

Learning to share: a study of sharing among home storage devices

Brandon Salmon, Frank Hady (Intel Corp.), Jay Melican (Intel Corp.)

CMU-PDL-07-107

October 2007

Parallel Data Laboratory
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Abstract

As an increasing number of home and personal electronic devices create, use, and display digitized forms of data, it is becoming more important to be able to easily share data among these devices. This paper discusses an observation of a home data sharing system constructed from currently available technologies. We deployed this system into two households for two and half weeks and studied the impact of the system on household data usage and device management. We focus on users' perceptions of the system's advantages, the reasons home users employed multiple devices, the ways home users managed their devices, and how some distributed file system concepts were ill-suited to the home environment.

Acknowledgements: We thank the members and companies of the PDL Consortium (including APC, Cisco, EMC, Google, Hewlett-Packard, Hitachi, IBM, Intel, LSI, Network Appliance, Oracle, Panasas, Seagate, and Symantec) for their interest, insights, feedback, and support. This material is based on research sponsored in part by the National Science Foundation, via grant #CNS-0326453. Brandon Salmon is supported in part by an NSF Fellowship, and in part by an Intel Fellowship. This work was done while Brandon Salmon was an intern at Intel Corporation.

Keywords: home storage, ethnography, user study, distributed file systems, management, security

1 Introduction

Digital content is now common in the home. An increasing number of home and personal electronic devices create, use, and display digitized forms of music, images and videos, as well as word processing documents. People are increasingly shifting important content from printed media, such as photographs and paper, to digital storage. The transition to digital homes is exciting, but does bring challenges.

One challenge is the need for users to access the same set of data from a variety of their household devices. For example, as users shift their photos into digital form, the ability to show these photos on a TV in addition to viewing and managing them on a laptop is very appealing. Similarly, as the number of computers found in the average home increases, the ability to share documents seamlessly between computers also becomes important.

This kind of sharing between computers has been standard practice in enterprise and business settings for many years through the use of distributed file systems. Thus, systems that provide such sharing do exist. File system researchers have also begun to look at extensions to these systems to match the unique aspects of the home [13, 14].

The literature is still lacking in studies exploring how current distributed file systems actually fare in a home setting, and what aspects of such systems do or do not match this new environment. To this end, we deployed a distributed file system into two households and used a combination of tracing and ethnographic techniques to explore how this new system changed the lifestyle of participants.

After outlining the design of our study, we will discuss various aspects of how these households used and managed their technology and how the system affected these patterns.

First, we will discuss the advantages household members found in using the sharing system. Participants noticed that easier access to their data increased the amount and frequency of the data that they accessed. They also found that the sharing system prompted them to spend more time in public, social areas.

Second, we will discuss why household members chose to use one device in the home over another for certain tasks. While the characteristics of the device itself were important, a variety of social and locational factors were also crucial in making these decisions.

Third, we will discuss the way each household managed their devices, both with and without the system. We found that participants did not use usernames and passwords, but instead used social and locational cues to provide similar functionality. We also found that access control was a key problem within the home, but that the home environment led to fundamentally different usage patterns than that of the enterprise. We also observed that despite the relatively small scale of the home, the number of administrators and mapping of administrators to devices was complex.

We conclude with a discussion of the implications of these findings for home system designers.

2 Previous work

While we are not aware of previous research that examines users' behavioral responses and adjustments to shared home data systems in the same manner as our study, a number of studies have employed similar techniques or investigated closely related issues. For example, our work uses ethnographic methods in the home, an area with considerable precedent [9, 7].

Grinter et al. [5] performed ethnographic studies of households with advanced networking setups in the home, exploring the management of the physical networks (e.g., wireless routers) installed by the owners themselves. They noted that data sharing between devices was a problem that households found vexing and did not know how to address. Our study explores the affect on household behavior of a researcher-installed distributed file system aimed at tackling this problem. Our findings concur with many of the observations in

the Grinter study.

O’Brien et al. described the impacts of a new desktop box on a set of households [11]. They found that users were not necessarily happy with having expanded functionality in a single device, and explored many of the social issues involved. Our study kept the devices separate, but allowed for data sharing between these devices. Still, many of our findings reinforce the social factors mentioned in the O’Brien study.

There have also been a variety of related studies of portable technologies [4, 13, 14]. For example, Aipperspach et al. studied the locations where notebook computers are used in the home, using a similar combination of tracing and trace-prompted interview techniques, noting stable patterns of usage [1]. O’Hara et al. studied the social implications of portable video, following up a rich set of prior work on TV usage [8, 3]. This study reinforced the fact that even within the home, users may employ multiple devices to access the same content [6].

A collection of studies have analyzed the filesystem behavior of enterprise settings. Oulasvirta et al. studied how users switch between different devices at work, and noted that users found many reasons to use a range of devices at work, and observed a variety of different methods for handling device swapping [12]. While not focused specifically on either home or business, several studies have also explored the rich social implications of music sharing in distributed environments [16].

A variety of researchers have also studied home automation technologies and their impact on the family [15, 2]. For example, Woodruff et al. described how home automation systems in Orthodox Jewish households play an important role in family worship [17].

3 Study description

To explore the effects of shared data on a household, we recruited two households from the Portland, Oregon area, installed a home data-sharing system in their homes, and observed them over two and a half weeks. Our analysis methods consisted of both system-level tracing and ethnographically-inspired interviews. While the short length of this study limits our ability to generalize all findings, we were able to observe many of the high-level challenges involved in adopting such a system.

We installed the system in each household and provided technical support as needed throughout the study. At the end of the study, we removed the technology and returned things as they were before. Each household was compensated monetarily for their participation.

3.1 System description

To test the effect of shared data in the home, we installed a home server system in our participants’ homes. We used a single Windows machine with Windows file sharing as the data server. We introduced two Linux-based digital video recorders (DVRs) to the house, which we connected to existing TVs. We also added a laptop-based stereo system running iTunes. In addition, we pointed the “My Documents” folder of all of the household computers into the system.

Each device in the household stored all of its data on the single server machine over Windows file sharing. This allowed all devices to access any data stored in the system. The following list explains the directory structure on the server and device targeting.

- **Music:** All music files were stored in this folder. The stereo, DVRs and iTunes applications on household computers pointed to this directory.
- **Video** Video content including:
 - **Movies:** All movies were stored in this folder. The DVRs had access to this directory for movies.
 - **TV:** The DVRs placed recordings of TV shows in this folder, which other devices could also access as mpeg streams.




Figure 1: **System usage form.** This shows an excerpt from the form we used to review system usage with each household. Each circle represents one day, across time, and different colors represent different types of access. This shows four days of access on Phoebe’s laptop.

- **Users:** Each user had their own sub-folder in this folder where the files from their “My Documents” were placed.

We provided users with 250GB of storage and backed up the full system in case of system failure. To avoid problems with network bandwidth, we put the server and DVRs on a wired gigabit network and connected this wired network to the household’s wireless router.

To allow household members to control access to their files, we let them specify a private folder only available to a user-specified set of devices and a public folder available to all machines.

3.2 Method

Questionnaires: To begin the study, we conducted a group interview with the collective memberships of each household. We also held one-on-one interviews during which we asked each household member to help us fill out a form detailing the devices they used, the data they stored on each, and their reliability, mobility, and privacy preferences for those data.

Technology tour: At the outset of the study, we had the household members walk us through their homes, showing us each digital storage or playback device, and describing what data was stored on it and how it was used. This allowed us to get a sense of the technology they used regularly, and the perceived differences between devices.

Tracing: We added a tracing module to the server to allow us to see which data users accessed from which devices at which times, and begin to discern weekly patterns in access behaviors.

Weekly interviews: To get a better sense of users’ interactions with the system, we held weekly interviews with each household. For these interviews, we brought visual representations of their weekly traces, to start the conversation about usage behavior, as shown in Figure 1. Both the diagramming technique and its application as interview prompt were borrowed from Aipperspach’s study of laptop usage [1].

3.3 Household selection

The goal of this study was to explore the effect of data sharing across devices on given households. In order to minimize the learning curve involved in adopting such a system, we recruited households that listened to music and that already had cable and Internet access. However, we avoided self-described computer experts, to avoid biasing the study towards this population.



Figure 2: **The Bradys.** Left to right: Mike, Carol, Oliver, Bobby and Peter.

We screened for households with at least three adult members, so that the behaviors and patterns would be more involved than what might be found in a two-person household. Since our system did not support taking devices and data out of the household, we targeted households that did not frequently take devices outside of their home.

We recruited two households representing different configurations by posting flyers and through snowball sampling following personal connections. One household consisted of a blended family, and the second consisted of a set of adult roommates.

Household 1 - The Bradys: Figure 2 shows the Brady¹ family. Mike and Carol Brady married 10 years ago, and both had children from previous marriages. The Bradys lived in a suburb of Portland. Two of Carol's sons still lived at home. Peter was in high school, and was interested in computers. Bobby was in junior high and was perhaps the least interested in technology in the family, although he played video games with his brothers. Carol's nephew Oliver also lived with the Bradys. Oliver was attending college and had recently started studying information management. Oliver was living with the Bradys for a year, and often worked somewhat separately from the family. Carol's oldest son, Greg did not live with them, but was a frequent visitor, along with his fiancé, Marcia.

The Brady's had two TVs that they connected to the system in the house: the upstairs TV in the living room and the downstairs TV in the den. A stereo system with large speakers was located in the family's living room. At their request, we hooked our new laptop stereo system through their existing system. Carol had a new laptop that she used to run a business out of her home and that the rest of the family used on occasion for personal activities. Oliver had his own laptop that he used for school and that his cousins used to play games. Peter was the primary user of an old desktop computer, which was located in the downstairs den. Bobby had an even older desktop in his room, which he only used for games, that was not connected to the Internet.

Household 2 - Piper, Paige and Phoebe: Figure 3 shows Piper, Paige and Phoebe. All three were young professional women in their mid to late 20s. Piper owned their house in another suburb of Portland, and had family nearby. She taught film and health at a nearby high school. Phoebe had lived with Piper for about a year and was an interior designer. Both Piper and Phoebe liked technology, and enjoyed tinkering with it. They did a lot of filming and photography, including a large number of home videos and images modified in Photoshop. Paige moved in a few months before the study. Paige had little interest in technology, and used it sparingly. The three roommates had many common friends and spent a lot of time together. They

¹All names have been changed to preserve anonymity.



Figure 3: Left to right, Paige, Piper and Phoebe.



Figure 4: **System setup.** This photo shows one DVR and the stereo setup in Piper's living room.

often threw parties and had friends over to their house.

Piper, Paige and Phoebe chose two TVs in the house to connect to the system: one in the main living room and one in the study (see Figure 4). At their request, we put the laptop-based stereo system in the living room. Piper owned a desktop computer that was placed in the study. Phoebe had a fairly modern laptop which she used throughout the house, including in her bedroom. Paige owned an older laptop that was not connected to the Internet and that she usually used in her bedroom. Piper used a portable hard drive to store and transport her films.

4 Advantages of data sharing

Current file sharing methods: Before the installation of the system, household members primarily used email and physical media to move data from device to device. They used CDs and DVDs (and, in one case, a portable hard drive) to transport large media, and USB drives and email for smaller files.

The common element of all these methods is the need to explicitly copy data from one device to another, often by moving a piece of media or a device to the location where the data is desired. We found that making data transparently available, thus removing the need to move data to the viewing device, improved the usage experience.

Despite the heavy use of distributed file systems in the enterprise, no one in either household used such file sharing systems at home or, for the most part, outside the home. Piper used two computers at work that were supposed to share data, but because of technical problems they did not. She explained that she “gets a lot of exercise” running data from one machine to the other.

Shared data increased usage: Intra-device data sharing led to a reported increase in data usage over the methods listed above. Study participants expressed great enthusiasm for having data from multiple users aggregated at shared devices, like the stereo and TV. For instance, when asked what she liked about the stereo, Phoebe replied:

***Phoebe:** Just coming home and turning it on and having it there. And we could all share our music and, you know, download it from somebody's computer and put it over to there. So yeah, that was fun just to have a central stereo in the house that we could all listen to.*

Phoebe also explained that the experience of viewing movies (both home movies and downloads from the web) and photos on the TV was a very different experience than viewing them on a laptop.

However, despite this enthusiasm for sharing data between devices, the study reinforced the fact that current methods limit the amount of data people will use. Phoebe explained that having the data readily accessible made them think about it more.

***Phoebe:** Our homemade movies we used more, pictures we were all able to view those. It was just something, because it was there, we thought about it more. You know, like I said before it was more accessible. So it's in front of you more.*

Factors that contributed to users' perceived contrast in the accessibility of their files — and, therefore, of their expressed likelihood of using them given the aggregated storage systems — included the simple removal of the need to compile media selections and physically retrieve storage media, the danger of damaging (and mess of managing) physical media, and the immediacy of shared data resources. Piper explained that the hassle of moving media made a large difference in the amount they used the stereo. She explained that they did not use the stereo much without the sharing system.

***Piper:** It's SO much easier when it's all just files on there. I mean the difference is ridiculous. 'Cause if you want to just switch out you can. And now you actually have to open up the thing, and take out the thing, put a new one in.*

For user-generated content, another inhibitor to sharing was the need to create the media in the first place. For example, despite her excitement at sharing her videos, Piper explained that even after a year of work, she had yet finish the DVD needed to share them with others.

Carol's treatment of her photos reinforces this point. She was very proud of her photos: In fact during our tour of her home, when we simply mentioned her digital camera, Carol stopped and spent five minutes showing photos to the interviewer on a small camera screen. While she had bought CDs to transfer the photos over, she was still showing people the photos on the tiny camera screen because of the hassle of transferring them to a computer.

Carrying media around also led to conflicts between household members. One of the big reasons Mike didn't want a TV in the living room was the way the kids handled the removable media.

***Mike:** [Downstairs the kids] will have 20 stacks of videos out, we don't understand that, why don't you put the video away after you finish? ... they still do the same thing with videos in [the living room]. No matter, nine years of telling them to put it away, they still pop 'em out and leave 'em out on ... the floor.*

In contrast, Mike and Carol were happy to have the stereo in the living room, because they did not feel that it brought clutter. But even with the stereo, moving CDs around caused some amount of discomfort. Mike explained that he knew the kids occasionally used the stereo because:

***Mike:** We'll flip the CD in [the stereo] and go “uh oh, this must be someone elses' ”.*

Household members also talked about the importance of having the media easily available right when needed.

Mike: So that's nice, you get an immediate resource to something you are talking about. You know, if you go "oh, let's go downstairs to look at the TV," and sometimes, people when they are tired, they don't wanna go. But if you have it right there, you can just pop it up. ... because it was more convenient. Like I say, if I hadn't had my laptop there I probably wouldn't have.

The timing was a critical aspect of media use because many of the items were very context-dependent. Piper, Paige, and Phoebe's discussion of Internet downloads further illustrates this point:

Piper: You know, we'll pull up stuff. Things from our email we get from people. [laughs] Oh my gosh. Some really funny stuff. ...

Paige: That's normally at 2am though, when we're really tired. [laughs]

Phoebe: Yeah, you think "that's so funny" but the next day you look at it and,

Piper: that's not so funny.

This extra convenience was important enough to both households that they expressed dismay at having to give up the system at the end of the study. Mike even asked us to tell him how the stereo system worked, because he wanted to try to build his own.

Shared data increased socialization: Interestingly, having data shared between devices increased reported social behavior. Because household members felt that they could access all of their data from the homes' communal spaces, they were more likely to spend time in these spaces. In addition, because they considered their private technology and the shared technology connected, they spent more time near the shared technology. When asked whether and how the system had changed her lifestyle, Phoebe replied:

Phoebe: ... I used my computer out [in the living room] more, probably, than in my room. You know, not at my desk so much. Out in the public area, just because we had so much out in the public area to use. So... I could sit there and download something onto the screen, and then everyone could pull it up, and that kind of thing. Yeah, it centralized things more, so I think we all spent a little more time there [in the living room], than in the office.

Similarly, we noticed that the printer associated with Carol's laptop moved from the Bradys' downstairs office up to the family room part way through the study, and the laptop spent more time upstairs.

5 Why multiple devices

During the study, household members were continually choosing a device to use for a task from a large set of potential devices. In this section, we discuss the various factors affecting this decision. While the characteristics of the device itself were important, a variety of social and locational factors were also crucial in making these decisions.

Hardware: One basic reason for data sharing was the difference in capabilities and interfaces for different devices. For example, the TV has a much bigger screen than the laptop, and thus lends itself to showing photos to large groups of people in a way the laptop does not.

However, even within the same device type capabilities varied widely. For example, the boys in the Brady household played games on Oliver's laptop because it was much faster than their desktop machines, which could not run their games. This was exacerbated by the fact that devices lived for a long time in both households, both for financial and convenience reasons [5]. For example, the Brady household contained three generations of family desktops, with the older versions going to each son.

Software: Unlike the enterprise setting, software also played a large role in differentiating computers. For example, Phoebe used to send some work documents to her home laptop because it had a more up-to-date version of Autocad than that on her work machine. While she was able to eventually convince her boss to upgrade the work version, this was a long and uncertain process.

Piper did her video editing work on a computer at her parents' home because this machine had Adobe Premier Pro installed, unlike her machine. Piper still used this computer partially because she had data and

programs on the computer that she didn't want to take the trouble to move to another device. Instead, she had been working on convincing her parents to give the computer to her.

Piper: *And so, they might give it to me, because they got a new computer [crosses fingers and laughs] But then they were saying that we had to clean it and everything, and I was like, "Nooo, just give it to me," 'cause I have SO much stuff on there.*

Location: Another differentiator for devices was their location. For example, nearby devices were easier to use. Paige and Phoebe both had small TVs in their rooms, which they used when getting ready in the morning or while working in their room. However, in general they watched the much larger TV in the living room. Similarly, Piper would sometimes check her email on Phoebe's laptop because it was close by in the living room, and she didn't have to go back to her desktop machine in the study.

The atmosphere of the location also influenced these decisions. For example, the Bradys were careful to keep the living room TV different from the downstairs TV, to keep the living room neat.

Carol: *The TV brings a lot more relaxed attitude, at least in my opinion. Oh, we've gotta sit and drink, we gotta have the popcorn, we gotta, you know, it's like, that's a family room, you know, a living room ... feels like somewhere where our friends come, we sit and talk, it's something nice that when someone walks in our home, boom, it looks nice.*

In contrast, they were happy to have the stereo in the living room, because they did not feel it brought clutter. Similarly, the office was primarily a work place for Piper, Paige and Phoebe, so when someone was working in the office they would not disturb them with entertainment.

User assigned function: Members of our households also differentiated otherwise similar devices by assigning particular devices to particular tasks. For example, Carol's laptop was purchased specifically for her business. While it was used for other purposes as well, all members knew that this was the device's primary function. Carol also wanted to purchase another laptop, to separate the work and personal devices. This finding is consistent with previous studies that showed that users preferred to use technology marked as specialized over those marked as general, even if there was no difference in the technology itself [10].

Household members would often use their personal devices over devices owned by other users. So, for example, Peter would primarily use his desktop machine for homework. However, when possible they sometimes switched to devices with better capabilities. To continue our previous example, Peter would switch to Carol's laptop to print documents, because it was much faster. In addition, household members would sometimes avoid performing tasks on certain machines for fear of breaking the device: For example, gaming in the Brady household was not allowed on Carol's laptop.

Failure: Household members would also use alternative devices when technical difficulties arose. For example, part way through the study Piper had trouble with her computer, so she temporarily used Phoebe's laptop as a replacement device until she could get the desktop fixed.

In a different take, when Oliver accidentally broke one of the DVRs while working with the wiring, Mike used the second (working) DVR as a template, and was able to discover and repair the problem.

Discussion: Recent technologies have focused on consolidation of multiple functions into single devices. While this will probably continue and may improve the user experience for home users, this study suggests that there are many reasons that users choose multiple devices, which will be disrupted by a consolidation of functionality into a single device. This reinforces similar findings in the O'Brien study [11] and suggests that home users will deal with data sharing between multiple home devices for the foreseeable future.

6 Administration

Administration was a challenge during the study. It took a full day of a Computer Engineering PhD student's time to install the system in each household, even after setting up several working installations in the lab.

Even with a PhD student as a system administrator, both households ran into challenges in understanding and keeping the system running. While we will not discuss the system-specific particulars of device failure and repair, we will discuss several major themes in the way home users managed their devices.

6.1 User accounts

Current computer systems, both for the home and enterprise, rely on the abstraction of a *user*. In this abstraction, each household member logs into the computer with a username and password. In a distributed setting, users carry the same username across devices. This allows users to sub-divide devices, customize their experience, and control access to important information. We found that this abstraction was not used in our households. The remainder of this section discusses the alternative way ownership was handled in our households and the implications for system designers.

Home users don't use user accounts: Neither of our households used the user abstraction on any of their devices. The Brady desktop, which was primarily used by Peter, had multiple users defined, but all family members signed on as Mike. They were surprised to learn (during the course of the study) that there were other users defined on the machine. All other machines had a single user defined.

Primary owners: A set of devices in each household, especially the TVs, were considered communal devices used by everyone. However, almost all other devices had a primary owner, even if they were shared between multiple users. For example, Carol's laptop was shared by all of the Brady household members, even occasional members like Greg, but was understood by all to belong to Carol for her work. Similarly Piper's desktop was shared by both Phoebe and Paige, but clearly acknowledged by both users as belonging to Piper.

Location: Location helped define who owned a particular piece of technology. Most obviously, devices located in a given user's bedroom were generally considered "owned" by that user. In the Brady household, the upstairs/downstairs split also served to define ownership. The downstairs was primarily used by the kids, and thus the TV and computer downstairs were mostly kids' technology. In contrast, the upstairs TV and stereo were primarily owned by the parents. Carol's laptop also spent a considerable amount of time upstairs, as did the printer that was placed in the living room. This division was explicitly made by the Brady parents to keep control over the appearance and atmosphere of the living room.

Occasional household members: In addition to the residents of the household, each household had instances of *occasional* household members, who while they did not reside in the home were frequent visitors and users of technology.

For example, Piper was still a household member at her parents' house. During the study Piper had trouble with the licensing for a piece of software on her computer, since she had inherited the computer from her parents but no longer lived at home. Their discussion reveals the ambiguity in Piper's status as a "household member."

Piper: It was my family's, ok, I'm still part of my family.

Paige: And that household. [laughs]

Piper: Hey, I still have stuff over there.

Paige: You still have a ROOM, and a spot at the dinner table, I still think that counts.

Similarly, while Greg and Marcia did not live with the Bradys, they were frequently in the home and were large users of the Bradys' technology, including Carol's laptop. Some of these occasional members were also temporary visitors, such as Carol and Mike's children who returned and visited over the holidays.

These occasional members provided a challenge to the typical user abstraction used in enterprise storage systems, which assumes the user base to be contained and relatively static.

Discussion: Our findings suggest that the user abstraction may not be an appropriate abstraction for the home. But, this study does not necessarily show that the user abstraction is not useful in the home; as home users become more accustomed to shared data and more familiar with the model the use of this primitive

could increase. However, it does suggest that many home users currently do not find the advantages of customization and fine-grained access control important enough to bother with the primitive. Designers of home technology should be careful to validate that users are interested in such features before designing systems around them. Instead, system designers may want to consider techniques that allow home users to more easily provide the needed functionality without relying on this abstraction.

6.2 Controlling access to files

One important task in managing data systems is *access control*, or controlling who can access which files. This is important to provide privacy for sensitive data, and security for important data.

On the first day the system was installed, a lack of adequate access control caused problems. While each user could choose the privacy of his local files, in order for music to be shared on other devices it needed to be shared on all devices. Peter listened to a variety of music on his desktop machine downstairs. Once the system was installed, all of this music was accessible from the DVRs and the stereo upstairs. Mike quickly noticed that this music contained songs with titles that concerned him, especially songs by several rappers Peter enjoyed. Mike immediately began deleting these songs, initiating a power struggle between Mike and Peter. Mike was concerned about the type of music Peter was listening to, and felt an obligation to enforce this rule. Peter expressed frustration about the event, and felt that Mike was unfairly judging his music.

While this problem had been less prevalent before the system, the addition of data sharing disrupted the previous social balance. When asked if the system had made this problem worse, Peter said:

Peter: Yeah, kinda, he just went on the TV and just like went through all the music and just saw what he didn't want and deleted it. And again on the stereo, too, cause he didn't want to see them on the stereo, I guess.

However, while access control was important in the home, it was also a very different environment than that found in most enterprises. We discuss these differences in the next sections.

Home users don't use passwords: Of the eight participants in our study, only Phoebe used a password on her laptop. When asked why she was concerned about password protecting her machine, Phoebe explained she has always been a private person. However, she also revealed that this was a habit she developed while working in a computer lab setting at school and carried it over from this environment because she had "got used to it."

This further suggests that the use of passwords, while prevalent in the enterprise, is not a natural fit for the home environment. In the remainder of this section we will explore why this may be the case, and alternative methods home users employed to protect their data.

Threat model is different: Perhaps the biggest difference appears to be in the threat model assumed in the home. In a business setting adversaries are assumed to be malicious, requiring strict mandatory control. However, in our households, there already existed a high degree of trust between household members, causing them to take a different approach to access control. While some households may not have such trust, it is reasonable to assume that a large number of households will fit the pattern we have observed.

Consider Piper's reply when asked why she was not concerned about her data on the desktop being unprotected:

Piper: 'Cause I trust [my roommates], I don't think they'd do that, go through stuff that was mine that I wouldn't want them to look at. And even if they could I doubt they could find anything I don't want them to see on there. [laughs] You know, so I don't really have much to hide. So it doesn't really bother me.

This quote hits many of the key points we found in home access control: When asked about access control, many of the household members gave similar replies to Piper's. In general, household members were fine with their data being accessible to other household members because they did not consider their information overly sensitive, and they had a high level of trust with other household members. While many

household members had certain, specific items that they did not want shared, in general they defaulted to “all data shared” except special cases.

While almost all household members said they had nothing to hide up front, when actually storing data in the system most discovered that there was indeed data that they did not want shared.

For example, initially Carol said she was fine with all the data on her laptop being shared with all household members. However, as we worked with Carol to import her data into the system, we ran across a set of journal entries that she stored on her laptop. She was very uncomfortable with this data being shared on other machines, and made sure to make it private to the machine. In fact, she was uncomfortable even talking about it with her husband nearby. Similarly, when questioned more closely about her data, Piper, who had previously said she was not concerned, confessed she had private data as well, which she did not put on her desktop machine.

Device as security boundary: While users did not utilize passwords to protect their data, they were skilled at using other techniques to protect their sensitive information. For example, while Carol was very concerned about the privacy of her data, she was comfortable with it being stored on her laptop, which was shared with a variety of household members, including occasional members like Marcia, without any kind of password protection.

The device boundary was sufficient for Carol to feel comfortable about her data. She owned the device, and could control who had access to it. This was the way users in both of our households controlled access to their data. They were concerned with the devices that had access to the data, but preferred to control access to these devices using location and control of the physical device, rather than passwords and users.

In part, this relied on the assumption that other users would not, or could not find the data stored on their devices. Piper’s previous comment states that she does not think her roommates can find her documents, later comments about her private data further elaborated:

Piper: Yeah, there are a couple things [I didn’t want people to see] [laughs] There are some pictures I don’t want people to see.

Interviewer: Would you be comfortable putting them in [the standard location for her files]?

Piper: [still laughing] Probably not.

Interviewer: So where do you keep those?

Piper: [still laughing] I just keep them on my D: drive. It’s not in [the standard location for her files] even on my C: drive.

Interviewer: Hidden.

Piper: Very hidden!

This assumption is not always true, of course. For example, Mike explained how he once caught a file that Peter tried to hide from him:

Mike: All these kids think they’re pretty computer smart. I’m not computer dumb [laughs] and I’ve got enough orneriness into me that I almost know what these guys do, and so its like he one time hid something under another file and I found it. “How did you know that?” well I’m not stupid, I’d probably do that myself.

However, despite his ability to find individual files when looking for them, Mike did not search through Peter’s desktop. When asked why he deleted Peter’s music files which he had not deleted previously, he replied:

Mike: Well, I never had access to ’em, see I wouldn’t know what he was listening to. ... On a system like this at least I know what they’re listening to.

This suggests that though he knew that he could search through Peter’s data, he did not consider this an option. Later comments suggest that he was interested in looking for specific, banned files, but didn’t want to peruse Peter’s data. This also seems included in the initial quote by Piper in this section; while she didn’t necessarily think her data would be safe from a dedicated attacker, she assumed her roommates would not

intentionally search for her private data, much as they would not rummage through her room for personal items.

Location: Location also had a significant impact on privacy concerns. For example, during one interview, Phoebe was uncomfortable discussing one video usage, where she watched a video on a laptop in her room, because she was watching it with a visitor who had dated Piper.

On the opposite side, Carol explains why Peter’s desktop is in the main room downstairs, and why she doesn’t want to buy him a laptop.

Carol: The boys need a new desktop or laptop downstairs. I’m kinda having a problem with Peter and games and attitudes. I’m NOT wanting to buy a laptop. On desktops, he kinda has to sit in the same place [rather than playing games in another room]. ... We’ve had to do some stuff [to limit his game playing]. Not that he’s a bad kid, just the addictiveness of it, and the attitudes that come from it.

Similarly, Peter said he would want his music available on the downstairs DVR, because it was only used frequently by the kids, but not on the upstairs DVR or stereo, since this area was primarily used by the parents.

Policy complexity: While participants utilized a much more coarse-grained method of access control, this is not to suggest that the policies enforced by such methods were simple. Especially in the Brady household, where the dynamic of parent and teenager was important, these policies were nuanced and complex. For example, when asked if he would allow Peter to have a file space over which Mike did not have control, Mike was not entirely sure how to answer.

Mike: [long pause] It depends on what files. ... I’d have to go by their honor in some case or another. ... I’d want them to have some privacy to writing letters to girlfriends and things like that. ... I’m just more concerned about the content, like the music. If there was some way of finding things by content, like some things of a sexual nature, ... I’d probably do that. I should trust them more but, you always should, but if something came up that I was concerned about it, I’d want to be able to check it up and discuss it or something.

Mike was balancing Peter’s need for privacy against his responsibility as a parent to watch out for him. This policy was nuanced, and required great thought on his part. However, like the other users in the study, he was more adept at manipulating locations and physical devices than computer-centric abstractions like users and permissions, and so used these approaches to create policy.

Physical analogues: While this approach may seem weak in contrast to methods used for enterprise access control, it is analogous to the access control methods found for many other items in the household already. For example, Carol and Mike separate their movies into an approved collection, which is stored downstairs next to the TV, and a restricted collection that they want to control more closely, which is stored in Mike and Carol’s bedroom. The kids are required to ask permission before using one of these restricted movies. However, these movies are not locked away; Mike and Carol rely on a combination of trust that the kids will obey the rules, and the increased scrutiny this location provides, to maintain access control.

Discussion: Our findings suggest that home system designers may want to support a weaker set of security primitives than might be provided in an enterprise setting. Just as much of the physical access control in the home is provided without locks, much of home digital access control appears to be provided without password protection. The high trust level inside the home may also suggest a “default to open” policy for sharing, in contrast to the enterprise. Either way, designers must keep in mind that the home environment appears to have very different characteristics than the enterprise.

6.3 Dividing up administration

In the enterprise, *administrators* are the employees responsible for setting up and repairing technological devices. In contrast, average users do not involve themselves in these repairs. While not institutionalized, we found similar distinctions in the home, where some users were administrators and repaired technology,

while others did not. However, in the enterprise a single administrator generally manages a large number of machines. In contrast, we found that the social structure surrounding administration in the home was much more complex relative to the number of devices.

Multiple administrators: For example, we found that both of our households had several administrators, despite only having a handful of devices. For example, Piper and Phoebe would work together to figure out problems, while Paige was not involved, and usually only heard about problems after they were fixed.

Similarly, in the Brady household Mike, Oliver and Peter all took turns at working with the technology when challenges arose. Carol and Bobby did not do administration, and instead would call for whichever of the administrators happened to currently be at home. This meant not only that the same device might be repaired by multiple administrators, but that previous repairs might have occurred without the current administrator's knowledge.

Multiple areas of authority: In addition to having multiple administrators, each home also had different administrators for different devices. The public devices were usually managed in combination, but many devices had private owners and were managed separately. For example, Mike Brady explicitly gave each son his own computer so that they could play their games on these devices without breaking the family computer.

Mike: I have lost programs in the past because of a crashed computer so it's always a concern. Especially when they get these games going on. That's what scares me, these things that they do, you know to performance optimize the ability of the computer. ... We kinda keep as much separated as possible.

Phoebe explained that this was another reason she password protected her laptop during college; a fellow student used her laptop for the Internet because it was available, and accidentally contaminated the machine with a virus. Piper even made an administrative division between her C: and D: drives.

Piper: So, I keep most of the stuff on the D: drive if I know that I don't want anyone to access it, cause I have had my C: drive crash before and lost a lot of stuff that way. So usually when I download I use the C: drive, the D: drive is never touched by outside sources. So, I never have a problem with that. That's usually where I keep my digital photos.

Discussion: This suggests that the design of coordination approaches for home administrators could be a fruitful area of exploration, because of the large number of administrators, and loosely coordinated activity. In addition, because it is important for the administration of some devices to be insulated from that of other devices, home system designers must be cautious not to destroy these important administrative boundaries, even as they connect devices.

7 Discussion

We conclude with a discussion of suggestions for system designers gleaned from our study. Our study further validates claims that sharing between home devices will help home users more fully integrate home technologies into their homes, including bringing technology into more public and social settings.

However, this study suggests that there are many reasons that users choose multiple devices which will not be solved by simple consolidation of functionality into more sophisticated devices. Thus, system designers should be careful when attempting to consolidate function that they do not interrupt the social and locational factors involved in multiple device usage.

One surprising aspect of our study is the suggestion that the current user abstraction may not be appropriate for the home. While this study's scope and setup do not allow us to make a strong claim against such an abstraction in the home, it does suggest that currently home users do not find the advantages of customization and fine-grained access control important enough to bother with the primitive. Thus designers of home technology should be careful to validate that users are interested in such features before designing systems around them. In addition, system designers may want to consider techniques that allow home

technology users to more easily provide the needed functionality without relying on this abstraction.

Similarly, this study shows that access control is an important problem in the home, but that the needs for home access control are very different from those of traditional enterprise settings. For example, our experiences suggest that home system designers may want to support a weaker set of security primitives than might be provided in an enterprise setting. Just as much of the physical access control in the home is provided without locks, much of home digital access control appears to be provided without password protection. The high level of trust in the home may also lead to a “default to shared” policy within the home, in contrast to the enterprise.

In addition, while the number of devices found in the home is relatively small, we found complex administrative social patterns. Each household had multiple administrators, and multiple administrators would work on the same machine at different times. Thus the design of coordination approaches for home administrators could be a fruitful area of exploration. In addition, it was important for the administration of some devices to be insulated from that of other devices, suggesting that home system designers must be cautious not to destroy these important administrative boundaries, even as they connect devices.

In summary, our study found that home users are excited about technologies that allow them to share data between their devices, but that in several aspects home social factors make traditional enterprise approaches awkward and ill-suited to the home context.

References

- [1] R. Aipperspach, T. Rattenbury, A. Woodruff, and J. Canny. A Quantitative Method for Revealing and Comparing Places in the Home. *UBICOMP* (Orange County, CA, Sep. 2006), 2006.
- [2] Genevieve Bell and Joseph Kaye. Designing technology for domestic spaces: A Kitchen Manifesto. *Gastronomica*, 2(2):46–62, 2002.
- [3] Barry Brown and Louise Barkhuus. The Television Will Be Revolutionized: Effects of PVRs and Filesharing on Television Watching. *CHI* (Montreal, Canada, 2006), 2006.
- [4] W Keith Edwards and Rebecca E Grinter. At home with ubiquitous computing: seven challenges. *Ubiquitous Computing International Conference (UBICOMP)* (Atlanta, GA, 2001), 2001.
- [5] Rebecca E Grinter, W Keith Edwards, Mark W Newman, and Nicolas Ducheneaut. The work to make a home network work. *European Conference on Computer Supported Cooperative Work (ESCW)* (Paris, France, 18–22 Sep. 2005), 2005.
- [6] Kenton O Hara, April Slayden Mitchell, and Alex Vorbau. Consuming Video on Mobile Devices. *CRYPTO* (San Jose, CA, 2007), 2007.
- [7] Debby Hindus. The importance of homes in technology research. *Cooperative Buildings (CoBuild)* (Pittsburgh, PA, 1999), 1999.
- [8] Robert J Logan, Sheila Augaitis, Robert B Miller, and Keith Wehmeyer. Living Room Culture – An Anthropological Study of Television Usage Behaviors. *Human Factors and Ergonomics Society 39th Annual Meeting*, pages 326–330, 1995.
- [9] Michael Mateas, Tony Salvador, Jean Scholtz, and Doug Sorenzen. Engineering Ethnography in the Home. *CHI* (Vancouver, Canada, 1996), 1996.
- [10] C. Nass, B. Reeves, and G. Leshner. Technology and roles: a tale of two TVs. *Journal of Communication*, 46(2):121–128, 1996.

- [11] Jon O'Brien, Tom Rodden, Mark Rouncefield, and John Hughes. At home with the technology: an ethnographic study of a set-top-box trial. *CHI*, 1999.
- [12] Antti Oulasvirta and Lauri Sumari. Mobile kits and laptop trays: managing multiple devices in mobile information work. *CHI* (San Jose, CA, 2007), 2007.
- [13] Daniel Peek and Jason Flinn. EnsemBlue: Integrating distributed storage and consumer electronics. *OSDI* (Seattle, WA, 06–08 Nov. 2006), 2006.
- [14] Brandon Salmon, Steven W. Schlosser, Lily B. Mummert, and Gregory R. Ganger. *Putting home storage management into perspective*. Technical Report CMU-PDL-06-110. Sep. 2006.
- [15] Alex S. Taylor, Richard Harper, Laurel Swan, Shahram Izadi, Abigail Sellen, and Mark Perry. Homes that make us smart. *Personal and Ubiquitous Computing*. Springer London, 2006.
- [16] Amy Vaida, Rebecca E Grinter, Nicolas Ducheneaut, W Keith Edwards, and Mark W Newman. Listening in: practices surrounding iTunes music sharing. *CHI* (Portland, OR, 2005), 2005.
- [17] A. Woodruff, S. Augustin, and B. E. Foucault. Sabbath Day Home Automation: It's Like Mixing Technology and Religion. *CHI* (San Jose, CA, 2007), 2007.