Greater Control and Transparency in Personal Data Processing

Havur, Giray; Vander Sande, Miel; Kirrane, Sabrina

Link to publication

Greater Control and Transparency in Personal Data Processing

Giray Havur[1,3][a], Miel Vander Sande[2][b] and Sabrina Kirrane[1][c]

[1]*Institute for Information Business, Vienna University of Economics and Business, Austria*
[2]*IDLab, Ghent University — imec, Belgium*
[3]*Corparate Technology, Siemens AG Austria, Austria*

Keywords: Usage Control, Consent, Transparency, Compliance, Trust, Decentralisation.

Abstract: Although the European General Data Protection Regulation affords data subjects more control over how their personal data is stored and processed, there is a need for technical solutions to support these legal rights. In this position paper we assess the level of control, transparency and compliance offered by three different approaches (i.e., defacto standard, SPECIAL, Solid). We propose a layered decentralised architecture based on combining SPECIAL and Solid. Finally, we introduce our usage control framework, which we use to compare and contrast the level of control and compliance offered by the four different approaches.

1 INTRODUCTION

The European General Data Protection Regulation (GDPR) is a game changer in terms of personal data management. In particular, the legislation affords *data subjects* control and transparency with respect to the processing of their personal data by *data controllers/processors* (i.e., product/service providers).

When it comes to GDPR compliance there are a variety of questionnaire based tools that enable data controllers/processors to assess the compliance of their products/services (cf., (Information Commissioner's Office (ICO) UK, 2017; Microsoft Trust Center, 2017; Nymity, 2017; Agarwal et al., 2018)). At the same time, researchers are looking into using technical solutions in order to: (i) enable data subjects to specify consent at a fine level of granularity; and (ii) make it possible to automatically check compliance of existing products and services with respect to the data subjects consent (cf., Bonatti and Kirrane (2019)).

In this position paper, we explore how technology can be used to provide stronger guarantees to data subjects with respect to the processing of their personal data. We start by defining a motivating scenario, which is subsequently used to examine the consent, transparency and compliance guarantees offered by three alternative approaches, namely: (i) the defacto standard where data subjects consent to very general processing by product/service providers; (ii) SPECIAL[1] which empowers data subjects by offering them flexible consent mechanisms and greater personal data processing transparency and compliance; and (iii) Solid[2] which decouples data from applications thus enabling data subjects to decide where their personal data resides and who gets access to this data.

Summarising our contributions, we: (i) provide a summary of existing policy languages, transparency and compliance techniques; (ii) assess the level of control, transparency and compliance offered by three different approaches in the context of our motivating scenario; and (iii) propose a control and compliance framework that can be used to assess different data processing and sharing architectures.

The remainder of the paper is structured as follows: Section 2 describes our motivating scenario and the corresponding high level requirements. Section 3 presents the necessary background. Section 4 examines three alternative personal data management approaches. Section 5 demonstrates how SPECIAL can be implemented in Solid. Section 6 proposes a framework for evaluating different personal data processing architectures. Finally, conclusions and directions for future work are outlined in Section 7.

[a] https://orcid.org/0000-0002-6898-6166
[b] https://orcid.org/0000-0003-3744-0272
[c] https://orcid.org/0000-0002-6955-7718

[1]https://www.specialprivacy.eu/
[2]https://solid.mit.edu/

2 MOTIVATION

We start by describing a concrete motivating scenario and the requirements used to guide our research.

Use Case Scenario. A fictitious company called BeFit (i.e. the data controller) is a producer of wearable appliances for fitness tracking. The device records parameters such as steps taken, active/inactive minutes, location, etc. In addition, the device can be used to monitor food and drinks consumed. The device owner (i.e. the data subject) uses the device in order to track activity, record workouts, and manage weight gain/loss. When it comes to data processing and collection, there are three specific purposes that we focus on in this paper:

(i) *Service Provision:* all of the data gathered by the Befit device is backed up on a server and is used to provide activity information to the device owner via the Befit fitness dashboard;

(ii) *Personal Data Sharing:* the device owner chooses to share data collected by the device with friends, followers or the general public via a third party social fitness network; and

(iii) *Secondary Use:* the device owner consents to their data being used by BeFit in order to optimise existing and future products and services.

Requirements. In order to enable scenarios such as those described above, within the context of the GDPR, the following three key requirements need to be facilitated:

Consent: BeFit needs to be able to specify what data is *desired* for which purposes. While at the same time the device owner needs to be able to specify which data should be used for which purposes.

Transparency: The device owner should be able to determine what data is collected, what processing is performed, for what purpose, where the data is stored, and with whom it is shared.

Compliance: When it comes to personal data processing a company needs to show that they are compliant with the device owners consent.

3 BACKGROUND

From a usage control perspective there are three broad bodies of research that need to be considered: (i) machine interpretable policy specification; (ii) personal data processing transparency; and (iii) compliance verification.

Consent. The traditional way to obtain consent is to ask for consent for all current and future personal data processing outlined in very general terms by clicking on an agree button. Acquisti et al. (2013) highlight that several behavioral studies dispute the effectiveness of such consent mechanisms from a comprehension perspective. A study by McDonald and Cranor (2008) indicates it would take on average 201 hours per year per individual if people were to actually read existing privacy policies.

In order to be able to support automated compliance checking it is necessary to encode consent in a manner that is interpretable by machines. Here policy languages play a crucial role. Over the years, several general policy languages that leverage semantic technologies (such as Rei (Kagal et al., 2003) and Protune (Bonatti and Olmedilla, 2007)) have been proposed. Such languages cater for a diverse range of functional requirements, such as access control, query answering, service discovery, and negotiation, etc. More recently the SPECIAL project has proposed a Description Logic based policy language that can be used to express consent, business policies, and regulatory obligations (Bonatti and Kirrane, 2019).

Sticky policies enable data providers to define polices (i.e., preferences and conditions) that state how their data can be used. For instance, a sticky policy can be used to govern data usage, for instance the purposes of the data use, whitelists and blacklists, obligations for data consumers, notification requirements, deletion periods, and trust authorities (Beiter et al., 2014). Given that data is initially sent in an encrypted form, encryption techniques play an important role in the sticky policies paradigm. Two different works (Tang, 2008; Beiter et al., 2014) summarise various encryption techniques used in sticky policy enforcement mechanisms.

Transparency. From a transparency perspective, Bonatti et al. (2017) identified a set of criteria that are important for enabling transparent processing of personal data at scale, and summarise existing literature with respect to the proposed criteria. Several of these works use a secret key signing scheme based on Message Authentication Codes (MACs) together with a hashing algorithm to generate chains of log records that can be used to ensure log confidentiality and integrity (cf., Bellare and Yee (1997)). When it comes to personal data processing, Sackmann et al. (2006) demonstrate how a secure logging system can be used for privacy-aware event encoding. In particular, they introduce the "privacy evidence" concept and discuss how logs can be used to ensure that privacy policies are adhered to. While, Pulls et al. (2013) propose a protocol, which is based on MAC secure logging

techniques, that can be used to ensure both confidentiality and unlinkability of events.

Compliance. From a GDPR compliance perspective, recently the British Information Commissioner's Office (ICO) (Information Commissioner's Office (ICO) UK, 2017), Microsoft (Microsoft Trust Center, 2017), and Nymity (Nymity, 2017) have developed compliance tools that enable companies to assess the compliance of their applications and business processes by completing a predefined questionnaire. When it comes to automatic compliance checking there is a large body of work that focuses on modelling and reasoning over legal requirements using semantic technologies (cf., (Boer et al., 2008; Bartolini et al., 2015; Pandit et al., 2018)). For instance, (Palmirani et al., 2011; Athan et al., 2013) demonstrated how LegalRuleML can be used to specify legal norms. More recently Bartolini et al. (2015) and Pandit et al. (2018) propose ontologies that can be used to model data protection requirements in a manner that supports compliance verification. While, De Vos et al. (2019) demonstrate how business policies and legal requirements can be represented using a flavor of ODRL, and checked automatically via the Institutional Action Language language.

4 ALTERNATIVE APPROACHES

Next we discuss different approaches to personal data management guided by our motivating scenario.

The Defacto Standard Approach. In the vast majority of cases when a user signs up for a new product/service the company presents them with a document where all possible current and future personal data processing is described in very general terms, and a checkbox that needs to be ticked in order to use the product/service.

Consent: The GDPR defines several potential legal bases (consent, contract, legal obligation, vital interest, public interest, exercise of official authority, and legitimate interest) under which companies can legally process personal data. In terms of consent companies should ask for consent if the data required goes beyond what is needed for other legal bases.

Transparency: The GDPR empowers data subjects with the right to obtain a copy of all personal data that a data controller/processor has concerning them. Following best practice companies should be transparent with respect to the information that will be collected for which purposes.

Compliance: The GDPR provides a legal framework for data subjects to lodge complaints with a supervisory authority if their rights have been infringed.

Considering common practices when it comes to handling personal data, in the standard case our BeFit use case scenario could be implemented as follows:

(i) Service Provision: BeFit should provide transparency with respect to the processing performed by the device by offering the device owner the ability to opt into all data processing that is necessary in order for the fitness device to function.

(ii) Personal Data Sharing: In order to benefit from existing cloud based analytic services the device owner would also need to opt into the 3rd party analytic service providers privacy policy and their terms and services. If integration with the desired third-party service is not possible the device owner can resort to data subject access requests to download their data such that it can be uploaded to the analytic service providers website.

(iii) Secondary Use: Article 5 of the GDPR states that personal data that is collected for specified, explicit and legitimate purposes should not be further processed in a manner that is incompatible with said purposes, unless there is a legal basis for doing so. For this reason it has become common practice for companies to ask for consent for secondary use separately.

The SPECIAL Approach. The SPECIAL platform, which is routed in Semantic Web technologies and Linked Data principles: (i) supports the acquisition of data subject consent and the recording of both data and metadata (consent, legislative obligations, business processes) as policies; (ii) caters for automated transparency and compliance verification; and (iv) provides a dashboard that make personal data processing comprehensible for data subjects, controllers, and processors.

Consent: The SPECIAL project has developed and evaluated several alternative consent user interfaces that enable data controllers to ask for consent for particular data points to be processed for explicitly stated purposes. The consent is subsequently translated into machine understandable policies (i.e., what data is collected, for which purposes, what processing is performed, where they data are stored for how long and with whom it is shared) that are encoded using the SPECIAL Policy Language[3].

Transparency: The SPECIAL log vocabulary[4] enables companies to record all data processing/sharing

---

[3]http://purl.org/specialprivacy/policylanguage
[4]http://purl.org/specialprivacy/splog

performed within their company. The log vocabulary builds upon the SPECIAL policy language ontology and reuses well known vocabularies such as PROV[5] for recording provenance metadata. While, the SPECIAL dashboard provides a uniform interface to let data subjects exercise their rights (i.e., access to data, right to erasure, etc.).

Compliance: The SPECIAL project supports three different types of compliance checking: (i) the data processing which a company would like to perform complies with the data subjects consent (i.e., ex-ante compliance checking); (ii) all personal data processing performed by the companies products and services are stored in an event log (i.e., the SPECIAL ledger) which is subsequently checked against the data subjects consent (i.e., ex-post compliance checking); and (iii) business processes are recorded as sets of permissions and checked against regulatory obligations set forth in the GDPR (i.e., business process compliance testing).

Our BeFit use case scenario could be implemented in SPECIAL as follows:

(i) *Service Provision:* The SPECIAL consent interface could be used to obtain fine grained consent for specific processing, for instance to derive calories burned, display route on map, back up data, etc. While the SPECIAL dashboard could be used to provide transparency with respect to the data processing performed on the device.

(ii) *Personal Data Sharing:* At the request of the device owner BeFit could share data with existing cloud based analytic services (e.g., Runkeeping and Strava). A sticky policy could in turn be used to tightly couple usage constraints and the personal data that it governs.

(iii) *Secondary Use:* At the request of the device owner BeFit could use the device owners personal data for the secondary purpose of improving BeFit's products and services. The device owner would have full transparency with respect to this processing and could elect to opt out at any point in the future.

The Solid Approach. The term "Social Linked Data" or *Solid* (Sambra et al., 2016) refers to a recommended set of tools, best practices, and predominately W3C standards and protocols, to build decentralised social applications based on Linked Data principles. Its main premise is establishing *pod-centric* platforms: data subjects maintain a personal domain and associated data storage, i.e. a data pod, from which they give applications permission to read or write personal data. The pod provides a set of

personal Web APIs, an identity provider using WebID, and an inbox to receive notifications based on Websockets or Linked Data Notifications (LDN) (Capadisli et al., 2017). Because of its open ecosystem and progressive stance, Solid is able to attract a significant developer community for improving the standards and tools, and building more applications.

Consent: Solid extensively decouples data from services, thus increasing the user's control over personal data, enhancing the mobility of data between services and lowering data duplication overall. The interoperability through Linked Data standards and protocols (i.e., the Resource Description Framework (RDF) data model) ensures any pod can provide data to any service application. A basic consent mechanism is present in the Solid pod as an access-control list (ACL), where data providers agree to let an application read or write certain resources that reside in their pod.

Transparency: Solid is able to achieve full transparency on primary data access: the data pod owner has a complete view on *who* is reading or writing *what* data and *when*, and whether they have the permissions to do so. These activities can be recorded by applying any log vocabulary; Solid currently does not provide a default one. Like in other approaches, transparency on secondary data access or data processing, i.e. data that does not directly originate from the data pod because it was copied, cached or inferred, requires additional measures. However, Solid does provides a notification system that enhances the implementation of transparency, for instance by alerting data subjects about how their data is used.

Compliance: Solid offers the standards necessary to connect to data pods, retrieve their data and use them. Ex-ante compliance checking is performed by enforcing the ACL rules: data access that is directly non-compliant will be blocked. Ex-post compliance checking can be performed by inspecting the pod's ACL log for patterns of misconduct.

Our BeFit use case scenario could be implemented in Solid as follows:

(i) *Service Provision:* The device owner owns a Solid pod, in which all data captured by the device is stored. In order to use BeFit's fitness dashboard application, the device owner registers with their pod. Upon registration, BeFit requests access to the personal data captured by the device and advises on the intended use. In the pod's management dashboard, the device owner can decide to grant or deny access, and specify the applied policy.

(ii) *Personal Data Sharing:* With the device's data residing in the device owner's pod, they can be shared
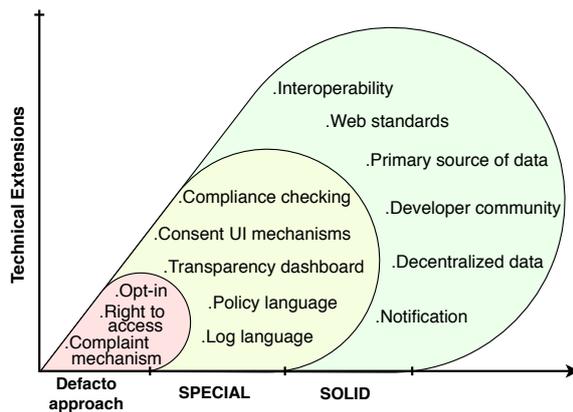
---

[5]https://www.w3.org/TR/prov-o/

Figure 1: Technical extensions.

with any existing service without the approval or interference of BeFit. The sharing process would be identical to the process for BeFit's fitness dashboard.

(iii) *Secondary Use:* BeFit can request additional data or intended use such as optimising existing and future products, when the device owner registers the pod or by sending a notification to the pod's inbox. In the pod's management dashboard, the device owner could opt-in or opt-out to these changes. When it comes to additional data access, full transparency is covered through the ACL log.

5 STRONGER GUARANTEES

In this section, we discuss how the consent, transparency and compliance methods from the SPECIAL project could be implemented in a Solid environment, resulting in great control and transparency. From a compliance perspective, we touch upon mechanisms that could establish trust between data providers and service providers in such a scenario.

Applying SPECIAL in Solid. In Section 4, we discussed three scenarios for personal data sharing and processing, each with a specific *technical architecture* and distinctive *consent, transparency and compliance* mechanisms. However, the layering of these approaches, as depicted in Figure 1, could in fact provide data subjects with stronger guarantees on personal data processing. The defacto standard approach sets the baseline where the data subject has legal guarantees originating from the GPDR: opt-in consent, the right to access, and a complaints mechanism. The SPECIAL project extends this scenario with machine understandable policies and logging, methods for automated compliance checking, a dashboard for transparency, and more control from a consent perspective.

With Solid, the mechanisms above can be embedded as follows: personal data now resides in a Solid data pod as part of a decentralized Web-based ecosystem under the full control of the data subject. This decoupling of applications and data provides data subjects with more leverage to co-determine the data usage policy. Service providers are not granted data access before the policy is decided upon, and moving data between services is significantly easier, allowing unsatisfied data subjects to go elsewhere. In addition to this paradigm, Solid offers the means to implement the SPECIAL consent, transparency and compliance mechanisms in an open, pod-centric, and decentralized Web environment: standard data exchange Web protocols, a notification system, open-source software and a growing developer community.

Figure 2 shows a possible implementation of SPECIAL using Solid. The service providers' applications and the data subjects' pods both adopt: (i) the *SPECIAL policy language* to increase the granularity of Solid's ACL-based access control; and (ii) the *SPECIAL log vocabulary* to record data usage events. Data subjects can register their personal data pod with applications from different service providers. Both parties use the policy language to decide on the policy to apply to the data, possibly with help of the SPECIAL's consent user interfaces. Consent can be given after a policy negotiation phase: (i) the service provider expresses the policy it desires in exchange for its service; (ii) the data subject responds with what is acceptable; (iii) after mutual agreement the data subjects consent is materialized as a sticky policy, and the application is granted access to the data. A Solid pod uses the log vocabulary to record all read or write activity in a local log, and a Solid application uses the log vocabulary to record all data retrieval, processing, and sharing activities. Depending on the trust mechanism in effect (more info on this later) the former is stored in the SPECIAL ledger, a local log, a distributed log, or with another trusted third-party.

Over plain Solid, the integration of SPECIAL offers service providers the means to communicate any further data processing (policy language and ledger); and the device owner has the means to monitor and manage data usage from its Solid pod in a comprehensible way (consent user interface and dashboard). The used policy and logs feed SPECIAL's automatic ex-post compliance checking process. A cross-check of logs from both parties can discover inconsistencies with the policy and thus detect compliance failure early. The ex-ante checking process can be adopted when Solid applications also describe their business logic, which can be displayed in the dashboard as well as enhance transparency even further. Finally,
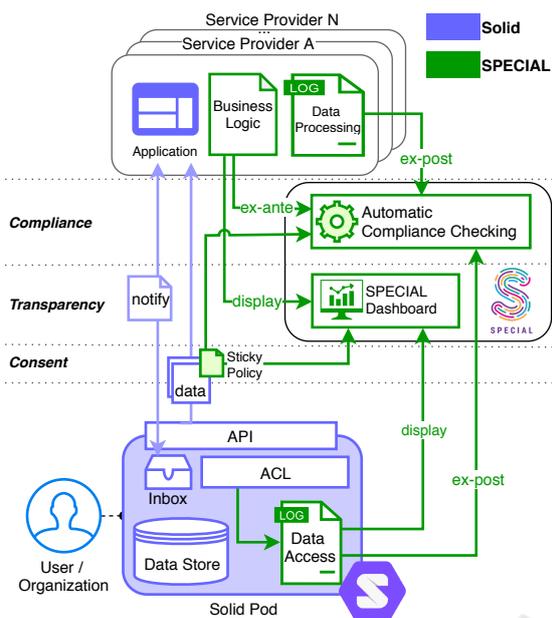
Figure 2: SPECIAL applied to Solid.

Solid's notification system allows a continuous interaction between data pods and applications, facilitating later changes such as policy updates, novel types of data usage, and opting out.

Trust Mechanisms. There are certain types of data use where strong compliance guarantees can only be given by means of *trust*. Therefore, we discuss several trust mechanisms adopted from (Jøsang et al., 2007) in the context of an open Web environment, with the purpose of obtaining *the degree of trustworthiness data subjects assign to a data controller/processor for adhering to a jointly agreed upon data usage policy*. Most existing works on trusted environments introduce strong ties to the OS and hardware layers, which makes them very applicable to centralized and distributed computing (Azzedin and Maheswaran, 2002). However, they loose most benefits and guarantees when this ecosystem is opened up, such as Solid. Protocols and environments with far-reaching trust guarantees such as Trusted Computing (Mitchell, 2005) even directly contradict an open platform and are criticized for encouraging vendor lock-in (Oppliger and Rytz, 2005). Within open, distributed and decentralized multi-agent systems like the Web and Solid, trust mechanisms are generally limited to softer guarantees in exchange for interoperability (Cofta, 2018). From the works herein, Pinyol and Sabater-Mir (2013) identify three main approaches:

Security or policy-based approaches rely on cryptography and digital signatures to ensure basic guarantees such as the authenticity and integrity of a specific party, which presumably leads towards trust. This approach, which is already adopted by the sticky policies from SPECIAL, does not increase trust in data usage. Hence, additional trust mechanisms are required.

Institutional approaches require a centralised third-party to reward or punish parties according to their reported activities. According to Golbeck (2006), these are most valuable in smaller data subject-processor subnets. For socio-economic and technical difficulties, it seems unlikely that Solid will ever span a Web-scale network. Instead, it is likely that data pods will be part of many small Solid subnetworks formed around a certain (type of) application, driven by the network effect (Hendler and Golbeck, 2008). For highly regulated applications with rather static user-bases such as banking, it is legitimate for data pods to trust a single auditing institution.

Social Social approaches establish trust based on past interactions qualified by the data subject or its peers (Bonatti et al., 2005), often coined as *referral trust* (Artz and Gil, 2007). One example involves building certificate chains to form a "Web of trust" (Backes et al., 2010). A member expresses belief in another member by singing his public key. This belief is transitive, therefore a member can trust public keys by verifying the existence of a chain. (Backes et al., 2010) show that this can also be done in anonymity by using non-interactive zero-knowledge proof of knowledge. Reputation-based approaches add *weightings* to referral trust to establish a softer, less binary decision. Wahab et al. (2015) identifies four common models: Feedback-based models calculate a trust value based on user reviews based on quality of service metrics. Major challenges are how to bootstrap trust in a new or modified network and to ensure the quality and credibility of reviews. Statistics-based models combine multiple sources of trust with objective statistical methods (e.g., a Bayesian network or PageRank); fuzzy-logic based models combine subjective feedback and objective quality measurements to indicate trust without computing a final trust value; and data-mining-based models use text mining to analyze reviews assuming user reviews are always credible. Because Solid redistributes leverage (i.e., the policy and data are not under control of the service provider by default), social trust approaches are significantly more powerful. Especially where a bad reputation can lead to data providers moving their Solid pod to a competing service.
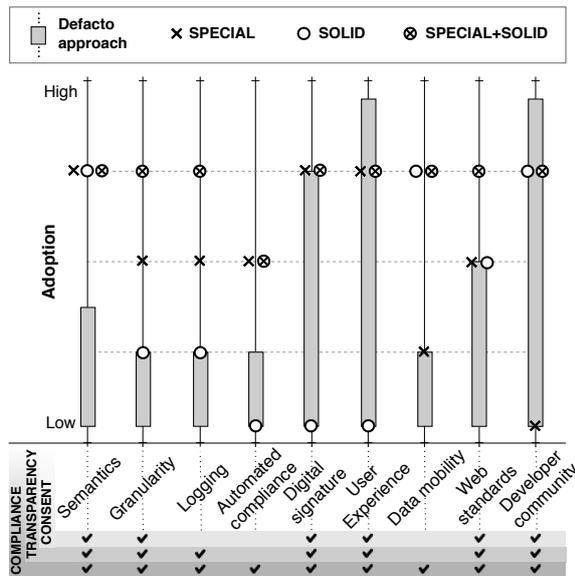
Figure 3: Usage Control Framework.

6 ASSESSMENT FRAMEWORK

Finally we propose a usage control framework that can be used to assess different data processing and sharing architectures. The framework consists of nine features whose adoption in architectures would provide stronger guarantees in terms of consent, transparency, and compliance. Aside from the *web standards* and *developer community* features (both of which are necessary for pushing this research agenda forward) all other features were derived from the literature presented in Section 3.

Semantics enables the machines to understand the data, which eases automation, data integration and interoperability across applications. Moreover, concepts, relationships between things and categories of things can be defined in semantic data models. Enabling such models for policies and logs would help machines not only to enforce the policies automatically but also to record the data access and usage related log entries in a machine-processable way.

Granularity refers to the levels of detail carried for describing policies and log records.

Logging is the act of recording events related to the execution of business logic and to access and usage of personal data.

Automated Compliance is the ability to automatically adhere to policies that are defined by data subjects and legal regulations while accessing or processing personal data.

Digital Signatures are mathematical schemes for ver-

ifying the authenticity of the source of data and for ensuring the integrity of data. These signatures support sticky policies and the integrity of the log.

User Experience (UX) deals with human cognitive limitations by improving human-computer interaction and system aspects such as utility, ease of use and efficiency for consent and transparency management.

Data Mobility is immediate and self-service access to personal data with regard to right of access defined in GDPR.

Web Standards are the formal, open standards and other technical specifications that define and describe aspects of the World Wide Web. The use of such standards would potentially standardize compliance checking mechanisms and facilitate transparency across corporate boundaries.

Developer Community is a group of programmers that are supported by APIs and proper documentation so that they can contribute to development efforts. The more the web standards are followed in an architecture the easier it would get for the developers to take develop applications.

In Figure 3, the y-axis denotes the degree of adoption of the described features by four different data processing architectures portrayed in Sections 4 and 5. The lower end of the figure categorizes the features with respect to their relation to the key requirements derived for usage control in Section 2.

7 CONCLUSIONS

SPECIAL affords data subjects more control over how their data is used, however given the data resides on the company servers SPECIAL assumes that they are working with companies who want to demonstrate compliance. Solid potentially provides for the greatest degree of control in terms of policy specification, however the enforcement of usage control in a decentralised setting is still an open research challenge. Thus, a combination of both complimentary approaches and a suitable mechanism to establish trust between parties, could provide a solid base for building environment with strong data usage control and compliance. In future work we plan to demonstrate how the Open Digital Rights Language can be used to specify Solid usage policies and to support negotiation between data producers and consumers and enhancing the Linked Data platform protocols with policy exchange and negotiation mechanisms.

ACKNOWLEDGEMENTS

REFERENCES

Acquisti, A., Adjerid, I., and Brandimarte, L. (2013). Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security & Privacy*, 11(4).

Agarwal, S., Steyskal, S., Antunovic, F., and Kirrane, S. (2018). Legislative compliance assessment: Framework, model and GDPR instantiation. In *Annual Privacy Forum*.

Artz, D. and Gil, Y. (2007). A survey of trust in computer science and the Semantic Web. *Web Semantics*, 5(2).

Athan, T., Boley, H., Governatori, G., Palmirani, M., Paschke, A., and Wyner, A. Z. (2013). Oasis Legal-RuleML. In *ICAIL*, volume 13.

Azzedin, F. and Maheswaran, M. (2002). Evolving and managing trust in grid computing systems. In *IEEE CCECE. Canadian Conference on Electrical and Computer Engineering.*, volume 3.

Backes, M., Lorenz, S., Maffei, M., and Pecina, K. (2010). Anonymous webs of trust. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*.

Bartolini, C., Muthuri, R., and Santos, C. (2015). Using ontologies to model data protection requirements in workflows. In *JSAI International Symposium on Artificial Intelligence*.

Beiter, M., Mont, M. C., Chen, L., and Pearson, S. (2014). End-to-end policy based encryption techniques for multi-party data management. *Computer Standards & Interfaces*, 36(4).

Bellare, M. and Yee, B. (1997). Forward integrity for secure audit logs. Technical report, Technical report, Computer Science and Engineering Department, University of California at San Diego.

Boer, A., Winkels, R., and Vitali, F. (2008). Metalex XML and the legal knowledge interchange format. In *Computable models of the law*.

Bonatti, P., Duma, C., Olmedilla, D., and Shahmehri, N. (2005). An integration of reputation-based and policy-based trust management. In *Semantic Web Policy Workshop*.

Bonatti, P., Kirrane, S., Polleres, A., and Wenning, R. (2017). Transparent personal data processing: The road ahead. In *International Conference on Computer Safety, Reliability, and Security*.

Bonatti, P. A. and Kirrane, S. (2019). Big data and analytics in the age of the GDPR. In *2019 IEEE International Congress on Big Data (BigDataCongress)*.

Bonatti, P. A. and Olmedilla, D. (2007). Rule-based policy representation and reasoning for the semantic web. In *Proceedings of the Third International Summer School Conference on Reasoning Web*.

Capadisli, S., Guy, A., Lange, C., Auer, S., Sambra, A., and Berners-Lee, T. (2017). Linked data notifications: a resource-centric communication protocol. In *European Semantic Web Conference*.

Cofta, P. L. (2018). Trust and the web – A decline or a revival? *Journal of Web Engineering*, 17(8).

De Vos, M., Kirrane, S., Padget, J., and Satoh, K. (2019). ODRL policy modelling and compliance checking. In *RuleML+RR*.

Golbeck, J. (2006). Trust on the world wide web: A survey. *Foundations and Trends in Web Science*, 1(2).

Hendler, J. and Golbeck, J. (2008). Metcalfe's law, web 2.0, and the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 6(1).

Information Commissioner's Office (ICO) UK (2017). Getting ready for the GDPR.

Jøsang, A., Ismail, R., and Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision support systems*, 43(2).

Kagal, L., Finin, T., and Joshi, A. (2003). A policy based approach to security for the semantic web. In *The Semantic Web - ISWC 2003*.

McDonald, A. M. and Cranor, L. F. (2008). The cost of reading privacy policies. *ISJLP*, 4.

Microsoft Trust Center (2017). Detailed GDPR Assessment.

Mitchell, C. (2005). *Trusted computing*, volume 6. IET.

Nymity (2017). GDPR Compliance Toolkit.

Oppliger, R. and Rytz, R. (2005). Does trusted computing remedy computer security problems? *IEEE Security & Privacy*, 3(2).

Palmirani, M., Governatori, G., Rotolo, A., Tabet, S., Boley, H., and Paschke, A. (2011). LegalRuleML: XML-based rules and norms. In *Workshop on Rules and Rule Markup Languages for the Semantic Web*.

Pandit, H. J., Fatema, K., O'Sullivan, D., and Lewis, D. (2018). GDPRtEXT-GDPR as a linked data resource. In *European Semantic Web Conference*.

Pinyol, I. and Sabater-Mir, J. (2013). Computational trust and reputation models for open multi-agent systems: A review. *Artificial Intelligence Review*, 40(1).

Pulls, T., Peeters, R., and Wouters, K. (2013). Distributed privacy-preserving transparency logging. In *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society*.

Sackmann, S., Strüker, J., and Accorsi, R. (2006). Personalization in privacy-aware highly dynamic systems. *Communications of the ACM*, 49(9).

Sambra, A. V., Mansour, E., Hawke, S., Zereba, M., Greco, N., Ghanem, A., Zagidulin, D., Aboulnaga, A., and Berners-Lee, T. (2016). Solid: A platform for decentralized social applications based on linked data.

Tang, Q. (2008). On using encryption techniques to enhance sticky policies enforcement. *DIES, Faculty of EEMCS, University of Twente, The Netherlands*.

Wahab, O. A., Bentahar, J., Otrok, H., and Mourad, A. (2015). A survey on trust and reputation models for Web services: Single, composite, and communities. *Decision Support Systems*, 74.