ANSWERING SECURITY QUESTIONS. QUESTIONING SECURITY'S ANSWERS

SEPTEMBER 2009
CLAIMS-BASED IDENTITY

COMPUTER SECURITY
ALERT

Computer Security Institute, 11 West 19th Street, New York, NY 10011, 212-600-3026, csi@ubm.com

Talk to Strangers

By fusing minimum data to maximum assurance, OpenID, infocards and SAML boost both privacy and security

Want to reduce your data security efforts? Have less data. Want to maintain your privacy? Don't give out so much personal information. Want to make sure a user is who they say they are? Then don't just ask the user; ask someone you trust to vouch for them. Want attackers to stop stealing your valuable data? Make your data less valuable to them.

The logic is sound, but historically, as it relates to electronic data, the practice has been difficult if not impossible. Making these logical actions both possible and relatively easy for everyone involved is the promise made by claims-based identity and access management—collectively, OpenID, information cards and SAML.

Call me a starry-eyed optimist, but I think it can and will deliver.

CSI Members-Only Web / Audio Conference on Claims-Based Identity
Wednesday, September 23, at 2 p.m. Eastern time

Discuss your identity management challenges with your fellow CSI members learn what they think about OpenID, information cards and SAML solutions. Register for the conference at:

https://cc.readytalk.com/r/wfp1oInnsgbh

During the meeting we'll discuss the results of this members-only survey on claims-based identity:

http://www.surveymonkey.com/s.aspx?sm=wDW3MDiPOn1IM5PCfR2t2w_3d_3d

The protocols each fill different and equally necessary niches—oversimply put, OpenIDs for pseudo-anonymous interactions, information cards for online purchases (and other secure transactions), SAML assertions for providing insiders with provisional access to confidential resources. They all, however, share two basic principles: the party requesting the authentication/access credential receives only the information they really need, and the information they receive is highly trustworthy.

These technologies can enable organizations to reduce, or in some cases eliminate, the need to transmit, receive and store large amounts of personally identifiable information, personal health information or payment card data—thus significantly easing an organization's privacy compliance troubles and saving them money. Organizations that issue OpenIDs and infocards for public use won't reduce their compliance efforts, but they will *make* money from the service, and most likely save money on fraud-related costs. For end users, these technologies can enable greater privacy, greater security and greater control over how one's personal information is used, while also simplifying users' lives by drastically reducing the number of online identities (usernames, passwords) they need to manage. Plus, while the high assurance level makes these types of credentials more valuable to the party requesting them, the manner in which some of the credentials are used makes them less valuable to attackers.

CONTENTS

So why aren't we all using it?

Until very recently, adoption of "identity 2.0" has been stymied by limited market awareness, insufficient product interoperability, insufficient credential portability, insufficient product penetration and a fundamental problem of chickens waiting on eggs and eggs waiting on chickens.

Now, however, those obstacles have been largely surmounted. The SAML community's efforts, championed by the Liberty Alliance, have eliminated most of SAML's interoperability troubles. You've probably already got all the technology you need to start using SAML 2.0 assertions. Microsoft's inclusions of CardSpace—Redmond's version of a client-side application necessary for using information cards—as a standard component of the Windows Vista operating system and the upcoming Windows 7, have equipped millions of users with the ability to use information cards (even though most of them haven't got a clue). The newly formed Kantara Initiative brings together major players from the OpenID, infocard and SAML communities to "foster identity community harmonization, interoperability, innovation and broad adoption through the development of open identity specifications, operational frameworks, education programs, deployment and usage best practices for privacy-respecting, secure access to online services," as it states on their Web site (http://www.kantarainitiative.org). The identity 2.0 landscape looks especially different in the light of last week's announcement that Yahoo!, PayPal, Google, Equifax, AOL, VeriSign, Acxiom, Citi, Privo and Wave Systems will provide OpenIDs or information cards for an open government pilot program being conducted by the United States Department of Health and Human Services, the National Institutes of Health and the Center for Information Technology.

Of course, not all the kinks have been worked out yet. Plus, there are new questions to be answered, particularly about liability issues.

Nonetheless, used appropriately, these identity 2.0 technologies are superior to the status quo by a cosmically wide margin. Now might be the time to start putting these technologies into active use, and now is *definitely* the time to prepare.

High assurance: a physical world parallel

A look at my United States passport shows my photo, my name, my date of birth, my place of birth, my nationality, a passport number, and the passport's issue and expiration dates. Not that much data, considering the number of hoops I had to jump through to get it. In order to receive my passport I had to go to a passport office in person. I had to present proof of my U.S. citizenship—in my case, this took the form of a "certified" birth certificate. (A certified birth certificate with a raised seal isn't something one carries around in one's wallet, so I lost track of it over the years and had to get a new one. To get my certified birth certificate, I went in person to the municipal vital records office in my city of birth, which was only open from 8 a.m. to 4 p.m. weekdays, presented my current, valid driver's license—which unfortunately did not contain my new address, so I also had to bring last month's utility bill as proof of address—and paid them $25 in cash.) I also had to provide a current, valid proof of identification that displayed both my photo and my written signature—in my case, this was my driver's license. I also had to provide a photocopy of both sides of my driver's license, on white letter-size paper. I filled out a hardcopy form that asked for my name, gender, date of birth, place of birth, Social Security number, current mailing address, phone number, e-mail address, any previous names I might have had and both my parents' names, dates of birth, places of birth and U.S. citizenship status. I had to sign that document in the presence of the passport office agent. I had my photo taken at the passport office. I paid them $114.85—$75 for the passport book, $25 for the execution fee, $14.85 for the overnight delivery. (I passed on the $60 "expedited" service.) In about four weeks, I received the passport—which had been made tamper-proof through a number of methods—in the mail.

The lesson here is that my passport is a *very* valuable credential, but what makes it valuable isn't what's *on it* but what went *into it*. Number one, the party who issued the credential (the "identity provider") was the U.S. Passport Office, not the bookstore down the street. Number two, the passport office wouldn't even issue me the credential until we'd completed a rigorous assurance process. If you want proof of certain claims I might make—my identity, my age, or, most importantly my U.S. citizenship status—you really can't get better proof than my passport.

So, if you feel that access to your organization's resources should require exceptionally high-assurance credentials, step one is to make sure that the credential was issued by a trustworthy source. That trustworthy source might be a party whom you already know directly and with whom you already have a trusted relationship—a colleague, business partner, contractor, or yourself. Or that trustworthy source might be some sort of an accredited organization that is also in a position to confirm the particular claims being made—for example, the U.S. passport office could provide the best proof of my identity, my bank could provide the best proof that I have an account with them, and Barnes & Noble could provide the best proof that I have a brand loyalty card from them (a claim that could not be upheld by my passport, regardless of how strong a credential it is).

So, although there are three claims-based IAM methods to discuss, it's useful to think of them as part of two categories: one category for IAM between parties with a direct, pre-existing trust relationship and one category for open IAM between parties that do not have a direct trust relationship.

Direct trust frameworks: SAML

The Security Assertion Markup Language, SAML, is a mature protocol, capable of supporting the highest level of assurance credentialing between parties that already have a trust relationship. SAML also enables superb resource provisioning, to ensure that users only have access to the resources they require.

Let's say you're a university research lab, working on a collaborative project with researchers at three other universities. It's mighty important to keep the intellectual property secure, but it's equally important that all the collaborators have access to your lab's work and you have access to theirs. Option one: send a lot of e-mails with a lot of attachments. Option two: convince all four universities to create individual logins for all the collaborators outside the university (so that each collaborator now has four logins to remember). Option three: SAML.

In option number two the universities only accept user credentials that they themselves issued. In option three, the universities choose to also accept user credentials that were issued by the other universities; SAML enables this "identity federation." So if a researcher from Tufts wants to access a resource at Rutgers, he proves his identity by providing the SAML assertion issued to him by Tufts (the same one he uses to access Tufts' resources), Rutgers accepts the assertion, then grants the researcher provisional access to Rutgers resources, based upon the access privileges Rutgers has granted him. Identity provided by Tufts; access provided by Rutgers.

Yes, in order for this to work, all the universities need to be using SAML. That does not, however, mean that they all need to be using the same product. Thanks in part to the Liberty Interoperable program led by non-profit Liberty Alliance, more and more identity and access management providers (CA, NTT Software, Oracle, Ping Identity, and others) are becoming truly interoperable, right out of the box. To see what works with what else, visit http://www.projectliberty.org/liberty/liberty_interoperable/implementations.

SAML can also be used internally for better provisioning access—for example to keep the work of the accounts payable department isolated from the work of the accounts receivable department. SAML assertions can be loaded onto a photo ID smartcard that could double as both a physical and logical access credential—making it a solid credential for any access requiring level four (top-secret) assurance. For case studies of a variety of SAML uses and users, see http://www.projectliberty.org/liberty/resource_center/case_studies.

Again, SAML is only appropriate for parties whose trust in one another's credentials is rooted in a pre-existing, direct relationship with one another. For everything else, there are OpenIDs or information cards.

Open trust frameworks: OpenID and infocards

If the Dalai Lama told me someone was a Buddhist and Albus Dumbledore told me someone was a wizard, I'd believe them. Me, personally, I've never met the Dalai Lama or Albus Dumbedore, but I nonetheless consider them to be enormously trustworthy sources of just that kind of information. It is in this type of open trust relationship that OpenID and infocards are based.

Let's back up. The current status quo for identity and access management online is that each Web site creates a direct trust relationship, of sorts, with each user—the site requires the user to create a username and password that are only good for that site...or so goes the theory.

In practice we know that people use the same exact username/password combination across many, many Web sites. Or they use very weak passwords because they're easier to remember and who can keep track of a bunch of unique combinations of lowercase, uppercase, numerals, and special characters? Or, they use one weak password. We know this. Attackers know this. We know that even if we do a tremendous job securing our public Web site, attackers will steal a database of login credentials from a poorly secured site, plug those same credentials into our site, and easily breach the accounts of some of the mutual users, who use the same login data for both accounts. We know this, so we add new factors of authentication—show them an image, ask them a question, send them a text message, give them a one-time token key fob—which could become cumbersome and expensive.

Enter OpenID. Instead of creating a direct trust relationship with all users, the Web site instead decides to accept OpenIDs from a third-party identity provider like Yahoo! (or from several identity providers). Users may obtain their OpenID from Yahoo! and manage it online through their browser. Users with a Yahoo! OpenID can use that one credential not only at that site, but at any of the sites that choose to accept Yahoo! OpenIDs.

Now, I know what you're thinking: How is that any better than the status quo? Isn't it, in fact, much worse?

No. Although a user only needs to remember or maintain the login credentials for one account—their OpenID account—and although the user presents the same OpenID to one site as she does to another, *each site sees a different OpenID.*

When a user presents an OpenID to a site for the first time, the site receives a newly generated, non-correlatable identifier for that user—meaning that the credential transmitted to and perhaps stored upon the receiving Web server can be used only on that site and no other. So a breach on one site won't result in a breach on another. Further, an attacker who's exploited a Web site's log-in page to redirect users to a site under the attacker's control will be presented with an OpenID that won't be usable on the legitimate site. (By "non-correlatable," this also means that none of the sites accepting the OpenID have any way of correlating a user's account on their site to that same user's account on another site. However, it is possible to disable the non-correlatable identifier, should you wish to share account data from one site with your account on another site.)

Further, “One advantage of both [OpenIDs and information cards],” aid Drummond Reed, executive director of the Information Card Foundation (ICF), ”is that, because that identifier is not something the user knows, it is largely impervious to social engineering attacks.”

So, in theory, the only way an attacker can obtain an OpenID credential that can be used at multiple sites, is to grab it from its place of birth—the user’s account at the OpenID provider’s Web site. And the log-in requirements for that account should be and can be more stringent than for your average Web site. The OpenID’s best practices for OpenID security (http://wiki.openid.net/OpenID-Security-Best-Practices) include that identity providers that use passwords to authenticate users “MUST require that their password verification form be displayed in an independent browser window or popup, with the address bar displayed,“ and “should deploy anti-password cracking defenses to prevent automated guessing of passwords.” They are also advised to “use HTTPS for their login screen. Password submission should always be over HTTPS.” Also, “OpenID Providers MUST not allow their Login or Approval screens to be framed by the relying party. Allowing the Login or Approval screens to be framed makes the approval flow vulnerable to clickjacking, and trains users to expect the URL Location bar to not have the OpenID provider’s URL, leaving them vulnerable to phishing.” The best practices for ID providers make no mention of multi-factor authentication, however their best practices for users include: “Users should consider taking advantage of additional authentication options offered by their OpenID Provider, including using a client certificate to sign in, two-factor authentication, or other stronger authentication options.”

Authentication methods for accessing OpenIDs vary in strength. For example, I signed up for a test Yahoo! OpenID, and while the password could be up to 32 characters, it only required six. The six-character password I tried, which did have one numeral and one capital letter, but did not contain any special characters earned a maximum password strength status—not strong enough for my tastes. Conversely, myOpenID (provided by JanRain) and VeriSign’s Personal Identity Provider (in beta) provide support for multi-factor authentication. myOpenID didn’t grant me “full strength” until I gave a 15-character password with several caps, one special character and several numerals. It required that I solve a CAPTCHA before granting me an account. It also gave me the added option of signing in with an SSL certificate. VeriSign, meanwhile, did not provide me any indicator of password strength. However after I opened the account, VeriSign provided me the options of installing a browser certificate, providing my cell phone number so I could use my Blackberry as a second factor, obtaining a one-time password fob, or obtaining a smartcard.

OpenIDs are adequate for simply maintaining “an account.” OpenIDs are well-suited to pseudo-anonymous activities for which the site needs assurance that the person using the Web service this time is the same person that used it last time, but for which it is unnecessary or undesirable that the site knows *who* this person is or anything else about them.

When things really start getting interesting is when a site needs to know something very important about a user, and isn't willing to simply take the user's word for it but rather requires it to be verified by a higher authority.

Enter infocards. Each user may obtain many infocards from many different identity providers who are each in a position to vouch for certain claims you might make about yourself when requesting access to a Web site—that you're over age 18, that you're *under* age 18, that you have a Citibank checking account, that there's available money in that account, that you're a resident of New York City, that you're a member of the AAA travel club (see sidebar on action cards), that you're a Buddhist, that you're a wizard, etc.

All of those infocards are stored centrally in the user's "card selector," which could be in the form of a Web application, like Azigo (formerly "Parity"), or could be an application that resides on a user's local machine, like CardSpace. When a Web site demands that the user log in, the user's card selector will only display those infocards that will satisfy the site's demands—meaning both that the card can vouch for the user's particular claim and that the "relying party" (RP) trusts the good word of the "identifying party" (IP) who provided the card. The user then chooses which of the acceptable infocards to present.

Infocards not limited to ID: Enter action cards

Information cards do not have to be used for proving identity alone—a fact that may ultimately prove a greater catalyst for adoption than better identity management.

Another way of using information card technology is for "action cards," which can be used to have richer relationships with customers, among other things.

For example, when conducting Web searches, a AAA membership infocard generates a symbol to indicate businesses that provide discounts to AAA members.

Information cards could also be used to encourage social responsibility. For example, Choix Vert (which means "choose green" in French) infocards can provide a similar indicator to point out organizations that exhibit superior environmental responsibility.

Then (just like it works in OpenID), the relying site receives a non-correlatable identifier representing the user's infocard, so that neither attackers nor the relying party can link the user's account on this site to their accounts on other sites. Further, because the infocards reside in a card selector controlled by the user—instead of a Web server controlled by the identity provider—the identity provider cannot make those correlations either. (Unless there was a need for the RP to communicate directly with the IP—to exchange public keys, for example.)

Also, card selectors provide users the ability to see how, where and when their credentials have been used. Thus they can potentially be used to detect and recover from fraud, enabling users to correct incorrect data or withdraw infocards previously submitted to an RP, for example.

Of course to *prevent* the fraud, that card selector full of all those rich credentials *must* be rigorously secured itself, through multi-factor authentication and preferably hardware-hardened in some way. For example, Microsoft recommends storing one's CardSpace infocard "wallet" in the TPM chip on one's local machine. (If you want your infocard store to be portable, but don't feel comfortable leaving them in the cloud, then you could store them on a USB stick and access them from another client machine that's equipped with a card selector—but just make sure that USB stick is heavy-duty encrypted and tamper-proof.)

"Yes [the card selector is] a single point of failure," said Reed, "but it's a strongly secured point." (I'd quibble that it *can be*, but not necessarily and automatically *is* strongly secured.) "A bank vault is a single point of failure, but nobody builds a bank out of a picket fence. People feel their money is safer in a bank, which is why they keep it there instead of under the mattress."

Let's get back to those identifying parties. If you're a relying party, a main reason for accepting an a user's infocard issued by a third party, is because that third party knows something *for certain* about the user that you don't know. If you really want to know what bank I use, ask my bank.

And once my bank tells you that my account's with them, what else do you *really* need to know? When the Dalai Lama tells me someone's a Buddhist, I'm not going to ask that person to recite The Four Noble Truths. When Albus Dumbledore tells me someone's a wizard, I'm not going to ask that person to perform a Patronus charm. In fact, I'm not going to ask them anything at all—not their date of birth, not their Social Security number, not even their name. My trust in Dumbledore and the Lama is that great.

Having this highest degree of assurance enables me to have the lowest degree of data...which brings us to the greatest thing about claims-based identity.

Least data = least compliance headaches

Let's say you are an online merchant. Let's assume that you are successfully paid for a product, know for certain that the purchase had been made by the legitimate credit account holder, and know that the buyer will be able to easily purchase things from your store in the future without ever needing to type in all their credit card information again. Would you still request, and even store, the user's credit card number, expiration date and CVV code? (Or, for that matter, their mother's maiden name, first pet's name, childhood hero or the last four digits of their Social Security number?)

No. Why would you? That collection of data is a decadently tempting target for attackers and is protected by a slew of security and privacy regulations, thereby demanding a huge security budget that you'll have to get approval for. Why invite that kind of trouble?

Taking it a step further, are you using a shipping service to send the buyer their new books? Then you don't necessarily need their address, as long as the shipping service knows it, which is information that the user can present them, through an infocard.

And it doesn't necessarily need to be a different infocard than they gave you. By default, the RP is not sent all the information on an infocard, only the minimum information they require. Users can choose to send the RP further information if they wish. What this means for the user is that they can get away with having a small number of high-assurance, information-rich infocards, instead of many infocards representing many specific combinations of data.

"This is one of those rare cases where you can get privacy and security and, for the trifecta, easier usability," said Reed.

Of course I'm oversimplifying.

The complex, uncharted territory

A major consideration is that not all identity providers follow the standards established by the larger identity 2.0 community to the letter.

You may still need more data than what could ideally be presented by these methods. Although the high assurance the bank gives you (the RP) may enable you to request less identity information (mother's maiden name, truncated S.S. number), unless the bank directly transfers the user's money into the your organization's account, you will still need the card data. Or, for accounting purposes you may need to store more data for at least a limited period of time. However, it is conceivable that a third-party service provider could handle that data for you, providing you access to information as need be. There are whispers that just this type of service may be available within 12 to 18 months.

Also, some of these interactions only work if your business partners also use OpenID, infocards or, in particular, SAML.

Other questions arise around legal liability. If the identity provider whose SAML assertions, Open IDs, or infocards you've decided to accept experiences a breach, does the provider shoulder all of the culpability, or do you, as an RP, share some of the legal responsibilities? What about if you are the identity provider who experienced the breach? Is this a situation you should work into your incident response plans and service agreements?

How many identity providers' OpenIDs and infocards should your Web site accept? How much, if anything, will it cost to accept them? If there is a charge, will it be a flat fee, or a per-transaction fee?

Fortunately (and unfortunately) these aren't questions you need immediate answers to, because there are very few organizations providing OpenIDs or infocards at the moment.

That may change sooner rather than later.

Open Government Initiative

There's reason to believe that adoption rates will pick up speed soon, because of the aforementioned announcement made Wednesday, Sep. 9 at the Government 2.0 conference in Washington, D.C. In an action that, hopefully, will decidedly break the ice, the U.S. government has joined forces with the few, brave private-sector identity pioneers to conduct an "open government" pilot program that will leverage OpenID and infocards to support greater government transparency and greater citizen/government interaction by providing greater privacy for citizens. The pilot program is a response to President Obama's open government memorandum that he issued earlier this year.

Yahoo!, PayPal, Google, AOL, VeriSign and Wave Systems have signed on to be OpenID providers. PayPal and Wave Systems will also be infocard providers, as well as credit bureau Equifax, global marketing service Acxiom, financial services giant Citi and children's privacy provider Privo. The U.S. Department of Health and Human Services, the Center for Information Technology, the National Institutes of Health (NIH) and related agencies are serving as the relying parties.

"Adoption [of OpenID and information cards] will be driven by what we've done here today," said Don Thibeau, executive director of the OpenID Foundation, the day of the announcement. "We want to make the pilot programs as open as possible. NIH will open it up to other agencies and we'll open it to the international partners of the OpenID Foundation and the Information Card Foundation. We want others to benefit from seeing the pioneers take the arrows."

Not just any OpenID or infocard will be accepted, however. The U.S. General Services Administration (GSA) has created a Trust Framework Adoption Process (TFAP), through which GSA will certify "trust networks" for government use. A "trust network" is a community of identity providers, all of which have been certified by the leader of the trust network. Two primary candidates for trust networks are the non-profit organizations the OpenID Foundation and the Information Card Foundation. The trust framework model, however will be based upon the InCommon SAML federation, a project (that grew out of the Shibboleth open-source single sign-on project that began in 2000) created to enable university students, faculty and researchers to access other universities' libraries and other resources. The InCommon federation is the first trust network that will be certified by the GSA.

"We're opening a new chapter," said Mary Ruddy, chair of the ICF Certification Committee. "We hope that what we're doing is flexible enough that we can easily extend it to new trust networks."

"This initiative is about the citizenry interacting with government," said Ron Carpinella, vice president of identity management at Equifax. "I think it's great that [Equifax is] an enabler of that, and that we're helping protect citizens' privacy." Carpinella said that providing assurance of identity through information cards is a natural extension of the services Equifax has provided for almost a decade. He said that Equifax is in a unique position to provide personally identifiable information—including name, current and previous addresses, date of birth, phone number, family information, etc.—on 90 percent of U.S. citizens. Equifax also developed "knowledge-based authentication" technology.

The hope is that the American government's support of these authentication and authorization methods will spur more adoption throughout the United States. According to Thibeau, OpenID use is increasing in Japan, Denmark, the United Kingdom, Australia and Canada. Reed says that Europe is the principal area of growth in information card use, especially in Germany; Reed says that the German chapter of ICF constitutes almost half of the entire organization.

What now?

'Tis true, all my claims-based identity dreams cannot yet be fully realized, and they will not be until more organizations begin providing and accepting OpenIDs, information cards and SAML assertions. As Roger Sullivan, president of the Kantara Initiative and vice president of Oracle identity management, explains it: "The technology is there, but the technology is way ahead of where the business standards are."

So what can you do now to ready your organization for the future?

Plenty. Number one, says Sullivan: "First establish what your business strategy is. What is your objective? What are you trying to facilitate?" Decide whether or not your organization has need for claims-based identity and access management. If so, what technology best suits those needs? Don't forget to consider the risk factors involved.

Next, find out what your capabilities already are. Are your client machines equipped with card selectors? Are your Web applications able to support OpenID? Can your current identity management product support SAML?

If SAML is your game, identify the organizations that your business might wish to partner with. Are they using SAML or at least equipped with products that support SAML? If so, what product? Is that product interoperable with the products you've already got? If not, should you purchase the same product as that potential partner organization, or is there another solution?

Before purchasing new software, make sure that it can support claims-based identity technologies, even if you're not immediately planning on kickstarting a complete claims-based identity management program.

Get yourself an OpenID, a card selector and an infocard or two and try them out. (Make a point of ascertaining whether or not the identity provider you fiddle with is completely compliant with the accepted best practices.) Learn how they work. If they suit your fancy, encourage your users to do the same.

When you're ready to become an identifying party or a relying party, go to the experts for help. Third-party services exist that will set you up as an IP or RP. Or if you'd like to set yourself up as an IP, the OpenID Foundation and the Information Card Foundation can provide software libraries. Or just call them up. The leaders of the OpenID Foundation, the Information Card Foundation, the Liberty Alliance and the Kantara Initiative are remarkably responsive and eager to help.

And remember: accepting OpenIDs, infocards and SAML assertions doesn't necessarily mean that you have to resolutely disband with your current IAM. What it could mean is that, if you get onboard before your business' competitors do, it could be one of those examples of how security can enable the business. —*Sara Peters*