Design Rationale behind the Identity Metasystem Architecture

Kim Cameron and Michael B. Jones

kcameron@microsoft.com, mbj@microsoft.com
http://www.identityblog.com/, http://research.microsoft.com/~mbj/

Abstract

Many of the problems facing the Internet today stem from the lack of a widely deployed, easily understood, secure identity solution. Microsoft's "InfoCard" project and the Identity Metasystem vision underlying it are aimed at filling this gap using technology all can adopt and solutions all can endorse, putting users in control of their identity interactions on the Internet.

The design decisions presented in this paper are intended to result in a widely accepted, broadly applicable, inclusive, comprehensible, privacy-enhancing, security-enhancing identity solution for the Internet. We present them and the rationale behind them to facilitate review of these design decisions by the security, privacy, and policy communities, so that people will better understand Microsoft's implementations, and to help guide others when building interoperating implementations.

1. Introduction

1.1. The Challenge: A Ubiquitous Digital Identity Solution for the Internet

By definition, for a digital identity solution to be successful, it needs to be understood in all the contexts where you might want to use it to identify yourself. Identity systems are about identifying *yourself* (and your things) in environments that are *not yours*. For this to be possible, both *your* systems and the systems that are *not yours* – those where you need to digitally identity yourself – must be able to speak the same digital identity protocols, even if they are running different software on different platforms.

In the case of an identity solution for the entire Internet, this is a tall order. It means that, to succeed, the solution will need to be adopted by the wide variety of operating systems, browsers, and web servers that collectively implement the phenomenon we know of as "the Internet".

1.2. Practical Considerations

To have any hope of such widespread adoption, we believe that any Internet-scale identity solution will need to satisfy these practical considerations:

- **Improved Security and Privacy:** To be widely adopted, platform and software vendors will need to be convinced that the solution results in improvements in the overall Internet security landscape. Likewise, consumers (and their advocates) will need to be convinced that the solution improves the consumer privacy landscape.
- **Inclusive of Technologies:** There are a number of identity technologies in widespread use today (Kerberos, X.509, SAML, etc.) with more being invented all the time. To gain wide acceptance, the solution should be able to leverage existing identity technologies and deployments, incorporating them as part of the solution and building upon their strengths, rather than calling for their wholesale replacement.
- **Inclusive of Scenarios:** The solution must be broadly applicable across a wide range of use cases, even accommodating those with conflicting requirements. For instance, in many cases users will want guarantees that their identity providers can't accumulate a record of the sites they visit. However, in some governmental and financial settings, an audit record of sites visited using an identity may be required. Both kinds of identities should be able to be accommodated. At an even more basic level, the solution must be applicable not just on workstations but also on different devices such as wireless mobile devices and cell phones.
- **Incrementally Deployable:** The solution must coexist with and complement existing authentication systems, rather than calling for a "forklift upgrade" or "flag day" where ex-

isting solutions must be replaced by the new one all at once.

1.3. Architecture of a Proposed Solution

Such a solution, the ***Identity Metasystem*** [Microsoft 05a], has been proposed and some implementations are under way. The Identity Metasystem is based upon a set of principles called the "***Laws of Identity***" [Cameron 05b]. The Laws are summarized in Appendix A. The Laws are intended to codify a set of fundamental principles to which a universally adopted, sustainable identity architecture must conform. The Laws were proposed, debated, and refined through a long-running, open, and continuing dialogue on the Internet [Cameron 05a]. Taken together, the Laws were key to defining the overall architecture of the Identity Metasystem.

While the Laws of Identity have undergone broad review and been met with significant acceptance, that's certainly not the end of the story. While the Identity Metasystem is designed in accordance with the Laws, there are also numerous practical design decisions that had to be made to translate the vision into working, interoperable systems.

The purpose of this paper is to publish the design decisions underlying the Identity Metasystem architecture and the rationale behind them. This is intended both to enable a deeper understanding of the problems that this solution addresses and to enable discussion of these design decisions by the security, privacy, and policy communities.

2. Identity Problems on the Internet and an Overview of the Proposed Solution

The section briefly describes the problems motivating the need for a new identity solution for the Internet and gives an overview of the mechanisms that the Identity Metasystem employs to do so.

2.1. The Internet's Problems are often Identity Problems

Many of the problems facing the Internet today stem from the lack of a widely deployed, easily understood, secure identity solution. Microsoft's "InfoCard" project and the Identity Metasystem vision underlying it are aimed at filling this gap using technology all can adopt and solutions all can endorse, putting users in control of their identity interactions on the Internet.

A comparison between the brick-and-mortar world and the online world is illustrative: In the brick-and-mortar world you can tell when you are at a branch of your bank. It would be very difficult to set up a fake bank branch and convince people to do transactions there. But in today's online world it's trivial to set up a fake banking site (or e-commerce site …) and convince a significant portion of the population that it's the real thing. *This is an identity problem.* Web sites currently don't have reliable ways of *identifying* themselves to people, enabling imposters to flourish. One goal of InfoCard is reliable site-to-user authentication, which aims to make it as difficult to produce counterfeit services on the online world as it is to produce them in the physical world.

Conversely, problems identifying users to sites also abound. Username/password authentication is the prevailing paradigm, but its weaknesses are all too evident on today's Internet. Password reuse, insecure passwords, and poor password management practices open a world of attacks by themselves. Combine that with the password theft attacks enabled by counterfeit web sites and man-in-the-middle attacks and today's Internet is an attacker's paradise.

The consequences of these problems are severe and growing. Last year the number of "phishing" sites was growing at over 1000% per year [Anti-Phishing 05]. Online banking activity is declining [Gartner 05]. The recent FFIEC guidance on authentication in online banking reports that "Account fraud and identity theft are frequently the result of single-factor (e.g., ID/password) authentication exploitation" [FFIEC 05]. Consumer trust of the Internet is low and dropping. The status quo is no longer a viable option.

2.2. "InfoCard" and the Identity Metasystem

The code-named "InfoCard" project at Microsoft is a joint effort with a diverse coalition of contributors across the computer industry to produce an authentication solution for the Internet that can:

- be widely accepted,
- work in a broad range of identity contexts,

- utilize existing authentication technologies, including multiple factors,
- incorporate new authentication technologies as they are invented,

and possibly most importantly,

- enable users to simply and consistently make informed and positive authentication decisions on their own behalf.

The result of this effort is known as the Identity Metasystem [Microsoft 05a], an overview of which is contained in this section. As previously mentioned, the Identity Metasystem is based upon a set of principles developed through an open industry dialog [Cameron 05a] called the Laws of Identity [Cameron 05b].

What do we mean by an "Identity Metasystem"? This concept is probably most easily introduced through an analogy.

Before 1982, the networking world was fragmented. If you wanted to write a network-enabled application you had to choose what network to write it for: Ethernet, Token Ring, ArcNet, X.25, etc. The invention of a Network Metasystem, the Internet Protocol (IP), changed all that. It made it possible to write networking applications that worked across networks without knowing the particulars of each network. It even enabled those applications to work with new networks that hadn't been invented yet, such as 802.11 wireless networks.

Digital identity is similarly fragmented today. If you want to write an identity-enabled application, you have to choose which identity system to write it for, such as Kerberos, SAML, X.509, Liberty, custom username/password systems, etc. The Identity Metasystem is intended change all that, just as IP did for networking. It will make it possible to write identity-enabled applications that can work across multiple identity systems and can even use new identity systems as they are invented and connected to the Identity Metasystem.

This analogy holds true in another way. IP didn't compete with or replace the individual networks such as Ethernet — it *used them*. Similarly, the Identity Metasystem doesn't compete with or replace individual identity technologies such as Kerberos, Liberty, X.509, SAML, etc. — it uses them. That's why it's called an identity ***metasystem*** —it's a system of systems, tying individual identity systems into a larger interoperable metasystem (see Law 5).

By allowing different identity systems to work in concert, with a single user experience, and a unified programming paradigm, the metasystem shields users and developers from concern about the evolution and market dominance of specific underlying systems, reducing everyone's risk and increasing the speed with which technology can evolve.

2.3. Roles within the Identity Metasystem

Different parties participate in the metasystem in different ways. The three roles within the metasystem are:

- ***Identity Providers***, which issue digital identities. For example, credit card providers might issue identities enabling payment, businesses might issue identities to their customers, governments might issue identities to citizens, and individuals might use self-issued identities in contexts like signing on to web sites.
- ***Relying Parties***, which require identities. For example, a web site or online service that utilizes identities offered by other parties.
- ***Subjects***, which are the individuals and other entities about whom claims are made. Examples of subjects include people, companies, and organizations.

2.4. Claims-Based Identities and InfoCards

In the Metasystem, digital identities consist of sets of claims made about the subject of the identity, where "claims" are pieces of information about the subject that the issuer asserts are valid. This parallels identities used in the real world. For example, the claims on a driver's license might include the issuing state, the driver's license number, name, address, sex, birth date, organ donor status, signature, and photograph, the types of vehicles the subject is eligible to drive, and restrictions on driving rights. The issuing state asserts that these claims are valid. The claims on a credit card might include the issuer's identity, the subject's name, the account number, the expiration date, the validation code, and a signature. The card issuer asserts that these claims are valid. The claims on a self-issued identity, where the identity provider and subject

are one and the same entity, might include the subject's name, address, telephone number, and e-mail address, or perhaps just the knowledge of a secret. For self-issued identities, the subject asserts that these claims are valid.

In the client user interface, each of the user's digital identities used within the metasystem is represented by a visual "Information Card" (a.k.a. "InfoCard", the source of this technology's codename). The user selects identities represented by InfoCards to authenticate to participating services. The cards themselves represent references to identity providers that are contacted to produce the needed claim data for an identity when requested, rather than claims data stored on the local machine. Only the claim values actually requested by the relying party are released, rather than all claims that the identity possesses (see Law 2).

2.5. Putting the User in Control

One of the fundamental tenets of the InfoCard work is that users must be in control of their identity interactions (see Laws 1 & 2). Among other things, this means that users must be given the choice of which identities to use at which services, they must know what information (which claims) will be disclosed to those services if they use them, and they must be informed how those services will use the information disclosed.

In the offline world, people carry multiple forms of identification in their wallets, such as driver's licenses or other government-issued identity cards, credit cards, and affinity cards such as frequent flyer cards. People control which card to use and how much information to reveal in any given situation.

Similarly, the Identity Metasystem makes it easier for users to stay safe and in control when accessing resources on the Internet. It lets users select from among a portfolio of their digital identities and use them at Internet services of their choice where they are accepted. The metasystem enables identities provided by one identity system technology to be used within systems based on different technologies, provided an intermediary exists that understands both technologies and is willing and trusted to do the needed translations.

Part of being in control that's all too often overlooked is that to be in control, you must be able to understand the choices you're presented with (see Laws 6 & 7). Unless we can bring users into the identity solution as informed, functioning components of the solution, able to consistently make good choices on their own behalf, we won't have solved the problem.

Many identity attacks succeed because the user was fooled by something presented on the screen, not because of insecure communication technologies. For example, phishing attacks occur not in the secured channel between web servers and browsers — a channel that might extend thousands of miles — but in the two or three feet between the browser and the human who uses it. The Identity Metasystem, therefore, seeks to empower users to make informed and reasonable identity decisions by enabling the use of a consistent, comprehensible, and self-explanatory user interface for making those choices.

One key to securing the whole system is presenting an easy-to-learn, predictable user interface that looks and works the same no matter which underlying identity technologies are employed. Another key is making important information obvious — for instance, displaying the identity of the site you're authenticating to in a way that makes spoofing attempts apparent. Likewise, the user must be clearly informed which items of personal information relying parties are requesting, and for what purposes. This allows users to make informed choices about whether or not to disclose this information.

2.6. Authenticating Sites to Users

To prevent users from being fooled by counterfeit sites, there must be a reliable mechanism enabling them to distinguish between genuine sites and imposters. Our solution utilizes a new class of higher-value X.509 site certificates being developed jointly with VeriSign and other leading certificate authorities. These higher-value certificates differ from existing SSL certificates in several respects.

First, these certificates contain a digitally-signed bitmap of the company logo. This bitmap is displayed when the user is asked whether or not they want to enter into a relationship with the

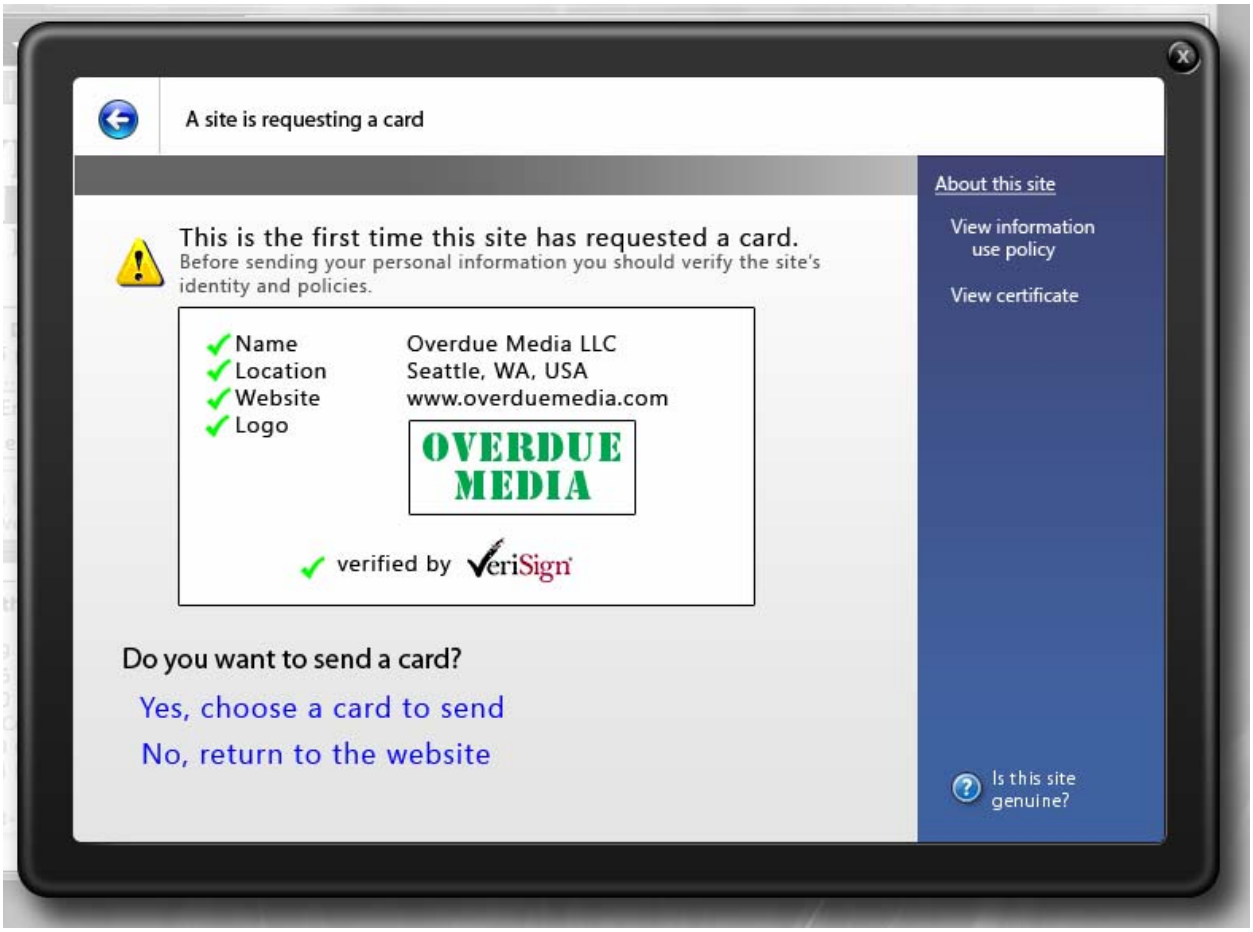

Figure 1: Site Verification Screen

site, the first time that the site requests an InfoCard from the user.

Second, these certificates represent higher legal and fiduciary guarantees than standard certificates. In many cases, all that having a standard site certificate guarantees is that someone was once able to respond to e-mail sent to that site. In contrast, a higher-value certificate is the certificate authority saying, in effect, "We stake our reputation on the fact that this is a reputable merchant and they are who they claim to be".

Users can visit sites displaying these certificates with confidence and will be clearly warned when a site does not present a certificate of this caliber. Only after a site successfully authenticates itself to a user is the user asked to authenticate himself or herself to the site.

To make this all more concrete, Figure 1 shows an example of what a screen displayed upon a user's first access to a relying party accepting "InfoCards" might look like. As this example shows, the screen can include the name, location, web site URL, and logo of the organization whose identity is being approved (such as Overdue Media). It can also include the name and logo of the organization that has verified this information (such as VeriSign).

To help the user make good decisions, what's shown on the screen varies depending on what kind of certificate is provided by the identity provider or relying party. If a higher-assurance certificate is provided, the screen can indicate that the organization's name, location, website, and logo have been verified, as shown in Figure 1. This indicates to a user that this organization deserves more trust. If only an SSL certificate is provided, the screen would indicate that a lower level of trust is warranted. And if an even weaker certificate or no certificate at all is provided, the screen would indicate that there's no evidence

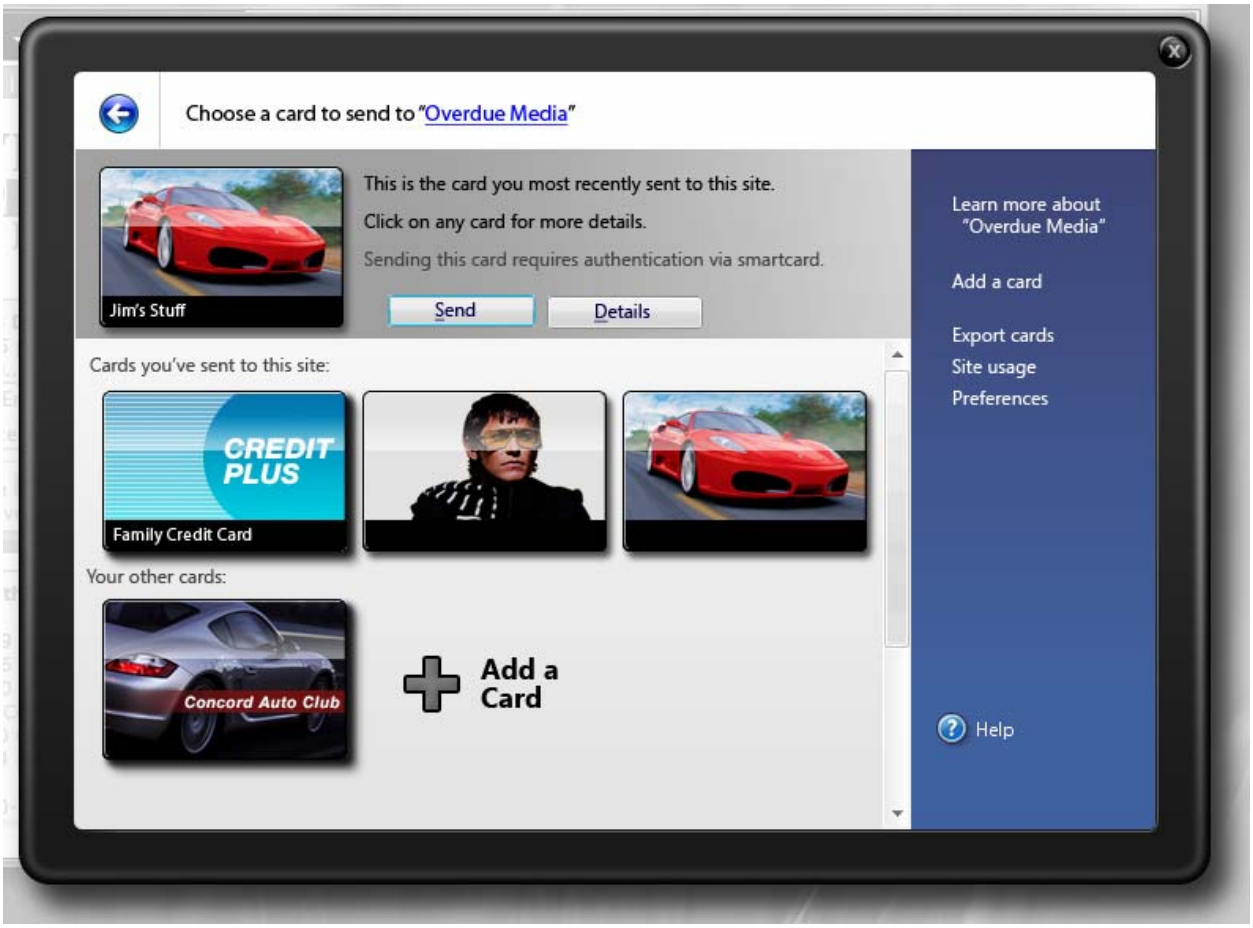

Figure 2: Identity Selector Screen

whatsoever that this site actually is who it claims to be. The goal is to help users make good decisions about which identity providers they'll let provide them with digital identities and which relying parties are allowed to receive those digital identities.

2.7. Authenticating Users to Sites

InfoCards have several key advantages over username/password credentials:

- Because no password is typed or sent, by definition, your password can not be stolen or forgotten.
- Because authentication is based on unique keys generated for every InfoCard/site pair (unless using a card explicitly designed to enable cross-site collaboration), the keys known by one site are useless for authentication at another, even for the same InfoCard.
- Because InfoCards will resupply claim values (for example, name, address, and e-mail address) to relying parties that the user had previously furnished them to, relying parties do not need to store this data between sessions. Retaining less data means that sites have fewer vulnerabilities. (See Law 2.)

InfoCard implements a standard user interface for working with digital identities. Perhaps the most important part of this interface, the screen used to select an identity to present to a site, is shown in Figure 2.

As this screen shot illustrates, each digital identity is displayed as an InfoCard. Each card represents a digital identity that the user can potentially present to a relying party. Along with the visual representation shown above, each card also contains information about a particular digital identity. This information includes what identity provider to contact to acquire a security token for this identity, what kind of tokens this identity provider can issue, and exactly what

claims these tokens can contain. By choosing to use a particular card, the user is actually choosing to request a specific security token with a specific set of claims created by a specific identity provider. But from the user's perspective, they're simply selecting an InfoCard to use at a site.

2.8. Protocols Behind the Identity Metasystem

The Identity Metasystem is built on a small number of interoperable Web Services (WS-*) protocols. Specifically, the encapsulating protocol used for claims transformation within the Metasystem is WS-Trust [WS-Trust 05]. Format and claims negotiations between participants are conducted using WS-MetadataExchange [WS-MetadataExchange 04] and WS-SecurityPolicy [WS-SecurityPolicy 05]. Finally, messages are secured using WS-Security [WS-Security 04].

These protocols enable building a platform-independent Identity Metasystem and form its "backplane". Like other Web services protocols, they also allow new kinds of identities and technologies to be incorporated and utilized as they are developed and adopted by the industry.

To foster the interoperability necessary for broad adoption, the specifications for these (and other) WS-* protocols are published and are freely available, have been or will be submitted to open standards bodies, and allow implementations to be developed royalty-free.

Deployments of existing identity technologies can be leveraged in the metasystem by implementing support for the small number of WS-* protocols above. Examples of technologies that could be utilized via the metasystem include LDAP claims schemas; X.509, which is used in Smartcards; Kerberos, which is used in Active Directory and some UNIX environments; and SAML, a standard used in inter-corporate federation scenarios.

3. Design Decisions behind the Identity Metasystem

This section lists many of the key design decisions behind the Identity Metasystem architecture and gives the rationale for them.

3.1. Protocol ≠ Payload

There are a number of forms of digital identity in use today such as Kerberos, X.509, SAML, and username/password systems, with more being invented all the time. Each typically represents identities in a different manner, and yet it is highly desirable to be able to utilize all these kinds of identities within the same identity solution. While some identity systems have developed custom communication protocols tied to particular identity formats, doing so results in little or no interoperability between the different systems using those incompatible protocols.

Instead, we decided to employ a single encapsulating protocol set capable of utilizing all identity payload formats in a common manner. Specifically, the protocol set was chosen to enable specification of requirements, negotiation of capabilities, transmission of payloads, and transformation of payloads, all in a format independent manner. This means that the encapsulating protocol remains stable even as the types of payloads used evolve.

3.2. Identity Selector ≠ Identity Provider

The Identity Metasystem employs software on each platform that lets users choose an identity from among their portfolio of identities to use for each authentication. This software is called the ***Identity Selector***, and is invoked each time the user needs to make a choice of identities. (Figure 2 shows a screen shot of an Identity Selector.) A key decision was to implement an Identity Selector that is independent of any specific identity provider, technology, or operator.

This enables an open architecture in which multiple identity providers using potentially multiple different identity technologies can all participate, with the user experience being the same each time. This open architecture allows both existing identity technologies and those yet to be invented to be used in the same ways.

Because identity providers are not tied to the identity selector but are instead communicated with by the selector using standard protocols, the identity providers can live anywhere: "in the cloud", at ISPs, on devices such as smart cards or USB keys, media players, cell phones, or on your PC, … anywhere reachable via the identity provider protocols. A corollary of this decision is that the simple self-issued identity provider that runs on your PC "out of the box" is just one among many and not "special" in any way.

3.3. Identity Selector ≠ Metadata Store

The identity selector software allows users to choose from among the identities in their portfolio of identities. This portfolio is represented by what we call the Metadata Store – the store of configuration info telling the identity selector how to contact an identity provider to obtain actual identity information. This metadata store also contains the pictorial representation of each identity, each "InfoCard".

We made the design decision to have the identity selector user interface software be separate from the metadata store software, with communication protocols connecting them. This decision provides significant flexibility that would otherwise not exist.

Specifically, it means that the identity selector user interface can run anywhere – not just on your workstation, but also on devices like your mobile phone or your media player. It also means that the metadata store can live wherever you want, for instance, on your phone, your media player, in the cloud, on a smart card or USB device supporting roaming, or on your PC. All of this contributes to giving the user control over how their identity is represented, stored, and released.

3.4. Guarantee Separation of Contexts

Many relying parties need a consistent handle to be presented each time an identity is used so they know that each use represents the same entity. But if this same handle is used at different relying parties, that gives them the opportunity to share data between them about how the same user has been using the different sites – all without the knowledge or agreement of the user.

A design decision was to mitigate this danger by supporting the use of "unidirectional identifiers" (see Law 4) so that the identifiers given to each relying party can be distinct from the identifiers given to others. The system is able to automatically generate pairwise identifiers for each combination of identity provider and relying party that is used. No common URL, GUID, etc. is sent that could serve as a correlation handle between sites.

Another way in which separation of contexts is facilitated is by ensuring that only those claims explicitly requested by a relying party are provided to it (as per Law 2). So, for instance, even though an identity provider might be capable of furnishing claims containing a subject's postal address and telephone numbers, unless they are requested the identity provider will not supply them to that relying party. Thus, the set of claims released varies on a per relying party basis.

3.5. Facilitate "Data Rejection"

Currently most sites retain a dossier of information about you: your "Customer Record". In the metasystem, a design decision was to have the selector remember what the user has released to a given site, and resupply that same information to the site whenever it requests it. The result of this decision is that sites can safely discard this information about you between sessions because it will be resupplied when next needed. Besides having privacy benefits for users, this option also has liability benefits for relying parties: Information that is not retained can not be stolen, meaning there cannot be data breaches for which a site can be held accountable.

3.6. Claims ≠ "Trust"

A design decision was to factor out trust decisions and not bundle them into the identity metasystem protocols and payloads. Unlike the X.509 PKIX [IETF 05], for example, the metasystem design verifies the cryptography but leaves trust analysis for a higher layer that runs on top of the identity metasystem.

3.7. Human Token ≠ Computational Token

For a human user to meaningfully control the information that would be released by selecting an identity, he or she must be able to view a human-readable and comprehensible representation of those claims. Hence, the identity selector must be able to display representations of claim values. However, because claims can be represented using any payload format, including new ones yet to be invented, it would be impossible to write identity selector code to meaningfully display claim values based only upon the payload's native representation of those claim values (unless we implemented potentially dangerous extension mechanisms, significantly increasing the vulnerability of the system).

Therefore a design decision was to have identity providers send claim values both in their

native format and in a human-readable format (the "display token"), with the two sets of values cryptographically bound together to allow auditing of an identity provider either by users or by relying parties that understand the claims.

3.8. Auditing ≠ Non-auditing Identity Providers

In many cases users will want guarantees that their identity providers can't accumulate a record of the sites they visit. Yet in some governmental and financial settings, an audit record of sites visited using an identity is absolutely required; both kinds of identities should be able to be accommodated. A design decision was to architect the identity metasystem such that it could accommodate identity providers exhibiting either of these mutually-exclusive requirements.

As a result, the system supports release of the identity of the relying party to "auditing" identity providers. But when interacting with non-auditing providers, it only releases a one-way function of the relying party's identity – computed on a per-user basis so the identity provider cannot deduce the identity of the relying party.

3.9. Authentication Goes Both Ways

Identity systems are typically used to prove the identity of the user to the relying party. But many forms of "phraud" are possible because the identity of the relying party is not adequately proven to users, meaning that imposter sites can pass as the real thing. A key design decision for the identity metasystem is to require that a site prove its identity to a user before the user ever supplies any information to the site.

3.10. Predictable, Protected Human Communication

Human beings are bad at handling complexity. Faced with unfamiliar choices, some fraction of the population will make the wrong decisions, even when those decisions are not in their best interests. Thus, a key design decision is to make the interactions that the metasystem has with its users as simple, familiar, self-explanatory, and predictable as possible.

This is achieved, in part, by making the communication channel *with the user* as narrow and constrained as possible, thus eliminating noise on the channel (complexity) that could increase the likelihood of the user misunderstanding the communication. Our user studies show that familiarity is a powerful weapon against social engineering attacks. When faced with the unfamiliar interactions caused by many forms of attacks within an otherwise familiar and predictable channel, the studies show that users will quickly and reliably realize that "something's not right here", decline to continue down the attack's path of choices, and thus thwart the attack.

Part of the familiarity comes from the design decision to represent all identities using the same InfoCard metaphor on the desktop, no matter what underlying identity technologies their providers use.

Another kind of familiarity derives from the user recognizing his or her portfolio of identities. Consider an analogy. If someone were to hand you a wallet that wasn't yours and try to get you to use it, you'd quickly look inside, see that the cards in the wallet were not your cards, and recognize that this wasn't your wallet. Similarly, if an attacker was to try to spoof the InfoCard user interface, they would be unlikely to convince many users to use it, because while they could put up the right sets of decorations on the window, the attacker wouldn't know your set of cards.

Finally, the InfoCard user interface is protected on the Windows implementation by running it in a separate secure desktop under a different user account. This means that unless malicious code is running with administrative privileges, it can't even see the InfoCard process, let alone control or communicate with it. All local secrets are stored in an encrypted form and no programmatic interface to the card store is provided.

Some might argue that these technical measures aren't foolproof (which is true). But compared to entering identity information such as passwords in a browser running in the user's context, they do significantly raise the bar.

4. Status and Plans

Microsoft has been actively working with innovators and industry players since 2004 developing both the principles behind the Identity Metasystem and interoperable implementations. For instance, in May 2005, we demonstrated interoperation with an open source Java identity pro-

vider written by Ping Identity [PingID 05]. Implementation guides [Microsoft 05b] have been published enabling (and encouraging) people on non-Windows platforms to build interoperable Identity Metasystem implementations. Several beta versions of Microsoft's implementations have been released [Microsoft 05b], with more to come.

Microsoft recognizes that, for the Identity Metasystem to succeed, it must be widely adopted, including on non-Windows platforms and by non-Microsoft browsers and web servers. We are heartened by the widespread recognition that, while Microsoft may be competing with other platforms and others' software offerings, we all share a common interest in seeing a viable, ubiquitous Internet authentication solution deployed.

Microsoft will be shipping its "InfoCard" client implementation as part of WinFX [Microsoft 06] — a set of managed code APIs that will be available on all of Windows Vista, Windows XP, and Windows Server 2003. WinFX will ship at the same time as Windows Vista.

While we are not at liberty to disclose others' implementation plans, we are excited at the possibilities of implementations on non-Microsoft platforms as well. Stay tuned for future developments!

5. Conclusions

Many of the problems on the Internet today, from phishing attacks to inconsistent user experiences, stem from the patchwork nature of digital identity solutions that software makers have built in the absence of a unifying and architected system of digital identity. The Identity Metasystem, as defined by the Laws of Identity, would supply a unifying fabric of digital identity, utilizing existing and future identity systems, providing interoperability between them, and enabling the creation of a consistent and straightforward user interface to them all. Basing our efforts on the Laws of Identity, Microsoft is working with others in the industry to build the Identity Metasystem using published WS-* protocols that render Microsoft's implementations fully interoperable with those produced by others.

The design decisions presented in this paper are intended to result in a widely accepted, broadly applicable, inclusive, comprehensible, privacy-enhancing, security-enhancing identity solution for the Internet. We present them and the rationale behind them to facilitate review of these design decisions by the security, privacy, and policy communities, so that people will better understand Microsoft's implementations, and to help guide others when building interoperating implementations.

We believe that many of the dangers, complications, annoyances, and uncertainties of today's online experiences can be a thing of the past. Widespread deployment of the Identity Metasystem has the potential to solve many of these problems, benefiting everyone and accelerating the long-term growth of the Internet by making the online world safer, more trustworthy, and easier to use. Microsoft is working with others in the industry to define and deploy the Identity Metasystem. We hope that you will join us!

References

[Anti-Phishing 05] Anti-Phishing Working Group. *Phishing Activity Trends Report*, February 2005. http://antiphishing.org/APWG_Phishing_Activity_Report_Feb05.pdf.

[Cameron 05a] Kim Cameron. *Kim Cameron's Identity Weblog*, May 2005. http://www.identityblog.com/.

[Cameron 05b] Kim Cameron. *The Laws of Identity*. Microsoft Whitepaper, May 2005. http://msdn.microsoft.com/webservices/understanding/advancedwebservices/default.aspx?pull=/library/en-us/dnwebsrv/html/lawsofidentity.asp.

[FFIEC 05] Federal Financial Institutions Examination Council. *Authentication in an Internet Banking Environment*, October 2005. http://www.ffiec.gov/press/pr101205.htm and http://www.ffiec.gov/pdf/authentication_guidance.pdf.

[Gartner 05] Gartner. *Gartner Survey Shows Frequent Data Security Lapses and Increased Cyber Attacks Damage Consumer Trust in Online Commerce*, June 2005. http://www.gartner.com/press_releases/asset_129754_11.html.

[IETF 05] IETF. *Public-Key Infrastructure (X.509) (pkix)*, December 2005. http://www.ietf.org/html.charters/pkix-charter.html.

[Microsoft 05a] Microsoft. *Microsoft's Vision for an Identity Metasystem*. Microsoft Whitepaper, May 2005. http://msdn.microsoft.com/webservices/understanding/advancedwebservices/default.aspx?pull=/library/en-us/dnwebsrv/html/identitymetasystem.asp.

[Microsoft 05b] Microsoft. *Windows Vista Developer Center: InfoCard.* http://msdn.microsoft.com/windowsvista/building/infocard/.

[Microsoft 06] Microsoft. *WinFX Developer Center*, January 2006. http://msdn.microsoft.com/winfx/.

[PingID 05] Ping Identity. *SourceID InfoCard STS Toolkit for Java*, August 2005. http://www.sourceid.org/projects/infocards/.

[WS-MetadataExchange 04] *Web Services Metadata Exchange (WS-MetadataExchange)*, September 2004. http://specs.xmlsoap.org/ws/2004/09/mex/WS-MetadataExchange.pdf.

[WS-Security 04] *Web Services Security: SOAP Message Security 1.0 (WS-Security 2004)*, March 2004. http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf.

[WS-SecurityPolicy 05] *Web Services Security Policy Language (WS-SecurityPolicy)*, July 2005. http://specs.xmlsoap.org/ws/2005/07/securitypolicy/ws-securitypolicy.pdf.

[WS-Trust 05] *Web Services Trust Language (WS-Trust)*, February 2005. http://specs.xmlsoap.org/ws/2005/02/trust/WS-Trust.pdf.

Appendix A – The Laws of Identity

The "Laws of Identity" [Cameron 05b] are intended to codify a set of fundamental principles to which a universally adopted, sustainable identity architecture must conform. The Laws were proposed, debated, and refined through a long-running, open, and continuing dialogue on the Internet [Cameron 05]. Taken together, the Laws were key to defining the overall architecture of the Identity Metasystem. They are:

- **User Control and Consent:** Identity systems must only reveal information identifying a user with the user's consent.
- **Minimal Disclosure for a Constrained Use:** The identity system must disclose the least identifying information possible, as this is the most stable, long-term solution.
- **Justifiable Parties:** Identity systems must be designed so the disclosure of identifying information is limited to parties having a necessary and justifiable place in a given identity relationship.
- **Directed Identity:** A universal identity system must support both "omnidirectional" identifiers for use by public entities and "unidirectional" identifiers for use by private entities, thus facilitating discovery while preventing unnecessary release of correlation handles.
- **Pluralism of Operators and Technologies:** A universal identity solution must utilize and enable the interoperation of multiple identity technologies run by multiple identity providers.
- **Human Integration:** Identity systems must define the human user to be a component of the distributed system, integrated through unambiguous human-machine communication mechanisms offering protection against identity attacks.
- **Consistent Experience across Contexts:** The unifying identity metasystem must guarantee its users a simple, consistent experience while enabling separation of contexts through multiple operators and technologies.

The Laws of Identity are discussed in more detail in The Laws of Identity whitepaper [Cameron 05b]. To join in the discussion of the Laws of Identity, visit www.identityblog.com.