Human Processes

Keith Harrison-Broninski
CTO Role Modellers (www.rolemodellers.com)

khb@rolemodellers.com

On the Internet, nobody knows you're a dog

In the last few years, Identity Management (IM) has matured yet remained - for many organizations – as bewildering as ever. Even setting aside technologies and considering only standards, are you confident that your IT department properly understands the relationship between, limitations of, and current direction of OpenID, OAuth, UMA, Open Social, Information Cards, XRDS, SAML, and WS-Trust?

Security professionals are well aware that their field is in rapid flux. For example, the semi-annual Internet Identity Workshop, which focuses on “the use of identity management approaches based on open standards that are privacy protecting”, deals with the pace of change by adopting a unique format:

> After the brief introduction on the first day, there are no formal presentations, no keynotes and no panels. After introductions we start with a blank wall and, in less than an hour, with a facilitator guiding the process attendees create a full day, multi-track conference agenda that is relevant and inspiring to everyone there. All are welcome to put forward presentations and propose conversations.
>
> We do this in part because the field is moving so rapidly that it doesn’t make sense to predetermine the presentation schedule months before the event.
> http://www.internetidentityworkshop.com/blog/

In this Column, I’m not going to try and navigate the shifting sands of IM in detail. Rather, I will take a very high-level view in order to discuss a limitation of **all** current approaches to IM. I will argue that typical IM implementations do not do enough to support the primary occupation of most knowledge workers - collaboration with colleagues. I will explain why this is an important problem, and show how to supplement any IM system in order to fill the gap.

What is Digital Identity?

Let’s start by looking at the commonly accepted definition of digital identity, as stated by Kim Cameron in his seminal paper, “The Laws of Identity”:

> We will begin by defining a digital identity as *a set of claims made by one digital subject about itself or another digital subject.*
> http://www.identityblog.com/stories/2004/12/09/thelaws.html

Cameron goes on to provide some illustrative examples of this definition:

- A claim could just convey an identifier - for example, that the subject’s student number is 490-525, or that the subject’s Windows name is REDMOND\kcameron. This is the way many existing identity systems work.

- Another claim might assert that a subject knows a given key – and should be able to demonstrate this fact.
- A set of claims might convey personally identifying information – name, address, date of birth and citizenship, for example.
- A claim might simply propose that a subject is part of a certain group – for example, that she has an age less than 16.
- And a claim might state that a subject has a certain capability – for example to place orders up to a certain limit, or modify a given file.

The objective of IM as a discipline is to provide technical infrastructure for managing such claims. An overview of some aspects of the current state of the art is provided in a helpful diagram by Eve Maler of Forrester Research:

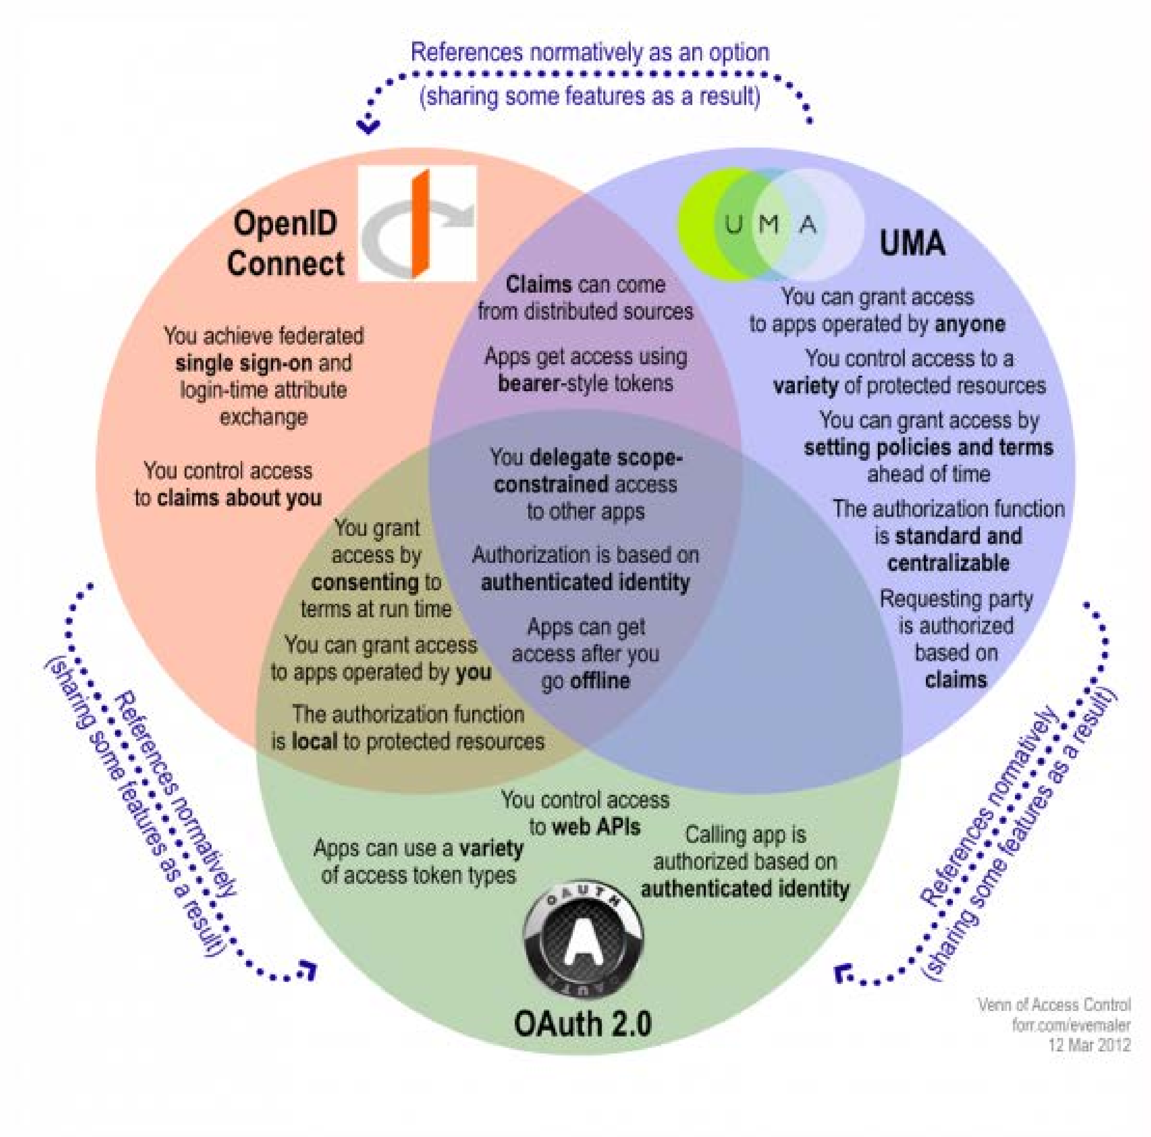
References normatively as an option
(sharing some features as a result)
OpenID Connect
UMA
You achieve federated single sign-on and login-time attribute exchange
You control access to claims about you
Claims can come from distributed sources
Apps get access using bearer-style tokens
You can grant access to apps operated by anyone
You control access to a variety of protected resources
You can grant access by setting policies and terms ahead of time
The authorization function is standard and centralizable
Requesting party is authorized based on claims
You delegate scope-constrained access to other apps
Authorization is based on authenticated identity
Apps can get access after you go offline
You grant access by consenting to terms at run time
You can grant access to apps operated by you
The authorization function is local to protected resources
You control access to web APIs
Apps can use a variety of access token types
Calling app is authorized based on authenticated identity
OAuth 2.0
References normatively
(sharing some features as a result)
References normatively
(sharing some features as a result)
Venn of Access Control
forr.com/evemaler
12 Mar 2012

Figure 1: Venn diagram of access control for the API economy, Eve Maler, http://blogs.forrester.com/eve_maler/12-03-12-a_new_venn_of_access_control_for_the_api_economy

This diagram is high-level, but not high-level enough. It is still a technologist's view that tells us nothing about the nature of the claims that are being managed.

This is deliberate, of course – IM is agnostic of claim contents, and can be used to manage claims of any kind. As Bob Blakley points out in the 2007 OECD Directorate for Science, Technology and Industry Working Paper, 'At A Crossroads: "Personhood" And Digital Identity In The Information Society':

> In discussions of digital identity one often sees "identity" defined as "a collection of attributes" or "a collection of claims" or "partial identity". The question naturally arises, "attributes of what?" or "claims about what?" or "partial identity of what?" These questions have no clear-cut answers. The answer "a subject" is too vague; the answer "a person" leads to arguments about how one distinguishes a natural person from a legal/juridical person, and perhaps even from a non-person; there may also be arguments about whether multiple persons can inhabit the same body, either at the same time or sequentially. And the answer "a human body" creates confusion about matters of intent and continuity of memory, which are important when making decisions about reward and punishment.
> "Current Conceptions Of IDM", p.39

This is heavy stuff – too heavy for the IDM industry, which deals with the issue like Monty Python's Brave Sir Robin, by bravely running away:

> In most practical cases it simply does not matter whether persons have immutable core identities; it is usually sufficient to answer an easier question: "Is this person the same person who did 'X' in the past?" While the (possible) lack of a core identity can create a small amount of doubt about whether our suspect previously did 'X', the evidence for or against the identification of the suspect is normally strong enough to justify confidence in the identification.
>
> This paper takes no position on the question of whether persons have core identities, but it does take the position that core identity is not observable by parties other than the subject himself – so identification systems need to operate on the basis of recognizing attributes or establishing the truth of claims.
> (ibid.)

This cutting of the Gordian knot may suit IM researchers by reducing their scope to something more manageable, but is unhelpful for organizations. If you are going to grant access to business-critical resources, you need to know **why** you are granting access and **what** will be done with those resources.

In other words, you need to understand the **business process context** in which access is being granted – and if, as is usual, you are granting access to a person rather than to an automated agent, this means understanding the work item that has caused the person to request access. In **Human Interaction Management** (HIM) terms, you need to understand:

- The **Activity** the person is carrying out;
- The **Role** they have been assigned to which the Activity belongs;
- The **Plan** to which the Role belongs.

Granting Access to Roles

The key element of identity here is the Role. In a workplace context, someone generally requests access to a resource not because of who they are (**Jane Smith**) but rather because they have been asked to play a Role in a Plan (**Quality Assurance** in **Social Care Transformation Programme 2013**) and hence to need to do a specific Activity (**Review Workforce Training Needs**). It is only on being assigned to the QA Role that Jane needs to review the documents – and if she is re-assigned, goes on leave, or is too busy, someone else will do it in her place.

Roles tell us how to think properly about confidentiality - something we rarely do, and something that is not encouraged by current approaches to IM. We only need to consider how the circulation of paper documents was traditionally controlled to understand the lack of thought that has carried through into current IM.

The classic approach to safeguarding the contents of a document is to use a protective marking – to write RESTRICTED, COMMERCIAL IN CONFIDENCE or CONFIDENTIAL AND PROPRIETARY TO XYZ CORPORATION somewhere on each page. However, in most cases this is an almost meaningless gesture, providing little guidance as to who is empowered to view the document – and to who will actually get to see it.

In practice, the copy list of a document is generally used mainly to determine who gets it initially. Each recipient may then need to copy it to others in their company, project or team in order to carry out the actions that arise from the document. Subordinates may be delegated tasks that require use of the document, or may read it while looking through files out of interest. Excerpts may be taken and sent to suppliers, copied to managers, or used in status reports. Moreover, there are people who see the document not as a result of any direct business need, but as part of processing it on behalf of others—system administration staff, proofreaders, secretaries, and so on.

Executives may not wish to admit it publicly, but this spread of information is almost impossible to manage. Protective markings provide no basis for structured control, since they deal only with the first of the “3 A’s” of security (Authentication, Authorization and Accounting) – yet protective markings are effectively the foundation of modern IM systems, which are generally used to handle “claims” such as “Jane Smith is currently a member of the Business Improvement Team”. What we really need to know is **why** Jane wants to see the document, and **what** she intends to do with it – claims with business meaning that can be verified after the fact.

To provide such information, it is necessary to appreciate that it is not the document that should carry the marking. A single document may be used in different ways at different times by different people. Further, the usages of a document will change over the life of a document – they cannot be predicted and set in stone at the time of writing. To **authorize** and **account** for the usages of an information item, access must be controlled by associating the document with specific Activities in specific Roles in specific Plans.

RBAC is back

The use of Roles as a guiding principle for security policy creation has a long and venerable history. Ross Anderson, in his authoritative 2001 survey of computer system security techniques, wrote:

> “The policy model getting the most attention at present from researchers is role-based access control (RBAC), introduced by David Ferraiolo and Richard Kuhn ... This sets out to provide a more general framework for mandatory access control than [Bell-LaPadula] in which access decisions don't depend on users' names but on the functions they are currently performing within the organization. Transactions that may be performed by holders of a given role are specified, then mechanisms for granting membership of a role (including delegation). Roles, or groups, had for years been the mechanism used in practice in organizations such as banks to manage access control; the RBAC model

> starts to formalize this. It can deal with integrity issues as well as confidentiality, by allowing role membership (and thus access rights) to be revised when certain programs are invoked. Thus, for example, a process calling untrusted software that had been downloaded from the Net might lose the role membership required to write to sensitive system files."
> Anderson, R., 2001, "Security Engineering: A Guide to Building Dependable Distributed Systems," Wiley Computer Publishing, p.145

However, it is not enough to define a Role using the sort of text attributes that Anderson suggests – membership of a group, basically – since "Roles" thus defined have no process context. The fundamental nature of a Role in the modern, flexible, distributed workplace is as *process participant*. Hence a Role definition for IM purposes must include essential aspects of collaborative process management, such as:

- Goals
- Responsibilities
- Information resources
- Structured, purposeful communication channels with colleagues
- Well-defined Activities, including the ability to start and manage other Plans
- A means to manage the status of Deliverables.

Critically, it is necessary to support the dynamic nature of a business process, which means that a Role must not only be able to modify the structure and contents of its own information resources, but also to adapt to circumstances by changing its own behavior and that of others.

Further, there are serious security questions related to *delegation of authority* that can only be solved via a process-based implementation of Roles. It is common for one process participant to offload a particular piece of work to another, perhaps adding a new person to the process specifically for this purpose. With process support for Roles, this can be done either by creating a new Role to which work is passed, or by re-assigning the current Role to a new person, or by starting a sub-Plan involving the new person, or by various other means.

Finally, people typically collaborate with colleagues in different teams or organizations, which means that participants in a single Plan may use different systems to do so. In other words, the Plan must live in several places at once, and part of the function of an IM system is to manage this. A Plan that is distributed in this way cannot be managed using a process system that requires all process participants to login to a central server – rather, different people will use different process systems, which talk to each other in order to synchronize the copies of the Plan held by the different participants. Some participants may never even login to a process management system at all, preferring to use simpler messaging technology such as email to collaborate with colleagues – yet this must also be handled transparently.

To handle Roles in this way, you need a **Human Interaction Management System** (HIMS) – and it is only by the integration of HIMS and IM technology that access to information resources in the workplace can be properly secured.

Conclusion

The new generation of IM systems aims to create an "Identity Metasystem":

1. A way to represent identities using **claims**. Claims are carried in security tokens, as per WS-Security.
2. A means for identity providers, relying parties, and subjects to **negotiate**. Dynamically negotiating the claims to be delivered and the security token format used enables the

Identity Metasystem to carry any format of token and any kinds of claims needed for a digital identity interaction. Negotiation occurs using WS-SecurityPolicy statements exchanged using WS-MetadataExchange.

3. An **encapsulating protocol** to obtain claims and requirements. The WS-Trust and WS-Federation protocols are used to carry requests for security tokens and responses containing those tokens.
4. A means to bridge technology and organizational boundaries using **claims transformation**. **Security Token Services (STSs)** as defined in WS-Trust are used to transform claim contents and formats.
5. A **consistent user experience** across multiple contexts, technologies, and operators. This is achieved via Identity Selector client software representing digital identities owned by users as visual **Information Cards**.

http://en.wikipedia.org/wiki/Information_Card#Components_of_the_Identity_Metasystem

However, the Identity Metasystem may do more harm than good if the claims that it exchanges cannot be properly authorized or accounted for.

Physically, confidential information may be **stored** in operating system files or database records, but logically, it **belongs** to the Roles that need it. Only by providing access to information through the mediation of process Roles will organizations be able to properly authorize and account for that access. Without Roles, access can only be granted to a bewildering sea of names – and in the digital world, it is no longer possible to associate names with faces, as in Peter Steiner's famous cartoon:

An enterprise infrastructure without an access control layer based on HIM Plans is shown in Figure 2:

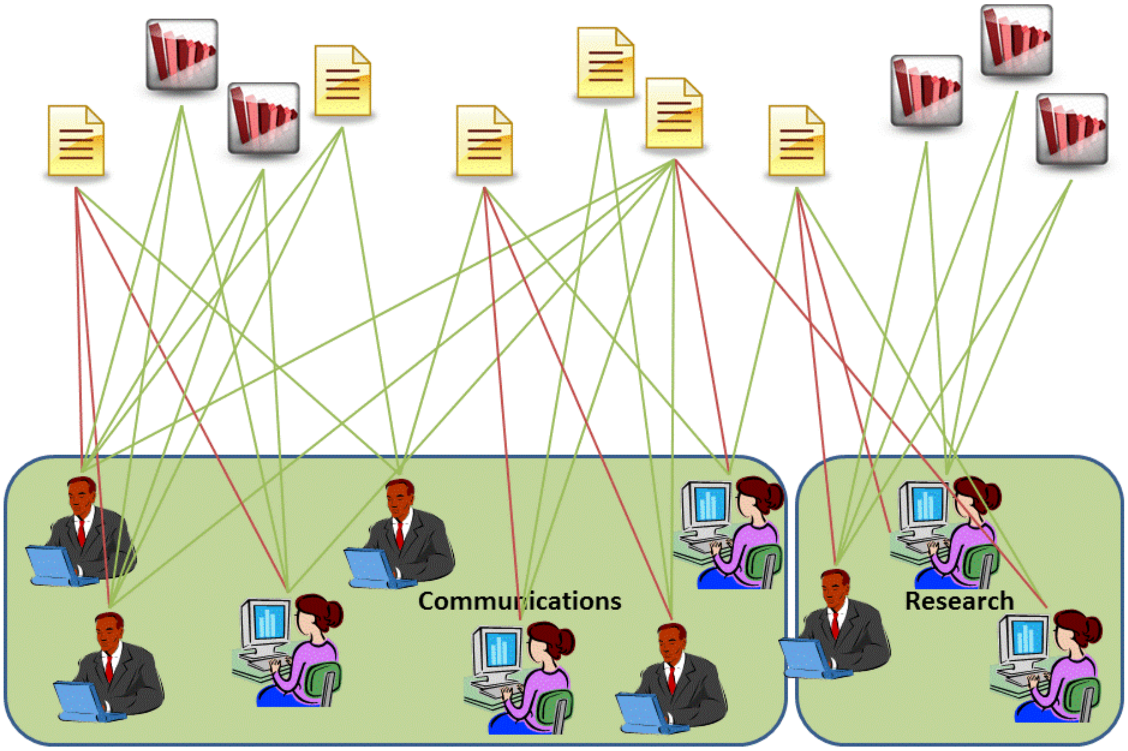

Figure 2: Identity "Management" without HIM Plans

Without some means of understanding why people need access to resources, and what they are doing with them, the situation is next to impossible to manage or audit. By contrast, Figure 3 shows how the picture simplifies if you introduce HIM Plans as an intermediate layer to facilitate authorization and accounting:

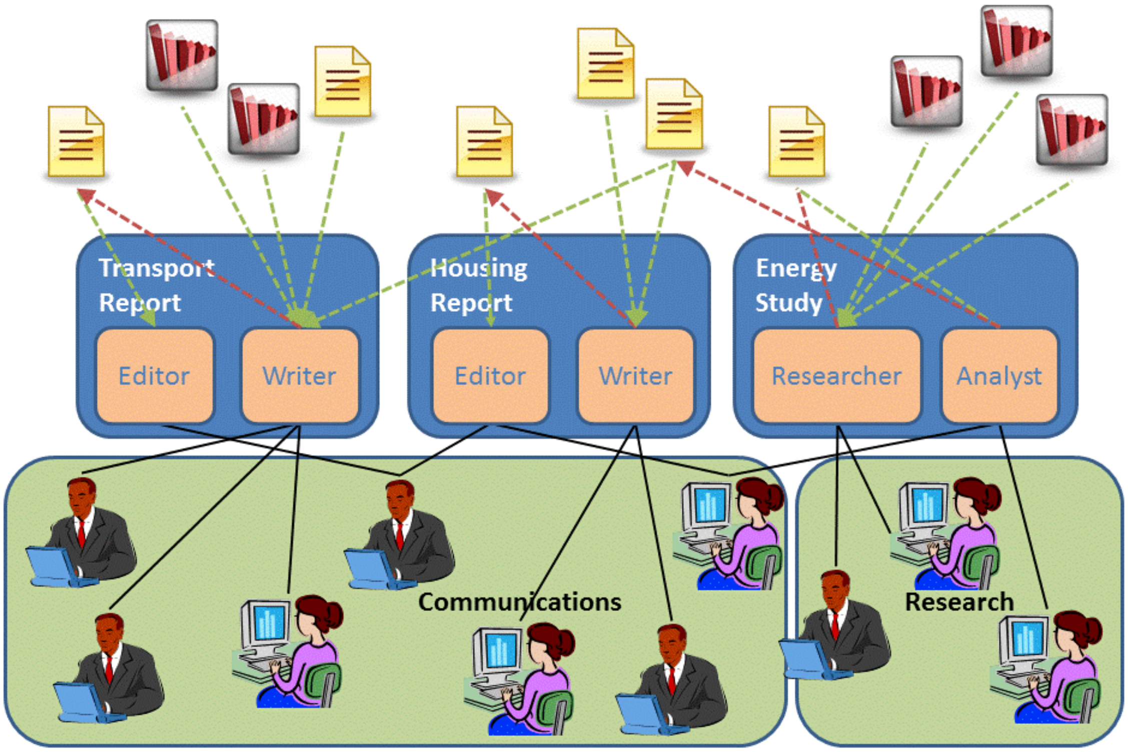

Figure 3: Identity Management with HIM Plans

Further, if access control is guided by Role assignment in HIM Plans, then sharing and delegation of access also falls under IM control. With Roles, data sharing becomes *conscious*, not accidental, and *tracked*, not invisible. This is not only because Roles mean something in a business context, but also because Roles exchange information via structured, purposeful communication channels. If a HIMS is used to manage workplace human interactions, messages are sent in context (by a **Role** as part of a **Stage** in a **Plan**) and automatically recorded for auditing. So not only is access authorized and accounted for, but also transfer of that access – which as we saw above, is a normal, fundamental part of the way that people use information.

Similarly, HIM Plans allow provision of access to people from multiple organizations to be managed without need for complex cross-boundary approval procedures. The organizational affiliation of each Role player becomes secondary to their assigned Role in a Plan of action that has already been agreed by all parties concerned. Organizations that implement IM in combination with HIMS technology will be able to open up their systems to partners safely and manageably. Other organizations may only find out long after the fact that their data is no longer confidential – and still be unable to work out what happened.

Author

Keith Harrison-Broninski is CTO of Role Modellers, a Gartner BPM Cool Vendor 2012. The company mission is to develop understanding and support of human-driven processes - the field that Keith pioneered. Its software product, the Human Interaction Management System (HIMS) **HumanEdj**, provides unique software support for collaborative, adaptive human work.

Keith has been regarded as an IT and business thought leader since publication of his 2005 book "Human Interactions: The Heart And Soul Of Business Process Management". Building on 20 years of research and insights from varied disciplines, his theory of Human Interaction Management (HIM) provides a new way to describe and support collaborative human work. Keith speaks regularly about HIM and the associated change management methodology Goal-Oriented Organization Design (GOOD) in keynotes to business, IT and academic audiences at national conferences, most recently in Poland, India, the Netherlands, the UK, Finland and Portugal.

More information about HumanEdj is available at www.rolemodellers.com and about Keith at http://keith.harrison-broninski.info.

BPTrends Linkedin Discussion Group

We created a BPTrends Discussion Group on Linkedin to allow our members, readers and friends to freely exchange ideas on a wide variety of BPM related topics. We encourage you to initiate a new discussion on this publication, or on other BPM related topics of interest to you, or to contribute to existing discussions. Go to Linkedin and join the **BPTrends Discussion Group.**