

Defining Identity Management



Identity management is the combination of business process and technology used to manage data on IT systems and applications about users. Managed data includes user objects, identity attributes, security entitlements and authentication factors.

This document defines the components of identity management, starting with the underlying business challenges of managing user identities and entitlements across multiple systems and applications. Identity management functions are defined in the context of these challenges.

Contents

- 1 Introduction** **1**

- 2 A variety of identity stores** **2**

- 3 Managing identities and entitlements across applications - the challenge:** **3**
 - 3.1 Different kinds of users 3
 - 3.2 Different kinds of identity data 3
 - 3.3 Identity life cycle 4
 - 3.4 Key identity challenges 6

- 4 Relevant technologies: the solutions** **7**
 - 4.1 Directories 7
 - 4.2 Meta Directories 7
 - 4.3 Web access management / Web single signon 8
 - 4.4 Password management 8
 - 4.5 Enterprise single sign-on 9
 - 4.6 User provisioning 10
 - 4.7 Role Based Access Control 11
 - 4.8 Access Certification 11

- 5 Identity management: a simple definition** **14**

- 6 Beyond the enterprise** **15**

- 7 Conclusions** **17**

- 8 References** **18**

1 Introduction

Identity management is the combination of business process and technology used to manage data on IT systems and applications about users. Managed data includes user objects, identity attributes, security entitlements and authentication factors.

This document defines the components of identity management, starting with the underlying business challenges of managing user identities and entitlements across multiple systems and applications. Identity management functions are defined in the context of these challenges.

The remainder of this paper is organized as follows:

- **A variety of identity stores:**

A description of why organizations manage user profile data in a diversity of systems.

- **Managing identities and entitlements across applications - the challenge:**

A step-by-step description of why managing user identity data is difficult in a large organization.

- **Relevant technologies - the solutions:**

How different technologies help to streamline and secure the identity management process.

- **Identity management - a simple definition:**

A definition for what constitutes identity management, given the preceding description of the business problem and its technological solutions.

- **Beyond the enterprise:**

How identity management technologies may soon extend beyond the boundaries of a single enterprise.

- **Conclusions:**

Some conclusions about the state of identity management today.

- **References:**

Where to learn more about identity management.

2 A variety of identity stores

Modern organizations run a complex mix of IT infrastructure, including:

- Network operating systems, used to share files and printers.
- Application servers, running web servers, databases and similar software.
- Mainframe and midrange servers, typically hosting legacy applications.
- E-mail and other collaboration software.
- User directories, publishing lists of users and other network objects.
- Human resources, payroll and contractor management systems.
- A variety of line-of-business applications.
- Customer relationship management (CRM) and enterprise resource planning (ERP) applications.
- Electronic commerce applications.

Many kinds of users access these systems, including:

- Employees.
- Contractors.
- Partners.
- Vendors.
- Customers.

Almost every system and application tracks its own users, how they sign in (i.e., their passwords) and their privileges (i.e., what they can see and do). This data about users must be managed, when users are hired, when their business roles or identifying information change and when they leave.

The diversity of these systems, each with their own security management user interface, administrators and change request processes creates complexity. This complexity impacts the IT operation – the same human user must be managed by different IT staff on different parts of the infrastructure. The complexity also impacts users – it can take a long time to make required changes and users are forced to memorize multiple login IDs, passwords and application sign-on processes.

This complexity leads to high IT cost, lower user productivity and security exposures.

Identity management technologies simplify the administration of this distributed, overlapping and sometimes contradictory data about users.

3 Managing identities and entitlements across applications - the challenge:

In this document, “enterprise” refers simply to medium to large organizations, with thousands of internal users.

3.1 Different kinds of users

Enterprises manage identity data about two broad kinds of users:

- **Insiders:** including employees and contractors.

Insiders spend most of their working hours engaged with the enterprise. They often access multiple internal systems and their identity profiles are relatively complex.

- **Outsiders:** including customers, partners and vendors.

There are normally many more outsiders than insiders. Outsiders generally access only a few systems (e.g., CRM, e-Commerce, retirement benefits, etc.) and access these systems infrequently. Identity profiles about outsiders tend to be less detailed and less accurate than about insiders.

The difference between insiders and outsiders and how this impacts identity management, may be illustrated by an example:

Consider a bank, with 15,000 employees, 5,000 contractors and 500,000 customers. Insiders at the bank are the 20,000 employees and contractors.

Insiders log into a network operating system, corporate Intranet, line-of-business applications, corporate mainframe, e-mail systems and Internet gateway. Their identity profiles include data relating to their employment and their many login IDs to internal systems. Insiders access components of their identity profile, in particular login IDs to various systems, many times each day.

Outsiders are primarily current and prospective bank customers. Their profiles may include from one to three login IDs and passwords – for Internet-, telephone- and ATM-based electronic banking. Their profiles also include customer information such as a mailing address and account numbers. Outsiders only access their login IDs occasionally. Personal profile data provided by outsiders, such as full name, home telephone number, or e-mail address may be inaccurate.

3.2 Different kinds of identity data

Just as there are different kinds of users whose identity an enterprise must manage, there are different kinds of data about these users that must be managed:

- **Personal information.**

This includes names, contact information and demographic data such as gender or date of birth.

- **Legal information.**

This includes information about the legal relationship between the enterprise and the user: social security number, compensation, contract, start date, termination date, etc.

- **Login credentials to managed systems.**

On most systems, this is a login ID and password. Identification may also use a PKI certificate and authentication may use tokens or biometrics or a set of personal questions that the user must answer.

3.3 Identity life cycle

As organizations deploy an ever wider array of IT infrastructure, managing that infrastructure and in particular managing users, their identity profiles and their security privileges on those systems becomes increasingly challenging.

Figure Figure 1 illustrates some of the challenges faced by organizations that must manage many users across many systems.

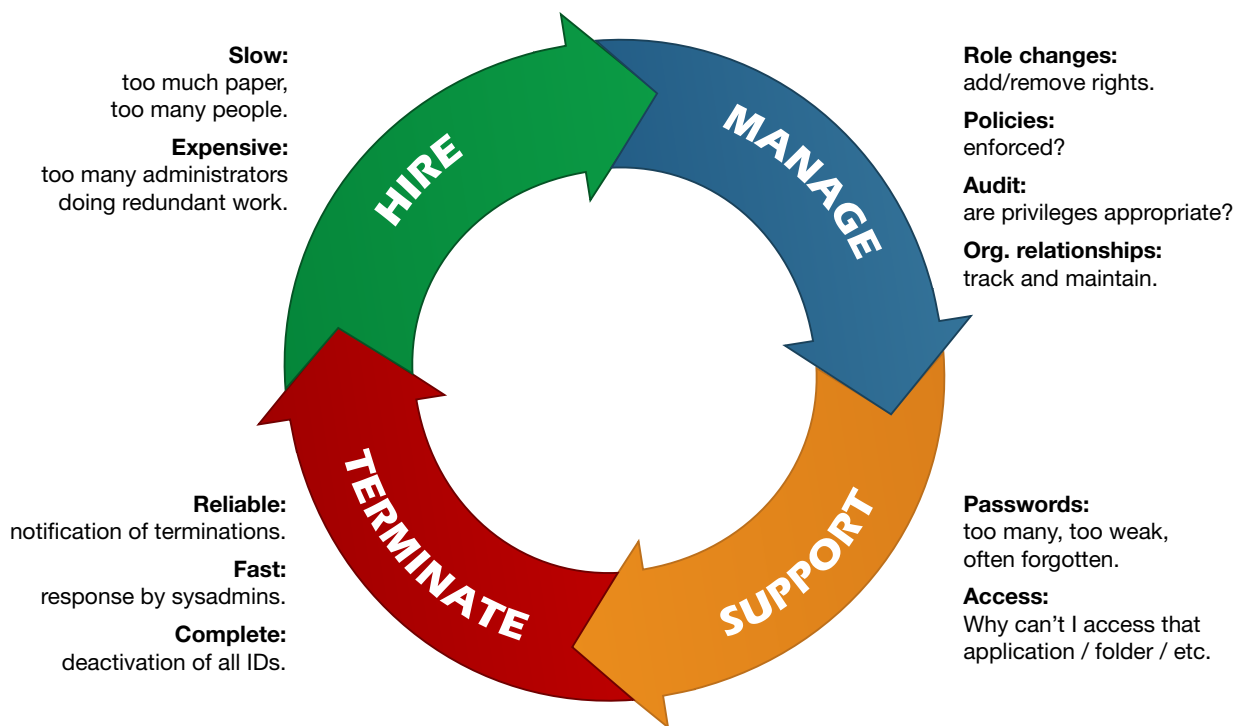


Figure 1: User Lifecycle Administration Challenges

In the figure, there are business challenges at each phase of the user lifecycle:

1. **Onboarding new users:**

- (a) **Delays and productivity:**

New users need to get productive quickly. Any delays in setting up access rights for new users cost money, in terms of lost productivity.

(b) **Requests and approvals:**

IT workers need to be certain that newly created accounts are appropriate. This usually means a paper process for requesting, reviewing and approving security changes, such as the creation of new users. This approval process may be hard to use, may require excessive effort on the parts of both requesters and authorizers and may introduce delays.

(c) **Redundant administration:**

Users typically require access rights that span multiple systems. A new user may need a network login, an e-mail mailbox, firewall access and login rights to multiple applications. These accounts are typically created by different administrators, using different tools. This duplication is expensive and time consuming.

2. **Managing change:**

Users often change roles and responsibilities within an organization. They may also change identity attributes (e.g., changes to a user's surname, contact information, department, manager, etc.). Such changes trigger IT work, to adjust user identity profiles and security rights.

Organizations face the same challenges in managing existing users that they face when creating new ones:

(a) **Delays:**

Reassigned users waste time waiting for IT to catch up with their requirements.

(b) **Change requests:**

Can be awkward to submit and may take time to approve.

(c) **Redundant administration:**

Similar changes are often required on different systems.

3. **IT support:**

In the context of routine use of systems, users often encounter problems that require technical support:

(a) Forgotten passwords.

(b) Intruder lockouts.

(c) Access denied errors.

Collectively, these problems typically represent a large part of an IT help desk's call volume. This means both direct cost (support staff) and indirect cost (lost user productivity).

4. **Termination:**

All users leave eventually. When they do, reliable processes are needed to find and remove their security privileges. These processes must be:

(a) **Reliable:**

If an organization fails to deactivate the access rights of a departed user, then that user or an intruder impersonating him might abuse the infrastructure or compromise sensitive data.

(b) **Timely:**

Access termination must be prompt, to minimize the time window available for the aforementioned exploits.

(c) **Complete:**

It is not enough to deactivate a departed user's login IDs on major systems. Every access right should be revoked, to eliminate the possibility of abuse by users inside the network.

3.4 Key identity challenges

Identity management presents several challenges in most organizations:

- **Security:**

Do user entitlements exactly match their needs? Are policies, such as segregation of duties rules, violated? Do access rights persist after they are no longer needed?

- **Consistency:**

User profile data entered into different systems should be consistent. This includes name, login ID, contact information, termination date, etc.

The fact that each system has its own user profile management system makes this difficult.

- **Efficiency:**

Setting a user to access multiple systems is repetitive. Doing so with the tools provided with each system is needlessly costly.

- **Usability:**

When users access multiple systems, they may be presented with multiple login IDs, multiple passwords and multiple sign-on screens. This complexity is burdensome to users, who consequently have problems accessing systems and incur productivity and support costs.

- **Reliability:**

User profile data should be reliable – especially if it is used to control access to sensitive data or resources. That means that the process used to update user information on every system must produce data that is complete, timely and accurate.

- **Scalability:**

Enterprises manage user profile data for large numbers of people. There may be tens of thousands of insiders and hundreds of thousands of outsiders.

Any identity management system used in this environment must scale to support the data volumes and peak transaction rates produced by large user populations.

4 Relevant technologies: the solutions

Several types of technologies are available to manage user identity data across the enterprise. In general, these systems focus on streamlining the identity management process and managing data consistently across multiple systems.

4.1 Directories

The cornerstone of many identity and access management infrastructures is a corporate directory.

Directories are network services which manage information about users, the organization and IT assets such as servers and printers. They perform a function similar to white pages or yellow pages phone books do for the public telephone infrastructure – enabling users of the network to find information about each other and about network services.

Most modern network directories are accessed using the lightweight directory access protocol (LDAP), which is based on the older, more powerful, but more complicated and less popular X.500 protocol.

A directory is just the starting point for identity and access management and provides no value in and of itself. To get value organizations must:

- Directory-enable their business applications, to eliminate silos of identity information.
- Implement effective technology and business processes to manage the contents of their directory.

Major platform vendors make inexpensive, robust and scalable directory products. These include:

- Microsoft Active Directory.
- Novell eDirectory (built on top of NDS).
- Sun ONE Directory (formerly Netscape and then iPlanet LDAP).
- IBM Directory (formerly Tivoli Directory).
- Oracle Internet Directory (OID).

There are also some open source directory products, such as OpenLDAP and Red Hat Directory Server.

4.2 Meta Directories

Meta directories are engines that synchronize data about users between different systems. A meta directory works as follows:

- Connectors to multiple target systems are configured, to read and write user profile data.
- Data streams from integrated systems are merged, to construct a master database of user profile information.

- Where a user's data in the master database differs from that user's profile on a lower-priority target system, the target system is updated to reflect the user's current information.
- Users may be added to or removed from target systems, based on changes detected on systems of record.

Meta directories simplify user administration by propagating changes from systems of record to managed systems, eliminating manual updates.

Since meta directories do not normally expose a user interface, or interact directly with users, they can be thought of as "plumbing" embedded in an enterprise identity and access management infrastructure.

An excellent meta directory product is ILM from Microsoft.

4.3 Web access management / Web single signon

A Web access management (WebAM) / Web single signon (WebSSO) system is middleware used to manage authentication and authorization of users accessing one or more web-enabled applications.

A WebSSO system intercepts initial contact by the user's web browser to a web application and either verifies that the user had already been authenticated (typically tracking authentication state in a cookie) or else redirects the user to an authentication page, where the user may use a password, token, PKI certificate or other method to authenticate himself.

Once a user is authenticated, the WebAM component of the system controls the user's access to application functions and data. This is done either by filtering what content the user can access (e.g., URL filtering) and by exposing an API that the application can use to make run-time decisions about whether to display certain forms, fields or data elements to the user.

WebSSO / WebAM products typically use an LDAP directory as a back-end repository, to identify all users. They often come tightly integrated with an "identity and access management" application, which enables delegated and in some cases self-service administration of the contents of that single directory.

4.4 Password management

Password management is a combination of password synchronization between systems and applications and self-service password reset.

Password synchronization is any process or technology that helps users to maintain a single password, subject to a single security policy, across multiple systems.

Password synchronization is an effective mechanism for addressing password management problems on an enterprise network:

- Users with synchronized passwords tend to remember their passwords.

- Simpler password management means that users make significantly fewer password-related calls to the help desk.
- Users with just one or two passwords are much less likely to write down their passwords.

There are two ways to implement password synchronization:

- Transparent password synchronization, where native password changes, that already take place on a common system (example: Active Directory) are automatically propagated through the password management system to other systems and applications.
- Web-based password synchronization, where users are asked to change all of their passwords at once, using a web application, instead of continuing to use native tools to change passwords.

Self-service password reset is defined as any process or technology that allows users who have either forgotten their password or triggered an intruder lockout to authenticate with an alternate method and repair their own problem, without calling the help desk.

Users who have forgotten or locked out a password may launch a self-service application using an extension to their workstation login prompt, using their own or another user's web browser or through a telephone call. Users establish their identity, without using their forgotten or disabled password, by answering a series of personal questions, using a hardware authentication token or by providing a biometric sample. Users can then either specify a new, unlocked password or ask that a randomly generated one be set.

Self-service password reset expedites problem resolution for users after a problem has already occurred and reduces help desk call volume. It can also be used to ensure that password problems are only resolved after strong user authentication, eliminating an important weakness of many help desks: social engineering attacks.

One of the core features of Hitachi ID Password Manager from Hitachi ID is self-service password reset.

4.5 Enterprise single sign-on

Users who log into many systems may prefer to sign into one master system and thereafter be able to launch applications having to type their ID or password again.

Most legacy and client/server systems cannot share authentication with modern infrastructures such as Kerberos or SAML. However, it is possible to store user credentials outside of the various applications and automatically enter them into applications when prompted.

Enterprise single sign-on (E-SSO) systems do just that: users sign into the E-SSO application, which stores every user's login ID and password to every supported application. Users launch various applications through the E-SSO client software, which opens the appropriate client program and sends keystrokes to that program simulating the user typing his own login ID and password.

Since they require the installation of client software, E-SSO systems are only appropriate for use by insiders.

E-SSO systems have had limited success in large production environments for a number of reasons:

- Deployment and integration costs.
- Concerns about security, due to the fact that the SSO system stores every user's password to every system.
- Concerns about availability, since if the SSO system fails, entire user populations will be unable to log into their systems and so will basically stop working.

Note that one E-SSO system does not store user passwords, but instead relies on password synchronization (Hitachi ID Login Manager – <http://Login-Manager.Hitachi-ID.com>).

4.6 User provisioning

A user provisioning system is shared IT infrastructure which is used to externalize the management of users, identity attributes and entitlements from individual systems and applications.

User provisioning is intended to make the creation, management and deactivation of login accounts and other user objects, which are spread across multiple systems, faster, cheaper and more reliable. This is done by automating and codifying business processes such as onboarding and termination and connecting these processes to multiple systems.

User provisioning systems work by automating one or more processes:

- **Identity synchronization:**
Detect changes to personal data, such as phone numbers or department codes, on one system and automatically make matching changes on other systems for the same user.
- **Auto-provisioning:**
Detect new users on an authoritative system (such as HR) and automatically provision those users with appropriate access on other systems and applications.
- **Auto-deactivation:**
Detect deleted or deactivated users on an authoritative system and automatically deactivate those users on all other systems and applications.
- **Self-service requests:**
Enable users to update their own profiles (e.g., new home phone number) and to request new entitlements (e.g., access to an application or share).
- **Delegated administration:**
Enable managers, application owners and other stake-holders to modify users and entitlements within their scope of authority.
- **Authorization workflow:**
Validate all proposed changes, regardless of their origin and invite business stake-holders to approve them before they are applied to integrated systems and applications.
- **Consolidated reporting:**
Provide data about what users have what entitlements, what accounts are dormant or orphaned, about change history, etc. across multiple systems and applications.

As well, a user provisioning system must be able to connect these processes to systems and applications, using connectors that can:

- Enumerate users and groups on the target system.
- Create new and delete existing login accounts.
- Read and write the identity attributes associated with a user object.
- Read and set flags, such as “account enabled/disabled,” “account locked,” and “intruder lockout.”
- Change the login ID of an existing account (rename user).
- Read a user’s group memberships.
- Read a list of a group’s member users.
- Add a user to or remove a user from a group.
- Create, delete and set the attributes of a group.
- Move a user between directory organizational units (OUs).

4.7 Role Based Access Control

Role-based access control (RBAC) is an approach to managing entitlements, intended to reduce the cost of security administration, ensure that users have only appropriate entitlements and to terminate no-longer-needed entitlements reliably and promptly.

In the context of a single system or application, RBAC means granting privileges directly to roles and attaching users to roles. Users acquire privileges through role membership, rather than directly. Within a single system, roles are sometimes called security groups or user groups.

Single-system RBAC is a time tested and successful strategy, as it allows administrators to group users, group privileges and attach groups of privileges to groups of users, rather than attaching individual privileges to individual users.

User provisioning systems extend RBAC beyond single applications. Roles in a user provisioning system are sets of entitlements that may span multiple systems and applications. The key element of roles is to replace many, technical entitlements with fewer roles that business users can understand. Business users can then a reasonable determination of which users should have which roles. This implicitly specifies which users should have which technical entitlements.

Roles consist of entitlements – login accounts and security group memberships. Roles are often also nested – i.e., one role can contain others. Nesting roles can reduce the cost of role administration.

4.8 Access Certification

Regulatory compliance requirements and security policies increasingly demand that organizations maintain effective controls over who has access to sensitive corporate information and personal data about employees and customers:

- Systems must limit access to just the right users, at just the right time.
- Organizations must be able to provide auditable evidence that these controls are in place and effective.

Section 404 of Sarbanes-Oxley specifically states that management must assess the effectiveness of internal controls on an annual basis.

- Organizations must be able to report which internal users currently have and had in the past, access to sensitive data.

Meeting these requirements can be challenging as users often have unique and changing business responsibilities, thus making their entitlements difficult to model using formal roles and rules.

The difficulty in modeling complex, heterogeneous entitlements is compounded by the fact that although users accumulate entitlements over time, they rarely ask IT to terminate old, unneeded rights. Moreover, it is difficult to predict when, after a change in responsibilities, a user will no longer function as a backup resource for his old job and so old entitlements can be safely deactivated.

These challenges together mean that it is difficult to model all of the entitlements that users need across multiple systems and applications at a single point in time and likely impossible to model those needs for thousands of users, over multiple systems, over an extended period of time.

Access certification is a process where business stake-holders are periodically invited to review entitlements, sign-off on entitlements that appear to be reasonable and flag questionable entitlements for possible removal.

There are several components to access certification:

- **Discovery:**

Before entitlements can be reviewed, they have to be collected from systems and applications and mapped to users. Technical identifiers should be replaced by human-legible descriptions that reviewers will understand. Since entitlements change all the time, discovery should be a regularly scheduled, automated process, not a one-time data load.

- **Who performs the reviews?**

Options include managers – asked to review their subordinates, application or data owners – asked to review lists of users who can access their applications or data or security officers – asked to review high risk entitlements.

- **When are reviews performed?**

The frequency may vary with the business risk posed by the entitlements in question.

- **What kinds of entitlements are reviewed?**

The highest level review is of employment status – should the in question still have access to any systems? Slightly more granular is a review of roles – should the user in question still have these roles? At the lowest level of granularity are basic entitlements – should the user in question have a login ID on this system or belong to this security group?

- **Which entitlements warrant a review?**

Not every entitlement poses a significant business risk. User membership in the social committee mailing list is not really worth reviewing, for example. Some determination must be made of the risk level posed by each entitlement, as this forms the basis for deciding whether to review it and how often.

- **What happens to rejected entitlements?**

Reviewers may flag entitlements as inappropriate, in which case something should be done. Does this raise a work order in an IT issue management system, or trigger a connector to revoke the entitlement immediately? Should further reviews take place before the entitlement is reviewed?

5 Identity management: a simple definition

With the above sections in mind, we propose a simple definition to encapsulate the various capabilities of enterprise identity management technologies:

Identity and access management is defined as a shared platform and consistent processes for managing information about users: who they are, how they are authenticated and what they can access.

6 Beyond the enterprise

Identity management can extend beyond a single organization:

- Customers would like to access multiple web sites without re-authenticating to each one.
- Employees would like to access vendor web resources without registering or re-authenticating.
- Companies would like to be able to provision their own users with access to partner and vendor resources automatically.

Federation enables applications in different domains to share information about users.

- Federated sites must have some pre-established relationship, bilaterally or in a group.
- Information about users is exchanged:
 - Identity: *Who is this user?*
 - Authentication: *How/when did the user sign in?*
 - Authorization: *What is the user allowed to do?*
- Federation enables single signon between sites:
 - User signs into one site (company A).
 - User clicks into another site (company B).
 - Site A passes information about the user to Site B.
 - The user is not prompted for his ID/password by site B.
- Federation reduces administrative burden:
 - Site B trusts Site A to name its own users.
 - Site B **does not** create its own objects for Site A users.

In order to work, federation requires that software at one site can communicate identity, authentication and authorization site to software at another site:

- Different organizations use different software products to manage user identity, authentication and authorization.
- To interoperate, different software products rely on standard protocols.
- There are multiple standards regarding federation:
 - Liberty Alliance ID-FF and ID-WSF.
 - Security Assertions Markup Language (SAML).

- WS-Federation.
- Shibboleth.
- CardSpace.
- The standards are complex, so it is reasonable to assume that different products implement them somewhat differently and may not be 100% compatible.

The problem with standards is that there are so many of them...

Hitachi ID Management Suite is intended to help an organization manage user identities and entitlements in a single domain – its own. It is compatible with federation products and protocols, but does not implement federation directly.

7 Conclusions

Identity management is a class of technologies intended to streamline the management of user identity information both inside and outside an enterprise. It includes:

- Directories, especially those using LDAP.
- Password management.
- Enterprise single signon.
- Web access management and web single sign-on.
- User provisioning.
- Federation.

8 References

- Basic functional definition of current technologies:
 - *What is Identity Management?*, Rutrell Yasin, Information Security Magazine, April 2002, http://www.infosecuritymag.com/2002/apr/cover_casestudy.shtml
 - *Identity Management: The Business Context of Security*, PriceWaterhouseCoopers, January 2002, <http://www.pwcglobal.com/extweb/manissue.nsf/DocID/2019770AA6282B3C85256B4A000ED4C7>.
- Various projects to make identity management span multiple systems on the Internet:
 - The Liberty alliance: <http://www.projectliberty.org/>.
 - W3C P3P Project: <http://www.w3.org/P3P/>.
 - *Identity Management Based On P3P*, Oliver Berthold, Marit Khntopp, January 2001, <http://www.koehntopp.de/marit/pub/idmanage/p3p/>.
 - Security Assertions Markup Language (SAML), <http://www.oasis-open.org/committees/security/#documents>.