

Using Patterns to Manage Governance of Solid Apps

Beatriz Esteves^{1,*}, Harshvardhan J. Pandit²

¹*Ontology Engineering Group, Universidad Politécnica de Madrid, Spain*

²*ADAPT Centre, Dublin City University, Ireland*

Abstract

Currently, the Solid Protocol and its specifications lack the necessary vocabulary and processes for ensuring transparency and accountability in the use of data. In particular, to deal with the obligations and requirements required by regulations related to (personal) data protection and privacy. In addition, the lack of a guiding vocabulary leads to no common mechanism through which apps can request data and how Solid maintains information about its use. To address these, we propose PLASMA – a policy language to describe the entities, infrastructure, legal roles, policies, notices, and records to understand and establish responsibilities and accountability within the Solid ecosystem. We present how ontology design patterns using PLASMA can provide a common interface to create structured policies, records, and logs within the diverse Solid use cases, and thereby solve challenges regarding the management and governance of apps and their privacy considerations.

Keywords

Solid, access control, policies, GDPR, regulatory compliance, ontology design patterns


1. Introduction


Solid is an initiative led by Sir Tim Berners-Lee to decentralise the web and provide its users with greater choice when it comes to the usage of their (personal) data. Solid builds upon web standards such as the Linked Data Platform (LDP) [1] or the Web Access Control (WAC) [2] specifications and, according to its Protocol [3] specification, relies on said standards to “realise a space where individuals can maintain their autonomy, control their data and privacy, and choose applications and services to fulfil their needs”. While it was designed with web’s ethical principles¹ in mind to share information in a secure and private way, Solid is currently lacking when it comes to its alignment with data protection regulatory efforts [4], such as the European Union’s General Data Protection Regulation (GDPR) [5].


In particular, Solid lacks practical mechanisms through which GDPR’s principles of transparency and accountability can be implemented, which leads to no common mechanism for users or apps to represent or record information about privacy notices, agreements or giving/withdrawing consent and exercising of rights. Solid’s access control specifications only record what entity has been permitted to access which data, but not the purpose or intent behind that access, nor what will happen subsequently to that data once accessed. This leads to issues

14th Workshop on Ontology Design and Patterns (WOP 2023@ISWC 2023), November 6–7, 2023, Athens, Greece

*Corresponding Author’s Emails: beatriz.gesteves@upm.es, me@harshp.com

 0000-0003-0259-7560 (B. Esteves); 0000-0002-5068-3714 (H.J. Pandit)

 © 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

¹<https://www.w3.org/TR/ethical-web-principles/>

with regulatory compliance and user disenfranchisement arising from the lack of actionable ‘records’, and permits the continuation of existing issues such as the use of dark patterns and manipulations to obtain access.

By taking this situation as a research challenge, we propose PLASMA – a Policy Language for Solid’s Metadata-based Access control, which enables responsible and accountable practices to be incorporated into existing Solid infrastructure and workflows in a way that addresses the above-mentioned challenges. The work is based on the use of the Open Digital Rights Language (ODRL), a W3C standard for expressing policies, along with an additional taxonomy for defining the relevant concepts present in Solid’s ecosystem. Through this, PLASMA forms a ‘*policy layer*’ on top of the existing Solid implementations where users, apps, and Pod providers can express their relevant information in the form of policies, and based on which appropriate tools can be developed to enable functions such as informed consent, data use logs, machine-readable notices, preference management for users, and their management through dashboards. In addition to providing the vocabulary, a series of ontology design patterns (ODPs) can be defined to ensure apps declare necessary data before being allowed the use of data, and further to create conformity mechanisms based on this where communities can assess and curate collections of actors and apps to create environments of trust.

The contributions of this work are based on the following research objectives:

- RO1. Creating a taxonomy of Solid’s entities and infrastructure to describe the actors and processes involved in the Solid ecosystem;
- RO2. Providing a metadata policy language (PLASMA), using the terms identified in RO1, to express information regarding legal roles and other compliance requirements in a jurisdiction-agnostic manner (while satisfying requirements from GDPR);
- RO3. Using PLASMA for defining a set of Solid-related ODPs regarding users and apps policies, data use logs and registries to provide easy access to data in Pods.

This paper is organized as follows: Section 2 provides background on Solid and describes relevant work in state of the art regarding Solid access control mechanisms and metadata expression, Section 3 presents an overview of the challenges to be addressed to have a legally-aligned Solid ecosystem, Section 4 presents the use of PLASMA as a metadata policy language and Section 5 its integration in the ecosystem using ODPs, Section 6 discusses the challenges of addressing legal compliance, and concluding statements are presented in Section 7.

2. Background & Related Work

The Solid Protocol specification [3] contains a set of conditions that resource server developers must fulfil in order to enable Solid servers to send and retrieve data and that app developers must implement for a Solid app to perform operations on resources. In addition to this, a set of Solid specifications² build on top of the Protocol are being developed by the W3C Solid Community Group. In the context of this work, we focus on the specifications related to the authorization and application interoperability systems as the background for the development

²<https://solidproject.org/TR/> provides an up-to-date list of Solid technical reports.

of PLASMA and present existing research on (i) the expression of machine-readable policies and auditing metadata in Solid, and (ii) the alignment of Solid with data protection requirements.

Currently, Solid has two specifications to determine access authorizations to resources stored in Solid Pods – WAC³ and Access Control Policy (ACP)⁴ [6]. In WAC, IRIs are used to represent resources and agents, and authorization statements are stored within Access Control Lists (ACLs) defined per resource or inherited from the parent resource. In ACP, access grants are provided to certain resources and agents based on an `acp:Context` that describes the agent's request for data and the access control resources (ACRs) that describe who is allowed or denied access to Pod resources using an `acp:Matcher`.

Additionally, the Solid Application Interoperability⁵ specification [7] is being developed in order to have an interoperability layer in Solid that allows users and agents to access and manage data stored in a Solid Pod using distinct applications. However, these specifications do not take into consideration legal and practical requirements for the access and management of data nor do they provide a way for apps and users to express their data handling practices. Regarding auditing, Inrupt's Enterprise Solid Server provides an auditing service⁶ which relies on the W3C Activity Streams 2.0 Recommendation [8] to document audit events in RDF.

To this end, Esteves et al. [9] developed an ODRL profile for Access Control (OAC)⁷ that allows the expression of machine-readable policies to express individual privacy preferences and apps' policies, using data terms from the Data Privacy Vocabulary (DPV)⁸. OAC was extended by Debackere et al. [10] to include the efforts made in the Application Interoperability specification as part of a policy-based architecture for access control in Solid. DPV [11] is a W3C community effort, that provides taxonomies for expressing machine-readable metadata about the processing of personal data based on legislative requirements, including legal entities, purposes for processing, legal basis, rights, personal data categories, and GDPR. In addition, previous work on policy-related ODPs has been published to describe personal data in privacy policies [12] and linked data licenses [13].

For alignment with data protection requirements, Pandit [4] identifies actors and roles, and investigates GDPR issues applicable to the use of Solid transparency of information, consent, and exercising of data subject rights. Janssen et al. [14] also provide an analysis of personal data stores, including Solid, as a 'privacy-enhancing technology' (PET) and discuss the uncertainties in the designation of GDPR responsibilities among the involved parties in personal data store systems, as well as the choice of a proper legal basis for processing personal data.

Externally, 'app stores' such as those maintained by Apple and Google and 'package repositories' such as those maintained by Linux distributions also feature the use of metadata that developers must provide in order for their apps to be enlisted in their stores. These platforms also use this metadata for governing apps, identifying compatibility, and enforcing guidelines.

³WAC uses the ACL ontology to describe access authorizations to resources. Its namespace is <http://www.w3.org/ns/auth/acl#> and preferred prefix is `acl`.

⁴ACP's namespace is <http://www.w3.org/ns/solid/acp#> and its preferred prefix is `acp`.

⁵The Solid Application Interoperability's vocabulary namespace is <http://www.w3.org/ns/solid/interop#> and its preferred prefix is `interop`.

⁶<https://docs.inrupt.com/ess/latest/services/service-auditing/>

⁷OAC's namespace is <https://w3id.org/oac#> and its preferred prefix is `oac`.

⁸DPV's namespace is <https://w3id.org/dpv#> and its preferred prefix is `dpv`.

3. Challenges and Hurdles to the Solid Vision

This section discusses the challenges of having a transparent, legally-aligned Solid ecosystem. The following issues were identified based on an analysis of the literature and the current state of Solid specifications:

- C1. **Identity of Solid actors and their roles is unknown** – Solid users are unaware of who are the entities proving their Pods or the apps they use.
- C2. **No metadata about Solid infrastructure** – Users do not have information on which Solid specification their Pod is running or which functionalities are installed within it.
- C3. **Availability/Discovery of categories of data** – To have granular access to data in Pods, based on their data category, Solid apps need to know if they exist and where are they stored in the Pod.
- C4. **Pod and app providers do not provide information on their data handling practices** – Pods and apps do not provide machine-readable privacy notices nor do they declare what data they need to function.
- C5. **Users cannot express their privacy policies** – Solid users do not have the tools to express their privacy preferences and requirements nor to manage incoming data requests or existing decisions for the use of data.
- C6. **No logging or record-keeping** – No provenance metadata is recorded for accountability in the Pods, e.g., users do not keep consent records, nor information regarding who has accessed their data, what they are doing with it, or changes to their data handling policies.
- C7. **No legal compliance checks** – Solid does not have tools for users to exercise their GDPR-related rights, such as withdrawing consent, nor does it provide authorities with the auditing information to perform investigations.

While we claim that this is not an exhaustive list of all challenges that need to be overcome to have a legally-aware Solid ecosystem, the work on PLASMA that we present in the following sections can be used to address them.

4. PLASMA: A Policy Language for Solid's Metadata-based Access Control

As covered in the previous sections, there is a gap in the representation of information related to the entities, infrastructure and processes involved in the Solid ecosystem. While there are existing efforts at providing legal vocabularies [11], these are not directly applicable to Solid use cases. Therefore, through PLASMA, we first provide relevant taxonomies to describe the actors, artefacts, and processes that are necessary to describe the use of Solid Pods. This enables us to later align these with legal terms (in our case – the GDPR). For example, the entity the data belongs/refers to is called *owner* in Solid, whereas the equivalent term would be *data subject* under GDPR. In addition, we also identify concepts without direct equivalence in GDPR or those that can take one of multiple roles. Through this, PLASMA is able to express a variety of use cases using terms familiar to the Solid community while also providing their relevance in terms of legal concepts (from GDPR).

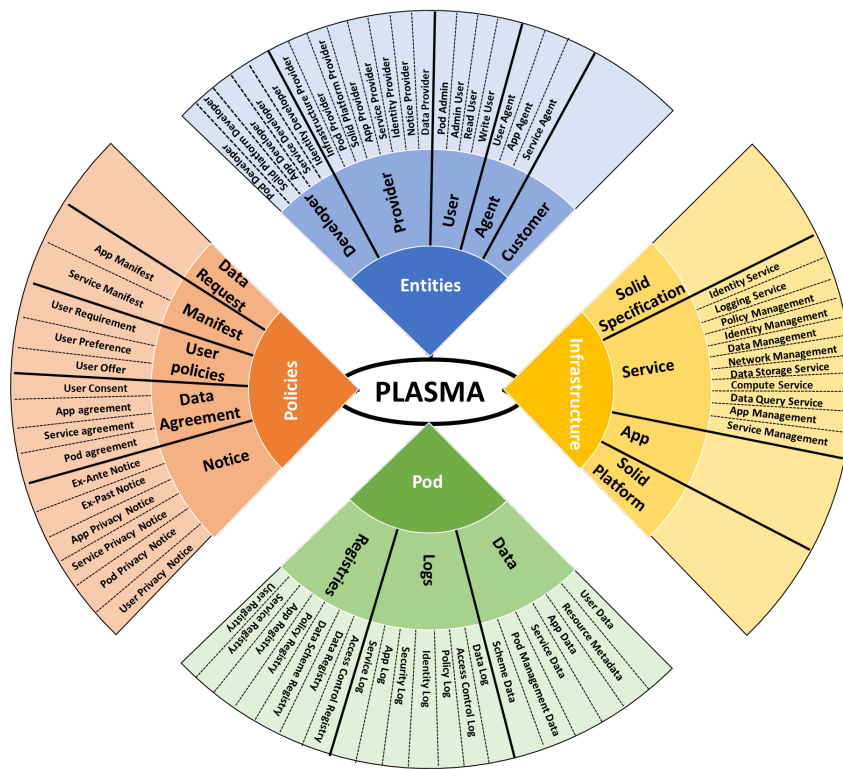


Figure 1: Overview of the main concepts included in PLASMA.

In order to understand and study Solid-related terms and their current definitions, a review of the existing Solid technical documents was performed by the authors, in particular of the specifications related to the authorization workflow, described in Section 2. While these specifications provide a definition for concepts such as Pods, apps or agents in their documentation, only a handful of them is actually defined in a machine-readable manner: the interop specification provides a definition for Application, Agent and Registry and the acp vocabulary provides a Policy term. However, no concepts were found in the base Solid vocabularies to define what is a Pod, a Service, an Entity or Data. As such, we provide a definition for each of these concepts in the base vocabulary⁹ of PLASMA and link them with the existing concepts mentioned before, if available. Thus, PLASMA relies on these base terms as its core concepts to be expanded into distinct taxonomies for entities, policies, services, registries, and data – Figure 1 provides an overview of these taxonomies. PLASMA is available at <https://w3id.org/plasma>, under the CC-By-4.0 licence, and reuses the available terms in the Solid vocabularies, in line with FAIR (Findable, Accessible, Interoperable, Reusable) principles. Each taxonomy is discussed below.

⁹PLASMA's base vocabulary documentation can be checked at <https://w3id.org/plasma#base>.

4.1. Entities

Currently, Solid vocabularies do not provide terms to describe entities beyond the existing properties in the `acl` and `acp` vocabularies to describe access conditions for agents, client applications or identity issuers. In PLASMA, we provide the terms to describe the providers and/or developers of Solid Pods, apps, services, data, identity and other Solid-related infrastructure, as well as terms to describe different types of users – we define `AdminUser` as a “User of a Pod that has the administrative capability to make decisions about Data on a Pod” and `PodAdmin` as a “User of a Pod that has the administrative capability to make decisions about the Pod (separate from Data in a Pod), such as deleting the Pod, changing identity or other resource providers”, to potentiate the distinction between users with full control over data and users with full control over the Pod, a feature currently missing from the Solid specifications. Regarding the specification of agents, PLASMA goes beyond the current definition of Solid in which such entities can be real, e.g., people or parents on behalf of children, or virtual, e.g., software agents – PLASMA agents are virtual agents, so as to distinguish them from an entity which represents a real entity, i.e., those that can be held legally accountable. These concepts address the C1 challenge identified in Section 3.

4.2. Infrastructure

PLASMA provides terms to specify the Solid specifications that the Pod conforms to in terms of behaviour and implementation details, as well as the specific Solid platform that is installed within a Pod. In addition, in PLASMA we make a distinction between Solid apps and Solid services – an app requires human intervention to perform an action on data with the aim of providing specific purposes or functionalities, while a service is a functionality that may utilise or interact with data within a Pod and may not require human intervention, e.g., is executed in the background. Users can choose services to run on their Pods, with the services having their own developers and providers, e.g., an identity verification service that checks the validity of certificates. As this is a new concept introduced by PLASMA, we provide a taxonomy of 11 Solid-related services, which can be further expanded to take into account new use cases. These concepts address the C2 challenge identified in Section 3.

4.3. Pod-related Data

Since Solid’s overall purpose is to provide individuals with a data storage service for their data and the choice of which applications or services to use for a particular purpose, it should also contain (meta)data regarding its users, apps, services, logs or registries. In addition, the recording of this provenance data in the Pods can also be of assistance to external entities auditing Solid-related activities, e.g., data protection supervisory authorities in Europe might require access to said data to investigate a personal data breach. To this end, beyond modelling metadata about Pod management, users, apps or services, PLASMA allows the expression of logs – provenance records associated with Solid processes such as adding/updating a data resource in a Pod, the reporting of a policy negotiation where the user gave consent to an app’s request for data, or of an error on the identity verification of a user. Moreover, the maintenance of a set of registries, e.g., data, data schema, policy, app or user registries, as an indexed record used to

provide information on the availability of data categories, supported schemas for data, apps, or services, relevant policies, or apps and users that have/had access to resources in the Pod. These concepts address the C6 and C3 challenges identified in Section 3, respectively.

4.4. Policies

In the context of PLASMA, policies are documents used to specify the requirements of users, apps and services regarding data handling practices over data stored or shared through Solid Pods. As such, PLASMA provides definitions for user policies, as well as for data requests, i.e., a request from an app, service, agent or user to access, use and perform other actions on Pod data. We envision integrating such requests in a manifest as defined by the W3C Web Application Manifest specification [15], which currently cannot be considered sufficient for an app's or service's manifest in the context of Solid Pods since it does not include information on entities, their identity or data handling policies, which is needed to have accountability and check for legal compliance. PLASMA also provides different types of data agreements, a concept totally missing from the Solid specifications as they only refer to access authorizations between users and apps. An agreement can be governed by a contract between the user and the entity responsible for the app, service or Pod, or can be based on the user's consent. Furthermore, current Solid specifications also miss the definition of notices – a document needed to provide context information about the entities, operations, or data involved in a particular process, e.g., notices can be specified to declare the providers and/or developers of apps and services or to describe data processing practices. In this context, PLASMA provides the terms to declare an *ex-ante* and *ex-past* notices, as well as privacy notices for apps, services, Pods and users. These concepts address the C4 and C5 challenges identified in Section 3.

5. Integrating PLASMA in the Solid Ecosystem through the Usage of Ontology Design Patterns

As previously discussed, the current Solid vision does not rely on a policy-based access control system, but on an access control list (WAC) or access grant (ACP) mechanism. In addition, it completely lacks the necessary information to perform an auditing activity on Pods, apps, services, entities and agents and a proper alignment with data protection regulatory requirements such as the ones brought on by the GDPR, which are necessary conditions for the lawful processing of personal data. In this section, we promote the usage of PLASMA and other identified semantic specifications to describe a set of ODPs for having commonly structured representations of policies, registries and logs in Solid, which can be used to establish responsibilities and promote transparency in personal data handling practices.

5.1. ODPs for Specifying Legally-aware Policies in Solid

PLASMA defines a set of 3 types of user policies, a `UserPreference` – a soft rule that may not be satisfied, a `UserRequirement` – a hard rule that should always be satisfied, and a `UserOffer` – a policy offer from the user regarding the purposes, entities and actions that can/cannot be performed on the data, in addition to the ones specified to characterize `DataRequests` and

DataAgreements, already discussed in Section 4.4. While a different ODP can be established for each type of policy, their main structure remains the same. For this reason and due to restrictions on the size of this publication, we present only the ODP for a data request policy¹⁰. The pattern for request policies aims to answer the following competency questions:

- (CQP1.) What is the unique identifier of the policy?
- (CQP2.) Who is the creator of the policy?
- (CQP3.) When was the policy issued?
- (CQP4.) Who is the assignee of the policy?
- (CQP5.) What application/service is being used to access the data?
- (CQP6.) What access mode is being requested?
- (CQP7.) What personal data is being accessed?
- (CQP8.) What is the purpose for accessing the data?

Since neither WAC nor ACP provide the necessary terms to express access policies for specific categories of data, with restricted purposes for usage, or with temporal or spatial constraints, PLASMA promotes the integration of ODRL, which has a convenient extension mechanism, into the Solid architecture. Moreover, ODRL already provides the terms to specify offers as an `odrl:Offer`, requests as an `odrl:Request`, and agreements as an `odrl:Agreement` and an ODRL profile for Access Control (OAC) [9], which invokes data protection-specific terms using DPV, that can also be integrated into the Solid ecosystem to express granular permissive and prohibitive policies related to the access to personal data stored in Solid Pods, with constraints over purposes, recipients or legal basis for processing. The DCMI Metadata Terms specification is also used to specify authorship of the policy, `dct:creator`, and its issuance date, `dct:issued`. A visualisation of the pattern is presented in Figure 2.

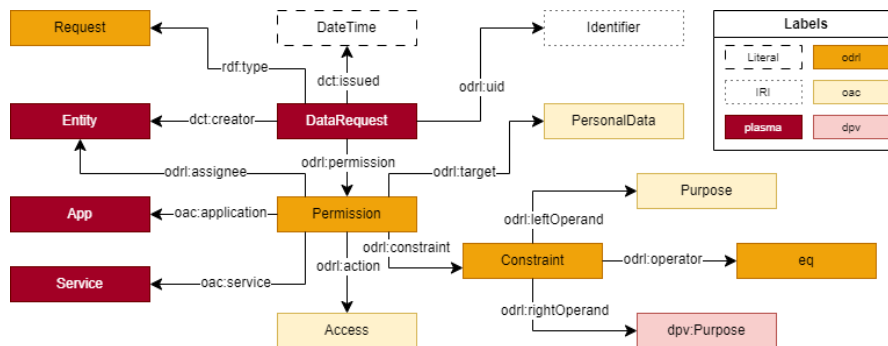


Figure 2: Ontology Design Pattern for DataRequests.

5.2. ODPs for Logging and Registries

In addition to pattern-based policies, as previously discussed in Section 4.3, PLASMA should be used to construct a variety of patterns for Solid logs and registries, to assist in the auditing of

¹⁰Examples of usage and ODPs for user policies and agreements are available at <https://w3id.org/plasma/odps>.

Pods, users, apps and services' activities as well as to have easy access to data in Pods. Due to restrictions on the size of the publication, we present only an ODP for a data log and an ODP for a data registry, while other ODPs will be available at <https://w3id.org/plasma/odps>. The pattern for a data log aims to answer the following competency questions:

- (CQL1.) What type of action, e.g., create, update, erase, is being performed on the data?
- (CQL2.) Who is the entity interacting with the data?
- (CQL3.) Who is the entity publishing the log?
- (CQL4.) When was the log issued?
- (CQL5.) Where is the data being stored?
- (CQL6.) What application/service is being used to generate the data, if any?

The pattern for a data registry aims to answer the following competency questions:

- (CQR1.) Who is maintaining the registry?
- (CQR2.) When was the registry created/updated?
- (CQR3.) What types of data are available?
- (CQR4.) Where is a specific type of data being stored?
- (CQR5.) What policy is associated with the data?

Since Solid does not currently support any vocabularies to express logs, the ActivityStreams protocol¹¹ can be used to describe such records. ActivityStreams supports the description of activity types, such as Create, Add, Update or Delete, and includes concepts to associate activities with the entity that performed them, as `:actor`, with applications/services that generated them, as `:generator`, or with the resources that they are connected with, as `:object`. In addition, data registries can be commonly structured as a DCAT¹² Catalog, where different datasets can be associated with the type of data they contain using DPV's personal data categories extension (DPV-PD¹³) and associate them with their storage location using `foaf:page`. A visualisation of the DataLog and DataRegistry patterns is presented in Figure 3.

6. Conformance and Legal Compliance

Using PLASMA, an app (and its controlling entity) is able to provide information or links to information regarding policies for who is the entity, what are they doing with the data, towards what ends, and other relevant information. When this information is machine-readable, it is possible to automatically assess it towards some objective – such as to ensure that usage purposes are explicitly acknowledged as being of a certain type (e.g., scientific research) or that data will never travel outside a jurisdiction. Such declarations can serve as valuable tools in automating or reducing the effort required to approve apps and requests where the use of data and Pods occurs within environments such as clinical research.

¹¹ActivityStreams 2.0 Terms are published under the namespace <https://www.w3.org/ns/activitystreams#>, with ``as`` as its preferred prefix.

¹²The Data Catalog Vocabulary (DCAT) is published under the namespace <http://www.w3.org/ns/dcat#>, with ``dcat`` as its preferred prefix.

¹³<https://w3id.org/dpv/dpv-pd>

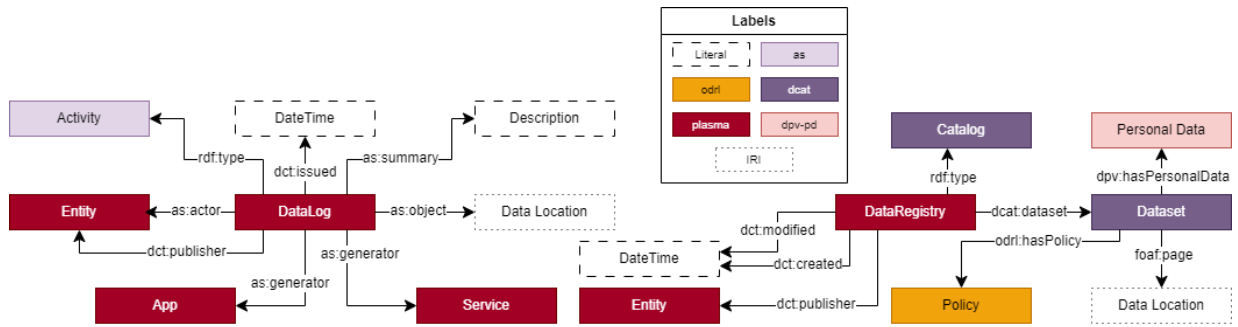


Figure 3: Ontology Design Pattern for a DataLog and DataRegistry.

Using PLASMA as the basis, we also foresee the possibility to realise *community* or *collective* efforts, such as those envisioned by the MyData¹⁴ movement, where apps must conform or satisfy requirements set forth by the collective and are only then allowed to approach the Pods and users. This is beneficial for apps as it reduces the number of individual notices and requests it has to manage, and also for users as they only receive requests that they know have been approved or vetted by the collective arrangement. Further, PLASMA also provides a way to use Solid Pods towards the new Data Governance Act (DGA) [16] where the data could reside in the Pod with a usage policy and a Data Intermediary or Co-operative or Altruistic Organisation could periodically collect only the policies available on the Pod to create a common registry of available data and how it can be reused, without providing direct access to the data itself.

In terms of legal compliance, more specifically the GDPR, PLASMA provides a forum for the provision and management of information comparable to the current privacy policies or notices where users can go and read details about how the service uses personal data and what are their options regarding it. It has been well studied and documented that such policies are hard for people to find, comprehend, and use due to their length and complex language and that users prefer not having to read them [17]. Similarly, the privacy notices, such as for consent, have also been documented to suffer similar issues and malpractices [18]. Without approaches such as PLASMA, the use of Solid will also suffer with these existing issues and additional ones that arise specifically from Solid’s lack of information or concrete guidance.

In order to make effective use of PLASMA, we emphasise that additional research and developments are required for the various possibilities highlighted throughout this article. For example, currently, there are no semantic specifications for recording or constructing privacy notices. Since legal requirements are specific to a jurisdiction, such mechanisms would invariably be tied to the interpretation of a specific law. Given that we are witnessing a global convergence of data protection and privacy laws with the general concepts and structure first emphasised by the GDPR, it is now the appropriate time to create a common legal vocabulary.

For this, we think that the DPV [11] is a suitable candidate that could be expanded to be used with PLASMA to declare privacy notices, consent records, and other pertinent documentation that uses Solid-friendly terms while still being tied to legal requirements (e.g., GDPR). For

¹⁴Information on the MyData movement available at <https://www.mydata.org/>.

example, using PLASMA, the use case can contain information that an actor is an App Developer, and with DPV they can assert that they are merely a Data Processor. Similarly, using PLASMA, the App can specify where to find its privacy notice, where the contents of that notice can be described using DPV. The use of ODPs, such as the ones described in Section 5, will ensure that the appropriate information based on its legal requirements, e.g., from GDPR's Article 13/14, is present in the Pod to be consulted. As such, the creation of patterns for the different types of policies, notices, logs and registries is of the utmost importance. We have undertaken this work within the DPVCG and are working towards developing machine-readable specifications based on standards of ISO/IEC 29184 Privacy Notices and ISO/IEC 27560 Consent Records, to help in addressing the C7 challenge identified in Section 3.

7. Conclusions

In this article, we first established the need to have common taxonomies to describe the entities, infrastructure and processes involved in the Solid ecosystem, in order to address current challenges that the Solid vision brings in terms of providing information about the identity of Pod, apps or other Solid-related services, information regarding their personal data handling practices or the generation and maintenance of logs related with important Solid processes, e.g., updating a Pod data resource, moving to a different Pod provider, or deleting a given access authorization from the Pod. To this end, we propose the integration of PLASMA, a metadata policy language, into the Solid ecosystem, with the purpose of expressing information regarding legal roles, agreements, policies, registries of data and logs, taking into consideration legal requirements, first in a jurisdiction-agnostic manner and then, in particular, considering GDPR's specific requirements. By integrating such a resource in Solid and providing guidelines on how to comply with it, we promote semantic interoperability in the terms used by users, Pods, apps, services and agents in order to help address the identified challenges to the Solid vision, while considering legal compliance with data protection laws.

We believe that our work further advances the field of decentralized storage of personal data, relying on existing Semantic Web standards, by promoting its integration with legal and social requirements and therefore providing a more trustworthy and private-friendly environment for Web users to share their data and gather benefits from it. As future work, we highlight the need to (i) evaluate the coverage of PLASMA to deal with the variety of different workflows and use cases, (ii) integrate the usage of PLASMA, and of ODRL and DPV as well, into the design of Solid servers, applications and services, (iii) develop SHACL shapes to check for compliance with the PLASMA specification and the developed PLASMA ODPs, including the usage of DPV to comply with legal requirements, (iv) align PLASMA with existing standardization efforts for privacy notices and consent records and (v) further develop PLASMA's patterns to cover different types of notices, logs and registries, while providing information on entities, policies, etc. as per GDPR's requirements.

References

- [1] S. Speicher, J. Arwe, A. Malhotra, Linked Data Platform 1.0, W3C Recommendation (2015). URL: <http://www.w3.org/TR/ldp/>.
- [2] S. Capadisli, T. Berners-Lee, Web Access Control Version 1.0.0, W3C Candidate Recommendation (2022). URL: <https://solidproject.org/TR/wac>.
- [3] S. Capadisli, T. Berners-Lee, R. Verborgh, K. Kjernsmo, Solid Protocol Version 0.10.0, W3C CG Draft Report (2022). URL: <https://solidproject.org/TR/protocol>.
- [4] H. J. Pandit, Making Sense of Solid for Data Governance and GDPR, Information 14 (2023) 114. doi:10.3390/info14020114.
- [5] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2018. URL: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.
- [6] M. Bosquet, Access Control Policy, W3C CG Draft Report (2022). URL: <https://solidproject.org/TR/acp>.
- [7] J. Bingham, E. Prud'hommeaux, e. Pavlik, Solid Application Interoperability, W3C CG Draft Report (2023). URL: <https://solid.github.io/data-interoperability-panel/specification/>.
- [8] J. M. Snell, E. Prodromou, Activity Streams 2.0, W3C Recommendation (2017). URL: <https://www.w3.org/TR/activitystreams-core/>.
- [9] B. Esteves, H. J. Pandit, V. Rodríguez-Doncel, ODRL Profile for Expressing Consent through Granular Access Control Policies in Solid, in: 2021 EuroS&PW, 2021, pp. 298–306.
- [10] L. Debackere, P. Colpaert, R. Taelman, R. Verborgh, A Policy-Oriented Architecture for Enforcing Consent in Solid, Workshop Proceedings of WWW '22, 2022.
- [11] H. J. Pandit, A. Polleres, B. Bos, R. Brennan, B. Bruegger, F. J. Ekaputra, J. D. Fernández, R. G. Hamed, M. Lizar, E. Schlehahn, S. Steyskal, R. Wenning, Creating A Vocabulary for Data Privacy, in: ODBASE2019, Rhodes, Greece, 2019, p. 17. doi:10/ggwx7x.
- [12] H. J. Pandit, D. O'Sullivan, D. Lewis, An Ontology Design Pattern for Describing Personal Data in Privacy Policies, in: 9th Int. Workshop on Ontology Patterns, 2018, pp. 29–39.
- [13] V. Rodríguez-Doncel, M. C. Suárez-Figueroa, A. Gómez-Pérez, M. Poveda-Villalón, License Linked Data Resources Pattern, in: 4th Int. Workshop on Ontology Patterns, 2013.
- [14] H. Janssen, J. Cobbe, C. Norval, J. Singh, Decentralized data processing: personal data stores and the GDPR, International Data Privacy Law 10 (2020) 356–384.
- [15] M. Cáceres, K. R. Christiansen, M. Giuca, A. Gustafson, D. Murphy, A. Kostianen, M. Lamouri, R. Dolin, Web Application Manifest, W3C WG Draft Report (2023). URL: <https://www.w3.org/TR/appmanifest/>.
- [16] Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European data governance and amending Regulation (EU) 2018/1724 (Data Governance Act), 2022. URL: <http://data.europa.eu/eli/reg/2022/868/oj/eng>.
- [17] M. Veale, F. Z. Borgesius, Adtech and Real-Time Bidding under European Data Protection Law, German Law Journal 23 (2022) 226–256. doi:10.1017/glj.2022.18.
- [18] M. Toth, N. Bielova, V. Roca, On dark patterns and manipulation of website publishers by CMPs, in: PoPETs, 2022, pp. 478–497. doi:10.56553/popets-2022-0082.