

Economic Tussles in Federated Identity Management

Susan Landau

Radcliffe Institute for Advanced Study

Harvard University

susan.landau@privacyink.org

Tyler Moore

Center for Research on Computation & Society

Harvard University

tmoore@seas.harvard.edu

Abstract

Federated identity management (FIM) enables a user to authenticate once and access privileged information across many disparate domains. It is a technology of great promise whose adoption has been disappointing. FIM's proponents include both governments and leaders in the IT industry. Many explanations have been given for its slow uptake, from disputes over liability assignment for authentication failures to concerns over privacy. We present an economic perspective on stakeholder incentives that can help shed light on why some applications have embraced FIM while others have struggled. By presenting seven use cases of successful and unsuccessful FIM deployments, we identify four critical tussles that may arise between stakeholders when engineering a FIM system. We show how the successful deployments have resolved the tussles, whereas the unsuccessful deployments have not. We conclude by drawing insights on the prospects of future FIM deployments.

1 Introduction

Federated identity management (FIM) provides a way to share user authentication information across a variety of domains.¹ Such systems allow a user to authenticate once – single sign-on (SSO) – and then use that identity to access information across multiple security domains. It is

¹Portions of the introduction appeared previously in [22].

a potentially powerful technology. Data sharing across domains creates efficiencies and can even provide increased privacy for the user (e.g., by authenticating an individual in a new domain as a member of a group authorized for access rather than the individual per se). But federated identity management blurs security boundaries and thus creates liability and privacy risks.

Around 2001, industry began developing federated identity systems for “single sign-on” online identity management; this included OASIS (Organization for the Advancement of Structured Information Standards), an international standards organization that was developing an XML framework for authentication and authorization [31], Microsoft, which was developing the Passport system, and the Liberty Alliance, which was developing set of open specifications for online single-sign on and identity federation. The Liberty work was folded into the OASIS effort, which coalesced around the Security Access Markup Language (SAML).

These early systems had problems. Passport system centralized all the data, which creates privacy and security risks, and was eventually abandoned; meanwhile the inter-industry Liberty Alliance effort was designed to satisfy the needs of the enterprise environment. Broader success was elusive. The problem of simple, easy, secure, privacy-preserving online authentication for everyday use remained unresolved.

As the Internet changed, the needs for user authentication shifted. On the one side, with blogging and its associated commenting, sites sought a lightweight identity system in order to exercise a modicum of control over commenters and avoid spam. OpenID filled this need. Frequently based on email addresses, the OpenID mechanisms were easy to use. But these electronic mechanisms came at a cost of being less than secure. At the same time, with high-level cyber exploitations of U.S. industry and government sites increasingly occurring, the need for robust online authentication had increased. OpenID did not fit this bill.

The US government has long perceived federated identity to be a crucial component of e-government [10] and in 2002 developed a federated PKI bridge for cross-department authentication [40]. But e-government solutions were slow in emerging. There were a host of explanations for this. Outstanding technical issues included the need for assurance levels for authentication, standards for the types of authentication technologies, and agreement on policies for issuing, retaining, and revoking authentication [40, p. 15]. There was a lack of interoperability between the private-sector systems [40, p. 19]. Furthermore, liability was seen to be a big stumbling block. Thus when Wells Fargo finally agreed to serve as an identity credentialer, it was only under a liability “holds harmless” arrangement for any resulting damages [15].

The lack of policies for data privacy created a real flash point [12]. The UK’s experience with identity cards demonstrates that if private data is not protected, the system may not be adopted; public concerns over the retention of transactional information into large databases [5] helped derail the UK government’s planned identity card scheme [26]. In the US, privacy issues are now

being explicitly addressed in the hope that they might be removed from the center of controversy. In late 2010, the US government released the “Privacy Green Paper,” [20], which emphasized privacy protections (“Privacy protections are crucial for maintaining the consumer trust that nurtures Internet growth” [20, p. iii]). Furthermore, the recently released National Strategy for Trusted Identities in Cyberspace strongly emphasizes the need to incorporate robust privacy protections into system design [33].

Meanwhile, some federated identity management systems have experienced modest success. Examples include Shibboleth in the higher education sector, SAML in the enterprise sector, and the National Institutes of Health federated identity management program. But federated identity management has not caught on in the broader Internet. In particular, federated identity management has functioned well in sectors in which the parties had first established contracts, but on the “open” Internet, where the Identity Providers (IdPs) and Service Providers (SPs) might not previously have had a relationship, federated identity management has experienced slow adoption. It is widely believed that the inability to solve the liability issue – who would bear the costs when federated systems inappropriately shared information or incorrectly authenticated a user – is at the root of the issue [12, 29, p. 22]. We believe the root of the problem lies in a more complex tangle of economic issues, of which liability is one piece.

The design of federated identity management systems is a classic case of economic tussle. When the systems have been successful, it has been because benefits accrue to both sides; when such systems have so far failed to achieve traction has been when the systems are weighted so that the benefits largely accrue to only one side. Rather than liability alone, the problem is actually one of maladjustment to the economic tussle. Consequently, if one can readjust the values in those systems so as to provide clear – and relatively balanced – benefits to all parties, then the federated system is much more likely to succeed.

We begin in Section 2 with a brief description of the players in a federated system. We continue in Section 3 with an examination of a number of use cases in federated identity management, both failures and successes. In Section 4 we present empirical evidence that lends support to our claim that a key reason why Facebook’s web authentication mechanism has attracted more adoption than OpenID is that Facebook shares more extensive user information than OpenID identity providers. In Section 5 we consider the tussles that may arise, while in Section 6 we re-examine the use cases with those tussles in mind. We close with a discussion of the insights made possible by examining federated identity management through the lens of economic tussle.

2 Players in federated identity management systems

Federated identity management systems have four main components:

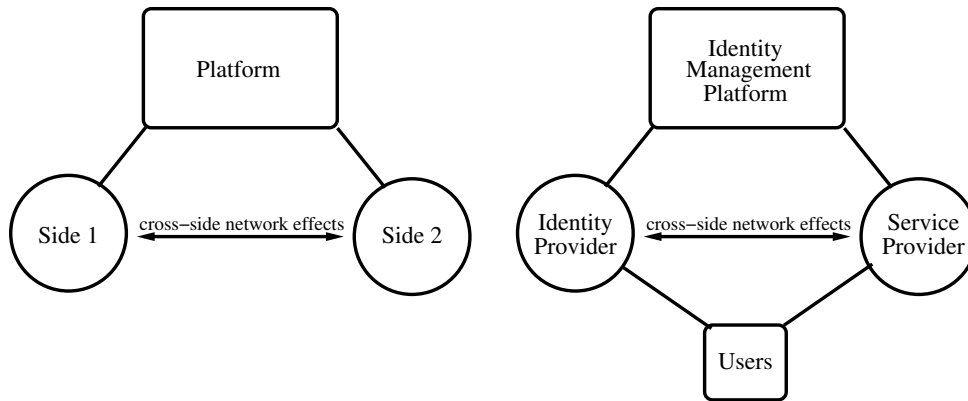


Figure 1: General diagram of a two-sided market (left), and a diagram of federated identity management as a two-sided market (right).

- The user (or user agent such as a browser) who has a particular digital identity that interacts with a network application;
- The Identity Provider (IdP), whose role is to authenticate the user and that may store attributes about the the user;
- The Service Provider (SP), an application which provides services to the user and which relies on the Identity Provider to perform user authentication.
- The Identity Management Platform, a framework or set of rules defining how IdPs, SPs and users interact.

The Service Provider is sometimes also called the Relying Party (RP).

The classic use case described in the original Liberty Alliance work is that of a user who logs into an airline reservation system and authenticates himself, makes bookings for a flight, then visits car-rental and hotel sites. If these sites are federated with the airline as an Identity Provider, the user is able to make reservations at the car-rental and hotel sites using the airline system's authentication system without any need to re-authenticate himself. Another example is that of a user at a corporation who has authenticated herself to the corporate system and then accesses an outsourced online service, such as a travel agent or health-insurance provider. The user is able to access the outsourced services without further authentication – this is the single sign-on provided by federated identity management – and, depending on arrangements, the services themselves may themselves be able to access protected corporate resources.

How authentication occurs varies: the SP can initiate the authentication request to an IdP that the user designates when signing on to the SP, or the user may first authenticate herself at the

IdP and then access an SP. In either case, the technology enables single sign-on (SSO), in which the IdP authenticates the user, thus allowing her access to protected resources at an SP.

Federated identity management is an example of a two-sided market [32], where two types of users are served by a common platform. As shown in Figure 1, the two sides in federated identity management are Identity Providers and Service Providers. Two-sided markets exhibit cross-side network effects: the value of the platform to one type of user depends on the number of users of the other type. This effect tends to yield very dominant platforms. Other examples of two-sided markets include payment-card networks (cardholders and merchants), newspapers (advertisers and subscribers), and operating systems for both PCs and phones (users and application developers).

One crucial aspect of identity management is the level of trust that can be placed in the identity claims. How much can the Service Provider rely on the assertions made by the Identity Provider? The degree of certainty a Service Provider has regarding an authentication after receiving an identity assertion from the Identity Provider is called “assurance.” Now authentication is a multi-step process that starts with a proofing mechanism that binds an identity (e.g., an email address or a user name) to a token (e.g., a hardware token or a password), uses a remote mechanism for authenticating, and has a mechanism for communicating the results of the remote authentication [11, p. 2]. Building on that, the U.S. National Institute of Standards and Technology (NIST) developed a set of four identity “Assurance Levels:”²

- Level 1: At level 1, there is no identity proofing; names are assumed to be pseudonyms. Authentication requires that the user demonstrate that she controls the token. The sole protection of user secrets comes from the requirement that user proofing data not travel in the clear, and the only thing that the level 1 mechanisms do is provide some assurance that it is the same user who is accessing the protected data. [11, p. 31] Level of Assurance (LoA) 1 gives minimal confidence about the user’s asserted identity.
- Level 2: At level 2, some identity proofing is required. (There are different requirements depending on whether the identity proofing is in person or remote; if in person, the user must show a valid current government identity document that has a picture as well as either nationality or address of record, while if remote, a financial account number is also required [11, p. 22].) Passwords and PINs are allowed for authentication, as are more secure forms of authentication (such as hardware tokens). There are system security requirements, e.g., there must be mechanisms to handle revocation of credentials, passwords must be a certain strength, etc [11, pp. 32-33]. Thus LoA 2 provides some assurance regarding the asserted identity.

²This is an abbreviated list of the requirements for the four levels; for the full set of requirements, see [11].

- Level 3: At level 3, the identity documents must be verified, with a higher level of proofing on the identity than for level 2, and two-factor authentication is required [11, p. vii]. In addition, the level 3 authentication mechanisms require cryptographic-strength protection of the primary authentication token [11, p. 33] (the token can be unlocked through a key or biometric [11, p. vii]). LoA 3 gives high confidence in the identity being asserted.
- Level 4: At level 4, identity proofing can only occur in person; the government ID is to be verified with the issuing agency [11, pp. 23-24]. The assertion mechanism is “hardened,” that is, only “hard” cryptographic tokens can be used, the FIPS 140-2 cryptographic module validation requirements are strengthened, and all critical data transfers are authenticated through a key bound to the authentication process. The user must prove that they control the hardware token [11, p. viii]. LoA 4 gives very high assurance in the asserted identity.

LoA 1 is often used as a “persistent identifier,” ensuring the user’s identity stays constant over the course of several visits to a site without actually providing information about who the user might be³. LoA 2 is used for low-value self assertions of identity; and sharing information resources across universities typically requires LoA 2 authentication. But accessing someone else’s data typically requires further assurance of the identity of the information accessor. The University of Wisconsin, for example, requires LoA 3 authentication for accessing restricted data other than one’s own [37]. LoA 4 would enable a law-enforcement official to access a controlled site such as a law-enforcement database containing criminal records.

The NIST assurance levels have also been adopted by other governments for providing e-government services, including Australia, Canada, New Zealand, and the UK, as well as for many other identity management systems.

The value of federated identity management is its simplification of function: the business of authentication is separated from the process of accessing resources and everyone – the user, IdP and SP – can benefit [24, p. 17]. The user only has to log in once with a single set of credentials. The Identity Provider can focus on improving the process of authentication, perhaps providing different modes of strengths of authentication, perhaps providing other services. The Service Provider no longer has to handle authentication – a messy, problematic business – and can focus instead on the provision of services. It would seem as if everyone benefits by this simplifying of roles. That is exactly the issue we wish to study further, so with this brief introduction, we turn to examining specific use cases.

³Of course, by “fingerprinting” the browser and doing other types of checking, the site could quite possibly determine a great deal about who the user’s might be; the point is that the identity information being provided need not reveal anything about the user other than the series of transactions in which the user has engaged.

3 Federated identity management systems: use case successes and failures

Tolstoy wrote, “Happy families are all alike but every unhappy family is unhappy after its own fashion,” [35, p. 13]. Federated identity management does not fit this pattern; successful deployments are all alike in their technology, but differ in what enables their success; unsuccessful deployments are all alike in their cause for failure to overcome economic tussles. We examine four successful federated deployments: InCommon, a higher-education-based resource-sharing system, the National Institutes of Health efforts in federated identity management with research institutions, Sun’s federation with its outsourced human resources partner, Hewitt, and Aetna’s system for managing medical billing. We then consider two less successful efforts: a federal government effort to promote information sharing across law-enforcement agencies and the OpenID framework for online authentication. We conclude this section by describing payment-card networks, which, while not strictly an FIM deployment, authenticate payments by individuals and have underpinned the successful rise of electronic commerce.

3.1 InCommon and online sharing of library resources

Universities are places of learning. Even in the Internet age, sharing of information resources form an integral part of this aspect of universities’ fundamental mission. The question arose of how to build an infrastructure supporting relatively frictionless information-sharing resource sharing between higher-education institutions. One solution was InCommon.

As of January 2011, InCommon serves 189 higher education institutions, eight government and non-profit research labs, centers, and agencies, and 69 “sponsored partners,” including publishers and medical libraries [19]. InCommon vets each institution’s credential-provisioning system: who provides them, how they are given out, how long they last, what information is made public in the identity database, what is kept private, and so on. Shibboleth⁴ is the technology used for sharing secured web resources and services among InCommon members. Identity Providers are typically higher-education institutions, while the Service Providers are libraries, commercial information providers (e.g., Lexis Nexis), and even individual labs and research groups.

From its inception, privacy was critical to the Shibboleth design. Why? First, the U.S. Family Educational Rights and Privacy Act⁵ protects the privacy of student educational records, possibly including library records. Second, librarians regard user privacy as inherent to the library’s

⁴Shibboleth is Hebrew and comes from the criterion used to distinguish one group, the Ephraimites, from another, the Gileadites.

⁵20 U.S.C. §1232g.

mission⁶. Shibboleth was developed so that “the users should control what personal information is released and to whom, and the resource provider should only receive as much user information as needed to make access control decisions unless the user chooses to release more” [25, p. 14]. Thus for example, the user login, which frequently functions as a user ID, is just an attribute in the Shibboleth design. Users are identified by their rights to the resources: as a member of a campus, as a member of a course, as a member of a cross-institution research group accessing shared resources. The user ID is simply another attribute of the user, to be shared only if access to the resource requires it [25, p. 14]. Consequently, the benefits of Shibboleth to the user are clear: simple privacy-protecting access to information resources.

The benefits to the Identity Providers are also clear, for InCommon provides broader access to resources in a privacy-preserving and extensible manner. IdPs do not need to build pair-wise relationships with each of the SPs; the Shibboleth infrastructure takes care of that.

There is also clear benefit to the Service Providers. One aspect of this is that SPs only have to handle authorizing groups of users from <this campus, that research group, this course> . Thus SPs are no longer responsible for authenticating users outside their own domain. In addition, Shibboleth provides better security than previous solutions for such remote access. Morgan et al. [25, p. 15] describe how JSTOR, a non-profit organization maintaining an archive of scholarly journals, had used a system of allowing access from a block of IP addresses allocated to an institution. But the interactive authentication system afforded by Shibboleth provided more nuanced and far better security. In addition, Shibboleth enabled JSTOR to personalize services to the user – even though it didn’t necessarily know whom the user was!

In sum, the InCommon system is a clear win for all participants: Identity Providers, Service Providers, and users.

A similar project is the Nordic Kalmar 2 Union, which is an inter-federation of national federated efforts (countries participating include Denmark, Finland, Iceland, Norway, and Sweden) [42], thus enabling cross-border authentication. Some IdPs are directly accessible through the Kalmar Union, while Danish users go through the Where Are You From (WAYF) service, which redirects the user to the appropriate IdP within Denmark. Thus authentication involves two steps: first the user is directed to the list of Kalmar countries; she chooses the correct one, then picks an IdP from that nation’s providers (e.g., Danish users are directed to `wayf.dk`). Once she authenticates, she is redirected back to the SP. All applications are authenticated at least to level of assurance 2.

An interesting aspect of the federation is that there is bank participation in WAYF as an IdP. However, that is an outgrowth of banks running the Danish NemLog-in, a single-sign-on system

⁶The American Library Association includes unrestricted access to information and guarding against impediments to open inquiry as part of its interpretation of the Library Bill of Rights [2].

for citizen access to public services (banks won the contract to run the service), and NemLog-in operates as an IdP for WAYF. There is currently no other role for banks with WAYF. While there is no objection to banks functioning as IdPs, thus far a business case has not been made for doing so [34].

3.2 InCommon and the National Institutes of Health

The National Institutes of Health (NIH) is the premier U.S. Department of Health and Human Services agency performing biomedical research. Or, from the point of view of some of its employees, NIH is a government institution, a set of research laboratories, and, on occasion, a patient-facing organization. These multiple roles mean that there are a plethora of requirements for NIH researchers to follow. This includes U.S. government regulations that applications and services should be online and that these should satisfy appropriate authorization requirements on secure identity management. If the application or service does not comply with federal requirements, research funding can be taken away (which has happened) [1, slide 6].

NIH researchers also often collaborate with researchers from universities and national laboratories both in the U.S. and abroad. The result is a remarkable array of authentication requirements that are quite difficult for an individual researcher to satisfy. NIH's solution has been to deploy an infrastructure that enables NIH staff to easily share information with outside colleagues by relying on credentials provided by collaborators' own organizations. In other words, NIH uses federated identity management to manage access to resources.

NIH relies on InCommon to vet member institutions but adopts a finer grained vetting of departments within a member institution (e.g., Duke Medical School, rather than Duke University). The NIH system relies on the application to determining the level of assurance (as outlined in Section 2) required. The NIH Login system manages the technology for accomplishing the login, while the NIH Chief Information Office manages the trust relationships (e.g., InCommon, relationships with the FDA and CDC, etc.) underlying the technologies [1, slides 7 and 9].

The result: the researcher tells NIH Login who her users are; the users authenticate with their home institution; the NIH federated login maintains a list of institutions trusted to perform authentication (this may be accomplished through the use of InCommon). It is a system in which everyone benefits. The researcher focuses on research, and her lab resources are not spent on vetting remote users; the Identity Providers vet their own participants, and the Service Providers are able to trust those vettings (and the Identity Providers and Service Providers often change roles in this system).

3.3 Sun Microsystems outsourced services

Sun Microsystems Inc. was one of the original leaders on the Liberty Alliance effort. In Silicon Valley fashion, Sun elected to implement federation identity management in some of its outsourced services, a way of testing the feasibility of such systems and, in the long term, a way of achieving cost savings.

In this federation model, the user was the Sun employee, the IdP was Sun Microsystems, and SPs were the various outsourced entities, including Hewitt, which handled Sun's Human Resources, and American Express, which handled Sun employees' travel arrangements.

The Liberty (later SAML) protocols were open standards-based specifications for identity federation and competed with Microsoft's Passport, which gave Sun a powerful incentive to support the work. Furthermore, if Sun's partners were able to succeed as Service Providers, they could use their experience to expand the effort to other companies. Thus both sides had strong business reasons for wanting the project to succeed. (The user, as a Sun employee, had little choice in this discussion.)

Much effort was expended on the initial legal contracts. Agreements were carefully drawn up on how to handle adding or withdrawing capabilities in the framework, upgrading authentication requirements in the future, and recovering from failures. The roles and responsibilities of the different institutions were carefully delineated. This included what Sun's business project's lead role was, what the SP business project's lead was, what business support Sun would provide, what the partner would provide, what level of IT support would be expected of each entity, how quickly the parties would respond in case of failures, the process to be followed during security breaches, and penalties for being late or system outages.

In the end, none of this legal infrastructure turned out to be important. When problems did arise (such as an employee getting the wrong access), both companies worked to solve the the problem without raising legal issues. A Sun employee described the tolerance by the fact that as an IT company, Sun knew that "such screw-ups did occur." The reluctance to seek legal remedy was that Sun understood, "There but for the grace of <deity of choice> go I. We (in IT) well understand that it is difficult to get everything right, especially when you're doing something for the first time. If the other party agrees to fix the problem, you just move on" [39].

The underlying reason for the spirit of cooperation was that Sun and Hewitt were both invested in the success of the project. So when difficulties arose both were motivated to solve the problem rather than to blame the other. Consequently, none of the carefully-crafted legal remedies were put to use. There have been numerous Liberty/SAML identity management implementations enterprise settings, and similar motivating factors have contributed to their success. We next discuss one such example, the billing system used by Aetna.

3.4 Aetna’s medical billing system

Aetna, the insurance company, has also deployed a federated system to manage the billing of medical practices. Aetna had been building its own identity assurance framework for online billing when the company discovered the NIST authentication guidelines [11]. They found that using a common reference model has reduced the cost of deployment, clarified requirements, simplified audit procedures, and has made working with partners more straightforward. Furthermore, Aetna saw using the NIST model as a way to increase identity assurance [14].

In the Aetna federated system, the Service Providers are medical billing offices (ranging from small practices to large). Aetna serves as an Identity Provider for credential provisioning (e.g., “this” practice is within our business network), as does NaviMedix, a provider of software for secure online systems, which manages the credentialing of the offices. Aetna tells NaviMedix that the practice is in network; NaviMedix credentials individual users within the system. NIST levels of assurance provided a standardized methodology for coordinating the procedures used by Aetna and NaviMedix.

Both Aetna and NaviMedix function as IdPs in the Aetna system; Aetna is also a Service Provider, as are the billing offices. The identity system is designed to be compliant with the Health Information Portability and Accountability Act (HIPAA). Within a particular office, users are granted different access to different types of data based on their roles. Front-office staff can only access appointment and check-in information, while accounting staff can access claim and payment information [27, p. 4]. Like the InCommon federation, the Aetna system is based on SAML (Security Access Mark-up Language) 2.0, an XML open standard for exchanging authentication and authorization information. SAML 2.0 represents a convergence of standards work in OASIS and the Liberty Alliance.

As of 2008, the Aetna system was used by three hundred thousand providers, with capacity for up to half a million [27, p. 3]. The success was enabled by a combination of factors. The NIST standards provided a common platform for Aetna and NaviMedix to handle identity credentialing. SAML provided a robust infrastructure for conducting the online transactions. The existing relationships between insurance companies and medical billing offices removed some natural friction and provided some necessary motivation to make the system work. Finally, HIPAA regulations clarified the responsibilities for different organizations within the system to protect the privacy of patient information, as well as liability for failing to do so.

3.5 Information sharing across law-enforcement agencies

In the federal government, the post September 11th world heavily relies on “information sharing.” One example is *Intellipedia*, an information-sharing site for the intelligence community

modeled on Wikipedia. The idea, proposed by the CIA Chief Technology Officer, is that any agent with classified clearance should be able to read or contribute relevant knowledge to the site. Another example is the joint Department of Justice and Department of Homeland Security Global Federated Identity and Privilege Management (GFIPM), a federated identity management system for the sharing of secure and trusted information.

GFIPM is a pilot project sponsored by the Criminal Information Sharing Alliance Network (CISA), the Pennsylvania Justice Network (JNET), and the Regional Information Sharing System Network (RISS) and run by the Georgia Institute of Technology. It has been somewhat successful, and numerous state and local agencies elected to participate⁷. But, as in any complicated system, a closer examination reveals some interesting details.

The first issue to note is that the information being shared regards state-level investigations, not federal ones. Thus, while the concerns include both criminal and national-security issues, they are more heavily weighted to the former. The second issue to note is an observation made by the Georgia Institute of Technology implementors: “IdPs are easier to integrate than SPs” [38, slide 47]. We believe that is key to understanding GFIPM’s limited success in deployment.

GFIPM clearly provides benefit to users: a single sign on gives access to multiple sites. It also provides clear benefit to the Identity Providers: their users then have access to multiple sites, without any extra work on the side of the IdPs (except for the effort of developing the initial architecture). But the same issue that benefits the IdPs creates a problem for the Service Providers, for in letting users from other domains seamlessly enter into their systems, the SPs lose control. As Nigriny and Sabett and numerous others have already observed, in the issue of security clearances, each agency trusts its own vetting process, but doubts the processes of other agencies [30, p. 8].⁸ They may also object to sharing secret information with outsiders, and resent the imposition from above that they share more extensively.

3.6 OpenID standard for online authentication

As noted earlier, Microsoft’s original identity management system Passport was centralized. Passport was designed so that Microsoft was the sole identity provider; unsurprisingly, very few Service Providers agreed to such terms.

⁷The list includes the Georgia Bureau of Investigation, the California Department of Justice, the Oklahoma State Bureau of Investigation, numerous Pennsylvania law-enforcement-related agencies, the County of Los Angeles.

⁸The Wikileaks of US State Department secret cables provides a striking example of this. These “cables” (they are, of course, no longer cables, but the name has stuck for historical reasons) were on the “Sipdis”, or Siprnet Distribution, the US military Internet (which is separate from the public Internet). This means that the information was available on US internal embassy websites and by US military – probably a much wider distribution than the State Department had ever intended [9].

The appeal of single-sign on for Internet services did not diminish following Passport’s failure, however. OpenID was established in 2005 as a non-profit platform using an open standard for federated identity management. In OpenID, notions of “identity” are weak – typically, a password-protected user account on a web-based email service. For many online interactions, such as posting a comment on a blog, such a weak level of authentication suffices.

OpenID has recruited many Identity Providers but very few Service Providers. OpenID is attractive to end users, who stand to gain by reducing the number of online accounts they must maintain. The benefits of OpenID to Service Providers are much less obvious. Most web services are supported by advertising, so collecting targeted demographic information on a website’s users is very valuable. While OpenID has developed an attribute exchange mechanism that allows for rich exchange of user demographic information, Identity Providers are free to choose which characteristics to share. To date, most IdPs have elected to share very few user details with service providers. For instance, with user consent, Google will share name, country, email address and language [17]. Upon request, Yahoo shares name, email address, profile picture and gender [43]. Given such limited information, many providers would rather use their own registration mechanisms to collect richer user profiles.

Meanwhile, the benefits accrued to OpenID Identity Providers are much more clear than those to Service Providers. OpenID encourages user loyalty to the Identity Provider, who learn quite a bit about the browsing habits of their users. For example, an Identity Provider is informed each time a user authenticates to a Service Provider. This information is valuable to Identity Providers since it can be used to create more tailored user profiles to target advertising better than would be possible without OpenID.

In one might be seen as Passport redux, Facebook has developed a centralized system where its users can log in to third-party websites using Facebook credentials. To attract wary Service Providers, Facebook shares social network information in addition to demographic information. Figure 2 illustrates the difference between the information shared with Service Providers using Facebook, compared to OpenID. The left screenshot shows the request `stackoverflow.com` issues for logins using Google credentials via OpenID, while the right screenshot shows the request from `nytimes.com` when logging in from Facebook. With OpenID, the Service Provider only learns the email address, while with Facebook the Service Provider learns the name, gender, list of friends, and all public information stored by Facebook. All profile information, including birthday, education and work history is also shared. Unsurprisingly, given the rich user data on offer, Facebook has managed to attract the participation of many more Service Providers than OpenID has.

In Section 4 we examine quantitatively the adoption of different online authentication mechanisms by top websites. The evidence we have collected indicates that IdPs that share the user’s

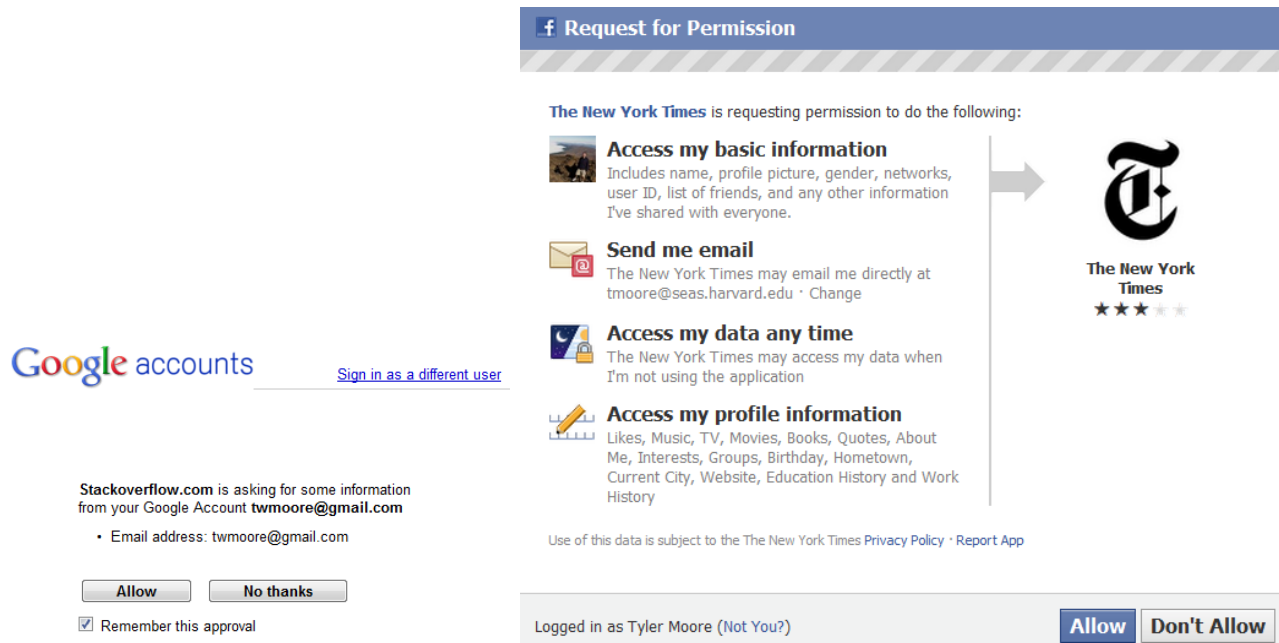


Figure 2: Screenshot of requests to share personal information when logging into `stackoverflow.com` with OpenID (left), and to `nytimes.com` with Facebook serving as Identity Provider (right).

social graph enjoy greater adoption.

3.7 Payment-card networks and e-commerce

While not strictly an identity management system, payment-card networks such as MasterCard and Visa share some characteristics with these systems, and have greatly contributed to the success of electronic commerce. They are thus worth examining in this context.

Payment-card networks offer a ready-made solution to processing payments, a prerequisite for many forms of e-commerce. The two sides of payment networks are issuing banks, which issue credit and debit cards to customers, and acquiring banks, which accept payments on behalf of merchants. In terms of identity management, issuing banks roughly correspond to Identity Providers while acquiring banks correspond to Service Providers. Issuing banks decide when to issue cards to consumers, including verifying the identity of the cardholder and assessing the credit risk.

Payment networks such as Mastercard and Visa set rules for interaction between issuing and acquiring banks. They are also responsible for promoting the growth of the overall network, attracting new cardholders and merchants. Payment networks were a natural fit for e-commerce because they provided an existing means of authenticating individual payments for a very large

number of consumers. Furthermore, payment networks appealed to existing businesses that already accepted credit-card payments for their traditional offline interactions.

It is instructive to consider how the rules of payment networks have been tweaked to fit the context of e-commerce. In particular, the rules for assigning responsibility for reimbursing and protecting against fraud have been adapted over time [23]. In the US, consumers are protected from liability for unauthorized charges on their accounts (credit cards are covered by the Truth in Lending Act of 1968, implemented by the Federal Reserve as Regulation Z, while debit card holders are covered by the Electronic Funds Transfer Act, implemented through Regulation E). Instead, the obligation to repay is allocated between issuing banks and merchants. For frauds occurring in face-to-face transactions, issuing banks normally foot the bill, rather than merchants. The rationale is that most face-to-face fraud cannot easily be prevented by merchants. Abnormally high levels of face-to-face fraud occurring at a single merchant raises suspicions that the merchant may be complicit in the fraud, but in other cases payment networks are happy to tolerate isolated incidences of fraud.

Because online transactions where the card is not present are inherently riskier than cases where the card is physically present, payment network rules often dictate that the merchant has to pay for fraudulently authorized transactions. Most merchants accept these less favorable terms because payment-card networks are the only viable option for processing payments, and a higher fraud bill is preferable to forgoing the opportunities provided by e-commerce. One consequence of this policy, however, is that payments originating from riskier sources (such as international payments) are frequently not authorized.

One might wonder why alternative methods of payment have not arisen to satisfy the particular needs of e-commerce (such as raising the level of authentication of cardholders in order to better mitigate fraud). One explanation is that the success of payment networks – with millions of participating merchants and cardholders – has created a significant barrier to new entrants offering innovations such as a more secure payment alternative. Having already invested heavily in a less secure payment technology and achieved market dominance, existing payment networks may be reluctant to invest further in security.

4 Empirical analysis of online authentication mechanisms

In Section 3.6 we argued that the OpenID standard for online authentication has attracted lower adoption rates than Facebook’s offering due to the extensive user information Facebook has been willing to share with service providers as enticement. We now provide empirical support to this claim.

Identity Provider	Logins on Top 300 Sites		% Internet Users	Shares Social Graph
	#	%		
Facebook	37	34.9	40.52	True
Google	12	11.3	11.07	False
Yahoo	10	9.4	5.03	False
MyOpenID	5	4.7	16.11	False
LinkedIn	5	4.7	4.11	True
Twitter	14	13.2	9.82	True
MySpace	5	4.7	1.36	True
Windows LiveID	2	1.9	5.30	False
AOL	4	3.8	0.77	False
Blogger	2	1.9	13.09	False
Flickr	1	0.9	2.26	False
Hyves	0	0.0	0.17	False
LiveJournal	2	1.9	1.17	False
Netlog	0	0.0	0.35	False
PayPal	0	0.0	2.35	False
Verisign	2	1.9	0.03	False
Wordpress	2	1.9	4.89	False

Table 1: Observed availability of online authentication mechanisms in the Alexa top 300 sites.

Data Sources There is no shortage of web-login systems on offer. Table 1 presents 17 such systems, along with additional characteristics gathered from different sources. We manually visited in April–May 2011 each of the 300 most popular websites, according to Alexa⁹. We first checked whether the site allowed users to login, and if so, whether they used their own system or allowed users to login using one of the 17 IdPs. We successfully classified 135 websites. We excluded from our analysis websites run by the IdPs (e.g., `google.com` and its country-specific variants), pornographic websites, and foreign-language websites where we could not assess login options. 106 of the 135 websites allowed users to login, and 102 of these offered their own login service.

The observations of each outside login service are given in Table 1. Facebook appears most often, in around 35% of websites with logins. While Google and Yahoo are both OpenID IdPs, we counted separately the times that Google, Yahoo, and OpenID were presented as the login options (since sometimes only Google or Yahoo were given as login options). Twitter was the second-most popular offering, followed by Google and Yahoo.

Table 1 also presents a measure of the popularity of each of the services, here listed as the percentage of Internet users who visit each service’s website monthly according to Alexa. For Google, Yahoo and Windows Live ID, we use the estimate for the number of visits to each com-

⁹<http://www.alexa.com/topsites/>

pany’s webmail service, since these serve as the basis for OpenID credentials. For MyOpenID, we combined the popularity of Google and Yahoo, since these two companies represent the bulk of the number of OpenID users.

The final column in Table 1 indicates whether the service shares its users’ social graph with websites. Unsurprisingly, each of the social networks – Facebook, LinkedIn, Twitter and MySpace – do share this information. While OpenID does offer the capability for other sites to share more extensive information, in practice this does not occur. For instance, only one site allowing Google logins requested the user’s contact information, while the rest only requested the email address, country and language of the user.

Empirical analysis We devised a linear regression using the prevalence of an authentication mechanism in the most popular websites as the dependent variable:

$$\text{PctTop300Logins} = \beta + \text{PctWebUsers}x_1 + \text{SharesSocialGraph}x_2 .$$

The two explanatory variables are the size of the IdP’s user base, measured by the percentage of web users who visit its website each month according to Alexa, and whether the service shares the social graph of its users with other websites. Both factors positively correlate with the percentage of sites running the login service. The results are statistically significant:

	Coefficient	Std. Error	Significance
PctWebUsers	0.664	0.099	$p < 0.001$
SharesSocialGraph	5.27	2.36	$p = 0.0335$
$R^2: 0.8445$			

Interpreting the coefficients, a one percentage point increase in the number of Internet users (corresponding to around 10 million users) using the login’s service corresponds to a 0.66 percentage point increase in the fraction of websites offering the service as an option for logging in. Likewise, when the IdP can offer the social graph as enticement to prospective service providers, the fraction of websites offering the login service corresponds to an increase of 5.3 percentage points. Notably, the R^2 is 0.8445, which suggests that these two factors explain over 80% of the variation in login-service uptake.

Note that while this analysis is consistent with our argument that service providers select identity providers based on how much user information they are given, other explanations are also possible. For example, websites could offer logins via social networks because they are currently “hot,” or out of a belief that the social features better engage their own users.

5 Tussles in federated identity management

In a seminal paper, Clark et al. identified a number of so-called “tussles” framing the evolution of the Internet’s design [13]. Tussles occur whenever the interests of stakeholders conflict in the design of an engineered system. In a highly distributed and highly valuable network such as the Internet, conflicting interests are inevitable. We argue that federated identity management systems are also subject to many tussles. This is due in part to the complex engineering task of designing and deploying a working system, but it is compounded by its two-sided market structure and by government interest in finding an acceptable solution.

Informed by the cases described previously, we now outline key tussles. When the incentives of stakeholders align on these tussles, then the identity management system has a fighting chance of success. When conflicts arise, then the given application is more likely to fail.

Tussle 1: Who gets to collect transactional data? Any identity management system generates rich evidence of transactions as a natural byproduct. Which stakeholder (if any) gets access to transactional data can be crucial to the success or failure of an identity management system.

In a few circumstances, user control over transactional data is explicitly guaranteed. For instance, to comply with US federal laws, Shibboleth includes privacy protections for users over what information is shared. Most private companies, however, would demand much weaker privacy guarantees to users in order to participate. Indeed, the tussle most private firms would prefer to engage in is between Identity and Service Providers over who controls transactional data.

Consider Internet single-sign-on services. OpenID Identity Providers have managed to keep user demographic information largely away from Service Providers, while at the same time Identity Providers are in a position to learn a great deal about user transactions from many Service Providers. This imbalance of transactional data control goes a long way towards explaining why few Service Providers participate in OpenID despite an enormous user base based on Identity Providers. By contrast, Facebook has enjoyed more success attracting Service Providers by agreeing to share more user data (namely the social graph) with them.

Government intervention could also change the dynamic on the handling of private information. For decades European regulators have used the Fair Information Practice Principles to protect privacy. Thus in the early days of federated identity management systems, Microsoft Passport’s centralized architecture drew significant scrutiny from the Article 29 Data Protection Working Group (a European Commission group whose membership consists of the “privacy commissions” of the E.U. Member States). The working group requested “radical changes in the information

flow data” and set deadlines for implementing various changes to the architecture [16, p.23]. By contrast, the group’s response to the Liberty Alliance’s federated architecture was positive: the working group simply asked to be kept informed of Liberty’s future steps [16, p. 23]. The efforts of European regulators have not abated; in recent years they have affected the amount of time search engines retained cookies [8], the deployment of Google’s StreetView, etc. More importantly, these efforts have impressed on U.S. technologists the need to be aware of — and design for — international requirements for protecting privacy.

An analysis by Bamberger and Mulligan describes a very different, and interesting, privacy protection mechanism developing in the U.S. [7]. Since 1996 the Federal Trade Commission (FTC) has been using its authority to act against deceptive trade practices to take an active role on privacy protection [7, p. 273]. Rather than issuing regulations directly, the FTC has relied on a combination of advisory committees, workshops, and reports, to help drive an emerging consumer privacy agenda [7, p. 286]. In this it has been aided by the development of state security breach notification laws and third-party advocates. These have created a dynamic that have caused corporations — at least the larger ones — to commit to adopting greater privacy protective stances [7, p. 250-252]. This is an important first step, but, of course, adherence is key. That is where the FTC has stepped in. The regulatory agency has been willing to take two key steps to improve enforcement: (i) fine for deceptive trade practices and (ii) determine what constitutes “unfair” trade practices. This creates a situation in which companies watch what is happening around them and continually seek to improve on current practice lest they become the poster child for poor privacy practice [7, p.274].

This dynamic should be considered in the context of a new U.S. government development on online identity management. In April 2011, the U.S. government introduced the *National Strategy for Trusted Identities in Cyberspace* (NSTIC), which provides a blueprint for private and public sector development of online trusted identities. This document strongly emphasizes privacy as one of the guiding principles (privacy is, in fact, the leading principle listed and is continually cited throughout the strategy document) [36, p. 3]. For example, the strategy emphasizes collecting and distributing only the information necessary for the transaction, keeping that information for a limited period of time [29, pp. 10-11]. In enforcement, it specifically lists legislation as one possible vehicle for increasing individuals’ privacy protections [36, p. 23]. But of course, the government already has one privacy enforcement mechanism. Given the FTC’s efforts on privacy protections, the NSTIC commitment to privacy in the development of online identity mechanisms should be taken seriously. In particular, within the U.S. developers of identity management systems should expect enforcement of the type described above.

Tussle 2: Who sets the rules of authentication? Sometimes, it is natural to determine which party should serve as Identity Provider, and which should serve as Service Provider.

Sometimes, though, multiple parties could serve both roles. The competition to serve as Identity Provider and thus to determine the framework's rules can impose unintended consequences on the methods of authentication that are ultimately adopted.

Identity management platforms offer a substantial first-mover advantage, which means that getting to market is more important than deploying the most robust technology for authentication. Once an established platform has taken hold, network effects can lead to lock-in. Payment networks offer an enlightening example. An entrenched payment network may be willing to tolerate a higher level of fraud so long as it can be recovered through fees. Once entrenched, the platform may be tempted to fiddle with the rules to shift liability on who has to pay for fraud once the switching costs for stakeholders have risen.

Related to this is who gets to select what is the appropriate level of authentication. If there is competition among Identity Providers to attract users, ease-of-use and simplicity are likely to be highly valued, even at the expense of more rigorous authentication mechanisms that are inherently less convenient and/or more costly (this was the reason for such broad early adoption of OpenID). In a competitive market for IdPs, security and privacy are initially likely to be less of a priority than growing market share, as in any other market with network effects [6]. In this scenario, higher levels of authentication or improved privacy guarantees are only feasible once a dominant IdP emerges.

However, because this is a two-sided market, to be successful an IdP also must also attract SPs, who may very well desire a higher level of authentication. Consequently, there could be countervailing pressure on IdPs to adhere to a baseline authentication level that is acceptable to prospective SPs but does not hinder adoption by users. Getting the balance right – one that does not favor onerous authentication requirements from SPs or watered-down IdP designs – is hard. Of course, the reason why different stakeholders are concerned with getting the authentication requirements right is that they worry about what might happen if authentication fails. This leads to our next tussle.

Tussle 3: What happens when things go wrong? There are two main ways failures can occur for an identity management system. First, the system could become unavailable to authenticate users, causing problems for SPs relying on its operation. Second, the authentication itself could fail, in that unauthorized users could be incorrectly authenticated as other users. In both cases, rules that determine which party is responsible are important and a potential source of conflict.

In some cases, it is clear where responsibility should lie. For Shibboleth, the library system serving as IdP is responsible for the consequences of incorrectly authenticating users, as well as for its own users behaving badly. It helps that the IdP and SP roles are symmetric, so that both

sides have responsibilities in each role (even though not at the same time). For Sun’s outsourced services, the SP and IdP already had an existing business relationship, which both simplified the process of dealing with the unexpected and provided motivation for resolving the issue. In such a situation, when things go wrong, because both parties want the partnership to succeed, there is a sense of shared responsibility that helps solve the problem without recourse to the courts.

Payment card networks provide an example of how disputes may arise over who is responsible for failures. There have been several attempts to shift liability among the different stakeholders in the payment system. In part this is because the cost of fraud is easy to measure, and there is a perception from each side that the other is at fault. As noted in Section 3.7, US regulations severely limit the liability of cardholders for fraud, whereas in the UK such consumer protections have historically been weaker. Consequently, UK banks have attempted to shift responsibility for fraud onto cardholders [4,28] where possible. In the US, since cardholders are not responsible for fraudulent activity on their cards, the tussles have arisen between the merchant banks and the issuing banks that represent the payment networks. Merchants blame outdated authentication mechanisms in payment card networks for fraud, while the payment networks blame merchants for lax operational security. The push for PCI compliance of merchants reflects this tussle between merchants and payment networks over who should pay for fraud.

What’s at stake also plays a role. For low levels of risk, clearly assigning liability for failures among players is less essential. Authentication failures for web-based credentials could disrupt access to online resources or leak private user information, but the potential financial impact is usually minimal. Some types of risk are financially significant but easy to measure, such as payment-card fraud. Here, liability arrangements must be clearly articulated and fairly distributed among stakeholders, but so long as there is an expected positive financial payoff, agreements are likely to be made.

A final class of risk is where the costs of failure are large and poorly understood. Introduction of a federated identity management system could elevate the risk of failure and introduce additional new liabilities for failures that may arise. Unfortunately, many of the federated identity use cases fall into this category. For example, critical infrastructures such as those in the energy sector are perceived to be at an elevated risk for attack by unauthorized insiders. Suppose a federated identity management system is developed to enable utility workers from across the nation to assist a region hit by a natural disaster using credentials issued by their home organization. Suppose furthermore that the credentials of one of the assisting workers is compromised and used by an attacker who prolongs an electricity outage rather than abates it. Might the assisting company be held liable for the actions of the attacker? Would the identity provider? Liability arrangements have to be drawn up in such a way that take into account these low probability but high impact events. Alternatively, some type of “hold harmless” agreement similar to the

one Wells Fargo agreed to before serving as credentialer to the Federal PKI may be required [15].

Tussle 4: Who gains and who loses from interoperability? One of the key advertised benefits of federated identity management systems is that users authenticated by one Identity Provider can be served by multiple Service Providers. The benefit or risk of such increased interoperability can vary by application and by stakeholder. When both Identity Providers and Service Providers see a clear benefit to increased interoperability, then the platform is more likely to succeed. If either IdP or SP does not view interoperability as beneficial, then the platform may be doomed.

Consider the case of federal security clearances discussed earlier. Policymakers decided that increased information sharing was important, and so encouraged the adoption of an identity management platform to facilitate sharing. For Identity Providers, this is an easy sell, because it could lead to increased access to intelligence information from other agencies. However, where IdPs see opportunity, SPs see risk due to increased access to sensitive information. To SPs wary of information leakage, lack of interoperability is a feature, not a bug.

OpenID provides a second example of when this tussle may arise. End users stand to gain from the convenience of single-sign on, while IdPs gain from collecting more information on user browsing habits. However, SPs do not stand to gain much from interoperability: they are unlikely to gain many new customers, only those very marginal customers who are only likely to visit once and not find it valuable to register.

6 Federated identity use cases revisited

We now return to the use cases outlined in Section 3 in light of the tussles just presented. Table 2 summarizes our findings. For the first three use cases — Shibboleth, Sun outsourcing, and the NIH identity-management platform — none of the four tussles presents an insurmountable obstacle for any of the stakeholders. Unsurprisingly, these three cases represent clearly successful applications of federated identity management.

Conflicts appear in the remaining three cases. Tussles 1 and 2 do not present a problem for federal clearances because the rules for transactional data and authentication have been externally imposed on prospective Service Providers and Identity Providers by the government. However, Tussle 3 may be a problem, since authentication failures means that the Service Provider suffers the problems even though it may be the Identity Provider that made the error. Tussle 4 is the biggest roadblock, since organizations may be reluctant to authenticate outside users because they may not want to share sensitive information with others.

	Tussle 1	Tussle 2	Tussle 3	Tussle 4	Success?
	Who Collects	Who Sets	When Things	Interoperability	
	Trans. Data	Auth. Rules	Fail	Gains/Losses	
Shibboleth	✓	✓	✓	✓	✓
NIH FIM	✓	✓	✓	✓	✓
Sun outsourcing	✓	✓	✓	✓	✓
Aetna's billing	✓	✓	✓	✓	✓
Clearances	✓	✓	✗	✗	✗
Open ID	✗	✗	✓	✗	✗
Payment networks	✓	✓	✓*	✓	✓

Table 2: Comparing use cases for their susceptibility to the tussles discussed in Section 5. A ✓ indicates that stakeholder interests are aligned so that the tussle can be overcome, while a ✗ indicates that the tussle is a source of conflict that may undermine the success of the IdM application. (*: Liability assignment for payment networks depends on the laws and regulations in the operating environment.)

Several tussles cause problems for federated identity management on the open Internet, for reasons explained throughout Sections 3 and 5. Notably, disagreements over who controls transactional data (Tussle 1) has undermined OpenID. Additionally, OpenID's reliance on very weak authenticators has attracted IdPs but very few SPs (Tussle 3). IdPs also gain more from interoperability than SPs (Tussle 4). In sum, it is not surprising that the OpenID effort has been a commercial failure. Non-federated alternatives such as Facebook may be successful by resolving Tussle 1 through sharing social-network data with Service Providers, though the level of authentication is very weak and below the level desired by some prospective Service Providers.

By some measures, payment networks have been a runaway success, with widespread adoption among users and merchants. E-commerce has relied on payment networks to complete most online financial transactions. However, the preceding discussion has also identified several undesirable outcomes of payment networks. Notably, disputes over who should pay for fraud arise frequently and continue to be contested. Additionally, the success of payment networks has created substantial barriers to entry for alternative arrangements in both payments and also identity-management solutions for online transactions.

7 Insights and concluding remarks

When seeking explanations for why federated identity management systems have not yet succeeded in the broad way anticipated at the beginning of the last decade, many have pointed to

the inability to sort out liability as a major cause of failure [29, p. 22]. We believe this is a mischaracterization of the problem. Instead, a more useful perspective is to examine the natural economic tussles that arise between the stakeholders in any engineered system.

All three parties in a federated identity management system must gain from the transaction, or there will be no incentive to use the system. A user has to gain through ease of use, access to more services, greater privacy, or improved security. A Service Provider has to gain by acquiring more user data (the Facebook model), in the ability to reach to larger markets, or by insulation from liability for failures (as happens in some instances of credit-card usage). An Identity Provider must also gain from the system. The gain in control of user data and of the user authentication process are obvious benefits to the Identity Provider, but those gains must be offset by granting some benefits to the Service Provider and user.

Looking at the situation in this light, it is clear that the early enterprise-oriented systems such as the Liberty Alliance protocols did not provide sufficient benefit to the individual so as to create widespread adoption (e.g., in the open Internet). However, certain instantiations such as InCommon or the NIH Federated system did provide these benefits, and uptake was high. (It can be argued that in those two instances the users did not have alternatives, but the fact remains that the systems provided clear advantages to users.)

Privacy, interpreted here as user control over personal data collection, should also be viewed from this perspective. Upon examining what has been produced in the market so far by OpenID and Facebook, users have been overlooked in the tussle between Identity and Service Providers over who controls user data. To handle this, some have proposed user-centric design in identity management systems [18,21], where control over transactional information is granted to end users who then decide what to share with Identity and Service Providers. The issue of control is a complicated one. Ease-of-use and user-data privacy are often in conflict; the success of the Kantara Initiative User Managed Access [21] and similar projects depends critically on easy methods for users to control their data.

Government regulators and policy makers also have a role to play if user privacy is to be included in successful systems. European data privacy commissioners have taken an active role in these issues; their negative response to Passport and positive one to the Liberty Alliance protocols were important in the early days of federated identity management systems. We suspect that the best prospect for achieving user privacy in future FIM deployments depends on a more active role by policy makers in advocating on behalf of users, who are largely voiceless in current debates over FIM proposals.

What are the lessons for the future?

Federated identity management systems exhibit a number of economic tussles, of which liability for failures is one. As in any complex engineered system, the tussles cannot be resolved sepa-

rately. Liability must be viewed as part of a larger set of economic tussles occurring between user, Identity Provider, and Service Provider. Resolving differences over liability in the context of other tussles may create opportunities for compromise. We are optimistic, therefore, that taking the broader view of all tussles may actually simplify the liability “problem” rather than complicate it.

Another way to put this is that if the Identity Provider accrues most of the benefits, it would be natural to also expect the Identity Provider to accrue most of the risk. At one level, that is obvious; at another, by isolating the various tussles, this begins to give us room to determine the bargaining that must arise between the three players. Of these, only two, the Identity Provider and Service Provider, are typically in the explicit negotiations; the users, of course, walk with their feet (or in this case, their fingers).

Another observation is that the payment-card networks have largely overcome liability issues between stakeholders and deployed a highly successful, if technically imperfect, system. When systems have failed to succeed commercially, it is usually caused by an unfair distribution of responsibilities and benefits between Identity Providers, Service Providers and users. Furthermore, one cannot expect any technology, including FIM, to solve irreconcilable incentive incompatibilities on its own. The key to success lies in setting the rules of the platform so that each stakeholder benefits from cooperation.

A key function of payment-card networks in e-commerce has been their ability to authenticate users for completing transactions. The early participation of American Express in the Liberty Alliance shows that there was initial interest by the payment-card industry, but there is less active participation now. Instead, much of their recent interest has focused on mobile payments, which may present an opportunity to fuse identity management and payments via mobile phones [3].

Payment-card networks already provide a usable solution to authenticating payments, which is the primary requirement for many e-commerce applications. Consequently, this weakens the business case for many aspiring identity-management solutions, particularly given the strong network effects present in two-sided markets and the high fixed costs of deployment. Furthermore, a widely-deployed FIM system might commoditize payment processing, particularly if their main competitive advantage is ubiquitous authentication of cardholders.

We conclude with an open question: can payment-card networks peacefully coexist with a successful, widespread deployment of a federated identity management system, or will the present success of payment-card networks prevent federated identity management systems from taking off in the open Internet?

References

- [1] Alterman, Peter, “U.S. Federal IdM/Federation Strategy and Your Apps”, NIH Federated Authentication Town Hall, November 29, 2007.
- [2] American Library Association, *Privacy: An Interpretation of the Library Bill of Rights*, <http://www.ala.org/ala/issuesadvocacy/intfreedom/librarybill/interpretations/privacy.cfm> [last viewed January 16, 2011].
- [3] Anderson, Ross, “Can we fix the security economics of federated authentication?”, 19th International Workshop on Security Protocols, Cambridge, UK, 2011, <http://www.cl.cam.ac.uk/~rja14/Papers/sefa-pr11.pdf>
- [4] Anderson, Ross, “Why cryptosystems fail”, ACM Conference on Computer and Communications Security, 1993, pp. 215–227.
- [5] Anderson, Ross, Ian Brown, Terri Dowty, Philip Inglesant, William Heath, and Angela Sasse, “Database State”, 2009, <http://www.cl.cam.ac.uk/~rja14/Papers/database-state.pdf>.
- [6] Anderson, Ross and Tyler Moore, “The economics of information security”, *Science*, Vol. 314, No. 5799, pp. 610–613, 2006.
- [7] Bamberger, Kenneth A. and Deirdre K. Mulligan, “Privacy on the Books and on the Ground”, *Stanford Law Review*, Vol. 63, No. 2 (2011), pp. 247-316.
- [8] BBC News, “Google Queried on Privacy Policy”, May 25, 2007, <http://news.bbc.co.uk/2/hi/technology/6692063.stm> [last viewed April 18, 2011].
- [9] Borger, Julian and David Leigh, “Siprnet: where America stores its secret cables”, *The Guardian*, <http://www.guardian.co.uk/world/2010/nov/28/siprnet-america-stores-secret-cables>, November 28, 2010 [last viewed February 23 2011].
- [10] Bolten, Joshua, Office of Management and Budget, Executive Office of the President, Memorandum to the Heads of all Departments and Agencies, “E-Authentication Guidance for Federal Agencies”, December 16, 2003.
- [11] Burr, William E., Donna F. Dodson, W. Timothy Polk, “Electronic Authentication Guideline”, NIST Special Publication 800-63, Version 1.0.2, National Institute of Standards and Technology, Technology Administration, U.S. Department of Commerce, April 2006.

- [12] Camp, Jean, “Identity Management’s Misaligned Incentives”, *IEEE Security & Privacy*, Vol.8, No. 6, pp. 90–94, 2010.
- [13] Clark, David, John Wroclawski, Karen Sollins, and Robert Braden, “Tussle in Cyberspace: Defining Tomorrow’s Internet”, *IEEE/ACM Transactions on Networking*, Vol. 13. Issue 3 (June 2005), 462-475.
- [14] Coderre, Mark, personal communication with Susan Landau, May 2, 2011.
- [15] Cross Certification Agreement between the United States Federal Public Key Infrastructure Authority and Wells Fargo Bank, November 26, 2008.
- [16] European Commission, “Seventh report on the situation regarding the protection of individuals with regard to the processing of personal data and privacy in the European Union and in third countries covering the years 2002 and 2003”, June 2004. http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2004/7th_report_prot_individs_en.pdf
- [17] Google, “Federated Login for Google Account Users”, <http://code.google.com/apis/accounts/docs/OpenID.html> [last viewed February 23 2011].
- [18] Higgins Open Source Identity Framework, <http://www.eclipse.org/higgins/> [last viewed February 23 2011].
- [19] InCommon, “Current Incommon Participants”, <http://www.incommon.org/participants/> [last viewed January 16, 2011].
- [20] Internet Policy Task Force, Department of Commerce, *Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework*, 2010.
- [21] Kantara Initiative, “UMA Work Group”, <http://kantarainitiative.org/confluence/display/uma/Home> [last viewed February 23 2011].
- [22] Landau, Susan, “NIST Leads the Charge on Online Authentication”, *Huffington Post*, January 12, 2011, http://www.huffingtonpost.com/susan-landau/post_1538_b_806394.html
- [23] MacCarthy, Mark, “Information Security Policy in the U.S. Retail Payments Industry”, 9th Workshop on the Economics of Information Security, Cambridge, MA, 2010, http://weis2010.econinfosec.org/papers/panel/weis2010_maccarthy.pdf
- [24] Maler, Eve, and Drummond Reed, “The Venn of Identity”, *IEEE Security and Privacy*, Vol. 6, No. 2 (2008), pp. 16-23.

- [25] Morgan, R.L. “Bob”, Scott Cantor, Steven Carmody, Walter Hoehn, and Ken Klingenstein, “Federated Security: The Shibboleth Approach”, *Educase Quarterly*, No. 4 (2004), 12–17.
- [26] Morris, Nigel, “ID cards go up in flames in first step to tackle ‘database state’”, *Independent*, February 10, 2011, <http://www.independent.co.uk/news/uk/politics/id-cards-go-up-in-flames-in-first-step-to-tackle-database-state-2209936.html>
- [27] Liberty Alliance Project, *Aetna Enhances Secure Provider Portal with SSO and SAML 2.0*, 2008.
- [28] Murdoch, Steven, Saar Drimer, Ross Anderson, and Mike Bond, “Chip and PIN is Broken”, *IEEE Symposium on Security and Privacy*, 2010, pp. 433–446.
- [29] National Strategy for Trusted Identities in Cyberspace, April 2011.
- [30] Nigriny, Jeff and Randy Sabett, “The Third-Party Assurance Model: A Legal Framework for Federated Identity Management”, *Jurimetrics Journal*, Vol. 50, No. 4 (2010) pp. 509–538.
- [31] OASIS, <http://www.oasis-open.org>
- [32] Rochet, Jean-Charles and Jean Tirole, “Platform Competition in Two-Sided Markets”, *Journal of the European Economic Association*, Vol. 1, No. 4, 990–1029 (2003)
- [33] Schwartz, Ari, “Identity Management and Privacy: A Rare Opportunity to Get it Right”, *Communications of the ACM*, to appear, April 2011.
- [34] Simonsen, David, personal communication to Susan Landau, April 28, 2011.
- [35] Tolstoy, Leo, *Anna Karenin*, Penguin Books, 1983.
- [36] United States Department of Commerce, *National Strategy for Trusted Identities in Cyberspace: Enhancing Online Choice, Efficiency, Security, and Privacy*, April 15, 2011.
- [37] University of Wisconsin, CIO and Vice President for Information Technology, “User Authentication and Levels of Assurance”, <http://www.cio.wisc.edu/security/initiatives/levels.aspx> [last viewed May 17, 2011].
- [38] Wandelt, John, “Global Federated Identity and Privilege Management”, August 2007.
- [39] Wilson, Yvonne, personal communication to Susan Landau, December 6, 2010.

- [40] United States General Accounting Office, Report to the Committee on Government Reform and the Subcommittee on Technology, Information Policy, Intergovernmental Relations and the Census, House of Representatives, *Electronic Government: Planned E-Authentication Gateway Faces Formidable Development Challenges*, GAO-03-952, September 2003.
- [41] WAYF Front Page, <http://wayf.dk/wayfweb/frontpage.html> [last viewed April 30, 2011].
- [42] WAYF Interfederation <http://wayf.dk/wayfweb/interfederation.html> [last viewed May 18 2011].
- [43] Yahoo!: “Yahoo! OpenID: Now with Attribute Exchange!”, December 4, 2009, http://developer.yahoo.com/blogs/ydn/posts/2009/12/yahoo_openid_now_with_attribute_exchange/ [last viewed February 23 2011].