

Comparing Identity Management Frameworks in a Business Context

Jaap-Henk Hoepman, Rieks Joosten, and Johanneke Siljee

jaap-henk.hoepman@tno.nl, rieks.joosten@tno.nl, johanneke.siljee@tno.nl
TNO, the Netherlands

Abstract. Several frameworks for identity management exist, each of them with its own distinguishing features. These frameworks are complex however, and their features not easily understood. This makes it hard for businesses to understand the intricacies, and difficult to select and deploy one. This paper develops business selection criteria and applies them to four popular identity management frameworks. The resulting score card (1) helps businesses to select and deploy an identity management system, and (2) provides valuable feedback to developers of identity management systems on the criteria that they should take into account when designing and implementing an identity management system that is useful for specific businesses.

1 Introduction

Businesses that provide a meaningful IT service require that only users with proper privileges, e.g. because of a subscription, can access this service. To check these privileges, the application providing the service must establish and verify the user's identity. Traditionally, applications handle this by themselves, meaning that many user registrations exist, each with its own ways of user authentication. While this is not user-friendly (users need to remember many passwords for example), it is also not efficient for the business as they cannot tell whether the same customer uses multiple services (which makes him a more interesting customer). Similar considerations apply when considering users that are in fact employees of a business, who need access to different sets of documents and business applications.

Identity management systems separate the act of identifying and authenticating the user from the act of providing a service to a user. This is attractive for large enterprises as it bears the promise of easier, more centralized management of users and their access rights. For users it promises to provide a uniform service access experience, without the need to enter usernames and passwords again and again.

Apart from the data needed to identify and authenticate users, services store additional information about their users in so-called user profiles. Most of that data is the same for many different services. By delegating (some of) the administration and storage of that data to the identity management system, that data is more easily kept up to date, and does not have to be entered by the user for every new service that he accesses.

Several frameworks for identity management exist: OpenID [13], Shibboleth [16], Liberty [10] and the Identity Metasystem [8], [9] (also referred to as CardSpace), to name but a few. While each of these systems has its own distinguishing features, at a high abstraction level they have several things in common, as shown in Figure 1. This figure shows that each framework incorporates:

- a technical component called a user agent¹ (UA), e.g. a web browser, that is operated by a person that wishes to access a service,
- a technical component called identity provider (IdP²), for instance a computer application or web service, that identifies and authenticates the person (user) that operates the UA and provides identity data, and
- a technical component called service provider (SP), again a computer application or web service, that offers the service the aforementioned person is interested in. As SPs rely on IdPs for user authentication, SPs are also called Relying Parties (RPs).

To accommodate communication between these components, identity frameworks (a) use common 'languages' (e.g. XML, SAML) for exchanging messages, (b) use common protocols (such as HTTP, SOAP and others) for exchanging messages between two individual components and (c) define the protocol(s) that govern the sequence in which components talk to one another and the types of data exchanged.




Fig. 1. Typical Identity Management architecture

As technical components cannot be held accountable, we introduce the notion of 'domain' to represent a legal entity (a business or individual person), that is responsible (and accountable) for the activities thereof. As bearing responsibility is associated with risk, businesses manage this by defining measures and policies for a domain. Identity systems in a domain must then implement such measures and follow the policies (for identity related risks). For example, a business in the Netherlands may trust banks and the Dutch government to provide identity data, but it may not trust telecom operators or a foreign government to do the same. It may state that data be digitally signed according to some Digital Signature Act, etc. Its identity

¹ Individual identity management systems may have slightly different terminology, e.g. 'user' for 'user agent'.

² We write IdP instead of the also used abbreviation IP, which is already used for Internet Protocol

management system must ensure this. As the risks that businesses face can be quite diverse, policies will differ from business to business, and identity management frameworks are challenged to provide a good match for that.

The fact that the merits and drawbacks of identity systems are to be judged by technical as well as business criteria, makes it all complex, hard to oversee, and difficult to make decisions about. In this paper we discuss how identity management can be applied in business contexts, thus giving a helping hand to future decision makers seeking to deploy one or more identity management components in their businesses.

Our contribution in this paper is the following. We describe the business context in which identity management systems need to operate and discuss the main business concerns that originate from that. These concerns are translated into business requirements against which we score the aforementioned four popular identity management frameworks. Running an identity management platform raises its own issues. We also discuss these operational requirements and score the four frameworks against those. To complete the picture, we also score the same four frameworks against the widely accepted 7 Laws of Identity [4], that are mostly user-centered (adding an 8th Law of Location Independence, as the final necessary user-centred requirement that was lacking in the former seven laws). This extends the work of Maler and Reed [12], and complements the comparison of identity management systems on the associated costs and organisational issues of Royer [15]. Our results are a useful tool helping organisations seeking to deploy an identity management system to choose the system that best suits their needs. They are also useful input for developers of next generation identity management systems that wish to improve current systems and broaden the range of application of their systems.

2 Identity Business

Traditionally (in IT), Identity is the answer to questions such as: 'Who is this customer?' or 'Who is this supplier?', and the answer was a name. Currently, Identity includes all information a business³ may need in dealing with its customers, suppliers etc. For example, if a business needs to send letters to an entity, then name and address will be part of its Identity. Note that as the business continuously improves its processes, its need for information changes over time, Identities change as well. For example, when email became available, Identities came to include one or more email addresses. Thus:

The Identity of a person or organisation, from the perspective of a given business, consists of all data (information) that this business needs or has at that particular point in time for dealing with that specific person or organisation⁴.

³ In this article, governmental organizations are also considered to run a business, with individual people as well as organizations playing the roles of customers, suppliers, etc.

⁴ A person or organization thus has as many Identities as there are businesses that have information about them.

This is not to say that businesses can gather, use, or provide identity data to others as they like. Laws and regulations, such as various EU Directives and domestic legislation that constrain the processing of personal data and the (free) movement thereof, must be complied with. Additional constraints may originate from e.g. supplier contracts that may impose restrictions with respect to the purpose for which the data may be used.

Businesses (and individuals alike) should have comprehensive policies for gathering, using and providing data. Such policies may state which individuals or organizations are trusted to obtain identity data from, or to provide such data to, or for what purposes certain identity data may be used. There may also be rules that govern the trustworthiness (integrity) of personal information, e.g. an email address can be decided to be trustworthy only after a response has been received to a message sent thereto.

Also, businesses and individuals may have concerns with respect to the possible consequences of correlating identity data over time. An individual may not want a web-shop to know what it has bought in earlier sessions, or he may not want the government to supply their address information to arbitrary businesses. If any organisation could freely collect identity data from other businesses, and aggregate and sell it to whoever pays for it, then this could for example facilitate identity theft. However, if identity data cannot be passed along, then people and businesses need to fill in the same information over and over again.

The identity business thus consists of specifying business objectives and policies regarding the processing of identity data and the exchange thereof, as well as managing them and realizing/enforcing them. Identity systems should accommodate not only for differences in identity data (types) and the way they are exchanged, but also for the management and realization of business goals and policies.

To be able to assess whether identity systems truly accommodate business goals and policies, we developed a set of business requirements. These requirements and the assessment of a number of currently popular identity systems can be found in section 3.

Note however that while business policies have impact on how identity management systems should operate, the converse is true as well: capabilities of identity systems may inhibit or enable businesses. An example of inhibition is identity systems that are susceptible to phishing attacks should not be used for commercial services as attackers could then use that service using someone else's account. An example of business enablement is given by identity systems that guarantee that identity data is only released to an SP with the user's consent, so that a business can act as an IdP for all identity related data that it has. An even further reaching idea is that of Identity Oracles; in which IdPs provide higher level information derived from personal data, as in "this person is at least 21 years old" [3].

3 Comparing Four Identity Management Frameworks

In this section we provide a set of requirements for identity systems that are useful for an organisation to assess which identity system to deploy.

3.1 Approach

By looking at current identity management systems and related work we derived a set of requirements. Part of the requirements are the widely accepted [1] 7 Laws of Identity [4], which is a set guidelines, aiming explaining the successes and failures of identity management systems from a user-centric perspective. In our opinion one important law is missing from this set, namely the requirement that a user should be able to access a SP using an identity management system not only from his PC, but also from a computer at a cybercafé in Hong Kong, for example. We call this the 8th Law of Location Independence [17]. This essentially means that the identity management system should not rely on any persistent data stored locally at the user's machine.

Dhamija and Dusseault [5] raise seven flaws of current identity management systems that need to be resolved before identity management systems will be adopted. Although these flaws can be translated into requirements as well, it is not useful to include them in our comparison as none of the current identity management systems fulfil them. An approach a bit similar to ours is presented in [12], where three popular federated identity protocols are profiled: the Security Assertion Markup Language (SAML), the OpenID specification, and the InfoCard specification underlying Microsoft's Windows CardSpace.

Furthermore, we add a set of requirements addressing business concerns, e.g. dealing with operationalisation of such systems, policy management, privacy concerns and known vulnerabilities. These requirements are derived from the discussion in the second section on Identity Business.

The total set of requirements, presented in the next section, is used to compare the four currently popular user-centric identity systems: OpenID 2.0, Shibboleth and Liberty (both based on SAML), and CardSpace 1.0.

3.2 Identity Management System Requirements

The first set of requirements are user-related, the first of which are the 7 Laws of Identity for which an underpinning is given in [4]:

1. **User Control and Consent** (LI1): The solution only reveals identity data with the user's consent.
2. **Minimal Disclosure for a Constrained Use** (LI2): The solution discloses no more than the necessary identifying information.
3. **Justifiable Parties** (LI3): The design ensures that disclosure of identifying information is limited to parties that have a necessary and justifiable place in a given identity relationship).
4. **Directed Identity** (LI4): The solution supports both "omnidirectional" identifiers for use by public entities and "unidirectional" identifiers for use by private entities, thus facilitating discovery while preventing unnecessary release of correlation handles.

5. **Pluralism of Operators and Technologies (LI5):** The solution channels and enables the interworking of multiple identity technologies run by multiple identity providers.
6. **Human Integration (LI6):** The solution defines the human user to be a component of the distributed system, integrated through unambiguous human-machine communications mechanisms offering protection against identity attacks.
7. **Consistent User Experience across Contexts (LI7):** The solution provides a simple consistent experience while enabling separation of contexts through multiple operators and technologies.

However, another requirement dealing with mobility is crucial for widespread acceptance of identity management:

8. **Location Independence (LI8):** the solution does not restrict users in access to their identity system to one location, such as one personal computer that holds specific data.

Then, from the discussion in section 2, we derive business-centred requirements for the SP and IdP. First, there are technical requirements that allow for technical implementation and usage:

9. **Use of standards:** The solution makes use of existing, well known and broadly used standards.
10. **Openness:** The solution itself should be freely usable, i.e. no patent fees or licenses required.
11. **Availability of (open) components:** The solution should consist of existing components that are usable in a wide variety of environments (Windows, UNIX, Linux, MAC, etc.) and preferably have an open source implementation for better evaluation of the correctness and security.
12. **Technical Interoperability:** the solution can interoperate (technically) with the other solutions.

Then, there are also operational requirements that relate to the business of running an SP and/or IdP:

13. **Pseudonymous and anonymous use:** the solution should provide means for users to use pseudonyms for identification, and/or remain completely anonymous towards SPs. This allows the system to be used in a more diverse set of usage scenarios (improving the business case by including the users that want or need to be anonymous) and potentially limits liability issues.
14. **Attribute semantics:** the solution should guarantee and/or provide means to unambiguously define the semantics of identity attributes.
15. **Validity and up-to-dateness:** the solution provides guarantees with respect to the validity of identity data, and the up-to-dateness thereof.
16. **Ease of local policy management:** the solution provides means to easily configure identity policies (i.e. without having to recompile code or create

(virtual) connections/adapters), in the event of regulatory changes, changes in business relationships, security incidents and so on.

17. **Business Case:** the solution should provide every party (domain) with a valid business case.
18. **Governance support:** the solution provides suitable means by which to achieve demonstrable compliance with (identity) legislation, policies.

3.3 Comparison

Currently, there are four popular user-centric identity systems: OpenID 2.0, Shibboleth and Liberty (both based on SAML), and CardSpace 1.0. We compare these four identity systems against the requirements.

In Table 1 the fulfilment of each identity system with respect all requirements is given. The scores in the table have the following meaning:

- ++ full support/compliance
- + reasonable support/compliance, but not to the full extent
- +/- support/compliance is subject to debate
- some support/compliance, but only very little
- no support/compliance

Table 1. Requirement fulfilment of OpenID, Shibboleth, Liberty, and Identity Metasystem

	OpenID 2.0	Shibboleth	Liberty	CardSpace (Identity Metasystem)
1. User control and consent	Users control which identity provider (IdP) they trust, what attributes the IdP may store, and which relying parties (RP) access which IdP in each session. Users do not control the actual identity data transferred in a session. Susceptible to phishing attacks, which is a violation of the first law	See Liberty	Federation of identity attributes without user consent is default after user gives permission for the first time. Also susceptible to phishing attacks, which violates this law.	Users control which identity provider (IP) they trust, what attributes the IP may store, and which relying parties (RP) access which IP in each session. Also, users can see which identity data is transferred in a session and decide not to, if necessary.
	+	-	-	+

4. Directed identity	3. Justifiable parties	2. Minimal disclosure	
<p>+ OpenID supports the use of multiple identities (pseudonyms) and leaves it up to the user to what end each identifier is used.</p>	<p>+/- OpenID is designed as a fully decentralized mechanism. Since RP and IdP (optionally) create an association, they know that the user is associated with the other and which information is shared. OpenID defines 'realms', i.e. RP-servers that may share the identity information. The IdP should inform the user what the realm is (which is outside the scope of OpenID).</p>	<p>+ OpenID Attribute Exchange 1.0 supports attribute exchange; however, minimal disclosure is not enforced, but decided by the RP.</p>	<p>OpenID 2.0</p>
<p>+ Both are supported (see Liberty)</p>	<p>+/- No mechanisms in place to enforce this; on the other hand, unknown parties are not part of the federation and will not obtain any useful information</p>	<p>+ Use of temporary handles ensures that users cannot be traced; SP can retrieve identity from the IdP later if allowed</p>	<p>Shibboleth</p>
<p>+ Pseudonyms are Liberty's unidirectional handles; cookies can be used as omnidirectional identifiers.</p>	<p>+/- The circle of trust is a logical context for the business partners involved; untrusted parties cannot be part of it. Other than that IdP and SP decide whether a new relationship is necessary and justifiable.</p>	<p>+ Uses pseudonyms and thereby limits disclosure of identity information</p>	<p>Liberty</p>
<p>+ Uses Private Personal IDs to create pseudonyms that are unidirectional for each individual RP. Claims may also contain omnidirectional handles.</p>	<p>++ The Identity Metasystem is designed such that RPs can be visually recognized by users (as opposed e.g. to having to interpret certificate attributes). The same holds for recognizing IPs. Also, there are some provisions that assist users to recognize RPs that it has used before.</p>	<p>+ While the user controls the data to be transferred, it does not control the actually provided information per se, i.e. more (detailed) information than strictly necessary could be asked for and provided. 'Personal identifiers' can be used that prevent RPs to link identity data to similar data of other RPs.</p>	<p>CardSpace (Identity Metasystem)</p>

7. Consistent user experience	6. Human integration	5. Pluralism of operators, techniques	
+/- User experience is consistent, but not simple for average users, as they need to check e.g. IdP certifications	-/- OpenID does not provide protection against phishing attacks or any other common security problems that stem from the human-machine communication. Also, users are required to use a URL or XRI as identifier, of which only part can be chosen by the user himself.	+ OpenID is an open, extensible standard, based on HTTP, and allows for multiple IdPs that each can provide attributes, but each needs to be "OpenID-enabled".	OpenID 2.0
+/- User experience is consistent. Separation of contexts unclear	- Not part of the specification	+ Uses SAML, but all parties must conform to the standard	Shibboleth
+/- User experience is consistent. Separation of contexts is unclear.	- No protection against phishing attacks or any other common security problems that stem from the human-machine communication.	+ Is open source, integrates SAML, and allows for multiple (Liberty-enabled) attribute providers.	Liberty
+ The user experience is consistent and user-friendly. Separation of contexts is supported. Interoperability has been demonstrated ⁵ .	+ Specifies interactions tailored for humans (using visual clues - see also Law 3). The Identity Selector runs in a separate desktop to resist spyware/malware attacks. It also resists phishing attacks. However, Identity Metasystem does not provide facilities that allow users to use their identities at other locations.	+ Is an open design, supporting Kerberos tickets, SAML tokens, X509 etc., and is extendable. Each needs to be "Identity Metasystem" enabled.	CardSpace (Identity Metasystem)

11. Availability of (open) components	10. Openness	9. Use of standards	8. Location independence	
+ Various implementations can be downloaded from the Internet, for Windows as well as Linux	++ The standards are open, and various implementations include sources, tutorials, help, etc.	++ HTTP, HTML, SHA, HMAC authentication, DH-key agreement, URL, XRI, Yadis, and more.	++ Users can enter their identifiers any place, any time.	OpenID 2.0
+ Open source, platform independent	++ Open standards, open source implementations available	+ SAML, HTTP	++ See Liberty	Shibboleth
+ Multiple open source projects exist (e.g. Open Liberty, Lasso, Cahill)	++ Open standards, open source implementations available	++ SAML 2.0, SOAP, HTTP	++ Interaction with SPs and IdPs are not dependent on the user's location.	Liberty
+ Example code (C#, PHP) downloadable from Microsoft. See also the Bandit project, Higgins, etc.	++ Open source (see Microsoft's Open Specification Promise (OSP) http://www.microsoft.com/interop/osp/default.mspx).	++ WS-* (Security, MEX, Policy, Trust, etc.), SAML, Kerberos, X509v3, etc.	- Currently, users can only export their InfoCards as XML files, transfer them to another device and import them there in the local Identity Selector.	CardSpace (Identity Metasystem)

15. Validity and currency	14. Attribute semantics	13. Pseudonymous and anonymous use	12. Technical interoperability	
+	+/-	+	Technical interoperability is to be evaluated on a product comparison basis, which is too detailed for this article. However, significant effort is being made to test implementations of frameworks against one another, for example the OSIS Interops [14]	OpenID 2.0
+	+	+	Pseudonyms: yes ('user-supplied identifier'); no protection against colluding SPs. Anonymous use: no	Shibboleth
+	+	+	See: Liberty	Liberty
+	+	+	Through OpenID Service Extension Attribute types are managed centrally. Semantics not always unambiguous.	CardSpace (Identity Metasystem)

18. Governance support	17. Business Case	16. Ease of local policy management	
<p>-</p> <p>While UA can be authenticated, the User cannot (due to phishing vulnerability). SP can change identity data. Any support must be build around OpenID</p>	<p>+/-</p> <p>OpenID is primarily useful for simple, non-security critical, IdM applications.</p>	<p>-</p> <p>out of scope for OpenID</p>	<p>OpenID 2.0</p>
<p>+/-</p> <p>Not enforced by standard</p>	<p>+</p> <p>See: Liberty; but easier management.</p>	<p>+</p> <p>SPs and IdP can update the Circle of Trust using the corresponding metadata.xml file.</p>	<p>Shibboleth</p>
<p>+/-</p> <p>Not enforced yet, work in progress through Id Governance Framework [7]</p>	<p>+/-</p> <p>Primarily useful for federating IdPs and SPs of already federated domains.</p>	<p>+/-</p> <p>Depends on Liberty-enabled implementation</p>	<p>Liberty</p>
<p>+</p> <p>Data used in CardSpace carries signatures of data providers and provides proof of consent of user</p>	<p>+</p> <p>SP can ask as much or as little as it likes; IdP can provide all data it has User is in control Suitable for a wide range of usage scenarios.</p>	<p>++</p> <p>SPs control the attribute types they request and the IdPs they trust to provide them. IdPs control whether or not identity data is to be audited, and if so, which SPs are entitled. Users see all requested and provided data pass, and may abort the protocol at any point.</p>	<p>CardSpace (Identity Metasystem)</p>

4 Conclusions and Recommendations

In this paper we have investigated the requirements on an identity management system from three different perspectives: User, Technical, and Business. We have formulated a set of important requirements from each of these perspectives, and have scored four existing, popular identity management systems against these requirements. The results show that each have their advantages and shortcomings, which can be summarised as follows:

- OpenID is highly location independent, and gives the user a lot of control, but scores badly with respect to the more business-oriented requirements.
- Shibboleth and Liberty are very similar, technologically wise. Within the limits of a browser-only (and hence location independent) IdM framework, they achieve a good overall score on most of the requirements.
- OpenID, Shibboleth and Liberty are susceptible to phishing and similar attacks. This is a common drawback of browser-only IdM frameworks.
- CardSpace fulfils many of the listed requirements. Currently, its major drawback is the fact that it is not location independent because Infocards are locally stored on the PC. This is a drawback of all IdM systems that rely on extra software beyond the browser.

For businesses seeking to deploy an identity management solution, we recommend that they first select the requirements most important to their business, and use the scorecard to select the solution that scores best on those requirements. This helps businesses taking balanced decisions.

References

1. Abhilasha Bhargav-Spantzel, Jan Camenisch, Thomas Gross, and Dieter Sommer, User Centricity: A Taxonomy and Open Issues, Journal of Computer Security, Volume 15, Number 5, 2007.
2. Bob Blakley, Identity and Community in Human Society, Catalyst Conference 2006, 15 June 2006 http://podcast.burtongroup.com/ip//2006/06/identity_and_co.html
3. Bob Blakley: Ceci n'est pas un Bob, 7 December 2006 <http://notabob.blogspot.com/2006/07/meta-identity-system.html>
4. Kim Cameron, The Laws of Identity, 21 May 2005, <http://www.identityblog.com/stories/2005/05/13/TheLawsOfIdentity.pdf>
5. R. Dhamija and L. Dusseault: The Seven Flaws of Identity Management: Usability and Security Challenges, IEEE Security & Privacy, March/April 2008 (Vol. 6, No. 2), pp. 24-29.
6. Information Card Foundation, <http://www.informationcard.net/>
7. Liberty Alliance Project, An Overview of the Id Governance Framework, Version: 1.0, 2007.
8. Thomas Daemen, Ira Rubinstein (eds): The Identity Metasystem: Towards a Privacy-Compliant Solution to the Challenges of Digital Identity, Microsoft white paper, October 2006 http://identityblog.com/wp-content/resources/Identity_Metasystem_EU_Privacy.pdf

9. Kim Cameron and Michael B. Jones: Design Rationale behind the Identity Metasystem Architecture. http://www.identityblog.com/wp-content/resources/design_rationale.pdf
10. Liberty Alliance Project, <http://www.projectliberty.org>
11. Susan Landau and Jeff Hodges: A Brief Introduction to Liberty, 13 February 2003, http://research.sun.com/liberty_intro/
12. E. Maler and D. Reed: The Venn of Identity: Options and Issues in Federated Identity Management, IEEE Security & Privacy, March/April 2008 (Vol. 6, No. 2), pp 16-23.
13. OpenID specifications, <http://openid.net/developers/specs/>
14. Open Source Identity Systems, <http://osis.idcommons.net/>
15. D. Royer: Assessing the Value of Enterprise Identity Management (EIdM) - Towards a Generic Evaluation Approach, In: Proc. 3rd Int. Conf. on Availability, Reliability and Security (ARES 2008), Barcelona, Spain, 2008, pp 779-786.
16. The Shibboleth project, <http://shibboleth.internet2.edu/>
17. Johanneke Siljee and Jaap-Henk Hoepman: Issues in Identity Management, Usability, Security and Privacy, TNO Whitepaper, 2008 (to appear).