ISR Institute for Software Research
University of California, Irvine

Security Day-to-Day: User Strategies for Managing Security as an Everyday, Practical Problem

Paul Dourish
Univ. of California, Irvine
jpd@ics.uci.edu

Jessica Delgado de la Flor
Univ. of California, Irvine
delgadoj@uci.edu

Rebecca E. Grinter
Palo Alto Research Center
beki@parc.com

Melissa Joseph
Univ. of California, Irvine
mjoseph@uci.edu

Brinda Dalal
Palo Alto Research Center
bdalal@parc.com

June 2003

ISR Technical Report # UCI-ISR-03-5

Institute for Software Research
ICS2 210
University of California, Irvine
Irvine, CA 92697-3425
www.isr.uci.edu

Security Day-to-Day: User Strategies for Managing Security as an Everyday, Practical Problem

Paul Dourish*, Rebecca E. Grinter+, Brinda Dalal+,
Jessica Delgado de la Flor*and Melissa Joseph*

*Institute for Software Research
School of Information and Computer Science
University of California, Irvine
Irvine, CA 92697-3425, USA
jpd@ics.uci.edu

+Palo Alto Research Center (PARC)
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
beki@parc.com

ISR Technical Report # UCI-ISR-03-5
June 2003

ABSTRACT

Effective security solutions depend not only on the mathematical and technical properties of those solutions, but also on people's ability to understand them and use them as part of their work. As a step towards solving this problem, we have been examining how people experience security as a facet of their daily life, and how they routinely answer the question, "is this system secure enough for what I want to do?" We present a number of findings concerning the scope of security, attitudes towards security, and the social and organizational contexts within which security concerns arise, and point towards emerging technical solutions.

Security Day-to-Day: User Strategies for Managing Security as an Everyday, Practical Problem

Paul Dourish[*], Rebecca E. Grinter[+], Brinda Dalal[+],
Jessica Delgado de la Flor[*] and Melissa Joseph[*]

[*]School of Information and Computer Science
University of California, Irvine
Irvine, CA 92697-3425, USA
jpd@ics.uci.edu

[+]Palo Alto Research Center (PARC)
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
beki@parc.com

ISR Technical Report #UCI-ISR-03-5, June 2003

Abstract

Effective security solutions depend not only on the mathematical and technical properties of those solutions, but also on people's ability to understand them and use them as part of their work. As a step towards solving this problem, we have been examining how people experience security as a facet of their daily life, and how they routinely answer the question, "is this system secure enough for what I want to do?" We present a number of findings concerning the scope of security, attitudes towards security, and the social and organizational contexts within which security concerns arise, and point towards emerging technical solutions.

1 Introduction

The security and integrity of computer and digital communication systems has always been a significant concern. The practical application of digital computing to real-world tasks in military, commercial and even academic settings has always required that storage, communication, and computation be protected by reliable security mechanisms. However, in recent years, the nature of the security problem has changed, due to the emergence and proliferation of online banking, electronic commerce, widespread electronic communication, and other activities associated with the rapid expansion of the Internet into daily life. Security is no longer merely a concern for system administrators, nor is it restricted to the relatively small, technically adept community of academics and scientists who, thirty years ago, had access to electronic communication systems. The pervasive spread of the Internet, and its decentralized nature, means that the security and integrity of computer and communication systems is now a practical, day-to-day problem for everyday, casual users of computer systems. Anyone who sends an email message, pays a bill online, or purchases a book from amazon.com must, on some level, be concerned with the security of the infrastructure that they use in carrying out those activities.

We believe that this transition implies a radical change in the way that security should be understood and studied. Specifically, we distinguish between two aspects of security as they arise in different research communities. First, we use the term "theoretical security" to describe the range of concerns that have traditionally been topics of investigation for the computer and communication security community. Theoretical security is the study of the mathematical basis of forms of protection and attack – the nature of cryptographic protocols, the problems of key exchange, the potential "leakage" of information via secondary channels, etc. Second, we use the term "effective" security" to refer to the degree of security that is practically achievable in real settings; in other words, how much and what kind of security people can actually use and understand. Theoretical security sets an upper bound for effective security, but, for a host of practical reasons, effective security typically lags theoretical security, often considerably. Consider two examples:

- *Turning it off.* The Placeless Documents project (Dourish et al., 2000) created an environment for associating active code with operations in a document repository, code that would travel with the documents. Clearly this required a strong security mechanism. An informal scheme had been initially developed, but was replaced by a more comprehensive, academically

rigorous system. However, in daily use, this system proved *too* rigorous – not only did it interfere with development activities, but the performance degradation that it imposed was judged too onerous. Most people, therefore, ran the system without the security features turned on – rendering it, effectively, less secure than it had been when the simple, ad hoc scheme was in place.

- *One-time Pad*. At one point, a research lab secured its firewall using a password mechanism based on one-time pads. Local software was provided to generate new one-time pads, protected by a password. The effectiveness of this solution hinged on the integrity of this password, and users were urged never to run the program and enter the password on anything other than their local computer, inside the firewall, to ensure that the password would never be broadcast over an insecure channel. However, this proved troublesome since many users had difficulty interpreting what might be meant by "their local computer." Users of older workstations reconfigured as X terminals, for example, naturally thought of the computer on their desk as "their local computer," despite the fact that it was running no client software; even users of regular workstations would have difficulty distinguishing those connections that might traverse an insecure network and those that might not, while, when people connected to their desktop machines from remote locations, the fact that they were making use of a network was sometimes hard to notice. Despite their PhDs in technical topics, many users found the terms of reference too difficult to interpret.

It is important and instructive to note that these are not cases involving naïve and non-technical end-users; rather, they both occurred in high-tech environments populated by well-qualified scientific personnel.

One response to the disparity between theoretical security and effective security is to dismiss it as a matter of the implementation of policy. It is ascribed to a lack of training, perhaps, or even to a willful stubbornness on the part of ill-educated users who, at best, don't understand their own best interests or, at worst, driven by petty spite, actively conspire to be revenged upon their employers for the apparent ignominies of daily working life by deliberately flaunting rules and procedures. However, we believe that to regard this disparity as purely a matter of procedure is to ignore some important issues – issues that, in fact, become more pressing as the reach of the Internet grows. We believe that it is important to treat the disparity between effective and theoretical security as a research problem.

Our working approach is this. Rather than focusing on the mathematical foundations of theoretical security, we want to examine the practical foundations of effective security. As they use computers and networks, people are continually confronted with the question, "is this system secure enough for what I want to do now?" Whether their task is to file a tax return, submit class grades, send email to a family member, or make a purchase, they must determine whether or not the configuration of technologies available to them meets the needs that they associate with this task. We regard every occasion on which someone enters a password or credit card number – or every occasion on which they decide not to – as an occasion on which this question has been answered (not always conclusively.) What is particularly notable about most solutions is that, in their attempts to provide security seamlessly and transparently, they deprive people of the resources they might need to answer this question. Our research goal is to examine just *how* people come to an answer.

In this paper, we report on a series of investigations into security as an everyday, practical problem, routinely encountered and resolved by computer users as they go about their activities. We have looked at how a range of users, in different organizational settings and with different degrees of technical skill, have addressed this question, and some of the strategies they use to resolve it. We begin by discussing related research. We then introduce the methodological approach that we have taken. Next, we present the results of our studies, and then conclude by considering their implications for the design of applications and infrastructures.

2 Background and Related Work

Security research has long acknowledged the role of human, social, and organizational factors in creating effective solutions.

In some cases, the complexity of making security work is as much a matter of interface design as anything else. Whitten and Tygar (1999) present a usability analysis of PGP 5.0, demonstrating the difficulties that users have in completing experimental tasks (in their user study, only 3 out of 12 test subjects successfully completed a standard set of tasks using PGP to encrypt and decrypt email.) The problems that they uncovered were largely problems of interface design, and in particular the poor matching between user needs and the structure of the encryption technology provided to meet these needs. Others such as Yee (2002) or Zurko and Simon (1996) have similarly explored the relationship between interaction design and security, and attempted to draw out a set of general principles concerning the interaction of the two.

In a series of studies, researchers at University College, London have explored some of the interactions between usability and security (Adams, Sasse and Lunt, 1997; Adams and Sasse, 1999). They have focused particularly on user-visible elements of security systems, such as passwords. Although many information systems professionals regard users as being uninterested in the security of their systems (and, indeed, likely to circumvent it by choosing poor passwords, etc), Adams and Sasse's investigations demonstrate that users are certainly motivated to support the security of the system, but often unable to determine the security implications of their actions. The specific problems that they identify with passwords have also led to interesting design alternatives (Brostoff and Sasse, 2000; Dhamija and Perrig, 2000).

One area at the intersection of usability and security that has received some attention is the role of access control in interactive and collaborative systems. For example, Dewan and Shen (Shen and Dewan, 1992; Dewan and Shen, 1998) have explored the use of access control and meta-access control models as a basis for describing and controlling degrees of information access and management in collaborative systems. This is not simply a technical matter, since the structure and behavior of these "internal" components can have a significant effect on the forms of interactivity and collaboration they can support (Greenberg and Marwood, 1994).

Many collaborative systems involve privacy issues and need to provide users with control over the disclosure of information to the other parties using the system. This has spurred a number of researchers to explore the development of privacy control systems that are tailored to the needs of end users. For instance, Dourish (1993) describes the relationship between three different security mechanisms for similar multimedia communication systems, each of which reflects assumptions and requirements of the different organizations in which they were developed. Bellotti and Sellen (1993) draw on experiences with multimedia and ubiquitous computing environments to identify the source of a number of potential privacy and security problems. Their primary concepts – disembodiment and dissociation – are both visibility problems, related to the disconnection between actors and actions that renders either actors invisible at the site of action, or actions invisible to the actor.

Based on their investigations of privacy problems in online transactions, Ackerman and colleagues propose the idea of *privacy critics*, semi-autonomous agents that monitor online action and can inform users about potential privacy threats and available countermeasures (Ackerman et al., 1999; Ackerman and Cranor, 1999). Again, this mechanism turns on the ability to render invisible threats visible.

Given our concern with how users manage security as an ongoing concern, one important related topic is control over the degree of security available. One of our criticisms of traditional security systems has been their "all or nothing" approach. However, there has been some work that attempts to characterize degrees of security provision, as embodied by the idea of "quality of security service." (Irvine and Levin, 2001; Spyropoulou et al., 2000). This builds on earlier work establishing a taxonomy of security service levels (Irvine and Levin, 1999). The fundamental insight is that organizations and applications need to trade-off different factors against each other, including security of various forms and degrees, in order to make effective use of available resources (Thomsen and Denz, 1997; Henning, 1999). While this work is directed towards resource management rather than user control, it begins to unpack the "security" black box and characterize degrees and qualities of security.

While the findings of these investigations are valuable and instructive, our goal here is somewhat different. We are not concerned with specific

security technologies or techniques such as passwords or quality of service, but rather with how users manage security as a practical, day-to-day problem. In this area, Freidman et al (2002) have investigated users' perceptions of security in Web-based interactions, and note that even in this restricted domain, problems arise in how people assess the security of settings they encounter. Rimmer et al. (1999) and Sheeran et al (2001) discuss the mental models that end users develop of network structure and behavior, and illustrate the impacts that these have on system use. Weirich and Sasse (2001) present a preliminary investigation of users mental models of security; their investigation is similar to ours, although they focus largely on technologically sophisticated users. In addition, our concern here is not with "mental models," per se, but with the practices through which these models operate – practices whose operation may be social rather than cognitive.

3 Methodological Issues

Our approach is broadly ethnographic in nature, based largely on semi-structured interviews that have been subjected to a qualitative analysis, drawing on the grounded theory approach (Glaser and Strauss, 1967). Grounded theory provides a set of procedures for developing analytic accounts of qualitative data, based on the iterative generation, validation, and refinement of coding schemes. A qualitative approach is more appropriate than a quantitative at this stage, given that our goal is not to provide definite answers to definite questions, but rather, to determine what questions we might want to ask in the first place. In particular, the goal of our investigations is not simply to document what users do, but rather to understand their *experience* of security as they encounter it. It is this concern – to be able to see security as it appears to the users we are studying – that drives our methodological approach. From this perspective, we gain more by understanding the experience of a small number of users in depth and detail than we would from a broader statistical account of the activities of a larger number of people.

The results presented here have been collected from a number of investigations conducted at three sites over the past twelve months or so.

Site A is an academic institution. Our interview subjects here are drawn from two pools – administrative staff members of both an academic department and a research institute, and graduate students in a management program. We were interested in these people because of their range of institutional affiliations and responsibilities. In addition, previous research (Sheehan, 2002) suggests that new graduate students should be an interesting group of people because of (as well as since they are likely to have recently reconsidered their infrastructure arrangements due to relocation.) We interviewed a total of eleven participants at Site A.

Site B is an industrial research lab. At site B we were particularly interested in finding people who, in addition to any end-user security problems, also had security needs that related to their jobs. In particular we were interested in what additional kinds of security needs these people had because their jobs required certain levels of confidentiality resulting from institutional (Federal Government) and organizational (the corporate owner of the site) rules and policies. At site B we conducted nine interviews with various members of staff in jobs that included media relations, human resources, executive administration, and legal.

Site C is a law firm. The motivation for interviewing lawyers came from a belief that they would be highly sensitized to issues of privacy, confidentiality, and trust, and because of that they would have interesting security practices. We conducted one interview in this firm.

In what follows, we discuss our findings from these sites. We will discuss these findings in three clusters of related topics. The first deals with the "scope" of security, or the range of concerns that manifest themselves when end users think about security issues. The second concerns the range of attitudes that people display towards security problems. Finally, we examine relevant aspects of the social and organizational contexts within which people encounter and solve security problems.

4 The Scope of Security

Although the mathematical and technical foundations of security systems delimit the scope of "security" for the research community, end users' see their encounters with security quite different and set the scope of concerns more broadly.

4.1 Security as a Barrier

One feature of our interview data that immediately stood out was the set of issues that arose under the auspices of security. Our questions were oriented around problems of security and information protection. However, we found that respondents would persistently turn to other issues that, to them, are intimately related to information security. Of these, perhaps the most prevalent is unsolicited email (spam). For our subjects, security and spam are two aspects of the same problem; as practical problems, viruses, network scanners, password sniffers, and unsolicited email form a "natural class," even though they may be technically quite different. What is more, conventional users not only regard these as the same problem, but also think of the same technologies as providing solutions; firewalls, for example, are described both as technologies to keep out unwelcome visitors but also unwelcome email messages. In other words, people seem to both imagine and seek unitary solutions to these problems. When we think of the real-world experiences on which people base their experiences, they think of security as a barrier, akin to a gate or a locked door. Security is, generically, something to "keep things out," and so the various threats – the things that are being kept out – become co-constructed as the common entities against which security protects.

There are three immediate implications of this observation. The first is that a security solution that solves only one problem but not others (e.g. a virus scanner) is likely to be seen as inadequate, and potentially to be rejected for that reason. Conversely, a second observation is that a technology deployed to solve one problem may be mistakenly interpreted as providing protection against the others; for example, one user at Site A talked of being protected from virii by a new filtering system installed by the network service (although, in fact, this was a spam filter with no virus detection facilities.) Third, the focus on barriers or "choke-points" diverts attention from channels, as exemplified, for instance, by users who install advanced firewalls but then run unencrypted 802.11b wireless networks.

4.2 Online and Offline

The relationship between online and offline experience is a complex one, but is also centrally important. Online conduct seems to be continually shaped by aspects of the offline world. This happens in a number of ways.

One is that a range of experiences in the offline world provide metaphors and analogies by which people understand the online world. The brand identity of large organizations, for example, is a significant factor in how people treat online entities, and particularly people felt more comfortable when dealing with organizations that had a physical presence. Institutional arrangements are perhaps even more important; banks, for instance, are seen as inherently more concerned about security, and therefore inherently more trustworthy.

A second relationship, one of greater import to our subjects, was the potential leakage of information between online and offline settings. While our initial expectation was that people would relate Internet security problems to internet-based fraud (e.g. forging email, identity theft, unauthorized financial transactions), a much more immediate concern for a significant number of our subjects was the possibility that inadvertent information disclosure online could create a threat offline. Most frequently, these were problems of personal security. Stalkers were an especially common reported threat, and much of people's attention to online information disclosure concerned information that might result in direct personal threat. We were struck by the regularity with which this issue arose, especially amongst women.

Online and offline come together in third way also: in the practical management of space and security. Computers inhabit not just a electronic world that needs to be protected but also a physical world that needs to accommodate them. Arranging the offline world to support practices in the online world was a constant source of practical security for a number of different people.

This relationship was made particularly clear to us by two separate people at Site B. The first person worked with many sensitive hard copy legal documents. A practice of locking away these hard copy documents was enforced. However, at the same time, many of these documents would be required for day to day online processing activities. The solution was to use a library cart, which was loaded with all the documents that would be required in the day's work and then wheeled into and out of

the locked area each day. The very presence of a cart full of documents (as many as 400 sheets of paper) was an indication of the relationship between online security and offline security.

The second person at Site B worked with employee data, but also received many employees as visitors to her office. She described several relationships between online and offline security. First, she positioned the screen to face her, and critically, away from the first point of entry into her office. If someone showed up to turn in a document or start a conversation they could not see what was on her screen, potentially (and often) data about an employee.

The office layout more generally was used as a means of protecting information. She had arranged her office (which she alone occupied) into two distinct sections. On entering the office, the first section was where office visitors would sit and talk with her. All the spare seats were in this part of the office, and there was a barrier (in the form of a desk) between the public visiting part of the office and the second—her private—part of the office. The front part was recognizable by the lack of paper of any sort. As she explained, the desk barrier was always kept clear, and the chairs were placed so that they would not make either the offline or online information visible to visitors.

By contrast, the back part of the office was covered in papers, her outstanding work assignments. With no spare chair in the back, along with the difficulty of walking around the desk, she had created a private part of the office. Although, in theory, someone who wanted to access the back of the office could, the social conventions of seating arrangements, barriers, were used to regulate visitors access to certain parts of the office and consequently certain types of data. Of course, both of these strategies were backed up by the presence of a locked file cabinet where files that were not being currently used were placed; these were the strategies that she used to protect data while it was being worked on.

The monitor and physical layout were not the only part of her office that needed this kind of protection. Her online work typically involved processing and referring to offline documents. For example, processing data about an employee's immigration status requires not just working on online documents but referral to physical copies of Immigration and Naturalization Service materials that need to be kept at the side of the monitor for easy reach. At the same time however, she could not have those documents seen by other employees who might potentially come into her office (since the immigration status of an employee is legally protected). Instead, she had devised a complex system of folders that allowed her to store documents away from the gaze of others and without revealing their contents by having to label them.

More than simply hiding the documents away, the folders were also used to sort the work out without making the sorting critieria visible to others through the use of written words. She used different colored folders for processing different types of information. For example, there were two colors associated with immigration (one for full time employees, one for interns—which require different kinds of processing work), and other colors for other types of confidential work. The colored folders serve two purposes simultaneously. First, they protect information from other people's sight, and second, they allow her to immediately understand what outstanding activities she has to work on without compromising the integrity of the data contained within the folders.

A fourth and final aspect of the relationship between online and offline aspects of security is particularly marked when people must deal with the physical manifestation of their networking service – the cables, routers, modems and other pieces of equipment through which their connection is maintained. Wireless networking technologies are perhaps especially interesting here due to their combination of tangible and intangible elements. Wireless networks offer people the same infrastructure interfaces that they conventionally associate with wired or point-to-point networks that conventionally carry with them obvious properties of physical security and accessibility. At the same time, however, they make the actual infrastructure of service provision intangible. One user we encountered spent some time trying to diagnose an apparently problem with a networked printer that refused to accept print jobs, before noticing that, in fact, he was connected through his neighbor's access point rather than his own (and so was connected from an unauthorized IP network). Another informant had resorted to wrapping his access point in aluminum foil in order to deal with apparent

interference problems, which turned out, in fact, to be intermittent DSL failures. The very intangibility of network infrastructure in these cases makes it harder to for people to relate online experiences to offline manifestations of technology.

In other words, what we see from our observations is that, for everyday users, security is not purely an online matter; it extends into the physical world. The information which is to be protected, the resources to be managed, and the work to be carried out all exist in a physical setting too. Security practices may draw as much on the arrangement of physical spaces as on the arrangement of technical resources and, again, providing people with technical solutions that cannot be understood or integrated into what people see as the "whole problem" will reduce their effectiveness.

4.3 Hackers, Stalkers, Spammers and Marketers

Our subjects have varying degrees of familiarity with technology, work in very different settings (physically and organizationally), and perform very different kinds of work. Similarly, a range of different situations are classed as threats, almost coextensively. We found four broad classes of threats that people brought up in discussion: hackers, stalkers, spammers, and marketers.

"Hackers" are individual threats who fit the expected image – people out to cause mischief and harm, generally highly skilled, but motivated by the same randomly violent impulse which leads to vandalism (rather than conducting targeted attacks). Hackers are broadly recognized as a threat, although they are not, in general, seen as a danger, but rather as a nuisance. The very image of hackers as "maladjusted nerds" seems to detract from concerns about identity theft, criminal intent, etc.[1]

"Stalkers" are those who, as described above, might use information gleaned online to pursue an offline threat. For many people, the protection of their online identity and information is an extension of the protection of their person and their property. The possibility for offline consequences of online activity – including but not limited to physical danger to ones own person or to ones friends and family – is a remarkably prevalent concern.

"Spammers" are organizations and individuals that advertise through unsolicited messaging, wasting people's time and using up organizational resources through an implicit denial of service. As we related earlier, unsolicited email is a major concern to people, and, critically, they see it as a security problem.

"Marketers" are people who invade individual privacy by surreptitiously collecting information about activities, purchasing patterns, and so forth. We found it interesting that the marketers are defined as threats in pretty much the same way as hackers, stalkers and spammers. People reported maintaining false identities, falsifying information, refusing to disclose identifiers, and other strategies by which they explicitly attempted to evade such tracking. To an extent, it seemed that older people seemed more likely to trust organizations and regard renegade individuals as threats; younger people were more likely to see organizations as potential threats. Age difference effects are noted by Sheehan (2002); we will have more to say about them in a moment.

5 Attitudes Towards Security

Clearly, experiences with security and systems vary between individuals. Our data suggest a range of attitudes that people display towards security.

5.1 Security as an Obstacle

We have already noted differences between younger and older participants in our study. A further interesting separation between our older and younger participants relates to this issue of experience. In general, our younger subjects, with a relatively longer exposure to computer systems (and in particular, it seems, childhood exposure) express a much greater confidence in their abilities with computer systems. In particular, they seem to have been more likely to encounter situations in which security services proved problematic, hindering rather than helping their activities. Getting files through firewalls, for instance, had been problematic for some, who found that they had to turn off or circumvent security technologies in order to get their work done. They were more likely to talk of security in terms of its costs as well as its benefits, and frame technical security measures as ones that can interfere

[1] Not that these concerns are not present, but they are less frequently associated with particular social groups.

with the practical accomplishment of work. This is, of course, the "barrier" argument when seen from the other side.

A similar example occurs in a related study of teen use of SMS (text messaging via GSM's Short Message Service) reported elsewhere (Grinter and Eldridge, 2003.) The teens studied never intentionally turned off their phones, which meant that they rarely if ever used their password to log back onto the phone after a reboot. This meant that when they accidentally let the battery run out, the teenagers described having to take their mobile phone to the nearest service center to get the password reset before they could resume receiving SMS's. This delay, in addition to creating inconveniences to have to go to the service center via public transportation, also caused consideration frustration when the teenagers realized just how many messages and potentially activities in the few hours or days it would take to get fixed.

In other cases, security appears as an obstacle in other ways. For much day-to-day use, security is not a primary concern for end users; people rarely boot their computer in order to deal with security configurations. The persistence of virus checkers, intrusion detectors, and other similar systems in interrupting current work in order to insist that something be done (new rules installed, ports blocked, or even just a button pressed) seemed to be problematic. This is, perhaps, another case of the difficulty of an "all-or-nothing" approach – security is either something unmentioned, or it is something to be dealt with suddenly and immediately.

5.2 Pragmatism

In broad terms, and in line with the previous observation, the younger respondents seemed more pragmatic about their security needs, expressing more nuance about the situations in which they might need security. For instance, they discussed using known insecure technologies in settings where they felt that the risks were justified (e.g. a machine that was known to be chock full of viruses, but was otherwise unused so it didn't matter.) This pragmatic orientation in younger subjects is in line with previous findings (Sheehan, 2002). Pragmatic users see security as a trade-off, one that must be continually struck as one balances immediate needs against potential dangers. For pragmatic users, then, systems need to be both flexible and translucent, so that these trade-offs can be made effectively.

5.3 Futility

However, even amongst those who expressed more confidence about their abilities and a more pragmatic orientation towards security, there is an overwhelming sense of futility in people's encounters with technology. This corroborates the similar observation was made by Weirich and Sasse (2001) in their investigations. Our subjects make repeated reference to the unknown others (hackers, stalkers, etc.) who will always be one step ahead, and whose skill with technologies will mean that there are always new attacks to be diverted. As a result, they talk repeatedly of security lying not so much in technology as in vigilance; the continual, active defense against new and evolving threats.

The results of this sense of futility vary depending on the setting and the forms of threat. With respect to the broad Internet, it certainly contributes to frustration and the sense that one is continually "running to stay in the same place"; it creates a fictive norm of adequate protection, against which people continually find themselves wanting. In organizational settings, it becomes manifest mainly as a concern with "due diligence" – the visible demonstration that one has done enough. As in the cases discussed earlier where security moves out of the computer and into the physical environment, the demonstration that one has taken due care to manage information and activities securely becomes important, even though subjects may not feel that these measures are likely to survive an assault.

6 Social Context

When we focus on security as a practical problem that people encounter and solve, we need also to consider the context within which it is encountered and solved. Frequently, the social and organizational context surrounding user activity is important.

6.1 Delegating Security

Unsurprisingly, most people, in the course of their daily work, have neither the time nor inclination to be continually vigilant for new threats; they are focused on getting their work done. One particularly interesting issue, then, is the various modalities by which people delegate responsibility for security.

Security is, to some extent, turned into someone else's problem, or at least, external resources are marshaled to deal with the problem. Four forms of delegation are identifiable in our interviews.

The first is to *delegate to technology*, which involves relying on some form of technology for protection. So, people might rely on SSL encryption for data connections, ssh tunneling for their email, or trust that switched Ethernet is more secure than a a traditional common medium. These are, of course, the solutions that the technical community is used to providing. Interestingly, though, this was perhaps one of the least common ways of managing security that we encountered. It is also interesting to observe that this delegation is an investment of trust, and we speculate that it depends on visible presence of technology to be trusted, which questions the idea of security as an invisible or transparent facet of a system. Arguably, the use of physical arrangements to embody security concerns is a case of delegation to technology (albeit less advanced.)

The second mode of delegation is to *delegate to another individual*, such as a knowledgeable colleague, family member, or roommate. Often, this might be someone who set the computer up in the first place; their knowledge and skill is cited as one element of a person's defense against potential threats. For people who feel limited in their ability to assess technology, known individuals may be more trustworthy.

The third mode is to *delegate to an organization*; like delegation to an individual, this delegates to others, but the others are organizationally defined and may not even be known personally. Essentially, this is the "we have a very good support group" argument. The skills and especially the vigilance of the organization is where people place their trust. In some cases, again due to the fictive norm associated with vigilance, more trust may be accorded to external organizations; and so, facilities run by a central service rather than by a local group, or facilities managed through an outsourcing arrangement, are seen as more secure.

Finally, we also found a mode in which people would *delegate to institutions*. So, our earlier examples in which financial institutions are seen as inherently more trustworthy because they are presumed to have a primary concern with security is an example of this trust in institutional arrangements and archetypes. Again, this is an online/offline relationship; impressions of the banks' concern with physical security (locked vaults and armed security guards) are carried over to online security, even though of course online interactions with a bank depend on a complex of intermediate technologies outside of any bank's control.

There is an important temporal aspect to this process of delegation. Essentially, once responsibility has been delegated, it becomes almost invisible; it seemed to be rare for these issues to be revisited. Individuals to whom responsibility had been delegated when they set up the computer sometimes disappeared from view (the former roommate or colleague, or the son who had left for college), and yet they were still invoked as the guarantor of security. In organizational settings, we found that, over time, some newer employees would not have any recollection of what kinds of access controls, for example, would be on their file systems. Delegation to the support group would have occurred several years prior to their arrival and they could not articulate what kinds of privileges existed. This is interesting for two reasons. First, the work practices of groups often "grow over" the underlying security while taking advantage of what it provides, until the security decisions are lost to conscious memory. Second, no-one concerned themselves with this. Between the initial security decision and the supporting work practices, the day-to-day configuration had disappeared but was still being enacted correctly.

6.2 Socially-Defined Security Needs

Another feature for several of the people that we interviewed was that security was not just defined as the means of securing the information and its transmission with respect to legal codes, but also with respect to how others perceived the relative sensitivity of the information. Decisions about how to handle and manage online (and even offline) data were sometimes made based on how someone else felt that the data should be treated.

For example, one person described how for certain information requests (where there was no clear institutional guidelines such as Federal laws) she would ask a number of people and then base her decision on the most conservative response. Deciding whether to grant electronic file access to

someone, or email documents, was not just a decision of what technical security parameters, but also a matter of determining what to apply to a situation.

6.3 Security as a Practice

The people we interviewed had a number of methods for managing online security, some of which involved using security protocols in what may seem like unusual ways, and others of which that appear to involve no security, but illustrate how people think about security in technology.

Whitten and Tygar's (1999) analysis of email discovered that users had incredible difficulties using PGP to secure their communications. However, what is clear is that people use email to communicate all the time, and even when they have information that needs to be protected. So, we wondered what they did to "secure" the information. We discovered two common strategies.

First, people use institutional means to secure communications. We see this each time we receive an email from someone that has an attached signature file that states the legal and illegal uses of the contents of the message. Rather than securing the communications, the purpose of these statements is to defend the contents if they become incriminated[2]. In other words, corporations often attempt to mitigate the risks of information leaks by securing the consequences of those leaks by marking the messages.

Although it may not be secure technically, it is not surprising that this approach is used. Marking documents has long been a means by which corporations sort and prioritize the contents of presentations and reports and so forth. The securing of email messages appears to coincide with the migration of email from an informal chatting technology to a formal means of corporate communications.

Second, we also found cases where people were using context to secure their email messages. By this we mean that we found cases where people described sending email that did not explicitly state what the subject was in the actual email itself, but used a shared working context to express the new information. For example, an email message that says, "I took the actions you requested" could refer to many types of activity including processing sensitive data such as updating someone's immigration status. Moreover, by removing any sense of time from the contents (other than the date stamp) no specific temporal information could be deduced.

Crucially, this arose not simply as happenstance; rather, it was an explicit strategy adopted to support secure communication as a part of a broader pattern of work. Using cryptic email rather than encrypted email offered two advantages. First, it was simply a lot easier to do than using a security tool to encrypt the information. By easier though, we do not just mean Whitten and Tygar's usability of various encryption software, we mean that context-based encryption may simply be the more visible form of security measures for many people working with email. The fact that it can be accomplished as part and parcel of the working activities, rather than as a separate and parallel activity, also makes it a more convenient feature of work practice (Smetters and Grinter, 2002).

The visibility of security systems (their presence and utility) let alone their usability is also illustrated by the use of media switching as a security measure. In our interviews, we found several occasions where people switched communications media based on security decisions. In particular, a number of people at site B described suggesting in email a switch to the telephone for the most private or sensitive discussions. Several interviewees said that they trusted the telephone for their most secure conversations, and introduced a media switch to the telephone from email when the most sensitive of topics came up.

Perhaps what is most surprising about this ability to rate technological mediums for security is the same phenomenon reported by teenagers (Grinter and Palen, 2002). Teenagers using Instant Messaging technologies also reported suggesting a media switch to the telephone for the most confidential of conversations. The telephone offers two related advantages. First, potentially, it is a more probablistically secure medium than email. Although it can be tapped and listened into, maybe it

[2] The inverse, perhaps, of the idea that it's easier to seek forgiveness than permission; it is easier to enforce through sanction than prevention.

is less statistically likely. Second, and the one more commonly articulated, the medium is ephemeral in the sense that nothing that is said is readily recorded and transmittable. Unlike electronic text, something that teenagers observed with clarity was that it was much harder to record and convince anyone else with absolute certainty what was said on the telephone. In other words, confidentiality and privacy of information can be more likely guaranteed in a medium that is not as readily recorded, and understanding the different between electronic conversation and electronic text, the audio word seemed more secure than the textual one.

Earlier, we discussed encryption and the need for secure communications and its relationship to media choice. In that and previous discussions, current security systems seemed almost to fail our users in the sense that they did not fit work practices. Moreover, they competed ineffectively with other alternatives such as simply switching media (something that was once just available to office workers is now something that teenagers at home can consider). However, we also found some occasions where security measures had been incorporated into working practices.

One example of this concerns access control to shared directories. In this case, two legal staff explained how they used the access control settings for online directories as a means of communications! Specifically, they both worked on files that once they had finished with their own notes needed to be sent to a centralized legal body for further processing. Rather than using email to do this, they used a system of shared online directories. While they worked together and locally on the files, they put them in an online directory that only they could access. When they had finished working on the file they would simply move the file to another directory, one whose access controls were set to allow other people to access it from a remote site. One advantage that the two local legal staff found with this scheme was that they did not have to know specifically who they had to send the files too (unlike email).

6.4 Managing Identity

In the interviews we discovered another challenge for security: that of identity. This manifested itself in two ways – the production of identity, and the interpretation of identity.

First, we found that our informants were very conscious of the ways in which they presented themselves online. Many, for instance, maintain many virtual identities (email addresses, online personas, etc) as a way of controlling their visibility. In addition, we found some evidence for the use of *partial* identities; by controlling which specific aspects of information (social security number, home zip code, etc) they gave to different entities, some respondents attempted to maintain control over the degree to which they could be identified and tracked.[3] When identity is seen as the sum of these factors, a subset seems secure.

The second issue is the interpretation of identity. In many of solutions proposed an individual secures him- or her- self individually. Personal firewalls, personal encryption, passwords, and so forth all place a primacy on the individual. However, we found a number of cases where seemingly individual people were in fact groups of people, and because of that security decisions and solutions became more problematic.

This problem was most clearly exhibited by individuals who had personal assistants. In many cases, the email address of an executive does not equate to the person themselves: it refers to them and their personal assistant. When we talked with assistants and people who emailed executives we discovered a mismatch in expectations and difficulties with this arrangement.

Although turning an executive's email into a group distribution list (that goes to the executive and their assistant) we found cases where people were surprised by this distribution. The metaphor implied by email is that the email goes to the individual identified by the handle that it is sent to. By contrast, group distribution lists are also visible because of the meaning conveyed in the handle used. For example *beki@company.com* is expected to go to a person by that name alone, but *football@company.com* could be considered as a group email for people interested in football.

[3] Whether they were successful in these efforts is, of course, entirely a different matter.

In our interviews we found that people were surprised when individual email addresses turned out to be group lists that contained a respondee that was not initially expected. In some cases, the revelation of an extra person raised security concerns for the surprised. We found cases where individuals who had other people read and respond to their email had to reassure surprised recipients by explaining that the other reader was a trusted individual.

In other cases, we discovered that occasions where these "extra" readers had been removed from the message. While this may have secured the information content, we also came to understand that when this involved scheduling the trade-off between security and organization was difficult. For example, we found that on occasion people would attempt to communicate with executives individually. Sometimes this would create problems for their assistants because private communications while being private would also mean that people needed to make travel, dining, or simply meeting arrangements would have no idea that commitments had been arranged.

The issue of identity management is a complex and challenging one for security systems. As Palen and Dourish (2003) observe, people act continually and simultaneously in multiple capacities – as individuals, as representatives of organizations or professional groups, as family members or members of some occupational groups, etc. The ways in which they act – the information that they chose to disclose, and how, when, and to whom they disclose it – mark them as affiliated with one or another set of people. Conventional separation into "roles" fails to capture the fluid and especially the simultaneous nature of these capacities in which one acts.

7 Dialectic Nature of Security and Privacy

The central element of our approach has been to examine security as a practical problem, as it manifests itself every day for the users of networked information systems, and as it is routinely solved by them in the course of their work. Across a range of users, a range of technologies, and a range of settings, we find a number of commonalities. One particularly relevant one, which relates to ongoing research into the nature of privacy in both online and offline settings, is the dialectic nature of security and privacy management (Altman, 1975; Palen and Dourish, 2003).

Conventional discussions of privacy have a number of common properties. Privacy is normally conceived of as a state of social withdrawal, for instance; and it is normally a normative issue. The dialectic model, though, suggests, first, that privacy may be better conceived of as a dynamic process of boundary management, characterized as much by a pressure for disclosure as by a pressure for withdrawal, and ultimately operating as a dynamic resolution between the two, managed according to circumstances and immediate needs; and second, that this dynamic process should be seen within a temporal context which is oriented both towards a history of past actions through which norms of conventional practice are defined, a future of potential needs and expectations.

We can clearly see security emerging in a similar light in our data. The very definition of what counts as "secure" is a contingent matter; "security" depends on the circumstances. Rather than being predefined and absolute, security is relative and must be continually adapted and managed in the course of system use. It is both responsive to social practice (which sets the conventions and patterns by which situations will be interpreted) and contributes to it (by creating new patterns that are shared within social groups and organizations). This perspective suggests a number of relevant design considerations.

8 Technical Responses

Clearly, as demonstrated by the data we have presented here, security is a significant concern for end users. Most importantly, though, it is a concern of a sort not typically explored by security research. Where security research has typically focused on theoretical and technical capabilities and opportunities, for end users carrying out their work on computer systems, the problems are more prosaic. They are concerned with getting their work done and with acting in ways appropriate to the settings in which they find themselves; they are concerned with communicating with colleagues; they are concerned with the details of their immediate tasks and generally not with more abstract concerns such as the specification, formulation, and configuration of security

mechanisms. Indeed, one of the things that we have learned is that it may be *inherently* implausible for typical users to specify, in advance of particular circumstances, what their security needs might be; those needs arise only as a result of specific encounters between people, information, and activities. Our goal in taking this broad look at security "as it happens" is not to specify a set of requirements for the redesign of specific security systems such as password systems, encryption mechanisms, or VPNs. Such usability-based explorations are valuable, but our concern is somewhat broader. Our investigations suggest to us that the problem of effective security will not be fixed by correcting design errors in existing technologies, but rather by respecifying the relationship between security facilities and everyday work.

As examples, consider the following implications for secure system design.

Our data points to the dialectic nature of security; that information protection and information sharing mutually define each other, and that the process of managing security is the dynamic process of managing that balance. This suggests that protection and sharing of information are two parts of the same task; they are always carried out together. Interestingly, though, most systems separate these two tasks. Typically, information sharing is a primary task, while information protection is a subsidiary task, specified in advance through control panels, configuration controls, abstract rule specification, etc. Our investigations suggest that we need to think of these as being the same tasks; I should use the same mechanisms to share information as to protect it, and, critically, they should be available to me at the same time. A split in which sharing information (adding files to a server, sending an attachment, logging into an IM service) is separated from the work of protecting it (specifying access control, encrypting the information, specifying availability) is ineffective. Essentially, practices such as maintaining multiple email addresses or AIM screen names is a practical solution that users have forged which allows them to conjoin information sharing and information protection as a unified task.

One critical issue that arises out of many of our observations is the extent to which people are able to monitor and understand the potential consequences of their actions. Since security requirements depend on the specific circumstances of action and are subject to continual reflection and revision, it is necessary to provide people with the means to understand the security implications of the current configuration of technologies at their disposal. Rather than being "transparent," then, security technologies need to be highly visible – available for inspection and examination seamlessly as a part of work. Interactive system design emphasizes that available functionality and courses of action should be continually available at-a-glance; security configuration should be available in just the same way. This, critically, is quite different from being able to "call up" security information when it's needed; the point is that it is needed as part and parcel of every activity in the system. Related research explores the technical infrastructure to develop these ideas (Dourish and Redmiles, 2002.)

Taking this one step further, we note that security is a mutual achievement of multiple parties. Like the people sending cryptic email to maintain the security of their information, people achieve security in the context of the activities that they carry out together. The security consequences of my actions depend not only on what I do but also on what my colleagues do. The scope of security, then, is a collaborative scope; it extends beyond the individual. Recognizing that the correct unit of analysis for security systems is groups rather than individuals, and that visualization approaches such as those suggested above might apply to groups, significantly changes how we conventionally think of security systems and security interfaces. Security is a collective accomplishment, an outcome of shared practices as well as shared configuration and technology.

9 Conclusions

While the research community has been extremely successful in developing the mathematical foundations of secure computing and communication services, we have, perhaps, been less successful in the more practical task of making day to day systems effectively secure. Any technology for secure communication is only as secure as the settings within which it is deployed. We have argued that a major obstacle to the development of more effective security strategies is that these systems

often match poorly to the ways in which people need to make use of them.

A small but growing group of researchers have begun to examine the usability of security technologies, and have noted a range of problems that interfere with the effective use of technologies currently employed for security in day-to-day settings. However, our argument here is broader. We believe that effective solutions will not come solely from repairing the usability problems associated with existing technologies, because the very nature of those technologies – the ways in which they conceive of the problems of security – is a source of trouble. When we look at those areas in which HCI can be said to have lead to radical improvements in usability – such as the development of graphical user interfaces and the development of the Internet into the Web – it is instructive to note that they did not arise through the incremental modification of existing technologies (e.g. systematically improving the usability of each command-line UNIX program.) Similarly, we believe that effective security will require that we examine the conceptual models on which our systems are built.

We have been exploring the conceptual foundations of users' experiences of security. Critically, we find these to be embedded in working practice, in physical settings, and in social and organizational arrangements. This seems to suggest an alternative approach to security – one that makes the security implications of actions visible and accessible in the same way that everyday actions are visible and accessible in the everyday physical and social environment. The physical and social world are organized in ways that are perceptible and rationalizable, while our security solutions tend to be invisible, unpredictable, and obscure for end-users. Rather than making security be about preventing users from doing things, the goal of our work is to make it rather be about enriching the experience of what they can do.

Acknowledgements

We would like to thank Tom Berson, Leysia Palen, David Redmiles, and Diana Smetters for their contributions to our thinking about these issues. We also gratefully acknowledge the patience and help of our interview subjects. This work has been supported in part by National Science Foundation award IIS-0133749.

References

[1] Ackerman, M. and Cranor, L. 1999. Privacy Critics: UI Components to Safeguard Users' Privacy. *Adjunct Proceedings of CHI'99* (Short Papers), 258-259.

[2] Ackerman, M., Cranor, L., and Reagle, J. 1999. Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences. *ACM Conf. on Electronic Commerce*, 1-8. ACM.

[3] Adams, A. and Sasse, M.A. 1999. Users Are Not The Enemy: Why users compromise security mechanisms and how to take remedial measures. *Comm. ACM*, 42(12), 40-46.

[4] Adams, A., Sasse, M.A., and Lunt, P. 1997. Making Passwords Secure and Usable. In Thimbleby, H. O'Connaill, B., and Thomas, P. (eds), *People and Computers XII: Proceedings of HCI'97*, 1-19. Springer.

[5] Altman, I. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory and Crowding*. Monterey, CA: Brooks/Cole Publishing Co. Inc.

[6] Bellotti, V. and Sellen, A. 1993. Design for Privacy in Ubiquitous Computing Environments. *Proc. European Conf. Computer-Supported Cooperative Work ECSCW'93*, 77-92. Kluwer.

[7] Brostoff, S. and Sasse, M.A. 2000. Are Passfaces more usable than passwords? A field trial investigation. In S. McDonald, Y. Waern & G. Cockton (Eds.): *People and Computers XIV - Usability or Else! Proceedings of HCI 2000*, 405-424. Springer.

[8] Dewan, P. and Shen, H. 1998. Flexible Meta Access-Control for Collaborative Applications Primitives for Building Flexibile Groupware Systems. *Proceedings of ACM Conference on Computer-Supported Cooperative Work CSCW'98*, 247-256. ACM.

[9] Dhamija, R. and Perrig, A. 2000. Deja Vu: A User Study. Using Images for Authentication. In *Proceedings of the 9th USENIX Security Symposium*, Denver, Colorado.

[10] Dourish, P. 1993. Culture and Control in a Media Space. *Proc. European Conf. Computer-Supported Cooperative Work ECSCW'93*, 125-137. Kluwer.

[11] Dourish, P., Edwards, K., LaMarca, A., Lamping, J., Petersen, K., Salisbury, M., Terry, D. and Thornton, J. 2000. Extending Document Management Systems with User-Specific Active Properties. *ACM*

Transactions on Information Systems, 18(2), 140-170.

[12] Dourish, P. and Redmiles, D. 2002. An Approach to Usable Security based on Event Monitoring and Visualization. *Proc. ACM New Security Paradigms Workshop NSPW 2002* (Virginia Beach, VA). New York: ACM.

[13] Friedman, B., Hurley, D., Howe, D., Felten, E., and Nissenbaum, H. 2002. Users' Conceptions of Web Security: A Comparitive Study. Short paper presented at ACM Conf. Human Factors in Computing Systems CHI 2002 (Minneapolis, MN.)

[14] Glaser, B. and Strauss, A. 1967. *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Chicago: Aldine.

[15] Greenberg, S and Marwood, D. 1994. Real-Time Groupware as a Distributed System: Concurrency Control and its Effect on the Interface. *Proc. ACM Conf. Computer-Supported Cooperative Work CSCW'94*, 207-218. ACM.

[16] Grinter. R. and Palen, L. 2002. Instant Messaging in Teen Life. *Proc. ACM Conf. Computer-Supported Cooperative Work CSCW 2002* (New Orleans, LA), 21-30. New York: ACM.

[17] Grinter, R. and Eldridge, M. 2003. Wan2tlk? Everyday Text Messaging. *Proc. ACM Conf. Human Factors in Computing Systems CHI 2003* (Ft Lauderdale, FL). New York: ACM.

[18] Henning, R. 2000. Security Service Level Agreements: Quantifiable Security for the Enterprise? *Proc. New Security Paradigm Workshop* (Ontario, Canada), 54-60. ACM.

[19] Irvine, C. and Levin, T. 1999. Towards a Taxonomy and Costing Method for Security Services. *Proc. 15th Annual Computer Security Applications Conference*. IEEE.

[20] Irvine, C. and Levin, T. 2001. Quality of Security Service. *Proc. ACM New Security Paradigms Workshop*, 91-99.

[21] Palen, L. and Dourish, P. 2003. Unpacking "Privacy" for a Networked World. *Proc. ACM Conf. Human Factors in Computing Systems CHI 2003* (Ft. Lauderdale, FL). New York: ACM.

[22] Rimmer, J., Wakeman, I., Sheeran, L., and Sasse, M.A. 1999. Examining Users' Repertoire of Internet Applications. In Sasse and Johnson (eds), *Human-Computer Interaction: Proceedings of Interact'99*.

[23] Sheehan, K. 2002. Towards a Typology of Internet Users and Online Privacy Concerns. *The Information Society*, 18, 21-32.

[24] Sheeran, L, Sasse, A., Rimmer J., and Wakeman, I. 2001. How Web Browsers Shape Users' Understanding of Networks. *The Electronic Library*, 20 (1), 35-42.

[25] Shen, H. and Dewan, P. 1992. Access Control for Collaborative Environments. *Proc. ACM Conf. Computer-Supported Cooperative Work CSCW'92*, 51-58. ACM.

[26] Smetters, D. and Grinter, R. 2002. Moving from the Design of Usable Security Technologies to the Design of Useful Secure Applications. *Proc. ACM New Security Paradigms Workshop NSPW 2002* (Virginia Beach, VA). New York: ACM.

[27] Spyropoulou, E., Levin, T., and Irvine, C. 2000. Calculating Costs for Quality of Security Service. *Proc. 16th Computer Security Applications Conference*. IEEE.

[28] Thomsen, D. and Denz, M. 1997. Incremental Assurance for Multilevel Applications. *Proc. 13th Annual Computer Security Applications Conference*. IEEE.

[29] Weirich, D. and Sasse, M.A. 2001. Pretty Good Persuasion: A first step towards effective password security for the Real World. *Proceedings of the New Security Paradigms Workshop 2001* (Sept. 10-13, Cloudcroft, NM), 137-143. ACM Press.

[30] Whitten, A. and Tygar, J.D. 1999. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. *Proc. Ninth USENIX Security Symposium*.

[31] Yee, K.-P. 2002. User Interaction Design for Secure Systems. *Proc. 4th International Conf. Information and Communications Security* (Singapore).

[32] Zurko, M.E. and Simon, R. 1996. User-Centered Security. *Proc. New Security Paradigms Workshop*. ACM.