

Privacy by Design in Federated Identity Management

Rainer Hörbe
Principal
Identinetics GmbH
Tulln, Austria
rh@identinetics.com

Walter Hötendorfer
Centre for Computers and Law
University of Vienna
Vienna, Austria
walter.hoetendorfer@univie.ac.at

Abstract—Federated Identity Management (FIM), while solving important scalability, security and privacy problems of remote entity authentication, introduces new privacy risks. By virtue of sharing identities with many systems, the improved data quality of subjects may increase the possibilities of linking private data sets; moreover, new opportunities for user profiling are being introduced. However, FIM models to mitigate these risks have been proposed. In this paper we elaborate privacy by design requirements for this class of systems, transpose them into specific architectural requirements, and evaluate a number of FIM models with respect to these requirements. The contributions of this paper are a catalog of privacy-related architectural requirements, joining up legal, business and system architecture viewpoints, and the demonstration of concrete FIM models showing how the requirements can be implemented in practice.

Keywords—identity management; federated identity management; privacy; privacy by design; security; data protection law; limited observability; limited linkability

I. INTRODUCTION

Today, due to the lack of a native “identity layer”, identity management (IM) on the Internet is still largely the operation of separated domains using userid/password accounts exclusively for the purposes of their own online services. To overcome this situation, different models of federated identity management (FIM) have been proposed that enable the use of identities across organizational borders, single sign-on, faster provisioning and linking of services and which would improve the situation for users and service providers alike. In particular, FIM has also the potential to mitigate the significant privacy flaws of the current situation (see e.g. [1]), but at the same time it introduces new privacy risks, primarily by centralizing user data and making it easier to track user behavior and to link data of the same user together.

To mitigate these privacy risks and to realize the full privacy-enhancing potential of FIM, the design process of FIM systems needs to take into account privacy requirements from the beginning. To put this into practice, this paper proposes a privacy by design approach based on earlier work by the Authors ([2], in German). This approach is a significant step in narrowing the gap between the highly abstract level of privacy legislation and the fairly concrete level of system architecture requirements.

This paper is structured as follows. Section II covers the questions that were the starting point of this work, the scope of FIM systems under discussion and the methodology applied in

this research. In section III we introduce FIM, privacy by design and privacy principles that are common in major data protection legislation and privacy guidelines. In section IV we apply the fundamental privacy principles to the FIM domain by deducing FIM-specific privacy by design requirements from them. Moving on to a more concrete level, we present eight architectural requirements that realize these privacy by design requirements. The last part deals with the existing models that have been proposed for improving privacy in FIM. We demonstrate how they fulfill the architectural requirements we came up with, which is at the same time an evaluation of existing models and a demonstration of how our architectural requirements can be applied in practice.

II. QUESTIONS, SCOPE AND METHODOLOGY

A. Questions

The primary motivation to look into privacy enhancing controls for FIM came from the observation that FIM projects are increasingly aiming at operating across sectors, such as: business-to-government and business-to-business, across unrelated supply chains, and across domains in smart cities.

The authors were involved in projects¹ with stakeholders questioning the sufficiency of current best practice with respect to FIM privacy controls. The ensuing discussions were the starting point for this paper, leading to the questions (a) whether there are feasible technical controls to improve privacy, in particular with respect to limiting the aggregation of metadata, and (b) how current FIM models map to these requirements.

B. The Scope of FIM Systems under Discussion

Our scope is on FIM models providing authentication and attribute assertion and claiming to respect the user’s privacy. We take SAML WebSSO [22] and OpenID (Connect) [33] as the predominant use case in these systems because of their respective ecosystems’ sizes [28] and deployments in certain sectors [29]; however, systematic market research is not publicly available [34]. We have excluded some well-known protocols because they provide Single-Sign-On only (Kerberos, authentication tokens) or authorization only (OAuth2 [30] and UMA [31]).

¹ These were commercial projects where the related documents cannot be referenced.

C. Methodology

As shown in Fig. 1, we obtained input from two sides. From the top we took general and abstract privacy principles from privacy legislation and guidelines (1). Based on our knowledge from the field and the literature, we examined what these principles mean for the FIM domain, and how they can be accomplished by design requirements (2). The result was a set of privacy by design requirements that can be realized with technical controls.

From the bottom we analyzed FIM models claiming improved privacy and recovered their architectural requirements (3). Finally, in an iterative process we joined both sides (4). We also took into account business requirements that are particularly relevant in our context. This resulted in a list of eight architectural privacy-related requirements for FIM systems.




Figure 1. Approach taken to elicit requirements

III. BACKGROUND

A. FIM Key Concepts and Risks

FIM is a middleware layer for electronic communication that provides Relying Parties (RP) with a means to identify users or devices (principals) without having to operate the related infrastructure for registration and authentication and its implied complexities for security, compliance and governance. Hence an RP relies on an identity representation including attributes of the user, which is asserted by an Identity Provider (IdP). A comprehensive overview on FIM can be found, e.g., in [3].

Although – as shown below – the introduction of FIM is able to increase privacy by featuring things like pseudonymous authentication and limited attribute release, FIM can also increase privacy risks for several reasons. Asserting attributes from a single source to multiple relying parties increases the data quality and thus the linkability. FIM systems might use a common identifier across multiple RPs. IdPs might profile users by aggregating authentication events. Even if FIM systems pool resources for information security, they pose a bigger target for cyber attacks at the same time [4].

This paper focuses on privacy (by design) requirements for FIM systems. A discussion of functional requirements for FIM systems and further related work can be found in [5].

B. Privacy by Design

Privacy by design can be briefly defined as the principle of including privacy in the system development life cycle from the beginning, but it is not yet fully clear what privacy by design means in practice [6]. The difficulty still is “figuring out how to translate the abstract principles, models, and mechanisms into comprehensive specific requirements for specific systems operating within specific contexts.” [7]

In our view, an important approach to this is to design a system in such a way that it infringes the privacy of individuals concerned by that system as little as possible, in particular by technically and architecturally precluding as far as possible that the system can be used in a privacy-infringing way.

Designing information systems in this way mitigates the common practical problem of illegitimate use of personal data inside the information systems of the data controller often happening without any consequences, because such illegitimate use can easily be hidden from the data subject concerned.

C. Privacy and Data Protection Principles

In this chapter, requirements for the design and implementation of FIM systems are deduced from data protection law and data protection guidelines. Primarily, a European perspective is taken, deriving the general principles from the EU Data Protection Directive (hereinafter “Directive 95/46/EC”) and the Council of Europe Convention for the Protection of Individuals with regard to Automatic Processing of Personal Data (hereinafter “CoE Convention 108”), but also from the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (hereinafter “OECD Guidelines”) and the ISO/IEC 29100:2011 privacy framework. We shall define the term “personal data” as “any information relating to an identified or identifiable natural person (‘data subject’)” (Article 2 (a) of Directive 95/46/EC).

PP1. Fairness and Lawfulness.² Personal data must be processed fairly and lawfully. Under European Union Law, processing of personal data is lawful only in cases explicitly permitted by law.

PP2. Finality.³ Personal data may be collected only for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes.

PP3. Proportionality.⁴ Personal data must be adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed.

PP4. Data Quality.⁵ Personal data must be accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that data which are inaccurate or incomplete, having regard to the purposes for which they were collected or for which they are further processed, are erased or rectified.

² Art. 6 a of Directive 95/46/EC; Art. 5 a of CoE Convention 108; Art. 7 of the OECD Guidelines.

³ Art. 6 b of Directive 95/46/EC; Art. 5 b of CoE Convention 108; Art. 9 of the OECD Guidelines; ISO/IEC 29100 (5.3).

⁴ Art. 6 c of Directive 95/46/EC; Art. 5 c of CoE Convention 108; Art. 8 of the OECD Guidelines; ISO/IEC 29100 (5.7).

⁵ Art. 6 d of Directive 95/46/EC; Art. 5 d of CoE Convention 108; Art. 8 of the OECD Guidelines; ISO/IEC 29100 (5.7).

PP5. Information Security.⁶ Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorized access, destruction, use, modification or disclosure of data.

PP6. Openness and Transparency.⁷ There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.

PP7. Individual Participation.⁸ An individual should have the right: (a) to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him; (b) to have communicated to him, data relating to him within a reasonable time (at a charge, if any, that is not excessive; in a reasonable manner; and in a form that is readily intelligible to him); (c) to be given reasons if a request made under subparagraphs (a) and (b) is denied, and to be able to challenge such denial; and (d) to challenge data relating to him and, if the challenge is successful, to have the data erased, rectified, completed or amended.

PP8. Accountability.⁹ A data controller should be accountable for complying with measures that give effect to the principles stated above.

Not explicitly in the list is the principle of Data Minimization, because it can be regarded as a super-principle. It is the underlying basis of the principles of Finality (PP2) and Proportionality (PP3). Data Minimization means to reduce the use of personal data to the extent absolutely necessary for the specified purposes [8]. This has several dimensions. It means that (a) if a problem can be solved in several ways, which are otherwise comparable, the way that requires the least amount of personal data and/or the least sensitive personal data must be chosen. It also means (b) that personal data may only be collected and stored in a form which permits identification of data subjects to the extent strictly necessary and (c) no longer than strictly necessary for the specified purposes¹⁰ and that (d) the number of people to whom personal data is disclosed or who have access to it should be minimized. Starting with data minimization is a necessary first step for any privacy by design approach [9].

Cameron [10] convincingly explains why data minimization is so important in identity management (and in general): “We should build [digital identity] systems that employ identifying information on the basis that a breach is always possible. [...] To mitigate risk, it is best to acquire information only on a ‘need to know’ basis, and to retain it only on a ‘need to retain’ basis. By following these practices, we can ensure the least possible damage in the event of a breach. At the same time, the value of identifying information decreases as the amount decreases. A system built with the

principles of information minimalism is therefore a less attractive target for identity theft, reducing risk even further”.

The general principles laid down in this chapter are – in this or in similar form – widely accepted and the foundation of many data protection and privacy laws around the world. But irrespective of whether or not they are strictly binding under the applicable national legislation – as they are in all EU Member States – they are general guidelines for implementing privacy into information systems by design. For the domain of FIM, this is done in the following section.

IV. REQUIREMENTS FOR FIM

We chose to categorize requirements along typical domain expertise: privacy engineering, business/IT-management, and system architecture. Other categorizations might be useful for other viewpoints.

A. Privacy by Design Requirements for FIM

One could now proceed with building a system that takes into account the data protection principles identified above and instruct operators that they must comply with these principles at runtime. However, in our privacy by design approach, a very important step follows before that.

It is based on the – trivial but effective – insight that operators and attackers not having the factual possibility to unnecessarily infringe privacy cannot do so. Hence the system must be designed in a way that minimizes the possibilities of using it in a privacy-infringing way while still fulfilling the purposes the system is built for. In other words, the system must be built in a way that its misuse is ruled out as much as possible. The basis for this is data minimization: Data that is unavailable cannot be misused.

To execute this, we have to come up with requirements that can be fulfilled by technical controls rather than by provisions that depend on compliance. Following this approach, we apply the fundamental privacy principles above to the FIM domain and deduce the following privacy by design requirements for FIM systems (the related fundamental principles are indicated in parentheses):

PDR1. Only use data in a form that permits identification of the data subject if absolutely necessary [11] (PP3). In cases where the identity of the data subject is not needed, it should not be disclosed to the relying party.

PDR2. Inhibit any entity from seeing and obtaining data it does not need for fulfilling its purposes [11] (PP1, PP2, PP3, PP5). A FIM system comprises several independent entities. As an entity cannot misuse data it does not have, this reduces privacy risk in the whole FIM system.

PDR3. Make the illegitimate linking of data across privacy domains as difficult as possible (PP1, PP2, PP3). Such undesired linking is facilitated by the use of unique identifiers and by the disclosure of commonly unique attributes, in particular the e-mail address. Generally, the more individual an attribute or a combination of attributes is, the easier is the re-identification of the data subject. Based on the principles of finality and accountability we define a *privacy domain* as the context where a particular controller processes personal data for a particular, specified purpose. Consequently, processing of the same personal data by the same controller for a different

⁶ Art. 17 of Directive 95/46/EC; Art. 7 of CoE Convention 108; Art. 11 of the OECD Guidelines; ISO/IEC 29100 (5.11).

⁷ Art. 10 and 11 of Directive 95/46/EC; Art. 12 of the OECD Guidelines; ISO/IEC 29100 (5.8).

⁸ Art. 12 of Directive 95/46/EC; Art. 13 of the OECD Guidelines; ISO/IEC 29100 (5.9).

⁹ Art. 22 et seq. of Directive 95/46/EC; Art. 14 of the OECD Guidelines; ISO/IEC 29100 (5.10).

¹⁰ Art. 6 e of Directive 95/46/EC.

purpose would already be considered as a different privacy domain.

PDR4. Make all flows of personal data transparent to the data subjects and provide them with the possibility to intervene (PP1, PP4, PP6, PP7, PP8).

PDR5. Implement the canonical information security protection goals, confidentiality, integrity and availability (PP5). In the context of a FIM system this means that unauthorized access of information is prevented (confidentiality), that unauthorized or unrecognized modification of information is prevented (integrity) and that authorized users are not prevented from the legitimate access of information or legitimate use of the system (availability, limited to the FIM context) [12].

Note that the relevance of the information security protection goals here is twofold: On the one hand, confidentiality, integrity and availability are requirements for IM systems, as for any other information security systems. On the other hand, the very purpose of IM systems is to ensure the confidentiality, integrity and availability (in the sense described above) of dependent information systems [12].

TABLE I. TRACING PRIVACY BY DESIGN REQUIREMENTS VS. PRIVACY PRINCIPLES

	PP1 Fairness and Lawfulness	PP2 Finality	PP3 Proportionality	PP4 Data Quality	PP5 Information Security	PP6 Openness and Transparency	PP7 Individual Participation	PP8 Accountability
PDR1 Minimal identification			↑					
PDR2 Disclose data based ...	↑	↑	↑		↑			
PDR3 Inhibit linking across...	↑	↑	↑					
PDR4 Transparency and co...	↑			↑		↑	↑	↑
PDR5 Information security					↑			

B. Business Requirements for FIM

FIM domain experts have stated that these requirements are particularly relevant, because the feasibility of FIM systems is delicate and depends on the balance of the privacy and business requirements. While BR1 is driving AR4, the other business requirements (BR2 and BR3) have to be applied directly to the FIM models.

BR1. While it is good practice to segregate services into separate privacy domains, there need to be links between them for specific purposes like mash-ups and integrated workflows. E.g., an electronic health record and an accounting system will quite likely reside in different privacy domains, but limited linking of personal data across privacy domains can be necessary for audit and billing processes.

BR2. Maximize compatibility with existing FIM protocol profiles to the extent that other requirements are not compromised: (i) Feasible implementation effort. The model shall make use of existing profiles and implementations as far

as reasonable; (ii) Feasible deployment effort. It shall be possible to use implementations of established technology stacks such as SAML2Int [32] and OpenID Connect [34] within current configuration limits.

BR3. Deployment effort must be feasible. As software distribution is still a key cost driver for not centrally managed user devices, deployment requirements such as browser plug-ins or local software should be avoided.

C. Architectural Requirements for FIM

We recovered architectural requirements pertaining to enhanced privacy from the models selected in the expert consultation and from good practice in large FIM systems [13] [14]. To improve their consistency and relevance we linked them with the privacy by design and business requirements laid down above. The related requirements (PDR/BR) are indicated in parentheses.

AR1. Limited observability (PDR1, PDR2). No entity shall be able to aggregate data about the usage of multiple services by users, which will keep it from being able to deduce personal interests or behavior.

AR2. Limited linkability (PDR1, PDR3). Relying parties shall not be able to aggregate personal data used in different privacy domains. Only if it is necessary for a legitimate purpose shall two relying parties processing data of a principal be able to link those data sets. Aside from unique identifiers, this concerns attributes that are identifying with high probability as well.

An important measure is the use of pseudonyms. If the full pseudonymization of user attributes is not feasible, then at least those attributes that identify a user (almost) uniquely shall be pseudonymized. This applies, e.g., to the ubiquitous e-mail address.

AR3. Prevent the unauthorized aggregation of attributes by central intermediaries such as gateways, brokers, etc. (PDR2). No actor shall be able to collect attributes beyond the specified purpose of a service and deduce personal information and behavior.

AR4. Constrained linking (BR1). Unidirectional links between instances of a principal in different privacy domains shall be possible either directly or mediated under control of the user or a third party.

AR5. Consent handling (PDR4). The flow of releasing attributes should regard the processing of user consent, where explicit consent is appropriate.

AR6. No supreme instance (PDR2). Actors managing trust roots must not have access to either attributes or transaction data.

AR7. Minimize the release of attributes (PDR1). The identity provider, in its role as data controller of a principal's identity information, must be assured that only those attributes deemed necessary for the purpose of the service are released to the relying party.

AR8. Uniqueness of identification (PDR2, PDR5). If the RP requires access control, e.g., to protect personal data, then the unambiguous identification of the principal is necessary, even if the user remains pseudonymous. Therefore, identifiers need to be immutable and non-reassignable in those cases.

TABLE II. TRACING ARCHITECTURAL VS. PRIVACY BY DESIGN REQUIREMENTS

	PDR1 Minimal identification	PDR2 Disclose data based on need to know	PDR3 Inhibit linking across privacy domains	PDR4 Transparency and control to the user	PDR5 Information security
AR1 Limited observability	↑	↑			
AR2 Limited linkability	↑		↑		
AR3 No unauthorized aggregation		↑			
AR4 Constrained linking				↑	
AR5 Consent handling					
AR6 No supreme instance		↑			
AR7 Minimized attribute release	↑				
AR8 Unique identification		↑			↑

V. COMPARING ARCHITECTURAL CHOICES

In this section we compare seven FIM models with respect to limited observability (AR1) and then outline options for the other seven architectural requirements.

A. Models for Limited Observability (AR1)

The following list of seven models compares the standard approach (using conventional organizational controls) against six privacy-enhancing approaches to limited observability within the scope of FIM. The initial list was considered representing the most relevant FIM models at a Kantara Federation Interoperability WG expert meeting [24]. FIDO U2F [23] was a further candidate system but was excluded because of the lack of attribute management, although it could improve privacy characteristics of other FIM models.

Other profiling methods to track users, such as IP addresses and device fingerprinting, are out of scope.

1) Organizational Controls

Regulation will mitigate the risk with preventive organizational safeguards, liability and legal enforcement. Providing choice with IdPs and enforcing transparent and efficient markets will drive providers to comply with privacy requirements. Note: this is standard practice but does not implement privacy by design.

Pro: As an analogy this works well in the financial services industry, where banks have a panoptical view on financial

transactions of their clients, but are still trusted due to the fact that the financial services industry is highly regulated.

Con: (a) Illegitimate behavior is not precluded on the factual level. (b) An attack on a single actor (IdP, or hub in the hub-and-spoke model) could lead to a data breach violating AR1.

2) Attribute-Based Credentials

The IdP is taken out of the interaction with the RP, using cryptographic technologies based on group signatures as in IBM's Idemix or blind signatures like Microsoft's uProve. This provides an assertion to the RP without the IdP knowing the actual RPs [25].

Pro: Strong technical control that satisfies the limited observability requirement AR1.

Con: (a) No implementation in mainstream products and no availability of deployment profiles for SAML or OpenID Connect; (b) Issues with other requirements (IdP business model, performance); (c) Increased complexity in crypto-technology.

3) Late Binding

Credential providers provide pseudonymous credentials to users, and RPs will bind attributes to those credentials. This model was proposed by the Government of Canada Cyber-Authentication Architecture. Their separation between credential service assurance and identity assurance implies that attributes are not released by the IdP, but obtained by the RP [15]. Note: The IdP does store identity attributes for account recovery and non-repudiation.

Pro: Straightforward architecture that goes well with existing technology based on common SAML profiles (BR2). Credential providers can easily federate with services, because there is only a minor privacy risk.

Con: (a) While mitigating AR2, it does not fulfill AR1. (b) The collection of identifying attributes like name, residential and e-mail addresses is still likely; thus, AR2 might not be fulfilled very well.




Figure 2. Late Binding model

4) Proxy Pool

Proxies that play RP to an IdP and IdP to an RP can significantly reduce the amount of data collection, if there are many of them operated by independent parties. Each proxy serving only a subset of RPs would not obtain the full profiles of all users.

Pro: This is a use of existing technology, because proxies and gateways for identity management are a well-established technology, e.g., part of the SAML specification (BR2).

Con: (a) Proxies would yield only a very limited improvement on AR1 until the number of proxies is quite large; thus, it would be difficult to overcome the hen-and-egg problem.

5) User-based IdPs

As proposed by IMI [16], the client would be the identity selector and could also hold the credentials locally. A similar concept has been proposed with personal authentication devices [17].

Pro: This architecture provides good support for AR1 with the possible exception of (b) below.

Con: (a) Deployment is hard because it is difficult to enhance web browsers (BR3) and (b) with PKI-based credentials there is still the tracking issue with OSCP responders (AR1). (c) Experience with the “Neuer Personalausweis”, the German national identity card (described in [1]), showed that complex deployment leads to the growth of cloud services that offload some deployment issues but violate AR1 in turn.

6) Constrained Logging Proxy (hub-and-spoke federation)¹¹

A proxy (hub) will hide the target RP from the IdP. The hub thus provides limited observability for IdPs, but is violating AR1 itself. To mitigate this, the gateway does not store log data on the local node, but sends it to a remote system where controls such as encryption and deletion after a short term reduce the risk of abuse.

Pro: Can be implemented without changes to FIM protocols (BR2).

Con: Only a partial technical control. While an adversary could cause only limited damage with a single data breach, a complete take-over of the proxy that would talk home to the adversary would violate AR1.

7) Blind Proxy

The Privacy-enhanced FIM model introduced by the authors in [18] enhances the hub-and-spoke model by offering technical controls that enforce limited observability and enable pseudonymous authentication. Its core property is that attributes are encrypted from the IdP to the RP, but the IdP cannot identify the RP. This is shown in Fig. 3, where a message exchange between RP and IdP is brokered via the blind proxy. As the RP’s encryption certificate is issued per transaction, the IdP can only identify groups of RPs. This model claims to have similar properties as attribute-based credentials in option (2), except that it is not resistant against a collusion of RP and IdP.

Pro: It proposes reasonably strong technical controls (AR1), works with any credential technology and is fairly easy to fit into hub-and-spoke federations (BR2). A SAML profile has been published [27].

Con: (a) Despite not requiring new technology, as attribute-based credentials (2) do, it is still not fully compatible with existing implementations; therefore, it will not run out of the box with existing products. (b) It requires RPs to participate in

a considerably large anonymity set, with each anonymity set having identical conditions for attribute release.




Figure 3. Blind Proxy model for the WebSSO use case

B. Models for Limited Linkability (AR2)

The use of opaque, pairwise identifiers is known as a targeted identifier in research & education federations [13][14], as a sector-specific identifier in government eIDs [19], and called persistent NameId in the SAML specification [20]. This concept is fairly easy to implement and widespread.

However, the problem of linkability using other attributes remains. E-mail address, credit card number, delivery address and name are frequently required and match individuals with high probability. The privacy-enhanced FIM model [18] proposes the use of proxy addresses for email, payment and physical delivery and user-selected pseudonyms for display names.

C. Options for Proxies (AR3)

Proxies used in hub-and-spoke¹² models could potentially aggregate user data. Controls to mitigate this are (a) not providing directly identifying data to proxies as in the Late Binding (3) and Blind Proxy (7) models, or (b) constraining logging as in the Constrained Logging Proxy model (6). A practical difficulty is that log information is essential for tracing technical problems. Therefore a process must be established to open up log information to trusted operators.

D. Options for Constrained Linking (AR4)

Unidirectional links have been defined, for example in the Austrian eID using encrypted sector-specific identifiers. This concept uses sector-specific pairwise identifiers encrypted for the target application. On a more general level, constrained links can be direct or mediated. Direct links, as in the Austrian eID system, are durable, whereas mediated links can be established by a broker for a specific transaction only. The latter would allow, for example, that a user consents to use a payment clearing service a single time, without leaving a possibility for either service to link the personal data later on.

E. Options for Consent Handling (AR5)

User consent is a function to notify the user about the intended release of attributes to an RP and to obtain a

¹¹ The Danish Academic Access Federation (WAYF) [26] does implement this model, but the details about the logging policy have not been published.

¹² Hub-and-spoke model: All interactions between SP and IdP are brokered via a hub.

permission to do so. Whereas systems with smart clients may use local storage for consent, hub-and-spoke systems that obey limited observability struggle with storing a link between an IdP and a service. The Blind Proxy (7) can store consent without risk, because it cannot relate the user to a real person; however, the user interface is constrained to display non-identifying data. In the Late Binding model (3), attributes are collected by the RP and any consent would be requested there.

F. Options for Avoiding a Supreme Instance (AR6)

It is necessary in any model to separate the organizational roles so that the risk of impersonation and subversion of encryption is minimized.

G. Options to Minimize Attribute Release (AR7)

Data minimization by releasing only attributes required for the purpose of the service is well-established in research and education federations and can thus be considered state-of-the-art from a research perspective. There are ongoing efforts to make attribute release saleable by categorizing services to minimize administrative efforts [21].

Further optimization can be achieved with deriving less identifying attributes, such as with age verification and locality. This concept has not yet entered mainstream adoption.

H. Options for the Uniqueness of Identification (AR8)

Unambiguous identification implies the creation of a unique identifier at the IdP during the registration process of a user, e.g., by collecting a sufficient number of attributes that will distinguish persons who share the name and date of birth. While this is common in countries with a central citizen registry, it might not be feasible in other places. The privacy impact needs to be assessed on a case-by-case basis.

VI. CONCLUSION

In this paper, we elaborated privacy by design requirements for FIM systems based on fundamental privacy principles, aiming at technical controls to preclude the privacy-infringing use of the system. The paper demonstrates that technical controls can be applied to the basic system architecture. While partial solutions have found their way into operational systems, solutions that are complete and sound are still in early stages.

Realizing limited observability is a key question in this paper. We found that controls range from strong to weak, with implementation effort correlating from high to low. A trade-off based on privacy risk, incentives and cost will have to be chosen case by case.

Finally, we noticed that current FIM systems deployments frequently implement pairwise identifiers to reduce linkability. While this is good practice, the problem of linkability through identifying attributes such as email address and name is rarely addressed. The resolution of this problem with respect to the ubiquitous email address as described in section V.B seems to be a low-hanging fruit. It is technically simple and suitable to several FIM models.

REFERENCES

- [1] E. Schweighofer and W. Hötendorfer, "Electronic identities – public or private," *International Review of Law, Computers & Technology*, vol. 27, no. 1–2, 2013, pp. 230–239.

- [2] R. Hörbe and W. Hötendorfer, "Privacy-by-Design-Anforderungen für das Federated Identity Management", in *Jahrbuch Datenschutzrecht* 2014, D. Jähnel, Ed. Wien and Graz, Austria: NWV, 2014, pp. 305–325.
- [3] D. W. Chadwick, "Federated identity management", in *Foundations of Security Analysis and Design V*, Lecture Notes in Computer Science, vol. 5705, A. Aldini, G. Barthe and R. Gorrieri, Eds. Berlin and Heidelberg, Germany: Springer, 2009, pp. 96–120.
- [4] E. Ghazizadeh, M. Zamani, J. L. Ab Manan, and A. Pashang, "A survey on security issues of federated identity in the cloud computing", in *CloudCom 2012 - Proceedings: 2012 4th IEEE International Conference on Cloud Computing Technology and Science*, 2012, pp. 562–565.
- [5] M. Ferdous and R. Poet, "A comparative analysis of identity management systems", *High Performance Computing and Simulation (HPCS)*, 2012 International Conference on, IEEE, 2012 pp. 454–461.
- [6] J. van Rest, D. Boonstra, M. Everts, M. van Rijn, and R. van Paassen, "Designing privacy-by-design", in *First Annual Privacy Forum, APF 2012*, Lecture Notes on Computer Science, vol. 8319, B. Preneel and D. Ikonomou, Eds. Berlin and Heidelberg, Germany: Springer, 2014, pp. 55–72.
- [7] S. S. Shapiro, "Privacy by design", *Communications of the ACM*, vol. 53, no. 6, 2010, pp. 27–29.
- [8] U. Dammann and S. Simitis, *EG-Datenschutzrichtlinie: Kommentar*. Baden-Baden, Germany: Nomos, 1997.
- [9] S. Gürses, C. Troncoso, and C. Diaz, "Engineering privacy by design", *Computers, Privacy & Data Protection* 14, 2011.
- [10] K. Cameron, "The laws of identity", Whitepaper, Microsoft Corp, 2005. Available at <http://myinstantid.com/laws.pdf>.
- [11] H. Zwingelberg, M. Hansen, "Privacy protection goals and their implications for eID systems", in *Privacy and Identity Management for Life*, J. Camenisch, B. Crispo, S. Fischer-Hübner, R. Leenes and G. Russello, Eds. Berlin and Heidelberg, Germany: Springer, 2012, pp. 245–260.
- [12] B. Schneier, *Secrets and Lies*. New York, NY: John Wiley & Sons, 2000.
- [13] SWAMID, Identity Provider Privacy Policy template, 2014. Available at <https://portal.nordu.net/display/SWAMID/SWAMID+template+Identity+Provider+Privacy+Policy>.
- [14] Internet 2, *InCommon Federation Recommended Practices*, 2013. Available at <https://spaces.internet2.edu/display/InCFederation/Recommended+Practices>.
- [15] Chief Information Officer Branch, Treasury Board of Canada Secretariat, "Federating Identity Management in the Government of Canada: A Backgrounder", 2013. available at: <http://www.tbs-sct.gc.ca/sim-gsi/docs/2011/fimgc-fgigc/fimgc-fgigcpr-eng.asp>
- [16] OASIS, OASIS Identity Metasystem Interoperability (IMI) Technical Committee Charter, 2008. Available at: <https://www.oasis-open.org/committees/imi/charter.php>.
- [17] A. Jøsang, M. Al Zomai, and S. Suriadi, "Usability and privacy in identity management architectures", in *Proceedings of the fifth Australasian symposium on ACSW frontiers - Volume 68 (ACSW '07)*, L. Brankovic, P. Coddington, J. F. Roddick, C. Stekete, J. R. Warren, and A. Wendelborn, Eds. Darlinghurst, Australia: Australian Computer Society, 2007, pp. 143–152.
- [18] R. Hörbe, "A model for privacy-enhanced federated identity management", *arXiv preprint arXiv:1401.4726*, 2014.
- [19] G. Aichholzer, S. Strauß, "The Austrian case: multi-card concept and the relationship between citizen ID and social security cards", *Identity in the Information Society*, vol. 3, no. 1, 2010, pp. 65–85.
- [20] OASIS, *Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0*, 2005. Available at <http://docs.oasis-open.org/security/saml/v2.0/>.
- [21] TERENA, "REFEDS Entity Category helps organizations release data securely", *TERENA News*, 2014. Available at https://www.terena.org/news/fullstory.php?news_id=3648.
- [22] OASIS, "Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0", ed, 2005.

- [23] FIDO Alliance. (2014). Universal 2nd Factor (U2F) Overview. Available: <http://www.fidoalliance.org/specs/fido-u2f-v1.0-ps-20141009/fido-u2f-overview-v1.0-ps-20141009.html>
- [24] Kantara eGov-WG, 2013-02-04 eGov Meeting Minutes, <https://kantarainitiative.org/confluence/x/vITEAw>
- [25] Mikkelsen, G. L., Damgård, K., Guldager, H., Jensen, J. L., Luna, J. G., Nielsen, J. D., ... & Zhang, H. (2015). Technical Implementation and Feasibility. In Attribute-based Credentials for Trust (pp. 255-317). Springer International Publishing.
- [26] D. Simonsen, Trusted third party based ID federation, enhancing privacy and lowering the bar for connecting http://www.wayf.dk/wayfweb/faq_-_ofte_stillede_sprgsmaal_attachmt/2008_11_23_tnc_abstract_trusted_third_party_based_id_federation_enhancing_privacy_and_lowering_the_bar_for_connecting.pdf, 2008
- [27] R. Hörbe, "SAML Profile for Privacy-enhanced Federated Identity Management", ed: Kantara Initiative, 2014.
- [28] Wikipedia, SAML-based products and services https://en.wikipedia.org/wiki/SAML-based_products_and_services
- [29] Kantara BCTF-DG. (2012). Global Trust Framework Survey. Available: <https://kantarainitiative.org/confluence/display/bctf/Global+Trust+Framework+Survey>
- [30] IETF, RFC 6759 The OAuth 2.0 Authorization Framework, 2012 <http://tools.ietf.org/html/rfc6749>
- [31] IETF, User-Managed Access (UMA) Profile of OAuth 2.0, 2015 <https://tools.ietf.org/html/draft-hardjono-oauth-umacore-12>
- [32] Kantara eGov-WG, SAML Interoperability and Deployment Profiles, 2014, <https://kantarainitiative.org/confluence/x/SgCgAw>
- [33] OpenId Foundation, OpenID Project Homepage, 2015, <http://openid.net/connect/>
- [34] J. Jensen, "Federated Identity Management in the Norwegian Oil and Gas Industry", 2014 (p.56)