevertheless, improvements are possible in the areas HCI and user guidance. PrimeLife Activity 4 investigated how policies could be made human readable. More information to this topic is provided (cf. PrimeLife deliverable D4.1.5).

No	Requirement	Achieved
37	User and service provider should be able to match the preferences and related policies.	Yes

This requirement has been met. The policy matching engine, which is part of the PPL policy engine, allows for matching of policy and preferences. The matching of policies is one of the core elements for what the obligation engine has been built for.

No	Requirement	Achieved
38	Matching of preferences and policies must be comprehensible.	Yes

This requirement has been met. The result of match and mismatch of preferences and policies is presented to the user. This work package investigated ways to visualize policy mismatches and ways to mitigate them.

No	Requirement	Achieved
39	A mechanism to express the anonymity set with regard to a specific data type should be supported.	N/A

This requirement is not addressed by the scenario. The nature of an job application scenario excludes anonymity.

3.3 Findings of the evaluation

Table 1 summarizes the assessment of each requirement as “fully achieved”, “partially achieved”, “not achieved” or “not applicable”. Please refer to the individual assessment given above for reasons why a particular requirement has been marked respectively. From a high-level perspective the table tells a number of things.

First and foremost, most of the requirements have been met. The demonstrator was thus well-chosen and the implementation was capable to show most of the interesting scenarios.

The logging requirements are mostly rated as “not applicable”. This is because secure logging did not play any role in the technical realization of the eCV scenario. Activity 1 did work on this topic, but couldn’t be regarded in this demonstrator since the work had been progressed to far already when the secure logging results became available. However, the generalization of our thinking about privacy-aware service composition, the “Abstract Privacy Framework”, features components for secure logging and even history information about data disclosure. That should be sufficient to achieve most of the requirements in this group.

Third, all the cross-domain specific requirements have been met except one that is linked to the reasoning above. This result is very positive, since the data controlling in downstream scenarios was the major aspect both in this demonstrator and in the PrimeLife Policy Language (PPL). Both achievements are strongly interwoven through intensive collaboration between the work-packages in course of the project. The key takeaway is, that the demonstrator enhances privacy even in multi-party, multi-domain service compositions.

Forth, core policy requirements have been mostly met as well. The demonstrator successfully shows how polices can be communicated in an unambiguous and well-defined form. They are rich and detailed enough to express all necessary facts.

Many of the not achieved or partially achieved requirements deal with user interfaces. Although Activity 6 collaborated with Activity 4 (HCI), the focus of the demonstrator was not on the presentation layer for policies. The demo conveys all the necessary information to the user – as stated by achieving the respective requirements – but we do not present them adequately. In fact, the user has to read and understand PPL, which is an XML based language that is (although human-readable) not self-explanatory to the majority of users.

Finally, the assessment of the assisting mechanisms – referring both to the third and fifth group of requirements – is indifferent. Most requirements have been met, but many are marked as partially achieved. This is again because we evaluated the concrete implementation of the demonstrator. The generalization of our work, the “Abstract Privacy Framework”, took more of these partially achieved requirements into account.

All in all the evaluation showed that the WP 6.3 demonstrator achieves the goal of being a working system that is in line with the privacy requirements developed. The current status of the demonstrator also illustrates that the requirements for privacy-enhancing Service-oriented architectures defined in [MeS09] are application oriented and not too difficult to be addressed.

Table 1: Overview of Assessment

Requirement			Achieved in Scenario			
Group	**No**	**Short description**	**Yes**	**Partially**	**No**	**N/A**
Core Policy Requirements	1	Unambiguous formalization	X			
	2	Nonrepudiation			X	
	3	Accessibility to users	X			
	4	Comprehensibility		X		
	5	Responsibility			X	
	6	Covered data		X		
	7	Level of detail		X		
	8	Recipients	X			
	9	Downstream policies	X			
Privacy Logging Requirements	10	Unambiguous fomalisation	X			
	11	Compliance checking		X		
	12	Nonrepudiation				X
	13	Logging made visible to user				X
	14	Originator of logging				X
	15	Accessibility to users				X
	16	Log Reference				X
	17	Covering relevant aspects				X
	18	Log recipients				X
Access to primary information	19	Unambiguous fomalisation	X			
	20	Nonrepudiation for granting access		X		
	21	Simple methodology		X		
	22	Link between data and policy	X			
	23	Data Minimization	X			
	24	Information about recipients		X		
Cross-Domain-specific Requirements	25	Maintain communicated policies	X			
	26	Renegotiation of policies	X			
	27	Adhere to upstream polices	X			
	28	Multi-level matching	X			
	29	History			X	
	30	Ex post notice	X			
Additional Mechanisms	31	Correction and erasure		X		
	32	Blocking of user data		X		
	33	Exercise of user rights			X	
	34	Guarantee compliance	X			
	35	Trust establishment			X	
	36	User preferences	X			
	37	Match preferences and policies	X			
	38	Comprehensible matching	X			
	39	Anonymity				X

Chapter 4

Closing remarks

This deliverable is a logical follow-up from project heartbeats/deliverables in work package 6.3, most notably H6.3.1 which outlined requirements for privacy-enhancing SOAs, the D6.3.2 which explains the technical details, and the two demonstrator heartbeats H6.3.2 and H6.3.3. It was also foremost in the partners' minds to build a bridge to the important work that was carried out in Activity 5 with the development of the Privacy Policy Language (PPL) and engine. Early discussions on the demonstrator were instrumental in providing requirements to the Activity 5 and subsequently had an impact on shaping their work particularly in the area of downstream usage.

The eCV demonstrator represents a realistic situation that includes actors from an employability situation such as Applicant, Employer, Headhunter and Legal Domain Experts. This type of scenario lends itself to legal/privacy issues on how data is handled and subsequently provides interesting challenges for service and policy composition. An applicant looking for a job needs to be provided which may or may not be of a sensitive nature; this person would like some assurances on how data will be treated. On the other side an employer needs to be able to recruit and verify the suitability of candidates but is also obliged to show how data will be handled.

The evolution of the scenario was influenced in many ways by the nature of the partnership structure; with EMIC, SAP and G&D being separate legal entities it was clear from the start that this would shape the overall structure of the cross-domain service composition both from a scenario and technical perspective. The clear division of tasks and application development ensured rich collaboration amongst the partners and enabled a role-playing environment.

The requirements are derived from an analysis of two central regulatory instruments within the European Union, namely the data protection directives 95/46/EC and 2002/58/EC. This analysis had previously been conducted in the context of SOA in Virtual Organization by project partners ULD who are also the evaluators of the eCV demonstrator.

Table 1 illustrates how the requirements have been grouped into five categories, namely, Core Policy, Privacy Logging, Access to Primary Information, Cross-domain specific and additional mechanisms. It was clear from the beginning of the project that not all of the requirements were achievable and consequently a decision had to be made as to what was feasible within the given project timeframe. This does not mean in a way that the requirements were selected on the basis of relevance to privacy, however the decision was based more on what the partners wanted to achieve from a technical perspective, namely, a data-centric service composition permitting the use of privacy policies. Furthermore topics not addressed were being investigated in parallel activities in the project and the integration at a late stage was not practical.

For example, the privacy logging group of requirements were not in scope and were dealt with to some extent in the context of Activity 1. Also anonymization in the additional mechanisms group and history in the cross-domain specific requirements were also not requirements that were addressed, however the former was dealt with in Activity 2.

Most of the requirements were however met and the evaluation highlights, in each instance, the reasons why this is the case. Where the requirements were deemed partially met or not met, the main instigators in the development were at hand to discuss with the evaluators the reasons why this was the case and to provide possible solutions for future work. It was also found that some requirements were not clearly defined and were therefore not able to be assessed.

The process of evaluating the requirements with the legal and technical expertise of ULD provided a means to complete the work that started at the beginning of the project. It assisted it showing the project partners how much variance, if any, existed with the original goals of the service composition work package. It is felt that it would have been more beneficial to have had a review cycle at more regular intervals during the development process.

Overall the assessment illustrates that the demonstrator has been closely aligned with the requirements and that the demonstrator has made good use of the PPL language and API of the PPL engine emanating from Activity 5 and the work carried out influenced the requirements of this Activity. Furthermore the evaluation is the first validation of PPL language in a “real” use case.

References

[CaVSa10] De Capitani di Vimercati, S. and Samarati, P.: Second research report on next generation policies; Deliverable D5.2.2 (2010) of the PrimeLife Project consortium, available at http://www.primelife.eu/images/stories/deliverables/d5.2.2-second_research_report_on_policies-public.pdf (2010)

[MeS09] Meissner, S. and Schallböck, J.: Requirements for privacy-enhancing Service-oriented architetcures; Heartbeat H6.3.1 (2009) of the PrimeLife Project consortium, available at http://www.primelife.eu/images/stories/deliverables/h6.3.1-requirements_for_privacy_enhancing_soas-public.pdf (2009)

[Pri09a] Pinsdorf, U. and Sommer, D.: Infrastructure for Privacy for Life; Deliverable D6.3.2 (2011) of the PrimeLife Project consortium, available at http://www.primelife.eu/results/documents (2011)

[Pri11b] Koschinat, S., Rannenebrg, K. and Bal, G.: Identity Management Infrastructure Protocols for Privacy-enabled SOA; Deliverable D6.1.1 (2009) of the PrimeLife Project consortium, available at http://www.primelife.eu/results/documents (2009)

[Tra10] Trabelsi, S.: Second release of the policy engine; Deliverable D5.3.2 (2010) of the PrimeLife Project consortium, available at http://www.primelife.eu/images/stories/deliverables/d5.3.2-second_release_of_the_policy_engine-public.pdf (2010)