Designing Deception Operations for Computer Network Defense

Jim Yuill, Fred Feer, Dorothy Denning
Fall 2004

Abstract: Deception is an appealing means for computer network defense (CND), as it pits the defender's strengths against the hacker's weaknesses. Hackers rely heavily, if not exclusively, on a single source of information—network data. The data is easily manipulated, and the hacker is highly vulnerable to deception. The defender has physical control of the network, and he knows the network well. Further, deception can be used to attack hackers' decision-making processes; thus deception provides an offensive security-measure-- something computer security defenders sorely lack. This paper explains how deception operations can be designed and developed for CND, including incident response, intelligence, detection, and prevention. Deception processes, principles and techniques are presented. They are based on the underlying nature of deception, and the extensive military deception-literature.

1 Introduction

This paper explains how deception can be used to advantage in computer security, including incident response, intelligence, detection, and prevention. It describes the process followed in deception operations, and it describes principles and techniques for developing and conducting deception operations. The paper focuses on deception principles that are of enduring use, and independent of current technologies. For instance, honeypots are currently one of the most widely used deceptions. The paper uses honeypots to illustrate principles, but honeypots are not the paper's focus. This paper is an abridgement of a larger work that we hope to publish as a book.[1]

Deception is an integral part of human nature and experience. However, few people use deception in the calculated manner needed for computer security. As military deception reveals, effectively deceiving an adversary is a job skill. The principles of military deception are well documented in the military deception-literature, and they are based on millennia of experience and thought. This paper adapts principles of military deception to computer security deception. In addition, one of this paper's authors has extensive experience in both military and intelligence deception.

In the following sections, section 2 describes the deception-operation process; section 3 describes deception-operation planning, and section 4 describes how to build the deception.

2 An overview of deception operations

In this section, basic deception concepts and terminology are presented, followed by a description of the deception-operation process.

[1] The larger work (80 pages) is currently available to DoD personnel. Contact Jim Yuill at: jimyuill at pobox dot com.

2.1 Basic concepts and terminology

Computer security deception is defined as being those actions taken to deliberately mislead hackers and to thereby cause them to take (or not take) specific actions that aid computer security.[2] Deception aims to mislead the hacker into a predictable course of action or inaction that can be exploited [Dew89]. Tricking the hacker, and making him think a certain way, is important only as a step toward getting him to make the decision that will result in the desired action [JDD96]. Thoughts without action are of little computer security value.

The scope of this paper is deception for computer security defense. It focuses on the tactical use of deception for a computer network. Some of the paper's computer and deception terms are defined as: 1) **CND:** computer network defense, 2) **deception planner**, or **planner:** the person who plans, develops and carries out the deception operation, 3) **deception operation:** the planned development and deployment of a deception-based computer security measure, 4) **target:** the person the deception operation seeks to deceive, 5) **intelligence:** information and knowledge obtained through observation, investigation, analysis, or understanding, 5) **ruse:** a trick designed to deceive.

Deception operations vary in the purposes they serve, the networks on which they are used, and the different types of hackers they target. Some deceptions are simple and predictable. For instance, ping scans can be easily and predictably deceived. Other deceptions are complex and uncertain. For instance, a honeynet can be large, and there can be many servers with extensive false content. Although deception operations vary widely, there are processes and principles that are applicable to many, or even all, deception operations. As a theoretician of military deception has observed:

> The basic principles and objectives of reinforcing the desires and perceptions of the deceived will not change, since human nature and the psychological mechanism of human perception are ever the same. In terms of its forms and the means employed, deception will, like war itself, change as new weapons and technologies appear. [Han85]

2.2 The deception-operation process

The deception-operation process involves complex adversarial relationships and complex engineering systems. Although the process of deception operations can be complex, there is a basic deception-process that is followed in almost all operations. This basic deception-process is shown in Figure 1, and it is described below.[3] (References to items in the figure are in **bold text**.) Due to the complexity of deception operations, this basic process is a simplified conceptual model. It is not meant to provide a complete description of all deception operations' elements and interactions.

[2] This definition is adapted from the U.S. DoD definition of military deception [JDD96].
[3] This basic deception-process was adapted from a draft written by our colleague Dr. Bowyer Bell.

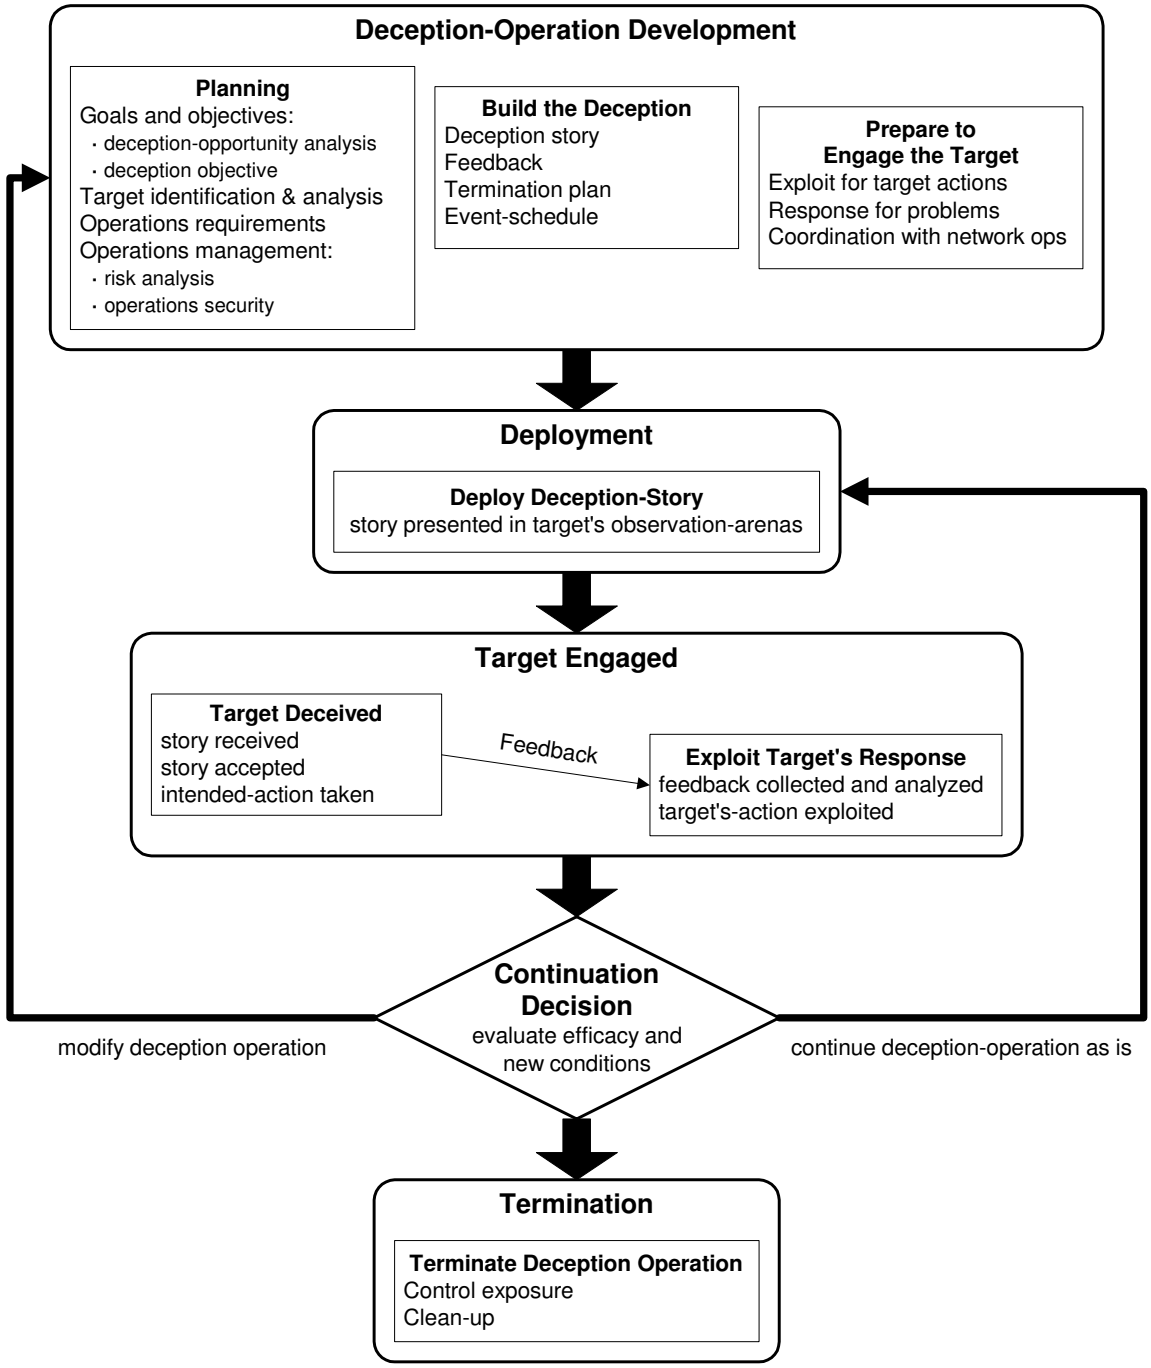

Figure 1 : The basic deception process

The deception operation begins with **Deception-Operation Development** (top-most box in figure). The deception operation's **plan**, **deception,** and means for **engaging the target** are developed, roughly in that order. **Planning** is an iterative process that is conducted throughout the deception operation. Its first step is recognition of the need or **opportunity** to deceive a target. What must be done in deception planning, and often is not, is to determine the result

desired from the deception. Mere acceptance of the deception may not be advantageous, and it may in fact prove costly. For example, a clever honeypot could attract an unwanted horde of script kiddies, and hiding a host's log files may make the hacker uncertain of the evidence he's left, prompting him to erase the entire file system, just to be safe (e.g., "rm –rf /").[4] Thus deception is a means, not an end. The **objective** of a deception operation is: **1)** to induce the target to take some specific action-- perhaps to do nothing, and **2)** to exploit that action, or otherwise use it to advantage.

Deception operations are ultimately against individual hackers, so planning includes **identification** of the deception **targets**, and **analysis** of their vulnerabilities to deception. Planning also involves **risk analysis** and **operations security**, to ensure the deception is not revealed to the target.

To induce the target to take the intended action, a **deception story** is designed, and it is implemented using various ruses. The deception story is presented to the target in his observation arenas. Typically, the most effective observation arenas are the target's intelligence sources. One of the primary ruses used in computer security are honeypots. They are computer systems that are designed to be probed, attacked or compromised by hackers [Spi03]. A honeypot contains servers and content that are attractive to hackers. A honeypot is placed on a network, where hackers are likely to encounter it. Honeypots are most useful for detecting attacks and for collecting intelligence about hackers. A honeynet is a network of honeypots.

The deception operation is **Deployed** (second box in figure) by presenting the deception story to the target, in his **observation arenas**. This is a key transition in the deception process, as the deception operation is now out the planner's control until the return of feedback that suggests an appropriate response. The deception story is maintained until it is received by the target. This can occur almost immediately, as with honeypots on a network's DMZ. Alternatively, the deception story might be maintained for months or years before being received, as might occur with an intranet honeypot used for detecting insider hacking.

The **Target is Engaged** (third box in figure) once he receives the deception story. The target is successfully **deceived** when he **receives** the deception story, **accepts** it, and, as a consequence, **takes the intended action**.

Feedback channels provide information about the target's reception of the deception story, and his response to it. The ultimate goal of deception operations is **exploiting the target's response**.[5] This occurs after the **feedback is collected and analyzed**, and it is known that the target has taken the intended action. For honeypots, feedback channels are an essential feature. For example, Symantec's ManTrap honeypot can record much of hacker's activity, including network traffic, process activity, and keystrokes. ManTrap can also detect hacker activity and send alerts.

The deception story exerts control on the target, manipulating him at a distance. Such manipulation may be intended to have a very short existence. For instance, BackOfficer Friendly (BOF) is a honeypot that can impersonate unauthorized remote-access servers, like BackOrifice. Such servers are installed by hackers via Trojan horses. BOF's impersonation is superficial and

[4] Unless stated otherwise, this paper's masculine pronouns refer to both men and women.

[5] Thanks to Fred Feer for showing us how the exploit is the deception operation's ultimate goal. In the military deception literature that we have read, the exploit's central role is under-emphasized.

its ruse can quickly be discovered by the hacker, but not before he is detected. Other deceptions may be intended to last indefinitely. For example, a fake VPN interface can be used to draw attention away from a network's real VPN interface. The deception is intended to last indefinitely.

A **continuation decision** (diamond in figure) is made for the deception operation, based on its efficacy and the current situation. The process can be **terminated**, **continued as-is**, or **modified**, in which case the process returns to deception-operation development. **Termination** (last box in figure) occurs when the deception story has achieved its purpose and is no longer needed, or when the target discovers the ruse. The target often discovers the ruse when his response to it is exploited. For example, hardware keystroke-loggers are dongles that attach to the keyboard cable. Their effectiveness depends on stealth: they are located behind the computer, appear to be a normal part of the cable, and few people know about them. When a hacker is confronted with evidence from a keystroke logger, the ruse will probably become apparent. Thus, it can no longer be used against him or his accomplices. Terminating the deception involves **controlling exposure** of the ruse, so it might be used again, as well as **cleaning-up** its affects upon computer systems and personnel. The remainder of this paper focuses on deception-operation planning and on building the deception.

3 Deception planning

A prince or general can best demonstrate his genius by managing a campaign exactly to suit his objectives and his resources, doing neither too much nor too little. Carl von Clausewitz

Deception-operation planning provides direction for the operation by developing its goals, objectives and requirements.[6] In conjunction, the targets are analyzed to learn their vulnerabilities to deception.

3.1 Deception opportunity analysis

Deception opportunity analysis identifies ways deception can be used to support CND. For the deception operation to be effective, it should be fully integrated with the overall CND effort. The deception operation must be compatible with, and coordinated with, the network's security and production operations. Deception is **not an end** in itself, and it should not be used simply because there are clever ways to trick hackers.

3.2 The deception objective

...it became a creed [among deception planners] to ask a General, "What do you want the enemy to do," and never, "What do you want him to think?" Dudley Clark, WWII deception planner[7]

The **deception objective** is the desired result of the deception operation; it consists of: **1)** the intended *target action*, and **2)** the *deception exploit.*[8] The **target-action** is a statement of what the hacker is to do (or not do) at some time and location. It is always stated in terms of specific actions, such as, "cause the targets' attacks against our server to be performed, instead,

[6] This planning process is adapted from the U.S. Joint Forces' deception process [JDD96].
[7] [Mur80]
[8] The *deception objective* is adapted from the U.S. Joint Forces deception manual [JDD96]. However, its *deception objective* only consists of the *target action*. We include the *deception exploit* with the deception-objective, as it is the deception-operation's ultimate objective.

against the honeypot server". A statement such as "have the hacker think that the honeypot server is the real server" is not a target-action, rather, it is a desired perception (described in section 4). Having the hacker think a certain way is important only as a step toward getting him to make the decision that will result in the intended action. Thoughts without action are of little security value.

The **deception exploit** is a statement of how the target-action will benefit CND, e.g., through attack detection, prevention, or response. The deception exploit may include actions to be taken against the target, following the target-action. For instance, the prior example's deception exploit would be, "for successful attacks against the honeypot, the honeypot will record the attack and send an alert." Some deception exploits do not require taking action against the target, e.g., when using a ruse to confound operating system (OS) fingerprinting, the deception exploit is thwarting attacks that depend on accurate OS fingerprinting.

The deception operation's ultimate goal is successful completion of the deception exploit. The deception-story and ruses are just means for inducing the target-action. After the story is deployed, feedback is analyzed to determine when the target-action is taken. The deception exploit can go into effect after the action is taken.

3.3 Target identification and analysis

It was so important to the deception work to be able to put oneself completely in the mind of the enemy, to think as they would think on their information, and decide what they would do.
WWII deception planner[9]

Deception attacks the target's perception and his thinking process, so effective deception requires **intelligence** on who the target is, how he works, and how he thinks. An understanding of how hackers work reveals their vulnerabilities to deception and how those vulnerabilities can be exploited. Fortunately, much is understood about how hackers work, as the complexity of hacking compels hackers to use publicly available tools and information. It is especially important to understand how the **hacker intelligence-process** works, as it is both a means and an end for deception operations. For instance, the hacker's intelligence collection is a means for communicating the deception story, and one use (end) of deception is thwarting the hacker's intelligence process.

There are specific **vulnerabilities** in the hacker's intelligence process that are helpful to know: **1)** the single sources of information that he may rely upon, as deception is easier when the ruse will not be cross-validated, **2)** the information he uses that is superficial and easily misrepresented, as with a ping scan, **3)** the investigations he performs when he is naive and thus easily duped, as during his initial network reconnaissance, and **4)** the intelligence processing of the hackers' automated agents, such as worms, as their simplicity and determinism may be easily duped.

Outsider hackers (non-insiders) are almost always **naive** about the networks they hack. A hacker's experience and skills are often asymmetric with the experience and skills needed for the network he is hacking. For example, hackers typically have never legitimately worked in a network or organization like the ones they are hacking. Even if a hacker has a high degree of technical skill, he may be naive about the network's topology and operations, as well as the

[9] [Mon78]

network personnel's language and culture.

4 The deception story

To induce the target to take the intended action, a **deception story** is designed, and it is implemented using various ruses. The deception operation's objective is to induce a specific target-action that benefits CND. The **desired perception** is what the target must believe in order for it to take the intended action [JDD96]. The **deception story** is an outline of how the computer system will be portrayed so as to cause the target to adopt the desired perception, and take the intended action. This section presents principles and techniques for developing the deception story.

In general, determining the desired perception can be difficult, as it requires an understanding of how the target works and thinks. Generally, it is much easier to **reinforce an existing belief** than to establish a new one [Heu81]. For example, a computer-savvy hacker will reasonably expect that a high-volume web site uses multiple web servers, load balancing, and a multi-tiered architecture. Also, to ensure the target-action is taken, the desired perception should make the target believe the target-action is in his **best interest**. Ideally, the target will perceive the intended action as compelling, and alternative actions as untenable.

4.1 Essential design-criteria

The deception story's essential design-criteria [JDD96,DH82b,USA78] are that it be:

- **plausible:** the story must be plausible from the target's perspective. Consequently, it should appear *appropriate* from both an engineering and operations perspective. Also, it must appear to be something the defender is *capable* of doing. The story should be *consistent* with real systems and operations, as well as being internally consistent.
- **receivable:** The story must be something the target's intelligence is capable of receiving and interpreting as intended.
- **verifiable:** If the target will verify the story through multiple intelligence sources, then the story should be portrayed through more than one source. For example, to avoid honeypot web sites, a target can verify web sites he discovers by searching for links to them from real web-sites.
- **efficacious:** For the story to be efficacious**,** it must be received, and it must effectively induce the desired perception and target-action.
- **implementable:** The story must be something the deception planner is capable of implementing.

4.2 Design principles

Some of the key design principles for the deception story are:

- **Inducing the target-action:**

The target-action is easier to induce if it something the target is predisposed to doing, such as: **1)** something he is already planning to do, **2)** something he normally does, or **3)** something he wants to do.

- **Making the story believable:**

In general, its easiest to persuade the target to believe something he already expects. Also, it can be easy to deceptively portray things that are normally **hidden** from the target. Often, the target only expects to find limited information about something that is hidden, in which case that is all that needs to be portrayed.

- **Preventing the target from uncovering the deceptions:**

The deception story's **falsehood** should be kept to a minimum. The truth is much stronger than a lie, and it can be difficult to maintain a lie over time. Also, minimizing falsehood makes the deception story easier to implement. Some techniques for minimizing falsehood are: **1)** make the story simple, **2)** weave the story into the truth, **3)** provide no more detail than is necessary, and **4)** impersonate things that are normally concealed from the target, as he will only expect to see bits and pieces of information about them, and only those pieces of information need to be portrayed. As an example of the latter technique, when impersonating a subnet that is protected by a stealthy firewall, only a few expected signatures may have to be shown.

Another way to prevent the deception story from being uncovered is to minimize the target's **scrutiny** of the deceptions. Three techniques for doing this are: **1)** the deceptions can be communicated to the target via his less scrutinizing intelligence capabilities. If the target cannot examine the deception closely, he will be less likely to detect it. **2)** Deceptions can be communicated to the target when he has little time to scrutinize them, and **3)** The deceptions can portray things of which the target has little understanding.

- **Ensuring the target receives the story:**

In the course of hacking, the target **eagerly seeks** particular information; this presents an opportunity for using the target's intelligence sources to communicate the deception story to him.

- **Revealing the story:**

A technique for revealing the deception story is to provide the story in bits and pieces and then let the target **piece the story together** by inference [USM89,Dew89]. The technique is consistent with the target's intelligence activities, as they normally acquire information in bits and pieces.

- **Implementing the story:**

Usually, only **parts** of the deception story will need to be implemented. Some of the story will be tied to the truth and portrayed by real systems and operations. Some of the story can be notional, implied by the parts of the story that are real and that are implemented.

To determine what parts of the story to implement, one must understand how the target receives the deception story, and what he expects to see [JDD96]. The target's **intelligence capabilities** determine how he receives the deception story. The deception planner must determine the things the target would expect to see, if the deception story was true. Having determined how the target discovers the deception story, and what he expects to see, the planner can then determine the parts of the story to implement.

- **Realism:**

For each part of the deception story that is implemented, the deception planner will need

to determine its degree of realism. The **realism needed** is a function of: **1)** the target's intelligence capabilities, and **2)** the time the target has available to analyze the situation and take appropriate actions [FN95]. Often, minimal realism is needed for deceptions that the target has little time to observe and analyze [FN95]. For example, hackers have little time to observe during extensive port scanning. When many ports are to be scanned, each probe must be quick, and thus superficial. Such scans are easy to deceive, and the deception is fairly reliable. In general, it is best to design the deception story so that the amount of realism needed is **kept to a minimum**.

5 Conclusion

After years of research and development, computer security remains an error-prone task and, in some respects, a loosing battle. Deception offers an opportunity to exploit the hacker's weaknesses and to attack his decision-making process. Deception is not useful in all situations. However, even when working with a small security budget, deception can be among the most cost-effective means for securing valuable assets. Conversely, not using deception may be an indication of over estimating the effectiveness of conventional security measures

Bibliography

[Dew89] Dewar, M. *The Art of Deception in Warfare*, David & Charles, 1989.
[DH82a] Daniel, D., K. Herbig, editors. *Strategic Military Deception*, Pergamon Press, 1982.
[FN95] Fowler, C., R. Nesbit. "Tactical Deception in Air-Land Warfare", *Journal of Electronic Defense*, June 1995.
[Han85] Handel, M. *Military Deception in Peace and War*, Magnes Press, 1985.
[Heu81] Heuer, R. "Cognitive Factors in Deception and Counterdeception", in [DH82a].
[JDD96] Joint Doctrine Division, *Joint Doctrine for Military Deception*, U.S. Joint Command, http://www.dtic.mil/doctrine, 1996.
[Mon78] Montagu, E. *Beyond Top Secret Ultra*, McCann & Geoghegan, 1978.
[Spi03] Spitzner, L. *Honeypots : Tracking Hackers*. Addison Wesley, 2003.
[USA78] *FM 90-2 Tactical Deception*, U.S. Army, 1978.
[USA88] *FM 90-2 Battlefield Deception*, U.S. Army, 1988.
[USM89] *FM 15-6 Strategic and Operational Military Deception: U.S. Marines and the Next Twenty Years*, U.S. Marine Corps, 1989.

Authors

Jim Yuill is a senior PhD student in the Computer Science Department at North Carolina State University. This paper is part of his dissertation. Jim previously worked at IBM in operating systems development. *jimyuill-at-pobox.com*
Fred Feer is retired from a career with the U.S. Army counterintelligence, CIA, RAND and independent consulting. Deception has been an interest and area of professional specialization for over 40 years. *ffeer-at-comcast.net*
Dr. Dorothy Denning is a Professor in the Department of Defense Analysis at the Naval Postgraduate School. She is an ACM Fellow, and the recipient of several awards, including the National Computer Systems Security Award. *dedennin-at-nps.navy.mil*

Acknowledgements: This paper was made possible by funding from The Office of Net Assessment, in the Office of the Secretary of Defense.