

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.

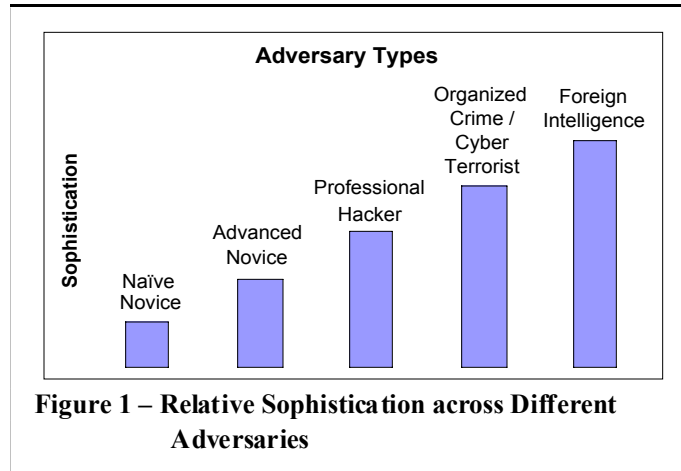


Figure 1 – Relative Sophistication across Different Adversaries

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.

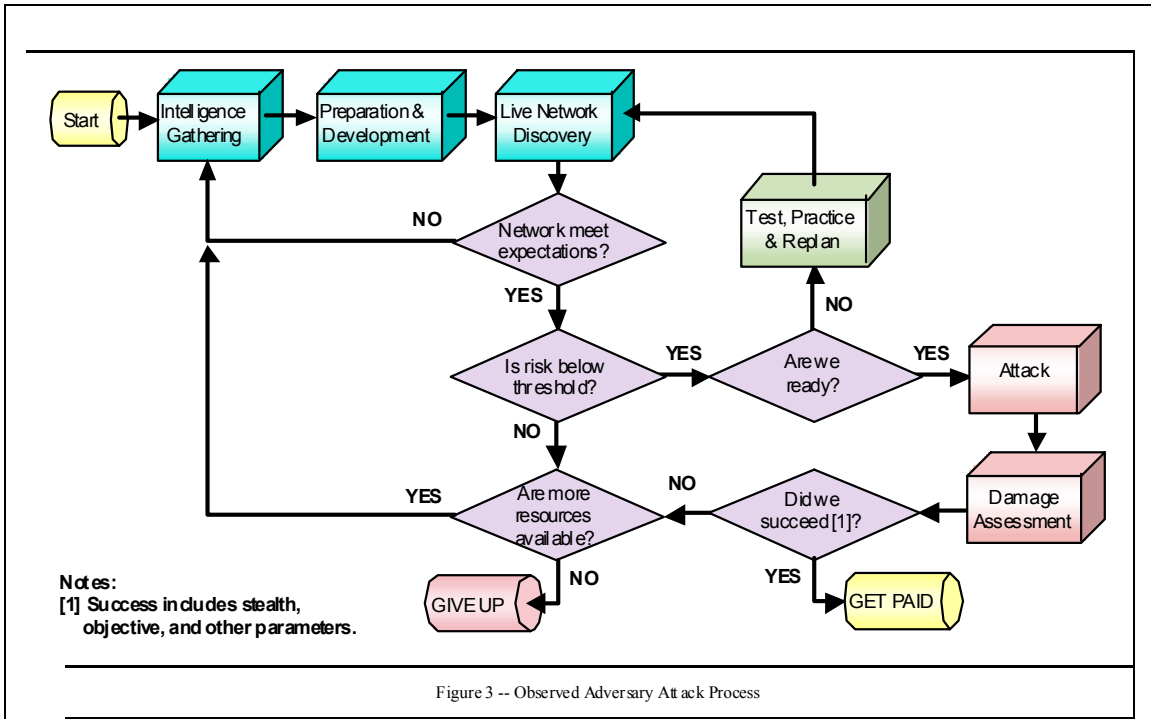
EARLY OBSERVATIONS

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.



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>