

Employing Deception in INFOSEC

Scott Gerwehr & Robert H. Anderson

An Ongoing Research Effort in RAND's National Defense Research Institute

BACKGROUND

During 1997-1998, RAND investigated the concept of a possible "minimum essential information infrastructure" (MEII) for U.S. defense systems. The results of that study are available in a RAND report, "Securing the U.S. Defense Information Infrastructure: A Proposed Approach," by R. Anderson et al (1999), available at <http://www.rand.org/publications/MR/MR993>. In that study, measures for increasing information security and assurance were placed into 13 categories, and 175 research projects underway were categorized accordingly. One category stood out as important, yet that showed little extant work: deception. A follow-on project has been under way since May 1998. This project explores applications of deception in information operations, as a possible aid to structuring research and development activities in this area, focusing on defensive techniques. Our tasks cover a wide range of activities, for example:

- Adapting existing (COTS or GOTS) hardware and software for deceptive purposes, in support of INFOSEC
- Developing explicitly deceptive tools for INFOSEC
- Devising methods by which deception may be used to divine identity and intentions of attackers, and improve indicators and warning (I&W) in cyberspace
- Developing guidelines for implementing deception on operational systems, at any level (from small LAN up to entire CINC-dom)
- Improving counter-deception capabilities (i.e., investigating how adversaries may employ deception to further their attacks, and how that deception may be combated)

This summary presents excerpts from the findings and developments of the first year's work on the subject.

PROJECT SUMMARY

The information security (INFOSEC) mission--protecting and maintaining the information assets of the US defense establishment--is a colossal and enduring challenge. The manifold difficulties faced by INFOSEC are likely to increase in the foreseeable future, as information technologies and the electronic environment evolve into greater complexity. Moreover, the stakes will continue to rise as US military and intelligence functions grow ever more reliant upon the defense information infrastructure. While INFOSEC research has been quite successful in producing numerous, formidable safeguards of information, information systems, and information infrastructure, many significant deficiencies remain. Tabulated in Table 1, these can generally be described as relating to the reactive nature of INFOSEC--building walls and plugging holes against dimly-perceived external threats--and the chaotic, noisy, evolving character of the electronic environment.

Table 1. Enduring Challenges to INFOSEC

Challenge
INFOSEC generally cedes the initiative to the adversary.
INFOSEC usually innovates in response to an adversaries' demonstrated capabilities or established friendly shortcomings.
There is a very large number of potential attackers, of many sorts.
The attackers may have many different motives, and many different objectives.
In most cases, attackers possess anonymity in the beginning; attacks emerge from under a blanket of secrecy.
There is a very large amount of data which might be relevant to the defense.
There is a very large amount of 'noise' surrounding the relevant data.
There are many 'locations' to defend, and adversaries may be insiders or outsiders.
Usually, the INFOSEC mission must be performed while keeping the protected systems up and running
Legal challenges may constrain (or even hamstring) optimal defense plans.

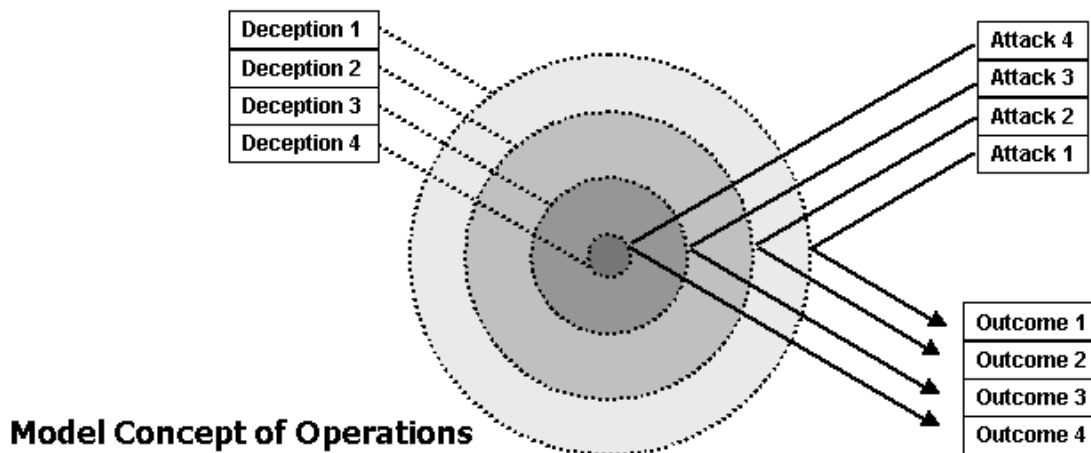
USING DECEPTION IN INFOSEC

Recent work suggests that INFOSEC research focuses perhaps too intently on traditional defensive measures (firewalls, encryption, biometrics, etc.) to the neglect of others. One such neglected measure, which the historical record suggests is a potent tool on both offense and defense, is deception. Used effectively, deception is also a valuable mechanism of intelligence-gathering, which is a critical piece of the INFOSEC puzzle. It is also proactive, in contrast to most other defensive measures. We believe deception in information systems, appropriately deployed and used, can address many of the INFOSEC challenges listed in Table 1. While often viewed as unsavory to practice and complicated to manage, there is little doubt that when employed successfully, deception is among the most powerful instruments of conflict. The investigators in this study feel that it is both reasonable and necessary to employ every legal means which may significantly advance the INFOSEC mission. Hence this research, which was conducted to explore how deception might contribute to INFOSEC in two general ways:

- What protective value deception measures can provide against a range of attacks on information infrastructure (paying close attention to any *unique* protective value which deception offers)
- What intelligence-gathering value deception measures can offer against a range of information infrastructure attackers

The authors' previous research suggests that deception offers powerful defensive and intelligence-gathering capabilities to INFOSEC. Moreover, deception is not a single tool: it is a diverse array of measures which may be employed individually or in depth, alone or in concert with more traditional defensive measures, as simple schemes or complex ruses. An essential part of employing deception is a well-honed CONOPS. Figure 1 illustrates the prototypical concept of operations which the authors are terming "deception in depth."

Figure 1. Deception in Depth



1. Defensive deceptions implemented in layers (weakest at periphery, strongest at core)
2. Intruders pierce or fall prey to deceptions differentially, depending upon their knowledge, experience, capabilities, determination, resources, etc.
3. Differential outcomes enrich the friendly intelligence picture of the enemy, as intruders reveal capabilities, MO, intentions, etc.

Deception methods which accomplish the same ends (e.g., diverting the attacker from his goal) may vary dramatically in their form (i.e., the diverting method may be decoy targets, disinformative bulletins, etc.). Further, while deception can have defensive value alone, deception methods employed in concert with other defensive measures can significantly enhance those other measures' effectiveness. There may thus be a synergistic effect in the use of deception alongside other defensive measures. And if the historical record is any guide, defensive deceptions may be gainfully employed from the tactical up to the strategic level with equal success. For INFOSEC, this means that deception can protect individual machines or local area networks (LANs), all the way up to the assets of an entire military command. In the

course of this study the authors mapped the rich diversity of military deception, criminal deception, deception in human psychology, animal and plant deception, and numerous other perspectives on deception, into the domain of INFOSEC. As an illustration of these efforts consider deception in the plant and animal kingdoms. In studies of animal behavior, it does not suffice to say that an animal employs camouflage. Rather, it is necessary to delve deeper, and ask questions about the particulars: how that camouflage is prepared; how tailored it is to specific adversaries; how adaptable to changing environmental conditions it is; and the like. Table 2 illustrates this principle as derived from studies of deception in animal biology. Similarly useful insights into deception and counterdeception were gained from the other disciplines investigated.

Table 2. Deception: Level of Sophistication

Level	Description	Example
STATIC	The deception method is in place irrespective of state, activity, or history of either deceiver or target.	Typical camouflage BDUs
DYNAMIC	The deception method is employed by the deceiver when circumstances trigger it	Tank coloration that changes with prevailing background
ADAPTIVE	The deception method is triggered as in DYNAMIC above, but the method or triggering event may be modified by feedback (i.e., trial-and-error)	Decoy emplacement improved based upon battle-damage assessment
PREMEDITATED	The deception method is designed and implemented by the deceiver based upon experience, knowledge of own vulnerabilities/strengths, and observations of the target	Choosing decoys instead of camouflage based upon analysis of enemy sensor vulnerabilities, decision-making weaknesses, etc.

TOOL DEVELOPMENT

After mining these varied resources for general principles and classification schemes, the second undertaking was the development of deception tools --the prototypical beginnings of an arsenal of defensive deceptions--for experimental testing and evaluation. These first generation tools (developed in a UNIX environment) certainly do not exhaust the possibilities, in fact they just scratch the surface. This was the authors' intent; we sampled the vast 'space' of possible deceptions in order to demonstrate the principles we have distilled, and to provide the audience with a heterogeneous set of tools for assessment. We have also been investigating the use of standard, COTS tools in deceptive ways. The new tools fielded in this first generation of development concentrated in four areas:

- **Varied methods for concealing processes deemed sensitive.** One method subverted a single standard process-reporting application (*ps*), while the other modified library routines that access */proc*, meaning that all applications which report on processes will yield deceptive results. The former technique might be considered of less sophistication and/or challenge than the latter, promoting both differential protection as well as some information-gathering capability in their use.
- **Varied methods for concealing directory contents.** One method subverted a single standard directory content-reporting application (*ls*), while the other generated a deceptive environment using an extant application (*chroot*). The former technique might be considered of less sophistication and/or challenge than the latter, also (as above) promoting both differential protection as well as some information-gathering capability in their use.
- **A method for generating a population of notional hosts on a "virtual network," once again exploiting an extant application (*chroot*) and an experimental software demon (*virtuald*) to novel ends.** This technique, considered complex and sophisticated in our classification scheme, offers the opportunity to plunge an attacker into a notional network topography. This provides both protective value as well as the opportunity to collect intelligence on the attacker.
- **A method for quickly and cheaply generating vast numbers of attractive, yet nonsensical, target documents.** This technique exploits a freely available tool (*dada*) to generate useful decoy and disinformation documents that might surround and camouflage key files.

The net result of these development efforts has been:

1. A set of hypotheses and research questions to be answered experimentally. These questions go the heart of this line of inquiry: Which kinds of deception are most effective? Against which kinds of attacker? Against insiders as well as outsiders? Against which kinds of attack? What are the benefits of deception versus other kinds of defensive investment (e.g., firewalls, etc.)? What is the protective value of deception in the short-term vs. the longer-term? Does deception in fact offer any unique defensive capabilities to INFOSEC? How valuable are these capabilities? How necessary?

2. A set of lessons and meta-lessons learned in the design and construction of these tools. Virtually no analytic work has previously been done in harnessing deception to the purposes of INFOSEC, and as such this research has produced a litany of findings on exactly how principle meets practice. For example, the notion of "wrapping" deception tools around extant applications is a tempting prospect. It would require less invasive intervention, would be more handily uninstalled, and would be less costly in management and maintenance. However, our studies to date have found numerous reasons why "wrapping" is an untenable method of implementing deception. The authors believe that these kinds of findings are unique among analyses of the use of deception, and yet are essential to any serious consideration of it use.

We view this present research as the first step in an intensive study of employing deception in INFOSEC. Later reports will document findings from experimental test and evaluation of the evolving prototype "toolkit" of deceptions we are in the process of developing.

EXPERIMENTAL SCENARIOS

We are attempting to field deceptive measures applicable against a wide range of adversaries, from trained, well-financed operatives acting on behalf of nation-states to script-kiddies operating out of their parents' home. The following three scenarios illustrate our attempt to parse the vast range of adversaries, capabilities and intentions into manageable and distinct subsets. They are only to be thought of as representative, and useful for setting up experimental environments to test friendly deception measures. Note: The term "friendly environment" used below refers broadly to friendly systems, capabilities, purpose/mission, personnel, topography, vulnerabilities, and the like.

Experiment One

Scenario	Outsider with little knowledge of friendly environment intrudes into target network and searches for targets of opportunity (files) to steal.
Deception	Hacked applications omit/edit out references to valuable files.
Subject	No prior knowledge of friendly environment. Resources limited to a few, widely-available, canned tools. No knowledge that deception may be present
Objectives	The subject is told that this network is restricted and may contain valuable data (measured in dollars or prestige). The data is said to relate to nuclear secrets. The subject is instructed to find and steal whatever they can and get out. The subject is told that vandalism or denial activities are likely to bring attention.
Performance Criteria	Finding something which appears valuable is worth x
	Stealing something which appears valuable is worth 3x
	Being discovered halves the subject's final score.
Access & Permissions	For the sake of argument, we assume that the subject has hijacked a typical user account and password, which gives him "mid-level access". The subject intrudes over a landline (56.6K) and may do whatever he wishes from there, using either indigenous instruments or the tools mentioned above as imports.
Timeline	Short. The subject is given 1 hour from start to finish.
Deception Implementation	A dozen files, a couple processes, and a handful of net-addresses are on the hide-list.
Other Defensive Measures	A firewall is in place, and good passwords are in use.

Experiment Two:

Scenario	Savvy outsider with some knowledge of friendly environment seeks to deny service or vandalize friendly systems.
Deception	virtuald- and chroot- based pseudo-environment ("sandbox")
Subject	Some knowledge of friendly environment. Resources generally unlimited in canned tools, and subject can improvise on the fly. No knowledge that deception is present, but knowledge that deception is possible.
Objectives	The subject is told that he should do as much damage as possible to the system, although at his discretion this may be immediate or "buried," visible or concealed.
Performance Criteria	Degrading service will be worth x
	Destroying service will be worth 2x
	Deleting valuable files is worth x
	Concealing evidence of the damage doubles its value
Access & Permissions	For the sake of argument, we assume that the subject has hijacked a typical user account and password, which gives him "mid-level access". The subject intrudes over a landline (56.6K) and may do whatever he wishes from there, using either indigenous instruments or the tools mentioned above as imports.
Timeline	Short. The subject is given 1 hour from start to finish.
Deception Implementation	A virtual network is in place, and chroot "sandboxes" are in place on each member of the population of phantom machines.
Other Defensive Measures	A firewall is in place, and good passwords are in use.

Experiment Three

Scenario	Insider with extensive knowledge of friendly environment is intent upon gaining restricted access and siphoning out valuable data (files) for profit.
Deception	virtuald- and chroot-based pseudo-environment ("sandbox"), plus massive data-generated pseudo-files larded in restricted directories.
Subject	In-depth knowledge of friendly environment. Resources unlimited in canned tools, and subject can improvise. Subject knows deception is present, but not its exact form.
Objectives	The subject is told that he should gain root or a specific user's access, and steal specific files of value.
Performance Criteria	Gaining target access (root or particular user) is worth x
	Stealing target file is worth x
	Spotting and avoiding deception triples values
	Giving information (target file, target user) to deception tool logging is worth $-x/\text{instance}$
Access & Permissions	The subject has an account already on the target system, with "mid-level access". The subject may act freely over an Ethernet connection using whatever tools he wishes, imported or indigenous.
Timeline	Long. The subject is given 3 days from start to finish.
Deception Implementation	A virtual network is in place, and chroot "sandboxes" are in place on each member of the population of phantom machines.
Other Defensive Measures	A firewall is in place, and good passwords are in use.

CONCLUDING REMARKS

Most members of the defense and intelligence communities readily acknowledge the importance of deception on both offense and defense, in the form of masking (anonymity, secrecy, concealment, etc.), misdirecting (diversions, disguises, etc.), or confusing ("noise" generation, purposeful deviation from established patterns, etc.). However, there are too many unanswered questions with regards to the effectiveness, importance, costs and risks of using (or falling prey to) deception for decision-makers to be considered fully-informed on the topic. This assertion applies as much or more with regards to the INFOSEC mission than it does to more conventional arenas of conflict. It is this lacuna in research and analysis which is the engine driving this project. Through a well-developed theoretic framework, tool development, wide-ranging experimentation, and thorough analysis the authors hope to provide operational guidance and instruments for both employing and combating deception in INFOSEC.

An extensive bibliography on deception is available from the authors. Please contact Scott_Gerwehr@rand.org.