

Access Control for Home Data Sharing: Attitudes, Needs and Practices

Michelle L. Mazurek, J.P. Arsenault, Joanna Bresee, Nitin Gupta, Iulia Ion[†], Christina Johns, Daniel Lee, Yuan Liang, Jenny Olsen, Brandon Salmon, Richard Shay, Kami Vaniea, Lujo Bauer, Lorrie Faith Cranor, Gregory R. Ganger, Michael K. Reiter[‡]

Carnegie Mellon University
Pittsburgh, PA

[†]ETH Zürich
Zürich, Switzerland

[‡]University of North Carolina
Chapel Hill, NC

ABSTRACT

As digital content becomes more prevalent in the home, non-technical users are increasingly interested in sharing that content with others and accessing it from multiple devices. Not much is known about how these users think about controlling access to this data. To better understand this, we conducted semi-structured, in-situ interviews with 33 users in 15 households. We found that users create ad-hoc access-control mechanisms that do not always work; that their ideal policies are complex and multi-dimensional; that a priori policy specification is often insufficient; and that people's mental models of access control and security are often misaligned with current systems. We detail these findings and present a set of associated guidelines for designing usable access-control systems for the home environment.

Author Keywords

Access control, Home computing, Privacy, Security

ACM Classification Keywords

D.4.6 Security and protections: Access controls; H.3.2 Information Storage: File organization; H.5.2 Information Interfaces and Presentation: User-centered design

General Terms

Human factors, Security

INTRODUCTION

Digital content is increasingly common in the home, as new content is created in digital form and people digitize their existing content. Devices such as digital cameras, mobile phones, and portable music players make creating and interacting with this content easy. Home users are increasingly interested in sharing content inside and outside their homes, across computers and other devices [1, 7]. New systems created to meet this need are converging toward a global file-

system model that will make home-centered data potentially available anywhere [10, 16, 21].

In this environment, it may be difficult to provide secure, usable *access control*—the ability to control which people and devices are able to read and change files. Studies repeatedly show that computer users struggle with specifying access-control policies [13, 22]. Home users are often technically inexperienced and impatient with complex interfaces. Large organizations have system administrators to set up and maintain access-control policies, but home users typically have only themselves, family members, and friends [15].

Not much is known about how people think about and interact with access control in the home [5]. It is not yet known how much or what kind of access control is required in order for home-centered data-sharing systems to be usable while providing the protections users desire. As a first step, we conducted semi-structured interviews with 33 non-technical computer users in 15 households to broadly examine their current access-control attitudes, needs and practices. We also discussed hypothetical scenarios with them to understand what their needs and preferences might be in a world where sharing digital files is routine and ubiquitous.

Key findings

Our interview data led to several findings. First, we found that people construct a variety of ad-hoc access-control mechanisms, but that these do not entirely allay their concerns about protecting sensitive data. Second, we found that people's ideal access-control policies can be complicated; they are not always defined exclusively in standard role-based terms but can also incorporate factors like who is present, where the access occurs, and what device is being used. Third, we found that a-priori policy specification is often insufficient, because it does not align well with social models of politeness and permission. In addition, many participants expressed a desire to update their policies iteratively in reaction to data access requests. Fourth, we found that people's mental models of access control and of computer security in general are often misaligned with current system designs in ways that could leave users vulnerable. From these findings, we distill a set of guidelines for designing usable access-control systems for digital data in the home environment.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10 – 15, 2010, Atlanta, Georgia, USA

Copyright 2010 ACM 978-1-60558-929-9/10/04...\$10.00.

Background and related work

Family dynamics and social norms are important in the home context. When working in small groups, as within a household, people often establish social rules that allow them to function without tight security [2]. Previous studies show that home users trust other members of their households and expect them not to pry beyond clearly marked boundaries. Instead of using technology to protect their files, users hide files or store them on devices identified as off-limits to others [20]. Little et al. found that family dynamics play an important role in household information control. Users expressed concern about the effects of ubiquitous technology on family social balance [12]. Our work expands on these ideas by focusing on home users’ current practices as well as future needs for access control.

Studies of academic and corporate environments have found that users have dynamic access-control policies that can change quickly [3, 8, 17]. Dourish et al. also found that younger office workers tended to articulate more complex security needs [8]. Razavi and Iverson found that students using a personal learning space had a strong need for privacy controls [17]. As documents moved through a life cycle, their privacy needs changed too. Users liked sharing some finished documents with wide audiences but wanted to control the visibility of private and work-in-progress documents. Users also found that managing access controls to match these preferences was too labor intensive. In a lab study, Olson et al. explored how comfortable people were sharing different types of data (including information such as age and salary, as well as digital files) with different types of people. They found that both people and data generally clustered into a small number of intuitive groups based on interpersonal trust relationships. However, they recommended that interfaces allow users to choose the granularity at which to specify access control and allow for exceptions [14].

Other studies have examined the use of computer accounts at home. Unlike in an enterprise environment, where each person has her own account, home users tend to share a single account on the family computer, which obscures the connection between users and data accesses [6, 20]. Account sharing is primarily driven by convenience; the ability to quickly access the computer outweighs the privacy and security concerns that can be addressed with multiple accounts [9].

During our study, we asked participants to think about using a *reactive policy-creation system*, which allows a file or resource owner to make a semi-real-time policy adjustment in response to an attempted access that cannot otherwise succeed. For example, in a deployment of Grey [4], users send messages with their cell phones to the “owners” of doors they wish to open in a university building. Door owners can respond with their own cell phones and grant access. Bauer et al. found the Grey system enabled users to construct policies closer to their ideal policies than did keys [3].

METHODOLOGY

We gathered data using semi-structured, in-situ interviews to increase our understanding of how home users think about

household	size	interviewees	laptop	desktop	storage	phone	mp3	DVR	dig. camera	other
C1	2	28m banking supervisor; 26f Spanish teacher	2			2				2
C2	2	27m forensic toxicologist; 28f healthcare recruiter	1	1		2	1	1	2	4
C3	2	27m student; 27f research assistant	1	1	1	2	2		2	1
C4	2	53f, 59m retired		1		2			1	
C5	2	27m police academy student; 24f nurse	1	1	1	2	2		2	1
F1	3	30m driver; 33f research associate	1			2	1	1	1	2
F2	4	53f pastor; 22f student	5	2	11	5	5		1	
F3	5	43f accountant; 9m student	4	1		2	2	1	2	
F4	3	50f admin. assistant; 18m, 15m students	1	1		3	2	2	1	3
F5	4	44m curator; 43f printer; 12f, 9f students		1	3	1	2			
F6	4	36m systems analyst	3	2	2	2	2	2	2	
R1	7	20f, 21f, 20f students	7		8	4	4		4	1
R2	2	25f, 26f students	2		4	2	2		2	
R3	4	20m student	4		2	3	2		2	1
R4	3	18f student; 24m video producer; 22m salesman	3	2	5	3	2		2	2

Table 1. Participant demographics. The columns identify household code; number of part- or full-time occupants; age, gender, and occupation of each interviewee; and number of household devices.

controlling access to their digital files. In order to address a broad sweep of possible scenarios, we decided to include any devices that reside at least part time in the home as well as any files that currently are or someday may be stored on these devices. We also considered possible access by anyone who might be able to use or connect to any of these devices, now or in the future: friends, family, colleagues, and even strangers. We did not start out with any hypothesis; instead, as we conducted and analyzed interviews iteratively, we developed theories about home users’ preferences, needs, and mental models for access control.

Participants

We recruited Pittsburgh-area participant households through Craigslist posts, emails to university distribution lists, and flyers hung at grocery stores and distributed to families at children’s soccer games. Households were prescreened to include those with a range of digital devices storing at least a moderate amount of personal data, but to exclude households that included computer programmers. We interviewed 33 people, ranging from elementary-school students to retirees. Participant households, which included five couples, six families with children, and four sets of roommates, were each paid \$50.

In this paper, we refer to participants using a naming scheme that identifies their household type (C for couples, R for roommates, or F for families), household number within that type, and member letter. For example, participant R2A belongs to the second roommate household (R2) and is the first member of that household to be interviewed (A). Table 1 lists some demographic information about the participants.

Households had different numbers of devices, ranging from

seven devices for four people (F5) to 29 devices for four people (F2). The most common devices included laptop and desktop computers, mobile phones, music players, DVRs, external storage devices, digital cameras, and video game systems. Many computers had passwords, and some had separate login accounts; PINs were also used on a few phones and music players and one DVR (for parental control), but not other types of devices. Overall, about one third of devices were primarily shared, while the rest were primarily for individual use. In general, couples and families shared more devices than roommates, but specific ways of sharing varied significantly.

Interview protocol

We conducted semi-structured, in-home interviews. We interviewed household members first as a group and then individually. All interviews were structured around a predetermined set of questions designed to cover a wide range of access-control-related topics. The questions were intended to encourage participants to discuss their past experiences along with their current behaviors, thoughts, and feelings. If a participant mentioned an interesting topic not in the questions, the interviewer probed further, but otherwise kept to the question list. At least two interviewers attended each session. The interviews were recorded, resulting in more than 30 hours of videotape.

Group interviews lasted approximately 30 minutes each and included all available household members. In these sessions, we asked participants how they currently protect important information both on paper and digitally. To guide participants' thinking, we asked them to draw maps of their homes and illustrate which devices and rooms they considered to be public or private. We also asked about current formal and informal policies for who can use which devices under which circumstances.

Individual interviews lasted 30 minutes to 1.5 hours per participant. The goal of the individual interviews was to understand how participants define their ideal access-control policies and what features they would find useful to implement desired policies. The interview protocol had three major components. First, we asked participants to describe past experiences when they were concerned that others might view or modify data in an unwanted way. This section was used to focus participants on why and when access-control policies would be important to them. Second, we used their home map to walk participants through a list of the types of digital data they own, asking generally about which types are more private and which types are more public. This list was used to guide the rest of the interview. Third, we asked participants to imagine that their data could be available to anyone, from any device, anywhere, and at any time. We presented ten scenarios in order to learn whether and to what degree participants would find various dimensions of access-policy definition useful. The scenarios tested policy-specification factors including: the identity and location of the accessor, the device used for access, whether the file owner is present during the access, the time of day access is attempted, the location of the file owner, and the incidence of social events.

We also asked participants to consider policy-management mechanisms including privacy indicators, a detailed access log, and reactive policy creation.

We prompted participants with specific events and people in an attempt to discern their general attitudes toward specific access-control mechanisms. For example, in order to assess whether participants wanted to restrict access based on person, we picked two people that the participant had mentioned—a close friend or family member and someone they were not close with—and asked: “Imagine _____ could view all of your files and data. What would you not want them to see or change?” We gauged the strength of these preferences by asking participants how upsetting a violation would be, using a five-point Likert scale ranging from ‘don’t care’ (1) to ‘devastating’ (5). We noted the scenarios that resonated with participants and elicited strong examples.

Data analysis

As we completed early interviews, we recorded whether each participant was very interested, somewhat interested, or uninterested in each axis of control, along with high-level explanations of their responses. This course-grained analysis revealed interesting patterns that helped focus later sessions.

Once interviews were complete, we iteratively coded each interview. The first round of coding was collaborative, with team members working together to validate each other’s results. We transcribed all the videos and then applied a more detailed process of topic and analytic coding [18, 19]. A single coder recorded fine-grained codes for each aspect of participants’ answers in a searchable database designed for easy cross-referencing of participants and topics. As new concepts emerged, the coder revisited previously analyzed transcripts to see how the new concepts related. Team members worked together to group individual codes into increasingly larger consensus categories. Using this process, we were able to formulate broader theories about participants’ access-control concerns, needs, and preferences.

Our results are purely qualitative. We sometimes report the number of participants who fall into a given category to provide context; this is not intended to imply statistical or quantitative significance. We have selected the most interesting or salient quotes and anecdotes to embody each point; these examples are meant to be illustrative but not exhaustive. All the results that we report are substantiated by data from several interviewees.

PEOPLE NEED ACCESS CONTROL

Unsurprisingly, we observed that people have data they consider important or sensitive and want to ensure this data is protected. We discuss this result here for completeness, as well as to shed light on specific concerns participants raised and on the sometimes-surprising ways in which they accomplish their access-control goals. In the first subsection, we demonstrate that people have data they classify as sensitive, and they find the idea of unauthorized people accessing this data disturbing. Next, we provide evidence that these concerns are not just hypothetical; several participants described

incidents where their data was put at risk or exposed. Finally, we discuss ways that people construct their own ad-hoc access-control mechanisms using both technical tools and social norms.

People have data they classify as sensitive

Almost all participants want to limit access to their personal data. When we asked participants to imagine a breach of their ideal policy, we found preferences for access limitations are often very strong. Eighteen participants across 13 households classified at least one hypothetical policy violation as a 4 or 5 on our Likert scale. These devastating or near-devastating scenarios included unauthorized access (read, modify, or delete) to financial data, schoolwork, email, hobby or activity files, work files, text messages, photos, home videos, journals, and home musical recordings.

Many participants considered unauthorized access by strangers, acquaintances, bosses, and teachers to be highly undesirable. Perhaps more surprisingly, several were equally disturbed by situations involving closer relationships like parents, children, family, friends, and even significant others.

Examples of these critical violations (along with their Likert scores) include F4A's children seeing her finances (4); C1B's boss seeing her photos (4); R4A's boyfriend seeing her personal documents or work files (4) or modifying them (5); and nine-year-old F5D's friends seeing her email (4).

People's concerns are not just hypothetical

Our results reveal that not only do people have data they want to protect, but their current protection mechanisms are sometimes inadequate (or perceived to be inadequate). We asked participants to recall situations where they were concerned their sensitive data might be at risk, as well as situations where their data was accessed improperly. Twenty-two participants could recall specific instances of concern; only six reported they had never had such concerns. Nine participants reported actual policy breaches of varying severity.

Participant F4A, a divorced mother of two teenage boys, reported concern about her sons accessing her email when she leaves her account logged in on a family computer. "Maybe someone sort of emails you a sexy email, or something, and I wouldn't want the kids to see it." R4B was upset when he caught a roommate in his bedroom, using his computer without permission. F2B said her roommate sometimes grabs her phone and looks through pictures on it without asking, which is "kind of uncomfortable." R4C has also had private photos exposed on more than one occasion, including one incident where his girlfriend "stumbled upon an ice skating video of me and my ex. And it wasn't anything, but it was an awkward moment."

R2A, a law student, once lent her computer to her adolescent sister, who inserted random words into a class assignment. R2A turned in the altered paper without noticing and had to apologize to the professor. Participant F1A reported a less serious instance of data modification: his wife acci-

dentally deleting shows from their DVR before he watched them. "It's frustrating, because you're expecting to see it... But what can you do, it's already done, it's gone."

People use a variety of access-control mechanisms

Because people are concerned about limiting access to their sensitive files, they take precautionary measures to reduce the risk of exposure. We found that while some people use standard access-control tools, others have developed ad-hoc procedures. These include using both technical and social mechanisms whose actual efficacy may vary, but which participants find reassuring. In total, 30 of 33 participants, including at least one in every household, reported using precautionary measures.

Use accounts, passwords, and encryption. Seven participants use passwords, encryption, or separate accounts for access control. Four said they are careful to log out or lock the computer when they walk away. R4C said, "I guess I'm a security junkie with my phone. Encrypting my text messages, it's not really necessary. But it makes me feel comfortable." Like most participants who used passwords, R4A protects her laptop rather than individual files. She said she uses the password "just in case when we have guests over, that nobody thinks that, 'Well, it doesn't have a password, that means I can use it.' Just to better my chances of not having my identity or secret information taken."

Limit physical access to devices. In most participants' configurations, data boundaries are device boundaries; anyone using a device has access to all the data stored on it. As a result, many participants are cautious about lending their devices to others, even for tasks like checking email or browsing the web. Most participants allow only people they trust to access their devices. As 15-year-old F4C said, "Obviously I don't let anyone who walks through the door on to my computer, but if someone's on my computer I trust them." A few participants allow others to use their devices only if they are present to supervise, and another few don't allow it at all. In households R1 and R3, devices left in common areas are considered available for general use, while devices stored in individual bedrooms are considered private. Some participants in these and other households shut down or put away devices to discourage others from using them. One participant keeps her most important data on an external hard drive, which she physically hides from her roommates.

Hide sensitive files. Participants also attempt to hide files within the file system: A few name sensitive files obscurely for concealment and others bury them in layers of directories. According to R2A, "If you name something '8F2R349,' who's going to look at that?" C2B said, "[My husband] is a good hider of things.... If someone was trying to find something specific and he had it hidden, it would take them a while."

Delete sensitive data. Six participants have deleted sensitive files to prevent others from seeing them. F1A has deleted pictures of his two-year-old daughter from his cell phone for this reason: "If I didn't want everyone to see them, I just

had them for a little while and then I just deleted them.” A few participants have closed Facebook accounts because of privacy concerns.

PEOPLE NEED FINE-GRAINED ACCESS CONTROL

In practice, many current access-control systems designed for home users favor simple, coarse-grained access policies. In Windows XP, the default “My Documents” and “Shared” folders divide a user’s files into those accessible only to her and those accessible to everyone on her network. Although more fine-grained controls are available, they may not be sufficiently usable, as evidenced by participants’ attempts to hide files. Apple’s iTunes offers options for sharing the user’s entire library, sharing only selected playlists, and requiring a password for the shared files. This configuration does not allow users to share different subsets of music with different people. Facebook supplies rich, customizable access controls for photo albums, but there is no differentiation between reading and writing. Any user who can view a photo can also tag it and leave comments on it. HomeViews, which is designed to enable easy data sharing for home users, is limited to read-only access [10].

Our results indicate people’s policy preferences may be incompatible with coarse-grained control mechanisms in several ways: 1) Some participants’ policies include fine-grained divisions of people and files. 2) Dimensions of policy specification beyond person and file are also important in some circumstances. 3) Even when individual policies are relatively simple, they differ significantly across participants; there is no small set of default policies that could meet most people’s needs completely. In the following subsections, we discuss each of these complicating factors.

Fine-grained division of people and files

Early in the individual interviews, we asked participants to explain which people they would allow to access which files. We found that many participants specified complex groupings for both dimensions.

For some participants, policy specification required many categories of photos, some of which had fuzzy boundaries. C5B, for example, made several kinds of distinctions among her photos. In her first attempt to categorize her photos, she divided them into photos she was willing to publish and those she wasn’t. After further thought, she divided the restricted photos into four categories: truly private photos as well as separate groups to share with family, sorority sisters, and general friends. Even these distinctions did not prove entirely adequate—there were some pictures she might only want to share with the people pictured. She also said her boyfriend could see some of the truly private photos, but not others, particularly those involving ex-boyfriends. R4C had a similarly complex division of photos into overlapping categories; he also mentioned photos that should carry different restrictions even though they were taken at the same event. Currently, both of these participants manage photo sharing by over-restricting; if they don’t feel they can control access to a photo precisely enough, they decline to share it at all. The need for multiple policy divisions is not unique to pho-

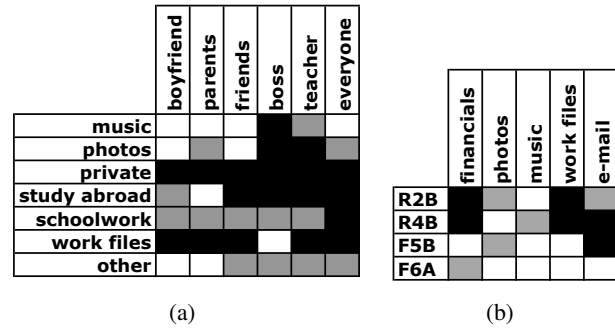


Figure 1. (a) A high-level view of participant R4A’s ideal policy. (b) Four participants’ ideal policies. White squares indicate willingness to share; black squares indicate restriction; and gray squares indicate a willingness to share some files under some circumstances.

tos; other participants specified similar distinctions within categories like music, videos, school files, and work files.

Our results also indicate people, like files, cannot be easily divided into just a few groups. Popular person designations included significant other, friends, family, co-workers, and strangers, but these groups often required additional subdivision. Several participants differentiated policy for one or two “best” friends; others made distinctions among close friends, casual friends, and acquaintances. Within families, policy varied for siblings, parents, and children. R2A said she is “far more willing to show my sister things than my parents.” Participants also make distinctions between bosses and colleagues as well as within groups of colleagues. C5A even differentiated among strangers: “I think I would feel less embarrassed if I knew someone 100 miles away was looking at it [a sensitive file] rather than someone on the bus.”

Figure 1(a) summarizes one participant’s ideal policy, indicating which files she would share (white), restrict (black), or sometimes share (gray) with which people. As this fairly-typical policy makes clear, access decisions are not binary across people or file types. The presence of gray squares indicates a finer level of detail would be required to completely specify this policy.

Dimensions beyond person and file

Factors beyond the person requesting access and the file being accessed also inform participants’ ideal policies. We asked participants to think about differences between read and write permissions, as well as whether or not the participant was present during the access, the participant’s location, the location of the accessor, the device used for access, and the time of day of the access. Each of these factors was meaningful to at least a few of the participants.

Distinguishing read access from write access. Many participants described important policy differences between read access and modify/delete access. F4C said no one else should ever be able to modify any of his files; C2A and F1B were willing to grant their bosses only read access to some files. A few participants described general categories of files they were not concerned about sharing, but that they would want

to protect from modification or deletion, including music, game files, schoolwork, and photos. This read-write distinction extends to highly trusted people such as family members and significant others. In one of several examples, middle-school-student F5C was willing to share almost all of her files with her family members, but did not want to grant modify or delete permissions. Similarly, R4A was willing to share highly sensitive files such as financial information and photos with her boyfriend, but did not want to grant him write access to any files.

On the other hand, a read-only system would not be sufficient for some participants, who see value in allowing others to edit their files sometimes. C5A wanted to let his mother improve his resume, and F2B would allow friends to provide feedback on scholarship essays. F5B would let clients update business files they send her. R2B expressed interest in allowing collaborators to edit files related to joint projects.

Presence. Policy specification based in part on whether or not the file owner is present resonated with a majority of participants. Participants believed that being present would allow them to exercise additional control over who accessed what, as well as providing social pressure to encourage good behavior. F3B, a nine-year-old boy, said, “If I was next to [my friend], I would know which files he would be bringing up, but away from him I wouldn’t have a clue what he was doing on my computer.” R4C said, “If you have your mother in the room, you are not going to do anything bad. But if your mom is outside the room you can sneak.” According to C3A, “If I’m in the house then it’s likely that I’m spending time with them. If I’m not with them, I can find them and say, ‘Hey! What are you doing on my computer?’”

For a few participants, being present provides additional benefits. Three said being present would allow them to make a last-minute decision to share something. R1B said she “might not remember what some of those files are,” but if she is present, she can look at them and make an informed decision. Others said they wanted to be there to witness the accessor’s reaction, explain things, and correct any misunderstanding. F6A wanted to make sure some opinionated journal entries wouldn’t be misunderstood: “I could explain myself! Totally! If only I had a crystal ball for all the times somebody got upset with me and I didn’t know it. If only I could have been there, then I could have told them: No, I am a lover, not a hater!” He also mentioned being present to explain things to his children: “Most movies I want to be there with [my son] ... in case he has questions or it’s too scary. I can calm him down.”

Location. We asked participants how location—their own or the file accessor’s—would affect their ideal policies. A slight majority said they felt safer sharing data in their home than in other environments. C5A said, “I don’t want them to look at my emails or texts. But if they were here, I wouldn’t care if they wanted to look at my email. I don’t know why, but I just feel more comfortable doing these things at home than being out in public with my information.” Eight participants did not want to share any files in public places like buses or

coffee shops. According to F1A, “Chang[ing] the settings as I move? That makes sense.... Going to work with the laptop vs. being at home—you might put it on extreme lockdown.” In general, participants’ responses to this question reflected their ideas about who was likely to be at a given location. R1C said, “At studio [at school] I am more hesitant to share my files if I am not there. In the apartment I can trust them with music or movie files. There is a mutual trust with people you live together with.”

To many people, the accessor’s location could be a proxy for trust: guests in the participants’ homes were presumed to be trusted. Several participants said they would share more with people in their house; a few others would share more with people who were in their bedrooms, an even higher marker of likely trust. According to C2A, “I feel that if they are in my house I can control them a little more. If they are in their house, they have a freedom to do whatever they want and there is not a chance of me walking in on them.”

On the other hand, some participants said their own location would not make a difference. According to F2B, “If there’s a way to have a certain setting for a specific individual and have that setting not change based on location, then I wouldn’t mind having the same access rights for my friends when I’m home or at school.” Policies based on the location of the accessor also didn’t make sense to many participants. As F6A said, “Just because you are inside my house, I would not categorize my files differently than if you were not there.”

Device. We also asked participants whether or not the device used for access would affect their ideal policy. Most said the device had no effect, but a sizable minority did find it meaningful. To several participants, including R4A, devices with smaller screens are preferable for accessing sensitive files, as “it feels more private on a smaller screen.” In contrast, others worried that a private device like a mobile phone might promote sneakier behavior than a public device like a television. According to F1B, “Maybe it’s worse doing it on the laptop [than the TV], because of being a bit more private about it.”

Time of day. We asked participants if their policies would vary according to the time of day when access was attempted. To a large majority, this idea did not make much sense; as C2B said, “It doesn’t matter the time of day.... The things that I don’t want you to see, I don’t want you to see at any time. And no time would be worse than another time.” A few, however, did find this possibility interesting. Some saw time of day as a proxy for presence or awareness; they did not want to share files while they were sleeping, because they could not know about or control the transaction. Said C3A, “If it’s bedtime and I’m in bed, then I don’t really get to see what people are looking at if I wanted to.” F3A wanted to restrict her young sons’ access to files at night, when they are supposed to be asleep.

Policies vary across people and households

As we have shown, some individuals’ ideal policies are complex. Even when our participants’ policies were relatively

simple, however, they rarely overlapped. Thus, no standard set of default rules can be expected to meet most people's needs. R2B wanted to tightly restrict financial and work files, was willing to share email and photos with most friends, and was not at all concerned about sharing music. R4B, by contrast, did not consider his photos private but was concerned about sharing email and music, even with friends. F5B was interested in restricting email, but not concerned about sharing financial or work files. F6A did not consider anything except some financial information private. Figure 1(b) illustrates some of these variations.

Participants also had different attitudes about privacy. Many participants, including R2A, C4B, C5B, and R4C, started from the presumption that everything should be private and then listed specific items to share with specific people. According to R2A, "Basically, it's my stuff; if I want you to have it I'll give it to you. If you want access to it, then ask." C5B said, "The files that I do have are private.... If I share it with you there's a specific reason." Other participants, including C3A, F4C, F5A, and F6A, started from the opposite position: sharing everything except a few specific exceptions. According to F5A, "I don't really have ... private files.... There's nothing that I am hiding from anybody." C3A said, "I'm not really that private. There's not a whole lot of stuff that I really want to keep from people aside from financial stuff."

In many cases, we found broad agreement on a general principle but enough variation in the details to make defining a satisfactory default policy difficult. For example, most participants identified one or two most trusted people—often a best friend or a spouse—to receive the most access. Within this group, about half were willing to grant this closest person complete access to everything; the other half wanted to grant access to most things but restrict access to some things. The specific exceptions varied and included everything from email, photos, and text messages to financial documents, work files, and even web-browser history.

AWARENESS AND CONTROL

In the previous section, we showed that participants responded positively to several policy dimensions beyond person and file, including location, presence, and device being used. Participants gravitated toward options they perceived as providing the most visibility into and control over accesses to their files. As C5B explained, "I guess I'm not a terribly private person, but I think if someone's going to be meddling in your things, you should be able to know what exactly they're looking at."

While it is not surprising that participants are looking for more control over their data, their ideas about what control means and how to achieve it show unexpected variety and depth. In the following subsections, we describe three specific manifestations of this desire for control: a preference for being asked, a need for iterative policy creation and refinement, and an interest in knowing not just who is accessing their data but why.

Permission and control

Participants often think of digital data sharing in terms of asking and granting permission. For some, setting policy a priori does not feel the same as granting permission. We found that mechanisms such as being physically present and responding to system-prompted access requests can provide a stronger sense of permission-based control and therefore increase people's comfort with data sharing.

Many participants wanted an access-control mechanism that reflected standard social conventions of asking before using someone else's things. Three participants said no one should access anything without their express permission. C2B said, "In general I want to be asked. I'd prefer to give [my files] to them. I would not want someone to just look at them." C3A agreed, saying "I'm very willing to be open with people, I think I'd just like the courtesy of someone asking me.... If you ask someone nicely for pretty much anything, people will be more than willing to help you out." When we asked R4B about who is allowed to use his devices, he answered, "Any friend of mine who asks." According to C4A, "Without my permission, without my directly sending it to you, I wouldn't like you to look at ... the financial files or my email. That's my personal stuff."

To many of these participants, the idea of specifying in advance which people can access which files does not seem to convey a sufficient sense of control, possibly because they don't understand the idea of policy specification or they don't trust that policy will be enforced correctly. Several participants, when asked to describe their ideal policies, responded that no one should be able to access their files without their permission. We asked R2A what her boss should be restricted from seeing, and she responded, "Ideally I wouldn't want him to see anything except what I give him access to." Along the same lines, as discussed previously, several participants expressed concern about allowing access while they were sleeping. As R3A said, "I can't be giving you permission while I sleep because I am sleeping." Responses like these suggest that, to many participants, setting an access policy does not seem equivalent to granting permission.

Five people said that when they are present they can control which files can be accessed. According to C2A, "We don't get company that much, and we are usually constantly with our company. If they were viewing something, I would be there at all times, guiding them through where they should go or not." C1B said being present would affect her policy, "because I could say, 'These are the things that you could see.'"

Participants responded positively to the idea of a reactive policy-creation system in part because they felt it would extend social conventions of permission into the digital world. C4A said a reactive system "sounds like the best possible scenario.... It would make me feel much more comfortable if people asked before they could modify or view the files at all. I like that a lot." Others said they would use a reactive system even for files they expected to rarely or never grant access to. C5A was open to making his finan-

cial documents—designated as highly restricted—available via such a system. “I don’t think I would mind, if it asked me permission first. Say if an employer needs to see it... I can’t imagine too many people want to look at my stuff.”

Iterative policy creation

For many participants, one important aspect of controlling access to their data was the ability to fine-tune policy easily and repeatedly. Some participants said they might want to make decisions about access at the last minute. R2B said she might change her policy “if there is something particularly relevant that I wanted to show, that I wouldn’t normally want everyone to have unfettered access to.” As discussed above, participants like C5A were interested in using a reactive policy-creation system to grant permission even to files they had not previously planned to share.

Three people placed particular emphasis on the ability to review policy and remove access. C2B said, “I would like to be able to go back on there and say, ‘You said yes to all these people to view these things,’ but if for some reason I no longer want them to do that, I could say ‘denied’ now and take them off the list.”

Participants were also interested in fine-tuning their policies based on observed activities. Nine participants were interested in checking a detailed access log for unexpected or undesirable access patterns and then changing policy accordingly. C1A said, “It’s nice to know who is accessing data more frequently. It opens the question: Are they the only ones viewing them, or are there other people standing next to them?” R1A added, “If someone has been looking at something a lot, I am going to be a little suspicious. In general, I would [then] restrict access to that specific file.”

We also found evidence that at least some ideal policies change over time. C2B wanted to temporarily limit her sister’s access when they fight. “She’s not talking to us right now.... She’s one of those people who, if you get mad at her, ... she’ll rip up all the pictures of you. ... She could erase stuff on my computer.”

Not just who, but why and for what purpose

Participants wanted to know not only who was accessing their files, but also why. C4B said, “Before you even touched anything, I would have to find out why you’re doing it.” F2A said she would like to use a reactive policy-creation system “if I know the purpose” for the request. F4B said a reactive policy-creation system “would be very useful, especially if maybe when they sent that they could add a message as to why they needed to see it.” This was especially true for write permission—C5B said that she might grant permission to modify a file, “but I think I’d probably have to get into contact with them and ask them why they wanted to.” C3B agreed: “I guess maybe if I got notified, like having the option of rejecting any of the changes, then that would be OK. Like if somebody changed it and it was better, then I could say OK, but if I didn’t like it then I could reject it.”

This interest extends to knowing how files will be used. F4B

said, “I feel more comfortable if they’re with me or I can see them, because then I have a better idea of what they’re doing with whatever files they’re seeing.” He also mentioned a similar concern related to the device used for access: “If it was something portable, if they’re using their phone, I might be worried about who else was watching.” F5A felt more comfortable sharing files in his home, where he assumed it would be impossible to show files to an unauthorized third party without his noticing. F3A wanted to limit the devices used for access out of concern about people making copies of her files: “Probably I wouldn’t want them to be able to save my information on their computers. ‘Cause from my devices they would be able to view it but not save it.”

MENTAL MODELS AND SYSTEM DESIGNS DON’T MATCH

Our interviews revealed several mismatches between people’s mental models related to access control and current system designs and operations. These often occur because users carry assumptions from the physical world into the digital world, where they may be invalid or inadequately supported by system designers. These assumptions affect the ad-hoc access-control mechanisms people create as well as the factors that make them feel secure.

Hiding files in the file system. As we discussed, some participants attempt to hide sensitive files. This idea comes from physical-world practices of hiding important items or mislabeling file folders to avoid suspicion. The couple in household C4, for example, keeps their most important papers in a small, hidden box; only less important papers are kept in the file cabinet, which is used as a decoy. The increasing availability of search tools, like Spotlight, Windows Search, and Google Desktop, that allow fast, accurate discovery of desired content regardless of file name or directory structure may invalidate this approach.

Preventing violations with presence. Based on physical-world experiences, many users believe being physically present can prevent policy violations. R4A, for example, said, “When I let people use my laptop, I’m usually near them, because it makes me feel comfortable that if anything were to happen, ... I’m right there to say, ‘OK, what just happened?’ So I’m not as worried.” Participants note that their presence may increase social pressure against behaving badly. They also believe they will be able to notice policy violations and react quickly enough to prevent problems. Computer policy violations, however, are often faster or less obvious than physical-world break-ins, which may complicate detection even if the file owner is in the same room as the offender.

Device boundaries. Many participants base their access-control measures on the idea that device boundaries and data boundaries are the same—anyone using a device can access all the files on it and no files can be accessed without physically touching the device where they are stored. As the increasing ubiquity of networking continues to blur distinctions between devices, this heuristic becomes less and less accurate. We also observed that users who subscribe to this model do not take advantage of tools like separate accounts or per-file encryption to segregate files within a device.

Location as a proxy. Some participants used the file accessor's location as a proxy for trust. For instance, based on the premise that only trusted people come into their homes, users would allow anyone within the home a high level of access to their data. It's not clear, however, that location is a particularly accurate proxy. C5B first said she would trust people in her house to access most files, but quickly changed her mind. "I guess originally my assumption would be ... if they were in the house, I'd know them, and they'd be close enough of a personal friend for them to actually be invited into my home. But then I was thinking, we've had plumbers here, guys laying carpet, stuff like that.... People are strange and might be snooping." In future work, it might be interesting to investigate whether the imprecision of this mechanism outweighs its convenience in real-world scenarios.

Infallible logs. Several participants wanted to use a detailed access log or notifications to verify enforcement of policy as well as to confront violators about their actions. F1A said, "[If] I all of a sudden got a thing [alert] on my phone, beep beep, somebody logged in to your account and is looking at it, yeah, I think that'd be great." According to C2A, a log would mean "I can call them on it [a violation] afterwards, and I would have proof of it." These statements rest on the assumption that even if the access-control system is sufficiently broken as to allow policy violations, the log or notification system would remain correct. This assumption seems dangerous, because an attacker sophisticated enough to bypass a reasonably robust access-control system may also be savvy enough to prevent her activities from being logged.

GUIDELINES FOR SYSTEM DESIGNERS

Based on our results, we have generated several guidelines for developers of access-control systems aimed at home users.

Allow fine-grained control. We found that participants' ideal policies were often complex and varied and were not always defined strictly in terms of files and people. It is important to keep in mind, however, that not all policies are fine-grained and not everyone wants to specify a detailed policy. An access-control interface should be designed to allow easy policy specification at multiple levels of granularity, according to the user's preference.

Plan for lending devices. We found that participants, especially those living with roommates, are often asked to lend computers to others who want to check e-mail or browse the web. Participants are often uncomfortable with these requests, because they worry that the borrower will access private files or overwrite important data, either accidentally or on purpose. Karlson et al. suggest lightweight, limited-access guest profiles for mobile phones, with an emphasis on switching to this mode discreetly to avoid the appearance of distrust [11]. We suggest applying a similar approach to laptops and other devices.

Include reactive policy creation. Response to a hypothetical reactive policy-creation system was overwhelmingly positive, with 27 participants expressing interest in using such a system in at least some circumstances. R3A said, "I'd like

that, it's useful. Only you can decide. That's something I would use." F4C answered, "That would be good.... Because then it would be easy access for them while still allowing me to control what they see."

Include logs. The majority of participants in our study also reacted positively to the idea of a detailed access log that would record all access attempts and their results. Some participants were interested in a log only out of curiosity, while others said that log contents might influence them to modify their policies. Six participants said they might share more if a log were available, including C4A, who said she would be "not a lot more open, but better than what I usually share." We recommend including a log or even a semi-real-time notification system designed to be human readable and to support policy changes based on log contents.

Reduce or eliminate up-front complexity. We found that although some participants' ideal policies are complex, defining fine-grained policies up front is difficult. Several participants, including C2A, reported that setting up a detailed access policy would be too much work. "If I had to sit down and sort everything into what people can view and cannot view, I think that would annoy me. I wouldn't do that." Even defining broad categories of access is seen as troublesome; participant R4C acknowledged he would not "go through the trouble of setting up a guest account" even to protect important files. As discussed earlier, some participants had difficulty specifying an ideal policy ahead of time and expressed interest in making last-minute policy decisions. We recommend reactive policy creation, either alone or in combination with preset policy, as one possible mechanism to reduce or even eliminate the up-front cost of setting fine-grained policies.

Acknowledge social conventions. A new design for an access-control system should take into account users' interest in the social convention of asking for permission. This is another instance where reactive policy creation could be helpful.

Another social convention for which we found strong interest was the idea of plausible deniability. Participants do not want to appear secretive or sneaky; as R4A said, "I don't want people to feel that I am hiding things from them." Several participants felt nervous about admitting they had private data and often felt compelled to justify it. C4A said, "Not that I have anything wrong or anything that can even be considered wrong, but I still want ... my privacy." Designers should take this into account and build into any new system a means of unobtrusively restricting access.

Support iterative policy specification. We found that ideal policies change over time and users need to be able to easily review and refine their policies. We recommend creating an interface to allow users to see their current policy, review the resulting access logs, and make changes as needed.

Account for users' mental models. We discovered many instances where users' mental models of computer security in general and access control in particular are not well aligned

with computer systems. New access-control systems should attempt either to fit into users' pre-existing mental models or to guide users to develop mental models consistent with the systems' behavior.

ACKNOWLEDGMENTS

We thank the members and companies of the PDL Consortium (including APC, DataDomain, EMC, Facebook, Google, Hewlett-Packard Labs, Hitachi, IBM, Intel, LSI, Microsoft Research, NetApp, Oracle, Seagate, Sun Microsystems, Symantec, and VMware) for their interest, insights, feedback, and support. This material is based on research sponsored in part by the National Science Foundation, via grants #CNS-0831407, #CNS-0756998, and #DGE-0903659; by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office; and by gifts from Intel.

REFERENCES

1. ABI Research. Home networking end-user snapshot: Consumer adoption of home and media networking, June 2008.
<http://www.abiresearch.com/research/1000323-Home+Networking+End-User+Snapshot>.
2. D. B. Baker. Fortresses built upon sand. In *Proc. NSPW*, pages 148–153, 1996.
3. L. Bauer, L. F. Cranor, R. W. Reeder, M. K. Reiter, and K. Vaniea. A user study of policy creation in a flexible access-control system. In *Proc. CHI*, 2008.
4. L. Bauer, S. Garriss, and M. K. Reiter. Distributed proving in access-control systems. In *Proc. IEEE Symposium on Security & Privacy*, 2005.
5. K. Beznosov and O. Beznosova. On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, 15:420–431, 2007.
6. A. Brush and K. Inkpen. Yours, mine and ours? Sharing and use of technology in domestic environments. In *Proc. Ubicomp*, 2007.
7. D. Dearman and J. S. Pierce. “It’s on my other computer!”: Computing with multiple devices. In *Proc. CHI*, 2008.
8. P. Dourish, E. Grinter, J. Delgado de la Flor, and M. Joseph. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal Ubiquitous Computing*, 8:391–401, 2004.
9. S. Egelman, A. Brush, and K. Inkpen. Family accounts: A new paradigm for user accounts within the home environment. In *Proc. CSCW*, 2008.
10. R. Geambasu, M. Balazinska, S. D. Gribble, and H. M. Levy. HomeViews: Peer-to-peer middleware for personal data sharing applications. In *Proc. SIGMOD*, 2007.
11. A. K. Karlson, A. B. Brush, and S. Schechter. Can I borrow your phone? Understanding concerns when sharing mobile phones. In *Proc. CHI*, pages 1647–1650, 2009.
12. L. Little, E. Sillence, and P. Briggs. Ubiquitous systems and the family: thoughts about the networked home. In *Proc. SOUPS*, 2009.
13. R. A. Maxion and R. W. Reeder. Improving user-interface dependability through mitigation of human error. *Int. J. Hum.-Comput. Stud.*, 63:25–50, 2005.
14. J. S. Olson, J. Grudin, and E. Horvitz. A study of preferences for sharing and privacy. In *Proc. CHI*, 2005.
15. E. S. Poole, M. Chetty, T. Morgan, R. E. Grinter, and W. K. Edwards. Computer help at home: methods and motivations for informal technical support. In *Proc. CHI*, 2009.
16. V. Ramasubramanian, T. Rodeheffer, D. B. Terry, M. Walraed-Sullivan, T. Wobber, C. Marshall, and A. Vahdat. Cimbiosys: A platform for content-based partial replication. Technical Report MSR-TR-2008-116, Microsoft Research, August 2008.
17. M. N. Razavi and L. Iverson. A grounded theory of information sharing behavior in a personal learning space. In *Proc. CSCW*, pages 459–468, 2006.
18. L. Richards. *Handling Qualitative Data: A Practical Guide*. Sage Publications, 2007.
19. L. Richards and J. M. Morse. *Readme First for a User’s Guide to Qualitative Methods*. Sage Publications, Thousand Oaks, 2007.
20. B. Salmon, F. Hady, and J. Melican. Learning to share: A study of sharing among home storage devices. Technical Report CMU-PDL-07-107, Carnegie Mellon University Parallel Data Lab, October 2007.
21. B. Salmon, S. W. Schlosser, L. F. Cranor, and G. R. Ganger. Perspective: Semantic data management for the home. In *Proc. FAST*, 2009.
22. S. Volda, W. K. Edwards, M. W. Newman, R. E. Grinter, and N. Ducheneaut. Share and share alike: exploring the user interface affordances of file sharing. In *Proc. CHI*, 2006.