

CONFERENCE PROCEEDINGS

RAND

Research on Mitigating the Insider Threat to Information Systems – #2

*Robert H. Anderson, Thomas Bozek,
Tom Longstaff, Wayne Meitzler,
Michael Skroch, Ken Van Wyk*

National Defense Research Institute

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

DTIC QUALITY INSPECTED 4

DTIC QUALITY INSPECTED 4
20010125 034

CONFERENCE PROCEEDINGS

RAND

*Research on Mitigating
the Insider Threat to
Information Systems – #2*

*Proceedings of a Workshop Held
August, 2000*

*Robert H. Anderson, Thomas Bozek,
Tom Longstaff, Wayne Meitzler,
Michael Skroch, Ken Van Wyk*

National Defense Research Institute

CF-163-DARPA

The conference papers described in this report were supported by the Defense Advanced Research Projects Agency (DARPA) and RAND's National Defense Research Institute.

ISBN: 0-8330-2962-2

The RAND conference proceedings series makes it possible to publish conference papers and discussions quickly by forgoing formal review, editing, and reformatting. Proceedings may include materials as diverse as reproductions of briefing charts, talking points, or carefully written scientific papers. Citation and quotation is permitted, but it is advisable to check with authors before citing or quoting because of the informal nature of the material.

RAND is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND® is a registered trademark. RAND's publications do not necessarily reflect the opinions or policies of its research sponsors.

© Copyright 2000 RAND

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from RAND.

Published 2000 by RAND

1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

1200 South Hayes Street, Arlington, VA 22202-5050

RAND URL: <http://www.rand.org/>

To order RAND documents or to obtain additional information, contact Distribution Services: Telephone: (310) 451-7002; Fax: (310) 451-6915; Internet: order@rand.org

PREFACE

The insider threat to critical information systems is widely viewed as being of the greatest concern. On August 30 through September 1, 2000, approximately 40 researchers and U.S. government research sponsors met in a three-day workshop at RAND's Arlington VA facilities to address three aspects of this issue: (1) Create specific R&D challenges and goals over the next 2-5 years; (2) Discuss and develop insider threat models; and (3) Develop near-term solutions, focusing on tailoring commercial off-the-shelf (COTS) products for near-term effectiveness against the malicious insider.

This workshop was one in a series of related efforts to address the insider problem. In August, 1999, a workshop was held in Santa Monica, CA, focusing on R&D initiatives for preventing, detecting, and responding to insider misuse of critical defense information systems (Anderson, 1999). A major report on mitigating the insider threat to U.S. defense information systems has been produced by an integrated process team, and coordinated by OASD/C3I (DoD-IPT, 2000). In July, 2000, a workshop in Oahu, HI on advanced network defense research devoted a breakout session track to prioritizing the research initiatives proposed in that DoD report (Anderson, Brackney and Bozek, 2000). The results from these previous workshops were made available to participants in the current workshop, whose goal was to move beyond these efforts in creating specific research directions to guide researchers and government research managers.

The workshop was sponsored by the Office of the Assistant Secretary of Defense for C3I (Thomas Bozek, OASD/C3I) and the Defense Advanced Research Projects Agency (Michael Skroch, DARPA/IA&S).

These proceedings summarize the findings and recommendations resulting from this workshop.

For further information regarding the content of this document, please contact Thomas Bozek at OASD/C3I <tom.bozek@osd.mil>, Michael Skroch at DARPA <mskroch@darpa.mil>, or Robert H. Anderson at RAND <Robert_Anderson@rand.org>.

RAND support for this workshop was provided within the Acquisition and Technology Policy Center of RAND's National Defense Research Institute (NDRI). NDRI is a federally funded research and development center sponsored by the Office of the Secretary of Defense, the Joint Staff, the defense agencies, and the unified commands.

CONTENTS

Preface.....iii

Figures.....vii

Tables.....vii

Summary.....ix

 Long-term (2-5 year) research challenges and goals.....ix

 Developing insider threat models.....xii

 Near-term solutions.....xii

List of Symbols.....xv

1. Background.....1

 Previous Related Workshops and Activities.....1

 Concurrent Related Efforts at DARPA.....2

 Plenary Presentations and Background Reading.....3

 Focus Areas.....3

 Structure of the Remainder of This Document.....4

2. Long-Term (2-5 Yr.) Research Challenges and Goals.....5

 Survivable Architecture Frameworks.....6

 Differential access controls.....9

 Provenance.....14

 Mobile code.....16

 Summary of Recommended Projects and Costs.....17

 Other Topics Worthy of Consideration.....18

3. Insider Threat Models.....21

 Terminology.....21

 Questions Addressed by This Group.....21

 Purposes and Intended Users of Insider Threat Models.....21

 Modeling Requirements.....23

 Model Components and Attributes.....24

 Data25

 Time Frame.....25

 Related Work.....26

 Insider Threat Models -- Conclusions and Recommendations.....26

4. Near-Term Solutions.....27

 Install vendor-supplied security patches.....27

 Review and monitor existing event logs.....28

 Use existing access control.....30

 Employ configuration management - the ability to map your
 network/hardware/software31

 Filter malicious code at system choke-points.....32

 Filter for future and unknown malicious code, and exercise
 mitigation and containment32

 Track data pedigree and integrity.....33

 Other Suggestions.....34

5. Concluding Remarks.....35

Appendices

A. An Insider Threat Model for Model Adversaries.....	37
B. An Insider Threat Model for Adversary Simulation.....	41
Introduction.....	41
Goals and Purpose.....	41
Attributes of the <i>Insider</i>	42
Observations.....	46
Editorial.....	46
Next Steps.....	47
Conclusion.....	47
The Author.....	47
Reference.....	48
C. Modeling Behavior of the Cyber-Terrorist.....	49
Abstract.....	49
Fundamental Hypothesis.....	50
An Approach to Modeling the Cyber-Terrorist Adversary.....	50
Early Observations.....	53
Early Experiences.....	55
Future Directions.....	56
Summary.....	57
References.....	58
D. Can Technology Reduce the Insider Threat?.....	61
1. Authentication.....	62
2. Access control.....	63
3. Audit trails.....	63
4. Protecting the Information from Observation.....	64
5. Protecting Information from unauthorized "Edit/Copy/Paste"....	64
Conclusion.....	65
E. The Insider Threat to Information Systems.....	67
Research in Progress.....	70
The Critical Information Technology Insider.....	71
Indispensable Role of the Insider.....	75
Personal and Cultural Vulnerabilities.....	76
Impact of Intervention.....	83
References.....	87
F. The Insider Espionage Threat.....	89
Preconditions for Insider Betrayal.....	90
Opportunity.....	90
Motive.....	91
Reduced Inhibitions.....	92
Triggers.....	94
Summary and Conclusions.....	94
References.....	95
G. Insider Threat - A Theoretical Model.....	97
H. Information Assurance Cyberecology.....	105
I. Workshop Agenda.....	107
J. Workshop Participants.....	109
References.....	111

FIGURES

3.1. Required Model Components.....23
3.2. Model Framework.....24

TABLES

2.1. Summary of Recommended Projects and Costs.....18
3.1. Purposes of Insider Threat Models.....22
3.2. Intended Users of Insider Threat Models.....22

SUMMARY

The August, 2000 invitational workshop reported on here, held in Arlington, VA, is one of a series of recent activities sponsored by DARPA, NSA, DoD/C3I and other government agencies addressing the insider threat to critical Defense and U.S. infrastructure information systems. That threat is considered one of the most significant, with attributes distinguishing it from others - and therefore needing special attention.

This three-day workshop provided participants with a number of background reading papers and a set of plenary presentations by experts that recapped previous workshops and reports, and highlighted relevant research underway. Participants were chosen from academia, industry, government, and the U.S. military; some had attended previous workshops, for continuity, others provided new contacts and perspectives.

The workshop was organized around three focus areas, each the subject of a set of breakout sessions:

- Long-term (2-5 year) research challenges and goals toward mitigating the insider threat
- Developing insider threat models
- Near-term solutions using commercial off-the-shelf (COTS) and government off-the-shelf (GOTS) products.

The main workshop findings and recommendations in each of these areas are given below. The workshop agenda is available in Appendix A; Appendix B contains the list of participants.

LONG-TERM (2-5 YEAR) RESEARCH CHALLENGES AND GOALS

This discussion group focused on four research areas as most important and promising: survivable architecture frameworks; differential access controls; provenance; and mobile code.

Survivable architecture frameworks

The group concluded that the most important research initiative - not only for mitigating the insider threat, but for security of

information systems in general - must be the development of an underlying system architecture designed explicitly with security and survivability in mind (unlike essentially all operating systems and network architectures in use today). Without this, we are patching flaws and vulnerabilities that will emerge indefinitely.

Such a "survivable architecture framework" would have the properties of:

- pervasive authentication, traceback, and accountability
- adherence to good software engineering principles
- specifically addressing the "need to know" principle
- addresses reliability, security, survivability, and so on when confronted with insider threats
- addresses operational implications and needs.

Critical components of such an overall architecture were listed as:

- thin clients and trustworthy servers
- multi-level security
- protocol issues
- end-to-end security
- separation of duties and roles
- layering of functionality
- includes all aspects of hardware, software, and network architectures.

Details on these concepts and other parts of this group's deliberations are given in section 2 of this report.

Differential access controls

Differential access controls for an information system grant access based on the concept of "least privilege:" the minimum needed by a user to accomplish his or her task. Research on the development of such controls must consider at least three separate components: the network, the operating system, and different access control models to be employed.

Previous research relevant to this topic includes the Multics and Kerberos systems developed at MIT; the Digital Distributed System Security Architecture; Carnegie Mellon University's TMach operating system; and a number of other products and projects spanning the last 40 years.

Provenance

By "provenance" is meant the ability to retain an accountability trail for all modifications to critical system components, such as the network, system and user files, and application programs. Such a trail or log can be used to tie misuse to a perpetrator, and be used as forensic evidence.

Specific research areas listed as important to this topic are:

- The ability to create a non-forgable, non-bypassable, non-subvertible record of alterations and/or access to system objects and files, with appropriate granularity. (That is, there should be flexible control over whether this record applies to all "temp" files, backup copies, and so on.)
- Study of the longer-term implications and utility of advances in, and evolution of, steganography and watermarking
- Creation of tools for unwrapping and analyzing objects and documents with embedded encrypted information
- Development of utility programs implementing provenance markings, to be embedded within the system's file system
- Study of concurrency issues in a distributed system
- Multi-locking protocols for individual (atomic) transactions.

The group recommended two specific project areas, one studying means of embedding provenance information in software objects; the other investigating relevant watermarking, steganography, and "obfuscation" techniques.

Mobile code

Code in the form of "applets", viruses, worms, or macros attached to documents is increasingly migrating around information networks. Uses and abuses of mobile code will become more prevalent in the future.

The projects deemed necessary in this area include investigation of the following areas:

- Obtaining security for mobile code
- Proof-carrying code (i.e., code that carries along with it its own proof of its integrity, authenticity, and reliability)
- Self-validating code
- Tamper-proof code (so that you have reliance that what was sent was what was received and executed)
- Confined environments (in which "alien" code may safely be executed) - an extension of the notion of a "sandbox" made popular by Java.

The group outlined specific projects in each of the four major research areas listed above, listing estimated cost/year and expected duration of projects. The resulting total research program is substantial: totals for the four research areas designated as longer-term (2-5 years) were:

Survivable architecture frameworks	\$19.55M
Differential access controls	\$73.95M
Provenance	\$11M
Mobile code	\$32M

DEVELOPING INSIDER THREAT MODELS

If various research groups are to address the "insider threat," it is important that they have a well-defined model of the malicious insider and the level of risk or threat he or she poses. The availability of such models can generate guides and metrics against which various mitigation and security approaches can be tested in a controlled manner.

The group listed three critical components of such models:

- People, including their behavior, knowledge, and motivation;
- Tools, including software, hardware, and networks;
- The environment, including the organizational culture.

The basic structure for models includes four elements:

- *Observable*, allowing for measurable parameters
- *Profile*, applying to the environment, people, and tools, and defining a framework for each of them
- *Behavior*, which defines characteristics
- A function, labeled $F(x)$, that incorporates the model's functioning.

The group highlighted the need for a database of insider incidents that would be vital in creating and testing any such models.

There are several ongoing efforts relevant to the development of such models; they are listed in section 3 of this report.

NEAR-TERM SOLUTIONS

This group was tasked with focusing on solutions that might be implemented within the six to 12 months, primarily through exploitation

of existing commercial and government off-the-shelf software and systems.

The group discussed seven near-term solution areas:

- Install vendor-supplied security patches
- Review and monitor existing event logs
- Use existing access control
- Employ configuration management - the ability to map your network/hardware/software
- Filter malicious code at system choke-points
- Filter for future and unknown malicious code, and exercise mitigation and containment
- Track data pedigree and integrity.

A number of the above recommendations can *technically* be implemented quickly, using available software. For example, many systems within DoD and other critical infrastructure elements don't apply security "patches" available on the manufacturer's website in a timely manner; yet hackers know of the vulnerabilities being patched, and can exploit them before those patches are made. In such cases, the roadblock to implementation may be more procedural and administrative than technical.

- - -

The workshop provided an excellent venue for researchers, research managers, and potential users of security techniques to establish priorities, and share information on available data and ongoing research programs. In most of the topics discussed, specific research approaches were listed, relevant background work highlighted, and scope/time/cost parameters for recommended research estimated. As a result, these results can form the basis for a targeted research program - addressing both near-term and longer-term approaches - addressing the insider threat to information systems.

LIST OF SYMBOLS

Symbol	Definition
AFS/DFS	Andrew File System / Distributed File System
ATIAS	Advanced Technologies for Information Assurance and Survivability (a DARPA research program)
BAA	Broad Area Announcement
C3I	Command, Control, Communications and Intelligence
CERT/CC	Computer Emergency Response Team / Coordination Center (at Carnegie Mellon University)
COTS	Commercial off-the-shelf (software or hardware systems)
DARPA	Defense Advanced Research Projects Agency
DIA	Defense Intelligence Agency
DoD	(U.S.) Department of Defense
FFRDC	Federally Financed Research and Development Center
FTP	File transfer protocol
GOTS	Government off-the-shelf (software or hardware systems)
IA&S	Information assurance and survivability
IET	Insider estimated threat
IETF	Internet Engineering Task Force
IOS	Internetwork Operating System
IPv6	Internet Protocol version 6
MIT	Massachusetts Institute of Technology
MLS	Multi-level secure (system)
NDRI	National Defense Research Institute (an FFRDC at RAND)
NSA	National Security Agency
OASD/C3I	Office of the Assistant Secretary of Defense for C3I (Command, Control, Communications and Intelligence)
OS	Operating system
PKI	Public key infrastructure
PNNL	Pacific Northwest National Laboratory
SAM	Secure Administrator Monitor
SBIR	Small Business Innovation Research program
SIAM	Situational Influence Assessment Model
SNL	Sandia National Laboratory
USSPACECOM	U.S. Space Command
VPN	Virtual private network

1. BACKGROUND

The workshop reported on here is one in a series of activities focused on mitigating the insider threat to critical Defense and infrastructure information systems. Its purposes were to bring together leading researchers, government officials, and industry leaders; to identify the evolving need for tools, techniques and products aimed at the insider threat; to help assess existing relevant tools and techniques; to prioritize the R&D activities required; and to provide awareness of ongoing efforts and allow self-coordination among researchers and managers and supporters of that research.

The invitational workshop was designed to provide continuity with past efforts by including some workshop attendees from two previous workshops (held in August, 1999 and July, 2000), combined with new participants from academia, industry, government, and the military.¹

PREVIOUS RELATED WORKSHOPS AND ACTIVITIES

Three specific earlier activities set the stage for the present workshop:

- A three-day workshop on "research and development initiatives focused on preventing, detecting, and responding to insider misuse of critical defense information systems," held August 16-19, 1999 in Santa Monica, CA. That workshop was primarily technology focused, and concentrated (as its title implies) on activities aimed at identifying the threat, and preventing, detecting, and responding to malicious insider activities. It included some auxiliary policy recommendations that were felt to be needed if the technology solutions were to be effective. Important themes from that workshop highlighted by Michael Skroch, in his introductory remarks, were:

- Identification of what is the insider, and what is malicious (i.e., being able to identify the threat;

¹ Much of the content of this section is taken from opening remarks and vugraph slides presented by workshop organizer Michael Skroch of DARPA.

- The need for finer-grained control of access, authentication, attribution, integrity, and so on. This needed to enable better prevention, detection, and response;
- The need for a trusted path with high level of assurance from the security system to personnel, or to application programs;
- Development of courses of action, and having the capability to carry out effective response to incident;
- The importance of policy definition and enforcement;
- Questioning the difference between insider personnel and code that exists (or has been introduced) inside the system;
- An understanding that the human element is the biggest unknown.

That workshop used as source material an early draft of a major DoD report on mitigating the insider threat (see next bulleted item, below). The August 1999 workshop is available in both hardcopy and online versions as Anderson (1999).

- *"DoD Insider Threat Mitigation Plan: Final Report of the Insider Threat Integrated Process Team* - a report coordinated by Tom Bozek of OASD/C3I. This report lists over fifty specific recommendations (both short-term and long-term) to be implemented regarding the insider threat. See DoD-IPT (2000).

- A workshop on *"Advanced Network Defense Research: Focusing Current Research on User Needs,"* held in Oahu, HI in July, 2000. A breakout session within that workshop focused on prioritizing the key recommendations in the above DoD-IPT report. This results of that workshop are documented in Anderson, Brackney and Bozek (2000).

All of the above documents were made available as prior reading material to participants in the current workshop.

CONCURRENT RELATED EFFORTS AT DARPA

After the August 1999 workshop, cited above, several seedling efforts were initiated at DARPA. They include:

- an SBIR program on "active profiling for insider threat." DARPA received 14 responses to this announcement; at the time of this writing (October 2000) decisions are pending;

- ATIAS BAA (broad area announcement) on "insider threat active profiling." Twenty-six responses were received; two contracts have been awarded as of October 2000, with a third pending.

PLENARY PRESENTATIONS AND BACKGROUND READING

The workshop's initial plenary session provided recaps of the three precursor workshops and documents listed earlier, plus the following invited presentations:

- "Insider Threats to Critical Information Systems: Typology of Perpetrators," by Dr. Jerrold Post and Dr. Eric Shaw of Political Psychology Associates, Inc. (See Appendix E for a related short paper by Shaw, Ruby and Post.)

- "Can Technology Reduce the Insider Threat?," by Michael Caloyannides and Carl Landwehr of Mitretek Systems (Appendix D).

At a later plenary session, Bill Leonard, Deputy Assistant Secretary of Defense for Security & Information Operations, provided remarks and engaged in a question/answer session with workshop participants.

In addition to documents and presentations described above, the following reading materials were provided to attendees:

- "An Insider Threat Model for Adversary Simulation," by Bradley J. Wood, SRI International (Appendix B)

- "Modeling Behavior of the Cyber-Terrorist," by Gregg Schudel, GTE/BBN Technologies and Bradley J. Wood, SRI International (Appendix C)

- "An Insider Threat Model for Red Teams," by Bradley J. Wood, SRI International (vugraph presentation, in Appendix A).

FOCUS AREAS

The workshop was structured around three focus areas:

1. *Specific longer-term R&D challenges and goals.* From earlier workshops and recommendations, the following three subjects were highlighted as worthy of consideration: (1) creation of tamperproof audit trails, unavailable to the insider; (2) watermarking of documents, unavailable to the insider; (3) the ability to identify critical information automatically. Participants were asked to start with a broader list, then focus on the top few items. They were to consider

the evolution of systems and technology, so that solutions proposed don't become obsolete in the foreseeable future.

2. *Insider threat models.* What types of models should be developed to assist in real-time control, assessment, experimentation, and design of tools to mitigate the insider threat? This group was asked to create guides and metrics against which various insider threat mitigation approaches can be tested in a controlled manner. Among the desired outcomes for this group was the production of one or more insider threat model definitions as examples, elaborated as time permits.

3. *Near-term solutions using COTS and GOTS products.* The insider threat is a reality today; near-term solutions are needed. This group was asked to consider tangible modifications to, and combinations of existing, widely-used COTS software products that can be effective in detecting and reacting to the malicious insider. Among the scenarios to be considered by this group were the commercial environment, the government environment (e.g., including classified data and national-level threats), and government+military needs (e.g., in tactical situations).

STRUCTURE OF THE REMAINDER OF THIS DOCUMENT

The following three sections describe the results of the three focus groups listed above. These are followed by brief concluding remarks, and appendices containing the agenda and list of attendees for the workshop.

2. LONG-TERM (2-5 YR.) RESEARCH CHALLENGES AND GOALS

This focus group's task was to develop what they considered to be the most important specific research proposals, to be carried out over a two- to five-year period, addressing the insider threat.

To accomplish this, they first generated topic areas in five categories: fundamental research, prevention, deterrence, detection, and response. As a result, 15 possible research topics were suggested and discussed. The group then voted on the most important ones, using the following rationale:

- The selections were based on choosing a portfolio of three or four research areas that will make dramatic improvements to our ability to cope with the malicious insider problem, with primary emphasis on long term;

- The portfolio of research topics comprises a comprehensive and highly synergistic set of solutions for the insider threat problem.

Other approaches identified were listed at the conclusion of this group's presentation, so that they might be given consideration in other venues.

As a result of this process, the "top four" research projects chosen for detailed examination were:

- Survivable architecture frameworks
- Differential access controls
- Provenance
- Mobile code

For each of these topics (described in more detail below), the group listed: its rationale for choosing this research topic, the expected outcome of this research, other factors to be considered, success criteria, key background work that forms a basis for this work, the recommended approach, scope, time, cost, and the specific applicability of this research to the insider threat.

Each of the four recommended research topics is described below.

SURVIVABLE ARCHITECTURE FRAMEWORKS

This focus group felt that *the most important* research initiative they could investigate and describe was what they called "survivable architecture frameworks." By this is meant creating an underlying system architecture that is designed explicitly with security and survivability in mind. Without such an architecture, there will always be flaws and vulnerabilities that can be exploited in any complex, multi-user, distributed information system. It is the basis for a long-term solution to security, integrity, and reliability in systems.

Baseline architectures

Key "baseline" architectures that would be critical components of such an overall architecture were listed by the group:

- *Thin clients and trustworthy servers.* "Thin" client computers contain only basic software that allow them to connect to a server and obtain needed applications and data from that server. As such, their code might be sufficiently simple that it could be "burned into" read-only memory (ROM), and not subject to malevolent modification. It might also be sufficiently simple that its operation could be understood to a level of detail that security flaws or vulnerabilities might be detected and corrected prior to its "hardening." The servers must of course be trustworthy sources for receiving applications and data.

- *Multi-level security.* The system must be capable of handling multiple levels of security in a manner that restricts a user's access to those levels for which he or she has been granted privileges.

- *Protocol issues.* There are important decisions to be made regarding which communication and other protocols will be used within the survivable architecture. Will "open" protocols (i.e., non-proprietary) be used throughout? Can existing protocols suffice, or are new ones needed?

- *End-to-end security.* The architectural design must explicitly consider all components, systems, and network in a transmission between users "end to end."

- *Separation of duties and roles.* Access authorization should be based on the duties and roles of a user, and those should be consciously

separated to the extent possible. (For example, a "sysadmin" role should not include a "security officer" role, so that those separate functions can be carried out in a "checks and balances" manner by two separate people.

- *Layering of functionality* (i.e., protection between objects and levels of abstraction). Specific levels of abstraction should be allowed and encouraged for system objects, with differing levels of protection at the various levels.

- *Hardware, software, and network architectures*. The system architecture must include all aspects of hardware, software, and network components within its purview.

Properties of survivable architecture frameworks

The necessary properties of a survivable architecture design include:

- pervasive authentication, traceback, and accountability
- adherence to good software engineering principles
- specifically addresses the "need to know" principle
- addresses reliability, security, survivability, and so on when confronted with insider threats
- addresses operational implications and needs.

Key background work

It was felt by the breakout group that there was much relevant previous research to be drawn upon - too numerous to be listed in the limited time available to the group at this workshop. It includes the pioneering work on the Multics operating system at MIT in the 1960s, and many other developments. Any group undertaking the proposed "survivable architectures" research should be thoroughly familiar with security architecture research developments of the past 40 years.

Approach

The breakout group recommended the following general approach to this research task:

- Create an Architecture Steering Group. It should choose the best alternatives, and elaborate them with regard specifically to addressing insider threats.

- Leverage open source (systems or concepts) by specifying that all research projects addressing this task be open source compatible.

- Make incremental developments available to other related research efforts. Don't wait years and then produce a "there it is, take it!" finale.

- Perform this task in parallel with other related research efforts.

- Create well-documented, specified, interoperable, reusable architectures.

- Start with a summer study, and define how organizations and projects will interact.

- Act as an integrator for related, more specific, research programs.

- Address problems of wireless connectivity as part of any legitimate framework.

Scope/time/cost

The group proposed a five-year research project designed to produce incremental results during that period. Expenditures were listed as:

- A Steering Group (7 to 10 people) meeting once each quarter for five years, providing broad policy and guidance on promising research approaches to the prime contractor -- \$450K/year.

- Support for the Summer Study to launch this effort -- \$2M.

- Contractor support: \$1M/year per subcontract.

- Support for a "deployment symposium" scheduled for year 2: \$300K.

Applicability to insider threat

The primary reasons given that this "survivable architecture frameworks" research program addresses the insider threat were:

- This framework assures that all other research programs address the insider threat, because it embodies solutions within the fundamental system architecture;

- It is not just relevant to the insider threat, but more broadly provides an approach to survivable, reliable systems for the future;
- It sets the agenda for addressing insider threats beyond the five years of this particular research program.

DIFFERENTIAL ACCESS CONTROLS

The second longer-term research program described by this breakout group involved differential access controls. By this is meant controls on access to an information system based on the principle of "least privilege" - granting the minimum access needed by a user to accomplish his or her task. Such access should be fine-grained, and apply to all system resources accessible within the network.

A factor to be considered in this research area is time scale: In the shorter term, one might prohibit access to such network resources as ActiveX and JavaScript; in the longer term, more specific access controls should be extended to all network resources.

The success criterion proposed for this project was: Within five years, have an operational ability to constrain system administrators and other insiders integrated into a major open source operating system.

The discussion of differential access controls research considered three separate components: the network, the operating system (OS), and access control models. Each are described separately below.

Network access controls

The key features desired for network access controls were listed as follows:

- They are distributed to all network elements
- A language is needed for specifying the policy regarding protocols, services, and roles. It must be dynamic and incremental to address new technology. Real-time analysis tools are needed to view and check for consistency.
- That descriptive language for specifying access control policies should be executable, to avoid a transcription step.
- Network access controls should be platform-independent
- Access controls must also be provided for users on servers
- A Kerberos-like capability should be focused on network services.

Key background work

Three prior research and development efforts were listed as background for the proposed research: Kerberos, the Digital Distributed System Security Architecture, and AFS/DFS.

Approach

One approach to network access controls is to extend the access and security features of IPv6 (version 6 of the Internet Protocol), perhaps producing an IPv7. The research should use open source components, including those for IOS, network stacks, and other devices.

Scope/time/cost

It was felt by the group that a five-year multi-organization program was needed, comprising four or five separate but cooperating research projects. These efforts would result in a working prototype. The resources would be \$3M per organization per year. It was noted that this effort requires the involvement of the Internet Engineering Task Force (IETF). It was hoped that the effort would result in early insights on how to differentiate roles and responsibilities, upon which differential access would be based.

Applicability to the insider threat

The importance of differential network access controls to the insider threat was highlighted by the following points:

- It permits identification and authentication of individual users for network objects and traffic, permitting network-wide differential access control;
- The network audit trail is pinned to a specific user
- It allows the system to distinguish between authentication and authorization for network services;
- It allows for the creation of zones of control (e.g., based on mission).

Operating system access controls

To implement differential access controls, it was felt that a research program was needed at the operating system (OS) level, to complement the work at the network level described above.

A restructured operating system (different from existing popular COTS systems) is needed to:

- support separation of duties among users;
- provide security beneath a less-secure OS; for example, by treating Microsoft Windows as an application rather than the basis for security;
- create the basis for a security architecture for servers, firewalls, gateways, routers, and other key hosts and nodes in a network;
- Provide audit trails that can reliably track the actions of sysadmins.

The architecture within which this next-generation more-secure OS might fit is one having one "thin client" system (hardwired, unchangeable) per user, communicating with a trusted server. Sysadmins would have access only to the server, and mobile code would be executed on trusted servers rather than on the client machines.

Key background work

The systems listed by the group as being important precursors to secure operating system R&D of the type proposed here are MIT's Multics, and Carnegie Mellon University's development of TMach.

Approach

Suggested approaches for the proposed secure operating system R&D were to:

- extend Linux/BSD to add or modify properties, functionality, and assurance;
- retain an open source philosophy;
- as mentioned above, consider an architecture based on thin clients and trustworthy servers.

Scope/time/cost

The breakout group's estimate of the resources and time required for the proposed research was based on two or three institutions working for two to four years at \$800K/year per organization. It was expected that this effort would result in a viable operating system ready to deliver to DoD and others at the conclusion of that period.

It is likely that the resulting system would be lower cost to DoD than current systems, because of the use of an open source OS as the starting point. If, as is hoped, the resulting system would be adopted by the commercial market (which is of course in need of better security also), DoD would be able to buy COTS support for the resulting system.

Applicability to insider threat

Only through the development of more secure operating systems can the following advantages related to the insider threat be obtained:

- an opportunity to catch a malevolent sysadmin;
- it removes the critical vulnerability of universal privileges now extended to the "root" user;
- the proposed R&D would provide lots of intermediate results and incrementally improved systems with regard to insider activity;
- the resulting system would, for the first time, allow the creation of protected, non-bypassable, non-alterable audit trails of a sysadmin's activity.

Differential access control models

The breakout session discussion on survivable architecture frameworks also considered various models for providing differential access control for users, depending on their role or activities.

Differing approaches considered were:

- Multi-level security (MLS) controls (confidentiality models). These would be needed for an MLS network. They would not, however, be needed for "system high" thin clients. Such controls would be a means of implementing "need to know" back into the networks.
- Differential integrity models.

- Security models specifically focused on the insider threat. For example, one should consider providing insider profiling directly within the operating system, not as an add-on application.

Key background work

Relevant background work mentioned include work by Clark and Wilson, in application integrity; Biba (a multi-level integrity architecture); SeaView (providing MLS database management); and object-oriented capability-based systems.

Approach

The proposed approach for this research initiative is to:

- Select and develop modifications to a widely-adopted open source operating system;
- Add the ability to enforce separation of duty policies;
- Integrate these developments into the secure operating system being developed under the research program described above.

Scope/time/cost

The breakout session felt that the proposed research would best be undertaken by an integrated team, comprising a maximum of three organizations at \$750K/year per organization, for two to four years. This team would integrate the developed models into the open source OS systems being developed independently, making them available to DoD and others with many intermediate deliverables. This research effort would also be coordinated with the project on network access controls.

Applicability to insider threat

The proposed research allows the implementation of fine-grained control of roles and responsibilities in open source operating systems. This is particularly important for controlling the authorities granted to sysadmins, but is also relevant to other users.

PROVENANCE

The dictionary definition of "provenance" is "origin, source." Used in the context of networked information systems, the term refers to an accountability trail for all modifications to critical system components, such as network, system and user files, and application programs. With such a trail, misuse can be tracked to a perpetrator. Provenance information about system objects or documents is also useful in providing records for intellectual property. Methods sometimes employed to record provenance include steganography and digital "watermarking" of documents.

Specific research areas

Specific research areas related to providing effective provenance for system objects include:

- The ability to create a non-forgable, non-bypassable, non-subvertable record of alterations and/or access to system objects and files, with appropriate granularity. (That is, there should be flexible control over whether this record applies to all "temp" files, backup copies, and so on.)
- Study of the longer-term implications and utility of advances in, and evolution of, steganography and watermarking;
- Creation of tools for unwrapping and analyzing objects and documents with embedded encrypted information;
- Development of utility programs implementing provenance markings, to be embedded within the system's file system;
- Study of concurrency issues in a distributed system;
- Multi-locking protocols for individual (atomic) transactions.

Key background work

The related background research relevant to this proposed task were listed as: the "Plan 9" operating system development at AT&T research; research on "atomic" transactions and concurrency; steganography and watermarking; obfuscation research; immutable file systems; cryptography, including public key infrastructures (PKI) and related public key systems; work in audit logging; and research on establishing trusted paths within information systems.

Approach

The group proposed two projects, each utilizing applicable elements emerging from the "architectural frameworks" research described earlier in this section:

Project 1: Provenance embedded in objects or stored in associated objects. The results of this work must be applicable in operating systems, applications, and networks.

Project 2: Watermarking, steganography, and obfuscation approaches. This research thrust does not rely on the surrounding environment, but - as above - must be applicable in operating systems, applications, and networks.

Each project should produce incremental deliverables during their course, and should be designed so their results may be integrated within open source frameworks.

Scope/time/cost

Project 1: The research could be performed by one group, funded at \$3M/year for three years. The first year of the project should be devoted to a look at all available techniques, disqualifying some. The second year would concentrate on ways of integrating these techniques with other approaches to information security.

Project 2: This would be performed by two groups, each funded at \$500K/year for two to three years. It would be expected that this project could show the significance and applicability of further research into steganography and watermarking by the end of the first year.

Applicability to insider threat

The proposed research thrusts address the handling and export of sensitive information. The research aims at providing significant accountability for all insiders. It also provides forensic evidence for prosecution of insiders, which is often not possible today. Such prosecution has a high deterrence value for other insiders.

MOBILE CODE

Information and computing environments in the future will increasingly rely on "mobile code" - application programs or applets that migrate from server to client, or that are sent and received attached to e-mail messages. They might be pushed to users, or pulled by them via a browser. A particularly prevalent, malicious form of mobile code is that which encodes computer viruses and worms.

It was felt by the discussion group that important research areas requiring attention in this regard are duals of each other: protecting mobile code from malicious environments; and protecting environments from malicious mobile code.

Projects that were deemed important related to mobile code were:

- Obtaining security for mobile code;
- Proof-carrying code (i.e., code that carries along with it its own proof of its integrity, authenticity, and reliability);
- Self-validating code;
- Tamper-proof code (so that you have reliance that what was sent was what was received and executed);
- Confined environments (in which "alien" code may safely be executed) - an extension of the notion of a "sandbox" made popular by Java.

Key background work

The background work deemed most relevant to research involving trusted mobile code was that dealing with: "type safety," confinement and restricting of privileges; in-line reference monitor(s) and wrappers; Java Virtual Machine and other interpreters; research on obfuscation; distributed systems protocols; proof-carrying code; self-validating code; and tamper-resistant code.

Approach

The group felt that an approach to this problem should embody the following principles:

- Create a steering group to structure the inquiry;

- If DARPA creates a research program on mobile code, this proposed work is essential to its success, and should be closely integrated with that program;

- This research must ensure a compatibility of approach with the other research tasks described in this section: survivable architecture frameworks; differential access controls; and provenance.

Scope/time/cost

It was felt that \$400K/year should be set aside for constituting the senior steering group (approx. 10 persons) meeting one week each quarter, for five years. This group should produce a research agenda during the first year, and assure that they are cross-connected with the Architecture Steering Group recommended above.

This effort should also allow for six to 10 individual and team research projects, for a total of \$6M/year for five years.

SUMMARY OF RECOMMENDED PROJECTS AND COSTS

Table 2.1 summarizes the projects and costs recommended above, comprising a serious attack on the problem of creating trustworthy, survivable information systems. It is felt that the breadth and depth of the proposed research is vital for building a basis on which both the insider and external threats to information systems can be resolved to a necessary degree.

Table 2.1
Summary of Recommended Projects and Costs

<i>Project</i>	<i>Cost/Year</i>	<i># of Years</i>	<i>Total Cost</i>
Survivable Architecture Frameworks			
Steering group	\$450K	5	
Summer Study (year 1)	\$2M	1	
Contractor support (\$1M x 3 contracts)	\$3M	5	
Deployment symposium (year 2)	\$300K	1	\$19.55M
Differential Access Controls			
Network access controls (4 projects at \$3M/year ea.)	\$12M	5	
Operating system (3 projects at \$800K/yr ea.)	\$2.4M	3	
Access control models (3 projects at \$750K/yr ea.)	\$2.25M	3	\$73.95M
Provenance			
Project 1	\$3M	3	
Project 2 (2 projects at \$500K/yr)	\$1M	2	\$11M
Mobile Code			
Steering group	\$400K	5	
Individual projects	\$6M	5	\$32M

OTHER TOPICS WORTHY OF CONSIDERATION

We listed above the four top-rated long-term (2-5 year) research projects deemed most essential for laying a proper basis for combating the insider (as well as external) threat to distributed, networked information systems.

However, the group felt that the following topics were also worthy of consideration. Research sponsors should consider projects in these areas in addition to those itemized above:

- Inspection. We need better code inspection, with automated tool support for such inspections. Aspects of this research include interpretation of code, response, and correlation.
- Anomaly and misuse detection. Such systems are in use today, but they need to accommodate insider misuse specifically. They need to address questions such as: How do you practice early detection (of insiders)? What is early detection? How do you perform notification and response? How can such systems provide autonomic response to the insider threat?

- Analysis of human factors, and psychological profiling. How can these tools be brought to bear for early detection of the insider "going bad?"

3. INSIDER THREAT MODELS

If one is to conduct research on countering or mitigating the effects of the "malicious insider," it is important to have insider threat models to generate guides and metrics against which various mitigation approaches can be tested in a controlled manner. This breakout session addressed the problem of defining the basis, need, and uses for models of the malicious insider.

TERMINOLOGY

For the purposes of this group, the following definitions were used:

- *Insider*: Any authorized user who performs unauthorized actions. Examples include users, privileged users, sysadmins, network administrators, facility support personnel, temporary employees, contractors.
- *Insider Threat*: Any authorized user who performs unauthorized actions that result in loss of control of computational assets.
- *Model*: An approximation of reality.

QUESTIONS ADDRESSED BY THIS GROUP

Among the questions addressed by the group were:

- What processes need a model? Should there be more than one type of model? Is there a spectrum of insider threats? What comprises a model? What will models do? Why do we need them? Who are consumers or users of models? What models now exist? What are the observables? What motives should we consider?
- Can or should we define a taxonomy of insider activities, a list of opportunities an insider needs, a taxonomy of motives, etc.?

PURPOSES AND INTENDED USERS OF INSIDER THREAT MODELS

This breakout session listed a wide variety of purposes to which an insider threat model might be put. They are listed in Table 3.1.

Table 3.1
Purposes of Insider Threat Models

Education and training
Studies of information systems
Standardized frame of reference
Support development of tools: predict, detect, respond
Indications and warning development
Creation of behavioral theory
Simplification
Simulation
Personnel management
Predict intent and behavior
Develop defensive courses of action (COAs)
Policy development
Hypotheses generation
Acquire new knowledge

In conjunction with the above table, the group also created a list of intended users of models or model results; they are shown in Table 3.2.

Table 3.2
Intended Users of Insider Threat Models

Red teams / experiment and test designers
Human resources
Researchers
System / software designers and developers
System & network administrators
Operators and planners (for both offensive and defensive IO)
Security analysts
Policy / decision makers
Counter-intelligence/law enforcement/intelligence personnel
Awareness trainers and educators

The functions that a model might perform for the above purposes and users were listed as: analysis, characterization, detection, risk assessment, prediction, prevention, and operational response.

MODELING REQUIREMENTS

The modeling requirements for the research and development, and operational communities, differ drastically. Likewise, requirements for macroscopic (high level) vs. microscopic (focused and detailed) models can vary; thus a suite of models will be required. The group referred to this as a "confederation of insider models." Defining common interfaces or elements for the various models would be very beneficial and allow some models to leverage others.

The group concluded that a comprehensive model of the insider threat needs a people component, tools component, and environment component as depicted Fig. 3.1.




Figure 3.1 - Required Model Components

The group identified a basic structure for the models that included four primary elements. The *observable* element provides for measurable

parameters. The *profile* element, which applies to the environment, people, and tools, and provides for defining a framework for each of them. The *behavior* element defines characteristics. The $F(x)$ provides for model functions. Inadequate time was available for the group to expand on the model functions. A high level diagram for the model structure is shown in Fig. 3.2.

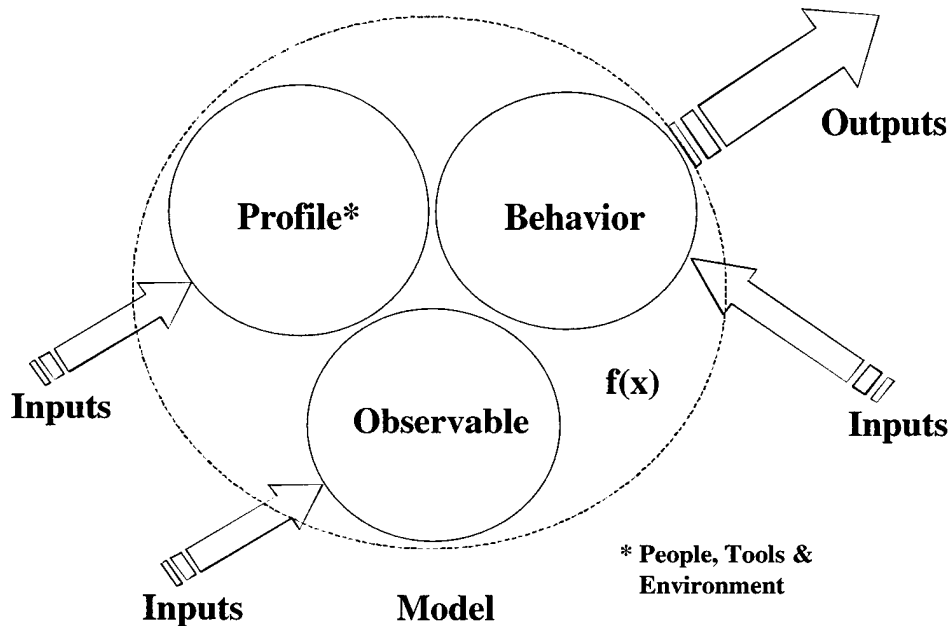


Figure 3.2 - Model Framework

MODEL COMPONENTS AND ATTRIBUTES

The group concluded that a comprehensive Insider Threat model would require as a minimum three major components. The components are People, Tools, and Environment. The People component would include parameters such as human performance, behavior, knowledge, and motivation. The Tools component would include all processes, automated or manual; procedures; data; computers; and other devices that people would use to address or solve the Insider Threat problem. The Environment component would include those parameters that apply collectively to both People and Tools or convey relationships between them. Figure 3.1 depicts graphically these Model components.

The group identified a basic model framework that included four primary elements. Any model framework element may contain one or more of the model components described in the previous paragraph. The observable element provides for measurable parameters. The profile elements, which apply to all four model components (people, tools, environment), provide for defining a framework or structure for each of them. The Behavior element defines characteristics, attributes, and relationships. The F(x) provides for model functions. Unfortunately, inadequate time was available for the group to expand on the model functions. Figure 3.2 provides a graphic depiction of the Model Framework.

The group estimates that the collective set of models to effectively address the insider threat will have the components and framework estimated above. Their implementation, application, and individual level of detail may vary drastically.

DATA

Little data exists today specific to the Insider Threat. What data exists is spread throughout many organizations and usually has distribution limitations preventing its use throughout the research community. Furthermore, the depth and bread of such data is inconsistent. Reference control data is needed for evaluating research progress and results. Such data is essential to developing and validating Insider Threat models. The only Insider Threat dataset cited during the workshop was one under development at Defense Security Research Center that contained approximately 25 cases; this dataset is expected to be made available to members of the workshop. Since Insider Threat data is essential to model validation, currently not available, and of immediate need to many of the workshop participants, the working group rated the development of such a dataset as a high priority for the successful development of Insider Threat Models.

TIME FRAME

While increased understanding can be gained during the Insider Threat Model development process, the full effect of such models is not

expected to be realized for a period of two to five years after development of such models is initiated.

RELATED WORK

The workshop session included introductions by each participant along with a few words of related activities. Ongoing efforts included modeling for red teams being performed by SNL and SRI (see Appendices A and G), the Secure Administrator Monitor (SAM) being developed at PNNL, the PPA Model being developed at Political Psychology Associates, Inc, the NSA Model (based on SIAM (Situational Influence Assessment Model)), and the IET Model and the insider dataset being developed by Defense Security Research Center (see Appendix H).

INSIDER THREAT MODELS -- CONCLUSIONS AND RECOMMENDATIONS

It was felt by the group that computer models for the Insider Threat could contribute to understanding this problem and subsequently build the foundation for future tools that ultimately may provide a solution.

The following recommendations were identified to allow Insider Threat Models to contribute to the defense and prevention of insider computer and network compromises:

- Develop a complete insider threat taxonomy
- Develop a plan to define and acquire insider data
- Develop a plan and strategy for a "confederation of insider models" that can work together synergistically. This requires a focused group of experts guiding this process.
- Develop metrics for success and quantify observables.

4. NEAR-TERM SOLUTIONS

The third of the three breakout sessions at this workshop focused on "near-term solutions" to the insider threat problem - ones that might be implemented within six to 12 months. Their aim was to use "low hanging fruit," primarily through exploitation of existing commercial and government off-the-shelf software (COTS and GOTS).

As a starting point for discussion, this group reviewed the suggestions related to mitigating the insider threat resulting from a July, 2000 workshop held in Oahu², which focused one breakout session on that same topic.

The group could not, in the limited time available, create solutions that were complete and comprehensive. They listed a set of factors that should be considered in any more thorough study of near-term solutions: training issues, separation of duties, acceptability, practicability and deployability, cost of deployment and maintenance, scalability, and sustainability.

With those caveats, the group concentrated on seven near-term solution areas:

- Install vendor-supplied security patches
- Review and monitor existing event logs
- Use existing access control
- Employ configuration management - the ability to map your network/hardware/software
 - Filter malicious code at system choke-points
 - Filter for future and unknown malicious code, and exercise mitigation and containment
 - Track data pedigree and integrity.

The following sections discuss each of these topics in more detail.

INSTALL VENDOR-SUPPLIED SECURITY PATCHES

In widely-used computer operating systems and application programs, there are literally hundreds of known flaws and vulnerabilities that are

² Anderson, Brackney, Bozek (2000)

widely posted on hacker bulletin boards and web sites, with new ones posted each day. For many of these vulnerabilities, "patches" that fix them are posted by the software manufacturers, or by such organizations as the Computer Emergency Response Team Coordinating Center (CERT/CC) at Carnegie Mellon University's Software Engineering Institute.

Unfortunately, many existing security patches aren't installed by system operators. The timely installation of such patches as soon as they become available is perhaps the single greatest "near-term" solution strategy that can be employed, and at relatively low cost compared with other solutions.

The direct impact of this solution on the insider problem is that it prevents exploits of known vulnerabilities, and will help prevent "privilege escalation" attacks by insiders.

Any such timely patch installation should be included in a robust countermeasures program that entails attributes and issues such as: assuring that such patches are obtained from a trusted source; operations testing of patches before installation; checks and balances during the installation process; support to system administrators in determining the priorities of various patches and fixes; incorporating the validation of the installation of patches within operational readiness reporting; and top-down validation of patch installation on mission critical systems.

A longer-term program should include research and implementation of tools for enterprise patch installation, and decision support systems for prioritization of patch installation based on mission criticality.

It was noted that a test-bed environment is highly desirable for testing patch implementation and distribution, and assessing the operational impact of any such system modifications.

REVIEW AND MONITOR EXISTING EVENT LOGS

Existing computer operating systems often have some facilities for logging various kinds of system events as they occur. These logs are not perfect: intruders or insiders might tamper with and alter the logs to hide malicious activity; they may be incomplete in some respects. But they are certainly better than nothing. Such event logs should

certainly be "turned on" in order to provide indicators of malicious activity to the extent they can. And the resulting logs need to be monitored by systems or security personnel for anomalous entries. Too often some logs are being kept by the system, but they are not reviewed by cognizant systems personnel.

Among the actions that should be taken now in this regard are: (1) configure existing systems to turn on a reasonable level of event logging; (2) divert and concentrate event log streams into hardened monitoring consoles; (3) reduce and analyze event streams for suspect activity; (4) create baselines of typical activity at the micro and macro level, then compare logs to this baseline; (5) provide training support for event log analysis; (6) mandate a dual review of event logs (e.g., by a sysadmin and an organization security official); and (7) tie event logging with the alerting and response process.

A longer-term research program in this area would include the following areas of activity:

- conduct research on what constitute reasonable levels of event logging;
- develop templates or profiles of malicious activity to support automated recognition systems;
- identify and develop tools to aid in the analysis of event streams;
- identify and develop means for making event log streams tamper resistant;
- identify and develop data mining tools to support trend analysis, norms, and templates;
- identify deviations from established profiles;
- establish event log data retention and storage policies.

More substantial (i.e., longer-term) research strategies would include these topics:

- continue to develop better recognition systems;
- develop means for correlation of heterogeneous event streams for trend analysis;
- develop data correlation tools to extend existing automated tools and information displays;

- determine data pruning needs in the area of event log data reduction and analysis.

In pursuing these recommendations, it is important to also address issues of classification and privacy in the collection and aggregation of event data within an organization's information system.

USE EXISTING ACCESS CONTROL

Current operating systems provide various mechanisms for protecting data and restricting access within their systems. Proper use of these existing access controls can inhibit and contain unauthorized insider access to resources. Use of these capabilities also facilitates maintaining an audit trail of access, and inhibits adversaries' mapping (i.e., surveying all system assets) capabilities within the system.

Among the activities that can be instituted now are these:

- set up and implement file access controls at network choke-points (e.g. at firewalls, routers);
- institute and use personal firewalls;
- enable individual file sharing by exception;
- implement appropriate tailored system/work unit/project access control policies;
- implement automated tools for setting up controls;
- audit file access violations;
- establish a process for granting and revoking system/work unit/project level access;
- use a virtual private network (VPN) to compartmentalize data;
- implement a procedure to be used as needed for removal of employees with high degrees of access (e.g., that doesn't allow them to retain such privileges after they have been notified that they are to be terminated).

If a longer-term research program were established in this area, it should include activities to:

- implement fine-grained access control;
- develop role-based access controls;
- produce and disseminate dynamic alert/warning banners;
- develop operating system-specific tools for controlling access.

Even longer-range research in the area would investigate means of ensuring mandatory access control, and development of multi-level secure (MLS) operating systems and networks.

EMPLOY CONFIGURATION MANAGEMENT - THE ABILITY TO MAP YOUR NETWORK/HARDWARE/SOFTWARE

In order to know if some "alien" host has joined your organization's network or that some new communication line has been established or disabled, it is important to know - dynamically - the current architecture of your system: what is there, and how it is connected. By this means, one can detect a wide range of insider misuse, identify current and future configuration exploits, and identify potential vulnerabilities.

Among the near-term steps that can be taken to enhance the security of existing systems in this regard are these:

- provide currently available mapping software (from vetted freeware to very expensive packages);
- employ existing scanning tools that are readily available, for network mapping, protocol/address monitoring, phone scanning, and application monitoring;
- obtain approval at appropriate DoD levels and develop trusted delivery mechanisms to implement these procedures;
- develop policy and process to reflect operational readiness and system baseline configurations, and perform validations;
- provide training support to the system administrator in using these tools and capabilities.

A longer-term research program would include such additional activities as:

- automate the system and network scans,
- develop tools for integrating, rationalizing and reporting results;
- utilize new data correlation software/techniques;
- evaluate existing tool solutions;
- investigate emerging technologies and their impact on the above approaches;

- provide tools to provide interpretive capabilities, and improve the visualization of reports.

Note that collected computer monitoring information is highly sensitive and must be handled appropriately. System administration authorities should be distributed, requiring cooperative effort for certain functions.

FILTER MALICIOUS CODE AT SYSTEM CHOKE-POINTS

Outsiders and insiders to an information system can both introduce malicious mobile code into the system. Such code (e.g., in the form of viruses, worms, Trojan horse software, executable macros attached to documents, and the like) can perform various monitoring, data access, and other activities once within the system.

Most systems have certain "choke points" through which data must flow during system operations. The emphasis in this recommendation is on using existing tools and techniques to monitor for and control mobile malicious code at those choke points. Specific recommendations are:

- apply ingress and egress filtering for known malicious code. This includes installing antivirus filters, firewalls, sandboxes, and byte code checkers;
- focus on such high-leverage choke-points as firewalls, email servers, FTP servers, web servers, and file servers.

Longer-term, activities would include enterprise-wide distribution and management of virus definitions and detection packages, and development of automated means to verify signature files enterprise-wide.

FILTER FOR FUTURE AND UNKNOWN MALICIOUS CODE, AND EXERCISE MITIGATION AND CONTAINMENT

This recommendation is similar to the previous one, but emphasizes the ability to detect previously-unknown malicious code. For these purposes, it is not sufficient to have a database of "signatures" of known viruses, worms, and malicious macros.

One the primary means of implementing this recommendation is the installation and use of "change detection" tools - ones that know the default configuration of application programs and the operating system,

and are capable of detecting any changes to those standard configurations. (Those changes may be valid, but at least if they are flagged they might then be checked for such validity by systems personnel.)

Longer-term strategies in this regard include the development and use of "proof-carrying code" as described in section 2 of this report.

TRACK DATA PEDIGREE AND INTEGRITY

Insiders have easy access to a number of data files within an organization. It would be extremely useful, and important for forensic purposes after an event occurs, to be able to track the "pedigree" of data - a log of that file's establishment (by whom? when? with what content?) and subsequent changes to it (by whom? when? with what content?). With such tools, it may be possible to prevent and detect unauthorized access, modification, destruction, propagation, and creation of data. As mentioned above, such logs - if one can assure their untamperability and integrity - can be extremely useful for forensic activities. These logs also assist in countering non-repudiation by insiders, and in the re-establishment of trust of data after an attack.

Among the short-term activities to be instituted in this regard are:

- use of digital signatures in combination with strong authentication;
- use of checksums on data files;
- use of off-site file backups.

Longer-term research strategies on determining data pedigree and integrity were discussed more fully under the topic "Provenance" in section 2, above. They include:

- identifying existing or emerging watermarking tools for data files;
- identifying what attributes of a data file should be tracked as part of its pedigree;
- identifying means for detecting use of known steganographic technologies;

- developing solutions to re-establish data trust, once it has been compromised.

OTHER SUGGESTIONS

Several other ideas for near-term solution strategies regarding the insider threat were mentioned during this group's deliberations, but these were not explored in greater detail. They are listed here for completeness:

- *Application wrappers.* Explore means of "wrapping" existing applications, so that their behavior may be circumscribed and monitored;

- *Dynamic alerting 'banners' as deterrents.* Existing software allows "banners" to be displayed to the user under various conditions. Use these banners to warn the user of their seemingly abnormal behavior, or reminding them that their actions are potentially monitored, etc. If such messages are tailored and dynamic, rather than mere "message of the day"s, they might have significant deterrent effect.

- *Sandboxes.* A "sandbox" is a defined, limited area within a system in which actions may be carried out, but such actions can not affect the system outside of that sandbox. To the extent that users can be placed into such sandboxes in their working environment without affecting productivity, it is a useful means of limiting potentially damaging actions that they might perform, and monitoring their actions while within those constrained areas.

- *Red teaming to detect insider models.* It was suggested that "red teams" be used to emulate current insider threat models, and test the resulting detection rates. That is, given what we believe a malicious insider "looks like," and given the solution sets proposed, use red teaming to test the effectiveness of those protection, detection, and response mechanisms. The intent is to try to measure how effective our solutions are.

5. CONCLUDING REMARKS


The purpose of this workshop was to create specific research and development recommendations that could be used to guide funding agencies in their information assurance programs, with special emphasis on mitigating the insider threat problem. It was felt by all participants that this goal was in general accomplished. A secondary outcome of the workshop was the creation and enhancement of personal contacts among key researchers in this area, and between the research community and potential users of the results of their work.

As mentioned earlier in this report, this workshop was one of a series of workshops and reports on mitigating the insider threat. It was felt by participants that successfully built upon that previous work, but should not be the last in this series. The problem is complex and requires continuing attention. However, we have come far enough in describing specific aspects of the problem and of specific mitigation strategies that focused but multifaceted research program - along the lines described in this proceedings - should be initiated as soon as possible.

Appendix

A. AN INSIDER THREAT MODEL FOR MODEL ADVERSARIES


Brad Wood, SRI International



An Insider Threat Model for Model Adversaries

Brad Wood
Cyber Defense Research Center
SRI International

[Redacted]



Purpose of This Model

[Redacted]

- Allow red teams and other model adversaries to accurately simulate the behavior of the “insider threat.”
- Goal of this effort is to
 - Develop a credible insider threat model
 - Use the model to discover weaknesses in and defenses against this adversary

7/12/00

Attributes of The Insider

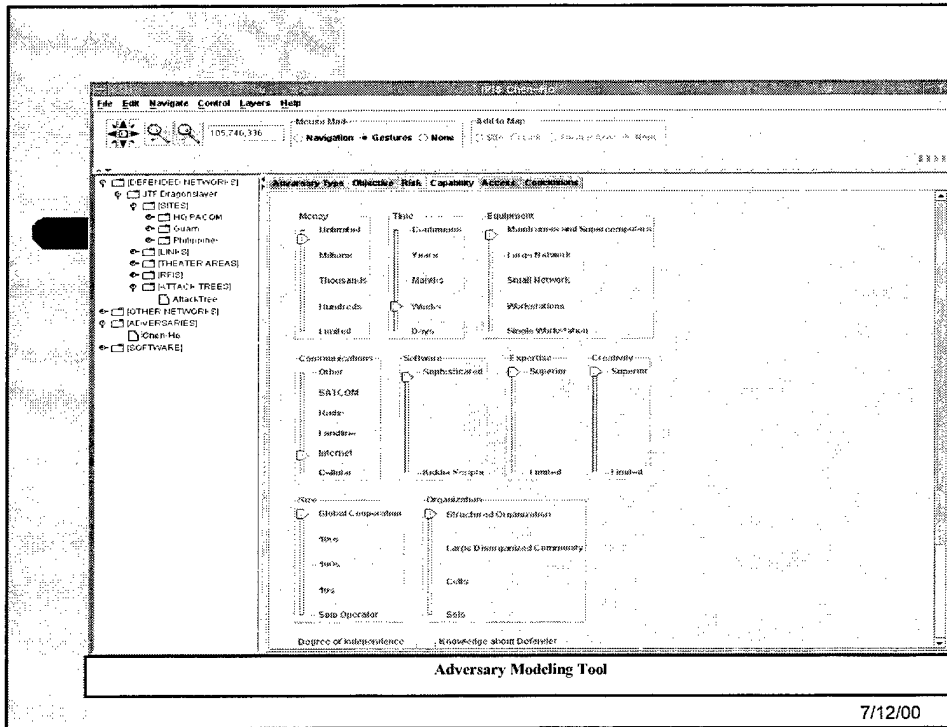
- Access
- Knowledge
- Privilege
- Skill
- Risk
- Tactics
- Motivation
- Process

7/12/00

Observations

- Who would do such a thing?
 - Someone with some character defect?
 - Operative from outside the enterprise. *
- Insider probably knows now NOT to get caught.
- Are cyber means the best way to catch this adversary?
 - Analysis should identify potential targets.
 - People with access to targets should be monitored.
 - Personnel reliability processes might identify these people.

7/12/00



Next Steps

- Socialize, review, challenge, and validate a credible model.
- Use model in operational red-team exercises to gather relevant data.
- Review data and results to identify “interesting” results:
 - Weaknesses, defenses, etc.

Summary

- We have a preliminary model of a malicious insider.
- Where do we go from here:
 - This model needs to be reviewed and validated.
 - Model needs to be implemented in an exercise to generate data.
 - Community must review data to develop new hypotheses on insider defenses.

7/12/00

Appendix

B. AN INSIDER THREAT MODEL FOR ADVERSARY SIMULATION

Brad Wood, SRI International

ABSTRACT

This position paper presents a preliminary analytical model for the malicious insider cyber adversary. The purpose of this model is to provide a simple but effective means of simulating the behavior this adversary. The purpose of this paper is to stimulate discussion and suggestion research activities to counter this adversary.

INTRODUCTION

The malicious *Insider* has long been identified as a serious threat to modern information systems. However, few attempts have been made to model and understand this adversary.

One way of developing defenses against this adversary is to model it using Red Teams and other model adversaries. Once a credible model is developed, some understanding can be gained about this adversary. Once we have some understanding, then we can develop effective defenses to thwart this adversary.

The purpose of this paper is to list some basic assumptions and assertions regarding the *Insider* so that this adversary can be modeled and studied. Ideally, these assumptions will be challenged or validated by the Information Assurance & Survivability community at large.

Ideally, the study of a credible *Insider* threat model could lead to valuable insights and other observations that may lead to effective mitigation strategies for the *Insider* Threat.

GOALS AND PURPOSE

The purpose for developing this adversary model include:

- Allow red teams and other model adversaries to accurately simulate the behavior of this threat.
- Provide researchers with some mechanism for testing countermeasures that are designed to mitigate this threat.
- Expose weaknesses that lead to effective countermeasures.

ATTRIBUTES OF THE INSIDER

The malicious *Insider* can be described from a variety of attributes. These attributes include: access, knowledge, privileges, skills, risk, tactics, motivation, and process.

For the purpose of this discussion, a *system* is the overall network within the scope of some relevant management domain. A *target* is the portion of a *system* that is subject to attack by the malicious *Insider*.

Access

The *Insider* has unfettered access to some part or parts of the system. One or more of the following assertions are assumed to be true:

- The *Insider* attacks the target from behind or inside a system's perimeter defenses.
- The *Insider* can breach a system's perimeter defenses without arousing the suspicion of network security managers.
- The *Insider* has physical access to the system that thwarts its perimeter network defenses.

One open issue is the case where an outside attacker gains access to the inside of a network by attacking or exploiting weaknesses in the network's perimeter defenses. For the purposes of this discussion, we will assume that the outside attacker who gains access is not considered an *Insider* unless they possess the other attributes of the *Insider*.

Knowledge

The *Insider* has extensive knowledge of both the system and the target. In particular:

- He or she has unfettered access to all documentation on the target and the system.
- The *Insider* can collect intelligence and perform discovery without arousing suspicion.
- The adversary has exquisite detailed information on the target.
- The adversary also has good intelligence on the entire system.

In some cases, the *Insider* may possess or be entrusted with the only accurate detailed information on the target.

Privileges

The *Insider* should have no problem getting the privileges they need to mount an attack. In particular:

- The adversary may not need *root* or administrator access to mount an attack.
- The *Insider* may already have privileged access to the target.
- The *Insider* may have to transition to some unauthorized privileged mode to mount a particular attack.
- The knowledgeable *Insider* may simply recruit someone who has the privileges needed to mount an attack.

The adversary may actually be the individual responsible for monitoring or enforcing the security policy on the target or system.

Skills

The knowledgeable *Insider* has the skills to mount a credible attack, subject to some limitations:

- The *Insider* may actually be the local domain expert on certain parts of the system.
- The *Insider* has at least some base-line familiarity with the target, and they can gather additional intelligence without arousing suspicion.
- A given adversary will not likely attack an unfamiliar target. (This assumption is based on the belief that the adversary will prefer to attack a familiar target rather than gain expertise with an unfamiliar target.)
- The adversary may actually be a local domain expert on the target.
- The adversary will generally work within their domain expertise.

Risk

The *Insider* is very risk-averse. Their ultimate defeat is to be discovered before they have mounted a successful attack. This leads to these assumptions:

- The *Insider* generally works alone.

- This adversary may recruit colleagues on an operation, but they will only be parties that the adversary trusts, and the adversary will employ others only to the extent necessary.
- The *Insider* may be able to co-op a colleague into enabling an attack without that person's knowledge.

Tactics

The tactics used by this adversary are completely dependent on the goals of the attack. These tactics might include:

- *Plant, run, & hit.* Here, the adversary is attempting to plant some malicious code (or the equivalent), leave the premises, and be well out of reach of any authority when the attack is launched. This seems like a logical tactic for cases where the objective of the attack is to disrupt or destroy part of the system.
- *Attack and (eventually) run.* In this case, the adversary launches their attack, but they are in no hurry to leave the scene of the crime. This might be an appropriate tactic for attacks that are not designed to disrupt the target. This gives the adversary time to leave the scene at a leisurely pace. This also gives the adversary an opportunity to assess the success or failure of their attack prior to leaving the enterprise.
- *Attack until caught.* Here, the adversary attacks a target repeatedly until some (unknown) event occurs. Eventually, we expect the adversary would be caught and prosecuted. However, it is not clear that this tactic is worthy of serious consideration.
- *Espionage.* In this case, the value of the adversary is measured in their ability to exfiltrate information from an enterprise. The adversary may choose to simply carry the information out with them at night, or they may attempt (the riskier option of) transmitting the information over the network. The only reason for the *Insider* to cease this campaign would be if they were discovered.

Motivation

What motivates the *Insider*? We expect the typical *Insider* is trying to force some sort of undesirable consequence within an enterprise to forward one of the following goals:

- *Profit* - Some party is paying the adversary to disrupt the target.
- *Provoke change* - The adversary is attempting to motivate some change in the enterprise. This could include invoking some sort of policy change or even blackmail.
- *Subversion* - The adversary may be trying to subvert the mission of the target organization.
- *Personal motive* - The adversary may be trying to exact some sort of revenge against the enterprise. The *Insider* may also be trying to demonstrate their prowess in controlling some portion of the enterprise.

Process

An *Insider* attack follows a basic, predictable process.

- a. *Someone becomes motivated to attack.* Either some event leads to the individual's discontent, or someone is planted who will act to subvert the enterprise. Some internal or external party may also recruit an individual.
- b. *Adversary identifies target.* Either the adversary identifies a target and mission that meets their personal need, or some outside party suggests a target of interest.
- c. *Adversary plans operation.* The malicious insider conducts some reconnaissance on the target. They plan an operation. They may even recruit associates to assist the operation. This process usually concludes only when the adversary has planned an attack that has some reasonable probability of success without prematurely disclosing the existence of the operation.
- d. *Launch attack.* It is not clear what the adversary will do once the attack is launched. Some options include: damage assessment, flee in a hurry, flee when convenient, or repeat the operation until they are either successful or caught.

OBSERVATIONS

Who would mount an *Insider* attack?

- Someone with a character defect.
- An operative from a competitive organization.

Ideally, there would be a large body of information in the counter-intelligence community that could assist in supporting or refuting this model.

A knowledgeable *Insider* probably knows how to mount an attack within a given system without getting caught, or they might believe this is true. In fact, our adversary may actually control the mechanisms that are supposed to thwart their attack.

Are cyber means the best way to thwart this adversary?

- Vulnerability analysis should identify potential targets in a given system.
- People with access to targets should be monitored.
- Personnel reliability methods might identify these people.
- It may be viable to counter this threat using traditional counter-intelligence methods.

What are the manifestations of an *Insider* attack? It is not clear that a typical cyber defender could identify *Insider* activity even it was known to exist.

EDITORIAL

There appears to be a vast gulf between the needs of the Operators in the field and the Researchers in the laboratories. Many Operators could benefit from the appropriate application of tools and techniques. Typically, the missing ingredient for the Operator (defender) is resources (time and funding). Applying only modest resources can help many current operations problems.

The real problems appear to be:

- a. Operators are severely limited in the resources they have to apply to Information Operations.
- b. Information Survivability is not an imperative in the government or private sectors.

This leads to an untenable situation where Operators are seeking "silver bullets" - solutions that cost little to develop but address hard problems. This essentially concedes a significant advantage to our potential adversaries until such time as some catastrophic event makes *Information Survivability* a national imperative.

NEXT STEPS

1. Knowledgeable authorities to need to validate or challenge the assertions contained in this model.
2. Researchers should seek hard data to support or refute the basis of this model.
3. The model should be implemented and studied by researchers and operators.
4. Observations, insights, and developments dealing with the *Insider* threat need to be widely communicated.
5. Defenses and strategies should be tested against this model.

A template for working with this type of adversary model can be found in [1]. Here, a model developed to simulate a *cyber terrorist* was used to develop innovative defenses against an entire class of sophisticated.

CONCLUSION

This paper proposes a model for simulating the malicious *Insider* threat. This model should be challenged and refined by the Information Assurance & Survivability community. After some refinement, experiments with this model could lead to new and innovative defenses against the malicious *Insider*.

THE AUTHOR

Bradley J. Wood is a Staff Scientist in the System Design Laboratory at SRI International. Brad leads the Research Red Team of the Cyber Defense Research Center at SRI. Brad's research interests involve using red teams and other model adversaries to improve cyber system assurance. Brad was formerly a Distinguished Member of Technical Staff at Sandia National Laboratories where he was the Program Manager

for the Information Design Assurance Red Team. You can reach the author via email at this address: bradley.wood@sri.com

REFERENCE

- [1] Schudel, Gregg and Wood, Bradley, *Modeling Behavior of the Cyber Terrorist*, (submitted to) 2000 National Information Systems Security Conference, Baltimore MD, October 2000. Also included as Appendix C of these proceedings.

Appendix

C. MODELING BEHAVIOR OF THE CYBER-TERRORIST

Gregg Schudel
(gschudel@bbn.com)
Senior Engineer
Information Assurance
Program Integration Team
GTE/BBN Technologies
2110 Washington Blvd., Suite 100
Arlington, VA 22204

Bradley Wood
(bjwood@sandia.gov)
Distinguished Member of Technical Staff
Information Design Assurance Red Team
Sandia National Laboratories
PO Box 5800, M/S 0449
Albuquerque, NM 87185

ABSTRACT

The Cyber-Terrorist is assumed to be a very real threat to modern information systems, especially those trusted to control the nation's defenses and critical infrastructure. Very little intelligence or solid data exist regarding this adversary.

This discussion chronicles the efforts by a team at the Defense Advanced Research Projects Agency to model and characterize this adversary. The ultimate goal of this research is to develop defenses against this new and sophisticated adversary.

It is not clear whether the Cyber-Terrorist is real or simply a theoretical class of adversary. Very little intelligence or other data exist in open literature that characterizes the behavior or existence of this class of adversary.

We will argue that the Cyber-Terrorist is a very real potential threat to modern information systems. Therefore, sophisticated defenders must understand the capabilities and behavior of this adversary in order to defend against it.

The Defense Advanced Research Projects Agency's (DARPA's) Information Assurance (IA) Program is attempting to incorporate the Cyber-Terrorist into a larger model of threats poised against information systems operated by the US Department of Defense. This paper lists some of the basic assumptions about this adversary, with the intent that these assumptions may be challenged within the research community.

FUNDAMENTAL HYPOTHESIS

This work is based on the fundamental hypothesis that the Cyber-Terrorist is a very real threat to modern information systems. Unfortunately, we are unaware of any research or other hard data that supports this hypothesis. Rather, this hypothesis is based on the following assertions:

- Terrorist threats still exist against the United States and their interests abroad. [1]
- Information systems that manage the nation's defenses and critical infrastructures are vulnerable to cyber attacks [2]
- Terrorists can forward their agenda by attacking the nation's critical infrastructures [3]
- Cyber attack costs (especially in proportion to their perceived relative effectiveness) asymmetrically favor the cyber-terrorist
- The ability for the cyber-terrorist to conduct attacks against US assets from foreign shores with little risk of consequence appears to be reality.

Therefore, it stands to reason that the nation is vulnerable to cyber-terrorism.

AN APPROACH TO MODELING THE CYBER-TERRORIST ADVERSARY

If we accept that this adversary exists, then how do designers defend against it without the benefit of documented attack cases? DARPA's IA program chose to study this adversary through the use of red teaming to simulate this adversary. Here, the adversary is modeled by the Information Design Assurance Red Team (IDART) [4] at Sandia National Laboratories [5].

Basic Assumptions

IDART's model of the cyber terrorist is based on the following assumptions:

Sophistication: The cyber-terrorist is believed to have a level of sophistication somewhere between that of a sophisticated hacker and a foreign intelligence organization (see Figure 1). The cyber-terrorist might even employ sophisticated or professional hackers in their operations. However, this adversary would not have access to any of the

very sophisticated attacks that are available to members of the intelligence community.

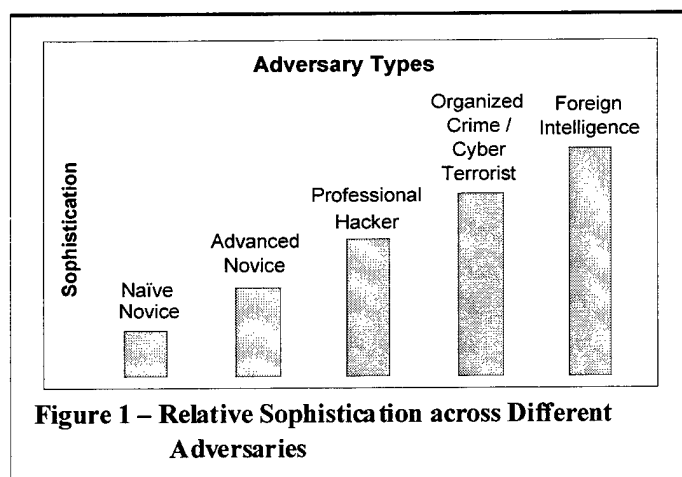


Figure 1 – Relative Sophistication across Different Adversaries

Resources: This adversary is believed to have access to all commercial resources that are generally available. These include:

- All publicly available information. This includes information on tools, attacks, and specific intelligence on a particular target.
- Consultants and other commercially available expertise.
- Any commercially available technology such as workstations, software, hardware, and diagnostic tools.
- Software developers, network developers, and other expertise required for developing their own attacks against a particular target.

This adversary is assumed to have limited funding. However, he is assumed to be able to raise funds on the order of hundreds of thousands to a few million dollars, and he is willing to spend these funds to accomplish his mission.

Intelligence: This adversary is assumed to be able to acquire all design information on a system of interest. This assumption is based on the following assertions:

- Much of the information is publicly available.
- Information that is not generally available is loosely controlled.
- Information that is controlled can be exfiltrated by bribing a trusted insider or through extortion.

Life Cycle: A sophisticated adversary could influence the life cycle of a particular product by influencing developers or individuals with access to the product's development. The cyber-terrorist may also attack product distribution channels in an effort to modify components before they are delivered for integration into the target system.

Risk Aversion: This adversary is assumed to be very risk averse. Premature detection is a serious negative consequence for the cyber-terrorist. This adversary may elect to mount an obvious or notorious attack on a system, but only at the time of their choosing. This has several important ramifications, some of which are illustrated in Figure 2:

- The cyber-terrorist is effectively neutralized if they are discovered before they attack.
- The cyber-terrorist will prefer quiet, stealthy, and passive techniques for attacking a system.
- An adversary will not attack a system if their perceived risk is above their tolerance or threshold.
- The adversary's risk tolerance actually decreases over time, because their exposure or risk increases over time.

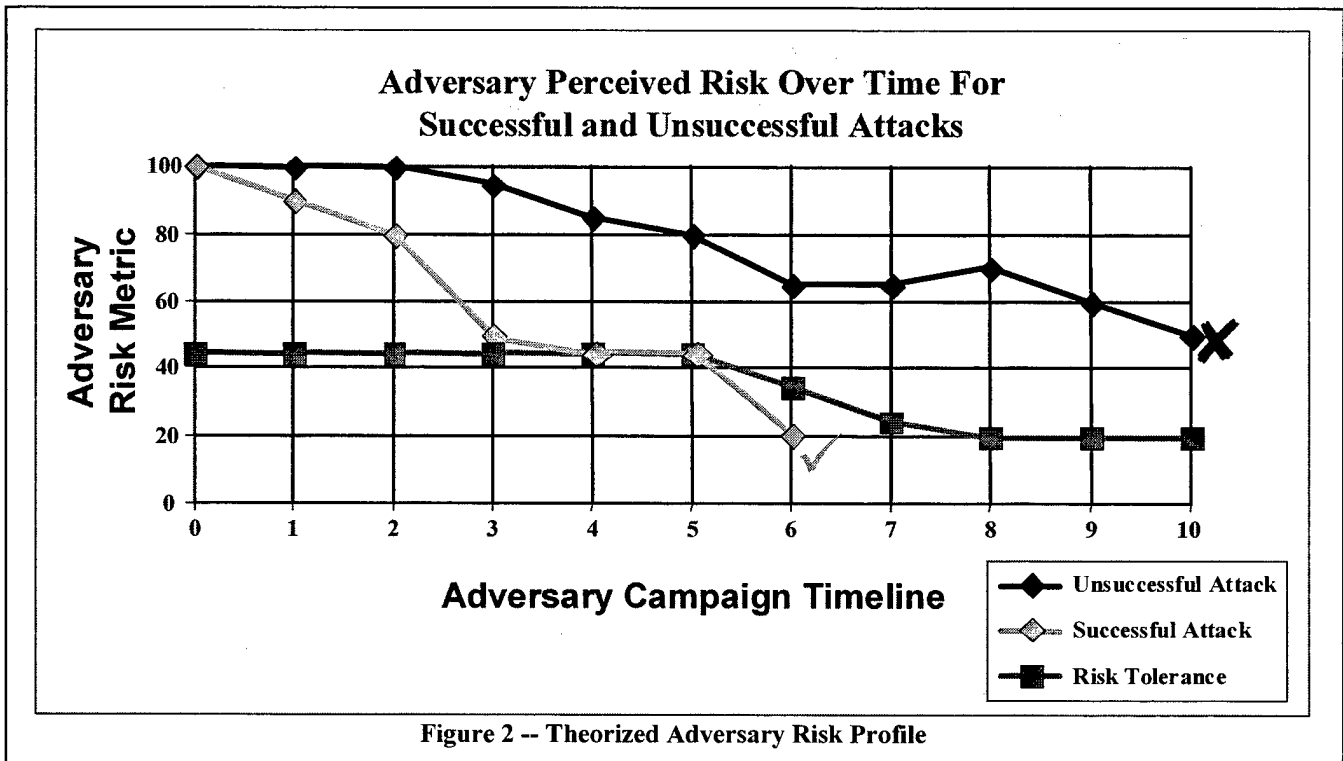


Figure 2 -- Theorized Adversary Risk Profile

Specific Targets: This adversary has specific targets or goals in mind when they attack a given system. Unlike hackers or naïve adversaries, the cyber-terrorist will attempt to target the exact host or system that must be compromised to accomplish their mission.

The adversary will also expend only the minimum amount of resources needed to accomplish their mission. They have no incentive to expend more resources than is absolutely necessary.

Other Behavioral Expectations: The cyber-terrorist is assumed to be professional, creative, and very clever. They will seek unorthodox and original methods to accomplish their goals. Individuals who are well-schooled in traditional information security techniques are not well suited to being a cyber-terrorist, simply because they have been exposed to or trained in classic security techniques and doctrine. The cyber-terrorist will seek to accomplish their mission by techniques not mitigated by classic security mechanisms.

EARLY OBSERVATIONS

IDART has served as a model cyber-terrorist for DARPA's IA program since July 1998. In April 1999, the GTE-led IA Integration Team held a review in Albuquerque, NM, to determine if any patterns had emerged from observations of the red team. Several patterns of significance were discovered and are summarized below.

Attack Process

The red team used essentially the same process for each mission as shown in Figure 3. A study of this process yielded these assertions:

- The red team spends most of their time gathering intelligence on the target system.
- The red team observes a target system until they can either (a) successfully attack the system or (b) they exhaust all available resources. Success for the red team includes both preserving stealth and meeting their mission objectives.
- The red team will give up before they will mount an attack that is above their risk threshold.
- The fact that this red team follows the same basic process could make it vulnerable to some countermeasures.

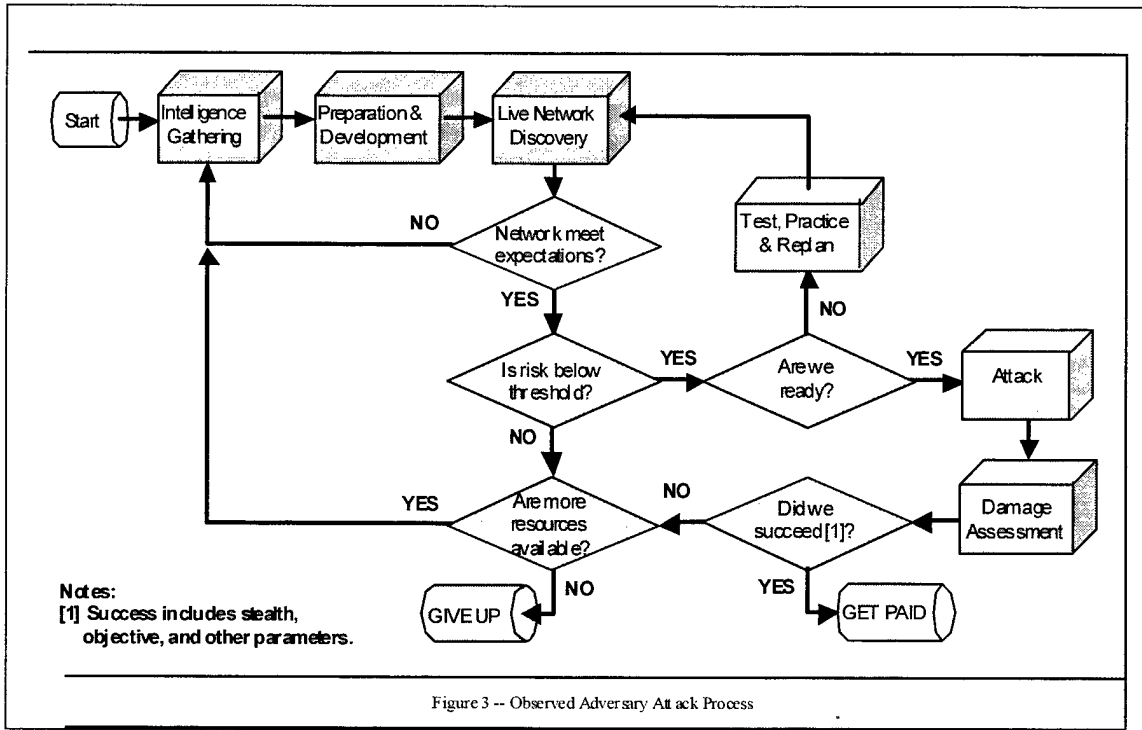


Figure 3 -- Observed Adversary Attack Process

Timing Analysis

GTE's IA Integration Team then studied how the red team spent their time in relation to the process described in Figure 3. These results are shown in Figure 4. This suggests that the red team spent the majority of their time gathering intelligence on target networks. This is consistent with other observations of cyber-adversaries [6].




Figure 4 -- Observed Adversary Time Expenditures

EARLY EXPERIENCES

IDART has played the cyber-terrorist in several DARPA exercises. Some of these exercises suggest some interesting and often unexpected results.

Information Superiority Technology Integration exercise 1998 (ISTI98) [7]- This was a large exercise that explored the hypothesis that war gaming could yield effective data to gauge the relative strengths and weaknesses of an information system. Here, IDART was one of two red teams attacking this network. Current assessments of the available ISTI data seem to refute the fundamental hypothesis. Current data suggest that careful experimentation yields better data.

DARPA IA Laboratory Exercise RT-1999-01 on Layered Defenses [8]- This exercise explored the fundamental hypothesis that information assurance technology layers add in overall defensive strength. Current assessments of the data from RT-1999-01 suggest that breadth is more important than depth, simply because the red team tended to work around the information assurance technologies that were deployed as obstacles to the adversary. RT-1999-01 data also suggested that unintended adverse interactions between layers could occur if they are not properly coordinated, and that the red team could exploit these interactions as a denial of service control surface. These results seem to suggest that the IA community needs a better understanding of the relationship between depth and breadth in developing and deploying layering strategies.

DARPA IA Analysis Exercise RT-1999-02 on Wrappers for Microsoft Windows NT [9]- This exercise was a study of how the red team might thwart the security services offered by Non-bypassable NT Security Wrappers [10]. The red team's report suggested that although this technology appears effective in preventing certain types of attacks, a clever adversary can still attain their goals on the target platform using different attacks that completely circumvent the supplemental security system.

DARPA IA Analysis Exercise RT-1999-03 on Adversary Behavior and Dynamic Defense [11]- During this analytical exercise, the GTE IA Integration team thoroughly debriefed the red team to characterize any

trends in their behavior. The motivation for this analysis was to determine if dynamic defense strategies had the potential for significantly impacting red team capabilities. Figures 2, 3, and 4 (above) of this appendix are examples of the results of this session.

DARPA IA Laboratory Exercise RT-1999-07 on Dynamic Defense [12]- This exercise resulted from the RT-1999-03 analysis, and was intended to explore the hypothesis that dynamic defenses - in this case, dynamic (on the fly) network reconfiguration) - can increase an adversary's work factor. Early interpretations of the data from this exercise support this hypothesis, although this assertion is supported in part by unexpected results from the experiment.

FUTURE DIRECTIONS

DARPA's experience suggests some improvements to the process that we are using to model the cyber-terrorist adversary.

Additional Red Teams - Additional red teams could generate more data that either supports or refutes the IDART results. Ideally, these red teams would provide some different perspective and some different results than the current team.

Improved Scientific Methods - One goal of the current DARPA effort is to improve the processes and procedures used to experiment with and gather data from red teams. Ultimately, each red team exercise should gather credible data that either supports or refutes some fundamental process. Gathering good data while preserving the possibility of the unexpected results is a constant challenge for the DARPA IA team.

Incorporate Verified Terrorist Behavior - No efforts have yet been made to research and incorporate models of actual terrorist behavior. Although there is little or no data on the behavior of the cyber-terrorist, it would be beneficial to attempt to incorporate traditional terrorist behavior in the cyber-realm.

War Game Cyber-Terrorist Scenarios - No efforts have yet been made to research the speed at which damage could be inflicted and its potential impact on defense and critical infrastructures. It would be beneficial to develop a few operational scenarios and then to run analytical cyber attacks.

Possible Approaches to Classical "Difficult Problems" - Work with red teams could lead to viable defenses against some classical IA problems that are currently believed to be difficult or impossible to solve. These include:

- **Life-cycle attacks** - Here, we assume that the adversary can influence the development of IA products. Credible defenses could evolve through studying the way adversaries mount these kinds of attacks.
- **Platform vulnerabilities** - It is widely held that IA designers can build robust networks to connect relatively insecure host platforms. Therefore, an adversary will likely attack the platforms if he can complete his mission. One red team member put it this way: "Why attack a hardened network when the same data is available on a nice juicy defenseless host?" Ideally, using red teams as adversaries might suggest approaches to improving data protections in this environment.
- **Users as adversaries** - It is widely held that it is difficult to build an information system that provides reliable access for a variety of mutually adversarial users. Studies of red teams as adversarial users may suggest approaches to this problem.
- **Knowledgeable insiders** - Conventional wisdom holds that designers cannot effectively protect themselves against a knowledgeable insider. However, current data suggests that this is a critical vulnerability in most high-consequence information systems. Red teams could be employed to study this problem and develop effective countermeasures.
- **Denial of service attacks** - It is widely held that designers cannot defend against an adversary who is intent on mounting a denial-of-service attack. Red teams could be used to study these attacks with the hope of developing effective defenses.

SUMMARY

It is not clear that the cyber terrorist actually exists. However, its existence has been hypothesized, and there is no data that clearly refutes the existence of this adversary. DARPA is attempting to study

this adversary in an attempt to proactively combat the potential threats posed by this adversary.

DARPA's IA Program has engaged Sandia's Information Design Assurance Red Team (IDART) to model this adversary. This paper discusses some of IDART's assumptions about this adversary as well as some of the early results of incorporating this adversary in DARPA's IA program. Finally, we theorize how red teams can be employed to develop credible defenses against some classically difficult IA problems.

REFERENCES

- [1] **Combating Terrorism:** Presidential Decision Directive 62, 22 May 1998)
- [2] **Protecting America's Critical Infrastructures:** Presidential Decision Directive 63, May 1998
- [3] **Critical Foundations:** Protecting America's Infrastructures, The Report of the President's Commission on Critical Infrastructure Protection, October 1997.
- [4] Information available on the World Wide Web at <http://www.sandia.gov/idart>
- [5] Information available on the World Wide Web at <http://www.sandia.gov>
- [6] Longstaff, T. A., et al, **Security of the Internet**, Froehlich/Kent Encyclopedia of Telecommunications vol. 15, Marcel Dekker, New York, 1997, pp. 231-256
- [7] Wood, B. J., **ISTI 98** After Action Report of Red Team # 2, December 1998, Sandia National Laboratories, to be published on the World Wide Web at <https://www.ests.bbn.com/>
- [8] Bouchard, J. F, Parks, R. C., Wood, B. J., **Attacking Layered Defenses**, Red Team Results from Exercise 1999-01, April 1999, Sandia National Laboratories
- [9] Obenauf, T., **NT Wrappers "Quick Look":** Red Team Results from Exercise 1999-02, May 1999, Sandia National Laboratories, to be published on the World Wide Web at <https://www.ests.bbn.com/>
- [10] Balzer, Robert, Nonbypassable NT Wrappers, FY98 DARPA Research Efforts. Further information is available directly from the author via email at balzer@isi.edu

- [11] Schudel, G, and Wood, B. J., **Adversary Behavior**, Results from Red Team Exercise 1999-03, April 1999, DARPA Information Assurance Program, to be published on the World Wide Web at <https://www.ests.bbn.com>

- [12] Duggan, D., and Wood, B. J., **Layered Defense**, Red Team Results from Exercise 1999-07, June 1999, Sandia National Laboratories, to be published on the World Wide Web at <https://www.ests.bbn.com>

Appendix

D. CAN TECHNOLOGY REDUCE THE INSIDER THREAT?

Michael Caloyannides and Carl Landwehr

Mitretek Systems

Assorted statistics have estimated the malicious insider to be responsible for 70%-80% of successful attacks on computing systems. Yet, most ongoing research on ways to either protect computing systems from attack or to ensure that they can function despite an attack has been focusing on the assumption that the threat is outside the computing system of interest.

This implicit fatalism about the insider threat is based on the multiple assumptions that:

- a) the insider knows everything about all of the protective measures in place, and
- b) is in a position to defeat them all, and,
- c) can erase all evidence

This paper shows that these assumptions need not be true and, therefore, that technology can help reduce the insider threat problem.

It is a time-tested axiom of military and intelligence operations that no one person who could be compromised for whatever reason needs to know *everything* about the operation (e.g. contingency plans, intelligence sources and methods, locations of caches, etc.), for obvious reasons. Access to particularly sensitive information often requires two authorized and qualified persons to be present; access to sensitive command centers often requires two or more individuals to concur and to act in unison before a momentous action can be taken.

Such common sense precautions are not viewed as an insult to the trustworthiness of any one person, but as an affirmation of human nature and the need to protect far-reaching equities despite the vagaries of human nature. Similarly, there is no reason why any one individual, such as the systems administrator, should have full knowledge of *all* security measures in place, nor is there any reason why a single individual, acting alone, should be able to trash an entire computing system.

Once it is accepted that no one individual should have full and unchecked access to all aspects of a computing system, then it is relatively easy to come up with a number of technological means of ensuring that, and of documenting -beyond the reach of any one person acting alone- all attempts to subvert a computing system. System administrators should not take offense to this; quite the contrary, they should welcome such a concept because it removes them from the unenviable position of the "prime suspect" and allows them to function professionally in their important role.

Protecting -or curing- a computing system from an attack is not much different in principle from protecting -or curing- a human being from "disease". There is no single "anti-illness vaccine" not a single magic pill that cures all illnesses. This obvious truism seems to be regularly forgotten by some researchers into information systems security who keep looking for "the" solution. There isn't one; there can't be one; there won't be one.

As with protecting any person or any facility, one needs a large repertoire of concurrent protective measures, each one intended mostly for a particular kind of threat. In the case of facilities to be protected, one protects the perimeter and hires guards and uses access control of some sort and uses the concept of concentric spherical layers of security "just in case" an intrusion gets past an outer layer or two, and does background checks on employees, and.. and.. etc.

Similarly, a successful means of protecting a computing system from a malicious insider must also involve a number of concurrent protective measures, such as those described below. In addition, however, it is essential that:

- a) no one insider knows about all of these measures, and
- b) no one insider is in a physical position to defeat them all, and
- c) no one insider is in a position to delete audit trails.

1. AUTHENTICATION

The primary means of authenticating insiders continues to be the password. Expecting humans to invent something that is random, changes frequently, and is to be remembered without being written down flies in

the face of common sense as well as human factors research. Biometric authentication in various forms is becoming inexpensive, accurate, and convenient enough to replace passwords, and spoofing a biometric is likely to be significantly harder for an insider wishing to masquerade than is stealing or guessing a password. A further consideration should be the use of tokens or other measures for continuous authentication, so that when an authenticated user leaves a workstation temporarily, his or her authenticated identity does not become available to anyone who happens to have physical access to the same workstation.

2. ACCESS CONTROL

Just as a system administrator does not need -nor should want- unquestioned access to all security related protective elements of a computing system, other users' access levels should be even more carefully circumscribed. The usual system-level controls as to who can "read" what, and especially who can "write" to what, should be strictly controlled.

The real problem is that today's complex software *is* buggy and allows the savvy unprivileged user (or even outsider) to exploit such "bugs" and gain access to the inner sanctum sanctorum of a computing system without any authorization at all. The exploitation of the many variants of the "buffer overflow" is a typical example that accounts for the vast majority of "hacks" into computing systems by insiders and outsiders alike.

Systems have been built that partition, for example, Unix superuser privileges, as have specialized systems that enforce the two person rule for specific operations. Role based access controls, properly implemented and configured, can provide a framework that can bring the principle of least privilege closer to realization.

3. AUDIT TRAILS

Audit trails have been plagued with a perennial problem: Nobody knows what constitutes "suspicious" conduct so as to record that only. If one opts to err on the side of caution and record almost everything, then the volume of audit trails recorded become unmanageably vast in very little time.

Here, too, there is no silver bullet. Any "solution" can only be optimized for each particular installation. What *is* important, though, is that audit trails should not be erasable by the perpetrator -or his/her accomplices-. Writing onto WORM writeable CDs that are remotely located is an option.

4. PROTECTING THE INFORMATION FROM OBSERVATION

Partly because encryption has been cumbersome in the past, rather than transparent to the authorized user as it should be, and partly because of the parochial pressures exerted by law enforcement to minimize the use of encryption, sensitive information in US corporate databases that should have been encrypted was not. Even the system administrator does not care (nor should have an interest in) the content of proprietary corporate information, let alone the average unauthorized employee or, even more so, the malicious outsider.

Databases should be encrypted by default; on the fly encryption and decryption should make the content accessible to authorized users only in a manner that is transparent to them. With today's high speed computers, the overhead of encryption and decryption is negligible if one uses modern computationally efficient encryption algorithms such as *Twofish* and *Blowfish*, rather than the slow and discredited DES.

For that matter, key system files and even executables can -and should be- encrypted as well, using on-the-fly encryption/decryption as well, but with different encryption keys. Even a lowly user of a typical PC today can get commercial software/hardware that allows one to encrypt one's entire hard disk on a track-by-track and sector-by-sector basis, in a manner which is transparent to the authorized user (but totally inaccessible to anyone else, including a computer forensics expert).

5. PROTECTING INFORMATION FROM UNAUTHORIZED "EDIT/COPY/PASTE"

Protecting information from unauthorized "edit/copy/paste" and even from unauthorized printed pages containing sensitive corporate information from leaving a building as either email or printed paper.

A document marked "proprietary", or otherwise intended not to be broadcast to unauthorized recipients can easily be electronically copied

onto another and emailed to the other side of the world or copied onto a floppy disk or printed out and be handcarried past any guard. Most can recall the US automaker's senior employees who walked out with numerous proprietary documents and joined a foreign automaker shortly thereafter.

There are numerous technical means to minimize such occurrences in the future:

a) One can use some of the recently commercialized systems whereby an email cannot be copied, nor printed out without the express authorization of the sender. These schemes usually require the installation of an additional IT infrastructure within an organization to handle such sensitive documents.

b) Documents can be electronically "watermarked" so that their passage through any electronic gate leaving the secure facility (e.g. as an attachment or an Edit/Copy/Paste operation) can be automatically detected and prevented. Such watermarks have to be robust enough so as to withstand attempts to remove them. Such technologies have been developed for digitized images and sound files, but not much for text files, which is where they are most needed. A promising way out is to handle each text as an image, rather than as ASCII file, and apply robust watermarks to it. The penalty of so doing is in an increase in the storage size of each text file; with today's low prices for hard disks of vast capacities, and also given today's rapidly dwindling costs of wideband communications, this should no longer be a real problem

CONCLUSION

Technology can, in fact, drastically reduce the "insider threat" problem to computer systems and networks, but only if it is preceded by a necessary change in the administrative level of access given so that no one person knows all security aspects protecting a computer system, nor is able to defeat them nor is able to erase all evidence of having done so.

The technological countermeasures involve a suite of solutions that should be used concurrently. Given the nature of complex software, however, where "bugs", such as buffer overflow schemes, have historically allows even total strangers to exploit weaknesses in the

code, and gain "root" access, the best one can hope for is to reduce the level of the "insider threat" to that of any outside hacker. This would be a major improvement, given that most computer system compromises of any consequence have been done by insiders.

Appendix

E. THE INSIDER THREAT TO INFORMATION SYSTEMS

Eric D. Shaw, Keven G. Ruby and Jerrold M. Post
Political Psychology Associates [1]

In the information age, as we have become increasingly dependent upon complex information systems, there has been a focus on the vulnerability of these systems to computer crime and security attacks, exemplified by the work of the President's Commission on Critical Infrastructure Protection. Because of the high-tech nature of these systems and the technological expertise required to develop and maintain them, it is not surprising that overwhelming attention has been devoted by experts to technological vulnerabilities and solutions.

Yet, as captured in the title of a 1993 conference sponsored by the Defense Personnel Security Research Center, *Computer Crime: A Peopeware Problem*, it is people who designed the systems, people who attack the systems, and understanding the psychology of information systems criminals is crucial to protecting those systems. [2]

- A Management Information Systems (MIS) professional at a military facility learns she is going to be downsized. She decides to encrypt large parts of the organization's database and hold it hostage. She contacts the systems administrator responsible for the database and offers to decode the data for \$10,000 in "severance pay" and a promise of no prosecution. He agrees to her terms before consulting with proper authorities. Prosecutors reviewing the case determine that the administrator's deal precludes them from pursuing charges.
- A postcard written by an enlisted man is discovered during the arrest of several members of a well-known hacker organization by the FBI. Writing from his military base where he serves as a computer specialist, he has inquired about establishing a relationship with the group. Investigation reveals the enlisted man to be a convicted hacker and former group member who had been offered a choice between prison and enlistment. While performing computer duties for the military, he is caught breaking into local phone systems.

- An engineer at an energy processing plant becomes angry with his new supervisor, a non-technical administrator. The engineer's wife is terminally ill, and he is on probation after a series of angry and disruptive episodes at work. After he is sent home, the engineering staff discovers that he has made a series of idiosyncratic modifications to plant controls and safety systems. In response to being confronted about these changes, the engineer decides to withhold the password, threatening the productivity and safety of the plant.
- At the regional headquarters of an international energy company, an MIS contractor effectively "captures" and closes off the UNIX-based telephonic switching system for the entire complex. Investigators discover that the contractor had been notified a week earlier that he was being terminated in part for chronic tardiness. Further investigation finds the employee to have two prior felony convictions and to be a member of a notorious hacker group under investigation by the FBI. The employee reports he is often up all night helping colleagues with their hacking techniques. Additional investigation reveals that he is the second convicted hacker hired at this site. An earlier case involved a former member of the Legion of Doom who had been serving as a member of a corporate information security team. He had been convicted of computer intrusion at a local phone company. Neither individual had disclosed their criminal history or had been subject to background checks sufficient to discover their past activities.

As these case summaries from the files of military and corporate security investigators demonstrate, growing reliance on information technology increases dependence on, and vulnerability to, those tasked with the design, maintenance and operation of these systems. These information technology specialists—operators, programmers, networking engineers, and systems administrators—hold positions of unprecedented importance and trust. Malevolent actions on the part of such an insider can have grave consequences. This is especially true for information

technology specialists operating within the critical infrastructure as identified in the 1997 President's Commission on Critical Infrastructure Protection's final report. [3]

These cases also demonstrate several points about the insider threat to the critical infrastructure. First, it is clear that insider problems already exist within the critical infrastructure, including the military, telecommunications, and energy sectors. Second, it appears that both inside and outside of our critical infrastructure, there is a tendency for managers to settle these problems quickly and quietly, avoiding adverse personal and organizational impacts and publicity. We do not really know how widespread the problems are. What is reported appears to be only the tip of the iceberg. Furthermore, we are at risk from repeat offenders, as perpetrators migrate from job to job, protected by the lack of background checks, constraints upon employers in providing references, and the lack of significant consequences for these offenses.

Finally, just as in organizations outside the critical infrastructure, the range of potential perpetrators and their motivations is broad. In many cases, acts of computer sabotage and extortion-like violence in the workplace have been committed by disgruntled employees who are angry about lay-offs, transfers, and other perceived grievances. Other cases involve employees who take advantage of their position of trust for financial gain, hackers who are employed within the critical infrastructure caught engaging in unauthorized explorations, and "well-motivated" employees who claim they are acting in the best interest of their organizations. [4] Other perpetrators include "moles," individuals who enter an organization with the explicit intent to commit espionage, fraud or embezzlement. Overall, case investigators report that the number of computer-related offenses committed by insiders is rising rapidly each year.

The extent of the insider threat has also been addressed in corporate and government survey results. According to WarRoom Research's *1996 Information Systems Security Survey*, 62.9 percent of the companies surveyed reported insider misuse of their organization's computer systems. The Computer Security Institute's *1998 Computer Crime Survey*

(conducted jointly with the FBI) reported the average cost of an outsider (hacker) penetration at \$56,000, while the average insider attack cost a company \$2.7 million. A comprehensive study conducted by the United Nations Commission on Crime and Criminal Justice which surveyed 3,000 Virtual Address Extension (VAX) sites in Canada, Europe and the United States, found that "By far, the greatest security threat came from employees or other people with access to the computers." While some researchers warn that survey data on computer crimes can be inaccurate due to unreported or undetected acts, such data are useful in characterizing a minimum level of threat and in drawing attention to the problem as a whole.

Paradoxically, in spite of the prevalence of the insider problem and the particular vulnerability of public and private infrastructures to the information technology specialist, there has been little systematic study of vulnerable insiders, while major investments are being devoted to devising technologies to detect and prevent external penetrations. Technological protection from external threats is indeed important, but human problems cannot be solved with technological solutions. Without a detailed examination of the insider problem and the development of new methods of insider risk management, such an unbalanced approach to information systems security leaves critical information systems vulnerable to fraud, espionage or sabotage by those who know the system best: the insiders.

RESEARCH IN PROGRESS

In response to the increasing recognition of the dangers posed by the insider threat to information systems, Political Psychology Associates, Ltd., under the auspices of the Office of the Assistant Secretary of Defense (Command, Control, Communications and Intelligence), have undertaken a study to improve understanding of the personality, motives and circumstances which contribute to information technology insider actions. By constructing psychological profiles of perpetrators and mapping their interactions with the organizational environment as they move over time toward the commission of violations, the goal of the study is to contribute to improvements in security, law

enforcement and counter-intelligence policies and practices. Specific applications for improving screening, selection, monitoring and management of information technology specialists are a primary goal of this research. The findings will also have implications for case investigation, information assurance audits, red team exercises, and information warfare.

THE CRITICAL INFORMATION TECHNOLOGY INSIDER

From the broad array of employees who have access to computers, we are focusing on the information technology specialists who design, maintain or manage critical information systems. Employees in this professional category are of particular concern because they possess the necessary skills and access to engage in serious abuse or harm. Typical jobs include systems administrators, systems programmers and operators and networking professionals. We are using the term Critical Information Technology Insiders (CITIs) to designate this professional category. [5]

Employment Contexts

The employment context is critical for understanding the relationship between the information technology specialist and the organization. The "insider-outsider" dichotomy is oversimplified, for in fact there is a spectrum of relationships between information technology specialists and organizations, which differentially affect loyalty and motivation.

Within the spectrum of "insiders," information technology specialists may serve as regular (full-time or part-time) staff employees, contractors, consultants or temporary workers (temps). In modern business practice, partners and customers with system access are also a source of exposure. In addition, former employees often retain sufficient access to the organization to remain an "insider" threat. Moles, information technology specialists who enter an organization with the intent to harm, are excluded from the current effort because they are potentially very different subjects from a psychological standpoint and present different screening and management problems. In this study we are primarily concerned with information technology specialists who develop their intent to harm the organization after being hired.

Employees (Full-Time and Part-Time)

Staff employees pose perhaps the greatest risk in terms of access and potential damage to critical information systems. As vetted members of the organization, employees are in a position of trust and are expected to have a vested interest in the productivity and success of the group. Considered "members of the family," they are often above suspicion—the last to be considered when systems malfunction or fail.

Among the several types of insider categories, organizations generally have the strongest influence and control over their own employees. To the extent that an employer is permitted by law to probe the background of a potential hire for security purposes, such investigations are much more likely to occur with prospective employees than with contractors, consultants, or temporary workers, whose roles in the organization are by design transient and who may or may not be vetted.

Employee CITIs who have caused damage have used their knowledge and access to information resources for a range of motives, including greed, revenge for perceived grievances, ego gratification, resolution of personal or professional problems, to protect or advance their careers, to challenge their skill, express anger, impress others, or some combination of these concerns. Three case examples serve to illustrate the employee threat:

Example 1: A senior MIS specialist at an international energy firm regularly created outages at Company sites around the world so that he could spend time abroad while gaining attention for his technical expertise.

Example 2: Michael Lauffenberger, a 31-year old programmer for the General Dynamics Atlas Missile Program, reportedly felt unappreciated for his programming work on a parts-tracking system. He planted a "logic bomb" in the system designed to erase critical data after he resigned. He then anticipated returning to rescue the company as a highly paid and valued consultant.

Example 3: Regional PC manager for the King Soopers supermarket chain Jay Beaman and two clerks were charged in an intricate computer fraud that cost the supermarket over two million dollars over two years.

The motives are described by investigators as beginning with financial necessity but quickly escalating into greed and ego. Among the strategies used was manipulating the computer accounting system to funnel certain purchases into a dummy account. At the end of the day, the perpetrators would take the amount funneled into the dummy account right out of the cash registers and then delete the account, also erasing any trace of their fraud.

In examples 1 and 2, the employees used their knowledge and access to a critical system to create crises, which would magnify their importance and worth within the organization. Jay Beaman was able to use his position to both commit and cover up his fraud, emphasizing the vulnerability of organizations to trusted employees.

Contractors, Partners, Consultants and Temps

Contractors, partners, consultants and temps are included as a category separate from employees because they are often not, in practice, subjected to the same screening and background checks. Moreover, a lesser degree of loyalty to the firm or agency would be anticipated. Many organizations within the critical infrastructure but outside the intelligence community have little control over the pre-employment procedures and hiring practices utilized by a contractor or consulting group. This is true even though contractors and consultants (and sometimes temps) often have highly privileged access to the organization's information assets due to the increase in outsourcing of programming and other information technology functions.

While the contracting organization is well within its rights to require contractors to screen the employees that will be working within the organization or provide a separate screening process for contracted employees, such steps are rarely taken, putting the organization at risk. The same goes for consultants and temps, though the transient nature of the consulting or temporary working relationship presents practical barriers to more rigid screening processes. The hiring of former hackers by some computer security consulting firms further increases the risk of security compromises. Employers have also

consistently underestimated the ability of contractors and consultants to take advantage of even limited access to important systems.

Example 4: A major international energy company recently discovered a logic bomb in software created by a contracted employee. It was installed as "job insurance" by the contracted employee with five prior convictions related to hacking. The contractor's firm failed to screen this employee who installed the code in anticipation of using it as leverage against his employer in case his criminal record was discovered.

Example 5: Zhangyi Liu, a Chinese computer programmer working as a subcontractor for Litton/PRC Inc., illegally accessed sensitive Air Force information on combat readiness. He also copied passwords, which allow users to create, change or delete any file on the network, and posted them on the Internet.

Example 4 illustrates the problems posed by poor screening measures and the vulnerability of organizations outsourcing their information technology functions. Example 5 demonstrates the espionage threat posed by contractors, though the motivations of this particular perpetrator are not yet clear. It also emphasizes the complex issues of loyalty in an international environment.

Former Employees

Former employees include individuals who no longer work at an organization but retain access to information resources directly -- through "backdoors" -- or indirectly through former associates. Anticipating conflict with an employer, or even termination, these perpetrators may prepare backdoor access to the computer system, alternative passwords, or simply stockpile proprietary data for later use. The number of cases in which separated employees have returned to extract vengeance on their former employers indicates a need for improved management of the termination process. This is particularly the case in episodes involving large numbers of layoffs. Such reductions can result in a pool of disgruntled employees and former employees with access and motivation for vengeance.

Example 6: Donald Burleson, a computer programmer for USPA & IRA Co., a Fort Worth securities trading firm, designed a virus after being reprimanded for storing personal letters on his company computer. The virus was designed to erase portions of the Company's mainframe and then repeat the process if a predetermined value was not reset in a specific location. After being fired, Burleson used a duplicate set of keys to return to the facility at 3 a.m. and employ an unauthorized backdoor password to reenter the system and execute the virus

INDISPENSABLE ROLE OF THE INSIDER

It is important to note that the efforts of "outside" groups (including foreign interests) could be aided significantly by the assistance of parties within the organization with access to, and knowledge of, critical information systems. For certain secure, self-contained systems, the insider's access will prove indispensable. Whether the insider is recruited directly, indirectly (e.g. "false flag" recruitment), coerced through blackmail, or through "social engineering" is manipulated while unaware that he is providing assistance to an adversary, his collaboration is a tremendous force multiplier. The potential damage an insider can now commit has also been increased within the last decade by two related trends in information systems -- consolidation and, for all intents and purposes, the elimination of the need-to-know principle. These changes, designed to improve information sharing, have removed obstacles to hostile collection. The hostile, sophisticated information technology professional now has many more opportunities to enter and damage larger systems. These vulnerabilities led one government information technology specialist, who focuses on system security, to refer to many allegedly secure government databases as "single point of failure systems."

Example 7: On the programming staff of Ellery Systems, a Boulder Colorado software firm working on advanced distributive computing software, was a Chinese national who transferred, via the Internet, the firms entire proprietary source code to another Chinese national working in the Denver area. The software was then transferred to a

Chinese company, Beijing Machinery. Ellery Systems was subsequently driven to bankruptcy by foreign competition directly attributed to the loss of the source code.

As illustrated by this case, the foreign connections of information technology specialists can increase their vulnerability to recruitment, manipulation, or independent hostile action.

PERSONAL AND CULTURAL VULNERABILITIES

Case studies and survey research indicate that there is a subset of information technology specialists who are especially vulnerable to emotional distress, disappointment, disgruntlement and consequent failures of judgment which can lead to an increased risk of damaging acts or vulnerability to recruitment or manipulation. Moreover, there are characteristics of the so-called "information culture" which contribute to this vulnerability. This report is not an attempt to cast suspicions on an entire professional category whose role in the modern computer-based economy has become so critical. However, we must better understand the motivations, psychological makeup, and danger signals associated with those insiders who do pose a threat to our information systems before we can really address this problem.

Reports of past research and our own findings based on interviews conducted so far, lead to the conclusion that there are several characteristics which, when found together, increase this vulnerability toward illegal or destructive behavior. These include: computer dependency, a history of personal and social frustrations (especially anger toward authority), ethical "flexibility," a mixed sense of loyalty, entitlement, and lack of empathy.

Introversion

According to a 1991 study by Professor Kym Pocius, the psychological testing of over fifteen hundred computer programmers, systems analysts, programmer trainees, and computer science students in seven separate studies consistently found these groups to be "overwhelmingly represented by introverts." Introverts differ from extroverts in being oriented toward the inner world of concepts and ideas rather than the outer world of people. They enjoy being alone,

prefer their own thoughts to conversation with others and may be socially unskilled. They also tend to be over-conscientious, secretive, pessimistic and critical. Authorities on the subject tell us that introverts are harder to distract than are extroverts, yet they are more reactive to external stimuli. According to H. J. Eysenck, a prominent personality psychologist, introverts tend to "shy away from the world while extroverts embrace it enthusiastically."

We wish to emphasize that, unlike the traits we are about to delineate, introversion is characteristic of computer technology specialists as a group, as well as scientists and other technology specialists. Indeed, some 40% of the overall population demonstrate this trait. One could not eliminate introverts from the ranks of computer technology specialists without eliminating the specialty. However, the preference for individual intellectual pursuits as opposed to interpersonal activity means that the signs of employee disaffection which would be apparent for extraverted employees may not be so readily visible. They may only occur, in fact, on-line, so the introvert poses challenges to management.

The following vulnerabilities have been identified in individuals who commit dangerous acts. They are associated with the vulnerable subgroup within computer technology specialists.

Social and Personal Frustrations

Surveys of computer professionals and computer science students indicate the presence of a subgroup whose entry into the field is motivated, in part, by frustrations getting along with others. According to a 1993 study by Professor R. Coldwell, this subgroup reports a history of conflicts and disappointments with family, peers and coworkers. They report preferring the predictability and structure of work with computers to the lack of predictability and frustrations of relationships with others. These experiences appear to have left them with a propensity for anger, especially toward authority figures. They also tend to be less socially skilled and more isolated than are their peers. Noting the high incidence of anger and alienation in these computer science students, Coldwell labeled it "revenge syndrome."

These traits create an increased vulnerability to feelings of alienation, disgruntlement, and disappointment on the job. Not only are such employees more likely to have innate antagonism for their supervisors, but they are less likely to trust and to deal directly with authorities when problems arise. In turn, these characteristics may also make some of these employees more vulnerable to recruitment and manipulation.

Computer Dependency

Two identified subgroups of computer users include individuals who exhibit an addictive-like attachment to their computer systems and those who manifest a similar attachment to the on-line experience offered by networks such as the Internet. Behavioral scientists studying these subgroups have found that they spend significantly more time on-line than is necessary for their work, frequently report losing any sense of the passage of time while on-line, and find that their on-line activities interfere significantly with their personal lives.

The "computer-addicted" individuals studied by researcher Margaret Shotten (1991) reported their primary interest as exploring networks, and viewed breaking security codes and hacking as honorable means of gaining emotional stimulation by challenging and beating security professionals. They did not consider pirating software unethical.

Computer dependents share a history of social failures and ostracization; and they admitted that the computer replaces direct interpersonal relationships. Their family histories include a high percentage of aloof, cool, and disinterested parents and authoritarian fathers. On formal psychological testing, this group contains a high percentage of well-informed, scientific, problem-solvers who enjoy intellectual pursuits. They are significantly more likely to be independent, self-motivated, aggressive loners, who make poor team players and feel entitled to be a law unto themselves. They reportedly tend to exhibit an unusual need to show initiative to compensate for underlying feelings of inadequacy.

Other researchers found that many members of the Internet-addicted subgroup are deeply involved in computer-mediated relationships,

including role-playing games. For many introverted, less socially skilled individuals, their computer-mediated social contacts are the least anxiety arousing of their interpersonal experience. In some cases, the sense of self, experienced on-line, becomes greatly preferred to the experience of self in the real world. Correspondingly, the on-line relationships of these individuals can displace affections and loyalties from real world ties. Noting the power of these relationships, many mental health professionals have characterized them as therapeutic building blocks that can help some people make the transition to subsequent real world contacts. However, for other more vulnerable individuals, these on-line relationships may also constitute an avenue for influence, recruitment or manipulation with security implications.

Ethical "Flexibility"

Concerns have been raised about looser ethical boundaries within the so-called "information culture." Surveys in recent years of current computer professionals indicate the presence of a subgroup whose members do not object to acts of cracking, espionage and sabotage against information resources. This subgroup appears to maintain the position that if an electronic asset, such as a limited access file, is not sufficiently secure, then it is fair game for attack. A disturbing aspect of these findings is the association between decreased ethical constraints and youth, suggesting that this perspective may be shared increasingly among new and future employees.

A number of social phenomena have been cited by several researchers as contributing to this dangerous trend. Lack of specific computer-related ethical training and lack of regulations within organizations have been implicated as contributing to lax employee ethical attitudes. Lack of similar ethical training in schools and at home by parents also contributes to this cross-generational trend. The boundary ambiguities of cyberspace, especially the lack of face-to-face connection, may also insulate perpetrators from the impact of their acts. The idea that exploring and even copying others' files inflicts no real damage has also been used to rationalize what would otherwise be considered privacy violations and theft in the outside world.

Finally, the computer industry has been implicated in the erosion of its own ethical standards. Some critics have suggested that the introduction of what they view as unrealistic and impractical restrictions on the use of purchased software produced contempt and disregard for these standards. Other critics suggest that the hiring and promotion of former hackers has sanctioned hacking and has even produced an incentive for this behavior.

Reduced Loyalty

Organizational loyalty among programmers and other professionals has been challenged increasingly by the high demand for their services and high rates of turnover in the profession. The resulting pressures to hire and retain computer professionals have also placed tremendous pressure on the security process.

Commenting on interviews with insider perpetrators of computer crime by the President's Council on Integrity and Efficiency, computer security expert Sanford Sherizan addressed the issue of distinct differences in programmer loyalty. Sherizan noted that there appear to be programmers who identify with the organization that pays them while others identify with the profession of programming itself. For these latter employees, their weak bond to the organization can lead to tensions in the workplace. Ambiguities about the "ownership" of intellectual properties in the form of source codes and other programs have also lead to a large number of conflicts between employers and computer professionals.

Entitlement

Our clinical investigations of vulnerable CITIs have consistently revealed two additional traits as risk factors, which have been alluded to but have not been emphasized. In assessments of CITI perpetrators from the energy and national security infrastructures, we have found that a sense of entitlement and anger at authority are consistent aspects of perpetrator motivation and personality.

A sense of entitlement, associated with the narcissistic personality, refers to the belief that one is special and owed corresponding recognition, privilege or exceptions from normal

expectations. This sense of "specialness" is often associated with a self perception of gifts or talents which are unrecognized by others. The perception that this specialness is not being recognized by authority figures often combines with a pre-existing anger at authority to produce feelings in these individuals that they have been treated unjustly and are entitled to compensation or revenge. Often, this sense of entitlement is supported by special arrangements or exceptions to rules granted to highly valued but "temperamental" MIS employees. Thus employers actually reinforce this belief, up the ante, and contribute to what often becomes an inevitable crisis. The current shortage of information technology personnel may also influence feelings of entitlement among older information technology employees, who may resent special treatment and bonuses paid to new hires.

According to a 1991 report by psychologists Robert Raskin and Jill Novacek, individuals with these narcissistic tendencies who are under higher levels of daily stress are prone to "power and revenge fantasies in which they see themselves in a powerful position able to impose punishment on those who have wronged them."

Our clinical sample helps validate a concern expressed by Coldwell about a group of programmers and computer science students who he characterizes as suffering from "revenge syndrome." Interviewees in this group appeared to present very similar perspectives and motives. As one interviewee in the previous study commented, when asked how he might utilize the power he was acquiring with his knowledge of programming, "I'll be getting my own back on the society that screwed me up."

Lack of Empathy

Disregard for the impact of their actions on others, or inability to appreciate these effects, has been a perpetrator characteristic noted consistently by investigators. It is also consistent with our clinical experience. Perhaps compounded by the impersonal layers of cyberspace, many computer perpetrators report never having considered the impact of their acts on other human beings. Many more appear incapable of placing themselves in their victim's shoes and imagining how the experience felt. This lack of empathy is a hallmark of individuals with

narcissistic and anti-social personalities, and is consistent with the traits of reduced loyalty and ethical flexibility.

Summary of Vulnerable CITI Personal and Cultural Characteristics

In summary, the research literature which we have surveyed identifies a coherent cluster of risk factors characteristic of a vulnerable subgroup of Critical Information Technology Insiders (CITIs). The negative personal and social experiences of a subgroup of information technology specialists tends to make them more vulnerable to experiencing the personal and professional frustrations which have been found to drive insider espionage and sabotage. Their social isolation and relative lack of social skills probably reduces the likelihood of their dealing with these feelings directly and constructively. Their reported vulnerability to ethical "flexibility," reduced loyalty to their employers, feelings of entitlement, anger at authority and lack of empathy probably reduces inhibitions against potentially damaging acts. At the same time, their loneliness, social naiveté and need to impress others may make them vulnerable to exploitation and manipulation.

The presence of any or all of these personal and cultural vulnerabilities does not, however, a perpetrator make. Indeed, it is more often the dynamic interaction between the vulnerable CITI's personal psychology (including the vulnerabilities enumerated above) and the organizational and personal environment that leads the vulnerable CITI down a slippery slope, at the end of which an act of information system aggression occurs. These critical pathways -- plural, for there are no set routes for the path to deviant, antisocial behavior -- that a CITI perpetrator might travel are being defined and explored further in the course of our research program.

What we do know already is that there is a complex interplay of personal and cultural or environmental factors which, over time, funnel an individual toward insider actions and that an understanding of this critical pathway has implications for personnel screening, monitoring, case management, and training. We also know that predisposing traits and situational factors are only part of the problem. What might be called acute situational stressors such as marital or family problems, episodes

of substance abuse, disappointments at work, threatened layoffs, or other stressful life events can trigger an emotional reaction leading to impaired judgment and reckless or vindictive behavior.

IMPACT OF INTERVENTION

Nevertheless, there are also mitigating forces that appear to reduce the likelihood of committing such acts or defuse a specific threatening situation. Highest on the list of mitigating factors is effective intervention by supervisors, co-workers, family members and close friends. Intervention might lead to counseling, involvement with support groups, or medical assistance. It is essential, however, that those who might intervene recognize and respond to significant warning signs and symptoms.

The Critical Pathway to Insider Espionage

A lucid description of the critical pathway to insider actions comes from Project Slammer, a major study of Americans convicted of espionage. Project Slammer mental health professionals conducted extensive interviews and formal psychological assessments with convicted perpetrators, most of whom were insiders. They also interviewed their coworkers, supervisors and families to identify not only the characteristics of perpetrators, but also the chain of events which led to their acts of treason. The results identified an interaction of factors, none of which alone was sufficient to result in an act of espionage. However, taken together and over time, these traits and experiences, common to many of the perpetrators, appear to have formed what we view as a common pathway to these acts. This pathway includes the following combination of events or "steps" which in some cases led to severe damage to national security:

- Predisposing Personal Traits
- An Acute Situational Stressor
- Emotional Fallout
- Biased Decision-making or Judgment Failures
- Failure of Peers and Supervisors to Intervene Effectively

As noted above, outside intervention is a critical mitigating factor on the path to insider acts. Unfortunately, in the insider espionage cases examined, it was often absent. Peers often assumed supervisors or others were aware of, and attending to, the problem. Supervisors often ignored the employee's problems, not wanting to deal with difficult individuals or not wishing to risk losing a valued member of the team. Often they attempted to manage the problem without considering the security risks involved. Sometimes the problem was pushed aside by transferring or firing the employee. It is interesting to note that a significant number of espionage offenders commit their acts after leaving their organizations. Abrupt termination does not appear to be a productive way to eliminate the security threat posed by such at-risk employees. Other supervisors incorrectly assumed that psychological referrals or on-going mental health counseling automatically took care of the problem and eliminated the risk of insider acts without requiring other intervention.

In the cases of destructive and criminal acts by vulnerable CITIs that we have analyzed to date, we are seeing a similar pattern in the sequencing of events. In a number of cases evaluated so far, we are confronted with examples of management failure to notice the problem, to accept the fact that a problem exists, or a willingness to tolerate dangerous behavior due to a desire to retain the services of a valued, technically competent employee. These findings have several implications for personnel management:

Pre-employment Screening

The critical path model views the probability of insider acts as the product of the interaction between predisposing traits, situational stressors and the organizational environment. Initial screening of employees should therefore emphasize the collection of information regarding traits, past and current behaviors (especially a criminal records check), and circumstances indicative of risk that is specifically tailored to the profile of the vulnerable CITI. Behaviors particular to the world of the computer professional should be central to this inquiry. Furthermore, successful screening will require that

human resources and information systems recruiters be sensitized to the factors contributing to CITI risk to guide them in the hiring process.

Improved Management of CITIs

Overall, the three most common management errors we have noted regarding CITI offenders have been (1) the failure to understand the personality and motivation of the at-risk employee; (2) the failure to have clear, standardized rules governing the use of company information systems with explicit consequences for misuse; and (3) the failure to punish rule violations. These problems often result in inadequate or even aggravating rules of conduct when constructive relief would be possible. Without organizational rules of conduct, employees have no guide to right and wrong and supervisors have no recourse to consequences when clear violations are discovered.

The company may also be held liable for illegal acts committed by employees in the absence of a well-defined and supported code of ethics. Solutions include specialized training for IT (information technology) managers to facilitate recognition of vulnerable CITIs and the selection of proper intervention techniques. The implementation of a comprehensive compliance program is also essential and should include a well-defined code of ethical behavior and support for employees facing ethical dilemmas or with questions regarding company policy.

Innovative Approaches to Managing At-Risk CITIs

For reasons discussed above, computer professionals present significant management challenges. In particular, monitoring their psychological state for risk using conventional observations is extremely difficult. As noted earlier, a subset of these individuals are likely to be more vulnerable to work-related stressors, while at the same time be much less likely to display overt signs of distress, complicating detection and delaying appropriate intervention by IT managers.

Compounding this problem is the shift of work-based communications toward computer-mediated communications in the workforce, a trend vastly accelerated among IT professionals in general, especially among those CITIs who find e-mail or chat rooms their preferred channel for

maintaining professional and personal relationships. The characteristics of the vulnerable CITI will inevitably require adapting traditional monitoring and intervention techniques to at-work electronic communications as the most effective means of understanding the psychological state and risk among these employees.

Innovative approaches for managing computer professionals include the creation of on-line environments designed to relieve work related stress by providing professional and constructive advice on dealing with problems in the office, e.g., on-line Employee Assistance Programs or job-stress hotlines. Electronic bulletin boards for logging anonymous complaints that can be monitored by management for purposes of addressing general grievances have also proven effective in some situations

One approach to effectively manage at-risk employees whose behavior has raised concern is to monitor their at-work electronic communications. This can be effectively used to detect changes in psychological state which warn of increased risk of destructive acts. While this approach raises privacy concerns, legal precedent has generally upheld the right of the employer to monitor their employees' use of company owned systems.

Comprehensive Information Security Audits

Finally, the critical path approach can also add a human element to the information security audit and its traditional emphasis on technological vulnerabilities and fixes. By reviewing the manner in which an organization selects, promotes, monitors, detects, manages and intervenes with problem CITIs, an investigator can gauge the organization's general sensitivity to insider risk and provide constructive solutions to managing the insider problem.

Only by adapting a comprehensive approach applying technological and human factors to information security can an organization adequately protect itself from both the outside threat of hackers and the more serious threat posed by the disaffected insider.

REFERENCES

1. This article is reprinted from *Security Awareness Bulletin* No. 2-98, published by Department of Defense Security Institute, September 1998. The research on which this article is based is part of a broader research program conducted by Political Psychology Associates, Ltd., for the Office of the Assistant Secretary of Defense (C3I).
2. Defense Personnel Security Research Center (PERSEREC) in Monterey, California, is now the Security Research Center of the Defense Security Service.
3. According to the PCCIP report, infrastructure is defined as "a network of independent, mostly privately-owned, man-made systems and processes that function collaboratively and synergistically to produce and distribute a continuous flow of essential goods and services." Critical components of the infrastructure, those affecting national security and the general welfare, include: transportation, oil and gas production and storage, water supply, emergency services, government services, banking and finance, electrical power, and information and communication infrastructures.
4. Our clinical experience indicates that seemingly simple cases of greed are rarely so simple when it comes to perpetrator motivation. Often there are other strong feelings and stressors behind the greed which complicate the motivational profile.
5. By definition, the term Critical Information Technology Insider (CITI) excludes the mass of end users who use computers as part of their jobs but for whom computers serve as a tool and not as a job in itself. While end users are associated with their own set of risks, we are specifically concerned with information technology specialists, whose job functions elevate them well above the average end-user in terms of skill, access and potential damage.

Appendix

F. THE INSIDER ESPIONAGE THREAT

Richards J. Heuer, Jr.

DSS/Security Research Center

To borrow a phrase from the former comic strip character, Pogo, "We have met the enemy, and he is us."

The initiative for most insider espionage comes from the insider, not from the foreign organization or group that receives the information. The overwhelming majority (about 75%) of Americans arrested for espionage during the past 20 years, and who had security clearance, were either volunteers who took the initiative in contacting a foreign intelligence service or were recruited by a close American friend who had volunteered to a foreign intelligence service. [1]

The initiative came from the foreign service in a comparatively small minority of the cases. It is difficult for foreign buyers of information to locate a willing American seller. They must proceed in secret, and with great care to avoid being caught, to identify one of the very few cleared Americans willing to betray their country.

It is easier for an American seeking to sell information to find a foreign buyer, although that, too, involves great risk. Twenty-six percent of the Americans arrested for espionage or attempted espionage during the past 20 years were caught by counterintelligence operations before they ever succeeded in compromising classified information, and 47% were caught during their first year of betrayal. [2]

Risk of betrayal of trust does not depend upon the presence of an implacable foreign adversary. It depends only upon an insider with the opportunity to betray, some combination of character weaknesses and situational stresses, and a trigger that sets the betrayal in motion. Common weaknesses include an arrogant attitude that the rules apply only to others, greed, impulsiveness, narcissism, feelings of entitlement, vindictiveness, alienation, paranoia, naiveté, and sensation-seeking.

There is reason to suspect that the number of insider spies today may be higher than in the past. One cannot know how many undiscovered spies are currently active or what the future will bring. Nevertheless,

we are not entirely in the dark when assessing the risk of undiscovered espionage. One can draw inferences from changes in American society and the international environment that may increase or decrease the propensity of cleared personnel to betray the Government's trust.

PRECONDITIONS FOR INSIDER BETRAYAL

As a general rule, four conditions must be present before a disaffected or troubled employee commits a serious betrayal of trust like espionage. The same conditions also apply to other insider crimes like embezzlement, sabotage, and procurement fraud, but those offenses are not discussed here. The four necessary preconditions for espionage are:

- An opportunity to commit the crime.
- A motive or need to be satisfied through the crime.
- An ability to overcome natural inhibitions to criminal behavior, such as moral values, loyalty to employer or co-workers, or fear of being caught.
- A trigger that sets the betrayal in motion.

The prevalence of these four conditions is influenced by changes in social and economic conditions in the United States and in our relations with the rest of the world. If the prevalence of these preconditions for espionage is increasing, the prevalence of insider betrayal may also be increasing. Analysis of changes in these preconditions for espionage gives some insight into what might be happening behind the scenes, without our knowledge, with respect to foreign espionage in the United States.

OPPORTUNITY

Opportunity is of two types:

- Access to information or materiel that can be exchanged for money or used to achieve some other goal.
- Personal acquaintance with, or easy access to, persons expected to be interested in obtaining such valuable information or materiel.

Starting with the widespread use of the Xerox copier in the 1950s, technological advances have made it increasingly difficult to control

the distribution of sensitive information. Today's large, automated databases and interconnected networks increase exponentially the amount of information that can be collected and compromised by a single, well-placed spy. Computer databases have greatly eased the spy's age-old problem -- how to purloin the exact information his or her foreign contact wants.

Opportunity equals temptation. It is now possible to commit crimes while sitting at one's computer engaged in what appears to casual observers as normal activity. More people have more access to more sensitive information than ever before. Like bank employees handling currency worth many thousands of dollars, not everyone is cut out to deal with that degree of temptation.

In today's increasingly open and interconnected world, it is also easier than in the past for an interested seller of information or materiel to find a foreign buyer. As compared with the Cold War days, there are many more countries to which a seller of information can turn in search of a buyer, but the risks are still great. In June 1996, the FBI had 800 open investigations of economic espionage involving 23 different countries. [3]

It is also dramatically easier for foreign intelligence services to take the initiative to spot, assess, and recruit knowledgeable Americans with exploitable weaknesses. The greatest change is in industry, where personnel involved in sensitive military R&D and production are increasingly in official business contact with their counterparts in foreign countries that are conducting espionage against the United States. The line between military and non-military, and between classified technology and unclassified technology sold to foreign countries, is increasingly blurred.

MOTIVE

When considering motives for espionage, it is useful to remember that the real motive may be different from the surface appearance. Although financial motivation is important, many people who commit espionage for money have more pressing emotional needs than financial needs. Espionage cases that appear to be financially motivated may

actually be motivated by out-of-control emotional needs. Money is valued not just for what it buys, but even more for what it symbolizes - success, power, influence and a route to happiness and self-esteem.

Espionage may also be an expression of power to influence events (satisfy a frustrated sense of self-importance), an outlet for anger (restore damaged self-image by outsmarting or punishing the bosses who failed to recognize one's talents), a means of revenge, or a source of excitement. It may also be motivated by divided loyalties or by an arrogant belief that one knows better than the U.S. Government what is in the best interests of the United States.

When looking at how social and economic changes in recent years affect motivation for espionage, two things stand out:

- Downsizing, outsourcing, transfer of jobs overseas, restructuring to adapt to the pressures of global economic competition, rapid technological change, and increased hiring of part-time workers to avoid paying benefits are all eroding many employees' sense of job security and loyalty to employer. At a minimum, this reduces the extent to which loyalty to employer inhibits misconduct. At worst, it provides a motive or rationalization for betrayal.

- About half of all the doctoral degrees in physics, chemistry and computer science granted by U.S. universities now go to foreign-born students. [4] One-third of all the engineers in Silicon Valley were foreign born. [5] This increasing internationalization of many high technology fields, combined with the increased number and variety of countries conducting intelligence operations against the United States, may increase the prevalence of conflicting loyalties.

REDUCED INHIBITIONS

Most personnel with access to classified information have the opportunity to betray, and many have a financial or other personal motive to do so. Betrayal is so rare only because it is deterred by basic moral values; loyalty to country, employer, or co-workers; and/or fear of being caught. Moral values, loyalty, and fear are the bedrock on which security is built. The stigma commonly associated with betraying

one's country also plays a role. Any social changes that erode these inhibitions to betrayal are likely to increase its frequency.

Morality is difficult to define and even more difficult to measure. This is not an appropriate place to pass judgment on the moral fiber of current American society from which our cleared personnel are drawn. Suffice it to note that the debate seems to be between those who see a serious degradation of moral values and others who view the state of morality in America as no worse than at other times in our history.

As noted under motives, loyalty is adversely affected by economic changes that devalue the long-term employer-employee relationship. Perceived inequities cause resentment. Feelings of entitlement to better treatment may be used to rationalize illegal behavior or may reduce inhibitions that otherwise deter illegal behavior. When people feel betrayed by their employer, it is easier for them to betray in return. Common rationalizations include: "I'm only getting back what they owe me." "It's their fault. They deserve it, because if they hadn't screwed me, I wouldn't be doing this."

The stigma of potentially being branded a traitor, or thinking of oneself as a traitor, also inhibits betrayal. This is somewhat diminished since termination of the Cold War ended the national "mission" to fight Communism and relieved the threat of nuclear holocaust. It is easier today for potential spies to rationalize the sale of classified information as a "purely business proposition" rather than a heinous activity that puts survival of country at risk. This is especially true when selling information to a "friendly" country or giving away information to a friendly country one wants to help.

The post-Cold War emergence of "friendly" countries as significant intelligence threats increases the prevalence of conflicting loyalties.

Although many people are honest because it's the right thing to do, others obey the law for fear of being caught. Fear of the unknown and fear of being caught are significant inhibitions to espionage, for the risk is indeed very high. There is no reason to believe that either fear has changed much in recent years.

TRIGGERS

Serious personal problems may fester indefinitely without leading to misconduct. The decision to betray will usually be triggered by some event in the individual's personal or professional life that pushes stress beyond that person's breaking point. The triggering event may be quite different from the underlying causes and motivation for betrayal.

Many people, perhaps most people, experience some form of stress that threatens their self-image at some time in their lives. They face serious financial problems combined with an available opportunity for illegal gain; failure to compete effectively with their peers; perceived injustice at the hands of an employer or supervisor; termination from a job under circumstances that prompt resentment; rejection or betrayal by a spouse or other close family member.

Emotionally stable and well adjusted individuals generally react to these experiences in positive ways--by learning from them, adjusting their expectations, working harder, or simply maintaining a stiff upper lip. Less stable or already troubled individuals sometimes react in ways that harm themselves or the organization. They may compound their problems by becoming less productive at work, turning to substance abuse or promiscuity, or attempting suicide. Or they may harm the organization by actions that range from absenteeism to self-serving decisions, theft, fraud, sabotage, or espionage.

There is no reason to believe the amount of stress in the lives of people in general is increasing. But many individuals do experience sharp changes in the amount of stress in their lives. The point is that stressful events are quite common, and that when they occur they can tip an otherwise weak, susceptible, or disturbed person over the edge.

SUMMARY AND CONCLUSIONS

The world is in the midst of an information revolution that many believe will have as far reaching an impact on politics, economics, and culture as that of the industrial revolution. It is surely affecting the manner in which nation states and other international actors compete economically as well as militarily, including the role of espionage in international competition and conflict. As a result of changes that have

already occurred in the domestic and international environment, the prevalence of insider betrayal may be greater today than during the Cold War.

Developments in information technology make it much harder to control the distribution of information. This greatly increases opportunities for espionage and the amount of damage that can be done by a single insider. A more open and interconnected world makes it easier for those interested in selling information to establish contact with willing buyers, as well as for those interested in buying information to spot, assess, and recruit willing sellers. Because U.S. national survival is no longer at stake since the end of the Cold War, personal interests are more likely than before to take precedence over national interests. It is easier to rationalize the sale of information to a "friendly" country as a "purely business proposition," rather than a heinous activity that puts survival of country at risk.

These social, economic and international trends may be creating uniquely fertile ground for the incubation and growth of espionage. They may infuse new vigor and intensity into the world's "second oldest profession," with the United States as the principal target.


REFERENCES

1. See Espionage by the Numbers. Information is from an unclassified database maintained by DSS/Security Research Center. For information on this database, see S. Wood & M. Wiskoff, Americans Who Spied Against their Country Since World War II. (Monterey, CA: Defense Personnel Security Research Center, 1992).
2. Ibid.
3. Geide, Kenneth (1996). "Economic Espionage: Looking Ahead." In Sarbin, Theodore (ed.), Vision 2021: Security Issues for the Next Quarter Century. Proceedings of conference sponsored by Defense Personnel Security Research Center and Security Policy Board Staff, June 25-25, 1996. Monterey, CA: Defense Personnel Security Research Center.
4. National Academy of Sciences (1995). Reshaping the Graduate Education of Scientists and Engineers. National Academy Press, p. 70.
5. Gilder, George (1995, Dec. 18) "Geniuses from Abroad," Wall Street Journal.

Appendix

G. INSIDER THREAT - A THEORETICAL MODEL

Ruth Duggan, Sandia National Laboratories



Insider Threat - A Theoretical Model

A Position Offered to the Insider Workshop
August 29-September 1, 2000


Ruth Duggan

Address:
PO Box 5800, MS 0449
Sandia National Laboratories
Albuquerque, NM 87185-0449



Email:
rduggan@sandia.gov

Phone:
505-844-9320

Program web site:
<http://www.sandia.gov/idart/>




Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

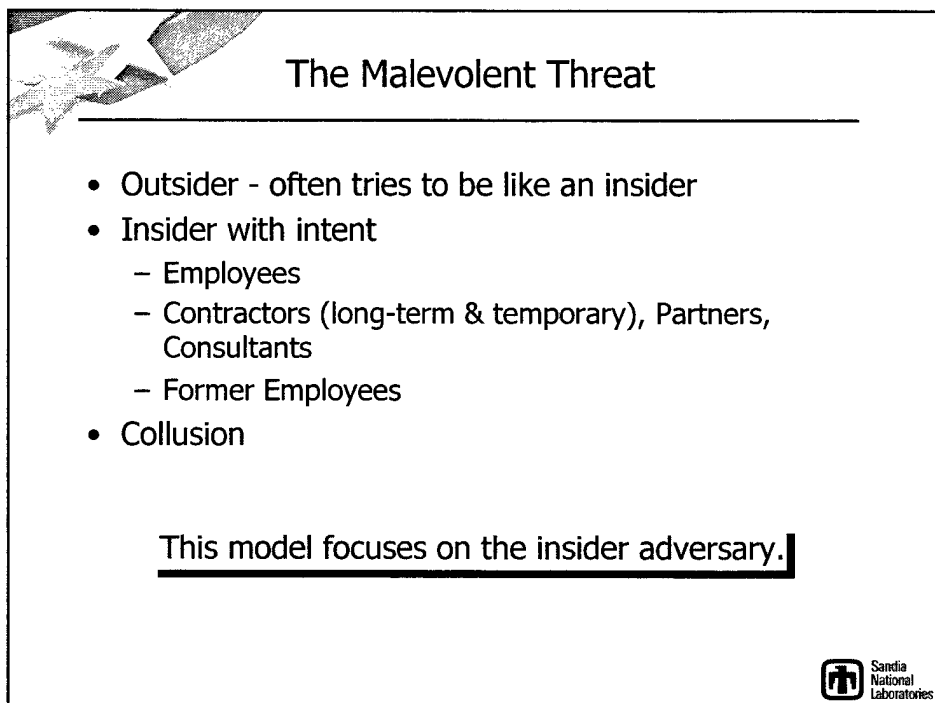
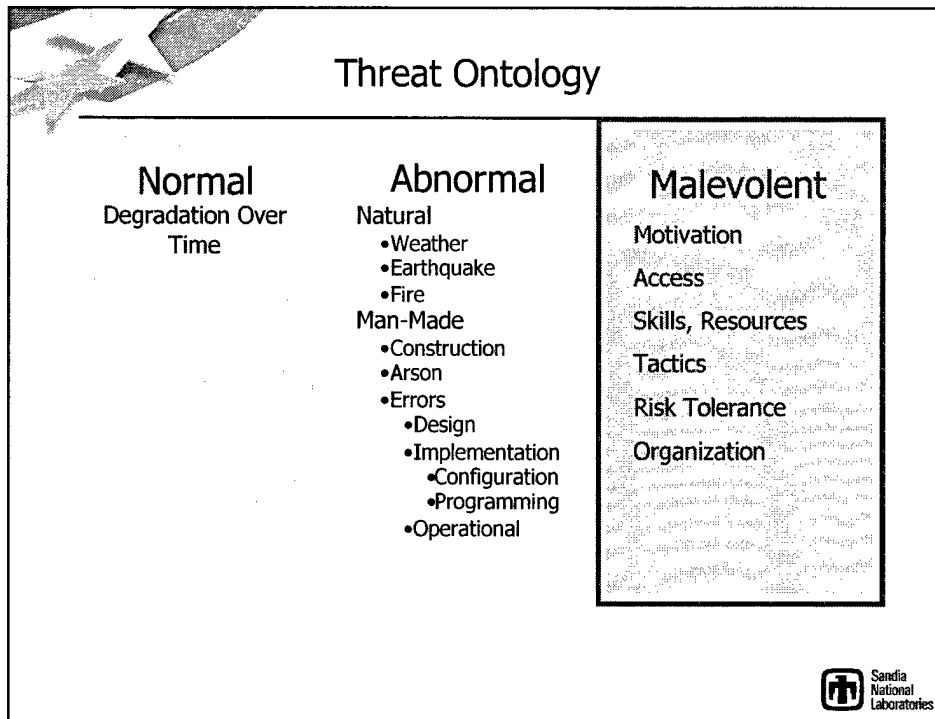



Why we use adversary models?

Primarily as a screening tool for attacks

- Attack space is large
- Certain adversaries may not be capable of different levels of attacks
- Not all adversaries have the same motivations or goals
- Help defenders plan and design better systems









Insider Definition

- Authorized user who performs unauthorized actions
- Levels
 - Physical access only
 - User level access
 - Privileged access

- Outsiders who obtain insider levels of access can operate like an authorized insider with that level of access.




How We Model Adversaries

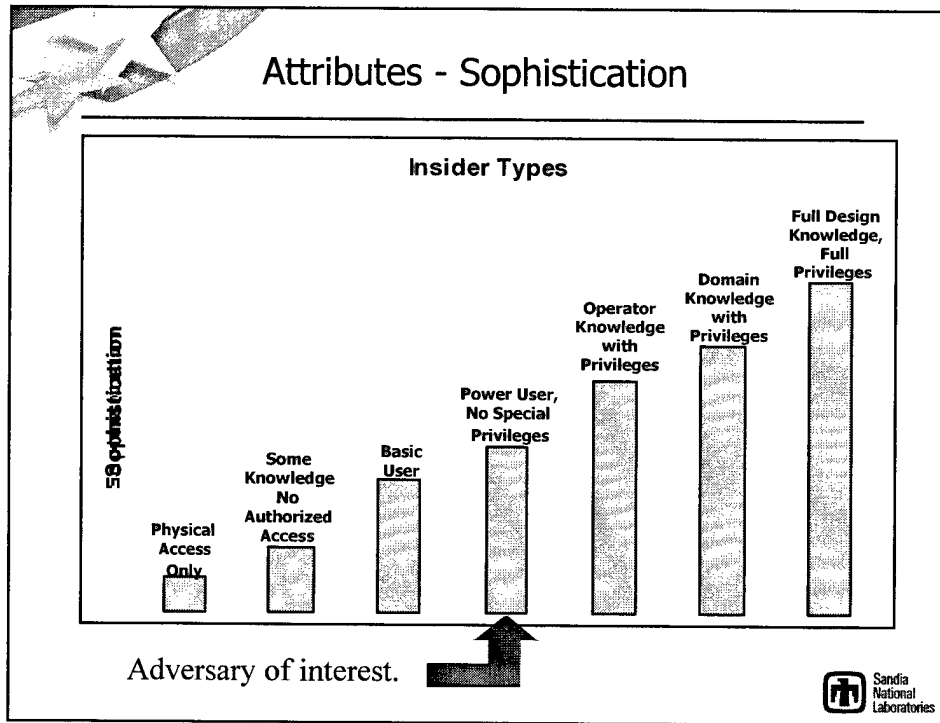
Variables in our models include:

- Level of sophistication - Access, Privilege, Knowledge
- Mission - What is the adversary's overall goal?
- Resources - Money & "magical powers"
- Risk Tolerance - How hard does the adversary avoid detection?

Assumptions used in models:


- Process
- Time






Mission - What will be attacked?



- What he knows about, probably outside their control
- Effects - nuisance through catastrophic






Resources

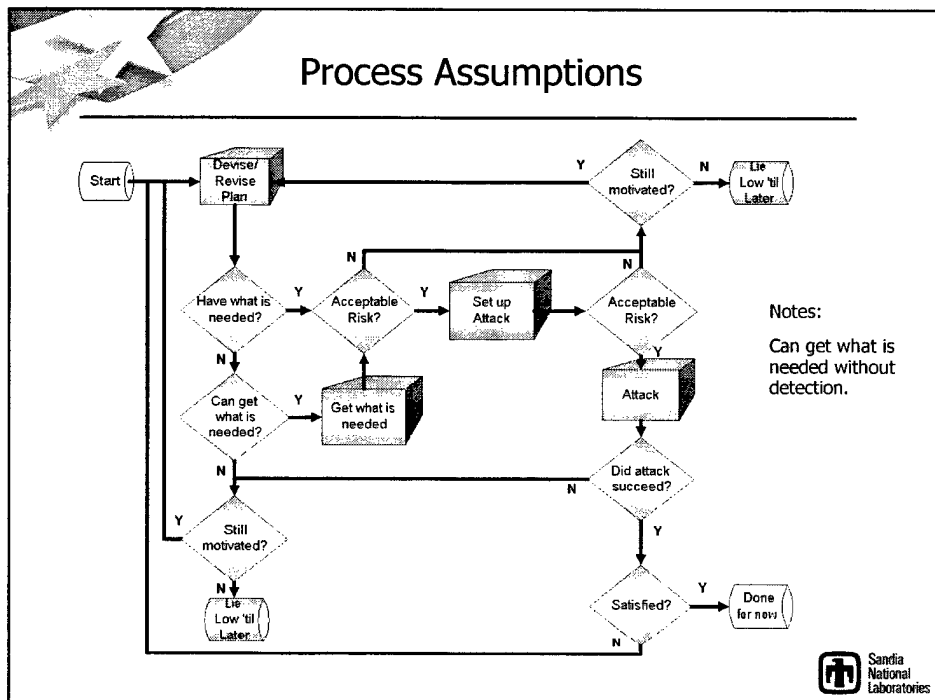
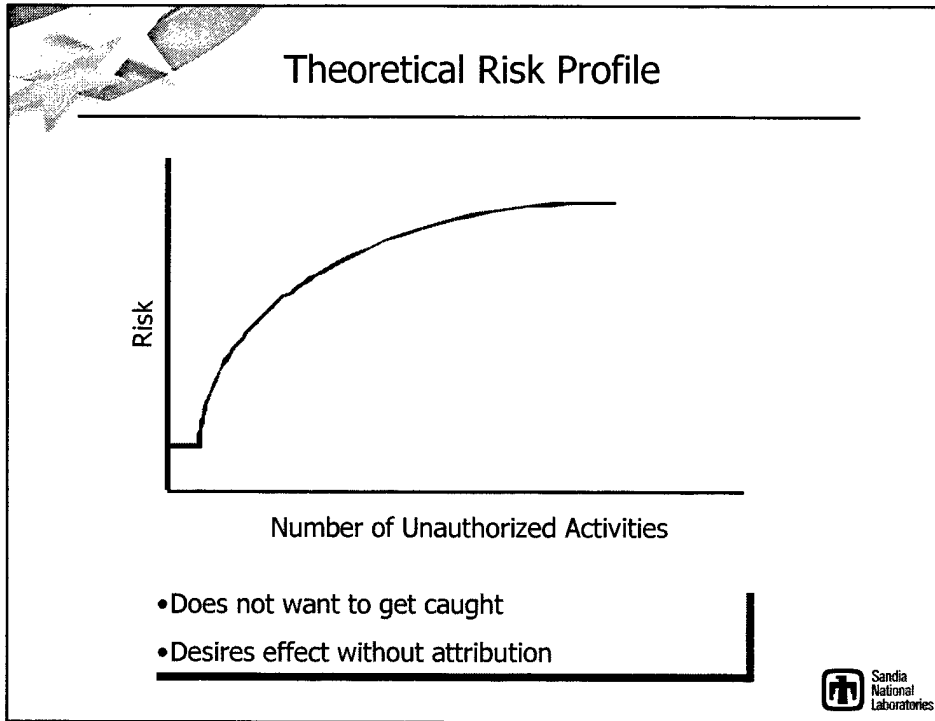
- Already has knowledge or has access to other insiders or inside information
- Works within knowledge base or what is anonymously available
- May or may not operate with a participating accomplice

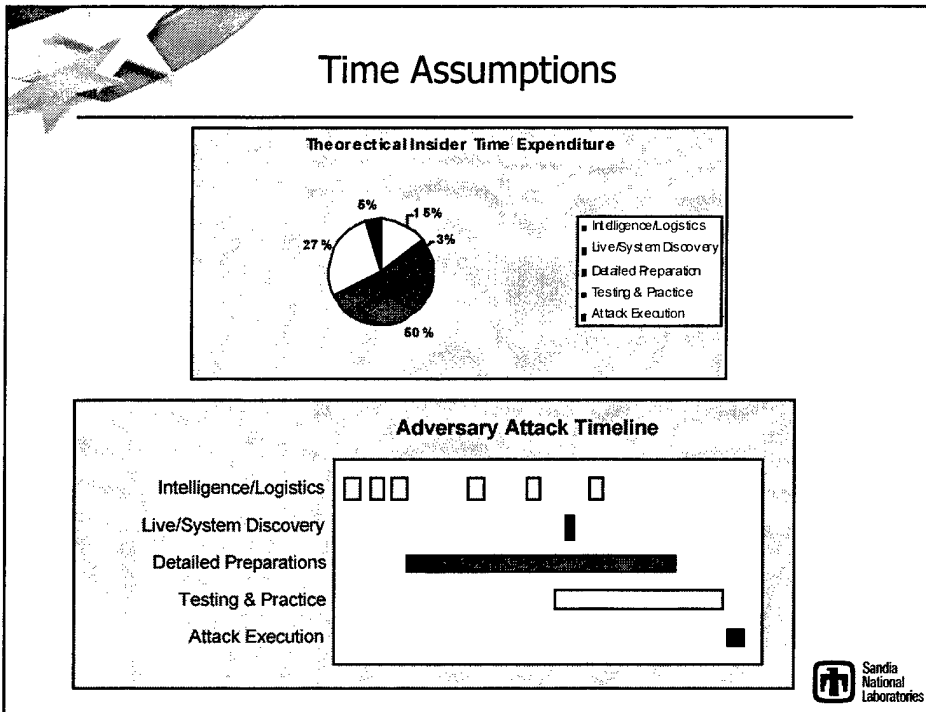


Methods


- Self-developed
- Attacks without attribution
- Use the Internet - exploits research
- Disinformation to supplement attack
- Take advantage of vulnerabilities known through insider status









- ### Basic Assumptions
- Inside to the level of power user, but not fully privileged
 - Can acquire some additional inside assistance if needed
 - Can afford to develop some attacks, but will take advantage of known vulnerabilities
 - Can influence software life cycle
 - Can learn most, if not all design information
 - Risk averse
 - Will employ quiet, stealthy attacks whenever possible.
 - Getting caught is generally unacceptable.
 - Specific targets and goals, based on opportunity
 - All open-source data
 - More likely to do anonymous trojan attacks, but other attacks possible...
- Sandia National Laboratories




Assertions

- The insider threat is recognized; however, the associated risk is often accepted while not being well understood
- Critical information systems are vulnerable
- A goal of outside adversaries is to look and operate like an insider to minimize detection



Conclusions

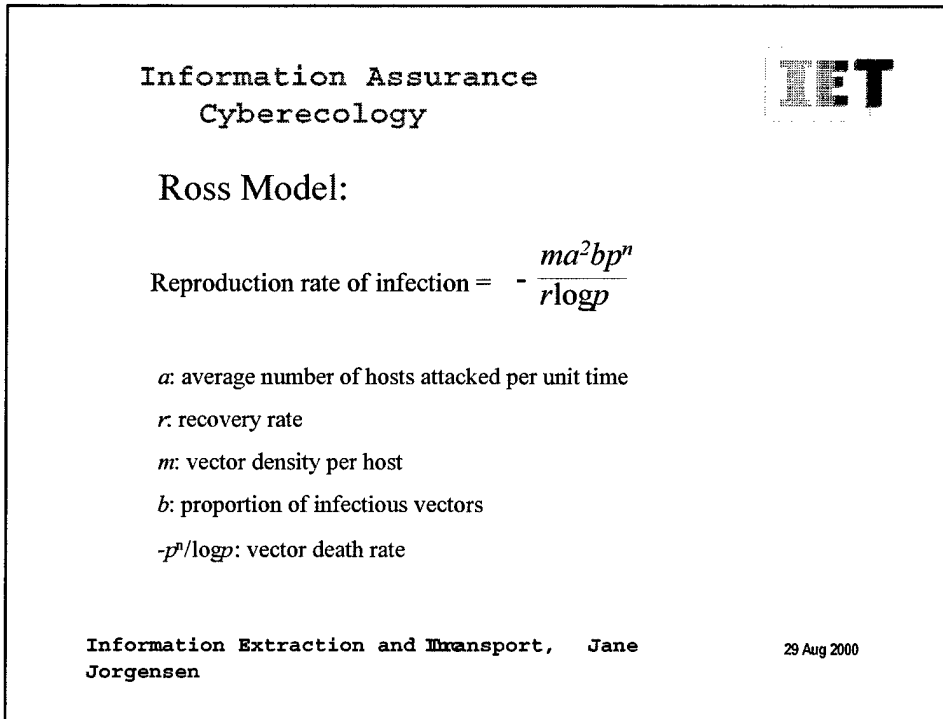
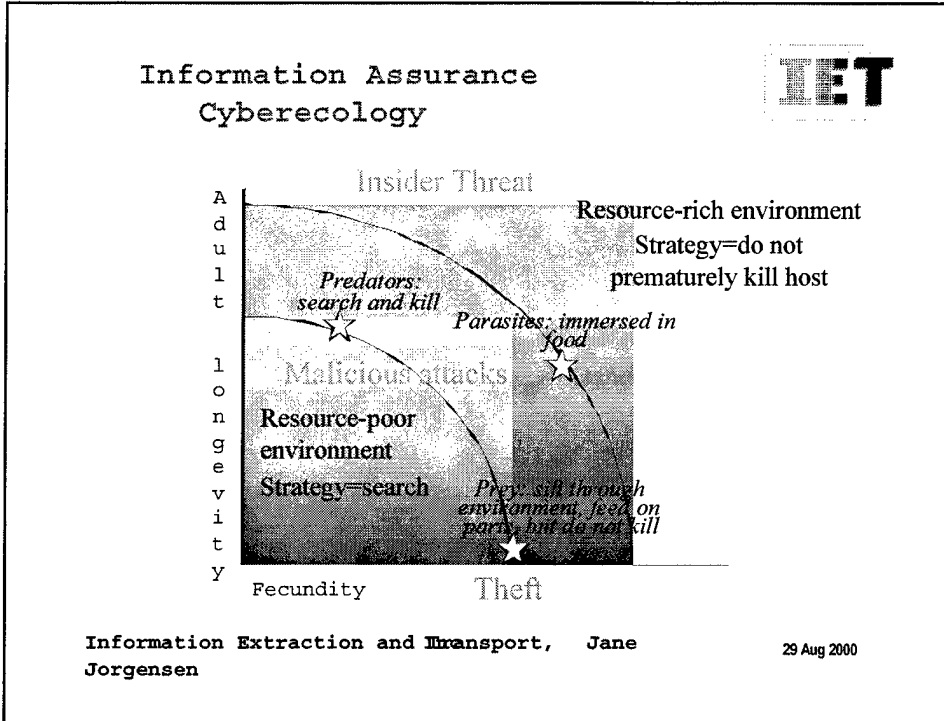
- The Insider Threat
 - is a recognized threat and exists at several levels
 - can be modeled
 - can be influenced
 - is unlikely to fall afoul of current detection and forensic methods



Appendix

H. INFORMATION ASSURANCE CYBERECOLOGY

Jane Jorgensen, Information Extraction and Transport Inc.



Information Assurance
Cyberecology



• Vector-borne diseases optimize different parameters:

- low n : Eastern equine encephalitis
- low r : onchocerciasis, Chagas
- high b : bubonic plague

• These parameters also define basis of control.

• Parallel parameterizations may apply to cyberagents.

Appendix

I. WORKSHOP AGENDA

Day 1 - August 30, 2000

0800	Registration / Continental breakfast		
0830	Plenary: Welcome by workshop sponsors (Michael Skroch, DARPA/ISO and Thomas Bozek, OSD/C3I); Introductions of all participants; Administrative details		
0915	Plenary: DoD Insider Threat Mitigation IPT Report and Hawaii workshop recap (Tom Bozek, OSD/C3I)		
0935	Plenary: August 1999 Insider Misuse workshop recap (Robert Anderson, RAND)		
1005	BREAK		
1020	Plenary: Insider Threats to Critical Information Systems: Typology of Perpetrators (Jerrold Post, Eric Shaw, Political Psychology Associates)		
1050	Plenary: Can Technology Reduce the Insider Threat? (Michael Caloyannides, Mitretek Systems)		
1115	Plenary: Tasking to the focus area groups: Purpose, approach, goals, and desired outcomes (Michael Skroch, DARPA/ISO and Thomas Bozek, OSD/C3I)		
1130	LUNCH		
1230	Focus 1: Long-Term Research Chair: Tom Longstaff	Focus 2: Insider Threat Models Chair: Wayne Meitzler	Focus 3: Near-Term Solutions Chair: Ken Van Wyk
1445	BREAK		
1500	Focus 1 (cont.)	Focus 2 (cont.)	Focus 3 (cont.)
1600- 1700	Plenary: Brief status reports by focus area leaders		

Day 2 - August 31, 2000

0800	Continental breakfast		
0830	Plenary: Guidance to focus groups (Michael Skroch, DARPA/ISO and Tom Bozek, OSD/C3I)		
0845	Focus 1 (cont.)	Focus 2 (cont.)	Focus 3 (cont.)
1030	BREAK		
1045	Focus 1 (cont.)	Focus 2 (cont.)	Focus 3 (cont.)
1200	LUNCH		
1300	Focus 1 (cont.)	Focus 2 (cont.)	Focus 3 (cont.)
1500	BREAK		
1515	Focus 1 (cont.)	Focus 2 (cont.)	Focus 3 (cont.)
1615- 1700	Plenary: Brief interim results by focus area		

Day 3 - September 1, 2000

0800	Continental breakfast		
0830	Plenary Speaker: Bill Leonard, Deputy Assistant Secretary of Defense for Security & Information Operations		
0900	Plenary: Briefing of final conclusions and recommendations by focus group leaders		
1145- 1200	Plenary: Summary and conclusions by workshop sponsors (Michael Skroch, DARPA/ISO and Tom Bozek, OSD/C3I)		

Appendix

J. WORKSHOP PARTICIPANTS

Sponsors

Michael Skroch	DARPA/ISO	mskroch@darpa.mil
Tom Bozek	OSD/C3I	tom.bozek@osd.mil

Support

Robert Anderson	RAND	Robert_Anderson@rand.org
Trina Labbe	Avenue Technologies	trina.labbe@avenuetech.com
Steve Sonnenberg	Avenue Technologies	steve.sonnenberg@avenuetech.com

Participants

Robert Anderson	RAND	Robert_Anderson@rand.org
Lee Badger	NAI Labs	lbadger@nai.com
Annette Benbow	CIA	aneterb@ucia.gov
Tom Bozek	OSD/C3I	tom.bozek@osd.mil
Michael Caloyannides	Mitretek Systems	Michael.Caloyannides@Mitretek.org
Ruth Duggan	Sandia National Labs	rduggan@sandia.gov
John Edwards	DIA	AFedwjm@dia.osis.gov
Lynn Fischer	Defense Security Research Center	fischelf@osd.pentagon.mil
Anup Ghosh	Reliable Software Technologies	anup.ghosh@computer.org
Paul Hulseberg	FBI/NIPC	phulseberg@fbi.gov
Terry Johnson	Army Research Lab	tjohnson@arl.army.mil
Jane Jorgensen	Information Extraction & Transport, Inc.	jorgenj@iet.com
Emily Joyce	NSA	edjoyce@alpha.ncsc.mil
Carl Landwehr	Mitretek Systems	carl.landwehr@mitretek.org
Susan Lee	Johns Hopkins U. Applied Physics Lab	sue.lee@jhuapl.edu
Lanark Lockard	Joint Task Force - Computer Network Defense	lockardl@jtfcnd.ia.mil
Tom Longstaff	Carnegie Mellon S/W Engineering Institute	tal@cert.org
John Lowry	BEN Technologies /Verizon	jlowry@bbn.com

Sylvia Mapes	DOD-CERT/DISA	stm@cert.mil
Sara Matzner	U. Texas, Applied Research Laboratories	matzner@arlut.utexas.edu
Roy Maxion	Carnegie Mellon U.	maxion@cs.cmu.edu
Wayne Meitzler	Pacific Northwest National Lab	wayne.meitzler@pnl.gov
Lynette Millett	National Research Council	lmillett@nas.edu
David Mitchell	Avenue Technologies Inc	david.mitchell@avenuetech.com
Robin Morel	Los Alamos Nat'l Lab	morel@lanl.gov
Peter Neumann	SRI International	Neumann@CSL.SRI.com
Jerrold Post	Political Psychology Associates, Ltd.	jmpost@pol-psych.com
Ron Schmucker	Lawrence Livermore National Lab	schmucker1@llnl.gov
Eric Shaw	Political Psychology Associates, Ltd.	eshaw@pol-psyych.com
Michael Skroch	DARPA/ISO	mskroch@darpa.mil
Steve Sonnenberg	Avenue Technologies	steve.sonnenberg@ avenuetech.com
Richard Szafranski	Toffler Associates	rsz@toffler.com
Maggie Vargas	DIA	AFvarmk@dia.osis.gov
Chuck Watterson	DTRA	chuck.watterson@dtra.mil
Brian Witten	DARPA/AIA	bwitten@darpa.mil
Ernest Wohnig	DIA	AFwohew@dia.osis.gov
Bradley Wood	SRI International	bradley.wood@sri.com
Ken Van Wyk	Para-Protect, Inc.	ken@para-protect.com
Lee Zimmerman	SPAWAR SSC SD	zimmer@spawar.navy.mil

REFERENCES

Anderson, R.H., 1999. "Research and Development Initiatives Focused on Preventing, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems: Results of a Three-Day Workshop." RAND CF-151-OSD, 1999. Available at: <http://www.rand.org/publications/CF/CF151>.

Anderson, R.H., R. Brackney, T. Bozek, 2000. "Advanced Network Defense Research: Proceedings of a Workshop" RAND CF-159-NSA, 2000. Available at: <http://www.rand.org/publications/CF/CF159>.

Insider Threat Integrated Process Team, Department of Defense (DoD-IPT), 2000. "DoD Insider Threat Mitigation" U.S. Department of Defense, 2000. Available at <http://www.c3i.osd.mil/>.