Workshop on Scalable Cyber-Security Challenges in Large-Scale Networks: Deployment Obstacles

Large-Scale Networking (LSN) Coordinating Group
of the
Interagency Working Group (IWG) for Information Technology Research and Development (IT R&D)
http://www.hpcc.gov/iwg/lsn.html

March 13-14, 2003
Landsdowne, Virginia

Table of Contents

1. Executive Summary

Internet security is universally acknowledged to be an extremely important problem. The scope of the problem is quite broad, encompassing all aspects of information security in networked environments, not just the security of a few widely deployed Internet protocols. Furthermore, the importance of the problem can only grow: Computers and networks are an increasingly central part of business, government, recreation, and almost all other aspects of daily life.

Security and cryptology have been the subjects of intensive study for at least three decades, and the technical community has developed some very elegant and widely admired solutions and methods. Indeed, some of the most successful computer science researchers in the world can be found in security and cryptology, and they have garnered at least four Turing Awards[1], one Gödel Prize[2], one Knuth Prize[3], and one Kanellakis Prize[4]. Furthermore, many of the elegant technical solutions that appear in the research literature have been implemented and tested, and some have even been commercially developed. One would think that this fortunate confluence of an important problem and a set of impressive solutions would lead to rapid adoption of security technology.

Reality has repeatedly confounded this expectation. The security field is littered with the carcasses of clever but largely unused solutions. Some people have made the (undoubtedly oversimplified and overstated) claim that security research has been an academic success and a real-world failure. Clearly, something is amiss.

In March of 2003, the Large-Scale-Networking Coordinating Group convened a workshop entitled "New Directions in Scalable Cyber-Security in Large-Scale Networks: Deployment Obstacles." The workshop was attended by 44 security professionals from a wide variety of academic, industrial, and governmental organizations. This report summarizes the workshop findings about security research needs. Information about the workshop location, program, and participants can be found at http://www.cs.yale.edu/homes/jf/LSN-worshop.html.

Workshop participants met in plenary sessions and break-out groups over the course of two days and had wide-ranging, stimulating discussions. Five broad, overlapping research themes were identified as important priorities:

1. Is there solid evidence that today's large-scale networks actually *are* too insecure to serve their intended purposes?

---

[1] Butler Lampson (1992); Manuel Blum (1995); Andrew Yao (2000); Ronald Rivest, Adi Shamir, and Leonard Adleman (2002).
[2] Shafi Goldwasser, Silvio Micali, and Charles Rackoff for "The Knowledge Complexity of Interactive Proof Systems" and Lázló Babai and Shlomo Moran for "Arthur-Merlin Games" (1993).
[3] Andrew Yao (1996)
[4] Leonard Adleman, Whitfield Diffie, Martin Hellman, Ralph Merkle, Ronald Rivest, and Adi Shamir for the invention of public-key cryptography (1996).

2. What are the appropriate *definitions*, *metrics*, and *models* of security for today's large-scale networks and for those currently being envisioned, designed, and developed?
3. To what extent are persistent security problems really the result of deployment obstacles? That is, would apparently good security technology that was *not* adopted actually have solved the problems of real users if it *had* been adopted?
4. How can mathematical methods and theoretical research contribute to our understanding of real-world security (as opposed to contributing solely to the creation of ever deeper and more complex theories of security)?
5. How can experimental research be conducted in the realm of large-scale-network security, which combines two realms in which experimental research is notoriously difficult: large-scale networking (in which one has to experiment on heterogeneous, ever-changing equipment that is owned and operated by a diverse set of parties) and security (in which adversaries are hard to model and experimental results are hard to interpret)?

The next five sections of this report flesh out the workshop findings in each of these areas. Several cross-cutting, high-level issues are relevant to all five sections and arose repeatedly in workshop discussions:

A. There is an urgent need for more and better data about actual network security. None of the pressing research questions covered in the following sections can truly be answered without ongoing access to data, and yet no effective data-collection infrastructure is now in place.
B. Economic aspects of network security and the lack thereof are extremely important. Many workshop participants echoed the recent observations of Ross Anderson and others that standard ideas from Econ textbooks, *e.g.*, the principal agent and moral-hazard problems, network externalities, liability dumping, and tragedies of the commons, go a long way toward demystifying lack of adoption of apparently good security technology. More generally, monetary aspects of security are clearly both very important and very poorly understood. Potential users of security-enhancing technology are entitled to ask "how much will it cost me to use this, and what benefits can I expect if I do?" Yet, with the tools now at their disposal, security experts cannot give meaningful answers to those questions.
C. The research questions raised at this workshop are inherently multi-disciplinary (incorporating unsolved problems in, *e.g.*, economics, business, statistics, ethics, psychology, and political science, as well as computer science and engineering) and multi-modal (requiring both theoretical and experimental research). Research teams will often have to span institutions and to include members from multiple academic departments as well as industry and government.
D. Educational reform is urgently needed in the security field. It has become common wisdom that computer-science educators must raise students' awareness of security and privacy issues and must teach standard cryptology and computer security techniques; however, this common wisdom is inadequate. IT professionals of all sorts will have to think critically and creatively about security, and today's curricula are inadequate. Funding-agency programs that address the

research needs outlined below should support educational innovation and curriculum development as well.

Good overviews of the subject of network security and the evolving research needs in the field can be found in [NRC1, NRC2, NRC3, NRC4].

The research activities suggested by this report will, if successful, not only yield solutions to specific technical problems that are widely recognized as important but will also yield general techniques, both theoretical and experimental, for analyzing security problems and comparing candidate solutions. Research suggested herein promises to broaden and deepen the nature of security expertise in a manner that (1) includes full consideration of adoption barriers and other social and economic issues and (2) integrates multiple intellectual disciplines and multiple research modes. As such, it promises to lead to greater actual network security, as opposed simply to more available security technology.

2. Towards a Realistic Assessment of the Problem

The subtitle of this workshop was "Deployment Obstacles," and these obstacles are certainly real – that is, there are a large number of security problems in deployed systems and a large number of "solutions" to these problems that have been rolled out but not adopted. However, the mere existence of problems that are fixable but have not been fixed does not necessarily imply that the Internet as a whole is so "insecure" as to pose a critical threat to the smooth functioning of the societies and organizations that rely on it. In fact, there are flaws and vulnerabilities in many prominent infrastructural systems in our society, and yet these systems are used extensively and productively, and the inconvenience and monetary losses attributable to known flaws are viewed as things to be worked on or just endured, not as crises demanding the attention of well funded research communities. This section makes some qualitative observations about the (so far) spotty adoption of security technology and identifies some research needs that must be met if we are to develop a realistic understanding of the Internet-security challenge.

2.1. Willingness to Pay for Universal Access

It is tempting to attribute the lack of adoption of known security technologies to ignorance or foolhardiness, but, in fact, there may be sound economic reasons to prefer a relatively insecure system to a competing secure one. The absence of accurate measures of security vulnerabilities argues for a humble attitude towards the adoption decisions made by individuals and businesses in the field. Although these decisions might not take into account all the costs of reduced security, we should start with the assumption that the people making them are responding rationally to their own assessment of the costs and benefits of adopting particular technologies.

Qualitatively speaking, one of the largest costs of many security technologies is that they stop you from dealing with people you don't know. While it is possible to build large systems that depend on pre-existing relationships between participants, the value of the most successful Internet technologies, including email and the World Wide Web,

ultimately depends on their easy, universal accessibility. Users are willing to pay a high price in security for such universal access (*e.g.*, by receiving spam email or by accepting innocuous-looking web links that might lead to bad places or, more likely, to nowhere at all), and, although there is clear demand for software that can reduce this price locally (*e.g.*, spam filters and web filters), there have been few serious proposals to solve the problem globally by replacing either the SMTP-based email system or the World Wide Web with closed systems usable only by those who have been certified as polite. This may be an inherent property of sufficiently large systems: As the number of participants grows, the strategy of excluding all potential bad participants becomes less and less cost-effective.

Given the choice between security and inclusion, many Internet-based businesses have opted for inclusion. For example, the HTTP protocol is notoriously vulnerable to snooping, and many Internet users are justifiably afraid to send sensitive information like credit-card numbers across unencrypted HTTP connections. Because of this fear, essentially all online retailers now use encrypted HTTPS connections for processing orders. However, many large Internet retailers (Amazon.com is perhaps the most successful example) are willing to process orders across unencrypted HTTP connections opened by users whose obsolete or misconfigured browsers prevent them from using HTTPS. Presumably, the cost of losing a sale is greater than the much more hypothetical cost of exposing such a user's credit-card number. (The cost of fraud is not borne directly by the retailer, but card-issuers who lose enough money have ways to pressure retailers to require encrypted connections.)

In contrast, security technologies that do not alter the relationships among participants in the system seem to be adopted much more easily. The relationship between the Telnet and SSH protocols is closely analogous to the relationship between HTTP and HTTPS; in both cases, the second is a drop-in secure replacement for the first. One difference is that Telnet was almost entirely used to allow remote users to connect to machines on which they already had accounts -- a pre-existing relationship; thus, system administrators take much less risk in replacing Telnet by SSH of inadvertently excluding ``new customers'' who would respond by taking their business elsewhere. Combined with this lower cost was a higher risk of not using security; a server that accepts unencrypted Telnet connections risks the theft of passwords that can then be used to compromise the machine. So, in this case, we might reasonably guess that the difference between the continuing survival of unencrypted HTTP and the near-total disappearance of Telnet is explained by the difference between the costs and benefits of replacing them.

The preceding discussion is an attempt to account for the separate and often undocumented decisions of millions of system administrators, and so it is necessarily somewhat speculative. But it illustrates the point that the decision to adopt a particular security technology appears to be highly sensitive to cost. In measuring the effectiveness of particular security technologies, therefore, we should start with a baseline assumption that users have already taken into account the costs and benefits that they personally derive from them and that the more sophisticated users are already using all technologies that are personally cost-effective. What remains is the question of predicting security

externalities - the costs to other users of a single user's decision to adopt a particular technology or not. Here we need good models, definitions, and metrics that can distinguish between failures that disrupt the activities of individual users and patterns of failures that threaten to bring down the entire system.

Economic analysis of security-technology adoption is currently an active area of study; see, for example, [AND1, VAR1] for further discussion.

2.2. Problem 1: "Insecurity" may be Inherent

The workshop began with a keynote talk by Andrew Odlyzko, director of the Digital Technology Center at the University of Minnesota and one of the few people (possibly the only person) in the country with substantial credibility in all of security, networking, and economics. The take-home message of his talk was clearly that the crisis mentality now dominating discussions of Internet (in)security and related research-funding priorities is both unwarranted and futile. The following abstract of [ODL1] summarizes the message in his words:

> Security is not an isolated good, but just one component of a complicated economy. That imposes limitations on how effective it can be. The interactions of human society and human nature suggest that security will continue being applied as an afterthought. We will have to put up with the equivalent of bailing wire and chewing gum, and to live on the edge of intolerable frustration. However, that is not likely to be a fatal impediment to the development and deployment of information technology. It will be most productive to think of security not as a way to provide ironclad protection, but the equivalent of speed bumps, decreasing the velocity and impact of electronic attacks to a level where other protection mechanisms can operate.

In the same paper, Odlyzko goes on to flesh out a point that he touched on in his talk at the workshop:

> The point is that we should be realistic about what can be accomplished. A productive comparison might be with auto safety. There has been substantial improvement in the past, and it is continuing. Greater crashworthiness of cars as well as better engineering of roads and more effective enforcement of drunk driving laws and more use of seat belts have made car travel far safer. In the United States, deaths per mile traveled by car fell at a compound annual rate of 4.7 percent between 1980 and 2000, by a cumulative factor of more than 2. However, because of growth in volume of travel (by 80 percent), the total number of deaths has only decreased from 50,000 per year to 42,000 per year. Moreover, in the last few years, the annual number of fatalities appears to have stabilized. Our society has decided (implicitly, without anyone ever voting on this explicitly) that we are willing to

> tolerate those 42 thousand deaths per year. Measures such as a drastic reduction in the speed limit, or devices that would constantly test the driver for sobriety or alertness, are not acceptable. Thus we manage to live with the limitation of the large masses of human drivers.

Equally importantly, he stresses in the same paper that the incompatibility of human beings and formal security technologies that he believes to be inherent may not be an obstacle to widespread, productive use of the Internet:

> The standard thinking in information security has been that absolute security is required. Yet we do have a rapidly growing collection of data that shows the value of even imperfect security. The experience of the pay-TV industry is certainly instructive. Although their systems have been cracked regularly, a combination of legal, technological, and business methods has kept the industry growing and profitable. Some more examples are offered by the applications of encryption technologies to provide lock-in for products, as in the replacement printer cartridge market [SCC1]. Very often, "speed bumps" is all that is needed to realize economic value.

Earlier in the paper, he substantiates this general point with some specific examples: "The really massive financial disasters of the last few years, such as those at Enron, Long Term Capital Management, or WorldCom, owed nothing to inadequacy of information security systems."

Indeed, there is good reason to conjecture that the Internet cannot be made "secure" in any of the rigorous senses considered in the security-research literature and that there is no compelling reason that it should be. The Internet is a vast, complex, evolving, heterogeneous, and universally accessible system. Examples of infrastructural systems that exhibit *some* of these characteristics include the highway system, the credit-card system, and the postal system. None of these systems satisfies any formal definition of "security" that are analogous to those found in the information-security literature; yet all serve society's needs to a large extent, and there has been no serious attempt to replace any of these with a "secure" alternative.

Similarly, the fact that well publicized attacks and breakdowns (*e.g.*, the "Slammer-worm" incident of early 2003 and the successful DDoS attacks on Amazon, eBay, and other ecommerce sites a few years ago) cause temporary disruption of popular Internet services does not mean that the network as a whole requires major redesign. Similar things happen in other contexts and do not provoke major redesign of infrastructural systems. For example, a flood on West 34th Street in Manhattan could shut down Macy's for a day or two; that would not cause people to redesign the retail-shopping system in the US. Periodically, snowstorms shut down the northern East Coast of the US, from Maine to Virginia, and people cannot get to work, to stores, to schools, or even to hospitals. Financial losses are considerable during these "outages," and loss of life is not

unknown or unexpected. Yet our society recovers quickly from all of them, and no major societal reorganizations are undertaken.

Perhaps a certain amount of major disruption is inevitable in the complex, modern world, and the Internet is just one more complex, modern world that we live in.

2.3. Challenging Issues

The preceding discussion provides an informal counter-argument to the oft-repeated (but never proven) claims that we are in the midst of an Internet-security crisis and that a massive increase in security-related spending and effort is warranted. However, it does not get us past the "devil's advocate" stage of the discussion and on to a realistic assessment of the Internet-security situation. Getting to that stage will require a more formal approach.

Research is needed to determine how serious a problem Internet insecurity really is in practice. For example, the following questions should be formulated precisely and answered as accurately as possible.

What fraction of current Internet activity is disrupted by security problems? Are most of these disruptions actually show stoppers, or are they less serious (*e.g.*, delays)? How much *potential* Internet activity is avoided because of security concerns? Are these numbers bigger or smaller than they were five years ago? Note that there are many ways to quantify losses and disruptions (*e.g.*, amounts of network traffic, amounts of employee time, overall dollar value), and whether or not meaningful answers can be obtained is a wide open question. However, it is clear both that there is anecdotal evidence that people are increasingly concerned about Internet security and that we do not know whether this concern is due to real increases in security problems or simply to increased media coverage of security problems.

Similarly, informal analogies between the Internet and other large, heterogeneous, universally accessible systems motivate us to take seriously Odlyzko's conjecture that we should not expect to have widely used systems that satisfy the type of stringent definitions of security found in the research literature. However, these informal analogies do not provide the basis for a realistic assessment of the Internet security situation; if comparisons with other large-scale infrastrucural systems are to lead to definitive conclusions, more systematic study is needed.

The Internet achieves scalable, unbrokered resource sharing through a modular, bottom-up design in which the critical network layer makes use of untrusted intermediaries and provides a datagram service that gives only the weakest guarantees. By contrast, security and authentication mechanisms often require heavyweight operations (*e.g.*, computationally intensive encryption algorithms) or make strong, shared assumptions that tend to break as the system is scaled up along various important dimensions. The question is whether or not scalable networked computing systems can be built that use

untrusted intermediaries and end-to-end protocols but nonetheless provide adequate trustworthiness for a significant spectrum of applications.

Research is needed to determine whether there is an inherent trade-off between the kind of scalability and universal access that makes the Internet so useful and the strong authentication mechanisms that make (typically smaller) systems "secure." Systematic and rigorous study of other large-scale infrastructural systems that have been in mass-market use longer than the Internet may be very instructive. In particular, the following questions should be formulated precisely and answered as accurately as possible.

In what sense, if any, are the large-scale infrastructural systems (*e.g.*, the financial or transportation systems) that we rely on "secure"? What security mechanisms are in wide-spread use in these systems, and how effective are they? How are security breaches dealt with in these systems? In what ways is the Internet-security problem analogous to security problems in other large-scale infrastructural systems, and to what extent can security mechanisms that are effective in other settings be carried over?

Note that the research needs identified in this section support the claims in Section 1 that there is an urgent need for data about actual Internet insecurity and that security research is an inherently multidisciplinary activity. In particular, collaboration with social-science researchers may help computer scientists better understand fundamental tradeoffs inherent in large-scale, universally accessible systems in which human behavior plays a key role.

3. Problem 2: Metrics, Models, and Definitions of Cybersecurity

A persistent theme of this workshop was that better metrics, models, and definitions would greatly aid the development of *deployable*, effective security technology. Recognition of the importance of the right abstractions is not new; computer-security professionals have traditionally respected this imperative. The following three major principles of cybersecurity are identified in [NRC1] and have been formalized and studied from many angles in the research literature:

- *Confidentiality*: controlling who gets to read information
- *Integrity*: assuring that information and programs are changed only in a specified and authorized manner; and
- *Availability*: assuring that authorized users have continued access to information and resources.

It was the consensus of the workshop participants that these principles, while important, are an inadequate foundation for the metrics, models, and definitions that are needed in today's large-scale networks. We explore this theme in this section.

3.1. Conventional Wisdom is Inadequate

Numerous reports by knowledgeable, respected groups of people have for years exhorted organizations to adopt "cybersecurity best practices" and force their employees to use them. These exhortations are exemplified by the following four recommendations from [NRC2]:

- Establish and provide adequate resources to an internal entity with responsibility for providing direct defensive operational support to system administrators throughout the organization. To serve as the focal point for operational change, such an entity must have the authority – as well as a person in charge – to force corrective action.
- Promptly fix problems and vulnerabilities that are known or that are discovered to exist.
- Mandate the organization-wide use of currently available network/configuration management tools, and demand better tools from vendors.
- Mandate the use of strong authentication mechanisms to protect sensitive or critical information and systems.

It is time to question this orthodoxy. If cybersecurity best practices have not yet been widely adopted, perhaps security professionals need to rethink our assumptions and reformulate our recommendations, rather than continuing to try to shove the old recommendations down everyone's throats. In particular, perhaps network managers should not be exhorted to fix *all* known problems and vulnerabilities promptly, because some are not worth fixing. Standard law-enforcement organizations do not strive to eliminate all criminality and unruliness in the real world, and simple, compelling economic models explain why it would be cost ineffective for them to do so (see, *e.g.*, [BEC1]). Instead of repeating exhortations to fix all problems and vulnerabilities promptly, security researchers could be more effective by providing tools for accurate assessment of the risks to information systems and of the costs and benefits of various risk-mitigation strategies.

Much of established cryptography and security theory focuses on providing worst-case guarantees, typically in a stylized model of the network and from the viewpoint of a single component of the network. Some beautiful results have been produced, *e.g.*, those on zero-knowledge proof systems and secure, multiparty function evaluation, and more will be said about some of the existing theory in Section 5 below. In this section, we focus on heretofore neglected modeling questions.

We believe that the security of computer networks is best viewed at the system level rather than at the component level and that it is better to provide quantitative measures of security with respect to a realistic model of user behavior rather than absolute guarantees of security with respect to a stylized model of behavior. Useful security metrics should

allow a comparison between the cost to deploy security measures and the benefit from system-wide security improvements.

Economics provides a natural framework within which to define metrics for systemic security. Consider network *A* and network *A+*, which is identical to *A* except that additional security technologies have been adopted (possibly only in a subnetwork of *A*). From an economic point of view, it is natural to measure the value of this additional security in terms of the expected increase in the *total utility* of users of the network. At one extreme, it might be that *A* was already "secure enough" and the additional effort made by the adopters of the security technology produced no economic value. At another extreme, *A+* may remain so insecure that the total utility of the system is negligible, and again the additional security has no economic value. Thus, a useful systemic metric of security should capture the aggregate opportunity cost to users of the system. Similar economic tools have been used by the networking-research community in its comparison of best-effort-only service and reservation-capable service (see, *e.g.*, [BrSh]), and that analysis may carry over to our setting.

If we had good quantitative metrics, decisions about the deployment of new security technology could be addressed not only in terms of absolute security for individual components but also in terms of a systemic cost-benefit analysis and return on investment. For further discussion, see Anderson's argument that, because there are many possible attacks, it can be better to develop methods to contain an attack and construct a patch than it is to anticipate and fix all attacks before they occur [AND1].

Threat models will be an essential part of good security metrics for today's networked information systems. This point was made many times throughout the workshop and is summarized in [NRC3] as follows: "Metrics for security could include number of attacks of different types, fraction of attacks detected, fraction of attacks repelled, damage incurred, and time needed to detect and respond to attacks. Note that making measurements on such parameters depends on understanding the attacks that do occur -- because many attacks are not detected today, continual penetration testing is required to establish such a baseline." We agree strongly that accurate system-level measurement and modeling that enables cost-benefit analysis will require a deeper understanding of the attacks that actually occur; the worst-case guarantees of invulnerability to all technologically feasible attacks that dominate existing security theory often require designers to work at too abstract a level and to ignore system features of practical importance.

There was also solid agreement among workshop participants that new metrics, models, and definitions should move from the research community's very heavy emphasis on prevention of security breeches and toward more emphasis on detection. Recently, the term "autonomic security" has been used to convey the importance of detection. The basic idea, as explained in [IBM1], is to "take the weakness of the Internet and turn it into an advantage, *i.e.*, use the large number of machines to provide enhanced security" in a three-pronged approach: compromise-tolerant distribution (partitioning applications and

data across many nodes so that compromise of a subset of the nodes does not destroy functionality), compromise detection, and compromise recovery.

3.2. Why Security is Hard to Measure and Model

Good mathematical models of secure networks have to incorporate something that is notoriously hard to model: human behavior. Collecting data about attackers is obviously difficult. However, it is not even straightforward to collect and analyze information about people and procedures within one's own organization and how they interact with security technology. Installing new security technology may have no effect (or even a negative effect) if the technology is ignored or used improperly or if necessary changes in procedures are not made at the same time. Furthermore, people may adapt their behavior if they think that security measures are preventing them from working efficiently. Adaptation of people to technology and *vice versa* have to be part of the overall measurement and modeling program.

The simplest goal one could set in this area would be accurate measurement of the occurrence of known security vulnerabilities in today's deployed systems and the extent of the damage that they cause. Even this goal is far from attainable today, because organizations are justifiably reluctant to share such data and do not necessarily have the expertise and the tools needed to gather them. Of course, new vulnerabilities are discovered every day, and threats are evolving. The basic, wide open question is how one can set up a measurement infrastructure that distinguishes "secure" network behavior from "insecure" network behavior. Some network states are obviously abnormal (*e.g.*, a denial-of-service attack at its peak), but some are ambiguous (*e.g.*, higher-than-normal but not crippling amounts of traffic from heretofore unseen sources), and some are deceptive (*e.g.*, the receipt of an email message with a heretofore unseen virus attached). How can one ever be sure that one is observing a "secure" network in operation? In fact, it is unlikely that any network is ever *completely* secure, and thus effective network-monitoring tools will have to make decisions based on the expected consequences of observed behavior, not on whether or not the behavior satisfies an abstract definition of security. Thus, data about the effects of security violations are important as well as data about the occurrence and frequency of such violations. Mathematical models of secure network behavior should evolve as network behavior evolves and as the quantity and quality of available data improve.

The technical challenges involved may seem overwhelming (both the engineering challenge of setting up an adequate measurement infrastructure and the mathematical challenge of modeling secure networks in a way that has predictive value), but similar challenges have been met in other scientific communities. Meteorological data are hard to collect, but collection technology and methods have improved steadily over the years, and weather prediction has benefited enormously both from improved access to data and from sophisticated mathematical modeling. There is an added complication in security measurement and modeling, namely the presence of adversarial behavior, but meteorology provides proof that complex, evolving, large-scale systems are amenable to

mathematical analysis and that the network-security community need not necessarily restrict itself to the (probably oversimplified) models now in the literature.

3.3. Challenging Issues

Research is needed in order to develop definitions and models of network security that are more realistic than the worst-case, component-oriented, extremely abstract definitions and models available today. Furthermore, research is needed into techniques for measuring the security state of a network so that new definitions and models can be validated or revised. A *systemic* approach is needed, in which the focus is on the state of the network as a whole, rather than on the state of individual hosts and links.

It is widely conjectured that organizations would put security procedures in place if they were held liable for security breeches that adversely affected their customers and business partners – or that they would at least buy insurance so that their liabilities for information-security disasters would not bankrupt them. However, assignment of liability is, at this point, putting the cart before the horse. Security professionals are not now in a position to assess risks accurately, to standardize risk-mitigation strategies, or to calculate appropriate insurance premiums. Research is needed in all aspects of *risk assessment* and *risk management* in networked information systems. The focus of this research should include how to predict, limit, contain, correct, and compensate people for the damage done by security failures, rather than solely on how to prevent such failures altogether. Success may require a *layered approach* to secure operations of networks and systems: Preventive measures will eliminate certain threats, detection and response will suffice for some others that cannot be prevented cost effectively, and some network activity will inevitably be disrupted by unforeseen or exceptionally powerful attacks. During these disruptions, back-up systems have to be in place so that mission-critical tasks can be accomplished without use of the compromised network. Successful research on modeling and risk assessment would allow organizations to provision these back-up systems appropriately.

Current theories often define security in terms of asymptotic bounds on computational resources such as time, space, and bandwidth or in terms of conformance with a formal models expressed as logical formulae or automata. Research is needed in order to develop definitions that are more quantitative and concrete and to develop measurement and monitoring techniques that determine how secure a particular network is at a given time. It is particularly important to be able to determine whether the security of a network is improved, impaired, or unaffected when a technological or procedural change is made.

Note that, as in Section 2 above, the research needs identified in this section support the claim that there is an urgent need for data about the current state of network security.

4. Problem 3: Towards a Deeper Understanding of Adoption Barriers

The preceding two sections argue that research is needed to determine the true extent of the Internet-insecurity problem and to develop better definitions, metrics, and models of security. The next two sections will consider specific research needs in theoretical and experimental computer science, respectively. In this section, we focus on deployment obstacles. Deployablity and adoptability have *not* often been explicitly identified as topics worthy of study; past workshops, conferences, and journal issues have focused almost exclusively on the design, analysis, and implementation of more security technology. This workshop's explicit identification of questions of deployability and adoptability bodes well for the future of security research and, with luck, will be looked back on as the beginning of a new, intellectually broad, and increasingly realistic period in the history of the field.

The nature of the adoption problem is far from clear and is probably not monolithic. It is possible that existing security technology, if it were consistently and correctly used, would solve the security problems that users of today's networked computers complain about. Where this is the case, the research community must strive to understand why these technologies aren't used. It is also possible that existing security technology, although intellectually satisfying, "solves the wrong problems." In these cases in which "our abstractions don't model our reality" [FFSS], the research community must strive for better abstractions.

4.1. Digital Signatures: A Cautionary Tale

Abstractly, digital signatures are enormously appealing, and it is quite easy to convince oneself that they are necessary for electronic commerce and for general migration of business correspondence from paper to bits. In reality, electronic commerce is flourishing, more and more business documents are created, sent, and stored only electronically, and digital signatures are playing a very minor role in this transformation.

Odlyzko [ODL1] compares digital signatures to fax signatures:

> The 1980s were the golden age of civilian research on cryptography and security. The seeds planted in the 1970s were sprouting and the technologists' bright hopes for a brave new world had not yet collided with the cold reality as clearly as they did in the 1990s. Yet the 1980s were also the age of the fax, which became ubiquitous. With the fax, we got fax signatures. While security researchers were developing public key infrastructures, and worrying about definitions of digital signatures, fax signatures became widespread, and are now playing a crucial role in the economy. Yet there is practically nothing as insecure as a fax signature, from a formal point of view. One can easily copy a signature from one document to another and this will be imperceptible on a fax. So what lessons can we draw from fax signatures, other than that convenience trumps security?

One plausible answer to this question is given in [ODL1]:

> One lesson is that our society somehow managed to function even when signatures became manifestly less secure with the spread of fax signatures. Moreover, it is easy to argue that fax signatures have contributed greatly to economic growth. How did this happen? This occurred because there is a context to almost every fax communication.
>
> Although fax signatures have become widespread, their usage is restricted. They are not used for final contracts of substantial value, such as home purchases. That means that the insecurity of fax communication is not easy to exploit for large gain. Additional protection against abuse of fax insecurity is provided by the context in which faxes are used. There are records of phone calls that carry the faxes, paper trails inside enterprises, and so on. Furthermore, unexpected large financial transfers trigger scrutiny. As a result, successful frauds are not easy to carry out by purely technical means. Insiders (as at Enron and WorldCom and innumerable other enterprises) are much more dangerous.

Another plausible explanation may be found in simple cost-benefit analysis. In situations in which digital signatures are desired, it is typically the receiver who wants to guarantee the provenance and preclude the repudiation of the message. However, the sender also incurs the cost of using digital signatures but receives little or no benefit from a positive adoption decision. Because it is usually quite easy to get by temporarily with plaintext, unsigned email exchanges and then use a signed paper letter or face-to-face meeting in order to "close the deal," neither senders nor receivers are strongly motivated to adopt PKI-based digital-signature schemes.

4.2. Challenging Issues

Research is needed into all aspects of security-technology deployability and adoptability. At least three types of systematic study could shed light on the apparently paradoxical situation that we now find ourselves in (incessant complaints about Internet insecurity coexisting with a vast array of impressive security technology built up over decades of research and development). The research community should be encouraged to expand upon these three.

First, systematic study is needed of past failed deployments. For example, there have been well funded, carefully designed proposals for public-key infrastructures, secure interdomain routing, secure domain-name service, and, in the application layer, encrypted email and secure digital-rights management. All were well motivated, but none was widely adopted. It could be quite instructive to determine whether these were "solutions to the wrong problems," solutions to problems that actually were not as severe as technology developers had been led to believe, victims of some other fundamental adoption barrier, or simply ahead of their time.

Second, research is needed on the effectiveness of the security technology that *has* been widely adopted. For example, firewalls may be seen as a ray of hope about the willingness of individuals and organizations to buy and install security solutions. Yet, there are mixed reports about how much security is actually provided by firewalls as they are currently used, and there is evidence that even IT professionals do not know how to manage and configure them. Some work along these lines has been done, *e.g.*, [BMNW], but more is needed.

Third, future security-technology research teams should pay far more attention to adoptability, deployability, and migration paths than previous teams have done. Again, there is some evidence that this is starting to happen; see, *e.g.*, the work of Goodell *et al.* [GAG+] on secure interdomain routing, in which incremental deployment is explicitly emphasized. As a community, security researchers should identify good "test cases" with which to assess future progress on adoptability. Candidates include spam fighting, privacy-preserving data mining, open-science environments, and DDoS fighting. Research needs within these areas are discussed elsewhere in this report. Here we wish to point out that there *is* consumer demand for solutions to these problems; thus, if technically good solutions are *not* adopted, the experience should provide a lesson about deployment obstacles.

As in Sections 2 and 3, the research needs identified in this section highlight the importance of data collection. Without ongoing access to accurate data about who is using what and how effective it is, how will adoption decisions ever truly be understood? The role of economic, political, social, and ethical issues in adoption decisions is further evidence that interdisciplinary research teams are needed.

5. Problem 4: Formal Techniques in Cyber Security

5.1. The Need for Proofs

Security is unusual among computer-science research areas in that virtually all researchers in the area, including system builders, acknowledge the need for formal models and proofs of security. This is because experimentation alone can never establish that a system is secure. Testing can reveal security flaws, but, once all known flaws have been fixed, there is no guarantee that adversaries will not find more flaws after the system is fielded. Thus there is great value in proofs that system designs satisfy formal security properties.

Various theoretical-cs research communities have consistently and energetically responded to this opportunity, and there is now an extensive body of knowledge about security theory. Indeed, as indicated in Section 1 above, some of the highlights of this theory are regarded as crowning achievements of computer-science research, and the leading researchers in the field are highly decorated. Two major (and largely disjoint) technical approaches to formal aspects of security are the algorithmic and complexity-theoretic approach (in which the emphasis is on reductions that prove that an adversary's

job is formally equivalent to an instance of a well-studied problem for which no efficient solution is known) and the logical approach (in which the emphasis is on formal specification of security properties and automated analysis of system designs with the goal of finding security vulnerabilities or proving compliance with specifications).

This extensive body of theoretical research has made some important practical contributions to network security. Most recently, zero-knowledge proof systems have been used in commercial trusted-platform development [BRI1].[5] Nonetheless, there was widespread agreement at the workshop that many more contributions are possible and that many opportunities have been missed over the years. These missed opportunities are bi-directional: Solutions to many practical security problems can be found in the theoretical-research literature, but often they are *not* found because much of the theory is not widely understood or appreciated; on the other hand, theoretical problems are often posed in ways that do not fully capture real-world security problems, because real-world security is also very poorly understood. This basic conclusion influenced the workshop findings in two major ways. First, the research needs identified in Sections 2, 3, and 4 above will clearly require theoretical as well as experimental work and will clearly be interesting to theorists; if they are addressed successfully, they will move security theory in fundamentally new directions. Second, established lines of theoretical research in security are of practical value and should be continued; because some of the techniques developed by the cryptographic-research community are extraordinarily powerful (almost paradoxically so), the research goals of that community should include a broad educational mission and a compelling story about the practical value of this theory in large-scale-network security.

5.2. Challenges of formal techniques in Cyber Security Research

The following research topics were discussed at the workshop and are likely to prove fruitful in the short-to-medium term. They are meant to be a *sample* of the many theoretical-cs research projects that could contribute to large-scale-network security (not an exhaustive list!) and to demonstrate the intellectual breadth of the theoretical-cs contribution. Some of the research needs identified here fall into the general category of needs for *further* research along fruitful lines already established in various theoretical-cs communities, and others represent entirely new directions.

Economic Theory and Security: Real-world aspects of user demand for security were discussed extensively at the workshop, because of the focus on "deployment obstacles." Workshop participants also raised the question of whether economic theory could be used to formulate new definitions, models, and theories of security that have some of the desirable properties identified in Section 3 above. For example, security *per se* is rarely a primary goal of users; instead, users typically want to accomplish a specific goal (*e.g.*, web searching or email), for which they use a specific network service or protocol and expect a certain level of performance at a certain cost. A plausible definition of a

[5] See Sections 5.2 and 6.2 below for an explanation of the term "trusted platform" and of research needs associated with these development efforts.

"secure" service or protocol is one that maintains good performance in the presence of adversaries. With such definitions in hand, the natural questions would be whether security raises the cost of network services and protocols, if so by how much, and whether users are willing to pay this increased cost. Research is needed in order to formulate and answer these questions for various classes of network services, notions of performance, and classes of adversaries. Preliminary work along these lines can be found in, *e.g.*, [MEA1] and [WEIS].

Formal Modeling and Analysis of IETF Standards: Because basic Internet protocols are specified in open standards documents, these protocols (both those already in use and those still in the proposal stage) provide excellent test cases for new protocol-analysis techniques. Research is needed into hybrid approaches to protocol specification and verification, *i.e.*, approaches that combine successful techniques from algorithms and complexity theory with those from logic. New definitions and models that capture system-level security properties instead of component-level properties (as discussed in Section 3) will be useful only if scalable, understandable protocols can satisfy them; IETF standards should be analyzed with this is mind, and the results should inform the development and refinement of new definitions and models.

Privacy-Preserving Data Mining: As concern for homeland security has grown, data-mining needs have become more distributed and more urgent. Various types of sensors are becoming more and more prevalent, both online and offline, and thus the volume of data that are both potentially relevant for security purposes and potentially damaging if leaked or misused is expected to increase greatly. There is a wide-ranging need for data-management techniques that balance the civil rights and property rights of data subjects and data owners with the rights of law-enforcement and other organizations to make legitimate use of sensitive data sets.

Although the phrase "privacy-preserving data mining" may at first seem like an oxymoron, there is actually no inherent contradiction between the goals of discovering a security-critical needle and ignoring a haystack full of security-irrelevant, private information. Computing exactly one relevant fact about a distributed data set while concealing everything else about it is precisely what cryptographic theory enables *in principle*. Research is needed into the questions of how to make the protocols in the theoretical-cs literature efficient enough for use on massive data sets and how to implement them on realistic network models. Specific data-mining algorithms that have proven useful should be studied from the viewpoint of privacy preservation, and new algorithms should be developed if the standard ones prove not to be efficiently implementable in a privacy-preserving manner. Principled use of approximation (*e.g.*, tradeoffs between accuracy and privacy) should be explored in this context. Preliminary work along some of these lines can be found in, e.g., [LiPi] and [FIM+].

Data-mining research needs may seem out of place in this report, because the workshop was about *network* security, not data security. However, one of the consistent themes of the workshop was that more and better data about the security of deployed systems are needed in order to meet LSN-research challenges. A great deal of these data will be

sensitive and company-proprietary, and organizations will not provide them unless they can be mined in a privacy-preserving fashion.

Security-Enforcing Languages: Language theory can contribute to large-scale network security on at least two fronts: programming languages and policy languages. Within the programming-language arena, research is needed to improve static and dynamic guarantees, including but not limited to type systems that support security requirements, distributed run-time environments, and support for resource bounding. Trade-offs between computational efficiency and security guarantees must also be studied within the programming-language framework. Research is also needed in the general area of programming languages for trusted platforms and other new computer architectures.

Within the policy-language arena, research is needed on the formal foundations of security-policy languages, on compilers that translate high-level policies into enforcement mechanisms, and into policy languages with good interoperability and modularity properties. Progress on policy languages should be integrated with work on algorithms and protocols, including work on privacy-preserving data mining; ideally, each data owner contemplating participation in a distributed data-mining exercise should have a machine-readable privacy policy, each data-mining system should have a machine-readable specification of what it does and does not reveal about a data set, and a policy-compliance-checking component should give data owners the information they need in order to decide whether they are willing to feed data to or receive data from the system.

A specific real-world problem that policy-language research can help to solve is network integration of the type that occurs when corporations or other administrative domains merge; how can administrators of the previously separate intranets ensure that security requirements that were previously enforceable remain enforceable when the networks are integrated, particularly if the policies were developed and expressed using different specification languages?

Trusted Platforms: The notion of a "trusted-computing base," or TCB, is central to standard cryptographic theory. Informally, the TCB is the part of the computing environment that has to be trustworthy; if it is corrupted, then no guarantees about the correctness or security of the rest of the environment can be made. The assumption that TCBs are available is so fundamental that it is hard to see how cryptographic theory could be constructed without it; for example, how could one use any cryptosystem without a private, correctly functioning environment in which to store private keys and perform encryption and decryption operations?

Yet, typical users of today's networked computers do *not* have TCBs. Our computing environments are riddled with software bugs that make them vulnerable to intrusions, denial of service, and other attacks. This situation is not likely to change in the absence of a major discontinuity in mass-market computing, because people have become accustomed to installing new applications and upgrades frequently, some of the provided by vendors and non-commercial websites that make no quality guarantees – even as some

known vulnerabilities are fixed with security patches, new bugs and vulnerabilities are installed. Indeed, the research need for new definitions and models, particularly "autonomic security," that was identified in Section 3 stems in large part from the realization that the textbook model of cryptographic protocols running on top of TCBs does not correspond to our computing reality.

A different way to approach this mismatch between theory and practice is to try to build TCBs and network them together. This is the goal of the "trusted-platform" initiatives [MIC1, TCPA] now underway in the computer industry. Informally speaking, an "attestable TCB" is one that can prove (*e.g.*, by supplying digitally signed statements) to a remote machine that it is running on the local machine and that the required properties of the local computing environment hold. The industry's "trusted-platform" vision is that attestable TCBs will be successfully mass marketed, that they will be networked together, and that lucrative applications that can make use of them, such as entertainment-content distribution, will be developed and adopted.

Research is needed in the theoretical foundations of trusted platforms. As of now, there is not even a formal definition of the term "trusted platform" that is widely accepted and captures the essence of the commercial-development efforts that are underway. All aspects of the theory of trusted platforms should be studied; interesting questions include but are not limited to the following. Is our current textbook cryptographic theory indeed directly implementable on the type of trusted platforms we may soon have available, or does the theory need to be adjusted? Do the trusted platforms now under development satisfy the formal definitions? Which tools are most appropriate for verification of trustworthiness as defined? Is there a good migration path from today's untrustworthy PCs to trusted platforms, or are trusted platforms likely to encounter the same deployment obstacles that plagued previous security-technology development efforts?

UI Aspects of Security: User-interface problems are widely recognized as significant practical barriers to effective use of security technologies and procedures, but they are essentially ignored in the extensive theoretical-cs literature on security. Research is needed in order to develop formal frameworks that require technology to be usable in order to be considered "secure." Definitions and proof techniques are needed to enable rigorous analysis of "ease of use" and "hardness of misuse" of security mechanisms.

6. Problem 5: Experimental-Research Needs for Security of Large-Scale Networks

Almost all of the research needs identified in the previous four sections have an experimental dimension. Because most of the necessary background and general discussion needed for this section can be found in the four sections above, we will devote all of this section to experimental-research needs. Specifically, we will make some high-level, methodological remarks about the experimental research that is needed, identify some specific projects that would be desirable, and explain the very pressing need for a large-scale data-collection infrastructure.

6.1. Methodological Aspects of Experimental-Research Needs

Intellectually Diverse Teams: It was widely acknowledged during this workshop that our discussions were enhanced by the intellectual diversity of the participants. The advantages of multidisciplinary research teams have been discussed in earlier sections. Here we remark on the advantages of collaboration by theorists and experimentalists. In particular, the experimental projects recommended in this section would benefit greatly from participation by theoretical–CS researchers. Solicitations and Calls for Proposals should explicitly encourage multi-PI projects in which at least one PI is primarily a theorist, and at least one is primarily an experimentalist.

Ab Initio Designs and Prototypes: Although the computer science world is justifiably concerned about the security of the current Internet and its constituent protocols, there is still a need for *ab initio* designs and implementations of "secure systems." This type of "pure research" in the secure-systems arena should include both systems that provide entirely new features and functionality (designed from the beginning with security in mind) and systems that provide fundamentally new, secure approaches to network functions that we are already using (*e.g.*, naming, routing, email, and file transfer).

Improved Security of Existing Systems: The following observation in [NRC4] is still a research imperative:

> Laudable as a goal, *ab initio* building of trustworthiness into a networked information system has proved to be impractical. It is neither technically nor economically feasible for designers and builders to manage the complexity of such large artifacts or to anticipate all of the problems that a [networked information system] will confront over its lifetime.

Although *ab initio* designs have traditionally dominated the research literature, and there is still a need for "pure research" along these lines, research is also needed in the general area of "securing our existing infrastructure." Incremental deployability, as discussed in Section 4, is part of this agenda; so is monitoring, data collection, assessment, and risk management.

Open Source: It is often said that open-source systems are more secure, because they undergo constant scrutiny by programmers who were not necessarily involved in the original design and implementation of the system (and who thus can spot problems that insiders may be blind to) and can benefit from those programmers' willingness to fix bugs and security vulnerabilities. Like many other pieces of common wisdom in security, this one has received very little systematic study. Research is needed in order to formulate the security advantages offered by the open-source development process, to improve the process if possible, to make these security-enhancing processes prominent parts of programmer training and software-development best practices, and to transfer as many benefits as possible to traditional, closed-source development of commercial software. For the last two years, DARPA has been funding open-source security research through the CHATS (Composable High-Assurance Trusted Systems) program. Early

results of this program have already had a positive influence on the OpenBSD, FreeBSD, and Linux communities. Unfortunately, it seems unlikely that DARPA will continue this program. Because early results are promising and much remains to be done, it would be highly desirable for some other agency to pick this up. One important research need that has not yet been addressed but that follows naturally from the earlier work in this program is the question of whether there are security problems that are *created* or *exacerbated* by the open-source development and distribution models and whether they can be avoided by security-conscious development communities.

Trustworthy Composition of Systems: The following observations from [NRC4] were echoed again and again at this workshop:

> Security research during the past few decades has been on formal policy models that focus on protecting information from unauthorized access by specifying which users should have access to data or other system objects. It is time to challenge this paradigm of "absolute security" and move toward a model built on three axioms of insecurity: insecurity exists; insecurity cannot be destroyed; and insecurity can be moved around. …
>
> Improved trustworthiness may be achieved by the careful organization of untrustworthy components. There are a number of promising ideas, but few have been vigorously pursued. "Trustworthiness from untrustworthy components" is a research area that deserves greater attention.

Research is needed to determine the viability of "trustworthiness from untrustworthy components." What does it mean to "trust a component" of a large, heterogeneous system? Must there be at least one "trustworthy component" in a system for the system as a whole to be trustworthy? Groups of people learn whom to trust over time and with experience. Can we build systems that "learn" what to trust over time and with experience and that adapt to variations in the trustworthiness of components? Note that composability and its effect on trustworthiness is likely to grow in importance, because large-scale collaborative computing efforts such as open science and peer-to-peer will require frequent recomposition and reconfiguration as the sets of people and machines involved in projects changes.

6.2. Challenging Issues

Privacy-preserving data mining: The need for theoretical research in the area of privacy-preserving data mining was explained in Section 5. Here, we stress that experimental research in this area is also essential. Privacy-preserving data-mining algorithms should be implemented and tested, and their effectiveness should be demonstrated through use on real data sets. Federal government data-mining needs present an excellent opportunity to position government agencies as technology leaders in this field. Multi-institutional, multi-year research proposals for a distributed experimental testbed for privacy-preserving data mining should be solicited. Note that success in this area could have a

significant positive impact on other research needs identified at the workshop: Owners and operators of data networks might ultimately be willing to contribute much-needed data about actual network insecurity if there were practical, scalable, privacy-preserving data-mining systems up and running, and their effectiveness had been demonstrated on valuable government data sets.

Configuration Management: Improper configuration of systems is well known to be a source of insecurity. For example, the following observation is made in [NRC2]:

> Many compromises of an information system or network result from improper configuration…. Because checking operational configurations is very labor-intensive if done manually, it is essential to have configuration management tools for both systems and networks that can automatically enforce a desired configuration or alert administrators when variances from known configuration are detected. Such tools are miserably inadequate today.

[NRC2] goes on to recommend that vendors of computer systems should:

- Drastically improve the user interface to security, which is totally incomprehensible in nearly all of today's systems. Users and administrators must be able to easily see the current security state of their systems; this means that the state must be expressible in simple terms.
- Develop tools to monitor systems automatically for consistency with defined secure configurations, and enforce these configurations.

Research is needed in all aspects of configuration management, from underlying principles (which are woefully underemphasized in academic computer-science research and largely absent from computer-science curricula) to deployable tools. Because most of the configuration-management tools now in use are both proprietary and inadequate, this is a good area in which to pursue *ab initio* designs and prototypes. Note that projects of this sort would fit well with other research needs identified at the workshop; for example, systematic study of the "ease of use" and "hardness of misuse" of configuration tools fits well with the identified need to improve UI aspects of security and to develop better security metrics, and systematic study of how best to specify secure configurations fits well with the identified need to improve policy languages and compilers that translate policies into configurations.

Trusted Platforms: Experimental research is needed to complement both the commercial-product development underway by Microsoft [MIC1] and the Trusted-Computing alliance [TCPA] and the theoretical-research needs in this area that were discussed in Section 5. At least three types of experimental research projects are needed. One is experimental assessment of the trustworthiness of the commercial "trusted platforms" that will soon be rolled out, followed by constructive suggestions by the research community about how to improve these platforms. By analogy with the flowering of

language-based security research that followed Sun's rollout of Java, we can expect interesting activity in the trusted-platform area during the next few years. A second type of research project that should prove fruitful is alternative designs and prototypes of the platforms themselves. A third type is design and prototyping of novel distributed applications that run on trusted platforms. The computer industry is quite focused on the potential use of trusted platforms for mass-market distribution of entertainment content and for enterprise-based document flow. Research is needed into their applicability in other contexts, including privacy-preserving data mining, open science, and other themes identified at this workshop.

Open-Science Environments: Science is increasingly conducted in cyberspace. Because the scientific enterprise is quite different from many others in which networks and computers are playing an increasingly prominent role, threat models are very poorly understood in the open-science area. Experimental research is needed to ensure that the cyber infrastructure supports collaboration by geographically (and sometimes culturally) separated team members as well as timely dissemination of results but that it also protects the integrity of scientific data and software. Security and privacy policies may govern the release of data, requiring, for example, that only aggregated data sets be released or that data undergo privacy-preserving transformations before release; these considerations are particularly relevant in scientific disciplines that use human subjects or classified government data. Research projects in cybersecurity for open science can leverage the results of other projects suggested in this report, including but not limited to those on policy languages, configuration management, privacy-preserving transformations of data sets, trusted platforms, and threat modeling. Federal-government support of scientific research provides an opportunity to position the government in general and National Labs in particular for technology leadership in the area of cybersecurity for open science.

New Internetwork Designs: Apart from "survivability in the face of failure," security was not among the original Arpanet design goals. Security mechanisms have been added as the network has evolved into a mass-market computing and communication system, but it is not surprising that security remains a concern, given that it was not a primary objective of the original design. Experimental research is needed into *ab initio* internetwork designs that prioritize security from the beginning. If perfect security is indeed incompatible with other fundamental Arpanet design goals, can specific security goals be met by new designs that also support the best properties of today's Internet? For example, is there a way to eliminate denial-of-service attacks without sacrificing too much? Experimental networking-testbed projects need to incorporate various sets of security properties as first-class design goals. Testbed projects that are successful from a security viewpoint will be candidates for commercial development, and research is also needed into migration paths from the current network architecture to the newer, more secure architecture.

6.3. Call for an LSN-Security Data Center

The most compelling and consistent theme of this workshop was that collection and analysis of data about network security is a central requirement for real progress on large-scale-network security. This is not the first time that this conclusion has been drawn, as the following recommendation from [NRC3] shows:

> Good information systems security requires an understanding of the types of threats and defenses that might be relevant. Thus, those responsible for information systems security need a vigorous ongoing program to monitor, assess, and understand offensive and defensive information technologies. Such a program would address the technical details of these technologies, their capability to threaten or protect friendly systems, and their availability.

Workshop participants recommend the creation of a unified research initiative, on the scale of an NSF Science and Technology Center, to address all aspects of LSN-security data. Successful proposals for research funding in this initiative must address the privacy concerns of network owners and operators who would be expected to contribute data. Research needs addressed by this initiative would include all aspects of measurement, modeling, monitoring, definitions, and cost-benefit analysis identified in this report.

7. References

[AND1] R. Anderson, "Why Information Security is Hard – An Economic Perspective," http://www.cl.cam.ac.uk/users/rja14/econ.pdf

[BEC1] G. S. Becker, "Crime and Punishment: An Economic Approach," in **The Essence of Becker**, R. Febrero and P. Schwartz (eds.), Hoover Institute Press, Stanford, 1995, pages 463-517.

[BMNW] Y. Bartal, A. Mayer, K. Nissim, and A. Wool, "Firmato: A Novel Firewall Management Toolkit," in *Proceedings of the 1999 IEEE Symposium on Security and Privacy*, pages 17-31.

[BrSh] L. Breslau and S. Shenker, "Best-Effort versus Reservation: A Simple Comparative Analysis," in *Proceedings of the 1998 ACM Sigcomm Conference*, pages 3-16.

[BRI1] E. Brickell, "The Direct Proof Method for Anonymously Proving Integrity," presented at the *2003 RSA Conference*.

[FFSS] J. Feigenbaum, M. Freedman, T. Sander, and A. Shostak, "Privacy Engineering in Digital Rights Management Systems," in *Proceedings of the 2001 ACM Workshop on Security and Privacy in Digital Rights Management*, pages 76-105.

[FIM+] J. Feigenbaum, Y. Ishai, T. Malkin, K. Nissim, M. Strauss, and R. Wright, "Secure Multiparty Computation of Approximations," in *Proceedings of ICALP 2001*, pages 927-938.

[GAG+] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing," in *Proceedings of the 2003 Symposium on Network and Distributed Systems Security*, pages 75-85.

[IBM1] IBM Research, "Autonomic Security: Compromise Tolerant Distribution, Compromise Detection, and Compromise Recover," LSN workshop presentation, March 2003.

[LiPi] Y. Lindell and B. Pinkas, "Privacy-Preserving Data Mining," *Journal of Cryptology* **15** (2002), pages 177-206.

[MEA1] C. Meadows, "A Cost-Based Framework for Analysis of Denial of Service in Networks," *Journal of Computer Security* **9** (2001), pages 143-164.

[MIC1] "Microsoft Next-Generation Secure-Computing Base-Technical FAQ," http://www.microsoft.com/Technet/security/news/NGSCB.asp

[NRC1] Computer Science and Telecommunications Board, National Research Council, **Computers at Risk: Safe Computing in the Information Age**, National Academy Press, Washington, DC, 1991.

[NRC2] Computer Science and Telecommunications Board, National Research Council, **Cybersecurity Today and Tomorrow: Pay Now or Pay Later**, National Academy Press, Washington, DC, 2002.

[NRC3] Computer Science and Telecommunications Board, National Research Council, **Realizing the Potential of C4I: Fundamental Challenges**, National Academy Press, Washington, DC, 1999.

[NRC4] Computer Science and Telecommunications Board, National Research Council, **Trust in Cyberspace**, National Research Council, Washington, DC, 1999.

[ODL1] A. M. Odlyzko, "Economics, Psychology, and Sociology of Security" to appear in *Proceedings of Financial Cryptography 2003*.

[SCC1] Static Control Corporation: Computer Chip Usage in Toner Cartridges and Impact on the Market: Past, Current and Future. White paper, dated Oct. 23, 2002, available at (http://www.scc-inc.com/special/oemwarfare/default.htm).

[TCPA] "Trusted Computing Platform Alliance," http://www.trustedpc.org.

[VAR1] H. Varian, "Managing Online Security Risks," Economic Science Column, *The New York Times*, June 1, 2000, http://www.nytimes.com/library/financial/columns/060100econ-scene.html

[WEIS] Workshop on Economics and Information Security, May 16-17, 2002, Berkeley CA. http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity

8. Participants

Joan Feigenbaum (Workshop Chair), Yale University

Mustaque Ahamad, Georgia Institute of Technology
Ian Alderman, University of Wisconsin – Madison
Jules Aronson, National Institutes of Health
Wu-Chun Feng, University of Illinois at Urbana-Champaign
Michael Fischer, Yale University
Mike Gill, National Institute of Health
Mike Greenwald, University of Pennsylvania
Carl Gunter, University of Pennsylvania
Stuart Haber, Hewlett Packard Laboratories
Sally Howe, National Coordination Office for Information Technology Research and Development
Russell Impagliazzo, University of California at San Diego
John Ioannidis, AT&T Labs
Stanislaw Jarecki, Stanford University
Admela Jukan, National Science Foundation
Rick Kennell, Purdue University
Angelos Keromytis, Columbia University
Hugo Krawczyk, IBM Research
Olga Kuchar, Pacific Northwest National Laboratory
Dirk Kuhlmann, Hewlett Packard Laboratories
Mike Langston, University of Tennessee
Wenke Lee, Georgia Institute of Technology
Ninghui Li, Stanford University
Grant Miller, National Coordination Office for Information Technology Research and Development
Doug Montgomery, National Institute of Standards and Technology
Terry Moore, University of Tennessee
Thomas Ndousse, Department of Energy
Andrew Odlyzko, University of Minnesota
David Parkes, Harvard University
Tal Rabin, IBM Research
Vijay Ramachandran, Yale University
James Rothfuss, Lawrence Berkeley National Laboratory
Dave Safford, IBM Research
Andre Scedrov, University of Pennsylvania

Carsten Schurmann, Yale University
Vitaly Shmatikov, SRI International
Frank Sibenlist, Argonne National Laboratory
Avi Silberschatz, Bell Laboratories
Gene Tsudik, University of California at Irvine
Rebecca Wright, Stevens Institute of Technology
Moti Yung, Columbia University
William Yurcik, University of Missouri-Kansas City
Steve Zdancewic, University of Pennsylvania
Sheng Zhong, Yale University
Taieb Znati, National Science Foundation