

Sawmill

Infrastructure for Distributed Collaboration in Detecting Network Attacks

PIs: Jay Aslam, David Kotz, and Daniela Rus

February 2002

This document addresses the reviewer comments for the proposed project “Infrastructure for Distributed Collaboration in Detecting Network Attacks”. Specifically, we clarify how this work meets government and industry needs at this time and address how this work is different than existing related projects.

Addressing Current Needs

Today hackers disguise their attacks by launching them from a set of compromised hosts distributed across the Internet. It is very difficult to defend against these attacks or to identify the hackers’ origin. These attacks have the power to paralyze operations and access restricted documents. Often, there is no way of identify that an attack is in progress until it is too late. Many times it is hard to identify what the attacker actually did. These are very important issues that need to be addressed by the computing community.

Current intrusion-detection systems, whether commercially available or research prototypes, can signal the occurrence of limited known types of attacks. New types of attacks are launched regularly but these tools are not effective in detecting them. Human experts are still the key tool for identifying, tracking, and disabling new attacks. Often this involves experts from many organizations working together to share their observations, hypothesis, and attack signatures. Unfortunately, today these experts have few tools that help them to automate this process.

Automated intrusion detection is difficult because new attacks are invented by hackers all the time. Existing intrusion detection systems, whether based on signatures or statistics, give too many false positives, they miss intrusion incidents, and they are generally difficult to keep current with all the attack signatures.

In this project we recognize that human experts will remain a critical part in the process of identifying, tracking and disabling computer attacks. We also recognize that an important part of the discovery, analysis, and defense against new distributed attacks is the cooperation that occurs between experts across different organizations. Furthermore, many installations do not have the expertize necessary to develop full attack analysis. Our goal is to build automated tools for computer experts and system administrators to

1. identify the characteristics of an attack given data from network sensors
2. develop a hypothesis about the nature and origin of the attack
3. share that hypothesis with security managers from other sites
4. test that hypothesis at those other sites and coordinate the results of testing
5. archive the data necessary for use as evidence in later law-enforcement actions.

So we propose to build two integrated *intrusion analysis* tools. The first allows humans (system administrators) to automatically examine logs and generate hypotheses for what is happening in the system. The tool helps the human iteratively refine a hypothesis about how the hacker got in, what the hacker did, and where s/he came from.

The second tool allows a system administrator to share that hypothesis with other system administrators, in a form that allows the receiving administrator to easily check the hypothesis at their site. It should prevent the leakage of proprietary information from the sender’s site, and allow the receiver to quickly verify that the testing of the hypothesis will not harm their own site.

Relation with previous work

Our proposed research project supports incident analysis and recovery, rather than on intrusion detection. Our work is complementary to Intrusion Detection Systems (IDSs). Our proposed project uses the output of intrusion-detection systems, such as SRI’s Emerald¹ and UCSB’s STAT², in two ways: (1) as a provider of events that will start an

¹<http://www.sdl.sri.com/projects/emerald>

²<http://www.cs.ucsb.edu/rsg/STAT/>

analysis; and (2) as data to be used in the analysis process. Indeed, we are collaborating with the UCSB team to develop a secure and flexible network-logging facility that can provide historical data to our analysis tool.

For us, an IDS provides the first alert that spawns a “backward” analysis aimed at identifying the sequence of actions (and associated evidence) that brought the system to the current (unsafe) state. During the incident analysis, IDS alerts will be used as supporting evidence, along with data from host and network logs.

The only relationship of the proposed research with the [IDWG of IETF](#)³ is that the analysis system must be able to understand the IDMEF alert format (and interoperate with the associated transfer protocol, IDXP) that is being standardized by the IDWG. We do plan to support IDMEF input.

Our proposed project is not closely related with MIT’s Lincoln Labs [IDS evaluation effort](#)⁴. The LL effort has been a three-year effort sponsored by DARPA. The goal of the LL project was to provide a means to evaluate IDSs, especially those funded by DARPA. Our proposed analysis system may use the test and training data produced as a byproduct of their project, to generate test cases for our analysis tools. Although the 1998 and 1999 data will be helpful, the data from the 2000 LL evaluation will be the most valuable because it contains data pertinent to multi-step attacks.

At NIST, [Computer Security Resource Center](#)⁵ has an intrusion-detection system based on mobile-agent technology. We do plan to use mobile code and possibly mobile agents to aid in distributed data collection, but not intrusion detection. The NIST [Common Criteria Evaluation Scheme](#)⁶ is another security-tools evaluation and certification project, but to the best of our knowledge they have not developed any tools like ours. The Federal Computer Incident Response Center ([FedCIRC](#))⁷, hosted at NIST, suggests tools for intrusion detection and is a forum for reporting attacks. This web site allows limited human-to-human collaboration about attacks but to our knowledge there is no specific software for forming, communicating, and automatic testing of hypotheses about attacks.

At CERT, the [AirCERT project](#)⁸ is another attempt to collect alerts from many sites around the country and to organize them into a Knowledgebase for broader analysis. AirCERT aims to collect intrusion information in real time, and to organize the information for CERT and others to analyse the data. AirCERT focuses on attacks that are known, or at least detectable by existing IDS technology.

CERT’s [Analysis Console for Intrusion Databases \(ACID\)](#)⁹ is a tool to analyse a database of alerts, log data, and packet data. This tool is perhaps closest to our project of any that we have seen, as it is a tool for exploring and analyzing intrusion data, and it can export information to email for informal collaboration. It does not, however, have any capability for hypothesis generation, refinement, or sharing, which are the core of our project.

A selection of other related literature is presented in the “References” section below, including our own short paper [[ACKR01](#)].

Our unique contributions

In short, our project has a unique approach to intrusion analysis, in that

1. we focus on intrusion analysis rather than intrusion detection,
2. our tools allow the human to collaborate with the system,
3. we propose a new process for iterating through hypotheses using log correlation,
4. we propose new algorithms for automated hypotheses refinement,
5. we propose distributed execution of log correlations,
6. we propose a collaborative process between humans via hypothesis sharing, and
7. we propose to develop new visualization methods for huge log data.

³<http://www.ietf.org/html.charters/idwg-charter.html>

⁴<http://www.ll.mit.edu/IST/ideval/>

⁵http://csrc.nist.gov/focus_areas.html

⁶<http://niap.nist.gov/cc-scheme/>

⁷<http://www.fedcirc.gov/>

⁸<http://www.cert.org/kb/aircert>

⁹<http://www.cert.org/kb/acid/>

References

- [ACF⁺00] Julia Allen, Alan Christie, William Fithenand, John McHugh, Jed Pickel, and Ed Stoner. [State of the practice of intrusion detection technologies](#). Technical Report CMU/SEI-99-TR-028, Software Engineering Institute, Carnegie Mellon University, January 2000.
- [ACKR01] Jay Aslam, Marco Cremonini, David Kotz, and Daniela Rus. [Using mobile agents for analyzing intrusion in computer networks](#). In *Proceedings of the Workshop on Mobile Object Systems at ECOOP 2001*, July 2001.
- [AT01] Accenture and The Center for Education and Research in Information Assurance and Security (CERIAS) at Purdue University. [CERIAS security visionary roundtable call to action](#). Accenture white paper version 1.0, January 2001.
- [Bas99] Tim Bass. [Multisensor data fusion for next generation distributed intrusion detection systems](#). In *Proceedings of IRIS National Symposium on Sensor and Data Fusion*, pages 24–27, May 1999.
- [Bel99] Steven M. Bellovin. [Distributed firewalls](#). *login:*, pages 39–47, November 1999.
- [CCD⁺99] Steven Cheung, Rick Crawford, Mark Dilger, Jeremy Frank, Jim Hoagland, Karl Levitt, Jeff Rowe, Stuart Staniford-Chen, Raymond Yip, and Dan Zerkle. [The design of GrIDS: A graph-based intrusion detection system](#). Technical Report CSE-99-2, Department of Computer Science, University of California at Davis, January 1999.
- [CPM⁺98] Crispian Cowan, Calton Pu, Dave Maier, Jonathan Walpole, Peat Bakke, Steve Beattie, Aaron Grier, Perry Wagle, Quian Zhang, and Heather Hinton. [StackGuard: Automatic adaptive detection and prevention of buffer-overflow attacks](#). In *Proceedings of the Seventh USENIX Security Symposium*, pages 63–78. USENIX Association, January 1998.
- [DCW⁺99] Robert Durst, Terrence Champion, Brian Witten, Eric Miller, and Luigi Spagnuolo. [Testing and evaluating computer intrusion detection systems](#). *Communications of the ACM*, 42(7):53–61, July 1999.
- [Den87] Dorothy E. Denning. An intrusion-detection model. *IEEE Transactions on Software Engineering*, SE-13(2):222–232, February 1987.
- [Den99] Denmac Systems. [Network based intrusion detection—a review of technologies](#). Denmac Systems, Inc., November 1999.
- [Fri00] Deborah Frincke. Balancing cooperation and risk in intrusion detection. *ACM Transactions on Information and System Security*, 3(1):1–29, February 2000.
- [Goa99] Terrance Goan. [A cop on the beat: Collecting and appraising intrusion evidence](#). *Communications of the ACM*, 42(7):46–52, July 1999.
- [HF99] Steven A. Hofmeyr and Stephanie Forrest. [Architecture for an artificial immune system](#). *Evolutionary Computation*, 7(1):1289–1296, 1999.
- [Ins98] The SANS Institute. [Building a network monitoring and analysis capability – step by step](#). Technical Report Version 1.1.5 980701, The SANS Institute, 1998.
- [JMKM99] Wayne Jansen, Peter Mell, Tom Karygiannis, and Don Marks. [Applying mobile agents to intrusion detection and response](#). Technical Report NIST Interim Report IR-6416, National Institute of Standards and Technology, Computer Security Division, October 1999.
- [LFM⁺00] Wenke Lee, Wei Fan, Matt Miller, Sal Stolfo, and Erez Zadok. [Toward cost-sensitive modeling for intrusion detection and response](#). In *Proceedings of the First ACM Workshop on Intrusion Detection Systems*, Athens, Greece, November 2000.

- [LNY⁺00] Wenke Lee, Rahul Nimbalkar, Kam Yee, Sunil Patil, Pragnesh Desai, Thuan Tran, , and Sal Stolfo. [A data mining and CIDF based approach for detecting novel and distributed intrusions](#). In *Proceedings of the Third International Workshop on Recent Advances in Intrusion Detection (RAID 2000)*, volume 1907 of *Lecture Notes in Computer Science*, pages 49–65. Springer-Verlag, Toulouse, France, October 2000.
- [LS98] Wenke Lee and Salvatore J. Stolfo. [Data mining approaches for intrusion detection](#). In *Proceedings of the Seventh USENIX Security Symposium*, pages 79–94. USENIX Association, January 1998.
- [LS00] Wenke Lee and Salvatore J. Stolfo. [A framework for constructing features and models for intrusion detection systems](#). *ACM Transactions on Information and System Security*, 3(4), 2000.
- [MHL94] Biswanath Mukherjee, L. Todd Heberlein, and Karl N. Levitt. Network intrusion detection. *IEEE Network*, pages 26–41, May–June 1994.
- [Pax98] Vern Paxson. [Bro: A system for detecting network intruders in real-time](#). In *Proceedings of the Seventh USENIX Security Symposium*, pages 31–52. USENIX Association, January 1998.
- [Pax99] Vern Paxson. [Bro: A system for detecting network intruders in real-time](#). *Computer Networks*, 31(23–24):2435–2463, December 1999.
- [PN98] Thomas H. Ptacek and Timothy N. Newsham. [Insertion, evasion, and denial of service: Eluding network intrusion detection](#). Secure Networks, Inc., January 1998.
- [RFR99] Steve Romig, Mark Fullmer, and Suresh Ramachandran. [Cisco flow logs and intrusion detection at the Ohio State University](#). *login:*, pages 23–26, September 1999.
- [Rom00] Steve Romig. [Correlating log file entries](#). *login:*, pages 38–44, November 2000.
- [Rui99] Dragos Ruiu. [Cautionary tales: Stealth coordinated attack HOWTO](#). *Digital Mogul*, 2(7), July 1999.
- [Rui00] Dragos Ruiu. [A DDoS proposal](#). Mailing List Posting, February 2000.
- [Sav99] Stefan Savage. [Sting: A TCP-based network measurement tool](#). In *Proceedings of the Second USENIX Symposium on Internet Technologies and Systems*, pages 71–79, October 1999.
- [SCH95] Stuart Staniford-Chen and L. Todd Heberlein. [Holding intruders accountable on the Internet](#). In *Proceedings of the 1995 IEEE Symposium on Security and Privacy*, pages 39–49, Oakland, CA, 1995. IEEE Computer Society Press.
- [Sch00] Fernando Schapachnik. [A DDoS defeating technique based on routing](#). Mailing List Posting, February 2000.
- [SF00] Anil Somayaji and Stephanie Forrest. [Automated response using system-call delays](#). In *Proceedings of the 9th USENIX Security Symposium*, Denver, Colorado, August 2000. pages not available from Usenix.
- [SK98] Bruce Schneier and John Kelsey. [Cryptographic support for secure logs on untrusted machines](#). In *Proceedings of the Seventh USENIX Security Symposium*, pages 53–62. USENIX Association, January 1998.
- [SMS99] Matthew Stillerman, Carla Marceau, and Maureen Stillman. [Intrusion detection for distributed applications](#). *Communications of the ACM*, 42(7):62–70, July 1999.
- [Sto99] Robert Stone. [CenterTrack: An IP overlay network for tracking DoS floods](#). In *Proceedings of the NANOG17 Meeting*, Montreal, Canada, 1999. NANOG.
- [SWKA00] Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson. [Practical network support for IP traceback](#). In *Proceedings of the 2000 ACM SIGCOMM Conference*, pages 295–306, Stockholm, Sweden, August 2000.
- [Tod00] Bennett Todd. [Distributed denial of service](#). *LinuxSecurity.com*, February 2000. Available only on the web.

- [VK99] Giovanni Vigna and Richard A. Kemmerer. [NetSTAT: A network-based intrusion detection system](#). *Journal of Computer Security*, 7(1):37–71, 1999.
- [YEA00] Jianxin Yan, Stephen Early, and Ross Anderson. [The XenoService— a distributed defeat for distributed denial of service](#). In *Proceedings of the 3rd Information Survivability Workshop (ISW2000)*, pages 195–200, Boston, October 2000. IEEE Computer Society Press.
- [ZL00] Yongguang Zhang and Wenke Lee. [Intrusion detection in wireless networks](#). In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking*, pages 275–283. ACM Press, August 2000.
- [ZP00] Yin Zhang and Vern Paxson. [Detecting stepping stones](#). In *Proceedings of the 9th USENIX Security Symposium*, Denver, August 2000. USENIX Association.