

# More Content - Less Control: Access Control in the Web 2.0

Michael Hart, Rob Johnson, Amanda Stent  
{mhart, rtjohnso, stent}@cs.sunysb.edu

## 1 Introduction

The popularity of social-networking sites, blogging and other content-sharing sites has exploded, resulting in more personal information and opinions being available with less access control than ever before [5]. Many content-sharing sites provide only the most rudimentary access control: a document can be either completely private or completely public. Other sites offer the slightly more flexible private/friends/public access-control model, but this still fails to support natural distinctions users need, such as separating real-world friends from online friends. The traditional response to these privacy concerns is to post anonymously or pseudonymously, but recent psychological research shows that some Internet users do not establish separate, online personae, but instead consider their online identity as an extension of their real-life self [3]. And although privacy expectations that users desire are easy to state, there is a large gap between the users' mental models and the policy languages of traditional access-control systems [2].

The consequences of poor access control are well-documented in the news media. Bloggers have lost their jobs when their employer discovered the employee's personal blog [9]. Sexual predators use social-networking sites to find victims [7]. Bloggers have been stalked based on the opinions and personal information placed on their blog [8]. Universities have disciplined students using photographs published on social-networking sites [1].

For all these reasons, we advocate that blogs and social networks need a policy mechanism that supports high-level policies that can be expressed succinctly, applied automatically, and updated easily. Current access-control systems fail to meet these goals. Users manually enforce and manage their policies, users, groups, and roles of the system. Furthermore, these systems lack intuitive tools and interfaces for policy generation. We propose to solve all these problems by specifying access-control policies in terms of the content being mediated, e.g. "Blog posts about my home-town are visible to my high school friends." The system will then automatically infer the posts that are subject to policy rules based

on the posts' contents. Similarly, the system can infer relationships and interests of the users based on the content of objects they create (see Section 3). Such policies will be intuitive and easy to specify, greatly enhancing usability for non-technical users.

We first discuss the current state of access control on content-driven sites and analyze approaches proposed in literature for implementing access control for the web. We then describe our proposed method of access control for content-sharing sites.

## 2 Current Access Controls

We surveyed 23 blogging and social-networking sites to determine what access control and privacy features are currently available. Our survey included sites like Blogger, Facebook, Flickr, YouTube, and MySpace.

The access control features implemented in these services fell into a few broad categories, with some sites offering minor extensions. Although some systems offer only basic access control (i.e. private/public) or unusual features like search engine invisibility, most employ the *friends* model. The friends model lets users create a list of friends and restrict content to be visible only to this group. Of the existing schemes, this scheme is the most successful because it strikes the best balance between ease-of-use and flexibility.

Despite its success, the friends model has several flaws. First, it does not let users segregate their disparate social groups. There are different degrees of intimacy between an individual and each one of his or her friends, but the friends notion is too coarse to capture these distinctions. For example, we found that MySpace users had a median of 115 friends<sup>1</sup>, suggesting that the friends notion is being stretched to cover a wide range of intimacy levels. Making matters worse, using the friends relation for access control forces users to choose between protecting their privacy and appearing popular, such as in MySpace's "Top Friends" feature. As the notion of friend loses

---

<sup>1</sup>Based on a random sample of 91660 MySpace users with public profiles that had logged-in between September 1st, 2006 and October 23, 2006.

its meaning, friends-based access control also becomes meaningless.

Because of the dynamic, administrator-less nature of social networks and blogs, we cannot solve these problems with traditional access-control schemes. Existing schemes, such as Role-Based Access Control, require an administrator to manage users, assign rights to them and maintain access-control lists on objects. No such administrator exists on social networks. Even if users were willing to perform these tasks, the frequent content updates and volatile nature of friendships would make this task even more difficult than in conventional computer systems.

Some researchers have attempted to solve the distributed access control problem using attribute-based access controls and credentials. The *MaX* project lets content publishers label their content with a list of attributes that viewers must possess to access that content, and viewers must prove that they possess the required attributes by presenting cryptographic credentials. Although this system provides a useful mechanism for making access-control decisions in a distributed environment, it still requires content publishers to manually specify the access-control policy for each object. We propose to automate policy specification as much as possible.

### 3 Usable Access Control

Our goal is to create an access-control system that is usable by non-technical users and can support the dynamic content of blogs, social networks, and other content-sharing sites.

There are essentially two requirements for such a system: usable policy specification and automatic policy application. There are already promising results in the design of usable policy interfaces that address the gap between users' mental models and actual policy. Most notably, Karat explored several intuitive user interfaces for policy acquisition, including natural language input, templates, and guides to facilitate machine-readable policies [4]. The template approach could allow users to succinctly specify natural policies, such as "Allow group *College Friends* to access entries on topic *Parties*." Such policies are short, fit users' mental models, and can be applied broadly to many documents.

Once the policy is specified, we need a mechanism for automatically applying it to existing and new documents based solely on their content, with minimal or no user guidance. The access-control system can implement content-based policies by reducing the content to tags on documents and users. To compute

these tags, we can exploit several light-weight techniques from machine learning and natural-language processing [6]. These methods extract document meta-information, named entities mentioned in the document, and other text phrases in the document that are statistically likely to summarize its content. Based on these document features, we can infer the appropriate tags of a document by comparing it to other documents with similar features and known tags. Such a system will make occasional mistakes. We therefore will also need a good, easy-to-use feedback mechanism for users to correct erroneous tags.

Beyond intuitive policy specification, these techniques lend themselves to new types of access controls not feasible before. One such policy is affinity-based access control, where posts on a topic  $T$  are visible to users that demonstrate sincere interest in  $T$  by, for example, posting repeatedly about  $T$ . This policy is well-suited to social-networking sites where users want to meet like-minded people without exposing their personal lives to anyone that might know them, such as their employer. This mechanism can also implement "Need-to-know" policies of the form, "Only people mentioned in this post may read this post."

In addition to bolstering privacy, content-based policies can help users maintain integrity constraints on their shared content. For example, many bloggers manually maintain the implicit integrity constraint that their blog should not reveal their real-world address. Using natural-language processing and text-summarization, we can build a system that can help bloggers detect when they accidentally violate this rule. Similarly, publicly-maintained knowledge-bases like Wikipedia have explicit rules about what sort of content updates are allowed. Although it is impractical to build a system that checks the veracity of content posted to Wikipedia, we can automatically prevent many other forms of vandalism by verifying that new content does not change the topic or sentiment of the post and is phrased in the form of a statement of fact.

### 4 Conclusion

As the Internet is absorbed into the social fabric of regular life, it must support the privacy expectations that non-technical users bring from their real-life experiences. The access control policies for today's social networking systems are simply not expressive enough, and this has already had increasingly serious social ramifications. Currently, we are working on two projects to address these problems. First, we are developing PLOG, a **P**rivacy/**P**olicy-aware

**bLOG**ging engine. The goal of PLOG is to facilitate access control that is automatic, expressive and convenient. We are also starting a Wikipedia Integrity Project, in which we hope to build systems to detect and deter vandalism and other malicious edits to Wikipedia. We believe this work will extend to other collaborative environments and knowledge-bases.

## References

- [1] Susan B. Barnes. Privacy paradox: Social networking in the United States. *First Monday*, 11(9), September 2006.
- [2] M. Benantar. *Access Control Systems: Security, Identity Management and Trust Models*. Springer, 2005.
- [3] D. A. Huffaker and S. L. Calvert. Gender, identity, and language use in teenage blogs. *Journal of Computer-Mediated Communication*, 10(2), January 2005.
- [4] C. Karat, J. Karat, C. Brodie, and J. Feng. Evaluating interfaces for privacy policy rule authoring. In *Proceedings of CHI 2006*, April 2006.
- [5] A. Lenhart and S. Fox. Bloggers: A portrait of the internet's new storytellers. [http://www.pewinternet.org//pdfs/PIP Bloggers Report July 19 2006.pdf](http://www.pewinternet.org//pdfs/PIP_Bloggers_Report_July_19_2006.pdf), July 2006.
- [6] C. Manning and H. Schutze. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA, 1999.
- [7] Kevin Poulsen. Myspace predator caught by code. <http://www.wired.com/news/technology/0,71948-0.html>, October 2006.
- [8] Darren Rowse. Blog stalkers - personal safety for bloggers. <http://www.problogger.net/archives/2006/02/07/blog-stalkers-personal-safety-for-bloggers/>, February 2006.
- [9] Ellen Simonetti. I was fired for blogging. [http://news.com.com/I+was+fired+for+blogging/2010-1030\\_3-5490836.html](http://news.com.com/I+was+fired+for+blogging/2010-1030_3-5490836.html), December 2004.