This product is part of the RAND Corporation conference proceedings series. RAND conference proceedings present a collection of papers delivered at a conference. The papers herein have been commented on by the conference attendees and both the introduction and collection itself have been reviewed and approved by RAND Science and Technology.

CONFERENCE PROCEEDINGS

Understanding the Insider Threat

Proceedings of a March 2004 Workshop

Richard C. Brackney, Robert H. Anderson

Prepared for the Advanced Research and Development Activity

NATIONAL SECURITY RESEARCH DIVISION

The work described here was conducted in the RAND National Security Research Division, which conducts research and analysis for the Office of the Secretary of Defense, the Joint Staff, the Unified Commands, the defence agencies, the Department of the Navy, the U.S. intelligence community, allied foreign governments, and foundations. These proceedings were supported by the advanced information research area in the Advanced Research and Development Activity within the U.S. intelligence community.

ISBN 0-8330-3680-7

Published 2004 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
201 North Craig Street, Suite 202, Pittsburgh, PA 15213-1516
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; Email: order@rand.org

Summary

A major research thrust of the Advanced Research and Development Activity (ARDA) of the U.S. intelligence community (IC) involves information assurance (IA). Perhaps the greatest threat that IA activities within the IC must address is the "insider threat"—malevolent (or possibly inadvertent) actions by an already trusted person with access to sensitive information and information systems.

This unclassified workshop, held March 2–4, 2004, focused on the insider threat and possible indicators and warnings, observables, and actions to mitigate that threat. The ARDA researchers participating gave special attention to the activities, processes, and systems used within the intelligence community.

A combination of plenary and breakout sessions discussed various aspects of the problem, including IC system models, vulnerabilities and exploits, attacker models, and characterization of events associated with an insider attack. A set of presentations by members of the IC and its contractors on Intelink (Appendix G) and such research activities as the development of "Glass Box" software (see Appendix H) and ARDA's "Novel Intelligence from Massive Data" (NIMD) research program (Appendix I) aided the workshop discussions. The present workshop built upon the availability of materials generated in an earlier workshop focused on the insider threat (Appendix F).

Several overall themes emerged from these deliberations, discussed below under the headings of "Research Questions and Challenges" and "Databases Needed" (by researchers).

Intelligence Community System Models

The overall intelligence process involves requirements, collection, processing and exploitation, analysis and production, dissemination, and consumption, with feedback loops at all steps, as shown in Figure S.1.

Variant models, such as the NSA Reference Model (NRM), also exist. Of key concern to this group of researchers was the question: What "observables"[1] can be obtained at all stages of this process that would allow comparison of normal analyst activity with abnormal activity—which is potentially, but not necessarily, malevolent? Figure S.2 provides an indication of the richness of the concept of "observable"; it is a taxonomy developed by the earlier insider threat workshop cited above. Similar taxonomies characterize IC "assets" and "users."

[1] An *observable* is anything that can be detected with current technology. A number of workshop participants argued that this definition should be broadened to include foreseeable future technological developments.

Figure S.1
Intelligence Process

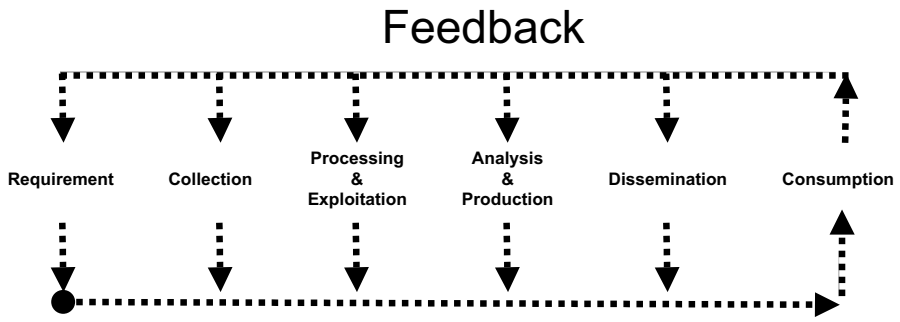

Figure S.2
Taxonomy of Observables

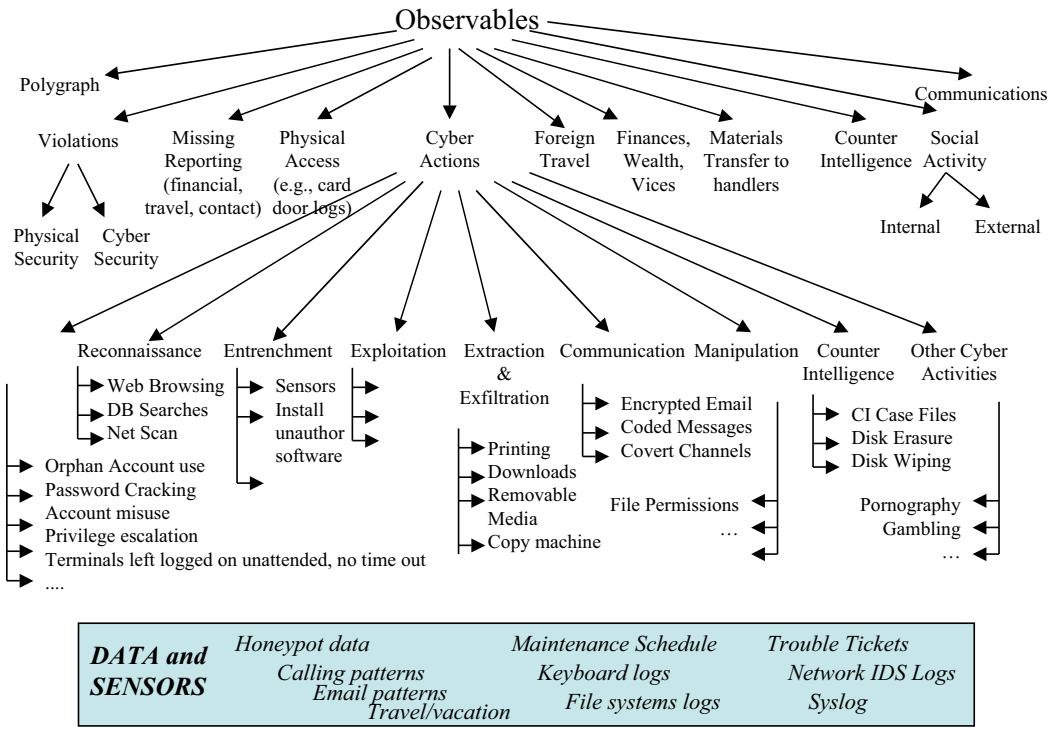

Vulnerabilities and Exploits

What types of exploits[2] might an insider use to obtain information, alter its integrity, or deny its availability to those who need it? This workshop concentrated on cyber-related

[2] The noun *exploit* is often used within the intelligence community to mean the development of a plan (and, usually, its subsequent execution—often surreptitiously) to obtain information or an advantage.

exploits because they were felt to be potentially the most damaging and most likely to increase in the future, as a new generation of analysts emerges with more computer skills than the previous generation.

Workshop participants generated a list of 33 example exploits. For each they listed a brief description, preconditions that would allow the exploit to happen, observables that might be generated during the exploit, and effects of the exploit (usually one of the following: a breach of confidentiality, integrity, or availability, or an enabler of other exploits). The short titles of the vulnerabilities are listed in Table S.1. Further details may be found in Chapter Three.

Attacker Models

Figure S.3 shows an overall model of the steps involved if a malevolent insider were to "mount an attack" against an IC asset. The attack might be as simple as obtaining access to information he or she does not have a need to know or as complex as disabling a key intelligence collection/processing/dissemination system.

Another way of depicting attacker actions is shown in Figure S.4. Here the attacker steps—motivation, benefit/risk assessment, acquiring the "client," collecting payment—were

Table S.1
Vulnerabilities and Exploits

1. Virus-laden CD and/or USB flash drive and/or floppy	18. Mislabeled paper
2. Administrator lockout	19. Netmeeting/WebEx controls
3. Social engineer passwords	20. "Day zero" attacks based on source code availability
4. Retry Internet attacks	21. Covert channels through steganography[a]
5. Smuggling out USB flash device or other media (exfiltration)	22. Copy and paste between classifications (from high to low)
6. "Missing" laptops/hardware	23. Internal e-mail that performs attacks
7. Targeted acquisition of surplus equipment	24. Wireless telephone cameras to capture information
8. Unpatched systems	25. Telephone tap recording onto removable media
9. Sabotaged patches	26. Telephone tap via hacking PBX telephone controller
10. False positives on anti-virus	27. Analyst changes workflow to exclude other analysts (dissemination)
11. Use of unattended terminal	28. Analyst changes workflow to include himself/herself
12. Targeting database "adjustments"	29. Insert bad content into report upon inception (e.g. translation)
13. Install software on host computer to capture keystrokes logger	30. Delete/withhold content into report upon inception
14. Extra copy of DB backups	31. Redirect analyst resources to support adversary's agenda
15. Wireless transmissions	32. Poor quality analysis/results/ reports
16. Cell phone/PDA/voice recorder in classified meeting	33. Get IC asset to collect info that benefits an unauthorized party
17. Suspicious activity on real systems (e.g., searching own name in databases)	

[a]Steganography is the hiding of information by embedding in an innocuous message or file, such as a digitized picture.

Figure S.3
Spiral Model Flowchart

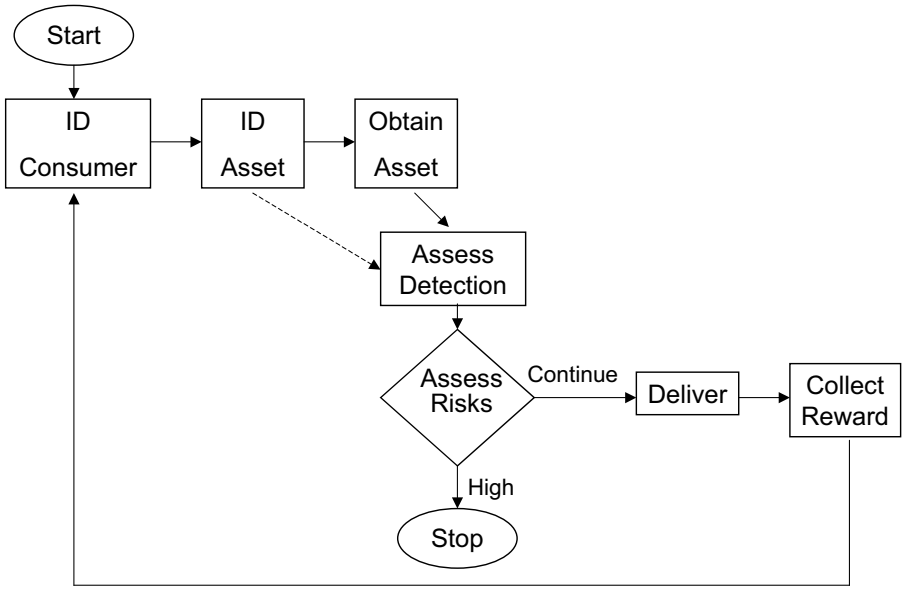

Figure S.4
Insider Attack Actions (white items not cyber observable)

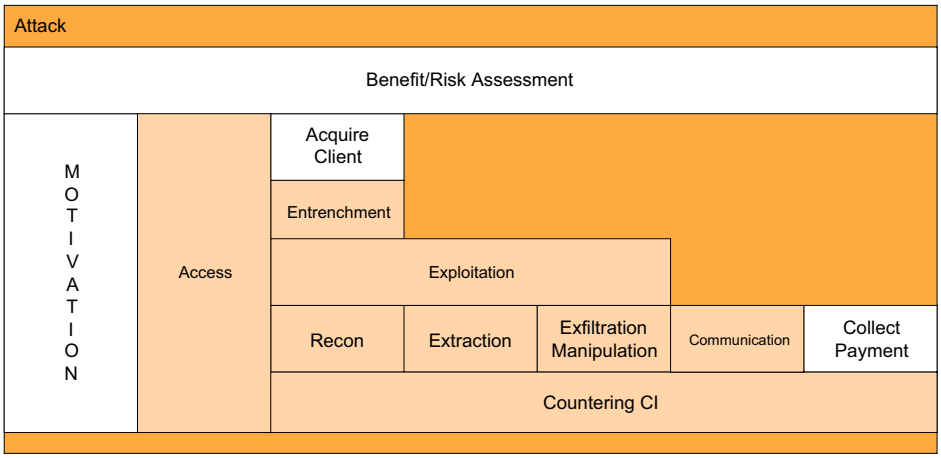

deemed *not* to generate cyber observables (that is, they would not be detected by information systems now in use or with enhancements planned by researchers and developers).

Given the various steps an attacker follows, as shown in Figure S.4, which steps are candidates for using the vulnerabilities and exploits shown in Table S.1? The answer is shown in Figure S.5, where the unitalicized insider actions have parenthesized numbers linking them to numbered entries in Table S.1. The parenthesized suffix letters C, I, A, E indicate whether the actions would lead to a breach of information Confidentiality, Integrity, Availability, or would be an Enabler of other attacks.

Figure S.5
Insider Actions Taxonomy Cross-Referenced with Vulnerabilities and Exploits (V&E) List

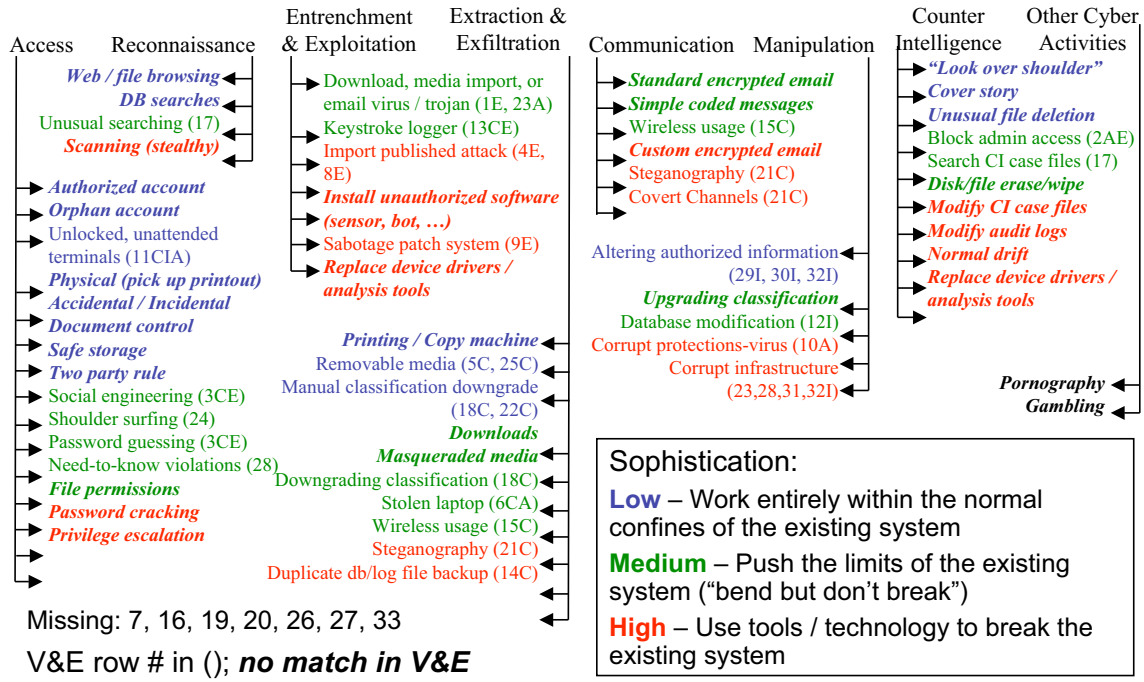

Event Characterization

As attacker actions generate observables through the operation of "detectors" of those observables, indicators of possible abnormal activity are generated. Those indicators can form a report; multiple reports can be fused into an "incident"; and multiple incidents then fused into a "case" of one or more incidents.[3] That process is shown graphically in Figure S.6.

Research Questions and Challenges

Each breakout group tried to formulate a set of research questions arising from its deliberations. Some groups stated these questions in the form of "grand challenges" to be addressed. We summarize the key questions and challenges below.

Six Categories of Research Questions

Research issues tended to fall within six categories:

1. User roles
2. Actions

[3] We assume that a *"case"* may be merely a collection of incidents having some commonality to be watched, or it could be the result of a post-facto analysis of source, cause, damage, etc.

Figure S.6
Data Collection Steps Regarding an Event

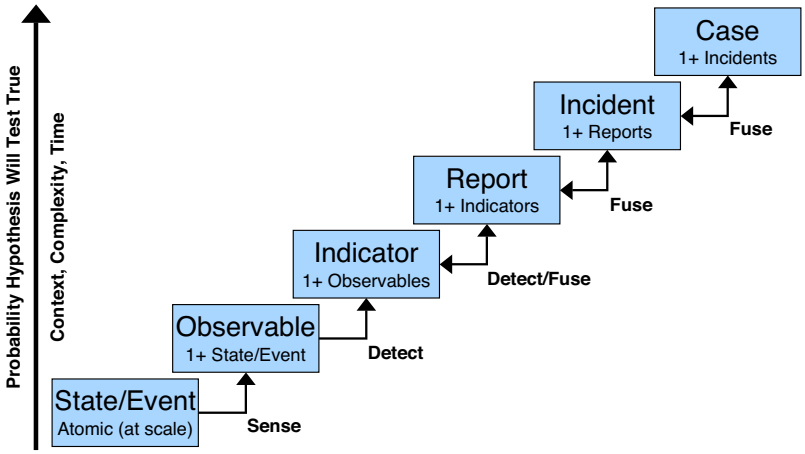

3. Observables (events)
4. Sensors
5. Fusion and analysis (both spatial and temporal)
6. "Triggers" (priorities, and level of certainty).

The first four categories each require *languages to describe them*, and *means for mapping each into the next* (i.e., from a description of user roles to a set of described user actions, which in turn lead to a set of potential observables. Those observables are then sensed and the sensed signals fed into fusion and analysis programs, which in turn create actions and alerts within the system).

An additional common thread is the need for *correlation and management tools* to correlate multiple events or triggers with an incident, to correlate multiple events with a case, and to correlate multiple cases into a coordinated attack.

The topic of sensors (item 4 in the above bulleted list) requires substantial research in at least the following areas:

- Identification of information that should go into an event record
- Development of sensors specific to particular applications
- Standardization of event record syntax and semantics; scales of severity and confidence; system interfaces; and means for establishing an inviolate "chain of evidence"
- Detection of "low and slow" attacks
- Optimization of selection, placement, and tuning of sensors
- Tradeoffs in adaptability: How do you recognize legitimate changes in user behavior? How do you resist the "conditioning" of sensors by a malicious insider (through a pattern of actions that "migrate" the sensor from a nominal setting to one that won't recognize the attack)?
- Development of validation and test data and techniques (see "Databases Needed," below).

Challenges

Participants stated several "grand challenges" for researchers:

- Define an *effective way of monitoring* what people do with their cyber access, to identify acts of cyber espionage. Focus on detection, not prevention. Such monitoring (or the perception of monitoring, which may suffice in some cases) can be an effective deterrent.
- Develop *policies and procedures* to create as bright a line as possible between allowed and disallowed behaviors (i.e., reduce the ambiguity).
- Consider *sociological and psychological factors* and create better cooperation between information systems personnel and human resources personnel (including security, medical, financial, and other support services). In short, broaden oversight of all aspects of a user's background and behaviors.
- *Combine events from one or more sensors* (possibly of various types or different levels of abstraction) to facilitate building systems that test hypotheses about malicious insider (MI) activity, to detect MI activity that is not detectable using a single event record, to develop a "calculus of evidence," to develop metrics for comparing and weighting diverse inputs, and to determine how "this fusion" can be used to create useful synthetic/compound events.

Databases Needed

Breakout sessions considered what databases would aid in this research if they were available. Researchers need databases containing examples of specific attacks, the characterization of normal behavior for users in different roles (including that of a system administrator), and artificial or real sensor data that include a mix of legitimate and malicious activity. Potential sources for the development of such datasets include a MITRE dataset of normal, and "insider threat" network activities; data from the ARDA NIMD[4] study; data obtained from use of the Glass Box[5] software; synthetically generated data from a simulator; and individual datasets developed by researchers that might be traded among projects.

A Concluding Remark

During a concluding plenary session, a senior member of the intelligence community, hearing the results from the various breakout session deliberations, made the comment, "What you're doing is important, but don't forget that IC analysts are people, too, and need a good work environment in which to stay motivated in their stressful jobs. When considering 'observables' and sensors and other means of keeping track of the activities of 'insiders,' please ask yourselves, 'Would I want to work in that (resulting) environment?'" It's important to keep this in mind, in the research enthusiasm for what *might* be monitored, and observed, and data-correlated. We must strike a balance between effectiveness in thwarting

---

[4] See Appendix I for information about the ARDA "Novel Intelligence from Massive Data" (NIMD) research thrust.

[5] See Appendix H for information about the "Glass Box" research effort.

insider exploits against intelligence assets and effectiveness in the process of generating and disseminating that intelligence information itself.