An Overview of Identity Management Systems Usage

Petar Djerasimović
University of Zagreb
Faculty of Electrical Engineering and Computing
Department of Applied Computing
Unska 3, Zagreb, Croatia
petar.djerasimovic@fer.hr

Gordan Gledec
University of Zagreb
Faculty of Electrical Engineering and Computing
Department of Applied Computing
Unska 3, Zagreb, Croatia
gordan.gledec@fer.hr

Abstract. *Identity management systems have been a staple of enterprise and academic digital landscapes for decades, practically since the first computer networks were organized. Today, the rising importance of online services in our daily lives, as well as growth of complexity and scale of online services and associated systems, mandates continuous development of identity management systems and their use in completely new scenarios. This paper provides an overview of the current state of identity management systems usage, challenges posed to successful development and deployment of these systems and related technologies; provides a brief overview of the most popular technologies used and estimates of their real-world use.*

Keywords. Identity management, identity, security, privacy, IMS

1 Introduction

Originally, the need for organized identity management arose in the enterprise. Solutions built upon protocols like LDAP and Kerberos met the needs of employees and administrators by providing centralized management of personal information and services like SSO within the enterprise, or more specifically, within a single domain. As time passed by, growth, splitting and merging of the enterprises increased complexity of the management of personal data and accompanied services. Also, new applications of identity management emerged, and at significantly greater scales. These came from public services where application domains like public healthcare and education posed new challenges to the management of personal data. Thus, the need was created for organized effort in providing best solutions to satisfy these new challenges, including enhanced privacy concerns, security at greater scale, easier and friendlier access for end users, mobility across domains in some cases and more complete isolation of domains in others. As the online activities of general population increase and online services enter our lives in more and more ways, the need to address these challenges to the management of our private data continues to increase. For the time being and in the near future, we can expect that in the changing landscape of our online lives we'll need to continually address these issues.

With that in mind, this paper provides a look at the current state of the field of identity management. The rest of this paper is organized as follows: in section 2 we list some of the traditional challenges met by identity management systems, as already established through previous research by various authors. To demonstrate how those challenges are currently handled, in section 3 we describe a selection of the most prevalent current technologies in use, especially with regard to the mentioned requirements. Section 4 discusses featured technologies and how they solve some of the previously introduced problems, as well as existing research into real-world data and estimates on current number of installations of discussed technologies. Finally, we conclude with a section that discusses expected future challenges.

2 Problem definition

In this section, we briefly describe a selection of the most commonly encountered challenges as described in the existing literature, e.g. in (Bertino et al., 2009), (Torres et al., 2013).

2.1 Modelling identity

To manage identities, firstly they must be expressed in a workable way and therefore precisely and unambiguously defined. However, basic phenomenon of identity is itself not as straightforward to define as it may appear. Usual approaches start by axiomatically asserting the existence of well-defined entities as a basic commonly understood term (alternatively one can also find some definitions of an entity, e.g. in (Cameron, 2005, pg. 7) "a thing with distinct and independent existence"). Within the target context entities are personae, computers, organizations, etc., and entities exhibit characteristics that differentiate

them from each other. Therefore, a common approach like in (Alpár et al, 2011) is to define the identity of an entity as a set of all characteristics attributed to this entity. An important realization is that identity is not absolute, but relative to a scope. For example, within a relatively small organizational unit, two people may be identified by their names alone – most often there are no namesakes within a small enough group. However, within a larger scope like a register of citizens, additional characteristics must be introduced to differentiate various entities (people), like dates and places of birth or ID card numbers. When dealing with computerized systems, these characteristics are more formally encoded as entity attributes – values from well-defined, representable and computable domains (i.e. sets of symbols).

Even if this definition may seem intuitive, some authors argue that there is no merit in observing entities and their characteristics. Cameron (Cameron, 2005) for example argues for a digital subject instead of entity as an identity-carrier, where a digital subject is "a person or thing represented or existing in the digital realm which is being described or dealt with". He makes a point that entities are of interest only if interacted with, and whether entities that aren't in any interaction exist or not is "a moot point". Common traits with the previous approach is that identities of subjects are both context-dependent and defined through their attributes. Additionally, Cameron argues that identities don't have to be unique even within a single context but that in certain scenarios it is even desirable for a single identity to represent more than one subject. The given example concerns groups of subjects being granted the same right to some resource in a foreign domain where from the viewpoint of that domain it's not of importance which subject of the group accesses the resource, therefore all the subjects are representable by the same identity.

2.2 Establishing trust

A simplified definition of trust is willingness of a party to assume a vulnerable position towards another party, expecting the other party to refrain from taking advantage it perceives (more formal definitions are available in (Alpár et al, 2011), (Meyer et al, 1995) and (Jensen et al, 2012)). IdM systems are built with trust assumptions on various levels. The most obvious trust assumption is that the relying party (the party that identity information is delivered to) trusts the identity providing party's (IdP) claims about a particular entity. Aside from this, the user is also assumed to "trust the IdP to make a particular claim about herself to a particular relying party" (Alpár et al, 2011). And then there is additionally the establishing of trust in more complex systems involving more actors and security domains, such as federated identity systems where more than one IdP or relying party interacts in a flow of data. Even harder to achieve is the dynamic establishing of trust, which is necessary in many scenarios where parties from different domains with no prior experience wish to interact. Unfortunately, this need for dynamic establishing of trust contradicts both the common wisdom that trust is earned with experience (which of course always applies) and also contradicts the desire for untraceabillity and unlinkability (described later in this paper) both of significant importance in many scenarios.

2.3 Privacy

In some contexts, privacy is not desirable, but on the contrary all actions performed are supposed to be visible within the context. Examples may include access to shared resources within an enterprise or to vital resources in a military environment. However, in other contexts where entities (usually users) have the right to privacy of their actions, the adoption of IdM solutions may significantly depend upon its ability to ensure user privacy. This is especially true in widely accessible solutions on the Internet whose adoption depends mostly on their reputation with the end-users. There are varying levels of privacy that must be ensured for user data, depending on the nature of the data and the context of use. Examples of the former may include health or financial data and the latter is illustrated by the difference of accessing a trusted personal resource (e.g. an online data storage facility) vs. accessing a site the user perceives as less reputable or trustworthy.

2.4 Linkability and traceability across domains

Two of the special privacy concerns are user data linkability across domains and traceability of user transactions across domains. In many scenarios users of an identity management scheme would prefer their actions not to be linkable or traceable. The most pervasive example is using the IdM system on the Internet for SSO purposes. Users performing purchases on a site that allows session establishment through any of the popular SSO systems (e.g. based on OAuth or SAML) would prefer the site not to be able to link the data it has with data from other similar sites. Additionally, users would prefer the IdP system used for establishing sessions to various such sites not to be able to link their data. The latter is expectedly especially hard to achieve since the IdP necessarily interacts in transactions to various relying parties.

2.5 Security

Security is of paramount importance in an IdM system. This stems from the fact that data kept is personal and misuse of it can lead to serious negative consequences. IdM systems are by design part of more complex environments where they interact with other systems on behalf of users whose identities are managed. That

means that breaches to security of IdM systems may result in identity theft or malicious data manipulation to some party's disadvantage.
By principle of proportionality, an IdM system will be more useful the more data it stores about its managed entities (usually personae) – various scenarios require various data about users, so the more different facts about users are available, the more scenarios the system can service. On the other hand, with more data the system stores about managed entities, it becomes more interesting to malicious parties wishing to subvert it and the consequences of its security breach become potentially worse. The risks rise and in ultimate case the IdM system may become a single point of failure. Unfortunately, to the best of our knowledge, there is no resolving of these contradicting requirements beyond a compromise between usability and risk.

3 State of the art solutions

As mentioned in the introduction, this section gives a brief overview of the most prevalent technologies in current use with regard to requirements explained in the previous section. Selected technologies are those currently in wide real-world use as described in the existing literature, e.g. (Vapen et all, 2016).

3.1 SAML 2.0

SAML is a well-established standard developed by the OASIS consortium for authentication and authorization. Developed during early 2000s, now in the version 2.0, it is an XML-based standard heavily influenced by the enterprise technologies popular at the time, especially the WS-* stack (which is a set of protocols standardized by the W3C that define how XML- and more precisely SOAP-based[1] web services should be implemented and includes publications defining SOAP, WS-Addressing, WS-Policy, WSDL etc.). Similar to some other XML-based standards, it defines a hierarchy of basic concepts – assertions, protocols, bindings and profiles. It was developed mostly for enterprise and business needs and is widely adopted in these domains as well as in public services mainly through Shibboleth, an open-source identity provider that builds upon SAML 2.0. Shibboleth is being developed and maintained by Shibboleth Consortium, whose members include educational, research, healthcare and other public organizations from more than 20 countries (including UK, Swiss, Ireland, Brazil, Czechia, Italy, Austria, etc.)
The most common SAML profile (simplified, profiles are essentially use-cases) is multi-domain web single-sign-on and includes interaction between three parties – a user, an Identity provider, and a Service Provider as illustrated in Fig. 1., taken from (Hughes, Maler, 2005.).

In this scenario, the user already has an established session with one web site, during which the user is redirected to a new web site. The original website acts as an IdP and relays information to the redirected website that acts as a Service provider.

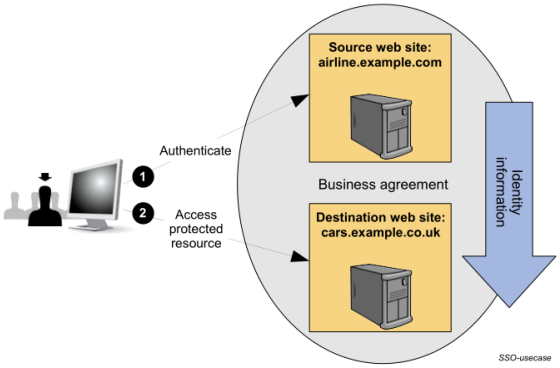

Figure 1. Basic OpenID flow.

The relayed information effectively authenticates the user through a SAML message allowing for a session between the user and the redirected website to be established. In SAML terms, the first website is generally referred to as the asserting party and the second as the relying party.
This use-case depends on various other interactions, including an agreement between parties what information to forward, what transport layer to use, whether to link user representation with both parties or to decouple them, whether to store user data or to rely on transient identifiers that get destroyed after the session is done, etc. SAML 2.0 provides mechanisms to address most of these questions on-the-fly. Additionally, there are many variations of this example flow (e.g. user-initiated instead of AP-initiated as in the example) as well as different flows supported by SAML 2.0, making it a very versatile and robust standard.
However, this also adds complexity which is often viewed as a drawback of the standard prohibiting it from use in some other scenarios, particularly those where some of the security concerns may be somewhat relaxed. Added to that the XML legacy of the protocol is also often viewed as cumbersome and out-of-date in scenarios demanding more lightweight solutions due to processing power and other technical limitations, e.g. smartphone and similar applications.

3.2 OAuth 2.0

OAuth 2.0 standard is published and maintained by IETF through a set of RFC documents, the core consisiting of RFC 6749 (Hardt, 2012 (1)), RFC 6750 (Hardt, 2012 (2)) and RFC 6819 (McGloin, 2013). OAuth 2.0 is foremost an authorization protocol. The user provides permission to a third party (a client in OAuth terms) to access some of user's data at the identity provider as shown in Fig 2., taken from (Hardt,

1

2012 (1)). For example, the user allows a site like fit4life.com read-only access the user's fitness data stored at google.com (an authorization server and resource server in OAuth terms), without necessarily divulging any of the user's other confidential data.

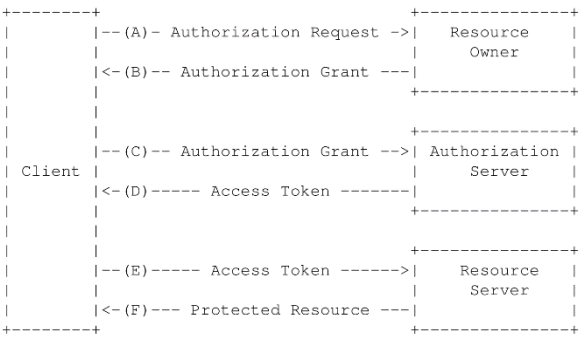

Figure 2. OAuth 2.0 Abstract protocol flow.

However, most of the real-world websites and services use this information to authenticate the user in a good-enough manner. For example, a blogging service or a news portal may use a user's email address or login name at a popular identity provider as a good-enough authenticator for the users posting comments on a blog or a news article, so obviously there's a small step to make between authorization of a client (the blogging service) to a good-enough authorization of a resource owner (the user).

Being a more modern protocol that targets not only web applications and services but also handheld devices, OAuth makes a clear distinction between public and confidential clients based on their ability to authenticate securely with the authorization server. This has important consequences on types of scenarios that are allowed (called grant types) for clients to obtain access tokens. OAuth specification defines four grant types as well as an extension mechanism for additional grant types. These grant types allow OAuth to be versatile in deployment to a heterogenous set of environments and use-cases including traditional web-based service access, as well as use by mobile apps and trusted third parties.

3.3 OpenID Connect

OpenID Connect is an authentication protocol specified and maintained by the OpenID Foundation. The foundation's members are mostly international corporations including well-known names like Google, Microsoft, Oracle, PayPal, Verizon, RSA, VMWare, Deutsche Telekom. The protocol is defined through a set of OpenID specifications including the core specification (Sakimura, 2014) and accompanying specifications detailing additional services (like dynamic provider discovery, dynamic registration with providers, response types, etc.) Unlike the previous versions of OpenID, the current one (final specification launched in 2014) is based on the previously described OAuth 2.0 protocol. The specification extends and additionally specifies parts of the OAuth specification that have been left open by OAuth, so OAuth-based mechanisms could be used in authentication scenarios. For example, one of the most important additions by OpenID is the definition of the ID token – defined are both the contents and the representation of a data structure containing claims made by the authentication server about an end-user's authentication. Some of the other extensions introduced are a list of standard claims about a user, subject identifier types (enabling an IdP to issue different subject value to each client for the same end-user with the purpose of disabling correlation of end-user information by different clients), list of allowed encryptions, etc.

4 Discussion and real-world data

The protocols mentioned in this paper are the most prevalent ones, judging both from the number of resources and size of literature (both academic and non-academic), as well as from usage claims by the groups promoting each of them. This is underlined by the fact that commonly known names in the online world like Facebook, Google and Twitter provide services based on these protocols.

There has been a number of efforts in recent years trying to ascertain the size of various protocols' install base. These efforts are unfortunately not as frequent as the changes observed in the field (primarily on the Internet) so there is to the best of our knowledge no single complete work estimating all the current technologies. A considerable obstacle here is the sensitive nature of the subject and most commercial entities using forms of identity management refrain for security reasons from divulging details of the backbone technologies they use. We speculate this to be the reason a lot of SAML installations aren't visible in existing research datasets.

The most directly observable part of identity management landscape is the Internet-wide public SSO landscape comprised of publicly available services primarily based on OAuth and OpenID/OpenID Connect. For concrete numbers, we'll turn to results of a recent study from 2016. (Vapen et all, 2016) that among several experiments did a web-crawling of a sample set of over 35.000 sites (a pick from top 1M sites by Alexa rank). In this set the study found 1865 relying parties and 50 identity providers. Of the top 10 global identity providers by number of attached relying parties, all were implementing OAuth or OpenID (in favour of OAuth) where the first five providers were facebook.com, twitter.com, qq.com, google.com and yahoo.com. As per the study, facebook.com dominates this list with 1293 found relying parties.

Even though SAML doesn't appear in datasets obtained through collection of publicly visible IdP services and connected relying parties, it's inclusion in this paper is justified by the number of known subjects using solutions based on the protocol. The most notable

implementation of SAML is the Shibboleth software provided by the Shibboleth Consortium (https://shibboleth.net/consortium/).
This implementation is known to be popular in academy and public services – many resources detailing the use of Shibboleth providers are available on the Internet, and the Consortium's website alone lists over 80 known installations mostly in academia and the public sector (and we know first-hand that the list isn't exhaustive since the existing installation at our own university isn't listed).
The brief description of these three standards hints at variations in how challenges introduced earlier in the paper are solved. As stated, the OAuth is primarily an authorization protocol, however, for a lot of websites in the wild the OAuth model suffices for basic authentication purposes – this is because for such websites a single user attribute (most commonly the user's email address or Facebook username) suffices to identify the user, making it a simple model of user identity. OpenID builds upon the OAuth capability by specifying (through the mentioned ID token structure) a set of attributes it considers useful and sufficiently defining of a user's identity. In SAML, the subject data is transferred encoded in XML elements like statements and assertions that are defined in respect to various scenarios (more specifically "profiles" in SAML) and in a loose manner (many elements are optional and of varying content). This allows certain freedoms to specific implementations, effectively letting the parties involved (identity providers and relying parties) to define the necessary attributes for user authentication. This allows for versatility of the protocol, but also introduces the possibility of implementation divergence – something OAuth protocol has been criticized for in the past. OAuth has also been criticized for leaving too many details outside its core specification scope. The most widely known of these critiques comes from one of the contributors to the original specification (Hammer, 2012). The problem of varying implementations and their interoperability observed in the current identity management landscape, was omitted in this paper in favor of more fundamental problems listed earlier, but was documented in earlier works (Jensen, 2012, Cameron, 2005).
Trust is mostly implied – the publicly available services mentioned earlier in the section, especially those from the mentioned top 10 list are trusted by their relying parties and end-users to the extent desired by their typical uses. We can reasonably argue that users will trust a Facebook or Google implementation in typical online scenarios like signing into a social application or picture sharing service, but would probably weigh again the risks involved in letting these public services access their health or financial records. The use of a solution based on such a service in an environment requiring highest trust (e.g. e-voting) is and probably will stay in foreseeable future only a hypothetical possibility.
Similarly, the privacy these services provide is at least questionable. Privacy was a goal in development of these standards - SAML's provision for pseudonym usage was mentioned, as well as OAuth's (and by implication OpenID's) end-user's conformance to sharing data to third parties. However, users often have no way of assuring the services based on them hold up to high privacy standards. Users are informed in variously specific detail about data they are sharing (e.g. "OAuth 2.0 Scopes for Google APIs" or "Permissions Reference – Facebook Login") and often have no real control in behind-the-scenes sharing of information. Experience shows that most users will implicitly accept some privacy disruption. E.g. the willingness to be traceable on the web is evidenced by wide success of online tracking technologies used in advertisement purpose, so even if aware of their traceability by IdPs and the linkability of their transactions across the Internet, the users still put enough trust in available IdPs to continue using them.
Similarly to trust, the security of the solutions is implicitly accepted by end-users. Here the earlier mentioned rule about trust and experience applies especially, as does the tradeoff (from user's perspective) between the security-related expectations and willingness to use the services to their full potential. A typical user will usually refrain from providing overly sensitive data to an online service even if it means abstaining from some online services requiring such data – it is acceptable to leave one's email address with an online IdP provider, but not necessarily a real-world address, no matter how many form autofills it could potentially provide with online shops.

5 Future research

Currently research regarding identity management is conducted in various directions. Projects like LIGHTTest, Aries, CREDENTIAL and many others are examining the suitability of blockchain technology to manage identities, the possibility of using the existing DNS infrastructure to support identity management as well as the possibility of implementing cloud-based yet privacy friendly identity providing services. These projects are still in early development so there is no definite literature available on the topic, but they have already garnered enough attention to receive institutional support (e.g. through EU's Horizon 2020 Programme).
In light of this interest in the topic of identity management and the open problems described in previous sections, we propose another, to our knowledge not yet examined method. The traditional model of attribute-based identity as described in section 2.1 deals with entity's characteristics (attributes) and considers them constant for each entity they are attributed to. We believe that by extending these constant attributes and viewing them as time- and

geolocation-dependent values (functions of time and location) we can create a model of identity that allows for a finer-grained control of information. Potentially this would allow for an easier and more natural "sharding" of identity information where different security domains would be able to keep and be responsible for information naturally belonging to those domains. A simple example would be a person's address after moving from a country to another one – each county assumedly possesses a system to store such data about their residents, so each of those systems is responsible for a person's address at various time intervals (pre- and post-move). This approach introduces various complexities regarding the data consistency, more complex security concerns, data federation concerns etc., but we believe those to be worth researching for the expected gains in overall accuracy of the stored data as well as the flexibility of systems based on such model.

6 Conclusion

As discussed in the previous sections, the technologies in current usage address the challenges identified and presented in previous academic body of work. Those challenges are, however, met in a limited manner. No current system can guarantee complete fulfillment of all requirements posed by end-users or other parties involved in all scenarios involving IdM systems. In our view, the reasons for this are twofold.

On the one hand, various stakeholders in the field have various expectations from identity management schemes. This is primarily observable with the privacy concerns. For example, we believe that to be the reason why the discussed technologies don't provide facilities to anonymize user transactions from the IdP system itself i.e. don't completely satisfy the user wish for untraceability, even though there exist known solutions that provide higher levels of privacy in this regard, e.g. (Camenisch, 2002).

On the other hand, even from the view of a single party, e.g. the end-user, some challenges seem to be in contradiction, as mentioned in sections on trust and security. These contradictions have, however, to our knowledge not been proven, and viewed from the current stand there may exist a solution that satisfies all the apparently contradicting demands.

Therefore, we conclude that further research will be required as long as the apparently contradicting expectations and requirements aren't either satisfied through a more complete solution or such a solution is definitely proven to be unattainable.

References

Alpár, G., Hoepman, J. H., & Siljee, J. (2011). The identity crisis. security, privacy and usability issues in identity management. arXiv preprint arXiv:1101.0427.

Bertino, E., Paci, F., Ferrini, R., & Shang, N. (2009). Privacy-preserving digital identity management for cloud computing. IEEE Data Eng. Bull., 32(1), 21-27.

Camenisch, J., & Van Herreweghen, E. (2002, November). Design and implementation of the idemix anonymous credential system. In Proceedings of the 9th ACM conference on Computer and communications security (pp. 21-30). ACM.

Cameron, K. (2005). The laws of identity. Microsoft Corp.

Dhamija, R., & Dusseault, L. (2008). The seven flaws of identity management: Usability and security challenges. IEEE Security & Privacy, 6(2).

Hammer, E. (2012). OAuth and the Road to Hell. Retrieved from https://hueniverse.com/oauth-2-0-and-the-road-to-hell-8eec45921529

Hardt, D. The OAuth 2.0 Authorization Framework. Request For Comments–RFC 6749, 2012.

Hardt, D., & Jones, M. The OAuth 2.0 Authorization Framework: Bearer Token Usage. Request For Comments–RFC 6750, 2012.

Hughes, J., & Maler, E. (2005). Security assertion markup language (saml) v2. 0 technical overview. OASIS SSTC Working Draft sstc-saml-tech-overview-2.0-draft-08, 29-38.

Hughes, J., & Maler, E. (2005). Security assertion markup language (saml) v2. 0 technical overview. OASIS SSTC Working Draft sstc-saml-tech-overview-2.0-draft-08, 29-38.

Jensen, J. (2012, August). Federated identity management challenges. In Availability, Reliability and Security (ARES), 2012 Seventh International Conference on (pp. 230-235). IEEE.

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. Academy of management review, 20(3), 709-734.

McGloin, M., & Hunt, P. (2013). OAuth 2.0 Threat Model and Security Considerations. Internet Engineering Task Force (IETF) RFC, 6819.

OAuth 2.0 Scopes for Google APIs. Retrieved from https://developers.google.com/identity/protocols/googlescopes

Permissions Reference – Facebook Login". Retrieved from https://developers.facebook.com/docs /facebook-login/permissions

Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., & Mortimore, C. (2014). Openid connect core 1.0. The OpenID Foundation, S3.

Torres, J., Nogueira, M., & Pujolle, G. (2013). A Survey on Identity Management for the Future Network. IEEE Communications Surveys and Tutorials, 15(2), 787-802.

Vapen, A., Carlsson, N., Mahanti, A., & Shahmehri, N. (2016). A look at the third-party identity management landscape. IEEE Internet Computing, 20(2), 18-25.