

# Adversary Work Factor as a Metric for Information Assurance\*

Gregg Schudel  
Mentor Technologies  
201 Defense Highway, Suite 200  
Annapolis MD 21401 (USA)  
1-703-861-4884  
(gschudel@mentortech.com)

Bradley Wood  
SRI International  
9508 Tasco Drive NE  
Albuquerque, NM 871111 (USA)  
1-505-798-9899  
(bjwood@sdl.sri.com)

## ABSTRACT

Adversary work factor is an informative metric for gauging the relative strengths and weaknesses of modern information systems. However, this metric can be difficult to measure and evaluate. This discussion reviews the efforts by a team at the Defense Advanced Research Projects Agency to evaluate this metric and to gauge its impact on a significant information security research effort.

## 1. INTRODUCTION

It is difficult to measure the relative strengths and weaknesses of modern information systems, especially when the safety, security, and reliability of those systems must be protected to some degree. Designers often apply information assurance or security technology to systems without the ability to evaluate the impact of those mechanisms to the overall system.

Classically, high-consequence information systems are designed to withstand some level of effort by a determined adversary. However, few efforts are directed at actually measuring the quantifiable impact of information assurance technology on the potential adversary.

The Defense Advanced Research Projects Agency's (DARPA's) Information Assurance Program has put a significant effort into modeling, observing, measuring, and evaluating the impact of information assurance technology on adversaries of interest. This effort has resulted in a large body of evidence that supports the assertion that adversary work factor is a quantifiable metric that yields valuable insights to complex information systems.

## 2. FUNDAMENTAL HYPOTHESIS

This work is based on the fundamental hypothesis that adversary work factor is a quantifiable metric that yields valuable insights into the relative strengths and weaknesses of modern complex information systems. This hypothesis is supported by the following assertions:

© 2001 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by a contractor or affiliate of the U.S. Government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.  
New Security Paradigm Workshop 9/00 Ballycotton, Co. Cork, Ireland  
© 2001 ACM ISBN 1-58113-260-3/01/0002...\$5.00

- Information assurance mechanisms are usually designed to frustrate some potential adversary. However, designers rarely make any attempt to quantify the impact of a given mechanism on a particular adversary.

- Classic approaches to gauging the impact of information assurance mechanisms [1] attempt to list a variety of requirements that a given information system must meet to guarantee assurance. However, there is little evidence to suggest that satisfying these requirements has any measurable impact on a determined adversary.

- Little is reported in the research community about potential adversaries, their capabilities, and their behavior.

Therefore, it stands to reason that designers could gain valuable insights into their systems if they could model, observe, measure, and evaluate the behavior of a potential adversary in response to various information assurance mechanisms.

## 3. AN APPROACH TO OBSERVING ADVERSARY WORK FACTOR

The DARPA Information Assurance (IA) Program has made a substantive effort to model, observe, measure, and evaluate adversary work factor in response to some of the technologies that are both planned and under development at DARPA.

### 3.1. The Model Adversary

DARPA began by developing a model adversary. DARPA tasked the Information Design Assurance Red Team (IDART) at Sandia National Laboratories [2] with the mission of providing a model adversary, pursuant to the following constraints:

- Represent the capabilities and behavior of a hypothetical cyber terrorist. This is based on the assertion that the nation is vulnerable to sophisticated cyber terrorism. However, little is discussed in the open literature about this adversary, its capabilities, and its impact.

- Rigorously document all tools, techniques, and behaviors that constitute the model adversary.

- Do not incorporate or disclose techniques that are uniquely held as national security information by the United States Government. All tools, techniques, and methods for this adversary must be

---

\* This work was performed under contract for the Defense Advanced Research Projects Agency's Information Assurance Program.

publicly available and could be subject to broad publication within the information security community.

- Focus on limiting the military effectiveness of the information systems under scrutiny.

The result was the formation of the IA Red Team. The goal of the Red Team was to model a modern cyber terrorist within the context of the DARPA research and development environment. The details of this model adversary are discussed in detail in a previous work [3]. However, other adversaries may be modeled in the future, pending development of detailed models of these other adversaries [4] [7].

### 3.2. Uses for the Model Adversary

The IA Red Team differs from an actual adversary in some important ways:

- The Red Team attempts to limit the actual damage to the information systems under scrutiny. The Red Team attempts no intentional physical damage to any system component. In addition, any accidental disclosure of information is limited to the members of the Red Team, who are actually trusted members of the development team.
- The Red Team is accountable to the DARPA IA Program Manager. There is little incentive for mischief on the Red Team.
- The Red Team discloses all tools, techniques, methods, and behaviors for a given exercise. This includes planning as well as the execution phases of any effort.
- The Red Team attempts to cooperate with the goals of the experiment and the overall program.

Generally, the goal of the Red Team is to assist in collecting data that either supports or refutes an experimental hypothesis. However, this goal can be realized through a variety of mechanisms:

**War Games:** The first engagement of the IA Red Team was to support a cyber-war game. The DARPA Information Superiority Technology Integration exercise for 1998 concluded in a week-long war game between a group of cyber-defenders (the Blue Team) and two different model adversaries (Red Team #1 and Red Team #2). This exercise was based on the hypothesis that war gaming is an effective means of gauging the robustness of a military information system. However, the results of this exercise were inconclusive [5] in that the data that was collected did not clearly support or refute this particular hypothesis.

**Experimentation:** The most prolific use of the IA Red Team has been to simulate a model adversary in the context of some limited experiment. These experiments are generally conducted in a controlled environment such as DARPA's Technology Integration Center [6]. Experiments also have a focused goal of collecting data that either support or refute some experimental hypothesis. Examples of topics explored in IA experiments include in Layered Defenses, Dynamic Defense, Dynamic Response, and Intrusion Detection system performance [7]. An example series of experiments will be included later in this paper.

**White-boarding:** *White-boarding* is a process that lets a variety of parties come together and discuss strategies for future model

engagements or exercises. This process began as a brainstorming tool for planning IA experiments. However, its value has been demonstrated in other venues. The primary purpose of this activity is to gather the model defenders (a Blue Team) and a model adversary (a Red Team) to conduct some sort of hypothetical limited engagement prior to committing resources to developing an actual experiment or exercise. The term *white-boarding* comes from the artifact that the most useful tool in this activity is a dry-erase board and a printer or digital camera to capture the results. A *white-boarding* session typically precedes any experiment (as previously described). However, this process alone has yielded some results in studying the IDART Cyber Terrorist Model and the placement of Intrusion Detection Sensors [7].

**Workshops:** The IA *White-boarding* process has evolved into a detailed brainstorming process for evaluating a variety of IA concepts. These workshops are unique in that they attempt to bring IA designers and developers together with the model adversaries. The results of this collaboration have been used to study Denial of Service Attacks, Other Adversaries, and the Malicious Insider threat [7].

**Other Activities:** The DARPA IA program has used their Red Team in a variety of other activities:

- Planning experiments that do not involve a model adversary.
- Scripting attacks for automated tools.
- Evaluating developmental IA technologies [8]
- Developing red team tools, techniques, and methods for other government agencies [9]

## 4. SUMMARY RESULTS

DARPA's IA program has had an aggressive program of experimentation. The IA Red Team has played a significant role in many of these exercises.

The following is a summary of some of the IA Red Team's activities conducted during the period of October 1998 through April 2000. The IA program does not currently publish the actual data from these and other activities. However, this data can be released to researchers by contacting DARPA directly [7]. The following is a summary of some of the IA Red Team engagements.

### 4.1. Red Team Experiment 1999-01

Red Team Experiment 1999-01 (RT9901) was an experiment to gauge the effect of multi-layer security middleware [7]. The fundamental hypothesis under evaluation was: "*Adding layers has at least a cumulative impact on adversary work factor.*"

The approach that the IA team used was to compare attacker work factor as more defensive layers were added in a typical client-server database network.

The Red Team was expected to attack four different configurations of a representative enclave as shown in Figure 1. Each configuration was designed to represent an increasing number of layers of security middleware. These configurations included:

**Level 1:** Baseline of Secure Sockets Layer, packet filtering firewall, TCP/IP firewall proxy

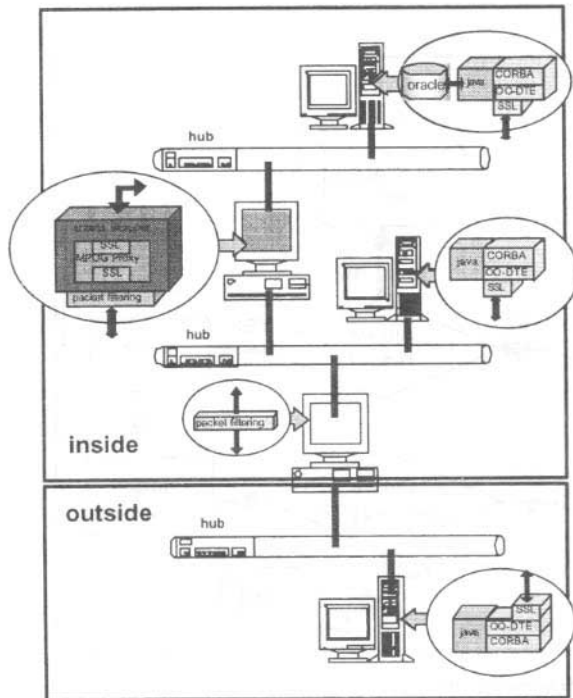


Figure 1 – RT9901 Experimental Enclave

- Level 2: Configuration 1 (baseline) plus Object-Oriented Data Type Enforcement
- Level 3: Configuration 2 plus Multi-protocol Object Gateway
- Level 4: Configuration 3 plus Generic Software wrappers.

Red Team Work Factor was compared for each configuration as the red team attempted to compromise three basic system properties: confidentiality, integrity, and availability. Red Team Work Factor was recorded as shown in Figure 2 and Figure 3.




Figure 2 - Red Team Work Factor for Confidentiality and Integrity Attacks

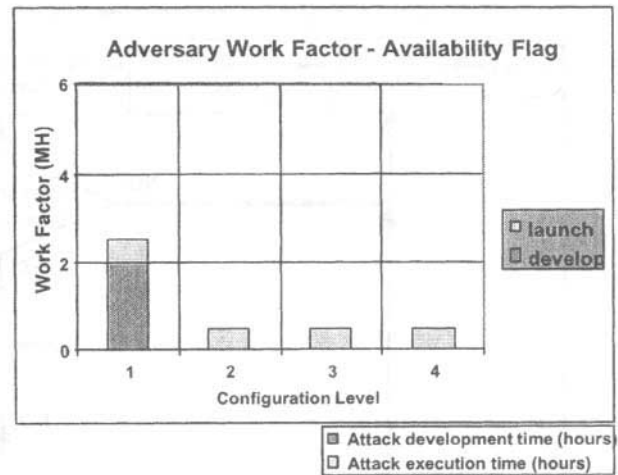


Figure 3 - Red Team Work Factor for Availability Flag

Analysis of this data suggests the following:

- Attack development time dominated attack execution time.
- Access control policy was not enforceable until Configuration Level 3.
- Insider access was required for Red Team success at Level 3 and 4. This was allowed under the rules of engagement for the exercise.
- Some IA mechanisms were actually used to frustrate the defenders.
- Layering to defend against one class of attack may not be effective in defending against different classes of attacks.

Some of the important conclusions drawn from this data are illus-

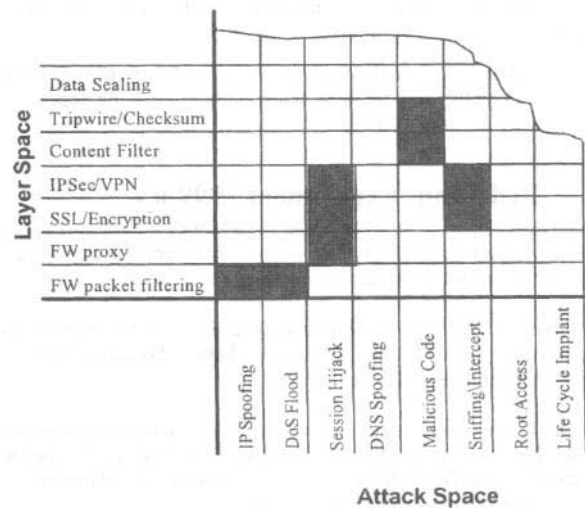


Figure 4 -- Layered Defense Illustration

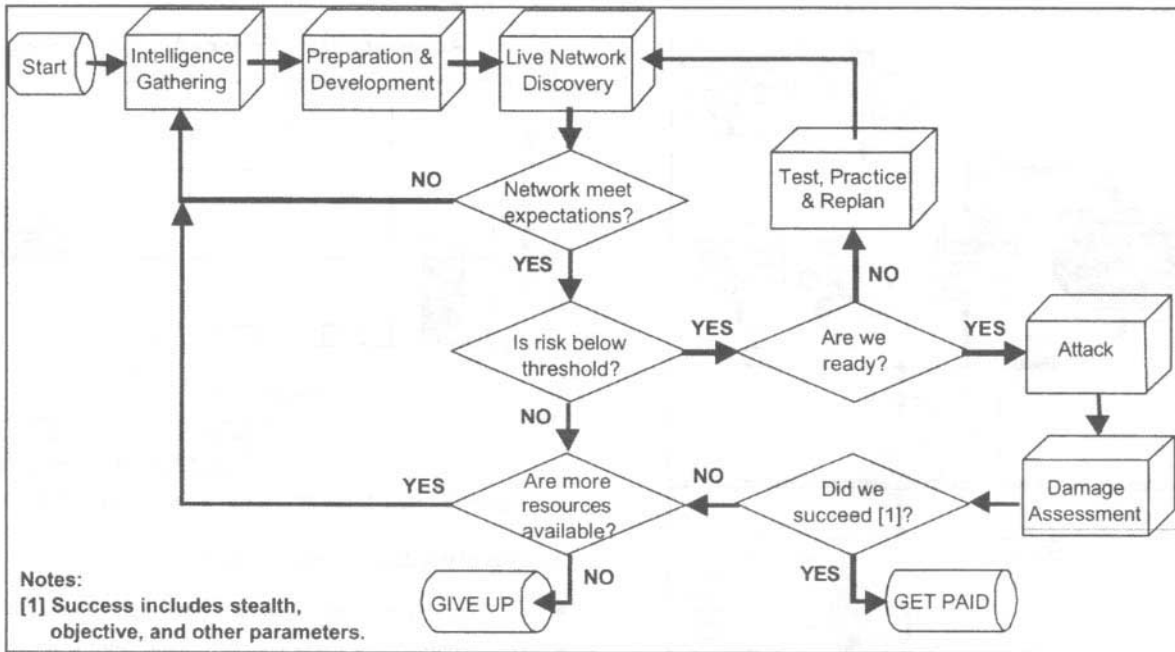


Figure 5 - Cyber-Terrorist Process Model

trated in Figure 4 and include:

- **Depth without breadth is useless.** It is more important to defend the broad range of attacks than to install multiple mechanisms to protect against a single type of attack. The adversary needs only one entrée into a system, whereas the defender must guard against all relevant attacks.
- **Individual layers may address specific attacks or classes of attacks.** It appears to be difficult to install a single defense mechanism that defends against a broad range of attacks.
- **Layers can move attack points to manageable places.** Defenders may actually be able to herd the adversary within a rational attack space.
- **Dependencies of individual layers must be managed.** This became apparent when the Red Team started to use the dependencies to frustrate the Defenders.

#### 4.2. Red Team Experiment 1999-03

Red Team Experiment 1999-03 (RT9903) was initiated to evaluate this hypothesis: "Dynamic defense mechanisms can have a significant impact on adversary work factor."

The exercise began as a white-boarding exercise to explore how the adversary spends their time and whether Dynamic Defense would impact their behavior.

The Red Team used actual accounting data to determine how they spent their time in previous exercises. The Red Team also allowed the Blue Team (the Defenders) to observe the planning of a Red Team exercise against a third party.

These observations resulted in documentation of the Red Team's Cyber-Terrorist behavior model [3]. These observations and assertions can be summarized as follows:

- The Adversary follows a well-defined process, as shown in Figure 5.
- The Adversary spends most of their time in intelligence-gathering activities, as shown in Figure 6.
- The Adversary is risk-averse, and they will not attack unless their perceived risk is below some threshold (as illustrated in Figure 7).





Figure 6 - Adversary Time Expenditure

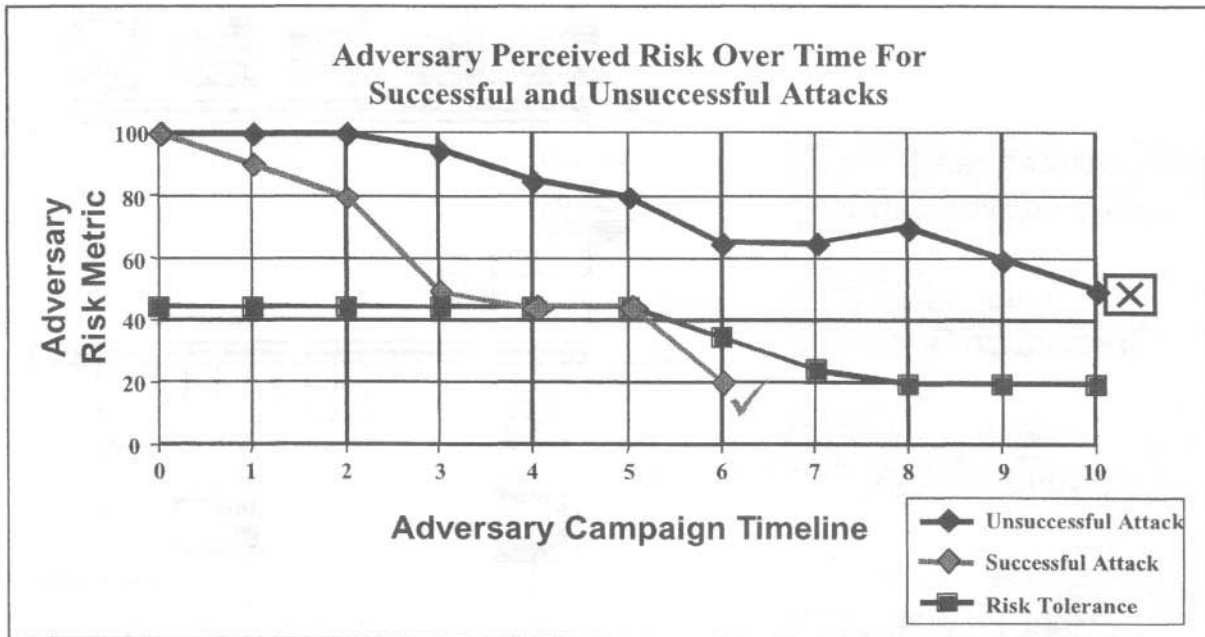


Figure 7 - Cyber-Terrorist Risk Profile

The group also agreed on the following assertions, based on this data:

- Dynamic deception / reconfiguration could hinder adversary intelligence gathering, potentially denying the adversary the information they need for a successful attack.
- Dynamic deception/reconfiguration can limit the time-value (life span) of adversary knowledge.
- Risk increases adversary resource expenditures. They will spend resources to reduce risk up to the limit of available resources. If no attacks can be created that meet both risk and resource constraints, the attack can not be accomplished.

Evaluation of this data led to the development of a targeted Red Team exercise to test these assertions.

### 4.3. Red Team Exercise 1999-07

Red Team Exercise 1999-07 (RT9907) was designed to test the hypothesis that “dynamic network reconfiguration effectively degrades the attacker’s ability to map the network, and hence increases attacker work factor and improves system assurance.”

The approach that the IA team used was to compare attacker work factors in the network discovery phase of an attack for static and dynamic reconfiguration architectures. The basic experimental network used for this exercise is shown in Figure 8.

The Red Team’s mission was to identify a critical server in the sample application and then target that server for an availability attack. The Red Team was not told in advance that a variety of deception techniques would be used, or that their work factor would be compared for successive tests of these different deception techniques.

The Red Team would attack a baseline network configuration, and then a configuration utilizing a Dynamic Network Translation bridge (DYNET) [10]. After DYNET was introduced, the Red Team was given increasing detail on how DYNET worked to determine whether intelligence gathering by the Red Team would have played a role in the Red Team’s ability to accomplish their mission.

The results of exercise RT9907 are summarized in Figure 9. Some of the observations from this data include:

- Learning by the Red Team could be observed for runs A1, A2, and B1.
- The Red Team was unable to accomplish their mission in runs C1, C2, D1, and D2. Therefore, only the data acquisition or network discovery times could be compared throughout the experiment.
- The Red Team was given more information about DYNET beginning in run D1.
- The Red Team was more successful at accomplishing their mission when the Secure Sockets Layer (SSL) encryption was turned off in run E1. Even with DYNET present, the Red Team had no problem accomplishing their mission when SSL was disabled.

Analysis of this data suggests the following:

- Dynamic remapping made *discovery* efforts more difficult and tedious for the Red Team. Work factor increases between 2:1 and 4:1 were observed.
- Blue Team (Defender) costs for the *on the fly* remapping scheme appears minimal. This implementation had NO impact to

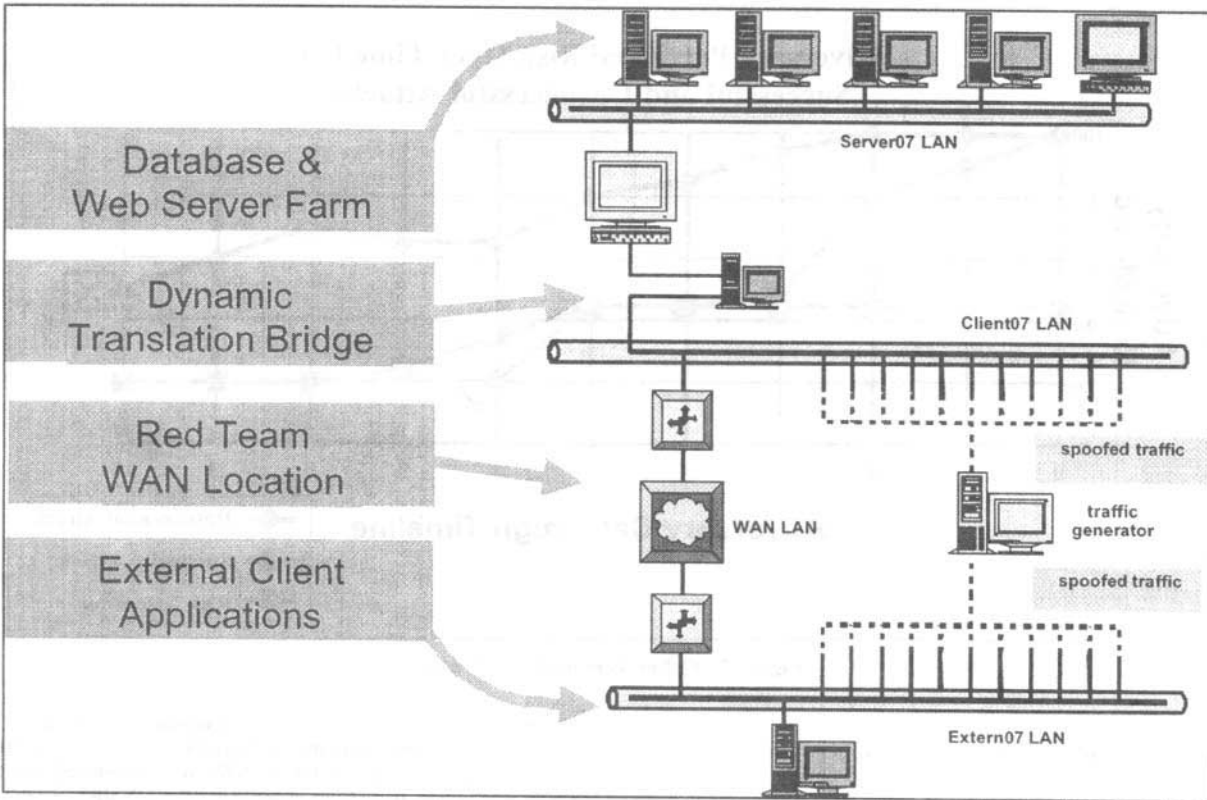


Figure 8 - Dynamic Network Experiment Network

client or server applications and negligible impact to the network architecture.

- Remapping changed adversary behavior:
  - If the adversary was just “poking around”, they probably would have moved on. Statistical *whiteness* of the traffic made the traffic uninteresting.
  - If the adversary had a particular mission to accomplish on this network, they would likely have been forced into a substantial data collection and analysis effort. This would have potentially required more time *in-place* increasing risk to the adversary due to exposure on the network.
  - Remapping could require the adversary to ship data off-platform for analysis, increasing the adversary's risk by requiring them to move larger data files throughout the network.

All of this data suggests that DYNET and similar forms of dynamic deception or remapping of network resources is worthy of further study.

## 5. LESSONS LEARNED

One of the primary benefits of experimentation with the IA Red Team has been the lessons learned from the process. These valuable lessons include:

**The best results are achieved through focused experiments.** This may seem like an obvious conclusion. However, the IA team has collected a large body of data at some considerable cost that supports this assertion. This data include the results from exercises IST198 [5], IFE 2.3 [7], and perhaps RT0001 [7].

As a result, the IA program has established guidance that discourages experimentation using the Red Team during large integrated tests or exercises. Priority for experimentation resources is given to limited Red Team experiments that are designed to study a focused hypothesis.

**Effective experimentation requires cooperation among a variety of diverse parties.** These groups typically include:

- **Principal Investigator(s)** – a small group or individual who proposes an experiment.
- **Experimentation Working Group** – a group of senior IA staff who critique the hypothesis, experiment, and data collection process for each experiment considered by the IA program.
- **Planning Team** – a group of IA professionals with some experience in conducting experiments who is assigned the responsibility for planning and executing a given experiment.
- **Integration Team** – a group of developers who must rapidly realize an experimental network configuration, including representative applications, in a timely and robust manner to support a particular experiment.

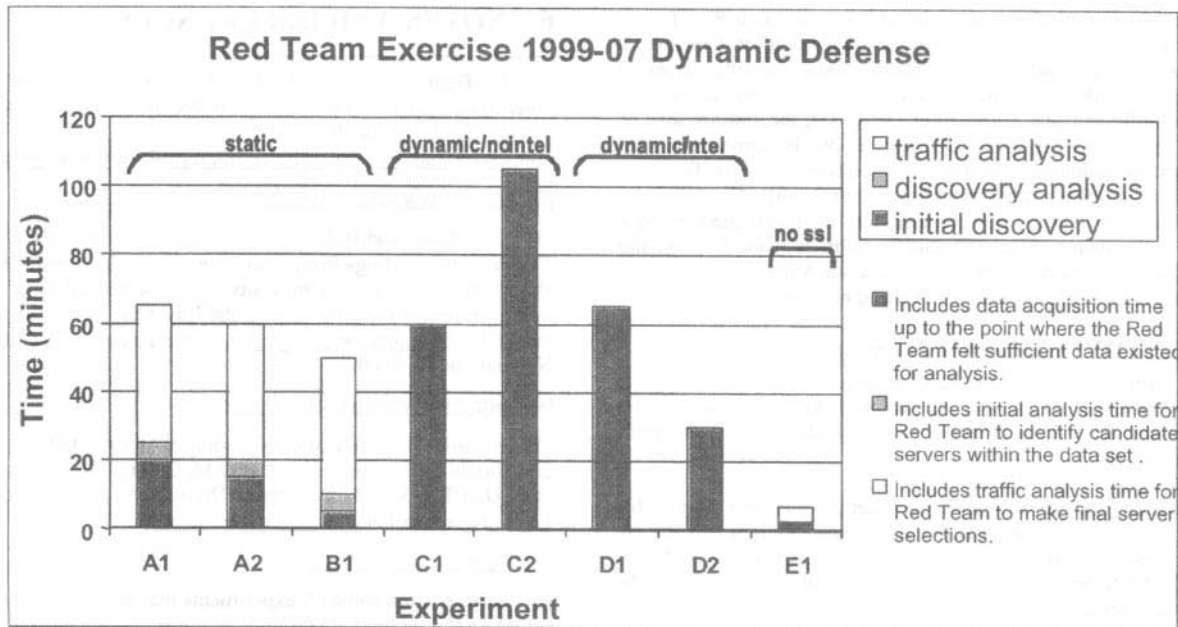


Figure 9 - Summary Results from RT99-07

- **Experiment Execution Team** – the team that must carry out a particular experiment. Ideally, this team consists of three to 12 people. However, the IA program has attempted experiments that involve as many as 40 people. These *Execution Teams* typically consist of defenders, adversaries, a test director, an experiment support group, and perhaps even a referee.

- **Data Analysis Team** – the group of independent analysts who study the data gathered from an exercise. This team may also attempt to explain any *unexpected results*.

**Many different activities benefit from participation by the Red Team.** In the IA program, the Red Team participates in most of the IA program's functions and teams. Red Team members serve as Principal Investigators, members of the Experimentation Working Group, most (experiment) Planning Teams, and most (experiment) Execution Teams. The Red Team rarely participates on the Integration Team, although there is close coordination between these groups. Red Team members have not yet been asked to serve as Defenders, although this concept has been considered for future experiments.

**The most interesting results of any experiment are the unexpected results.** For example, RT9901 suggested that the IA community (in general) does not understand how *layers* compose. RT9907 suggested that the Red Team was frustrated more by SSL than by DYNET. Therefore, experiments must be planned to observe these unexpected results. This requires that the participants gather enormous amounts of data, simply because the teams cannot anticipate these results. In addition, exercises must be run in a controlled fashion that can still adapt to contingencies and new information.

**Integration before Experimentation – always.** Every reasonable effort must be made to insure that the test environment is operating as expected before the Red Team arrives and experimentation

begins. Often, integration and testing of the test bed is occurring at the same time that the experiment is scheduled to begin. The risk is that the test bed will not perform as expected, and lengthy troubleshooting and repair efforts are required. Unfortunately, this can occur while the Red Team is sitting idle at a rate of up to \$200 / man-hour. As a result, the IA program has adopted guidance that requires some sort of end-to-end test of the environment before the Red Team arrives.

**It can be difficult to constrain even a cooperating adversary.** Typically, Defenders and Developers want to constrain the Red Team to attack very specific portions of a network. Conversely, Adversaries usually want to attack anything that stands between them and their target, regardless of whether this is of interest to the designers or the Principal Investigator. It would be unrealistic for the parties to agree that the Adversary would only attack certain IA mechanisms, because a real adversary would not be so cooperative. So, how do the Defenders constrain the model adversary without unintentionally skewing the results of an experiment? Ideally, Defenders constrain the model adversary with the same mechanisms they use to constrain an actual adversary – by use of various IA technologies and other defenses. However, it can be difficult to develop a test network that contains all for the needed defenses to properly constrain the Red Team. Therefore, the IA program has adopted the approach that the Defenders may stipulate certain defenses for a given network. The Red Team is then expected to behave as though these mechanisms are actually in place. This has some risk of skewing the results, simply because the Red Team may not accurately model the response of an adversary to the stipulated defenses. However, this approach appears to be an acceptable risk compared to the additional risk of requiring that each experiment exhibit a full set of robust defenses just to constrain the model adversary

**Relative measures of Red Team Work Factor yield the most information.** One can argue that the IA Red Team is only one model adversary, and they are not necessarily representative of all

of the adversaries of interest. In addition, different Red Teams might exhibit very different behaviors, depending on their preparations, training, and talents. Therefore, the absolute values of adversary work factor measured as a result of any experiment contain little valuable information. However, the comparisons of work factor between different exercises by the same team may yield more valuable information. Therefore, IA Red Team experiments usually require establishing some sort of baseline for Red Team and system performance, and then multiple runs of a given experiment. Every reasonable effort is made to limit the variables between each run. Still, the Data Analysis Teams must account for learning by the red team and other variables.

## ACKNOWLEDGEMENTS

IA experimentation with the Red Team has been a stimulating and rewarding experience. The authors wish to thank some of the people who have made this effort such a productive and enjoyable environment:

**Sami Saydjari** was the DARPA Program Manager responsible for establishing the Information Assurance program and the IA experimentation effort. Sami's wisdom, vision, support, and clarity of thought have been the guiding principals responsible for the group's success.

**The IA Integration Team** at the Technology Integration Center, including Mike Dean, Dorene Kewley, and Tom Hash (all with BBN Technologies division of GTE [1]), for their boundless effort and ingenuity in supporting an aggressive and ambitious experimentation schedule.

**The Information Design Assurance Red Team** at Sandia National Laboratories [2], including Julie Bouchard, David Duggan, Ray Parks, and Dave Farrell, for their ingenuity and creativity in challenging the fundamental assertions that constitute our current understanding of Information Assurance.

## 6. NOTES AND REFERENCES

- [1] US Dept. of Defense *Standard 5200.28-STD, Trusted Computer System Evaluation Criteria*, 26 December 1985 (a.k.a. *The Orange Book*), on the Internet at <http://www.radium.ncsc.mil/tpep/library/rainbow/index.html>
- [2] <http://www.sandia.gov/idart>
- [3] G. Schudel and B. Wood, *Modeling Behavior of the Cyber-Terrorist*, proceedings from *Countering Cyber-Terrorism Workshop*, June 22-23, 1999, University of Southern California, Information Sciences Institute, Marina del Ray, CA. Proceedings available at <http://www.isi.edu/cctws> or via email from Clifford Neuman ([bcn@isi.edu](mailto:bcn@isi.edu))
- [4] <http://www.csl.sri.com/~biwood>
- [5] Inquiries for information regarding DARPA's ISTI 98 exercise should be directed to Ms Cathy McCollum, Program Manager, DARPA, 3701 North Fairfax Drive, Arlington VA 22203-1714 or via email at [cmccollum@darpa.mil](mailto:cmccollum@darpa.mil)
- [6] <http://tictdc.isotic.org>
- [7] Results from some IA experiments may be published in the future on the Internet at <http://www.iaands.org>; Otherwise, researchers should direct inquiries to Mr. Michael Skroch, Program Manager, DARPA, 3701 North Fairfax Drive, Arlington VA 22203-1714 or via email at [mskroch@darpa.mil](mailto:mskroch@darpa.mil)
- [8] R. Duggan and B. Wood, *Red Teaming of Advanced Information Assurance Concepts*, proceedings from DARPA Information Survivability Conference and Exposition (DISCEX 2000), January 25-27, 2000, page Vol. II, page 112
- [9] Proceedings from *Assessments for Information Assurance, a Strategic Roadmap Workshop*, sponsored by DARPA & the Joint Information Operations Command, March 14-15, 2000, Technology Integration Center, Arlington, VA. Information available on the Internet at [http://schafercorp-ballston.com/ia\\_roadmap](http://schafercorp-ballston.com/ia_roadmap) and <http://www.iaset.org/news.htm>
- [10] Questions about DYNET should be addressed to Mr. Mike Dean via email to [mdean@bbn.com](mailto:mdean@bbn.com)
- [11] <http://www.bbn.com>