Federated Authorization over Access to Personal Data for Decentralized Identity Management

Thomas Hardjono

Abstract

The digital identity problem is a complex one in large part because it involves personal data, the algorithms which compute insights on the data, and the management of the identifiers that are linked to personal data. The reality of today is that personal data of an individual is distributed throughout the Internet, in both private and public institutions, and increasingly also on the user's devices. In order to empower individuals to have a say in who has access to their personal data and to enable individuals to make use of their data for their own purposes, a coherent and scalable federated authorization architecture is required as a fundamental component of the decentralized identity solution. This federation must allow an individual to easily manage access policies, and to grant and retract consent for access to data distributed across multiple repositories. This article describes the User Managed Access architecture and protocols that provide the foundation for scalable federated authorization.

Introduction: Data, Identity and Trust

The advent of Bitcoin and blockchain technology generally has illustrated the potential use of public key cryptography in the context of transactions conducted directly between parties over the peer-to-peer network of nodes in an unmediated fashion. This has raised interest in the potential use of blockchain technology in the area of digital identity. However, the question of digital identity pre-dates blockchain technology and is closely intertwined with the complex problem of personal data and privacy.

The increasing awareness of the public regarding the power of big data and the recent cyber-attacks on organizations holding large amounts of consumer personal data (e.g., Anthem data breach in 2015; Equifax breach in 2017) has shaped public trust in recent years. Over the last decade there has been a continuing decline in trust on the part of individuals with regard to the handling and fair use of personal data [1]. Pew Research reported that 91 percent of Americans agree or strongly agree that consumers have lost control over how personal data is collected and used, while 80 percent who use social networking sites are concerned about third parties accessing their shared data [2]. The Webbmedia Group, writing in the *Harvard Business Review*, has identified data privacy as one of the top 10 technology trends of 2015 [3]. Related to the loss of trust — and perhaps as a consequence of it — is the recent development of new regulations aimed at addressing data privacy. The enactment of the EU General Data Protection Regulation (GDPR) has in turn influenced the data privacy discourse in the United States and elsewhere (e.g., the California Consumer Privacy Act [(CCPA]).

The state of declining trust was already reported by the World Economic Forum in 2014. The WEF report [4] was the culmination of a multi-year initiative with global insights from various high-level leaders from different sectors of society (industry, governments, civil society and academia). A theme running through the 2014 WEF report is the *need to strengthen individual trust*. The WEF report suggests three means to address this problem [4, p. 14]. First, increase *transparency* by focusing on engagement and response, and by providing individuals with insight and meaningful control. This is instead of the current approaches focusing on disclosure and providing details (which often overwhelm individuals). Second, improve *accountability* by orienting throughout the value chain (front-end to back-end) with risks being equitably distributed. This is in contrast to the current industry practices that are oriented toward the front-end of the value chain with risks and responsibilities residing with the individual. Third, *empower individuals* by way of giving them a say in how data about them is used by organizations and giving individuals the capacity to use data for their own purposes. Empowerment should be distributed with shared incentives for empowering individuals and distributing value closer to the source of data production (the individual). This is in contrast to the current approaches, which are focused on maintaining information differentials among a concentrated set of actors. Echoing the WEF Report [4], we believe that individuals need *meaningful control* (Fig. 1) over their personal data, which is increasingly distributed across various entities on the Internet [3].

Today, the reality is that personal data typically does not reside with the individual, for practical reasons. The GDPR recognized this reality, and reflected it by using the notion of *data control-*

Digital Object Identifier:
10.1109/MCOMSTD.001.1900019

Thomas Hardjono is with MIT Connection Science & Engineering.

2471-2825/19/$25.00 © 2019 IEEE

lers and *data processors*. In order to provide an individual with true meaningful control over their personal data, the controllers (of an individual's personal data) must collectively provide an easy way for the individual to configure access policies (consent rules) that will apply to the personal data located at each of the controllers. We refer to this as *authorization federation*. The overall goal of federated authorization is to empower the individual to set access policies (i.e., consent) at one location (e.g., at one data controller) and have the access policies propagated automatically to other data controllers and be enforced there also. In this way, the individual is relieved from having to log-in to many sites for the purpose of configuring the access policies multiple times. This is consistent with and follows from the WEF recommendations.

From the perspective of data minimization and privacy, industries which hold and process personal data should adopt a data handling philosophy that favors *sharing insights* instead of exchanging raw data. We refer to this as the *open algorithms* paradigm that was first developed at MIT [5]. The open algorithms approach advocates that:

- Data should never leave its repository.
- The vetted algorithms are instead transmitted to the data repository to be executed there.
- Only aggregate answers are returned, which do not permit the re-identification of individuals.

Any algorithm execution yielding a response that goes deeper or finer-grained than aggregate results must first obtain explicit consent from the individuals concerned.

The goal of this article is to explore the notion of a *federated authorization* model for the decentralized management of the various aspects of an individual's identity. Following from the open algorithms approach, the term "authorization" in this article is generally taken to mean an individual's permission or consent to have a vetted algorithm be executed over their personal data. In order to ground the discussion in real-world applications, we use a specific example called User Managed Access (UMA), which provides the foundation for the authorization federation among the data controllers.

Decentralized Identity Management

As mentioned previously, the question of individual digital identity has always been related to personal data and privacy. In discussing digital identity management, we adopt a more encompassing view that incorporates not only digital identifiers (e.g., email addresses, Social Security numbers, public keys), but also:

- Personal data that are distributed throughout the Internet
- The algorithms that are applied to personal data resulting in insights
- The assertion or claim structures that convey the various degrees of insights to external entities
- The digital identifiers that are linked or bound to the assertions

We define *identity* generally as the collective aspect of the set of characteristics or features by which a thing (e.g., human, device, organization) is recognizable and distinguishable one from another [6]. In the context of a human person, the individuality of a person plays an important role in that it allows a community of people to recognize the distinct characteristics of an individual person and consider the person as a persisting entity. It is this individuality that leads people to form social networks, and allows them to interact with members of their social network and exchange ideas [7]. Data regarding an individual's social network interactions is an important part of the corpus of personal data of that individual. Thus, we define the *core identity* of a person as a collective aspect of the set of characteristics (derived through personal data, algorithms, and insights) by which a person is uniquely recognizable.

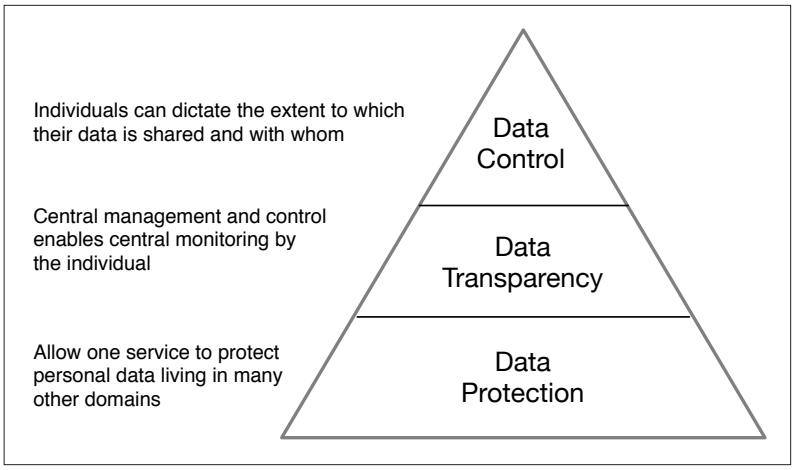

Figure 1. The vision of privacy 2.0 (after [3]).

Proper identity management, therefore, becomes the broader challenge of how an individual can better control and manage the various aspects of their digital identity life cycle, including the creation and usage of their personal data, algorithms, insights and assertions that are distributed over the Internet. Thus, correctly viewed, the problem of *decentralized identity management* is in reality the challenge of achieving decentralized control by individuals over their digital characteristics, expressed through personal data, algorithms, insights, and assertions made about them. We refer to this as *individual-centric identity management*, which is in contrast to the traditional centralized institution-centric identity management. The emphasis on the individual-centricity means giving individuals the capacity to use data for their own purposes and bringing value closer to the source of data production, namely the individual person. The fact that today much personal data are distributed across the Internet means that a federated authorization model for individual-centric access is an important component of the decentralized identity solution.

For the nascent field of blockchain technology, there are several potential applications of blockchains beyond the binding of public keys to claims via a blockchain [8, 9]. For example, blockchains could be employed as the foundation technology for the logging, tracking, and auditing of the algorithms that were applied to data by the relevant entities involved. As another example, the granting of consent (e.g., to run algorithms)

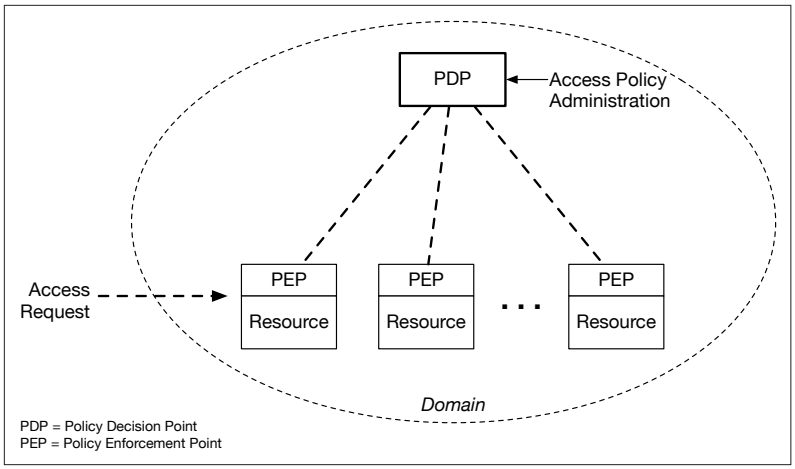

FIGURE 2. Overview of policy-based access control with PDPs and PEPs.

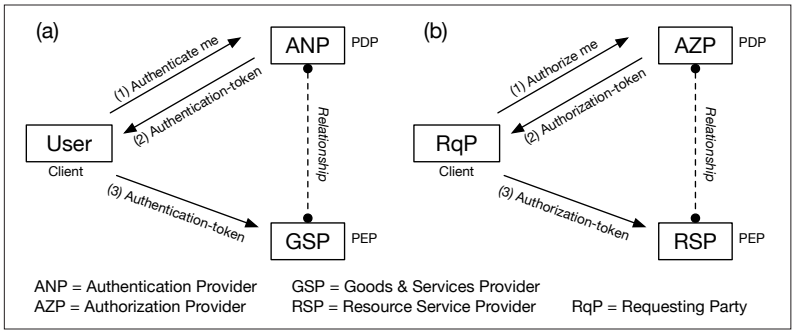

FIGURE 3. Overview of a) authentication provider; b) authorization provider.

and their corresponding receipts can be captured using the consent-receipt standard protocol, recorded on a blockchain. Finally, the "smart contracts" (stored procedures) capability of some blockchain systems could be the basis for expressing algorithms in the form of smart contracts, with remuneration going directly to the individual owners of the personal data and the authors of the algorithms [10].

Policy-Based Access Control and Authorization

The issue of controlling access to multi-user resources has been an important theme since the mid-1960s, with the rise of time-share mainframe computers. Generally, the term *access control* is applied not only to physical access (to the computer systems) but also to system resources (e.g., memory, disk, files). Notable among the early efforts in the 1970s was the Multics system. In the context of government and military applications, there was the further issue of access based on a person's rank or security clearance. Here, the concept of *mandatory* and *discretionary* access control in multi-level systems came to the forefront in the form of the Bell and LaPadula Model (BLM) [11].

In the BLM, access control is defined in terms of subjects possessing different *security levels*, seeking access to *objects* (i.e., system resources). Thus, for example, in the BLM a subject (e.g., user) is permitted to access an object (e.g., file) if the subject's security level (e.g., "Top Secret") is higher than the security level of the object (e.g., "Secret"). The notion of roles or capacities was added to this model, leading to the Role-Based Access Control (RBAC) model. Here, as a further refinement of the BLM, a subject (user) may have multiple roles or capacities within a given organization. Thus, when the subject is seeking access to an object, he or she must indicate the role within which the request is being made. The formal model for RBAC was subsequently defined by NIST in 1992 [12].

Access control to resources is also a major concern for enterprises and corporations. This need became acute with the widespread adoption of local area network (LAN) technology by enterprise organizations in the 1990s. The same RBAC model also applies to corporate resources attached to the corporate LAN. This problem was often referred to as authentication, authorization, and auditing (AAA) in the 1990s. Part of the AAA model developed during the 1990s was an abstraction of functions pertaining to *deciding* access rules from functions pertaining to *enforcing* them. Entities that decided on access rules were denoted as policy decision points (PDPs), while entities that enforced these access rules were denoted as policy enforcement points (PEP). Figure 2 summarizes this abstraction.

The policy-based access control model is foundational to many systems deployed within enterprises today. Many solutions, such as Microsoft's Active Directory (AD), are built on the same model of policy-based access control. In the case of AD, a fairly sophisticated cross-domain architecture was developed that allows an enterprise to logically arrange itself into dozens to hundreds of interior domains (e.g., each department as a different AD group). Permissions and entitlements for subjects (employees) in AD are expressed in a comprehensive Privilege Attribute Certificate (PAC) data structure. Interestingly, the main authentication mechanism within Microsoft AD and many similar products is the MIT Kerberos authentication system (RFC 1510).

Mediated Authentication and Authorization

Today, there are a number of entities on the Internet that mediate transactions between an individual user and an online service provider. One key offering of many of these entities is *mediated authentication* services. Here, in discussing mediated authentication we use a slightly modified terminology for clarity of discussion. This is to avoid employing industry jargon, which is often inaccurate and a product of historical development. For example, instead of using the generic term "service provider" for the entities that offer a broad range of offerings, we use the more specific terms of goods and services provider (GSP) and resource service provider (RSP). The first denotes entities that offer goods (e.g., Amazon), while the later denote entities that offer computer-related resources including cloud-based storage (e.g., DropBox), compute capabilities (e.g., AWS/EC2), and others. In many cases the role of resource service providers is to support the business of the GSPs (e.g., online merchants).

The functions of mediated authentication and authorization are delivered by two kinds of third-party providers (Fig. 3).

Authentication Provider (ANP): A mediating authentication provider has the task of managing and validating a user's credential (e.g., password, keys) on behalf of a GSP entity (e.g., online merchant). This allows the GSP to be relieved of the task of authenticating the user (e.g., customer of the merchant). Consequently the ANP also has the task of managing the credentials belonging to users (customers) on an ongoing basis. Typically, the GSP must have a business relationship with the ANP before the customer can perform authentication to the ANP. This is shown in Fig. 3a.

There are several variations of the protocols for mediated authentication. Generally, the ANP issues an *authentication-token* as proof of the user's successful authentication event at the ANP. The authentication-token can be delivered to the GSP via the client software (front channel) deployed by the user, or the token can be delivered directly from the ANP to the GSP (back channel). An example of these tokens are Kerberos tickets and SAML2.0 login assertions.

Today, the ANP function is fulfilled by a category of providers referred to as identity providers (IdP). The typical consumer-facing IdP issues an identifier (e.g., email address) and manages the credentials of the user (e.g., change password). When the user seeks to access services offered by the GSP, the user is temporarily redirected to the IdP for authentication. The IdP issues an authentication-token, which can then be validated by the GSP.

Authorization Provider (AZP): A mediating AZP has the task of managing access policies pertaining to access to resources such as files, documents, and media. The resources typically reside at one or more resource service provider (RSP) entities, and the owner of the resource sets the access policies at the AZP entity (Fig. 3b). A back-channel typically exists between the AZP and the RSPs, permitting the policy rules and configuration settings (set by the resource owner) to be communicated from the AZP to the RSPs. Looking at Fig. 3b, the AZP implements function of the PDP, while the RSP implements the function of the PEP.

When a third party — referred to as the Requesting Party (RqP) — seeks to access a given resource at an RSP, it must first be authenticated by the relevant ANP entity who issues it an authentication-token. The ANP is assumed to have a business relationship with the AZP. The requesting party presents the authentication-token to the AZP entity as proof that the requesting party has been authenticated. The AZP in turn issues an *authorization-token* (e.g., OAuth2.0 token, Microsoft PAC) as a means to convey the access privileges assigned to the requesting party for given resource at the relevant RSP. Currently, most consumer-facing resource sharing providers (e.g., photo or calendar sharing sites) merge together the functions of the AZP and RSP.

In order to scale up services, over the years a number of ANPs in the consumer space have banded together to form consortiums that provide their members with a broader reach for their services collectively. We use the term "authentication federation" for this kind of consortium arrangement. The goal of an authentication federation is essentially to help the GSP entities to ensure that a new or returning user (i.e., customer) can be authenticated quickly. To achieve this efficiency, a GSP enters into a business relationship with either an ANP who is a member of the federation or the federation organization directly. An authentication federation typically operates under a set of bylaws and contracts, referred to as the federation Legal Trust Framework (LTF) for the consortium (e.g., OpenID-Exchange or OIX).

Federation of Mediated Authorization Services

Similar to authentication federation, in order for authorization architectures in the consumer space to be able to scale up, an authorization federation among the providers is needed. To place authorization federation in the proper context, we use the classic policy-based resource access control model [12] as our starting point (shown earlier in Fig. 2). This is applied to a collection of *domains*, each representing distinct data controllers (holding personal data of various individuals). In Fig. 4, both Domain 1 and Domain 2 hold resources associated with an individual, which we refer to as the *data subject* (or simply *subject*) following the GDPR definition. The subject as the resource owner has data located at both Domain 1 and Domain 2. A third party, denoted as the *requesting party*, seeks access to the subject's data located in Domain 1 (e.g., to execute an algorithm on the data in Domain 1).

There are at least three (3) goals for a scalable federated authorization model.

Cross-Domain Policy Propagation and Enforcement: A subject (resource owner) must be able to set access policies in one domain, and have the policies automatically propagated to all domains in the federation that contain the subject's resources and have those policies enforced locally by each relevant domain.

An example is illustrated in Fig. 4, where the subject sets access policies at PDP2 in Domain 2, while enforcement also occurs in Domain 1 to resources at PEP1.1, where the subject's resources reside.

Decentralization of Enforcement: Once an access policy is decided at one PDP in one domain, enforcement within all domains in the federation that contain the subject's data/resources must occur automatically without the subject's further involvement. Each PEP in each relevant domain must operate independent of other PEPs in the same domain or other domains.

Legal Trust Framework for Authorization Federation: A legal trust framework must be agreed upon by all domain owners in the federation, one that defines, among others, the agreed behavior of PDPs and PEPs in propagating access policies and enforcing them.

In the next section we discuss the UMA architecture as one of the embodiments of federated authorization concept.

Federated Authorization for Personal Data: UMA

The goal of the User Managed Access (UMA) architecture is to provide individual-centric control over "resources" (e.g., personal data, algorithms,

In order to scale up services, over the years a number of ANPs in the consumer space have banded together to form consortiums that provide their members with a broader reach for their services collectively. We use the term "authentication federation" for this kind of consortium arrangement.

Once an access policy is decided at one PDP in one domain, enforcement within all domains in the federation that contain the subject's data/resources must occur automatically without the subject's further involvement. Each PEP in each relevant domain must operate independently of other PEPs in the same domain or other domains.

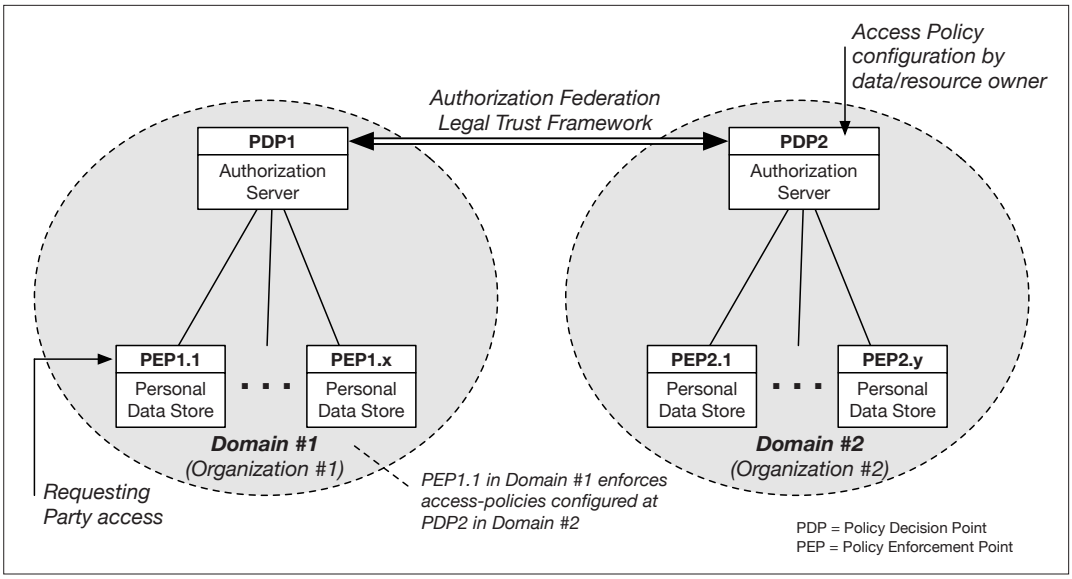

FIGURE 4. Overview of an authorization federation of two domains or institutions.

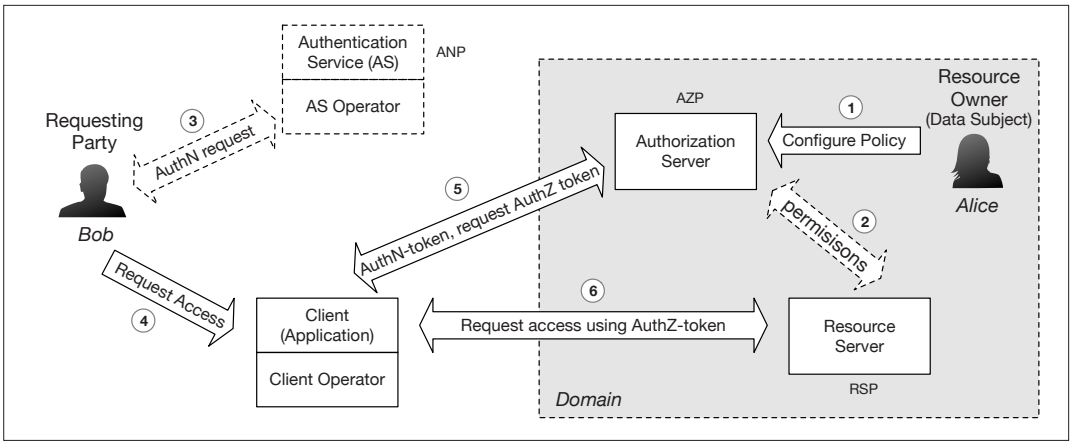

FIGURE 5. Overview of UMA architecture (single domain).

assertions) that may be distributed across multiple locations, each employing a resource server as the RSP entity (Fig. 3b). The basic idea of UMA is that the data-subject as the resource owner (RO) would set access policies at one AZP entity, and for the access policies to be propagated automatically to all RSPs who hold resources (i.e., data) belonging to the data-subject and be enforced by each of the RSPs independently. When a requesting party (RqP) seeks access to a given resource protected by an RSP entity, the requesting party must first obtain an authorization-token from the AZP, and deliver it to the RSP with its access request.

In other words, the UMA architecture is an embodiment of the mediated authorization function discussed in the previous two sections and shown in Fig. 3b. The UMA standard began to be developed in the Kantara Initiative in 2009, with the UMA Version 1.0 specifications being completed in 2014 [13] and the Version 2.0 specifications published in 2017 [14]. Throughout its decade-long development, the UMA philosophy has been consistent with much of the data privacy discourse in the World Economic Forum [1, 4] and with the GDPR notion of privacy and consent.

Due to the popularity of the OAuth2.0 framework [15] among social media providers (which already supported the sharing of low-value resources such as photos and calendars), the UMA architecture started by employing the basic OAuth2.0 terminology and technical constructs, including the token structures and access grant flows. Compared to the classic RBAC model, the OAuth2.0 framework is fairly rudimentary in that it only recognized three entities in its ecosystem. These were the *Client* (understood commonly to be a web-based application or mobile app), the *Authorization Server* (as the authorization provider or AZP in Fig. 3b), and the *Resource Server* (as the resource service provider or RSP in Fig. 3b). UMA adopted these three OAuth2.0 entities as the starting design.

However, in addition, UMA crucially introduced the requesting party (RqP) and the resource owner (RO) as the data-subject (data owner) in Version 1.0 of the UMA specifications [13]. These additions placed UMA in line with the well-understood RBAC model discussed earlier [11, 12]. UMA explicitly defined the requesting party as a separate legal entity from the client-operator, whereas OAuth2.0 only recognized the client as a piece of software. The explicit separation

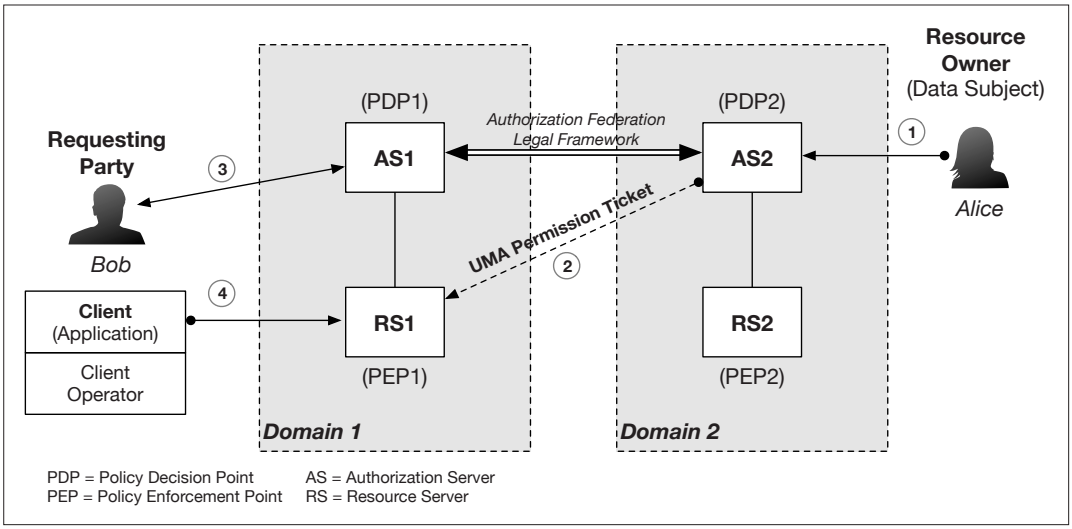

FIGURE 6. Overview of federated authorization in UMA 2.0 (two domains).

of the requesting party from client-operator has dramatic implications in that it forced the recognition that in real-world scenarios there is always someone (i.e., a person or organization) that is deploying or manipulating a piece of client application software — which is accessing the data/resources located at the resource server (RS). This means that the privacy-related legal obligations stemming from the data-subject's consent applies to both participants as legal entities.

In terms of the UMA protocol flows (Fig.5), this means that both the requesting party and the client-operator must obtain distinct authorization-tokens:

- The requesting party must be authenticated by an ANP and be issued an authentication-token (Fig. 5, Step 3).
- Both the requesting party and the client-operator must obtain separate authorization-tokens from the authorization server Fig. 5, Step 5) prior to requesting access to the resource server (Step 6).

The UMA architecture is designed to support federated authorization across multiple domains (Fig. 6). The standardized data structure used to propagate the access policies between the authorization server (as the PDP) and the multiple resource servers (as the PEPs) is the *permissions ticket*. UMA itself is agnostic to the policy expression syntax or language and can support any policy syntax (e.g., XACML).

Figure 6 illustrates the case where Alice as the resource owner has personal data at both RS1 and RS2. Alice sets her access policies at one location only, namely at AS2. Bob as the requesting party is seeking access to Alice's resources at RS1. The UMA architecture supports the propagation of the permissions ticket from the AS2 (PDP2) in the originating domain (Domain 2) to the RS1 (PEP1) in the enforcing domain (Domain 1). The signed permissions ticket can be propagated from Domain 2 to Domain 1 directly from AS2 to RS1, or it can be propagated cross-domain from AS2 to AS1, followed by a local transfer from AS1 to RS1. This allows RS1 to enforce Alice's access policies even though Alice may subsequently be off-line.

Conclusions

In this article we have focused on the notion of a federated authorization model as applied to personal data held by various data controllers, and the importance of decentralized control by an individual over access by external parties to personal data distributed across these controllers. In discussing digital identity management, we have adopted a more encompassing view that includes not only digital identifiers, but also personal data, the algorithms that are applied to personal data yielding insights, the assertion or claim structures that convey the various insights to external entities, and the digital identifiers that are linked or bound to the assertions.

When correctly understood, the problem of decentralized identity management is in reality the challenge of achieving decentralized control by individuals over their digital characteristics, which are expressed through personal data, algorithms, insights, and assertions about them. We refer to this as individual-centric identity management. The emphasis on the individual-centricity means giving individuals the capacity to use data for their own purposes and bringing value closer to the source of data production, namely the individual person.

We have reviewed the concepts of mediated authentication and mediated authorization as the basis to understand federated authentication and federated authorization. We discuss the User Managed Access architecture, which implements a federated authorization model. UMA provides individual-centric control and policy setting capabilities that gives individuals better control over their personal data. The fact that today much personal data are distributed across the Internet means that a federated authorization model for individual-centric identity management is a crucial component of any decentralized identity solution.

Acknowledgments

We acknowledge the following for their tremendous support for the UMA effort since its inception in 2009: Eve Maler, Maciej Machulak, Domenico Catalano, George Fletcher, Mike Schwartz, Justin Richer, Sal D'Agostino, Tim Reini-

When correctly understood, the problem of decentralized identity management is in reality the challenge of achieving decentralized control by individuals over their digital characteristics, which are expressed through personal data, algorithms, insights and assertions about them. We refer to this as individual-centric identity management.

The fact that today much of personal data are distributed across the Internet means that a federated authorization model for individual-centric identity management is a crucial component of any decentralized identity solution.

ger, Mark Lizar, Colin Wallis, Sandy Pentland, and Justin Anderson.

References

[1] World Economic Forum, "Personal Data: The Emergence of a New Asset Class," 2011; http://www.weforum.org/reports/ personal-data-emergence-new-asset-class.
[2] M. Madden, "Public Perceptions of Privacy and Security in the Post-Snowden Era," Nov. 2014; http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/.
[3] E. Maler, "Extending the Power of Consent with User-Managed Access: A Standard Architecture for Asynchronous, Centralizable, Internet-Scalable Consent," *Proc. 2015 IEEE Security and Privacy Workshops*, San Jose, CA, May 2015. DOI: 10.1109/SPW.2015.34.
[4] World Economic Forum, "Rethinking Personal Data: A New Lens for Strengthening Trust," May 2014; http://reports.weforum.org/rethinking-personal-data.
[5] T. Hardjono and A. Pentland, "MIT Open Algorithms," *Trusted Data — A New Framework for Identity and Data Sharing*, T. Hardjono, A. Pentland, and D. Shrier, Eds. MIT Press, 2019.
[6] The Jericho Forum, "Identity Commandments," The Open Group, 2011; www.opengroup.org.
[7] A. Pentland, *Social Physics: How Social Networks Can Make Us Smarter*, Penguin Books, 2015.
[8] D. Reed and M. Sporny, "Decentralized Identifiers (DIDs) v0.11," W3C, Draft Community Group Report 09 July 2018; https://w3c-ccg.github.io/did-spec/.
[9] M. Sporny, D. Longley, and D. Chadwick, "Verifiable Credentials Data Model 1.0," W3C Candidate Rec., Mar. 2019; https://www.w3.org/TR/verifiable-claims-data-model.
[10] T. Hardjono, K. Erhardt, and A. Pentland, "Open Algorithms as Smart Contracts: Enabling Future Data Markets Using Blockchain Technology," *Proc. ICIS Wksp. Opportunities and Challenges of Blockchain Technology*, Seoul, Korea, Dec. 2017.
[11] D. E. Bell and L. J. LaPadula, "Secure Computer Systems: Mathematical Foundations," The MITRE Corp., Tech. Rep. MTR-2547 I ESD-TR-73-278 (Vol. I-II), Nov. 1973.
[12] D. F. Ferraiolo and D. R. Kuhn, "Role-Based Access Controls," *Proc. 15th National Computer Security Conf.*, Baltimore, MD, Oct. 1992, pp. 554–63; https://csrc.nist.gov/CSRC/media/Publications/conference-paper/1992/10/13/role- based-access-controls/documents/ferraiolo-kuhn-92.pdf.
[13] T. Hardjono *et al.*, "User-Managed Access (UMA) Profile of OAuth2.0, Specification Version 1.0," Kantara Initiative, Kantara Published Spec., Apr. 2015; https://docs.kantarainitiative.org/uma/rec-uma-core.html.
[14] E. Maler, M. Machulak, and J. Richer, "User-Managed Access (UMA) 2.0," Kantara Initiative, Kantara Published Spec., Jan. 2017; https://docs.kantarainitiative.org/uma/ed/uma-core-2.0-10.html.
[15] D. Hardt, "The OAuth 2.0 Authorization Framework," Oct. 2012, RFC 6749; http://tools.ietf.org/rfc/rfc6749.txt.

Biography

Thomas Hardjono is currently the CTO of Connection Science and technical director of the MIT Trust-Data Consortium, located at MIT, Cambridge, Massachusetts. For several years prior to this he was the executive director of the MIT Kerberos Consortium, championing the Kerberos protocol to become the most ubiquitously deployed authentication protocol in the world today. Over the past two decades, he has held various industry technical leadership roles, including Distinguished Engineer at Bay Networks, principal scientist at VeriSign-PKI, and CTO roles at several startups. He has been at the forefront of several industry initiatives around identity, trust, and cybersecurity. His areas of interest include IoT security, trusted computing, decentralized identity, personal data privacy, P2P networks, and blockchain systems. He has authored several technical papers, patents, and books covering cryptography, network security, identity, and blockchain security.