

**By Members of
the DETER and
EMIST Projects**

CYBER DEFENSE TECHNOLOGY NETWORKING AND EVALUATION

As the Internet has become pervasive and our critical infrastructures have become inextricably tied to information systems, the risk for economic, social, and physical disruption due to the insecurities of information systems has increased immeasurably. Over the past 10 years there has been increased investment in research on cyber security technologies by U.S. government agencies (including NSF, DARPA, the armed forces) and industry. However, a large-scale deployment of security technology sufficient to protect the vital infrastructure is lacking. One important reason for this deficiency is the lack of an experimental infrastructure and rigorous scientific methodologies for developing and testing next-generation cyber security technology. To date, new security technologies have been tested and validated only in small- to medium-scale private research facilities, which are not representative of large operational networks or of the portion of the Internet that could be involved in an attack.

To make rapid advances in defending against attacks, the state of the art in evaluation of network security mechanisms must be improved. This will require the development of large-scale security testbeds [3] combined with new frameworks and standards for testing and benchmarking that make these testbeds truly useful. Current deficiencies and impediments to evaluating network security mechanisms include lack of scientific rigor [6]; lack of relevant and representative network data [5]; inadequate models of defense mechanisms; and inadequate models of the network and both the background and attack traffic data [1]. The latter is challenging because of the complexity of interactions among traffic, topology, and protocols [1, 2].

To address these shortcomings, we will create an experimental infrastructure network to support the development and demonstration of next-generation information security technologies for cyber defense. The Cyber Defense Technology Experimental Research network (DETER network) will provide the necessary infrastructure—

Creating an experimental infrastructure for developing next-generation information security technologies.

networks, tools, and supporting processes—to support national-scale experimentation on emerging security research and advanced development technologies. In parallel, the Evaluation Methods for Internet Security Technology (EMIST) project will develop scientifically rigorous testing frameworks and methodologies for representative classes of network attacks and defense mechanisms. As part of this research, approaches to determining domains of effective use for simulation, emulation, hardware, and hybrids of the three are being examined.

The goal of this joint effort¹ is to create, operate, and support a researcher- and vendor-neutral experimental infrastructure open to a wide community of users. It is intended to be more than a passive research instrument. It is envisioned to serve as a center for interchange and collaboration among security researchers, and as a shared laboratory in which researchers, developers, and operators from government, industry, and academia can experiment with potential cyber security technologies under realistic conditions, with the aim of accelerating research, development, and deployment of effective defenses for U.S.-based computer networks.

Information Security Challenges

To develop a testbed framework for evaluating security mechanisms, the project focuses on a select subset of the overall problem space. Several different types of attacks and defenses will be studied with two goals: to elevate the understanding of the particular attack or defense by thoroughly evaluating it via different testing scenarios; and to further the understanding of the degree to which these evaluations can be unified into a single framework that spans the diversity of the problem space.

¹There are nine teams involved in the joint effort: U.C. Berkeley, U.C. Davis, University of Southern California-Information Sciences Institute (USC-ISI), Pennsylvania State University, NAI Laboratories, International Computer Science Institute (ICSI), Purdue, SPARTA Inc., and SRI International. The project also includes an industrial advisory board consisting of equipment vendors, carriers, and ISPs including AOL, Cisco, Alcatel, Hewlett-Packard, IBM, Intel, Juniper, and Los Nettos.

Three different classes of attacks are focus areas for our research: denial-of-service, worms, and attacks on the Internet's routing infrastructure, as well as attacks that are coordinated combinations of these three types. Together they span a broad range of general types of attacks. In addition, the project will closely monitor new Internet security breaches in order to analyze how new attack scenarios can be incorporated into the developing testing methodology. The focus of this effort will be on attacks targeting network infrastructure, server end-systems, and critical end-user applications. Such attacks are difficult to accurately simulate using existing testing frameworks because of the major challenges in accurately simulating Internet phenomena in general [2, 4].

Security Testing Methodologies

Testing frameworks will be adapted for different kinds of testbeds, including simulators such as NS (see www.isi.edu/nsnam/ns), emulation facilities such as Emulab [8], and both small and large hardware testbeds. The frameworks will include attack scenarios, attack simulators, generators for topology and background traffic, data sets derived from live traffic, and tools to monitor and summarize test results. These frameworks will allow researchers to experiment with a variety of parameters representing the network environment including attack behaviors, deployed defense technology, and the configuration of the defense mechanisms under test. It will be critical to make progress on the very difficult problems, particularly:

- How to construct realistic topologies, including bandwidth and inter-AS policies,
- How to generate realistic cross-traffic across these topologies,
- How to quantify how accurate the models need to be, and
- How to select the best metrics for evaluating various defense mechanisms.

Conducting these tests will require incorporating defense mechanisms into a testbed (either as models or as operational code), and applying and evaluating the frameworks and methodologies. Conducting these tests will also help to ensure the testbed framework allows other researchers to easily integrate and test network defense mechanisms of their own design. Furthermore, the documentation of the tests will serve as a tutorial for users of the testbed framework as they confirm their results or evaluate their own mechanisms and techniques.

Testbed Architecture and Requirements

The preliminary requirements for the DETER Testbed are drawn from four sources: a DARPA-funded study of security testbed requirements [3], input from network security researchers, general considerations on network research testbeds through a NSF workshop [4], and experience with a variety of earlier experimental and test networks. High-level requirements are briefly described here.

The general objectives for the testbed design require that it must be fully isolated from the Internet and all experiments must be soundly confined within the DETER network. Furthermore, it is expected the network will be subjected to destructive traffic and that experiments may temporarily damage the network. Therefore, there must also be mechanisms for rapid reconstitution of the testing environment.

The scale of the testbed will be approximately 1,000 PCs, each with multiple network interface cards, and a significant number of commercial routers and programmable switches. Within this environment, the network must provide sufficient topological complexity to emulate a scaled down but functionally accurate representation of the hierarchical structure of the real Internet, and to approximate the mixing of benign traffic and attack traffic that occurs. Initially, the network will be formed using a homogeneous network of existing technology. Carefully chosen hardware heterogeneity—commercial router boxes—will be added as the effort progresses. Finally, conducting experiments with large-scale denial-of-service attacks and defense technologies to protect the Internet infrastructure will require high-bandwidth componentry.

In addition to the preceding infrastructure requirements, there are various requirements for software to facilitate experimentation. The utility of DETER will depend on the power, convenience, and flexibility of its software for setting up and managing experiments including registration, definition, generation control, monitoring, check-pointing, and archiving. An important aspect of the management software will be the requirement for sophisticated network monitoring and traffic analysis tools for both experimenters and DETER network operators. Experimenters will also require traffic generation software to generate attack traffic and typical day-to-day (legitimate use) traffic.

Preliminary Architecture. DETER will be built as three permanent hardware clusters, located at ISI in Los Angeles, ISI-East in Virginia, and UC-Berkeley. To provide the earliest possible service to experimenters, initial development during the first six months focuses on building software and configurations for cyber security experimentation on PlanetLab and/or Emulab [8].

The architecture will also deploy aspects of the X-bone (see www.isi.edu/xbone) to allow topologies with revisitation, where, for example, a 10-node ring can be used to emulate a 100-node ring by visiting the same node multiple times. During the early stages of the testbed, this will enable the simulation of topologies that are larger than can be supported with one-to-one mapping of physical resources. Meanwhile, a phased development effort, moving from carefully controlled emulation environments to a mix of emulation and real network hardware will occur.

Conclusion

The development of testing methodologies complemented by an experimental infrastructure will support the realistic and consistent evaluation of mechanisms purported to mitigate large-scale attacks. This is an extremely challenging undertaking—no existing testbed or framework can be claimed to be effective. The research described here requires significant advances in the modeling of network attacks and the interactions between attacks and their environments, including deployed defense technology, background traffic, topology, protocols, and applications. It will also require advances in the understanding of metrics for evaluating defense mechanisms.

Our results will provide new scientific knowledge to enable the development of solutions to cyber security problems of national importance. This will be accomplished through experimentation and validation of cyber defense technologies using scientific methods. The lack of open, objective, and repeatable validation of cyber defense technologies has been a significant fac-

tor hindering wide-scale adoption of next-generation solutions. Results obtained using the DETER testbed will contribute to the development of innovative new technologies that increase commercial availability and viability of new production networks and services, providing true cyber protection. **G**

REFERENCES

1. Floyd, S. and Kohler, E. Internet research needs better models. *Hotnets-I* (Oct. 2002).
2. Floyd, S. and Paxson, V. Difficulties in simulating the Internet. *IEEE/ACM Transactions on Networking* 9, 4 (Aug. 2001), 392–403.
3. Hardaker, W. et al. *Justification and Requirements for a National DDoS Defense Technology Evaluation Facility*. Network Associates Laboratories Report 02-052, July 26, 2002.
4. Kurose, J., Ed. *Report of NSF Workshop on Network Research Testbeds* (Nov. 2002); gaia.cs.umass.edu/testbed_workshop.
5. McHugh, J. Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system valuations as performed by Lincoln Laboratory. *ACM Transactions on Information and System Security* 3, 4 (Nov. 2000), 262–294.
6. Pawlikowski, K., Jeong, H., and Lee, J. On credibility of simulation studies of telecommunication networks. *IEEE Communications Magazine* (Jan. 2001).
7. Peterson, L., Anderson, T., Culler, D., and Roscoe, T. A blueprint for introducing disruptive technology into the Internet. In *Proceedings of the 1st ACM Workshop on Hot Topics in Networks (HotNets-I)* (Oct. 2002), 4–140.
8. White, B. et al. An integrated experimental environment for distributed systems and networks. In *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI02)*, (Dec. 2002).

MEMBERS OF THE DETER AND EMIST NETWORK PROJECT INCLUDE R. BAJCSY, T. BENZEL, M. BISHOP, B. BRADEN, C. BRODLEY, S. FAHMY, S. FLOYD, W. HARDAKER, A. JOSEPH, G. KESIDIS, K. LEVITT, B. LINDELL, P. LIU, D. MILLER, R. MUNDY, C. NEUMAN, R. OSTRENGA, V. PAXSON, P. PORRAS, C. ROSENBERG, J.D. TYGAR, S. SASTRY, D. STERNE, AND S.F. WU.

This project is funded jointly by the U.S. Department of Homeland Security (DHS) and the National Science Foundation (NSF) under grant ANI-0335241.

© 2004 ACM 0002-0782/04/0300 \$5.00

Coming Next Month in Communications

Etiquette for Human-Computer Relations

A growing community of computer scientists, researchers, sociologists, psychologists, educators, and industry practitioners are taking the “etiquette perspective” in designing, building, and analyzing human interaction with computers and other forms of advanced automation. And they are finding, among its many advantages, the etiquette approach facilitates user acceptance of systems and products and, more importantly, improves the accuracy and speed with which the users develop trust in such products.

It’s a practice in its infancy and next month’s special section introduces the concept of etiquette and provides a variety of perspectives on its use and importance, including a number of examples of current research and applications.

Also in April Software Project Risks and their Impact on Outcomes • Cross-Cultural Issues in Global Software Outsourcing • Who Should Work for Whom? • Managing Academic E-Journals • Economics of Wireless Networks • Managing Knowledge in Distributed Projects • Information Cascades in IT Adoption