

Access Control Requirements for Web 2.0 Security and Privacy

Dr. Carrie E. Gates
CA Labs, CA, Islandia, NY 11749
carrie.gates@ca.com

The increased social networking capabilities provided by Web 2.0 technologies requires a review of what we consider “private” and what we consider “personal” information, and will subsequently drive a new way of limiting and monitoring the information that we release online. Web 2.0 applications are creating large, complex conglomerations of personal data and so we need new approaches to describe and execute access control on that data.

“Private” information currently tends to be loosely defined by legislation, rather than by what individuals consider to be personal. Generic information such as a person’s home address and phone number are normally considered personally identifiable information (PII) and are to be protected when collected and stored by an organization – additionally, the use and release of specific data, such as financial or medical information, is controlled legislatively. However, there also exists information that an individual may consider to be *personal*, and want to release only to particular people (such as close friends) or people meeting a particular criteria (such as people attending the same school). Thus a person might want to control portions of their digital life in the same manner that they control what information is released in their analog life. In the analog world, a person can choose to tell someone or some group some piece of information about themselves. However, it is often the case that in the online world these controls do not exist, leading to de facto public disclosure.

Approaches, such as password protection, have nearly always been available for standard web pages, blogs, webmail, and bulletin boards. However, as aspects of Web 2.0 continue to be adopted, the ability to protect information *within* the same page will be required. For example, a blogger might maintain a single blog, but wish to control access to particular entries based on the reader’s relationship to the blogger. The ability to perform this type of fine-grained access control will not only become essential in the world of Web 2.0, it will largely determine the success or failure of many social, political, and economic realms in the Web 2.0 world.

As automated tools become available and more popular, this kind of access control must “follow” content. We don’t want to re-invent a digital rights management scheme - we understand that in the digital world, copying content is simply a given. Where we need to go with this approach is to hinder *inadvertent* disclosures and aggregations of data: the case where a person who has access to particular content inadvertently makes this content available to others.

These new forms of interactions generate new technical requirements, particularly regarding access control mechanisms. The following four requirements are key to developing a system that addresses the issues:

1. *relationship-based*: Previous access control models have been based on business interactions (e.g., role-based access control [3], coalition-based access control [1]), however a new paradigm of access control needs to be developed that is based on interpersonal relationships — relationship-based access control (ReBAC). Thus the data owner can control the release of their personal information in the same manner he would control it in the analog world — based on their relationship with the data receiver rather than the receiver’s role. One result is that people can hold multiple relationships with someone (e.g., both sister and close friend), and can even be present in what might be considered to be conflicting relationships (e.g., a mother might generally be considered to be a friend, yet a daughter might not want to reveal everything she reveals to her friends to her mother as well). Some social networking sites, such as FaceBook (<http://www.facebook.com>), have started to develop these forms of control, however the relationships that they can represent are still limited. For example, FaceBook allows a user to join various “networks” (e.g., home university, home city) and control what information is released to each network. Further, a user can specify if a particular person should be “limited” from seeing particular material or blocked entirely from seeing any material. However, a user cannot define their own relationship groups (such as family or close friends).
2. *fine-grained*: In addition to being relationship-based, access control must also be available in a fine-grained format, protecting, for example, individual blog entries, particular personal information, or specific photos. A user may even wish to protect specific words or phrases within an article, such as desiring the ability to redact someone’s name so that only particular people can see it. While some work has been done in this area (for example, FaceBook allows a user to state what networks can view particular personal information such as relationship status or favourite activities), it is not available regardless of the social network tool employed nor does it necessarily provide the granularity that a user might want.
3. *interoperability*: Users may access many different Web 2.0 sites, showing no favour to any one of them. For example, a user might have accounts on FaceBook, Flickr (<http://www.flickr.com/>), LinkedIn (<http://www.linkedin.com/>) and Blogger (<http://www.blogger.com/>). Thus any access control system that is developed should be interoperable between the multiple sites. Ideally, the access control policies and relationship groups defined by the user should *follow the user*, rather than be redeveloped for each individual site. There are two possible approaches to meeting this requirement: (1) a central access control site is created that can be used by all Web 2.0 sites, or (2) an easy mechanism for exporting access groups from one site to another is provided.
4. *sticky policies*: Policies that are created by a user need to not only be interoperable amongst the different sites, but should also follow the data to which they apply. This concept was introduced by Mont *et al.* [2] as “sticky policies”, where the access policy was associated with the data. However, their solution has strong requirements for the underlying software and hardware architecture, requiring, for example, TCPA. If these requirements are relaxed, it may be possible to deploy a system that provides policies that, while it will be possible to intentionally subvert them, will at least prevent inadvertent disclosure of personal information.

We need to develop new paradigms for protecting privacy in this evolving environment. We currently have the opportunity to envision security as a user-driven process; one that enables them to choose what information they consider to be personal and to protect it.

References

- [1] Eve Cohen, Roshan K. Thomas, William Winsborough, and Deborah Shands. Models for coalition-based access control (cbac). In *SACMAT '02: Proceedings of the seventh ACM symposium on Access control models and technologies*, pages 97–106, New York, NY, USA, 2002. ACM Press.
- [2] Marco Casassa Mont, Siani Pearson, and Pete Bramhall. Towards accountable management of identity and privacy: Sticky policies and enforceable tracing services. In *Proceedings of the 14th International Workshop on Database and Expert Systems Applications*, pages 377–382, 2003.
- [3] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinnsstein, and Charles E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38 – 47, 1996.