LBL-38669
UC-400

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Information and Computing Sciences Division

Presented at the *Internet Security Workshop*,
Haystack Observatory, MA, May 20–21, 1996,
and to be published in the Proceedings

Multi-Cultural Network Security

D.F. Stevens

April 1996

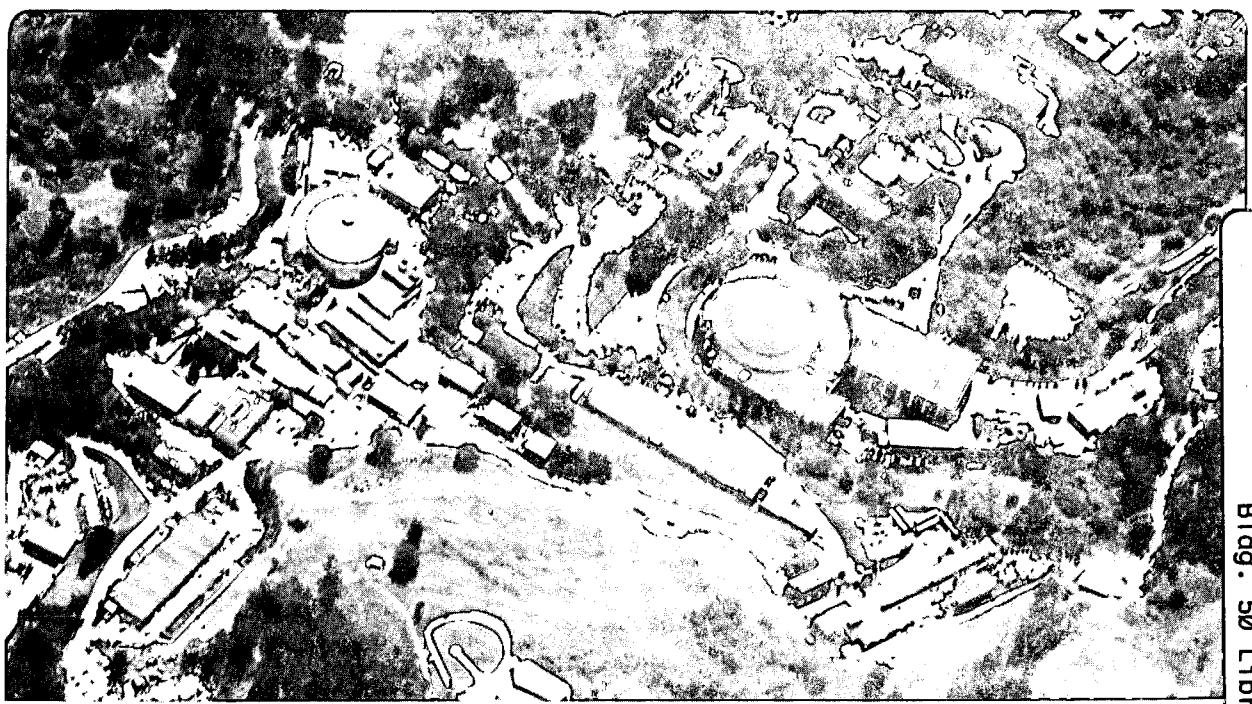

REFERENCE COPY
Does Not
Circulate
Bldg. 50 Library.
LBL-38669
Copy 1

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBL-38669
UC-400

Multi-Cultural Network Security

David F. Stevens

Information and Computing Sciences Division
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, California 94720

April 1996

Prepared for the
Internet Security Workshop
Haystack Observatory, Massachusetts
May 20-21, 1996

This work was supported by the Director, Office of Energy Research, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Recycled Paper

Multi-cultural Network Security*

D. F. Stevens
Ernest Orlando Lawrence Berkeley National Laboratory
Berkeley, CA 94720

Premises: 1. Computer security in an open environment depends upon informed people acting responsibly.

2. The Internet user community incorporates many cultures, with different attitudes towards computer and network security.

3. To be effective, security procedure must be tailored to the culture of the user community involved.

4. Current network security policy is tailored to the culture of the security community that created it.

Conclusion: Current network security policy will not be very effective if it is applied to the Internet.

1. Introduction

Education and awareness are widely acknowledged to be among the fundamental issues of Internet security, but only in the sense of making Internet users more security conscious. For the Internet to achieve its promise as an information highway, however, a complementary education effort is needed. If adequate Internet security is to be achieved, we must also increase the awareness of the professional security community of the requirements, attitudes, and habits of the many different cultures that participate in the Internet.

Discussions of "the Internet" encourage the misapprehension that there is a single, uniform user community instead of a loose alliance of many cultures that differ in many fundamental aspects. This is true even if we limit our consideration to ethical cultures. At this Workshop alone we have representatives of administrative and military cultures, Governmental and commercial cultures, profit-cultures and non-profit cultures, research and operational cultures.

Internet cultures are united in their desire to exploit the connectivity, flexibility, and rapidity of communication provided by the net, but differ greatly in their motivations, their attitudes towards authority, their willingness to cooperate within their own communities, their interest in technical arcana, and the patience with which they will put up with—or the enthusiasm with which they will embrace—the growing list of procedures deemed necessary for acceptable security. They even differ in how they define "acceptable security".

* This work was supported by the U. S. Department of Energy under contract No. DE-AC03-76SF00098.

This problem has been addressed to some extent by the introduction of *enclaves* into the Internet security model. While this can be expected to mitigate the problem, it seems unlikely to eliminate it, because many enclaves will themselves be multi-cultural in the sense used here. There is significant anecdotal evidence that the imposition of excessive procedural requirements on populations that do not respect such approaches will lessen security rather than increase it. Thus it is incumbent upon the designers of Internet security policy to understand the influence of culture on security, and to design their policies to accommodate the expected diversity.

An exhaustive discussion of all the cultures and their attitudes towards network security is probably beyond the scope of the Workshop and certainly beyond the purpose of this paper. My intent here is to take a single example—the collaborative research culture with which I am familiar and the computer and network security culture with which it must deal—to illustrate the tensions that arise when security policy is defined without regard for the culture upon which it is imposed. The paper concludes with a few aphoristic reminders intended to help the policy makers to retain the desired flexibility of approach.

2. Cultures in Conflict: Collaborative Research and Network Security

The characteristic research environment today includes personal workstations networked locally to each other and to various servers, with further direct or indirect connections to national and international networks. "Big" research is characterized by broader and more diverse collaborations, with individual scientists using facilities, and even controlling experiments, at institutions located hundreds or thousands of miles away.

Each scientist is essentially autonomous, acting as his own system manager, and free to join a large number of research communities that collectively embody a wide range of formality, coherence, and observance of sound security practice. Such central policies as exist, do so by consent rather than by decree. To the scientist, "the Network" is the sum of all the communities to which he belongs. Within this extended "Network" there is no overall authority, no common set of security precautions, and not even a general agreement about which aspects of security are properly the responsibility of "the Network" and which are the responsibility of the individual nodes. Few of the users have more than the most rudimentary knowledge of the topology of the connecting infrastructure, or where the responsibility for its administration lies. They have little knowledge of where the responsibility for security lies within any node other than their own. There is no generally recognized point of contact for the dissemination of information about attempted penetrations. In sum, the research scientist treats "the Network" much as he would the telephone system: as a medium of communication about which he needs to know nothing except how to address his colleagues.

The computer security professional cannot take this black-box approach to "the Network". He must know both the physical and administrative topologies—the nodes and connections, and the Authorities in charge of each—and his inability to do so is a source of great frustration. Instead of rejoicing at the breadth of possible connections, he worries about the number of unknown "users" who have access to the many gateways into his domain.

These, and other, considerations give rise to two cultures that differ in many fundamental ways. While these differences need not lead to overt conflict, they are bound to create significant tension. We list here a few of the many examples of the differing viewpoints that contribute to this tension.

- The research community tends to be oblivious to the possibility that any of their data or systems might be the targets of annoying or malicious penetration attempts. They have no secrets, and so they believe they are not worth attacking; they assume that all persons with network access will always act ethically, correctly, helpfully, and harmlessly. Computer security professionals know that all systems are under more or less constant attack, simply because they are there, and the attackers are capable of inflicting both accidental and deliberate damage.

- Collaborative research environments are collegial. They are loose confederations, spread over many organizations, often in several countries, and there is no single chain of responsibility or command. There is no common authority to whom all the participants in a collaborative research environment report. By contrast, the organizations that develop traditional computer security procedures are strongly hierarchical, and the procedures they produce assume the existence of a single, *strong* chain of command. Traditional security organizations do not understand the collegial style, and their procedures cannot cope with it.

- Collaborative research environments are characterized by a lack of routine. Researchers are free, within broad limits, to establish their own procedures, and to change them essentially instantaneously to take advantage of new insights, new equipment, new technology, and the talents of new personnel. Standard security practice, on the other hand, is designed to force operations to be routine, and to remove freedom and discretion from the manner in which they are performed.

- Collaborative research environments are designed to facilitate connectivity, the sharing of knowledge among all the members, and the cross-pollinization of ideas throughout the whole domain of enquiry. By contrast, traditional security measures include the principle of *division of responsibility*, which partitions each domain into single-function regions and erects impenetrable fences between them.

- Policy and practice in collaborative research environments are generally aimed at maintaining and increasing access, to the broadest posible extent. Traditional computer security is founded on a basic distrust of everyone not known, validated, and authenticated to/by the system, and is aimed at restricting access to the narrowest possible population.

- Collaborative research often involves the presence at a site of a constant flow of visiting workers, whose time on the spot is brief and whose need to share information is great. The imposition of even relatively simple security precautions, such as the prohibition of

password-sharing, or the elimination of any password or system ID known to a departed visitor, can place a significant procedural burden on the stable core of the group. On the other hand, failure to observe these precautions can make the achievement of even a minimally acceptable level of security impossible.

3. Reconciling the Worlds of Security and Research

The problem in a nutshell is to allow unlimited connectivity while providing adequate security. For this to be accomplished, both the researchers and the security professionals are going to have to modify their accustomed ways of doing things to some extent. If this is to be done without a full-scale clash of cultures, the initial adjustment must come from the security establishment, primarily because, whereas connectivity can be achieved without security or the cooperation of the security professionals, security cannot be achieved without the cooperation of the research community. The research community contains many extremely creative people whose culture is to view constraints—whether natural or administrative—as challenges to be overcome rather than as limits within which to work. An unthinking traditional approach to security is as likely to motivate them to circumvention as to obtain their cooperation. On the other hand, if the research community can be convinced that a modicum of security is in their best interest, they will turn their creative energies to supporting security instead of subverting it.

How, then, can their cooperation best be achieved? By establishing a security program that is designed for the milieu in which it is set. Such a program cannot depend upon rigid enforcement of centrally-imposed rules, but must distribute responsibility for security in the same manner that responsibility for the quality of the research is distributed. It must recognize a realistic span of control. It must admit of the existence of multiple domains with differing requirements, but with the need to communicate with each other. It must assume a reluctant population, and minimize the burden that accompanies the implementation of the security program. It must provide each user with the knowledge and the tools necessary to maintain adequate security. It must make it simple to install and use elementary precautions. It must make it difficult or awkward to maintain known security flaws.

Traditional security wisdom suggests that trusted *systems* are the highest form of the security practitioner's art; I suggest rather that even trusted systems depend upon proper operation by people, and that a successful program in the collaborative research environment must depend upon trusted *people*.

4. Words to the Wise (being a brief compendium of useful precepts and exegetical remarks addressed to designers of network security policies)

The following is not intended to be a complete course on how to design computer and network security policies for use in the research community, but rather a suggestive list of hints as to why the traditional approach might not be successful. All of these precepts have validity beyond the present arena, but most of them have particular pungency in this context. (They all appear, with full references, in *Why Computer Security Systems Don't Work* (LBL-30352; D. F. Stevens, 1991.)

A ship in a harbor may be safe, but that's not what ships are built for.
Computer security is about as useful as bones in a potato.

In other words, computers are not installed to be secure, they are installed to do work. If a choice must be made, the researcher will invariably choose work, and will resent whatever interference is afforded by the security policy.

Procedures convert means to ends.

It seems to be in the nature of Government to create procedures that accrete and solidify to thwart creativity. These procedures may start out as "mere" formalities to regularize repetitive activities, but they eventually achieve a life of their own. Thus what began as a tool to achieve some human (or bureaucratic—not necessarily the same thing) goal eventually becomes a goal in its own right. A guideline for the development of a computer security program can easily solidify into a rigidly defined Computer Security Plan whose contents and level of detail are independent of the threat being addressed. The plan becomes the goal, and the protection it is presumed to enhance becomes a mere side effect: the procedural goal replaces the human goal.

Loose systems last longer and work better.

Loose systems allow for the fuzziness of human activity, and provide us all an opportunity to recover from error. Computer security systems are often oriented towards a specific culture. As a result, they tend to be rather rigid in their approaches to both what is permitted and how it may be done. We have seen some awesome demonstrations of the eventual fragility of rigid systems in Central and Eastern Europe in the recent past. A large, strong, rigid system can stand against the shifts of the underlying world for a long time, but you can't bend an inflexible rod, and eventually the strain becomes too great and the system shatters.

Efficient systems are dangerous to themselves and to others.

Many of us have fallen victim to a classic example of ruthless system efficiency: the *unix* command *rm* *. This command requires only five keystrokes (including the carriage-return) to eliminate all of one's files. Looseness, of course, is not the only attribute of a well-running system. Another is illustrated by a story once told of the Lawrence Livermore National Laboratory. It seems that someone had the bright idea of providing free bicycles for on-site use, in the hopes of reducing the strain on the intra-site bus and taxi services, saving fuel, reducing air pollution, improving muscle tone, and all those good things. So bicycles were provided, in sufficient numbers that one could just pick up the closest one, ride it to wherever one was going, leave it there for someone else, pick up another for the next journey, etc. Well, this sort of worked. The trouble was, the bicycles didn't stay randomly distributed around the Lab; a great preponderance seemed to come to rest at a certain building, that we shall call "Building N". They would be re-scattered periodically, but like the monarch butterflies to Pacific Grove, they would eventually find their way back to Building N. This remained a great mystery until someone

happened to be looking at a topographical map of the site, and noticed that the entrance to Building N was located at the lowest spot in the Laboratory. They had just discovered that

Systems run better when they run downhill. That is, they work better when they are aligned with human motivational vectors.

Ay, there's the rub. Security systems are rarely aligned with human motivational vectors. Not only—or even primarily—because people are malicious, but because, by and large, they don't care for bony potatoes.

To err on purpose is human.

Another reason why security systems fail to align well with human motivation is that people, especially people in the sorts of environments we serve, like to experiment: they push the envelope of the permitted. This means that everything possible will be tried, even when it is explicitly forbidden. Instruction and advice alike are often ignored. While this says something about people, systems, and documentation, it also points out the fact that

Time spent learning the system is perceived as time wasted.

The traditional complaint about documentation is that is too late, too obscure, and too inaccurate. While this may be true, the real truth is that even if it were on time, correct, and clear, it would not help matters a whole lot, because

Nobody reads the manual to prevent errors.

Documents are better reminders than teachers. This is as true of security documentation as of any other sort. People have tried to finesse the problem by providing various forms of on-line and automated assistance and advice. This, too, tends to fail, for a number of reasons, two of which are particularly instructive:

Automated advice fails because it assumes the user knows what he did.
Automated advice fails because it assumes the user will follow it.

In other words, most automated advice treats users as though they were computers: i.e., possessed of perfect recall and willing to follow instructions. Automated advice fails in the first case because it doesn't address the user who was hitting keys randomly to see what might happen, or the user who meant to do one thing but actually did another, or the user who was in a totally unfamiliar context. The second statement emphasizes that automated advice just doesn't cope well with people. The user might not be trying to do what the system thinks he had in mind, for instance, or may have decided to solve the problem a different way, or simply decided to do something else instead. One hopes the system will let him, but many systems do not.

This advice was *not* written under the assumption that the user will follow it. It was written under the assumption that we need frequent reminders about how the policies we develop in the ideal world of system design can conflict with the oddities of real people in real situations. When our attention is narrowly focused on our special task and discipline of computer security it is easy to forget the general truths presented here. We must keep them in mind, however, if we are to have any hope of producing systems that are both acceptably secure and reasonably useful.

Ernest Orlando Lawrence Berkeley National Laboratory
Technical and Electronic Information Department
University of California | Berkeley, California 94720