

Briefing Document:

**Interoperability
In the New Digital Identity Infrastructure**

for the

Workshop on eInnovation and ICT Interoperability

19 & 20 January, 2007 -- Weissbad (Switzerland)

Introduction

This paper maps out multiple dimensions of interoperability in the emerging digital identity management infrastructure. It elaborates on a document that has been made available in wiki form to allow further iterations as the technology develops.¹

The emerging identity infrastructure involves many kinds of components, exchanging many types of information, over many possible interconnection paths. Interoperability here concerns the ability of any given component successfully to interact with any other. This paper attempts to define this multidimensional interoperability space with a view to promoting openness in the infrastructure.

The paper is therefore structured to include: A background to the new digital identity management infrastructure (identity infrastructure); an explanation of components, items exchanged, and interconnection flows; an account of interoperability in this infrastructure; a look at the costs and benefits of interoperability; and a brief consideration of factors that influence interoperability and innovation, such as the market, law, and self-regulation.

It should be noted at the outset that many of the components, items of exchange, and interconnection flows mentioned in this paper are under development. Developers in the industry are still grappling with how even to define interoperability in this space.

¹ The paper is based on the contents of the wiki document as of January 8, 2007, which Paul Trevithick authored except where noted. The original purpose of the document was to support an ongoing discussion on the topic of interoperability in the group called the Open Source Identity System (OSIS – for additional information, see http://osis.netmesh.org/wiki/Main_Page).

As such, this paper is a work in progress;² it is an effort both to understand and to offer insights into this emerging infrastructure.³

This analysis is offered in the hope of furthering the overarching goal of a user-centric identity infrastructure – that is, one that promotes human dignity and puts decision-making as close to the individual as possible.

Background to the Identity Infrastructure

In recent years increasingly sophisticated cyber attacks have caused growing numbers of people to fall victim to data loss, phishing, and pharming.⁴ With this spate of attacks, the public has grown leery of entrusting sensitive data to others.⁵

While people do not know how to safeguard data themselves, they nonetheless are cautious about the idea of entrusting one large entity with that data. Lessons have already been learned along this line: Although Microsoft expected that people would flock to it when it offered to help manage people's data for online interactions, those hopes backfired with Hailstorm and its successor Passport.⁶ As explained by Microsoft's chief identity architect, people simply do not want one large company to hold the keys to all their trust relationships, especially if that same company holds the keys to everyone else's relationships as well.⁷

To respond to today's authentication problems without these drawbacks, developers such as Microsoft are creating new user-centric solutions that are more open, less centralized in their design, more user-centered and offer a consistent user experience for all authentication interactions. Industry analysts anticipate widespread use of these new solutions, especially given that one example of such, Microsoft's new CardSpace™ system, could ultimately reach hundreds of millions of users (i.e. those running Windows Vista and those who download and install this capability on Windows XP). The expectation is that the infrastructure will facilitate transactions involving major e-commerce players like Amazon and eBay and micro-enterprises alike, as well as non-commercial interactions among countless participants.

² This paper will be posted on a comment-enabled web site to allow easy editing by the identity community and the most important actors in the identity infrastructure, namely, individual users.

³ The authors would like to thank John Clippinger for his highly valued feedback in the drafting process.

⁴ As described in an article by David Bank and Riva Richmond, "Information Security:

Where the Dangers Are," Wall Street Journal, July 18, 2005: "In 'phishing' scams, fraudsters send emails that appear to come from a trusted source, like Citibank or eBay. Click on a link in the email, and you're directed to a fake Web site, where you're asked to reveal account numbers, passwords and other private information... Then there's 'pharming,' where hackers attack the server computers where legitimate Web sites are housed. Type in the address of the legitimate site, and you are redirected to a look-alike."

⁵ Statistics have shown a sharp drop (42%) in the number of consumers who feel comfortable participating in e-commerce, as well as a large decline (28%) in the number of people who feel safe engaging in Internet banking. See Riva Richmond, "Internet Scams, Breaches Drive Buyers Off the Web, Survey Finds," Wall Street Journal, June 23, 2005, p. B3, reporting on a Gartner study of 5000 online consumers.

⁶ In the past several years the market has largely rejected Microsoft's "Passport" identity management system for cases in which Microsoft had no direct role as a party in the transactions.

⁷ See Kim Cameron's blog at <http://www.identityblog.com>. Kim has been one of the most effective proponents of a user-centric identity system and is credited with articulating "The Seven Laws of Identity."

Other developers are going a step further. Recognizing their interdependence in this regard, a group of industry developers in 2006 came together and created the Open Source Identity System (OSIS) initiative. They have been working together to build a infrastructure that will allow systems using different underlying network protocols interoperate – including, for example, those protocols used by Microsoft CardSpace, the Liberty Alliance, and OpenID – and to do so on any platform including Linux, Unix and OSX. Participants include representatives from Microsoft, IBM, Novell, VeriSign, Oracle, Red Hat, Higgins, CA, and Sun, as well as some notable independent, open-source developers. As noted on the OSIS wiki, “OSIS brings together many identity-related open-source projects, and synchronizes and harmonizes the construction of an interoperable identity layer for the internet from open-source parts. Its first deliverable is interoperability with Microsoft CardSpace’s protocols, although OSIS also encompasses alternate network protocols such as OpenID and SAML.”⁸

In order for the identity infrastructure to be accessible by all and to become a true “layer” of the Internet, a critical factor will be the interoperability of components and interconnection flows. For example: a person might delegate responsibility for overseeing transfers of his data to an “Identity Agent” such as Microsoft offers in its new CardSpace system; the person might run this CardSpace Identity Agent on a Windows Vista operating system; he might click on a CardSpace icon to send the most limited set of information possible to a “Service Provider” such as an e-commerce site to complete an online transaction. The Microsoft components that he and others like him use will gain value if compatible, open-source components are available, so that users with different operating systems supporting other Identity Agents (other than those of CardSpace) can exchange data with Service Providers; similarly, open-source technologies in this area benefit from work Microsoft has done to promote its particular Identity Agent (CardSpace) and its associated network “protocol” standards for transferring information and related infrastructural components. After all, an e-commerce site is much more likely to invest in the technology to participate in this infrastructure if the company knows it will be able to reach all users rather than just a subset. In other words, it is in the interest of all the actors to work together so that the products of each can enjoy a fully functioning infrastructure within which to perform.

With these new components, users will no longer need to memorize dozens of passwords and repeatedly type in payment information for online transactions; they will also be apprised of the trustworthiness of actors on the other end of a potential transaction; and they can enjoy these benefits without feeling they are sacrificing autonomy to one data giant.

⁸ See http://osis.netmesh.org/wiki/Main_Page (viewed January 10, 2007).

The Transfer of Identity Information in the CardSpace Model

The new identity infrastructure entails a sequence of steps to allow parties to interact online and to obtain reliable information about each other. In this interaction, a person is largely unaware of what is taking place technically in the various steps of an exchange of identity information, and instead the simple user interface spares him from having to deal with all the complexities of the underlying system. For the sake of illustration, we describe an authentication interaction using only the CardSpace sequence of steps.

So, for example, if a person wishes to purchase a bicycle online, he does not need to type in passwords or fill in an array of fields, but instead he simply chooses among electronic identity cards (“I-Cards”) and picks one that offers the type of information needed (e.g., payment and shipping details).

The example below outlines the sequence for an order placed online using the identity infrastructure. It should be noted that before any of the steps in the sequence can happen, the User must have acquired I-Cards from his various Identity Providers (flow A below); he must also have installed (imported) these cards into his Identity Agent that will act on his behalf (flow B below). The example also assumes that the person has authenticated himself directly to his Identity Agent.

Steps in this sample sequence, defined by Microsoft CardSpace and followed by compatible open-source projects such as Higgins and others, are as follows:

1. A certain user named Alberto uses the Firefox browser (or, rather, Firefox with an extension⁹) to go to the Best Buy web site. This site acts as the “Service Provider”. (See flow #1 in the diagram, below.)
2. Best Buy’s web page contains special HTML tags that are recognized by the Firefox extension as indicating that it is possible to sign-in using an I-Card, and that the site requires a certain set of information, or “claims” (e.g., name, email address, minimum age, etc. asserted to be true).

The Firefox extension reads Best Buy’s “policy” (i.e. what that Service Provider site requires in terms of claims and acceptable “token” types for secure packaging). (See flow #2, below.)

3. The Firefox extension conveys the site’s policy to Alberto’s Identity Agent and requests a token that conforms to this policy.

⁹ Some Identity Agents require an extension to operate on some computer platforms and/or browsers.

Alberto's Identity Agent then begins the "authentication" user experience. If this is the first time that Alberto has visited the Best Buy site, a page is displayed showing information about that Service Provider, including the site's level of security.

Alberto next sees a page pop up displaying his various I-Cards. Each I-Card represents a certain combination of data, or a claim. His collection of I-Cards might include, among others, one containing information from his driver's license and car insurance policy, another with his health-club membership information, and yet another with payment information and a shipping address. Unless Alberto's I-Cards were self-issued,¹⁰ they each have an associated Identity Provider (e.g., a bank, government agency, etc.) that Alberto has designated to fill in the actual data (the "data values").

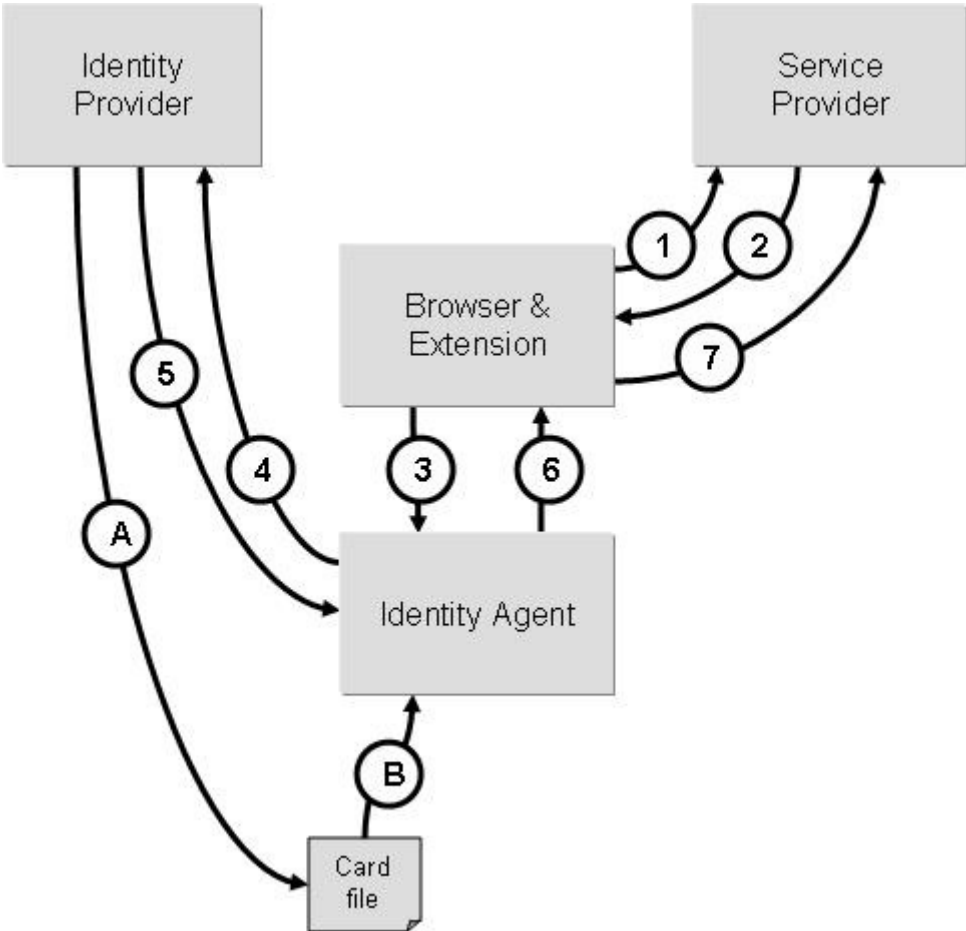
Alberto's Identity Agent searches his collection of I-Cards to find those whose claims would match what is required by Best Buy. It then grays out (disables) the I-Cards that do not have the required claims and displays only those cards that fit the bill.

Alberto selects the I-Card he wishes to use and clicks on it. He can also choose to push a button to preview the data elements associated with a card, and thereby review his name, age, current bank balance, etc. before releasing this information to a Service Provider like Best Buy.

4. When Alberto picks an I-Card and clicks on it, his Identity Agent sends a request over the Internet (flow #4 in the diagram below) to the I-Card's associated Identity Provider (in this case, the Bank of Canada), requesting it to provide the data values which Alberto has entrusted to it (e.g., "Albert" for first name, "over 18" for age, etc.).
5. The Bank of Canada as Identity Provider gathers the relevant data elements and wraps them in a cryptographically signed security token, which it then sends to Alberto's Identity Agent (flow #5 in the diagram below).
6. Alberto's Identity Agent sends the requested token to the Firefox extension (flow #6 below).
7. The Firefox extension sends the token to Best Buy (flow #7 below). Finally, Best Buy unwraps the token and takes out the information that is needed for the transaction.

¹⁰ For self-issued cards, this Identity Provider resides on a User's computer or device, although in most cases the Identity Provider will be a hosted service on the Internet.

The diagram below shows the relevant system components and interconnection flows in a transfer of identity information, numbered according to the steps just outlined.



Opportunities for Interoperability

For interoperability in this identity infrastructure, essentially the vision is that all components, interconnection flows, and items exchanged should work together. In other words: any Identity Agent should work with any Service Provider and any Identity Provider; any Identity Provider should work with any Service Provider and any Identity Agent; and any Service Provider should work with any Identity Agent and any Identity Provider. In addition, these components should each support the major protocols, claim types, and token types, plus they should communicate their policies in a shared language. Meanwhile, the user experience should be consistent throughout, independent of the underlying systems.

This section discusses opportunities for interoperability in the components, interconnection flows, and items exchanged in the identity infrastructure as currently con-

ceived. In addition, the section also offers some commentary on the role of standardization (*set off by italics*) in achieving these goals.

Components

Identity Agents in the CardSpace-Defined Flow: Known choices in Identity Agents that are compatible with the CardSpace-defined flow include:

1. Microsoft CardSpace - runs on Windows as a rich client application. (Status: in production – bundled within Windows Vista and available for WinXP as part of IE7 and NET 3.0.)
2. Higgins local Identity Agent - client runs on Linux, OSX or Windows as rich client application. (Status: under development.)
3. Higgins hosted Identity Agent. (Status: under development.)
4. OpenInfoCard - runs in Firefox on Linux, OSX or Windows. (Status: under development.)
5. Ian Brown's Identity Agent - runs in Safari on OSX. (Status: under development.)

Choices of Identity Agents (and required extensions, if any) can also be viewed in light of what various browsers will support. Popular browsers on different operating systems support the following Identity Agents:

1. Internet Explorer 7
On Windows: With no extension required uses the CardSpace Identity Agent.
2. Firefox
 - On Windows: With new Microsoft Extension, uses the CardSpace Identity Agent.
 - On Windows, Linux or OSX: With Higgins Browser Extension (HBX), uses a Higgins Identity Agent (Web or Client variant).
 - On Windows, Linux or OSX: With OpenInfoCard Extension, uses the OpenInfoCard Identity Agent.
3. Safari
On OSX: With Ian Brown's Extension, uses the Ian Brown Identity Agent.

In choosing an Identity Agent, a User may be influenced by the fact that the following types of Identity Agents are capable of importing and exporting the CardSpace I-Card format:

1. CardSpace Identity Agent
2. Higgins Identity Agent
3. OpenInfoCard Identity Agent
4. Ian Brown's Identity Agent

The Higgins Identity Agent is the only Identity Agent that currently exists for creating, importing, and exporting the Higgins I-Card format. However, as an open,

documented format, the Higgins I-card format may be processed by any Identity Agent.

Finally, the possibility for indirect transfers between Identity Providers and Service Providers via a user's trusted Identity Agent can prevent those parties from knowing who each other are and what history of information transfers exists between them. By keeping transactions unlinked, the system can help the user remain in control and enhance privacy. This separation is theoretically made possible in part¹¹ by the overall level of interoperability in the identity infrastructure and in part by cryptographic technology in the Identity Agent. Through standardization and interoperability, the infrastructure can better achieve its user-centric quality.

Identity Providers: Identity Providers are relatively new in the emerging identity infrastructure, and the business model is just developing. (Banks, credit card companies, and government agencies have traditionally played a similar role, storing personal information for use in various transactions; the idea here is to allow specialized provision of this service in a way that affords the user more control over the flow of data pertaining to him.)

Choices of Identity Providers that are CardSpace-compatible include:

1. Ping Identity Provider/STS
2. Microsoft Identity Provider/STS (Active Directory)
3. Higgins Token Service

Choices of non-CardSpace-compatible Identity Providers include:

1. SAML
2. OpenID Server
3. LDAP (Lightweight Directory Access Protocol)
4. Many others

Items Exchanged

Claim Types, I-Cards, and I-Card Formats: As indicated above, the Service Provider stipulates what types of claims it requires to conduct a transaction. Similarly, the Identity Provider indicates what claim types it is able to issue. (This should be no surprise to the Identity Agent since the Identity Provider has been chosen by the User in the first place for its services.)

There are different types of claims, each representing a different piece of data (e.g., surname, postal code, etc.). Instead of being able to tailor claims to the specific needs

¹¹ Anonymity in identity management demands attention at other layers of the communication stack as well. See Mary Rundle and Ben Laurie, "Identity Management as a Cybersecurity Case Study," (Better User Control over Personal Data), Berkman Center Publication Series, September 2005 (available at http://cyber.law.harvard.edu/home/uploads/521/2006_01_Rundle_IdentityManagement_CybersecurityCaseStudy.pdf).

of a transaction, the identity infrastructure currently enjoys prescribed selection of allowed data elements, in the form of a small set of generic, well-known claim types (e.g., surname, etc.). Current claim types may leave the Service Provider unable to request a specialized piece of information in any standardized way. While it is true that these systems (including CardSpace) do allow any Identity Provider and Service Provider to agree bilaterally on any claim types that they both find useful, such claims would only be understood by these two parties – other Service Providers would not understand them. For interoperability in claim types between Identity Providers and Service Providers that don't have a prior relationship, a broader, standardized set of claims are required. An open standardization process is required to define a broader set of claim types.

For managing claim types, there is an ongoing discussion in the OpenID community on how to manage claim types (or, in OpenID jargon, attribute types).

Meanwhile, with I-Cards, for the user to know what type of information he is passing to a Service Provider, he can use the visual representation of an information card, or I-Card. I-Cards are essentially data objects that are created and exported by Identity Providers; for self-asserted information they may also be created by the User using their Identity Agent. For I-Cards issued by Identity Providers, a User first acquires cards from his various Identity Providers and installs them into his Identity Agent.

The value of an I-Card is in its ability to connote something familiar to the User. Familiarity may rest in part on the range of claim types' not being too extensive – in other words, for there to be a certain standardization of them. However, just because the user may be better off with fewer, more familiar choices, does not mean there is not an important role for flexibility in claim types.

There are two choices for I-Card data formats¹² to choose from:

First, there is the CardSpace Format (XML), which has been defined by Microsoft for used in its CardSpace. This format is used in two somewhat different ways, i.e. for a personal profile (to describe personal cards, whose information a person asserts himself) and for a managed profile (to describe managed cards, whose information is vouched for by others).

Second, there is the Higgins Format (XML), which can be thought of as a superset of the CardSpace format. CardSpace format data can be converted without loss to the Higgins format.

¹² CardSpace "Card format" issues include the following: Should the card format be maintained by a proper standards body? Or should Microsoft maintain its version, and Higgins another (with Higgins's being a proper superset, for example)? What are the extensibility points of that format? Does CardSpace tolerate elements in the file that it does not understand?

As noted above, the Higgins Identity Agent is the only Identity Agent that currently exists for creating, importing, and exporting the Higgins format. However, as an open, documented format, the Higgins format may be processed by any Identity Agent.

It would seem that the infrastructure would improve if the production of these more detailed formats was not limited to a select group or technology. Developing such technologies in an open standards process enables more actors to contribute to building a thriving system.

Interoperability in Claim Schemas: Different identity systems may use different terminology for referring to the same data elements that go into claims. For example, a value classified as “last name” in one system may be classified as “surname” in another. The Identity Commons working group, IdentitySchemas.org, is focused on claim interoperability.¹³

Tokens: Prior to the transfer of a claim, the Service Provider indicates what kind of token it will accept. There are two main groups of token types, only the first of which are formal standards namely:

1. OASIS WSS Token Profiles
 - SAML 1.1
 - SAML 2.0
 - Kerberos v5
 - UN/PW
 - X509 v3
 - Idemix

2. OpenID Token Types
 - OpenID 1.1
 - OpenID 2.0
 - LID 2.0

When it comes to this token acceptability, it would seem that a Service Provider should be willing and able to accept a range of token types, so long as they succeed in transferring data securely. All parties should be equipped to deal with all the major types of tokens. Standardization efforts especially at OASIS mentioned above have already done much good in this area—there are several standard kinds of tokens. The newer identity protocols allow transport of any kind of token as well and this is a major step forward for interoperability.

¹³ See http://idschemas.idcommons.net/moin.cgi/List_Of_Schemas (viewed January 10, 2007).

Interconnection Flows

Service Provider's Policy Expression: In interacting with the User's browser or that of the Identity Provider, the Service Provider uses a particular format to express its policy—what information it requires, what data formats (e.g. tokens) it accepts and what protocol it supports. Choices include:

1. CardSpace-compatible HTML object tag
2. CardSpace-compatible XHTML informationCard tag
3. OpenID 1.1
4. OpenID 2.0
5. SAML 2.0
6. RSS+SSE
7. HTML Scraping (including microformats)
8. WS-Federation

Regarding protocols and claims transfers, it would not make sense for a transfer of claims to be prevented just because the parties could not agree on a protocol for interconnection. Protocol standardization makes sense, and the less that decisions about the appropriate protocol factor into claims transfer, the smoother the system.

With particular reference to the way that tokens pass between an Identity Agent and a browser/extension, it would be logical for all parties to be equipped with the technology for a smooth hand-off of tokens between one endpoint and another. Such standardization, again, would help keep the user at the center of his identity management since there would not be exclusionary languages and systems developed among select parties.

Even if policies were different among the distinct parties, having a common language would make it easier to get things done. Standardization is thus important in the area of policies as well, especially as policies can either facilitate or obstruct the flow of information via innovative methods.

Protocols for interaction between the Identity Agent and the Identity Provider: To request a claim data, the Identity Agent uses a protocol to request it from the Identity Provider. Choices for the protocol used here include the following:

1. WS* (WS-Trust, WS-MetadataExchange, etc.)
2. OpenID 1.1
3. OpenID 2.0
4. Liberty
5. SAML 2.0
6. LDAP

An Identity Agent should be able to support these different network protocols. Again, standardization in the protocol allows innovation in the technologies that flow over it.

Claim Type Namespaces: The claim types used in card formats are identified by globally unique identifiers or, in this case, uniform resource identifiers (URIs; a URI is a string of characters that names a resource, or tells its whereabouts, to enable interaction with a representation of it over a network¹⁴). Although there is no requirement to do so, in practice the well known claim types share a common root URI, which here is called the “claim type namespace”.

At present there is only one choice for a well-known claim type namespace – namely, that which Microsoft has currently defined within the claim types of CardSpace.¹⁵

In terms of additional choices here, other well known sets of claim types *could* be mapped into claim type URIs if the industry would agree on a standard for mapping them to URIs. Examples include the Liberty People Service and other Liberty specifications. Over time the industry will likely converge on other well-known claim types, and there will need to be an open process for standardizing these. At present there is no such process.

General Effects of Standardization

The account above addressed some of the benefits that standards might bring for interoperability in specific dimensions of the emerging identity infrastructure. More generally, open standards processes can help guard against the tendency of competitors to limit choice and favor the adoption of their own “standards” in a way that gives them leverage in the market for services. The fear is that an advantage in one dimension – e.g., an operating system – could allow a major industry player to secure an advantage in another dimension – e.g., Identity Agents.

As noted at the outset, all players have an interest in seeing an interoperable infrastructure evolve. Nonetheless, there will always be competitive economic pressures that will cause players to want to drive each other out of the market.

In cultivating an environment that is favorable to interoperability, policymakers will need to consider what the market already favors, given technologies that are already deployed and have a good foothold. For example, CardSpace enjoys a significant advantage in that has been designed to work seamlessly with Microsoft Windows. Given the prevalence of Windows and other factors that will affect the market for

¹⁴ See the entry for URI in Wikipedia, at http://en.wikipedia.org/wiki/Uniform_Resource_Identifier. The article notes that a “Uniform Resource Locator (URL) is a URI that, in addition to identifying a resource, provides means of acting upon or obtaining a representation of the resource by describing its primary access mechanism or network ‘location’” (as viewed on January 16, 2007).

¹⁵ All of the 14 defined CardSpace claim type URIs share the common root URI. See <http://schemas.xmlsoap.org/ws/2005/05/identity/claims/> (viewed January 10, 2007).

identity management services, it is worth asking whether it would be beneficial to the Information Society's identity infrastructure wholeheartedly to embrace CardSpace. On the one hand, CardSpace could offer the best hope of achieving critical mass and network effects; on the other hand, it is possible that more room needs to be made for other technologies in order to foster the kind of competition that brings flourishing innovation.

So many questions are yet to be raised. What would really happen if the market were to default to a single protocol, for example WS*? Would this shift lock in bigger players and lock out smaller ones, or would this choice in protocol not really matter for openness? Do there need to be alternatives in every dimension of the identity infrastructure? Could a "metasystem" such as that envisioned by Microsoft¹⁶ or a protocol- and platform-agnostic platform like that proposed by Higgins¹⁷ cause these problems to dissolve?

Given the potentially critical role of the identity infrastructure for a person's interactions in the Information Society, the interrelationships of interoperability, standardization, and innovation deserve vast exploration. While issues raised here represent only a start to this endeavor, they perhaps can alert policymakers to some of the dynamics triggered by emerging identity management technologies.

The Need for Debate

Before religiously embracing the concept of an interoperable identity infrastructure, it would be wise to ask if there are costs to an infrastructure where interoperability pervades. To spur initial reflection, some ideas are set out below.

With respect to a single, dominant protocol, one argument often advanced in favor of the idea is that it would help create a critical mass and allow the identity infrastructure to take off as a new layer of the Internet. The critical mass would bring significant network effects, where the infrastructure would enjoy increased value with each additional participant. Similarly, an identity infrastructure with a standardized protocol would presumably allow for terrific innovation, helping to ensure the entire system's flexibility and relevance over time. New technologies could easily be plugged in and work throughout the infrastructure, similar to the way that the simple TCP/IP protocol for the Internet has allowed endless innovation.

Of course, the same could be argued for interoperable interconnection flows.

Another argument aimed at supporting the homogeneity cause might relate to the very purpose of the identity infrastructure – that is, the degree to which it promotes

¹⁶ See Kim Cameron's blog entry at http://www.identityblog.com/?page_id=355.

¹⁷ See <http://www.eclipse.org/higgins/higgins-charter.php>.

privacy and data security. Proponents of homogeneity could contend that a single protocol should be dictated. According to this line of thought, the fault with the multi-protocol approach is that the more protocols that the infrastructure tries to support, the more vulnerable it is to attack. A secure system thus would require strong security throughout the whole chain. If security were weak at any point, the entire system could break down. Since additional computer code and complexity would make for more links in the chain, any one of which could then be the weakest link, the infrastructure would arguably be more robust with one simple protocol. Therefore it would be better to concentrate resources in making all the links in that single chain strong.

Of course, one could take the opposite stance and argue that for security reasons, the identity infrastructure should support multiple, competing protocols for interconnection flows as this would allow multiple lines of defense should one protocol be attacked. Put another way, interoperability leads to increased security since it results in reinforcements.

With respect to user interface, critics of interoperability might say that a chaotic market confuses users and does not allow for a consistent practice to emerge. By this same logic, a standard, homogeneous system would not only comfort disoriented users, but would also bring the security benefit of helping people recognize fraudsters when they encountered them. According to this logic, then, a well orchestrated identity infrastructure would help put the Internet back on track for continued innovation in e-commerce and online interactions generally.

Here again, champions on the other side would counter that interoperable systems afford flexibility and openness, and so spur the type of competition that leads to increasingly better products' being offered to the user.

These arguments are all marked by hyperbole. They do, however, point to the need for true debate on these subjects.

With regard to the specific issue of whether there should be one, standard protocol for interconnection flows, it would seem that a wider community should be debating the merits of this idea. Any regulatory stance that favors a monopoly should be scrutinized, whether the policy would be a result of a government rule or industry self-regulation. It would make sense for experts in competition policy, particularly those specialized in network effects, to contribute to the assessment. Law could play an important role here, either by creating a favorable climate for a natural monopoly, or by ensuring that competition remains in protocols.

Conclusion

This paper has highlighted some of the tensions regarding interoperability in the emerging infrastructure for digital identity management. Because the identity infrastructure promises to have a major impact on the Information Society, and because interoperability will play a major role in shaping this infrastructure, policymakers must devote urgent attention to the subject.

While it is in everyone's interest to construct an open system where new components can be added, there will nonetheless be a tendency for pockets to develop in a way that drives out others in balkanization efforts. The infrastructure thus runs the risk of losing its best qualities: that is, a welcome environment for innovation, flexibility on the part of participants, and the potential for wide adoption of improvements.

As this initial exploration has suggested, there should be an ethos of interoperability pervading the new identity infrastructure. Can those with the greatest capacity to help create the commons be convinced of the good faith of others, especially as they put aside short-term proprietary interests to build this ecology? Who is to maintain interoperability in practice? Who is to enforce it? With nobody designated as the keeper of the commons, and with no strong ethos that people adhere to, interoperability will collapse, and innovation will no longer thrive. In other words, with no agreement to preserve the commons, the parties to the infrastructure will atomize it, and the joint investors in it will suffer loss.

At their very heart, these issues pose a problem of the tragedy of the commons. If there could be a common commitment to keep the infrastructure as diverse, distributed, and open as possible, users as a whole would achieve a much higher optimization point. If the commitment could be one that was naturally calibrated by the market, this would be even better for keeping the identity infrastructure open. Perhaps what is needed is a set of agreements that the infrastructure is a public good and should be maintained as such, with these agreements building in natural incentives for parties to resist compromising interoperability for proprietary advantage.

Naturally, in considering possible approaches for areas where interoperability is deemed desirable, policymakers could usefully consider those points where it is relatively straightforward to achieve interoperability, as opposed to those where this task is more difficult.

As suggested, one tool for affecting interoperability is the law. In addition to the law's ability to influence market through anti-trust policy, it also has potential to influence through other means the uptake of the identity infrastructure as a whole. For example, the law can spell out liability for data loss, know-your-customer require-

ments, and other rules that may cause market participants to begin to rely on new identity management technologies.

Similarly, certain basic principles should apply in industry self-regulation. By way of example, although the identity infrastructure may already have its major contenders¹⁸ for a standard protocol, generally speaking, such a process best serves the greater good if it is open for wide participation (with the cost of participation being low), and if it allows for rigorous review in which people may comment on proposed variations. In this sense, interoperability and the innovation it affords are closely related to a commitment to the open development of standards.

While the means are debatable, one thing is clear: The Information Society has everything to gain from an identity infrastructure that is privacy-enhancing and that puts decision-making as close to the end-user as possible.

(Contact: Mary Rundle, Berkman Center for Internet & Society, Harvard Law School, mrundle@cyber.law.harvard or +1.617.495.7547; Paul Trevithick, ptrevithick@alum.mit.edu)

¹⁸ The two most notable are the “WS-*” web services stack developed in the OASIS standardization body (largely as a collaborative effort between Microsoft and IBM) and that produced by the Liberty Alliance.