Identities Management for E-Commerce and Collaboration Applications

Michael Koch and Kathrin M. Möslein

ABSTRACT: The development of the Internet assumed that users would remain anonymous, but more and more services now need to identify users in order to provide personalized services or introduce users to other users. As in real life, users interact with services hosted by different providers and thus have to provide and update their personal information for each service separately, resulting in cold-start problems for new services and inconvenience for users. This paper argues the need for user-centric global identities management in future e-commerce and collaboration applications. It reviews the state of the art in this area, discusses needs and possibilities for future development, and proposes a novel solution for identity management.

KEY WORDS AND PHRASES: Electronic commerce, identities management, personalization, privacy, user profile.

Personalization and community support are increasingly important ingredients of successful applications for e-commerce and collaboration.

Personalization

Personalization techniques are used to tailor information services to the needs of individual users. The support of personalization for one-to-one marketing should increase customer share over a lifetime. What used to be possible in the corner shop, where the shopkeeper knew the customers personally, will be even more so in the electronic medium through the storage of profiles and automatic evaluation based on predefined rules.

Technically, personalization is about using personal information to select or filter information objects or products for an individual (*see Figure 1*). There are several methods for performing this selection, ranging from content-based filtering with rules or vector similarities to automated collaborative filtering (see [28] for more information). Aside from the method, the ability to deliver personalization is always based on the acquisition of an electronic representation of the user, or user profile. Every method of personalization has different

This work was partly funded by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) as part of the "SFB582: Marktnahe Produktion individualisierter Güter" program (www.sfb582.de) and by the German Ministry of Research and Education as part of the "COSMOS–Community Online Services and Mobile Solutions" project (BMBF FKZ 01HW0107 - 01HW0110, www.cosmos-community.org).

The authors thank Johann Schlichter, Ralf Reichwald, Roger Clarke, and Rolf Wigand for their feedback and suggestions on earlier drafts of the paper. Elizabeth McArdle and Christian Stewart provided helpful support during the final reading.

1086-4415/2005 $9.50 + 0.00.

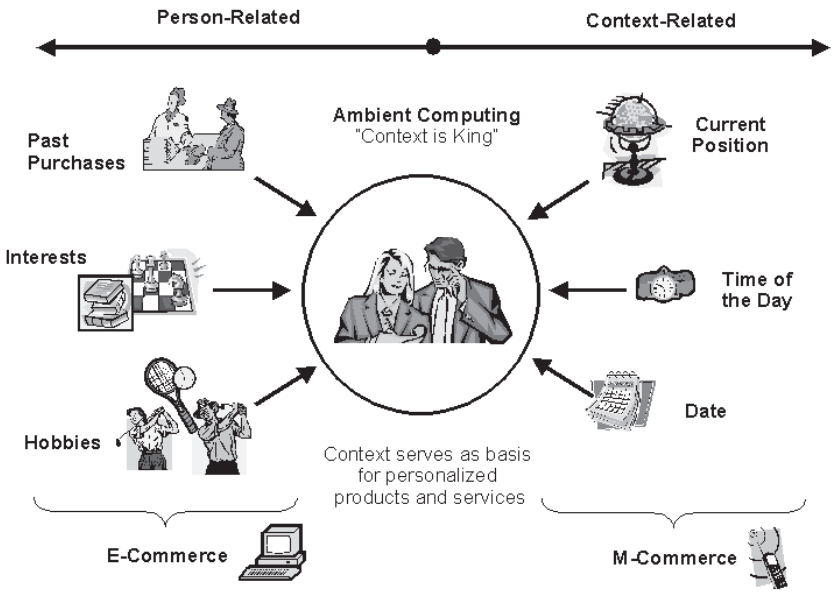

Figure 1. Personalization as Key to Ambient Computing [28]

requirements for the user profile. For content-based filtering, it is necessary to store information about preferred content and relationships to content objects. For collaborative filtering, relationships to other users and explicit ratings or comments have to be managed.

Community Support

Personalization is not the only reason for service providers to collect information about users. More and more often, Web-based services offer community support functionality of some kind. Users are not supported independently but are encouraged to contact one another. The services provide support for exchanging information, getting in contact, and communicating.

Three major possibilities are created by bringing communities of users together:

1. the building of trust
2. the collection and effective use of (trusted) community information
3. the economic impact of accumulated purchasing power

The economic impact of communities for accumulating purchasing power has already been discussed by Hagel and Armstrong [13]. They focus on groups of Internet users that are drawn together around products or companies and use the extended possibilities of the on-line medium to cooperate and gain advantages they would not have if they were acting as isolated customers

(better information, discounts). Such groups are often referred to as virtual communities of transaction [27]. In addition to the accumulation of purchasing power, communities of this type are sources for valuable data about products and about the community members. Thanks to techniques like collaborative filtering, data mining, and personalized user interfaces, the additional information about the users is often the basis for personalization.

In addition to the use of information about users for personalization, user profiles are needed in community support platforms for introducing users to one other. Communication is the primary activity in communities, and knowing the identity of those you communicate with is essential for understanding and evaluating an interaction and for building trust [8]. This means that the community members have to know one another and recognize one another's attributes. The user representation in the community need not be linked to the user's real-world identity, but it has to be persistent. Persistent pseudonymous user representations can, for example, be found in on-line auction platforms like eBay, where reputation information is stored as part of the user's identity [20].

In sum, for modern applications, user representations have to be available for personalization and for introducing users to one another.

User Representation and Identities

Identity

The discussion to this point has treated "user representation," "user profile," and "identity" as synonymous. Before proceeding further, it will be useful to clarify these terms.

Webster's New World Dictionary defines "identity" as "1) the condition or fact of being the same or exactly alike (sameness, oneness); 2a) the condition or fact of being a specific person or thing (individuality); b) the condition of being the same as a person or thing described or claimed" [32]. Accordingly, the two important features of identity are the process of verifying that a person is who he (or she) claims to be, and the idea of identity as the sum of all the information describing a specific person.

In e-commerce applications, the aspect of identity in terms of verification of a specific person plays an important role and already has been addressed broadly. For personalization and for collaboration support, however, the aspect of identity as all the information that describes a specific person is much more important. This article, therefore, will consider identity mainly in the context of a user profile, a set of attributes describing a person or role in the digital world.

Identities and Entities

The dictionary definition of "identity" embodies a claim for equality or "oneness." A person has one identity, but this is not true for user representations. A

person may use different representations during a life span, and even maintain several representations at once. Such multiple roles are neither illegal nor used primarily for illegal purposes. Since the term "identity" is to be reserved in this article for user representations, a new term is needed for "real-world identity." In the literature, the term "entity" is used for this purpose [6]. An "entity" is a real-world thing of any kind, including natural persons. An entity does not necessarily have a single identity but may have many. An identity is a particular presentation of an entity.

Several papers by Clarke offer a deeper discussion of "digital persona" and of the differences between entities and identities in the context of (id)entification [3, 4, 6].

Modeling Identities

Different approaches are taken when trying to capture identities in data structures. The most general approach is to model an identity as a set of attribute value pairs. A user profile model provides a schema for the attributes present in a profile.

Figure 1 shows different types of information about a user that can be used for personalization. The information ranges from the user's name, demographic attributes, and history of past transactions to dynamic attributes like the current position as used in location-based personalization.

Several standards are available for coding simple user profile information, such as address or payment information. Examples are the vCard standard and the profile schema included in the World Wide Web Consortiums P3P specification [14, 24]. These approaches are mainly based on hierarchically structured sets of attribute value pairs. For more complex information like interests or browsing histories, personalization applications currently define proprietary codings based on the application and on the algorithms operating on the information. Some first attempts to standardize codings for complex information can be found in the work on artificial intelligence and knowledge management research. See the papers by Essing and by Fink and Kobsa for information on abstract modeling of user profiles and on user profile servers [10, 11].

Acquiring Identities

The methods for capturing user representations engage the user to different degrees. One usually distinguishes asking the user (fill-in-profile, explicit feedback, or ratings) from observing the user (click-stream-analysis or transaction analysis) [28]. While these methods are important, they do not address some basic issues in user profile acquisition:

1. Users often do not trust services that collect and use profile information (and therefore try to provide no information or false information).

2. Even if users cooperate, the effort required of them results in a time delay until enough information is collected to provide appropriate recommendations. Small or infrequently visited sites may not have a chance to compile a detailed enough user profile.

The latter issue is called a cold-start problem. Users expect good recommendations from the outset, but the system is only able to provide recommendations after the user has answered its many questions or after it has watched the user for some time. This issue is of special importance for small and medium-sized enterprises (SMEs).

Identities Management

Managing the availability of information about users for applications is called identity management. Because the present discussion concerns managing several identities for one single user or entity, identities management is a more precise designation.

Identities management can help to solve the two key problems mentioned in the preceding section (distrust and cold-start) by

- giving control of identity information back to users (thereby solving the distrust issue), and
- allowing the reuse of identity information by the different personalization services (thereby solving the cold-start problem).

To gain these advantages, identities management solutions first have to separate

- the usage of identity information (in the personalization or collaboration services), and
- the storage of identity information (e.g., in a central user profile or identity server).

The separation of identities usage from identities storage opens the possibility of identity reuse and provides a single location where user access control and user awareness (of who is using which information) can be implemented in a way the profile owners trust.

Identities management is something one does in normal conversation every day when deciding what to tell others about oneself. In interactions with others, one considers the situational context and the role one is currently acting in as well as one's respective relationships with the communication partners. This consideration results in different sets of information being released to different interaction partners. Sometimes this leads to a situation in which a person is known by different names in different contexts—for example, a special name, nickname, or pseudonym suiting the occasion [19].

In the electronic world, an identities management system would allow people to define different identities and roles, associate personal data to them, and decide which users to give data to and when to act anonymously. An identities management system would empower users to maintain privacy and control their digital identities.

Components of Identities Management

A close look into identities management indicates that it entails the following tasks: (1) storing identity data, (2) authenticating identity owners, (3) allowing the identity owners to define access rights for other users, and (4) evaluating access rights when answering queries.

In order to perform these basic tasks, an identities management system has the following main components:

- *directory service (repository):* storing information about registered users
- *authentication:* establishing the validity of user identities (linking a user to an identity)
- *authorization (policy control):* controlling access to user profile information (for other services and other users)

Virtually every service (and device) today holds a directory with information about one or more identities. Most services also have some authentication functionality. The main issue for the future is to reduce this variety to a few directories, or possibly a single directory, and add authorization functionality in a user-centric way.

Authorization in identities management can be characterized as defining and enforcing access rights to user profile attributes. Users do not want to share their information with all services, because they may regard some services as more trustworthy than others. Requirements for authorization in identities management include flexibility, arbitrary level of detail and granularity, purpose binding of requests, control over sharing and distribution of data, and dependency on the privacy policy of the service.

Federated Identities Management

Identities management is at present a core component of system security environments inside companies, where it is used to maintain account information for log-in access to a system or a limited set of applications. The rise of inter-company collaboration, the need for identity data about customers, and the emergence of Web services architectures are driving the need for companies to understand and manage inter-company dependencies. Centralized Internet-wide identities management services and federated identities management networks are solutions addressing these issues.

Federated identities management systems provide mechanisms for companies to share identity information between discrete domains. The central attribute of such systems is that no single site holds the information about all identities (*see Figure 2*). Federated systems support multiple identity providers and a distributed, partitioned store for identity information. Clear operating rules govern the participants in a federation—both the operators of components and the operators of services who rely on the information provided by the identities management system. In a federated environment, users can log on through their identity provider and then leverage that authentication state to easily access resources in external domains.

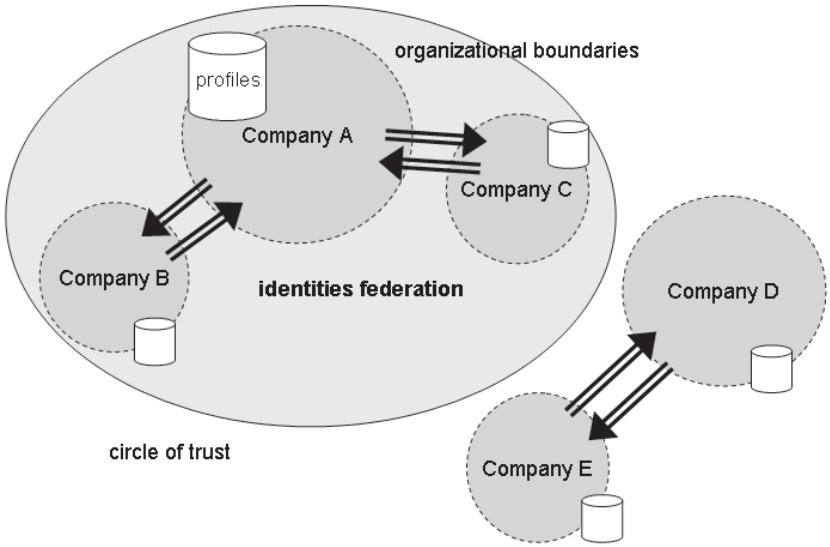

Figure 2. Identities Federation

The federation, or decentralization, of identities management can be seen as a contribution to minimize the risk of a single point of failure. This is actually more a way to bridge organizational boundaries without forcing any organization to give up control of its own information.

Federated identities management will be discussed in more detail in connection with the solutions covered in the following section.

Privacy

Information privacy is the position holding that information about individuals should generally not be available to other people or to organizations, and that one should be able to exercise a substantial degree of control over any personal data about oneself in the possession of another party [5].

One important way to maintain privacy has already been mentioned: Having different views on your entity in the form of different identities. Therefore, using different identifiers contributes to privacy. When interacting with a service, one provides access to only a single particular identity. Providing means for defining (service-specific) access rights to the user profile (authorization) also contributes to privacy. However, a single identifier is used for the dynamic views in this case, and services could get information they are not supposed to by combining their knowledge with the knowledge held by other services.

In the context of privacy, there are several different levels of identity:

- *veronymous:* It is possible to derive the real entity or an entifier from the information in the identity.
- *pseudonymous:* The identity (identifier) is persistent—that is, it is used several times to indicate that this is the same person, but it is not possible to derive the entifier from the attributes of the identity.

- *anonymous:* The identity is only valid for one transaction or page access, and it is not possible to derive identifiers or entifiers of formerly used identities.

Clarke has coined the term "nym" to distinguish (id)entifiers for (id)entities from keys referencing pseudonymous or anonymous user representations: "A 'nym' is one or more data items relating to an (id)entity that are sufficient to distinguish it from other instances of its particular class, but without enabling association with a specific (id)entity" [6].

The importance of pseudonymous and anonymous identities was shown in a survey conducted by Ackerman, Cranor, and Reagle [2]. Respondents were less inclined to provide information when personally identifiable information was requested. In a scenario involving a banking Web site, 58 percent of the respondents said that they would provide information about their income, investments, and investment goals in order to receive customized investment advice. However, only 35 percent said they would also supply their name and address so that they could receive an investment guide booklet by mail [2, p. 5].

On the technical side, the different levels of identity can all be reduced to authorization. The main point in pseudonymous and anonymous identities is concealing user profile attributes or identifiers, thus preventing a party from finding out about the real entity or about other identities of the person.

There is an obvious need for mechanisms that would allow users to specify and enforce their personal preferences regarding privacy (authorization). In earlier work on identities management, the authors collected the following requirements for a privacy infrastructure:

- Flexible access rights control system (e.g., through rules and negotiation)
- Monitoring of access rights and accesses
- Option to use a pseudonym instead of a real identity
- Data access rights bundled by purpose
- Ability to allow access for temporary use
- Ability to revoke granted access rights
- Control of whether user data can be distributed to other services (and users)
- Integration of cryptographic techniques for anonymous data transfers

On-line services and businesses would benefit from powerful user-centric privacy architecture in identities management. Users who are less fearful about risking their privacy are likely to make more and better personal information available to services [19].

Existing Identities Management Solutions

Work related to exchanging user profiles and to identities management can be found both in product development and in research. The solutions can be separated into two basic categories. First, there is work on client-based profile management. These infomediary solutions store identity information on the user's computer. Second, there are server-based identities management solutions.

Infomediaries

Identity information can be stored on the user's computer rather than in different services and can be provided to services when needed. Locating personal information near the user in this way would increase trust, especially because the user could control and monitor the usage of profile information.

Client-side user profile storage is implemented by so-called infomediaries. These are (small) applications on the client computer that manage user profiles and offer services such as automatic fill-in of Web forms or P3P interfaces for exchanging the information with services [7]. Jotter (www.jotter.com) and Persona (www.persona.com) are two examples of infomediaries.

The main problem with client-side storage of user profile information is that it is not portable [23]. Personal information stored on one computer (e.g., at work) cannot be easily transferred to another one (e.g., at home or a mobile device). An additional problem with today's infomediaries is that the definition of access rights is far too complex for everyday usage.

Single-Sign-On and Server-Based User Profile Databases

Whereas infomediaries focus on user control of identity information (authorization), server-side solutions are often more service-centered and focus solely on authentication. These solutions relate back to multi-server authentication solutions like Kerberos [29]. In such single-sign-on (SSO) solutions, different servers or services share one service to authenticate users.

Today different software vendors offer single-sign-on solutions for intranets. The solutions are mainly based on (X.500) directory services or are accessible via the LDAP directory access protocol.

While the single-sign-on solutions mentioned before are tailored for intranet usage, global solutions like DigitalMe from Novell (www.digitalme.com) or Microsoft Passport (www.passport.com) extend this approach to a service for Internet usage [9]. The core of these services is a central user profile directory. Users can store and maintain their personal data in these directory servers via Web interfaces. Services certified by the profile storage operator can get access to the authentication and profile information when a user tries to log in at these services.

Microsoft .NET Passport

Passport is the service that made identities management well known. Passport lets you sign up with a minimum of (unverified) personal data: a working e-mail address and a password. Therefore, Passport provides "persistent pseudonyms"—as many as you want. In its pure form, Passport is just authentication. It does not assert attributes other than the correspondence between identity and e-mail address. The main service it offers is an SSO that provides the authentication for different services. Microsoft is using its market

power to propagate Passport by making Passport accounts obligatory for all users of Microsoft on-line services, including the Hotmail e-mail service. As a result, the Passport service currently reports several hundred million accounts and several billion authentications per month.

On the technical side, "http redirect" is used for authentication. Thereby, the user's Web browser is redirected from the service to the identity server if no authentication token is present. The identities management system then handles the authentication of the user and sets the authentication token (usually a cookie) that can be used to authenticate the user at further services.

On the privacy side, Passport provides an "opt in." You have to give explicit permission to have your profile information shared. However, once you opt in you can no longer control the sharing of the information, nor do you get any indication of how your information is used. Microsoft can forward it to all partners and accept new partners without notifying you.

In addition to the weak privacy, Microsoft Passport has been criticized for security problems and lack of privacy considerations [21]. Among other things, the problems with "http redirect" include eavesdropping on the transmission of authentication information and illicit use of stolen authentication tokens.

Liberty Alliance

The hundreds of millions of accounts in Microsoft Passport pale beside today's large-scale production authentication systems like Visa and MasterCard.

To allow interoperability, different identities management providers and companies that operate such large identification services have joined in the Liberty Alliance to develop a standard for connecting their identification and user profile storage services in a federated way. There are several articles providing more information about the industry's viewpoint on federated identity services [1, 30, 31].

Currently the Liberty Alliance is working to define an open standard for representing identities, authenticating users, and authorizing access to user profile information. The goal is to make it easy for services that store user profile information to exchange the information with one another [22]. The idea of the Liberty Alliance is a balance-of-powers notion that has all the companies competing to be the customer's first point of contact, whereas the Microsoft approach simply takes it for granted that Microsoft is the primary point of contact (for authentication, at least). The object is not to create a platform for sharing personal data, but rather to pass and link unique identifiers and confirm that they have been authenticated. Therefore, identities are not available to all members instantly, but information is forwarded from one member to another, based on bilateral contracts.

Federated identities management standards, like those being created by the Liberty Alliance, form an "encapsulation layer" around the local identity and security environments of diverse domains. This encapsulation layer is what provides interoperability between disparate security systems inside and across domains, thus enabling federation. The group of service providers that share linked identities and have business and operating agreements in place

is known as a *circle of trust* (*see Figure 2*). The attribute-sharing policies within a circle of trust are typically based on

- a well-defined business agreement between the service providers
- user notification of the information being collected
- user granting of consent for the type of information collected.

However, the Liberty Alliance is still focused only on authentication. As yet, there is no real user control of authorization built into the proposal.

WS-Federation

In parallel to the efforts of the Liberty Alliance, Microsoft and IBM have also started to design for federated identities management. Their effort is aligned with their promotion of Web services and the development of a security backbone for Web services. The first standard, WS-Security, offers mechanisms for attaching security tokens to messages, including tokens related to identity. WS-Federation builds on WS-Security and related standards like WS-Trust. It defines protocols for identity providers to link user and machine identities for implementing single-sign-on and distribution of user profile attributes.

Profile Information Exchange

Other related work in the commercial field pertains to exchange and synchronization of user profile information among users or among the applications of one user.

Examples of the replication of user data among users are business card exchange services (e.g., www.cardxchange.net). Users of these services can store their contact information (and any additional attributes) and make a subset (view) of it explicitly available to other users. When the owner changes the information, the electronic business card changes in all places, or the people replicating the information are notified by e-mail. Similar functionality is often built into community support platforms.

User-Centric Identities Management

The solutions described above are mainly concerned with authentication and not authorization. They fail to satisfy the need for a generic *user-centric* identities management solution. In short, these existing systems may provide powerful solutions for making identities accessible, but they do not really take the interests of the identity owner into consideration. The distrust issue is not solved, because there is no real control or awareness for the user.

The trust issue can only be addressed through a better orientation toward the identity owner. In particular, this includes adding access right definition (authorization) possibilities and providing awareness and transparency of access.

In the Cobricks project, the authors are piloting an identities management solution that is clearly focused on user empowerment and supports complex user profile attributes, the IDRepository [16, 17, 18].[1]

Design Requirements

The IDRepository's design addresses the following requirements derived from the discussion in the preceding sections:

- separation of user profile storage and user profile usage
- federated architecture for user profile storage
- user profile information model
- trusted attributes in user profiles
- user profile owner authorization of profile access
- awareness of profile access
- focus on usability

The design decisions taken for this architecture, for the information model, and for authorization to reflect these requirements will now be briefly outlined.

Architecture

Following the idea of separating user profile storage and user profile usage, the core component of the architecture is a user profile repository server (IDRepository) that stores identities and offers the owners of the identities and authorized services interfaces to access this information (*see Figure 3*). The server offers the functionality to store multiple identities and to link identities to one another. The linking is done by defining data propagation paths—that is, defining how attributes that change in one identity are propagated to other identities.

Each IDRepository server is designed to enable the central storage of the identities of several people. At the same time, it is possible to have several IDRepository servers for one person, each storing a different identity. All the server instances can work together in a federated network. The network of IDRepository servers can also be a subnetwork in a larger set of Liberty Alliance nodes. Using the Liberty Alliance circle of trust concept, the IDRepository nodes can forward information to non-IDRepository nodes that fulfill certain basic requirements toward user-centric identities management.

Figure 4 shows the setup of several IDRepository servers storing the identities of one user. The figure highlights two features of the IDRepository: (1) Identities stored on one or more servers can be linked (by defining data propagation paths). Such a link will ensure that changes to selected attributes in one profile will be forwarded automatically to the linked profiles. This offers an alternative to having one profile with different views defined by access rights and makes it possible to define different identities that cannot easily be mapped to the same person by external watchers. (2) Services may cache

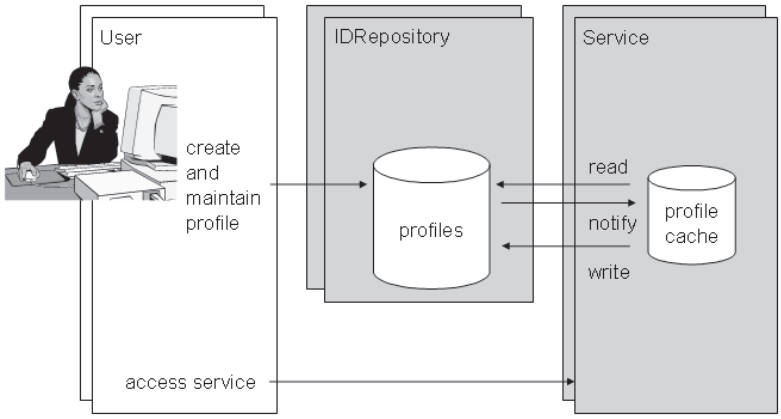

Figure 3. IDRepository Architecture Components

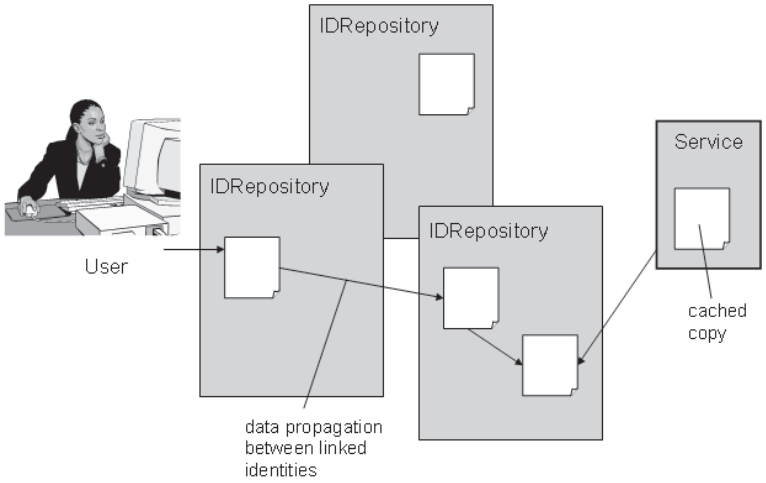

Figure 4. Identities Management Network

information retrieved from the IDRepository network. To keep the cached copies of the user profiles up to date, the IDRepository provides update notifications to services that cache information.

Certification of Attributes

Because attributes can be set by different sources, it may be necessary to associate meta-data to them (e.g., who last modified the attribute). In addition to knowing who wrote a particular piece of information, there has to be a means to verify this assertion. If an attribute is intended for use in only one application, the service could store it locally (and only store other attributes globally in the identity management service), but there are scenarios in which attributes

must be exchanged among services and still have to be trusted (e.g., an attribute that a user has bought for more than $1,000 at one e-commerce site that conveys entitlement to special discounts in other services). For this purpose, the IDRepository servers offer the functionality to sign any subhierarchy or subset of attributes in the repository and store the signatures. Anyone then can check the origin and the integrity.

Authorization

Authorization or access control has to address the issue of what platform should be provided with what type of information (to personalize), and which users should be able to see what information (for matchmaking and trust building).

As was discussed earlier in this article, authorization in identities management has to support flexibility, an arbitrary level of detail and granularity, the purpose binding of requests, and the dependency on the privacy policy of a service. In addition to primary access control, it also has to provide control over sharing and distribution of data (again based on the privacy policy of the service).

Traditional access control systems, such as role- or group-based solutions, do not sufficiently support these requirements (see [25, 26] for an overview of traditional access control systems). These solutions cannot deal properly with unknown services by making the access decision dependent on the privacy policy of the service or the purpose of the access.

The IDRepository uses a two-step access-token-based approach developed from P3P ideas [35]. In the first step, access requests are exchanged and negotiated (possibly with the participation of the profile owner), and an XML-based access ticket is transmitted to the requesting service. The access ticket contains information about the granted access rights (including the contexts in which they can be exercised). Using the access ticket, the service can efficiently request information from the IDRepository.

The access rights (and access tickets) may be dependent on

- the requesting service
- the privacy policy of the requesting service
- the purpose of the request
- the requested information (i.e., its granularity)
- the communication channel
- the identity level (veronymous, pseudonymous, anonymous)

More information on the access control scheme used in the IDRepository is available in several articles [33, 34, 35].

The access token solution addresses the issue of controlling which information should be provided to which platform. A rule-based approach has been developed for the second aspect of access control (controlling access to information by other users on platforms allowed to store copies). The user (with the help of tools and standard rule-sets) can define rules for which personal information other users can see. These rules may depend on any attributes in the identities of the requesting users and the requested user. Using

different functions on the profile attributes, it is also possible to allow access to user profile attributes in different levels of detail instead of just denying access. This feature can be used, for example, to determine in how much detail other users can see a person's current location in location-aware systems. Being part of the profile (as meta-data of user profile attributes), the rules are distributed whenever the user profile is replicated.

Limitations of User-Centric Identities Management

At first glance, the user-centric identities management concept has no drawbacks. It acts in the interests of the user, builds up trust in identities management solutions, and thereby increases their usefulness. User-centric identities management also generates new business possibilities for services providing trusted hosting of identities.

The main limitation of the user-centric identities management concept is, however, that user control can be exercised properly only when appropriate user interfaces for defining and managing the various access rights exist.

There already has been some work done on user interfaces for identities management (e.g., [12, 15]). The authors are building on this work and employing it in the IDRepository described above. More precisely, they are working on solutions for supporting users in defining access rights for other platforms and for other users. Some ideas about how this can be done are outlined below.

There has already been some work done on defining access rights for other platforms from the P3P project of the World Wide Web Consortium. In this context, the APPEL language was developed to express user privacy preferences that can be automatically matched with the privacy declarations of requesting services. However, APPEL was not designed for nonexpert users. The same is true for a powerful rule language developed by the authors to define access rights for other users.

The authors are currently providing predefined sets of rules to make the functionality available for nonexpert users. This involves compiling sets of rules that serve specific purposes (e.g., full privacy for all attributes, public visibility for name and e-mail address) and making them available to users. The effort to ascertain what sets of rules users need is still in progress. In this context, work is also in progress on determining means for distributing the provision of such presets among different providers, so that users can select presets from their trusted providers.

The issue of awareness is an important feature of "real-world" identities management (i.e., how people manage what they reveal to other people in direct encounters). People are aware of what information other people receive about themselves and can directly react on the information—both in dialogue and in shaping their presence as a whole. This suggests that providing awareness would be an important feature for an electronic identities management—awareness of what is revealed to others, and awareness of who is conscious of what information.

Providing this awareness in a way that users can manage will be an important issue of future work on user interfaces for identities management. The authors are currently working to extend the broad body of work on awareness issues in computer-supported cooperative work in this direction. This includes employing different means for visualizing the information gathered in the identities management network to be consumed in a peripheral way.

Summary and Conclusions

The availability of identity information for user representation will be important for future Internet-based e-commerce and collaboration applications. Information about users is needed for performing transactions, providing personalized services, and introducing users to one another.

Identities management and central user profile repositories might

- motivate users to make user profile information available (because they have control and awareness about who is using it)
- enable services to provide effective personalization without cold-start problems.

These two effects could help boost the use of personalization in on-line services.

The ultimate benefit of activities like Microsoft Passport and the Liberty Alliance project will be to make authentication and data-sharing practices open and visible. The review of market development shows that federated identities management solutions will soon become a market standard for enabling e-business and e-collaboration.

To win the trust of the profile owners, however, a solution has to clearly support

- definition and handling of different access rights and subidentities
- provision of awareness of access to the profile information.

The functionality has to be provided in an intuitive way and has to cover the emerging mobile applications that also need user profile information to perform their services. Therefore, users should be able to choose among different (trusted) operators of identity management servers.

The IDRepository presented in this paper extends the possibilities of these solutions by providing user-centric identities management while preserving interoperability with emerging networks wherever possible.

There are still some challenges to be addressed. From the technical point of view, the most important issues are

- how to specify (and enforce) access rights (especially including usability and user interface issues)
- how to represent user profile data to make it usable by different services.

Another challenge, which cannot be solved through technology alone, is the issue of access rights enforcement. As already discussed in the P3P project

of the World Wide Web Consortium (W3C), a certification is needed to ensure that the services make accurate statements about their privacy policy. An issue linked with service certification is the selection of a trusted operator for the identities management service. Different options have to be available for this choice, and the question of business models for future identities management providers has to be addressed. The first ideas for such business models are drawing upon an analogy between "user profile banks" and traditional banks that have gained trust and provide access to money from everywhere.

When this issue is solved, it will be possible to extend the scope of central identities management solutions to appliances—that is, to have personal appliances load user profile information for personalization from the central repositories.

NOTE

1. See www.cobricks.org for more information.

REFERENCES

1. Aberdeen Group. Federated identity systems: An executive white paper. Technical Report, Aberdeen Group, Boston, June 2002.
2. Ackerman, M.S.; Cranor, L.F.; and Reagle, J. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the ACM Conference on Electronic Commerce.* New York: ACM Press, 1999, pp. 1–8.
3. Clarke, R. The digital persona and its application to data surveillance. *Information Society*, *10*, 2 (June 1994), 77–94.
4. Clarke, R. Human identification in information systems: Management challenges and public policy issues. *Information Technology & People*, *7*, 4 (December 1994), 6–37.
5. Clarke, R. Internet privacy concern confirm the case for intervention. *Communications of the ACM*, *42*, 2 (February 1999), 60–67.
6. Clarke, R. Authentication: A Sufficiently Rich Model to Enable e-Business. Working Paper, 2001 (www.anu.edu.au/people/Roger.Clarke/EC/AuthModel.html).
7. Cranor, L.F. Agents of choice: Tools that facilitate notice and choice about Web site data practices. In *Proceedings of the 21st International Conference on Privacy and Personal Data Protection.* Hong Kong: Office of the Privacy Commisioner for Personal Data (PCO), 1999, pp. 19–25.
8. Donath, J.S. Identity and deception in the virtual community. In P. Kollock and M. Smith (eds.), *Communities in Cyberspace.* London: Routledge, 1998, pp. 29–59.
9. Dyson, E. Digital Identity Management. Release 1.0, 6, 20 (June 2002).
10. Essing, N. Anwendungen der Benutzermodelle aus der künstlichen Intelligenz: Repräsentationsmöglichkeiten für Benutzerprofile. Technical Report, Institut für Wirtschaftsinformatik, Westfälische Wilhelms-Universität Münster, 2001.

11. Fink, J., and Kobsa, A. A review and analysis of commercial user modeling servers for personalization on the World Wide Web. *User Modeling and User-Adapted Interaction*, *10* (2000), 209–249.
12. Gerd tom Markotten, D.; Jendricke, U.; and Mül, G. Benutzbare Sicherheit—Der Identitätsmanager als universelles Sicherheitswerkzeug. In G. Müller and M. Reichenbach (eds.), *Sicherheitskonzepte für das Internet*. Berlin: Springer, 2001, pp. 135–146.
13. Hagel, J., and Armstrong, A. *Net Gain: Expanding Markets Through Virtual Communities.* Boston: Harvard Business School Press, 1997.
14. Howes, T.; Smith, M.; and Dawson, F. *MIME Content-Type for Directory Information (vCARD Specification).* RFC 2425, 1998.
15. Jendricke, U., and Gerd tom Markotten, D. Usability meets security—The identity-manager as your personal security assistant for the Internet. In *Proceedings of the 16th Annual Computer Security Applications Conference.* Los Alamitos, CA: IEEE Computer Society Press, 2000, pp. 344–355.
16. Koch, M. Interoperable community platforms and identity management in the university domain. *International Journal on Media Management*, *4*, 1 (2002), 21–30.
17. Koch, M. An architecture for community support platforms: Modularization and integration. In H. Luczak, A.E. Cakir, and G. Cakir (eds.), *Proceedings of the 6th International Conference on Work with Display Units: World Wide Work.* Berlin: Ergonomic Institut für Arbeits- und Sozialforschung, 2002, pp. 533–535.
18. Koch, M., and Wörndl, W. Community-support and identity management. In *Proceedings of the European Conference on Computer-Supported Cooperative Work.* Dordrecht: Kluwer Academic, 2001, pp. 319–338.
19. Köhntopp, M., and Bertold, O. Identity management based on P3P. In *Proceedings of the Workshop on Design Issues in Anonymity and Unobservability.* Berlin: Springer, 2000, pp. 141–160.
20. Kollock, P. The production of trust in online markets. In E.J. Lawler, M. Macy, S. Thyne, and H.A. Walker (eds.), *Advances in Group Processes*, vol. 16. Greenwich, CT: JAI Press, 1999, pp. 99–123.
21. Kormann, D.P., and Rubin, A.D. Risks of the passport single sign-on protocol. *IEEE Computer Networks, 33* (2000), 51–58.
22. Liberty. *Liberty Architecture Overview: Version 1.0.* Technical Report, Liberty Alliance Project, July 2002.
23. Mulligan, D., and Schwartz, A. Your place or mine? Privacy concerns and solutions for server and client-side storage of personal information. In L.F. Craner (ed.), *Proceedings of the Tenth Conference on Computers, Freedom and Privacy: Challenging the Assumptions.* New York: ACM Press, 2000, pp. 81–84.
24. P3P: The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, World Wide Web Consortium (W3C) Candidate Recommendation, December 2000.
25. Park, J.; Sandhu, R.; and Ahn G.-J. Role-based access control on the Web. *ACM Transactions on Information and Systems Security, 4*, 1 (February 2001), 37–71.
26. Sandhu, R., and Samarati, P. Access control: Principles and practice. *IEEE Communications Magazine*, *32*, 9 (September 1994), 40–48.

27. Schubert, P. *Virtuelle Transaktionsgemeinschaften im Electronic Commerce: Management, Marketing und soziale Umwelt.* Lohmar and Cologne: Josef Eul Verlag, 1999.
28. Schubert, P.; and Koch, M. The power of personalization: Customer collaboration and virtual communities. In *Proceedings of the Americas Conference on Information Systems.* Atlanta: AIS, 2002, pp. 1953–1965.
29. Steiner, J.G.; Neuman, B.C.; and Schiller, J.I. Kerberos: An authentication service for open network systems. In *Proceedings of the Winter 1988 Usenix Conference.* Berkeley: USENIX Association, 1988, pp. 191–201.
30. Sun. *Strategic Implications of Network Identity.* Technical Report, Sun Microsystems, Palo Alto, CA, 2002 (wwws.sun.com/software/sunone/wp-identity.pdf).
31. Sun. *How to Implement Network Identity.* Technical Report, Sun Microsystems, Palo Alto, CA, 2002 (wwws.sun.com/software/sunone/wp-implement_ni.pdf).
32. Webster. *Webster's New World Dictionary of American English.* 3d college ed., Cleveland: Webster's New World, 1988.
33. Wörndl, W. Privatheit bei dezentraler Verwaltung von Benutzerprofilen. Ph.D. thesis, Department of Informatics, Technical University of Munich, August 2003.
34. Wörndl, W., and Koch, M. Privacy in distributed user profile management. In *Proceedings of the 12th International World Wide Web Conference.* Budapest: World Wide Web Consortium, May 2003, pp. 20–21.
35. Wörndl, W. Using P3P to negotiate access rights to user profiles. Paper presented at the Workshop on the Future of P3P, Washington, DC, November 2002.

MICHAEL KOCH (kochm@in.tum.de) is assistant professor of applied informatics in the Department of Informatics at the Technical University of Munich in Germany. He is working on different projects in the area of collaboration support, ranging from the application of groupware solutions in enterprises to support for leisure interest communities with mobile devices.

KATHRIN M. MÖSLEIN (kmoeslein@london.edu) is associate professor at the Advanced Institute of Management Research at the London Business School, and senior research fellow at the Institute for Information, Organization, and Management at the TUM Business School, Munich, Germany. Her research focuses on the strategic implications of new information and communication technology on organizations and markets.