Toward Better Usability, Security, and Privacy of INFORMATION TECHNOLOGY: REPORT OF A WORKSHOP

Steering Committee on the Usability, Security, and Privacy of Computer Systems

Computer Science and Telecommunications Board

Division on Engineering and Physical Sciences

NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

THE NATIONAL ACADEMIES PRESS
Washington, D.C.
www.nap.edu

THE NATIONAL ACADEMIES PRESS 500 Fifth Street, N.W. Washington, DC 20001

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This project was supported by the National Science Foundation under Grant No. CNS-0841126 and the National Institute of Standards and Technology under Grant No. 70NANB8H8126. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the view of the organizations or agencies that provided support for this project.

International Standard Book Number-13: 978-0-309-16090-2
International Standard Book Number-10: 0-309-16090-1

Copies of this report are available from:

The National Academies Press
500 Fifth Street, N.W., Lockbox 285
Washington, DC 20055
(800) 624-6242
(202) 334-3313 (in the Washington metropolitan area)
Internet: http://www.nap.edu

Printed in the United States of America

THE NATIONAL ACADEMIES

Advisers to the Nation on Science, Engineering, and Medicine

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Ralph J. Cicerone is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Charles M. Vest is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Harvey V. Fineberg is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Ralph J. Cicerone and Dr. Charles M. Vest are chair and vice chair, respectively, of the National Research Council.

www.national-academies.org

STEERING COMMITTEE ON THE USABILITY, SECURITY, AND PRIVACY OF COMPUTER SYSTEMS

NICHOLAS ECONOMIDES, New York University, *Chair*
LORRIE FAITH CRANOR, Carnegie Mellon University
JAMES D. FOLEY, Georgia Institute of Technology
SIMSON L. GARFINKEL, Naval Postgraduate School
BUTLER W. LAMPSON, Microsoft Corporation
SUSAN LANDAU, Radcliffe Institute for Advanced Study
DONALD A. NORMAN, Northwestern University
CHARLES P. PFLEEGER, Pfleeger Consulting Group

Staff

JON EISENBERG, Director, Computer Science and Telecommunications Board
NANCY GILLIS, Program Officer (through January 2010)
SHENAE BRADLEY, Senior Program Assistant

Preface

Usability has emerged as a significant issue in ensuring the security and privacy of computer systems. More-usable security can help avoid the inadvertent (or even deliberate) undermining of security by users. Indeed, without sufficient usability to accomplish tasks efficiently and with less effort, users will often tend to bypass security features. A small but growing community of researchers, with roots in such fields as human-computer interaction, psychology, and computer security, has been conducting research in this area.

With sponsorship from the National Science Foundation and the National Institute of Standards and Technology, the National Research Council's Computer Science and Telecommunications Board conducted a 2-day workshop in July 2009 to identify promising research directions that would help advance usability, security, and privacy. It was also intended that the workshop would build awareness—in the research community as well as in federal agencies and the broader technical community responsible for the design, development, and deployment of information systems—of the challenges at the nexus of usability and security/privacy, the trade-offs that exist today, and the opportunities for making advances. A single workshop of this sort cannot be comprehensive; indeed, important topics such as the special usability considerations faced by those with impairments were not covered.

The Steering Committee on the Usability, Security, and Privacy of Computer Systems was convened to plan the workshop (biosketches of the steering committee members can be found in Appendix C). The work-

BOX P.1
Statement of Task

An ad hoc committee will plan and conduct a public workshop on ways to advance the usability, security, and privacy of computer systems. The workshop will feature invited presentations and discussions on the state-of-the-art in usability, security, and privacy and how usability contributes to security and privacy. The agenda should include topics on ways to mutually advance objectives in usability and security/privacy especially in cases that replace trade-offs (e.g., between usability and security) with win-win scenarios. It should also include topics on research opportunities and potential roles for the federal government, academia, and industry and ways to embed usability considerations in research, design, and development related to security, privacy and vice versa. A report of the workshop will be issued.

shop was designed to identify research opportunities and potential roles for the federal government, academia, and industry and ways to embed usability considerations in research, design, and development related to security and privacy, and vice versa (the formal statement of task appears in Box P.1).

This report summarizes the workshop. As a workshop report, it does not necessarily reflect the consensus views of the committee or the workshop participants, and the committee was not asked to provide findings or recommendations.

The workshop was structured to gather suggestions from experts on computer security, privacy, and usability, as well as from economists and sociologists on new research topics within the intersection of usability, security, and privacy. It also involved a number of federal government representatives interested in usability, security, and privacy research. A detailed agenda can be found in Appendix A, and a list of workshop participants can be found in Appendix B.

The workshop featured two overview presentations, the first addressing computer security and the second addressing usability (summarized in Chapter 2). It also included six presentations intended to provide an overview of current and prospective research topics (summarized in Chapter 3). Following these talks, workshop participants split into smaller groups that discussed research needs and opportunities, addressing the topics listed in Appendix A. They were provided in advance with a set of potential research questions developed by the steering committee. The committee's summary of results from the breakout sessions is presented

in Chapter 4. Chapter 5 discusses overarching questions in advancing research in usability, security, and privacy.

The committee thanks the workshop participants for their thoughtful presentations and discussion. It also acknowledges the financial support provided by the project's sponsors, the National Science Foundation (NSF) and the National Institute of Standards and Technology (NIST), and it appreciates the encouragement and support of Mary F. Theofanos (NIST) and Karl N. Levitt and C. Suzanne Iacono (NSF).

Nicholas Economides, *Chair*
Steering Committee on the Usability, Security, and Privacy of Computer Systems

Acknowledgment of Reviewers

This report has been reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the National Research Council's (NRC's) Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the institution in making its published report as sound as possible and to ensure that the report meets institutional standards for objectivity, evidence, and responsiveness to the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process. We wish to thank the following individuals for their review of this report:

Steven M. Bellovin, Columbia University,
Bob Blakley, Gartner, Inc.,
Tadayoshi Kohno, University of Washington,
Eric Sachs, Google, Inc., and
Stuart E. Schechter, Microsoft Research.

Although the reviewers listed above have provided many constructive comments and suggestions, they were not asked to endorse the views expressed, nor did they see the final draft of the report before its release. The review of this report was overseen by Joseph F. Traub, Columbia University. Appointed by the NRC, he was responsible for making certain that an independent examination of this report was carried out in accor-

dance with institutional procedures and that all review comments were carefully considered. Responsibility for the final content of this report rests entirely with the authoring committee and the institution.

Contents

1

Overview of Security, Privacy, and Usability

This overview briefly discusses computer system security and privacy, their relationship to usability, and research at their intersection. The chapter is drawn from remarks made at the National Research Council's (NRC's) July 2009 Workshop on Usability, Security, and Privacy of Computer Systems as well as recent reports from the NRC's Computer Science and Telecommunications Board (CSTB) on security and privacy.[1]

SECURITY

Society's reliance on information technology (IT) has been increasing simultaneously with the ability of individuals, organizations, and state actors to conduct attacks on computer systems and networks. IT has become essential to the day-to-day operations of companies, organizations, and government. People's personal lives also involve computing in areas ranging from communication with family and friends to online banking and other household and financial management activities. Companies large and small are ever more reliant on information systems to support diverse business processes, including payroll and accounting, the tracking of inventory, the operation of sales, manufacturing, and research

[1] National Research Council, *Toward a Safer and More Secure Cyberspace*, Seymour E. Goodman and Herbert S. Lin, eds., The National Academies Press, Washington, D.C., 2007; and National Research Council, *Engaging Privacy and Information Technology in a Digital Age*, James Waldo, Herbert S. Lin, and Lynette I. Millett, eds., The National Academies Press, Washington, D.C., 2007.

and development—that is, computer systems are increasingly needed for organizations to be able to operate at all. Critical national infrastructures—such as those associated with energy, banking and finance, defense, law enforcement, transportation, water systems, and government and private emergency services—also depend on information systems and networks. The telecommunications system itself and the Internet running on top of it are critical infrastructure for the nation. Information systems play a critical role in many governmental functions, including national security and homeland and border security.

The conventional definition of computer security relates to the following attributes of a computer system: confidentiality (the system prevents unauthorized access to information), integrity (information in the system cannot be altered without authorization), and availability (the system is available for authorized use). Authentication—the verification of identity using some combination of something that one knows (such as a password), something that one has (such as a hardware token), and something that one is (such as a fingerprint)—is often thought of as an additional essential security capability. Reliability is a closely related concept—a reliable system performs and maintains its functions even in hostile circumstances, including but not limited to threats from adversaries.

Nearly all indications of the severity of the security threat to computer systems, whether associated with losses or damage, type of attack, or presence of vulnerability, indicate a continuously worsening problem.[2] The potential consequences fall into three broad categories:

- *Economic drag*—To counter security problems, organizations are forced to spend in order to defend and strengthen insecure IT systems.
- *Avoidance*—Because of the perceived security risks of computing, individuals or organizations avoid using IT systems, thereby missing the potential benefit of their use.
- *Catastrophe*—Failure of an IT system causes major economic loss and perhaps even loss of life. A catastrophe could be the result of a cyberattack, a serious software design or implementation flaw, or system misuse.

Despite advances that have been made in both practice and technology, cybersecurity will be a concern into the foreseeable future. More and more sensitive information will be stored in systems whose security does not necessarily increase in proportion to the value of the assets they contain. The threats will continue to evolve both on their own and as defenses against them are discovered and implemented. New vulnerabilities will emerge as previously unknown weaknesses are uncovered and as innova-

[2] NRC, *Toward a Safer and More Secure Cyberspace*, 2007, p. 2.

tion leads to the use of IT in new applications and the deployment of new technologies. The growing complexity of IT systems and the fast-growing importance of network access and network-intermediated computing are likely to increase the emergence of new vulnerabilities.

PRIVACY

Information privacy concerns the protection of information about individuals and other entities. The environment for privacy is dynamic, reflecting societal shifts (e.g., increases in electronic communication), varying and evolving attitudes (e.g., across generations or cultures), and discontinuities (e.g., events and emerging conditions that rapidly transform the national debate, such as the September 11, 2001, attacks and the global response to them) as well as technological change. The decreasing cost of storage combined with the increase in communications devices, including, and especially, mobile ones, has led to remarkable impacts on personal privacy within a very short period of time. Private information can be compromised by attacking networks and computers directly or by tricking users into revealing the information or the credentials required to access it.[3] Protecting privacy often occurs in the face of competing interests in the collection or use of particular information, and addressing privacy issues thus involves understanding and balancing these interests.

USABILITY

Usability may be thought of narrowly in terms of the quality of a system's interfaces, but the concept applies more broadly to how well a system supports user needs and expectations. The International Organization for Standardization (ISO) 9241-11 standard defines usability as "the extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use."[4] A framework attributed to both Nielsen[5] and Shneiderman[6] describes usability in terms of learnability, efficiency of use, memorability, few and noncatastrophic errors, and subjective satisfaction. Usability relates not only to understanding what taking a particular action means in

[3] One example of the latter is *phishing,* which refers to attempts to acquire sensitive information such as passwords by pretending in an e-mail or other communication to be a trustworthy entity.

[4] International Organization for Standardization (ISO), *Ergonomics of Human System Interactions: Guidance on Usability* (Part 11), ISO, Geneva, 1998.

[5] Jakob Nielsen, *Usability Engineering,* Academic Press, San Diego, Calif., 1993, p. 26.

[6] Ben Shneiderman, *Designing the User Interface: Strategies for Effective Human-Computer-Interaction,* Addison-Wesley, Reading, Mass., 1992.

the context of a particular interaction, but also to whether the user understands the implications of his or her choices in a broader context. Information system design and development inevitably embed assumptions and values, both implicit and explicit, that have impacts on a system's users; these considerations may be thought of as another aspect of usability.

USABILITY, SECURITY, AND PRIVACY

Despite many advances, security and privacy often remain too complex for individuals or enterprises to manage effectively or to use conveniently. Security is hard for users, administrators, and developers to understand, making it all too easy to use, configure, or operate systems in ways that are inadvertently insecure. Moreover, security and privacy technologies originally were developed in a context in which system administrators had primary responsibility for security and privacy protections and in which the users tended to be sophisticated. Today, the user base is much wider—including the vast majority of employees in many organizations and a large fraction of households—but the basic models for security and privacy are essentially unchanged.

Security features can be clumsy and awkward to use and can present significant obstacles to getting work done. As a result, cybersecurity measures are all too often disabled or bypassed by the users they are intended to protect.[7] Similarly, when security gets in the way of functionality, designers and administrators deemphasize it. Workshop participant Don Norman quipped, "The more secure a system, the less secure the system"—that is, when users find that security gets in their way, they figure out ways to bypass it.[8] Indeed, some participants suggested, it may be the dedicated workers who are most highly motivated to defeat security measures.

The result is that end users often engage in actions, knowingly or unknowingly, that compromise the security of computer systems or contribute to the unwanted release of personal or other confidential information. For example, industry reports, such as the one issued in 2008 by the

[7] A recent paper by Herley explains that "security advice is a daily burden, applied to the whole population, while an upper bound on the benefit is the harm suffered by the fraction that become victims annually." C. Herley, "So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users," New Security Paradigms Workshop 2009, Oxford.

[8] This observation was published following the workshop in D.A. Norman, "When Security Gets in the Way," *Interactions* 16(6): 60-63, 2009; a similar observation ("More onerous security requirements can lead to less secure situations") appears in D.A. Norman, *Living with Complexity*, MIT Press, Cambridge, Mass., 2010, Chapter 3, in press.

Verizon Business RISK Team, have highlighted the impact that end users have on system security. As the Verizon report observed:

> [L]oosely defined, error is a contributing factor in nearly all data breaches. Poor decisions, misconfigurations, omissions, non-compliance, process breakdowns, and the like undoubtedly occur somewhere in the chain of events leading to the incident.[9]

Usability and security are thus attributes that can trade off against each other. For example, requiring users to change their passwords periodically may improve security but places a greater burden on users. (Poor usability may also reduce security by driving users to workarounds, such as when users tape hard-to-remember passwords to their workstations.) Or, a password may be replaced by a hardware token; this relieves the user of having to remember a password but imposes a new burden on the user to carry the token wherever that access is required.

Poor usability is also an impediment to privacy protection. For example, a privacy policy or privacy settings that are difficult to understand or navigate make it difficult for users to know what privacy choices they have made or to change the settings to best reflect their preferences.

Usability, security, and privacy are all especially challenging aspects of system design. For example, although well-established techniques exist for testing the usability of a system, at least in the narrow sense of the quality of the system's interface, much less is known about how to effectively embed usability considerations in a specification. Better user models might help in the identification of usability requirements and more generally speed development. More sophisticated models might make it easier to strike the right balance between usability and risk mitigation. Moreover, usability, security, and privacy have all come to be understood as attributes that must be addressed throughout a system's development life cycle. Early decisions about architecture, data structures, and so forth can have a large impact on what sorts of usability aspects are even feasible. Finally, both usability and security/privacy considerations are not finished once a product or system is released, but need to be kept in mind through the life cycle of use—assumptions, norms, and expectations may change over time. Data about these factors can be gathered and taken into account during system updates and revisions.

[9] Verizon Business RISK Team, *2009 Data Breach Investigations Report,* Verizon business. Available at http://www.verizonbusiness.com/products/security/risk/databreach; accessed February 16, 2010.

USABILITY, SECURITY, AND PRIVACY: AN EMERGING DISCIPLINE

A small but growing research community has been working at the intersection of usability, security, and privacy—one that draws on expertise from multiple disciplines including computer security, human-computer interaction, and psychology. Participants noted that as an emerging and multidisciplinary discipline, it is sometimes viewed as too "soft" by some engineers and scientists and that it does not always have buy-in from those responsible for managing the development and operation of computer systems. There has, however, been growing interest in the field from the more traditional disciplines. Papers at the intersection have appeared occasionally at traditional security conferences for many years, but until recently there have been few sustained research efforts in this area. Exploratory workshops held in 2003 and 2004 led to the organization in 2005 of the first formal conference on this topic, the Symposium on Usable Privacy and Security (SOUPS), which has been held annually since then. Increasingly, usable security and privacy papers are also appearing at traditional security conferences and human-computer interaction conferences, more academic and industry researchers are focusing their research in this area, several universities now offer courses in this area,[10] and the National Science Foundation's Trustworthy Computing program highlights usability as an important research area.

[10] For example, courses have been offered by Carnegie Mellon University ("Usable Privacy and Security"; see http://cups.cs.cmu.edu/courses/ups.html), and Harvard University ("Security and Privacy Usability"; see http://www.seas.harvard.edu/courses/cs279/syllabus.html).

2

Framing the Security and Usability Challenges

Talks by Butler Lampson and Donald Norman provided workshop participants with an overview of key challenges related to security and usability. Lampson's presentation discussed the current state of computer security and its relationship to usability considerations. Norman's remarks centered on the issue of design as it relates to usability, security, and privacy. The following sections summarize these remarks.

AN OVERVIEW OF THE STATE OF COMPUTER SECURITY (BUTLER LAMPSON)

Computer security today is in bad shape: people worry about it a lot and spend a good deal of money on it, but most systems are insecure. The primary reason for this poor state of computer security, Lampson argued, is that users do not have a model of security that they can understand. Lampson suggested that research is needed to decide whether appropriate models can be elicited from what users already know, or whether there is a need to invent and promote new models.

Metrics play an important role in addressing the state of computer security. Security is about risk management: balancing the loss from breaches against the costs of security. Unfortunately, both are difficult to measure. Cost is partly in dollars budgeted for firewalls, software, and help desks but mostly in the time that users spend typing and resetting passwords, responding to warnings, finding workarounds so that

they can do their jobs, and so forth. Frequently the costs and risks are unknown, and there are no easy ways to estimate them.

A proper allocation of economic incentives is essential to improving computer security. Users, administrators, organizations, and vendors respond to the incentives that they perceive. Users just want to get their work done. Without an appropriate understanding of the risks involved and how proper security may help avoid those risks, they view security as a burden, causing them to ignore it or to attempt to work around it. Organizations do not measure the cost of the time that users spend on security and therefore do not demand usable security. Vendors thus have minimal incentive to supply it.

Many people think that security in the real world is based on locks. In fact, real-world security depends mainly on deterrence and hence on the possibility and severity of punishment. The reason that one's house is not burgled is not that the burglar cannot get through the lock on the front door; rather, it is that the chance of getting caught, while small, together with a significant punishment, makes burglary uneconomic. It is difficult to deter attacks on a computer connected to the Internet because it is difficult to find "the bad guys." One way to fix this is to communicate only with parties that are accountable, that one can punish. There are many different punishments: money fines, ostracism from some community, firing, jail, and other options.

Some punishments require identifying the responsible party in the physical world, but others do not. For example, to deter spam, one might reject e-mail unless it is signed by someone known to the receiver or unless it comes with "optional postage" in the form of a link certified by a trusted third party, such as Amazon or the U.S. Postal Service; if one clicks the link, the sender contributes a dollar to a charity.

The choice of safe inputs and the choice of accountable sources are both made by *one's own* system, not by any centralized authority. These choices will often depend on information from third parties about identity, reputation, and so forth, but which parties to trust is also one's own choice. *All trust is local.*

To be practical, accountability needs an ecosystem that makes it easy for senders to become accountable and for receivers to demand it. If there are just two parties, they can get to know each other in person and exchange signing keys. Because this approach does not scale, there is also a need for third parties that can certify identities or attributes, as they do today for cryptographic keys. This need not hurt anonymity unduly, since the third parties can preserve anonymity except when there is trouble, or accept bonds posted in anonymous cash.

This scheme is a form of access control: you accept input from me only if I am accountable. There is a big practical difference, though, because

accountability allows for punishment or the possibility to undo things that should not have been allowed to occur. Auditing is crucial, to establish a chain of evidence, but very permissive access control is acceptable because one can deal with misbehavior after the fact rather than preventing it up front.

One obvious problem with accountability is that one often wants to communicate with parties about whom one does not know much, such as unknown vendors or gambling sites. To reconcile accountability with the freedom to go anywhere on the Internet, one should, Lampson suggests, use two (or more) separate machines: a *green* machine that demands accountability and a *red* one that does not.

On the green machine one keeps important things, such as personal, family, and work data, backup files, and so forth. It needs automated management to handle the details of accountability for software and Web sites, but one chooses the manager and decides how high to set the bar: like one's house or like a bank vault. Of course the green machine is not perfectly secure—no practical machine can be—but it is far more secure than what is generally available today.

On the red machine one lives wild and free, not putting anything there that one really cares about keeping secret or not losing. If anything goes wrong, the red machine is reset to some known state.

Things are so bad for usable security, Lampson concluded, that it will be necessary to give up on perfection and focus on essentials. The primary cause of the problem is metrics and incentives: the costs either of getting security or of not having it are not known, so users do not care much about it. Therefore, vendors have no incentive to make security usable.

To fix this, it is necessary to measure the cost of security, and especially the time that users spend on it. Simple models of security that users can understand are needed. To make systems trustworthy, accountability is needed, and to preserve freedom, separate green and red machines are needed, to protect things that one really cares about from the wild things that can happen on the Internet.

USABLE SECURITY AND PRIVACY: IT'S A MATTER OF DESIGN (DONALD NORMAN)

Among the recurring questions at the workshop were these: Does added security make things more difficult to use? Will people always resent the extra steps? The answer to both questions is the same: Not necessarily. Consider the physical world of doors and locks mentioned earlier: one can see that they can get in the way of easy access but are tolerated because they seem necessary and because the amount of effort required to open them usually seems reasonable. This example highlights

two key design issues: (1) the importance of users (and vendors) understanding the necessity for protection and (2) the reasonableness of the effort required.

Different groups are involved in ensuring the security of a computer system, each group requiring a different form of design assistance. System developers provide the underlying mechanisms, but the information technology (IT) administrators at the various sites determine just how those policies are to be enforced. The IT staff is under considerable pressure from its own administration to reduce security and privacy concerns, but to do so it must be well versed in technology, in the law, in the needs of the user community, and in the psychology of both the legitimate and the illegitimate users. What the community needs, Norman suggested, is a set of standardized scripts, templates, and system tools that allows them to implement best practices in ways that are both effective and efficient, standardizing interactions across systems in order to simplify the life of users but still tailoring the requirements to any special needs of the organization. These tools do not exist today.

In the absence of standard guidelines and adequate tools, different systems implement the same policies with very different philosophies and requirements, complicating life for people who must use multiple systems. Developers who lack an understanding of real human behavior tend to impose logical rules and requirements on a bewildered, overwhelmed audience. The users, either not understanding the rationale or simply disagreeing with the necessity for the procedures imposed on them, see these as impediments to accomplishing their jobs. Moreover, the system developers may lack understanding of the clever ruses and social engineering skills of the illegitimate users, who break into systems the easy way: by lying, stealing, and deceiving. The strongest locks in the world do not deter the clever social engineer.

Security and privacy are difficult problems. Norman suggested that a way to improve security is to design systems that are easy to use for their intended purposes or by the intended people, but difficult for non-authorized people or uses. For these purposes, Norman added, one needs to consider components not normally considered in simple product design: means of authenticating identities or authority, needs, and permissions.

It also means undertaking research to ensure that systems are accompanied by a clear and understandable conceptual model, Norman concluded. Individuals do appear willing to adapt to the inconvenience of locks that seem reasonable for protection, but not to those that just get in the way. If people understand why they are required to implement security protocols, they might be more willing to pay a reasonable penalty of inconvenience.

3

Current Research at the Intersection of Usability, Security, and Privacy

Six workshop speakers who work at the forefront of usability, security, and privacy and associated fields were asked to discuss the challenges, applicable research, and potential research needs associated with usability, security, and privacy. Their remarks are summarized below.

USABLE PRIVACY (LORRIE FAITH CRANOR)

Privacy has been described as an "adjustment process" in which humans continuously adjust the views of themselves that they present to others. In the online world, humans often rely on software tools to help them manage this process. However, many currently available privacy tools are difficult to use. Lorrie Faith Cranor's presentation addressed areas in which usability research is needed in order to provide more effective privacy protection and explored areas in which some privacy goals may appear to conflict with other privacy goals, usability goals, or security goals.

Cranor began her talk by observing that privacy is hard to define, and quoted from a paper by Robert C. Post in the *Georgetown Law Journal:* "Privacy is a value so complex, so entangled in competing and contradictory dimensions, so engorged with various and distinct meanings, that I sometimes despair whether it can be usefully addressed at all."[1] She went on to provide a variety of definitions of privacy that have been offered

[1] Robert C. Post, "Three Concepts of Privacy," *Georgetown Law Journal* 89: 2087, 2001.

by public figures and in legal and other academic literature. The myriad definitions have at their core the basic notions limiting access to and providing control over personal information or contact with individuals.[2]

Access and control can be provided through either technical or legal and regulatory measures. Access can be limited using either laws that prohibit or limit the collection and disclosure of information or technology that facilitates anonymous transactions or otherwise minimizes disclosure. One way to provide control over personal information is through laws and regulations that mandate choice, the choice either to opt in or to opt out. Another is the use of technology that facilitates informed consent, such as tools to keep track of and enforce privacy preferences.

Although work in the past has often focused on information collected by Web sites, a wide array of current and emerging technologies will have significant impacts on privacy, including behavioral advertising, social networks, deep packet inspection, server log files, and location sharing. All of these technologies raise questions about how to communicate meaningfully about the effects that these technologies will have on privacy and about how to help people understand privacy risks that may seem distant or not relevant to them today. Related to this, different rates and patterns of use and the acceptance of these technologies suggest that different types of communication may be necessary to reach people in different age groups, of different genders, or in different cultures.

Cranor drew the connection between privacy and usability, observing that the privacy concerns that people often express seem inconsistent with their actual behavior—that is, people say that they want privacy but do not always take the steps necessary to protect it. There are many possible explanations—for example, people may not actually care all that much about privacy, or they may favor short-term benefits that may come at the cost of privacy over the long-term consequences for their privacy. But there are other possible explanations for the gap between expressed concerns and behavior: people may not understand the privacy implications of their behavior; the cost of privacy protection may be too high (including the cost of figuring out what steps should be taken to protect their privacy); or users might think that they have taken steps to protect their privacy but misunderstood those steps and actually did not. All three possibilities directly implicate usability.

One case where usability issues impede privacy protection is the use of privacy policies, which are intended to inform consumers about pri-

[2] Two recent CSTB reports explored these definitional issues: see National Research Council, *Who Goes There: Authentication Through the Lens of Privacy,* The National Academies Press, Washington, D.C., 2003; and National Research Council, *Engaging Privacy and Information Technology in a Digital Age,* The National Academies Press, Washington, D.C., 2007.

vacy practices and to help them decide whether those practices are acceptable or whether to opt out. However, Cranor observed that most policies are difficult to read, long, and subject to frequent change, with the result that few people read privacy policies; this suggests that privacy policies of the sort common today do not really enable consumers to exercise effective control over their personal information. Meaningful control is only possible if individuals can understand what their options are and what the implications of these options are, if they have the means to exercise the options, and if the costs (in terms of money, time, convenience, and cost versus benefit) are reasonable. Cranor described a research effort in which she is involved that aims to address these issues through the development of standardized, easy-to-read labels akin to nutritional labeling on food.

Another case is privacy configuration management. How can the creation of privacy rules be simplified even though the context may be very complex? How can people be allowed to establish privacy preferences easily up front for a range of applications? How can people be helped to realize when adjustments to these settings are needed and to adjust them easily or automatically? Cranor described a research effort studying some of these privacy configuration issues: the location-finding service, Locaccino, developed at Carnegie Mellon University. The application includes capabilities for defining with whom, when, and where location information is shared. It also provides information about who has asked to view a user's location and who can view that information currently, and it is instrumented to collect feedback on how comfortable users are with this information.

The compelling functionality as well as the significant privacy impacts of location-finding services is illustrative of the conflicts that can arise. How can the need to store information be balanced with the need to discard information to provide privacy? Examples of such conflicts involve not only information used to improve application functionality but also information used to automate privacy configurations. Similar tensions arise between privacy and other interests, such as the need to store access data for auditing purposes versus the need to protect employee privacy, or the needs of law enforcement versus the need to discard information to protect privacy. Are there technical solutions that can preserve privacy while enabling these functions?

Anonymity tools can enhance privacy in certain situations. These tools typically hide users in cover traffic or send traffic by way of a circuitous route that is difficult to trace back. Users typically give up speed, convenience, or functionality in exchange for this anonymity. The tools must also be turned on and off, which is cumbersome and requires explicit user action. Are there ways of providing anonymity without degrading the user experience?

Cranor ended her talk by presenting a series of slides listing a number of the research questions discussed above. She closed by posing three questions with broad implications for privacy and usability as well as future research on these topics:

- As today's youth grow up with their lives online, will they come to expect less privacy?
- As we increasingly trade off privacy for convenience and functionality, are we doomed to a slow erosion of privacy that eventually leaves us with minimal expectations of privacy?
- Can "usable privacy" be designed into technology to provide convenience and functionality without sacrificing privacy?

ECONOMIC ISSUES OF USABLE SECURITY AND PRIVACY (NICHOLAS ECONOMIDES)

The talk by Nicholas Economides addressed how the incentives of both users and companies with respect to usable security and privacy are not currently structured to maximize social benefit.[3] Most users do not have sufficient incentives to secure their computers to prevent network-wide catastrophic events, and they might find it very difficult to implement sufficient security even if they had sufficient incentives. What economic and legal policies can be implemented to change the incentives of users, software and hardware companies, firms conducting electronic commerce, and companies providing online services such as search so that they are closer to maximizing social benefit? What are some possible economic motivators for usable security and privacy from the perspective of the end user, private companies, and society? How do economic incentives change when viewed domestically versus globally?

Economides began by noting the significant security deficiencies of computing devices and software today, the complexity of the interfaces that define security functionality, and the poor knowledge that users typically have about the level of privacy present in the software and services that they use. The Internet is widely understood to have both multiplied the security problems of connected devices and highly increased the global impact that results from a local lack of security. Indeed, typical users have a very limited understanding of the network capabilities of their computers and the possibilities of abuse in a network setting.

[3] A similar phenomenon was noted in Don Davis, "Compliance Defects in Public-Key Cryptography," *Proceedings of the 6th Usenix Security Symposium*, San Jose, Calif., 1996, pp. 171-178, available at http://world.std.com/~dtd/#compliance.

The question of incentives can be approached from a number of perspectives, such as those of the individual or residential user; private companies (which have different perspectives depending on the nature of their business); the overall network or societal interests; vendors of hardware, software, and services; and Internet service providers.

Even individual users face a myriad of choices with respect to their activities that depend on computing, communications, and storage capabilities. It is not clear that users do—or even can reasonably be expected to—understand the financial or other consequences to themselves or others from poor security in any of these choices. Do users have sufficient economic incentives (either rewards or penalties) to use sufficient security? Improved usability of security would make it possible for at least those users who aim for higher security to achieve it at reasonable cost.

Private firms' views on security and privacy vary widely. Some firms, such as banks, investment brokers, and electronic commerce firms, generally desire higher levels of security and have found various private solutions to make their transactions more secure. (The level of security achieved and the investment that they make reflect such firms' view of the costs and benefits and will not necessarily provide a level of security demanded by broader societal interests.) Other firms, such as online advertisers, tend to favor more retention or disclosure of private information so that they can use this information to identify products and services that better match consumer preferences. Economides observed that, as a result, a very secure online world in which users are made fully aware of the impact of disclosures of their private information would cut into the profits of these firms. Other firms that produce operating systems and other software have not fully adjusted to today's world in which the exploitation of even small security flaws can have global consequences. Operating systems' producers do not face full liability for the damage that may be caused by security flaws. Once sold, many systems will persist for years; security issues and questions about incentives apply not only at the time of purchase but also throughout the useful life of the product. Internet service providers (ISPs) have an interest in furthering the security of end users, given that breaches can affect their own networks; ISPs may also view security as an attractive value-added business. Given these diverse perspectives, Economides observed that a consensus among companies on security and privacy is unlikely.

From a societal point of view, the value of security is much higher for the network than it is for an individual user. That is, users, left on their own, will generally tend to achieve lower security than what society desires. Low security at the nodes can lead to catastrophic network events that are much more damaging to society than to the individual node. The owner of the node does not face the network-wide financial and other

liability that low security at the node causes. The lack of security at a node is, therefore, a negative externality to the network. Similar considerations apply to the vendors of hardware, software, and services.

Economides posed related questions about the incentives for security:

- What legal and economic policy changes would help improve the usability of the security of operating systems, Web sites and services, or Internet service providers?
- How can the usability of security be improved (and thus its cost reduced) so that users who aim for higher security are better able to achieve it?
- When usable security is available, how can economic incentives be created so that users will aim for sufficient security?

A variety of potential incentives might be considered. These include positive monetary incentives, awards and other nonmonetary positive incentives, and punishments. Negative incentives would include end-user liability for damage caused by insecure nodes, liability for vendors, or regulation. For example, regulations could prohibit computers that fail a basic security test from being connected to the Internet, or they could prohibit systems from being shipped with known insecure default settings. There are also thorny policy issues that apply in individual sectors. For example, blocking access on the basis of a security test limits to some extent the rights of computer owners. Also, there may be a tension between asking ISPs to play a greater role in limiting or preventing some attacks and ensuring that carriers comply with network neutrality principles such as not prioritizing content.

Economides closed by posing the following key questions regarding incentives for security and privacy:

- How can society best deal with the negative externality for the network and society that is created by the lack of usable security of individual network nodes?
- How can positive and negative, monetary, and nonmonetary incentives be provided to both users and private-sector firms to reduce or eliminate the negative externality?
- How can the usability of security be improved so that the costs are lowered for users who aim to achieve higher security?

WHAT WOULD USER-CENTERED SECURITY LOOK LIKE? (ANGELA SASSE)

Angela Sasse started with the observation that user-centered approaches to designing technology start with understanding user requirements. To do that, researchers and developers try to establish the following:

- The needs of the target users, plus specific capabilities or limitations that they have;
- The tasks and business processes that the users have to perform; and
- The physical, cultural, and situational context in which the interaction takes place.

However, since security is not a primary goal of users (protecting data, transactions, and systems is secondary to "getting the job done"), users often experience security as something that gets in the way of their activities as opposed to being something that is valuable. How can security be made less of a "barrier" that gets in the way of user goals? How can the user effort required be reduced? When is it reasonable to expect users to expend extra effort on security? What are existing user needs and values that could be connected to security?

Sasse then turned to the reasons that usability is important for security. She observed that the results of failure to make security usable are much more widespread than is generally realized. For users, this failure manifests itself as errors, frustration, annoyance, and individual loss of productivity. For organizations, there are the risks of system failure, the alienation of customers, and damage to organizations' reputations and impacts on their business processes and performance. For society, security ends up being seen as an annoyance or obstacle rather than as something that should be valued. Poor security makes possible attacks that undermine trust and confidence.

Sasse offered a framework for thinking about usability that includes the following elements: the users and actors (including individuals and organizations), the activity (the goals of the interaction [the "what"] and the tasks and processes to be carried out to achieve those goals [the "how"]), and the context (including physical, situational, and cultural aspects). In addition, one must consider the system or technology platform in question.

In terms of users, one must understand their requirements and capabilities, which include not only such factors as human memory, propensity to make errors, fatigue, biases, and the like, but also what the users are trying to achieve. Another consideration is the specific capabilities of users; because it is often essential to use security capabilities in order to

gain access to services, accommodations should be made for user groups that have particular requirements.

In terms of activity, it is important realize that security is a secondary or enabling activity. From a user's perspective, security at best slows down the completion of a task, and at worst it can prevent the user from achieving a goal. From an organization's perspective, security consumes resources and slows down business processes; at worst it may stop business processes altogether. As a result, the needs of business processes and user tasks impose performance requirements on security tasks.

A number of contextual factors have a bearing on usability and privacy. These include the physical environment, situational factors such as the impact of interactions and failures, and cultural factors. Cultural factors include behavioral norms such as the acceptability of touching equipment, or reactions to the prohibitions on smiling associated with some face-recognition systems.

Security has both costs and benefits. Individual costs include the physical workload (e.g., additional keystrokes or mouse clicks) and mental workload (e.g., remembering passwords). Both actual and perceived costs are relevant. Organizational costs include the cost of operating security capabilities (including training and maintenance) and the cost when these capabilities fail. The impacts of security extend beyond business efficiency to employee behavior, trust, and goodwill. These costs and benefits are weighed in each decision about whether or not to comply with security measures. Such decisions are affected by the design of the security system, the organizational culture, and the extent of monitoring and the possibility of sanctions for noncompliance.

Sasse closed by listing the following as key research challenges:

- Identifying and understanding trade-offs,
- Developing ways to quantify and compare costs for different usability and security criteria and for different stakeholders,
- Identifying and reconciling individual and collective goals with respect to security, and
- Developing a better understanding of the short- and long-term impact of security measures on individuals, businesses, and society.

SECURITY IN VIRTUAL WORLDS (FRANK GREITZER)

Social media such as blogs, microblogs (e.g., Twitter), social networking sites (e.g., Facebook), and virtual worlds provide new tools for individuals to communicate, play, and work. Because these virtual communities are being used for many of the same things that people do in real life, they are becoming plagued by many problems and crimes of

the real world—including theft of identities and virtual assets. Identity and access management is a particular challenge in virtual environments because it is difficult to establish that an online identity is in fact the real-life person that it claims to be. Moreover, online tools do not necessarily provide protection that is strong enough to protect confidential discussions (and it may be appropriate today to shift such activities to a private environment).

This suggests the need for a better understanding of the security issues that threaten trust and privacy in these environments and for a better understanding of the role played by usability. Frank Greitzer noted several conventional cybersecurity challenges that may play out in different ways in virtual environments. These include what sorts of authentication and credentials are most appropriate in virtual worlds, who should be responsible for managing credentials and verification, and how authentication and identification can best be manifested in a virtual environment.

One of the most important research questions concerns the human factors and usability implications of proposed solutions. How can someone trust that the person (avatar) with whom he or she is interacting is accountable? For any particular solution, how can the solution be made usable and trustworthy for individuals who participate in virtual worlds? Finally, Greitzer underscored that validation—how to evaluate the effectiveness of proposed solutions—is essential.

FEEDING PRACTICE BACK INTO RESEARCH (MARY ELLEN ZURKO)

Mary Ellen Zurko discussed how to integrate lessons learned from practice into research thinking, noting that not only should research results inform practice, but practice and real-world experience with development, deployment, and use also should inform research. Issues that can only be understood in this context include scaling; performance; usability, accessibility, and user experience; and the total cost of ownership and return on investment.

For example, the security weaknesses of text passwords have been revealed by understanding their use and changes in their use. In the early days, passwords were used primarily by a handful of professionals to access a single computer. Today, people make use of passwords for a wide array of services, each of which has different strength requirements and management policies. The result is that almost all forms of deployed security using passwords are weak in terms of both usability and the security that results. Researchers are exploring alternatives to passwords for authentication, but these have many barriers to deployment, such as those

associated with the scale of enrollment and the need to retrofit complex infrastructures that only support passwords.

Another connection between practice and research is the real-world constraints that affect the deployment of research results. For example, a researcher might come up with a better way of presenting a user with information about how much trust to place in the claimed sender of an e-mail message. In the real world, the space available for presenting this information may be significantly constrained by an e-mail client's user interface. Products routinely have a number of features competing for space in the user interface, with designers making decisions based on factors such as primary use cases, sales criteria, organizational politics, esthetics, technical difficulty, and maintenance. Such trade-offs, commonplace in practice, need to inform research so that researchers can successfully transfer their results into practice and products. Such technology transfer depends on the development of tools and best practices that allow practitioners to incorporate research results on user-centered security into the systems that they design, build, and operate. It also depends on the development of criteria and approaches for evaluating how usably secure a system or approach is likely to be. The transfer into practice can be facilitated through standards groups such as the Web Security Context Working Group of the World Wide Web Consortium. Intellectual property concerns can also be a barrier to uptake.

Zurko proposed a number of ideas that would encourage a greater emphasis on technology transfer concerns within the context of the research environment. Most obviously, funding specifically targeted at usable security research addressing uptake issues would drive progress in that area. Venues for publishing the results of such research are critical, as one of the main activities of researchers is to publish. Framing devices such as use cases, frameworks, and challenges can inspire and structure potential research and its results.

Zurko suggested several opportunities for research to be informed by experiences with deployed systems, including the following:

- Conducting user studies of deployed technology, including contextual analysis;
- Measuring changes in user behavior in response to changes in services;
- Using open-source and free-product betas as a source of information on user behavior; and
- Studying the characteristics of deployed security, through such techniques as tiger-teaming.

The presentation closed with the observation that although there is no substitute for the ground truth of real-world experiments, there are also constraints on what can be done in these settings. One should not be able to make changes that deliberately impair materially the security of an operational system. As a result, experiments with security in real-world settings require controls and oversight, much as efficacy and safety considerations govern the conduct of drug trials.

CYBERSECURITY INSIDER THREAT (DEANNA CAPUTO)

Deanna Caputo began her presentation by discussing the problem posed by trusted insiders. Espionage, intellectual property theft, and sabotage involving computer networks are among the most pressing cybersecurity challenges that threaten government and the private sector. Surveys reveal that current or former employees and contractors are the second-greatest cybersecurity threat, exceeded only by hackers. The insider threat is manifested when human behavior departs from compliance with established policies, regardless of whether it results from malice (malicious insiders) or a disregard for security policies.

Because insiders can make use of the privileges that they have been granted, they do not need to engage in behaviors that break explicit rules, making it difficult to detect these actions. What are the possible signatures of lawful but suspicious activities? How can these detection mechanisms be made usable by security analysts? How can the interests in detecting suspicious behavior be balanced with the privacy interests of employees?

Caputo went on to describe work being done at the MITRE Corporation to address these questions. This work includes the development, testing, and piloting of a prototype detection system known as Exploit Latent Information to Counter Insider Threats (ELICIT). It uses sensors to collect information used to detect and prioritize potential threats. It is based on a characterization of how trusted insiders use information, and it uses information about both the user and the information context to differentiate malicious and legitimate activities. Caputo commented that the resulting information allows time-consuming and costly threat validation and forensic investigation to be concentrated on a small number of prioritized cases.

Work on ELICIT prompted a team of social scientists and engineers to explore experimentally how malicious insiders use information differently from how a benign baseline group uses information.[4] Caputo

[4] Deanna D. Caputo, Marcus Maloof, and Gregory Stephens, "Detecting Insider Theft of Trade Secrets," *IEEE Security and Privacy* 7(6): 14-21, November/December 2009.

discussed preliminary results from the double-blind study of malicious insiders, which revealed some counterintuitive results. One surprise was that malicious insiders tend to grab and go—favoring quantity over quality—contrary to expectations that insiders would be "low and slow," working meticulously to avoid raising suspicions. Caputo also offered some essential aspects gleaned from these efforts of approaches for detecting insider threats. The work has also involved the development of test data to represent both malicious and benign users.

The work has also informed practical guidance developed by MITRE for handing these threats.[5] The following measures can be used by organizations to defend against the insider threat:

- *Make employees the first line of defense.* Educate them about spotting suspicious behavior. Understand that satisfied workers are less likely to be disgruntled insiders.
- *Pay attention to employee behavior.* Look for signs of vulnerability, unexplained wealth, and so on.
- *Prioritize assets.* Concentrate monitoring resources where it matters most.
- *Know what baseline behaviors on the network look like* so that anomalies can be recognized. Enumerate trust relationships with other organizations because their insiders can become your insiders.
- *Divide responsibilities.* Separate duties for key functions to reduce exposure.
- *Grant least privileges,* and audit for privilege overentitlement.
- *Prepare for recovery* through continuity of operations and data backup plans.

Caputo also described work on the insider threat by several other research groups. Shari Lawrence Pfleeger and Joel Predd at RAND have developed a framework for understanding the insider threat and a taxonomy for describing insider actions, and they are developing a framework for response to the insider threat. Frank Greitzer at the Pacific Northwest National Laboratory is looking at behavioral data to support predictive modeling and analysis in order to improve situational awareness for the security analyst, facilitate response coordination, and help the analyst focus on the highest-risk activities. A prototype system is under development that provides enhanced visual analytics and a multilayered user

[5] Mark Maybury, *How to Protect Digital Assets from Malicious Insiders*, The MITRE Corporation and Institute for Information Infrastructure Protection. Available at http://www.thei3p.org/research/mitremi.html; accessed February 25, 2010.

interface encompassing displays for high-level status as well as detailed monitoring.

In terms of areas for further research, Caputo posed the following questions:

- What trade-offs associated with insider threat monitoring are there between the individual's right to privacy and the organization's need to protect its assets?
- What are the implications of pre-interventional activities such as monitoring and the collection of data and predictive modeling? How might they affect morale or violate employee trust or legal guidelines? What is the potential for false accusations or misuse?
- What is the impact of user profiling, and what are the ethical and legal issues surrounding this approach?

Finally, Caputo noted that research on the insider threat would be aided by good operational data samples.

4

Some Potential Research Directions for Furthering the Usability, Security, and Privacy of Computer Systems

A principal goal for the workshop was to identify research questions and areas within the emerging field of usability, security, and privacy that would assist in increasing the security of computer systems used by individuals and organizations. Limiting the discussion to research questions was, perhaps not surprisingly, a challenge. Participants approached the problem from a multitude of perspectives, reflecting the many disciplines represented at the workshop and the involvement of academic, industry, and government researchers as practitioners. And because many participants were very engaged with the usability-security-privacy challenge, there was a natural temptation to explore possible solutions as well as fruitful research areas. The following sections summarize research directions that emerged from the questions posed to workshop participants and from breakout sessions, reports back from the breakout sessions, and plenary presentations and discussion.

DIMENSIONS OF USABILITY, SECURITY, AND PRIVACY

Definitions

Breakout session participants spent a considerable amount of time grappling with how to define usable security, working under the belief that one cannot improve something that cannot be measured and that one cannot measure something without a good definition for what one seeks to measure. Indeed, definitions were discussed in every breakout

session in some form, leading the committee to identify the need for better agreement on terminology and definitions as one of the four overarching research challenges at the intersection of usability, security, and privacy (see Chapter 5).

Usability for Whom?

Although usability is often equated with the experience of end users of IT systems—and this was indeed the focus of many presentations and discussions at the workshop—usability concerns for other groups were also discussed. Notably, administrators of IT systems also contend with systems that are difficult to understand and configure. The security or privacy consequences of a misconfiguration or other error by a system administrator can, of course, be much more serious and wider in scope than the consequences of an error of a single user. However, the line between administrator and end user is somewhat blurry because every home user is in effect the administrator of his or her own home network and the computers and other devices attached to it, which suggests that both system administrators and home users stand to benefit from improvements aimed at either group.

Usability also matters for system developers. More usable tools would make it easier for them to avoid or detect design and coding errors that affect security and privacy. Moreover, there is an opportunity to improve the usability and security of systems by introducing better usable security and privacy features to development environments and libraries.

To what extent do demographic and cultural differences affect usability, security, and privacy? One particular question that came up repeatedly during the workshop was whether it was true that younger generations are more security-savvy and less privacy-sensitive. A related question, assuming that younger users are less privacy-sensitive, was whether they would retain that perspective as they grew older.

Finally, participants cautioned that academic studies of usability are not necessarily representative of the user population. They typically employ small groups of college students, which reflects poor experimental design for two reasons: the group sizes are too small, and they are not drawn from a group that is representative of the broader population. Companies can also make the same mistake with respect to usability studies used to test new services.

Is Usability for Security and Privacy Special?

How might usability for security and privacy be distinct from the broader topic of usability of information technology? One difference of

possible significance is that security inherently involves an actor other than the user—the active adversary who will try to take advantage of usability flaws and may also attempt to mislead the user through "social engineering." Another is that security involves focusing the user's attention not only on the task at hand but also on the future consequences and aftereffects of the task. Yet another is that security is generally not the end user's primary concern. Further investigation of the similarities and differences might yield insights as to what lessons can be transferred directly from other usability work and where the issues are in fact different.

METRICS, EVALUATION CRITERIA, AND STANDARDS

Metrics—that is, measures of how usable or secure a system is—are important to assessing progress (e.g., how much better is this system than another one?) and making rational decisions about investment (e.g., is this system "good enough" or is further investment in improvements warranted?). Workshop participants observed that security has long resisted precise measurement—let alone in combination with usability. That is, there are few good ways to determine the effectiveness or utility of any given security measure, and the development of metrics remains an open area of research.[1] With respect to usability, participants noted a multitude of potentially relevant measures for usability (which might be measured in terms of user errors, time required to configure or modify a system, time to master a system, or user satisfaction ratings) and system effectiveness or utility. Further research would help identify which of these measures, or what others, are most useful.

Related to metrics is the question of what criteria should be used in evaluating and accepting the usability and security of an IT system and how one might go about certifying a system as aligning security, privacy, and usability. How might such criteria be instantiated as future guidelines? Are there exemplar software applications that could be identified as benchmarks for security and usability and therefore serve as a source for creating a set of criteria for usable, yet secure, systems? Several discussions considered how such criteria might vary accordingly to application, context, or perspective. For example, how might one divide applications into categories in which similar weights would be given to security and usability? Despite the likely differences among the categories, might it be possible to develop a common checklist that contains a core set of usability and security criteria that would cover 80 percent of all applications?

[1] For a detailed discussion of the challenges associated with cybersecurity metrics and possible research directions, see NRC, *Toward a Safer and More Secure Cyberspace,* 2007, pp. 133-142.

Workshop participants also grappled with the question of perspective. How might criteria for a usable and secure system differ for people in different roles, including system administrators, security professionals, system owners, end users, security designers, and developers?

Another question raised was whether compliance with usability and security standards might become a condition of connecting to enterprise or public networks. Finally, with respect to the development of standards, it was observed that such efforts would be challenging today given the limited understanding of what constitutes a system that is usable and secure and that appropriately protects personal information. What would be required to develop useful standards? What organizations and institutions are best positioned to develop them?

UNDERSTANDING USERS

Central to the topic of usability is a better understanding of users. An approach known as user-centered design addresses the needs, desires, and limitations of users. The related field known as human-centered computing concerns information technology artifacts and their relationship to people. Both approaches are informed by and depend on observation of human behavior. Workshop presentations and discussions approached this topic from several perspectives: user mental models, risk perception and communication, and user incentives. (Incentives, another important topic with respect to understanding users and their motivations, are considered separately below, because they also apply to other actors.)

User Mental Models

"Mental models" describe people's thought processes and understanding. (A related term used by some speakers was "user metaphors.") Workshop participants suggested that work to understand and enhance models of security and privacy would be valuable.

A first research topic and logical starting point is to gain a better understanding of the mental models that people apply to security and privacy today. What are the best ways to elicit these current mental models? What do they tell us that could be used to make improvements in today's systems and in the design of future systems? What specifically do system designers and developers need to know about user mental models to design systems and applications that are usable yet secure?

A second research topic is the development of better models that could be adopted in system design. For example, are there models for security or privacy that have the concreteness and usefulness of the now-familiar desktop and folder scheme? This nearly ubiquitous metaphor

has been enormously successful in making computing accessible to a broad population. What abstractions might make security and privacy more usable?

A third topic is how to deploy better models—that is, how best to introduce new models to users and incorporate these new models in future system design. (This issue also relates to the topic of user education, discussed below.) One specific suggestion was that it might be useful to develop "user stories" describing appropriate use of IT that highlights the importance of security and privacy. Such user stories could be created after the development of a better understanding of how users make use of security indicators and interfaces. Taking an epidemiological perspective, it would be useful to understand how many individual users' mental models would have to be changed to make a noticeable impact in improving computer security "for the masses."

A fourth topic is to study how well users understand their own user model. Can they assess their technical proficiency well enough to understand whether or not they are capable of making informed security decisions? One suggestion for how to assess this understanding is to compare the results of self-reporting with testing, to determine the proficiency of different user types to make informed security decisions.

Risk Perception and Communication

Do people understand how secure (or insecure) their computers are? Do they understand the concept of risk—that is, the probabilities and consequences—and the risks associated with particular actions? Do people understand the implications for themselves and others of a lack of attention to security? Do they understand the risks associated with system failure, disclosure of confidential information, or the release of their private information? Do people care less about damage to others if they themselves do not pay or incur damages for security breaches caused to others? What role does the information source play in getting users to change their behavior? What impact do disclosures about the use of personal information have on the use of security functionality? How might the literature on risk communication developed largely in other domains be applied and extended to enhancing security and privacy? How can what is known about how people understand and react to risk be used to induce them to do things that are good for them and for society?

A closely related set of issues involve what languages and processes can best be used to communicate with users, including those within particular organizations as well as the general public. How can best practices be transferred to those who compose training materials, documentation, and user messages?

Learning About and From Mistakes

A number of comments at the workshop related to the importance of understanding users' mistakes that play a role in security incidents—mistakes that are often a direct result of usability problems. A better understanding of these mistakes could be fed back into better designs and better user education. One suggestion was to develop a taxonomy of human security errors and mistakes, which would help with identifying general classes of problems and thus a set of general solutions that would influence behavior. A first step would be to conduct a literature review and meta-analysis of past studies.

Participants noted that it is not easy to gather information on user mistakes. How does one get users to figure out that they have made mistakes? How can users be convinced to report mistakes (and how are the associated privacy issues to be dealt with)? How does one create an environment in which users are motivated to report errors (so that design and user education can be improved), yet maintain a culture of user accountability? What can be learned from the records that organizations (e.g., enterprises or Internet service providers) keep about security incidents?

It was also noted that individual users may have quite different definitions of what constitutes an error; there may be many security incidents in which the end user would not view something as an error even though others might. What are useful definitions for developers, managers, and users to adopt?

User Education

Users who better understand how to use systems and appreciate the security and privacy implications of their actions are better positioned to protect security and privacy. Better education can help overcome usability challenges; however, workshop participants cautioned that an emphasis on education not be used as an excuse for not improving usability. One suggested area for research is to achieve a better understanding of the knowledge that users currently have and how they attained that knowledge.

User education was also suggested as a way of influencing values associated with security and privacy. How can one influence norms for acceptable and/or appropriate behavior with respect to security and privacy? How is a "culture of security" to be created among different user groups? What can be learned from such fields as social psychology or social marketing?

Participants also suggested examining the limits of user education as a way of improving security and privacy. For example, to what extent is it valid to assert that "if they understood why they are being inconve-

nienced, users would follow the directions"? The discussion of incentives, below, suggests that there are significant limits. Another limiting factor may be that security is generally not the end user's primary concern.

A final set of questions relates to curriculum and institutionalizing education. What are core concepts that one should teach? How could user education best be incorporated into specific settings such as kindergarten through grade 12 education or employee training programs? How might user education be introduced into informal learning settings such as libraries? How might other informal learning techniques be used—techniques such as videos that play while software is loading or online games that teach about security and privacy? Under what circumstances should user education be mandated, and by whom?

INCENTIVES FOR BETTER SECURITY AND PRIVACY

Many workshop participants observed that incentives are an important force in shaping the behavior related to security and privacy. Incentives can be applied to different actors. (For example, should the onus for security be placed on a home Internet user or on that user's ISP or on both?) One might even consider how incentives apply to adversaries. (For example, if the cost of mass-scale attacks is increased, will adversaries instead conduct targeted attacks?)

Incentives can take both positive and negative forms. For example, employees can be given positive incentives through the use of awards for maintaining good security, or they can be given negative incentives through reprimands or poorer evaluations for security failures. In the marketplace, positive incentives might include favorable reviews of products with better security, whereas negative incentives would include liability for inadequate security or negative reports in the press.

Importantly, incentives for usability, security, and privacy are not necessarily aligned. To take a simple example, an employee who faces pressure to accomplish a task to meet a deadline may choose to sidestep security measures that slow his or her work. However, if a system administrator fears being sanctioned for a possible security breach, he or she may impose on user activity onerous restrictions that reduce usability.

Externalities play an important role in considering incentives. Individuals can easily take steps that have little consequence for themselves but negatively affect many others. For example, household computer users do not face the cost of damage that poorly secured computers may have across the Internet when those household users fail to take simple steps to prevent their computers from being infected. Nor does an employee incur the total cost of allowing a virus to infect a corporate network. The result is that individual users will tend to pay less attention

to security than is desirable from an organizational or societal perspective. How can the right incentives be created so that users choose a level of security that better protects everyone else? What fraction of such failures can be attributed to inadequate incentives, a lack of information, or the poor usability of today's security tools?

More generally, participants noted a misalignment of individual, corporate, and societal incentives. Modern computer systems, especially in a network setting such as the Internet, exhibit very significant differences between the effects of an insecure computing environment on an individual and the effects on society. In particular, often each individual faces small negative consequences from a lack of security for his or her computer system, but when such a lack of security is widespread, the consequences are exponentially large negative effects, even catastrophic ones. The divergence between private and public incentives with respect to exerting effort to secure computing systems leads with mathematical certainty to a less secure IT environment as computing systems become more interconnected and more complex, making better alignment of private and public incentives on security an important challenge for policy makers.

With respect to incentives for businesses, participants asked where the money is in usable security. How might business models be adjusted to make usable security profitable? How might regulatory models be adjusted to make unusable security less profitable? They also pointed to a particular problem of users who continue to use systems even though their subscriptions to security updates have expired. Are there viable business models in which security subscriptions never expire?

Behavioral aspects surfaced repeatedly in the workshop discussions, notably in the observation that sometimes individuals seem not to act in a fully rational way in protecting security. Such seemingly irrational behavior can have multiple explanations—actors not being well informed, actors considering a wider range of outcomes than have been anticipated by the system designer, or such ideas as "bounded rationality" that have been developed in behavioral economics.

Finally, participants observed that it is hard to develop appropriate incentives when little is known about costs or impacts. For example, relatively little is known about the cost of identity theft or cybersecurity breaches. This is due in part to the inherent difficulty in obtaining access to the relevant data. Neither private firms nor the government is incentivized to share such data (see Chapter 5). The ironic result is that it may be necessary to address the issue of incentives to share data in order to acquire a better understanding of how to increase incentives to enhance security and privacy.

APPROACHES TO CONSTRUCTING SYSTEMS WITH "USABLE SECURITY"

Automation

One specific approach to improving the usability of systems is to reduce the burden on the end user through automation. People may be more satisfied with systems when they have more control; but in the context of security, it may be that the more control allowed the user, the greater the opportunities for introducing vulnerabilities or security breaches. To what extent and when should usable security aim to automate security decision making and remove the human from the loop entirely, versus providing a more usable interface for the human to interact with? Despite the appeal of taking the human out of the loop, participants cautioned that there are limits, because automation cannot handle unexpected, novel events—and the one thing that is known about such events is that they are certain to occur at some point.

Several specific ideas were proposed. One was to use machine learning from context to come up with an acceptable security policy for a user without the user's directly having to adjust security or privacy parameters. Another idea was to have a user establish policy by specifying desired outcomes and having the system express those outcomes as a set of security rules. The system would then verify that the rules derived from those outcomes are consistent and complete, and only ask the user for additional instructions in the event that they are not. Research could help shed light on the feasibility of such approaches.

Authentication Beyond Passwords

Many participants noted the well-known shortcomings of passwords with respect to security and usability. Simply, the effort spent entering passwords and recovering or resetting them when they are forgotten was noted to be a significant waste of time. Passwords that are easy to remember are also easy to guess, but passwords that are hard to remember are more easily forgotten or subject to compromise if they are written. Systems often require users to change passwords periodically, which may also lead to users' writing them down or using guessable mnemonic schemes for generating their passwords. Systems typically require their own passwords, often with conflicting rules about acceptable user names and passwords, meaning that users must keep track of a wide array of credentials.

Alternatives that address these shortcomings have been developed. They are used for certain applications but have not enjoyed widespread support and use. These alternatives include hardware token authentica-

tion, which provides stronger authentication than do passwords, and (primarily within enterprises) federation-based authentication schemes that free users from keeping track of multiple passwords. Several barriers to these alternatives were mentioned, including a lack of awareness about alternatives, the cost of implementing a new approach, and the lack of off-the-shelf "drop-in" replacement technology. Another barrier is the potential impact on privacy arising from the potential for the use of alternatives to link activities across multiple systems. Several techniques have been proposed to reduce the likelihood of such linkage, but they may nonetheless be susceptible to determined attack.[2]

Participants offered a number of open questions that research could address:

- What obstacles have been encountered to the deployment of alternatives to passwords? What can be learned from data and research collected by industry groups such as the OpenID Foundation?
- What have been the barriers to the adoption of federation-based authentication schemes? Would standardizing the rigor of systems used for authenticating help?
- Suppose that authentication schemes were to be considered as they related to the needs of sets of users: How would one even begin to classify what the different sets of users are?
- There are populations of users that have already been issued strong authenticators (e.g., the federal government's Personal Identity Verification card and its predecessor, the Common Access Card). What has prevented their use outside the workplace?
- Suppose that users had a single authenticator that could be used universally. Would they prefer to have that supplied by the government or by private industry? How aware and concerned are people about the potential for the linkage of activities across multiple systems? What approaches are best suited for preventing linkage across multiple systems, and what would it take for them to be widely deployed?

Processes and Tools

Participants suggested a number of development and management processes and tools that would help advance usable security and privacy—as well as associated research challenges:

[2] One commercial example of a technology for preventing the linking of visits across multiple parties that rely on a common identifier is "U-Prove," offered by Credentica, a firm recently purchased by Microsoft. It relies on a zero-knowledge scheme developed by Stefan Brands and colleagues.

- *Creating better developer support tools*. Guidelines, principles, and design patterns can all help support developers in building systems that provide usable security and privacy. Research questions include how well usable security can be built into such elements as integrated development environments or libraries and how one would evaluate the effectiveness of support tools.
- *Dealing with dynamic threats that develop between design iterations*. Security threats involve adversaries who seek to exploit weaknesses—often more rapidly than the typical design-cycle time. How are threats to be dealt with that arise between typical design iterations? Can the design process be sped up?
- *Making recovery more usable*. Recovery from security breaches, where the extent of the damage done may be difficult to determine, is a major challenge. How can recovery processes be made more secure and usable?
- *Simplifying user decisions.* Complexity impedes usability. How can one make the best use of such approaches as establishing useful bundles of security settings or secure default settings in order to reduce the burden on users?
- *Redesigning infrastructure*. Are there ways that key infrastructure such as the Internet or operating systems (which can be difficult to change in major ways given their enormous installed base) might be redesigned to provide more usable security and privacy? How might barriers to making such changes be overcome?

Usability Through the "Stack"

Computer systems are often thought of in terms of layers—for example, the commonly used Open System Interconnection model for communications networks consists of the physical, data link, network, transport, session, and presentation layers. Similarly, software runs on top of operating systems that provide abstractions for accessing computing, storage, and display resources. Such layering hides details below each layer from the layers above. Much of the work in usable security has focused on advances at the presentation layer—in user interfaces. But it was suggested at the workshop that one should consider whether changes to this conventional model might enhance usable security and privacy. Participants suggested several questions regarding how these conventional abstractions might be reconsidered in order to enhance usable security and privacy.

How "far down the stack"—that is, how far down into the design of the underlying system—is it necessary to go to provide usable security?

Can one enhance usable security by tweaking the abstractions that are used today? What are possible improvements that might result from rethinking the abstractions? How might lower layers be redesigned to support metaphors that would improve usable security? What ambitious new usable security goals could be achieved by redesigning the stack?

What information is needed from lower levels to interact with the user about security errors? How can the application developers at upper levels be helped to understand and use the security information from lower levels?

What if the abstractions were to be changed, say from hosts in the network to user data? How would one express protocols in those terms? Would this help with users' control of their information? Does moving the security abstractions to the data make them safer? How can a life-cycle view of user data be incorporated—that is, who can it be sent to, who can store it, how is it protected, and how is it controlled?

Other Opportunities for Improving Systems

Presentations and discussions advanced a number of specific opportunities for improving the usability, security, and privacy of IT systems:

- *Distinguishing green and red machines.* Butler Lampson's talk (Chapter 2) suggests enhancing security and privacy by using separate "green" and "red" machines for conducting activities that are safe and not safe. The green machine, used for important things, would demand accountability, whereas the red one would not. This approach immediately raises a usability question: How does a user readily identify green and red machines and understand their distinct purposes? More generally, what are the potential advantages of more-specialized machines, and what are the usability challenges associated with using multiple machines?
- *"Scarlet letter" option.* Is it helpful to inform users that they are interacting with a system or service that is following unsafe practices? What can be learned about the effectiveness of such capabilities that have been included in browsers and search engine result pages? How does one deal with the risk of spoofing? How does one address the privacy issues introduced because the service identifying unsafe activities knows what systems or services the user is interacting with?
- *Building systems that assume worst-case scenarios.* No matter how usable computer systems are, no matter how well users are trained and motivated, and no matter what precautions are taken, errors will occur and systems will be compromised. How should systems be built to cope with these inevitable problems?

- *Whitelisting versus blacklisting.* Whitelists (lists of approved entities) and blacklists (lists of prohibited entities) can both be used to authorize access or to grant privileges. In which cases does each approach provide better security and usability?

5

Overarching Challenges to Advancing Research in Usability, Security, and Privacy

Four overarching challenges facing researchers working in the field of usability, security, and privacy were apparent in the presentations and discussions at the workshop. Although these challenges apply to many emerging research areas, they are particularly relevant to research on usability, security, and privacy.

INCONSISTENT TERMINOLOGY AND DEFINITIONS

Participants in the breakout sessions devoted considerable time and attention to terminology and definitions. "Usable security" was the term frequently used to capture the notion of security measures developed with attention to usability considerations. Another commonly used term was "HCI-SEC" (human-computer interaction–security). Whatever the specific term used to describe the intersection of usability, security, and privacy, each participant tended to define the area in relation to his or her own background. Interestingly, usability practitioners tended to stress security issues, and security practitioners tended to stress usability issues.

Adding "privacy" to the mix complicated matters still further, as definitions of privacy were frequently based on personal philosophies and experience, perhaps reflecting the deeply personal way in which many individuals approach privacy issues. Moreover, some workshop participants noted that although some activities, such as the annual Symposium on Usable Privacy and Security mentioned above, explicitly call out both terms, neither "usable security" nor "HCI-SEC" explicitly invokes issues

related to privacy, despite the technical and policy links between the two concerns. Some may immediately associate privacy issues with the term "security," but this is not universally true. Agreeing to a common definition or term that was inclusive of the concept of privacy proved challenging throughout the workshop.

LIMITED ACCESS TO DATA

Several workshop participants cited the need for more and better empirical data and commented on the difficulties that they faced in gaining access to such data. For example, data on industry or government computer system security breaches are generally unavailable—corporations are hesitant to disclose this information owing to the potential threat to reputation, stock price, and ongoing business; and information about breaches to government computer systems is frequently treated as sensitive or classified. Even data on matters less touchy than security breaches cannot be readily obtained. Participants noted, for example, the difficulty in obtaining data on the productivity impacts of security measures. Even when researchers are able to obtain data, nondisclosure agreements may restrict their ability to publish their results. If researchers do gain the ability to work with corporate data, an additional challenge is that of conducting research in a way that enables repeatability.

SCARCENESS OF EXPERTISE AND UNFAMILIARITY WITH EACH OTHER'S WORK AT THE INTERSECTION OF USABILITY, SECURITY, AND PRIVACY

Many of the workshop participants commented that working in the area of usability, security, and privacy is especially challenging because of the need for researchers who are familiar with both computer security and human-computer interaction. These were, at least until recently, considered distinct disciplines—most security researchers have traditionally ignored usability issues, and vice versa (and likewise for usability and privacy).

One consequence is unfamiliarity with each other's work. Throughout the workshop, there were frequent instances in which either a computer security or a usability expert would identify a research question outside his or her area of expertise, only to receive immediate feedback from relevant experts that this particular question had already been addressed. "I did not know that that research existed" was a common lament heard at the workshop. Although this immediate feedback was useful to the workshop participants, it also suggests there may be a significant lack of knowledge about usability-related work among security researchers and

about security-related work among usability researchers (with a similar situation existing with respect to usability and privacy). Another consequence pointed out by workshop participants is that valuable resources may be spent re-researching questions that are already well understood.

Still another consequence is that although a few interdisciplinary research collaborations have emerged, there remain few individuals in either area with sufficient expertise to identify their counterparts on the other side—and fewer still with expertise in both areas. Research funding at the intersection would foster the development of such expertise by training graduate students and attracting young faculty.

Appendixes

A

Workshop Agenda

USABILITY, SECURITY, PRIVACY OF COMPUTER SYSTEMS: A WORKSHOP

July 21–22, 2009
National Academy of Sciences, 2100 C St., N.W., Washington, DC

July 21, 2009

9:00 a.m. **Welcome**
Nicholas Economides
- Introduction of Committee Members and Provocateurs
- Purpose and Goals of Workshop
- Review Workshop Agenda
- Logistical Items

9:30 **Framing the Usability, Security, and Privacy Research Challenge**
Butler Lampson

10:00 **Perspectives on Current and Prospective Research**

Security in Virtual Worlds
Frank L. Greitzer

Usable Privacy
Lorrie Faith Cranor

Feeding Practice Back into Research
Mary Ellen Zurko

Cybersecurity and Insider Threat
Deanna D. Caputo

Creating a Hierarchy of Categories of User Interactions
Angela Sasse

Framework of Economic Issues on Usable Security
Nicholas Economides

12:15 p.m. Working Lunch

1:30 **Breakout Sessions I**

How Do We Measure Usable Security?
Frank L. Greitzer and Charles P. Pfleeger, session leads

Approaches to Usable Security
Lorrie Faith Cranor and Don Norman, session leads

Developing a "Usable Security" Standard
Butler Lampson, session lead

Economic Issues for Usable Security and Policy Changes
Nicholas Economides and Susan Landau, session leads

Beyond Phishing 1: Improving Systems
James Foley and Simson Garfinkel, session leads

3:00 Break

3:30 **Breakout Sessions II**

Approaches to Usable Security
Lorrie Faith Cranor and Don Norman, session leads

Developing a "Usable Security" Standard
Butler Lampson, session lead

Beyond Phishing 2: Alternatives to Passwords
Simson Garfinkel and Susan Landau, session leads

Human Factors and Security Incidents
Deanna D. Caputo and Charles Pfleeger, session leads

Usable Security Through the Stack, Its Life Cycle, and All Its Users
Angela Sasse and Mary Ellen Zurko, session leads

Report Back from Session Leads

July 22, 2009

9:00 a.m.	**Welcoming Remarks** *Nicholas Economides*
9:30	**Moving from Usability to Understandability** *Don Norman, Co-Founder, Nielsen Norman Group*
10:00	**Breakout Sessions: Identifying Short- and Long-term Research Projects Related to Usability, Security, and Privacy of Computer Systems**
11:30	Lunch
1:00 p.m.	**Session Leads Report Back**
2:00	**Closing Remarks**

B

Workshop Participants

Alessandro Acquisti, Carnegie Mellon University
Gail-Joon Ahn, Arizona State University
Lujo Bauer, Carnegie Mellon University
Amy Baylor, National Science Foundation
Richard Beckwith, Intel Corporation
Richard Beigel, National Science Foundation
Genevieve Bell, Intel Corporation
Steven Bellovin, Columbia University
Konstantin Beznosov, University of British Columbia
Sameer Bhalotra, Senate Select Committee on Intelligence
Duane Blackburn, Office of Science and Technology Policy, Executive Office of the President, The White House
Bob Blakley, The Burton Group
Matt Blaze, University of Pennsylvania
Robert Bohn, National Coordination Office for Networking and Information Research and Development
Roy Boivin II, IT Masterminds
Tanya Brewer, National Institute of Standards and Technology
Desiree Campbell, Cigital Federal
Deanna Caputo, Mitre Corporation
Bill Cheswick, AT&T
Yee Yin Choong, National Institute of Standards and Technology
Douglas E. Comer, Purdue University
Greg Conti, Rumint

Alissa Cooper, Center for Democracy and Technology
Earl Crane, Department of Homeland Security
Lorrie Faith Cranor, Carnegie Mellon University
Renee Crews, Department of Justice
Anita D'Amico, Applied Visions
Patrick Dempster, CSC
Rachna Dhamija, Harvard University
Roger Dingledine, TOR Project
Donna Dodson, National Institute of Standards and Technology
Cathy Dunaway, Department of Education
Stephen Duncan, General Services Administration
Nicholas Economides, New York University
Keith Edwards, Georgia Institute of Technology
Sherri Eillis, Department of Transportation
Carl Ellison, Microsoft Corporation
Jeremy Epstein, SRI International
Nicholas Feamster, Georgia Institute of Technology
James Fisher, Noblis
Heather Foley, Peace Corps
James Foley, Georgia Institute of Technology
Myisha Frazier-McElveen, CitiGroup
Jeffrey Friedberg, Microsoft Corporation
Simson Garfinkel, Naval Postgraduate School
Carrie Gates, Computer Associates International, Inc.
Nathaniel Good, Palo Alto Research Center
Chris Greer, National Coordination Office for Networking and Information Research and Development
Frank Greitzer, Pacific Northwest National Laboratory
Wendy Grossman, Independent Consultant
Lawrence Hale, General Services Administration
Gillian Hayes, University of California, Irvine
Marty Herman, National Institute of Standards and Technology
Haym Hirsh, National Science Foundation
Jason Hong, Carnegie Mellon University
Darren Kall, Kall Consulting
Jason Kerben, Department of State
Joseph Kielman, Department of Homeland Security
Larry Koved, IBM
Mike Lake, IBM
Butler Lampson, Microsoft Corporation
Susan Landau, Privacylink.org
Carl Landwehr, Office of the Director of National Intelligence
Ji Sun Lee, Department of Homeland Security

Richard Lempert, Department of Homeland Security
Bill Lewis, United States Navy
Susan Lightman, Executive Office of the President
Patrick Lincoln, SRI International
Roy Maxion, Carnegie Mellon University
Ernest McDuffie, National Coordination Office for Networking and Information Research and Development
Gary McGraw, Cigital Federal
Ross Micheals, National Institute of Standards and Technology
Richard Morris, National Institutes of Health
Pat Muoio, Office of the Director of National Intelligence/Science and Technology
Elaine Newton, National Institute of Standards and Technology
Brand Niemann, Environmental Protection Agency
Donald Norman, Nielsen Norman Group
Lucy Nowell, Department of Energy
Jennifer O'Connor, Department of Homeland Security
Andrew Patrick, Carleton University
Hetel Petel, Peace Corps
Charles Pfleeger, Pfleeger Consulting
Shari Lawrence Pfleeger, Rand Corporation
Gary Phillips, Symantec
Walt Polansky, Department of Energy
Jules Polonestky, Future of Privacy Forum
Rob Reeder, Microsoft Corporation
Holly Rensvold, Department of Homeland Security
Tom Rhodes, National Institute of Standards and Technology
Marc Rogers, Purdue University
Charles Romine, National Institute of Standards and Technology
Marc Rotenberg, Electronic Privacy Information Center
Norman Sadeh, Carnegie Mellon University
Angela Sasse, University College London
Stuart Schechter, Microsoft Corporation
Diane Smetters, Palo Alto Research Center
Darren Smith, National Oceanic and Atmospheric Administration
Jim Sorace, Department of Health and Human Services
Sylvia Spengler, National Science Foundation
Brian Stanton, National Institute of Standards and Technology
Tim Stanton, Naval Postgraduate School
Joe Steele, Adobe Systems
Brock Stevenson, Department of Justice
Michael Sulak, Department of State
Denise Tayloe, Privo

Mary Theofanos, National Institute of Standards and Technology
Joseph Trella, Truestone
V.N. Venkatakrishnan, University of Illinois at Chicago
Daniel Weitzner, Massachusetts Institute of Technology
Geoff Willcher, Bellevue College
Jeannette Wing, National Science Foundation
Irene Wu, Federal Communications Commission
Lenore Zuck, National Science Foundation
Mary Ellen Zurko, IBM

C

Biosketches of Steering Committee Members and Staff

Nicholas Economides, *Chair,* is a professor of economics at the Stern School of Business at New York University. He is an internationally recognized academic authority on network economics, electronic commerce, and public policy. His fields of specialization and research include the economics of networks, especially of telecommunications, computers, and information; the economics of technical compatibility and standardization; industrial organization; the structure and organization of financial markets and payment systems; antitrust; application of public policy to network industries; strategic analysis of markets; and law and economics. Professor Economides has published more than 100 articles in top academic journals in the areas of networks, telecommunications, oligopoly, antitrust, and product positioning, and on the liquidity and the organization of financial markets and exchanges. He is editor of *Information Economics and Policy, Netnomics, Quarterly Journal of Electronic Commerce, Journal of Financial Transformation,* and *Journal of Network Industries;* he is on the advisory board of the Social Science Research Network, editor of *Economics of Networks Abstracts* by SSRN, and former editor of the *International Journal of Industrial Organization*. His Web site on the Economics of Networks has been ranked as one of the top four economics sites worldwide by *The Economist* magazine. Professor Economides is the executive director of the NET Institute, http://www.NETinst.org, a worldwide focal point for research on the economics of network and high-technology industries. He is an adviser to the U.S. Federal Trade Commission; the governments of Greece, Ireland, New Zealand, and Portugal; the Attorney

General of New York State; major telecommunications corporations; a number of the Federal Reserve Banks; the Bank of Greece; and major Financial Exchanges. He serves on the advisory board of *The Economist* Intelligence Unit. He has commented extensively in broadcast and in print on high-technology, antitrust, and public policy issues. Previously, he taught at Columbia University (1981-1988) and at Stanford University (1988-1990). He holds a PhD and MA in economics from the University of California at Berkeley, as well as a BSc (First Class Honors) in mathematical economics from the London School of Economics.

Lorrie Faith Cranor is an associate professor of computer science and of engineering and public policy at Carnegie Mellon University, where she is the director of the CyLab Usable Privacy and Security Laboratory (CUPS). She is also chief scientist of Wombat Security Technologies, Inc. She has authored more than 80 research papers on online privacy, phishing and semantic attacks, spam, electronic voting, anonymous publishing, usable access control, and other topics. She has played a key role in building the usable privacy and security research community, having co-edited the seminal book *Security and Usability* (O'Reilly, 2005) and founded the Symposium on Usable Privacy and Security (SOUPS). She also chaired the Platform for Privacy Preferences Project (P3P) Specification Working Group at the W3C and authored the book *Web Privacy with P3P* (O'Reilly, 2002). She has served on a number of boards, including the Electronic Frontier Foundation board of directors, and on the editorial boards of several journals. In 2003, she was named one of the top 100 innovators 35 or younger by *Technology Review* magazine. She was previously a researcher at AT&T-Labs Research and taught in the Stern School of Business at New York University. Dr. Cranor received her doctorate degree in engineering and policy from Washington University in St. Louis in 1996.

James D. Foley is a professor in the College of Computing, and a professor in the School of Electrical and Computer Engineering at the Georgia Institute of Technology (Georgia Tech). A leading international figure in two major disciplines of computer science (graphics and human-computer interaction), Dr. Foley has received lifetime achievement awards in both fields from the Association for Computer Machinery's special interest groups (SIGGRAPH in 1997 and SIGCHI in 2007). Dr. Foley was one of the computer graphics pioneers who went on to help establish HCI as a discipline. The co-author of three books, he is the first author of what many consider the definitive text in computer graphics, *Fundamentals of Interactive Computer Graphics,* which has sold 400,000 copies in 10 translations. Dr. Foley arrived at the College of Computing in 1991 and founded the GVU Center. Four years later, *U.S. News and World Report* ranked the

center No. 1 for graduate computer science work in graphics and user interaction. Active in industry, Dr. Foley became the director of MERL (Mitsubishi Electric Research Laboratory) in 1996 and then CEO and chair of Mitsubishi Electric Information Technology Center America in 1998. He returned to Georgia in late 1999 to head up the state's Yamacraw economic development initiative in the design of broadband systems, devices, and chips. For 4 years (2001-2005), Dr. Foley chaired the Computing Research Association (CRA), which represents more than 200 research universities, corporate research laboratories, and professional societies. In February 2008, he was elected to the National Academy of Engineering. A few months later, he received the 2008 Class of 1934 Distinguished Professor Award, the highest honor that Georgia Tech bestows on faculty. Of all his awards, Dr. Foley says that he most treasures the one given him by computing graduate students who named him "Most Likely to Make Students Want to Grow Up to Be Professors."

Simson L. Garfinkel is an associate professor at the Naval Postgraduate School in Monterey, California, and an associate of the School of Engineering and Applied Sciences at Harvard University. His research interests include computer forensics, the emerging field of usability and security, personal information management, privacy, information policy, and terrorism. Dr. Garfinkel is the author or co-author of 14 books on computing. He is perhaps best known for his book *Database Nation: The Death of Privacy in the 21st Century.* His most successful book, *Practical UNIX and Internet Security* (co-authored with Gene Spafford), has sold more than 250,000 copies and has been translated into more than a dozen languages since the first edition was published in 1991. Dr. Garfinkel received three bachelor of science degrees from the Massachusetts Institute of Technology (MIT) in 1987, a master of science in journalism from Columbia University in 1988, and a PhD in computer science from MIT in 2005.

Butler W. Lampson is a technical fellow at Microsoft Corporation and an adjunct professor of computer science and electrical engineering at MIT. He was on the faculty at the University of California, Berkeley, and then at the Computer Science Laboratory at Xerox PARC and at Digital Systems Research Center. He has worked on computer architecture, local area networks, raster printers, page description languages, operating systems, remote procedure call, programming languages and their semantics, programming in the large, fault-tolerant computing, transaction processing, computer security, WYSIWYG editors, and tablet computers. He was one of the designers of the SDS 940 time-sharing system, the Alto personal distributed computing system, the Xerox 9700 laser printer, two-phase com-

mit protocols, the Autonet Local Area Network, the SDSI/SPKI system for network security, the Microsoft Tablet personal computer (PC) software, the Microsoft Palladium high-assurance stack, and several programming languages. He holds a number of patents on networks, security, raster printing, and transaction processing. At Microsoft he has worked on anti-piracy, security, fault-tolerance, and user interfaces. He was one of the designers of Palladium and spent 2 years as an architect in the Tablet PC group. Currently he is in Microsoft Research, working on security, privacy, and fault-tolerance, and kibitzing in systems, networking, and other areas. He is a member of the National Academy of Sciences and the National Academy of Engineering and a fellow of the Association for Computing Machinery and the American Academy of Arts and Sciences. He also served on the Computer Science and Telecommunications Board of the National Research Council. He received an AB from Harvard University, a PhD in EECS from the University of California at Berkeley, and honorary ScD's from the Eidgenössische Technische Hochschule, Zurich, and the University of Bologna.

Susan Landau is a fellow at the Radcliffe Institute for Advanced Study during the academic year 2010-2011. She recently completed a book on security risks of building surveillance into communications infrastructures (to be published by MIT Press in the spring of 2011). From 1999 to 2010 Dr. Landau was a Distinguished Engineer at Sun Microsystems; there she concentrated on the interplay between security and public policy. She has briefed government officials both in Washington, D.C., and in Europe on such disparate issues as security risks in surveillance mechanisms, digital rights management, and cryptographic export control; she has written numerous articles and op-ed pieces on these issues. Most recently she testified for the House Science Committee on Cybersecurity Activities at the National Institute of Standards and Technology's (NIST's) Information Technology Laboratory. She and Whitfield Diffie wrote *Privacy on the Line: The Politics of Wiretapping and Encryption.* Dr. Landau is a member of the Commission on Cyber Security for the 44th Presidency, established by the Center for Strategic and International Studies, and serves on the Computer Science and Telecommunications Board of the National Research Council and on the advisory committee for the National Science Foundation's Directorate for Computer and Information Science and Engineering. Before joining Sun, Dr. Landau was a faculty member at the University of Massachusetts and at Wesleyan University. She is the recipient of the 2008 Women of Vision Social Impact Award, a fellow of the American Association for the Advancement of Science, and an ACM Distinguished Engineer.

Donald A. Norman is the Breed Professor of Design at Northwestern University where he co-directs MMM, the dual-degree MBA and engineering program offered jointly by Northwestern's schools of Management and Engineering that focuses on managing products and services from design to execution. He is also co-director of the Segal Design Institute. He is Distinguished Visiting Professor at KAIST, the Korea Advanced Institute of Science and Technology, in the Department of Industrial Design. He is co-founder of the Nielsen Norman Group and has been vice president of Apple Computer and an executive at Hewlett Packard. He serves on many advisory boards, such as the editorial advisory board of *Encyclopedia Britannica* and KAIST. He has received honorary degrees from the University of Padova (Italy) and the Technical University of Delft (the Netherlands), the "Lifetime Achievement Award" from SIGCHI, the professional organization for Computer-Human Interaction, and the Benjamin Franklin Medal in Computer and Cognitive Science from the Franklin Institute (Philadelphia). He is well known for his books *The Design of Everyday Things* and *Emotional Design*. His most recent book, *The Design of Future Things,* discusses the role that automation plays in such everyday places as the home and the automobile. He is currently working on a new book called *Sociable Design* that combines the lessons of his previous works, extending them to cover social networks and social interaction. He earned a PhD in psychology from the University of Pennsylvania.

Charles P. Pfleeger is an independent consultant for Pfleeger Consulting Group specializing in computer and information system security. Among his responsibilities are threat and vulnerability analysis, system design review, certification preparation, training, expert witness testimony, and general security advice. His customers include government and commercial clients throughout the world. Dr. Pfleeger was previously a master security architect on the staff of the chief security officer of Cable and Wireless, and Exodus Communications, and before that he was a senior computer scientist and director of research for Arca Systems, director of European Operations for Trusted Information Systems, Inc. (TIS), and a professor in the Computer Science Department of the University of Tennessee. Dr. Pfleeger was chair of the IEEE Computer Society Technical Committee on Security and Privacy from 1997 to 1999 and has been a member of the executive council of that committee since 1995. He is on the board of reviewers for Computers and Security, is a book review editor for *IEEE Security and Privacy,* and is on the board of advisers for OWASP, the Open Web Application Security Project. Dr. Pfleeger has lectured throughout the world and published numerous papers and books. His book *Security in Computing* (of which the fourth edition—co-authored with Dr. Shari Lawrence Pfleeger—was published in October 2006) is the

standard college textbook in computer security. He is the author of other books and articles on technical computer security and computer science topics. He holds a PhD degree in computer science from Pennsylvania State University and a BA with honors in mathematics from Ohio Wesleyan University. He is a Certified Information Systems Security Professional (CISSP).

CSTB STAFF

Jon Eisenberg is director of the Computer Science and Telecommunications Board of the National Research Council. He has also been study director for a diverse body of work, including a series of studies exploring Internet and broadband policy and networking and communications technologies. In 1995-1997 he was a AAAS (American Association for the Advancement of Science) Science, Engineering, and Diplomacy Fellow at the U.S. Agency for International Development, where he worked on technology transfer and information and telecommunications policy issues. Dr. Eisenberg received his PhD in physics from the University of Washington in 1996 and a BS in physics with honors from the University of Massachusetts at Amherst in 1988.

Shenae Bradley is a senior program assistant at the Computer Science and Telecommunications Board of the National Research Council. She currently provides support for the Committee on Sustaining Growth in Computing Performance, the Committee on Wireless Technology Prospects and Policy Options, and the Computational Thinking for Everyone: A Workshop Series Planning Committee, to name a few. Prior to this, she served as an administrative assistant for the Ironworker Management Progressive Action Cooperative Trust and managed a number of apartment rental communities for Edgewood Management Corporation in the Maryland/DC/Delaware metropolitan areas. Ms. Bradley is in the process of earning her BS in family studies from the University of Maryland at College Park.