Privacy and Identity Management

When developing an identity management system, designers must consider the system's purpose and particular privacy needs. A set of guidelines and advice can help them make these determinations.

Marit Hansen
Independent Centre for Privacy Protection Schleswig-Holstein, Germany

Ari Schwartz and Alissa Cooper
Center for Democracy and Technology

Creating and managing individual identities is a central challenge of the digital age. As identity management systems—defined here as programs or frameworks that administer the collection, authentication, or use of identity and information linked to identity—are implemented in both the public and private sectors, individuals are required to identify themselves with increasing frequency. Traditional identity management systems are run by organizations that control all mechanisms for authentication (establishing confidence in an identity claim's truth) and authorization (deciding what an individual should be allowed to do), as well as any behind-the-scenes profiling or scoring of individuals. Recent work has looked toward more user-centric models that attempt to put individuals in charge of when, where, how, and to whom they disclose their personal information.

Identity management technologies can help realize the potential of the digital age, whether by making e-commerce exchanges more seamless, tying together information on multiple devices, combating fraud, or enabling yet unimagined services. However, the digitization of information—by facilitating the collection, storage, and sharing of large amounts of data—can exacerbate privacy risks inherent in identity management systems.

Privacy in context

System designers with limited exposure to the concepts of identity and privacy might be tempted to apply blanket privacy rules to identity management systems to address the privacy risks that those systems create. For example, "collect as little information from individuals as possible" might seem like a rule that could help protect the privacy of participants in an identity management system. Although this approach's simplicity is appealing, in practice, the relationship between identity management and privacy is nuanced, and what might seem intuitive might not always apply. Designers must evaluate how a particular identity management system protects privacy in context—that is, accounting for the system's purposes, participants, and potential abuses.

With regard to minimizing data collection, consider an identity-risk-analysis system as an example. Identity-risk analysis involves determining the probability that an individual engaged in a particular transaction is using a stolen or forged identity. To make this determination, you'd want to gather as much information as possible about the individual involved so you can compare the transaction to the individual's history or profile. If the credit card involved in the transaction is suddenly being used to make purchases in countries where it's never been used before, for example, someone might be using the individual's identity fraudulently.

Although gathering and maintaining a rich profile of an individual and his or her transactions might seem antithetical to privacy interests, in this case it might actually help protect the individual's privacy by raising a red flag about suspected identity theft. So, although less data collection can often mean more privacy, in this case the opposite might be true.

The importance of understanding and accommodating the context in which an identity management system will be used extends beyond considerations for the amount of data collected. The "less data collected equals more privacy" idea also fails to account for the type and sensitivity of the identity information involved. An identity management system that collects and stores a person's single fingerprint can be more invasive than a system that stores a person's entire credit history. Likewise, a small amount of identity information that's shared with numerous parties or isn't properly secured might put an individual's privacy at greater risk than a large amount of information that's properly secured and accessed only by authorized parties. These nuances ultimately point to the need to evaluate identity management systems with respect to privacy in context.

Privacy guidance

There is no shortage of principles and guidelines for establishing and maintaining privacy in identity management systems. Determining how to apply them to a particular identity management system requires a solid understanding of the environment in which the system operates and of the risks and benefits that the system must balance.

Fair Information Practice Principles

Designing and choosing a privacy-protective identity management system requires a solid grounding in foundational privacy principles. The most widely accepted set of such principles is the Fair Information Practice principles (FIPs), which were first developed in the 1970s and have been adapted by many government agencies, public interest groups, and private companies around the world (see www.cdt.org/privacy/guide/bsic/fips.html). The Organization for Economic Cooperation and Development (OECD), for example, has issued a set of guidelines based on the FIPs that focus on privacy as personal data flows between its 30 member countries.[1]

These principles apply broadly to the collection and use of personal data in the traditional sense—names, addresses, government-issued identifiers, and so on. Insofar as identity management systems are concerned, the seven FIPs are highly instructive:

- *Openness.* The existence of systems containing personal data should be publicly known, along with a description of the system's main purposes and uses of the personal data in the system.
- *Individual participation.* Individuals should have a right to view all information that's collected about them. They should also be able to correct or remove data that isn't timely, accurate, relevant, or complete.
- *Collection limitation.* Limits to the collection of personal data should exist. Personal data should be collected by lawful and fair means and, where appropriate, with the individual's knowledge or consent.
- *Data quality.* Personal data should be relevant to the purposes for which it's collected and used. It should be accurate, complete, and timely.
- *Finality.* The use and disclosure of personal data should be limited. Personal data should be used only for the purposes specified at the time of collection and shouldn't be otherwise disclosed without the consent of the individual or other legal authority.
- *Security.* Personal data should be protected by reasonable security safeguards against such risks as loss, unauthorized access, destruction, use, modification, and disclosure.
- *Accountability.* The keepers of personal data should be accountable for complying with fair information practices.

These principles are the logical starting point for anyone designing an identity management system. Because the FIPs were developed before the dawn of the digital age, however, they might be inadequate for many new environments that require identity management. In the new digital environment, massive data collection is inexpensive and efficient, databases are seamlessly networked together, and the data collected goes beyond traditional notions of personal data. In the face of these changes, designers of cutting-edge identity management systems and technologies might find three additional principles helpful:

- *Diversity and decentralization.* Enrollment and authentication options in identity management systems should function like keys on a key ring, letting individuals choose the appropriate key for a specific need. Designers should resist centralizing identity information or using a single credential for multiple purposes. If linking several identity management systems and databases together proves necessary, designers should implement appropriate safeguards to limit the associated privacy and security risks.
- *Proportionality.* The amount, type, and sensitivity of identity information collected and stored by an identity management system should be consistent with and proportional to the system's purpose. Some systems might require greater amounts of data or more sensitive data than others, but each system should match its information collection limits to its goals.
- *Privacy by design.* Privacy considerations should be incorporated into the identity management system from the outset of the design process. Considerations include safeguards for the physical system components as well as policies and procedures that guide the system's implementation. Incorporating these considerations at the beginning will save time and effort in the long run.

Often, not all the principles will apply to a given system equally. System designers should consider each principle and how to maximize it within a given system, but might conclude that it's more appropriate to focus on some principles while downplaying others.

Regulations and guidelines worldwide

Identity management system designers must also respect the privacy laws and regulations within their jurisdictions. In some areas of the world, such as Europe, a strong legal framework has provided fertile ground for privacy guidance and tools that go beyond the FIPs. The following subsections describe the legal frameworks in Europe, the US, and Canada, along with other notable privacy initiatives in those areas.

European Union. In 1995, the EU developed harmonized data-protection legislation to be applied across all 27 EU member states.[2] The harmonization aimed to remove potential obstacles to cross-border flows of personal data and to ensure a high level of protection within the EU. Unlike the US's more sectoral approach, the European Data Protection Directive forms an overarching privacy regulation that all data controllers within the EU must adhere to.

The EU Data Protection Directive doesn't permit processing personal data at all, except when a specific legal basis explicitly allows it or when the individuals concerned consented prior to the data processing. Generally speaking, the FIPs apply in the legal context of Europe, in particular the paradigms of transparency, individual participation, and legitimate purpose. EU data-protection law also stresses the commonly accepted principle of data minimization, limiting the collection and processing of personal data to the extent necessary for the given purpose.

In Europe, identity management systems must comply with the law, so in theory they fulfill the principles we've described. With the conversion to digital processing and storage of personal data in identity management solutions, designers could implement the law's transparency requirements directly in the system technology. Similarly, the new crop of user-controlled identity management systems can help users maintain and exercise their privacy rights by technologically implementing legal obligations and even enhancing user privacy by going beyond what the law requires.

In this spirit, a wide range of industry, academic, and governmental organizations from across the EU have joined forces through the Privacy and Identity Management for Europe (Prime) project to develop working prototypes of privacy-enhancing identity management systems. (Early work in privacy-enhancing techniques appears elsewhere.[3,4]) These solutions support users' sovereignty over their private spheres and help enterprises with privacy-compliant data processing. The EU's Sixth Framework Program funds Prime, which is acknowledged as a flagship for privacy technology development by the European Commission.[5] Some of the concepts discussed in the following section are based on Prime's work.

US. With the rapid advances in information technology beginning in the 1990s, the US Congress came under increasing pressure to establish regulations to protect information privacy. The resulting laws have followed a largely sectoral approach, with distinct regulations for many kinds of consumer data, but no overarching framework to secure consumer privacy across the board. Today, the US has separate privacy laws for medical information (the Health Insurance Portability and Accountability Act), financial information (the Gramm-Leach-Bliley Act), data related to children (the Children's Online Privacy Protection Act), and a slew of others.

For identity system designers, this patchwork of regulations provides little baseline guidance for building privacy-protective systems. Designers will likely find standards such as the OECD principles or the European framework more helpful in building privacy protections into their systems, although they'll have to consider US law for systems involving data covered by any of the myriad US regulations.

Canada. Canada has what the US lacks—a baseline privacy law governing the use of personal data. The Canadian regime is roughly equivalent to the EU regime.

Identity system designers will likely find work by Ann Cavoukian, Ontario's Information Privacy Commissioner, to be helpful in understanding the Canadian view of privacy. Her 2005 paper, "7 Laws of Identity: The Case for Privacy-Embedded Laws in the Digital Age,"[6] gives a unique interpretation of an earlier paper by Microsoft's Kim Cameron, "The Laws of Identity."[7] Cameron's laws of identity describe the basis for a "unifying identity metasystem" that can be applied to identity on the Internet. Cavoukian's work teases out the privacy implications intertwined in this new vision for digital identity.

Building blocks for privacy and identity management

In the digital world, two core informational privacy concerns are:

- *Observability.* The possibility that others (potential observers) will gain information. Observers might include the parties communicating (for example, two people emailing back and forth), the service providers facilitating the communication (for ex-

Table 1. Different parties' sufficient knowledge in an online shopping scenario.

	NAME/IDENTIFIER	PURCHASED GOODS	SHIPPING ADDRESS	FINANCIAL INFORMATION
Vendor	Pseudonym 1	+		
Delivery service	Pseudonym 2		+	
Payment service	Pseudonym 3			+

ample, email or Internet service providers), and eavesdroppers (for example, attackers sniffing email content or Internet traffic).

- *Linkability*. The potential to link between data and an individual as well as potential links between different data sets that can be tied together for further analysis. Controlling linkability involves both maintaining separate contexts so observers can't accumulate sensitive data and being cautious when identity information is requested to keep track of information disclosure.

How much (or little) observability and linkability are desirable in a specific situation depends on its context as well as on the perspectives of the parties involved.

For some services, information is disclosed with the express purpose of making it observable—on social networks, for example. But even in such situations, designers can tailor observability in a fine-grained way (for example, letting users control which of their friends can see certain information on their social network profiles).

As for linkability, consider a social networking site that lets users set up multiple profiles. These profiles' linkability should be a key concern for the site designers—profiles could be publicly linked, linked only on the site's back end, or not linked at all. The social network's users might have different preferences from those of the site itself. For example, they might want to keep their work and personal profiles unlinked, whereas the site might view the creation of combined profiles as richer targets for marketing or other purposes. However the social network is designed, linkability should be a core consideration.

Several mechanisms and tools for identity management systems can help designers control observability and linkability. Whichever mechanisms a designer uses, they must be implemented in an easily understandable and user-friendly way. The Prime project's white paper demonstrates and illustrates these concepts for user-controlled identity management.[8]

Separating workflows

Incorporating linkability control into the design of an identity management system should entail a separation of contexts (which is in line with Helen Nissenbaum's concept of "privacy as contextual integrity").[9] A designer could do this by, for example, preventing globally unique identifiers (strings pointing to individuals) and instead limiting the identifiers' scope to the necessary domain. Using different pseudonyms in different contexts could prevent undesired context-spanning linkage and profiling by third parties.

Existing workflows could be "delinked" by separating domains that don't necessarily need to be linked. In some cases, specific service providers who are responsible for only a subset of tasks could perform this separation. An obvious example is an online shopping scenario in which a company selling goods uses a payment service and a delivery service. Table 1 divides this scenario into three subprocesses that the different parties can perform, thereby separating knowledge of the buyer's information. The subprocesses relating to the same purchase case must communicate status information to each other, but not the buyer's personal data, as long as everything runs smoothly. The Liberty Alliance project, which is developing specifications for federated identity and identity-based Web services, proposes a similar separation.[10]

This means, for example, that the delivery service would have to know the shipping address, but not the goods to ship. Of course the three processes aren't fully independent—a link must exist between the purchase, the payment, and the delivery; and delinking only works if the services involved agree not to share information. Still, this link could be realized under the control of the user who, for example, might send all data encrypted for the appropriate recipients. However, in the traditional world, the shipping address and the financial account information would typically contain the user's real name. Still, the purchase itself doesn't necessarily require a real name—today's online auction platforms commonly use pseudonymous accounts, and almost everyone has made cash purchases at a bakery or bookstore where real identities are unimportant. In fact, the use of pseudonyms in transactions is generally legally permissible as long as it doesn't harm others.

The separation of workflows is already in common practice in cases in which the use of personal data is heavily regulated (for example, only particular parties can process medical data under the US HIPAA regulations). But the practice is also useful when applied to online identity management systems and other forms of data collection that aren't necessarily subject to strict legal rules in all jurisdictions.

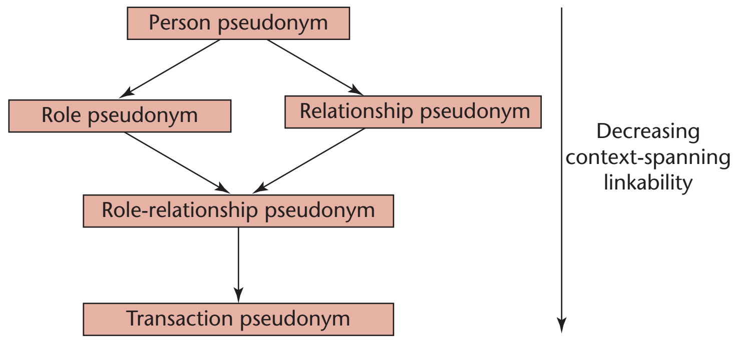

Figure 1. Pseudonyms according to their usage. Person pseudonyms are typically used as substitutes for real names in many contexts. Role pseudonyms are used with respect to a person's current role, such as a customer or patient. Relationship pseudonyms are used with respect to specific communication partners. Buying goods in two different bookshops, for example, would result in different relationship pseudonyms, regardless of whether the books belong to the private or professional context. Role-relationship pseudonyms combine the role and relationship pseudonyms and differ by role and communication partner.[11]

Choosing appropriate pseudonyms

From a technological perspective, all individual identifiers—except for real names—can be regarded as pseudonyms, even if they belong to hardware or software in the individual's possession. This can encompass IP addresses, cookie identifiers, hardware or software serial numbers, RFID tags, or other bit strings that are related to a person and might identify individuals within a certain scope.

Three main questions are relevant when discussing pseudonyms' privacy properties:

- Who knows (or can find out) a person's pseudonym?
- How strong is the link between the pseudonym and a specific individual? That is, does the individual possess the pseudonym uniquely and securely, or can different people consecutively or even simultaneously act under the same pseudonym?
- How much information can be gathered by linking data disclosed under the same pseudonym (that is, the content of a pseudonymous profile)? In other words, is the pseudonym used in a context-spanning or context-specific way, thus providing more or less information to be linked?

Figure 1 shows how pseudonyms might vary in aiding or restricting linkability.

For all situations, designers can tailor pseudonyms according to the required properties. For users, proper pseudonym handling in the online world to separate contexts isn't always trivial; user-controlled identity management systems should provide more effective mechanisms for achieving separation. In principle, the goal should be to manage all possible identifiers that might enable linkage, including the identifiers that correspond to the data trails in the digital world that most users aren't even aware of.

Private credentials

Private credentials (also called minimal disclosure tokens) let individuals prove their authorization (for example, that they're over 18 years old) without revealing information that might identify them.[12,13] In the encryption context, these private credentials derive from a certificate issued on different pseudonyms of the same person. Equipped with special cryptographic software, users can create multiple private certificates from a single master certificate that a credential provider has issued. These private certificates are linkable neither to each other nor to the issuance interaction in which the master certificate was obtained, and the credential issuer is rarely involved when the derived private certificates are used. Private credentials ensure users' accountability without giving away their privacy, as long as they behave according to the agreed-upon rules. Victims of misuse can revoke the user's anonymity with the credential provider's help.

Other types of private credentials exist. E-coins, for example, use credential providers that don't keep identity information. Although these credentials can't guarantee accountability, they can detect or even prevent misuse (for example, double-spending) in some cases.

Privacy policies

Organizations are familiar with displaying their privacy policies on their Web sites. But providing privacy policies that users truly understand and that serve as rules for automated data processing within the organization continues to be a challenge. Privacy policies are often the baseline for informed consent, which is needed before the organization can process users' identity information. In theory, machine-readable privacy policies (standardized in Platform for Privacy Preferences format, for example), should be a good way to match against (or possibly negotiate with) configured preferences on the user's side. The semantics of privacy policies need further international harmonization, and organizations need incentives to implement machine-readable policies. Currently, the lack of implementation makes the noble goal of greater transparency through the use of these polices an unlikely outcome.

The same is true for making privacy policies more accessible and understandable as we move into a world of ubiquitous connectivity, tiny mobile devices, and similar technological advances. Graphical (or even multimedia) expressions of privacy policy content,

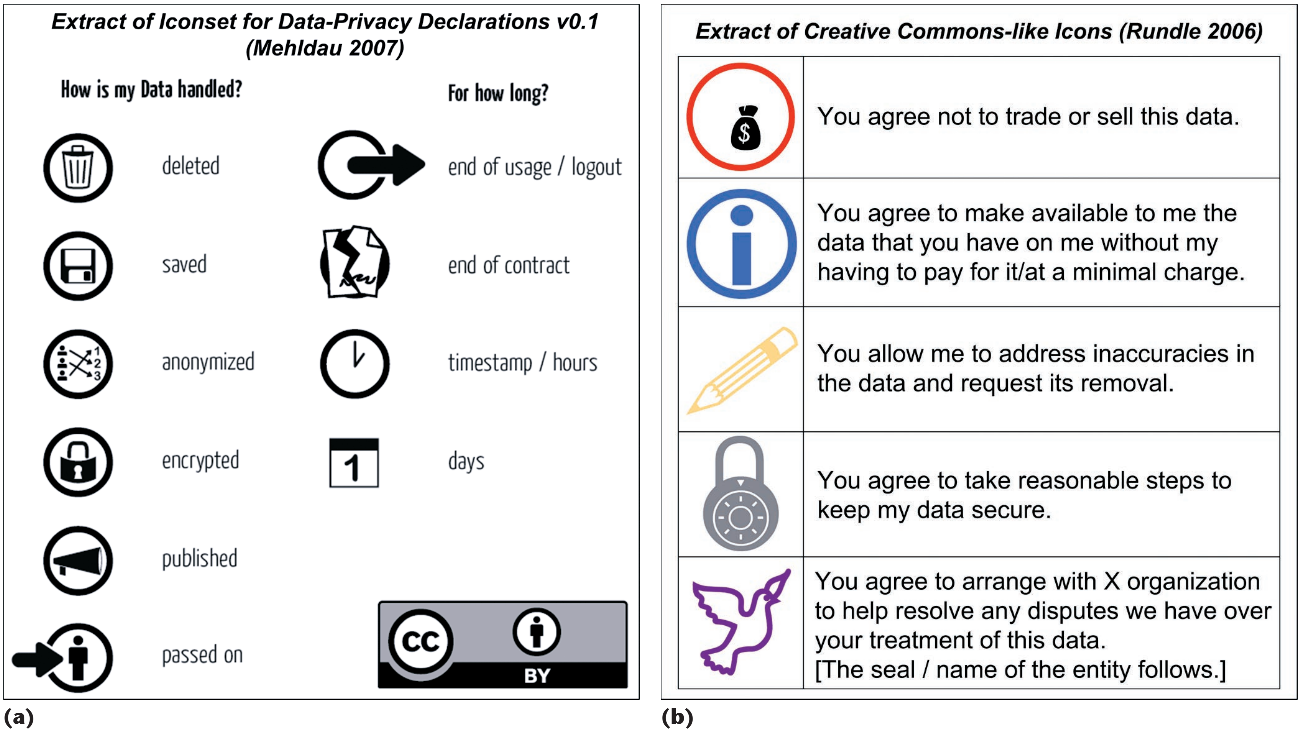
Extract of Iconset for Data-Privacy Declarations v0.1 (Mehldau 2007)
How is my Data handled?
deleted
saved
anonymized
encrypted
published
passed on
For how long?
end of usage / logout
end of contract
timestamp / hours
1
days
CC
BY
(a)
Extract of Creative Commons-like Icons (Rundle 2006)
$
You agree not to trade or sell this data.
You agree to make available to me the data that you have on me without my having to pay for it/at a minimal charge.
You allow me to address inaccuracies in the data and request its removal.
You agree to take reasonable steps to keep my data secure.
You agree to arrange with X organization to help resolve any disputes we have over your treatment of this data.
[The seal / name of the entity follows.]
(b)

Figure 2. Snippets from proposed icon sets for expressing privacy policies. (a) Matthias Mehldau developed a set of pictograms for data-privacy declarations (see the full icon set at http://asset.netzpolitik.org/wp-upload/data-privacy-icons-v01.pdf). (b) Mary Rundle proposed a set of "Creative Commons-like icons" (see her presentation on data- and identity-protection tools at http://identityproject.lse.ac.uk/mary.pdf).

such as simple and recognizable icons, can spare people from having to read lengthy texts in legal jargon. Figure 2 presents two example privacy policy icon sets.

Sticky policies

Are users sold down the river after releasing their identity information? Not necessarily. Current data-processing systems usually can't guarantee the binding between the data collection's purpose and the data's actual uses. However, researchers have proposed leveraging cryptography and other mechanisms to "stick" policies to data, similar to how digital rights management (DRM) tries to stick copyright policies to content.[14,15] These "sticky policies" together with data-management systems can guarantee privacy-compliant processing by enforcing the rules on how the data may be processed even after the information has been disclosed and left the user's control.

Transparency tools

What do others know about me? Knowing the answer to this question is a prerequisite for maintaining privacy. History functions such as the Prime project's Data Track store all relevant information from online transactions, including a record of what identity information has been disclosed to whom and under what conditions. The stored data also includes information from the privacy policies of services requesting the data. Users can review this information later to understand what exactly they've consented to. The Data Track doesn't only provide transparency (clear visibility) for users, but also lets them later ask data controllers whether they really treated the data as promised. In Europe, this would mean exercising users' privacy rights to access, rectify, or erase data and would let them possibly withdraw consent. In addition, the Data Track helps users choose the appropriate pseudonym and password for a particular context, keeping them separate unless otherwise desired.

Another aspect of transparency is information on current security vulnerabilities or reported privacy-related misuses. The Prime project has proposed security and privacy RSS feeds to alert users of potential risks or misuse. These RSS feeds could get the information from Computer Emergency Response Teams (CERTs), but also from companies that must act according to security breach notification laws, as required in many US states and planned in the upcoming revision of the EU ePrivacy Directive.

Usable system design

Users should be able to control their private spheres in an identity management system. Otherwise, they might blindly trust the system and unwittingly release more identity information than they intended. User interfaces must provide all necessary information without overwhelming users, a particularly tricky task in the complex field of privacy regulation. Because few users configure their IT systems, the systems' default privacy settings are critical. A single universal default setting won't suit all individuals, so users should be able to configure identity management systems according to a trusted party's recommendations, such as a privacy commissioner, a consumer protection organization, or simply a skilled peer. Existing usability research can help inform the construction of these mechanisms.

Systems can accomplish many goals without using an identity component at all, dramatically lessening the time and effort required to safeguard privacy.

Advice for practitioners

These building blocks are in different stages of development within a wide range of initiatives and products. Even when choosing among available identity management products and services, system designers face an array of choices and interoperability scenarios for software, hardware, and the protocols that define interactions within a system. We've developed some advice to help designers navigate the landscape of these choices.

Determine whether identity is necessary

The first consideration should always be whether you need an identity management system to solve the problem at hand. Systems can accomplish many goals without using an identity component at all, dramatically lessening the time and effort required to safeguard privacy. System designers shouldn't assume that adding an identification element to a system will make it more robust. The advantages of collecting and using identity information should be weighed against the need—and possibly legal requirements—to protect privacy.

Identify risks

Developers of all kinds of systems commonly plan only for regular workflows and processes, without considering the possibility of failure or attack. Understanding all risks to an identity management system, whether they're likely to occur daily or are highly unlikely to occur, is fundamental to protecting privacy in the system. Threat-analysis tools in the IT security field, such as attack trees, are well-known among experts, yet underused in identity management settings.[16] These tools are suitable for identifying privacy risks.

Discourage unnecessary linkages

In a networked world, the urge to link identity management systems and databases together will always exist. Linking together disparate identity data might improve convenience, efficiency, and even security (in cases such as fraud detection, in which linking information can help detect and deter fraudulent activity). System designers should choose components that let them easily erect strong safeguards to ensure that unnecessary linkages—between databases, communications channels, and personnel—don't occur. These safeguards should be built in during an identity management system's design phase.

For example, in the earlier online shopping scenario, you could design the database of identity information controlled by the delivery service to only store shipping information and pseudonyms. Although this doesn't prevent later linkages to other identity information, the fact that you'd need a new database schema to add this information later might discourage linkages down the line.

Implement security during design

Data security products have been in use for decades and should be one of the most straightforward features for designers to include. A comprehensive security plan should be developed from the outset to ensure that encryption, automatic deletion of identity information, network security processes, physical security safeguards, and the like are inherent to the system.

Adopt trust-enhancing measures

Even the most secure identity management systems must gain user trust. Many simple mechanisms are available to help enhance trust in the system and make users more comfortable. In accordance with the FIP openness principle, providing a clear, simple, layered privacy policy will provide the baseline information that users need to evaluate the system. Offering users a way to give feedback about the system and responding to that feedback in a timely and helpful manner will help build user confidence. Users should be able to easily access, correct, and in some cases delete information about themselves, and there should be a structured procedure for challenging conclusions drawn from that information. System designers should also consider applying for a privacy seal or publishing the results of a third-party privacy audit.

All of these measures will help build user trust and acceptance of the system.

The urge to identify individuals will only grow as new technological advances make identification easier and more cost effective. Perhaps the greatest challenge is to make privacy considerations an inherent part of the design process. Although they're frequently considered mutually exclusive, privacy, efficiency, and security often go hand-in-hand when they're considered from the outset.

We've explored an array of privacy principles, tools, and tips for identity management system designers looking to build privacy-protective systems. By determining which of these is appropriate for a particular system and grounding the system in a solid privacy framework, system designers will be on their way toward safeguarding privacy as they tackle the ever-increasing push toward individual identification. □

References

1. Organization for Economic Cooperation and Development, *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, 1980; www.oecd.org/document/18/0,3343,en_2649_34255_1815186_1_1_1_1,00.html.
2. European Commission, *EU Data Protection Directive 95/46/EC*, Oct. 1995; http://ec.europa.eu/justice_home/fsj/privacy/law/index_en.htm.
3. D. Chaum, "Security Without Identification: Transaction Systems to Make Big Brother Obsolete," *Comm. ACM*, vol. 28, no. 10, Oct. 1985, pp. 1030–1044.
4. B. Pfitzmann, M. Waidner, and A. Pfitzmann, "Secure and Anonymous Electronic Commerce: Providing Legal Certainty in Open Digital Systems without Compromising Anonymity," IBM research report RZ 3232, no. 93278, IBM Research Division, Zurich, May 2000.
5. Commission of the European Communities, *Comm. from the Commission to the European Parliament and the Council on Promoting Data Protection by Privacy Enhancing Technologies (PETs)*, COM(2007) 228 final, May 2007; http://eur-lex.europa.eu/LexUriServ/site/en/com/2007/com2007_0228en01.pdf.
6. A. Cavoukian, "7 Laws of Identity: The Case for Privacy-Embedded Laws of Identity in the Digital Age," Office of the Information and Privacy Commissioner/Ontario, Oct. 2006; www.ipc.on.ca/images/Resources/up-7laws_whitepaper.pdf.
7. K. Cameron, "The Laws of Identity," Microsoft Corp., May 2005; www.identityblog.com/?page_id=352.
8. R. Leenes, J. Schallaböck, and M. Hansen, eds., "Privacy and Identity Management for Europe," Prime whitepaper, ver. 2, June 2007; www.prime-project.eu/prime_products/whitepaper.
9. H. Nissenbaum, "Privacy as Contextual Integrity," *Washington Law Rev.*, vol. 79, no. 1, 2004, pp. 119–157.
10. S. Clauß and M. Köhntopp, "Identity Management and its Support of Multilateral Security, *Computer Networks*, vol. 37, no. 2, 2001, pp. 205–219.
11. A. Pfitzmann and M. Hansen, "Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management—A Consolidated Proposal for Terminology," ver. 0.31, 15 Feb. 2008; http://dud.inf.tu-dresden.de/Anon_Terminology.shtml.
12. J. Camenisch and A. Lysyanskaya, "Efficient Nontransferable Anonymous Multishow Credential System with Optional Anonymity Revocation," research report RZ 3295, no. 93341, IBM Research, Nov. 2000.
13. S.A. Brands, *Rethinking Public Key Infrastructures and Digital Certificates*, MIT Press, 2000.
14. G. Karjoth, M. Schunter, and M. Waidner, "Platform for Enterprise Privacy Practices: Privacy-Enabled Management of Customer Data," *Proc. 2nd Workshop Privacy Enhancing Technologies* (PET 2002), LNCS 2482, Springer, 2002, pp. 69–84.
15. M. Casassa Mont, S. Pearson, and P. Bramhall, *Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services*, tech. report, Trusted Systems Laboratory, HP Laboratories Bristol, HPL-2003-49, 2003; www.hpl.hp.com/techreports/2003/HPL-2003-49.pdf.
16. B. Schneier, *Beyond Fear: Thinking Sensibly about Security in an Uncertain World*, Springer, 2004.

***Marit Hansen** is deputy privacy commissioner of Land Schleswig-Holstein, Germany and head of the Privacy-Enhancing Technology (PET) department at the Independent Centre for Privacy Protection. Her research interests include identity management, anonymity, pseudonymity, transparency, and end-user empowerment. Hanson has a diploma in computer science from the University of Kiel, Germany. She's a member of the ACM and Gesellschaft für Informatik, where she serves as chair of the Special Interest Group on PETs. Contact her at marit.hansen@acm.org.*

***Ari Schwartz** is vice president and chief operating officer of the Center for Democracy and Technology. His research interests include online privacy, increasing individual control over personal information, and access to government information. Shwartz has a bachelor's degree in sociology from Brandeis University. He's a member of the Harvard Berkman Center's Stopbadware project Advisory Board and the State of Ohio Privacy Advisory Committee. Contact him at ari@cdt.org.*

***Alissa Cooper** is the chief computer scientist at the Center for Democracy and Technology. Her research interests include online privacy and security, Internet neutrality, and digital copyright. Cooper has a master's degree in computer science from Stanford University. Contact her at acooper@cdt.org.*