Enhancing Tools for Community Privacy

Dennis Kafura, Tom De Hart, Manuel Pérez-Quiñones,
Denis Gracanin, Andrea Kavanaugh
Department of Computer Science
Virginia Tech
{kafura, tdehart, perez, gracanin, kavan}@vt.edu

COMMUNITY PRIVACY

In addition to enhancing the ability of individuals to protect their personally identifying information we believe it also important to enhance tools to protect the privacy of a community. By "community" we mean a group of people sharing sensitive information that they have professional, legal, ethical, or personal obligations to safeguard. By "privacy of a community" we mean the ability of the community to control the confidentiality of the sensitive information that they share. We envision applications enhanced to (1) support the definition and enforcement of community-defined confidentiality over its resources, and (2) interact with other similarly enhanced applications to provide ubiquitous interpretation and enforcement of community-defined confidentiality over their collective resources.

Threats to community privacy include accidental disclosure of sensitive information violating known expectations [1], lack of awareness of expectations to restrict disclosure in a given context [2] , inability to determine if a given disclosure meets the expectations of the community, inability to use the system's interfaces to achieve the correct sharing [3, 4], and inability of the system to guarantee the desired restriction. These threats can lead to significant harms to individuals and the community. One report indicates that the vast majority (80%) of improper disclosures come from inside the organization [5]. There is also evidence that accidents and errors, rather than malicious intent, account for the majority of information leaks [6]. Such inadvertent leaks can be costly, as illustrated by the case of an email error that led to a $1 billion loss [1]. We believe that many, though not all, of these threats can be prevented or ameliorated by community-enhanced applications.

Our preliminary studies [7, 8] shed some light on community privacy. First, we noted many behaviors exhibited "privacy awareness": individuals seeking to limit the disclosure of information sensitive to others in their organization, to their customers, or to the organization itself. The sensitivity resulted from a variety of factors ranging from legal mandates to simple concern for the dignity of their co-workers. Second, the studies revealed various types of "situated disclosure": extraordinary circumstances that motivated the disclosure of sensitive information outside of the usual community. These disclosures are important because they require tools to accommodate spontaneous, ad-hoc, and unforeseen needs for managing community privacy. Third, we saw that three privacy behaviors previously seen in individuals [9] can also be seen in communities. Behaviors termed "avoidance" (taking steps to circumvent situations where privacy is at risk), "modification" (taking steps to reveal as little as possible) and "alleviatory" (preventing the spread of information or reducing its consequences) were identified. While the specific actors in these cases were individuals, the behavior was grounded in the needs and culture of the community.

We have developed a model of community privacy [10] to guide future development of tools for community privacy. The model (Figure 1) is illustrated by this scenario. Bob, a manager, and Alice, in human resources, are in *agreement* to create a two-person *community* for private email communication about a member of Bob's staff. Their agreement results in a *tag* representing their community. When Bob sends a new email to Alice on this subject he annotates the email with the tag. The tag is permanently affixed to the email and all email derived by forwarding and replying. The privacy-enhanced email system enforces the rule that only Bob and Alice can see their tagged emails, thus creating a privacy boundary. At some point Alice needs the advice of another person, Rhonda, in her department and forward's one of Bob emails to Rhonda. This is a privacy violation about which Bob is informed via a *notification*. Realizing the privacy mistake, Alice asks Bob for an *exception* to the privacy boundary explaining her reasons. Through agreement Bob and Alice allow the email to be forwarded to Rhonda.

CMAIL: TESTING DESIGN ISSUES

A user interface for a community-enhanced email application, CMail (Figure 2), was prototyped to evaluate two elements of the privacy model: tags and agreement. Exceptions played an indirect role in the evaluation and notifications were not used. The interface contains facilities to create tags and their associated communities (see the "tags" menu item in Figure 2), to send, receive, and forward tagged and untagged emails (tagged emails in the inbox are shown with lock icons in Figure 2), and to reach agreement on adding or removing members from community tags, and on exceptions. The status of agreement issues can be seen by using the mini-status window at the top of the screen or by viewing the "ballots" menu item in Figure 2.

We chose email for testing because: (1) there is evidence that loss of confidentiality in email communication can have significant adverse economic consequences [1]; (2) email is a familiar tool to a broad cross section of the population; (3), email tagging is also familiar to many people (e.g., in Gmail) for organizing and searching email and, thus, the extension to deal with email confidentiality may be a more natural step; (4) email infrastructure is universal and, thus, there is the potential for this research to have broad impact and adoption; (5) it is possible to build with manageable effort a prototype email system for research purposes; and (6) email exhibits confidentiality concerns that are comparable to or more complex than those in other domains.

A critical design issue is the user's mental model [11] of a system. The user's mental model influences how the user develops plans to achieve a desired outcome, predicts the behavior of the system, and explains the result of the system operation. Whether planned or not, a user interface inherently reflects the designer's understanding of the user's mental model. The collection, organization and naming of user interface elements embodies this understanding. Mismatches between the user's mental model and the model embodied in a user interface are the causes of fundamental usability problems. A number of mental models users have for security have been studied [12]. Our own work demonstrates the difficulty of finding a good mental model for community privacy.

We conducted an IRB-approved study involving 22 users from the university student community and 8 users from a local software development company. In the experiment the participant takes the role of an employee working in the Human Resources department at a fictitious company and performs several email tasks. Data was collected through screen recordings, audio recordings, and post-study questionnaires. Our analysis used both grounded theory [13] and critical incident analysis [14].

Two mental model mismatches were discovered centered on the notion of "tagging":

- *Content vs. confidentiality*. Many participants evidenced a mismatch between a tag used for describing content and a tag to keep emails private. A few participants were able to articulate that this resulted from overriding the meaning of "tag".
- *Inclusion vs. exclusion*. A mismatch occurred between a community tags that excludes non-community members with an email distribution list that includes all members on the list. Many study participants, both students and professionals, considered community tags as a way to send email to many recipients. In this case the sense of "community" was confused with the idea of "group" used in listservs.

In addition to these mental model issues we identified two important problems related to community agreement. In CMail all members of the community must approve decisions on membership and exceptions. The two problems are:

- *Timeliness*: Reaching agreement took too long, especially for time critical tasks. Given their role in a fictitious human resources department, participants argued that tasks like employee termination need to happen quickly.
- *Relevance*: When is agreement truly necessary? In some cases agreement was not relevant when the system model did not match the social structure of a community. In these cases, agreement was seen as entirely unnecessary in the context of their community. In other cases, implicit confidentiality was expected with the need for agreement.

Changes to the user interface to cope with mental model issues could involve different terminology (to avoid overloading "tag') or changing the workflow to separate the email editing process from the privacy annotation process. Changes to the agreement process could involve closer matching to a community structure and agreement practices (e.g., proxies, differential authority), and better handling of exceptions (e.g., allowing emergency disclosures that trigger a notification to the community).

REFERENCES

[1] P. Zilberman, *et al.*, "Analyzing Group Communication for Preventing Accidental Data Leakage via Email," presented at the Proceedings of the 2010 Workshop on Collaborative Methods for Security and Privacy (CollSec'10), Washington, D.C., 2010.

[2] A. Lampinen, *et al.*, "We're in it together: interpersonal management of disclosure in social network services," presented at the Proceedings of the 2011 annual conference on Human factors in computing systems, Vancouver, BC, Canada, 2011.

[3] S. L. Garfinkel and R. C. Miller, "Johnny 2: a user test of key continuity management with S/MIME and Outlook Express," presented at the Proceedings of the 2005 symposium on Usable privacy and security, Pittsburgh, Pennsylvania, 2005.

[4] A. Whitten and J. D. Tygar, "Why Johnny can't encrypt: a usability evaluation of PGP 5.0," presented at the Proceedings of the 8th conference on USENIX Security Symposium - Volume 8, Washington, D.C., 1999.

[5] D. Liu, *et al.*, "Mitigating Inadvertanent Insider Threats via Incentives," presented at the Financial Cryptography and Data Security FC'09), Barbados, 2009.

[6] *Homeland Defense Journal*, 2007.

[7] S. Codio, *et al.*, "A Case Study of Community Privacy," presented at the International Conference on Social Informatics, Washington, D.C., USA, 2012.

[8] S. Codio, *et al.*, "Identifying Critical Factors of Community Privacy," presented at the International Confernce on Privacy, Security, Risk and Trust (PASSAT'12), Amsterdam, Netherlands, 2012.

[9] K. Caine, *Exploring Everyday Privacy Behaviors and Misclosures*. Atlanta, GA, USA: Georgia Institute of Technology, 2009.

[10] D. Kafura, *et al.*, "An Approach to Community-Oriented Email Privacy," presented at the Third IEEE International Conference on Information Privacy, Security, Risk and Trust (PASSAT 2011), MIT, Boston, MA, 2011.

[11] S. J. Payne, "Users' Mental Models: The Very Ideas," in *HCI Models, Theories, and Frameworks*, J. M. Carroll, Ed., ed San Francisco: Morgan Kaufman Publishers, 2003, pp. 135-154.

[12] L. J. Camp, "Mental Models of Privacy and Security," *IEEE Technology and Society Magazine,* vol. 28, pp. 37-46, 2009.

[13] A. Strauss and J. M. Corgin, *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*: SAGE Publications., 1998.

[14] J. C. Flanagan, *The critical incident technique.*: American Psychological Association, 1954.

Appendix

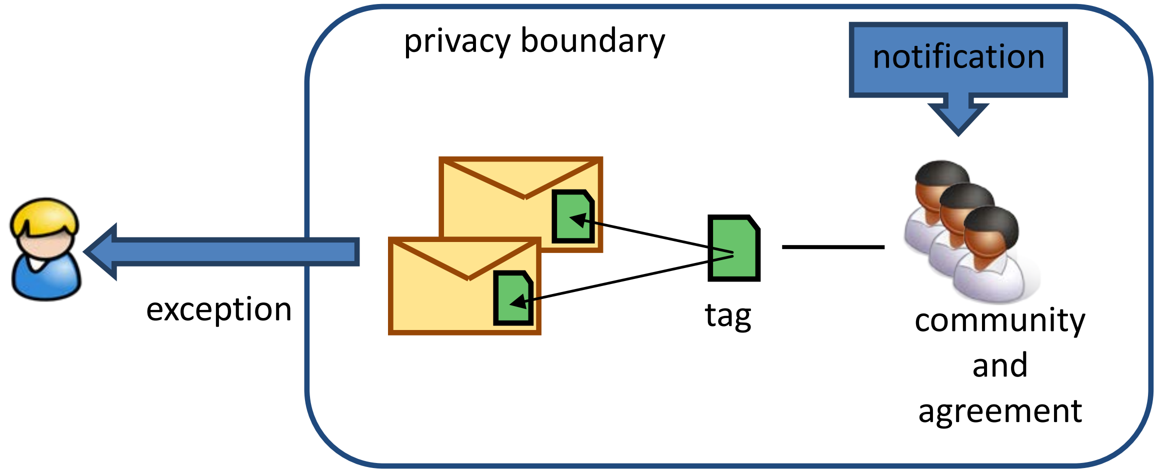

Figure 1: Community Privacy Model

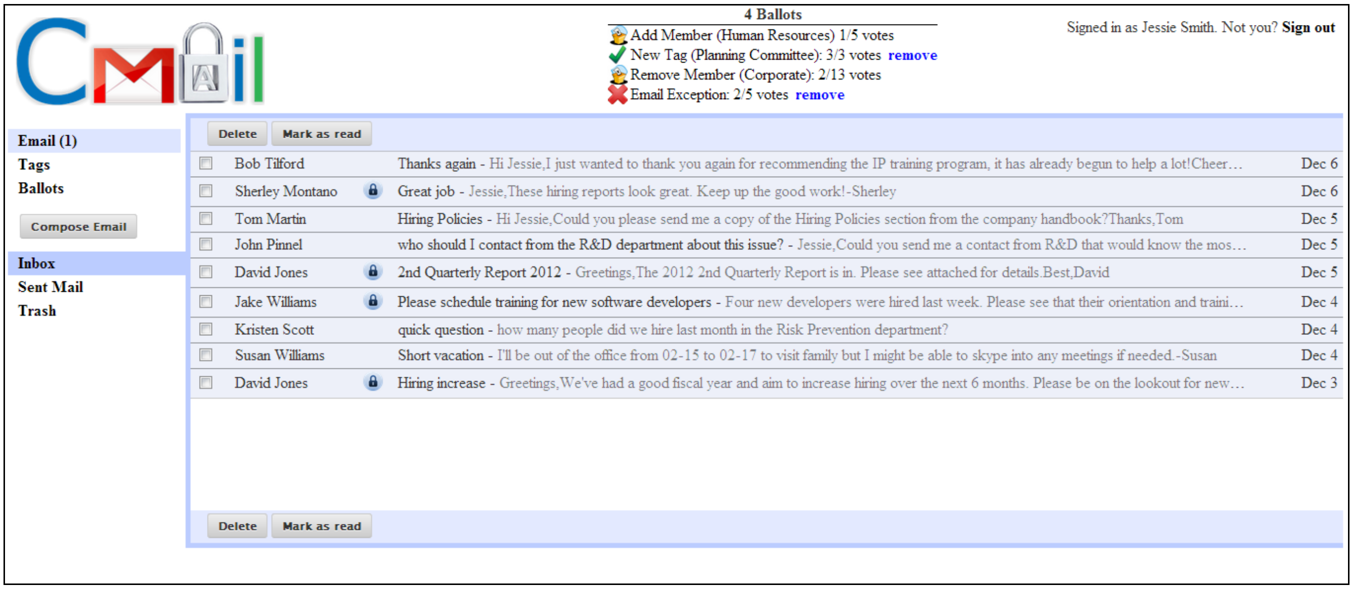

4 Ballots

- Add Member (Human Resources) 1/5 votes
- ✓ New Tag (Planning Committee): 3/3 votes **remove**
- Remove Member (Corporate): 2/13 votes
- ✗ Email Exception: 2/5 votes **remove**

Signed in as Jessie Smith. Not you? **Sign out**

□	Bob Tilford		Thanks again - Hi Jessie,I just wanted to thank you again for recommending the IP training program, it has already begun to help a lot!Cheer…	Dec 6
□	Sherley Montano	🔒	Great job - Jessie,These hiring reports look great. Keep up the good work!-Sherley	Dec 6
□	Tom Martin		Hiring Policies - Hi Jessie,Could you please send me a copy of the Hiring Policies section from the company handbook?Thanks,Tom	Dec 5
□	John Pinnel		who should I contact from the R&D department about this issue? - Jessie,Could you send me a contact from R&D that would know the mos…	Dec 5
□	David Jones	🔒	2nd Quarterly Report 2012 - Greetings,The 2012 2nd Quarterly Report is in. Please see attached for details.Best,David	Dec 5
□	Jake Williams	🔒	Please schedule training for new software developers - Four new developers were hired last week. Please see that their orientation and traini…	Dec 4
□	Kristen Scott		quick question - how many people did we hire last month in the Risk Prevention department?	Dec 4
□	Susan Williams		Short vacation - I'll be out of the office from 02-15 to 02-17 to visit family but I might be able to skype into any meetings if needed.-Susan	Dec 4
□	David Jones	🔒	Hiring increase - Greetings,We've had a good fiscal year and aim to increase hiring over the next 6 months. Please be on the lookout for new…	Dec 3

Figure 2: Example of CMail User Interface